Report Series

“OASIS3 User’s Guide”

Number 2

Edited by:

Sophie Valcke, Arnaud Caubel, Damien Declat, Laurent Terray

OASIS3
Ocean **A**tmosphere **S**ea **I**ce **S**oil

User's Guide

August 2003

Sophie Valcke, Arnaud Caubel, Damien Declat, Laurent Terray

PRISM Project WP3a
Leading partner: C.E.R.F.A.C.S.
42 avenue Coriolis, 31057 TOULOUSE cedex 1

Abstract

This User's Guide contains a detailed step-by-step description on how to realize a coupled simulation with OASIS3.

The aim of OASIS3 is to provide a flexible and user friendly tool for coupling independent general circulation models of the atmosphere and the ocean (A/O-GCMs) as well as models of the other climate components (sea-ice, land, atmospheric chemistry, ocean biogeochemistry, ...) and regional models. The resulting coupled models are necessary tools to tackle current climatic paradigms such as the natural variability, El Niño Southern Oscillation (ENSO) or the greenhouse gas global warming effect.

OASIS3 synchronizes the exchanges of coupling fields between the models being coupled, and performs 2D interpolations and transformations needed to express, on the grid of the target model, the coupling field produced by the source model on its grid. Modularity and flexibility have been particularly emphasized in the OASIS3 design.

.

Contents

1 Overview of step-by-step use of OASIS3 **7**

2 Obtaining OASIS3 sources **8**
2.1 The OASIS3 package 8

3 Interfacing a model with the PSMILe library **9**
3.1 Initialisation 10
3.2 Partition definition 10
3.2.1 Serial (no partition) 11
3.2.2 Apple partition 11
3.2.3 Box partition 12
3.2.4 Orange partition 12
3.3 I/O-coupling field declaration 14
3.4 End of definition phase 14
3.5 Sending and receiving actions 15
3.5.1 Sending a coupling field 15
3.5.2 Receiving a coupling field 16
3.5.3 Auxiliary routines 16
3.6 Termination 17
3.7 Coupling algorithms - SEQ and LAG concepts 18
3.7.1 The lag concept 18
3.7.2 The sequence concept 21
3.7.3 A mix of lag and sequence: the sequential coupled model 21
3.7.4 An experiment mixing sequential and parallel coupled model runs: the use of `prism_put_restart_proto` 24

4 The OASIS3 configuration file *namcouple* **25**
4.1 An example of a simple *namcouple* 25
4.2 First section of *namcouple* file 27
4.3 Second section of *namcouple* file 29
4.3.1 Second section of *namcouple* for `EXPORTED`, `AUXILARY` and `EXPOUT` fields 30
4.3.2 Second section of *namcouple* for `IGNORED` and `IGNOUT` fields . . . 31
4.3.3 Second section of *namcouple* for `INPUT` fields 31
4.3.4 Second section of *namcouple* for `OUTPUT` fields 32

5 The transformations and interpolations in OASIS3 **33**
5.1 Using OASIS3 in the interpolator-only mode 33
5.2 The time transformations 33
5.3 The pre-processing transformations 34
5.4 The interpolation 36
5.5 The "cooking" stage 43
5.6 The post-processing 44

6 OASIS3 auxiliary data files **46**
6.1 Grid data files 46
6.2 Coupling restart files 47
6.3 Input data files 48
6.4 Transformation auxiliary data files 48
6.4.1 Auxiliary data files for `EXTRAP/NINENN`, `EXTRAP/WEIGHT`, `INTERP/SURFMESH`, `INTERP/GAUSSIAN`, `MOZAIC`, and `SUBGRID` 49
6.4.2 Auxiliary data files for `FILLING` 50
6.4.3 Auxiliary data files for `SCRIPR` 51

7 Compiling and running with OASIS3 **52**
7.1 Compiling OASIS3 and the toy coupled model 52
7.2 Running OASIS3 53

A The grid types for the transformations **54**

B The SCRIP 1.4 copyright statement **56**

1 Overview of step-by-step use of OASIS3

To use OASIS3 for coupling models and/or perform I/O actions, one has to follow these steps:

1. Obtain OASIS3 sources (section 2).

2. Identify the coupling or I/O fields and adapt the component models to allow their exchange with the CLIM communication technique and its associated PSMILe library, based on MPI1 or MPI2 message passing[1]. The PSMILe library is interfaced with the `mpp_io` library from GFDL [2] and therefore can also be used to perform I/O actions from/to disk files. A practical example of a toy coupled model is given in the `/prism/util/mod/toyclim` directory. (Section 3).

3. Define all coupling and I/O parameters and the transformations required to adapt each coupling field from its source model grid to its target model grid; prepare OASIS3 configuring file *namcouple* (sections 4 and 5). OASIS3 supports more interpolation algorithms for more different grids than before as it is now interfaced with the SCRIP 1.4 library [1] (see appendix B).

4. Generate required auxiliary data files (section 6).

5. Compile OASIS3, the component models and start the coupled experiment (section 7).

The toy coupled model in `/prism/util/mod/toyclim` has been successfully run on on Fujitsu VPP5000, NEC SX5, SGI Octane and 03000, COMPAQ Alpha cluster and Linux PC cluster. If you need extra help, do not hesitate to contact us (oasishelp@cerfacs.fr)!

[1] In OASIS3, the PIPE, SIPC, and GMEM communication techniques should still work but are not maintained anymore and were not tested.

2 Obtaining OASIS3 sources

2.1 The OASIS3 package

OASIS3 and all related libraries are available on CERFACS CVS server elnino.cerfacs.fr. The repository is /home/valcke/PRISMCVS. To obtain the CVS login and password as well as the most recent OASIS3 tag, please contact us (oasishelp@cerfacs.fr). OASIS3 directory structure follows the PRISM standard one:

```
- prism/src/mod/oasis3/src OASIS3 main code
                      /doc OASIS3 documentation

                /toyatm     toy model code 1
                /toyoce     toy model code 2
                /toyche     toy model code 3

- prism/src/lib/anaisg/     GAUSSIAN interpolation library

                /anaism/    SURFMESH interpolation library

                /clim/      CLIM/MPI1-MPI2 communication library

                /fscint/    INTERP interpolation library

                /pipe/      CRAY PIPE communication library
                            (not maintained anymore)

                /psmile/    PRISM System Model Interface Library

                /scrip/     SCRIPR interpolation library

                /sipc/
                /svipc/     SIPC communication library
                            (not maintained anymore)

                /sxgmem/    GMEM communication library
                            (not maintained anymore)

- prism/util/mod/toyclim    a toy coupled model environment

              /testinterp   an environment to test OASIS3 in the
                            interpolator-only mode
```

3 Interfacing a model with the PSMILe library

At run-time, OASIS3 acts as a separate mono process executable which drives the coupled run, interpolates and transforms the coupling fields. To communicate with OASIS3, or directly between the component models, different communication techniques have been historically developed. The technique used for one particular run is defined by the user in the configuration file *namcouple* (see section 4). In OASIS3, the CLIM communication technique based on MPI1 or MPI2 message passing is used.

To communicate with OASIS3, or directly between them, or to perform I/O actions, the component models need to be interfaced with the PRISM System Model Interface library (PSMILe), which sources can be found in `prism/src/lib/psmile` directory. PSMILe supports:

- parallel communication between a parallel component model and OASIS3 main process,
- direct communication between parallel component models when no transformations are required,
- automatic sending and receiving actions at appropriate time following user's choice indicated in the *namcouple*,
- time integration or accumulation of the coupling fields,
- I/O actions from/to files.

To adapt a component model to PSMILe, specific calls of the following classes have to be implemented in the code:

1. Initialisation (section 3.1)
2. Partition definition (section 3.2)
3. I/O-coupling field declaration (section 3.3)
4. End of definition phase (section 3.4)
5. I/O-coupling field sending and receiving (section 3.5)
6. Termination (section 3.6)

Finally, in section 3.7, different coupling algorithms are illustrated, and explanations are given on how to reproduce them with PSMILe, defining the appropriate indices of lag and sequence for each coupling field.

A practical example of a toy coupled model based on OASIS3 and PSMILe is given in `prism/util/mod/toyclim` (see the README_toyclim therein).

3.1 Initialisation

All processes initialise the coupling and, if required, retrieve a local communicator for the component model internal parallelisation.

- `USE mod_prism_proto`

  Module to be used by the component models. The MPI include file, mpif.h, also needs to be included in the component models (`INCLUDE 'mpif.h'`).

- `CALL prism_init_comp_proto (compid, model_name, ierror)`

  - `compid [INTEGER; OUT]`: component model ID
  - `model_name [CHARACTER*6; IN]`: name of the calling model (as in the *namcouple*)
  - `ierror [INTEGER; OUT]`: returned error code.

  Routine called by all component model processes, which initialises the coupling.[2]

- `CALL prism_get_localcomm_proto (local_comm, ierror )`

  - `local_comm [INTEGER; OUT]`: value of local communicator
  - `ierror [INTEGER; OUT]`: returned error code.

  Routine called by all model processes to get the value of a local communicator to be used by the model for its internal parallelisation. This is done following a "colouring scheme" developed by Reiner Vogelsang from Silicon Graphics GmbH. With MPI1, this call is mandatory for parallel models; with MPI2, the communicator MPI_COMM_WORLD will be returned as local communicator.

3.2 Partition definition

When a component of the coupled system is a parallel code, each coupling field is usually scattered among the different processes. With the PSMILe library, each process sends directly its partition to the OASIS3 main process or directly to the other component model if no transformation nor repartition is required. To do so, each process implied in the coupling has to define its local partition in the global index space.

- `CALL prism_def_partition_proto (il_part_id, ig_paral, ierror)`

  - `il_part_id [INTEGER; OUT]`: partition ID
  - `ig_paral [INTEGER, DIMENSION(:), IN]`: vector of integers describing the local partition in the global index space
  - `ierror [INTEGER; OUT]`: returned error code.

The vector of integers describing the process local partition (argument `ig_paral` of `prism_def_partition_proto`) has a different expression depending on the type of the partition. In OASIS3, 4 types of partition are supported: Serial (no partition), Apple, Box, and Orange.

[2] The model may call MPI_Init explicitly, but if so, has to call it before calling prism_init_comp_proto; in this case, the model also has to call MPI_Finalize explicitly, but only after calling prism_terminate_proto.

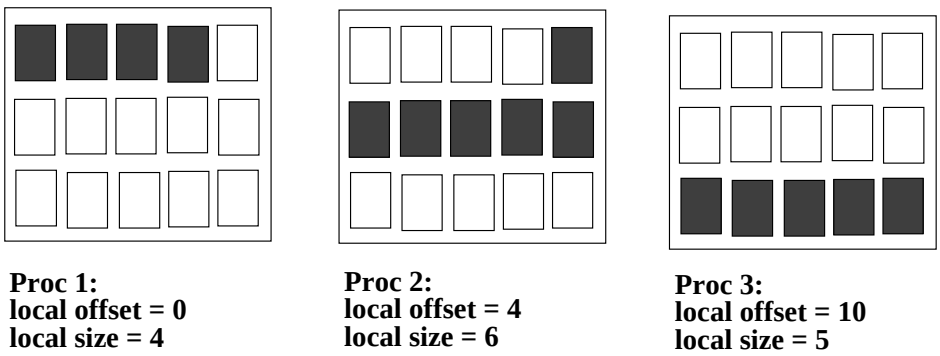

Figure 1: Apple partition

3.2.1 Serial (no partition)

This is the choice for a monoprocess model. In this case, we have `ig_paral(1:3)`:

- `ig_paral(1)` = 0 (indicates a Serial "partition")
- `ig_paral(2)` = 0
- `ig_paral(3)` = the total grid size.

3.2.2 Apple partition

Each partition is a segment of the global domain, described by its global offset and its local size. In this case, we have `ig_paral(1:3)`:

- `ig_paral(1)` = 1 (indicates an Apple partition)
- `ig_paral(2)` = the segment global offset
- `ig_paral(3)` = the segment local size

Figure 1 illustrates an Apple partition over 3 processes.

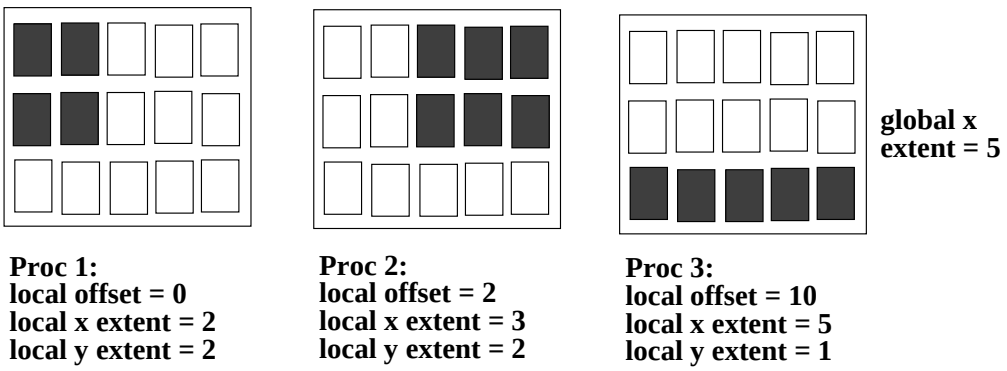

Figure 2: Box partition

3.2.3 Box partition

Each partition is a rectangular region of the global domain, described by the global offset of its upper left corner, and its local extents in the X and Y dimensions. The global extent in the X dimension must also be given. In this case, we have `ig_paral(1:5)`:

- `ig_paral(1)` = 2 (indicates a Box partition)
- `ig_paral(2)` = the upper left corner global offset
- `ig_paral(3)` = the local extent in X
- `ig_paral(4)` = the local extent in Y
- `ig_paral(5)` = the global extent in X.

Figure 2 illustrates a Box partition over 3 processes.

3.2.4 Orange partition

Each partition is an ensemble of segments of the global domain. Each segment is described by its global offset and its local extent. In this case, we have `ig_paral(1:N)` where `N = 2 + 2*number of segments`.

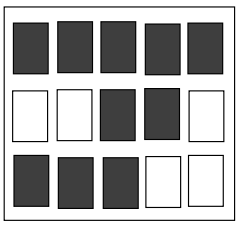

Figure 3: Orange partition for one process

- `ig_paral(1)` = 3 (indicates a Orange partition)
- `ig_paral(2)` = the total number of segments for the partition
- `ig_paral(3)` = the first segment global offset
- `ig_paral(4)` = the first segment local extent
- `ig_paral(5)` = the second segment global offset
- `ig_paral(6)` = the second segment local extent
- ...
- `ig_paral(N-1)` = the last segment global offset
- `ig_paral(N)` = the last segment local extent

Figure 3 illustrates an Orange partition with 3 segments for one process. The other process partitions are not illustrated.

3.3 I/O-coupling field declaration

Each process implied in the coupling declares each field it will send or receive during the simulation.

- `CALL prism_def_var_proto(var_id, name, il_part_id, var_nodims, kinout, var_actual_shape, var_type, ierror)`
  - `var_id [INTEGER; OUT]`: coupling field ID
  - `name [CHARACTER*8; IN]`: field symbolic name (as in the *namcouple*)
  - `il_part_id [INTEGER; IN]`: partition ID (returned by `prism_def_partition_proto`)
  - `var_nodims [INTEGER, DIMENSION(2); IN]`: var_nodims(1) is the rank of field array (1 or 2); var_nodims(2) is the number of bundles (always 1 for OASIS3).
  - `kinout [INTEGER; IN]`: `PRISM_In` for fields received by the model, or `PRISM_Out` for fields sent by the model
  - `var_actual_shape [INTEGER, DIMENSION(2*var_nodims(1)); IN]`: vector of integers giving the minimum and maximum index for each dimension of the coupling field array; for OASIS3, the minimum index has to be 1 and the maximum index has to be the extent of the dimension.
  - `var_type [INTEGER; IN]`: type of coupling field array, i.e. `PRISM_Real` or `PRISM_Double`. PRISM standard is to exchange coupling fields declared `REAL(kind=SELECTED_REAL_KIND(12,307))`. By default, all real variables are declared as such in OASIS3. In this case, both `PRISM_Real` or `PRISM_Double` are valid for `var_type`[3]. No automatic conversion is implemented; therefore, all coupling fields exchanged through OASIS3 main process must share the same kind[4].
  - `ierror [INTEGER; OUT]`: returned error code.

3.4 End of definition phase

Each process implied in the coupling closes the definition phase.

- `CALL prism_enddef_proto(ierror)`
  - ierror [INTEGER; OUT]: returned error code.

[3] To exchange single precision coupling fields, OASIS3 has to be compiled with the pre-compiling key use_realtype_single, the coupling fields must be declared `REAL(kind=SELECTED_REAL_KIND(6,37))` in the component models, and `var_type` must be `PRISM_Real`.

[4] Coupling fields exchanged directly between two component models still have to be either `PRISM_Real` or `PRISM_Double`, but can have a kind different from the ones exchanged through OASIS3 main process, as long as it is the same in both models.

3.5 Sending and receiving actions

3.5.1 Sending a coupling field

In the model time stepping loop, each process implied in the coupling sends its part of the I/O or coupling field.

- `USE mod_prism_put_proto`

  Module to be used by the component model to call `prism_put_proto`.

- `CALL prism_put_proto(var_id, date, field_array, info)`

  - `var_id [INTEGER; IN]`: field ID (from corresponding prism_def_var_proto)
  - `date [INTEGER; IN]`: number of seconds in the run at the beginning of the timestep
  - `field_array [REAL, IN]`: I/O or coupling field array
  - `info [INTEGER; OUT]`: returned info code i.e.
    * PRISM_Sent(=4) if the field was sent to another model (directly or via OASIS3 main process)
    * PRISM_LocTrans (=5) if the field was only used in a time transformation (not sent, not output)
    * PRISM_ToRest (=6) if the field was written to a restart file only
    * PRISM_Output (=7) if the field was written to an output file only
    * PRISM_SentOut (=8) if the field was both written to an output file and sent to another model (directly or via OASIS3 main process)
    * PRISM_ToRestOut (=9) if the field was written both to a restart file and to an output file.
    * PRISM_Ok (=0) otherwise and no error occurred.

This routine may be called by the model at each timestep. The sending is actually performed only if the time obtained by adding the field lag (see 3.7) to the argument `date` corresponds to a time at which it should be activated, given the coupling or I/O period indicated by the user in the namcouple (see section 4). A field will not be sent at all if its coupling or I/O period indicated in the *namcouple* is greater than the total run time.

If a local time transformation is indicated for the field by the user in the namcouple (INSTANT, AVERAGE, ACCUMUL, T_MIN or T_MAX, see section 5), it is automatically performed and the resulting field is finally sent at the coupling or I/O frequency.

For a coupling field with a positive lag (see 3.7), the OASIS3 restart file (see section 6.2) is now automatically written by the last `prism_put_proto` call of the run, if its argument `date` + the field lag corresponds to a coupling or I/O period. To force the writing of the field in its coupling restart file, one can use `prism_put_restart_proto` (see below).

This routine can use the standard blocking send MPI_Send (by default as before) or the buffered MPI_BSend (if pre-processed with CPP flag `key_BSend`) to send the coupling fields (see section 7).

3.5.2 Receiving a coupling field

In the model time stepping loop, each process implied in the coupling receives its part of the I/O-coupling field.

- `USE mod_prism_get_proto`

  Module to be used by the component model to call `prism_get_proto`.

- `CALL prism_get_proto(var_id, date, field_array, ierror)`

  - `var_id [INTEGER; IN]`: field ID (from corresponding prism_def_var_proto)
  - `date [INTEGER; IN]`: number of seconds in the run at the beginning of the timestep
  - `field_array [REAL, OUT]`: I/O or coupling field array
  - `info [INTEGER; OUT]`: returned info code
    * PRISM_Recvd(=3) if the field was received from another model (directly or via OASIS3 main process)
    * PRISM_FromRest (=10) if the field was read from a restart file only (directly or via OASIS3 main process)
    * PRISM_Input (=11) if the field was read from an input file only
    * PRISM_RecvOut (=12) if the field was both received from another model (directly or via OASIS3 main process) and written to an output file
    * PRISM_FromRestOut (=13) if the field was both read from a restart file (directly or via OASIS3 main process) and written to an output file
    * PRISM_Ok (=0) otherwise and no error occurred.

This routine may be called by the model at each timestep. The `date` argument is automatically analysed and the receiving action is actually performed only if `date` corresponds to a time for which it should be activated, given the period indicated by the user in the namcouple. A field will not be received at all if its coupling or I/O period indicated in the *namcouple* is greater than the total run time.

3.5.3 Auxiliary routines

- `CALL prism_put_inquire(var_id, date, info)`

  - `var_id [INTEGER; IN]`: field ID (from corresponding prism_def_var_proto)
  - `date [INTEGER; IN]`: number of seconds in the run at the beginning of the timestep
  - `info [INTEGER; OUT]`: returned info code.

This routine may be called at any time to inquire what would happen to the corresponding field (i.e. with same `var_id` and at same `date`) below the corresponding `prism_put_proto`. The possible value of the returned info code are as for `prism_put_proto`:

- PRISM_Sent(=4) if the field would be sent to another model (directly or via OASIS3 main process)
- PRISM_LocTrans (=5) if the field would be only used in a time transformation (not sent, not output)
- PRISM_ToRest (=6) if the field would be written to a restart file only
- PRISM_Output (=7) if the field would be written to an output file only
- PRISM_SentOut (=8) if the field would be both written to an output file and sent to another model (directly or via OASIS3 main process)
- PRISM_ToRestOut (=9) if the field would be written both to a restart file and to an output file.
- PRISM_Ok (=0) otherwise and no error occurred.

This is useful when the calculation of the corresponding `field_array` is CPU consuming and should be avoided if the field is not effectively used below the `prism_put_proto`.

- `CALL prism_put_restart(var_id, date, ierror)`
  - `var_id [INTEGER; IN]`: field ID (from corresponding prism_def_var_proto)
  - `date [INTEGER; IN]`: number of seconds in the run at the beginning of the timestep
  - `info [INTEGER; OUT]`: returned error code (should be PRISM_ToRest=6 if the restart writing was successful)

This routine forces the writing of the field with corresponding `var_id` in its coupling restart file (see section 6.2). If a time operation is specified for this field, the value of the field as calculated below the last `prism_put_proto` is written. If no time operation is specified, the value of the field transferred to the last `prism_put_proto` is written.

3.6 Termination

- `CALL prism_terminate_proto(ierror)`
  - `ierror [INTEGER; OUT]`: returned error code.

  Each process must terminate the coupling by calling `prism_terminate_proto`[5] (normal termination).
- `CALL prism_abort_proto(compid, routine_name, abort_message)`
  - `compid [INTEGER; IN]`: component model ID (from prism_init_comp_proto)
  - `routine_name; IN]`: name of calling routine
  - `abort_message; IN]`: message to be written out.

[5] If the process called `MPI_Init` (before calling `prism_init_comp_proto`), it must also call `MPI_Finalize` explicitly, but only after calling `prism_terminate_proto`.

If a process needs to abort (abnormal termination), it must do so by calling `prism_abort_proto`. This will ensure a proper termination of all processes in the coupled model communicator. This routine writes the name of the calling model, the name of the calling routine, and the message to the job standard output (stdout).

3.7 Coupling algorithms - SEQ and LAG concepts

Using PSMILe library, the user has full flexibility to reproduce different coupling algorithms, without modifying the component model codes themselves. In the component codes, the sending and receiving routines, respectively `prism_put_proto` and `prism_get_proto`, can be called at each model timestep, with the appropriate `date` argument giving the actual time (at the beginning of the timestep), expressed in "number of seconds since the start of the run". This `date` argument is automatically analysed by the PSMILe[6] and depending on the coupling period, the lag and sequencing indices (LAG and SEQ), chosen by the user for each coupling field in the configuration file *namcouple*, different coupling algorithms can be reproduced **without modifying anything in the component model codes themselves.** The lag and sequence concepts and indices are explained in more details here below. These mechanisms are valid for fields exchanged through OASIS3 main process and for fields exchanged directly between the component models.

3.7.1 The lag concept

If no lag index or if a lag index equal to 0 is given by the user in the *namcouple* for a particular coupling field, the sending or receiving actions will actually be performed, below the `prism_put_proto` called in the source model or below the `prism_get_proto` called in the target model respectively, each time the `date` arguments on both sides match an integer number of coupling periods.

To match a `prism_put_proto` called by the source model at a particular date with a `prism_get_proto` called by the target model at a different date, the user has to define in the *namcouple* an appropriate lag index, LAG, for the coupling field(see section 4). The value of the LAG index must be expressed in "number of seconds"; its value is automatically added to the `prism_put_proto` date value and the sending action is effectively performed when the sum of the date and the lag matches an integer number of coupling periods. This sending action is automatically matched, on the target side, with the receiving action performed when the `prism_get_proto` date argument equals the same integer number of coupling periods.

1. LAG concept first example

   A first coupling algorithm, exploiting the LAG concept, is illustrated on figure 4.

   On the 4 figures in this section, short black arrows correspond to `prism_put_proto` or `prism_get_proto` called in the component model that do not lead to any

[6] With the PIPE, SIPC, GMEM and previously with the CLIM communication techniques, no such analysis was performed. For PIPE, SIPC, and GMEM, the sending actions on the source side would automatically match the receiving actions on the target side on a FIFO (First In First Out) basis.

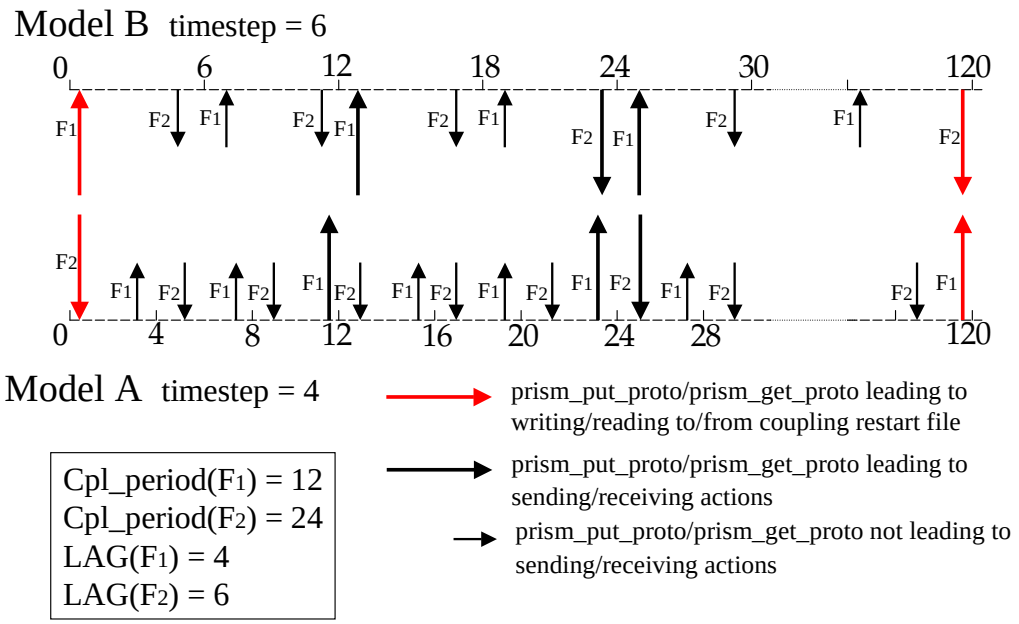

Figure 4: LAG concept first example

sending or receiving action; long black arrows correspond to `prism_put_proto` or `prism_get_proto` called in the component models that do effectively lead to a sending or receiving action; long red arrows correspond to `prism_put_proto` or `prism_get_proto` called in the component models that lead to a reading or writing of the coupling field from or to a coupling restart file (either directly or through OASIS3 main process).

During a coupling timestep, model A receives F_2 and then sends F_1; its timestep length is 4. During a coupling timestep, model B receives F_1 and then sends F_2; its timestep length is 6. F_1 and F_2 coupling periods are respectively 12 and 24. If F_1/F_2 sending action by model A/B was used at a coupling timestep to match the model B/A receiving action, a deadlock would occur as both models would be initially waiting on a receiving action. To prevent this, F_1 and F_2 produced at the timestep before have to be used to match respectively the model B and model A receiving actions.

This implies that a lag of respectively 4 and 6 seconds must be defined for F_1 and F_2. For F_1, the `prism_put_proto` performed at time 8 and 20 by model A will then lead to sending actions (as $8 + 4 = 12$ and $20 + 4 = 24$ which are coupling periods) that match the receiving actions performed at times 12 and 24 below the `prism_get_proto` called by model B. For F_2, the `prism_put_proto` performed at

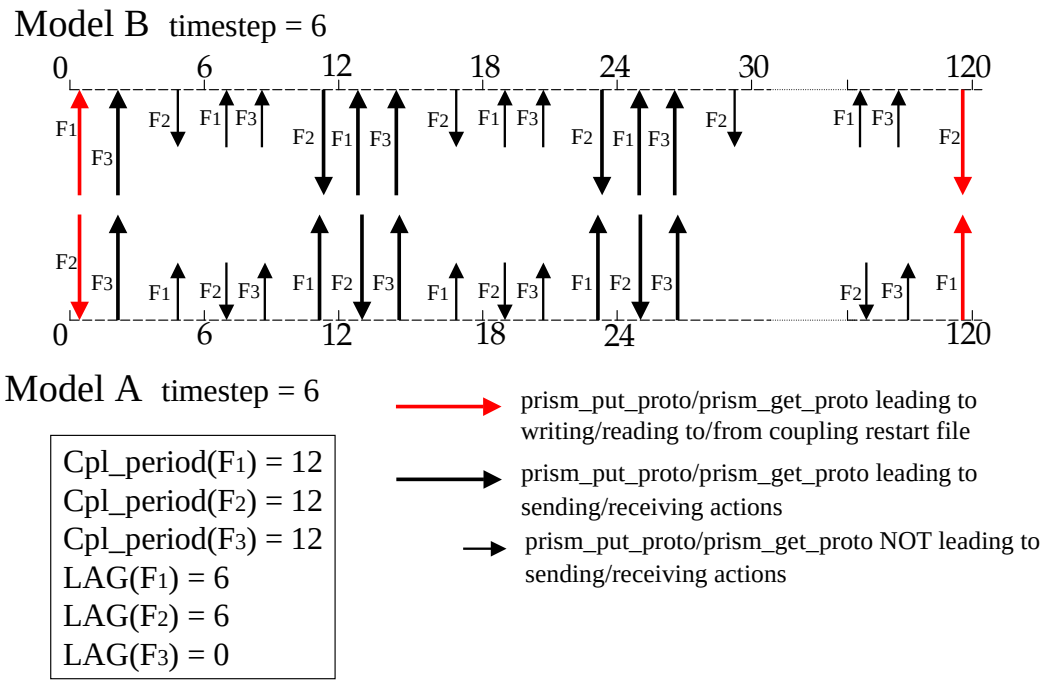

Figure 5: LAG concept second example

time 18 by model B then leads to a sending action (as 18 + 6 = 24 which is a coupling period) that matches the receiving action performed at time 24 below the `prism_get_proto` called by model A.

At the beginning of the run, as their LAG index is greater than 0, the first `prism_get_proto` will automatically lead to reading F_1 and F_2 from their coupling restart files. The user therefore have to create those coupling restart files for the first run in the experiment. At the end of the run, F_1 having a lag greater than 0, is automatically written to its coupling restart file below the last F_1 `prism_put_proto` if the `date` + F_1 lag equals a coupling time. The analogue is true for F_2. These values will automatically be read in at the beginning of the next run below the respective `prism_get_proto`.

2. LAG concept second example

A second coupling algorithm exploiting the LAG concept is illustrated on figure 5. During its timestep, model A receives F_2, sends F_3 and then F_1; its timestep length is 6. During its timestep, model B receives F_1, receives F_3 and then sends F_2; its timestep length is also 6. F_1, F_2 and F_3 coupling periods are both supposed to be equal to 12.

For F_1 and F_2 the situation is similar to the first example. If F_1/F_2 sending

action by model A/B was used at a coupling timestep to match the model B/A receiving action, a deadlock would occur as both models would be waiting on a receiving action. To prevent this, F_1 and F_2 produced at the timestep before have to be used to match the model A and model B receiving actions, which means that a lag of 6 must be defined for both F_1 and F_2. For both coupling fields, the `prism_put_proto` performed at times 6 and 18 by the source model then lead to sending actions (as $6 + 6 = 12$ and $18 + 6 = 24$ which are coupling periods) that match the receiving action performed at time 12 and 24 below the `prism_get_proto` called by the target model.

For F_3, sent by model A and received by model B, no lag needs to be defined: the coupling field produced by model A at the coupling timestep can be "consumed" by model B without causing a deadlock situation.

As in the first example, the `prism_get_proto` performed at the beginning of the run for F_1 and F_2, automatically read them from their coupling restart files, and the last `prism_put_proto` performed at the end of the run automatically write them to their coupling restart file. For F_3, no coupling restart file is needed nor used as at each coupling period the coupling field produced by model A can be directly "consumed" by model B.

We see here how the introduction of appropriate LAG indices results in receiving, below the `prism_get_proto` in the target model, coupling fields produced, below the `prism_put_proto` by the source model, the timestep before; this is, in some coupling configurations, essential to avoid deadlock situations.

3.7.2 The sequence concept

To exchange the coupling fields going through OASIS3 main process (i.e. with status EXPORTED, AUXILARY, or EXPOUT, see section 4), in a given order at each coupling timestep, a sequence index SEQ must be defined for each of them. This is not required for I/O fields or for coupling fields exchanged directly between the component models, i.e. with status IGNOUT, INPUT or OUTPUT. SEQ gives the position of the coupling field in the sequence.

A coupling algorithm, showing the SEQ concept, is illustrated on figure 6. All coupling field produced by the source model at the coupling timestep can be "consumed" by the target model at the same timestep without causing any deadlock situation; therefore, LAG = 0 for all coupling fields. However, at each coupling timestep, a particular order of exchange must be respected; F_1 must be received by model A before it can send F_2, which in turn must be received by model B before it can send F_3. Therefore, SEQ = 1, 2, 3 must be defined respectively for F_1, F_2 and F_3. As all fields can be consumed at the time they are produced (LAG=0 for all fields), there no reading/writing from/to coupling restart files.

3.7.3 A mix of lag and sequence: the sequential coupled model

One can run the same component models simultaneously or sequentially by defining the appropriate LAG and SEQ indices. In the example illustrated on figure 7, the models perform their `prism_put_proto` and `prism_get_proto` calls exactly as in the first lag example above: model A receives F_2 and then sends F_1; its timestep length

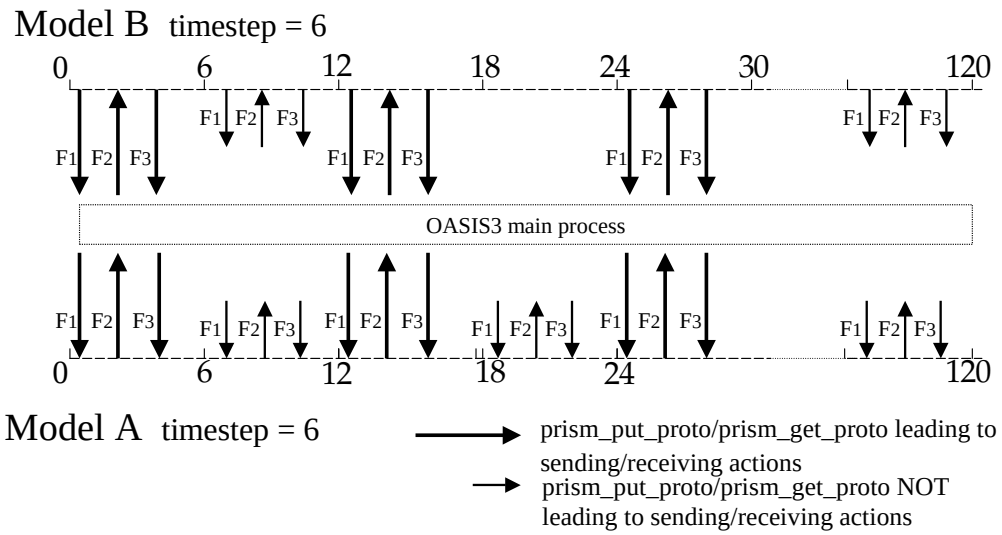

Figure 6: The SEQ concept

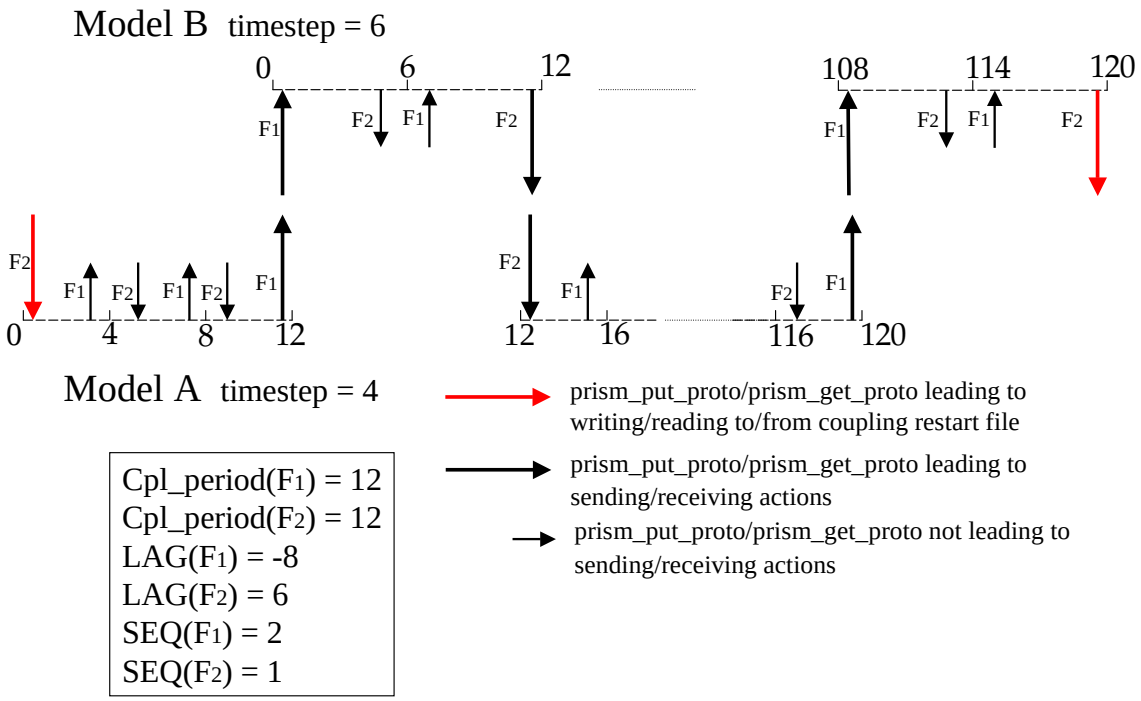

Figure 7: Mix of LAF and SEQ concepts

is 4. During a coupling timestep, model B receives F_1 and then sends F_2; its timestep length is 6. F_1 and F_2 coupling periods are both 12. By defining a LAG index of -8 for F_1, the models will now run sequentially.

As the LAG for F_2 is positive (6), a reading of F_2 in its coupling restart file is automatically performed below the initial `prism_get_proto`. As the LAG for F_1 is negative (-8), no reading from file is performed initially and model B waits; at time 8, a sending action is effectively performed below model A F_1 `prism_put_proto` (as 8 + LAG (-8) = 0 which is the first coupling timestep) and matches the initial model B F_1 `prism_get_proto`. Below the last model A F_1 `prism_put_proto` of the run at time 116, a sending action is effectively performed, as $116 + LAG(-8) = 108$ is a coupling period (as the LAG is negative, the field is not written to its coupling restart file). Below the last model B F_2 `prism_put_proto` of the run at time 114, a writing of F_2 to its restart file is performed, as $114 + LAG(6) = 120$ is a coupling period and as the LAG is positive.

If the coupling fields are transformed through OASIS3 main process, it is important to indicate a sequence index. In fact, at each OASIS3 main process coupling timestep, F_1 is necessarily treated after F_2. Therefore, $SEQ(F_1) = 2$ and $SEQ(F_2) = 1$.

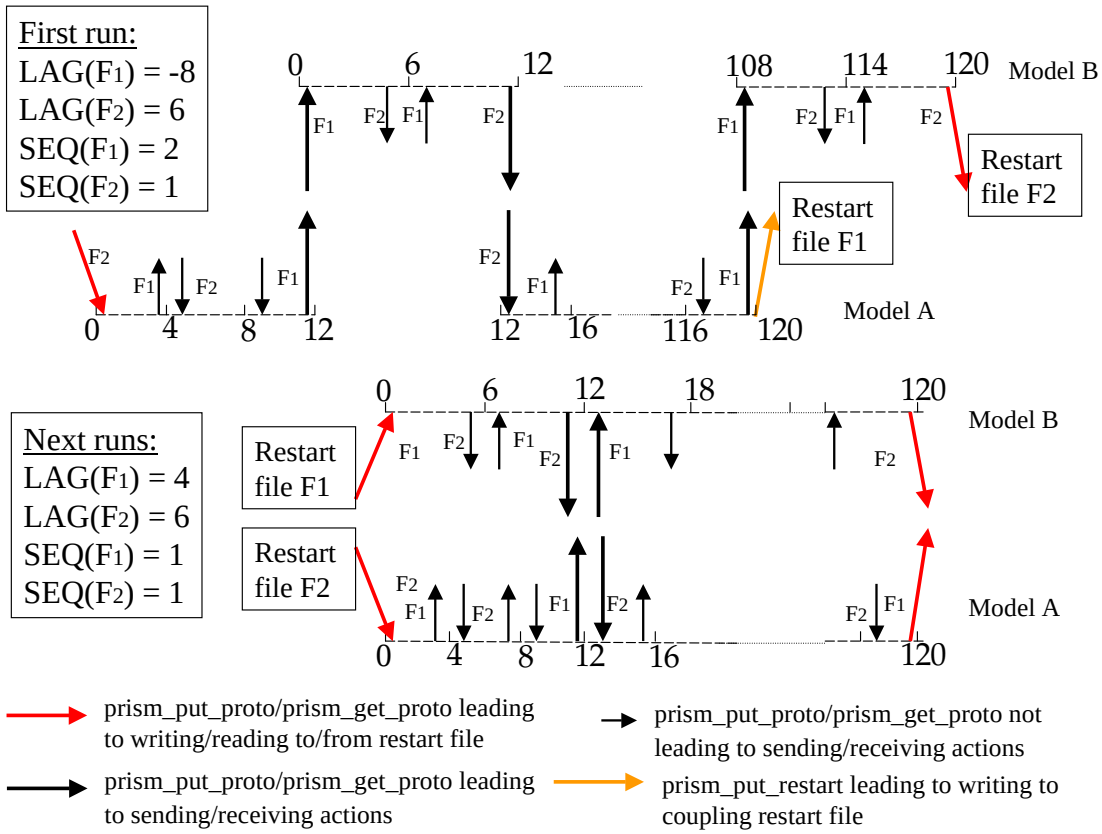

Figure 8: An example using prism_put_restart_proto

3.7.4 An experiment mixing sequential and parallel coupled model runs: the use of `prism_put_restart_proto`

In the example illustrated on figure 8, the models run sequentially for the first run only and then run simultaneously. For the first run, the LAG and SEQ indices must be defined as in section 3.7.3. After the first run, the situation is similar to the one of section 3.7.1, and positive LAG must be defined for F_1 and F_2. As their lag is positive, their corresponding first `prism_get_proto` will automatically lead to reading F_1 and F_2 from coupling restart files. In this case, model A has to write F_1 to its restart file explicitly by calling `prism_put_restart_proto` (illustrated on the figure by an orange arrow) at the end of the first run; in fact, F_1 lag being then negative, such writing is not automatically done below the last `prism_put_proto` of the first run.

4 The OASIS3 configuration file *namcouple*

The OASIS3 configuration file *namcouple* contains, below pre-defined keywords, all user's defined information necessary to configure a particular coupled run. The *namcouple* is a text file with the following characteristics:

- the keywords used to separate the information can appear in any order;
- the number of blanks between two character strings is non-significant;
- all lines beginning with # are ignored and considered as comments.

Note that **blank lines are not allowed.**

The first part of *namcouple* is devoted to configuration of general parameters such as the number of models involved in the simulation, the number of fields, the communication technique, etc. The second part gathers specific information on each coupling or I/O field, e.g. their coupling period, the list of transformations or interpolations to be performed by OASIS3 and their associated configuring lines (described in more details in section 5 , etc.

In the next sections, a simple *namcouple* example is given and all configuring parameters are described. The additional lines containing the different parameters required for each transformation are described in section 5. Examples of realistic *namcouple* can be found in directory `/prism/util/mod/toyclim/Cpl/Nam`.

4.1 An example of a simple *namcouple*

The following simple *namcouple* configures a run in which an ocean, an atmosphere and an atmospheric chemistry models are coupled. The ocean provides only the SOSSTSST field to the atmosphere, which in return provides the field CONSFTOT to the ocean. One field (COSENHFL) is exchanged directly from the atmosphere to the atmospheric chemistry, and one field (SOALBEDO) is read from a file by the ocean.

```
#############################################################################
# First section
#
 $SEQMODE
    1
#
 $CHANNEL
    MPI2
    1    1   arg1
    3    1   arg2
    3    1   arg3
#
 $NFIELDS
```

```
    4
#
 $JOBNAME
    JOB
#
 $NBMODEL
    3  ocemod   atmmod  chemod  55  70   99
#
 $RUNTIME
    432000
#
 $INIDATE
    00010101
#
 $MODINFO
    NOT
#
 $NLOGPRT
   2
#
 $CALTYPE
   1
#
#############################################################################
# Second section
#
 $STRINGS
#
# Field 1
#
 SOSSTSST SISUTESU 1 86400  5  sstoc.nc  EXPORTED
 182  149  128  64  toce  atmo   LAG=+14400  SEQ=+1
 P 2 P 0
 LOCTRANS CHECKIN MOZAIC  BLASNEW CHECKOUT
#
  AVERAGE
  INT=1
  at31topa   91    2   48
  CONSTANT     273.15
  INT=1
#
# Field 2
#
 CONSFTOT SOHEFLDO 6 86400  7   flxat.nc  EXPORTED
 atmo   toce  LAG=+14400  SEQ=+1
 P 0 P 2
 LOCTRANS  CHECKIN  SCRIPR CHECKOUT
```

```
#
  ACCUMUL
  INT=1
  CONSERV  LR  SCALAR  LATLON   10  FRACAREA   FIRST
  INT=1
#
# Field
#
 COSENHFL  SOSENHFL  37  86400   1  flda3.nc  IGNOUT
 LAG=+7200  SEQ=+1
 LOCTRANS
 AVERAGE
#
# Field 4
#
 SOALBEDO SOALBEDO  17  86400  0  SOALBEDO.nc  INPUT
#
#############################################################################
```

4.2 First section of *namcouple* file

The first section of *namcouple* uses some predefined keywords prefixed by the $ sign to locate the related information. The $ sign must be in the second column. The first ten keywords are described hereafter:

- `$SEQMODE`: On the line below this keyword is the maximum number of fields that have to be, at one particular coupling timestep, necessarily exchanged sequentially in a given order. For `$SEQMODE` ≥ 1, the position of each coupling field in the sequence has to be given by its SEQ index (see below and also section 3.7).

- `$CHANNEL`: On the line below this keyword is the communication technique chosen. Choices are `MPI1` or `MPI2` for the CLIM communication technique and related PSMILe library, using `MPI1` or `MPI2` message passing. To run OASIS3 as an interpolator only, put `NONE` (see also section 5.1). `PIPE`, `SIPC` or `GMEM` should still work but are not officially supported anymore and were not tested.

  To use the CLIM/MPI2 communication technique, the lines below `$CHANNEL` are, e.g. for 3 models:

  ```
  $CHANNEL
  MPI2
  1    1  arg1
  3    1  arg2
  3    1  arg3
  ```

  where `MPI2` is the message passing used in CLIM and PSMILe and the following lines (one line per model listed on the `$NBMODEL` line) indicate for each model the total number of processes, the number of processes implied in the coupling, and

possibly launching arguments. Here the first model runs on one process which is of course implied in the coupling and the argument passed to the model is "arg1"; the second and third models run on 3 processes but only one process is implied in the coupling (i.e. exchanging information with OASIS3 main process), and the argument passed to the models are respectively "arg2" and "arg3".

To use the CLIM/MPI1 communication technique, the `$CHANNEL` lines are e.g.:

```
$CHANNEL
MPI1
1    1
3    1
3    1
```

where `MPI1` is the message passing used in CLIM and PSMILe and the following lines (one line per model listed on the `$NBMODEL line`) are as for `MPI2` except that there is no launching argument.

To use the `PIPE`, `SIPC` or `GMEM` communication technique, the user has to write only `PIPE`, `SIPC` or `GMEM` below the `$CHANNEL` keyword.

- `$NFIELDS`: On the line below this keyword is the total number of fields exchanged and described in the second part of the *namcouple*.
- `$JOBNAME`: On the line below this keyword is a `CHARACTER*3` or `CHARACTER*4` variable giving an acronym for the given simulation.
- `$NBMODEL`: On the line below this keyword is the number of models running in the given experiment followed by `CHARACTER*6` variables giving their names. Then the user should indicate the maximum Fortran unit number used by the models. In the example, Fortran units above 55, 70, and 99 are free for respectively the ocean, atmosphere, and atmospheric chemistry models. If no maximum unit numbers are indicated, OASIS3 will suppose that units above 1024 are free. If `$CHANNEL` is `NONE`, `$NBMODEL` has to be 0 and there should be no model name and no unit number.
- `$RUNTIME`: On the line below this keyword is the total simulated time of the run, expressed in seconds.
- `$INIDATE`: On the line below this keyword is the initial date of the run. The format is `YYYYMMDD`. This date is important only for the `FILLING` transformation and for printing information in OASIS3 log file *cplout*.
- `$MODINFO`: For SIPC and GMEM techniques only, the line below this keyword indicates if a header is encapsulated in the binary coupling restart file (see section 6.2): it can be `YES` or `NOT`.
- `$NLOGPRT`: The line below this keyword refers to the amount of information that will be written to the OASIS3 log file *cplout* during the run. With 0, there is practically no output written to the cplout; with 1, only some general information on the run, the header of the main routines, and the names of the fields when treated appear in the *cplout*. Finally, with 2, the full output is generated.

- `$CALTYPE`: This new keyword gives the type of calendar used. For now, the calendar type is important only if `FILLING` analysis is used for a coupling field in the run and for printing information in OASIS3 log file *cplout*. Below this keyword, a number (0, 1 or n) must be indicated by the user:
  - 0 : a 365 day calendar (no leap year)
  - 1 : a 365 or 366 (leap years) day calendar A year is a leap year if it can be divided by 4; however if it can be divided by 4 and 100, it is not a leap year; furthermore, if it can be divided by 4, 100 and 400, it is a leap year.
  - n : $n \geq 1$ day month calendar.

4.3 Second section of *namcouple* file

The second part of the *namcouple*, starting after the keyword `$STRINGS`, contains coupling information for each coupling or I/O field. Its format depends on the field status given by the last entry on the field first line (`EXPORTED`, `IGNOUT` or `INPUT` in the example above). The field status may be:

- `AUXILARY`: sent by the source model, received and used by OASIS3 main process for the transformation of other fields; this choice is supported for all communication techniques.
- `EXPORTED`: exchanged between component models and transformed by OASIS3 main process; this choice is supported for all communication techniques.
- `EXPOUT`: exchanged, transformed and also written to two output files, one before the sending action in the source model below the `prism_put_proto` call, and one after the receiving action in the target model below the `prism_get_proto` call; this is supported only by the `CLIM/MPI1-MPI2` communication technique and its associated PSMILe.
- `IGNORED`: exchanged directly between the component models without being transformed by OASIS3 main process. The grid and partitioning of the source and target models have to be identical. This is supported only by the `CLIM/MPI1-MPI2` communication technique and its associated PSMILe.
- `IGNOUT`: exchanged directly between the component models without being transformed by OASIS3 main process and written to two output files, one before the sending action in the source model below the `prism_put_proto` call, and one after the receiving action in the target model below the `prism_get_proto` call. The grid and partitioning of the source and target models have to be identical. This is supported only by the `CLIM/MPI1-MPI2` communication technique and its associated PSMILe.
- `INPUT`: simply read in from the input file by the target model PSMILe below the `prism_get_proto` call at appropriate times corresponding to the input period indicated by the user in the *namcouple*. This is supported only by the `CLIM/MPI1-MPI2` communication technique and its associated PSMILe. See section 6.3 for the format of the input file.

- OUTPUT: simply written out to an output file by the source model PSMILe below the prism_put_proto call at appropriate times corresponding to the output period indicated by the user in the *namcouple*. The name of the output file (one per field) is automatically built based on the field name and initial date of the run ($INIDATE). This is supported only by the CLIM/MPI1-MPI2 communication technique and its associated PSMILe.

4.3.1 Second section of *namcouple* for EXPORTED, AUXILARY and EXPOUT fields

The first 3 lines for fields with status EXPORTED, AUXILARY and EXPOUT are as follows:

```
SOSSTSST SISUTESU 1 86400  5  sstoc.nc  sstat.nc EXPORTED
182  149    128  64  toce  atmo   LAG=+14400 SEQ=+1
P 2 P 0
```

where the different entries are:

- Field first line:
  - SOSSTSST : symbolic name for the field in the source model (CHARACTER*8). It has to match the argument name of the corresponding field declaration in the source model; see prism_def_var_proto in section 3.3.
  - SISUTESU : symbolic name for the field in the target model (CHARACTER*8). It has to match the argument name of the corresponding field declaration in the target model; see prism_def_var_proto in section 3.3.
  - 1 : internal label used by OASIS3 for corresponding CF standard name (see prism/src/mod/oasis3/src/inipar.F).
  - 86400 : coupling and/or I/O period for the field, in seconds.
  - 5 : number of transformations to be performed on this field.
  - sstoc.nc : name of the coupling restart file for the field (CHARACTER*8); it may be a binary of netCDF file (for more detail, see section 6.2).
  - sstat.nc : name of the field output file, required for NONE (and PIPE) communication techniques only. It may be a binary of netCDF file (see section 6.2).
  - EXPORTED : field status.
- Field second line:
  - 182 : number of points for the source grid first dimension (optional if a netCDF coupling restart file is used).
  - 149 : number of points for the source grid second dimension (optional if a netCDF coupling restart file is used).
  - 128 : number of points for the target grid first dimension (optional if a netCDF coupling restart file is used).
  - 64 : number of points for the target grid second dimension (optional if a netCDF coupling restart file is used).

  - toce : locator prefix (CHARACTER*4) for the source grid parameters in the grid data files (see section 6.1).
  - atmo : locator prefix (CHARACTER*4) for the target grid parameters in the grid data files (see section 6.1).
  - LAG=+14400: (optional) lag index for the field expressed in seconds (see section 3.7).
  - SEQ=+1: (optional) sequence index for the field expressed in seconds. (see section 3.7)

- Field third line
  - P : source grid first dimension characteristic ('P': periodical; 'R': regional).
  - 2 : source grid first dimension number of overlapping grid points.
  - P : target grid first dimension characteristic ('P': periodical; 'R': regional).
  - 0 : target grid first dimension number of overlapping grid points.

The fourth line gives the list of transformations to be performed for this field. There is then one or more additional configuring lines describing some parameters for each transformation. These additional lines are described in more details in the section 5.

4.3.2 Second section of *namcouple* for IGNORED and IGNOUT fields

The first 2 lines for fields with status IGNORED and IGNOUT fields are as follows:

```
COSENHFL SOSENHFL 37 86400 1 flda3.nc IGNOUT
LAG=+7200 SEQ=+1
```

where the different entries are as for EXPORTED fields, except that there is no output file name on the first line, and that there are no locator prefix on the second line. Note that a second line is mandatory; put at least LAG=0.

The third line is LOCTRANS if this transformation is chosen for the field. Note that LOCTRANS is the only transformation supported for IGNORED and IGNOUT fields (as it is performed directly in the PSMILe below the prism_put_proto call). If LOCTRANS is chosen, a fourth line giving the name of the time transformation is required.

4.3.3 Second section of *namcouple* for INPUT fields

The first and only line for fields with status INPUT is:

```
SOALBEDO SOALBEDO  17  86400  0  SOALBEDO.nc  INPUT
```

where the different entries are:

- SOALBEDO: symbolic name for the field in the target model (CHARACTER*8 repeated twice)
- 17: OASIS3 internal label

- 86400: input period in seconds
- 0: number of transformations (always 0 for `INPUT` fields)
- `SOALBEDO.nc`: `CHARACTER*32` giving the input file name (for more detail on its format, see section 6.3)
- `INPUT`: field status.

4.3.4 Second section of *namcouple* for `OUTPUT` fields

The first two lines for fields with status `OUTPUT` is:

```
SYMBNAME SYMBNAME  99  86400  0   OUTPUT
LAG=+0
```

where the different entries are:

- `SYMBNAME`: symbolic name for the field in the source model (`CHARACTER*8` repeated twice)
- 99: OASIS3 internal label
- 86400: output period in seconds
- 1: number of transformations (0 or 1 for `LOCTRANS`)
- `OUTPUT`: field status

Note that a second line is mandatory; put at least LAG=0.

The third line is `LOCTRANS` if this transformation is chosen for the field. Note that `LOCTRANS` is the only transformation supported for `OUTPUT` fields (as it is performed directly in the PSMILe below the `prism_put_proto` call). If `LOCTRANS` is chosen, a fourth line giving the name of the time transformation is required.

5 The transformations and interpolations in OASIS3

Different transformations and 2D interpolations are available in OASIS3 to adapt the coupling fields from a source model grid to a target model grid. They are divided into five general classes that have precedence one over the other in the following order: time transformation (with `CLIM/MPI1-MPI2` and PSMILe only), pre-processing, interpolation, "cooking", and post-processing. This order of precedence is conceptually logical, but is also constrained by the OASIS3 software internal structure.

In the following paragraphs, it is first described how to use OASIS3 in an interpolator-only mode. Then a description of each transformation with its corresponding configuring lines is given.

5.1 Using OASIS3 in the interpolator-only mode

OASIS3 can be used in an interpolator-only mode, in which case it transforms fields without running any model. It is recommended to use OASIS3 in this mode before hand to test different transformations and interpolations without having to run the whole coupled system. In the interpolator-only mode, all transformations, except the time transformations, are available.

To run OASIS3 in an interpolator-only mode, the user has to prepare the *namcouple* as indicated in sections 4.2 and 4.3. In particular, `NONE` has to be chosen below the keyword `$CHANNEL`, and `0` (without any model name and Fortran unit number) below the keyword `$NBMODEL`. Finally, the fields to transform and interpolate have to be present in their coupling restart file (6th entry on the first line). After their transformation, OASIS3 writes them to their output file which name is the 7th entry on the first line. Note that, in the interpolator-only mode, OASIS3 may treat different fields but only one temporal instance for each. A practical example on how to use OASIS3 in a interpolator-only mode is given in `prism/util/mod/toymodel/testscrip`.

The configuring parameters that have to be defined in the *namcouple* for each transformation in the interpolator-only mode or in the coupling mode are described here after.

5.2 The time transformations

`LOCTRANS` can be chosen as first transformation if CLIM/MPI1-MPI2 communication and the PSMILe interface are used. `LOCTRANS` requires one configuring line on which a time transformation, automatically performed below the call to PSMILe `prism_put_proto`, should be indicated:

- `INSTANT`: no time transformation, the instantaneous field is transferred;
- `ACCUMUL`: the field accumulated over the previous coupling period is transferred;
- `AVERAGE`: the field averaged over the previous coupling period is transferred;

- T_MIN: the minimum value of the field for each source grid point over the previous coupling period is transferred;
- T_MAX: the maximum value of the field for each source grid point over the previous coupling period is transferred;
- ONCE: only one prism_put_proto or prism_get_proto will be performed; this is equivalent to giving the length of the run as coupling or I/O period.

5.3 The pre-processing transformations

The following transformations are available in the pre-processing part of OASIS3, controlled by preproc.f.

- **REDGLO** (not recommended anymore as interpolations now exist directly for Reduced grids):

  REDGLO (routine redglo.f) performs the interpolation from a Reduced grid to a Gaussian one. The interpolation is linear and performed latitude circle per latitude circle. When present, REDGLO must be the first pre-processing transformation performed. The configuring line is as follows:

```
# REDGLO operation
    $NNBRLAT  $CDMSK
```

  where $NNBRLAT is NOxxx where xxx is half the number of latitude circles of the Gaussian grid. For example, for a T42 with 64 latitude circles, $NNBRLAT is "NO32". In the current version, it can be either NO16, NO24, NO32, NO48, NO80, NO160.$CDMSK is a flag indicating if non-masked values have to be extended to masked areas before interpolation ($CDMSK = SEALAND) using the Reduced grid mask (see section 6.1) or if the opposite has to be performed ($CDMSK = LANDSEA). If $CDMSK = NOEXTRAP, then no extrapolation is performed.

- **INVERT**:

  INVERT (routine invert.f) reorders a field so that it goes from south to north and from west to east (the first point will be the southern and western most one; then it goes parallel by parallel going from south to north). INVERT should be used only for fields associated to A, B, G, L, Z, or Y grids (see annexe A) but produced by the source model from North to South and/or from East to West. INVERT does not work for Reduced ('D') or unstructured ('U') grids (see annexe A).

  The generic input line is as follows:

```
# INVERT operation
    $CORLAT  $CORLON
```

  where $CORLAT = NORSUD or SUDNOR and $CORLON = ESTWST or WSTEST describes the orientation of the source field in longitude and latitude, respectively.

- **MASK**:

  MASK (routine masq.f) is used before the analysis EXTRAP. A given REAL value VALMASK is assigned to all masked points following the source grid mask (see section 6.1), so they can be detected by EXTRAP.

  The generic input line is as follows:

```
# MASK operation
    $VALMASK
```

  Problems may arise if the value chosen approaches the maximum value that your computing platform can represent; choose a value well outside the range of your field values but not too large.

- **EXTRAP**:

  EXTRAP (routine extrap.f) performs the extrapolation of a field over its masked points. The analysis MASK must be used just before, so that EXTRAP can identify masked points. Note that EXTRAP does not work for Reduced ('D') or unstructured ('U') grids (see section A).

  Two methods of extrapolation are available. With NINENN, a N-nearest-neighbour method is used. The procedure is iterative and the set of remaining masked points evolves at each iteration. The configuring line is:

```
# EXTRAP operation for $CMETHOD = NINENN
    $CMETHOD  $NVOISIN  $NIO  $NID
```

  where $CMETHOD = NINENN; $NVOISIN is the minimum number of neighbours required to perform the extrapolation (with a maximum of 4)[7]; $NIO is the flag that indicates if the weight-address-and-iteration-number dataset will be calculated and written by OASIS3 ($NIO= 1), or only read ($NIO= 0) in file *nweights* (see section 6.4); $NID is the identificator for the weight-address-iteration-number dataset in all the different EXTRAP/NINENN datasets in the present coupling.[8]

  With $CMETHOD = WEIGHT, an N-weighted-neighbour extrapolation is performed. In that case, the user has to build the grid-mapping file, giving for each target grid point the weights and addresses of the source grid points used in the extrapolation; the structure of this file has to follow the OASIS3 generic structure for transformation auxiliary data files (see section 6.4).

  The configuring line is:

```
# EXTRAP operation for $CMETHOD = WEIGHT
    $CMETHOD  $NVOISIN  $CFILE  $NUMLU  $NID
```

[7] For some grids, the extrapolation may not converge if $NVOISIN is too large.

[8] An EXTRAP/NINENN analysis is automatically performed within GLORED analysis but the corresponding datasets have to be distinct; this is automatically checked by OASIS3 at the beginning of the run.

where $CMETHOD = WEIGHT; $NVOISIN is the maximum number of neighbours required by the extrapolation operation; $CFILE and $NUMLU are the grid-mapping file name and associated logical unit; $NID is the identificator for the relevant grid-mapping dataset in all different EXTRAP/WEIGHT transformations in the present coupling.

- **CHECKIN**:

  CHECKIN (routine chkfld.f) calculates the mean and extremum values of the source field and prints them to the coupler log file *cplout*.

  The generic input line is as follows:

```
# CHECKIN operation
    $NINT
```

  where $NINT = 1 or 0, depending on whether or not the source field integral is also calculated and printed.

- **CORRECT**:

  CORRECT (routine correct.f) reads external fields from binary files and uses them to modify the coupling field. This transformation can be used, for example, to perform flux correction on the field.

  This transformation requires at least one configuration line with two parameters:

```
# CORRECT operation
    $XMULT  $NBFIELDS
```

  where $XMULT is the multiplicative coefficient of the current field, and $NBFIELDS the number of additional fields to be combined with the current field. For each additional field, an additional configuring line is required:

```
# nbfields lines
    $CLOC  $AMULT  $CFILE  $NUMLU
```

  where $CLOC and $AMULT, $CFILE and $NUMLU are respectively the symbolic name, the multiplicative coefficient, the file name and the associated logical unit on which the additional field is going to be read. The structure of the file has to follow the structure of OASIS3 binary coupling restart files (see section 6.2).

5.4 The interpolation

The following interpolations, controlled by interp.f, are available in OASIS3.

- **BLASOLD**:

  BLASOLD (routine blasold.f) performs a linear combination of the current coupling field with other coupling fields or with a constant before the interpolation *per se*.

  This transformation requires at least one configuring line with two parameters:

```
# BLASOLD operation
    $XMULT     $NBFIELDS
```

where $XMULT is the multiplicative coefficient of the current field, and $NBFIELDS the number of additional fields to be combined with the current field. For each additional field, an additional input line is required:

```
# nbfields lines
    $CNAME  $AMULT
```

where $CNAME and $AMULT are the symbolic name and the multiplicative coefficient for the additional field. To add a constant value to the original field, put $XMULT = 1, $NBFIELDS = 1, $CNAME = CONSTANT, $AMULT = value to add.

- **SCRIPR**:

  SCRIPR is new in OASIS3 and gathers the interpolation techniques offered by Los Alamos National Laboratory SCRIP 1.4 library[9][1]. SCRIPR routines are in prism/src/lib/scrip. For more details, see also the SCRIP 1.4 documentation in prism/src/mod/oasis3/doc/SCRIPusers.pdf. Linking with NetCDF library is mandatory when using SCRIPR interpolations.

  The following types of interpolations are available:

  - DISTWGT performs N nearest-neighbour interpolation (N neighbours). All grid types are supported. The configuring line is:

    ```
    # SCRIPR/DISWGT
        $CMETHOD  $CGRDSRC  $CFLDTYP  $RESTRIC  $NBRBINS $NVOISIN
    ```

    where:

    * $CMETHOD = DISTWGT.
    * $CGRDSRC is the source grid type (LR, D or U)- see annexe A.
    * $CFLDTYP is the field type (SCALAR or VECTOR[10]).
    * $RESTRIC is the search restriction type: LATLON or LATITUDE (see SCRIP 1.4 documentation).
    * $NBRBINS the number of restriction bins (see SCRIP 1.4 documentation B).
    * $NVOISIN is the number of neighbours used.

  - GAUSWGT performs N nearest-neighbour interpolation weighted by a gaussian function. All grid types are supported. The configuring line is:

    ```
    # SCRIPR/GAUSWGT
        $CMETHOD  $CGRDSRC  $CFLDTYP  $RESTRIC  $NBRBINS  $NVOISIN $VAR
    ```

    where all entries are as for DISTWGT, except that:

    * $CMETHOD = GAUSWGT

[9] See the copyright statement in annexe B.

[10] No particular treatment for vector interpolation is performed presently.

  * $VAR, which must be given as a REAL value (e.g 2.0 and not 2), defines the weight given to a neighbour source grid point as inversely proportional to $exp(-1/2 \cdot d^2/\sigma^2)$ where d is the distance between the source and target grid points, and $\sigma^2 = \$VAR \cdot \overline{d}^2$ where $\overline{d}^2$ is the average distance between two source grid points (calculated automatically by OASIS3).

- BILINEAR performs bilinear interpolation.
- BICUBIC performs a bicubic interpolation.

  For BILINEAR and BICUBIC, Logically-Rectangular (LR) and Reduced (D) source grid types are supported and the configuring line is:

```
# SCRIPR/BILINEAR or SCRIPR/BICUBIC
    $CMETHOD  $CGRDSRC  $CFLDTYP  $RESTRIC  $NBRBINS
```

  where:

  * $CMETHOD = BILINEAR or BICUBIC
  * $CGRDSRC is the source grid type (LR or D)
  * $CFLDTYP and $NBRBINS are as for DISTWGT.
  * $RESTRIC is as for DISTWGT, except that only LATITUDE is possible for a Reduced (D) source grid.

- CONSERV performs 1st or 2nd order conservative remapping, which means that the weight of a source cell is proportional to area intersected by target cell.

  The configuring line is:

```
 # SCRIPR/CONSERV
$CMETHOD  $CGRDSRC  $CFLDTYP  $RESTRIC  $NBRBINS  $NORMA  $ORDER
```

  where:

  * $CMETHOD = CONSERV
  * $CGRDSRC is the source grid type: LR, D and U are supported for 1st-order remapping if the grid corners are given by the user in the grid data file which is, in this case, necessarily a netCDF file (grids.nc, see section 6.1); only LR is supported if the grid corners are not available in the grid data file and therefore have to be calculated automatically by OASIS3. For second-order remapping, only LR is supported because the gradient of the coupling field used in the transformation has to be calculated automatically by OASIS3.
  * $CFLDTYP, $RESTRIC and $NBRBINS are as for DISTWGT.
  * $NORMA is the normalization option:
    - FRACAREA: The sum of the source cell intersected areas is used to normalise each target cell field value: the flux is not locally conserved, but the flux value itself is reasonable.
    - DESTAREA: The total target cell area is used to normalise each target cell field value even if it only partly intersects the source grid cells:

local flux conservation is ensured, but unreasonable flux values may result.

  - FRACNNEI: as FRACAREA, except that the source nearest neighbour is used for target cells that do not intersect any unmasked source cells.

* $ORDER: FIRST or SECOND[11] for first or second order remapping respectively (see SCRIP 1.4 documentation).

- **INTERP**:

  INTERP gathers different techniques of interpolation controlled by routine fiasco.f. The following interpolations are available:

  – BILINEAR performs a bilinear interpolation using 4 neighbours.

  – BICUBIC performs a bicubic interpolation.

  – NNEIBOR performs a nearest-neighbour interpolation.

    These three interpolations are performed by routines in /prism/src/lib/fscint and support only A, B, G, L, Y, or Z grids (see annexe A).

    The configuring line is as follows:

```
# BILINEAR or BICUBIC or NNEIBOR interpolation
   $CMETHOD $CGRDSRC $CFLDTYP
```

    where

    * $CMETHOD = BILINEAR, BICUBIC or NNEIBOR
    * $CGRDSRC is the source grid type (A, B, G, L, Y, or Z, see annexe A)
    * $CFLDTYP the field type (SCALAR or VECTOR). VECTOR has an effect for target grid points located near the pole: the sign of a source value located on the other side of the pole will be reversed.

  – SURFMESH (routines in /prism/src/lib/anaism) is a first-order conservative remapping from a fine to a coarse grid (the source grid must be finer over the whole domain) and supports only Lat-Lon grids. For a target grid cell, all the underlying source grid cells are found and the target grid field value is the sum of the source grid field values weighted by the overlapped surfaces. No value is assigned to masked cells. Note that it is not recommended to use this interpolation anymore, as the more general SCRIPR/CONSERV remapping is now available. The configuring line is as follows:

```
# SURFMESH remapping
    $CMETHOD $CGRDSRC $CFLDTYP $NID $NVOISIN $NIO
```

    where

    * $CMETHOD = SURFMESH
    * $CGRDSRC and $CFLDTYP are as for BILINEAR

[11] CONSERV/SECOND has not been tested in detail.

    * $NID is the identificator for the weight-address dataset in all the different INTERP/SURFMESH datasets in the present coupling. This dataset will be calculated by OASIS3 if $NIO= 1, or only read if $NIO= 0.
    * $NVOISIN is the maximum number of source grid meshes used in the remapping.

  – GAUSSIAN (routines in /prism/src/lib/anaisg) is a gaussian weighted nearest-neighbour interpolation technique. The user can choose the variance of the function and the number of neighbours considered. The configuring line is:

```
# GAUSSIAN interpolation
$CMETHOD $CGRDSRC $CFLDTYP $NID $NVOISIN $VAR  $NIO
```

    where

    * $CMETHOD = GAUSSIAN
    * $CGRDSRC and $CFLDTYP are as for the BILINEAR interpolation.
    * $NID is the identificator for the weight-address dataset in all the different INTERP/GAUSSIAN datasets in the present coupling. This weight-address dataset will be calculated by OASIS3 if $NIO= 1, or only read if $NIO= 0.
    * $NVOISIN is the number of neighbours used in the interpolation.
    * $VAR is as for SCRIPR/GAUSWGT (see above).

- **MOZAIC**:

  MOZAIC performs the mapping of a field from a source to a target grid. The grid-mapping dataset, i.e. the weights and addresses of the source grid points used to calculate the value of each target grid point are defined by the user in a file (see section 6.4). The configuring line is:

```
# MOZAIC operation
     $CFILE  $NUMLU  $NID  $NVOISIN
```

  where

  – $CFILE and $NUMLU are the grid-mapping file name and associated logical unit on which the grid-mapping dataset is going to be read),
  – $NID the identificator for this grid-mapping dataset in all MOZAIC grid-mapping datasets in the present coupling
  – $NVOISIN is the maximum number of target grid points use in the mapping.

- **NOINTERP**:

  NOINTERP is the analysis that has to be chosen when no other transformation from the interpolation class is chosen. There is no configuring line.

- **FILLING**:

  `FILLING` (routine `/prism/src/mod/oasis3/src/filling.f`) performs the blending of a regional data set with a climatological global one for a Sea Surface Temperature (SST) or a Sea Ice Extent field. This occurs when coupling a non-global ocean model with a global atmospheric model. `FILLING` can only handle fields on Logically Rectangular grid (LR, but also A, B, G, L, Y, and Z grids, see section A.

  The global data set has to be a set of SST given in Celsius degrees (for the filling of a Sea Ice Extent field, the presence or absence of ice is deduced from the value of the SST). The frequency of the global set can be interannual monthly, climatological monthly or yearly.

  The blending can be smooth or abrupt. If the blending is abrupt, only model values are used within the model domain, and only the global data set values are used outside. If the blending is smooth, a linear interpolation is performed between the two fields within the model domain over narrow bands along the boundaries. The linear interpolation can also be performed giving a different weight to the regional or and global fields.

  The smoothing is defined by parameters in `/prism/src/mod/oasis3/src/mod_smooth.F90`. The lower smoothing band in the global model second dimension is defined by *nsltb* (outermost point) and *nslte* (innermost point); the upper smoothing band in the global model second dimension is defined by *nnltb* (outermost point) and *nnlte* (innermost point). The parameter *qalfa* controls the weights given to the regional and to the global fields in the linear interpolation. *qalfa* has to be $1/(nslte - nsltb)$ or $1/(nnltb - nnlte)$. For the outermost points (*nsltb* or *nnltb*) in the smoothing band, the weight given to the regional and global fields will respectively be 0 and 1; for the innermost points (*nslte* or *nnlte*) in the smoothing band, the weight given to the regional and global fields will respectively be 1 and 0; within the smoothing band, the weights will be a linear interpolation of the outermost and innermost weights.

  The smoothing band in the global model first dimension will be a band of *nliss* points following the coastline. To calculate this band, OASIS3 needs *nwlgmx*, the greater first dimension index of the lower coastline and *nelgmx*, the smaller first dimension index on the upper coastline. The parameter *qbeta* controls the weights given to the regional and to the global fields in the linear interpolation. *qbeta* has to be $1/(nliss - 1)$. The weights given to the regional and global fields in the global model first dimension smoothing bands will be calculated as for the second dimension.

  The user must provide the climatological data file with a specific format described in 6.4. When one uses `FILLING` for SST with smooth blending, thermodynamics consistency requires to modify the heat fluxes over the blending regions. The correction term is proportional to the difference between the blended SST and the original SST interpolated on the atmospheric grid and can be written out on a file to be read later, for analysis `CORRECT` for example. In that case, the symbolic name of the flux correction term read through the input file *namcouple* must correspond in `FILLING` and `CORRECT` analyses.

In case the regional ocean model includes a coastal part or islands, a sea-land mask mismatch may occur and a coastal point correction can be performed if the field has been previously interpolated with `INTER/SURFMESH`. In fact, the mismatch could result in the atmosphere undesirably "seeing" climatological SST's directly adjacent to ocean model SST's. Where this situation arises, the coastal correction consists in identifying the suitable ocean model grid points that can be used to extrapolate the field, excluding climatological grid points.

This analysis requires one configuring line with 3, 4 or 6 arguments.

1. If FILLING performs the blending of a regional data set with a global one for the Sea Ice Extent, the 3-argument input line is:

   ```
   # Sea Ice Extent FILLING operation
        $CFILE  $NUMLU  $CMETHOD
   ```

   where `$CFILE` is the file name for the global data set, `$NUMLU` the associated logical unit. `$CMETHOD`, the `FILLING` technique, is a `CHARACTER*8` variable: the first 3 characters are either `SMO`, smooth filling, or `RAW`, no smoothing ; the next three characters must be `SIE` for a Sea Ice Extent filling operation; the last two define the time characteristics of the global data file, respectively `MO`, `SE` and `AN` for interannual monthly, climatological monthly and yearly. Note that in all cases, the global data file has to be a Sea Surface Temperature field in Celsius degrees.

2. If FILLING performs the blending of a regional data set with a global one for the Sea Surface Temperature without any smoothing, the 4-argument input line is:

   ```
   #Sea Surface Temperature FILLING operation without smoothing
        $CFILE  $NUMLU  $CMETHOD  $NFCOAST
   ```

   where `$CFILE`, `$NUMLU` are as for the SIE filling. In this case however, `$CMETHOD(1:3) = RAW`, `$CMETHOD(4:6) = SST`, and the last two characters define the time characteristics of the global data file, as for the SIE filling. `$NFCOAST` is the flag for the calculation of the coastal correction (0 no, 1 yes).

3. If FILLING performs the blending of a regional data set with a global one for the Sea Surface Temperature with smoothing, the 6-argument input line is:

   ```
   #Sea Surface Temperature FILLING operation with smoothing
        $CFILE  $NUMLU  $CMETHOD  $NFCOAST  $CNAME  $NUNIT
   ```

   where `$CFILE`, `$NUMLU` and `$NFCOAST` are as for the SST filling without smoothing. In this case, `$CMETHOD(1:3) = SMO`, `$CMETHOD(4:6) = SST`, and the last two characters define the time characteristics of the global data file, as for the SIE filling. `$CNAME` is the symbolic name for the correction term that is calculated by OASIS3 and `$NUNIT` the logical unit on which it is going to be written.

5.5 The "cooking" stage

The following analyses are available in the "cooking" part of OASIS3, controlled by `cookart.f`.

- **CONSERV**:

  `CONSERV` (routine `/prism/src/mod/oasis3/src/conserv.f`) performs global flux conservation. The flux is integrated on both source and target grids, without considering values of masked points, and the residual (target - source) is calculated. Then all flux values on the target grid are uniformly modified, according to their corresponding surface. This analysis requires one input line with one argument:

  ```
  # CONSERV operation
       $CMETHOD
  ```

  where `$CMETHOD` is the conservation method required. In this version, only global flux conservation can be performed. Therefore `$CMETHOD` must be `GLOBAL`.

- **SUBGRID**:

  `SUBGRID` can be used to interpolate a field from a coarse grid to a finer target grid (the target grid must be finer over the whole domain). Two types of subgrid interpolation can be performed, depending on the type of the field.

  For solar type of flux field (`$SUBTYPE = SOLAR`), the operation performed is:

  $$\Phi_i = \frac{1-\alpha_i}{1-\alpha} F$$

  where Φ_i (F) is the flux on the fine (coarse) grid, α_i (α) an auxiliary field on the fine (coarse) grid (e.g. the albedo). The whole operation is interpolated from the coarse grid with a grid-mapping type of interpolation; the dataset of weights and addresses has to be given by the user.

  For non-solar type of field (`$SUBTYPE = NONSOLAR`), a first-order Taylor expansion of the field on the fine grid relatively to a state variable is performed (for instance, an expansion of the total heat flux relatively to the SST):

  $$\Phi_i = F + \frac{\partial F}{\partial T}(T_i - T)$$

  where Φ_i (F) is the heat flux on the fine (coarse) grid, T_i (T) an auxiliary field on the fine (coarse) grid (e.g. the SST) and $\frac{\partial F}{\partial T}$ the derivative of the flux versus the auxiliary field on the coarse grid. This operation is interpolated from the coarse grid with a grid-mapping type of interpolation; the dataset of weights and addresses has to be given by the user.

  This analysis requires one input line with 7 or 8 arguments depending on the type of subgrid interpolation.

  1. If the the `SUBGRID` operation is performed on a solar flux, the 7-argument input line is:

```
# SUBGRID operation with $SUBTYPE=SOLAR
  $CFILE  $NUMLU  $NID  $NVOISIN  $SUBTYPE  $CCOARSE $CFINE
```

where $CFILE and $NUMLU are the subgrid-mapping file name and associated logical unit (see section 6.4 for the structure of this file); $NID the identificator for this subgrid-mapping dataset within the file build by OASIS based on all the different SUBGRID analyses in the present coupling; $NVOISIN is the maximum number of target grid points use in the subgrid-mapping; $SUBTYPE = SOLAR is the type of subgrid interpolation; $CCOARSE is the auxiliary field name on the coarse grid (corresponding to α) and $CFINE is the auxiliary field name on fine grid (corresponding to α_i). These two fields needs to be exchanged between their original model and OASIS3 main process, at least as AUXILARY fields. This analysis is performed from the coarse grid with a grid-mapping type of interpolation based on the $CFILE file.

2. If the the SUBGRID operation is performed on a nonsolar flux, the 8-argument input line is:

```
# SUBGRID operation with $SUBTYPE=NONSOLAR
  $CFILE $NUMLU $NID $NVOISIN $SUBTYPE $CCOARSE $CFINE $CDQDT
```

where $CFILE, $NUMLU, $NID, $NVOISIN are as for a solar subgrid interpolation; $SUBTYPE = NONSOLAR; $CCOARSE is the auxiliary field name on the coarse grid (corresponding to T) and $CFINE is the auxiliary field name on fine grid (corresponding to T_i); the additional argument $CDQDT is the coupling ratio on the coarse grid (corresponding to $\frac{\partial F}{\partial T}$) These three fields need to be exchanged between their original model and OASIS3 main process as AUXILARY fields. This operation is performed from the coarse grid with a grid-mapping type of interpolation based on the $CFILE file.

- **BLASNEW**:

  BLASNEW (routine /prism/src/mod/oasis3/src/blasnew.f) performs a linear combination of the current coupling field with any other fields after the interpolation. These can be other coupling fields or constant fields.

  This analysis requires the same input line as BLASOLD.

- **MASKP**:

  A new analysis MASKP can be used to mask the fields after interpolation. MASKP has the same generic input line as MASK.

5.6 The post-processing

The following analyses are available in the post-processing part of OASIS3, controlled by /prism/src/mod/oasis3/src/postpro.f.

- **REVERSE**:

  REVERSE (routine /prism/src/mod/oasis3/src/reverse.f) reorders a field.

This analysis requires the same input line as INVERT, with $CORLON and $CORLAT being now the resulting orientation. REVERSE does not work for U and D grids (see annexe A).

- **CHECKOUT**:

  CHECKOUT (routine /prism/src/mod/oasis3/src/chkfld.f) calculates the mean and extremum values of an output field and prints them to the coupler output *cplout*.

  The generic input line is as for CHECKIN.

- **GLORED** (not recommended as coupling fields can be directly interpolated to a target Reduced grid, if needed):

  GLORED performs a linear interpolation of field from a full Gaussian grid to a Reduced grid. When present, GLORED must be the last analysis performed.

  Before doing the interpolation, non-masked values are automatically extrapolated to masked points with EXTRAP/NINENN method (see above); to do so, the masked grid points are first replaced with a predefined value. The required global grid mask must be present in data file masks or masks.nc (see section 6.1).

  The generic input line is as follows:

```
# GLORED operation
     $NNBRLAT  $NVOISIN   $NIO   $NID
```

  where $NNBRLAT is as for REDGLO (see REDGLO description above). The next 3 parameters refer to the EXTRAP/NINENN extrapolation (see EXTRAP/NINENN description above). The value assigned to all land points before interpolation is given by amskred in /prism/src/mod/oasis3/src/blkdata.f; as for the $VALMASK in MASK analysis, it has to be chosen well outside the range of your field values but not too large to avoid any representation problem.

6 OASIS3 auxiliary data files

OASIS3 needs auxiliary data files describing the grids of the models being coupled, containing the fields coupling restart values or input data values, as well as a number of other auxiliary data files used in specific transformations.

6.1 Grid data files

The grid of the models being coupled must be given by the user in grid data files. These files can be all binary or all NetCDF files. NetCDF examples can be found in `/prism/util/mod/toyclim/Cpl/Grids`. The arrays containing the grid information are dimensioned `(nx, ny)`, where `nx` and `ny` are the grid first and second dimension, except for Unstructured ('U') and Reduced ('D) grid, for which the arrays are dimensioned `(nbr_pts,1)`, where `nbr_pts` is the total number of grid points.

1. *grids* or *grids.nc*: contains the component model grid point longitudes and latitudes in single or double precision `REAL` arrays (depending on OASIS3 compilation options). The array names must be composed of a prefix (4 characters), given by the user in the *namcouple* on the second line of each field (see section 4.3), and of a suffix (4 characters); this suffix is '.lon or '.lat for respectively the grid point longitudes or latitudes (see `/prism/src/mod/oasis3/src/mod_label.F90`.)

   If the `SCRIPR/CONSERV` interpolation is used for a grid, the grid data file may contain longitudes and latitudes for model mesh corners as arrays dimensioned `(nx, ny, nc)` or `(nbr_pts,1, 4)` where `4` is the number of corners; in this case, it must necessarily be a NetCDF file (*grids.nc*). For Logically Rectangular 'LR' grids, the grid corners will be automatically approximately calculated if they are not given in *grids.nc*. The names of the arrays must be composed of the grid prefix and the suffix '.clo or '.cla for respectively the grid corner longitudes or latitudes.

   File *grids* or *grids.nc* must be present with at least the grid point longitudes and latitudes for all component model.

2. *masks* or *masks.nc*: contains the masks for all component model grids in `INTEGER` arrays (0 -not masked- or 1 -masked- for each grid point). The array names must be composed of the grid prefix and the suffix '.msk'. This file, *masks* or *masks.nc*, is mandatory.

3. *areas* or *areas.nc*: this file contains mesh surfaces for the component model grids in single or double precision `REAL` arrays (depending on OASIS3 compilation options). The array names must be composed of the grid prefix and the suffix '.srf'. The surfaces may be given in any units but they must be all the same (in `INTERP/GAUSSIAN`, it is assumed that the units are m^2 but they are used for statistics calculations only.) This file *areas* or *areas.nc* is mandatory for `CHECKIN`, `CHECKOUT` or `CONSERV`, and used for statistic calculations in `INTERP/GAUSSIAN`; it is not required otherwise.

4. *maskr* or *maskr.nc*: this file contains Reduced ('D') grid mask in `INTEGER` arrays dimensioned `array(nbr_pts)` where `nbr_pts` is the total number of the Reduced

grid points (0 -not masked- or 1 -masked- for each grid point). This file is required only for grids to which the `REDGLO` or `GLORED` transformation is applied. As mentionned above, these transformations should not be used anymore as interpolations are now available for Reduced grids directly. If used, the mask array name must be 'MSKRDxxx" where "xxx" is half the number of latitude circles of the reduced grid (032 for a T42 for example).

If the binary format is used, *grids*, *masks*, *areas*, and *maskr* must have the following structure. The array name is first written to the file to locate a data set corresponding to a given grid. The data set is then written sequentially after its name. Let us call "brick" the name and its associated data set. The order in which the bricks are written doesn't matter. All the bricks are written in the grid data file in the following way:

```
...
WRITE(LU) array_name
WRITE(LU) auxildata
...
```

where

- `LU` is the associated unit,
- `array_name` is the name of the array (CHARACTER*8),
- `auxildata` is the REAL or INTEGER array dimensioned `(nx, ny)` or `(nbr_pts,1)` containing the grid data.

6.2 Coupling restart files

At the beginning of a coupled run, some coupling fields may have to be initially read in from their coupling restart file (see section 3.7). The name of the coupling restart file is given by the 6th character string on the first configuring line for each field in the *namcouple* (see section 4.3). Coupling fields coming from different models cannot be in the same coupling restart files, but for each model, there can be an arbitrary number of fields written in one coupling restart file. Note that in the NONE techniques, output files with the same format are also created for writing the resulting field after transformation (see `sstac.nc` in section 4.3).

Both binary and NetCDF formats are supported; for NetCDF file the suffix .nc is not mandatory. If the coupling restart file for the first field is in NetCDF format, OASIS3 will assume that all coupling restart files (and output files for `NONE` communication techniques) are NetCDF[12].

In the coupling restart files, the fields must be single or double precision REAL arrays (depending on PSMILe and OASIS3 compilation options) and, as the grid data files, must be dimensioned `(nx, ny)`, where `nx` and `ny` are the grid first and second

[12]Note that even if the grid auxiliary data files are in NetCDF format, the restart coupling files may be in binary format, or vice-versa.

dimension, except for fields given on Unstructured ('U') and Reduced ('D) grid, for which the arrays are dimensioned `(nbr_pts,1)`, where `nbr_pts` is the total number of grid points. The shape and orientation of each restart field (and of the corresponding coupling fields exchanged during the simulation) must be coherent with the shape of its grid data arrays. The only exception is for A, B, G, L, Z, or Y grids for which the field may be oriented from North to South and/or from East to West, in which case, `INVERT` transformation will have to be used - see section 5.3.

In the NetCDF restart files, the field arrays must have the source symbolic name indicated in the *namcouple* (see section 4.3.1).

In binary restart file, each field is written in the following way:

```
...
WRITE(LU) array_name
WRITE(LU) restartdata
...
```

where

- `LU` is the associated unit,
- `array_name` is the source symbolic name of the field (CHARACTER*8),
- `restartdata` is the restart field REAL array dimensioned `(nx, ny)` or `(nbr_pts,1)`[13]

6.3 Input data files

Fields with status `INPUT` in the *namcouple* will, at runtime, simply be read in from a NetCDF input file by the target model PSMILe below the `prism_get_proto` call, at appropriate times corresponding to the input period indicated by the user in the *namcouple*.

The name of the file must be the one given on the field first configuring line in the *namcouple* (see section 4.3.3). There must be one input file per `INPUT` field, containing a time sequence of the field in a single or double precision REAL array (depending on PSMILe compilation options), named with the field symbolic name in the *namcouple* and dimensioned `(nx,ny,time)` or `(nbr_pts,1,time)`. The time variable as to be an array `time(time)` expressed in "sec_since_beginning_of_run". For a practical example, see the file `prism/util/mod/toyclim/Cpl/Restart/SOALBEDO.nc`.

6.4 Transformation auxiliary data files

Many transformation need auxiliary data files, such as the grid-mapping files used for an interpolation. Some of them are created automatically by OASIS3, others have to be generated by the user before starting the coupled run.

[13] If REDGLO is the first transformation applied on a Reduced grid field, the Reduced field must be given is an array `restartdata(nx*ny)` where `nx` and `ny` are the global Gaussian grid dimensions and the Reduced field is completed by trailing zeros.

6.4.1 Auxiliary data files for EXTRAP/NINENN, EXTRAP/WEIGHT, INTERP/SURFMESH, INTERP/GAUSSIAN, MOZAIC, and SUBGRID

The auxiliary data files containing the weights and addresses used in these transformations have a similar structure; their names are given in Table 1.

File name	Description
nweights	weights, addresses and iteration number for EXTRAP/NINENN interpolation
any name	weights and addresses for EXTRAP/WEIGHT extrapolation
mweights	weights and addresses for INTERP/SURFMESH interpolation
gweights	weights and addresses for INTERP/GAUSSIAN interpolation
any name	weights and addresses for MOZAIC interpolation
any name	weights and addresses for SUBGRID interpolation

Table 1: Analysis auxiliary data files

The files *nweights*, *mweights* and *gweights* can be created by OASIS3 if their corresponding `$NIO` = 1 (see `EXTRAP/NINENN, INTERP/SURFMESH, INTERP/GAUSSIAN` in sections 5.3 and 5.4).

The name of the (sub)grid-mapping files for `MOZAIC, EXTRAP/WEIGHT` and `SUBGRID` analyses can be chosen by the user and have to be indicated in the *namcouple* (see respectively sections 5.3 and 5.4 and 5.5). These files have to be generated by the user before starting the coupled run.

The structure of these files is as follows:

```
...
CHARACTER*8 cladress,clweight
INTEGER iadress(jpnb,jpo)
REAL weight(jpnb,jpo)
OPEN(unit=90, file='at31topa', form='unformatted')
WRITE(clweight,'(''WEIGHTS'',I1)') knumb
WRITE(cladress,'(''ADRESSE'',I1)') knumb
WRITE (90) cladress
WRITE (90) iadress
WRITE (90) clweight
WRITE (90) weight
```

where

- `jpnb` is the maximum number of neighbors used in the transformation (`$NVOISIN` in the *namcouple*)
- `jpo` is the total dimension of the target grid
- `at31topa` is the name of the grid-mapping data file (`$CFILE` in *namcouple*)
- `knumb` is the identificator of the data set (`$NID` in *namcouple*)

- `cladress` is the locator of the address dataset
- `clweight` is the locator of the weight dataset
- `iadress (i,j)` is the address on the source grid of the i^e neighbor used for the mapping of the j^e target grid point. The address is the index of a grid point within the total number of grid points.
- `weight(i,j)` is the weight affected to the i^e neighbor used for the transformation of the j^e target grid point

For file *nweights*, there is an additional brick composed of a `CHARACTER*8` variable (formed by the characters `INCREME` and by the data set identificator) and of an `INTEGER array(N)` which is the iteration number within `EXTRAP/NINENN` at which the extrapolation of the n^e grid point is effectively performed.

6.4.2 Auxiliary data files for `FILLING`

For the FILLING analysis, the global data set used can be either interannual monthly, climatological monthly or yearly (see 5.4). The name of the global data file can be chosen by the user and has to be indicated in the *namcouple* have to be given to OASIS through the input file *namcouple*. In case of monthly data, the file must be written in the following way:

```
      ...
      REAL field_january_year_01(jpi, jpj)
      ...
      WRITE(NLU_fil) field_january_year_01
      WRITE(NLU_fil) field_february_year_01
      WRITE(NLU_fil) field_march_year_01
      etc...
      WRITE(NLU_fil) field_december_year_01
C
C if climatology, one stops here
C
      WRITE(NLU_fil) field_january_year_02
      etc...
```

where

- `field_...` is the global dataset
- `jpi` and `jpj` are the dimensions of the grid on which FILLING is performed
- `NLU_fil` is the logical unit associated to the global data file and is defined in the input file *namcouple*

Note that the first month needs not to be a january. This is the only file within OASIS in which the fields are not read using a locator.

6.4.3 Auxiliary data files for SCRIPR

The NetCDF files containing the weights and addresses for the SCRIPR remappings (see section 5.4 are automatically generated at runtime by OASIS3. Their structure is described in detail in section 2.2.3 of the SCRIP documentation available in prism/src/mod/oasis3/doc/SCRIPusers.pdf.

7 Compiling and running with OASIS3

7.1 Compiling OASIS3 and the toy coupled model

OASIS3 uses the PRISM standard directory structure and compiling environment. To compile OASIS3, execute the script `COMP_oasis3` in `prism/src/mod/oasis3`. This script automatically executes `COMP_libs_oasis3` in `prism/util`, which compiles appropriate libraries in `prism/src/lib`.

A practical example of a toy coupled model using OASIS3 and PSMILe with the CLIM/MPI1 or MPI2 communication technique is given in `prism/util/mod/toyclim`. To compile its different components, execute the scripts `COMP_toyatm`, `COMP_toyoce`, `COMP_toyche` in `prism/src/mod/toyatm`, `/toyoce`, or `/toyche` respectively.

To compile OASIS3 and the toyclim components on another platform than those already included, adapt the parts "user specifications" and "configurable specification" of those scripts. OASIS3 and the toy coupled model have been compiled and successfully run on Fujitsu VPP5000, NEC SX5, SGI Octane and 03000, COMPAQ Alpha cluster and Linux PC cluster.

In the "user specifications" part of those scripts, the user has to indicate in particular his choice of message passing (`MPI1, MPI2` or `NONE`). For `MPI1` and `MPI2`, one also as to indicate whether or not the buffered `MPI_BSend` should be used to send the coupling fields instead of a standard blocking send `MPI_Send` (`use_key_BSend = BSend` or `NoBSend`).[14] This will activate, or not, the pre-compiling key `key_BSend`.

Other pre-compiling keys are automatically activated by those compiling scripts:

- To indicate which communication technique will be used (`CLIM/MPI2` by default):
  - use_comm_MPI2 : CLIM/MPI2 for OASIS3 and PSMILe, or
  - use_comm_MPI1 : CLIM/MPI1 for OASIS3 and PSMILe, or
  - use_comm_NONE: no communication technique for OASIS3 (interpolator-only mode).
- Other pre-compiling keys for the PSMILe I/O library (routines in `prism/lib/mpp_io`):
  - use_netCDF (mandatory)
  - use_libMPI (mandatory)
  - use_LAM_MPI, if LAM MPI version is used.

The following keys are also available but not activated in those compiling scripts:

[14] Use the standard blocking send `MPI_Send` if the coupling fields are necessarily sent and received in the same order, or on architectures on which `MPI_Send` is implemented with a mailbox (e.g. VPPs); in this second case, make sure that the size of the mailbox is sufficient. Use the less efficient buffered send `MPI_BSend` on architectures on which `MPI_Send` is not implemented with a mailbox (e.g. NEC SX5) if the coupling fields are not sent and received in the same order. This pre-compiling option affects `prism/src/lib/clim/src/CLIM_Export.F` and `prism/src/lib/psmile/src/mod_prism_put_proto.F90`.

- `use_realtype_single` to exchange single precision REAL variables as coupling fields (see section 3.3)
- `key_noIO` for compiling PSMILe without compiling the I/O library(i.e. routines in `prism/src/lib/mpp_io`, and `mod_psmile_io.F90` and `mod_psmile_io_interfaces.F90` in `prism/src/lib/psmile`.

7.2 Running OASIS3

To run the toy coupled model which set-up is in `/prism/util/mod/toyclim`, first compile OASIS3 and the toyclim component models (`toyatm, toyoce`, `toyche`) with `MPI1` or `MPI2` communication technique (see above). To start the run, adapt the "User Interface" part of `sc_run_toyclim` and execute it. For more details, please read carefully the `README_toyclim` therein.

To run OASIS3 in an interpolator-only mode, first compile OASIS3 with `NONE` as message_passing and then execute `/prism/util/mod/testinterp/sc_run_testinterp`. For more details, please read carefully the `README_testinterp` therein.

A The grid types for the transformations

As described in section 5, the different transformations in OASIS3 support different types of grids. The characteristics of these grids are detailed here.

1. Grids supported for the `INTERP` interpolations (see section 5.4)

   - `'A' grid`: this is a regular Lat-Lon grid covering either the whole globe or an hemisphere, going from South to North and from West to East. There is no grid point at the pole and at the equator, and the first latitude has an offset of 0.5 grid interval. The first longitude is 0^o (the Greenwhich meridian) and is not repeated at the end of the grid (`$CPER` = P and `$NPER`= 0). The latitudinal grid length is 180/NJ for a global grid, 90/NJ otherwise. The longitudinal grid length is 360/NI.
   - `'B' grid`: this is a regular Lat-Lon grid covering either an hemisphere or the whole globe, going from South to North and from West to East. There is a grid point at the pole and at the equator (if the grid is hemispheric or global with NJ odd). The first longitude is 0^o (the Greenwhich meridian), and is repeated at the end of the grid (`$CPER` = P and `$NPER`= 1). The latitudinal grid length is 180/(NJ-1) for a global grid, 90/(NJ-1) otherwise. The longitudinal grid length is 360/(NI-1).
   - `'G' grid`: this is a irregular Lat-Lon Gaussian grid covering either an hemisphere or the whole globe, going from South to North and from West to East. This grid is used in spectral models. It is very much alike the A grid, except that the latitudes are not equidistant. There is no grid point at the pole and at the equator. The first longitude is 0^o (the Greenwhich meridian) and is not repeated at the end of the grid (`$CPER` = P and `$NPER`= 0). The longitudinal grid length is 360/NI.
   - `'L' grid`: this type covers regular Lat-Lon grids in general, going from South to North and from West to East.. The grid can be described by the latitude and the longitude of the southwest corner of the grid, and by the latitudinal and longitudinal grid mesh sizes in degrees.
   - `'Z' grid`: this is a Lat-Lon grid with non-constant latitudinal and longitudinal grid mesh sizes, going from South to North and from West to East. The deformation of the mesh can be described with the help of 1-dimensional positional records in each direction. This grid is periodical (`$CPER` = P) with `$NPER` overlapping grid points.
   - `'Y' grid`: this grid is like 'Z' grid except that it is regional (`$CPER` = R and `$NPER` = 0).

2. Grids supported for the `SCRIPR` interpolations

   - `'LR' grid`: The longitudes and the latitudes of 2D Logically-Rectangular (LR) grid points can be described by two arrays `longitude(i,j)` and `latitude(i,j)`, where i and j are respectively the first and second index dimensions. Streched or/and rotated grids are LR grids. Note that A, B, G, L, Y, or Z grids are all particular cases of LR grids.

- 'U' grid: Unstructured (U) grids do have any particular structure. The longitudes and the latitudes of 2D Unstructured grid points must be described by two arrays `longitude(nbr_pts)` and `latitude(nbr_pts)`, where nbr_pts is the total grid size.
- 'D' grid The Reduced (D) grid is composed of a certain number of latitude circles, each one being divided into a varying number of longitudinal segments. In OASIS3, the grid data (longitudes, latitudes, etc.) must be described by arrays dimensioned `(nbr_pts,1)`, where `nbr_pts` is the total number of grid points. There is no overlap of the grid, and no grid point at the equator nor at the poles. There are grid points on the Greenwich meridian.

B The SCRIP 1.4 copyright statement

The SCRIP 1.4 copyright statement reads as follows:

"Copyright 1997, 1998 the Regents of the University of California. This software and ancillary information (herein called SOFTWARE) called SCRIP is made available under the terms described here. The SOFTWARE has been approved for release with associated LA-CC Number 98-45. Unless otherwise indicated, this SOFTWARE has been authored by an employee or employees of the University of California, operator of Los Alamos National Laboratory under Contract No. W-7405-ENG-36 with the United States Department of Energy. The United States Government has rights to use, reproduce, and distribute this SOFTWARE. The public may copy, distribute, prepare derivative works and publicly display this SOFTWARE without charge, provided that this Notice and any statement of authorship are reproduced on all copies. Neither the Government nor the University makes any warranty, express or implied, or assumes any liability or responsibility for the use of this SOFTWARE. If SOFTWARE is modified to produce derivative works, such modified SOFTWARE should be clearly marked, so as not to confuse it with the version available from Los Alamos National Laboratory."

References

[1] http://www.acl.lanl.gov/climate/software/SCRIP/scrip_frames.htm

[2] http://www.gfdl.gov/~ vb/mpp_io.html