UNIFIED MODEL DOCUMENTATION PAPER

C2: Atmosphere - Ocean coupling: technical overview

BY

R. Rawlins

VERSION 2

19/11/98

MODEL VERSION 4.5

Numerical Weather Prediction
Meteorological Office
London Road
BRACKNELL
Berkshire
RG12 2SZ
United Kingdom

Modification Record

Document version	Author	Description....................
1.0	R.Rawlins	First version. 30/06/97
1.1	Chris Jones	Add Carbon-cycle coupling fields. 19/11/98

Atmosphere - Ocean coupling: technical overview

---

1. Introduction

This paper is an overview of the method used to couple atmosphere and ocean models within the Unified Model system. Descriptions are confined to technical aspects of software components: no scientific issues or rationale for the adopted coupling technique is given here.

A requirement of the Unified Model system is that atmosphere and ocean models can run independently, or sequentially in coupled mode. For independent integrations, state variables such as sea-surface temperature in the atmosphere model, or radiative heat flux into the sea surface for the ocean model, are supplied as *ancillary fields* derived from external files. In the coupled model, these constitute *coupling fields* which are supplied from the ocean or atmosphere integrations alternately, and can be derived from either prognostic or diagnostic variables.

The UM shares the same control code for atmosphere and ocean components: hence the top-level code is capable of integrating either independently, or in coupled mode. This is achieved partly by the use of namelist variables as input, and partly by compile time *DEF switches, both generated from the UMUI job set-up. A number of changes were required to enable the coupling method to work for both MPP and non-MPP architectures. A description of changes for MPP is gathered into a single section in this paper.

2. Submodels

Atmosphere and ocean code comprise separate *submodels*, ie each has distinct data structures inside the model and separate data files (model dumps) used to define the model state for restarting integrations. A lower level of modularity is described by *internal models*, which can be integrated for one or more timesteps within a submodel. For atmosphere-ocean coupling this distinction is not required, since atmosphere and ocean submodels only contain one internal model each, ie atmosphere and ocean internal models. [The extra complexity is needed for handling the slab internal model, where slab and atmosphere internal models coexist within the atmosphere submodel, and coupling is handled within the submodel.]

The sequence of control through multiple timesteps of the coupled model is illustrated in Figure 1. A coupling period needs to be an integral number of atmosphere or ocean timesteps and the code should be general enough to allow a range of values. However a period of 1 day has been chosen for all practical models and other periods will not have been subjected to extensive testing. In particular, climate diagnostic meaning routines within the model have traditionally assumed coupling every day.

3. Compilation

Submodel specific code is compiled under *DEF compile switches ATMOS and OCEAN. Hence coupling code is usually protected by

```
*IF DEF,ATMOS,AND,DEF,OCEAN
*ENDIF
```

Other compile (and run-time) switches within coupling routines are given in Table 1.

Atmosphere and ocean routines are compiled together, in a single implementation of nupdate followed by the f90 compilation step, thus ensuring that atmosphere and ocean COMDECKs are consistent, and that any control routine shared by both atmosphere and ocean submodels are only compiled once.

4. Data structures

State variables are held in atmosphere and ocean model dumps, defining *primary* initial fields for the start of the integration, consisting of prognostic and ancillary fields. Forecast dumps after integration are written out periodically with the same format, except that accumulated, mean and other diagnostic fields are written at the end of each file if required to preserve restartability from interim dumps. The complete set of fields are accompanied by descriptions of each field, held in lookup tables.

Data from the initial dump is read sequentially as a series of horizontal fields into a large array - the D1 array - which constitutes the data structure within the model. Atmosphere and ocean D1 arrays are separate: each submodel data structure is read into D1. For non-MPP implementations the D1 array is then immediately written out to a dedicated local file in order to reduce memory requirements. Hence moving between atmosphere and ocean data values requires a swap from the current contents of D1 to a copy held elsewhere. Arrays containing the lookup tables are held in memory simultaneously - A_LOOKUP and O_LOOKUP - for atmosphere and ocean submodels.

The D1 array in the model contains primary and diagnostic fields from the corresponding dump, but also secondary space for intermediate fields, which can include coupling variables. The complete set of pointers to each field within D1 is calculated by the STASH addressing routines - which also control the contents of the dump. Access to coupling fields is obtained by finding pointers in D1, for atmosphere and ocean submodels separately. Coupling fields derived from prognostic variables are simply accessed from the primary pointers calculated in routines SET_ATM_POINTERS and SET_OCN_POINTERS. For diagnostic variables, access to coupling fields is achieved through searching a STASH array (STLIST) for locations of the correct fields. These fields are introduced into the job by the addition of a STASH macro generated by the UMUI for coupled models. A STASH macro tag (=10 for A-O; =11 for O-A) ensures unique identification from potentially multiple diagnostics. [See UM Documentation paper C4: STASH]

Note that atmosphere and ocean dumps, although covered by the same generic formats, have different implementations. Atmosphere fields are on the full horizontal domain, except for some land only fields, which are compressed to land points. Ocean fields can be held compressed on to ocean points, but are now mostly converted to the full horizontal domain. In addition, horizontal domains are represented differently: the first row of an ocean array is located at the S pole, whereas atmosphere arrays start at the N pole.

5. Coupling fields

5.1 Ocean to atmosphere

A list of ocean to atmosphere coupling fields is given in Table 2. The number of fields depends on whether sea-ice is predicted by the ocean submodel and whether the full interactive carbon cycle model is selected. Each field is a prognostic variable in the ocean model and is uniquely determined by the stash code (internal model,section,item). Equivalent variables in the atmosphere submodel are primary variables, and there is a one to one correspondence from ocean to atmosphere fields.

5.2 Atmosphere to ocean

A list of ocean to atmosphere coupling fields is given in Table 3. The number of fields needed also depends on whether sea-ice is predicted in the ocean submodel and also if the full interactive carbon cycle model is selected. In contrast to above, some coupling fields are derived from atmosphere submodel diagnostics, and are used in combination to obtain moisture and heat fluxes. Hence a direct correspondence of atmosphere and ocean fields is not obtained for all variables.

6. Coupling transformations

For atmosphere-ocean models, coupling only occurs at the ocean surface, therefore only surface fields are involved. Transformations are required to:

Provide horizontal interpolation from the atmosphere to ocean grid and vice versa. For the standard HADCM2 configuration the 2 grids are congruent, and no interpolation is required. In the standard HADCM3 configuration the ocean grid is finer than the atmosphere grid and horizontal interpolation is necessary. Bi-linear or area weighted interpolation is enabled under *DEF compile switches, with a choice over individual coupling fields, as in Table 1.

Mask land points from the coupling with coastal adjustment of fields if needed.

Conversion of multiple diagnostics into a single field and changes from rates to fluxes.

Invert order of rows between atmosphere and ocean grids.

7. Modifications for MPP

i) Data structure

In the MPP version of the code each processor has a D1 array for the limited area of the domain dedicated to that processor. For the atmosphere submodel the domain decomposition is 2-D, so that a rectangular portion of the global grid is represented. For the ocean submodel a 1-D domain decomposition is applied and each processor represents contiguous rows of the global grid. Hence, each processor has more memory to cope with extra arrays and it is no longer necessary to swap ocean and atmosphere D1 arrays between memory and disk storage.

Extra copies of D1 for atmosphere and ocean submodels are held in memory in arrays D1_A and D1_O. Swapping between these arrays and the main array D1 is retained, but is now simply a local copy operation.

Note that this technique wastes memory by duplication of storage. A more efficient method would employ a submodel pointer. This could point to D1_A or D1_O as required, eliminating D1, or enlarge D1 to accomodate atmosphere and ocean fields, eliminating D1_A and D1_O copies. The ability to write D1 to disk is retained as an option, in the event of serious memory shortcomings. This option is not available through the UMUI, but could be implemented by setting AO_D1_MEMORY=.false. in COMDECK MPPTRANS.

ii) Coupling transformations

Since the cost of coupling transformations is small compared to the model integration, these are performed on a single processor. This requires gathering coupling fields from the distributed D1 arrays, performing interpolations on global fields, and then scattering the new primary fields back into D1 arrays over the distributed set of processors. This simplifies modifications to the coupling transformation routines, which are essentially unchanged from the non-MPP version, at the cost of losing scalability for this portion of code. New DECKs SWAPA2O2 and SWAPO2A2 were created for MPP interfaces to the transformation routines.

iii) Other changes

Extra calls were required to change parameters describing the domain decomposition - atmosphere or ocean - to ensure that local and global MPP sizes were appropriate for the submodel. In particular, each time a swap between atmosphere and ocean data partitions is performed there is an implicit change of domain decomposition and the routine CHANGE_DECOMPOSITION was added to subroutine TRANSIN.

Expansion of the ocean entry point index array to a global index array was needed, requiring access to the land sea mask via the IOVARS comdeck.

Changes associated with the move to holding full ocean arrays, ie where horizontal fields are no longer compressed onto ocean points alone.

8. Code structure

Routines relevant to atmosphere-ocean coupling are shown in Figure 2. The structure is identical for MPP and non-MPP versions but descriptive comments refer to MPP operation.

Control of submodel and internal model sequences is initiated from the general routine UM_SUBMODEL_INIT called from the top level UM_SHELL, with input from a UMUI generated namelist. Atmosphere and ocean data addresses are calculated using generic routines also called from the top level. Initialisation routines are called from INITIAL, where atmosphere and ocean start dump files are accessed sequentially, with data first read into the D1 array. TRANSOUT/TRANSIN routines are used to copy data out to and back from the appropriate D1_A or D1_O array. Locations of coupling fields are determined in INITA2O: primary fields have pointers assigned directly; pointers for diagnostic fields are found using the generic FINDPTR routine to locate STASH tags.

At the end of initialisation the coupling routine SWAPO2A (or SWAPA2O) is called to populate coupled fields for the first submodel - the atmosphere submodel - to be integrated. Iteration of timesteps is controlled by SETGRCTL, switching to a new submodel and calling the relevant SWAP routine when a coupling period is complete if the cycle is to be continued. See Figure 1 for the normal calling sequence. Note that timestep control should be sufficiently general that ocean steps could precede atmosphere steps. However, this is not an enabled option since it would require special treatment to make atmosphere diagnostic fields available for coupling before integration of the atmosphere submodel.

9. Other coupling considerations

A different method is provided for coupling between atmosphere and slab internal models. Here both atmosphere and slab primary fields are held within the same submodel data partition, ie the same D1 array within the model. Further, the location of prognostic and coupling fields are shared within D1. This avoids some of the complexity of the atmosphere-ocean system, since atmosphere and slab horizontal fields are congruent by definition and explicit swapping of fields unnecessary. However, atmosphere-slab coupling has not been enabled for MPP at this version and no further details are given.

An alternative approach to coupling is under development in which separate executable files are built for each submodel. The separate submodels then run independently but with synchronisation provided at the unix level, with piped communications of coupling fields.

Acronyms

MPP: Massively parallel processors: ie distributed memory architecture.

UM: Unified Model.

UMUI: Unified Model User Interface.

Figures

- **Figure 1. Timestep sequence in coupled model**

```
Initialise Atmosphere and Ocean
                 |
                 |--------------------------------------
                 |                                          |
                 |                                          |
Couple Ocean - Atmosphere                                   |
                 |                                          |
                 |                                          |
                 |----------------------           |
                 |                          |                |
Timestep control for Atmosphere            |Iterate |
then                                       |over a |
Atmosphere timestep                        |coupling        |
                 |                          |period          |Iterate
                 |----------------------           |over
                 |                                          |coupling
                 |                                          |periods
Couple Atmosphere - Ocean                                   |
                 |                                          |
                 |                                          |
                 |----------------------           |
                 |                          |                |
Timestep control for Ocean                 |Iterate |
then                                       |over a |
Ocean timestep                             |coupling        |
                 |                          |period          |
                 |----------------------           |
                 |                                          |
                 |--------------------------------------
```

- **Figure 2. Top level Atmos - Ocean coupling code**

```
UM_SHELL

initialise addresses, pointers

- U_MODEL
   |
   |- INITIAL
   |   |
   |   |- INITDUMP   Read dump        ]
   |   |                              ]  loop over atmos, ocean
   |   |- TRANSOUT   D1 - D1_A/D1_O  ]
   |   |
   |   |- TRANSIN    D1_A/D1_O - D1
   |   |
   |   |- SETGRCTL   control atmos, ocean stepping order
   |   |
   |   |- TRANSIN                     ]
   |   |                              ]
   |   |- initialisation routines     ]  atmos then ocean
   |   |                              ]
   |   |- TRANSOUT                    ]
   |   |
   |   |- INIT_A2O   find tagged pointers to addresses in D1
   |   |             (FINDPTR)
   |   |- TRANSIN    D1_A - D1           ]
   |   |                                  ]  (if atmos last)
   |   |- SWAP_A2O   swap coupling fields ]
   |   |             atmos - ocean grid  ]
   |   |
   |   |- TRANSIN    D1_O - D1           ]
   |   |                                  ]  (if ocean last)
   |   |- SWAP_O2A   swap coupling fields ]
   |   |             ocean - atmos grid  ]
   |
   |- ATM_STEP      atmos timestep                                ]
   |                                                              ]
   |- OCN_STEP      ocean timestep                                ]
   |                                                              ]
   |- SETGRCTL      switch between atmos/ocean timesteps          ] iterate
   |                                                              ]
   |- SWAP_A2O       (when atmos steps completed 1 cpl period)  ]
   |                                                              ]
   |- SWAP_O2A       (when ocean steps completed 1 cpl period)  ]


SWAP_A2O
   |
   |--           copy atmos coupling fields to local arrays
   |
   |- TRANSOUT   D1 (atmos)-  D1_A
   |
   |- TRANSIN    D1_O      -  D1 (ocean)
   |
   |- TRANSA2O   transformation of atmos to ocean fields

SWAP_O2A
   |
   |--           copy ocean coupling fields to local arrays
   |
   |- UNPACK     extract surface fields (only if compressed)
   |             multi-layer ocean fields in D1
   |
   |- TRANSOUT   D1 (ocean)- D1_O
   |
   |- TRANSIN    D1_A      - D1 (atmos)
   |
   |- TRANSO2A   transformation of ocean to atmos fields


TRANSA2O
   |
   |- H_INT_CO   generate bi-linear interpolation coeffs. A - O grids
   |
   |- H_INT_BL   interpolate           ]
   |                                   ]  winds (u then v)
   |- POST_H_INT write to target grid  ]
   |
   |IF(CYCLIC) - CYCLICBC   overwrite cyclic E-W boundaries
   |
   |- H_INT_CO   generate bi-linear interpolation coefficients
```

```
|- COAST_AJ   coastal adjustment of interpolation
|
|IF area-average - PRE_AREAVER   prepare coefficients       ] windmix
|                                                           ] blue sol
|                - DO_AREAEVER   interpolate area average   ] heat flux
|                                                           ] Ppn - Evp
|                - COPYA2O       fill in array on mask      ] river out
|ELSE                                                       ] snowfall
|                - H_INT_BL   interpolate                   ] sublimn
|                                                           ] ice diff
|                - POST_H_INT set coast,write to target grid] ice melt
|                                                           ]
|IF(CYCLIC) - CYCLICBC   overwrite cyclic E-W boundaries    ]
```

```
TRANSO2A

similar to TRANSA2O but with extra
IF(INVERT) - ROWSWAP inverts order of rows to facilitate coupling of
             atmos model (first row at south edge) with an
             ocean model (first row at north edge).
```

Tables

Table 1. Code switches in coupling routines

Compile *DEF:

- TRANGRID: Ocean/atmos grids not congruent - transformations required.
- RIVERS: River outflow - get the runoff from land points into ocean entry points.
- AVER_TAO: Area average coupling fields between grids (TRANSA2O).
- AVER_T: Area average T field (TRANSO2A)
- AVER_U: Area average u,v field (TRANSO2A)
- SEAICE: Include sea-ice model.

Run-time logical

- CYCLIC: periodic ocean.
- GLOBAL: global ocean model.
- INVERT: ocean model S-N.

Table 2. Ocean to atmosphere coupling fields

Fields generated by the ocean submodel (* if sea-ice included in ocean):

Sea surface temperature (2,0,101)

Zonal, meridional sea currents (2,0,121),(2,0,122)

Snowdepth (2,0,141) *

Sea-ice concentration and depth (2,0,146),(2,0,147) *

Air-Sea CO2 flux (2,30,249) (if carbon-cycle selected)

and fields received as coupling fields into the atmosphere submodel:

Surface temperature (1,0,24)

Zonal, meridional sea currents (1,0,28),(1,0,29)

Snowdepth (1,0,23) *

Sea-ice fraction and depth (1,0,31),(1,0,32) *

Air-Sea CO2 flux (1,0,250) (if carbon-cycle selected)

Table 3. Atmosphere to ocean coupling fields

Diagnostic fields generated by the atmosphere submodel (* if sea-ice included in ocean):

Surface u,v windstress components(1,3,219),(1,3,220)

Wind mixing power (1,3,224)

Net downward integrated solar radiation (1,1,203)

Net downward blue-band radiation (1,1,204)

Net downward longwave radiation (1,2,203)

Surface evaporation weighted by leads (1,3,232)

Sensible heat flux over open sea (1,3,228)

Sublimation (1,3,231) *

Melting of bottom and top of sea-ice (1,3,201),(1,3,235) *

Large scale rainfall and snowfall rates (1,4,203),(1,4,204)

Convective rainfall and snowfall rates (1,5,205),(1,5,206)

Fast and slow river runoff of net fresh water into coastal ocean outflow points (1,8,204),(1,8,205)

Surface atmospheric CO2 concentration (1,0,252) (if carbon-cycle selected)

Primary field stored by the atmosphere submodel:

[Sea-ice concentration (1,0,31) for calculating snowfall on sea-ice] *

ocean entry point index (1,0,93) used to calculate river outflow.

and fields received as coupling fields into the ocean submodel:

Surface u,v windstress components (2,0,150),(2,0,151)

Wind mixing energy flux into ocean near-surface layer (2,0,152)

Penetrative solar flux into ocean (2,0,161)

Non-penetrative solar flux plus heat flux into ocean (2,0,162)

Precipitation minus evaporation fresh water flux into ocean (2,0,165)

River runoff of net fresh water into coastal ocean outflow points (2,0,166)

Snowfall flux onto ocean and sea-ice (2,0,171) *

Sublimation from sea-ice (2,0,172) *

Net heat flux from atmosphere into sea-ice and sea-ice melting (2,0,190),(2,0,191) *

Surface atmospheric CO2 concentration (2,0,200) (if carbon-cycle selected)

Rick Rawlins / frrr@fr0800 / frawlins@meto.gov.uk