


PRISM Support Initiative



Coupler and I/O work group The current status

*A. Caubel, D. Declat, J. Ghattas, R. Redler, H. Ritzdorf,
T. Schoenemeyer, S. Valcke and R. Vogelsang*

PRISM Community Meeting 2005
CERFACS, Toulouse, November 16-17, 2005

CRAY

CECMWF



NEC



SMHI





Outline

- Historical background
- Some key notes
- OASIS Community today
- OASIS3:
 - Coupled model configuration
 - Communication
 - Interpolations/transformations
 - Conclusions
- OASIS4:
 - PSMILe API
 - XML description and configuration
 - Driver, communication
 - Transformations, grids supported
 - Conclusions



Historical background

OASIS: developed since 1991 to couple existing GCMs

1991

2001

|--→

|--- PRISM →

OASIS 1 → OASIS 2

→ OASIS3 (demo runs) →

→ OASIS4 →

OASIS1, OASIS2, OASIS3:

- low resolution, low nbr of 2D fields, low coupling frequency:
 - flexibility very important, efficiency not so much!

OASIS4:

- higher resolution parallel models, 3D fields
 - need to optimise and parallelise the coupler



Some key notes

- Main developers:
 - OASIS3: CERFACS, Fujitsu/Fecit, SGI
 - OASIS4: NEC CCRLE, CERFACS, SGI, NEC HPCE
- Programming language: Fortran 90 and C
- Public domain; open source license (LGPL)
- Public domain libraries; vendor optimized versions may exist:
 - MPI1 and/or MPI2; NetCDF/parallel NetCDF; libXML
 - mpp_io; SCRIP



NEC HPCE



OASIS3&4 community today

• CERFACS (France)

ARPEGE3 - OPA 8.1

ARPEGE4.5 - NEMO
ARPEGE3 - ORCA2LIM

• METEO-FRANCE (France)

ARPEGE4 - ORCA2
ARPEGE3 - OPA8.1

ARPEGE medias - OPAMED
ARPEGE2 - OPA TDH

• IPSL - LODYC, LMD, LSCE (France)

LMDz - ORCA4
LMDZ - ORCA2LIM

LMDz - NEMO
LMDz - OPA ATL3/ATL1
ECHAM4 - ORCA2

• MERCATOR (France)

• MPI - M&D (Germany)

ECHAM5 - MPI-OM
PUMA - C-HOPE
ECHAM5 - E-HOPE

ECHAM5 - C-HOPE
EMAD - E-HOPE
ECHAM4 - E-HOPE

• ECMWF (UK)

IFS Cy23r4 - E-HOPE
IFS - CTM

IFS Cy15r8 - E-HOPE



OASIS3&4 community today

•CGAM-Reading (UK)	HadAM3	-	ORCA2
•Southampton University (UK)	Inter. Atm	-	OCCAM lite
•UCL (Belgium)	LMDz	-	CLIO
•SMHI (Sweden) ECHAM - RCA	RCA(region.)	-	RCO(region.)
•U. of Bergen (Norway)	MM5	-	ROMS
•KNMI (Netherlands)	ECHAM5	-	MPI-OM
•DMI (Danemark)	ECHAM	-	HIRLAM
•INGV (Italy)	ECHAM5	-	MPI-OM
•IRI (USA)	ECHAM4	-	MOM3
•JAMSTEC (Japan)	ECHAM4	-	OPA 8.2
•BMRC (Australia) BAM - MOM4	BAM3	-	ACOM2
•U. of Tasmania (Australia)	Data Atm.	-	MOM4
•RPN-Environment Canada (Canada)	MEC	-	GOM



OASIS3: coupled model configuration


- In text file *namcouple*
 - read by OASIS3 main process
 - distributed to component model PSMILes at beginning of run
- Configuration parameters:
 - total run time
 - component models
 - number of coupling fields
 - for each coupling field:
 - source and target names (end-point communication) (*var_name*)
 - grid acronym (*grid_name*)
 - coupling and/or I/O status
 - coupling or I/O period
 - transformations/interpolations




OASIS3: communication

PSMILE based on MPI1 or MPI2 message passing


- Parallel communication between parallel models and Oasis3 interpolation process



- Direct communication between models with same grid and partitioning



- I/O functionality (automatic switch between coupled and forced mode): GFDL mpp_io library






OASIS3: interpolations/transformations

➤ separate sequential process Oasis3

- ✓ neighbourhood search
- ✓ weight calculation
- ✓ interpolation per se during the run



➤ on 2D scalar or vector fields

- Interfacing with RPN Fast Scalar INTerpolator package
 - nearest-neighbour, bilinear, bicubic for regular Lat-Lon grids
- Interfacing with SCRIP1.4 library
 - nearest-neighbour, 1st and 2nd order conservative remapping for all grids
 - bilinear and bicubic interpolation for «logically-rectangular» grids
- Bilinear and bicubic interpolation for reduced atmospheric grids
- Other spatial transformations: flux correction, merging, etc.
- General algebraic operations



OASIS3: Conclusions


- Since 01/05, ~10 new groups asked for latest OASIS3: (France, Italy, UK, Canada, USA, Japan, Australia, China)
- Maintenance and user support as usual.
- A new official version planned beginning of next year.
 - Few bug fixes
 - Adaptations for NAG compiler
 - Additional communication technique (Unix sockets/RPN)
- Hopefully, community will progressively switch to OASIS4



The OASIS4 coupler

Situation today:

- Higher resolution, parallel and scalable models
- Higher coupling frequencies desirable
- Higher number of models and (3D) coupling fields
- Massively parallel platforms



Need to optimise and parallelise the coupler

OASIS4 is composed of:

- a Driver
- a Transformer
- a new PRISM System Model Interface Library



OASIS4: PSMILe API (1/2)

• Initialization:

```
call prism_init_comp (comp_id, comp_name, ierr)
```

```
call prism_get_localcomm (comp_id, local_comm, ierr)
```

• Definition of grid (3D) and grid points

```
call prism_def_grid (grd_id, grd_name, comp_id, grd_shape, type, ierr)
```

```
call prism_set_corners (grd_id, nbr_cnr, cnr_shape, cnr_array, ierr)
```

```
call prism_set_mask (msk_id, grd_id, msk_shape, msk_array, ierr)
```

```
call prism_set_points (pt_id, pt_name, grd_id, pt_shape, pt_lon,  
pt_lat, pt_vert, ierr) ...
```

• Coupling or I/O field declaration

```
call prism_def_var (var_id, var_name, grd_id, pt_id, msk_id,  
var_nbrdims, var_shape, var_type, ierr)
```



OASIS4: PSMILe API (2/2)

- End of definition

```
call prism_enddef (ierr)
```

- Coupling or I/O field sending and receiving:

- in model time stepping loop
- end-point communication; follows user externally-defined configuration

```
call prism_put (var_id, date, date_bounds, var_array, info, ierr)
```

```
call prism_get (var_id, date, date_bounds, var_array, info, ierr)
```

- Termination

```
call prism_terminate (ierr)
```



OASIS4: component model description (XML)

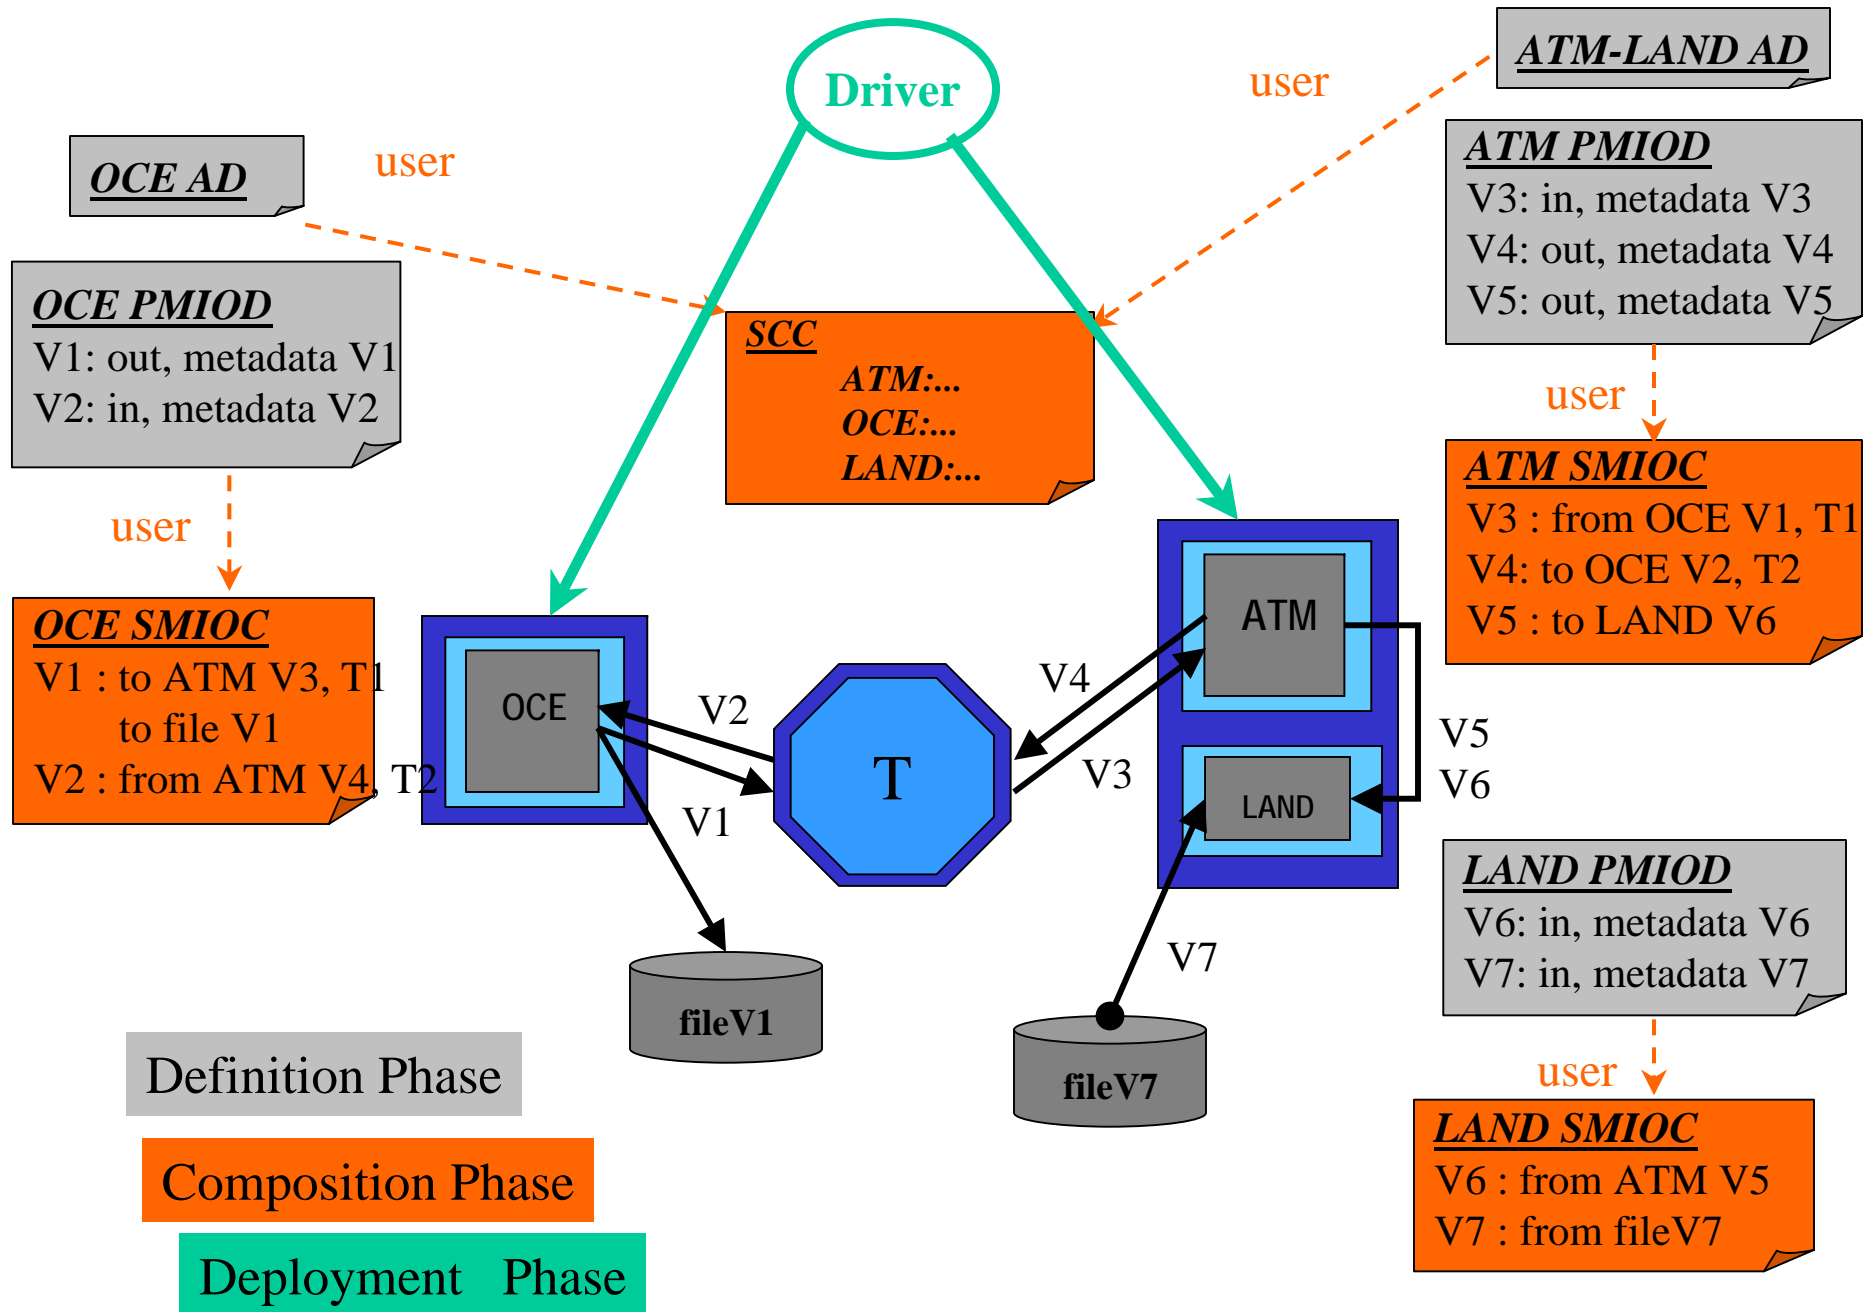
- *For each application (code):*
 - one Application Description (AD):
 - possible number of processes, components included, etc.
- *For each component in the application:*
 - one Potential Model Input and Output Description (PMIOD)
 - component general characteristics
 - grid information
 - potential I/O or coupling variables (prism_put/prism_get):
 - name (NetCDF CF convention)
 - units
 - numerical type
 - associated grid
 - intent -input and/or output



OASIS4: coupled model configuration (XML)

- (Through a GUI) , the user produces
 - a Specific Coupling Configuration (SCC):
 - start date and end date
 - start mode (MPI1, MPI2)
 - applications, components
 - host(s), etc.

- *For each component,*
 - a Specific Model Input and Output Configuration (SMIOC)
 - grid information: chosen resolution, ...
 - I/O or coupling variables:
 - name, units, valid min max, grid
 - activated intent -input and/or output
 - source and/or target (component and/or file)
 - coupling or I/O dates
 - transformations/interpolations






OASIS4: the Driver

Driver


- Manages a static coupling:
 - launches models at the beginning of the run (MPI2)
or
 - models are started directly in the run script (MPI1)
- Participates to communicator set-up
- Reads the SCC and SMIOCs XML user-defined configuring information
- Distributes it to the component PSMILes

OASIS4 communication (1/2)

- Model interface library: PSMILe based on MPI1 or MPI2
- Parallel communication **including repartitioning**:
 - based on geographical description of the partitions
 - parallel calculation of communication patterns in source PSMILe




Same grid, different decomposition
 ⇒ direct repartitioning:



Different grid and decomposition
 ⇒ via parallel Transformer:

OASIS4: communication (2/2)


- end-point communication
- parallel 3D neighbourhood search in each source process PSMILE
- parallel multigrid search algorithm: high efficiency
- extraction of useful part of source field only
- one-to-one, one-to-many
- parallel I/O (vector, bundles, vector bundles) : GFDL mpp_io, parNetCDF





OASIS4: interpolations/transformations

- Local transformations in PSMILE (below prism_put, prism_get):
 - Source time averaging, accumulation
 - Target time interpolation (I/O only)
 - Addition/multiplication by scalar
 - Statistics
- Interpolations in parallel Transformer (loop on PSMILE requests):
 - nearest-neighbour 2D, "none" in the vertical
 - nearest-neighbour 3D
 - bilinear, "none" in the vertical
 - bicubic, "none" in the vertical
 - trilinear





Grids supported by OASIS4

- **Regridding, repartitioning, I/O:**
 - **Regular in lon, lat, vert ("Reglonlatvrt"):**
 - lon(i), lat(j), height(k)
 - **Irregular in lon and lat, regular in the vert ("irrlonlat_regvrt"):**
 - lon(i,j), lat(i,j), height(k)
 - **Irregular in lon, lat, and vert ("irrlonlatvrt") (*not fully tested*)**
 - lon(i,j,k), lat(i,j,k), height(i,j,k)
 - **Gaussian Reduced in lon and lat, regular in the vert ("Gaussreduced_regvrt")**
 - lon(nbr_pt_hor), lat(nbr_pt_hor), height(k)
- **Repartitioning and I/O only:**
 - **"Gridless" fields**
 - no geographical information attached
 - local partitions described in the global index space (prism_def_partition)
- **I/O only:**
 - **Unstructured grids ("unstructlonlatvrt")**
 - lon(npt_tot), lat(npt_tot), height(npt_tot)



OASIS4: conclusions

- OASIS4 tested and run with toy examples on:
 - Intel Pentium 4 Workstation Cluster
 - SGI O3000/2000 server with MIPS 4 processors and IRIX64
 - SGI IA64 Linux server Altix 3000
 - NEC SX6
 - AMD 2800 Cluster
 - IBM Power 4
- OASIS4 now being used in a reduced number of real applications:
 - MOM4 + toy atmosphere
 - At SMHI, for regional coupling
 - At IFM-GEOMAR (Kiel) in pseudo-models to interpolate high-resolution fields.
 - GEMS: coupling between atmospheric chemistry and atmospheric dynamics
- OASIS4 full public release planned within a year:
 - Global search
 - Full validation of existing interpolations
 - 2D conservative remapping



The end