DD, Chapter 7, draft version

1. Introduction

Demonstration runs consist of a set of short experiments (most of them covering 1 model year of integration), documenting the technical functioning, the usability and the correctness of the approach. As a first step we demonstrate the portability of PRISM coupler by integrating a “toy” coupled configuration on different platforms. After that we demonstrate that the system is capable of coupling two “real” general circulation models (GCMs). We document that combinations run successfully and at the same time fulfill several technical requirements, concerning coupling functionalities and the scripting environment. These technical aspects are:

- The portability
- The reproducibility, employing different numbers of CPUs
- The frequency of communication
- The frequency for archiving output files
- The restartability

On the other hand, those experiments extend the test of portability to some real model combinations which are integrated on several platforms. Later, some specific coupling configurations are demonstrated, such as an atmospheric regional downscaling over the Northern Europe and two ocean biogeochemistry models on a general ocean circulation model.

2. Experiments design

In the present section we introduce a series of experiments to demonstrate the above mentioned items.

2.1 Experiments using Toy Models (TOYCLIM system)

The TOYCLIM system is a combination of toy global models (i.e., empty models, in which no physics is included) for which both coupling communication and interpolation are realistic, representing the atmosphere (ATM), the ocean (OCE) and the atmospheric chemistry (CHE), and using OASIS3 software (Valcke 2004). An essential technical aspect of a flexible earth system modeling software system is its portability, that was demonstrated by integrating the OASIS3 TOYCLIM system on different platforms as described below.

E1. The portability experiment:

OASIS3 ATM model coupled with CHE and OCE models were integrated for a period covering 6-day of model integration on the following platforms: NEC SX-6 (node barolo), SGI IRIX-64 (node p1), VPP5000 (node tora), IBM Power-4 (node hcpa) and SUN E15k (node yang). Information about platforms is in Section 3. The experiment was defined as follows:

- 1 processor for ocean toy model
- 3 processor for atmospheric toy model
- 3 processor for atmospheric toy chemistry
- exchanging fields every ocean model time step
- extensive standard output for exchange control

The sole difference between these runs was the use of different platforms. Details of models settings are provided in Table 1. In particular, toy models carry only one vertical level, the surface level, since they were built to illustrate how effective and user-friendly the coupling software OASIS3 is in exchanging variables and performing all necessary transformations to the fields sent from source to target model (e.g., interpolations, masking, change of units, etc).

Table 1: Main settings of TOYCLIM models

OASIS3 TOYCLIM models	TOYOCE	TOYATM	TOYCHE
Grid	Non parametric grid 182x152	T31 Gaussian grid 96x48	T31 Gaussian grid 96x48
Time steps to end of run (6-day run)	36	144	72
Time step in seconds	14400 sec 6 time steps/day	3600 sec 24 time steps/day	7200 sec 12 time steps/day

Detailed information about the OASIS3 TOYCLIM system settings behavior and demonstration is in Carril (2004).

2.2 Experiments using Real GCMs

Second step focus on combinations made by two "real" component models: an atmospheric GCMs coupled with an ocean GCMs as detailed in table 2.

Table 2: GCMs combinations tested on diverse platforms. The node were the experiments were integrated and the number of CPUs used to integrate the atmospheric component are indicated in every box. Ocean components were always integrated using 1CPU.

Platform vs Models Combination	NEC SX-6	SGI IRIX-64	Fujitsu VPP5000
ECHAM5 + MPI-OM	- Node barolo - m=1	- Node p1 - m=3	- Node tora - m=1
ECHAM5 + ORCA-LIM	- Node barolo - m=3		
ARPEGE4 + ORCA			- Node tora - m=3
LMDZ + ORCA-LIM	- Node mercure - m=1		- Node tora - m=1
HadAM3 (SRES) + ORCA	- Node cs - m=3		

E2. The standard experiment:

For every combination indicated in Table 2 (real atmosphere model + real ocean model + Oasis3), the experiment was defined as follows:

- 1 proc for real ocean model
- m proc for real atmospheric model (as indicated in table 2)
- exchanging fields once a day
- 1 model year of simulation
- re-starting integration monthly
- extensive standard output for exchange control
- saving output files as monthly means

Basing on the "standard experiment", a set of "deviation experiments" were done to test diverse technical aspects. E3 to E6 describe the deviation experiments.

E3. Deviation experiment testing the reproducibility to the number of CPUs:

For the combination ECHAM5 + MPI-OM integrated on the NEC SX-6 (node barolo), the experiment was defined as E2, except by using

- *4 proc for the atmospheric model*

E4. Deviation experiment testing the functionality of exchanging frequencies:

For the combination Arpege4 + ORCA-LIM integrated on the VPP5000 (node tora), the experiment is defined as E2, except by

- *exchanging fields every 2 days*

E5. Deviation experiment testing the flexibility to I/O fields:

For the combination LMDz-ORCHIDEE + ORCA-LIM integrated on the NEC SX6 (node mercure), the experiment is defined as E2, except by

- *saving output files twice-a-day for the atmosphere and every 5 days for the ocean*

E6. Deviation experiment testing the restartability:

For the combination ECHAM5 + MPI-OM integrated on the SGI IRIX64 (node p1 at KNMI), 1 model year of simulation. The experiment is defined as E2, except by

- *re-starting integration once a year*

Table 3: Main model settings

Models	Grid and Resolution	Time step
ECHAM5	T21L19 64x32x19 ~5.625°x5.625° ; 19 vertical levels	24600 sec. 36 time steps/day
LMDZ	Arakawa grid (type C) 96x72x19 3.75°x2.5° ; 19 vertical levels	1800 sec. 48 time steps/day
ARPEGE-Climat	T63L31 128x64x31 2.8125°x2.8125° ; 31 vertical levels	1800 sec. 48 time steps/day
HadAM3	Arakawa grid (type B) Regular lat-lon grid N48L19 96x73x19 3.75°x2.5°; 19 vertical levels	1800 sec. 48 time steps/day
ORCA_LIM	Arakawa grid (type C) No parametric grid with two poles in the NH 182x149x31 2°x ~ 1.5°to 0.5°; 31 vertical levels	4800 sec. 18 time steps/day
MPI-OM	Arakawa grid (type C) No parametric grid with one pole in the NH 120x101x20 3°x1.8° ; 20 vertical levels	8640 sec. 10 time steps/day

Table 3 illustrates with the models resolution, grid details and time stepping of integrations. All the experiments (E2 to E6) were conducted starting from the same input files for the individual models e.g., initial 3-dim ocean variables (temperature and salinity), 3-dim atmospheric variables (vorticity, divergence, temperature, humidity), surface salinity, surface SST, surface boundary conditions for the atmosphere (sea, glacier, lake masks, topography, albedo, roughness length), the hydrological parameters, the ozone climatology, vegetation information, ocean bathymetry, basin mask, ocean grid description. Also input files for OASIS containing information about the latitudes and longitudes of the grids involved in the coupling, the grid cell surfaces, see-land masks, as well as weights and addresses used for extrapolation are common for every particular combination. Most of the experiments were performed using the SCE and SRE. Exceptions are ARPEGE+ORCA and HadOPA which are still not adapted to the PRISM standard environments. PRISM version used to run experiments E2 to E6 is prism_2-2 (except for EOL which was integrated using PRISM version target prism_2-1). Detailed information about demonstrations using real GCMs is in Carril et al. (2004).

2.3 Downscaling Experiment

The models participating in this demonstration are the Rossby Centre Atmosphere (RCA) regional model and a toyGCM (a small program pretending to be a GCM). As displayed in Figure 1, the toyGCM reads global data from disk files (data stemming from an ECHAM5 pre-industrial control run), and passed to RCA via the PRISM interface Psmile. This data communication is initiated and controlled by OASIS3. In this setup, all interpolation from the GCM grid to the finer RCM grid is taken care of within the regional model.

E7. Downscaling experiment was defined as follows:

- RCA combined with a toyGCM, no other components,
- 1 model year of simulation on the SGI Origin 3800 at NSC/Linköping
- one way from global to regional, off-line coupling
- coupling is done every 6 hours (the output frequency of the global model)
- regionalization is over the Northern Europe

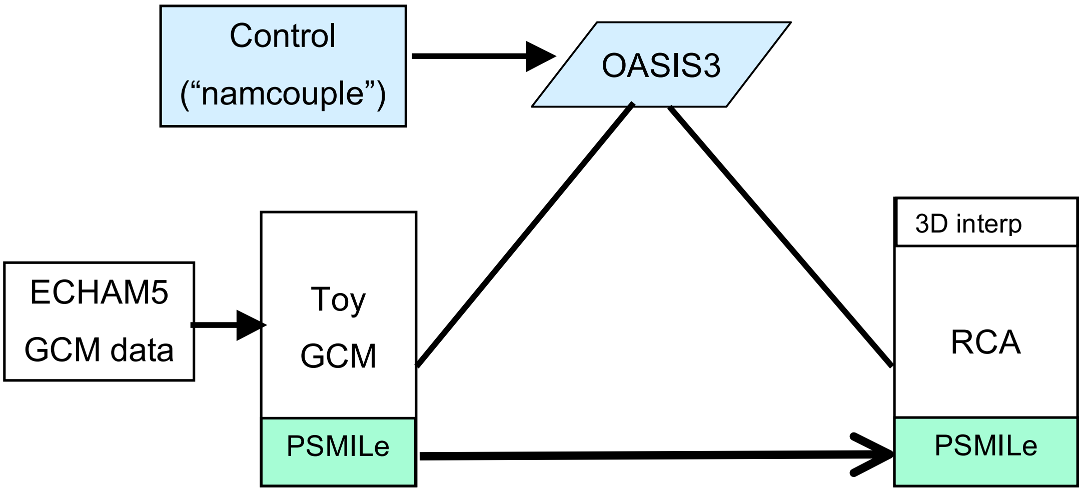

Figure 1. Setup of the demonstration of GCM-RCM coupling

Main models settings are summarized in table 4. That system is level-one compliant with the PRISM system, i.e. the technical interface is used. In addition, a physical interface is used. This would imply a level-two compliance according to an earlier preliminary standard definition before the PRISM project abandoned standard physical interfaces. The system is not yet adapted to the PRISM SCE and SRE. Detailed information about the downscaling demonstration is in Carril et al. (2004).

Table 4: Main models settings for regionalization experiment

Prism version target prism_2_1	RCA -regional atmospheric model	toyGCM
Grid and resolution	Rotated latitude/longitude 106x102 points 0.44°x0.44°	Latitude/longitude 97x39 points 1.875°x1.875°
time step	model: 30 minutes coupling: 6 hours	coupling: 6 hours
nproc	4	1

2.4 OBG Experiments

In this suite of experiments we demonstrate the modularity of the PRISM system for models that consist of a main model and one of several sub-models. The main model is the ocean model MPI-OM, which can be run standalone or in combination with one of the biogeochemistry sub-models: the Hamburg Ocean Carbon Cycle biogeochemistry model Version 5 (HAMOCC5) and the Pelagic Interaction Scheme for Carbon and Ecosystem Studies (PISCES) biogeochemistry model. The biogeochemistry sub-models themselves cannot run standalone.

E8. Ocean biogeochemistry experiment is defined as follows:

- PISCES (HamOcc) OBGC model combined with MPI-OM ocean GCM,
- no other components except sea-ice,
- 1 model year of simulation on the NEC SX-6 (node cs)

In contrast to experiments where several separate program units are mutually communicating through the PRISM coupler, the model components in these experiments are hosted in the same program unit and thus do not need a coupler. Technically, this is done by issuing subroutine calls from the ocean host model to the biogeochemistry schemes. The current model versions are such that both biogeochemistry models employ the same ocean transport routines. Also, for these tests we implemented consolidated initial conditions for the different biogeochemistry model. Thus, we arranged both biogeochemistry models to basically start from the same initial concentrations. Some extra setting information is in Table 5. For the demonstrations, the model sources were made compliant with the prism_2_2 version of the Standard Compilation Environment (SCE) and Standard Run Environment (SRE).

Table 5: Main settings of the experiments

Prism version target prism_2-2	MPI-OM	BGC models
Grid and resolution	Arakawa grid (type C) staggered 122x101x20	Inherited
Time step	8640 seconds	Inherited
Calendar	Equal months of 30 days	Inherited
nproc	1	1

In these experiments, we tested the setup conditions both for newstart as well as for continuation runs. In a newstart run, the ocean and biogeochemistry models initialize with data from files or with model specific preset data. On completion of a run, intermediate output is written to certain restart files. These files are then pre-loaded in a continuation run. Detailed information about the OBG demonstrations is in Carril et al. (2004).

3. Platforms

In the present section we present the platforms which were available to conduct demonstration runs. Table 6 describe for every node in a particular platform, the main characteristic of the system as well as version of installed libraries.

Table 6:Details about the platforms in which the demonstrations were performed

Platform (node and institution)	NEC SX-6 (node Barolo at INGV)	NEC SX-6 (node mercure at CEA)	NEC SX-6 (node cs at DKRZ)	VPP5000 (node tora at MeteoFrance)
Fortran 90	FORTRAN90/SX Version 2.0 for SX-6, Rev. 285 2003/09/25	FORTRAN90/SX Version 2.0 for SX-6, Rev. 285	FORTRAN90/SX Version 2.0 for SX-6, Rev. 274 2003/06/20	Fortran 95 UXP/V V20L20 2002/09/03
C++	C++/SX Compiler for SX-4 Version 1.0, Rev. 061 2004/01/06	C++/SX Compiler for SX-4 Version 1.0, Rev. 058	C++/SX Compiler for SX-4 Version 1.0, Rev. 057.3 2003/08/04	C++ Fujitsu UXP/V V20L20 2000/11/22
NetCDF library	NetCDF library version 3.5.1	NetCDF library version 3.5.0	NetCDF library version 3.5.0	NetCDF library version 3.5.1
MPI library	MPI-Library Version is MPI/SX, SUPER-UX R13.1 V6.7.8 (LC310052)	MPI-Library Version is MPI/SX, SUPER-UX R13.1 (LC310039)	MPI-Library Version is MPI/SX, SUPER-UX RXX.X VX.X.X (LC310039)	Fujitsu MPI UXP/V 2003/06/12
Comments	"Barolo" is the NEC SX-6 at INGV. Its configuration is 8 CPUs at 64 GFlops and 64 GB of Main Memory. For program development and compiling there is a powerful cross compiler environment, which allows the compilation of SX-6 executables on the SUN server (node SFV880; 8 CPUs, 900 MHz and 16GB with SunOS 5.8).	"Mercure" is the NEC SX-6 at CEA. Its configuration is 44 CPUs at 8 GFlops and 352 GB of Main Memory. For program development and compiling there is a powerful cross compiler environment, which allows the compilation of SX-6 executables on the Linux server (mercure; 8 CPUs, 1GHz, 16GB with Itanium2).	"Hurrikan" is the NEC SX-6 at DKRZ. Its configuration is 24 nodes of 8 CPUs at 64 GFlops and 64 GB of Main Memory per node. For program development and compiling there is a powerful cross compiler environment, which allows the compilation of SX-6 executables on the NEC TX-7 Itanium2 server (node ds8; 16 CPUs, 1 GHz and 32GB with Linux 2.4.18).	"Tora" is a 64 processors VPP5000. Each processor has a peak performance of 9.6 GFlops, and a memory of 4 GB or 8 GB. Each processor has 4 ports to connect on the high-speed crossbar. Point to point bandwith is 1.6 GB/sec, bidirectionnal. This machine has one primary PE, 7 IO/PEs and 56 secondary PEs. A sister VPP5000 : kami, is also installed at Meteo France with 60 PEs. It is reserved for Meteo France production work.

Table 6 (cont.):Details about the platforms in which the demonstrations were performed

Platform	SGI Origin 3800 at NSC/Linköping	SGI IRIX-64 (node p1 at KNMI)	IBM Power 4 (node hpca at ECMWF)	SUN E15k (node yang at ???)
Fortran 90	MIPSPro Compiler, Version 7.41	MIPSPro Compiler, Version 7.41	F90 Compiler, Version 7.1.1.3	Sun Fortran 95, 7.1 Patch 112762-05

				2003/10/27
C++	MIPSPro Compiler, Version 7.41	MIPSPro Compiler, Version 7.41	VisualAge C++ Professional / C for AIX Compiler, Version 5	Sun C 5.5 Patch 112760-04, 2003/10/03
NetCDF library	NetCDF library version 3.5.0	NetCDF library version 3.5.1	NetCDF library version 3.5.0	?
MPI library	SGI MPI 4.3, as component of SGI MPT 1.8	SGI Message Passing Toolkit 1.9.1	MPI library is MPI1, part of the parallel environment 3.2.0.17	mprun: Sun HPC ClusterTools 5 11 Feb 2003 CRE 2.0,REV=2003.02.10.5.37
Comments	www.nsc.liu.se/systems/sgi3k 128 CPUs, each with 1 GFlops peak performance and 1 Gb memory	"TERAS" is a 1024-CPU system consisting of two 512-CPU SGI Origin 3800 systems. This machine has a peak performance of 1 TFlops per second and it is fitted with 500MHz R14000 CPUs, organized in 256 4-CPU nodes and it is equipped with 1 TByte of memory (10 TByte of on-line storage and 100 TByte of near-line Storage Tek). TERAS consists of 45 racks, 32 racks containing CPUs and routers, 8 I/O racks and 5 racks containing disks. The operative system is IRIX64 6.5.23m (node p1)	IBM Power 4, LPAR 8 CPUs, 16 GB Main Memory; hpca.ecmwf.int	Sun E15k, 72 CPUs, 144 GB Main Memory, executor.zmaw.de

4. Results

We present a series of diagnostics that illustrates the functionality of the overall PRISM system on different configurations and demonstrate some technical and physical aspects of the exchange between the model components and with the coupler. Nevertheless, there are not scientific diagnostics, since these are beyond the scope of the project; the correct performance of the individual models remains the task of the individual model developers. Hereafter, diagnosis to assess the PRISM performance are divided into “technical”, “physical” and ‘system” diagnostics.

4.1 Technical Diagnostics

Because individual models were designed separately, testing the interface between the components is a crucial point. A simple set of *technical diagnostics* consist on plotting the exchanged fields before and after they are passed through the coupler. Hereafter, we present technical diagnostics from some particular experiment, in which the fields plotted are the ones saved by the coupler in output files thanks to the option EXPOUT in the namcouple.

We have first selected the experiment E2 integrated using the combination EOL, on the NEC SX-6 (node barolo). List 1 presents the list of transformations defined in the namcouple file for the SST at every time step of exchange. SST undergoes a mask and an extrapolation over land, then a SCRIPR interpolation, a change of units (Celsius to Kelvin) and an invert transformation north/south – west/east before being sent to the atmosphere. SST field before and after to be exchanged at a particular time step of coupling is in Figure 2. Table 7 complements the figure with some statistics as collected from the log file cplout. These values are written there thanks to the options CHECKIN for the source grid and CHECKOUT for the target grid specified in the namcouple.

List 1: Transformations defined in the namecouple for the SST

```
# Field 1 Sea_surface_temperature (o->a 1)
#
SOSSTSSW SSTATMOS 1 #Dto2a  7  #Cnfileow  #Exp
#Lono #Lato  #Lona #Lata  or1t  at21  SEQ=1  LAG=#Lago2a
P  2 P  0
CHECKIN MASK EXTRAP SCRIPR BLASNEW CHECKOUT REVERSE
INT=1
9999.99e+06
NINENN 2 1 2
DISTWGT LR SCALAR LATLON 10 4
1.0 1
CONSTANT 273.15
INT=1
NORSUD WSTEST
```

Table 7: Extreme values for the SST in the source and target grid at a given time step of coupling

SST (o2a)	**Checkin**	**Checkout**
Global average	0.8526744E+01	0.2834310E+03
Global maximum	0.3039561E+02 Pt i-j= 28 61	0.3032602E+03 Pt i-j= 24 14
Global minimum	-0.2100092E+01 Pt i-j=119 137	0.2710499E+03 Pt i-j= 53 27
Ocean grid points	16751	1584
Ocean average	0.1377373E+02	0.2864425E+03
Ocean maximum	0.3039561E+02 Pt i-j= 28 61	0.3032602E+03 Pt i-j= 24 14
Ocean minimum	-0.2100092E+01 Pt i-j=119 137	0.2710499E+03 Pt i-j= 53 27
Land grid points	10367	464
Land average	0.4865439E-01	0.2731500E+03
Land maximum	0.2679043E+02 Pt i-j=152 91	0.2731500E+03 Pt i-j= 1 1
Land minimum	-0.2088089E+01 Pt i-j= 82 10	0.2731500E+03 Pt i-j= 1 1

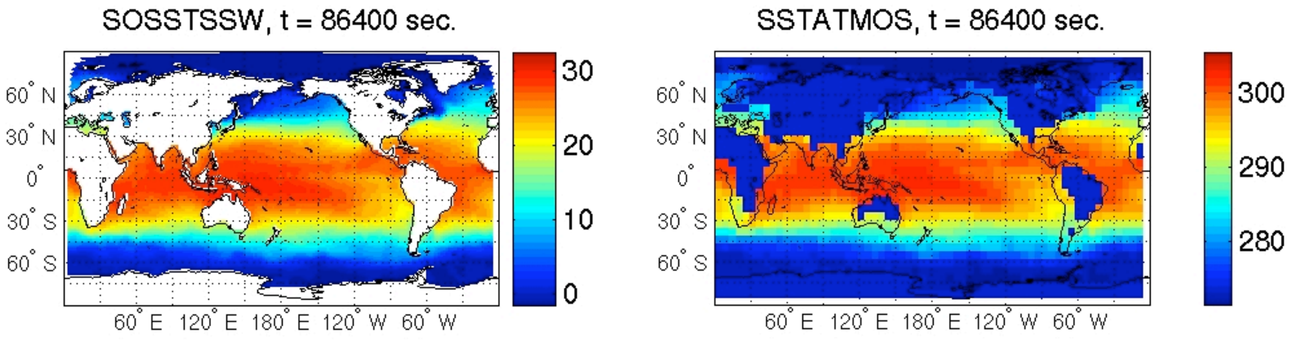

Figure 2: The SST field in the source (ORCA_LIM, left) and target (ECHAM5, right) grid. Units are ºK.

Moreover, the PRISM system is technically efficient interpolating fields using other methods and applying diverse transformations. Figure 3 (from E2 integrating LMDz+OPA on the NEC SX-6, node mercure at CEA) presents an atmospheric field (the U component wind stress along X axis) which undergoes an invert transformation north/south, a mask and an extrapolation over land. Then, a bicubic interpolation before being sent to the ocean.

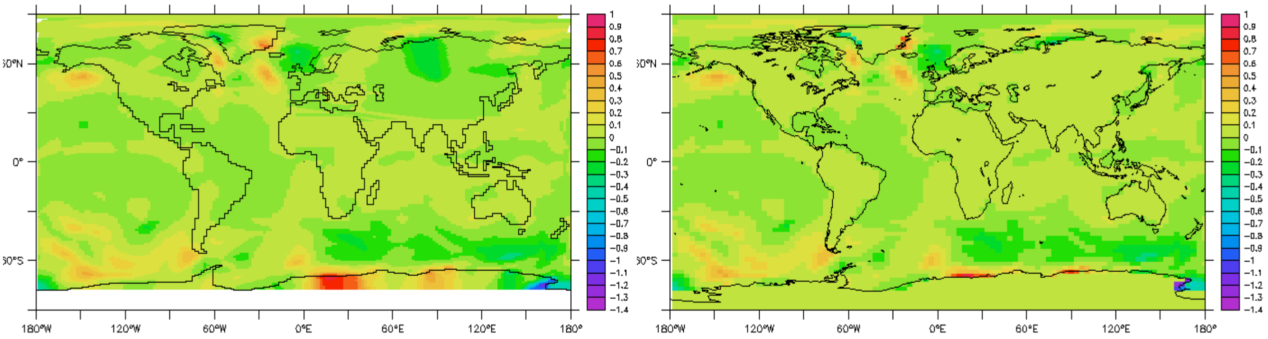

Figure 3 : The wind stress field in the source (LMDZ, left) and target (OPA, right) grid. Units are Pa.

In essence, technical diagnostics demonstrate that the coupling interface is physically consistent, the components are able to communicate with each other adequately, models communicate using common units, or appropriate change of units are done by the coupler, the fields are interpolated adequately and the frequency of communication is defined properly. For more details see Carril et al. (2004).

4.2 Physical Diagnostics

Let us now plotting some particular fields after performing a short integration to physically ensure the correctness of their treatment. As a first example we present a physical diagnostic from a GCMs coupling. The plots in Figure 5 cames from the integration of the experiment E2 and E4, using the combination Arpege+OPA on the VPP5000 (node tora). The figure displays zonal-mean profiles of the sea-level pressure for the year mean and the extreme seasons. Extremely low-pressure values, on the order of 985hPa, are found in the belt around Antarctica. The seasonal shifts in the high-pressure belts near 30^{0}N and 30^{0}S and their intensification during the winter season are clearly shown. The greater differences in the zonal mean sea level pressure between E2 (Figure 5, left panel) and E4 (Figure 5, right panel) are located to higher latitudes where the variability of this field is greater.

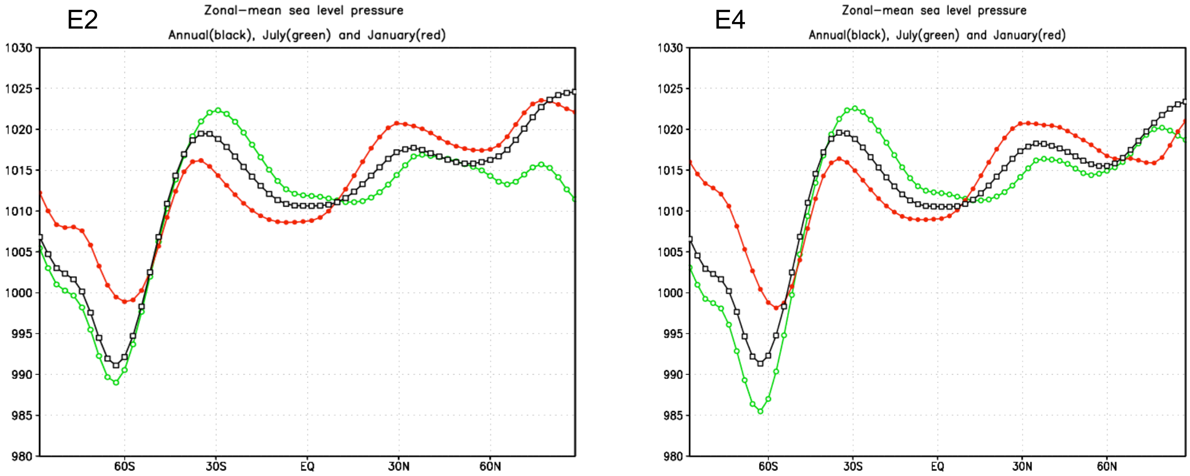

Figure 5: Meridional profiles of the zonal mean sea level pressure for annual mean conditions (black) and for January (red) and July (green) in hPa. Left panel is for E2 while the right one is for E4.

We move now to the regional downscaling experiment. Figure 6 shows a comparison of observations with the downscaling demonstration, after 1-year of integration. High resolution observations from pre-industrial time are not available, thus recent climate observations must serve to indicate plausibility of results. The CRU observations used here, represent 30 years of recent climate (1961–1990) while the model were applied to an arbitrary year of pre-industrial climate. Still, many similarities are evident. In northern Europe the annual mean temperature of this one year is within one standard deviation from the mean in the 1961-1990 climatology. In southern Europe, where this particular year is warmer than the climatology, the annual mean temperature lies within two standard deviations of the 1961-1990 average.

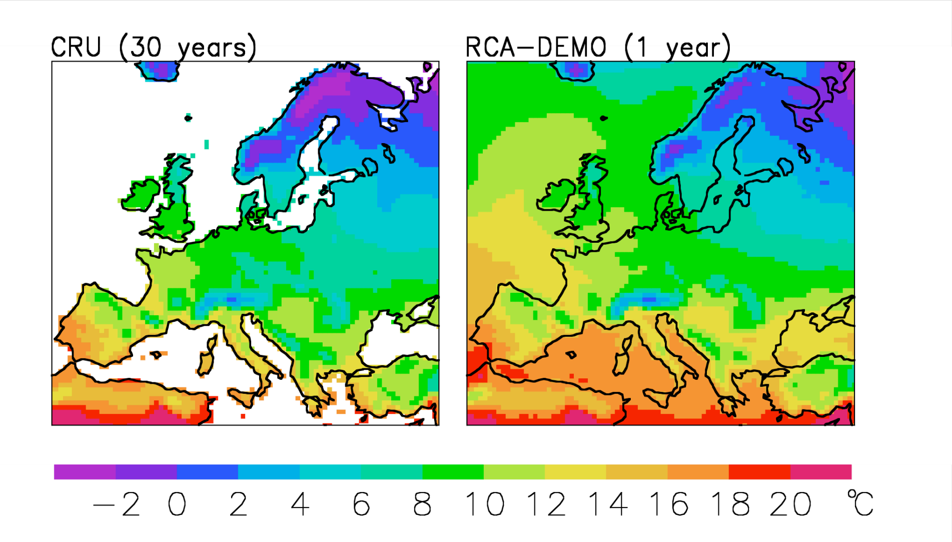

Figure 6: 30-year mean of 2-m temperature observations from the CRU database (left) and 1-year mean of 2-m air temperature in the regional model RCA (right)

Finally, we illustrate with a comparison of selected tracers as part of OBGC demonstrations. Most features of the biogeochemistry models are equilibrium values and thus cannot be obtained at a reasonable reliability from short simulations such as the demonstration runs performed for the demonstrations. However, visual inspection of the results may serve as a means to evaluate the basic functioning of the models. We considered all biogeochemistry tracers and selected three lead variables that are defined in both BGC models.

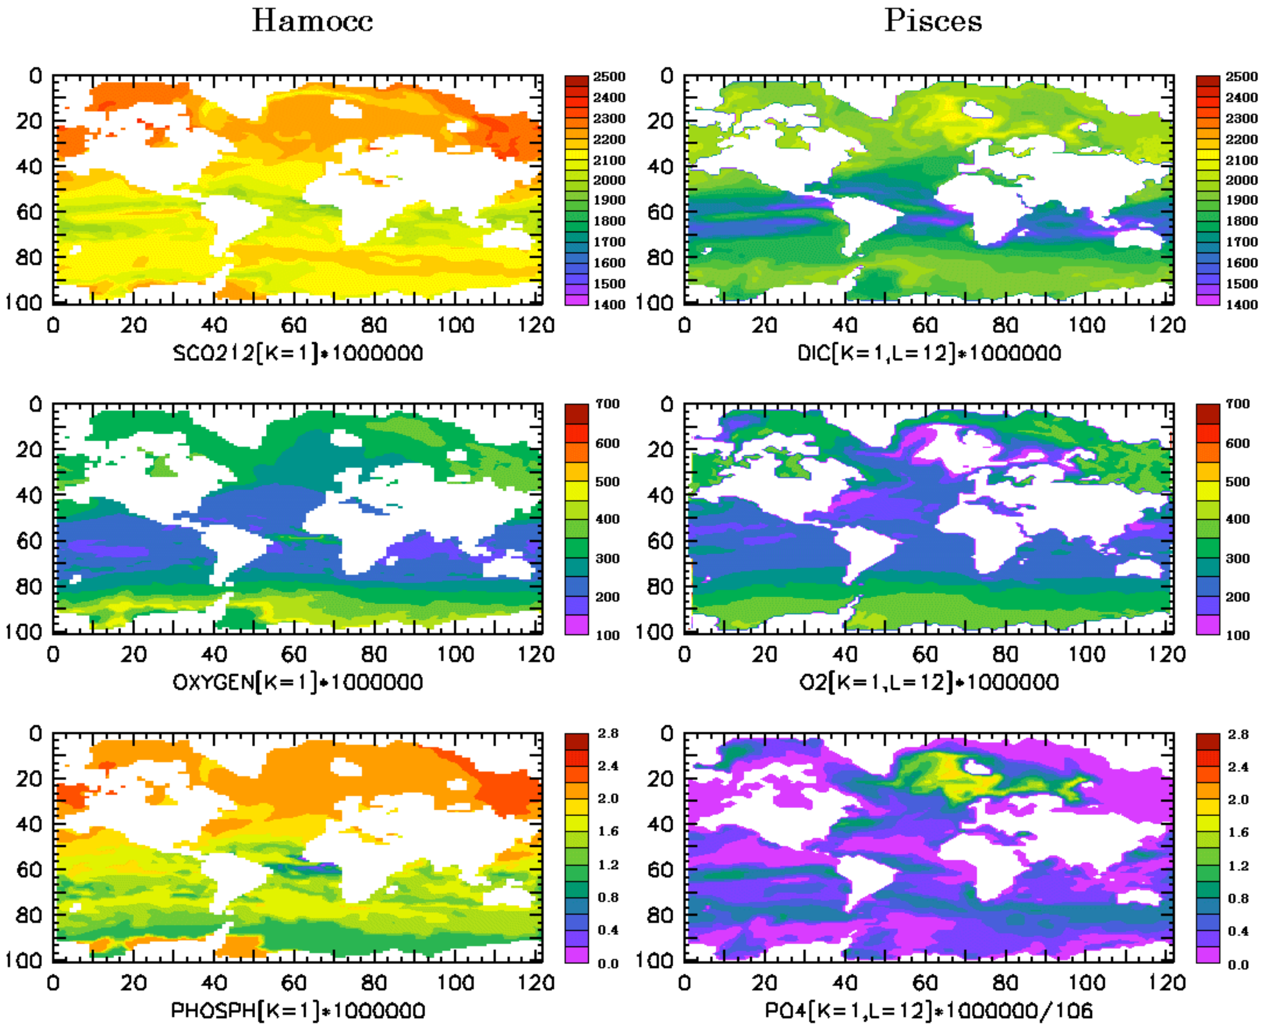

Figure 7: Comparison of biogeochemistry results for 12-month runs. The Hamocc and Pisces models were started from the similar initial conditions.

As illustrated in Figure 7 , after one year of integration, all the modeled variables remain within reasonable bounds of observations, for both model combinations. The Hamocc MPI-OM combination remains closer to initial values, as expected because Hamocc was originally developed within MPI-OM. The Pisces MPI-OM, although still within reasonable values, shows low concentrations for most surface tracers. This could be due simply to the differences in physical transport between MPI-OM and the OPA model, which is the physical model in which Pisces has been developed.

4.3 System Performance

To illustrate the PRISM system performance we focus on E2 experiment using ECHO combination on the VPP5000 (node xbar). Table 8 presents the standard job accounting for one of the 12 jobs of the E2 experiment. Details for each processor are in the table.

Table 8: System performance for the a job of 1-month of E2 integration using the combination ECHO on the VPP500 (node xbar).

```
Operating System              : UNIX_System_V xbar00 4.1 ES 3 5000
NQS Queue Name                : mlexpr
NQS Request ID                : 46908.xbar00
Job Name (ID)                 : echo (29554)
Job Execution Mode            :       SHAR
Job Class                     :          4
Job Created                   :   07/22/04 11:49:49
Job Queued                    :   07/22/04 11:49:50
Job Started                   :   07/22/04 11:49:51
Job Ended                     :   07/22/04 12:14:35
Elapsed Time (Secs)           :      24:44 (      1484     )
Time Spent for Alloc (Secs)   :          0 (         0     )
User CPU Time (Secs)          :      32:55 (      1975.6923)
User Vector CPU Time (Secs)   :      11:22 (       682.5708)
System CPU Time (Secs)        :         20 (        20.5054)
Allocated Memory on each VP   :       1216 Mbytes
Maximum memory used           :        448 Mbytes
MRFS                          :          0 Mbytes
I/O Chars Transferred         :       6746 Mbytes (   7074239608 Bytes)
VPs Allocated / Used          :          4 /   4
PEs Allocated / Used          :          4 /   4
Job Efficiency                : CPU = 33.62%, Memory = 36.84%, Vector = 34.19%

VP Usage Report
===================
ID PEID T System CPU(Vector)    User CPU(Vector)       Total CPU   MEM I/O Chars
== ==== = ================== ======================== =========== ==== =========
 0  2.0 s       6.600( 0.00)     947.380(      0.009)     953.980  192   831910k
 1  d.0 s       9.021( 0.00)      10.033(      0.501)      19.054  192   781994k
 2 12.0 s       1.359( 0.00)     858.968(    583.371)     860.328  448  5239626k
 3 17.0 s       3.524( 0.00)     159.309(     98.688)     162.834  384  56224212
```

Note that the ECHAM5 model runs on Virtual Processor (VP) 2 and MPI-OM on VP 3. OASIS3 is on VP 1 and the VP 0 is a server process used by the mpiexec command. This process uses only scalar CPU and waits for the actions (like spawning new processes). It is not mandatory to use mpiexec. OASIS3 can do the spawning of processes, but then its CPU usage will be close to that of mpiexec.

Note also that the Vector Ratio time of ECHAM5 is 583/860 = 67.8% while that of MPI_OM is 99/162 = 61.1%. Of course more optimizations can be done through directives in order to enhance the vectorization. The code has NOT been modified at all, it is the original code as it runs on NEC. The overall Vector ratio: 34.19% refers to all processes in the job, so it includes the mpiexec scalar process and OASIS3, also mostly scalar.

5. Conclusions and Future Work

Aim of PRISM demonstrations was to probe a series of tasks regarding the feasibility and the correctness of the PRISM approach, based on different model combinations integrated on diverse platforms. Technical and physical diagnostics have reveal the good functioning of the system according to the different levels of development reached for particular combinations, as described in along the document. An extended report on demonstrations could be read in Carril (2004) and Carril et al. (2004). Nevertheless, we now that the system could be still develop and hereafter we briefly enumerate some steps to be done in the future.

Concerning coupling GCMs, some particular combinations are one-level PRISM compliant, but some extra work is still needed to adapt those components to the PRISM Standard Environments. Moreover, some new components could be further adapted to the PRISM coupler, while new combinations could be assembled and others could be ported to new platforms.

About regional downscaling, in the present stage all interpolation from the GCM grid to the finer RCM grid is taken care of within the regional model. This is the traditional way. In the long term a complete outsourcing of interpolation into the coupler is envisaged, however it was not possible within the limited time of the PRISM project. As a further step, OASIS4 (not available at the time of the demonstration) will contain a multiple 2D interpolation, which allows outsourcing parts of the 3D interpolation. On the other hand, the system is not yet adapted to the PRISM SCE and SRE due to incompatibilities with the unified modeling approach at SMHI. This point will be addressed at a later time.

Regarding the concept of sub-models embedded in container models (as the case of the OBGC) future plans concern a completion of the current model combination, e.g. in terms of grid resolution. Also planned is the integration of a third OBM, namely the Hadocc model. Finally, the scientific exploit of the models awaits its realization.

It is suggested to use the runs and combinations tested for PRISM demonstrations as a base of longer runs for scientific exploitation. All PRISM partners were invited to participate in the “Open House Experiments”, and they were encouraged to coordinate these experiments within the PRISM community, incorporate PRISM combinations into the scientific plans of their institutions, and let the PRISM system keep track of these experiments. Checking out every possible combination under realistic operational conditions will contribute to a final debugging and optimization of the system.

Add some words about CRAY and MOM?

6. References

Carril A. F. (2004), The wp5 TOYCLIM Demonstration Run report. PRISM Report Series #13, n pp.

Carril A. F. et al. (2004), PRISM Demonstration Runs. PRISM Report Series #14, n pp.