

Coupling General Circulation Models on a Meta-computer

Wolfgang Joppich¹ and Johannes Quaas²

¹ Fraunhofer Institute for Algorithms and Scientific Computing (FhG-SCAI), Schloss Birlinghoven, D-53754 Sankt Augustin, Germany

² Laboratoire de Météorologie Dynamique du CNRS, 4 place Jussieu, F-75252 Paris, France

Abstract. Sophisticated climate models, describing the circulation in atmosphere and oceans, respectively, have been coupled together. This has been done using two different remote supercomputers in an efficient way for the first time.

Two state-of-the-art climate models have been chosen, a coupling tool has been adapted, and a new developed message-passing library has been used.

1 Introduction

One of the most important research topics at the moment is the understanding and the prediction of the climate change due to anthropogenic influences. The most useful instrument therefore are comprehensive models, which describe the climate system. At least the circulation in its two main subsystems, the atmosphere and the ocean, has to be considered in climate research.

The most efficient way to take into account several climate subsystems is to couple together models of the respective systems. Models of different parts of the climate system generally are developed independently by different research groups. Coupling the existing codes of subsystems is reasonable, because the only interaction between two subsystems occurs at their common surface. Only small changes within the approved models then have to be made.

The integration of climate models requires extremely large computing capacities. To become more and more reliable, the resolution of the model grids has to be increased and more physical processes in the subsystems and their interaction have to be taken into account. Both approaches lead to the need of even larger computing power. Beside the development of more efficient numerical solutions for the integration it is the application and adaption of new technical resources that promises to satisfy the necessary computing requirements.

In the latter area massively parallel systems have been introduced very successfully during the last years. In climate modeling, the integration of the physical equations is distributed to a large number of processors using domain decomposition technique.

If the models of two or more climate subsystems are coupled, one can consider

to distribute the computation of the two models, not only on different subsets of the processors of one computer, but use different computers for the different models. In the general case, each of the models should be allowed to run in parallel. This technique is called “meta-computing”. This could have at least three advantages: First, in the ideal case the overall integration time will be reduced to that of the most time-consuming component. Secondly, for each model the suitable hardware can be chosen. This can improve largely the efficiency, if e.g. one model runs more efficient on a vector machine, and the other one e.g. on a parallel system. If codes are written for different hardware, a transfer of the codes to a common hardware can be avoided. Thirdly it would be easier for different research centers, which develop models of different parts of the climate system, to work together. While coupling the codes, their respective models can reside on their own machines.

The goal of our work presented here is to test the hardware needed and to develop or adapt, respectively, the suitable software needed for the coupling of climate models on a meta-computer. We developed a new coupled model using two approved general circulation models of the atmosphere and the ocean to gain the experience with the coupling of climate models on a meta-computer. We did this in the framework of a project called “Distributed Computation of Weather- and Climate Models”, carried out by the Institute for Algorithms and Scientific Computing of the German National Research Center for Information Technology (GMD-SCAI, St. Augustin) together with the Alfred Wegener-Institute for Polar and Marine Research (AWI, Bremerhaven). This project was part of the “Gigabit Testbed West”-project (GTBW) funded by the German Research Network (DFN) [1][2]. The GTBW provided an ATM-network with a bandwidth of 2.4 GBit/s and a latency of 7ms, connecting the GMD and the Forschungszentrum Jülich (FZJ) crossing a distance of about 100km [3]. Beside our project, a variety of other scientific applications have been tested on the testbed. Connected to the gigabit network were an IBM SP2 at GMD’s site and a Cray T3E at the FZJ.

Known to the authors are three other works on the topic of the distributed computation of coupled climate models.

The first was the work of C. R. Mechoso et al. 1993 at the University of California, Los Angeles [4]. Not only the atmosphere- and the ocean model ran distributed, but the atmosphere model was also splitted in two parts. Although this work was a big step concerning the idea of meta-computing in climate research, the results have been disappointing because of the insufficient performance of the used technique. The respective programs ran each sequential, and the model needed 193s for the computation of one day. The time for the communication between the models consumed 1930s, exactly the ten-fold time. The standard deviation for the communication time had a value of 1220s (10 measurements) and therewith the same order of magnitude as the mean value.

1998 a coupling of French climate models has been carried out with a similar coupling technique as used by us [5]. Also in this work the models ran each se-

quential. An integration over one year has been carried out. But as the network was not reliable, some phases of the integration had to be done several times. In the third work concerning this topic by Météo-France an integration of the coupled model had only been simulated. This work focused on the network security [6].

2 Description of the Coupled Models

We have chosen two well-known climate models. As atmosphere model we used the Integrated Forecast System (IFS) of the European Centre for Medium-Range Weather Forecasts (ECMWF), and as ocean model the Modular Ocean Model (MOM) of the Geophysical Fluid Dynamics Laboratory (GFDL), in a version which is maintained by the AWI.

2.1 The Atmosphere Model IFS

The IFS has been developed by the ECMWF jointly with Météo-France from 1993 on [7]. In the present work we used the model version “CY16R2”.

The governing equations of the IFS are the so-called “primitive equations” which describe the general circulation in the atmosphere. Prognostic variables are the horizontal velocity, the temperature, the humidity and the surface pressure. The land surface parametrization of the IFS contains a bulk model, which represents processes such as evapotranspiration, snow formation and melting, precipitation and runoff using four layers below the earth’s surface [8]. As boundary condition the IFS needs the sea surface temperature (SST) for climate runs.

The IFS applies a spectral technique for the solution of its differential equations of dynamics. Several computations, e.g. the radiation processes, are carried out in grid-point space. Also the coupling to the ocean model takes place in the grid-point space. We used the model in a spectral resolution with the triangular truncation T63, which corresponds to a horizontal mesh of $1.9^\circ \times 1.9^\circ$. To avoid a very high resolution at high latitudes, the IFS uses a so-called “reduced grid”. As schematically shown in figure 1, near the poles less grid points along one latitude are used compared to the low latitudes.

In the vertical, the model has 31 layers, and a time-step of 1 hour has been used. The version used in our experiments is parallelized using the message-passing-interface (MPI)-library.

IFS is the operational model for the numerical weather prediction of the ECMWF. It has also been used for several climate runs, e.g. in the AMIP [9]. The IFS has been coupled with some ocean models, e.g. by the ECMWF for its seasonal forecast [10][11]. With MOM it hasn’t been coupled so far.

2.2 The Ocean Model MOM

Based on works of K. Bryan, the MOM is developed by the Geophysical Fluid Dynamics Laboratory (GFDL, Princeton) since 1969 [12]. The version MOM 2




Fig. 1. Scheme of the “reduced grid”.

used by us has been released by the GFDL 1995 [13]. Our MOM-version has been further developed by AWI, in particular a sea-ice model has been included. MOM uses the primitive equations to describe the circulation in the oceans. It’s prognostic variables are the horizontal velocity, the temperature and the salinity. The sea-ice model simulates the formation and melting of ice, it’s drift and the radiative balance at it’s surface [14]. The MOM needs input of horizontal wind stress (τ_x and τ_y), the wind speed (V), the temperature (T), the dew point temperature (T_d), the total cloud cover (N), and the precipitation (P) at it’s surface as boundary conditions.

We used the MOM in a horizontal resolution of 194×92 grid points with 29 layers in the vertical. The time step has been 4 hours. The code is parallelized using the Cray-SHMEM-library.

MOM is widely used in the climate modeling community, and has been coupled to several atmosphere models [4][15].

3 The Coupling Method

We used the so-called “external coupling” approach, in which the respective models are kept almost unchanged [16]. In each of the models coupling interfaces are defined, in which the data is sent and received, respectively. The data transfer between the models is managed by an additional program, the so-called “coupling tool”. We wanted to use the message passing technique to exchange

the data between the models, in particular the MPI library to achieve an efficient and portable way of coupling.

3.1 The Adapted Coupling Tool: The NCAR CSM-Coupler




Fig. 2. Exchange of data between IFS and MOM

As coupling tool, the CSM-coupler of the National Center for Atmospheric Research (NCAR, Boulder) has been adapted.

In general, such a coupling tool receives the data from one model, performs computations, if required, interpolates the fields between the different grids, and sends it to the other model. In NCAR's use for it's Climate System Model (CSM), the CSM-coupler performs the data exchange between four models describing atmosphere, ocean, sea-ice and land surface, respectively [17]. As in our work the sea-ice model was coupled internally to the ocean model and we decided to keep the land surface parametrizations of the atmosphere model, the communication constructs for the land and the sea-ice in the coupler had to be eliminated. Further, at NCAR the coupler computes the fluxes between the models of NCAR's CSM. The computation of the input fluxes is a relevant part of a particular model. As the formulas used by the CSM-coupler are very different from those used by IFS and MOM, this parts of the coupler had to be removed.

The remaining tasks to be handled by the coupler are receiving data from one model, interpolating it between the grids, and send it to the other model, as schematically illustrated in figure 2. The atmosphere time step is longer than the ocean time step, and so the coupler averages the data from the atmosphere's time step to the ocean's time step. As the CSM-coupler needs regular grids for the interpolation routine and the interfaces are designed for the use at NCAR, the coupler had to be adapted according to the needs of our models.

However, the CSM-coupler had the advantage to use MPI. Other coupling tools like the in the climate modeling community well-known coupler "OASIS" of the Centre Européen de Recherche et de Formation en Calcul Scientifique (CERFACS, Toulouse) didn't support the use of the MPI-library at the time we started our work (OASIS does so from the version 3.0 on, released in July 2000) [18].

3.2 The New MPI-Library “Meta-MPI”

Supercomputer vendors supply efficient MPI-libraries only for their own hardware. The use of a public-domain implementation of MPI, such as MPICH, doesn’t provide an efficient library for massively parallel systems [2]. In the framework of the GTWB-project therefore a new MPI-implementation, the “Meta-MPI”, has been developed. Meta-MPI uses for the communication between several programs router-nodes. If one program sends data to another one, the data is sent by the router node of this program to the router node of the other one. So for the use of n programs $n \times (n + 1)$ router-nodes are necessary. The configuration of Meta-MPI is done by an input-file, comparable with the “hostfile” used by MPICH. The configuration can be done using a graphical user interface. Meta-MPI has an efficiency comparable to the hardware vendor’s MPI-implementation [19].

3.3 Interfaces in the Models

In both models modular interfaces for the coupling have been defined. The routines for sending and receiving data have been implemented similar to the writing of output and the reading of input by the models, respectively. As the coupling is managed by the serial coupler, the data is received by one processor and distributed to all the others. To send data to the coupler, it has to be gathered by one processor, which sends the data to the coupler. The conversion of different units and, in IFS, the calculation of current values of accumulated precipitation fields is carried out by the interface routines. In the IFS an additional interpolation between the “reduced grid” and the regular grid had to be introduced, and in the MOM the different grids of ocean and sea-ice had to be handled.

4 Results

To test the coupled model and our coupling technique, we carried out some test runs. In these runs, the IFS has been computed on the IBM SP2, the MOM and the coupler on the Cray T3E. The SP2 is the less powerful machine, so there all available nodes have been used. From the 34 nodes of the SP2, two were necessary for the communication managed by Meta-MPI, so 32 were left for the integration of the model. To achieve load-balance between the two models, we used eight nodes of the Cray T3E for the integration of the MOM. With this configuration the integration of the MOM on the T3E consumed slightly less time than the IFS on the SP2. The coupler ran on one processor of the T3E (see figure 3).

The coupled model has been integrated for eight months. For these integrations, the IFS has been initialized using analysis data of the ECMWF for January 1, 1997. For the MOM, a “spinup”-run in uncoupled mode has been carried out




Fig. 3. The models and the underlying hardware.

over 100 years integrated time to achieve a climatological equilibrium. In this spinup-run, the model has been forced using climatological boundary values for the wind forcing. For the initialization for the temperature- and salinity-fields climatological data has been used. Starting with the begin of January, the models have been coupled together.

4.1 The Results of the Coupled Model

The main goal of our work was the development the tools and the technique for the coupling of climate models, and the results of the coupled model shall only be reported briefly in a qualitative manner.

The two models have been coupled together for the first time, and a first goal was therefore to develop a stable running coupled model. This has successfully been reached.

In figure 4 the sea surface temperature is plotted, as it is seen by the IFS. The treatment of the land-/sea-mask suppresses the use of the SST provided by the ocean model in the IFS, if the IFS assumes any land below a mesh box. This leads to the particular land-distribution shown in figure 4.

Note that the temperature over sea ice is too cold over the arctic. The reason for this is the special representation of the ice temperature in the sea ice model, which makes it difficult to get the real ice surface temperature. Both facts mentioned require further development for climate scenario runs. Figure 5 shows the temperature as diagnosed in 2m above the surface in the atmosphere, averaged over July. In general, it agrees with the expected temperature distribution.

4.2 Technical Results

Every time step, eight physical fields in the grid-point space had to be sent – seven from the atmosphere model to the coupler, and one in the opposite direction. The horizontal grid consisted of 96 points in latitude and 192 points in longitude, so for a 8-Byte-representation of the data results a data volume of 1.3 MB. To send such a volume of data through the used network, a time in the range of less than 10ms were necessary. The computing time consumed for the integration of one time step of the atmosphere model has been on average 11s. During the test-runs of the model never any failure of the network has been occurred.




Fig. 4. Average of the sea surface temperature in July.

The comparison with test runs in uncoupled mode showed, that in our case it was even much faster to run the ocean model on a remote computer and send the boundary values via the network to the atmosphere model than to read it simply from the SP2's hard disk.

5 Summary and Conclusion

Using state-of-the-art general circulation models for atmosphere and ocean, we developed a new coupled climate model. The two models, the atmosphere model IFS of the ECMWF and the ocean model MOM in a version extended by an internally coupled sea-ice model by the AWI, have been coupled together for the first time. The coupled model runs stable and produces in general reliable results. For scenario climate runs further development would be necessary.

We coupled the models using an adapted version of NCAR's CSM-coupler. Because many changes had to be introduced to adapt this tool, probably the choice of a more general coupling library would be reasonable for future works. The MpCCI-library for example provides the facilities needed, and supports additionally a coupling in parallel, so that the data does not need to be gathered by one processor for sending, but processors of the two models working on the same domain can directly communicate together [20].

The models ran on two remote computers, one Cray T3E in Jülich, the other one on an IBM SP2 at the GMD St. Augustin. The data exchange took place on a




Fig. 5. Average of the temperature in 2m in july.

gigabit test-network supplied by the DFN, crossing the distance of about 100km. We used a new implemented message passing-library, “Meta-MPI”, which has the advantage of a high-performance connection of different supercomputers and a quite easy handling. Neither bandwidth of the network used nor it’s latency have been a constraint for our simulation. The data volume of 1.3MB has been sent in some ms, a factor of 1000 faster than the computation time. The network has been perfectly reliable. In comparison to former studies of distributed computation of climate models our work has been the first which does it in an efficient way. As networks of the quality used by us will become more and more available to every research center, we hope that our coupling technique can be used by scientists to couple their climate models together.

Projects concerning this topic are actually carried out by our partners of the AWI, and are planned by the GMD together with several climate research centers in Germany.

References

1. T. Eickermann, P. Wunderling, R. Niederberger and R. Völpe, Aufbruch ins Jahr 2000 – Start des Gigabit Testbed West, DFN-Mitteilungen 45 (1998) 13–15.
2. Eickermann, T., J. Henrichs, M. Resch, R. Stoy and R. Völpe, Metacomputing in gigabit environments: networks, tools, and applications, *Parallel Comp.* 24 (1998) 1847–1873.
3. F. Hommes and R. Niederberger, Gigabit Testbed West Technik, in: *DFN-Symposium*, DFN, Berlin, 1999.

4. C. Mechoso, C.-C. Ma, J. Ferrara and J. Spahr, Parallelization and distribution of a coupled atmosphere–ocean general circulation model, *Mon. Wea. Rev.* 121 (1993) 2026–2076.
5. C. Cassou, P. Noyret, E. Savault, O. Thual, L. Terray, D. Beaucourt and M. Imbard, Distributed ocean–atmosphere modeling and sensitivity to the coupling flux precision: the CATHODE project, *Mon. Wea. Rev.* 126 (1998) 1035–1053.
6. TEN34, Final report on advanced application monitoring, Deliverable D13.2, DANTE Ltd., Cambridge, 1998.
7. M. Déqué, C. Dreveton, A. Braun and D. Cariolle, The ARPEGE/IFS atmosphere model: A contribution to the french community climate modelling, *Clim. Dyn.* 10 (1994) 249–266.
8. ECMWF, Documentation of the ECMWF Integrated Forecast System (IFS), Technical Report, ECMWF, Reading, U.K., 1999
9. W. Gates, J. Byle, C. Covey, C. Dease, C. Doutriaux, R. Drach, M. Fiorino, P. Gleckler, J. Hnilo, S. Marlais, T. Phillips, G. Potter, B. Santer, K. Sperber, K. Taylor and D. Williams, An overview of the results of the atmospheric models intercomparison project (AMIP), Technical Report 45, PCMDI, Livermore, California, USA, 1998
10. T. Stockdale, M. Latif, G. Burgers and J.-O. Wolff, Some sensitivities of a coupled ocean–atmosphere GCM, *Tellus* 46A (1994) 367–380
11. T. Stockdale, Coupled ocean–atmosphere forecasts in the presence of climate drift, *Mon. Wea. Rev.* 125 (1997) 809–818
12. K. Bryan, A numerical method for the study of the circulation of the world ocean, *J. Comp. Phys.* 4 (1969), 347–376
13. R. Pacanowski, MOM 2 documentation, user's guide and reference manual, Technical Report 3, GFDL, Princeton, USA, 1995
14. H. Fischer, Vergleichende Untersuchungen eines optimierten dynamisch–thermodynamischen Meereismodells mit Beobachtungen im Weddelmeer, Technical Report 166, Alfred–Wegener–Institut for Marine and Polar Research, Bremerhaven, Germany, 1995
15. M. Fischer and A. Navarra, GIOTTO – A coupled atmosphere–ocean general circulation model: The tropical pacific, *Q.J.R. Met. Soc.* 126 (2000) 1991–2012
16. K. Cassirer, T. Horie and W. Joppich, On the use of coupling interfaces for the transfer of data from the global model GME to the local model LM of the German Weather Service, Technical Report, GMD–SCAI, St. Augustin, Germany, 1998
17. F. Bryan, B. Kauffman, W. Large and P. Gent, The NCAR CSM flux coupler, Technical Report, NCAR, Boulder, Colorado, USA, 1997
18. L. Terray, S. Valcke and A. Piacentini, OASIS 2.3 user's guide and reference manual, Technical Report TR/CGMC/99–37, CERFACS, Toulouse, France, 1999
19. J. Henrichs, Meta–MPI, Technical Report, Pallas GmbH, Brühl, Germany, 1998
20. R. Ahrem, M. G. Hackenberg, P. Post, R. Redler and J. Roggenbuck, Specification of MpCCI Version 1.0, GMD–SCAI, St. Augustin, 2000