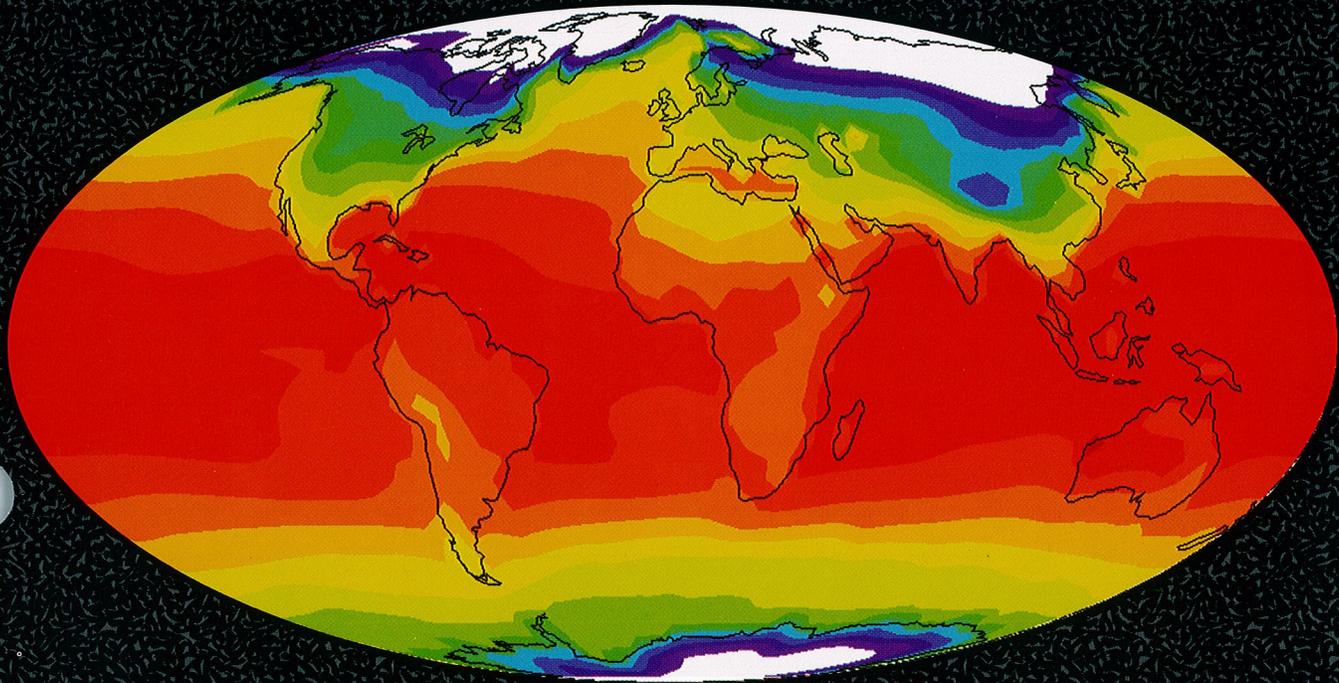




BUFFER

National Energy Research Supercomputer Center



- Climate System Modeling on High-Performance Computing Systems, with an Application to Modeling the Atmosphere
- An Introduction to Cyberspace
- Network PVM

1995 Buffer Index

We are developing an advanced generation of parallel climate models together with a computational framework that will provide a comprehensive climate systems modeling capability on high-performance computing systems.

Climate System Modeling on High-Performance Computing Systems, with an Application to Modeling the Atmosphere

Arthur A. Mirin and Michael F. Wehner*
Lawrence Livermore National Laboratory

Understanding the response of the Earth's climate system to both natural and anthropogenic (man-made) stimuli, together with understanding the natural variability of the climate system, is part of the science base needed for rational formulation of energy policy. Developing this knowledge requires a combination of field observations and numerical modeling as well as theoretical advances. A comprehensive climate system model must describe the behavior of the atmospheric, oceanic, cryospheric, and biospheric systems and the interactions between them.

The principal tools used to simulate the behavior of the global atmosphere and oceans are known as general circulation models. Large numbers of cells or degrees of freedom are needed to adequately resolve the circulation on the global scale. The computational resources required to provide significant predictive value with this class of fully coupled models in decadal to century time scales is well beyond current capability.

A Comprehensive Climate System Model

The Climate Systems Modeling Group of the Atmospheric Sciences Division at Lawrence Livermore National Laboratory (LLNL) is developing a comprehensive climate system model specifically designed to exploit distributed-memory, massively

parallel processing computers such as the CRAY T3D.¹ This model includes submodels of atmospheric general circulation, oceanic general circulation, sea-ice dynamics and thermodynamics, atmospheric chemistry and transport, ocean biogeochemistry, and soil and plant biogeochemistry. A coupling framework that allows for concurrent execution of multiple models, each with their own domain decomposition and grid structure, is also being developed.²

Parallelism within each model is achieved by partitioning the global spatial domain into a collection of subdomains and time-advancing the various subdomain solutions concurrently. Processors are assigned to subdomains in a deterministic manner, and variables local to a given subdomain are stored on the memory of the processor assigned to that subdomain. Data are transmitted between subdomains in the form of messages. The independent spatial variables are latitude, longitude, and a vertical coordinate. A two-dimensional latitude/longitude domain decomposition is implemented whereby each subdomain consists of a number of contiguous columns having full vertical extent appropriate to that package.

Process-level concurrence is attained by partitioning the machine among select physical packages. In some cases it is advantageous for submodels to execute concurrently, and in other cases it is more efficient for the same processors to execute a set of submodels. For example, the ocean general circulation, sea-ice dynamics, and ocean biogeochemistry are typically advanced using the same domain decomposition on the same set of processing elements, whereas the atmospheric general

* Other scientists contributing to the computational framework and atmospheric modeling efforts include J. H. Bolstad, B. Chan, U. E. Creach, W. P. Dannevik, P. B. Duffy, P. G. Eltgroth, C. M. Matarazzo, D. E. Shumaker, and H. Wang.



circulation is typically advanced on a separate set of processors. Package execution is based on the data flow concept; that is, a package is time-advanced until it is due to receive data from a different package.

The coupling aspects are nontrivial from both the physical and computational standpoints. In addition to identifying the data to be exchanged, one must assure conservation of key physical quantities. Also, one must take into account the different mesh sizes and domain decompositions of the various packages. We have addressed some of these issues by creating separate coupling modules, in particular for the ocean/atmosphere interface.

Our codes are transportable, yet also take advantage of particular features of each architecture. The use of preprocessing and macro constructs allows the maintenance of a single source code. We have encountered two issues that affect portability—dynamic memory management and interprocessor communication. These issues are addressed by the MICA package,³ which invokes standard UNIX preprocessors (CPP and M4). Conditional compilation also allows the inclusion of optimization constructs particular to the given platforms. Our codes have run on a number of architectures, including the CRAY T3D, Meiko CS-2, Intel Paragon, TMC CM-5, IBM SP2, CRAY Y-MP C90, BBN TC-2000, and Sun, IBM, SGI, and DEC clusters (using the PVM, MPI, and P4 message-passing libraries).

Efficient Use of the CRAY T3D

The CRAY T3D has been our primary compute engine over the past year and a half. We have significantly improved the throughput of our models by making use of direct access to remote memory (the SHMEM communications library). For point-to-point message passing there are two alternatives provided under SHMEM. One can either “get” data from a remote processor, or one can “put” data into the memory of a remote processor. When executing the latter procedure, one must make sure that the remote

processor cache is invalidated before the remote processor uses the data. Otherwise, the values previously cached as opposed to the values more recently put into its local memory could be used. We have chosen to utilize “puts” because the associated performance is typically about twice as fast as with “gets.” We have encapsulated this logic as an option in MICA so that from the user’s standpoint the remote memory access appears as message passing. Our encapsulation has minuscule overhead, and the “message-passing” calls can be placed virtually anywhere in the code. Invocation of this procedure has allowed attainment of optimal T3D communications performance with minimal source code changes.

Another optimization feature on the CRAY T3D is the reduced arithmetic library of transcendental functions, called “benchlib.” This library, provided by Cray Research, Inc., can significantly improve performance on the T3D at the risk of reduced accuracy. We have realized significant speed increases at essentially no loss in accuracy, particularly for calculations in which “rtor” (a real number to a real power) is intensive.

Modeling of the Atmosphere

Our atmospheric general circulation model (AGCM) was developed from the UCLA model of Arakawa and coworkers.⁴ The parallelization was carried out at LLNL.⁵ The AGCM has two major components—hydrodynamics and column physics. The hydrodynamics portion uses an explicit finite-difference algorithm in which key physical quantities are conserved. The “hydrostatic approximation” reduces the problem from fully three-dimensional to a set of coupled two-dimensional equations closely related to the shallow water equations. The finite difference mesh is staggered in both the vertical and horizontal directions (the Arakawa “C-grid”). Because of the mesh singularities at the poles, an unrealistically small time-step is dictated for



numerical stability. Rather than use such a small time-step, a more efficient procedure, developed by Arakawa, is to use a larger time-step and filter out the unwanted numerical modes in the polar regions. The basic hydrodynamics algorithm requires just nearest neighbor communications, whereas the filtering requires nonlocal communications along each latitude line.

The other major component of the AGCM, the column physics, consists of a collection of physical processes that operate along vertical columns (hence the name). Many such processes are included, one of the most important being radiation. Because the domain decomposition is only in latitude and longitude and not vertical, each column is contained

within the memory of a given processor. Hence, the column physics parallelization is trivial.

The AGCM has been run on a number of parallel platforms, two of the most recent being the CRAY T3D and the IBM SP2. A performance comparison is shown in Figure 1. On a per-processor basis, the SP2 is faster than the T3D, as expected, but the T3D demonstrates superior parallel scaling as seen by comparing the concavities of the two curves. Using 256 processors on the CRAY T3D, the AGCM operates over 11 times faster compared to running on a single processor of the CRAY Y-MP C90. We see that our performance is roughly comparable to that of an efficiently multitasked code on the C90.

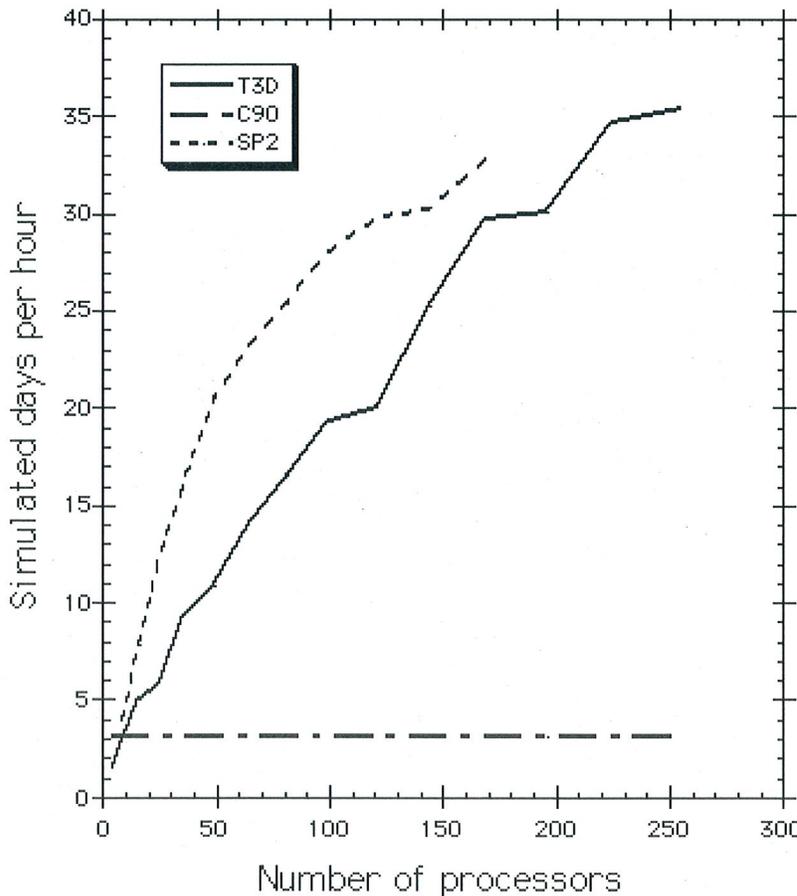


Figure 1. Performance of the Atmospheric General Circulation Model on the CRAY T3D, IBM SP2, and CRAY Y-MP C90, in units of simulated days per machine hour, on a 2-deg (latitude) by 2.5-deg (longitude) by 9-level mesh. The performance on the massively parallel processors is roughly comparable to that of an efficiently multitasked code on the C90.



A standard means of validating models of atmospheric general circulation is provided by the Atmospheric Model Intercomparison Project (AMIP). The test problem is defined by simulation of the 10-year period from 1979 to 1989, with prescribed sea surface temperatures. Over thirty AGCMs worldwide have been compared using this test problem.

Our acquired parallel processing technology has enabled us to execute an ensemble of twenty such realizations and to calculate the model variability of the atmosphere. This provides an estimate as to the natural variability of the atmosphere, without which we could not assess human-induced climate change. We believe this ensemble is the largest yet performed by any one group. Figure 2 (also on the cover) shows the mean surface temperature during the northern hemisphere winter months. Figure 3 shows the average standard deviation of the ensemble mean surface temperature. Note that the highest degree of

inter-realization variability occurs over regions of sea ice and snow-covered land, which is suggestive of a radiative feedback mechanism. Both Figures 2 and 3 were created using the Visualization and Computation System (VCS) developed by the Program for Climate Model Diagnosis and Intercomparison (PCMDI) at LLNL.⁶

Conclusions

This work provides one important example of massively parallel processors having come of age for grand challenge scientific calculations. Our success at simulating the atmosphere and the oceans, along with recent progress in the area of coupled ocean/atmosphere calculations, suggests that comprehensive climate simulations in the teraflops range should be feasible in the coming years.

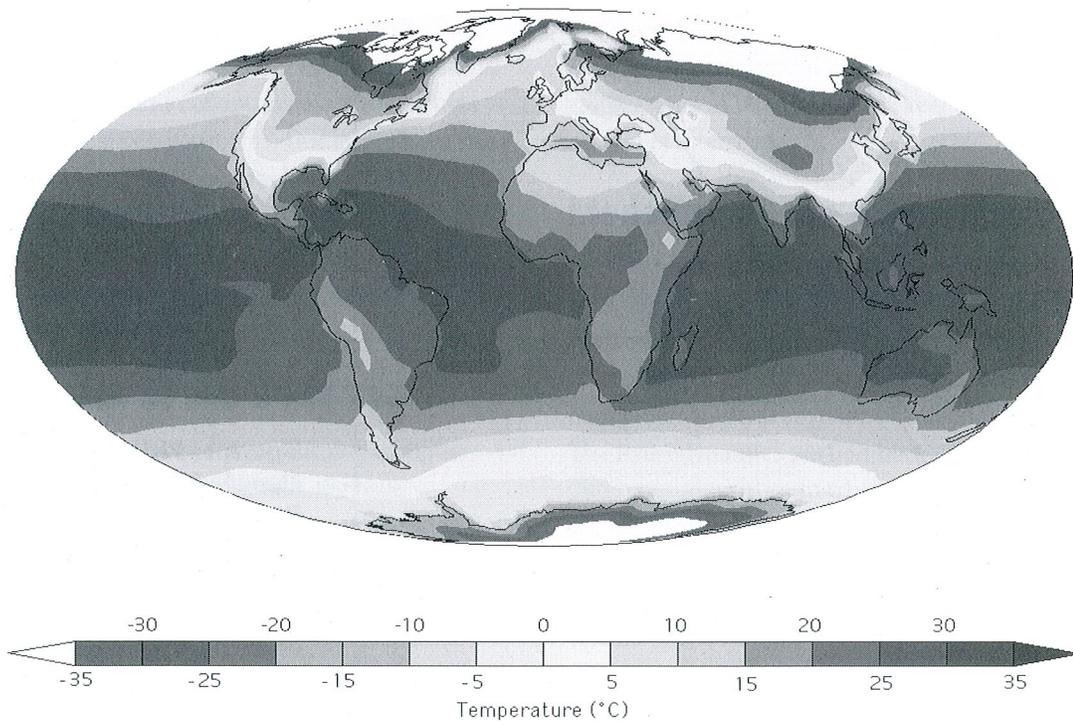
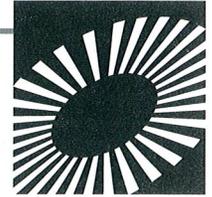


Figure 2. Mean surface temperature during the northern hemisphere winter months.



References

1. A. A. Mirin, et al., "Climate System Modeling Using a Domain and Task Decomposition Message-Passing Approach," *Computer Physics Communications* 84 (1994): 278–296.
2. M. F. Wehner, et al., "Parallel Coupled Oceanic-Atmospheric General Circulation Model," *Proc. Sixth Workshop on Use of Parallel Computers in Meteorology*, ECMWF, Reading, England (1994): 512–522.
3. Jonathan Brown and Art Mirin, "MICA, a Facility to Achieve Portability for Message Passing and Dynamic Memory Management in Fortran," *Buffer* 18, no. 2 (February 1994): 12–18.
4. A. Arakawa and V. Lamb, "Computational Design of the Basic Dynamical Processes of the UCLA General Circulation Model," *Methods in Computational Physics* 17 (1977): 173–265.
5. M. F. Wehner, et al., "Performance of a Distributed Memory Finite Difference Atmospheric General Circulation Model," *Parallel Computing* 21 (1995): 1655–1675.
6. VCS is computer software for the selection, manipulation, and display of climate data (see the December 1995 *Buffer*, p. 5). VCS is available free of charge to nonprofit institutions. For more information about VCS, access URL <http://www-pcmdi.llnl.gov> or contact Dean N. Williams (williams13@llnl.gov).

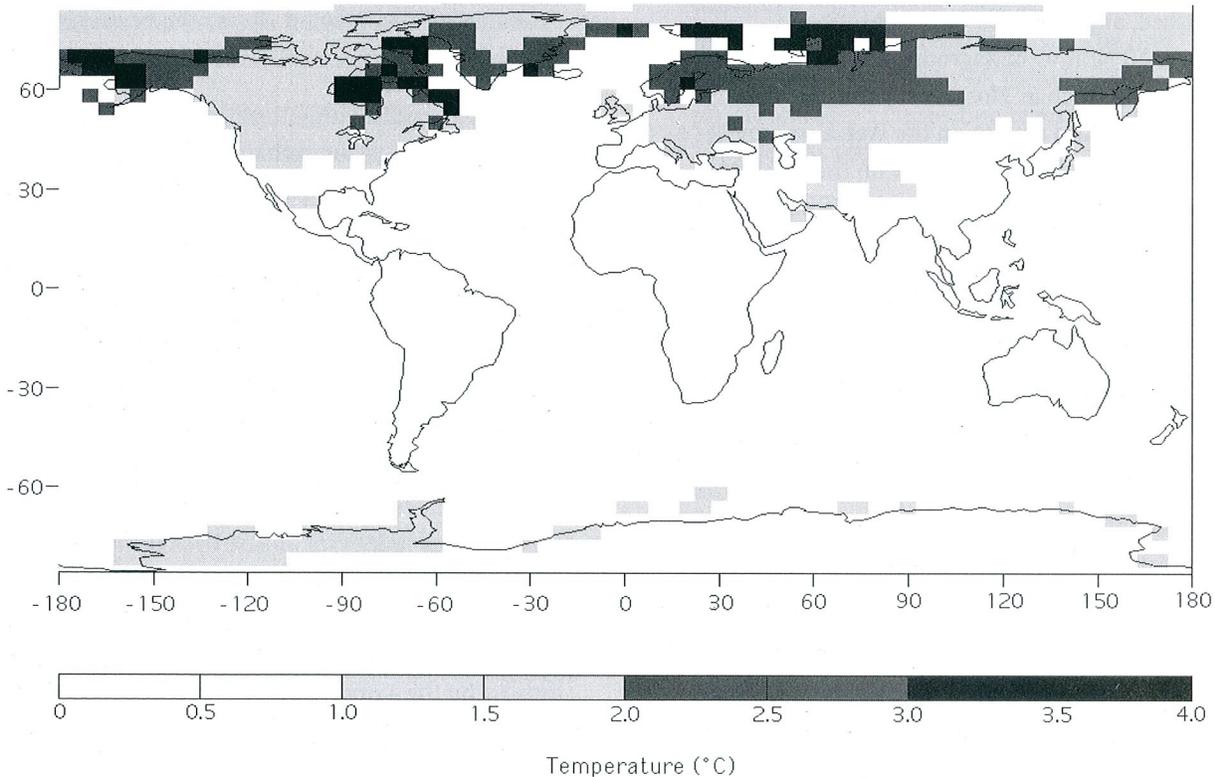


Figure 3. Average standard deviation of the ensemble mean surface temperature during the northern hemisphere winter months. The variation averages 0.4 deg and is largest over regions of sea ice and snow-covered land.