




Climate, the Ocean, and Parallel Computing

Robert C. Malone,
Richard D. Smith,
and John K. Dukowicz

The Laboratory's most powerful Connection Machine, a CM-5

The climate of the earth is controlled by an interplay among many competing physical processes operating in the atmosphere and the ocean and on land. Of the many questions facing today's climatologists, two seem particularly urgent. Is the balance among those processes being affected by human activities? And if so, how large are the resulting climatic changes relative to natural climatic variations? The high priority accorded internationally to answering those questions is impelled in part by the growing list of troubling environmental problems such as greenhouse warming, ozone depletion, pollution of the atmosphere and ocean, and tropical deforestation.

Past observations of the atmosphere and the ocean have contributed greatly to our knowledge of the climate system but still constitute only a short and incomplete baseline of data with which future changes can be compared. More extensive observations are in progress or being planned: NASA's Earth Ob-

serving System satellite will produce the most comprehensive picture to date of the present state of the earth's climate; the Tropical-Ocean Global-Atmosphere Program is investigating the impact of El Niño and the Southern Oscillation on weather patterns in mid-latitude regions; and the decade-long international World Ocean Circulation Experiment will probe to great depths the circulation in all the major ocean basins.

Such projects are an essential aspect of climate studies, but a theoretical framework is also needed as a basis for interpreting the accumulated data. The computer models known as general-circulation models (GCMs) provide such a basis by simulating the temporal evolution of the atmosphere or the ocean in three dimensions (latitude, longitude, and altitude or depth). In addition, GCMs are indispensable tools for investigating parts of the climate system, such as the deep ocean, that are very difficult to observe and for estimating the effects of natural and human-in-

duced environmental changes on climate. Three-dimensional GCMs were first developed in the 1960s. Their fidelity has since been greatly increased but is still limited by a shortage of data for validating the models and by the capabilities of the computers on which they are implemented.

Atmospheric and oceanic GCMs each contain mathematical representations of the dominant relevant physical processes. Included in an atmospheric GCM are transport of heat and moisture by the winds; exchange of momentum, moisture, and heat between the atmosphere and oceanic and terrestrial surfaces; condensation of moisture into clouds and precipitation; and absorption and scattering of incident sunlight and emission and absorption of infrared radiation by clouds, atmospheric gases, and oceanic and terrestrial surfaces. Also included are factors that affect those processes. For example, sea ice, snow, and vegetation affect energy exchange through their influence on the fraction

of incident sunlight absorbed by the ocean and land masses, and the earth's rotation strongly influences the circulation patterns of the winds. Included in an oceanic GCM are interaction of the ocean surface with the winds and with solar and infrared radiation; exchange of fresh water and heat with the atmosphere through evaporation and precipitation; convection driven by temperature and salinity variations; interaction with the edges of continents and islands and with the ocean bottom; and the effect of the earth's rotation on the ocean's circulation. Clearly a reasonably complete description of the climate system—one that couples the dominant physical processes operating in both the atmosphere and the ocean—is enormously complex, and therefore climate simulation taxes the capabilities of even the most powerful of today's supercomputers.

Development of detailed and realistic atmospheric GCMs has been spurred by their use in weather prediction and by the extensive array of satellite- and ground-based equipment put in place over the last few decades to observe atmospheric conditions. Improvements in atmospheric GCMs for weather prediction are directly applicable to GCMs for climate prediction (and vice versa) because the physical processes involved in both are the same. However, prediction of climate requires simulations extending over much longer time intervals (decades to centuries) than does prediction of weather (days to weeks). Therefore the ocean, which varies much more slowly than the atmosphere, can be held fixed in a weather model but must be treated as a dynamical component in a climate model.

Climate can be thought of as the statistical aspects of weather averaged over a period of many years. For example, a weather forecast might tell us whether rain is likely in Los Alamos

during the next several days, whereas a climate forecast might tell us whether the springtime precipitation in the midwestern United States, averaged over a decade, will increase or decrease and by how much compared with the present-day average. Although an atmospheric GCM can address both of those questions, the model must be used differently in each case. An atmospheric GCM calculates the temporal evolution of various atmospheric variables (such as temperature, wind velocity, and humidity) at a number of regularly spaced grid points. When an atmospheric GCM is used for weather prediction, the number of grid points must be large (that is, the grid points must be closely spaced, typically less than a hundred kilometers apart horizontally and a kilometer or less apart vertically) to achieve the most accurate prediction possible for a region the size of an average American state. Achieving such fine spatial resolution is very costly in terms of computer time, but a weather prediction need extend only a short time, say a week, into the future. When an atmospheric GCM is used for climate studies, the computing time is kept within reasonable bounds by sacrificing spatial resolution in favor of simulating time intervals of decades or longer.

The atmospheric component of a climate model is essentially a weather-prediction model with a coarse horizontal spatial resolution, typically several hundred kilometers. The atmospheric model generates a time sequence of simulated atmospheric states that can be analyzed statistically to obtain the time averages, variances, and covariances of the atmospheric variables used to describe climate. However, as noted above, the atmosphere is strongly influenced by and interacts with other components of the climate system that evolve more slowly (primarily the

ocean but also soil moisture, vegetation, sea ice, and glaciers). One of the challenges of climate modeling is to develop and validate models that adequately represent the more slowly varying components, which can then be coupled to the atmospheric model to form a complete model of the interactive climate system. Another challenge is to determine what spatial scales must be resolved to realistically model the long-term dynamics of climate.

Oceanic GCMs calculate the temporal evolution of oceanic variables on a three-dimensional array of grid points spanning the global ocean domain. Validation of ocean models has been hampered by the scarcity of oceanic data; that situation will be greatly improved by the data on temperature, salinity, and currents to be gathered by the World Ocean Circulation Experiment. (Primarily because they are relatively easy to observe, much is known about the ocean's surface currents. But currents at lower depths are less well explored. Figure 1 shows the major surface currents of the ocean.) Another difficulty of ocean modeling is that the dynamics of the ocean involves longer time scales and smaller spatial scales than does the dynamics of the atmosphere. The deep ocean takes far longer (decades to centuries) to react to external changes than does the atmosphere (months), and oceanic eddies are much smaller horizontally (less than 100 kilometers) than are atmospheric eddies (1000 kilometers or larger). Atmospheric eddies are known to play an important role in the transport of heat and momentum; oceanic eddies are believed to play a similarly important role, although the evidence for that is less conclusive. Therefore ocean simulations not only must extend over long time intervals but also should probably be finely resolved spatially. A question that remains to be answered is whether eddy-resolving




Figure 1. Major Oceanic Surface Currents

Prominent features in this map of the major oceanic surface currents include the subtropical gyres centered on 30° latitude in each of the major ocean basins. The earth's rotation and the change in wind direction with latitude (from the east in the tropics and from the west at mid latitudes) cause the circulation of the gyres to be clockwise in the Northern Hemisphere and counterclockwise in the Southern Hemisphere. The well-known Gulf Stream in the Atlantic and its counterpart in the Pacific, the Kuroshio Current, are strong currents that carry heat northward from the tropics. Both currents are evident in the model simulations displayed in Figures 3 and 4. The Antarctic Circumpolar Current flows around Antarctic in a band of ocean centered on latitude 60°S that is uninterrupted by continents. The Antarctic Circumpolar Current can be clearly seen in Figure 4. (Adapted with permission from a figure in *Principles of Ocean Physics* by John R. Apel, Academic Press, 1987.)

ocean models are necessary to understand climate.

That the ocean is a major component of the climate system is well documented. For example, as shown in Figure 2, the amount of heat transported from the tropics to the polar regions by the ocean is comparable to the amount transported by the atmosphere. In addition, because the ocean, particularly the deep ocean, has such a tremendous heat capacity, it acts as a “thermal flywheel” for the climate system by moderating

changes occurring in the atmosphere. For example, by sequestering heat trapped by greenhouse gases such as carbon dioxide and methane, the ocean may be delaying the onset of global warming due to production of those gases by human activities. The ocean also acts as a reservoir for carbon dioxide. It is estimated that the ocean presently holds fifty times as much carbon dioxide as the atmosphere and takes up half of the carbon dioxide released into the atmosphere each year. The ocean's buffering of carbon diox-

ide plays a central role in the carbon cycle of the earth.

Atmospheric and oceanic GCMs must be coupled to model interactions between the atmosphere and the ocean. Although ocean conditions can be held fixed for short-range weather prediction, for seasonal forecasting (three to six months in advance) and longer-term variations such as El Niño and the Southern Oscillation, atmosphere-ocean interactions must be modeled by treating the ocean as a dynamical entity. El Niño is the name given to the

dramatic warming of surface waters of the eastern tropical Pacific Ocean that occurs every three to seven years. Such warmings are now understood to be closely linked physically to episodic shifts, called the Southern Oscillation, in atmospheric circulation linking the Indian and eastern Pacific Ocean regions. The combined El Niño–Southern Oscillation phenomenon is an unforced free oscillation of the atmosphere-ocean system. The interactions between ocean and atmosphere that are responsible for an El Niño event are not fully understood, nor are the precursor conditions that initiate an event. El Niño events are of great interest because they not only affect weather in the tropical Pacific region but also seem to influence atmospheric conditions beyond the tropics in significant and potentially predictable ways.

Paleoclimatic data and computer simulations both suggest that shifts in ocean-circulation patterns are associated with changes in climate. At present a global-scale ocean-circulation pattern known as the conveyor belt warms the climate of northern Europe by carrying warm water into the North Atlantic Ocean via the Gulf Stream. That circulation pattern is driven by thermohaline effects (effects related to oceanic temperature and salinity gradients). The warmth of the water being carried into the North Atlantic Ocean enhances its evaporation, which increases the water's saltiness and hence its density. It cools and sinks, forming North Atlantic deep water (NADW). A current of NADW flows southward in the Atlantic Ocean, around Africa, and into the Indian and Pacific oceans, where it slowly mixes and rises toward the surface. Water near the surface flows back through the Indian Ocean, around Africa and northward in the Atlantic Ocean, thus completing the global conveyor belt. Evidence from ice cores




Figure 2. Heat Transport by the Atmosphere and the Ocean

Shown here are the annual mean values of the heat transported by the atmosphere ($Q_{\text{atmosphere}}$), by the ocean (Q_{ocean}), and by the atmosphere and the ocean together (Q_{total}), each as a function of latitude. The $Q_{\text{atmosphere}}$ values are measured values; Q_{ocean} cannot be measured and must be determined indirectly. To do so, Q_{total} is deduced by computing the total heat transport required to offset the imbalance at various latitudes between incoming solar radiation (dominant in the tropics) and outgoing infrared radiation (dominant in polar latitudes). Q_{ocean} is then obtained by subtracting $Q_{\text{atmosphere}}$ from Q_{total} . Positive and negative heat-transport values indicate northward transport and southward transport, respectively. (Adapted with permission from a figure in *Physics of Climate* by José P. Peixoto and Abraham H. Oort, American Institute of Physics, 1992.)

and deep-sea sediments suggest that the North Atlantic branch of the thermohaline circulation became active at the end of the last ice age and that fluctuations in its geographical extent and intensity are correlated with changes in atmospheric conditions. Thus the paleoclimatic data, as well as computer simulations with coupled atmospheric and oceanic GCMs, suggest that more than one mode of ocean circulation may exist. In the present-day mode NADW is formed in the North Atlantic and transported to other oceans by the conveyor belt. In another, glacial mode, production of NADW is either partially or completely shut off. An important question is how the northward extent of the conveyor belt might be affected by the warming induced by

greenhouse gases. Answering that question demands the best computer models that can be developed.

Using coupled atmospheric and oceanic GCMs to simulate climate is expensive because the simulations must extend over the long times required for the deep ocean to adjust. Many simpler models were developed in the past when computers were much less powerful, and research on simpler models continues today. Unfortunately, simpler models typically include ad hoc assumptions that are often difficult to justify or validate. Furthermore, they may fail to include the effects of feedback within or between components of the climate system that arise from the complexity or nonlinearity of the processes. Thus, although simpler models can be

very useful in preliminary investigations, their predictions must ultimately be compared with those of the most realistic models available. Therefore GCMs are still the tool of choice for simulating climate.

The advent of massively parallel computers such as the Connection Machine has provided a new and potentially much more powerful approach to computing. The Department of Energy established the CHAMMP (Computer Hardware, Advanced Mathematics, and Model Physics) Program to pursue development of a new generation of global climate models to be implemented on such computers. It is anticipated that the power and capacity of massively parallel computers will enable future models to include more realistic representations of a greater number of climate-system components. However, fully utilizing the potential of massively parallel computers may require that the mathematical representations of the processes included in climate models be extensively reformulated or that entirely new representations be developed.

A brief explanation of what a massively parallel computer is and how it differs from traditional computers may be helpful here. The words "massively parallel" refer to the fact that such a computer contains hundreds or thousands of processors, all performing their allocated share of the computational work more or less simultaneously. A local memory unit attached to each processor holds the data on which the processor is operating, and a high-speed network connects each processor to the others so that data can be exchanged among the processors whenever required. In contrast to a massively parallel computer, a traditional supercomputer contains only a small number of powerful processors (four to six-

teen), all of which share direct and equal access to a global memory bank through a very-high-speed network. Unfortunately, as the number of processors increases, the cost of such a "shared-memory" network becomes prohibitive. Therefore designers of massively-parallel computers are forced to distribute memory so that each processor has fast access to its local memory unit but slower access to all other memory units. Thus, the large number of cooperating processors and the distribution of memory and data across all of the processors are the key features that distinguish a massively parallel computer from traditional supercomputers.

Designing codes to run efficiently on massively parallel computers is more difficult than designing codes to run efficiently on traditional computers, but the possibility of using a very large number of processors is a strong incentive for the extra effort. Codes for massively parallel computers must use mathematical algorithms that divide the work as equally as possible among all the processors; the data must be organized so that most of the data needed by each processor is stored in its local memory; and when data must be exchanged between processors, the least data necessary must be transmitted as efficiently as possible. Developing computer codes with those characteristics is a challenging task.

The availability several years ago of a Connection Machine (a CM-2) in the Laboratory's Advanced Computing Laboratory motivated a long-term project to develop the first global ocean model for massively parallel computers. Albert Semtner of the Naval Postgraduate School in Monterey, California, and Robert Chervin of the National Center for Atmospheric Research in Boulder, Colorado, generously gave us a copy of their ocean model, which had

been designed to run on traditional Cray supercomputers. Semtner and Chervin have used their model to perform what are, to date, the highest-resolution simulations of global ocean circulation. The simulations were performed on a grid whose points are 0.5 degree apart in latitude and longitude and located at twenty vertical levels. Such a grid, hereafter referred to as the 0.5-degree grid, is sufficiently fine to begin resolving the oceanic eddies. The Semtner-Chervin model is a variant of the highly regarded and widely used Bryan-Cox-Semtner model, which was originally developed in 1969 by Kirk Bryan of the Geophysical Fluid Dynamics Laboratory/NOAA in Princeton, New Jersey.

Our approach was to first develop a version of the Semtner-Chervin model for use on the CM-2 without changing any of the basic numerical algorithms. Such an approach would allow us to verify that the resulting model functioned properly after being moved to a computer with a different architecture and would provide us with a performance baseline with which future improvements could be compared. However, moving the code from a computer with a few processors and shared memory to one with thousands of processors and distributed memory made certain changes obligatory. The data structures were completely reorganized to improve usage of the CM-2's processors, and the code was entirely rewritten in data-parallel FORTRAN to improve its organization and structure. (Data-parallel FORTRAN extends the old standard, FORTRAN 77, to include new features such as array syntax, which is a simpler and more compact way of expressing operations on data arrays. These features are part of the new standard, FORTRAN 90.) After making the necessary modifications to the code, we found that the performance of the model on the

2048 floating-point processors of the CM-2 was about the same as its performance on a 4-processor Cray X-MP. However, the portion of the code that calculates the vertically averaged (“barotropic”) velocity field did not function efficiently on the CM-2, so we were led to reformulate the equations for that portion. We also implemented more efficient algorithms for solving the reformulated equations on parallel computers.

Details of the changes to the model’s “barotropic solver” are presented in the sidebar “New Numerical Methods for Ocean Modeling on Parallel Computers.” Only their benefits to the physical realism and computational efficiency of the model will be discussed here. First, all eighty of the islands that can be resolved on the 0.5-degree grid can be included in the revised model at the same computational cost required to include the three “islands” used in the original model (Antarctica, Australia, and New Zealand). Second, unlike the original model, the revised model can be executed without smoothing the topography of the ocean bottom to remove steep depth gradients. And third, the revised model does not impose an artificial condition on the ocean surface (the “rigid-lid” condition) that was needed in the original model to eliminate surface waves.

As indicated above, our revisions to the model’s barotropic solver have also increased its efficiency. The revised barotropic solver is many times faster than the original (when each is executed on the 0.5-degree grid) even though it treats eighty islands and the original treats only three. (The difference in running times would be even greater, of course, if the original model treated more than three islands.) In its entirety the revised code runs about two and a half times faster on the CM-2 than did the original implementation. The 0.5-

degree simulation now runs on 512 floating-point processors (one-fourth) of the CM-200 (an upgraded Connection Machine that is more than 40 percent faster than the CM-2) at about the same speed as the original code runs on a 4-processor Cray X-MP.

We have begun using the new model in global ocean simulations. The simulations are initiated by setting the temperature and salinity to approximate climatological values and the velocity to zero everywhere. The real ocean is driven at the surface by the atmospheric winds and by exchange of heat and fresh water with the atmosphere. We would like to drive the model with accurately measured values of both wind velocities and fluxes of heat and fresh water. The velocity data are available but the flux data are not, at least not on a global basis. However, reasonably complete measurements of temperature and salinity have been made across the surface of the global ocean and, with less accuracy, even in the deep ocean. Therefore measured wind velocities are applied at the model’s upper surface, and the temperature and salinity values in the model’s topmost layer are continually “nudged” toward climatological values to compensate for the lack of data about heat and fresh-water fluxes. The nudging, which forces the solution toward observed values over a time scale of a month, causes the model’s predictions for slowly changing aspects of surface temperature and salinity to correspond closely to the climatological data, but more rapidly changing aspects are hardly affected by the nudging. Because the deep ocean evolves so slowly, a common practice is to nudge the model’s temperature and salinity fields there also toward observed values but to do so much more slowly than in the surface layer. That technique was used to produce the simulations presented here because it greatly re-

duces the computer time required to obtain fairly realistic conditions in the deep ocean.

Figures 3 and 4 show examples of output from the revised ocean model, as executed on the 0.5-degree grid and after 12 simulated years. Figure 3 shows the simulated temperature of the ocean surface. As expected, the large-scale temperature distribution closely resembles the climatological distribution imposed by nudging. But the velocity field and smaller, rapidly evolving features of the temperature distribution can be considered to be predicted by the model. Examples of such features that are evident in Figure 3 are two narrow, meandering western-boundary currents: the Gulf Stream off the eastern coast of the United States and the Kuroshio Current east of Japan. Figure 4 displays the magnitude of the vertically integrated velocity. Predictions of the model that correspond to known ocean features are pointed out in the captions to Figures 3 and 4.

Much more work must be done to thoroughly characterize the circulation patterns predicted by the model and to compare them quantitatively with observations. We hope to use the model to gain insight into aspects of the real ocean that are difficult to observe and, in the future, to couple it to an atmospheric model to study some of the atmosphere-ocean interaction phenomena described earlier.

Although our revisions to the ocean model were motivated primarily by the desire to improve its performance on massively parallel computers, most of the revisions are advantageous even when an ocean model is executed on a traditional supercomputer. Our free-surface representation, for example, is currently being implemented in versions of the Bryan-Cox-Semtner model used at the Naval Postgraduate School and the Geophysical Fluid Dynamics

New Numerical Methods for Ocean Modeling on Parallel Computers

The Bryan-Cox-Semtner ocean model is a three-dimensional model in Eulerian coordinates (latitude, longitude, and depth). The incompressible Navier-Stokes equations and equations for the transport of temperature and salinity, along with a turbulent eddy viscosity, are solved subject to the hydrostatic and Boussinesq approximations. The model includes a rigid-lid approximation (zero vertical velocity at the ocean surface) to eliminate fast surface waves; the presence of such waves would require use of a very short time step in numerical simulations and hence greatly increase the computational cost. The equations of motion are split into two parts: a set of two-dimensional "barotropic" equations describing the vertically averaged flow, and a set of three-dimensional "baroclinic" equations describing temperature, salinity, and deviation of the horizontal velocity components from the vertically averaged flow. (The vertical velocity component is determined from the constraint of mass conservation.) The barotropic equations contain the fast surface waves and separate them from the rest of the model.

The baroclinic equations are solved explicitly; that is, their solution involves a simple forward time-stepping scheme, which is well suited to parallel computing and presents no difficulty on the Connection Machine. On the other hand, the barotropic equa-

tions (two-dimensional sparse-matrix equations linking nearest-neighbor grid points) must be solved implicitly; that is, they must be solved at each time step by iteration. For historical reasons the barotropic equations in the Bryan-Cox-Semtner model are formulated in terms of a stream function. Such a formulation requires solving an additional equation for each island, an equation that links all points around the island. The extra equations create vectorization difficulties when the model is implemented on a Cray and serious communication difficulties when it is implemented on a Connection Machine because a summation around each island is required for every iteration of the implicit solver. Therefore all but the three largest islands had been deleted from the original model, even though eighty islands are resolvable at the horizontal resolution employed (0.5 degrees latitude and longitude). Even so the barotropic part of the code consumes about one-third of the total computing time when the model is executed on a Cray and about two-thirds of the total computing time when the model is executed on a Connection Machine.

The above considerations led us to focus our efforts on speeding up the barotropic part of the code. We developed and implemented two new numerical formulations of the barotropic equations, both of which involve a surface-pressure field rather

than a stream function. The surface-pressure formulations have several advantages over the stream-function formulation and are more efficient on both parallel and vector computers.

The first new formulation recasts the barotropic equations in terms of a surface-pressure field but retains the rigid-lid approximation. The surface pressure then represents the pressure that would have to be applied to the surface of the ocean to keep it flat (as if capped by a rigid lid). The barotropic equations must still be solved implicitly, but the boundary conditions are simpler and much easier to implement. In addition, islands then require no additional equations, and therefore any number of islands can be included in the grid at no extra computational cost. Furthermore, and perhaps more important, the surface-pressure, rigid-lid formulation, unlike the stream-function, rigid-lid formulation, exhibits no convergence problems due to steep gradients in the bottom topography. The matrix operator in the surface-pressure formulation is proportional to the depth field H , whereas the matrix operator in the stream-function formulation is proportional to $1/H$. As a result, the latter matrix operator is much more sensitive than the former to rapid variations in the depth of waters over the edges of continental shelves or submerged mountain ranges, where the depth may change from several thousand meters to a few tens of meters

within a few grid points. Because such a rapidly varying operator may prevent convergence to a solution, steep gradients were removed from the stream-function formulation by smoothing the depth field. The surface-pressure formulation, on the other hand, converges even in the presence of steep depth gradients. Smoothing of the depth field could significantly affect the accuracy of a numerical simulation of the interaction of a strong current with bottom topography. For example, the detailed course and dynamics of the Antarctic Circumpolar Current (the strongest ocean current in terms of total volume transport) is greatly influenced by its interaction with bottom topography.

As we worked with the surface-pressure, rigid-lid model, we noticed a problem in shallow isolated bays such as the Sea of Japan. In principle, we should have been able to infer the elevation of the ocean surface (relative to the mean elevation) from the predicted surface pressure. We found, however, that the surface heights so inferred were quite different from those expected due to inflow or outflow from the bays. Removing the rigid lid solved that problem, but of course it also brought back the undesirable and unneeded surface waves. We were able to overcome that new difficulty by treating the terms responsible for the surface waves implicitly, which artificially slows down the waves, whereas the rigid-lid approximation artificially speeds up the waves to infinite velocity. (Either departure from reality is acceptable: Climate modeling does not require an accurate representation of the waves because they have little effect on the ocean circulation.)

Those considerations led us next to abandon the rigid-lid approximation

in favor of a free-surface formulation. The surface pressure is then proportional to the mass of water above a reference level near the surface. The benefits of the surface-pressure, free-surface model are greater physical realism and faster convergence of the barotropic solver. In particular, the revised barotropic part of the code, including eighty islands, is many times faster than the original, including only three islands (when both are implemented on the 0.5-degree grid). In addition, the surface pressure is now a prognostic variable that may be compared to global satellite observations of surface elevation to validate the model, and satellite data may now be assimilated into the model to improve short-term prediction of near-surface ocean conditions.

None of our revisions, of course, changed the fact that the large matrix equation in the barotropic solver must be solved implicitly. We chose to use conjugate-gradient methods for that purpose because they are both effective and easily adapted to parallel computing. Conjugate-gradient methods are most effective when the matrix is symmetric. Unfortunately, the presence of Coriolis terms (terms associated with the rotation of the earth) in the barotropic equations makes the matrix nonsymmetric. By using an approximate factorization method to split off the Coriolis terms, we retained the accuracy of the time-discretization of the Coriolis terms and produced a symmetric matrix to which a standard conjugate-gradient method may be applied. We also developed a new preconditioning method for use on massively parallel computers that is very effective at accelerating the convergence of the conjugate-gradient solution. The method exploits the idea of a local approxi-

mate inverse to find a symmetric preconditioning matrix. Calculating the preconditioner is relatively expensive but need be done only once for a given computational grid.

Further Reading

Kirk Bryan. 1969. A numerical method for the study of the circulation of the world ocean. *Journal of Computational Physics* 4: 347–376.

John K. Dukowicz and Arkady S. Dvinsky. 1992. Approximate factorization as a high order splitting for the implicit incompressible flow equation. *Journal of Computational Physics* 102: 336–347.

J. K. Dukowicz, R. D. Smith, and R. C. Malone. 1993. A reformulation and implementation of the Bryan-Cox-Semtner ocean model on the Connection Machine. *Journal of Atmospheric and Oceanic Technology* 10: 195–208.

Albert J. Semtner, Jr. 1976. Finite-difference formulation of a world ocean model. In *Advanced Physical Oceanographic Numerical Modeling*, edited by J. J. O'Brien, 187. Dordrecht, The Netherlands: Reidel.

Albert J. Semtner, Jr., and Robert M. Chervin. 1988. A simulation of the global ocean circulation with resolved eddies. *Journal of Geophysical Research* 93: 15502–15522.

R. D. Smith, J. K. Dukowicz, and R. C. Malone. 1992. Parallel ocean general circulation modeling. *Physica D* 60: 38–61.







Figure 3. Simulated Oceanic Surface Temperature

The sea-surface temperatures shown here were simulated with our revised ocean model on the CM-200. The temperatures are color-coded from dark red for the hottest to dark blue for the coldest. Continents and islands are black. Meanders and eddies are evident in the warm water being transported poleward in the Gulf Stream along the east coast of North American and in the Kuroshio Current east of Japan. Another interesting feature is the progression of waves in the tropical Pacific Ocean; in movies we have made from the model output, those waves are seen to propagate westward. Similar westward-propagating waves have been observed in satellite measurements of sea-surface temperature. The spatial resolution of the computer model is 0.5° in latitude and longitude with 20 vertical levels; realistic ocean bottom topography is used.




Figure 4. Simulated Vertically Integrated Ocean Currents

Shown here is the magnitude of the vertically integrated horizontal velocity field simulated with our revised ocean model on the CM-200. The speeds are color-coded from dark red for the highest to dark blue for the lowest. Continents and islands are black. Intense flows are evident in the Antarctic Circumpolar Current, the Gulf Stream, and the Kuroshio Current. The influence of submerged topography on the flow is particularly evident east of Australia, south of Alaska, and at several locations along the path of the Antarctic Circumpolar Current.

Laboratory. And, in collaboration with scientists at the latter institution, we are developing a more comprehensive data-parallel version of the model that includes more options for physical parameterizations and is designed to run on a variety of computer architectures.

Further Reading

Wallace S. Broecker. 1987. Unpleasant surprises in the greenhouse? *Nature* 328: 123–126.

Wallace S. Broecker, Dorothy M. Peteet, and David Rind. 1986. Does the ocean-atmosphere system have more than one stable mode of operation? *Nature* 315: 21–26.

Christopher C. Charles and Richard G. Fairbanks. 1992. Evidence from Southern Ocean sediments for the effect of North Atlantic deep-water flux on climate. *Nature* 355: 416–419.

Arnold L. Gordon. 1986. Interocean exchange of thermocline water. *Journal of Geophysical Research* 91: 5037–5046.

Scott J. Lehman and Lloyd D. Keigwin. 1992. Sudden changes in North Atlantic circulation during the last deglaciation. *Nature* 356: 757–762.

S. Manabe and R. J. Stouffer. 1988. Two stable equilibria of a coupled ocean-atmosphere model. *Journal of Climate* 1: 841–866.

S. G. H. Philander. 1983. El Niño Southern Oscillation phenomena. *Nature* 302: 295–301.

Warren M. Washington and Claire L. Parkinson. 1986. *An Introduction to Three-Dimensional Climate Modeling*. Mill Valley, California: University Science Books.



Robert C. Malone (center) received his B.A. in physics from Washington University and his Ph.D. in theoretical physics from Cornell University. He has been a Staff Member at the Laboratory since 1973, carrying out research on laser-fusion and magnetic-confinement-fusion theory, atmosphere and climate modeling, and computer science. He has also fulfilled various managerial roles. He has been honored by the Laboratory with a Distinguished Performance Award and currently directs model development for the DOE's Computer Hardware, Advanced Mathematics, and Model Physics Program.

Richard D. Smith (right) received his B.A. in liberal arts from St. John's College and his Ph.D. in theoretical physics from the University of Maryland. He was a Postdoctoral Fellow at Lawrence Livermore National Laboratory before coming to Los Alamos in 1987. He became a Staff Member in 1990 and has since been honored by the Laboratory with a Distinguished Performance Award. His research interests include computational fluid dynamics, parallel computing, and climate modeling.

John K. Dukowicz (left) received his B.A.Sc., M.A.Sc., and Ph.D in aerophysics from the University of Toronto. After advancing to the position of Principal Aerophysicist at Cornell Aeronautical Laboratory, he transferred first to the Mitre Corporation as a Technical Staff Member and then to General Motors Research Laboratory as a Senior Research Engineer. He joined the Laboratory as a Staff Member in 1977 and has since been honored by the Laboratory with a Distinguished Performance Award and an appointment as a Laboratory Fellow.