from *SIAM News*, Volume 33, Number 1

Modeling the Global Climate System On High-Performance Computers

By Robert L. Higdon

The prospect of global warming has received great attention in the popular press in recent years. An increase in average global temperatures could bring rising sea levels, due to thermal expansion and the retreat of polar ice caps; global disruption of weather patterns, including more frequent hurricanes and stronger El Niño events caused by higher sea-surface temperatures in the tropics; and larger populations of disease-spreading organisms, due to increased heat and humidity in certain regions. All this is blamed on increasing concentrations of carbon dioxide and other "greenhouse gases" in the atmosphere, which in turn are blamed on deforestation and the burning of fossil fuels. At the same time, however, skeptics assert that global warming is a fiction and that there is no justification for expensive remediation procedures. What does science have to say about all this?

This issue was addressed by Albert Semtner in a talk titled "Realistic Climate Modeling on US-Made High-End Parallel Computers: Formulations, Machine Considerations, and Results," which was given in San Antonio in March 1999 during a joint session of the Fifth SIAM Conference on Mathematical and Computational Issues in the Geosciences and the Ninth SIAM Conference on Parallel Processing.

Semtner is a professor of oceanography at the Naval Postgraduate School in Monterey, California, and an affiliate of the National Center for Atmospheric Research in Boulder, Colorado. In the late 1980s he collaborated with Robert Chervin of NCAR to produce the first fine-resolution simulations of global ocean circulation on decadal time scales. This award-winning work has been cited widely, and it is an essential step toward the larger goal of simulating the global climate system.

Semtner is now a key figure in the development, testing, and application of the Parallel Climate Model (PCM), a major climate modeling project sponsored by the U.S. Department of Energy and supported in part by the National Science Foundation. This project involves researchers from NCAR, the Naval Postgraduate School, Los Alamos and Oak Ridge National Laboratories, and the University of Texas at Austin. The model includes representations of atmospheric circulation, oceanic circulation, sea-ice dynamics, and river runoff. The principal investigator on this project is Warren Washington of NCAR. In his talk in San Antonio, Semtner described some of the physical, mathematical, and computational aspects of the PCM, and he also gave some preliminary results obtained with this model.

Historical Record

Semtner began by mentioning existing observational data related to mean global temperatures. In 1997, Cavalieri et al. [1] reported that over the previous 20 years there had been a general downward trend in the area of the earth's surface covered by polar sea ice, with a decrease in Arctic sea ice outweighing a slight increase in Antarctic sea ice. (Some theories of climate change predict a decrease in Arctic sea ice because of atmospheric warming and an increase in Antarctic sea ice because of a slowdown in ocean circulation.)

In addition, Mann, Bradley, and Hughes [2] have developed a history of temperatures dating back to 1400. For the distant past, they used paleoclimatic data records from sources like tree rings and core samples from ice packs and ocean sediments. Their results show that hemispheric mean temperatures have oscillated in time, due to natural variability and to external forcing from such sources as variations in solar radiation, shading due to explosive volcanism, and recent increases in greenhouse gases. From 1400 to 1900, the mean temperatures in the Northern Hemisphere oscillated irregularly over a range of a few tenths of a degree Celsius, without displaying any long-term trends. Since 1900, however, the average for the Northern Hemisphere has shown a definite upward trend, with a net increase of roughly half a degree. In addition, the three warmest years since 1400 have occurred in the 1990s. After examining correlations between the temperature record and known data about the forcings mentioned earlier, Mann, Bradley, and Hughes concluded that greenhouse gases are now the dominant external forcing of the climate system. Skeptics on global warming, Semtner pointed out, are a distinct minority within the cognizant scientific communities.

Where do we go from here? It is obviously of great interest to try to predict the future evolution of the climate system, based on current conditions and practices, and to assess the potential impacts of remediation strategies. Since physical laboratory experiments are out of the question, computational experiments must play a key role. Various scenarios of future energy usage and carbon dioxide production, for example, could affect the earth's climate in different ways; experiments performed with a numerical model of the global climate system might be used to estimate the possible outcomes.

Ocean Modeling

In San Antonio, Semtner focused primarily on the role of the ocean in climate and on the numerical modeling of oceanic circulation. This reflects not only his personal experience, but also the profound role played by the ocean in the global climate. The ocean, he said, can be seen as the "flywheel" of the global climate system, as it serves to moderate short-term fluctuations caused by atmospheric activity; in particular, the ocean covers most of the earth's surface, and the heat capacity of its top two and a half meters is comparable to that of the entire atmosphere. Strong currents along the western boundaries of ocean basins, such as the Gulf Stream in the North Atlantic and the Kuroshio in the North Pacific, transport heat away from the tropics and thus moderate the temperature differences between the tropics and the polar regions. In addition, the famous El Niño–La Niña cycle is manifested by anomalies in ocean temperature along the equatorial Pacific.

Perhaps the first widely used numerical model of ocean circulation was that created in the late 1960s by Kirk Bryan and Michael Cox of the Geophysical Fluid Dynamics Laboratory in Princeton and subsequently developed further by Bryan, Cox, Semtner, and others. A genealogy of various versions of this model is given in Semtner's introduction to Bryan's original 1969 paper as reprinted in the 30th-anniversary edition of the *Journal of Computational Physics* [3].

The Bryan–Cox model is based on the Navier–Stokes equations of fluid dynamics, modified and scaled to account for the particular properties of oceanic flows. For large-scale flows, for example, the depth is small relative to the horizontal length scale, which implies that the flow is approximately hydrostatic. In addition, some of the motions in the ocean are rapidly moving external gravity waves, which are approximately independent of depth. For reasons of computational efficiency, these motions are split off into a two-dimensional, vertically averaged subsystem that is solved separately from the remaining three-dimensional system. The governing equations are discretized with finite differences. One of the recent variants of the Bryan–Cox model is the Parallel Ocean Program, which was developed at Los Alamos National Laboratory by Richard Smith, John Dukowicz, and Robert Malone. This version is especially adapted for massively parallel processing, and it is the ocean model being used in the PCM.

Numerical modeling of the time-dependent ocean circulation presents great computational demands, due to the space and time scales involved. The western boundary currents mentioned earlier, for example, have widths on the order of tens of kilometers, but they must be resolved properly because of their importance to climate. In addition, ocean currents are typically unstable and develop meanders, eddies, and turbulent patterns on scales of the same order as boundary current widths. Corresponding structures in the atmosphere are an order of magnitude larger. In addition, the ocean responds much more slowly than the atmosphere to external forcing, which typically implies long integration times for climate-scale simulations.

Semtner gave some specific examples of the spatial resolution required for ocean modeling. These examples were obtained from numerical simulations using the Parallel Ocean Program (POP) and the Parallel Ocean Climate Model (POCM); the latter, Semtner's version of the Bryan–Cox model, is adapted to parallel vector processors. In a POCM global simulation with 20 levels in the vertical direction and an average horizontal grid spacing of 1/4 degree of latitude and longitude, many of the features of the real ocean were reproduced. However, this simulation yielded incorrect locations of the positions where the Gulf Stream and Kuroshio separate from the boundaries of their respective basins and travel into the ocean interior. The Parallel Ocean Program produced similar results in a 1/6-degree global simulation. However, in a recent POP simulation of the North Atlantic with a horizontal resolution of 1/10 degree and 40 levels, the Gulf Stream followed the correct path and displayed realistic instabilities, as measured by the time variability of sea-surface height.

Another example is a simulation of sea-ice anomalies in the Antarctic. Variations in the sea-ice extent travel in a wave around Antarctica, with a period of about 10 years. This phenomenon was captured in a simulation that used a horizontal grid spacing of 1/4 degree. With this resolution, the simulated anomaly agreed closely with observations, whereas with lower resolution the anomaly did not move at the correct speed. The key was to give an adequate resolution of the Antarctic Circumpolar Current, a strong, eastward-moving current that encircles Antarctica and drives the sea-ice wave. Another numerical simulation with 1/4-degree grid spacing produced sea-level variations in the tropical eastern Pacific that correlated well with observations.

It is easy to imagine that this level of resolution, applied on a global scale for integrations over many years, would lead to extremely heavy computational demands. Table 1 shows Semtner's estimates of computation times for several parallel machines with various grid resolutions in global simulations. The estimated computation times for the PVP machines were obtained by scaling the performance of Semtner's Parallel Ocean Climate Model, and the times for the other two categories are based on the performance of the Los Alamos Parallel Ocean Program. The last entry in each row of the table is the estimated computation time for one year of simulated model time. The computational demands of such simulations—performed for, say, a few decades or centuries of model time—would clearly be great.

Machine/ No. of processors (year)	Sustained gigaflops	Horizontal grid (degrees)	No. of vertical levels	Machine hours per model year
PVP				
XMP/4 (1986)	0.5	1/2	20	70
YMP/8 (1989)	1.2	1/4	20	60
C90/16 (1992)	5	1/4	20	15
T-90/32 (1995)	15	1/4	40	10
SX-4/32 (1996)	25	1/8	40	50
SX-5/16×24 (1998)	1000	1/10	40	3
MPP				
CM5/1024 (1993)	10	1/6	40	30
T3E/1024 (1997)	30	1/8	40	25
VPP700/256 (1997)	150	1/10	40	10
DSM				
Origin2000/128×2 (1997)	10	1/8	40	75
Origin2000/128×12 (1999)	10	1/10	40	15

Table 1. *Semtner's estimates of computation times on various machines, at various resolutions, for simulations of global ocean circulation with recent parallel versions of the Bryan–Cox ocean model. PVP indicates parallel vector processor; MPP, massively parallel processor; and DSM, distributed-shared-memory machines. The first four machines in the PVP category are from Cray and the last two from NEC. The MPP machines are from Thinking Machines, Cray, and Fujitsu; those in the DSM category are Silicon Graphics Origin2000 machines. In each case, the average horizontal grid spacing is expressed in terms of degrees of latitude and longitude. The last column gives hours of computation time per simulated year. Climate studies will require simulations of the coupled atmosphere–ocean system over periods of hundreds of years.*

In the version of the Parallel Ocean

Program currently used in the PCM, the average horizontal grid resolution is 2/3 degree. The grid is not uniform in latitude and longitude, but instead has been deformed so that the north coordinate pole is on a land mass, thereby avoiding a coordinate singularity within the fluid domain. This deformed grid has higher concentrations of grid points in regions where greater resolution is needed, including the Gulf Stream and Kuroshio, the Arctic, and the area near the equator (for resolution of El Niño events). A recently developed version of the Parallel Ocean Program with an average resolution of 1/3 degree is projected for inclusion in the PCM.

Other Aspects of Climate Modeling

In the PCM, the dynamics of the atmosphere are represented by version 3 of the Community Climate Model (CCM3), a widely used general circulation model developed at NCAR over the past two decades. The model includes representations or parameterizations of many aspects of atmospheric physics, including convection and condensation; cloud coverage and the optical properties of clouds; solar radiation and the reflection and absorption of radiation; exchanges of heat, moisture, and momentum between the atmosphere and the surface; and vertical diffusion and the atmospheric boundary layer. The PCM employs a recently parallelized version of CCM3. The model uses a pseudospectral spatial discretization in the horizontal dimensions, with a resolution of 42 wavenumbers after truncation of higher wavenumbers. Vertical discretization is based on finite differences, with 18 vertical levels.

The PCM also includes components that describe river runoff and sea-ice dynamics. The atmospheric component of the climate model represents precipitation, some of which is deposited on the land. Some of this water must then make its way to the ocean. This process is modeled by a river transport model developed at the University of Texas at Austin. The model uses information about mass flux and river direction in order to transport water from grid cell to grid cell, and ultimately to the ocean component of the model. The dynamics of sea ice are currently represented by a model developed at the Naval Postgraduate School in which a rectangular grid is centered over each pole. An alternative model developed at Los Alamos may be incorporated into the PCM in the future for reasons of computational efficiency.

In the real climate system, quantities like momentum, heat, and freshwater are exchanged among the atmosphere, ocean, sea ice, and land. In the PCM, these exchanges are managed by a flux coupler developed at NCAR. Because the different components of the climate model use different spatial grids, the flux coupler includes a scheme for interpolating between different grids while enforcing numerical conservation of quantities that are conserved physically.

A remaining question concerns the uses to which a climate model can be put once it has been constructed. Because the coupled model of atmosphere, ocean, ice, and land describes the evolution of the system through time, we might think in terms of predicting the future from prescribed initial and boundary data. Problems arise here, however, because the physical system being modeled is inherently chaotic. Hydrodynamic instability in the atmosphere leads to the variations in weather that we all experience, and hydrodynamic instability in the ocean leads to the meanders, eddies, and turbulent behavior mentioned earlier.

In a numerical simulation, the instability makes the computed solution extremely sensitive to perturbations in the initial conditions. Compounding the problem are limitations on the accuracy and density of the available initial data, especially for the state of the ocean. Improvements in empirical data, coupled with techniques of data assimilation, could lead to a small degree of predictability of time-averaged conditions on seasonal to perhaps decadal time scales. For purposes of long-term climate modeling, however, it simply is not feasible to start with present conditions and integrate forward in time to predict the specific state of the earth's atmosphere and ocean in, for example, a specific month many decades in the future.

Instead, climate simulations regard the climate primarily as a system that is forced by solar radiation and boundary conditions. From an initial state based on climatological data, the system is integrated in time until it reaches a quasi-equilibrium state. This state of the model can then be used as a starting point for numerical experiments involving the global climate, which are interpreted in a statistical sense.

With the PCM as now used, the startup process begins with the integration of the atmospheric model alone for 10 model years. During this process, the atmospheric model is forced by solar radiation and by climatological values of sea-surface temperatures and sea-ice concentrations. The ocean and sea-ice models are then integrated for 100 model years. This integration is forced by atmospheric data obtained from the last five years of the atmospheric startup, used in repeating cycles.

A significant problem with the startup of the ocean model is the very slow response of the deep ocean to changes in surface forcing. To compensate for this in the model, the heat and salt capacities in the deep ocean are reduced to a tenth of their true physical values during startup, and during the 100-year simulation the behavior of the deep ocean mimics evolution over 1000 years. This acceleration device is then removed, and the models of the atmosphere, ocean, sea ice, and land are coupled and integrated together for an additional 50 model years. At this stage the coupled model is at an approximate equilibrium, as measured by net transports of quantities like heat, mass, and momentum. It is then possible to begin some experiments.

Experiments and Predictions

The PCM group has performed several experiments to date, with many more anticipated for the future. One such experiment is a control run in which the coupled system is integrated further in time without any change in the composition of the atmosphere. This control has been run for 300 model years and will be extended further. One result of this run is illustrated by the lowest curve in Figure 1, which shows the global average of sea-surface temperature (SST) as a function of time. For a 300-year period, the average SST fluctuates over a range of about 0.2° Celsius, with no discernible long-term trends. Another diagnostic obtained from this run is the mass of Arctic and Antarctic sea ice. These quantities vary annually and over periods of a few years, but no long-term trends are evident in the computed results.

The lack of trends is itself significant. In past climate modeling efforts by several groups worldwide, the computed results have tended to display "climate drift"; that is, the simulated climate conditions have drifted away from current conditions without any changes in physical forcing that would account for the drift. To counteract climate drift, various groups have employed "flux correction"—the introduction of artifical transports between the atmosphere and surface in order to force the model climate toward present-day conditions. As a result of higher resolutions and improvements in model physics, climate drift has become less of a problem in recent years. The present model uses no flux corrections, and the results to date indicate little or no climate drift.

The PCM group has begun to use its model to investigate the effect on global climate of increasing concentrations of carbon dioxide in the atmosphere. In the experiments performed to date, the amount of CO_2 in the atmosphere is increased by 1% per year, which is a good estimate of the rate currently found in the real atmosphere. In one experiment, this rate of increase is continued until the amount of CO_2 is twice its present value, which takes approximately 70 years. According to the model, the average global sea-surface temperature then rises by approximately 1° Celsius. This is illustrated by the curve in Figure 1 that levels off in the middle of the figure, at about 18.8°. In another experiment, the 1% annual increase in CO_2 is continued until the amount has quadrupled. In that case, the model predicts an increase in sea-surface temperature of a little more than 2° Celsius. This is illustrated in the figure by the curve that nearly reaches the top of the graph frame.

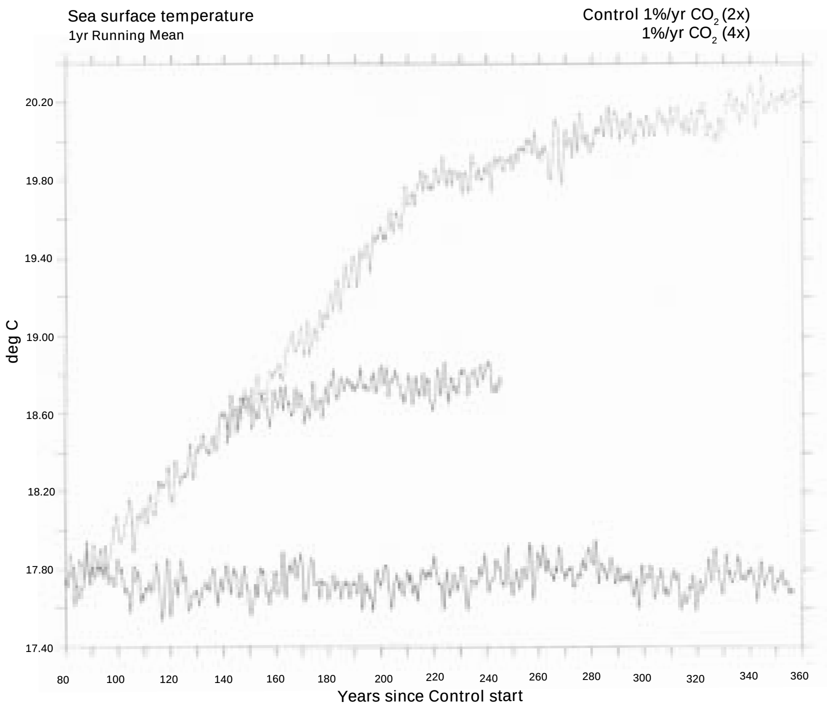

Figure 1. *Plots of globally averaged sea-surface temperatures obtained with the Parallel Climate Model. The lowest curve is obtained from a control run, and the others are based on increases in the amount of carbon dioxide in the atmosphere by 1% per year until it has been doubled or quadrupled relative to present values. The model also indicates that the average air temperature over land will increase by two to three times as much as will the average temperature of the sea surface. (Figure courtesy of the National Center for Atmospheric Research.)*

A global increase in sea-surface temperature is just one measure of change in global climate, and the model results obtained to date suggest other changes as well. For example, the model predicts that the temperature increase will be, on average, two to three times greater over land than over the ocean. One possible explanation for the difference is that the humidity over land will increase, and the increased amount of water vapor will itself provide an additional greenhouse effect. Because the air above the ocean is already closer to saturation, there is less room for change in ocean temperature through this mechanism.

The model also predicts that global warming will be greater in the higher latitudes than in the tropics. Because of warming, the amount of sea ice will decline; at the higher latitudes, the earth's surface will consequently reflect less sunlight than before. This change, in turn, will result in additional warming in those regions.

Another prediction from recent model runs is a weakening of the thermohaline circulation, a circulation in the global ocean that is generated by variations in buoyancy. In this circulation system, water moves northward in the Atlantic Ocean along the Gulf Stream, after which some of the water sinks because of decreasing temperatures. The water then spreads slowly across the global ocean basin and eventually returns to the surface at various locations. A warming of the northern latitudes would weaken the convec-tive events that drive this process.

More generally, according to Semtner, the modes of natural climate variability could change as CO_2-induced climate change takes place. This modification in natural variability could produce greater local extremes of temperature and precipitation, in addition to the overall increase in average background temperature. This disruption of recent climate patterns may be the most harmful effect of global warming.

The model runs obtained to date are by no means the end of the experiments planned for the PCM. In particular, one run of the model is just one realization of a physical process that is inherently statistical. Accordingly, Warren Wash-ington's NCAR group is conducting ensemble calculations, in which the initial states are varied within reasonable limits of uncertainty while the other aspects of the model are kept fixed. The resulting outputs will then suggest the range of possible outcomes for the scenario being investigated.

A set of additional control runs, for example, is needed for a better understanding of the natural variability of climate. Additional ensemble calculations must be performed to assess the outcomes of numerous scenarios involving CO_2 emissions, sulfate emissions, energy-use strategies, and the effects of biomass. Moreover, adding chemical and biological components to a climate model can increase its computational requirements many times over. All the calculations will involve centuries of model time per run, and the computations will combine for a total of tens of thousands of years of model time.

The proposed experiments give a sense of the enormous computational demands imposed by modern climate research. Semtner points out that all the components of the PCM are designed to be maximally parallel on evolving U.S. machines. Progress in climate modeling, he says, now involves a continuing adaptation to new computer architectures as they become available.

Further Information

An article of this nature cannot possibly describe the full scope of present-day efforts to model the atmosphere, ocean, and climate, nor can a verbal description do justice to the many intricate and colorful outputs obtained from numerical models of these systems. As a starting point for further explorations of these areas, readers are referred to the following (admittedly incomplete) list of Web sites:

The site for Semtner's group in Monterey is www.oc.nps.navy.mil/~braccio, and the NCAR site for the Parallel Climate Model is goldhill.cgd.ucar.edu/pcm. An alternative approach to ocean modeling is isopycnic coordinate modeling, in which the vertical coordinate is a quantity related to density. More information can be found at panoramix.rsmas.miami.edu/micom, the site for the Miami Isopycnic Coordinate Ocean Model. This site has links to the sites of other isopycnic modeling groups.

Another major institute for climate research is the Hadley Centre for Climate Prediction and Research in the UK. Its Web site, www.meto.govt.uk/sec5/sec5pg1.html, includes the Centre's own projections of global warming, based on computer models. The site also includes plots of historical data that show a distinct upward trend in global average temperatures over the past century.

References

[1] D.J. Cavalieri, P. Gloersen, C.L. Parkinson, J.C. Comiso, and H.J. Zwally, *Observed hemispheric asymmetry in global sea ice changes*, Science, 278 (1997), 1104–1106.

[2] M.E. Mann, R.S. Bradley, and M.K. Hughes, *Global-scale temperature patterns and climate forcing over the past six centuries*, Nature, 392 (1998), 779–787.

[3] A.J. Semtner, *Introduction to "A numerical method for the study of the world ocean,"* J. Comput. Phys., 135 (1997), 149–153.

Robert Higdon, a member of the program committee for the Fifth SIAM Conference on Mathematical and Computational Issues in the Geosciences, is a professor of mathematics at Oregon State University.