CRAY CHANNELS

VOLUME 17 NUMBER 2 • A CRAY RESEARCH, INC. PUBLICATION

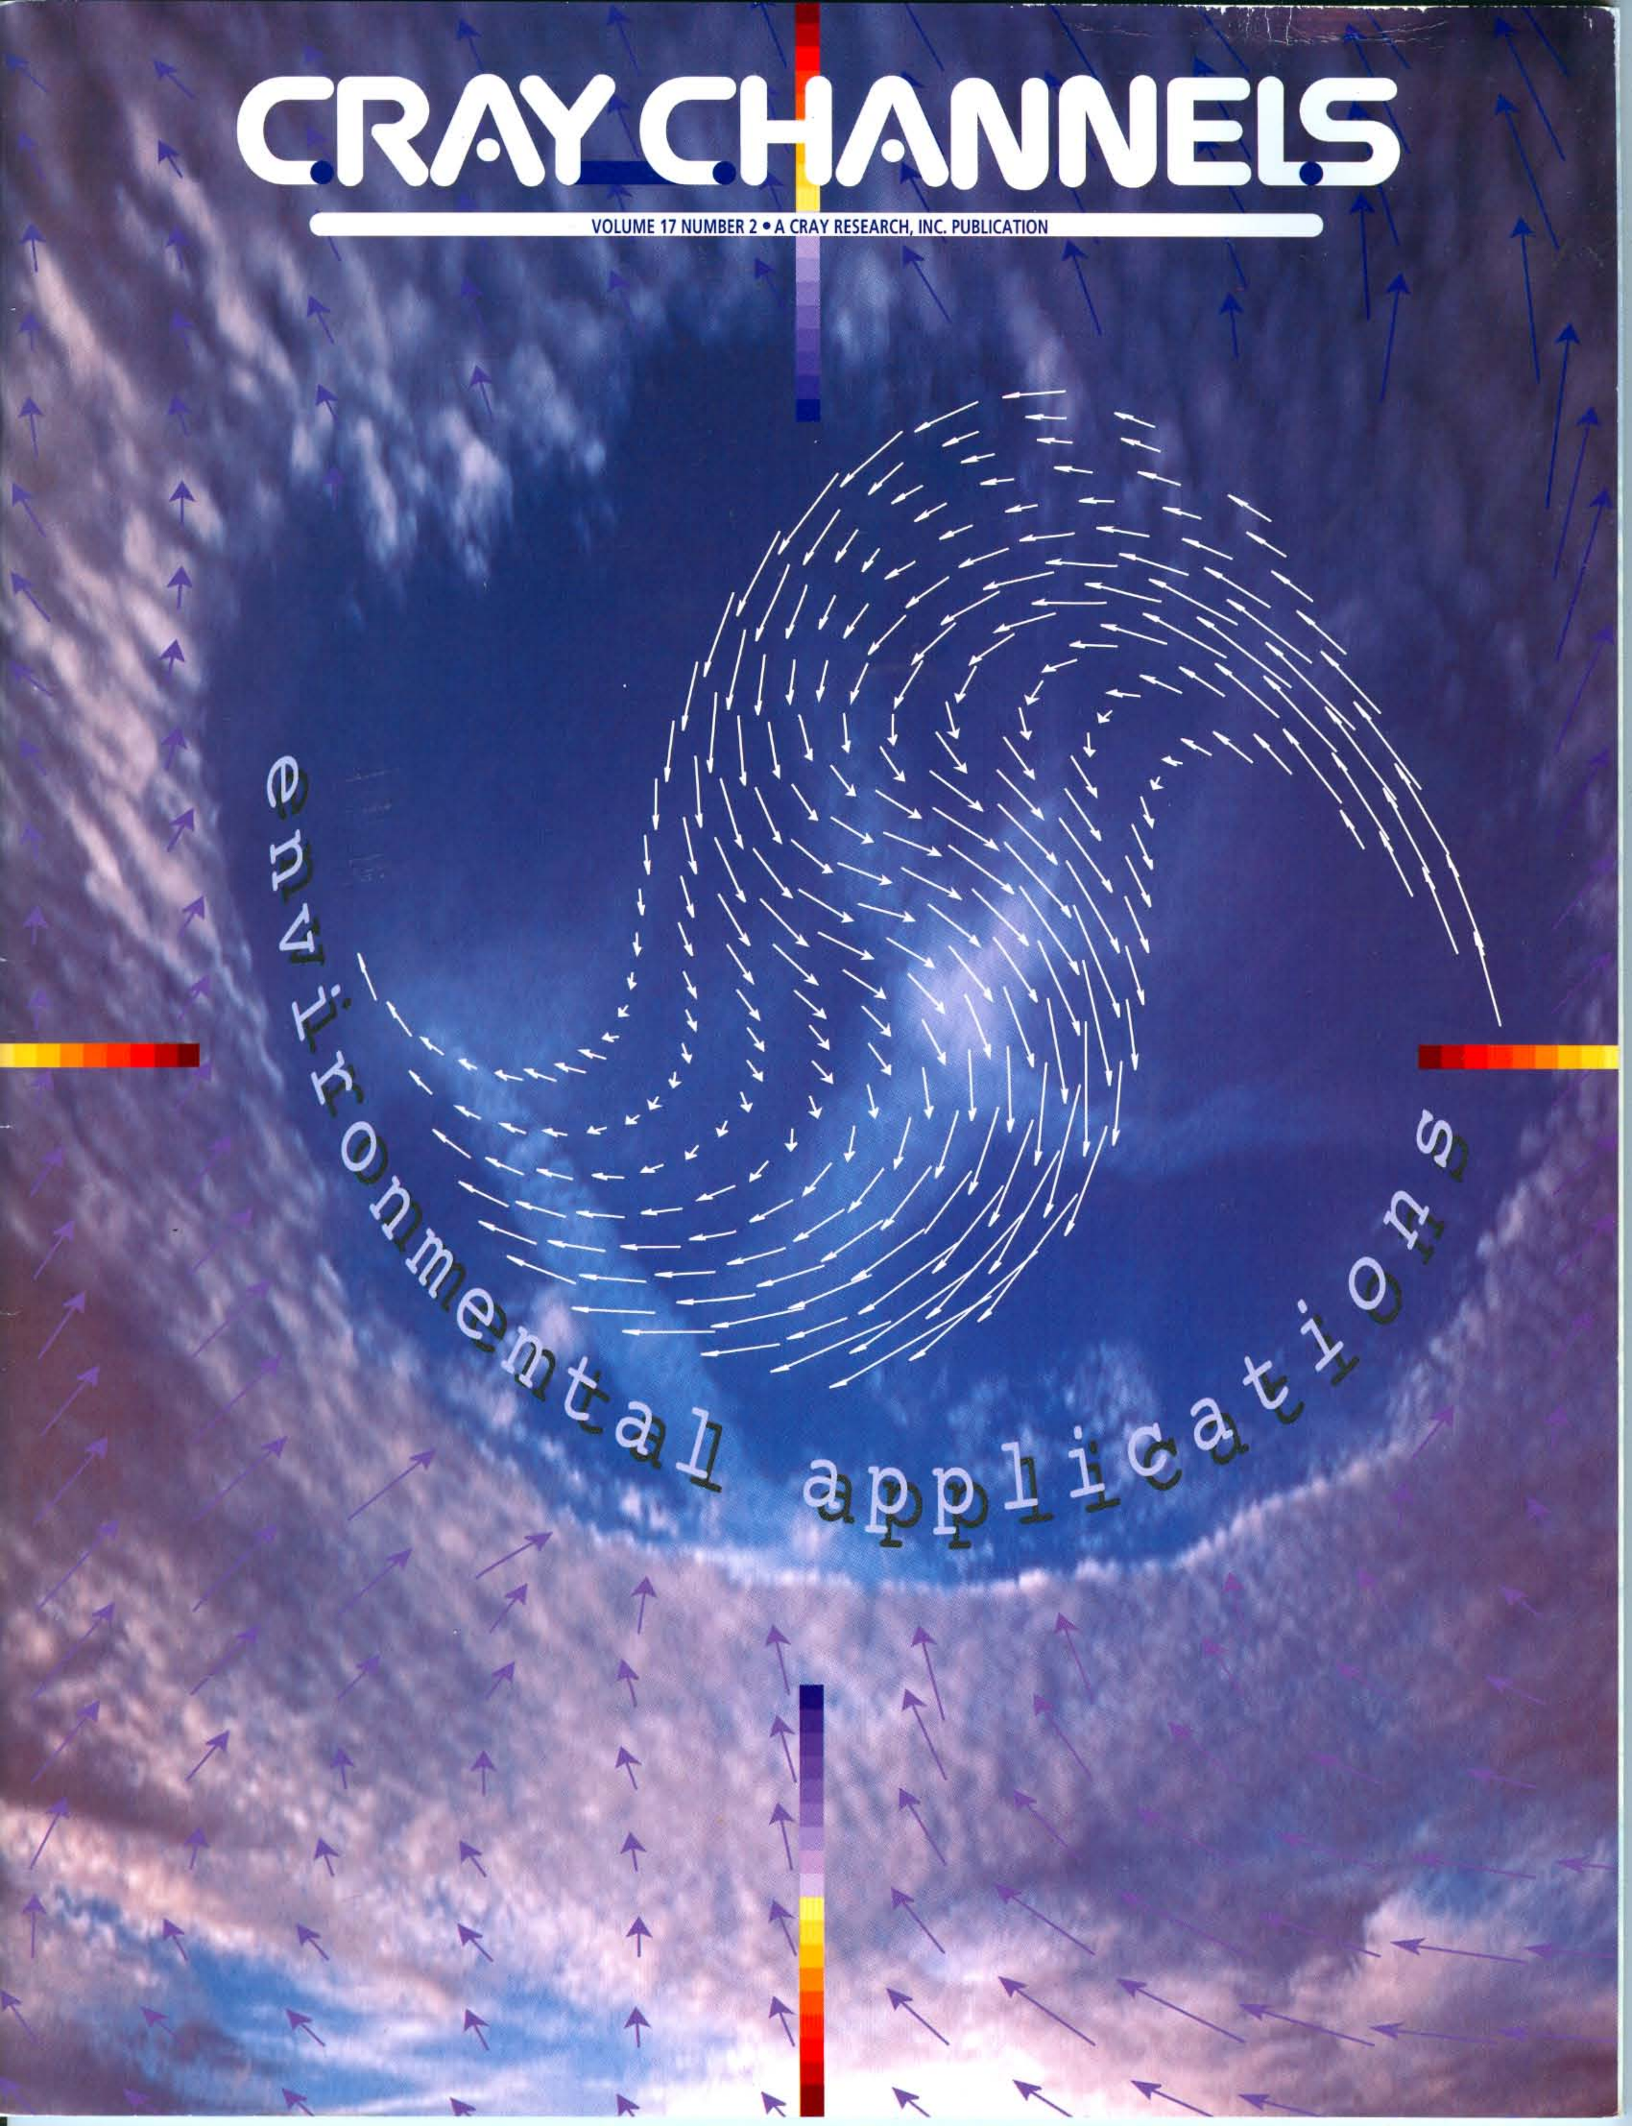

C O N T E N T S

IN THIS ISSUE

Will the summer of 2015 be hot and dry or cold and rainy? Where will wheat grow best in 2002? What will toxic waste of forty years ago do to the groundwater tomorrow, and what does it mean to the health of our children?

If scientists can accurately predict the global climate 50, 10, even 5 years from now, world leaders can make informed decisions and, presumably the best use of taxpayer dollars, to keep countries at optimal agricultural and industrial levels and maintain the environmental health of the planet. Credible climate prediction depends on powerful computers. Today, the majority of environmental work is done on Cray Research systems. Cray Research systems provide the capabilities needed for timely and cost-effective modeling in climate research and other environmental applications.

In this issue of CRAY CHANNELS, we highlight the role of Cray Research systems in weather forecasting and climate research at several laboratories in the United States and Europe. Included are reports on groundwater contamination remediation, improved weather forecasting, improved management of satellite data, and land-atmosphere modeling on the highly parallel CRAY T3D computer system. Also in this issue we introduce the CRAY J932 system—the latest member of the CRAY J90 family of systems—fast becoming a major player in environmental applications.

Our regular departments include information on Earthvision, a computational science program for elementary and high school students devoted exclusively to environmental research; arctic ocean modeling being done by scientists at the Arctic and Antarctic Research Institute in St. Petersburg, Russia; as well as updates on recent application releases and an interview with Cray Research's new CEO, J. Phillip Samper.

FEATURES

DEPARTMENTS

CRAY CHANNELS

CRAY CHANNELS is a publication of the Cray Research, Inc. Marketing Communications Department. It is intended for users of Cray Research computer systems and others interested in the company and its products. Subscription inquiries and address changes should be sent to Department D, CRAY CHANNELS, Cray Research, Inc., 655D Lone Oak Drive, Eagan, Minnesota 55121. Other comments and queries may be emailed to channels@cray.com. For more information about Cray Research and its products, visit http://www.cray.com/ on the World Wide Web.

Volume 17, Number 2

Editorial staff
Jeff Paulson, editor
Cindy Bielke, Heidi Hagen, Ken Jopp, and Molly Sachse, contributing editors

Design and production
Barbara Cahlander, electronic production artist
Eric Hanson, art direction and graphic design

Atmospheric reanalysis project *at the* National Meteorological Center

Robert Kistler
Masao Kanamitsu
John Woollen
Suranjana Saha
and Eugenia Kalnay

U.S. National Meteorological Center
Washington, D.C.

Computer weather forecasts issued by the U.S. National Meteorological Center (NMC) in Washington, D.C., guide forecasts from the U.S. National Weather Service (NWS). NMC forecasts are performed running (integrating in time) computer models of the atmosphere that can simulate, given today's weather observations, the evolution of the atmosphere in the next few days. The time integration of an atmospheric model is an initial value problem; a reliable forecast depends on the computer model being a realistic representation of the atmosphere, with accurate initial conditions.

NMC, to be renamed the National Centers for Environmental Prediction pending congressional approval, has performed operational computer weather forecasts since the 1950s. From 1955 to 1973, the forecasts included only the Northern Hemisphere; they have been global since 1973. Over the years, the quality of the models and methods for using atmospheric observations has improved continuously, resulting in the forecast improvements shown in Figure 1. Unfortunately, every time the model or the analysis scheme is improved, the change also causes spurious changes in the model output. For example, changes in resolution have produced large jumps in the estimated air temperature over the ocean and the average surface pressure over the globe, which are obviously spurious. Other more subtle changes are more difficult to separate from true climate variations. This problem is addressed by performing a reanalysis of past data with a frozen analysis and forecast system.

NMC and the National Center for Atmospheric Research (NCAR), Boulder, Colorado, have started the Reanalysis Project, supported by the National Oceanic and Atmospheric Administration (NOAA) Office of Global Programs, the National Science Foundation, and the NWS. This reanalysis will provide easily accessible data on the state and processes of the global atmosphere for the last 40 years.

Figure 1 shows the longest available score for the skill of numerical weather prediction; it measures the 36-hour forecast error in the horizontal gradient of the geopotential height at 500 hPa (in the middle of the atmosphere) over North America. The empirical experience of human forecasters indicated that a value of 70 corresponded to a useless forecast, and a value of 20 was an essentially perfect forecast; it was the relative difference between analyses hand-made by different forecasters fitting the same observed data.

The improvement in skill over the last 40 years is due to three factors:

- ☐ Improved resolution of the atmospheric model
- ☐ Better incorporation within the model of small-scale physical processes (clouds, precipitation, and turbulent transfers of heat, moisture, momentum, and radiation)
- ☐ Use of more accurate methods of data assimilation, which have resulted in improved initial conditions for the model.

Availability of advanced supercomputing power is critical to these advances. At NMC, the operational forecasts are performed on a 16-processor CRAY C90 supercomputer that has 256 Mwords of memory. The operational global model that will be implemented in early 1996 has a horizontal resolution of about 77 km (triangular truncation T170 in spherical harmonic domain) and 42 vertical levels. The model used in 1985 had a much coarser resolution of about 350 km and 12 vertical levels. In just over one decade, the number of degrees of freedom of the model will increase by a factor of about 20, and the number of computations needed to make a one-day forecast will be about 40 times larger. The model now includes many detailed—and computationally expensive—parameterizations of effects, such as that of clouds on the amount of solar radiation impinging the Earth's surface, or that of soil moisture and vegetation on air temperature and relative humidity.

Improvement of the initial conditions for the model integrations is equally important. The amount and diversity of atmospheric observations gathered every day and used for numerical weather prediction are staggering: about 1100 balloons measuring temperature, moisture, and winds are launched at 00:00 and 12:00 Greenwich meridian time (GMT). There also are about 16,000 satellite-derived temperature soundings; 11,000 satellite wind

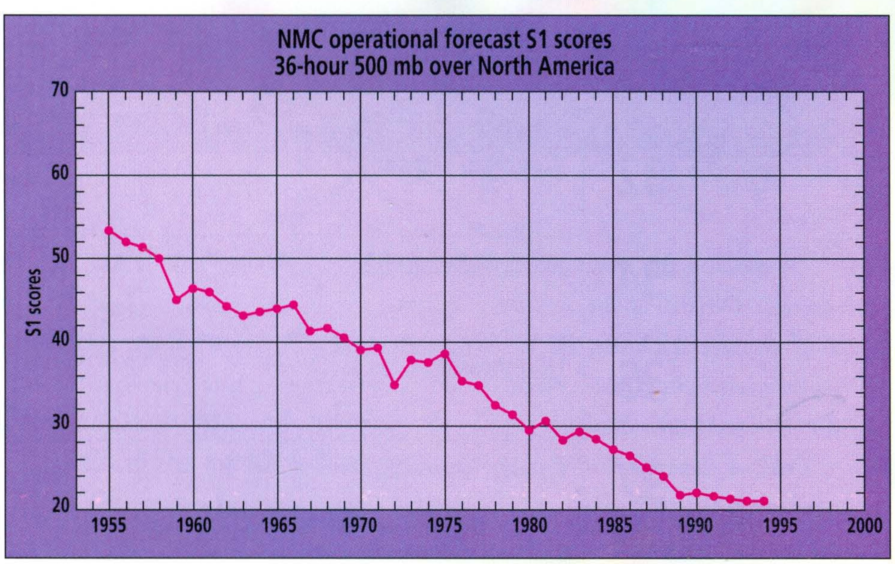

Figure 1. Evolution of the S1 score (percentage relative error in the horizontal pressure gradient) over North America since the beginning of numerical weather prediction. 70 represents a useless forecast; 20 represents a forecast that, subjectively, fits well with observations.

measurements derived from the motion of clouds; 28,000 surface observations of pressure, temperature, wind, and humidity; and, currently, 16,000 aircraft observations of temperature and winds. These data are distributed unevenly throughout the world and vary in accuracy. To start the global atmospheric model integration, initial conditions are needed for all five prognostic variables of the model (two horizontal wind components, temperature, humidity, and surface pressure) at every point for every level of the model. In other words, to perform forecasts, we need well over 1,000,000 pieces of information four times a day, but the observations provide two orders of magnitude less data than needed.

Four-dimensional data assimilation

To solve this problem of underdetermination, operational weather forecasting centers perform "four-dimensional data assimilation." This is done through a six-hour analysis cycle that is executed for 00:00, 06:00, 12:00, and 18:00 GMT every day. Suppose that the current time is 12:00 GMT of any given day, and a six-hour forecast is made by using the global atmospheric model. This short forecast is a first-guess estimate of the complete state of the atmosphere six hours later, at 18 GMT. The data corresponding to 18 GMT observations (within a window of –3 to +3 hours) are gathered. To produce the best estimate of the state of the atmosphere at 18 GMT, the 18 GMT observations and the 18 GMT forecast are combined. This statistical interpolation (analysis) of the data and the model is then used as the initial conditions for the model, and the next six-hour cycle can begin. This process never stops; four times a day atmospheric observations are assimilated into the global model through the analysis cycle. The model is essential to solving the problem of underdetermination. It transports information about the state of the atmosphere from data-rich to data-poor regions, and it produces a more accurate estimate of the state of all atmospheric variables, even those that are not observed.

At NMC, the statistical method used in the analysis cycle to create initial conditions for the model is performed in a variational scheme denoted "Spectral Statistical Interpolation" (SSI, also known as "three-dimensional variational data assimilation," or 3-D VAR).[1] This method finds the 3-D state of the global atmosphere closest to both the six-hour forecast and the 20,000 or so global observations gathered over approximately three hours. This very large minimization problem requires inverting, every six hours, a matrix of the size of the model (about 1,000,000 elements).

The output of the analysis cycle is a gridded best estimate of the complete state of the atmosphere every six hours. Its main use is to provide initial conditions for the atmospheric model. The operational analyses also are used in many other research areas, such as atmospheric dynamics, predictability studies, and, very importantly, the monitoring of atmospheric climate anomalies. The NMC operational analysis has been one of the most widely used tools for monitoring short (weeks through interannual) climate variations.

Reanalysis of past data

In the NMC/NCAR Reanalysis Project, a 40-year reanalysis of past atmospheric and oceanic data will be performed over a period of three years, using a state-of-the-art but frozen data assimilation system.[2] The observations, including many datasets that were not available in real time, have been gathered at NCAR. The old data will be reassimilated within an analysis cycle, but the statistical analysis and model modules are kept unchanged throughout the reanalysis. Since January 1995, the same system of data assimilation also is being continued into the future as a "Climate Data Assimilation System" (CDAS); this system will allow climate researchers to compare the present climate anomalies with a very long climatology obtained with the same frozen reanalysis system. The project should be completed by 1997. At that time, 40 years of reanalysis/CDAS information (1957 to 1996) will be available to the research community.

To perform this huge task in a reasonable amount of time, we must carry out the reanalysis at a much faster rate than the operations (which proceed at a rate of just one day of analysis per clock day). This has required the development of a new system that can perform one month of reanalysis per clock day, and for which humans can monitor the quality of the input data and of the output products at a much faster rate than is possible in operations. The reanalysis system was developed over three years; execution began in June 1994.

System configuration

The CDAS/reanalysis is executed at the NOAA Central Computer Facility in Suitland, Maryland. Unlike the operational NMC system, which currently is based on IBM systems running the MVS operating system and Cray Research supercomputers running UNICOS (Cray Research's implementation of the UNIX operating system), all processing in the CDAS/Reanalysis System is done in UNICOS. This new system eliminates the transmission of large amounts of data across a variety of systems, has many new advanced features, and will be adopted in NMC's normal operations.

We will conduct the reanalysis by using our 8-processor CRAY Y-MP system with 128 Mwords of memory and a CRAY EL supercomputer. Other hardware includes a robotic silo, upgraded to 4490 STK, with storage capacity of 0.6 Gbytes/tape. Because the CRAY Y-MP system was saturated, we could not begin the reanalysis until a new CRAY C90 system was installed at NMC in early 1994 and the operational system was migrated out of the CRAY Y-MP system in April. Software used includes UNICOS 7.0.5, workstation NFS mount of Cray complex files, Bourne shell UNIX scripts, Fortran, some C, some X Window System software, Cray Research's Data Migration Facility, the Cray/REELlibrarian, and the graphics system GrADS (COLA). We plan to install UNICOS 8.0 in the near future and replace the Bourne shell with the POSIX (Korn) shell. In 1995, we also expect to replace the CRAY Y-MP and CRAY EL systems with two CRAY J916 systems.

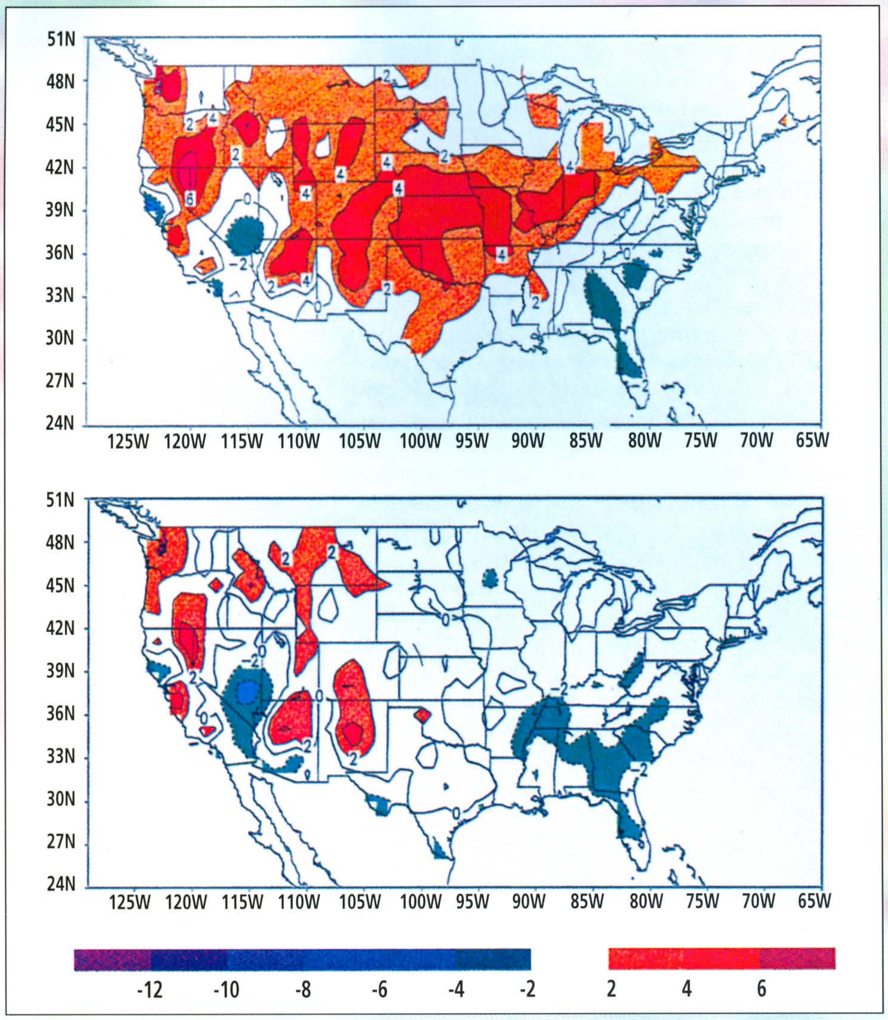

Figure 2a (top). Original reanalysis with unrealistically high plant evaporation resistance.

Figure 2b (above). New reanalysis with corrected plant resistance.

Figure 3 (below). Four-times-a-day reanalysis for December 29, 1990, plotted over the United States. Low-level (850 hPa, about 1500 m) winds in red, high-level (250 hPa, about flight level) winds in blue, and accumulated precipitation in green.

The Reanalysis System

The NMC/NCAR Reanalysis System has three major components: a data preparation module, a data assimilation module, and output and data distribution.

In the data preparation module, observations from many sources and formats are decoded, sorted, and checked for duplicates. The observations are encoded into binary universal format representation (BUFR). This BUFR representation has been expanded to include not just the data itself, but also its history as it is processed and quality controlled in the reanalysis system. For example, each piece of data will include the climatological average for that variable at its space and time location, as well as the first guess and the final analysis value. When communication problems cause simple errors (such as digit transposition), the complex quality control system at NMC frequently corrects the observations. The BUFR data will contain the original value, the corrected value, and the results of other quality control decisions.

The data preparation module preprocesses data for sequences of a year or longer, and the results are presented to human monitors. If major problems are encountered, such as missing or mislocated data, specialists have a few days to correct them before executing the computationally costly data assimilation module.

The data assimilation module is the core of the system, and it includes a model and analysis system. The model has a horizontal resolution of about 210 km and 28 vertical levels. This model is identical to the NMC global model implemented operationally in January 1995 except for horizontal resolution, which is double for the operational model.[3] The analysis is performed with the spectral statistical interpolation analysis. There is complex quality control of rawinsonde data, with confident corrections of heights and temperatures,[4] and optimal interpolation-based complex quality control of all other data.[5,6]

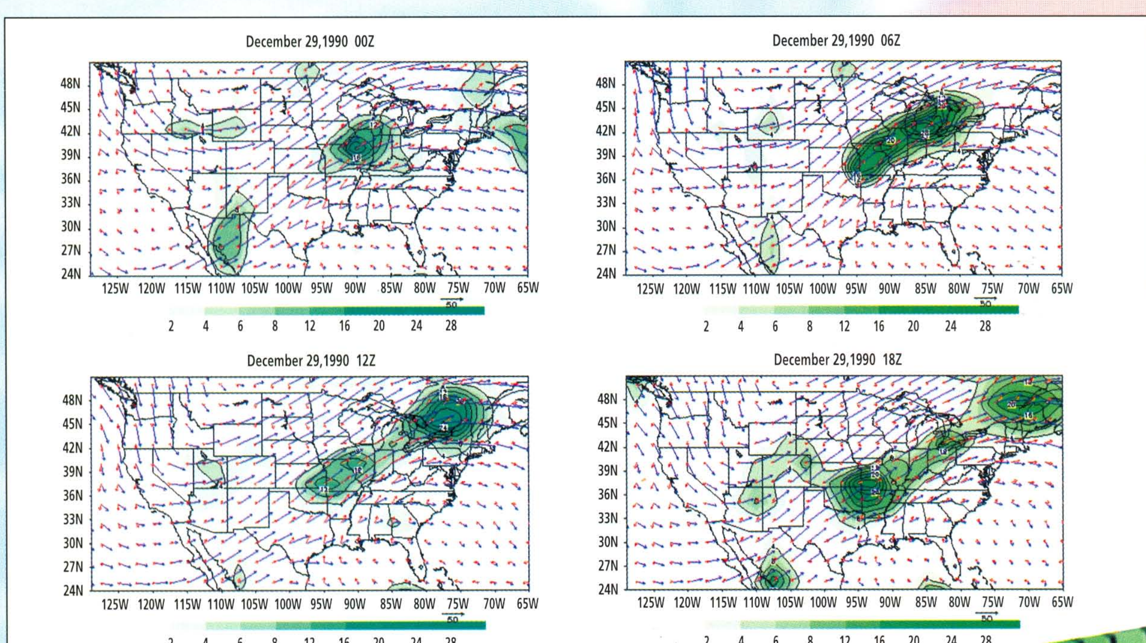

The third and final component of the NMC/NCAR Reanalysis System is output and data distribution. It has become clear that many applications exist for the reanalysis, and that some of them (for example, transport of greenhouse gases, which needs, in principle, all turbulent transports between any two layers) have storage requirements that far exceed what the project can handle. For this reason, we decided that each unit of reanalysis output (one month) will include restart files so that, for special purposes, shorter reanalyses with extended output can be performed a posteriori.

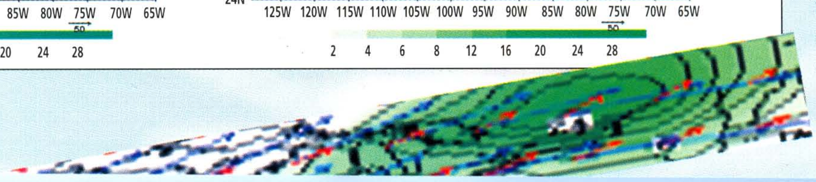

We designed the reanalysis archive to satisfy two major user requirements: the output should be comprehensive, allowing, for example, the performance of detailed water budget studies; and the output should be easily accessible to the user interested in a long time series of data. A single archival format cannot satisfy both requirements; the output module includes several different archives:

- Level-2 observational data in BUFR, including quality control, climatological, analysis, and six-hour forecast information.
- Comprehensive analysis, first guess, and diagnostic fields presented in "synoptic" form (all fields every six hours) in the model sigma coordinates, as well as in pressure and isentropic coordinates, in gridded binary (GRIB) format. A restart file is included once a month to allow rerunning shorter periods with enhanced diagnostics.
- A time series archive in which each field is available for all times, including standard pressure level fields, precipitation, surface fluxes, and other widely used diagnostic fields. This format will be the most useful for many users.
- A "quick look" archive on CD-ROMs, one per year, including the most widely used fields: daily values of variables at selected pressure levels, surface and top of the atmosphere fluxes, precipitation, monthly and zonal averages of most quantities, covariances, and isentropic level variables.
- Eight-day forecasts performed every five days, which should allow predictabiltiy studies and estimates of the impact of inhomogeneities in the observing systems coverage, with anomaly scores.

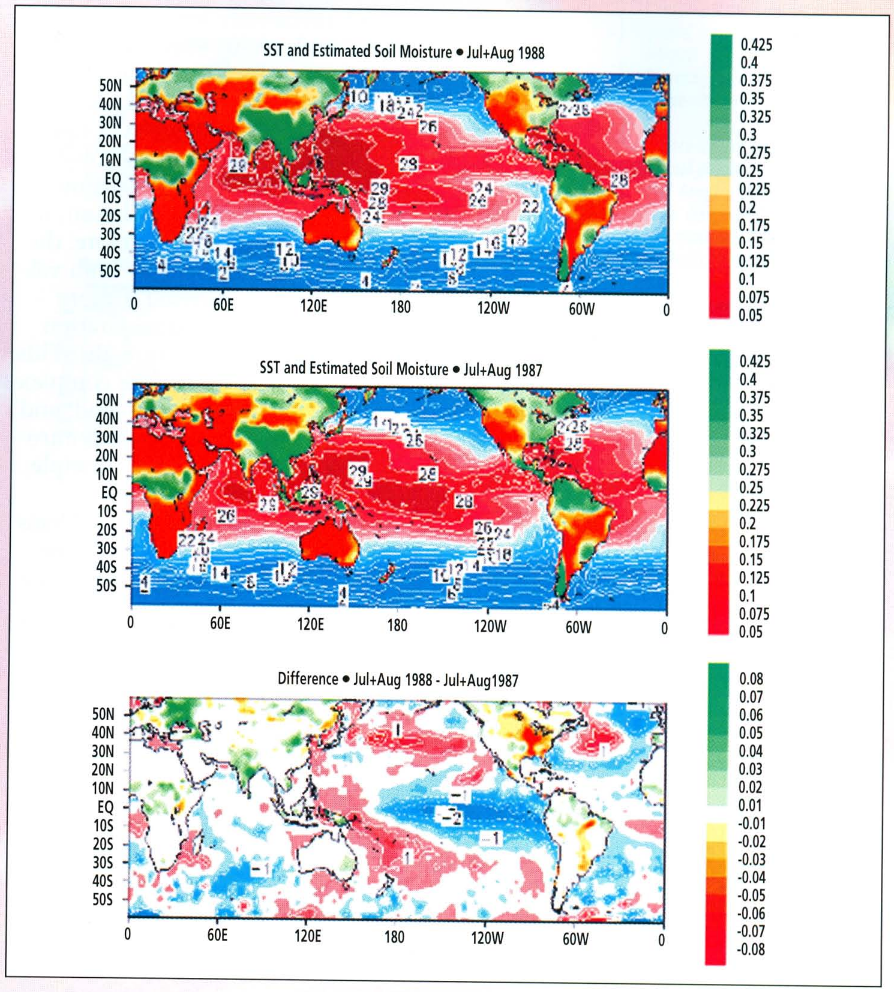

Figure 4. Soil moisture and sea surface temperature for July and August 1987, 1988, and their difference.

Part of the output is posted in the NMC public server NIC and is available through anonymous ftp. NCAR, the National Climatic Data Center, and the Environmental Research Laboratories/Climate Diagnostic Center will distribute the bulk of the reanalysis data.

Execution

In June 1994, the execution phase of the reanalysis started on the CRAY Y-MP system NMC provided for this project. About 24 hours of the CRAY Y-MP supercomputer (2 to 7 processors) are needed to perform one month of reanalysis and forecasts per day. By July 1995, 10 years (1985 to 1994) were completed, in addition to several years of reruns performed to assess the impact of changes in observing systems. The period 1979 to 1985 is being reanalyzed next and will be completed around December 1995, followed by the 1957 to 1978 decades. We expect to complete the 40 years of reanalysis (1957 to 1996) by early 1997. The extension into 1948 to 1957, if feasible, will be done during 1997.

In the reanalysis system, the processing of observations is I/O- and single-processor-bound. To reduce the resulting inefficiency, we make use of the available eight CPUs by running several single-processor batch jobs in parallel whenever possible. For example, preprocessing, data assimilation, postprocessing, and tape copies are performed simultaneously on the CRAY Y-MP system. We have reserved 18 Gbytes of disk space for this project to minimize movement of data and to facilitate file transposes that go from the original "every six hour" output to the order most convenient for data access used in the backup tapes. In addition to executing the reanalysis, the CRAY Y-MP system is used for long-term climate simulations and accessed by non-NMC researchers working in the area of coupled ocean-atmosphere seasonal forecasting.

This first phase of reanalysis will be followed by a second phase in which a 1998 state-of-the-art system will be used for reanalysis. NMC plans currently call for an updated reanalysis every five years or so.

Results

After the reanalysis of 1985 and 1986 was completed, it became clear that the reanalysis surface temperature over the midwest and eastern sections of North America was unrealistically high (Figure 2a). Considerable investigation traced this puzzling result to the new soil model incorporated in the reanalysis model, which also includes the effect of vegetation resistance on the surface evapotranspiration. NMC had been using a geographical distribution of vegetation that assumed that the eastern half of North America is cultivated with winter wheat, and that in the early summer months

We designed the reanalysis archive to satisfy two major user requirements: the output should be comprehensive, allowing, for example, the performance of detailed water budget studies; and the output should be easily accessible to the user interested in a long time series of data. A single archival format cannot satisfy both requirements; the output module includes several different archives:

- ☐ Level-2 observational data in BUFR, including quality control, climatological, analysis, and six-hour forecast information.
- ☐ Comprehensive analysis, first guess, and diagnostic fields presented in "synoptic" form (all fields every six hours) in the model sigma coordinates, as well as in pressure and isentropic coordinates, in gridded binary (GRIB) format. A restart file is included once a month to allow rerunning shorter periods with enhanced diagnostics.
- ☐ A time series archive in which each field is available for all times, including standard pressure level fields, precipitation, surface fluxes, and other widely used diagnostic fields. This format will be the most useful for many users.
- ☐ A "quick look" archive on CD-ROMs, one per year, including the most widely used fields: daily values of variables at selected pressure levels, surface and top of the atmosphere fluxes, precipitation, monthly and zonal averages of most quantities, covariances, and isentropic level variables.
- ☐ Eight-day forecasts performed every five days, which should allow predictabiltiy studies and estimates of the impact of inhomogeneities in the observing systems coverage, with anomaly scores.

Part of the output is posted in the NMC public server NIC and is available through anonymous ftp. NCAR, the National Climatic Data Center, and the Environmental Research Laboratories/Climate Diagnostic Center will distribute the bulk of the reanalysis data.

Execution

In June 1994, the execution phase of the reanalysis started on the CRAY Y-MP system NMC provided for this project. About 24 hours of the CRAY Y-MP supercomputer (2 to 7 processors) are needed to perform one month of reanalysis and forecasts per day. By July 1995, 10 years (1985 to 1994) were completed, in addition to several years of reruns performed to assess the impact of changes in observing systems. The period 1979 to 1985 is being reanalyzed next and will be completed around December 1995, followed by the 1957 to 1978 decades. We expect to complete the 40 years of reanalysis (1957 to 1996) by early 1997. The extension into 1948 to 1957, if feasible, will be done during 1997.

In the reanalysis system, the processing of observations is I/O- and single-processor-bound. To reduce the resulting inefficiency, we make use of the available eight CPUs by running several single-processor batch jobs in parallel whenever possible. For example, preprocessing, data assimilation, postprocessing, and tape copies are performed simultaneously on the CRAY Y-MP system. We have reserved 18 Gbytes of disk space for this project to minimize movement of data and to facilitate file transposes that go from the original "every six hour" output to the order most convenient for data access used in the backup tapes. In addition to executing the reanalysis, the CRAY Y-MP system is used for long-term climate simulations and accessed by non-NMC researchers working in the area of coupled ocean-atmosphere seasonal forecasting.

This first phase of reanalysis will be followed by a second phase in which a 1998 state-of-the-art system will be used for reanalysis. NMC plans currently call for an updated reanalysis every five years or so.

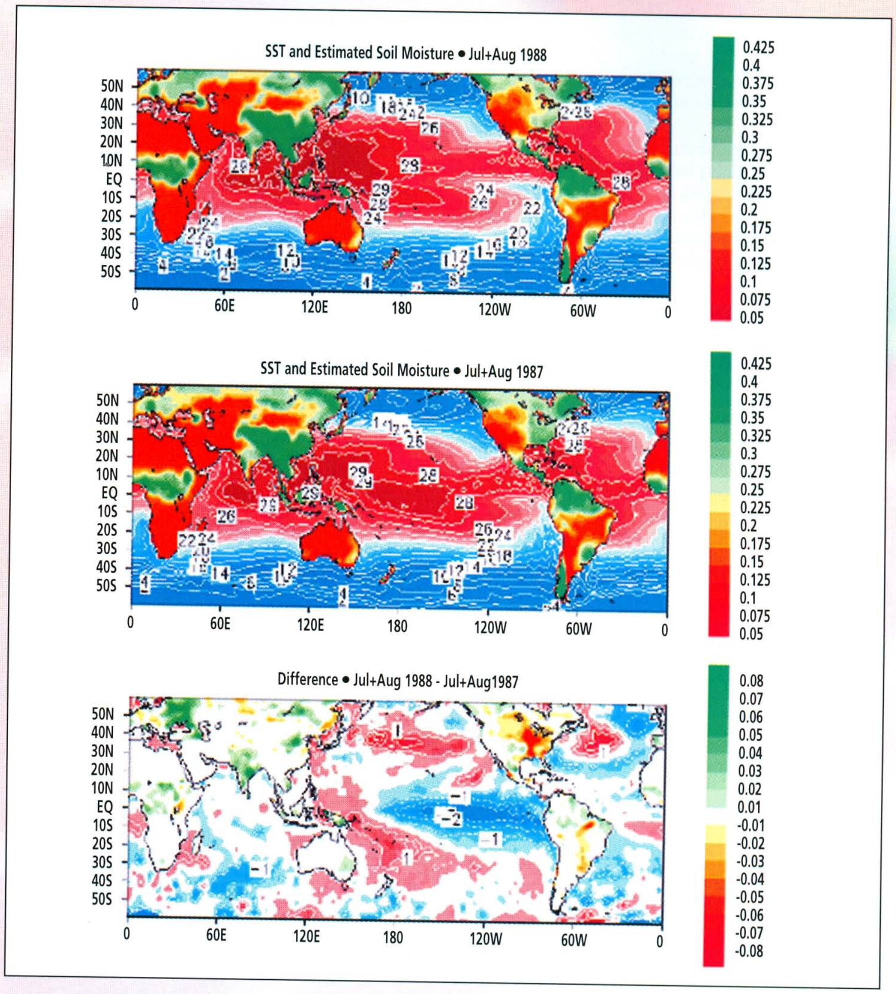

Figure 4. Soil moisture and sea surface temperature for July and August 1987, 1988, and their difference.

Results

After the reanalysis of 1985 and 1986 was completed, it became clear that the reanalysis surface temperature over the midwest and eastern sections of North America was unrealistically high (Figure 2a). Considerable investigation traced this puzzling result to the new soil model incorporated in the reanalysis model, which also includes the effect of vegetation resistance on the surface evapotranspiration. NMC had been using a geographical distribution of vegetation that assumed that the eastern half of North America is cultivated with winter wheat, and that in the early summer months this crop is harvested, essentially leaving half of the continent covered with straw. This unrealistic vegetation cover did not produce problems with the simpler soil model used at NMC until January 1995 because it was rather insensitive to the effects of vegetation. With the new soil model, however, the result of covering the land with straw meant that when the ground was heated by insolation, it could not be cooled by evaporation. Therefore, the ground temperature rose to unrealistically high values, ruining the two years of reanalysis! A more realistic definition of the plant evapotranspiration resistance eliminated this problem (Figure 2b). This was a sobering lesson on how carefully the complex interactions between the atmosphere, the land, and the ocean must be considered, especially when introducing a more complex model which, in principle, should lead to better results.

As the reanalysis proceeded, other problems were found, many of which were corrected; some cannot be solved with current state-of-the-art systems. For fields for which no data is incorporated into the analysis, the model-produced results can be considered only estimates.

Figure 3 shows the accumulated precipitation in green, 850 hPa winds (about 1500 m above sea level) in red, and 250 hPa (about flight level) in blue, for four six-hour periods of December 29, 1990. The winds are derived from data, and the precipitation is estimated from the model. The CD-ROMs will make accessible this and much more information about the weather patterns for the whole world every 12 hours for the 40 years of reanalysis.

The next set of figures presents fields that are not directly measured or assimilated into the model (except for the sea surface temperature field). Although there is no guarantee that these fields are realistic, the results are quite encouraging.

Figure 4 shows the sea surface temperature (SST) analyzed over the oceans and the soil moisture content over land for July and August of 1987, 1988, and their difference. Figure 5a presents the outgoing long-wave radiation (the radiation emitted by the Earth and the atmosphere) as observed by satellites, and Figure 5b shows the outgoing long-wave radiation derived from the reanalysis for the same two years, as well as their difference. During the summer of 1987, an El Niño, or "warm episode," was occurring. This large anomaly of the coupled ocean-atmosphere system is characterized by unusually warm SST in the central Pacific (near the dateline). The atmosphere responds to this anomaly by shifting the massive clouds and precipitation normally concentrated over the western Pacific toward the east, where the anomalous warm SSTs are located. During the summer of 1988, the situation was almost reversed; a "cold episode" made the SST in the central Pacific unusually cold. These effects are quite apparent in the SST analysis. The changes in the tropical, and to some extent the extratropical, climate between these two extreme years also are clear, and there generally is a good agreement between the fields that the reanalysis produces and the satellite observations of outgoing long-wave radiation, which were not used in the reanalysis.

Finally, Figures 6a and 6b show the precipitation fields for the same months, both for the reanalysis and for microwave satellite estimates. In general there is also good agreement between estimates made from space and in the reanalysis. For example, the summer of 1988 was much wetter than the summer of 1987 in the Indian monsoon area, but drier over the United States. This is apparent in both precipitation fields as well as in the soil moisture estimated from the reanalysis (Figure 4). More information about the reanalysis can be found at http://sgi62.wwb.noaa.gov:8000/research/reanl.html.

Figure 5a (right). Outgoing long-wave radiation (OLR) for July and August 1987, 1988, and their difference, observed by satellite.

Figure 5b (far right). Outgoing long-wave radiation (OLR) for July and August 1987, 1988, and their difference, estimated by the model during the reanalysis without using the satellite OLR data.

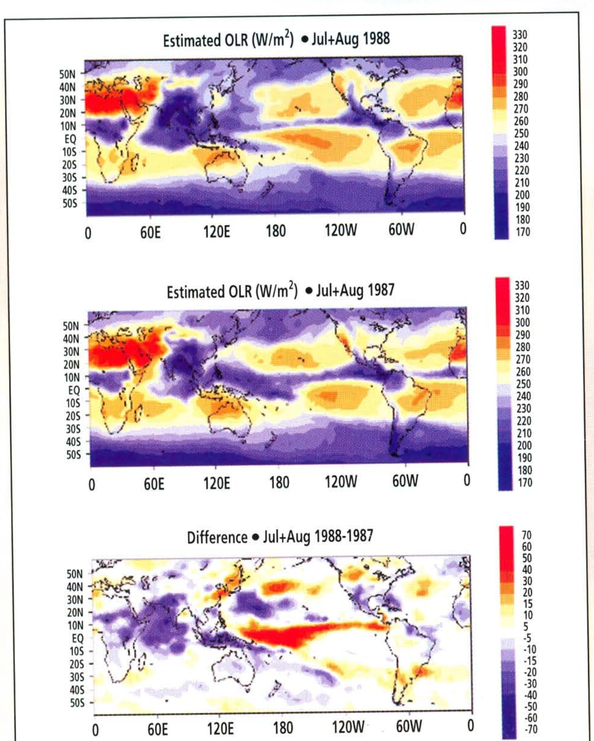

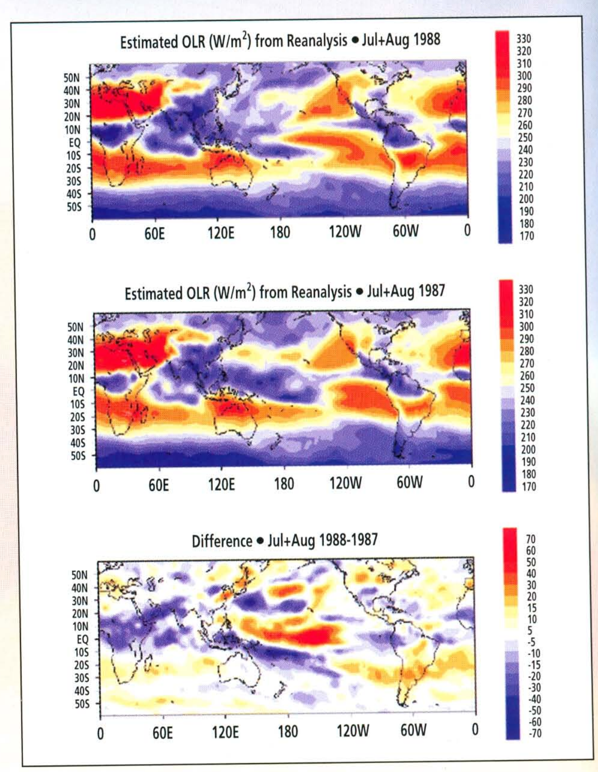

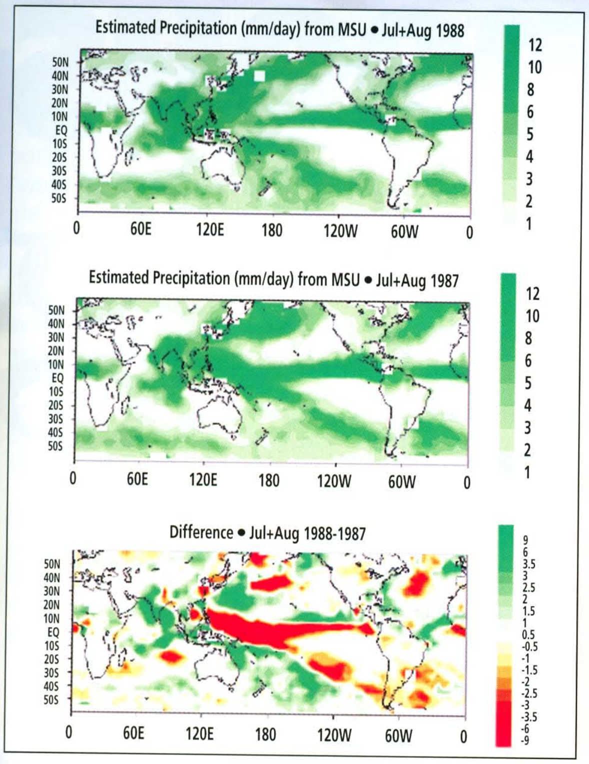

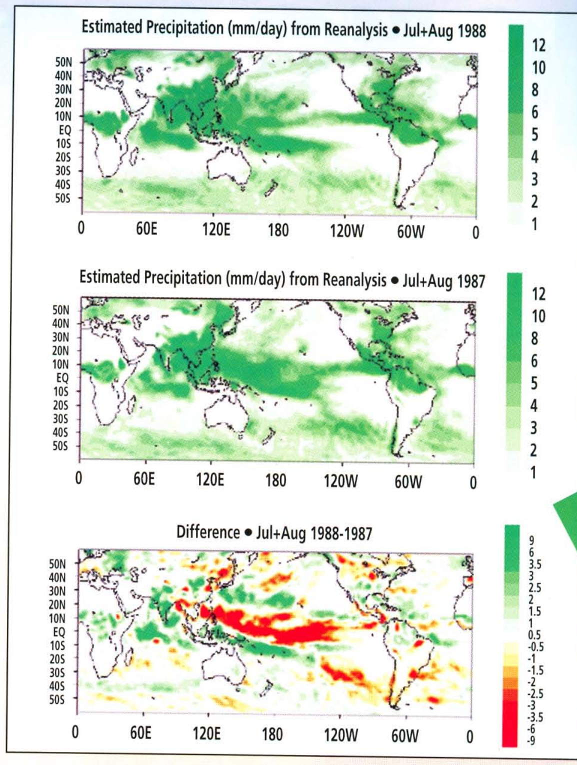

Figure 6a (far left). Accumulated global precipitation for July and August 1987, 1988, and their difference, estimated from microwave data.

Figure 6b (left). Accumulated global precipitation for July and August 1987, 1988, and their difference, estimated from the model during the reanalysis without using precipitation data.

Acknowledgments

The dedicated effort of many people (including the co-authors of the Kalnay et al., 1995 paper) at NMC and at NCAR has made the Reanalysis Project possible. The project has benefited from the enthusiastic support of Ron McPherson, director of NMC, the active participation of Development Division and Climate Analysis Center scientists, and the help of the Automation Division of NMC and Cray Research personnel. The NOAA Office of Global Programs and the National Science Foundation have contributed to the funding of the project at NMC and NCAR, respectively. Mallse C. Dick and Huug van den Dool improved the readability of the article. The guidance of the Reanalysis Advisory Committee also is gratefully acknowledged.

About the authors

Eugenia Kalnay received an undergraduate degree from the University of Buenos Aires in 1965 and a Ph.D. degree in meteorology from the Massachusetts Institute of Technology in 1971. She was a professor at the University of Montevideo and at MIT before joining the NASA Goddard Laboratory for Atmospheres in 1979 and NMC in 1987. She and Roy Jenne of NCAR are co-principal investigators of the reanalysis project.

Masao Kanamitsu did his undergraduate studies at Hokkaido University and received a Ph.D. degree in meteorology from Florida State University in 1975. He worked at the Japanese Meteorological Agency and at the European Centre for Medium-Range Weather Forecasts before joining NMC in 1985. He is the chief of the Global Weather and Climate Modeling Branch of the development division and the principal scientist of the reanalysis project.

Robert Kistler is a meteorologist in the development division, Global Weather and Climate Modeling Branch. He received B.S. and M.S. degrees in meteorology from Penn State University. He joined the development division in 1973 and is the manager of the NMC/NCAR Reanalysis Project.

Suranjana Saha obtained a Ph.D. degree in meteorology from the Jawaharlal Nehru University in New Delhi, India in 1983 and became a postdoctoral research associate at the University of Maryland in 1984. She joined the development division in 1987. Involved in the reanalysis system since its start, she is responsible for the development of the automatic quality control of the reanalysis output.

Jack Woollen is a systems analyst with General Sciences Corporation at the development division. Before joining NMC in 1989, he was at the NASA Goddard Laboratory for Atmospheres. He has developed the new BUFR archiving system, which allows the addition of quality control information to each observation as it is processed by the reanalysis system.

References

1. Parrish, D. F. and J. C. Derber, "The National Meteorological Center's Spectral Statistical Interpolation Analysis System," *Monthly Weather Review,* Vol. 120, pp. 1747-1763, 1992.
2. Kalnay, E., M. Kanamitsu, R. Kistler, W. Collins, D. Deaven, L. Gandin, M. Iredell, S. Saha, G. White, J. Woollen, Y. Zhu, M. Chelliah, W. Ebisuzaki, J. Janowiak, K. C. Mo, C. Ropelewski, J. Wang, A. Leetmaa, R. Reynolds, and Roy Jenne, "The NMC/NCAR Reanalysis Project," submitted to *Bulletin of the American Meteorological Society,* 1995.
3. Kanamitsu, M., "Description of the NMC Global Data Assimilation and Forecast System," *Weather and Forecasting,* Vol. 4, pp. 334-342, 1989.
4. Collins, W. G. and L. S. Gandin, "Comprehensive Hydrostatic Quality Control at the National Meteorological Center," *Monthly Weather Review,* Vol. 18, pp. 2754-2767, 1990.
5. Woollen, J. S., "New NMC Operational Optimal Interpolation Quality Control," Ninth Conference on Numerical Weather Prediction, American Meteorological Society, pp. 24-27, 1991.
6. Woollen, J. S., E. Kalnay, L. Gandin, W. Collins, S. Saha, R. Kistler, M. Kanamitsu, and M. Chelliah, "Quality Control in the Reanalysis System," Tenth Conference on Numerical Weather Prediction, American Meteorological Society, pp. 13-14, 1994.

Heterogeneous computing

for climate research

Carlos R. Mechoso, John D. Farrara, and Joseph A. Spahr

Department of Atmospheric Sciences
University of California, Los Angeles

Numerical simulation of the Earth's climate is a "grand challenge" problem that requires a high level of computing resources. Gathering the required resources may be possible through the creation of "metacomputers" consisting of several supercomputers connected by a high-speed network. One example of such a "metacomputer" is a CRAY C90 system connected by a gigabit-per-second network to a CRAY T3D system. Such an arrangement may provide superlinear speedups in run times for climate models when the model tasks are distributed among the metacomputer's component systems. Obtaining these speedups requires the use of special techniques for code distribution. The CASA Gigabit Network Testbed project, which is sponsored by the National Science Foundation and the Advanced Research Projects Agency, is developing these techniques.

Climate is the quintessential example of a "nonlinear" system. In this context, nonlinearity means that doubling the forcing, the amount of atmospheric carbon dioxide (CO_2), for example, does not necessarily result in a doubling of the parameters that describe global climate, such as temperature and precipitation. The climate system also is characterized by complex interactions and feedbacks among its components. Growing concerns over a deteriorating environment are encouraging the development of numerical models of the climate system that can represent these nonlinearities, interactions, and feedbacks. Research groups throughout the world use such models to produce "scenarios"—or predictions—to guide the planning that will be required to mitigate the deleterious effects of climate change on Earth's biosphere.

At the University of California, Los Angeles (UCLA), we are actively developing an Earth System Model (ESM) to study problems in climate, climate change, and climate/chemistry interactions, including the general circulation of the coupled atmosphere-ocean system, global distributions of greenhouse gases, and global ozone perturbations. The UCLA ESM is based on the following comprehensive models of the atmospheric and oceanic circulations, and chemical tracers:

- The UCLA general circulation model of the atmosphere (AGCM)
- The GFDL/Princeton University general circulation model of the ocean (OGCM)
- The NASA Ames/UCLA atmospheric chemical/aerosol tracer model (ACTM) and
- The UCLA oceanic chemical/aquasol tracer model (OCTM).

A unique feature of the UCLA ESM is that the model is being configured to run either on a single supercomputer or distributed across several computers with different architectures and connected via local or wide-area networks with broad bandwidths. This highly experimental work is part of the CASA Testbed project, which is investigating whether a "metacomputer" that consists of distributed supercomputers connected by a high-speed (gigabit-per-second) network is a viable concept for large scientific applications. A particularly challenging aspect of the problem addressed by CASA is to determine whether the metacomputer paradigm can result in superlinear speedups of execution despite the added overhead of distributed systems due to latency and communication delays. The CASA testbed is an integral part of the Gigabit Project sponsored by the National Science Foundation (NSF) and the Advanced Research Projects Agency (ARPA).

Table 1. Timings (seconds/simulated day) of the atmospheric GCM.

	CRAY C90 (8 processors)	CRAY T3D (242 processors)	CRAY C90(8)-T3D(242)
AGCM/Dynamics	x	29	x
AGCM/Physics	17	7	7
Total	$x + 17$	36	$x + 7$

CASA, which also addresses the formidable technical issues associated with gigabit-per-second bandwidths, comprises nearly 60 computers at the San Diego Supercomputer Center (SDSC), California Institute of Technology (Caltech), NASA's Jet Propulsion Laboratory (JPL), and Los Alamos National Laboratory (LANL) linked by approximately 3000 km of fiber optic cable.

Speedup of execution of the AGCM

The UCLA AGCM is a complex code representing many physical processes.[1] Despite the complexity of the code, one can identify the following two major components:

- AGCM/Dynamics, which computes the evolution of the fluid flow governed by the appropriate equations (the primitive equations) written in finite differences and
- AGCM/Physics, which computes the effect of processes not resolved by the model's grid (such as convection on cloud scales) on processes that are resolved by the grid (such as the flow on the large scales).

The results obtained by AGCM/Physics are supplied to AGCM/Dynamics as drivers for the flow simulated by this component. When the model domain is divided in subdomains that consist of vertical columns from the Earth's surface to the top of the atmosphere for parallel computation in a multiprocessor machine, AGCM/Physics scales very well because almost no communication is required between columns.[2] AGCM/Dynamics, on the other hand, scales less well because substantial amounts of interprocessor communication are required by the horizontal finite differencing scheme. The AGCM components can be run as separate processes on different computers connected by a network. This type of task decomposition allows the different model components to be assigned to the machines with the most suitable architecture.

To illustrate the potential for speedup of the AGCM, we consider the distribution of the code across a metacomputer consisting of an 8-processor CRAY C90 system and a 242-processor CRAY T3D system. We refer to a metacomputer by the name of the individual machines followed by the number of processors in parentheses and separated by a dash, for example, CRAY C90(8)-T3D(242). Table 1 shows timings in each of the individual computers of a version of the code with a resolution of 5° longitude by 4° latitude in the horizontal and nine layers in the vertical. The "x" in Table 1 corresponds to the time required by AGCM/Dynamics to simulate one day for a particular optimization of the code in the CRAY C90 system. We will neglect all overheads due to the distributed system, a hypothesis that will be discussed later in this article.

The expressions for speedup on the CRAY C90(8)-T3D(242) metacomputer with respect to each individual computer are given by the following equation:

$$S_{C90} = \frac{x + 17}{x + 7}; \quad S_{T3D} = \frac{35}{x + 7} \qquad (1)$$

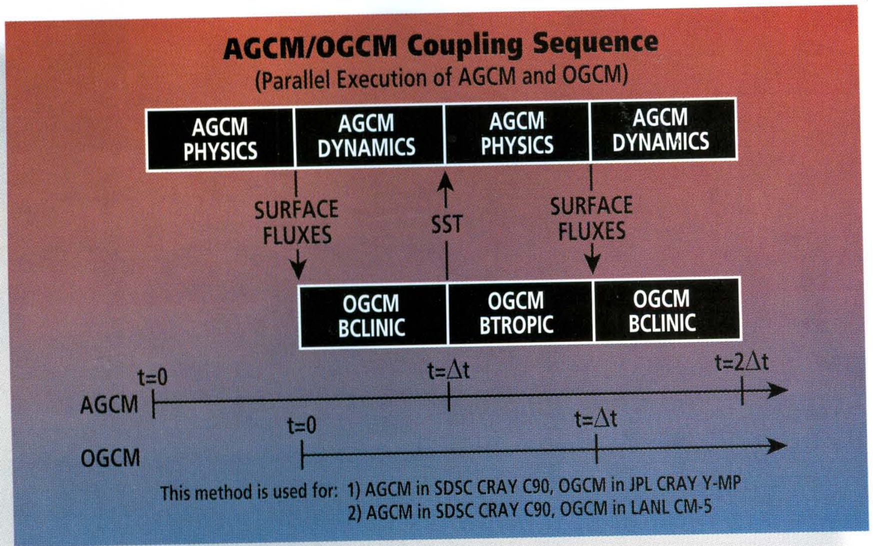

Figure 1. Schematic of the execution sequence of the task decomposed coupled atmosphere-ocean GCM.

We can reasonably expect that using the C90(8)-T3D(242) metacomputer will speed up the code with respect to any one of the individual computers. Firstly, $S_{C90} > 1$ for all values of x. Secondly, $S_{T3D} > 1$ if $x < 29$ s, which is possible without a major restructuring of the code (Table 2 shows that the current version of AGCM/Dynamics requires 29 s per simulated day in one processor of the CRAY C90 system).

Achieving superlinear speedup with respect to individual computers by using the CRAY C90(8)-T3D(242) metacomputer is not possible without restructuring the code. $S_{C90} > 2$ requires that AGCM/Dynamics be executed in less than 3 s per simulated day, while the smallest possible value for x with the current version of the code is about 29 s/8 = 3.625 s. $S_{T3D} > 2$ requires that AGCM/Dynamics be executed in eight processors of the C90(8) in less than 11 s, while the minimum timing we have obtained with this machine is 17 s per simulated day. Improving this performance requires restructuring the code so that it can be more effectively microtasked. Note that these conclusions depend on the model's resolution, and that a finer grid probably would allow for superlinear speedup of the code with respect to at least one of the individual computers as the ratio between communication and computation decreases.

Speedup of execution of the coupled GCM

When run on a single node, the AGCM and the OGCM codes execute sequentially and exchange information corresponding to the air-sea interface. The AGCM is first integrated for a fixed period of time and then transfers the time-averaged surface wind stress, heat and water fluxes to the OGCM. This component is then integrated for the same period of time and transfers the sea surface temperature to the AGCM. The data transfers, including the interpolations required by differences in grid resolution between model components, are performed by a suite of coupling routines.

Our OGCM also has two major components: 1) OGCM/Baroclinic determines the deviations from the vertically averaged velocity, temperature and salinity fields, and 2) OGCM/Barotropic determines the vertically averaged distributions of those fields.[3,4] The coupled atmosphere-ocean GCM, therefore, can be decomposed into four components. Because AGCM/Dynamics does not exchange data with the OGCM, these components can run in parallel. Further, AGCM/Physics can start as soon as OGCM/Baroclinic completes its calculation, because this includes the sea surface temperature, and can run in parallel with OGCM/Barotropic. Figure 1 is a schematic of this running strategy, under the assumption that overheads due to distribution can be neglected and model components running in parallel are perfectly balanced.

Distribution overheads and load imbalances among model components, however, cannot be neglected a priori. Consider the execution of the coupled GCM code using two processors of the CRAY C90 at SDSC in dedicated mode. Figure 2 shows a timeline for one simulated hour of computation (Carl Scarbnick and Gary Hanyzewski, personal communication). The times in Figure 2 correspond to the AGCM with a resolution of 5° longitude by

Figure 2. Timeline for one simulated hour of dedicated execution of the coupled atmosphere-ocean GCM on the SDSC CRAY C90 system using PVM for intertask communications. Time is given in seconds. Data transferred to the OGCM includes surface heat and water fluxes and surface wind stress. Data transferred to the AGCM includes sea surface temperature.

AGCM Processor	Communication to OGCM	AGCM/Dynamics	Wait	Communication from OGCM	AGCM/Physics
Elapsed time (seconds)	0.8	1.0	3.0	0.05	2.1

Time →

OGCM Processor	Communication to AGCM	OGCM/Baroclinic	Communication from OGCM	AGCM/Physics	Wait
Elapsed time (seconds)	0.8	4.0	1.05	0.95	1.15

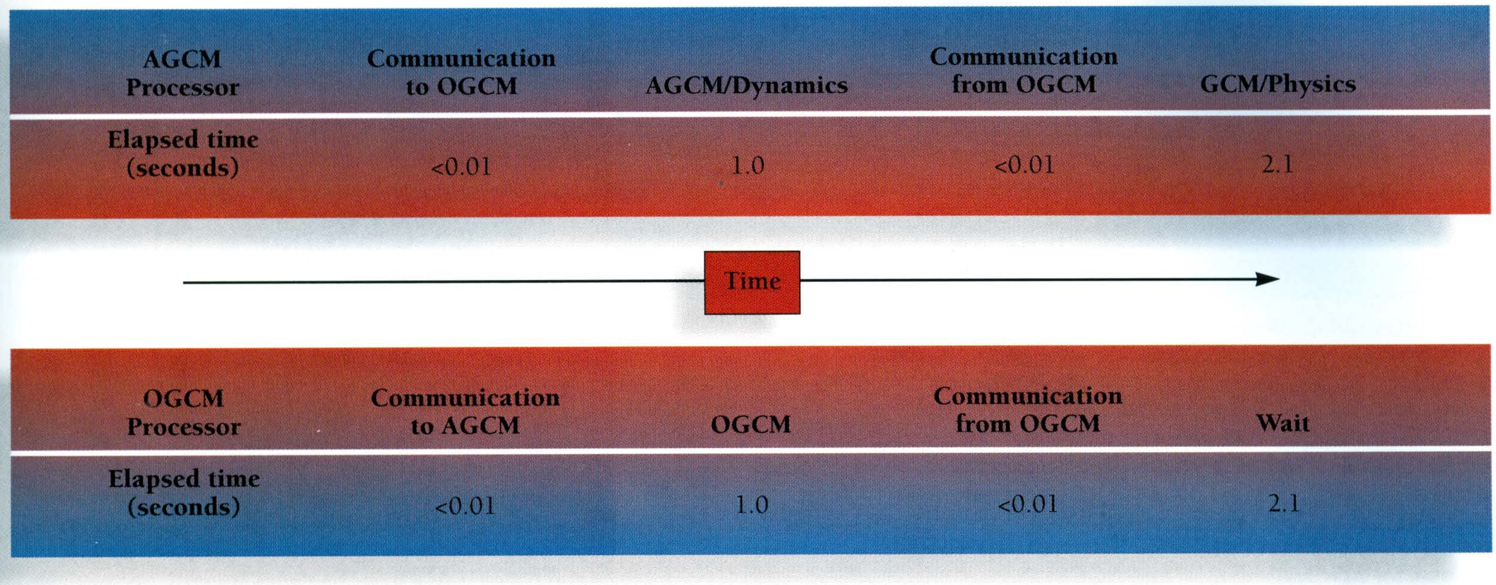

Figure 3. Timeline for one simulated hour of dedicated execution of the coupled atmosphere-ocean GCM on the SDSC CRAY C90 system (AGCM) and Caltech Intel Paragon system (OGCM) using EXPRESS for intertask communications. Time is given in seconds.

4° latitude and nine layers, coupled to the Tropical Pacific version of the OGCM, with a resolution of 1° longitude by 1/3° latitude between 10° S and 10° N, increasing gradually polewards, and 27 layers. The AGCM and OGCM exchange information each simulated hour. Intertask communication was performed using PVM.

According to Figure 2, when run sequentially on a single processor the code requires the following equation to complete one hour of simulation:

$$T^{(s)} = 8.05 \text{ s.}$$

Using two processors of the CRAY C90, which allows for partial overlap in the execution of AGCM/Dynamics and OGCM/Baroclinic as well as AGCM/Physics and OGCM/Barotropic, resulted in a reduction of the time required to simulate one hour to

$$T = 6.95 \text{ s,}$$

which implies a speedup of

$$S = 8.05/6.95 = 1.158$$

over the sequential code. The minimum execution time, if all communication time and wait time were masked by computation, would be that required by the two slowest components (AGCM/Physics and OGCM/Baroclinic). That is,

$$\min\{T\} = 6.1 \text{ s,}$$

which implies that the maximum possible speedup for this configuration of the code in two processors of the CRAY C90 system is

$$S_{max} = 8.05/6.1 = 1.320.$$

A slightly higher speedup was obtained by the same SDSC researchers running the coupled GCM in a CRAY C90(1)-Paragon(144) metacomputer, with the CRAY C90 system at SDSC connected by the Gigabit Network to the Intel Paragon at Caltech. Interprocessor and intertask communications were performed using NX and EXPRESS, respectively.

The OGCM running on 144 nodes of the Intel Paragon system requires 1.0 s to simulate one hour, while the AGCM requires 3.6 s. On one processor of the CRAY C90 system the OGCM requires 1.5 s to simulate one hour, while the AGCM requires 3.1 s (see Fig. 2). Thus, the sequential execution time on both machines is the same: 4.6 s per simulated hour. With the code executing as shown in Figure 3, the time required to complete one simulated hour was measured to be 3.1 s. According to Figure 3, the intertask communication times in this experiment were much smaller than those measured within the CRAY C90 system using PVM (see Figure 2), suggesting that EXPRESS imposes a much lower overhead cost than PVM. Also, the 1 s of OGCM execution was completely masked by the execution of AGCM/Dynamics. The model speedup in this heterogeneous environment is given by

$$S_{C90} = S_{Paragon} = 4.6/3.1 = 1.483.$$

So far, we have considered a straightforward parallelization of model components and obtained only modest speedups primarily due to distribution overheads and load imbalances among model components. The availability of multiple nodes in the computer environment allows for the application of a scheme specifically designed for the overlapping of communications with computations in coupled GCMs. In this scheme, or I/O decomposition,[5] the exchange of data between different model tasks is carried out in subdomains consisting of latitude bands from the Earth's surface to the top of the atmosphere. The transmission of the data produced by each task for one subdomain to the other tasks can then be masked by computations in another subdomain. As the end result, the wallclock time for execution of the coupled GCM is reduced to that required by the slowest model component. If this scheme is applied to the coupled GCM divided in two tasks (AGCM and OGCM) running on the

	CRAY C90 (1 processor)	CRAY T3D (121 processors)	CRAY T3D (242 processors)
AGCM/Dynamics	29*	36	29
AGCM/Physics	54	12*	7
OGCM	76	29*	23
Total	159	77	59

*denotes optimal distribution of model components for fastest overall simulation of one day.

Table 2. Timings (seconds/simulated day) of the coupled atmosphere-ocean GCM.

CRAY C90 system, the time to simulate one hour could be reduced to that of the OGCM, which represents a speedup of 1.636 over the sequential code.

Next we consider the coupled GCM running on a CRAY C90(1)-T3D(242) metacomputer, in which the CRAY T3D system is divided into two partitions of 121 processors each. Table 2 gives the timings for AGCM/Dynamics and AGCM/Physics on the individual computers. We expected AGCM/Dynamics to run on the CRAY C90 system and AGCM/Physics on the CRAY T3D system. Therefore, we chose the highest resolution version of the OGCM that can execute on 121 processors of the CRAY T3D system in a wall-clock time corresponding to that for AGCM/Dynamics on one processor of the CRAY C90 system.

According to Table 2, the time required for sequential execution of one simulated day is 159 s on one processor of the CRAY C90 system and 59 s on 242 processors of the CRAY T3D system. If the I/O decomposition scheme is used, and the model components are distributed as shown by the asterisks in Table 2, the time required to complete one day of simulation is 29 seconds. The corresponding speedups are given by

$$S_{C90} = 159/29 = 5.483; \quad S_{T3D} = 59/29 = 2.034,$$

both of which are superlinear. The next section of this article discusses the possibility of using the I/O decomposition method for overlapping communication with computations.

Masking communications with computation

The estimates presented earlier in this article demonstrate the potential for substantial speedups under the assumption that the different tasks of the AGCM or the coupled GCM can be run in parallel so that the timing of the full model is reduced to that of its slowest component. This running strategy is possible by using I/O decomposition. To explore whether the conditions for the I/O decomposition can be met by the distributed model, we designed a version of the model that retains all the major control structures of the full model but excludes all computations and external file I/O. We refer to this version as the NGCM. The NGCM timing, therefore, represents the overhead of the distributed system.

We have run an NGCM that consists of AGCM/Physics and AGCM/Dynamics, with a resolution of 5° longitude by 4° latitude and nine layers, and *m* I/O subdomains. This NGCM was run in two environments, one local and one distributed. In the local environment, we used two processors of the SDSC CRAY C90 system. In the distributed environment, we used one processor of the CRAY Y-MP system at JPL and one processor of the Intel Delta computer at Caltech, which are connected by the Gigabit Network; the AGCM/Dynamics and AGCM/Physics components of the NGCM ran on the CRAY Y-MP and the Delta systems, respectively. Intercomponent communication was performed by using utilities provided by EXPRESS. If the I/O decomposition is designed such that data transfers are sequential and data exchanges between model components occur at the end of every simulated hour (the time-step for the AGCM/Physics calculation), the method[5] requires that

$$[\text{NGCM}]_m < \max\{[\text{AGCM/Phys}], [\text{AGCM/Dyn}]\}, \quad (6)$$

where square brackets denote the wall-clock times required to run the corresponding components for a simulated period of time.

Table 3 gives the values of [NGCM] for one simulated day and several values of *m* obtained in dedicated ($[\text{NGCM}]^D$) and production ($[\text{NGCM}]^P$) mode. In both the local and distributed environments, the values of $[\text{NGCM}]^{D_m}$ are generally much smaller than the 54 seconds per simulated day required to execute the AGCM/Physics in one processor of the CRAY C90 system for $m < 4$. Thus, in dedicated mode it would be possible to use the I/O decomposition method and completely mask communication with computation. In production mode, the values of $[\text{NGCM}]^P$ are larger than that of [AGCM/Physics] for $m > 3$. In the distributed environment, the wall-clock times are comparable to those obtained in the production environment on the CRAY C90 system at SDSC, although in this case, the model components are running in computers approximately 10 kilometers apart. For $m = 5$, the minimum number of I/O subdomains required,[5] $[\text{NGCM}]^P$ is 45.2 s, which is less than the 61 s required to execute AGCM/Dynamics on the CRAY Y-MP system at JPL. Thus, it also would be possible in this production environment to completely mask communication with computation by using the I/O decomposition method.

Table 3. Wall-clock time (seconds/simulated day) for the NGCM on a dedicated ($[\text{NGCM}]^D$) and production ($[\text{NGCM}]^P$) CRAY Y-MP system, using different numbers of I/O subdomains, *m*, for AGCM/Physics calculation.

m	$[\text{NGCM}]^D$ CRAY C90-C90	$[\text{NGCM}]^P$ CRAY C90-C90	$[\text{NGCM}]^P$ CRAY Y-MP-Delta
1	10.1	38.4	31.4
3	10.4	40.8	37.6
4	10.4	96.0	45.8
5	10.5	88.0	45.2
8	–	–	68.5

Gigabit network performance

We also have used the NGCM to analyze the performance of the Gigabit Network of the CASA Testbed. Sites in the testbed are linked via multiple OC-3 SONET channels provided by US

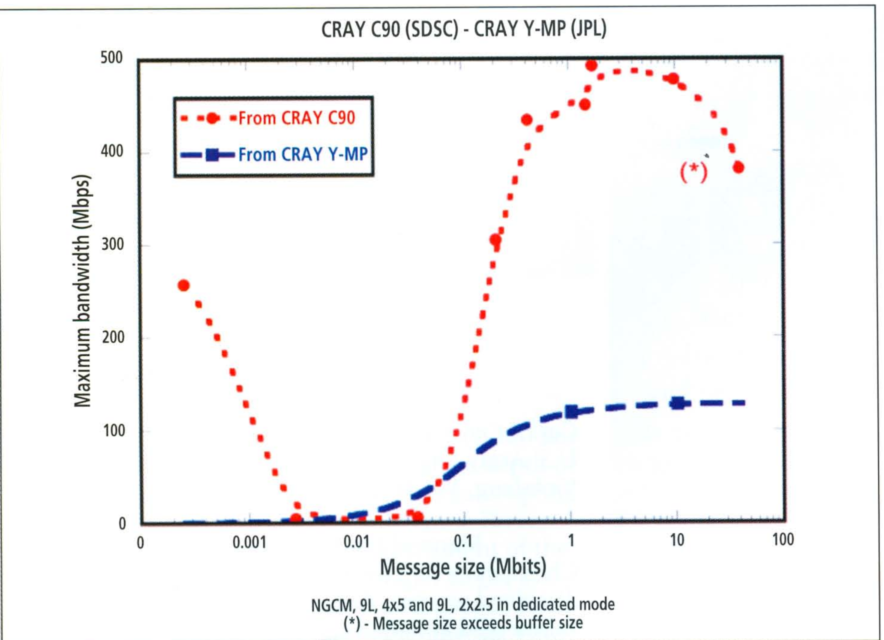

Figure 4. The effective bandwidth (Mbits/s) as a function of message size for messages sent by the 5° longitude by 4° latitude, 9-layer and 2.5° longitude by 2° latitude, 9-layer versions of the NGCM running distributed between the JPL CRAY Y-MP and SDSC CRAY C90 systems.

West, MCI, and Pacific Bell. Gateways developed at LANL link up to seven OC-3 channels to HIPPI local area networks (LANs) at each site. By striping the message traffic across the seven OC-3 channels, 800 Mbit/s bandwidths are achieved between HIPPI LANs.

The bandwidth obtained as a function of message size for distributed NGCM simulations between the CRAY C90 system at SDSC and the CRAY Y-MP system at JPL is given in Figure 4. The "effective" bandwidth in this figure is defined as the ratio between the message length and the time required for the sending process to complete a blocking write to the receiving process. The data points are for all the different size messages sent by the 5° longitude by 4° latitude, and by the 2.5° longitude by 2° latitude versions of the NGCM, both with nine layers.

In this experiment, data transfers originating from the CRAY Y-MP system were noticeably slower than those originating from the CRAY C90 system; therefore the data was plotted in two separate curves, one for transfers originating from the CRAY C90 system, the other for transfers originating from the CRAY Y-MP system. The peak effective bandwidth obtained when sending data from the CRAY C90 system is approximately 480 Mbits/s for messages of approximately 16 Mbits. For transfers originating in the CRAY Y-MP system, the peak effective bandwidth was about 128 Mbits/s.

Conclusions

This work shows that systems comprising computers that have relatively few powerful processors and shared memories and computers that have large numbers of less powerful processors and distributed memories can provide important advantages to climate modelers. Such an arrangement can produce superlinear speedups in at least some climate modeling applications, despite the added overhead due to latency and communications. In any event, climate models that deal with major problems of interest to society at large, and that include not only computational aspects, but also aspects of data management and output visualization, will provide ideal testbeds for advances in science and technology.

Because the four supercomputer centers sponsored by the NSF have coordinated a "metacenter," and because there seems to be support for gigabit networks, the availability of a "metacomputer" appears to be certain in the not-too-distant future. The running of codes in support of research on problems of vital importance will require a great deal of agreement and coordination. Nevertheless, the interest in distributed computing suggests that this agreement and coordination, which has been achieved by CASA on an experimental basis, might also be possible in production mode.

Acknowledgments

The authors are grateful to the San Diego Supercomputer Center for providing dedicated CRAY time for the performance of the experiments described in this article. A similar acknowledgment is due to the NASA/Caltech Jet Propulsion Laboratory. This work was supported by NSF and ARPA under cooperative agreement NCR8919038 with the Corporation for National Research Initiatives and NASA under Grant NAG 5-2224.

About the authors

Carlos R. Mechoso is a professor in the Department of Atmospheric Sciences at the University of California, Los Angeles. He is a member of the CASA Testbed Project Executive Committee. He received a Ph.D. degree in geophysical fluid dynamics from Princeton University in 1978.

John D. Farrara is a postdoctoral researcher in the Department of Atmospheric Sciences at the University of California, Los Angeles. He received a Ph.D. degree in Atmospheric Sciences from UCLA in 1989.

Joseph A. Spahr is a senior computer programmer in the Department of Atmospheric Sciences at the University of California, Los Angeles. He received a B.S. degree in Atmospheric Sciences from Cornell University in 1974.

References

1. Suarez, M. J., A. Arakawa, and D. A. Randall, "The parameterization of the planetary boundary layer in the UCLA general circulation model: Formulation and results," *Monthly Weather Review*, Vol. 111, pp. 2224-2243, 1983.
2. Cox, M. D., "A primitive equation, three-dimensional model of the ocean," GFDL Ocean Group Technical Report No. 1, 143 pp., 1984.
3. Ma, C.-C., C. R. Mechoso, A. Arakawa, and J. D. Farrara, "Sensitivity of a coupled ocean-atmosphere model to physical parameterizations," *Journal of Climate*, Vol. 7, pp. 1883-1896, 1994.
4. Mechoso, C. R., C.-C. Ma, J. D. Farrara, J. A. Spahr, and R. W. Moore, "Parallelization and distribution of a coupled atmosphere-ocean general circulation model," *Monthly Weather Review*, Vol. 121, pp. 2062-2076, 1993.
5. Mechoso, C. R., J. D. Farrara, and J. A. Spahr, "Achieving superlinear speedup on a heterogeneous, distributed system," *Parallel & Distributed Technology*, Vol. 2, 57-61, Summer 1994.
6. Bryan, K., "A numerical method for the study of the circulation of the world ocean," *Journal of Computational Physics*, Vol. 4, pp. 347-376, 1969.

environm

Chesapeake Bay, the United States' largest and most productive estuary, is plagued by a host of problems that accompany agricultural and industrial development and population growth along its shores and headwaters. Recently, a mathematical model was used to examine pollution-reduction strategies for the bay. The model study advanced the state of the art in environmental modeling in three ways: coupling a water-quality model to a three-dimensional (3-D) time-variable hydrodynamic model; coupling the model of the water column with a fully predictive sediment oxygen demand and nutrient flux model; and continuous, multi-year application of the model on an intertidal timescale.

Carl F. Cerco, Billy H. Johnson, Mark Dortch, Ross Hall
U.S. Army Engineer Waterways Experiment Station
Vicksburg, Mississippi

Many of the problems associated with the Chesapeake Bay, including a perceived decrease in dissolved oxygen, decline in fisheries, and loss of submerged aquatic vegetation, are related to the release of excess nutrients into bay waters. In 1987, the chief executives of Pennsylvania, Maryland, Virginia, and the District of Columbia pledged to reduce by 40 percent the amounts of nitrogen and phosphorus entering the mainstem of Chesapeake Bay. Part of this agreement called for reevaluation of the nutrient-reduction goal following a study of the feasibility and effects of the planned cleanup. An environmental model of Chesapeake Bay was the principal tool in the reevaluation effort.

Negotiations between the U.S. Environmental Protection Agency Chesapeake Bay Program Office, the U.S. Army Engineer District (Baltimore), and the U.S. Army Engineer Waterways Experiment Station (WES) led to the selection of WES to conduct the environmental modeling of the bay. Researchers assembled a package of interactive models to meet the requirements of the study. This package included a 3-D hydrodynamic model,[1] a 3-D water-quality model,[2] and a predictive model of sediment-water interactions.[3]

The Chesapeake Bay system (Figure 1) consists of the mainstem bay, five major western-shore tributaries, and a host of lesser tributaries

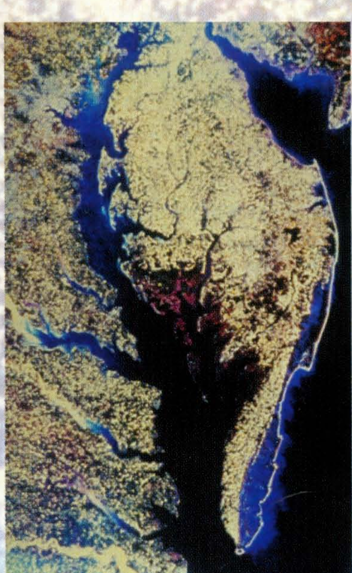

Figure 1. Chesapeake Bay.

ental modeling of Chesapeake Bay

and embayments. The mainstem is roughly 300 km long and 8 to 48 km wide, with an average depth of 8 m. A deep trench with depths to 50 m runs up the center of the mainstem. The Susquehanna River, which empties into the northernmost extent of the bay, is the primary source of freshwater to the system (approximately 62 percent of total gauged freshwater flow). Other major sources include the Potomac and James Rivers located along the western shore. The bay is a classic example of a partially mixed estuary in which long-term average circulation is upstream along the bottom and downstream near the surface, although local and distant meteorological events frequently alter this pattern. Major urban centers along the bay and tributaries include Norfolk and Richmond, Virginia, Washington, D.C., and Baltimore, Maryland.

The hydrodynamic model

The hydrodynamic model, CH3D-WES (Curvilinear Hydrodynamics in Three Dimensions-Waterways Experiment Station), is an extensively modified version of a model originally developed by Y. Peter Sheng.[4] CH3D-WES makes hydrodynamic computations on a curvilinear or boundary-fitted platform grid. The boundary-fitted-coordinate feature of the model provides enhancement to fit the deep navigation channel and irregular shoreline of the bay and permits adoption of an accurate and economical grid schematization. The Chesapeake Bay grid contains 729 surface cells and a maximum of 15 vertical layers, resulting in 4073 computational cells. Grid resolution is 1.52 m vertical, approximately 10 km longitudinal, and 3 km lateral. Tides, wind, density effects, freshwater inflows, turbulence, and the influence of the Earth's rotation, all of which affect baywide circulation and vertical mixing, are some of the physical processes modeled.

To verify the hydrodynamic model, researchers compared it with extensive data collected during three short-term periods. A key simulation was an autumn 1983 wind-mixing event. Wind-generated turbulence and rapid cooling of the surface waters diminished the vertical density gradient, resulting in destratification of the bay. Nearly uniform vertical salinity distribution occurred in the central bay at the peak of the wind-mixing event (Figure 2).

The researchers applied the environmental model to the years 1984 through 1986. They simulated this period in three one-year hydrodynamic model runs. The integration timestep was five minutes, and each run consumed roughly 10 CPU hours on a CRAY Y-MP system installed at the U.S. Department of Defense HPC Major Shared Resource Center within the Information Technology Laboratory of WES. The hydrodynamic model was verified further by comparing predicted and observed tides and salinity during the year-long simulations. In the fall of 1985, a major storm resulted in approximately 200-year flood flows in the James River. Near the peak of the storm, the numerical model

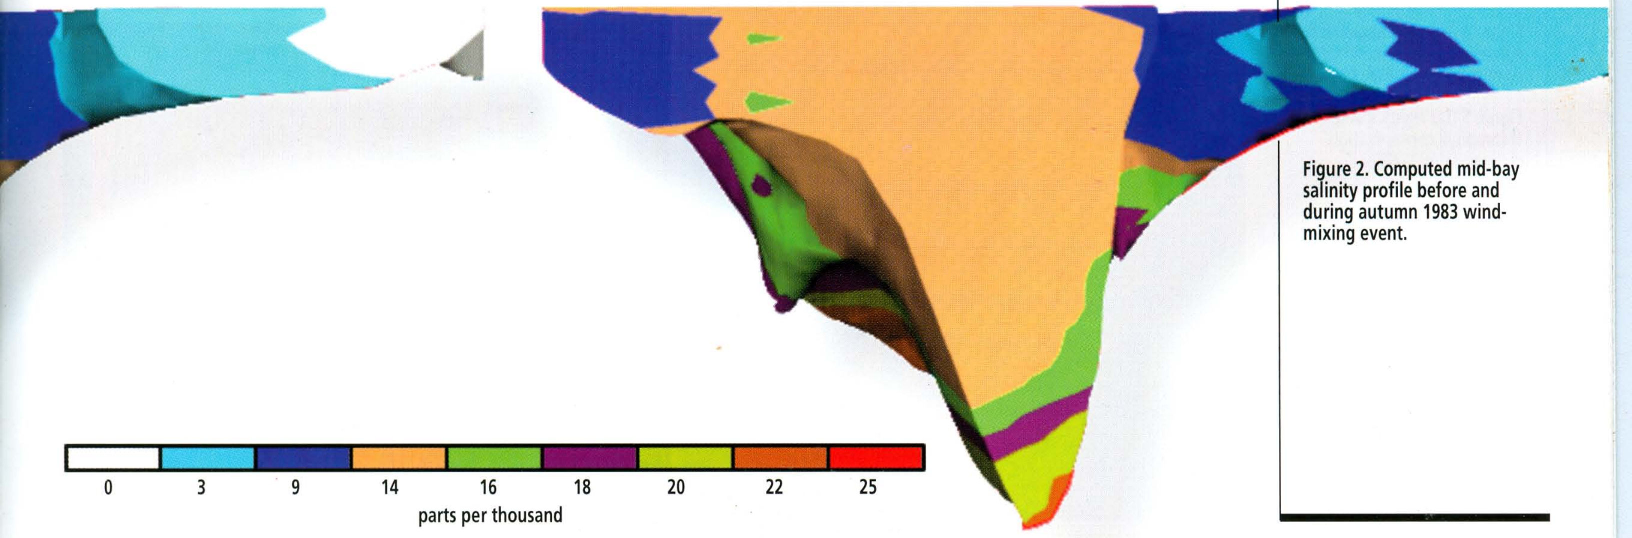

Figure 2. Computed mid-bay salinity profile before and during autumn 1983 wind-mixing event.

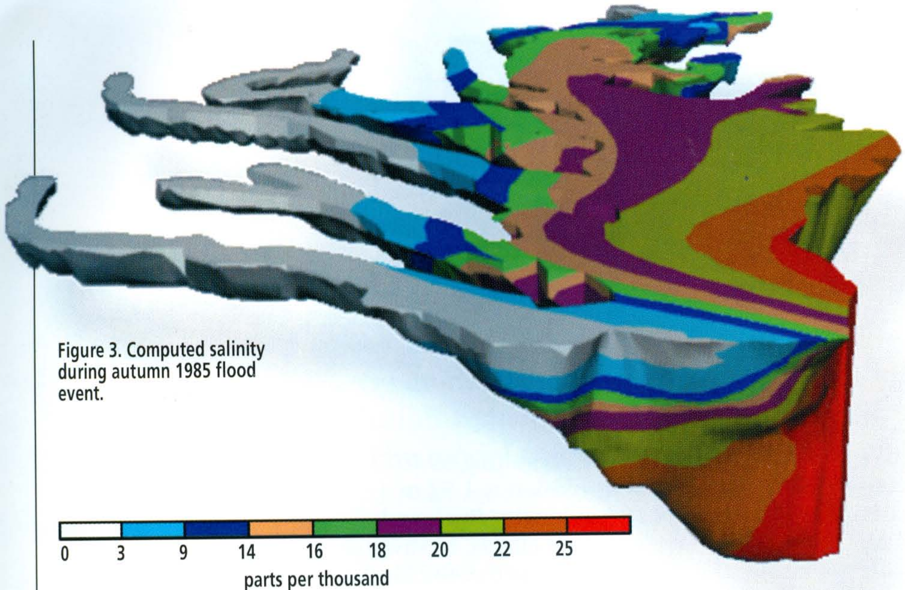

Figure 3. Computed salinity during autumn 1985 flood event.

computed virtually salt-free surface waters for the entire James River (Figure 3). Field data collected just before and after the peak flow substantiated these results.

The hydrodynamic model/water-quality model interface

Water quality is modeled by an integrated-compartment model, CE-QUAL-ICM, developed for the study. Researchers selected the compartment structure to allow maximum flexibility for adaptation of the model to alternate hydrodynamic models. For the Chesapeake Bay application, transport in the water-quality model was driven by output from the CH3D-WES model. CE-QUAL-ICM operates on an unstructured grid in which compartments correspond to cells in the structured CH3D-WES grid. A major effort in this study was the development of the linkage between the hydrodynamic and water-quality models.[5] The linkage involved temporal averaging and storage of hydrodynamic information, as well as mapping of the structured and unstructured grids.

Because storage of hydrodynamic information at every timestep consumed unacceptable volumes of disk space, the interfacing procedure included temporal averaging of hydrodynamics prior to storage. The averaging period was equivalent to a tidal cycle (12.4 hours). Lagrangian averages, rather than simple arithmetic (Eulerian) averages, were computed. The Lagrangian average current is the net displacement of a water parcel divided by the elapsed time. Lagrangian residual currents may differ substantially from arithmetic mean currents due to nonlinear interactions of the tides.

Researchers incorporated interface subroutines into the CH3D-WES code so that intertidal, Lagrangian average, hydrodynamic information was computed and stored while the hydrodynamic model was executing. The major benefit of using intertidal hydrodynamics was the reduction in computer disk storage requirements. One year of intertidal hydrodynamics (in binary form) for the Chesapeake Bay model required 85 Mbytes of disk space. More than 1 Gbyte was required when hydrodynamics data was stored at two-hour intervals. The use of intertidal hydrodynamics also resulted in reduced computational effort, because larger water-quality model timesteps were possible, and hydrodynamic updates were required less frequently.

The water-quality model

The water-quality model incorporates 22 state variables. Kinetic variations among the state variables are described in 80 partial differential equations that employ over 140 parameters. The state variables can be categorized into six groups or cycles: a physical group (salinity, temperature, and inorganic suspended solids), the carbon cycle, the nitrogen cycle, the phosphorus cycle, the silica cycle, and the dissolved oxygen cycle.

The water-quality model is directly coupled to a predictive benthic sediment model. Particle deposition, temperature, nutrient, and dissolved oxygen concentrations are passed from the water-quality model to the sediment model. The sediment model computes sediment-water fluxes of dissolved nutrients and oxygen based on computed decay and on conditions in the sediments and water. The computed sediment-water

Figure 4 (right). Computed dissolved oxygen concentration during summer 1986.

Figure 5 (far right). Computed chlorophyll concentration during spring 1986.

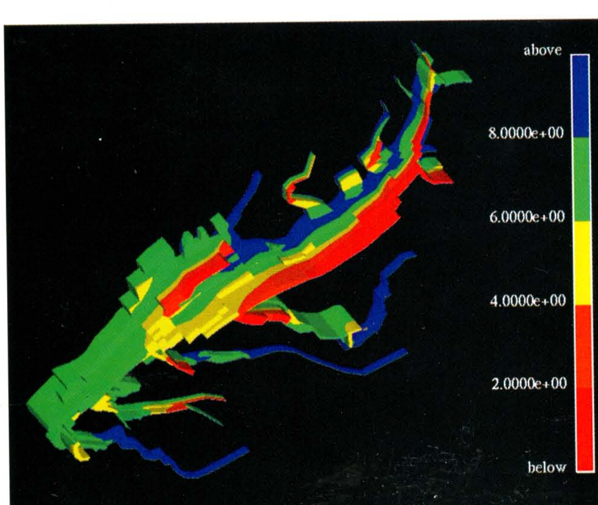

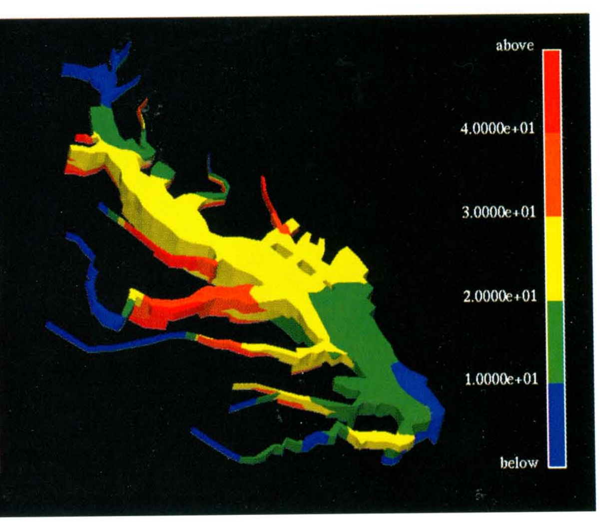

fluxes are passed to the water-quality model and incorporated into appropriate mass balances and kinetic reactions.

Simulation of eutrophication processes

One focus of the study was representation of hypoxic (low dissolved oxygen) conditions that are especially harmful to fish and other valued living resources in the bay. Bottom-water hypoxia occurs at recurrent, predictable intervals. The onset is in late May when spring warming enhances respiration in benthic sediments. Decay of organic matter deposited in spring and in previous years removes oxygen from bottom water. Density stratification prevents mixing of oxygenated surface water downward. Low-oxygen conditions continue through the summer (Figure 4), maintained by respiration in bottom water. In mid-September, autumn winds end the hypoxic period by mixing surface water down to the bottom.

Dissolved-oxygen computation required simulation of complex interactions among the physical, chemical, and biological environments. One key process leading to hypoxic conditions is the recurrent spring phytoplankton bloom (Figure 5). The bloom usually commences in February, reaches a maximum in April, and ends precipitously in May. High chlorophyll concentrations throughout the water column characterize the bloom. A spring peak in carbon deposition to sediments occurs simultaneously with the algal bloom. The deposition and decay of fresh organic matter contribute to oxygen demand during the onset of the hypoxic period.

A subtle—and potentially more important—link is through a nutrient-trapping mechanism. Nutrients in spring runoff are taken up by algae during the bloom. Predation and algal mortality result in the transfer of nutrients, in particulate organic form, to benthic sediments. In the summer, the nutrients are mineralized in the sediments and released to the water column. Nutrients released from the sediments support summer algal production. Carbon produced by algae settles to bottom waters, decays, and consumes oxygen. Diminished oxygen in bottom water enhances sediment nutrient release, especially of ammonium. The nutrient release continues the cycle of benthic release, algal production, and oxygen consumption.

Additional developments

Following application to 1984 to 1986, the model package was used in a simulation of bay conditions from 1959 to 1988.[6] Each 30-year water-quality simulation consumed 90 CPU hours on a CRAY Y-MP system and was especially valuable in evaluating long-term trends in water quality. The simulations revealed the dominance of physical processes in determining the hypoxic volume of the bay.

The Waterways Experiment Station conducted 30 10-year pollutant-reduction scenarios that required 30 CPU hours on the CRAY Y-MP system. These scenarios confirmed the basic strategy of a 40 percent nutrient load reduction. The model package subsequently was passed to the U.S. Environmental Protection Agency and installed on EPA's Cray Research system in Bay City, Michigan. The model package is available for further examination of pollutant control strategies and for additional environmental research purposes.

Acknowledgments

Development of the Chesapeake Bay Model Package was sponsored by the U.S. Army Engineer District, Baltimore, and the Chesapeake Bay Program Office, U.S. Environmental Protection Agency. This information is published with the permission of the Chief of Engineers.

About the authors

Carl F. Cerco is a research hydrologist in the Environmental Laboratory, U.S. Army Engineer Waterways Experiment Station (WES), Vicksburg, Mississippi. He has been engaged in water-quality modeling since beginning graduate study in 1969. He is a registered professional engineer with a Ph.D. degree in marine science from the College of William and Mary.

Billy H. Johnson is the special assistant for numerical modeling in the Hydraulics Laboratory, WES. He is a registered professional engineer with a Ph.D. degree in general engineering from Mississippi State University.

Mark Dortch is a supervisory research hydraulic engineer and chief of the Water Quality and Contaminant Modeling Branch, Environmental Laboratory, WES. Since arriving at WES in 1972, he has conducted a variety of hydraulic and water-quality model studies. He is a registered professional engineer with a Ph.D. degree in civil engineering from Colorado State University.

Ross Hall is a research limnologist in the Environmental Laboratory, WES. He has over 20 years of experience in the study of water quality. He holds an MS degree in limnology from Oklahoma State University.

References

1. Johnson, B., K. Kim, R. Heath, B. Hsieh, and L. Butler, "Validation of a Three-Dimensional Hydrodynamic Model of Chesapeake Bay," *Journal of Hydraulic Engineering*, Vol. 199, No. 1, pp. 2-20, 1993.
2. Cerco, C. and T. Cole, "Three-Dimensional Eutrophication Model of Chesapeake Bay," *Journal of Environmental Engineering*, Vol. 119, No. 6, pp. 1006-1025, 1993.
3. DiToro, D. and J. Fitzpatrick, "Chesapeake Bay Sediment Flux Model," Contract Report EL-93-2, U.S. Army Engineer Waterways Experiment Station, Vicksburg, Mississippi, 1993.
4. Sheng, Y., "A Three-Dimensional Mathematical Model of Coastal, Estuarine, and Lake Currents Using Boundary Fitted Grid," Report 585, Aeronautical Research Associates of Princeton (A.R.A.P) Group of Titan Systems, Princeton, New Jersey, 1986.
5. Dortch, M., R. Chapman, and S. Abt, "Application of Three-Dimensional Lagrangian Residual Transport," *Journal of Hydraulic Engineering*, Vol. 118, No. 6, pp. 831-848, 1992.
6. Cerco, C., "Simulation of Long-term Trends in Chesapeake Bay Eutrophication," *Journal of Environmental Engineering*, in press.

Supercomputer modeling of groundwater contamination by nonaqueous phase liquids

Alex S. Mayer, Michigan Technological University
Houghton, Michigan

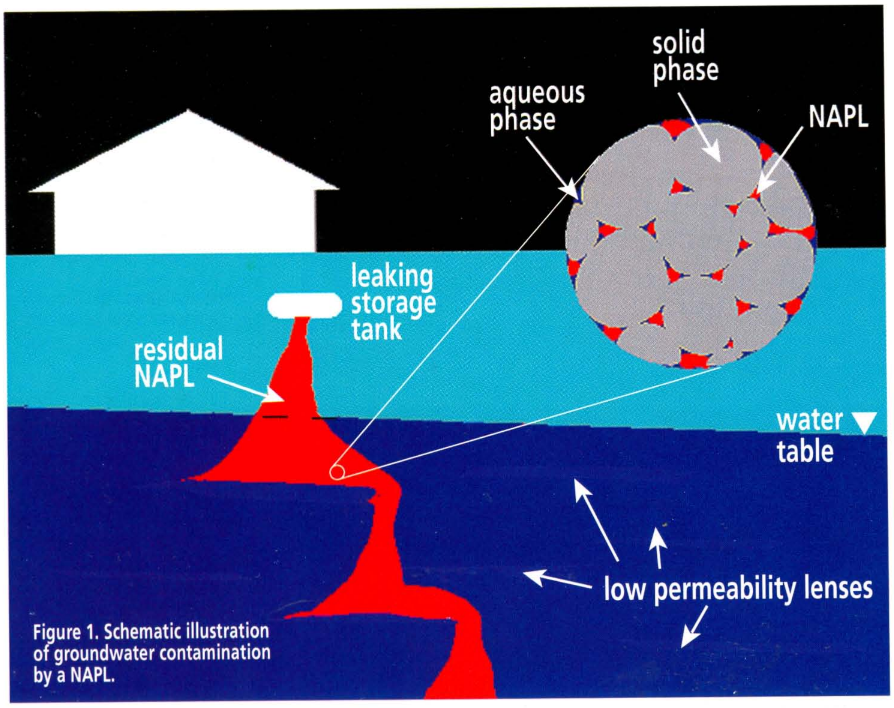

Figure 1. Schematic illustration of groundwater contamination by a NAPL.

Groundwater is a critical part of the world's water resources. In the United States, half of the population depends on groundwater for drinking water; this precious resource is threatened by contamination. Many groundwater contamination scenarios involve the release and subsequent migration of immiscible, organic liquids into the subsurface. These liquids, referred to as nonaqueous phase liquids, or NAPLs, consist of a wide range of hazardous compounds. Petroleum fuels and chlorinated solvents are typical examples of NAPLs. Although NAPLs migrate as a phase that is separate from the groundwater, NAPL compounds can transfer into the gas, solid, or aqueous (groundwater) phases. Figure 1 is a schematic illustration of the contamination of a groundwater aquifer by a dense NAPL.

NAPL dissolution, the exchange of NAPL compounds with the aqueous phase, is an important route of interphase mass exchange. NAPL compounds are frequently quite toxic when ingested in drinking water. A liter of trichloroethylene, a chlorinated solvent, can contaminate up to 300 million liters of groundwater at the recommended maximum concentration of trichloroethylene in drinking water (0.005 milligrams per liter).[1] Unfortunately, the NAPL dissolution process has not been well characterized. It has not been established whether NAPL dissolution is a mass-transfer-limited process or may be assumed to be an equilibrium process under typical subsurface conditions. The question of mass transfer limitations versus equilibrium has implications for the amount and rate of groundwater contamination by NAPLs and also how quickly NAPLs can be removed from the subsurface by groundwater remediation technologies. Another missing link in understanding NAPL dissolution is the behavior of this process under heterogeneous conditions, that is, conditions where the physical and chemical properties of the porous media change from point to point in the subsurface.

Researchers have investigated some of these issues in the laboratory, but most of the experimental studies have been conducted in small-scale, homogeneous systems. Mathematical models, on the other hand, can be used to investigate the behavior of NAPL dissolution over a wider range of subsurface conditions. While mathematical models are useful tools, sophisticated algorithms are required to generate the large number of computationally intensive simulations required to resolve some of the preceding issues. This article describes the use of mathematical models on Cray Research supercomputers to provide at least a partial understanding of the NAPL dissolution process. The research discussed in this paper was conducted in collaboration with C. T. Miller of the University of North Carolina at Chapel Hill.

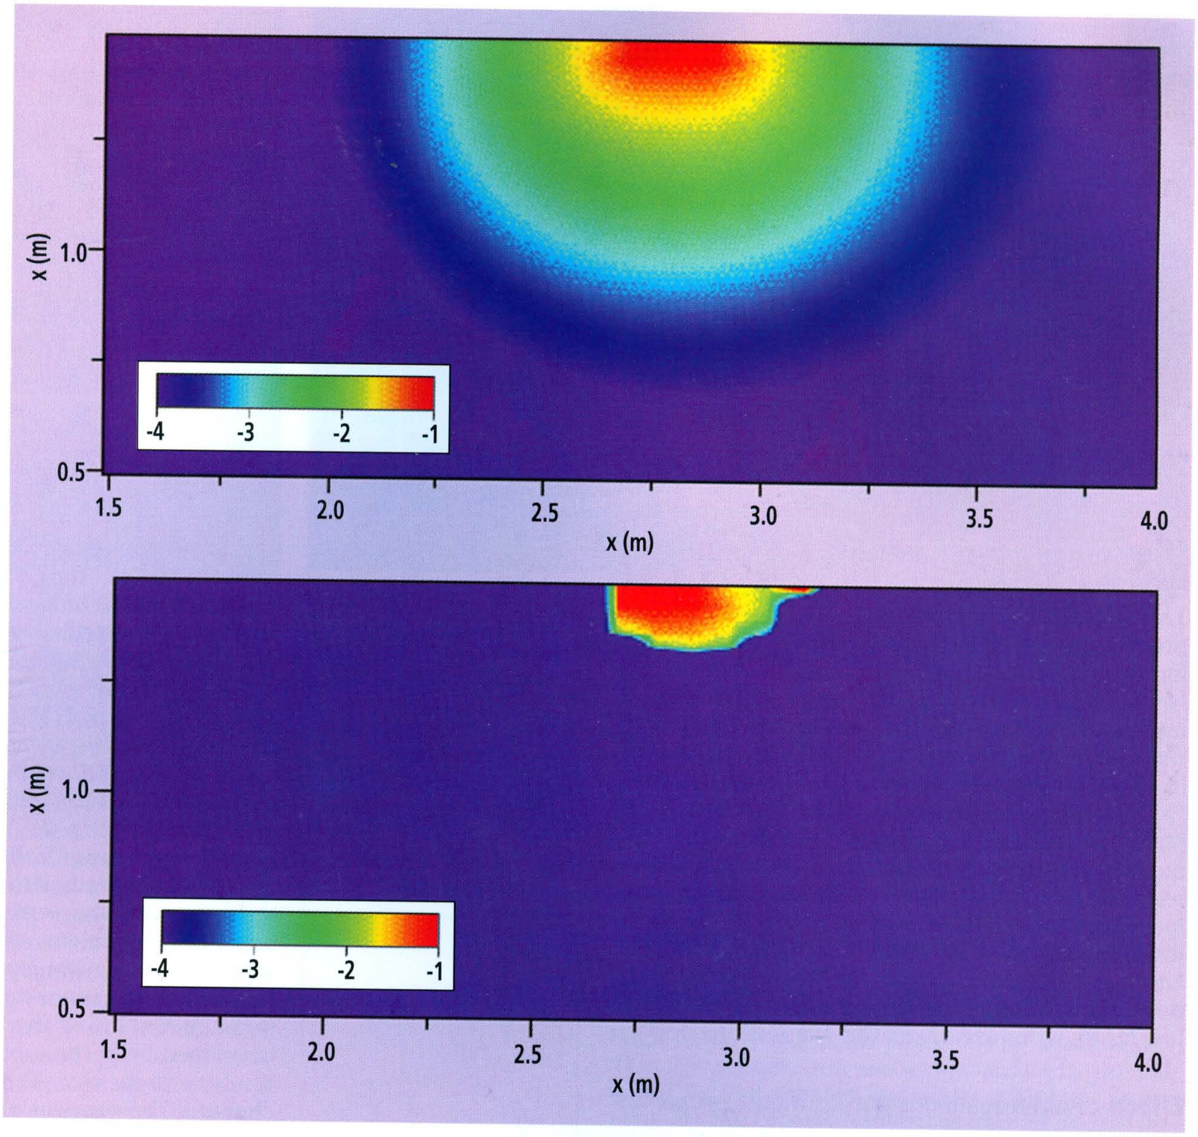

Figures 2a and 2b. Simulations of NAPL dissolution into a homogeneous system at about (a, top) one day and (b, above) five years: color bar indicates logarithm of NAPL volumetric fraction.

Modeling NAPL dissolution

Although several groundwater flow and transport codes are available in the public and private domains, NAPL migration and dissolution cannot be described by conventional groundwater codes. A new, two-dimensional code was developed to simulate multiphase flow, interphase mass exchange, and transport of the dissolved NAPL in the aqueous phase. The code predicts aqueous phase and NAPL pressures, volumetric fractions of each phase, and aqueous phase concentrations as a function of space and time. The volumetric fraction of a fluid phase is the fraction that the fluid phase occupies in a unit volume of porous media.

Each simulation involves the introduction of a pulse of a dense NAPL below the water table, migration of the NAPL until it reaches a residual, or immobile, state, and dissolution of the NAPL into a flowing aqueous phase. The domain in the simulations was 7.11 m long by 1.47 m deep and consisted of 4000 grid points. A horizontal, aqueous phase pressure gradient was imposed on the domain, resulting in an aqueous phase velocity of 0.1 m/day. This velocity is typical of groundwater aquifers. The simulated NAPL has physical and chemical properties similar to trichloroethylene, while the porous media properties are typical of a medium sand. Figures 2a and 2b show the simulated NAPL volumetric fractions at about one day and five years, respectively, in a vertical section through a hypothetical aquifer. Figure 2a shows the NAPL when it has reached the residual state and occupies the largest area during the simulation. Figure 2b illustrates that the NAPL was partially depleted as it dissolved into the aqueous phase. In this simulation, the NAPL is depleted completely in 23 years, involving about 10,000 time steps. Typical run times on a single-CPU CRAY Y-MP system for a complete simulation would be on the order of several days.

Sensitivity to mass transfer rates

Several researchers have conducted laboratory studies to determine whether NAPL-aqueous phase mass exchange can be assumed to be an equilibrium process, or whether mass transfer limitations are significant. Some of these studies have produced models that relate mass transfer rates to system properties such as aqueous phase velocities and NAPL volumetric fractions. However, the published models exhibit a wide range of quantitative relationships between mass transfer rates and system properties. A series of simulations was conducted to test the sensitivity of the NAPL dissolution process to the magnitude of mass transfer rates, the dependencies of mass transfer models on system properties, and the assumption of equilibrium between phases.

Figure 3 shows the results of these simulations for four mass transfer models[2, 3, 4] and a simulation where the equilibrium assumption was employed. Figure 3 displays the total, or global, mass of NAPL remaining in the system as a function of time. The global NAPL mass is normalized by the initial NAPL mass in the system. Figure 3 shows that, with the exception of the Powers et al. "b" simulation, the mass transfer models gave results which were almost identical to the equilibrium simulation. These results imply that three of the four models produced high mass transfer rates, leading to rapid equilibrium between the NAPL and the aqueous phase. However, the dissolution rate simulated with the Powers et al. "b" model was significantly lower after about 15 years, resulting in longer persistence of the NAPL. The Powers et al. "b" model produced mass transfer rates that are much lower than those obtained with the other models, at low NAPL volumetric fractions. These results indicate that mass transfer rates at low volumetric fractions could be very important in the NAPL dissolution process. However, since little research has been conducted at lower NAPL volumetric fractions, the validity of the mass transfer models at low NAPL volumetric fractions is questionable.

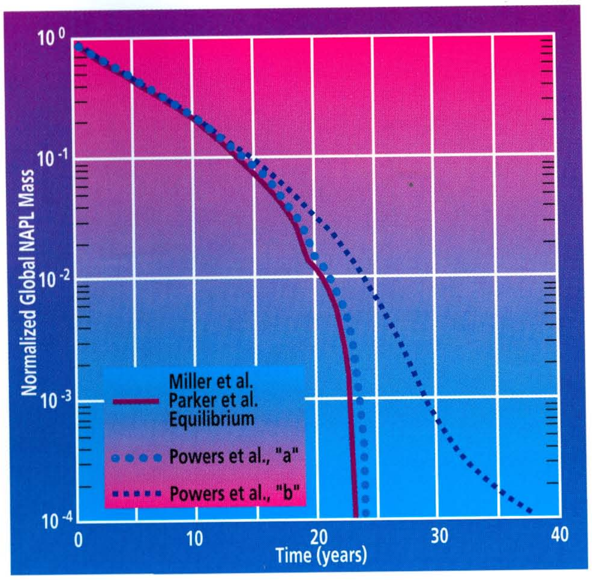

Figure 3. NAPL mass as a function of time for various mass transfer models.

Effects of heterogeneous porous media properties

A set of simulations was performed to study the influence of heterogeneous porous media properties on NAPL flow and transport. The idealized heterogeneous simulation consisted of a single, 9 cm thick, low-permeability lens located 9 cm below the top of the domain. The intrinsic permeability of the lens was one-half of the surrounding material. The other features of the domain and the sequence of the simulations were similar to the previous simulations.

Figure 4 shows the normalized, global NAPL mass as a function of time for the low-permeability lens simulation and an equivalent, but homogeneous, simulation. This figure indicates that NAPL dissolution for the low-permeability lens simulation proceeded at a significantly lower rate than the homogeneous simulation. The slower dissolution may be explained by the NAPL configuration. Figure 5 shows the NAPL volumetric fractions at the time when the NAPL reached a residual level. The effect of the low-permeability lens is indicated by the pooling of NAPL on the top of the lens. The majority of the NAPL is contained in the pool above the low-permeability lens and inside the low-permeability lens. These portions of the NAPL were subjected to low aqueous phase flow rates and correspondingly low mass removal rates.

Figure 4. NAPL mass as a function of time for homogeneous and heterogeneous systems.

It should be emphasized that this result was obtained with a highly idealized heterogeneous system and is only one of an infinite number of heterogeneous configurations. A significant amount of research remains to be done before a complete understanding of NAPL dissolution in heterogeneous systems can be achieved. This means that many more simulations need to be conducted over a wide range of porous media properties.

Improving model efficiency

A potentially large number of simulations will be required to gain a better understanding of the NAPL dissolution process in realistic systems. Simulations in three dimensions may eventually be needed, and more sophisticated physical, chemical, and biological mechanisms will be incorporated into the simulations. However, multiphase flow and transport simulations are computationally intensive from both a CPU-time and memory standpoint. New computational techniques for making the simulations' algorithms more efficient will be required to conduct these simulations in a reasonable amount of time.

One way to improve code efficiency is to recognize that the NAPL and aqueous phase often are found simultaneously only within a small area relative to the full simulation domain. The solution of the flow equations is much more intensive when both the aqueous phase and NAPL are present. This type of problem is ideally suited to the application of domain decomposition (DD) techniques. DD techniques involve separating the problem into a series of subproblems and solving each subproblem separately. The advantages of DD are that the two-phase flow equations are solved only in the subdomain containing the NAPL. In the algorithm developed here[5] the subdomains are resized at each timestep as the size of the region containing NAPL changes. The aqueous phase solutions from each subdomain are updated in a sequential iterative process until the solutions converge at the boundary shared by the subdomains. We also developed a parallel DD scheme that involves decoupling the subdomains by lagging the one subdomain solution behind the solution of the other subdomain. The decoupling allows for solving each subdomain problem on a separate processor, but presumably at the cost of slower convergence.

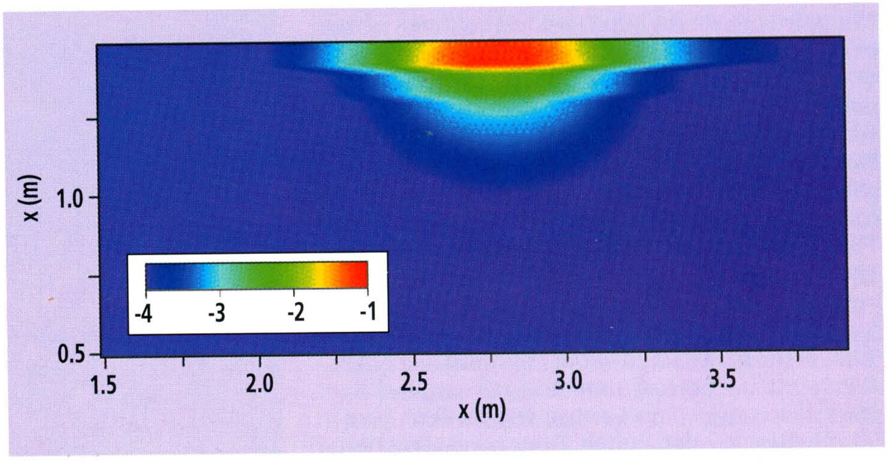

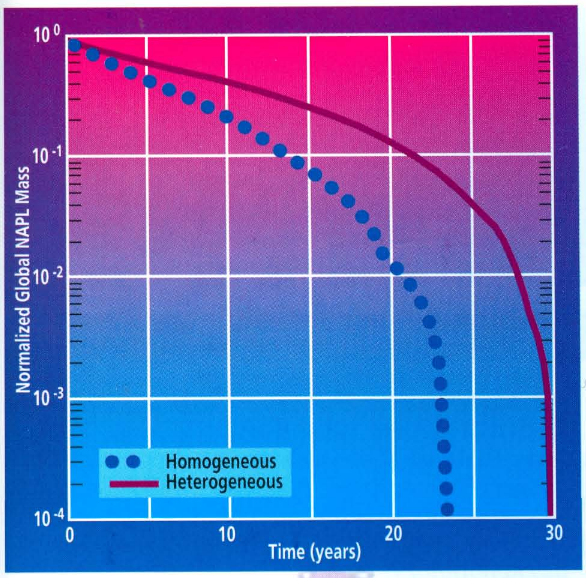

Figure 5 (left). Simulation of NAPL infiltration into a heterogeneous system: color bar indicates logarithm of NAPL volumetric fractions.

Figure 6 (right). Results of test problem for domain decomposition schemes.

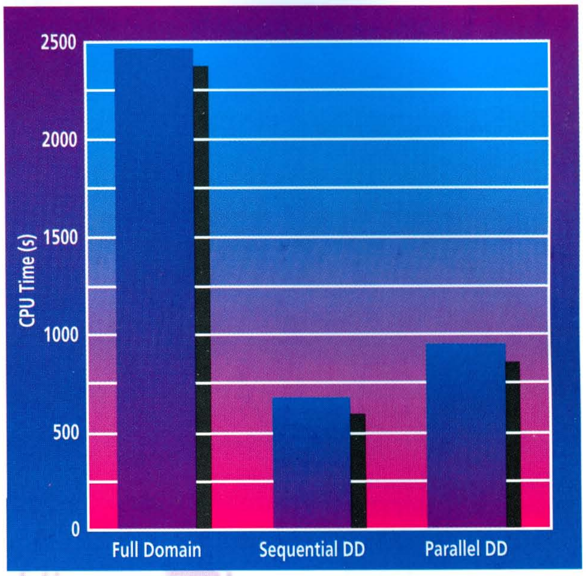

A sample problem was used to test the sequential DD scheme, the parallel DD scheme, and a full domain scheme. The size and properties of the test problem were similar to the problems used in the previous simulations. Figure 6 shows the results of the test problem in terms of the CPU time required to solve it. Application of the two DD schemes resulted in at least twofold CPU-time reductions over the full domain scheme. Application of the parallel DD scheme resulted in longer CPU times when compared to application of the sequential DD scheme on a multiprocessor CRAY Y-MP system. A speedup of 1.3 (measured as the ratio of CPU time to wall clock time) was achieved with the parallel DD scheme, but the number of iterations required for convergence increased significantly. It is likely that parallel schemes can be developed that have superior convergence properties and take better advantage of multiple-processor platforms.

Summary

Simulations of NAPL dissolution revealed that, in all cases but one, differences in NAPL-aqueous phase mass transfer models are not important and that the models gave similar results to a simulation using an equilibrium approach. The exception was a mass transfer model that is more sensitive to NAPL volumetric fractions. Simulations conducted with this model exhibited significantly slower dissolution when the NAPL volumetric fractions were low. This result indicates that NAPL dissolution at low NAPL volumetric fractions may be crucial to predicting the length of time required to dissolve NAPL contaminants in the subsurface.

A simulation conducted with a stratified heterogeneous porous media produced dissolution rates that were significantly lower than in the corresponding homogeneous case. However, this configuration of heterogeneous porous media is idealized and is only one of many possible configurations that might be found in the subsurface. Additional modeling efforts must be conducted, over a wide range of porous media heterogeneity, before the effects of heterogeneity on the NAPL dissolution process are completely understood. The efficiency of the computational algorithms must be improved so that these and more complex simulations can be conducted in a reasonable amount of time. The development and application of domain decomposition algorithms represents a significant enhancement over conventional algorithms. Even better performance is expected when improved parallel domain decomposition algorithms are implemented.

About the author

Alex S. Mayer is an assistant professor in the Department of Geological Engineering and Sciences and the Department of Civil and Environmental Engineering at Michigan Technological University. His research and teaching interests are experimental and mathematical modeling studies of groundwater flow and contaminant transport and remediation. He received a B.S. degree in civil engineering from Brown University and M.S. and Ph.D. degrees in environmental engineering from the University of North Carolina at Chapel Hill.

References

1. U.S. Environmental Protection Agency, *National Primary Drinking Water Standards*, U.S. Environmental Protection Agency Office of Drinking Water, Paper 810-F-94-001, 1994.
2. Miller, C. T., M. M. Poirier-McNeill, and A. S. Mayer, "Dissolution of Trapped Nonaqueous Phase Liquids: Mass Transfer Characteristics," *Water Resources Research*, Vol. 26, pp. 2783-2796, 1990.
3. Powers, S. E., L. M. Abriola, and J. W. Weber, "An Experimental Investigation of NAPL Dissolution in Saturated Subsurface Systems: Steady-State Mass Transfer Rates," *Water Resources Research*, Vol. 28, pp. 2691-2705, 1992.
4. Parker, J. C., A. K. Katyal, and J. J. Kaluarachichi, "Modeling Multiphase Organic Chemical Transport in Soils and Groundwater," U.S. Environmental Protection Agency, 1990.
5. Mayer, A. S., "Application of Domain Decomposition Techniques for Multiphase Groundwater Problems," Xth International Conference on Computational Methods in Water Resources, Kluwer Academic Publ., Dordrecht, Germany, pp. 951-958, 1994.

PARALLELIZATION
of an ocean general circulation model on the CRAY T3D system

Josef M. Oberhuber
Deutsches Klimarechen-zentrum GmbH (DKRZ)
Hamburg, Germany

Klaus Ketelsen
Cray Research, Inc.

Some evidence suggests that increasing amounts of atmospheric pollution are threatening the delicate balance of the Earth's ecosystems. High-resolution ocean models used in coupled atmosphere-ocean-land simulations are vital to predicting climate changes and their time scales. Their size and complexity demand such vast amounts of computing resources that costs have been prohibitive and turnaround times impractical to date. A cooperative project completed recently between DKRZ, the German Climate Computing Center in Hamburg, and Cray Research, during which the ocean general circulation model OPYC—already implemented on a CRAY C90 system—was ported to a CRAY T3D massively parallel processing system, meant an important step forward in finding a practical, affordable solution for computationally intensive ocean modeling.

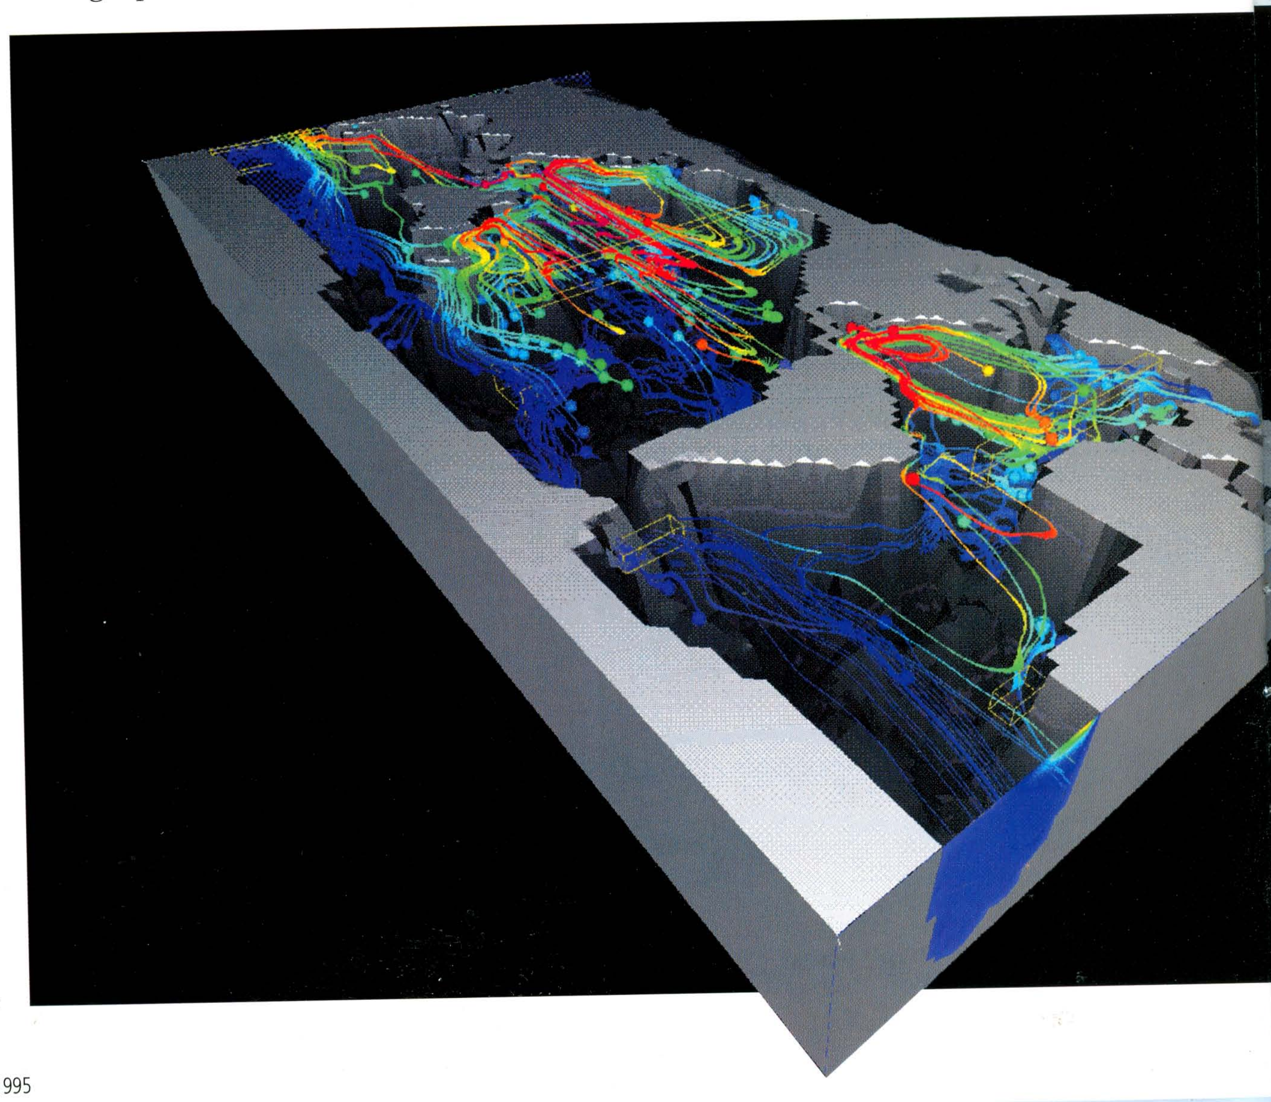

Figure 1. OPYC, a state-of-the-art system to simulate ocean general circulation, reveals a three-dimensional flow in the global ocean basin visualized by particle trajectories, where particles have been released at specified locations. Trajectories in blue indicate low temperature, red indicates high temperature.

Rising sea levels and global warming caused by increasing amounts of greenhouse gases released into the atmosphere are examples of pressing global environmental issues for which decision makers must find quick solutions to preserve life on planet Earth. To document the delicate balance of the Earth's ecosystems and provide decision makers with predictions about the consequences of disturbing that balance, environmental researchers use computer simulation models that span time scales ranging from days to millennia.

In recent years, ocean models have become an important part of these environmental studies because of the ocean's pivotal role in climate variability. Ocean models are used to study the variability of ocean circulation across years, centuries, and millennia in order to understand the physical causes of El Niños (warm events in the equatorial Pacific) or of warm or cold periods in the Earth's history, for example, and to study their influence on global ocean circulation. Extremely high-resolution ocean models are capable of resolving oceanic eddies and allow studying their influence on global ocean circulation or the natural variability of the ocean, while low-resolution models are used to study climate variability on extremely long time scales. Used in coupled atmosphere-ocean-land models, ocean models are vital to predicting climate changes and their time scales resulting from any kind of increased atmospheric pollution.

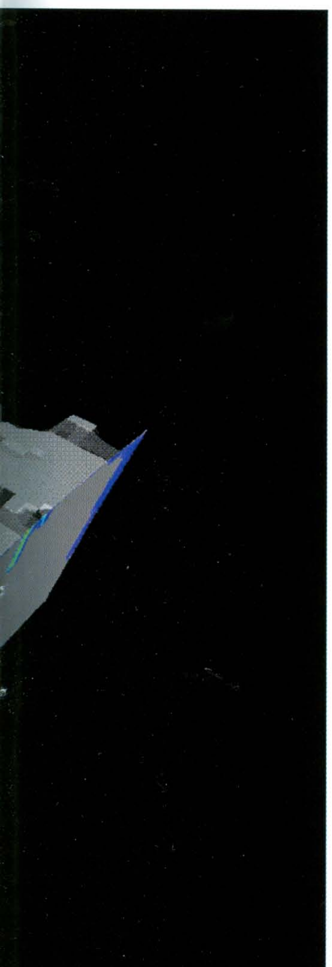

Even when run on highest-performance parallel-vector supercomputers, the sheer size and complexity of high-resolution ocean models make computation costs prohibitive and turnaround times impractical when used for time scales of centuries. For example, a high-resolution global ocean model requires a resolution of 0(10 km), which results in a 3600 x 1800 horizontal grid size times approximately 40 levels vertically.

Today, massively parallel processing (MPP) systems, such as Cray Research's CRAY T3D system, hold the promise of providing the vast amounts of computing resources required for ocean modeling. To begin tapping these powerful resources, researchers at the Deutsches Klimarechenzentrum (DKRZ), the German Climate Computing Center in Hamburg, and computer specialists from Cray Research are cooperating on a project to port OPYC[1] (Ocean model with isoPYCnal coordinates), an ocean general circulation model already implemented on DKRZ's CRAY C90 system—and on many other Cray Research systems worldwide—to the CRAY T3D system.

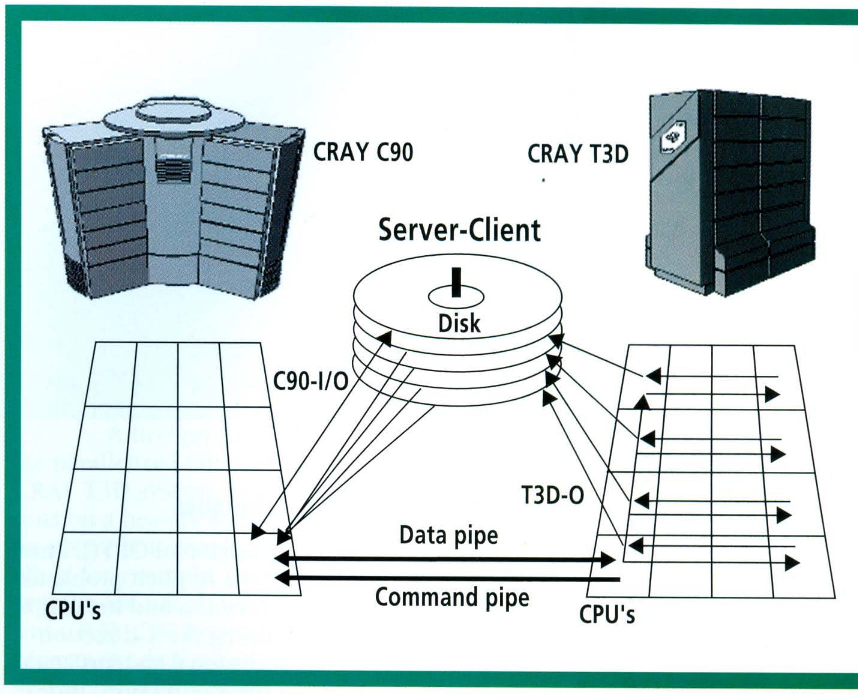

Figure 2. A client-server concept was used during parallelization of the ocean general circulation model OPYC on the CRAY T3D system. The diagram shows the flow of various data streams. A data and a command pipe are used to initialize model variables on the CRAY T3D system and to extract the final ocean state for a subsequent restart of the model. Intermediate data frequently generated by the ocean model are saved on disk. These data, which are still partitioned in one geographical direction, are asynchronously read by the CRAY C90 system and converted to a final file format.

The ocean model

The ocean general circulation model OPYC is a state-of-the-art system for simulating large-scale ocean circulation, as shown in Figure 1. It uses primitive equations in flux form with free surface formulation. The density field is computed from temperature and salinity via a general state equation for seawater. Horizontally, the model uses a regular, staggered grid on spherical coordinates. Vertically, isopycnals (planes of constant potential density) are used as Lagrangian coordinates. In addition to standard parameterizations for convection, isopycnal and diapycnal mixing, a mixed layer model simulates the depth of a turbulent surface layer. A dynamic-thermodynamic sea-ice model predicts ice and snow cover according to a comprehensive ice rheology. The OPYC model uses a fully implicit time integration scheme. This leads to nonlinear equations for the three-dimensional wave equation, for advection and diffusion of momentum, mass, heat, and salt, and for thickness and concentration of ice and snow and the related momentum.

Parallelization on the CRAY T3D system

The basic strategy for parallelizing the time-stepping part is the (x,y)-domain decomposition. Columns (i.e., the z-direction) are located on the same processing element (PE). Each PE carries additional latitudes and longitudes for boundary values, as does the CRAY C90 version for the single domain. In addition, the client-server concept is used as a guideline (Figure 2).

Table 1. Benchmarks for T106 model without and with postprocessing.

PEs	npe_x	npe_y	*Without postprocessing* Time (sec)	MFLOPS	*With postprocessing* Time (sec)	MFLOPS	MFLOPS/PE	Speedup
16	4	4	2459	273	2474	271	16.9	16.0
32	4	8	1234	541	1250	534	16.7	31.6
64	8	8	675	990	680	983	15.3	57.9
128	8	16	346	1930	346	1930	15.1	114.4
256	16	16	198	3368	200	3338	13.0	196.9
512	16	32	97	6885	101	6646	13.0	393.8

Changes in the iterative techniques

As in the CRAY C90 version of OPYC, linear equation systems for each of the implicit problems are solved directly in the x-direction and iterated in the y-direction. After partitioning the x-direction into npe_x subdomains, the tridiagonal systems could be solved across all PEs in the x-direction. Independent of the chosen algorithm, this would result in a number of sequential operations across all PEs in the x-direction to solve the linear equations. Consequently, one may expect that such an algorithm will not scale well when the code is run on a truly massively parallel computer with 512 PEs, for example. This was demonstrated by Eltgroth,[2] who developed an efficient technique for the solution of linear equations for MPP systems.

Such a technique has already been applied to OPYC's sea-ice model, which, as part of the coupled atmosphere-ocean-land model at Lawrence Livermore National Laboratory (LLNL) in Livermore, California, is coupled to the Modular Ocean Model (MOM) from Geophysical Fluid Dynamics Laboratory (GFDL) in Princeton, New Jersey. Eltgroth showed that such a scheme scales reasonably only up to 0(100) PEs.

To achieve good scaling for a larger number of PEs, we developed an engineering type of approach. Instead of solving one tridiagonal equation across all PEs in the x-direction, we formulated a complete linear equation system for each PE and coupled the equations by a back-and-forth shifting of the grid point data relative to the virtual address space. This data shifting is carried out after each iteration. Thus, boundary values located near the interface between adjacent PEs will be located in the center of the PE grid point space and thus in the center of the linear equation system during the next iteration. This method results in a convergence rate similar to a fully direct method and yields surprisingly few changes to the version implemented on the CRAY C90 system.

Heterogeneous environment

The code was split into client and server components. The server remains on the CRAY C90 system and performs the initialization phase and postprocessing. In addition, the server supplies the client with atmospheric forcing data required periodically each month. The server reads these data and sends them via a named pipe to PE(0,0). PE(0,0) distributes the data to all other PEs. The client runs on the CRAY T3D system and carries out the computation. During the initialization phase the server fills all COMMON blocks and sends a copy to PE(0,0) on the CRAY T3D system via a named pipe. PE(0,0) partitions the data in the x- and y-directions and distributes the respective parts to the other PEs. At the end of the client program all modified COMMON blocks are returned to the server. These data are used to generate a restart file. Using the client-server concept for porting the code offers the following advantages:

- ☐ The number of routines that had to be ported was minimized. Only the computationally intensive part of OPYC runs on the CRAY T3D system. On the other hand, if the model is initialized once from scratch, the expensive and hardly parallelizable routines for data preprocessing, etc., run on the CRAY C90 system.
- ☐ The CRAY T3D version of OPYC is transparent to the user. All files that communicate with the outside world are created on the CRAY C90 system. From the user's point of view there is no difference using either the combined CRAY C90/T3D version or the full CRAY C90 version.

Communication

The parallel part of OPYC is running on the CRAY T3D system as a single program multiple data (SPMD) code. Explicit shared memory message passing is used for communication on the CRAY T3D system. Communication between different PEs is required in three different cases:

- ☐ Zonal and meridional boundary conditions
- ☐ Data shift in the x-direction to improve convergence of the solvers
- ☐ Global sums

An analysis of the network activity on 32 PEs showed that 1.5 percent of the total CRAY T3D time is used for global sums, 2 percent for zonal and meridional boundary conditions, and 8 percent for the data shift mechanism. While the network is active, 964.4 Mbytes/s are moved between all PEs.

Data postprocessing

The output of postprocessing data has been optimized to meet the requirements of a large CRAY T3D system with 512 PEs. To take better advantage of the I/O bandwidth of the CRAY T3D system, the first PE, i.e., PE(0,y), of all y-groups collects the data of its group and writes them onto an intermediate disk file. The server collects the postprocessing data from the intermediate files, converts the format from IEEE to Cray-binary format, and merges them into the final postprocessing file. This asynchronous procedure offers the advantage that the server may pick up the data from distributed files whenever the CRAY C90 system's scheduling provides the server with some resources.

Optimization for the DEC Alpha chip

To improve the performance of OPYC, optimization work has been done to take into account the architecture of Digital Equipment Corporation's Alpha chip. Alpha is a RISC processor with a slow hardware floating point divide. The main goals of the optimization work have been

1. Reducing memory traffic. The following techniques were used:
   - □ Loop collapsing
   - □ Storing intermediate results in temporary variables
   - □ Reordering the sequence of computation
2. Better use of cache, for example, due to swaps of loop indices to achieve an increment of one in the innermost loop.
3. Reducing the number of divides by rearranging loops or exchanging divides by a constant for a multiply by a constant.

Benchmark tests

Runs with a variety of partitions (Table 1) for the global ocean model with T106 resolution (320 x 160 horizontal grid points) show that the performance does not drop significantly for up to 512 PEs. There are only small differences when data postprocessing is switched on or off. After completing the Alpha chip optimization we achieved 292 MFLOPS on 16 PEs, 593 MFLOPS on 32 PEs, 1029 MFLOPS on 64 PEs, and 2173 MFLOPS on 128 PEs for a lower-resolution T42 (128 x 64 horizontal grid points) model. This means that for this problem 23 PEs on the CRAY T3D system are equivalent to one CRAY C90 CPU. The T106 model ran with 7.4 GFLOPS on the CRAY T3D system with 512 PEs.

Conclusion

During parallelization of the ocean general circulation model OPYC, the client-server layout was found to be useful because of the small amount of work involved in porting the entire code and because of the user transparency achieved.

Use of an implicit time step scheme in combination with the line-relaxation method and the data shift mechanism resulted in a code that involves surprisingly few changes relative to the CRAY C90 version. The network activity caused by the data shift mechanism was found to be tolerable on the CRAY T3D system and scales very well.

Data postprocessing was found to be crucial for production runs. An asynchronous write from the client and read into the server was found to be the best strategy. The server reconstructs all full-domain data in a CRAY C90-compatible format. Pipes are used only to initialize the CRAY T3D system, to return data that the server requires for generating the restart file, and to supply the client with monthly data. It was not found useful to transfer data frequently across pipes because the server on the CRAY C90 usually is not running in dedicated mode.

A first production run demonstrated that the parallelized OPYC runs efficiently on a 32-PE CRAY T3D system. Despite the fact that the server runs on a heavily used CRAY C90 system at DKRZ, no severe bottlenecks, caused by CRAY C90 scheduling, for example, have appeared. This scientific experiment, which used approximately 1500 equivalent CRAY C90 hours and produced 100 Gbytes of permanent data, confirmed the reliability of Cray Research's MPP system.

Porting the general ocean circulation model OPYC to the CRAY T3D system is an important step toward finding a practical, affordable solution for computationally intensive ocean modeling. Improved turnaround and greater model precision are indispensable for finding timely answers to critical environmental problems.

Acknowledgments

Josef Oberhuber thanks Cray Research for its support for this project, both for access to a CRAY T3D system at Cray Research's computing facility and for the time provided by Klaus Ketelsen. The authors acknowledge the support of Michael Böttinger from the Visualization Group at DKRZ for generating the three-dimensional image using the Data Visualizer from Wavefront Technologies, Inc. in Santa Barbara, California.

About the authors

Josef M. Oberhuber is a modeler within the Model Application and Development group at DKRZ in Hamburg, Germany, who focuses on ocean modeling and coupled ocean-atmosphere-land models for climate predictions. Oberhuber received a Ph.D. degree in meteorology in 1984. He developed an ocean general circulation model at the Max Planck Institute for Meteorology and the Meteorological Institute of the University, both in Hamburg, Germany.

Klaus Ketelsen is an applications analyst in the Marketing and Sales Support group at Cray Research GmbH, Germany, focusing on porting application codes to the CRAY T3D system. He received an advanced degree in electrical engineering from the University of Hannover in 1973 and has worked in the petroleum industry for nine years.

References

1. Oberhuber, J. M., *The OPYC Ocean General Circulation Model*, Technical Report No. 7, Deutsches Klimarechenzentrum, Hamburg, 1993, p. 130.
2. Eltgroth, P. G., *Two Portable Parallel Tridiagonal Solvers*, UCRL-JC-118017, Lawrence Livermore National Laboratory, 1994, p. 9.

FLENUMME

Supercomputing and database management at Fleet Numerical Meteorology and Oceanography Center

R. Michael Clancy, and James L. Copeland
Fleet Numerical Meteorological
and Oceanography Center

Operating around the clock, distributing more than 30,000

TOCCEN

Figure 1. Significant wave height in meters (contours and color) and primary wave direction (arrows) from a 24-hour forecast valid at 0000 GMT 28 September 1994. The region of high waves (greater than 7.5 m in height) south of Japan is a result of Typhoon Orchid, which was moving northward toward eventual landfall in southern Japan. The relative minimum in wave heights in the center of the high-wave region reflects the eye of the typhoon. This was a routine operational product made at FLENUMMETOCCEN with a high-resolution (0.2° latitude by 0.2° longitude) regional implementation of the WAM Third-Generation Wave model.[2] In this instance, the wave model was forced by surface wind stress at 1° latitude by 1° longitude resolution from the Navy Operational Global Atmospheric System (NOGAPS) meteorological model.[3] Both the WAM wave model and NOGAPS meteorological model run operationally on the POPS-2 CRAY C90 system.

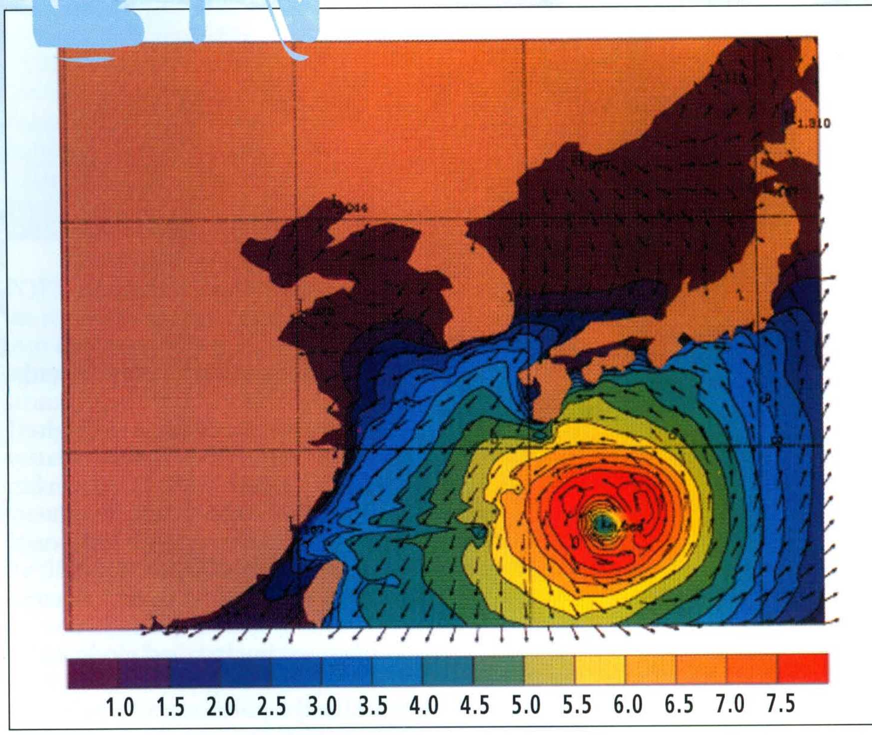

The U.S. Navy's Fleet Numerical Meteorology and Oceanography Center (FLENUMMETOCCEN) is the Department of Defense (DoD) central production site for all standard, fully automated real-time meteorological and oceanographic prediction products. FLENUMMETOCCEN fulfills this role through a suite of sophisticated global and regional meteorological and oceanographic models, extending from the top of the atmosphere to the bottom of the ocean.[1] These models are supported by one of the world's largest real-time meteorological and oceanographic observational databases, consisting of over one million new observations each day from satellites and other global observing systems. FLENUMMETOCCEN operates around the clock, 365 days per year, and distributes more than 30,000 meteorological and oceanographic products per day to DoD and civilian users around the world. Figure 1 shows a sample FLENUMMETOCCEN product which indicates significant wave height in meters (contours and color) and primary wave direction (arrows) of September 1994's Typhoon Orchid.

In the past year, FLENUMMETOCCEN upgraded its meteorological and oceanographic modeling infrastructure dramatically with the installation of the Primary Oceanographic Prediction System (POPS-2) supercomputer system (Figure 2). POPS-2 is complemented by a similar system referred to as POPS-1, at the Naval Oceanographic Office at the Stennis Space Center in Mississippi. The Naval Oceanographic Office system is used to generate human-machine, special automated, and real-time oceanographic prediction products, as well as for research and development support. POPS-1 and POPS-2 are linked via high-speed (1.5 Mbits/s) T1 data communications; together they provide the means to develop and operate state-of-the-art meteorological and oceanographic models.

Hardware

In June 1994, POPS-2 became the primary production platform for all operational meteorological and oceanographic prediction models at FLENUMMETOCCEN. POPS-2 replaced an aging and obsolete Control Data Corporation CYBER 205 computer, which had been in use at the center since 1982.

The heart of the POPS-2 system is a CRAY C90 supercomputer on which all of the meteorological and oceanographic prediction models run (Figure 3). In its current configuration, the machine has 8 processors, 128 Mwords of central memory, a 256-Mword SSD solid-state storage device, 46 Gbytes of DD-60 disk space, and 45 Gbytes of DA-301 Redundant Array of Independent Disk (RAID) space.

The CRAY C90 system is complemented by a CRAY Y-MP 2E system, which acts as a data server for the CRAY C90 production platform. The CRAY Y-MP 2E system has 2 processors, 32 Mwords of central memory, a 128-Mword SSD, 46 Gbytes of DD-60 disk space, and 10 Gbytes of DA-301 RAID space. The CRAY Y-MP 2E system drives

Figure 2 (below). Dedication of the POPS-2 CRAY C90 system at FLENUMMETOCCEN. From left to right: RADM Geoffrey Chesbrough, then Oceanographer of the Navy; CAPT Robert Plante, Commanding Officer of FLENUMMETOCCEN; Congressman Sam Farr, D-CA; RADM John Chubb, then Commander of the Naval Meteorology and Oceanography Command; and RADM Thomas Mercer, Superintendent of the U.S. Naval Postgraduate School.

meteorological and oceanographic products per day

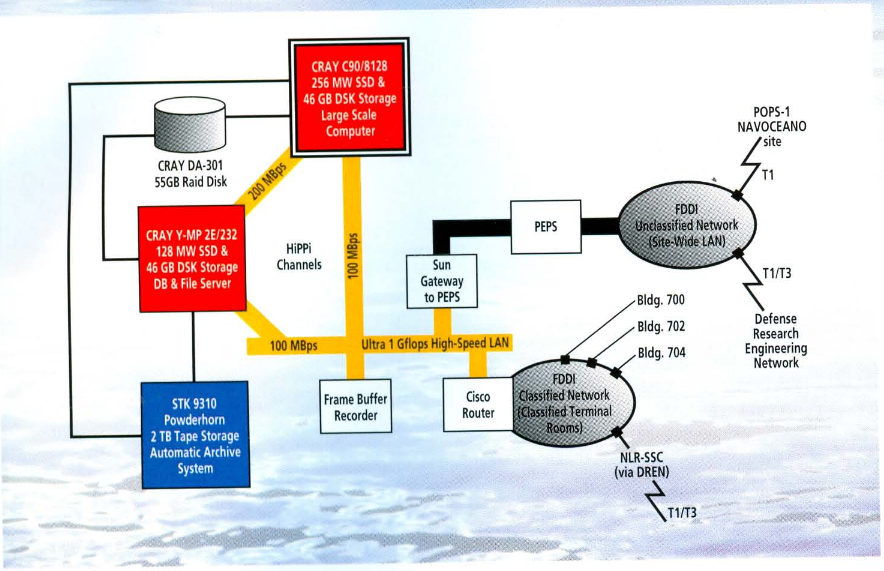

Figure 3. Architecture of the POPS-2 System at FLENUMMETOCCEN.

a StorageTek 9310 tape archive system, which supports 2 Tbytes of storage in its current configuration.

As shown in Figure 3, the CRAY C90 and CRAY Y-MP 2E systems are connected to an Ultra 1-Gbyte/s high-speed LAN via 100 Mbyte/s High-Performance Parallel Interface (HIPPI) channels. The CRAY Y-MP C90 and CRAY Y-MP2E systems also are linked directly by a 200-Mbyte/s double-wide HIPPI channel. A Sun Gateway provides a path from POPS-2 to a suite of CDC CYBER 860 mainframes, collectively referred to as the Primary Environmental Processing System (PEPS). The POPS-2 system is operated in classified mode 24 hours per day to allow processing of classified data in the meteorological and oceanographic models. PEPS provides for a trusted one-way flow of unclassified products out of the POPS-2 system and onto an unclassified site-wide LAN of personal computers and workstations for further distribution. Access to POPS-2 is possible only through classified terminal rooms linked via a classified network (Figure 3).

Finally, POPS-2 is fully supported by an uninterruptible power supply system with a diesel generator backup. This makes the supercomputer system impervious to commercial power fluctuations and outages.

Software

With the installation of the POPS-2 hardware suite came sweeping and revolutionary changes in FLENUMMETOCCEN's software support infrastructure. Key aspects of this included a nearly complete transition to the UNIX operating system, heavy reliance on commercial off-the-shelf software, and a fundamental commitment to the philosophy of open systems architecture. The replacement of the PEPS system with high-performance workstations over the next several years will complete FLENUMMETOCCEN's full transition to UNIX and open architecture.

The Supervisor Monitor Scheduler (SMS) software, developed by the European Centre for Medium-Range Weather Forecasts (ECMWF) to provide centralized monitoring and scheduling of complex job streams, controls the operational job stream on POPS-2. FLENUMMETOCCEN is now using a formal, structured process for development, documentation, and implementation of new general-purpose utility routines as required by the meteorological and oceanographic models. Software standards for all applications as well as new procedures and tools for performing configuration management of all of the POPS-2 software also are in place.

Database management

Perhaps the most far-reaching and revolutionary change in the FLENUMMETOCCEN software infrastructure, however, was the implementation of the Integrated Stored Information System (ISIS), which uses a relational database management system to manage the data ingested and produced by the meteorological and oceanographic models on POPS-2.[4] This system has established FLENUMMETOCCEN as a leader in meteorological and oceanographic data management, and it is providing unprecedented flexibility in meeting new and more complex meteorological and oceanographic modeling requirements.

ISIS evolved from the Naval Environmental Operational Nowcasting System (NEONS), which was developed by the Naval Research Laboratory.[5] As in NEONS, ISIS data are organized into generic types (for example, gridded fields, latitude-longitude-time, geographic, and image) and represented by a set of cross-referenced data tables, which contain much information about the data (metadata) in addition to the data itself. ISIS extends the capabilities of NEONS to handle the much larger data volumes and many more data types required to support operations at FLENUMMETOCCEN.

Figure 4 shows the role of ISIS in the FLENUMMETOCCEN software hierarchy. Meteorological and oceanographic model applications interface with the ISIS software shell, which in turn uses embedded Structured Query Language (SQL) to interface with the relational database management system (RDBMS). The RDBMS interfaces with the UNIX operating system, which ultimately performs input/output (I/O) operations to disk or tape. The hierarchy is very strict; applications are not permitted to bypass the ISIS layer to reach the database tables, and the RDBMS is not permitted to bypass the UNIX layer to perform storage and retrieval. These rules allow for the technology at any level to change independently.

EMPRESS, a commercial off-the-shelf client-server database management system, was selected as the RDBMS. EMPRESS is compliant with U.S. government standards for portability and open-systems architecture design, including ANSI standards for SQL and U.S. military standards for TCP/IP communication protocols. The ISIS/EMPRESS combination provides locking, integrity checking, access and concurrency controls, sophisticated data indexing, storage and query optimization, interactive access, archiving, and backup and recovery of data across distributed UNIX platforms.

ISIS/EMPRESS currently runs operationally on both the CRAY C90 and CRAY Y-MP 2E systems, and it will soon be extended to a network of high-performance Sun workstations. ISIS currently manages over 15 Gbytes of highly dynamic data organized into approximately 3000 data tables, with these numbers expected to double in the next year as new data types and applications are implemented.

Database technology arose and matured in the business community, where the principal need was to deliver small quantities of randomly selected data. Meteorological and oceanographic models typically must process very large amounts of data at high speed, and the records in meteorological and oceanographic databases are significantly larger than their counterparts in business applications. Thus, efficient management of very large data records, which are referred to as binary large objects (BLOBs), is a fundamental and important requirement for ISIS. BLOBs may be gridded field data output from meteorological and oceanographic models, satellite images, meteorological or oceanographic observational data, geographic data, computer programs, or anything the user chooses them to be. BLOBs may, and often do, have internal structure; it is merely unknown to the RDBMS.

Most ISIS data is stored as BLOBs. Within the BLOBs, ISIS uses the World Meteorological Organization-supported gridded binary (GRIB) and binary universal form for data representation (BUFR) self-referential formats for storage of gridded fields and observational data, respectively. ISIS uses "true" GRIB and BUFR, which means that data is stored in its packed binary transmission format, allowing inspection down to the bit level.

ISIS uses Cray Research's Data Migration Facility (DMF) to archive data on the tape silo system. Some data also may be archived on smaller systems such as jukeboxes. ISIS presents a seamless interface to the user, with appropriate controls to prevent users from clogging the system with large data requests. To save processing time, ISIS archives database tables rather than extracting the data into another archive format.

Recent ISIS enhancements include optimization of block sizes, dynamic indexing during production processes, SQL query optimization, data partitioning by 12-hour intervals to prevent tables from becoming too large, and data caching of small tables. Recent EMPRESS enhancements include a rewrite of the client-server software to take advantage of fast TCP/IP service on both the Cray Research and Sun platforms and considerable reduction in I/O overhead in local database performance.

Summary

The POPS-2 system became the primary production platform at FLENUMMETOCCEN in June 1994, replacing an obsolete CYBER 205 computer. The CRAY C90 supercomputer, which is the heart of POPS-2, provided the computing power to produce improved operational meteorological and oceanographic prediction products in support of all branches of DoD.

As an integral part of the POPS-2 project, revolutionary and far-reaching changes were ushered into the meteorological and oceanographic modeling support infrastructure at FLENUMMETOCCEN. Transition to the UNIX operating system, heavy reliance on commercial off-the-shelf software, a fundamental commitment to the philosophy of open-systems architecture, and implementation of the ISIS relational database management system were the cornerstones of these changes.

The use of ISIS to manage the heavy volumes of data characteristic of FLENUMMETOCCEN's supercomputer-based meteorological and oceanographic modeling operation has been revolutionary. To make ISIS a viable solution in this environment, FLENUMMETOCCEN extended RDBMS technology to the high-performance computing and very large data records typical of meteorological and oceanographic modeling applications. This technology is working well to meet the needs of FLENUMMETOCCEN, and ISIS is providing unprecedented flexibility in meeting new and more complex meteorological and oceanographic modeling requirements.

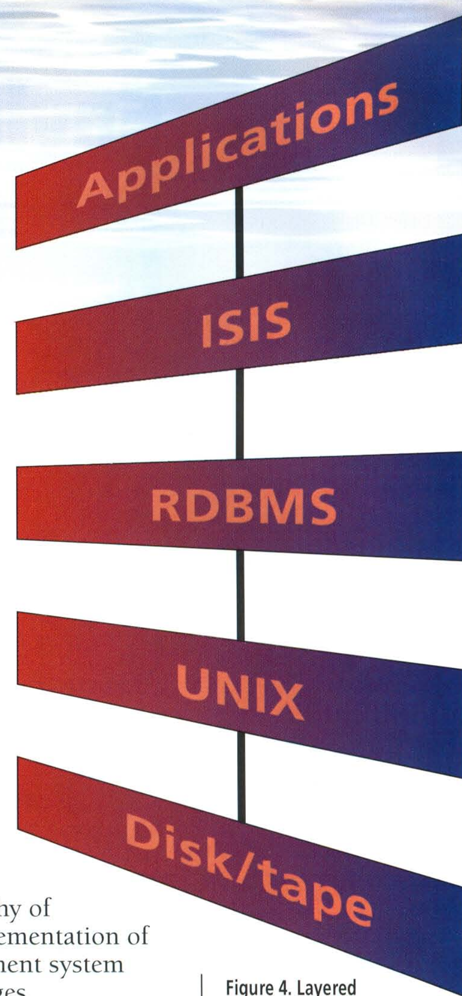

Figure 4. Layered implementation of the database at FLENUMMETOCCEN.

About the authors

R. Michael Clancy heads the Ocean Models Division at Fleet Numerical Meteorology and Oceanography Center.

James L. Copeland heads the Data Administration Division at Fleet Numerical Meteorology and Oceanography Center.

References

1. Plante, R. J. and R. M. Clancy, "An Overview of Operational Prediction Capabilities at Fleet Numerical Meteorology and Oceanography Center," *Proceedings of the MTS '94 Conference*, September 7-9, 1994, Washington, D.C., Marine Technology Society, Washington, D.C. 20036, pp. 352-358, 1994.
2. WAMDI Group, "The WAM Model—A Third-Generation Ocean Wave Prediction Model," *Journal of Physical Oceanography*, Vol. 18, pp. 1775-1810, 1988.
3. Hogan, T. F. and T. E. Rosmond, "The Description of the Navy Operational Global Atmospheric Prediction System's Spectral Forecast Model," *Monthly Weather Review*, Vol. 119, pp. 1786-1815, 1991.
4. Copeland, J. L. and R. J. Plante, "The Development of an Integrated Environmental Database for Navy Operational Production Computers into the 21st Century," *Proceedings, Tenth International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology*, American Meteorological Society, Boston, pp. 278-282, 1994.
5. Jurkevics, A., R. Titus, and J. Clark, "Environmental Database for the Naval Environmental Operational Nowcasting System," *Proceedings, Sixth International Conference on Interactive Information and Processing Systems for Meteorology, Oceanography, and Hydrology*, American Meteorological Society, Boston, pp. 80-83, 1990.

Environmental research gets a boost from new CRAY

Given the complexity of current environmental challenges—from reducing pollution to predicting climate changes—it's no wonder that Cray Research systems are used at 82 percent of all supercomputer-based environmental sites around the world. They are used for tasks such as weather prediction, climate modeling, ecosystem simulations, and analysis of groundwater transportation of contaminants. A case in point: in 1994 the National Weather Service began making climate predictions up to one year in advance, thanks to the power of their CRAY C90 and CRAY Y-MP systems. These forecasts are helping water management officials and transportation managers make informed decisions on such matters as irrigation planning and stockpiling of salt for winter road maintenance.

Environmental research isn't just confined to the largest supercomputer systems anymore. Cray Research's latest air-cooled supercomputers, the CRAY J90 series—including the CRAY J916 and the CRAY J932 systems—already are making inroads into this important market. Climate researchers at the Scripps Institute of Oceanography (SIO) are using a CRAY J916 system in climate research aimed toward seasonal and annual climate predictions. The CRAY J916 system also will be used soon at the National Meteorological Center (NMC) to provide support for climate modeling research, and pre- and postprocessing for models already running on a CRAY C916 system.

The CRAY J932 system features high-speed processing, high memory bandwidth of 51.2 Gbytes/s, and 8192 Mbytes of shared memory—twice the memory capacity of the CRAY J916 system. The system is available with up to 32 processors implemented in CMOS technology in a shared-memory architecture that maximizes workload throughput and parallel processing efficiency.

Retaining full compatibility with the successful CRAY J916 system, the CRAY J932 system provides new opportunities for cutting costs while accelerating research and model refinement. In addition, automatic parallelizing compilers and the UNICOS operating system help deliver the maximum benefit from the CRAY J932 system.

A software environment that maximizes productivity

Its open, standards-based software environment ensures that the CRAY J932 system will integrate easily into most existing distributed computing networks. The UNICOS operating system, a standard UNIX environment enhanced for supercomputing, provides efficient parallel processing,

J932 system

production-quality resource management, security, and network connectivity. This is of particular value to organizations like the U.S. National Meteorological Center (NMC), which hope to achieve the goal of an all UNIX environment by 1996.

All CRAY J932 systems are air-cooled to simplify installation and maintenance and reduce operating costs. They require no special computer-room environments or special power or cooling arrangements. The CRAY J932 processor, memory, and I/O components are field upgradable, making it easy to expand computing resources as job mixes and workloads evolve and grow. The systems elevate midrange supercomputing to a new level of productivity and ease of use. Such "departmental" capabilities were pivotal in a decision to purchase a CRAY J90 system by the NCAR Mesoscale and Microscale Meteorology Division for use in advanced research.

The CRAY J932 system is a true scalable multiprocessor supercomputer that includes a scalar cache to complement its shared-memory bandwidth. The system´s powerful processors, combined with fast, efficient interprocessor communication and synchronization, allow the system to operate equally well as a throughput server and as a time-to-solution server. Both types of capability are required at Scripps Institute of Oceanography (SIO) and NMC. Research requirements necessitate high job throughput to continually advance model skill. The operational forecast demands high performance to complete each run in a constrained time frame.

As a throughput server, the system's multiple CPUs process multiple-job workloads quickly. To collapse the time-to-solution of very large jobs, Cray Research's state-of-the-art parallel-processing software minimizes parallel overhead. The CRAY J932 system's 51.2 Gbytes/s of memory-to-CPU bandwidth exceeds by a factor of 20 that of most competing systems in this price class and maximizes the computational efficiency of both multi-user and memory-intensive workloads. As a result, the system meets the most demanding computing needs regardless of workload mix.

To further expand throughput, CRAY J932 systems can be clustered using the Cray Super-Cluster Environment software. A supercluster environment provides flexible options for configuring a high-performance network, while balancing workload and allocating the most appropriate resources to each job. In addition to the batch processing, data availability and management, resource allocation, and security features that UNICOS provides at the node level, Cray SuperCluster Environment software provides automatic job-level workload distribution, distributed file access and programming, closely coupled connectivity for distributed programs, and high-performance shared data access, among other features. A clustered environment can include a wide range of systems, from supercomputers and servers to workstations. Its high-bandwidth I/O makes the CRAY J932 system ideally suited for use as a central cluster server.

To provide maximum connectivity, the CRAY J932 system is equipped with 1 to 16 VME-64 I/O subsystems and can accommodate up to 2.4 Gbytes/s of total I/O bandwidth. This industry standard design provides access to a wide range of price and performance disk products, third-party peripherals, FDDI and Ethernet networks, and I/O devices from vendors such as E-Systems and Storage Technology. For hierarchical storage management (HSM), the system supports Cray Research's Data Migration Facility (DMF). Connecting to existing small computer system interface (SCSI) peripherals is easy, too, because the CRAY J932 system has easily accessible SCSI bulkhead connectors. The CRAY J932 HIPPI interface provides high-performance connectivity to HIPPI-based peripherals and for distributed processing applications such as Parallel Virtual Machine (PVM). This high-performance connectivity also facilitates the transfer of large files to CRAY Y-MP, CRAY C90, and CRAY T3D systems. The CRAY J932 system can be configured with up to eight 100-Mbyte/s HIPPI-to-memory channels. The CRAY J932 system is also designed to allow for future performance enhancements to be incorporated, allowing you to preserve your investment as future enhancements are developed.

The CRAY J90 series is well suited for growing production and complex environmental computing needs, from weather forecasting to ocean-atmosphere models. By providing cost-effective, sustained application performance, the series offers departmental or enterprise-wide capabilities to meet the most challenging requirements. For more information on the CRAY J916 or CRAY J932 system, call 1-800-289-CRAY, visit our World Wide Web server at http://www.cray.com, or email us at crayinfo@cray.com.

From weather forecasting, to ocean-atmosphere models, the CRAY J90 series is well suited for growing production and complex environmental computing needs.

"KNOW THE CUSTOMER"

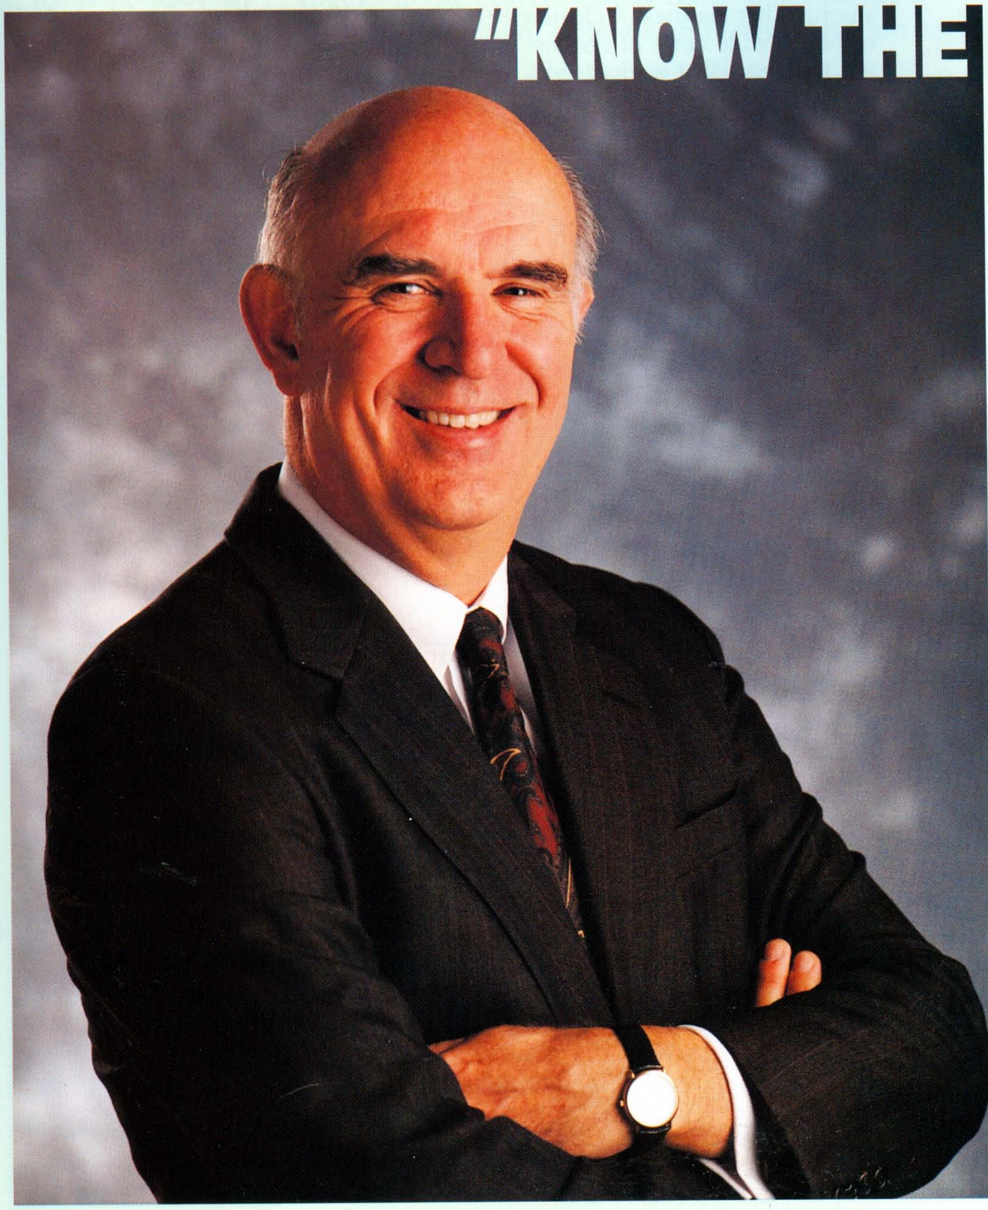

A conversation with J. Phillip Samper new Cray Research Chairman and CEO

When Phil Samper—former Eastman Kodak vice chairman and president of Sun Microsystems Computer Corp., and newly named Cray Research Chairman and CEO—is asked about his marketing philosophy, he doesn't hesitate: "Know the customer's business, be in contact with the customer, get out with the sales organization to understand customers and potential customers."

Samper, 60, is a man of his word.

During his tenure at Sun Microsystems—a company the Wall Street Journal credits Samper with transforming from a "late-stage entrepreneurial firm to a well-managed business"—he made a point of personally contacting customers. A call to the CIO of Fingerhut Companies, a catalogue merchandising leader, was par-

ticularly appreciated for its timeliness. At the height of Fingerhut's Christmas rush season, Samper phoned to make sure the customer's computing resources were keeping pace with Fingerhut's business.

It's that attention to customers' business concerns—along with extensive management leadership that includes a career spanning 28 years with Eastman Kodak and a stint as president and CEO of Kinder-Care Learning Centers, Inc.— that led Cray Research's board to name Samper as successor to John F. Carlson, who retired as chairman in May.

"Phil Samper is the ideal person to lead Cray Research at this important point in the company's history," says Robert G. Potter, a Cray Research board member and chairman of the board's executive committee. "His strong record of business leadership in the computer industry and other industrial technology sectors will help maintain Cray's supremacy in the high-end supercomputing market and drive the company's growing success in new markets. Phil will advance Cray's business strategy and build shareholder value, working closely with Cray President and COO Bob Ewald and other members of the management team."

Ewald shares the board's enthusiasm for Samper. "I knew Phil before he came to Cray," Ewald says. "He's a person of high integrity and he gets things done. He fits the bill perfectly for what we need at Cray Research."

Samper serves on the boards of Armstrong World Industries, Inc., Lancaster, Pennsylvania; Interpublic Group of Companies, New York; Sylvan Learning Systems, Inc., Columbia, Maryland; and on the advisory board of the University of California Business School. An MIT Sloan Fellow (1972-73), he is married to the former Gail Harrison and has two children.

Why did you agree to join Cray Research?

I am excited to assume this leadership role in a company of Cray's stature. The Cray Research name is synonymous with technology leadership and strong customer relationships based on delivering unique value.

What kinds of things did you do at Sun?

We brought out products in record time. We focused on marketing and marketing communications, pumped up public relations, put more emphasis on the international area and on sales.

What from your Sun experience is relevant to Cray?

The value of teamwork. The importance of brand and company image. The importance of shortening the product time cycle. The ability to remain in touch with market opportunities. The importance of successful financial performance.

Why did you leave Sun?

For the reasons we stated at the time. I was on the Sun board, and the plan was for me to run the computer company for 12 to 24 months and bring some elements together. We laid out our objectives, and after 12 months those objectives were met and, in some cases, significantly surpassed. I told Scott McNealy it was time to find a replacement. He initially tried to dissuade me, but I knew the time was right. So at the end of 1994 we announced that I would be leaving Sun. My last day at Sun was March 1, 1995. My discussions with Cray started after that. I'm an admirer of Sun, and Scott McNealy and I remain great friends.

Is your plan at Cray similar?

No. The board and I haven't placed any time limit on my tenure at Cray.

Will the Cray-Sun relationship become stronger?

Cray and Sun established a good and growing relationship before I arrived at Cray. The two companies collaborate on a number of fronts. Anything further we do with Sun will be because it is good for Cray. I do expect opportunities for Cray and Sun to do more together, though, and I think my relationship with Sun will be useful there.

Cray's market has been tough. Do you see prospects for returning to growth?

Not just prospects, I see progress—moving into new technical and commercial markets. The key is getting people to work together. We did a lot of this at Kodak and Sun. Success requires trust, honesty, and teamwork. Another key is putting the infrastructure in place to move quickly into new markets.

How important is the high-end supercomputer market for Cray?

Cray Research enjoys a 70 percent global share of the supercomputer business. We don't want to give that up. At the same time, Cray needs to expand into new growth markets by leveraging its good products and respected technology into new markets here and abroad. It's a question of execution and moving more aggressively. And we need a superior scalable platform that carries applications forward with superior performance and price-performance.

What applicable experience do you bring to Cray?

I have experience managing companies and operations, experience in all areas of management, including manufacturing, marketing, and sales. I think that experience can be helpful at Cray.

What about your international experience?

I spent 17 of my 28 years at Kodak outside of the United States. At Sun, I was based in Mountain View, California but traveled the world. I grew up in South America.

What about cost-cutting? Will you be doing that at Cray?

Cost-cutting is a logical step for any company today, including Cray. You have to have good cost-performance to be competitive.

Do you have a vision for Cray?

Cray Research has a great image in technology. I want to help the company move more aggressively into new markets, improve cost-performance, productivity, and ultimately shareholder value.

APPLICATIONS UPDATE

FAIM software on a CRAY T3D system slashes chip masking time

Photolithography encompasses all of the steps necessary in the process of manufacturing computer chips. Until recently, this difficult process was accomplished through costly and time-consuming trial-and-error methods. But now, Vector Technology has developed software for Cray Research supercomputers that enables electrical engineers to simulate this process much more efficiently and cost-effectively.

The process of designing and producing computer chips involves the creation of delicate masks, which are then transferred to the surface of silicon wafers to form an integrated circuit. Creating masks is a complex and costly process. Production of a single standard mask can cost as much as $5,000, and a single phase shift mask can cost as much as $40,000. It is not uncommon for more than 20 masks to be required to produce one integrated circuit. The trial-and-error nature of this process can also cause delays in production. As a result, improvements in the photolithographic process that reduce the number of mask trials required can significantly reduce production costs and time to market.

Vector Technology's Fast Aerial Image Modeling (FAIM) software, makes these benefits a reality. A newly developed algorithm allows simulation of the photolithographic process for full-sized chips. On a 256-processor CRAY T3D system, FAIM can simulate 4-Mbyte DRAMS in less than one hour. Such a simulation might have required hundreds of years of CPU time to complete with the algorithms previously available to electrical engineers. With the new algorithms, work and storage requirements scale linearly with the number of processors to take full advantage of the massively parallel architecture of the CRAY T3D system. Performance of the FAIM software on the 256-processor CRAY T3D system has reached nearly 6 GFLOPS. This performance allows the electronics industry to perform simulations that simply cannot be performed in any other way.

For more information about FAIM on the CRAY T3D system, contact Mike Aamodt at 612/683-3665 or send email to maa@cray.com.

HSPICE version H95.2 released for Cray Research supercomputers

Version H95.2 of HSPICE, Meta-Software's industrial grade analysis code for the simulation of electrical circuits in the steady-state, transient, and frequency domains, is now available for Cray Research supercomputers.

HSPICE users with Cray Research systems will realize improved performance with this new version of a versatile electronics code. MOSFET models have been optimized and now run up to three times faster than the original code for these types of analyses.

For more information about the new release of HSPICE, contact Mike Aamodt at 612/683-3665 or send email to maa@cray.com.

Cray Research releases CRI/*Turbo*Kiva version 2.2

CRI/*Turbo*Kiva, Cray Research's enhanced engineering version of the KIVA-3 computer program for the calculation of transient, two- and three-dimensional, chemically reactive fluid flows with sprays is now available in version 2.2. CRI/*Turbo*Kiva can be used to simulate air flow, fuel injection, combustion, and pollutant emission in continuous combustors, as well as reciprocating two-stroke and four-stroke engines featuring multiple moving valves. Multiple fuel injectors with multiple injection rates may be used. With CRI/*Turbo*Kiva, more than 30 types of liquid and gaseous fuels may be used to study turbulent combustion, soot, carbon monoxide, and nitric oxide emissions. CRI/*Turbo*Kiva also features a powerful Motif-based input processor that greatly facilitates data preparation. CRI/*Turbo*Kiva relies on externally generated meshes.

For more information about CRI/*Turbo*Kiva, please contact Ting-Ting Zhu at 612/683-3631 or send email to tingting@cray.com.

Cray Research announces UniChem 3.0

UniChem 3.0 is the latest version of Cray Research's complete, easy to use, molecular modeling package. Highlights of this release include new functionality, increased accuracy and performance in the chemistry simulation codes, and major improvements to the user interface. These new features enhance UniChem's usability and extend the range of practical calculations. UniChem supports the leading quantum mechanics chemistry codes DGauss 3.0, MNDO94, CADPAC 5.2, and *Gaussian 92*. The UniChem package provides a graphical user interface for building molecules, initiating simulation jobs, and examining the results using 3-D visualization techniques. New features of DGauss 3.0 include the use of analytic second derivatives to improve performance and numerical stability, all new basis sets, NMR chemical shift calculations, and point charge fitting. DGauss 3.0 and MNDO94 now support solvent effect studies via a new reaction field model. Other enhancements to MNDO94 include support for four new transition metals, addition of a direct molecular dynamics method, and implementation of analytic gradients that improve accuracy and a tenfold performance improvement over previous methods.

Molecule building with the UniChem 3.0 user interface has been improved with the addition of 3-D scaffolding, a new 3-D cursor, and a one-step "undo" feature. A new "Quick Cleanup" feature optimizes molecular geometry by using an empirical force field valid for all atoms. Adding to the current support for X Window System graphics on SGI workstations, a version for the Sun Solaris environment has been developed. A 2-D plotting package and new animation features round out the UniChem 3.0 user interface enhancements.

For more information about UniChem 3.0, please contact Lee Roll at 612/683-3552 or send email to lee.roll@cray.com.

USER NEWS

EarthVision instructs students and teachers about environmental research

EarthVision is the only computational science program for elementary and high school teachers and students devoted exclusively to environmental research. Sponsored by the U.S. Environmental Protection Agency, EarthVision is a professional development and educational program for high school teachers and students. The two primary goals of the program are to train teams of teachers and students to do environmental research by using computational science and high-performance computing and to establish environmental research and computational science programs, as well as high performance computing capabilities at high schools. Teacher-student teams are provided with an intensive two-year academic program.

High schools receive a scientific workstation and a telecommunications link to the National Environmental Supercomputing Center (NESC) in Bay City Michigan. Each team is associated with the science community through mentors who are professional researchers from business and industry, academia, and government. These mentors work with the teams in the design and implementation of their research projects. Each team constructs scientific theory, translates the theory into a mathematical model, writes the code in Fortran to run the model in a high-performance computing environment, infuses data, and analyzes the output of the model by using scientific visualizations. Each team reports the results of its research in scholarly forums, professional journals, or meetings such as the annual IEEE Supercomputing conference, 14th World Congress of the International Association for Mathematics and Computers in Simulation (IMACS), and the Second International Workshop on Scientific Visualization.

Research activities at high schools

Following are summaries of the research activities at some of the EarthVision high schools.

The 1992-94 EarthVision team from the Center for Arts and Sciences High School in Saginaw, Michigan conducted a project titled "A Summary of Heavy Metal Distribution Data in the Saginaw Bay and a Computer Model to Interpret These Data." They hypothesized that the most important factors influencing the transportation of heavy metals include water direction and velocity, ion concentration, topography of the bottom, sediment type, and biota. During the investigation, the team of two teachers and four students created color-coded maps of existing historical data showing heavy metal distribution, determined how metal ions interact with water sediments and living organisms, created a mathematical model that described and visualized these interactions in the Saginaw Bay, and compared its model's predictions with actual historical maps.

The 1993-95 EarthVision team from Caseville High School in Caseville, Michigan conducted a project titled "A Comparison of Weather Data and Crop Water Demands for Shebeon-Kilmanagh Soils in Huron County." They hypothesized that, based on available weather data, a majority of current crop water demands will exceed available precipitation in Huron County. The team investigated at least 30 years of weather data and crop yields from Shebeon-Kilmanagh soil to determine if seasonal crop water demands exceed precipitation.

The Caseville team is working closely with the agricultural extension agent in their area and with local farmers, both of whom are interested in their research efforts. The results of the investigation are likely to have an impact on regional farming policy through the determination of the most effective crops to plant from both economic and ecological perspectives.

The 1993-95 EarthVision team from Saginaw High School in Saginaw, Michigan conducted a project titled "A Study of the Webber Street Combined Sewer Retention Basin." They hypothesized that the new Webber Street Sewer Retention Basin will prevent the release of untreated sewage from the basin's associated drainage area in the Saginaw River. Their research goal is to determine if the capacity of the new retention basin meets the NPDES requirement of the retention of a one-half inch, one-hour rain.

The Saginaw team is collaborating with the contracted engineering firm assigned to the design and construction of the Webber Street retention basin. Their calculations will be incorporated into the actual design of the facility.

Implementing education and computational science programs at the high schools

Each participating high school must develop an educational plan for the implementation of an environmental research and computational program at its school as well as conduct a specific research project of their own design. Central High School in Bay City, Michigan has established a School of Environmental Research (SER), which enrolled 60 students in the fall of 1994. The program is based on the EarthVision experience and uses an integrated curriculum, including science, sociology, English, computer science, and math. Students in the program design and conduct environmental research, which is evaluated, supported, and taught by the teachers in each of the preceding disciplines. Professional scientists provide external support and serve as mentors. They use the scientific workstations and have networked 100 PCs with X Window System terminal emulation software so that they can use AVS for visualization. They are connected, through the Internet, to the Cray Research supercomputer at the NESC. The administration of the high school has provided a common meeting time for the teachers to plan curriculum and activities.

The SER is a "school within a school." Central High School is an urban high school that has about 1400 students; eventually the whole school will participate. The district and the state are examining this model for replication in other settings. The implementation of the education plans is a key element in the EarthVision strategy to expand the impact of the program and to sustain, institutionalize, and establish environmental research and computational science at the high schools. Each school designs an educational plan based on the requirements and circumstances of its particular setting.

EarthVision originated in fiscal year (FY) 1993 as a pilot project, drawing participants in 14 counties surrounding the NESC in Bay City, Michigan. In FY 1994, participants were drawn from the entire state of Michigan; in FY 1995, EarthVision-Cleveland will begin, and participation will be extended from the Great Lakes region. It is anticipated that by 1996, nationwide participation will be possible. Participants have represented rural, suburban, and urban communities.

Modeling of a three-dimensional structure of the Arctic Ocean M2 tide with a high spatial resolution

A set of numerical experiments modeling the three-dimensional structure of M2 tide in the Arctic Ocean in a coarse grid (55.6 km) was performed by Igor Polyakov and Nicolay Dmitriev in the Arctic and Antarctic Research Institute, St. Petersburg, Russia. They concluded that continental slope with a large horizontal depth gradient plays a very important role in formation of tidal current vertical structure. On the other hand, a low spatial resolution did not allow them to reproduce effects of tidal energy trapping by irregularities of bottom topography. In 1993 Zygmund Kowalik and Andrey Proshutinsky (Institute of Marine Science, University of Alaska, Fairbanks) used a two-dimensional numerical model of the Arctic Ocean with a high spatial resolution (13.89 km) to demonstrate that in many locations of the Arctic Ocean, trapped waves of tidal origin are generated. These waves are a very important element of the ocean tidal dynamics; therefore, investigation of the three-dimensional structure of tides with high spatial resolution is advisable. This is a new approach for understanding real ocean processes.

A three-dimensional, primitive equations, numerical model is used to model tides. The model includes momentum, heat and salt balance equations, equations of state and continuity, and hydrostatic approximation. Processes of vertical turbulent exchange are described by variable coefficients related to vertical gradient of velocity and a scale of turbulence. The model has a free surface and a split timestep. The external mode is two-dimensional and uses a short timestep. It allows taking into account the dynamics of surface gravity waves. The internal mode is three-dimensional and uses a longer timestep.

The Arctic Ocean domain is approximated by a grid developed by Kowalik and Proshutinsky with a spatial resolution of 13.89 km. The model has 25 vertical levels. Tide forcing is done through the open boundaries of the computational domain and by astronomical forcing. The astronomical forcing includes not only tide-generating potential but also various corrections due to earth tide and ocean loading. As a first approach, the barotropic tide was simulated—temperature and salinity were constant in space and time. The simulation covered 40 tidal periods (approximately 20 days). During this time the problem has a quasi-stable solution. Cotidal maps (amplitudes and phases of sea level) and three-dimensional structure of currents for the semidiurnal principal lunar tide M2 were obtained as a result of the tide simulation.

A comparison model of results against the results of Kowalik and Proshutinsky's two-dimensional solution shows that both sea level data sets are practically identical. Modeling accuracy is estimated by comparison against 250 coastal and pelagic stations (some stations from the available database are not included in the comparison because they are located in the small bays and rivers not resolved by our grid step). To demonstrate the proximity of calculated and measured amplitudes, coefficients, intercept, and standard deviation of the linear regression line were used. The coefficients of correlation between two sets of data are calculated as well. Agreement is obtained for the coastal station with a coefficient of correlation equal to 0.95. Very good agreement also is achieved for the pelagic stations located in the North Atlantic where the correlation coefficient is 0.962.

The most interesting result is related to a structure of tidal currents. The model demonstrated that there are many regions with local enhancement of tidal currents. These are limited zones (about 200 km) located close to shallows of islands, banks, seamounts, and continental slopes with large depth gradients. This phenomenon may be observed in the M2 tide characteristics only to the north from a critical latitude (74.5 degrees for the M2 tide) and was explained (Kowalik and Proshutinsky) as the shelf waves that are trapped or partially trapped by the irregularities of the bottom topography. According to the simulation, the maximal vertical variability of tidal current occurs in these regions. On the other hand, in the deep portion of the Arctic Ocean, the tidal velocities are uniform from the surface to the bottom and close to a mean average value.

Future numerical experiments will consider real temperature salinity distribution, and ice cover. High spatial resolution will allow tide-ice interaction and internal waves generation to be described.

Acknowledgments

Nikolay Dmitriev, Igor Polyakov, and Andrey Proshutinsky gratefully acknowledge support for this work from Cray Research and Mike East, Paul Over, and Crispin Keable, all of Cray Research.

GALLERY

This skull image uses rendering software developed at the Arctic Region Supercomputer Center for their CRAY T3D system. A ray casting technique was used to visualize three-dimensional medical data obtained from magnetic resonance imaging. The software was designed and implemented by visualization manager Dr. John Genetti and research assistant Greg Johnson.