CRAY CHANNELS

WINTER 1991 · A CRAY RESEARCH, INC., PUBLICATION

Environment

CRAY CHANNELS

In this issue

Pilots, farmers, picnickers, and others long have dreamt of controlling the weather, but such a capability seems unlikely to emerge in the near future. Meanwhile, Cray Research supercomputers are helping forecasters predict the inevitable more accurately and precisely. This capability helps planners in agriculture, industry, and government allocate resources most effectively for irrigation, transportation, disaster relief, and other needs. And although direct weather control remains beyond human reach, human activities might affect the weather in unintented ways. Many researchers believe that emissions of carbon dioxide and other so-called greenhouse gases could raise atmospheric temperatures significantly in the coming decades. Other types of industrial emissions also affect atmospheric chemistry in ways that could have widespread ecological effects. Researchers in several countries are using Cray Research supercomputers to model the effects of industrial, automotive, and other chemical emissions on the biosphere to help understand the potential dangers that such emissions pose.

In this issue of CRAY CHANNELS, we survey applications of Cray Research computer systems in weather forecasting and climate research at several laboratories in the United States and Europe. Our regular departments include reports on Cray Research's new DD-60 and DD-61 disk drives, engineering and process modeling software packages, and some of the 1990 Cray Research Gigaflop Performance Award winners.

The need to balance industrial and environmental interests poses difficult problems for public policy makers. Cray Research computer systems provide the modeling capabilities needed to evaluate policy options quickly and economically. In addition to their long-standing uses in weather forecasting and climate research, Cray Research systems can play an important role in helping to formulate fair and workable environmental protection policies.

Features

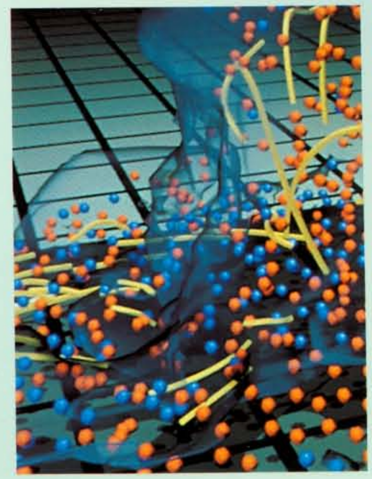

10

Departments

CRAY CHANNELS is a quarterly 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. Please mail feature story ideas, news items, and Gallery submissions to CRAY CHANNELS at Cray Research, Inc., 1440 Northland Drive, Mendota Heights, Minnesota 55120.

Volume 12, Number 4

Editorial staff
Ken Jopp, editor
Elizabeth Knoll, associate editor
Vicky Frank, contributing editor
Heidi Hagen, contributing editor
Rich Brueckner, contributing writer

Design and production
Barbara Cahlander
Eric Hanson
James Morgan
Cynthia Rykken

On the cover: The Earth's biosphere is a complex system of integrated chemical and physical processes. Short- and long-term fluctuations within this system can affect almost all spheres of human activity. To help anticipate nature's surprises and to predict our own impact on the biosphere, researchers are using Cray Research computer systems to model the Earth's atmoshphere, oceans, and the interactions between them. The timelier, more accurate weather forecasts and better understanding of pollution effects that result from this research provide economic and other benefits to all sectors of society. (photos, top to bottom: ©Craig Aurness, Bill Ross, Guy Motil/West Light)

Numerical weather prediction at the French Weather Service

Jean-Pierre Bourdette, Jean Clochard, Jean Coiffier, and Jean-François Geleyn, French Weather Service Paris, France

The French Weather Service, or Direction de la Météorologie Nationale (DMN, also known in France as Météo France), is well known in Europe for its pioneering activities in the area of numerical weather prediction. In 1983, a number of French universities and public research organizations including DMN formed the non-profit Center for Vector Computing in Research, (Centre de Calcul Vectoriel pour la Recherche, CCVR), a consortium that provides its members with a high-speed computing facility in Paris. Because DMN already had researched numerical weather prediction on the CRAY-1 system at the European Centre for Medium-range Weather Forecasts (ECMWF) in Reading, England, it was ready to implement a new numerical weather prediction suite on CCVR's CRAY-1 system in the spring of 1984. This system was replaced by a CRAY-2 system in 1986.

At the end of 1991, DMN will be moving its operational branch and CCVR's CRAY-2 system from Paris to Toulouse in the south of France. Because of the center's growing need for CPU time for operational weather forecasting and climate studies, its share of the CRAY-2 system has doubled in four years from 12 to 24 percent of system resources. In Toulouse, the CRAY-2 supercomputer will be fully dedicated to the wide range of meteorological applications at DMN, including forecasting and climate research.

The evolution of DMN's forecasting models

The numerical models EMERAUDE and PERIDOT, which DMN runs on CCVR's CRAY-2 system, evolved from pioneering work conducted during the 1960s, when computers became available for operational forecasting. EMERAUDE, DMN's large-scale numerical weather prediction system, and PERIDOT, a finer-scale system nested within EMERAUDE, each consist of an analysis model and a forecast model. During the analysis phase of the run, the initial situation for the forecast is computed on a regular three-dimensional grid from observational data collected from ground-based measurements, upper-air readings, and satellites, which provide a kind of mathematical snapshot of the atmosphere. Based on the results of the analysis, the model then produces forecasts of fields, such as surface pressure, temperature, wind, and humidity, at different times. Each forecast model is "re-entrant," that is, early-range forecasts are used as first guesses for new analyses in a cyclic process called "data assimilation." The quality and reliability of the data assimilation process greatly influences the meteorological quality of the whole system.

EMERAUDE's computational grid covers the whole Earth, and its forecast model uses a spectral technique with isotropic truncation at wave number 79 on 15 vertical terrain-following levels. The horizontal grid associated with the model has an approximate resolution of 1.5 degrees both in latitude and longitude. The analysis grid has a resolution of 1.5 degrees in latitude and 2 degrees in longitude and covers 16 constant pressure levels.

In contrast, PERIDOT covers an area of about 3300 by 3300 km including Western Europe and part of the Atlantic, a square projected on a stereographic representation, and uses a finite difference technique both in the horizontal and in the vertical dimensions. PERIDOT has 95 by 95 grid points, 15 vertical terrain-following levels, and a horizontal mesh size of about 35 km. Unlike EMERAUDE, PERIDOT's analysis and forecast models have the same horizontal and vertical grid resolutions.

The horizontal boundary conditions of PERIDOT are provided by EMERAUDE at every six-hour interval of the forecast. While EMERAUDE deals with large-scale phenomena in a very large horizontal domain, PERIDOT, the system nested within it, is dedicated to mesoscale modeling, which uses finer scales in a more restricted domain. Also, a version derived from EMERAUDE is being used as a climate model. Because of its high horizontal resolution and its sophisticated technique for data assimilation, PERIDOT is considered a pioneer mesoscale model. It also serves as the basis for a successful line of research models.

While the EMERAUDE and PERIDOT models were being developed to run on CCVR's CRAY-1 system, particular interest was paid to vectorization of the codes. The 1-Mword memory still presented a big constraint, and many codes ran out of core. The EMERAUDE model had its historical data packed and written/read asynchronously and concurrently on four work files that resided on separate disks assigned to different controllers. Despite this sophisticated I/O scheme, about one-third of the time required to run a job was spent waiting for I/O. PERIDOT had to pack its historical data to remain in core.

However, compared to earlier work on DMN's mainframe computer, the use of CCVR's CRAY-1 system provided a factor of 10 gain in CPU power. When the CRAY-2 system replaced CCVR's CRAY-1 system at the end of 1986, its four CPUs and 256 Mwords of common memory made it possible for the first time to leave historical data unpacked in central memory, initially by using this facility as a big buffer, then using it directly. Internal work files also disappeared. In addition, the large central memory made it possible to rewrite an extensive part of the EMERAUDE model, the Legendre transforms, using calls to the SCILIB matrix-multiply routine MXMA, which reduced CPU time by about 30 percent. Hardware updates and switching from version 2.0 of Cray Research's UNICOS operating system to version 5.0 corrected initial setbacks caused by longer vector loops and memory request conflicts. Recent measurements taken on the EMERAUDE model show a factor of 10.8 improvement in execution time with vectorization when excluding routines that make calls to the SCILIB routines RFFTMLT (Fourier transforms) and MXMA/MXVA (matrix multiplication).

The main codes also were modified to take advantage of multitasking on the CRAY-2 system. Each analysis or forecast model may use a number of logical CPUs determined at run-time by a directive file. For example, the operational EMERAUDE analysis uses four CPUs once daily and two CPUs in other operational runs. As a consequence of operational tuning that can be performed at the script level, the EMERAUDE model always uses two CPUs. The PERIDOT model, which uses the results of the EMERAUDE system as its boundary conditions, begins to run concurrently with EMERAUDE on two CPUs, and continues with four CPUs after the large-scale model has stopped running. Because PERIDOT requires more CPU time than EMERAUDE to reach the same forecast range, no waiting for synchronization occurs beyond the very beginning. Figure 1 shows the operational scheme used on the CRAY-2 system for daily forecasting.

Codes were adapted for multitasking by using the macrotasking facilities TASK, LOCK, and EVENT. Emphasis has been placed on assessing reproducibility of results, except for some printed diagnostics to guarantee bug-free codes, and to make meaningful comparisons in test mode possible. For both forecast and analysis models, reproducibility has been achieved for any number of logical CPUs, although the feature is optional for the PERIDOT analysis and this number has to be fixed for the EMERAUDE analysis. The large memory size of the CRAY-2 system has greatly facilitated the adaptations.

Dynamic partitioning of the workload has been used widely. Some parts of each code have decreasing granularity as the synchronization point approaches, which improves work balance and reduces overhead. In both forecast models, pairs of EVENTs rather than TASK features minimize overhead to achieve task synchronization.

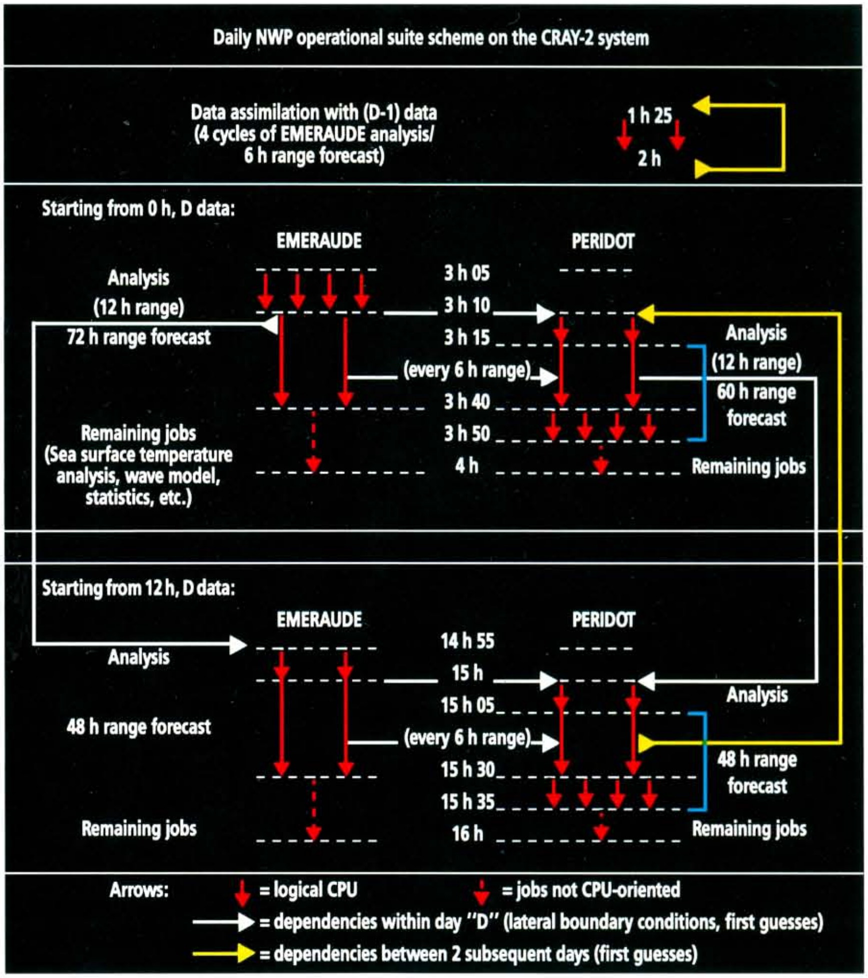

Figure 1. Daily numerical weather prediction operational suite scheme on the CRAY-2 system.

Using the ratio of CPU time to wall-clock time as a measure of multitasking efficiency, the following efficiencies were achieved: 3.91 for a single time-step of the EMERAUDE model, as measured within the program exclusive of system CPU time, and 3.72 for the main program of the EMERAUDE model. These efficiencies were achieved using four CPUs in a batch context where the job was the only user batch job, but with system jobs also present.

Research and development of future operational systems

The availability of a large memory on CCVR's CRAY-2 system allowed DMN to start designing an operational system that would make better use of this memory size and would benefit from the positive and negative aspects of the EMERAUDE and PERIDOT models. Launched as a project in 1987, the ARPEGE model is scheduled for full operational implementation in the spring of 1992 at DMN's new center in Toulouse. Because DMN and ECMWF are sharing a substantial amount of the preliminary software, the new system will be tested on both the CRAY-2 computer at CCVR and on ECMWF's CRAY Y-MP system.

The system of nested models in the previous systems is replaced in the ARPEGE system by a single global spectral model run on a fictitious sphere,

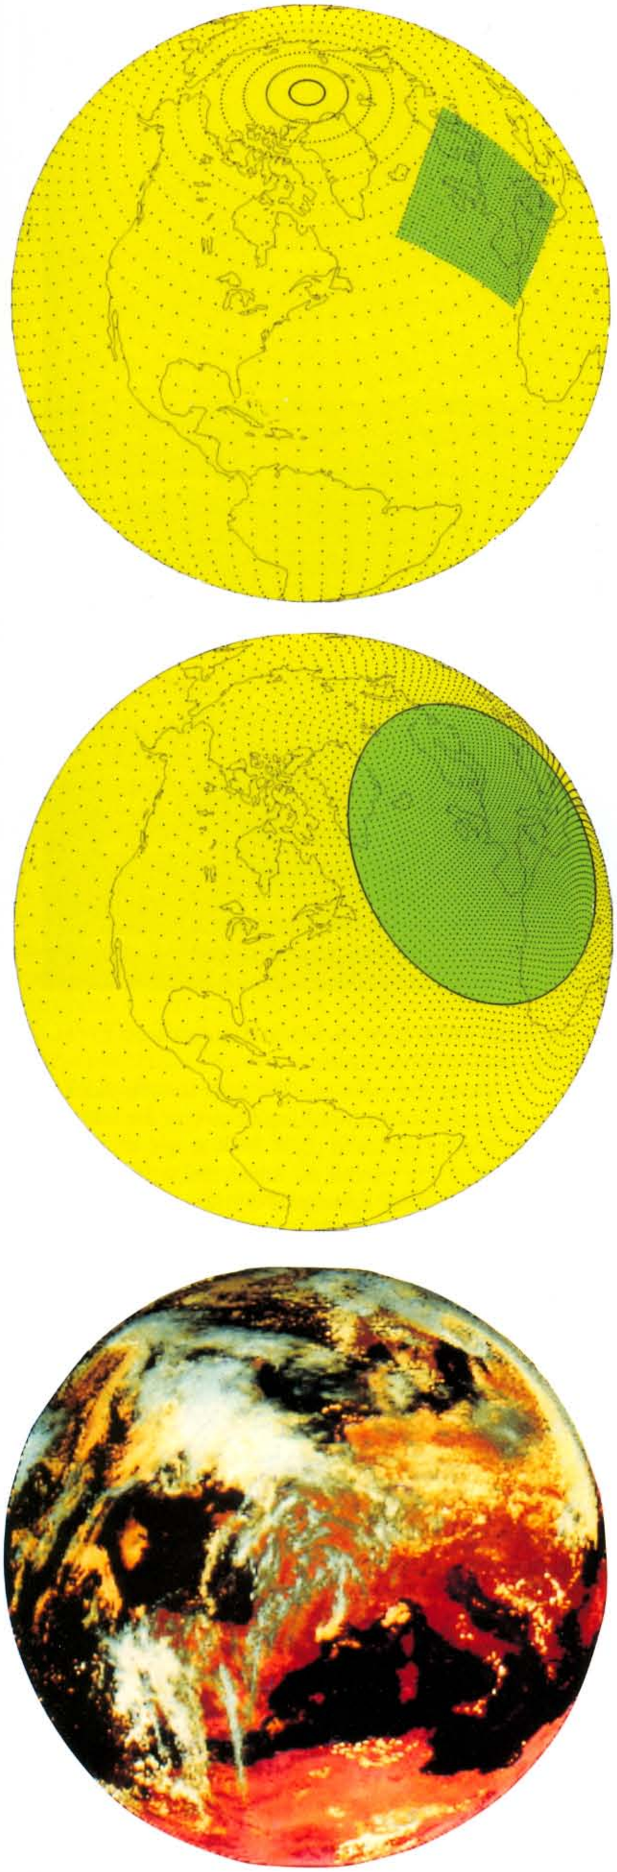

Figure 2. Computational grids for the EMERAUDE and PERIDOT models (top). The finer-resolution PERIDOT model is nested within EMERAUDE. The ARPEGE model, under development, replaces the regular grid of EMERAUDE with a variable-mesh grid that concentrates points around the dilatation point (center). Processed satellite image shows a pseudoEarth on which the ARPEGE grid is stretched to uniformity (bottom).

transformed from the real Earth through a conformal change of coordinates. This transformation expands the area around France, or any point of the globe, by an arbitrary factor, the antipode being contracted in the same proportion with a continuous transition in-between. In practice this amounts to replacing the classical regular Gaussian grid of spectral models like EMERAUDE with a variable-mesh grid that has points concentrated in the pseudo-hemisphere around the dilatation point. In addition, the poles of the computational grid are rotated away from the pole of dilatation, and a reduction in the number of points on each line of pseudo-latitude is performed when going toward these "computational poles" (Figure 2).

To avoid time steps that are too short because of this particular geometry, semi-implicit-semi-Lagrangian techniques are needed for the time integration problem. They will replace the classical semi-implicit techniques but are far more complicated to implement in spherical geometry than in plane geometry, even more so when the above-mentioned deformation is included.

Results of the physical routines, such as radiation; water phase changes; turbulence; moist convection; and surface exchanges of heat, moisture, and momentum, also are likely to be strongly influenced by the new geometry. A research effort aims at reducing this dependency as much as possible through more general and more integrated algorithms, especially for the hydrological cycle, and involves close cooperation with teams working in the research branch of DMN on mesoscale modeling. Simultaneously, efforts are under way to create a low-resolution, nondistorted version of the ARPEGE code as a basis for a French Climate Community Model, promising a rapid and positive scientific feedback to numerical weather prediction, especially in areas such as land surface

Figure 3: Surface absolute vorticity map after ten days of an idealized integration of a mid-latitude wave. Taking into account the variation with latitude of the Coriolis force leads to the formation of the more realistic so-called T-bone structure (in contrast with the classical "lambda" shape) for the frontal zones where vorticity is concentrating.

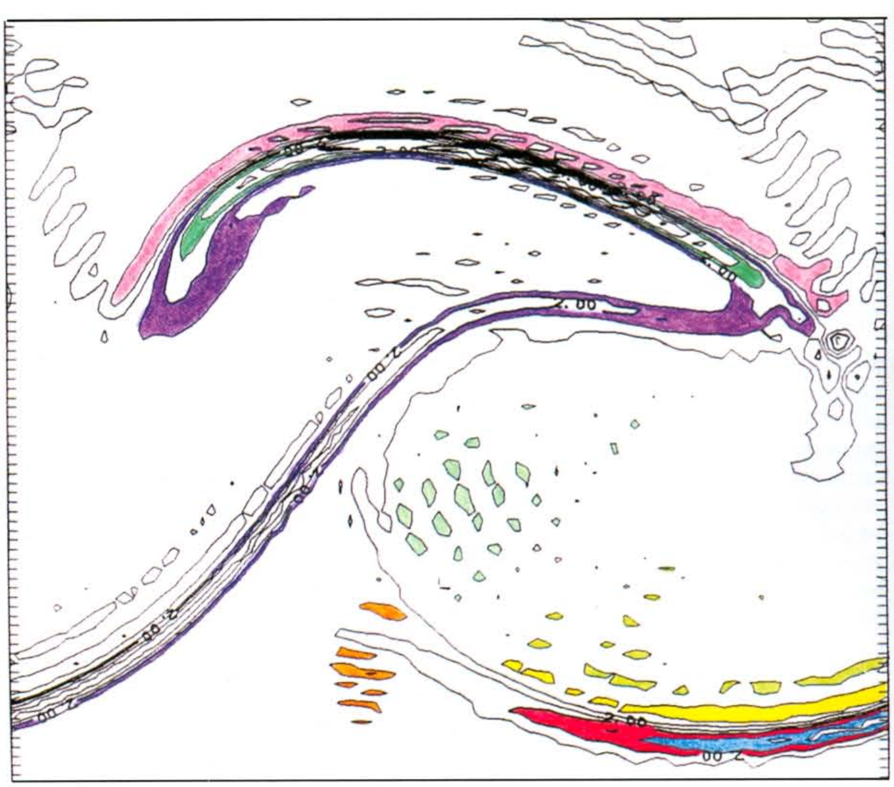

processes, radiative transfer, and, in the longer term, atmosphere-ocean interactions.

Initially, the data assimilation part of the project will be based on tools very similar to those used in EMERAUDE and PERIDOT. It will aim at an optimal blend of their characteristics and will give special attention to the new geometry (one single analysis code but with "randomly" spaced evaluation points). Subsequently, this system will evolve toward so-called variational approaches beginning with the intermittent mode — alternating modeling and analysis — but ultimately achieving a continuous mode. This evolution, and particularly the last step, requires the availability of the "tangent linear model" and the "adjoint model" of any direct computation. Both the tangent linear and the adjoint models of the direct code are undergoing parallel development within the ARPEGE project.

Finally, rewriting the code enabled further optimization for the CRAY-2 system, building on what already had been achieved for EMERAUDE. New optimizations include

- ☐ Totally dynamic memory allocation, which is useful when testing new ideas because it suppresses the need to recompile when parameters change, such as truncation or the number of vertical levels
- ☐ A full implementation of "local" use of automatic arrays based on specific routines as interface to implement these features
- ☐ A transfer of the relevant information into specially arranged arrays for computationally intensive operations, such as Legendre transforms
- ☐ Dynamic macrotasking with tasks ordered so that the smallest ones are last
- ☐ Modification of the MXMA operator to ensure vectorization along the longest loops in matrix product operations

For example, the increased cost of computing Legendre transforms due to longer vectors could be lowered from its theoretical threshold of n^3 to $n^{2.8}$, where n is the spectral truncation. The full adiabatic code has a dependency around $n^{2.16}$, which is bound to decrease even more toward n^2 as diabatic calculations are included. As a consequence the cost/benefit superiority of the spectral method against grid point/finite element methods can be maintained up to higher resolutions than currently assumed.

Although the ARPEGE code is still being developed, it has proven sufficiently robust to be used in an ongoing scientific study aimed at achieving a theoretical understanding of the formation of frontal structures in atmospheric flows and requiring flexibility of access to a "primitive equations" model as well as to efficient number-crunching performances for very high resolution calculations (Figure 3).

In looking back on the short history of numerical weather prediction, it appears that major milestones in supercomputing such as vectorization, parallel processing, and large memories, are correlated to advances in numerical weather prediction techniques. The theoretical limit of progress for these techniques is still far beyond reach, and many challenges lie ahead for both supercomputing and numerical weather prediction. ▬

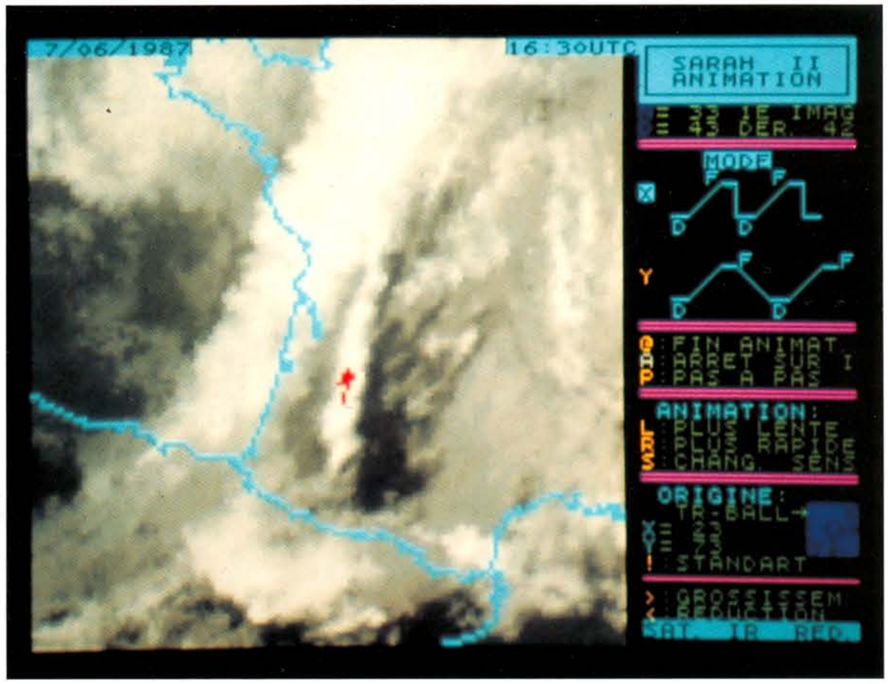

Processed images of remote-sensed (above) and model-generated data (right) of a squall line reaching the Southwest coast of France on June 7, 1987.

Acknowledgments

The authors thank all those who participated in the preparation of this paper: Eric Backelandt and Valérie Morcet for preparing the documents and drawing the figures, Jean Tardieu who produced the imagery, and Yves Ernie who coordinated the input for this work.

About the authors

Jean-Pierre Bourdette worked as head of the computer section at DMN for a number of years and is now deputy director of the operations department.

Jean Clochard is in charge of organizing and optimizing model codes on Cray Research supercomputers at DMN.

Jean Coiffier participated in the development of operational numerical models and is now head of the general forecasting branch at DMN.

Jean-François Geleyn worked at the European Centre for Medium-range Weather Forecasts for many years before managing the DMN research team for numerical weather prediction.

Simulation of the greenhouse effect with coupled ocean-atmosphere models

Ulrich Cubasch, Max Planck Institute for Meteorology Hamburg, Germany
Robert Sausen, Josef Oberhuber, and Frank Lunkeit Meteorological Institute of the University of Hamburg, Germany
Michael Böttinger, German Climate Computing Center (DKRZ), Hamburg, Germany

Industrial emissions of the so-called greenhouse gases, such as carbon dioxide, methane, and chlorofluorocarbons, have sparked numerous debates on the potential for global warming, which could occur as a result of these gases accumulating in the atmosphere. In several countries computer models are used to evaluate various policy options that pertain to greenhouse emissions and global warming. Researchers at the Meteorological Institute of the University of Hamburg, and the Max Planck Institute for Meteorology in Hamburg, Germany, are using a CRAY-2 computer system to run coupled ocean-atmosphere models to study the potential impacts of greenhouse emissions.

Until recently, such models contained only rudimentary representations of the ocean as a "mixed-layer" body of water with a maximum depth of 50 m. The models therefore failed to take into account the heat storage changes in the deep ocean and the changes in heat transport via oceanic currents that would be created by the climatic changes. For some of the more advanced models, the horizontal and/or vertical transportation of heat and salt was predetermined by climatology, but was not allowed to fluctuate in case of a climate change, thus severely biasing the climate response toward present conditions.[1]

In July 1990, a group of German scientists from the Max Planck Institute for Meteorology in Hamburg and the Meteorological Institute of the University of Hamburg, submitted an intermediate progress report to the German Ministry of Research and Technology on a new method for predicting climate changes induced by CO_2 emissions. In this cooperative study the previous models have been expanded to global coupled ocean-atmosphere circulation models. In addition to confirming some results of the mixed-layer model simulations, the inclusion of a comprehensive ocean circulation revealed a number of features not shown by the earlier models.

The models

The atmosphere model (ECHAM) was developed from the ECMWF forecasting model into a climate model.[2] Changes made to the original model include a decrease of the horizontal resolution to T21 (approximately 5.6 on a Gaussian grid) and a vertical resolution of 19 layers, a different radiation parameterization, the introduction of cloud liquid water as a prognostic variable, a more scale-selective horizontal diffusion, a ten-layer soil temperature parameterization with a zero flux assumption at a depth of 10 m, and a temperature-dependent snow albedo. Those parameters of the original model that remained unchanged were fine-tuned to produce an optimal fit with the altered parameterizations.

The first ocean model used in these experiments is the "large-scale geostrophic" (LSG) ocean general circulation model developed by Maier Reimer et al.[3] with 11 layers in the vertical, and a horizontal

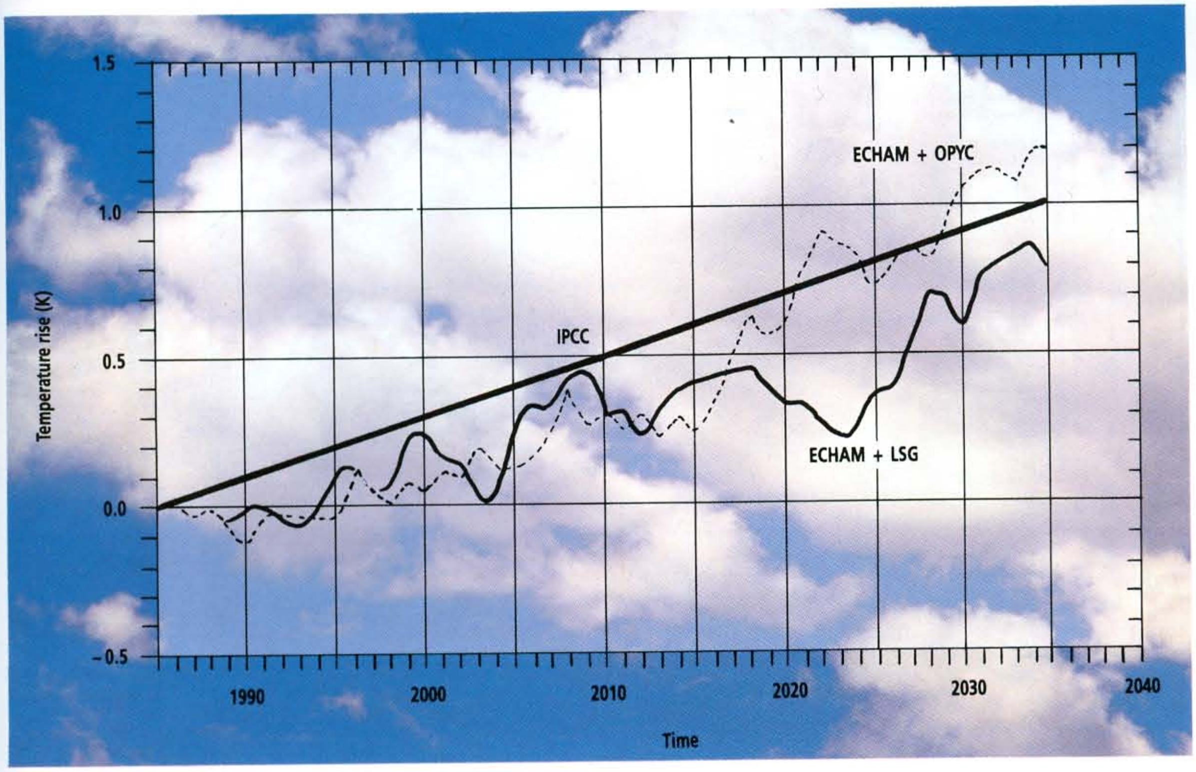

Figure 1. Global mean surface temperature change for scenario A for both coupled global ocean-atmosphere models, as well as the "best estimate" of the IPCC.

resolution chosen to fit the grid of the atmosphere model as closely as possible. It includes a thermodynamic sea-ice model.

The second is an "isopycnic" ocean model (OPYC) developed by Oberhuber[4] with a horizontal resolution of 2.8 degrees and 9 layers in the vertical. It runs on isopycnic coordinates and includes a prognostic mixed layer and a sea-ice parameterization according to Hibler (1979).[5]

The coupling

To reduce climate drift, both ocean models were coupled synchronously with the atmosphere model using the flux correction method.[6] The atmosphere model drives the ocean model with heat flux, fresh water flux, and wind stress, while the ocean model provides the sea surface temperature and ice distribution as input for the atmosphere model. In the coupled version, the atmosphere model includes a parameterization to calculate the temperature over the ice.

Coupled to both ocean models, the atmosphere model has been run with a time step of 40 minutes. A time step of one day was used for the top level of the LSG ocean model for the mixing of heat and fresh water flux, while the remainder of the model was run with a time step of 30 days. The OPYC ocean model used a time step of one day. In this combination, a one-year integration requires about 8 hours (12 for OPYC) of CPU time on one processor of the CRAY-2 system. The coupled model produces about 400 Mbytes (600 for OPYC) of data per year of integration time, which are stored in packed form (GRIB) on IBM cassettes.

The experiments

The experiments included a number of scenarios that eventually will be integrated into a prediction time of 100 years. Two of the four scenarios proposed by the Intergovernmental Panel on Climate Change (IPCC) were used. Scenario A, the IPCC "business-as-usual" scenario, assumes that the increase in CO_2 production continues at the same rate as in the past. In contrast, IPCC scenario D, the "accelerated policies" scenario, assumes a strong reduction in fossil carbon burning, and a resulting drastic reduction of greenhouse gas emission, so that a doubling of the pre-industrial CO_2 level will never be reached. In addition, one set of experiments was carried out assuming an instantaneous doubling of greenhouse gases. A total of four experiments was run:

- ☐ A 50-year control experiment, with CO_2 levels held constant
- ☐ A 50-year integration with IPCC scenario A
- ☐ A 30-year integration with IPCC scenario D
- ☐ A 30-year instantaneous doubling experiment

All experiments were run on the CRAY-2 system with both the ECHAM-and-LSG and the ECHAM-and-OPYC coupled ocean-atmosphere models and eventually will be expanded up to 100 years.

Summary of results

For scenario A both models predict that the global mean surface temperature will rise by approximately 0.8 K for the ECHAM-and-LSG model (1.2 K for the ECHAM-and-OPYC model) during the first 50 years (Figure 1). This value is approximately the same as that predicted with box diffusion models (IPCC). However, the initial increase is slower due to the different thermal structure of the deeper ocean and the more rapid convective transport of heat into the deeper ocean layers in the general circulation model (GCM) than in the box model.

Figure 2. Annual mean temperature difference between the control and the scenario A experiments for the years 41 to 50: cross section through the Atlantic for the ECHAM-and-LSG model.

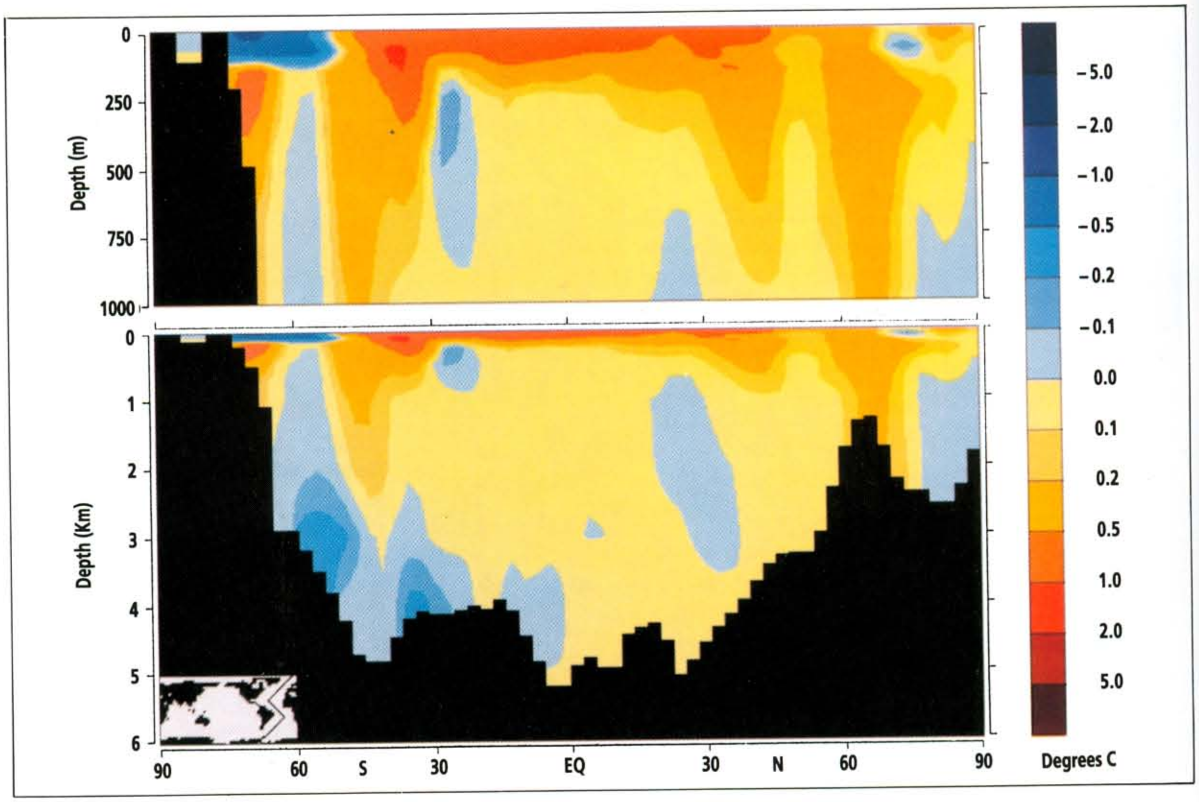

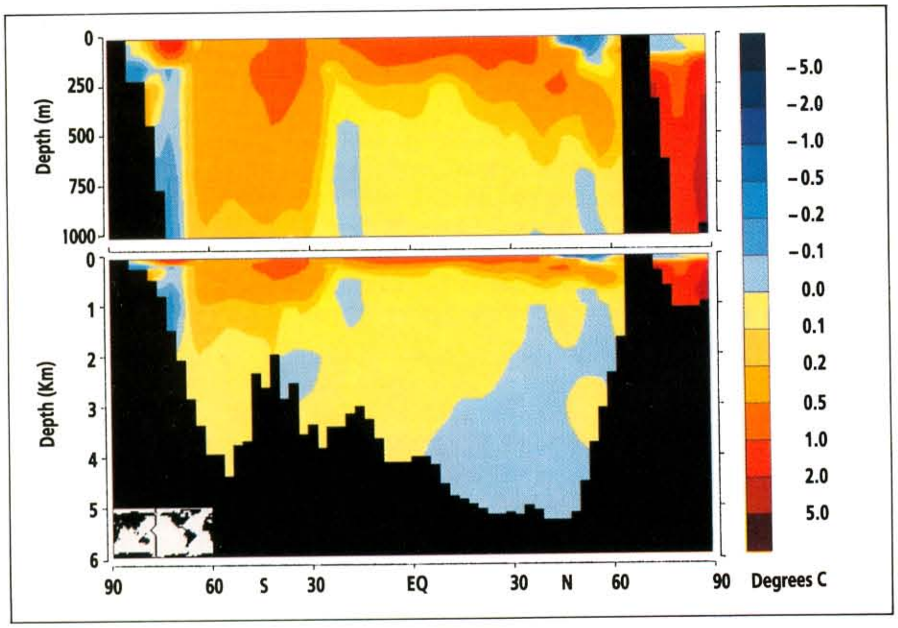

Figure 3. Annual mean temperature difference between the control and the scenario A experiment for the years 41 to 50: cross section through the Pacific for the ECHAM-and-LSG model.

A large portion of additional heat is sequestered by the ocean (Figures 2 and 3). After relatively short integration times an increase of ocean temperature is found in actively convective areas even at depths of 1000 m. Particularly in the North Atlantic, the ocean absorbs heat very efficiently. This casts serious doubts on results obtained with the models quoted in the IPCC report that represent the ocean with a depth of only 50 m. For the analysis of the deep ocean, box diffusion models are not capable of simulating this fast heat transport into the deep ocean and consequently do not reflect the delayed temperature rise.

Both models predict a hemispherically asymmetric warming of the surface (Figure 4). The middle and high latitudes of the Northern Hemisphere warm faster than the rest of the globe. In the Southern Hemisphere, regional cooling can be observed, which is caused by changes in oceanic circulation as a response to the changed distribution of heating by the greenhouse effect. Mixed layer model simulations (IPCC) predict an interhemispheric symmetrical warming with maximum warming occurring at the arctic and antarctic ice limits.

Maximal warming is observed in the troposphere and in the higher levels of the tropical troposphere, but the stratosphere becomes colder. A similar pattern can be found in calculations with mixed layer models (IPCC).

The tropospheric water vapor content is increased in both model runs, particularly in the tropics. This reinforces the greenhouse effect of CO_2 in a similar way as in mixed layer model calculations. Precipitation is increased slightly in the tropics. Soil moisture is changed significantly but is regionally inconsistent in both model simulations.

Both models predict a partial melting of the arctic ice, while the Southern Hemisphere exhibits no uniform behavior. During the 50-year period the mean sea level rises by about 4.5 cm with both models, though variations of up to 15 cm are observed regionally. The delay of the sea level rise is even stronger than the delay in the surface temperature. The sea surface has to become warmer before the deep ocean can warm up. Since the sea level change represents a vertical integral of the temperature rise, the inertia of the system delays the sea level rise twice. The long delay in the sea level rise could be simulated by an uncoupled ocean model driven by the heating predicted by mixed layer models.[7]

Differences between the models emerge mainly in the small-scale structure as well as in the

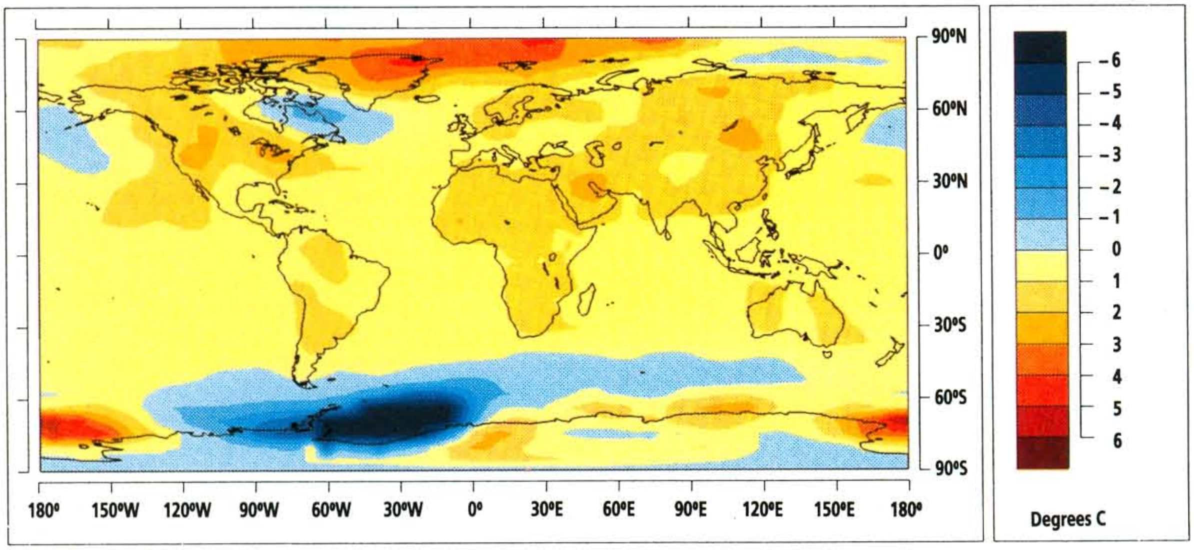

Figure 4. Annual mean surface temperature difference between the control and scenario A experiments for the years 41-50 for the ECHAM-and-LSG model.

amplitude of the pattern. For example, the tropical oceans are warmed more in the ECHAM-and-LSG model than in the ECHAM-and-OPYC model, while the second model predicts a larger warming in the polar region.

Scenario D produces results similar to those of scenario A. The global mean temperature rises as much as in scenario A, and the mean sea level rises by 1 cm for the time period integrated so far. Only the melting of the arctic ice is significantly reduced. The current prediction time of 30 years is obviously too short to show in detail the effects of a reduction in CO_2 emissions.

The CO_2 doubling experiments predict a larger climate change than scenarios A and D. For the ECHAM-and-LSG model we find an increase in surface temperature of 1.5 K (2.6 K for the ECHAM-and-OPYC model). This indicates that the signals of scenarios A will become stronger as integration time progresses. The ECHAM-and-OPYC model also has a higher sensitivity than the ECHAM-and-LSG model.

On a regional scale the spread between the results of both models becomes quite large, which causes severe problems for a regional climate forecast. Agreement between both models only can be found in the mean and on large interhemispheric structures.

Acknowledgments

The authors thank K. Hasselmann for supervising the project, the staff of the Max Planck Institute for Meteorology, the staff of the Meteorological Institute of the University of Hamburg, and the German Climate Computing Center (DKRZ) for their kind and patient support. This research is sponsored by the German Ministry for Research and Technology, the Max Planck Society, and the University of Hamburg.

About the authors

Ulrich Cubasch received degrees in meteorology (major) and oceanography (minor) from the Universities of Frankfurt and Kiel. From 1977 to 1987, he worked as a scientist at ECMWF in Reading, England, where he focused on numerical modeling and experimentation, as well as long-range weather prediction. Since 1987, Cubasch has been a climate research scientist at the Max Planck Institute for Meteorology in Hamburg. He specializes in the coupling of global ocean-atmosphere models and conducts climate change experiments. He is principal author of Section 3 (Processes and Models) of the IPCC report WG1 entitled "Scientific Assessment of Climate Change."

Robert Sausen received a degree in meteorology from the University of Darmstadt. From 1982 to 1986, he worked at the Max Planck Institute for Meteorology in Hamburg before joining the Meteorological Institute at the University of Hamburg as assistant professor. His work provides the theoretical basis for coupling ocean and atmosphere models.

Josef Oberhuber joined the Max Planck Institute in 1984 after receiving a degree in meteorology from the University of Munich. He has been working at the Meteorological Institute of the University of Hamburg since 1989, where he developed the isopycnic ocean model for use in climate change studies.

Michael Böttinger received a degree in geophysics in 1988 and has been working on the development of graphics software at the Max Planck Institute since that time. He has recently joined the German Climate Computing Center (DKRZ).

Frank Lunkeit is currently preparing his doctoral thesis on "Coupled Models and Scenario Calculations."

References

1. *IPCC 1990: Climate Change,* Cambridge University Press, Cambridge, p. 365, 1990.
2. Roeckner, E., L. Dümenil, E. Kirk, F. Lunkeit, M. Ponater, B. Rockel, R. Sausen, and U. Schlese, "The Hamburg Version of the ECMWF Model (ECHAM)," GARP Report No. 13, WMO, Geneva, WMO/TD No. 332, 1989.
3. Maier Reimer, E., K. Hasselmann, D. Olbers, and J. Willebrand, "An Ocean Circulation Model for Climate Studies," technical reports, MPI für Meteorologie, Hamburg, 1982.
4. Oberhuber, J. M., "Simulation of the Atlantic Circulation with a Coupled Sea-ice, Mixed Layer, Isopycnal General Circulation Model," 1990, to be submitted to the *Journal of Physical Oceanography.*
5. Hibler, W. D., "A thermodynamic sea-ice model," *Journal of Physical Oceanography,* Vol. 17, pp. 815-846, 1979.
6. Sausen, R., K. Barthels, and K. Hasselmann, "Coupled Ocean-Atmosphere Models with Flux Correction," *Climate Dynamics,* Vol. 2, pp. 154-163, 1988.
7. Mikolajewicz, U., B. D. Santer, and E. Maier Reimer, "Ocean Response to Greenhouse Warming," *Nature,* Vol. 345, pp. 589-593, 1990.

Environmental studies using Cray Research supercomputers at NCSA

Robert Wilhelmson, National Center for Supercomputing Applications, University of Illinois, Urbana-Champaign Illinois

The National Center for Supercomputing Applications (NCSA) was established in 1985 with a National Science Foundation grant to provide resources to the national research community to carry out computationally demanding research. Additional funding comes from the state of Illinois, the University of Illinois at Urbana-Champaign, and industry via industrial partnerships. The center has provided resources for over 5000 users at more than 200 universities and corporations. Currently, NCSA has two high-performance Cray Research systems, a CRAY-2S/4128 system and a CRAY Y-MP4/464 system, which was recently installed to replace a CRAY X-MP/48 system. These supercomputers are being used by researchers from a variety of disciplines including the environmental sciences. Here an overview of four environmental projects is presented.

Ozone: the chemical cycles

Atmospheric ozone forms a blanket around the Earth at a height of 10 to 30 miles and has the unique ability to absorb ultraviolet (UV) light from the sun. Without the ozone layer, life on Earth would be dramatically different because unblocked UV rays can kill plant and animal cells. Even small amounts of UV light may cause cell damage and induce skin cancer. The creation and destruction of ozone also is related to the presence of UV radiation. Ozone (O_3) can form from the combination of oxygen (O_2) molecules and unstable oxygen atoms freed from O_2 molecules by UV radiation. Similarly, UV radiation can break ozone apart. Through these and other reactions involving nitrous and hydrogen oxides, the ozone layer is maintained in natural chemical equilibrium.

Manmade chemicals that reach the stratosphere, however, alter this balance through depletion of some of the ozone with a resultant increase in O_2 molecules. For example, chlorofluorocarbons (CFCs), which are composed of chlorine, fluorine, and carbon, are extremely stable near the ground, but when they drift into the stratosphere where sunlight is intense, they break apart into highly reactive species, with the chlorine reactions accounting for approximately 30 percent of ozone depletion. Nitrous oxide (N_2O) diffusing upward from auto exhaust breaks into destructive nitrogen oxide (NO) in the stratosphere and accounts for an even greater percentage (50 percent) of ozone depletion. The discovery of the CFC problem in the late 1960s resulted in a ban on aerosol propellants containing CFCs.

The urgent need for a solution to the problem of pollutants and ozone depletion gained new momentum in 1985 when scientists discovered a large hole in the global ozone layer over Antarctica. It closes in the summer and fall but returns each spring in response to more intense UV radiation associated with the hottest months in the Southern Hemisphere.

The primary ozone depletion cycles in the stratosphere are the reactive chlorine (ClO_x) cycle, the reactive nitrogen (NO_x) cycle, and the reactive hydrogen or hydroxyl (HO_x) cycle. The species in these cycles may contain either one or two oxygen atoms, so x may be 1 or 2. Each cycle results in the destruction of ozone by converting it to O_2. But more importantly, each cycle operates catalytically — that is, each cycle concludes by again freeing the destructive radical to attack more ozone. For example, each chlorine molecule can destroy as many as 100,000 ozone molecules before it becomes inactivated by combining with a substance outside of the depletion cycle. (Chlorine is the free radical born of the breakdown of CFCs.)

Almon G. Turner, professor of physical chemistry at the University of Detroit, has turned his attention to refining current ozone depletion models using Cray Research supercomputers at NCSA. He believes an important step toward solving the dangerous ozone depletion problem is gaining a thorough understanding of the chemical reactions involved. His work will augment field observations and laboratory studies and aid researchers modeling the complex chemical interactions that occur in the stratosphere by providing additional information on the rates for these interactions. In addition, Turner and his colleagues calculate the physical properties of the reactant and product molecules.

Specifically, Turner and his associates use state-of-the-art quantum mechanical methods with GAUSSIAN88 software on NCSA's Cray Research systems to examine the stability of various molecules and the topology of the associated potential energy surfaces. With this information, Turner can calculate rate constants for proposed coupling reactions. The rate constant tells him how rapidly a reaction occurs with respect to other reactions in the cycle. If the rate is low, the reaction probably does not occur with enough frequency to influence the depletion cycle. If it is high, the reaction is likely to happen, and it becomes a candidate for further study. "We're using iterative computational techniques to study the chemical cycles of ozone depletion," says Turner. "You can't begin to do this without a supercomputer. The methods that we use lead to CPU-intensive calculations. The calculations also require large computer memory."

Turner and his research team are looking at not only the primary ozone depletion cycles, but also other reactions that can slow down the depletion. For example, molecules such as nitrosyl chloride (NOCl) and its valence isomer chloride nitrosyl (ClNO), may serve as "sinks" for the removal of active species in the cycles, or may themselves undergo reactions that modify the depletion of ozone. Sinks tie up chlorine and reduce ozone depletion. "From studying energetics we know NOCl is a sink," says Turner. "But we must also study the photodissociation cross sections to know the whole picture."

The photodissociation cross section examines the stability of a compound under bombardment by stratospheric sunlight. At ground level, NOCl would tend to tie up both free radicals NO and Cl, making them unavailable for ozone depletion. But the high-energy sunlight of the stratosphere may be able to break this molecule apart and free the radicals again for ozone destruction. Turner's goal is to determine whether this breaking, or photolysis, is occurring.

Rainfall: fluctuations over central North America

The need to document and understand the variability of climate is now well known. The occurrence, rate, and duration of precipitation in a continental region are related to the existence and strength of different atmospheric features and flows over the region and how they interact. Michael B. Richman and Peter J. Lamb from the Office of Climate and Weather Analysis at the Illinois State Water Survey in Champaign-Urbana are seeking to understand the intraseasonal and interannual fluctuations of growing season rainfall in central North America (between the Rocky and Appalachian Mountains and between the Gulf Coast and 55 degrees north latitude in Canada) in terms of the relative importance of several possible contributing climate system processes. These include the local presence and external supply of water vapor and the presence of vertical motion that triggers rain-producing clouds such as nimbostratus and cumulonimbus clouds, resulting in squall lines and tornadic storms. These features and flows are important in the development of physically based seasonal climate prediction schemes and for the generation of plausible scenarios for greenhouse-gas induced climate change projected to occur during the next 10 to 50 years. Central North American agriculture is important nationally and internationally, and projections of the impact of the region's climate change on agricultural production are very important.

In their studies, Richman and Lamb have identified characteristic patterns for the variation of rainfall in Central North America. According to Richman, Cray Research supercomputers at the National Center for Supercomputing Applications have been essential for manipulating the large amount of data involved and for performing the necessary statistical analyses. The database used was created especially for this project, as no pre-existing set had the fine temporal and spatial resolutions needed for treatment of the rainfall largely associated with convective (not stratus) precipitation. The raw data were obtained from the National Climate Data Center on over 100 magnetic tapes. These were downloaded into the Cray Research system's mass storage, and incorrect data were removed. The final data set contained daily rainfall totals for a gridlike network of 557 locations in the central United States and southern Canada (typical station separation was 65 miles) for the May to August periods of 1949 to 1980, which total more than 2.19 million observations and occupy 184 megabytes of storage. The subsequent research employed several advanced eigenvector and correlation techniques, available in the IMSL software package, to further compress the data set without sacrificing physically meaningful information on the key spatial scales related to studying rainfall variations. The use of the CRAY-2 system was vital in the diagonalization of the numerous 557 by 557 matrices generated. Linear transformations (rotations) were applied to the scaled eigenvectors to identify the individual spatial patterns in which rainfall was coherent over various totaling periods (3, 7, 15, and 30 days) for certain intraseasonal intervals (May, June, and August).

A key result of this work was the documentation of the morphology and scale of individual areas of coherent rainfall (it rains in all parts of the area sometime during any three-day period) as the growing season progressed. Richman and Lamb reported a decrease in the importance of the more organized,

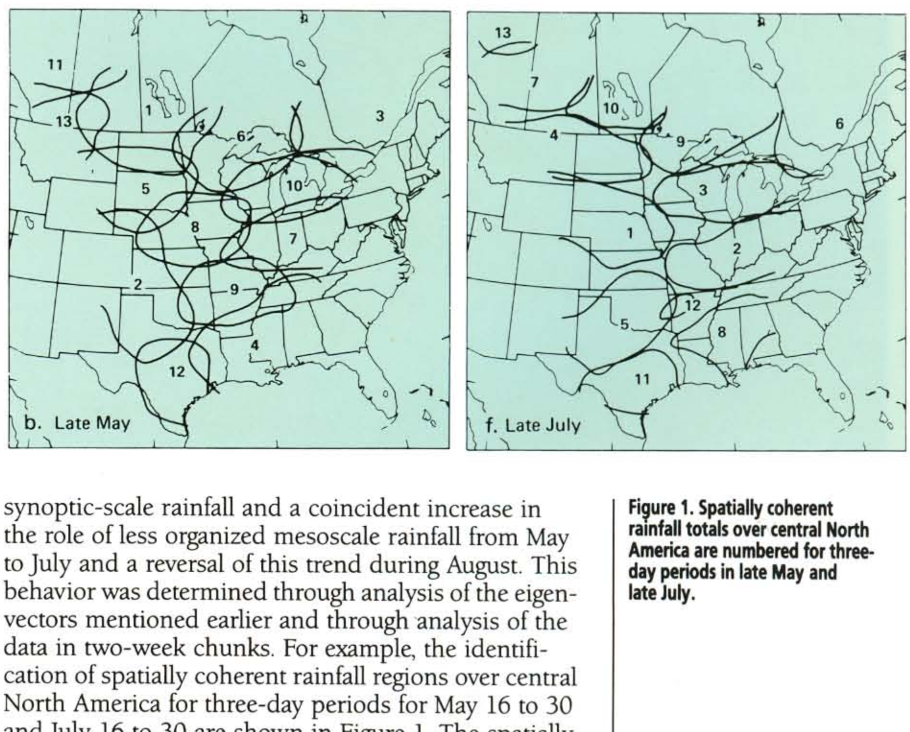

Figure 1. Spatially coherent rainfall totals over central North America are numbered for three-day periods in late May and late July.

synoptic-scale rainfall and a coincident increase in the role of less organized mesoscale rainfall from May to July and a reversal of this trend during August. This behavior was determined through analysis of the eigenvectors mentioned earlier and through analysis of the data in two-week chunks. For example, the identification of spatially coherent rainfall regions over central North America for three-day periods for May 16 to 30 and July 16 to 30 are shown in Figure 1. The spatially coherent regions are numbered. Noteworthy features identified for the first time include the decrease in regional overlap from May through July and the subsequent increase by late August, and the more irregular or elliptical orientation of the regions near the Great Lakes in July and August. These results have provided an important framework for the physically based investigation of the processes and conditions that contribute to the region's intraseasonal and interannual growing season rainfall variations.

Microbursts: an aircraft hazard

One of the more significant aviation hazards associated with thunderstorms is the generation of

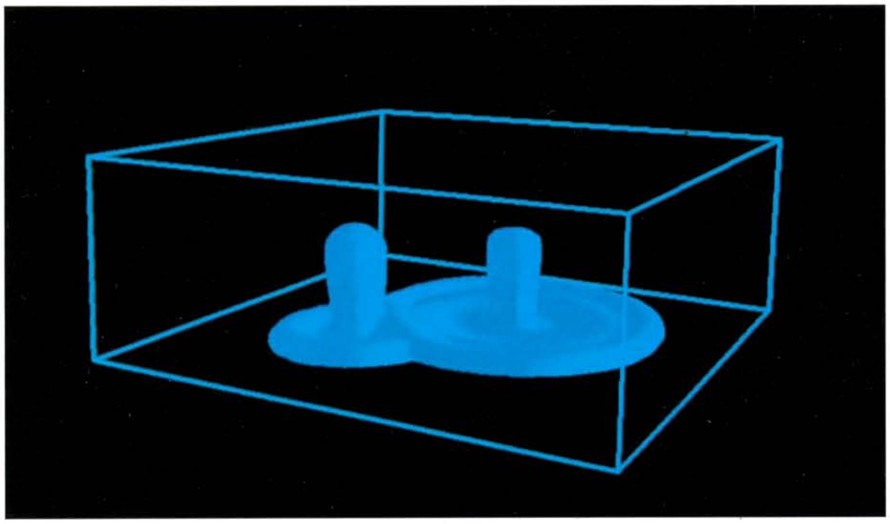

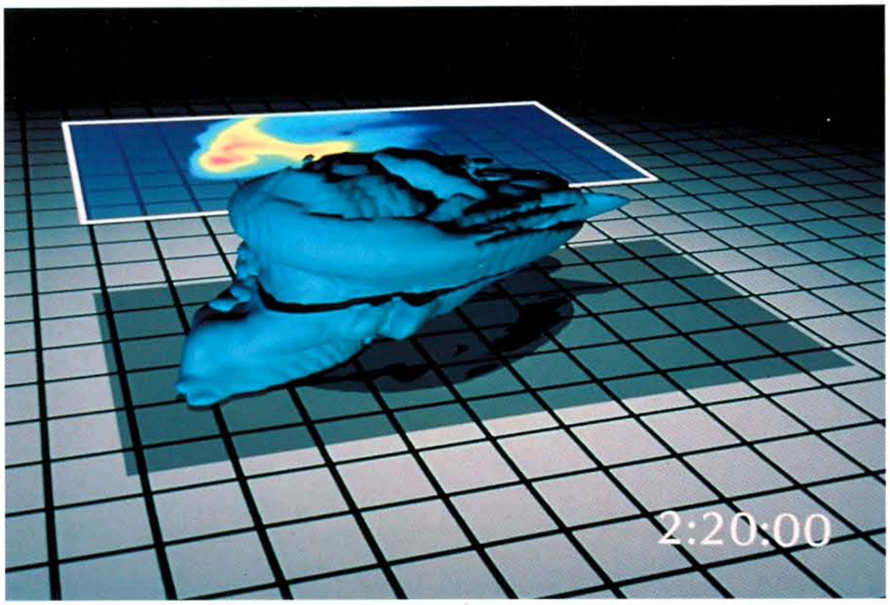

Figure 2 (top). The cold temperature regions within two downbursts that have collided near the ground are enclosed in the surfaces shown.

Figure 3 (bottom). The precipitation region within a modeled severe storm is shown during the storm's potentially tornadic phase.

small-scale, intense downdrafts, known as microbursts. The narrow downdrafts (less than 4 km wide) are formed in convective clouds between 2 and 4 km above the ground, and they descend to the ground in a short time during which they splash outward and generate surface outflows with wind velocities of up to 20 or 30 miles per hour. Sometimes downdrafts are wet (containing cloud or rain drops) while other times they are dry. When aircraft encounter these outflows during takeoff or landing they can experience an extremely rapid loss of airspeed, which results in the loss of lift and dangerous drops in altitude. This loss in air speed is a result of sudden wind changes.

One modeling effort being conducted by University of Wisconsin researchers John Anderson and Jerry Straka (now at the University of Oklahoma) using the Cray Research supercomputers at NCSA is aimed at understanding the environmental conditions that are conducive to microburst formation. These conditions include changes in temperature, pressure, moisture, and wind relative to altitude. They are used as input into a storm cloud model in which changes in temperature and other variables are predicted using time-dependent, spatially three-dimensional flow equations. These partial differential equations are used to compute changes to the initial input conditions every few seconds over a selected period of time and at selected points in space that form a three-dimensional grid. The motion of small and large drops, ice, and hail also are computed using parameterizations for their distribution and growth at each point because modeling the complete detailed behavior of liquid and solids in clouds is, as of yet, too computationally expensive.

Many cloud model simulations have been performed for environments that characterize clouds that do and do not produce microbursts. In some cases, model cloud structure has been compared against clouds observed with research Doppler weather radars. This past summer this was done predictively in collaboration with weather radar researchers from MIT-Lincoln Laboratory stationed in Orlando, Florida. Afternoon operational forecasts were made with the model using temperature, humidity, and wind data from a morning rawinsonde sounding with modifications to account for boundary layer mixing effects after the sounding and before clouds formed in the afternoon. The CRAY-2 system was used in the morning to make the model simulations, which take between 20 and 40 minutes of CPU time. "The forecast of observed microbursts based on these simulations has been encouraging," said Anderson. The simulations are being analyzed for further information helpful in distinguishing between microburst and non-microburst days, and for predicting maximum microburst outflow strengths.

A second effort related to microbursts involves exploring the hypothesis that the generation of elevated wind hazards may be due, at least in part, to the interaction between the outflows associated with more than one thunderstorm microburst. This would explain why a substantial number of radar and aircraft measurements indicate wind hazards above the typical 100 to 200 m (above the ground) danger zones associated with single microbursts. Further, colliding microbursts have been implicated in many of the air carrier accidents attributed to microbursts. Very high-resolution simulations of colliding microburst outflows are needed. These simulations can be achieved in an hour of CRAY-2 CPU time if microburst behavior is simulated only near the ground. Results from one simulation are shown in Figure 2 in which cold air within the microbursts is enclosed in the three-dimensional surfaces. Two dry microbursts were initiated at the cloud base at different times using prescribed heat sinks with magnitudes and spatial structures that approximate those in model simulations of microbursts in which the storm cloud is fully represented. The elevated cold air region between the microbursts is evident in the middle of the figure.

Results of this study indicate that the interaction region between the microbursts can indeed produce regions of intense wind shear and pose aviation hazards that can be even more severe than a single outflow.

Severe storm structure: an inside view

With the aid of supercomputers, four-dimensional storm models (three spatial dimensions and time) have been used to study the evolution of clouds from small cumulus to giant storms for the past two decades. As computer power and memory have increased, researchers have been able to simulate more of the detailed development of these storms. The resulting growth in the amount of data produced during simulations has stimulated exploration into ways of analyzing the data to better understand storm dynamics and physics. The wind, temperature, and moisture (or humidity) used to initialize the storm model in a recent study were taken from observed information near a severe storm that occurred in Oklahoma on April 3, 1964. This storm produced a tornado near Wichita Falls that killed 7 and injured 111 people and caused $15 million in damage. The storm (not the tornado) was simulated using 1.0 km horizontal and 0.5 km vertical separation between grid points within a 100-by-54-by-16 km domain. Using this grid domain, the storm model takes roughly the same amount of CPU time on the CRAY-2 system as the actual storm took to reach its mature stage (about two hours). The speed of the Cray Research system and its large memory enable rapid development and testing of new models and model physics.

In the April 3, 1964 simulation, the model data were stored every 12 seconds to study the evolutionary characteristics of the storm in detail. The resulting database consisted of 9 billion bytes of data. These data were used to visualize the structure and flow within the simulated storm in a collaborative effort that resulted in a seven-minute video entitled "Study of a Numerically Modeled Severe Storm" produced at NCSA. The animaions were performed with a Wavefront Technologies rendering package in conjunction with custom software developed within the group. The final version of the video took about 200 hours of computer time on a Silicon Graphics IRIS 4D/240 GTX workstation, using Wavefront software and special software written by the visualization group, and was the culmination of 11 months of work. The video tape accompanied a recent article by Wilhelmson et al.[1]

Different ways of displaying three-dimensional flows were explored using the data produced on the Cray Research system. The storm structure during its mature stage (the time when high surface winds or a tornado are expected) is shown in Figure 3. The precipitation field (filled with large raindrops, ice, and hail) is enclosed in the surface shown. The view is from the south and above the storm. The horizontal extent of the model domain, which moves with the storm, is represented at the surface by the medium gray region. The surface grid lines are spaced 10 km apart. The darker gray regions are shadows obtained by darkening the region directly below the storm. A two-dimensional slice from the precipitation region reveals the internal structure (qualitatively similar to radar reflectivity) at 4.75 km above the ground (highlighted on the precipitation surface by the black band). The interior red region contains the maximum concentration of rain drops in the two-dimensional slice. The model time for this and the following figures appears in the lower-right corner in hours, minutes, and seconds.

To understand the airflow within a simulated severe storm, weightless tracer particle motion is shown in Figure 4. This snapshot of storm-relative flow results from the frequent release of weightless tracers as the storm evolved from a horizontal plane that is 1.0 km above the ground. The tracers (balls) are initially spaced at 2.0 km intervals in this plane, and the color indicates whether the tracer is moving up (orange/red) or down (blue) at viewing time. This coloring technique can be used to follow wave propagation in a direction different from that of individual tracers when animated. Yellow ribbon tails are used on selected particles to help clarify the tracer paths being taken, representing the path over a 500-second period. The combined cloud-ice and precipitation surface (qualitatively similar to the visual appearance of a storm) is displayed using transparency so that the tracers can be seen both inside and outside the storm.

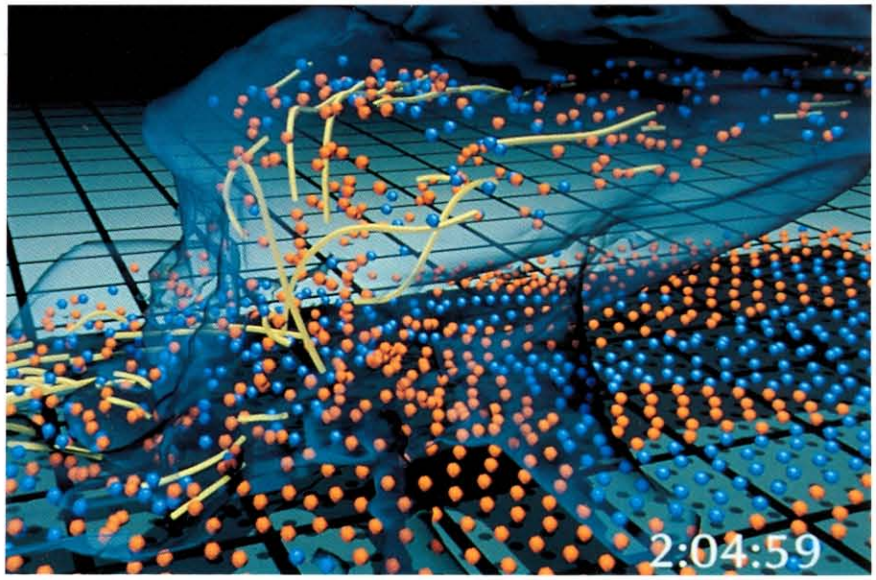

Figure 4 (top). Weightless balls move with the wind, depicting air motion within and surrounding a modeled severe storm.

Figure 5 (bottom). The ribbons depict the separation of upward (orange/red) and downward (blue) moving air within a modeled severe storm.

Tracers also can be released from selected points in the storm updraft and downdraft. In Figure 5 tracers along two lines are followed backward and forward in time with their paths indicated by the ribbons. Air rises upward after approaching the storm near the ground from the southeast (lower right), following the orange/red ribbons. Air entering the downdraft approaches the storm from the south and from several kilometers above the ground, wrapping around the

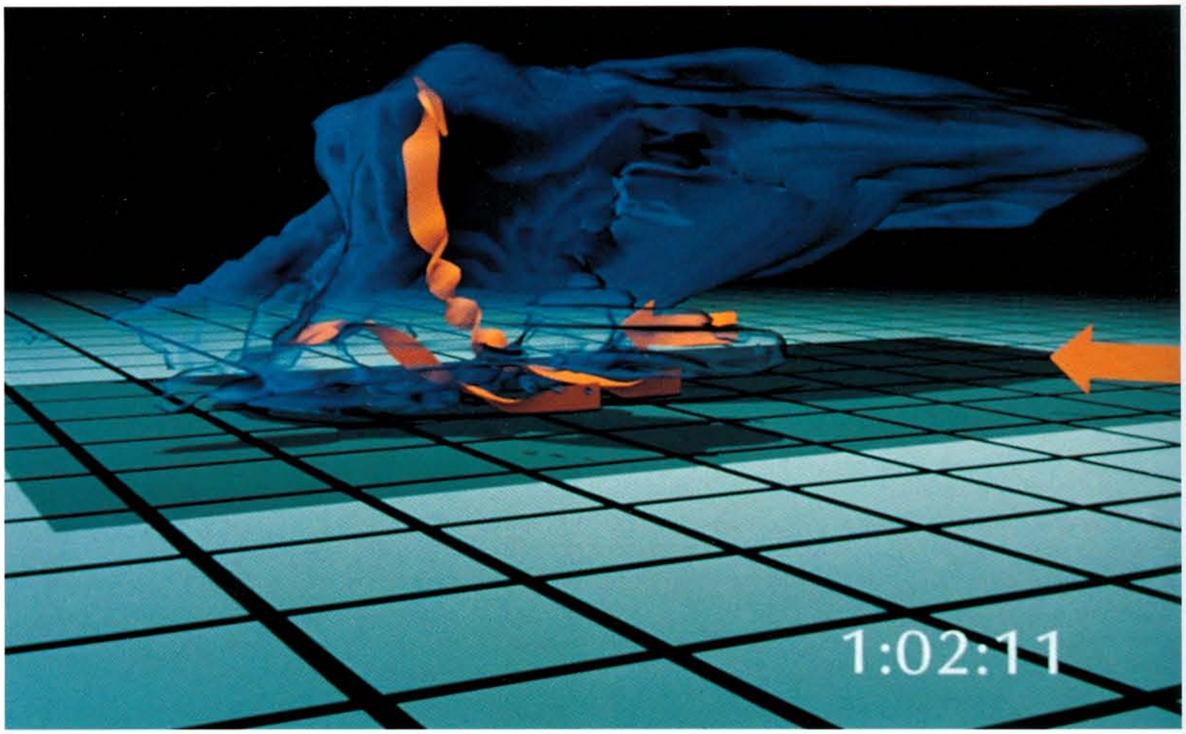

Figure 6. The twisting ribbons indicate the rate of rotation of air around the path of air particles moving through a modeled severe storm.

updraft as it sinks to the ground. The figure clearly indicates that the warm updraft and cool downdraft are intertwined as previously noted using wind data from a single time during the model simulation.

As tracers move along with the flowfield, it is important to be able to display changes in the physical properties of the flow. One such property is the vorticity, which is a three-dimensional vector representing the local rotation rate of the flow. The vorticity vector can be broken into two components, the "streamwise" component, which is parallel to the velocity vector, and the "crosswise" component, which is normal to and left of the velocity vector. The streamwise component can be directly related to the rotation of a perfectly spiraling football. This is important because it has been argued that an updraft fed with high streamwise vorticity can support a long-lived tornado-producing storm. The twisting ribbon in Figure 6 indicates the presence of substantial streamwise vorticity within the lower half of the storm updraft. The ribbon length corresponds to a 20-minute period. The amount of rotation indicates the relative magnitude of streamwise vorticity. Earlier in the life of the storm this ribbon was rotating very little as it approached the storm. This is also the case as it exits into the anvil.

The ability to perform a scientific study of storms and other flows using visualization and other software tools is being enhanced as technology improves. In the everyday working environment, software tools enable scientists to capitalize on these improvements to explore the growing amount of data from their simulations in a timely fashion. The video mentioned earlier took 11 months to complete from its inception. To analyze model data on a daily basis, a prototype software package for exploring storm behavior has been built at NCSA. It uses the CRAY-2 system and a Silicon Graphics workstation for analyzing and displaying data from an active model simulation or from previously stored data.

Acknowledgments

Scientific research for the video, "Study of a Numerically Modeled Severe Storm," was performed by NCSA and Department of Atmospheric Sciences staff including the author, Harold Brooks, Brian F. Jewett, Crystal Shaw, and Louis J. Wicker; and visualization staff at NCSA including Matthew Arrott, Mark Bajuk, Jeffrey Thingvold, Jeffery Yost, Colleen Bushell, Dan Brady, and Bob Patterson. The author thanks John Anderson of the University of Wisconsin's Space Science and Engineering Center, Jerry Straka of the University of Oklahoma's School of Meteorology, Michael Richman and Peter Lamb of the University of Illinois State Water Survey, and Almon Turner of the University of Detroit for their contributions. These researchers have used supercomputer time on NCSA's Cray Research supercomputers. Also, thanks to Randall Graham, who wrote an article about Turner's work for NCSA's Access newsletter that was used heavily in this article. NCSA is supported in part by the National Science Foundation.

About the author

Robert B. Wilhelmson is professor of meteorology in the Department of Atmospheric Sciences and a research scientist at the National Center for Supercomputing Applications (NCSA) at the University of Illinois. He received his Ph.D. degree in computer science from the University of Illinois in 1972, and has since been using supercomputers to simulate severe weather. In the mid-1980s he participated in establishing NCSA and served as assistant and associate director from 1985 to 1987. His interests include severe weather, the development and evaluation of numerical methods, and the development of an integrated environment for studying fluid flow problems involving parallel processing, distributed computing, databases, and visualization.

References

1. Wilhelmson, R. B., B. Jewett, C. Shaw, L. Wicker, M. Arrott, C. Bushell, M. Bajuk, J. Thingvold, J. Yost, "A Study of the Evolution of a Numerically Modeled Severe Storm," *International Journal of Supercomputing Applications*, MIT Press, 4 (2), pp. 20-36, 1990.

Shallow sea hydrodynamic models in environmental science

Alan M. Davies and Roger Proctor
Proudman Oceanographic Laboratory
Birkenhead, Merseyside, England
Mike O'Neill, Cray Research, Inc.

Figure 1. Historic lighthouse adjacent to Proudman Oceanographic Laboratory building.

The Proudman Oceanographic Laboratory (POL) evolved from the Liverpool Observatory built on Bidston Hill in 1866. Early work involved astronomical measurements to determine an accurate time origin for the setting of chronometers, which were used by early mariners to aid navigation on their voyages from Liverpool to North America. Although the telescopes are long gone, the domes in which they were placed remain, and the observatory and adjacent buildings are still local landmarks (Figure 1).

Tidal prediction

With the arrival of accurate timepieces, the astronomical work was no longer required. However, the laboratory remained involved in maritime activities through predicting the tides caused by the gravitational forces of the moon and sun. Tidal predictions were accomplished using the harmonic method. Observed tides at a port are decomposed into constituent parts (tidal harmonics) related to the astronomical motions of the Moon-Sun-Earth system. Once these harmonics are determined (an amplitude and phase lag for each tidal constituent) tidal predictions can be made for that port. The predictions were calculated first by hand, then by using an analog computer, a tide-predicting machine consisting of pulleys, a continuous wire, and a pen; this produced a time series of tidal heights on a drum recorder. Up to 42 tidal constituents would be used, with a year's tidal prediction completed in about 1.5 days. Using modern digital computers the number of constituents has been increased to above 100, significantly improving the predictions, particularly at shallow water ports. The time required to complete the predictions is now less than one minute.

Numerical modeling of storm surges

The widespread availability of digital computers and increased understanding of finite difference methods in the 1950s led scientists to predict the movement of the sea surface, under the influence of tidal and wind-driven forcing, by solving the hydrodynamic equations that describe the motion of the sea. Although this motion is three-dimensional, it is necessary to compute only the depth mean flow to determine changes in sea level.

The major coastal flooding of southeast England and the Netherlands during January and February 1953, caused by exceptionally strong northerly winds acting over the North Sea, spurred efforts to develop storm surge prediction models. These sea models, when forced by predictions of wind and surface pressure from an atmospheric model, give advanced warning of potential flooding problems at the coast. Research on these storm surge models has proven beneficial. The models have been validated successfully by hindcasting many storms, including the January-February 1953 surge[1] and surges in the Bristol Channel.[2] The models also can reproduce accurately the storm surge elevation (the height of sea level above the predicted tidal level) along the east coast of England and the European continental coast.

A storm surge model developed at POL has been run routinely at the U.K. Meteorological Office for the past 12 winters. This model, running twice daily, covers the northwest European continental shelf and uses predictions of wind and sea level pressure from the Meteorological Office atmospheric 15-level model to predict storm surges up to 36 hours in advance. The storm surge level at high water was predicted successfully during a severe storm on February 26, 1990, which resulted in extensive flooding at Towyn on the North Wales coast and the evacuation of 2000 people. Observed levels during this storm reached those expected only once in 200 years, and the flooding was compounded by large waves damaging sea

Figure 2 (right). Surge heights in meters and corresponding surge currents at 12 p.m. on February 26, 1990.

Figure 3 (far right). Significant wave heights in meters and corresponding wave directions at 12 p.m. on February 26, 1990.

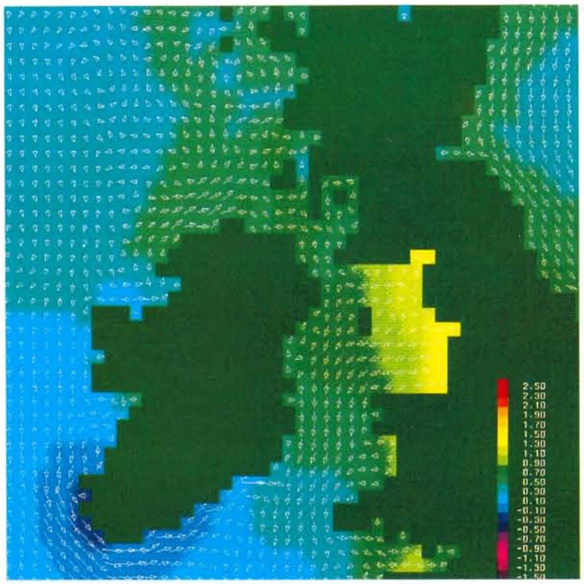

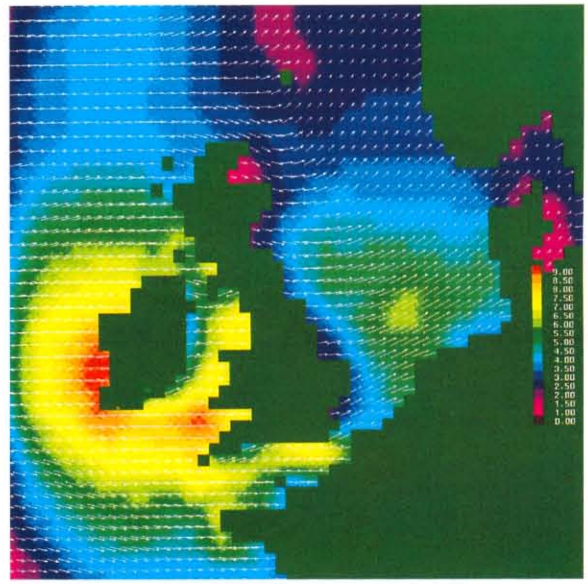

Table 1. Performance of the three-dimensional model on CRAY X-MP/48 and CRAY Y-MP/832 systems using masking and strip mining. CPU and wall-clock times are in seconds.

	CRAY X-MP/48 microtasked for four processors			CRAY Y-MP/832 automatically multitasked for eight processors		
	CPU	Clock	MFLOPS	CPU	Clock	MFLOPS
Masking	1269	318	548	780	98	1768
Strip mining	1298	325	449	1157	146	999

Masking

```
CMIC$  DO GLOBAL
       DO 1 J = 1, M
       DO 2 I = 1, N-GRID—POINT
          memory access    (MASK(I)*FIELD-ARRAY(I,J))
       2 CONTINUE
       1 CONTINUE
```

Strip mining

```
CMIC$  DO GLOBAL
       DO 1 J         = 1, M
       DO 2 I-STRIP  = 1, N-STRIP
       DO 3 I         = I-START(I-STRIP), I-END(I-STRIP)
          memory access    (FIELD-ARRAY(I,J))
          3 CONTINUE
          2 CONTINUE
          1 CONTINUE
```

Table 2. Programming strategy for masking and strip mining the three-dimensional model.

defenses. The computed storm surge at high water at 12 p.m. on February 26, 1990 shows surge heights in excess of 1.5 m in Liverpool Bay (Figure 2). Wave model calculations of significant wave heights are shown in Figure 3 for the same time.

The surge prediction model is modified when changes may enhance the quality of the predictions. This winter, the surge model will be run on a CRAY Y-MP8/832 system. Higher resolution meteorological forcing and a refinement of the sea model grid should bring improved surge forecasts to the areas around the U.K. most susceptible to coastal flooding.

Three-dimensional models

With advances in computing memory and processing power over the past 10 years, solving the full three-dimensional hydrodynamic equations is now possible. Also, in recent years interest has increased in a number of environmental problems such as oil spill prediction, which requires time-evolving surface currents, and movement of bed sediments, which often have toxic heavy metals adhered to them and require time-evolving near bed currents. These problems necessitate the computation of the full three-dimensional flow field.

A computationally efficient means of solving these equations[3] is to use a finite difference grid in the horizontal and a spectral/modal approach in the vertical.[4,5] Such a representation is ideal on multiprocessor vector computers, such as the CRAY X-MP or CRAY Y-MP series of supercomputers. A high level of parallelism is achieved on these computers by assigning a processor to each vertical mode and integrating forward in time. In the horizontal domain, a long vector length and maximum performance of more than 1.7 GFLOPS is achieved by performing the calculation over both land and sea regions (Table 1), then masking off the land (Table 2). However, when the land region is a considerable fraction of the domain, this approach uses excessive computer time. Only performing the calculation over each sea strip reduces the computation time even though the processing speed is reduced.

Wind-induced currents

This mixed finite difference/modal approach has been used very successfully in a number of three-dimensional simulations of tidal- and wind-induced flows. For example, the strong surface and bottom currents driven by the intense westerly winds associated with storm surges off the west coast of Britain are shown in Figures 4 and 5. Figure 4 shows that surface currents are aligned in the direction of the wind; and hence any surface oil spillage in the Liverpool Bay area at that time would be rapidly moved on shore. In

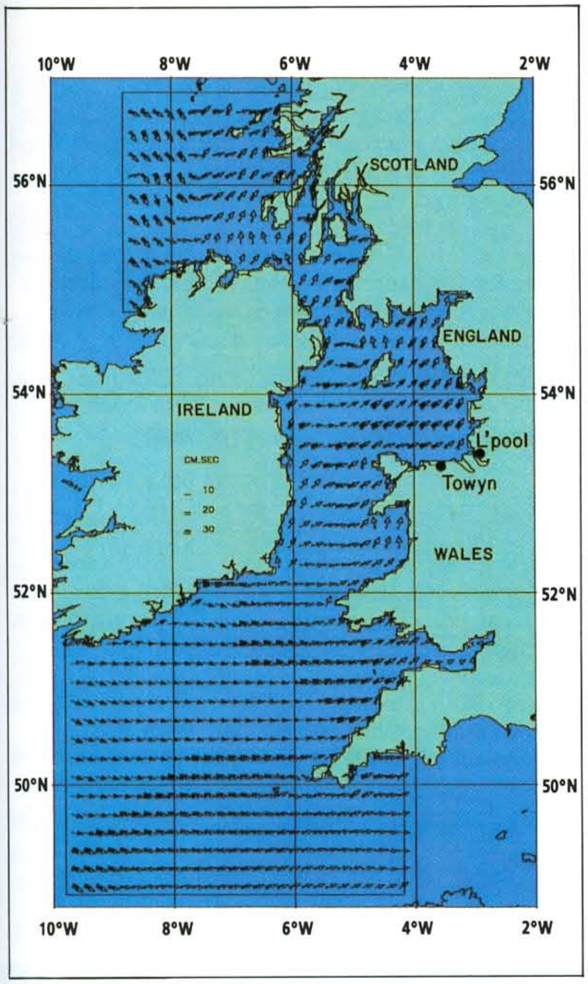

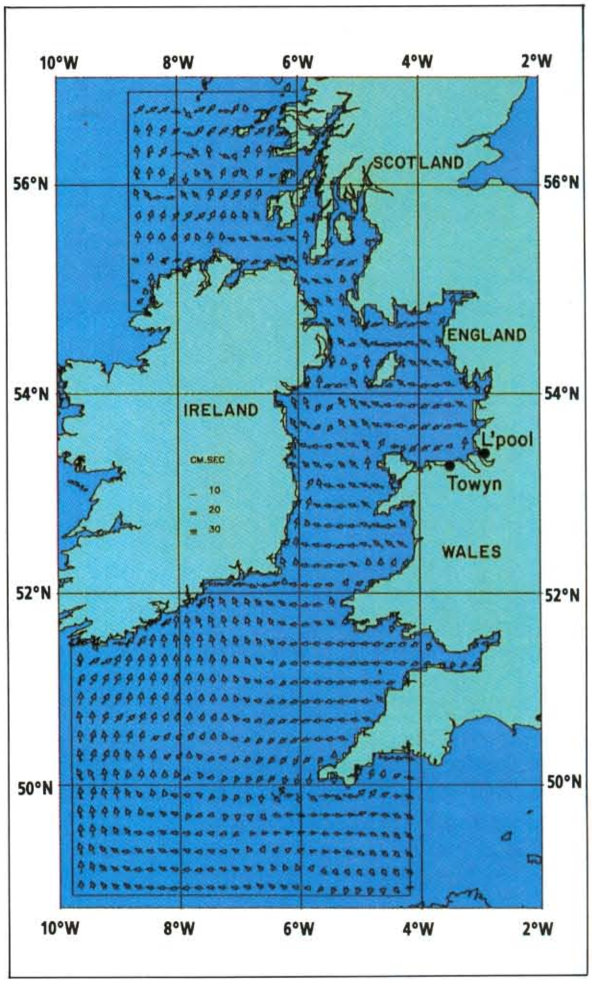

Figure 4 (far left). Computed wind-induced surface currents at 6 p.m. on November 11, 1977.

Figure 5 (left). Computed wind-induced bed currents at 6 p.m. on November 11, 1977.

Liverpool Bay, down-welling occurs against the English coast, with offshore currents at depth in the opposite direction to the wind (Figure 5). These bottom currents are strong and readily can transport sediments or dumped material along the sea bed.

In general, results from the three-dimensional models agree quite well with observed currents. In the future more complex physical processes within the models will be necessary, in particular wave-current interaction and turbulence energy closure schemes, both of which are important in determining the movement and fate of pollutants. Improving the physics in the model necessitates grid refinement, and such developments will increase the computational demands of the models in terms of main memory and processor time.

About the authors

Alan Davies, a physical oceanographer at the Proudman Oceanographic Laboratory, develops and applies three-dimensional hydrodynamic models. He received a B.Sc. degree in chemistry in 1972 and a Ph.D. degree from Sheffield University for work on simulating the electromagnetic properties of small molecules. He recently edited a book on numerical modeling in physical oceanography, and jointly with R. B. Grzonka was presented with a 1990 Cray Research Gigaflop Performance Award for work on optimizing a three-dimensional shallow sea model.

Roger Proctor is a physical oceanographer at the Proudman Laboratory, with interests in tides and storm surge forecasting and three-dimensional circulation modeling. He received a B.Sc. degree in mathematics from Teesside Polytechnic in 1975 and a Ph.D. degree from Liverpool University in 1981 for work on modeling the long-term circulation in the Irish Sea.

Mike O'Neill is a consultant analyst with Cray Research, U.K., Ltd., where he is a specialist in the application of parallel processing. He began his career in computing at the Institute of Computer Science at London University in 1965 and has worked in the field of large-scale parallel array and vector processing since 1972.

References

1. Flather R. A., "A Numerical Model Investigation of the Storm Surge of 31 January and 1 February 1953 in the North Sea," *Quarterly Journal of the Royal Meteorological Society,* Vol. 110, pp. 591-612, 1984.
2. Proctor, R. and R. A. Flather, "Storm Surge Prediction in the Bristol Channel — the Floods of 13 December 1981," *Continental Shelf Research,* Vol. 9, pp. 889-918, 1989.
3. Davies, A. M., R. B. Grzonka, and C. V. Stephens, "The Implementation of Hydrodynamic Numerical Sea Models on the CRAY X-MP," to appear in *Advances in Parallel Computing,* Vol. 2, edited by D. J. Evans.
4. Davies, A. M., "Formulation of a Linear Three-Dimensional Hydrodynamic Sea Model using a Galerkin-Eigenfunction Method," *International Journal of Numerical Methods in Fluids,* Vol. 3, pp. 33-60, 1983.
5. Davies, A. M. and R. Proctor, "Developing and Optimizing a 3D-Spectral/Finite Difference Hydrodynamic Model for the CRAY X-MP," *Computers and Fluids,* Vol. 18, pp. 259-270, 1990.

The Regional Acid Deposition Model

A new tool for policy assessment

Henry Lansford, Atmospheric Sciences Research Center
State University of New York at Albany
Hugh Patrick, Cray Research, Inc.

Among the environmental issues that concern public policy-makers, atmospheric acid deposition continues to be one of the more vigorously debated at every level of society. Most differences of opinion among environmentalists, industry representatives, and government officials concern the type, level, and timing of controls to be imposed on the chemical emissions that are the major causes of acid deposition. For policy assessment, the most important species are compounds containing sulfur, nitrogen, and organic compounds emitted primarily from coal-fired power plants but also from automobiles and factories. Some of these compounds remain unchanged in the atmosphere and some are neutralized, but others are oxidized into more acidic forms through a complicated series of chemical, meteorological, physical, and biological interactions.

Acid deposition refers to the process by which these acidic compounds are transported in the atmosphere, then deposited on the Earth's surface. To better understand acid deposition, the U.S. Environmental Protection Agency (EPA) established an atmospheric modeling project in 1983. The project researchers developed the Regional Acid Deposition Model (RADM) on the CRAY X-MP/48 computer system at the National Center for Atmospheric Research (NCAR) in Boulder, Colorado. The model already is being used in policy analysis and will be used more extensively as it becomes refined through comparisons with field data.

Increasing acidity increases problems

The atmospheric processes that produce acidic compounds are driven by sunlight and fueled by pollution. The process begins when a photon of sunlight cleaves an ozone molecule (O_3) into a diatomic oxygen molecule (O_2) and a single oxygen atom. The oxygen atom combines with a water molecule to form two highly reactive hydroxyl radicals, each of which can transform one molecule of the pollutant nitrogen dioxide (NO_2) into nitric acid or catalyze the oxidation of thousands of sulfur dioxide (SO_2) molecules into aqueous sulfuric acid.[1]

Although the term "acid rain" is used more widely than "acid deposition," only about half of acidic deposition is in the form of aqueous precipitation. The remainder occurs as dry deposition in the form of small particles or gases comprised of unreacted pollutant. Dry deposition usually occurs within 480 km

Flowchart of atmospheric chemistry involved in acid deposition. Adapted from Reference 1.

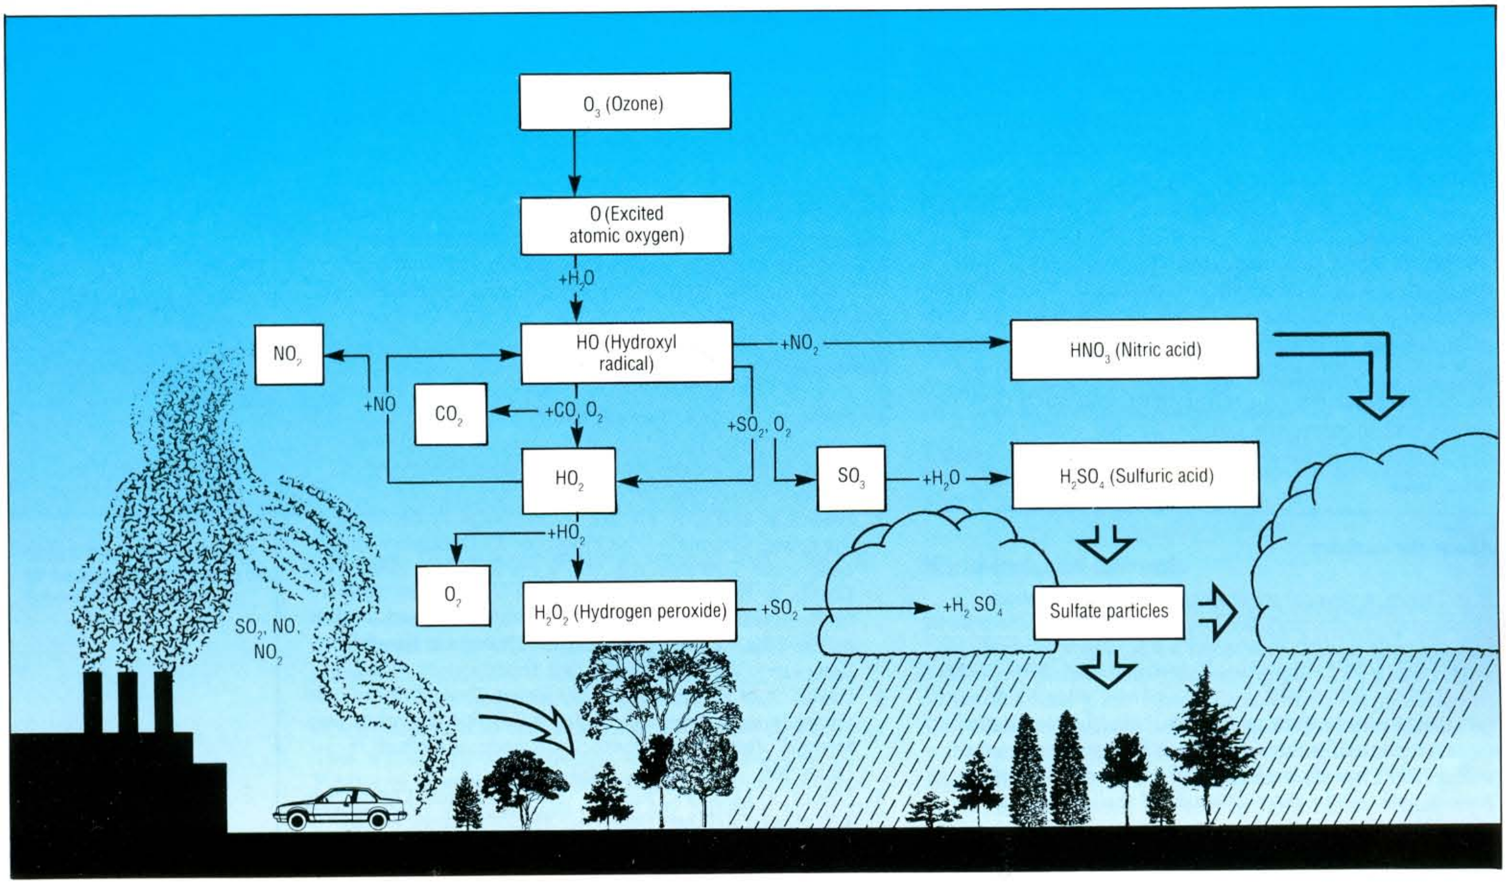

of the source, while the more damaging aqueous precipitation of reacted products can occur up to 2400 km away.[2] In the United States, the areas most affected by acid deposition are those downwind of high-sulfur coal-burning power plants, such as those along the Ohio River Valley. Studies have shown acid deposition to be the cause of many environmental problems, including declining fish populations in streams and lakes, damage to forests, and corrosion of buildings and other structures.

According to Robert Bruck, a professor of plant pathology at North Carolina State University who has studied acid rain effects in the Smoky Mountains, the average pH of rainfall today in the eastern United States is 4.1, compared to a pre-industrial level of 5.0. (The pH scale is a tool for measuring the acidity of aqueous solutions in terms of hydrogen ion concentration. A pH of 7 indicates a neutral solution. A pH of 6 is ten times as acidic as a pH of 7, and a pH of 5 is one hundred times as acidic.) The typical pH of water at the base of summer clouds in the eastern United States is 3.6, 25 times as acidic as pre-industrial rain, and can drop as low as 2.5. This causes severe damage to plant life at high elevations when mountain forests extend up into these low-pH clouds.

Three-dimensional output from the RADM program. The "clouds" collectively represent an iso-surface of sulfur dioxide over the eastern United States. The iso-surface is created by connecting all points that have the same concentration of the pollutant.

Scientific input for public policy

To provide useful input for rational decisions about control of acid deposition, policy assessments must consider present and future emissions. Simple empirical relationships derived from current observations are of limited use, because human contributions to the atmosphere might change substantially in the future. Also, empirical relationships cannot be used to distinguish the impacts of individual sources from the total effect of emissions from many sources.

The most useful tool for assessing improvements in air quality and reductions in acid deposition that will result from proposed emission reductions is an integrated theoretical modeling system that includes relevant physical and chemical coupling and feedback. Such a system can be used to identify and analyze processes involved in the formation, transport, and deposition of acid in the atmosphere and to compare the merits of alternative control strategies.

The Regional Acid Deposition Model

In 1983, the EPA, one of a number of federal agencies cooperating in the National Acid Precipitation Assessment Program, established the Acid Deposition Modeling Project (ADMP) at NCAR.[2] The project's primary mission was to design, develop, and implement a comprehensive regional modeling system to provide accurate quantitative assessments of relationships between acid-deposition sources and receptors. The project was based at NCAR until 1987, when it moved to the State University of New York (SUNY) at Albany, as the focus of its work shifted to model testing and application. In August 1990, ADMP became the Atmospheric Modeling Section, Atmospheric Sciences Research Center, SUNY. Julius Chang has directed the group since its inception.

The Regional Acid Deposition Model (RADM), developed by Chang and his colleagues, is an evolving system of computational submodels that describe physical and chemical processes involved in acid deposition. The RADM development effort had three objectives:

- ☐ RADM would be a state-of-the-art modeling system suitable for carrying out source-receptor assessment studies.
- ☐ The models would be sufficiently flexible to allow integration of current and developing representations of relevant physical and chemical processes.
- ☐ The models would describe the spatial and temporal distribution of pollutants from known sources.

In response to these requirements, the RADM system was designed as a hierarchical set of models. Scientists in many research groups conducted basic studies of individual physical and chemical processes such as atmospheric transport, dry deposition, wet scavenging, and gas-phase chemistry to provide a solid foundation for the system. ADMP's interdisciplinary team synthesized the results into a set of submodels and integrated them into a coupled, comprehensive system.

The major components of RADM are emissions, meteorology, gas-phase chemistry, transport and mixing of chemical species, dry deposition, and cloud processes. The model describes the influence of each component on a particular chemical species and its concentration in the geophysical domain in properly balanced computational sequence, mathematically representing the simultaneous interactions of species at a given location and time.

The meteorological component of the ADMP system was developed at Pennsylvania State University and NCAR. It uses standard global weather data for the initial data set, with additional North American data interpolated. This model produces self-consistent simulated three-dimensional wind and temperature

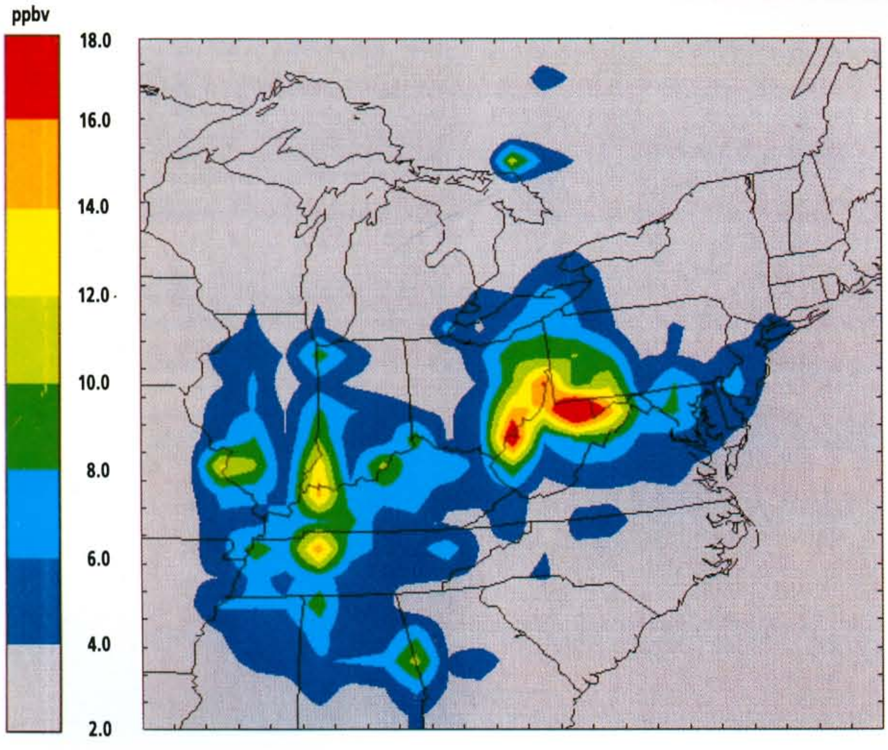

RADM output of the sulfur dioxide (SO_2) mixing ratio (the ratio of SO_2 to air) in parts per billion volume (ppbv) averaged over four days in September 1988 and averaged vertically over six RADM levels from the surface to about 1.5 km. Courtesy Julius Chang.

fields, precipitation and other meteorological variables for any North American region and any time period.

RADM currently divides the atmosphere over the eastern United States into six or fifteen vertical levels and its surface domain into a 35 by 38 horizontal grid with a resolution of approximately 80 km. RADM uses the meteorological model with analyzed emission inventories of sulfur oxides, nitrogen oxides, volatile organic compounds, and ammonia at each grid point to simulate episodes lasting several days. At each point, it predicts time-varying concentrations of 48 trace gases, together with wet and dry deposition fields, by solving equations that include the relevant physical and chemical processes. Pollutants from point sources, such as power plants and smelters, and from area sources, such as automobile exhaust emissions, are transported by the modeled winds. At the same time, RADM simulates their chemical transformations and redistribution through cloud interactions and other mixing processes. All these meteorological and chemical factors, as well as terrain and land-use characteristics, determine how much wet and dry acid deposition a particular area of the Earth's surface within the modeling domain will receive.

Because of the emphasis on modularity in computer coding to preserve flexibility in incorporating new findings and advanced submodels, the first version of RADM was not designed for speed. Initially, a six-level model with a 30 by 30 horizontal grid and 24 chemical species required more than two hours of CRAY X-MP processor time per three days of simulation. As the code matured and vectorization increased, execution speeds improved by a factor of more than eight. But with the second, more complex version of RADM (currently in use), the speed decreased again, requiring over four hours of CPU time on a single CRAY Y-MP processor per three simulated days. In single-processor mode, RADM currently executes at speeds of up to 150 MFLOPS, but work is underway to parallelize the code. The researchers hope for significant improvements in execution time with the addition of more vector processors. Future versions of the code are anticipated to be three times as large, with a commensurate increase in CPU requirements. Acid deposition modeling clearly is a problem that demands the most advanced supercomputing resources.

Evaluation and application

To establish the scientific credibility of RADM more conclusively, comparison of simulation and field data will continue until 1992. It is already the most thoroughly tested model ever developed for the agency, according to a senior EPA manager. Since 1990, RADM has been used in national acid precipitation policy analyses, and it will provide valuable information to policy makers for oxidant and visibility strategy assessments and for studies of the role of atmospheric chemistry in regional climate change.

In addition to providing scientific insight into atmospheric acid formation and deposition, the predictive capability of RADM has far-reaching economic implications. It has been estimated by the White House Council of Economic Advisors that meeting the acid rain precursor emission reduction requirements of the recently passed national clean air legislation in the United States could cost U.S. industries \$4 billion to \$5 billion annually by the year 2005. As such regulations are refined and revised in years to come, RADM will provide accurate estimates of the changes in acidic deposition that would result from various levels of emission reductions. In addition, its potential to define source-receptor relationships will help to determine where controls will be most effective. Finally, because the interactions between acid rain's chemical precursors also are suspected of playing a role in ozone depletion and climate change, RADM's ability to predict interactions should help to prevent policy makers from inadvertently worsening one pollution problem when taking action to alleviate another.[3]

About the authors

Henry Lansford received degrees from Tulane University and the University of Southern Mississippi. He is a scientific writer and editor with the Atmospheric Modeling Section, Atmospheric Sciences Research Center, State University of New York at Albany. He is based at the National Center for Atmospheric Research in Boulder, Colorado, and recently received a special award from the American Meteorological Society "for distinguished writing on meteorology, leading to increased public awareness and understanding of the subject."

Hugh Patrick received degrees in microbiology and electrical engineering from the University of Florida, and an MBA degree from Clemson University. He is employed by Cray Research at Research Triangle Park in North Carolina.

References

1. Mohnen, Volker A., "The Challenge of Acid Rain," *Scientific American,* Vol. 259, No. 2, pp. 30-38, 1988.
2. Frankel, Richard A., et al., *Air Pollution,* Report to the House of Representatives Committee on Energy and Commerce, GAO/RCED-90-14, 1989.
3. Schindler, D. W., "Effects of Acid Rain on Freshwater Ecosystems," *Science,* Vol. 239, January 8, pp. 149-157, 1988.

CORPORATE REGISTER

Cray Research sells first CRAY Y-MP2E systems

Chrysler Corporation installed a CRAY Y-MP2E computer system in the fourth quarter at the company's engineering center in Highland Park, Michigan. The new system replaces a CRAY X-MP/14se system and will be used for general automotive applications, including structural analysis and design, computational fluid dynamics, and crash simulation. The CRAY Y-MP2E system runs Cray Research's UNICOS operating system, which is based on the industry standard UNIX System V operating system from UNIX System Laboratories, Inc. "This new CRAY Y-MP2E system is needed for Chrysler to remain competitive in the 1990s," said Bill McVinnie, manager of Chrysler's computer-aided engineering department. "It will assist Chrysler engineers in the design of new vehicles to meet the challenging requirements for safety, emissions, fuel economy, and customer satisfaction. Without this new Cray Research system, the job would be significantly more difficult."

Mazda Motor Corporation in Hiroshima, Japan, has ordered a CRAY Y-MP2E computer system, with 32 Mwords of memory, and an SSD solid-state storage device. The computer system will run Cray Research's UNICOS operating system. This was the second order of Cray Research's new CRAY Y-MP2E system and the first CRAY Y-MP2E system order from outside the United States. The system will be used for automotive computer-aided engineering applications such as structural analysis and crash simulation. This system is Mazda's second Cray Research supercomputer; the company acquired a CRAY X-MP system in April 1989.

Honda Research & Development Company, Ltd., of Tochigi, Japan has installed a multiprocessor CRAY Y-MP supercomputer at its research and development facility in Tochigi. The CRAY Y-MP system replaces a CRAY X-MP system that was installed in 1987. The new system runs Cray Research's UNICOS operating system. The CRAY Y-MP system will be used for computer-aided engineering applications such as structural analysis, crash simulation, computational fluid dynamics, and other advanced design applications.

Tohoku University in Miyagi, Japan, installed a four-processor CRAY Y-MP system in the fourth quarter of 1990 at the university's Institute of Fluid Science. Tohoku University is a new Cray Research customer. The system runs Cray Research's UNICOS operating system. "The system installed at Tohoku University is a major milestone in Cray Research's further entry into the Japanese public sector market," said John Rollwagen, Cray Research chairman and chief executive officer. "We are extremely pleased that Tohoku University recognizes that Cray Research can provide it with the finest technology available to further develop computational fluid science and thus enhance its position in the academic community."

The NASA Langley Research Center in Hampton, Virginia installed its second Cray Research supercomputer, a CRAY Y-MP system, in the fourth quarter of 1990. The order included Cray Research's recently announced file server technology, the UNICOS Storage System, which will be used to manage more than one trillion bytes of storage. NASA will run the file server capability on its CRAY Y-MP system. The system will be used primarily for general aerospace applications, including structural analysis and design and computational fluid support for research in atmospheric modeling and computational chemistry. "This CRAY Y-MP system will be the second supercomputer delivered by Cray Research to NASA Langley under a contract established in 1988 through a competitive procurement process," said Mickey G. Rowe, assistant division chief of NASA Langley's Analysis and Computation Division. "The first supercomputer, a CRAY-2 system, has been an immensely successful experience, and we look forward to working with Cray Research to make the CRAY Y-MP system equally successful in support of NASA's research programs."

Shell Oil Company has installed a CRAY Y-MP system at its information center in Houston, Texas. "Shell has been a valued Cray Research customer for the past five years," said John Rollwagen, chairman and chief executive officer of Cray Research. "We look forward to continuing our relationship with Shell in the future." The new system will run the UNICOS operating system.

PRAKLA-SEISMOS AG, a private German geophysical and seismological research center, has ordered a CRAY Y-MP system for installation at the company's facilities in Hannover, Germany. This is PRAKLA's second Cray Research system. The center purchased a CRAY X-MP system in August 1989. The new system will run GEOSYS, PRAKLA's seismological interpretation and modeling software. This software analyzes large amounts of field data recorded acoustically on magnetic tape. PRAKLA and other geophysical companies have a need for magnetic tape software processing, which allows the supercomputer to read acoustical data at high speeds. PRAKLA pioneered the use of seismological magnetic tape processing with Cray Research's UNICOS operating system.

The National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign installed a four-processor CRAY Y-MP computer system in the fourth quarter of 1990. NCSA is one of four National Science Foundation supercomputing centers in the United States founded to help establish a national supercomputing arena to benefit U.S. academia and industry. NCSA's new Cray Research system will be used to advance knowledge in a variety of disciplines, including meteorology, chemistry, physics, and astrophysics. The new system replaces a CRAY X-MP system; NCSA also has a CRAY-2 system. "The new CRAY Y-MP supercomputer promises much in terms of speed and memory, both of which are important to the research community's ability to run and solve more advanced, computationally intensive problems," said David Curtis, NCSA's associate director of information and publications.

Cray Research offers upgrade credit for CRAY XMS minisupercomputers

Cray Research announced in November that purchasers of CRAY XMS minisupercomputer systems will receive full upgrade credit on orders for Cray Research's follow-on minisupercomputer system, which is scheduled to be available in the second half of 1991.

"The CRAY XMS system is designed to fulfill customers' immediate need for a cost-effective, Cray-compatible minisupercomputer solution," said John A. Rollwagen, Cray Research's chairman and chief executive officer. "This financial upgrade credit provides new and existing customers with an easy transition to our follow-on minisupercomputer system, which will be compatible with our current CRAY Y-MP supercomputer product series."

The financial upgrade credit "allows customers to purchase a Cray Research product now, with full confidence that their purchase value can be applied toward additional computer power they may need in the future," said Mike Lindseth, vice president and general manager of the company's Entry Level Systems Division, a new business unit formed to manage the design, manufacture, and marketing of Cray Research's minisupercomputers.

Lindseth said the trade-in arrangement will apply to orders of CRAY XMS systems taken by the end of April 1991 for delivery in 1991. Orders for the follow-on minisupercomputer must be taken, and delivery made, by December 31, 1992.

The CRAY XMS system, based partly on technology obtained through Cray Research's acquisition of Supertek Computers, Inc., is a 64-bit minisupercomputer compatible with the company's earlier CRAY X-MP supercomputer series. The CRAY XMS system runs Cray Research's UNICOS operating system, CF77 Fortran and Standard C compilers, and networking software. Existing UNICOS applications run on the CRAY XMS without modification, and applications developed on the CRAY XMS can be run on CRAY-2, CRAY X-MP, and CRAY Y-MP computer systems.

Minisupercomputers are a vital element in the company's network supercomputing strategy, Rollwagen said. "When networked with a high-end Cray Research supercomputer and other products, the Cray Research minisupercomputer systems make the entire computing environment more time- and cost-effective, by allowing users to select the level of Cray Research computing performance required for a given application."

Cray Research introduces new disk drives for the CRAY Y-MP2E supercomputer

In November Cray Research introduced the DD-60 and DD-61 disk drives. These disk drives are designed to complement the performance of the new CRAY Y-MP2E supercomputer. The new drives use highly reliable eight-inch disk technology to give users access to large amounts of data at high or low speed, with reduced requirements for floor space, power, and cooling. Software support for the new drives is available under release 6.1 and later versions of the UNICOS operating system. The new drives and version 6.1 of the UNICOS system will be available in the second quarter of 1991.

The DD-60 disk drive is a 24 Mbyte-per-second drive with a formatted storage capacity of 1.96 Gbytes and sustained transfer rates of 16 to 20 Mbytes per second, which is twice as fast as previous products. The lower priced DD-61 disk drive offers a larger formatted storage capacity of 2.23 Gbytes and a sustained data transfer rate of 2.3 to 2.6 Mbytes per second. A fully configured CRAY Y-MP2E system can support more than 400 Gbytes of DD-60 or DD-61 storage. Both the DD-60 and DD-61 can be daisy chained to allow up to eight disk drives to be connected to a single channel.

"We developed these drives to deliver high performance and high reliability," said Les Davis, Cray Research executive vice president for Chippewa Falls operations. "Their design includes data parity recording and other new engineering advancements such as closed-loop micropositioning, which ensures reliability in the most demanding Cray Research system applications."

The new disk drives also complement Cray Research's UNICOS Storage System, a high-performance network storage management system that provides high-speed access to large numbers of files on a variety of storage media. With the introduction of the DD-60 and DD-61 disk drives, this file server feature of the UNICOS operating system can be supported by large high-speed data storage, which can be combined with automated bulk tape storage to provide total system data capacity of over 19 trillion bytes.

The DD-60 and DD-61 disk drives consume less power than previous models, and are equipped with dual 50/60 Hz power supplies to accommodate most standard power sources. Floor space requirements are reduced significantly with a new disk cabinet design that contains up to eight disk drives and with the elimination of a separate disk controller cabinet.

"These products are important to our network supercomputing strategy," said Davis. "With these new disks, the UNICOS Storage System, and the CRAY Y-MP2E supercomputer, we have all the components for high-performance network data management at very appealing price/performance ratios. We can store the large data files that our customers need and move data at rates that efficiently use the processing power of the CRAY Y-MP2E supercomputer."

APPLICATIONS UPDATE

DADS analysis package available on Cray Research systems

The DADS (Dynamic Analysis and Design System) software package from CADSI is a mechanical computer-aided engineering (MCAE) package that performs nonlinear large-displacement transient analysis and simulation. The package runs on all Cray Research systems under the UNICOS operating system. It allows the user to model real-world behaviors of complex systems. The DADS package solves for joint position, displacement, velocity, acceleration, total and potential energy, and internal reaction forces of models. It performs several functions:

- Assembly analysis determines whether all parts of a model can be connected by the defined joints.
- Static analysis calculates the static equilibrium position and potential energy of the modeled system and components.
- Kinematic analysis calculates the relative motions of bodies in a mechanism without regard for the mass effects of the components or any forces in the system.
- Inverse dynamic analysis calculates the forces necessary to move a model through a predetermined path or given motion.
- Dynamic analysis calculates the relative motion of bodies in a mechanism, including mass effects of the components and any forces in the system.

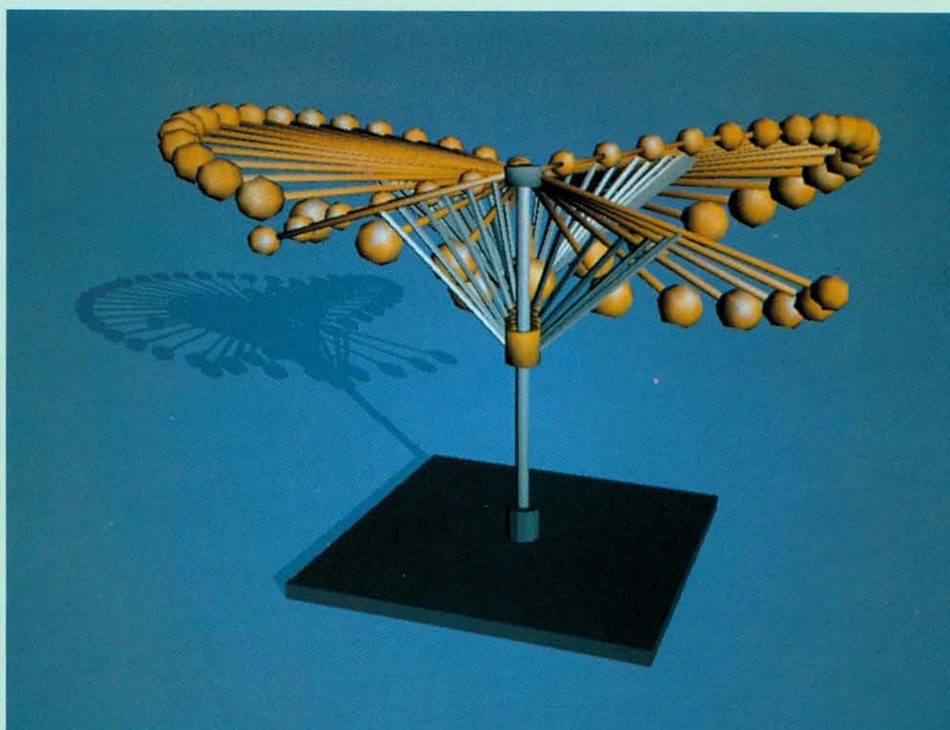

Flyball governor simulation created with the DADS software package. The simulation shows a time-lapse visual effect possible with the software.

The DADS package includes a library of control and hydraulic components, which enables users to model dynamics of feedback, control, and hydraulic subsystems; their coupling with a mechanical system; and forces and torques that are fed back to and act on components in the mechanical system. DADS also combines flexible-body analysis capabilities with traditional rigid body dynamics, using natural modes of vibration and static correction modes that are associated with kinematic constraint reaction forces. It includes an interface processor for use with finite element structural analysis codes to calculate stiffness, mass, and deformation properties.

The DADS package includes a preprocessor and a postprocessor that presents the results of analyses in tabular or graphical form. DADS can be integrated with other MCAE applications, such as FEA/FEM and CAD/CAE, as part of a total design system. Interfaces are available for the AutoCAD, PDA/PATRAN, MSC/NASTRAN, COSMIC/NASTRAN, SDRC/I-DEAS, and ANSYS programs. For more information about using DADS on Cray Research computer systems, contact the Marketing Department, CADSI, P.O. Box 203, Oakdale, IA, 52319; telephone: (319) 337-8968; or Doug Petesch, Cray Research, Inc., 655-E Lone Oak Drive, Eagan, MN, 55121; telephone: (612) 683-3654.

ASPEN PLUS with ModelManager available on Cray Research systems

The ASPEN PLUS program with ModelManager, from Aspen Technology, Inc., is a process modeling system available on Cray Research computer systems under the UNICOS operating system. ASPEN PLUS is used to simulate, design, and revamp process industry plants that have a continuous flow of materials and energy through a network of process units. ASPEN PLUS allows engineers to simulate manufacturing processes, thereby reducing the need for pilot plant experimentation, allowing products to be brought to market faster.

ASPEN PLUS offers a complete library of generalized unit operation models to simulate any kind of process, including processes with solids, electrolytes, or complex substances. ASPEN PLUS unit operation models do not have any dimensional limitations and can have any number of components, inlet streams, stages, and reactions. These unit operation models include

- ☐ Mixers, splitters, and separators
- ☐ Flash separators
- ☐ Distillation, absorption, and extraction
- ☐ Heaters and heat exchangers
- ☐ Continuous and batch reactors
- ☐ Pumps and compressors
- ☐ Solids handling operations
- ☐ Stream manipulators

ASPEN PLUS also has an extensive library of state-of-the-art physical property models, which are key to obtaining useful simulation results. The physical property system calculates all of the thermophysical properties needed to run unit operation models, size equipment, and produce reports and tables. ASPEN PLUS includes data for more than 1500 components, including the DIPPR Data Compilation.

To solve large and difficult flowsheeting problems, ASPEN PLUS has powerful convergence and control features, including

- ☐ Automatic sequencing, which automatically detects recycle loops, selects tear streams, generates convergence blocks, and sequences calculations independently of block input order
- ☐ Fortran blocks, which allow users to insert their own in-line Fortran statements into flowsheet computations
- ☐ Design specification control, which allows users to specify a desired value for any block result, stream flow, component flow or purity and designate any input variable to be manipulated to meet the specification
- ☐ Transfer blocks, which enable users to copy streams, stream variables, and block variables easily from one part of the flowsheet to another
- ☐ Report scaling, which allows users to scale reports so that a designated simulation result has a desired value
- ☐ Heat and work streams which allow ASPEN PLUS to represent heat flow from one distillation column to another

ASPEN PLUS allows users to change virtually any problem specification interactively and to build a simulation model one block at a time. The latest version of ASPEN PLUS includes ModelManager, Aspen Technology's interactive, expert system, graphical interface. ModelManager works on workstations and personal computers networked with Cray Research systems running ASPEN PLUS. ModelManager provides a complete environment for performing and managing process modeling. With ModelManager, process engineers receive expert-system assistance in translating real-world engineering problems into error-free computer terms for flowsheet simulation and process optimization using ASPEN PLUS.

RATEFRAC, a rate-based nonequilibrium separation process model used with ASPEN PLUS, also is available on Cray Research systems. Engineers can use RATEFRAC to design and develop operating strategies for multicomponent, multistage separation processes such as distillation, absorption, and desorption. The model can be used with both tray and packed columns. RATEFRAC enables engineers to predict the behavior of separation processes accurately under nonequilibrium conditions. This knowledge is essential for the development of more reliable process designs.

ASPEN PLUS with ModelManager and RATEFRAC are used in the chemical, petroleum, mineral processing, food processing, pulp and paper, and other processing industries. For more information about using ASPEN PLUS with ModelManager and RATEFRAC on Cray Research computer systems, contact Mary K. Deery, Aspen Technology, Inc., 251 Vassar Street, Cambridge, MA, 02139; telephone: (617) 497-9010; or Steve Zitney, Cray Research, Inc., 655-E Lone Oak Drive, Eagan, MN, 55121; telephone: (612) 683-3690.

ANSYS engineering program includes new features

The ANSYS program, from Swanson Analysis Systems, Inc., is a general-purpose, finite element analysis program that includes preprocessing, problem solving, and postprocessing. ANSYS is used in a wide range of disciplines to solve mechanical, thermal, and electronic problems. The program runs on all Cray Research computer systems under the UNICOS and COS operating systems. Revision 4.4A"+" of ANSYS offers several new features, including

- ☐ Improved modeling efficiency with an enhanced picking and working plane functionality that allows more graphical interaction and less command manipulations
- ☐ Enhanced user convenience with a simplified ANSYS Parametric Design Language (APDL)
- ☐ Increased mesh reliability with a new solution-level command that controls automatic mesh refinement and error norm convergence
- ☐ Improved manageability and flexibility with network licensing, a user optimizer, an external file converter, and enhanced wavefront reordering capabilities
- ☐ Product quality enhancement capabilities that allow ANSYS to model unsymmetric matrices, two-dimensional viscoelastics, nonlinear composite materials, tetrahedron and hexahedron linear elastics, shear panels, contact surfaces, and inertia relief

For more information about using ANSYS on Cray Research computer systems, contact John Twerdok, Swanson Analysis Systems, Inc., Johnson Road, P.O. Box 65, Houston, PA, 15342-0065; telephone: (412) 746-3304, extension 412; or Sifon Eng, Cray Research, Inc., 655-E Lone Oak Drive, Eagan, MN, 55121; telephone: (612) 683-3663.

USER NEWS

Gigaflop awards honor top performers

Cray Research computer systems have made possible new levels of computing performance across a range of industrial and scientific applications. To recognize researchers who are achieving the highest levels of performance while solving practical problems on Cray Research systems, Cray Research has instituted the Gigaflop Performance Award program. At the 1990 Gigaflop Performance Award ceremony, held in November at the IEEE Supercomputing '90 conference in New York City, Cray Research honored 30 researchers and research teams who have achieved a sustained processing rate of at least 1.5 billion floating-point operations per second (1.5 GFLOPS) on Cray Research computer systems. The qualifying speed for the 1990 awards was 50 percent higher than the 1.0 GFLOPS requirement of the 1989 award program. Computer performance at GFLOPS speed can enhance significantly the cost-effectiveness, timeliness, and accuracy of numerical results in engineering and scientific problem solving. The following descriptions highlight a few of the 1990 Gigaflop Performance Award winners. For a brochure that describes the work of each of the winners, contact Diane Ciardelli, Marketing Communications, Cray Research, Inc., 1440 Northland Drive, Mendota Heights, MN, 55120; telephone: (612) 683-7215.

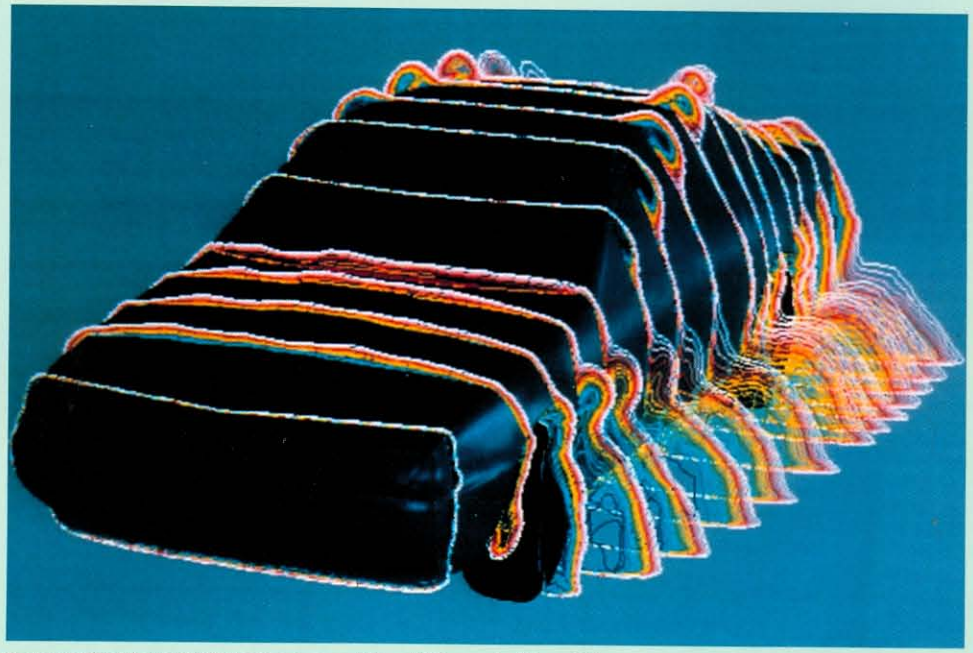

Contours of total pressure in the time-averaged flow field around a Nissan Cefiro.

Automotive aerodynamics at 1.5 GFLOPS

Researchers Ryutaro Himeno and Katsuro Fujitani of Nissan Motor Company Ltd. developed the DRAG4D program to predict aerodynamic characteristics, such as drag and lift coefficients, of production automobiles. The code models unsteady, incompressible, viscous three-dimensional air flow, which has a measurable effect on aerodynamic performance.

The work for which Himeno and Fujitani received a Gigaflop Performance Award represents a breakthrough in the application of computational fluid dynamics (CFD) techniques in automotive aerodynamics. On an eight-processor CRAY Y-MP system, the researchers achieved performance of 1.544 GFLOPS, which translates into a 30 percent cost savings compared with wind tunnel tests. DRAG4D was written in Fortran and initially achieved 1.4 GFLOPS on the CRAY Y-MP system. The major enhancement that enabled performance to exceed 1.5 GFLOPS was a reorganization of operations to obtain more chaining and functional unit overlaps.

DRAG4D uses a third-ordered upwind difference scheme. This grid system was generated by using a multiblock transformation and a trans-finite method. The problem included 1.3 million grid points and was solved over 5000 time steps. The code solved both a Poisson equation with 1.3 million unknown pressure values and linear equations with 3.9 million velocity components and used time integration in each time step. The successive over-relaxation method was used to solve both types of equation. Himeno plans to continue working with DRAG4D and improve its speed by 20 percent within a year.

Elastic wave model aids petroleum exploration

Computational techniques that can handle large problems with minimal approximating assumptions are increasingly important to geologists and geophysicists in the petroleum industry. Accurate computer modeling of geological structures can significantly speed up and improve interpretations of field-gathered data. EWAVE, developed by mathematician Kirk E. Jordan and colleagues while on staff at Exxon Research and Engineering Co., uses an algorithm that models the earth as an elastic medium by solving a full approximation to the elastic wave equations. Using an eight-processor CRAY Y-MP system, Jordan and Cray Research analyst John P. Flory achieved a processing speed of 1.773 GFLOPS with the EWAVE code.

The code uses a fourth-order-accurate in space and second-order-accurate in time, finite difference scheme for the computation of waves in an elastic medium. It provides more accurate solutions than can be obtained with more conventional approximation schemes, which treat the Earth as an acoustic medium.

"For modeling wave propagation in the Earth, this approach is more realistic than methods based on solving acoustic waves, since the Earth for the most part is an elastic medium," Jordan said.

The elastic wave equations, in contrast to acoustic wave equations, are complicated by the addition of stress tensors that must be solved along with the velocity terms. In the EWAVE code, a system of five partial differential equations is solved numerically. The model problem solved is further complicated by incorporating a region of fluid, simulating a lake, in the problem domain,

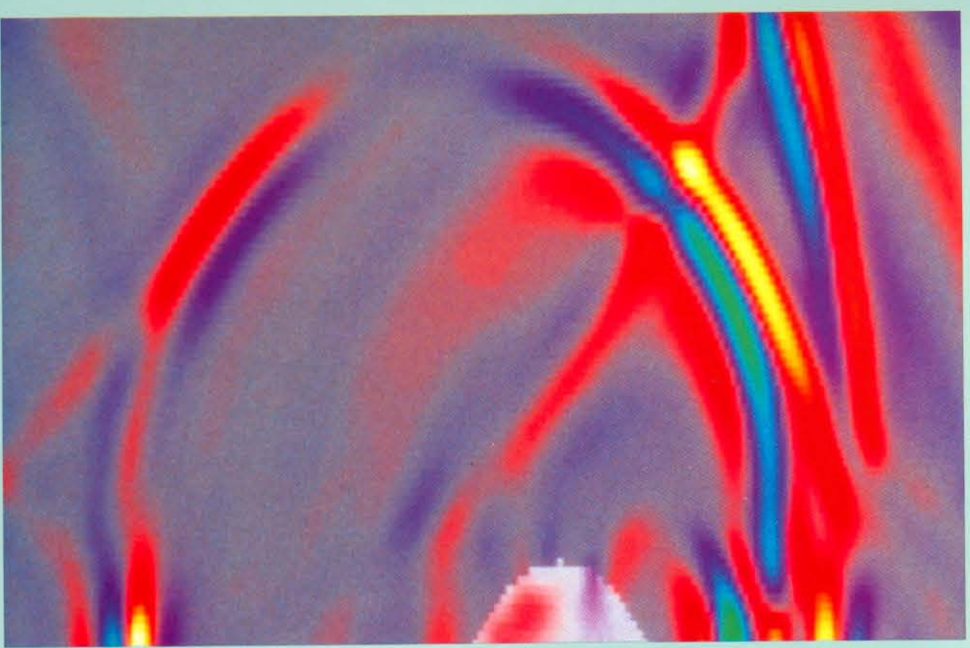

Snapshot of vertical velocity obtained for the problem of a Rayleigh wave scattered from fluid.

which gives an elastic-fluid interface from which the Rayleigh waves are scattered.

"EWAVE holds interesting possibilities for the extension of these same techniques to three dimensions where fast, accurate solutions are essential," Jordan said.

Team uses GFLOPS performance to reveal three-dimensional turbine flows

Researchers Viet-Nam Nguyen, M. Vish Bhat, and Martin F. Peeters of Pratt & Whitney Canada, and Wagdi G. Habashi of Pratt & Whitney Canada and Concordia University in Montreal have collaborated for the past 13 years on a finite element computational fluid dynamics (CFD) methodology for the analysis and design of gas turbines. Recently, the team developed a Navier-Stokes code to analyze three-dimensional viscous flows in complex internal flow passages. The problem for which the team received a Gigaflop Performance Award involved flow through a gas turbine diffuser located between the low and high pressure centrifugal compressors of the Pratt & Whitney Canada PW-124 series engines. In a very short distance, this pipe diffuses the flow and guides it into the second centrifugal compressor.

"Understanding the complex flow in such diffusers is crucial to improving their aerodynamic design because they account for a significant part of the losses in the overall compression," Habashi said.

The finite element method developed by the team for the Navier-Stokes equations uses a regularization procedure based on a pressure dissipation term in the continuity equation, allowing equal order elements for velocity and pressure. The equations

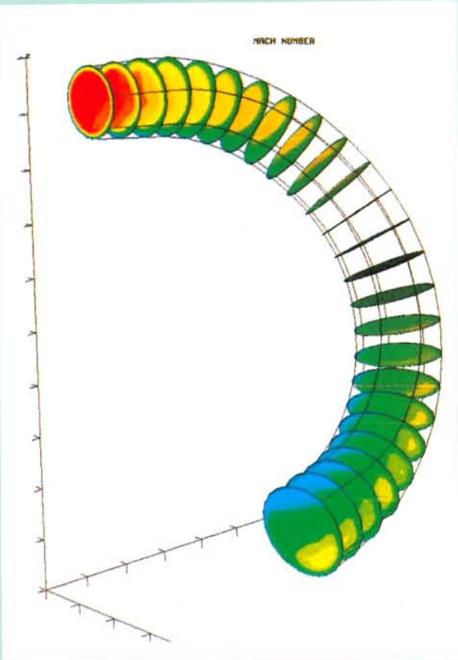

Three-dimensional Navier-Stokes analysis of the gas turbine diffuser by the finite element method.

are linearized by Newton's method and the velocities and pressure are solved simultaneously in a fully coupled manner. While efficient for two-dimensional flows, such a direct approach was previously considered prohibitively expensive for three-dimensional flows, in terms of execution time and memory.

The team broke through that barrier by using the largest-memory eight-processor CRAY Y-MP system. The nonsymmetric solver algorithm was developed by Nguyen in Fortran, by taking advantage of dynamic assignment of equations to CPUs, dynamic loop unrolling, and a dynamic elimination procedure. The team achieved a processing speed of 2.3 GFLOPS on the Cray Research system.

"With this advance," Habashi said, "future testing or parametric design of diffusers can be accomplished with powerful computers like the CRAY Y-MP system instead of prototypes. It will save hundreds of thousands of dollars per design."

Supercomputer simulates 300 years of fire history at Yellowstone

The flames that engulfed enormous portions of Yellowstone National Park and dominated national news for weeks in 1988 did more than just ignite precious forest land: they also sparked a controversy about the role of fire in forests. Since 1872, the United States had operated under a fire suppression policy that required forest fires to be extinguished. But in 1972 the fire suppression policy was lifted for a "let burn" policy, in which natural fires were allowed to burn. As the Yellowstone fires blazed during the summer of 1988, this "let burn" policy was being questioned by the public.

With a growing amount of data to support their understanding of how fires affect a forest's evolution, experts have determined that natural fire is beneficial and even necessary to a forest's vitality — turning fallen plant material into essential nutrients, releasing seeds from pine cones, and allowing sunlight to penetrate to seedlings. Fires also assure plant and wildlife diversity in forests that otherwise would become stagnant groups of similar trees, home for a relatively narrow range of wildlife. However, this notion has not been communicated successfully to the public.

As data was collected about Yellowstone's fire history, the University of Illinois' David Kovacic became interested in presenting the information visually. With a CRAY X-MP supercomputer he was able to simulate the results of various fire management scenarios. "There is a tremendous amount of information about the 1988 fire in Yellowstone, as well as the park's fire history during the last 300 years," Kovacic said. "But the only method used previously has been to look at single frames of information — what happened in a single year, and so on. We thought it would be interesting to take all of the information and look at it in a more visual, dynamic context."

An ecologist in the Department of Landscape Architecture at the University of Illinois in Urbana-Champaign, Kovacic first became interested in visual techniques several years ago. Impressed with

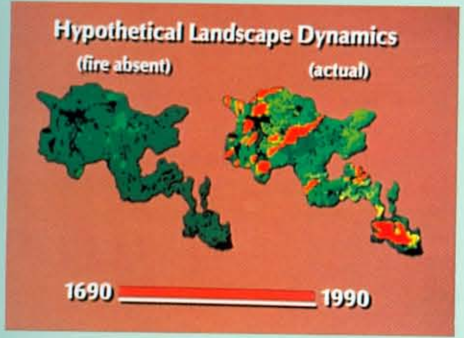

Comparison of vegetative growth in Yellowstone National Park (1690-1988) after fires (right) and result of absence of fire during same period (left) as simulated on a CRAY X-MP system at the National Center for Supercomputing Applications. Red indicates areas of 1988 fires.

Legend with color codes for vegetative growth comparisons.

an elk migration study that used video animation to communicate data, Kovacic tried the same technique in one of his own talks.

Kovacic used data compiled by Bill Romme, associate professor of ecology at Fort Lewis College in Colorado, and Don Despain, an ecologist at Yellowstone National Park. Their data described the percentage of trees of various ages for each period of their study, which went back to the year 1735. They charted the occurrence and severity of fires during the same period. They then plotted this information mathematically to correlate the fire data with data for tree age groupings.

Yet this methodology didn't go far enough, Kovacic explained. "You can look at the individual computer printouts for certain time intervals, but trying to get a feel for what happened over that long period of time is difficult. Looking at the information in a dynamic visual format gives you a much better feel for what is actually happening."

Working with Alan Craig, a visualization specialist from the National Center for Supercomputing Applications (NCSA) in Urbana-Champaign, Kovacic was able to create an animation of the data compiled by Romme and Despain.

"When the original data came from Kovacic, I created the preliminary images using NCSA Image on the Macintosh computer." Craig said. Craig and Kovacic then converted the original code to run on a CRAY X-MP/48 supercomputer.

The data matrices were so big that Romme and Despain could run their analysis only at large intervals of 20 or 25 years on the Macintosh computer. "When we used the supercomputer, we could look at the data on a yearly basis," says Kovacic. "Not only did we visualize the past 300 years of fire history at Yellowstone, we also showed what would have happened to the biological diversity if there had been no fires.

"Using the Cray Research supercomputer enabled us to make certain images, look at them, then try a different run, and make more changes," Craig said. The Cray Research system gave us the speed we needed so that we could do 'what ifs' as we went along. It also gave us a convenient working platform for generating the images."

The supercomputer allowed the researchers to create hypothetical scenarios to describe what would happen if a fire had been withheld from a certain area, or if the fire suppression program had been in place all along. "The project was interesting," Craig said, "because while we wanted to see images of the actual data, we also wanted to be able to see what the landscape would look like if we had suppressed all the fires during the last 300 years. By doing this, we were able to better understand how important natural fires are, and how our policies might have affected the park by interfering with the natural fires that occur.

Kovacic said that they were also able to drive the model backwards. "We were able to get an estimate of what was there as far back as 40 years prior to our earliest point of information."

The results of the output were written as NCSA HDF (Hierarchical Data Format) files, which were used with NCSA Image software on a Macintosh II system to create the images that went into the final animation. Color represents stands of different ages in the images.

Kovacic said that the sophisticated technology they were able to use in producing this video enabled the researchers to put together high-quality graphics and produce one of the first animations to look at this type of natural resource problem. He said, "This work is valuable to researchers, but also will help to bridge the gap between researchers and the public."

Model shows which way the wind blows

Among the considerations that architects, civil engineers, and city planners must address when designing individual buildings or large development projects is wind. Wind gusts and vortices can impart loads that can damage buildings, and winds can cause significant heat loss by infiltrating and modifying a building's internal air flow. Wind flow also can create air pockets behind buildings that can trap pollutants from a kitchen or chimney. In downtown areas, vortex shedding and the wind-tunnel effect caused by tall buildings can produce discomfort and even hazards for pedestrians.

To predict the loads induced by winds on buildings and to assess the effects of winds on indoor and outdoor energy performance and comfort, University of Minnesota professor Charles C. S. Song and research assistant Jianming He have developed an unsteady three-dimensional numerical model of wind flow. The model uses the weakly compressible flow approach to calculate unsteady flows at conditions characterized by small Mach numbers and large Reynolds numbers. The researchers are using the CRAY-2 and CRAY X-MP/416 computer systems at the Minnesota Supercomputer Center (MSC) to compute unsteady

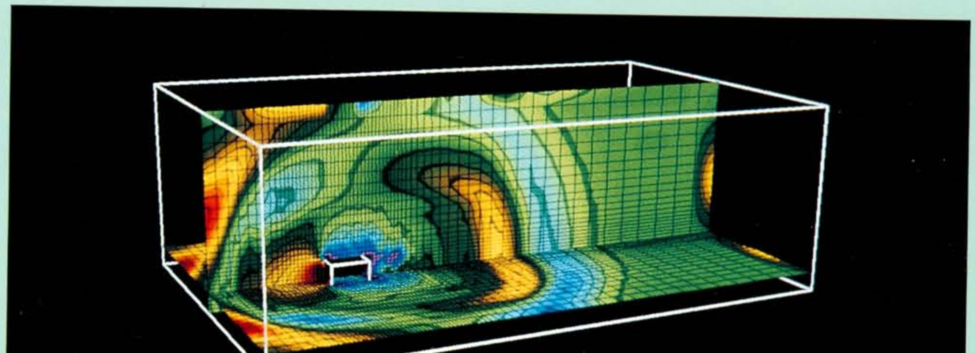

Turbulent wind flow around a building as simulated on the CRAY-2 supercomputer at the Minnesota Supercomputer Center. The image shows a three-dimensional view of time-averaged pressure distribution in a vertical and a horizontal plane. The work was commissioned by the state's Legislative Commission on Minnesota Resources.

CUG reports

Users of Cray Research computer systems established the Cray User Group (CUG) in 1977 to provide a forum for the exchange of ideas related to Cray Research systems and their applications. The group holds two general meetings each year. Its second meeting of 1990 was held October 1 - 5 in Austin, Texas, and was hosted by the University of Texas Center for High Performance Computing. Below, CUG president Mary Zosel of the Lawrence Livermore National Laboratory, offers her comments on the meeting and other CUG-related business.

The recent Cray User Group meeting held in Austin, Texas, was a large and busy meeting with about 400 people registered and about 90 sites represented for the elections. User requirements in several areas were surveyed, and discussions about accounting and NewQS design took place. The User Services area was formally recognized as a Mutual Interest Group (MIG), the environmental group got off to a good start, the Operations Special Interest Committee (SIC) was reported to be "on a roll," and the attendance at the Mass Storage sessions continued to prove the importance of this topic to the sites. The theme of the meeting, "Social Benefits of Supercomputing Research," produced some very interesting talks about computational medicine and dentistry.

Representatives from Cray Research not only reported on progress and plans for hardware and software, but also set up live demonstrations of a number of their tools, installed a CRAY XMS minisupercomputer in the hotel, and demonstrated T3 line access from the University of Texas.

As a result of elections held at the meeting, the new CUG Board of Directors is as follows: president: Mary Zosel of Lawrence Livermore National Laboratory; vice president: Karen Sheaffer of Sandia National Laboratory, Livermore; secretary: Gunter Georgi of Grumman; treasurer: Howard Weinberger of Lockheed; at-large directors: Claude Lecoeuvre of CEA-CEL, Ken Neves of Boeing, and Dave McWilliams of the National Center for Supercomputing Applications.

CUG is continuing to address a number of issues arising from the changing supercomputing environment. A small group has been formed to work with Gary Sparks of Cray Research to address the functionality of the user-accessible SPR database, CASPR. We expect this functionality to be increasingly important as more sites operate without resident analysts. CUG is asking Cray Research to work on an information package to aid sites in their strategic planning, in particular with regard to the role of large, centralized computing resources in evolving, distributed, workstation-based environments. The board also is having discussions with Cray Research to try to plan a symposium for computer center executives to discuss strategic computing issues.

CUG asked Cray Research to address IEEE arithmetic, and company representatives responded at the meeting with their plans for future architecture support. CUG also has asked Cray Research to take a stand on the issue of support for software tools that run on platforms other than Cray Research systems in network environments. Cray Research, in turn, has asked CUG to find a way to get better feedback about issues related to software service.

As Cray Research's customer base grows, we are becoming increasingly concerned with maintaining international communications. We will be addressing the question of ensuring international representation in CUG management and better ways to serve sites that are geographically remote from almost all of the CUG meetings. Another issue facing CUG is the best way to serve our membership as Cray Research diversifies its product line into a broader range of sizes, which could result in many more potential member sites for CUG.

CUG has not yet adopted formal sponsorship of a software-sharing mechanism, but it is clear that many sites would find this useful. We would like to thank Reagan Moore of the San Diego Supercomputer Center for helping to address this need by making facilities at the center available for software sharing.

The CUG meeting schedule for the next few meetings includes spring 1991 in London (April 22 - 26), fall 1991 in Santa Fe, New Mexico (September 23 - 27), and spring 1992 in Berlin (April 6 - 10). We hope that representatives from as many sites as possible will join us.

flow components such as three-dimensional velocity, vorticity, and the pressure fields that result from large-scale flow structures, such as vortex shedding and horseshoe vortices. Prior to this work, studies of wind flow around buildings typically were limited to steady-state or time-averaged conditions, because of the difficulty in measuring or calculating time-dependent quantities.

"Continuing support in the form of computer time on the CRAY-2 and CRAY X-MP/416 computer systems from the MSC makes this study possible," said Song. To produce sufficient resolutions and statistically meaningful results, a single building model with 94,864 (77 by 28 by 44) grid points and 6000 time steps requires 2.86 hours on the CRAY-2 system and requires over 15 million words of central memory. Most of the CPU time and memory are used for calculating the 27 flux components of the three-dimensional turbulent fluxes around the surfaces of the discretized finite volumes. The input flow profile at the upstream boundary is based on the real-world situation.

"The time-averaged calculation agrees very well with experimental data, as do the calculated instantaneous flow features," said Song. "We are able to simulate the pressure distribution acting on the building, vortex shedding from the building, and horseshoe vortex phenomena successfully on the CRAY-2 system."

The instantaneous numerical output data are converted on the CRAY-2 system to color table indices with values from 0 to 255, using the *fltras* program developed at MSC.

The frame images are created using the Starbase software package on a Hewlett-Packard 350 graphics workstation and are recorded continuously on video tape. The visualization was rendered by Elaine Larson of MSC. The video tape enables the researchers to visualize the detailed flow structure and its evolution through time. The image shows the time-averaged pressure distribution in a vertical and a horizontal plane. It reveals increases in pressure as flow approaches the building, attaining its maximum value on the front surface of the building. The flow separates from the building immediately after crossing the front edge, where pressure drops sharply. In addition, a low-pressure cavity is created behind the building, indicated by blue. On the horizontal plane, an apparently symmetrical low-pressure zone around the building is caused by a horseshoe vortex. Unlike these time-averaged quantities, the instantaneous flowfield is quite asymmetrical due to unsteady vortex shedding from the building. These results demonstrate the feasibility of using Cray Research systems to simulate complicated, unsteady, three-dimensional turbulent wind flow around a building.

GALLERY

"A Garden on Vega Nine" was rendered on a CRAY Y-MP computer system using the SCOPE program written by Melvin Prueitt. The image illustrates an application of texture mapping. Image courtesy Melvin L. Prueitt, Computer Graphics Group, Los Alamos National Laboratory.

CRAY CHANNELS welcomes Gallery submissions. Please send submissions to the address inside the front cover.