A CRAY RESEARCH, INC. PUBLICATION

CRAY CHANNELS

Winter 1988

FEATURE ARTICLES:

Multitasking ECMWF's weather model

Environmental research at NCAR

Computational forecasting in Canada

Winning winglets

Comparing tridiagonal solvers

DEPARTMENTS:

Corporate register

Applications in depth

User news

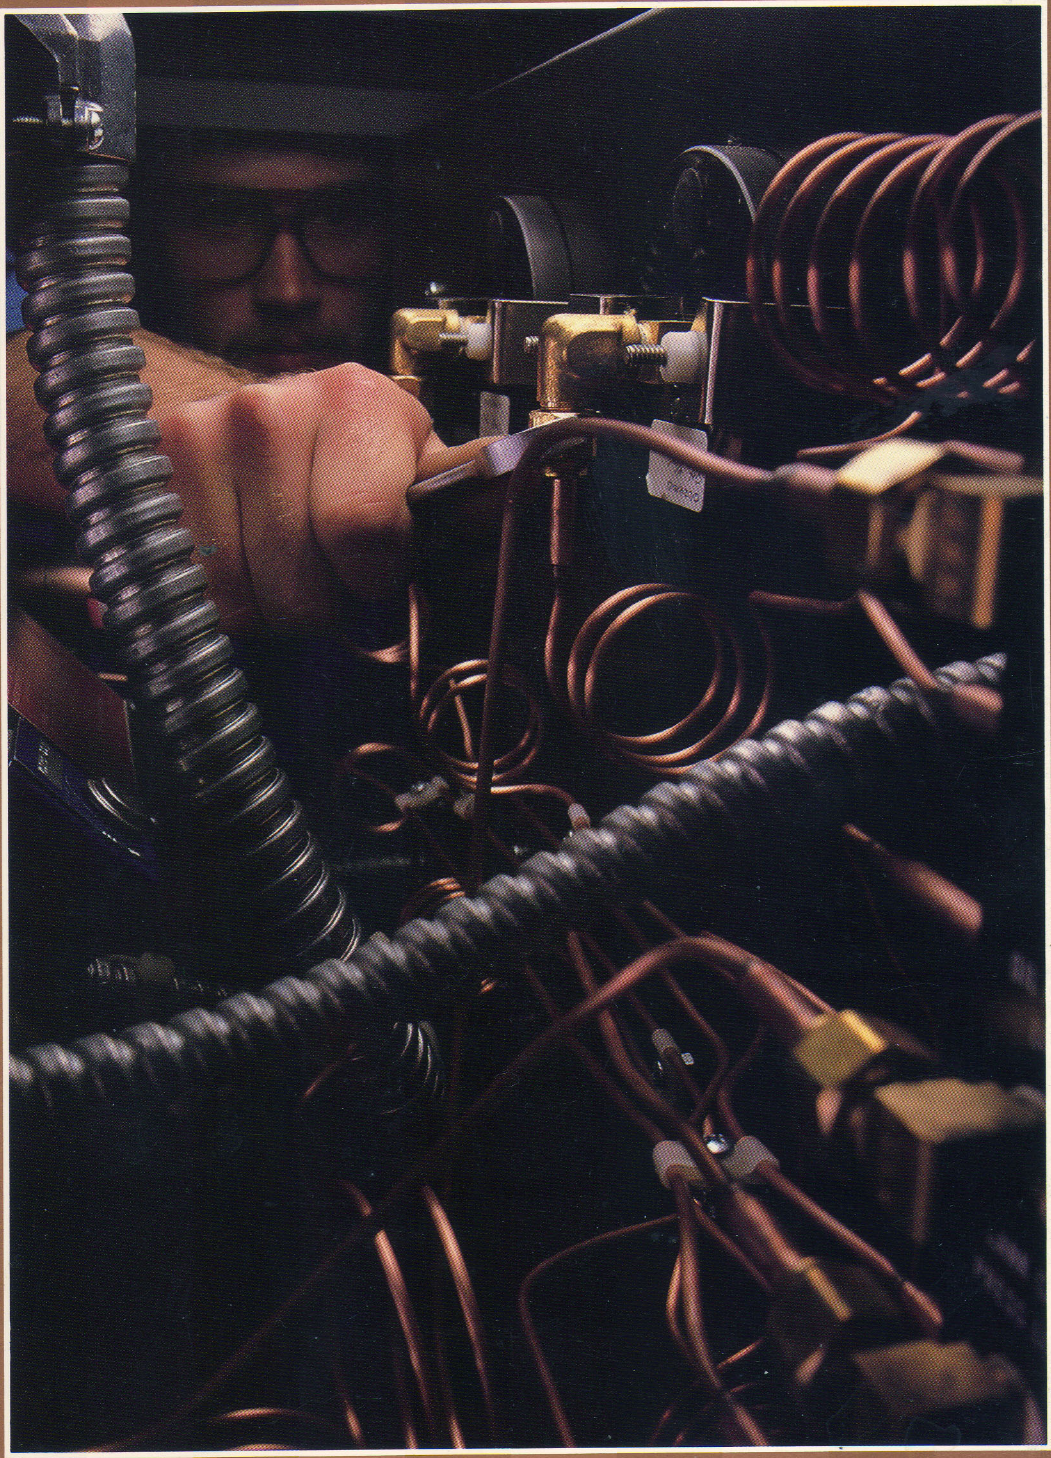

IN THIS ISSUE

Exploring the interacting forces of the winds, oceans, and sun would be nearly impossible without the power of today's supercomputers. By bringing the forces of nature into the laboratory, computational models enable meteorologists and climatologists to examine all ends of the Earth. Throughout the world, Cray systems provide the means to map ozone depletion, recreate ancient climates, probe the atmosphere of space, and even assess weather patterns of the future.

This issue of CRAY CHANNELS spotlights some unique uses of Cray systems for environmental science. Researchers at the European Centre for Medium-Range Weather Forecasts (ECMWF) describe multitasking methods for maximizing the power of their CRAY X-MP system. Thanks to efficient use of the Cray architecture, ECMWFs five- to six-day forecasts of today are more accurate than the three-day forecasts of 1977. We also travel to the National Center for Atmospheric Research in Boulder, Colorado, where scientists are using Cray supercomputers to model complex systems and find answers to questions about climate on Earth and in space.

The power of today's computers has become essential to the advancement of meteorology. In fact, computers have changed the way forecasting is done. In this issue, the chief of the Canadian Meteorological Center's computer center summarizes advancements in Canadian meteorological research that have been the result of the acquisition of a Cray system. Design consultants for the 1987 America's Cup winner, *Stars & Stripes*, reveal some of their computational design strategies for hydrodynamic optimization. We also take a look at performance comparisons of tridiagonal solvers. Our regular departments feature new Cray system orders, a look at Cray Research's environmental sciences support team, and computerized sculpture at the National Center for Supercomputer Applications.

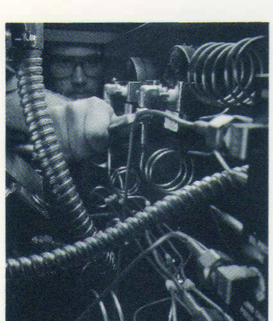

On the cover a Cray technician at Cray Research's mechanical assembly facility in Chippewa Falls, Wisconsin makes final adjustments to a 50-ton condensing unit that will cool a CRAY X-MP system.

CRAY CHANNELS

A Cray Research, Inc. publication

Winter 1988

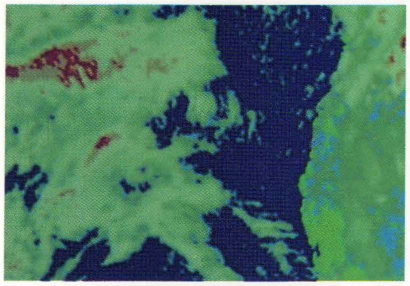

14

FEATURES

DEPARTMENTS

CRAY CHANNELS is a quarterly publication of Cray Research, Inc. intended for users of Cray computer systems and others interested in the company and its products. Please mail subscription requests, feature story ideas, and news items to CRAY CHANNELS at Cray Research, Inc., 608 Second Avenue South, Minneapolis, MN 55402.

Volume 9, Number 4

Editorial staff
Tina Bonetti
Ken Jopp
Elizabeth Knoll
Carol Smith
John Swenson

Design and production
Barbara Cahlander
Eric Hanson
Cynthia Rykken

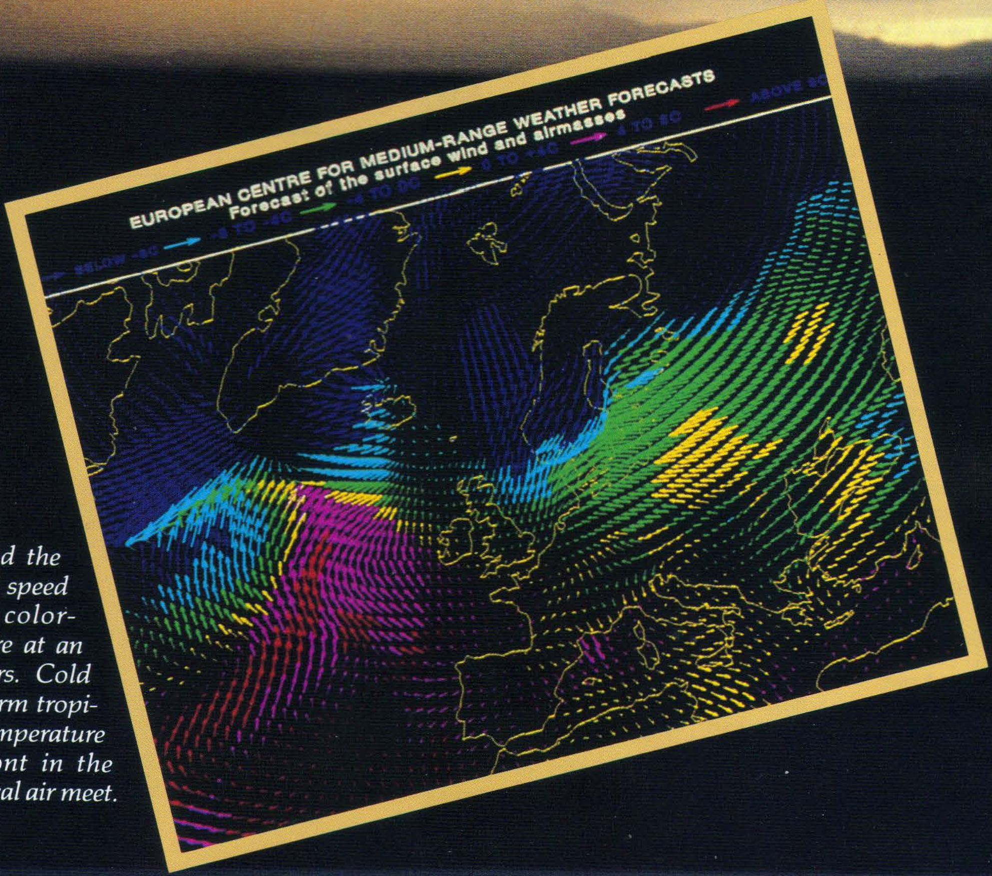

Computer map of Europe and the North Atlantic showing wind speed and direction. Arrows are color-coded to show air temperature at an altitude of about 1500 meters. Cold arctic air is shown in blue; warm tropical air in red. Note large temperature gradient across the cold front in the Atlantic where arctic and tropical air meet.

Multitasking the ECMWF weather forecasting model

David Dent and Rex Gibson
European Centre for Medium-Range Weather Forecasts
Reading, Berkshire, England

Computational weather forecasting today requires the fastest computers available, along with numerical models that fully exploit the available hardware power. During its 10-year history, the European Centre for Medium-Range Weather Forecasts (ECMWF) has advanced the science and art of weather forecast modeling such that a five- to six-day forecast in 1987 was more accurate than a three-day forecast in 1977. To accommodate the ever more complex numerical models used to prepare forecasts, ECMWF has upgraded its top-end computer systems successively from a CRAY-1 system acquired in 1979 to a CRAY X-MP/22 system in 1984 to the CRAY X-MP/48 system used at the center since 1986.

The current ECMWF production weather forecasting model first went into daily production use on a CRAY-1 system in 1983. At that time the model, which uses the spectral technique, had a resolution of 63 waves (T63 in our terminology) with 16 vertical levels. This corresponds approximately to a regular grid with points every 1.87° latitude and longitude. Three vertical levels are in the lowest kilometer of atmosphere. The remainder are spaced more widely, up to a maximum altitude of about thirty kilometers. In 1985 the model was moved to a CRAY X-MP/22 system and multitasked to run on the system's two processors. This allowed ECMWF to increase the model's resolution to T106, and still run it within operational time

Change date	Computer system	Forecast model	Vertical levels	Approximate resolution	Wall-clock time (hours)
7/79	CRAY-1	N48 grid point	15	1.875° × 1.875°	4
5/83	CRAY-1	T63 spectral	16	1.875° × 1.875°	5½
1/85	CRAY X-MP/22	T63 Spectral (single tasked)	16	1.875° × 1.875°	4
5/85	CRAY X-MP/22	T106 spectral (2-way multitasked)	16	1.125°×1.125°	5½
12/85	CRAY X-MP/48	T106 spectral (2-way multitasked)	16	1.125° × 1.125°	4
1/86	CRAY X-MP/48	T106 spectral (4-way multitasked)	16	1.125° × 1.125°	2½
5/86	CRAY X-MP/48	T106 spectral (4-way multitasked)	19	1.125° × 1.125°	2¾

Figure 1. Run-time comparisons of various resolutions of the ECMWF operational forecast model run on various computer systems.

constraints. In 1986 this model was moved to the present CRAY X-MP/48 system. Run times for various resolutions on the various computer systems are compared in Figure 1.

In anticipation of multiprocessor computer architectures, ECMWF began designing its production forecast model in 1979 so that it could be adapted for multitasking later. Care was taken to ensure that the model would be modular, with routines clustered into treelike calling sequences. Special code was introduced to manage the use of central memory and to address variables through the management package. Standard Fortran 77 was used throughout, except for addressing variables. Here, a Cray construct called POINTER, similar to BASED variables in PL/1, was used to address variables. This nonstandard feature was designed to simplify the separation of variables between tasks and to enable the use of simple, dynamic, and direct array structures.

By the beginning of 1986, the model had been developed to use the four processors of a CRAY X-MP/48 system efficiently. The number of vertical levels was increased to 19 and additional physical parameterizations were included.

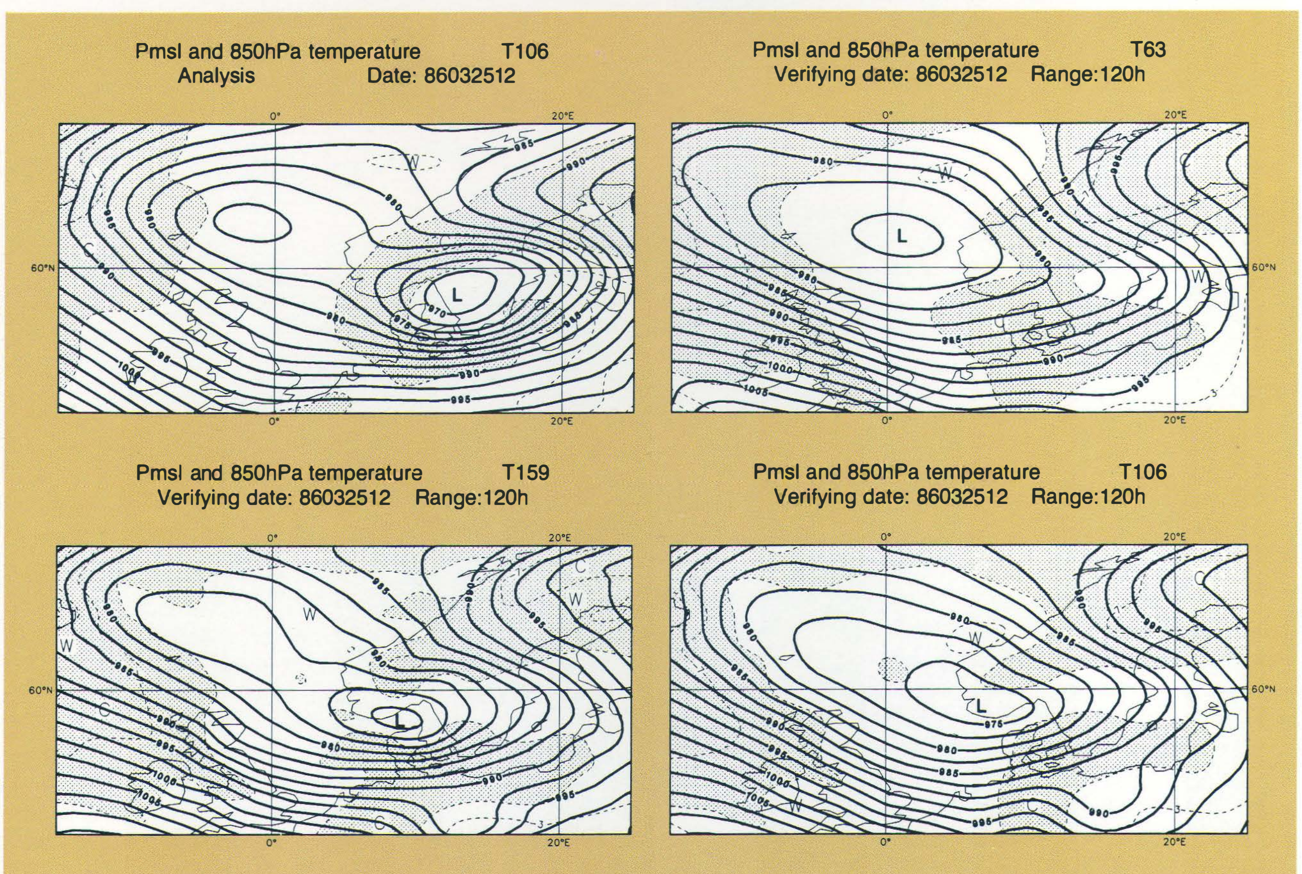

Figure 2. Relationship of model resolution to forecast accuracy. Upper left image shows the actual location, intensity, and elongated shape of a low pressure region. Clockwise from top right are images from models of increasing resolution. T159 represents an experimental model not yet in production use.

The T106 model with 19 vertical levels represents the current state of the production code, although development work continues. The most significant path of development, from the perspective of future computer architectures, is achieving finer resolution. At present, some experimentation is taking place at T159 that just can be made to fit into the CRAY X-MP/48 system. Examples of the more realistic fields that result from this finer resolution are shown in Figure 2.

Model structure

The model is organized so that work file data are processed in two scans for each time step. Within each scan, the I/O and computation are inside a loop over lines of latitude.

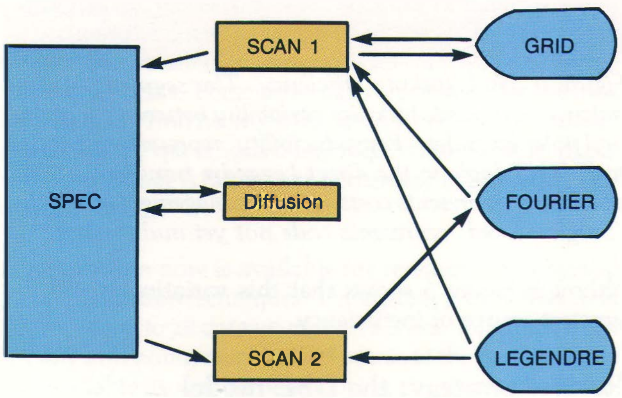

Figure 3. Overview of data flow within the ECMWF model.

That is, data belonging to all vertical levels for one line of latitude are collected into one record and stored in the work files. At the operational resolution (T106), there are 160 rows between poles and 960 time steps in a 10-day forecast. The work files occupy about 16 million words in an SSD solid-state storage device. Between the scans some computation occurs in spectral space for diffusion and semi-implicit time stepping. Data in the work files are split into three types: Fourier data, grid-point data, and Legendre coefficients. The relationship between the work files and the two scans of the model is illustrated in Figure 3.

The two scans comprise three processes. Scan 1 begins with dynamical calculations in Fourier space (Process 1) and requires input from the Fourier work file. Conversion to grid-point space is followed by all the physical parameterization and results in the grid-point work file being read and rewritten. Another Fourier transform (Process 2) then generates contributions for a Legendre transform. When all lines of latitude have been processed, the scan is complete and the Legendre transforms have generated the fields in spectral form (the memory space labeled SPEC in Figure 3). Scan 2 is relatively simple and computes the inverse Legendre transform (Process 3). Output is in the form of the rewritten Fourier file.

Multitasking strategy

Because multitasking the ECMWF forecast model first was carried out on the CRAY X-MP/22 system, memory use was a dominant consideration. Hence, a strategy was chosen that provided efficient computation combined with minimum memory requirement. This strategy involves splitting scan 1 into two processes, each of which runs on a separate processor. The first process includes the dynamics, the physics, and the subsequent FFT (Process 1). The second consists of the direct Legendre transform (Process 2).

Process 1 can be applied simultaneously to a Northern Hemisphere row and to its equivalent Southern Hemisphere row with a synchronizing point to ensure that the next process does not proceed until the data from Process 1 are ready. Each north-south pair of Fourier transforms provides components that are symmetric and antisymmetric about the equator. These are combined in various ways to contribute to the Legendre transform. This becomes Process 2, which can be multitasked so that the combinations are carried out simultaneously. Each copy of Process 2 updates one-half of the spectral arrays so that no danger exists of writing to the same memory simultaneously. Consequently, no need exists to use the Cray locking routines.

The inverse Legendre transform (Process 3) executed in scan 2 also can be distributed across two processors; the resulting data constitute the Fourier record that represents a north-south latitude pair. The small amount of work between scans also can be multitasked but is not considered here because it is not inside a latitude loop and has only a minor effect on timing.

The multitasking strategy for two processors is illustrated in Figure 4. This scheme can be generalized to any even number of processors by a logical splitting of the work at a higher level. The scheme involves giving each pair of processors a north-south latitude pair of data to work on independently. The loop over latitude rows now is shared by N processors. Generalizing scan 2 is trivial.

Using four processors, however, presents a potential hazard in the direct Legendre transform because two northern rows may update the same part of the spectral array simultaneously. Similarly, two southern rows may conflict with each other. This danger may be removed by the use of the Cray locking routines. These routines enable code to be protected against simultaneous execution on more than one processor. The Cray locking routines are found in UTLIB, a library containing multitasking and other kinds

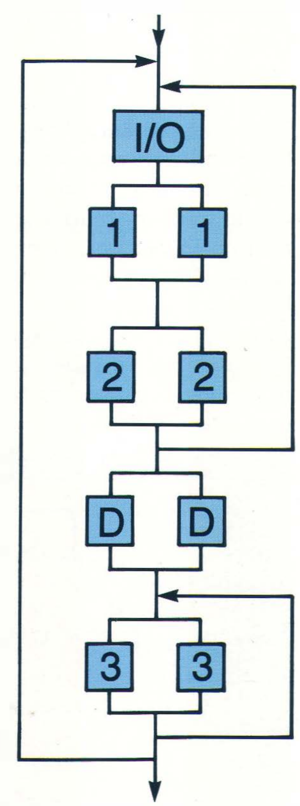

Figure 4. Multitasking strategy for two processors.

of routines. However, the locking routines must be used carefully to avoid either excessive multitasking overhead from too much lock control or substantial idle processor time when one processor is prevented from executing while a lock is held by another task.

Unfortunately, an indeterminacy now exists in the results because the order in which latitude rows update the spectral array no longer is predetermined by the code. The output from the Legendre transform is no longer unique because, on any finite-word-length computer, $A + B + C$ is not identical to $A + C + B$.

The differences are small and physically insignificant, but nevertheless extremely irritating from the testing point of view. EVENT logic may be used to force the computation into a certain order and hence re-establish reproducibility. This turns out to be an inexpensive remedy in an operational environment, and therefore is implemented as a default, although it easily can be disabled. All of the complicated logic involved in this procedure can be avoided, however, if the model data are completely memory resident, as would be possible with the large memory of a CRAY-2 system. In such a case, the transform can be reorganized so that the reproducibility problem does not occur.

Current performance

Using the CRAY X-MP hardware performance monitor, measurements of the production model environment reveal the following statistics:

Multitasking efficiency	= 90 percent
Execution rate	= 335 MFLOPS
Vectorization	= 99 percent of all floating-point operations

Time stamps inserted into the code at suitable points allow the granularity and the variability of task sizes to be measured. The small variations in Process 2 and Process 3 are due to memory bank conflicts. Process 1 shows a variation of 15 percent, mainly due to inequalities in work content (Figure 5). The workloads for different latitude lines are not equal, particularly in the physical parameterization of convection. Figure 5 clearly shows greater cost in equatorial latitudes compared to polar areas, where little convection occurs. Fortunately, the pairing of north-south latitude lines helps minimize this variation,

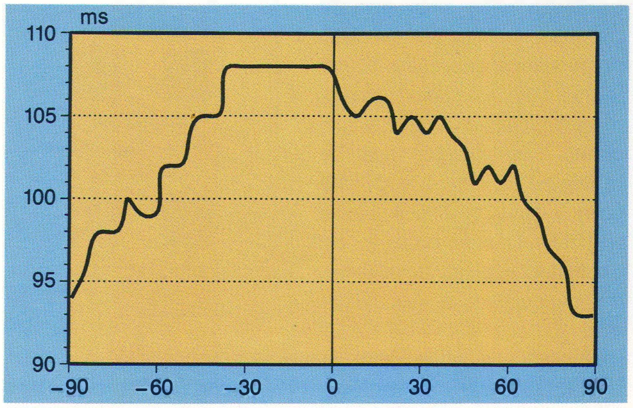

Figure 5. Computation time versus latitude for Process 1.

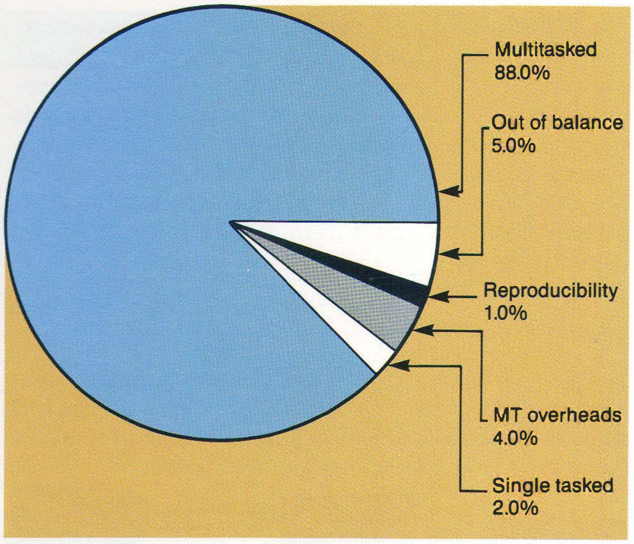

Figure 6. Multitasking efficiency. The segment "out of balance" represents task size variability between equatorial and polar latitudes. "Reproducibility" represents waste due to EVENT logic in the direct Legendre transforms. "MT overheads" represents costs of task management and locks. "Single tasked" represents code not yet multitasked.

although Figure 6 shows that this variation is still the greatest source of inefficiency.

Revised strategy: the DSC model

Having identified several places where inefficiencies exist, the multitasking strategy can be revised. Because the CRAY X-MP/48 system has a larger memory than the CRAY X-MP/22 system, a dynamic strategy becomes attractive. This revised approach has been labeled the DSC dynamic scheduling technique.

The DSC technique retains the north-south pairings of latitude rows for computational economy, but computes them sequentially, that is, north followed by south. To be more precise, Process 1 is executed sequentially, first for the northern row, then for the equivalent southern row. All data then are available for Process 2, the direct Legendre transform, which can be executed twice to complete the contributions to the spectral array from the pair of latitude rows. Once this series of calculations is complete, it can be repeated for another, essentially independent, pair of rows. This approach offers the advantage that no synchronization is necessary after each pair of rows because the necessary data must be available. Because individual processors can be given independent pairs of rows to compute, the only synchronization necessary is at the end of the scan to ensure that all rows are complete. The decision as to which pair of rows to compute can be made at run time. Therefore, the strategy is completely flexible, allowing for any number of available processors. A modest penalty of increased memory use exists with this approach, however, because each task must retain a Fourier record and a pair of symmetric-antisymmetric components from the Fourier transforms. The static multitasking strategy requires these data only for each pair of tasks.

The I/O scheme for the model work files must be revised because it is based on the static multitasking strategy.

Because version 1.15 of the Cray operating system COS offers an alternative access method that achieves savings by queuing I/O requests to minimize system overhead, the revision was planned to make maximum possible use of this facility. This approach is attractive particularly for random I/O requests to SSD storage in circumstances where several requests may be outstanding at one time. The revised scheme, therefore, always attempts to keep the available buffer space full by issuing as many READ requests as possible, perhaps even in advance of requests. If sufficient buffer space is allocated, by means of a run-time parameter, then double buffering is achieved.

Because the I/O master routines are called from within the multitasked code, locks are used to enforce single-threading. However, this does not cause much inefficiency because the current execution of the I/O routine may do work on behalf of subsequent calls to the same routine from other tasks. These other calls are delayed by the lock but normally find no work left to do when the lock finally is released. By queuing as many requests as possible, maximum advantage is taken of the economy in operating system overheads.

A new option now is available for research into physical parameterization techniques. The current production code allows access to all data on the current line of latitude only. For any grid point, neighboring points to the east and west are available as well as neighboring points vertically below and above. However, because of the I/O structure, data to the north and south cannot be accessed. As the model's grid resolution becomes finer, this limitation probably will become more serious. However, the revised I/O scheme allows records, and hence lines of latitude, to be retained in buffer space until the running tasks signal that they no longer are needed. Providing that sufficient buffer space exists, data to the north and south for any desired number of latitude lines can remain available. A run time parameter controls this option.

The performance analysis has been repeated on short tests using the DSC method. However, standard I/O techniques were used because the queued I/O access method was not yet available at ECMWF. The performance gains so far are

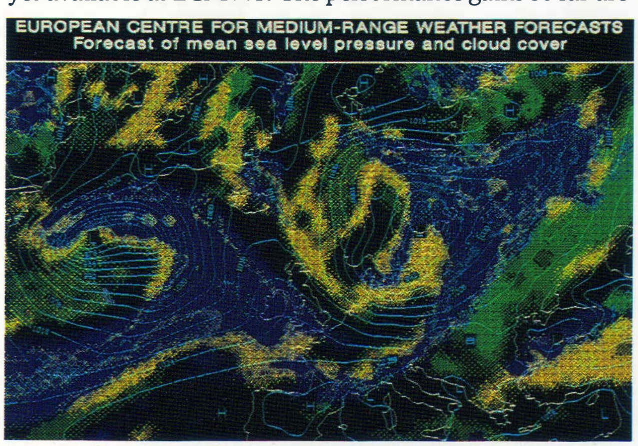

Computed map of Europe and the North Atlantic showing pressure contours and cloud fields for same situation as shown in the photograph inset on page 2. Upper clouds are shown in blue, midlevel in green, and low in yellow. Convective clouds are shown in brown.

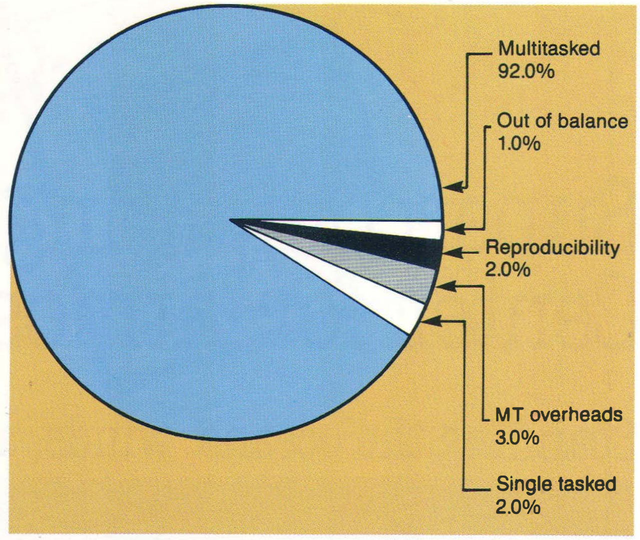

Figure 7. Multitasking efficiency of the DSC model.

small but are expected to improve with further code enhancements. The multitasking efficiency has increased to 92.5 percent. Figure 7 shows the out-of-balance overhead reduced to 1 percent and the multitasking overhead reduced to 3 percent. The multitasked portion has increased to 92 percent.

Summary

Although the inefficiencies in the current operational code are relatively small, they likely will become more objectionable in the future when the model is run on computers with many more processors. The revised DSC scheme, however, may achieve a reasonably high multitasking efficiency on future multiprocessor computer systems. Additional gains may be made by eliminating many of the locks currently present only to enable statistics gathering. The residual singletasked code can be tackled either by the conventional macrotasking techniques or by use of microtasking where granularity is small.

The problems of the remaining out-of-balance waste and idle time due to forcing reproducibility are being tackled by using microtasking within the existing macrotasking structure. The judicious use of microtasking facilities will enable ECMWF to take advantage of idle processors that become available during the direct Legendre transform and to speed up the computations in which waste is occurring. □

About the authors

David Dent has a background in mathematics and has worked for a number of years in ECMWF's operations department. He now deals with computing matters on behalf of the research department, with special responsibility for technical aspects of the ECMWF forecasting model.

Rex Gibson is an associate fellow of the Institute of Mathematics and its Applications. As head of the meteorologial applications section, he has overall responsibility for the operational codes that support numerical weather prediction at ECMWF. He was responsible for the original design and logistics of the computer code for ECMWF's forecast model, and developed the first multitasking version, which ran on a CRAY X-MP/22 computer system.

Supercomputing at the National Center for Atmospheric Research

Scanning the oceans, winds, and skies with Cray systems

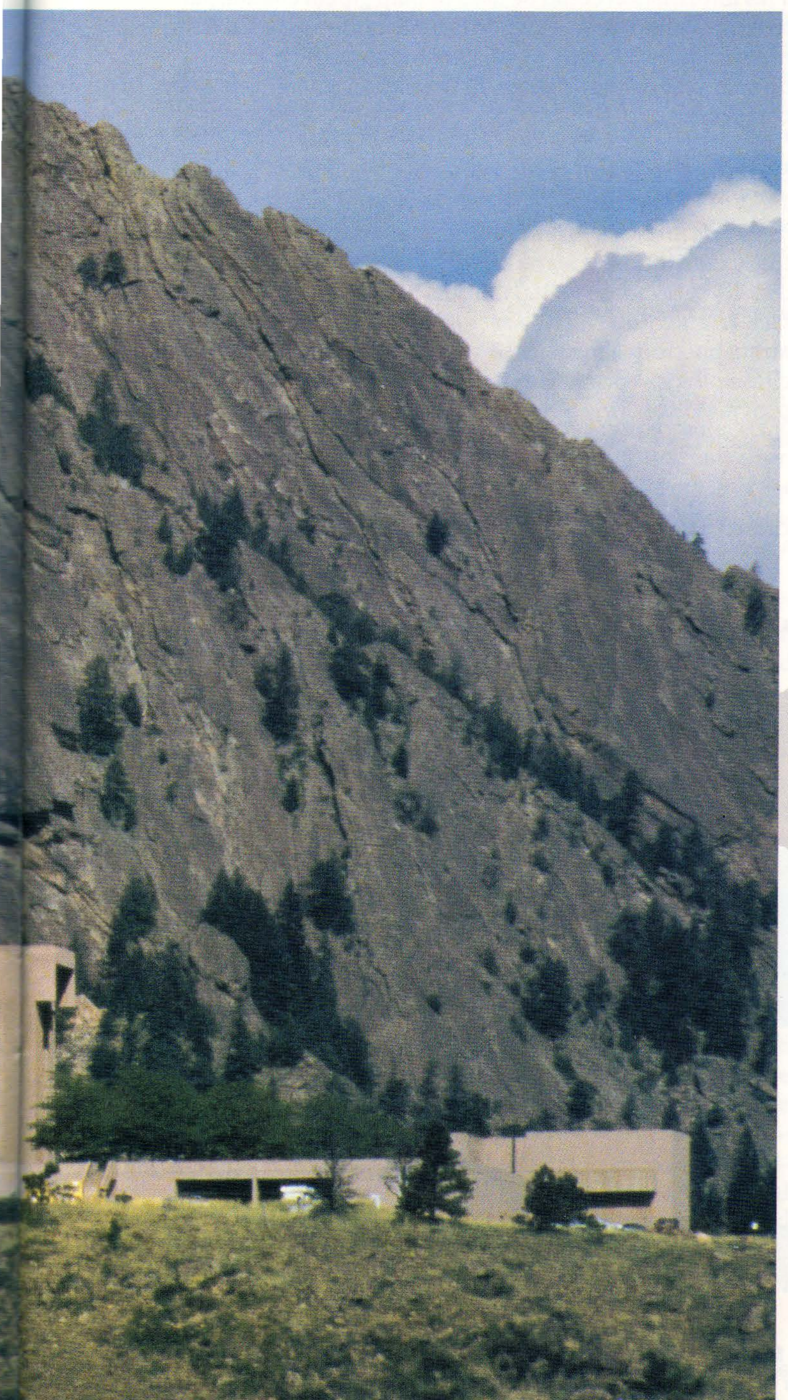

Perched atop a Colorado mesa, the slablike stone walls framing the National Center for Atmospheric Research (NCAR) form a monument to the idiosyncrasies of nature. Designed by world-renowned architect I. M. Pei, the coral towers of the Mesa Laboratory are surrounded by 380 acres of forestland. While deer graze on the wild grass outside NCAR, inside the geometric fortress scientists analyze the complexities of the atmosphere using tools such as the center's two Cray supercomputers.

Probing the globe

For 27 years, NCAR researchers have been employing sophisticated scientific instruments to unfold the mysteries of the atmosphere, looking into the dynamics of nature's physical processes, examining interactions of the sun, oceans, plants, and animal life. The speed and memory of NCAR's Cray systems allow researchers to model large-scale phenomena with great accuracy — and to develop coupled models that simulate the interaction of intricate systems, such as the oceans and winds. Climate specialists apply computational models to simulate the dynamics of clouds, solar radiation, and carbon dioxide.

To examine the interacting systems of our planet, NCAR's four research departments harness the most powerful supercomputing systems available. NCAR's Scientific Computing Division (SCD) serves each area by processing and archiving massive amounts of collected data, as well as assisting with the development of complex numerical models of the atmosphere's fluid motions. "A single processor of the CRAY X-MP system is one hundred times faster than typical minicomputers and at least one thousand times faster than typical desktop computers," says SCD Director Bill Buzbee. He adds that 75 percent of the CRAY X-MP system time at NCAR is spent performing calculations that require more than one hour of computer time, problems that would be almost impossible to solve on other computer systems.

Case study: exploring the El Niño mystery

One such complex problem involves modeling the global atmosphere and oceans to unravel the scientific mystery

The Colorado Rockies provide a backdrop for NCAR's Mesa Laboratory. Photo courtesy of the National Center for Atmospheric Research/National Science Foundation.

of El Niño, a climatic phenomenon blamed for peculiar weather around the globe. By simulating El Niño with NCAR's CRAY X-MP system, Bob Chervin of the Climate and Global Dynamics Division is studying the validity of assertions that unusually warm Pacific water causes bizarre worldwide climatic changes.

Some have linked El Niño to droughts in Australia, Indonesia, and southeast Africa, as well as floods in Argentina and the United States. El Niño has been blamed for brush fires, monsoons, famine, heavy snowfalls, and thunderstorms all over the world. Some even have claimed that this phenomenon indirectly has caused snakebites, back injuries, and the bubonic plague.

El Niño, which in Spanish means the Christ Child, originally referred to the December warming of the normally cool Pacific water off the west coast of South America. Today scientists use the term to describe severe weather anomalies caused by the interaction of the tropical Pacific Ocean, the atmosphere, and the annual cycle of seasons.

El Niño events begin when warm surface water travels from the western equatorial part of the Pacific Basin to the equatorial regions off South America's west coast. During El Niño events, the warm water disrupts the usual coastal upwelling of nutrient-rich cool water. Not only is the warming of the Pacific linked to tropical storms, but the reduced nutrient content of the coastal waters leads to the widespread mortality of plankton, fish, and guano-producing birds, which disrupts the economies of countries such as Ecuador, Peru, and Chile.

Chervin has developed a highly parallelized version of the NCAR Community Climate Model and also is optimizing an ocean model to determine if there is something distinct about the atmosphere at the time of El Niño events. This information could be used to map global climate patterns and improve worldwide climate outlooks. Chervin's model relates 30 years of varying ocean surface temperatures to year-to-year changes in time-averaged atmospheric conditions.

By collecting ocean and atmospheric data from January 1950 to December 1979, Chervin is able to study seven moderate to strong El Niño events that occurred during this time. One model integration simulates atmospheric conditions in response to 360 months of changing ocean surface temperatures. The second determines the range of atmospheric response to 30 identical annual cycles of climatological ocean forcing. Thus, the first integration provides year-to-year anomalous forcing through realistic patterns of global ocean surface temperatures.

"This comparison determined if the variability of the ocean caused a broader spread of atmospheric states. I found that

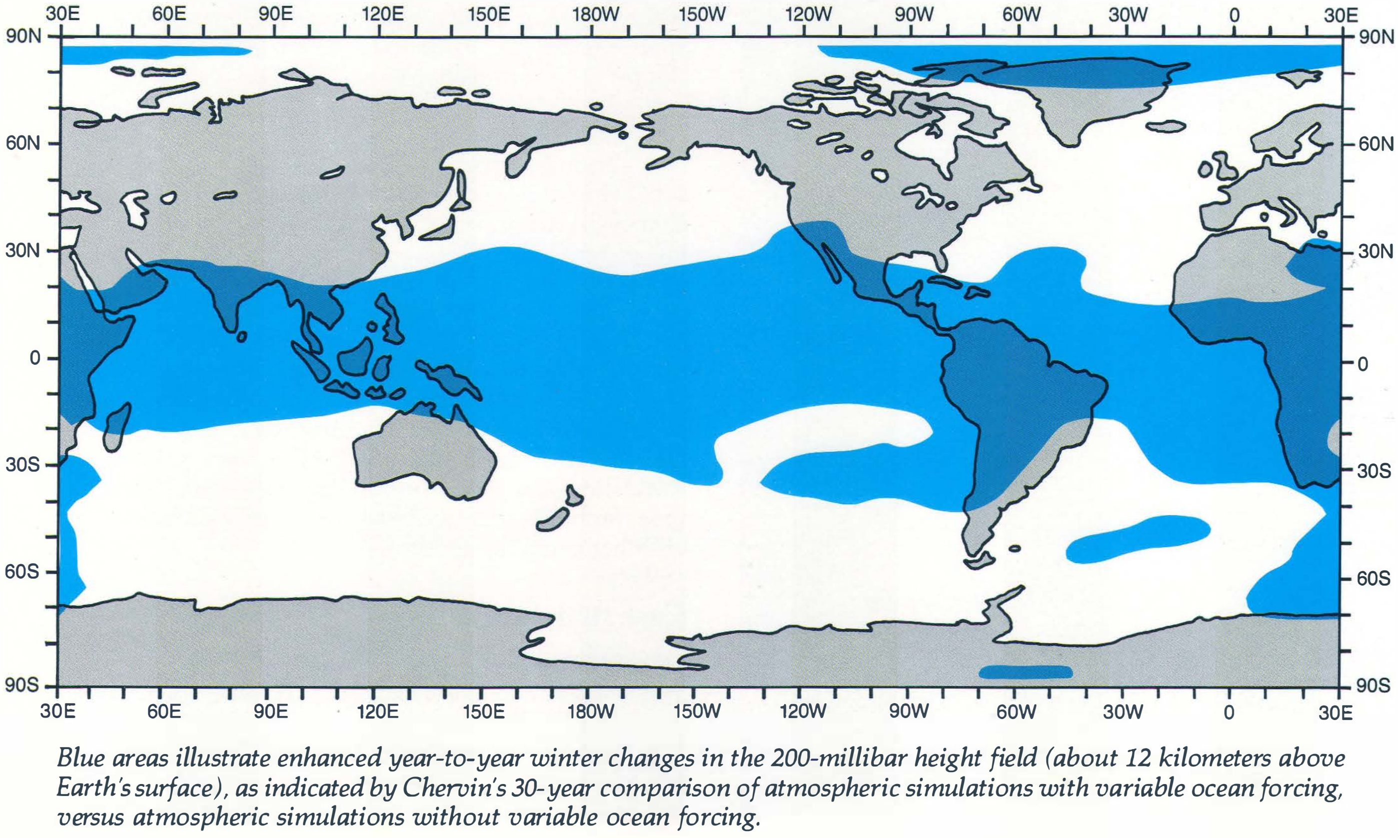

Blue areas illustrate enhanced year-to-year winter changes in the 200-millibar height field (about 12 kilometers above Earth's surface), as indicated by Chervin's 30-year comparison of atmospheric simulations with variable ocean forcing, versus atmospheric simulations without variable ocean forcing.

more extreme events occurred only in certain geographic regions with the variable tweaking of the atmosphere by the ocean," says Chervin. His findings indicate there is no typical El Niño pattern occurring in the middle latitudes, although a distinct pattern does exist in the tropical region.

"A widespread heating of the entire tropical band occurs because of this additional energy source. But as you move to the higher latitudes, apparently scrambling mechanisms confuse any El Niño-type of signal that might propagate from the Tropics," says Chervin. While El Niño has been linked to severe weather in the United States, Chervin sees no strong evidence to confirm such links. "For example, extremely cold temperatures on the East Coast were attributed to an El Niño event that occurred in 1976-77," says Chervin, explaining that while there have been attempts to link the East Coast weather to the El Niño that was taking place, the next time an El Niño took place, in 1982-83, temperatures were warm on the East Coast. "It would be nice if there were distinct links, but I'm afraid that is not the case," he says.

Chervin optimized the model to speed up computation. "Although each integration required approximately 180 CPU hours to simulate 30 years of atmospheric behavior, the matching of the parallelized model structure and the Cray system architecture allowed the research to be completed about 3.7 times faster than would be possible in a purely sequential mode," says Chervin.

Chervin is collaborating with Albert Semtner of the Naval Postgraduate School to develop another multitasked model — a primitive-equation global ocean model. "It has become a hot-performing model," says Chervin. "Preliminary results are impressive, as indicated by the model's sustained performance rate of 450 MFLOPS." His ultimate aim is to develop a highly efficient, highly parallelized coupled ocean and atmospheric model to look further into the dynamics of El Niño events and other questions related to our coupled climate system.

Forecasting space weather

While scientists like Chervin are modeling the global atmosphere and oceans, a group from NCAR's High Altitude Observatory is working on a less earthly task. The group is predicting weather in the outer reaches of the atmosphere near space. Senior scientist Ray Roble, a theoretical investigator for the NASA Dynamics Explorer satellites, has developed a model that illustrates the general circulation of the thermosphere. He compares model and satellite data to study the outer layer of the atmosphere, 97 to 500 kilometers above Earth's surface, using NCAR's CRAY X-MP system.

The Thermospheric General Circulation Model calculates the thermosphere's temperature, circulation, composition, and chemistry, as well as its response to extreme ultraviolet variations in solar flux and auroral activity. These simulations describe the impact of solar variability on the upper atmosphere and how these effects can be transferred to the lower atmosphere. "Solar variability probably does not affect weather near the surface of Earth, but it does affect man's activities in space," says Roble. In 1979 solar activity increased, causing the outer atmosphere to expand, resulting in increased orbital drag on Skylab. This caused it to re-enter Earth's atmosphere sooner than expected, creating considerable uncertainty about its re-entry location.

The thermospheric model was derived from a numerical weather prediction model for the lower atmosphere. "We stripped out all the mountains, clouds, and radiation processes that were in the model designed to represent lower atmospheric physics, then moved it to the upper atmosphere and modified it to include auroral processes, ionospheric interactions, some fast diffusive properties, and chemical reactions that occur in the upper atmosphere," explains Roble. The model divides Earth into a grid with five-degree spacing in latitude-longitude. Twenty-five vertical layers extend upward from 97 to 500 kilometers in altitude.

According to Roble, the complexity of this model requires the power of a Cray system. "The Cray system has offered speed and versatility that allow us to incorporate more physical processes, providing a deeper understanding of the atmosphere," he says. After simulating the impact of solar ultraviolet radiation and solar wind variability on the circulation, temperature, and compositional structure of the thermosphere, Roble then compared his model results with data collected from NASA satellites and ground-based observatories to gain a better understanding of the physical processes taking place. "There's no shortage of data from space," says Roble, explaining that data sources include satellite systems, ground-based incoherent scatter radars that probe the upper atmosphere, and radio and optical techniques that include a network capable of measuring wind speeds and temperatures 300 kilometers above Earth.

So far, Roble's model has described strong interactions of the thermosphere's neutral gases and solar variability. It also has been used to study the impact of particles of the solar wind that bombard Earth, causing the aurora borealis. Variability in solar wind properties regulates the intensity of the aurora during geomagnetic storms. Roble plans to take advantage of available resources to increase the

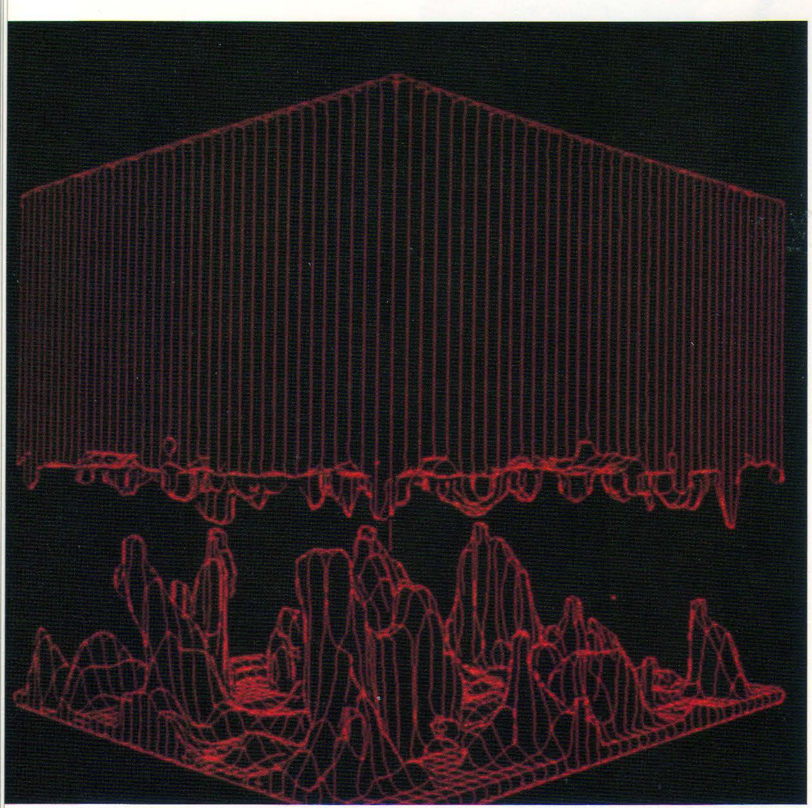

A three-dimensional plot of a curve of constant temperature (302° K) made directly from large eddy simulations. This plot reveals the turbulent structure of planetary boundary layer convective thermals. Photo courtesy of the National Center for Atmospheric Research/National Science Foundation.

number of physical processes in the thermospheric model. "We're putting in an interacting ionosphere and we plan to extend the model downward into the mesosphere to 65 kilometers above Earth." Also, the thermospheric model will be coupled with a magnetospheric model that will describe how solar wind energy is tranferred into Earth's atmosphere. "Eventually, we hope the model will describe realistically how solar wind plasma and solar ultraviolet radiation variability will affect Earth's dynamical, radiational, and chemical processes," explains Roble. This will help scientists better understand the propagation of radio communication signals, satellite drag, and the transport of chemicals produced in the upper atmosphere.

"One chemical produced in the aurora, nitric oxide, can be transported to the vicinity of the polar night region where it does not come in contact with sunlight. It then can be transported downward to the stratosphere where it can interact catalytically with ozone. One of the theories proposed to explain the ozone hole in Antarctica links the problem to the downward transport of this chemical," says Roble. While he does not believe nitric oxide is a likely mechanism for ozone depletion in Antarctica, Roble thinks it may affect ozone in the upper atmosphere.

Although research in the thermosphere may seem far removed from Earth, the basic applications are extensive. Says Roble, "It is an exciting area of research. It gives me an opportunity to describe numerically the natural processes that occur in the upper atmosphere to give us a better understanding of their effects on the whole atmosphere."

Mapping turbulent eddies

Another critical component in worldwide climate forecasts is the transport of heat and moisture throughout the atmosphere. The traditional method of investigation involves using research aircraft to gather data. This procedure is expensive and inefficient, however, because the full three-dimensional flow field cannot be measured physically. Chin-Hoh Moeng of the Mesoscale and Microscale Meteorology Division prefers using NCAR's CRAY X-MP system to take snapshots of the atmosphere's complex eddy patterns. "Instead of flying out there to collect the velocity and temperature fields, I can sit at the terminal and analyze data generated by the Cray system," she says.

Solar radiation is the heat source for atmospheric motion, though the sun does not directly heat the atmosphere. Instead, the sun heats Earth's surface and turbulent eddies in the planetary boundary layer (PBL) transport this heat aloft. These eddies also carry moisture from Earth's surface to the atmosphere, creating clouds. "We don't know much about this type of turbulence because it is very complicated, so we study the turbulent flow field with simulated model data," says Moeng. "This type of three-dimensional flow structure, which is otherwise impossible to obtain, allows us to study and parameterize the characteristics of the turbulent PBL."

According to senior scientist John Wyngaard, the power of today's high-speed computers can be credited for advances in turbulent eddy simulation. "We have had the equations since the mid-nineteenth century, but you can do very little with them mathematically," he says. "The supercomputer is allowing us to progress now, through very accurate numerical approximations."

The flow structure of the simulated eddies is modeled by solving the Navier-Stokes equations. These equations describe a turbulent flow field for a heated Earth surface, a process much like heating the bottom of a cup of water to generate heat circulation. "Because of the warmer surface below, the air becomes unstable. The less dense warm air tends to move upward, while the heavier cold air tends to move downward. This disturbs the air and generates turbulence," explains Moeng.

Moeng developed the turbulent eddy model five years ago using a three-dimensional grid with 40^3 points. The increased power and memory of NCAR's CRAY X-MP system enabled her to expand the lattice to 96^3 grid points.

Moeng uses 2 million words of internal memory, and reduces I/O time by employing 30 million words of storage from the SSD solid-state storage device. Pure processing time for the larger grid is 33 seconds per time step. "As a result of the refined grid, smaller eddies are resolved; consequently, the results are more reliable," says Moeng.

The results of the turbulent flow model will help refine the larger forecast models by adding proper heat and moisture transfers. Next, Moeng plans to use the turbulent eddy model to study the effects of stratus clouds that form above coastal California during the summer. "These clouds can affect the climate by blocking the sun from heating the ocean, which changes the atmospheric radiation," Moeng explains. By simulating cloud data otherwise difficult to gather, Moeng's results will help others perfect larger global forecast models.

Reaching research goals

Just as supercomputers are increasing in power, the climatic models created at NCAR are growing in complexity. Scientists can look at global processes as linked systems that involve the dynamic interactions of the air, seas, plant, and animal life. Answering the need to understand Earth's interrelated systems, NCAR has grouped its research areas into four divisions based on goals of the National Science Foundation. The four divisions rely heavily on NCAR's supercomputing resources.

In the Climate and Global Dynamics Division, researchers model the dynamics and thermodynamics of the atmosphere and oceans and their interactions with land surfaces and the biosphere. To develop better forecasting models, they compare worldwide weather forecasts with physical events, scrutinize the forces of tropical oceans, the impact of cloud changes on climate, and the influences of increased amounts of carbon dioxide in the atmosphere. In the Atmospheric Chemistry Division, scientists study the atmosphere's chemical reactions. After collecting data from satellites, high-altitude aircraft, and optical and chemical instruments, they study chemical reactions under varied atmospheric conditions, and develop models to analyze the causes of acid rain and ozone depletion. The Mesoscale and Microscale Division investigates medium-sized storms, as well as smaller phenomena such as cloud processes and hailstone formation. Scientists determine the evolution of storms and explore banding, precipitation formation, the water content of clouds, and the patterns of turbulent eddies. The High Altitude Observatory is NCAR's solar research division. Here scientists investigate the nature of solar-terrestrial interactions and the behavior of the sun as a star. An example of such research is Roble's general circulation model of the thermosphere developed with the CRAY X-MP system. Serving these four research divisions is the Atmospheric Technology Division, a facility that develops and maintains state-of-the-art data collection equipment, such as satellites, aircraft, and radar.

Scientific computing

NCAR's Scientific Computing Division supports the four research divisions in five main areas of study: climate, cloud physics, oceanography, upper atmosphere, and regional modeling of phenomena such as thunderstorms and tropical storms. "These disciplines correlate with the National Science Foundation's long-range research plans," SCD Director Bill Buzbee says. He added that of the division's 1000 active users in 1987, two-thirds were from universities, one-third from NCAR.

In 1977, NCAR became the second site to receive a CRAY-1 system. "The fact that NCAR purchased one of the first Cray systems shows this community's need for high-speed computing," Buzbee says. Due to increasing demands for supercomputing power, NCAR recently purchased a CRAY X-MP/48 system. The facility today includes three components: high-speed computers, communication systems, and network servers. The CRAY-1 system and CRAY X-MP/48 system with 256-million-word SSD solid-state storage device are host systems to an IBM 4381 front end that provides interactive computation for local and remote users. The SSD makes three-dimensional simulation possible, a feature that is essential to atmospheric science, according to Buzbee.

Mission: desktop supercomputing

Looking toward the future, NCAR's supercomputing division is dedicated to supporting large models and data sets. A principal mission is to provide distributed computing, which allows local users to work primarily from desktop systems. Using this process, local users carry out simulations on supercomputers, access NCAR's mass storage unit, or use SCD's film recorder and laser printer.

"Our dream is to have the supercomputer appear as an attached processor to each scientist's desktop. We would like our mass storage system to be available as an auxiliary filing system to each desktop computer, and we would like our output servers to be optional devices to the desktop computers," Buzbee says. He says this dream may materialize within three years, adding that the emergence of a UNIX standard operating system, along with the development of standard protocols such as TCP/IP, will pave the way for widespread distributed computing. The result: easier access for scientists, which in turn may provide more answers about our planet and its interactions with living things, the sun, oceans, winds, and skies.□

Computational forecasting
The expanding role of supercomputing in Canada

Iain B. Findleton, Canadian Meteorological Center, Dorval, Quebec

In the fall of 1983 the government of Canada took an important step toward reinforcing the safety and security of the Canadian public and advancing Canadian scientific research capabilities. As part of the Numerical Weather Prediction program of the Atmospheric Environment Service, Environment Canada installed a CRAY-1/S computer system at its Dorval, Quebec site. At the time of its installation, this system was one of the most advanced computers of its type. But developments in the computer industry proceed at a rapid pace. The CRAY-1/S system has since been replaced by a CRAY X-MP system that delivers, at lower cost, much improved performance.

Supercomputer impact

The installation of the CRAY-1/S system in 1983 marked the first phase of a planned series of enhancements to the computing facilities available to the Atmospheric Environment Service (AES) through 1990. The system was installed to make available the necessary computer power to implement a series of advanced weather forecasting programs. The cost of the computer facilities was justified by the promise of more accurate and timely weather services delivered to the Canadian public.

In any country, a supercomputer represents valuable scientific resource. Recognizing this fact, the government decided to make available to Canadian universities 10 percent of the supercomputer resources at Dorval. The policy was intended to stimulate Canadian excellence and leadership in selected areas of basic research.

CRAY-1/S system performance

In evaluating the performance of the CRAY-1/S system, it is useful to examine the activities that took place on the system during its lifetime. Overall resources on the system were allocated in five areas: meteorological production (51 percent), meteorological research (19 percent), climatological research (15 percent), Canadian universities (10 percent), and other government research (5 percent).

Figure 1 shows the highly smoothed history of the rate of job submission to the Cray system by all users of the system. This curve has the classic shape of the logistics curve that describes a facility's saturation. While an exact definition of what would constitute saturation of a supercomputer is a matter of some debate — noting that until April of 1986 users essentially were unconstrained in their access to the system — it appears from the curve that the facility became saturated from the user point of view around April 1985. Considering that the machine was installed in November of 1983 and commissioned for operational use in June of 1984, the rate of saturation of the Cray system was rapid indeed. The average saturation rate of supercomputer facilities has been studied by

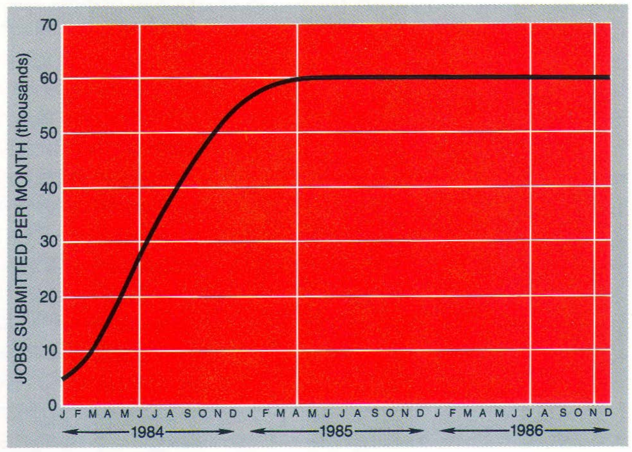

Figure 1. Development of Dorval supercomputer workload.

market researchers working for supercomputer manufacturers, and the results indicate about an 18-month lifetime for new facilities before they require enhancements of some kind.

Weather forecasting

The main use of the CRAY-1/S computer system by AES operational components was the development and implementation of new technology. When the Cray system first was commissioned for use in June 1984, the debut coincided with the movement toward routine production of an improved weather prediction model. The new weather model incorporated scientific developments that had been available in the research community and in the operations of other countries for a few years. Canada had not been able to exploit these developments because of inadequate computer power, a barrier that was removed by the installation of the CRAY-1/S system.

In early 1986, the AES introduced another more advanced weather prediction model that had been in development since 1976. This improvement in weather forecasting technology is the start of an effort to provide accurate local short-range weather forecasts. Weather events that have the greatest impact on the safety and security of the general public tend to develop rapidly and to affect limited geographical areas.

Climate research

Global climate research is another meteorological application of supercomputer power that has been advanced through the use of the CRAY-1/S system at Dorval. Aside from efforts to achieve a fundamental understanding of the way climate has developed over history, climate research primarily supports policy analysis relating to the effects of human activity on the atmosphere. The Cray system provided Canadian scientists with the basic simulations of atmosphere behavior to support research projects addressing particular phenomena. The system also was important in developing research programs addressing issues such as the "greenhouse effect" caused by the build-up of carbon dioxide in the atmosphere.

Longer-term activities in the area of climate prediction are underway. The computational requirements of climate-change prediction systems are enormous and cannot be met without the use of supercomputers. The range of activities connected with climate prediction models includes studies on the effects of manmade chemicals on the ozone layer and the long-term atmospheric effects of various types of chemical pollutants.

Environmental emergency response

A significant activity on the CRAY-1/S system was the development of programs to track hazardous airborne materials released into the atmosphere either inadvertently or purposely. Environmental emergency response models deal with two aspects of such events: the local short-range dispersion and the long-range transport of the dangerous material.

Environmental emergencies include releases of toxic chemicals through transportation accidents or mishaps at production and processing installations and releases of radioactive materials from decayed satellites, nuclear power installations, and weapons tests. Civil defense authorities can profit from the availability of timely knowledge of the expected track, the rate of dissipation, and the intensity of clouds of dangerous substances. Generally, the gain in lives saved and the effectiveness of prevention and recovery measures are related directly to the accuracy, timeliness, and horizon of the available advice about the hazard. All of these elements of the emergency response system depend on the sophistication of the computer model used to track the hazard. Without the use of supercomputers, only the most elementary advice could be provided.

Global response

Major geopolitical issues such as the long-range transportation of pollutants, acid rain, the effects of large-scale construction on local climate, global surveillance of food production, and the study of drought confront the Canadian public and its government on a daily basis. International responses to these issues have become increasingly technical during the past 20 years. The broad-scale international interdependence of human beings is causing governments to seek scientifically sound policy options that protect the interests of their citizens and alleviate the onerous consequences of human activity on the environment. In this arena, the use of supercomputer technology and large-scale atmospheric models has become de rigueur. Results of policy studies that do not have a basis in soundly formulated scientific analyses are overshadowed in the international scientific and political forum by studies that have sound support.

University-based research

The CRAY-1/S system was an integral part of the basic research carried out at a number of Canadian universities. University professors and their research assistants used the system to investigate high energy physics, molecular biology, combustion, molecular dynamics, astrophysics, quantum chemistry, and fluid dynamics.

The impact of supercomputer resources is such that wide areas of basic research simply cannot be carried out without the computational power such resources provide. State-of-the-art research could not have been carried out in Canada in these fields without the availability of the Dorval Cray systems.

Nonweather government research

The major nonmeteorological clients of the Dorval Cray systems from within the ranks of the Canadian government have been research workers investigating either hydrological or oceanographic problems. The hydrologists are working on reservoir and flood control problems using computer simulation models of river drainage basins. The oceanographers are studying ocean circulation and atmosphere-ocean interaction problems. Both types of study require extensive computations on highly evolved computer models that cannot be carried out on conventional computers in reasonable time.

Goals accomplished with the supercomputer

The direct benefits to Canada of the acquisition and installation of the CRAY-1/S and CRAY X-MP supercomputers at Dorval can be summarized as

- ☐ A direct improvement in the accuracy of weather forecasts
- ☐ Implementation of a sophisticated emergency response system
- ☐ Attribution of international credibility in climate research
- ☐ Enhancement of Canadian capability in environmental policy analysis
- ☐ Stimulation of Canadian excellence in research

The main program goal was the delivery to the Canadian public of more accurate weather forecasts. The rationale followed by the AES is that using the most advanced weather prediction technology in routine forecast production results in fewer errors in the basic forecasts used to prepare weather advisories for the public. Based on the rather long previous experience with this process, the AES was able to predict an improvement in forecast accuracy. Continued performance also is safely predicted to result from currently projected enhancements to the supercomputer facility.

Another dimension of the expectations for improved weather forecasts relates to the time horizon of the forecasts. Weather forecasts are termed as being short-range, medium-range, or long-range. The time horizons referred to by these labels vary with the particular community of meteorologists using them. In the Canadian context, short-range forecasts refer to a one-hour to two-day time horizon, medium-range forecasts refer to a three- to ten-day time horizon, while long-range forecasts refer to seasonal and climatic prediction activities.

Direct improvements in short- and medium-range forecasts come from improvements in the scientific quality of the main synoptic forecast models. The technology used for the time horizon beyond two days has improved to the extent that forecasts for day five now are showing significant skill, a marked change from the time prior to the supercomputer acquisition. Improvements in the short-range forecasts with a one-day time horizon are being pursued through the implementation of a new mathematical approach that promises remarkable improvements in forecast accuracy. Before the Cray system installation at Dorval, computer models for long-range forecasting did not exist in Canada.

Implementing a sophisticated environmental emergency response system presupposes an existing advanced meteorological analysis and prediction system. For this reason, to the extent that it has made possible improvements in the accuracy of weather forecasts, the Cray system also has contributed to the protection of the public in the face of hazards transported by the atmosphere.

The Dorval upgrade to a CRAY X-MP/28 computer system

The CRAY-1/S system was replaced in November 1986 by a CRAY X-MP/28 system with SSD solid-state storage device. The CRAY X-MP/28 system offers much higher performance than the CRAY-1/S system and supports a much larger memory. The larger memory is significant in enabling users to run larger, more refined, and more accurate weather prediction models. The benefits of the larger memory and greater computational rate also extend directly to the other activities of the university and research communities that make use of the Dorval computer facility. Direct benefits include improved modeling of Earth's climate, acid rain potential, and the atmospheric transport of toxins for developing emergency response procedures.

The supercomputer program as originally approved extends into the 1990s. The evidence so far suggests that the program not only is meeting its stated objectives, but also is generating significant additional benefits to the scientific community. The increasing demand for access currently is controlled by priority-based rationing, a policy that will continue to be needed over the lifetime of the current program. Beyond the current program, economic and technological considerations will determine the future direction of the Dorval supercomputer installation.□

About the author

Iain Findleton is chief of the Canadian Meteorological Center's computer center, the Centre de l'Informatique de Dorval (CID), located in Dorval, a suburb of Montreal, Canada. CID operates a CRAY X-MP computer system in support of the operational and research programs of the Atmospheric Environment Service of Environment Canada. A meteorologist with 18 years of experience in weather forecasting, Findleton has administered the computer center since January 1986.

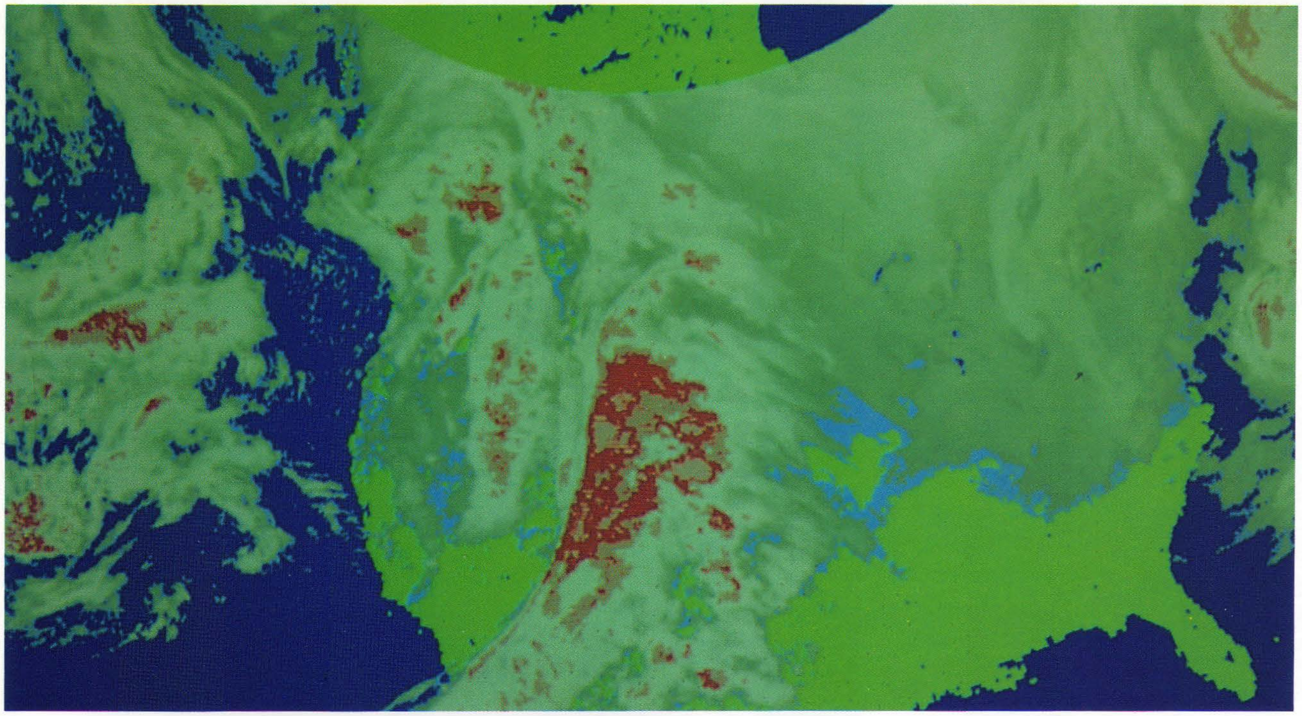

Forecast map of North America produced at the Canadian Meteorological Center on the center's CRAY X-MP computer system.

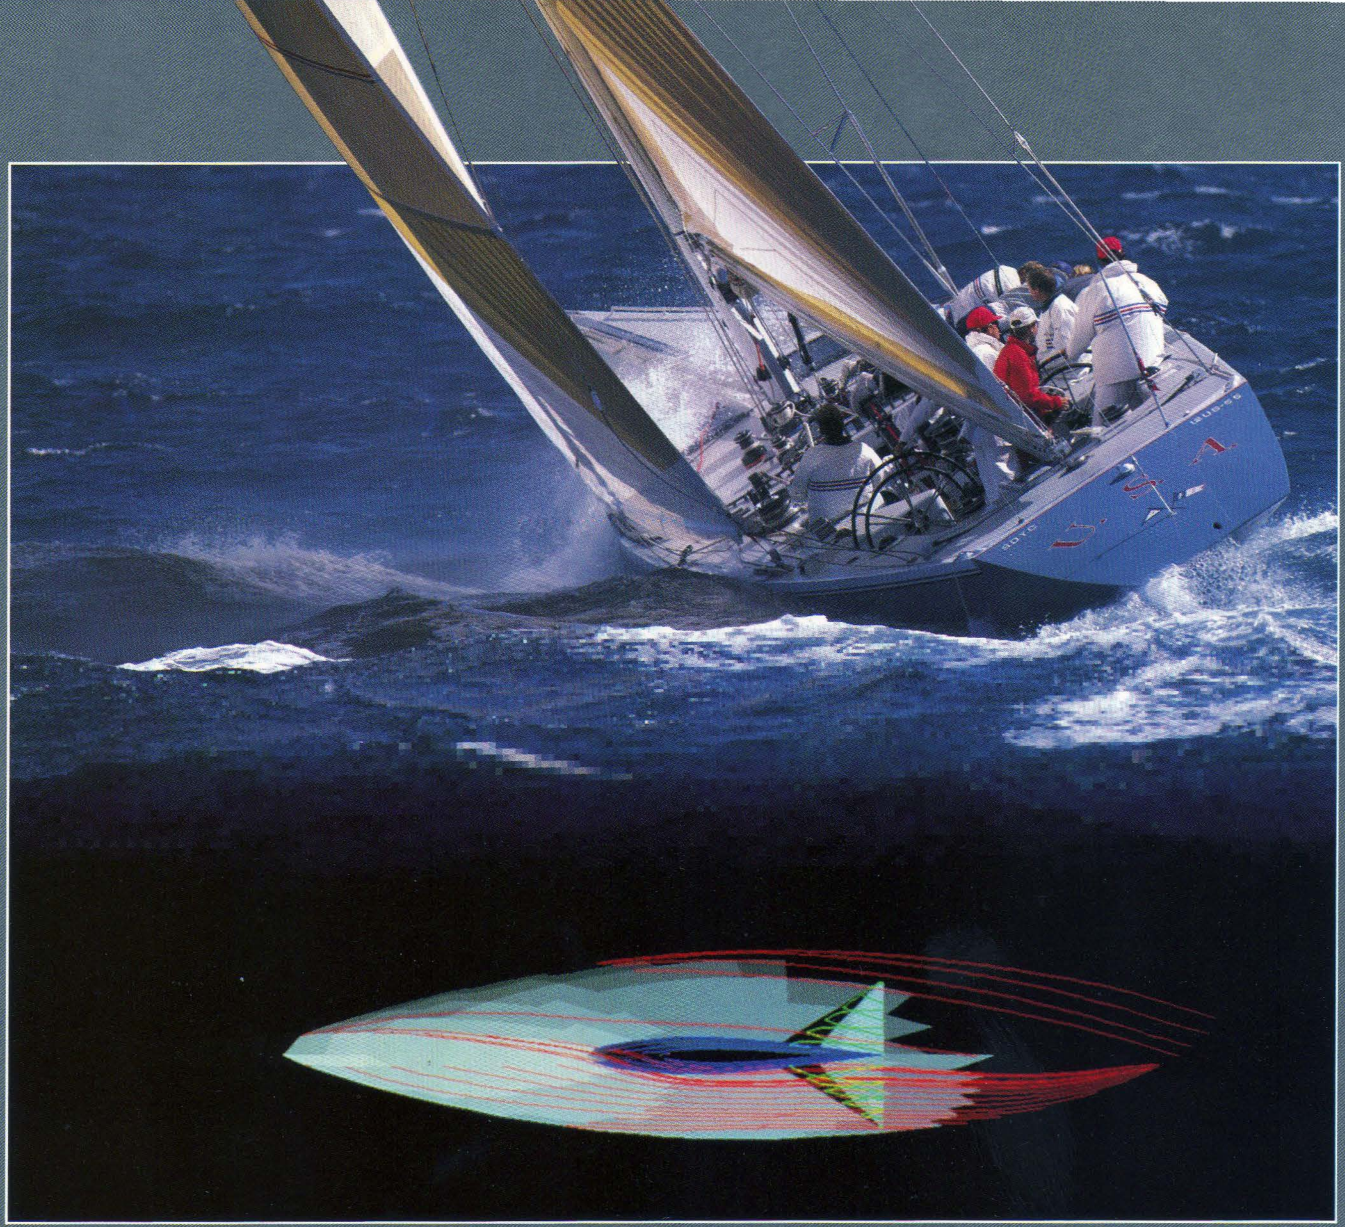

Supercomputer design for Stars & Stripes

Charles W. Boppe and Bruce S. Rosen, Engineering Design & Analysis Consultants, Smithtown, New York

Cray supercomputers and computational fluid dynamics (CFD) combined to optimize the 1987 America's Cup winner *Stars & Stripes*. Clearly, this "World Series" of sailboat racing captured the interest of more people than at any other time in Cup history. Technically oriented enthusiasts are intrigued by this particular sporting event's blend of athletic and engineering skills. For some, the races have evolved into technology demonstrations.

Twelve-meter yacht design for the America's Cup has been pushed to a high level of sophistication. Traditional skills of the naval architect now must be combined with design guidance derived computationally. Observing these three-hour, 24-mile races in which competitors are separated by intervals typically no greater than 100 seconds, one concludes that a one-half percent difference in average sailing speed is important. In some cases, a very small fraction of one percent makes the difference. Towing tank facilities, the traditional approach to design optimization, are incapable of resolving boat speed differences at this level of precision. As a result, naval architects depending solely on sub-scale towing tank tests for final decisions can develop inefficient, uncompetitive full-scale designs.

The challenge

Dennis Conner's Sail America Foundation initiated a design team effort in 1984 with the task of catching up to and surpassing the performance demonstrated by the 1983 America's Cup winner *Australia II*. The radical design of *Australia II* featured wings at the end of the keel. During the first year of the Sail America design program, hydrodynamic computer methods were developed and engineering design was conducted using Grumman Corporation's CRAY-1/M computer system. In the summer of 1985, the Sail America design team began to use Cray Research's CRAY X-MP system in Mendota Heights, Minnesota. The CRAY X-MP system eventually altered the course of the design program. Studies showed that computationally derived drag levels required for performance predictions and design optimization were sensitive to finite element (singularity panel) density. High-density, high-resolution models were prerequisites for predicting accurate performance trends. This factor, coupled with the computer power required for simulating the free air-water (wave) surface and viscous effects, led to relatively complex computational analyses.

The design team met every six weeks to review progress and make final decisions on fabrication or field test items. Often, during the meeting day, "What if..." situations surfaced. The power of the Cray system then was used to evaluate appropriate matrices of variations that provided trending information. This information was available during the evening of the day it was requested. Many novel concepts were evaluated computationally and eliminated, based on identified idiosyncrasies, before model fabrication and testing.

The aircraft design connection

The very low aspect ratio, hull-dominated yacht flow field compromised standard aircraft drag build-up techniques. Thus, it became necessary to push computational prediction accuracy beyond levels used in aircraft design. Automated computer model set-up procedures were essential for achieving this goal. Component definitions of hull, keel, and winglet shapes were processed to yield finite-element models. These were input to a version of the VSAERO CFD code, which was modified for simulating the free surface. Panelization was executed in the same manner for each case of interest. In this way, human modeling errors essentially were eliminated and the probability that a computed increment represented a hydrodynamic effect (as opposed to a numerical error) was enhanced.

Ships generate water waves in the same manner that supersonic aircraft generate shock waves. In both cases, the waves represent energy left behind, which registers as wave drag or resistance on the configuration. Computational simulation of ship waves for Sail America was accomplished by refining and implementing numerical techniques typically used for aircraft flow field predictions. Accuracy and stability levels were enhanced by a tailored five-point upwind operator for setting the free-surface boundary condition. Far-field boundary damping was used, and "rotated" differencing decoupled the numerics from the computational grid. The free-surface technique developed for the Sail America Challenge '87 effort has features that make it the most advanced in the marine industry. Unlike many other approaches, this technique is potentially applicable to the nonlinear free-surface problem.

While wave resistance prediction/minimization was the primary driver for developing this method, it was also important to include free-surface effects for flow angularity predictions, particularly in the plane of the keel's winglet. Winglets misoriented with the local flow present a situation similar to that of dragging an anchor. Properly aligned winglets reduce drag only when the yacht is carrying side-force to work its way to windward (into the wind). Downwind, even the best winglet design must pay a friction drag penalty, due to wetted area, in the same fashion as the keel. If improperly aligned, the winglets' resultant load will produce lift-induced drag with no positive side effect. Calculations eventually revealed that a yacht winglet differs from an aircraft winglet in one important way: speed effects are appreciable. Thus, optimum winglet orientation and spanwise "twist" will be a function of boat speed. This computer code eventually was used to set *Stars & Stripes*' winglet camber, twist, and juncture attachment angle. Experimentally evaluated shape variations could not produce better performance than the shape derived computationally.

A winglet swap

Construction of *Stars & Stripes* was initiated with the keel shape identified as optimum by computational analysis. This keel provided a three-boat-length advantage over an *Australia II*-type keel design. The winglets attached to the keel essentially performed like those of *Australia II*. Empirically designed, these winglets were the best available after 18 months of testing.

By the spring of 1986, modeling accuracy had advanced to the point where subtle drag increments produced by winglet shape changes could be predicted with high levels of confidence using the CRAY X-MP system. Sizing/shape optimization studies were conducted in June of 1986 to determine the best winglet possible for the *Stars & Stripes* hull-keel combination. The computationally designed winglet was smaller (less wetted area) than the basic installed winglet. This improved downwind performance and, at the same time, dramatically improved lifting efficiency by reducing lift-induced drag.

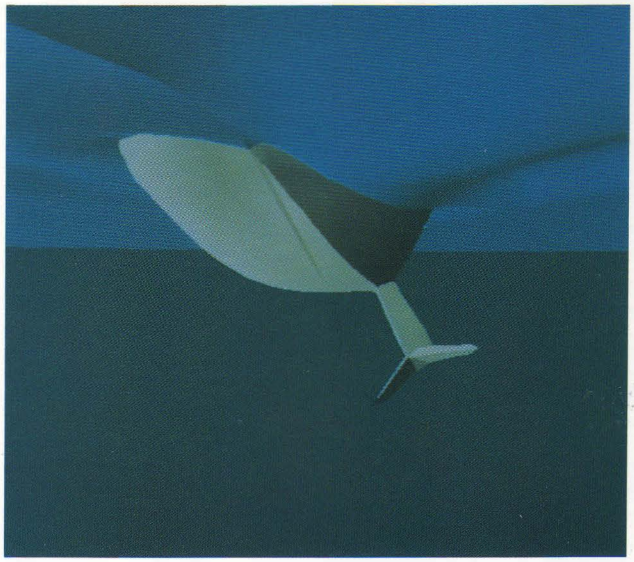

Figure 1. Underwater Australia II *computer model with wave surface studied over one year period (1985) to understand how the winged keel worked (results reported on in Reference 1).*

Reduction in course time with new winglets

Our analyses indicated the new winglets would reduce course time by an average of 40 seconds at high wind speeds and 120 seconds at low wind speeds. Nonetheless, the new winglets were not immediately accepted. First, there was a perception that the Foundation could win without them, and second, implementation resources were not available.

The Challenger Elimination Series began on October 5, 1986. After only a week of racing, it was clear that *Stars & Stripes* was facing several fast boats. On October 15, Richard Walling, a friend of Dennis Conner, called to offer support for testing and construction of the new high-efficiency winglets. Towing tank tests confirmed the computer-derived benefit. Only one hour after the last race of the Challenger Elimination Series, *Stars & Stripes* was hoisted from the water for a winglet swap. Performance now was critical as the Foundation was to face two high-performance boats, *USA* and *Kiwi Magic*, in the Challenger semi-final and final races. Both of these yachts had a 2 win/1 loss record against *Stars & Stripes* in the Elimination Series. In the nine races that followed, *Stars & Stripes* registered only one loss, and that was attributable to an equipment failure.

Riblets

Application of riblets to the underwater surfaces of *Stars & Stripes* received considerable publicity. These microscopic grooves etched in tape are designed to reduce viscous drag by organizing minute eddies very close to the surface. Flow angularity and biofouling, however, can

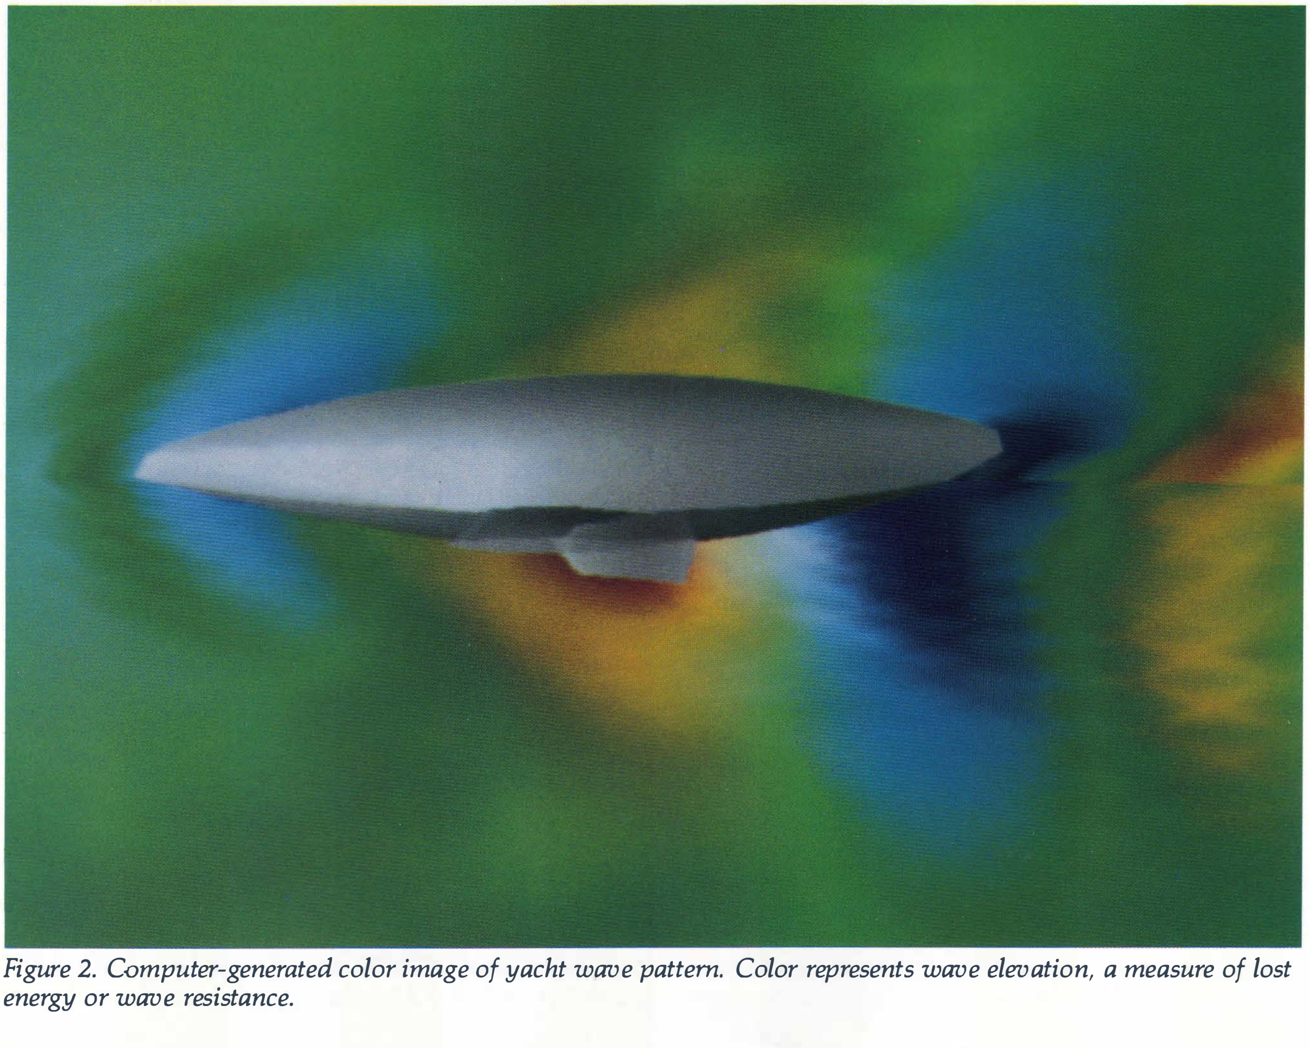

Figure 2. Computer-generated color image of yacht wave pattern. Color represents wave elevation, a measure of lost energy or wave resistance.

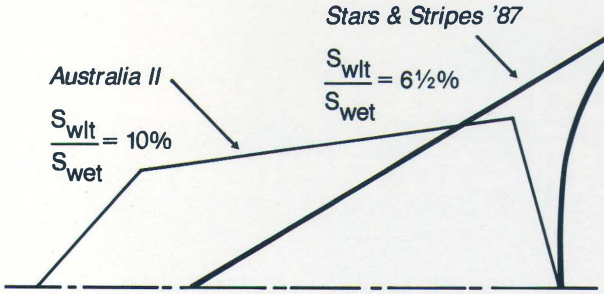

Figure 3. Comparison of winglet planforms for Australia II *and* Stars & Stripes. *The planforms cannot be compared directly because boat sizes are different. In terms of percentages, however, winglets for* Australia II *represent about 10 percent of the total configuration wetted area, while winglets for* Stars & Stripes *represent only six and one-half percent of the total wetted area. The smaller winglet reduces downwind penalty (the winglet serves no purpose when the keel is unloaded) and substantially reduces lift-induced drag levels.*

impair performance. NASA tests have demonstrated that flow/groove angularity greater than 20° can impair drag reduction. Computational flow patterns provided by Cray Research indicated the upright orientation for applying the riblet tape grooves. Analyses also showed that for heeled/yawed tacking conditions, regions of 20-30 degree flow angularity existed. *Stars & Stripes* field testing revealed the total riblet benefit about the course to be quite small, perhaps 10 seconds at best (within field test accuracy limits). This might be attributed to flow angularity problems or to the highly unsteady character of sailing off the coast of Australia.

Enhanced maneuvering performance

Throughout the Challenger Series of races, *Stars & Stripes'* speed loss during maneuvers proved to be a major weakness. When an equipment failure occurred during race three of the Challenger Finals against *Kiwi Magic*, skipper Chris Dickson used this weakness to prevent Dennis Conner from passing with *Stars & Stripes'* superior straight-line speed. It was during this period that Navier-Stokes (ARC 3-D) analysis of hull afterbody shapes began to yield results on the CRAY X-MP system. While arriving late in the program, there was still time for this type of computation to play an important role.

Concern about the maneuvering deficiency intensified. On December 24, 1986, naval architect Bruce Nelson contacted our team hoping to unravel the maneuvering problem. Field test tuft photographs were available, but tuft density was insufficient for diagnostics. Computer-generated flow separation patterns for the afterbody revealed a football-shaped/sized separation bubble just upstream of the rudder. Severe speed loss in the field was identified with rudder deflections greater than 4°. It was postulated that the rudder deflection was aggravating the separation region, causing it to grow significantly in volume. A design modification was recommended. It involved reducing rudder chord length at the top (near the position of the bubble) and adding the area back by extending the rudder tip deeper into clean flow. Thus, smaller deflections would be required to generate a given turning moment. Also, any deflection of the rudder would result in smaller bubble disturbances because reduced areas were involved. A new rudder was constructed along these lines to replace the original during mid-January. Maneuvering performance was enhanced appreciably, as was apparent during the America's Cup series against *Kookabura III*. *Stars & Stripes'* starting line maneuvering characteristics were impressive, and tacking duels to the first wind ward mark revealed that turning deficiencies had been eliminated.

The bottom line

Cray supercomputers played a critical role in the *Stars & Stripes* design campaign. Computer modeling exposed details of flow mechanisms that could not be obtained by any other means. Most important, the computational analyses permitted a detailed understanding of yacht hydrodynamic drag components; these included wave- and lift-induced elements, as well as flow separation. This understanding is not easily generated by subscale or field testing, yet it is required to truly optimize a design. The next "technology demonstration" probably will be held in 1991. The stakes will be higher, and supercomputers no doubt will play an even larger role. □

Acknowledgments

The authors would like to thank Cray Research for providing computing resources and assistance during design of the 12-Meter yacht Stars & Stripes. *In particular, our appreciation is expressed to Kent Misegades for generating yacht flow pattern maps showing flow orientation for riblet taping, and to Carlos Marino for suggesting and implementing the "Cray connection."*

About the authors

Charles Boppe and Bruce Rosen are engineers in the aerodynamics section of Grumman Corporations's Aircraft Systems Division. Their background is in high-speed wing design and the development and application of computational methods. Early in 1985, they formed an independent consulting company to provide computational design services for Sail America Foundation. The effort described herein was conducted over a two-year period.

References

1. Boppe, C. W., B. S. Rosen, J. P. Laiosa, and B. Chance, "*Stars & Stripes*; Computational Flow Simulations for Hydrodynamic Design," *The Eighth Chesapeake Sailing Yacht Symposium Proceedings*, pp. 123-146, March 6-7, 1987.

2. Boppe, C. W., "Detailed Keel-Winglet Design Considerations for Sailing Yachts," Ancient Interface XVII Sailing Symposium, AIAA/SNAME, October 1987.

3. "Water Wings," *Scientific American*, p. 73, April 1987.

4. "Supercomputers, the America's Cup, and winning," *CRAY CHANNELS*, pp. 20-23, Summer 1986.

Comparing tridiagonal solvers on the CRAY X-MP/416 system

Phuong Vu and Chao Yang, Cray Research, Inc.

The solution of tridiagonal systems lies at the heart of many problems in scientific computation, such as the numerical solution of the Helmholtz, Laplace, Poisson, and diffusion equations by finite-difference approximations. This article compares performance of variants of Gaussian elimination, two existing Cray assembly language (CAL) implementations of cyclic reduction, and some of our own implementations of cyclic reduction. Different factors affecting the performance of the cyclic reduction algorithm are tested.

Many methods exist for solving tridiagonal systems such as standard Gaussian elimination, recursive doubling, cyclic reduction, the partition algorithm, and twisted factorization methods. All of the algorithms tested here assume no pivoting is necessary to ensure stability. All tests were performed using the CRAY X-MP/416 system in dedicated mode running the UNICOS 3.0 operating system with the CFT 1.15 Fortran compiler. An additional performance improvement of approximately 5 percent can be obtained using the Cray Research CFT77 1.3 Fortran compiler.

Gaussian elimination

The general tridiagonal system may be written as

$$b_{i-1}x_{i-1} + a_i x_i + c_i x_{i+1} = y_i \quad \text{for } i = 1, 2, \ldots, n$$

with $b_0 = c_n = 0$ where b_i, a_i, *and* c_i represent the elements of the subdiagonal, main diagonal, and superdiagonal respectively, while x_i and y_i represent the elements of the unknown vector and the right-hand side respectively. The standard Gaussian elimination algorithm may be stated as follows:

$$d_i = (a_i - l_{i-1}c_{i-1})^{-1} \quad \text{for } i = 1, 2, \ldots, n$$

and

$$l_i = b_i d_i \quad \text{for } i = 1, 2, \ldots, n-1$$

followed by

$$w_i = y_i - l_{i-1}w_{i-1} \quad \text{for } i = 1, 2, \ldots, n$$

and

$$x_i = d_i(w_i - c_i x_{i+1}) \quad \text{for } i = n, n-1, \ldots, 1$$

where $l_0 = 0$.

This algorithm is recursive and cannot be vectorized. Three different implementations of Gaussian elimination will be considered: TRIDA, TRID1, and TRID2, in which TRIDA is a straightforward scalar Fortran code, TRID1 is a burn-at-both-ends (BABE) Fortran code, and TRID2 is a CAL implementation of the BABE algorithm. The BABE algorithm simultaneously starts the elimination of elements off the main diagonal at both the top of the subdiagonal and at the bottom of the superdiagonal. It proceeds until the elimination meets in the middle. The back substitution then simultaneously starts from the middle, and works outward to the top and bottom of the matrix. The Fortran code for TRIDA is given by

```
      A(1) = 1.0/A(1)
      DO 10 I = 2, N
      TEMP = A (I-1) * B(I)
      A(I) = 1.0/(A(I) - TEMP * C(I - 1))
      X(I) = X(I)-TEMP * X(I-1)
10    CONTINUE
      X(N) = X(N) * A(N)
      DO 20 I = N-1, 1, -1
      X(I) = (X(I) - X(I+1) * C(I)) * A(I)
20    CONTINUE
```

A simple operation count provides a cost of $9n-7$ floating point operations to solve a single tridiagonal system using Gaussian elimination (as implemented in TRIDA). In comparing the performance of the different solvers, TRIDA is used as the reference case, and the speedup of the other solvers is computed with respect to this code. The average observed rate of TRIDA in runs with sizes ranging from $n = 32$ to $n = 4096$ is about 7.4 MFLOPS.

Cyclic reduction

Cyclic reduction is a highly efficient method for vectorizing the solution of tridiagonal systems. This is accomplished by eliminating all the odd variables first, then all the odd multiple-of-2 variables. The process then is repeated recursively for *ncycle* levels of reduction, or until the

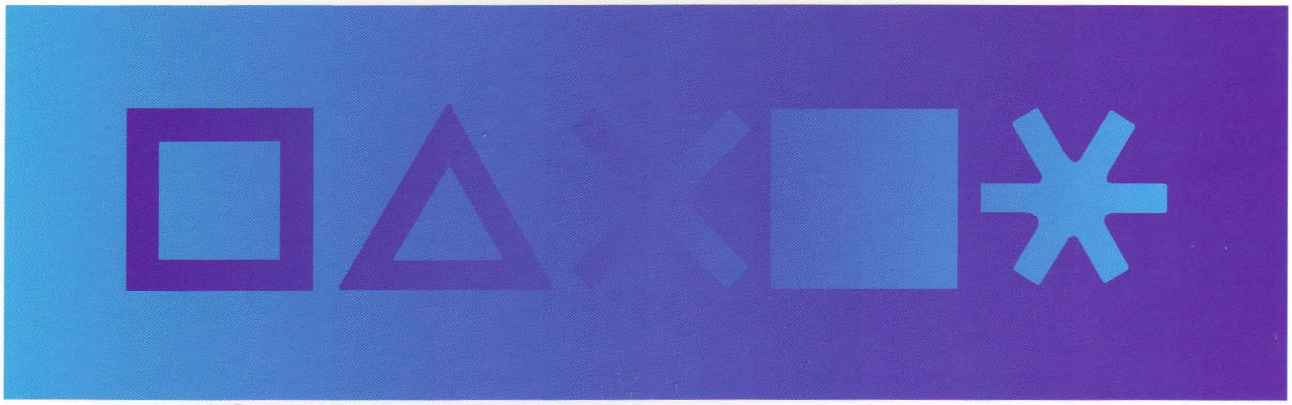

reduced tridiagonal system at the current level becomes smaller than a specified cutoff size, *ncutoff*. The final reduced system is then solved using a scalar algorithm. The remaining unknowns subsequently can be found from a filling-in procedure. For example, if $n = 3584$, ncycle = 10, and ncutoff = 32 are given, then the algorithm will perform seven levels of reduction, giving a final reduced system of size 28. It is easy to show that cyclic reduction is just Gaussian elimination applied to a permuted system for a particular permutation matrix P.

As a result, the cyclic reduction algorithm is numerically stable for matrices for which Gaussian elimination is stable without pivoting, such as symmetric positive-definite or diagonally dominant matrices.[1,2] The result is a vectorizable algorithm with vector lengths $n/2$, $n/4$, $n/8$,.... However, this gain in vectorization is not free. The number of operations in the vectorizable algorithm is greater than the number of operations in the scalar algorithm. Also, a possible memory bank conflict exists as memory is accessed in increments of 2^i at the $i\text{-}th$ level of reduction and the $i\text{-}th$ level of solution fill. This may be the reason that most available cyclic reduction codes for the CRAY X-MP system have been implemented with ncycle equal to 1.

The Cray system memory is organized in a set of banks. Each bank has a cycle time of four clock periods. As long as four or more clock periods occur between accesses to the same bank, memory operations go at full speed — one access per clock period. For the first cycle of cyclic reduction, the memory stride is 2, for the second it is 4, for the third it is 8, and so on. When the memory stride is equal to half the number of banks, memory references go at half-speed. When the memory stride is a power of 2, and equal to or greater than the number of banks, memory references go at one-fourth speed. For a CRAY X-MP/416 system with 64 banks, we expect to see a degradation when ncycles = 5, or a stride of 32. However, as shown from testing, this stride problem can be alleviated by using extra work space for the three diagonals, and possibly the right-hand side. This extra work space permits elements generated by the cyclic reduction steps to be stored in successive locations so that subsequent reads are at a memory stride of 2, at most.

To illustrate the effect of different factors on the performance of the cyclic reduction algorithm, 11 different implementations are considered here, including two existing CAL implementations. All of our codes implement the general k-level cyclic reduction algorithm, allow a user to specify the cutoff point for the final reduced system, and give nonunit increments between elements of the diagonals and the right-hand side. Therefore, these codes are more flexible than most currently available codes. The two codes considered are David Kershaw's DECOMP/SOLVE and Dennis Kuba's TRIDSOL. Both implement a one-level cyclic reduction algorithm entirely in CAL.

Codes were written in recursive and nonrecursive Fortran. Some use CAL versions for the inner kernels, and some use extra memory for either the diagonals or for the diagonals and the right-hand side to alleviate the stride problem. Table 1 summarizes the different versions of the cyclic reduction algorithm tested.

Code name	Non-recursive Fortran	Recursive Fortran	With CAL kernels	With extra memory
CYCRED1		X		
CYCRED2		X	X	
CYCRED3	X			
CYCRED4	X		X	
CYCRED5	X			X
CYCRED6	X		X	X
CYCRED7*	X			X (+RHS)
CYCRED8		X		X
CYCRED9		X	X	X
DECOMP/SOLVE*			X	X (+RHS)
TRIDSOL			X	

*For CYCRED7 and DECOMP, extra memory is used for both the diagonals and the right-hand-side vector.

Table 1. A summary of the tested versions of cyclic reduction.

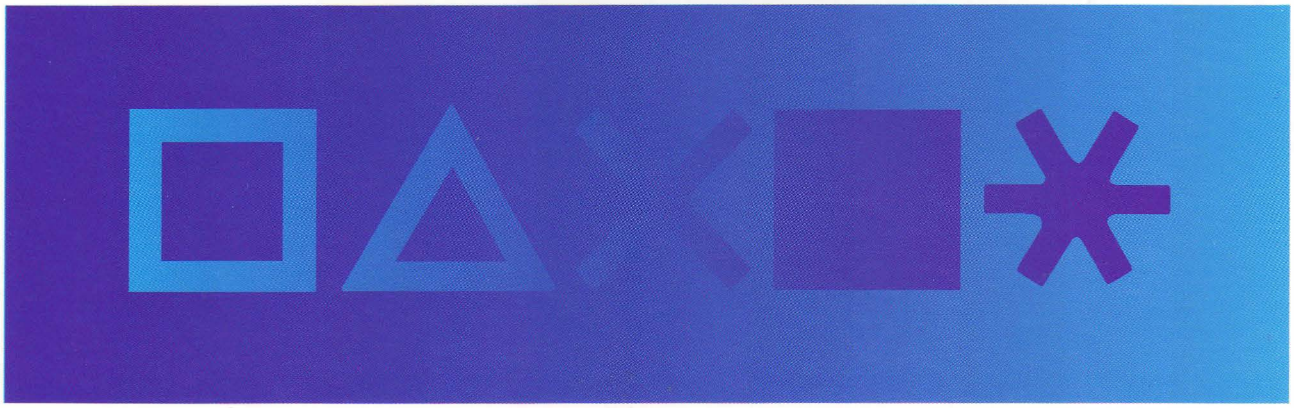

All tests are run with two different ncutoff values that are determined experimentally and correspond to the crossover point in the performance between a scalar Gaussian elimination code and a cyclic reduction code using ncycle = 1. The ncutoff value equal to 32 (respectively 128) comes from comparing TRID1 versus CYCRED3 (respectively TRID2 versus CYCRED4). Other values also were tried, but the results obtained did not differ significantly from those using either ncutoff equal to 32 or 128. This indicates that ncutoff is not a significant parameter. Consequently, we restrict the discussion of our results to the case of ncutoff = 32. To test the effect of the number of reduction levels, all our cyclic reduction codes are run with ncycle = 1, ..., mxcycle, where *mxcycle* is the maximum number of reduction levels corresponding to the size of the problem, *n*, and a given ncutoff.

From the test results of CYCRED1, ..., CYCRED9 cyclic reduction codes, CYCRED6 consistently performs best. Hence, we have chosen to use CYCRED6 in comparison against other solvers. Figure 1 displays speedup results of TRID1, TRID2, TRIDSOL, DECOMP/SOLVE, and CYCRED6 versus TRIDA. Figure 2 compares the performance of CYCRED6 and its Fortran counterpart CYCRED5. For each problem size, the results for CYCRED5 and CYCRED6 are chosen from among the best speedups for different ncycles. CYCRED6 is asymptotically about 33 percent faster than either of the two existing CAL implementations, and about eight times faster than the straightforward Fortran scalar code TRIDA. Surprisingly, with a good choice of ncycle, the Fortran version of CYCRED5 is about 12 percent faster than either of the two existing CAL implementations when the problem size increases. In

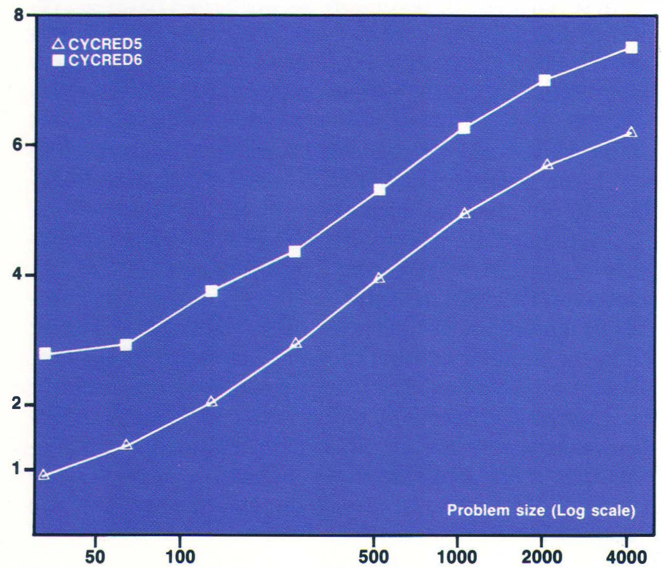

Figure 2. CYCRED5 versus CYCRED6, ncutoff = 32.

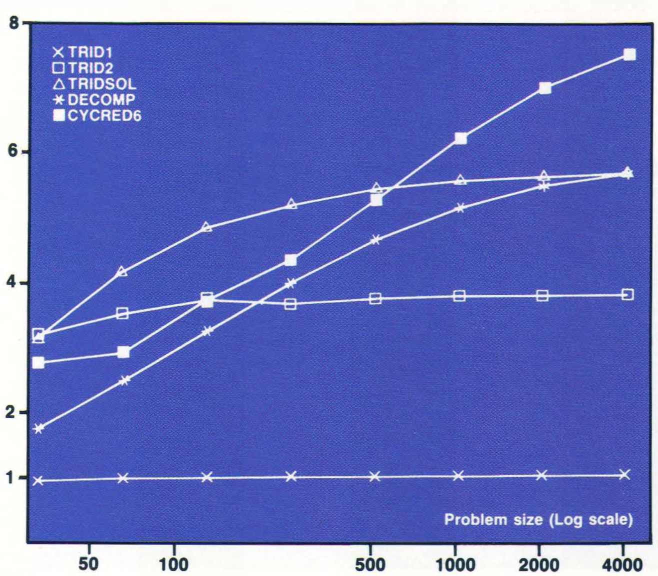

Figure 1. Speedup relative to TRIDA, ncutoff = 32.

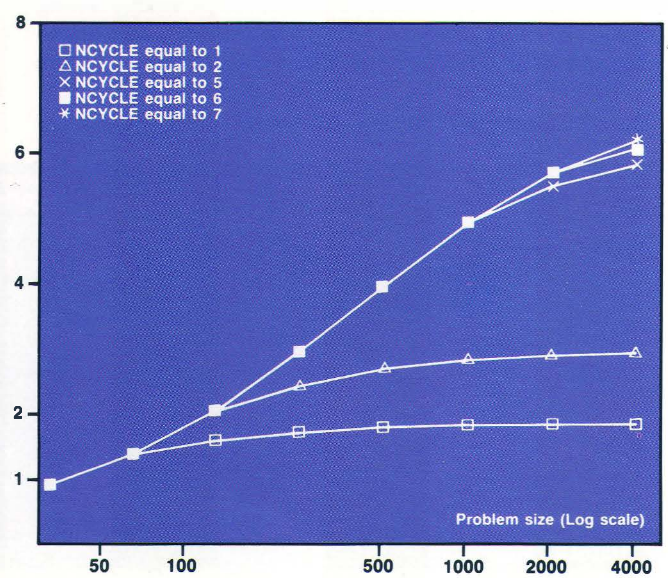

Figure 3. Variation of CYCRED5 due to ncycle, ncutoff = 32.

Figures 3 and 4, the effect of ncycle on the performance of the cyclic reduction algorithm is illustrated by showing the speedup results of CYCRED6 and its Fortran version CYCRED5 for several values of ncycle. Finally, Figure 5 illustrates the advantages of the extra memory used to alleviate the stride problem. The speedup results of CYCRED6 are compared with the equivalent nonextra memory version CYCRED4.

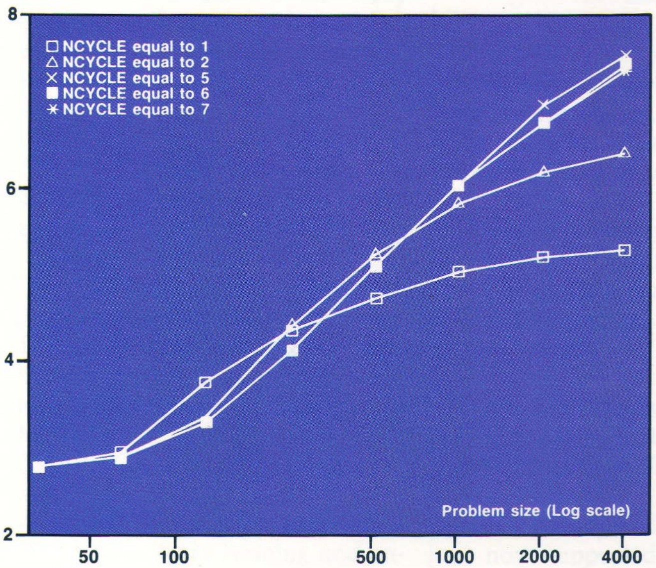

Figure 4. Variation of CYCRED6 due to ncycle, ncutoff = 32.

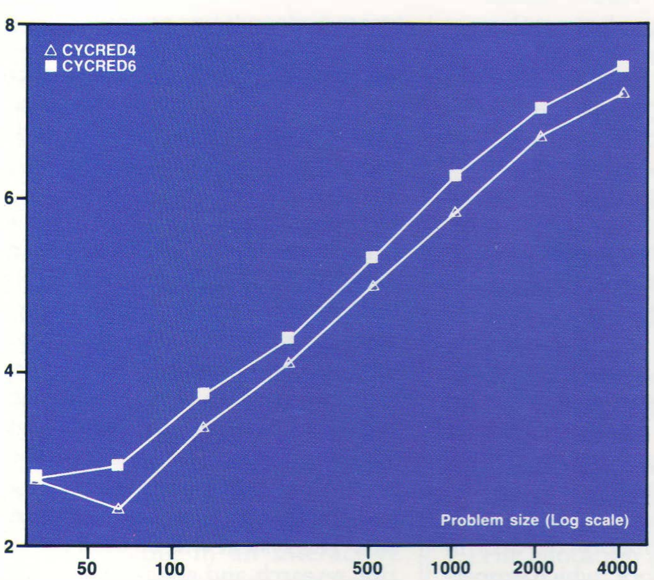

Figure 5. CYCRED4 versus CYCRED6, ncutoff = 32.

Evaluating performance

As demonstrated, the best asymptotic performance of the cyclic reduction algorithm for solving tridiagonal linear systems comes from an implementation that uses k levels of reduction until the size of the reduced system is smaller than a given cutoff size. The best implementation with CAL versions of the inner kernels is asymptotically about eight times faster than a straightforward Fortran implementation of the scalar algorithm. Moreover, a k-cyclic reduction algorithm implemented entirely in Fortran can outperform a good CAL implementation using a single level of reduction for k as small as 2. □

Acknowledgments

The authors thank Tom Hewitt for providing a copy of his periodic tridiagonal solver, which was extremely helpful in the early stage of this project. Those interested in additional information may contact Phuong Vu or Chao Yang at Cray Research, Inc., 900 Lowater Road, Chippewa Falls, WI, 54729; telephone: (715) 726-1211.

About the authors

Phuong Vu is a senior applications engineer with Cray Research's High Speed Computing Group located in Chippewa Falls, Wisconsin. His primary area of expertise is numerical software. Vu joined Cray Research in 1986 from J.S. Nolen & Associates. He received his Ph.D. degree in mathematics from the University of Houston in 1983.

Chao Yang is an applications engineer with Cray Research's High Speed Computing Group in Chippewa Falls, Wisconsin. Yang joined Cray Research in 1985 from Digicon Geophysical Corporation. He worked in the SCILIB group in his first two years at Cray Research. Yang received his M.S. degree in mathematics from Ohio State University in 1980, and his M.S. degree in computer and information science from Ohio State University in 1981.

References

1. Jordan, T. L., "A Guide to Parallel Computation and Some CRAY-1 System Experiences," *Parallel Computations*, P. G. Rodrigue, ed. 1982.

2. Kershaw, D. S., "Solution of Single Tridiagonal Linear Systems and Vectorization of the ICCG Algorithm on the CRAY-1 System," *Parallel Computations*, P. G. Rodrigue, ed. 1982.

CORPORATE REGISTER

Cray adds new customers

In October Cray Research announced that Argonne National Laboratory had ordered a CRAY X-MP/14 computer system. The system was installed in the fourth quarter of 1987 at Argonne's research laboratory in Argonne, Illinois. The laboratory is run by the U.S. Department of Energy. "The new computer will provide the foundation for a long-term computing capability that can be upgraded to meet Argonne's future scientific computing demands," said David Weber, the laboratory's manager of computing services. Argonne National Laboratory is a new customer for Cray Research.

Cray Research also announced in October that the U.S. Army had ordered a CRAY-2 computer system. The system will be installed in the first quarter of 1988 at the U.S. Army's tank command base (TACOM) in Warren, Michigan. The system will be used to plan, design, and develop the next generation of U.S. Army tactical vehicles.

General Electric (GE), acting as a prime contractor for the Department of Energy, has ordered a CRAY X-MP/48 computer system with SSD solid-state storage device. The system was installed in the fourth quarter of 1987 at the Knolls Atomic Power Laboratory (KAPL) in Schenectady, New York, which is operated by GE for the Department of Energy. DOE/KAPL is a new customer for Cray Research, which announced the order in October.

Cray Research also announced in October that the Japanese Ministry of International Trade and Industry (MITI) had ordered a CRAY X-MP/216 computer system with SSD solid-state storage device. The system will be installed in the first quarter of 1988 at the Agency of Industrial Science and Technology in Tsukuba, Japan. The system will be the ninth Cray system installation in Japan and the first at a Japanese government site. The system will be used for scientific research and education.

In November Cray Research announced that Texaco, Inc., had installed a CRAY X-MP/28 computer system at Texaco's geophysical center in Bellaire, Texas. The system will be used for general petroleum research, including seismic analysis and reservior modeling. Texaco is a new customer for Cray Research.

UNICOS release 3.0 provides enhancements

Release 3.0 of the Cray operating system UNICOS provides users with an enhanced interactive operating system designed to run on all Cray computer systems. This release reflects Cray Research's commitment to provide its customers with software technology that takes full advantage of Cray hardware technology.

UNICOS 3.0 provides the following enhancements for all Cray computer systems:

- ☐ The UNICOS kernel and Network Queuing System (NQS) support checkpointing and recovery of batch requests during a scheduled shutdown, including enhancements to the Network Configuration Tool, QMAPMGR. Also, NQS initiation is enhanced so that all queues appear with a STOPPED status, inhibiting any requests from running and enabling the operator to inspect the queues.
- ☐ Job accounting and the jobs concept feature are improved. The jobs concept feature allows accounting and resource limits by groups of processes rather than by single processes. The job accounting feature improves information available for user jobs, including interactive sessions and NQS jobs.
- ☐ Many commands, system calls, library routines, and products are new or upgraded for this release, including the addition of many UNIX System V.3 enhancements.
- ☐ The new CRAYPERF feature provides graphical displays for system performance monitoring.
- ☐ The X Window System now allows users working in an interactive UNICOS environment to write graphics output directly on a bit-map display terminal.
- ☐ The system activity report utility is now available on all Cray systems for monitoring system performance.
- ☐ Memory management (swapping and scheduling) is improved.
- ☐ The new library building tool is implemented.
- ☐ The multigroups feature for verifying group identification is new with this release.
- ☐ UNICOS signal capabilities are upgraded to AT&T UNIX System V.3 level.
- ☐ On-line diagnostics include more testing features and error reporting.

In addition to the enhancements to UNICOS that apply to all Cray computer systems, UNICOS 3.0 also includes the following enhancements that apply only to CRAY-2 computer systems:

- ☐ The new File System Switch feature allows multiple file system implementations simultaneously.
- ☐ A new method is provided for faster deadstarts.

Release 3.0 includes the following enhancements that apply only to CRAY X-MP and CRAY-1 computer systems:

- ☐ A system-level tape subsystem is now supported for Cray computer systems configured with an I/O Subsystem (IOS). This provides multivolume and multifile capabilities for C programs and standard UNIX utilities, a set of user and administrator commands, and an operator interface.
- ☐ The HSX high-speed external communications channel is supported.
- ☐ The boot-time configuration feature allows disks, SSD solid-state storage device, and Buffer Memory Resident to be configured at boot time using an IOS parameter file, thus simplifying system configuration.
- ☐ File backup provides file dumping and reloading.
- ☐ A real-time scheduling queue, with special handling capability, is implemented for real-time systems.
- ☐ File allocation is improved
- ☐ SSD support is improved.
- ☐ The Cray Simulator is improved over the previous version, including features for tape simulation and symbolic tracing.

Enhancements to the transmission control protocol/internet protocol (TCP/IP) include improvements in both performance and software maintenance. Station support is provided through the UNICOS Station Call Processor (USCP), which includes improvements in performance and reduction in the amount of memory used.

UNICOS 3.0 supports many languages, libraries, and utilities. To help current COS operating system customers migrate to UNICOS, Cray Research continues to offer the necessary environment and tools. Documentation based on AT&T UNIX System V documentation also is provided with the UNICOS operating system. A primer on the UNICOS operating system and new publications are available to system users.

Seymour Cray extends agreement, defines CRAY-4 system

Cray Research announced in November that Seymour Cray has extended his design and development agreement with the company. The agreement now extends to December 31, 1992. The extension relates to Cray's decision to begin work on the design of a new supercomputer, the CRAY-4 system.

Seymour Cray said he anticipates that the new system will have 64 processors and will run with a clock speed of one nanosecond (one billionth of a second). The system's performance objective is to provide throughput more than 1000 times that of the CRAY-1 computer system, which was introduced in 1976. The CRAY-4 system will be based on gallium arsenide circuitry.

Under the design and development agreement, Seymour Cray acts as an independent contractor to Cray Research, furnishing development work for the company's advanced computer systems. The agreement may be extended further by additional projects.

APPLICATIONS IN DEPTH

New team supports environmental sciences

Weather forecasters take their lumps for drenched picnics and other minor disasters, though their predictions save the world's economy billions of dollars each year. From the role it plays in agricultural planning and aircraft routing to emergency storm warnings, forecasting proves its worth in dollars and lives saved. Perhaps second only to advances in numerical weather models, supercomputers have played a decisive role in advancing environmental research and weather prediction.

Atmospheric scientists were among the first supercomputer users. Over the years their interests have expanded to include not only local forecasting, but also basic studies of atmospheric phenomena, more comprehensive (global) forecasting, and environmental issues. To ensure that Cray Research understands the needs of this research community, the company established an environmental sciences support team.

"Atmospheric and meteorological researchers are sophisticated users of Cray computer systems because of their long experience with them," says David Blaskovich, Cray's director of environmental systems marketing. "For this reason, providing them with sheer processing power isn't enough anymore. They want a reliable hardware and software supercomputer solution."

The environmental sciences team consists of Blaskovich; Eric Pitcher, a senior analyst in Cray Research's applications department who has extensive circulation modeling experience; and Bob Wellck, a senior sales analyst. Collectively, they represent nearly 60 years experience with research laboratories and operational centers.

"We are working to show how Cray systems meet environmental modeling requirements," says Blaskovich. "We're also trying to influence product development within Cray so that new products and enhancements will complement the environmental industry's needs."

According to Pitcher, the environmental research and weather forecasting industries make up a unique market for Cray Research. "Unlike some other users of our systems, environmental scientists don't have off-the-shelf application programs available to them. They generally have to use programs developed in-house, and these programs undergo continual development. We don't necessarily want to change that situation, but we want to understand these researchers' special needs."

"The market is expanding in various ways," Blaskovich adds. "The international community is showing a growing interest in global models, as well as in smaller, regional models. Also, most operational weather forecasting today occurs in centers in Europe and North America. We expect an interest to emerge in other parts of the world as well. This industry often is one of the first to develop for us in new geographical markets."

The environmental sciences research community is growing in other important ways as well. One emerging trend is the use of supercomputers for interdisciplinary research. Researchers want to develop models that link Earth's atmosphere and oceans into an integrated system. Coupling oceanographic and atmospheric models is a natural development because roughly 40 percent of the solar heat Earth receives travels from the equatorial zones to the poles via the oceans. Such large coupled models eventually would be expanded to include Earth's biosphere as well. Much larger computers than are now available will be needed to run these models.

Even within the field of oceanography, many unique applications of large-scale modeling are being explored. The U.S. Navy, for example, is considering development of a worldwide ocean forecasting model that, among other things, could be used to predict ocean circulation and chemical and biological fields. Other emerging oceanographic applications include thermal studies of the oceans to predict the effect of water temperature on sonar operation.

The environmental sciences market for Cray Research also is changing due to changing operational mandates of traditional industry customers. "Their mandates no longer cover just operational meteorology," says Wellck, "but include basic and applied research as well." Researchers increasingly want to use computer modeling to study air and water pollution, global heating (the greenhouse effect), atmospheric anomalies (such as the ozone depletion over Antarctica), and the prospect of "nuclear winter."

"Atmospheric models were among the first codes to be multitasked on Cray systems, and they lend themselves to multitasking fairly readily," says Blaskovich. "Some experiments show that 12 to 16 processors could be used effectively, so this industry should be able to move smoothly into the next generation of Cray products and take full advantage of the new systems' features. We see growing needs for accurate and timely weather forecasting and for advanced research in climate dynamics, oceanography, and pollution monitoring. We plan to make sure that Cray systems provide the best solution for these computational problems."

Analysis package available on Cray systems

The Dynamic Analysis and Design System (DADS) is a package of general-purpose computer programs for kinematic and dynamic analysis. DADS now is available to run on Cray computer systems under Cray Research's UNICOS operating system. The DADS package is a product of Computer Aided Design Software, Inc. (CADSI).

The DADS software package can be used to model and predict the motion of a variety of real-world mechanisms. The package includes the main analysis program, a data entry preprocessor, and a postprocessor for evaluating simulation results. DADS offers several analysis options:

- ☐ Kinematic analysis to calculate positions of physical bodies
- ☐ Dynamic analysis to compute positions, velocities, accelerations, and reaction forces of bodies at discrete time intervals
- ☐ Static analysis to calculate a model configuration for static equilibrium
- ☐ Inverse dynamic analysis to calculate forces necessary to generate a prescribed motion
- ☐ Assembly analysis to assemble all bodies into a configuration that satisfies all joint and constraint connections between bodies

DADS includes control systems to construct transfer functions that represent feedback, control, and hydraulic systems. The transfer functions are automatically added to the set of differential equations in DADS, and are solved using numerical integration. The control systems can be coupled directly to the dynamic analysis of rigid and flexible bodies throughout a simulation.

DADS uses modal synthesis techniques to account for the effects of flexible bodies. This approach allows users to combine finite element analysis with dynamic analysis, and eliminates the approximations of rigid body analysis.

For more information on using DADS with Cray computer systems, contact CADSI, P.O. Box 203, Oakdale, IA 52319; telephone: (319) 337-8968.

Informix offers database management tools

The Informix line of database management tools, from Informix Software, Inc., now is available to run on Cray computer systems under the UNICOS operating system. The Informix family of products, from Informix Software, Inc., includes several database management tools.

INFORMIX-SQL is a relational database management system based on IBM's Structured Query Language (SQL). It provides the tools for building database applications, including a menu creation facility, an interactive schema editor, an ANSI-standard SQL-based query and data definition language, a forms generator, a report writer, and a set of utilities for database administration.

INFORMIX-4GL is a fourth-generation database programming language. It is a C- or COBOL-caliber procedural language, not an applications generator. INFORMIX-4GL allows users to create complex customized applications that are easy to modify and maintain. The

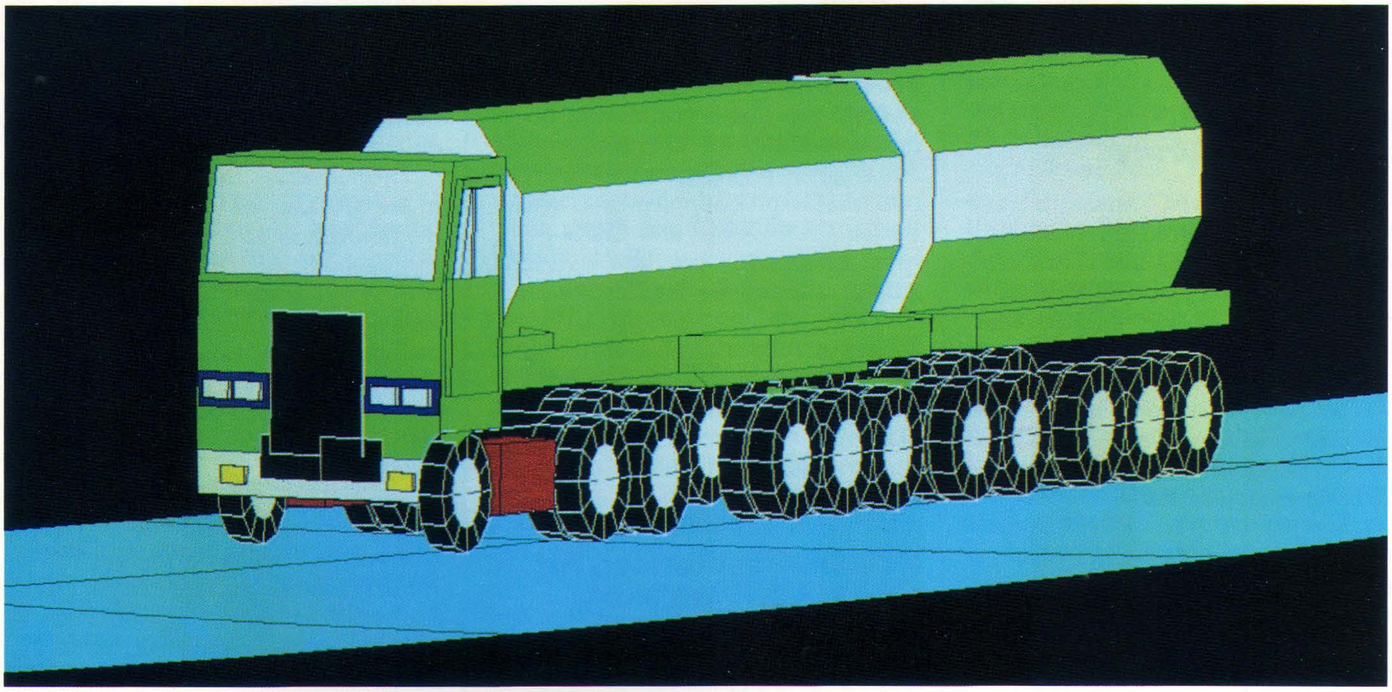

Analysis model created with the DADS software package.

INFORMIX-4GL language consolidates SQL syntax for data definition and retrieval and its own proprietary syntax for creating menus, screens, and reports.

INFORMIX ESQL/C allows programmers to embed industry-standard SQL syntax directly into custom database applications built with C. This greatly reduces programming time by reducing the amount of code necessary for file or database access or manipulation.

C-ISAM is a library of C functions for creating and manipulating indexed files. It is the standard access method for UNIX systems in the United States and Europe, and is the internal access and retrieval method for Informix products. Programmers can use C-ISAM in conjunction with the C programming language to build completely customized applications that do not require the screen, report, or menu capabilities offered by database systems, application generators, and fourth-generation languages.

For more information on using Informix database management tools with Cray systems, contact Neil Blumenfield, Informix Software, Inc., 4100 Bohannon Drive, Menlo Park, CA 94025; telephone: (415) 322-4100; FAX: (415) 322-4571.

NAG offers numerical library

Mark 12 of the NAG Fortran Library now is available from the Numerical Algorithms Group, Inc. The NAG Library is a library of numerical algorithms designed to meet the mathematical and statistical needs of computer users. Implementations of the library are available for Cray computer systems running under Cray Research's COS or UNICOS operating systems or under the CTSS operating system.

Mark 12 of the NAG Fortran Library includes 688 user-level routines and an additional 732 auxiliary routines used by the user-level routines — a total of approximately 200,000 lines of Fortran source text. Mark 12 of the NAG Fortran Library includes 97 new linear algebra support routines for commonly occurring scalar, vector, and matrix-vector operations of linear algebra.

Documentation includes example programs that illustrate the use of the routines. Data and results are provided, together with Fortran source code for the example programs, printed in manuals and in machine-readable form for each implementation. Users frequently modify the example programs to solve specific problems.

Application programs developed on a system using the NAG Fortran Library can be ported with minimal effort to other systems using the NAG Fortran Library. Each implementation of the library is specific to an environment (computer, operating system, compiler, and precision). The library comprises many "chapters" that cover a wide range of mathematical applications. General areas covered include

- ☐ Ordinary differential equations
- ☐ Partial differential equations
- ☐ Curve and surface fitting
- ☐ Minimizing and maximizing functions
- ☐ Matrix operations, including inversion
- ☐ Eigenvalues and eigenvectors

The optional Online Information Supplement provides access to user documentation and an interactive NAG HELP system for selecting routines. The optional Graphic Supplement is available for CRAY X-MP and CRAY-1 systems. This option provides high-level graphic routines for data display.

For more information on using the NAG Fortran Library with Cray computer systems, contact Cynthia L. Mohns, Numerical Algorithms Group, 1101 31st Street, Suite 100, Downers Grove, IL 60515; telephone: (312) 971-2337.

PCGPAK vectorized for CRAY X-MP systems

The Preconditioned Conjugate Gradient Package (PCGPAK) from Scientific Computing Associates, Inc., is a portable library of Fortran 77 subroutines for efficiently solving large sparse linear systems. Portions of PCGPAK have been vectorized to provide optimum performance on CRAY X-MP computer systems running the UNICOS operating system.

To demonstrate the performance improvements resulting from vectorization, Scientific Computing Associates conducted a benchmark test to compare the performance of PCGPAK on scientific computers. The 8000-equation model problem was a seven-point central difference discretization on a 20 by 20 by 20 grid of the following partial differential equation on the unit cube $[-((e^{xy}u_x)_x + (e^{xy}u_y)_y + (e^{xy}u_z)_z) + 80(x + y + z)u_x + (40 + (1 + x + y + z)^{-1})u = f]$ with Dirichlet boundary conditions and with f chosen so that the exact solution is $u = (1-x)\ (1-y)\ (1-z)\ (1-e^{-x})\ (1-e^{-y})\ (1-e^{-z})$. The convergence criterion was a reduction in the two-norm of the residual by a factor of 10^{-15}, starting from a zero initial guess.

Solving the model problem with the banded Gaussian elimination method, without pivoting, would require about 50 Mbytes of memory and the execution of about 2.5 billion floating point operations Using PCGPAK, solving the problem required only 1 Mbyte of memory and the execution of 14 million floating point operations. Computation time on the CRAY X-MP system with the nonvectorized version of PCGPAK required 3.54 seconds to solve the model problem. Running the vectorized version of the package, the Cray system solved the problem in 0.73 seconds.

For large problems for which standard dense or banded Gaussian elimination methods require excessive computing resources and time, PCGPAK's sparse data structures and iterative techniques can make solving such problems practical.

For more information on the PCGPAK software package vectorized for CRAY X-MP computer systems, contact David Foulser, Scientific Computing Associates, Inc., 246 Church Street, Suite 307, New Haven, CT 06510; telephone: (203) 777-7442.

USER NEWS

'Renaissance team' sees higher dimensions with Cray system

Take the visual and creative abilities of an artist, add the scientific background of a mathematician and the practical skills of a computer programmer. The sum is computer sculpture that animates the nonvisual.

At the National Center for Supercomputing Applications in Champaign, Illinois, Donna Cox, George Francis, and Ray Idaszak have linked their talents and the technology of a CRAY X-MP/48 system to create a kind of interactive sculpture machine. Using a Cray system and an IRIS Superworkstation, they can visualize mathematical and scientific data as shadows projected from higher-dimensional forms.

Besides producing a new type of art, this collaboration has led to the discovery of a single aesthetic form that represents a mathematical theorem describing the transformation of one multidimensional geometric figure into another. These curvilinear figures differ from more hard-edged computer graphics, resulting in sensuous and organic shapes, explains Cox, an assistant professor of art and design at the University of Illinois.

The project, which was supported by a grant from the National Science Foundation, grew from the varied interests and backgrounds of each group member. Cox refers to the group as a "renaissance team," with each specialist providing a broad spectrum of skills in the quest for discovery. Looking for project input, Cox and Idaszak, a computer programmer, approached Francis, a mathematician at the University of Illinois. Francis teaches topology, which can be described as "rubber sheet geometry" because it is a form of mathematics that describes geometric shapes that can be transformed into other shapes without being cut or pasted. This process of transformation is called *homotopy*.

Francis suggested a project to recreate visually the Romboy homotopy, which is the transformation of the Roman surface, a shape that looks somewhat like a three-sided inflated raft, into the Boy surface, a twisted, pretzel-like surface. The equations had been worked out previously by French mathematician Francois Apery. In 1984, Apery discovered that by changing parameters gradually, he could transform the Roman surface smoothly into the Boy surface. Apery's work revealed that both surfaces can be regarded as three-dimensional shadows from higher-dimensional forms viewed from various vantage points.

Before Francis was approached by Cox and Idaszak, he was working on a similar project on a smaller scale, using Apery's findings and a personal computer to create a crude rendering of the transformation. When Francis joined the team, he, Cox, and Idaszak worked together to convert his Apple program into supercomputer code. With the power of the CRAY X-MP system, they were able to enhance the program by adding shading, color, and a skinlike surface to the geometric shapes. "I continually sculpted and explored those topological surfaces with the interactive programs Ray updated according to specific needs," says Cox.

Each frame of the resulting 600-frame animation follows the deformation of an organic, archetypal form moving through time and space. The surfaces created by this technique look ovalesque, since they are created by a moving ellipse, explains Cox. "Imagine a semicircle spinning around 360° in space; it will trace out a sphere in three-dimensional space. Essentially, every one of these mathematical surfaces was created in a similar way," says Cox. She adds that 10 parameters in the program actually control the wobbling ellipse that traces out the surfaces. "This is the first time the Romboy homotopy has ever been animated," says Cox.

While the group's main intention was to create the Romboy homotopy, they discovered a new surface when they initially computed Apery's equations. "We sort of knew what some of these surfaces were going to look like in the transformation, but we really didn't know it was going to look like a Venus shape. It looks like an icon or a fertility object from Paleolithic times, so we dubbed our effort the 'Venus Project,'" says Cox.

Cox explains that the complexity of this project required the use of a Cray system. "Without the Cray system we never would have realized the animation. It just takes too long to calculate all the points in space that create one of these surfaces. It simply made life easier," she says. "The surface is very multidimensional and weaves in and out of itself. Every point on that surface has to be calculated, including color and light intensity." Each second, 30 of these calculations were solved to create the two-minute animation.

Just as each member of the renaissance team brought unique talents to the project, each gained valuable insights. Francis discovered a new mathematical surface that had not been anticipated by other topologists. He generated new mathematics and demonstrated the significant role computer graphics can play in mathematical research. Idaszak used the Cray system to build his own ray-tracing algorithm, and to develop new computer graphics techniques that depict surfaces that penetrate themselves and quickly change direction.

By altering the program's parameters, Cox now has a method of developing new shapes and unusual art. "I just keep making images with this wonderful sculpture machine that makes beautiful transparencies," she says. "This is the first time organic sculpture has ever been done with a supercomputer, and the first supercomputer art that was done by an artist supported by the National Science

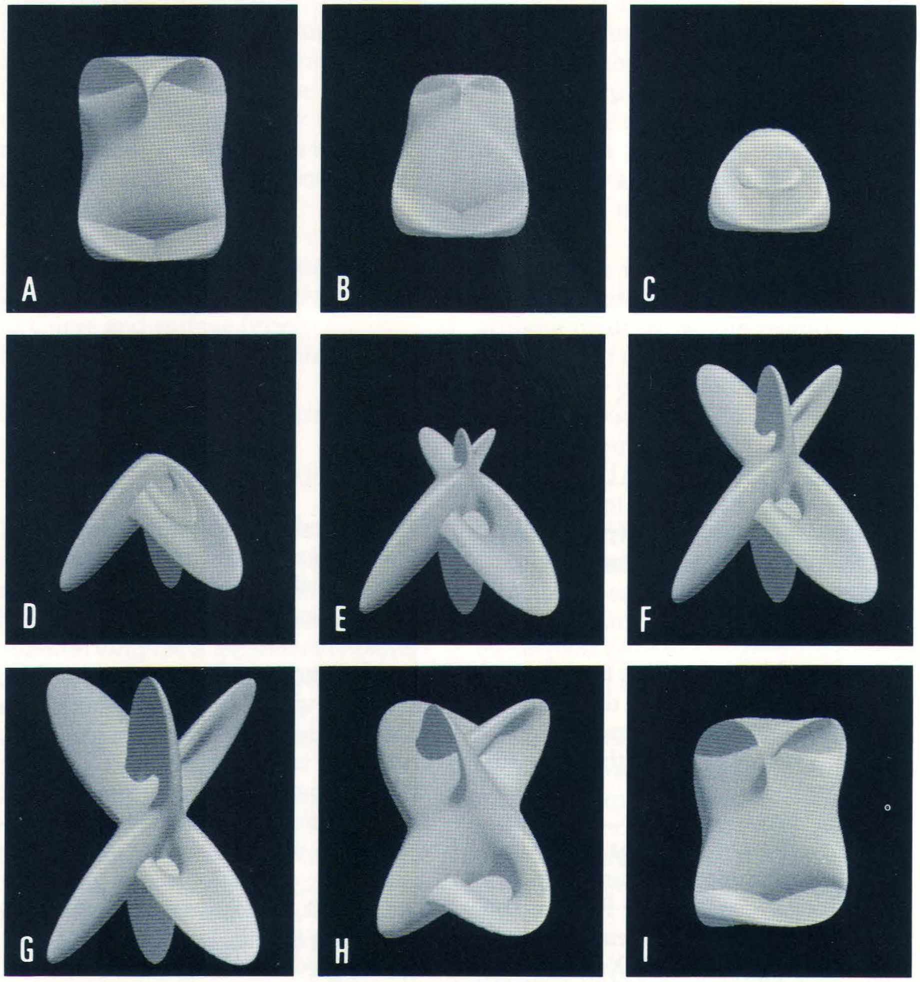

Selected frames from "Metamorphosis: Shadows from Higher Dimensions," produced by Cox, Francis, and Idaszak at the National Center for Supercomputing Applications. The sequence illustrates the Romboy homotopy starting with the transformation of Venus (A) into the Roman surface (C). Next, the Roman surface evolves into the Boy surface (D), which then turns into a new surface called Ida (G). Ida then evolves into the original Venus figure.

Venus figure as illustrated by Cox, Francis, and Idaszak. This shape represents shadows from a higher dimension.

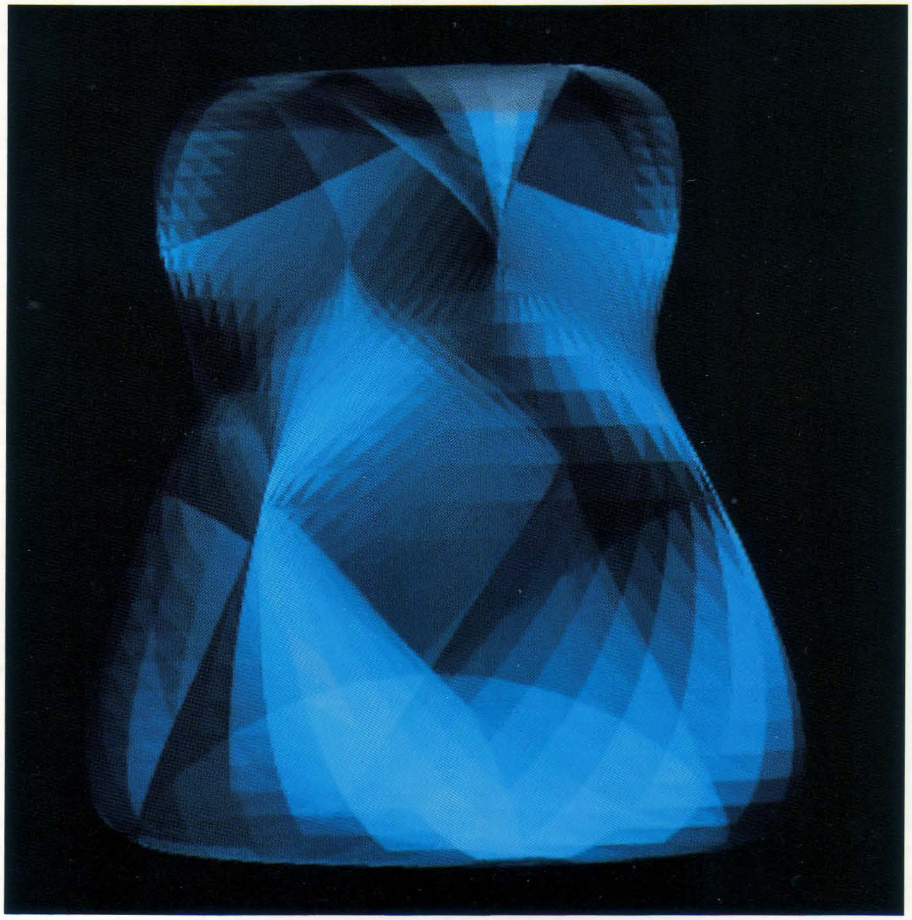

Transparent view of Venus.

Foundation." Cox has been working with two artists in Chicago to create three-dimensional images of these new shapes. Because of the unique nature of this art, Cox even has been approached by art collectors. "In the art world these images are hot stuff," she says.

All three contributors have opened doors to future paths of discovery, setting an example of how computers can connect the talents of scientists and artists to help others visualize the abstract. "Supercomputers and computer graphics systems are providing the advanced tools; interdisciplinary teams will provide the creative intelligence," says Cox.

Air Force saves fuel dollars with Cray system

By using a CRAY X-MP system to provide wind forecasts for use in computer flight plans, the Air Force Global Weather Central (AFGWC) has made it possible for the U.S. Air Force to reduce fuel costs by $14 million over the past two years, according to Captain John Pace, assistant chief of the Numerical Models Section.

"The wind forecasts the Cray system produces are about 20 percent better than the older system's," says Pace, explaining that the Cray system's accurate wind forecasts enable Air Force pilots to save $7 million per year by efficiently using tail winds and avoiding head winds. "This saves a great deal of money because they can plan their routes more efficiently," says Pace.

The AFGWC could be considered the hub of worldwide weather services, providing data for the U.S. Air Force, U.S. Army, and other governmental organizations. The AFGWC's CRAY X-MP/22 system is used primarily for running two models that analyze and forecast weather around the world. "We have a model called the High Resolution Analysis System, which takes in weather observations from all over the world to produce a detailed weather analysis. Based on that analysis, we run a forecast model called the Global Spectral Model that produces 96-hour forecasts four times daily," says Pace.

Since its installation in 1985, the Cray system has provided data to forecasters stationed around the world. They use the information to support space shuttle launches, to prepare severe weather warnings for over 450 specific locations in the continental United States, and to support operational readiness inspections for the Strategic and Tactical Air Command and the Military Airlift Command. In fact, since 1985, the AFGWC has improved the average lead time for U.S. severe weather warnings from 30 minutes to 75 minutes. Operational commanders now have more that twice as much time to implement resource protection measures.

"The Cray system is powerful enough to provide a forecast for the entire globe. Our older models were forced to split the world into three regions: a tropical region, a Northern Hemisphere region, and a Southern Hemisphere region," says Pace, adding that these forecasts were often poor because artificial boundaries impeded the simulation of realistic interactions between regions. "With the new model we do not have these problems any more," says Pace. "Forecasters in the Pacific region (Hawaii, Japan, Guam, and the Philippines) have complimented the AFGWC on improvements in forecasts."

Pace says AFGWC's models initially represent the atmosphere as a three-dimensional grid. "The more grid points you have, the better the resolution, and the better the forecasts can be. Therefore, we need a large and fast computer to accomplish the forecasts quickly so we can make use of the results."

AFGWC's models then use a spectral technique, in which the grid points are replaced by functions that describe the global weather patterns. Then the models solve differential equations that predict how these functions will change as time passes. This approach produces accurate forecasts, but requires a very large number of calculations. To manage this problem, AFGWC's models use mathematical techniques such as Fast Fourier Transforms, Gaussian quadrature, and Legendre polynomials to reduce the number of calculations.

A Shinto ceremony in Chippewa Falls to help ensure prosperity.

But Pace says the need for the power of a Cray system goes far beyond three-dimensional modeling. "We want to include as many atmospheric effects as we can. There are many that occur on scales smaller than the grid spacing of the model, so we parameterize them." These effects include cloud formation, precipitation, and interactions between oceans, land, and the atmosphere.

Next, AFGWC plans to extend its models to provide high-resolution regional forecasts and to support electrical-optical weapons systems. They are developing models to forecast conditions in the stratosphere and near-space, and provide data to support the National Aerospaceplane development. "We will have to take into account more data sources than we now have," says Pace. New data sources include Doppler radar, improved battlefield weather observations, new types of satellite observations such as data derived from microwave sensors, and a wind profiler that provides high-resolution wind observations.

Vectorization is used extensively at the AFGWC, resulting in reduced CPU run time of the main section of the Global Spectral Model from 75 minutes to 50 minutes. Vectorization also helped computer users shorten one subroutine of the program from 90 seconds to 3 seconds. "It's important to optimize this model because it allows us to do more work. Now we're able to incorporate 35 instead of 20 observations around each grid point for better analysis," says Pace.

East meets West in Chippewa Falls

Cray Research's rigorous system of reliability checks helps ensure that Cray systems stay up and running. But customers sometimes choose to take matters into their own hands. In September a traditional Shinto ceremony was performed at Cray Research's systems check-out building in Chippewa Falls, Wisconsin, to introduce a protective spirit into a new CRAY X-MP/24 computer system. The ceremony was requested by Century Research Center Corporation (CRC), the Japanese service bureau that ordered the system. In 1980 CRC became Cray Research's first customer in Japan when the bureau ordered a CRAY-1 system.

The ceremony held in Chippewa Falls had a twofold purpose: to help protect the system during shipping and to ensure that it will run smoothly once it is installed on site."The ceremony places a spirit in the computer that ensures the company will prosper as customers use the system," explains Tony Hagiwara, manager of the Minneapolis/St. Paul office of Kintetsu World Express (KWE), the shipping firm that will deliver the system to Japan.

The Shinto religion predates Buddhism in Japan; the two religions coexist peacefully in modern Japanese culture. The ceremony traditionally is performed in Japan as a kind of blessing for significant events, such as an important purchase or the dedication of a new building.

In preparation for the ceremony, plates holding seaweed, rice, vegetables, salt, and water, along with a bottle of sake and several small sake cups, were set on a platform in front of the computer. The food and drink served as an offering to the spirit, so that all forms of nourishment would be available to it.

Yoshi Tani, an employee of KWE Japan, led the ceremony. Although he is not a Shinto priest, he has received special training. The other participants were Tomoo Takahara, president of CRC; Mitsuru Maruyama, director of CRC's systems engineering division; Yoshikazu Hori, president of Cray Research Japan, Limited; Jim Otis, Cray Research's coordinator of country services; and Carl Diem, Cray Research's director of marketing support. Each participant laid an evergreen branch on the table before the computer system, bowed, clapped hands twice, and bowed again. When all had finished, the six participants each had a cup of sake and the brief ceremony was over.

Tani brought from Japan a small wooden easel on which were written, in Japanese, the names of Cray Research and CRC. The easel will travel to Japan with the Cray system and sit nearby the system once it is installed, so that the spirit will continue to guard the system, helping keep its operation trouble-free.

Cray systems decipher chemical reactions

Catalysts play an important role in a variety of chemical reactions, including detoxifying poisonous gases, converting raw petroleum to octane, and synthesizing fertilizers. Improved catalysts mean more efficient results and higher yields for chemical companies such as Dow Corning of Midland, Michigan, a major producer of silicone polymers. Researchers at Northwestern University, Dow Corning, RCA Laboratories, and Cray Research have used Cray supercomputer systems to gain a clearer picture of the atomic-level process of catalysis.

Specifically, the scientists wanted to explore bonding between a silicon surface and the noble metals copper and silver. The scientists wondered why copper and silver, two similar elements, react differently with a silicon surface. Although they understood that copper is the superior catalyst, they wanted to pinpoint qualities that make it unique.

To answer this question, the scientists calculated the adsorption geometry and energetics of copper atoms on silicon and of silver atoms on silicon using a fully quantum mechanical approach. The researchers treated the problem by solving the local spin-density equation self-consistently with the DMOL/86 program using approximately 40 hours of CPU time on the CRAY X-MP/48 and CRAY-2 systems at the Cray Research computer center in Mendota Heights, Minnesota.

"Only a Cray system can handle this type of calculation," says Shih-Hung Chou, a post-doctoral research scientist at Northwestern University working with professor Arthur Freeman. Processing speed, fast I/O, and large memory were essential, according to Chou, who carried out the calculations as a visiting scientist at Cray Research. To better analyze the data, they produced three-dimensional images of their findings on various display devices using the OASIS ray-tracing program. They also used UNIRAS software to create two-dimensional images.

As a result, the scientists found a clear difference in the way that copper and silver atoms react with a silicon surface. Most strikingly, they discovered that copper has a higher chemisorption energy and a significantly lower penetration barrier than silver. This means that at elevated temperatures, copper atoms can diffuse easily into the silicon surface and weaken the bonds between silicon atoms, whereas silver atoms stay on top of the silicon surface. Because the silver atoms do not penetrate the surface, they have a smaller weakening effect on silicon-silicon bonding. The scientists believe these structural and electronic differences may account for copper's higher catalytic activity because the weakened silicon bonds ease silicon's reaction with chemicals such as methylchloride.

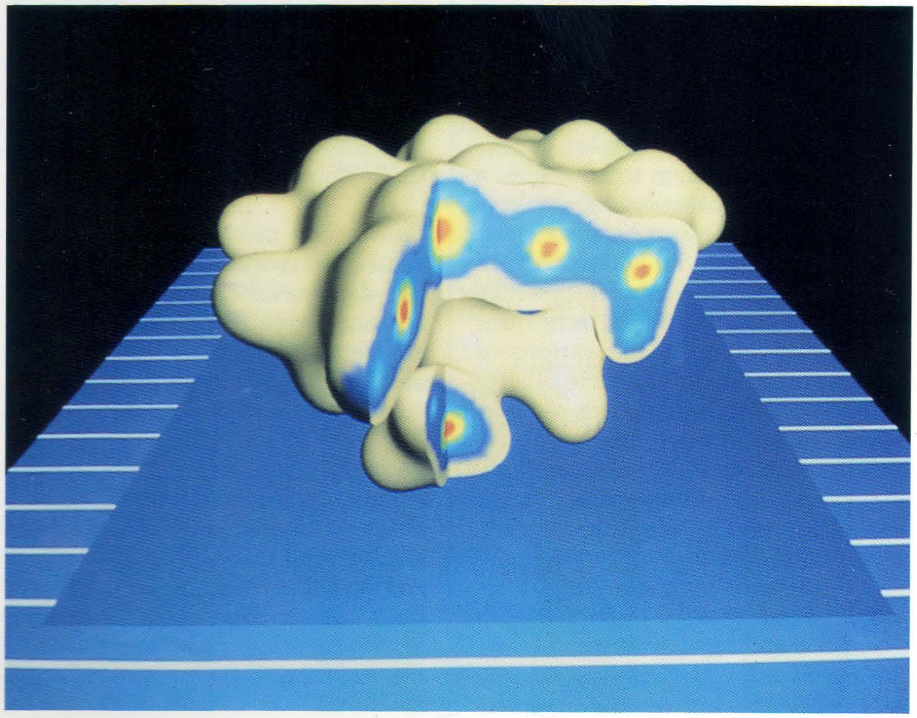

A cluster of silicon atoms shows the interaction between copper atoms and a silicon surface. The electron density is shown with the copper atom in the calculated equilibrium position. The cluster is cut open to reveal the increasing electron density, as indicated by the colors changing from blue to red.

"Chemical companies are excited because this information was not available before; the calculations simply were not possible," says Erich Wimmer, chemistry applications manager at Cray Research. Wimmer, who initiated the project with Dow Corning researchers, says that by using these results to design new experiments, scientists can better understand catalytic activity.

Tom Gentle and Stelian Srigoras from Dow Corning explain, "The type of process we studied was introduced over 40 years ago, but no one has understood it. This is a step toward understanding the process on the atomic level."

Companies share resources to break molecular-design bottleneck

Cray Research and a group of companies with varied proprietary interests are pooling their resources to achieve a unified goal — to develop consistent molecular parameters for more accurate computer-aided simulations. The companies hope to relieve a bottleneck that has inhibited molecular design work: greater supercomputer power, but inconsistent estimates of parameters.

The consortium was founded in February 1986 by Biosym technologies, a research company that develops molecular modeling software. Arnold Hagler, co-founder of Biosym, noticed tremendous advances in computer technology and in molecular simulation techniques, yet few advances in derivations of molecular energy surfaces. He believes consistent, reliable depictions of molecules' energy surfaces are essential to achieve accurate models, since all applications depend on representations of energy as functions of molecular coordinates and distances.

John Wendoloski, DuPont senior researcher and consortium member, encountered a similar bottleneck. "Molecular mechanics and dynamics have a lot of potential in the entire life-science area. The lack of parameters has been one of the factors that has prevented other scientists from using them," he says. Since available parameters are as varied as the researchers who derive them, Wendoloski describes the practice as an "art form."

In fact, Hagler heard those sentiments echoed from many scientists. "Now that chemical, polymer, and drug companies have become interested in molecular modeling, I saw an opportunity to make a very extensive effort to advance the scientific rigor of these calculations," explains Hagler, who is also head of the biophysics department at the nonprofit Agouron Institute in La Jolla, California. "It was in the mutual interest of the companies because they need these functions and parameters for more accurate simulations, but for a single company to fund it would be a very expensive venture. It's basic science, so they are willing to share the results," he says.

The consortium now comprises 11 members representing a diverse base of interests. Their pooled resources support the team of five scientists and four programmers and research assistants at Biosym who pursue the development of potential functions and parameters. Besides Cray, Biosym, and DuPont, members include: Abbott Laboratories, Battelle Pacific Northwest Laboratory, ETA Systems, Imperial Chemical Industries, Merck Sharpe & Dohme, Monsanto, Rohm and Haas, and Upjohn.

Consortium members work independently, but twice annually they attend two-day workshops to share information and designate areas of study. "They are learning state-of-the-art techniques and fine points in applying energy calculations," Wendoloski says.

Each company contributes molecular parameters obtained from experimental data and quantum mechanical calculations. Wendoloski is using DuPont's Cray system to calculate potential functions, while Hagler is using Cray systems at Mendota Heights, Minnesota, and at the San Diego Supercomputer Center.

Deriving parameters for molecules is a problem that expands very quickly, using large amounts of computing time. "Without computers like Cray systems, this consortium would not be possible," Wendoloski says. Also, intermediate numbers are generated that require the large storage capacity of a Cray system.

The quantum mechanics equations produce second derivative, or force-field matrices for the molecules. Consortium scientists at Biosym are continuing to develop software, which is needed to fit the parameters into potential functions used by molecular mechanics programs.

"So far, we have made tremendous progress," Hagler says. "The companies have contributed quite a bit of manpower and results." Consortium scientists have developed a program to test the functional form of the potential energy and to obtain its parameters. They also have compiled a database of experimental properties such as internal coordinates from gas phase measurements, vibrational frequencies, and lattice constants of crystals for approximately 250 compounds. This growing database also includes results of ab initio calculations, which will be used to test functional forms and constants.

Although consortium scientists have completed thousands of calculations, published research papers, compiled an extensive database of experimental properties, and provided members with software for molecule fitting, they only have scratched the surface of deriving parameters, according to Hagler. "There is a tremendous amount of work involved in doing this, and we are only at the first generation of force fields," he says.

Based on its success to date, Hagler plans to expand the consortium to involve additional participants, including members of academia who would carry out experimental projects such as measuring vibrational frequencies, sublimation energies, and various other experimental properties. The consortium aims to derive rigorous parameters for all organic functional groups, and derive some parameters for inorganics and catalysts as well. Says Hagler, "If we can do all of that, it will be one of the most important contributions to molecular simulation and design that can be made today."

WORLDWIDE BUSINESS CENTERS

Domestic

Corporate Headquarters
608 Second Avenue South
Minneapolis, MN 55402
(612) 333-5889, telex 6879144

**Government Marketing/
Government Relations Office**
National Place
1331 Pennsylvania Avenue, Suite 1331 N
Washington, DC 20004
(202) 638-6000

Central Region Office
3375 Mitchell Lane
Boulder, CO 80301
(303) 447-8702
- Albuquerque
- Chicago
- Colorado Springs
- Detroit
- Minneapolis/St. Paul
- St. Louis

Eastern Region Office
4041 Powder Mill Road, Suite 600
Calverton, MD 20705
(301) 595-5100

Atlanta	Huntsville
Bernardsville	New York
Boston	Pittsburgh
Cincinnati	Tampa
Darien	Washington, DC

Petroleum Region Office
5847 San Felipe, Suite 3000
Houston, TX 77057
(713) 975-8998
- Dallas
- Tulsa

Western Region Office
5776 Stoneridge Mall Road, Suite 350
Pleasanton, CA 94566
(415) 463-2800
- Huntington Beach
- Los Angeles
- Salt Lake City
- Seattle
- Sunnyvale

International

Cray Asia/Pacific Inc.
Suite 1003
Sutherland House
3 Chater Road Central
Hong Kong
(5)-214669, telex 89765

Cray Canada Inc.
4211 Yonge Street, Suite 210
Willowdale, Ontario
Canada M2P 2A9
(416) 229-2729

Cray Research France S.A.
7 Rue de Tilsitt
75017 Paris, France
(01)-766-01-55, telex 660568

Cray Research GmbH
Perhamerstrasse 31
8000 Munich 21, West Germany
(089)-56014-0, telex 5213211

Cray Research S.R.L.
Via Vivaio, 15
20122 Milan, Italy
(02)-70-98-27, telex 315403

Cray Research Japan, Limited
Ichibancho Eight One Building
6-4 Ichibancho
Chiyoda-Ku, Tokyo 102, Japan
(03)-239-0711, telex 28893

Cray Research (UK) Ltd.
Cray House
London Road, Bracknell
Berkshire RG12 2SY
United Kingdom
(0344)-485971, telex 848841

CRAY RESEARCH, INC.