

CRAY CHANNELS

VOLUME 15, NUMBER 1 • A CRAY RESEARCH, INC. PUBLICATION

E N V I R O N M E N T



In this issue

Concern about environmental conditions is growing in both the industrialized and nonindustrialized worlds. Deforestation, industrial emissions, and other factors threaten to trigger far-reaching environmental changes, according to some scientific studies. Determining the likelihood and extent of such changes requires further research into environmental systems and their interrelationships. Cray Research supercomputers provide the capabilities needed for timely and cost-effective environmental modeling in climate research and other applications.

This issue of CRAY CHANNELS surveys environmental modeling on Cray Research systems conducted by researchers around the world. Included are reports from Australia, Japan, The Netherlands, and the United States. We also cover research in Canada to improve the reliability of electric power grids and in the United Kingdom to understand the dynamics of star formation. This issue also describes the newest version of Cray Research's UNICOS operating system.

The environmental feedback loops that link the atmosphere, oceans, cryosphere (snow and ice), and terrestrial ecologies make comprehensive climate modeling a formidable technical challenge. But only integrated modeling will reveal the weaknesses and resiliencies of the biosphere as a whole. And only Cray Research computer systems provide the processing speed, storage capacity, and communications bandwidth needed to make Earth-system modeling a realistic goal.

Editor's note: Beginning with this issue, CRAY CHANNELS will be published three times yearly.

Features



2

4

8

12

17

20

23

26

Departments

CRAY CHANNELS is a publication of the Cray Research, Inc. Marketing Communications Department. Published three times per year, it is intended for users of Cray Research computer systems and others interested in the company and its products. Subscription inquiries and address changes should be sent Attention Department D, CRAY CHANNELS, Cray Research, Inc., 655D Lone Oak Drive, Eagan, Minnesota 55121.

Volume 15, Number 1

Editorial staff

Ken Jopp, managing editor
Mary Miller, associate editor
Heidi Hagen, contributing editor
Chris Kruell, contributing editor
Jeff Paulson, contributing editor

Design and production

Eric Hanson, graphic designer
James Morgan, art director

29

30

31

33

Announcing UNICOS 7.0: Peak performance at your fingertips

Cray Research's latest release of the UNICOS operating system achieves a new standard of reliability to deliver the most advanced and feature-rich UNIX environment in the computer industry.

Hydrology and topography mapping from digital elevation data

Brian Davis, Hughes STX Corporation, USGS/EROS Data Center, Sioux Falls, South Dakota

When scientific workstations prove inadequate for large image-processing applications, Cray Research systems provide a capable alternative.

Hydrodynamic and transport modeling of estuarine and lake waters

Y. Peter Sheng, University of Florida, Gainesville, Florida

Large-scale computational models help predict the spread of water pollutants.

Global climate change modeling in Australia

Bryant J. McAvaney, Australian Bureau of Meteorology Research Centre, Melbourne, Australia

Australian researchers play an important role in international efforts to compare global climate models and refine the models' predictive abilities.

Seasonal variability simulation using a general ocean circulation model

Tomonori Matsuura and Toshikazu Morohoshi, National Research Institute for Earth Science and Disaster Prevention, Tsukuba, Japan

Investigating global water cycles as part of a ten-year research project, NIED researchers are able to simulate the seasonal variability of the general ocean circulation.

Optimization of shallow sea models

Arnold Heemink, Tidal Waters Division, Rijkswaterstaat, The Netherlands

Theo van Stijn, Informatica Centrum voor Infrastructuur en Milieu (ICIM), Rijswijk, The Netherlands

Water level forecasts provided by a computational storm surge forecasting system developed at ICIM help protect the Dutch coast from destruction by the North Sea.

Supercomputing for power system analysis at Hydro-Québec

Gia Tong Vuong and Roger Chahine, Centre d'Analyse Numérique de Réseaux, Hydro-Québec, Montréal, Canada

Luc Ostiguy and Steve Behling, Cray Research, Inc.

At Hydro-Québec, a Cray Research supercomputer is dedicated to power system analysis to help deliver stable and reliable service.

Simulating the birth of a star

A. H. Nelson and A. P. Whitworth, University of Wales College of Cardiff, Cardiff, Wales

New simulation techniques lend understanding to the formation and evolution of stars.

Corporate register

Applications update

User news

Gallery



Announcing UNICOS 7.0— peak performance at your fingertips

Continuing its tradition of leadership in software development, Cray Research announces release 7.0 of the UNICOS operating system—the most advanced and feature-rich UNIX programming environment available for any computer system.

Since the first release of UNICOS in 1986, Cray Research has worked to make UNICOS a powerful, mature, reliable, production-quality supercomputing environment. Based on UNIX System V, Release 3—with selected Release 4 features—UNICOS 7.0 was developed and tested to achieve a new standard in reliability and lay the foundation for Cray Research's massively parallel systems. In keeping with Cray Research's commitment to standards, the UNICOS system conforms to the POSIX 1003.1 standard.

Performance enhancements

Release 7.0 of the UNICOS operating system includes significant performance enhancements:

- New scientific libraries (LIBSCI) include highly efficient algorithms and code tuned specifically to the Cray Research architecture. This includes 3:1 cyclic reduction routines that are 3-10 times faster than a standard tridiagonal solver. Virtual

basic linear algebra subroutines (BLAS) offer high-speed out-of-core solutions of linear equations for problems that are too large to fit in memory.

- A cache layer and double buffering have been added to the Flexible File I/O (FFIO) facility, the feature that supports on-the-fly data conversion, to increase I/O speed as much as tenfold.
- Shared text is now possible for multiprocessed applications, enabling better use of system memory.

Improved visual tools

The 6.0 version of the UNICOS operating system included the Application Development Toolset, providing graphical and verbal interpretations of application performance data and offering suggestions to improve performance. With UNICOS 7.0, the toolset has been expanded to include xbrowse, xfm, and xprof tools, and many previously existing tools have been enhanced. To support equivalent capabilities for line-oriented and batch usage, these tools are provided in a nonvisual format where applicable.

- The Cray Research browser, xbrowse, provides a visual, interactive environment for viewing



and editing Fortran code. Xbrowse provides language-sensitive source code analysis, such as variable cross reference, variable tracing, and call tree information.

- The Cray Research File Manager (xfrm) is a new graphical user interface to UNICOS that lets users execute, copy, and remove files, and navigate through directory structures without typing UNICOS commands. Files and directories are represented by icons, and commands are executed by clicking a mouse button on the icons or pull-down menu selections.
- The Cray Research process monitor, xproc, provides a visual interface that allows users to monitor the status of all their processes running on a Cray Research system. Xproc visually displays information from the Network Queueing System (NQS) qstat and UNICOS ps commands and allows the user to customize the display in a variety of ways.
- Jumptrace and jumpview tools collect and process simulation and runtime statistics to help users optimize their programs. With UNICOS 7.0, these tools now deliver precise timing information at the instruction level and help in selecting candidates for inlining.

- Profview, which provides a statistical profile of program execution times, now splits individual functions from C programs so they appear to be separate modules. This allows users to view statistics for each function separately. The memory graph provides further functionality, allowing users to zoom into the smallest level of detail.
- Atscope, which provides data scoping assistance, now aids cross-module parallelism and helps users apply Autotasking directives manually. Additionally, atscope has been incorporated into the new UNICOS browsing tool, xbrowse.
- The cdbx debugger has been enhanced to support C-style casting in expressions and has a new fuzzy-match heuristic for locating mistyped symbol names.

With UNICOS 7.0, other development tools also have been added or enhanced:

- Lint is a tool that performs static analysis of Standard C code and reports potential errors in syntax, grammar, and type usage for variables, functions, function arguments, and more. With UNICOS 7.0, the Fortran version of the lint utility is available via cflint.
- The cflist tool now is available with UNICOS 7.0. This tool lists Fortran programs with cross-references, loop and parallel marks, and static analysis. It embeds syntax error messages (if any) with the listing and indicates jump and target directions.

Binary option

Many customers with CRAY Y-MP and CRAY X-MP systems have no need to modify UNICOS source code. These customers now can choose to receive the software in binary form only, enabling them to enjoy the quickest installation possible and save storage media without acquiring a source license from UNIX System Laboratories. Binary licenses are available through Cray Research.

Publications and training

To better meet customer needs, the UNICOS 7.0 publications set includes ready references for software documentation and for the UNICOS shells, environment variables, and the installation menu system. Additionally, a software documentation map now appears in each UNICOS manual. Organized by user function, this map highlights available documents and indicates which documents also are available online through the Docview utility.

Training workshops and an update workbook have been developed to acquaint customers with new UNICOS features. Videotapes are available for those who cannot attend the workshops.

For more information on UNICOS 7.0, refer to the *UNICOS 7.0 Release Overview*, publication number RO-5000 7.0; contact your Cray Research representative; or dial 800/284-CRAY in the United States. ■

Hydrology and topography mapping from digital elevation data

Brian Davis, Hughes STX Corporation
USGS/EROS Data Center
Sioux Falls, South Dakota




Figure 1. Shaded relief of the Puget Sound input data.

The study of global climate change requires the computer processing of very large datasets. Organizations such as the U.S. Geological Survey's (USGS) Earth Resources Observation Systems (EROS) Data Center routinely process remotely sensed satellite imagery to produce detailed geographic maps. This processing, however, has been limited to relatively small datasets, representing relatively small geographical areas. Global change studies, by contrast, require the processing of enormous datasets too large to be processed with conventional computing resources. Although satellite image processing is not a traditional supercomputer application, processing the datasets involved in global change studies will require supercomputing resources.

To explore the feasibility of using supercomputers for this application, the EROS Data Center has been processing large datasets using Cray Research systems. These runs have involved

the extraction of hydrology and topography information from digital elevation models (DEMs), and the initial results are very encouraging. A moderate amount of software optimization, in addition to the hardware capabilities of the systems, has provided substantial reductions in the runtimes needed to process the datasets, relative to the speeds obtained using scientific workstations.

Software optimization

A DEM is an ordered array of numbers that represents the spatial distribution of elevations above some arbitrary point in a landscape. It may consist of elevations sampled at discrete points over the average elevation or a specific segment of the landscape. Digital elevation data are available for selected areas in several forms varying in x-y resolution and accuracy.

Drainage analysis is among the applications for which DEMs are useful. Researchers at EROS have developed software tools to extract topographic structures and delineate watersheds and overflow paths from DEMs.¹ These tools have even been used to analyze the topography of the Martian landscape using DEMs constructed from Viking images.² The software, called Drainage, is a set of special-purpose Fortran programs that can be interfaced with general-purpose raster or vector spatial analysis packages. The output datasets created by Drainage can be used by relational database management packages. Although Drainage is not a typical remote sensing application, it is an example of an algorithm developed for small-area study sites that demands high-performance computing when applied to a large-area dataset. This example exposes many of the limitations of conventional computing resources when applying proven techniques to datasets of a continental or global scale.

Among the projects for which the Drainage software has been used was an analysis of a dataset describing Puget Sound in Washington state (Figure 1). The resolution was 30 m between pixels and the dataset was 8977 lines by 7616 samples in size. The knowledge gained while optimizing the processing of this dataset can be used for current projects like the UCM (an acronym for the U.S. and parts of Canada and Mexico) dataset and possible future projects, including a proposed digital raster elevation dataset of the globe.

The Drainage software was ported to the IRIX operating system to process the Puget Sound data on a Silicon Graphics 4D/480 Power Series system. Because of continual emphasis on portability, the transfer involved only changing the statements controlling the data record size such as OPEN, READ, and WRITE.

Over one-half of the speedup obtained for the Silicon Graphics system was accomplished by a new release of the IRIX operating system which included new Fortran I/O. The old Fortran library I/O calls were replaced by simple calls to the operating system's C I/O calls, enabling the new Fortran I/O calls to be only slightly slower than C I/O calls.

Additional speedup was accomplished through normal scalar optimization: breaking up

the code; analyzing the logic of the inner loops of a typical algorithm; and applying any applicable changes to similar loops, algorithms, and modules.

Performance analysis tools revealed that parallel processing opportunities were very limited and would have required extensive loop rewrites. Even if the loops were rewritten, however, the amount of work in each loop would have increased, and the speedup would have been only a fraction of the potential (potential speedup being roughly the runtime on one processor multiplied by the number of processors). With the number of varying algorithms involved, the fraction of potential speedup could have been very small.

In the data-dependent algorithms, a pixel receives its new value from one of its boundary pixels, which may in turn have received its value from a boundary further out on a previous pass through the data. The Drainage module CONCAVE, which fills depressions, is a very inefficient algorithm to run in parallel because each new pixel value is dependent on many others. A pixel cannot be assigned its new value until those pixels immediately surrounding it have received their new values.

These initial attempts at parallelization indicated that a platform with a fast CPU and large memory would produce a faster turnaround time. Therefore, the next logical step was to investigate the performance of the Drainage software on a CRAY Y-MP supercomputer.

CRAY Y-MP optimization

The initial port of the software to Cray Research's UNICOS operating system involved the same I/O considerations as the IRIX port: OPEN, READ, and WRITE statements. Additionally, a data conversion routine was written to provide data to the UNICOS version of the code in eight-byte words, allowing for more efficient reading and writing. A complementary conversion routine was written to convert the output data for validation because the read and write conventions were unlike the original schemes. The software should not only run very fast, but also produce results identical to those from other platforms. This approach preserved the general-purpose nature of the software. It also allowed for variable input data and parameterization, thereby avoiding the creation of a single-use version of the software.

After the porting issues were resolved, the software was optimized for Cray Research's CF77 Fortran compiling system, the UNICOS operating system, and CRAY Y-MP-specific hardware capabilities. Because the dataset was large (66 Mwords) and some processing steps required multiple interim datasets of equal size to be used simultaneously for input to subsequent processing steps, optimization was attempted by using a 1-Gword SSD storage device. Some additional scalar optimizations for the Cray architecture also were incorporated. A large data array was created as a look-up table to replace one section of if-then-else logic. Eight Cray vector management calls plus one statement replaced a potential 32 statements in an innermost loop. A similar reduction of a potential 16 state-


ments down to 8 was accomplished by using vector management calls in another innermost loop. Also, for those cases that use only one input dataset, the entire image was read into memory so that multiple passes over the data could be done from array rather than disk I/O.

These optimization steps dramatically improved runtime. The improvement probably would have been acceptable for the UCM dataset (Figure 2) using the Puget Sound parameters. Analysis using the Drainage software, however, is not a canned procedure. As products are analyzed, parameters may need to be adjusted, and some or all of the processing steps repeated. Some modules

Figure 2. Relief map of the area included in the UCM project. Puget Sound presented the problems associated with processing high-resolution data covering a large area. UCM presents a new problem: processing data on a much larger, continental scale. The UCM data have a one-half km resolution, lower than the resolution of the Puget Sound data, but cover a much larger area, 24,318,695 km², resulting in an image size of 10,200 lines by 10,000 samples. The yellow polygon defines the Mississippi River drainage basin dataset, a subset of the UCM data.

need to make multiple passes over the data to resolve relationships, depending on the features contained in the data. In the case of the Puget Sound dataset, one of the modules required 15,911 passes. Another and more important consideration is that data anomalies and defects are uncovered as higher level products are assessed (Figure 3).

With all of the variables in addition to dataset size that affect runtime, the goal of the research was not to make the entire process interactive, but to condense the elapsed time of the process enough for efficient analysis, defect correction, and reprocessing.



CRAY Y-MP C90 optimization

To make the code more efficient, a Fortran main routine was written that called each of the Drainage modules as subroutines and passed the data as arrays in memory (or SSD) instead of using intermediate disk files to pass the data on to subsequent processing modules. Previously, disk reads and writes had been mimicked when possible by SSD I/O. The very large and fast memory of the CRAY Y-MP C90 system was underused, however, because the device performing the I/O merely had been replaced. For full memory use, the I/O logic needed to be modified. Because some modules require multiple interim datasets, the 128 Mwords of memory available on the CRAY Y-MP C90 system were insufficient to hold all the necessary data in memory arrays. The system was able to hold the single dataset being processed, and all other arrays were retained only as far away as the SSD. Because the I/O logic had not been modified, some interim datasets were still being stored on disk. By modifying the I/O logic, all data could reside in either memory or the SSD for the processing duration.

Disk reading would take place only at startup, to ingest the input data, and disk writing would take place to save the final image only after all processing was complete. Essentially all data traffic would take place within memory or at the SSD-to-memory bandwidth speed. This data handling was sped up due to the higher SSD-to-memory bandwidth of the CRAY Y-MP C90 system over the CRAY Y-MP system.

We currently are planning to use a CRAY Y-MP M90 system configured with 2 Gwords of memory. The input data will go directly from disk to memory, and no calculations or data manipulations will involve disk I/O or SSD I/O traffic. The data will reside totally in memory until all processing is complete and the resulting images are ready to be saved in a disk file.

Preliminary results are promising enough to state that this version of the Drainage software running on a CRAY Y-MP M90 system will allow a user to process a dataset comparable to that of the Puget Sound dataset in about 20 minutes. This runtime would allow for essentially uninterrupted analysis between reprocessing phases.

Conclusion

A satisfactory turnaround time has been achieved for a dataset of the Puget-Sound-class problem. Currently, attempts are being made to process the Mississippi River Basin dataset (outlined in Figure 2), which is an interim step to processing the UCM dataset. Both of these datasets are significantly larger than the Puget Sound dataset. Furthermore, the USGS is finalizing plans with the Defense Mapping Agency to create a digital raster elevation dataset of the entire globe, known as the Digital Chart of the World. The creation of this dataset is an ambitious task, and subsequent processing of a dataset of this size could require another generation of both software and hardware. ■



Figure 3. Anomalies, such as that in the lower left corner, must be identified as either a defect in the DEM data, an artifact of processing, or a valid feature. After identification and possible correction, processing is restarted.

Acknowledgments

The author thanks LeAnn Dix, Sue Jenson, Stuart Doescher, Dan Hogberg, John Lacher, Aleksander Mazur, and Pat Bruzek. Work by the author was funded by the U.S. Geological Survey; however, the views expressed are entirely those of the author. Computer time was provided by Cray Research. Software originally was developed by Susan K. Jenson of the USGS.


This article was condensed from the author's paper, "Using Cray Computers to Extract Hydrology and Topography from Digital Elevation Data," which was presented at the fall 1992 Cray User Group Conference and published in the proceedings.

About the author

Brian Davis received a computer science degree from Augustana College in Sioux Falls, South Dakota. He works for Hughes STX Corporation, contractor to the U.S. Geological Survey's Earth Resources Systems Data Center near Sioux Falls. He is a member of the software development staff, supporting scientists using satellite remote-sensing analysis to conduct global change research.

References

1. Jenson, S. K. and J. O. Domingue, "Extracting Topographic Structure from Digital Elevation Data for Geographic Information System Analysis," *Photogrammetric Engineering and Remote Sensing*, Vol. 54, No. 11, pp. 1593-1600, November 1989.
2. Jenson, S. K., "Applications of Hydrologic Information Automatically Extracted from Digital Elevation Models," *Hydrological Processes*, Vol. 5, pp. 31-44, 1991.



Hydrodynamic and transport modeling of estuarine and lake waters

Y. Peter Sheng, University of Florida, Gainesville, Florida

Most of the world's water bodies (estuaries, lakes, and continental shelves) are becoming increasingly polluted due to growing populations and point-source and nonpoint-source discharges from municipal, industrial, and agricultural run-offs. This has led to a decline in water quality and a reduction in the abundance and diversity of aquatic plant and fishery species. The scope of the problem in any particular location is influenced by various hydrodynamic processes (circulation, wave climate, and turbulent mixing) and transport processes (sediment transport, salinity transport, nutrient dynamics, and oil spill movement). Figure 1 shows several important hydrodynamic processes that contribute to the dispersion and fate of contaminants in an estuarine or coastal environment.

Circulation and salinity transport in estuaries

Most of the world's population centers around numerous estuaries, which provide vital water, food, and recreational resources. Numerous contaminants (waste heat, agricultural runoff, municipal sewage, and industrial waste) are discharged into the estuaries through rivers, canals, and nonpoint sources along the shorelines. Large-scale circulation, driven by tide and wind, transports these contaminants through the estuaries while small-scale turbulence mixes them vertically and horizontally. Although flushing due to large-scale circulation will lead to the seaward transport of some of the contaminants, many may reside in the bottom sediments over extended periods. Some of the contaminants in the bottom sediments, however, can be brought back into the water column and degrade the water quality during high-energy episodic events, such as the passage of a storm or weather front.

Estuarine circulation is quite complex and must be studied by means of numerical modeling and field monitoring. Due to recent advances in field instruments, it is possible to obtain long-term in-situ measurements of wind, wave, current, salinity, water temperature, and other water quality parameters in estuaries. Laboratory models are not suited to the study of turbulent flow and sediment transport due to inherent problems in scaling the field conditions. As a first step toward understanding the circulation and transport in estuaries, it is necessary to quantify the circulation associated with tide (with dominant time scales of 12 hours and 24 hours), wind (with time scales ranging from a few hours to a few days), and waves (with a dominant time scale of a few seconds). While large-scale circulation provides the fundamental mechanism for basin-wide flushing of materials within an estuary, small-scale (on the order of a few centimeters to tens of meters) turbulence within an estuary affects the vertical and horizontal mixing of water and other materials. Both the large-scale circulation and the small-scale turbulence are influenced strongly by the complex geometry and bathymetry of an estuary, in addition to the effect of the Earth's rotation.

Tidally averaged circulation in an estuary consists generally of a two-layered "gravitational"




Figure 1. Various components and processes affecting the circulation, transport, and fate of sediments and contaminants in estuarine and lake environments. The red arrows indicate boundary fluxes; the black arrows indicate internal fluxes.

circulation: salinity from the ocean is transported into an estuary primarily through the bottom water, while fresh water from the rivers is transported seaward through the surface water. Salinity in bottom water often is higher than in the surface water, thus producing a vertical stratification that tends to suppress the vertical mixing. During the passage of a cold front, however, cooler surface water and warmer bottom water produce an unstable stratification which tends to create large-scale convective overturning in the water column.

Numerical models of estuarine circulation and transport

Numerical models of estuarine circulation and transport are needed to gain a scientific understanding of complex estuarine processes and to provide tools for management. Based on long-term data obtained from various estuaries, it is possible to construct reliable numerical models for estuaries. The starting point of the circulation models is Newton's law of motion, which can be expressed by three-dimensional, time-dependent, partial differential equations that describe the motions of fluid particles. Basically, the acceleration of a fluid particle in an estuary is influenced by the following processes: horizontal advection, vertical advection, Coriolis acceleration, barotropic pressure gradient (surface slope), baroclinic pressure gradient (density gradient), horizontal turbulent mixing, and vertical turbulent mixing. Typically, an estuary is divided into a number of grid points, with horizontal spacing ranging from tens of meters to a few kilometers and vertical spacing ranging from the

entire water depth to one-tenth of the water depth. Simulations on the order of one month to one year are needed for understanding the long-term circulation and transport in an estuary. Thus, an estuarine circulation model often is expected to reproduce the observed circulation (due to tide, wind, wave, and density gradient) with spatial scales ranging from a few centimeters to the basin scale and with time scales ranging from one tidal cycle to one year. This puts a severe demand on the numerical model. Despite the advances in workstations and microcomputers, many of the environmental modeling problems require the use of Cray Research supercomputers. With the availability of CRAY Y-MP systems, many interesting estuary simulations have been produced.

For example, funding provided by the U.S. Environmental Protection Agency's National Estuary Program permitted a study of circulation and salinity transport in Sarasota/Tampa Bay in Southwest Florida.¹ Five instrument platforms obtained long-term data of wind, water level, current, salinity, and temperature in Sarasota Bay. These data have provided valuable insight into the estuarine circulation driven by tide, wind, and density gradient. Starting from the CH3D (Curvilinear-Grid Hydrodynamic 3-D) model developed for Chesapeake Bay² and James River,³ a robust three-dimensional circulation model for Sarasota Bay was developed and validated using the field data. Figure 2 displays the results of a three-dimensional simulation model that reproduced the formation and movement of the observed front in the James River. An example of the instantaneous circulation pattern in Sarasota Bay is shown in Figure 3. Recently,





Figure 2. Results of a three-dimensional simulation of near-surface currents and salinity distribution in the James River estuary. While flow in the upstream portion is still ebbing, flow near the entrance to the Chesapeake Bay already is flooding, thus creating a front with flow convergence and a sharp salinity gradient.

Figure 3. Simulated instantaneous (ebb cycle) near-surface currents and water level in Sarasota Bay, Florida. Strong currents are found in the tidal inlets, and a null zone is found between two major inlets.



the Sarasota Bay model was extended to include the neighboring Tampa Bay (Figure 4), where a large amount of field data have been collected by the National Ocean Service. In general, the CH3D model gives quite accurate estimations of water level, current, and salinity in Tampa Bay and Sarasota Bay over time periods of one to two months. Based on comparisons in terms of long-term root mean square statistics, peak amplitudes, and phases, water level can be simulated to within 10 percent error, and currents can be simulated to within 20 percent error. Accuracy of the model is improved significantly when a curvilinear grid is used to resolve accurately the complex shoreline and internal features, such as navigation channels.

Applications of models of estuarine circulation and transport

After being calibrated and validated, these models will provide hydrodynamic information required by water quality or oil spill models to be used by resource management agencies. For exam-

ple, the Chesapeake Bay model is being used by the U.S. EPA Chesapeake Bay program to develop strategies to regulate the loading of nutrients from tributaries. The Tampa Bay circulation model was coupled recently to a three-dimensional Lagrangian particle trajectory model to simulate the movement of an oil spill in Tampa Bay which took place in June 1991. Due to the accurately simulated flow field, results of the particle trajectory model agreed quite well with the observed oil spill movement in Tampa Bay.⁴ In another example, the circulation and salinity transport model was used to estimate the potential impact of the proposed opening of the Midnight Pass, a closed tidal inlet, on the water quality of the Little Sarasota Bay. Using the numerical model, the flushing times in various parts of the bay were computed for two conditions: the present condition with the Midnight Pass closed and the proposed condition with the Midnight Pass opened. Results of the simulation suggested that if the Midnight Pass were opened, water quality in the majority of Little Sarasota Bay would be improved significantly. However, water quality in certain parts of Little Sarasota Bay would deteriorate due to the creation of a null zone, a region with very little flow.

Circulation and transport in lakes

Due to increased loading of nutrients, such as phosphorous and nitrogen, many of the world's lakes have experienced increased algal growth, reduced oxygen levels, and increased fish kill. These problems are affected strongly by lake circulation, which is primarily wind-driven and influenced by thermal stratification. A comprehensive field and modeling study of the wind-driven circulation, wave climate, sediment transport, and phosphorous dynamics recently was completed for Lake Okeechobee, a large shallow (2-3 m average depth) lake in South Florida. Long term data on wind, wave, current, temperature, and suspended sediment concentration were measured with a 2 Hz sampling frequency over long time periods in 1988 and 1989. Phosphorous concentrations in the bottom sediment and water column were measured less frequently. The results indicated that bottom sediment is fluidized readily and resuspended by the diurnal wind, which generally produces significant wave conditions in the large lake. Large-scale wind-driven currents produce horizontal advection and sufficient shear to cause vertical mixing of the suspended sediments. During periods of strong winds and waves, sediments are resuspended readily, leading to increased phosphorous concentrations in the water column.

Numerical model of wind-driven circulation and transport in lakes

As a result of the Lake Okeechobee study, a number of robust models were developed. These include a rectangular-grid three-dimensional circulation/sediment/phosphorous model,⁵ a curvilinear-grid three-dimensional circulation model,⁶ and



Global climate change modeling in Australia

Bryant J. McAvaney, Australian Bureau of Meteorology Research Centre, Melbourne, Australia

Australian atmospheric scientists long have been active in numerical weather prediction and climate modeling. The majority of climate models throughout the world use the spectral modeling technique, which was first developed in Australia in the early 1970s¹ and was applied successfully to climate simulation for the first time in 1978.² After the Australian Bureau of Meteorology formed the Bureau of Meteorology Research Center (BMRC) in 1985, a number of the scientists involved in the earlier work joined this new organization. The Bureau of Meteorology subsequently acquired an in-house supercomputing facility which includes a CRAY X-MP/14se system and a CRAY Y-MP 2E system. Building on the vectorizing capability developed at other computing sites, scientists within the BMRC quickly harnessed this array of computing power for numerical weather prediction and climate research.

In 1989, the Australian government funded additional research into the potential impact of enhanced greenhouse gas concentration. Since then, an accelerated program of climate model development and experimentation has ensued at the BMRC. The ever increasing demands of climate modeling research have meant that computing resources within the Bureau of Meteorology have had to be supplemented with Cray Research systems in the United States: a CRAY-2 system at Lawrence Livermore National Laboratory and a

CRAY Y-MP/2 system sponsored by the Model Evaluation Consortium for Climate Assessment (MECCA) and operated by the National Center for Atmospheric Research (NCAR).

The evolution of climate change modeling at BMRC

The global numerical modeling system, Global Assimilation and Prediction (GASP), used operationally by the National Meteorological Center, Melbourne, was the basis of the current climate change model. In a broad definition, the model used for climate change research is simply a lower horizontal resolution R21 (approximately 5.6° longitude by 3.3° latitude) and a lower vertical resolution (nine levels) version of the R31L19 model (approximately 3.75° longitude by 2.25° latitude and 19 levels) used for numerical weather prediction research.

Our current climate model development fails to take into account possible changes in both heat storage in the deep ocean and in oceanic heat transport that would be created by climatic changes. In response, we have extended our modeling effort to include a coarse resolution global general circulation version of the Modular Ocean Model originally described by Lau et al.³ We are in the process of synchronously coupling this ocean model with the atmospheric model and aim to

begin experiments with this coupled system by late 1992.

The experiments

The main concerns in modeling climate change are the accurate simulation of the present climate and, simultaneously, the simulation of any climate change in a scientifically supportable manner. We can validate any simulation of the present climate against observations, but we have greater difficulty validating simulations of the individual physical processes that shape climates. We have a still harder task validating a climate change simulation.

At BMRC we have concentrated on validating the climate simulations produced by the model against the available observations. Increasingly, we have been involved in various international exercises in which we compare the ability of various models to simulate the present climate and compare the individual feedback processes in models.

Feedback intercomparison experiments

BMRC promotes a range of experiments following the methodology for understanding and comparing atmospheric climate feedback processes in general circulation models, which originally was proposed by Cess and Potter.⁴ The experiments have been simplified to maximize participation from modeling groups. BMRC has joined a group of climate modelers working under the acronym FANGIO (Feedback Analysis for GCM Intercomparison and Observation) who have compared feedbacks due to clouds and snow and have compared climatic forcing fields (for example, doubling of CO₂).

In a standard climate change experiment, some type of external perturbation is made to the climate system, such as changes in solar flux or atmospheric CO₂ content. This perturbation produces an imbalance ΔQ in the net radiation at the top of the atmosphere. To restore equilibrium, where the net downward radiative flux at the top of the atmosphere Q returns to zero, the variables of the climate system, including the surface temperature T_s , will change. We then can define a climate sensitivity parameter λ through the expression $\Delta T_s = \lambda \Delta Q$. As FANGIO experimenters, we specify the climate response by perturbing the sea surface temperature by ± 2 K (ΔT_s) and use the model to determine the magnitude of the forcing that would be in balance with such a response. With this specification all models yield approximately the same mean global surface temperature difference, 4 K, so the differences in the ΔQ with each model imply differing climate sensitivities.

In an experiment designed to compare the cloud feedback, perpetual July conditions were imposed with both snow and sea ice cover held fixed to suppress any surface albedo feedback. As part of the experiment, a sensitivity parameter was computed for all cloud-free points. For the BMRC model, the cloudy climate sensitivity parameter was $0.54 \text{ K m}^2 \text{ W}^{-1}$, whereas for clear conditions the sensitivity parameter was $0.52 \text{ K m}^2 \text{ W}^{-1}$. The slightly

greater value for cloudy conditions implies a very small positive feedback due to clouds. Figure 1 shows the clear-sky-sensitivity parameter, which differs from the cloudy sensitivity parameter by a factor of almost three.⁵ This wide variation is due primarily to differences in the cloud feedback processes in the various models.

Another experiment was designed to enable comparison of the feedback due to snow. The same SST perturbation was used, but this time perpetual April conditions were chosen, and two sets of climate change simulations were performed: one for which the snow was permitted to vary for the $-2 \text{ K } \Delta \text{SST}$ to $+2 \text{ K } \Delta \text{SST}$ climate change, and the other where the snow condition was held fixed at that for the $-2 \text{ K } \Delta \text{SST}$ simulation. By analyzing the difference between the climate sensitivity parameter in these two cases, a measure of snow feedback can be determined.⁶ The BMRC model produced a moderate positive snow feedback. Further analysis has shown that snow feedback involves two distinct climate responses: a direct effect corresponding to the classical snow albedo effect in the visible part of the spectrum and indirect effects due to cloud interactions and feedback in the infrared.

The direct radiative perturbation caused by an instantaneous increase in CO₂ acts as a forcing of the climate system. The climate system responds by changing the climate to restore radiation balance. The details of the response are governed by the feedback mechanisms which can vary widely between models. It also is important to compare the CO₂ radiative forcing between models. Another experiment has been conducted recently by FANGIO participants to enable such comparison. The BMRC model produced a 2^*CO_2 radiative forcing at 200 hpa of 4.1 W m^{-2} ; the spread of results among the various models is shown in Figure 2. It now has been demonstrated that this variation is due mostly to differences in the long-wave radiation codes in the models. A direct com-




Figure 1. The clear-sky and cloudy-sky climate sensitivity parameter for 19 general circulation models.

begin experiments with this coupled system by late 1992.

The experiments

The main concerns in modeling climate change are the accurate simulation of the present climate and, simultaneously, the simulation of any climate change in a scientifically supportable manner. We can validate any simulation of the present climate against observations, but we have greater difficulty validating simulations of the individual physical processes that shape climates. We have a still harder task validating a climate change simulation.

At BMRC we have concentrated on validating the climate simulations produced by the model against the available observations. Increasingly, we have been involved in various international exercises in which we compare the ability of various models to simulate the present climate and compare the individual feedback processes in models.

Feedback intercomparison experiments

BMRC promotes a range of experiments following the methodology for understanding and comparing atmospheric climate feedback processes in general circulation models, which originally was proposed by Cess and Potter.⁴ The experiments have been simplified to maximize participation from modeling groups. BMRC has joined a group of climate modelers working under the acronym FANGIO (Feedback Analysis for GCM Intercomparison and Observation) who have compared feedbacks due to clouds and snow and have compared climatic forcing fields (for example, doubling of CO_2).

In a standard climate change experiment, some type of external perturbation is made to the climate system, such as changes in solar flux or atmospheric CO_2 content. This perturbation produces an imbalance ΔQ in the net radiation at the top of the atmosphere. To restore equilibrium, where the net downward radiative flux at the top of the atmosphere Q returns to zero, the variables of the climate system, including the surface temperature T_s , will change. We then can define a climate sensitivity parameter λ through the expression $\Delta T_s = \lambda \Delta Q$. As FANGIO experimenters, we specify the climate response by perturbing the sea surface temperature by $\pm 2 \text{ K}$ (ΔT_s) and use the model to determine the magnitude of the forcing that would be in balance with such a response. With this specification all models yield approximately the same mean global surface temperature difference, 4 K , so the differences in the ΔQ with each model imply differing climate sensitivities.

In an experiment designed to compare the cloud feedback, perpetual July conditions were imposed with both snow and sea ice cover held fixed to suppress any surface albedo feedback. As part of the experiment, a sensitivity parameter was computed for all cloud-free points. For the BMRC model, the cloudy climate sensitivity parameter was $0.54 \text{ K m}^2 \text{ W}^{-1}$, whereas for clear conditions the sensitivity parameter was $0.52 \text{ K m}^2 \text{ W}^{-1}$. The slightly




Figure 1. The clear-sky and cloudy-sky climate sensitivity parameter for 19 general circulation models.

greater value for cloudy conditions implies a very small positive feedback due to clouds. Figure 1 shows the clear-sky sensitivity parameter, which differs from the cloudy sensitivity parameter by a factor of almost three.⁵ This wide variation is due primarily to differences in the cloud feedback processes in the various models.

Another experiment was designed to enable comparison of the feedback due to snow. The same SST perturbation was used, but this time perpetual April conditions were chosen, and two sets of climate change simulations were performed: one for which the snow was permitted to vary for the $-2 \text{ K } \Delta \text{SST}$ to $+2 \text{ K } \Delta \text{SST}$ climate change, and the other where the snow condition was held fixed at that for the $-2 \text{ K } \Delta \text{SST}$ simulation. By analyzing the difference between the climate sensitivity parameter in these two cases, a measure of snow feedback can be determined.⁶ The BMRC model produced a moderate positive snow feedback. Further analysis has shown that snow feedback involves two distinct climate responses: a direct effect corresponding to the classical snow albedo effect in the visible part of the spectrum and indirect effects due to cloud interactions and feedback in the infrared.

The direct radiative perturbation caused by an instantaneous increase in CO_2 acts as a forcing of the climate system. The climate system responds by changing the climate to restore radiation balance. The details of the response are governed by the feedback mechanisms which can vary widely between models. It also is important to compare the CO_2 radiative forcing between models. Another experiment has been conducted recently by FANGIO participants to enable such comparison. The BMRC model produced a 2^*CO_2 radiative forcing at 200 hpa of 4.1 W m^{-2} ; the spread of results among the various models is shown in Figure 2. It now has been demonstrated that this variation is due mostly to differences in the long-wave radiation codes in the models. A direct com-




Figure 2. The double CO₂ radiative forcing (W m⁻²) at the "tropopause" (defined to be 200 hpa here) for 11 general circulation models.

parison of the 2*CO₂ long-wave forcing profile using the mean temperature and moisture profile has been made against detailed line-by-line radiative transfer calculations. The BMRC profile generally is within 5 percent of the line-by-line calculations. Improvements in the radiative code in the BMRC model already are being tested. These improvements should reduce further this relatively small discrepancy.

Validation of the model

In any climate simulation experiment, an enormous amount of information can be produced, and validation can proceed in many ways. Because the greatest amount of observed data either is obtained through the use of satellites looking at the earth from outside its atmosphere or from measurements made at ground level, climate modelers often choose to validate fields at the top of the atmosphere or at the ground.

To illustrate the capabilities of the BMRC model, we have chosen a ground-based form of validation that uses two fields, annual precipitation and annual biotemperature, to calculate a "life zone" based on the Henderson-Sellers simplification of the Holdridge life zone classification system.⁷ The biotemperature is defined as the sum of the surface temperatures during a year, with each temperature value (daily, weekly, monthly, or seasonal) set to 0 °C if it is less than or equal to 0 °C. The sum then is divided by the total number of values, for example, 12 for monthly temperatures. The average biotemperature represents conditions that are important during the growth season of plants.

Henderson-Sellers has proposed a generalized nine-class life zone scheme that produces a reasonable representation of present-day ecosystem types using global datasets of observed monthly mean surface air temperatures and monthly mean total precipitation. The life zone groups are as shown in Figure 3 calculated on a scale grid similar to that used in the climate model simulation (2.5°

longitude by 2.5° latitude). This coarse grained distribution of life zone groups looks intuitively reasonable and corresponds quite well to the distribution of biomes identified by Emanuel.⁸ The life zones computed from the simulated surface air temperature and precipitation produced by the BMRC climate model (including a mixed layer ocean) with current levels of carbon dioxide (330 ppm) are shown in Figure 4. The simulation has produced a reasonable representation of present-day ecosystem types. The dry areas, especially the *hot desert* and *other desert areas* are well distributed, but there are some problems. For example, the *rainforest* region in South America is restricted to the west coast, and *tundra* has appeared in the Himalayas. However, because a highly simplistic land surface parameterization package has been used, the simulation is extremely encouraging.

Viewed from the top of the atmosphere, one of the many fields we use for validation is the outward long-wave radiation (OLR), which is measured readily by satellite instrumentation. Figure 5(a) shows the observed OLR pattern for the month of January as determined from 10 years of satellite measurements. This includes the OLR values as they respond to cold temperatures. These values typically are related to the emission from high cirrus clouds in the tropics, which correspond to very deep convective activity. Higher values are associated with long wave emission from the surface and low clouds at higher temperatures. Figure 5(b) shows the ensemble average for January from 15 years of model simulation. In general, the model has performed reasonably well but still has flaws, as evidenced by the extremely low values in the area of the Philippines and the slightly more extensive areas of high emission. Discrepancies such as these spark further model development and improvement.

Equilibrium 2*CO₂ experiments

Although coupled ocean-atmosphere models are being used more frequently for enhanced greenhouse gas concentration simulations, doubled CO₂ experiments with atmospheric models coupled to mixed-layer ocean models continue to be of interest. The computer time to arrive at equilibrium is much less than for a coupled ocean-atmosphere model; therefore, model sensitivity studies are more feasible.

We have used the BMRC atmosphere model together with a mixed-layer ocean model to simulate the equilibrium climate resulting from an instantaneous doubling of CO₂ from 330 ppm to 660 ppm. The simulated globally averaged increase in surface temperature was 2.1 K accompanied by a 3 percent increase in precipitation. These simulated increases are less than in most other models. The difference in sensitivity may be related to the relatively small positive feedback due to clouds as discussed earlier.

Accompanying the simulated change in surface temperature and precipitation are changes in the life zones. The computed distribution of life zones for the 2*CO₂ simulation looks somewhat similar to that simulated for the present day




Figure 3. The "life zones" computed from fields of the observed monthly mean surface air temperature and precipitation.




Figure 4. The "life zones" computed from the simulated monthly mean surface air temperature and precipitation fields produced by the BMRC climate model (with a mixed layer ocean) for current conditions.

climate (Figure 4), with minor changes in area coverage for most classifications. The decrease in the area of *tundra* is the most obvious change. The slight changes in life zone class are due to changes in both precipitation and temperature.

Modeling a coupled ocean-atmosphere-cryosphere-biosphere

As noted earlier, we are in the process of constructing and testing a coupled ocean-atmosphere model. However, such a model would

remain incomplete because the representation of land surface processes in our current generation model is simple and represents only the gross characteristics of the interaction between the soil-snow-atmosphere system. As a result of collaboration with scientists from Macquarie University, a more complex land surface scheme incorporating vegetation—the Bare Essentials of Surface Transfer (BEST)⁹—has been incorporated into the BMRC atmospheric model. Results from the preliminary experiments have been extremely encouraging. Further refinements and more extensive analysis




Figure 5. (a) The observed outward long-wave radiation (OLR) pattern (W m^{-2}) for the month of January derived from the Earth Radiation Budget Experiment (ERBE) and (b) the ensemble average for January from 15 years of model simulation with the BMRC model.

of the results are expected to position BEST as a major link in a coupled ocean-atmosphere-biosphere model.

Another important component of a more comprehensive climate model is a sea-ice model. The BMRC model currently has a simple thermodynamic sea-ice model with no allowance for leads. It has been suggested that inclusion of sea-ice dynamics is required to model the surface boundary conditions over the high latitude oceans satisfactorily. The addition of a dynamic-thermodynamic sea-ice model to the BMRC model is planned in the near future. However, it is recognized that the sensitivity of such a model to the simulation of the low level wind structure may require subsequent improvements in the parameterization of boundary layer processes in the polar regions. ■

Acknowledgments

The author thanks fellow members of the BMRC greenhouse modeling group who have participated in the

work described in this paper and James Fraser who programmed the life zone computation and prepared the color diagrams.

About the author

Bryant J. McAvaney is a principal research scientist with the Bureau of Meteorology Research Centre within the Australian Bureau of Meteorology. He has had a long career in the application of supercomputers to climate modeling. He currently leads a research group devoted to modeling climate change, especially the impact of an enhanced greenhouse effect.

References

1. Bourke, W. P., "An Efficient, One-Level, Primitive Equation Spectral Model," *Monthly Weather Review*, Vol. 100, pp. 683-689, 1972.
2. McAvaney, B. J., W. P. Bourke, and K. K. Purl, "A Global Spectral Model for Simulation of the General Circulation," *Journal of Atmospheric Science*, Vol. 35, pp. 1557-1583, 1978.
3. Lau, N.-G., R. C. Philander, and M. J. Nath, "Simulation of ENSO-like Phenomena with a Low-Resolution Coupled GCM of the Global Ocean and Atmosphere," *Journal of Climate*, Vol. 5, pp. 284-307, 1992.
4. Cess, R. D. and G. L. Potter, "A Methodology for Understanding and Intercomparing Atmospheric Climate Feedback Processes in General Circulation Models," *Journal of Geophysical Research*, Vol. 93, pp. 8305-8314, 1988.
5. Cess, R. D., G. L. Potter, J. P. Blanchet, G. J. Boer, S. J. Ghan, J. Hansen, J. T. Kiehl, H. Le Treut, Z.-X. Li, X.-Z. Liang, B. J. McAvaney, V. P. Meleshko, J. F. B. Mitchell, J.-J. Morcrette, D. A. Randall, L. J. Rikus, E. Roeckner, U. Schlese, D. A. Sheinin, A. Slingo, A. P. Sokolov, K. E. Taylor, W. M. Washington, R. T. Wetherald, I. Yagaii, "Intercomparison and Interpretation of Cloud-Climate Feedback Processes in Nineteen Atmospheric General Circulation Models," *Journal of Geophysical Research*, Vol. 95, 16,601-16,615, 1990.
6. Cess, R. D., G. L. Potter, J. P. Blanchet, G. J. Boer, R. Colman, S. J. Ghan, J. Hansen, J. T. Kiehl, H. Le Treut, Z.-X. Li, X.-Z. Liang, B. J. McAvaney, V. P. Meleshko, J. F. B. Mitchell, J.-J. Morcrette, D. A. Randall, L. J. Rikus, E. Roeckner, U. Schlese, D. A. Sheinin, A. Slingo, A. P. Sokolov, K. E. Taylor, W. M. Washington, R. T. Wetherald, I. Yagaii, "Intercomparison and Interpretation of Snow-Climate Feedback Processes in Seventeen Atmospheric General Circulation Models," *Science*, Vol. 253, pp. 888-892, 1991.
7. Henderson-Sellers, A., "Predicting generalized ecosystem groups with the NCAR CCM: First steps towards an interactive biosphere," *Journal of Climate*, Vol. 9, pp. 917-940, 1990.
8. Emanuel, W. R., H. H. Shugart, and M. P. Stevenson, "Climatic Change and the Broad Scale Distribution of Terrestrial Ecosystem Complexes," *Climatic Change*, Vol. 7, pp. 29-44, 1985.
9. Pitman, A. J., Z.-L. Yang, J. G. Cogley, and A. Henderson-Sellers, "Description of Bare Essentials of Surface Transfer for the BMRC Atmospheric General Circulation Model," *BMRC Research Report No. 32*, Melbourne, Australia, 1991.

Seasonal variability simulation using a general ocean circulation model

Tomonori Matsuura and Toshikazu Morohoshi, National Research Institute for Earth Science and Disaster Prevention, Tsukuba, Japan

The National Research Institute for Earth Science and Disaster Prevention (NIED) in Tsukuba, Japan, started research on global water and energy cycles in 1991 as part of a 10-year plan. The ocean study's objective is to investigate the effect of fresh water (precipitation-evaporation, sea ice, river runoff) on general ocean circulation.

To perform this research, the institute installed a Cray Research supercomputer in March 1992. The CRAY Y-MP 2E system has been operating without problems since its installation and can run high-speed image processing routines that allow real-time processing of the simulation results.

The Modular Ocean Model (MOM) software application developed by Geophysical Fluid Dynamics Laboratory (GFDL) was installed on the CRAY Y-MP 2E supercomputer, allowing NIED researchers to perform sensitivity experiments on global ocean circulation. The institute also installed global atmospheric circulation software, developed by the Japanese Meteorological Agency, which precisely models vegetation. Using MOM, researchers were able to get the following simulation results of the seasonal variability of the general ocean circulation.

The model


The computational area of the model is the entire globe. The grid size between longitude and latitude is $2^\circ \times 2^\circ$, with 15 vertical layers. Although the geographical features of the continents, islands, and bottom topography have been represented as accurately as possible, the bottom topography has been smoothed to avoid numerical instability.¹

For physical parameters, a lateral viscosity coefficient of $A_m = 1 \times 10^9 \text{ cm}^2 \text{ s}^{-1}$, a vertical viscosity coefficient of $K_m = 2 \times 10^1 \text{ cm}^2 \text{ s}^{-1}$, a lateral diffusion coefficient of $A_t = 1 \times 10^7 \text{ cm}^2 \text{ s}^{-1}$, and a vertical diffusion coefficient of $K_t = 1 \text{ cm}^2 \text{ s}^{-1}$ have been used. For wind stress the climatological values derived by Hellerman and Rosenstein were used.²


In the simulation, the model's surface temperatures and salinities are continuously adjusted toward observation (the Levitus³ data), using the restoring term in equations of both temperature and salinity conservation, in an attempt to replicate the boundary fluxes of heat and fresh water at the interface between ocean and atmosphere. The restoring parameter has a value of $1/(50 \text{ days})$ for the data averaged annually and of $1/(25 \text{ days})$ for the data averaged monthly.



a



b



c


Figure 1. Streamlines at various depths (steady state): 19 m (a), 1040 m (b), 5020 m (c).

The steady state was established by computing a time span of 2000 years using an annual mean climatological value as an external force. Then 1000 years was computed using monthly mean climatological values. The results from the 2000-year computation were used to analyze the steady state. For the seasonal variability analysis, another year was run after the 1000 years with seasonal variability, and the results for each of the 365 days were used. It took about 12 minutes of CPU time to simulate one year using both CPUs of the CRAY Y-MP 2E system.

Results of the experiments

The results of the steady-state simulation are in close agreement with those of Cox⁴ and Fujio

Figure 2. Water temperature distribution and velocity vector of the first layer in February (top) and in August (bottom).



et al.,⁵ except for the Antarctic circumpolar current, which has a larger value of 260 Sv. The streamlines in Figure 1a show a gyre field caused by the wind stress at a depth of 20 m. The streamline pattern at a depth of 1000 m agrees well with the transport function pattern (Figure 1b). At a depth of 5000 m, a deep boundary current flowing north from New Zealand can be observed (Figure 1c).

An animation of the study was produced using 62 frames from every six days for one year to see the sea surface temperature and the seasonal variability of the velocity vector of the entire ocean dynamically. Figure 2 shows one scene of the month of February (winter) and August (summer).

A comparison of the resulting sea surface temperatures in winter and summer reveals that temperatures of both the Japanese current Kuroshio and the Gulf stream as the western boundary current are much higher in the summer than in the winter. The surface current velocity vector distribution shows larger velocity at low latitudes of the western Pacific in the winter and a strong western offshore current in the Indian ocean in the summer.

Figure 3 shows the seasonal variability of the Kuroshio (30° N) and the Indonesian through-flow (0° N). Kuroshio's current velocity increases around January and June and decreases abruptly in October. The transport of the Indonesian

throughflow reaches its maximum between July and August (southbound) and becomes minimal in January. The average transport of Kuroshio (50 Sv) is about five times larger than the Indonesian throughflow (10 Sv), but their seasonal transport variability is almost the same (Figure 3b). The simulation results showed that the subtropical gyre system, including Kuroshio, affects the transport variation of the Indonesian throughflow.

The Indonesian throughflow acts as a pipe of seawater flow between the Pacific and the Indian Ocean. We think that the seasonal variability of this flow affects ENSO (El Niño and Southern Oscillation). Therefore, it is an important topic of study to analyze the cause of seasonal variability of the Indonesian throughflow.

Visualization of the results

The visualization software HISCORE produced the graphics. HISCORE is a high-speed image processing software that uses the UltraNet framebuffer. It can display a 1280 x 1024 x 3-byte image with a transfer speed of 20 scenes per second. This performance is achieved when 60 scenes are stored in the 64-Mword main memory. Using the 256-Mword SSD and disks (about 50 Gbytes), approximately 240 frames, or a few thousand frames, can be processed, respectively, although the display rate decreases.

The research objectives at NIED are to simulate potential disaster, to obviate disaster, and to cope with the Earth's environmental problems in order to respond to apprehension about abnormal phenomena associated with environmental changes such as the greenhouse effect. NIED researchers plan to develop a combined atmosphere-ocean model to obtain the hydrologic and atmospheric elements required for estimating potential disaster. NIED intends to use a CRAY Y-MP supercomputer to optimize the general atmospheric circulation model to include the hydrological process of the Earth's surface, and to optimize the ocean general circulation model to simulate the response to fresh water. ■

Acknowledgments

The authors thank R. Pacanowsky of GFDL who agreed to the laboratory's use of the MOM code. They also thank Cray Research Japan, the Department of Advanced Analysis Technology Research, and the Department of Atmospheric and Oceanic Science Technology Research of NIED for their help in using the Cray Research supercomputer and in rendering the visualization.

About the authors

Tomonori Matsuura became assistant of the Department of Technology at Ibaragi University in 1983. In 1987, he received a Ph.D. degree from Kyushu University, Japan. He became a researcher at Texas A & M University in 1988. Since 1990, Matsuura has worked as chief researcher in the Atmospheric and Oceanic Science Technology Department at the National Research Institute for Earth Science and Disaster Prevention in Tsukuba, Japan.




Figure 3. Seasonal transport variability for Kuroshio (30° N, 124° E - 144° E) and the Indonesian throughflow (0° N, 126° E - 132° E), whole transport (a) and difference from annual average (b).

Toshikazu Morohoshi has been chief researcher at the Advanced Analysis Technology Research Department of the National Research Institute for Earth Science and Disaster Prevention in Tsukuba, Japan, since 1974.

References

1. Killworth, P. D., "Topographic Instabilities in Level Model OGCMs," *Ocean Modelling*, Vol. 75, pp. 9-12, 1987.
2. Hellerman, S. and M. Rosenstein, "Normal Monthly Wind Stress over the World Ocean with Error Estimates," *Journal of Physical Oceanography*, Vol. 13, pp. 1093-1104, 1983.
3. Levitus, S., "Climatological Atlas of the World Ocean," NOAA Professional Paper 13, 173 pages, U.S. Government Printing Office, Washington, D.C., 1982.
4. Cox, M. D., "A Baroclinic Numerical Model of the World Ocean: Preliminary Results," in *Numerical Model of Ocean Circulation*, R. O. Reid, ed., pp. 107-120, National Academy Press, 1975.
5. Fujio, S., T. Kadowaki, and N. Imasato, "World Ocean Circulation Diagnostically Derived from Hydrographic and Wind Stress Fields," Part I., "The Velocity Field," *Journal of Geophysical Research*, Vol. 97, pp. 11,163-11,176, 1992.

Optimization of shallow sea models

Arnold Heemink, Tidal Waters Division, Rijkswaterstaat, The Netherlands
Theo van Stijn, Informatica Centrum voor Infrastructuur en Milieu (ICIM),
Rijswijk, The Netherlands



Studying the impact of water on the environment has been a geographical necessity—and often a question of survival—for the Netherlands throughout its history. As witnessed by a particularly devastating storm in 1953, the delta area in the southwestern part of the Netherlands that lies below or just above sea level is particularly vulnerable (Figure 1). Accurate water level forecasts for the entire Dutch coast are vital to determine the precautionary actions that should be taken prior to a storm surge to protect the dikes and whether the storm surge barriers in the Eastern Scheldt estuary and in the New Waterway to Rotterdam harbor should be closed.

To improve the accuracy of water level forecasts in the Netherlands, the Dutch Water Management Agency's (Rijkswaterstaat) Tidal Waters Division, the Computing and Information Center for Civil Engineering and the Environment (ICIM, Informatica Centrum voor Infrastructuur en Milieu), and the Royal Meteorological Institute (KNMI) are cooperating to develop a new storm surge forecasting system. Operation of the very large storm surge barrier now under construction in the New Waterway to Rotterdam (Figure 2) will depend directly on the precision and timeliness of the storm forecasting system's output. The new storm surge forecasting system is based on a detailed numerical model that describes the tidal flow on the entire Northwest European Continental Shelf. The meteorological input for this hydro-

Figure 1. The disaster of February 1, 1953.

dynamic model is determined using a numerical model of the atmosphere. The resulting system is not perfect, however, as errors are introduced by inaccuracies in the meteorological input and by uncertain parameters in the model. In addition, considerable uncertainty is associated with the boundary conditions along the Atlantic Ocean. Therefore, an algorithm has been developed for the optimization of the uncertain parameters in the models using the available measurements. A data assimilation scheme also is being implemented to incorporate online measurements of the water level into the hydrodynamic model, thereby improving the water level forecast on a routine basis.

The storm surge forecasting system is run routinely four times a day on a minisupercomputer. However, to solve the time-consuming parameter estimation problem and to determine an optimal data assimilation scheme—a problem that requires even more computational effort—the forecasting system also is run on the CRAY Y-MP 2E system at ICIM.

The hydrodynamic model

The hydrodynamic model is based on shallow water equations that consist of depth-integrated motion and continuity equations written in spherical coordinates.¹ These grid-defining equations are approximated numerically using an alternating direction implicit (ADI) method. This ADI method is unconditional stable and lacks artificial viscosity. The numerical model is vectorized and computed on the CRAY Y-MP 2E system. A long vector length and a performance of approximately 40 MFLOPS are achieved by performing the calculations over both land and sea regions. The longitudinal grid spacing is one-fourth degree and the latitudinal spacing one-sixth degree. The time step for the integration is 10 minutes.

Parameter estimation

Before a tidal model can be used to predict water movement accurately, its reliability has to be determined by iteratively adjusting the model so that it reproduces a series of near-optimal observations. This so-called model calibration process, which had to be performed by experts until recently, now can be done automatically using a parameter estimation procedure developed by Ten Brummelhuis and Heemink in 1990.²

The calibration process consists of the iterative minimization of a so-called error function $I(p)$, where p denotes the unknown parameters. This function is a measure of the difference between model output and observed data. An adjoint tidal model is used to determine the gradient of $I(p)$, and a quasi-Newton method is used to find the optimal parameter values that minimize I . The parameter estimation algorithm is employed to correct the space-varying bottom friction coefficient, the uncertain parameters in the open boundary conditions, the wind stress coefficient, and even some depth values in the tidal model. To determine the 50 uncertain parameters, the algorithm requires approximately 250 simulations of either

the tidal model or its adjoint model. To take into account both spring tide and neap tide, the chosen calibration period is 28 days. The total CPU time for this massive calibration run is approximately 15 hours on the CRAY Y-MP 2E system.

Data assimilation

Data assimilation can be defined as a procedure to incorporate real-time data into a model to improve the predictions. However, incorporating real-time data into a nonlinear numerical model for storm surge forecasting is far from trivial. The most common data assimilation technique used in numerical weather prediction is optimal interpolation. However, the correction produced by optimal interpolation is not consistent with the underlying numerical model. As a consequence, when describing a tidal model's complicated flow pattern that is due to an irregular geometry, the use of optimal interpolation yields unrealistic corrections or even introduces instabilities.

A new data assimilation approach is Kalman filtering. To use a Kalman filter to incorporate data into a numerical shallow-water flow model, the tidal model is embedded into a stochastic environment by introducing noise processes. This allows taking into account the inaccuracies of the underlying deterministic model and of the boundary conditions. A Kalman filter combines the information provided by the model and the noisy measurements taken from the actual system to obtain an optimal (least squares) estimate of the state of the model. Kalman filtering is similar to optimal interpolation. Unlike optimal interpolation, however, the correction produced when using a Kalman filter is guaranteed to be consistent with the underlying model, even in the case of very irregular flow patterns, ensuring that numerical problems are avoided.

In the last decade Kalman filtering has gained acceptance as a powerful framework for data assimilation. At least for linear or weakly nonlinear problems, Kalman filtering is the optimal data assimilation method. However, the computational requirements for Kalman filtering are usually very large.

Heemink³ and Heemink and Kloosterhuis⁴ developed a suboptimal filter algorithm, in which the coefficients of the data assimilation scheme are assumed to be constant, allowing them to be computed once offline. This is the most time-consuming part of the Kalman filter algorithm, in which 50 to 1000 model simulations are required to solve the filtering problem. However, unlike the parameter estimation algorithm, these model simulations are independent of each other, allowing a definition of the vectors all across the models. In this way the model runs very efficiently on the CRAY Y-MP 2E system, achieving a performance of more than 180 MFLOPS. In addition, the useless computations over land can be avoided. Total CPU time to determine the optimal data assimilation scheme for the continental shelf model is only about three minutes. In fact, in this case it is not the CPU time but the memory requirements that become prohibitive.




Figure 2. Artist's view of the very large storm surge barrier now under construction in the New Waterway to Rotterdam harbor, one of the world's largest. Closing time for the barrier will be three hours, shutting down traffic for a minimum of six hours. The precision and timeliness of the storm forecasting system will have a direct economic impact.

Storm surge forecasting results

The performance of the storm surge forecasting system is illustrated in Figures 3 to 5. Figure 3 shows the wind-induced water level rise determined by the hydrodynamic model at midnight on February 1, 1953. In Figure 4 the water level prediction is compared with the measurements at the Dutch measurement location IJmuiden.

The storm surge of February 1989 could not be predicted very accurately by the model without data assimilation. To demonstrate the improvement of the water level predictions that can be obtained by incorporating the data into the model, the predictions of the system with and without data assimilation are compared in Figure 5. In this case

Figure 3. Wind-induced water level rise on February 1, 1953, 12:00 midnight.






Figure 4. Water level prediction at IJmuiden.






Figure 5 (below). Water level prediction at IJmuiden with and without data assimilation.



the data assimilation scheme was provided only with the data available up to 2:00 a.m. on February 14. The short-term predictions using the data are clearly much more accurate, ultimately ensuring that adequate—and life saving—measures can be taken. Of course the improvement obtained by filtering the available data decreases with the prediction interval.

The storm surge forecasting system has been used routinely since 1991. As a result, there has been a major improvement in the accuracy of the water level forecast along the entire Dutch coast. The system output also is essential for the operation of the storm surge barriers in the Eastern Scheldt and in the New Waterway.

Conclusion

In the near future the research group plans to focus on environmental problems such as pollutant dispersion in the Dutch coastal waters. These environmental simulations eventually will be run on a routine basis but will require much more detailed models with respect to hydraulics and the environment. Optimization runs for the planned environmental models will be possible only on powerful supercomputers. ■

About the authors

Arnold Heemink, head of the mathematics group of the Tidal Waters Division of Rijkswaterstaat, received a B.Sc. degree in applied mathematics in 1980 and a Ph.D. degree in 1986 from the University of Twente. Heemink develops mathematical models for shallow water flow, dispersion processes, and sediment transport. Recently he was appointed professor of applied mathematics at the Delft University of Technology.

Theo van Stijn, head of the mathematical modeling group of the Computing and Information Center for Civil Engineering and the Environment (ICIM), received an M.Sc. degree in numerical mathematics from the Delft University of Technology and a Ph.D. degree in applied mathematics from Groningen University. Van Stijn and his colleagues develop mathematical models for fluid flow problems, traffic and transportation problems, and environmental problems. The models are developed for routine use in real life situations.

References

1. Flather, R. A. and R. Proctor, "Prediction of North Sea Storm Surges using Numerical Methods," in *North Sea Dynamics*, I. Sündermann and W. Lenz, eds., Springer Verlag, Berlin, 1983.
2. Ten Brummelhuis, P. G. J. and A. W. Heemink, "Parameter Identification in Tidal Models with Uncertain Boundary Conditions," *Stochastic Hydrology and Hydraulics*, Vol. 4, pp. 193-208, 1990.
3. Heemink, A. W., "Two-dimensional Shallow Water Flow Identification," *Applied Mathematical Modelling*, Vol. 12, pp. 109-118, 1988.
4. Heemink, A. W. and H. Kloosterhuis, "Data Assimilation for Nonlinear Tidal Models," *International Journal for Numerical Methods in Fluids*, Vol. 11, pp. 1097-1112, 1990.

Supercomputing for power system analysis at Hydro-Québec

Gia Tong Vuong and Roger Chahine,
Centre d'Analyse Numérique de Réseaux, Hydro-Québec, Montréal, Canada
Luc Ostiguy and Stephen Behling, Cray Research, Inc.

Designing, installing, and operating a power transmission network and ensuring its stability and reliability are complex challenges. Each network is a dynamic system subjected to oscillations which can lead to costly equipment failures, network separation, and eventually blackouts. Each year, an average of 40 events, such as lightning bolts, ice storms, and tornadoes, disturb and threaten to disrupt power network operation in Québec. In addition, the growing demand for electrical power and the complexity of interconnected, expanding networks require prudent operational planning and the ability to predict a power system's behavior under various conditions. Testing and predicting scenarios under variable contingencies therefore have become essential operational and planning activities.




Figure 1. Hydro-Québec's transmission network.

Dedicated supercomputing for power system analysis

As a member of the Northeast Power Coordinating Council (NPCC), the Canadian public utility Hydro-Québec has extensive power exchanges with neighboring networks. Its system and operation planning studies therefore must take into account not only its own network (Figure 1) but the whole integrated NPCC network. To analyze the behavior of this network, which comprises up to 12,000 buses (electrical nodes in the transmission network), Hydro-Québec installed a Cray Research supercomputer at its computing center for power system analysis (PSA) in March 1990.

The CRAY X-MP/216 system installed at Hydro-Québec's Centre d'Analyse Numérique de Réseaux (CANR) is the only supercomputer dedicated to PSA activities to date. The PSS/E integrated software package, designed by Power Technology, Inc. to study power systems in steady-state and dynamic regimes, is the first power system analysis program to be converted and optimized to run on a Cray Research supercomputer. The optimization modifications made to the code, in a cooperative effort between Hydro-Québec and Cray Research, focused on vectorization and other performance improvements. In addition, Hydro-Québec's transient stability program (ST600) and the electromagnetic transient program (EMTP) originally developed by the Bonneville Power Authority are the other application programs ported to the Cray Research system. One particularly difficult aspect of power system simulation using vector and parallel processing is solving sparse matrix systems. A new algorithm was developed that improves vector lengths and processing speeds, and integrating the new algorithm in existing programs should further increase their performances.

CANR's supercomputer, over 80 workstations, numerous peripherals, and three interconnected high-speed Ethernet networks serve three administrative units concerned with operations and planning at Hydro-Québec. The planning division performs long-term network simulation studies; the production, transmission, and telecommunications division runs present network simulations; and the research group develops and ports software applications for power system analysis.

Network stability: a constant challenge at Hydro-Québec

As depicted in Figure 1, the Hydro-Québec network is V-shaped. Loads are concentrated in the Montréal area, located at the bottom of the V. Each of the branches, referred to as the east and west corridors, consist of longitudinal transmission systems of more than 1000 kilometers.

A vital concern in PSA activities is that this network remain in a steady-state and not become unstable. Figure 2 shows the evolution of power generation at three power houses following a three-phase fault and the loss of an important line, highlighting the importance of fault-clearing time





Figure 2. Evolution of power generation following a three-phase fault: clearing after 5 cycles (top); permanent loss of a major line (stable case) and clearing after 11 cycles (bottom); permanent loss of a major line (unstable case).

which makes the difference between stable and unstable situations. A case is considered stable if all the state variables return to their initial values or settle down to acceptable ones after a major disturbance; otherwise it is considered unstable.

The importance of transient stability simulation in power system planning is illustrated in another example in which two of three transmission lines in the west corridor are lost due to a tornado sweeping through the area. The anticipated double-contingency perturbation is very severe, and without proper corrective actions it would cause excessive power swings and the loss of the entire network. To investigate this contingency,

simulations were performed on the entire network with the system load set at 50 percent of peak load—a condition that normally prevails in October when tornadoes are most frequent. Several scenarios were simulated, leading to the design of a remedy action that consists of a forced load shedding in the Montréal area combined with a load rejection in the La Grande generation complex. The remaining power imbalance is expected to be absorbed by low-frequency load-shedding relays already installed at different points in the network.



Figure 3 shows the evolution of power flows in the corridor. The figure also shows that the power on two lines successively dropped to zero at a 0.5 cycle interval, causing power swings on the remaining line(s). Thanks to proper corrective actions, however, the system stabilized at a new equilibrium point, as illustrated by the curve corresponding to the power flow on the third line, of which the final portion is flat, reflecting a stable condition.

Operational and planning studies

To evaluate network behavior, a disruptive event is simulated under existing network conditions. As the possible scenarios are simulated, various disturbance risks at all sensitive points in the network are reviewed. These simulation studies benefit not only operation engineers who want to know the limitations of the network but also system planners who want to find optimal ways to integrate future changes into the network. In particular, evaluating the behavior of the NPCC system requires detailed representations of very large networks. On typical workstations, transient-stability studies using PSS/E require about 45 hours per simulation. With the supercomputer, CPU time is reduced to less than two hours, yielding turnover times of two to three hours depending on system loading. Moreover, many studies require approximately 100 transient-stability simulations. Without supercomputing, it would be impossible to complete these simulations within the usually expected timeframe of two to three months. Even with enough workstations available to process scenarios in parallel, important parts of the work remain sequential: analyzing the results, modifying the study parameters, and running complementary simulations.

Extreme contingencies studies

Before using the supercomputer, turn-around for this kind of simulation using ST600 was four to five hours per case. Consequently, it was common to submit many simulations concurrently to reduce idle time waiting for results. On the other hand, limited computing resources restricted investigations to the most critical contingencies, creating the risk of potentially damaging oversights. With the supercomputer, response times were reduced to between 10 and 15 minutes, making it more practical to submit simulations one by one, each being selected on the basis of preceding results. Because resuming work after an



interruption usually requires an adjustment period for the investigator to regain full control of the situation and begin useful work, continuity leads to better productivity as well. In turn, reducing the response time allows a larger number of contingencies to be considered and indirectly contributes to more reliable solutions. In responding to an urgent request for information on the network's security for a particular configuration, an engineer was able to provide answers within two days compared to the two to three weeks required previously.

Statistical transient simulation

In 1989, Hydro-Québec decided to strengthen its system by means of series compen-

sation. Because this action would affect transient recovery voltages all over the system, it was decided to reevaluate rating requirements for all 735 kilovolt circuit breakers.

The first phase, which simulated the east corridor linking the Churchill Falls complex to Montréal, began in September 1990 and lasted until April 1991. Studies were done on Hydro-Québec's transient network analyzer (TNA) and included three months to set up the simulator, two months for preliminary parametric studies, and four additional months for statistical studies. One hundred simulations of breaker openings at random timing were performed for each case.

At the beginning of 1991, when studies for the second phase covering the west corridor that connects the James Bay complex to Montréal were planned, the TNA was fully booked. Waiting for the TNA to be free would mean a delay of six to nine months. Therefore, simulations with EMTP were adopted as an alternative solution. Because EMTP has not yet been optimized for vector processing, speedup ratios were not the main concern. However, supercomputing provided a solution to a critical situation, and the feasibility of simulation for a type of study that had belonged exclusively to the realm of TNAs was established for the first time.

The pioneering work done by Hydro-Québec in using supercomputing in its PSA activities is a significant step toward enhancing the stability and reliability of large and complex power transmission networks. As changes to these interconnected networks occur faster, and as their size and complexity increase, the time-critical power and speed of supercomputers become indispensable. ▀

About the authors

Roger Chahine received a B.S. degree in computer science from the Ecole Supérieure, Beirut, in 1977 and a B.S. degree in electrical engineering from the University of Québec at Trois Rivières in 1981. At Hydro-Québec since 1982, Chahine works in the planning department, where he manages the power system analysis group and the CANR computer center.

Gia Tong Vuong received B.S. and M.S. degrees from the Ecole Polytechnique, Montréal, in 1966 and 1968 respectively. He received a Ph.D. degree from INRS-Energie, University of Québec, in 1978. Having joined Hydro-Québec in 1969, Vuong now is senior engineer with the planning department in the power systems analysis group.

Luc Ostiguy received a B.S. degree in physics from the University of Montréal in 1983 and an M.S. degree in atmospheric physics from the University of Québec at Montréal in 1989. Since 1990, Ostiguy has worked for Cray Research (Canada) Inc. as an applications systems analyst.

Stephen Behling received a B.S. degree in nuclear engineering from the University of Wisconsin-Madison in 1974. From 1975 through 1986 he specialized in computer applications for nuclear reactor safety at the Idaho National Engineering Laboratory. Behling manages the industry applications group at Cray Research.

Figure 3. Simulating the effect of a tornado, evolution of line flows in the west corridor.

Simulating the birth of a star

A. H. Nelson and A. P. Whitworth
University of Wales College of Cardiff, Cardiff, Wales

The Milky Way Galaxy, in which our solar system is located, is just one of many flat disc galaxies in the universe. The disc consists of stars—seen as the Milky Way spanning the sky from horizon to horizon on dark nights—and interstellar gas, which fills the space between the stars.

Our sun was formed from such gas nearly four billion years ago, and today stellar systems still are being formed in the remote spiral arms of the Milky Way. Observations of young stars in both our own galaxy and in other flat disc galaxies show that stars are forming continuously from the interstellar gas, often in a spiral pattern. Tenuous in nature, the interstellar gas is not spread evenly in the galaxy but clumped with large clouds of molecular gas scattered in the galactic disc. Stars form in these clouds when the density of the clumps becomes high enough that the self-gravitational field of the gas overcomes the internal pressure, leading to rapid collapse. The collapse continues until the density and temperature at the center of the protostar initiate nuclear fusion and internal pressure halts the collapse—thus is a star born.

Despite decades of observational and theoretical research, the exact details of how sufficient mass is collected in a single clump or cloud remained unknown. Recently a group of astrophysicists at the University of Wales College of Cardiff, led by the authors, made a significant breakthrough in simulating the process of star formation using a CRAY X-MP supercomputer at the Rutherford Laboratory of the Science and Engineering Research Council (SERC) in the U.K.

Colliding clouds

The computer model replicated the collision of two interstellar clouds—a process thought to occur in the interstellar medium, specifically in the spiral arms of galaxies. When two clouds collide at supersonic velocities, a shocked compression region occurs in the collision interface, leading to very high gas densities and gravitational collapse to protostellar densities.

(Left) Optical image of the Whirlpool Galaxy. Star formation is widespread over the disc and is particularly enhanced in the spiral arms. The disc is approximately 60,000 light years across.

Such a collision process has been modeled numerically before but with only limited success. Almost all previous simulations have been based on fixed grid methods limited to changes of scale of less than 100. Yet the material density from interstellar to stellar conditions increases by 20 orders of magnitude during star formation. Hence the scale of the cloud must decrease by a factor of just over 10^6 , or one million, to get to stellar densities.

The star formation collapse calculations performed by the Cardiff group simulated a decrease in scale of a factor of 1000 and a density increase of 10 orders of magnitude, one million times greater than any previous calculation. The Cardiff group achieved a much larger collapse factor and a more faithful representation of the physics of the final protostellar objects by combining two relatively new and very powerful numerical techniques—neither of which involves a fixed grid.

Smoothed particle hydrodynamics, in which gas elements are represented by particles that interact with each other by a short-range repulsive force, is a technique that mimics the effects of gas pressure. Tree-code gravity obtains the gravitational force on any one particle by summing the individual interactions with nearby particles, while distant clusters of particles are lumped together to approximate their combined force.


Together these new techniques provide a very high dynamic range of spatial resolution. To date, most runs have been performed using 4000 gas particles. A few runs were executed that used 50,000 particles, and one run has been performed with 200,000 particles on the newly delivered CRAY Y-MP system at Rutherford. Runs with larger numbers of particles confirm the lower-resolution 4000-particle results and enable simulations of the formation of several stars on the same molecular cloud collision event.

A protostar is born

Figure 1 depicts the time sequence of the head-on collision of two spherical, initially uniform, interstellar clouds. The high density region in the collision interface is the scene of the star formation, but the protostar itself, a very small object at the center of the clouds, cannot be seen on this scale.

In Figure 2 the velocities of the particles in the protostar are projected onto the three orthogonal planes, demonstrating that it is a spinning disc. The origin of the spin is connected with the noise in the original random distribution of particles in the initial clouds, and it mimics the spin known to exist in the interstellar gas.

The collapse has halted because the centrifugal forces of the spinning disc are counteracting gravity and preventing further collapse. This is exactly what is expected to happen to protostellar objects, which in order to collapse further have to get rid of their spin—probably by torques due to the magnetic field expected to thread the protostars. All stars, including the sun, have a magnetic field, and the twisting effect of the field lines extending out from the protostar into the surrounding inter-







Figure 2 (left). Velocity vectors for each particle in the high density object, plotted at the particle position in the corresponding plane.

Figure 3 (right). Particle positions, color coded by density, for the near final state after a glancing collision. The density increases in the sequence blue, green, red, white. The two objects form a binary protostellar system.



form a binary star system with respective masses twelve and six times the solar mass.

Conclusion

These new simulation techniques running on Cray Research supercomputers not only yield some of the most spectacular examples of spatial dynamic range in computational physics to date, but also will help astrophysicists understand better the statistical distribution of stellar masses from cloud-to-cloud collisions and other possible scenarios for forming stars. This knowledge will lead to an estimate of the rate at which mass is transferred from the gas to stars. The rate at which the interstellar gas mass is transferred to stars by star formation (together with the inverse processes in which stellar material is ejected back into the interstellar medium by stellar winds or supernovae) is a very important quantity in determining the overall evolution of a galaxy. These calculations will lead to a deeper understanding of both the evolution of the universe of galaxies as a whole, and of the eventual fate of our own galaxy. ■

About the authors

Alistair Nelson is an astrophysicist at the University of Wales College of Cardiff. He received a Ph.D. degree in plasma physics from Imperial College, London, and works on the computer simulation of galaxies and star forming regions.

Anthony Whitworth is an astrophysicist at the University of Wales College of Cardiff. He received a Ph.D. degree in the physics of the interstellar medium from Manchester University. His research interests include the interstellar medium and star forming regions.

stellar medium is almost certain to account for the slowing down of the spin and the subsequent continued collapse to stellar densities. However, it is known from the highly bound state of the high density object in this calculation that it is certain to form a star eventually. Its mass from the bound mass at the end of the calculation is estimated to be 16 times that of the sun.

When the collision of the two clouds is not a head-on, but a glancing, collision, then two collapsed objects are obtained, and sometimes more due to fragmentation of the compressed layer. Figure 3 shows two protostellar objects formed in such a way. They are bound together to

CORPORATE REGISTER

Demand continues for high-end systems

Cray Research announced in December that the European Centre for Medium-Range Weather Forecasts (ECMWF), Reading, England, had installed and formally accepted a CRAY Y-MP C90 system, marking the first European installation and the first installation outside the United States of the high-end supercomputer. Initial CRAY Y-MP C90 system benchmarks showed that ECMWF's 10-day weather forecasting model would achieve a five-fold speedup over the performance of the eight-processor CRAY Y-MP system it is replacing. With the improved performance, ECMWF scientists will be able to refine model resolution, conduct more runs, and research new techniques to further improve forecasting accuracy.

In September, the National Aeronautics and Space Administration (NASA) announced that it had selected Cray Research for negotiations leading to the award of a contract for a new supercomputer for NASA's Ames Research Center, Mountain View, California. "Cray is an acknowledged pioneer in the supercomputer field," said Kristin Hessenius, Director of Aeronautical Research at NASA Headquarters, Washington, D.C. "They were chosen because of the clearly superior capabilities of the CRAY Y-MP C90 and Cray's history of providing NASA with reliable, cost-effective facilities over the years."

In September, the U.S. Department of Energy's Knolls Atomic Power Laboratory (KAPL), Schenectady, New York, ordered a CRAY Y-MP C90 supercomputer. KAPL will use the CRAY Y-MP C90 system to provide design and engineering support for the Naval Nuclear Propulsion Program.

Lawrence Livermore National Laboratory's National Energy Research Supercomputer Center (NERSC), Livermore, California, accepted a CRAY Y-MP C90 system in October. The system will be used for energy science applications.

Germany's Federal Waterways Engineering and Research Institute (Bundesanstalt für Wasserbau), based in Karls-

ruhe, ordered a CRAY Y-MP 2E system in October. The new system will be used to model estuaries and shorelines throughout Germany.

In October, the **Compagnie Générale de Géophysique (CGG)** installed and accepted a CRAY Y-MP 4E system at its London facility, located in the suburb of Feltham. The system will be used for CGG's North Sea business and UK customers.

Mobil Corporation ordered and installed a CRAY Y-MP 4E system in October at the Mobil Exploration and Producing Technical Center, Mobil's technical computing facility in Dallas. The supercomputer will be used for seismic processing related to Mobil's exploration business.

In October, **Amoco Oil Company** and **Landmark Graphics Corporation** each installed a CRAY Y-MP EL system. Both companies' systems are housed in Calgary, Alberta. Amoco's new system will be used for seismic processing involving the complex foothills geology of Northeastern British Columbia. Landmark, which provides software and turnkey solutions for geophysicists, geologists, and petrophysicists, will use its system for internal research and development involving 3-D seismic migration capabilities for the petroleum industry.

Tokyo-based **Mitsui Construction Company** installed a CRAY Y-MP EL system in October. Mitsui's new Cray Research system will be installed at the company's information systems facility, where it will be used for advanced technical analysis of structural and civil engineering problems related to the construction of buildings and other structures.

In December, **The Academy of Sciences** in Prague, the Czech and Slovak Federal Republic (CSFR), ordered a CRAY Y-MP EL system, the first order for a Cray Research system from an organization in Eastern Europe. The system will be installed at the academy's Institute of Physics. The Academy of Sciences, one of the oldest institutions in Europe, is an assembly of about 1000 faculty members from universities throughout the republic. The new system will be used for physics research.

"The CRAY Y-MP EL system will open up a new world for us," said Jaroslav Nadrchal, head of the computer department at the Academy's Institute of Physics. "Acting as a node on our computer network, the system will be used for large-scale computations of atomic and electronic structures of complex solids, simulation of kinetic processes, and in high energy physics."

The University of Magdeburg, Magdeburg, Germany, located in the new federal state of Sachsen-Anhalt in the former East Germany territory, accepted a CRAY Y-MP 2E supercomputer in December. The new system will be used for general academic research in physics, chemistry, and engineering with special emphasis on problems involving fluid dynamics and structural mechanics. "We hope to demonstrate the positive impact of supercomputing technology on the education and economy within the new federal states," said Rolf Knocke, director of the computer center at the university.

Department of Energy, Cray Research team up to boost competitiveness

As part of the U.S. Federal Government's High Performance Computing and Communications Initiative, the Department of Energy (DOE) and Cray Research in October entered into a preliminary agreement calling for two of DOE's premier research facilities—the Los Alamos and Lawrence Livermore National Laboratories—to work with the supercomputer-maker on developing massively parallel processing (MPP) systems that will result in greater capabilities for the national labs and enhanced competitiveness for U.S. firms in the global market. Through the agreement, to be funded jointly at \$70 million over a three-year period, the two DOE labs and Cray Research will work together to create operating systems and capabilities for use on the company's MPP systems. Four areas—environmental modeling, defense systems, materials design, and advanced manufacturing—will be among the most direct beneficiaries.

APPLICATIONS UPDATE

MSC/NASTRAN Version 67R2 available for Cray Research systems

Two important enhancements in V67R2 of the MSC/NASTRAN structural analysis package substantially reduce the time required to solve large finite element problems on CRAY Y-MP systems. The two enhancements are reduced memory requirements for the sparse decomposition module and a Cray Research-specific I/O buffering mechanism. (MSC/NASTRAN V67 previously introduced a sparse decomposition algorithm that is three to ten times faster than the traditional method; however, it required a large amount of memory.)

The sparse algorithm in V67R2 reduces the amount of required memory by a factor of two or three. Now, most models that contain up to a half million degrees of freedom (DOF) can be run with fewer than 32 Mwords. The additional I/O associated with using the minimum amount of memory can be large, so it is still beneficial to use more than the minimum. V67R2 also uses a Cray Research-optimized numeric kernel for the sparse decomposition module which is up to 30 percent faster than V67.

The Cray Research-specific I/O buffering mechanism in V67R2 takes advantage of the flexible file I/O (FFIO) mechanism in the UNICOS operating system to realize more of the potential bandwidth of current and future Cray Research I/O devices and larger memory systems. Two software layers developed specifically for the type of I/O done by MSC/NASTRAN are available in addi-

tion to the standard UNICOS layers. The Enhanced Intelligence Engineering (EIE) layer provides management of a smart paged cache using either memory or secondary data segments (SDS) for the cache space. It detects sequential file access, either forward or backward, and uses asynchronous I/O to prefetch data that is likely to be required soon and write cache pages that have been filled (read-ahead and write-behind).

The stripe layer lets the user spread the I/O across multiple devices and channels to make efficient use of the site-specific disk configuration. When used with the EIE layer on a file system in which the partitions are independent devices, the stripe layer allows multiple asynchronous reads and writes to proceed in parallel, which can dramatically increase the I/O throughput. I/O wait time can be reduced by a factor of 10 using the stripe layer and a relatively small amount of memory or SDS space for the EIE cache. The overall elapsed time required to solve many jobs has been cut in half.

In a recent test conducted by Cray Research applications group analysts, a full car body model that included 680,000 DOF was run as a single super-element. Running on a CRAY Y-MP M92 system using MSC/NASTRAN V67R2 and 140 Mwords of main memory, the problem took only six CPU hours and eight Gbytes of disk space to run to completion. When this job ran using an earlier release of MSC/NASTRAN on a CRAY Y-MP 4E/232 system, where the user was restricted to 12 Mwords of main memory, the job was split into 11 super-elements. The job took 15 CPU

hours and 20 Gbytes of disk space to be completed.

In another recent test, a 1.1 million DOF static model was run on a CRAY Y-MP C90 system for a different automotive customer. The auto body model contained 182,000 grid points and 192,000 elements. The standard solver could not be used because the fill-in caused by this method would create a data block that exceeds the MSC/NASTRAN limit of about 1.2 billion words. The estimated time for the decomposition using this method is 80,000 CRAY Y-MP CPU seconds. The analysis, including 75 Gbytes of I/O, took less than two hours using the sparse solver in MSC/NASTRAN V67R2 on a CRAY Y-MP C90 system.

CRAY Y-MP C90 and CRAY Y-MP M90 systems and MSC/NASTRAN V67R2 give automobile and airplane manufacturers the competitive edge in structural analyses by providing

- ☐ The ability to run large MSC/NASTRAN models as one super-element
- ☐ More efficient processing of larger, more complex models
- ☐ Reduction of engineering time
- ☐ Substantial cost savings

MSC/NASTRAN V67R2 requires UNICOS 6.1 or later.

For more information on using MSC/NASTRAN V67R2 with CRAY Y-MP systems, contact Dennis Nagy, The MacNeal-Schwendler Corporation, 815 Colorado Blvd., Los Angeles, CA 90041; telephone: 213/259-3814; or Doug Petesch, Cray Research, Inc., 655E Lone Oak Drive, Eagan, MN 55121; telephone 612/683-3654.

Eli Lilly scientists win supercomputing award

A CRAY-2 supercomputer played a key role in the winning of the 1992 Industrial Grand Challenge award by Robert Hermann and David Herron of Eli Lilly & Co. The award is presented annually by the National Center for Supercomputing Applications (NCSA) to the industrial scientists with the most significant supercomputing success story of the year.

Hermann was recognized for his seminal work on the structure of human phospholipase A2, an enzyme which was derived from bovine pancreas and rattlesnake venom. Calculations performed on NCSA's CRAY-2 system resulted in the subsequent development of a humanized form of the enzyme on Eli Lilly's own CRAY-2 supercomputer. This simulation was used to solve the x-ray crystal structure of the human enzyme, which is implicated in toxic shock.

Herron was acknowledged for his research on leukotrienes, the inflammatory agents in asthma. Three key leukotrienes were analyzed using scientific visualization. With the visualization acting as a guide, highly active antagonists against the leukotrienes were synthesized by Lilly asthma researchers. Some have been tested in asthma sufferers and were found to be medically

effective. Herron's work has contributed to the development of new compounds that are potential new treatments for asthma.

Lilly, who joined NCSA's Industrial Program in January 1988, was one of the first companies to use supercomputers for drug design. Shortly after Lilly's affiliation with NCSA, two new molecular codes, Ovid and Super, were created jointly. Both codes were ported to NCSA's CRAY-2 system. By the spring of 1990, Lilly had installed its own CRAY-2 supercomputer at its research headquarters in Indianapolis.

1992 Supercomputer Challenge results

Months of research and programming paid off as five Española Valley High School students watched an artificial organism evolve inside a CRAY Y-MP supercomputer. The creation of the organism earned the Española Valley team first prize in the second annual New Mexico Supercomputing Challenge. The "Challenge" is designed to encourage high school students to participate in a yearlong program to devise computational solutions to scientific problems by using over \$100 million of supercomputing power available in the state.

"Programs like the Supercomputing Challenge are designed to enable the youth of today to become the giants

of tomorrow in science and engineering," said Irene Qualters, vice president for software development at Cray Research.

The Española Valley team competed against 127 teams of budding computer scientists from more than 50 New Mexico schools. The teams researched scientific problems, learned basic computer functions and programming languages, and ran their programs on the world's most powerful computers at Los Alamos and Sandia National Laboratories and the U.S. Air Force's Phillips Laboratory. Scientists from the national laboratories, university and college faculties, and the private sector provided coaching.

The Española Valley team programmed the artificial mind for its organism and the rules governing its behavior into one of Los Alamos' CRAY Y-MP supercomputers. Using a genetic algorithm in a program known as the Learning Classifier System, the organism, called "The Critter" by the students, evolved by learning survival skills such as finding food and avoiding predators within a simulated environment.

In addition to Cray Research, sponsors of the Supercomputer Challenge include the three laboratories, the University of New Mexico, and New Mexico Technet, a nonprofit, statewide computer network that links New Mexico's national laboratories, universities, and many private businesses.

CUG reports

Users of Cray Research computer systems established the Cray User Group (CUG) in 1977 to provide a forum for the exchange of ideas related to Cray Research systems and their applications. The group holds two general meetings each year. Its second meeting of 1992 was held September 14-18 in Washington, D.C., and was hosted by Technology Applications, Inc. Below, CUG President Gary Jensen of the National Center for Atmospheric Research, offers his comments on the meeting and other CUG-related business.

CUG meetings just keep getting better. TAI did a wonderful job organizing the Washington meeting in September. CUG is an all-volunteer organization, and the volunteers are to be commended for their hard work. Without them, there would be no CUG.

The new Board of Directors elected at the meeting is listed at the end of this article. The previous Board opened up some new ways of doing things, including scheduling meetings by "tracks." For the next meeting, the entire schedule for the technical program has been changed, so if you plan to attend, please read your meeting announcement carefully.

Sites receive various forms of payback when they send representatives to CUG meetings. They acquire valuable technical information and are given an opportunity to influence Cray Research's direction. Another benefit is the opportunity to meet directly with representatives from Cray Research to discuss problems or offer suggestions. All areas of Cray Research are represented by people who can make decisions.

Perhaps the most valuable information is exchanged during the breaks, lunches, and outings. During these times participants are making friends with similar problems and challenges. This networking develops business relationships that last between meetings. You may not have to reinvent the wheel if you find someone that just solved your problem.

The following information about CUG was organized by CUG Secretary Gunter Georgi and is presented as an overview for readers who may not be familiar with the organization.

What is CUG?

The Cray User Group (or CUG) is a not-for-profit, international organization of sites that have procured computers from Cray Research, Inc. The mission of CUG is to provide the high-performance computing community with leadership and information exchange to enable the development and effective use of Cray Research computational tools in achieving the business and research objectives of CUG members.

What benefits does CUG offer members?

Once a site joins CUG, it can send representatives to CUG conferences. Conference attendees can establish contacts with Cray Research personnel and other CUG members. CUG generates publications that are sent automatically to each installation representative and to conference attendees. In addition to meetings, CUG sponsors several activities related to information exchange between member sites. These projects include the unicos-l electronic mailing list and CUGDUS, a centralized catalog of UNICOS system software (utilities, libraries, scripts, and mods) and applications available from the various UNICOS sites.

How often are meetings? When? Where? Who can attend?

CUG meetings of one week duration are held twice a year, usually in April and September. Because CUG is an international organization, the meetings are scheduled in North America, Europe, and Asia. The registration cost per person varies between US \$300 and \$600, depending on location and exchange rate. Attendees of the meeting must be employees or consultants of regular CUG member sites. Every CUG meeting includes a reception for new CUG attendees and a session that explains the functioning of CUG.

What are SICs and MIGs?

SICs are Special Interest Committees that have special tracks of sessions unique to particular areas. MIGs are Mutual Interest Groups that are special-topic or industry-oriented organizations. Cray Research sends representatives to all SICs and MIGs. SICs currently schedule sessions on applications and algorithms, graphics, management, mass storage systems, networking, operating

systems, operations, performance and evaluation, and user services. Mutual Interest Groups (MIGs) have sessions on aerospace, environment, massively parallel, and security.

How do I join CUG?

Sites that have procured computers from Cray Research can join. Application forms are available from the CUG secretary. After approval of the application, the CUG treasurer will send an invoice. The current annual membership fee per installation is US \$400 per year. The CUG fiscal year is from October 1 until September 30 of the next year.

Who do I contact for more information?

Listed below are the current Board members, together with their company affiliations and e-mail addresses.

President: Gary Jensen, National Center for Atmospheric Research, guido@ncar.ucar.edu

Vice President: Jean Shuler, National Energy Research Supercomputer Center, shuler@nersc.gov

Secretary: Gunter Georgi, Grumman, georgi@ncar.ucar.edu (primary information contact)

Treasurer: Howard Weinberger, Technology Applications, weinberger@plk.af.mil


Regional Director-Europe: Claude LeCoeuvre, Commissariat à l'Energie Atomique/CEL-V, lecoeuvr@glai.eul, saclay cea.fr

Regional Director-Asia/Pacific: Kyukichi (Eugene) Ohmura, CRC Research Institute, k-oomura@crc.co.jp

Regional Director-Americas: Francis J. Pellegrino, Westinghouse Electric Corporation, pellegrino@a.psc.edu

Past President: Karen L. Sheaffer, Sandia National Laboratories, Livermore, karen@snll-arpagw.llnl.gov

The next CUG meeting is in Montreux, Switzerland, March 29-April 2, 1993. The fall meeting is in Kyoto, Japan, September 20-24, 1993.



Volume data set (brain) with surface geometry was created with a hybrid ray tracing system that supports both volumetric and surface-based data sets (eyes, environment). The software, Denga, runs on the CRAY Y-MP8/864 system located at the Ohio Supercomputer Center. Image courtesy of Don Stredney, Michael Torello, and Stephen May of The Advanced Computing Center for the Arts and Design, The Ohio State University.