A CRAY RESEARCH, INC. PUBLICATION

CRAY CHANNELS

Volume 6, Number 1

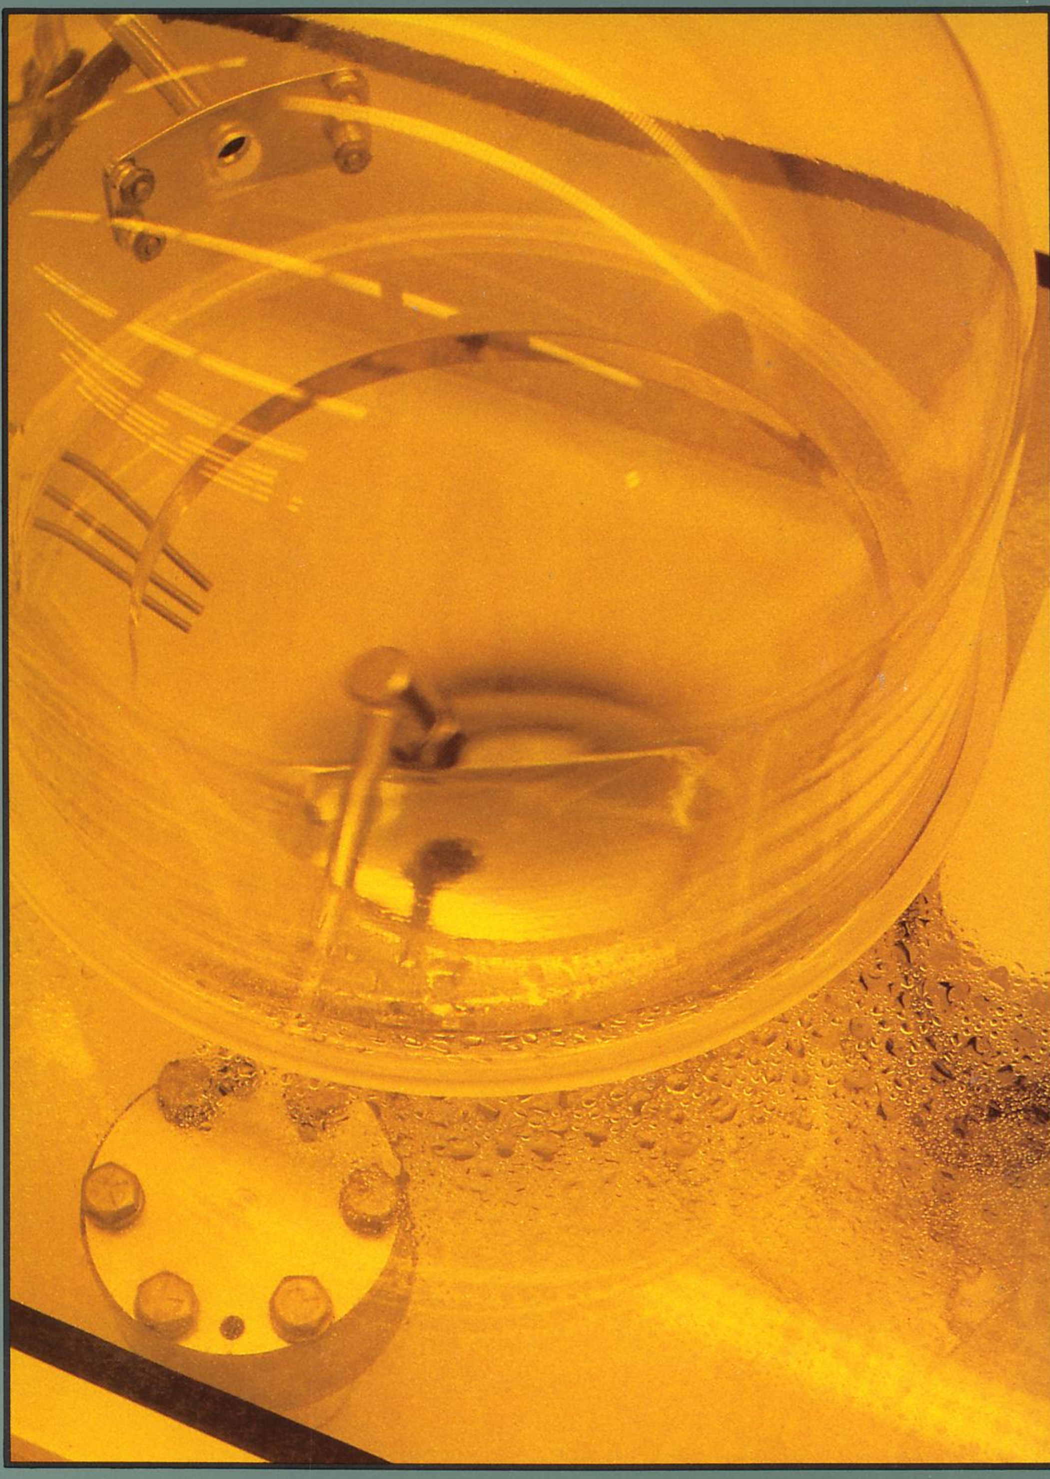

FEATURE ARTICLES:

REGULAR COLUMNS:

IN THIS ISSUE

CRAY computers around the world process information that can and does effect each and every one of us daily. Would you have guessed that the weather forecasts heard this morning may have been formulated with a CRAY in the background? Or that mathematical research possible with the CRAY may impact the security and privacy of electronic funds transfers, auto-bank tellers, and the like? Would it ever cross your mind that a CRAY computer is involved in developing the United States' atmospheric emergency response capability that helps guarantee our safety? And did you know that the spiffy car commercial on TV that may have caught your eye recently was created with the help of a CRAY?

The articles in this issue of CRAY CHANNELS discuss work CRAY computers are doing that touches our lives in ways that would never have crossed your minds. We hope you enjoy this issue that includes topics ranging from weather forecasting to chess championships. And please keep in mind that if you have ideas for topics you'd like to see covered in future issues of CRAY CHANNELS, call or drop us a note. We're interested, as our other readers are, in what you're thinking and doing!

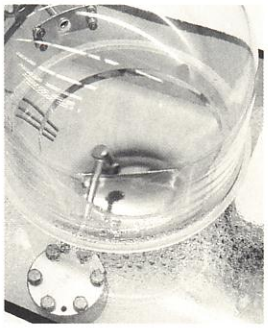

On the cover a photo resist spinner installed at Cray's IC fabrication facility is shown. The device is used in applying a thin uniform layer of photo resist on silicon wafers. To do so, a wafer is placed on a chuck (holding mechanism) in the spinner; resist is dispensed onto it. The wafer is spun at high speed — about 3,000 RPMs to obtain the thin uniform coating of resist. It is then ready for the next step in fabrication.

CRAY CHANNELS

A Cray Research, Inc. publication

Volume 6, Number 1

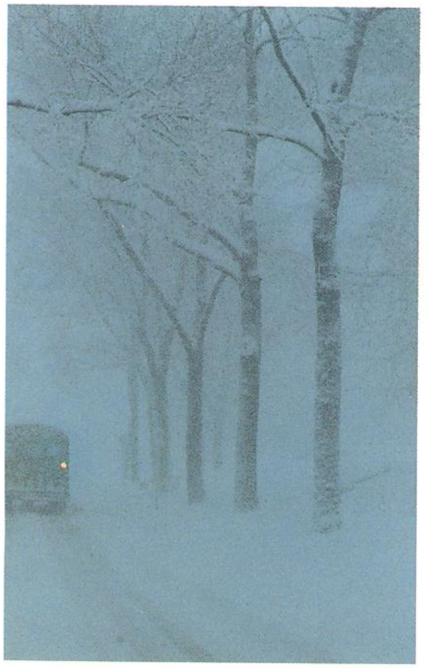

2

FEATURES

DEPARTMENTS

Credits

Editor
Carol Smith

Design
Kathleen Van Kirk
Diane Thomas

CRAY CHANNELS is a quarterly publication of Cray Research, Inc., 608 Second Avenue South, Minneapolis, MN 55402. It is intended for users of Cray computer systems and others interested in the company and its products. Subscription requests, feature story ideas and news items submitted to the editor at the above address are welcomed.

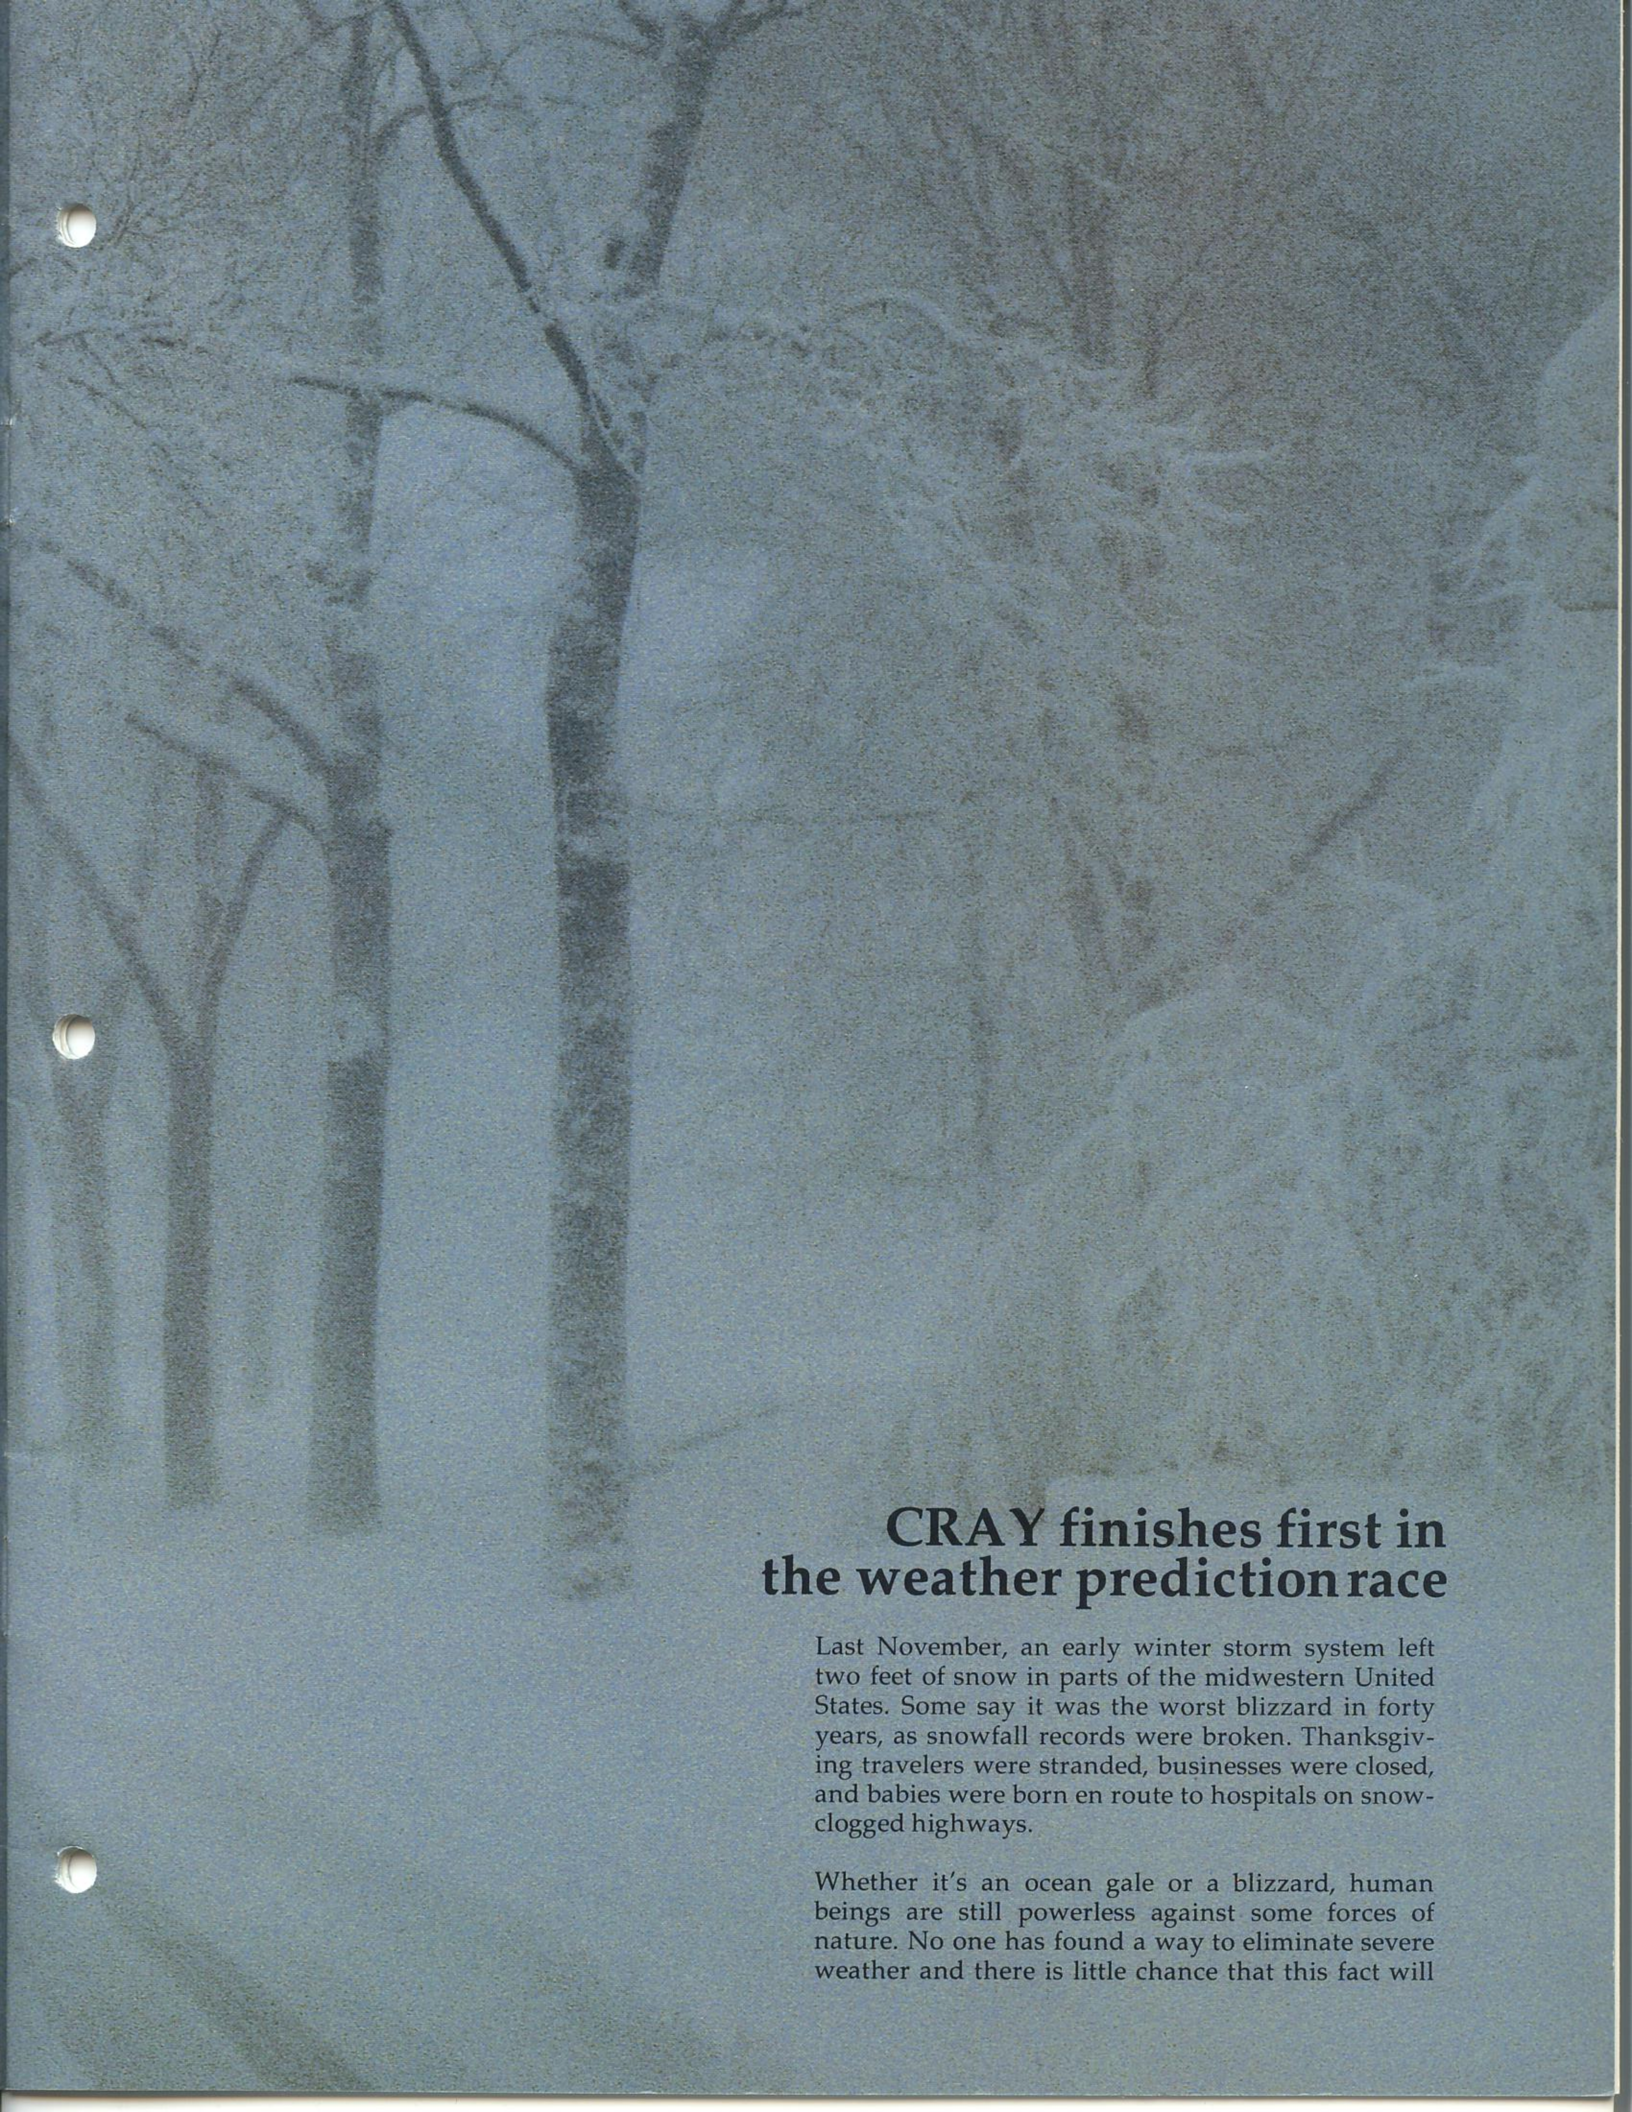

CRAY finishes first in the weather prediction race

Last November, an early winter storm system left two feet of snow in parts of the midwestern United States. Some say it was the worst blizzard in forty years, as snowfall records were broken. Thanksgiving travelers were stranded, businesses were closed, and babies were born en route to hospitals on snow-clogged highways.

Whether it's an ocean gale or a blizzard, human beings are still powerless against some forces of nature. No one has found a way to eliminate severe weather and there is little chance that this fact will change in the near future, so the best we can hope for is to have time to prepare for it. With a bit of warning, human and economic tragedy could be avoided or minimized.

Accurate weather forecasting based on thorough meteorologic and atmospheric research methods is our best protection against surprises from the elements. Over the years, scientists have developed exacting procedures for gathering, analyzing, and disseminating weather information. Since World War II, computers have played an increasingly important role in predicting the weather.

Large-scale computing in weather centers

Today, supercomputers are essential in getting weather information out — and more CRAY systems are pitching in than ever before. Supercomputers can process more data, more accurately, and in less time, thus enabling meteorologists to make better forecasts. CRAY systems are also heavily used in climatology and atmospheric research. Because the models for this type of research are so large and complex, computers with power like the CRAY are essential. From large-scale climatic and atmospheric studies to predicting tomorrow's weather, CRAY computers are helping scientists around the world to anticipate nature's brusque ways.

When CRAY systems first were introduced, their potential application to weather research was recognized almost immediately. The second CRAY computer ever built was installed in a weather research organization. The National Center for Atmospheric Research (NCAR) in Boulder, Colorado, which is involved in many aspects of atmospheric research, installed its first CRAY in 1977. After running the system at capacity for a number of years, NCAR added a second CRAY in 1983.

In 1978, the European Centre for Medium-range Weather Forecasts (ECMWF) in Reading, England installed a CRAY, and recently upgraded to a CRAY X-MP. Lennart Bengtsson, Director of the Centre commented, "With the first CRAY, our six-day winter time forecasts were as accurate as the best winter two- or three-day forecasts available in Europe 12 years ago. With the CRAY X-MP, we are confident that we will be able to predict even more accurately, and to issue useful forecasts for even longer ahead."

Late in 1983, the Canadian Atmospheric Environment Service of which the Canadian Meteorological Centre (CMC) is a part, installed a CRAY-1/S. Two years before, CMC realized that in order to contribute more valuable weather forecasts, computer power must be increased from the two to three million floating point operations per second (MFLOPS) of its existing system to a system that could sustain 50 MFLOPS over long periods of calculation. CMC selected the CRAY to address its current needs, and will migrate to a CRAY X-MP in 1986. Officials at CMC estimate that the system increases their computing capability by about 1000%.

"We have contracted to upgrade our CRAY computer power over the next six-and-a-half years so that a ten-day forecast in 1990 will be as accurate as the current five-day forecast," Zavie Miller, CMC Project Manager for the CRAY installation said.

Iain Findleton, Chief of Meteorological Operations at CMC, explained that CMC has a vast and changeable weather climate to monitor. "Weather systems can change significantly as they move across Canada," says Findleton. "Canadian weather can include snow squalls, gales and tornadoes — all of which not only threaten human life, but have an economic impact as well. In addition to better forecasting, the CRAY also opens the door to improving climatological research being done by the Atmospheric Environment Service, and for ongoing research relating to acid rain."

Forecasting

The steps involved in predicting the weather haven't changed much over the past 50 years; large amounts of data must be collected and analyzed to produce a representation of atmospheric conditions at a given time. That information is then projected into the future using a scientific model. Meteorologists calibrate the results of the model with other available information and derive a forecast.

Developing a weather forecast for a period in excess of 24 hours involves five phases: data acquisition, evaluation, analysis, forecasting, and dissemination. The World Meteorlogical Organization (WMO), representing services of many countries, and forecasting organizations on national, regional, and local levels all play a part in the weather prediction process.

Gathering the data

How much data is needed to produce an accurate forecast? The answer is easy: as much as possible. The more data and atmospheric variables that can be considered, the greater the chances of developing an accurate forecast. Today, very sophisticated data acquisition procedures are in place around the world to gather atmospheric information.

Collecting data involves measuring temperature, barometric pressure, humidity and wind. Twice each day at 0000 Coordinated Universal Time (UTC) and 1200 UTC, data is obtained by radiosondes (weather balloons) that record information at more than 750 locations around the world. In addition, surface data is collected four to eight times a day, at almost 9000 locations. Additional data is recorded throughout the day at various times and locations by aircraft, satellites and ships. Other useful data includes readings of sea level pressure, upper air temperature at different altitudes, geopo-

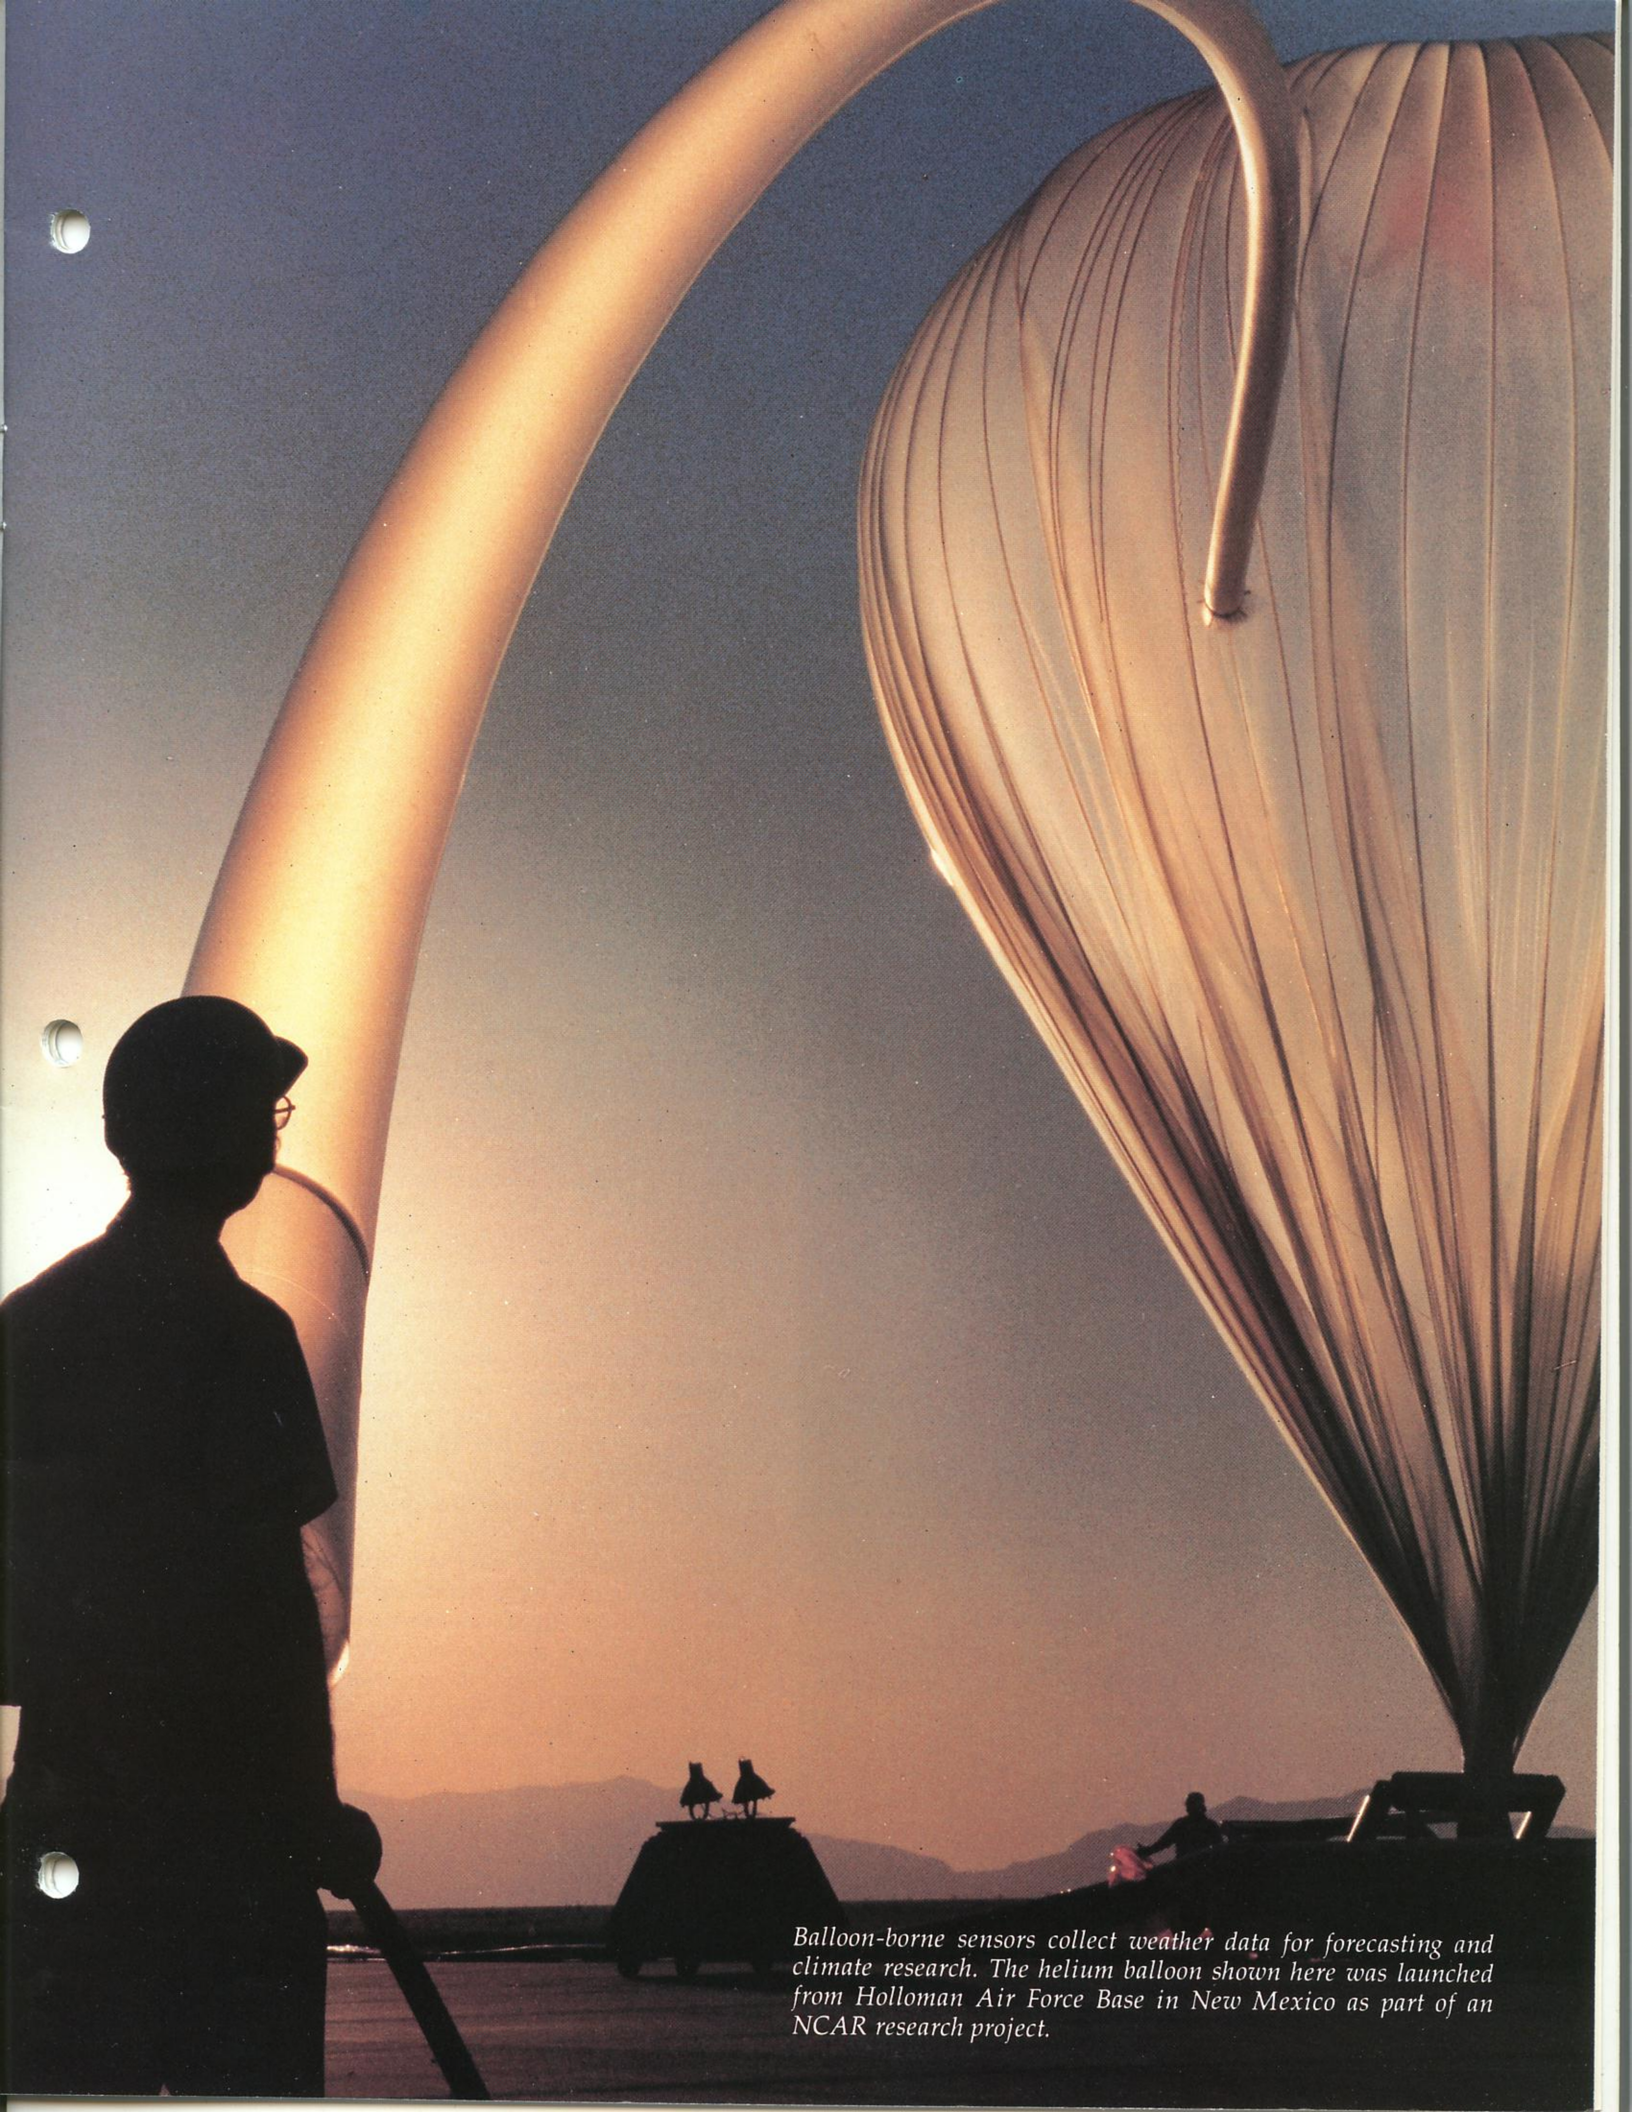

Balloon-borne sensors collect weather data for forecasting and climate research. The helium balloon shown here was launched from Holloman Air Force Base in New Mexico as part of an NCAR research project.

tential heights and winds, precipitation, snow cover, clouds and more. The observations are broadcast by the WMO's dedicated communications network, the Global Telecommunications System, which makes the data available to nearly every nation shortly after it is recorded.

Analyzing the data

Weather centers around the world analyze the data necessary to produce a forecast for a particular area. Most national weather bureaus attempt to issue forecasts within six hours after data observations are received, hence time becomes an important factor at this point. In Canada, for example, the recorded data is ready at CMC within one-and-a-half hours of the observation time, leaving just four-and-a-half hours for analysis, interpretation, processing, and dissemination.

Modeling the weather

Mathematical forecasting models are composed of algorithms representing the laws of physics that affect the atmosphere such as fluid motion, thermodynamics. In addition, the large and sophisticated programs take into account such important

A prototype computerized global model of the earth's atmosphere was produced by a CRAY computer at the National Center for Atmospheric Research.

physical effects as solar radiation, latent heat release, friction as the wind blows over irregular terrain. The equations are applied against a grid point network; at each grid point, a set of meteorological values are computer, simulating the change of the weather for a single time step (typically on the order of 15 minutes). By repeating this process over and over, forecasts are made for one, two, or ten days ahead.

The resolution of the weather model contributes to the overall accuracy of the forecast. It is conceivable that an influential weather system may not even be considered in an analysis if it is positioned between grid points — and its potential effects will not be represented in the forecast. Iain Findleton explained, "If our model grid points are 150 miles apart, we may not even pick up major changes in a weather system." Increasing the number of grid points improves model resolution, but significantly increases the computational task. To illustrate, at CMC computation would take 16 times longer than it does now if resolution were doubled in both horizontal and vertical directions ($2^3 = 8$), and twice as many time steps are required.

Already the computing requirements for processing state-of-the-art weather models are very high. At ECMWF, approximately one-half trillion computations must be executed to produce a ten-day forecast. Out of necessity, a computer capable of maintaining performance at 50 million instructions per second is used. The organization's CRAY system easily handles this type of workload. "Introducing models which represent the atmosphere with greater precision, taking into account more and more physical processes and representing ever more accurately the effects of mountains, oceans etc, will continue to require vastly more powerful computers. Our scientific staff are pointing to the need for ever greater number-crunching power," Dr. Bengtsson, at ECMWF says.

Currently, ECMWF is in the process of implementing a new type of model that will extend the usefulness of forecasts by about one day and provide even more accurate forecast data. Development of new models such as this is very costly and places a heavy load on weather centers' resources. It is this type of research that has earned ECMWF its reputation as a premiere weather center.

Disseminating the forecast

As the forecast is pushed forward in time on the model, the results are output at regular intervals, e.g. every 12 hours. The output is interpolated from the model grid to the analysis grid and written to the database, which then allows the forecast information to be processed further or displayed.

Nearly all forecast data is graphically displayed for easy evaluation by meteorologists. CRAY computers are used at all three weather installations to generate satellite photos and contour maps. The centers then send forecast maps to regional weather centers, where local meteorologists compare regional data with the model forecast and derive a forecast for their area.

Several factors can inhibit the quality of the final weather forecast: the number of observations included in the analysis, the quality of the observations, the soundness of the mathematical techniques, and the degree to which the model reflects conditions in the real world. Of course, the ele-

ment of time may also limit the thoroughness of the analysis. The amount of computation which a given set of data undergoes in combination with the factors listed above ultimately translates into the accuracy of the forecast.

The impact of forecasts

Nature's influence on the way we live is more pervasive when it comes to weather than other natural phenomena. The conditions on the outside can make our lives infinitely easier or cause untold misery depending on nature's whim. If we can outwit the elements and prepare for what is in store for us in a timely way, the number of lives lost, crops damaged, money wasted, and strife inflicted can be lessened. That is the reason that so much effort is devoted to predicting and understanding weather patterns. The complexity of weather is why supercomputers are needed.

It is difficult to measure the exact economic value of current or improved forecasts. However, recent studies conducted under the supervision of the World Meteorological Organization have concluded that a benefit/cost ratio of at least 20 to 1 would apply for the national weather service of any industrialized country.

"Accurate medium-range forecasting provided by the CRAY can be crucial to the Canadian economy," Iain Findleton at CMC commented. Weather forecasts can save farmers millions of dollars. Based on anticipated weather, farmers decide about planting, spraying, fertilizing and harvesting crops. For instance, for spraying to be effective, calm winds the day of spraying and no rain for two days following are needed. At a cost of about $250 per acre, wasted spraying can mean big losses for farmers. In 1977, it was estimated that 15% of all fertilizers used by U.S. farmers was washed away by heavy rain. Canadian officials have estimated that reliable weather information could have saved farmers $22 million during the same time period.

The ramifications of the fertilizer runoff as described above are more wide-reaching than most of us would think. If farmers use less fertilizer, herbicides, or pesticides over a growing season, natural gas consumption declines, since ammonia, which is a basic ingredient in most fertilizers, is derived from gas. Using fewer herbicides and pesticides also reduces the ecological impact of these products.

Better weather forecasts can help minimize air travel delays and provide information that would allow planes to fly at optimum altitudes. An estimate in 1977 showed that if fuel consumption could be reduced 1.8%, a saving of $10 million could be realized annually. Preventing weather-related aircraft accidents is another example of how improved forecasting can save time, money, and even human lives.

Utility companies use weather information to determine peak loads in order to plan for the most economically efficient power generation. Officials at Canadian Meteorological Centre determined that if the accuracy of forecasting the temperature could be improved by one degree Centigrade, the annual savings in Canada would be approximately $2 million.

Accurate weather forecasts help the transportation industry by lowering operating costs, reducing damaged and lost cargo, and reducing the cost of rescue operations. Shipping companies rely heavily on weather forecasts for planning transcontinental routes and avoiding delays.

These are just a few examples of the effects of accurate weather forecasting At ECMWF, it was estimated that the value of reliable forecasts of up to ten days was $250 million per year in 1970 — and the savings is considerably greater now.

Research

There is a fine distinction between forecasting and climatic research. While weather and climate forecasting both use the same type of computer model as a research tool, weather forecasting describes the instantaneous state of the weather, while climate research reflects the average weather patterns over a long-term period. Atmospheric and climate research findings offer scientists information about the long-range weather trends that evolve from diverse phenomena, such as increased carbon dioxide in the atmosphere, fluctuations in ocean surface temperatures such as El Nino, volcanic dust and air pollution.

Using two CRAY 1-A systems at NCAR, scientists explore such topics as: why severe weather occurs more regularly over certain areas; solar activity and its effect on the atmosphere; the shifting of substances in the atmosphere, such as ozone or carbon dioxide; and climatic patterns and trends that may cover a few years or a few centuries.

CRAY computers make it possible for scientists to perform the advanced modeling and data analysis that such research requires. One example is a model that was created when the Soviet Union announced it was considering a plan to change some of the flow of four major rivers by diverting the flow southward, away from the Arctic Ocean where the rivers normally empty. The plan was designed to increase Soviet agricultural productivity, but many scientists were concerned that such a plan might reduce Arctic sea ice and indirectly result in a climate change. The Soviets maintained that the effects would be minimal.

To study the potential effects of a river diversion, a numerical model was developed at NCAR and run on one of the CRAY's. After simulating the possible conditions if such a plan were followed, the results showed that sea-ice extent would be virtually unaffected. It was decided that the Soviet river diversion plan would probably not disturb the climate of the Arctic.

NCAR scientists also study mesoscale weather, which includes phenomena such as blizzards, hurricanes, tornados, and hailstorms. Since these patterns of severe weather are not well understood, NCAR scientists developed models to study the formation and the reaction of such storms with other weather systems. The models were used to study theoretical conditions and to investigate real-data situations. Scientists have been able to ascertain the validity of their weather models in this way.

Canadian researchers at CMC also plan to use the CRAY for simulations that will provide input into government decisions. For example, if the government were researching the possibility of increasing the use of coal as a primary fuel, a model could actually "place" coal plants at several appropriate locations and with data such as winds, other pollutants in the area and atmospheric moisture, determine the long-range effects of burning coal at those locations. Climatic models already exist for this type of application, but the CRAY-1/S will allow CMC to refine the models.

Without research into improving models, there would be no hope of increasing either the accuracy or usable period of a forecast. Before a change is made to a model, much experimentation takes place to determine if the change will actually improve overall model results. For example, ECMWF discovered that Atlantic low pressure systems were not following the path delineated by the model. The systematic error was finally traced to the heights of the Rocky Mountains and the Alps. These mountain ranges tend to deflect and influence the movement of weather systems all over the world, but the grid would average and flatten out the mountain heights into a plateau, so the effect of the mountains wasn't taken into consideration. Revising the model to account for mountain heights extended experimental forecasts an average of six to twelve hours — and sometimes as much as two or three days.

The research that goes into developing a new forecast model can take years to complete. A new model must remain operational for a minimum of four or five years so that many sets of data can be tested, and the model can be altered and continually developed. The improvements contained in a new model are usually caused by refining the mathematical equations, which increases the number of computations and requires a substantial upgrading of the computer system.

CMC, ECMWF, and NCAR share the results of their efforts, since a new model can be used for a wide variety of applications in both weather forecasting and climatic research. Communicating new techniques and the results of experiments provide important information to all scientists working on weather-related research.

Conclusion

You can be sure, however, that scientists have no information indicating that Mother Nature has any intention of taming her pedantic personality. Nor have meteorologists found the key to controlling the many forms of her unabashed rage. But, today, the timely and more accurate forecasts generated with the help of supercomputers are as important (if not more so) as the coats on our backs in protecting us from the elements.

CRAY computers are powerful tools helping to improve the timeliness and accuracy of weather forecasts. Perilous weather like snowstorms and

CMC

The Canadian Meteorological Centre (CMC) in Montreal is the center for numerical weather prediction for the Atmospheric Environment Services (AES). At CMC, objective analyses are performed on current weather conditions and weather forecasts are produced and disseminated to Regional Weather Centres across Canada and to the Canadian Forces Weather Service. Canada's size and varying climate make accurate weather forecasting especially crucial. CMC has used computer-based numerical weather prediction for the past 20 years. Recently CMC became the first CRAY-1 customer in Canada.

ECMWF

The European Centre for Medium-range Weather Forecasts (ECMWF) was established in 1975 to develop medium-range weather forecasts (four to ten days ahead) on an international basis. ECMWF is located in Reading, England and serves the following European member states: Austria, Belgium, Denmark, Finland, France, Federal Republic of Germany, Greece, Ireland, Italy, Netherlands, Portugal, Spain, Sweden, Switzerland, Turkey, the United Kingdom, and Yugoslavia.

In addition to developing forecast models and preparing medium-range weather forecasts for distribution to Member States, ECMWF also helps implement the programs of the World Meteorological Organization (WMO). ECMWF was Cray's first international customer in 1977 and is in the process of installing a CRAY X-MP.

NCAR

The National Center for Atmospheric Research (NCAR) in Boulder, Colorado is sponsored by the National Science Foundation and operated by the nonprofit University Corporation for Atmospheric Research, comprising 53 universities with graduate programs in the atmospheric sciences. NCAR was created in 1960 to focus efforts in atmospheric research, concentrating on problems of national and global importance to produce knowledge that can lead to better-informed policy decisions and actions concerning the atmospheric environment. NCAR's research efforts include these major areas: climate, weather prediction, atmospheric quality and solar research. The Scientific Computing Division at NCAR has two CRAY-1 computers.

hurricanes can be anticipated better, thus, people are better able to prepare for their effects. Workers in industries as diverse as agriculture and shipping can schedule activities to avoid the costly impact adverse weather can have on their businesses. And the important research into the long term atmospheric trends is providing information that helps us plan for the subtle climatic changes that can mean the difference between life and death and the preservation of our delicate environment.

Acknowledgements

This article was prepared with the kind assistance of the following people: Bruce Atfield, Canadian Atmospheric Environment Service; Zavie Miller and Iain Findleton, both of CMC; Joan Vandiver Frisch, NCAR, Lennart Bengtsson, ECMWF; and Robert Wellck and Jacques Belanger, Cray Research, Inc.

Photo credits: Page 2, Janet Robidoux, Cray Research, Inc.; pages 4 and 6, NCAR; page 9, Everett Anderson.

References

Besen, Linda, "Scientific Computing at the National Center for Amospheric Research." *Documentation Project of the Scientific Computing Division (NCAR)*, June, 1983.

"Using Computers to Model the Weather," *Weatherwise*, Vol. 36, No. 5, October, 1983, pp 224-238.

Kerr, Richard A., "The Race to Predict Next Week's Weather," *Science*, Vol. 220, April, 1983, pp 39-41.

Knighting, E., "*The European Centre for Medium-range Weather Forecasts,*" 1982.

Wolkomir, Richard, "Scientists Are Doing Something About the Weather," *Smithsonian*, Vol. 13, No. 1, April, 1982, pp 88-96.

Emergency response capability enhanced with new perspectives

Patrick P. Weidhaas & Hoyt Walker,
Lawrence Livermore National Laboratory

The pictures were produced on a CRAY-1, requiring approximately 30 seconds of CPU time each at a resolution of 1580x2048 pixels. All three images are different perspective views of the San Francisco Bay Area and are described from the lower left, rotating clockwise. The first is a view northwards with San Jose and the southern tip of the Bay at the lower left, and the chain of the Eastbay Hills extending parallel to the Bay from the lower left corner to the center of the picture. The second photo is a view from the Pacific northwards; Half Moon Bay is in the foreground, Marin County at the upper left edge, the Eastbay further to the right. The third image is a view of Richardson Bay in Marin County. The snow-covered mountain is Mt. Tamalpais; Sausalito would be on the lower left part of the shore.

Imagine yourself sitting at a color graphics terminal, touching a few keys, and within seconds looking at a three-dimensional image of the San Francisco Bay Area viewed from the south at an altitude of 4,000 feet. Touch a few more keys, and the image shifts, showing the same area, but this time viewed from the northeast at 6,000 feet. Parts of the hills are brightly lit by the incoming sunlight, other parts are in shade. Now imagine that with a few additional commands you could generate a realistic shaded color image of any other area within the continental United States, no matter how remote. Every hill, mountain, valley and canyon is available and can be viewed from any position. If a particular image strikes your fancy, just hit a button and obtain a color hardcopy of the scene within seconds.

The realization of such a system is one of the ambitious goals of the Atmospheric Release Advisory Capability (ARAC). Research on the necessary tools is in progress (relying heavily on the CRAY-1 computers) at the Lawrence Livermore National Laboratory, home site of the ARAC project.

Discussion of the graphics techniques would be incomplete without including the motivation for such a capability within the ARAC system. Thus, this article initially focuses on ARAC, followed by a description of the graphics technique employed in implementing the system.

The ARAC system provides real-time assessments of the consequences which could result from an atmospheric release of radioactive material. The system is an integration of data acquisition systems, numerical models, data analysis techniques, and professional staff supporting the Department of Energy and the Department of Defense in their role of providing an emergency response capability to Federal and State agencies. Situations that may require ARAC's participation include accidents at nuclear power reactors and other facilities handling nuclear materials. Some of the key responses during the ten years of ARAC's existence include the Three Mile Island Power Plant accident in 1979 and the Titan II accident of 1980.

The essential elements of an ARAC response include gathering all available data describing an incident, e.g., the source material, real-time collection of meteorological data for the area during the course of an accident, and the generation of various spatial data bases describing the region in which the accident occurred. This information is then fed into a computer model that predicts dose levels and surface contamination in the region near the release. The computer results are presented in graphical form, and then sent to emergency response personnel via telephone data links.

Several types of geographic information are important in responding to an atmospheric release, the most significant being terrain data. Winds near the earth's surface are strongly affected by the local terrain, and in order to generate credible wind field information for the accident site, ARAC's computer model requires digital topography data in addition to the meteorological data.

Terrain data can be extracted, using a system designed by one of the authors (Hoyt Walker), from one of two data bases in the ARAC system that cover the continental United States. One data base has a resolution of about 65 meters and is used for building grids of various sizes and resolutions. That data was produced by the Defense Mapping Agency. The other data base was generated from the Defense Mapping Agency data and is tailored to the needs of a real-time ARAC response. This data base is used to quickly create a standardized terrain grid at a 500 meter resolution, normally in three to five minutes.

Apart from input to ARAC's numerical model, the terrain data can also be used to create a picture of a region, thereby providing ARAC's meteorologists with valuable clues about the types of atmospheric phenomena to expect. The behavior of the atmosphere in the complex terrain of a mountain range can be very different from the behavior above a plain. The current system only produces contour maps for picturing topography. While contours provide a great deal of useful information, they have limitations and may not allow clear visualization of an area. For example, in areas of rugged terrain the contours tend to crowd each other, creating confusing displays — hills and valleys may be hard to tell apart.

In the fall of 1982, Craig Upson, a computer scientist at LLNL, began investigating alternate means of displaying terrain data. His first products were color shaded relief maps generated on a CRAY-1. Many of the ambiguities inherent in contour maps were eliminated by this presentation. First, the color of the terrain surface was assigned by altitude; for example, low elevations were green with higher elevations becoming brown. Second, the slopes of the terrain were shaded according to their angle relative to the rays of an imaginary sun. This shading effect gave a feeling of depth to the map.

Early in 1983, Mr. Upson and Pat Weidhaas, one of the authors, took on the more challenging task of providing a general graphics capability to produce three-dimensional images of a digital terrain surface from any point of observation. In contrast to the overhead views of contour and relief maps, this capability allows viewing the terrain from the side. However, introducing this generalization also introduced a significant complication: hidden surfaces. In the vertical (parallel) projection used for the relief maps, all points on the terrain surface are visible to the overhead observer. This is no longer true as one looks sideways into the terrain, e.g. a distant chain of mountains may be hidden by other hills that are located closer to the observer. Elimination of such hidden surfaces is one of the fundamental problems in three-dimensional computer graphics.

The most natural technique for solving this dilemma seemed to be the method of ray tracing, which identifies only surface points that are visible to the observer, and does not deal with invisible areas. Conceptually, one should think of the observer as a point (an "eye") somewhere above the terrain with many straight lines, or view rays, emanating from it. Some view rays might be directed upward and never hit the terrain; these rays represent the sky. Many other view rays, however, intersect the terrain surface, where the ray terminates. This point is visible to the observer. Every visible point on the terrain surface is the termination of a view ray, while hidden surfaces are never reached by a view ray. The ray tracing algorithm steps along each view ray, starting at the eye and proceeding until the intersection with the terrain is found. The distance from the observer, height of the terrain above sea level, and the shading of the surrounding slope are stored in memory for each visible point. The algorithm then repeats the process for the next view ray.

Of course the computer cannot handle the infinitely many view rays that exist mathematically, so a finite subset must be chosen. This is done by placing a window between the eye and the terrain. The window is partitioned into picture elements (pixels), similar to the raster on a television screen. A typical partition may consist of a rectangular array of 400 by 500 pixels; a high resolution image may contain 4000 by 4000 pixels. Only those view rays that pass from the eye through a pixel are dealt with; thus, for high resolution, as many as 16 million view rays may have to be processed. The program written by Upson and Weidhaas then "paints" each pixel, using the information determined by the ray tracing. Color is assigned to a pixel as a function of the height of the associated surface point. The brightness of the pixel is determined by the shading of the surface element that surrounds the point. The program can realistically simulate the effect of haze by changing the colors of the surface as a function of distance from the eye. Also, the sky above the terrain can vary in color depending on the height of the pixel above the horizon.

The images accompanying this article are examples of the current capability of the program and were produced on one of LLNL's CRAY-1 computers. In the future, additional information needed by ARAC will be included. For example, base map information — roads, streams, political boundaries — should be superimposed on the terrain surface. ARAC is also interested in using land-cover data to select the surface color according to what is really covering the terrain — forest, cities, desert. Additional realism may be achieved by generating a texture on the surface appropriate to the type of surface cover. The realistic presentation of topographic and geographic information is a highly complex and challenging endeavor offering a multitude of problems waiting for ingenious solutions.

Acknowledgement

This work was performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.

About the authors

Patrick P. Weidhaas, a native of West Berlin received his Master's degree in mathematics from the University of California at Berkeley in 1966. He joined Lawrence Livermore National Laboratory in 1972 as a Computer Scientist. His work at Livermore has involved software development of numerical models in the area of fluid mechanics and hydrodynamics. He currently works on the design of new graphics techniques for the ARAC project.

Hoyt Walker has worked at Lawrence Livermore National Laboratory since 1974, and has been involved with the ARAC project since 1977. For several years, Hoyt's work was concentrated on the atmospheric wind model used in most ARAC responses. Currently he is working on the development and implementation of large geographic data bases required by the ARAC project. Hoyt received his M.S. in Computer Science in 1978 from the University of California at Davis.

CORPORATE REGISTER

CRAY to be installed in West Berlin

Cray Research recently signed a contract with the Ministry of Science and Research in West Berlin for the purchase of a CRAY computer system. The system, which was installed in the first quarter of 1984, will be used by the Technical University, the Free University, and a number of research institutes located in West Berlin. Applications running on the system will cover a broad spectrum, including basic university research, engineering and weather modeling.

Aramco orders CRAY-1/M

The Arabian American Oil Co. (ARAMCO) placed an order recently for a CRAY-1 M/4400 computer system. The system is scheduled to be installed at ARAMCO's Exploration and Petroleum Engineering Center in Dhahran, Saudi Arabia, in the third quarter of 1984, subject to export license approval. ARAMCO's CRAY-1/M will be the first CRAY system installed in the Middle East and the eleventh CRAY system harnessed by the petroleum industry. ARAMCO's CRAY-1/M will be used in petroleum reservoir simulation.

LLNL, Sandia install CRAYs

Cray recently announced that it was installing two additional CRAY computers in Livermore, California. Lawrence Livermore National Laboratory Computing Center (LCC) received a four-million word CRAY X-MP/24 during the first quarter of this year. The X-MP joins four CRAY systems already installed at the LCC and two CRAY systems at the Magnetic Fusion Energy Computer Center at the same site.

Sandia National Laboratories, also located in Livermore, installed a CRAY-1 S/1000 computer system during the same time period. This brings the total number of CRAY systems installed for Sandia to three. Two are located at Sandia's Livermore laboratory. Sandia National Laboratories and the Lawrence Livermore National Laboratory are operated for the Department of Energy by the University of California.

French Aerospace Institute orders CRAY supercomputer

Cray Research, Inc. recently announced that it will install a CRAY supercomputer for the Office National d'Etudes et de Recherches Aerospatiales (ONERA) in Chatillon (near Paris), France, through a joint proposal with the French company Bull. A CRAY-1 S/2000 computer system will be installed at ONERA in the third quarter of 1984.

The mission of ONERA is to develop, direct and coordinate aeronautics research in cooperation with other French scientific and technical research organizations. Acquisition of the system is part of a major program to provide an advanced computational institute for aerospace research. The CRAY will be the major node of ONERA's computer network, including seven local and remote front-end stations installed at four French aerospace companies: DASSAULT, MATRA, SNECMA and SNIAS. The Bull-Multics will be locally connected to the Cray and used as a file management station.

Schlumberger orders CRAY-1/M

In February, Schlumberger Technology Corporation placed an order for a CRAY-1 M/2200 computer system. The system is scheduled to be installed at Schlumberger's research facility in Ridgefield, Connecticut, in the third quarter of 1984.

Schlumberger is engaged in the evaluation of hydrocarbon reservoirs throughout the world. The company's new CRAY-1/M will be used for the design and development of tools and techniques that implement Schlumberger's efforts to locate and characterize underground hydrocarbon reservoirs for the oil and gas industry.

U.S. Air Force to install CRAY

The U.S. Air Force has ordered a CRAY-1 M/2200 that will be installed at the Arnold Engineering Development Center in Tullahoma, Tennessee. The system will be used for engineering calculations in support of system-development tests, for test data processing, and computational fluid dynamics processing. The system will be installed during the fourth quarter of 1984.

Corrections: CRAY CHANNELS Vol. 5 No. 4, pg 24, MOVIE.BYU: NASA-Ames is not a MOVIE.BYU user. Those interested in the Cray version should contact Sara Graffunder, Cray Research, Inc.; tel.: (612) 452-6650.

USER NEWS

CRAY BLITZ: 1983 World Computer Chess Champion

It's not unusual to accept congratulations on behalf of someone else — but on behalf of a computer and a program?

That is exactly what happened last October in New York at the fourth triannual World Computer Chess Tournament. Harry Nelson and Robert Hyatt stepped ahead of 22 competitors from the United States, West Germany, Great Britain, France, Austria, Sweden, Norway, Holland and Canada, to accept the title of 1983 World Computer Chess Champion on behalf of the CRAY X-MP computer and the CRAY BLITZ chess program.

Both Nelson and Hyatt have lived with the anticipation of this win for a long time — in fact, years. Hyatt, an instructor and Chief of Systems at the University of Southern Mississippi, wrote the original version of BLITZ back in 1976. He combined his knowledge of programming with the help of a colleague who just happened to be a competent chess player, Albert Gower, and created the forerunner of the CRAY BLITZ. Harry Nelson, a computer scientist at Lawrence Livermore National Laboratories, joined the effort in 1980. He converted Hyatt's program from FORTRAN to Cray Assembly Language (CAL), streamlined it for speed, and voila! — the CRAY BLITZ chess program emerged.

What makes CRAY BLITZ a winner? Speed is definitely a factor. In game play at the World Computer Chess Tournament, a game consists of 40 moves made within two hours. This allows three minutes for each move — not a long time when the possible terminal positions number in the billions for a typical game. To speed up the search process, Nelson and Hyatt CAL-coded and revectorized significant portions of the program. The result was a BLITZ program with over 8,000 lines of assembly language operating at a search speed of 10,000 nodes per second.

But CAL-coding and revectorizing weren't enough to satisfy Nelson and Hyatt in their quest for a top-seated chess program. Just four weeks before the New York tournament, Hyatt and Nelson had the opportunity to work with a CRAY X-MP, and they scrambled to modify the CRAY BLITZ for the larger system. To them, multiprocessing capability spelled S P E E D. The two weeks prior to tourney play were spent debugging the results. So 20,000 lines of assembly language operating at 25,000 nodes per second later, CRAY BLITZ and the CRAY X-MP captured the 1983 World Chess Champion Tournament title.

Currently, the CRAY BLITZ is rated at 2,300, placing the program in the Master Chess Player rank, a higher rank than either Hyatt or Nelson has achieved. According to Nelson, even if the program operated on a par with its programmers' chess skills, its speed and calculation abilities when executing on the CRAY would give it a clear advantage.

Does this imply that the coveted titles of future World Chess Champions will become the calculated, unappreciated spoils of computer victories? Will the David Levys (today's Grand Master) fall prey to the computer's capacity to perform exhaustive searches at lightning speed?

Maybe. But not assuredly. For one thing, improvements still need to be made in the computer's ability to recognize an inevitable loss — a loss which will necessarily follow when a temporarily safe chess piece is forced to move and be captured. Nelson thinks that the ability to recognize inevitability would bring the computer one step closer to being able to exhibit rational thought. On the other hand, there is still no computer program available today which duplicates human ingenuity and intuition. In the end, our unpredictability and our willingness to play the ever-tempting hunch still distinguish us from machine — for better or for worse.

Try factoring that

We may all have done some factoring of numbers in high-school algebra and thought nothing about it since, but to mathematicians, factoring is a very hot — and significant — topic. (You may recall that to factor a number is to find two numbers whose product is equal to it; 15 is factorable into 5 x 3.)

Mathematicians in Sandia's Applied Mathematics Department at Albuquerque are in the forefront of this work worldwide. They, with the help of a Cray analyst and a CRAY, recently found some clever ways to break through the formidable barriers that have prevented factoring numbers much beyond 50 digits in length. As a result, they successfully factored a 67-digit number in December 1983.

It is hard for the layman to appreciate the significance of factoring; believe it or not, it is more than simply a matter of mathematical muscle-flexing. Data integrity, which today involves everything from automatic bank tellers to seismic analysis, ultimately depends on cryptographic-like schemes, mostly on two-key techniques. The main contender for a national standard for a two-key cryptographic algorithm, Sandia's RSA algorithm, depends on the difficulty of factoring for its cryptosecurity. In two-key cryptography, one number, or "key", is used to encode information. Once either one of those keys is factored, it is relatively simple to deduce the other key, says Gus Simmons, manager of Sandia's applied mathematics department. So it is important that the numbers used as keys be difficult to factor. Research on factoring is the only way to determine how large the numbers must be in order for data encrypted with the RSA technique to be secure.

Factoring requires fast computing. Even at a billion divisions a second, it would take 100 million years to factor a 50-digit number, and 10^{19} years, a billion times longer than the age of the universe, to factor a 71-digit number.

The algorithm used at Sandia breaks the problem down into smaller parts: it factors hundreds of thousands of smaller numbers on its way to the big solution. In doing so, it takes many time-saving shortcuts. If one of the numbers doesn't factor easily, it instantly goes on to the next. Particular numbers may have to be factored only partially; it doesn't waste time factoring them completely.

Jim Davis and Diane Holdridge, both mathematicians at Sandia, used a particular factorization algorithm, called the quadratic sieve, and found that it may be faster than any other one known. This algorithm substitutes a faster subtraction process for the multiple precision division process necessary in the other leading contender.

The big breakthrough came while a group of Sandia and CRAY mathematicians were discussing the enormous difficulties in factoring large numbers. The killing problem was that very long number components, or vectors, had to be modified many millions of times. The time it took to make each change was proportional to the length of the component even though only a small number of places on each was involved.

Tony Warnock, a Cray analyst, said that the CRAY's vector architecture made such a shortcut possible. Why was that? Well, suppose you had a roll of stamps and you wanted to mark every seventh one. You might count each stamp out — 1, 2, 3, and so on. Or you might fold the stamps into groups of seven and simply mark each seventh stamp down the edge. The latter method is obviously preferable, which resembles the way the CRAY handles the same problem.

The Sandia mathematicians first ran the basic quadratic sieve algorithm without modification on the CRAY and factored a 55-digit number in 4.4 hours. After modification for the CRAY, the same operation took only 1.8 hours.

On December 4, 1983 the world-record 67-digit number was factored in 13.7 hours. Scientists now have their sights set on a 71-digit number. For that factoring project, the Sandia mathematicians will "borrow" Los Alamos National Laboratory's CRAY X-MP computer.Gus Simmons said he expects numbers with digits in the mid-70s will soon be within the CRAY's reach.

Obviously, 80-digit cryptographic keys are no longer secure, although Simmons speculates that the keys now used by industry will be safe for about ten years. He explained, "We're trying to determine how big is big enough. Larger numbers are more secure, but their cost (time, complexity of equipment, and money) increases exponentially with the size of the numbers."

This story was based on text appearing in Sandia National Laboratories LAB NEWS, Vol. 35, No. 12.

USER NEWS

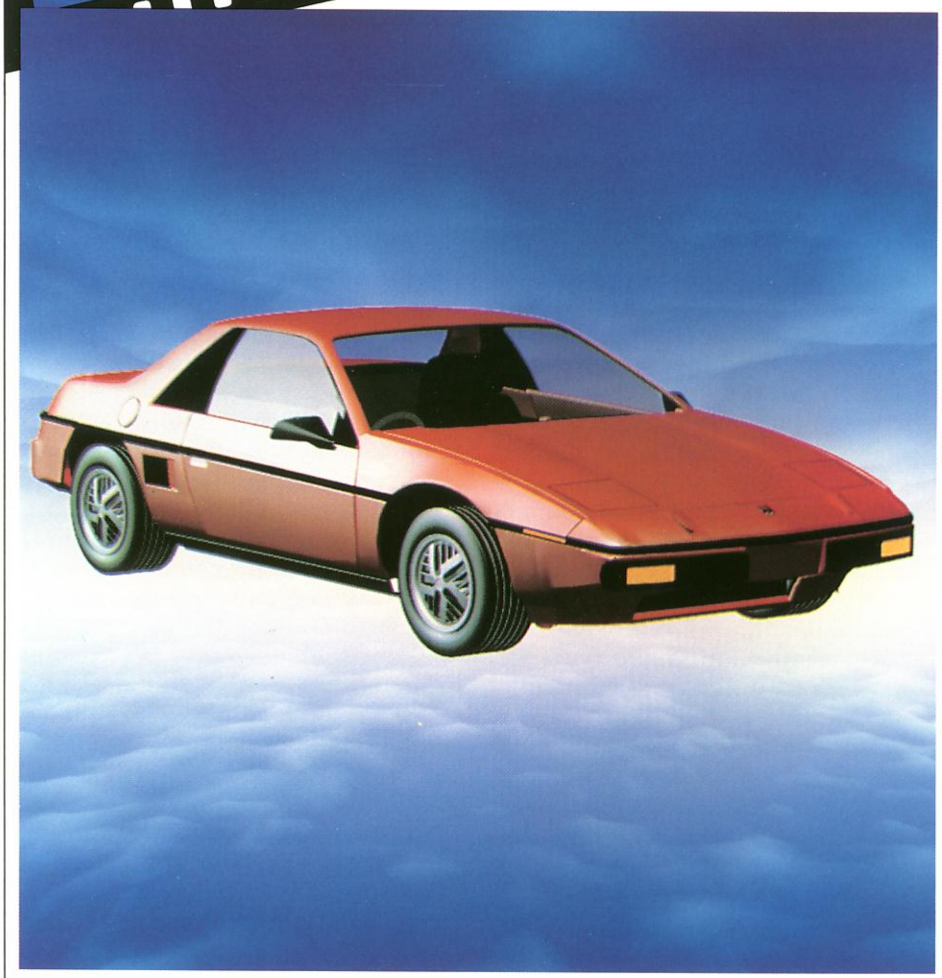

Lights, camera, action!!

Los Angeles, CA — Those who watch TV may recognize some of Digital Productions' (DP) recent commercial work illustrated here. DP is a Cray customer involved in generating sophisticated computer graphics for motion pictures, television commercials and programs, and special effects. Pontiac excitement is captured in the new computer-generated Fiero TV commercial, while Sony precision is detailed in another ad. The whimsical high-tech town (located in the midwest, they say) was created for a video game manufacturer. Credit for all photos: Digital Scene Simulation sm by Digital Productions, Los Angeles, CA © Copyright, 1984. All rights reserved.

WORLDWIDE BUSINESS CENTERS

Domestic

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

Government Relations Office
1828 L Street NW
Suite 201
Washington, DC 20036
(202)775-0155

Central Region Office
5330 Manhattan Circle, Suite F
Boulder, CO 80303
(303)499-3055
- Albuquerque
- Chicago
- Minneapolis

Eastern Region Office
11710 Beltsville Drive
Suite 500
Beltsville, MD 20705
(301) 595-5100
- Atlanta
- Boston
- Laurel
- Pittsburgh

Petroleum Region Office
5858 Westheimer, Suite 500
Houston, TX 77057
(713) 975-8998
- Dallas
- Tulsa

Western Region Office
5776 Stoneridge Mall Road
The Atruim, Suite 350
Pleasanton, CA 94566
(415) 463-2800
- Los Angeles
- Seattle

International

Cray Research France, S.A.
21 Avenue Victor Hugo
75016 Paris, France
01-501-52760

Cray Research GmbH
Perhamerstrasse 31
8000 Munich 21, West Germany
89-56014-0

Cray Research Japan, Limited
W 1154, Shin Aoyama Building
1-1 Minami-Aoyama 1-chome
Minato-Ku, Tokoyo 107, Japan
03-403-3471

Cray Research (UK) Ltd.
Cray House
London Road, Bracknell
Berkshire RG12 25Y
United Kingdom
0344-485971

CRAY RESEARCH, INC.

Corporate Headquarters
Cray Research, Inc.
608 Second Avenue South
Minneapolis, MN 55402
Tel: (612) 333-5889