

# GRAY CHANNELS

Summer 1987

Announcing new CRAY-2  
systems, enhanced memories

## FEATURE ARTICLES:

**Ultra-speed graphics for computational science**

**Graphics and remote communications at NASA Ames**

**Engineering graphics**

**Thunderstorm outflows and microbursts**

**Supercomputing in The University of Texas System**


**Fortran optimization**

## DEPARTMENTS:

**Corporate register**

**Applications in depth**

**User news**



---


## IN THIS ISSUE

---

Supercomputers often generate amounts of data too large to interpret easily when presented as pages of printed numbers. But computer graphics can be used to translate numerical data into images that are easily interpreted. By exploiting the superb pattern recognition capabilities of the human brain-eye system, computer graphics helps researchers turn the raw power of supercomputers into insight and knowledge. For this reason, computer graphics has become an integral tool for computational research.

This issue of CRAY CHANNELS presents articles on the growing role of computer graphics in supercomputing environments. It also provides a look at supercomputing in The University of Texas System, and at atmospheric research into windshear microburst, a menace to low-flying aircraft.

Dazzling pictures aside, computer graphics has acquired a critical supportive role in computational research. Increasingly, the goal in configuring graphics equipment around supercomputers will be to create interactive environments in which researchers can observe and intervene in computations as they occur. When such a situation becomes commonplace, the new age of computational science will be well established.




**On the cover** is a portion of a test board used to test the DC characteristics of 2500-gate arrays. The board was designed in-house at Cray Research and is used to test gate arrays that will be used in a Cray system under development.

# CRAY CHANNELS

A Cray Research, Inc. publication

Summer 1987



18

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

CRAY, CRAY-1, SSD, UNICOS, and CRAY CHANNELS are registered trademarks and COS, CRAY-2, CRAY X-MP, and SUPERLINK are trademarks of Cray Research, Inc. The UNICOS operating system is derived from the AT&T UNIX System V operating system. UNIX is a registered trademark of AT&T. UNICOS is also based, in part, on the Fourth Berkeley Software Distribution under license from The Regents of the University of California. Alliant FX-8 is a trademark of Alliant Computer Systems Corporation. AEGIS is a trademark and Apollo is a registered trademark of Apollo Computer Inc. CDC is a registered trademark and NOS is a product of Control Data Corporation. DEC, Q-bus, VAX, VAXstation II/GPX, VMS, and UNIBUS are trademarks of Digital Equipment Corporation. Gould is a registered trademark and Power Node is a trademark of Gould Inc. IBM is a registered trademark, DACU is a trademark, and VM and MVS are products of International Business Machines Corporation. HYPERchannel is a registered trademark of Network Systems Corporation. ProNET is a registered trademark of Proteon, Inc. Vitalink TransLAN is a trademark of Vitalink Communications Corporation. Ethernet is a registered trademark of the Xerox Corporation. IRIS is a trademark of Silicon Graphics, Inc.

Volume 9, Number 2

**Editorial Staff**  
Tina Bonetti  
Ken Jopp  
Carol Smith  
John Swenson

**Design and Production**  
Barbara Cahlander  
Eric Hanson  
Cynthia Rykken

## FEATURES

### 2 Expanded CRAY-2 series offers new technology, configuration options

The series now includes four models, three with SRAM memory technology, all with online tape support.

### 4 Ultra-speed graphics for computational science

*Karl-Heinz A. Winkler, Stephen W. Hodson, Jay W. Chalmers, Michael McGowen, and Donald E. Tolmie, Los Alamos National Laboratory, Los Alamos, New Mexico; Paul R. Woodward, University of Minnesota; Norman J. Zabusky, University of Pittsburgh*

A leading-edge effort links supercomputers with interactive animation.

### 10 Graphics workstations, supercomputing, and remote communications

*Thomas Lasinski, NASA Ames Research Center, Moffett Field, California*  
Three technologies dovetail into a premier aerospace design environment.

### 14 Engineering graphics on Cray computer systems

*Mike Long, Cray Research, Inc.*

A generous menu of graphics options is available to engineers using Cray systems.

### 18 Numerical simulation of thunderstorm outflows and microbursts

*Kelvin K. Droegemeier, University of Oklahoma, Norman, Oklahoma*  
Computer modeling reveals the dynamics of an aviation hazard.

### 24 A new kind of gusher for the Lone Star state

*Supercomputing in The University of Texas system*

This university system provides Cray power to staff and students.

### 28 Fortran optimization techniques: reducing memory bank conflicts on CRAY X-MP computer systems

*Pete Sydow, Cray Research, Inc.*

Investing the time to accommodate memory architecture pays off in performance.

## DEPARTMENTS


### 32 Corporate register

### 34 Applications in depth

### 37 User news

# Introducing the enhanced CRAY-2 series of computer systems

Once again, Cray Research has extended the price/performance and functionality of Cray computer systems. New models of the CRAY-2 system bring enhanced performance alternatives to Cray users, and tape support for the CRAY-2 systems adds yet another new dimension to this expanding product line. In addition, Cray Research has reduced prices on some models of the SSD solid-state storage device.



## Enhanced CRAY-2 systems

The CRAY-2 series of computer systems has been expanded to a line of four supercomputers offering either two or four processors, a variety of memory size options, and a choice of memory chip technologies. Three new CRAY-2 models feature static random-access memory (SRAM). SRAM offers significantly improved performance over dynamic RAM (DRAM, which was used in the original CRAY-2 system) and complements the extremely fast 4.1-nanosecond clock cycle of the CRAY-2 systems.

The new CRAY-2S/4-128 system combines four background processors with 128 million words of SRAM common memory; the new CRAY-2S/2-128 system combines two background processors with 128 million words of SRAM common memory. For users with smaller memory requirements, a third new model, the CRAY-2S/2-64, combines two processors with 64 million words of SRAM common memory.

SRAM CRAY-2 systems offer performance improvements of 15 to 25 percent over comparably configured DRAM systems. Some performance improvement is realized through faster SRAM chip speed, while additional performance gains come through reduced memory contention. For applications that experience a great deal of memory contention, the new SRAM CRAY-2S/4-128 system will yield performance up to 40 percent greater than a comparable DRAM CRAY-2 system.

The original CRAY-2/4-256 system with four processors and 256 million words of memory has been upgraded with faster DRAM memory chips. These new chips will improve overall system performance by 15 to 25 percent over previous DRAM CRAY-2 systems. The DRAM CRAY-2/2-128 and CRAY-2/4-128 systems have been discontinued.

All four CRAY-2 systems have the innovative compact architecture, liquid immersion cooling, and powerful I/O that have become trademarks of the product. All DRAM CRAY-2 systems also feature pseudobanking, which allows faster memory access and reduces interprocessor memory contention. Table 1 details the CRAY-2 series.

## CRAY-2 tape support

All CRAY-2 systems now offer online tape support through the new Cray Tape Controller (CTC). The CTC-1 tape controller supports up to eight IBM 3480 cartridge tape drives. The CTC-2 supports up to 16 tape drives, and the CTC-4 supports up to 32 tape drives.

Support for online tapes adds a new I/O capability to the powerful CRAY-2 series. On four-processor CRAY-2 systems, up to 40 peripheral devices can be configured. That includes up to 36 Cray Research DD-49 disk drives; up to eight HSX high-speed external channels for high performance graphics and other applications; and up to 16 external I/O controllers for connecting front-end systems, network adapters, or online tape units. CRAY-2 systems with two background processors (and thus half as many I/O channels) can accommodate up to 20 peripheral devices, with maximum configuration options one-half those of the four-processor CRAY-2 systems.

## Improved SSD price/performance



Cray Research has reduced prices on certain models of the SSD solid-state storage device by up to 20 percent. The SSD storage device complements CRAY X-MP system performance by offering fast access to large datasets and temporary storage for system programs. The entire line of five SSD storage devices offers from 32 million to over 512 million words of random-access secondary memory.

## Another step ahead

"We have expanded the CRAY-2 product line twice in 1987, largely because new classes of users are looking to exploit the unique capabilities of the systems," said John A. Rollwagen, chairman of Cray Research. "What started as a single, unique computer system has grown to a line of four systems. As the product line has grown, we have also steadily improved performance; the CRAY-2 systems we offer today are up to 40 percent more powerful than at their introduction only two years ago. I am pleased to see the growth and diversity of the supercomputer market, and am proud of our ability to continue to deliver effective solutions to our users."

Model	CRAY-2S/4-128	CRAY-2S/2-128	CRAY-2S/2-64	CRAY-2/4-256
Background processors	4	2	2	4
Common memory (in Mwords)	128	128	64	256
Memory type	SRAM MOS	SRAM MOS	SRAM MOS	DRAM MOS
Disk storage units	4-36	4-18	4-18	4-36
6- or 12-Mbyte/sec of channels	4-16	2-8	2-8	4-16
Magnetic tape channels	0-16	0-8	0-8	0-16
100-Mbyte/sec channels	0-8	0-4	0-4	0-8

Table 1. CRAY-2 system configuration options.



## Ultra-Speed Graphics

for  
computational  
science

Karl-Heinz A. Winkler, Stephen W. Hodson,  
Jay W. Chalmers, Michael McGowen, and  
Donald E. Tolmie, Los Alamos National Laboratory  
Paul R. Woodward, University of Minnesota  
Norman J. Zabusky, University of Pittsburgh


When the first electronic computers were introduced in the 1940s, some intuitive minds already knew that these high-speed calculating machines would have a spectacular impact on scientific research. The pioneering theorist John von Neumann, in particular, anticipated the use of computers to model physical systems and thus to overcome some of the limitations of analytical methodology. In 1946 he wrote, "Our present analytical methods seem unsuitable for the solution of the important problems arising in connection with nonlinear partial differential equations and, in fact, with virtually all types of nonlinear problems in pure mathematics." Using computers to solve these problems numerically, von Neumann realized, would advance scientific knowledge across many fields.

Von Neumann's insight has since been borne out, and most dramatically since the introduction in the mid-1970s of high-speed, large-memory vector processors, or supercomputers. Today, research is underway to realize fully von Neumann's hope that computer modeling would enable a researcher to intervene while a model is running and to exercise intuitive judgement as the calculation develops. The Ultra-Speed Graphics (USG) project at the Los Alamos National Laboratory is one such effort.

The Los Alamos Laboratory houses the most concentrated assemblage of scientific computing equipment in the

world. At the heart of the Laboratory's computing resources is a cluster of eight Cray supercomputers. By linking the supercomputers to state-of-the-art graphics, storage, and networking equipment, we hope to reduce the constraints on computational research from those of hardware performance to those of human aptitude.


In the 1940s computer scientists also saw the eventual need for adequate data output devices. Card punches and printers would be too slow to handle the volumes of data that scientific computers would generate. The need for faster output was recognized by von Neumann also, when he wrote, "The most commonplace objection against a very high speed device is that, even if its extreme speed were achievable, it would not be possible to introduce the data or to extract (and print) the results at a corresponding rate. Furthermore, that even if the results could be printed at such a rate, nobody could or would read them and understand (or interpret) them in a reasonable time." As a remedy for this potential bottleneck, he suggested the use of electronic devices such as oscilloscopes to display computed data. Recent advances in computer graphics and image processing have led to rapid and elegant data display capabilities that probably far exceed those envisioned by von Neumann. Indeed, the spirit of the USG project is grounded in his belief that obstacles limiting output and interpretation are ultimately surmountable.



Final density contours for 28 numerical experiments performed on Cray computer systems at the Los Alamos National Laboratory. Each simulation required 30 to 400 processor hours. The seven Mach numbers used (left to right columns) are  $M = 1.5, 3, 6, 12, 10^2, 10^3, 10^4$ . The four density ratios used (top to bottom rows) are  $\eta = 10^2, 10^3, 10^4, 10^5$ . These parameters were selected for their astrophysical interest. The cocoon thickness primarily depends on the density ratio  $\eta$  and secondarily on the Mach number  $M$ . In contrast, separation of the bow shock from the injected material is mostly a function of Mach number  $M$ . By inspection, one can discover the existence of an asymptotic self-similar solution for the Euler equations in the three right-most columns, that is, for  $M = 10^2, 10^3$ , and  $10^4$ . Here the beam gas is injected into the ambient medium so forcefully that the bow shock cannot significantly detach from the cocoon and propagate

way ahead of it, as in the cases with lower Mach numbers. This mechanism seems to provide a natural boundary for the self-similar shape of the cocoon and bow shock region for a given density ratio  $\eta$ . Although the actual values of the physical quantities vary widely in the region, they can be scaled approximately linearly in time to match each other. Effective Mach numbers inside the beam are limited and never larger than 30 for  $\eta = 10^2$  and 10 for  $\eta = 10^5$ , in contrast to the very high initial Mach numbers of the injected gas. This mechanism, converting a tiny fraction of the initial kinetic energy of the injected gas into heat in the first shock system it encounters, seems to provide a natural and self-consistent way of limiting the Mach numbers to be expected from observations of extragalactic radio jets.

Correction to line 3 on preceding page: The four density ratios used (top to bottom rows) are  $\eta = 10^2, 10^3, 10^4, 10^5$ . Correction to line 4 above: Mach numbers inside the beam are limited and never larger than 30 for  $\eta = 10^2$  and 10 for  $\eta = 10^5$ .



< Figure 1. Three snapshots of the evolution of a propagating pressure-matched supersonic jet. Density contours are shown. Color bar shows increasing density values from left to right. A beam of gas (grey part near the center of symmetry) is continuously admitted into an ambient gas (yellow) with an inflow velocity of Mach number  $M=100$ . The density ratio of the injected gas with respect to the ambient gas is  $\eta = P_{\text{beam}}/P_{\text{ambient}} = 10^{-2}$ . In the so-called working surface near the head of the jet, the gas that shot down the beam turns around and discharges into the cocoon (also grey) which separates the beam from the ambient medium. Because of the impact of the injected gas on the ambient gas, a bow shock (red) propagates ahead of the injected gas in the ambient medium. Internal structures in the gas cocoon region include shock systems, contact discontinuities, and vortices. The underlying numerical simulations for all the jets shown here were performed with the piecewise-parabolic method PPM for the solution of the Euler equations on several of the eight Cray supercomputers at the Laboratory. This particular jet alone required 46 hours of CPU time on a CRAY X-MP/48 system. The radius of the beam is resolved with 10 gridpoints. The length of the jet is 100 beam radii.

## The human brain-eye system

Displaying and interpreting computed scientific data invariably relies on the sophisticated data processing capabilities of the human brain-eye system. But asking researchers to interpret large volumes of data presented as columns and rows of numbers is a very inefficient use of the system. The human brain-eye system is designed to process data most efficiently by means of pattern recognition. Humans also have the ability for abstraction and generalization, enabling us to compress data by separating important from unimportant information. These capabilities are best exploited when output devices translate data into animated, digital, color-coded video displays. Through the USG project we hope to fully exploit the brain-eye system's unique capabilities by achieving an impedance match between the display rate of graphic output devices and the input capacity of the human brain-eye system.

We began this project by estimating the highest input data rate discernible to human vision. For this estimate, we need to know only the angular resolution, field-of-view, color resolution, and viewing frequency of the brain-eye system. The angular resolution, or dynamic acuity, of the brain-eye system is about 1 arc minute up to a field of view of not more than  $60^\circ$  by  $60^\circ$ . In contrast, detectability of a dot, a star twinkling against the night sky for example, is approximately 0.5 arc seconds, more than two orders of magnitude more sensitive than we actually can resolve in a dynamic context. Combining the values of angular resolution and field of view leads to a maximum requirement of 3600 by 3600 pixels for a display device. Color resolution of the human eye is limited to about 5, 6, and 7 bits for each of the three (red, green, blue) color channels, resulting in a total depth requirement of at least 18 bits per pixel for a display device. Viewing frequency of the eye is about 8 to 12 Hz. At higher display rates people gradually lose the ability to distinguish individual frames, and perceive continuous motion.

Putting the numbers together provides an estimate of approximately 2.8 Gbits/sec for the maximum input data rate of the human brain-eye system. Requiring a movie-like display rate would double the baud rate to 5.6 Gbaud. On the other hand, working within a screen limitation of 2048 by 2048 pixels results in a baud rate of 0.9 Gbaud,



Figure 2. Space time diagram of density contours near the central regions of the beam for the  $M=100$ ,  $\eta=10^{-2}$  supersonic jet. Color bar is similar to the one in Figure 1. We see that the bow shock (red) never propagates far ahead of the working surface or terminating shock system inside the beam. Inside the upper left part of the diagram, blue and white colors exhibit strong rarefaction waves. Sharp transitions to black indicate the presence of multiple shock systems. By inspection we find almost stationary, forward- and backward-propagating shock systems.

very close to the performance of the 100-Mbyte/sec (or 0.8 Gbaud) Cray HSX channel, and definitely within the reach of current technology.

By contrast, current display times (rates) of equipment commonly used in scientific laboratories range from two hours (1.2 Kbits/sec) at the low end, to one minute (150 Kbits/sec) at the high end for one 1024 by 1024 by 8-bit image. Even the best of these times leads to excessive waiting on the part of the researcher. Consequently, the working environment of a numerical continuum physics researcher is structured along the lines of one in which a theoretician is allowed to write for one minute, then forced to stop for five, allowed to write for one, stop for five, and so on. Much talent and time is wasted in such an environment. And much scientific data is left unanalyzed.

## A productivity revolution

Our belief is that, if researchers can access an interactive graphics environment that fully exploits the capabilities of the human brain-eye system, the result will be an unprecedented, revolutionary gain in individual user productivity. The impact will be felt in areas such as numerical gas dynamics, radiation hydrodynamics, and other related and unrelated fields such as chemistry, biology, and medicine.

The particular physics application being used to benchmark the project's progress involves simulating a pressure-matched, supersonic gas jet penetrating an ambient environment that is much denser than the jet, a situation encountered in certain astrophysical phenomena. As the speed and quality of graphic output devices begin to approach the input capacity of the human brain-eye system, a researcher can maintain a continuous train of thought concerning the physics issues involved in the problem under investigation. Previous overriding concerns about the usefulness of time-consuming raster display techniques become irrelevant. A substantially more productive research mode and a better understanding of physics problems have already resulted from the current, preliminary build-up of the USG project.

## Evolution of the USG configuration

With the present configuration of the USG project, we already have obtained an immediate performance improvement of three to 4 orders of magnitude over widely used data rates of 9.6 Kbits/sec, connecting Cray supercomputers to display devices for raster color images. This approach enables us to display 800 raster images of 1024 by 1024 by 8 bits at up to 8 frames per second or 3200 images of 512 by 512 by 8 bits at up to 30 frames per second. We achieved this data rate by separating data generation on centrally located supercomputers from the truly interactive visualization process on a locally available and controllable image support processor. In this context we can consider a program running on a supercomputer merely as an information pump.

Our image support processor consists of the dual-processor Gould PowerNode 9080, the Gould image array processor IP9516, a set of four real-time disks and 16 normal system disks of about 12-Gbyte storage capacity. The

image support processor is connected to a CRAY-1 worker machine over a 300-foot cable in full duplex mode, using two sets of high-speed port interfaces. The image support processor is used primarily to receive large sets of compressed solution numbers from the CRAY-1 system, to turn them into individual images at a rate of up to 100,000 images per day for 20 different primary and secondary flow variables, and to display the images as movies on the digital display device. Storing the images in digital form — and therefore completely decoupled from a particular color representation — allows for tremendous flexibility in digging out the flow structures that are hidden in the numbers. Also, compared to the NTSC TV standard, digital display devices typically feature better display quality of a factor of 15 to 40.

The main tools for code development are a variety of scientific workstations, which are connected to the image support processor via an Ethernet link. The image support processor, with its substantial disk capacity, also acts as a file server for these workstations. UNIX is our operating system of choice on all machines. The network has been complemented with the ProNET-80 token ring network, which operates at a burst data rate of 80 Mbits/sec. We are hopeful that its performance will eventually allow us to display a 1024 by 1024 by 8-bit image in one second on the workstations, compared to the approximately ten seconds required via the Ethernet network. The interface from a workstation to the ProNET-80, not the network itself, is the limiting factor.

An important addition needed to further increase user productivity will be the acquisition of a high-capacity, high-speed data storage and retrieval system. With the appearance of new tape drives, it is possible to combine conventional, low-cost, half-inch video tape (with up to 5 Gbytes storage capacity per tape) with high-speed (3 Mbyte/sec) data rates, and thus to load and off-load data from the image support system at the existing disk controller speeds. This addition will remove the remaining bottleneck in the system by providing an impedance match between on-line and off-line storage. This will allow us to deal with almost limitless databases. Currently, optical digital laser disks don't seem to fulfill our storage needs because of low data rates and high storage costs.

## Future development

Many issues remain unresolved with the configuration just described. For example, we have yet to arrive at an interactive proficiency that would realize von Neumann's ideal, so that a researcher watching a simulation "can then intervene whenever he sees fit." As we have already seen, approaching the human physiological visualization limit requires another order of magnitude performance increase over our current project.

Another important reason exists for achieving these higher data rates soon. Class 7 supercomputers, with billions of words of random-access memory and performance two orders of magnitude faster than a CRAY-1 CPU, will arrive in research laboratories in the coming years. Such systems will easily saturate any computer network now in existence. Unless we upgrade the network surrounding

supercomputers accordingly it will be impossible to derive optimal benefit from the increased processing power of the new machines. For these and related reasons, we find it is now time to prepare for the final phase of the USG project.

This phase will require a new facility within the laboratory. Essential to its success is the availability of an image support processor with one Gbyte of random-access memory, a 300-Mbyte/sec system bus, and with sufficient parallel and vector performance to absorb and transform into images the data from multiple 100-Mbyte/sec channels of class 6 or 7 supercomputers. As real-timedisks, with many parallel read/write heads, are very complicated and limited in functionality, it is desirable and we hope soon will be cost effective to instead use massive amounts of readily available random-access memory. This not only will permit the playback of longer sequences of digital movies at significantly higher resolution, but also will enhance the functionality of the entire system for many other purposes. Anticipating the future availability of display devices with the data rates needed for the display of time-dependent three-dimensional images is probably realistic. Incorporating a wall-size high-resolution flat panel display into such an environment would allow for visually cross referencing many different numerical solutions in parameter space.

It is quite natural to follow our train of thought and apply it not only to the interfacing of two computers, but also to the whole computer networking problem. The current Laboratory Integrated Computing Network (ICN) connects multiple supercomputers to a large fileserver, a batch server, multiple terminal servers, network gateways, and other general-purpose systems. Machines communicate through the High Speed Parallel Interface at data rates of up to 50 Mbit/sec. Approximately 4000 terminals with data rates of up to 333 Kbit/sec are available to the user community. The problem we face is how to extend the capabilities, demonstrated within the USG project, to that community.


The diagram at right shows a spin-off of the conceptual work done so far within the USG project. Our proposed network has a 20-Gbit/sec crossbar switch called CP\* (sea-pea-star) at its core. It connects the Cray system and other supercomputers, massively parallel machines for example, by way of an Ultra-High-Speed Parallel Interface (UHSPI). Data is stored in a 32-Gbyte network storage system. Workstations and display devices are connected to the central CP\* by links through other CP's. The Ultra Keyboard Graphics Concentrators (UKGC) connect individual displays and workstations with data rates ranging from 9600 bits/sec to 640 Mbit/sec.

The many users at the Laboratory require the simultaneous transfer of data and pictures to multiple displays with speeds approaching the channel limits. Thus we need independent data paths to provide smooth data presentation. CP\* is a crossbar switch that meets the requirement for multiple, noninterfering, simultaneous data transfers. If the basic structure of a crossbar is compared to a bus, ring, or backplane, one can see that the crossbar allows simultaneous transfers, while only a single path is available in the other configurations.

Crossbar switches often are avoided in computer networks because of wiring problems and the large number of circuit elements required for each crosspoint. Very Large Scale Integrated (VLSI) circuits have reduced the physical complexity of this problem. CP\* can be built with a central core composed of 48 crosspoint switch VLSI circuits, which interconnect 16 UHSPIs. The Fairchild F100716 16 by 16 crosspoint switch is capable of data rates in excess of 100 MHz. By switching multiple 32-bit lines in parallel, the core of CP\* could achieve burst transfer rates on a single UHSPI channel of 3.2 Gbits/sec. With each of 16 channels sending and receiving simultaneously, total system throughput could reach 100 Gbits/sec. Within these theoretical limits, operating the core at 25 MHz seems realistic. It is worth noting that another chip, the Fairchild CMOS FGC6000 gate array, interconnects 32 UHSPIs with data rates of 50 MHz. Even higher speeds can be obtained with the GaAs 1-GHz gate array, 14 by 16 crosspoint switch from Gain Electronics, New Jersey. Total network throughput of 500 Gbits/sec seems achievable.

We also foresee special interfaces on CP\*. The UKGCs could supply 80-Mbit/sec to 640-Mbit/sec service on fiber optic cables to displays and workstations. Other interfaces will provide multiplexing and lower-speed capability.

The UHSPI network channel will be full duplex and symmetrical, use a flow-control handshake signal for each 245-word parcel, and have two 32-bit data paths. The 640-Mbit/sec rates for the channel are based on a 32-bit-wide data path with 20-MHz switching. The control portion of the UHSPI will be microsequencer based. The types of functions to be executed include examining the header, checking security fields, checking for data errors and assembling/disassembling packets between the host computer and the rest of the network. In addition, higher levels of protocols can be implemented in the UHSPI, such as TCP/IP or a real-time graphics protocol. The intent is to off-load as much of the protocol work from the hosts as possible. A 50-Mbit/sec data rate from a low-speed I/O channel of a CRAY X-MP/416 system directly into a frame buffer and display device already has been realized at Los Alamos. We plan to connect a prototype UHSPI directly



*Proposed upgrade concept for ICN-network at Los Alamos, utilizing a crossbar switch CP\* with 20-Gbit/sec throughput at its heart.*

to the CRAY X-MP HSX channel. The UHSPI also can be built with commercially available chips.

Critical to the idea that the UHSPI can indeed sustain throughput at the indicated rates was the realization that the unique architectures of high-resolution image capture systems are directly applicable to the computer networking problem. A packet network interface faces the same bandwidth bottleneck as a graphics system. Large amounts of data must be moved in and out of memory while the CPU processes the network headers. By associating network packets with lines of video the problem is very similar to the graphics system except less demanding, because once a packet is sent it typically will not be necessary to output it again. It is worthwhile to point out a specific component of the UHSPI that provides a very elegant and efficient method for input/output. Although it was initially developed for a graphics frame buffer, the multi-ported video ram is ideally suited for full-duplex, synchronous data entry. In addition, it allows concurrent CPU access with minimal bus contention. The video ram architecture eliminates the bus contention by decoupling the video data bus from the processor data bus. A typical video ram con-

sists of a 64K by 1 DRAM interfaced to an on-chip 256-bit parallel-load shift register. The shift register is loaded with a complete row of the DRAM array and can then be accessed serially at frequencies of up to 50 MHz while the DRAM array is accessed simultaneously and asynchronously by the processor via the random access port. Future research should deal with the architectural concept of adding increased functionality to the memory components themselves. Put another way, the creation of smart memories may reduce the need for direct memory access.

Such a network, with a total bandwidth of 20 Gbaud, would increase the intersystem throughput by about three orders of magnitude, and in particular would allow for the specialized equipment of the USG project to be simultaneously connected to numerical experiments running on more than one supercomputer. In fact, the proposed network would be so fast, that for high complexity algorithms, that is, those with many algebraic operations per memory fetch, the entire computational resource in the network could be used in a parallel processing mode across multiple machines. Such a system would allow us to deal with otherwise intractable problems. It certainly would be




Figure 3. For the same  $M=100$ ,  $\eta=10^{-2}$  supersonic jet, four further physical quantities are shown. Pressure (upper left) shows very nicely the alternating rarefaction waves and shock systems inside the beam. However, pressure varies smoothly across the contact discontinuity, separating the ambient gas from the cocoon, where density discontinuously changes by two orders of magnitude (see Figure 1). Divergence of velocity contours (lower left) shows the filigree network of shock systems (yellow, red, and dark colors) and the existence of extensive rarefaction waves (white and blue colors). Azimuthal component of curl of velocity, also called vorticity (lower right), exhibits regions of strong shear and demonstrates the turbulent character of the flow. Absolute value of velocity (upper right) shows the forward-moving gas inside the beam (yellow) and the return flow in narrow regions in the cocoon.

of great importance in times of a national disaster, such as Three-Mile Island or Chernobyl, when the entire computational power of the Los Alamos National Laboratory could be utilized to predict future developments inside or outside the reactor. For a successful and timely completion of the proposed network and the USG project, collaboration between the Los Alamos Laboratory, universities, and industry is essential. We have to encourage product developments and foster the exchange of ideas and system specifications among vendors, because no single vendor offers the system solution we seek. Specifications for the interfaces in the network should be openly available so that future machines from a variety of vendors can be easily integrated into the system. Unless these conditions are met, completion of the project will be difficult, if not impossible. One sign of a good working relationship with a university is the cloning of our present system configuration by P.R. Woodward at the Minnesota Supercomputer Institute at the University of Minnesota. An extensive exchange of specially developed software has been mutually beneficial.

## Benefits and applications

Of course, the central question is whether the USG project's payoffs will justify its human and material costs. We have little doubt they will. Already, with the current system configuration we have made some discoveries that we had not anticipated. We are more aware of the amazing accommodation of which the human brain-eye system is capable, and we are using the USG equipment in unexpected ways. For example, we find that by watching a 20- or 30-second movie of a rather complex phenomenon over and over again (perhaps hundreds or thousands of times), we are able to memorize individual structures of the flow and accumulate cross references. The amazing thing is that in any one viewing one never seems to be able to see all the features. Only in the mind's eye does the whole puzzle come together, which should not be surprising. One can focus on specific features of a display or take a more relaxed view, and so fully exploit the fantastic dynamic range of the human eye, free of any software development costs.

In another learning experience we were able to recognize new physical features and eliminate subtle numerical errors in a computation that was several years old. Without the digital movie capability of our system, those effects would never have attracted our attention and would have remained unnoticed. Because we now can see many more details of the simulations, new confidence is developing in the validity of our gas dynamical experiments. We hope to develop the numerical experimental art to such a high degree that we can separate physical and numerical effects completely in fairly complicated flows.

Another important implication of this work is that it makes available, in the form of color images, the entire database of a numerical experiment. As a result, other researchers will better be able to verify or refute a given interpretation of a numerical experiment. The end result of numerical experimentation should be insight and knowledge, not numbers. Through the immensely powerful new visualization tools of our project, we can do things

in a routine fashion that we could not even dream of doing before.

A system such as that described here, when combined with adequate temporary and archival storage capability and interfaced to a class 7 supercomputer, will for the first time permit the execution of real-time numerical experiments along the lines conceived by John von Neumann for complex physical systems. As a result, a new world of numerical experimental science will open up. Researchers will be able to produce computations as food for thought and further analysis, without having to worry about how long it will take to see a particular result once it has been computed. A lively dialogue between individual researchers and the computational model in the machine can be maintained. Extraordinary convenience and simplicity in performing computational tasks will not support laziness but will be a prerequisite for creativity and productivity. □

## Acknowledgments

*Getting a project such as this off the ground in a matter of months requires the help of many dedicated people. We acknowledge the help and support of Robert W. Selden for initiating the funding, E. Staley Hadden, Lawrence B. Warner, Suzanne W. Peterson, and Carla R. Walsh for providing operational support, Eldon J. Linnebur for staff support, William Spack for taking care of DOE-related matters, Norman R. Morse for financial and technical support from the Computing and Communications Division at the Los Alamos National Laboratory, and Hoyt E. Hart of Cray Research, Inc. for customer support.*

## About the authors

Karl-Heinz A. Winkler, Stephen W. Hodson, Jay W. Chalmers, Michael McGowen, and Donald E. Tolmie are researchers at the Los Alamos National Laboratory. Paul A. Woodward is a professor of astronomy at the University of Minnesota. Norman J. Zabusky is a professor of mathematics and engineering at the University of Pittsburgh.

## References

1. Goldstine, H. H., and J. von Neumann, "On the Principles of Large Scale Computing Machines," in *John von Neumann, Collected Works*, Vol. V, Pergamon Press, 1963.
2. Thompson, P., "Visual Perception: An Intelligent System with Limited Bandwidth," in *Fundamentals of Human-Computer Interaction*, A. Monk, editor, Academic Press, 1984.
3. Winkler, K.-H. A., M.L. Norman, and J.L. Norton, "On the Characteristics of a Numerical Fluid Dynamics Simulator," in *Supercomputers: Algorithms, Architectures, and Scientific Computation*, F.A. Madsen and T. Tajima, editors, University of Texas Press, Austin, 1986.
4. Winkler, K.-H. A., and M.L. Norman, "MUNACOLOR: Understanding High-Resolution Gas Dynamical Simulations Through Color Graphics" in *Astrophysical Radiation Hydrodynamics*, Reidel, 1986.
5. Woodward, P. R., "PPM: Piecewise-Parabolic Methods for Astrophysical Fluid Dynamics," in *Astrophysical Radiation Hydrodynamics*, K.-H. A. Winkler, and M. L. Norman, editors, Reidel, 1986.
6. Zabusky, N. J., "Computational Synergetics," *Physics Today*, July 1984

# Graphics workstations, supercomputing, and remote communications

Thomas Lasinski, NASA Ames Research Center, Moffett Field, California

The NASA Ames Research Center has long been at the forefront of aerodynamic analysis. The research problems addressed by researchers at the center have grown in size and complexity over the years, as have the researchers' analytical tools. The Numerical Aerodynamic Simulation (NAS) program is the latest advance in computational technology at the center. The NAS environment consists of more than 30 Silicon Graphics Inc. IRIS graphics workstations. These are networked to a CRAY-2 computer system and other computers via HYPERchannel and Ethernet networks. This article discusses the current approach to computational and visualization systems, as well as projected plans for graphics workstations, supercomputers, and remote communications in the NAS program over the next few years.




Particle traces over a fighter-like configuration calculated on the CRAY-2 computer system at the NASA Ames Research Center.

## Overview of the NAS system

NASA established the NAS program to achieve three principal goals. The first goal of the program is to provide continuing national leadership in computational fluid dynamics (CFD). The program's second goal is to act as an agency pathfinder in the integration and use of state-of-the-art computer hardware and software technologies. The third goal is to provide a strong research tool for NASA's Office of Aeronautics and Space Technology.

The CRAY-2 computer system and Silicon Graphics IRIS workstations are critical parts of the NAS network. The CRAY-2 computer system has a 4.1-nsec clock cycle and 256 million 64-bit words (2000 million bytes) of common memory. A growing number of remote sites connect to the NAS system through terrestrial lines with transfer rates ranging from 56 to 224 Kbits/sec. NASA's Langley and Lewis Research Centers have had IRIS workstations connected to the NAS system for the last two years.



## Workstation requirements

The NAS graphics workstations were chosen through competitive procurement in the spring of 1984 with a few basic requirements. The workstations had to have good, direct communications to the CRAY-2 system, stand-alone CPUs with performance equivalent to that of a VAX 780, and large disk storage capacity. They also had to be able to provide for real-time dynamic graphics, and video animation with tools to produce journal-quality output.

These requirements are common to all potential workstation users. The potential for real-time, dynamic graphics was of particular interest in that it distinguished the workstation required by NAS from most commercial workstations available in 1984. While bit-map graphics was a common feature among workstations, the requirement called for the capability to manipulate graphic objects in near-real time. Few vendors today can deliver this capability in a workstation environment.

Most NAS users are working in the area of CFD. This discipline attempts to solve the differential equations of fluid flow numerically to understand flows over actual three-dimensional objects. Historically, fluid dynamicists have studied flows with photographs using a variety of techniques to enhance the salient features of the flow. The use of computer graphics is a natural step for the CFD scientist. The supercomputer serves as a "wind tunnel," performing numerical simulations while workstations with dynamic graphics capability provide numerical flow visualization. Many examples of such applications exist.<sup>1</sup>

## Present workstation hardware

The graphics workstation used in the NAS project is the IRIS (Integrated Raster Imaging System) 2500 Turbo. A Motorola 68020 chip and Weitek chip set give the workstation floating-point capability greater than that of a VAX 780 with floating-point-assist.

In addition to the usual complement of RS232 and Ethernet ports, the IRIS workstation also supports a HYPERchannel interface. This is simply a multibus board made by IKON which permits connection to a Network Systems Corporation A400 HYPERchannel adapter. This interface allows the IRIS workstation to communicate directly with the CRAY-2 system.

## The network environment

The three principal networks associated with the NAS system are Ethernet, HYPERchannel, and the long-haul network, NASnet. The local area network based on Ethernet links all computers with the exception of the CRAY-2 computer system. The CRAY-2 system is linked to the network by HYPERchannel. Both of these local-access networks connect computers located within the NAS facility and computers located in other buildings at Ames. These campus-wide connections are made with fiber optics technology, repeaters, and bridge connections for the Ethernet and HYPERchannel networks.

NASnet provides Ethernet access for more than 70 percent of NAS users. This network connects the local NAS

Ethernet to remote site Ethernets through Vitalink Trans-LAN communication bridges. The necessary terrestrial communication links are provided by NASA's newly implemented Program Support Communication Network. This long-haul network has been in a prototype mode for more than two years. In particular, IRIS workstations located at NASA's Langley and Lewis Research Centers have Ethernet access to the NAS network.

All the computers at NAS run TCP/IP communication protocols. In addition, Berkeley-style networking commands are supported on all systems, which greatly facilitates the addition of new computers to the network. The use of internet protocols provides remote users transparent access to the CRAY-2 computer system through the Ethernet and HYPERchannel gateways.

### The CRAY-2 system as a graphics coprocessor

An application program known as RIP (Remote Interactive Particle tracer) was written at NASA Ames to make use of the large-scale scientific computing tools described above. The RIP application involves two processes, one on the IRIS workstation and the other on the CRAY-2 computer system. These processes communicate over the Ethernet or HYPERchannel networks using TCP/IP protocols. The process on the IRIS workstation controls a graphics database for the model under study, such as a space shuttle or airplane design. The IRIS workstation can rotate and

zoom through this database in near-real time, independent of the CRAY-2 system. The IRIS workstation also provides the principal interface to the user in this application. The process on the CRAY-2 system controls the solution database for the model. This database is generated by solving partial differential equations that describe fluid flow. It is typically 50 Mbytes in size and can require from 1 to 20 hours of processing time to calculate on the CRAY-2 system.

In the RIP system, the CFD scientist uses a mouse to indicate a point at which a test particle is to be released on or near the object. The solution database has the information to show how the test particle flows past the object. When the mouse is clicked, the location of the point is sent to the CRAY-2 system, which then calculates the particle flow from the solution database. This technique is referred to as *particle path tracing*. The search procedure is very CPU-intensive but in roughly one-tenth of a second, the CRAY-2 system returns to the IRIS workstation a series of perhaps 400 short vectors that geometrically define the particle trajectory.

In a few minutes, the CFD user can define and build a visualization of the flow field over the model. At any point, the user can rotate the model and the traces to study the flow from different orientations. In this way, the user can correlate one area of flow with another. Thus, RIP provides an interactive tool for visualizing and exploring flow



field calculation results. Since RIP sends display list information, not image data, to the IRIS workstation, researchers view and manipulate geometry and traces independently of the supercomputer.

### Remote usage

RIP can be used effectively by remote users since it does not send huge amounts of data to the user's workstation. One trace is typically made up of about 400 vectors. The tip of each vector is defined by 12 bytes ( $x, y, z$ ). Another 3 to 4 bytes are needed for control and color information. At about 15 bytes per vector, one trace amounts to approximately 6 Kbytes of data that must be transmitted. This takes about one second at 56 Kbits/sec and one-thirtieth of a second at 1.5 Mbits/sec. An interactive response is one-fourth second or faster for most users.

The problem with particle traces is that they reveal only a fraction of the information available in a simulation. In two dimensions, CFD scientists often use contour plots of density or Mach number to gain more detailed insight into the dynamics of a flow field. Such plots typically are on the order of 40 to 100 traces. Users also can stack two-dimensional contour plots to obtain some sense of the three-dimensional dynamics of a flow field.

Eventually more sophisticated visualization techniques will be developed. For now, one can reasonably estimate that the equivalent of about 1000 traces will be involved. Generally, CFD users would like to view their data with single graphics frames corresponding to about 1 Mbyte of graphics data. Such views would require 2.5 minutes to transmit at 56 Kbits/sec or about 5 seconds at 1.5 Mbits/sec. These image densities (1 Mbyte) are much more typical of what the CFD scientist would like to study interactively. In this regard RIP is a compromise. The CFD scientist can get some very useful information from the particle traces, but the technique may be of little or no use to the visualization of other physical problems.

CFD scientists also would like to study unsteady or time-dependent flows. As a dramatic example of unsteady flow, consider the problem of store-separation from a highly maneuverable aircraft. In many cases, new store designs or new tactics lead to the store destroying a portion of the aircraft wing. The experimental study of such phenomena is very expensive. The time-dependent nature of this flow is clear. Visualizing it would require hundreds of timesteps, with each frame requiring about one Mbit of data — as discussed above. The relative motion of various components of this flow is also a critical feature of such a problem. Unsteady flow applications require animation on the order of 10 frames per second and data rates on the order of 10 Mbytes/sec.

As in other applications, the CFD user also wants to be able to rotate, zoom, and pan through these images either "on the fly" or in a temporarily paused state. The work of Karl-Heinz Winkler and colleagues<sup>2</sup> at Los Alamos National Laboratory illustrates the power of animating CFD results at rates of 60 Mbits/sec. In this environment, graphics images are stored on magnetic disk and transferred as quickly as possible to a frame buffer, providing animation of about

15 frames per second. This approach, however, does not permit the user to interact with the visualization. New external bus technology gives the promise of streaming data from the CRAY-2 system itself to frame buffers at the rate of 100 Mbytes/sec. With appropriate software, interactive use could then be controlled by the CRAY-2 system.

### T1 networks and beyond

The NAS experience demonstrates that applications such as RIP can be quite successful in wide area networks running on T1 (a channel with a speed of 1.5 Mbits/sec). Although faster rates are always desirable, a truly robust T1 network would not only handle RIP-style applications, but also permit sizable solution files to be shipped to the remote user's site for more intensive study. Scientific users want to move large (50 Mbyte) files routinely to and from their home sites. The impression of many researchers using existing 56-Kbit/sec networks is that they work fine for handling electronic mail but are inadequate for moving files.

While T1 networks may be suitable for the next few years, there are several factors pushing for greater capabilities. As graphics workstations become more common, researchers will develop new techniques to visualize scientific applications. These techniques will require wide-area rates in excess of 1 Mbyte/sec. Graphics workstations in the next few years will have significantly greater hardware capabilities. Processing speed will go from the 1 MIPS range to 10 MIPS or better. Instead of 100,000 vector transformations per second, systems will exist that will transform 300,000 polygons per second including z-buffering. The advent of 4-Mbit RAM chips will bring display memories of 256-bit planes. It is clear that future graphics workstations will have the fast buses and large memories needed to handle intensive supercomputer output. Unresolved, however, is whether future wide-area networks will have the hardware and software capability to sustain communications between future supercomputers and graphics workstations. □

---

### About the author

Thomas Lasinski received his Ph.D. degree in physics from the University of Chicago in 1971. He was a research associate in particle physics at the Lawrence Berkeley Laboratory from 1970 to 1974, and a research associate in experimental particle physics at the Stanford Linear Accelerator Center from 1974 to 1979. In 1979 Lasinski joined NASA and since has worked at the NASA Ames Research Center as an aerospace engineer in the applied computational aerodynamics branch and as an electronics engineer in the systems engineering branch of the NAS projects office. He presently is manager of the workstation subsystem at the NAS systems development branch.

### References

1. Bunning, P., and J. Steger, "Graphics and Flow Visualization in Computational Fluid Dynamics," AIAA-85-1507-CP, AIAA 7th Computational Fluid Dynamics Conference, July 15-17, 1985.
2. Winkler, Karl-Heinz A., J. W. Chalmers, S. W. Hodson, R. L. Judd, M. McGowen, D. E. Tolmie, "The Ultra-Speed Graphics Project of the Advanced Computing Initiative at Los Alamos National Laboratory," Los Alamos National Laboratory, 1986.


# ENGINEERING GRAPHICS ON CRAY COMPUTER SYSTEMS

*Mike Long, Cray Research, Inc.*

Cray supercomputers are reshaping the engineering enterprise in industrial laboratories around the world. Using Cray systems, engineers can test ideas and optimize designs faster and more cost-effectively than would be possible by any conventional means. As a result, large-scale engineering projects now are being pursued that previously were unapproachable. Traditional kinds of engineering projects also are reaching completion sooner and at less expense than they would if developed with traditional engineering tools. By solving large, complex problems in aerodynamics and hydrodynamics, thermal and structural analysis, and design optimization, supercomputers enable engineers to apply their expertise and skill in the most efficient manner possible.

But progress in this area relies to a great extent on a related technology: computer graphics. The unsurpassed calculating capacity that Cray systems provide can generate quantities of data beyond the ability of researchers to analyze quickly. By translating that data into images, however, users can more easily interpret and identify what is significant in their computed results. Computer graphics provide an effective way to reveal information that otherwise would remain buried in an avalanche of numbers.

To most effectively translate supercomputer data into useful information, computer graphics environments must satisfy three criteria: they must be fast, able to handle large-scale models, and easy to use. Many strategies exist for implementing computer graphics in ways that meet these criteria. Following are descriptions of some of the methods available to Cray computer users.



*Example of engineering graphics generated on a Cray system. Shaded automobile body display courtesy Steve Westin, Ford Motor Company.*

## Graphics using the COS operating system

Users of Cray computer systems that run under the Cray operating system COS have several options available for converting their results into graphics. These include the use of batch and interactive stations, use of Cray Research's SUPERLINK integrated support processor, and use of dedicated graphics channels.

### Batch graphics

Cray Research Station software connects front-end computers with Cray systems and provides several methods for transferring graphics data from the COS environment to a graphics device in a batch mode. Two of these methods are *metafile* and *direct dispose* graphics.

The metafile method is a simple but effective way to transfer graphics. It involves first creating on a Cray system a generic graphics file, that is, a dataset containing pure graphics data. This file is then disposed to the front-end computer, and reformatted with a program running on the front-end computer. The data is reformatted for a particular graphics device, and then forwarded to that device for display. Because this method removes much of the interactive processing needed for engineering and scientific research, it is best reserved for batch-type operations. And although the method has other drawbacks, such as its impact on front-end disk space (graphics files contain tremendous amounts of data) and its consumption of front-end computational cycles, it often is used for creating device-independent graphics that are processed on the front-end system.


Direct dispose is an efficient method that transfers graphics data directly to a graphics device, providing a tremendous tool for engineering analysis. With this method, the graphics data are intermingled with device instructions and formatted in a sequence acceptable to the graphics device connected to the Cray system. The combined data are transferred directly to the graphics device. The direct dispose method is available to interactive and batch users and is well-suited for large amounts of raster data. The dispose of the graphics data can be activated internally from the application program. Users running interactively can access the batch portion of the station software, allowing the rapid transfer of many images at the user's command.

Direct dispose graphics is available via several front-end stations. It is supported on the DEC VAX/VMS, IBM/MVS, and IBM/VM operating systems and on the Apollo stations. In most cases the graphics data are transferred directly to the device, bypassing the front-end disk, and thereby retaining performance that would be lost if the data were written to and later read from disk. Using a Ramtek 9400 graphics device connected by direct memory access to a VAX 11/785 computer, station overhead is minimal and overall display speed is slightly less than that of a dedicated connection.

Using the direct dispose method with the MVS and VM operating systems can be accomplished by several means, including disposing data from a Cray system to an IBM system, in which case the data are sent to an Auscom 8911A channel interface which converts IBM Block Mux channel data to a DEC Q-bus or UNIBUS adapter connected to a frame buffer. An alternate to the Auscom channel interface is a DACU converter manufactured by IBM.

### Interactive graphics

The ability to run programs interactively, which the interactive stations provide, is a tremendous asset to the engineering community. Users can log directly onto a Cray system and interact directly with the user program.



Interactive graphics allows users to view their output in as close to real time as possible. The use of interactive stations provides a flexible method for routing graphics from a Cray system to a graphics device. Interactive stations can separate graphics data from user prompts, sending the graphics data to an alternate graphics terminal and the prompts to the user's terminal, or graphics and text can be sent to the same terminal.

By designating the Fortran logical unit number as a transparent dataset, the station bypasses any terminal drivers, and sends data byte-for-byte to the terminal. This method works especially well for vector images where response must be immediate or the quantity of data is small. For larger amounts of data, typically raster data, users can use the direct dispose option of the batch station, even though they may be running interactively.

Interactive station capabilities are available through the following operating systems: VAX/VMS, IBM/MVS, IBM/VM, Cyber/NOS, Apollo/AEGIS, and on Sun, Hewlett-Packard and other systems.

#### SUPERLINK graphics

Graphics via Cray Research's SUPERLINK integrated support processor offers an alternative to the graphics facilities provided by the interactive and batch stations of the MVS operating system. SUPERLINK transports large amounts of sequential data between the IBM/MVS and Cray/COS environments, and provides the capability to perform memory-to-memory transfers. When this tool is used, graphics data are computed on a Cray system and downloaded to an IBM mainframe via SUPERLINK. The

graphics data then are transferred from IBM memory to a graphics device such as an IBM 5080.

Typically, the methods described above are the options available in a production environment. Obviously, the fastest way to display graphics information using a Cray system running COS is through a dedicated channel.


#### Graphics using the UNICOS operating system

Cray Research's UNICOS operating system provides opportunities for efficient transfer to graphic display devices of engineering data from a CRAY-2 or CRAY X-MP computer system running UNICOS. Users running UNICOS can send generic data directly to a front-end or workstation; no drivers are needed.

#### Graphics via TCP/IP

Two basic methods exist for moving graphics data using TCP/IP in the UNICOS environment. The first method involves using standard output to the user terminal; the second involves using sockets from the UNICOS operating system. Sockets are programming constructs available with TCP/IP, and are used for inter-program communication.

Using the standard output method, graphics data are channeled back to a user terminal and the graphics data are intermingled with the prompt data. Using sockets, the graphics data are piped back to an alternate location, which could be a window of a workstation or another graphics device. Sockets can be used in conjunction with the Cray Graphics Primitives, a set of primitives that are device-independent. In this environment a user would run



*Four views of an automobile engine section modeled on a Cray computer system and displayed with the graphics program MOVIE.BYU. Images courtesy Ford Motor Company.*

Brigham Young University's MOVIE.BYU graphics package (or a similar application), and the graphics data would be piped to a workstation in a generic format independent of the destination. The data are ported to a window of the workstation with a graphics server that adds the appropriate control commands to the specific workstation and then channels the data to the correct location. Sun, Apollo, Silicon Graphics IRIS, and DEC VAXstation II/GPX workstations are supported by these primitives.

Sockets also can be used to distribute the application across computer systems. The Minnesota Supercomputer Institute, for example, has distributed the more interactive portions of the graphics program MOVIE.BYU to an IRIS workstation while running the more computationally intensive portions on a CRAY-2 system. In this situation, the local hardware capabilities of the IRIS workstation — such as zoom, pan, and window — are used to position models; the computationally intensive calculation of the model is performed by the Cray system. The image then is sent back to the workstation.

This turns out to be a very efficient engineering environment, as the intelligent local capabilities of the workstation can be used to their fullest potential and the computationally intensive portions of the program are executed quickly on the Cray system.

#### Raw HYPERchannel connections

Transfers to workstations are possible via raw connections to a Network Systems Corporation HYPERchannel, rather than via TCP/IP. Workstations generally are unable to load bit-mapped images directly into memory. But using a

HYPERRchannel connection, a server program on the front-end calls the appropriate system graphics routine as graphics data are transferred to the workstation. Display speeds are improved because the TCP protocol is bypassed.

#### HSX channel graphics


Cray Research recently released the HSX-1, a high-speed external communications channel. The HSX-1 can be connected to devices such as an Ultra Corporation frame buffer. The Ultra graphics system provides animation of time-dependent simulation fields from 15 to 60 frames for precomputed images. Precomputed high-resolution 1280 by 1024, 24-bit images can be displayed in one-fifteenth of a second. The 90 Mbyte/sec transfer rate allows real-time display of the precomputed data.

Methods in addition to those outlined here exist for converting supercomputer results into graphics. Users may have to experiment to determine the route that best meets their needs. But an efficient configuration of hardware and software to generate graphic output will significantly increase the versatility and ease of use of any Cray system used in an engineering environment. □

---

#### About the author


*Mike Long received his B.S. and M.S. degrees in civil engineering from Brigham Young University. He joined the applications department at Cray Research in 1983. Long currently is a senior applications analyst in the department, primarily responsible for developing and testing graphics methods that use various stations, networks, and channels, and for developing design optimization strategies using Cray systems.*



# Numerical simulation of thunderstorm outflows and microbursts

*Kelvin K. Droegemeier, University of Oklahoma, Norman, Oklahoma*

Although thunderstorms are familiar events to most people, they continue to challenge scientists who want to understand better their structure and evolution. This knowledge can be used not only to help predict storms, but also to help explain specific storm-related phenomena such as microburst wind shear. Microbursts are known to have played a role in several commercial aircraft accidents, and thus are of particular importance to the scientific community and the general public. By running large numerical models of thunderstorms on Cray computer systems, researchers are gaining insight into the physical processes that give rise to and characterize these phenomena.



*Figure 1. Schematic showing the relationship of the cold-air outflow to its parent storm.*

## Storm dynamics

The dynamics of convective thunderstorms are controlled largely by events in the region between the cloud base and the ground. In particular, as a thunderstorm reaches maturity, rain-cooled air within the cloud produces a strong downdraft that impacts the ground and spreads radially in all directions, much like a miniature cold front (Figures 1 and 2). As this cold-air outflow moves along the ground, it forces unstable low-level air upward. Under certain conditions, this forced lifting can trigger intense new storms at the outflow's leading edge (gust front), or sustain an existing storm by continually supplying its updraft with unstable air.

Although the basic roles of outflows in storm morphology are well understood, details of outflow structure have eluded investigation. One can gain some insight into the turbulent structure of outflows by examining their laboratory counterpart, submerged density currents. As illustrated in the time-sequence photographs of Figure 3, laboratory density currents<sup>1</sup> often are characterized by turbulent Kelvin-Helmholtz eddies at the interface between the dense (white) and lighter (black) fluids. If these eddies are present atop real outflows, they could play a significant role in convective dynamics, particularly by altering the propagation speed of the gust front due to its entraining, or pulling along, ambient air.




Figure 3. A time sequence of photos, spaced one second apart, showing a vertical cross section through a laboratory density current (salt water is colored white) moving from left to right in a tank of fresh water. Note the development of Kelvin-Helmholtz eddies at the interface between the two fluids.



Figure 2. Photograph of the leading edge of a Florida outflow, moving from left to right. Note how topsoil is carried aloft by the strong winds and curves upward in a counterclockwise fashion just behind the front.

A more transient yet equally important subcloud phenomenon is the microburst<sup>2</sup>, a narrow (1 to 4 km diameter), intense downward-flowing jet of air that emanates from the base of a cloud (Figure 4). Although most microbursts last for only three to seven minutes and often are associated with rather innocuous looking clouds, they can produce downdraft and radially outflowing velocities along the ground in excess of 30 meters/sec. Since the associated wind shears are strongest near the ground, microbursts pose a significant threat to aircraft during takeoff and landing.

Although outflows and microbursts occur in a readily accessible region of the atmosphere, they are not amenable to observation using conventional instruments. Therefore, a high-resolution, two-dimensional numerical model has been developed to simulate thunderstorm outflows and microbursts. By focusing only on the subcloud region and neglecting the dynamics of the parent storm, high spatial resolution is used to depict explicitly the details of outflow and microburst structure, including turbulent mixing.

### The numerical model

Using conventional finite difference techniques, the numerical model efficiently solves equations for momentum, pressure, thermodynamic energy, and three categories of water (water vapor, cloudwater, and rainwater).<sup>3,4</sup> The highly vectorized 3000-line hydrodynamical code was de-

veloped originally on a CRAY-1 computer system at the National Center for Atmospheric Research (NCAR), and was recently converted for microtasking on CRAY X-MP and CRAY-2 systems.

Approximately 30 billion calculations are needed to simulate one hour of real time for a typical model configuration. To extract meaningful information from these complex simulations, data are written to history tapes and analyzed interactively using a software package developed by Joseph Klemp at NCAR (see CRAY CHANNELS, Summer 1985). Using color graphics and one-, two-, and three-dimensional display formats, one can interactively compute and overlay various fields within any region of the computational domain.

This analysis package recently was modified to animate user-selected data fields automatically. Viewing model data in color animation sequences lends a great deal more insight into complex physical processes than do standard contour plots. This comes as no surprise, because real thunderstorms are anything but quiescent. Thus viewing numerical representations of them as static images is counterproductive.

### Outflow simulation

To establish a basic understanding of outflow dynamics, and to allow for direct comparison between model results




Figure 4. Schematic showing the life cycle of a microburst.




Figure 5. Vertical cross sections through a laboratory density current (top) and the numerically simulated outflow.

and laboratory density current experiments, the simulations described in this article were conducted in a somewhat idealized setting. The ambient environment was calm, isentropic (at constant entropy), and completely dry; the outflow was initialized as a 2-kilometer-deep, horizontally-propagating pool of cold air entering the domain through a lateral boundary. The lower model boundary is a rigid, free-slip plate, the lateral boundaries allow for radiation of gravity waves, and the upper boundary is a non-rigid, constant-pressure surface. The slab-symmetric model domain is 40 kilometers long and 10 kilometers tall with a uniform grid spacing of 100 meters. Approximately 40 minutes are required for the outflow to travel the length of the domain.

Figure 5 shows a vertical cross section through a laboratory density current<sup>1</sup> (top panel) and the perturbation potential temperature (entropy) field from the numerical simulation (bottom). To facilitate comparison, a uniform color is assigned to all temperature values within the simulated outflow. A striking similarity is evident between the Kelvin-Helmholtz eddies in the model and those produced in the laboratory experiment, and a detailed analysis reveals that their nondimensional growth rates and aspect

ratios also are quite similar. Unfortunately, the ultimate breakdown of the eddies into small-scale turbulence, which is a highly three-dimensional process and is readily apparent in the laboratory experiment, is absent in the model results. Fully three-dimensional simulations now are being conducted to represent this process more accurately.

It is well known from surface observations that gust-front passage is heralded by a sudden drop in temperature, a rapid rise in pressure, a wind shift, and an increase in wind speed and gustiness. Figure 6 shows a split-screen presentation of the physical relationships among these events, consisting of vertical cross sections through the perturbation potential temperature field (lower two panels) and the perturbation pressure field (upper two panels). The entire computational domain is shown in the center two windows, while the extreme upper and lower windows show close-up views of the frontal region. High and low pressures are indicated by reds and blues, respectively, in the upper panels, while blue indicates the coldest air in the lower two panels. The pressure in the undisturbed body of the outflow (left portion of the domain) is in hydrostatic balance with the temperature field, while just ahead of the front, a region of high pressure exists due to the collision between the cold outflow and its ambient environment. Kelvin-Helmholtz eddies are evident immediately behind the gust front and are accompanied by regions of low pressure aloft (green and blue) that are manifest as periodic oscillations in pressure at the ground. Although observational studies have hinted at correlations between surface pressure fluctuations and organized turbulence aloft, this numerical study was the first to provide clear supporting evidence for such a relationship. In fact, shortly after the discovery of Kelvin-Helmholtz eddies in these simulated outflows, their presence was verified observationally in a Doppler radar study of a Colorado outflow.

To investigate the kinematics of the outflow, several thousand massless tracer particles are implanted in the computational domain at various locations and times, and are tracked individually by a Lagrangian technique. By continuously releasing particles at three vertical levels in the cold-air source region (Figure 7), one finds that air near




Figure 6. Vertical cross section through the model pressure field (upper two windows) and the perturbation potential temperature field (lower two panels).







Figure 7. Massless particles continuously released at three vertical levels describe the internal kinematics of the simulated outflow.



*Figure 8.* Rows of massless tracer particles, placed 0.5 km apart in a vertical stack beginning at a height of 2 km above the ground, illustrate how air passing over the gust front becomes entrained by the turbulent Kelvin-Helmholtz eddies.

the top of the outflow (green and red particles) never reaches the front, but is instead entrained by the developing Kelvin-Helmholtz eddies. Air close to the surface (blue) travels the fastest and, upon reaching the front, curves upward in a counterclockwise circulation induced by pressure-density solenoids. Evidence of such a circulation is revealed by topsoil carried aloft behind the gust front in Figure 2.

To determine if an outflow is capable of initiating strong thunderstorms, it is important to understand how ambient



*Figure 9.* Pillars containing several thousand tracer particles are placed along the ground to illustrate the destiny of low-level air upon being lifted by the gust front.

air is lifted by the oncoming gust front. By placing several rows of tracer particles at various heights in the model domain (Figure 8), it can be observed that only air near the top of the outflow (2 km above the ground) is entrained by the turbulent eddies. The remainder passes smoothly over the gust front, and thus would have an opportunity to feed the updraft of the parent storm (which would be to the far left in the figure). More important to the process of storm initiation is the lifting of near-surface air by the gust front. Figure 9 shows a time sequence in which pillars containing several thousand tracer particles are placed along the ground in advance of the outflow. After being lifted by the gust front, nearly all of the particles are entrained by the Kelvin-Helmholtz eddies and thus never have an opportunity to reach the parent storm updraft. It is clear from this particle sequence that a great deal of turbulence attends the simulated Kelvin-Helmholtz eddies, and that it would be undesirable to attempt an aircraft landing through this particular outflow.

If condensation were allowed to occur in this experiment, the associated buoyant lift due to latent heat release could cause the low-level air to continue rising, thereby avoiding entrainment by the turbulent eddies and instead resulting in a new storm at the gust front. This storm might capture much of the low-level air destined for the parent storm updraft, thus causing it to die. Further work with a version of the model that includes moisture is underway to investigate the unique periodic development and decay exhibited by these types of storms.

### Microburst simulation

In contrast to the linear two-dimensionality of outflows, microbursts tend to be axisymmetric. The model used in these experiments therefore is axially symmetric with a uniform grid spacing of 75 meters and an 8-kilometer vertical and 10-kilometer radial dimension. The ambient initial environment is completely dry with a temperature structure characteristic of the microburst-producing environments in Colorado. Although the former condition is somewhat unrealistic, it provides the greatest potential for evaporative cooling and thus represents an upper limit to microburst severity.

The microburst is created by placing a 2-kilometer-wide and 4-kilometer-deep Gaussian distribution of rain (maximum value of 8 grams of water per kilogram of dry air) on the axis of the domain. As the integration begins, the rain falls and evaporates, thereby cooling the surrounding air and initiating an intense downdraft that impacts the ground and spreads radially. After only four minutes, the downdraft speed reaches 19 meters/sec, and is accompanied by a dome of high pressure at the surface. This high contains not only a hydrostatic component due to the weight of the cold downdraft air, but also a dynamic component due to the impact of the downdraft on the rigid ground. The pressure dome deflects the downdraft air radially, and is responsible for the rapid acceleration of radial winds which, at this time, are in excess of 15 meters/sec at the ground. By seven minutes, nearly all of the rain has evaporated, leaving behind a shallow pool of cold air which has a noticeably elevated "head" immediately behind its leading edge. Although the downdraft has

weakened somewhat, a strong updraft (maximum value of 8 meters/sec) becomes evident at the cold-air front.

An aircraft attempting to land on a glidepath through this updraft would encounter strong upward motion. To compensate, the pilot might lower the nose of the aircraft and reduce power to recapture the glideslope. However, the aircraft soon would exit the updraft and begin descending rapidly in the downdraft where increasingly strong tail winds (maximum intensity of 36 meters/sec at the ground) would be encountered after crossing the downdraft axis. Because the aircraft already is in a nose-down attitude at decreased power and has its flaps, leading edge slats, and landing gear extended, the continued loss in lift, which becomes more severe as the aircraft continues to descend, may prove impossible to overcome.

## Outlook for the future

Two basic limitations exist in the experiments described above. First, the two-dimensionality of the model limits the realism of the simulations because turbulent processes are represented incorrectly.<sup>5</sup> Furthermore, increased spatial resolution is needed to adequately resolve small-scale turbulent structures, particularly the three-dimensional characteristics of thunderstorm outflows. High-resolution experiments now are being conducted with two- and three-dimensional versions of the model on the CRAY X-MP/48 computer system at NCAR, and on the CRAY-2 system at the University of Minnesota Supercomputer Institute. Adaptation of the code written originally for the CRAY-1 system was accomplished in approximately 30 minutes, and microtasking was successfully implemented in a few hours.

As a practical step toward addressing the issue of aircraft safety, the model data are being considered for implementation into the fully-configured B-727 flight simulator at the Federal Aviation Administration Training Academy in Oklahoma City. By allowing pilots to fly through numerous datasets, and thus physically different scenarios, we hope to not only identify characteristics that distinguish potentially deadly shears from the more innocuous situations, but also to provide pilots with insight into how an aircraft might respond in various wind shear situations. Because the model data are fully time- and space-dependent, pilots could choose to fly through different regions of the flow field at various times in the microburst's brief but dangerous life cycle.

Additional research involves examining the initiation of convection along gust fronts, and particularly examining the role of low-level wind shear in providing an environment favorable for severe storm generation. These experiments are being conducted on a CRAY X-MP/48 system with the Klemp and Wilhelmson<sup>6</sup> three-dimensional cloud model. On a somewhat larger scale, the Colorado State University Regional Atmospheric Modeling System (RAMS), which was designed on Cray supercomputers at NCAR, is being used to examine the characteristics of dry-line thunderstorms in western Oklahoma. These storms are somewhat anomalous in that they typically produce little precipitation, but can become severe and even tornadic as they move eastward into central Oklahoma.

Finally, a long-range effort is underway to develop a next-generation mesoscale model and data analysis system at the Norman, Oklahoma Weather Center (which is composed of the University of Oklahoma School of Meteorology, the National Severe Storms Laboratory, the Cooperative Institute for Mesoscale Meteorological Studies, and the National Weather Service Forecast Office). Known as the Central Oklahoma Mesoscale Modeling and Analysis Project (COMMA), the initiative is using advanced numerical techniques and Cray supercomputers to build an atmospheric model capable of representing individual clouds and the larger-scale environment. Through close interaction with the operational component of the Weather Center, the COMMA model will also serve as a prototype system for producing short-term local weather forecasts. □

## About the author

Kelvin Droegemeier is an assistant professor in the School of Meteorology at the University of Oklahoma. He received his B.S. degree in meteorology from the University of Oklahoma in 1980, and his M.S. and Ph.D. degrees in atmospheric science from the University of Illinois in 1982 and 1985, respectively. He organized and is serving as the principal coordinator of the Central Oklahoma Mesoscale Modeling and Analysis Project, and was named a Presidential Young Investigator by the National Science Foundation in 1987.

## Acknowledgments

Figures 1 and 5-9 were taken from a movie produced by the author on a CRAY X-MP/22 computer system at Digital Productions, Los Angeles (copies are available from the author upon request). Figure 2 was provided by Ron Holle of the National Oceanic and Atmospheric Administration in Boulder, Colorado, and Figure 4 was made available by Professor T. Theodore Fujita of the University of Chicago. Figure 3 is from reference 1. This research was supported by the National Science Foundation under grants ATM84-15222 and ATM86-04402.

## References

1. Simpson, J. E., "A Comparison Between Laboratory and Atmospheric Density Currents," *Quarterly Journal of the Royal Meteorological Society*, Vol 95, 758-765, 1969.
2. Fujita, T. T., *The Downburst - Microburst and Macroburst*, University of Chicago, 1985.
3. Droegemeier, K. K. and R. B. Wilhelmson, "Numerical Simulation of Thunderstorm Outflow Dynamics Part I: Outflow Sensitivity Experiments and Turbulence Dynamics," *Journal of Atmospheric Sciences*, Vol. 44, No. 8, pp. 1180-1210, 1987.
4. Droegemeier, K. K. and R. P. Davies-Jones, "Simulation of Thunderstorm Microbursts with a Supercompressible Numerical Model," *5th International Conference on Numerical Methods in Laminar and Turbulent Flow*, Montreal, Quebec, Canada, 1987.
5. Deardorff, J. W., "The Development of Boundary-Layer Turbulence Models for use in Studying the Severe Storm Environment," Proceedings, SESAME Meeting, Boulder, NOAA-ERL, 1975, p. 251-264.
6. Klemp, J. B. and R. B. Wilhelmson, "The Simulation of Three-Dimensional Convective Storm Dynamics," *Journal of the Atmospheric Sciences*, Vol 37, 1978, pp. 1070-1096.



# A new kind of gusher for the Lone Star state

## Supercomputing in The University of Texas System



One literary image cannot adequately describe the vast and diverse state of Texas. Its cosmopolitan cities are bustling centers of oil trade and high finance. The Texas rangelands are a seemingly endless carpet of green and brown, dotted with Longhorn cattle and the faithful bobbing of oil derricks. Texas is The Alamo, rich in a history of local pride and staunch independence. Texas is also a state in transition, working to forge new industries in the wake of the current glut in the world oil market. Against this backdrop is The University of Texas System, one of the largest university systems in the world, and one of the growing number of institutions that have looked to supercomputing to help keep them at the forefront of academic research.

The University of Texas System Center for High Performance Computing (UT System CHPC) was established to serve the research and instructional needs of the seven academic and six health component institutions of the UT System. The UT System was the fourth U.S. university to acquire a Cray supercomputer and currently operates without the benefit of National Science Foundation (NSF) funding.

### Serving the Lone Star state

Unlike many university supercomputer centers, the CHPC was established almost solely to serve the needs of the UT System. Although access is available to other users, computational work conducted at the CHPC generally has some relationship to a UT System institution's project. The CHPC does not sell time commercially and does not solicit use by out-of-state researchers. Because of the growing demand for computation at the 13 UT System institutions, very few surplus cycles are even available.

The UT System had three primary goals for the CHPC: to provide state-of-the-art computational resources for research by faculty and students, to encourage the use of computational techniques and expand the computer "literacy" of students at UT System institutions, and to provide tools that would attract the finest researchers and teachers to the UT System campuses.

Dr. Gerhard Fonken, chairman of the UT System CHPC executive committee and executive vice president and

provost of UT Austin, explains, "The utility of a supercomputer center is obvious for a large graduate and research institution like UT Austin, but a little less obvious for a smaller institution such as UT Tyler, which has relatively little in the way of graduate programs and research. They do, however, have a computer science program that now has access to a supercomputer. As they begin to exploit that access, they can provide excellent advanced training to students, who could then leave UT Tyler with a bachelor's degree that includes a supercomputer sequence, making them more valuable to graduate programs and industry."

Similar opportunities exist for the health institutions. "The UT System health component institutions are using the CHPC to educate medical research students in computational methods, as well as to educate medical faculty," notes Dr. Ken Tolo, vice provost of UT Austin. "The most profound initial impact in this area has been among institutions that were already primed; those that had begun to develop faculty research capabilities but lacked the computational resources. UT Dallas and UT Arlington, in particular, have become very active users of the CHPC."

On the instructional side, The University of Texas at San Antonio has initiated a course development project that has helped to develop curricula in a serious effort to teach students the fundamental "hows and whys" of supercomputing. The university leaders behind the CHPC believe that the courseware being developed there may have broader application at other UT System institutions and institutions in other states as well. "In general, the availability of the CHPC also means that students being taught Fortran, particularly at the smaller schools, can use the system to pursue their discipline and learn vector computation or multitasking (executing a single job on two or more processors of a multiprocessor system simultaneously) at the same time they are learning computational methods generally," explains Dr. Charles Warlick, executive director of academic information systems for the UT System. "Eventually we could have a whole generation of students who would never know that vector computation or parallel processing is not a normal part of computational science."

Allocation of CHPC computing cycles is administered by the CHPC executive committee. The exact amount of computer time available to each UT System institution varies, roughly determined by the size of the institution and its projected computational needs. Regardless of size, however, every institution is allocated at least 200 CPU hours annually. Each institution has complete autonomy to distribute its allocation among its various research and educational programs. "Each institution has control of its own computational resource — no questions asked," said Fonken. "What we have found is that with local autonomy comes a kind of pride of ownership, and a willingness to try unusual or high-risk programs that these institutions might not otherwise undertake. This approach gives them maximum freedom for experimentation."

### The UT System network

The CHPC is located at the Balcones Research Center north of the UT Austin campus. The center's current hard-



ware configuration consists of a CRAY X-MP/24 computer system with 32-million-word SSD storage device, IBM 4381 and DEC VAX 8600 front-ends, and a variety of peripherals and tape and disk mass storage. The VAX is linked to four local high-performance graphics workstations and to a network that connects to the 13 UT System institutions and external networks (see map below).

The primary goal of the UT System network is to give supercomputer users at each institution the perception that the CHPC is located on their campus. The network posed several implementation challenges, not the least of which was a geographically large service area where component institutions were separated by as much as 700 miles. William Bard, director of the UT System office of telecommunication services, designed and implemented the statewide network. "To accomplish our goals, we knew the network would need to provide basic network services, including virtual terminal interaction, file transfers, and remote process communication," Bard said. "And, to the largest extent possible, it would have to do so utilizing computing resources already in place on the campus of each component institution. This posed a further challenge, because it meant that the hardware and software used within the network had to be versatile in adapting to the various computers, operating systems, local networks, and communications systems found on the campuses."

The CHPC's role as a service center for the component institutions defined other network characteristics. It had to provide network services without placing additional administrative, fiscal, or technical demands on each institution. A tight implementation schedule and the limited size of the network staff imposed additional constraints on the network's design. The network had to be assembled using off-the-shelf hardware and software products that had a proven performance history. There was no time to develop special network systems or experiment with untested products.

The network was designed to adhere to a single set of consistent standards and protocols implemented in hardware and software supplied by a single vendor. A network interface processor (NIP) terminates the network on the campus of each component institution. Each NIP is custom configured to provide services and interfaces required by each particular campus. For scheduling purposes, the UT System was divided into six geographic areas. Initial projections of the network traffic in each area suggested


that the Dallas, Houston, and San Antonio network routers be connected via 56 Kbit/sec circuits, while the Tyler and west Texas routers are served with 9.6 Kbit/sec circuits. All circuits connecting NIPs to routers operate at 9.6 Kbits/sec. Data communication between the UT Austin campus and the CHPC is provided via a microwave link.


In addition to system-wide access to the CHPC, the network provides each institution with access to ARPANET (the network of the Defense Advanced Research Projects Agency), BITNET, and an evolving network of Texas institutions called TEXNET.

According to Bard, future network development will focus on transmission cost and performance improvements, increased reliability, network management, and additional network applications. "Because of the long distances involved in the CHPC network, satellite transmission appears to be a particularly attractive alternative for the future," said Bard. "For increased reliability, we're primarily looking at installing additional data circuits to provide alternate routing within the network."

## Two perspectives on space

Drs. Bob Schutz and Byron Tapley, professors of aerospace engineering and engineering mechanics at UT Austin, are using the CHPC's CRAY X-MP system to analyze satellite tracking data in an effort to understand various geophysical and geodetic phenomena. A key part of this effort is modeling the Earth's gravity field, which has variations in longitude and latitude, as well as time. Modeling the gravity field, as inferred by observations of the motions of man-made satellites, requires the estimation of several thousand parameters based on several million satellite observations. The observations are mainly laser ranging measurements collected by NASA tracking stations, including the McDonald Observatory operated by UT Austin in west





*Figure 1. Color-contoured shaded relief map of sea surface topography created from altimeter measurements collected by the satellite Seasat. The area illustrated is in the Western Pacific and the colors represent contour intervals ranging from -9.9 meters (purple) to 52.5 meters (pink), relative to an ellipsoidal surface. Bathymetric features are enhanced by the shading in this sea surface contour image, such as the Japan Trench and various seamounts. Gray areas are land features, such as Japan and China, for which no data exist since the altimeter operates in ocean areas only. This image represents one aspect of research using the Cray system in space geodesy.*

Texas and several cooperating stations in Europe, Asia, and Australia.

One result of such modeling has been a better understanding of how the Earth's shape is changing over time. The slow deglaciation of the polar regions over tens of thousands of years has lightened the load on the Earth, and thus subtly changed the oblate spheroid shape of the planet. In addition, laser ranging data identifying the movement of tectonic plates help researchers analyze continental drift. Analysis of satellite ranging data has also revealed new details about the topography of the ocean surface. As can be seen from the color map (Figure 1), deep ocean trenches cause low points on the ocean surface, just as sea mounts (underwater mountains) cause high points on the surface.

While Schutz and Tapley are using space-based instruments and the CHPC to gain a better understanding of worldly events, another team of UT scientists is using the CHPC to gain a better understanding of the cosmos. Drs. Robert Harkness and J. Craig Wheeler, of the astronomy department at UT Austin, are using the CHPC's Cray system to model supernovae, the explosive deaths of stars. Their findings are not only shedding light on this spectacular cosmic event, but also are helping to answer long-standing questions about the size and age of the universe.

Effectively modeling a supernova requires combining theoretically predicted luminosities of the nuclear burning of various elements with the theory of stellar atmospheres and radiative transfer. As the supernova burns, it rapidly expands into a dispersed gaseous mass of varying opaci-

ty. It is this expansion that makes solution of the radiative transfer equations so demanding. The equation is essentially four-dimensional, and the unique composition of relatively heavy elements with complex atomic structures requires an extremely dense frequency sampling to approximate the opacity of the matter. The final result is a theoretical prediction of the development of the spectrum over time from a given theoretical model of the explosion. Comparison with observed supernovae has led to refinement of allowable physical models and to important insights into a new class of supernovae. The work has also provided a special tool to explore the recently discovered supernova in a nearby galaxy, the brightest in almost 400 years.

## Oil, of course

A story about supercomputing in Texas would be incomplete without a report on computational work relating to reservoir simulation; that is, the modeling of petroleum reservoirs to determine the best recovery techniques, typically by injecting agents into the reservoirs to release trapped hydrocarbons. The Petroleum Engineering Department at UT Austin primarily works with three reservoir simulation codes: UTCOMP, a carbon dioxide flooding simulator; UTTHERM, a compositional/thermal simulator; and UTCHEM, a chemical flooding simulator. Drs. Gary Pope and Kamy Sepehrnoori, professors of petroleum engineering, have focused considerable attention on optimizing UTCHEM for the CRAY X-MP system. Through vectorization and a faster solver, they have achieved speedup factors of up to 22.8 over unoptimized versions of the code. The Cray system's four-megaword memory has also allowed the development of larger models and even field-scale chemical flooding simulations.

According to Pope and Sepehrnoori, one of the most significant applications of UTCHEM during the past year has been its use to simulate the Big Muddy (a Conoco oil field) surfactant plot. Although they had done field-scale studies previously, this was their first attempt to simulate an actual field test. Pope and Sepehrnoori achieved a simulated value for the final oil recovery within one percent of the actual field value, indicating that they had accounted for all the major phenomena and had no serious errors in the controlling physical property data measured in their laboratory for the surfactant system Conoco used.

As part of their field-scale process simulation studies, Pope and Sepehrnoori have made a large number of three-dimensional simulations of chemical floods in heterogeneous reservoirs to evaluate such things as crossflow between permeability strata, and have correlated these results with other process characteristics. Pope and Sepehrnoori believe that their studies could be of significant value in predicting the best conditions for chemical flooding as well as enhancing overall understanding of the significance of the many complex coupled phenomena affecting the oil recovery in chemical flooding. Before the CHPC was established, the accuracy of such studies was suspect due to the constraint of relatively small mesh sizes imposed by less powerful computers. Thanks to the CHPC and their code optimization work, Pope and Sepehrnoori have been able to systematically evaluate the accuracy of their model for



large-scale three-dimensional cases. "While the accuracy of the simulations depends on the type of problem," Pope said, "we are confident that our results are reasonably accurate given the coarse grids we are using for routine process cases."

## Chemistry

Dr. Robert Wyatt, professor of chemistry at UT Austin, is using the large-scale computational capability of the CHPC to model several problems in chemical dynamics, including laser-molecule interaction, the transport of shock waves in crystals, quantum mechanical studies of chemical reactions, and neural networks.

Wyatt's neural network studies involve integrating the analog equations of motion for a large number of neuron-like elements, and then analyzing the capability of the systems to store or recognize neural patterns. In effect, Wyatt's model treats the neural system as a large parallel processor. During the "training" period, he can feed patterns into the system, and the system cooperatively develops into groups that are actually tightly coupled processors (the individual neurons). A learning algorithm in the program helps the system develop couplings among the processors, which become stored memory. Wyatt can even feed in a "bad" pattern, and the system will correlate it with the stored pattern — a sort of associative memory. Figure 2 shows a plot of firing neurons for one such study. "My approach is to put as much physiology into the model as possible," Wyatt says. His findings could result in new ways to store, analyze, and retrieve information.

Wyatt is also working on solving the Schrodinger time-dependent wave equation for a laser interacting with a molecule. A recursive technique is used to find quantum mechanical transition probabilities for systems much larger than any systems previously analyzed. While the largest systems studied previously had no more than about 1200 quantum mechanical states, Wyatt's calculations have been possible on systems of up to 500,000 states. For some of these systems, the molecular energy plotted as a function of time is linear, while in other cases the energy oscillates and does not permit the molecule to absorb much energy. For such studies, Wyatt says, the speed and memory of the Cray system at the CHPC is critical.

## Medical applications

Much of the work described here is being conducted at UT Austin with support from Cray Research's University Research and Development Grant Program. The Cray grant program is helping to fund computational work at other UT system institutions as well. Researchers at the UT Medical Branch in Galveston are using the Cray system to process magnetic resonance images of the human heart, measuring parameters such as ventricular wall thickness, ventricular volume, and tissue relaxation times. At The UT System Cancer Center, Dr. E. Neely Atkinson is working on simulation techniques to study the macroscopic course of breast cancer.

Other programs at UT System institutions range from developing new algorithmic approaches to laser beam analy-




Figure 2. Neuron voltages (in millivolts) vs. time (in milliseconds). Neuron 55 (solid curve) receives information from receptor neurons for the first 50 msec. Neuron 1 fires after neuron 55 and continues to fire after 50 msec. due to 'reverberations' set up in the network of background neurons. There are altogether 100 neurons in the network, and 22 of them (including neuron 55) receive a pattern from the receptors. The neural network is modeled by integrating a system of differential equations for the voltages.

sis. At UT Arlington, studies are underway to simulate the flow produced by a helicopter rotor during hover, while researchers at UT Dallas are busy investigating the possibility of developing an environment in which parallel algorithms can be mapped into the vector architecture of Cray systems.

## Santa Rita #2?

The UT System CHPC is bringing new perspectives to old problems such as oil recovery, and it is helping to break new ground in fields such as chemistry and astronomy. Just as important, the CHPC is creating a new generation of graduates, graduates who will have an understanding of a technology that is sure to be a part of their future.

The first oil discovered on UT System land was at a site called Santa Rita #1, a discovery that helped the UT System build what has become one of the largest institutional endowments in the United States. Soon after the CHPC began operation in the spring of 1986, the then governor of Texas called the new supercomputer "Santa Rita #2." As the initial hopes for the CHPC are realized, the UT System may indeed find itself with a new form of black gold. □

## Acknowledgments

Many thanks to the users and administrators of the University of Texas System Center for High Performance Computing, whose warm hospitality and candor helped to make this article possible.

*Editor's Note:* Dr. Charles Warlick, executive director of academic information systems for the UT System, was interim director of the UT System CHPC at the time this article was written. The UT System recently announced the appointment of Dr. James Almond as director of the UT System CHPC.



# Fortran

# optimization

# techniques

## Reducing memory bank conflicts on CRAY X-MP computer systems

*Pete Sydow, Cray Research, Inc.*

The processing speed delivered by any computer system depends to some degree on the skill with which programmers structure their codes. Sometimes minor programming changes that enable codes to exploit hardware features more efficiently will improve code performance considerably. An understanding of memory organization, for example, can be used to structure codes to avoid memory conflicts, which are a common obstacle to maximizing code performance. This article describes the basic memory organization of Cray computer systems and offers suggestions for avoiding memory conflicts when accessing memory in vector mode.

Although the vector architecture of CRAY X-MP systems is the most forgiving of memory conflicts when compared to other vector architectures currently available, several potential sources of memory conflict nonetheless exist when accessing interleaved memory in vector mode. Most other architectures penalize a code if it accesses memory locations in increments other than 1, that is, if the code's memory stride is not consecutive but incrementally skips some locations as it accesses memory. Such architectures impose a performance penalty on codes in this situation even if the incremental stride does not cause memory conflicts. For example, if the memory stride is 2 (that is, ev-

0	1	2	3	4	5	6	7
0	1	2	3	4	5	6	7
8	9	10	11	12	13	14	15
•	•	•	•	•	•	•	•
•	•	•	•	•	•	•	•
•	•	•	•	•	•	•	•
0 + 8n	1 + 8n	2 + 8n	3 + 8n	4 + 8n	5 + 8n	6 + 8n	7 + 8n

Table 1. Locations of memory addresses in an eight-bank memory.

every other bank is accessed), then the access time may be doubled. CRAY X-MP system architecture, however, will not impose any such penalty for nonunit memory strides, so long as no memory bank conflicts result.

When memory conflicts do occur, they can be due to any of several causes:

- References within a given access can conflict with each other. That is, a vector load or store can conflict with itself.
- Different memory accesses from the same processor can conflict with each other through multiple memory ports.
- Memory accesses from different processors can conflict with each other on multiprocessor systems.
- Memory accesses can conflict with I/O traffic, such as data flowing between memory and an SSD solid-state storage device.

This article will discuss ways in which programmers can avoid the first and second sources of conflict listed above. It will concentrate primarily on the first source because that is the most common source of conflict and the one most easily resolved.

## Hardware

The interleaved memory of Cray systems is arranged so that consecutive addresses are in consecutive banks. Assuming that memory address 0 is assigned to bank 0, the bank number  $N$  for any other address  $M$  is given by

$$N = M \text{ modulo } B$$

where  $B$  equals the total number of memory banks. On Cray systems, which have a total number of banks that always is a power of 2 ( $B=2^k$ ), the bank number is given

directly by the low-order  $k$  bits in the binary representation of the memory address. Table 1 shows which addresses go in which banks in an eight-bank system. The same principle applies to 16-, 32-, and 64-bank systems.

Once a memory bank is accessed, the access will keep the bank busy for a fixed number of clock cycles. This duration is called the *bank busy time*, or memory cycle time, and is the most important factor to consider in detecting and resolving memory bank conflicts. The bank busy time for CRAY X-MP systems with ECL bipolar memory circuits is four clock cycles; the bank busy time for CRAY X-MP systems with MOS memory circuits is eight clock cycles. If a reference is made to a bank while it is busy with a previous reference, the result is a memory bank conflict, and the conflicting reference must wait for the busy condition to clear. The wait time will be one to three clock periods with ECL memory, or one to seven clock periods with MOS memory. Although this explanation of memory hardware is simplified, it identifies the important points regarding bank conflicts.

## Software

To write code that minimizes memory bank conflicts, programmers must understand how arrays and matrices are stored in memory, based on the rules of the programming language (Fortran in this case). For example, Table 2 shows the storage arrangement for the arrays  $A$  and  $B$  declared in the dimension statement

*DIMENSION A(8,4), B(4,4)*

Notice in the example for an eight-bank machine shown in Table 2 that if the program uses these arrays by accessing the first dimension, consecutive references are in consecutive banks. Notice also that if the program uses these arrays by accessing the second dimension, then bank

0	1	2	3	4	5	6	7
a(1,1)	a(2,1)	a(3,1)	a(4,1)	a(5,1)	a(6,1)	a(7,1)	a(8,1)
a(1,2)	a(2,2)	a(3,2)	a(4,2)	a(5,2)	a(6,2)	a(7,2)	a(8,2)
a(1,3)	a(2,3)	a(3,3)	a(4,3)	a(5,3)	a(6,3)	a(7,3)	a(8,3)
a(1,4)	a(2,4)	a(3,4)	a(4,4)	a(5,4)	a(6,4)	a(7,4)	a(8,4)
b(1,1)	b(2,1)	b(3,1)	b(4,1)	b(1,2)	b(2,2)	b(3,2)	b(4,2)
b(1,3)	b(2,3)	b(3,3)	b(4,3)	b(1,4)	b(2,4)	b(3,4)	b(4,4)

Table 2. Memory bank locations of arrays  $A$  and  $B$  declared by the dimension statement *DIMENSION A (8,4),B(4,4)* for an eight-bank memory.

```

DIMENSION X(N,N),Y(N,N),Z(N,N)
DO 200 I=1,N
DO 200 J=1,N,2
  X(I,J) = X(I,J) + Y(I,J) * Z(I,J)
200 CONTINUE

```

Figure 1. Sample code used in example to calculate bank conflicts.

conflicts are certain because consecutive references use the same bank for matrix *A* and alternate between two banks for matrix *B*. Such an access will conflict with its own previous references. In addition, if the program uses these arrays by accessing the second dimension, then consecutive vector loads stand a good chance of conflicting with each other because of the small number of memory banks involved.

### Detecting memory bank conflicts

Programmers can use a simple analytical principle to detect memory bank conflicts. Consider the expression

$$\frac{\text{number of banks}}{\gcd(\text{memory stride}, \text{number of banks})}$$

where  $\gcd$  is the greatest common divisor, and memory stride can be determined by taking the product of the leading dimensions and the skip increment of the array access. If the number given by this expression is less than the bank busy time, the vector memory access has bank conflicts with itself. Consider the sample code shown in Figure 1, where  $N=20$ , being run on an eight-bank machine with an eight-clock-period bank busy time. In this loop the leading dimension is 20 and the skip increment of the array access is 2, so the memory stride is 40. The gcd of 40 and 8 is 8, and the number of banks (8) divided by 8 equals 1. This means that a reference is going to the same bank every clock period. Since this (one clock period) is less than the bank busy time of eight clock periods, memory bank conflicts will occur.

### Resolving memory bank conflicts

Sometimes simple program changes will help resolve memory bank conflicts. If a program accesses memory in some pattern, generally a software solution exists. Codes such as Monte Carlo and table look-up problems that randomly access memory do not offer simple ways to resolve bank conflicts. In such cases, memory bank conflicts can be addressed by an algorithm change or a hardware upgrade to obtain more banks. But for programs that do access memory in a pattern, guidelines exist for resolving memory bank conflicts:

1. When using multiple-dimensioned arrays, ensure that the innermost DO loops reference the leading dimension. In many cases this can be accomplished by inverting DO loops. In some cases, inverting DO loops alone will not help because the procedure violates guideline 3 below.
2. When using multiple-dimensioned arrays, use odd num-

bers for leading dimensions. For example in an eight-bank system, an array *A* dimensioned (8,4) will be stored in this arrangement

BANK NUMBER							
0	1	2	3	4	5	6	7
1,1	2,1	3,1	4,1	5,1	6,1	7,1	8,1
1,2	2,2	3,2	4,2	5,2	6,2	7,2	8,2
1,3	2,3	3,3	4,3	5,3	6,3	7,3	8,3
1,4	2,4	3,4	4,4	5,4	6,4	7,4	8,4

If the program uses this array by going down the second dimension, the first four memory access calls will be to the same bank, resulting in memory bank conflicts. If, however, the array is redimensioned to (9,4) then it will be stored in memory in this arrangement

BANK NUMBER							
0	1	2	3	4	5	6	7
1,1	2,1	3,1	4,1	5,1	6,1	7,1	8,1
9,1	1,2	2,2	3,2	4,2	5,2	6,2	7,2
8,2	9,2	1,3	2,3	3,3	4,3	5,3	6,3
7,3	8,3	9,3	1,4	2,4	3,4	4,4	5,4

and so on. In this case reference to the second dimension will not cause bank conflicts, since the memory stride is 9. In this way more memory is used, but where memory is not a constraint it provides a simple way to avoid bank conflicts. In many cases this method can be used to detect memory bank conflicts as well. It will be of benefit in cases where a single access conflicts with itself and where multiple accesses conflict with each other. In a recent case this solution resulted in a factor-of-two improvement in total code performance.

3. Make the largest iteration the innermost DO loop. This will increase vector length and decrease loop overhead. This point actually is not directly related to bank conflicts, but is generally a good practice to improve code performance.
4. Reduce memory traffic. This can be accomplished by several means:

- Loop fusion — combining two loops into one
- Loop unrolling
- Use of matrix-vector or matrix-matrix kernels instead of vector-vector kernels; that is, use of MXV, MXM, or BLAS2 kernels instead of the old BLAS kernels
- Use of more common expressions

Table 3 shows the resulting performance improvements when the first two strategies listed above are applied as shown in Figure 2.

Taking the DO 200 loop in Figure 2 as the original coding, the two loops following it represent attempts to avoid memory bank conflicts. The DO 210 loop was created from the original by inverting the inner and outer loops. The DO 220 loop was created by changing the leading dimension. The numbers in Table 3 represent the run time ratios of the modified loops to the original. The most significant improvements can be seen for  $N=40$  and  $N=80$ .

### Interprocessor memory conflicts

It is worth noting a few observations about a different type of memory conflict: conflict among processors in multiprocessor systems. The fluctuating code performance experienced when interprocessor conflict is present reveals information about the ability of the memory bandwidth to accommodate memory traffic. A small fluctuation indicates sufficient bandwidth; a large fluctuation indicates insufficient bandwidth.

Even on CRAY X-MP systems containing the maximum number of memory banks, some degradation in code performance will occur due to interprocessor memory conflicts when heavily vectorized codes run on all processors. Scalar performance is somewhat immune to bandwidth considerations, and fluctuation due to interprocessor conflicts is usually negligible for scalar codes. Typically, on a CRAY X-MP/48 system with 64 banks, the degradation of vector-dominated code running in a batch environment is from three to five percent. In a worst-case scenario, on the same system running four vector-dominated codes, performance degradation can be 10 percent or more.

A convenient way to determine whether memory bandwidth is sufficient is to divide the number of banks by the product of the number of processors, the number of memory ports per processor, and the bank busy time. If the resulting number is greater than or equal to 1, then the

n	DO 210 invert loops	DO 220 change dimension
10	.97	1.00
20	.92	.96
30	.94	1.02
40	.65	.67
50	.94	1.00
60	.87	.89
70	.95	.99
80	.32	.32
90	.97	.98
100	.87	.89

Table 3. Performance improvements when the DO loop given in the text example is optimized by the first two strategies discussed in the text. The numbers are ratios of optimized to original run times.

memory bandwidth is generally sufficient even if all processors are running vector-dominated codes.

### Conclusion

Although some memory conflicts may be unavoidable for a given code on a given system, an understanding of hardware organization can be helpful in minimizing memory conflicts. Some memory bank conflicts can be resolved relatively easily through programming. Structuring code to take into account potential memory conflicts will in most cases improve total code performance, and sometimes the improvement can be substantial. □

### About the author

Pete Sydow is manager of benchmarking analyst services at Cray Research, a position he has held since 1984. Sydow received his B.S. and M.S. degrees in mathematics from Mankato State University in Minnesota. He joined Cray Research in 1979 and taught user-level software training courses until he joined the benchmarking department in 1981.

```

DIMENSION X(N,N),Y(N,N),Z(N,N)
DIMENSION XX(N+1,N+1),YY(N+1,N+1),ZZ(N+1,N+1)

-----(original)
DO 200 I = 1 , N
DO 200 J = 1 , N
      X(I,J) = X(I,J) + Y(I,J) * Z(I,J)
200 CONTINUE

-----(first modification)
DO 210 J = 1 , N
DO 210 I = 1 , N
      X(I,J) = X(I,J) + Y(I,J) * Z(I,J)
210 CONTINUE

-----(second modification)
DO 220 I = 1 , N
DO 220 J = 1 , N
      XX(I,J) = XX(I,J) + YY(I,J) * ZZ(I,J)
220 CONTINUE

```

Figure 2. Sample code and its modifications used to arrive at ratios given in Table 3.

# CORPORATE REGISTER

## Demand for Cray systems continues

Cray Research announced in March that Standard Oil Production Company has ordered a CRAY X-MP/28 computer system with SSD solid-state storage device. The system will be installed at Standard Oil's Technical Data Center in Dallas, Texas. It replaces a CRAY X-MP/24 computer system in operation since 1984. Also in March, Cray Research announced that Los Alamos National Laboratory in Los Alamos, New Mexico, installed a CRAY X-MP/416 system

during the last quarter of 1986.

In April Cray Research announced the installation of a CRAY X-MP/24 computer system at Lockheed Missiles & Space Company. The system replaced a CRAY-1/S system in operation since 1983. The new system will be used for the design and analysis of missiles and spacecraft.

Also in April, Cray Research announced that British Petroleum had ordered a CRAY X-MP/24 computer system. The system is scheduled for installation at the

London headquarters of BP Exploration Co. Ltd. in the fourth quarter of 1987. The CRAY X-MP/24 system will replace a CRAY X-MP/12 system in operation since 1985. The system will be used mainly to support the oil production activities of BP Exploration including research in reservoir simulation and seismic processing.

In April, Cray Research also announced the order of a CRAY X-MP/24 computer system by Electricite de France (EDF). The system is scheduled for installation in the fourth quarter of 1987. EDF is the

French national utilities company in charge of power plant design and construction as well as electricity production and distribution. The Cray system will be installed at EDF's research facility in Clamart, France. EDF will continue to operate a CRAY X-MP/216 computer system previously installed at its Clamart facility.

Also in April, Cray Research announced that Sandia National Laboratories in Albuquerque, New Mexico, installed a CRAY X-MP/416 computer system with SSD solid-state storage device. Cray Research announced in April that Lawrence Livermore National Laboratory also had ordered a CRAY X-MP/416 computer system and SSD solid-state storage device.

Cray Research announced in May the order of a CRAY X-MP/216 computer system with SSD solid-state storage device by the Deutsche Forschungsund Versuchsanstalt fuer Luft- und Raumfahrt (DFVLR), the West German national aerospace agency. The system will be installed in the third quarter of 1987 and will replace a CRAY-1/S computer system in operation since 1983.

Cray Research also announced in May the order of a CRAY X-MP/14 computer system by the Aerospace Corporation, a new Cray customer that provides general systems engineering and integration services for military space systems, principally to the Space Division of the U.S. Air Force Systems Command. The system will be installed in the fourth quarter of 1987 at the company's computer facility in El Segundo, California.

In May Cray Research announced that CINECA, a consortium of 13 Italian universities, had installed a CRAY X-MP/48 computer system with SSD solid-state storage device. The system was installed in the second quarter of 1987 in CINECA's data processing center in Casalecchio di Reno, Bologna.


In May Cray Research announced that Aerospatiale, a French aerospace company, has ordered a CRAY X-MP/14se computer system. The system will be the first CRAY X-MP/14se system installed

at a customer site. The computer system will be installed in the fourth quarter of 1987, pending export license approval, at Aerospatiale's facility in Toulouse, France. Aerospatiale designs and manufactures commercial aircraft such as the Caravelle, Concorde, and Airbus.

### **Contest recognizes quality research**

The Concours Seymour Cray is an annual research proposal competition sponsored by Cray Research France, S.A. to acknowledge excellence in research by French scientists. Although the independent jury felt that last year's proposals were not of sufficient quality for a grand prize, the contest organizers are not deterred. "Next year we will sponsor a contest in Switzerland as well as in France," said Robert Levy, president of Cray Research France S.A. "Last year we awarded consolation prizes to four young people. Our goal is to encourage — not discourage — participation."

Cray Research France limits its participation in the contest to financing and logistical support, while an independent jury develops the competition rules, reviews the submissions, and determines the award recipients. Prizes are given in three disciplines: computer architecture,



Poster announcing the Seymour Cray Competition in France, sponsored by Cray Research France S.A.

numerical simulation, and microelectronics. A research project titled "A Method of Automatic Conception for the Systolic Architectures" received the top prize awarded in 1986, given in the computer architecture category.

"Cray Research will continue to grow and improve as a company largely because customers explore new applications of our products," Levy said. "The Seymour Cray Competition is one way we can acknowledge the researchers who do that exploration. From them we can learn how our systems are being used now and how they will be used in the future. And that knowledge will help us provide better products and service to our customers."


### **Hori named president of Cray Japan**

Cray Research announced in April that it has named Yoshikazu Hori president of Cray Research Japan, Limited, the company's Japanese subsidiary. As president, Hori will be responsible for all Cray Research activities in Japan. He previously was senior director and sales and marketing division manager of Nippon Fairchild K.K., a division of Fairchild Semiconductor Corporation, a subsidiary of the Schlumberger Company.


Hori replaces James Otis, who served as president of Cray Research Japan, Limited for four years and recently returned to Minneapolis to serve in a corporate position.

"Current Cray customers in Japan, including customers in the automotive, electronics, and service bureau industries, are very pleased with their Cray systems," Hori said. "And the market is growing rapidly because of the industrial and economic environment in Japan. But the international trade situation has created stresses — some good, some bad. As a result, I think Japanese industry will have to become even more involved in research and development in the future. That will provide us with new opportunities to inform Japanese researchers about the value and performance of Cray systems."

# APPLICATIONS IN DEPTH



Seismic data displayed with the UNIRAS SEISPAK package.



Map projections displayed with UNIRAS GEOPAK.

## UNIRAS graphics for scientific applications

UNIRAS graphics software from UNIRAS Inc. serves a variety of computer graphics needs. Applications include business graphics, two- and three-dimensional mapping, image processing, and seismic display. UNIRAS graphics software runs on Cray computer systems under Cray Research's COS and UNICOS operating systems.

UNIRAS software is device-independent and written in Fortran. UNIRAS packages contain an internal color table that can, at user choice, be loaded with RGB, CMY, or HLS. The Hue-Lightness-Saturation color module enables users to easily define an infinite range of colors. To accommodate devices such as ink jet plotters and pen plotters, which do not support this intensity variation, UNIRAS uses an advanced dither pattern technique to simulate various intensity levels for ink jet devices and a cross-hatching technique for pen plotters. This adaptability makes the varie-

ty of colors in a UNIRAS program device independent.

Packages in the UNIRAS system that are of particular interest to Cray computer users include several designed for petroleum-related applications. Among them is GEOPAK, an advanced contour and grid mapping package capable of producing full-color, high-resolution, smooth-shaded maps in two and three dimensions with little computer programming. The package's features include determination of subsurface structure, faulting, and smooth shading between contours.

GEOPAK can be used to generate many types of maps, including

- Two-dimensional contour maps with annotated contour curves and shading between contour levels
- Three-dimensional contour maps, with color related to z-values
- Four-dimensional contour maps, with color representing a fourth variable

- Projected ceiling contour maps
- Projected time slices
- Three-dimensional block diagrams
- Three-dimensional polygon area maps (perspective)
- Fast two- and three-dimensional contour maps based on polygonal methods

The GEOPAK routines are written in ANSI Fortran, and may be called from any language supporting linkage to a Fortran library. Extensive system defaults are utilized, enabling the novice user to obtain a first map with only a few lines of code. GEOPAK can display any variables over Cartesian coordinates or in three-dimensional space. Its applications include topological, geological and geophysical mapping; meteorologic, thematic and demographic surveys; and representation of complex technical data.

Another UNIRAS package for petroleum-related applications, SEISPAK, is a collection of Fortran-callable subroutines for the construction of custom ap-

plication programs in seismology, geophysics, and well-logging. SEISPAK offers many options for accommodating input data, display techniques, annotating features, and postprocessing. The package supports dual attribute plotting of seismic sections and construction of seismic cubes. It can be called from any language supporting linkage of a Fortran library, and calls to subroutines from other UNIRAS packages can be combined freely with SEISPAK calls.


To enhance the versatility of UNIRAS packages, UNIRAS drivers for seismic rasterizing devices are "tunable" for plot speed and quality optimization. Options for optimization include

- The use of formatted seismic primitives, passing them between different computer systems before rasterization by an external device
- The support of double buffering and readback mode, sending the resultant (hardware-generated) rasters back through the rasterizer to a plotter
- The ability to use the hardware rasterizer in readback mode but save the rasters for subsequent plotting by a UNIRAS software device driver
- The use of specifically tuned drivers for binary I/O between different computer systems, such as between Cray and VAX systems, enabling the best use of memory and I/O

For more information on using UNIRAS graphics packages with Cray computer systems, contact Jan McDaniel, UNIRAS, Inc., 5429 LBJ Freeway, Suite 650, LB 144, Dallas, TX 75240; telephone: (214) 980-1600.

### Oasis graphics on Cray systems

The Oasis graphics package is a set of tools for producing still and animated computer imagery. The heart of the system is the Clockworks utility, which takes both a geometric and a kinematic description of a scene and, using one of several rendering techniques, produces finished raster images. A ray tracer and A-buffer are provided in addition to facilities for producing simple motion test imagery. Tools also are provided to



*Architectural study created with the Oasis graphics package.*

model geometry, develop and preview scripts, and display the generated raster images. A wide range of graphics devices are supported including Silicon Graphics IRIS, Evans and Sutherland PSS300, Ramtek 9400, 9460, and 9465, Raster Technologies Model One, Tektronics 4014, and Sun and Apollo workstations.

Oasis is built on the concept of an open and extendable framework of tools and utilities. Programs communicate with each other through a standard set of file formats, simplifying the development of new utilities. An added benefit is that importing information from an external source is accomplished simply by writing a "liaison" program to convert the incoming data from its original format into that of the Oasis system. The use of files also allows different parts of the animation task to be performed on different machines, such as modeling on a workstation and rendering on a Cray system.

A scene consists of a script written in an object-oriented language and a set geometry to be acted upon by the script. The geometry is created using a constructive solid geometry system in which the primitive objects are polygons, iso-valued surfaces, and superquadrics. Any number of light sources may be defined and may be shuttered and finite-

volumed. Objects can have two- or three-dimensional texture maps applied to them and may be viewed by any number of user-defined cameras.

Oasis is available for demonstration and benchmarking at Cray Research's applications department in Mendota Heights, Minnesota. The package runs on the COS, CTSS, and UNICOS operating systems. For more information on using Oasis with Cray computer systems contact Gray Lorig, Applications Department, Cray Research, Inc., 1333 Northland Drive, Mendota Heights, MN 55120; telephone: (612) 681-3645.

### The PATRAN System available on Cray systems

The PATRAN System is an open-ended, general purpose, three-dimensional, mechanical CAE environment available on Cray systems from PDA Engineering, Inc. The PATRAN System links engineering design, analysis, and results evaluation functions, and is composed of PATRAN Plus, Application Modules, Application Interfaces, and Gateway Utilities.

Every PATRAN user receives PATRAN Plus, the core of the PATRAN System.

# APPLICATIONS IN DEPTH

PATRAN Plus includes an advanced hybrid solid modeler, sophisticated graphics imaging capabilities, a finite element modeler, interactive color graphic representation of analysis results, and a unique open-ended "gateway" architecture that permits access to virtually any design, analysis, or manufacturing software.



The PATRAN System's closely coupled modules provide a comprehensive range of engineering capabilities, including

- Geometry construction using parametric cubic representation of primitives, Boolean operations and freeform lines, surfaces and solids
- Finite element modeling, including node and element generation throughout a surface and/or solid model
- Graphic display of analysis results
- Imaging including hidden-line plots and light source shading
- A gateway capability to interface with other software packages and analysis codes.

Configuring optional Application Modules can provide linear statics and dynamics analysis; thermal analysis, composite material design and analysis; rigid body dynamics modeling, X-Y and bar chart plotting of analysis results, section property calculations and engineering tomography using X-ray/CATSCAN pixel data for image processing and non-destructive evaluation. One Application Module is an expert system for modeling offshore oil structures.

Extensive construction, viewing, and editing features have been integrated into the PATRAN System, providing a friendly interface for design analysis modeling and evaluation. Fast efficient use of color graphics enhances model verification and results evaluation. More than 26 programs interface with the PATRAN System.

For more information on using the PATRAN System with Cray computer systems contact Rick Caselli at PDA Engineering, 2975 Redhill Avenue, Costa Mesa, CA 92626; telephone: (714) 540-8900.



Renderings of a bobbin (top) and piston/connecting rod/crankshaft model (bottom) created with the PATRAN System.

# USER NEWS

## Cray system tunes radio images

Given a good eye and a clear night sky, anyone can practice the ancient science of optical astronomy. But visible light coming from space, even when concentrated by large optical telescopes, can only partially reveal the workings of the heavens. This is because much astronomical information is carried by electromagnetic waves that lie outside the visible light spectrum. Radio waves, for example, carry important astronomical information, and today radio astronomers are analyzing radio emissions to learn about distant objects and events. The world's most powerful radio telescope is the Very Large Array (VLA) in New Mexico, which is operated by the National Radio Astronomy Observatory (NRAO). By combining radio astronomical data from the VLA with Cray computing power, NRAO astronomers have been able to produce some of the most detailed radio astronomical images to date.

Because the VLA is an *aperture synthesis* radio telescope, NRAO astronomers must process extensively the data that it

gathers. This type of telescope comprises many antennae distributed over a large area. The array samples incoming wave fronts at discrete points, and the sampled data then must be computer-processed into a continuous image. (A single-antenna radio telescope, as opposed to a multiple-antenna synthetic one, would require an aperture many kilometers in diameter to achieve acceptable resolution. Even if technically feasible, building such a telescope would be prohibitively expensive.)

The VLA consists of 27 antennae, each 25 meters in diameter, arranged in a Y configuration. In a typical observation, the array measures the correlations between all pairs of antennae, producing up to 200,000 bits of information every 10 seconds. The data must be edited, calibrated, and Fourier transformed into images of the radio sky, each having up to 4096 pixels on a side. An astronomer's observing run is typically eight hours long and may produce anywhere from one to hundreds of such images.


Because the wavefront is sampled at discrete points, the resulting images are degraded by the gaps between the anten-

nae. For a large array, the processing needed to compensate for this degradation requires advanced computing hardware to complete the job in a reasonable time. Many operations involved in processing synthesis-array data require iterative use of large, two-dimensional Fourier transforms. Although relatively few arithmetic operations are performed per data element, a very large amount of data is involved — more than can be kept in central memory.


"Data we recently processed to create an image of the supernova remnant Cassiopeia A required optimizing four million variables, in this case the pixels for a final image," explains Tim Cornwell, an NRAO astronomer who uses Cray systems to process VLA data. "It's a constrained optimization because we are looking for a result that not only agrees with the observed data but also has maximum entropy, that is, has the lowest contrast."

The problem requires that data and intermediate computational results be moved from the CPU to disk and back, often repetitively. "An optimal supercomputer configuration for handling this

# USER NEWS



Radio image of the supernova remnant Cassiopeia A. The original "dirty" image was processed in 20 minutes of CPU time on a CRAY X-MP computer system to produce this image. (Courtesy NRAO/AUI. Acknowledgment: P. E. Angerhofer, R. Braun, S. F. Gull, R. A. Perley, R. J. Tuffs)



Radio images of the galaxy Cygnus A. The initial map produced after the collected data were Fourier transformed is shown at left. This image then was deconvolved, or "cleaned," to produce the center image. The right image was produced by self-calibrating the center image to remove atmospheric distortion. The right image clearly reveals the filamentary structures connecting the galaxy to the lobes. (Courtesy NRAO/AUI. Acknowledgment: R. A. Perley, J. W. Dreher)

problem would have the largest practical amount of fast memory, and a disk-to-CPU transfer rate comparable to the speed of its arithmetic pipelines," explains Ronald D. Ekers, NRAO assistant director. "Researchers often must see the intermediate image iterations as the image processing proceeds, to control the path taken by the algorithm. Hence, a need exists for interactive supercomputing and the ability to display large volumes of data. Front-ended supercomputer configurations are not well suited to this application."

To help meet its supercomputing needs, the NRAO received an NSF grant to develop software for processing images on a CRAY X-MP computer system. "We can now take images from the raw observational data and process these images with our two primary algorithms," explains Paul A. Vanden Bout, NRAO director. "The CLEAN algorithm removes the visual effects of sidelobes from the images and a self-calibration algorithm removes atmospheric distortion."

The NRAO's application of compute-intensive image processing to radioastronomical data has generated some significant results. A processed image of the radio galaxy Cygnus A revealed important information about the structure of radio galaxies. For years, astronomers have been aware of radio galaxies, which often were observed in the form of two radio-emitting lobes, equidistant from a central galaxy visible in normal light. Astronomers speculated that the lobes were either ejected by the central object or powered by it in some way.

A cleaned and self-calibrated image of the galaxy reveals a tenuous thread of matter that connects the galaxy to the lobes — but this is seen only in the image that has been cleaned and self-calibrated. "Scientifically, such an image is priceless," says Vanden Bout. "It unequivocally demonstrates the role that the central galaxy plays in continuously fueling the radio lobes. The extraordinary energy flowing from the galactic center leads to the critical inference that a black hole exists in the center of this particular radio galaxy."


In the case of the supernova remnant Cassiopeia A mentioned previously, NRAO astronomers used a Cray system to produce the most detailed radio image yet of any object. The supernova that produced the remnant may have been first detected in 1680, and is still visible from Earth. Last year NRAO astronomers processed radio images of the remnant that reveal previously undetected structural features. Scientists had assumed that gases thrown out by the explosion of a dying star formed a huge, smooth shell. But the processed pictures of Cassiopeia A show the shell was mottled by bumps and holes as denser blobs of heavier gases punctured it from within. Studying the structure of the gas shell should reveal information about the motion of debris from the supernova explosion.

"Thanks to the NSF we will continue to receive access to Cray systems," says Cornwell. "In the near future we will be running on a system at Westinghouse Electric Corporation and we hope eventually to make use of the CRAY-2 system at the NASA Ames Research Center. As long as the access is available to us, we can continue to explore the limits of large-scale scientific image processing." (Portions of this article appeared in *ZeroOne Supernet*, Spring 1986.)

## University center emphasizes graphics

The National Center for Supercomputer Applications (NCSA) at the University of Illinois at Urbana-Champaign is developing advanced graphics capabilities to help researchers make the most of the center's Cray computer system. The advanced visualization facility will enable researchers to interact with and explore computed data in real-time through graphical animation and to produce videotapes and films of research results.

The hardware that handles graphics at the facility is configured around an Alliant FX-8 minisupercomputer. Attached to this system are two Raster Technology frame buffers that will display color images at twice the resolution of regular television screens. Connected to these



Phase one hardware configuration of the NCSA advanced visualization facility.

frame buffers on a video pipeline is an Abekas digital video storage device for real-time animation. Planned additions to the visualization facility include a high-resolution film recorder for 16mm, 35mm, and 4 by 5-inch film, and a Pixar Image Computer.

Software for the Alliant system includes Wavefront Technologies' motion choreography and image rendering software to produce images of time-dependent three-dimensional data; Precision Visuals' DI-3000, an integrated system of graphics software tools based on the ACM/SIGGRAPH CORE graphics system; and OASIS, a three-dimensional rendering package courtesy of Cray Research, Inc.

The visualization facility is configured so that researchers need not understand computer graphics technology or the details of scientific movie production. Artistic consultation also is provided to help render data effectively. A technical staff assists with the translation of data between the simulation and the available software. Researchers thus are freed from the task of having to learn quickly how to use a visualization package. Post-production staff produce the final film or video, and other staff are available to modify or develop special-purpose software.

"Our future plans include upgrading the Alliant system to about a 200-MFLOP machine," explained Nancy St. John,

# USER NEWS

manager of the scientific visualization program at NCSA. "We also will be working to integrate our software and hardware, to make data transfers among machines transparent to users."

Establishment of the NCSA's advanced visualization facility symbolizes the importance of graphics for virtually all applications of advanced scientific computing. As computer graphics technology becomes faster, more highly resolved, and more interactive, computational methods will provide ever-greater opportunities for engineers and research scientists to solve problems and make discoveries.

## Graphic ad nominated for Clio

A promotional television ad for Bell Atlantic that has been nominated for a Clio


award features graphics generated on a CRAY-2 computer system. The ad was produced by R. Greenberg Associates of New York for Ketchum Advertising of Philadelphia. The graphics were generated at the University of Minnesota using proprietary software that R. Greenberg Associates previously had run on the university's CRAY-1 system. Moving to the CRAY-2 system has added versatility and speed to the company's computer graphics capabilities.

"In-house, we use Sun Micro workstations networked to a DEC VAX system in a UNIX environment, and we do our programming in C," explained Chris Woods, senior producer at R. Greenberg Associates. "Working with the CRAY-2 system has been very exciting. The Cray implementation of UNIX has worked out very well, and the Cray C compiler enables us to modify our software quick-

ly. Last year's work was much less painful than the previous work we did on the CRAY-1 system using the COS operating system and non-supported C. We are working in a much better, stronger environment now."

R. Greenberg Associates used the CRAY-2 system to produce three television ads for Bell Atlantic last year. The spot titled "Integration" is a finalist for this year's Clio awards, given for excellence in broadcast advertising. It received nominations in two categories: graphics and computer graphics. "The graphics area judgement is very subjective, and the computer graphics category is judged primarily on technique; it's a technical category," Woods said.

The award winners were announced as this issue of CRAY CHANNELS went to press.



Single frame from "Integration" television advertisement for Bell Atlantic produced on a CRAY-2 computer system.

---

## **WORLDWIDE BUSINESS CENTERS**

---

### **Domestic**

**Corporate Headquarters**  
608 Second Avenue South  
Minneapolis, MN 55402  
(612) 333-5889, telex 4991729

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

**Central Region Office**  
5350 Manhattan Circle  
Boulder, CO 80303  
(303) 499-3055  
Albuquerque  
Chicago  
Colorado Springs  
Detroit  
Minneapolis/St. Paul  
St. Louis

**Eastern Region Office**  
4041 Powder Mill Road, Suite 600  
Calverton, MD 20705  
(301) 595-5100  
Atlanta              Huntsville  
Bernardsville        Laurel  
Boston                Pittsburgh  
Cincinnati          Tampa  
Darien

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

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

### **International**

**Cray Asia/Pacific Inc.**  
Suite 1003  
Sutherland House  
3 Chater Road Central  
Hong Kong

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

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

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

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

**Cray Research Japan, Limited**  
W 1154, Shin Aoyama Building  
1-1 Minami-Aoyama 1-chome  
Minato-Ku, Tokyo 107, Japan  
(03)-403-3471, telex 28893

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

**CRAY**  
**RESEARCH, INC.**