SUMMER 1993
VOLUME 14, NUMBER 4

SCD COMPUTING NEWS

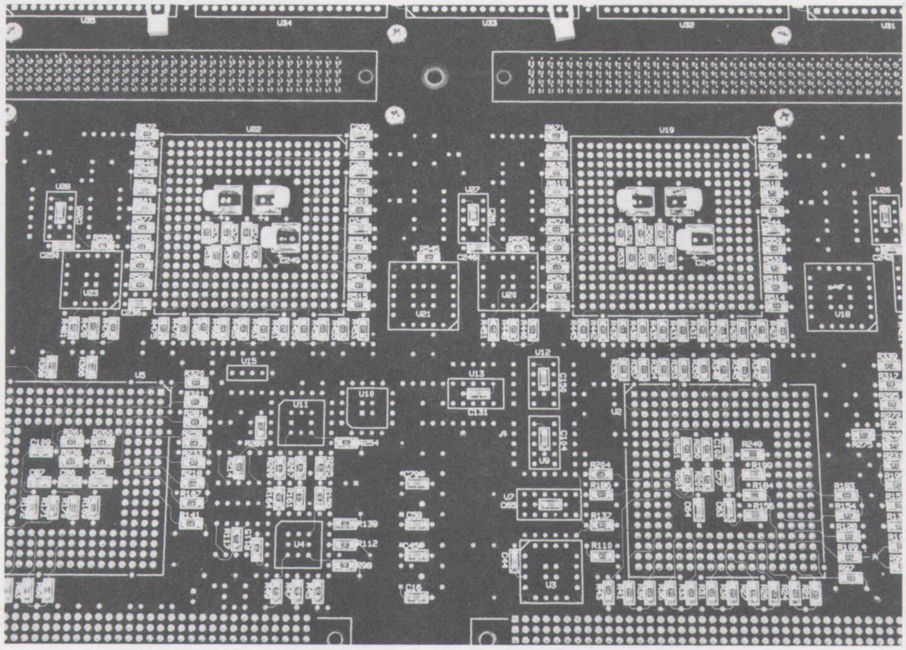

Special Issue: Ninth SCD User Conference

- **HIGH-PERFORMANCE COMPUTING IN THE '90s**
- **SCD EXPLORES LEADING-EDGE COMPUTER ARCHITECTURES**
- **NCAR GRAPHICS 3.2 IS HERE!**
- **USER INTERFACE TO IMPROVE**

SCIENTIFIC COMPUTING DIVISION
NATIONAL CENTER FOR ATMOSPHERIC RESEARCH

About the cover

Looking like an aerial view of a city at night, the patterns in the cover image are formed by vector processors on a Connection Machine circuit board. This particular board is a spare part for the CM-5 in the NCAR Computer Room. (Board courtesy Thinking Machines Corporation; photo by Carlye Calvin.)

SCD Computing News welcomes cover graphic submissions. Please send your graphic, along with a description including SCD facilities or software used, to Christine Guzy, Mesa Lab; or send e-mail to guzy@ncar.ucar.edu; or call Christine at (303) 497-1826.

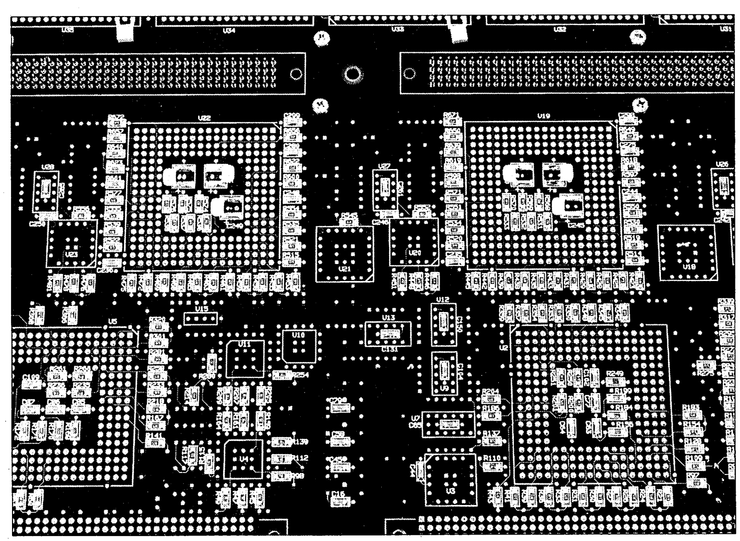

Volume 14, Number 4

SCD Computing News is published seven to ten times per year by the Scientific Computing Division (SCD) of the National Center for Atmospheric Research (NCAR) in Boulder, Colorado. NCAR is operated by the University Corporation for Atmospheric Research and is sponsored by the National Science Foundation. Please submit articles, photos, graphics, and letters to *SCD Computing News*, NCAR/SCD, P.O. Box 3000, Boulder, CO 80307-3000.

Staff

Lynda Lester, managing editor
Brian Bevirt, writer/editor
Nancy Dawson, writer/editor
Christine Guzy, production coordinator
Jacque Marshall, production assistant
Juli Rew, writer/editor
Jim Petruzzelli, writer/editor

Δ Software change articles: The delta symbol (Δ) appears next to the title of articles that may affect how you run your jobs.

Conventions: For a list of conventions used in this newsletter, see the Documentation department.

Printed on recycled paper (50% mill waste, 10% post-consumer waste) with soy inks.

PRINTED WITH SOY INK™

FEATURES

3 **Information bandwidth high at Ninth SCD User Conference**
In the supercomputing world of the nineties, technology makes vector strides

4 **Weekend update with Pete Peterson**
New machine architectures at NCAR, and other highlights from FYs 1992–93

7 **Nontrivial pursuit: Supercomputing in the nineties means fast action**
Rate of change poses challenge, opportunity

10 **CM-5 alive!**
Connection Machine is up and running in the NCAR Computer Room

13 **Don't try this at home: The making of a CRAY-3**
It's sleek, it's grey, it's a four-foot-tall postmodern cabinet

15 **Are cluster farms green?**
Uh HUH!

18 **The MSS: NCAR's 'Big Gulp'**
The scientific computing community has a powerful thirst for data

19 **NCAR's data archive keeps growing and improving**
Data Support Section upgrades its repository of meteorological and climate data

20 **Good news: User interfaces to become easier**
Access to SCD's distributed computing environment will become more UNIX-like

22 **Quick takes**
A photo gallery of conference highlights

24 **NSFNET: Big changes on the horizon**
High-bandwidth links between supercomputing centers will support advanced applications

26 **New and improved: NCAR Graphics Version 3.2**
Many utilities have new functionality, better output quality

30 **SCD video facilities launch a new era in scientific visualization**
How to make a good video, and more

33 **Why'd you do it? Users cite rationale for attending SCD conference**
"It was the CM-5 thing that pushed me over the edge . . ."

SCD Calendar

October

14–15 SCD Advisory Panel

November

1–3 NCAR Graphics Class

15–19 Supercomputing '93, Portland, Oregon

22 SCD Users Group (SCDUG)

December

13–15 NCAR Graphics Class

January

5–7 NCAR Graphics Class

For more information about these calendar items, send e-mail to scdinfo@ncar.ucar.edu.

SOFTWARE NEWS

36 **High-Performance Fortran includes Fortran 90**
HPF is based on Fortran 90, with added functionality for parallelism

38 **IMSL Fortran Library Version 2.0 is more efficient, accurate Δ**
Improvements made to a number of routines

40 **IRJE and MIGS users: Avoid Monday bottlenecks on MSS jobs**
No more "giving up waiting for status" messages

DEPARTMENTS

43 **Hints from the SCD Consultants**
Get date/time information with Fortran-callable routines on Y-MPs

45 **Systems News**
Obtain detailed usage reports with new 'gaus' utility • Memory limits increased for small memory subqueues

45 **MSS News**
More powerful than 'msls': 'msinfo' • The 'lread' and 'lwrite' commands are mothballed

46 **Graphics News**
Note this TAGS correction

46 **Training News**
Learn NCAR Graphics

47 **Documentation**
Philosophy, structure, strategy: "NCAR Graphics Fundamentals" is here • NCAR Graphics tutorial shows how to contour and map • How to access the SCD Daily Bulletin • Obtain SCD documentation via anonymous FTP • Conventions used in this newsletter

Information bandwidth high at Ninth SCD User Conference

by Lynda Lester

Imagine a great supercomputing room: production machines share the floor with strange new architectures, reliability meets experimentation, past greets future. At Control Central in the heart of NCAR, CRAY Y-MPs are stationed near "killer micros"; a monolith with flashing red lights stands opposite a device that seems to be half circuitry, half aquarium. This is the high-performance computing world of the nineties, where change is nontrivial and technology makes vector strides.

To investigate the challenges and opportunities therein, 126 participants came to the Ninth SCD User Conference, held June 11, 12, and 14 at the Mesa Lab. Scientists and software engineers heard the latest on data caches and network topologies, drank coffee, and took periodic computing room tours.

The conference, titled "High-Performance Computing at NCAR in the Nineties," covered:

- Trends in high-performance computing
- Recent and projected developments in SCD
- Experiments with leading-edge computer architectures
- What's happening with mass storage and data archives
- The creation of a new user interface for SCD's distributed computing environment
- Network futures
- What's new with NCAR Graphics

Information bandwidth was high. In the end, this conference on change in the nineties seemed to illustrate the adage: "Success is a journey, not a destination."

This issue's feature section highlights talks given at the conference.

Lynda Lester is a writer/editor in the Documentation Group within the SCD User Services Section.

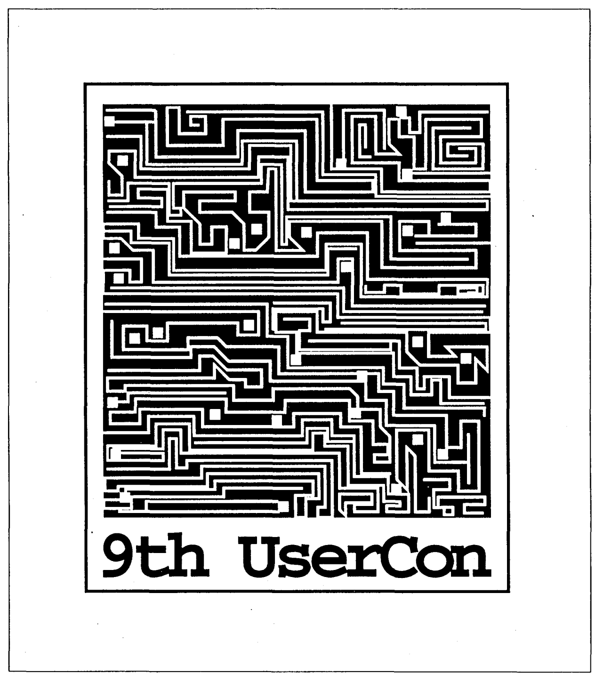

Weekend update with Pete Peterson

by Lynda Lester

On Friday, June 11, SCD deputy director Pete Peterson opened the first-ever SCD user conference held over a weekend. "We chose this schedule so more of our university friends could come," Pete said. Reduced air fares available with Saturday-night stayovers were appreciated by people attending the last part of the conference, a special CM-5 overview class that was held on Monday.

Leading off with an update on SCD, Pete reviewed highlights of Fiscal Years 92–93 and previewed what to expect in FYs to come.

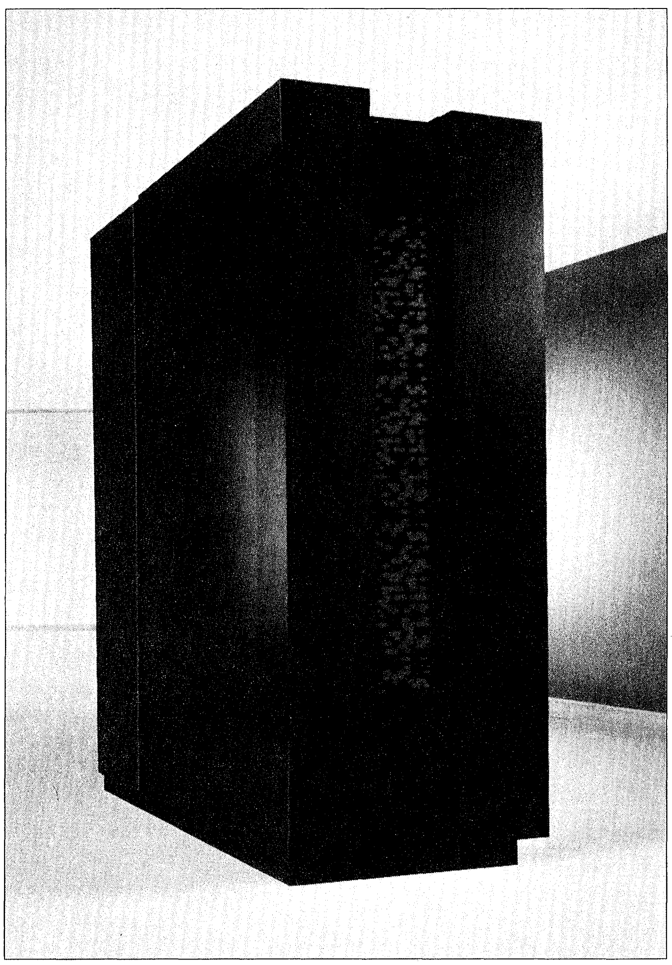

In April, NCAR acquired a 32-node CM-5. (Photo courtesy Thinking Machines Corporation.)

FY past

Big jobs for big iron: "We believe that we use supercomputers for jobs that require supercomputers," Pete said. While 70% of the wallclock time on the CRAY Y-MP8/864 (shavano) goes to jobs that need one CPU hour or more, over half goes to jobs needing four CPU hours or more. About a year ago SCD established a special queue on shavano for parallel processing; for five hours each day the entire machine is turned over to single jobs in multitasked mode. "We've heard from users that this queue is very effective in accomplishing the type of modeling we need to do at NCAR," Pete said, "and we've had requests to increase the number of hours dedicated to this queue."

MSS amasses data: One area of constant concern is the Mass Storage System (MSS). SCD has had a tremendous rate of growth in the number of bytes stored, and within a couple years we'll reach the system's maximum capacity. In 1992 SCD made a number of hardware and software enhancements to improve MSS operation and turnaround, but we need additional capacity—plus advances in storage technology—to help us out of the crunch.

Cluster passes muster: Microprocessor performance has advanced rapidly, overtaking the performance of both mainframes and minicomputers, and SCD expects this trend to continue. We are currently in production mode with an IBM RS/6000 cluster consisting of four nodes and a gateway, and have been offloading single-processor jobs from shavano to this cluster. Users can submit jobs interactively during prime shift and run batch jobs during off hours. Other benefits of the cluster include a software-rich environment and High-Performance Parallel Interface (HIPPI) access to the MSS. "At $2K per sustained megaflop, clusters appear to be cost-effective," Pete said. "We don't believe they will supplant or obviate the need for supercomputers, but we certainly feel they have a place in super-computing centers."

Figure 1. SCD distributed computing environment

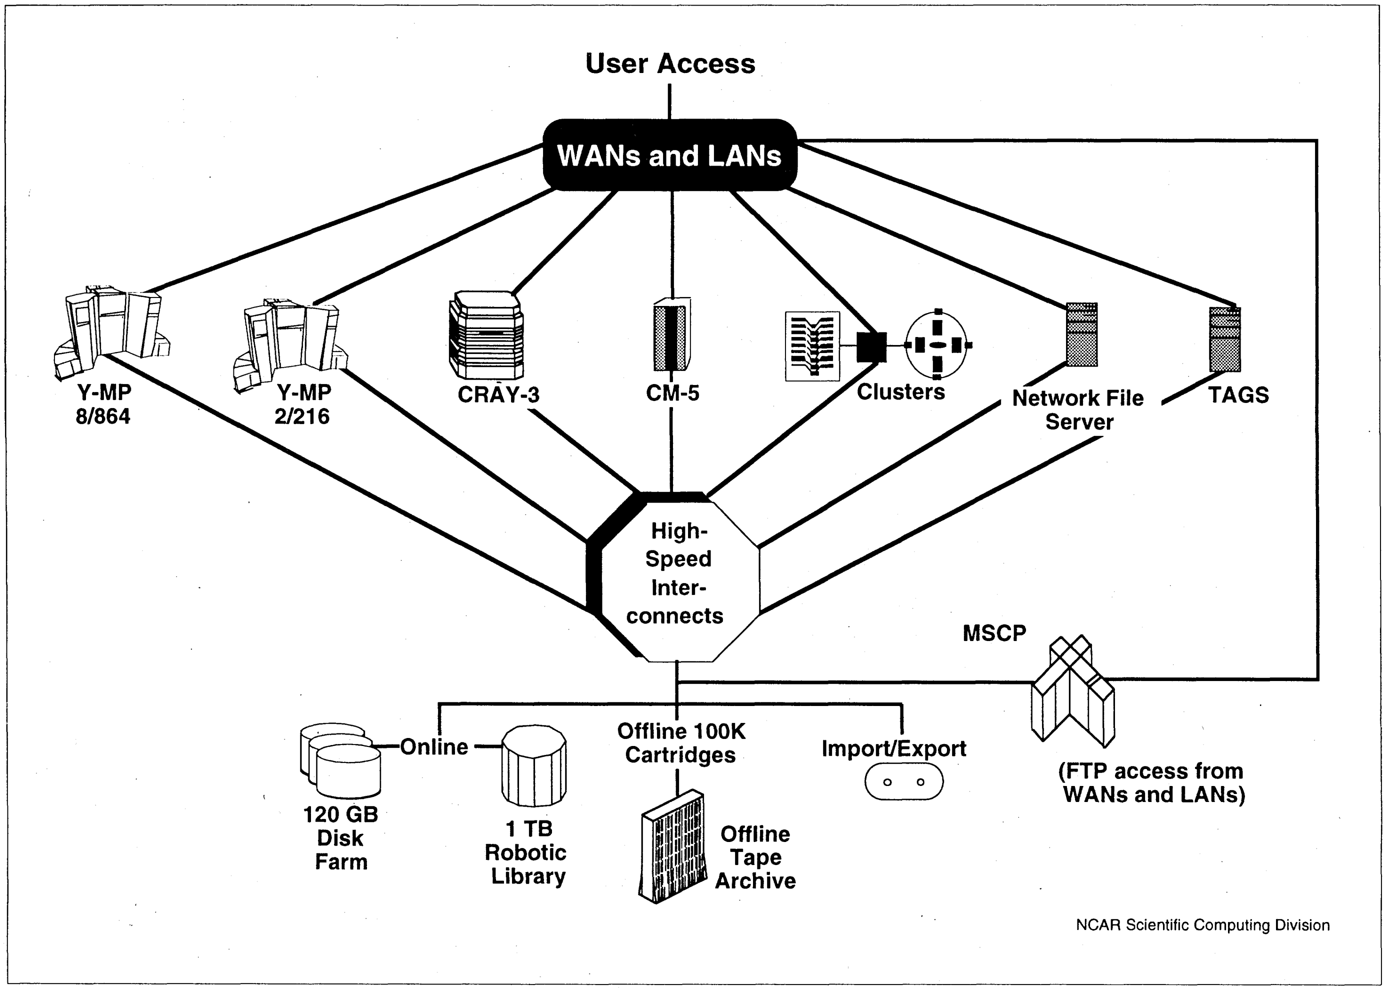

Quick bits: As part of a cooperative agreement with Cray Computer Corporation, NCAR users logged 15,000 hours on the CRAY-2 in Colorado Springs during FY92, Pete noted. He also described NCAR network upgrades in the past year, gave statistics on SCD output services, and talked about acquisition of the UNIX front-end computer (meeker) and the UNIX central file server.

FY present

Demand and supply: Pete explained that since fall 1989, SCD has tripled the amount of computing power available to NCAR users. This was necessary because leading-edge simulations have a huge appetite for CPU hours. For example, on the CRAY Y-MP2/216, which runs under the auspices of the Model Evaluation Consortium for Climate Assessment (MECCA), a typical allocation is 1,000 to 2,000 CPU hours. Each year, SCD receives more requests for large allocations: in 1990, users asked for 12,000 GAUs for long-running simulations; in 1992, they wanted 30,000. And that's only an indication of what's to come with next-generation coupled climate models.

To meet the ever growing demand, NCAR is constantly upgrading its roster of supercomputers (see Figure 1). In April SCD acquired a 32-node Connection Machine-5 (CM-5), and in May, a 4-processor CRAY-3 joined the lineup. An IBM 4-node cluster is now running in production mode,

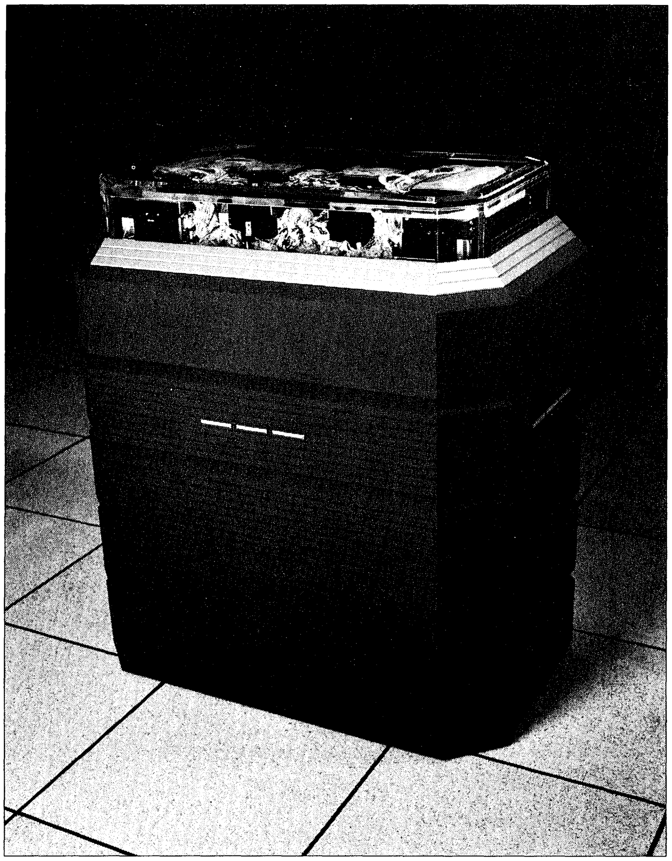

In May, a four-processor CRAY-3 joined the roster of NCAR supercomputers. (Photo courtesy Cray Computer Corporation.)

and we will soon install an IBM 8-node SP-1 cluster. *(Editor's note: The SP-1 was installed in July.)*

Reaching out over the Internet, NCAR has made even more computing capability available to users. As a member of the National Consortium for High-Performance Computing—NCHPC, or N-chips, as it's called—NCAR has access to parallel supercomputers at other locations, including a Kendall Square Research machine, an Intel Paragon, and three large CM-5s. Additionally, as a participant in the Department of Energy's Computer Hardware, Advanced Mathematics, and Model Physics (CHAMMP) initiative, NCAR has access to the 1,024-processor CM-5 at the Los Alamos National Laboratory.

Since fall 1989, SCD has tripled the amount of computing power available to NCAR users.

"What does this all mean? Ha ha!" laughed Pete. "It means an even greater load on the MSS!"—which, he noted, SCD will continue to develop to keep pace with the output from all these computers.

FY future

Quick bits: NCAR Graphics 3.2 has just been released, and Version 4.0 will come out next year. Other focal points for SCD in the near future include expanding the use of CD-ROM for popular datasets, further progress on a 35-year data reanalysis, the annual Winter Icing and Storms Project (which addresses real-time weather forecasting), and a complete refurbishment of the networking capability in the Mesa Laboratory.

The wrap-up: Pete concluded by noting that the rapid changes in technology are both a challenge and an opportunity for the computing community. (For further discussion of this idea, see "Nontrivial pursuit: Supercomputing in the nineties means fast action," next page.)

Lynda Lester is a writer/editor in the Documentation Group within the SCD User Services Section.

Nontrivial pursuit: Supercomputing in the nineties means fast action

by Lynda Lester

Bill Buzbee: The pulse of technology is going so fast—
Audience member: How fast is it?
Bill Buzbee: REAL fast.

The current rate of change in computing is stressing most established organizations, said SCD director Bill Buzbee in his talk, "Trends in supercomputing technology and potential benefits to climate modeling." Our task is to figure out which way things are going, and—if we do that right—how to keep up. But it is true that ultimately, challenge equals opportunity.

Four trends will characterize the nineties, Bill predicted:

- Single-processor applications will migrate from supercomputers to microprocessor systems—the so-called killer micros.
- Parallel processing and supercomputing will become synonymous.
- Storage technology will provide an order-of-magnitude increase in capacity and efficiency.
- When these advances are combined with networking, a framework will develop for global collaboration.

Beep-beep!

Like roadrunners, new workstations are small and fast; these machines will provide cost-effective augmentation to supercomputers in the nineties. Tests have shown that a five-day Community Climate Model simulation that executes in 1 wall-clock hour on a single CRAY Y-MP8/864 processor takes 10.7 wall-clock hours on an IBM RS/6000 node. However, Bill said, as we are painfully aware, 10 hours is about equivalent to actual turnaround time for 1-hour jobs on the Y-MP. That's one reason why serial codes are moving from supercomputers to killer micros.

Speed zone

The Y-MP in multiprocessor (parallel) mode is a solid gigaflop machine, Bill said. (A gigaflop is one billion flops, or floating-point operations per second.) The Y-MP is about 100 times faster than a typical workstation, and 40 times faster than the RS/6000. But even more speed—50 to 100 gigaflops—is needed for high-resolution coupled climate models that include data on atmospheric chemistry and dynamics, ocean dynamics, marine biogeochemistry, land surfaces, and the terrestrial ecosystem. And running decadal simulations at the regional level will take up to 1,000 gigaflops—the power of a teraflop computer.

But traditional shared-memory supercomputers ultimately hit a speed limit: 186,000 miles per second. Flop rate is limited by the fact that electrons can't go faster than the speed of light. At that point, the only way to improve performance is by parallel processing, where multiple pieces of a model are processed simultaneously rather than sequentially.

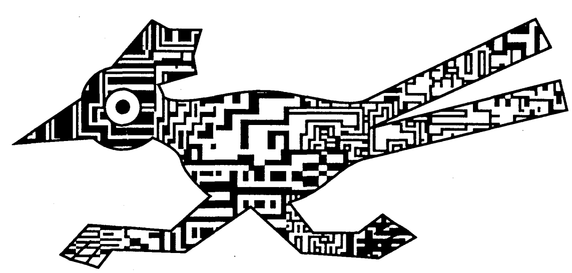

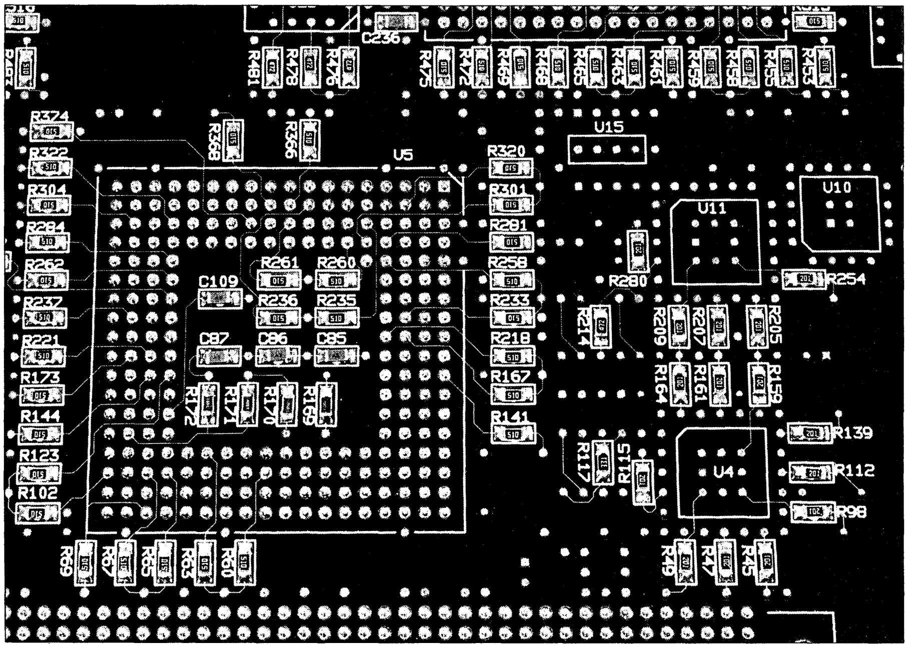

The large square pattern, left, is the underside of a vector unit on a CM-5 processor board. (Board courtesy Thinking Machines Corporation; photo by Carlye Calvin.)

NCAR is well on track to parallel processing, Bill said. The acquisition of the 32-node Thinking Machines CM-5 (a massively parallel machine) is an important step, and we've dedicated a significant part of the day to multitasking on the Y-MP. However, we've still got a lot to learn.

In the meantime, in view of the experience with parallel systems at NCAR, NASA Ames, Los Alamos, and other laboratories, SCD believes that shared-memory supercomputers will continue to offer the highest levels of cost performance available until 1994–95. By then, the next generation of parallel systems may be substantially outperforming traditional machines. (However, to outperform a traditional machine, the microprocessors in a parallel system must sustain at least 100 megaflops each.)

35 TB of data, with no end in sight

The amount of data in NCAR's Mass Storage System (MSS) is growing. In 1990 the MSS held 15 terabytes (TB); in 1991, 20 TB; and in 1992, 29 TB. This year it hit 35 TB, with no end in sight. Since users store about 500 bytes of data output per megaflop, with a next-generation computer, the archival rate could be as high as 100 TB per year—the total maximum capacity of NCAR's current MSS!

But storage is only half the problem, Bill said; access is the other half. Recently two scientists needed a year's worth of MSS data—daytime only, South Pacific. They had to search 1,300 satellite tracks on 400 cartridges, and it took several days. If the information had been online in the MSS "Silo," the job would have taken two to three hours. That's still not good enough: for the high-resolution decadal simulations of the future, we need a response time in minutes.

Advances in technology will allow for denser storage (about 50 gigabytes per cartridge as opposed to the current maximum of 0.4 gigabytes) and faster access. The new technology should be commercially available in three to five years—although needless to say, it's going to cost.

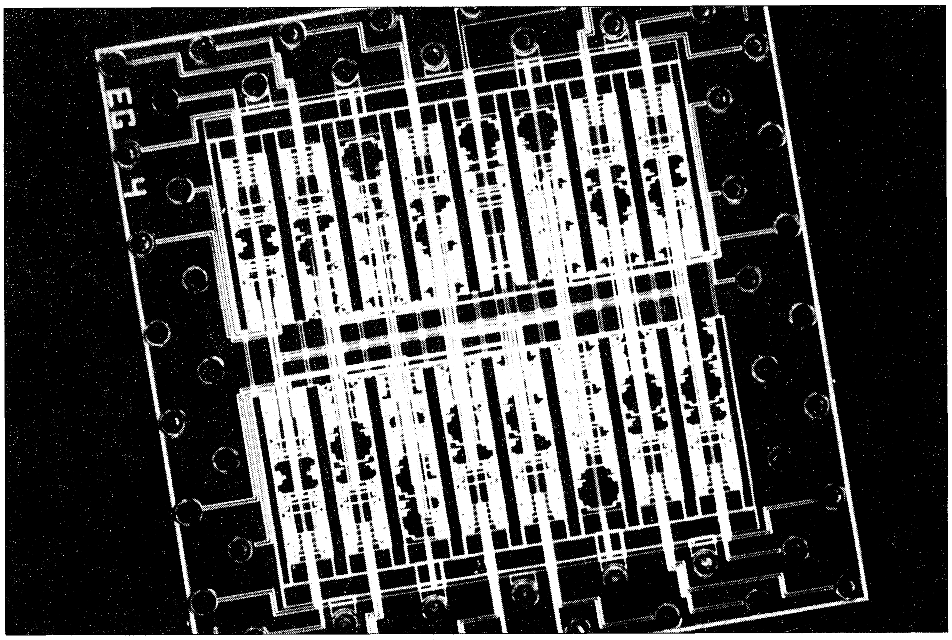

Gallium arsenide circuitry on a CRAY-3 logic die. (Photo courtesy Cray Computer Corporation.)

Meanwhile, out on the Internet . . .

Communication bandwidth on the networks is going up, up, up. From simple text-file transfer to video teleconferencing, interactive visualization, and composite imaging, network technology is evolving; multimedia database access and distributed computing have become common.

What does this mean for environmental science? The research community of the future will be handling huge volumes of data: observational data collected by space-borne systems as well as output data produced by large simulations on high-performance computers. Both types of data will be archived in mass storage systems.

Remote university researchers and scientists throughout the world will be able to access these archives, pull data back to their workstations, and conduct distributed analysis via the Internet. Feedback loops in all directions will provide a framework for collaboration in global change research.

Remote scientists will conduct distributed analysis via the Internet.

"It's my belief that we'll see this type of thing developing in years to come, not only in climate research but in other disciplines—high-energy physics, medicine, biology, the genome project, and so forth," Bill concluded.

"So in the end, the challenge of rapidly changing technology—a nontrivial challenge, to say the least—does in fact become our opportunity."

Lynda Lester is a writer/editor in the Documentation Group within the SCD User Services Section.

CM-5 alive!

by Lynda Lester

A 32-node Connection Machine-5 (CM-5) from Thinking Machines Corporation arrived at NCAR on April 27, passed acceptance testing on May 17, and is now available to users for conducting parallel experiments. (See the CM-5 photo on page 4.) Richard Loft, Thinking Machines site representative at NCAR, introduced the machine at the SCD User Conference in a one-hour talk and a one-day class.

Quick specs

Within the CM-5, computational and disk storage nodes are integrated into a single architecture that features scalable input/output (I/O), communications, memory, and processing. Each of its 32 nodes has four vector units and a Sun SPARC microprocessor that can perform 22 million instructions per second. The machine has 1 gigabyte of internal memory (32 megabytes per node) and performs at a peak speed of 128 megaflops per node.

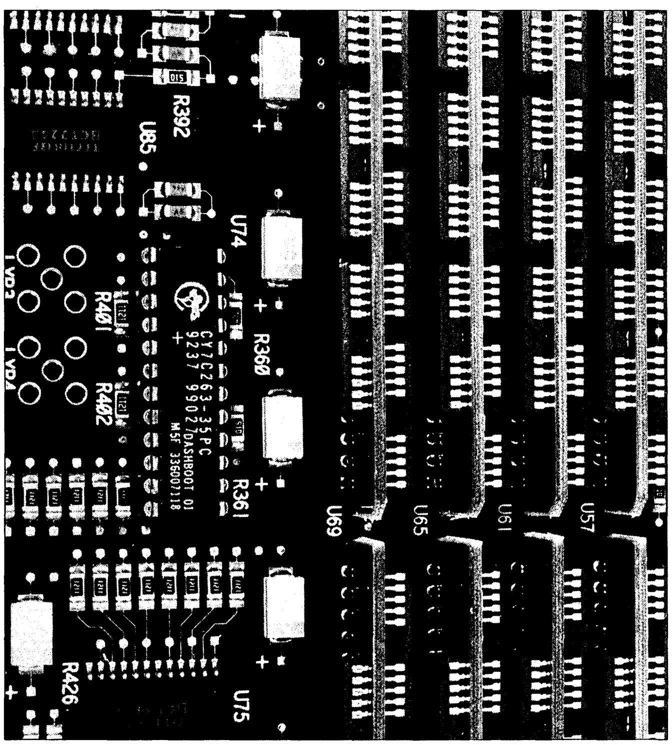

Close-up view of a CM-5 board shows circuitry, left, and memory chips, right. The chips provide 32 megabytes of memory per node for the CM-5 vector units. (Board courtesy Thinking Machines Corporation; photo by Carlye Calvin.)

"Prism is one-stop shopping for your code-development needs."

The CM-5 runs CMost, a Sun-based UNIX operating system that can service multiple users on a time-shared basis. ("CMost stands for CM Operating System—T," Rich said. "I suspect the T was added, since most is a superlative.")

Massively parallel processing programming paradigms

The CM-5 supports two programming models, data parallel and control parallel:

- The *data parallel* model has been primarily identified with SIMD (single instruction, multiple data). In this model, a single program runs on all the nodes, which execute the same instructions at the same time. This is a global method of programming that utilizes CM Fortran, C*, and *Lisp. (However, it can also exploit SPMD—single program, multiple data—and message passing.)

- The *control parallel* model has been primarily identified with MIMD (multiple instruction, multiple data). In this model, each node acts on sets of instructions or entire programs that it executes at its own speed. This is a nodal method of programming that utilizes message passing and FORTRAN 77 or C. (However, it can also exploit SPMD and CM Fortran.)

A prismatic approach to code development

Prism is the CM-5's parallel programming environment. From within Prism, you can set breakpoints in your programs and run performance evaluations; you can also inspect large arrays of data, either numerically or visually. "Prism is a debugger and more," Rich said. "It's one-stop shopping for your code-development needs."

Speed's good

CM Fortran, which is FORTRAN 77 with Fortran 90 extensions, is the language most CM-5 users at NCAR will want to use. "Fortran 90 extensions are what you use to access the CM-5 vector units; Fortran 90 array syntax is how you get performance out of the machine," Rich said. "If you're going to be using the CM-5, you should learn CM Fortran. Speed's good."

Libraries: What's up

The following libraries are up and running on the CM-5:

- The CM Scientific Software Library (CMSSL) consists of highly optimized computational kernels for data parallel applications.
- The CMX11 library supports parallel data output in graphical form to X11 servers.
- The CM/Application Visualization System (CM/AVS) library has a Prism visualizer interface and supports integration with AVS environments.
- The CMMD message-passing library provides MIMD communications primitives and global synchronization routines.
- The Parallel Virtual Machine (PVM) message-passing library is available for compatibility with other message-passing platforms.

Several portable libraries written in FORTRAN 77 will be also available on the CM-5. SCD does not expect them to perform well, however, since they run on the CM-5's front-end control processor and not on the nodes.

Local vs. nonlocal, or the inherent problems of grid topology

CM-5 nodes are connected by two internal networks, Rich said. The *control network* is optimized for operations such as broadcasting, reduction, and error reporting—operations that involve cooperation of all nodes. The *data network,* on the other hand, is optimized for node-to-node data transfers—regardless of the distance between the nodes or the number of source/destination pairs.

"Most massively parallel processors (MPPs) do well with local communications, but when you go nonlocal, performance suffers," Rich said. "Mesh architectures are popular because they can compute very fast in local transfers. That's great if people live near where they work. But a lot of data applications are like Los Angeles, where no one lives anywhere near where they work. Applications get wedged, traffic stops, and so forth. It's an inherent problem of grid topology. In general, transfers of large amounts of data among random, node-to-node pairs require special network design.

"TMC found by experience that the algorithms people use, such as spectral methods and semi-Lagrangian transport methods, have nonlocal communication patterns. Therefore, we tried to design a machine so that local and nonlocal communications would have similar interprocessor bandwidths. Ideally the absolute bandwidth would be as high as possible, so communication would vanish as a problem. We're not there yet, but we are on the way."

Take CMAXimum advantage of the CM-5

CMAX, which stands for CM array translator, is a tool that helps convert FORTRAN 77 into CM Fortran; it also provides a method for identifying data usage. "CMAX is automatic in the same way driving a car is automatic," Rich said. "You have to hold the steering wheel or you'll run into an abutment." Even with CMAX, he explained, codes will need modification to take best advantage of the CM-5.

Questions?

Rich is available for online conversation about CM-5 hardware and software. His e-mail address is loft@ncar.ucar.edu.

In addition, a CM-5 user group has been formed at NCAR. The group meets regularly to exchange technical information and discuss CM-5 issues. To receive meeting notices and subscribe to the mailing list, send e-mail to cm5ug-request@ncar.ucar.edu.

Lynda Lester is a writer/editor in the Documentation Group within the SCD User Services Section.

The CM-5 supports two programming models: data parallel and control parallel.

Bears—where?

Following the tradition of naming its supercomputers after peaks in the Colorado Rockies, SCD named the CM-5—specifically, the CM-5 control processor—"littlebear." It seemed then only natural to call the CM-5 compile server "teddybear"; soon, the CM-5 batch queue drew the nickname "carebear." But, said Thinking Machines site representative Rich Loft, he drew the line at "yogi."

Don't try this at home: The making of a CRAY-3

by Lynda Lester

It's sleek, it's gray, it's a four-foot-tall postmodern cabinet. Through a clear acrylic case on top, interconnect wires can be seen waving like kelp in a sea current. The current is made of liquid coolant, which dissipates integrated-circuit heat; the computer is the world's first CRAY-3, designed by Seymour Cray and installed at NCAR in May. (See the CRAY-3 photo on page 6.)

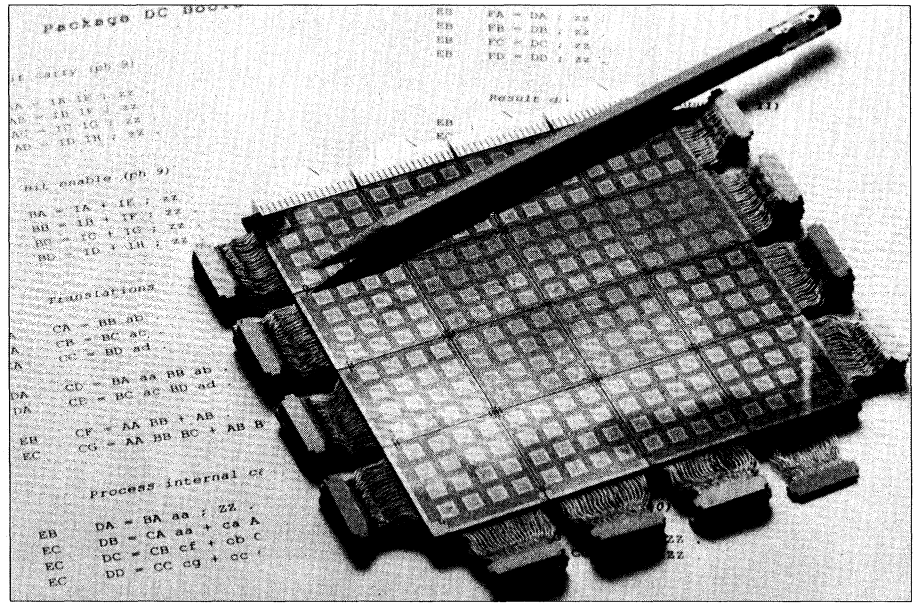

Old and new computational tools: pencil and CRAY-3 module. (Photo courtesy Cray Computer Corporation.)

In his talk "Introduction to the CRAY-3," Ron Larson of Cray Computer Corporation (CCC) used slides to show the manufacturing process for the electronic modules that make up the core of the CRAY-3.

Unlike conventional supercomputers, the CRAY-3 uses gallium arsenide (GaAs) for its logic circuits. CCC fabrication plants in Colorado Springs produce these logic circuits, as well as the printed circuit boards that go into the machine. During the production process, ultrasonic die bonders attach gold leads (tiny wire legs) to GaAs integrated circuit die. The leads of the die are then inserted through tiny holes in a one-inch-square printed circuit board and crushed, in a process that joins together die and boards. Sixteen die go on a board. Sixty-four boards, arranged in 16 stacks that are sandwiched around a set of four-inch-square plates, form a module; this packaging produces a logic density of 100 gates per cubic inch. Four modules comprise a processor. There are four processors in the CRAY-3 at NCAR.

In addition to its processor modules, NCAR's CRAY-3 has 64 memory modules and 4 input/output modules. All these modules have a combined volume of 405 cubic inches and burn 88 kilowatts of power.

"Eighty-eight kilowatts is a lot of energy to dissipate," Ron said, noting that the machine is filled with cooled Fluorinert that is pumped right through the component modules. "Now you know why Seymour is as good at refrigeration design as he is at computer logic."

The computer's two-nanosecond clock represents a tremendous advance in miniaturization and speed since the CRAY-1 hit the market in 1976—but it's only a preview of things to come. "I saw a quote from Seymour last week," Ron said, quoting the legendary inventor: "'As long as we can make them smaller, we can make them faster.'"

Lynda Lester is a writer/editor in the Documentation Group within the SCD User Services Section.

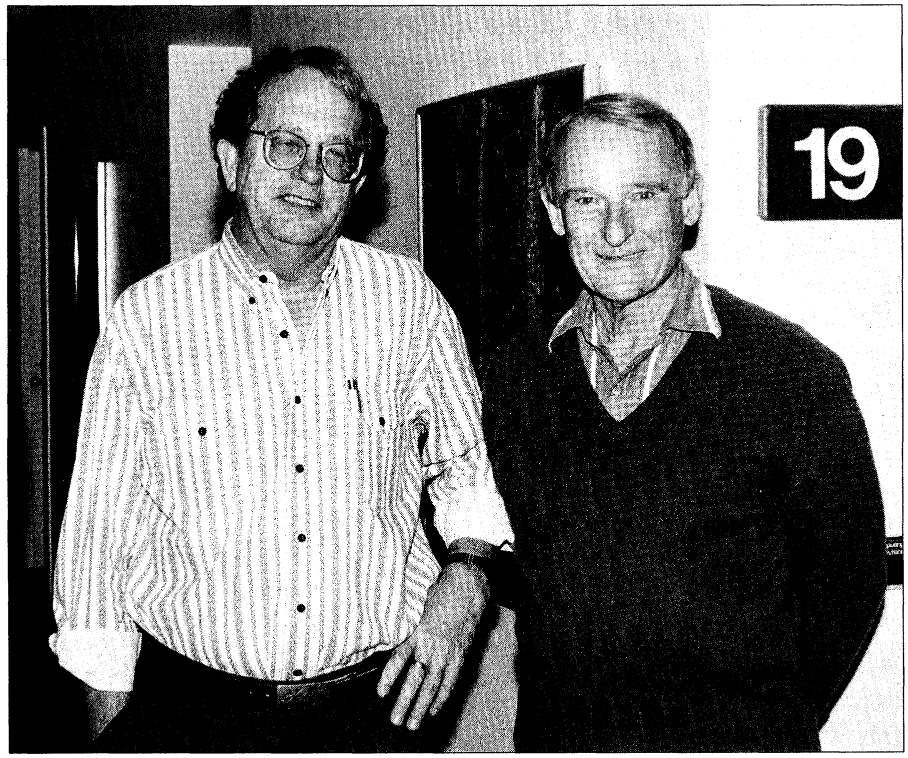

Bill Buzbee, left (SCD director) and Seymour Cray, right (chairman and executive officer of Cray Computer Corporation) discuss the CRAY-3 supercomputer installed at NCAR in May. In recent months, NCAR has also installed a Connection Machine-5 from Thinking Machines Corporation and an IBM SP-1 cluster. (Photo by Bob Bumpas.)

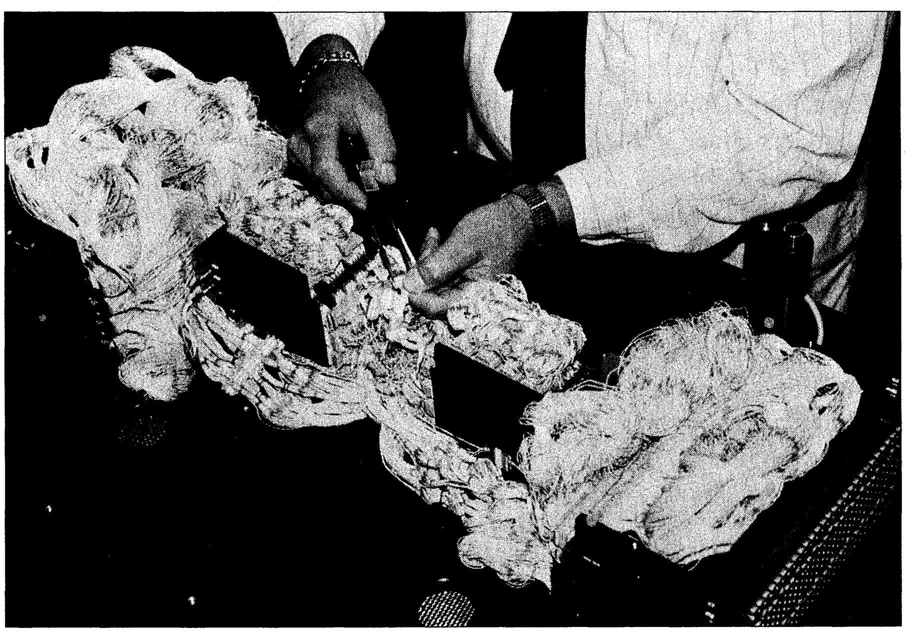

This photo shows the processors and memory in the CRAY-3. When the machine is running, the entire assembly is immersed in liquid coolant. (Photo courtesy Cray Computer Corporation.)

Are cluster farms green?

by Lynda Lester

"Boulder is a self-proclaimed center for political correctness often referred to as the People's Republic of Boulder," said Dan Anderson, who has been called the Hunter S. Thompson of SCD. "So in putting this talk together, I tried to keep with the themes of the area, and titled this talk 'Are cluster farms green?'

"Some of the architectural features of microprocessors, workstations, and cluster farms are quite different than those of the conventional supercomputers you've been dealing with till now. And in spite of the advantages, there are some problems you should be aware of. What niche do clusters have in the diversified computing environment of the nineties? Are they effective cost performers? Even if you're using supercomputers, you may be concerned with these issues in the next few years."

In early 1992, Dan explained, NCAR started an experiment with an IBM cluster farm. Goals were to provide an inexpensive source of interactive, single-processor flops (floating-point operations per second); a source of batch flops during off hours; and a source of multiprocessors for parallel experiments.

The cluster farm consists of five IBM RS/6000 550 workstations networked together—four elements and a gateway that's used for compiling, managing files, and providing access to the elements. The cluster is part of an increasingly distributed computing environment at NCAR that includes two CRAY Y-MPs and a CRAY-3. "We also have an MPP [massively parallel processing] starter kit in the CM-5," Dan said. "In the future you'll see more of these networked heterogeneous compute servers."

Cache flow surplus

The big question is, how well do microprocessors in general perform as compute servers? Microprocessor architecture is different from Cray architecture, which provides relatively random-address, low-latency memory with a fast bus; on a Cray, Dan said, there's no penalty for randomly reaching out for a word. This is not true with microprocessors, where memory chips have a high latency and a slow bus compared with CPU speed. In order to "feed" the CPU at its higher speed, manufacturers interpose a cache between memory and CPU. The idea is to put memory data the CPU needs for calculation into the cache and, hopefully, reuse the data many times. If this is successful, performance can approach that of a supercomputer.

Also, in some microprocessors the memory data and the input/output (I/O) data both flow through the data cache. If you don't have good locality of data—that is, if you haven't organized your data to map efficiently in the cache—you get a lot of movement in and out of the cache; this creates cache turbulence, and flops take a nosedive. Therefore, Dan advised, you want to think about how your data map to the cache. Even if you have a good compiler, the better you plan, the better your code will run.

With cache turbulence, flops take a nosedive.

Why is this important? Some supercomputer manufacturers are already moving toward a cached architecture—Fujitsu has already done so, for instance, and Cray Computer Corporation is thinking about it. Cache turbulence, therefore, is a problem that may occur not only in clusters but in "big iron."

Parallel, no. SPMD, yes.

Meanwhile, how are the parallel experiments coming along? "First of all, we're not actually working with parallel codes," Dan said. "We're doing *very* coarse-grained parallelization—which means we're running

Now available on a cluster near you

During his talk, Dan described the software available on the IBM RS/6000 cluster. A variety of environmental utilities such as shells and editors are supported, he said—"Not ones that UNIX-geek administrators would provide, but those you're familiar with." The cluster also features general utilities such as timers, Mathematica, X Window System Releases 4 and 5, NCAR Graphics, and the GF graphics package.

Math libraries such as the International Mathematical and Statistical Library (IMSL), the Numerical Algorithms Group (NAG) Fortran library, the Linear Algebra Package (LAPACK), and a host of math utilities are installed; this allows users to move codes back and forth between the CRAY Y-MP8/864 (shavano) and the cluster. Number and file conversion are also supported, so users can convert Cray numbers to both single- and double-precision IEEE format.

Access to the Mass Storage System (MSS) is via High-Performance Parallel Interface (HIPPI); this means that **msread** and **mswrite** are available on the cluster elements, and users are reporting good data transfer rates.

"For those folks who really want to torture themselves, we provide parallelization tools," Dan said. The message-passing libraries Parallel Virtual Memory (PVM) and Parasoft Express are both alive and well on the cluster. Linda, a software package that makes a set of disparate machines appear to have shared memory, may become available. One problem with Linda, however, is that it doesn't yet support asynchronous reads and writes.

Fortran 90 with high-performance extensions is becoming the de facto standard for people trying to write conventional parallel codes, Dan noted; accordingly, the NAG Fortran 90 compiler is supported on the cluster.

The cluster is part of an increasingly distributed computing environment at NCAR.

different parts of the 'outer loop' on the different cluster elements. This is called data chunk parallelization, or SPMD [single program, multiple data], and it has three caveats; I call these 'Anderson's axioms.'

"First, the interprocessor communication bandwidth must be great enough that the CPUs don't fall idle. Those of you subjected to the proselytizing of MPP manufacturers have heard talk about high bandwidth and low latency of the communication networks connecting processors. For the computational fluid dynamics codes we've been running, latency is not a problem—the problem is the bandwidth of the communication pipe.

"Second, the communication bandwidth must go up as computation requirements go down—that is, if the flops take less time, so should the communication, or the CPU has to wait.

"Third, the solution for all problems would be the network bus running as fast as the memory bus."

In view of these axioms, can cluster parallelization experiments be generalized to MPP—that is, if you parallelize to the cluster, can you easily port your code to an MPP supercomputer like the CM-5? Yes, Dan said—*if* you have an SPMD model with message passing and the necessary communication rates are attainable on the MPP machine.

Definitely maybe

What conclusions can we draw about clusters as general-purpose compute servers? As a source of interactive flops and flops for appropriately sized serial jobs, Dan said, clusters are cost effective. When clusters are used for parallelization, however, cost effectiveness varies depending on what part of the system we look at. For the price, CPU performance is excellent and memory performance is good; however, communication performance is only fair. In a well-designed, next-generation cluster, cost effectiveness for parallelization may increase with the addition of fast-kernel support for communication and larger caches that will help keep the CPU busy.

Green grow the clusters-O

Clusters fit well into a diverse, heterogeneous, "multicultural" computing environment. They're distributed, decentralized, and democratic. They're small. And, next to a $35 million supercomputer, they're cheap.

So are they green?

Even by Boulder standards.

Lynda Lester is a writer/editor in the Documentation Group within the SCD User Services Section.

Think about how your data map to the cache.

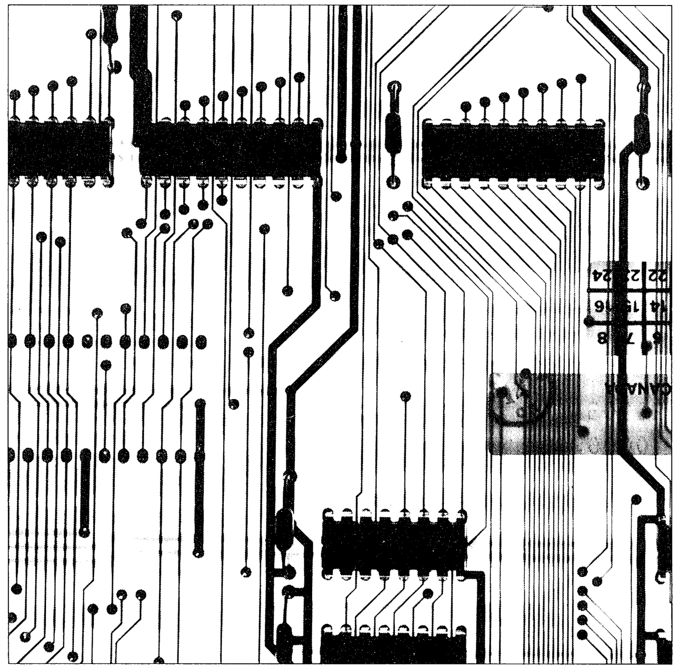

Electronic circuits weave a geometric design. (Photo by Bob Bumpas.)

The MSS: NCAR's 'Big Gulp'

by Juli Rew

In his talk, "Current plans for the Mass Storage System," John Sloan of the SCD High-Performance Systems and Networking Section described SCD efforts to satisfy the computing community's powerful thirst for scientific data. He began by noting that data in active use are kept on the disk farm and in the Storage Technology Corporation (STK) automated cartridge system known as the "Silo." Less frequently accessed data are kept in the IBM 3490 tape archive.

SCD has been trying strategies for caching data active on the Mass Storage System (MSS). Caching involves placing data in high-speed buffers. For small file sizes, we have been able to meet our target for residency times. However, we need at least twice the amount of Silo space in order to be able to cache the large files. John said that a Silo upgrade using double-length, double-density tapes is planned; this will increase the Silo capacity from one to four terabytes (TB) and improve the situation greatly.

Today, the MSS archives 35 TB of data and is approaching nearly a TB per month of new data. (Photo by Bob Bumpas.)

On the tape archive, we will continue to use IBM 3490 tapes for small files, but we may also be able to move to D-2 or D-3 (helical scan) tapes for very large files. The demand already exists for handling file sizes in excess of a gigabyte—users are simply breaking files up to get them onto several cartridges on the MSS.

Today, the MSS archives 35 TB of data and is approaching nearly a terabyte per month of new data. On average, SCD has found that for every gigaflop of computing done on the supercomputers, users generate about 1/2 megabyte of data. John predicts our data storage demands will continue to accelerate as massively parallel systems come into use. He says the only networking technology that presently can handle the required data rates is High-Performance Parallel Interface (HIPPI), with a peak rate of 800 megabits per second. SCD has already set up HIPPI connections to two of the supercomputers in NCAR's network and is planning to install them for most major systems.

Other future technologies on the horizon include:

- Redundant arrays of inexpensive disks (RAID)
- Helical scan media
- The Fibre Channel (an optical fiber communications standard under development)
- Replacement of the existing IBM 3090 Mass Storage Control Processor with UNIX file server front ends
- Andrew File System–like access to the various major components on the network

Juli Rew is head of the Documentation Group within the SCD User Services Section.

NCAR's data archive keeps growing and improving

by Brian Bevirt

A vast repository of meteorological and climate data, NCAR's data archive is being updated and upgraded by SCD's Data Support Section (DSS). Dennis Joseph provided an overview of the data holdings in the archive and highlighted recent additions. He described current DSS projects for improving the archive and making the holdings more easily accessible to users. He also outlined user, media, and data-format trends affecting the archive.

Archive holdings

At present, 3.7 terabytes (TB) of original data are stored in the archive; this includes satellite data but not intermediate or backup products. These data include:

- Station data—Surface and upper-air observations from many sources
- Satellite data—2.5 TB of level 1b data from National Oceanic and Atmospheric Administration (NOAA) polar orbiters and derived data from various satellites
- Analyzed fields—Gridded analyses and forecasts

The DSS FTP facility lets you download datasets, documents, and software.

- Climate model products—Gridded output from various climate models
- Supporting sets—Station catalogs, elevation sets at a variety of resolutions, vegetation and soil information

Recent additions to the archive

9th UserCon

In the last two years, many new gridded and spectral products with higher resolution that use more advanced analysis techniques have been added. Most of these data cover recent times. Significantly more and better ocean datasets are now available from NCAR; some examples of these are:

- Comprehensive Ocean Atmosphere Data Set (COADS)—An ongoing effort to make the best possible set of ocean surface observations from 1850 to 1992
- Global Sea Surface Temperature Analysis—Global gridded fields derived by blending observations from satellites, ships, and buoys, and incorporating sea-ice edge data (previously available monthly on a 2-by-2-degree grid, it is now available weekly on a 1-by-1-degree grid)
- ERS-1 Satellite Collected Scatterometer Data—Global ocean surface wind speed and direction

Current DSS projects

Current DSS projects include a major update to the COADS ship data archive; more extensive data distribution on CD-ROMs, helping to make near-real-time data available to NCAR users; and major input to the reanalysis project. The reanalysis project will improve the reliability and consistency of analyzed data because a single analysis technique is being used, and because previously unavailable raw data are being included. This project should eventually provide data from 1958 to the present. Various analysis techniques have been used in the past, and at the times when these techniques changed, the derived data products may show abrupt, artificial discontinuities. The reanalysis project will eliminate these artifacts and provide more complete and consistent analysis products for users.

Data access enhancements

To improve the accessibility of archive holdings for users, DSS now operates an anonymous File Transfer Protocol (FTP) facility. This facility allows users to browse the archive and download datasets, documentation, and software. Various ways to view the archive holdings are provided by a set of catalogs. Users can access inventories in the dataset directories to determine if a particular dataset contains the data they need.

Current trends

Trends affecting the archive keep DSS staff busy. Archive users are changing: the number of users is growing, there is a marked increase in nontechnical users who need more help, users employ a wider variety of computers, and they typically have less experience with data transfers than past users. The media used for data transfers are also changing rapidly. About five years ago, most data were shipped from NCAR on half-inch magnetic tapes. With the increase in popularity of IBM 3480 cartridges, floppy diskettes, and eight-millimeter videotape cassettes, less than 50% of current data requests now use half-inch tape. Improvements in networks and increases in their carrying capacities are making FTP and e-mail the fastest growing method for obtaining data from the archive.

To access the archive via anonymous FTP, use the address ncardata.ucar.edu (128.117.8.111). You can send e-mail to datahelp@ncar.ucar.edu to request data, or you can phone (303) 497-1219. You can also make fax requests using the number (303) 497-1137.

Brian Bevirt is a writer/editor in the Documentation Group within the SCD User Services Section.

Good news: User interfaces to become easier

by Juli Rew

Don Middleton-Link of the SCD Scientific Visualization Group reported on proposals by the Distributed Computing Environment (DCE) Committee on how to facilitate computing at NCAR by making the user interfaces more consistent, informative, and easier to learn. Don presented interfaces for three major SCD services: the Mass Storage System (MSS), printing, and batch job submittal.

Facing up to interface problems

Don outlined problems with the current computing environment, including:

- Too many interfaces
- Inadequate error reporting
- Interfaces that do not follow industry or de facto standards
- Interfaces that are difficult to use
- File size limitations

These problems have developed over the years, as SCD gradually developed a home-grown, high-speed network called the Mainframe and Server Network (MASnet). MASnet software was written locally. The number of interfaces has proliferated, notable examples being the MASnet/Internet Gateway Server (MIGS) and the Internet Remote Job Entry (IRJE) system. Now, however, the MASnet hardware is rapidly becoming obsolete, and standard interfaces are more readily achievable.

Guidelines for a new user interface

In developing a new user interface, Don said the committee would follow certain guidelines, such as:

- Making all services available on all systems
- Making interfaces the same for all systems
- Making services behave the same on all systems (synchronous systems are most desirable)
- Making services transparent to the user
- Basing interfaces on either de facto or industry standards, for example, POSIX
- Making naming conventions consistent

Solutions for the Mass Storage System

Don said that building a UNIX file system on top of the MSS would relieve many problems related to computing on this service, but that this is a large, difficult project. The DCE Committee is thus proposing a three-phase plan to reach that goal:

- Phase I: Treat the MSS as a non-UNIX archive. Programmatic, command line, and graphical user interfaces (GUIs) that parallel existing UNIX commands would be provided. Existing hardware would be used. This phase is targeted for completion in about a year.
- Phase II: Provide a restricted UNIX file server whereby only metadata functions would be supported (for example, **ls** for listing files).
- Phase III: Provide full, unrestricted UNIX file server access whereby the archive is functionally transparent to the user.

Import and export functions for the MSS are also in progress.

SCD plans to make user interfaces more consistent and easier to learn.

And more . . .

In addition to the MSS interface, more standard interfaces for printing and batch job submission are being planned. A UNIX-like interface is planned for printing that would have commands specific to the type of hardcopy output desired, including paper, film, slides, fiche, and videotape. A preview capability is also envisioned. An interface similar to Network Queueing System (NQS) is being contemplated for batch job submission to all machines, similar to what is presently running on the Cray computers.

You may obtain a copy of the proposed plan from Laura Morreale (e-mail: laura2@ncar.ucar.edu).

Juli Rew is head of the Documentation Group within the SCD User Services Section.

Quick takes: A photo gallery of conference highlights

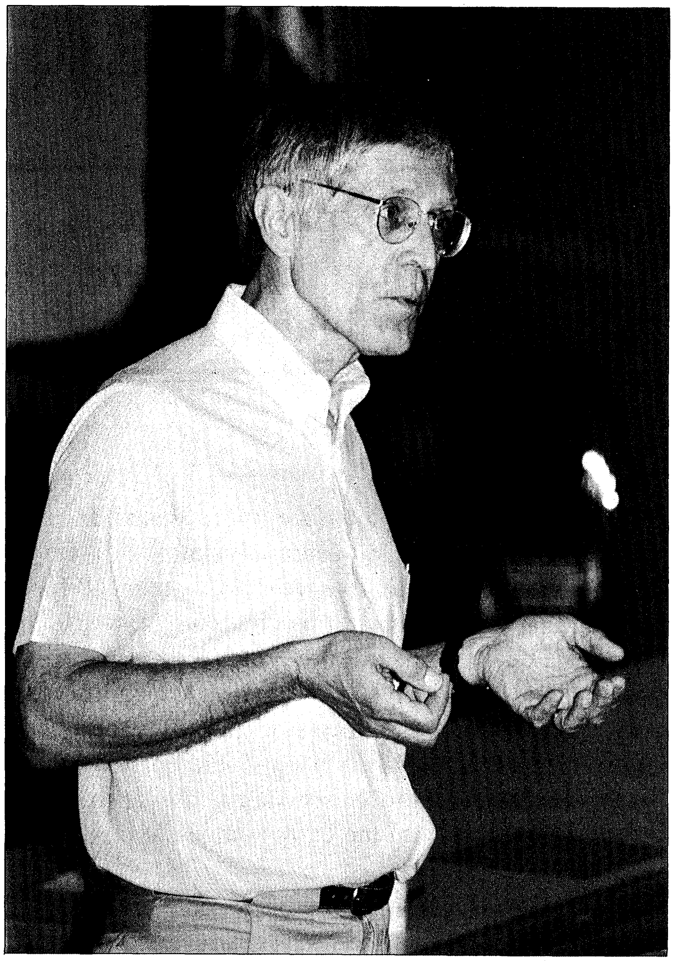

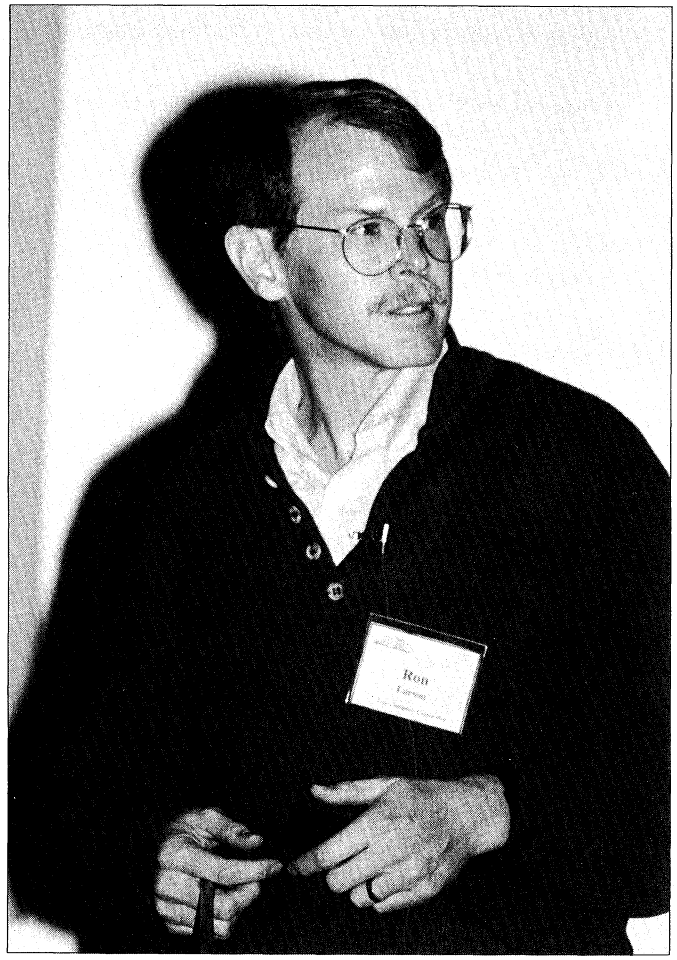

SCD's Don Morris, upper left, spoke on NSFNET plans; Ron Larson, upper right, of Cray Computer Corporation, showed slides on the manufacture of a CRAY-3. Lower right, SCD's Jeff Kuehn explains a point of interest to Shengbai Zhang of the National Renewable Energy Laboratory during a tour of the NCAR Computer Room. (Photos by Curt Zukosky.)

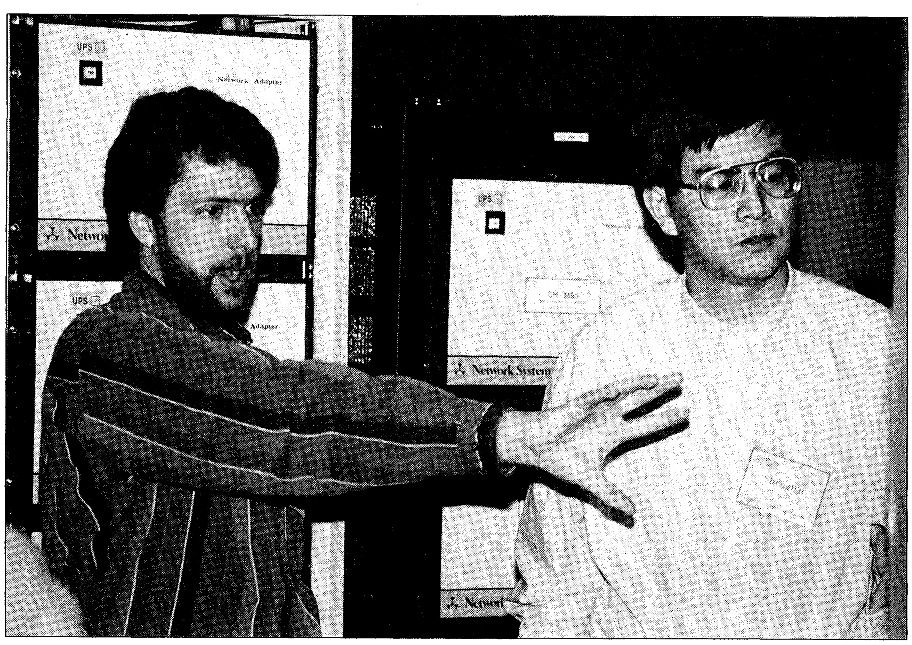

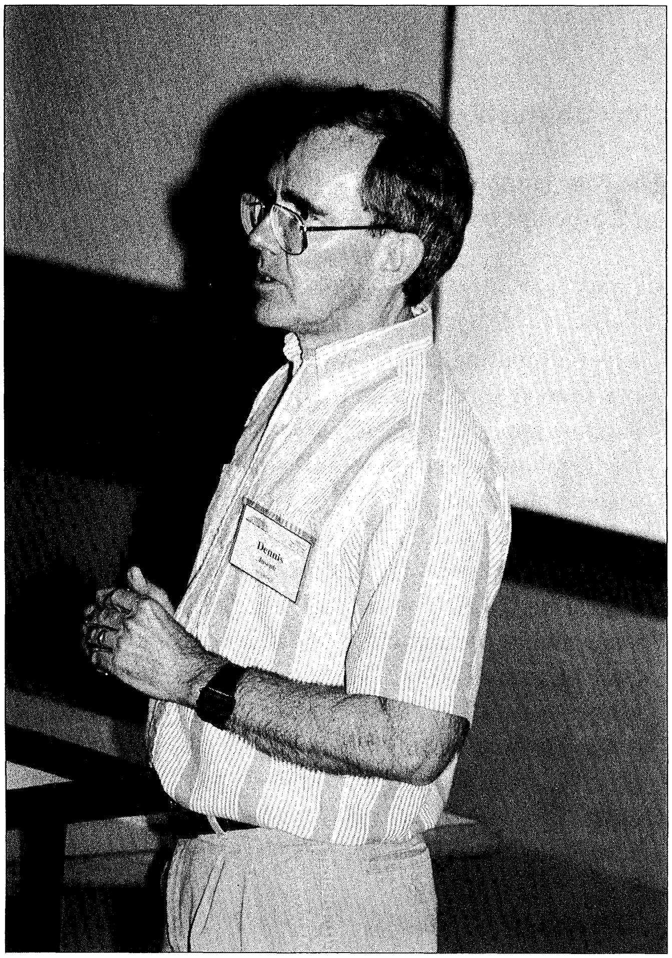

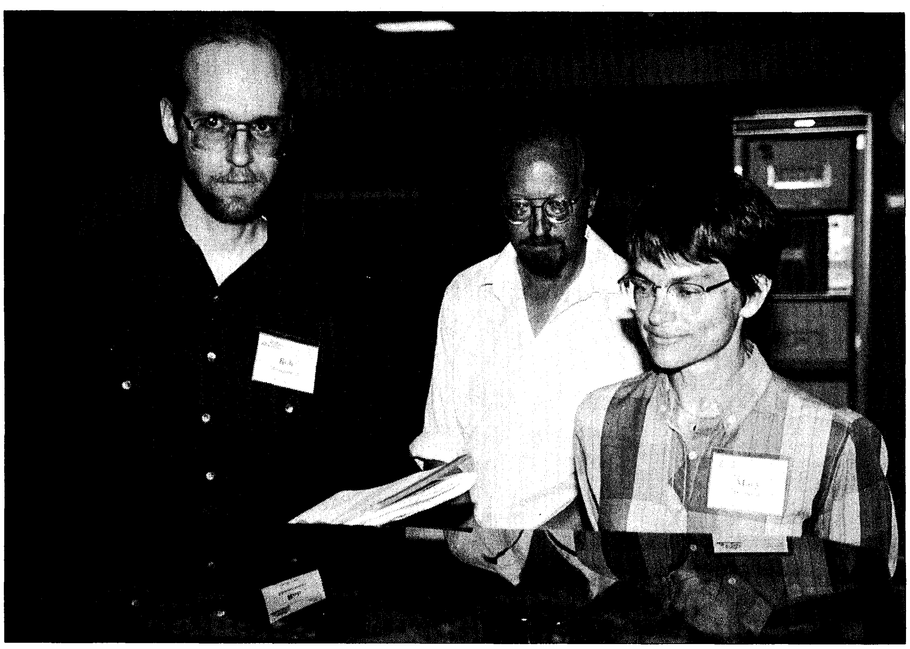

SCD's Dennis Joseph, upper left, reviewed the data holdings in the Data Support Section archive; SCD's Dan Anderson, upper right, talked about the "greenness" of clusters. Lower left, Rob Montgomery (NCAR's High Altitude Observatory) and Mary Downton (NCAR's Environmental and Societal Impacts Group) check out the new CRAY-3, while a visiting scientist looks on from behind. (Photos by Curt Zukosky.)

NSFNET: Big changes on the horizon

by Juli Rew

Don Morris of SCD's Networking and Data Communications Group spoke on "National network futures." Don focused primarily on plans for the National Science Foundation Network (NSFNET). NSF has issued a solicitation for a new network architecture upon expiration of the current cooperative agreement with Advanced Network Services, which currently operates the NSFNET backbone.

Since it was established in 1986, NSFNET use has grown at a high rate, paralleling the growth of the worldwide collection of interconnected networks called the Internet. The new solicitation recognizes this growth, along with advances in networking technology in the commercial sector, and appears to encourage more participation by commercial vendors and communications carriers in providing network service. It calls for a "transition to a networking infrastructure that is increasingly provided by interconnected network service providers operating in a competitive environment."

A high-speed backbone would connect NSF supercomputing centers with at least 155 Mbps service.

Don reported that NSF's solicitation has raised fears in some circles that scientists will no longer have the free access to the NSFNET they presently enjoy. Don allayed these fears somewhat, saying that NSF may actually be spending more to get the new infrastructure in place than it is spending now, and that for the next four years regional networks will continue to be subsidized for connections to NSFNET.

New network infrastructure

The four components of the new NSFNET infrastructure are:

- *vBNS (very-high-speed Backbone Network Service)*—A backbone connecting five NSF supercomputing centers with at least 155 megabits per second service. The high bandwidth would support applications such as distributed high-performance computing and isochronous visualization. NCAR will be on the vBNS and have a connection to a NAP (defined below). Traffic on the vBNS must be in support of research and education (R&E), consistent with NSF's current "acceptable use policy" (AUP). Excess bandwidth may be used for nonresearch and educational purposes, but the vBNS provider must be able to distinguish between NSF customer traffic and that of others. Technology to accomplish this does not yet exist, Don noted. "We can't color some packets red and some blue," he said.

- *NAPs (Network Access Points)*—Major spots where the vBNS, commercial network service providers (for example, Sprint and MCI), and other appropriate networks may interconnect. The solicitation is very broad in what may be proposed as a NAP, which could be anything from an Asynchronous Transfer Mode (ATM) switch to a Fiber Distributed Data Interface (FDDI) ring or a metropolitan area network. Three priority NAP locations are specified, with five additional desirable locations. It appears that NSF wants one company to be the overall manager of the NAPs.

- *RA (Routing Arbiter)*—An organization that will establish and maintain databases and routing services that may be used by attached networks to obtain routing information (such as network topology, policy, and interconnection information)

Figure 1.

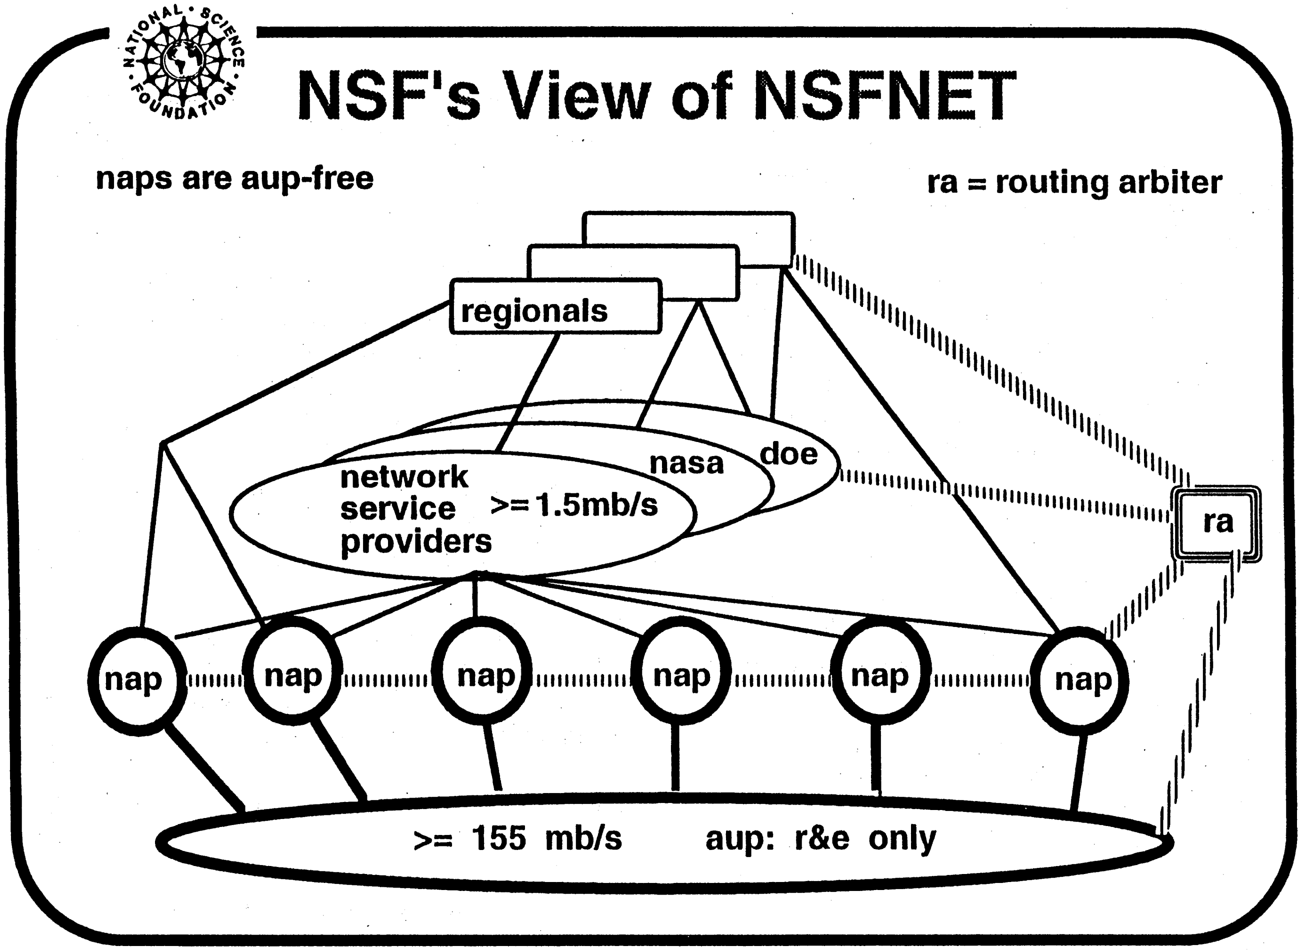

with which to construct routing tables. The object of the RA is to promote routing stability and manageability and advance routing technology.

- *Regional networks*—NSF proposes to pay only for connectivity to the NAPs, or to network service providers that connect to NAPs. Under this scheme, regionals may be required to make their own arrangements with network service providers to obtain full interregional connectivity. After four years, the regional networks will be asked to become self-supporting.

Adding Internet addresses

Participants asked about the future of the Internet Protocol (IP) addressing system. Don observed that the number of available IP numbers is running out. One proposal is to use Open System Interconnection (OSI) addresses (which are 48 bits instead of 32 bits for IP); this would allow for needed expansion. The IP protocol would continue to be used.

Juli Rew is head of the Documentation Group within the SCD User Services Section.

New and improved: NCAR Graphics Version 3.2

by Brian Bevirt

Bob Lackman, head of the Scientific Visualization Group, announced many improvements in the latest release of NCAR's scientific visualization software, Version 3.2 for UNIX computers. (Generic and VMS ports are not available for Version 3.2.) This article presents detailed information from Bob's talk; for a quick summary, see the accompanying sidebar ("At a Glance: New features in NCAR Graphics 3.2").

Utility upgrades

Bob highlighted new functionality in the Version 3.2 utilities listed below. Other utilities have also been upgraded; many improvements in usability and output quality have been made.

Conpack *(contour drawing).* Conpack can now produce raster contouring using cell array output in the Computer Graphics Metafile (CGM). This routine has been integrated with all of Conpack's other features, such as transformations, map overlays, grids, labels, labelbars, legends, and so on. Hachuring, or placing tick marks on contours to indicate slope, can now be done by the utility.

Ezmap *(map drawing).* Inverse map transformations are now available for all ten map projections. A great-circle interpolator is provided.

Isosurface *(surface visualization).* Isosurface is now capable of interpolating a surface to a higher resolution. This utility can now be controlled by parameters that offer increased functionality.

Streamlines and Vectors *(field flow visualization).* Streamlines and Vectors now contain built-in mappings for all ten Ezmap map projections and for use with polar coordinate systems. Both utilities use an improved process for overlays with other utilities and can now be controlled by parameters that offer increased functionality. The Vectors utility now has the ability to color vectors according to vector magnitude or a scalar field.

Plotchar *(text drawing).* Eight commonly used fonts are now available in a high-quality color-filled form, and the bold equivalents for three of these are included. Two of the new filled fonts have not been available through Plotchar before; they produce all the standard meteorological weather symbols. You can now control character color, line width, shadows, and outlines with parameters. You can also transform characters so they appear to wrap around any surface.

NCAR GKS-0A *(NCAR's subset of Graphical Kernel System drawing and workstation-control routines).* The new GKS-0A library allows you to select one of several output types. Options include:

- An NCAR CGM (the only output type available in previous versions)
- American National Standards Institute (ANSI)–standard CGMs
- X11 windows driven directly by your program (avoiding the intermediate step of producing an NCAR CGM and viewing it with **ctrans**). NGPICT is a new utility for controlling multiple, simultaneously active X11 windows.

The idt viewer now supports workstation animation.

New applications for raster files

Four new tools allow NCAR Graphics users to manipulate raster files in various formats:

- The **rascat** filter concatenates raster images, converts between raster image formats, resamples images, and performs color quantization.

- The **rasview** application is a raster file previewer.

- The **rasls** tool lists basic information about raster files and image-specific information like resolution, encoding, and frame count.

- The **rassplit** tool splits multiframe raster files into single-frame raster files.

The supported formats include Sun; X11; NCAR Raster Image Format (NRIF); Advanced Visual Systems (AVS); Hierarchical Data Format (HDF); Abekas; and Silicon Graphics, Incorporated (SGI).

ANSI-standard C bindings for all user entry points

All user entry points in NCAR Graphics and in the supported subset of GKS routines can now be called directly from a C program. All GKS C bindings comply with ANSI standards.

Improved CGM viewing and output tools

The **idt** viewer now allows users to zoom in repeatedly on an image. It also supports desktop animation of NCAR CGMs. Its point-and-click interface has been streamlined; there are now only two interactive windows. The **ctrans** and **ictrans** tools now offer more capabilities. Users can output raster files, simulate background colors for PostScript devices, redirect output for graphcap and raster devices, and specify a color palette file to override any color table specified by the metafile. Further, **ctrans** and **ictrans** will support the Hewlett-Packard LaserJet PCL language and many new raster output devices: Sun, xwd, NRIF, AVS, HDF, Abekas, and SGI.

Many new examples

The software developers have updated the existing set of examples and added examples that demonstrate the new functionality in Version 3.2. Also, more examples have been generated for the new user manuals. All examples can be accessed via the existing **ncargex** command, which was upgraded with more options that provide greater user control. Examples showing the use of the new C bindings are also available; these can be accessed via the **ncargcex** command. Users can now see a complete list of the examples by typing the UNIX commands **man ncargex** and **man ncargcex**.

At a glance: New features in NCAR Graphics 3.2

- Many utilities provide new functionality and contain improvements in usability and output quality.

- New applications work with raster files.

- C language bindings now exist for all user entry points in the package.

- CGM viewing and output tools have been improved.

- Many new Fortran examples have been included, and C examples are also provided.

- Full binaries are being distributed for the first time.

- A new distribution medium is now available: CD-ROM.

- The installation script has improved, and it is more flexible than past versions.

- New installation tools help customers smoothly upgrade to new versions on their local systems.

- Directories of prototype and unsupported software are available.

- User documentation has been completely redesigned.

Distribution and installation enhancements

Full binaries, in addition to the full source code, are now available for all supported systems running their current operating systems: CRAY Y-MP, DECstation, IBM RS/6000, HP 9000/Series 7*xx*, SGI, Sun-3, Sun-4, and Sun SPARCstation. The full source distribution contains almost all the available documentation in PostScript form. Distribution is now available on the following media: CD-ROM, 8-millimeter tape, QIC-150 1/4-inch cartridge, and 1/2-inch tape at 6250 bits per inch. With the release of full binaries, the installation script was modified to allow customers to install only needed parts of the release. Finally, new environment variables are provided so customers can upgrade smoothly to newer versions of NCAR Graphics, change the default name and/or individual names of output metafiles, and control the redirection of output.

C language bindings exist for all user entry points in the package.

New directories of prototype and unsupported software

A prototype version of NCAR Graphics' next-generation High Level Utilities (HLU) library is included with the Version 3.2 release. This prototype software contains a small subset of the functionality planned for the Version 4.0 release, and it will run only on Sun systems. The new **nhlex** command runs HLU examples. The HLU library provides a simplified interface to the package, and it will underlie the more advanced interfaces of the upcoming NCAR Interactive release. Users can experiment with this prototype to become familiar with this next-generation interface, and any user of the prototype can influence the future direction of NCAR Graphics' interfaces by providing comments and feedback to the developers. The goal of the NCAR Interactive project is to make NCAR Graphics functionality available to a broader range of users by simplifying the programmatic interface and offering an advanced graphical user interface.

Various examples showing how NCAR Graphics can be used in a data flow environment are also provided as unsupported software. A data flow environment is a system with an advanced graphical user interface where programs are built by connecting modules in a pipeline. The supplied examples are all done with the Explorer package from SGI; this package is included with SGI workstations, so these examples will only run for customers with SGI equipment. Others can study these examples to see how NCAR Graphics enhances the functionality offered by other vendors' data flow environments.

There is also a source of unsupported NCAR Graphics programs, documentation, and data files available in the **ncarg/softshare** directory via anonymous File Transfer Protocol (FTP) from the computer named ftp.ucar.edu. This service was started by the NCAR Graphics User Group and is administered by a volunteer from the user community.

An e-mail group has been set up for NCAR Graphics users who want to discuss problems and solutions with other users. To subscribe or unsubscribe to this group, send e-mail to ncarg-talk-request@ncar.ucar.edu. Once you are on the list, you can send e-mail to the entire group by sending e-mail to ncarg-talk@ncar.ucar.edu.

Conpack can produce raster contouring using cell array output in the CGM.

All-new user documentation set

The documentation set has been completely revised for Version 3.2. A training guide for the heart of the package is now ready: the *NCAR Graphics Contouring and Mapping Tutorial* is a step-by-step guide to using the Areas, Ezmap, and Conpack utilities. Along with the examples that accompany it, this manual helps new users become proficient at using and combining advanced utilities. The *NCAR Graphics Fundamentals* guide, an introduction to NCAR Graphics programming, provides an introduction to all the other utilities and applications in the package. Both of these new manuals are designed to get you up and running with NCAR Graphics as fast as possible.

New **man** pages are now provided for all user entry points in the package: all routines, user-modifiable routines, and parameters in all utilities. Each utility-overview **man** page provides the name and description of each routine in that utility, and the **ncargintro man** page provides an overview of all utilities.

Another new type of document is introduced with Version 3.2: the programmer documents. These PostScript files on the distribution are the software developers' reference notes; they are provided for users who want to explore the information that developers use to maintain a utility's source code. Programmer documents assume that readers have a strong background in NCAR Graphics programming.

Obtaining NCAR Graphics Version 3.2

NCAR Graphics Version 3.2 is designed for installation on UNIX-based computer systems. NCAR Graphics is sold to both nonprofit and for-profit groups. For details about ordering, licensing, and pricing, contact:

NCAR Graphics Information
NCAR Scientific Computing Division
P.O. Box 3000, Boulder, CO 80307-3000
Phone: (303) 497-1201
E-mail: scdinfo@ncar.ucar.edu

Brian Bevirt is a writer/editor in the Documentation Group within the SCD User Services Section.

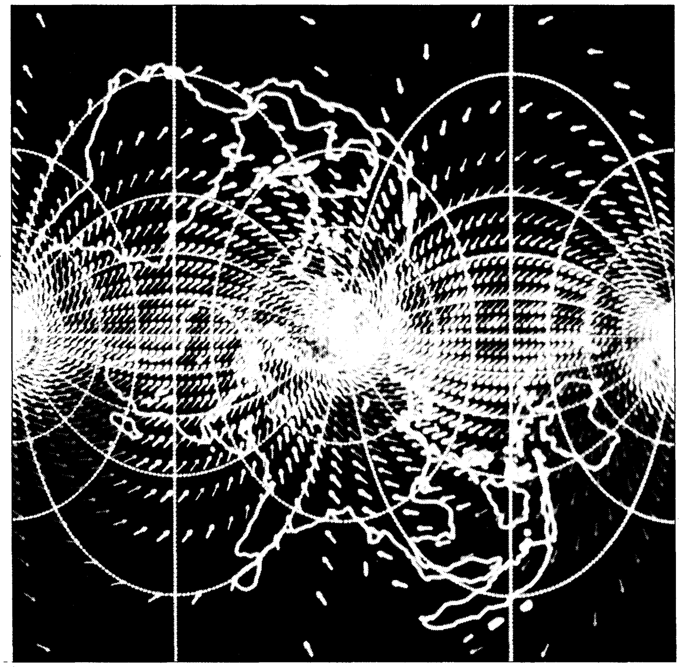

Field flow utilities that generate wind vectors and streamlines have been upgraded for Version 3.2. Both vectors and streamlines can be transformed through all ten of the standard NCAR Graphics mapping transformations. This graphic uses dummy data to show a set of wind vectors plotted as a transverse Mercator projection.

SCD video facilities launch a new era in scientific visualization

by Brian Bevirt

Don Middleton-Link and Tim Scheitlin, both of SCD's Scientific Visualization Group, described SCD's efforts to help researchers create video presentations for scientific purposes. In their talk, Don and Tim noted the expanding roles of video in scientific animation, outlined the video facility goals, and described the video printers. They also provided a brief tutorial on producing animations with the video printers, showed SCD's new videotape that demonstrates how to produce good video animations, and offered a glimpse of the enhancements planned for SCD's video labs.

Videotapes are becoming common in scientific presentations at conferences and symposia.

New roles for video

Don opened the presentation with an overview of the new roles being filled by animation in scientific communication. The first new role is publication. Don showed the July 10, 1992 edition of *The Astrophysical Journal;* many of the graphical images for this issue were provided on a videotape that was included with the magazine. This use of video in scientific publications is expected to increase. Videotapes are also becoming a common medium in scientific presentations at conferences and symposia, and are a convenient medium for peer communication—especially when collaborators are working at different sites. Because patterns and trends in data are sometimes easily spotted simply by animating a sequence of images, video is an excellent tool for data analysis and exploration. The demand is also growing for scientific animations used in education at all levels. Finally, videos are making scientific information accessible to the general public through broadcast media such as television news programs.

A caveat

Video presentations have quality limitations—especially those on VHS, the most common tape format. Users can readily notice differences in resolution and color fidelity between their videos and their workstation screens. Don cautions researchers not to expect results that rival the effects produced by major movie studios, but he also notes that video output is still a useful and effective analysis tool. With the production convenience and reasonable cost offered by SCD's video facilities, this visualization method adds a new dimension to research.

SCD video facility goals

There are two parts to SCD's video facilities: the automated "video printers" now integrated into the Text and Graphics System (TAGS) and the two interactive video labs (one at the Mesa Lab and one at the Foothills Lab).

The goals of SCD's video facilities are:

- To further scientific research through visual means
- To make it easy to produce the video equivalent of a scientific journal article

TAGS video printers

Don avoided describing all the hardware, software, and interconnections in the video printers; he encourages researchers to think of video animation as a new TAGS output option. The video printers are not tied to any one visualization environment; they are automated production services that accept many image formats. SCD's video printers are also growing in popularity: at one peak, 70 jobs were completed in one day.

Think of video animation as a new TAGS output option.

While taking advantage of automated processes, the video labs have not become depersonalized: Don and Tim make time to help new users get started, consult on tricky animation problems, and provide support for educational projects. With John Clyne, also of SCD, they wrote a new SCD user document, "A Guide to the Production of Computer-Generated Video Animations Using TAGS." This document supplies all the information new users need to make a video animation, provides specific procedures for avoiding common problems, and offers tips for producing optimal results within the limitations of the medium. The video titled "A Guide for Producing Computer-Generated Video Animations at NCAR" demonstrates how to apply these tips and improve the quality of scientific animations.

A video animation tutorial

Tim presented the important steps for using SCD's new "video printer" in a brief tutorial format. He offered a simplified overview of the process: create a sequence of raster or page-description images with any visualization tool, store the images in one or more files on SCD's Mass Storage System or another system connected to TAGS, then send the file(s) to the video printer via TAGS.

Tim listed the numerous tools provided by NCAR Graphics for sizing images, combining them into a single file, splitting them into different files, and converting between image formats. Three interfaces are available to submit the image file to TAGS: Mainframe and Server Network (MASnet), MASnet/Internet Gateway Server (MIGS), and Internet Remote Job Entry (IRJE). Tim described the advantages and disadvantages of the four supported video formats: VHS, SVHS, U-matic SP, and Betacam SP. He then detailed numerous methods for using TAGS options to optimize the finished product. For example, users must plan for the video play rate of 30 frames per second; without using one of the appropriate options, a 90-image file would only run for three seconds.

Tim also encourages users to take advantage of the full amount of run time available on each tape by repeating the animation sequence or printing multiple animations on a single tape. The Scientific Visualization Group provides tools that help users make the calculations necessary for proper timing.

Supplemental services are offered by the SCD Film Room. They can recycle old videotapes, make high-quality copies of tapes, and duplicate an NTSC-standard tape to produce a PAL or SECAM standard tape for playback in other countries.

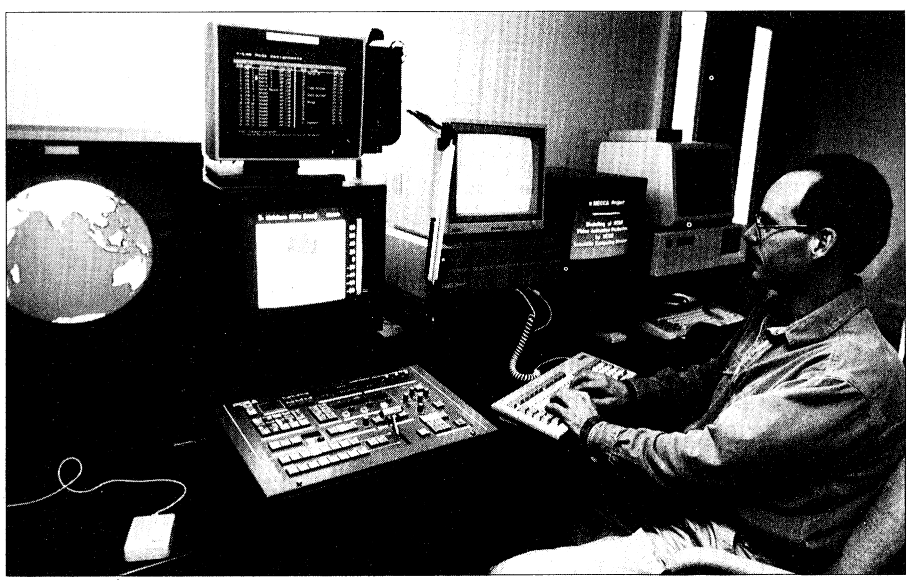

Tim Scheitlin of the Scientific Visualization Group demonstrates animation capabilities in one of the interactive video labs. (Photo by Carlye Calvin.)

One of SCD's goals is to make it easy to produce the video equivalent of a scientific journal article.

Video lab enhancements

Some enhancements that Don and Tim are planning for the video labs include:

- An Animation Scripting Language (ASL) that offers a shortcut to the lengthy and complicated TAGS scripts that may be needed to produce longer or more refined videos
- Digital image filtering algorithms for improving video quality
- Industry-standard tape headers
- The ability to handle larger files on TAGS
- Multimedia publications
- Advanced television (ATV) and high-definition television (HDTV) output
- Larger video lab facilities

Videotape and documentation guides

Demands are increasing on the video labs and on Don and Tim's time. Prospective users are encouraged to begin learning about this new visualization option by ordering the videotape "A Guide for Producing Computer-Generated Video Animations at NCAR," the user document "A Guide to the Production of Computer-Generated Video Animations Using TAGS," and the appropriate user document for the TAGS interface to be used. (MASnet users need "Using the Text and Graphics System from UNIX and UNICOS Computers," MIGS users need "Using the Text and Graphics System via the MASnet/Internet Gateway Server," and IRJE users need "Using the Text and Graphics System from UNIX and UNICOS Computers" and "IRJE: Using the NCAR Internet Remote Job Entry System.") To order the videotape and any SCD user documents, see "SCD Documentation Orders" on the back cover of this issue.

Brian Bevirt is a writer/editor in the Documentation Group within the SCD User Services Section.

Why'd you do it? Users cite rationale for attending SCD conference

by Lynda Lester

From doetzl to doty, cheng to jing, swerdlin to swindle and nanda and zhang, they came to the Ninth SCD User Conference at NCAR—users with logins on computers all over the world. Multiple were their reasons for attending, which they shared when cornered in the lobby.

Steve Baum, *Applied Research Corporation, College Station, Texas:* "I do climate and oceanographic modeling. I use shavano, I use the Data Support Section's FTP archive; also the MSS. I use NCAR Graphics daily, hourly—every five minutes practically. But it was the CM-5 thing that pushed me over the edge to come. I want to get into MPP—I see that as the way to go, because you run into a speed-of-light limit with single processors.

We flog NCAR Graphics to the limit.

"To solve the problems we've set ourselves to solve—global warming and such—you need a coupled model; you can't look at just the atmosphere or the ocean and get answers that will correspond to what we see in the real world. And for that, you need high resolution, which translates to more compute cycles. So you must build models to run efficiently on MPP machines. The ideal situation would be to have a processor for every degree of freedom in the model—to have a processor for each finite-difference grid point, for instance. Of course, there'd be computational overhead from the message passing, but that's really the way to go. We need to change the algorithms we're currently using to solve problems. Like what the guy was saying this morning [Richard Loft in his talk, 'Introduction to the CM-5']—the algorithms used on single-processor machines usually aren't the best choice for MPP machines."

Real users don't expand acronyms

When real users speak, they use jargon and acronyms—which they *don't* spell out in parentheses during conversation. The following list expands the shorthand terms used in this article.

CM-5	The Connection Machine-5
FTP	File Transfer Protocol
IRJE	Internet Remote Job Entry
LTER	Long-term ecological research
MASnet	Mainframe and Server Network
MIGS	MASnet/Internet Gateway Server
MPP	Massively parallel processing
MSS	Mass Storage System
NSFNET	National Science Foundation Network
PCs	Personal computers
PI	Principal investigator
T3	A network service that transmits a signal at 45 megabits per second
shavano	CRAY Y-MP8/864
Y-MP	CRAY Y-MP8/864

Philip Politowicz, *Department of Atmospheric and Ocean Sciences, University of Wisconsin—Madison:* "At Wisconsin, we're sorting out the chemical effects of high-speed aircraft on the ozone. We're trying to take a 2-D model and make it do microphysics—it's a way of simulating the sulfite aerosol layer in the stratosphere. The simulation has to do with

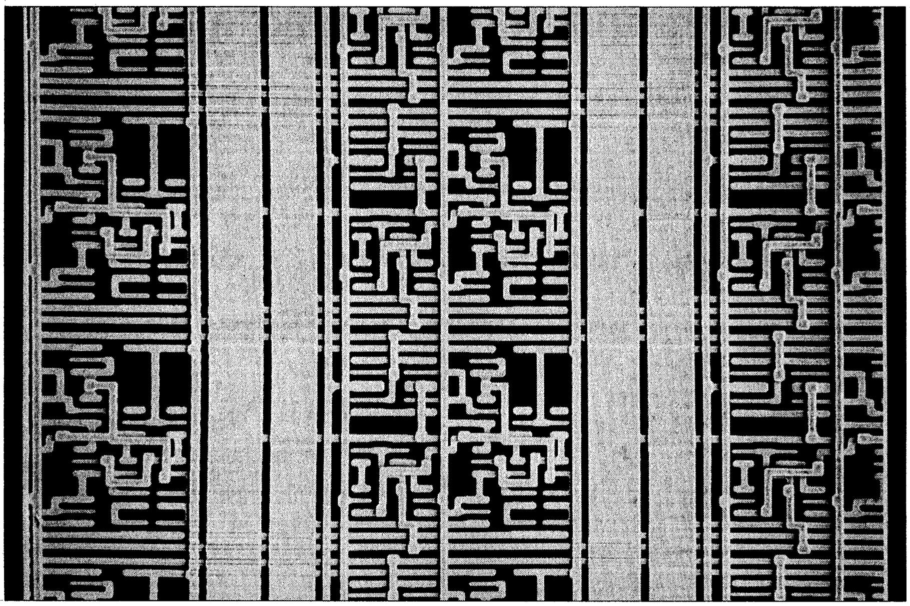

These robotic earthworm trails are really a CRAY-3 circuit viewed through an electron microscope. (Photo courtesy Cray Computer Corporation.)

quasibiennial oscillation; it's a theoretical model to help us determine the behavior of the middle atmosphere. So we do a lot of things with aerosol distribution, and that data is put on the MSS. We flog NCAR Graphics 3.1 to the limit.

30 terabytes of data—unimaginable, really.

"The new graphics package coming out may be very important for what we're doing, so we want to find out how it works. Personally, I'm interested in this video stuff. We generate a lot of pictures in sequence, and video might be a good way to make our data more perceivable—of lowering our level of ignorance, if you want to put it that way."

David Newman, *Astrophysical Planetary and Atmospheric Sciences Department, University of Colorado, Boulder:* "I use the computers at NCAR quite a bit. Basically I wanted to know what will be happening in the next year—the changes, the new machines. I'm working in space physics—specifically, computational plasma physics applied to the ionosphere, magnetosphere, and solar wind. I'm especially interested in some of the new things coming up in NCAR Graphics, and the CRAY-3 and CM-5 seem potentially useful."

Kimberly Schaudt, *Department of Atmospheric Physics, University of Arizona, Tucson:* "My boss [well-known climate modeler Robert Dickinson] wanted me to learn more about what's going on with the computers, especially the new CM-5. I'm doing work with radiative transfer in rainforest canopies, trying to find if the signatures of stress in leaves that we identify in the lab will be visible via satellite. I may be using the CM-5—or maybe telling everyone else in my department how to use it."

Roy Wessel, *W. Roy Wessel & Associates, Boulder:* "I do consulting for scientific modeling on high-performance computers. I worked at NCAR in the late 1970s, and at Control Data Corporation for eight years after that, using vector supercomputers.

"I thought the CRAY-3 talk was quite interesting, as was the whole discussion of traditional computer architecture versus MPP. The physical laws that limit the performance of computer hardware indicate that Bill Buzbee is right: in the future, supercomputing will have to be highly parallel. I would go farther than his statement that supercomputing and MPP will be synonymous, because I think that in not too many years we'll see real parallel processing even in high-end PCs. By the end of the decade, you'll go into a store and buy a PC for a month's pay that will have a good deal of parallel processing in it—user transparent, of course—and the machine as a whole will be roughly equivalent to a CRAY-1."

Paulo Nobre, *Brazilean Numerical Weather Prediction Center, Instituto Nacional de Pesquisas Espaciais/Centro de Previsao de Tempo e Estudos*

Climaticos (INPE/CPTEC): "The information I got here helped break the myth of a cheap solution for supercomputing. The general idea is that clusters are cheaper: put a machine together, it costs a few bucks and you've the equivalent of a Cray—but we're not there yet. Supercomputing still has its place. The cluster also has its place, I'm happy to see. Each architecture has its place, and what I learned most here was what hardware has a value for us in our arena of work. I was also interested to see that SCD is making the user interface less complex and more standard—that will be a good thing for us."

The information I got here helped break the myth of a cheap solution for supercomputing.

Jordon Hastings, *Desert Research Institute, Reno, Nevada:* "The NSF has recently awarded us a new LTER site in the McMurdo dry valleys, Antarctica. Robert A. Wharton, Jr. is the PI; I am assisting him with data and information management for the project. In addition, there are seven co-PIs scattered from Bozeman, Montana to Tuscaloosa, Alabama; four are located on the Colorado Front Range. Our field studies involve hydrogeology and biogeochemistry in an extremely cold desert environment. We are also developing numerical models of the principal glaciers, streams, and lakes. Climate change is a potential issue.

"I have some familiarity with NCAR from times past: I was employed here in the 1970s, and have consulted back here occasionally since. Because of the co-PIs in Colorado, particularly, I am interested in exploring the computing resources at NCAR for the McMurdo LTER. The Data Support Section already has a large collection of Antarctic data, and we might like to use the MSS as a central repository for our own data over the lifetime of the project (five-plus years). The fact that the MSS is now being connected to the main T3 Internet backbone is very significant. We also have a potential interest in the Cray and IBM machines for our modeling work."

Jonathan Streeter, *U.S. Department of Commerce/ International Trade Administration:* "I'm an economist. I work for the International Trade Administration, and our job is to promote U.S. industrial competitiveness. Some of it has to do with helping U.S. companies export their products, some with making recommendations to the federal government about policies that affect the economy. Specifically, I look at high-performance computer companies and manufacturing groups.

"Government spending accounts for 30 percent of supercomputer sales, and the intelligence value that government labs have provided in terms of developing hardware, software, and so forth, has been instrumental in making the industry the most competitive in the world. With the end of the Cold War, one of the biggest concerns now is the defense build-down—it's placing a big question mark on the future of supercomputing companies.

"But a place like NCAR does work that's not directly related to defense issues, and it's clearly a place of immense computing needs. Bill Buzbee said 100 to 1,000 gigaflops was the performance range for running certain climate simulations; I'd say that's a remarkable indication that there's a great demand out there for compute cycles. I've been most impressed by the huge amount of data that's processed and stored in this institution. It's absolutely amazing to me, 30 terabytes of data—unimaginable, really."

Lynda Lester is a writer/editor in the Documentation Group within the SCD User Services Section.

High-Performance Fortran includes Fortran 90

by Jeanne Adams

HPF Fortran
SECOND IN A SERIES

Editor's note: This is the second article in a series on the development of new Fortran 90 language extensions for High-Performance Fortran (HPF) on parallel machines. The first article, "Behind the scenes with High-Performance Fortran: Some history," appeared in the January/February 1993 issue of SCD Computing News. *Future articles will discuss the FORALL statement, directives for distributing data, and directives to control synchronization.*

HPF Language Specification is released

The Version 1.0 draft of the HPF Language Specification was released in January 1993, after the December meeting of the High-Performance Fortran Forum (HPFF). Comments were encouraged from the general public during the review period, which closed March 1. Many editorial and technical changes were made at the March HPFF meeting based on the comments received.

The finalized Version 1.0 HPF Language Specification was released in June 1993. (It is available via anonymous File Transfer Protocol [FTP] from the computer named research.att.com in the directory **netlib/hpf**; the compressed PostScript file is named **hpf-v10-final.ps.Z**.) The HPFF plans to continue development of the specification, and Version 1.0 contains an appendix of topics to be considered for continued research.

The HPF Language Specification includes a full Fortran 90, except for some limitations on COMMON and EQUIVALENCE that are constraints on the alignment of data. (These limitations appear for both FORTRAN 77 and Fortran 90.) The specification also includes added functionality to control distribution of arrays on parallel computers with distributed memory. In brief, the specification includes:

1. The Fortran 90 standard
2. The FORALL statement
3. Directives for distributing data
4. Directives to control synchronization

Both the High-Performance Fortran Language Specification and the Fortran 90 standard are available for reference in the Mesa Lab and Foothills Lab SCD consulting offices.

Language Specification includes HPF subset

The HPF Language Specification also includes a smaller HPF subset. This subset consists of a selected number of Fortran 90 features, the FORALL statement, and a selected number of directives. The subset will provide a standard interim capability for initial implementation and testing by the user community. The HPFF group felt this subset would encourage early development of compilers and precompilers, and use of the new features would occur sooner. (Chapter 9 of the Version 1.0 HPF Language Specification contains a definition of the subset.)

The finalized Version 1.0 HPF Language Specification was released in June 1993.

First, I would like to say that I do not favor subsets for language standards, primarily based on the experience with the subset to FORTRAN 77. Although it has always been considered a proper subset, the FORTRAN 77 subset was never a truly useful language, because various extensions beyond the subset were always present. The result was various nonstandard combinations of the subset language, none of which contained the full language. The Fortran 90 included in the HPF subset, I believe, would not be a viable language in itself, even though it includes most of FORTRAN 77. Most vendors will be moving to the full Fortran 90 language. The HPF subset may delay the introduction of the full HPF language.

This article describes the selected Fortran 90 features included in the HPF subset. The FORALL statement and the directives contained in the subset will be presented in future articles.

Fortran 90 features in the HPF subset

The following Fortran 90 features are included in the HPF subset of the HPF Language Specification.

Array features. The most important Fortran 90 feature in the HPF subset is the array language defined in Fortran 90. Most all of the array definitions are included, except for the *character-type* array features. The following is a list of these features with a reference to a section number in the Fortran 90 standard. In general, this list includes the arithmetic and logical array features, excluding *character-type* arrays.

- Array section notation, including subscript triplets and vector valued subscripts (6.2.2.3)
- Array constructors limited to one level of implied DO (4.5)
- Array assignment on whole arrays and array sections (2.4.3, 2.4.5, 7.5)
- Masked-array assignment, the WHERE block, and the WHERE statement (7.5.3)
- Array-valued external functions (12.5.2.2)
- Automatic arrays (5.1.2.4.1) and assumed-shape arrays (5.1.2.4.2)
- ALLOCATABLE arrays and the ALLOCATE and DEALLOCATE statements (5.1.2.4.3)

FORTRAN 77. The Fortran 90 in the HPF subset includes the full FORTRAN 77 language with the exceptions noted above for sequence and storage association and the use of COMMON and EQUIVALENCE.

The Military Standard (MIL-STD-1753). The Military Standard has been in use for many years, and the features have been extensions to a good many compilers. Users for the most part expect these additions to Fortran 90 to be already available, and the HPF subset includes them. They include the DO WHILE statement, the END DO statement, IMPLICIT NONE, the INCLUDE line, the bit manipulation intrinsic procedures, and binary, octal, and hexadecimal constants in DATA statements. These features were discussed in the article "CF77 5.0 offers Fortran 90 bit manipulation" in the January/February 1992 issue of *SCD Computing News.*

Type specification attributes. The *type-specification* statement does not allow a kind selector and excludes user-defined types. The allowable attributes in a *type-specification* statement are: DIMENSION, ALLOCATABLE, DATA, PARAMETER, INTENT, OPTIONAL, SAVE, EXTERNAL, and INTRINSIC. Note that the FORTRAN 77 statements such as DIMENSION and EXTERNAL are allowed, since all FORTRAN 77 statements are in Fortran 90.

Miscellaneous. Interface blocks facilitate passing the directives across procedure boundaries. Module procedures and generic specifications are not included. Optional arguments (5.2.2) in a subprogram or interface body are allowed, as are keyword actual arguments. (The keyword, if included, is the dummy argument name.)

Certain syntax items are included in the subset: 31-character names, lowercase letters where the processor has the capability, underscore (_) as part of a name, and exclamation mark (!) to indicate comments to follow or to indicate a full comment line.

Intrinsic procedures. All the array functions have been included in the HPF subset, as well as the bit manipulation functions mentioned in the Military Standard section, above. All FORTRAN 77 procedures are part of the subset, which include all the numeric and mathematical elemental functions. The new argument and result types extended by Fortran 90 are included, except for DIM (dimension) arguments that do not provide a known shape at compile time. These routines are discussed in Chapter 13 of the Fortran 90 standard, and are listed in Chapter 9 of the HPF Language Specification.

Features not included in the HPF subset

Many features that are relatively simple to implement, such as CASE and NAMELIST, have not been included if the feature is not needed by the HPF subset. Modules were omitted because the group felt that the added complexity would delay compilers supporting only the subset.

An ideal subset would include some of these relatively simple features. However, this issue was discussed and rejected by the Fortran Standards Committee (X3J3), because the full Fortran 90 language would be delayed and a large number of nonstandard programs would result from varying combinations of features in different subsets.

The full HPF Language Specification includes the full Fortran 90 standard, which, in my opinion, is a better implementation.

Jeanne Adams is a member of SCD's Computational Support Group and coauthor of the Fortran 90 Handbook *(New York: McGraw-Hill, 1992.) She was chair of the American National Standards Institute committee that developed Fortran 90, and is a past chair of the International Programming Languages Committee of the International Standards Organization.*

IMSL Fortran Library Version 2.0 is more efficient, accurate Δ

Compiled by Dick Valent

MATH SOFTWARE LIBRARIES

This article highlights some of the improvements and additions made to the International Mathematical and Statistical Library (IMSL) Fortran Library for Version 2.0, which became the default version on the CRAY Y-MP8/864 (shavano) on June 28.

With the exception of the information about library access on shavano, this information comes with minor modifications from a larger document provided by Visual Numerics, the IMSL Vendor. You can see the complete document on shavano by typing:

news imsl | more

IMSL Version 2.0 contains approximately 160 new routines, covering a wide range of applications. A number of routines are more efficient and accurate. Enhancements include clarification, additional examples and graphics, and fully typeset manuals. Errors in both code and documentation have been corrected.

Three in one

With Version 2.0, the three IMSL libraries IMSLMA, IMSLSF, and IMSLST have been combined into one library named IMSL that contains the routines for math, special functions, and statistics. To access this library on shavano, type:

cf77 -L/lib,/usr/lib,/usr/local/lib -limsl job.f

or

segldr -L/lib,/usr/lib,/usr/local/lib -limsl job.o

Version 2: It's better

Here are some enhancements to Version 2.0:

- Level 3 Basic Linear Algebra Subroutines (BLAS) have been included. Level 2 and 3 BLAS are now documented (that is, they are user callable).
- Redirection of error messages is now allowed so users can separate error messages from printed results.
- An options manager has been implemented to make some routines easier to use and provide for more user control of the computations without increasing the complexity of the argument list.
- Performance enhancements have been made in several areas, but most notably in linear algebra solvers, eigensystem analysis, fast Fourier transforms (FFTs), and random number generators.
- BLAS have been incorporated in more routines.
- Workspace allocation has been improved. Dynamic workspace has been implemented where feasible to provide access to the maximum amount of physical memory available.
- The efficiency of a number of routines has been enhanced as a result of optimization for specific architectures. Routines have been selected and benchmarked.

Math routines

Enhancements to the math routines include:

- Sparse linear solvers and a constrained, linear, least-squares solver
- A spline routine that allows various kinds of constraints, including monotonicity and convexity
- Routines for convolution and correlation, 3-D FFT (forward and back), and linearly constrained optimization
- Performance enhancements
- Striking gains in efficiency for many routines in the linear systems chapter
- Improvements in virtually every routine in the area of eigensystem analysis

In all, 80 subroutines including 30 level 3 BLAS have been added, and 46 level 2 BLAS are now documented, for a total of 126 newly documented subroutines.

Special functions

Enhancements to the special functions include:

- Bessel functions of real order and complex argument
- Mathieu functions, Jacobi elliptic functions, and Fresnel integrals

In all, 15 functions have been added, and 21 functions corresponding to Fortran intrinsics have been deleted.

Statistical routines

Enhancements to the statistical routines include:

- Robust statistical methods
- Multivariate general linear hypothesis tests and linear regression
- Multiple comparisons in analysis of variance
- Analysis of correlation, canonical correlation, categorical data, survival, and principal components and factors
- Goodness-of-fit tests, time series, and density estimation
- Multidimensional scaling, data matrix sorting, and computations with singular matrices
- Random number generators and inverses of distribution functions
- Additional printing options to several routines
- Additional options to some routines dealing with linear models

Comments?

Please address questions or comments to SCD math librarian Dick Valent (e-mail: valent@ncar.ucar.edu).

Dick Valent is head of the Math Libraries Support Group within the SCD User Services Section.

IRJE and MIGS users: Avoid Monday bottlenecks on MSS jobs

by Nancy Dawson

If you use the Internet Remote Job Entry (IRJE) system or the MASnet/Internet Gateway Server (MIGS) to manipulate your files on the Mass Storage System (MSS), you may have noticed long turnaround times or "Giving up waiting for status" messages in your output during the past few months—particularly on Mondays. Why Mondays? Because many users submit long jobs to update many of their MSS files on Mondays, in response to the MSS purge notices they received over the weekend. Here's what you can do to alleviate the overload and get your jobs back quicker:

- Avoid submitting long **.msc** jobs (IRJE users) or long **nrnet mstouchm** jobs (MIGS users) on Mondays. Submit **.msc** or **nrnet mstouchm** jobs during evening or weekend hours for faster turnaround.
- Do not resubmit a job you think may have been lost. Instead, check the newly created **.gwsnap** file to see if your job is in the queue. If your job isn't listed, follow the instructions under "Finding your job" at the end of this article.

Getting the status file

The **.gwsnap** file (which stands for "gateway snapshot") lists all jobs currently in the queue of the IRJE/MIGS gateway computer. This file is like the other status files you probably already use, such as the Daily Bulletin file and the Cray status file. The **.gwsnap** file is updated every five minutes.

IRJE users: To get a copy of the most current **.gwsnap** file, use File Transfer Protocol (FTP) to access your IRJE account just as you do when submitting a job. Then type the following FTP **get** command:

get .gwsnap

Example 1. A typical .gwsnap status file

(See sidebar, next page, for an explanation of the report columns)

```
                              IP Initiated Requests            Page:    1
  SEQN    TRANS   STATE   SCI    RC   IP   REMAIN   NEXT                      SIZE
 ------   -----   -----   ----   --   --   -------  -----------------  ---------
 M66760    GET   que'd   2350   MS   M6   216,192  93/05/07 14:36:01
 IR3250@   GET    netf   4523   MS   IR   216,192  93/05/07 14:07:52   26083328
 IR3217@   GET   que'd   7472   MS   IR   216,192  93/05/07 12:53:00   40128512
 IR3209@   GET    netf   7472   MS   IR   216,192  93/05/07 12:38:56  128671744
 IR3212@   GET    ftpf   7472   MS   IR   216,192  93/05/07 13:26:47  128659456
 IR3220@   GET   que'd   7472   MS   IR   216,192  93/05/07 12:58:17   38965248
 CG8827    GET   que'd   4631   MS   CG   216,192  93/05/07 14:40:23
 M66774    GET   que'd   2350   MS   M6   216,192  93/05/07 14:37:45
 CI2822@   PUT    ftpf   6404   SH   CI   216,192  93/05/07 14:45:01    2020058

                             MAS Initiated Requests            Page:    1
  SEQN    TRANS   STATE   SCI    RC   IP   REMAIN   NEXT                      SIZE
 ------   -----   -----   ----   --   --   -------  -----------------  ---------
 IR3248    DIS   *ftpf   4081   SH   IR   216,192  93/05/07 14:47:25     459991
 IR3097@   DIS    ftpf   4105   SH   IR   216,192  93/05/07 14:37:24     345600
 IR3245    DIS   *ftpf   4346   CA   IR   216,192  93/05/07 14:46:28       2044
 IR3189    DIS   *ftpf   4036   SH   IR   216,192  93/05/07 14:46:31       3974
 IR3252    DIS   *ftpf   4164   SH   IR   216,192  93/05/07 14:47:07       6526
 IR3114@   DIS    ftpf   4105   SH   IR   216,192  93/05/07 14:45:32     298080
 IR3251    DIS   *ftpf   4195   SH   IR   216,192  93/05/07 14:46:14      23278
 IR3167@   DIS    ftpf   3362   SH   IR   216,192  93/05/07 14:37:11    3159604
 IR3167@   DIS    ftpf   3362   SH   IR   216,192  93/05/07 14:41:18    2966471
```

If you want to rename the file before it arrives on your local host, add the new name to the end of the command. For example, to rename the file **aug24.status**, type:

get .gwsnap aug24.status

MIGS users: To get a copy of the most current **.gwsnap** file, use the **gwsnap** verb (or the **gwstat** verb, which performs the same function) with the **nrnet** command. You must add a name for the file on your local host at the end of the command:

nrnet gwsnap aug24.status

Finding your job

Look for your scientist number under the "SCI" column in the **.gwsnap** file to verify that your job is in the queue. If the job you submitted is not listed, check the status file for the MSS to see if the job is queued there. (IRJE users do this with the **get .msstat** command; MIGS users use the **nrnet msstat** command.) If you still can't locate your job, contact the consultant on duty (303-497-1278 or e-mail to consult1@ncar.ucar.edu) to help you trace it.

Nancy Dawson is a writer/editor in the Documentation Group within the SCD User Services Section.

Explanation of .gwsnap report columns

(See Example 1, previous page, for a sample report)

SEQN	Sequence number. A trailing at sign (@) indicates an active job.
TRANS	Transaction type
GET	Get a file
PUT	Put (send) a file
DIS	Outbound Cray file
NUL	No data transfer, such as a **touch** request
STATE	Systems information, not relevant for users
SCI	Scientist number. Look for yours here.
RC	Remote computer
XX	A **.msc** command file
MS	Mass Storage System
SH	Y-MP8/864 (shavano)
CA	Y-MP2D (castle)
TG	Text and Graphics System (TAGS)
NG	Gateway job
IP	Internet node
IR	IRJE job
Anything else	MIGS job
REMAIN	Gauge of attempts to process job
216,192	Not yet processed
NEXT	The time the job is scheduled to be processed. Not a guaranteed time! The system delays processing because (1) the maximum number of jobs—currently 14—are already active, (2) the user already has three active jobs, or (3) the remote computer already has six active jobs.
SIZE	Size of job file in bytes

HINTS FROM THE SCD CONSULTANTS

Get date/time information with Fortran-callable routines on Y-MPs

by Tom Parker

A variety of Fortran-callable subroutines and functions for obtaining date and time information are available on the CRAY Y-MP8/864 (shavano) and the CRAY Y-MP2/216 (castle). This article describes some of the more common routines; for details on usage, please see the **man** pages.

Routines that return the current date/time in ASCII

DATE	Date (mm/dd/yy)
JDATE	Julian date (yyddd)
CLOCK	Time (hh:mm:ss)

Routines to compute elapsed wall-clock time

These routines return the current wall-clock time. By calling the routine twice and computing the difference, you can derive the elapsed wall-clock time.

SECONDR	Elapsed wall-clock time in seconds
TIMEF	Elapsed wall-clock time in milliseconds

Routines to compute CPU time

TREMAIN	Remaining CPU time in seconds
SECOND	CPU time in seconds. By calling this routine twice and computing the difference, you can derive the elapsed CPU time.

Other routines

RTC, IRTC	Elapsed wall-clock time in "real-time-clock ticks"
ICPUSED	Elapsed CPU time in "real-time-clock ticks"
TSECND	Elapsed CPU time for a calling task during a multitasked program
NDAYS, NDYIN	Convert between date and number of days from January 1, 1900. These two routines are in the International Mathematical and Statistical Library (IMSL) Fortran Library. See the IMSL manual for more information (there are no **man** pages).

For example . . .

Example 1 (see next page) shows a sample Fortran program that calls some of these routines.

Tom Parker is a consultant in the SCD User Services Section.

Example 1. A sample Fortran program

```
C DATE
      write(6,'(" Today''s date is: ",a8)') DATE()
C Might print: Today's date is: 05/03/93

C JDATE
      write(6,'(" Today''s Julian date is: ",a8)') JDATE()
C Might print: Today's Julian date is: 93123

C CLOCK
      write(6,'(" Current time is: ",a8)') CLOCK()
C Might print: Current time is: 11:43:11

C SECONDR
      t0 = SECONDR()
      call mysub1
      t1 = SECONDR()
      elapsed = t1 - t0
      print *,'Elapsed wall-clock time in seconds was: ',elapsed
C Might print: Elapsed wall-clock time in seconds was: 1.0703903972171E-2

C TIMEF
      t0 = TIMEF()
      call mysub2
      t1 = TIMEF()
      elapsed = t1 - t0
      print *,'Elapsed wall-clock time in msecs was: ',elapsed
C Might print: Elapsed wall-clock time in msecs was: 11.52453

C TREMAIN
      call TREMAIN(x)
      print *,'Seconds of CPU time left in my job is: ',x
C Might print: Seconds of CPU time left in my job is: 1800.985287796

C SECOND
      t0 = SECOND()
      call mysub3
      t1 = SECOND()
      elapsed = t1 - t0
      print *,'Elapsed CPU time in seconds was: ',elapsed
C Might print: Elapsed CPU time in seconds was: 10.036974

      end
```

SYSTEMS NEWS

Obtain detailed usage reports with new 'gaus' utility

With a new **gaus** utility accessible by Telnet from your front-end computer, you can now obtain detailed reports on computer resource usage for your project or scientist number. Reports can be summed by job class for one day, consecutive days in a month, or monthly for the most recent six months.

You can request two kinds of reports. The first kind of report lists overall General Accounting Unit (GAU) usage for all the following systems: the CRAY Y-MP8/864 (shavano), the Mass Storage System (MSS), the Text and Graphics System (TAGS), and 1-800 Connect Account calls. The second lists information for one system only (shavano, castle, the MSS, or TAGS), but shows GAU usage divided into categories such as CPU hours used, Cray disk drive activity, and MSS disk I/O and tape cartridge activity.

To obtain a report, Telnet from your front-end computer (*not* from shavano) by entering:

telnet gaus.ucar.edu

When prompted for a login, enter:

gaus

You do not need a password.

The new **gaus** utility is interactive and will prompt you with a system of menu items from which to choose. For more information, type **man gaus** on shavano or contact Rosemary Mitchell (e-mail: rosemary@ncar.ucar.edu; phone: 303-497-1235).

Memory limits increased for small memory subqueues

The per-request and per-process memory limit have increased to three megawords for the small memory subqueues on the CRAY Y-MP8/864 (shavano). The new limit affects the following premium (pr), regular (rg), and economy (ec) subqueues:

- pr_Ms_Ts
- pr_Ms_Tl
- rg_Ms_Ts
- rg_Ms_Tl
- ec_Ms_Ts
- ec_Ms_Tl

This change was made because some optimization and program combinations require more memory for the Fortran compiler.

MSS NEWS

More powerful than 'msls': 'msinfo'

A new command named **msinfo** is now available on the CRAY Y-MP8/864 (shavano), the CRAY Y-MP2/216 (castle), and the SCD UNIX front-end computer (meeker). The **msinfo** command lets you get various information about your Mass Storage System (MSS) files and is more powerful than the similar **msls** command. For details, please see the **man** page for **msinfo**.

The 'lread' and 'lwrite' commands are mothballed

The Mass Storage System (MSS) file-transfer commands **lread** and **lwrite** were removed from the CRAY Y-MP8/864 (shavano) and the CRAY Y-MP2/216 (castle) on July 12. These commands have not been supported by SCD since the improved **msread/mswrite** MSS file transfer commands were released approximately two years ago. Since castle has been upgraded to UNICOS 6, there is no longer a need to keep the less efficient **lread/lwrite** on either system.

The syntax of **msread/mswrite** is different from that of **lread/lwrite**; the **msread/mswrite** set has a more UNIX-like format. The differences will require changes for all users who currently have either **lread** or **lwrite** in their code or scripts.

We recommend that **lread/lwrite** users read the **man** pages for **msread** and **mswrite**, available on either castle or shavano. Please contact the SCD Consulting Office with any questions or problems; send e-mail to consult1@ncar.ucar.edu or call (303) 497-1278.

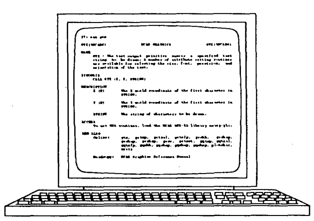

See the man pages for more detailed information about **msread** and **mswrite**.

GRAPHICS NEWS

Note this TAGS correction

The article "Jumbo TAGS combo matrix: Your guide to TAGS capabilities" in the March 1993 *SCD Computing News* contained two inaccuracies on page 19. The **videoumatic** parameter was defined in the text as U-matic when it should have been defined as U-matic SP. Also, the **videobetacam** parameter was defined as Beta when it should have been defined as Betacam SP.

For questions on the Text and Graphics System (TAGS), contact the SCD consultant on duty by sending e-mail to consult1@ncar.ucar.edu or calling (303) 497-1278.

TRAINING NEWS

Learn NCAR Graphics

Are you having trouble with NCAR Graphics? Come to an NCAR Graphics class November 1–3, December 13–15, or January 5–7.

This class will move from the basics—such as how to use the example programs and how to view your plots—to more advanced techniques, such as how to draw X/Y plots, geographic maps, and vector fields; how to fill areas in black and white or color; and how to draw and fill contours. This class not only demonstrates the software in a classroom environment, but gives you hands-on experience.

For further details, including information on cost and location, call (303) 497-1201 or send e-mail to scdinfo@ncar.ucar.edu. Class sizes will be limited, so enroll soon.

DOCUMENTATION

This department lists recently released SCD documents, which you can order by phone or electronic mail. To obtain printed copies of SCD documents, see "SCD Documentation Orders" on the back cover of this newsletter. For information on obtaining documents online, see "Obtain SCD documents via anonymous FTP," below.

Philosophy, structure, strategy: "NCAR Graphics Fundamentals" is here

NCAR Graphics Fundamentals, UNIX Version, Version 1.0, May 1993 (509 pages) is an introductory user guide designed to get you up and running with NCAR Graphics as fast as possible. The manual includes information on the philosophy of NCAR Graphics, the structure of NCAR Graphics programs, and strategies for learning the software. It also includes instructions for compiling, linking, and running programs; instructions for producing and editing program output; and guidelines for using NCAR Graphics utilities.

NCAR Graphics tutorial shows how to contour and map

The NCAR Graphics Contouring and Mapping Tutorial, Version 2.0, May 1993 (440 pages) provides a step-by-step guide to the functionality of the NCAR Graphics geographic mapping and contouring utilities. The manual also covers the Areas utility, an area-processing utility that allows you to fill regions, draw masked lines, and perform other useful functions. The tutorial requires that users have basic Fortran skills and knowledge of how to generate and view graphics.

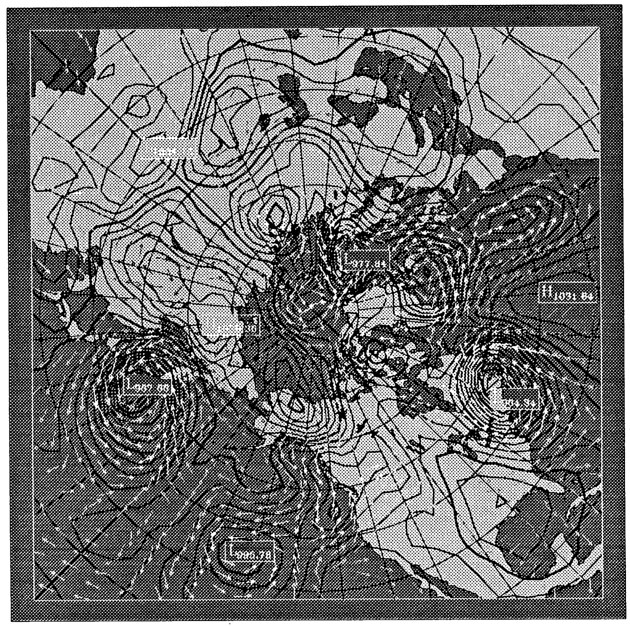

Four NCAR Graphics utilities created this image using data from the U.S. National Meteorological Center.

How to access the SCD Daily Bulletin

The Daily Bulletin is an online daily status report of all SCD computing systems. It is the most current source of news about computing at NCAR, giving information about hardware, software, documentation, communication links, and scheduled and unscheduled computer downtime.

The Daily Bulletin is prepared weekdays by the SCD Consulting Office between 08:45 and 09:00 Mountain Time and again, if necessary, at 16:00.

Interactive access

If you log into the CRAY Y-MP8/864 (shavano), the UNIX front-end computer (meeker), or one of the NCAR divisional computers, type:

dailyb

IRJE access

If you use the Internet Remote Job Entry System (IRJE), use File Transfer Protocol (FTP) to access the computer named windom.ucar.edu. Once you have logged into your account, type:

get .dailyb *filename*

where ***filename*** is the name you assign to the file on your computer.

MIGS access

If you use the MASnet/Internet Gateway Server (MIGS) from your local computer, type:

nrnet dailyb *filename*

where ***filename*** is the name you assign to the file on your computer.

FTP access

The Daily Bulletin is now also available via anonymous FTP from the computer named ftp.ucar.edu in the top-level directory with the filename **dailyb**. For directions on how to obtain the Daily Bulletin and other SCD documents via anonymous FTP, see "Obtain SCD documentation via anonymous FTP," below.

If you have questions, please contact the SCD consultant on duty by sending e-mail to consult1@ncar.ucar.edu or calling (303) 497-1278.

daily

Watch the Daily Bulletin to stay current with new developments in the Scientific Computing Division.

Obtain SCD documentation via anonymous FTP

A growing number of SCD documents is available via anonymous File Transfer Protocol (FTP) on the computer named ftp.ucar.edu. The *User Documentation Catalog*, which includes descriptions of UNICOS documentation, and the yearly indexes of *SCD Computing News* are also available online. The complete list of online documents is in the **README** file in the **docs** subdirectory.

To obtain copies of online documents, follow the steps below.

1. From your local computer connected to the Internet, type:

   ftp ftp.ucar.edu

   or

   ftp 128.117.64.4

2. When prompted for a login name, type:

   anonymous

   Note: If your local computer is a Digital Equipment VAX running VMS, you may need to type:

   "anonymous"

3. Enter your login ID at the password prompt and wait for the ftp> prompt.

   You can obtain a **README** file with a list of the documentation categories (subdirectories) currently available by typing:

   cd docs
   get README
   quit

You can read the **README** file using your own system tools.

If you already know the subdirectory you want, you can use the **dir** (or **ls**) command within directories to list the contents.

4. To transfer a file to your present working directory on your local computer, change directories to the desired subdirectory of **docs** and use the **get** command. For example:

    cd cray
    get ***filename***

    where ***filename*** is the name of the file you want to transfer.

    Caution: If your local computer already has a file with a name identical to the one you want to transfer, your existing file will be replaced with the new file. To give a file a new name on your local computer, type:

    get ***filename newfilename***

5. To terminate the anonymous FTP session, type:

    quit

Conventions used in this newsletter

Bold represents command names, options, filenames, pathnames, directories, and other items that must be typed as shown.

Bold italics represent variables where you provide the substitution (such as ***filename***). In text, *italics* are used to highlight terminology being defined.

`Courier` is used for Fortran programs, shell scripts, and screen displays.

The \ (backslash) followed by pressing RETURN when you are entering a command allows you to continue the command on the next line. In UNIX, a \ escapes (ignores) the next character.

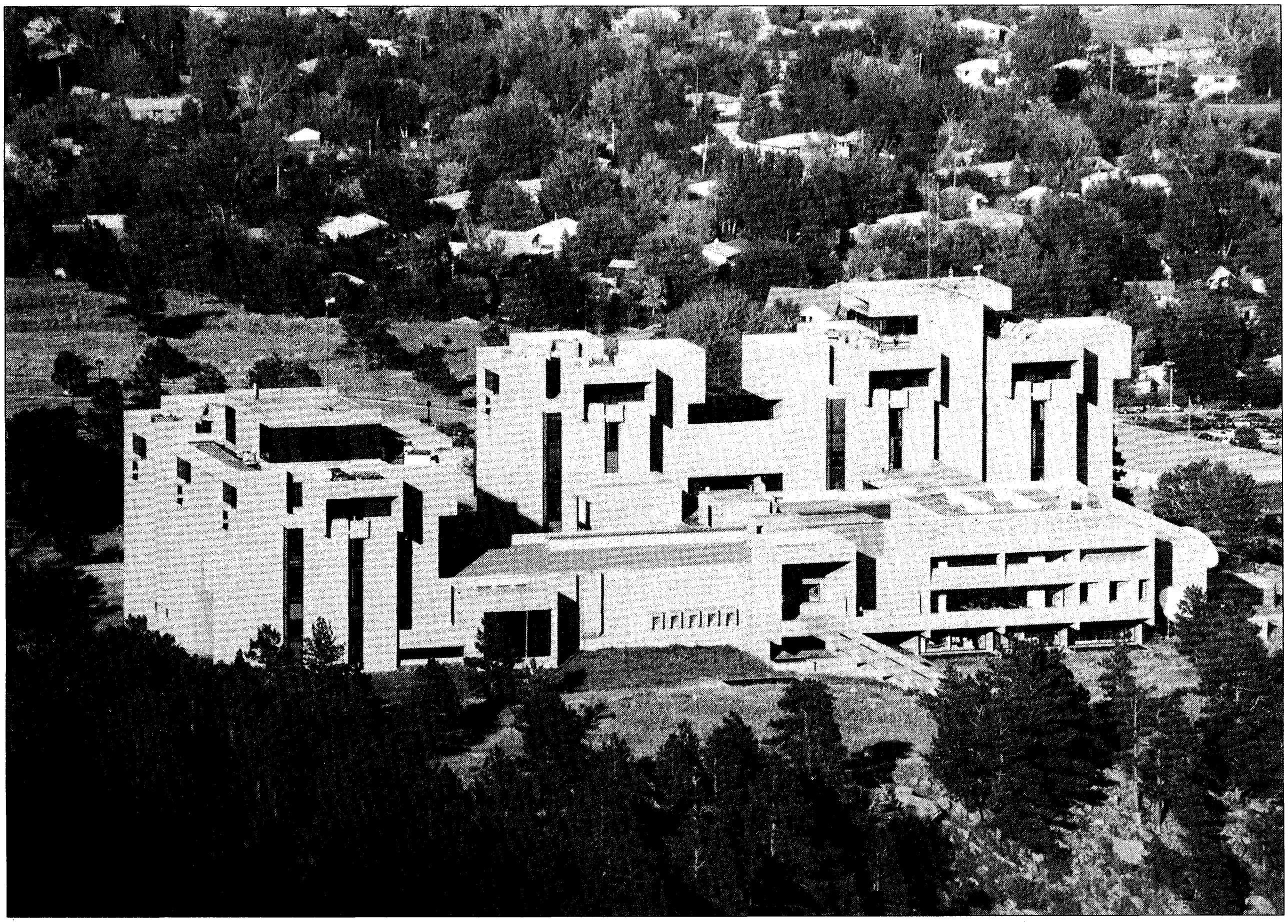

Behind the scenes at NCAR: The rear of the Mesa Laboratory in Boulder. (Photo by Carlye Calvin.)

SCD computer schedule

The SCD computers run continuously, except for scheduled maintenance times and unforeseen equipment or power failures. The computers may be unavailable during the following times:

Computers	Days	Times (Mountain Time)
All computers	Daily as needed	06:00–08:30
CRAY Y-MP8 (shavano)	2nd Monday of month	05:00–08:00
CRAY Y-MP/2D (castle)	3rd Wednesday of month	05:00–08:00

The latest version of the operations schedule is available via anonymous FTP on ftp.ucar.edu under the pathname **/docs/other/op.schedule**. Read the online version of the Daily Bulletin for daily information on the status of all SCD equipment.

Change of address form

☐ Add to mailing list ☐ Delete from mailing list ☐ Change existing entry

Name: ______________________________

New address: ______________________________

____

____

User number ______________ Phone number ______________

Send to: User Information, Scientific Computing Division, NCAR, P.O. Box 3000, Boulder, CO 80307-3000. Or include the above information in e-mail addressed to sylvia@ncar.ucar.edu on the Internet.

SCD SERVICES DIRECTORY

	Contact	Phone (303)	E-mail
General Information			
SCD Consulting Office	Consultant on Duty	497-1278	consult1
SCD Visitor/User Information		497-1225	scdinfo
SCD Course Enrollment		497-1225	scdinfo
Graphics Software Purchase Information		497-1201	scdinfo
Research Data Archive Access	Data Support	497-1219	datahelp
Project & User Number Assignment	Rosemary Mitchell	497-1235	rosemary
Computing Resource Applications	Jana L. Jones	497-1205	jana
SCD Computing News Editor	Lynda Lester	497-1285	lester
NCAR Switchboard		497-1000	
Networking Information			
Data Communications/Networking Information	Consultant on Duty	497-1278	consult1
Network Trouble Reports (24 hour)	Operations Supervisor	497-1200	opns
NCAR/UCAR Network Work Requests	Belinda Housewright	497-1310	scdmg
Operations Information			
Computer Operations	Bob Niffenegger	497-1240	niff
Machine Room	Operations Supervisor	497-1200	opns
Graphics Operations	Andy Robertson	497-1241/42	andy
Tape Librarian	Sue Jensen	497-1245	sue
Output Mailing	Mary Buck	497-1232	docorder
Documentation			
SCD Documentation Information		497-1225	scdinfo
SCD Documentation Orders	Mary Buck	497-1232	docorder

Sending Electronic Mail to NCAR Staff: Use the addresses in the e-mail column above and the appropriate network information below to send e-mail to SCD staff:

- Internet address: *name*@ncar.ucar.edu (IP node 128.117.64.4)
- BITNET users: Use the BITNET address *name*@ncario. Please consult your system administrator for the exact syntax.
- SPAN address: NSFGW::"*name*@ncar.ucar.edu" (DECNET node 33.599 or 34391)
- Telemail/OMNET address: *name*@ncar.ucar.edu (for the Internet address)

For further information on e-mail, please see the UserDoc "Using the NCAR E-mail System."

SCD COMPUTING NEWS
Scientific Computing Division
National Center for Atmospheric Research
P. O. Box 3000, Boulder, CO 80307-3000

Nonprofit Organization
U.S. POSTAGE PAID
Boulder / Colorado
Permit No. 558