Table of Contents

July - August 1994 -- Volume 2, Number 4

IBM Consortium Awarded $22.4 Million for CAS Research Project

By Elisabeth Wecshler

On May 2, a consortium headed by IBM Corp. was selected for an award under the High Performance Computing and Communications Program (HPCCP) Testbed-1 Cooperative Research Announcement (CRA). The consortium will receive $22.4 million from the NASA Computational Aerosciences (CAS) project over three years to carry out research and test new computer systems using highly parallel computer technology.

The NAS Facility at Ames Research Center will be the principal computational site for this new program. NASA's Langley Research Center (LaRC), Hampton, VA, also will participate in the research.

In addition to the IBM T. J. Watson Research Center, Yorktown Heights, NY, and IBM Federal Systems Co., Houston, the consortium includes two aerospace companies (Boeing Computer Services Co., Seattle, and Lockheed Missiles and Space Co. Inc., Palo Alto, CA); two universities (Rensselaer Polytechnic Institute, Troy, NY, and Rice University, Houston); and two independent software vendors (Centric Engineering Inc., Palo Alto, and Intelligent Aerodynamics Inc., Princeton, NJ).

Research Will Be Jointly Managed

The selection, announced by Wesley Harris, NASA associate administrator for aeronautics, Washington, D.C., marks the first time NASA has used the CRA as a joint venture to conduct computer research and testing. The research program will be jointly managed by IBM and NASA. The performance period was scheduled to start in June, when initial configurations of the computer systems were to be delivered to the NAS Facility.

"All parties involved in the consortium are considered to be cooperative partners and will be making contributions to the overall effort -- including cost sharing," said David Bailey, of the NAS Applied Research Branch.

The CRA evaluation team consisted of ten members, seven from Ames, and one each from NASA centers at Langley, Lewis, and Goddard. For background information about the CRA and the evaluation process, see " Cooperative Research Agreement -- A New Approach," in the May-June 1994 issue of *NAS News*.

System Hardware

Project scientists will use the new IBM SP-2 parallel computer systems. IBM will provide three computers for the cooperative research: A 160-node SP-2 system will be located at the NAS Facility, a 48-node system at LaRC, and an eight-node system at Rice. Under this agreement, NASA does not actually purchase any computer equipment -- the computer systems will be removed in June 1997, at the end of the performance period.

"The IBM SP-2 will outperform the conventional mainframe computers we now have at a significantly lower cost," Bailey said. "One thrust of this effort is to learn more about parallel computers," he added. "Our goal is to have a highly parallel computer system serve as our production supercomputer here at NAS."

Research Projects

An important component of this cooperative agreement is a package of research projects in the areas of computational aerosciences and system software for highly parallel computers. "We were very pleased by the quality of the research proposals," Bailey said.

The institutions, principal investigators, and titles of research projects are, as follows:

- Boeing Computer Services; K. Neves, *et al.*; "Rotorcraft Design and Optimization" and "Wing Aerodynamics and Aeroelastics"
- Rensselaer Polytechnic Institute; M. S. Shephard; "Scalable Parallel Automated Adaptive Finite Element Methods"
- Lockheed Missiles and Space Co. Inc.; Y. Tassa, *et al.*; "Parallel Structure and Aerothermodynamics Modules for Multidisciplinary Design Optimization"
- Centric Engineering Systems; T. Hughes, *et al.*; "Multiphysics Product Simulation"
- Intelligent Aerodynamics; T. Jameson; "Computational Aerodynamics on Massively Parallel Processor Systems"
- Rice University; J. E. Dennis, *et al.*; "Scalable Parallel Direct Search Methods for Multidisciplinary Design Optimization"
- IBM T. J. Watson Research Center; V. K. Naik; "Adaptive and Dynamic Parallel Job Scheduling"

- IBM Federal Systems Co.; R. Coyne, *et al.*; "High Performance Storage System"

In addition to these research collaborators, CAS-funded scientists at Ames, LaRC, other NASA centers, universities, and aerospace companies around the country, will also be using the HPCCPT-1 computer systems.

"Accounts will be granted and computer time will be allocated to CAS users nationwide. People within NAS will be given accounts as their duties require," Bailey said.

NASA Ames scientists participating in the CAS program (mostly in the Fluid Dynamics Division) will run their applications on this system, Bailey said. "Another part of NASA's direct involvement in this effort is that a number of people at all three NAS branches will be analyzing and evaluating the effectiveness and performance of the system," he added.

Areas likely to be investigated are: possible performance bottlenecks, difficulties with operating system software, weak spots in the overall system management software, and I/O facilities, he said.

After evaluating the system, NAS researchers and staff may identify certain aspects that may need extra attention, Bailey noted. At that point, NAS may enter into joint development efforts with IBM -- still within the funding and overall plan of the CRA -- to develop tools for enhancing the effectiveness of the system.

Computer Architecture

Each node of the SP-2 systems at the NAS Facility and LaRC is an RS6000/590 workstation module. "At the present time, the IBM 590 workstation is arguably the most powerful single-processor RISC [Reduced Instruction Set Computer] workstation, particularly in terms of actual sustained performance on large memory scientific calculations," Bailey said, adding that in the latest NAS Parallel Benchmark (NPB) report, "the 590 outperforms some vector and highly parallel supercomputers on certain benchmarks."

Significant are the two separate integer computation units and the two 64 bit floating point computation units for each module, Bailey said. But what really sets the system apart is the main-memory-to-cache bandwidth of four 64 bit words per clock period.

All of the nodes in the SP-2 systems at the NAS Facility and LaRC will have at least 128 megabytes (MB) of main memory, at least two gigabytes of disk space, and two Ethernet connections. "Some nodes have more than these amounts of memory and disk space, and some have HiPPI [High Performance Parallel Interface] and other I/O interfaces. Each system has an external disk system in addition to the disks on each node," Bailey said.

All nodes are connected with a specially designed, high-speed switch. This switch is an outgrowth of

IBM's TF-1 project, Bailey explained. The network topology consists of an omega network with a hierarchy of crossbar switches. Latency is less than one microsecond, and approximately 30 microseconds with software. The node-to-node bandwidth is 40 MB per second bidirectional and approximately 30 MB per second with software.

The overall system statistics, together with those for the NAS CRAY Y-MP C90 for comparison, are shown in Figure 1.

System reliability is expected to be good because of the use of commodity workstation modules (except for the switch) and off-the-shelf AIX/6000 software. "We'll be able to compare loosely coupled workstation clusters with parallel computers," Bailey noted.

NPB Performance Results

Several months ago, IBM submitted a complete set of NPB results for a single 590 workstation and, more recently, released some preliminary results for a 16 node and 64-node SP-2 system. These timings are listed in Figure 2.

Estimates for the sustained single-job performance on the NAS-based SP-2 can be made on the basis of currently available NPB results for the SP-1 and the SP-2. "Comparing the three CFD [Computational Fluid Dynamics] application benchmarks, it appears that a 128-node SP-2 will perform roughly on a par with a full [CRAY Y-MP] C90/16," Bailey said, noting that since the NAS Facility will have a 160-node SP-2 "there's room for error in these estimates."

"It is clear that the NAS-based SP-2 system will be a very powerful computer system -- with more than twice the main memory and disk space of the C90," Bailey said. "If it does outperform the C90 on the NPBs, this will be an important milestone in the highly parallel computing field."

Next Article | Contents | Main Menu | NAS Home

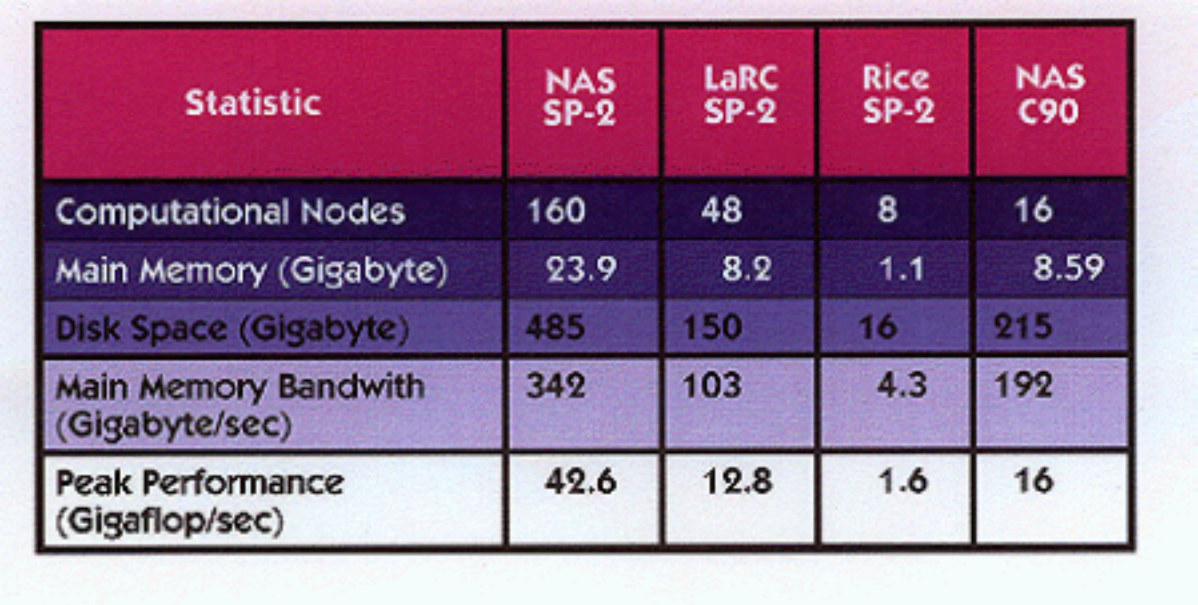

Statistic	NAS SP-2	LaRC SP-2	Rice SP-2	NAS C90
Computational Nodes	160	48	8	16
Main Memory (Gigabyte)	23.9	8.2	1.1	8.59
Disk Space (Gigabyte)	485	150	16	215
Main Memory Bandwith (Gigabyte/sec)	342	103	4.3	192
Peak Performance (Gigaflop/sec)	42.6	12.8	1.6	16

Figure 1 - This table compares the hardware capabilities of the SP-2 that will be provided by IBM and the NAS CRAY Y-MP C90. The prefix "giga" means 10^9 in each case.

to article

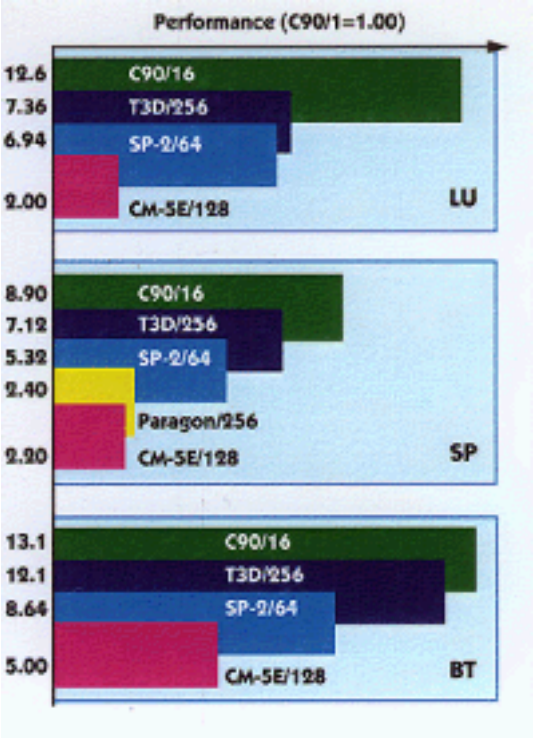

Figure 2 - This chart compares the performance of a 64- node SP2 with other parallel supercomputers on the three CFD applications of the NAS Parallel Benchmark suite. The rates have been scaled so that a single processor of a C90 scores 1.00.

to article

Next Article | Contents | Main Menu | NAS Home

New Technique Reduces Filesystem Recovery Time

by Nicholas Cardo

A set of three NAS-developed software programs has dramatically reduced the amount of time and the number of tape cartridges required to back up data on */m*, the near-term file storage system for the NAS CRAY Y-MP C90, **vonneumann**. For NAS users, this means that the meantime to recovery from a major disk failure has improved, and the filesystem can be returned to service in significantly less time.

The programs, which have improved the recoverability rate of */m* data by as much as 91 percent, were developed by members of the High Speed Processor group in the NAS Computational Services Branch.

Currently, the */m* filesystem has approximately 130 gigabytes of disk cache and about one terabyte of data managed by the Cray Data Migration Facility (DMF). This large filesystem allows data to reside online and, correspondingly, increases the amount of data that must be backed up. Usually, backing up filesystems of this size takes lots of time (8 to 12 hours) and introduces possible recoverability problems. In the event of a catastrophic filesystem failure, the meantime to recovery is long due to the physical amount of data to restore.

Previously, backups of */m* required eight hours to complete and 45 1000-megabyte tape cartridges. When the backup completed, it was no longer an accurate representation of the filesystem, because over the course of eight hours the contents of the filesystem changed significantly.

To attain the objective of obtaining an accurate backup of this large volatile filesystem, it was determined that a backup should complete within two hours. This would improve -- by at least 75 percent -- the ability to restore an accurate representation of the filesystem in case of a catastrophic failure. Another objective was to assure that users see no change in the residency of their files. By utilizing DMF to manage the data, the size of nightly manual backups is reduced.

When a file is migrated, the data blocks are assigned to a temporary file and a migration request is sent to DMF, which controls the duplication of the data blocks to tape. If space is needed on the filesystem, the data blocks are released from this temporary file. When a migrated file is retrieved, DMF first checks to see if a temporary file still exists with the data blocks; if it does, the data blocks are assigned back to the original file, otherwise DMF retrieves the data from tape. When a file is retrieved, it is marked as "dual state," meaning that the data blocks on tape match those on disk. Since the data blocks for dual-state files already exist on tapes controlled by DMF, only the inode (information about the files) is backed up,

making the amount of data to back up smaller.

By taking advantage of DMF's features, the three programs, **dmmighit**, **dmmigput**, and **dmmigget**, were developed to migrate data blocks to tape. First, **dmmighit** produces a list of inodes whose data blocks have not yet been migrated. Then, **dmmigput** sends migration requests to DMF to duplicate the data blocks for the inodes selected by **dmmighit** to tape. Finally, **dmmigget** sends recall requests to DMF, which assigns the data blocks back to the original file. The file is now marked as dual-state, with the data blocks assigned to the original file.

Backups of */m* now usually complete in under 90 minutes and require six cartridges. So far, the record for the shortest backup is 45 minutes, using three cartridges. Besides improving recoverability and saving computer time and tape media, the NAS Control Room staff has gained back about 200 hours per month that were used to monitor and mount tapes for backups -- freeing them to do other needed tasks.

For further details on the techniques used to improve data recoverability at NAS, send email to **cardo@nas.nasa.gov**.

Nicholas Cardo is a member of the NAS High Speed Processor group.

Next Article | Contents | Main Menu | NAS Home

PLOT3D Team Awarded Fourth Largest-ever Prize By NASA Space Act Program

By Elisabeth Wecshler

The fourth largest Space Act Award in the history of the NASA program was presented to Pieter Buning, research scientist in the Computational Technology Branch of the Ames Fluid Dynamics Division, and three support staff members for their work on PLOT3D. This software program is credited with revolutionizing scientific visualization and analysis of three-dimensional Computational Fluid Dynamics (CFD) solutions.

The $35.5K award recognized PLOT3D's substantial commercial potential, proven cost savings, and innovation, said Paul Kutler, Chief, Fluid Dynamics Division, who presented checks to the four recipients on May 9 at the NAS Facility.

Innovativeness and Commercial Use

Award submissions were judged by the NASA Headquarters Inventions and Contributions Board to determine the technology's innovativeness and commercial viability. Winners were eligible for cash awards ranging from $250 to $100,000.

In addition to Buning, who received $25,000, Pat Elson and Larry Pierce, both of Sterling Software, and Pam Walatka, Computer Sciences Corp., each received $3,500 for their work supporting the PLOT3D development.

Buning began writing PLOT3D in 1982. The PLOT3D team worked together from 1987 to 1992 under the direction of Val Watson, currently senior scientist in the Fluid Dynamics Division. The software was released to COSMIC, NASA's Software Technology Transfer Center, University of Georgia, in 1991.

Linked Wind Tunnels, CFD Work

The impact of PLOT3D has been widespread and significant. One of its first important contributions was "to establish a link between the experimental (physical) solutions of wind tunnels and CFD simulations," Buning said. The software, available on several computer platforms, has been used in a variety of NASA

and military aircraft, space, and missile programs, as well as in private industry.

PLOT3D is credited with enhancing the development of graphical workstation technology "by demonstrating a sophisticated solution to an essential need in the area of scientific visualization for complex flow field topologies," according to a NASA document describing the award.

Precursor of FAST

PLOT3D, as the precursor of FAST (Flow Analysis Software Toolkit), has contributed substantially to automobile aerodynamics, computational chemistry applications, bio-fluids, hydrodynamics, atmospheric weather simulations, and the petroleum industry.

PLOT3D can calculate any one of 74 grid, scalar, vector, particle-trace, and shock-wave functions; output can be displayed on-screen from any angle, printed, plotted, and animated, according to Walatka.

Although PLOT3D lacks a menu interface common in current-generation visualization programs, its command-line interface is fairly straightforward to learn and has allowed it to be ported to different machines more easily, Buning said.

Cost Savings are $12 Million So Far

Estimated cost savings to government and industry are $12 million to date, with a projected $2 million a year in future savings, according to NASA. Development costs were approximately $600,000.

Elson, currently NAS User Interface Manager, distributed PLOT3D beta test code and helped with the documentation. Pierce, who last month rejoined the Computational Technology Branch as programmer-analyst, served as technical facilitator and user contact for PLOT3D, incorporating a number of enhancements to the program, as well as directing the COSMIC submission. Walatka, technical writer for the NAS Systems Development Branch, wrote the PLOT3D documentation.

Next Article | Contents | Main Menu | NAS Home

Space Act Award recipients (from left) Pieter Buning, Pat Elson, Pam Walatka, and Larry Pierce are presented checks totaling $35,500 by Paul Kutler (right), Chief, Ames Fluid Dynamics Division, for their work on PLOT3D.

to article

NAS Consultants Help New and Seasoned Users

By Jill Dunbar

Nearly every day of the year, experts in the NAS user consulting groups solve both complex and basic problems for users of high-speed processors (HSP) and massively parallel processors (MPP). The consultants bring almost two centuries of combined computing experience, along with an assortment of advanced degrees, to address technical needs and solve problems for NAS users.

The consultants share their expertise with users -- comprised of scientists from academia, industry, other NASA centers, and national research labs -- primarily through phone calls, but they also spend time one-on-one with local users. In addition, the consultants optimize and "parallelize" user codes, give training classes, provide technical information for the *NAS User Guide* and other online documentation, and write articles for *NAS News*.

Optimization Increases Productivity

The HSP consultants specialize in Fortran optimization, vectorization, and multitasking on code for the NAS HSP (**vonneumann**) and the Aeronautics Consolidated Supercomputer Facility system (**eagle)**, both CRAY Y MP C90 systems. HSP users' computing problems are primarily related to Computational Fluid Dynamics (CFD), and to a lesser degree, computational chemistry, space sciences, and biology.

By helping users optimize code, the consultants can significantly improve the code's efficiency. Users who have made their codes available to the consultants for optimization have found an average decrease in overall CPU time of 40 percent -- some as much as 68 percent -- according to George Myers, HSP consulting group lead. Total CPU can be divided into two categories: user and system CPU time, Myers said. Code optimization can affect either or both of these categories. Improvements in I/O performance will affect system CPU time.

Another key indicator of improvements in code efficiency is the megaflops (MFLOPS) rate, Myers said. "One notable success in this area was a case in which we increased megaflops by 513 percent!" Increases in MFLOPS are usually accompanied by decreases in user CPU time. Any decrease in CPU time is a gain for the users because it means they can get more done in less time, Myers explained. "Optimizing code is one area where we can help users," he said, adding that users who prefer to do their own optimization can still benefit from the consultants' experience. "We can point users in the right direction and give them lots of suggestions."

Myers is particularly proficient in optimizing and multitasking code, as are Clayton Guest (who focuses

on I/O optimization) and Terry Nelson. Other members of the group are Bob Hirsch, with skills in HSP graphics libraries, and R.K. Owen, who specializes in math libraries and teaches an introductory numerical methods class, which he describes as "how to do numerical computations without blowing things up."

Cutting-edge Parallel Techniques

The consultants in the parallel systems science support group devote much of their time to keeping up with cutting-edge technology on parallel architectures and their uses. They focus on making code parallel for the Intel IPSC/860 and Thinking Machines CM-5, as well as for the Paragon XP/S, a closed system with a limited number of users. The consultants then apply this knowledge to assist scientists and researchers with parallel application code development and algorithms, according to Chuck Niggley, parallel systems science support group lead. As with the HSP users, most parallel systems users are engaged in CFD research, although a substantial amount of electromagnetics work is done on the Paragon system. The consultants' latest challenge is to "come up to speed" on the IBM SP-2 system, recently selected as the High Performance Computing and Communications Program testbed system (see *IBM Consortium Awarded $22.4 Million for CAS Research Project* in this issue).

Niggley is adept at using compilers and converters to help users make their codes parallel. The other parallel systems consultants are Ed Hook, who specializes in Fortran and C optimization; Subhash Saini, with unique knowledge in parallel mathematical libraries, and Bill Saphir, who concentrates on PVM (the Parallel Virtual Machine) and other message-passing libraries.

In addition to helping new users become familiar with the NAS parallel systems, the consultants research and isolate software and library "bugs," assist with system modifications, help test new software releases, and take their turn at giving training classes, including parallel programming strategies such as load balancing and data parallelism, and High Performance Fortran techniques.

All of the consultants participate in developing and teaching classes at the NAS Facility, ranging from basic classes for new HSP and parallel systems users during each New Operational Period (NOP) to courses on the latest computing techniques. Last month, the parallel systems consultants presented two days of classes for the Parallel Systems NOP, which began June 6.

Many of the courses are videotaped and available to off-site users through the NAS Documentation Center; selected courses can also be taken through a video teleconferencing capability *(see the May-June 1994 issue of* NAS News *for more information).*

In addition, the consultants informally train the NAS User Services analysts, who are usually the first to answer users' calls, providing system status and basic system information.

Varied Projects Fill `Spare' Time

In their "spare" time, consultants in both groups also find time to: write online tutorials and examples, run benchmarks for newly installed operating systems and compilers, investigate new compilers, develop test suites for new software, write and present papers at parallel and supercomputer conferences, and study new methods for solving user problems.

They also track resource utilization to find out which tools and libraries are being used, then train scientists to use those resources more efficiently. For example, by gathering usage information on the Session Reservable File System (SRFS), consultants discovered that over 80 percent of the users who made reservations for space on the HSP computers typically used less than 80 percent of their reservation -- keeping unused space from being available to others. The consultants now notify these users and explain how to more accurately estimate their space requirements upfront.

Help Available Around-the-clock

Each month the NAS User Services analysts receive an average of 1,350 calls -- of these, about 100 calls are passed to the HSP consultants and approximately 25 calls to the parallel group (which continue to increase with the growing number of parallel systems users). User Services analysts are available 24 hours a day, year-round, except for government holidays. Currently, the HSP group provides support weekdays from 5:00 a.m to 6:00 p.m., Pacific time. Parallel consultants are available from 6:00 a.m. to 4:00 p.m., Pacific time.

To reach a NAS consultant, call User Services at (415) 604-4444 or (800) 331-USER (8737), or send email to **nashelp@nas.nasa.gov**.

Next Article | Contents | Main Menu | NAS Home

Hierarchical Network Management System Offered Through COSMIC

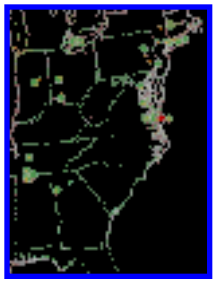

By Elisabeth Wecshler

The NAS-developed Hierarchical Network Management System (HNMS) 2.0 was released in April to COSMIC, NASA's Software Technology Transfer Center, and is now available to other organizations.

"HNMS enables an administrative staff to monitor a very large Internet Protocol network," said Jude George, of the network development group. Since the entire network can be viewed as a whole, HNMS is designed to alert the administrator when something on the network becomes unreachable. HNMS also automatically collects statistics on network traffic to allow the administrator to analyze network utilization, and stores the data in one convenient place.

Extensions to HNMS Are Encouraged

George hopes that other network administrators will download the nonproprietary HNMS software from COSMIC and use it in their organizations. NAS encourages the development of extensions to HNMS that can be shared with the entire Internet community -- for example, a new type of display for new classes of networks such as Asynchronous Transfer Mode.

Before HNMS was developed, a user would call NAS User Services when a problem occurred, and the staff would run tests to determine if any network links were malfunctioning. Now, HNMS will inform the administrator within one minute if there's a problem, George said. "This makes the administrator's job easier and gets information about network problems out to the user faster."

"In most cases, the administrator will be able to see the problems before users are impacted," said Doug Sheppard, NAS network operations team lead. Special tools were designed to send out warnings when errors reach a certain level, he added. "For instance, users benefit because HNMS alerts the staff of increasing error counts, allowing us to repair circuits before the network goes down."

The network operations group and the User Services staff have assumed responsibility for maintaining HNMS at NAS. The development team, including George, group lead Jim McCabe, and Leslie Schlecht, will continue to add program enhancements.

End Users Normally Won't Have Access

NAS users currently do not have direct access to HNMS, George said, adding that this policy would be reviewed by other organizations' operations staffs. One advantage to users having direct access to HNMS includes being able to troubleshoot their particular network problems. The disadvantages include giving users access to passwords and certain security-based information (such as permitting the reconfiguring of the network) that wouldn't be appropriate, George said.

In special cases, some users or groups of users will be able to run HNMS, but "it's CPU intensive and the load increases for each user interface added, so you want to limit users with direct access," Sheppard said. "If users are working on a special development project, which NAS normally wouldn't monitor, we can work things out to give them access to HNMS." A large-scale overhead projector is scheduled to be installed in the NAS control room so that staff and special user groups can view the network's status while they work. This will increase HNMS network monitoring from the current eight hours a day, five days a week to 24 hours a day, seven days a week.

Easily Adaptable To Other Programs

Most networks currently use Simple Network Management Protocol (SNMP), a data-collection protocol. HNMS also uses Hierarchical Network Management Protocol (HNMP), developed at NAS to make network data easily available to other software programs.

Sheppard emphasized that the HNMS architecture is "easily adaptable to other network configurations because it's environment independent."

Customized versions of HNMS were written for Supercomputing '93 (Portland, OR); for Bellcore, (Redbank, NJ), to test telecommunications equipment; and for Wellfleet Communications (Boston), to test network routers.

The beta testing of HNMS this spring involved Boeing Co., Bellevue, WA; United Technologies Research Center, East Hartford, CT; Cygnus Support, Mountain View, CA; National Center for Supercomputing Applications, University of Illinois, Champaign-Urbana; and Oregon State University, Corvallis.

Obtaining More Information

For more information about HNMS, the following articles are available from the NAS Documentation Center:

- "The NAS Hierarchical Network Management System," *Integrated Network Management III*, H.

G. Hegering and Y. Yemini (editors), Elsevier Science Publishers, Amsterdam, 1993

- "NAS Develops New Approach to Managing Local and Wide Area Networks," by Jude George and Leslie Schlecht, in *u/nas/news* Vol. 7 No. 7, August 1992

- *The NAS Hierarchical Network Management System Version 2.0* (documentation), March 1994

Next Article | Contents | Main Menu | NAS Home

Figure 1

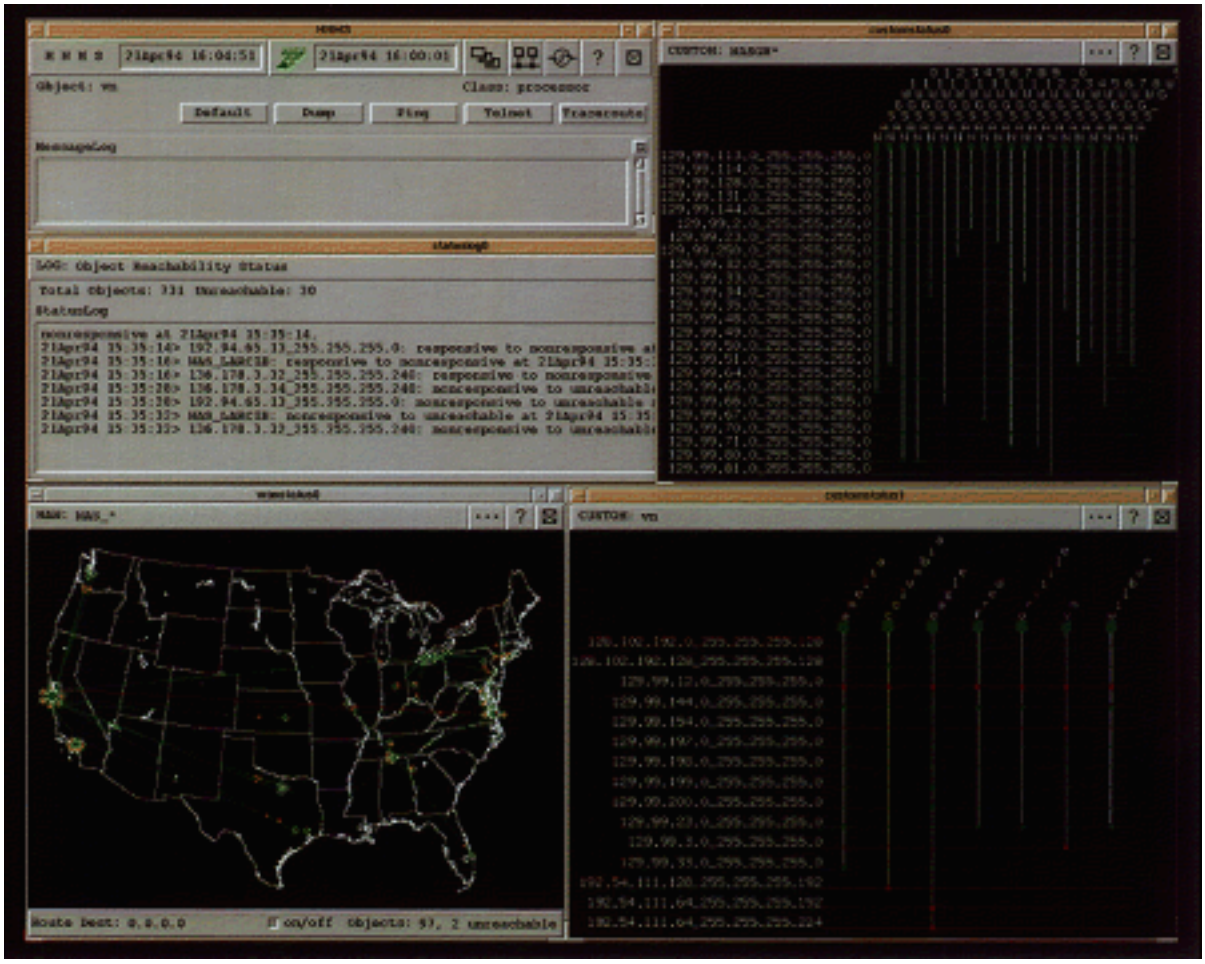

Hierarchical Network Management System (HNMS) 2.0 displays selected portions of the NAS local- and wide-area networks. The log collates data from the diagrams and alerts the network administrator to events occurring at any location on the network.

to article

Figure 2

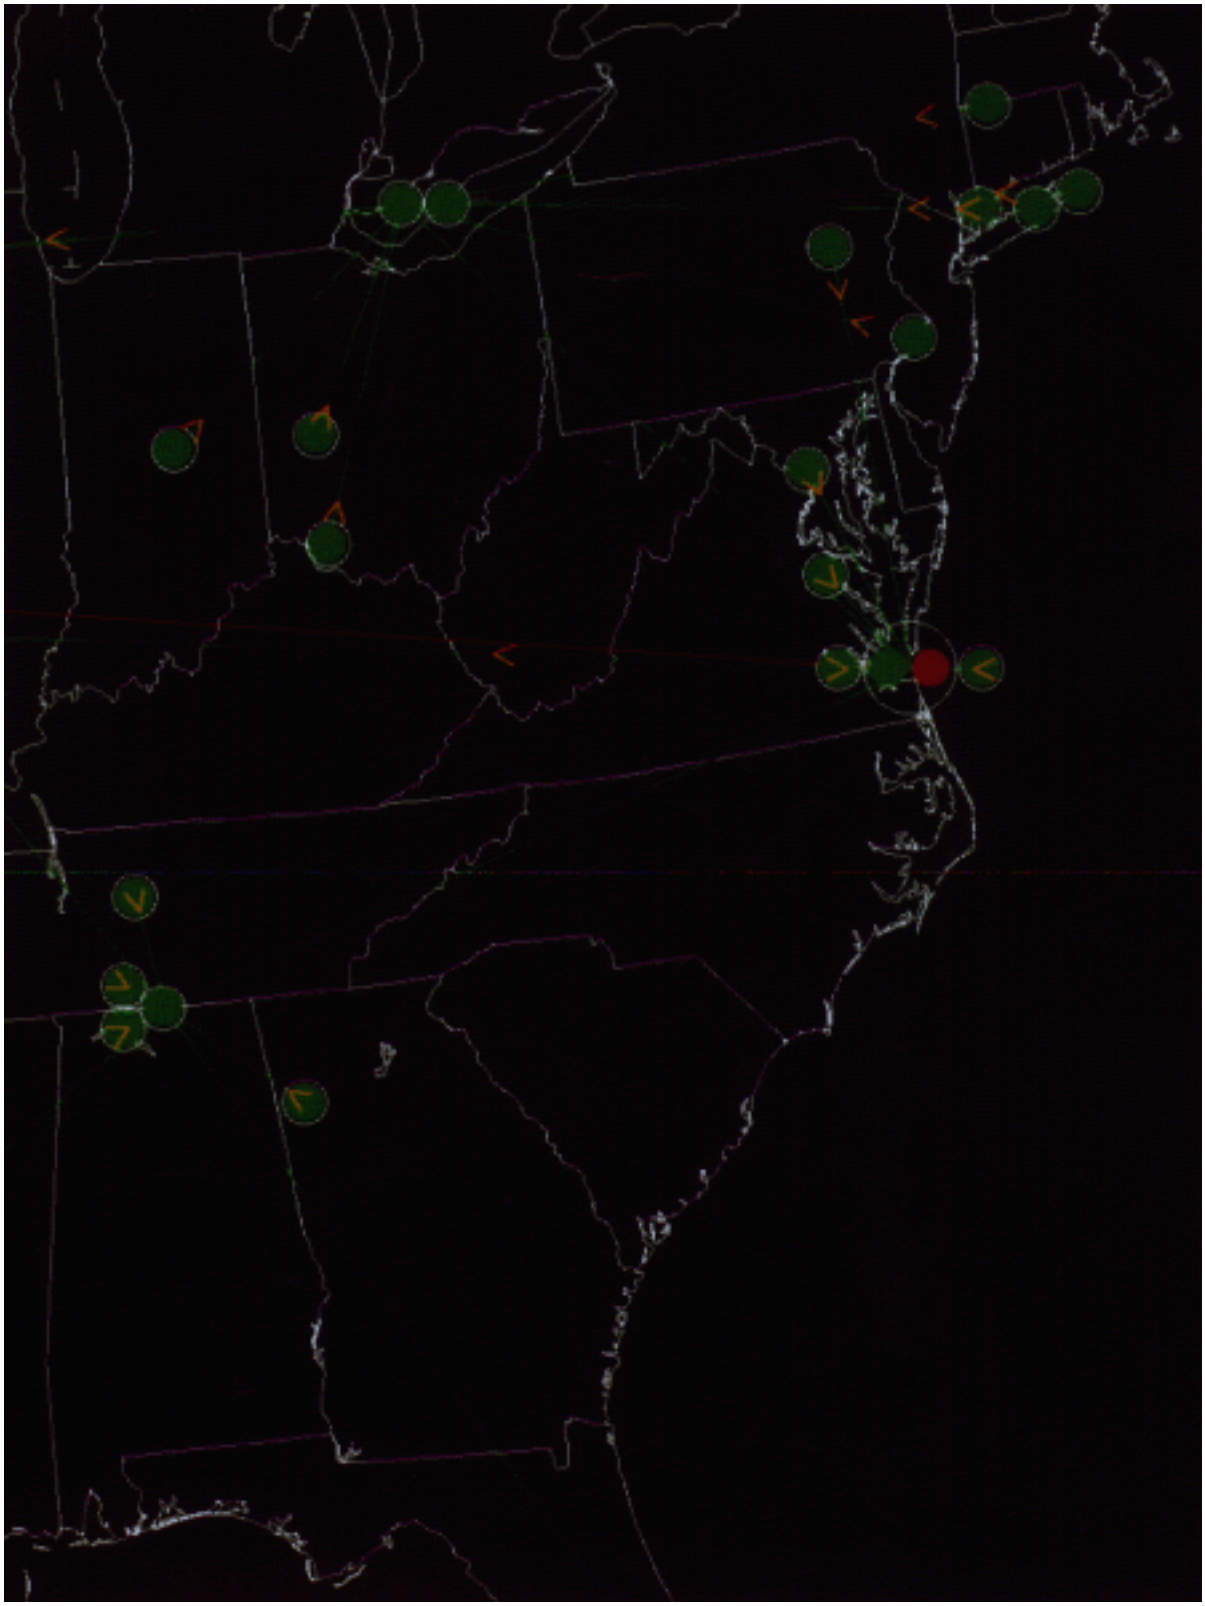

A close-up of the eastern portion of the NAS wide-area network, AERONet, shows that one of the links from NAS to NASA's Langley Research Center (LaRC) is down. Windows (not shown) display a portion of the LaRC network, as well as names and phone numbers for staff serving as network points-of-contact at that location. The failed link was later traced to a malfunctioning interface on a data service unit -- a piece of network hardware.

Figure 2

 to article

AERONET II Increases Network Bandwidth and Efficiency

By Elisabeth Wecshler

The NAS network development team recently upgraded AERONet, NASA's Aeronautic Research long-haul network, to improve bandwidth for individual scientists' requirements, according to Anthony Lisotta, network development engineer in the NAS Systems Development Branch and AERONet team member. Lisotta said that the upgrade will also ensure more efficient bandwidth use across the shared, or "backbone," portion of AERONet.

AERONet is the high-performance wide-area data communications network that provides access to the NAS supercomputing facility for users conducting aerospace research at government, university, and commercial sites throughout the U.S. (*See "Background and Accomplishments of AERONet I."*)

Implementation of AERONet II

Implementation of the upgrade, called AERONet II, was scheduled for June, Lisotta said. This upgrade represents phase two of scheduled improvements to AERONet; the entire upgrade project is expected to be finished by September 1995.

Two user models were developed to guide the NAS team's work, said Jim McCabe, network development group lead. One model consists of scientists working independently -- transferring files and geometry data between systems at NAS and their local sites. This model is used primarily at aerospace companies, universities, and some NASA centers, where Computational Fluid Dynamics (CFD) solutions tend to generate less than ten gigabytes (GB) of data. Also, scientists at these sites are more concerned with network access restrictions because of the proprietary nature of their work, he explained.

"For this model, cross-country delays on the order of 100 milliseconds [ms] are acceptable, as are capacities of between 1.5 and 45 megabits per second [Mb/sec]," McCabe said.

Second User Model Developed

The second model, used by scientists at NASA centers Langley (LaRC) and Lewis (LeRC), provides a high-performance environment for collaborative and distributed visualization of CFD solutions that approach one terabyte (TB) of data. Network requirements for this environment are more stringent, McCabe explained, because of delay requirements to support full-motion video (50-60 ms) and capacity

requirements for the transfer of TB-size solution files and megabyte-size geometry data.

For this reason, network capacities "in the hundreds of megabits per second" range are needed, McCabe said.

In July, scientists at LaRC and LeRC should see an increase in throughput and speed, and a decrease in time required to transfer files, Lisotta said. Scientists working at other sites won't see major improvements until later.

Customized Bandwidth in Future

"Once we begin evaluating network traffic to the users, we'll be able to provide more bandwidth on a case-by-case basis," Lisotta said. "In the past, we haven't been able to customize bandwidth and other network services" for individual remote sites. Lisotta estimated that customizing bandwidth requirements for industry sites could begin after the completion of AEROnet III.

Switched Multimegabit Data Service (SMDS) and Frame Relay are the two technologies that enable improvements to AEROnet's bandwidth. SMDS relies on statistical averaging to interleaf, or weave, multiple connections across a single circuit, Lisotta explained. The Frame Relay service maximizes bandwidth efficiency for a backbone network by guaranteeing the committed information rate, which the scientist specifies before connecting to the network, and providing a predetermined additional amount of bandwidth (called the burst rate) to the user, depending on network utilization at that time, he added.

Performance Results `Impressive'

The performance results for AEROnet's bandwidth capacity have been impressive, according to McCabe:

- For NAS-based leaf (or industry) sites, performance improved from 56 kilobytes per second to DS1 [Digital Signal] circuits, a type of T1 circuit that transmits 1.5 Mb/sec -- representing a 2,700 percent increase.

- For NASA's Jet Propulsion Laboratory and Marshall Space Flight Center backbone sites, performance improved from 2xDS1 to 2xE1, a type of circuit that transmits 2.03 Mb/sec -- representing a 133 percent increase.

- For LaRC and LeRC backbone sites, performance improved from 4xDS1 to DS3, a type of circuit that transmits 45 Mb/sec -- representing a 750 percent increase.

In addition, the NAS network interface to AEROnet II was upgraded to a Synchronous Optical NETwork (SONET) OC-12 transport node, operating at 622 Mb/sec, in order to support DS3 circuitry and prepare for future upgrades.

Plans For AERONet III

The next upgrade, called AERONet III, will incorporate Asynchronous Transfer Mode (ATM) technology and SONET-level services. An ATM prototype was built between NAS, LaRC, and LeRC this spring to evaluate ATM as a potential long haul network technology suitable for distributed visualization (including distributed virtual reality) applications. If the technology proves to be successful, a vendor-provided ATM service offering high-speed, nationwide data communications to the general AERONet community will replace the prototype.

Other planned service upgrades include SONET at the OC-3c (155 Mb/sec) rate to LaRC and LeRC, to support Computational Aerosciences communications requirements, and making DS3 available to JPL, McCabe said. SMDS may be expanded to include more sites in the JPL and MSFC regions, as well as additional sites in the NAS, LaRC, and LeRC regions, he added.

As Frame Relay, SMDS, and ATM are implemented in AERONet III, issues such as these products' ability to work together, as well as their compatibility with IP routing, will be examined closely. McCabe fully expects that these services will predominate in AERONet III, replacing most of the point-to-point circuits.

For more information about the AERONet upgrade, send email to **lisotta@nas.nasa.gov**.

Next Article | Contents | Main Menu | NAS Home

Visualization Lab Video Highlights Case Studies

By Elisabeth Wecshler

A video highlighting case studies of diverse types of problem-solving using the NAS visualization lab is scheduled for release this summer and will be available at no charge from the NAS Documentation Center.

Four case studies are presented in the professionally produced video, according to Kevin McCabe, manager of the lab:

- Bob Meakin, formerly of the Ames Computational Aerosciences Branch and now a principal at Overset Methods Inc., used Unsteady Flow Analysis Toolkit (UFAT) and PLOT3D to create a series of ARCGraph files, showing particle paths and pressure variations. These files were animated in the visualization lab using FAST (Flow Analysis Software Toolkit); each frame was recorded on the Abekas digital video disk recorder, which enabled the simulation to be played back in real time.

  Animations of unsteady simulations require "tedious work," said McCabe, "but it's about the only way to visualize the dynamics of problems of this magnitude."

- Merritt Smith, Computational Aerosciences Branch, used the NAS virtual wind tunnel as well as stereoscopic viewing to evaluate the complex traces that result from Harrier aircraft datasets. "It's very difficult to see what's in front of what on a flat surface," McCabe pointed out. The special three dimensional screen on a lab workstation creates a more realistic environment for examining the traces.

  The virtual wind tunnel host used for Smith's work is a Silicon Graphics Inc. (SGI) 4D/440 Skywriter. The workstation used for the ster eoscopic viewing is an SGI Crimson RE, and the stereo monitor is a Tektronix model.

- S. V. "Mani" Subramanian, ASE Technologies Inc. and Dave Cherry, of General Electric Co.'s Turbo Technology Division -- both of Cincinnati -- collaborated with Nateri Madavan, of the MCAT Institute working in the Ames Computational Algorithms and Applications Branch, to create an unsteady simulation and visualization of rotor-stator interaction.

"This case underscores NAS's unique capability as well as its commitment to collaborate with industry in solving aerosciences problems," McCabe said.

- Akil Rangwalla, also of the MCAT Institute at Ames, used sonification techniques at the visualization lab to analyze CFD solutions with acoustics. Sonification is data-driven sound or sound produced from data, explained McCabe, who wrote the software enabling Rangwalla to listen to his data on a visualization lab workstation.

  "In the study of tonal acoustics, a scientist usually looks at a series of X-Y plots, which represent a filtered portion of the signal," McCabe said, adding that sonification techniques allow the scientist to comprehend the whole signal at once. The main advantages are to give an overview of the data, for verification against other techniques (such as X-Y plot techniques), for direct simulation, and the possibility for data reduction.

  "Our [human] aural system can integrate two channels of information into an understanding of an entire volume -- something not easily done by examining X-Y plot pairs, McCabe said.

  "As with virtual reality, where one is provided with visual clues, audio can add a more realistic environment for data analysis," McCabe explained, predicting that researchers increasingly will be interested in the acoustic components of their data.

  "We are looking at ways to make data components audible to enhance the understanding already provided visually," he said.

To get a copy of the video, send online request to the NAS Documentation Center.

Next Article | Contents | Main Menu | NAS Home

Next Article | Contents | Main Menu | NAS Home

UNICOS 8.0 Beta Test Results a `Pleasant Surprise'

The *UNICOS 8.0 Beta Test Report*, published in May, documents the benefits and problems that NAS experienced with the latest version of UNICOS, Cray Research Inc.'s (CRI) operating system based on UNIX. NAS undertook the UNICOS 8.0 test with the goal of exploring the benefits to NAS users by putting beta software into full service, while limiting the potential negative impact to them.

Staff in the NAS Computational Services and Systems Development branches conducted the tests on **vonneumann**, a production CRAY Y-MP C90 supercomputer. Beta software was received in early October 1993 and was placed in production on December 2. Previously, the C90 ran in production with UNICOS 7.C.2.

The primary reason for deciding to move directly to version 8.0 was its multi-threaded kernel, in which all CPUs can execute kernel code simultaneously, as opposed to only one CPU executing large areas of code at a time.

Another reason was the ability to designate primary andsecondary devices within a given filesystem, allowing maximum utilization from NAS's RAID (Redundant Array of Inexpensive Disks) device, a collection of off-the-shelf, often small disk drives with a sophisticated programmable controller, which is cheaper than traditional disk devices.

Through a cooperative effort, NAS and Cray kept administrative details from delaying tests by several weeks. CRI agreed to provide early access to UNICOS 8.0 source code and on-site staff to port NAS's local modifications, and committed to respond to reported problems quickly. For its part, NAS agreed to provide adequate dedicated test time and to allow open testing to users. It was also agreed that version 8.0 would not be placed into production service until all local modifications, primarily the NAS-developed Session Reservable File System (SRFS) software, worked. Logging many hours, CRI's staff ported SRFS and -- most importantly -- made it work with the multi-threaded kernel. The official version of UNICOS 8.0 was released on March 11.

NAS system activity information gathered before and after the beta test period shows that the C90 experienced a 61 percent reduction in system time, based on the NAS configuration of 16 CPUs in general production use (actual user workload), *not* in dedicated mode.

At NAS, the UNICOS 8.0 beta version appears to be just as stable as release 7.C.2 -- a pleasant surprise, according to the *UNICOS 8.0 Beta Test Report* -- considering the significant number of changes made to

the operating system to accommodate multi-threading.

To request a copy of the *UNICOS 8.0 Beta Test Report* by NAS staff members Bob Ciotti, Jim Craw, Alan Powers, and Phil Tam, send email to **craw@nas.nasa.gov**.

Next Article | Contents | Main Menu | NAS Home

Next Article | Contents | Main Menu | NAS Home

July - Aug 94, Vol. 2, No. 4

Executive Editor: Marisa Chancellor

Editor: Jill Dunbar

Senior Writer: Elisabeth Wechsler

Contributing Writers: Nicholas Cardo, Jean Clucas, Harry Waddell

Other Contributors: David Bailey, Pieter Buning, Karen Castagnera, Bob Ciotti, Jim Craw, Dan DePauk, Pat Elson, Jude George, Steve Harding, Vee Hirsch, Sandy Johan, Bill Kramer, Paul Kutler, Tom Lasinski, Anthony Lisotta, Kevin McCabe, George Myers, Chuck Niggley, Cathy Parks, Alan Powers, Marcia Redmond, Doug Sheppard, Steve Smith, Phil Tam, Val Watson

Editorial Board: Marisa Chancellor, Jill Dunbar, Chris Gong, Serge Polevitzky, Jim Ruppert, Pamela Walatka, Elisabeth Wechsler, Rita Williams