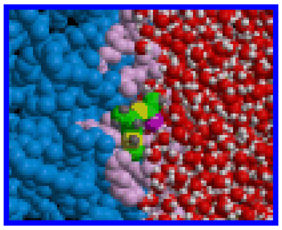 Chicken or Egg?

May - June 1995 -- Volume 2, Number 10

Chicken or Egg?

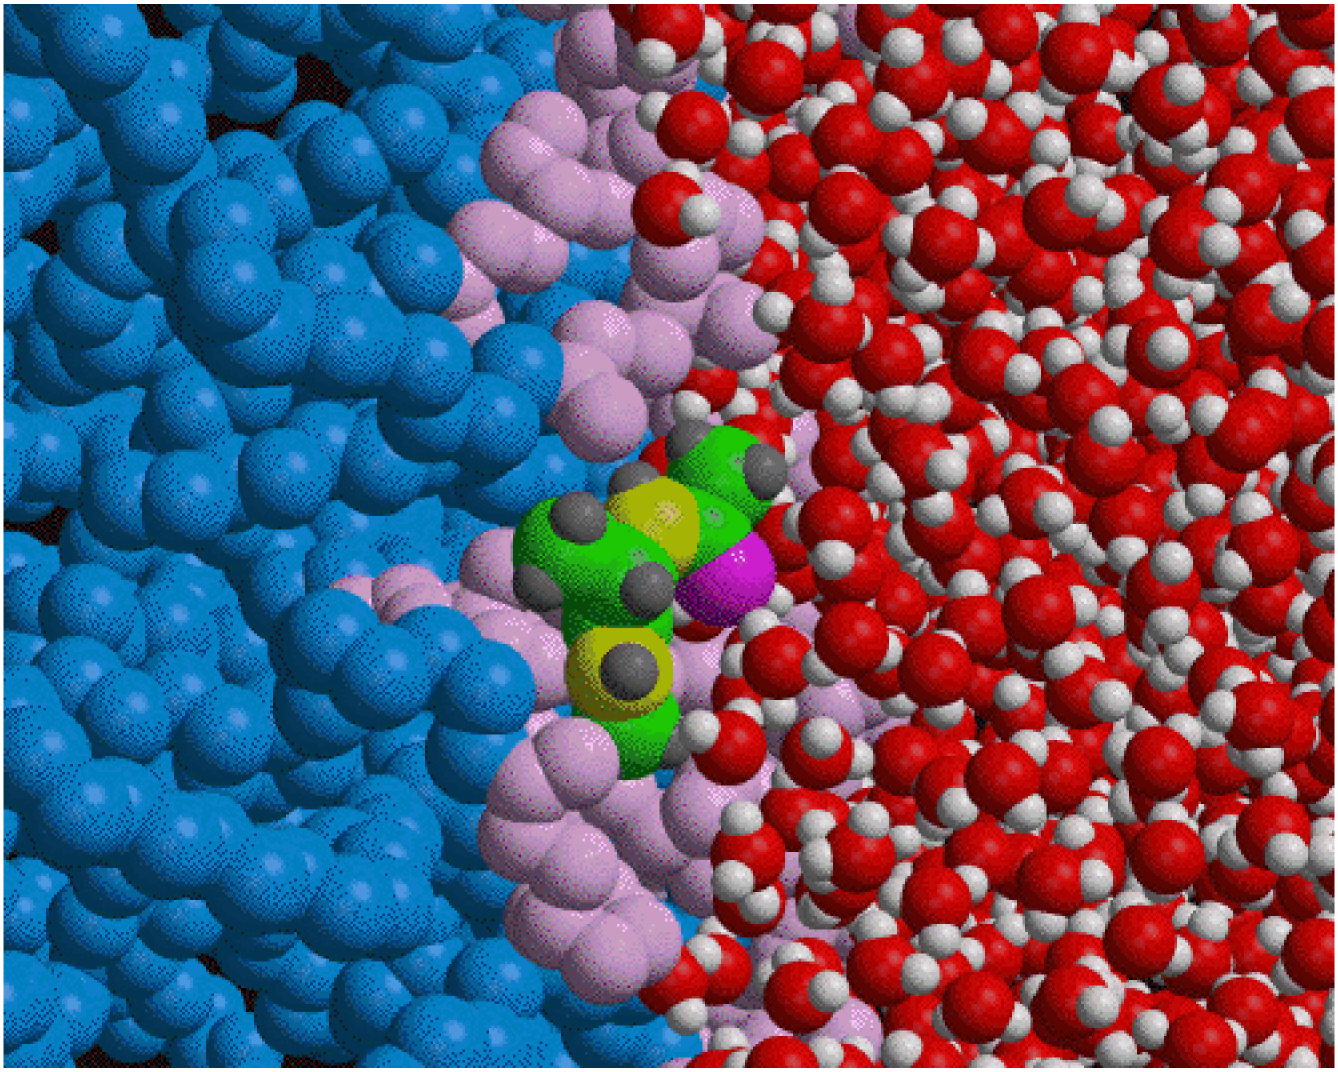

NASA Ames researchers Andrew Pohorille and Michael Wilson study the age-old question of the origin of life. They are attempting to show how simple cell like structures, called vesicles -- which some scientists believe came to earth in meteorite fragments -- might have been the first self-replicating systems in our world. Shown here is a dipeptide molecule at the boundary between membrane and water (red and white atoms). In contemporary cells, long protein chains catalyze the chemical reactions needed to sustain life. Catalysis requires the proteins to assume ordered structures (such as helices), but small peptides, such as the one illustrated here, are usually random. Pohorille and Wilson have shown that peptides can assume ordered structures if located at the interface between water and membrane. From the 1993-1994 *NAS Technical Summaries*. For a free copy, send an email request to **doc-center@nas.nasa.gov**.

Contents | Main Menu | NAS Home

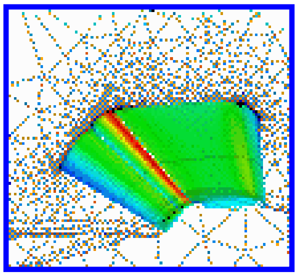

NASA's Future Aeronautics Role Debated at CAS Workshop

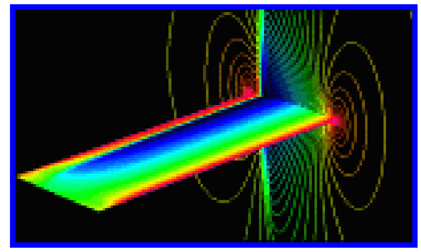

by Elisabeth Wechsler

Just as the U.S. aeronautics industry faces heightened competition from overseas, NASA's budget is being severely cut. With even fewer federal resources to draw upon in the future, the timing could hardly be worse.

At the Computational Aerosciences (CAS) Workshop, held March 7-9 at the Santa Clara (CA) Convention Center, the question of NASA's appropriate role in such an environment sparked heated discussion. Speakers and panelists from NASA headquarters, the aeronautics industry, universities, and federal laboratories, as well as audience members, offered their collective wisdom and actively debated how NASA's priorities should be set.

Wesley Harris, formerly Associate Administrator for the Office of Aeronautics at NASA headquarters, underscored a sense of urgency by noting in introductory remarks that U.S. export of airplanes and jet engines declined in 1994 -- the first time in six years. Referring to this $30 billion manufacturing sector as providing "the largest export credit to the U.S. economy," Harris said that "our way of life as a nation is truly at stake." He is currently Special Assistant to the Associate Deputy Administrator (Technical).

CAS Work `Absolutely Essential'

For the U.S. to retain its world market share, the CAS work -- funded by NASA's High Performance Computing and Communications (HPCC) Program -- is "absolutely essential," Harris emphasized. "We need affordable airplanes that are efficient." He added that this objective requires "new technologies." Harris told the audience of 200: "We need your help and participation in this," and asked for "more team play between NASA, industry, and universities."

After reviewing the program's goals, William J. Feiereisen, CAS project manager for the agency and workshop chair, told the audience: "We want to reduce the risk to industry of implementing parallel processing into the design processes by focusing on design applications that motivate aerospace computing. Two areas where we'll increase support are the Advanced Subsonic Transport and High Speed

Research programs."

Toward this end, CAS has focused more funding on networked workstation clusters, such as the Silicon Graphics Inc. testbed at the NAS Facility, and by cooperative work with aerospace companies and independent software vendors.

Short-term vs. Long-term Goals

Feiereisen invited the audience to provide feedback during the workshop, the major portion of which involved technical paper presentations. There was ambivalence as to whether NASA should fund short-term goals (for example, clusters providing production-quality systems, software, and tools) or take a more long-term approach, such as massively parallel processors (MPPs) performing basic research and developing or testing high-performance, high-risk systems.

An unidentified participant from NASA headquarters commented that "in communication with Congress, NASA walks a fine line between subsidizing U.S. industry and setting industrial policy."

"Industry can't afford to do basic research any more -- the national labs are the last hope," said Robert MacCormack, of Stanford University's Department of Aeronautics and Astronautics, in his panel session entitled, "How Can NASA Better Advance High Performance Computing?"

Another panelist, Antony Jameson of Princeton University and Intelligent Aerodynamics Inc., seemed to agree: "NASA should focus more on basic research and not take over the role that companies should be doing themselves -- like commercial software." He believes that NASA should "back out before (the research) reaches the point of producing commodity products."

R&D Budgets Under Pressure

"R&D budgets must lead to products within 12 months to survive cuts," said Ray Cosner, CFD applications manager at McDonnell Douglas, in his panel session entitled, "High-performance Cost-effective Heterogeneous Computing." "NASA should be taking a three-to-five year vision on its research -- 12 months isn't enough -- and integrating the technology and creating interface standards," he added.

Cosner sees cycle time and affordability as the key issues. "We get marching orders to cut cycle time from top management, and then we scramble to fit within the new budgetary constraints. If there's time for only one wind tunnel test, the question is, when do you run it -- at the beginning, middle, or end of the design cycle?"

"NASA's value added is pushing the envelope in very fundamental ways in terms of scalability -- storage, bandwidth, I/O -- and telling everyone," said panelist Irving Wladawsky-Berger, General Manager, IBM Power Parallel Systems.

Reducing `Time to Market'

Without federal laboratories putting systems into production for their own research needs, the time required to bring these technologies to market would be pushed out significantly, Wladawsky-Berger explained. "A huge amount of learning will be achieved" through cooperative development efforts with NASA.

"Companies need a world-class product development process to succeed," said panelist Robert Melnik, of Northrop Grumman Research and Development Center. "We need to worry more about integration and we need computer infrastructure. Numerical simulation is a key ingredient -- multidisciplinary simulation tools that are responsive, robust, reliable, accurate, and cheap."

Cost Must Be Factored

Pradeep Raj, of Lockheed Aeronautical Systems Co., also a panelist, noted serious trade-offs in producing high-quality products at affordable prices. "NASA must put `cost' in its list of metrics" for high-performance computing and also "needs to develop a better appreciation of economic and market issues," he said.

Raj would like to see more focus "on the total computing environment, not just more computing power." The hardware and system software infrastructure, application software, and people are all important, he said, but added that "NASA needs to put emphasis on the *process*, not just the components."

Another panelist, Alan Egolf, of United Technologies Research Center, said that "companies need to capitalize on all resources, not just traditional supercomputers or workstations." He'd like to see the "transparent use of disparate computers on a network" with individual workstations offering "a window to the universe," and said that NASA can help by integrating networked heterogeneous computers into a usable environment.

Egolf also thinks that NASA can help with automatic grid generation. "We spend 90 percent of a new project on grid generation. We need automatic grid generation for CFD problems," he said.

High-Risk, High-Payoff Research

Robert Ni, a panelist from Pratt & Whitney, unequivocally stated that "NASA should take on high-risk, high-return projects."

"NASA's pathfinding role in providing high-risk high-payoff research adds tremendous value," Raj said, adding that it's "natural for the agency to assume a leadership role, since everyone talks with NASA but not (necessarily) to each other."

An audience member commented that "PLOT3D or FAST wouldn't have been done by industry."

Harris cautioned the audience that "there's no premium (at NASA) to publish research. The goal is to regain world market share in commercial aviation, because if we fail in that, we won't be free to publish."

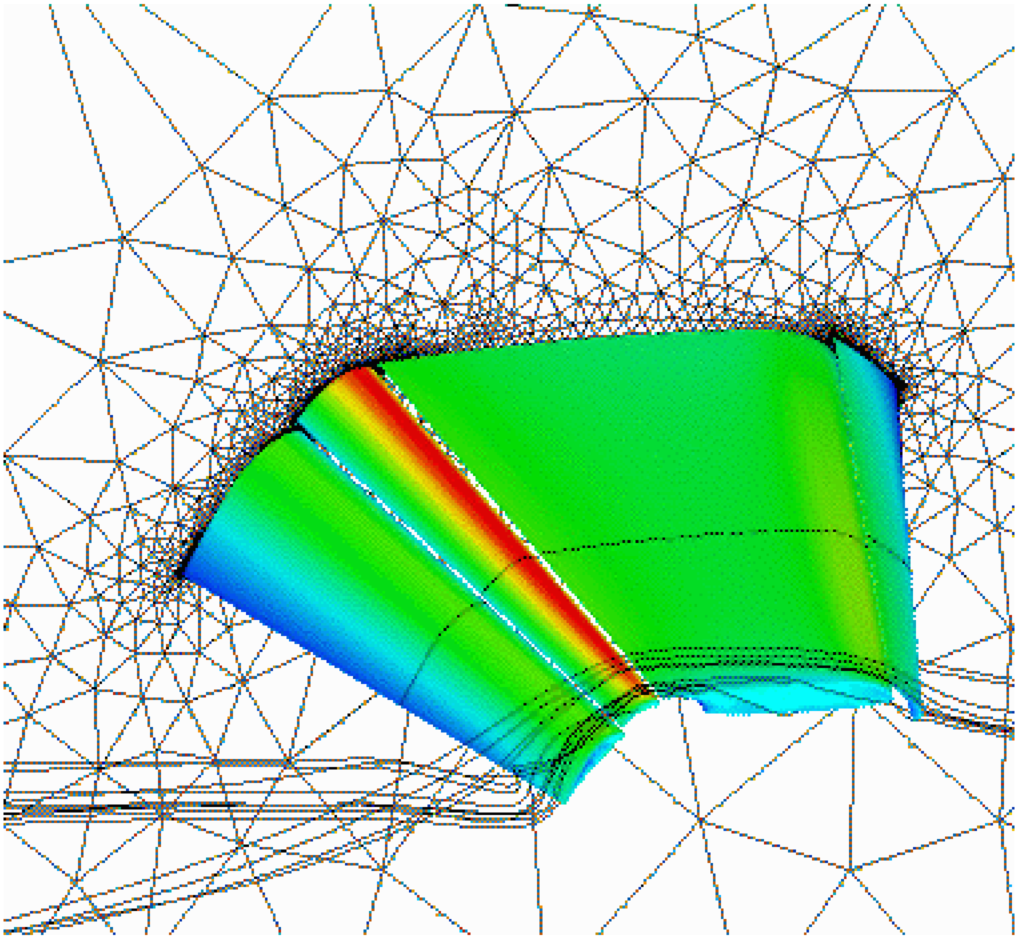

This graphic, part of a technical presentation by Tim Barth at the Computational Aerosciences Workshop held March 7-9, depicts inviscid compressible fluid flow past a multiple-component wing geometry with flaps deployed, in which the fluid flow equations are discretized and solved using an inexact Newton method. Parallel efficiency is obtained by distributing the three-dimensional unstructured mesh and flow field computation among 16 processors of the IBM SP2. The numerical solution, mesh geometry, viewing software, and further details are available online.

to the article

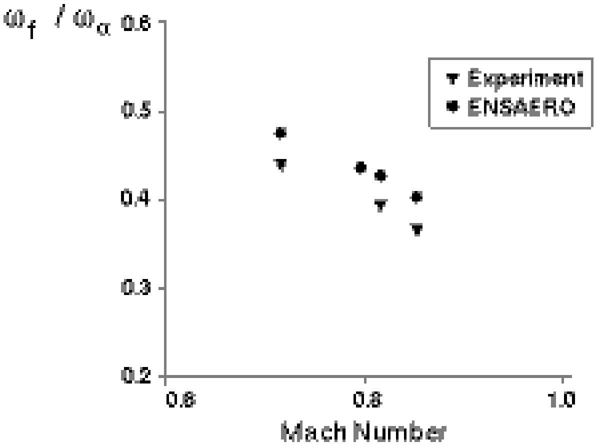

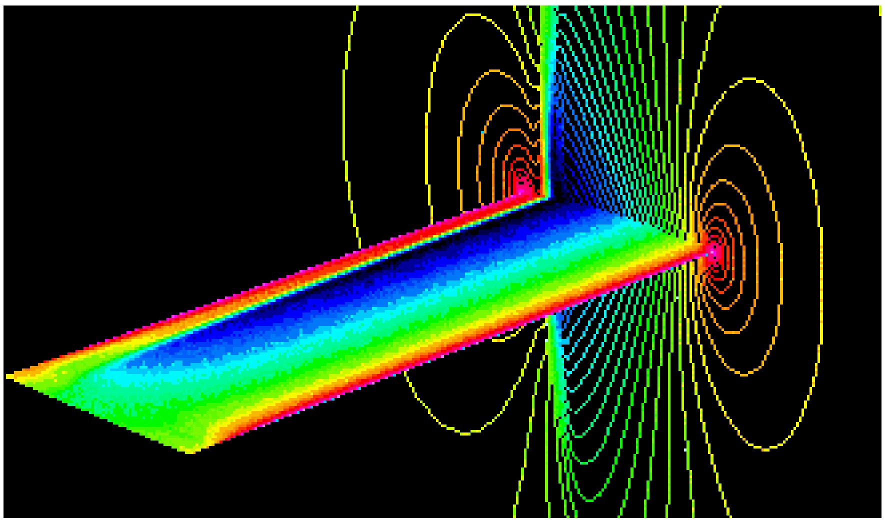

This graphic shows the results of an efficient parallel flutter-computation capability, based on the aeroelastic code ENSAERO, which was developed and first demonstrated on the NAS IBM SP2. Computational results of the flutter boundary for this rectangular wing compare well with measured data. Nine SP2 nodes gave the equivalent performance of one CRAY C90 processor. Graphic is from a technical presentation by Guru Guruswamy and Chansup Byun at the CAS Workshop.

to the article

HPCCP-funded Workstation Testbed Assesses Performance/Resource Trade-offs

by Elisabeth Wechsler

The NAS Systems Division has focused a major development effort on testing the resource potential of loosely coupled, heterogeneous workstation clusters to determine whether they are a cost-effective alternative to supercomputers. The work is funded in large part by NASA's High Performance Computing and Communications (HPCC) Program. (For information on aeronautics industry cluster experiments, see "Industry Harnesses Clustered Workstations To Squeeze Extra Cycles," *NAS News* March-April '95.)

"The NAS pathfinding role convinced NASA headquarters to consider workstation clusters as an alternative to supercomputing," said William Kramer, former NAS Computational Services Branch Chief and now a team leader for the Advanced Air Transportation Technologies Integrated Planning Team at Ames Research Center. "A loosely coupled configuration is more difficult to achieve but of great interest to our industry customers" he added.

Vision of Effort Is Defined

According to Eric Barszcz, who until recently headed the workstation cluster effort in NAS's Scientific Computing Branch, the vision of an integrated computing environment at NAS is: "To be able to log in from any workstation and have access to your files as if they were local. Also, to be able to submit a job from any workstation, giving resource constraints -- time, cost, disk space, and so forth -- and have the system figure out the best place to execute it."

To meet these goals, NAS has set up a testbed for system software integration, tasked to assemble a unified body of software in a common operating environment -- including networking, message passing, job scheduling, and file system management. The idea for the NAS testbed began two years ago, with the Distributed Computing Team activities spearheaded by Kramer. The projected sunset date for the testbed is September '97.

Four Network Technologies

The hardware consists of 16 Silicon Graphics Inc. (SGI) Power Challenge L workstations: 14 are single-processor, 75 MHz R8000 machines with 256 megabytes (MB) of memory each; one is a dual-processor, 75 MHz R8000 machine with 640 MB of memory; the last is a 4-processor, 75 MHz R8000 machine with 1280 MB of memory. All machines have 4 gigabytes (GB) of local disk and are connected with ATM OC3, Ethernet, FDDI, and HiPPI network technologies.

Barszcz hopes to persuade developers of PVM (Parallel Virtual Machine), MPI (Message Passing Interface), HPF (High Performance Fortran), and job scheduling software to define a common application programming interface between message passing libraries, parallel languages, and job schedulers.

On December 1, NAS opened the SGI testbed to a few outside users, who have been evaluated on a case-by-case basis, Barszcz said, emphasizing that the system is not set up for large-scale production.

Focus on Propulsion Applications

One project that will use the HPCCP--NAS testbed is the Cooperative Agreement Notice (CAN), managed by Lewis Research Center (LeRC). The CAN is funded through the Affordable High Performance Computing Project of the HPCC Computational Aerosciences (CAS) program.

The objective of the CAN, which is targeted for commercial propulsion vendors, is "to accelerate the development of affordable high-performance computing utilizing networked engineering workstations." The project will focus on "propulsion applications that guide the research development of system software technology directed at affordable and reliable distributed computing," according to the executive summary of the CAN's request for proposal. (The deadline for proposal submissions was January 6, and awards -- totaling $4 million -- were expected to be announced in April.)

Reduce Supercomputing Costs to 25 percent

The proposed CAN research "must meet or exceed" the following HPCC milestone, according to Barszcz: "By September '96, demonstrate cost-effective high-performance computing at performance and reliability levels equivalent to 1994 vector supercomputers at 25 percent of the cost."

To make the cluster heterogeneous, NAS plans to add two IBM and two Hewlett Packard workstations to its testbed later this year. The original configuration specified eight computers, but budgetary restrictions reduced this number to four, Barszcz said.

Other NASA centers have different hardware: Langley Research Center (LaRC) uses DEC Alpha, Sun Sparc 10s (or higher), and SGI R4400 work-stations, while LeRC uses primarily IBM 560 and 590 workstations, Barszcz said.

"The goal is to be able to use the same basic software, whether it's run from Ames, LaRC, or LeRC, by the end of the project," Barszcz said. "The most hopeful achievement is that you won't know you're running the job on a cluster. This is what some people call transparent computing."

Solutions Applicable to MPPs

"Generally, we'd expect to have more difficulties with workstation clusters than with massively parallel processors (MPPs), so if you solve those problems on clusters, the solutions are applicable to MPPs," he said. Among the limitations for workstation clusters from the user's point of view, according to Barszcz, are:

- no high-performance file system
- higher latencies, which mean that fewer message-passing algorithms run effectively
- no nonproprietary scientific libraries
- minimal debugging capability
- current system administration tools

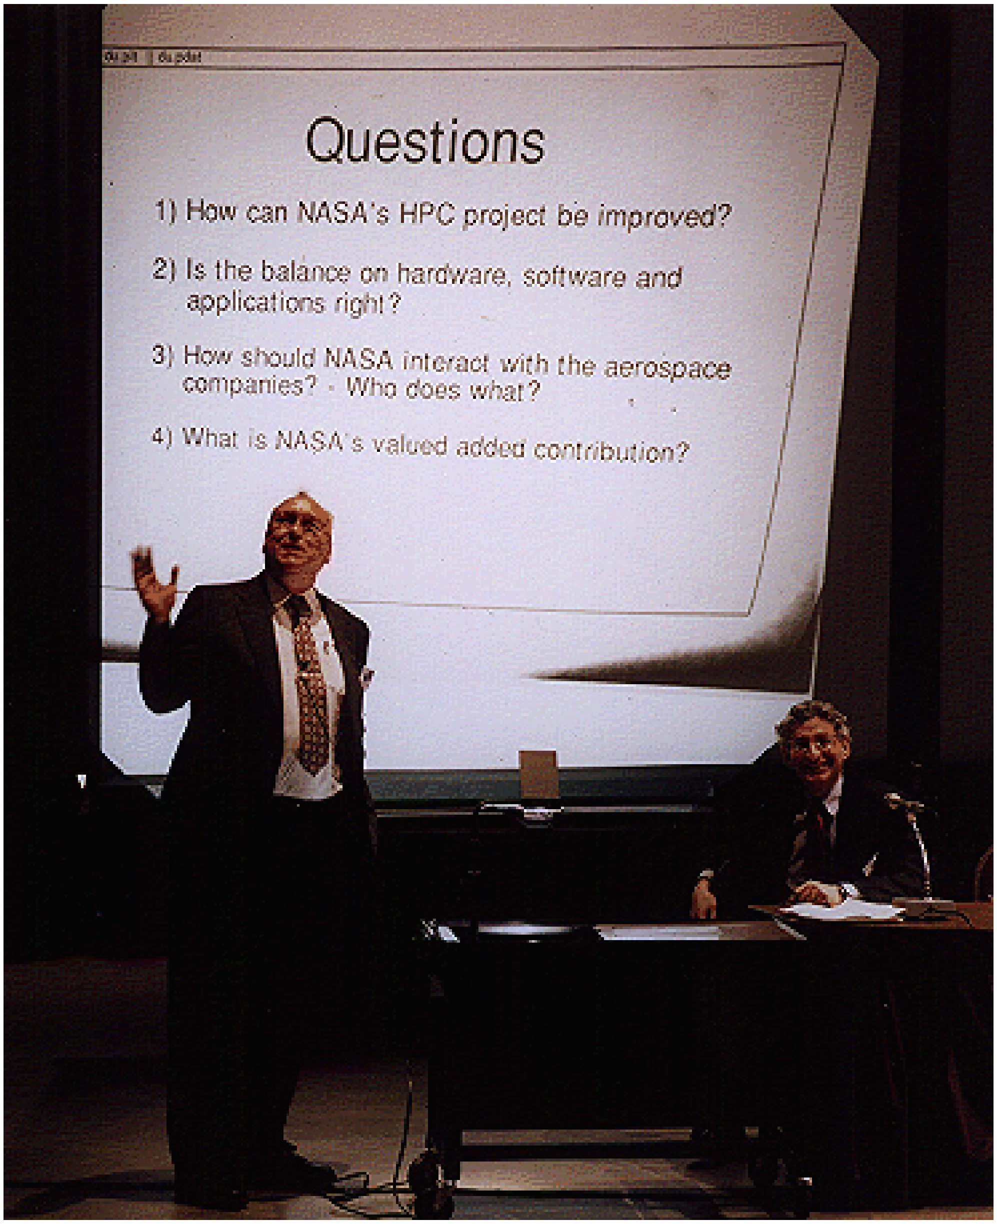

Robert MacCormack, of Stanford University, introduces questions for Computational Aerosciences Workshop panelists on March 7. Irving Wladawsky-Berger (IBM) is pictured at right. Also on the panel were Wen Huei-Jou (Boeing), Robert Melnick (Northrop Grumman), and Pradeep Raj (Lockheed).

to the article

PBS Offers Flexible Job Scheduler for the IBM SP2

by Elisabeth Wechsler

Release 1.0.4 of Portable Batch System (PBS), developed specifically for the 160-node IBM SP2 at NAS, was put into production in March. It features a parallel job scheduler initially developed by NAS and Lawrence Livermore National Laboratory (LLNL). The ability to implement scheduling policies within the Cooperative Research Agreement (CRA) consortium using the SP2 was a "top priority," said Bob Henderson, PBS group lead in the NAS Scientific Computing Branch.

Includes BASL and Tcl Languages

"PBS offers the flexibility to schedule jobs as the CRA researchers see fit," explained Dave Tweten, manager of the PBS and usage accounting groups. The first beta version of the scheduler framework was extended to include two languages for the SP2 version of PBS: Batch Scheduling Language (BASL), an easier-to-use but less flexible high-level language that describes how to schedule jobs, and Tool Control Language (Tcl), a more complicated but also more flexible format, used to specify scheduling algorithms.

NAS contributed the concept and high-level design for the BASL language, and Clark Streeter of National Energy Research Supercomputer Center at LLNL wrote and implemented the detailed design specification. BASL is well-suited for a site with a fairly standard scheduling policy, and "where they don't want to do anything oddball," Tweten said.

Tcl, developed by John Ousterhout at the University of California, Berkeley, provides the power to address more unusual requirements -- for example, said Tweten, the ability to look at user bulletins and figure out when the machine will be down for dedicated time, noting that "this is not a traditional part of a scheduler's job."

Big Improvement Over NQS

One advantage of PBS over other batch systems is that it lets users start and terminate jobs from their workstations. For system administrators, PBS offers "incredibly increased flexibility" especially in job scheduling, as compared with Network Queuing System (NQS) -- which is now about ten years old, Tweten explained.

PBS 1.1, adapted for the CRAY C90, is expected to be ready later this spring. A beta testing program for this version, begun last August, is operating at nine sites so far, including Cornell University Theory Center, Cray Research Inc., Geli Engineering (Sunnyvale, CA), Hughes Information Technology Co., IBM (Kingston, NY), Pavilion Technologies Inc. (Austin), Purdue University, University of Maryland, and University of New Hampshire. In addition, a minor release of PBS -- to handle interactive jobs -- was expected in late March.

Supporting parallel jobs for UNIX-based cluster computing is a future goal of the NAS PBS team, according to Tweten. Currently, the SP2 is treated as single system. Release 1.2 is being developed to "run on heterogeneous clusters -- all the workstations at NAS" -- and will include protocol improvement, distributed job management, interactive batch support, and a graphical user interface, he said.

IEEE Standard Superset

The guiding development principles for PBS include: avoiding proprietary code, while ensuring modularity of design; introduction of generic resources; command and graphical user interfaces; support for asynchronous communication; a wide range of scheduling algorithms; and the creation of a superset of the IEEE standard for batch queuing extensions, Tweten said.

Bernard Traversat, a member of the NAS parallel systems group that conducted the evaluation and configuration of PBS, commented that "PBS has been a really useful and reliable tool to add to our SP2 to more efficiently manage its resources."

Still, he feels that there is one important issue to be resolved by PBS: enforcing scheduling policies, a system management issue. "You want to be able to ensure a fair sharing of resources among users. Being able to reliably kill any job that exceeds its policy limits and restricting access to resources (such as direct login on a node) is very difficult to do on the SP2." (*A new NQS scheduler for the CRAY C90 does just this. See "How to Manage Resources With the New CRAY C90 Batch Scheduler."*)

Wants Resource Management Under PBS

Traversat would like to see a standard interface for PBS that interacts with parallel programming tools such as MPI (Message Passing Interface) and PVM (Parallel Virtual Machine) to put resource management for all processes of a parallel application under PBS control. Such an interface would ideally also support dynamic process management (such as the user's ability to add or release nodes, as needed, during computations), he said.

For more information about PBS, contact Tweten at **tweten@nas.nasa.gov**, or see the new PBS home page. For more information about the NAS-LLNL partnership, see *NAS News* March-April '94.

Next Article | Contents | Main Menu | NAS Home

How to Manage System Resources with the New CRAY C90 Batch Scheduler

by George B. Myers

A new batch scheduler for the NAS CRAY C90 was put in place on March 6. The motivation for the new scheduling scheme, developed by Nick Cardo of the high speed processor (HSP) group, was threefold: to give users maximum utilization of system resources, to minimize user impact on those resources, and at the same time ensure a fair distribution of resources among active users.

How the Scheduler Works

To accomplish these goals, the scheduler uses the following scheme: A batch job goes through three states during it's life in the system -- pending, queued, and running. Any user can have an unlimited number of jobs in the pending state.

A job enters the queued state if none of the user limits are exceeded.The user limits for the "production complex" (batch and mtask queues) are: five jobs queued, 320 megawords (MW) of memory, and 12 hours of CPU time. (Entries in the debug and mtasklg queues are not counted against these limits.) If a pending job exceeds one of the limits, it remains pending. The status field of the **qstat** command identifies the reason that a job is held in the pending state:

> **QcJ** - Exceeds maximum number of eligible jobs
> **QcM -** Exceeds maximum amount of memory
> **QcT -** Exceeds maximum amount of CPU time

A global set of limits is used to determine entry into the running state. The scheduler uses a "best-fit highest-priority" scheme to select an eligible job from the queued state. Priority is based on age in the queue. When resources become available, the job with the highest priority and that fits within the resources available at that time is selected to run. Global limits are set on a queue by queue basis. Currently, only the production complex is managed this way; however, this queue handles all job sizes up to 256 MW of memory and 28,800 CPU seconds (eight hours). The default job size is 4 MW of memory and 300 CPU seconds. (Users must now specify these parameters, which formerly defaulted to the queue maximum.) If no resource requirements are given, they will be set to the default.

The scheduler can more accurately manage system resources if it knows the job's requirements in advance. A feature in the scheduler assists users in determining what resources are needed for a job. At job termination, a message that lists resource usage and resource requests is appended to the NQS job's standard output file.

User Input Must Be Accurate

To maximize usage and assist the scheduler in accurately managing system resources, users must accurately specify resource requirements for each job. Giving accurate amounts of memory, time, and other required resources -- such as Session Reservable File System (SRFS) storage space -- will increase the scheduler's accuracy and improve a job's chance of running. Requesting more system resources than are needed will keep a queued job waiting longer, and once it's running, the "excess" resources won't be available until the job finishes. The impact of this is twofold: first, other jobs -- your own and others' -- may be blocked from running while waiting for resources; second, one of your other jobs may be unable to enter the queued state because it would exceed the user limits. The overall effect is reduced job throughput for everyone.

The scheduler provides accurate information about each job's resource requirements. Take advantage of this information by using the flags provided by the *qsub* command to specify all resources.

For example:

> **lM** - Memory requirements for the job (request)
> **lT** - CPU time requirements for the job
> **lr** - Specify SRFS requirements

These flags can be typed either on the command line when submitting the job or within the job script using the "#QSUB -l..." format. See the man page for **qsub** for details on these flags.

Although it isn't necessary to specify resource requirements for the $BIGDIR filesystem, do it anyway for batch jobs; otherwise, if space needed is not available while the job is executing, the job will abort.

`Manage Your Resources'

Another resource users should manage is mass storage. NAS's "superhome" file system gives each user 10 gigabytes of storage. In addition, each user has an account on NAStore for long-term storage. Keep current files in your home file system and move older files to NAStore. Files larger than 1 MB are candidates for data migration. To speed up a job, execute the **dmget** command at the beginning of a job script for all files likely to be migrated.

A command called **qorder** allows users to manage the order in which their jobs appear in the queue. See the man page for **qorder** for more details.

The message here is "manage your resources." Doing so will enhance the effectiveness of the scheduler to manage limited system resources for everyone.

More information on the batch scheduler is available from the NAS HSP home page.

Next Article | Contents | Main Menu | NAS Home

NAS's Experience with Early Versions of High Performance Fortran

by Jeffrey C. Becker

Because High Performance Fortran (HPF) has great potential to simplify programming of parallel machines -- which will allow scientists to concentrate on problem-solving, rather than get bogged down in low-level details -- a group of NAS technical staff and scientists has been involved in an effort to acquire and evaluate compilers from different vendors since December 1994. Focusing on compilers for the IBM SP2 and Silicon Graphics Inc.'s Power Challenge Cluster, the team has also begun porting codes, with initial positive results.

Another important reason for this work was to allow CM Fortran users to start porting code to other platforms before the Thinking Machines Inc. (TMC) CM-5 was decommissioned in March.

HPF at a Glance

HPF has been evolving since 1992 as a superset of the Fortran 90 standard, with extensions for data distribution and data parallel constructs (such as FORALL). Version 1.1 of the HPF Specification was completed in November 1994 by the HPF Forum. For background on HPF features, see *NAS News*November December '93 and January-February '94.

Among the vendors who have recently announced HPF products are: Digital Equipment Corp. (DEC), Hitachi, Intel Computer Corp., Meiko World Inc., Motorola Inc., and NEC Corp.

For the most part, these vendors have implemented Subset HPF, a subset of both Fortran 90 and HPF 1.x, which is designed to encourage early HPF compiler releases. Features contained in Subset HPF include:

- static data-mapping features (ALIGN, DISTRIBUTE, TEMPLATE, PROCESSORS directives)
- the single-statement FORALL
- FORTRAN 77 standard conforming features, except for storage (COMMON, EQUIVALENCE) and sequence association (argument passing)

Subset HPF does *not* include:

- dynamic mapping features (REALIGN, REDISTRIBUTE, DYNAMIC directives)
- the (multi-statement) FORALL construct
- some FORTRAN 90 features, including CASE statement, derived types, modules, and some I/O extensions

Early Experience Positive

So far, NAS has installed two compilers: "pghpf," developed by The Portland Group Inc. (PGI), has been available on the SGI workstation cluster since December (the SP2 version is currently being tested by NAS staff); and Applied Parallel Research Inc's "xhpf" has been available on both the cluster and the SP2 since January. A beta copy of IBM's HPF for the SP2 will be installed soon.

NAS's early experience with these compilers has been positive. Several local users have been porting codes and experimenting with HPF. Despite the fact that they haven't been getting very good performance with these early versions -- compared to message passing Fortran implementations -- some NAS benchmarks and other scientific benchmarks scale well on both the SGI cluster and the SP2. The table above shows performance of the APR compiler on the SP2 compared to the Cray Research Inc. T3D and the Intel Paragon.

HPF has the potential to make scientific programming easier and to perform well. Not only does it support data parallelism, but it allows the programmer to specify data placement, which can minimize communication time. In addition, HPF's extrinsic feature allows the programmer to drop into message-passing code (or code in a different language, for that matter) if needed.

Interest in Scientific Math Libraries

In addition to giving feedback on performance problems, which the vendors are addressing, the NAS team lets vendors know about new features that would be useful. Early NAS users have expressed an interest in scientific math libraries, especially linear solver and transpose routines that exist in TMC's CMSSL library. The NAS team is considering a development project on a CMSSL-like library.

NAS is also considering the DEC HPF compiler for the DEC Alpha running OSF/1.

Online Information Available

Information on how to use the PGI and APR compilers is available online. Select the Cluster or SP2 icon for details. Two useful publications are *The High Performance Handbook* and (since HPF is Fortran 90-based) *Fortran 90 Explained*. The HPF 1.1 specification document is available from the NAS Documentation Center.

For more information on HPF, contact Bill Saphir at **wcs@nas.nasa.gov**.

Jeffrey C. Becker, formerly a member of the NAS parallel systems group, has left NAS. Others involved in HPF work include Subhash Saini and Bill Saphir (of NAS's Algorithms, Architectures and Applications Group), and Rob Schreiber (Research Institute for Advanced Computer Science), who has been a member of the HPF Forum since its inception.

Next Article Contents Main Menu
NAS Home

Visualization Research Grants Awarded for FY95

by Jill Dunbar

Four innovative visualization projects have been selected to receive grants sponsored by the NAS Systems Division through an FY95 NASA Research Announcement (NRA). Two more projects were chosen to continue research funded last year, including Brown University's work on user interface techniques.

Virtual Workbench for Scientists

Advancing the idea that virtual reality technologies should adapt to specific human work environments, Stanford University is building the next generation Responsive Workbench, a virtual environment for scientists and engineers. The workbench is a virtual space in which a scientist can interact with life-size three-dimensional objects, using a 3D input device and a glove, and simultaneously view an actual physical work area, as well as other objects and colleagues.

Led by principal investigator Patrick Hanrahan, the Stanford team will improve software tools and develop algorithms for NAS's Virtual Windtunnel. They will also port the windtunnel system to the workbench and create interactive techniques tailored specifically to tasks performed by scientists and engineers using Virtual Windtunnel software.

Focus on Tools, Methods

To address users' needs for high-level visualization tools and methods for handling large datasets, Advanced Visual Systems (AVS), Waltham, MA, will expand AVS/Express, its visually-based software development tool, to encompass building distributed visualization software. Gary Oberbrunner will lead this project; in its first phase, a transport-independent message-passing and communication layer and distributed high-performance visualization modules will be developed.

Researchers at the University of California, Santa Cruz, will develop visualization tools for multidimensional datasets, focusing on those arising from time-varying computational fluid dynamics (CFD). In previous years, principal investigators Jane Wilhelm and Allen Van Gelder received grants to create volume visualization techniques for CFD.

Arizona State University's Gregory M. Nielson will develop new methods for analyzing and visualizing flow fields. Specific research includes multiresolution for very large, time-dependent, curvilinear grids.

FY94 Projects Continue

In addition to Brown University's continuation project, Stanford University principal investigator Lambertus Hesselink will complete the last of a three year effort in tensor field visualization. Researchers will animate tensor texture and tensor topological skeletons. See *NAS News*, March-April, 1995 for more details.

For more information on NAS visualization grants send email to **vaziri@nas.nasa.gov**.

Next Article | Contents | Main Menu | NAS Home

NAS Sponsors Research in 3D, Direct-Manipulation User Interfaces

by Kristina D. Miceli

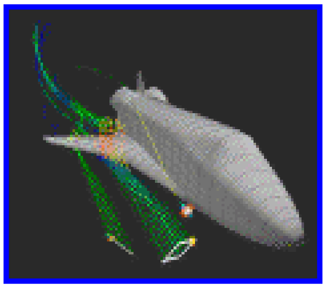

Researchers at Brown University are designing innovative user interface techniques for the intuitive exploration of three-dimensional (3D) computational fluid dynamics (CFD) datasets. This research, led by principal investigator Andries van Dam, focuses on 3D direct-manipulation techniques for interactively controlling visualization tools using conventional 2D hardware. The work is sponsored by a NAS visualization grant.

3D Widgets Suited for CFD

Many systems available today use 2D widgets and/or text input for investigating 3D datasets, but these methods can be difficult and tedious. For example, a scientist may be interested in exploring streamlines in a velocity vector field. Most current systems require the scientist to enter seed locations for the streamlines via text or a slider. This process is often done by trial and error and can be frustrating and time consuming. For intuitive exploration, a more appropriate choice would be 3D widgets that can be placed directly on the rendered scene by the scientist.

Positioning Widgets is Necessary

A necessary first step for exploring a dataset using direct manipulation is the positioning of a probe (widget). While this may sound simple, many complicating factors arise. For example, when manipulating a 3D dataset through a 2D computer screen, it is often difficult to place a probe due to a lack of depth information. "Shadow" widgets attempt to solve this problem by projecting shadows of themselves in the scene, which provides helpful depth cues.

Another difficulty arises when scientists lose track of their position in the 3D Cartesian space. To help solve this problem, the Brown team designed "object handles" to maintain widget orientation. These handles are line segments that orient themselves to the coordinate axes for reference. When dragged, an object handle translates an object along its associated coordinate axis. The team also created "grid-aligned handles," which are similar to object handles but are aligned to the curvilinear grids found in CFD simulations. When a grid-aligned handle is dragged, the selected object is forced to move along its associated grid line, allowing scientists to translate objects along complex surfaces whose geometry is reflected in the computational grid.

Probes Control Sample Points

The research team also designed generalized probes to control collections of sample points, providing information necessary to compute visualization techniques such as streamlines, rakes, hedgehogs, cutting-planes, and isosurfaces.

Generalized probes can be zero-, one-, two-, or three-dimensional, depending on the amount of information the scientist wants to view. Zero-dimensional (0D) probes provide information about a single point in the dataset and can be the starting location for a streamline. One-dimensional rake probes, similar to the smoke rakes seen in wind tunnels, define a set of sample points at regular intervals along a line. Visualization techniques such as streamlines can be controlled using the rake widget. The 2D plane widget controls a set of points arranged in a regular planar grid. These points provide the information to compute visualization techniques such as with streamlines and cutting planes. In some cases, such as streamlines, scientists can adjust the concentration of points to avoid visual clutter.

The 3D volume probe generates a volume of sample points. A common technique controlled by the volume probe is a set of cutting planes that can be moved as a unit. These widgets can all be translated and rotated by grabbing them and moving accordingly.

More information about the Brown Research group is available online.

To contact van Dam about this research, send email to: **avd@cs.brown.edu**.

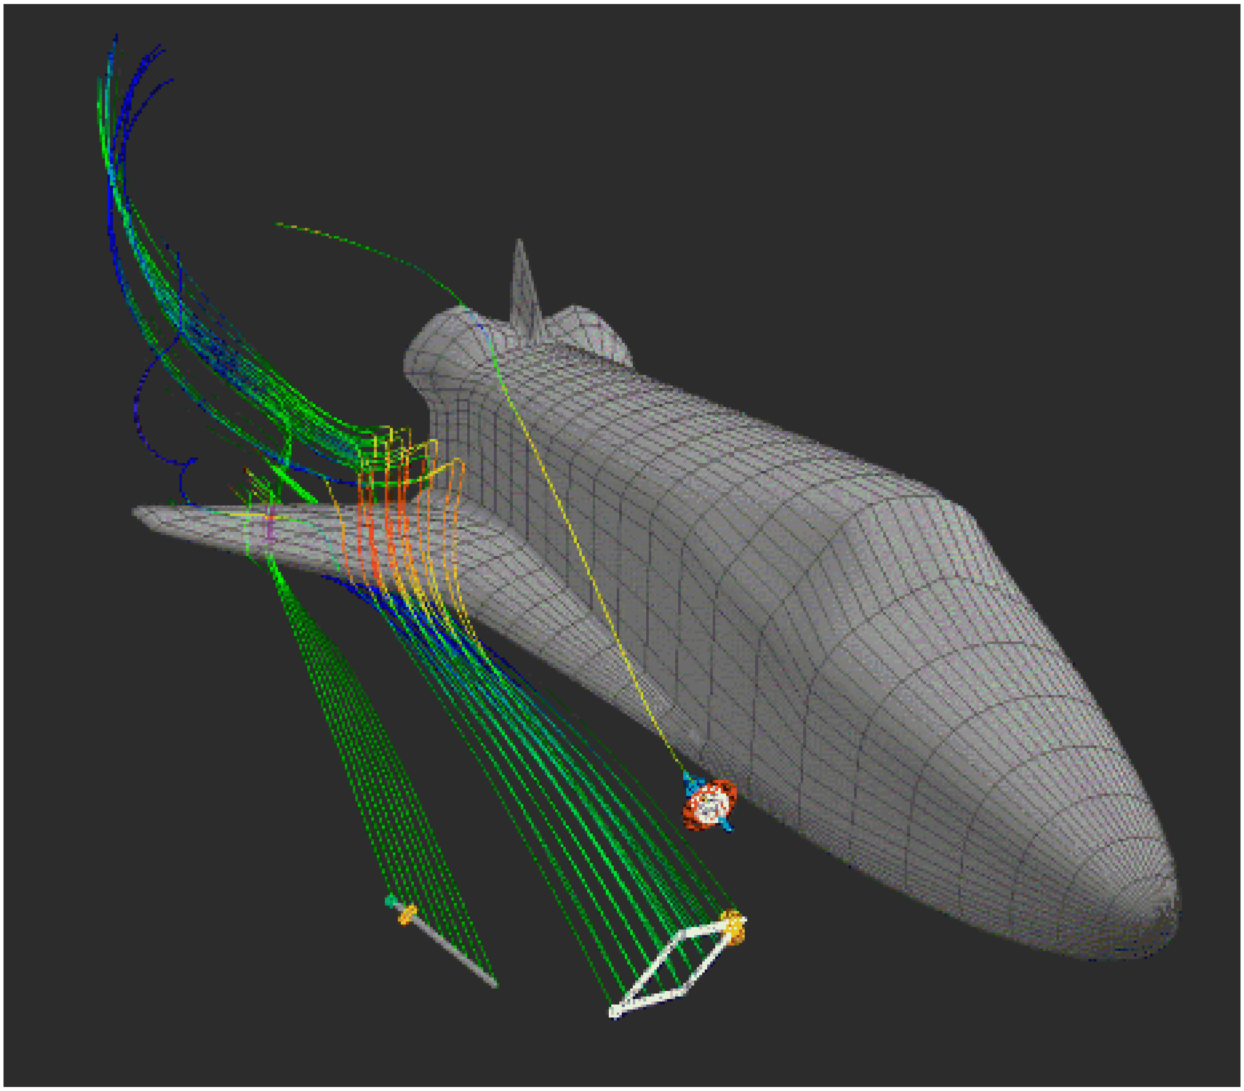

Zero-, one-, and two-dimensional probes are shown displaying streamlines about the wing of the space shuttle. The 0D probe defines a single location in the dataset, which is used as the starting point for a streamline. The 1D probe is a rake of streamlines containing several points sampled at discrete locations along the rake. The 2D probe defines a rectangular set of points. Users can control the number of points and spacing between them.

to the article

Users Save Time With UFAT 3.0's Built-in Scripting

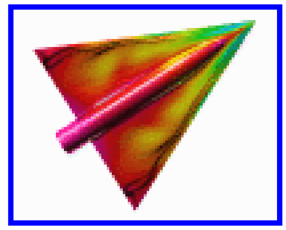

by Jill Dunbar

Thanks to a built-in scripting capability in the latest version of NAS's Unsteady Flow Analysis Toolkit (UFAT), users can save significant time in applying time-independent flow visualization techniques to their unsteady CFD flow solutions. And in tests, a new algorithm produced "dramatic improvement" in particle-tracing speeds, according to David Kenwright, who worked with UFAT developer David Lane on several design features. The upgrade is scheduled for release this sometime this month.

Lane listened carefully to users when designing this latest release. Neal Chaderjian, research scientist at NASA Ames Research Center's Applied Computational Aerodynamics Branch, said that Lane has "worked with me and others -- he's very responsive. There's a nice synergism between (us), which results in a much more useful tool."

New features include:

- automatic color contouring of surfaces
- specifying physical seed locations
- particle trace subsetting
- generating stream ribbons and surfaces, and grid surfaces
- surface vector (tuft) plotting

Color contouring and surface vector plotting, in particular, save users significant time. Without these features, users would have to write custom scripts, then apply them to each timestep, which typically consists of 200-1,000 unsteady flow files. Once the PLOT3D scalar functions (such as temperature or pressure) are computed, UFAT assigns values to each grid node to generate color contour surfaces. PLOT3D, an interactive program used by scientists to visualize CFD data, was developed at Ames.

Chaderjian, who is working with the U.S. Air Force to advance the understanding of flow physics on a delta wing, said that the ability to render grid surfaces colored by pressure is especially useful. "Before (this version of UFAT), I used `brute force' to write a FAST shell script to generate surface pressures from 600 files, which took two full days," compared to a similar job with 1,000 files, which took less than an hour using UFAT 3.0.

With the physical seed feature, users can release particles from locations in the physical grid, and UFAT automatically finds the corresponding computational grid location. Users are no longer limited to releasing particles exactly at grid points -- now, they can specify fractional locations from those points. As a result of this fine-grained control, users can better detect flow features.

For more information, send email to **lane@nas.nasa.gov** or **davidk@nas.nasa.gov**. Lane and Kenwright are part of NAS's visualization algorithms and tools group, and have begun designing a UFAT home page on the World Wide Web.

Next Article | Contents | Main Menu | NAS Home

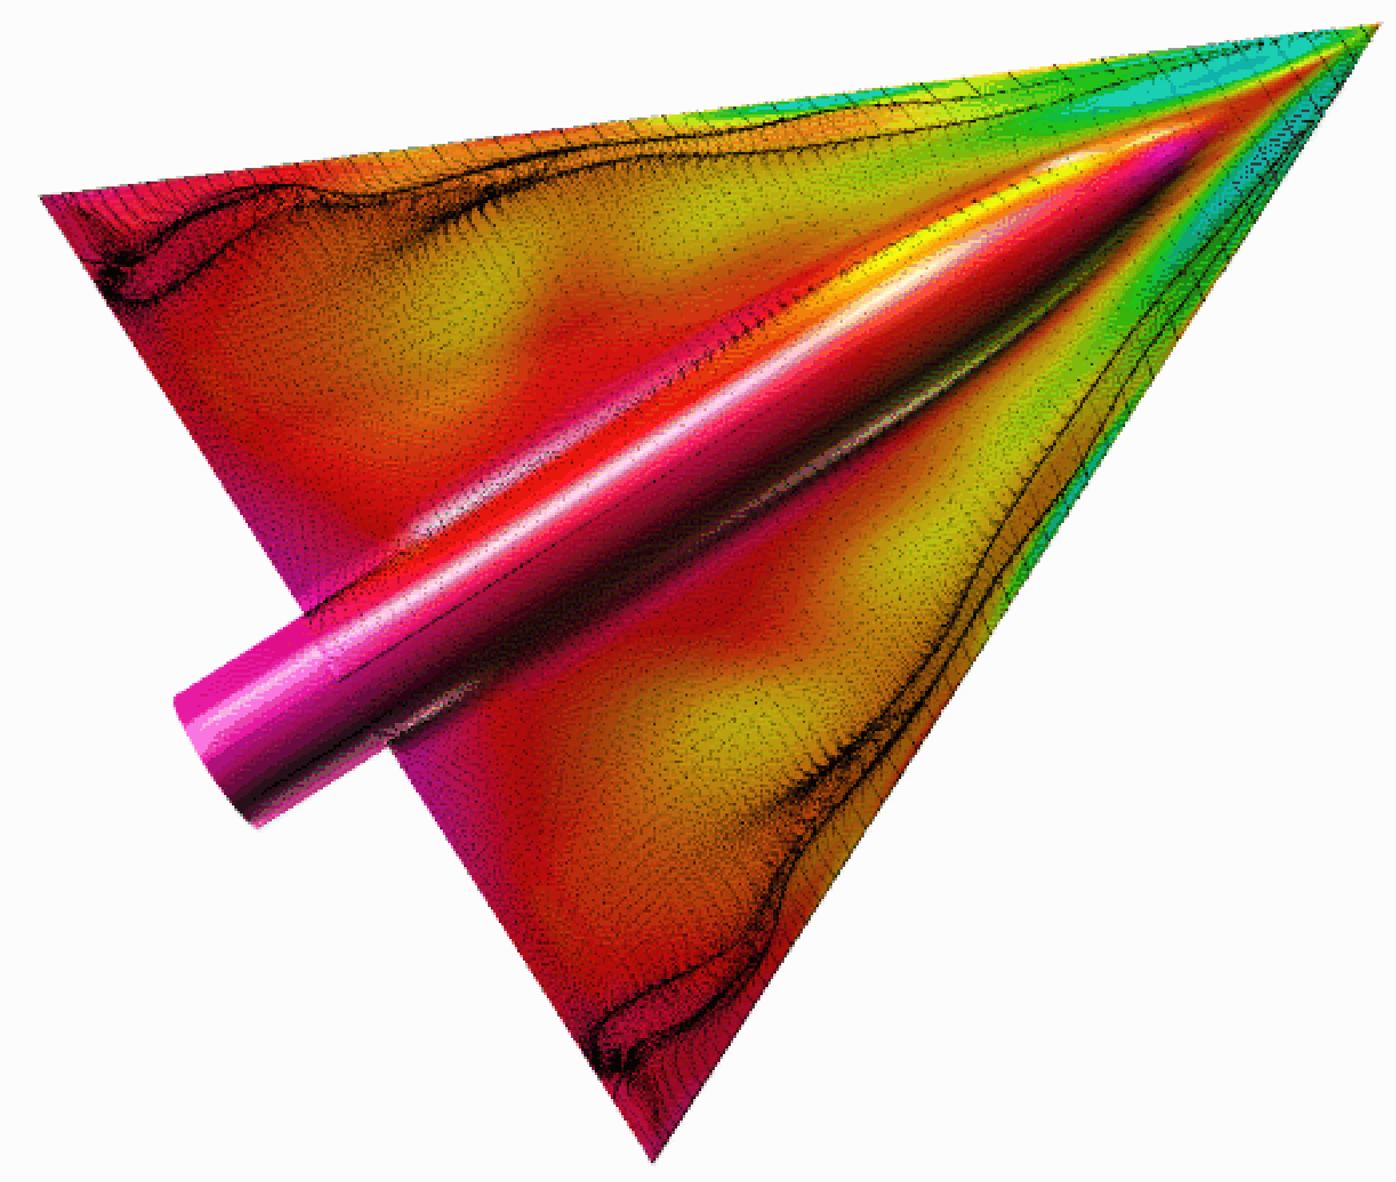

Top view of an unsteady flow surface for a 65-degree sweep delta wing at 30 degrees angle of attack. Particles are restricted to the delta wing surface and were computed on a Convex C3240 using UFAT 3.0. The grid surface, generated by UFAT, is colored by pressure (low = blue, high = magenta).

Flow computation of the delta wing by Neal Chaderjian and Lewis Schiff. Image by Chris Gong and David Lane.

to the article.

CRA-SP2 Partners Give Research Updates at NAS Meeting

by Elisabeth Wechsler

Representatives from the eight Cooperative Research Agreement (CRA) partners met at the NAS Facility on March 6 to give semiannual reports on their IBM SP2 system development projects. David Bailey, who coordinates the CRA research project at NAS, and Richie Jacobovits, IBM program manager for the CRA, chaired the all-day meeting.

Several topics prompted debate among participants, including system usability, job schedulers and policies, making "legacy" code parallel, job turnaround time, scalability results, solving aeroelastic problems, increasing memory, and achieving portability.

David Serafini, of Rice University, believes that "portability is as important as scalability." Others seemed to agree.

"For a production environment, we'd like to see the ability to dynamically allocate nodes," said David Young, Boeing Computer Services, Seattle.

Parallel vs. Vector

When Jacobovits asked the audience if parallel systems can provide sustained performance equivalent to vector machines, a lively discussion ensued.

"It appears on cost benchmarks that the SP2 has achieved the same results as a CRAY C90," said Mark Shephard, of Rensselaer Polytechnic Institute, Troy, NY.

Toby Harness, of the parallel systems group at NAS, commented that "with the cost component factored in, the answer may indeed be `yes.'"

"Support costs are still going to be significant with parallel systems, even if the hardware becomes much cheaper," said Mike Eldredge, of Centric Engineering Systems Inc., Santa Clara, CA.

Intelligent Aerodynamics and Lockheed were also represented. The next meeting is planned to coincide with Supercomputing '95 in December.

Toby Harness Pushes the Envelope -- With Harmony and Precision -- in All Pursuits

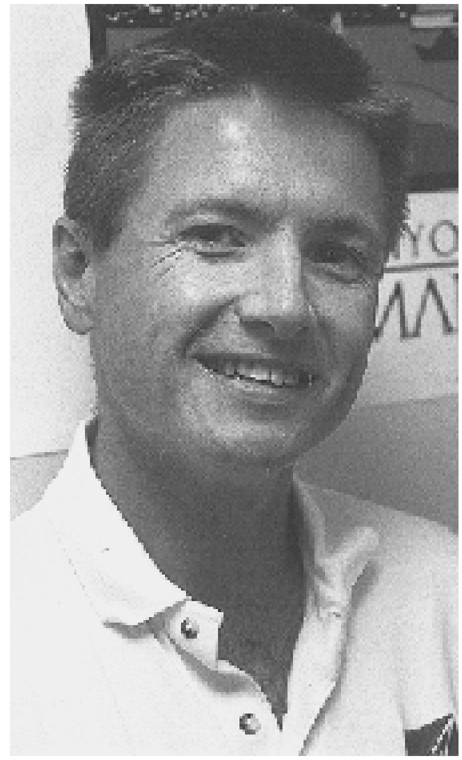

by Elisabeth Wechsler

Testing limits and not flinching -- whether it's skydiving, finishing a marathon, or making system software deliver parallel computing performance -- seems to be the underlying theme in Toby Harness's life. He also appreciates the precision work required in fine-tuning a process or system.

As manager of the NAS parallel systems group for two years, Harness sees his role as that of conductor, helping all parts of the process work in harmony by setting "some priorities on particular projects, then stepping back and letting the group do it." He described the 13 civil servants and contractors in his group as "self actuating -- both as individuals and as a team." Harness enjoys the challenge of helping the group select which projects to work on, as well as "monitoring projects that we can't do ourselves."

(Editor's note: At press time, Harness decided to focus on the technical aspects of parallel systems development, and has passed responsibility for group leadership to Leigh Ann Tanner.)

System Software for CRA

Part of his job involves coordinating systems software development with other partners of the NASA Cooperative Research Agreement (CRA) on the IBM SP2 system. (For more information about the CRA or the NAS SP2, see the previous four issues of *NAS News*.)

The CRA represents "a different way to interact with vendors and software suppliers," Harness said, adding that he appreciates the close association with IBM research centers. "One of our goals is to help IBM develop better software for the industry at large."

Harness believes that NAS has a major benefit to offer IBM as a CRA partner: "We have a group of sophisticated users who can put a demanding load on the system and who are accustomed to working with experimental software. We also have a technical staff that understands about making a product bullet proof."

There has been "good interaction" between the parallel systems staff and the Message Passing Interface (MPI) and Parallel Virtual Machine (PVMe) groups at IBM. Contact between NAS and two other IBM groups, High Performance Fortran (HPF), and Parallel File System (PFS), has been less productive, Harness said. "We haven't made as much progress as we'd like, which somewhat limits what we can do," because of internal delays at IBM.

In addition to the NAS part of the SP2 system and Silicon Graphics Inc. (SGI) cluster projects, Harness oversees parallel systems software development for MPI, MPI-I/O, HPF, PFS, and network device drivers. Most of this software is designed to run on multiple platforms.

Solicits Feedback on Policies

Harness regularly solicits feedback within the group to set policy for the computer systems he oversees, such as how to apply for dedicated time, how often a certain size job should run, or criteria for accepting users.

What's most challenging to him is "keeping all these separate efforts coordinated, and not losing track." He tries to ensure that "there's enough work in the pipeline to keep everyone happy and content."

He'd like less "paperwork and bureaucracy -- which unintentionally contributes to the slowdown." Harness acknowledges that the current NAS budget freezes are "a problem" and will "prohibit some aspects of the projects from moving forward." His solution is to keep focused on the highest priorities.

From Math to Behavioral Sciences

Harness received his Bachelor of Science degree in behavioral sciences from the University of Chicago, having started there as a math major and then getting "sucked up into computer sciences courses." After exhausting the limited computer science offerings at the time, his interest gravitated toward artificial intelligence. Thinking "it would be good to learn about human intelligence first," Harness enrolled in several behavioral sciences courses, and ultimately switched majors.

After working at the University of Chicago for five years, Harness was recruited in 1988 by General Electric Co. (GE), to work in the NAS Control Room. After the GE contract expired, Harness worked one year for Computer Sciences Corp. in the High Speed Processor (HSP) group. Harness resigned from this position in 1990 and joined the federal civil service to manage the HSP group in what is now the NAS Scientific Computing Branch. Harness became manager of the parallel systems group in 1993.

Exploring New Technology

What he likes most about his job is "exploring new technology." One of Harness's responsibilities is "playing a strong role in the acquisition of hardware and principal software components for NAS," and he's served on various evaluation teams for parallel systems and vector processors.

Outside of work, Harness pursues a number of intense, limits-testing interests -- among them: skydiving (both participating and teaching), racing 30 ft. sailboats, snow boarding, in-line skating, and scuba diving. For the water sports, he regularly travels 85 miles south to Monterey from his home in Mountain View and admitted that one of his dreams is to move closer to Monterey Bay.

A poster in Harness's office is testimony of his participation in the 1994 San Francisco marathon -- his first. "The challenge for every first-time marathon runner is finishing," Harness said. So far, he's made no promises about entering the race this year.

In spite of his quiet, low-key manner and obvious embrace of both personal and professional challenge, one gets the impression that Harness only takes on those commitments he's sure he can keep.

Next Article | Contents | Main Menu | NAS Home

May - June 1995, Vol. 2, No. 10

NAS NEWS

Volume 2, Number 10 May – June 1995

NASA's Future Aeronautics Role Debated at CAS Workshop

by Elisabeth Wechsler

CAS Work 'Absolutely Essential'

> "Industry can't afford to do basic research any more—the national labs are the last hope."

Short-term vs. Long-term Goals

Continued on page 2

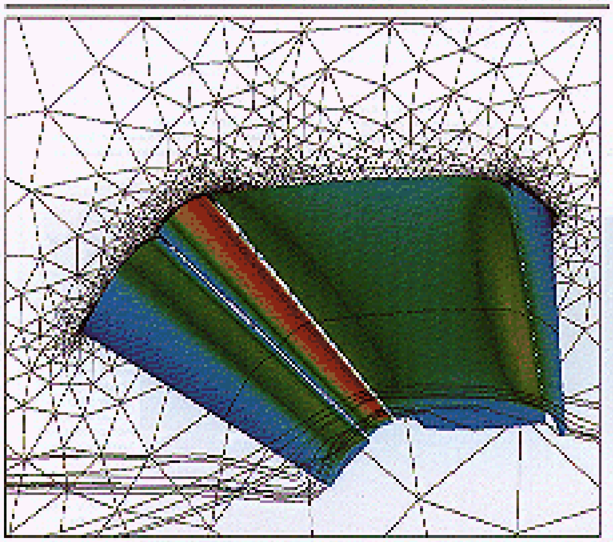

PBS Offers Flexible Job Scheduler for the IBM SP2

by Elisabeth Wechsler

Includes BASL and Tcl Languages

Big Improvement Over NQS

Continued on page 2

THIS ISSUE