CRAY CHANNELS

VOLUME 15, NUMBER 2 • A CRAY RESEARCH, INC. PUBLICATION

Announcing new hardware and software systems

CRAY CHANNELS

In this issue

Aerospace design and engineering are longstanding applications of Cray Research's supercomputer technology. Supercomputer modeling and simulation have saved the aerospace industry many millions of dollars and untold work-hours over the past decade. Originally used to evaluate aircraft designs to minimize expensive wind-tunnel testing, supercomputer modeling in the aerospace industry now ranges from traditional external aerodynamic studies and structural analysis to electromagnetic simulation and cabin airflow studies to improve passenger comfort.

This issue of CRAY CHANNELS surveys both new and traditional aerospace industry applications. The topics covered fall short of a comprehensive list of aerospace applications but convey the breadth of Cray Research's continuing contributions to the industry. This issue also introduces new hardware products at the high and low ends of the Cray Research supercomputing spectrum, along with the latest version of the MPGS engineering postprocessing software system.

The exacting demands of aircraft design and engineering can challenge even the most advanced computational resources. To meet the needs of the aerospace and other industries, Cray Research continually develops new products to solve the largest and most difficult problems in industrial research and development. Later this year the company will announce the CRAY T3D massively parallel computer system. This system will provide new opportunities for improved design and engineering efficiency, as well as basic research opportunities, throughout industry. As it has been with Cray Research's computer systems to date, the aerospace industry will be among the primary beneficiaries of this new technology. The result will be shorter design cycles for new aircraft, improved fuel efficiency and passenger safety, and an enhanced bottom line for companies that take advantage of the latest in Cray Research supercomputing technology.

Features

10

Departments

CRAY CHANNELS is a publication of the Cray Research, Inc. Marketing Communications Department. Published three times per year, it is intended for users of Cray Research computer systems and others interested in the company and its products. Subscription inquiries and address changes should be sent Attention Department D, CRAY CHANNELS, Cray Research, Inc., 655D Lone Oak Drive, Eagan, Minnesota 55121

Volume 15, Number 2

Cray Research expands successful CRAY Y-MP C90 system into series

The world's most powerful general-purpose supercomputer system, itself enhanced with four times its previous memory capacity, is now also available in smaller, lower-priced versions. So popular is the industry-leading CRAY Y-MP C90 supercomputer—19 orders received and 14 systems shipped since its November 1991 introduction—that Cray Research expanded this system's advanced technology to four new supercomputers with list prices starting at $3.25 million in the United States. In addition to the original 16-processor system, now called the CRAY C916 supercomputer and available with up to one gigaword of real memory, customers can choose from the following new systems:

- ☐ The CRAY C92A supercomputer, an air-cooled system available with one or two processors. This system requires no special cooling arrangements and no motor generator set and operates on standard 50 Hz and 60 Hz commercial electrical power, so it can be installed virtually anywhere in the world.
- ☐ The CRAY C94A system, an air-cooled system with the same features as the CRAY C92A system and available with up to four processors. The C94A system provides increased flexibility in terms of I/O and CPU capacity.
- ☐ The CRAY C94 system, a liquid-cooled system with two to four processors.
- ☐ The CRAY C98 system, a liquid-cooled system with four to eight processors.

These systems strengthen Cray Research's product line at the midrange and complement the success enjoyed at the top and low ends by the CRAY Y-MP C90 and the CRAY Y-MP EL systems, respectively. Each new system features the fastest memory technology available—four-megabit SRAM (static random access memory). The smaller systems are offered with up to twice the memory capacity of their counterparts in the earlier CRAY Y-MP product line.

More than 600 leading software applications currently available on the CRAY C916 and CRAY Y-MP systems will run without modification on the new series of systems. Each of the CRAY C90 systems makes use of the one-gigaflops CPU used in the original CRAY C916 system, which, when fully configured, has a peak system performance of 16 GFLOPS. On the ANSYS structural analysis code from Swanson Analysis Systems, Inc., the CRAY C916 system has sustained single-job performance of 5.9 GFLOPS and, when running two jobs simultaneously, has exceeded 7 GFLOPS.

Left to right: The CRAY C916, CRAY C98/C94, and CRAY C94A/C92A computer systems.

Model	CPUs	Central memory Mwords	I/O clusters	Optional SSD Mwords	Cooling	Peak Performance GFLOPS
CRAY C92A	1 or 2	32 - 128	1 - 2	32 - 2048	Air or water	2
CRAY C94A	2 or 4	64 - 128	1 - 3	32 - 2048	Air or water	4
CRAY C94	2 or 4	64 - 256	1 - 4	512 - 2048	Water	4
CRAY C98	4 or 8	128 - 512	1 - 8	512 - 2048	Water	8
CRAY C916	8 or 16	128 - 1024	2 - 16	512 - 4096	Water	16

In addition, 15 well-known third-party application programs now operate at sustained speeds of 6 to 11 GFLOPS on the CRAY C916 system. The highest sustained performance for any third-party code by a current massively parallel processing (MPP) system in the 1992 IEEE-sponsored Gordon Bell prize competition was 5.4 GFLOPS.

Beyond the more well-known applications available for the CRAY C90 series, many Cray Research customers are successfully writing their own codes. In many cases these proprietary codes have run in excess of 8 GFLOPS.

Cray Research intends to turn the strengths of its parallel-vector architecture into massively parallel processing (MPP) success. The thrust will be to ensure that the wide variety of third-party codes that run on Cray Research's parallel-vector systems today will run on Cray Research's MPP system tomorrow. The CRAY T3D system, Cray Research's first MPP system due out later this year, will closely couple microprocessor technology with the company's parallel-vector supercomputer systems, including all of the CRAY C90 series systems. All CRAY C90 systems feature the established modular Model E I/O subsystem, enabling the systems to work in conjunction with the CRAY T3D system. Cray Research expects the CRAY T3D system to outperform competing MPP systems.

The most advanced supercomputing architecture available

To provide unmatched problem-solving capabilities, CRAY C90 systems feature a balanced architecture with the following advantages:

- *1 to 16 powerful CPUs.* Each CPU has a peak performance of 1 GFLOPS. Combined with the most advanced parallel processing software in the industry and the highest level of memory bandwidth available, CRAY C90 systems deliver sustained gigaflops on a wide range of applications.
- *Up to 1 Gword (8 Gbytes) of high-speed central memory.* The CRAY C90 systems enable users to run large jobs in-memory with up to four times the central memory capacity of the original CRAY Y-MP C90 system.
- *Up to 256 I/O channels providing up to 13.6 Gbyte/s I/O bandwidth.* To provide high performance connections to peripheral devices and networks, the CRAY C90 systems feature flexible I/O subsystems with the highest I/O rates in the industry.
- *MPP connectivity.* The CRAY C90 systems can be coupled closely to the massively parallel CRAY T3D system. For highly parallel applications, this heterogeneous architecture can deliver an unprecedented level of performance to a broad spectrum of users.
- *Optional SSD solid-state storage device for increased throughput.* An optional SSD provides very-high-speed secondary memory with up to 4 Gwords (32 Gbytes) of storage capacity and a maximum bandwidth of up to 13.6 Gbyte/s.

For now, demand for the small CRAY Y-MP C90 systems remains strong. When this issue of CRAY CHANNELS went to press, the company already had received eight orders for these supercomputers. General Atomics will replace the San Diego Supercomputer Center's CRAY Y-MP system with a CRAY C98 supercomputer to be used for a variety of scientific and engineering applications by more than 3000 users at over 100 universities and research institutions affiliated with the center. The Instituto Nacional de Meteorología, Spain's weather forecasting and climate research center, will be the first weather organization to install one of the systems, a CRAY C94 supercomputer, in the new series. The Space Data and Computing Division of the NASA Goddard Space Flight Center will install a CRAY C98 system at the NASA Center for Computational Sciences (NCCS) in Greenbelt, Maryland. The new system will replace the center's CRAY Y-MP supercomputer and be used for earth and space sciences research involving global climate change studies, global data assimilation, atmospheric/oceanic modeling, geodynamic earth surface modeling, space physics theory, astrophysics, and earth observing systems simulation.

Cray Research doubles performance, quadruples memory capacity of successful entry level system*

Upgrading made easy for existing CRAY Y-MP EL customers

Cray Research's strong-selling CRAY Y-MP EL entry level supercomputing system—130 orders booked in its first full year—got even better this spring. In March the company introduced the CRAY EL98 system, with twice the peak performance of the CRAY Y-MP EL configurations at the same prices and with four times the central memory capacity. Because the CRAY EL98 system uses the same chassis as Cray Research's previous entry level system, customers who already have a CRAY Y-MP EL system can easily add or exchange modules to upgrade their processing power. Cray Research currently has 11 customers upgrading to CRAY EL98 systems; seven of these organizations are upgrading to fully configured eight-CPU systems.

With up to eight processors and 4096 Mbytes (512 Mwords) of central memory, the CRAY EL98 supercomputer provides a peak performance of over 1 GFLOPS. To deliver the highest levels of sustained performance, the balanced architecture of the CRAY EL98 system provides 4.2 Gbytes/s of total memory bandwidth in combination with 1.05 Gbytes/s of I/O bandwidth.

The CRAY EL98 system maximizes data throughput with up to four integral I/O subsystems (IOSs) per processor module, which allows up to 16 IOSs in a maximum configuration. Each IOS handles data transfers with disk subsystems, tape drive subsystems, and networks. The CRAY EL98 system also can be configured with up to four 100 Mbyte/s HIPPI-to-memory channels. The HIPPI interface provides high performance for network connectivity applications.

The large, real memories of CRAY EL98 systems provide unique problem-solving capabilities and enhanced throughput on large-memory jobs. The system is highly effective when out-of-memory solutions are not practical. By running large jobs in memory, turnaround time can be reduced by 3 to 10 times over comparable out-of-core memory solutions. Third-party scientific applications are available that take full advantage of this unique real memory feature on the CRAY EL98 system.

The CRAY EL98 system is binary compatible with CRAY Y-MP and CRAY Y-MP M90 supercomputer systems and runs Cray Research's powerful and feature-rich UNICOS 7 operating system and CF77 and Cray Standard C compilers. As problems become more complex, a CRAY EL98 application can readily be scaled to run on any of the larger systems. More than 600 leading software applications currently available on these systems, and on the original CRAY Y-MP EL supercomputer, will operate without modification on the CRAY EL98 system. Its strong parallelization capabilities and large real main memory give the CRAY EL98 system scalable application solution capabilities unique to its price range.

Cray Research's entry level systems have brought in more than 70 new customers in the aerospace, automotive, chemical, financial, construction, utilities, and electronics industries, as well as universities and environmental and general research centers. The Federal Home Loan Mortgage Corporation uses a CRAY Y-MP EL system to analyze risks and returns on its mortgage-backed securities business, computing most valuations at least 50 times faster than in the past. Amoco Oil Company and Landmark Graphics Corporation are two new petroleum industry customers.

Cray Research is advancing its long-term commitment to entry level products with a project aimed at delivering the next-generation system in the second half of 1994. This entry level system will be based on aggressive CMOS integrated circuit technology and provide substantial speed and performance gains.

Like all Cray Research systems, the CRAY EL98 supercomputer supports connectivity to a wide variety of computer systems and adheres to industry standards, providing easy integration into heterogeneous computing environments. As a result, the CRAY EL98 system can be used as a stand-alone compute server, a file server, or as a node in a heterogeneous network.

The CRAY Y-MP EL98 system runs on standard 200 - 240 volt, single-phase power and can easily be installed virtually anywhere, including on ships, at remote processing sites, or in an office environment. Both 50 hz and 60 hz power are supported to allow installation worldwide.

***Without raising price**

CFD environment for aerodynamic design at Aerospatiale Avions

Vincent Rivoire and Yvon Vigneron, Aerodynamics Department, Aerospatiale Avions, Toulouse, France

To optimize products and reduce development time and costs, Aerospatiale Avions continually refines the process that leads to the aerodynamic shapes of its aircraft. This process is traditionally based on an iterative scheme combining theoretical analysis and experimental validation. However, computing and algorithmic methods developed over the past 30 years have profoundly changed the aerodynamicist's approach.

The wide spectrum of fast and reliable computational methods available today allows early optimization of proposed configurations and better targeted and more accurate wind tunnel or flight tests. The diversity of these methods lets engineers choose the numerical tool that best addresses a given problem, making the design process more reliable and economical.

Aerospatiale Avions relies on a broad range of numerical methods and computational resources to reduce design cycle time, improve modeling accuracy, and achieve computational transparency. A CRAY X-MP vector supercomputer performs large-scale computations. A CDC CYBER 2000 front-end computer prepares computations and stores results. Workstations handle interactive processing and graphics.

Flexible, high-speed tools for preliminary design

In the preliminary design phases of a new aircraft program, the aerodynamicist assesses possible technical solutions as they relate to general design problems, such as positioning the engines on the airframe. These are often long parametric studies, during which the geometric data are refined and sometimes redefined as the program progresses. This process, along with budgetary and deadline constraints, requires flexible and economical

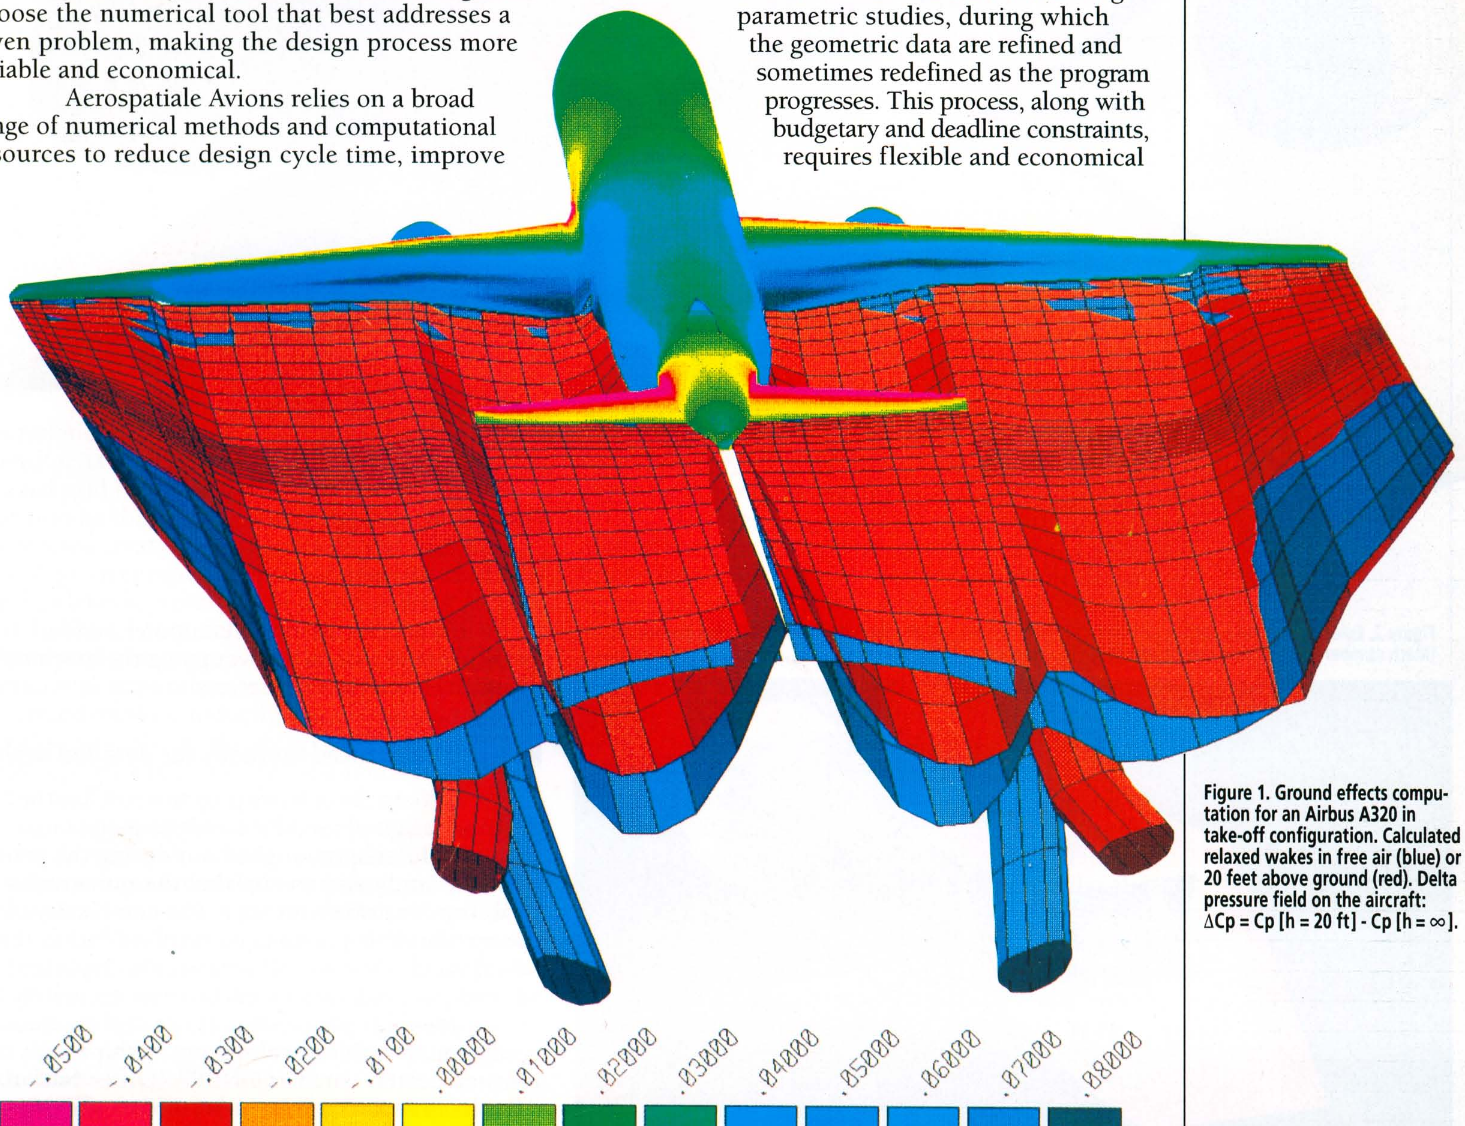

Figure 1. Ground effects computation for an Airbus A320 in take-off configuration. Calculated relaxed wakes in free air (blue) or 20 feet above ground (red). Delta pressure field on the aircraft: $\Delta Cp = Cp\,[h = 20\ ft] - Cp\,[h = \infty]$.

methods accessible to all users without special assistance.

For these reasons, the surface panel methods that heralded the advent of computational fluid dynamics (CFD) methods 30 years ago still are highly attractive for the preliminary study of complex geometries. Requiring only surface meshing of the aircraft, they greatly facilitate wide parametric aerodynamic analyses related to exploratory variations of the shapes and positions of aircraft components.

Figure 1 (previous page) illustrates a surface panel application developed by Aerospatiale Avions for calculating ground effects, taking the vertical speed of the aircraft into account.[1] Supplementing wind tunnel tests, this application helps characterize the aerodynamic behavior of the aircraft in certain flight phases. From a computational standpoint, Aerospatiale's use of its surface panel method is characterized by

- ☐ Temporary high-speed disk space requirements exceeding several gigabytes to store the aerodynamic coefficient matrices
- ☐ Simultaneous sets of a great number of computations in batch mode requiring limited user control and often providing automatic processing of results
- ☐ Possible access from all terminals via a user-friendly interface

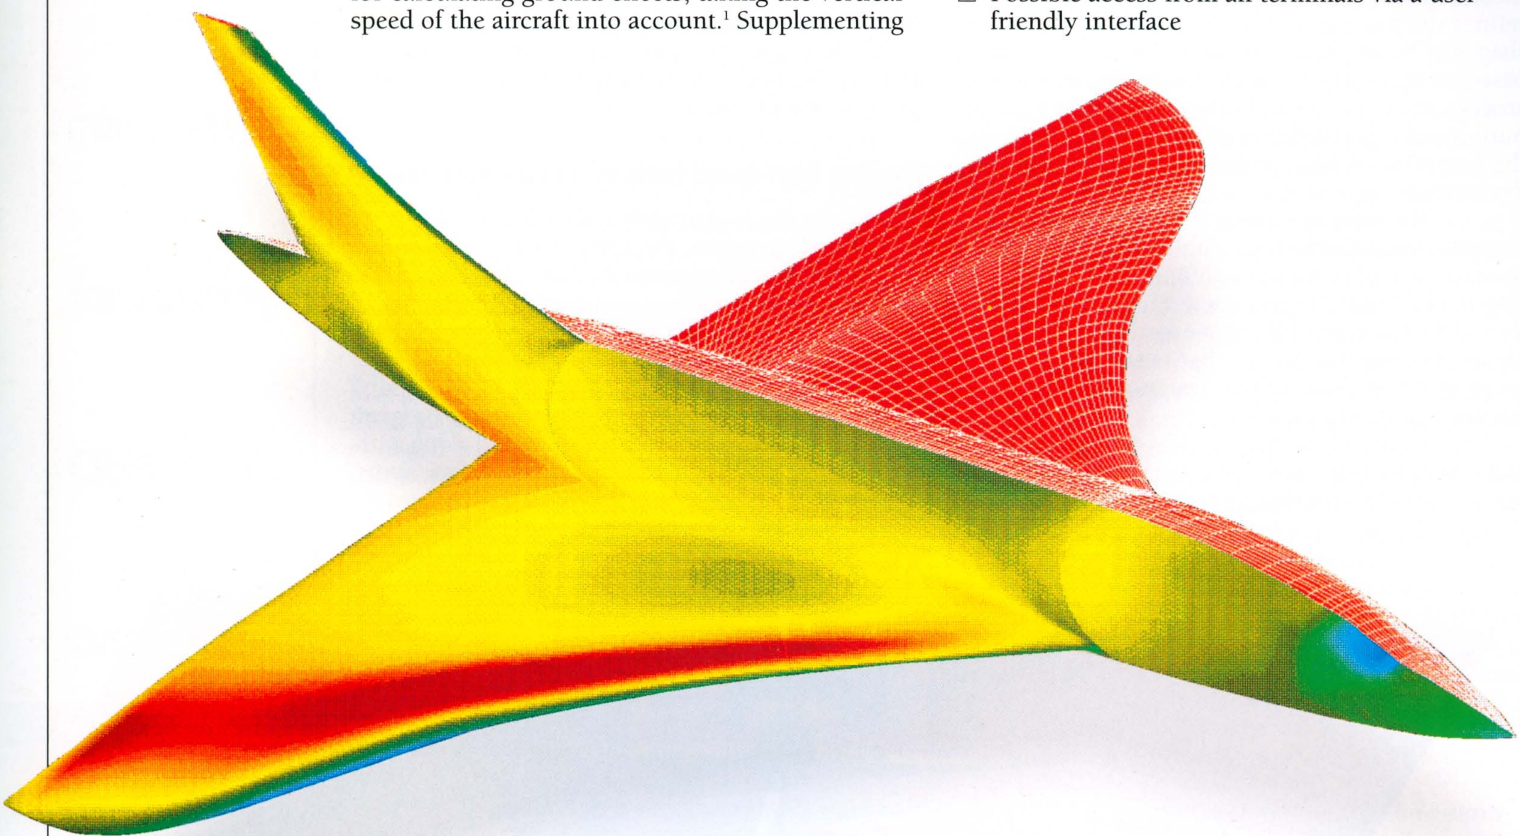

Figure 2. Euler computation for the Concorde airplane (2a, above) and its supersonic air intake (2b, below) (Mach number contours and streamlines in vertical plane).

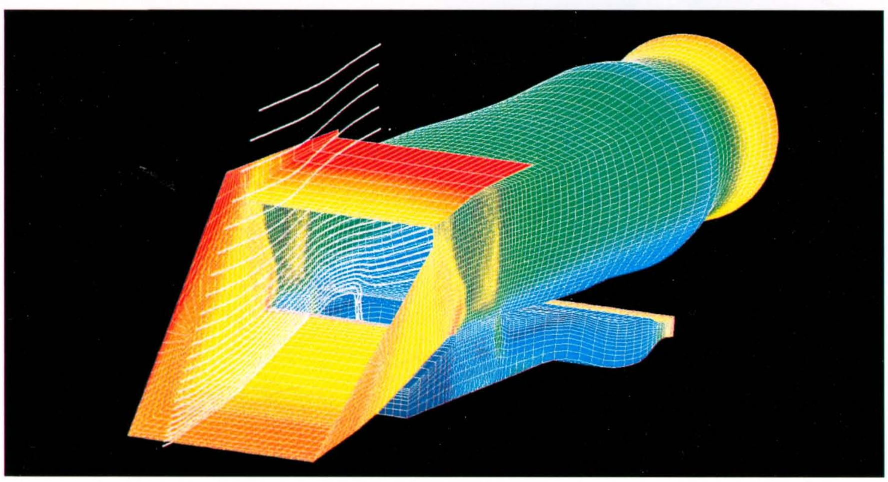

The Cray Research computer system meets these constraints. Occupying the machine's off-peak periods minimizes costs.

Powerful, advanced methods for detailed design

When the aircraft program reaches the detailed design phase, the aerodynamicist must optimize the shapes roughed out during the preliminary study and ensure that the numerous flyability constraints are met. The complexity of the aerodynamic phenomena involved makes the task difficult. Numerical methods that represent all transonic and viscous phenomena accurately on complete aircraft configurations still are undergoing considerable development with partner research centers, such as ONERA (Office National d'Etudes et de Recherches Aérospatiales), the national office for aerospace research and development.

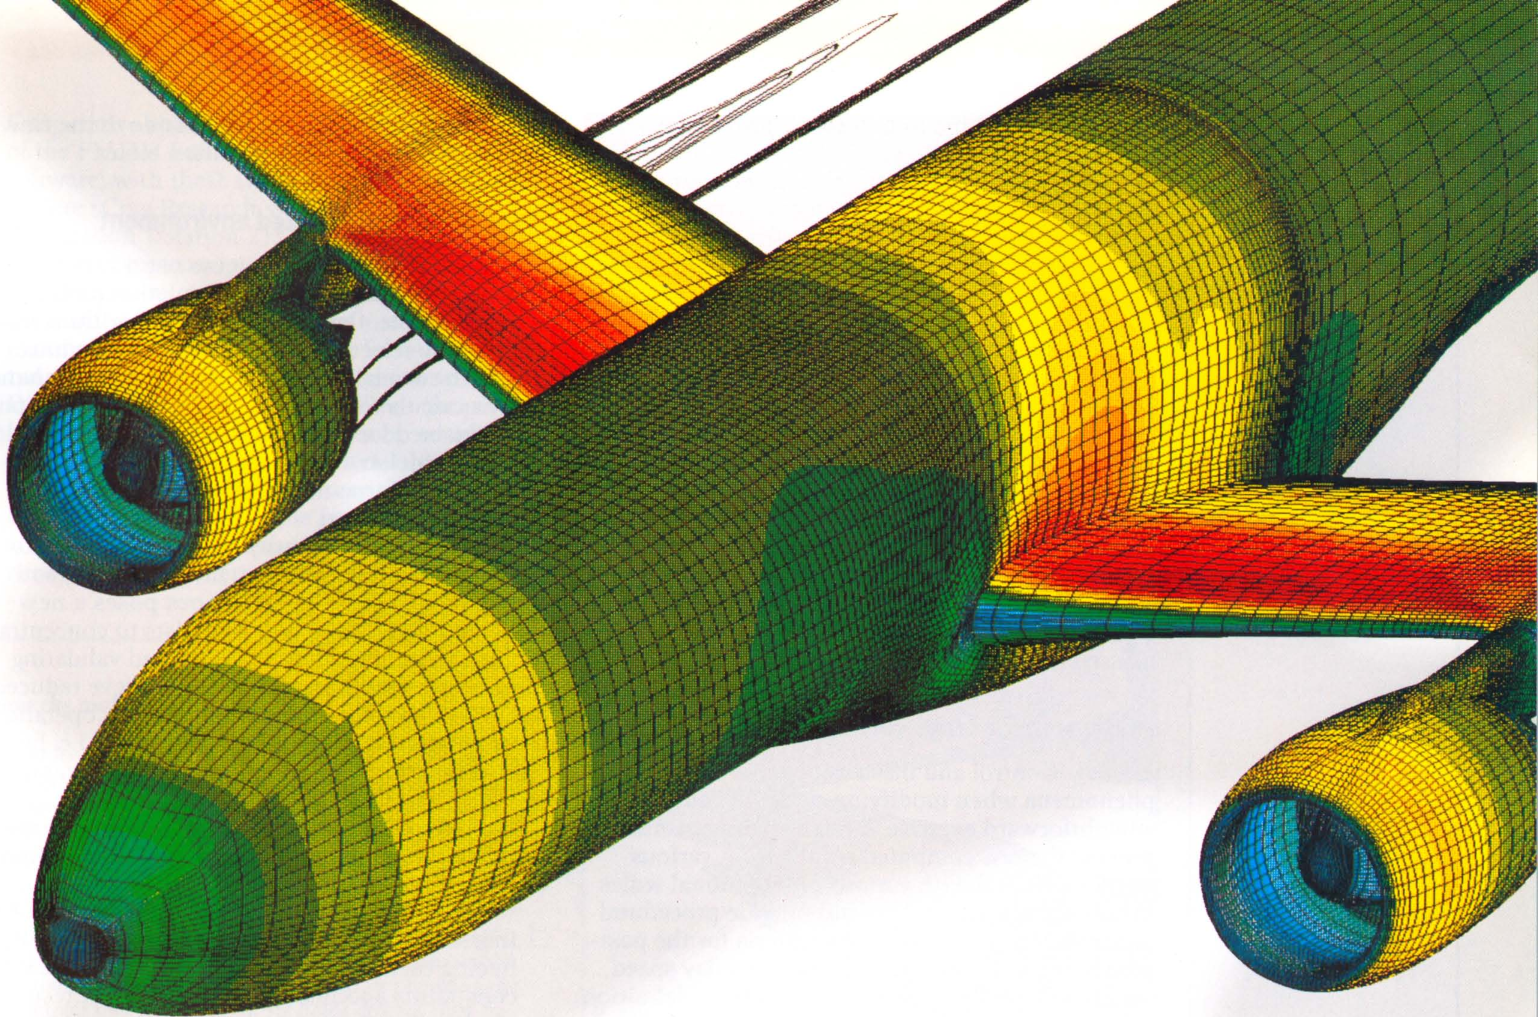

Figure 3. Euler computation for an engine integration study, Mach = 0.82, 1.4 million nodes (pressure contours and iso-total pressure lines in the jets).

The preliminary step in the use of refined computational methods that take these nonlinear phenomena into account is to define a computational grid covering the complete fluid domain surrounding the aircraft. For a long time, this meshing process created a bottleneck in the computation of these complex aircraft shapes. To reduce that bottleneck, Aerospatiale developed the software ICEM-CFD in collaboration with Control Data several years ago, which meets Aerospatiale's operational requirements: production of block structured meshes; reduction in production time (thanks especially to the ability to reactivate the meshing commands on similar aircraft configurations); and efficient control of mesh quality.[2]

A computational software package currently used at Aerospatiale Avions[3] solves the full potential equation using a finite element technique[4] coupled with a local boundary layer method developed at ONERA-CERT (Centre d'Etudes et de Recherches de Toulouse). It is accessible via a workstation interface that ensures a direct link to the volumetric mesh and graphical postprocessing. This interface also allows design engineers to manage all results with database-type functions.

The methods for solving Euler equations constitute a second generation of CFD tools. Compared with the full potential model described above, they take rotational flow into account and therefore allow for better models of turbulent wakes and engine jets at various energy levels.

Figure 2 shows Euler computations around a supersonic aircraft and its air intake with boundary layer bleed slot.[5] Optimal air intake design represents one of the key points in the Alliance supersonic transport aircraft program. For this application, many cycles of the solver FLU3M, developed by ONERA, needed to be computed during the same day. Complementary to computations on Cray Research systems, a complete implementation of the FLU3M solver was used on a workstation.

Another use of Euler codes at Aerospatiale Avions is illustrated in Figure 3. This exploratory study of a new engine installation on an aircraft of the Airbus family demonstrates in particular the effects of the engine core and fan jets on the flow around the wing at various engine speeds.[6] Because these computations require main memory capacity of 50 to 100 Mwords and high CPU power, they were run on CRAY Y-MP or CRAY-2 systems at partner sites. Close cooperation with ONERA, the developer of the Euler solvers used by Aerospatiale, allows these codes to become basic design and engineering tools.

Beyond the Euler model, only Navier-Stokes equations allow simulation of all aerodynamic phenomena brought into play around the aircraft, especially three-dimensional separations. The codes used by Aerospatiale were developed in cooperation with the research centers ONERA and CERFACS (Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique). To obtain sufficient accuracy for normal transport aircraft flight conditions, these codes require extremely fine computational meshing close to the airplane surface. The use of vector supercomputers such as the CRAY-2 and CRAY Y-MP systems is indispensable for reasonable computing times and in laying the groundwork for future standard industrial use. Figure 4 shows the relatively simple computation of a wing in a wind tunnel, demonstrating massive separation induced by a shock. Efforts are under way to adapt these methods to massively parallel computers.

Postprocessing large volumes of information

Control and mastery of aerodynamic phenomena when modifying shapes is rarely a straightforward exercise. Design engineers must process the raw computed results from various standpoints and with various observational scales to substantiate key points and provide procedural guidance. For this reason, the criteria for the postprocessing tools are execution and display speed, screen and hardcopy peripheral graphical definition quality, a high level of interactivity, and extension capabilities for new functions.

Figure 5 illustrates Aerospatiale Avions' QUICKVIEW postprocessing and visualization software. This software allows real-time processing of results for meshes comprising more than one million nodes.

The QUICKVIEW interface lets users create various superimposable graphical objects and move them by scanning the computed volume using the mouse. This also allows animation and video recording of unsteady phenomena, adding the time variable to the flow display and making interpretation easier.

Developing an optimized environment

Design aerodynamicists use a variety of increasingly complex digital simulation tools to fulfill their role. On the one hand, algorithms must be specially selected to complement the architecture of the available computers. On the other hand, today's calculation power allows detailed solutions to be obtained for problems of increasing complexity. Aerodynamicists thus handle an ever-growing quantity of information to study and validate the shapes they produce.

The large quantity of data that needs to be handled in the modeling, interpretation, and synthesis of aerodynamic shapes poses a new challenge. To allow aerodynamicists to concentrate on the essential tasks of studying and validating these aerodynamic shapes and to achieve reduced design cycles and costs, data processing operations must be made as transparent as possible, i.e. both hardware and software tools need to be organized in an optimal fashion.

Aerospatiale Avions is developing an integrated software system called Atelier Aérodynamique (aerodynamic workshop) to

- ☐ Improve aerodynamicists' productivity by freeing them of purely data processing-related constraints and making the quantity of results produced easier to handle
- ☐ Improve the quality and reliability of the tools used and reduce maintenance costs
- ☐ Promote self-training of new users
- ☐ Optimize the management and volume of the numerical data

Development priorities include

- ☐ Improving communications between the software applications
- ☐ Improving the ergonomics of the applications
- ☐ Developing computer engineering techniques (validation, industrial standardization, and maintenance of applications)

In parallel with this project, pilot applications have been launched to assess possible contributions from specific data processing techniques.

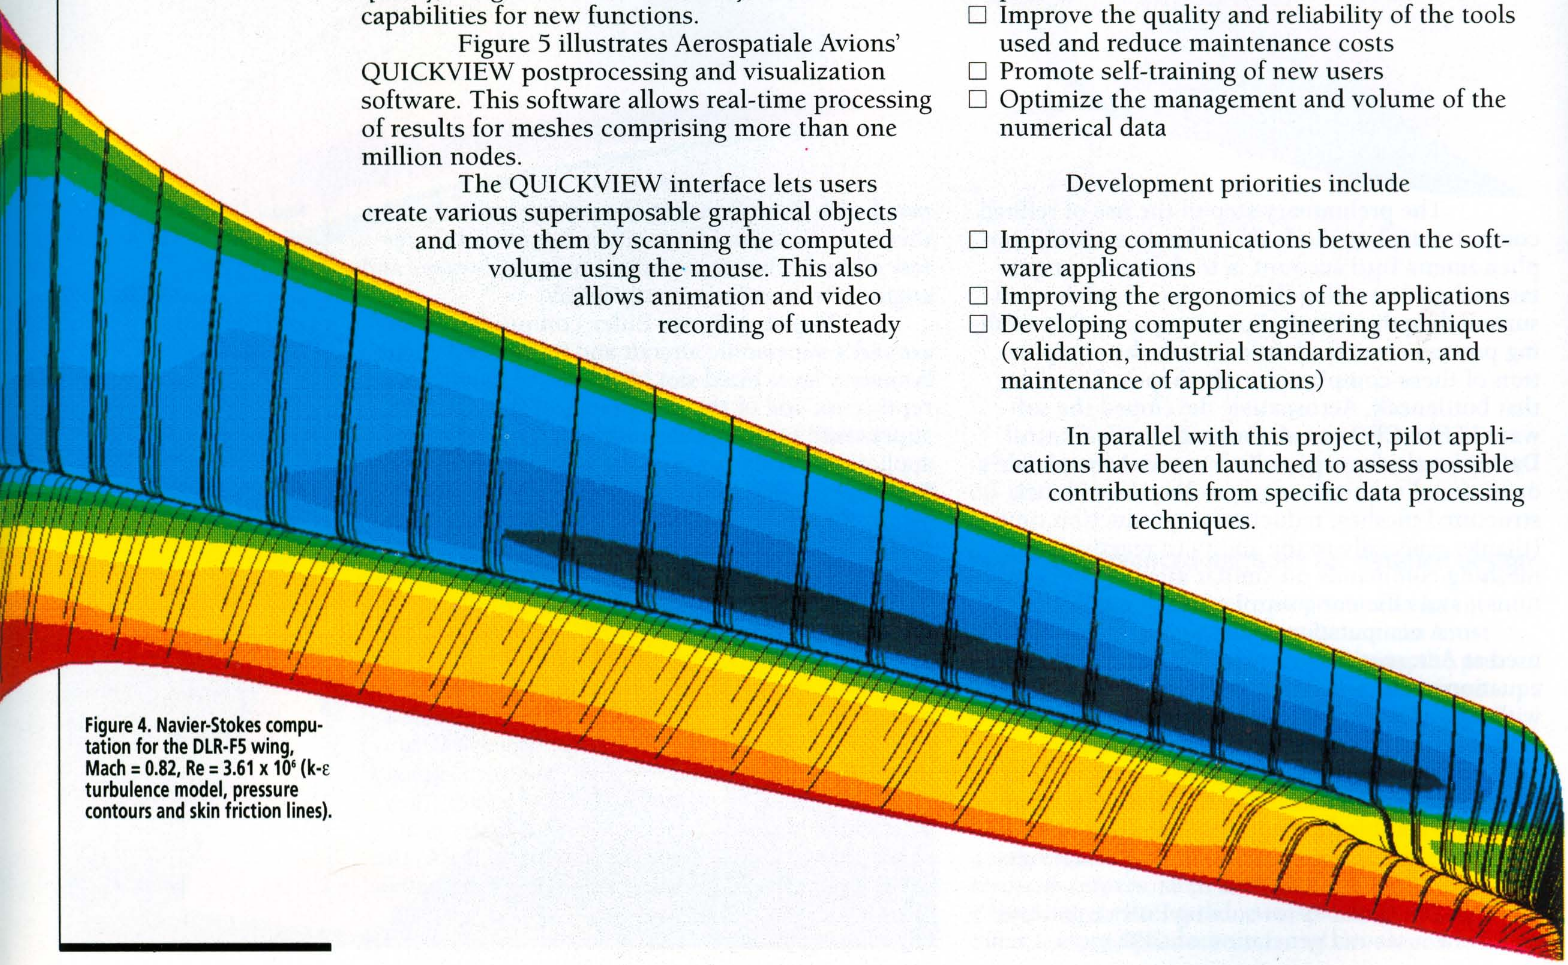

Figure 4. Navier-Stokes computation for the DLR-F5 wing, Mach = 0.82, Re = 3.61×10^6 (k-ε turbulence model, pressure contours and skin friction lines).

□ Cooperative applications with integral control of the FLU3M Euler solver by the QUICKVIEW software; with these two processes being either remote (Cray Research system/station) or on site (station 1/station 2).[7] This application lets users interrupt computations at any time, modify computational parameters, and display the results during the computational process, making it particularly attractive for initial fine-tuning of computations on new geometries.

□ Distributed applications. This allows the parallelization of computational codes on several systems and optimal use of all available system resources. A study is under way using the PVM library to measure the feasibility and gains achievable by distributing FLU3M across a network of computers.

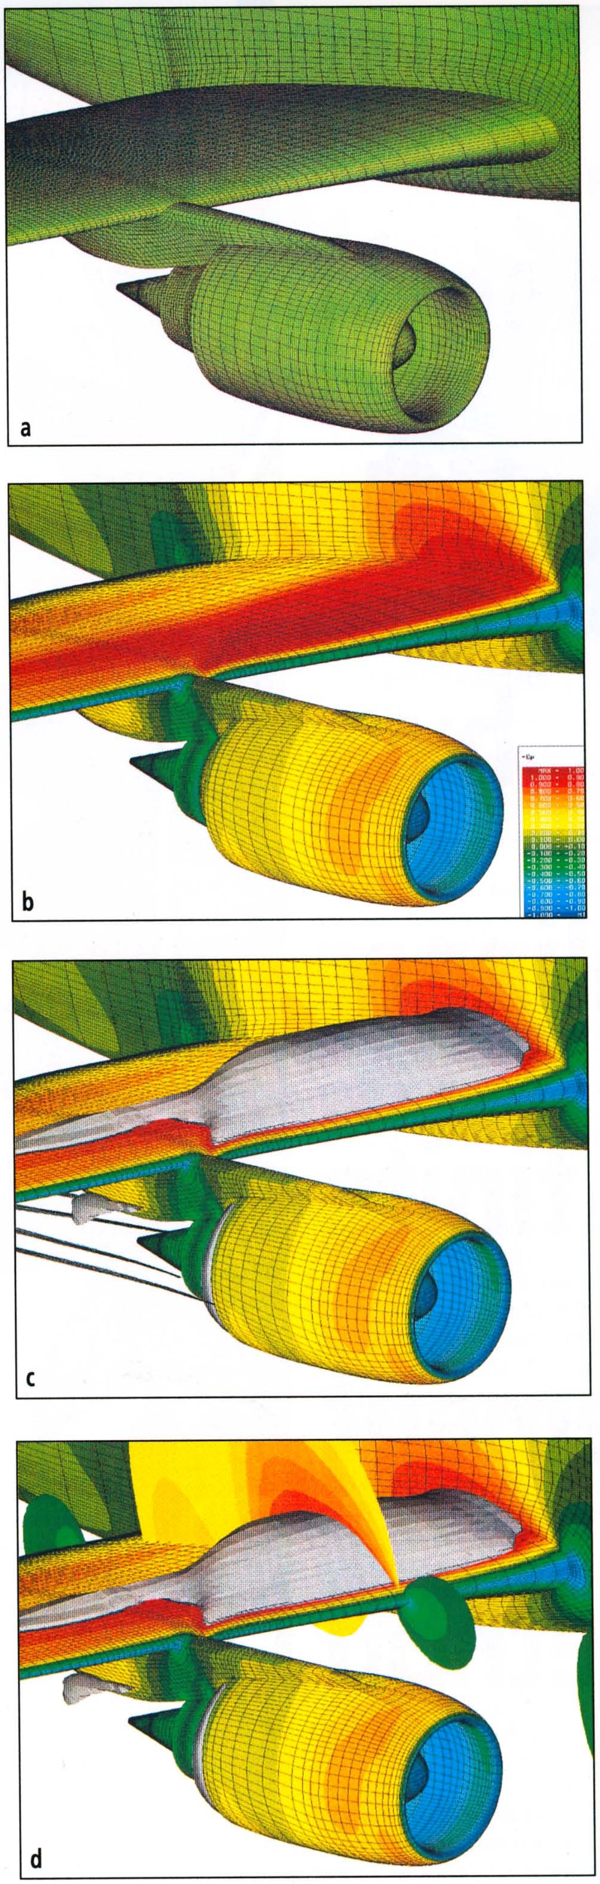

Figure 5. Images created with the QUICKVIEW postprocessing software (Euler computations): (a) Surface mesh, (b) Pressure contours, (c) Iso-total pressure lines in jets and iso-Mach = 1.2 pocket surface, (d) Pressure contours in vertical cut plane.

About the authors

Vincent Rivoire is head of the computational fluid dynamics engineering group at Aerospatiale Avions, in charge of providing advanced and reliable software tools to the design teams. He is a graduate of the Ecole Nationale Supérieure de l'Aéronautique et de l'Espace, Toulouse.

Yvon Vigneron is head of the Aerodynamic Design Department, in charge of aerodynamically defining the aircraft parts for which Aerospatiale Avions is responsible within the Airbus and ATR product lines as well as providing the related tools and technologies. He holds degrees from the Ecole Nationale Supérieure de l'Aéronautique et de l'Espace, Toulouse, and Iowa State University.

References

1. Prel, P., V. Rivoire, and J. Bardou, "Prévision des effets de sol sur un avion de transport commercial (Prediction of ground effects for a civil transport aircraft)," 29th Applied Aerodynamics Colloquium, CEL Biscarosse, September 1992.
2. Bertin, D., J. Lordon, and V. Moreux, "A new Automatic Grid Generation Environment for CFD Applications," AIAA paper no. 92-2720, Palo Alto, 1992.
3. Moreux, V., D. Bertin, and V. Rivoire, "Simulation of Transonic Flows over Complex Transport Aircraft Configurations," to appear in La Recherche Aérospatiale, Vol. 2, 1993.
4. Vigneron, Y., T. Lejal, and R. Collercandy, "Simulation d'écoulements par éléments finis (Flow simulation by finite elements)," 6th International Symposium on Finite Element Methods in Flow Problems, Antibes, June 1986.
5. Bousquet, X., D. Gisquet, and T. Surply, "Numerical Simulation of the Airflow in Supersonic Air Intakes: Application to the Future Supersonic Transport Aircraft," ICAS paper 92_3.3.2, 18th ICAS Congress, Beijing, September 1992.
6. Colin, P., P. Mogilka, and A. Molbay-Arsan, "CFD Analysis of Engine/Airframe Interference for Advanced Integration Studies on Transport Aircraft," to be presented at RAS 1993, the European forum on recent developments and applications in aeronautical CFD, to be held in Bristol, September 1993.
7. Soulard, A., "Application coopérative entre un Cray et une station de travail CDC910 (Cooperative application between a Cray and a CDC910 workstation)," presentation to the CRAY Latin User Group, Toulouse, November, 1992.

Gear system analysis by a nonlinear finite element method

Mark J. Valco, Vehicle Propulsion Directorate
U.S. Army Research Laboratory
NASA Lewis Research Center, Cleveland, Ohio

In the ongoing pursuit of improved aircraft performance, propulsion system designers share a common goal: reduce the weight, size, and cost of propulsion systems while increasing their power, efficiency, and reliability. This objective applies to the design of helicopter and fixed-wing vehicle drivetrains alike, but the difference in requirements is significant. Unlike fixed-wing aircraft, which generally can glide in the event of a propulsion system failure, helicopters rely totally on the integrity of their lightweight high-power-density drivetrains for their airworthiness. Although helicopter drivetrains exist in a variety of configurations, they all use gears to transfer the engine's high-speed rotating shaft power to the aircraft's lower-speed rotor systems.

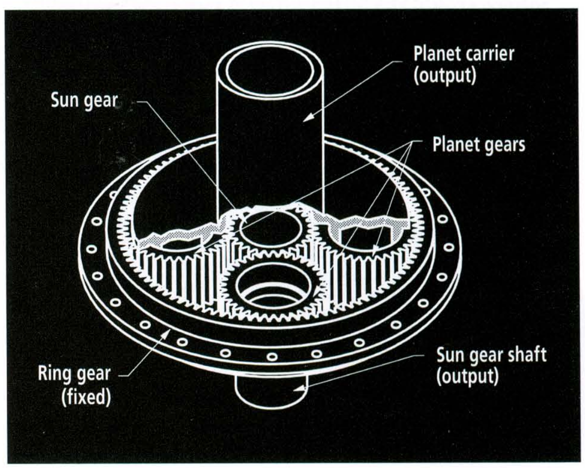

Figure 1. Planetary gear train.

Most contemporary helicopter drive trains have a planetary gear train (Figure 1) which functions as the low-speed, high-torque final reduction stage in helicopter main transmissions. The finite element method (FEM) of structural analysis, well known for accurate deformation and stress analysis of complex forms, is especially advantageous in the study of gear meshing action at the National Aeronautics and Space Administration (NASA) Lewis Research Center (LeRC).

Past research

While gears have been used for thousands of years, modern gear design methods still require the use of empirical and experience-based techniques to overcome limitations in analytical procedures, complicated gear tooth geometry, and the statically indeterminate load sharing in meshing gear teeth. Since the early 1970s, many researchers have performed FEM analysis of gear teeth, with a primary focus on the calculation of the gear tooth root and fillet stresses. The finite element gear models typically consist of a single gear tooth or a sector of a gear. Estimated point loads are applied to a tooth to simulate the gear mesh loading, and reasonably good agreement with experimental and a variety of analytical methods has been reported.[1] Values for the direction, magnitude, and location of the tooth load are important parameters for conducting these analyses.

More recent FEM gear analysis research uses gap elements to simulate tooth pair contact conditions and solve the statically indeterminate multiple tooth pair load sharing problem.[2] A complication here is that proper application of the tooth contact forces depends on the gap elements having tooth-to-tooth nodal alignment in the tooth pair contact region in the deformed state. As gear teeth rotate through the meshing cycle, the tooth surfaces roll and slide through contact. Advances in FEM provide deformable-body-against-deformable-body contact detection algorithms that can be used to establish the nonlinear boundary conditions to solve contact problems.

For the gear modeling and analysis at LeRC, Version 2.5 of the PATRAN[3] mechanical computer-aided engineering system was used to pre- and postprocess the FEM model data. PATRAN was run on a Silicon Graphics IRIS 4D/340VGX workstation and on a Digital VAX 9410 node of the NASA LeRC Scientific VAX Cluster. The FEM analysis was performed using version K4.1 of the MARC Finite Element Program[4] running on the NASA LeRC CRAY Y-MP8/6128 supercomputer.

The FEM advantage

A study was conducted prior to analyzing gear teeth in contact to verify that the automated FEM contact algorithms can provide reasonably accurate estimates of the localized deformations and stresses associated with elastic body contact.[5] The FEM contact results were compared with Hertzian contact semi-infinite body theoretical solutions, and the results indicate that FEM contact analysis can yield solid accuracy, provided an appropriate finite element mesh density and contact tolerance are specified.

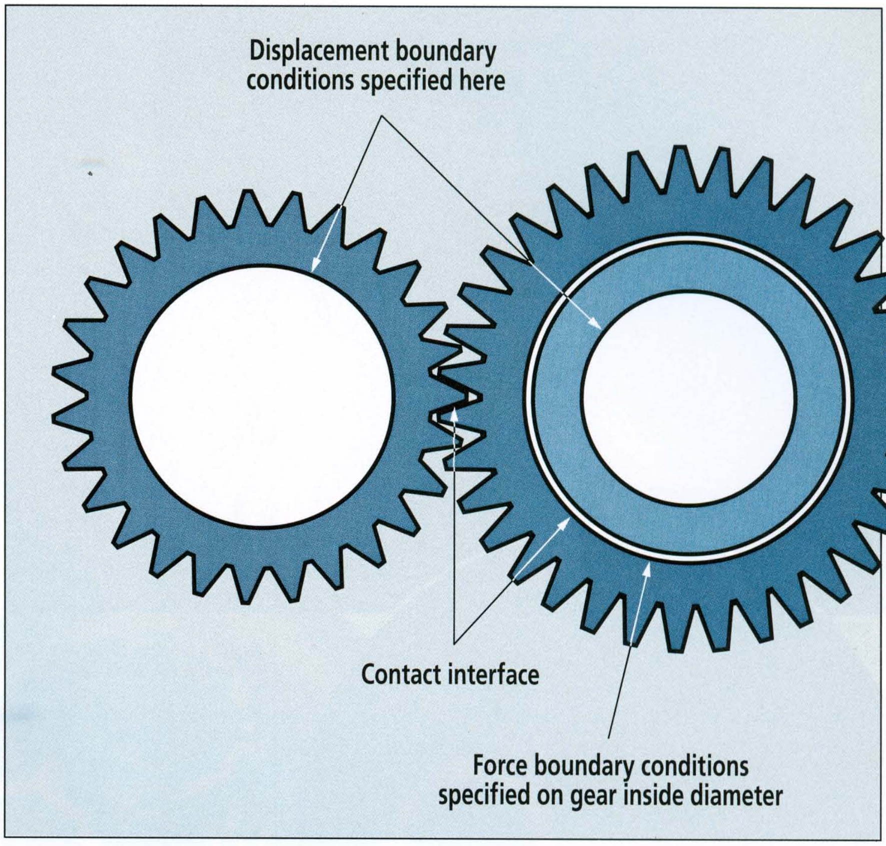

Figure 2. Gear pair model boundary conditions.

FEM's suitability for analyzing components with complex geometry, in conjunction with its ability to perform deformable-deformable body contact analysis, are key features in applying the method to gear analysis. The SPURPAT computer program[5] was used to create a PATRAN-readable input file of commands to generate each spur gear finite element model. SPURPAT can model the three types of spur gear configurations: gears with rims, gears with rims and webs, and gears with rims, webs, and hubs.

The ability to perform finite element contact analysis allows user-specified boundary conditions to be placed further from the critical sections of the gear teeth being analyzed. Instead of modeling a rigidly supported gear tooth or gear sector, a complete gear with an elastic supporting structure is modeled. Figure 2 shows a diagram of a gear pair arrangement for finite element analysis. One gear is used to apply incremental displacement boundary conditions for gear rotation, and the other gear is used to apply incremental force boundary conditions for the gear torque. These boundary conditions are specified and controlled by the MARC user subroutine interface FORCDT.

A quasistatic incremental nonlinear finite element gear analysis is then performed using the updated Lagrangian formulation allowing for small strains with large displacements. Nonlinear analysis is required due to the large displacements associated with the gear rotation and the changing boundary conditions associated with contact analysis. At each increment of rotation the gear bodies are evaluated for surface contact engagement or

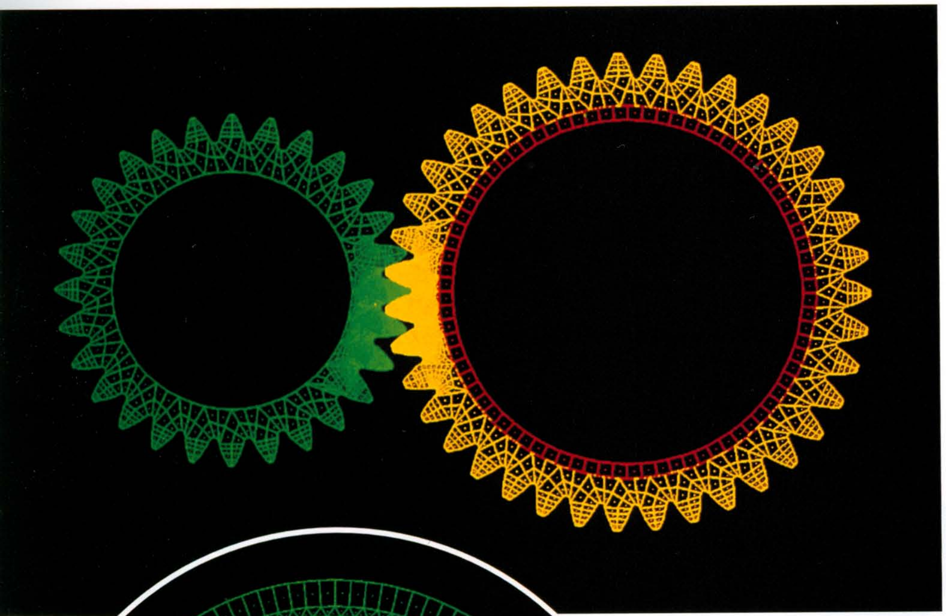

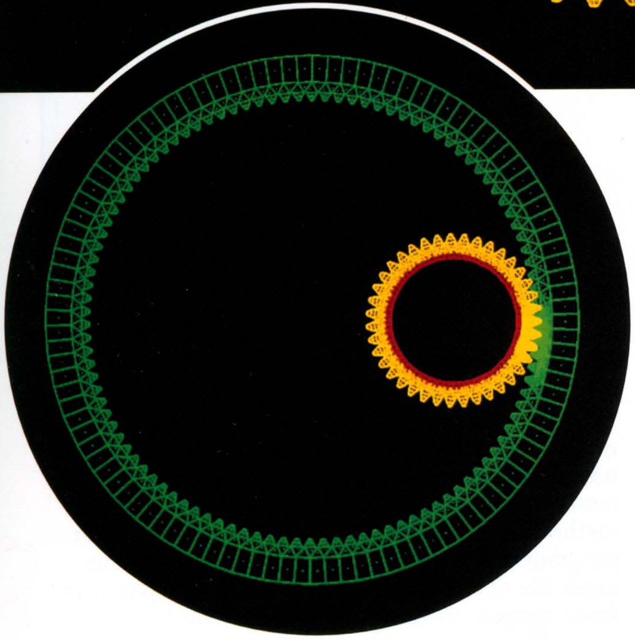

Figure 3 (above). External/external gear pair FEM model.

Figure 4 (left). External/internal gear pair FEM model.

separation, and the contact boundary conditions are recycled until the equilibrium convergence criteria are satisfied. At the start of a gear pair analysis the mating gears may not be in contact due to the gear tooth backlash clearance. Therefore, a system of body-connecting spring elements is required to prevent rigid body modes. After torque is applied and the mating gears enter into contact, the stiffnesses of the body-connecting springs are reduced to relatively insignificant levels.

Requirements and results

Analysis of gear pair meshing action using finite element contact methods requires an accurate geometrical representation of the gear tooth surfaces. The tooth pair contact points and the level of load sharing with multiple tooth pairs in contact is influenced by the tooth profile geometry. Achieving an accurate representation of the gear tooth profile geometry requires a relatively refined finite element mesh. Furthermore, very refined finite element meshes are necessary to produce reliable stress predictions in the tooth contact and root/fillet regions. Unfortunately, as the finite element model becomes more refined, the size of the mathematical model increases. A large finite element model in conjunction with the solution requirements for a nonlinear incremental analysis leads to very substantial computer processing time. To demonstrate the analysis method and obtain reasonable computer turnaround times, some compromises in the level of mesh refinement were necessary.

Finite element analysis was performed on an external/external gear pair (3747 elements, 4603 nodes, 9206 degrees of freedom, 1100 rotation increments, 65 CPU hours) and an external/internal gear pair (3777 elements, 4829 nodes, 9658 degrees of freedom, 1560 rotation increments, 152 CPU hours). Figures 3 to 6 show the finite element meshes and stress contour plots. The results from

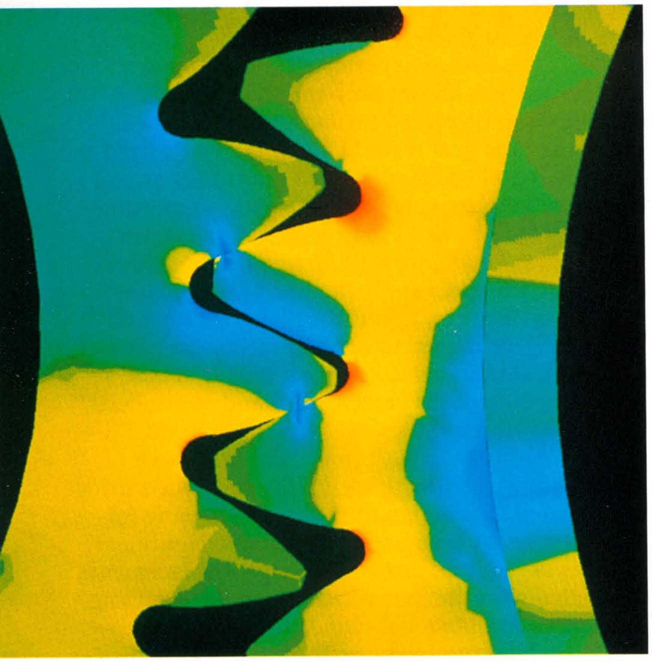

Figure 5. External/external gear pair stress plot.

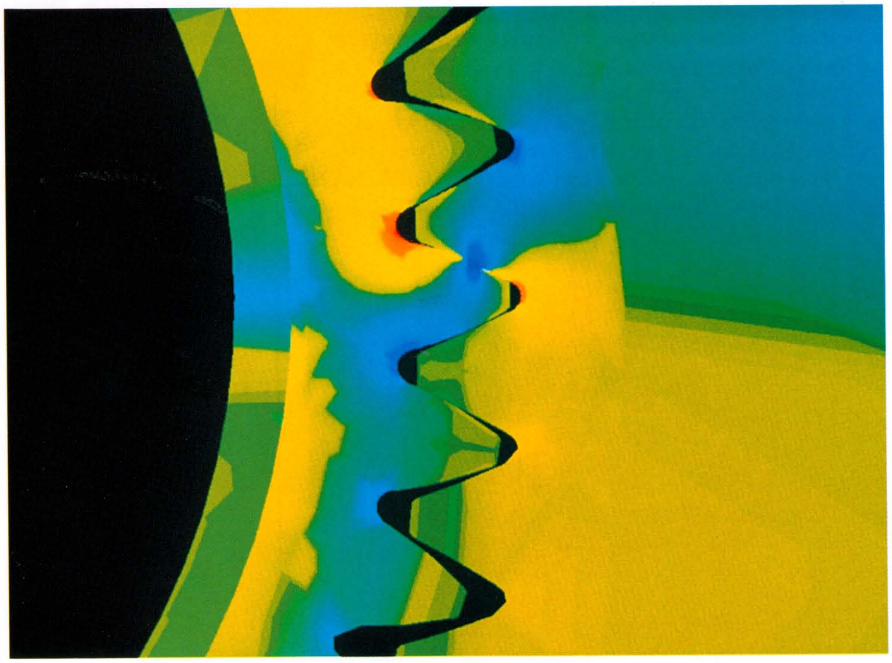

Figure 6. External/internal gear pair stress plot.

the analyses demonstrate the gear analysis method's ability to obtain deflection and stress data as the gears roll through mesh. Although a high angular rotation rate is desired to advance the gears through mesh in the fewest increments, the automated contact system must track the body interactions as the incremental gear rotation proceeds. The gear rotation rates ranged from 0.0005 to 0.0250 degrees per increment. The gear analysis procedure can be used to determine tooth pair engagement locations and the load distribution for single and multiple tooth pairs in mesh.

The gear pair modeling procedure was expanded to model a planetary gear train. Once again the SPURPAT and PATRAN software combination was used to model and locate the sun, ring, and planet gears along with a planet carrier into a planetary gear train configuration. In a planetary gear train the ring gear is a nonrotating member, and in the finite element model it is constrained and supported by radial and tangential elastic foundations. The elastic foundation simulates the elastic structural support provided to the ring gear by the transmission housing. For the planetary gear train model, incremental force boundary conditions which specify the torque are applied to the sun gear, and incremental displacement boundary conditions which specify the angular rotation are applied to the planet carrier. MARC user subroutine boundary condition control (FORCDT) and rigid-body-mode-preventing body-connection springs are used in the planetary gear train analysis in a manner analogous to the method for gear pair analysis.

A four-planet planetary gear train was modeled and analyzed. The finite element model is shown in Figures 7 and 8 (12,952 elements, 15,862 nodes, 31,724 degrees of freedom, 416 rotation increments, 216 CPU hours). The planetary gear train model has ten deformable bodies involved in contact interactions. Each of the bodies has simultaneous contact interactions with three or four different bodies. The results from the planetary gear train analysis demonstrate the ability of the nonlinear finite element method with automated contact to estimate deflections, strains, and stresses throughout the planetary gear train system.

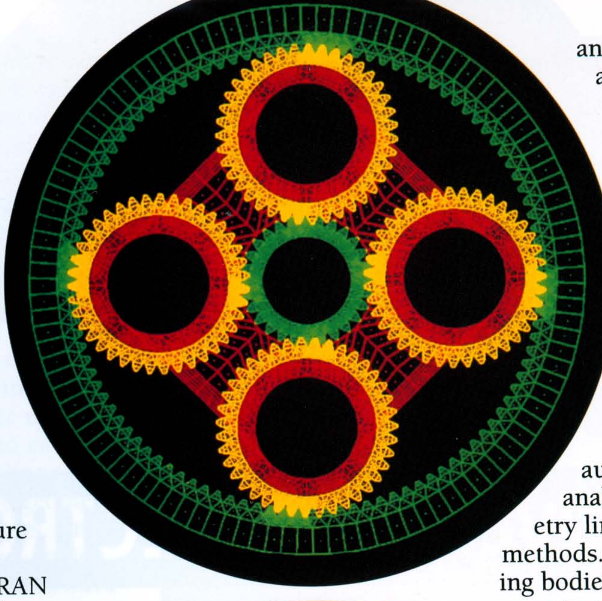

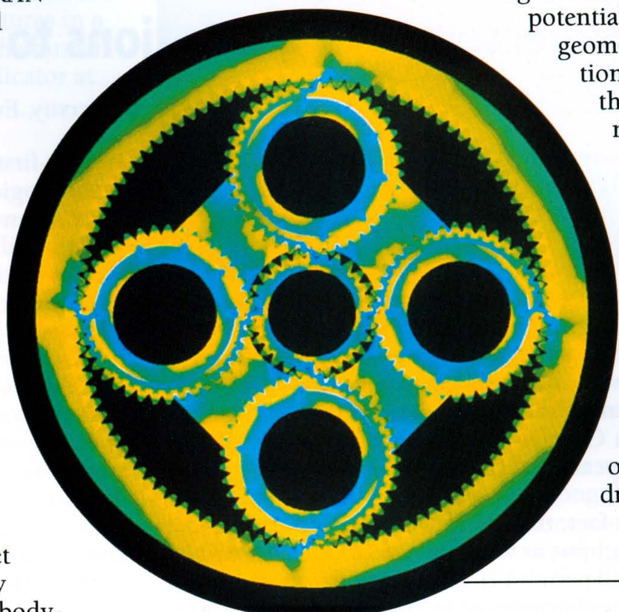

Figure 7 (top). Planetary gear train FEM model.

Figure 8 (above). Planetary gear train stress plot.

Summary

This research advances gear analysis technology by applying nonlinear finite element techniques to the study of spur gear pair meshing action. The data indicates that accurate FEM contact results can be obtained if very refined finite element meshes are used. A finite element gear analysis procedure was developed for application to spur gear systems.

The method was applied to an external/external spur gear pair, an external/internal spur gear pair, and spur planetary gear train. The gear analysis method can be used to study the influence of tooth geometry and a gear support structure on gear tooth load sharing and gear system transmission error. In addition, gear tooth and gear body deflections, strains, and stresses can be analyzed.

Application of the FEM with automated contact capability to gear analysis overcomes the gear tooth geometry limitations of most other gear analysis methods. The strength of the FEM in evaluating bodies with arbitrary geometric shapes has potential for application to other gear tooth geometries and gear system configurations. Future work in this area involves the addition of greater detail to the models, improved bearing supports, and three-dimensional gear geometry. While the computer processing requirements increase with more complex models, software algorithm improvements in the latest release (K.5) of the MARC program promise to enhance the computational efficiency of gear models. Additionally, future computer hardware performance improvements will provide greater opportunities for analysis of rotorcraft drive systems.

About the author

Mark Valco is an aerospace engineer conducting research in the area of rotorcraft drive systems for the U.S. Army Research Laboratory, Vehicle Propulsion Directorate, located at the NASA Lewis Research Center. He received a Ph.D. in engineering from Cleveland State University.

References

1. Chang, S. H., R. L. Huston, and J. J. Coy, "A Finite Element Stress Analysis of Spur Gears Including Fillet Radii and Rim Thickness Effects," *Journal of Mechanical Design,* Transactions of the ASME Paper No. 82-WA/ED-35, American Society of Mechanical Engineers, New York, November 1982.
2. Drago, Raymond J. and Ravi N. Margasahayam, "A Non-linear Finite Element Approach to the Gear Tooth Load Sharing Problem," MSC/NASTRAN World User's Conference, Los Angeles, California, March 1989.
3. *PATRAN Plus,* Release 2.5, PDA Engineering, Costa Mesa, California, 1991.
4. *MARC Finite Element Program,* Revision K.4-1, MARC Analysis Research Corporation, Palo Alto, California, 1990.
5. Valco, Mark J., "Planetary Gear Train Ring Gear and Support Structure Investigation," Ph.D. dissertation, Cleveland State University, June 1992.

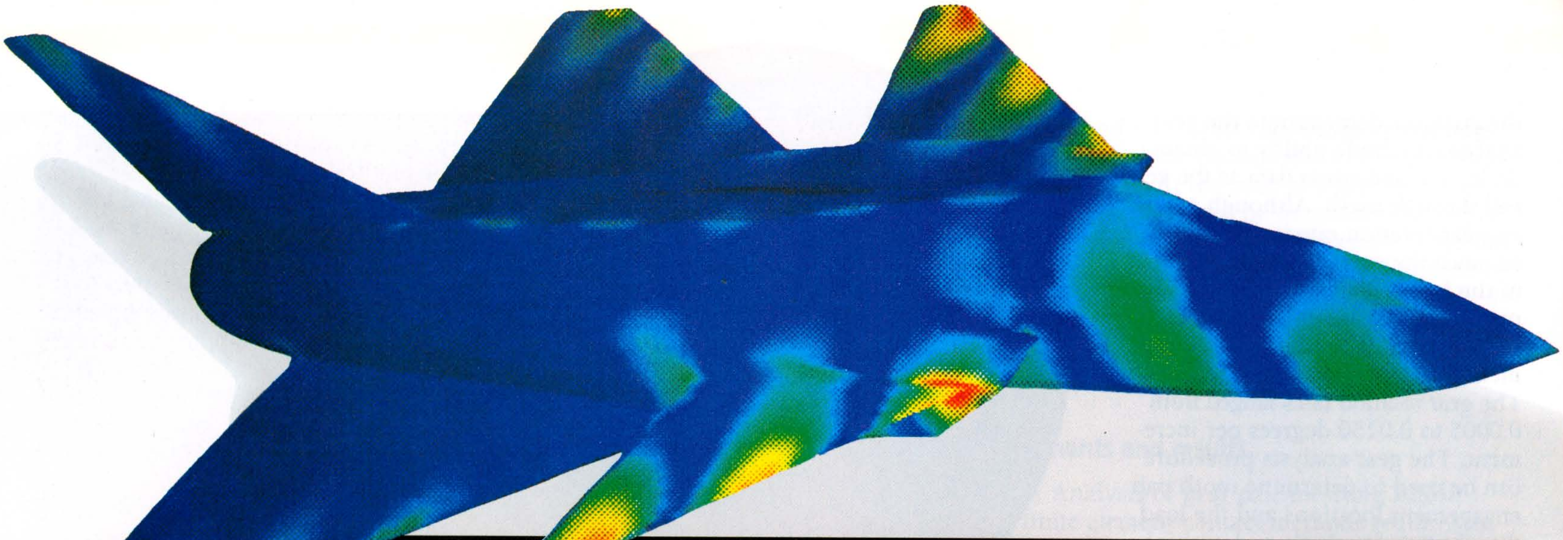

REINVENTING ELECTROMAGNETICS

New supercomputing solutions to Maxwell's equations

Allen Taflove, Northwestern University, Evanston, Illinois

Figure 1 (above). Snapshot of the induced surface electric currents on the VFY-218 prototype Lockheed fighter aircraft for an illuminating radar frequency of 100 MHz at nose-on incidence. This model was implemented using a beta version of the Cray Research FD-TD software, EMDS (ElectroMagnetic Design System).

Until 1990, defense requirements for low radar cross section aerospace vehicles drove the development of large-scale methods in computational electromagnetics. Interest in finite-difference time-domain (FD-TD) and other direct space-grid Maxwell's-equations solvers for this purpose has grown to challenge the previously dominant frequency-domain integral equation approaches. For example, at the July 1992 IEEE Antennas and Propagation Society International Meeting in Chicago, 92 papers were presented on various aspects of FD-TD and other space-grid techniques.

In fact, the emergence of supercomputers with throughput in the range of 10 GFLOPS to 1 TFLOPS will permit FD-TD and similar approaches to model the dynamics of billions of field unknowns. Novel FD-TD algorithms incorporating lumped or distributed nonlinear effects over extremely large instantaneous bandwidths may be the key to understanding many types of technology at the core of twenty-first century electrical engineering. These technologies include microwave and millimeter wave antennas, subnanosecond electronic packaging, picosecond transistors, Josephson junctions and electro-optic devices, and femtosecond all-optical logic elements. This is because Maxwell's equations, with nonlinearities and dispersions properly modeled, provide an overarching framework for the physics of electromagnetic wave transport phenomena from direct current to light, and all high-speed devices of interest to modern society have such wave transport behavior as a critical operating factor.

The models discussed here were constructed at Northwestern University and run on CRAY Y-MP8 and CRAY Y-MP C90 supercomputer systems and solved for up to 60 million vector field unknowns in three dimensions.

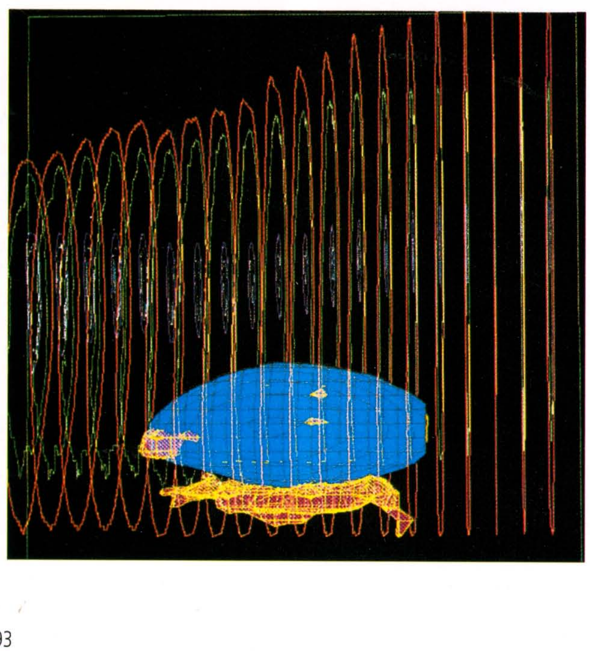

Figure 2. FD-TD-computed isoloci of specific absorption rate (yellow-magenta = 25 percent of peak power absorption) and induced temperature (blue= 42°C) in a patient-specific model of the human thigh exposed to a waveguide hyperthermia applicator at 918 MHz. (No water bolus is used between the waveguide and the thigh.) A side view of thigh is shown, with the stack of CT planes seen edge-on.

Radar cross section modeling

It currently is feasible to embed an entire jet fighter within an FD-TD space grid to compute its induced surface electric currents and radar cross section for radar frequencies up to at least 500 MHz. For example, Figure 1 is a snapshot of the induced surface electric currents on the VFY-218 prototype Lockheed fighter aircraft for an illuminating radar frequency of 100 MHz at nose-on incidence. This model was implemented using a beta version of Cray Research's EMDS (ElectroMagnetic Design System) FD-TD software, which the Northwestern group helped develop. EMDS incorporates the Lockheed ACAD computer-aided design software to enable engineers to generate complex aerospace structures. The software automatically generates a conformal smooth-surface electromagnetics model of the structure using a highly structured finite-difference mesh. Cray Research's MPGS software, which is integrated into EMDS, provides the color visualization of the computed surface currents.

Bioelectromagnetic systems

FD-TD Maxwell's-equations solvers are being applied extensively in clinical settings to help design electromagnetic hyperthermia therapies. This technology uses electromagnetic wave absorption at radio, ultra-high, or microwave frequencies to heat cancerous tumors inside the human body, rendering them more vulnerable to ionizing radiation or chemotherapy. The Northwestern group has helped pioneer electromagnetic field hyperthermia therapies tailored to individual patients. The group uses computed tomography (CT) imaging to establish a three-dimensional dielectric medium database for the FD-TD solver unique to each patient's tissue structure,[1] thereby modeling the field physics unique to the patient's tissue geometry and selection of electromagnetic applicators. An example of this work is shown in Figure 2, which depicts the FD-TD-computed absorbed microwave power distribution and induced temperatures in a CT-generated patient-specific model of the human thigh for a waveguide hyperthermia applicator at 918 MHz.

Packaging and metallic interconnect design for digital circuits

The area of packaging and metallic interconnect design for digital circuits involves engineering problems in the propagation, crosstalk, and radiation of electronic digital pulses. This area has important implications in the design of the multilayer circuit boards and multichip modules widely used in modern digital technology. Existing computer-aided circuit design tools can be inadequate when digital clock speeds exceed about 250 MHz. These tools cannot handle the physics of UHF/microwave electromagnetic wave energy transport along metal surfaces such as ground planes, or in the air away from metal paths, that predominate above 250 MHz. Electronic digital systems develop substantial analog wave effects when clock rates are high enough, and full-vector (full-wave) Maxwell's-equations solvers become necessary to understand these effects.

In perhaps the most complex modeling of these effects so far, the Northwestern group constructed three-dimensional FD-TD models of subnanosecond digital pulse propagation and crosstalk behavior in a module consisting of a stack of four 22-layer circuit boards connected by three 50-pin connectors.[2] A uniform resolution of 0.004 inch permitted each layer, via, and pin of the circuit boards and connectors to be modeled. A maximum of 60 million vector field unknowns was solved per modeling run.

Figure 3(a) shows the plan view of an outwardly propagating electromagnetic wave within the top circuit board of the stack generated by the passage of the pulse down the via pin. Although the relatively intense magnetic field (shown by the yellow color) adjacent to the excited via is quite localized, moderate magnetic fields (shown by light blue) emanate throughout the entire transverse cross section of the board and link all of the adjacent via pins, shown as dark dots in a diamond pattern. Figure 3(b) depicts the early-time coupling of magnetic fields from the excited via pin to the adjacent unexcited via pins as seen in a vertical cut through the top 22-layer board and connector of the stack.

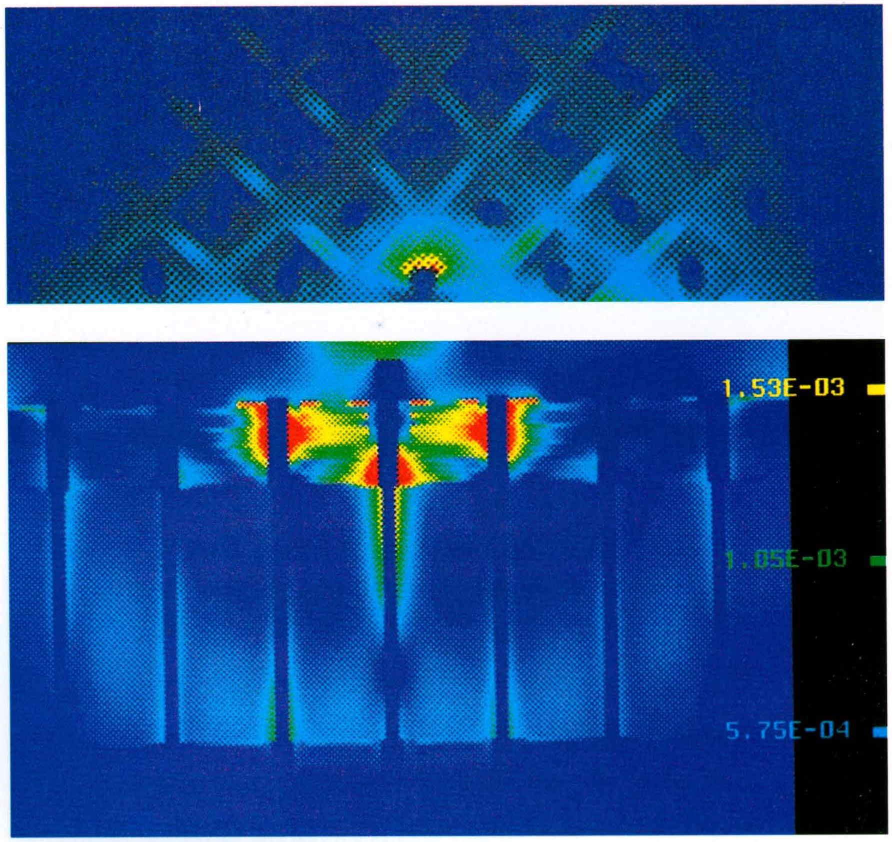

Figure 3. FD-TD-computed electromagnetic wave effects in the top 22-layer circuit board of the four-board/three-connector stack generated by the passage of a subnanosecond pulse down a single vertical via pin. (a) Plan view of outwardly propagating electromagnetic wave within the board. Color scale: yellow = maximum; light blue = moderate; dark blue = negligible. (b) Early-time coupling of magnetic fields from the excited via pin to the adjacent unexcited via pins as seen in a vertical cut through the top board and connector. Color scale: red = maximum; yellow = moderate; green = low-level; dark blue = negligible.

Figure 4 shows the magnitude and direction of late-time currents flowing along the vertical cross section of the four-board/three-connector stack for a subnanosecond digital pulse assumed to excite a single vertical via pin in the top 22-layer board. The currents were calculated in a postprocessing step by numerically evaluating the curl of the magnetic field obtained from the three-dimensional FD-TD model. Red denotes downward-directed current, while green denotes upward-directed current. At the time of this visualization, current had proceeded down the excited via through all four boards and all three connectors. Upward-directed

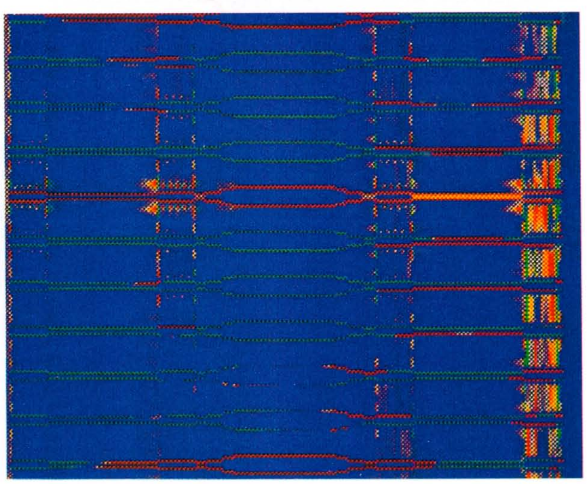

Figure 4. FD-TD-computed magnitude and direction of late-time currents flowing along the vertical cross section of the complete connector module for a subnanosecond digital pulse assumed to excite a single vertical via pin in the top 22-layer board. Color scale: red = downward-directed current; green = upward-directed current; dark blue = negligible.

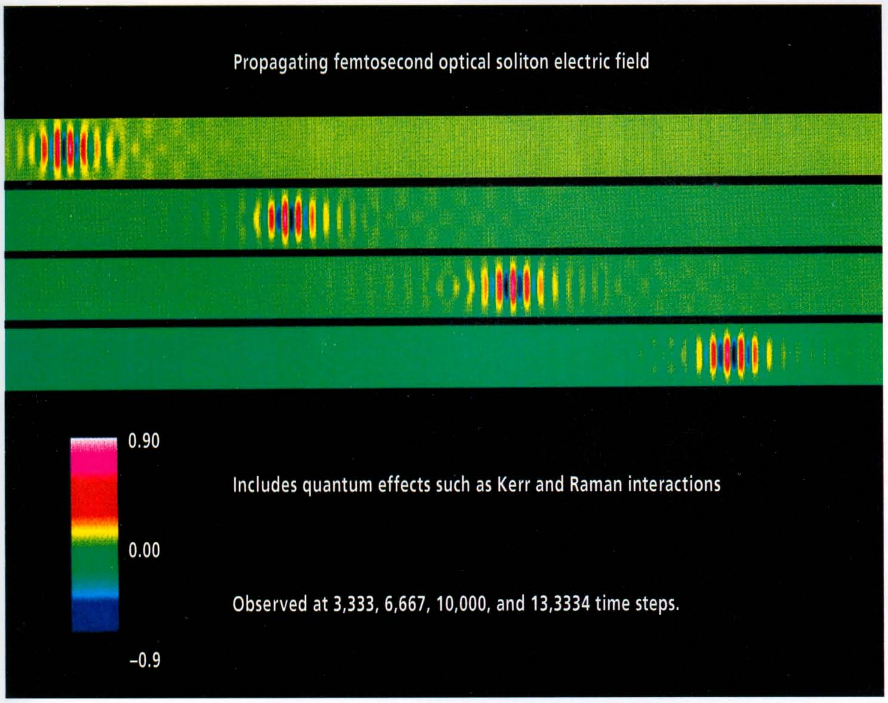

Figure 5. FD-TD-computed electric field of an initial 30-femtosecond duration optical carrier pulse (wavelength = 2.19 microns) propagating in a 1-micron thick nonlinear dispersive dielectric waveguide, showing the formation of a temporal soliton.

current, however, is seen to flow on the adjacent vias. This represents an undesired ground-loop coupling to the digital circuits using these vias that is capable of jamming the operation of these circuits.

This work is leading to the direct time-domain Maxwell's-equations modeling of the metallic interconnects and packaging of general-purpose digital circuits operating at clock speeds above 250 MHz. The example clearly shows that the analog coupling effects for such devices can be so complex that there may be no way to design them and make them work in a timely and reliable manner without such modeling.

Application to all-optical devices

In the last two years, the Northwestern group, working with Peter Goorjian of the NASA-Ames Research Center, has pioneered FD-TD solutions of the vector nonlinear Maxwell's equations for femtosecond optical soliton propagation and scattering, including carrier waves, in one and two dimensions.[3,4] Our FD-TD approach efficiently implements second-order (Lorentzian) linear and nonlinear dispersive convolutions for the electric polarization, with the nonlinear convolution accounting for two key quantum effects, the Kerr and Raman interactions. By retaining the optical carrier, the new method solves for fundamental quantities—the optical, electric, and magnetic fields in space and time—rather than a nonphysical envelope function as did all previous approaches.

In this spirit, the researchers present first-time calculations from the vector nonlinear Maxwell's equations of temporal optical solitons propagating in a two-dimensional dielectric waveguide.[4] Here, a dielectric waveguide 1 micron thick and surrounded by air is considered. The calculations are for pulses with a carrier frequency of 1.37 x 10^{14} Hz (wavelength = 2.19 microns) and an initial width of only about 30 femtoseconds. (One femtosecond = 1 millionth of one nanosecond). Figure 5 depicts snapshots of the FD-TD-computed electric field for such a pulse as it travels down the dielectric waveguide. The dielectric is assumed to have anomalous linear dispersion due to a single Lorentzian relaxation and a dispersive nonlinearity due to the Kerr and Raman quantum interactions. Figure 5 clearly shows that the optical pulse retains its modal distribution as it propagates in the nonlinear dispersive waveguide, a property characteristic of a temporal optical soliton. Numerical experiments indicate that this is due to a balance of nonlinearity and dispersion; slightly reducing the dielectric nonlinearity causes the computed pulse to spread out progressively and attenuate as it propagates.

First-time calculations from the vector nonlinear Maxwell's equations of light switching light are presented last.[5] The group uses FD-TD to model spatial soliton propagation and scattering, including carrier waves, in a two-dimensional homogeneous nonlinear dielectric medium. Figure 6 shows the FD-TD-computed electric field of two initially parallel, equal-amplitude optical carrier beams co-propagating in a two-dimensional homogeneous nonlinear medium. Numerical experiments indicate that one such beam propagating alone retains its transverse profile due to self-focusing effects caused by the nonlinear medium, a property characteristic of a spatial optical soliton. Two co-propagating spatial solitons mutually interact via alternating attraction and repulsion and separate at a substantial deflection angle. After spatial filtering of either soliton, this deflection mechanism can provide the basis for a fast-acting photonic (all-optical) switch. The Northwestern group is investigating just how fast this switching action might ultimately be and whether the power

Figure 6. FD-TD-computed electric field of two initially parallel, equal-amplitude optical carrier beams co-propagating in a two-dimensional homogeneous nonlinear medium. Note the mutual interaction of the beams (spatial solitons) via alternating attraction and repulsion, followed by separation at a substantial deflection angle. This illustrates the phenomenon of light switching light.

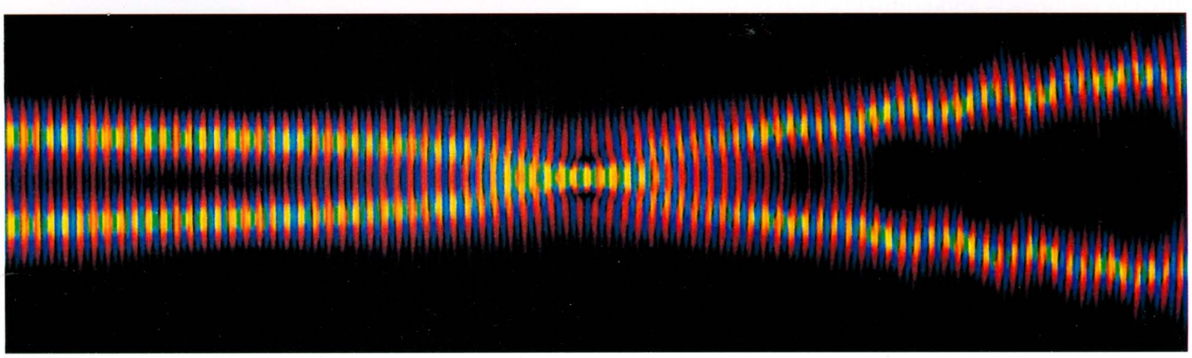

in one beam can be reduced well below that of the other. The ultimate goal would be femtosecond-regime angular deflection (switching) accompanied by optical gain. This would provide the basis of a femtosecond all-optical transistor.

The Northwestern group's novel approach to computational nonlinear optics achieves robustness by rigorously enforcing the vector field boundary conditions at all interfaces of dissimilar media in the time scale of the optical carrier, whether or not the media are dispersive or nonlinear. As a result, the new approach is almost completely general. It assumes nothing about the homogeneity or isotropy of the optical medium, the magnitude of the nonlinearity, the nature of the material's dispersive properties, or the shape, duration, or vector nature of the optical pulse(s). It has the potential to provide an unprecedented two- and three-dimensional modeling capability for millimeter-scale integrated optical circuits having submicron-engineered inhomogeneities.

If successful, this work may lead to the modeling of all-optical digital logic devices switching potentially in 10 to 50 femtoseconds at room temperature. This is about 1000 times faster than the best transistor today and 100 times faster than a Josephson junction operating under liquid helium. The implications may be profound for the realization of "optonics," a proposed successor technology to electronics in the twenty-first century that would integrate optical fiber interconnects and all-optical processors into systems of unimaginable information processing capability.

Conclusions

Supercomputers of the late-1990s, which promise to achieve rates from 0.1 to 1 TFLOPS, will permit us to attack some grand challenges in electromagnetics. One such challenge remains from the radar cross section technology side—the airplane-in-the-grid. In fact, using this new class of supercomputers and the new class of space-grid time-domain Maxwell's solvers, it certainly will be possible to obtain whole stealth fighter models in the 1 to 3 GHz range, and jet-engine-inlet models up to perhaps 5 to 10 GHz, with predictive dynamic ranges up to 70 dB.

Moreover, the same algorithms implemented on the same computers could attack other grand challenges in electromagnetics:

- ☐ Patient-specific, whole-body electromagnetic hyperthermia cancer therapies
- ☐ Sub-nanosecond passive metallic interconnects, packaging, and multi-chip modules for electronic digital circuits
- ☐ Logical operation of transistors and other active digital elements at clocks above 1 GHz self-consistently incorporating device nonlinearities and electromagnetic effects
- ☐ Novel picosecond transistors exploiting interactions of electromagnetic and charge waves
- ☐ Novel femtosecond all-optical logic elements or even more exotic nonlinear wave species yet to be discovered

In fact, the ultra-large-scale solution of Maxwell's equations using time-domain grid-based approaches may be fundamental to the advancement of technology as we continue to push the envelope of the ultra-complex and the ultra-fast. Simply speaking, Maxwell's equations describe the physics of electromagnetic phenomena from direct current to light, and accurately modeling them is essential to understand high-speed signal effects exhibiting wave transport behavior. The goal is the computational unification of full-vector electromagnetic waves; charge transport in transistors, Josephson junctions, and electro-optic devices; surface and volumetric wave dispersions (including those of superconductors); and nonlinearities due to quantum effects. Then we can attack some computational grand challenges to markedly advance electrical engineering and directly benefit society.

Acknowledgments

The author thanks Evans Harrigan and John Ahnert of Cray Research; Peter Goorjian of the NASA-Ames Research Center; V. Sathiaseelan of Northwestern Memorial Hospital; graduate students Rose Joseph, Dan Katz, Melinda Piket-May, Chris Reuter, and Eric Thiele; and undergraduate research students Peter Anvin and Susan Hagness. Support for this work was provided by NASA-Ames University Consortium Joint Research Interchange Grants NCA2-561, 562, and 727; ONR Contracts N00014-88-K-0475 and N00014-93-0133; Northrop Electronics Systems Division; and Cray Research, Inc.

About the author

Allen Taflove is a professor of electrical engineering and computer science at Northwestern University McCormick School of Engineering. In 1990, he was named an IEEE Fellow for his pioneering FD-TD methods in computational electromagnetics since 1973.

References

1. Piket-May, M. J., A. Taflove, W.-C. Lin, D. S. Katz, V. Sathiaseelan, and B. Mittal, "Initial results for automated computational modeling of patient-specific electromagnetic hyperthermia," *IEEE Transactions on Biomedical Engineering*, Vol. 39, pp. 226-237, March 1992.
2. Piket-May, M. J., A. Taflove, and J. Baron, "FD-TD modeling of digital signal propagation in 3-D circuits with passive and active loads," submitted to *IEEE Transactions on Microwave Theory and Techniques*, January 1993.
3. Goorjian, P. M., A. Taflove, R. M. Joseph and S.C. Hagness, "Computational modeling of femtosecond optical solitons from Maxwell's equations," *IEEE Journal of Quantum Electronics*, Vol. 28 (Special Issue on Ultrafast Optics and Electronics), pp. 2416-2422, October 1992.
4. Joseph, R. M., P. M. Goorjian, and A. Taflove, "Direct time integration of Maxwell's equations in two-dimensional dielectric waveguides for propagation and scattering of femtosecond electromagnetic solitons," *Optics Letters*, Vol. 18, pp. 491-493, April 1993.
5. Joseph, R. M., A. Taflove, and P. M. Goorjian, "Calculations of photonic switching in a nonlinear medium with the vector Maxwell's equations," paper submitted to *Optics Letters*, March 1993.

Command and control modeling with FLAMES

Kenneth D. Watts and Michael D. Cash, Ternion Corporation, Huntsville, Alabama

Modern military forces have fielded many new weapon, sensor, communication, and support systems that make use of highly advanced technologies. Efficiently and effectively harnessing the capabilities of these systems is primarily a command and control (C^2) problem. Thus, C^2 doctrine development and training have emerged as major challenges to the military forces of the 1990s. These challenges are exacerbated with the reduction in military resources and with the military threat changing from a well known, established force to many smaller forces that are geographically dispersed, less predictable, and more volatile.

In the past, C^2 doctrine and procedures often were developed and studied in large field exercises where commanders simulated a battle using real assets. Funding reductions have done away with this concept and have greatly increased the role of simulation in new concept development, system effectiveness analyses, and employment training. As a result, C^2 simulation must now represent more systems with greater fidelity to analyze the many intricate relationships that influence employment of high technology weapons in remote, diverse environments. Past simulations attempted to address only small, isolated areas of military performance and were not designed to address the broad and rapidly changing areas faced by today's military forces. Modern C^2 simulations must utilize a new design paradigm—a software system that accounts for current and future hardware capabilities—if the simulations are to respond to complex and continually evolving military requirements.

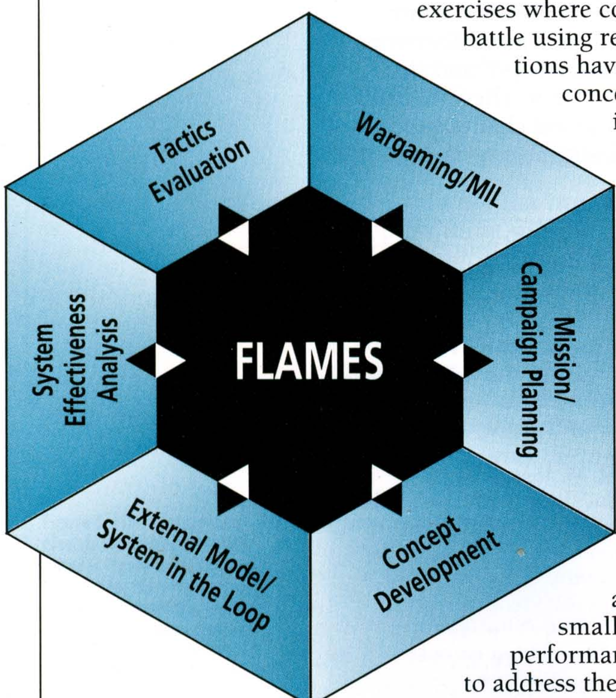

The Force Level Analysis and Mission Effectiveness System (FLAMES) is a simulation environment designed and developed using this new paradigm. From the outset, FLAMES has been constructed as an object-oriented modeling framework that can represent the full theater environment, with its diverse weapons and systems, as well as the human C^2 role. Knowing that the human role can never be considered as simply a series of fixed responses, FLAMES incorporates an innovative and unique technique for representing the human situation perception and C^2 actions. This technique allows rapid C^2 modeling modification by the user and for the insertion of human controllers into the simulation to play the role of any of the thousands of units being modeled in the theater scenario. FLAMES' advanced object-oriented design also supports the integration of existing models into its framework.

Cray Research systems networked with low cost graphics workstations provide the ultimate execution platform for FLAMES. Memory capacities, I/O bandwidth, and an orientation toward multiuser operations give the FLAMES user an almost dimensionless capacity designed specifically for computationally intensive tasks.

FLAMES as a simulation framework

FLAMES provides a common modeling environment capable of representing tactical land, sea, air, and space combat activities in theater-sized scenarios—each with its unique equipment, environments, C^2 architectures, information flow, and human decisions and actions. Flexible analysis tools and a man-in-the-loop (MIL) interface enable FLAMES to support a broad range of user requirements.

FLAMES is actually a family of programs. The principal programs include FORGE, the FLAMES scenario creation tool; FIRE, the scenario execution tool; and FLARE, a postprocessing tool. With FORGE, users create units, configure them with equipment, provide them with movement capabilities and signature characteristics, assign them to a task organization, give them missions, and define their C^2 responsibilities. Scenarios created by FORGE are stored in the FLAMES database, a sophisticated, custom database management system (DBMS) developed specifically for FLAMES to manage large amounts of scenario data shared among numerous users. The core program in the FLAMES family is FIRE. FIRE executes a scenario created with FORGE, allows runtime updates to that scenario, and supports interfaces for external consoles and models. Data that is output by FIRE is used by FLARE, the postprocessing tool. FLARE uses a subset of the industry standard structured query language (SQL) to let the user specify the data calculations and the format of the results. Additionally, FLAMES provides a graphical playback tool with flexible two- and three-dimensional views to help the user visualize resultant scenario activities, determine time periods, and identify units and activities for further data analysis using FLARE.

The FLAMES design also addresses the model developer. Within FLAMES there are essentially two parts. The first is a software kernel that provides basic simulation and modeling functionality, and the second part provides specific simulation models. The kernel maintains the simulation calendar, provides user data management, controls the execution of all internal application models, handles file management, and supports interfaces for external consoles and models. The object-oriented implementation of the kernel establishes the FLAMES class hierarchy and structure and therefore completely describes the interface for the model developers, allowing them to quickly and easily incorporate new models into the FLAMES framework. Models provide basic equipment functionality, human decision making capabilities, communication capabilities, and C2 functionality.

Technological advances

One of the most innovative aspects of the FLAMES design is its representation of human decision making. The model developer can create models that represent individual activities performed by a human. The FLAMES user can select the desired models, associate specific functions with them, and sequence them into threads of execution. Each FLAMES unit can have an unlimited number of these threads executing in parallel to represent an individual task or decision sequence that a human operator would perform. FLAMES' method for binding the models to the functions and then to each other is innovative as well. FLAMES allows the user to establish this binding in data. The binding does not occur at compile time or at runtime, but is specified in the scenario configuration established by users to represent their problems. Therefore, FLAMES' C^2 functionality can be completely reconfigured with no source code modifications. Since the binding is specified in data, FLAMES allows the user to alter these bindings and threads during the execution of a scenario.

FLAMES is the first simulation and modeling framework of its kind with a complete object-oriented design and implementation. Its implementation makes use of all of the primary aspects of the object-oriented paradigm. FLAMES supports inheritance, modularity, encapsulation, abstraction, and extensibility. The well established, robust class hierarchy within FLAMES provides the structure required for modeling all aspects of modern, combined arms exchanges between forces at all levels.

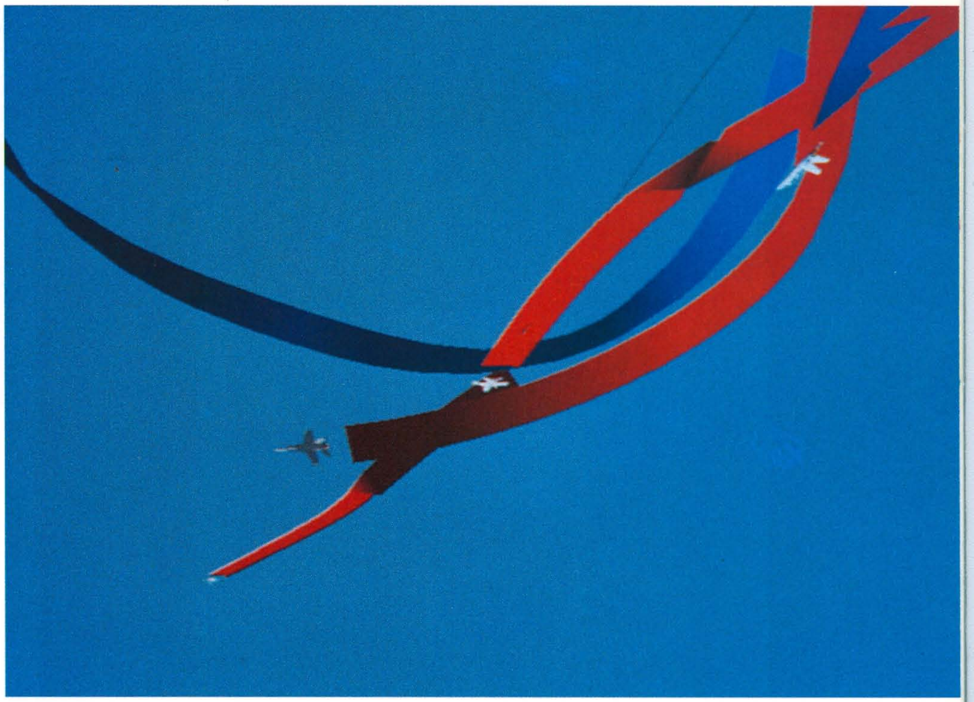

Figure 1. Defensive counter air operations, as simulated by FLAMES.

FLAMES makes extensive use of relational database management system (RDBMS) constructs in both traditional and nontraditional fashion. RDBMS technology has been incorporated directly into the FLAMES kernel for use by models, allowing full scale, relational tables to be created, updated, manipulated, and queried—all with standard SQL. However, in FLAMES, standard SQL has been extended to allow a developer-defined function to manipulate the data gathered via the query before returning control to the model. The combination of these powerful techniques provides more functionality and allows the model developer to write less code.

Current applications

FLAMES can create an explicit representation of individual simulated units and exercise it in a full theater context with all of its counterparts. Friendly and opposing forces can be configured to represent their order of battle, geographic deployment, and C^2 capabilities. Defensive aircraft, for example, can be positioned to defend friendly airspace and maintain air superiority. The air defense C^2 architecture, long-range surveillance, fighter configuration, stationing, and tactics can be examined in realistic detail to ensure that air superiority can be achieved and maintained. Every aspect of thousands of engagements, from flight tactics to individual aircraft maneuvering, can be examined in detail, as demonstrated in Figure 1. Likewise, the C^2 activities for offensive strike missions can be

Figure 2. Offensive air strike capabilities, as simulated by FLAMES.

analyzed. These activities include target identification and location, strike package tasking, in-flight C^2, threat avoidance, target updates, electronic combat, aircraft threat reactions, target engagements, and post-strike damage assessment, as shown in Figure 2. The full theater complement of land, sea, air, and space-based units can be represented within the same common framework.

FLAMES currently is being used or evaluated for several types of analysis. One of the largest FLAMES applications represents a current real-world contingency theater scenario. The software is used to evaluate the contribution of various real-time and near-real-time intelligence-collection systems to theater combat air operations. FLAMES is being used to evaluate alternative concepts for both air- and space-based sensor systems to collect target and threat information, rapidly process it at different C^2 echelons, and disseminate it to operational units, including en route strike aircraft. By modeling all units involved in this global C^2 architecture, FLAMES provides a means of evaluating these concepts with measures of effectiveness that include target damage and losses to friendly aircraft. To capture a real-world baseline scenario for these evaluations, FLAMES was deployed to a theater command center and operated throughout a multiday exercise. FLAMES operated continuously in this environment, with new units, equipment, and missions being added to the scenario, in concert with the battle staff as they performed the C^2 actions in the exercise.

Other FLAMES applications include prototyping and evaluating candidate C^2 consoles for a large air command and control system. FLAMES is used to represent the C^2 echelons from the commander of all air forces down to the individual aircraft and its parent squadron. Prototype consoles at key C^2 echelons are built on graphics workstations and connected to FLAMES to execute in a long duration theater scenario. Techniques for filtering and displaying data to battle staff officers and eliciting their input will be comparatively analyzed. Methods for determining the human versus computer-automated role in short timeline decision making for theater missile defense are analyzed as part of this effort. Additionally, FLAMES is used to evaluate sensor configurations and C^2 systems for a shipborne terminal defense system design. Several ongoing FLAMES applications have specific system contributions being measured in the context of full theater threat environments.

FLAMES in the supercomputing environment

Ternion Corporation has ported FIRE, the runtime program in the FLAMES family, to the CRAY Y-MP system. Initial benchmarking has demonstrated a remarkable improvement in runtime and throughput efficiency. Because C^2 scenarios are a combination of a number of modeled units, the interactions among the units, and the geographic volume in which they operate, it is difficult to establish a single C^2 modeling benchmark scenario. Additional improvement in runtime is expected as Ternion engineers work with Cray Research analysts to further optimize the appropriate portions of FLAMES.

As the number of FLAMES users increases, the true benefit of its environment continues to become more obvious. The ability to run diverse models and scenarios in a common environment and to share them among development sites is a capability that military analysts have demanded but never had. The stopgap measures to interface outdated and dissimilar simulations to create multiservice, theater C^2 models by exchanging huge amounts of generic data can now be abandoned. The FLAMES development environment provides a robust, object-oriented implementation into which the core models of these outdated simulations can be integrated far more efficiently. Because FLAMES runs on the most powerful Cray Research supercomputers, users can access the computational power to tackle theater C^2 analytical problems that once were unsolvable.

About the authors

Kenneth D. Watts is a Lt. Colonel in the U.S. Marine Corps Reserves with extensive joint service experience. While on active duty, he served as an airborne ECM Officer and Reconnaissance Systems Officer flying in the EA-6A and RF-4B aircraft. He has participated in numerous live and command post exercises and has served on the Operations Department (J3) staff at the Headquarters, U.S. Atlantic Forces (LANTCOM). He has an M.S. degree in electrical engineering from the U.S. Naval Postgraduate School and has managed projects which adapted existing simulation tools to model large theater level scenarios.

Michael D. Cash is a graduate of Louisiana Tech University with a B.S. degree in electrical engineering. He has 14 years of experience in the design, development, and management of large-scale force-on-force simulations. His software experience ranges from real-time man-in-the-loop simulators and trainers to graphics intensive user interface applications to embedded avionics programs.

Simulating current collection by an electrodynamic tether in space

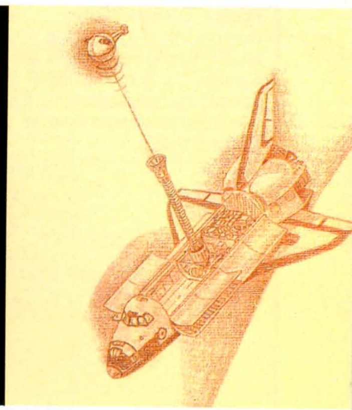

Above, the tethered satellite system showing the shuttle orbiter, the tether, and the satellite.

Nagendra Singh, Bharat Ishwar Vashi, Wing-Ching Leung, Department of Electrical Engineering, The University of Alabama in Huntsville

In the summer of 1992, the U.S. National Aeronautics and Space Administration (NASA) conducted its first experiment on the Tethered-Satellite System (TSS-1) in low earth orbit. The goal was to develop a better understanding of the circuit consisting of the shuttle, tether, tethered satellite, and the space plasma. However, when technical difficulties prevented the tether from extending to its full length of 20 km above the shuttle, the system's total voltage fell short of expected levels, and the tether's operation was hampered.

Researchers at the University of Alabama in Huntsville set out to develop a better understanding of the electrical contact between the tethered satellite and the space plasma. The study of satellite-plasma interaction is analytically intractable. By developing their own three-dimensional simulation model based on new algorithms for solving Poisson's equation, the researchers successfully reduced—by 33 percent in one instance—the CPU time necessary for the simulation. The simulation, in turn, yielded important results about the satellite's electron-collection process and the impedance between the satellite and plasma.

Satellite-plasma contact impedance

The basic principle of an electrodynamic tether is that when a conducting wire cuts across magnetic field lines, an electromotive force is generated. If the wire is a part of a circuit, electrical power can be delivered to a load, or, conversely, thrust can be generated for propulsion by using a power source to pass a current through the wire. Operation of an electrodynamic tether depends on the closure of an electrical circuit. It is believed that the circuit closes through the space which is full of plasma having a density of about 1×10^{11} particles per cubic meter and a temperature of about 2000 K. The plasma is an ionized gas, which can conduct electricity.

In the tether project, the tethered satellite attains a positive potential with respect to the space plasma. The magnitude of this potential depends on the overall performance of the circuit, including the impedances in the circuit and the current-voltage characteristic of the satellite-plasma contact. Because the satellite is positively charged relative to the plasma, it collects electrons from the plasma. The basic question is, how much electron current is collected for a given potential on the satellite?

Three-dimensional model of satellite-plasma interactions

The researchers' three-dimensional model consists of a spherical satellite at the center of the simulation system and plasma reservoirs at the boundaries. The satellite could be of any shape or size; these simulations are limited to spherical and cylindrical shapes. The satellite is at a potential Φ_o with respect to the boundaries. The plasma flow is treated by the particle-in-cell (PIC) method,[1] according to which plasma is modeled by a large number of computational macroparticles. These particles carry charges much larger than those of a real electron and ion, and the appropriate scaling is obtained by balancing the real and computational charges in the system. The flow of plasma is advanced in time by solving the equation of motions of the charged macroparticles in self-consistent fields obtained by solving Poisson's equation. Because solving Poisson's equation in three dimensions is the most CPU-consuming aspect of the numerical model, a brief discussion of the algorithms for solving it is appropriate.

In the cylindrical coordinate system (r, θ, z) Poisson's equation can be written as

$$\frac{1}{r}\frac{\partial}{\partial r}\left(r\frac{\partial\Phi}{\partial r}\right)+\frac{1}{r^2}\frac{\partial^2\Phi}{\partial\theta^2}+\frac{\partial^2\Phi}{\partial z^2}=-\frac{\rho}{\varepsilon_o} \qquad (1)$$

where Φ is the potential and ρ is the charge density, which is obtained by sharing the charges of the macroparticles on grid points according to a volume-weighing method.[1] The researchers developed two

Figure 1. Current-voltage characteristic of a cylindrical satellite having equal length and diameter D = 30 cm. Numerical currents are compared with theoretical predictions from an idealized theory.

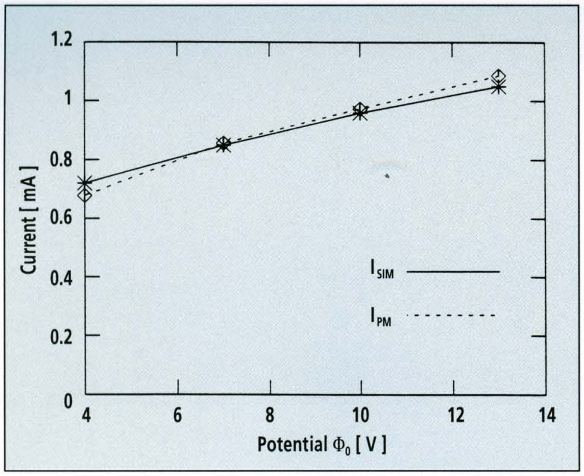

main algorithms for solving equation (1). In one algorithm they directly convert the equation into a difference equation by using a seven-point differencing scheme. The resulting set of linear equations can be written as$[[A]][\Phi]=[\rho]$. The matrix $[[A]]$ turns out to be a sparse matrix, and the solution of the difference equation is obtained by the preconditioned conjugate-gradient method.[2]

An alternative to this direct solution is to utilize the problem's azimuthal symmetry. Fourier transforms for Φ and ρ with respect to the azimuthal angle θ can thus be performed. This reduces the three-dimensional Poisson's equation to a set of two-dimensional equations for each azimuthal harmonic.[3] The latter equations are solved by converting them into difference equations using the preconditioned conjugate gradient method. After obtaining the solutions for the azimuthal harmonics of Φ, an inverse Fourier transform is performed to obtain $\Phi(r,\theta,z)$.

The researchers performed runs using both the direct and the Fourier algorithms. A comparison, based on 60,512 grid points, showed that the algorithm based on the Fourier transform saved about 33 percent CPU time on the CRAY X-MP/24 supercomputer.

Code validation

To check the three-dimensional code's validity, researchers performed several simulations adopting the ideal conditions for which an analytical theory has been determined.[4] They found an excellent agreement between the theoretical and the numerical predictions on the time-averaged collection of electrons. Figure 1 shows the comparison between the theoretical and numerical electron currents collected by a cylindrical satellite of equal height and diameter D = 30 cm and shows the current-voltage characteristic of the satellite. However, the current from the simulations showed slow oscillations determined by the ion transit time from the simulated plasma sources to the edge of the positive potential sheath around the satellite. Such oscillations are artifacts of simulating plasma sources.

Highlights of new results

To develop an understanding of the current collection process, researchers performed several simulations that have been analyzed in detail.[3] When the orbital motion is included in the simulations, the sheath structure around the satellite extends along the magnetic field. An example of sheath structure around a spherical satellite is shown in Figure 2, which gives the equipotential surfaces in the x - y, x - z, and y - z planes of the simulation system, respectively. In

Figure 2. Potential distributions in the sheath around a spherical satellite of radius 25 cm and at a potential of 10 V in (a) x - y, (b) x - z, and (c) y - z planes passing through the center of the satellite.

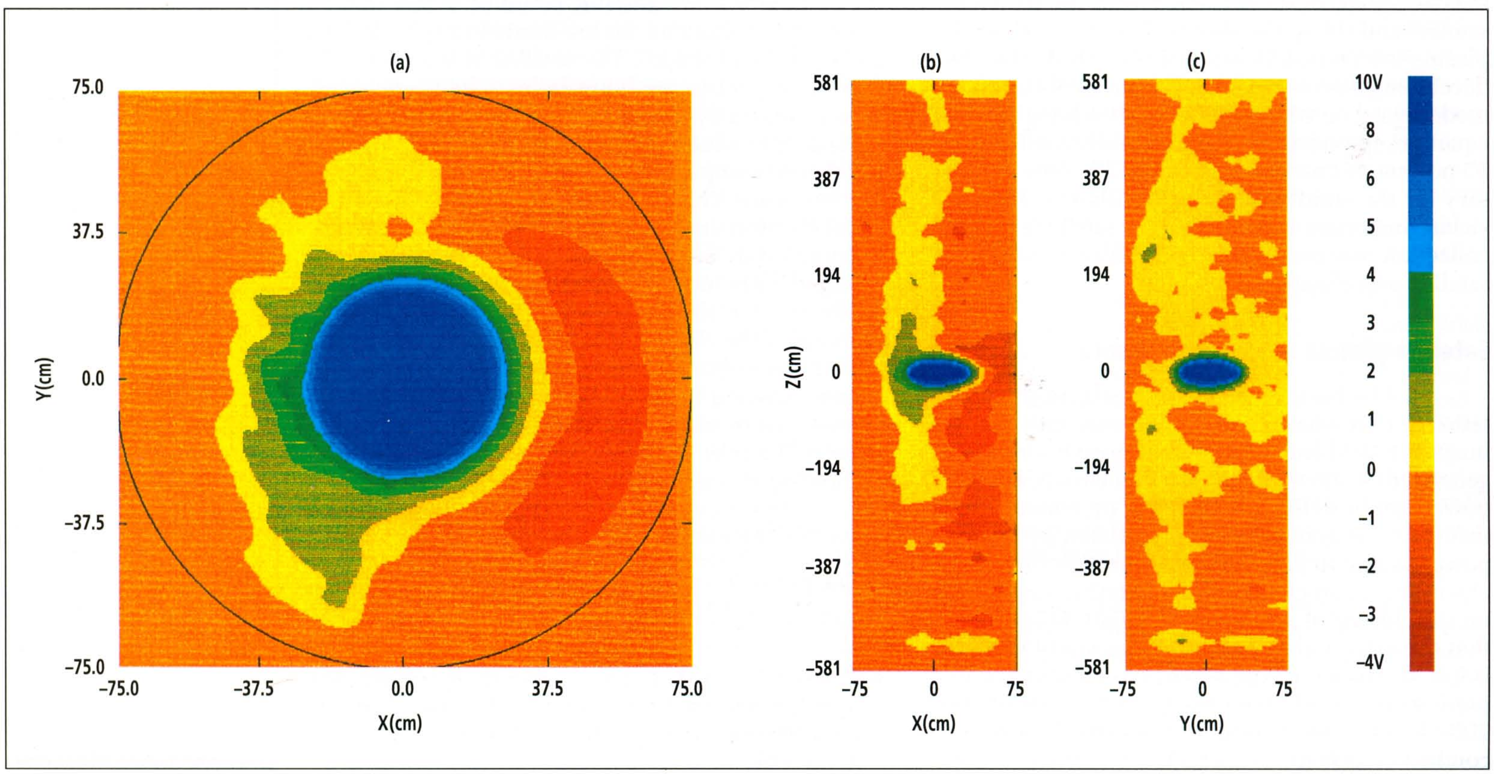

this example the satellite's diameter is 50 cm, and the satellite is at 10 V. Figure 2(a) clearly shows the ram-wake structure. In the rest frame of the satellite, the flow is coming from the left. The plasma impinging on the satellite causes the relatively high potentials to extend over a larger volume in front of the satellite than that in the wake behind it, where the potentials even become negative. The spatial extension of the potential in the ram direction is attributed to a slowing of the ions in the flowing plasma by the positive potential on the satellite.

The sheath structures in the x - z and y - z planes (Figures 2(b) and (c)) show the sheath's extension along the z-axis, along which the magnetic field is oriented. This extension is caused primarily by the combined effects of the electron collection from a magnetic-field-aligned volume, and the slowing of the ions in the flow (coming from the left) by the positive potentials.

The sheath structures shown in Figure 2 have a profound effect on the satellite's collection of electrons. In a nonflowing plasma situation when the orbital motion is ignored, the only source of electrons being collected by the satellite is a cylindrical volume extending along the magnetic field, both above and below the satellite.[4] When the relative flow due to the orbital motion is included, an additional source of electrons for collection is introduced; these are the electrons entering the sheath with the flow. Some of these electrons, when acted upon by the electric fields in the sheath, are collected by the satellite. This additional collection increases the collected current and hence lowers the contact impedance between the plasma and the satellite. Its implication for the tether mission is that a relatively larger part of the generated electromotive force is available for driving other elements of the circuit, including the useful load.

Figure 3 shows the time history of the current collected by the satellite during the course of a simulation. The quasisteady current after $t >$ 25 µs is about 6.5 mA. The simulation without the orbital motion gives a current of about 1.22 mA, in agreement with an idealized theory.[4] The former current is about five times larger than the latter one. Thus the simulations show that the orbital motion has a significant effect on the current collection process and on the impedance between the satellite and plasma.

The three-dimensional results presented here were for an orbital velocity of 60 km/s, which is too large compared to the 8 km/s in the ionosphere. However, simulations with smaller orbital velocities were performed using two-dimensional simulation models,[3] confirming the feature of the current enhancement and the modifications in the sheath structure around the satellite caused by the orbital motion.

Future research

A typical run using the three-dimensional model requires about 30 hours of CPU time when the potential on the satellite is relatively small. For a larger potential on the satellite, a correspondingly larger volume around the satellite must be simulated. This increases CPU usage considerably.

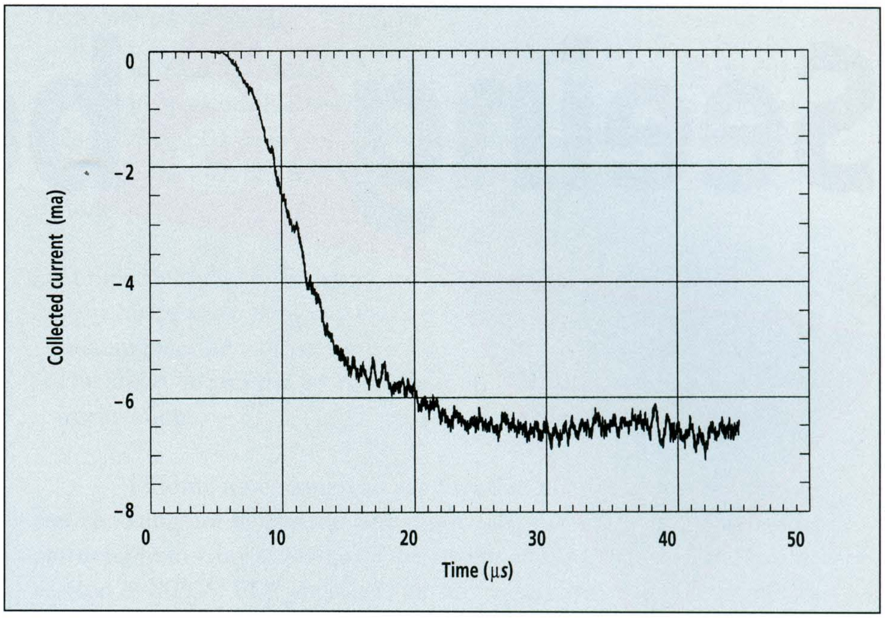

Figure 3. Temporal evolution of current collected by the satellite in the course of a simulation run. The current during the buildup phase is an artifact of initializing the simulation run. The quasisteady current after t = 25 µs is relevant.

For this reason, researchers at the University of Alabama in Huntsville have not performed simulations for large potentials, which are more appropriate for the experiments in the TSS-1 mission. If the TSS mission flies again, researchers hope to perform simulations with large potentials.

About the authors

Nagendra Singh is a professor in the Electrical and Computer Engineering Department of the University of Alabama in Huntsville. He received a B.Tech. degree in electrical engineering from the Indian Institute of Technology, Kanpur, India, and M.S. and Ph.D. degrees in electrical engineering from the California Institute of Technology, Pasadena.

Wing-Ching Leung is a graduate student working toward a Ph.D. degree in electrical engineering from the University of Alabama in Huntsville, from which he received B.S. and M.S. degrees.

Bahrat Vashi is a postdoctoral fellow in the Electrical and Computer Engineering Department of the University of Alabama in Huntsville, from which he received a Ph.D. degree. Vashi received a B.S. degree in electronic engineering from the University of Kent, Canterbury, Kent, England. He received an M.S. degree in electrical engineering from the Florida Institute of Technology, Melbourne.

References

1. Morse, R. L., "Multidimensional Plasma Simulations by Particle-In-Cell Methods," *Methods in Computational Physics,* Vol. 9, pp. 213-239, B. Adler, S. Fernbach, and M. Rotenberg, eds., Academic, New York, 1970.
2. Hestenes, M. R. and E. Stiefel, "Methods of Conjugate Gradients for Solving Linear Systems," *Journal of Research of the National Bureau of Standards,* Vol. 49, pp. 409-436, December 1952.
3. Vashi, B. I., "Numerical Simulations of the Electrodynamic Interactions between a Satellite and Plasma," Ph.D. dissertation, The University of Alabama in Huntsville, September 1992.
4. Parker, L. W. and B. L. Murphy, "Potential Buildup of an Electron-Emitting Ionospheric Satellite," *Journal of Geophysical Research,* Vol. 72, pp. 1631-1636, March 1967.

Announcing release 5.0 of the MPGS engineering postprocessing system

Seeing is believing

Engineering simulations in the automotive, aerospace, and other manufacturing industries often produce large, complex datasets. To postprocess and visualize these datasets, engineers around the world turn to Cray Research's MPGS engineering postprocessing system. The broad capabilities of the MPGS system can guide engineers at each step in the simulation process. Users can interpret their results with a variety of interactive tools and create presentation materials easily with the system's hardcopy and video output capabilities. MPGS is an invaluable postprocessing system for engineers working in structural analysis, computational fluid dynamics, electromagnetics, and other applications.

Capabilities enhance engineering productivity

The MPGS engineering postprocessing system enhances the productivity of engineers by minimizing the time required to analyze and interpret simulation results. Release 5.0 includes many new features critical to engineers, including

- ☐ Data query and FFT analysis
- ☐ Integrated 2-D plotting tool
- ☐ Input and display of experimental data
- ☐ Selective geometry replication
- ☐ Multiple coordinate reference frames
- ☐ Multiple scalars and vectors and computation of new variables

Because it assumes that model geometries are unstructured, MPGS works with a wide range of engineering analysis methods, including finite element, finite difference, finite volume, and boundary element. MPGS is also fully compatible with emerging analytical technologies that generate new geometries and connectivities at each time step.

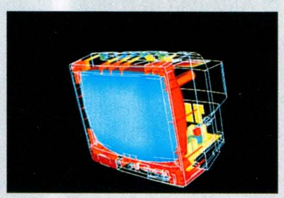

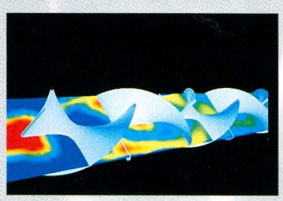

MPGS visualizations, top to bottom:

Airflow around a car body, courtesy Honda Motor Company

Television shock test simulation

Species concentration through a laminar flow static mixer, courtesy Chemineer, Inc.

Versatile interface works with most popular codes

MPGS enhances the value of dozens of widely used analysis codes, including those that provide postprocessing as a secondary function, and supports locally written codes. Because MPGS is the only postprocessing software needed for many popular codes, it saves time that otherwise would be spent learning multiple postprocessors.

In addition to files saved in generic MPGS formats, results files in MOVIE.BYU, N3S, and ESTET formats are supported directly, and MPGS can postprocess data from the following commercial packages:

ABAQUS	KIVA	PAM-CRASH
AIRPLANE	LS-DYNA3D	PATRAN
ANSYS	MARC	PLOT3D
FIDAP	MOLDFLOW	RADIOSS
FLOW3D	MSC/EMAS	RAMPANT
FLUENT	MSC/NASTRAN	STAR-CD

Several commercial analysis codes provide MPGS-readable data as an output option, while other formats may be used by passing the data through translators. For maximum user flexibility in handling unique features of analysis codes, source codes for popular translators are provided (at no charge) wherever possible. The MPGS data format is expressly designed for easy conversion of analysis code data.

Networking and standalone options serve diverse environments

MPGS 5.0 can operate between a Cray Research system and a user's workstation or reside solely on the workstation and perform all operations there. MPGS will be supported on the following workstation platforms:

- ☐ Silicon Graphics
- ☐ IBM RS6000 300, 500, 600, 700 Series, with Sabine, GTO, or GT4 graphics board options
- ☐ Hewlett-Packard 700 series
- ☐ Sun Microsystems SPARC (GS and GT)
- ☐ Digital Equipment/Kubota Alpha AXP/Denali

Ease-of-use features

Since its initial release, Cray Research has refined and enhanced the MPGS system to make it as easy to use as possible. Key MPGS ease-of-use features include:

- ☐ *Dynamic graphical transformations.* MPGS makes examining datasets fast and intuitive by using a workstation's graphics hardware features to perform manipulations directly. Users do not have to wait for time-consuming and network-clogging screen refreshes to rotate, translate, or zoom to areas of interest. Up to four viewports can be used to present multiple views of results simultaneously.
- ☐ *User interface.* Users have several options for controlling graphics transformations: a mouse interface, slider bars, programmable macro keys, and command language input. All menus are ergonomically designed, combining suggestions from interface design specialists at Electricité de France (EDF) and users around the world, with careful adherence to familiar MOTIF standards.
- ☐ *Backup and restore.* An active MPGS session can be saved as a file and restarted at the user's convenience.

Comprehensive postprocessing capabilities

MPGS offers the most complete set of postprocessing features and capabilities available. These include:

- ☐ *Data query and plotting.* MPGS can specify information about all geometric entities and variables used in a simulation. Users can ask MPGS to report the values of specific variables or use the 2-D plotting tool to graph variables across time steps. Variable values along lines or other geometric entities also can be plotted. Fast Fourier transforms allow further analysis of the resulting information. Annotation, curve smoothing, and automatic linear, logarithmic, and quadratic graph scaling capabilities add powerful investigation and presentation capabilities.

Customer partnership sees results

Release 5.0 of the MPGS system was designed in collaboration with engineers from Electricité de France (EDF), the French national power utility. When EDF managers researched postprocessing options for their engineers, they drafted specifications for an organization-wide system that included

- ☐ Compatibility with the many analysis packages used at EDF
- ☐ Distributed processing capabilities between their Cray Research systems and engineering workstations
- ☐ The ability to run the postprocessor on workstations alone or between workstations

Finding no commercial product that met these specifications, and choosing not to develop a postprocessor in-house, EDF proposed a partnership to Cray Research in the spring of 1991 to help develop a new version of MPGS. EDF engineers already were experienced users of Cray Research computer systems and the MPGS postprocessor and determined that MPGS could best be adapted to their needs. At the time the proposal was made, MPGS developers were drafting plans for a new release. In exchange for resource support, Cray Research developers agreed to meet EDF's specifications in the 5.0 release. The partnership accelerated development of the 5.0 release by about 18 months, or about half of the projected development cycle.

EDF formally accepted MPGS 5.0 in July 1993.

MPGS visualizations, top to bottom:

Car crashworthiness model, courtesy Kia Motors Corporation

Pressure on a ducted fan surface, courtesy Allison Gas Turbine Division of General Motors Corporation

Metal tube forming simulation, courtesy Mannesmann Demag

- ☐ *Integrated video recording.* MPGS provides a powerful, easy to use, video recording capability that allows users to save and present animated images on videotape. Because MPGS supports the most popular recording devices, no additional video recording software is required.
- ☐ *Parts.* Model parts can be copied or split in two, with the option to remove an undesired portion. By translating between reference frames, parts can be moved to visually advantageous positions, such as side by side for comparison.
- ☐ *Particle tracking.* Streamlines, streaklines, and pathlines can be tracked through multiple time-steps and variable geometries as needed, and streamlines may be integrated in positive and negative directions. Particle tracks can be displayed in a ribbon format, and animated particles that follow particle tracks can be used to dramatize flow characteristics.
- ☐ *Comparison to measured results.* MPGS lets users visually compare measurements taken from physical models with results from simulations. MPGS uses information from measured results as auxiliary input data that can be queried and plotted, enhancing the program's analysis and presentation capabilities.

The MPGS system's broad range of capabilities makes it easy to see why engineers at many major corporations trust MPGS for their postprocessing and visualization needs. With its broad functionality, interactive tools, and flexible hardware support, MPGS will make a believer out of you.

CORPORATE REGISTER

Demand strong for new CRAY C90 systems

General Atomics ordered a CRAY C98 supercomputing system, to be installed at the San Diego Supercomputer Center in the fourth quarter of 1993. The CRAY C98 system is the eight-processor model in the newly expanded CRAY C90 series, which now includes systems with 1 to 16 processors based on Cray Research's industry leading CRAY C90 technology.

Operated by General Atomics, the San Diego Supercomputer Center is one of four National Science Foundation supercomputing centers. The new CRAY C98 supercomputer will replace the center's CRAY Y-MP system and will be used for a variety of scientific and engineering applications by more than 3000 users at over 100 universities and research institutions affiliated with the center.

"We are very pleased to make Cray Research's newest, most advanced supercomputer technology available to our users," said Sidney Karin, director of the San Diego Supercomputer Center. "Cray Research supercomputing technology historically has been the most usable in the industry. The CRAY C90 is the next generation of what we've come to expect from Cray Research, and we look forward to making use of the advanced technology for important engineering and scientific research work."

The **Instituto Nacional de Meteorología (INM)**, Spain's weather forecasting and climate research center, has ordered a CRAY C94 system. INM will be the first weather organization to install one of the systems in the new CRAY C90 series.

The Space Data and Computing Division of the **NASA Goddard Space Flight Center** ordered a CRAY C98 system to be installed in the second half of 1993 at the NASA Center for Computational Sciences (NCCS) in Greenbelt, Maryland, which is managed by Goddard's Space Data and Computing Division.

The new system will replace the center's CRAY Y-MP supercomputer and be used for earth and space sciences research involving global climate change studies, global data assimilation, atmospheric/oceanic modeling, geodynamic earth surface modeling, space physics theory, astrophysics, and earth observing systems simulation.

"With 256 million words of central memory, atmospheric modelers will be able to perform multi-year data assimilation experiments with complex dynamical, chemical, interactive global, atmospheric circulation models requiring more than 100 million words of memory, something they could not practically do today," said Milton Halem, chief of the Space and Computing Division at NASA Goddard.

Martin Marietta Energy Systems, Inc., Oak Ridge, Tennessee, awarded Cray Research a subcontract to supply six CRAY APP (Attached Parallel Processor) systems with related computer equipment and system integration. Cray Research Superservers (CRS), Beaverton, Oregon, a wholly owned subsidiary of Cray Research, supplied the systems.

Cray Research announced that the U.S. Department of Energy's **Bettis Laboratory,** West Mifflin, Pennsylvania, placed an order for a CRAY Y-MP C90 system. Bettis is a government-owned laboratory operated by Westinghouse Electric Corporation for the U.S. Naval Nuclear Propulsion Program.

PSA Peugeot-Citroën accepted a CRAY Y-MP 4E supercomputer in PSA's Data Processing Center in Velizy, France. The system will be used for crash analysis, structural analysis, and computational fluid dynamics applications for automobile design and engineering. PSA's new Cray Research system displaced an earlier CRAY Y-MP system the organization installed in mid-1991. "The speed and memory size of the CRAY Y-MP 4E system will help us increase significantly our ability to do more complex numerical simulations," said Bernard Girard, the PSA director of product and manufacturing information technologies strategy. "This, in turn, will help us reduce our design cycle, improve quality, and save on costs related to prototype experiments."

NASA's Ames Research Center, Mountain View, California, ordered two Cray Research supercomputing systems. Under an agreement with the NASA Ames Division of Sterling Software Inc., a CRAY Y-MP 2E system was installed at the Central Computing Facility of the NASA Ames Research Center in Moffett Field, California. The agreement calls for a CRAY Y-MP C90 system to be delivered to the same facility this year.

The **Pittsburgh Supercomputing Center (PSC)** accepted a 16-processor CRAY Y-MP C90 supercomputer system, which will be used by academic and industrial researchers nationwide. "The power and broad applicability of the C90 system will enable the nation's researchers to master new challenges important for U.S. scientific leadership, and it provides a powerful resource for industrial research and product development," said Robert Mehrabian, president of Carnegie Mellon University. The Pittsburgh Supercomputing Center, a joint project of Carnegie Mellon University and the University of Pittsburgh, together with Westinghouse Electric Corporation, was established in 1986 by a grant from the National Science Foundation with support from the Commonwealth of Pennsylvania.

Forschungszentrum Jülich GmbH, also known as KFA, the largest of 16 national research establishments in Germany, ordered a CRAY Y-MP M94 system for various research applications. The new system will replace a CRAY X-MP system at the center.

The **University of Alaska Fairbanks (UAF) Arctic Region Supercomputing Center (ARSC)** accepted a CRAY Y-MP M98 system. The new supercomputer, named Denali, the Alaska native name for North America's highest mountain, is the first Cray Research system for UAF and the first of the company's systems in Alaska. It will be housed at Cray Research's Eagan facilities until the completion of ARSC's new supercomputer center. "Denali will provide a tremendous resource for the federal environmental research and development program through the center's focus on arctic system science research," said Tom Healy, director of ARSC.

The **Leibniz Computer Center of the Bavarian Academy of Sciences (Leibniz Rechenzentrum der Bayerischen Akademie der Wissenschaften–LRZ)** Munich, Germany, ordered an eight-processor CRAY Y-MP system. The system replaces a four-processor CRAY Y-MP supercomputer. The contract also calls for a central memory upgrade, scheduled for third quarter 1993, that will double the supercomputer's memory capacity. The upgraded system will be the central supercomputing resource for all universities and research institutes in the State of Bavaria.

The **Queen's University of Belfast** has installed a CRAY Y-MP EL entry level supercomputer system at the university's computer center in Belfast, Northern Ireland. This is the first of the company's entry level systems to be installed at a UK university. The system will be used as a local resource for large-scale computation. It also will be linked to the more powerful CRAY Y-MP 8I supercomputer in the UK Atlas Supercomputing Centre at the Rutherford Appleton Laboratory near Oxford, England.

New Cray Research product simplifies supercomputer applications development

Earlier this year, Cray Research announced a new product that will help programmers create easy-to-use applications software for Cray Research supercomputing systems. The product, Cray Visualization Toolkit 2.0 (CVT 2.0), consists of the most recent releases of six emerging standard windowing and graphics tools that software programmers use to create the user interface for their programs. The product also makes it easy for applications that use these standards to be ported to Cray Research systems.

The CVT 2.0 tools consist of libraries and toolkits used to create common elements of user interfaces (that is, windows and graphics elements such as menus and icons). It includes the following components:

- ☐ The X Window System Release 5 (X11 R5)
- ☐ Sun Microsystem's XView Toolkit 3.0
- ☐ Open Software Foundation's (OSF) Motif Toolkit 1.2
- ☐ Silicon Graphics' Distributed Graphics Library (DGL)
- ☐ The PEXlib library
- ☐ The Tk/Tcl toolkit command interpreter

Programmers can use these components as building blocks to reduce the development time for applications. For example, familiar components such as buttons, menus, and scrollbars can be created with a minimum of library calls; programmers customize the button for the application.

Cray Research has ported each of the standard toolkits in CVT 2.0 to the company's entire product line. With this Cray Research-supported product, customers can choose from these toolkits to produce portable applications software for Cray Research system users.

X11 R5 provides an essential foundation for CVT 2.0. This release supports diverse client applications and allows users to run X11 R5 applications on their Cray Research system and to connect to any workstation, PC, or terminal that runs the corresponding X Window System server. The other standard toolkits provide the following advantages:

- ☐ XView Toolkit 3.0 lets XView applications be easily ported to Cray Research systems.
- ☐ Motif 1.2 Toolkit provides a library of graphical objects such as menus, scroll bars, push buttons, and text windows.
- ☐ DGL delivers interactive three-dimensional graphics to the workstation and allows users to rotate, zoom, and transform their on-screen images.
- ☐ The Tk/Tcl toolkit and command interpreter provides high-level access to X Window System programming, which is much like writing shell scripts.
- ☐ The PEXlib library of low level graphics routines gives programmers access to the PEX protocol, allowing end users to display and interact with three-dimensional graphics from any workstation that has an X Window System server with the PEX 5.1 protocol extension.

Cray Research enters high performance RAID market

Cray Research announced in March its newest disk array technology based on the company's DA-60 and DA-62 disk arrays and its new DCA-3 channel adapter. This technology provides RAID (Redundant Array of Inexpensive Disks) level 3 reliability features and a fourfold increase in data transfer. This is the first RAID level-3-based storage product from Cray Research.

The company also announced a price reduction on the DD-60 and DD-62 disk drives, the components that make up the DA-60 and DA-62 arrays. The DD-60's pricing now is reduced by 20 percent; the DD-62's pricing is reduced by 28 percent. Cray Research formally announced these disk drives in 1991 and 1992, respectively.

The DA-60 array is capable of a sustained sequential I/O transfer rate of 80 million bytes (Mbytes) per second and a storage capacity scalable from 7.84 billion bytes (Gbytes) to 62.72 Gbytes.

The DA-62 array has a sustained sequential I/O transfer rate of 32 Mbytes per second and a capacity scalable from 10.92 Gbytes to 87.36 Gbytes.

The DCA-3 channel adapter is based on Cray Research's latest CRAY C90 integrated circuit technology. The new hardware product provides a high I/O performance direct disk channel connection between the I/O (input/output) subsystem of Cray Research supercomputers and the DA-60 and DA-62 disk arrays. This new technology provides improved disk storage reliability, enhanced data integrity, and high I/O performance. The disk arrays can be used with all CRAY Y-MP supercomputers with an IOS module and the more recent CRAY C90 systems.

The DCA-3 channel adapter supports four data drives and a dedicated drive for storing parity information, which is used in data reconstruction if a drive failure or uncorrectable media error occurs. Cray Research's new disk array hardware and software (included in the most recent release of Cray Research's UNICOS operating system) provides one of the industry's fastest data reconstruction environments. If a drive fails, for example, the array hardware continues to operate while the failed drive is replaced. The data reconstruction software rebuilds the data on the replacement drive while the array is still operating and provides several orders of magnitude faster reconstruction than the company's previous disk array technology.

Cray Research announces workstation version of UniChem product

Cray Research announced in March the availability of a workstation version of the company's UniChem molecular modeling program originally designed for its supercomputers. This was the first time Cray Research made its supercomputer software available on workstations.

The new software initially is available on Silicon Graphics workstations and will be available on most UNIX based workstations later this year. It has the same look and feel as UniChem for Cray Research supercomputers and contains all of the same features as the supercomputer version, which was developed by a consortium of Cray Research and leading chemical and pharmaceutical companies.

The common graphical user interface of the workstation and supercomputer versions of UniChem lets users access simulation codes that run on local workstations and on any Cray Research system. Currently the workstation version of UniChem is available through Cray Research.

Cray Research previews software for enhanced distributed processing

At the spring 1993 Cray User Group meeting in Montreux, Switzerland, Cray Research announced key features of its new Network Queuing Environment, a suite of software tools designed to make it even easier for Cray Research supercomputer users to send computer jobs from their workstations to high-performance computers on today's UNIX based networks. Product availability is slated for late 1993.

With Network Queuing Environment software and high-speed hardware connections, desktop workstation users can more easily access the computer resources needed to solve their problems. With this new software environment, users need not understand the actual computing resource they are using or concern themselves with computer programming when submitting jobs. The software manages file transfers, job queuing, and load balancing on the network, letting users focus on the scientific and engineering problems they must solve.

The Network Queuing Environment will include two components: Cray Research NQS and a set of Network Queuing Tools that make it easier to distribute and monitor workloads on networked UNIX systems. The tools are

- ☐ Network Queuing Client, an easy-to-use interface for job submission to NQS
- ☐ File Transfer Agent (FTA), which queues outbound and inbound file transfers over ftp
- ☐ Network Load Balancer, which monitors jobs in queue and sends them to the most appropriate resource on the network for processing

Cray Research NQS software is included in the UNICOS operating system software. It has long provided supercomputing features not found in other NQS versions, such as checkpoint/restart, tape and SSD storage support, password encryption, enhanced status display support, global limits, and a single administrator interface. In the UNICOS 7.0 release, Cray Research NQS introduced significant new features, including fair share scheduler support, guaranteed network output return using FTA, and enhanced scheduling support for jobs needing immediate execution.

Spectrum available on Cray Research systems

Cray Research users now have access to *Spectrum,* a new software package that enables engineers to model the multiple and interacting fluid, solid, and thermal physics inherent in real-world engineering designs. Based on proven finite element analysis concepts, it moves beyond the simplified component analysis capabilities of traditional tools to accurate, multiphysics simulation of complete product designs.

Most mechanical designs are determined by a wide range of physics—aerodynamic structural response, noise vibration and harshness characteristics of automobiles, various types of metal forming processes, and air cooling of semiconductor devices, to name a few. Accurate simulation of such designs is possible only with a tool that comprehensively captures the interacting effects of the multiple structural, fluid, and thermal physics involved.

Spectrum is built on CENTRIC's Full Spectrum Physics capability, the fundamental technology enabling the simulation of multiple, interacting physical forces and processes. To ensure optimal solutions in complex simulations, *Spectrum* includes multiple numerical solution techniques that the software automatically selects based on the design problem's characteristics and the computer hardware environment. Users interact with *Spectrum* through a unique graphical user interface, powerful command language, or popular complementary modeling and results interpretation tools.

For more information on using *Spectrum* on Cray Research systems, contact Bill Booth, CENTRIC Engineering Systems, Inc., 3801 East Bayshore Road, Palo Alto, CA, 94303; telephone: 415/960-3600; or Jeff Zais, Cray Research, Inc., 655E Lone Oak Drive, Eagan, MN 55121; telephone: 612/683-3656; email: jeffz@cray.com.

PerfStat real-time performance monitor available on Cray Research systems

PerfStat, an easy-to-use real-time performance monitor from Instrumental, Inc. is available on Cray Research supercomputers. PerfStat enables system administrators, performance analysts, user support staff, and end users to

- ☐ Improve system throughput
- ☐ Reduce individual or workgroup turnaround time
- ☐ Make more informed upgrade decisions
- ☐ Prepare more meaningful management briefings
- ☐ Improve productivity

PerfStat collects real-time, process-level UNICOS performance data and estimates some unreported data. PerfStat then interprets, validates, and graphs these data, creating usable information that can be accessed from a workstation running the X Window System. Because each site has unique data analysis requirements, PerfStat includes a commercial UNIX spreadsheet called WINGZ, and data may be exported to other analysis applications.

PerfStat will

- ☐ Collect, display, and graph data at the process, job, and system levels, providing a more accurate picture of resource usage
- ☐ Allow workload classifications based on resource usage or user information
- ☐ Calculate CPU wait time, which is not tracked by UNICOS
- ☐ Allow reconstruction of system behavior

For more information on using PerfStat on Cray Research systems, contact Robbie Cordo, Instrumental, Inc., 4500 Park Glen Road, Suite 390, Minneapolis, MN 55416; telephone: 612/920-6188; email: rtc@instrumental.com.; or Rebecca Koskela, Cray Research, Inc., 655F Lone Oak Drive, Eagan, MN 55121; telephone: 612/683-5756; email: koskela@cray.com.

FLUENT 4.2 includes more physical models, increased solver speed, and improvements to user interface

FLUENT 4.2 incorporates many enhancements to this flow modeling software, including the addition of physical models, increased solver speed, and improvements to the user interface and graphic display. Cray Research helped complete the extensive code optimization and vectorization work in this version. With the use of improved equation solvers, this work translates into greatly improved performance on both scalar and vector machines. FLUENT 4.2 routinely runs six to ten times faster than the previous release, FLUENT 4.1. On some benchmark problems, speedup factors of up to 25 have been achieved.

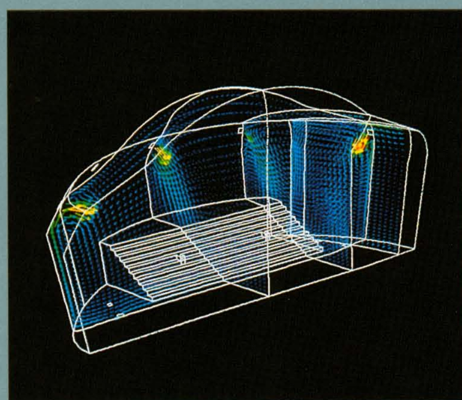

Velocity vectors show the ventilation pattern in an enclosed stadium. Computed using the CFD code FLUENT.

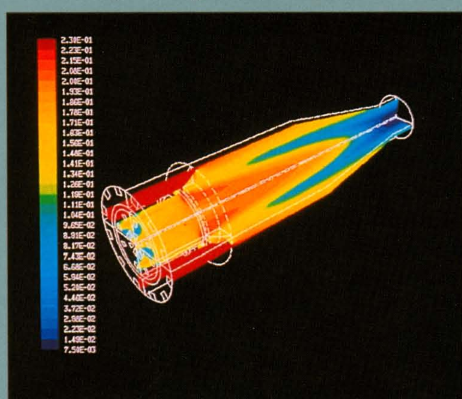

Oxygen concentration in a high velocity gas burner. Computed using the CFD code FLUENT.

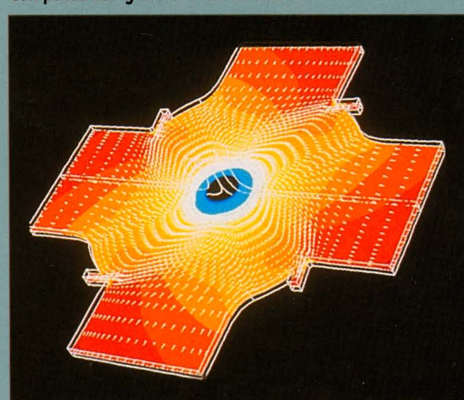

Turbulent flow in a vortex amplifier showing the static pressure and velocity field. The swirling flow is accurately predicted by the RNG turbulence model in FLUENT.

New features in the FLUENT 4.2 solver include a renormalization group (RNG) theory turbulence model, an improved differential Reynolds stress model (RSM), an option to use cylindrical velocity components, specified mass flow rate, periodic boundary conditions for fluid flow and heat transfer, and chemical vapor deposition models. In addition, particle trajectory algorithms have been improved, a fixed-variables option enabled, time-dependent calculations in rotating coordinates can be performed, and improvements have been made to droplet evaporation/heat transfer models. FLUENT 4.2 also includes a special co-located velocity-pressure formulation that is considerably more robust than conventional co-located procedures.

FLUENT 4.2 now offers a choice of three solver types: line Gauss-Seidel with block correction, an additive-correction multigrid method, and GMRES (generalized minimal residual method). Use of the latter two methods can yield substantial performance improvements on many difficult problems. Interface and graphics improvements include additional grid manipulation utilities, new alphanumeric postprocessing features, and more powerful boundary condition options. A greatly improved preprocessor with three-dimensional graphics and on-screen selection capabilities also is being shipped with FLUENT 4.2.

For more information on using FLUENT 4.2 on Cray Research systems, contact Dipankar Choudhury, Fluent Inc., Centerra Resource Park, 10 Cavendish Court, Lebanon, New Hampshire, 03766; telephone: 1-800-445-4454; or Dick LaRoche, Cray Research, Inc., 655E Lone Oak Drive, Eagan, MN 55121; telephone: 612/683-3696; email: laroche@cray.com.

Cray Research, Moldflow team up to advance supercomputer simulation in plastics industry

Cray Research and Moldflow, Pty, Ltd. announced a partnership to advance supercomputing capabilities for the plastics injection molding process widely used in the plastics industry.

Moldflow is the developer of the MOLDFLOW suite of plastics injection simulation software. This suite includes software modules that accurately model and optimize the plastics injection molding process and its effects on the final molded product. MOLDFLOW software analyzes polymer flow within the mold, the effects of mold cooling on flow, and the resultant shrinkage and warpage of the product. Moldflow will use its newly installed CRAY Y-MP EL system to boost the software's performance on Cray Research's full line of supercomputing systems, including a targeted tenfold performance gain in plastics shrinkage and warpage analysis.

The partnership also calls for converting the MOLDFLOW software for use on Cray Research's T3D system, the company's first-phase MPP system due out later this year. In addition, Moldflow plans to do next-generation software development on the CRAY Y-MP EL system to add capabilities for fiber orientation and process optimization, which are highly compute-intensive analyses.

The Cray Research version of MOLDFLOW contains unique features that take advantage of the supercomputer's capabilities. Plastics engineers can simulate a wide range of production scenarios in minutes, compared with hours on a workstation. Apple Computer, the Samsung Advanced Institute of Technology in South Korea, and The Ohio State University have licenses for the Cray Research version of MOLDFLOW. Apple used Moldflow on its CRAY Y-MP system to analyze nearly all molded components of the PowerBook and the PowerBook Duo laptop computers prior to committing huge sums of money for tooling. The simulation also saved tooling engineers time by eliminating the need for the typical seven or eight tooling tryouts.

For more information on using MOLDFLOW on Cray Research systems, contact Roland Thomas, Moldflow, Pty, Ltd., 259-261 Colchester Road, Kilsyth, Victoria 3137, Australia; telephone: 61-3-720-2088; or Mike Obermier, Cray Research, Inc., 655E Lone Oak Drive, Eagan, MN 55121; telephone: 612/683-3650; email: mikeo@cray.com.

CMF.Baseline release 4.4 provides enhanced performance monitoring and capacity management for Cray Research systems

TeamQuest Corporation recently announced a graphical performance analysis and capacity planning tool for Cray Research systems. Capacity Management Facility Baseline (CMF.Baseline) release 4.4 includes the capability to

A sample CMF.Baseline window showing rates of system calls, CPU utilization, and most-active disks.

- ☐ Monitor system performance in real time
- ☐ Report on past performance from a performance database
- ☐ Perform automated correlation analyses to find possible causes of performance bottlenecks
- ☐ Perform long-term trend analysis to plan for future system needs

CMF.Baseline supports two types of performance reporting: real-time monitoring and historic performance reporting. Real-time monitors report current system performance data and are updated on a timed basis to reflect the most current data available. This provides the earliest warning of problems, allowing users to investigate bottlenecks as they occur. Historic reporting enables users to look back at data, find and examine performance and system usage trends, and predict future system needs.

CMF.Baseline uses a client/server architecture. The client software runs on the desktop and handles all user interface tasks. The server software runs on the mainframe and handles all data collection, data maintenance, and correlation analysis. To generate a report, the client software formats a request and sends it to the server software for processing. Data packets are returned to the client and used to create line graphs, area graphs, bar charts, scatter plots, pie charts, or textual reports.

Multiple report windows from a variety of systems can be opened within the desktop, enabling users to examine facets of performance on a single system or compare the performance of more than one mainframe with side-by-side reports. In addition, predefined report definitions are installed with the product; these definitions can be modified at any time, and new reports can be added. Users can copy and paste report windows into other desktop applications. Reports also can be sent to hardcopy devices such as laser printers.

For Cray Research customers, easier access to system performance and capacity management information will enhance decision making by both technical staff and management and allow them to more readily tap into Cray Research expertise for advice and help. Cray Research's system performance and capacity planning specialists can use the TeamQuest products to provide performance management services faster and more efficiently by basing system analysis on easy-to-explain graphical representations of performance data.

For more information on using CMF.Baseline with Cray Research systems, contact TeamQuest Corporation at 1-800-551-TEAM or 515/357-2700; email: norb@TeamQuest.com; or Doug Wiedder, Cray Research, Inc., 655F Lone Oak Drive, Eagan, MN 55121; telephone: 612/683-5503; email: dlw@cray.com.

USER NEWS

Human birth process a welcome addition to computer simulation family

Thanks to Ernst Nalepa, Arthur Wischnik, and a CRAY Y-MP supercomputer, human births could become less risky for mother and child. Nalepa, consultant engineer for computer aided engineering at Electronic Data Systems (EDS) GmbH, and a development team at EDS worked with Wischnik to create a computer model to simulate childbirth. They have been running the simulation on the Cray Research system. By providing a pre-labor look at the birth process, the model ultimately could help physicians identify potential problems and take the necessary steps to address them.

"The human birth process is inherently much more complicated than in quadruped mammals, for example," said Wischnik, head physician at the University Women's Clinic in Mannheim, Germany. "On the one hand, when humans became erect-walking creatures they developed a more complicated pelvic structure. On the other hand, evolution outfitted humans with a larger brain mass requiring more space. The result is a bigger, rounder cranium, which is why normal vaginal birth is possible for *Homo sapiens* only through tricks of nature."

Two such tricks are the soft cranial structure of the fetus, which allows the head to be geometrically deformed during birth, and the ability of the birth canal to expand several thousand percent—something no known artificial material could withstand. Despite these accommodations, mothers and babies who undergo a difficult birth process can be subject to harmful immediate and later consequences. Mothers often suffer massive difficulties for months and may experience infirmities during or after menopause. For babies, a difficult birth may lead to neurological ailments such as hyperactivity, spastic episodes, difficulty in concentrating, and reduced intelligence.

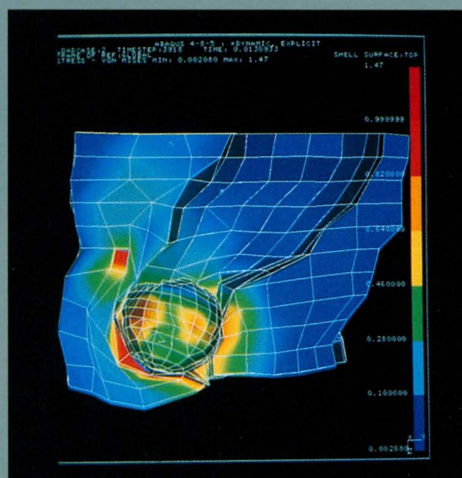

Simulation of baby's head emerging from the birth canal. Color scales indicate the degree of mechanical stress.

Watching a television interview in which Nalepa discussed computer simulations of car crashes, Wischnik was struck by the connection to childbirth: both events involve an interaction between two bodies in which the impacted materials are stretched to, or beyond, the limits of their physical integrity. Knowing that automotive engineers use computers to model car crashes, Wischnik tried to calculate the risks of the birth process with a personal computer, but lacked sufficient computing power, appropriate software programs, and a way of properly evaluating the measurement data. Nalepa provided access to the necessary technological resources; a CRAY Y-MP supercomputer in EDS's computing center in Rüsselsheim was used to perform the complex calculations of the birth simulation model.

Led by Nalepa, an EDS development team created a finite element model from representations of fetal and maternal bone and related soft tissue structures that are most frequently subject to injuries. The team used the ABAQUS EXPLICIT software package from Hibbitt, Karlsson, and Sorenson, Inc. to convert the properties of the materials, contact surfaces, boundary, and other conditions, as well as the stresses of the birth process, into calculable numerical values.

The model uses medical imaging data to reconstruct the anatomy of the mother and baby and portray the fetus's movement through the birth canal. The model takes into account all of the essential factors: physical posture, movement capacity of the pelvic joints, extensibility of the birth canal, and the ability of the child's head to deform safely. It also is possible to test the head movement under various conditions, such as different size assumptions.

All of this can be watched on a computer screen. Color scales indicate the degree of mechanical stress at each point in time and show how much stress the head and soft tissues can withstand. This information can be visualized precisely for each phase of the birth. The effects of an imbalance between the mother's birth canal and the unborn child's head size are assessed and quantified on the computer screen.

Once the team could visualize the process, they modified and refined the model, improving the numerical stability of the data analysis. Finally, the team was able to represent the results graphically as static images, which led to the development of a series of image sequences.

GALLERY

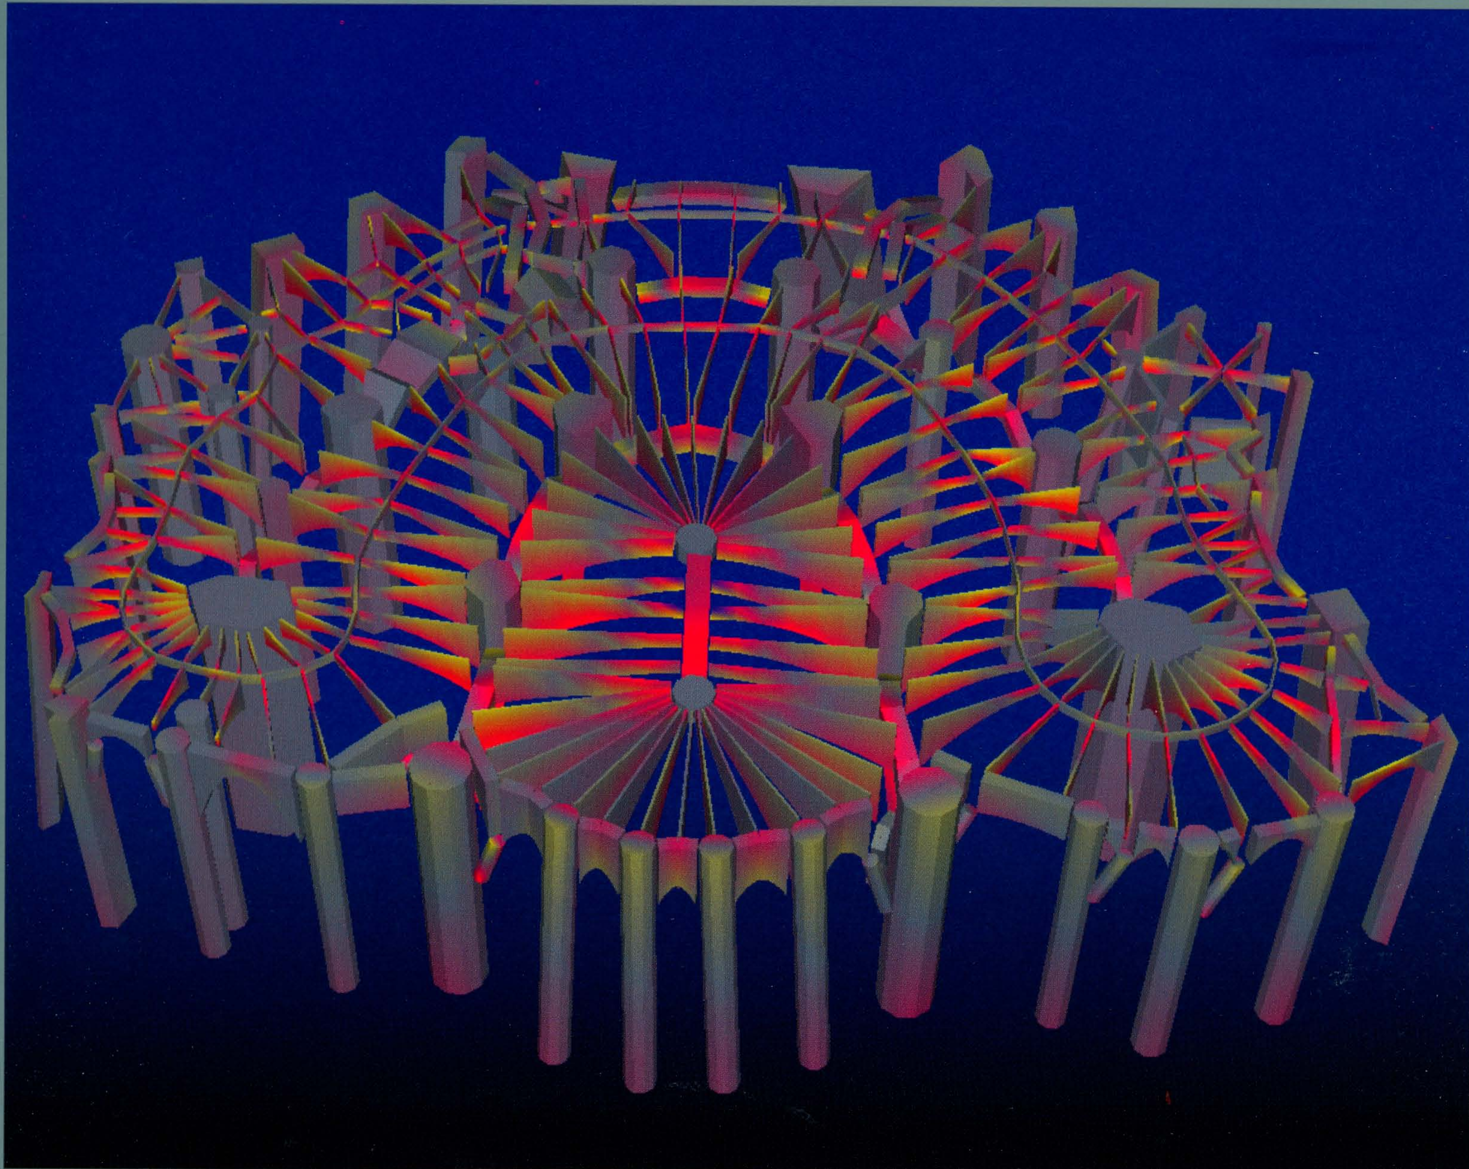

A structural model showing the stress distribution of the Crypt of Colonia Guell, located near Barcelona. Designed by Catalan architect Antoni Gaudi, the original plan included a church to be built on top of the crypt. Gaudi designed the church and crypt to achieve perfect stress equilibrium. However, funds ran out before the structure could be completed. As a result, the crypt has been damaged by some loads that it is unable to support.

Universitat Politècnica de Catalunya researchers Climent Molins and Pere Roca conducted a numerical study of the crypt using their own locally developed software. The software runs on the CRAY Y-MP/232 system at the Supercomputer Center of Catalonia. Molins and Roca are analyzing the stress on the structure and hope to devise ways to compensate for the loads. The image was generated with Cray Research's MPGS engineering postprocessing software.

Funds for the project were provided by the Architectural Heritage Department of the District of Barcelona (Servei de Patrimoni Arquitectonic de la Diputacio de Barcelona).