


CRAY CHANNELS

SPRING 1992 • A CRAY RESEARCH, INC. PUBLICATION



E M E R G I N G
A P P L I C A T I O N S

CRAYCHANNELS


In this issue

Supercomputers streamline research and product development in numerous industries by minimizing the guesswork involved in R&D. By giving scientists and engineers information faster and more economically than traditional R&D methods, supercomputer modeling helps companies in competitive industries save time and cut costs. From automotive design to petroleum recovery, from molecular modeling to medical imaging, and from aerodynamic optimization to complex structural analysis, supercomputing has revolutionized the way scientists and engineers tackle their technical challenges.

But supercomputing is expanding into new applications as well. In this issue of CRAY CHANNELS, we survey some emerging uses for large-scale computational modeling, including the modeling of complex phenomena in metal forming, plastic injection molding, and aircraft cabin ventilation. Our regular departments include reports on new software for processing-plant and electromagnetic modeling and new research into the AIDS virus.

As long as scientists and engineers have access to the highest levels of computing power, they will continue to forge new pathways in computational science. Their efforts ultimately benefit diverse sectors of society by improving safety, efficiency, and economy across many product areas. Their efforts also help businesses run more competitively by bringing the newest technologies to R&D.

Features



6

2

6

10

13

16

20

26

Departments

CRAY CHANNELS is a quarterly publication of the Cray Research, Inc. Marketing Communications Department, Tina M. Bonetti, Director. It is intended for users of Cray Research computer systems and others interested in the company and its products. Please mail feature story ideas, news items and Gallery submissions to CRAY CHANNELS at Cray Research, Inc., 1440 Northland Drive, Mendota Heights, Minnesota 55120. Subscription inquiries and address changes should be sent Attention Department D.

Volume 14, Number 1

Editorial staff

Ken Jopp, editor
Mary Miller, associate editor
Heidi Hagen, contributing editor
Anne Hazelroth, contributing editor
Chris Kruell, contributing writer

Design and production


Barbara Cahlander, typesetter, production artist
Eric Hanson, graphic designer
James Morgan, art director
Cynthia Rykken, graphic designer

30

32

33

37



Global change research: Penn State and the Earth Observing System

Eric J. Barron, Earth System Science Center, Penn State University, University Park, Pennsylvania

Equipped with a CRAY Y-MP 2E system, researchers at Penn State University are developing large-scale hydrological models for NASA's global climate research project, the Earth Observing System.

Advances in metal forming simulation at Mannesmann

Hans Joachim Pehle and Peter Thieven, Mannesmann Demag Hüttentechnik, Mönchengladbach, Germany

Robert Zeller, Cray Research GmbH, Munich, Germany

Process simulation, network supercomputing, and finite element modeling help engineers improve the design of metal forming processes.

Apple Computer PowerBook computer molding simulation

Brian Guinn, Apple Computer, Inc., and Mike Obermier, Cray Research, Inc.

Plastics engineers at Apple Computer use supercomputing to cut tooling costs and ensure high quality parts.

Network supercomputing with database management systems at EDF

Michel Demuynck and Philippe Bédu, Electricité de France, Clamart, France

EDF engineers achieve true network supercomputing with transparent file transfer between their IBM DB2 database and Cray Research systems.

Numerical analysis of airflow in aircraft cabins

Fred Aboosaidi, The Boeing Company, Seattle, Washington

Supercomputer modeling of cabin airflow distribution systems helps ensure airplane passenger comfort.

Quantum ballistic transport in a model semiconductor nanostructure

Cengiz Besev, Zhen-Li Ji, and Karl-Fredrik Berggren, Linköping University, Linköping, Sweden

Quantum-level simulations suggest new approaches to the design of ultraminiaturized electronic devices.

Parallel processing with ATExpert, a case study

Margaret L. Simmons, Los Alamos National Laboratory, Los Alamos, New Mexico

Kenneth M. Lord, Cray Research, Inc.

Cray Research's Autotasking expert system, ATExpert, detects obstacles in user codes that hinder multiprocessor execution. This tool improves productivity by helping users take advantage of multiprocessing.

Corporate register

Applications update

User News

Gallery

GLOBAL CHANGE RESEARCH

Penn State and the Earth Observing System

Eric J. Barron, Earth System Science Center, Penn State University, University Park, Pennsylvania



Recently the Earth System Science Center (ESSC) at Penn State University installed a CRAY Y-MP 2E system to advance ESSC research as part of NASA's Earth Observing System (EOS). The EOS project is part of an international effort to study the interaction of Earth's environmental systems and measure and predict the human effect on global climate change. Studying the Earth as a whole system presents one of the greatest technical challenges ever, both to the scientific community and to the computer industry, whose engineers are challenged to provide computing power several orders of magnitude greater than is provided today. This article focuses on this challenge and the people who have accepted it.

The size of the challenge

Some elements of the human "fingerprint" on the Earth system are as easily observed as the brown clouds of pollution that hang over many large cities. Other elements are much more difficult to detect, however, due to the complexity of the system's components and their interactions. Today's most sophisticated general circulation models (GCMs) of the atmosphere predict a global warming over the coming decades in response to an increased concentration of greenhouse gases such as carbon dioxide and methane. The model predictions for a globally averaged temperature increase typically are within the range of 1.5 to 5 °C. In comparison, the last natural climate change of this magnitude was the most recent ice age 18,000 years ago, during which the Earth cooled by

3 to 5 °C. While the ice age changes occurred over thousands of years, the human-induced changes could occur over a period ranging from decades to a century. Such rapid global change would have an enormous impact on water and water resources, agriculture, land-use, industry, and ecology.

Computer simulations offer the only feasible way to predict the effects of human activities and natural cycles on the planetary environment. With increasingly sophisticated Earth system models, scientists hope to provide reliable simulations and predictions, including scenarios that would illustrate the consequences of potential policy actions. The world's most capable supercomputers are being used in this effort, including Cray Research systems at the National Center for Atmospheric Research (NCAR) in Boulder, Colorado, and other international climate research centers and the newest system devoted to global change research, the CRAY Y-MP 2E supercomputer at the ESSC. However, current model-based predictions suffer from several shortcomings because we lack many observations necessary to verify and validate models, and we lack a sufficient understanding of many of the processes that link the oceans and atmosphere. Predictive models also have been limited by the availability of computer resources.

In response to the challenge of predicting global climate change, the international scientific community moved rapidly to develop global change programs through the World Climate Research Programme headquartered in Geneva and the International Geosphere-Biosphere Programme headquartered in Stockholm. In the United States, the Committee on Earth and Environmental Sciences was created to formulate a strategy for U.S. Global Change Research. NASA's Mission to Planet Earth was proposed as the space-borne contribution to the U.S. Global Change Program. In 1989, U.S. President George Bush called on NASA to work internationally, particularly with Europe and Japan, to launch this mission. The approval of the 1991 federal budget marked the beginning of this comprehensive program to study Earth from space.

Mission to Planet Earth

Mission to Planet Earth includes three near-term components: the EOS, the EOS Data and Information System (EOSDIS), and Earth Probes. In combination, these components are designed to address four critical shortcomings of current Earth studies:

- ☐ Current observations suffer from a lack of consistency or longevity (measurement devices have changed or observations are intermittent). Consequently, separating human-induced global change from natural variability is difficult. EOS is designed to provide a fifteen-year set of consistent and continuous observations of critical physical variables.
- ☐ Previous space-based platforms typically were planned to derive measurements of specific physical quantities. In global climate change research, however, detailed information is needed about the more general processes that link the

oceans, atmosphere and land surface, as well as the evolution of the system interactions. If the ultimate goal is to predict environmental changes, then the scientific tradition of focused disciplines must be replaced by a comprehensive interdisciplinary view. EOS will address this need by providing a broad, integrated set of observations using suites of sensors on single platforms in addition to satellites with specific objectives such as measuring tropical rainfall.

- ☐ Global climate change prediction has been limited by the coarse spatial and temporal resolution of global models. EOS seeks to provide broad high-resolution spectral and spatial coverage that will give the information needed for developing predictive capabilities on time and space scales appropriate for human activities. Extremely high-performance computing systems will be essential for the advancement of this effort.
- ☐ Much of the data about Earth is not easily accessible because of its high cost, lack of knowledge about its availability, or its inaccessible formats. The EOSDIS will provide low-cost, timely access to EOS data and to the complementary ground-based data. It also will support research and analysis of existing data as well as data collected before the existence of EOS. Within the first years after the launch of the mission, EOSDIS has the potential to exceed in volume all Earth observations ever collected.

The first ten years of the EOS program, including the launch of the first two platforms, is expected to cost close to \$11 billion, with the launch of the first platform expected in 1998.

Climate and hydrology: major elements of EOS

The hydrologic cycle is one of the most critical challenges in understanding the Earth system. EOS's prime focus on water is hardly surprising, given that water is an essential element of most processes that govern global climate change. The heat fluxes that determine climate largely depend on the transport, evaporation, condensation, and freezing of water. This crucial study is the aspect of the EOS project on which the ESSC will focus.

Despite the significance of water, our knowledge of the hydrologic cycle as part of the global interacting system is meager. We lack fundamental information about the size of the cycle's reservoirs and the rates of flux among them. Consequently, we are far from understanding how the hydrologic cycle will respond to a changing Earth system and the nature of water's influence across the components of the system.

The lack of comprehensive understanding of this fundamental cycle is readily explained: it is the outcome of dividing knowledge into separate disciplines. Adding to the researchers' difficulties has been the lack of consistent long-term measurements for critical quantities, such as those of precipitation, evaporation, soil moisture, atmospheric water vapor, snow cover, sea ice, glaciers, vegetation, cloud cover, temperature and winds—all essential for characterizing the climate and hydrologic




Figure 1. Topography of the United States at the resolution of a typical circulation model.

cycle. The fact that hydrologic processes cross an enormous range of time and space scales, from global phenomena to individual thunderstorms, poses another problem. EOS will provide the first ever long-term consistent global measurements of many of the key variables in the hydrologic cycle.

Penn State's role in EOS

In June 1988, 17 ESSC faculty associates proposed an EOS interdisciplinary investigation, "The Global Water Cycle: Extension Across the Earth Sciences." After a successful review, a formal work plan covering the period 1991 through 2000 was accepted, and the 17 faculty members, with four additional scientists from NASA's Marshall Space Flight Center, were confirmed as EOS investigators.

The research strategy for the investigation is based on the use of computational models as devices to summarize knowledge, convert patterns observed from space into physical variables and rates, and provide quantitative results that can be compared with results from field studies for model verification and improvement. Models relating physical variables to remote observations will produce the information needed to transform EOS observations into accurate models that reveal the role of water and its interactions with a changing Earth system. The central objective of the Penn State investigation is the development of a hierarchy of models from global to regional scales that will become the prime tools for an integrated predictive understanding of the hydrologic cycle. The prediction of global climate change is the ultimate goal of this research.

Modeling the hydrologic cycle

The problem of predicting climatic change on scales that range from the global to the regional is easily illustrated. We use as an example the first modeling and observation experiment planned for the EOS investigation—the prediction of changes in the hydrologic cycle of the Susquehanna River




Figure 2. Susquehanna River Basin dimensions relative to the topography in a typical general circulation model.




Figure 3. View of a four-mile segment of the Mahantango watershed of the Susquehanna River Basin, and the hydrologic flow paths associated with this 30 m topographic resolution.




Figure 4. Susquehanna River Basin at regional resolution (16 km) showing topography and river flow. (The white dot indicates the city of Baltimore.)




Figure 5. The Mahantango region of the Susquehanna River Basin and hydrologic flow model at a resolution appropriate for coarse-resolution basin-hydrologic models (80 m).

Basin. The weather patterns that influence the water balance of this region are clearly governed by global phenomena. However, current general circulation models are of relatively coarse resolution, often with the globe represented horizontally by only about 2000 points (Figure 1) and vertically by 9 to 12 levels. This resolution is inadequate to represent even large river basins. Figure 2 illustrates this problem, showing the Susquehanna Basin dimensions relative to the topography in a typical general circulation model.

If we consider the opposite end of the hydrologic spectrum and ask, at what spatial scale would a hydrologist want to simulate the physical processes of surface and ground water flow associated with particular storm events, we encounter very different requirements. Figure 3 shows a view of a subbasin of the Susquehanna, a four-square-mile segment of the Mahantango watershed. At this scale, preferred by the hydrologist, the grid squares are 30 meters across. To map the entire Susquehanna at this resolution would require 108 million grid points. Given that the computational time in a three-dimensional GCM goes up by the cube of the horizontal resolution, global models reaching the spatial and temporal resolution required to simulate the hydrologic cycle on a scale appropriate for human activities is not possible in the foreseeable future. Yet enhanced predictive capability is essential for researchers to predict global climate change.

Our response to this problem has been the development of nested predictive models. Initially, we are working to link the GCM with the Penn State/NCAR mesoscale climate model, which has a horizontal resolution ranging from 2 to 200 kilometers. The mesoscale model will be linked with a "lumped" or coarse-resolution basin hydrologic model (Figure 4) and with higher-resolution hydrologic models (Figure 5). Higher-resolution models of the boundary layer and of fluxes such as evapotranspiration will be used to explore the linking of the atmosphere and the land surface.

The Penn State program is computer intensive, even when using a nested modeling approach to overcome the computational resource problem associated with high-resolution climate studies. Recognizing this, NASA EOS became the major sponsor for the CRAY Y-MP 2E supercomputer at the ESSC. In October 1991, ESSC became the first individual academic research unit with its own Cray Research system out of the 44 systems located at academic institutions worldwide. Already the system is almost fully utilized in simulating past, current, and future climate and hydrology.

The project also will generate global and regional-scale data sets for the documentation of global change, particularly for hydrologic variables and for model validation and verification. The digital elevation models illustrated in Figures 3 to 5 represent layers in a geographic information system designed to provide the basis for either model improvements or model validation. Together, the model experiments and data collection will serve two purposes: to explore the potential of nested models in making regional scale prediction of the impact of global change and to assess the vulnerability of societies to changes in the water balance. ■

New system promotes interdisciplinary study

For those accustomed to the large laboratory-like spaces that usually house central computing systems, the unique installation of the ESSC's CRAY Y-MP 2E system may come as something of a surprise. The system totally fills a small office with the cooling system placed on the roof overhead.

This installation is temporary, however, and a new facility is being planned. Meanwhile, the need for supercomputing power to respond to questions of global change is so urgent that expediency was the only answer.

"We were ready for the Cray Research system," said William Peterson, a geoscientist and ESSC computing facilities manager. "We've been running close to full utilization from the start. Even at Christmas, when the entire university was on vacation, the Cray Research supercomputer operated at capacity for the whole term break."

"Now we can think about coupling the mechanics of mountain growth with processes of denudation," said geologist Rudy L. Slingerland. "The Cray Research system has opened up a whole new class of problems for me. Previously it simply wasn't possible to study an evolving mountain belt—you can't build experimental equipment, and observations are impossible at the scale we need. The set of equations was knowable but unsolvable. Now it just may be solvable for simplified sets of equations that focus on the most important processes. Having a supercomputer next door to do this research is very exciting."

For long-time Cray Research system users such as meteorologist Tom Warner, the new computer brings relief rather than excitement. "The availability of computational resources is a critical limitation in mesoscale modeling research," he said. "This dedicated machine will help keep us at the forefront of atmospheric modeling." But Warner also stresses the benefits of a local CRAY Y-MP system for interdisciplinary work. "It allows swifter transfer of large data sets over the fiber optic networks between the computer and the scientist who must interact closely with the data output, and it allows scientists from a variety of fields to work together on coupled computer models of complex physical systems."

ESSC's Cray Research system has had a dramatic impact for some researchers in the College of Earth and Mineral Sciences. In the Department of Geography, Peter Gould and his students are modeling the geographic spread of the AIDS epidemic. Exploring the spatio-temporal structure of map sequences with spatial adaptive filtering and parametric tracking used to be confined to the state level, with 60 to 100 counties a feasible maximum. With the help of the Cray Research system, doctoral student Joseph Kabel is now predicting what the national distributions will look like over the next five years.

The research opportunities for students are just beginning. "The Cray Research supercomputer is a lifeline for my students," says polymer scientist Sanat Kumar. "They homed in on it immediately. I estimate they will get their Ph.D. degrees more than a year earlier now."


The Cray Research system will allow researchers at the ESSC to produce many more simulations of past climates that can be used to evaluate climate sensitivity. The system not only makes work on the EOS project easier, but also brings diverse researchers closer together. The insight gained from the cross fertilization of ideas across disciplines is a major benefit of supercomputing in an academic environment. ■

About the author

Eric J. Barron is a professor of geosciences and the director of the Earth System Science Center at Penn State University. He received a Ph.D. degree in oceanography from the University of Miami. Prior to joining the university in 1986, he was a scientist in the climate section of the National Center for Atmospheric Research.

ADVANCES

in metal forming simulation at Mannesmann



Comparison of tube deformation in experiment (top) and simulation (bottom).

Challenged by international market demands and increasingly complex production requirements, a growing number of heavy industries worldwide are exploring new technologies to optimize manufacturing processes. Research and development efforts at the Mannesmann Demag Metallurgical Plant and Equipment Division of the German industrial conglomerate Mannesmann center around the development and optimization of pipe and tube milling equipment with particular focus on the control of tube forming processes. Recently Mannesmann teamed up with Cray Research to develop advanced computer models of various metal forming processes. The collaboration demonstrated many clear advantages of supercomputer modeling over traditional process design methods.

Metal forming processes

The complex metal forming processes performed in a tube forming plant, such as mandrel rolling, stretch reducing, cold pilgering, and cross rolling, require sophisticated tools and powerful equipment to ensure optimal use of the parametric variations involved in the production of these diversified industrial tube products.

Problems arising during these tube forming processes include

- ☐ Material defects, tears, and cracks
- ☐ Deviations of finished tube size from the required diameter and wall thickness
- ☐ Stresses and deformations of tubes along stands of a rolling mill
- ☐ Interstand stress distribution on tubes
- ☐ Extreme material spread into roll gaps during tube elongation
- ☐ Inner ovalness or "unroundness"

Tolerance deviations in wall thickness occurring during the deformation process, for

*Hans Joachim Pehle and Peter Thieven,
Mannesmann Demag Hüttentechnik,
Mönchengladbach, Germany
Robert Zeller,
Cray Research GmbH, Munich, Germany*

example, or material defects can lead to tube products that cannot be sold. Therefore, an understanding of the interdependence between process parameters and the resulting tube deformation is essential to the cost-effective production of tube products.

Most processes on rolling mills occur hidden from view, and because many of them are at high speeds and temperatures of around 1000 °C, it is difficult to observe or measure them by traditional experimental approaches. The experimental approach for improving forming processes, and consequently tube products, basically relies on numerous guesses or variations based on observed material flaws and defects. This often means lengthy trial-and-error experimentation. In addition, a plant manufacturer has to back-rent milling equipment at customer sites for the duration of the testing if no appropriate laboratory equipment is available. Earlier models for the computation of metal forming processes rely on simplifications of the basic material equations and tool geometry.

Project "Supercomputing in metal forming"

With the goal of exploring the application of computer simulation in process technology in general, and visualization as a tool to optimize tube forming processes in particular, Mannesmann Demag and Cray Research joined hands in a cooperative project called "Supercomputing in metal forming" in the fall of 1989. Two project phases were completed in the two years from 1989 to 1991.

During the first phase, the mandrel rolling process was simulated for one stand, the stretch reducing process for a two-stand configuration (with reduced distance between stands to lower computation time), and 30 percent of one stroke of the cold pilgering process was simulated. Stresses and deformations exerted on tubes during these forming processes were visualized computationally during this phase for the first time. Commercially available finite element (FEM) programs were used to simulate the forming process. These first phase results amounted to a technology breakthrough comparable to introducing computational crash analysis in automotive design.¹

Simulations during the second phase included several tube forming processes as well as

a number of parametric variations and tool configurations. Two of the processes simulated during the first phase, mandrel rolling and stretch reducing, were expanded considerably. This time, the mandrel rolling process simulation included two roll stands. For the stretch reducing process, the actual distance between individual stands was simulated with process variations. In addition, a third, and one of the most complex metal forming processes, cross rolling, was simulated for the first time.

Feasibility of process simulation established by network supercomputing

Both phases of the "Supercomputing in metal forming" project are excellent examples of network supercomputing. The computations for the second phase of the project, for example, ran on a CRAY Y-MP464 supercomputer at Cray Research's manufacturing facilities in Chippewa Falls, Wisconsin, via trans-Atlantic network connections with graphical output to two Silicon Graphics IRIS workstations in Munich, Germany. For the computations of both phases, the finite element programs MARC, from MARC Analysis Research Corporation, and ABAQUS, from Hibbitt, Karlsson & Sorensen, Inc., were used.


The complexity of these metal forming processes is evidenced by the 400 CPU hours required to run the computations during the second phase of the project and the 200 CPU hours used during the initial phase. Project participants on both sides of the Atlantic reported smooth operation of both hardware and software throughout all simulation runs.

The initial computations established the feasibility and practical applicability of process simulation in metal forming, moving Mannesmann Demag toward the company's long-term goal of ultimately using a standard process that would include process simulation as a development tool. The computations performed during the second year of the project not only confirmed beyond any doubt the feasibility of visualizing complex metal forming processes, but also moved Mannesmann Demag even closer to using process simulation during process and plant design, with the goal of shortening development and design cycles as well as achieving substantial cost and time savings. With practical applicability no longer in doubt, product and process optimization will be the main focus of future process simulations at Mannesmann Demag.


Interdependence of set parameters and tube deformation

The computations performed for the mandrel rolling, stretch reducing, and cross rolling processes demonstrate the interdependence between set parameters and tube deformation. Figures 1, 2, and 3 show how deformation develops depending on the predefined roll design and tool setting.


During the mandrel rolling process (Figure 1), for example, the tube is rolled on a mandrel—an inner tool—while its wall thickness is



1a




1b




Radial Strain^{1c}


Figure 1. Wall thickness reduction and tube elongation are visualized in two adjacent stands of a mandrel rolling assembly using MPGS. By partial removal of the tools, the simulation offers a direct view of the usually unobservable rolling process.



2a



2b



2c

being reduced and elongation occurs. To expose a full view of this reduction and elongation process in the simulation, these tools, or rolls, are partially removed, using a special feature of Cray Research's Multipurpose Graphic System (MPGS) software. The red areas of the tube passing through the two-stand mandrel rolling assembly (Figure 1c) clearly show the axial strain exerted on the tube during elongation.

The stretch reducing simulation visualized in Figure 2 involves two adjacent stands of a stretch reducing assembly. The simulated distance between the two stands is identical to the actual distance in the rolling assembly. The variable rolling speeds in adjacent stands exert axial tension on the tube. The exposed view in Figure 2b shows radial displacement, a measure of the diameter change, and Figure 2c reveals the tangential stress on the circumference of the tube when pressure and friction are applied by the rolls.

The deformation of a tube in a cross rolling assembly is shown in Figure 3a. As in the mandrel rolling and stretch reducing processes, the simulation presents an exposed view of the cross rolling stand assembly to allow direct observation of the forming process. Figure 3b shows the axial displacement of the tube under the force of the cross rolls. Zones of greatest deformation form a spiral around the tube, showing the rotation of the tube during its passage through the rolls. The tangential stress applied to the tube by the cross rolls is visualized in Figure 3c.

An understanding of the economics involved in tube rolling equipment is gained by looking at a single stand of a stretch reducing mill. One stand contains three rolls per casing, along with internal gears and other parts representing a monetary value of between DM 30,000 and DM 60,000. A total of 28 of these stand configurations constitutes a single reduction series with a value of DM 1 million or more. Obviously, substantial savings can be achieved if one complete equipment set can be eliminated by optimizing the roll shape and a number of parameters.

Results and future goals

As a result of the joint project "Supercomputing in metal forming," the feasibility of applying process simulation to tube technology has been confirmed. The processes chosen to be included in these simulations cover the main application areas for Mannesmann Demag's broad spectrum of tube mill equipment. By simulating the most difficult processes, it was proven that all tubing processes can be numerically simulated with the finite element method (FEM).

Mannesmann Demag's goal for all processes is to expand production ranges of its machines, covering thick and thin walls with minimal production expenditure. By optimizing the forming processes, product quality is increased, and research

Figure 2. In this configuration of two adjacent stands of a stretch reducing assembly, a tube is rolled at variable speeds by means of applied pressure and friction forces. Radial displacement (Figure 2b), a measure of diameter change, and tangential stress (Figure 2c), a measure of the forces applied to the tube by the rolls, are clearly visualized.

Figure 3. During cross rolling, one of the most complex of the tube forming processes, the axial displacement (Figure 3b) of the tube is shown by spirals formed around it, showing the tube's rotation during the rolling process. The tangential stress (Figure 3c) applied by the rolls corresponds to the forces forming the tube.

and development expenditures are reduced considerably.

The ability to test rolling processes with the finite element method before experimental equipment is available will reduce the experimental testing of numerous parametric variations. As soon as the equipment for actual testing of the results becomes available, targeted testing, based on the computed results, can be done, eliminating much of the traditional trial-and-error experimentation. Performing computations in the pre-development and pre-experimental stages before the physical capabilities even exist to perform experiments in the traditional sense amounts to jumping way ahead of experimentation. Supercomputer simulation allows Mannesmann to perform tests on rolling mills that currently do not exist.

Having opened a window onto these hard-to-observe processes during the "Supercomputing in metal forming" project, developmental engineers at Mannesmann now are working toward integrating computer simulation into development and project work as a standard procedure. The challenge now is to convert these computational results to practical application in actual production runs to achieve the ultimate goal of applied cost-effectiveness. ■

Acknowledgments

The authors thank Hibbitt, Karlsson & Sorensen, Inc., and their German subsidiary Abacom, for their support in using ABAQUS, and MARC Analysis for its support in using MARC K4. The animations and visualizations for the second phase of the project were performed by Rainer Johanni of Cray Research GmbH.

About the authors




Hans Joachim Pehle holds a Ph.D. degree in metallurgical engineering from the Technical University in Aachen, Germany. Since 1982, he has been head of the Technological Computations Department at Mannesmann Demag Hütten-technik in Mönchengladbach, Germany.


Peter Thieven holds a mechanical engineering degree from the University of Aachen, Germany, with emphasis on metal forming processes. He has been working at the Technological Computations Department of Mannesmann Demag Hütten-technik in Mönchengladbach, Germany, since 1988.

Robert Zeller has been a presales analyst with Cray Research GmbH, Germany, since 1987, focusing on simulation as an industrial technology. He holds an advanced degree in theoretical physics and a Ph.D. degree in theoretical solid state physics from the Technical University of Braunschweig, Germany.

Reference

1. Passmann, R., H. J. Pehle, P. Thieven, and R. Zeller, "Aspects of the Industrial Application of Process Simulation in Metal Forming," *MPT - Metallurgical Plant and Technology International*, No. 3, pp. 58 - 67, 1991.





Apple Computer PowerBook computer molding simulation

Brian Guinn, Apple Computer, Inc. and Mike Obermier, Cray Research, Inc.

Cooling line layout, material feed system, and finite element model for the PowerBook top cover mold.

"Do it right the first time" is a way of life for quality improvement devotees. An even better strategy is to "do it right before you do it," as demonstrated by Apple Computer engineers when they designed the manufacturing process for the PowerBook laptop computer plastic casing. Apple recently introduced the PowerBook, a portable version of the Macintosh personal computer, to the marketplace. By using the company's CRAY Y-MP 2E supercomputer to simulate the injection molding of the four major components of the PowerBook plastic case, Apple plastics engineers saved time and minimized production costs.

Injection molding simulation is not new to Apple. One of the first injection simulations Apple engineers performed was for the Macintosh SE computer bezel; the engineers used the MOLD-FLOW and MOLDTEMP application packages on a CRAY X-MP/48 system (CRAY CHANNELS, Fall 1989). The accuracy of the simulations and continuing successful results for other plastic


components throughout the Apple product line, including printers, keyboards, and monitor and CPU housings, easily justified an upgrade to a CRAY Y-MP 2E system in 1991. According to Rod Archer, manager of mechanical design technology, "Simulation using the CRAY Y-MP 2E system is like an insurance policy—it saves us time, money, and improves our product design. Plastic parts will never keep us from selling a machine."

Plastic injection simulations must be done thoroughly and quickly if they are to be of any value to the design process. Before Apple plastics engineers started using a Cray Research system to simulate the molding of plastic parts, the plastics analyses typically were incomplete and did not reach the tooling engineers in time to meet production deadlines. The tooling engineers therefore did not use the analyses from the plastics engineers, even if an analysis suggested a different tooling design. Now, the tooling engineers can work with color-coded profiles of parts, complete

with simulated material feed systems and cooling line placements because the power of the CRAY Y-MP 2E system allows the plastics engineers to complete their work within a project's design phase.

The PowerBook simulations were particularly noteworthy, however, because the engineers simulated the flow, cooling, warpage, and shrinkage of all four major components of the PowerBook's plastic case before finishing the detailed design drawings. Other plastic parts manufacturers might run a flow analysis; occasionally, they will add a cooling analysis, but never have they gone all the way through to the warpage simulations. And no other company has approached Apple's ability to run all four simulations before the final product design drawings are finished. The Apple engineers influenced the design of the PowerBook plastic components by running several iterations of all four simulations using computer models averaging about 7500 elements. Companies limited by computing power, however, typically are limited to much smaller models, usually less than a third the size of Apple's models. The fewer elements in a model, the less effective the simulation. Because it is limited to simple geometries, such a model may leave out design details that have an impact on the manufacture of the part. The PowerBook models were relatively simple compared to models of more than 25,000 elements that Apple engineers have used for other products.


The design of a plastic part has to take into account many variables: the material used, the design of the mold, the design features of the part, the processing conditions for the mold, how the mold (and therefore the part) is cooled, and the overall production goals for the tooling, so that




the proper number of molds can be built. Too many molds, and money is wasted. Too few, and production goals cannot be met.

The molding simulation involves designing a finite element mesh, modeling the plastic material feed systems, modeling the cooling water line layout, and calculating the cooling effects. A well-designed process might turn out parts that initially are well outside of print specifications, but they soon cool to their intended shapes—if the part, mold, and injection process were properly designed.

PowerBook top-cover mold filling simulations. Clockwise from upper left: fill time, temperature profile, warpage prediction, shaded image of completed part.



Colored temperature profile of plastic in display bezel mold.



Colored temperature profile of plastic in display housing mold.

The simulation process

Apple engineers used six codes from the MOLDFLOW suite of applications to simulate the injection process: MFL4 for filling; FCLP for cooling; MFLP for filling and packing; SHRINK for shrinkage; WARPIN as the preprocessor for FINAS; and FINAS for the structural solver.

The engineers used MFL4, an injection mold flow analysis code, as a starting point to develop data that would be further refined by the rest of the codes. They typically ran five or six iterations of MFL4 per part. With each iteration, the engineers fine tuned assorted parameters such as gate location, cavity fill rate, and melt temperature. The MFL4 results provided input data to the FCLP code, a cooling analysis code, which the engineers typically ran three or four times per part. FCLP starts with material heat load information from MFL4; the simulation calculates the inlet water temperatures and cooling water flow rate required to maintain the optimal mold surface temperatures suggested by the plastic material manufacturer. The results of the cooling analysis gave Apple engineers a better idea of the actual heat load profiles for the entire mold for each part. As a result they could rerun the fill analysis using one or two iterations of the MFLP code, which combines finite element and finite difference techniques to provide an even better simulation of filling and packing the mold cavity. In the past, tooling engineers used educated guesses to place cooling lines in the mold. This approach typically led to seven or eight tooling tryouts for each mold used, adding considerable manufacturing time and costs to new products.

The Apple engineers fed the results of the MFLP simulation into the SHRINK code to isolate the causes of warpage in the final, cooled (but still simulated) part. The biggest causes of incorrect warpage tend to be improper material orientation

based on the fill locations, an ineffective cooling system for the mold, or improper packing pressure (the pressure maintained on the filled cavity).

The next simulation used WARPIN as a preprocessor to prepare the code for FINAS. Of the six codes used in the overall design analysis, FINAS is the most CPU intensive, requiring four to five CRAY Y-MP system CPU hours for each 7500-element simulation. FINAS predicts the final warped shape of the plastic part in its most relaxed state, which needs to be at room temperature.

As a result of running the simulations on the Cray Research system, the plastics engineers saved the tooling engineers the typical seven or eight tooling tryouts. The colored heat load profiles and the confirmation of cooling line placement give tooling engineers concrete data to work with when designing their molds and injection process; the simulations helped avoid fundamental flaws in the mold. In addition, a properly designed part does not need to be put into a cooling fixture to force it into its intended shape. Using a cooling fixture adds time to the process, and it is only a short-term solution because plastic tends to return to its as-molded shape. Tooling tryouts cannot be completely eliminated because the tooling engineers need one or two tryouts to make sure the fit and finish of the mold components are correct.

Once production began on the parts for the PowerBook case, the parts appeared to be out of specification as they came out of the mold. However, as the parts cooled, they flattened into the shapes and tolerances required by the parts drawings.

By the time production is about to begin, even minor changes to a design are expensive. A design change at this time affects many more people and takes more time than if the product had been properly designed in the beginning. A product design is much easier to change when it is still on the drawing board.

The future will see computer-aided design and engineering (CAD and CAE) connected to the simulations, which will save even more time. As more and more plastics manufacturers begin to perform these kinds of simulations, they, too, will discover what Apple has already learned — injection molding simulation saves time, money, and is an invaluable aid in turning out superior products. ■

About the authors

Brian Guinn has worked as a Moldflow analyst at Apple Computer, Inc. for the past two years. Previously, he worked for Moldflow, Pty., Ltd. Guinn received a B.S. degree in industrial technology in 1988 and has a two-year degree in plastics engineering technology, both from Ball State University, Indiana.

Mike Obermier is a polymer processing specialist at Cray Research, Inc. He has spent the past nine years working in the injection molding industry, focusing on injection molding simulations on supercomputers. Obermier received a B.S. degree in plastics engineering in 1983 from Ferris State College, Michigan.



Network supercomputing with database management systems at EDF

*Michel Demuynck and Philippe Bédu
Electricité de France, Clamart, France*

Until recently, engineers at Electricité de France's (EDF) Research and Development Division (DER) had to be file transfer experts as well as engineers. To take advantage of both the data stored in their IBM DB2 relational database management system (RDBMS) and the computational power of their Cray Research supercomputer, they had to collect the data from the database in flat files, transfer the files to a Cray Research system, and write filters that would enable the particular applications they were using to read the data. This process was time consuming and error prone. Administrators at EDF's Clamart Computer Center wanted a faster and easier alternative. Specifically, they wanted to make information stored in their IBM DB2 relational database accessible to users of their two CRAY Y-MP supercomputers, without the users having to run intermediate jobs. Additionally, to implement a client/server model with DB2, EDF's administrators used an ORACLE interface, one of the most familiar RDBMS to EDF engineers.

EDF is France's sole supplier of pressurized water reactors used in the nuclear power industry, which supplies about 70 percent of France's total electrical output. The company's computer resources are used for a variety of applications, including post-accident simulation, and trying to predict the amount of electricity that must be produced to meet a particular day's requirements. This environment is computationally demanding, and huge amounts of data must be shared by several departments.

DER's network is extremely heterogeneous. Although this configuration satisfies a variety of user needs, achieving data transparency at every network node is a sizable challenge. At the heart of the configuration are four IBM 3090 computers supporting the company's strategic data. They run software products including IBM's proprietary database, DB2; one of them also runs the ORACLE relational database from Oracle Corporation. Heavy scientific computing is performed using the center's two CRAY Y-MP supercomputers. Engineers can perform most of their scientific calculations locally or remotely.

The computers at the Clamart Computer Center communicate via Local Area Networks (LANs). These LANs use the UltraNet and HYPER-channel communication devices, from UltraNetwork Technologies and Network Systems Corporation, respectively, and supply additional connectivity with the corporate wide area network, RETINA, and the site LAN, allowing access to the engineers' workstations. The Cray Research systems are connected to DER's IBM systems via FDDI and Ethernet, with TCP/IP as the communication transport protocol.

Relational database management systems in a Cray Research network supercomputing environment

Denton Olson, Cray Research, Inc.

Figure 1 (below). Oracle Corporation's SQL*Net and SQL*Connect gateway products enable transparent data access to IBM's DB2 database from ORACLE client applications running on the CRAY Y-MP computer system.

Table 1 (bottom). GEODE runtime results in CPU seconds comparing IBM DB2 centralized architecture to the CRAY Y-MP client/server architecture.

To meet the goals of data transparency and accessibility, the center implemented an RDBMS client/server architecture (see sidebar). This was not an obvious solution, given that IBM's DB2 system is proprietary and does not normally support a client/server architecture. Additional challenges

The network supercomputing goal at Cray Research is to make access to Cray Research supercomputers easy, reliable, and cost-effective for users anywhere on a network. Through the implementation of networking standards, Cray Research provides connectivity to virtually all UNIX-based mainframes (and, as in the case of EDF, non-UNIX mainframes), minicomputers, and workstations. The many communication products and protocols offered by Cray Research allow applications to be distributed within the network, enabling the most efficient use of resources.

The need to share data in a heterogeneous computing environment is leading more and more scientific and engineering companies to look for ways to standardize their data management across all hardware platforms, including supercomputers. A solution for many customers is a relational database management system (RDBMS).



were posed with the requirements that the DB2 and ORACLE databases, acting as servers, could reside essentially anywhere on the network, with their locations and data formats transparent to the user. Finally, the client running on the Cray Research supercomputer had to be able to use data generated by the databases without becoming involved in complicated file transfer processes or incurring performance degradation due to data management.

EDF chose Oracle's SQL*Connect gateway product to permit transparent data access to the DB2 database from ORACLE front-end clients running on the CRAY Y-MP computer systems. Functioning as an RDBMS "translator," SQL*Connect on the IBM makes DB2 look like ORACLE to the front-end ORACLE clients on the Cray Research supercomputers (Figure 1). Using Oracle's pre-compilers, Pro*Fortran and Pro*C, EDF engineers can embed structured query language (SQL) statements in Fortran and C programs on their Cray Research computer systems, allowing them to perform writes and reads to and from the DB2 databases over the network.

EDF wanted its engineers to use data stored in the DB2 systems and ORACLE systems without having to know the location of the data. An additional goal was to improve usage of network resources by distributing applications away from central database servers to allow client applications to run on dedicated CPUs for a faster response rate. By using ORACLE's SQL*Net network software, the data could be distributed over multiple, incompatible networks and computers, yet appear to users as if it resided in a single database on a single computer.

Performance

Once the client/server architecture was established, EDF wanted to test a real-world code's performance under this architecture and measure the results quantitatively, based on CPU performance, and qualitatively, based on the goals of transparency, flexibility, and integrity.



Today's RDBMS support Cray Research's vision of network supercomputing. Networked access allows the front-end tools to reside on one system while the server is located elsewhere in the network. Distributed databases allow scientists and engineers to use several databases residing on a variety of hardware platforms as if they were part of a single conceptual database.

The client/server RDBMS architecture is available on Cray Research systems through products including the INGRES/NET product from ASK Computer Systems, and Oracle's SQL*Net. The SQL client (often called a "front-end" tool) can run on a Cray Research or other system while the database server (often called a "back-end") runs across the network on a remote Cray Research or other system. This capability is an excellent example of the Cray Research network supercomputing envi-

ronment, in which the user has complete flexibility to determine where the RDBMS back-end or front-end should reside.

RDBMS gateways permit access to data residing in one vendor's database using another vendor's front-end tools. EDF, for example, uses Oracle's SQL*Connect to communicate with ORACLE front-end tools running on Cray Research supercomputers and DB2 on IBM systems. A similar product for INGRES, INGRES Database Gateway to IBM/DB2, exists for Cray Research supercomputers from Ingres. All of these products help make the Cray Research vision of network supercomputing a reality—delivering easy, reliable, and cost effective access to the Cray Research systems from any node in a network.

To test this architecture, EDF's engineers chose a large data volume application, GEODE, an electric power consumption forecast application. GEODE simulates the balance between power consumption and production on a national scale. It factors in variables, such as the predicted weather, to decide, for example, how much power to produce and when maintenance can be performed. The more accurate and faster this simulation, the greater the savings or avoided losses.

Prior to the implementation of this architecture, the GEODE application had to be run on the IBM system, or file transfers had to be performed to move the processing to the Cray Research system. EDF wanted to be able to easily take advantage of the speed of the Cray Research system without awkward file transfers.

Test results

The quantitative results were drawn from a comparison of two architectures:

- ☐ A client/server architecture with one computation program running on a CRAY Y-MP system with data managed by a DB2 database on an MVS IBM machine with communications handled by SQL*Net and SQL*Connect
- ☐ A centralized architecture where the computation is processed on the dataserver (The file transfer is not taken into account in the centralized architecture to enable CPU execution time comparisons.)

Quantitative results in Table 1 show the client/server architecture to have a performance advantage over the DB2 system alone. The results are somewhat tentative, however, in that Cray Research system CPU time and IBM system CPU time are being compared, and the CPU time used for the software communication module based on TCP/IP is not included.

Additionally, significant amounts of data, up to 200 Megabytes of integer and floating point

numbers at a time, are extracted and transferred from IBM's DB2 system for processing on the Cray Research system. Some of these results are then reinserted into the DB2 database.

Qualitative results are more definite, however. The client/server architecture was able to deliver

- ☐ Flexibility. Data can be accessed conveniently without running intermediate jobs.
- ☐ Transparency. The remote IBM/DB2 data are processed as if they were local.
- ☐ Integrity. The RDBMS efficiently handles concurrency control and transaction processing.

Future directions

Today, interest in other implementations of database client/server architecture in the Clamart Computer Center is growing. An application for nuclear fuel loading also is in development using this capability between Cray Research supercomputers and IBM/DB2 systems. More and more engineers at EDF believe that the client/server architecture is the most rational solution for using the Cray Research system in a DB2 RDBMS environment.

About the authors

Michel Demuynck is the group leader and research engineer in the Database Research Group of the Department of Studies and Research at EDF. He is a graduate of the National University for Computing and Applied Mathematics of Grenoble.

Philippe Bédou is a research engineer in the Database Research Group of EDF. He is a graduate of the Faculty of Science at Pierre and Marie Curie University, Paris.

Denton Olson is a senior applications analyst in the Database Management Systems (DBMS) group of the Industry, Science, & Technology Department at Cray Research, Inc. He received his M.S. degree in Computer Science from North Dakota State University.

Numerical analysis of airflow in aircraft cabins

Fred Aboosaidi, The Boeing Company, Seattle, Washington

Accurate predictions of airflow in aircraft cabins are of prime importance in the commercial aerospace industry because passenger comfort depends on the precise control of temperature, air distribution, and air quality within the cabins. The design of an airplane's environmental control system involves the specification of supply and return airflow quantities, speeds, and locations to yield an appropriate air distribution system. Full-size airplane cabin mockups continue to be used in the design process, but the expense and time required for testing limits the number of airflow systems that can be analyzed in this way.


Computational fluid dynamics (CFD) techniques accelerate and reduce the expense of evaluating designs, thereby improving the potential for complementary design studies. Unlike exterior aerodynamic surfaces, where the solution of simplified defining equations can yield a great deal of helpful information, the Navier-Stokes equations are the appropriate equations for defining interior ventilation flows, because the flows are viscous and recirculate extensively. The calculations described here were performed on the CRAY Y-MP system at Boeing Computer Services. Visualizations were performed on a Silicon Graphics IRIS workstation using Cray Research's MPGS software and the PLOT3D and FAST public-domain programs.

Although two-dimensional analysis sometimes can be used, we were concerned with the application of a Navier-Stokes code to three-dimensional aircraft cabin configurations that were also the subject of a series of experimental tests. The results have been used to evaluate the accuracy of the predictions, so that the appropriate system

of CFD analysis and experiment can be devised to best assist in the timely and cost-effective design of the environmental control system in an airplane cabin.

The prediction of the flow field in an aircraft cabin is a substantial challenge for a CFD code. The configurations typically are complicated by the cabin geometry and the presence of passenger

Figure 1. Schematic of the airplane cabin interior.







Figure 2. Detail of the nozzle configuration for the cabin air supply.

Figure 3. Computational grid for the nozzle.



seats. In addition, large differences exist between the dimensions of the supply air nozzles and those of the main cabin. A meaningful length scale in a supply nozzle can differ from that of a cabin by a factor of 100 or more. The flow paths include complex flow features such as free and wallbounded shear layers and recirculation regions. Thus, to represent the appropriate level of physics, the full form of the Navier-Stokes equations should be solved on a mesh of sufficient resolution to predict the turbulent flow characteristics near the jet mixing regions as well as the extensively recirculating flow elsewhere in the cabin. The use of periodic boundary conditions allows the solution of one seat row as a subset of a configuration with a large number of seat rows. The CFD code used in these applications was the FLUENT code from Fluent, Inc.

The experimental configuration

A 3.5 m mockup, shown in Figure 1, with an experimental stowage bin/ceiling configuration, was the subject of a comprehensive experimental testing program. Seating in the mockup was nine abreast with an 86 cm pitch. A nozzle configuration supplied air downward from the ceiling to be exhausted through floor outlets. Air velocities were measured in a matrix of measurement densities required to resolve all flow features properly. The air is supplied from a plenum and locations in several longitudinal (along the aisle) cross sections. Finally, smoke was released in various locations to facilitate visualization of the airflow patterns and identification of any stagnant pockets. This visualization allowed the recording of velocity vector

directions in the airflow at the predetermined measurement locations.

The numerical analysis focused on the investigation of the flow distribution in the aircraft cabin exclusive of heat and contaminant transport. Passengers and luggage also were neglected in the numerical model. The dimensions of the air supply




Figure 4. Slice of the computational grid for the full cabin model.

nozzle array are much smaller than the overall dimensions of the aircraft cabin. For this reason, a separate analysis of the flow field in the vicinity of the nozzles was performed. The results of this small scale analysis near the vicinity of the nozzles were analyzed in their own right, then input to a larger-scale model of the main cabin.

Nozzle analysis

The nozzle configuration is shown in Figure 2. Although the configuration is three-dimensional, with the array of inlets arranged in an alternating pattern along the length of the aircraft cabin, the high aspect ratio of each nozzle justifies a two-dimensional approximation. The two-dimensional model can predict the relative flow distribution and velocity magnitudes in the multiple flow paths created by the nozzle array. Because the alternating inlet patterns are mirror images of each other, only one arrangement was studied. In the main cabin airflow analysis, inlet velocity components for the two nozzle patterns can be supplied as boundary conditions based on the analysis of a single nozzle configuration.

The inlet nozzle configuration was defined using geometry modeling tools. Based on the geometry description, a body-fitted grid was set up using the solution to a set of elliptic grid generation equations (Figure 3). The complex, multiply connected domain was discretized using a 121 by 63 grid. The choice of grid size arose from a grid dependency study in which solutions based on a range of grid densities were used to examine flow structures and determine where the grid goes through the nozzles before exiting to the cabin. Consequently, the inlet velocity at the plenum is treated as uniform with a vertical velocity of 1.29 m per second. The flow is modeled as fully turbulent. The pressure in the main cabin is nearly constant, leading to the choice of a pressure-specified boundary condition for the outlet in the numerical model.

Analyses performed of the nozzle flow field, rendered as velocity vectors and streamlines,

revealed several airflow characteristics, some of which could be applied to an improved nozzle design. The total air supply is divided about equally among the three nozzles; this is expected because the three nozzles have approximately the same cross section and flow resistance. The sharp change in flow direction at the entrance to each nozzle passage creates an adverse pressure gradient. Consequently, the flow separates and forms recirculation zones. Additional separation regions also form in the nozzle passages due to sudden area changes. As the flow leaves the nozzles, the air jets entrain the surrounding air and two recirculating regions are formed on either side of the jets. In the two side nozzles, the air streams are further accelerated due to the converging geometry, whereas the air stream in the middle nozzle expands just before it exits.




The computed flow distribution suggests several conclusions about this nozzle design. First, the expansion regions in the middle nozzle appear to be unnecessary. Because flow recirculation entails loss of energy, more pumping power may be required in the current design than is necessary. Second, the side nozzles have a baffle near the exit that prevents the air from entering the passenger cabin more uniformly. As a result, all three air jets enter the cabin in the same downward direction. Finally, smoothing the nozzle entrances to eliminate the sharp corners could reduce the amount of recirculation.

Main cabin analysis

The second part of the numerical simulation models the airflow in the main aircraft cabin. The results of the two-dimensional nozzle airflow simulation were used to provide the inlet boundary conditions for the three-dimensional model of the aircraft cabin.

As in the two-dimensional nozzle analysis, a two-equation turbulence model was used in which transport equations for the kinetic energy of turbulence and its rate of dissipation are solved in addition to the momentum, pressure, and continuity equations. A 58 by 59 by 33 mesh was used to model half of the cabin (using the symmetry inherent in the geometry). Due to the periodicity of the seat arrangement, only one row of seats is modeled. A slice of the grid is shown in Figure 4, at a section corresponding to a location between the seats. The grid lines are concentrated in the vicinity of the inlet jets to better capture the shear layers formed by the high velocity jets and the lower velocity air circulation within the cabin.

Figure 5 shows the results of the three-dimensional simulation presented in vector plots for this cross section and the streamlines determined from the experiment. While the experimental result lacks resolution, the predicted flow patterns closely match those determined from the experimental test configuration. Figure 6 shows contours of velocity magnitude from the solution at the same cross section and the corresponding experimental results. The results indicate a good agreement between experiment and prediction, especially in the critical regions encompassing



draft-sensitive areas around the passengers. Large regions of stagnant air appear to be absent in any of the areas surrounding seated passengers. The results show that the flow field in the cabin is characterized by two large vortices—one on the inboard side and the other on the outboard side.

Figure 7 shows the solution (in the form of velocity vectors) for a section that corresponds to the seat cushion, ahead of the seat back and at a cross section that cuts through the back of the seat. The solution in the first case reflects a no-slip condition for the flow at the seats; the flow patterns in the second case are affected somewhat by the presence of seats, but the effect appears to be significant only in areas very close to the seats. These images, in addition to pointing out the predominant flow features and velocity levels, show the behavior of the flow in the vicinity of the nozzle configuration. The flow enters the domain through the array of inlet nozzles and exhausts through an outlet located on the outboard wall below the seat next to the wall. The vector plots in the three slices show

that the staggering of the inlet array leads to different magnitudes and directions of the flow at different locations. Smaller secondary vortices are formed in the cross section near areas such as the overhead luggage compartment. Each component of the multicomponent nozzle is evident; this configuration appears to contribute to a favorable flow pattern without developing unreasonably large velocities in the vicinity of passengers.

Conclusions

This analysis predicted the overall flow patterns in the cabin configuration, while velocity magnitude was predicted to an acceptable level over the majority of the cross section. The analysis also showed the applicability of the CFD analysis in the microanalysis of the nozzle configuration. It was shown not only that the flow features within the nozzle region can be predicted in detail through preliminary analysis, but also that analysis can be used in the full cabin application by predicting the characteristics of the jet nozzle flow where it enters the main cabin, allowing their use as boundary conditions for the main cabin flow prediction.

The potential exists for future application to cases where pollutant dispersal and/or heat transfer are present and important. Examples of the latter might include hot or cold walls and hot or cold supply jets in conjunction with baseline velocity conditions at the jets. In addition, more realistic modeling could be analyzed, with possible consideration of end effects (such as from galleys), below-seat luggage stowage, and passenger occupancy. Including real passengers would require an examination of the effects of heat transfer (such as from hot persons to cooler cabin air) and contaminant dispersal (for example, exhaled breath). These additional physical processes and geometric models represent goals for future CFD analyses as software tools and computer resources continue to mature.

Acknowledgment

The author thanks N. A. Nelson, chief engineer of mechanical/electrical engineering at the Boeing Commercial Airplane Group, for his continuing support and encouragement, and Scott Ellsworth, Alan Dejager, Raymond Horseman, and Matt Warfield of Boeing for their assistance in the cabin air experimental program and analysis/measurement comparisons for the reported case.


About the author

Fred Aboosaidi is a computational fluid dynamics engineer in the environmental control systems research and development department at the Boeing Commercial Airplane Group. He previously worked in hypersonic and supersonic research at Boeing. Aboosaidi received a B.S. degree in physics from Michigan State University and M.S. degrees in aerospace and nuclear engineering from the University of Michigan at Ann Arbor. He also was a Ph.D. candidate in aerospace engineering at the University of Michigan and conducted research in aerospace engineering at the University of Washington, Seattle.

Quantum ballistic transport in a model semiconductor nanostructure

Cengiz Besev, Zhen-Li Ji, and Karl-Fredrik Berggren, Linköping University, Linköping, Sweden

Figure 1. A schematic diagram of the $\text{Al}_x\text{Ga}_{1-x}\text{As}/\text{GaAs}$ heterojunction FET used to create a high-mobility two-dimensional electron gas (2DEG) at the semiconductor interface.



A growing challenge in semiconductor physics and microelectronics is to design semiconductors that are smaller than ever before. In our effort to meet this challenge, we at the National Supercomputer Centre at Linköping University in Linköping, Sweden, are relying on the computational power of the CRAY X-MP supercomputer. With the help of this powerful tool, we can visualize animated sequences for electron densities and velocity distribution data, allowing us to gain insight into the complex and volatile current flow present in quantum nanostructures.

Steady improvements in thin-film growth techniques have made it possible to grow the high-quality semiconductor heterostructures necessary for this research. Techniques such as molecular beam epitaxy and metal-organic chemical vapor deposition enable the growth of layered structures of semiconductor materials. Potential profiles and impurity distributions may be designed with a dimensional control close to interatomic spacing and with virtually defect-free semiconductor interfaces in a lattice-matched case such as the alloy $\text{Al}_x\text{Ga}_{1-x}\text{As}$ grown on GaAs (typically $x \sim 0.3$). Technological advances of this kind are of great importance in device applications. In transistors, for example, as the lateral feature size is made smaller, interfaces and doping profiles must become more precise.

The gradual improvement in making sharper semiconductor interfaces has opened a new research area: the study of artificial two-dimensional quantum structures in which carriers move with extremely high mobilities at low temperatures.¹⁻³ Mobilities in excess of $10^6 \text{ cm}^2/\text{Vs}$ may be achieved, implying that the carriers travel several μm before colliding with an impurity or imperfection. In semiconductor structures or devices with

small features, quantum effects will emerge at the forefront. Much of the new physics and new device concepts associated with artificial, layered structures such as superlattices, quantum wells, and resonant tunneling diodes, have been demonstrated first with $\text{Al}_x\text{Ga}_{1-x}\text{As}/\text{GaAs}$. From a technological point of view, it is important that high quality $\text{Si}_{1-x}\text{Ge}_x/\text{Si}$ is now explored in the same context, but we will limit our discussion to $\text{Al}_x\text{Ga}_{1-x}\text{As}/\text{GaAs}$.

Although investigations up to this time have used basic semiconductor structures that obey simple physical laws, simplicity does not imply triviality in terms of the outcome. On the contrary, these very complex and volatile phenomena demand visualization for displaying and further analyzing the physics involved. We networked SUN SPARCstations and CA-DISSPLA software together with the CRAY X-MP system at Linköping to visualize animated sequences for both electron densities and velocity distributions data. The visual output, which was videotaped, shows electron densities and current flow as a function of electron energy.

The tunable two-dimensional electron gas in a high-mobility field-effect transistor (HEMT)

Although the general framework should apply to other high-quality semiconductor interfaces as well, here we discuss the case of $\text{Al}_x\text{Ga}_{1-x}\text{As}/\text{GaAs}$. Figure 1 shows schematically a single heterojunction field effect transistor (FET) used to create a high-mobility two-dimensional electron gas at the semiconductor interface. Nominally undoped GaAs is grown on a substrate, followed by a thin undoped $\text{Al}_x\text{Ga}_{1-x}\text{As}$ layer and a heavily n-doped region. The purpose of the thin layer is to physically separate the electron gas at the interface from the scattering centers, mainly ionized donors, in the n-region. The metal gate makes it possible to vary the density of the electron gas at the interface. At sufficient voltage, all electrons are actually removed.

Due to a mismatch of bandgaps in the two semiconductor materials, electrons from the n-doped $\text{Al}_x\text{Ga}_{1-x}\text{As}$ region will spill over to the GaAs side where they will be trapped at the interface by restoring electrostatic forces. Electrons then are forced to move only along the interface. The electron gas formed at the interface therefore behaves dynamically as a two-dimensional system. Typical electron density is from $\sim (4-5) \cdot 10^{11} \text{ cm}^{-2}$ to zero depending on the gate voltage V_g that may be applied to the metal gate. The electron density thus may be varied continuously, simply by changing the V_g . We may even remove all electrons under the gate by applying an appropriate gate voltage. This opens possibilities to create a new type of nanostructure based on a lateral confinement of the two-dimensional electron gas (2DEG) trapped at the interface.

Creation of quantum wires and other nanostructures by lateral confinement in a HEMT

The structure in Figure 1 forms the basis for the more recent developments that have led to

new physical discoveries and high expectations on potential device applications. The key to this is the additional confinement that may be induced by means of current lithographic and etching techniques in semiconductor technology. Figure 2 illustrates in schematic terms one of the principal ideas for achieving a lateral confinement as well. If parts of the metal gate are removed by the above techniques, one can tune the gate voltage and confine the electrons to the ungated regions. If a cross is etched into the gate (as in Figure 2), one is able to create two very narrow, crossing




Figure 2. A schematic diagram of the $\text{Al}_x\text{Ga}_{1-x}\text{As}/\text{GaAs}$ heterojunction FET, with a cross etched in the metal gate.

channels at the interface. Because of the high quality of the semiconductor interface, these channels will be fast. In the same way one can shape the 2DEG into other, more complicated nanopatterns like rings, dots, cavities, lattices, and networks by writing the corresponding pattern on the gate. Because these structures are so small, quantum mechanical effects are strongly manifested and electrons then move as matter waves rather than as classical particles. For example, in a narrow channel the wavelength of such a particle can be of the same order as the width of the channel itself. Hence quantization of the transverse motion becomes an important issue; there will be distinct modes or subbands corresponding to each transverse standing wave. Electrons then propagate in one-dimensional subbands, of which only a few may be occupied at low temperatures. For high-mobility samples, such as $\text{Al}_x\text{Ga}_{1-x}\text{As}/\text{GaAs}$, we then can compare quantum wires or electron wave guides with optical wave guides. Such discussion approaches the future goal of fabricating nanostructures that obey quantum mechanical laws and integrating such structures




Figure 3. Schematic representation of 2DEG shaped into a crossbar structure achieved by means of a patterned gate as in Figure 2.

into larger units, perhaps a "quantum chip." Transistors in such a chip could be a thousand times smaller and orders of magnitude faster than they are today. They also would consume less power and allow scientists to use interference effects and tunneling in cases when an electron moves without the usual interconnects. Other potential applications involve lasers and optoelectronics based on one-dimensional quantum wires and zero-dimensional quantum dots.

The case of two intersecting quantum wires

Many of the proposals for new devices are speculative or merely enthusiastic overstatements. It probably will be many years before these applications become practical. Much work is required to improve fabrication techniques and to understand the fundamental physics involved. Actually the physics aspect is most interesting because we are dealing with a new low-dimensional system with unique properties. The physics associated with such systems is worth studying on its own merits. For this reason we have chosen to study an elementary structure consisting essentially of two intersecting quantum channels. A remarkable feature of a simple cross of this kind is that bound states reside at the intersections of the channels.^{4,5} From a classical point of view the potential is open and therefore could never trap a classical particle.

The existence of bound states in such a potential is thus due entirely to the wave-mechanical nature of the electrons. In principle, the presence of trapped electrons and interference effects in intersecting channels opens the possibility of novel devices based on quantum mechanical tunneling and/or interference effects. We therefore have investigated electric transport in a model semiconductor structure or device that exploits the quantum aspects of a cross.

Mathematical modeling

Our model device is shown in Figure 3. It consists of two short intersecting, perpendicular channels, one of which connects to two 2DEG regions acting as emitter and collector regions. This crossbar structure is large enough to embrace the bound states residing at the intersection. The crossbar area connects two 2DEG regions which serve as emitter and collector when a weak electric field is applied. Then electrons incident on the channel area are either reflected or transmitted as indicated by the arrows. The result is a net current. When a weak potential difference is applied between emitter and collector, an electric current will flow in the crossbar region. The dimensions of this region are ~ 50-100 nm and therefore smaller than the typical distance between collisions or the mean free path of the two-dimensional gas mentioned above. Hence the transport through the channel region is expected

to be ballistic at low temperatures, meaning no collisions take place. A first choice for the confining potential can be fairly simple because we need not concern ourselves with quantitative aspects to observe how the physics works. We therefore solve the one-electron Schrödinger equation exactly, assuming infinite square wells for the confinement. For an electron with energy E and mass m^* we thus solve the two-dimensional Schrödinger equation ($\hbar = h/2\pi$ is Planck's constant)

$$-\frac{\hbar^2}{2m^*} \left\{ \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \right\} \Psi(x,y) = E\Psi(x,y) \quad (1)$$

The wave-function $\Psi(x,y)$ has to vanish at boundaries. The solution $\Psi(x,y)$ also should express that electrons impinging on the crossbar structure from the emitter region may either penetrate the structure or be back-scattered. For the region to the left in Figure 3 we therefore write⁵

$$\Psi_L(x,y) = e^{ik_x x + iky} + \int_{-\infty}^{\infty} dk' A_L(k') e^{-ik'x + ik'y} \quad (2)$$

where the wave vectors $\mathbf{k} = (\kappa, k)$ and $\mathbf{k}' = (\kappa', k')$ for incident and reflected waves, respectively, are related to the energy as $E = \hbar^2 \mathbf{k} \cdot \mathbf{k} / 2m^* = \hbar^2 \mathbf{k}' \cdot \mathbf{k}' / 2m^*$ and A_L to the probability for back-scattering. Similarly, an electron entering the collector, the right region of Figure 3, is represented by

$$\Psi_R(x,y) = \int_{-\infty}^{\infty} dk' A_R(k') e^{ik'x + ik'y} \quad (3)$$

Within the cross itself it is convenient to define different sections as indicated in Figure 3. Within each section the solution is expanded in terms of the exact solutions satisfying the boundary conditions. For example, in channel (1) of width w , we would have

$$\Psi_1(x,y) = \sum_n (B_n^1 e^{iq_n x} + C_n^1 e^{-iq_n x}) \sin\left(\frac{n\pi}{w}(y+w/2)\right) \quad (4)$$

and similarly for the other sections. The sinusoidal function describes transverse standing waves with energies $E_n = \hbar^2 (n\pi/w)^2 / 2m^*$ with $n=1,2,\dots$. The same type of modes are found in other types of wave guides, such as optical wave guides. The exponential functions in Equation 4 describe motion along the channel. The energy of the electron is therefore $E = E_n + \hbar^2 q_n^2 / 2m^*$. If the energy of the incoming electron is less than the subband threshold E_n , the quantity q_n will be imaginary, and there will be exponentially damped and growing solutions. Typically such solutions appear when tunneling takes place. The coefficients in the expansions above are determined from the requirement that functions and derivatives must be continuous at the boundaries between the different sections. If there are N_{max} modes included in the expansions of $\Psi(x,y)$, there will be $10 N_{max}$ linear equations determining the expansion coefficients for each impact angle.⁵

We next consider the current induced by a voltage drop between emitter and collector regions, as shown in the two 2DEG reservoirs in Figure 3. The conductance $G=J/V$, the total current divided by voltage V , is obtained by evaluating the expect-






Figure 4. The conductance G at different energies of the carriers for the model quantum structure of Figure 3.

tation value of the quantum mechanical current operator, using the wave-functions described above.⁵ Since we limit our discussion here to small voltages and low temperatures (0 Kelvin, so far), only electrons at the Fermi energy E contribute to the net current. (The Fermi energy is the energy of the highest occupied energy level.)

Electrical conductance and effects of quantum mechanical tunneling and interference

When the energy E of the incoming electrons is varied, the conductance undergoes strong variations. In an experimental situation we would rather keep the energy fixed and vary the width of the channels by means of the applied gate voltage. In a qualitative way the physics is, however, the same in this case. Thus we find an intricate pattern of resonances or tunneling associated with bound states and antiresonances due to interference effects as shown in Figure 4. The voltage drop V across the channel region is assumed to be small, and the conductance at 0 Kelvin is evaluated using the exact solutions to the Schrödinger equation. We first note that the conductance vanishes at the lowest energies. We normally might expect this to happen below the lowest subband threshold $E_1 = \hbar^2 (\pi/w)^2 / 2m^*$. However, just below this mode we observe a strong resonance peak (denoted E_{B1}) which is due to resonant tunneling via the bound state at the center of the cross as discussed above. The tunneling characteristics are similar to resonant tunneling in a double barrier bipolar transistor but with the remarkable difference that there are no

Figure 5. Images of electron and velocity distributions obtained by solving the Schrödinger equation in a model of two intersecting quantum channels of finite length. The left column of four pictures shows the contours of electron densities. The highest probability region has been designated as dark red and the lowest region as magenta, by scaling down 45 levels from the maximum relative values: 45.89, 21.94, 5.79, 61.08 from top to bottom pictures, respectively. The middle column shows the associated quantum mechanical velocity distributions where the highest vector velocity has been designated as violet and the lowest velocity as blue, by scaling down seven levels from the maximum relative velocity value as indicated in the right side of each picture. The density and velocity distribution for a given electron energy and its associated conductance has been indicated (red shaded curve) in the right column.



physical barriers to be penetrated. Another remarkable feature is that the conductance peaks at $G = 2e^2/h$, where h is Planck's constant. The resonance at position E_{B2} is associated with resonant tunneling via the second bound state associated with an infinite cross. This resonance is found just below the second subband threshold $E_2 = \hbar^2(2\pi/w)^2/2m^*$.

In Figure 4 we also notice other sharp resonance peaks ($E_{R1}, \dots, E_{R4}$ and higher). These peaks appear in addition to E_{B1} and E_{B2} because the cross is finite. Thus the resonances at $E_{R1}, \dots, E_{R4}$ are related to standing waves associated mainly with the vertical stubs. In contrast to the two first resonances, the position of these resonances exhibits a very sensitive dependency on the extension of the stubs.

The resonances are superimposed on a broad background that corresponds to normal transport in the different subbands. Due to quantum mechanical interference, the conductance vanishes at certain energies E_{I1} , E_{I2} and higher. The energies at which destructive interference occurs are sensitive to geometry. This feature may be of interest for device applications, possibly in the form of a quantum interference transistor.

Electron distributions

The series of graphs in Figure 5 shows some typical electron densities associated with the current flow. The first two cases show the tunneling via the bound states characterizing an infinite cross. Clearly the wave functions are well localized to the interior region of the crossbar structure. The third graph shows the electron density for an energy just slightly less than E_{I2} at which the channels are switched off due to interference effects. The electrons accumulate essentially in the channel to the left because interference is partly blocking the channel. At E_{I2} the channel is completely blocked, but electrons still penetrate the channel to the left as above. The fourth graph shows the electron cloud at the resonance E_{R3} . This resonance is due to a vertical standing wave across the stubs.

Velocity distribution transition from laminar to vortex flow

The corresponding quantum mechanical current distributions show that the flow of carriers is laminar at low energies. As the energy is increased to the extent that interference becomes important, vortices appear in the current, a new observation. Upon further increase of the energy, the vortices will form an increasingly complex and volatile pattern. The current distributions accompanying the different charge distributions discussed above also are given in Figure 5.

Although the origin of the vortices is quantum mechanical in this case, we are dealing with a situation similar to the transition from laminar to vortex flow in classical liquids and gases. In such cases the Cray Research system is an integral tool for analysis, allowing us to visualize the results for further study. In particular, very narrow, δ -function-like resonances are very hard to find without access to a supercomputer. As the energy is varied

in small steps, visualization of spatial current distributions and electron densities shows clearly in advance of an anticipated resonance.

Conclusion

* Our investigation of a simple semiconductor structure that obeys simple physical laws has yielded results that are complex and intricate indicators of future applications in semiconductor technology. At this juncture, access to a supercomputer is indispensable because of the large number of frames that have to be generated in a model semiconductor nanostructure. Supercomputer resources at the National Supercomputer Centre at Linköping University have allowed us to visualize animated sequences for both electron densities and velocity distribution data obtained by solving the Schrödinger equation on a CRAY X-MP/416 supercomputer. By using SUN SPARCstations 1/370 (Sunview) and CA-DISSPLA software, we visualized the output and stored the sequences in intermediate storage in the form of Component Recording Video (CRV), using a Sony LVR-6000/LVS-6000P laser videodisc recorder and processor. ■

About the authors

Cengiz Besev is a graphics software analyst for scientific visualization at the National Supercomputer Centre at Linköping University. He studied physics at the University of Ankara, Turkey, and continued his studies at Uppsala University, Sweden. Prior to joining the National Supercomputer Centre, he was a staff member at the European Centre for Medium-Range Weather Forecasts, Reading, England.

Zhen-Li Ji is on leave from the Department of Physics, Guangxi Teachers University, Guilin, Guangxi, People's Republic of China. He has been involved in the research of semiconductor nanostructures during the past two years as a research student at Linköping University.

Karl-Fredrik Berggren received a Ph.D. degree from Uppsala University in Uppsala, Sweden, and is now professor of theoretical physics at Linköping University. He has served as director of the National Supercomputer Centre, and is a member of the Swedish Engineering Research Council.

References

1. Beaumont, S. P. and C. M. Sotomajor Torres, eds., *Science and Engineering of One- and Zero-Dimensional Semiconductors*, Plenum, 1990.
2. "Nanoscale and Ultrafast Devices," a special issue of *Physics Today*, February 1990.
3. Berggren, K.-F., "Quantum Phenomena in Small Semiconductor Structures and Devices," *International Journal of Quantum Chemistry*, Volume 33, p. 217, 1988.
4. Schult, R. L., D. G. Ravenhall, and H. W. Wyld, "Quantum Bound States in a Classically Unbound System of Crossed Wires," *Physics Review B*, Vol. 39, p. 5476, 1989.
5. Berggren, K.-F. and Zhen-Li Ji, "Resonant Tunneling via Quantum Bound States in a Classically Unbound System of Crossed, Narrow Channels," *Physics Review B*, Vol. 43, p. 4760, 1991.

CASE STUDY

Parallel processing with **ATExpert**

*Margaret L. Simmons, Los Alamos National Laboratory, Los Alamos, New Mexico
Kenneth M. Lord, Cray Research, Inc.*

When the first multiprocessor supercomputers became available from Cray Research in the mid-1980s, researchers at the Los Alamos National Laboratory set out to parallelize application codes to take advantage of the parallel processing capabilities of the new machines. The task was difficult for several reasons, including the lack of software tools and the many potential paths that could be followed to partition a given code for parallel execution. Cray Research addressed some of the difficulties in 1988 when it included the Autotasking automatic parallelization feature in the company's Fortran compiler. Now Cray Research has made available an Autotasking performance tool, called the Autotasking Expert System (ATExpert), to help users diagnose obstacles in their codes that prevent them from getting the maximum benefit from Autotasking. We have used this tool to help parallelize the Los Alamos hydrodynamics code PUEBLO with impressive results. Our goal was to parallelize the code to obtain the best possible speedup by using information provided by the tools and under the constraint that neither the algorithm nor the code could be altered.

When used in conjunction with Cray Research's CF77 Fortran compiling system and the Cray Research tools SCOUNT, PROF, and PROFVIEW, ATExpert produces parallelized code plus information to help programmers further

optimize the code. The tool uses an X Window System interface to present the information in an easy-to-read graphical format, making the large quantity of data easier to understand. It also provides an ASCII interactive interface and a batch user interface. ATExpert is included in version 6.0 of Cray Research's UNICOS operating system.

Overview of PUEBLO

The PUEBLO code is highly vectorizable, but not parallelized. In using ATExpert to parallelize it, we found that the tool is indeed easy to use, and without a large effort we obtained (predicted) speedups of 6.4 to 7.2, depending on problem size, on an eight-processor CRAY Y-MP system. To test the accuracy of this tool, we ran a large problem and a small problem on a dedicated CRAY Y-MP8/128 system using various numbers of processors and measured the actual speedup using the CF77 compiling system with the Autotasking feature enabled. When we compared the results of these tests with ATExpert's predictions, we found that the differences were less than 10 percent at all levels of optimization.

The PUEBLO code numerically solves the equations of transient, multimaterial, compressible fluid dynamics. It can treat material strength and plasticity, high explosive detonations, and a k-ε model of turbulence. The code uses a time-explicit,

conservative finite-volume numerical technique in which all hydrodynamic variables, including velocities, are cell centered. This technique is based on a form of the Gudunov method, which uses a Riemann solver. The method is used to achieve tight coupling between velocity and pressure fields. The hydrodynamics cycle is split into a Lagrangian phase and a rezone-advection phase in which conserved quantities are transferred from the Lagrangian mesh to an arbitrarily specified mesh. PUEBLO is highly vectorizable, with primary vector lengths equal to the cube of one dimension of the mesh. We ran our problem using two different mesh sizes: 32 by 32 by 32 and 64 by 64 by 64.

Overview of Cray Research tools

SCOUNT is a benchmarking utility that counts the number of times each statement in a Fortran program is executed and produces a source listing with an execution tally next to each line of code. We used SCOUNT to ensure that during the initial phases of optimization our efforts would be concentrated on those loops that the problem actually executed.


Through a method of timing by address range, the PROF utility indicates the time spent in various segments of code within routines. At regular intervals, the operating system records the address of the instruction being executed. Addresses are grouped in "bins" or "buckets" with selectable sizes; these bins can be associated with labels internal to a program.

The PROFVIEW utility generates reports in various formats from the raw data generated by PROF. Beginning with release 6.0 of the UNICOS operating system, PROFVIEW provides an X Window System interface.

ATEXpert measures and displays information about the Autotasking performance of a job that is run on an arbitrarily loaded system. It can predict speedups that would result during dedicated execution from data collected during a single run on a nondedicated system. It provides a variety of information about the code under consideration, which enables programmers to find code segments that may be contributing to performance bottlenecks.

ATEXpert comprises three phases: an instrumentation phase, a data-gathering phase, and an analysis phase. During the instrumentation phase, the FMP preprocessor (from the CF77 compiling system) adds timing code to parallel regions that is compiled into the user's program. During the data-gathering phase, the program is executed, and raw timing information is gathered. The instrumentation also records the number of untasked scalar iterations for each loop, the number of concurrent iterations for each parallel loop, and other relevant information associated with each loop. When the program terminates, this information is written to a file. In the analysis phase, this file is read by ATEXpert, which displays the collected Autotasking performance data.

ATEXpert divides the execution time of a program into parallel region time and preceding serial time (time spent outside of parallel regions).



Parallel region timing is obtained for both unitasked and multitasked execution of the region. Multitasked execution is subdivided into a measurement of the distribution of work among processors and a measurement of overhead costs for parallel execution.

ATEXpert obtains the fraction of code that can execute in parallel from exploited parallelism rather than existing, or potential, parallelism. A user can change this fraction from potential to detected parallelism by inserting the appropriate directives. According to ATEXpert, overhead is the difference in speedup obtained by the application of Amdahl's Law and the difference measured (or predicted) from an Autotasking run. Overhead calculated (or projected) in this way is subdivided into various contributing factors such as the overhead associated with the beginning of a parallel region, slave arrival, convoy time, and other events. The definition of overhead as the difference between predicted and measured speedup was proposed by Buzbee as an extension to the Ware model.¹

Method and results

Initial runs using SCOUNT provided information about which loops were actually being executed by our problem. After studying this information, we began using PROF and ATEXpert. Our first PROF run showed that the ISMIN routine consumed 11.7 percent of the run time (Figure 1). When we allowed the compiler to replace our ISMIN with the more efficient SCILIB version, the time spent in the routine dropped to an insignificant amount and the total execution time improved by 10 percent. Figure 2 shows the output from PROF that resulted from this change. Next, we ran this improved version through the CF77 compiling system with the Autotasking feature enabled. Figure 3 shows that the subroutines RIEMAN and ADVECT still were heavily used, but LAGHYD and VOLUME had replaced TIMSTP and AREAPJ as the routines in which the code spent most of

Figure 1. Partial graphic output from prof/profview showing the percent of execution time spent in the primary modules when using the original ISMIN routine.





Figure 2 (above). Partial graphic output from prof/profview showing the percent of execution time spent in the primary modules when using Cray Research's ISMIN routine.

Figure 3 (above right). Partial graphic output from prof/profview showing the percent of execution time in the primary modules when using Autotasking.

its time. When subroutines are run in parallel, they account for less of the total execution time, and serial subroutines come to dominate the remaining execution time. Figure 3 showed us where we needed to look to improve parallelism further. In this code, RIEMAN dominates in serial and parallel modes; it is clearly the most important subroutine in the code in terms of execution time.

Armed with this information, we ran the code under ATExpert. The results are shown in Figure 4. The plot on the left side of the figure shows that the predicted speedup is only 2.5 out of a

performing poorly. Clicking on "Source Files" in the command menu allowed us to bring up a window containing source code for RIEMAN, which led us to target loop 20. Figure 5 shows the fragment of code that represents this loop. As soon as we saw the code, we saw the problem. By default, the Autotasking system tries to vectorize inner loops and multitask outer loops. The outer loop in this piece of code had an upper limit of 3. The inner loop, however, had an upper limit of 32,768! If we could vectorize the inner loop concurrently—sending "chunks" of the inner-loop vector to each of the processors—the performance would improve. We saw in the Autotasking manual that a micro-tasking directive would let us do just that. By inserting a directive of the form *cmic\$ do parallel vector* we got the long inner loop partitioned across the eight processors, thus allowing both vectorization and full parallelization, which improved our granularity and produced a better speedup.

Because PUEBLO is a three-dimensional code, many of its DO loops probably have an upper limit of three, and a check of other subroutines that contribute heavily to the run time seemed like a good idea. Doing so revealed several more instances of the same problem. Adding directives to these subroutines gave us the results seen in Figure 6. On the 32 by 32 by 32 mesh, at the highest level of optimization that we achieved, ATExpert predicted a speedup of 6.4, out of a possible 7.2, or 90 percent of the Amdahl's Law prediction. To verify the ATExpert predictions, we ran the code on a dedicated CRAY Y-MP8/8128 system. We measured a speedup of 6.1 based on the sequential




Figure 4. Graphic output from ATExpert showing the predicted speedup of 2.5 percent for the initial Autotasking run. Also shown are measurements of the overheads associated with running the code.

possible 4.8 (assuming eight processors). Because a highly vectorizable code implies many loops that should be amenable to parallelization, this result was puzzling and disappointing and needed to be investigated further. Figure 4 suggests that the problem began when more than three processors were used. By inspecting additional information provided by ATExpert, we found that loop 20 in the most heavily used subroutine, RIEMAN, was




Figure 5. Code fragment from the subroutine RIEMAN showing a loop that contributes to the low performance figure.



code using our version of a routine called ISMIN and 5.9 with the SCILIB version of ISMIN. This is a difference of 5 and 8 percent, respectively. By the judicious use of the directives, both the predicted and measured parallelisms were improved. (Recall that the fraction of the code that can execute in parallel is obtained from detected parallelism not from potential parallelism.) The box at the bottom left of Figure 6 provides us with potential problems that may be inhibiting parallelism in the code. By making use of this additional information we may be able to continue to improve the Amdahl's Law prediction as well as the actual speedup.

Figure 7 shows the ATExpert results for the 64 by 64 mesh with a predicted speedup of 7.2 over a single processor. Note that both the Amdahl's Law prediction and the ATExpert prediction for speedups have improved. This is due to the fact that the vector lengths are now 65,536, and thus a larger percentage of the execution time is spent in the parallel parts of the code.

These longer vector lengths also enabled us to amortize more of the overhead associated with multitasking, resulting in a decrease in predicted overhead from 0.8 CPUs for the smaller problem to 0.6 CPUs for this one, a 25 percent improvement. During dedicated time on this problem, we measured an actual speedup of 6.7 on eight processors, a difference of 7 percent over that predicted by ATExpert. The difference probably stemmed from several causes. One probable cause is the variable amount of work done in some of the loops in PUEBLO, which affects the accuracy of the ATExpert predictions. Another is the magnitude of the effect of the memory contention in the code. The first probable cause could raise or lower ATExpert's predictions relative to the measured result; the second effect would increase the measured time and therefore lessen the speedup.

Conclusions

ATExpert is an excellent tool for parallelizing large codes. The fact that the PUEBLO code is nearly 100 percent vectorizable helps because the analysis needed to determine vectorizability is essentially the same as that needed to determine loop-based parallelism. ATExpert's analysis provides users with much more information than any previous tool. Furthermore, this information is presented in several easy-to-understand formats. We were able to take a large, sequential, but highly vectorizable code, and with a modest amount of effort, parallelize it to obtain predicted speedups of between 6.4 and 7.2 across eight processors (actual measurements



were from 6.1 to 6.7). The effort that would have been required without the CF77 compiling system with Autotasking and diagnostic tools such as ATExpert would have been orders of magnitude greater. ■

About the authors

Margaret L. Simmons is a staff member in the Computer Research and Applications Group at the Los Alamos National Laboratory, where she currently is involved in supercomputer performance evaluation. Her research interests include performance instrumentation and measurement and parallel algorithms and architectures. She received a B.S. degree in mathematics from the University of Houston and an M.S. degree in computer science from the University of New Mexico.

Kenneth M. Lord is project leader for the Americas Market Analysis and Forecasting Group at Cray Research. Previously, as a Cray Research analyst, he provided benchmarking support and optimization assistance for the Los Alamos National Laboratory. He received a B.S. degree in geology from Michigan State University.

Reference

1. Buzbee, B. L., "The Efficiency of Parallel Processing," *Computer Design*, June, 1984.

Figure 6 (top). Graphic output from ATExpert showing the predicted speedup of 6.4 percent for the optimized Autotasking run on the 32 x 32 problem. Also shown are measures of the overheads associated with running the code.

Figure 7 (bottom). Graphic output from ATExpert showing the predicted speedup of 7.2 percent for the optimized Autotasking run on the 64 x 64 problem. Also shown are the top subroutines and the percentage of time they take.

CORPORATE REGISTER

Cray Research receives first commercial order for CRAY Y-MP C90 system

Ford Motor Company ordered a CRAY Y-MP C90 system to be installed at the automaker's Engineering Computer Center in Dearborn, Michigan. The new system, which will replace a CRAY Y-MP supercomputer that Ford installed in 1989, will be used for structural analysis, crash simulation, computational fluid dynamics and other problems related to automotive design and engineering. Ford was the first U.S. automotive company to purchase a CRAY Y-MP system. The CRAY Y-MP C90 system is the follow-on to that system and continues Ford's commitment to using leading-edge computational solutions for computer-aided engineering.

Lawrence Livermore National Laboratory (LLNL) has signed an agreement to purchase a CRAY Y-MP C90 supercomputer for delivery to LLNL's National Energy Research Supercomputer Center (NERSC) in 1992. LLNL's order for the CRAY Y-MP C90 system follows the cancellation of a system order from Cray Computer Corporation of Colorado Springs, Colorado.

The Republic of China (Taiwan) **Central Weather Bureau (CWB)** has ordered a CRAY Y-MP 8I system and a CRAY Y-MP EL system. The CRAY Y-MP 8I system, an upgrade replacement for a Control Data Corporation CYBER 205 system, will be used by the Taipei-based weather bureau for daily weather forecasting. The CRAY Y-MP EL system, Cray Research's entry level supercomputer, will be used as a file server. The CRAY Y-MP 8I system also will operate

the EMPRESS relational database management system with a unique binary large object (BLOB) attribute. This attribute is ideally suited for storing data that has no easily defined structure, such as satellite data, analyzed fields, forecast fields, and satellite images. The use of EMPRESS on CWB's new system is the first instance of a Cray Research weather customer running a standard database application package on a Cray Research system.

The Institut Méditerranéen de Technologie (IMT), a leading teaching and research institution in Marseilles, France, ordered and installed a CRAY Y-MP 2E system for general academic and industrial research, primarily in the area of computational mechanics applications with emphasis on aeronautics, astronautics, and ocean engineering.

The **Max Planck Society** has installed a CRAY Y-MP 4E supercomputer at its Institute for Plasma Physics in Garching, Germany, a suburb of Munich. The institute is one of the world's premier high energy physics research laboratories. The CRAY Y-MP 4E system, the third Cray Research supercomputer installed at the institute, replaces a CRAY X-MP system. It will be used primarily for research in plasma physics, astrophysics, and solid state physics. The system also will be used by other Max Planck Institutes in Germany for a variety of research applications.

The **Australian Bureau of Meteorology**, Australia's weather agency, has ordered and installed a CRAY Y-MP 2E supercomputer at its computer center in Melbourne. The new system joins a CRAY X-MP system, which was the

bureau's first Cray Research system and was installed in July 1990. The CRAY Y-MP 2E system, which will be used primarily for operational weather forecasting for Australia, will be networked to the CRAY X-MP system. The Australian Bureau of Meteorology operates one of the major processing centers within the World Meteorological Organization's Weather Watch program and is the leading center of its kind located in the Southern Hemisphere.

The **Swiss Federal Institute of Technology (EPFL)** located in Ecublens, Switzerland, was the first customer worldwide to install a CRAY Y-MP EL system, Cray Research's entry level supercomputer. EPFL has a nationwide network of approximately 3000 computer systems including workstations, servers, mainframes, and Cray Research supercomputers. EPFL also ordered a CRAY Y-MP 4E system that will operate as a central compute server. The CRAY Y-MP 4E supercomputer will be networked via the 800-Mbit/s HIPPI interconnection to the CRAY Y-MP EL system, which will act as a high-speed file server for the central computers and all the departmental compute servers on the network.

Three other European organizations also have ordered CRAY Y-MP EL systems: **The Defence Research Agency, Maritime Division, Admiralty Research Establishment (ARE)** Haslar, a UK Ministry of Defense (MOD) site located at Gosport in Hampshire; **Lotus Engineering**, an automotive consultant in the UK that is part of Group Lotus, the manufacturer of high-performance sports cars; and **Ecole Centrale de Paris (ECP)**, an engineering school in Paris, France.

ARE researches all aspects of Naval Operations, with a special emphasis on advanced defense technology. The new CRAY Y-MP EL system will be used for computational fluid dynamics applications in ship and submarine hydrodynamics. Lotus Engineering will use its CRAY Y-MP EL system to support the facilities and expertise offered by its Design Analysis Group in Hethel, England. The ECP system will be used by the school's Mechanical Engineering Laboratory, which works to optimize codes and applications for the French automotive, geophysical, civil engineering, and other industries.

Sun Microsystems and Cray Research to cooperate on development of hardware and software environment

Sun Microsystems, Inc., and Cray Research announced in January that the firms will cooperate in the creation of a hardware and software environment that will allow Cray Research systems and Sun workstations to work together more easily and efficiently.

"There is a very fast growing installed base for SPARC systems, with more than 3800 third-party hardware and software products," said John A. Rollwagen, chairman and CEO of Cray Research.

"We're committed to the concept of network supercomputing so that individual users can easily access the power of our supercomputers." Cray Research has joined SPARC International, an organization of more than 200 SPARC vendors that supports binary compatibility among products through the SPARC Compliance Definition (SCD).

"Cray Research has the ability to build the ultimate accelerator for SPARC applications," said Scott McNealy, Sun chairman and CEO, "and by choosing SPARC, Cray Research can take advantage of the benefits of the SPARC community. SPARC users have access to a wide range of compatible systems at many price points from laptops to supercomputers."

Cray Research and Digital Equipment Corporation reach worldwide marketing agreement

Cray Research and Digital Equipment Corporation announced in January a comprehensive agreement that grants Digital worldwide rights to market, sell, and distribute the CRAY Y-MP EL supercomputer system. The CRAY Y-MP EL

air-cooled, entry level supercomputer is fully compatible with Cray Research's entire CRAY Y-MP supercomputer product line and integrates easily into Digital's open distributed computing environment. Under terms of the agreement, Digital will sell the system in conjunction with Digital's own high-performance technical computing (HPTC) products: the VAX vector systems and the DECmpp line of massively parallel systems.

The agreement provides for an orderly marketing and sales transition. For the first six months of the agreement — through June 30, 1992 — both Digital and Cray Research will market and sell CRAY Y-MP EL systems. The agreement does not apply in certain limited cases, such as CRAY Y-MP EL systems sold in Japan and to certain government customers. Starting July 1, 1992, with those limited exceptions, customers will purchase CRAY Y-MP EL systems exclusively through Digital sales representatives. Cray Research will offer back-up sales support and applications expertise for the CRAY Y-MP EL system. Officials of both companies said they consider 1992 a building year for this relationship and that they expect the benefits of the agreement, including sales volume increases, to become more evident in 1993 and succeeding years.

Cray Research completes acquisition of selected FPS assets, creates wholly owned subsidiary

In December, Cray Research completed its acquisition of selected assets of Floating Point Systems (FPS), producers of a minisupercomputer based on SPARC microprocessor technology. Cray Research has formed a wholly owned subsidiary, Cray Research Superservers, Inc., to market and support current products and to explore technologies for creating software compatibility between the company's supercomputers and SPARC-based systems. Cray Research acquired rights to FPS technologies, system and parts inventories, and equipment to support manufacturing and development efforts. Selected customer and vendor contracts also were assumed as part of the transaction.

Cray Research announces first graduate fellowships in computational chemistry

In December, Cray Research, in collaboration with the American

Chemical Society's (ACS) Division of Physical Chemistry, awarded three graduate fellowships in computational chemistry for the 1991-92 academic year. The Cray Research Fellowship in Computational Chemistry was established to encourage doctoral degree studies in computational chemistry; the fellowships are to be awarded annually on a competitive basis to graduate students in the research and dissertation stage of their doctoral programs.

The award focuses on innovative developments and applications of methods in theoretical chemistry or related disciplines, with emphasis on large-scale numerical simulations on supercomputer systems. Recipients of the 1991-92 Cray Research Fellowships in computational chemistry are

- Zhihong Deng, Department of Chemistry, University of Pennsylvania
- Scott M. Le Grand, Department of Chemistry, Pennsylvania State University
- Randall Q. Snurr, Department of Chemical Engineering, University of California, Berkeley

"These fellowships allow Cray Research to recognize and support some of the finest emerging talent in computational chemistry," said Robert A. Eades, manager of Cray Research's chemical research and engineering group. "We are delighted to be helping these talented graduate students continue their important and innovative research work."

Funded by the Cray Research Foundation, the fellowships provide a cash stipend of \$5000 for each student per school year as supplementary financial support. In addition, the students were recognized personally by representatives of Cray Research and the ACS Subdivision of Theoretical Chemistry in April at the 1992 ACS Spring National Meeting in San Francisco.

Applicants interested in the 1992-93 Cray Research Fellowship in Computational Chemistry should submit applications no later than June 30, to Dr. John Tully, AT&T Bell Laboratories, Room 1D346, 600 Mountain Avenue, Murray Hill, NJ 07974. Applications should not exceed five typewritten pages and must include a definition of the proposed research project and an explanation of its scientific importance, the proposed use of the grant money, a description of the importance of high performance scientific computing to the proposed research, and two letters of recommendation.

APPLICATIONS UPDATE

PRO/II Version 3.02 available on Cray Research systems

Simulation Sciences, Inc., has released Version 3.02 of PRO/II, its fourth-generation process simulator, for use on Cray Research systems running the UNICOS operating system. PRO/II performs complete heat and material balances for processing plants in the oil and gas, refining, chemical, and petrochemical industries. Program applications include plant process design, troubleshooting, trend analysis, equipment monitoring, and offline and online optimization.

Key to the program's enhanced simulation capabilities is CHEMDIST, a non-ideal distillation algorithm that includes a non-ideal initial estimate generator. The algorithm is well-suited for the solution of tough three-phase distillation systems.


PRO/II also contains a suite of solids-handling unit operations including crystallizers and dissolvers. These unit operations are coupled with PRO/II's state-of-the-art stream structure for storing attributes and substream data. Together these permit the simulation of industrial crystallization processes which up to now were outside the capabilities of commercial process simulators.

PRO/II Version 3.02 also includes the following enhancements:

- ☐ A proprietary simulation database that allows easy data transfer to other programs
- ☐ Entry of engineering, thermodynamic, and property data in all standard dimensional units
- ☐ Output reporting options including graphical analyses of simulation results
- ☐ A library of physical property data, mixture data, and thermodynamic models to simulate a wide range of systems

PRO/II simulates processes of unlimited size and complexity, without the need to compile and relink. It has no limits on the number of components, column trays, unit operations, assay streams, reactions, and recycle streams.

The program's open architecture reduces development time to implement new technology, provides transparent links to the user's own programs and third party software, and improves program reliability and robustness.



Relief contours show the electric field propagation of a dipole antenna array computed with MSC/EMAS and rendered with the MSC/XL pre- and postprocessor.

For more information on using PRO/II with Cray Research computer systems, contact Katherine M. Thurston, marketing assistant, Simulation Sciences, Inc., 601 S. Valencia Avenue, Suite 250, Brea, CA 92521; telephone: 800/827-7999, in California, 714/579-0412; or Richard LaRoche, Cray Research, Inc., 655E Lone Oak Drive, Eagan, MN 55121; telephone: 612/683-3696.

The MacNeal-Schwendler Corporation announces Version 2 of MSC/EMAS

The MSC/EMAS program from The MacNeal-Schwendler Corporation (MSC) is a finite element analysis program for modeling electric and magnetic fields and is available for use on Cray Research computer systems running the UNICOS operating system. MSC/EMAS is a general purpose, two- and three-dimensional program for simulating static, frequency-dependent, transient and resonant fields. Typical applications include rotating and linear machines, switch-gear and transformers, pulsed power apparatus, eddy current and magnetic sensors, and microwave components.

Version 2 of MSC/EMAS, released in November, offers users unlimited problem size, new solver technology to reduce solution time and storage requirements, and availability of complex materials. Absorptive boundary conditions allow more accurate models to be built with less detail and allow effective nearfield analysis of antennas and radiating structures. The program also includes user interface improvements that allow new users to become productive in less time.

MSC/EMAS is suited particularly for two- and three-dimensional linear and nonlinear eddy currents and microwave cavity applications. The program solves the complete and general forms of Maxwell's equations and has been tuned to address the needs of electrical engineers. It currently is being used in nondestructive evaluation to analyze the electromagnetic signature of down-hole sensors in the oil industry. Other applications include heating optimization analysis in consumer microwave ovens being conducted by food-packaging researchers and consumer-oven industries and the analysis of losses generated by turbine generators being conducted by suppliers to the electrical utility industry.

The new release includes the option of using absorbing boundary conditions for static and dynamic analyses, which lets users simulate free space without having to terminate the field at the model boundary of radiation applications.

The technology used in MSC/EMAS is an extension of that used in the industry-standard MSC/NASTRAN structural analysis package. The program is tightly coupled with MSC/XL, the company's graphical pre- and postprocessor. Results from MSC/EMAS can be used for subsequent analysis with MSC's other products for stress, vibration, thermal, and acoustic analysis.

For more information on using MSC/EMAS with Cray Research systems, contact Perry Grant, The MacNeal-Schwendler Corporation, 815 Colorado Blvd., Los Angeles, CA 90041; telephone: 213/258-9111; fax: 213/259-3838; or Mike Aamodt, Cray Research, Inc., 655E Lone Oak Drive, Eagan, MN 55121; telephone 612/683-3665.


Protein structure simulation advances understanding of HIV virus

Using molecular dynamics (MD) simulations, S. Swaminathan and David L. Beveridge of Wesleyan University, Middletown, Connecticut, and William E. Harte, Jr., of Bristol Myers Squibb in Wallingford, Connecticut, have developed a dynamical model for the structure of the human immunodeficiency virus 1 (HIV-1) protease dimer in aqueous solution. The dynamical structure also suggests information that ultimately may be useful in understanding and advancing the development of specific inhibitors of HIV-1 protease.

The Wesleyan/Bristol Myers Squibb team used the CRAY Y-MP system at the Pittsburgh Supercomputer Center to study HIV-1 protease dimer, an enzyme essential for the replication of the AIDS virus. This enzyme, a significant macromolecular component of the AIDS virus, and one for which detailed structural information is available, is a target of considerable pharmaceutical interest in the quest for AIDS therapies.

The MD simulation was based on the crystal structure of the enzyme reported in 1989 by Alex Wlodawer and coworkers at the National Institutes of Health (NIH). The structural model for HIV-1 protease developed from crystallographic studies established the essential secondary and tertiary structures of the enzyme.

MD calculations were performed with the Monte Carlo (MC) and MD simulation program WESDYN, a program that was optimized to take advantage of the vectorization capabilities of the CRAY Y-MP system. The simulation, one of the largest to date reported in the field of protein dynamics, included 6990 water molecules of solvation along with the 198 amino acid residues of the protein and required over 100 hours on the CRAY Y-MP system. Inclusion of solvent was shown to be essential for obtaining accurate results. "The capability of doing calculations of this scale was not available until the advent of supercomputers," said Beveridge.



Dynamical structure of HIV-1 protease as modeled on the CRAY Y-MP system at the Pittsburgh Supercomputing Center.

The simulation demonstrated dynamical motions of the "flap" region adjacent to the active site, which are likely to be crucial in the enzyme mechanism and inhibitor complexes. The simulation predicts that motions in two distinct regions of the HIV-1 protease are correlated with each other although the regions are separated by a considerable distance in space. One domain, named the "cantilever" by researchers, acts through a "fulcrum" to cause motion in a "flap" region near the enzyme's active site. Domains such as the cantilever that affect the flap also could play a part in enzyme function and therefore could be targets for enzyme inhibitors—and potentially of use in developing AIDS therapies. "These changes emphasize the need for a consideration of structural dynamics in understanding HIV-1 protease action," said Beveridge.

However, says Beveridge, his work is not so much a breakthrough in AIDS research as a new way of looking at the dynamical structure of proteins.

"One would never know about this," he explained, "and would have no means to investigate it without large-scale molecular dynamics simulations. No experiment is going to give you all this information on detailed atomic motions. Here, computational theory has a unique vantage point in elucidating the details of molecular motion."

In recent follow-up work on this project, Harte's Wesleyan doctoral dissertation project, simulations have been performed on a mutant form of the HIV-1 protease in which a glutamine residue replaces the naturally occurring asparagine at position 88 in the sequence. This form of the protease is known to be inactive due to dimer dissociation, even though the change is somewhat removed from both the dimer interface and the active site. The simulations demonstrated that the protein dimer is destabilized in the mutant form due to an alteration in internal hydrogen bonding that causes the monomeric units to be forced out of articulation. A series of additional simulations compared with the crystal structure data have been used to determine the protonation state of the active site residues in an HIV-1 protease inhibitor complex. The researchers expect future MD simulation projects to be highly useful in providing an interpretation of observed behavior of the enzyme, as well as uncovering aspects not directly accessible to experimental measurement.

Scientific visualization provides insight at U.S. Army Ballistic Research Laboratory


The U.S. Army Ballistic Research Laboratory (BRL), at Aberdeen Proving Ground, Maryland, serves as the Army's primary laboratory for research in ballistics, vulnerability, and vulnerability reduction. Scientific visualization at BRL, now a cornerstone of the laboratory's strategic effort, began in the early 1980s with BRL-CAD, an internally developed program that features solid-modeling applications support, ray tracing, image

analysis, and image manipulation. Today, BRL researchers and engineers have several visualization software packages available, including Cray Research's Multipurpose Graphic System (MPGS), Wavefront Technologies' Data Visualizer and Advanced Visualizer, and BRL's in-house packages, BRL-CAD and ShAYD.


A recently formed Scientific Visualization Team helps researchers choose the packages that are most appropriate for their applications. The combination of scientific visualization and high-performance computing carried out on the lab's CRAY-2 and CRAY X-MP/48 computer systems has contributed to a number of significant accomplishments that directly serve BRL's mission. Following are some examples.

Using a computational fluid dynamics model, researchers Jubaraj Sahu, Charles J. Nietubicz, and Karen Heavey of BRL's Launch and Flight Division analyzed changes in the flight characteristics of a projectile after its base had been modified. The changes had become apparent during range firing tests. The researchers computed airflow around the projectile at 370,000 grid points in three dimensions and collected data sets for several velocities. Using the MPGS package to study this large body of data interactively, the researchers generated a set of surfaces to represent areas where the flow expands from subsonic to supersonic speeds and then shocks back down to subsonic speeds. The most significant result observed was the formation of a small region of supersonic flow directly behind the base area. This base flow region affects the drag and lift characteristics of the projectile. Until these three-dimensional images were created, the existence and structure of the small supersonic region was not readily apparent.


Doug Savick, also of BRL's Launch and Flight Division, routinely uses three-dimensional visualization as an initial analysis tool for evaluating potential designs of projectile components. A class of projectiles known as long rod penetrators typically are much smaller in diameter than the gun tube from which they are fired. A thrust-transmitting carrier, known as a sabot, surrounds each projectile to position it correctly and seal it in the tube during launch. After exiting the gun tube, aerodynamic forces cause the sabot to break away, and the projectile then proceeds unhindered. However, improperly designed sabots can affect the trajectory of the projectile by striking it at separation or even by remaining attached.




Model used at the U.S. Army's Ballistic Research Laboratory (BRL) to study airflow around a projectile. The study revealed a region of supersonic flow (in red) directly behind the redesigned base of the projectile.



Model used at BRL to study the separation of a projectile carrier, a sabot, from the projectile.



Pressure levels generated by a liquid propellant and distributed throughout the combustion chamber of an experimental gun. Computational studies such as this can reveal information unobtainable by experiment alone.



Long rod penetrator impacting plate. Color indicates pressure on plate surfaces.

Savick models sabot discard and uses Advanced Visualizer to convert the model output to a sequence of images. Viewing the images in series as an animation lets Savick analyze many new designs quickly for potential disturbances between the sabot components and the projectile.

Researcher Terry Coffee, of BRL's Interior Ballistics Division, is researching the characteristics of liquid propellants in an experimental gun. The primary phenomena of concern are the pressure waves generated as the propellant enters a combustion chamber, is converted into a gas, and drives a projectile down the gun tube. Experimentation can measure only the pressure on the walls of the

chamber, but the pressures acting on the projectile are critically important.

To improve the understanding of this combustion process, Coffee developed a computer model that comprised nearly 35,000 cells and ran in just over 6 CPU hours on the CRAY X-MP system. The simulation results were made into a series of three-dimensional visualizations using BRL/ShAYD, which plotted the pressure levels throughout the chamber as a surface. Pictures were made at every 0.002 ms over a 6 ms interval, resulting in a total of 3000 frames. Each frame can be viewed individually or in series as an animation. This visualization not only provided insights about

the propellant's properties in the chamber but also suggested improvements to the model itself.

Researcher Steven Segletes of BRL's Terminal Ballistics Division serves on a team that uses several simulation codes, known as hydrocodes, to model ballistic impacts and interactions between explosives and metals. Visualizations produced with the MPGS package revealed a suspected instability of an equation of state that is used widely under certain conditions and was causing an adverse effect on the simulations.


The ongoing revolution in scientific visualization is leading to improved productivity and effectiveness at all levels of research and engineering. Visualization also provides a convenient communication channel to help explain conclusions to colleagues and peers as well as management. "The cost effectiveness of visualization is clear, and BRL will aggressively support the efforts of its researchers to explore the applications of this rapidly developing technology," said Richard C. Angelini, a member of the Scientific Visualization Team at BRL. "The use of these techniques is expanding, and through the efforts of the 'Vis' team we are finding new, and sometimes unexpected, ways to use them to solve our research and engineering problems."

Modeling provides a new notion for textile and clothing manufacturers

When apparel and textile manufacturers create clothing designs in the future, they may well pattern their methods after those of researcher Jeff Eischen of North Carolina State University. Eischen is using a CRAY Y-MP system at the North Carolina Supercomputing Center to determine how limp materials drape over rigid objects.

His work through NCSU's College of Textiles is sponsored by the federally funded National Textile Center. The Fabric Mechanics Analysis project has two potential areas of application: computer-aided clothing design and textile manufacturing automation.

The present challenge is to understand the physics of motion as it applies to fabrics. Using finite element analysis software, Eischen divides the surface of a material, which is draped over a block, into a series of grid points. Next, by using the Wavefront animation program, material attributes such as color and texture are assigned, and the fabric is "put into motion" by NCSC animation



Advanced computer rendering techniques are used to visualize the dynamics of fabric at the North Carolina Supercomputing Center.

specialist Chris Landreth. The customized code reads in data sets and renders frames that are recorded to video for playback as an animated sequence.

"There's a real push toward automation by textile and apparel manufacturers," Eischen said. "The manufacturers would like to have robots on the factory floor that pick up, align, fold, and sew fabric parts together automatically." More importantly, due to quick-change requirements, the automations would have to distinguish among many types of fabrics. "What's needed are computational models to set up the robots," Eischen said. He envisions a day when body measurements will be taken electronically, and dressing rooms will be replaced by three-dimensional imaging systems. Supercomputer modeling is providing a first step, by showing manufacturers how clothes will drape on the human form.

Greenhouse Earth exhibit debuts

In January 1992, the Franklin Institute Science Museum in Philadelphia, in collaboration with the Association of Science-Technology Centers, Washington, D.C., opened a major exhibition on global climate change. The 4000-square-foot exhibition features interactive computers, hands-on exhibits, video, and live demonstrations. Over a three year period, Greenhouse Earth will visit

museums in ten major cities throughout North America, reaching more than two million people.

The exhibition organizers hope to increase visitors' understanding of current scientific knowledge of global climate change and of forecasts for the coming century.

The exhibition received funding support from Cray Research, the National Science Foundation, and the U.S. Department of Energy. "Global climate change is an issue of extreme importance and one in which high-performance computing is making significant contributions to our understanding," said David Blaskovich, Cray Research's director of environmental systems marketing. "Educating people and heightening public awareness is crucial, particularly when you consider that we've really only got one chance to test the hypothesis of climate change."

Through exhibit devices that include interactive computer and live demonstrations, visitors will come to better understand major areas of uncertainty in global climate change research and the ways in which mathematical models are helping resolve those uncertainties. An interactive computer demonstration allows visitors to explore segments of the supercomputer climate change model from the National Center for Atmospheric Research and explains the uses of modeling as a tool in climate change research. Cray Research also supported produc-

tion of a video included in the exhibition that features interviews with experts and discussions of the roles clouds, oceans, and biological feedback mechanisms play in regulating the physical environment.

The Greenhouse Earth exhibition will reside at the Boston Museum of Science through the summer before moving on to the SciPort Museum in Shreveport, Louisiana in the fall. For more information, please contact ASTC, 1025 Vermont Avenue, NW, Suite 500, Washington, D.C. 20005; 202/783-7200.

Learjet engine pylon design

Due to the rocketing costs of flight tests, the emphasis at Learjet is to design prototype aircraft as close to optimum as possible, requiring sophisticated computer analysis. To that end, Learjet's corporate parent, Bombardier, Inc., has made an enormous investment in computer technology. In addition to the resources already invested in high technology leadership at Learjet, Bombardier backed the construction of a new Flight Test and Development Center adjacent to Learjet's engineering facilities. The center is equipped with the latest state-of-the-art ground station for support and is staffed by 150 engineers and service technicians. The center will be utilized by all three of Bombardier's aerospace companies: Canadair, Shorts, and Learjet.

Learjet visionaries are convinced that the buildup of their four analysis systems, AERO, CAD/CAM, DAS and NASTRAN, will pay off as the confidence level in these tools and systems significantly reduces the need for developmental, structural, and flight testing to optimize the design of new aircraft.

In large measure, the validity of analysis tools has been proven already at Learjet. For example, it took less than three weeks for NASA Ames and Learjet aerodynamicists to design an advanced engine pylon for the Learjet Model 60. The search for a shock-free, minimum-drag pylon required using the most advanced computational methods available, according to Mike Hinson, Learjet chief aerodynamicist for the project.

"The project would not have been possible without the CRAY Y-MP system," said Hinson. "Before utilizing the Cray Research supercomputer we could only look at pressure distribution for pylons in subcritical conditions. Experience and gut feel were not


enough to predict the drag. It was critical that the shock waves be considered. With the Cray Research system, we were able to look at the real problem and get the right answer."

The computing capabilities available at the NASA Ames research facility in California focus on the NASA CRAY Y-MP supercomputer, where two hundred billion calculations can be performed in less than two hours to arrive at one Tranair solution for a complete aerodynamic configuration.

"It is critical that the Tranair code runs on a machine like the Cray Research supercomputer for engine pylon design," said Hinson, who explained that there were 500,000 points in the computational grid that defined the aerodynamics of the Model 60. Each of these half-million points need to be evaluated individually and iterated hundreds of times. The CRAY Y-MP system repeatedly calculated the differences in air flow qualities at each point as the project team continually changed the shape of the pylon in pursuit of a minimum drag, shock-free profile.

The computational analysis of the Learjet Model 60 engine pylon design required about 40 hours on the CRAY Y-MP system. Project team members obtained an optimum design configuration, resulting in a significant reduction in drag and an increase in range translatable to 50 nautical miles for the Learjet Model 60.


Learjet aerodynamicists plan some flight testing of the design. Surface pressures will be measured in flight on the prototype Model 60, according to Hinson, who is confident that the computerized predictions will match the test flight findings. "Since the success we have already experienced with the pylon was achieved with the active assistance of top NASA experts and with the use of the world's most advanced computational system, we believe that what we have is what we'll get. We don't expect any surprises."



(Above). Cross-section of air flows in region of baseline and improved pylons as computed using NASA's Tranair code. The improved pylon used on the Learjet Model 60 shows substantially lower local Mach numbers.

(Below). Aerodynamic shock waves across the wings and fuselage as modeled on Learjet Model 60. Yellow represents barely supersonic and red is Mach 1.3.





The surface of constant amplitude (96 percent) from one of a degenerate pair of molecular orbitals in a gold complex [the $Au_{45}(SH)_6$ system]. Performed on a CRAY-2 system at the National Center for Supercomputing Applications (NCSA), such calculations demonstrate sulfur-sulfur bonding between adsorbates when arranged in the "knight move" array on the Au(100) surface. Image courtesy of Harrell Sellers, NCSA.