CRAY RESEARCH, INC.
Marketing Communications
1333 Northland Drive
Mendota Heights, MN 55120
612/681-3438

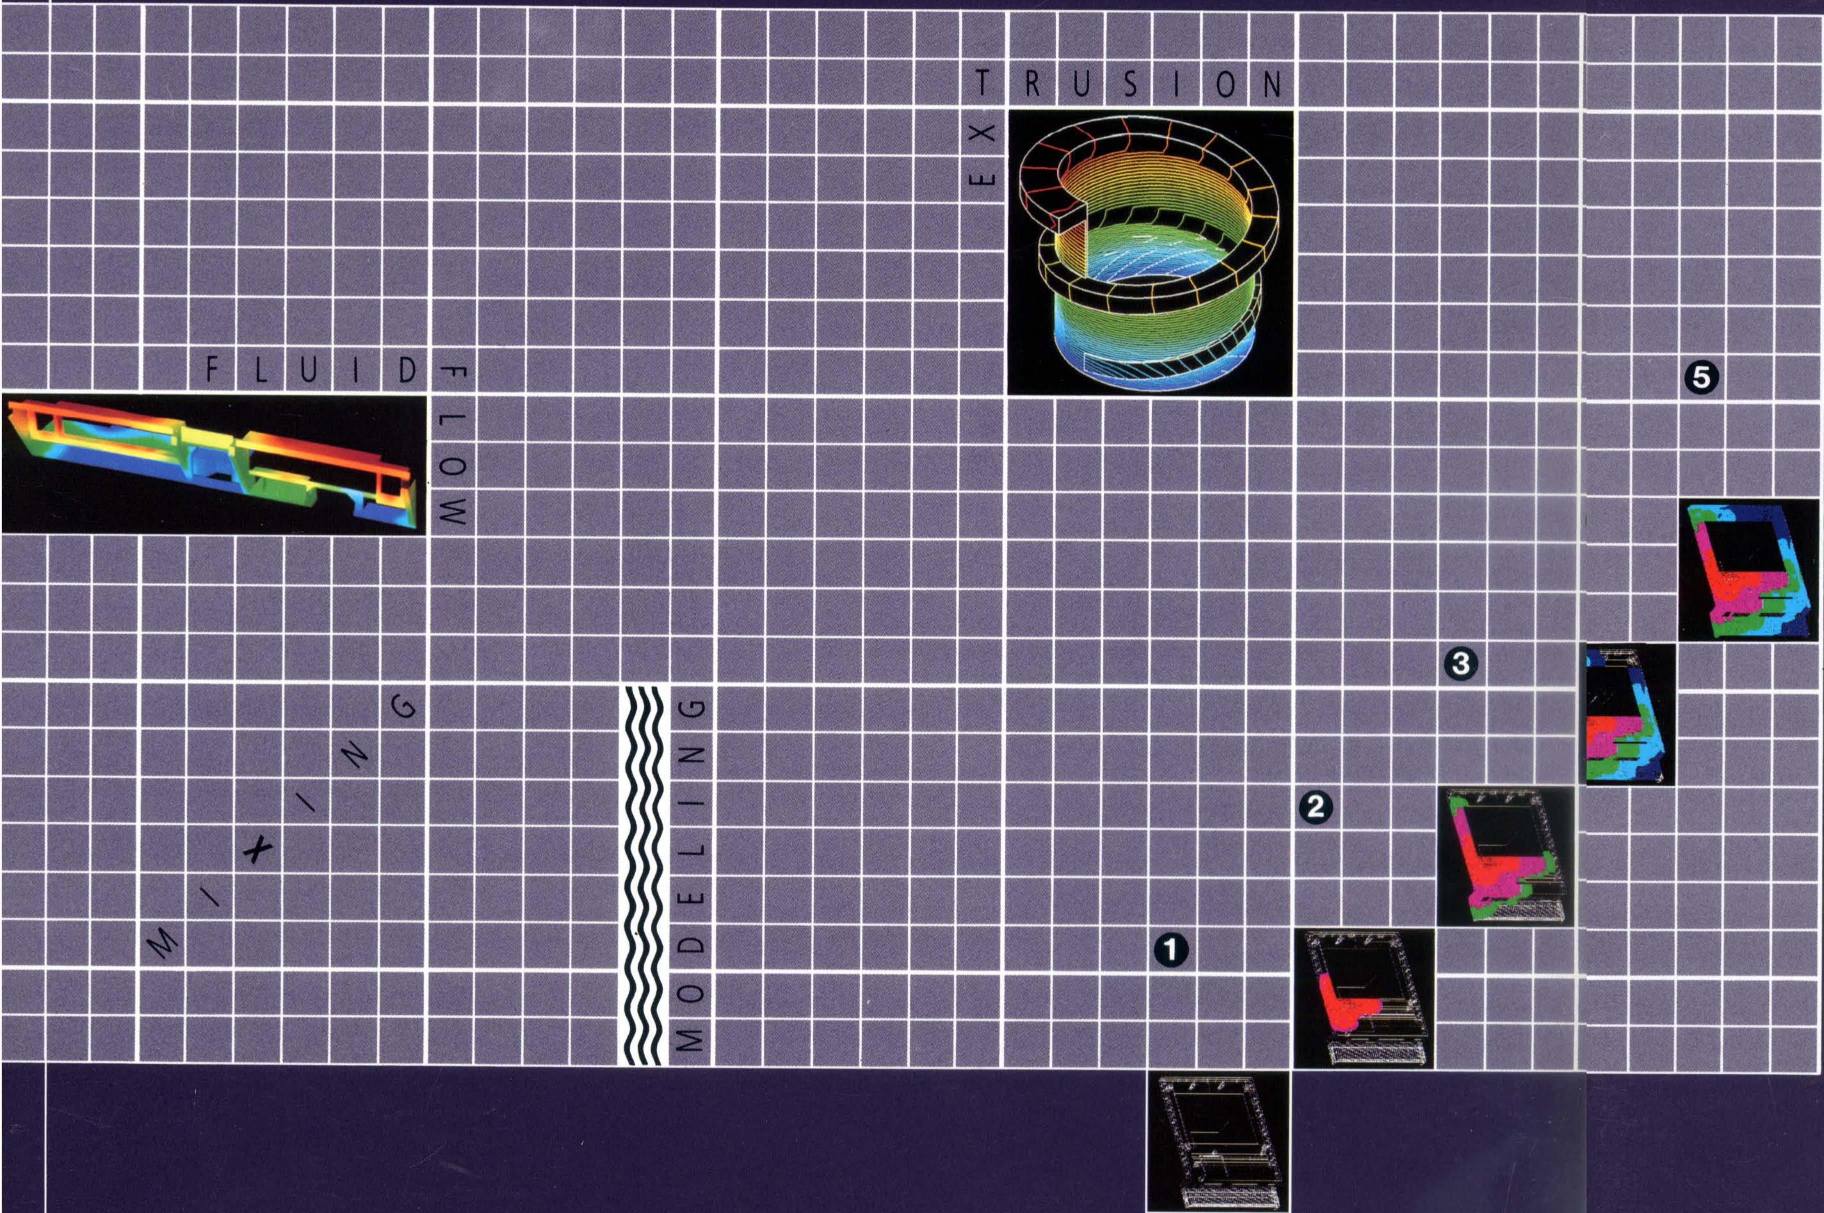




# CRAY CHANNELS

FALL 1989 · A CRAY RESEARCH, INC., PUBLICATION

## MATERIALS PROCESSING

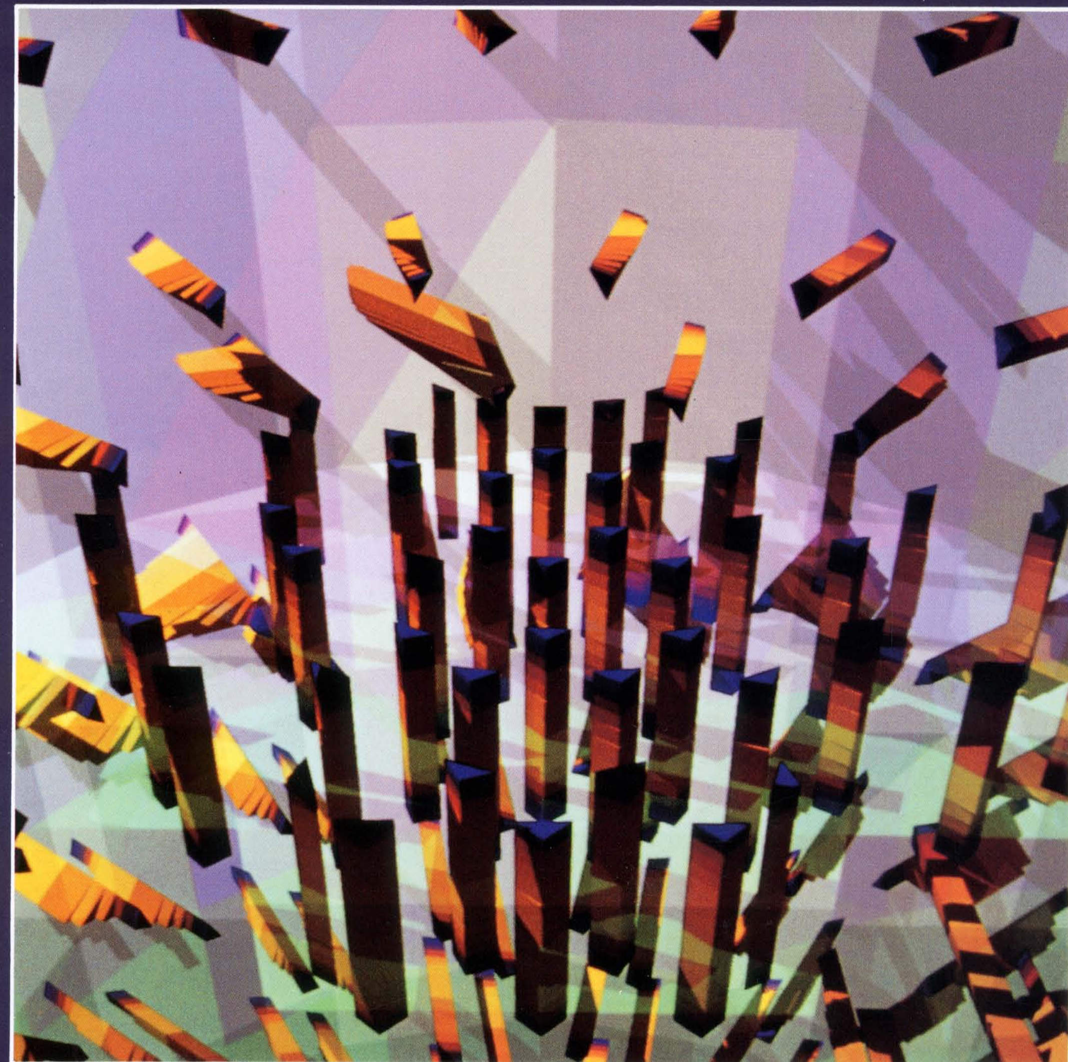

INJECTION MOLDING

CRAY CHANNELS

## In this issue

Optimizing production tools and processes is one way that manufacturing companies can conserve their resources. But optimizing the production of some types of products, such as plastic film spools or computer cabinets, can require a costly and time-consuming process of prototyping and refinement. Companies such as Eastman Kodak, Apple Computer, and others are optimizing production tools and processes in a more cost-effective way, through the use of supercomputers. These companies are using Cray Research computer systems to model the behavior of molten materials used in molding and casting. In this way, researchers are able to optimize mold configurations and injection and material handling conditions to maximize productivity. This issue of CRAY CHANNELS features several articles that describe these materials processing applications of Cray systems.

This issue also looks at the use of a Cray system to help understand the spread of a fatal fire at London's King's Cross underground railway station. In addition, our regular departments include stories on computer chess champions, astrophysical research, and the Migration Tools software package, which supports migrations from Cray Research's COS to UNICOS operating systems.

Optimizing the production of various products requires an understanding of the ways in which materials behave during various processing steps. The ability to solve materials processing problems quickly and precisely makes Cray computer systems cost-effective tools for production optimization. Cray computer systems provide unsurpassed modeling capabilities to improve efficiency in research, design, engineering, and production.

## Features

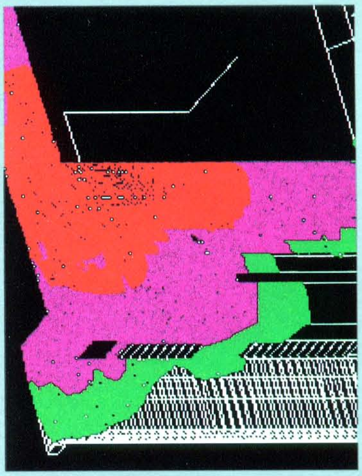

6



## Departments



CRAY CHANNELS is a quarterly publication of Cray Research, Inc., intended for users of Cray 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., 1333 Northland Drive, Mendota Heights, Minnesota 55120.

Volume 11, Number 3

**Editorial staff**
Ken Jopp, editor
Elizabeth Knoll, associate editor

**Design and production**
Barbara Cahlander
Eric Hanson
James Morgan
Cynthia Rykken

**On the cover:** Modeling materials processing operations on Cray computer systems helps manufacturers apply their production resources more efficiently. Modeling provides a fast, economical, and precise way to design production tools and improve manufacturing processes. Injection molding images courtesy of Eastman Kodak Company and Apple Computer, Inc.

# Injection molding

## Supercomputing and supergraphics

*Richard N. Ellson, Eastman Kodak Company*
*Rochester, New York*
*T. Marc Olano, University of Illinois at Urbana-Champaign*

High-volume manufacturing of complex, thin-walled structures often relies on the injection molding process. This technology allows functional items to be formed in a single, economical step. Familiar injection-molded items include slide carousels, cases for video or audio tapes, ice-cube trays, and telephone handsets. Successfully manufacturing parts by injection molding requires an understanding of the motion of molten plastic with high pressure gradients and knowledge of the ways in which the shape of a particular part influences its flow field.

By simulating plastic flows computationally, designers can develop prototypes of mold configurations and recommend processing conditions without cutting molds into steel. Simulation lowers the time and cost of converting a mold concept into a mold capable of producing quality parts — provided designers have access to sufficient computational resources for performing simulations and access to tools for analyzing simulation results.

The Eastman Kodak Company meets these needs with supercomputers and "supergraphics." The supercomputers, the CRAY X-MP/48 and CRAY-2 systems at the National Center for Supercomputing Applications (NCSA) of the University of Illinois at Urbana-Champaign, perform the computational fluid dynamics with a version of the Cornell Injection Molding Program modified by Kodak. The supergraphics, which visually communicate the results of supercomputer computations, are performed on an AT&T PIXEL machine connected to a Sun Workstation. The graphics software, the Kodak Glyph Visualization System for Injection Molding, interacts with the AT&T PIXEL machine through the PIClib and RAYlib graphics libraries. Supergraphics enables interactive, three-dimensional, visual, simultaneous access to temperature, velocity, and pressure data via a data compositing technique.[1]

Figure 1. In each of the circles area correlates with time, showing the impact of the progression of tools on the productivity of mold designers.

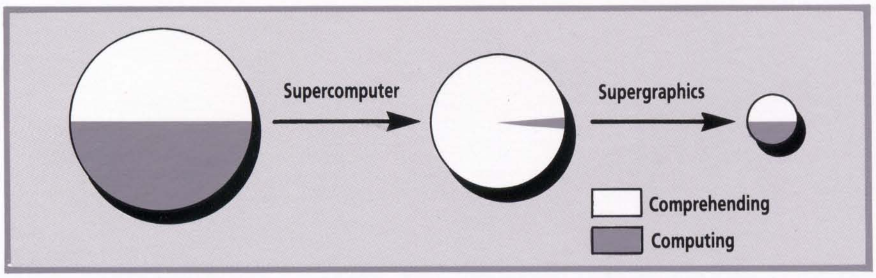


### Injection molding

The injection molding cycle begins when a steel mold cavity is filled with molten plastic. Once the mold is filled, the plastic cools and solidifies. When the outer skin of the newly formed part becomes solid, the mold opens and ejects the part. The cycle then is repeated. Simulations to date have focused on the filling phase because it has the greatest impact on the final characteristics of a molded part and serves as the initial condition for the later processes of packing and cooling. Hence, only the filling phase is discussed here.

During the filling phase, pressures of up to 1000 atmospheres push the viscous molten plastic through narrow gaps between the steel walls of a water-cooled mold. The viscosity of the plastic resists the pressure-driven flow, causing some of the kinetic energy of the plastic to convert to thermal energy. This heat conducts to the mold walls to be carried off by the cooling water. Heat production influences the flow of the material as the viscosity of the molten plastic reduces with increasing temperature. Thus, the plastic has a higher "fluidity" when heated. In addition, fluidity depends on the velocity gradients in the flow. These gradients are indicative of the shear between liquid layers, which untangles the long polymer chains of the plastic, raising the fluidity. This makes predictions of flow behavior difficult and sensitive to initial conditions.

The formation of a complete part, a basic goal of any molding process, requires the plastic to enter all portions of the mold. Incomplete parts arise when excessive local cooling causes some of the plastic to solidify prematurely and block flow to an empty region of the mold. Undercooled parts take a long time to solidify (lowering the productivity of the molding machine), and if the plastic reaches the thermal degradation temperature, the long-chain plastic molecules will break into pieces and weaken the part. Balancing the heating and cooling is just one concern of mold designers who attempt to produce high-quality, yet inexpensive, parts, while meeting specifications for strength, size, tolerances, appearance, and weight. Knowledge of the flow behavior during molding plays a central role in design decisions.

### Supercomputing

The mathematical model of mold filling that we use accounts for mass, momentum, and energy with the characteristics of the flowing plastic introduced through the constitutive relation for viscosity. Generalizations drawn from typical cavity geometries, material viscosities, and physical properties allow some simplifying assumptions to be made; the formulation of the model and its computer implementation with a hybrid finite-element and finite-difference method originate with the research of the Cornell Injection Molding Program and the FLOW3D simulation code.[2]

The computation of the flow behavior thus consists of solving a transient heat transfer and fluid dynamics problem for a highly viscous, temperature-sensitive, shear-thinning fluid. Like most computational fluid dynamics codes, the injection molding simulation vectorizes well. As a result, the move to a supercomputing environment reduced computation times by two orders of magnitude. Because the mold designer's typical problem can be solved within a few minutes on a Cray system, designers can complete many runs per workday. Prior to supercomputers, overnight batch runs were the norm.

Before supercomputing resources were implemented in the design cycle, approximately equal time had been spent on computing as on analyzing the results. The move to a supercomputing environment nearly doubled the speed of the design cycle by eliminating the computing time and leaving only the analysis time. In a situation analogous to vectorization, in which the speed of the entire program becomes dominated by the code run in scalar mode, the speed of the design cycle now depends primarily on the time required to analyze the data and not on the time to produce it. The supercomputer, by drastically reducing simulation time, sets the stage for similar improvements in the speed of the entire design cycle (Figure 1). These improvements are realized through supergraphics.

## Supergraphics

With supergraphics, we strive to present the viewer with an understanding of the computational results in the shortest possible time. To earn the prefix "super," graphics must keep pace with supercomputer output. This goal motivated our selection of the PIXEL machine because it provides 820 MFLOPS of graphics power, complementing the Cray system's high performance.

Given sufficiently fast graphics hardware, the next endeavor is to determine the important aspects of the simulation and to compose them into a visual image. Ideas for compositing can be found in the data representation techniques of print graphics such as graphs and charts.[3,4] Although print graphics remain effective in communicating the results of many simulations, for dynamical systems with many variables, computer graphics surpass print graphics with increased control of color and motion. By modifying print graphics for computers, the complex, global, dynamic behavior in simulations can be visualized with an artful use of symbols, color, motion, and geometry.

In the filling phase of the injection molding process, the key variables to visualize are pressure, temperature, and velocity, with one value of pressure for each finite element, and a velocity vector and temperature for every finite difference of each element. The glyph, an entity that combines abstract information into a graphical symbol, derives its shape from the velocity field of the fluid flow in the domain of a finite element at one snapshot in time. The length and orientation of the glyph shape indicate the velocity magnitude and direction of flow. The glyph geometry represents a velocity profile, a bar chart of the velocity

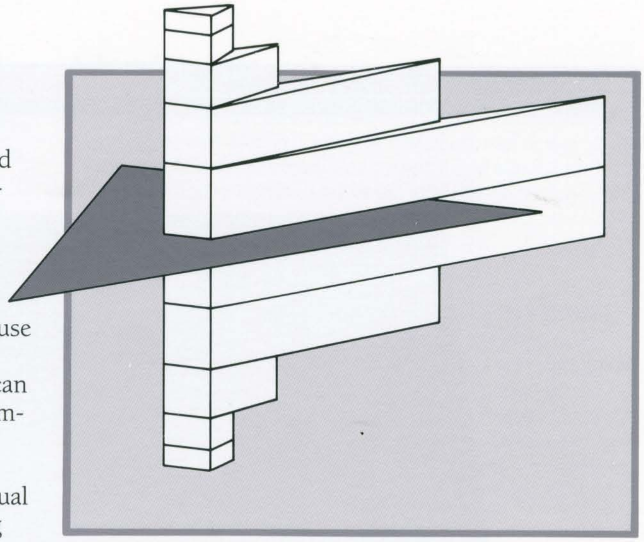

Figure 2. Structure of a glyph, which represents the plastic flow through a cross section of the mold. Each level shows flow direction and speed. The color of a level indicates its temperature. The pressure of the plastic is shown by the color of the base plane in the center of the glyph.

data plotted across the gap between the mold walls (Figure 2). The velocity determines the length of an isosceles triangle that is extruded into a three-dimensional pie wedge. This wedge forms a layer of the glyph. Each glyph rests on a triangular base, which is the domain of a single finite element. Because flow in a given finite element is constrained to the same direction (although potentially different magnitudes), the glyph swivels, acting as a weather vane to indicate flow direction.

Pressure and temperature values are conveyed with color. The injection-molding glyph is painted to indicate both the temperature variation over the vertical profile and the pressure indicated by the coloration of the finite element on which the glyph rests. Since two variables are represented with color, care must be taken in selecting color maps to avoid visual confusion. Hence, the color maps selected for the glyphs as shown at the top of Figure 3 are not spectral, the most common choice of color maps. Both maps vary from dark colors at the low values on the left to light colors at the high

Figure 3. Screen from an interactive session for a simple mold. The left side of the mold is 3 mm thick and the right side is 1 mm thick. The flow exhibits hesitation as the plastic does not enter the thinner side of the mold.

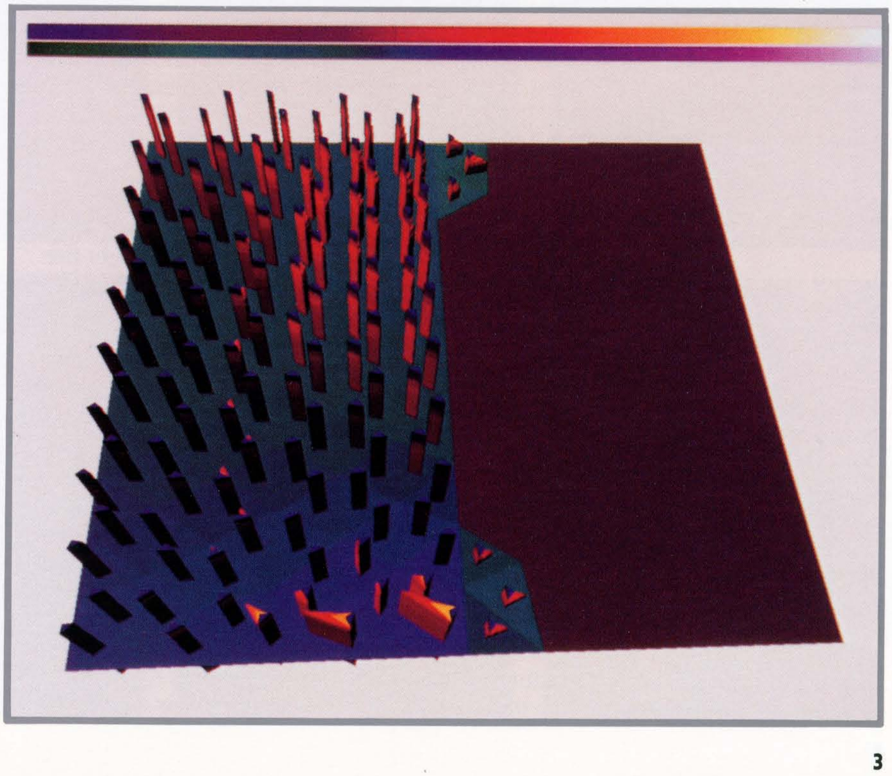

Figure 4. A later time step in the same mold-filling simulation as that shown in Figure 3. This illustrates some of the interactive tools available in the Kodak Glyph Visualization System. The mold wall is shown as a semitransparent surface. To highlight fast flow, glyph levels for plastic moving at less than 10 percent of the maximum velocity are not shown.

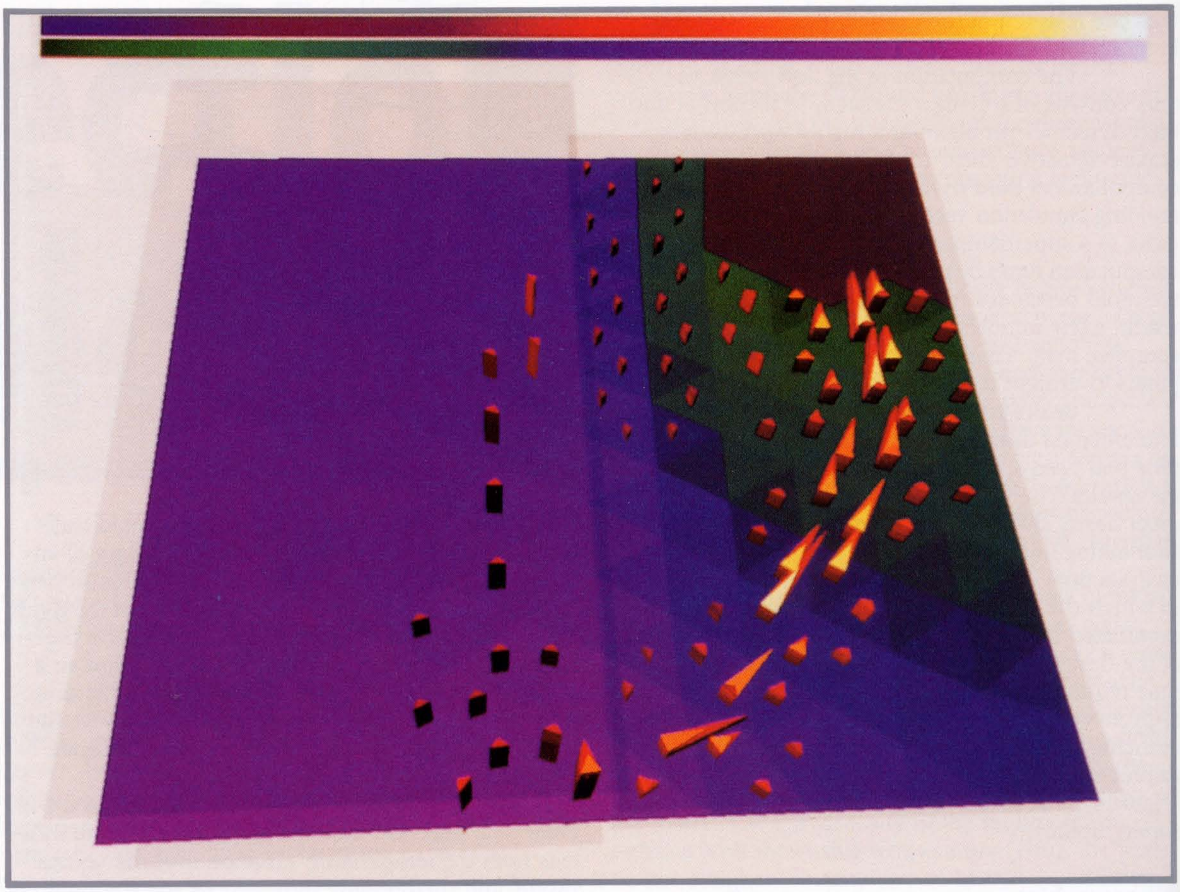

values on the right; however, the colors on the two maps do not overlap. The top map, corresponding to the temperature, has a larger perceptual color gradient to accent the high spatial variation of temperature values.

Isolated views, however, are not sufficient for understanding the complex flow in a three-dimensional model. Interactive control provides the ability to watch the model fill or stop on a single frame while moving the viewpoint to the best position to observe the flow. In addition to complete view and frame control, the graphical interface provides control of other variables to aid in the visualization of the model. Because the mold thickness may be difficult to perceive, even though it is illustrated by the height of the glyphs (Figure 3), a transparent surface can be used to show the outline of the mold (Figure 4). To highlight the areas of most rapid flow, layers of fluid moving slower than a user-defined value are not displayed.

## Applications

The simple mold shown in Figures 3 and 4 illustrates the phenomenon of flow hesitation, a common behavior in injection molding. Hesitation refers to the reluctance of plastic to enter thinner regions of a mold where its fluidity will be lower. Notice that the plastic first fills the entire thick portion of the mold before the thin one. The white bands in the glyphs of Figures 3 and 4 indicate the presence of high-temperature plastic in the mold heated by viscous dissipation. This generation of heat appears more prominently on the thin side of the mold. The hot plastic acts as a lubricant

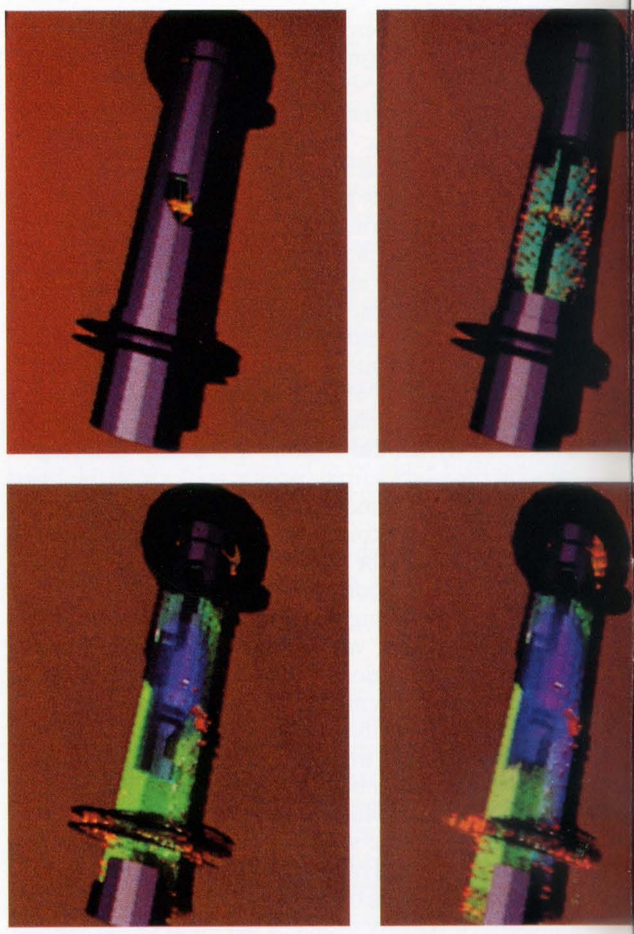

and concentrates the flow in the central region of the thin side when it begins to fill. These images were generated at a few frames per second in an interactive mode. The viewer has control (through sliders in SunWindows) over viewing-point and position, scaling of the glyphs, time during the simulation, presence of the transparent outline of the mold, and the velocity below which glyph layers will not be displayed.

The flat mold illustrates characteristic injection molding flow behaviors, but many molds have more complicated shapes. Injection molded parts often have tabs, slots, structural ribs, hinges, pins, cooling fins, screw holes or threads, because all these features can be molded at once. These complex molds may require more sophisticated graphics. The color plates of Figure 5 show frames from the filling of a more realistic mold shape — the spool at the core of a 35 mm film roll. In this case, because the object was cylindrical with interior parallel ribs, ray-tracing and transparencies were added to allow views of the entire flow. Interactivity must be sacrificed in these cases, because nearly a minute is required to generate images at a 1024-by-1024 resolution.

## Conclusion

In many fields that rely on computer simulation, such as injection molding flow analysis, the supercomputer has become a necessary tool for lowering turnaround time for large simulations. Supergraphics add to productivity by communicating the results of simulations clearly and in less time.

### *Acknowledgement*

*The visualization software owes its conceptual development to a joint effort with Donna Cox, professor of art and design at the University of Illinois at Urbana-Champaign and adjunct professor of supercomputing applications at the university's National Center for Supercomputing Applications.*

### *About the authors*

*Richard N. Ellson received a B.S. degree in fluid and thermal science and an M.S. degree in mechanical engineering from Case Western Reserve University. He researches flow simulation techniques for application to manufacturing processes at the Eastman Kodak Company. As a recipient of a Kodak Doctoral Award, Ellson is completing a Ph.D. degree in mathematics at the University of Illinois, Urbana-Champaign, while on assignment for Kodak at the university's National Center for Supercomputing Applications.*

*T. Marc Olano is completing a B.S. degree in electrical engineering at the University of Illinois, Urbana-Champaign. Olano currently is a consultant to Kodak on applications of color and real-time graphics interfaces to scientific visualization.*

### *References*

1. Ellson, R. and D. Cox, "Visualization of Injection Molding," *Simulation,* 51:5, 1988, pp. 184-188.
2. Wang, K. K., et al., "Computer-Aided Injection Molding System," Progress Reports Nos. 1-16, Cornell University, 1975-1989.
3. Tufte, E. R., *The Visual Display of Quantitative Information,* Graphics Press, 1983.
4. Sacco, W., et al., *Glyphs: Getting the Picture,* Janson Publications, 1987.

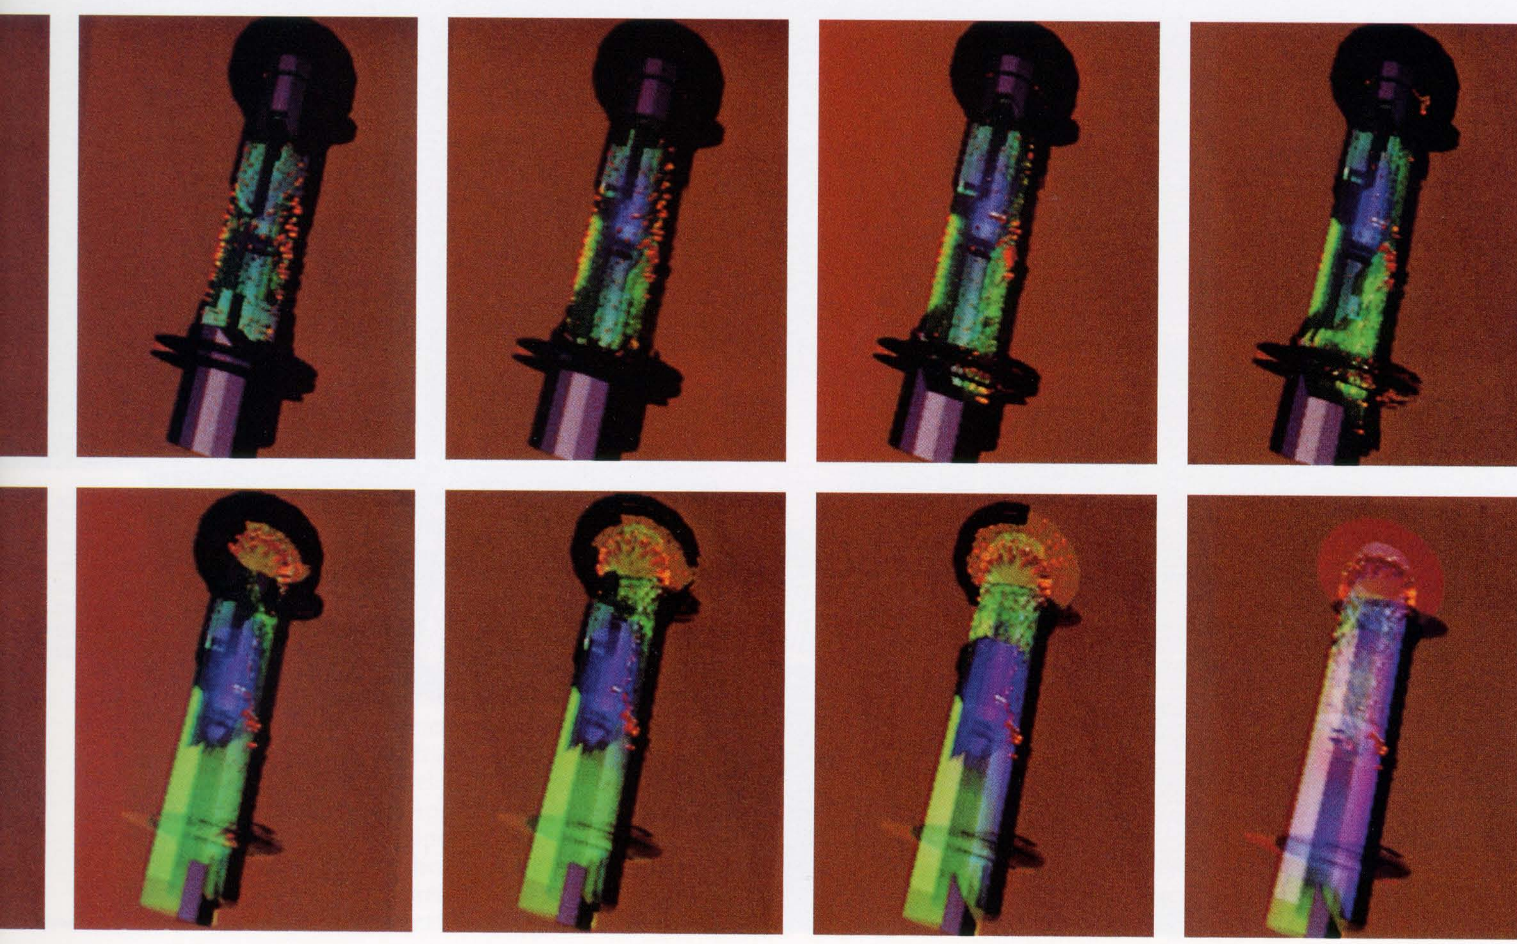

**Figure 5. A sequence of 12 exposures from the simulation of the molding of a 35 mm film spool. The colors are the same as those shown in the molding of the slab in Figure 3. Under the molding conditions in this simulation, the mold does not fill completely. A single element remains at the bottom of the spool, which does not fill. While the upper half of the part was filling, the lower half cooled sufficiently and the fluidity declined. The effective solidification of the plastic blocks the more fluid material from reaching the empty region.**

# Supercomputer research and engineering applications at Apple

*Mike Obermier, Gus Pabon, Malcolm Slaney, and Larry Yaeger*
*Apple Computer, Inc., Cupertino, California*
*Steve Nowlan, University of Toronto and Carnegie-Mellon University*

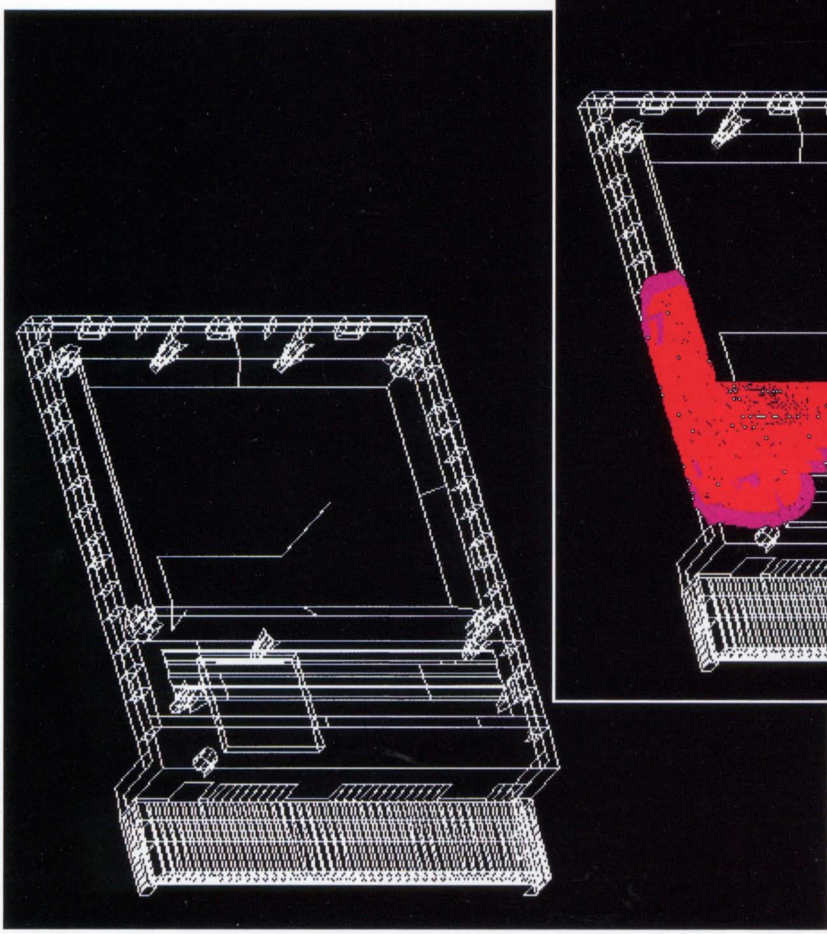

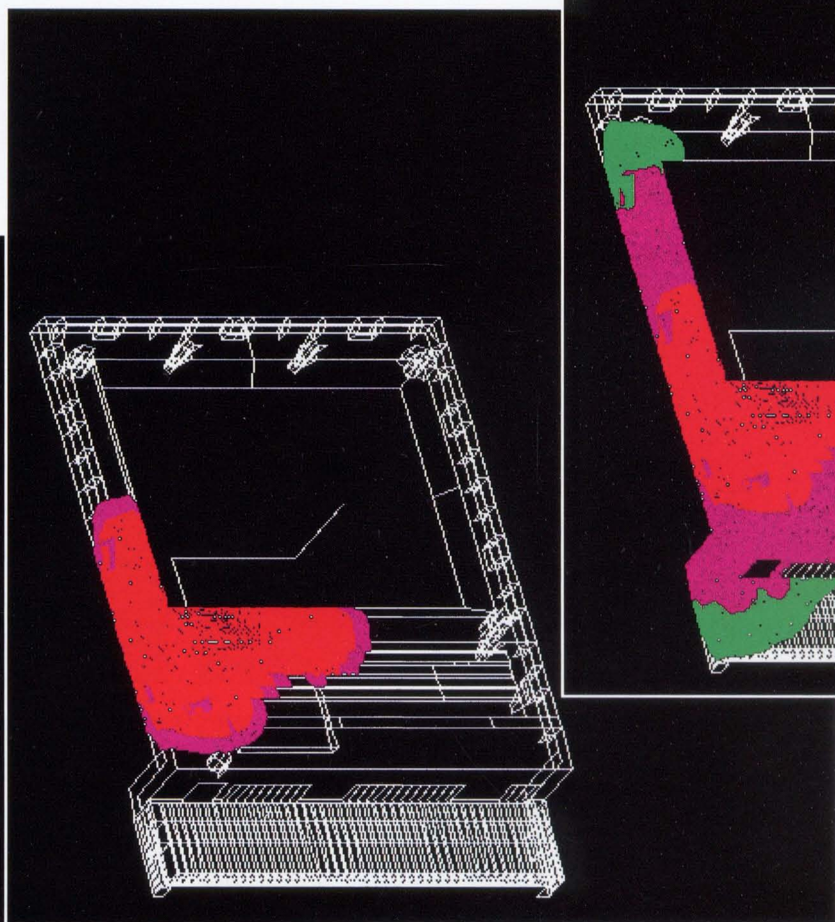

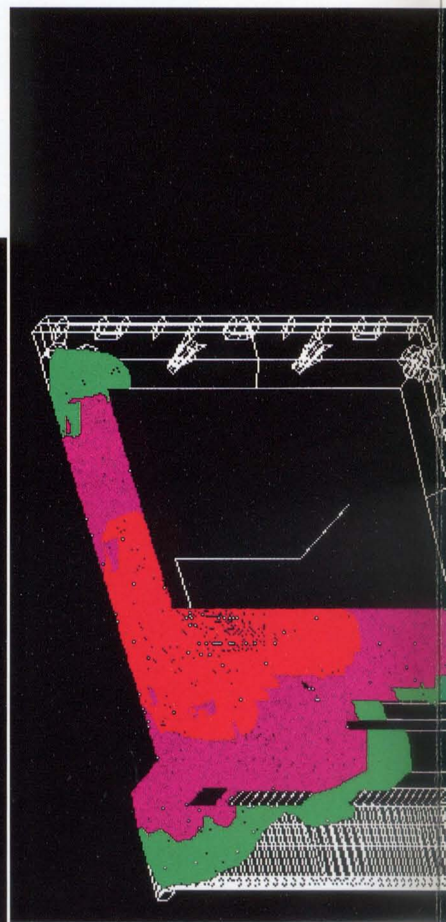

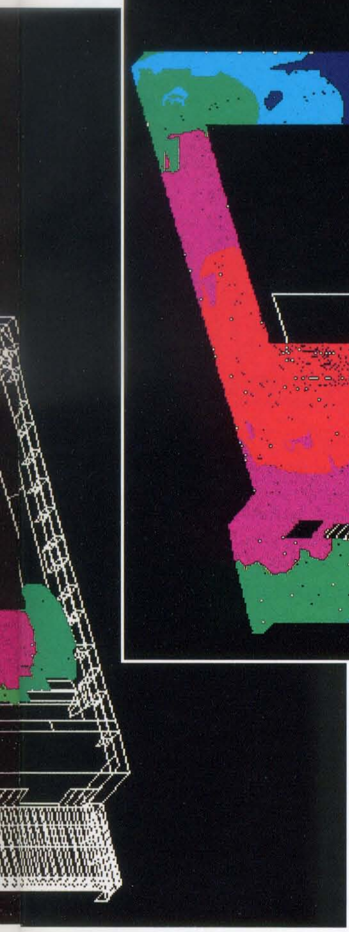
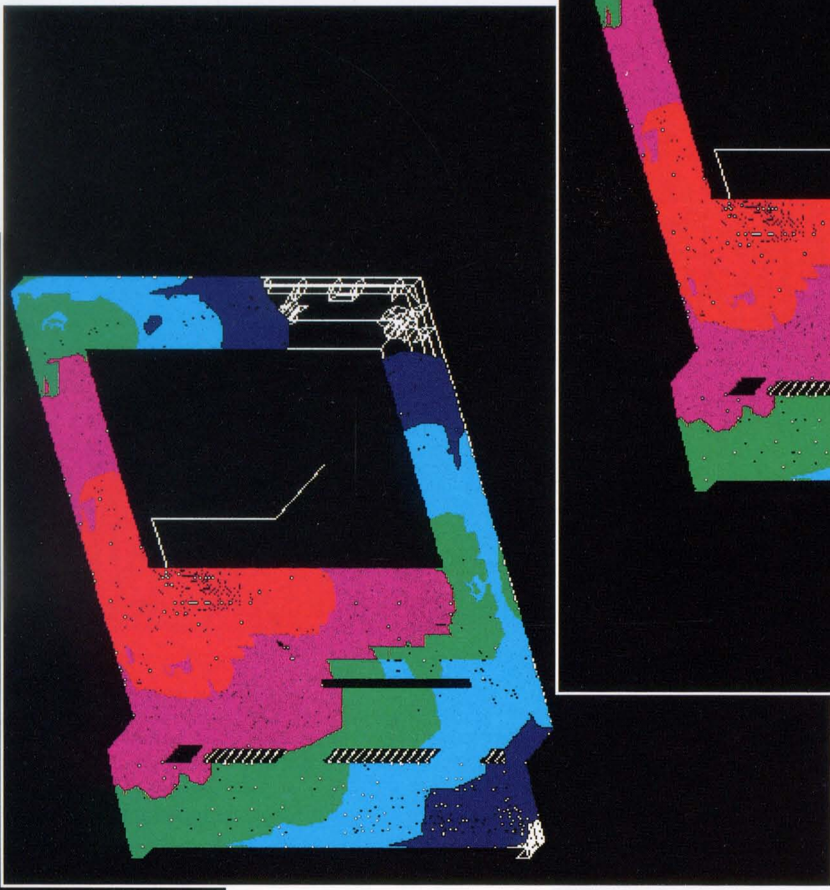
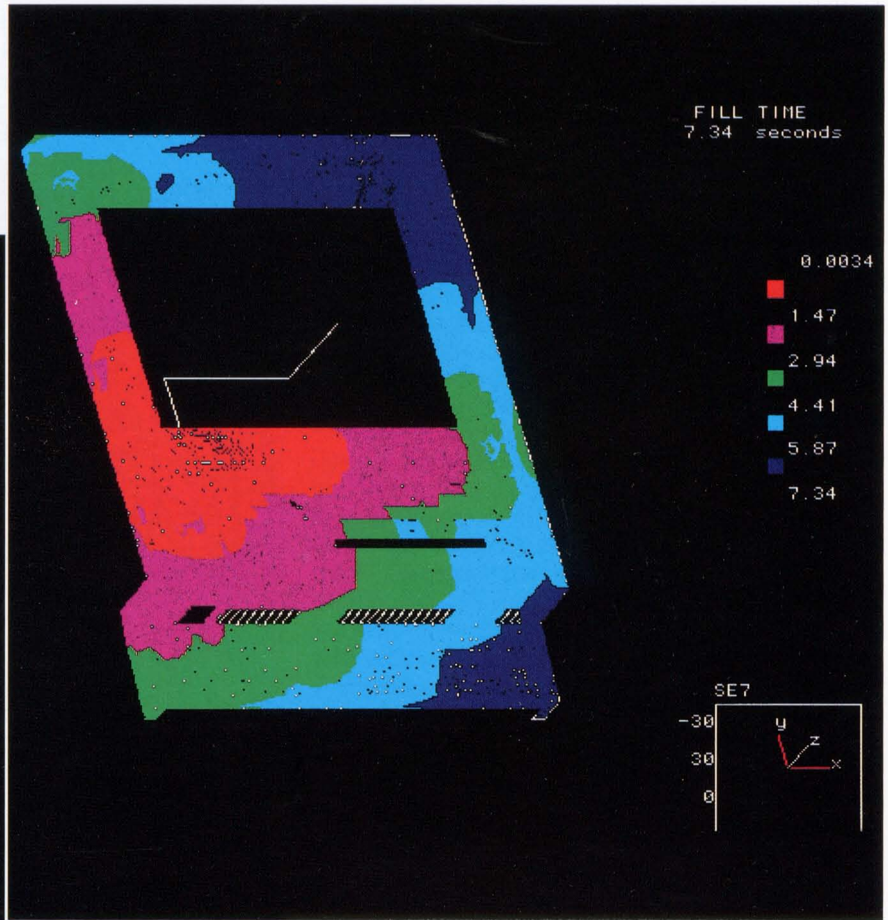


Injection molding simulation for the Macintosh SE computer bezel. The left image shows the model geometry. The following four images show, left to right, the extent of the flow field within the mold at various times up to completion.

In 1986, Apple Computer acquired a CRAY X-MP/48 computer system to support a wide range of advanced engineering and research projects. As in most scientific computing environments that include supercomputers, the Cray Research system at Apple is used primarily to simulate physical phenomena. The system's processing speed and memory size enable it to solve large sets of equations that describe the behavior of complex devices and processes. This method of research in turn enables engineers and research scientists to test and refine many ideas before committing resources to prototyping and production.

This article focuses on four applications that run on the Cray system at Apple. Plastic injection molding and advanced power supply development are engineering applications that address immediate product development and production needs. Artificial neural network simulation and speech recognition research are long-term research applications aimed at the development of more versatile user interfaces for future Apple products. The Cray Research system is a general-purpose supercomputer that enables Apple researchers to apply supercomputing technology to many areas of research, development, and production. The benefits to Apple include cost savings in research, faster product development cycles, and more reliable and aesthetically pleasing products. The benefits to Apple's customers include more economical, attractive, and highly functional products that become available sooner than otherwise would be possible.

## Plastic injection molding

The plastic cabinet that houses a computer's electronics serves many important functions, including representing the product to the user. As a result, computer manufacturers invest considerable time and material resources designing cabinets and the production processes associated with them. Manufacturers want to minimize cabinet flaws, which can be very visible and can be costly to a company if they result in reduced production levels.

Mike Obermier, a plastics engineer in the Mechanical Design Analysis Department at Apple, is using the Cray system to model the injection molding process that produces the cabinet components used in various Apple products. Simulating the process on the Cray system enables Obermier to pin down precisely the optimal mold design and molding conditions required for new parts before committing a particular design to the machine shop for tooling. The time and expense involved in retooling a mold is saved each time a design iteration is evaluated on the computer instead of being physically machined.

Aesthetic defects common to injection molded parts include weld lines, which form where flow fronts meet, and gate blushes, which are blemishes

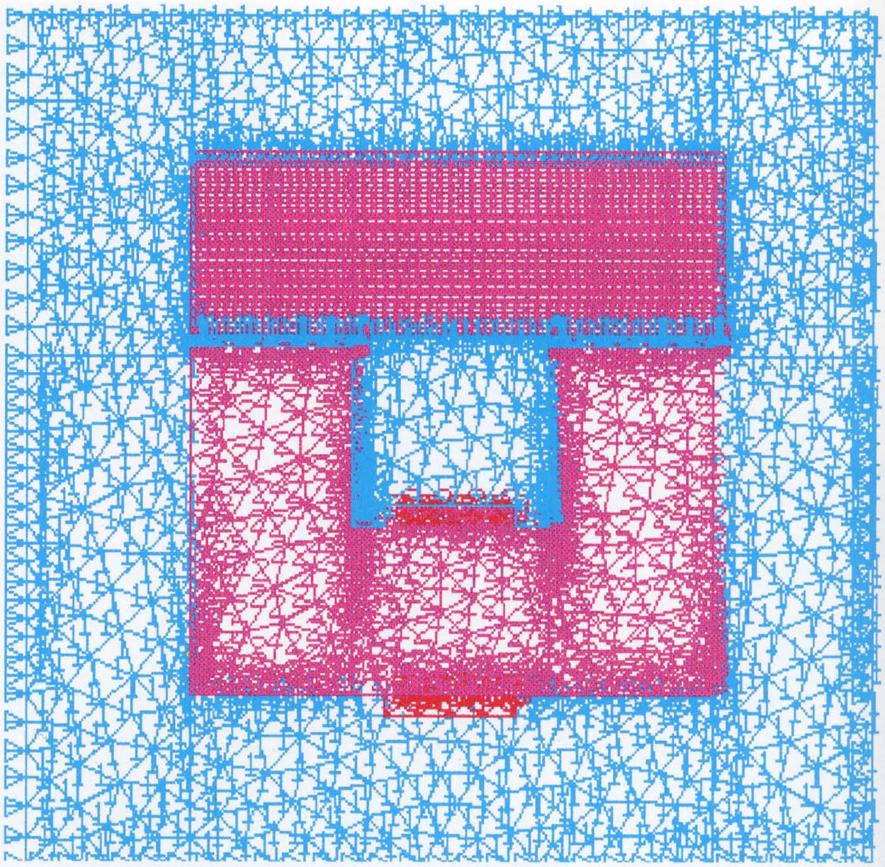

that occur at the points where injected plastic enters a mold. These types of aesthetic defects are virtually unavoidable in injection molding, but by using the Cray system, Obermier has been able to reduce the number of these types of blemishes and to place the remaining ones inconspicuously. In the case of the bezel (the frame surrounding the monitor) of the Macintosh SE computer, Obermier was able to evaluate several mold designs under various injection temperatures and pressures and improve the appearance of the parts by adjusting wall thicknesses, relocating the material entry points, and reducing their number of entry points. Ultimately, one weld line was eliminated, one was moved to a corner of the bezel, where it is hardly noticeable, and one was moved to another inconspicuous location. The number of gate blushes was reduced from two to one, with the remaining one relocated out of sight, hidden in a radius, a curved corner of the bezel. The simulation technique has been applied to parts across the entire family of Macintosh products, including the new 15-inch Portrait monitor and the Apple extended keyboard.

As a result of modeling on the Cray system, new molds can be tooled optimally the first time. The models provide templates and processing guidelines for each type of tool, so that approaches for similar parts will be designed from a starting point very near the optimum. By making tooling more efficient, this methodology not only saves time and money, but because Apple relies on tool makers in Ireland and Japan, it also saves overseas travel and communication expenses related to quality control.

The time savings involved in this type of design methodology in particular makes it ideal for high-volume production. If modeling saves one week in production that otherwise would be spent retooling, which is a reasonable estimate based on our experience, then the product can be brought to market one week sooner. In the personal computer business, one week's worth of sales for a given product more than justifies the cost of the supercomputer and modeling methodology for the entire year.

This work is conducted using the MOLDFLOW software package to model the flow of molten plastic and the MOLDTEMP package to perform cooling analyses to optimize the mold cooling systems and to study and minimize warpage of the molded parts. These software packages are products of Moldflow Pty. Ltd. of Melbourne, Australia.

## Power converter design

Lowering the power needs and increasing the efficiency of electronic components is an important objective for computer development engineers. Inefficient, high-power subsystems generate heat, which may accumulate and overheat the system if not adequately removed. Macintosh computers include fans to help maintain the desired operating temperature, but ultimately Apple would like to eliminate the fans, which can be loud, often are the first components to fail, and are relatively expensive to build into systems. One way to eliminate the fans would be to reduce the heat generated by the computer's electric components, including the hard disk drive and power converter. Gus Pabon, an advanced development engineer at Apple, is using the Cray system to design and implement a completely surface-mountable power converter with more than 90 percent conversion efficiency for use in future Apple products. The new power converter not only will reduce heat loss, but also has the potential

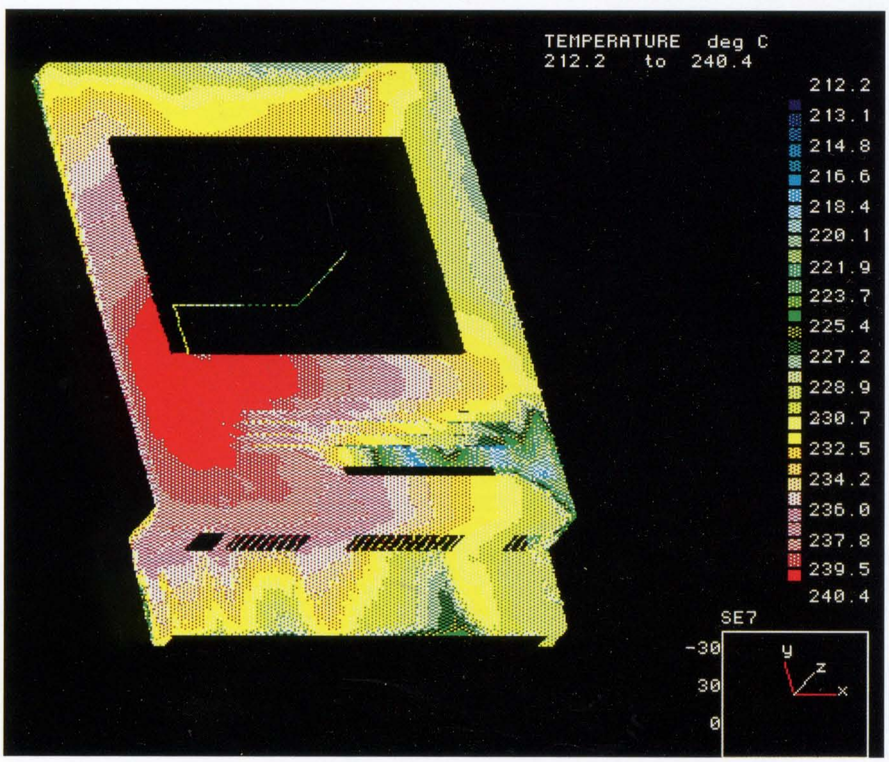


**Temperature distribution within the bezel mold at filling completion.**

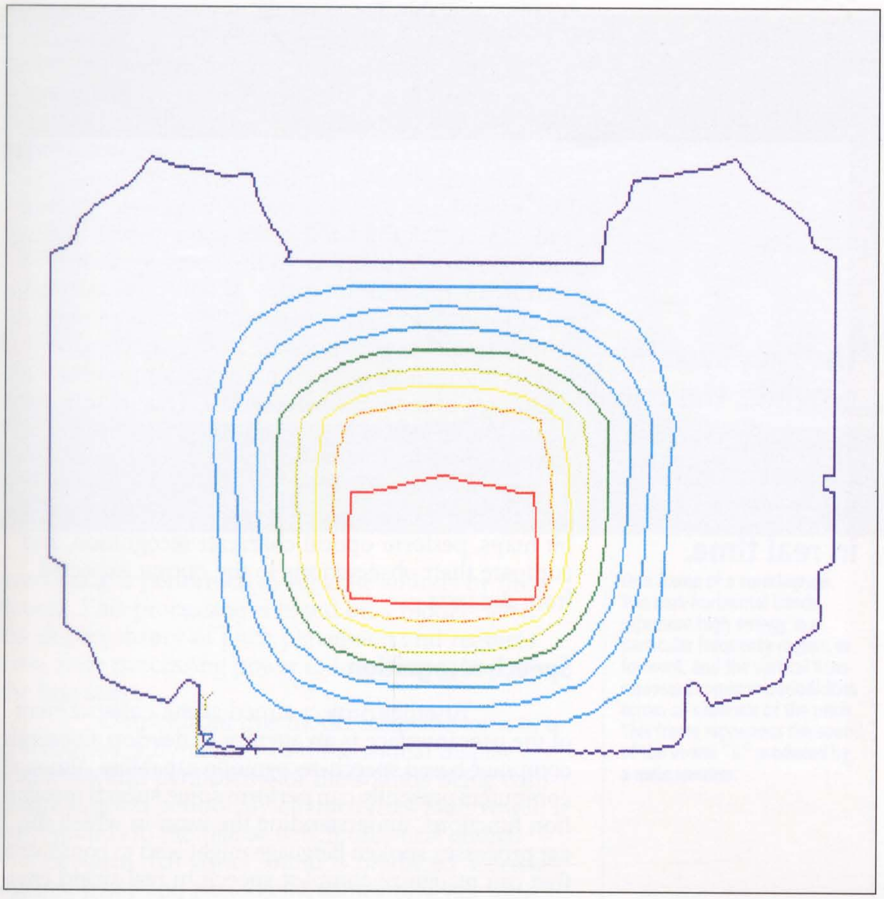

relatively large and weighed up to several pounds, which made them too heavy for automated pick-and-place machines to handle. The size of a transformer is determined largely by the frequency at which it operates, and older transformers were large and heavy because they operated at relatively low frequencies. However, during the past few years, advances in materials have made higher-frequency transformers practical for use in personal computers. As a result, transformers can be built that weigh as little as two ounces, light enough for automated equipment to handle. The higher frequency transformers also are more efficient and dissipate less power than the older ones.

Pabon is running the ANSYS finite-element program from Swanson Analysis Systems, Inc., on the Cray system to model high-frequency transformer designs. The models predict the direction and magnitude of the magnetic flux lines and the flux densities of the transformer designs for various core geometries. These characteristics reveal the efficiencies of the transformers described by the various design specifications. By modeling many designs in this way, the most efficient design can be arrived at quickly, while minimizing the time, materials, and other resources spent during the process. Finite-element analysis is the only way to conduct this kind of transformer core design research, and the Cray system's architecture is ideally suited to solving finite-element problems. Pabon hopes to demonstrate a prototype of the new power supply to Apple upper management later this year, and if it is approved, the design will be assigned to a specific Apple product.

**Transformer core simulation for a new power supply design. The images show the finite-element mesh structure of the core (above left), lines of equipotential magnetic flux (above), and flux density magnitude plots (below right).**

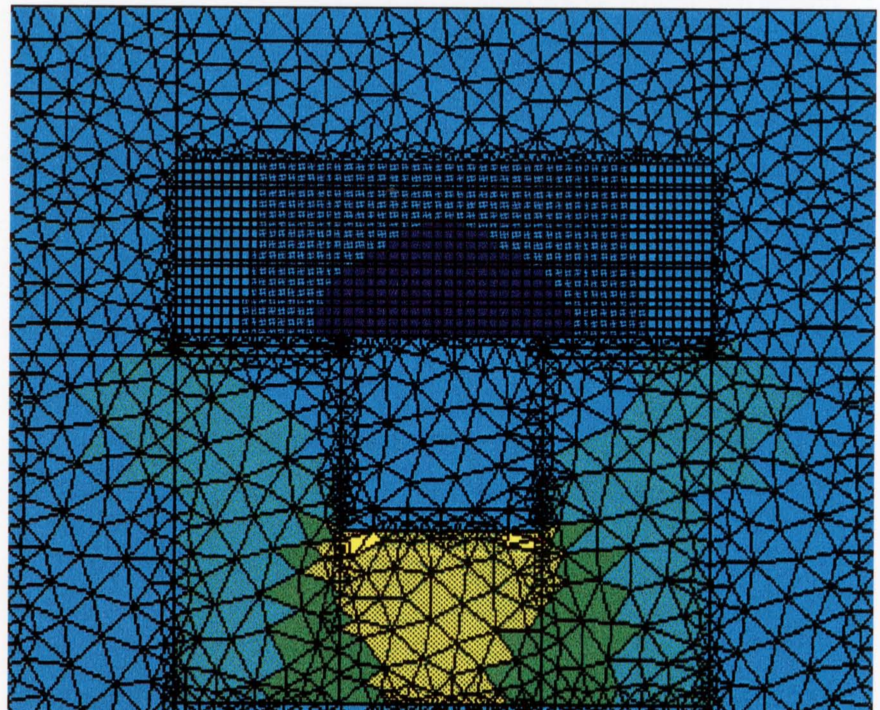

to be much smaller than present converters. As a result, its use will free-up space inside the computer that then can be used for additional logic.

The power converters currently used in Macintosh computers operate at about 75 percent efficiency. This means that 25 percent of the power is dissipated as heat during conversion from wall-socket AC to the DC current used internally by the computer. By redesigning the circuit topology from a conventional switched-mode converter to one that produces nonconventional waveforms, Pabon has been able to design on/off switches that minimize power loss by operating at a higher conversion efficiency.

The circuit designs that Pabon has developed for power converters are complex to analyze and require the solution of large sets of second-order differential equations. He is using the analog circuit simulation package SPICEPLUS from Analog Design Tools, to optimize the design of the circuitry. Pabon is using the Cray system to cut development time for the new circuitry by rapidly solving highly iterative time-domain simulations. Pabon most recently has concentrated on the design and analysis of transformers for use in power supplies. Transformers are a major source of efficiency loss in power supplies. In addition to being more efficient, the new power supplies must be surface-mountable; that is, automated pick-and-place machines must be able to handle them during manufacturing. Traditionally, transformers used in power supplies were

## Artificial neural networks

Along with using the Cray system in applied research aimed at near-term product development and production, Apple engineers are using the system for basic research projects that are longer term. Among these are projects aimed at increasing the flexibility of the Macintosh user interface. The interface, based on a desk top, mouse, and pull-down screen menus, is perhaps the most distinguishing feature of Apple's Macintosh computers and is largely responsible for the success of the Macintosh line. One project aimed at enhancing the interface involves the development of artificial neural networks that will enable future systems to perform optical recognition of handwriting and to extract phonemes from printed text. Larry Yaeger, principle engineer of Apple's Vivarium project, which is an attempt to design technologies that incorporate principles of biology, and Steve Nowlan, a Ph.D. student working jointly at the University of Toronto and Carnegie-Mellon University, are developing and training neural networks on the Cray system for these applications.

Artificial neural networks are computing systems designed according to principles that also govern biological nervous systems. Although many types of artificial neural network architectures exist, all include weighting factors that are adjusted to channel data selectively through the network as the network is trained to perform a particular task. And although the goal of neural network research ultimately is to implement artificial neural networks in hardware, researchers perform their design work on software models that run on existing computer hardware. Artificial neural network models are large sparse-matrix applications that are highly vectorizable, and therefore take good advantage of Cray system capabilities.

To develop networks capable of recognizing handwriting, Yaeger and Nowlan are beginning by training a network to recognize handwritten examples of the digits, 0 through 9. The network optically scans several handwritten versions of each digit repeatedly, each time computing the likeliest candidate and comparing that guess to the digit actually represented. By cycling through a sample of handwritten digits many times, the network may be able to "learn" to recognize handwritten digits by recognizing distinguishing features of each digit. The network used in this work is a moderately sized supervised back-propagation network with 256 input units, 50 hidden units, and 10 output units (one for each digit).

The phoneme-extraction network is a larger supervised network with 203 input units, 50-100 hidden units, and 57 output units (one for each phoneme). By microtasking the code for the CRAY X-MP system's four CPUs, Yaeger and Nowlan were able to improve the performance of the model by a factor of about 3.2, and the model now runs at about 6.8 million interconnections per second. The network is being trained on a vocabulary of 20,000 words. It cycles through the word list, assigns a number of phonemes to each word, and compares its guesses to the actual phonemes that the words represent. This type of network eventually may be able to read printed text aloud in near-real time as the text is entered.

Once the capabilities of neural networks are better understood, many other applications may become available. For example, as a user interacts with a Macintosh, the computer could "learn" the user's window placement preferences and automatically place windows on the screen where the user would want them. A more far-reaching potential application for neural networks within Apple is their use in hardware and software reliability testing. During testing, many Macintoshes are driven remotely to exercise hardware and software, but in some cases the only way to get the necessary feedback from the tested systems is to observe their screen bit maps. Although memory and other functions can be driven and monitored remotely without interference, the need to actually observe screen bit maps to monitor output compromises the usefulness of remote quality testing. Quality tests can run for hundreds of hours, so human screen-watchers are not a practical solution; the process must be more automated. However, properly designed and trained neural networks may be able to "watch" many screen bit maps, perform optical character recognition, and compare their observations to the output expected from the tests.

**The Cray system provides the only means by which the correlagram model might run in real time.**

## Speech recognition

Another project aimed at the enhancement of the user interface is an attempt to develop a versatile computer-based speech-recognition capability. Although computers presently can perform some speech recognition functions, understanding the ways in which the ear processes spoken language might lead to computers that can recognize complex speech in real-world environments. Apple research scientist Malcolm Slaney is working on a project to model sound processing in the human auditory system to help develop better speech recognition capabilities for computers. Among the phenomena he hopes to understand better is the ability of humans to attend selectively to a particular voice among many. If user interfaces for computers are to make wide use of voice interactions, then computers, like humans, will have to understand speech in noisy environments, such as offices or parties. Slaney describes the verbal interface as an adjunct to the keyboard and mouse that will make certain types of interactions easier.

The hearing research group at Apple has implemented two models to study sound processing in the inner ear. The models simulate the cochlea, a spiral fluid-filled tube in the inner ear. When we hear, the outer and middle ears transduce sound waves mechanically to the cochlea. This energy creates pressure waves in the cochlear fluid. The waves move hair-cells that line the cochlea, which in turn are attached to nerves that relay signals to the brain's auditory cortex.

The first inner-ear model is a black-box model in which the cochlea's internal structure is ignored. The second model includes representations of the fluid dynamics within the cochlea and of hair-cell physiology. The immediate goal of the modeling is to represent cochlear functioning accurately enough to predict correctly the probabilities of auditory nerve firings at each point along the cochlea. The Cray system at Apple has been especially useful in studying the processing of neurons at higher levels of the brain. After sound is converted from acoustic pressure waves to nerve firings by the cochlea, a second stage of

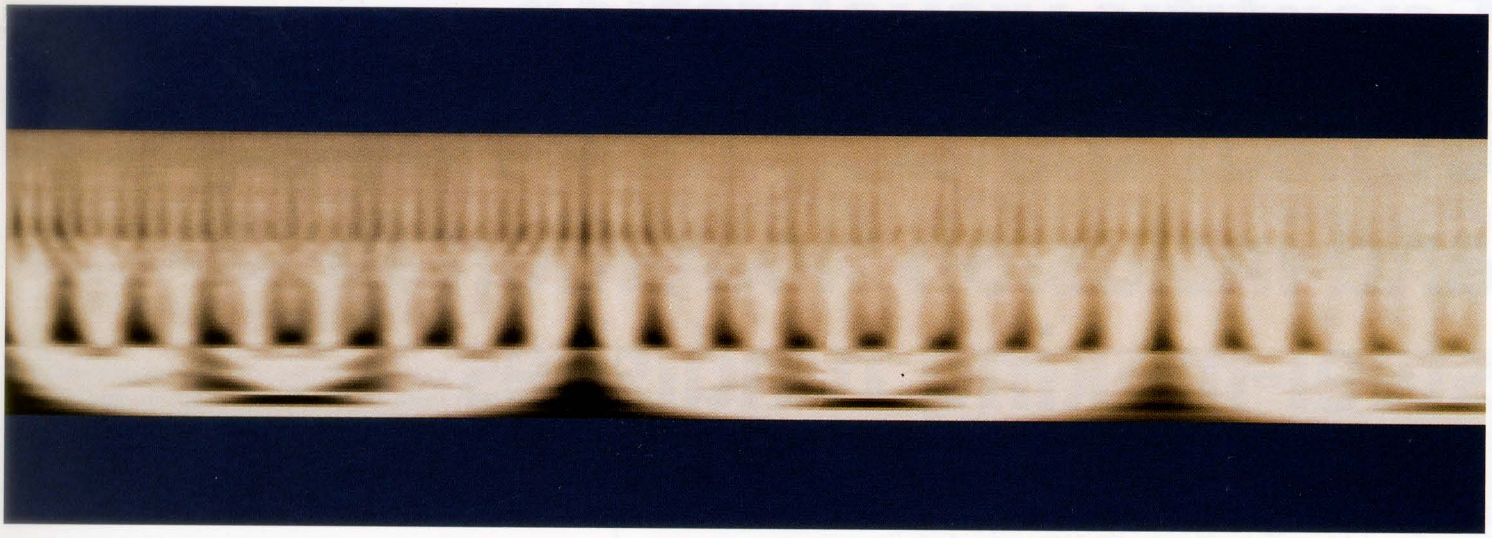

One frame of a correlagram. The dark horizontal bands represent high energy in a particular frequency region, or formant, and the vertical lines represent common periodicities across all channels of the pitch. This frame represents the sound of the vowel "u" produced by a male speaker.

processing is performed to separate sounds by their source. This processing is based on a model called the duplex theory of pitch perception and requires even more processing power to implement than does the first stage.

Output from this model takes the form of a correlagram, a dynamic two-dimensional display of frequency and periodicity for the sounds being computed. A given periodicity characterizes each vowel; it corresponds to the rate at which the vocal chords vibrate when that vowel sound is produced. Correlating frequency and periodicity, Slaney believes, may represent accurately the neural activity that occurs in the auditory system when spoken language is processed. Using the Cray system, Slaney is able to compute correlagrams in about half real time. He hopes to be able to produce correlagrams in real time with some code restructuring to make more effective use of the Cray system's multiple CPUs.

The correlagram model requires from 100 to 200 MFLOPs, and the Cray system provides the only means by which it might run in real time. The Cray system enables Slaney to run more data through the model in one week than had been run through it during the past few years using other computer systems. Scientists conducting basic research must be able to run many experiments in a short time, to refine their models through successive approximations. Without the fast turnaround provided by the Cray system, individual experiments of this type become painstakingly long to execute and can discourage researchers from exploring many alternative approaches.

A foreseeable milestone in the application of this research might be the development of systems that can participate in simple telephone conversations. For example, a person might be able to call his or her office computer and ask a calendar program to list upcoming appointments or to schedule new ones. As the technology becomes refined it will become able to handle increasingly complex interactions.

## General-purpose supercomputing

The flexibility of the Cray system makes it a valuable tool for Apple's engineering and research departments. By modeling physical processes on the system during early stages of research and development, engineers and research scientists are able to fine-tune their ideas; they can explore many options in design, development, and production that would be too time-consuming or expensive to evaluate by traditional means. This methodology increases the chances that the optimal solution to a problem will be found.

The applications discussed in this article represent only a sample of those that run on Apple's Cray system. Because the Cray system is a general-purpose computer system, we anticipate that many additional applications will be brought to it as more engineers and research scientists become familiar with its capabilities. This expansion in the system's use will involve both the transfer of existing applications to it and the development on it of new applications that would not be practical to develop on other systems. The point of making this level of computing power available to Apple research scientists and engineers is to minimize the amount of resources that are spent to develop new products by making the product-development process as cost-effective as possible.


***About the authors***

*Mike Obermier is a plastics engineer at Apple Computer. He received a B.S. degree in plastics engineering in 1983 from Ferris State College in Michigan.*

*Gus Pabon is an advanced development engineer in the analog group at Apple Computer, Inc. He received B.S. and M.S. degrees in electrical engineering from the University of Tennessee, Knoxville, in 1981 and 1983, respectively.*

*Malcolm Slaney is a member of the Advanced Technology Group at Apple. He received a Ph.D. degree from Purdue University in 1985. He is the author, with A. C. Kak, of* Principles of Computerized Tomographic Imaging, *IEEE Press, 1988.*

*Larry Yaeger is the principal engineer for Apple Computer's Vivarium project. He researches computer graphics and artificial neural networks for applications in user interfaces and artificial intelligence.*

*Steve Nowlan is a Ph. D. student working jointly at the University of Toronto and Carnegie-Mellon University. His thesis area is learning algorithms for artificial neural networks.*

# Scientific visualization of heat transfer, fluid flow, and inclusion floatation in steel-making tundishes

*Roderick Guthrie and Sanghoon Joo, McGill University, Montreal, Canada*
*Hany Greiss, Anders Grimsrud, and Kent Misegades, Cray Research, Inc.*

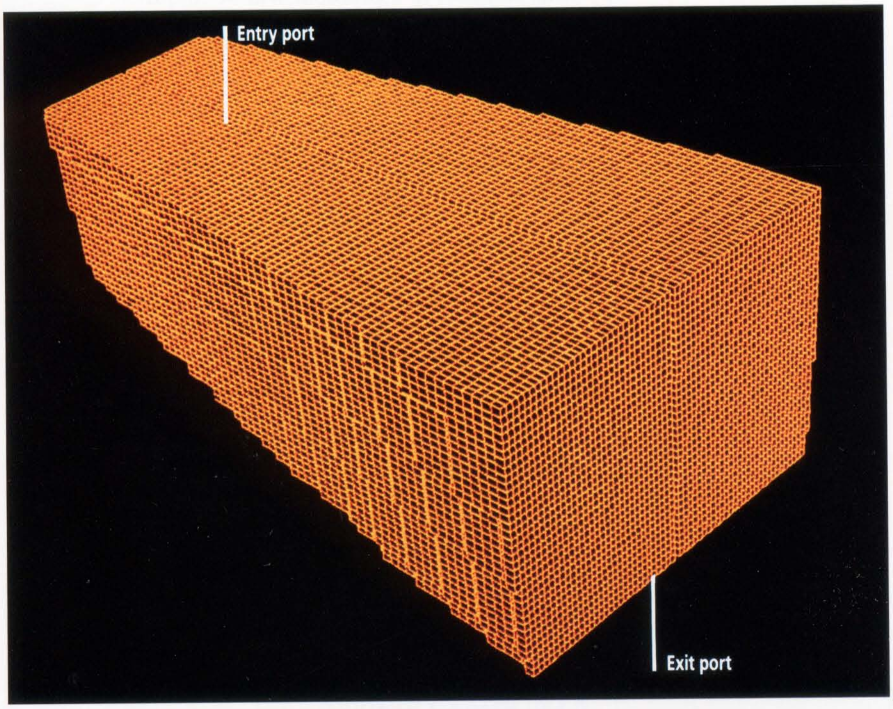


**Figure 1. Geometry of the tundish model.**

The Port Kembla works of BHP, Australia, is an integrated steel plant producing about four million tons of flat-rolled product per year. Among the plant's components are two types of tundishes — large containers that hold molten metal and deliver it to casting molds. To better understand the role of tundish design and flow modification devices in determining steel quality, studies of full-scale water models of the two types of tundishes used at Port Kembla were performed. The results of these investigations recently were reported by Dobson et al.[1] together with the findings of computational fluid flow analyses using the PHOENICS code.[2] To continue the mathematical analysis of these vessels, we have extended the description of fluid flow, heat transfer, and inclusion float-out, using the computer code METFLO, which was developed at McGill University to describe such transient, three-dimensional flow systems, and associated heat- and mass-transfer phenomena.

In this work, the METFLO code was extended to take into account the role of thermal natural convection, which typically has been neglected in computational models used in tundish studies. Although the computational task involved is greater when accounting for thermal natural convection in conjunction with convective heat transfer flows, the discrepancy in the modeling of low-Prandtl-number liquids with higher-Prandtl-number liquids is eliminated, as is the problem of matching Rayleigh numbers. Computations were performed on a family of Cray computers, including the CRAY-1S system at the Dorval Weather Center in Montreal, and a CRAY Y-MP system at Cray Research's computer center in Mendota Heights, Minnesota.

The tundish study reported here addressed a longitudinal wedge-shaped vessel with a downward-sloping base, and sloping side-walls diverging in the direction of the steel flow (Figure 1). This relatively unusual vessel design allows for a single exit port along the bisecting longitudinal plane for casting wide slabs. Alternatively, it can be fitted with two exit nozzles for the simultaneous casting of two narrow slabs. The ladle shroud for the entry jet is located close to the narrow, shallow end of the tundish.

## Mathematical model

To describe fluid flow, heat transfer, and particle (inclusion) float-out in such vessels, the partial differential equations that need to be solved are the equations of continuity, momentum, energy, and species conservation expressed in Cartesian form. A rectangular grid of elemental fluid volumes, or cells, was used to discretize these equations for numerical solution. Cell blockage procedures were adopted to allow for the inclined surfaces of the tundish. Similarly, relatively fine grids were chosen to ensure that the results were independent of grid size. The elements occupied a 40-by-80-by-40 matrix for the *x*, *y*, and *z* vectors. A set of typical boundary conditions was chosen for the extension of the METFLO code to include thermal natural convection phenomena. These included steady-state flows, heat losses, and an overlying slag that is wetting to inclusions.

### Heat losses

Steady-state heat conduction was assumed in the modeling of heat losses through the side-walls and the surface of steel in the tundish. Various constant heat flux conditions for the upper, side, and bottom surfaces were specified. Heat flux losses through the ladle base and side-walls were calculated to be about 2.6 $kW/m^2$ based on thermocouple implant tests and the thermal conductivities of the insulating materials. To estimate heat transfer through the slag, the upper surface heat flux was taken to be 31.0 $kW/m^2$, based on a stagnant molten layer of slag 30 mm thick conducting and radiating heat to the atmosphere.

### Fluid flows

A nonslip condition for steel flow was chosen for the bottom and side-wall surfaces, together with a free-slip condition at the slag/metal interface. The kinetic energy of turbulence within the entering jet was set at 3 percent of the entering jet's kinetic energy.

### Inclusion behavior

For particle/inclusion float-out, the number of particles separating to the surface of the melt was assumed to follow Stokesian behavior, wherein the inclusions, being wetting to the overlaying slag, were totally absorbed at the slag/metal interface. This leads to the inclusion flux equation

$$\dot{n}'' = u_s C^*$$

where $C^*$ represents the stagnant boundary-layer number density of mono-sized inclusions with a Stokes rising velocity of $u_s$. The analysis further assumed that the side-walls and bottom of the tundish were nonwetting (reflecting) to inclusions. Similarly, potential agglomeration/coalescence phenomena within the tundish were not modeled.

## Scientific visualization

In conjunction with this work, a video was prepared to visualize the massive amounts of data produced by the computational models. To create the video images, a distributed processing system consisting of a Silicon Graphics IRIS 4D graphics workstation and a CRAY Y-MP supercomputer provided the hardware platform. Cray Research's MultiPurpose Graphic System (MPGS) software package was used. This package consists of two parts, a portion that resides on the Cray system for CPU- and memory-intensive tasks, and a workstation-resident portion, which provides the user interface.

Certain calculations, such as multiple particle traces, required extensive computing. The MPGS package essentially used the CRAY Y-MP system for these operations, while the workstation was used for local graphic manipulations and control. The communication between the workstation and the Cray system was carried out over the network without user intervention. The MPGS package also is able to record commands for a sequence of images. Once an animation sequence is prepared properly, a standard analog video signal can be generated directly from the workstation onto a video recording system. The video segments of the movie were produced in this way.

Using this method, thermal profiles within the tundish were color-coded for temperature, while streak lines were used to illustrate the transient passage of massless particles between the entry and exit ports. This visualization revealed the rotational features present within the flow. These features led to a stochastic spiraling, and looping back of particles.

Similarly, the MPGS package allowed the visualization of a transient scalar field (envelope of isodensity inclusion levels) to be viewed three-dimensionally. The simulations illustrated the value of flow modification devices in reducing the number of larger inclusions entering the effluent steel from the tundish into the slab caster mold.

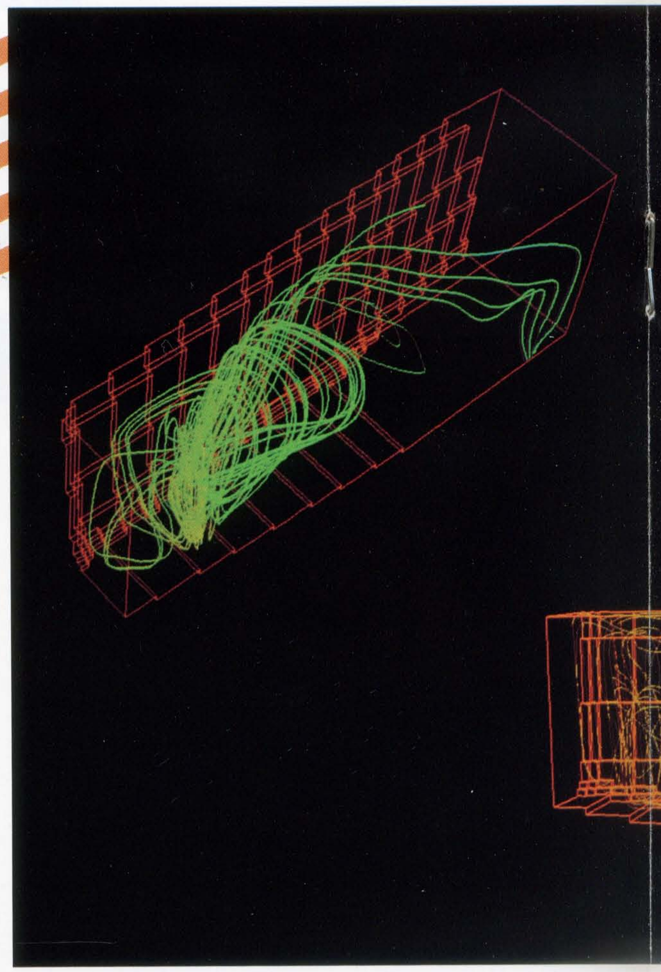

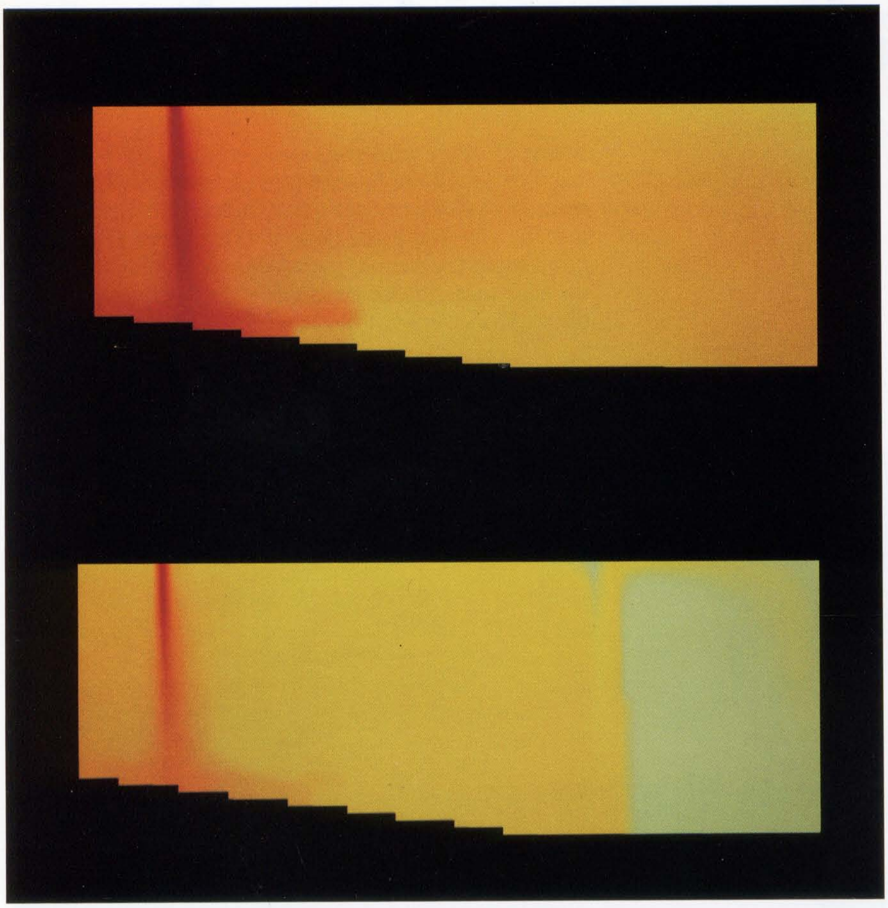

## Results

Computed temperature distributions within the tundish illustrate the boundary of the hot, vertical jet of steel, which enters at 1570°C (red), and the cooler steel (yellow) adjacent to the outside end-wall near the surface, which has a predicted temperature of 1562°C (Figure 2, top). A flow modification device (a refractory dam and weir) creates a clear delineation of temperature (orange to yellow) in the steel on either side (Figure 2, bottom). These calculations predict temperature drops of 10°C and 8°C respectively, with and without flow modification devices, and correspond well with drops typically observed.

Computed flows were presented visually by means of particle traces (Figure 3). These computations illustrate the marked effect of a dam and weir placed at three-fourths the length of the tundish as measured from the entry end.

Figure 4 shows an isodensity surface following the continuous input of large heavy inclusions. The flow control device was shown to be ineffective for small inclusions with $u_s \leq 0.5$ mm/s, but very effective for inclusions of 120 $\mu$m ($u_s$ = 4.5 mm/s), where a 70 percent improvement in effluent quality for that size would be expected. However, the complete elimination of these and larger particles is not achieved with this tundish design. At the time of this writing, the tundish is being operated with a ported dam arrangement.

**Figure 2. Computed temperature profile of the tundish without (top) and with (bottom) a dam-and-weir flow modification device.**

**Figure 3 (above right). Particle traces showing the circulation of molten metal through the tundish without (top) and with (bottom) a dam-and-weir flow modification device.**

## Conclusions

Several conclusions can be drawn from the results of these computational studies:

- ☐ Thermal natural convection phenomena in large, deep tundishes can generate strong downflows adjacent to side- and end-walls.
- ☐ Thermal convection can generate secondary recirculating flows and increased fluid motion near tundish exit ports. These flows reduce separation efficiencies for inclusions, but also reduce thermal cold spots.
- ☐ Flow modification devices can lead to significant improvements in steel quality for the intermediate (50 $\mu$m) and larger inclusions (120 $\mu$m), with effluent inclusion ratios typically exhibiting 50 percent improvements in metal cleanliness over nonmodified tundishes.
- ☐ The dam is the critical component of a dam-and-weir arrangement for enhancing inclusion separation.

# Modeling complex fluid flow with finite elements

*Philippe A. Tanguy, René Lacroix, and François H. Bertrand*
*Rheotek, Inc., Cap-Rouge, Quebec*

Supercomputing plays a key role in technological progress and innovation in many fields of science and engineering. In the field of materials engineering, supercomputing technology enables designers to develop new processing equipment or improve existing tools to deliver better products in less time. Many material-processing operations involve the flow of materials in a molten state through channels or reservoirs of complex shapes. These fluids are complex because their rheological behavior generally is nonlinear and strongly influenced by temperature.

## A finite-element approach

Since the mid-1960s, modeling fluid flow computationally has been recognized as a cost-effective alternative to experimental modeling, giving birth to a new field of scientific investigation called computational fluid dynamics (CFD). The finite element method is one of the most efficient mathematical methods that can be used in CFD. Its basic principle is the transformation of a set of partial differential equations governing mass, momentum, and energy conservation into a linear (or nonlinear) system of algebraic equations using variational principles. Each unknown in the matrix system is associated with a degree of freedom located at a point in the domain, that is, a velocity component, a temperature, a concentration, or a pressure.

The use of finite elements for modeling complex industrial fluid flow problems is an extremely challenging field of investigation because it involves the complexity of fluid flow simulation, the solution of very large matrix problems, and the treatment of stiff nonlinearities originating from the fluid rheological behavior. The POLY3D software package from Rheotek, Inc., was developed to apply finite element methodology to complex fluid flow modeling.

POLY3D is a finite element software package for the computation of three-dimensional rheologically complex fluid flow with thermal effects. It is based on an augmented Lagrangian formulation and uses an incomplete Uzawa solver. The former formulation enables users to compute shear-thinning, yield stress, and viscoelastic fluids by decoupling the computation of the velocity and pressure fields from the computation of the stress field. This algorithm has proven robust and efficient in many applications.[1,2] The incomplete Uzawa solver is an extension of the classical Uzawa and Arrow-Hurwicz algorithms and uses an incomplete Choleski factorization of the matrix as a preconditioner; this solver has provided outstanding performance in terms of memory requirements and CPU time.[3] POLY3D uses enhanced hexahedral and tetrahedral elements for the meshing of the computational domain.[4] The following examples present modeling capabilities of POLY3D running on a CRAY-2S/4-128 supercomputer.

## Three-dimensional mold-filling simulation

Injection molding is one of the key manufacturing technologies used in the thermoplastics industry. This process includes forcing a polymer melt into one or several cavities shaped like the finished product, and cooling the material once the mold has filled. This technology is used widely, and is applied to the development of various thin-walled parts for the automotive and aerospace industries.

The injection molding process is a cyclic operation, with each cycle composed of five stages: feeding, filling, packing, cooling, and ejection of the part. Though cooling is by far the longest cycle stage, the filling stage is likely to be the most important because at this time the molded part acquires its shape and physical properties. An accurate simulation of the filling stage involves the solution of both fluid flow and energy equations in thin layers. POLYFAN software, a specialized subset of POLY3D, was developed to model the filling stage.

In injection molding, the mold walls typically are 100°C to 200°C colder than the polymer melt as it enters the cavity, and a thin layer of plastic solidifies along the walls. The prediction of this skin is of great importance because the skin restricts the flow passage and eventually may freeze the polymer melt before the mold fills. Moreover, due to the rather poor thermal conductivity of the melt, the frozen skin constitutes an additional resistance to heat transfer that must be accounted for when devising the overall thermal control strategy of the molding process. The treatment of three-dimensional mold filling can be achieved either by laying flat the domain topology under crude simplifying assumptions or by directly considering the three-dimensional frame of reference. The latter approach is used in the POLYFAN software.

Numerical tests were conducted for the molding of a high-density polypropylene fish crate manufactured by IPL, Inc. The volume of this mold is 0.07 $m^3$. Five injection points located at the bottom were used to fill the mold. The finite element mesh is composed of 3760 hexahedral elements corresponding to 7588 equations. A subgrid of 8330 additional nodes, a total of 15,918 equations, is used for the thermal simulation. Figures 1-3 show the position of the filling front after 15 percent, 35 percent, and 90 percent of filling, respectively (short shots). These results are useful for detecting the location of the weld lines, the regions filled last, and more generally, air entrapment. This information will be used to position the vents in the mold matrix.

The position of the five injection points can be identified clearly in Figure 1. The formation of weld lines also can be anticipated at the bottom of the crate. Figure 2 reveals the nonsymmetry of the filling pattern. The delay of the melt front in the middle of the crate is due to the "choking" of the central injection point. At 90 percent, the mold filling is close to completion (Figure 3). This series of short shots was compared to experimental data provided by IPL and the results proved to be in excellent agreement.

Figure 4 shows the percentage of solidification in the mold just after filling has completed. It can be seen that the most sensitive region is around the stiffeners at the upper part of the mold. This was expected since the contact surface with the matrix wall is very large due to the shape of the stiffeners promoting heat transfer. Such a prediction is helpful in avoiding premature solidification during filling, which would foul the mold, as well as in devising a correct temperature control strategy.

The above computation required three hours on a CRAY-2S/4-128 system and about 12 Mwords of memory. To fill the cavity and determine the frozen skin thickness, 295 thermal fluid finite element problems on the same number of meshes were solved. For each problem, two coupling iterations were necessary to ensure stability. The average number of equations for each problem was 5000.

## Modeling of a stirred mixing device

Mixing is a unit operation routinely used in industry to homogenize fluids of different viscosities. Mixing devices consist of a vessel and a motorized shaft with a turbine. The design of mixing devices is governed mainly by three parameters: power, mixing time, and circulation time. Depending on the viscosity level of the fluids to be homogenized, various blades

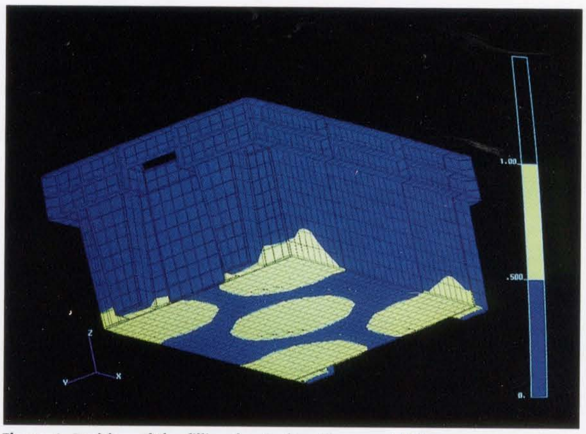

Figure 1. Position of the filling front after 15 percent filling.

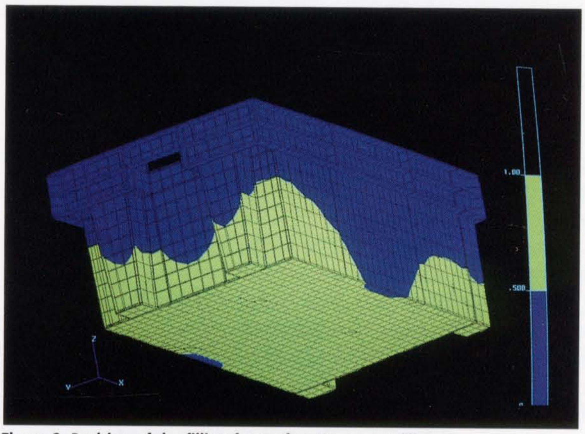

Figure 2. Position of the filling front after 30 percent filling.

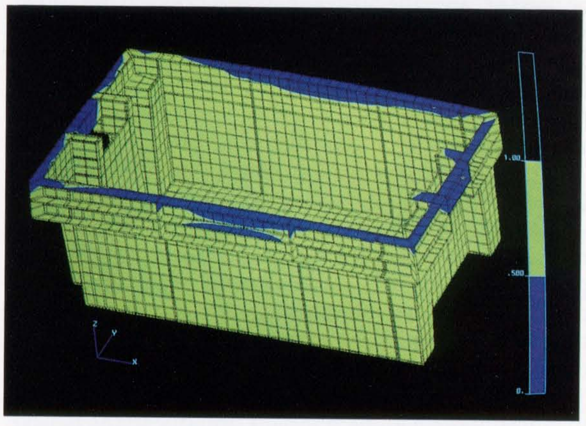

Figure 3. Position of the filling front after 90 percent filling.

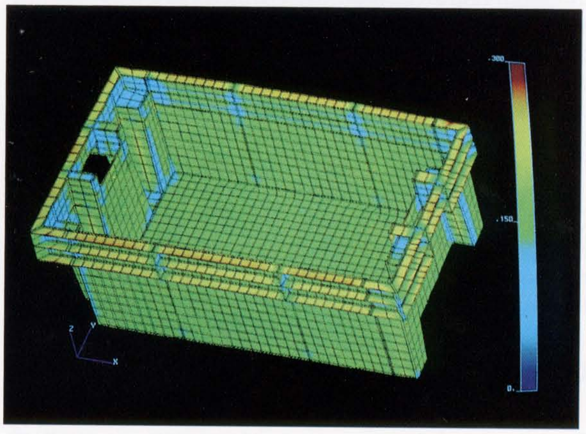

Figure 4. Percentage of solidification after filling completion.

**Supercomputing technology enables designers to develop new processing equipment or improve existing tools to deliver better products in less time.**

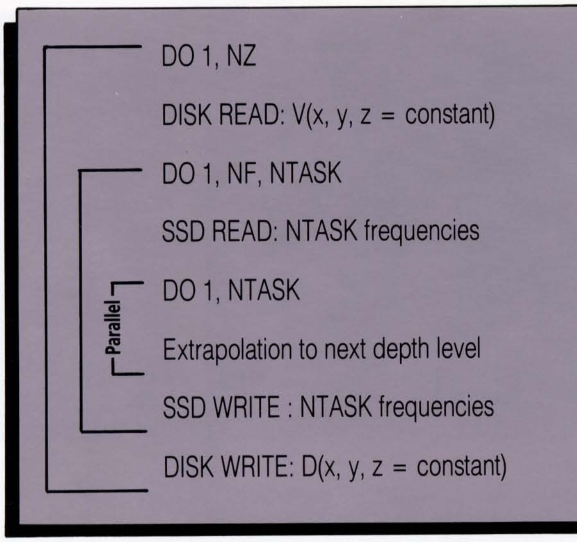


**Figure 1. Computational scheme for depth migration problem. *NZ* represents the number of depth levels, *NF* represents the number of frequencies, and *NTASK* represents the number of tasks.**

each frequency independently. Figure 1 is a simple representation of the computational scheme employed. Here, *NZ* is the number of depth levels, *NF* is the number of frequencies, and *NTASK* is the number of tasks. Almost all of the computations are performed in the highly vectorized extrapolation section that includes options that specify several types of approximation. The number of calculations is related directly to the accuracy required in the extrapolation. I/O operations, on the other hand, are the same for all extrapolation levels.

For most extrapolation options, each frequency requires the same number of computations. The parallel computational part is therefore easy to load balance with the above scheme, provided that *NTASK* is a multiple of the number of processors. Macrotasking subroutine calls are used to exploit parallelism in this high-level parallel scheme, although microtasking also has been used on this problem.

## Data management

The computational scheme shown in Figure 1 indicates that all I/O operations are performed serially. Obviously, this strategy eliminates the extra effort usually associated with optimizing I/O procedures. However, this scheme is justified only if the total I/O time, including system overhead, is insignificant compared to the computational time. An easy way to avoid I/O overhead is by overlapping I/O transfers and CPU work. The major drawback of this approach is that a "double-buffer" memory scheme must be used. For high-level parallel schemes, in which each task is assigned to work on a large portion of the global data, this memory scheme may become too expensive. The scheme presented in the figure also leads to extra memory allocation over the serial version of the code, since private arrays such as FFT work arrays must be duplicated on the stack of each processor participating in multitasking. For realistic problems, this means that each processor used requires a private array space of 0.5-1.0 Mwords. It is clear that in this case a large number of processors results in a large additional memory allocation, although this scheme uses much less memory than the double-buffer scheme.

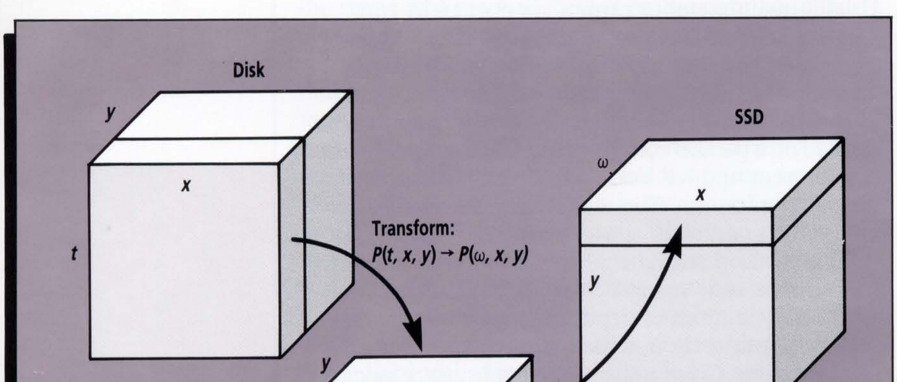


**Figure 2. Data organization during transposition.**

The extrapolation scheme implemented in this work is a one-step scheme. Arrays that contain the transformed data may be updated at each depth level. Since the frequency loop is the inner loop (typically 300-600 frequencies), only a single, horizontal "slice" (constant depth) of the velocity array and the output depth image are kept in memory, shared by all tasks. In addition, these arrays can reside on disks because they are accessed only once during the entire migration. The I/O-intensive element of this task is reading all of the frequency-transformed data and writing the data out at every depth level. To optimize this part, the input data are reordered and loaded to the SSD in a way that will permit consecutive data access in execution time. Figure 2 describes the data organization. Assuming the input data reside on disks in trace order, several *X-T* planes are read into memory and the data are transformed from time to frequency domains. The transposition to the SSD orders the data in frequency planes instead of traces. Optimal transposition requires two passes from disk to SSD. The first pass (described in Figure 2) is slow because in the general case the transfers cannot be performed in sector boundaries and the data are not stored consecutively. In the second pass, large blocks of *X-Y* planes (*NTASK* planes per block) are read in sector boundaries and stored in this order on disks. Next, the entire data set is returned to the SSD for execution. The very large I/O requests during execution minimize the system overhead. Sustained transfer rates exceeding two billion bytes per second usually are obtained. All I/O operations are performed using the word-addressable I/O routines. As the following examples show, the use of the SSD for large out-of-core applications enables effective machine use without sophisticated coding.

## Performance test

For the performance test, the most computationally intensive extrapolation scheme was used. The input data consisted of 135,000 traces ($NX = 375$, $NY = 360$), each for 1500 samples in time. After the data were transformed from the time domain to the frequency domain, 400 frequencies ($NF = 400$) were migrated. On a dedicated system running the UNICOS 5.0 operating system, 12 Mwords of main memory

| CPUs | CRAY X-MP system Time (sec) | CRAY X-MP system MFLOPS | CRAY X-MP system Speedup factor | CRAY Y-MP system Time (sec) | CRAY Y-MP system MFLOPS | CRAY Y-MP system Speedup factor |
|---|---|---|---|---|---|---|
| 1 | 280.60 | 147 | 1.00 | 199.61 | 206 | 1.00 |
| 2 | 143.65 | 287 | 1.95 | 100.90 | 408 | 1.98 |
| 4 | 75.14 | 548 | 3.73 | 51.71 | 797 | 3.86 |
| 8 | | | | 26.95 | 1529 | 7.41 |

CRAY X-MP system: 8.5 nsec clock, CRAY Y-MP system: 6.0 nsec clock

**Table 1. Results of the performance test. Figures are wall-clock times required to migrate all data to a specific depth level.**

and 110 Mwords of SSD space were allocated for the test. Table 1 summarizes the results measured on CRAY X-MP/416 and CRAY Y-MP8/832 systems. The table shows wall-clock times required to migrate all data to a specific depth level. During this time 1728 Mbytes of data were transferred to and from the SSD. Using very large I/O requests (69.1 Mbytes per request), the total I/O time was 0.794 seconds, with an average transfer rate of 2176 Mbytes/sec. Results indicate that parallel processing was used efficiently on both computer systems. Although I/O operations were performed serially, the test was more than 99 percent parallel, promising efficient use of a larger number of processors. For the same number of processors, the CRAY Y-MP system's performance is somewhat better than that of the CRAY X-MP system. This may be the result of lower memory contention on the CRAY Y-MP system.

The same migration was performed with DD49 disks (single streaming) on the CRAY X-MP system. The four-CPU test took 257.84 seconds, making the code completely I/O bound. Clearly, in this case, using even two processors would be difficult to justify. The above numbers show that a 60-70 Mbyte/sec transfer rate is needed to make the eight-CPU job computationally bound. I/O in this case must be overlapped by computations, and local memory demand will exceed 21 Mwords. Thus, the CRAY X-MP system's memory would not provide enough storage for the above problem. In many cases less-accurate extrapolation schemes may be used. Since only the amount of computations is reduced, higher transfer rates (100-300 Mbytes/sec) may be required to attain a CPU-bound job.

Time required to prepare the data on the SSD was not included in the test. With a single disk streaming, the initiation of the above data on the SSD required 275 seconds. To complete a migration of this type, several hours would be needed. The initiation time in this case can be neglected.

## One million traces problem

The feasibility of depth migrating one million traces on the CRAY Y-MP8/832 system was tested. The input array was 1000-by-1000 traces, with 2000 samples per trace. Since the frequency-transformed data are complex, 500 frequencies require one Gword of storage. With a 512-Mword SSD, it was necessary to pack the data into 32-bit words while residing on the SSD. The memory needed for eight frequencies was 16 Mwords. An additional array of 8 Mwords was assigned to occupy the packed data. (Unlike a double-buffer scheme, this array may be used as a work array during computation.) For each depth step, 1000 Mwords of data had to be packed and unpacked. This operation required about 25 seconds, and therefore was performed in parallel.

| CPUs | Time (sec) | MFLOPS | Speedup factor |
|---|---|---|---|
| 1 | 853.04 | 188 | 1.00 |
| 2 | 429.90 | 371 | 1.98 |
| 4 | 218.80 | 731 | 3.89 |
| 8 | 116.16 | 1377 | 7.34 |

**Table 2. Results of migrating one million traces to a specific depth level on a CRAY Y-MP8/832 system.**

Table 2 presents the results of migrating the above data to a certain depth level. Eight thousand Mbytes of data were transferred to and from the SSD (64 Mbytes per request). The effectiveness of the parallel scheme is well demonstrated. For such large problems, the overhead introduced by storing the data in 32-bit words becomes minor compared to the reduction in storage space and I/O operations. This example shows that with the current configuration of the CRAY Y-MP system, out-of-core problems of more than 1000 Mwords can be executed at a sustained rate well above 1 GFLOPS.

*CRAY CHANNELS regularly includes technical articles that offer insights into the Cray environment. The editors thank Chris Hsiung and Jim Schwarzmeier for their technical advice.*

***About the author***

*Moshe Reshef received a B.Sc. degree in geology from the Hebrew University of Jerusalem and M.Sc. and Ph.D. degrees in geophysics from Tel Aviv University. He is a consultant to Cray Research and a senior researcher at the Institute for Petroleum Research and Geophysics in Tel Aviv.*

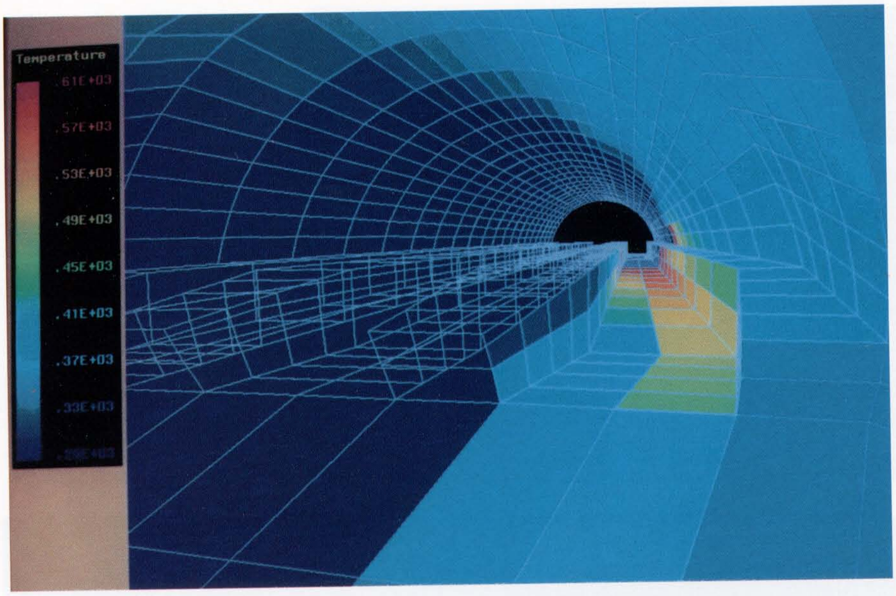


# Modeling the King's Cross station fire

Figure 1. A view from the computer model, looking down the escalator tunnel. Colors correspond to various temperature ranges. The tendency of the hot gases to lie between the escalator handrails, the "trench effect," is responsible for the hot surfaces on the floor of the escalator.

*Ian P. Jones, Suzanne Simcox, and Nigel S. Wilkes*
*Harwell Laboratory, Oxfordshire, United Kingdom*

On the evening of November 18, 1987, a fire erupted from an escalator tunnel into the ticket hall of London's King's Cross underground railway station, killing 31 people and injuring many others. The formal investigation of the tragedy concluded that computational work carried out on the CRAY-2 supercomputer at the U.K. Atomic Energy Authority's Harwell Laboratory in Oxfordshire revealed an important and unsuspected phenomenon, the "trench effect," that accounted for the rapid spread of the fire.[1,2]

During the formal investigation, a consensus quickly emerged as to the probable reason for the start of the fire: a match dropped by a smoker fell through a gap between the escalator treads and the side panels on the right hand side of the escalator. Such a match easily could have ignited grease underneath the escalator. The fire then would have spread above the escalator and ignited the side panels and treads.

Experts could not agree, however, on the reason the fire spread as rapidly as it did. It grew from one described by an experienced fireman as "a fire such as might be produced by a large cardboard box" to one that within as short a time as half a minute engulfed the ticket hall in flames. Such a rapid engulfment is called a flashover. A plausible reason for the flashover was the effect of the fire-retardant paint on the ceiling of the escalator shaft. Volatile gases from the paint layers could have formed a fuel-rich zone that in turn would have led to a fireball. Alternative theories argued that the role of the paint was secondary and that instead a large wood fire spread rapidly up the escalator. The proponents of the fireball theory argued that a wood fire would not spread rapidly enough to create a sudden flashover.

In December 1987, Harwell Laboratory was asked by the scientific investigation team of the Health and Safety Executive of the British government to attempt to model the aerodynamics of the flow caused by the fire during the period leading to the flashover. Harwell Laboratory has a commercial research program that is largely concerned with heat transfer and fluid flow, and the laboratory has considerable expertise in the numerical simulation of turbulent and combusting flows. The team wanted to know if any important three-dimensional flow effects could have influenced the spread of the fire. Investigating this problem was a formidable task. The model would have to simplify the problem, yet retain its essential features. An important constraint was the need to model in some detail the flow coming out of the escalator shaft into the ticket hall.

The model that was adopted included the total area of the ticketing hall and the tunnel enclosing the three Piccadilly line escalators. This tunnel was 42 meters long and had a diameter of nearly 7 meters with the escalators descending at an angle of 30° to a floor level some 16 meters below the ticketing hall floor. A simplified representation of the other openings into the ticketing hall from the Victoria line escalator and from three openings into the orbital passageway outside the booking hall were taken into account. Finite volume techniques were used; the region of interest was divided into cells, and within each cell equations were solved that represented conservation of mass, momentum, and energy. The industry-standard $k$-$e$ model of turbulence was used to model recirculation in some regions of the flow. The $k$-$e$ model has many well-known deficiencies, but it remains the industry standard for practical use. With this model, additional equations for the transport of turbulence are solved.

The HARWELL-FLOW3D software package was developed to model flows of this kind in complex three-dimensional geometries. The package uses a structured grid system that can be bent or twisted, but not torn, into a cubical grid. This permitted easy vectorization on the laboratory's CRAY-2 system. Certain cells were blocked off to represent the ticket office and the structure of the escalator. The model comprised 15,000 cells. The tight time scale for the work imposed by the formal investigation made it necessary also to simplify the physics of the problem. The various stages of the fire's development were modeled by its heat output and extent, and therefore the effects of combustion and radiation were ignored.

A clear picture emerged from the results: the hot gases in the simulations did not climb the side wall of the escalator shaft as might be expected. Instead, they lay within the "trench" between the escalator hand rails. Figure 1 shows a view looking down the escalator tunnel with the surfaces color coded according to temperature. The figure clearly demonstrates the trench effect, with the hot surfaces on the floor of the escalator.

Two physical effects give rise to this situation. The first is the chimney effect, which is the tendency of hot gases to rise and suck up air from below to replace them. On its own, the upward air flow would bend the hot gases away from the vertical. The second effect is related to the Coanda effect, which is the tendency of a jet or flame near a wall to stick to the wall because it can entrain air from all sides except that nearest the wall. The simulations indicated that the combined influence of these effects caused the hot gases to be blown flat and to flow up into the ticket hall along the floor of the escalator. This event would preheat the wood of the escalator significantly, permitting the fire to spread more rapidly than otherwise would be expected from a wood fire on an inclined surface.

Because the numerical results had to be available before the end of the formal investigation, it was essential to obtain convergence from overnight runs. The criteria were tightened to guarantee convergence of the overall implicit iteration strategy. A conservative choice of time step also was taken, with the solution advanced in increments of one-tenth of a second. This choice made the overall solution times quite long, on the order of 48 hours per run, but it also guaranteed good results after a weekend, or overnight. The interactive capability of the UNICOS operating system was extremely useful, enabling runs to be checked as they were executed. A Silicon Graphics IRIS 3130 workstation that was connected to the CRAY-2 system via an Ethernet network was a tremendous help in analyzing the results quickly and presenting the main features to a nonspecialist audience. In addition, a fluid dynamics program, RIP, from the NASA Ames Research Center was distributed between the CRAY-2 and IRIS computers. Time-consuming particle traces were calculated on the CRAY-2 system, then displayed on the IRIS workstation. This procedure allowed us to show graphically the path of the smoke.

The idea of the "trench effect" created considerable interest and controversy when the results were presented. Within a week or two after the Harwell results became available, Dougal Drysdale from Edinburgh University's fire safety unit and a team from the Health and Safety Executive had carried out experiments on one-tenth scale cardboard channels. These experiments showed the same features presented by the computer simulation.

The Health and Safety Executive had taken close interest in a computer-generated video that was shown at the inquiry. They had determined to pursue unresolved questions raised by some scientists who still believed that the paint was the primary cause of the sudden increase in energy responsible for a fireball. A large-scale test was required to settle conclusively the cause of the rapid fire spread. The Health and Safety Executive built a one-third scale model of the escalator tunnel and the ticket hall in a 25-foot shaft on the side of a hill outside their laboratory in Derbyshire. The ensuing experiments showed the features predicted by the simulations, with the flames lying flat. After these experiments, Ted Osborne, scientific adviser to London Underground, stated, "The demonstration, by computer simulation and fire modelling, of a 'trench effect' has shown that a mechanism exists for a fire within the escalator trough to develop very rapidly indeed. This is a newly discovered phenomenon, not previously identified in any previous fire situations or tests and not anticipated even in expert circles."[1]

**Computational work carried out on the CRAY-2 system revealed an important and unsuspected phenomenon that accounted for the rapid spread of the fire.**

The report from the formal investigation of the disaster also cited the importance of the computer simulations and concluded, "The computational work carried out by Harwell first drew attention to an important and unsuspected phenomenon in the form of the trench effect. In the computer simulation the airflow resulting from the fire in the trench formed by the balustrades and steps, instead of rising more or less vertically to the ceiling and flowing up the apex of the ceiling, flows up the trench. Further up the trench the flow separated into two streams; the top stream rose out of the trench, spiralled in a clockwise direction up the fascia board and across the ceiling, as viewed from the bottom of the shaft. The second stream remained in the trench and continued up the escalator shaft into the tube lines ticket hall.

"The experimental work on scale models carried out by Dr. Drysdale and the Health and Safety Executive at Buxton served to confirm the existence of a trench effect in which the flames rapidly extend up the trench until they erupt into the ticket hall, as postulated by the Harwell computational work."[1]

The report also concluded that the trench effect was "the proper scientific explanation" for the transformation of a modest escalator fire into a flashover. The CRAY-2 supercomputer service at Harwell Computing Centre played a vital role in arriving at this conclusion. Without the computing power to carry out the analyses, and the associated ability to visualize the results, much more time would have been required to perform the simulations and give the investigators the vital clue as to the reason the fire spread so rapidly.

***About the authors***

*Ian P. Jones is group leader of the Applied Mathematics Group of the Computer Science and Systems Division at the Harwell Laboratory. He received a B.Sc. degree in mathematics in 1970 and an M.Sc. degree in fluid dynamics in 1971 from Bristol University and a Ph.D. degree in aerodynamics in 1974 from the University of East Anglia.*

*Suzanne Simcox is a member of the Applied Mathematics Group of the Computer Science and Systems Division at the Harwell Laboratory. She received a B.Sc. degree in mathematics in 1985 from Exeter University.*

*Nigel S. Wilkes is the section leader of the Computational Fluid Dynamics Development Section of the Engineering Sciences Division at the Harwell Laboratory. He received a B.A. degree in mathematics in 1972 and a Ph.D. degree in mathematics in 1975 from Oxford University.*

***References***

1. Fennell, D., *Investigation into the King's Cross Underground Fire,* HMSO, 1988.
2. The King's Cross Underground Fire: Fire Dynamics and the Organisation of Safety, seminar organized by the Institution of Mechanical Engineers, June 1, 1989.

# CORPORATE REGISTER

## Demand grows for CRAY Y-MP systems

The **Ford Motor Company** has ordered a CRAY Y-MP8 computer system and peripheral equipment. Ford will install the system at its Engineering Computer Center in Dearborn, Michigan, in the third quarter of 1989. The system will be used for crash simulation, computational fluid dynamics, and other types of vehicle testing. Ford will use Cray Research's UNICOS operating system, which is based on AT&T UNIX System V. "Since 1985 we have made extensive use of our CRAY X-MP system in analyzing new vehicle designs — including the 1989 Thunderbird, Motor Trend's Car of the Year," said F. Gordon Willis, Ford Motor Company's director of product and manufacturing systems. "Our usage requirements now substantially exceed the CRAY X-MP system's capacity, and we are looking forward to having the greater power of the CRAY Y-MP system."

The **U.S. Army** has ordered a CRAY Y-MP8 supercomputer and peripheral equipment to be installed at the **Waterways Experiment Station** (WES) in Vicksburgh, Mississippi. The system will be installed in the newest WES laboratory, the Information Technology Laboratory, in the fourth quarter of 1989. The system will be used for military and civil works projects related to national defense, flood control, and navigation. WES will use the UNICOS operating system, which is based on AT&T UNIX System V. "The CRAY Y-MP system will complement the three existing Army Cray systems and will provide much needed additional supercomputing capacity for Army researchers nationwide," said N. Radhakrishnan, chief of the Information Technology Laboratory at WES.

**British Aerospace Plc**, has ordered a CRAY Y-MP2/132 supercomputer. British Aerospace, a new customer for Cray Research, will install the system in the fourth quarter of 1989. The system will be used for aerospace applications, including computational fluid dynamics, structural engineering, and impact analysis. The supercomputer also will be used by a group of companies owned by British Aerospace, including Austin Rover, which will have access to the system for automotive design.

**Renault**, a French automaker, has ordered a CRAY X-MP/18 computer system. Renault is a new customer for Cray Research. Renault will install the system at the Renault Technical Center in Rueil Malmaison, France. The company will use the supercomputer for structural analysis, combustion research, acoustic analysis, and process simulation. Renault will use Cray Research's UNICOS operating system.

**Continental AG**, a West German tire manufacturer, has ordered a CRAY X-MP EA/18se computer system. Continental AG, a new customer for Cray Research, is the fourth largest tire manufacturer in the world. Continental AG will install the system at its main computer center in Hannover, West Germany, in the fourth quarter of 1989. The company will use the Cray system for tire research and design. Continental AG will use the UNICOS operating system.

The **California Institute of Technology** has installed a CRAY X-MP/18 supercomputer at its Jet Propulsion Laboratory (JPL) in Pasadena, California. The institute is a new customer for Cray Research. The system will be used for research in areas including

Earth and planetary climatology, geodynamics, oceanography, scientific visualization and image processing, spacecraft navigation and design, radio astronomy, and solar system physics. The institute will use Cray Research's UNICOS operating system. "The acquisition of the Cray system will provide JPL and Caltech researchers with a state-of-the-art, high-performance computing environment that will facilitate solving computationally intensive problems," said A. Kukkonen, director of JPL's supercomputing project.

**Mitsubishi Electric** has ordered a CRAY Y-MP4/132 computer system to be installed in the fourth quarter of 1989 at Mitsubishi's main Large-Scale Integration Laboratory near Osaka, Japan. Mitsubishi Electric is a new customer for Cray Research. The company will use the Cray system for electronics applications. The system will run under Cray Research's UNICOS operating system.

**Phillips Petroleum Company** has ordered a CRAY Y-MP2/116 supercomputer and peripheral equipment. Phillips will install the system in the fourth quarter of 1989 at its main computer center in Bartlesville, Oklahoma. The CRAY Y-MP system replaces a CRAY-1 system that was installed in 1984. The new system will be used for petroleum exploration and production activities. Phillips plans to migrate to Cray Research's UNICOS operating system by the third quarter of next year.

**PRAKLA SEISMOS AG**, a West German geophysical services company, has installed a CRAY X-MP/18 computer system at its computer center in Hannover, West Germany. PRAKLA SEISMOS, a new customer for Cray Research, will run its GEOSYS software package under the UNICOS 5.0 operating system. "GEOSYS operation depends on supercomputer performance capabilities, as provided by CRAY X-MP systems. The code was especially designed for flexible operation on modern vector computers and provides processing capabilities necessary for two- and three-dimensional onshore and marine seismic data," said Emil Hinrichs, PRAKLA SEISMOS computer center manager.

General Atomics, which operates the **San Diego Supercomputer Center** (SDSC) at the University of California, San Diego for the National Science Foundation, has ordered a CRAY Y-MP supercomputer. The system, which will be installed in the fourth quarter of 1989 at SDSC, will replace a CRAY X-MP/48 system that was installed in the first quarter of 1986. The system will be used for scientific computing in computational chemistry, high-performance graphics, physics, and electrical engineering. SDSC also is implementing a world-class scientific visualization center as an adjunct to the CRAY Y-MP system. "The increased capabilities of the CRAY Y-MP system, together with the new scientific visualization center, will significantly enhance the ability of SDSC's users to lead the world in research accomplishments and high technology products," said Sid Karin, director of SDSC and General Atomics vice president for advanced computing.

**Texas A&M University**, a new customer for Cray Research, has installed a CRAY Y-MP2/116 system at its main campus at College Station, Texas. The system will be applied to research in areas including computational chemistry, agriculture, biochemistry, meteorology, and oceanography, as well as nuclear, mechanical, and aerospace engineering. Cray Research will provide Texas A&M with grants over five years to further supercomputer research in areas such as artificial intelligence and biochemical modeling related to agriculture. "Texas A&M is ranked eighth among all U.S. universities by the National Science Foundation in terms of total research funding, with $230 million. The availability of supercomputer resources is critical to the future of our research programs," said William H. Mobley, Texas A&M president. John J. Dinkel, associate provost for computing and information systems, added, "Acquisition of the CRAY Y-MP system is part of Texas A&M's commitment to an enhanced research and academic computing environment."

The **NASA Ames Research Center's Central Computing Facility**, Moffett Field, California, has installed a CRAY Y-MP8/832 supercomputer. Sterling Software's ZeroOne System Division was the prime contractor for the order. The system will be applied to research in areas including physics, fluid dynamics, and chemistry. The new CRAY Y-MP8/832 system will replace a CRAY X-MP/48 system that was installed in February 1985.

## Tools support UNICOS migration

To help users migrate from the Cray operating system COS to Cray Research's UNICOS operating system, Cray Research provides a set of software migration tools. Release 5.0 of the migration tools package now is available. This release includes new tools and updates previous ones to take advantage of the most current versions of the two operating systems, UNICOS 5.0 and COS 1.17. The tools can be used with CRAY Y-MP, CRAY X-MP EA, CRAY X-MP, and CRAY-1 computer systems. New software tools in this release include

### COS

- ☐ The $DMYCLB library, which contains entry points from $DMYLIB and the libCOS library available with UNICOS 5.0 ($DMYCLB is either taken from the COS release tape or produced under UNICOS by **putpdt**)
- ☐ The COSPAR utility, which reports on the number of times that selected keywords for specific JCL commands are used

### UNICOS

- ☐ The gc program, which allows users to access archived COS datasets
- ☐ The administrative utility **gcld**, a shell script that provides a global interface to the following routines, which collectively create the gc data base:
  - — **gcbcd**, which reads the COS BCD and BVCD datasets and adds corresponding catalog entries to the gc data base
  - — **gcmap**, which adds COS to UNICOS user mappings to the data base
  - — **gcmcat**, which creates the data base structure
  - — **gcpcat**, which copies the COS BCD and BVCD datasets from the COS archive catalog tape

### Guest operating system (GOS)

- ☐ The UNICOS-to-COS Software Link, which is used while running GOS to facilitate intramemory data transfer between COS and UNICOS

Migration notes for system administrators and users are included with release 5.0 of the migration tools.

### Migration training

Training for migrating from COS to UNICOS is available either at Cray Research's software training center in Mendota Heights, Minnesota, or at region or customer sites. Courses are offered in migration planning, migration tools usage, COS/GOS operation and internals, and application conversion.

For more information about Cray Research's UNICOS migration tools, contact the nearest Cray Research sales office.

UNICOS

# APPLICATIONS UPDATE

## CPlex: solving large linear programming problems

CPlex is a linear optimization system from CPlex Optimization, Inc., for solving large linear programming problems on all Cray systems. CPlex is written in C and vectorized to take advantage of Cray system architecture. Typical applications include refinery operation optimization, airline scheduling, and financial portfolio optimization.

CPlex efficiently handles highly degenerate problems and problems for which an initial feasible solution is difficult. State-of-the-art factorization routines, an automatic crash procedure, and a piecewise linear phase I approach contribute to overall solution efficiency. Innovative pricing routines all but eliminate the need to hand-tune the algorithm to obtain best performance.

There is virtually no limit to the number of variables, constraint equations, or nonzero matrix elements when run on Cray computer systems. To date, problems with more than two million nonzero coefficients have been solved.

CPlex features include

- ☐ Ability to read and write industry standard MPS files or accept input directly from a user-supplied matrix generator
- ☐ Interactive entry and editing of linear programming problems
- ☐ Efficient restart capability from previous optimal solution or from a user-selected advanced basis
- ☐ Objective function and right-hand side ranging information upon request
- ☐ Output directed to either terminal or logfile
- ☐ Simple command structure with online help
- ☐ Easy customization to meet user-specific requirements

CPlex optimization routines and functions are integrated easily into user-developed applications. The routines are available optionally in linkable object module form, callable from Fortran, C, and other standard languages. Common applications include links to front-end modeling languages and matrix generators, downstream report writers, or graphics packages.

For more information about using the CPlex package on Cray computer systems, contact Todd A. Lowe, CPlex Optimization, Inc., 7710-T Cherry Park, Suite 124, Houston, TX, 77095; telephone: (713) 550-9763; or contact John Gregory, Cray Research, Inc., 1333 Northland Drive, Mendota Heights, MN, 55120; telephone: (612) 681-3634.

## NSPCG linear solver available on Cray systems

NSPCG (Nonsymmetric Preconditioned Conjugate Gradient) is a software package that solves the linear system $Au = b$ by various iterative methods. The package is written in Fortran 77. It is structured for vectorization and runs on Cray computer systems under the UNICOS, COS, and CTSS operating systems.

The coefficient matrix $A$ can be passed in one of several matrix data storage schemes. These sparse data formats allow matrices with a wide range of structures, from highly structured ones, such as those with all nonzeros along a relatively small number of diagonals, such as a band matrix, to completely unstructured sparse matrices. Alternatively, the package allows the user to call the accelerators directly with user-supplied routines for performing certain matrix operations; one can use the data format from an application program and not be required to copy the matrix into one of the package formats. This is particularly advantageous when memory space is limited.

The main entry point into the package is through the subroutine call. The various methods are accessed by using a particular naming convention for the first two parameters that in turn select a preconditioner, an accelerator, and a data storage scheme. Some of the basic preconditioners that are available are Jacobi, incomplete LU decomposition, and symmetric successive over-relaxation, as well as block preconditioners. The user can select from a large collection of accelerators such as conjugate gradient, Chebyshev, generalized minimal residual, biconjugate gradient squared, and many others. The package is modular so that almost any accelerator can be used with almost any preconditioner. One of the main objectives in the development of the package was to provide a common modular structure for research on iterative methods.

The degree of vectorization attained in a particular application depends on many factors, including the particular iterative method, the underlying structure of the matrix, the data storage format, the ordering of the equations, and the architecture of the computer. The NSPCG package permits several sparse matrix data structures that are suitable for regularly or irregularly structured matrices, various orderings for enhanced vectorization, and different vectorizing philosophies, such as register-to-register on Cray systems.

The NSPCG package has been applied successfully to matrix problems arising from many applications:

- ☐ Oil reservoir simulation problems from J. S. Nolen and Associates in Houston, Texas
- ☐ Finite difference analysis used in convection and diffusion problems
- ☐ Finite element analysis used in fluid flow and structural analysis problems
- ☐ Spectral methods for Laplace's equation

The NSPCG package also has been used in conjunction with Newton's method for nonlinear finite element applications to Navier-Stokes equations.

The NSPCG package was developed as part of the ITPACK project directed by David M. Young and David R. Kincaid at the Center for Numerical Analysis at the University of Texas at Austin. More detailed information about the package can be found in the NSPCG user's guide Version 1.0 written by Thomas C. Oppe, Wayne Joubert, and David R. Kincaid and available through the Center for Numerical Analysis.

For more information about NSPCG on Cray computer systems, contact David Kincaid, Center for Numerical Analysis, University of Texas at Austin, RLM Hall 13.150, Austin, TX, 78713-8510; telephone: (512) 471-1242; or contact John Gregory, Cray Research, Inc., 1333 Northland Drive, Mendota Heights, MN, 55120; telephone: (612) 681-3634.

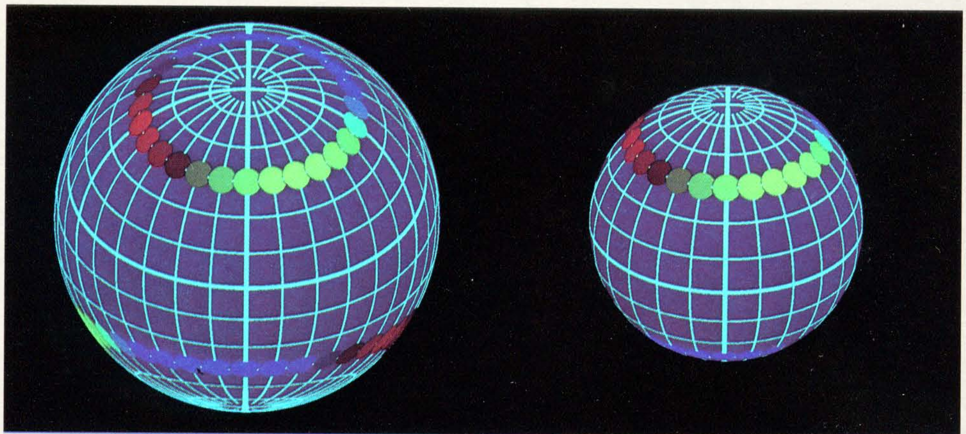

Visualization of relativistic light deflection. The images show a polar grid on a neutron star with two corotating "hot spots." The colors correspond to different phases of the star's rotation. The left image accounts for light deflection, and the right image does not.

The images of the polar grid were computed in two steps. First, any pair of neighboring points was checked to see if a line of the coordinate grid passed between them. If a line did pass between them, then the point closest to the line was set. In the second step, the lines were given a "physical" width, and the points on the screen that correspond to a point on the star touched by such a line were set. These broadened lines give the picture a much more "realistic" appearance; they also contain some information about the physics of the situation. The resulting images are shown above; they have a resolution of 1000-by-1000 points, and each required about three minutes of CPU time on the CRAY-2 system.

The researchers' work on cosmic x-ray sources extends beyond x-ray pulsars and includes models of accretion disks around cataclysmic variables, which are similar to x-ray pulsars but contain white dwarf stars rather than neutron stars. They also are modeling atomic systems to refine the theory of matter in strong magnetic fields.

## Cray Blitz tops international grand master

Chess anyone? How about a match with a supercomputer?

Earlier this year, 12 of the best human chess players from the 1989 Dutch chess championship challenged the 12 best computer chess programs from around the world to a match. Among the invited players was Cray Research's favorite chess team, including Robert Hyatt of the University of Alabama at Birmingham, Harry Nelson of Lawrence Livermore National Laboratory, and the Cray Blitz program, along with a CRAY Y-MP8/832 supercomputer. The event was sponsored by the AVRO TV network in the Netherlands and taped for local and international distribution.

Although the human players were expected to sweep the tournament, the Cray Blitz program and two other computer programs won their matches. The most significant computer win was the Cray Blitz victory over an International Grand Master who had a rating of over 2500. This class of chess player is quite rare — only 100 to 200 exist worldwide. The other top-rated chess computers, Deep Thought and Hi-Tech, both from Carnegie-Mellon University in Pittsburgh, are special-purpose machines that contain an entire program on one VLSI chip. Deep Thought lost, while High Tech drew its match.

The Cray Blitz program, composed of 30,000 lines of Fortran and 10,000 lines of assembly code, searches over 200,000 positions per second using a highly sophisticated parallel searching algorithm. The large memory and eight processors of the CRAY Y-MP system provided a speed edge against the grand master.

"It was an amazing performance," said Hyatt, who has been competing in computer chess tournaments with the Cray Blitz program for nine years. "My goal is to play in one human tournament per year and see how well we compete against people. Our goal has always been to become the best chess competitor in the world — human or nonhuman." The Blitz team captured the North American championship title for three consecutive years — 1982 through 1984, and the world championship for six consecutive years — 1983 through 1988. The program has been ranked among the top three computer chess programs since it tied for first place with Belle from AT&T Bell Laboratories in the 1982 Annual Computer Chess Tournament.

"We've been talking about using a Cray HSX high-speed communications channel to link two to four Cray systems. We are already using all eight processors of the CRAY Y-MP system — it's just a matter of distributing part of the chess program over the network using the HSX channel." But Hyatt admits, no matter how much computer power is dedicated to a chess game, an element of luck always is involved. He adds, "I wouldn't mind betting money on one of these games — but not my life!"

## Aerospace application sustains 2.1 GFLOPS

Since Cray Research introduced the eight-processor CRAY Y-MP system in 1988, researchers have been attaining GFLOPS performance with increasing frequency. Recently, Cray Research analysts benchmarked a large electromagnetics problem for an aerospace industry customer that achieved a sustained performance of more than 2.1 GFLOPS. Analysts used a technique called lower-upper (LU) decomposition, a matrix method for solving a very large series of linear equations. The LU decomposition method commonly is applied to problems that are important in the design of radar evading aircraft as well as in medical imaging technology.

The 40,000-by-40,000 complex number problem comprised 3.2 billion 64-bit words of data. Over two trillion bytes of data were transferred between the CRAY Y-MP system and Cray Research's DS-40 disk subsystem, with I/O overhead accounting for only 3 percent of the total run time. The solution to the problem was based on a matrix multiply, and an algorithm based on the BLAS 3 (Basic Linear Algebra Subroutines) kernels was used. Cray Research analysts adapted a routine that uses a block-oriented method from Argonne National Laboratory's LAPACK test called CGETRF to run out-of-core. To achieve this, they divided the matrix into slabs. The eight-processor Cray system worked on slab pairs, and the matrix then was decomposed from left to right, one slab at a time. To compute a new leading slab, all preceding slabs were brought into memory one at a time for computation. Three slab-sized buffers were used in the code to allow for asynchronous I/O. To enhance the performance of the complex matrix multiplication, a kernel requiring a reduced number of operations was used.

The ability to solve complex electromagnetics calculations with the LU decomposition method enables engineers to design advanced tactical aircraft and other vehicles that have electromagnetic properties important to detection or the avoidance of detection. The same technology is applied to medical imaging, allowing researchers to determine safe radiation exposure levels.

# GALLERY

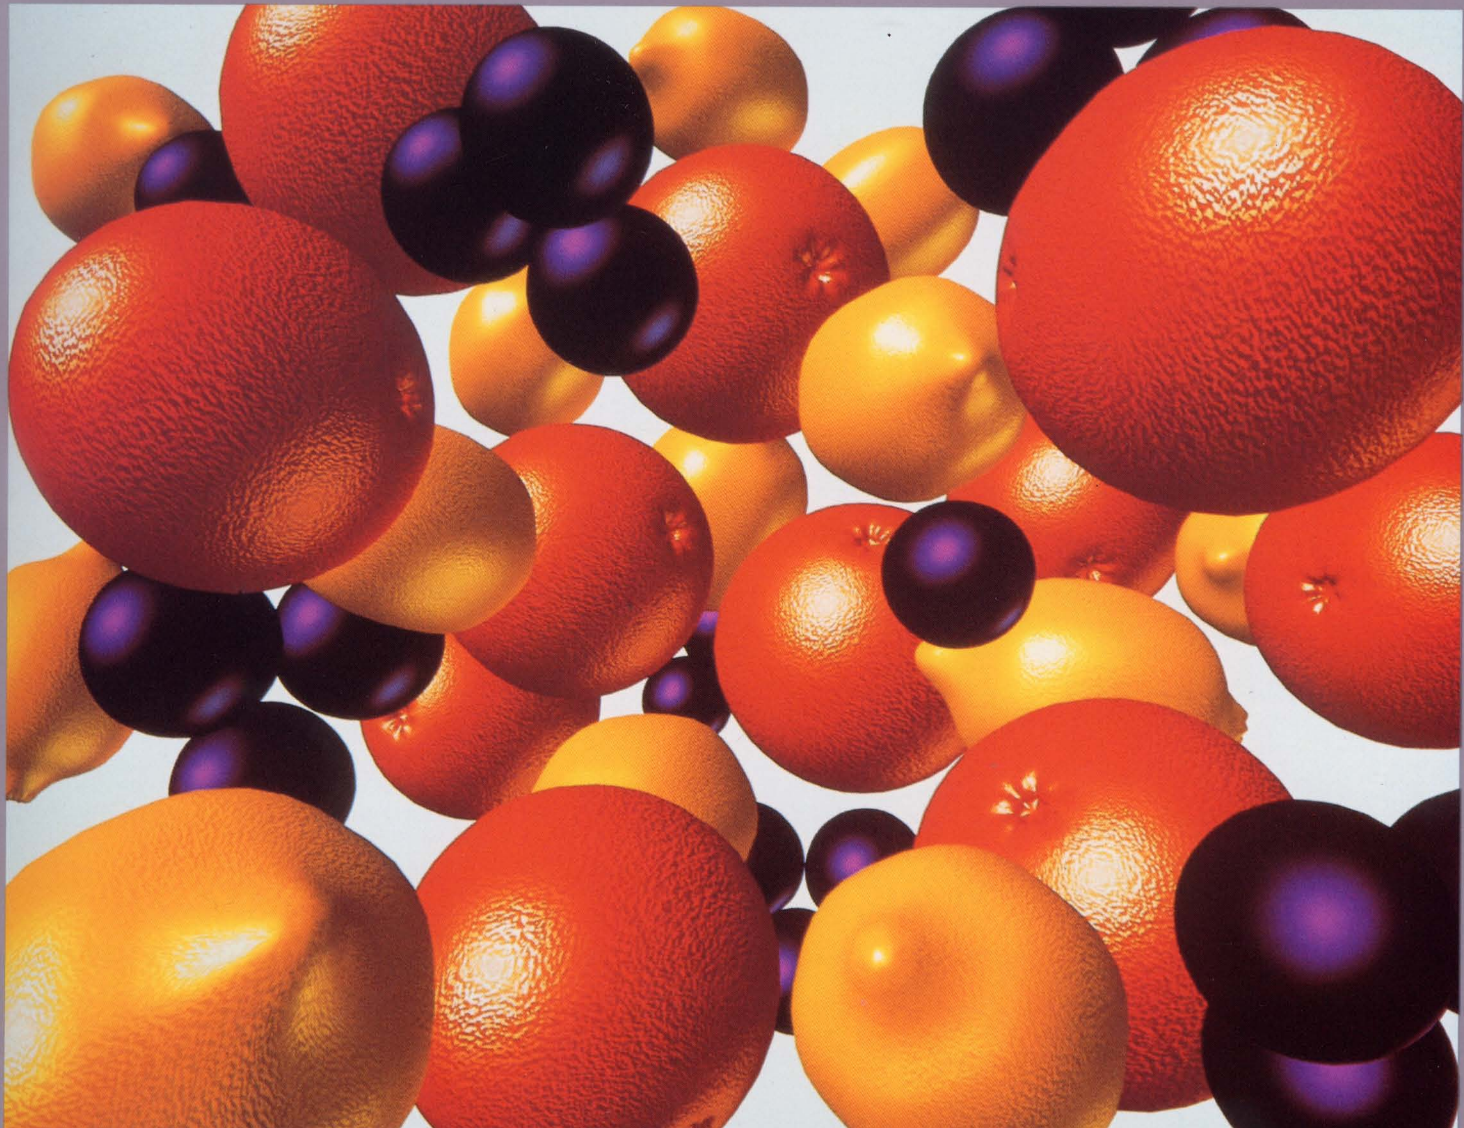

Oranges, lemons, and grapes defy gravity in "Fruit" by Shelly Lake. The image, a still from a television cereal commercial, was produced on a CRAY X-MP/22 system. The lighting effects, reflectivity, color, and texture for the 5120-by-3072 image were computed in about three minutes.

CRAY CHANNELS welcomes Gallery submissions. Please send submissions to the address inside the front cover.