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

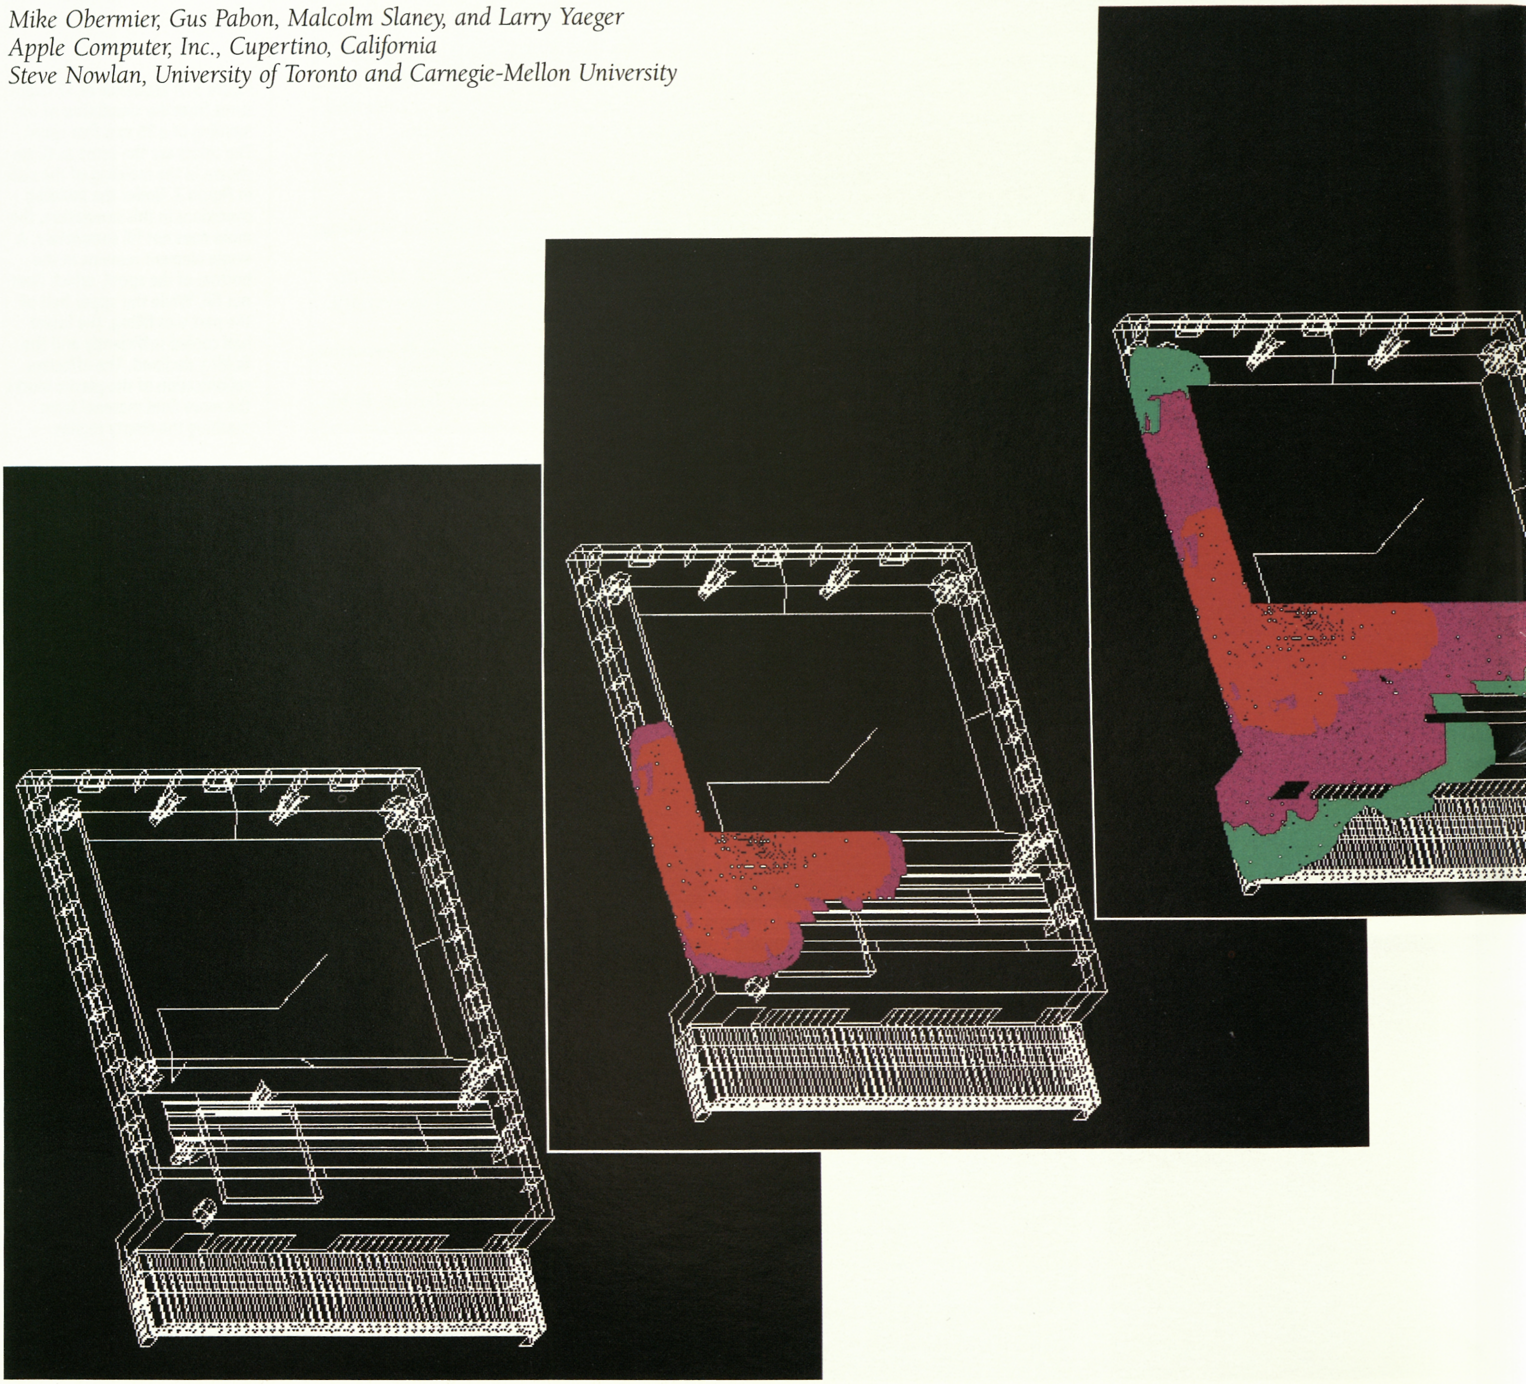

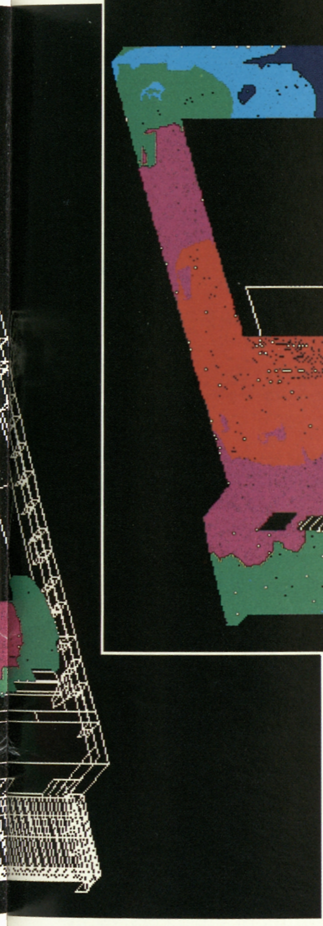

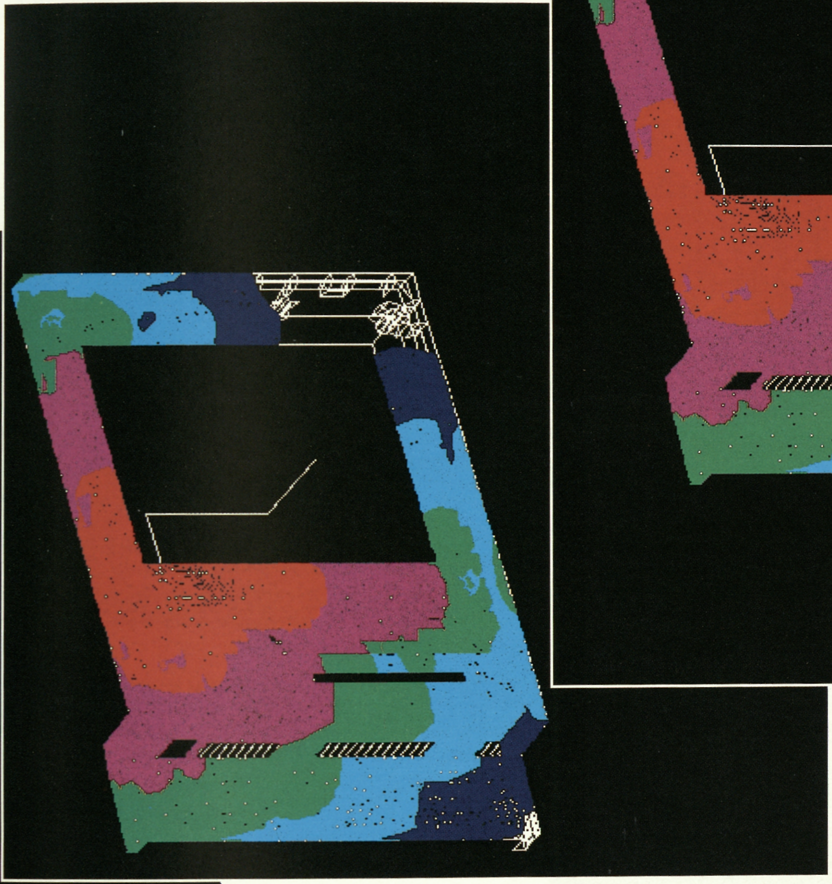

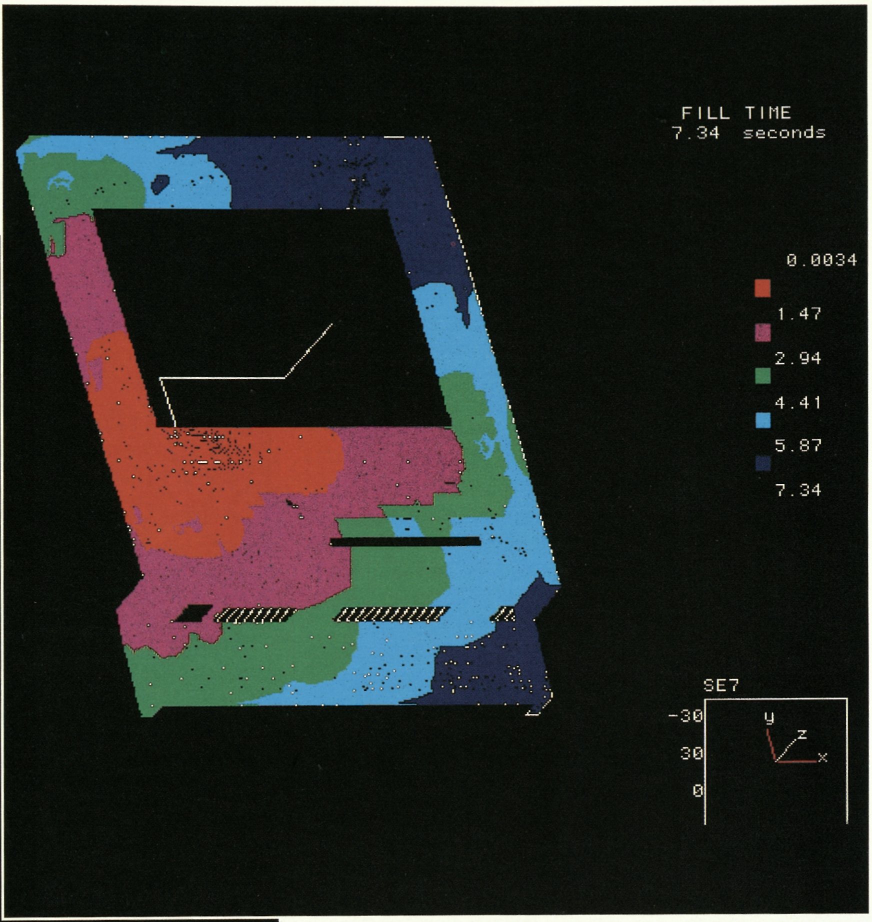

Injection molding simulation for the Macintosh SE computer bezel. The left image shows the model geometry. The following four images show, left to right, the extent of the flow field within the mold at various times up to completion.

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.

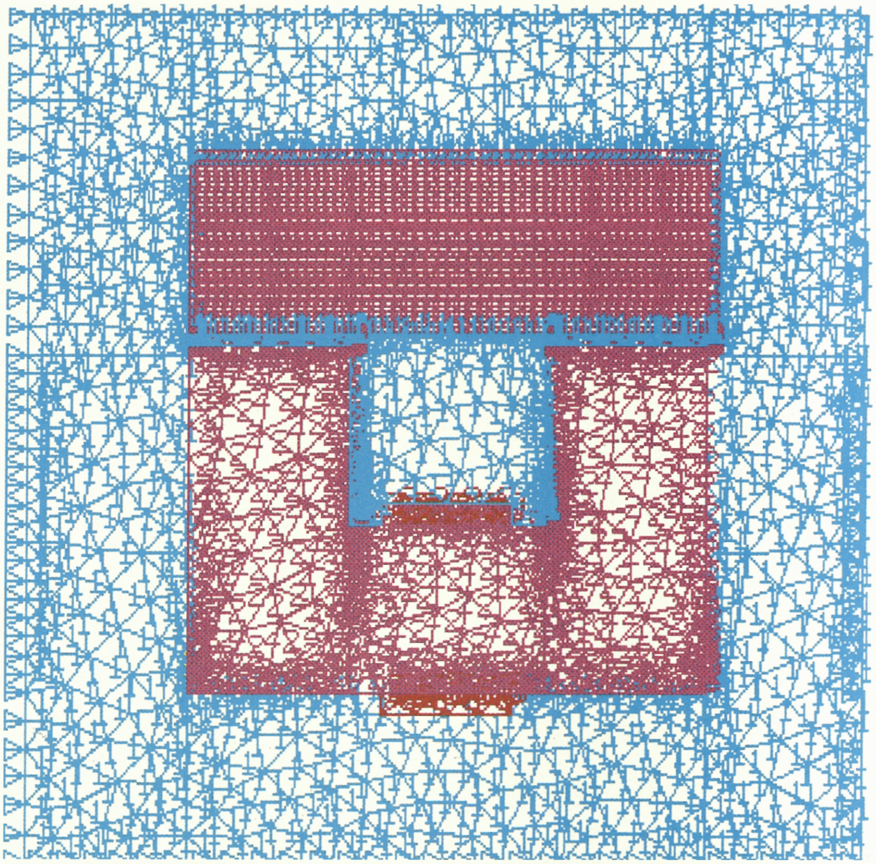

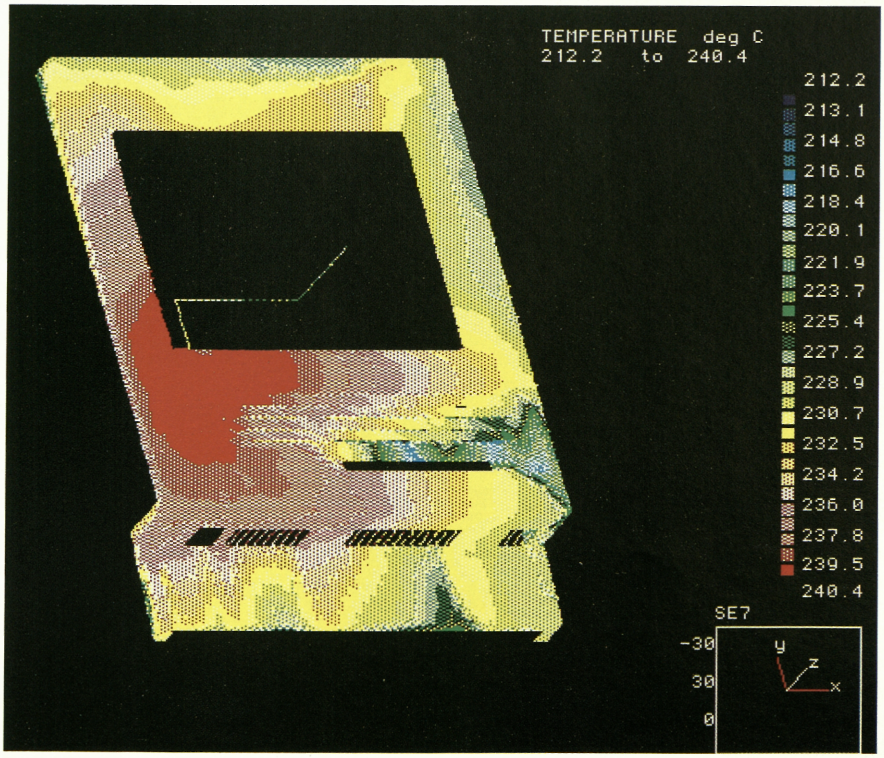

Temperature distribution within the bezel mold at filling completion.

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

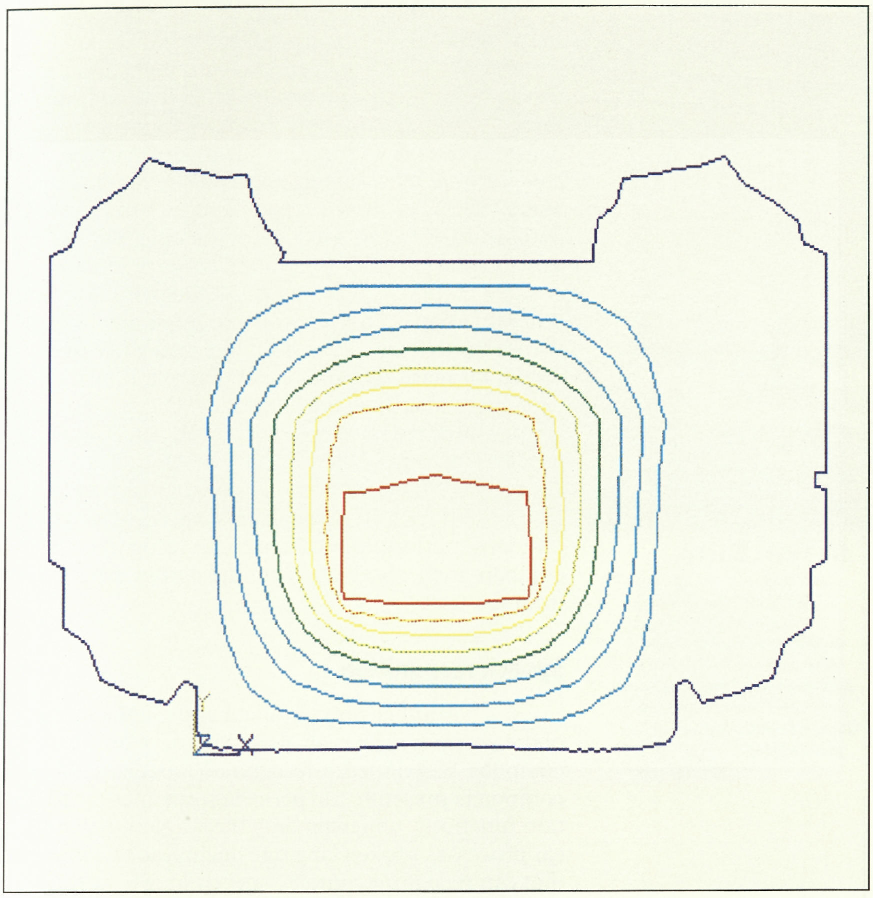

Transformer core simulation for a new power supply design. The images show the finite-element mesh structure of the core (above left), lines of equipotential magnetic flux (above), and flux density magnitude plots (below right).

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 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.

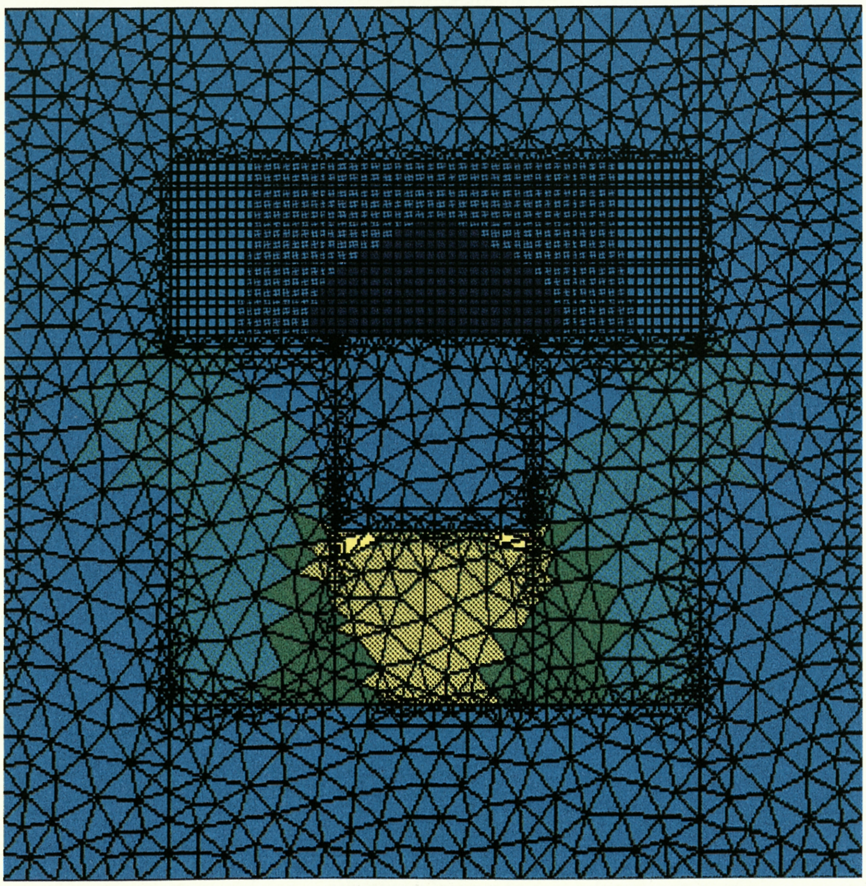

Artificial neural networks

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

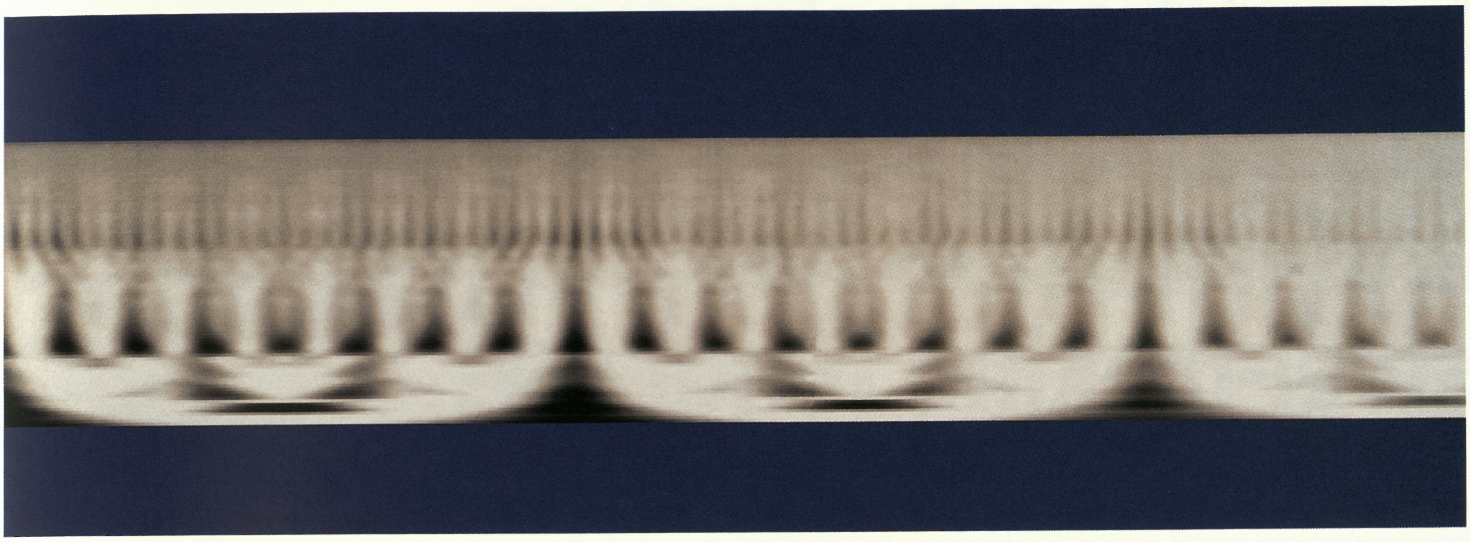

One frame of a correlagram. The dark horizontal bands represent high energy in a particular frequency region, or formant, and the vertical lines represent common periodicities across all channels of the pitch. This frame represents the sound of the vowel "u" produced by a male speaker.

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.