A CRAY RESEARCH, INC. PUBLICATION

CRAY CHANNELS

Winter 1986

FEATURE ARTICLES:

DEPARTMENTS:

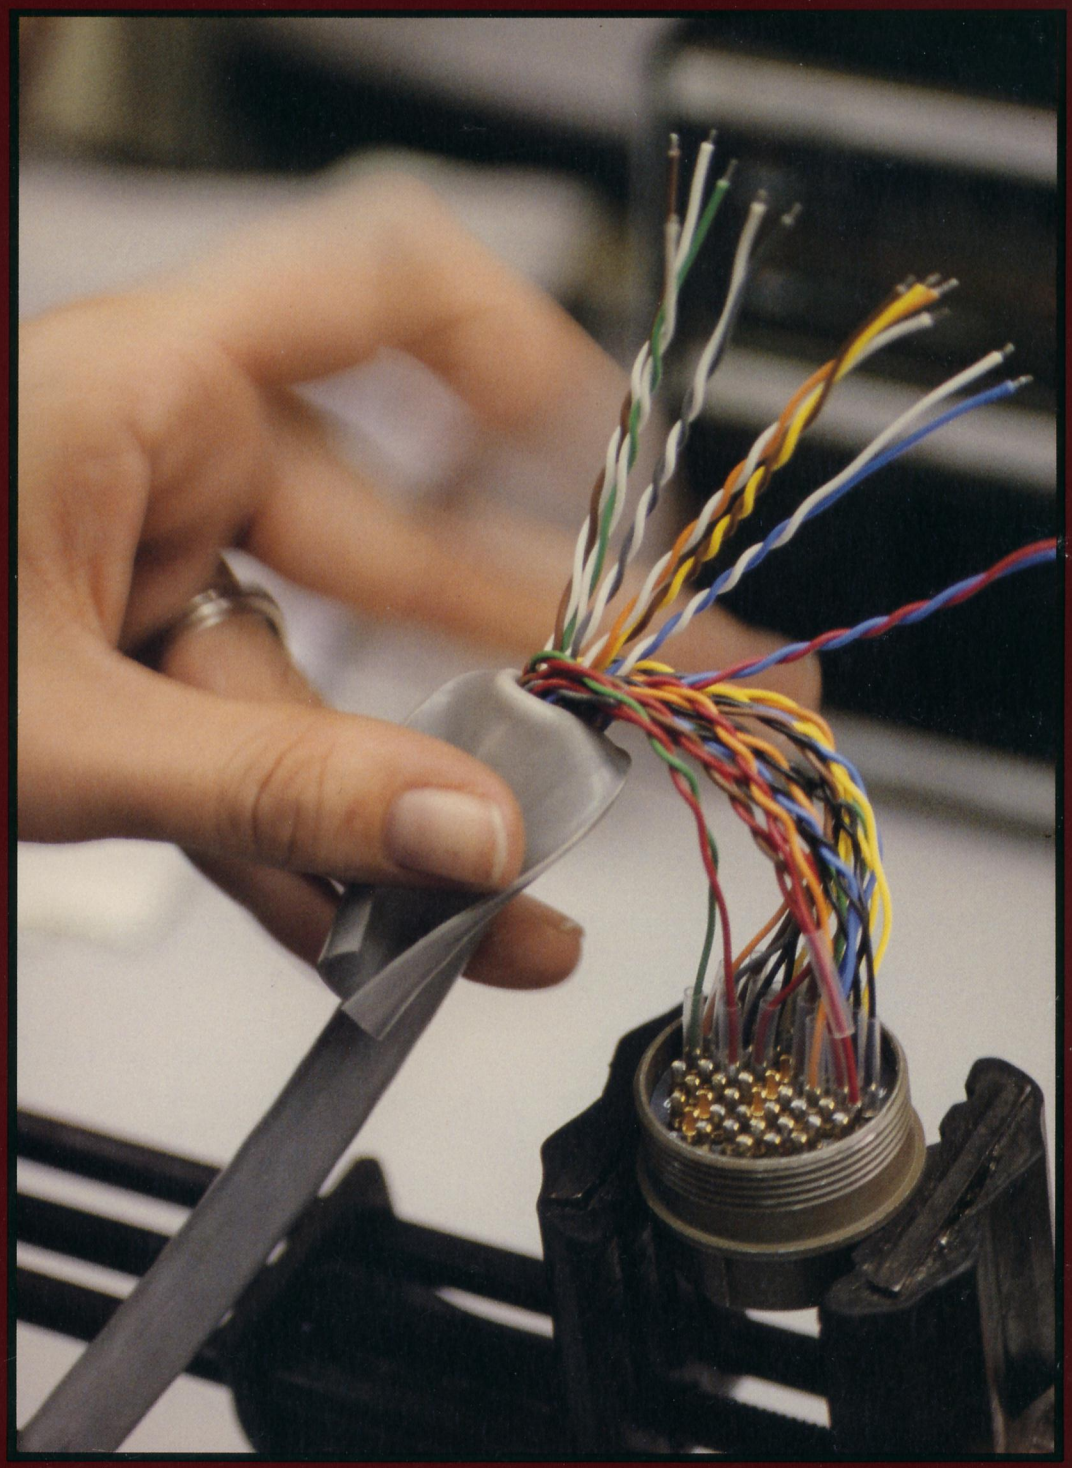

IN THIS ISSUE

For years chemists have relied on traditional experiments to discover the structure, stability, and physical properties of molecules. Sometimes it was impossible or impractical to search for a chemical means of creating a new compound for study. With the availability of supercomputers, however, chemists can determine with striking accuracy the characteristics of new molecules without having to create them. The computational power of a CRAY system gives researchers the ability to model the complex interactions among thousands of electrons and atomic nuclei in a proposed compound.

This issue of CRAY CHANNELS also looks at a very humanistic application of supercomputer technology: medical imaging. Our regular departments discuss new application software packages and how prehistoric climate modeling can provide practical information to aid in petroleum exploration.

This issue initiates what we hope will be a regular feature in CRAY CHANNELS: how to maximize the performance of CRAY computer systems. The premier article talks about the CRAY X-MP hardware performance monitor, a valuable source of information on program performance that can be used to optimize programs of all sizes. We welcome your input — feature ideas or articles written by users — to help us tailor the information to meet your needs.

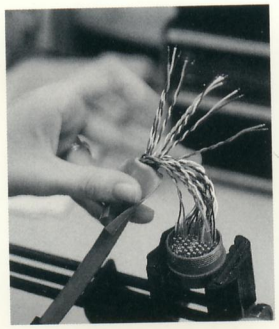

On the cover is a logic cable under construction. Each of the 55 color-coded wires in the cable is being soldered into a corresponding receptacle in the connector shown in the clamp. When complete, the cable will be used to connect a CRAY mainframe to a peripheral device such as a disk drive or Solid-state Storage Device. Cabling is performed at Cray Research's manufacturing facilities in Chippewa Falls, Wisconsin.

CRAY CHANNELS

A Cray Research, Inc. publication

Winter 1986

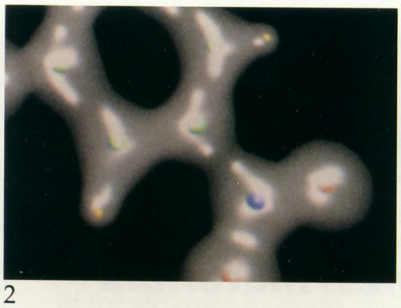

2

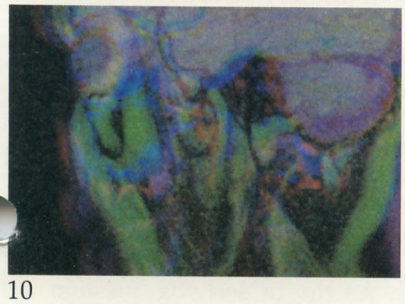

10

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 subscription requests, feature story ideas and news items to **CRAY CHANNELS** at Cray Research, Inc., 608 Second Avenue South, Minneapolis, MN 55402.

Volume 7, Number 4
Winter 1986

Editorial Staff
Tina Bonetti
Jean Egerman
Anne Hazelroth
Ken Jopp
Carol Smith
Jack Thompson

Design
Cindy Erickson
Eric Hanson

FEATURES

DEPARTMENTS

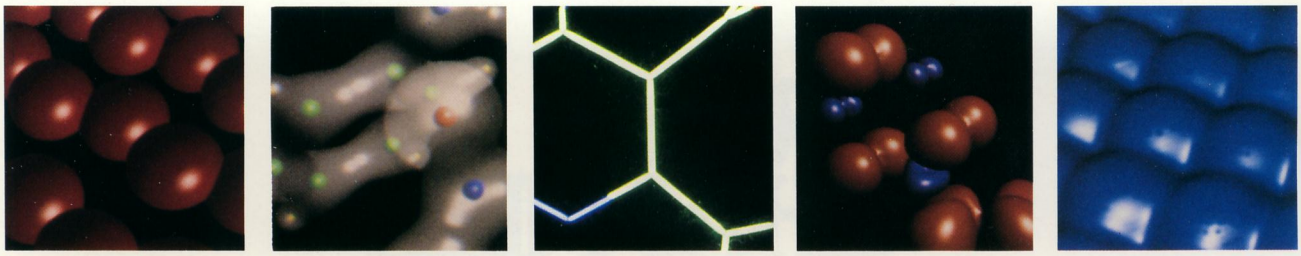

Computational chemistry by supercomputer

Steps toward computer-aided design of new materials

Erich Wimmer, Cray Research, Inc.
Arthur J. Freeman and Chong-Long Fu, Northwestern University, Evanston, Illinois
Bernard Delley, University of Neuchatel, Switzerland

What stone was to the Stone Age, iron to the Iron Age, and bronze to the Bronze Age, a diversity of materials are to us today. Advanced societies rely on everything from steel for cars and light alloys for aircraft to fertilizers for agriculture and semiconductors for electronic devices. Many of our most commonly used materials are synthetic products such as certain medical drugs, polymers for clothing, and colors used in painting and printing. Although much of this material wealth had chance origins, supercomputers are making it possible to consciously guide the design and development of new materials.

The classical approach

In ancient societies, most new compounds and alloys were discovered accidentally. Later, empirical methods were employed to improve on chance. These methods involved a tedious trial-and-error process. Scientists made and tested a series of related compounds, then dismissed most of them because they lacked the desired properties. Using this procedure, scientists developed many empirical and semi-quantitative rules to describe the relationships between chemical composition and material properties.

Although scientists have learned to make general predictions about the properties of new compounds, the predictions are not always reliable and are restricted to closely-related compounds. Progress has been made with this methodology, yet the underlying question has persisted: What determines the properties of a material?

The role of the electron

Scientists describe and classify materials according to properties such as optical appearance, electrical conductivity, magnetic properties, melting point, hardness, and chemical reactivity. Theoretically, if we knew the mathematical equations that determine these properties, we could predict the properties for any material by simply solving the equations. But to find the equations we must first discover the physical laws that govern the properties of materials.

A breakthrough in our understanding of these physical laws occurred in the late nineteenth and early twentieth centuries, when scientists discovered that the electrons in a material play a crucial role in determining its physical properties. Electrons create and break chemical bonds, and their distribution determines the chemically important properties of a material, including:

- the shapes of molecules,
- the arrangements of atoms in solids,
- the forces between atoms, and
- the electrical properties of materials (whether they are insulators, semiconductors, or metals).

Even the complexity of biomacromolecules in living organisms is finally governed by the electrons of the atoms forming these molecules. Clearly, an understanding of the physics governing electrons is necessary to understand, describe, and predict material properties.

Compared with the spatial extension of electrons, atomic nuclei are very small. Although they contain most of the mass of an atom, for chemical purposes nuclei are usually considered merely as points of a certain mass and positive charge.

The quantum mechanical approach

In the mid-1920s, several scientists, primarily the Austrian physicist Erwin Schroedinger and the German physicist Werner Heisenberg, formulated the theory of quantum mechanics. This theory allows a quantitative prediction of the distribution and energetics of small particles such as electrons. The fundamental equation of quantum mechanics, called the Schroedinger equation, contains the key to understanding the properties of matter. It gives a complete description of the electrons in atoms, molecules, and solids. Solving Schroedinger's equation for a given arrangement of atoms reveals the distribution of the electons and the corresponding energy of the molecule or solid. From there, one can determine things such as the arrangement of the atoms in a molecule or solid in their ground state configuration (the arrangement with the lowest total energy) and how to calculate the forces binding the atoms together (Figure 1). The ground state configuration is generally the most stable form of the molecule and thus the form it naturally assumes.

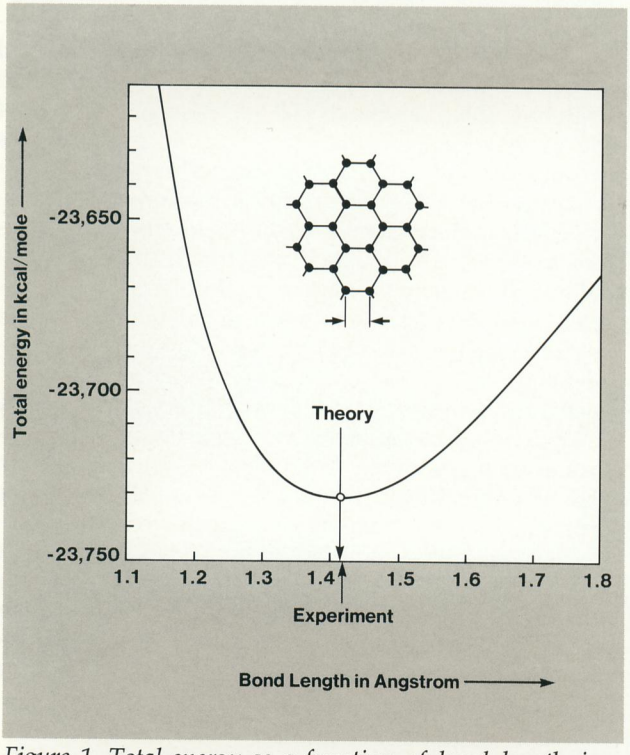

Figure 1. Total energy as a function of bond length in a single hexagonal layer of carbon atoms (graphite). The lowest total energy corresponds to the natural bond length. The slope of the curve is related to the force between the atoms.

To determine the ground state configuration of a molecule or solid, researchers choose an arbitrary atomic configuration and calculate the total energy. They then change the geometry, recalculate the total energy, and compare it with the previous value. If the new value is lower, this procedure is repeated until the structure of lowest total energy is found. This structure is taken as the molecule or solid's actual structure.

From there, one can proceed to pull the molecule or solid apart, step by step, recalculating the total energy for each step until the molecule or solid falls apart. This series of calculations reveals how strongly the molecule or solid holds together and how much energy it takes to break it. Alternatively, one can start with isolated atoms or molecules and bring them together through a series of calculations to form a new molecule, thereby revealing the new molecule's stability. These two procedures — making and breaking chemical bonds and forming new compounds — are at the heart of chemistry.

In practice, the analysis process can be made more efficient by calculating not only the total energy at each point, but also the forces acting on the atoms. The forces will indicate the appropriate direction one needs to move to obtain the ground state configuration.

The difficulties

All of this sounds very straightforward, but several attendant problems exist. First, even small molecules have many electrons. Furthermore, each electron interacts with all the other electrons and with all the atomic nuclei in a molecule or solid. For example, a water molecule, a relatively small molecule, contains 10 electrons and three atomic nuclei. In the metal tungsten, used as a filament in light bulbs, each individual atom carries 74 electrons, and it takes about 10^{23} atoms to form even a small piece of metal.

A second difficulty arises because the motion of electrons is not governed by the laws of classical mechanics, which have been applied so successfully in the engineering disciplines. The motion of electrons is governed by the laws of quantum mechanics, sometimes also including the theory of relativity. At this point one might resign in despair: 74 electrons per atom, 10^{23} atoms in a solid, and all treated by relativistic quantum mechanics — this might be too much, even for a CRAY supercomputer!

Does this mean that researchers must abandon their dreams of calculating the properties of complex materials and limit their research to calculating the already-known properties of small molecules? Fortunately, it is possible to tackle complex chemical problems by combining creativity with the power of supercomputers.

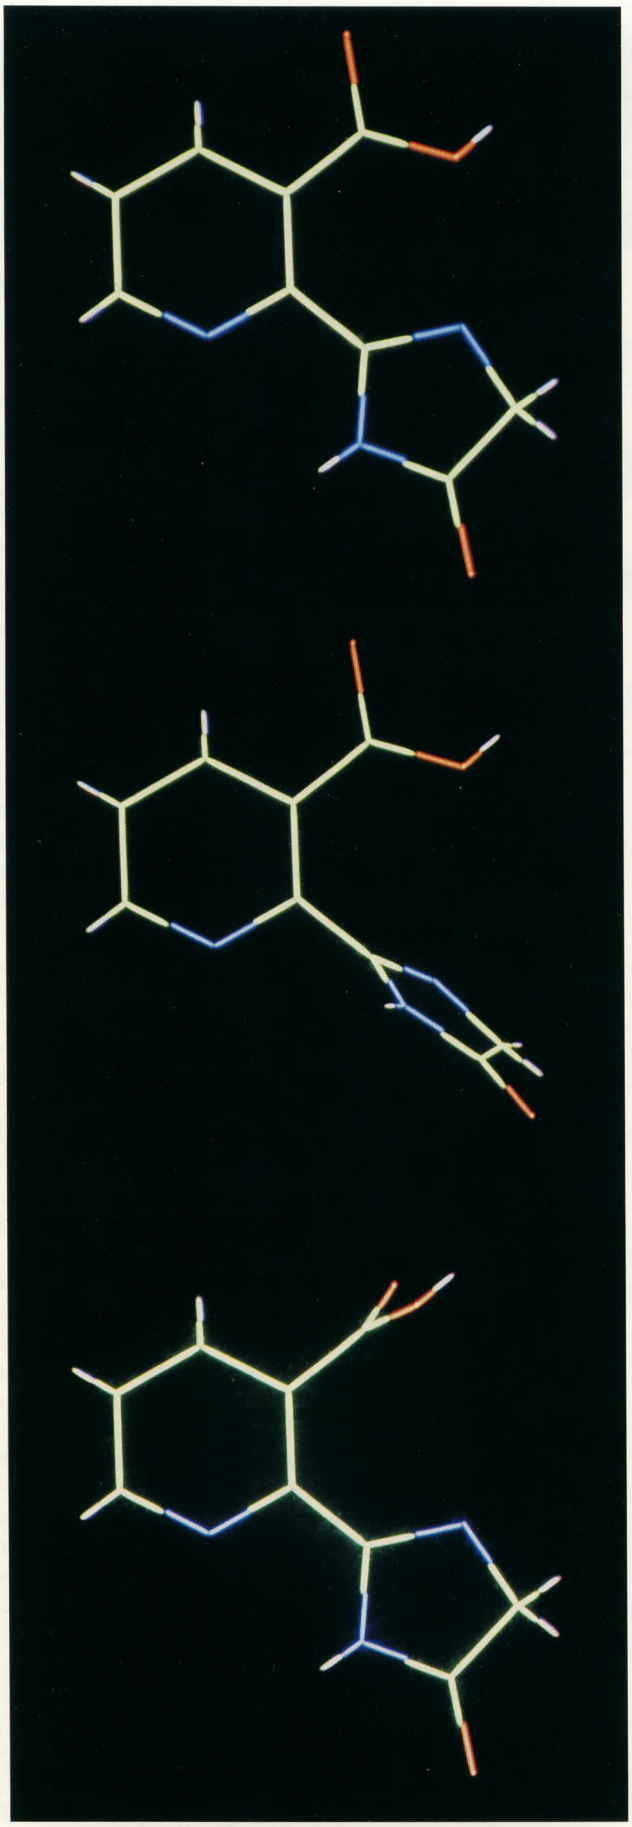

Figure 2. Possible three-dimensional conformations of a molecule that acts as a herbicide.

Since the formulation of quantum mechanics in the mid-1920s, scientists have been struggling with these challenges, and many brilliant ideas have been formulated. Without going into detail on these theories, this article looks at some examples that demonstrate how far researchers have come in attempting to calculate and predict the properties of materials. The problems discussed below were solved using all-electron local density methods.

Herbicide conformation

The molecule shown in Figure 2 is a herbicide in which the six-member ring and the five-member ring could be oriented in many ways with respect to each other. The two rings could be coplanar, the planes of the two rings could be orthogonal, or they could have any angle with respect to each other. The -COOH group, which itself defines a plane, can also be oriented in an arbitrary way with respect to the six-member ring. In other words, there are many possible orientations for the molecule's ground state configuration. In such a situation, chemists want to know the molecule's geometric structure (what it looks like) and how much energy is needed to deform it.

Aside from appeasing scientific curiosity, there is a practical reason to answer these questions. The herbicide shown blocks a critical part of an enzyme in plants via a key-lock mechanism. To fit into the active site in the enzyme, the various groups of the molecule have to be in a certain orientation with respect to each other.

To design another herbicide with similar enzyme-blocking effects, but perhaps with different residual groups or with some modification in the six-member ring, one must be able to predict the conformation of the new molecule. One way this could be done is through difficult and expensive experiments by actually synthesizing and testing new molecules. Alternatively, the molecule could be designed with an interactive graphics device and its conformation then calculated. If the newly designed molecule assumes a different structure, it is very likely to be inactive and can be ruled out.

As part of the visiting scientists program sponsored by Cray Research's applications department, the conformational problem for the molecule shown in Figure 2 was solved using a CRAY X-MP/48. In the process, about 40 different geometries were investigated. Each geometry required about 750 seconds on one processor to obtain the distribution of the electrons and the corresponding total energy. As shown in the figure, two critical rotational angles were varied independently in the search for the configuration with the lowest total energy.

The result is quite remarkable and would have been difficult to predict. The angle between the five-member ring and the six-member ring is 20.2 degrees and the angle between the -COOH group and the six-member ring is 69.1 degrees. The calculations produced inter-

esting numbers, but perhaps more importantly, they provided insight into the bonding mechanism that causes this particular geometric arrangement. It turns out that the delocalized electrons in the six-member ring couple with the electrons in the five-member ring to form a bond that wants to keep the two rings coplanar. No such mechanism is found between the -COOH group and the six-member ring. For this reason it is harder to rotate the five-member ring than the -COOH group. However, if the five-member ring and the -COOH group were coplanar, the nitrogen atoms (shown in blue in Figure 2) of the five-member ring would come too close to the oxygen atoms (red) in the -COOH group and they would repel each other. Therefore, both groups tend to rotate to increase this distance. However, because it is harder to rotate the five-member ring than it is the -COOH group, the latter rotates more than the former. This explains the peculiar arrangement.

Electron density and spin

As mentioned earlier, the behavior of electrons is described by quantum mechanics, which predicts certain properties of electrons that do not correspond to any properties of the macroscopic world. One such property is the so-called electron spin; electrons can only have "spin up" or "spin down." In many molecules, such as the water molecule, the number of electrons with spin up equals the number with spin down, and the spatial distribution of the spin up electrons is the same as that of the spin down electrons. In certain molecules, however, there are more electrons of one type than of the other.

The molecule shown in Figure 3 has one more electron with spin up than with spin down. Chemists call such an electron "unpaired." Molecules with unpaired electrons tend to show great reactivity, because the unpaired electron wants to find a partner. Schematically, the unpaired electron is represented by a dot over one of the nitrogen atoms, but this is a relatively arbitrary convention to represent the electron's position. A quantum mechanical calculation performed on a CRAY X-MP/48 revealed the true situation: the unpaired electron is spread over several atoms. The difference in the distribution between the spin up and the spin down electrons is shown in Figure 4 (page 6).

Surface chemistry

The understanding of metallic surfaces is of enormous technological importance. Wear and tear of pistons, for example, is determined to a large extent by processes taking place on metallic surfaces, particularly the interaction between the surface and a lubricant. Another important process is corrosion, a primarily electrochemical process that annually costs consumers and industry millions of dollars in property damage. Corrosion is similarly governed by the reactivity of metallic surfaces. Catalysis, another important chemical effect, is used to create many fundamental chemicals, such as ammonia for fertilizers. Many industrial catalytic processes use metallic surfaces as catalysts.

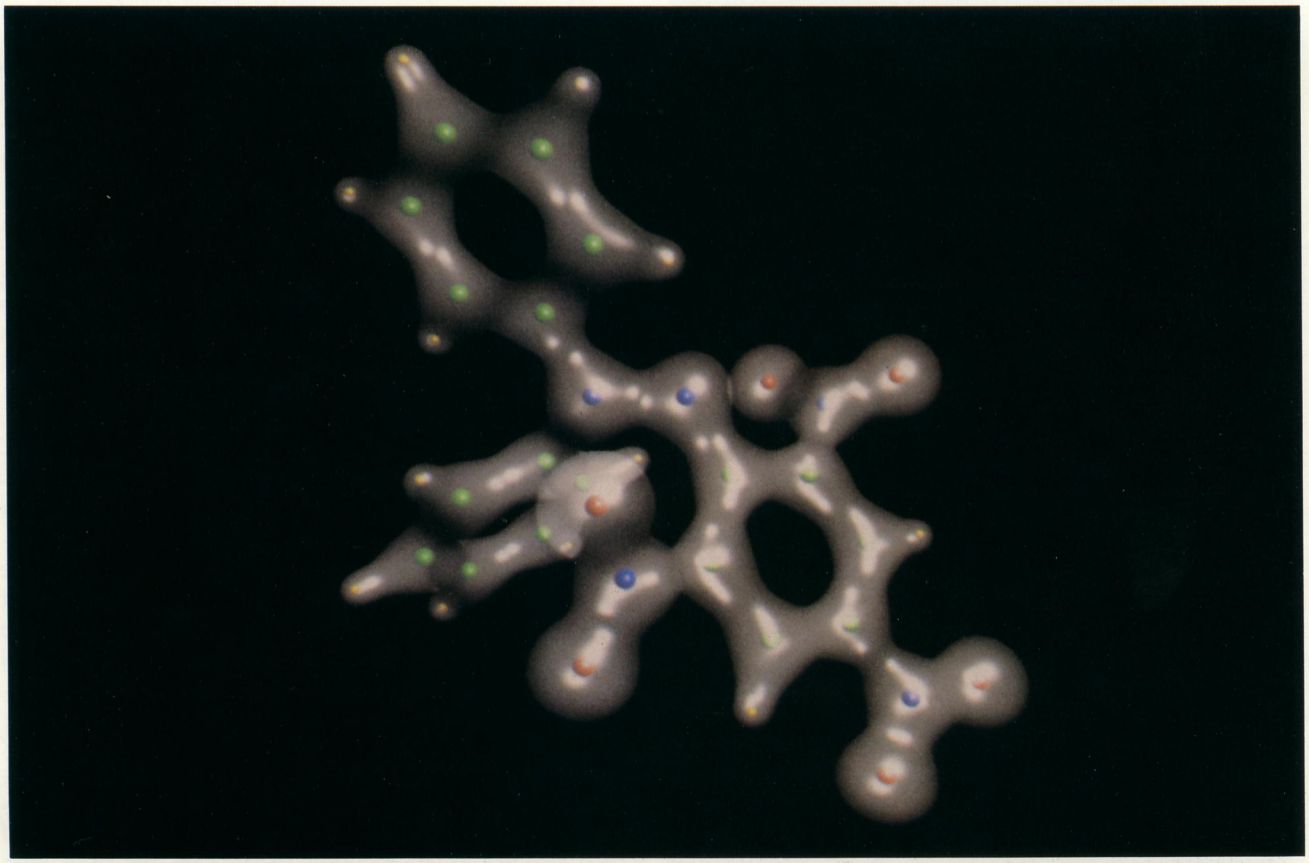

Figure 3. Distribution of the electronic charge in a molecule called DPPH obtained from first-principles density functional theory. The three-dimensional surface connects points in space having a certain value of electron density.

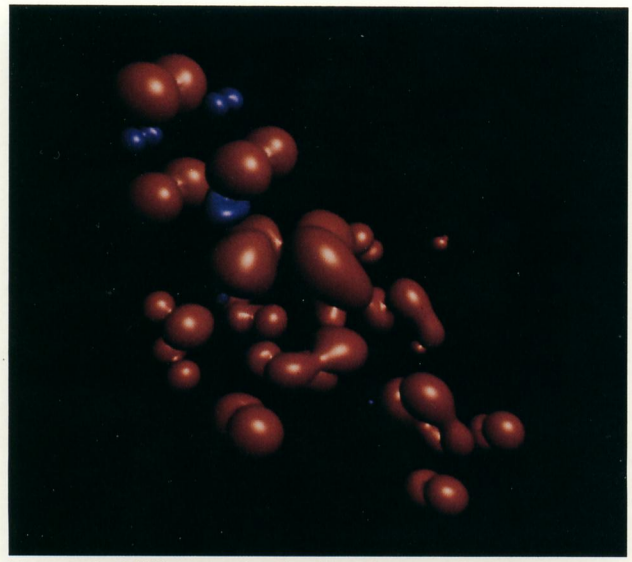

Figure 4. Difference in the distribution of electrons with "spin up" and "spin down" obtained from self-consistent local spin density functional theory for the DPPH molecule. The three-dimensional surface connects points in space where the difference has a certain positive (red) and negative (blue) value.

The understanding of these important processes on an atomic level is still in its infancy. However, a better understanding of catalytic processes will certainly lead to the design of better catalysts, providing higher yields and perhaps less pollution of our environment. Another field where knowledge of surface chemistry is of key importance is the manufacture of electronic devices, or chips. As our industries move from silicon technology to compound semiconductors such as gallium arsenide, an understanding of catalytic processes becomes even more important. Many fundamental questions remain to be answered.

Atomic displacement

The metal tungsten is one of the most studied but perhaps least understood metallic surfaces. Experimental evidence indicates that the surface is not simply a cut-off of the bulk (Figure 5a), where the position of the atoms is known fairly accurately. Rather, it appears that the atoms at the surface rearrange themselves. Several possible rearrangements have been suggested, as shown in Figure 5.

As already discussed, calculating the total energy for a variety of different arrangements will reveal the correct one. Using a CRAY X-MP/24, scientists from the group at Northwestern University were able to calculate the answer by solving the quantum mechanical problem for the tungsten surface. The results show that a vertical displacement (Figure 5b) does not do the job; as the atoms are displaced, the total energy increases. On the other hand, if the atoms are displaced laterally (Figure 5c), the total energy first increases very slightly and then starts to drop; the system goes into its equilibrium geometry at a lateral displacement of 0.18 angstroms. As it turns out, recent experimental data agree with this theoretical result. Researchers can feel confident that the structure shown in Figure 5c represents the real structure of this surface.

All electrons of the tungsten atoms were included in these calculations and the actual value for the total energy at the equilibrium is 4,395,270.51 electron volts (the preferred energy unit of solid-state scientists). By contrast, the energy differences involved in the rearrangement of the atoms is only in the range of 0.01 electron volts. It is significant that this kind of numerical precision can be used to obtain reliable results. (Another check of the reliability of the theoretical/computational approach is the calculation of the distance between tungsten atoms in the bulk of the material. The calculation gives agreement to within one-half percent of the measured value.)

As the example of the herbicide molecule already discussed, the calculation of this "reconstruction" of the tungsten surface not only gives reliable numbers, but also provides insight and understanding. The analysis of the results clearly demonstrates why a vertical displacement cannot lead to a lowering of the total energy, whereas the lateral displacement of the atoms is accompanied by pronounced changes in the bonding between the atoms at the surface and the subsurface layer. This change in bonding favors the lateral displacement.

Ultra-thin metal layers

It has long been recognized that coating a metal such as iron or steel with a thin layer of another metal, such as chromium, can protect the surface of the original material. The shiny chromium parts on cars are examples; without these protective overlayers, the bare steel would soon rust. Again, the chemist is interested in underlying questions: How thin can an overlayer be to modify a metal surface? Can new materials be made by exploiting the properties of ultra-thin layers?

In recent years, these questions have received considerable attention, but preparing and characterizing these overlayer systems experimentally can be quite tedious. However, the power of supercomputers allows the quantum mechanical calculation of such systems.

For example, examination of bulk chromium metal reveals that some atoms have a majority of spin up electrons, whereas neighboring atoms have a majority of spin down electrons. On average, however, throughout the piece of metal the two types of electrons have equal representation. Is this still the situation if just one atomic layer of chromium atoms is deposited onto another metal such as gold, which has an equal number of spin up and spin down electrons?

Recently, this situation has been simulated on CRAY computers, again using quantum mechanical calculations, with interesting results. If just one layer of chromium atoms is deposited on a gold surface, then all of the chromium atoms have more spin up than spin down electrons. The average over the whole system is

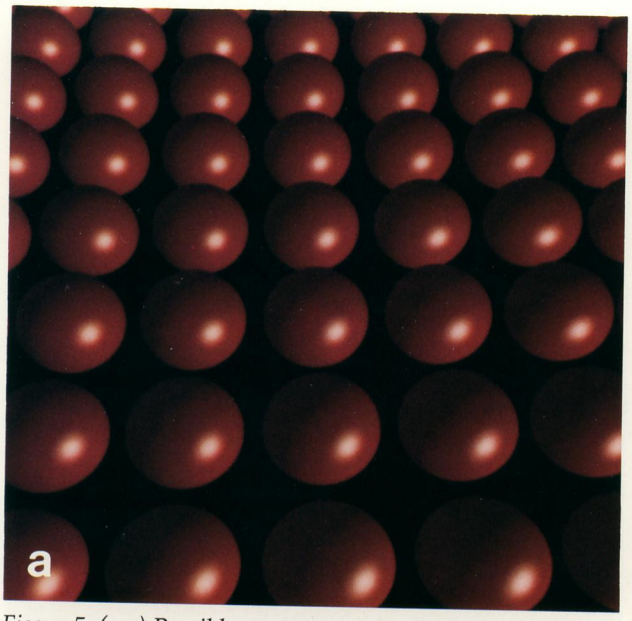

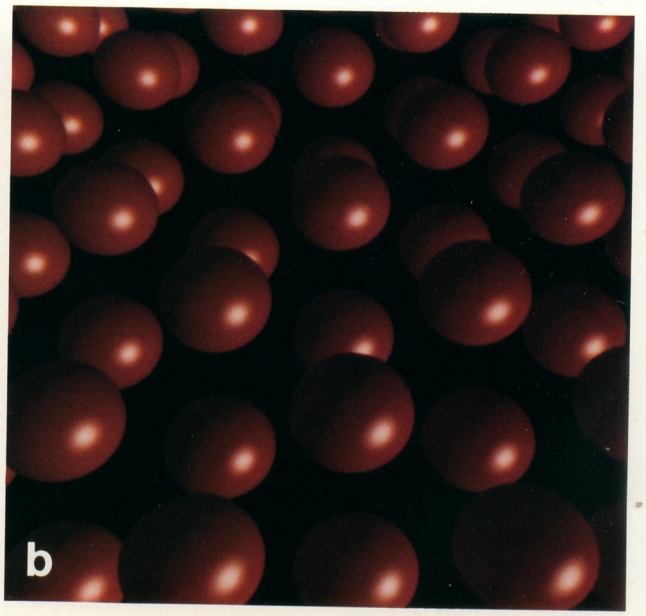

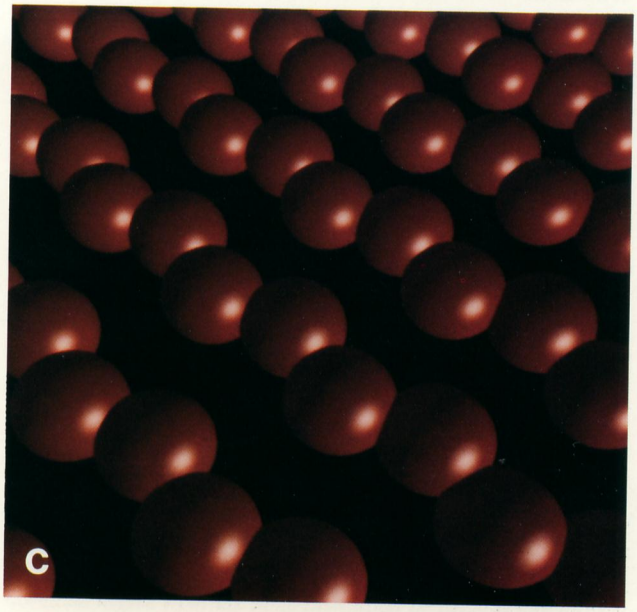

Figure 5. (a-c) Possible arrangements of the atoms on a tungsten surface; sophisticated quantum mechanical calculations allow a prediction as to the natural arrangement. (d) The line drawing of the so-called energy band structure indicates the computational complexity.

then no longer zero, and the result is a large magnetic moment which is six times larger than that for a chromium atom in the bulk. Significantly, the gold atoms adjacent to the chromium atoms also become magnetic. This is the first time that gold atoms have been determined to become magnetic.

Conveniently, one can change the gold substrate to silver merely by changing a few lines in the program input. In some sense, this ability to "transmute" metals is a realization of the ancient alchemical dream. Using computers, researchers can build materials and study and modify their properties before they are synthesized. Today, chemists can use computation to pioneer a completely new field: the computer-aided design of materials.

Catalytic promotion

As mentioned previously, the understanding of catalysts on an atomic level is in its infancy. The next example typifies some of the remaining problems.

Suppose, for example, one wants to produce a hydrocarbon such as gasoline from coal. As a first step, one would let the coal react with oxygen to form carbon monoxide (CO), a very stable compound. In a later step, one would want to initiate a reaction between CO and hydrogen to make hydrocarbons, with water as a by-product. But this last reaction does not proceed by itself, because CO is too passive to react spontaneously with hydrogen. Thus, a catalyst that activates CO is needed. Certain metal surfaces such as that of nickel can perform this function.

The addition of potassium to the reaction further enhances the activation and dissociation of CO. In other

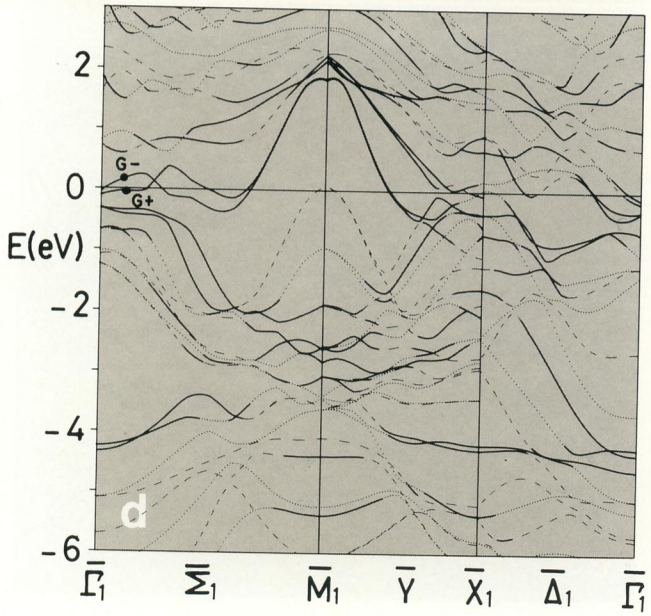

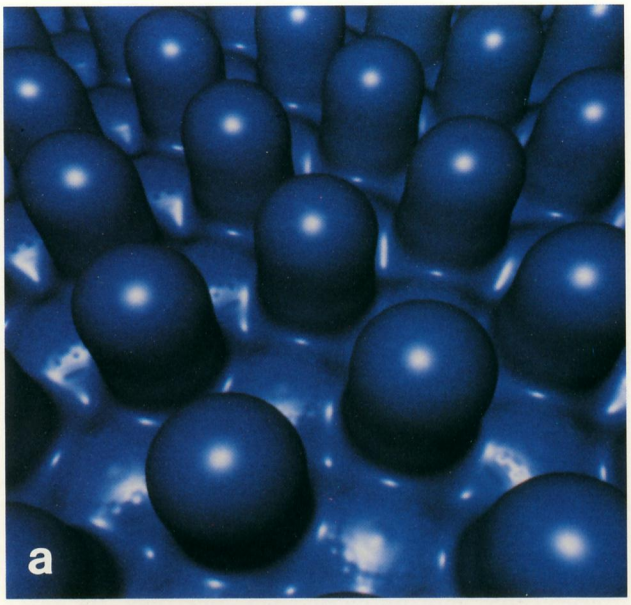

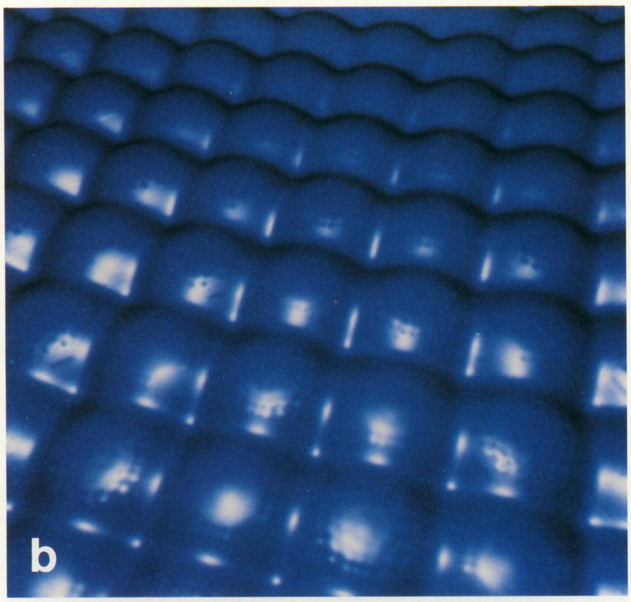

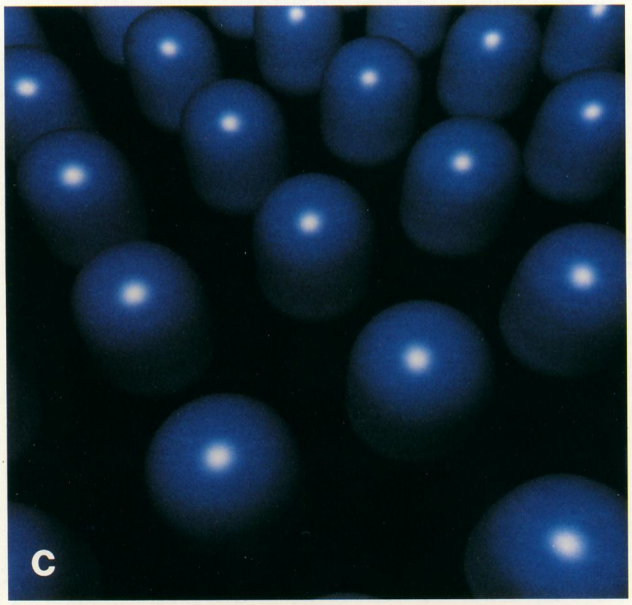

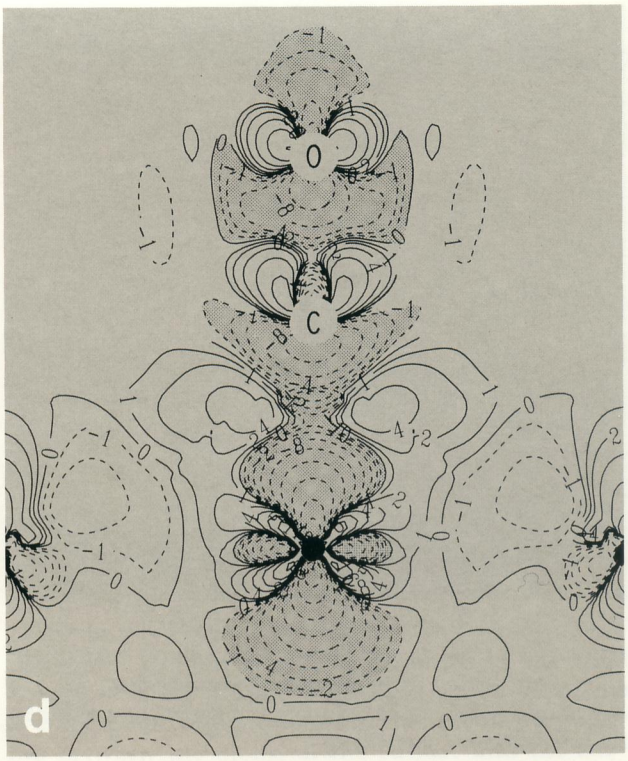

Figure 6. (a-c) Three-dimensional surfaces created by connecting points in space having a certain value of electron density: (a) CO molecules chemisorbed on a nickel surface, (b) a clean nickel surface, and (c) isolated CO molecules. (d) The difference between (a) and (b + c), showing how electrons redistribute themselves when CO is chemisorbed on nickel.

words, potassium acts as a catalytic promoter. Although these examples of catalysis and promotion have been observed, their mechanisms were not understood until recently. Elucidating the mechanisms is important because such understanding can suggest methods to improve or modify the catalytic surface systematically.

Using a CRAY X-MP computer, researchers recently made quantum mechanical calculations of a nickel surface with chemisorbed CO molecules with and without coadsorbed potassium atoms. Such a theoretical/computational study presents the possibility of investigating intermediate steps that are experimentally inaccessible but important nonetheless in understanding the catalytic mechanism.

In the case of CO molecules chemisorbed in a dense overlayer on a nickel surface, an outstanding problem was quantifying the interaction between the molecules and the metal surface. Experimentally, it would have been rather difficult to study these interactions separately. But by performing simulations on a CRAY X-MP computer, it became relatively easy.

The simulations were performed by first forming a layer of CO molecules spaced as if they were chemisorbed on the nickel surface, but omitting the nickel substrate. This model revealed the interactions between the CO molecules. The bare nickel surface was

similarly studied. In a further step, the CO molecules were put down on the nickel surface and the changes of the electronic distribution were monitored by taking the total charge density of the system CO on nickel (Figure 6a) and subtracting the charge density of the bare nickel surface (Figure 6b). The charge density of the isolated CO molecules (Figure 6c) was then subtracted. This difference revealed how the electrons redistribute themselves when CO is chemisorbed on nickel (Figure 6d). This redistribution shows the bonding change, which enables the theoretical chemist to deduce that the bonding between the carbon and oxygen atoms is weakened when CO is chemisorbed on the nickel surface. As shown in Figure 6d, there is less electronic charge between the carbon and oxygen atoms, meaning the carbon-oxygen bond is weakened and thus easier to break. This demonstrates that the metal surface catalyzes the dissociation of the CO molecule. If potassium is now added, the carbon-oxygen bonding is further weakened. Thus, potassium acts as a catalytic promoter. In contrast, the coadsorbtion of sulfur was found to act as a catalytic poisoner.

Clearly, the surfaces of real catalysts are complex rather than flat and many chemical processes are taking place simultaneously. Computer simulation provides an in-depth picture and understanding of crucial aspects of the catalytic process described by a simplified geometric model. Knowing the details of this idealized situation provides new insight that can trigger the imagination.

Only the beginning

Computational chemistry opens many doors for pure science as well as for industry. Chemists are now in a position to make quantum mechanical calculations for systems containing complicated organic molecules with unpaired electrons. Chemists can also accurately predict the arrangement of atoms on the surfaces of metals and explain the mechanism that drives these arrangements. Additionally, researchers can now form new systems such as an atomic monolayer of chromium atoms on a gold surface and predict the resulting properties, such as magnetism. For certain catalytic processes, theoretical/computational simulations based solely on the laws of quantum mechanics provide a detailed understanding of the catalytic mechanism. Insights can also be obtained as to the mechanism of catalytic promotion by the addition of potassium atoms.

Using sophisticated theoretical/computational approaches and the power of CRAY supercomputers, chemistry is entering an era where the simulation of real-world chemical systems is a reality. Theoretical chemistry is no longer confined to the study of small systems where most of the answers have been found through experiment.

Chemists can now tackle complex systems that are difficult or even impossible to study experimentally. Molecules, solids, and surfaces can be constructed on graphics workstations. By performing complex calculations on a CRAY supercomputer linked to a graphics workstation, chemists can visualize and study the electrons as they form and break bonds. Intermediate steps that are never accessible experimentally can also be studied in detail, and thus a control over chemical processes can be developed that would otherwise be impossible. In this way, the computer-aided design of new materials is becoming a reality that will revolutionize the way chemists develop new materials to satisfy human needs.

Applications for computer-aided design of materials include ceramics and synthetic polymers for engineering, semiconductors for electronic devices, color pigments for dyes and paints, fertilizers and herbicides for agriculture, and drugs for medicine. The ability to visualize and model biomacromolecules by supercomputer will deepen our understanding of living systems and enable us to perform genetic engineering computationally. Equipped with supercomputers, computational chemists will explore new dimensions in the world of chemistry. But these efforts will always go hand-in-hand with experiment. No matter how sophisticated a simulation model might be, it can never capture the full reality of nature. Nonetheless, the deeper and more accurate our insight into chemical processes, the more we will respect and marvel at the complexity of nature, particularly within the realm of biological systems. Exciting times are ahead of us! □

Acknowledgement

The three-dimensional computer-generated images in this article were computed on a CRAY X-MP/48 using the OASIS image synthesis package. OASIS was written, and the images were rendered, by Gray Lorig of Cray Research's applications department.

About the authors

Erich Wimmer, a native of Austria, received his engineering degree from the Technical University of Vienna in 1974 and his doctorate degree from the same institution in 1977. Before joining Cray Research's applications department in 1985, he was associate professor of theoretical chemistry at the University of Vienna. Wimmer's research interests center on the electronic structure of surfaces, atomic overlayers on surfaces, and molecules chemisorbed on surfaces.

Arthur J. Freeman, a native of Poland, is Morrison professor of physics at Nortwestern University. He received his B.S. and Ph.D degrees from Massachusetts Institute of Technology in the 1950s. Freeman heads an active research group at Northwestern and has published extensively in the scientific literature.

Chong-Long Fu, a native of Taiwan, currently is a research associate in the physics and astronomy department at Northwestern University. He received his bachelor's degree from Fu-Jen University in Taiwan, and a Ph.D. in condensed matter physics from Iowa State University. His research interests center on the physical properties of bulk materials, surfaces, and interfaces.

Bernard Delley is an assistant professor at the University of Neuchatel in Switzerland. He received his doctorate from the University of Basel in 1979. His research centers on the electronic structure theory of molecules, metallic clusters, and solids.

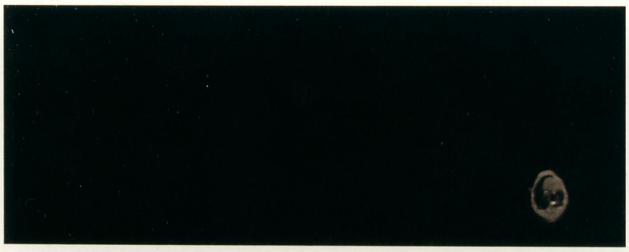

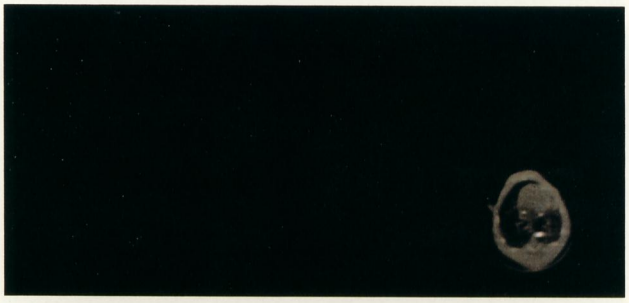

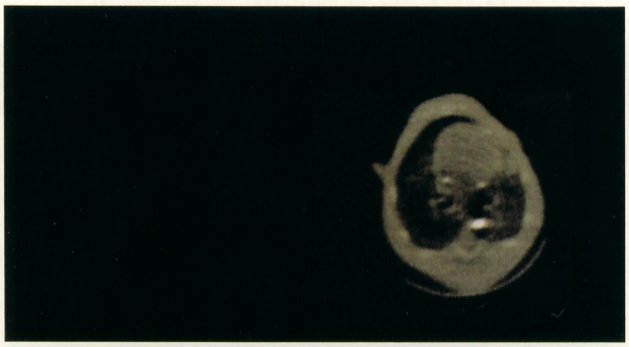

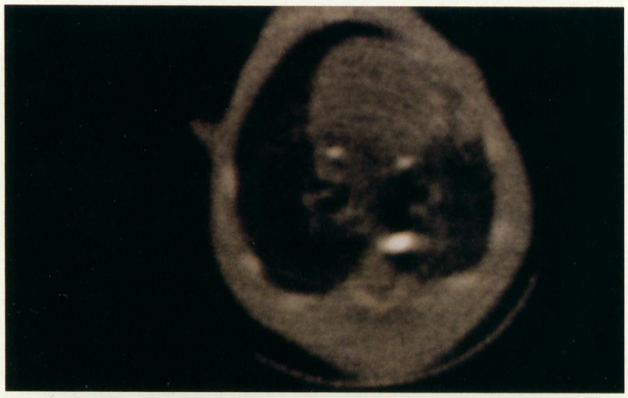

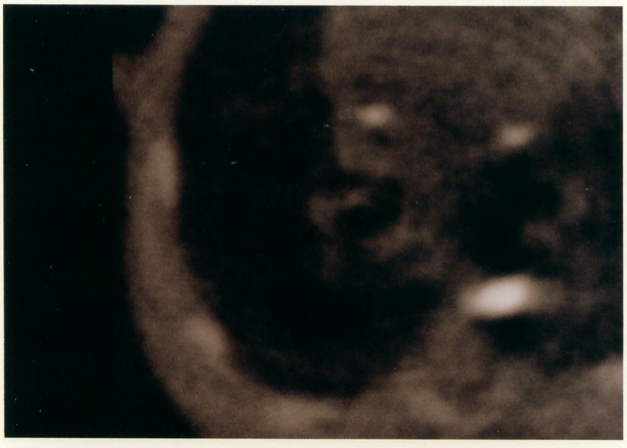

Medical imaging with CRAY computer systems

Supercomputers offer new possibilities in imaging techniques

Few people, when boarding a new jetliner, or noticing a drop in the price of gasoline, or feeling the smooth ride of a new car, stop to realize that a CRAY computer system may have been involved in making each of these situations possible. Although it is easy to identify areas where CRAY computers have contributed to far-reaching advances, it is somewhat difficult to determine their impact on daily life. Recently however, CRAY computers have aided research in areas that could touch our lives on a very personal level, by giving health professionals the ability to diagnose and treat with new precision and insight.

Medical imaging is fast becoming an area in which CRAY computers can make a vital contribution. Through complex image processing techniques, physicians can obtain accurate, detailed pictures of internal structures in the human body, allowing non-surgical diagnoses where exploratory surgery was once required. Even when an operation is indicated, the surgeon may use computer reconstruction to perform a "practice" procedure to determine, for example, the best angle of attack on a brain tumor.

Figure 1. Five cross-sectional CT images of a dog thorax. The brightest spot at the center is a catheter containing a special dye. Two smaller catheters can be seen above it. The intersection of the rib cage with the viewing plane can be observed as the six larger, bright areas around the outside of the image. Shown in correct proportion, the images contain 128 x 128, 256 x 256, 512 x 512, 1024 x 1024, and 2048 x 2048 pixels respectively. The last two images were made practical by the computational power of the CRAY X-MP/48. (The data for the images is courtesy of Dr. Richard Robb at the Mayo Clinic, Rochester, Minnesota.)

The technology to gather the data used in medical imaging has been available for some time. The tradeoff has always been time versus accuracy: producing timely graphics for use by health care workers meant that only limited processing could be performed on the data. The resulting graphics, while valuable, could not provide the kind of detail or the three-dimensionality that was ultimately desirable.

Today, two areas of medical imaging are benefiting from the availability of CRAY computers. The first, computerized tomography (CT), uses limited amounts of x-ray radiation and reconstructs images based on repeated readings taken at several different angles. The second method, magnetic resonance (MR, formerly known as nuclear magnetic resonance or NMR) offers two great advantages. First, it is able to differentiate between adjacent tissues that may have very similar x-ray densities (and would therefore be difficult to distinguish in a conventional x-ray or even a CT scan). Secondly, it does not involve radiation, depending instead on the ability of a magnetic field to spin the electrons of certain atoms.

Going x-rays one better

Conventional x-rays use photographic film and reduce a three-dimensional patient to a two-dimensional picture. The image depends on the total absorption of x-rays through each path. As a result, the two-dimensional picture can be difficult to interpret in three dimensions, particularly if there are juxtaposed or overlying structures that have similar x-ray densities. This problem can be particularly acute in areas such as the brain, which is surrounded by bone, or in an organ like the liver, where a tumor can closely resemble the healthy tissue surrounding it.

CT eliminates some of this uncertainty. It uses a thin strip of electronic sensors rather than film to record the amount of x-ray that passes through the body. Several images are recorded at various angles in the plane being studied. Successive planes can be analyzed in a like manner, producing a set of cross sections. The data is digitized, and a numerical image of each cross section can be calculated.

The resolution of the resulting image depends on the spacing of the sensing devices along each strip. A continuous set of data would be required to determine the x-ray image exactly in each plane, but the amount of data involved would overwhelm even the most modern CT processing equipment. The optimal solution, therefore, is to gather as much information as can be efficiently manipulated by the computer system being used.

Radiology without radiation

Vastly different physical principles are used to produce MR images. Instead of monitoring the absorption of x-rays as they pass through a structure in the body, MR uses a magnetic field pulse to induce certain atoms in the body to produce a detectable magnetic field of their own. Through complex calculations, different tissues in the body can be differentiated by the strength and frequency of their responses.

The magnetic resonance of particular atomic nuclei is the key to MR imaging and the origin of its name. Certain atoms have an intrinsic angular momentum or "spin" and also tend to become aligned in the presence of a static (that is, constant) magnetic field. Hydrogen, one of the most prevalent atoms in the body, is among those exhibiting this property.

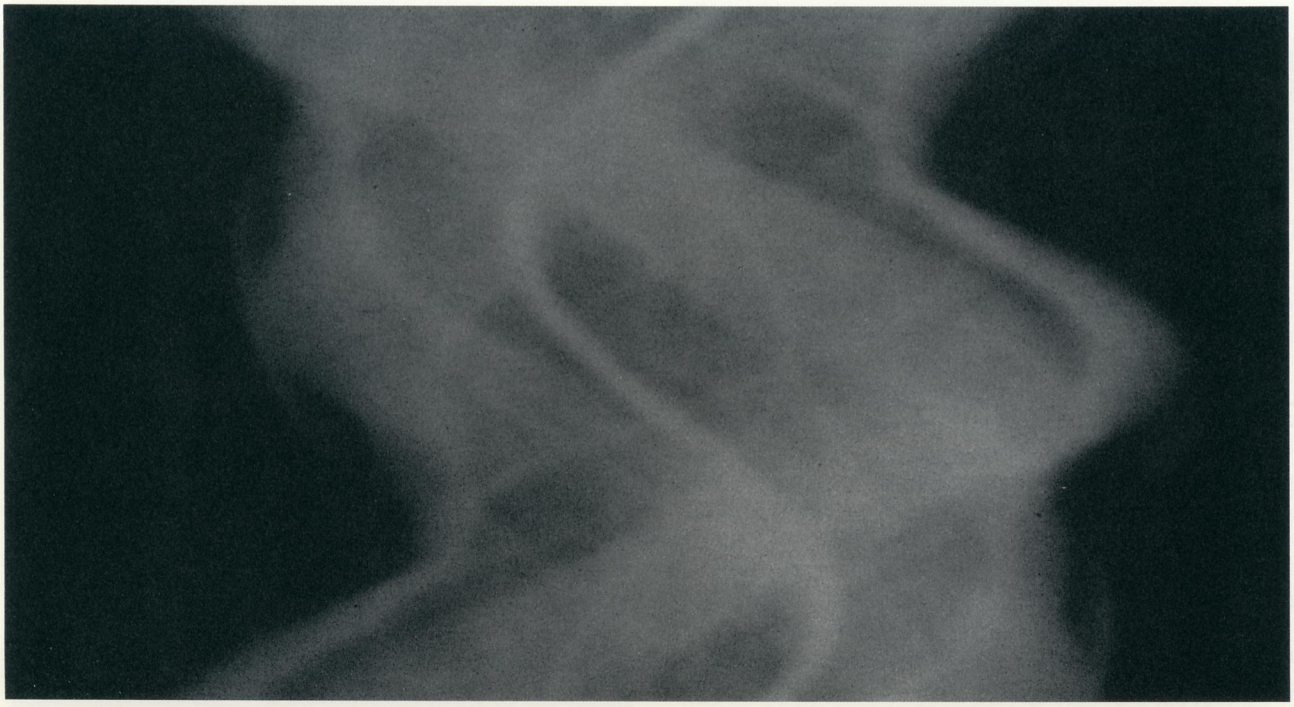

Figure 2. Input data for the dog thorax shown in Figure 1. The collected data, before processing, contains one identifiable feature: the bright S-shaped curve through the center will form the large catheter in the finished images.

The atoms have another characteristic vital to MR: they emit a signal (also known as "exhibiting a resonance") at a frequency directly proportional to the strength of the static magnetic field. When exposed to a second magnetic field that oscillates at this particular frequency, the atoms are able to absorb enough energy to realign themselves away from the first field. If then left alone, the atoms have a characteristic time, called the T1 relaxation time, during which they will return to the orientation they had before the pulse. During this return, the atoms themselves produce a magnetic field that oscillates at the resonance frequency. The field can be detected by sensors outside the body and will diminish with a second characteristic time known as the T2 relaxation time.

A problem arises in that the receiver coil for the MR signal cannot be focused very well. Atoms from the entire area emit a signal at the same frequency, making it impossible to distinguish specific structures. Fortunately, the resonance frequency can be controlled by the strength of the static magnetic field. An additional field, of varying strength across the section under study, is added as a kind of "booster" to the first field. As a result, all the atoms in one plane perpendicular to the superimposed field will have a resonance frequency different from their neighbors in the next plane. The atoms will still be emitting signals simultaneously, but when the signal is decomposed, the strength at each frequency will represent the sum of the signals emitted from one particular plane.

While serving to create distinct planes in the area under study, the concept of varying frequency between planes has a helpful side-effect: data for several planes can be gathered at the same time, reducing the number of recordings that must be taken to cover a relatively large area. As with CT, the procedure must be repeated at several angles, but more data can be garnered each time. After decomposition, the data produced in MR scans is analogous to CT data and is processed similarly. The collection of sums is converted into numerical pictures of cross sections and processed for viewing.

Computing with CT and MR

Both methods of medical imaging produce two-dimensional arrays that contain numerical data corresponding to the image at particular points in space. If the sections are taken close enough, these arrays can be stacked, allowing for three-dimensional representation of the images or the computational creation of a new cross-sectional view. In three dimensions, images can be rotated in space, providing a realistic tour about a specific area.

Another promising area for medical imaging is the shaded surface display. Regions of particular interest can be isolated, and various image processing techniques can be employed to highlight particular data. This would allow for easier distinction between adjacent tissue areas that appear homogenous at first glance.

Most CT images are displayed as 512 x 512 discrete picture elements (pixels), requiring (theoretically) 262,144 evaluations of a complicated double-integral. Fortunately, many of the calculations can be shared, but the task nonetheless requires the use of a minicomputer with the attachment of array processing hardware. For each angle at which the data was collected, the processed data is back-projected across a 512 x 512 pixel area. The summation of the back projections yields the final image.

Processing MR data is even more complex. The digitized signal data must be filtered to decompose the data into the separate frequencies that comprise each plane. The data is then transformed using an algorithm similar to the CT processing technique, and back-projected. The extra processing time for the filtering can be partially offset by the MR's ability to gather data in multiple planes simultaneously, an advantage over CT's rather serial data acquisition technique.

The role of a CRAY system

For a hospital whose CT scanner is used on up to 20 patients per day, the amount of processing that can be done is clearly limited. When working in two dimensions, a doubling of an image's resolution requires a four-fold increase the number of pixels to be calculated. In this area, CRAY systems have been used to take images from 512 x 512 pixels to 1024 x 1024 or even 2048 x 2048. While this kind of precision is seldom necessary in a short-term medical situation, it has the potential to profoundly affect medical research capabilities. The same holds true for MR, which is not yet in as widespread use as CT, but promises to gain popularity in the future because of the quality of its images and its non-reliance on ionizing radiation.

As an example, Figure 1 contains five images of a dog thorax in cross section. True to scale, the images range from 128 x 128 pixels up to 2048 x 2048 pixels. All were processed on the CRAY X-MP/48 at Cray Research's computer center in Mendota Heights, Minnesota. The third image (at 512 x 512 pixels) shows the approximate computational limit of a current generation CT scanner. The two largest images would not have been practical to generate with less than the power of a CRAY computer.

CRAY systems have also been used in MR image processing to provide results that are impractical to compute by any other means. At the University of Minnesota Underground Space Center, Research Associate Mike Mixel used the university's CRAY-1 to build a three-dimensional image from sets of MR data. As Mixel explains, "With our conventional equipment, the image took over two hours to compute. Using the university's CRAY system, that time was cut to about a minute. This kind of turnaround time is vital in the development stage when techniques are being refined. It is very difficult to conduct research any other way."

Frequently in benchmarking, the CPU time required for completion of a job tells only half the story — a

problem requiring large amounts of I/O can have relatively low CPU time, but can take an inordinately long time to run nonetheless. Thus the time elapsed during job running (usually called "wall-clock time") is at least as important as CPU time in judging the true efficiency of a computer system and the code running on it.

The smallest image in Figure 1 required 1.8 CPU seconds and 4 wall-clock seconds to compute, the difference resulting from a combination of I/O and system overhead. The largest reconstruction, which required the calculation of 4,194,304 separate pixels, required 307.7 CPU seconds and 344 wall-clock seconds. The close correlation between CPU and wall-clock time is solid evidence of the value of a CRAY system. Not only were the millions of calculations carried out with characteristic speed, but the data was able to reside entirely in memory, avoiding the costly I/O that would be required on smaller computer systems.

Future directions

Supercomputers are just beginning to play a role in medical imaging, but there are many reasons to expect that role to increase dramatically in the future. The CRAY X-MP/24 and 128-million-word Solid-state Storage Device (SSD) to be installed at the National Cancer Institute and the National Institutes of Health will be used, in part, to expand the current image rendering capabilities and support important research in tissue classification and diagnosis methods.

University of Minnesota researcher Mixel plans to use the CRAY-2 to produce high-resolution, three-dimensional images. "We hope to be able to build images at any desired angle quickly enough to be of use to surgeons. The CRAY-2 will be an invaluable tool in this research."

In both CT and MR, the computational power of a CRAY system can provide the first practical three-dimensional image reconstruction system. While it may be some time before CRAY systems are standard equipment at medical research facilities, they are becoming important diagnostic and research tools for health care professionals. □

Acknowledgement

Special thanks to Bill Samayoa of the Cray Research applications department for his help in preparing this article.

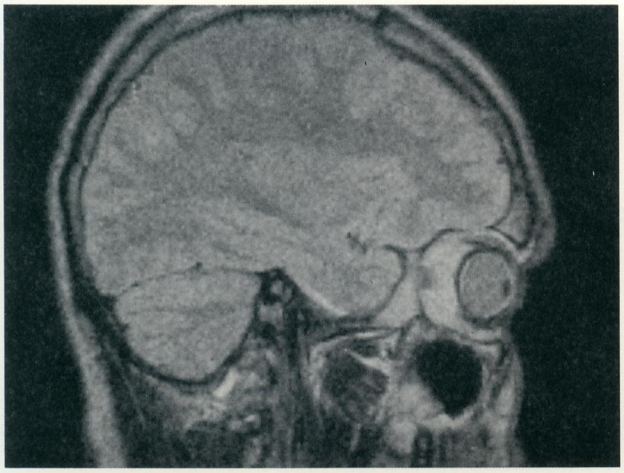

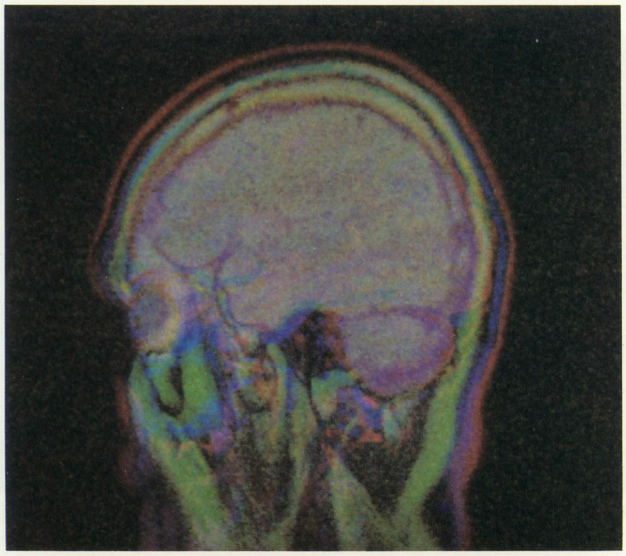

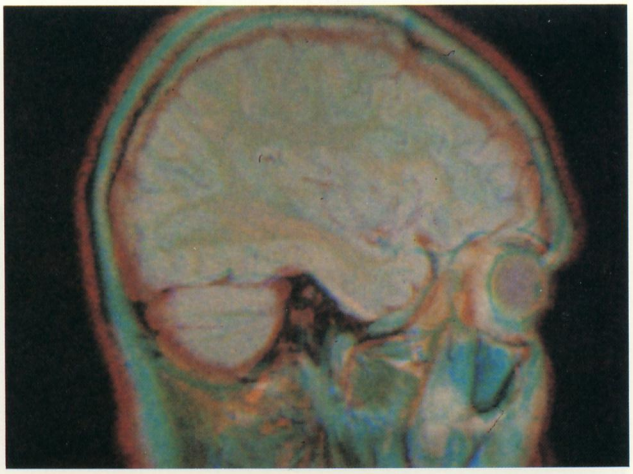

Figure 3. Three MR images recorded by the Siemens MR scanner at the University of Minnesota Hospitals and processed in part on a CRAY X-MP/48. The color images are formed by superimposing three adjacent, monochromatic cross sections (one red, one blue, one green). In the black-and-white image, the benefits of MR technology can be observed in the detail found in the skull region. The scalp is seen as the white edge; the bone is composed of two dark bands beneath the scalp. The MR image is able to detail the marrow within the scalp bone, a feature not generally possible with CT scanning. Note also the reversal of how the bone is represented: in CT scanning, bone is x-ray-dense and therefore white. In MR, bone does not tend to have a high magnetic resonance and appears dark in most images. (Images courtesy of Mike Mixel, University of Minnesota Underground Space Center, and Dr. Gunnar Lund, University of Minnesota Hospitals. Besides being involved in moving the data to the CRAY system, Mixel also happens to be the subject in the MR images.)

Gallium arsenide, the supersemiconductor

That reliable workhorse of the computer industry, silicon, is entering a race to keep its long-held position. The coming years will see alternative semiconductors challenge silicon's premier status, at least in supercomputers. Having reached the physical limits of circuit density and speed possible with silicon, circuit designers have been forced to investigate alternative semiconductors. Among them, gallium arsenide (GaAs) is emerging as the likeliest replacement for silicon in high-speed digital circuitry applications.

Unlike the single element silicon, GaAs is a compound of two elements, gallium and arsenic. Its high electron mobility gives it a factor-of-four speed advantage over silicon, making it attractive for use in high-speed computer systems. In addition, GaAs has electrical properties much more resistant to extremes of temperature and radiation than silicon's. GaAs devices also require less power than comparable silicon devices, allowing greater circuit density without overheating.

Cray Research is pioneering the application of GaAs technology in supercomputers — the company's GaAs research project is currently producing functional GaAs devices and Seymour Cray is well into the design of a next-generation supercomputer based on GaAs technology.

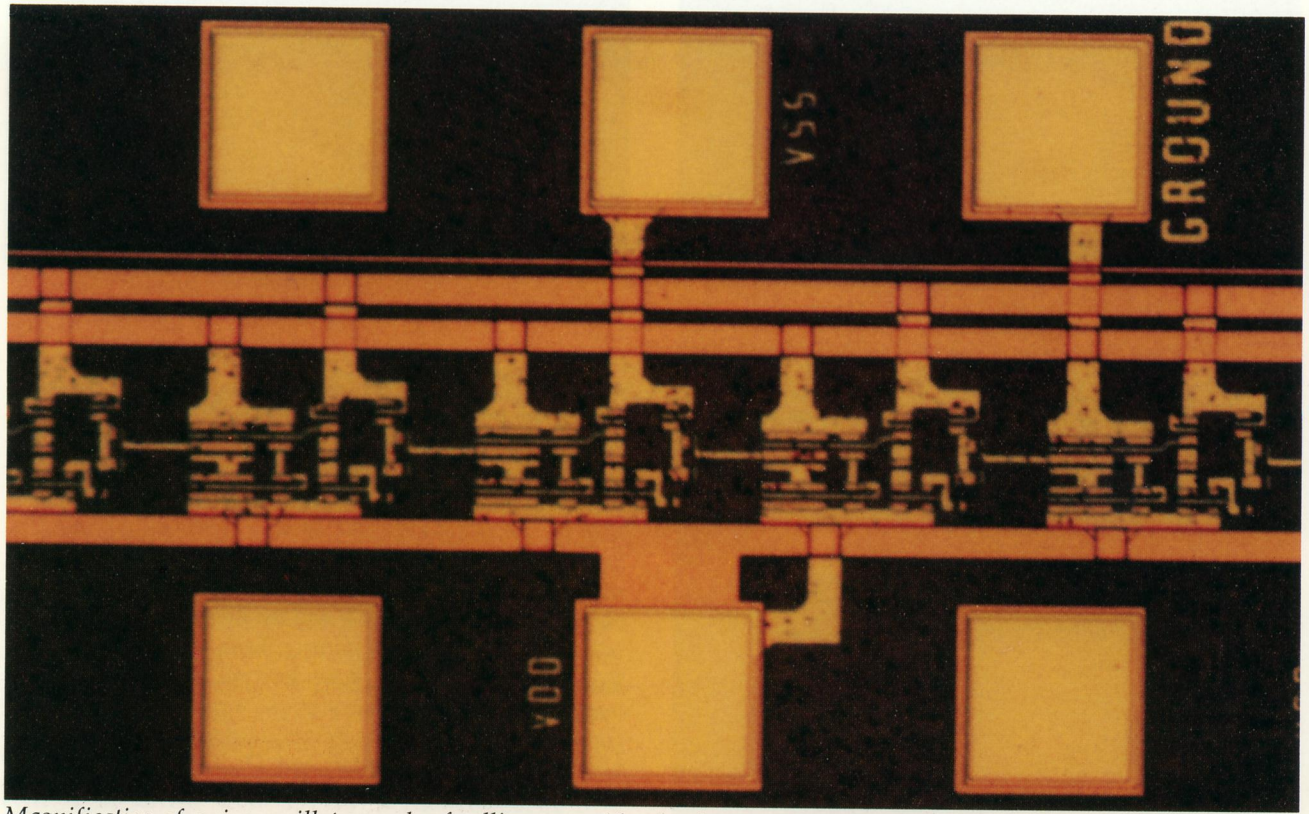

Magnification of a ring oscillator made of gallium arsenide transistors with one-micron gates. Ring oscillators are used at Cray Research to test the performance of new semiconductor technologies.

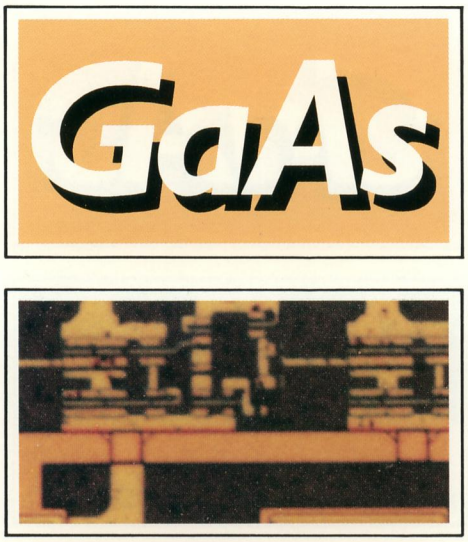

History

Gallium arsenide can be considered a third-generation semiconductor, having been preceded by germanium and silicon. In 1947, germanium was used in the first transistor. In 1958, the first integrated circuit was made by putting several transistors on a single piece of germanium. However, the process used did not lend itself to mass production. Then in 1959, the first integrated circuit of silicon was made, igniting the microelectronics revolution. The important difference between germanium and silicon is that only on silicon is it possible to grow a stable oxide film. This silicon oxide layer is used as a mask, allowing the semiconductor to be selectively treated with impurities. This treatment, called *doping*, is a necessary step in the manufacture of integrated circuits. Germanium, by contrast, grows an oxide layer that is water-soluble and not useful as a mask.

During the 1960s, engineers experimented with many semiconductor materials, mostly for optoelectronic devices. These electronic devices, such as light emitting diodes (LEDs), required new semiconductors to create the color of light desired. The color of light emitted by a semiconductor is closely related to the binding energy of its electrons, called the *bandgap*. Engineers discovered that a compound of gallium, arsenic, and phosphorous had the correct bandgap to emit red light. This discovery led to the development of red LEDs found on many types of electronic equipment. The commercial manufacture of these devices also spurred the development of new techniques for processing compound semiconductors.

In 1966, the first GaAs transistor was made. However, the transistor did not attract much attention, and no practical use was made of the device. It was initially disregarded because of the difficulty in doping GaAs with standard silicon processing techniques. In addition, early gallium arsenide tended to be unstable. Other than optoelectronics, the only early use for compound semiconductors was in simple two-terminal diodes for microwave signal generation.

However, in the 1970s, engineers began addressing the performance limits of silicon-based integrated circuits. Computers became smaller, with more transistors on each integrated circuit, but the intrinsic time delay of each circuit element had changed by a factor of only two in ten years. The speed with which electrons move through silicon was a fundamental barrier to making faster circuits.

Research ensued to develop high-speed integrated circuits. IBM and other companies invested heavily in research exploring Josephson junctions. This type of silicon device requires cryogenic temperatures, such as that of liquid helium, to achieve a superconducting high-speed transistor. However, problems associated with Josephson junctions, including the problems of operating a computer at cryogenic temperatures, hindered acceptance of this technology, and research in the United States has since been greatly scaled back.

During this period electronics firms such as Hewlett Packard and Rockwell International also explored the use of alternative semiconductor materials, primarily GaAs. In 1974, Hewlett Packard produced the first GaAs integrated circuit, demonstrating the material's ability to operate at four to five times the speed of silicon.

Recognizing the potential of GaAs for high-speed applications, the Japanese electronics industry is also developing this technology. At least four Japanese electronics and computer firms are working under contract to Japan's Ministry of International Trade and Industry to produce GaAs components for supercomputer use.

GaAs at Cray Research

Intrigued by GaAs's potential for high-speed applications, Cray Research began investigating the material in the early 1980s. The company's commitment to explore the technology was cemented in June 1982, when it organized a formal GaAs development team. Since then, a research facility has been built housing a

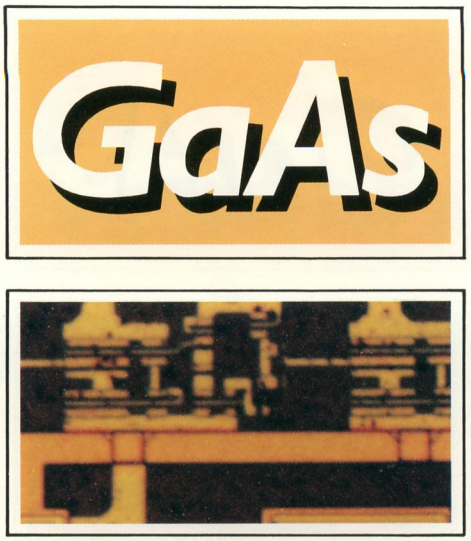

fully equipped state-of-the-art clean room. In December 1984, the facility produced its first functional GaAs integrated circuits. Having recently expanded operations to two shifts, the development team is now concentrating on improving yields and reducing production costs and turnaround time.

GaAs design considerations

Cray Research's exploration of GaAs technology is motivated by the material's potential speed advantage over silicon. But to best realize that advantage, circuit designers have to address the material's speed/power tradeoff. As Steve Nelson, director of Cray Research's GaAs research project, explains, "We can make devices that are about as fast as those we typically work with and save on power, or we can increase the power and get faster devices. Seymour Cray exploits that flexibility in his circuit designs. Rather than always striving for the fastest individual circuit, he may design one to save on power."

However, the power savings GaAs affords in turn translates into a speed advantage. A smaller power consumption produces less heat, allowing circuits to be packed more densely. Denser circuitry speeds signal propagation by shortening circuit interconnections. As a result, maximizing the speed of individual devices needn't be a primary design goal.

Another consideration relating to GaAs's speed/power tradeoff is the integration scale possible with current technology. Large-scale integration of GaAs devices is not yet feasible. "We're looking at relatively small-scale integration and relying on Seymour's ability to package the devices tightly," Nelson said. "The Japanese, on the other hand, have taken the opposite approach. They are aiming directly at large-scale integration. But we believe our goals are feasible now."

Manufacturing considerations

Although Cray Research has established working relationships with several GaAs vendors and has benefited from exchanges of technical information, the company's manufacturing process is its own. Cray Research's completely custom designed GaAs devices are created by a process involving nearly 150 steps. Unlike conventional silicon gate arrays, which are often preprocessed by vendors because the initial manufacturing steps are standardized, the company's GaAs devices are made entirely in-house from blank wafers.

The first GaAs integrated circuits produced at Cray Research were ring-oscillator circuits that consisted of about 79 transistors. "The complexity of these devices is very similar to that of the 16-gate arrays currently used in our products," explained Nelson. "Even though the design isn't optimized — it's much more typical of what you'd find in the real world — these were faster than any circuits we'd seen at Cray Research by about a factor of three."

In September 1985, Cray Research's GaAs research project successfully fabricated fully functional 1024-bit GaAs memory circuits. The circuits, designed by Seymour Cray, contain approximately 11,000 transistors. "The number of such circuits ever made by anyone could be held in one hand," commented Nelson.

The company's GaAs devices are custom designed and manufactured for good reason. "Commercially available devices are always a compromise," explained Phil Gerskovich, head electrical engineer at the GaAs research project. "Commercial devices have to serve many different customer needs. In particular, commercial manufacturers have attemped to match standard voltage levels to interface with standard circuit types. Our plan is to customize a computer around gallium arsenide completely without compromise. We don't have to be concerned with anybody else's integrated circuit needs."

The original CRAY-1 computer was based on a standard integrated circuit used throughout the machine. The company's current products, the CRAY X-MP and CRAY-2 computers, are based on approximately

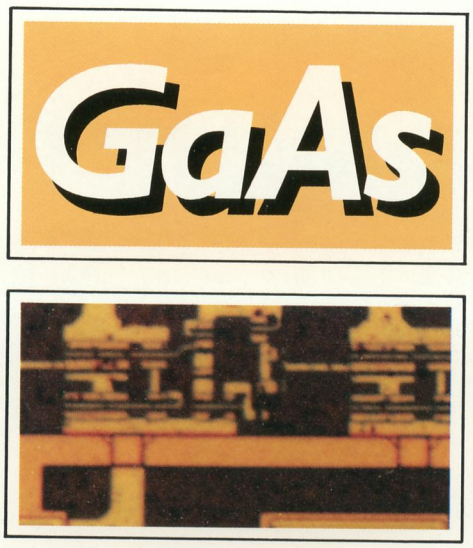

20 different kinds of circuits that vary only in their final metallization layers. "A CRAY computer system based on gallium arsenide will be made of many different totally customized logic and memory circuits," Gerskovich explained. "The circuits are becoming so complex that each has to be customized from the start."

Not only is Cray Research pioneering GaAs circuitry, but the company is doing so at a much finer lithographic resolution than is used for silicon circuits. Lines on silicon circuits are typically two to four microns wide, whereas Cray Research is successfully producing GaAs devices with lines one micron wide, a considerable technological achievement.

Commitment to the future

Cray Research's custom designed GaAs circuits typify the company's "start from scratch" approach to computer design and assure that CRAY systems using GaAs will exploit the technology to its full potential. In a relatively short time, a dedicated team of Cray Research scientists and technicians has laid the necessary groundwork to design and build a supercomputer around GaAs technology. The GaAs research project has already fine tuned its manufacturing process and is shifting into production.

Scientific computer users have come to rely on CRAY systems to handle their largest problems. Today, GaAs is a technology that promises to expand the range of supercomputer capabilities. Cray Research is pioneering GaAs circuitry because only by continually refining next-generation technologies can the company meet the needs of scientists and engineers through the 1980s and beyond. □

Acknowledgement

Special thanks to Phil Gerskovich and Steve Nelson of Cray Research's gallium arsenide research project for their help in preparing this article.

What is a semiconductor?

The chemical elements can be divided into two broad catagories, metals and nonmetals. Metals are typically good electrical conductors because electrons are not rigidly bound to metal atoms. The electrons in a piece of metal flow freely throughout the piece. Most nonmetals, by contrast, are poor electrical conductors because the electrons are bound to the atoms and are not free to flow.

However, the strength with which electrons are bound to the atoms of a nonmetal varies greatly. In some nonmetals, the electrons can break free from their atoms with energy from light or heat. When this happens, the freed electrons can flow through the material, enabling it to conduct. The energy binding the electrons to the atoms in a nonmetal is called the *bandgap*. A nonmetal with a small bandgap needs only a small energy input to become a conductor. However, if the bandgap is sufficiently large, it may be almost impossible to supply enough energy to free electrons and convert the material into a conductor. Nonmetals with small enough bandgaps to allow some conduction at room temperature are called semiconductors.

A semiconductor's conductivity can be increased significantly by adding a small amount of impurity to the material. If, for example, a small amount of arsenic is added to silicon, the conductivity can be increased many orders of magnitude. This happens because the arsenic atom has five outer electrons to chemically bond with silicon's four outer electrons. The fifth electron from the arsenic then becomes free to flow through the silicon and contribute to its conductivity. Adding such an impurity to a semiconductor is called *doping*, and is a standard process in the manufacture of integrated circuits. Ironically, silicon is an effective dopant for gallium arsenide.

CRAY X-MP hardware performance monitor

John Larson, Cray Research, Inc.

A variety of tools exist on CRAY systems to help the user understand the behavior of programs. These tools cover a wide spectrum of user needs. For example, static analysis tools such as FTREF and VMARK provide the user with information on the references to FORTRAN variables and vectorization potential at compilation time. After unsuccessful executions, utilities including DEBUG and DUMP are invaluable in removing programming errors. During program execution, dynamic analysis software such as FLOWTRACE and SPY supply the user with timing information. This article describes FLOP TRACE and PERF TRACE, two additional dynamic analysis tools available through BENCHLIB on CRAY X-MP systems.

A unique hardware feature of CRAY X-MP computer systems is the hardware performance monitor. It is a set of counters that monitors certain hardware-related events. These performance monitor counters provide the user with detailed information about the speed, execution requirements, and characteristics of a program. Such information is valuable to the applications programmer for code analysis and optimization. In addition, it assists the performance modeler in the development and verification of computer models.

Each CPU of a CRAY X-MP computer system contains eight 46-bit counters that track certain events in one of four event groups referred to as Groups 0, 1, 2, and 3. Group 0 includes floating-point operations, instruction issue and hold issue, instruction buffer fetch, and I/O and CPU memory references. Group 1 consists of detailed conditions that prevent instruction issue. The events in Group 2 are central memory references and conflicts. Instruction types and vector operations are in Group 3.

The hardware counters monitor events only while the CPU is in user mode. In this mode, the counters are incremented each clock period by the number of events that occur. The performance monitor counters operate in parallel with the execution of the user's program and do not interfere with or delay the events being monitored.

```
PROGRAM  RANK=  0    %AGE= 100.0  ACCOUNTED CP= 28677215970     CALLED= 1
         INDX=  0  ACCUM%= 100.0   CPU SECONDS= 272.43355172  AVE TIME= 0.27E+03
 12.16 MIL. INSTR./SEC (MIPS)        3312144296 INSTRUCTIONS ISSUED          8.66 AVE. CP PER INSTRUCTION
                                    24551730045 CP HOLDING ISSUE            85.61 % OF CP HOLDING ISSUE
  0.09 MIL. INSTR. BUF. FET./SEC       23907224 INSTR. BUFFER FETCHES
  0.07 MIL. I/O MEMORY REFS/SEC        19534556 I/O MEMORY REFERENCES          0.08 % I/O REFS IN TOTAL REFS
 90.90 MIL. CPU MEMORY REFS/SEC     24764737579 CPU MEMORY REFERENCES         0.70 MEMORY REFERENCES/FLOP
 59.52 MIL. FL. PT. ADDS/SEC        16215981208 FLOATING POINT ADDS          45.56 % ADDS IN TOTAL FLOPS
 68.29 MIL. FL. PT. MULTS/SEC       18603590392 FLOATING POINT MULTIPLIES    52.27 % MULTS IN TOTAL FLOPS
  2.83 MIL. FL. PT. RECIPS/SEC        770826896 FLOATING POINT RECIPROCALS     2.17 % RECIPS IN TOTAL FLOPS
130.64 MIL. FL. PT. OPS/SEC(MFLOPS) 35590398496 FLOATING POINT OPERATIONS    100.0 % FLOPS IN PROGRAM FLOPS
```

RANK	INDX	NAME	CALLED	TIME(SEC)	AVE-TIME	%AGE	ACCUM%	ADDS	MULTS	RECIPS	FLOPS	MEM/FLOP	MMEM/SEC	MFLOPS
1	11	UPDH	36864	1.03E+02	2.79E-03	37.73	37.73	5.66E+09	6.77E+09	5.25E+08	1.30E+10	0.43	53.98	126.07
2	13	PLANLOOP	1152	5.27E+01	4.57E-02	19.34	57.07	4.07E+09	4.38E+09	3.32E+05	8.45E+09	0.94	150.59	160.35
3	10	PLAQIN	36864	3.81E+01	1.03E-03	13.99	71.06	2.96E+09	3.06E+09	0.00E+00	6.02E+09	0.68	107.97	157.94
4	18	FUNC	18432	2.07E+01	1.12E-03	7.60	78.66	9.45E+08	1.47E+09	2.15E+08	2.63E+09	0.21	27.07	126.71
5	8	BICORN	27648	1.82E+01	6.58E-04	6.68	85.34	1.31E+09	1.53E+09	0.00E+00	2.84E+09	0.70	109.26	156.28
		...												
		...												
0	0	PROGRAM	1	2.72E+02	2.72E+02	100.00	100.00	1.62E+10	1.86E+10	7.71E+08	3.56E+10	0.70	90.90	130.64

Figure 1. Shown at the top is abbreviated output from PERF TRACE showing only the statistics for an entire program. Below that, sample FLOP TRACE output gives condensed data on individual subroutines.

System software associates a single event group with the hardware counters and manages the counters for the user. The library routine PERF provides the user interface to the system, allowing the user to specify the desired event group. Calls to this routine may be inserted into the user's code to turn the counters "on" and "off," to provide the user with the values in the counters, and to indicate the elapsed time. Cray Research's *Library Reference Manual* for the COS Version 1.14 release describes the use of this routine and lists the monitored events.

New software now allows use of the hardware performance monitor without program modification. The first step is to compile the code with the FLOWTRACE option. This causes the compiler to insert calls automatically in the user program to the FLOWTRACE system software in preparation for execution flow analysis. The first statement in every routine becomes a call to FLOWENTR, and a call to FLOWEXIT is placed before each RETURN statement. A STOP statement is replaced by a call to FLOWSTOP. These names are also the entry points in the new software, so at load time the user may specify that the association of entry points be made with the new software rather than with FLOWTRACE.

The use of FLOP TRACE and PERF TRACE is described in the BENCHLIB documentation. These tools complement the timing and flow information from FLOWTRACE by providing additional statistics obtained from the performance monitor counters.

PERF TRACE is a collection of four analysis programs, one for each performance monitor event group. The output of each PERF TRACE program is a detailed listing of the eight monitored events for each of the subroutines in the user code. The actual counts of the monitored event described above, the average rate at which they occurred, and the percentage of some event aggregate represented by the single counter are displayed. Program statistics are also shown. A sample of the output is shown in Figure 1 where the events in Group 0 are displayed for the entire program.

FLOP TRACE is an analysis program that provides selected information about Group 0 events. The output of FLOP TRACE is a table containing one line of information about each of the subroutines in the program. The statistics include the actual number of floating-point operations performed by the subroutine, the average memory reference rate, the ratio of memory references to floating-point operations, and the average MFLOP rate. The table is sorted so that the most time consuming routines are at the top of the list. The same statistics are produced for the program as a whole. A sample of the output is shown in Figure 1. The BENCHLIB documentation provides assistance in interpreting the quantities displayed and offers optimization suggestions.

Through the use of the hardware performance monitor and supporting software, the user may identify the important subroutines in a program and concentrate optimization efforts there. Thus the user can learn if subroutines are executing at a high MFLOP rate, and if not, many clues are provided as to the cause. Moreover, this information can be collected by simply accessing a new library. No program modifications are required.

The hardware performance monitor is among the utilities available in BENCHLIB. Those interested in ordering the software should contact Cray Research, Inc., Applications Department, 1333 Northland Drive, Mendota Heights, MN 55120, or their nearest Cray Research regional office. □

About the author

John Larson is a senior applications engineer with the Cray Research development group in Chippewa Falls, Wisconsin. He joined Cray Research in 1982 after working for Burroughs Corporation from 1978 to 1982 on the development of the Burroughs Scientific Processor (BSP) and Flow Model Processor (FMP) systems. He received his Ph.D. in computer science from the University of Illinois at Urbana-Champaign in 1978.

Notice to all CRAY supercomputer users

We would like to include a forum in CRAY CHANNELS for addressing supercomputer performance issues. You can help us in this effort by submitting ideas for topics that you would like to see covered. Your technical questions will provide the direction for this forum, and we will do our best to answer as many of your questions as possible.

We are also interested in your experience using CRAY computers in your computing environment. If you have developed techniques for improving system performance, let us know. Through this forum in CRAY CHANNELS we can share the information with other users who might benefit.

The level of interest you, our readers, have in this effort will determine whether it becomes a permanent feature of CRAY CHANNELS. In the meantime, we plan to offer articles with practical "how to" advice for getting the most out of CRAY systems. Such articles planned for future issues will describe and offer advice for implementing multitasking, microtasking, and new I/O and compiler features.

Quantum chromodynamics on a CRAY X-MP/48 with an SSD

Philippe de Forcrand and John Larson, Cray Research, Inc.

Quarks are believed to be the fundamental constituents of matter. Quantum chromodynamics (QCD) is the leading theory that describes how quarks interact by exchanging "colored" massless objects called gluons. Unfortunately, QCD is not amenable to an analytic solution. Therefore, theoretical physicists have turned to a numerical approach using computers.

Using numerical simulation, the candidate QCD theory is simulated on a computer. Measurements taken during the simulation are compared with experimentally derived results. Such comparisons help researchers decide whether to accept or reject the theory, particularly as they improve control over the biases due to modeling simplifications — a challenging assignment given the formidable computer resources required.

Nevertheless, the increased predictability of a more refined model fully justifies the endeavor. Eventually, by adjusting a single free parameter of QCD theory, all of the properties of matter should be reproducible on a computer. For instance, experimental physicists have yet to observe an isolated quark. However, the unique property of quarks to remain confined has now been observed in simulation. The predictive power of QCD could also be tested by looking for new particles seen only on the computer so far, for example, "glueballs" made of gluons.

A full QCD simulation includes the modeling of both quarks and gluons. Research continues on how to simulate light quarks efficiently. The standard approach taken here addresses the case of gluons and static (infinitely heavy) quarks. The Monte-Carlo simulation generates snapshots of a four-dimensional space-time box of gluons, where the chromo-electric and chromo-magnetic gluon fields are subject to quantum fluctuations. Averages over this ensemble of configurations are then compared with experimental quantities, assuming the box size and the number of lattice points in the box are both large enough.

The most fundamental quantity to measure is the potential energy between a static quark and a static antiquark. These particles are linked by a string of gluons and propagate in time without moving, thus producing a rectangular loop with sides representing the space separation R and the propagation time T. The integral of the gluon field around this R by T rectangle is a number W(R,T) called the Wilson loop. For large T, the average value of W(R,T) has a known functional form, from which one can extract the quark-antiquark potential energy for any selected distance.

Recently, a computer simulation was executed on a CRAY X-MP/48 with a 128-million-word Solid-state Storage Device (SSD). The simulation modeled a box of 24 x 24 x 24 x 48 lattice points representing over 21 million degrees of freedom. The multitasked program consisted of two sections repeated nearly 4000 times: generation of a new configuration (48 million words of data stored on the SSD) and measurement of loop quantities.

Natural parallelism in the first part of the program stemmed from the local character of the interaction among gluons — only neighboring data were needed to evolve each gluon from the old configuration to the next one. The updating algorithm, of the quasi-heatbath type, was efficiently vectorized over non-interacting regions by a checkerboard partitioning in two dimensions. Parallelism was exploited over a third dimension with perfect distribution of the work among processors. The fourth dimension, time, was ascribed to the I/O operations. To update a given "time-slice," information was extracted only from the neighboring two time-slices. Three time-slices stayed together in central memory with a buffer space for a fourth. The buffer space alternatively contained updated or old information to be asynchronously sent to or from the SSD for the next iteration. Each degree of freedom was read only once from the SSD, updated, and written back, thus minimizing the I/O. Even with this optimal strategy, the speed of the computation was such that conventional disks would have compromised the execution time severely — I/O time would have totaled over a week using Cray Research's fastest disk drives, the DD-49s. (The DD-49 data capacity is 1200 Mbytes and it has a sustained transfer rate of 9.8 Mbytes/sec measured at the user job level). As it was, the SSD instantly eliminated I/O wait time.

The second part of the program, which measures Wilson loops, had to be done simultaneously with the snapshot generation part, since the data storage requirements otherwise would have been prohibitive. The gluon field was only read, so that vectorization and parallelism could be exploited easily over the entire lattice. An arrangement similar to the first part

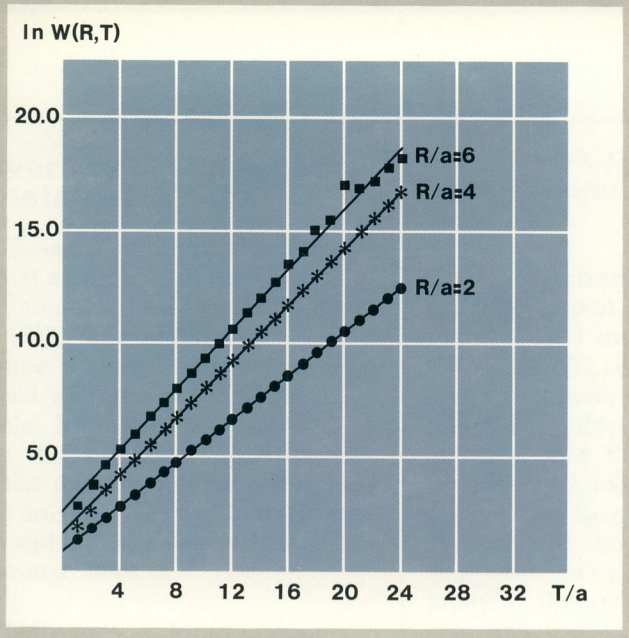

Figure 1. The potential V(R) can be extracted as the asymptotic slope of a straight line fit to the logarithms of the loops of width R.

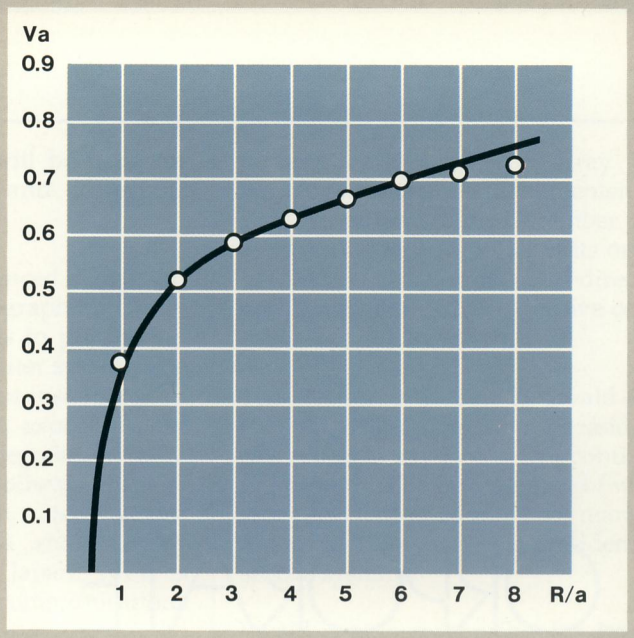

Figure 2. The dependence of the potential V(R) on the quark separation.

of the program was retained, with vectorization over a full two-dimensional plane, parallelism over four such planes, and I/O along the fourth dimension. Various methods like gauge transformation (elimination of spurious, non-physical degrees of freedom) and local averaging were used to decrease the number of operations and the statistical noise.

Multitasking exploits the natural parallelism of the problem to produce a factor of 3.77 speedup on the CRAY X-MP/48 over a single-processor execution. In this instance, multitasking becomes a necessity for practical reasons; in dedicated mode, the problem requires 200 hours of CPU time. Using multitasking, this can be reduced to 53 hours of wall-clock time. In batch mode, multitasking minimizes the impact of this job's requirements (5 million words of central memory to roll in and out, 48 million words on the SSD) on overall system throughput.

Optimization of the program has been pursued with the help of the hardware performance monitor (see the article on page 18), concentrating the effort on the weakest and most important routines. For example, FLOP TRACE output (from the hardware performance monitor) revealed that routine PLANLOOP was slower than expected. Careful data reshuffling led to a longer vector length while still avoiding memory conflicts. As a result, the performance of the routine jumped from 120 to 160 MFLOPS. Figure 1 on page 18 summarizes the present performance of the code for the generation and analysis of one snapshot on a single CPU. The sustained processing rate is 130 MFLOPS. With the speedup factor of 3.77 folded in, a sustained 490 MFLOPS is realized with all four processors on the X-MP/48. The link-update time, by which such programs typically are rated, is less than six microseconds.

The two figures above show a sample of the results. In Figure 1, the average Wilson loops W(R,T) are plotted versus T on a semi-logarithmic scale. For a fixed R and large T, they tend to fall on a straight line. The slope of the line is the potential V(R). As the measurements extend further along T (up to T=24 instead of 8 as in previous simulations), the estimate of the true slope becomes more exact. The fitted values V(R), as well as the simplest model (solid line) $V_M(R) = V_0 - \alpha/R + \sigma R$, are plotted in Figure 2. The simulation agrees rather well with this model where the force between quarks tends to a constant σ — not to zero! — at large distance. This constant, called the string tension, can be evaluated experimentally, and fixes the single adjustable parameter of the simulation.

Simulations of this type do not strain the machine despite their spectacular size. We hope this example will encourage researchers to gauge the latest CRAY systems at full power and prompt them to tackle problems previously considered too complex to attempt. □

About the authors

Philippe de Forcrand is a senior applications engineer for Cray Research in Chippewa Falls, Wisconsin, having joined the company early in 1985. Formerly, he was pursuing post-doctoral research at Ecole Polytechnique in Paris, gaining experience with a CRAY-1/S system there. He received his Ph.D. in theoretical physics in 1982 from the University of California, Berkeley.

John Larson is a senior applications engineer with the Cray Research development group in Chippewa Falls, Wisconsin. He joined Cray Research in 1982 after working for Burroughs Corporation from 1978 to 1982 on the development of the Burroughs Scientific Processor (BSP) and Flow Model Processor (FMP) systems. He received his Ph.D. in computer science in 1978 from the University of Illinois at Urbana-Champaign.

CORPORATE REGISTER

Use of CRAY systems expanding

Cray Research has passed a number of significant milestones through the years: the first order from a commercial customer, the first order from an aerospace company, and the first order from a petroleum company, to name a few. Each of these achievements was the beginning of a wider relationship between Cray Research and a group of new customers. In the past few months, the company announced a pair of new milestones: the receipt of orders from customers in two new industry areas.

In August, Cray Research announced it had received an order from the Du Pont Company for a CRAY-1 computer system. The system will be the first supercomputer installed in the chemical industry and will be housed at Du Pont's Experimental Station near Wilmington, Delaware. The CRAY-1 will be used in a number of research areas, including life sciences, electronics, catalysis, and polymer sciences. It is currently being installed.

H.E. Simmons, vice president of the Central Research and Development Department at Du Pont, said, "With its ability to solve far more complex equations in far shorter times than conventional computers, we believe it is essential to our R&D effort that Du Pont scientists have full access to state-of-the-art computer power."

Du Pont is a diversified manufacturer and marketer of chemistry-based products and a supplier of energy resources. The company currently spends in excess of $1 billion annually on research and ranks fifth among U.S. manufacturers in R&D expenditures.

Shortly after the Du Pont order was disclosed, Cray Research announced that Fairchild Camera and Instrument has ordered a CRAY-1 S/2000 supercomputer. The system has been installed at Fairchild's Gate Array Division in Milpitas, California. This order represents the first CRAY system in the U.S. dedicated to semiconductor research and development.

Lanny Ross, General Manager of Fairchild's Gate Array Division, said: "The CRAY system will be integrated into the existing CAD environment for use by Fairchild's customers, and standard cell and gate array simulation and layout software will be moved to the CRAY system." Ross noted that the new capability will reduce design turnaround by a significant factor, and stated, "Advanced research in semiconductor development done on the CRAY supercomputer will provide Fairchild with another major edge over our competition."

CRAY X-MP to be installed for National Cancer Institute

Cray Research announced late in November that a CRAY X-MP/22 computer system will be installed at the National Cancer Institute (NCI) in Bethesda, Maryland, in the first quarter of 1986. The system will be purchased by Falcon Systems, Inc., a Bethesda-based provider of systems integration services. NCI will use the system for complex bioengineering and molecular modeling in its cancer research.

John Rollwagen, chairman of Cray Research, said, "We have received many important and exciting orders for CRAY computer systems in the past, but none has filled us with as much good feeling as this one. We share with everyone the hope that the research efforts of the National Cancer Institute will bear much fruit, and it will be wonderful if a supercomputer can contribute to that."

University of Minnesota to install CRAY-2

Cray Research announced in September that an affiliate of the University of Minnesota, Research Equipment, Inc. (REI) has ordered a CRAY-2 computer system. The system will be shipped to the REI facility in Lauderdale, Minnesota by the end of 1985. REI provides supercomputer facilities and services to researchers at the university, as well as to government and industrial users throughout the United States.

The University of Minnesota Supercomputer Institute will be one of the major users of the CRAY-2 system. Dr. Peter C. Patton, director of the institute, said: "We see the CRAY-2 as an essential enabling technology for research programs in computational science and engineering. Our efforts to develop large-scale computational science as a complement to theoretical and experimental science require a research tool of this magnitude. We are excited by the opportunities the CRAY-2 architecture provides."

John A. Rollwagen, chairman of Cray Research, commented: "The University of Minnesota has been an academic pioneer in this field. It was the first university to acquire a CRAY-1 computer and now it will be the first to have access to the CRAY-2. We are delighted to be a part of the university's continued leadership."

More news about orders

In October, Cray Research announced an order from the Societe Nationale ELF Aquitaine Production for a CRAY X-MP/12 computer system and a 32-million-word Solid-state Storage Device. The system will be installed at the customer's facility in Pau, France in the first quarter of 1986, subject to export license approval. It will replace the CRAY-1/M computer system that was installed in 1983. The Societe Nationale ELF Aquitaine Production is France's leading petroleum company. The new CRAY X-MP will be used for seismic research and reservoir modeling.

Cray Research announced in September that Nippon Telegraph and Telephone (NTT) intends to purchase a CRAY X-MP/1 computer system. The system is scheduled for installation in mid-1986, subject to export license approval and the negotiation of a final contract. John Rollwagen, chairman of Cray Research, noted that this system will be the sixth CRAY computer going to Japan. NTT is Japan's domestic communications company.

Organizational update

Cray Research announced several important internal events in October. The board of directors elected Steve Chen senior vice president and approved the initiation of a new project for Chen's research and development team. In addition, Seymour Cray extended by two years his agreement to work as an independent contractor to Cray Research. The board also elected two new vice presidents, Robert H. Ewald and Paul W. Dillingham, Jr.

The charter of Steve Chen's new group is to design the most powerful supercomputer it can, developing both hardware and machine-specific software and applications. Chen was the principal designer of the CRAY X-MP Series of Computer Systems.

"The new research project will enable Steve and his team to explore new architectures and technologies," said company chairman John Rollwagen. "This decision is significant for the company because it gives us two teams working at the frontiers of supercomputing." The other such team, of course, is headed by Seymour Cray. Rollwagen added that the company would continue to devote the largest part of its technical resources to the extension of its current product line.

Meanwhile, Seymour Cray has elected to extend his research and development agreement with Cray Research once again. The agreement is now in effect until December 31, 1989 and continues in all of its original terms. Cray remains a director and a member of the executive committee of Cray Research.

The election of Robert H. Ewald and Paul W. Dillingham to vice presidencies reflects Cray Research's continuing growth. The Board named Ewald vice president of commercial marketing and Dillingham vice president of special systems.

Ewald is now responsible for assessing opportunities in the commercial marketplace and identifying ways to further expand the company's existing commercial business. Since joining Cray Research in mid-1984, Ewald has directed education industry and commercial marketing efforts. Prior to joining Cray Research, he worked for seven years at the Los Alamos National Laboratory, in Los Alamos, New Mexico. His most recent position at the laboratory was leading the computing and communications division.

Dillingham's new role gives him responsibility for assessing marketing opportunities to government classified customers and for helping direct a field organization of sales representatives and customer support personnel. Dillingham joined the company in 1984 as director of the special systems office, after a distinguished 30-year Naval career in intelligence and cryptology.

Next CUG meeting to discuss UNIX

The spring meeting of the Cray User Group (CUG) will be held May 5-8 in Seattle. The theme for the meeting will be UNIX™. A call for papers is being handled by David Lexton, University of London Computer Centre, 20 Guilford Street, London WC1N 1DZ, England. Dennis Ritchie, one of the original designers of UNIX at AT&T Bell Laboratories, will be the keynote speaker.

CORPORATE REGISTER

Cray Research's new marketing support services building in Mendota Heights, Minn.

Marketing services expand in MH

What do Cray Research's applications, benchmarking, contracts, customer services, diagnostics, field support, and proposals departments have in common? They're all in the same building! These departments are the residents of the company's new marketing support services building in Mendota Heights, Minnesota. Completed last summer, the new facility enables marketing personnel to better serve customer needs. Along with the added elbow- and legroom, the new facility houses a well-equipped computer graphics laboratory. In addition, a high bandwidth fiber optic connection, via a Network Systems Corporation HYPERchannel®, provides the facility with access to the company's computer center, also in Mendota Heights. The computer center, equipped with a CRAY X-MP/48, a CRAY X-MP/216, and a prototype single-processor CRAY-2, supports benchmarking, software development, and code conversion efforts.

The new marketing facility features an enlarged presentation room equipped with a remote audio-visual system, including video projection equipment and several slide projectors. "In the future we expect to accommodate visitors with multimedia 'tours' of our engineering and manufacturing facilities in Chippewa Falls," said John Aldag, manager of the applications department.

Expansion of the software division forced the construction of the new building, and software now enjoys sole residence of the site that formerly housed both software and marketing personnel. With desks in hallways not unheard of when both divisions shared the same building, all concerned agree that the expansion was necessary. The company's software and technical operations buildings are within walking distance of the new marketing building, so employees and visitors can still enjoy the campus-like atmosphere of the Mendota Heights facilities.

New AT&T UNIX Station Software Service released

Since its inception, Cray Research has worked consistently to integrate its computer systems into existing customer environments, communicating with a diversity of equipment and a host of operating systems. Now users of the UNIX™ operating system can have immediate access to the power and performance capabilities of the CRAY X-MP and CRAY-1 computer systems running the CRAY Operating System (COS).

Written in the C programming language, the AT&T UNIX Station Software Service for CRAY systems provides fully functional batch and interactive access from VAX computers manufactured by the Digital Equipment Corporation. Future releases will support a wide variety of hardware and software from other manufacturers.

The station links CRAY systems running COS Version 1.13 or later and operating systems based on UNIX System V Version 2.0. The connection is handled through the Network Systems Corporation HYPERchannel® network.

The AT&T UNIX Station Software Service provides a state-of-the-art interactive link to optimize user productivity and performance. Users may specify redirection of interactive input and output and can observe the results of each CRAY job control language command before proceeding to the next.

In addition to the interactive capabilities, the station provides for data transfer between UNIX and COS operating systems, CRAY job input and output, and COS job monitoring. Files may be obtained from or created on the front-end system using the COS ACQUIRE, FETCH, and DISPOSE control statements. A TEXT parameter on each of these commands provides the user with a means to identify any file on the UNIX system. Files on the front end may be sent directly to CRAY mass storage with the SAVE station command. On-line documentation is available to assist users with the station commands.

The AT&T UNIX Station Software Service is leased independently of CRAY system software. For additional information, contact the nearest Cray Research sales office.

APPLICATIONS IN DEPTH

UM-SPICE simulates GaAs ICs

The electronics industry is sprouting wings; demand for compound semiconductors is soaring. These exotic new materials may someday challenge silicon's dominant semiconductor status. Gallium arsenide (GaAs), in particular, promises to usher in a new age of compound semiconductors that will substantially outperform silicon in digital circuitry applications (see related article on page 14). But as compound semiconductors such as GaAs grow in popularity, integrated circuit designers must "retool" their workstations to accommodate and take full advantage of the new materials.

Such a retooling has been initiated by University of Minnesota electrical engineering professor Michael Shur and graduate students Bruce Bernhardt and Choong Hyun, who have adapted the popular circuit simulation program SPICE to simulate GaAs circuits. SPICE was developed at the University of California to simulate silicon circuitry and is used in circuit design worldwide. Although SPICE can be used to model GaAs circuits, the model parameters do not relate to actual GaAs processing parameters. This discrepency limits the program's usefulness in optimizing the design of GaAs devices.

The researchers solved the problem by replacing certain routines in SPICE with routines they developed specifically for GaAs. The new program, UM-SPICE, includes routines that model MESFETs (MEtal-gate Schottky Field Effect Transistors), and other routines that model the ungated GaAs field effect transistor and the aluminum gallium arsenide/GaAs MODFET (MOdulation Doped Field Effect Transistor). The new routines incorporate more accurate models of GaAs's capacitances and current-voltage characteristics than were available in the original SPICE.

"The improved capacity to model MODFETs that UM-SPICE provides will be increasingly important as MODFETs come into commercial use," said Steve Nelson, director of Cray Research's GaAs research project. "We expect to use the code, or similar ones, when we investigate gallium arsenide MODFETs."

The University of Minnesota research team used an IBM PC for the initial development work on UM-SPICE. They later converted the program for the university's CRAY-1. "The PC was ideal for simulating small circuits," Bernhardt explained. "But a microcomputer is relatively slow and has limited memory. A much more powerful computer is needed for realistic simulations of large circuits."

"To simulate a 25-stage ring oscillator, the program creates and solves a set of equations in a 101 x 101 matrix," Bernhardt continued. "We know that as an inverter chain grows by a given amount, the matrix grows by the square of four times that amount. So, as the circuits grow, the computing power required increases substantially. This is why we converted UM-SPICE to the university's CRAY-1. We were primarily interested in the speed advantage offered by the system's vector capability, though its large memory also helps us to simulate larger circuits."

Using UM-SPICE, simulation of a 25-stage ring oscillator takes 46 minutes on the IBM PC compared to 2.4 seconds on the CRAY-1, Bernhardt noted. The largest ring oscillator the researchers have simulated on the CRAY-1 so far contained 75 stages, requiring approximately one fourth of the system's one-million-word memory.

The tremendous leap in memory and processing capacity the CRAY-1 offers is being exploited by the research team to analyze the effects of threshold voltage variations on yield and reliability. GaAs circuits designed using MESFETs require tight control of the threshold voltage. But this is difficult to accomplish because the threshold voltage across a GaAs wafer, and from wafer to wafer, has greater variance than the threshold voltage across silicon wafers.

Simulations of large circuits of inverter chains will be used for the threshold voltage studies. The simulations will include an average threshold voltage and a threshold voltage variance of a given standard deviation about the average, variables which can be set by the researchers. As the standard deviation is varied, the researchers will calculate the chains' best, average, and worst case noise margins and propagation delays. This information will enable them to estimate the tolerances necessary for the threshold voltage if a given circuit yield is desired. The studies will also provide information on circuit reliability, given a certain deviation about a given threshold voltage for a particular logic family.

For more information on UM-SPICE, contact Bruce A. Bernhardt, Department of Electrical Engineering, University of Minnesota, Minneapolis, MN 55455; telephone: (612) 373-4102.

NSF centers open for business; resource services expanded

In March 1985, the National Science Foundation (NSF) established four national supercomputer centers at U.S. universities. These "Phase II" centers are intended to address the long-range need of university researchers for supercomputer access. In addition, the NSF recently expanded its list of Phase I supercomputer resource centers. NSF purchases services from Phase I centers to provide researchers with interim access to supercomputers.

Two of the designated Phase II centers will provide access to CRAY systems: the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign, and the San Diego Supercomputer Center administered by GA Technologies at the University of California, San Diego. Although the Phase II centers are not scheduled to open to remote users until January 1986, the NSF began accepting requests for computer time in the fall of 1985. Any researcher, whether NSF-supported or not, is eligible to apply for time. Proposals should be submitted to the appropriate NSF research program in one of three categories:

- Requests associated with currently-funded NSF research. These requests should receive very quick turnaround, perhaps one week or less, because the research has already passed the peer review process.
- Requests that are components of standard NSF research proposals. These requests will be reviewed in conjunction with the proposal. The peer review process typically takes about three months.
- Requests for supercomputer resources only, where the research is not supported by the NSF. These requests must find a host NSF research program, and will be treated similarly to new research proposals. If the research is funded by an agency other than NSF, peer review will be required for the supercomputer request itself.

Once review procedures are established at the centers, researchers will have the option of applying either to the NSF or to the center directly. A document with instructions for applying to the four centers is currently being prepared by the NSF's Office of Advanced Scientific Computing, and will include an application form. The locations and configurations of the centers and corresponding contact persons are listed below:

- University of California at San Diego, CRAY X-MP/48, contact Dr. W. Pfeiffer, telephone: (619) 455-3467.
- University of Illinois, CRAY X-MP/24 with 32-million-word SSD, contact Dr. R. Wilhelmson, telephone: (217) 244-0072.
- Cornell University, IBM 3084QX with FPS 164 and 264 processors, contact Mr. W. Schrader, telephone: (607) 256-8686.
- Princeton University, CYBER 205 with two pipelines and four-million-word memory, contact: Dr. A. Brenner, telephone: (609) 734-8191.

In addition to establishing the four new centers, the NSF expanded access to supercomputers through its resource center program. The original resource centers, Purdue University, the University of Minnesota (soon to be offering time on a CRAY-2), and Boeing Computer Services, have been joined by three additional participants. The new resource centers are Bell Laboratories, Colorado State University, and Digital Productions. These centers offer the services of a CRAY X-MP/24, CYBER 205, and a CRAY X-MP/22, respectively.

"Within the last six months the number of requests we've received for supercomputer access has skyrocketed," said Jim Bottum, staff associate at the centers program. Bottum attributed the surge in demand to several causes, including the academic calendar, centers hooking up to national communications networks, and increased publicity. "Until recently, NSF had been allocating time at the center requested by the researcher," Bottum said. "But NSF is no longer allocating time at Boeing and Purdue because of the high volume at those centers." Questions regarding procedures for submitting requests to either Phase I or Phase II centers should be directed to the NSF's Office of Advanced Scientific Computing; telephone: (202) 357-9776 or -9717.

Cray Research hosts NASTRAN workshop

In October, representatives from a number of Cray Research customer sites attended a three-day workshop on MSC/NASTRAN and CRAY computers. MSC/NASTRAN is a structural analysis program used extensively in the automotive, aerospace, and civil engineering industries. The workshop, which was sponsored by the company's applications department and the MacNeal-Schwendler Corporation, was intended for people with experience running MSC/NASTRAN on computers other than CRAY systems. "The point was to share our experience with customers and show them the most efficient ways to run MSC/NASTRAN on CRAY computers," explained Greg Clifford, a software analyst in the applications department.

Attendees were invited to bring their own problems to run on the CRAY X-MP/48 at Cray Research's computer center in Mendota Heights, Minnesota. John Clinard from Oak Ridge National Laboratory in Tennessee submitted a problem that has been running on a CRAY-1/S at the National Magnetic Fusion Energy Computer Center at Lawrence Livermore National Laboratory in Livermore, California. The problem involves calculating stresses created by the magnetic fields of six superconducting electromagnets. Each magnet weighs 40 tons and produces a magnetic field of 8 tesla. The problem

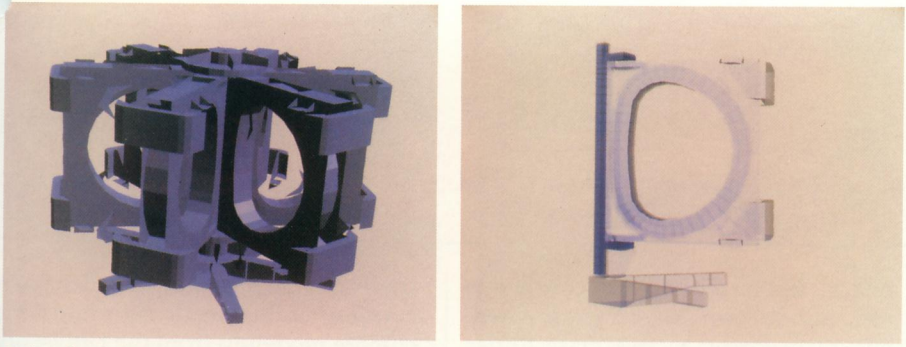

Left, shaded image of electromagnet array rendered with MOVIE.BYU. Right, the "GLASS" command invokes MOVIE.BYU's transparency feature, used here to display the inner structure of a single coil of the six-coil array.

is being used to help set safe operational levels for the facility housing the electromagnet array.

To help analyze the problem, a NASTRAN finite element model of one coil was passed via a translator to MOVIE.BYU and rotated about a vertical axis to create the additional five sections. The entire model was displayed with MOVIE.BYU as a hidden line plot, a continuous-tone shaded image, and then "x-rayed," using the program's transparency feature. The MOVIE.BYU rendering is the largest to date using that software — the model contains 38,000 polygons.

"I came to the workshop to get information about and experience handling large databases," Clinard said. "Our current database exceeds 100 million words. The workshop examined the use of large buffers, double buffering, striping, and other techniques involving the SSD. It was very worthwhile — I was able to conclude that with the proper hardware configuration and SSD utilization, we can manage our I/O and solve the problem. Currently, we are installing MSC/NASTRAN on our CRAY X-MP/12 at Oak Ridge. Using this system, we can apply our workshop experience to maximizing throughput for large problems."

"We feel the workshop was valuable to all the participants," Clifford added. "We're glad to share our expertise with customers, and help them get the most use out of their CRAY systems." The workshop concluded with tours of the computer center and the computer graphics laboratory in Cray Research's new marketing support services building.

Petroleum users and Cray Research talk

The exchange of information was the goal of a three-day conference sponsored by Cray Research's petroleum region in Houston, Texas in October. The conference, titled "The Efficient Use of Supercomputers in Petroleum Applications," was attended by about 150 representatives from virtually all the major oil companies — CRAY system users and prospects alike. In addition, several seismic processing companies and researchers attended the conference. The conference provided a forum for petroleum industry supercomputer users to share information about using their systems and for Cray Research to receive feedback about what the company can do to further address petroleum industry needs.

The first day was dedicated to seismic processing issues. Speakers included Don Larson, GeoQuest International, Inc.; Doug Spragg, Exxon Production and Research; and Dr. Dan Kosloff, Tel Aviv University. Perhaps the highlight of the day was the panel discussion on the efficient use of CRAY systems in seismic processing and modeling. Straightforward questions and candid answers offered by the panel made up of Cray Research customers in the petroleum industry made the discussion productive.

The second day focused on practical aspects of high-speed computing on CRAY computers. Hardware, software, and implementation techniques addressing petroleum industry supercomputing concerns were discussed. Technical personnel from all areas of Cray Research spoke on a variety of topics including networking, performance tools, CFT programming techniques, and multitasking and microtasking implementations. Steve Chen, senior vice president of Cray Research, spoke to the group following the evening banquet about past, present, and future supercomputer development activities at Cray Research. He encouraged users to "work hard" at using multiprocessor systems, stressing the increasing role parallel systems will play in the future of high-speed computing.

The final day of the conference was devoted to reservoir simulation computational developments. Dr. John Wallis of J.S. Nolen and Associates spoke about reservoir simulation on CRAY computers, while other speakers from the petroleum companies included Dr. Leslie Smith, Exxon Production and Research; Phil Stephenson, Mobil; and John Killough, ARCO Research. With a panel of industry experts, a lively discussion of reservoir simulation issues concluded the conference.

George Stephenson, petroleum region general manager, was pleased that so many representatives from the different oil companies participated in the conference. Stephenson commented, "I hope that the participants came away with a better understanding of how CRAY computers can be used more effectively in the petroleum environment. At the same time, we welcomed the opportunity to hear the honest comments and suggestions about how the company can continue to meet petroleum industry supercomputing needs."

USER NEWS

NCAR offers insights to petroleum researchers

More than 50 percent of today's oil deposits are found in Cretaceous rocks — rocks that formed 65 to 140 million years ago. Eric Barron, formerly a research scientist at the National Center for Atmospheric Research (NCAR), now working at the University of Miami, has been using NCAR's CRAY-1 systems to simulate the climate conditions responsible for the formation of petroleum source rock. So fascinating are the findings that the work has attracted the attention of several petroleum companies. ARCO, Exxon, Texaco, and Mobil have all contributed gifts to NCAR to support this research.

By simulating prehistoric climate conditions, researchers can identify probable areas of coastal upwelling — areas that are conducive to the burial and preservation of organic compounds and thus to the formation of petroleum deposits. Upwelling results from a surface divergence, which occurs when surface water is replaced by underlying water. The most important observation pertaining to surface divergence is the deflection of surface water by wind forces, giving a mean direction for water motion at right angles to the wind. Along coastlines, prevailing winds with the appropriate direction can cause persistent, large-scale upwelling. Complicating an otherwise simple relationship between wind direction and coastal upwelling are the effects of bottom friction, stratification, local currents, topography, and shoreline configuration.

Reconstructing prehistoric climate conditions also involves simulating movements of the Earth's tectonic plates, upon which the continents rest. Subduction zones, areas where one plate forces a land mass under another, are the primary locations of earthquakes, volcanoes, and deep-sea trenches. The movement of continents, carbon dioxide from volcanoes, and many related factors combine to affect atmospheric conditions and climate. With mathematical models for testing the sensitivity of many climate variables, researchers are able to determine the extent to which these factors affect the atmosphere.

In simulating prehistoric climates, Barron performs sensitivity experiments on NCAR's Community Climate Model (CCM), a three-dimensional, time-dependent, general circulation model of the atmosphere. "I use the CCM on the CRAY-1 to flood the continents, then remove the floods and compare the climate changes," Barron said. "But even with the two CRAY systems at NCAR, this is still a computer-limited problem."

The CCM is coupled to mixed-layer and fully-resolved ocean models capable of simulating the seasonal cycle. Most recently the studies have been testing theories suggesting that the tilt of the Earth was 10 degrees in prehistoric times rather than the 23-degree tilt we see today. "The more variables we find that influence our studies, the more computationally intensive the problem becomes," Barron said. "The minute you want to work with all complexities, three-dimensional grids, and all of the variables and factors involved in climate studies, you need to consider the computer power available for analysis. Supercomputers are absolutely necessary for handling these complex problems." Barron is continuing his work on prehistoric climate modeling from the University of Miami via a satellite link to NCAR's CRAY systems.

Numerical climate models based on the fundamental physical laws that govern atmospheric and oceanic circulation have the potential to be a major new tool in locating petroleum source rock. With the capabilities provided by supercomputers, researchers can address the controls of organic productivity and preservation and offer new insights into the buried treasures of today's fossil fuels. Supercomputers allow scientists to work with increasingly complex models — they are tools of discovery that allow scientists to view their work, and the world, in a new dimension.

CFD challenge flows through CRAY X-MP/48

Some people can't resist a contest, including certain aerodynamicists. When West Germany's GAMM aerospace research institute issued a challenge for an international computational fluid dynamics (CFD) competition, James Scott of the University of Dayton Research Institute responded, armed with a CRAY X-MP/48. The challenge posed by GAMM was to calculate supersonic flow through a two-dimensional double-throated nozzle. The problem described a highly specialized flow situation which provided a test of program efficiency and computer hardware performance.

Scott worked on solving the problem while visiting Cray Research's applications department in Mendota Heights, Minnesota, last fall. He

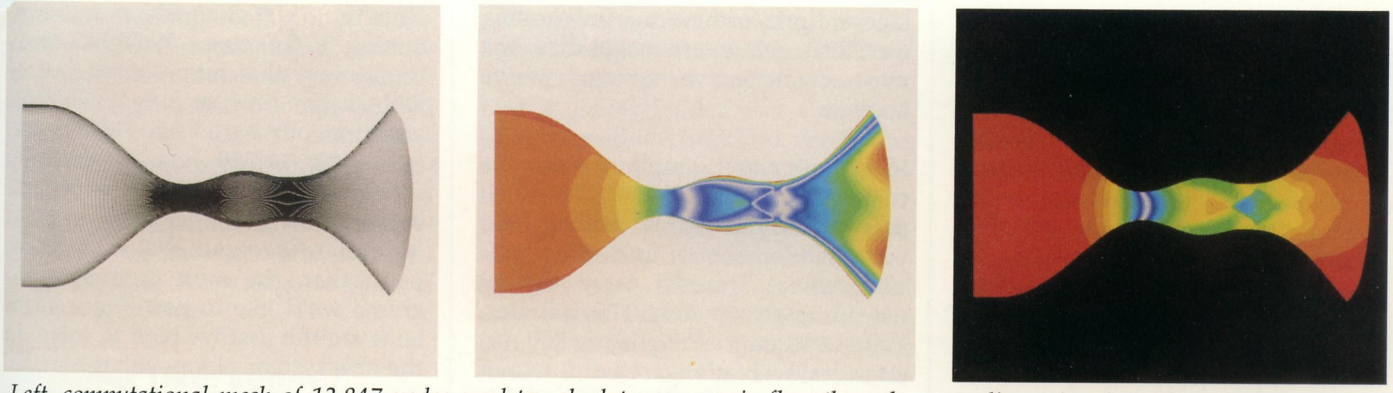

Left, computational mesh of 13,847 nodes used to calculate supersonic flow through a two-dimensional, double-throated nozzle. Center, plot of flow velocity up to Mach 3.3. Red represents slowest flow; white, fastest. Right, plot of static pressure up to one pound per square inch. Red represents highest pressure; white, lowest. Flow is from left to right.

solved the problem using a two-dimensional Navier-Stokes code that he had developed in collaboration with researchers at Wright-Patterson Air Force Base. Scott had been running the code on a CYBER 845 at Wright-Patterson prior to visiting Cray Research. Without modifications, the code ran at about 100 MFLOPS on a single processor of the CRAY X-MP/48, allowing the work for five configurations to be completed in about four days. "This same job would have taken about seven months using the CYBER 845," Scott commented.

Not only did the CRAY X-MP/48 run the code approximately 60 times faster than would have been possible on the CYBER mainframe, but it also enabled Scott to solve the problem using a very fine computational mesh. "Memory constraints on the CYBER limited my mesh resolution to 3000-4000 nodes," Scott said. "The CRAY X-MP system's larger memory allowed me to use a mesh of nearly 55,000 nodes."

The fine mesh computation (54,813 nodes) was performed to verify the validity of medium (13,847 nodes) and coarse (3627 nodes) mesh computations. Performing calculations with a finer-than-necessary grid mesh, as was done here, is a common CFD practice for verifying results obtained from coarser meshes. Because the accuracy of a Navier-Stokes (viscous) flow solution is directly proportional to the number of nodes used for the computational mesh, a fine mesh will reveal small-scale flow phenomena that escape the detection of a coarser mesh.

"The problem GAMM posed specified a Reynolds number that determined a high viscosity for the fluid," Scott explained. "I suspect they were primarily interested in seeing how people would deal with the high viscosity — it's a unique flow situation. There were no winners in the competition; its real purpose was to get people to compare their methods to see what they could learn from each other." Scott presented his findings in December at a GAMM workshop, "Numerical Simulation of Compressible Navier-Stokes Flows."

Genetic sequencers bank on CRAY systems

Comparing database entries is a computer application one probably wouldn't associate with supercomputers. Few databases are large enough, or contain complex enough entries, to require a supercomputer for conducting comparisons. But there is at least one that does — the DNA database.

As molecular biologists have probed deeper into the workings of DNA, they have identified thousands of genetic sequences that specify protein structures and govern the maintenance and reproduction of biological cells. To help researchers centralize and catalog the sequence data, the National Institutes of Health (NIH) established a computer database at the Los Alamos National Laboratory in Los Alamos, New Mexico. The database, called GenBank, is available on magnetic tape, floppy diskettes, on-line, and in hardcopy form from Bolt, Beranek and Newman, Inc., of Cambridge, Massachusetts. GenBank was founded in 1982 and is managed and updated by laboratory personnel who use several computers for searching the database and comparing genetic sequences. For quantitatively comparing very long sequences, the laboratory's CRAY computers are the staff's tools of choice.

GenBank's mission is to collect all known genetic sequences containing at least 50 nucleotide bases, the chemical subunits of DNA. Bases are of four types, abbreviated A, T, C, and G. Their sequential arrangement determines the function of a given strand of DNA. With known sequences growing at a rate of close to two million bases per year, there is a clear need for high-speed computing to make database searches and sequence comparisons practical.

Genetic sequence comparisons can provide clues about the evolutionary relatedness of different species. But the greatest importance of sequence comparison information is the insight it provides medical researchers.

USER NEWS

"Finding similarities between a sequence of unknown function and one of known function aids molecular biologists working with the unknown sequence," explained Walter Goad, GenBank's scientific director. "For example, although the entire genome (genetic makeup) of a given virus may have been identified, the function of a particular sequence within the genome may not be known. Knowing the function of similar sequences can help medical researchers discover how the virus causes disease or how to make an effective antigen for it."

However, comparing sequences is a complex statistical operation that must account for the probability of coincidental similarities among sequences and three kinds of localized changes in sequence: replacement of one base by another, deletion of a base, and insertion of a base. Conducting such comparisons is an onerous task that has been made much more bearable by GenBank's access to the Los Alamos laboratory's CRAY computers.

To enable and encourage searches of the entire database for similarities to a newly identified sequence, the GenBank staff relies on a "distance" algorithm using the following strategy: for every comparision of the query sequence with a known sequence, the similarity score for the best local alignment of the two sequences is saved; after a run through the database, the statistically significant scores are printed out, together with the names of the corresponding sequences. This list can then guide a more focused examination of the similarity of the query sequence to others in the database.

"Comparing many long sequences can be very time-consuming and expensive," commented Christian Burks, a theoretical biologist and GenBank staff member. "For local reasons, we've recently been using mainframes, but we plan to transfer all sequence comparison functions to CRAY computers. The speed advantage we get from vector processsing on CRAY computers makes them the most cost-effective systems we've applied."

In a particular example of sequence comparisons run on a CRAY-1, all pairwise comparisons among 204 vertebrate sequences (including their complement strands) were carried out in approximately 170 minutes. Pairs of sequences, averaging 800 nucleotides, were compared at a rate of 240 per minute. (A discussion of this work appears in *Nucleic Acids Research*, Vol. 13, No. 2, 1985, pp. 645-656.)

The medical applications of the GenBank database are well appreciated by the NIH, which will soon have access to a CRAY system installed at the National Cancer Institute. Comparing newly identified sequences with those already catalogued in GenBank will be a primary use of the new CRAY X-MP system.

"Our main emphasis will be on translating genes into proteins and running sequence comparisons of the amino acid chains that make up the proteins," said Jacob Maizel, Chief of the Laboratory of Mathematical Biology at the Frederick Cancer Research Facility, NIH's largest divsion. "We've been working the past year on sequencing proteins translated from the DNA of an adenovirus containing 36,000 bases. We'd like to do the same with a herpes virus, but its DNA is five times as long. Doing this work on our two VAX 11/780s and one VAX 11/750 requires dividing the sequence into small segments and compromising on detail by using a 'skimming' program that is less thorough than the more rigorous programs we'd like to use. These problems are the first we plan to solve on the new system, along with the analysis of cancer viruses and oncogene sequences."

Distance algorithms are also useful for comparing the sequence of one strand of DNA with that of its own complement. High similarities in this type of comparison can be used to trace regions of potential "hairpin" structures in the RNA transcribed from the DNA. Such structures, where the RNA folds and bonds to itself, are in some cases known to be important signals to enzymes and may regulate the rate of protein synthesis. (For a discussion of computer modeling of RNA structures, see "RNA modeling for biotechnology," CRAY CHANNELS, Spring 1985.)

Although CRAY systems are used primarily to simulate physical phenomena, any computationally demanding problem is a potential

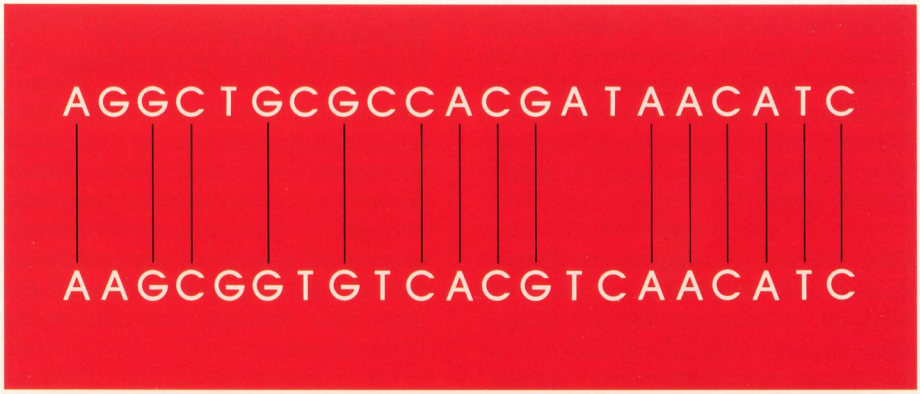

Example of nucleotide base sequence comparison with 15 matched nucleotides. Similarity scores are determined by a coefficient times the number of matches minus a coefficient times the number of misses. The programs to be used at NIH can realign sequences if doing so improves the similarity score. Such realignment accounts for isolated deletions, insertions, and replacements of bases that throw off the score of otherwise similar sequences.

candidate for supercomputing. In the case of genetic sequencing, a virtually unmanageable database is yielding its fruits thanks to the computational proficiency of CRAY computers. Giving creative scientists and engineers access to supercomputers brings us nearer to not only understanding the natural world, but also applying that knowledge to solving human problems.

Another (Mersenne) prime candidate

More than 2500 years ago, the Greek mathematician Euclid proved that an infinite number of primes exist. They do not occur in a regular sequence, and today, even with the help of supercomputers, it is impossible to predict where and when one will pop up. Father Marin Mersenne, a 17th-century French monk, developed the first sequence of prime numbers in 1644 by raising the number 2 to successively higher prime powers then subtracting 1 from the result: 2^p-1, where p is prime. Since then, the search for larger prime numbers has become an elite challenge among computer scientists and mathematicians. All you need is a few hours time on the world's fastest computer and a little luck.

That was the scene in Houston last September when the 30th Mersenne was discovered during a test on Chevron Geosciences' CRAY X-MP/24 system. The machine was running a confidence test using Prime Finder, a program created by Cray Research software engineer David Slowinski of Chippewa Falls. As the creator of the Prime Finder program, Slowinski will go down in the record books for the discovery of the 30th Mersenne, but it was the Chevron analysts running the program who realized the discovery. Under the direction of Bill Bartz, vice president of Chevron's computer services, they were looking over integrity testing results when they saw that the CRAY system had made an exciting discovery: 2^{216091}-1 is a prime number—all 65,050 digits of it.

Although there is no practical use for a 65,050-digit number, there is a very practical use for the program that can find it. Prime Finder uses a method called "residue check" to verify the correctness of each of the thousands of 80,000-digit integer multiplications used to test potential primes. This residue check is executed dozens of times each second and immediately recognizes a hardware fault. (See "Searching for the 27th Mersenne Prime" by David Slowinski, CRAY CHANNELS, Vol. 4, No. 1.) The time required to hunt for each successively higher Mersenne prime increases exponentially. On average, it is as much work to discover one Mersenne prime *now* as it would be to rediscover all the others four times over.

If finding a Mersenne prime is a challenge, testing its primeness is another. In 1876, French mathematician Edouard Lucas published a method for determining whether or not a Mersenne number is a prime. The method was refined in 1930 by University of California, Berkeley, mathematics professor, Derrick Lehmer. The Lucas-Lehmer test relies on multiplication and addition and eliminates time-consuming division.

The task of re-checking the 30th Mersenne prime for primeness fell to Steve McGrogan of Elxsi Computer in San Jose, California. While the CRAY X-MP/24 required a little over three hours to discover the 30th Mersenne, it took McGrogan 216 hours of computer time to verify the discovery using Elxsi systems and his own program. McGrogan's program includes a systematic search through lower number space to determine whether any Mersenne primes exist between those already known. He is also developing an algorithm that may significantly speed up the testing for prime numbers. Commented McGrogan on independently searching for and testing Mersenne primes, "It's a learning experience and a challenge. It's like Mount Everest — why do people climb mountains?"

Superkids visit MFE

A program hosted by the National Magnetic Fusion Energy Computer Center (NMFECC) at Lawrence Livermore National Laboratory tapped the energy of 52 of the country's brightest high school students in August. The program, sponsored by the U.S. Department of Energy, familiarized the students with supercomputers and their capabilities.

The students were selected by the governors of each of the 50 states, the District of Columbia, and Puerto Rico as outstanding students based on scholastic aptitude and competence in science and mathematics. Although the participants are students now, in ten years they will be among the most sought-after young scientists and researchers in the country.

"As you can see, we have a selfish interest in attracting bright young people into science," said James Decker, deputy director in the U.S. Department of Energy's Office of Energy Research, in his opening address. "Our country owes its high standard of living to its pre-eminence in science and technology. In order for us to maintain that leadership, we must attract young people with the very best minds into scientific careers."

Along with hearing presentations by representatives from Lawrence Livermore National Laboratory, Los Alamos National Laboratory, the Department of Energy, the University of California, the California Governor's office, NASA, and Cray Research, the students were paired with tutors from NMFECC who helped them use the world's most powerful computers, including the CRAY-2. MFE's computer center has four CRAY computers (a CRAY-1, a CRAY-1/S, a CRAY X-MP/22, and a CRAY-2) and provides services to over 4000 researchers working in fusion and basic energy research.

By tackling selected projects with their tutors, the students obtained

Fifty-two outstanding high-school students learned about supercomputing at the National Magnetic Fusion Energy Computer Center.

firsthand knowledge of the value of supercomputing in scientific research. For an introduction, the students ran a quadratic equation solver, which gave them an opportunity to access the system, use the editors, and debug the system. A second project involved a particle physics code which, because of its complexity, demonstrated the need for supercomputers in a scientific environment.

The students' third project examined a mathematical model known as a Mandelbrot set, one of the most complex objects in mathematics. A Mandelbrot set involves a series of numbers and fractals in a complex plane. As each number reaches a limit, it is assigned a color. When an iterative operation is applied to these numbers, the ones outside the set flee to infinity, and the numbers inside mark the onset of instability. Graphing the results creates a complex and beautiful representation of the model. The Mandelbrot matrices quickly filled the available memory on the CRAY-2, further demonstrating the need for continued development in supercomputing for mathematical research.

Also introduced in the NMFECC summer program were projects on the distributed plotting of utilities and a number of small programs that were interspersed between the major projects. Students also came prepared with programs of their own, which they eagerly adapted to run on the CRAY systems. One student worked on a problem to project population growth, some were intent to find Mersenne Primes, and 27 played chess simultaneously against the CRAY BLITZ chess program.

Summer institutes offer supercomputer training

The National Science Foundation (NSF) sponsored three summer supercomputer institutes this year to train U.S. researchers in state-of-the-art computing. The National Center for Atmospheric Research (NCAR), the University of Minnesota, and Boeing Computer Services were selected to hold the forums, which included coursework, seminars, speakers, tutorials, and hands-on experience with supercomputers.

"This type of program is extremely important because much of the material is not available as standard curriculum in the universities," noted Paul Swarztrauber, director of the institute at NCAR. "Supercomputing is expanding at an increasing rate, and it is important that we provide an educated user base for industry as well as academia."

Demand for training in supercomputing clearly is escalating. Each NSF summer institute had approximately 80 applicants vying for fewer than 30 openings. Participants included graduate students, postdoctoral fellows, and young faculty members from around the country. "It has been our experience that young faculty and graduate students are prime candidates for supercomputer training," said Tom Walsh, Scientific Director of the Supercomputer Institute at the University of Minnesota. "Because these people are just beginning their research and thesis programs, they have a need and desire to learn about the technical capabilities available to them."

Applied problems in the basic sciences were a key part of the summer programs. The Boeing Computer Services program had participants with backgrounds in the life sciences, control systems, fluid dynamics, and physics, whose research projects included problems ranging from the fluid dynamics of a heart valve to predator-prey parameter estimation.

"It appears that the program will affect the 25 participants in 25 different ways," explained Mike Markley, project manager of the Boeing Computer Services summer institute. "We had people with no supercomputer experience and others with extensive backgrounds in vector processing. The students each received ten hours of CRAY time for attending the institute and they are continuing their research projects at their respective universities. Many learned about techniques and algorithms for structuring their problems to take advantage of the hardware and software available today."

WORLDWIDE BUSINESS CENTERS

Domestic

Corporate Headquarters
608 Second Avenue South
Minneapolis, MN 55402
(612)333-5889

Government Marketing/
Government Relations Office
National Place
1331 Pennsylvania Avenue, Suite 1331 N
Washington, DC 20004
(202)638-6000

Central Region Office
5330 Manhattan Circle
Boulder, CO 80303
(303)499-3055
Albuquerque
Chicago
Detroit
Minneapolis/St. Paul
St. Louis

Eastern Region Office
11710 Beltsville Drive, Suite 500
Beltsville, MD 20705
(301) 595-5100
Atlanta
Boston
Laurel
Pittsburgh
Rochester
Tampa

Petroleum Region Office
5847 San Felipe, Suite 3000
Houston, TX 77057
(713) 975-8998
Dallas
Tulsa

Western Region Office
5776 Stoneridge Mall Road
The Atrium, Suite 350
Pleasanton, CA 94566
(415) 463-2800
Los Angeles
Seattle

International

Cray Canada Inc.
4211 Yonge Street, Suite 210
Willowdale, Ontario
Canada M2P 2A9
(416) 229-2729

Cray Research France S.A.
7 Rue de Tilsitt
75017 Paris, France
(01)-766-01-55

Cray Research GmbH
Perhamerstrasse 31
8000 Munich 21, West Germany
(089)-56014-0

Cray Research Japan, Limited
W 1154, Shin Aoyama Building
1-1 Minami-Aoyama 1-chome
Minato-Ku, Tokyo 107, Japan
(03)-403-3471

Cray Research S.R.L.
Via Vivaio, 15
20122 Milan, Italy
(02)-70-98-27

Cray Research (UK) Ltd.
Cray House
London Road, Bracknell
Berkshire RG12 2SY
United Kingdom
(0344)-485971

CRAY RESEARCH, INC.

Corporate Headquarters
Cray Research, Inc.
608 Second Avenue South
Minneapolis, MN 55402
Tel: (612) 333-5889