A CRAY RESEARCH, INC. PUBLICATION

CRAY CHANNELS

Volume 5, Number 4

FEATURE ARTICLES:

Computational physics on the CRAY

Everything you wanted to know about GaAs

Cray addresses U.S. House

Supercomputing conference at Los Alamos

Los Alamos turns forty

REGULAR COLUMNS:

Corporate register

Applications in depth

User news

In this issue

Over the past year, the supercomputer industry has received unprecedented attention. The importance of supercomputers for research as well as the need for advancement of supercomputer development has gained worldwide recognition. While not seeking out the limelight, Cray Research has been the target of much of this attention.

After this issue of CRAY CHANNELS was completed, it became apparent that, ironically, all of the articles contained here either augment or recap supercomputer issues that have been discussed repeatedly over the past 12 months. We hope that you'll enjoy our feature articles that relate Cray's participation in conferences held to explore the status and future of supercomputer technology in the United States. Another story provides insight into a new technology being developed at Cray that is expected to significantly impact our next-generation system. And finally, two other articles look at how large-scale computing capability impacts progress in research. One article discusses the contributions supercomputing is making in the field of computational physics, and our article about Los Alamos National Laboratory describes how advanced computing capability has affected its research efforts.

The issues surrounding the supercomputer industry today are still heating up. We hope this issue offers you an interesting perspective of Cray's approach to supercomputer development and the pivotal role our systems have in research environments.

About the cover

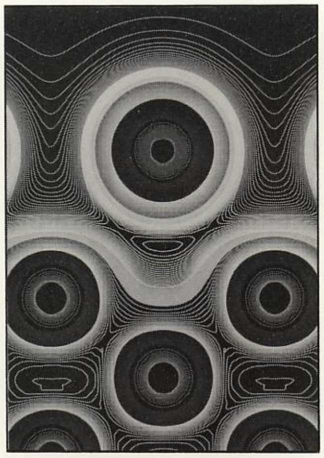

Contours of a different color

The charge density contours for an atomic overlayer of cesium on tungsten are illustrated here. Using a CRAY at Livermore for both computation and graphics generation, physicists have investigated the electronic structures of these materials and obtained results impossible to determine analytically. For further explanation of their findings and more about computational physics, see page 2.

CRAY CHANNELS

A Cray Research, Inc. publication Volume 5, Number 4

CONTENTS

FEATURE ARTICLES

2 Computational physics on the CRAY: a happy union of software and hardware

Arthur J. Freeman, Henri J. F. Jansen, Erich Wimmer, all of Northwestern University

Using CRAY power and state-of-the-art computer programs, our understanding of computational physics is expanding rapidly.

8 Everything you wanted to know about gallium arsenide (but didn't dare ask)

This primer offers a look at the semiconductor material that holds promise for faster computing.

12 Cray addresses U.S. House

In June, John Carlson, Executive Vice President of Cray appeared before two House subcommittees. Excerpts of his speech are contained here.

16 Supercomputing conference at Los Alamos

A summary of the five-day conference is punctuated with comments by John Rollwagen, Cray Chairman.

18 Los Alamos turns forty

We celebrate Los Alamos' 40th birthday, gaining an appreciation for its history and ongoing computational requirements.

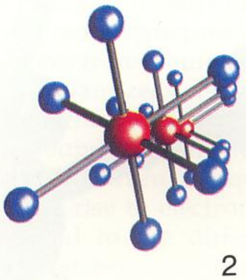

2

REGULAR COLUMNS

22 **Corporate register**

24 **Applications in depth**

25 **User news**

Staff

Editor	Carol Smith
Art/Design	Kathleen Van Kirk Diane Thomas

CRAY CHANNELS is a quarterly publication of Corporate Communications, Cray Research, Inc., P.O. Box 154, Minneapolis, Minnesota 55440. **CRAY CHANNELS** is intended for users of Cray Research computer systems and others interested in the company and its products.

Requests for copies of **CRAY CHANNELS** may be addressed to the Corporate Communications Department of Cray Research at the above address. If you have a feature article or news item to submit for publication, please contact the editor of **CRAY CHANNELS** either by mail or by telephone at (612) 333-5889.

Computational physics on the CRAY: a happy union of software and hardware

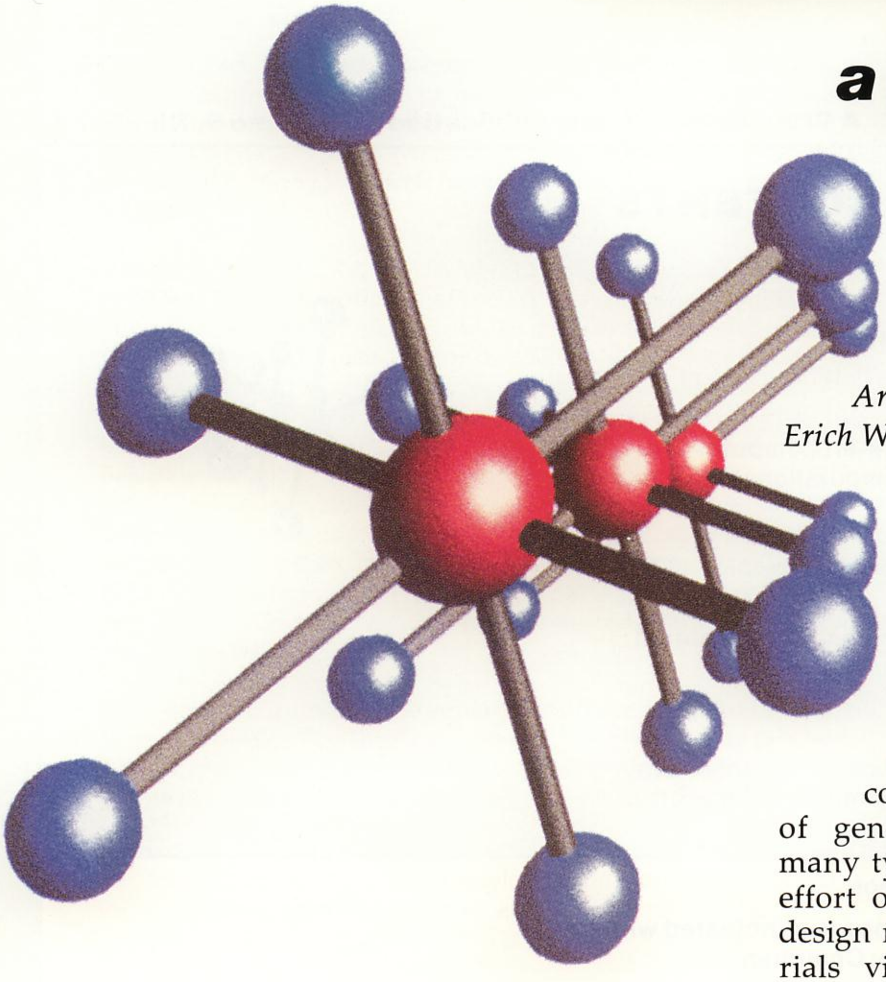

Arthur J. Freeman, Henri J. F. Jansen, and Erich Wimmer, all of Northwestern University

Computational physics has spurred major advances in all areas of physics. In particular, in condensed matter physics, our theoretical understanding of the structure of large molecules, complex solids, surfaces, and interfaces has been significantly impacted. These developments have, in large part, been made possible by the fruitful interaction between experimentalists working on "clean" systems, and theorists carrying out realistic numerical modeling of these systems.

A number of highly reliable state-of-the-art computer programs, such as those developed by the authors and their colleagues at Northwestern University, provide significant confrontation between theory and experiment. Several of these programs have been used in an ongoing collaborative research effort using the CRAY computers at the Magnetic Fusion Energy Computer Center at Livermore, California for over three years. The group at Northwestern has developed theoretical approaches and computational algorithms capable of generating accurate descriptions for many types of systems and structures. The effort of theorists and experimentalists to design new electronic and structural materials via computational physics is very important, relating directly to work being done at the academic, industrial, and national laboratories, and contributing to the thrust of the basic sciences.

Increased computational power at the supercomputer level will continue to be required for this class of problems. For instance, in Dr. Freeman's group's work, the algorithms do not make any gross physical or numerical approximations. Large matrices, up to 800 dimensional, have been diagonalized. Iterative procedures are used to carry out solutions of complex integral-differential equations.

We were interested in understanding the nature of these problems that can easily require tens of CRAY CPU hours to solve. In this article the authors explain the theoretical approaches and computational methods of condensed matter physics and the "explosive opportunities" which they feel CRAY supercomputers present.

Introduction

As fields of science advance, there is a natural evolution in theory from analytic to numerical solutions. Once satisfactory theoretical formulations have been established and their validity demonstrated for simple test cases, the next natural step is the application of the theory to realistic, i.e., more complex systems. At this point, large-scale computation on computers like the CRAY becomes necessary; one moves from investigating systems with a few degrees of freedom to those with many degrees of freedom. Thus, complexity is the key to understanding computational science.

The step from the simple analytical to the complex numerical domain is particularly interesting in physics, where theories such as relativity and quantum mechanics have revolutionized our understanding of the macrocosmos and microcosmos. Hence, the application of such powerful theories has greatly stimulated the development of computational physics.

By its nature, methods, and needs, computational physics is so different and separate from the two well-known and well-recognized branches of physics — analytic theoretical physics and experimental physics — as to constitute the third branch of physics. Basically, computational physics is the numerical investigation of complex physical systems; it permits theory to be extended well beyond analytic limitations. As with other areas of computational science, the development of this new field continues to be intimately connected to the existence and evolution of supercomputers.

Condensed matter theory is an area where computational physics has had a wide impact. This field includes such diverse subdisciplines as the electronic and geometric structure of perfect and imperfect solids, disordered and amorphous materials, liquid metals and alloys, surfaces, interfaces, and artificial (laboratory-made) exotic materials such as superlattices. Condensed matter theory ranges from direct applications of principles for device fabrication to complete formal theories to explain and predict new properties of matter. It is one of the rare areas where the science is as basic as science gets, yet the ideas developed are often easily applicable to practical problems. Since all these developments are based on the underlying quantum theory of matter, we describe this first before discussing numerical approaches and examples of applications.

The quantum theory of matter

The laws governing the interactions of electrons and nuclei in materials are the well-known laws of quantum mechanics. It is possible, in principle at least, to determine theoretically all the properties of matter — electromagnetic, mechanical, thermal, and chemical. In practice, however, there are formidable difficulties in deriving the properties of any but the simplest material systems consisting of only a few electrons and nuclei, just as there are for any but the simplest classical planetary systems. Interestingly, the motion of an electron in the quantum theory is described by an orbital in the way that the trajectory of a classical particle — planet or satellite — is described by an orbit. The fundamental difference between the two is that an orbital does not represent a trajectory but rather a distribution, which gives the probability that an electron can be found in a given region of space.

This distinction did not exist in the early successful semi-classical theory of atomic structure by Niels Bohr. He postulated that electrons could exist only in certain stationary or quantized energy states having stable orbits around the nucleus, just as planets have stable orbits around the sun. However, transitions of the electrons between the quantized electron orbits actually give rise to electromagnetic radiation with energy equal to the difference in energy of the two quantized states.

In the quantum theory of atomic structure, the electron orbitals are found as solutions of the Schroedinger wave equation called wave functions, whose absolute magnitude squared yields the position probability density (or orbital density). Thus, the orbital density for the lowest energy state (also called the ground state) of the simple hydrogen atom has a radial spatial distribution about the nucleus shown in Figure 1, which gives the probability of finding an electron at any distance from the nucleus. Whereas its most probable position — the peak in the distribution in Figure 1 — is precisely at the Bohr orbit distance, there is a strong likelihood of finding the electron at other radial positions. Of course, the probability of finding the electron any-

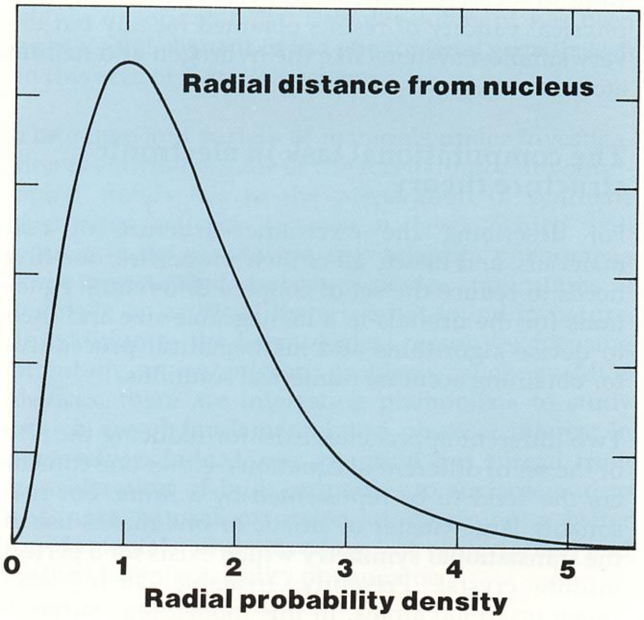

Figure 1. The radial charge density which gives the probability of finding an electron in a hydrogen atom at a distance from the nucleus.

where in space is unity, corresponding to there being just one electron. One may think of this probability density as a "swarm of one bee." Determining these probability densities for "simple" systems such as s atoms is important because it relates to how these systems will interact with each other to form molecules and solids.

The simple case of the hydrogen atom with only one electron was solved analytically by Erwin Schroedinger himself in the mid-1920's, demonstrating the correctness of his formulation of quantum mechanics. Unfortunately, even a system with two electrons outside a nucleus, such as the helium atom, poses severe difficulties if one attempts to solve the problem analytically. Helium is a "3-body" system (one nucleus and two electrons). These seemingly simple systems still defy exact analytical solution. However, very accurate results can be achieved numerically.

For more complex systems, such as many-electron atoms, molecules, and solids, even numerical techniques are challenged. The great complexity of finding these probability distributions from the orbitals, or the relative density of the charge distribution, lies in the fact that electrons have very strong interactions with each other and with the nuclei, so the orbital of any particular electron depends strongly on the position of all the other electrons. In order to determine these orbitals, one would have to solve a large set of coupled partial differential equations for the roughly 10^{23} electrons in the system. Clearly, such a direct approach is not feasible; one has to resort to symmetry arguments and to make approximations which reduce the set of equations to a manageable size. Until the advent of supercomputers like the CRAY, with enormous speed and computational power, the approximations themselves were so severely restricted as to cast doubt upon the physical validity of results obtained for any but the very simplest systems like the hydrogen and helium atoms.

The computational task in electronic structure theory

For describing the electronic structure of real materials, and hence, all of their properties, one first needs to reduce the set of coupled differential equations for the orbitals to a manageable size and then to devise algorithms and mathematical procedures for obtaining accurate numerical solutions.

Two different approaches exist for reducing the size of the set of differential equations: either one considers the solid to be represented by a finite, but reasonably large cluster of atoms, or one makes use of the translational symmetry which exists for a perfect infinite crystal. Presently, treatable cluster sizes range up to 100 atoms. In this model, the "surface" or boundary of the cluster has an important effect, which is proportional to the ratio of the number of atoms on the surface to that in the bulk. By comparing larger and larger clusters, this surface effect is scaled out and the results obtained will converge to those appropriate for a bulk material. Fortunately, it is possible (and presently necessary) to use theoretically derived correction terms to cancel approximately these important surface effects for even small clusters of atoms. On the other hand, the cluster approach is directly applicable for the investigation of a number of important problems, including superfine particles, whose strong catalytic effects are due to the large surface-to-volume ratio.

In the alternative model, the so-called energy band approach, one makes use of the fact that the size of a macroscopic solid is very large compared to atomic dimensions to justify the mathematical description of a solid as an infinite repetition of small, identical unit cells. Since the time average of the electronic orbitals in all these unit cells is the same, one only has to solve the differential equations for the orbitals in one of the unit cells. In the case of perfect crystals, the unit cell contains relatively few atoms; however, when the effects of impurities, vacancies, defects, or dislocations are included to investigate more realistic situations, the size of the cell rapidly becomes very large.

A major development in electronic structure theory was the discovery and formulation of density functional theory by Hohenberg, Kohn, and Sham in the mid-1960's. This is based on the important theorem that all ground state properties of matter can be derived from the charge density distribution. All material systems are considered as formed from atomic nuclei and swarms of strongly interacting electrons. The electrons are defined at any point in space by a collective probability density given by the charge distribution arising from all electrons in the system. The conceptually crucial aspect of density functional theory is this collective approach. From this and the so-called local density approximations, one can derive a new set of simplified decoupled differential equations, which is formally equivalent to the original set of equations of motion for the orbitals. The great advantage is that these new (local density) equations are completely determined by the charge density distribution and can be solved independently for each orbital.

Solution of the local density equations

Local density theory permits one to solve a relatively small number of decoupled partial differential equations to obtain the orbital charge densities. The catch in all this is that the orbital charge densities are required to construct the local density equations themselves. The way out of this dilemma is to set up an iterative procedure until "self-consistency" is achieved.

The basic idea in such an iterative cycle is to start with assumed approximate values for the orbitals and to use this input to decouple the differential equations. The solution of this simplified system of

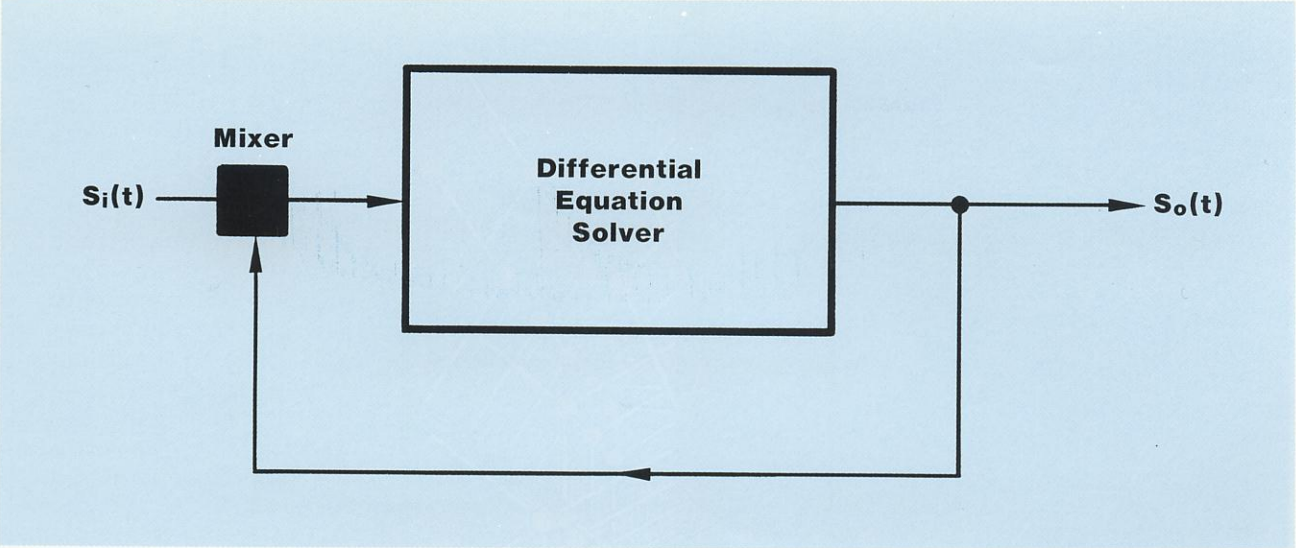

Figure 2. Feedback system used to solve the local density equations.

equations then provides us with new approximate values for the orbitals (i.e., charge density), and the iteration procedure is continued to self-consistency (i.e., until the input and output charge densities are identical within a prescribed precision). This iterative technique has a nice analogue in electronic circuits. To see this, one replaces the set of differential equations by some electronic device (a "black box") which takes an input signal $S_i(t)$ (input orbitals) and produces an output signal $S_o(t)$ (output orbitals) (see Figure 2). One is looking for a special value of the input signal, S_c, for which the output signal is also equal to S_c. Unfortunately, this device is an amplifier; any deviation from S_c in the input will give rise to a larger deviation of the output. However, the sign of the deviation is reversed and hence one can use standard feedback techniques to drive the device to the constant self-consistent solutions S_c (see Figure 2). In our case, this feedback is equivalent to some mixing of input and output charge densitities.

At this point, the set of coupled differential equations has been replaced by simple equations, plus an iterative feedback loop. The main numerical effort is in solving the differential equation; the feedback implies that this step has to be repeated many times. Both the number of iterations and the time per iteration depend on the complexity of the system under consideration. The ten or so iterations needed for a simple mono-atomic solid like aluminum increases to 50-100 for more complex systems. The time per iteration has ranged from ten minutes to more than one hour on the CRAY. The final results obtained are considered to be accurate, self-consistent solutions which are then independent of any starting assumptions. Obviously the complexity of the physical system to model dictates a massive computational requirement regardless of the numerical approach. CRAY systems are now being made more available to meet this requirement. Here multitasking on the CRAY X-MP will have a dramatic effect on the research capabilities.

Electronic structure of real materials

The last decade has witnessed dramatic advances in condensed matter theory, driven in large part by new and sophisticated experiments on high-purity materials which have been well and carefully characterized. In most areas of condensed matter theory, and particularly in electronic structure, these advances may be attributable directly to the close collaboration of theoretical and experimental research. Indeed, the new-found ability to apply fundamental theoretical concepts to real materials (rather than to simple model systems) made possible by utilizing the continued rapid development of computer power, has served to fill the increasingly urgent demand of experimentalists for theoretical interpretation of their data. Also, in some cases, these computational efforts can be used to provide data which would be impossible or impractical to obtain experimentally. This development has been an essential element in the phenomenal growth seen in this area of materials science.

The range and variety of materials under investigation is enormous. One of the recent important developing trends lies in the preparation of synthetic structures on the submicron level, which will permit, in the near future, new scientific phenomena to be investigated, and novel device applications to be made on artificial materials not found in nature. In addition to the technological reason for focusing attention on submicron problems, such as MOS devices, there are interesting phenomena to study which are of fundamental and physical interest in themselves. Indeed, our attention has spread from consideration of bulk properties to surfaces and to the next natural extension of obtaining a better fundamental understanding of interfacial and reduced-dimensionality phenomena.

Results obtained using the CRAY

We illustrate the power of these contemporary

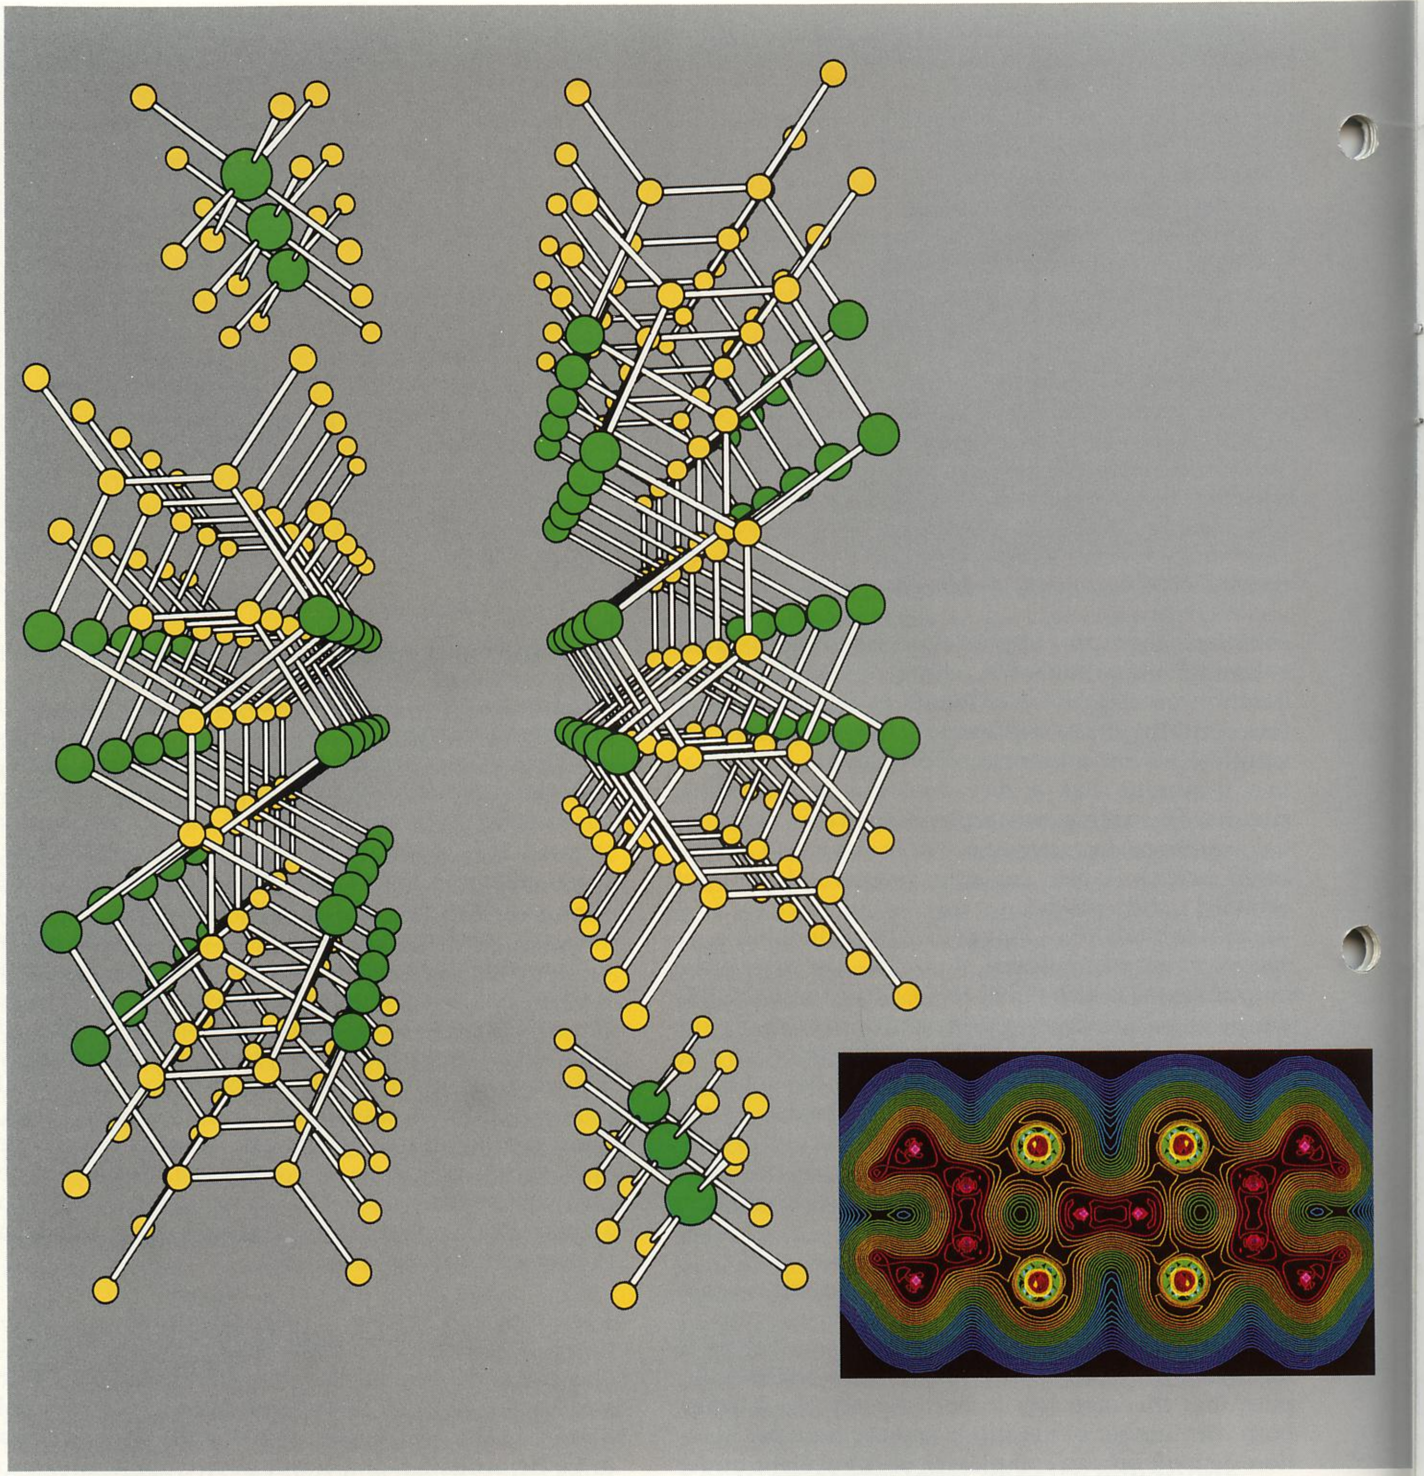

Figure 3a. A ball and stick model of $(TMTSF)_2ClO_4$*, the only known organic superconductor at ambient pressure. Conduction takes place between each TMTSF molecule in the zig-zag stack along the corridors (after Thorup et al.). Figure 3b. (Inset) Charge density contours for a planar molecule of TMTSF. Blue is low density, pink is high density. The hydrogen atoms are not shown.*

computational/theoretical methods by describing two examples of our recent work using the CRAY computer. In both instances, the power and speed of a CRAY were essential to obtaining solutions. In the first case, we consider the tetramethyltetraselenafulvalene (TMTSF) molecule, which is the structural building block of a recently discovered organic superconductor. Until now, this ability to carry an electric current with zero resistance had only been known in metals. In the crystal, these flat molecules are stacked in a regular array, which permits conduction electrons to move readily only along the corridor formed by the stack, as shown in Figure 3a. The self-consistent electronic probability charge distribution for the TMTSF molecule using the local density molecular cluster approach cited above is illustrated in Figure 3b. The origin of superconductivity in a salt of TMTSF is now being investigated by calculating the electronic structure of a stack of these molecules using an energy band approach—at a substantial increase in computational time.

Our other example arises out of our work (in collaboration with J. R. Hiskes and A. M. Karo of the Lawrence Livermore National Laboratory) dealing with the generation of negative H ions in particle-surface collisions on composite surfaces that are selected for their minimum work function. (The work function is the minimum energy required to remove an electron from a metal.) Negative ion sources are an important means of producing high-energy neutral beams as fuels for fusion machines. As a step toward calculating a precise value for the probability of production of negative H ions backscattered from selected composite surfaces, we have investigated the electronic structure of a cesium overlayer on a tungsten surface. We used the local density energy band approach and a newly developed algorithm which allows one to treat all electrons in the system without any shape approximation imposed on the charge density (or its resultant potential energy functions). The presence of the cesium surface layer causes a considerable change in the electronic charge density distribution at the interface (see Figure 4), and leads to a lowering of the work function by almost a factor of two, a result which is impossible to determine analytically. The calculation predicts a lower work function and the production of more back-scattered negative H ions. The production of negative ion beams provides a means for obtaining high-energy neutral beams that can be applied in the power sources of the future — fusion reactors.

Conclusion

Today's important applications of the powerful theories of physics owe much to recent developments in computational physics. Complementing the strides in numerical methods are tremendous advances in computer power. The availability of CRAY computers has made an enormous difference in the kind of problems that can be attempted. The results of this basic research provide exciting insight to many applications, ranging from the development of metal-oxide-semiconductor (MOS) devices to fusion energy research. Today, we have found a key — computational physics on the CRAY — that is providing the answers to many questions. We are convinced that as computational physics and supercomputers to evolve, the expansion of research on even more complex systems will become possible. □

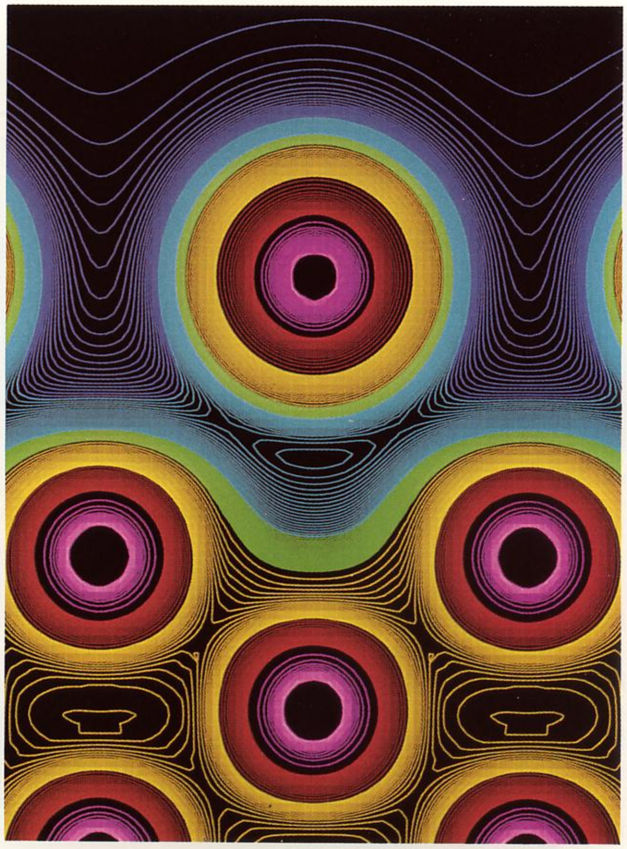

Figure 4. Charge density contours for an atomic overlayer of cesium on tungsten.

Acknowledgement

This article is based on a presentation made by Professor Freeman at the Cray Research 1983 Internal Technical Symposium.

About the authors

Having hailed originally from Poland, Arthur J. Freeman received both his B.S. and Ph.D. degrees from Massachusetts Institute of Technology in the 1950's. Since that time, he has held positions in both academia and government agencies. In 1967 he went to Northwestern University as Chairman of the Physics Department. Professor Freeman is a member of several professional and governmental committees and has received numerous honors including Guggenheim and Fulbright Fellowships (1970-71) and the Alexander von Humboldt Senior Fellowship (1977-78). He heads an active research group at Northwestern, has published well over 250 papers, and was recently appointed as Morrison Professor of Physics.

Henri J. F. Jansen has been conducting research in physics at Northwestern University since 1981 as a Post-doctoral Research Associate. Before that time he had been working in the Solid-State Physics Laboratory of the University of Groningen, Netherlands, where he obtained his Ph.D. in 1981. Dr. Jansen has received several honors, and in February 1982, Shell Research, B.V. in Holland awarded him their annual prize for his Ph.D. research. Subjects of his current research include valence fluctuations, total energy calculations for bulk solids, and magnetism and superconductivity.

Erich Wimmer began his physics research career in 1974 at the Technical University of Vienna, Austria where he received his bachelor's degree in inorganic chemistry. He was awarded his doctorate in theoretical solid-state chemistry/physics from the the same institution in 1977. Dr. Wimmer spent the years from 1979 to 1982 at Northwestern University. He is now an Assistant Professor at the Technical University of Vienna and spends four months of a year at Northwestern. His current research centers on the electronic structure of surfaces, atomic overlayers on surfaces, and molecules chemisorbed on surfaces.

Everything you wanted to know about gallium arsenide

(but didn't dare ask)

In some ways it appears that gallium arsenide (GaAs) will be to computing in the 1980's what silicon (Si) was to the industry in the late 1950's and early 1960's. As recently as a year and a half ago, little was heard about semiconductor materials other than silicon. But by now, most people are aware that gallium arsenide promises to play an important role in the next generation supercomputer.

Those familiar with Cray Research know that earlier this year, Seymour Cray announced that the CRAY-2 follow-on development is exploring the use of gallium arsenide integrated circuits. Toward that end, Cray Research has been moving ahead to develop GaAs circuits for use in the new system. In June 1982, Cray organized a development team for the GaAs project. Earlier this year the company signed a joint agreement with Harris Microwave Semiconductor, Inc. to develop gallium arsenide integrated circuits. (For some time, the Harris research group has specialized in the manufacture of GaAs products for communication and signal processing systems.) During the summer of 1983, the group of 13 moved into a new gallium arsenide research facility in Chippewa Falls, Wisconsin, just up the road from Seymour Cray's Riverside Project. New prototype GaAs wafers are already being tested.

GaAs background

In the computer industry, the use of gallium arsenide for semiconductors has been minimal, although the potential value of the material has been recognized for a decade. Because of the difficulties in working with GaAs and its extremely high cost, most U.S. firms conducted minimal GaAs research throughout the 1970's. A single square inch of an unprocessed GaAs wafer easily costs $30, and it is not unusual for vendors to market very pure gallium arsenide for $45 per square inch. It is interesting to note that the two areas where gallium arsenide has been successfully used are the satellite communications industry (as mentioned earlier) and digital LED device manufacture, before that technology was replaced with digital LCD's.

In 1970, the first GaAs transistor was made, and in 1974 Hewlett Packard (HP) introduced the first gallium arsenide integrated circuit. At that time, HP researchers were plagued with problems such as contamination of crystals, uneven consistency of the GaAs mixture through the crystals, and unevenly shaped crystals. Now, nine years later, much more is known about gallium arsenide and engineers are able to create three-inch round wafers (a feat unheard of in '74), limit the impurities, and minimize the defects in the crystals. Alas, the cost of GaAs is still very high.

GaAs characteristics

What, you may ask, is it about gallium arsenide that capitivates computer scientists? In a word — speed. GaAs integrated circuits can operate at speeds up to five times that of silicon-based equivalents. But because of the inherent difficulties associated with GaAs, scientists have contented themselves with squeezing higher speeds from silicon circuits by packing more and more devices on each chip. However, size reduction is approaching its practical limits in conventional device fabrication. For instance, circuit networks created with conventional contact photolithography have conductors only 1 μm wide and 1 μm apart (a human hair is about 100 μm wide). Beyond the limitations of proven fabrication technology, other major improvements are limited by silicon's properties.

Looking at the Periodic Table of Elements, particularly columns 3a, 4a, and 5a, helps one understand what makes GaAs such a good semiconductor material (see Figure 1). All of these elements, but especially those in column 4a, exhibit a strong binding force between electrons. This is important in semiconductor work. By selectively adding different impurities to these elements, many semiconductor permutations can be derived. By adding different dopants to silicon, for example, a variety of different electrical characteristics are derived to create transistors, resistors, and other types of devices. The combination of elements from columns 3a and 5a provides about the same binding strength as those in column 4a and the potential for making semiconductors. Several of the 3-5 compounds tested have shown properties that make them useful for high-speed circuits.

Periodic Table of Elements

3a	4a	5a
B	C	N
Al	**Si**	P
Ga	Ge	**As**
In	Sn	Sb
Tl	Pb	Bi

Figure 1. Relative positions of gallium, arsenide, and silicon on the periodic table.

Phil Gerskovich, Senior Design Engineer for Cray Research, explained, "The primary characteristics worth considering when choosing materials for circuit design are electron mobility and insulation. The electron mobility of GaAs is about four to five times faster than that of silicon. Not only that, but these electrons can be moved at high speeds and at lower voltages than silicon. That factor is very important in designing supercomputers that demand tremendous amounts of energy."

Phil went on to explain that because of gallium arsenide's insulating properties, the switching time between devices is extremely fast. "It turns out that if you are attempting to run signals over a substrate, the capacitance is lower with gallium arsenide circuits. Another advantage of gallium arsenide's insulating properties is that the need for special insulating designs is eliminated, and at the same time superfluous interaction between devices is minimized."

The gains of using GaAs do not come without their costs, and literally, the compound's high cost is of some concern. Whereas a silicon wafer will cost about $6, the same GaAs wafer is worth about $300. However, most of the integrated circuit's value is in its processing and packaging, not in the raw material. But researchers are hopeful that as gallium arsenide becomes more readily available, the costs will come down as they have with silicon.

Gallium arsenide's brittleness is another troublesome characteristic. Engineers at Cray practice their material handling skills before they ever come near an actual piece of GaAs.

And unlike silicon, with GaAs one must be concerned with the behavior of two elements, not just one — and those two elements may behave very differently in the same environmental conditions. Later on we will see how this impacts wafer processing. But in spite of all this, computer scientists are very excited about the emerging semiconductor material. Figure 2 illustrates the major differences between Si and GaAs semiconductor materials.

Characteristics	Silicon	Gallium Arsenide
Electron mobility	$800\ cm^2/vs$ (at $1 \times 10^4 v/cm$)	$5000\ cm^2/vs$ (at $3 \times 10^3\ v/cm$)
Substrate	conducting	insulating
Cost	$\$1.00/in^2$	$\$30.00/in^2$

Figure 2. Comparison of gallium arsenide and silicon semiconductor characteristics.

Gallium arsenide MESFETS

Essentially, a MESFET (MEtal-gate Schottky Field Effect Transistor) consists of a small area on a semiconductor crystal that has been selectively doped with impurity atoms and provided with three parallel conductors. Ironically, silicon is a major dopant in GaAs circuits, and arsenic a major dopant in Si circuits. Either of the two wide outer conductors can be a low-resistance source of electrons (the second conductor acts as the drain), and the narrow electrode between them serves as the gate (or control) electrode. At low applied voltages, the current is proportional to the voltage difference between the drain and the source. Both the source and the drain are in good ohmic contact with the crystal, but the gate electrode makes rectifying contact, that is, gate current would flow through it only if it were given a positive potential. The current between the drain and source then is controlled by the gate voltage.

Today, there are two types of GaAs MESFETs that hold promise for the production environment: depletion- and enhancement-mode devices. At Cray Research, it now appears that the depletion-mode device will be the first device made in any significant quantity. Of the two devices, the depletion-mode device is less risky to construct because of its simplicity. It is distinguished by the fact that it normally exists in an "on" state, meaning that current is flowing through it. Consequently, the device requires higher voltages for operation.

On the other hand, the enhancement-mode GaAs device is a more precise device, but more complicated to fabricate. The enhancement-mode device normally exists in an "off" state, akin to typical silicon circuits. Less power and voltage swing is required to operate enhancement-mode GaAs circuits, thus making the job of cooling a computer system much easier than with the former. Certainly within the next two years our understanding of enhancement-mode devices will increase and hold more promise for the production environment. Some Japanese companies have already made good progress in working with these devices.

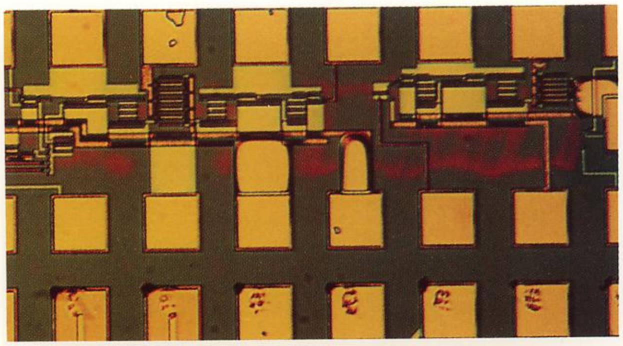

Cray-designed test gallium arsenide depletion-mode transistor.

Silicon integrated circuits are complex structures when compared to a gallium arsenide circuit (see Figure 3). Whereas a Si chip is made of many layers of impurities, GaAs chips include only one or two impurity regions. About half the number of masking steps occur in GaAs fabrication. And unlike silicon, gallium arsenide devices must undergo a complex metallization process. Metallization is important in determining the function of a given circuit element in the realm of GaAs technology. For example, one metal may work well with GaAs, another may be needed to transfer the signal, and a third metal may be needed to connect the two other metals.

Gallium arsenide processing

To construct MESFETS, one would have little notion of the intricate procedure involved in translating a basically simple concept into a working device. Special washes, rinses, etching processes, ion implantation, annealing, and above all, extreme cleanliness contribute to making a successful MESFET. The slightest change in operating procedures or chemicals can increase the rejection rate from a few percent to 50% or more. It is expected that early GaAs circuits will use small-scale integration to ensure higher yields.

One reason processing is so critical is that the semiconductor's fixed surface charge density is a function of the mechanical and chemical history of the chip's surface. Hence, small changes in processing can produce big changes in a circuit's operating characteristics.

Silicon bipolar integrated circuit:
Complex structure, simple metal connect

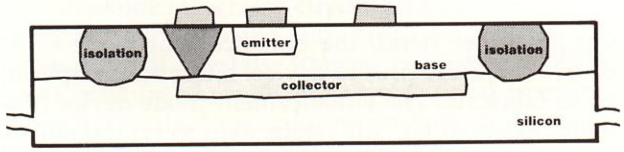

Gallium arsenide integrated circuit:
Simple structure, complex metal connect

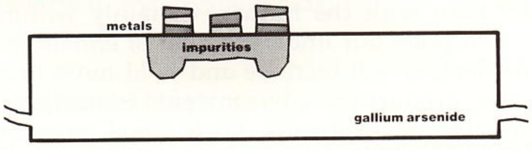

Figure 3. Basic Si and GaAs circuit layout.

GaAs processing involves ion implantation. Ion implantation is a way that dopant materials are embedded into a semiconductor material; in the case of GaAs semiconductors, it is the only way. The process requires that the wafer be placed in a vacuum chamber in which impurities are accelerated in an electric field. The impurities move rapidly and upon striking the wafer surface, the atoms are embedded in the gallium arsenide wafer.

Unfortunately, when the ions are accelerated in the electric field, they tend to crash into the wafer and actually damage it. Recall that GaAs is a very brittle material. That damage must be repaired. With silicon, repairs can be accomplished very nicely by placing the doped wafer into a furnace heated to 850 °C to 950 °C, and regrowing the Si crystal. However, when the same technique is used with GaAs wafers, the arsenic tends to evaporate from the wafer before the crystal is regrown. Since this results in a defective wafer, different annealing techniques are used. One method called "capless annealing" involves pressurizing the furnace with arsine gas so that the arsenic on the wafer cannot migrate. Another method of preventing the arsenic from evaporating is to place a cap of silicon nitride over the wafer. This process is called "annealing with an encapsulant."

Summary

Development of new technologies such as gallium arsenide are often spurred by necessity to meet Cray Research's standards for the next-generation supercomputer. And while silicon will always be important in the world of semiconductors, brute force supercomputers require much faster circuits. Realizing that the performance of silicon circuits has just about been maximized, scientists are turning to other semiconductor materials like gallium arsenide. Even though GaAs research has not been a priority for many companies, important strides have already been made in understanding and practical use of the material. Researchers at Cray Research are confident that the next-generation supercomputers can be made with GaAs circuits. Because Cray is in the business of defining the supercomputer standards, it really comes as no surprise that the company is involved in GaAs development. And it can be expected that someday soon, gallium arsenide semiconductors will become common in this specialized industry. □

Acknowledgements

This article is based on a presentation made by Phil Gerskovich, Cray Research, at the 1983 Cray Research Internal Technical Symposium. Thanks for their help in the preparation of this article are extended to Phil and Steve Nelson, Director, Advanced Projects for Cray Research who, began Cray's GaAs development work in 1982.

Gallium Arsenide's

special properties

To understand the special properties of gallium arsenide (GaAs), one must know something about electron transport in semiconductors. The key parameters are bandgap and the semiconductor's electron mobility and excess-carrier decay time.

In a simplified view, all the outer electrons surrounding the nucleus of an isolated atom are trapped in valence states (that is, they have nowhere else to go). If the atom is incorporated into a crystal, the overlapping of the attractive potentials of adjacent atoms creates additional states (conduction states) in which electrons can move in response to electric fields. In describing mass motions of electrons, we can group the many electronic states into bands — a valence band and one or more conduction bands.

In a metal, the overlap of the conduction and valence bands creates a plentiful supply of electrons for current flow. In an insulator, a large energy gap (the bandgap) exists between the valence and conduction bands, no electrons cross over, and no current flows. In a semiconductor, the bandgap is narrow, various common processes can boost electrons across it, and applied voltages can affect the motion of the electrons. This ability to change resistance in response to a signal is what makes semiconductors so useful.

In most semiconductors, there are multiple conduction bands. The principal conduction band (the one containing the largest number of mobile electrons) is the one that is lowest in energy. The conduction band may be thought of as a surface in electron energy-momentum space, and the electron mobility is strongly related to the curvature of the band surface at the state occupied by an electron.

So far, we have been describing electron transport in pure crystals. It is possible also to make a semiconductor material more conductive by doping it with selected impurities that contribute charge carriers.

One advantage of gallium arsenide over silicon (Si) as a semiconductor is that is has a larger bandgap (1.43 eV compared to 1.12 eV). What this really means is that the upper-end of the speed limit of the electrons as they race through GaAs circuits is much faster than in Si. And even more importantly, the force with which one is required to move these electrons is much lower than with silicon. In the case of a low electric field of 1 kV/cm (obtainable at 0.1 V across a gap only 1 m wide), the electron velocity in high-purity silicon is 15.5 km/s. In high-purity gallium arsenide, it is 85 km/s.

In the case of high electric fields (above 3.5 kV/cm), the differences between silicon and gallium arsenide are even more striking. In silicon's single conduction band, the electrons become "hot" (depart from thermal equilibrium with the crystal lattice) and begin to excite phonons, dissipating some of their energy. Instead of going faster in proportion to the field strength, they tend to approach a limiting velocity.

Gallium arsenide has two conduction bands separated by an energy gap of 0.38 eV. Electrons drift much faster in the lower conduction band than in the upper band, where they drift relatively slowly. An energy gap of 0.38 eV is almost 15 times as large as the thermal (300 K) excitation energy of 0.0259 eV. Hence, few electrons move into the upper conduction band by accident. It takes energy to put them there. Also, the conductivity of a gallium arsenide circuit can be increased in selected spots, and the rest of the crystal remains almost insulating. This greatly reduces parasitic currents that do nothing but generate unwanted heat and slow the circuit. Consequently, switching time between devices on gallium arsenide circuits is very fast requiring little power to go from point A to point B. Another advantage of GaAs's insulating characteristics is that it simplifies the processing of circuits since additional insulation between transistors is not required.

In addition to being narrower, silicon's bandgap is indirect. A conduction electron cannot decay, that is, drop directly back to the valence band; to conserve momentum it must simultaneously interact with a crystal lattice phonon (a quantum unit of lattice vibration). Gallium arsenide's bandgap is direct, enabling electron decay without a phonon. Disposing of excess electrons, as at the end of a pulse, is therefore much faster in gallium arsenide than in silicon. □

Acknowledgement

"Gallium arsenide's special properties" is based on text originally appearing in Energy and Technology Review, *"Predicting integrated-circuit performance," Lawrence Livermore National Laboratory, August 1981.*

United States
of America

Cray addresses U.S. House

WASHINGTON D.C., JUNE 14 and 15, 1983

Prompted by the findings of the Lax Report, the U.S. House subcommittees on Energy Development and Applications and Energy Research and Production called hearings this year to investigate the current status and future needs of computers and their role in energy research.

The hearing charter explained:

Japanese industry and government have designed two cooperative research and development programs that are intended to result in a 100-fold improvement in large-scale computer performance. If successful, they could dominate the international markets for mainframe computers within 20 years. These Japanese programs have focused the attention of Congress, the Administration, and U.S. computer manufacturers on the technical, national security, and trade implications of maintaining dominance in computer technology......

It has become clear that hardware, architectures, language, and software have highly interdependent roles to play,.......and that tradeoffs exist depending on uses. Thus, what works for a home user or in an automated factory will not work for the materials scientist or nuclear weapons design physicist. For mass markets and robotic applications, the most efficient software may well depend on research in artificial intelligence. For the scientific users, research in applied mathematical analysis will be of critical importance. The consequences for the scientific user will be slow in coming and will depend ultimately on spin-offs from more commercial development unless the Federal Government becomes a major partner in new computer development. The Department of Energy (DOE) has a vital interest in the widespread availability of supercomputers.....In certain areas scientific progress has been paced by the availability of supercomputers..... Thus, the DOE has a critical interest in making sure that access to and development of new supercomputers is not only timely, but that the development of supercomputer software and mathematical analysis is an integrated part of its research and technology programs.

Six speakers from the computer industry and government agencies that have large-scale computing needs were invited to address the group on Capitol Hill. Each participant was asked to address the issues regarding the current and future status of supercomputing in the United States, the government's role in fostering such development, and scientists' and engineers' research needs. On June 14, 1983, John Carlson, Executive Vice President of Cray Research, spoke to members of the U.S. House of Representatives. We thought you might be interested reading excerpts from John Carlson's presentation addressing these concerns.

Since it was founded in 1972, Cray Research has grown to over 1400 employees located across the United States and in five foreign countries. By the end of 1983, we will have added two more countries — Canada and Sweden — to our foreign-installed base. We currently manufacture four major products, including two types of central processors, the CRAY-1/M and the CRAY X-MP. The CRAY-2 is in final stages of development and will be available in late 1984. To date, we have installed 52 CRAY-1 supercomputers, approximately 80% of all current generation supercomputers worldwide. By year-end, we expect to install 16 new systems, including six of the new CRAY X-MP's.

Importance of supercomputers in research

Almost 40% of our systems are used to support some type of energy research. Most of these systems support energy research conducted at DOE National Laboratories and similar facilities in Germany, France, and the United Kingdom.

A second area of energy research for which supercomputers are the key is support of exploration for and development of petroleum resources.....Using advanced modeling and seismic signal processing techniques, oil companies expect to demonstrate dramatic dollar savings in exploration and reservoir management.

In addition to direct energy research, there are a number of related opportunities which are highly dependent upon current and future generations of supercomputers. One of the most important areas is computational fluid dynamics. In addition to the significant impact upon the entire process of engineering design, recent results from the use of the CRAY-1 by the aerospace community have demonstrated the potential for designing air frames with significant savings in total fuel consumption, which translates into direct savings to both the public and private sectors and preservation of national fuel reserves. We expect the impact of these savings to grow as supercomputers are introduced into the automobile and shipbuilding industries.

The cost trade-offs (of supercomputers) are dramatic. If an oil company can be spared drilling a single dry hole, the entire supercomputer installation will be paid for several times over. If a major aircraft or a new automobile model can be made as little as two or three percent more fuel-efficient, the savings to the nation in domestic fuel reserves could be staggering. We have not even begun to reap the potential benefits from the use of supercomputers in biotechnology, medical health-care delivery systems, socio-economic modeling, and a vast array of operational defense problems.....Even in the area of artificial intelligence.....the payoff in certain areas could be substantial if the (artificial intelligence) algorithms could be developed in a reasonable amount of time.......

Supercomputers have made significant and well-documented contributions to energy and general scientific research. However, we have only begun to realize the potential impact that advanced computing could have upon our nation's economy and security. The results we have obtained using this generation of computers have spawned new methods fundamental to emerging technologies.....Access to state-of-the-art systems will be key to advances in technologies of societal and political importance, including biotechnology and medical research. Recognizing these facts, other western countries are moving to capitalize upon the potential that supercomputing offers.......

Continued world leadership a priority

The world is on the verge of a major computing revolution as supercomputing is applied across a broad industrial and scientific base. The only question at this juncture concerns which nation will choose to lead this revolution.

As we approach the environment of the year 2000, we should not be surprised that Japan, France and the United Kingdom have all recognized the importance of supercomputing to the extent that they have initiated national programs of varying scopes to deal with the issue. Many have questioned why the Japanese would enter supercomputing since their normal strategy is to target industries with broad commercial significance. While it is possible that the venture into supercomputing has been more for prestige, I believe that history will record the venture as a very shrewd one. When Fujitsu and Hitachi announced their supercomputers last year, they projected deliveries to approximately 30 domestic customers in the first 18 months, or more new customers than Cray Research expects worldwide during the same period. At this time, there are only two current generation supercomputers in Japan. If these Japanese manufacturers are successful, we can expect to see a computing revolution throughout Japanese industry that could further threaten U.S. industrial competitiveness.

As these facts indicate, the importance of leadership in supercomputing technology goes far beyond already well-established links to defense and national security.....Leadership in supercomputing technologies, and in the aggressive implementation of that technology, could play a significant role in the future of both our basic industries and new high technology frontiers.

The United States, as a nation, and Cray Research, as a private enterprise, have demonstrated their ability to achieve world leadership. But leadership does not guarantee success. True success is sustained achievement. Our challenge is to ensure that we can continue to keep that leadership........We must define the problems to be solved and focus upon the solutions rather than the competitors......Cray Research, and we believe, the nation, are ready to meet the supercomputing needs of our research and industrial community. Today, however, other terms of the formula — national economic and trade policy, semiconductor availability, and domestic supercomputer policy — have risen in importance. With the serious challenges from the Japanese, British and French, many of our efforts hinge upon technology-related economic policy.

Need for technology-oriented economic and trade policy

Since we (Cray Research) devote large financial resources to development, our efforts have benefited from the recently enacted temporary research and development tax incentives. These incentives need to be strengthened and enacted permanently.

Because our financial success and our position of world leadership are dependent upon foreign sales, we are very sensitive to changes in export administration and West-West technology transfer policy. We need a consistent, well-managed and reasonable policy which will protect national security without harming our national economic assets.

Need for advanced semiconductors

Lack of high-speed semiconductors is currently the Achilles' heel of the U.S.supercomputer effort. Product development projects have been delayed or even suspended due to the lack of new high-speed, reliable semiconductor parts.......We are concerned that we will have to rely upon the Japanese semiconductor industry for the special chips requiredto build our computers. If this happens, U.S. leadership would begin to depend upon the willingness of Japanese suppliers to sell us the most advanced chips at the same time that these chips are available to our Japanese competitors. Our second concern with this situation is if we buy our logic chips from the Japanese, we will be forced to transfer to our supplier the proprietary mask design which reveals exactly how we will be using the chip. This type of technology transfer could be detrimental to both U.S. competitiveness and national security.

Finally, we are concerned that the government sponsored research is not addressing the semiconductors most useful to high-speed supercomputers. We are not hopeful that either the VHSIC or the DARPA gallium arsenide programs will yield any advantage to the U.S. supercomputer industry.

Having addressed what we consider to be the primary impediments to U.S. leadership in supercomputing, I will now turn to the final area, that of the Federal role in fostering supercomputer development.

The Federal role in fostering supercomputer development

We are currently in an exciting phase at Cray Research. By the end of next year, we will have delivered a major new product each year for four years.....In 1984, we expect to deliver the first CRAY-2 system, which will usher in a number of new technologies required to meet our goals for the second half of the decade. We are very encouraged with the current state of supercomputer development. If the market continues to be robust, and if some of the software and applications difficulties can be overcome, it is not overly optimistic to assume that the system which Cray Research delivers in 1990 will perform at speeds of at least two orders of magnitude faster than the CRAY-1; memory sizes are expected to increase at the same rate.

In the applications area, we are not quite as encouraged. The record of the last five years indicates that in environments other than the Department of Energy laboratories, the U.S. research and engineering community has been slow to implement supercomputing technology. Even the defense and aerospace communities, which have in the past aggressively applied new computers, have lagged seriously behind their counterparts overseas. Since the success of new computing technologies is heavily dependent upon their application in a user environment, lack of aggressive use of new generations of supercomputers could impede attempts to speed up the development cycle.

In the software area, there is a need for parallel processing breakthroughs in software technology.....Now that the CRAY X-MP is available, it is possible for researchers to simultaneously investigate parallel processing and high speed vector processing. Unfortunately, unless action is taken in the near future, this unique tool will be largely inaccessible to the vast university research community.....

In the United States, supercomputers are largely inaccessible to the vast majority of defense, industrial and academic scientists and engineers. The picture is quite different in Europe and Japan.....

The Department of Energy laboratories have played a crucial partnership role in the development of supercomputers since Mr. Cray began his career......In many respects, this particular vendor-customer relationship has been a crucial key to the success of Cray Research's efforts in developing another generation of supercomputer beyond the original CRAY-1.....

Although we applaud the efforts of the Department of Energy and others to assist in supercomputer development, we are concerned that this zeal could result in efforts to attack the wrong problems. Specifically, I offer the following recommendations as you consider the proper course:

- Ensure that the DOE is sufficiently funded to pursue its charter. Using supercomputers, the DOE is able to leverage the taxpayer investment better than with any competing technology. Under no circumstance should computer equipment budgets be curtailed in favor of development projects or other research.
- Enhance the DOE's capability to acquire computers strictly for software and algorithm research. The most advanced system available should always be available for general research....
- Encourage and provide extra appropriations for university access under DOE sponsorship.....
- Encourage continued close partnership between universities, industry and the laboratories.

In conclusion, Mr. Chairman, I am pleased to appear before you as a representative of our company and our customers. Even more important, however, I am appearing as the representative of our technology and those who are most directly responsible for American dominance of this vital technology, Seymour Cray, the chief designer of the CRAY-1 and its predecessor systems; Les Davis, Cray Research Executive Vice President for Manufacturing, Research and Development, and the individual who has delivered the Cray machines to the users; and Dr. Steve Chen, the chief designer of the CRAY X-MP, which is currently the world's fastest supercomputer......

Although we may address issues of importance to individual companies or customers, the success of these hearings and any action which the Congress or Administration takes will ultimately be measured by how that action fosters an atmosphere in which these individuals, and others waiting to emerge, can flourish and deliver new generations of system to the open marketplace. If Cray Research can provide an environment where new creative talent can emerge, and at the same time provide an environment where an exceptional individual like Seymour Cray can continue to be creative, we will indeed have done something very special. We believe we have begun to reap the benefits of that accomplishment, and we are committed to helping our Nation to do the same. □

Supercomputing conference at Los Alamos

In August, Los Alamos National Laboratory and the U.S. National Security Agency sponsored a five-day conference on "Frontiers of Supercomputing" at the Laboratory. The conference was organized to discuss how the United States could maintain leadership in supercomputing, which is considered vital to the country's defense and technological progress. The 162 participants hailed from different segments of the high-technology arena. They were movers and shakers in the field — from a Nobel Laureate, to two chairmen of major American supercomputer corporations, to research policy makers from the Departments of Energy and Defense.

Among the speakers was John Rollwagen, chairman of Cray Research. CRAY CHANNELS spoke to John Rollwagen about the week's proceedings upon his return from the conference. This article recaps the larger ideas that came out of the conference: As participant and observer, Rollwagen offers an insightful perspective here.

What they said

No one disagreed on the "compelling need" for more and faster supercomputers. Participants felt that for leadership in research, for industry to be able to compete in international markets, for defense, and for a thriving economy supercomputers are absolutely necessary. The experts agreed on the urgency of the problems and the roadblocks that may be encountered once action is taken. However, they agreed less about how the problem might best be addressed. The unanswered question that remains is whether U.S. industry, government, and universities will take appropriate action to meet the challenge.

Many people felt industry needs government support to compete with other countries' government supported supercomputer projects. Their reasons were that development costs are high and the small market makes the potential for profits low. Additionally, others recommended that anti-trust laws be modified to encourage competition in the marketplace rather than in research. The net result would foster industrial cooperation in sharing research. Along the same line, the argument was made that industry must pool talent from competing companies on the Japanese model. On the subject of system development, experts emphasized that supercomputers must be developed as systems along with languages, software, and applications. Others added that government and academia must collaborate in the national interest, and that one possibility might be for government to provide universities with funds to acquire and gain access to supercomputers for software and hardware research. Virtually all concurred that development of high-speed components and new computer architectures are critical.

The U.S. government recognizes that its action or inaction in this area will have tremendous impact. U.S. Senator Jeff Bingman from New Mexico told the group that he saw the conference "as an attempt to define an industrial policy for this critical industry." He went on to add that, "Congress would like to be out in front for a change, helping our most vibrant industries compete in world markets, not bailing industries out of trouble." Keynote speaker, Bobby Inman, president of Microelectronics and Computer Technology Corporation, added, "The first requirement is a consensus on national strategic policy, and decisions on priorities."

One of those priorities must be to rectify the educational deficiencies in the United States. Nobel Laureate Kenneth Wilson from Cornell University argued that the chief barrier to market growth — software development — could be overcome on university campuses. Wilson said computer scientists on campuses were isolated from industry and therefore were not working on industrial applications for supercomputers. "Markets will develop when the ideas of computer scientists flow into industry," Wilson said.

Robert Cooper, director of the Department of Defense Advanced Research Projects Agency (DARPA), stated that the development of faster and more intelligent supercomputers is essential "to hold off the Soviet threat." DARPA will spend between $500 and $800 million on joint projects with universities and industry to develop "intelligent" computers and software systems. Cooper commented, "With present-day supercomputers the Navy can replace the decision-making of a team of 30 people who fuel, load, launch and land an airplane. With machines 100 times more powerful, we could make decisions for an entire carrier."

John Rollwagen

William Norris, chairman of Control Data Corporation, called for direct aid from the government to the computer industry and for changes in the antitrust laws to allow corporations to pool their efforts. Another speaker added that we can't afford to have people recreate the same technologies in many companies.

Jack Worlton, nationally respected expert in supercomputing, commented, "The American supercomputer industry is in danger of being overwhelmed by the Japanese. They have awakened a sleeping giant. Business-as-usual on our part won't work, not when we're in competition with 115 million Japanese whose leaders have a "determination to succeed." Mr. Worlton went on to say, "The Sputnik effect may keep the Japanese from succeeding. The United States will stay ahead if it responds to the Japanese challenge. Rather than react, we should have a vision of our own."

The Cray perspective

Worlton's call for "a vision of our own" epitomizes Cray Research's philosophy as espoused by John Rollwagen. While several major research efforts, including DARPA and MCC, are being organized in response to the technology challenge, Rollwagen believes that Cray Research's approach to supercomputer development will continue to be successful, as it has been over the past 11 years. That approach is founded in the belief that small groups of individuals working independently, develop more creative, innovative ideas than large consortium-type efforts.

"Quite frankly," Rollwagen said, "I don't know whether the Japanese approach to R&D will be successful in the United States. It may or may not work. All I'm saying is that Cray Research represents an alternative to the other projects that are starting up. We know that our approach to research works. It was successful in the the past and has continued to set the supercomputing standards for the world today. And we are about to redefine those standards again in the very near future."

Rollwagen added that he hopes that groups like MCC are successful, and to the extent that they are, he is hopeful that individuals between organizations will share ideas. As he explained, "There's a basic difference in that type of exchange versus the large group effort. For example, when people like Seymour Cray and Gene Amdahl meet, they share ideas about their work on separate projects that have entirely different goals. It is precisely that meeting of diverse thought processes that sparks the really innovative idea. On the other hand, large groups of people working on the same goal tend to create a uniformity and consistency in the thought process that can hamper creativity."

Rather than accept government money for supercomputer development, Rollwagen prefers to see Cray fund its research through the sale of systems. "Just give us a clientele that can afford to install and use our computers," Rollwagen said. "That way two objectives are satisfied — the systems are made available to those who really can use them, and Cray can fund development projects on its own." With regard to Cray's ability to fund R&D, Rollwagen pointed out that Cray Research already spends as much as or more than the Japanese each year on supercomputer R&D.

"Using supercomputers is much harder than making them," Rollwagen stated. "Applications and systems software are gating factors. Ken Wilson is right, we need the universities to help develop algorithms to use the power we already have available. There is always a lag time between the introduction of new high-powered hardware and the development of software that takes advantage of it. Providing universities access to supercomputers would help close the gap."

Regarding the fifth-generation supercomputer projects, Rollwagen speculated that it would be decades before an intelligent computer would be available. "Today's computers are left-brain machines. A right-brain oriented computer is far beyond where we are today. You know, Seymour has always been successful by taking one step at a time and working with elements that he knew he could use."

Rollwagen closed by saying, "The conference was very good because it brought very important issues to the forefront. We appreciate the government's concern. The real significance of a conference like this is that it heightens and nurtures people's awareness of supercomputer status and needs."□

Los Alamos turns forty

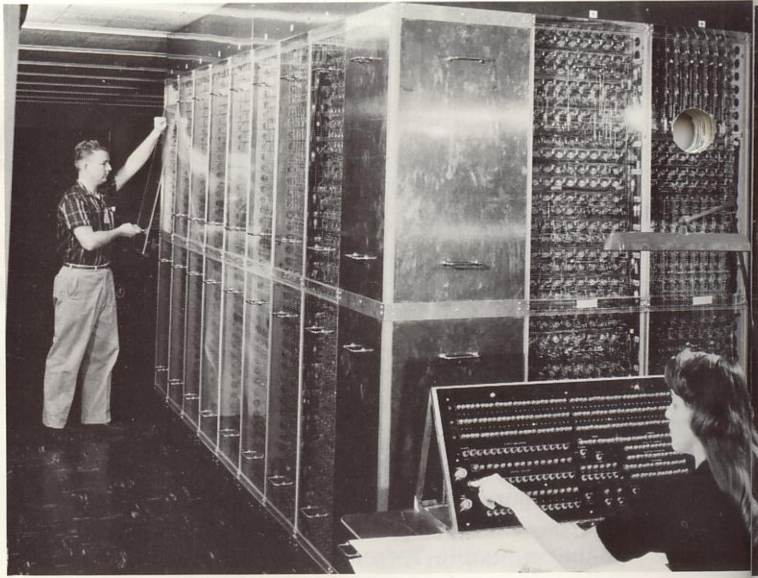

This year marks the 40th anniversary of the founding of one of the world's preeminent scientific laboratories. In 1943, Los Alamos National Laboratory was established to develop technologies for national security against the backdrop of World War II. Since that time, Los Alamos has always led the nation in scientific research that has had major implications for both the security of the United States and the way we live. In support of that research, the Laboratory has always relied upon and nurtured state-of-the-art computing capability, from the earliest desk-top mechanical calculators to today's CRAY computers.

The Lab's relationship with Cray has grown over the past seven years and today, Los Alamos houses more CRAY computers that any other of our customers; it really comes as no surprise that the Lab received the very first CRAY-1 system in 1976. In 1982 Los Alamos installed Cray's fiftieth system, and just recently, installation of a CRAY X-MP was completed, bringing the total of "CRAYs-in-residence" to five. Approximately 10% of all CRAY computers in existance can be found at Los Alamos National Laboratory.

In this article, we celebrate Los Alamos' anniversary by exploring its roots, computing history, and current computing capabilities, and consider what that has meant to the Lab's success over 40 years.

The first 40 years

The story of Los Alamos began in 1943 amidst uncertainty, urgency, and utmost secrecy. The United States and its allies were fighting World War II. The Lab's founding mission was to create a weapon that would end the war, and this mission was fulfilled with the successful explosion of the first atomic bomb in 1945.

For some years before Los Alamos was established, nuclear research had been conducted by various groups around the country. But the bombing of Pearl Harbor plunged the United States into war, and nuclear energy suddenly became more than a matter for scientific curiosity. The United States was in a race to develop a nuclear bomb before its foes, who, it was thought, may have had a head start.

President Roosevelt appointed General Leslie Groves to take charge of the American nuclear effort. The project was called Project Y of the Manhattan Engineer District. J. Robert Oppenheimer, who was at the University of California at Berkeley conducting theoretical research on the feasibility of a bomb, was named head of the project. He was to direct scientific experts such as Enrico Fermi, Niels Bohr, and John von Neumann. The University of California agreed to manage the program. A remote site in the mountains of northern New Mexico, the

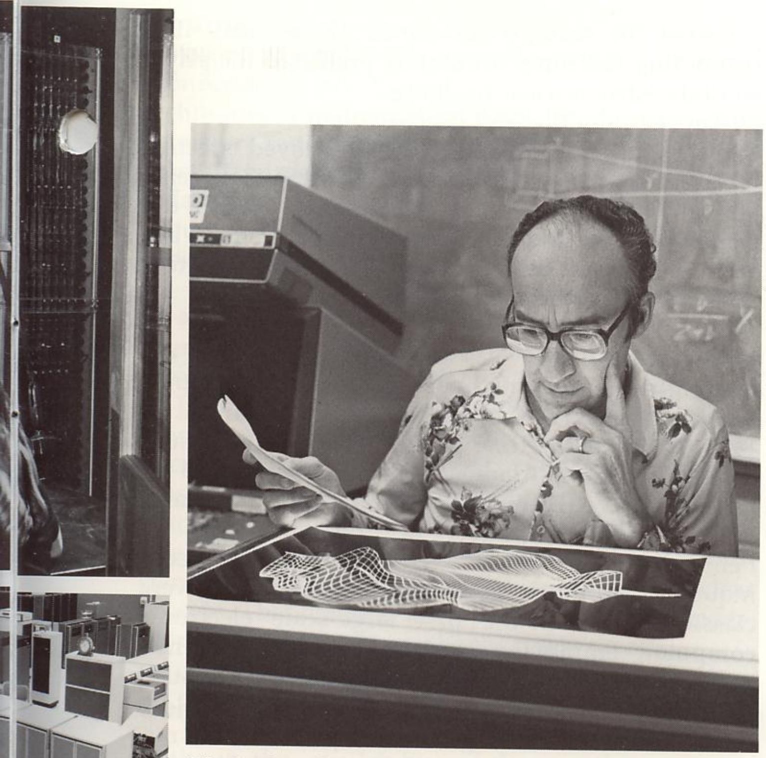

Clockwise from bottom left: Hastily constructed buildings housed wartime labs. MANIAC II, in use from 1957 to 1977, was powerful and easy to use because it included floating point arithmetic. A graphic image of an exponential function is analyzed. CRAY computers are an integral part of research.

Los Alamos Ranch School for boys, was suggested by Oppenheimer as the location for the secret project.

Originally, it was thought that about 100 scientists would work on the Manhattan Project. But by 1945, more than 3000 people were at Los Alamos. They worked in haste in makeshift laboratories with borrowed equipment in the isolated countryside. Living at Los Alamos was difficult; secrecy was the rule. Famous names were disguised, drivers' licenses were issued to numbers rather than names, and babies were born at Post Office Box 1663. Mail was censored throughout the war and telephone calls were monitored. But the people were young and flexible; the average scientist was only 29 years old.

Then on July 16, 1945, just 27 months after the original group of scientists made its way to Los Alamos, the first atomic explosion, named the Trinity Test, shattered the desert sky. "We knew the world would not be the same," Robert Oppenheimer reflected. His feelings after the explosion were shared by all who had worked with him on the project. Just as World War II had changed their lives, so their work at Los Alamos had propelled the world into a totally new era. Within a month, World War II was over.

Los Alamos' original mission was also over, and many, including Oppenheimer, left the Laboratory. Norris E. Bradbury took Oppenheimer's place as director, a position that he would hold for the next 25 years. However, it was not until late in 1946 that the Lab's fate was decided with Congress' establishment of the Atomic Energy Commission. A first priority of the commission was the stabilization and revitalization of the Los Alamos Laboratory. The Bradbury philosophy that the Lab should engage in "the most nearly ideal project to study the uses of nuclear energy" — which in effect meant it would be a laboratory combining ongoing nuclear weapons development with wide-ranging fundamental research — continued through Harold Agnew, the third director, and does so today under Donald M. Kerr's directorship.

Los Alamos continued advancement of fission and hydrogen weapons after the war, and in 1952 the first full-scale thermonuclear explosion was conducted in the Pacific. Until the Limited Test Ban Treaty of 1963, eight weapons test series were conducted at the Pacific Proving Ground. Beginning in 1951, other tests were conducted at the Nevada Test Site, where underground testing continues today. Los Alamos has remained at the forefront of weapons development: and about two-thirds of the weapons currently in the U.S.'s stockpile were designed at Los Alamos.

Multi-faceted research

Over the years Los Alamos has entered dimensions of science few imagined at its beginning. Today, energy and basic research are undertaken alongside national security projects.

The Lab began to broaden the scope of its research with the first stirrings of fusion research in 1947 which led to experiments beginning in the 1950's. The 1950's saw Los Alamos investigating nuclear propulsion for space vehicles; the Lab continues to work on many projects with NASA today, including the recent space shuttle. In response to the energy crisis of the 1970's, the Laboratory investigated new energy sources and refined existing ones, including nuclear, fossil, solar, and geothermal power. In basic research, Los Alamos remains a preeminent scientific institution. For instance, our understanding of particle physics has shed light on the origins of the universe, and technology enabling physicians to correct vision defects has evolved from studies of heat conducted at the Laboratory.

Los Alamos is engaged in researching radiobiology and radiotherapy, and in producing medical radioisotopes for biomedical research. The biomedical research program, originally begun to study radiation effects and oncology, has branched out into various aspects of molecular biology. Recently Los Alamos was selected to organize a national nucleic acid data bank where all known DNA sequences are stored. The Laboratory is also a national resource for flow cytometry.

Computing history

But what of the computers that have made all this research possible? Computer development and the success of numerical analysis owes much to research conducted at Los Alamos over the last 40 years. Numerical analysis at Los Alamos helped develop the whole logic of computers out of its wartime need for increasingly complicated calculations. Practically speaking, the age of modern computing, which has changed the world as much as anything, began at Los Alamos.

Numerical computation at Los Alamos began modestly with the use of various models of desk calculators back in 1943. But it soon became clear that this method would be totally inadequate as the machines often broke down and the flow of them back to the factory for repair was alarming to the scientists. (For secrecy's sake, repairmen were not sent to Los Alamos.) It was proposed that IBM's electromechanical business machines be considered for the group. The Laboratory went about the task of acquiring the IBM equipment: three multipliers, a tabulator, a reproducer-summary punch, a verifier, a keypunch, a sorter, and a collator.

At the same time, a formal hand-computing group whose members were primarily recruited from the wives of the local scientists was organized. Most of their work focused on the more complicated algorithms that were not tractable by the electromechanical machines. These problems required computation that took anywhere from a few minutes to a few days for one person to complete. In addition, the group executed calculations to verify that the electromechanical machines were operating correctly. For this, a production line was set up, with each member of the group computing one step of the calculation corresponding to one machine operation. This was turned into a lively competition, and the group was spurred to an early lead over the machines.

The business machines' highest priority involved integrating a coupled set of nonlinear differential equations through time from a prescribed initial configuration. The numerical procedure used a punch card for each point in space and time; a deck of cards represented the state of an implosion at a specific instant of time. Processing a deck of cards through one cycle in the calculation effectively integrated the differential equations ahead one step in the time dimension. This one cycle required processing the cards through about a dozen separate machines with each card spending one to five seconds at each machine.

The machines were fairly complex and the constant New Mexico dust frequently caused intermittent errors — at least one in every third integration step. Fortunately, the computational procedure was stable and insensitive to small errors. The machines were operated three shifts a day, and although primitive by today's standards, those early IBM computing machines accurately predicted the physics of the early nuclear reactions.

The simulation approach to design proved effective — many scientists first learned of numerical methods and computing doing this research. Before the war, only a few had considered such techniques, but after 1945 many began to apply the techniques to a wide variety of scientific problems.

John von Neumann was an important figure in extending numerical methods and applying computing capabilities to the research at Los Alamos. In 1945, at about the time of the Trinity Test, he suggested that Los Alamos problems be moved to the ENIAC (Electronic Numerical Integrator and Calculator) project at the University of Pennsylvania. Originally, the ENIAC, its techniques, capabilities, even its very existence, were classified. It was the first large-scale electronic computer, employing electron tubes rather than relays or mechanical counters. The machine bristled with 18,000 electron tubes, and a half-million solder joints, and filled one large room. It performed in minutes computations that took days on the business machines. The computer age began.

Los Alamos was responsible for designing one of the first supercomputers, the Mathematical Analyzer, Numerical Integrater, and Computer in 1952. The system was dubbed "MANIAC." In spite of its name, MANIAC was a major step in computer development because with it, technology fully progressed from electromechanical to electronic operation.

In 1953 research on programming languages and an operating system began at Los Alamos. Language development was critical since a computer's potential could not be realized without an easy way to communicate with it. In 1957, with MANIAC II, experiments on man-machine interactions were started in which the programmer could direct the computer during the course of a programmed calculation.

Whenever possible, Los Alamos has taken advantage of the computer industry's most powerful equipment offerings. The Laboratory has also collaborated with manufacturers such as IBM, Westinghouse, and even Cray in developing innovative systems. Many of the Lab's early systems came from IBM, including the IBM 701 and STRETCH systems. Later, in the 1960's the Laboratory turned to Control Data Corporation (CDC) for its state-of-the-art computers when it acquired Seymour Cray's CDC 6600 system. By the middle of that same decade, Los Alamos had installed Mr. Cray's next machine, the CDC 7600.

Numerical method and program development, optimization, and program conversion have been ongoing activities that provide reason for the powerful array of computing capability at the Lab. For instance, in the 1940's von Neumann and Fermi pioneered the development of Monte Carlo analysis techniques, a mathematical discipline for predicting transport radiation. It has been an extremely successful approach for studying radiation physics, and today Los Alamos has the largest Monte Carlo group in the U.S. for development of production-type computer codes.

In another example, Kenneth Duerre of the Engineering Design Group at the Lab completed conversion of a 370,000-line FORTRAN code for operation on the CRAY-1 this year. The effort took about three months to complete along with the help of another scientist. The program called ASPEN originated at Massachusetts Institute of Technology and is important because of its ability to analyze solid streams. The goal of Los Alamos' work with ASPEN is to efficiently study chemical, fossil fuel, and nuclear process. These studies are conducted for both external contractors and Los Alamos researchers.

Computing at the Lab

Without the availability of large-scale computer power, most of the work done at Los Alamos would have taken much longer to do or could not have been undertaken at all. The Laboratory's dependence on superior compute power was a given almost from its beginning.

Today, Los Alamos supports an unparalleled powerhouse of computing capabilities. The Computing Division, established in 1968, supports a vast and varied array of equipment including some experimental computers. More CRAY computers are installed under one roof at Los Alamos than at any other location in the world. Five CRAY systems are installed at the site, including a CRAY X-MP that began operation in November 1983. In addition, four of Seymour Cray's earlier large-scale computers developed while he was at CDC are still operating at Los Alamos.

Seventy percent of all the computing resources at Los Alamos are committed to the National Security programs — most of these programs are classified. Bob Ewald, Computing Division Leader, explained that the Inertial Confinement Fusion program is the second largest user of computer resources, and the Defense Nuclear Agency and the Nuclear Regulatory Commission are the next largest users. All but two of the CRAY systems are virtually dedicated to classified research projects.

The Los Alamos system has been tailored to address the needs of its 5,000 users (4,000 internal, 1,000 external), with advanced networking, internal and external communications, printing and graphics capabilities, and transportability of programs between systems. The graphics system at Los Alamos is among the most sophisticated in the world. Numerous output devices are available for the users. Ewald commented, "Each of our high-speed printers can print 10,000 lines a minute here. Our users generate two to three million pages of print and frames of microfiche output each month."

He went on to say, "We have automated our computer operations and network functions. For instance, it now takes one or two people to manage the supercomputer operations per shift, whereas before we had 20. Our common file system contains 13 trillion bits of information in 1.3 million files."

The most active files are stored on disk, larger files on an IBM 3850 mass storage system, and less active and larger files are stored offline. Files are moved between these classes of storage by a migration program that analyzes file activity and file size. Careful tuning of the file migration program has resulted in approximately 85% of all file accesses to disk; another 14% of the online accesses are to cartridges. Only 1% of all file accesses are to offline cartridges. The result is very cost-effective. The average response time to disk is five seconds and cartridges are accessed within 1.5 minutes. This means that within 1.5 minutes, 99% of the files requested are returned to one of the large computers for further processing.

Daytime use of the CRAY systems involves timesharing with the CTSS system. Ewald said, "On the one million-word CRAY computers 78% of CPU time is delivered to the users during the day. The four-million word systems typically deliver 85%; at night we deliver about 80% and 89% CPU-time per system, respectively. On the larger CRAY computers we get better CPU usage and we can run larger codes." Los Alamos-size jobs can easily require 70 CPU-hours on a CRAY.

And of course, the Computing Division is involved in computing research. Major areas of research center on parallel processing, man-machine interfaces, application of numerical methods, computer architecture, languages, artificial intelligence, and symbolic computing.

Summary

For 40 years, Los Alamos has been involved in addressing the most difficult scientific questions ever posed. It has been conscientious not only in answering those questions, but also in continuing to question and obtain even better answers. That type of effort has earned Los Alamos a solid reputation.

Out of necessity, advances in computer science have developed at the Lab as scientific problems have become more and more complex. Throughout its history, Los Alamos has been one of the largest and most sophisticated users of computer power in the world. Its computing capabilities today attest to that. Cray has been proud to provide systems for the Los Alamos computing network in recent years, and looks forward to working with the Laboratory in the future. □

CORPORATE REGISTER

Grumman orders a CRAY-1/M

Recently, Grumman Data Systems Corporation installed a CRAY-1 M/2200 that will be used exclusively by Grumman Aerospace Corporation. Grumman Data Systems now has two CRAY-1 systems installed at the Bethpage, New York center.

The new CRAY-1/M is being used in the development and testing of designs. Among the initial users of the system are engineers and scientists from flight test, aerodynamic design, mission analysis, flight control simulation, and electronic circuit design.

General Motors installs a CRAY

Cray announced that a two-million word CRAY-1/S was installed at the General Motors Research Laboratories during the fourth quarter of 1983. The computer is housed at GM's Warren, Michigan technical center and will be used primarily for research and development.

Cray installs X-MP at NASA

A CRAY X-MP/22 computer was recently installed at NASA-Ames Research Center in California through an arrangement with Technology Development of California (TDC). The system, which was operational by the end of 1983, replaces a CRAY-1/S that was installed by TDC for NASA two years ago. The new system is part of a major program to provide advanced computational capability for NASA-Ames.

CRAY X-MP delivered to New Mexico

In November 1983, Cray installed a CRAY X-MP/24 computer at Los Alamos National Laboratory. The Lab is engaged in scientific and engineering research and development and will use its newest CRAY computer for the same. John Rollwagen, Cray chairman, said, "The first CRAY system delivered to a customer was installed at Los Alamos in 1976. It's fitting that Los Alamos has now installed our newest system, a CRAY X-MP."

Westinghouse orders second CRAY-1/S

A CRAY-1 S/2300 system will be purchased and installed by Westinghouse Electric Corporation in the first quarter of 1984. The new system will supplement the capacity of the CRAY-1/S installed at the site in 1981.

The second CRAY computer will provide additional engineering and scientific computer resources for the firm's Energy and Advanced Technology group. Westinghouse engineers throughout the United States and Western Europe use the computer to design and analyze high-technology products, including turbine generators, nuclear reactors, electronic circuits and equipment for future energy sources such as fusion, solar energy and coal gasification.

Circuit Tools, Inc. — Cray sub formed

In September, Circuit Tools, Inc. was established as a subsidiary of Cray Research, Inc. The company is to be headed by John C. May, who has extensive background in the electronics design automation field. Circuit Tools will develop and market software programs for use in the design of VLSI circuits and provide related electrical engineering support. The company is located in Santa Clara, California. John May commented, "The close proximity of Circuit Tools to many circuit design and manufacturing companies will help us provide the timely service and support that those companies need."

The company initially will offer two software programs, CSPICE, a circuit simulation program and RSIM/C, a timing simulation program. CSPICE was developed at Cray Research by May and is based on SPICE, a University of California, Berkeley, public domain program. RSIM/C is based on RSIM, also a public domain program from the Massachusetts Institute of Technology. Both software packages feature distribution between the environments of the supermini or micro workstation and that of the supercomputer. The benefits to the user will be fast-turnaround verification of an integrated circuit design before it is committed to silicon.

NTT to order CRAY

Cray Research was recently informed that Nippon Telephone and Telegraph (NTT) intends to purchase a CRAY X-MP computer. NTT officials visited Cray Research in Minneapolis to present their letter of intent. Cray notes that the contract is undergoing negotiation. The system would be installed in Tokyo during the second half of 1984 pending export license approval and negotiation of a successful contract.

Over the summer, NTT acknowledged that it was opting for a U.S. supercomputer because the Japanese-announced systems either would not be ready as soon as needed or were not as long established in the market. A spokesman for NTT said the U. S. supercomputer manufacturers have far more sophisticated software already available. NTT needs a computer to run complex programs in structural analysis, dynamic analysis, and other basic research calculations.

Kaora Kubo, Senior Staff Engineer at Nippon Telephone and Telegraph's New York office translates conversation between Cray and NTT officials.

Seymour Cray extends contract

Cray Research, Inc. announced in September that the existing Design and Development Agreement with Seymour Cray has been extended by two years until December 31, 1987. Under this agreement Mr. Cray continues as independent contractor to the company, furnishing development work for its advanced computer systems. The extension evolved from Mr. Cray's decision to explore gallium arsenide circuitry to replace or augment silicon technology and may be extended further by additional projects.

In 1981 the company founder announced that he would step down as chairman of Cray Research in order to devote full time to the CRAY-2 project as an independent contractor. Since that time he has been developing the next-generation supercomputer from his laboratory in Chippewa Falls, Wisconsin. Mr. Cray has remained a director and a member of the executive committee of Cray Research.

APPLICATIONS IN DEPTH

Circuit design programs from Circuit Tools

The development of CSPICE, the circuit simulation program optimized for use on CRAY systems, was announced in CRAY CHANNELS Vol. 5 No. 3. The product is now available and fully supported by Circuit Tools, Inc., a new subsidiary of Cray Research. The company will soon add enhancements to CSPICE so that it will remain a key resource for state-of-the-art circuit designers.

In addition, Circuit Tools is offering timing simulation software used in obtaining an accurate overview of a circuit's performance. RSIM/C allows the designer to determine functional and approximate timing characteristics with greater accuracy than gate-level simulation. This can be done on larger circuits than can be accommodated by circuit analysis programs.

In the near future, two additional products — CSPICE Extensions and RSIM/C Extensions — will be available from Circuit Tools. Both of these programs will replace the input and output modules of the parent programs and add the capability to accomplish circuit and timing simulation tasks interactively. CSPICE and RSIM/C must be available and installed before CSPICE Extensions and RSIM/C Extensions, respectively, is added.

All four products operate on CRAY computers systems running the CRAY Operating System (COS) Version 1.11 or later, and VAX systems running VAX/VMS Version 3.0 or later. They also can be adapted easily to fit a variety of other manufacturers' equipment.

These products are fully supported and available from: Circuit Tools, Inc., 3130 Crow Canyon Place, Suite 310, San Ramon, CA 94583; telephone: (415) 838-8463.

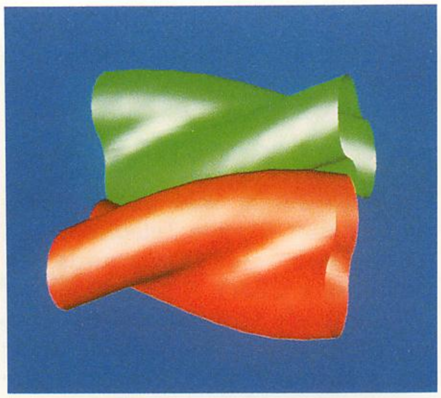

MOVIE. BYU CRAY-generated image.

MOVIE.BYU enhancements

A new release of the MOVIE.BYU computer graphics software has been converted for operation on CRAY systems. The new capabilities include transparency, multiple light sources (up to five), shadowing, dithering, and antialiasing routines.

MOVIE.BYU is recognized as a powerful set of FORTRAN programs for display and manipulation of data representing mathematical, topological, or architectural models whose geometry may be described in polygonal elements or contour line definitions. Many CRAY users at sites such as the Air force Weapons Laboratory, Kirtland AFB in New Mexico, and NASA-Ames Research Center are heavy users of the graphics system.

Those interested in additional information about MOVIE.BYU should contact: Hank Christiansen, Civil Engineering, 368 CB BYU, Provo, UT 84602.

AOS/MAGNETIC™ analyzes on the CRAY

The AOS/MAGNETIC Analysis Program that calculates magnetic fields and performance of a wide variety of magnetic and electrical products has been converted for operation on CRAY systems. It uses the finite element method, extending that technology to the analysis of electromagnetic devices. The program is useful to engineers maximizing electrical efficiency of devices such as: AS, DC, and permanent magnet motors, transformers, AS and DC solenoids, switches and relays, loudspeakers with ferrite, ALNICO, or rare earth magnets, alternators and generators, and busbars and transmission lines. Interactive pre- and post-processing and computer graphics aid the user in creation, verification, and modification of input data required for analysis and evaluation.

Persons interested in additional information about AOS/MAGNETIC should contact: A.O. Smith Corporation Data Systems, 8901 N. Kildeer Court, Milwaukee, WI 53209; telephone: (800)558-6980, or in Wisconsin: (414) 357-2956.

USER NEWS

Dutch Shell computing center opened

A CRAY-1/S computer was installed for Shell Research BV at the Koninklijke/Shell Exploratie en Produktie Laboratorium (KSEPL), the Netherlands, in April 1983. While the system has been operational for some time, the official opening of CRAY services at KSEPL took place in October.

Of the laboratory staff of 700, more than 100 are already registered users of the CRAY service. Most of these individuals are geophysicists or reservoir engineers. The laboratory provides service to other Shell operating companies around the world, who are either linked directly to the computer network or send in exploration and production data to be processed on the CRAY. Access to the CRAY allows the users to process exploration and production data with mathematical methods requiring fewer simplifying assumptions than with previous systems. The fewer the assumptions that must be made, the more accurately data can be handled.

In a report issued to the press about the CRAY opening, KSEPL explained that astronomical amounts of computation are involved in processing seismic data. Typically a seismic recording is made by firing a shot, the sound waves of which are recorded by 96 or more recording stations. Data may be recorded for about five seconds, and may be digitized at intervals of two milliseconds.

"Our single shot record therefore consists of 96 x 2500 numbers. Shots are typically fired every 25 meters. In 1982 the Shell Group shot some 100,000 kilometers of seismic line. To compound the computing problem, it is necessary to sample seismic data over a whole area, rather than along separate lines, increasing the amount of data to be handled enormously. About 30% of our seismic work is now concerned with areal coverage of this kind."

KSEPL officials pictured from left to right, Mr. van Engelshoven, Group Managing Director for Exploration and Production; Mr. Helfrich, President Director; and Mr. van Dam, Head of Production; celebrate the CRAY opening festivities.

The report went on to say, "It is appropriate to point out that the methods we apply are always compromises. Clearly we have to cut corners, make simplifying assumptions and try various mathematical tricks to get acceptable solutions within a reasonable timeframe. Each time a bigger and faster machine is made available, we are able to cut fewer corners and hence get better results."

In a speech at the opening ceremony, Mr. J. M. H. van Engelshoven, Group Managing Director for Exploration and Production at Shell stated, "Reducing the risk from the unexpected in oil exploration and production depends on the race between technology and the growing problems posed by the nature of finding smaller oil reservoirs. New technology, provided with appropriate software, will help us win this race. To increase the percentage of recoverable oil requires better quantitative analysis. To this, the CRAY can contribute."

Research scientists can now focus the seismic data better in the subsurface, which leads later to more accurate guidance of the drill bit. In reservoir modeling, the new generation of computers means that the complexity of the simulated reservoir can be increased. It is also possible to predict more accurately how a field will produce under different technical scenarios.

Mr. J. C. van Wijnen, Director of KSEPL, noted at the opening ceremony that, "Technical developments are often made in leaps. You are witness of such a moment in our continuous striving to sustain the provision of energy at the current level."

The Los Alamos data security system works

"You can knock at the door, but you'll have a heck of a time stepping over the threshold. Los Alamos' Integrated Computer Network (ICN) is as impenetrable as a system could be and still do the job for which it was designed." Those reassuring words came from Jimmy McClary, Division Leader and Dorothy Camillo, Group Leader for Computer and Telecommunications Security at Los Alamos National Laboratory.

The comment was made in mid-August after the international media event that centered on the unauthorized access of a Lab computer by a group of young people in Milwaukee. Most know the story by now: Using home computers and telephones hooked to a commercial data communications network the kids penetrated, briefly, the Lab's small VAX computers in the unclassified partition of the three-partition network. The invasion was quickly spotted and reported to the Department of Energy. That was in June. The incident has been investigated by the FBI since then, and the Lab has obtained a new telephone connect code number.

Despite allusions to the "War Games" movie, penetration of the ICN is no laughing matter, McClary said. What's really important is the way the network is protected from harmful invasion, the way it is partitioned into open, sensitive and classified segments so that the more sensitive the data, the more carefully it can be secured.

The raiders entered the system through the open partition that can be accessed through several computers with dial-up connections. These machines function in the perimeter of the open partition to serve researchers around the country who need to do business with Los Alamos.

Despite this protection, information in the open partition is limited to unclassified scientific calculations, general correspondence, reports, and other non-sensitive material. Beyond that are two more partitions: the sensitive, or administrative, partition contains protected information such as payroll, personnel, medical and financial records. The classified partition is just that: classified. Most of Los Alamos' CRAY systems are in the classified section.

The partitioned system was implemented in 1979 to deliver as much of the Lab's computing resources to the users as possible, while at the same time, providing adequate data security — a pair of tough and conflicting goals.

Until the partitioned system was developed, the Lab could offer no open computing service. But the nature of the Laboratory makes it beneficial for cleared and uncleared people to work together on some projects. Because working together usually means computing together, the partitioned system solved a major problem by making it possible for uncleared people to use the Lab's computing facilities.

This article is based on text appearing in the Los Alamos Newsbulletin, *Vol. 3 No. 33, August 19, 1983.*

NCAR's CRAY takes on acid rain modeling

There will come soft rains and the smell
of the ground,
And swallows circling with their
shimmering sound;
Not one would mind, neither bird nor
tree,
If mankind perished utterly...
—Teasdale, 1920

Today, Teasdale's words carry a sad irony. Because of acidic precipitation, rain is no longer soft; it can often be like vinegar. Swallows can no longer be blithely indifferent; they, and not humans, may perish of acid rain. But thankfully, many people are very concerned, and significant resources are now committed to understanding this vexing condition.

This past summer, the National Center for Atmospheric Research (NCAR) began a three-year, $3.5 million acid rain research program principally funded by the Environmental Protection Agency with support from the National Science Foundation. Virtually all computational models used in studying acid rain will be executed on the center's CRAY computers.

The phenomenon commonly known as acid rain was not detected until the late 1950's and early 1960's by Scandinavian and English researchers, who associated it with emissions of sulfur and nitrogen oxides from hydrocarbon combustion. When emitted into the atmosphere and carried up to hundreds of miles, the chemicals often are converted to sulfuric and nitric acids before they are finally washed out in rain or other precipitation.

Today, North America is also plagued with acid rain. The ecological woes associated with it are staggering. When acid rain runs through soils, it may flush away nutrients and leave toxic metals in its wake.

A number of extensive studies have been undertaken since the mid-1970's to study this problem. Many studies have provided inconclusive results because of acid rain's global scope in addition to the fact that it affects every part of an ecosystem. Acid rain is a silent and invisible marathoner — its effects are cumulative and may not surface for decades. Researchers need to know exactly how resilient the environment is to acid rain. One of the biggest gaps in acid rain knowledge is atmospheric chemistry. Most researchers are convinced of the link between air pollution and acid rain, but the actual connection eludes them.

To address the questions that continue to plague scientists, the Acid Deposition Modeling Project will develop a computer model for studying the acid rain problem in the United States. Dr. Julius Chang, director of the project at NCAR, said that the work is practically locked on the CRAY systems located at the center's Boulder, Colorado facility. The goal of the project is to design an analytical tool based on mathematical modeling to portray the interactions among the various physical and chemical processes in the atmosphere which lead to or contribute to acid deposition. The analyses are based on three-dimensional Eulerian methods.

The researchers will use an established meteorological model, which will significantly improve the understanding of the role of transport in acid deposition. Fundamental chemical process equations are also being applied to test scientists' understanding of the transformation processes by comparison with observations. Particular effort will be devoted to the interpretation of modeling results and the relationship between pollutant source and the affected area. "Basic research going into the development of the models is considerably more detailed than previous efforts," said Dr. Chang.

Dr. Chang who is a leader in the field of large computer models, commented, "Even with the CRAY, we will not be able to run all the analyses we would like. Both the volume of data and the size of the models are overwhelming."

"The first region to which we will apply the model will be the northeastern United States and southeastern Canada. At the end of three years, we expect to have a model available for the Environmental Protection Agency to use as an assessment tool," Chang explained.

He went on to say, "The philosophy of the project is to make every effort to examine all relevant physical and chemical processes and include them in a balanced manner in the computer model. Considerable emphasis is placed on illuminating the uncertainties in the model with regard to winds, chemical reaction rates and other atmospheric components."

The NCAR group will collaborate with researchers around the world, and participation by scientists in universities and federal laboratories is welcome. Because of the history of acid rain in Scandinavia, scientists from those countries will be invited to participate.

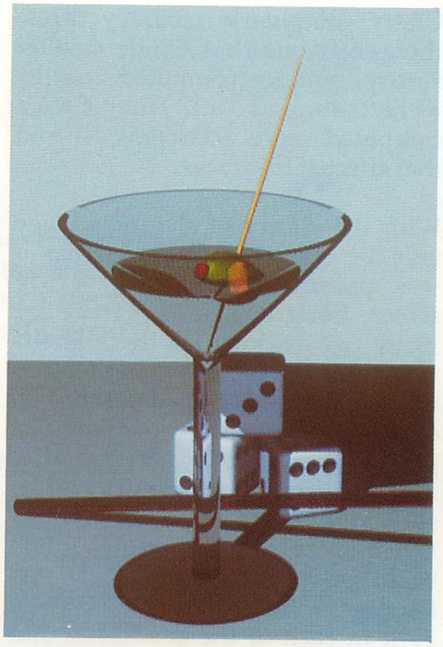

The CRAY acquires new talents

Mixing a martini normally calls for a little ice, plenty of gin and vermouth, and your favorite garnish. But the martini in the inviting image above was made of an entirely different concoction. "Martini", as the image is called, was generated by a CRAY executing a FORTRAN code based on raytracing methods.

Gray Lorig and Al Barr of Rensselaer Polytechnic Institute (RPI) in New York originally produced "Martini." Lorig recently received his master's degree in Computer and Systems Engineering, and Barr his Ph.D. in Applied Mathematics both from RPI. The image was first generated using the school's Prime 500 system.

On the Prime, the image took ten hours to compute. When the computation was done on the CRAY, the image was generated in three minutes. Lorig commented, "It was really impressive. The program was running on the CRAY within half a day, although it took some time to work out the bugs. Part of the problem was that the

CRAY computer's accuracy threw things off somewhat. On the smaller system where computed results were rounded off, the CRAY computed more accurately — and that affected the image."

With the clarity of the image, one may be surprised to learn that it was produced on a 512 x 512 screen. Lorig explained, "There are really two reasons for that. One is that the raytracing technique enables very effective shadowing. The other reason is that we subdivided pixels where needed to minimize aliasing effects. Each pixel could be subdivided into four quadrants. Whether or not it would be was based on two criteria: 1) the difference in intensities between the light in each of the quadrants, and 2) whether of not any of the four quadrants within the pixel was traced to a different object than the other corners. If it was, we subdivided; each quadrant would then be subdivided if necessary also."

Lorig went on to say, "By subjectively increasing the overall resolution, we could produce a pretty sharp image on a 512 x 512 screen. By contrast, most people increase the overall resolution of the screen and then pixel average down."

Lorig is quick to acknowledge that there is still a long way to go in computer graphics. The tremendous amounts of time needed to create high quality images is staggering, although that is changing as CRAY-level power becomes more prevalent. In animation, orchestrating and choreographing movement still challenges computer graphics experts. As an example, Lorig explained that the blurring of the back portion of objects as they move forward is very difficult. One typically wouldn't notice this effect, but in its absence, movement just doesn't look right. There is a need to have a language that is able to easily describe the relationships between objects and the motions they go through such as the rate of acceleration, speed, and the path of an object's movement.

Her Royal Highness Princess Anne and Mr. Ken McKenzie at ULCC.

FOUND: New Mersenne Prime

In September, a CRAY X-MP in Chippewa Falls, Wisconsin discovered what should prove to be the 29th Mersenne Prime. It certainly seems that CRAY systems have found the key in finding these esoteric numbers. Actually, Dave Slowinski, Cray researcher who has been involved in searching for the primes, was as surprised as anyone that the number was found so quickly. The 28th Mersenne Prime was discovered earlier this year by another CRAY. The new prime number, which has yet to be verified, is $2^{132,049} - 1$. The number has 39,751 digits. If printed, the number would fill more than one entire newspaper page.

ULCC's regal opening

In June of this year, the new extension to the University of London Computer Centre (ULCC) was opened by the Chancellor of the University, Her Royal Highness Princess Anne. The new extension houses the CRAY-1 S/1000 system recently moved from the Science and Engineering Research Council, Daresbury Laboratory.

The centre provides supercomputer service both to London and some 50 other universities in the United Kingdom. It houses several other computer systems including an Amdahl 470/V8 which front-ends the CRAY system.

Her Royal Highness Princess Anne inspects the CRAY computer with Mr. Ken McKenzie in the photograph above. Mr. McKenzie is the Deputy Director of the University of London Computer Centre.

Corporate Addresses

Worldwide Business Centers

Domestic

Central Region Office
5330 Manhattan Circle, Suite F
Boulder, CO 80303
Tel: (303) 499-3055

Albuquerque, NM
Chicago, IL
Minneapolis, MN

Eastern Region Office
11120 New Hampshire Avenue
Silver Spring, MD 20904
Tel: (301) 681-9514

Atlanta, GA
Boston, MA
Laurel, MD
Pittsburgh, PA

Government Relations Office
1828 L Street, Suite 201
Washington, DC 20036
Tel: (202) 775-0155

Petroleum Region Office
5858 Westheimer, Suite 500
Houston, TX 77057
Tel: (713) 975-8998

Dallas, TX

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

Los Angeles, CA
Seattle, WA

International

Cray Research France, S.A.
21 Avenue Victor Hugo
75016 Paris, France
Tel: (331)-501-5270

Cray Research GmbH
Perhamerstrasse 31
8000 Munich 21, West Germany
Tel: (49)-89-56014-0

Cray Research Japan, Limited
W 1154, Shin Aoyama Building
1-1 Minami-Aoyama 1-chome
Minato-Ku, Tokyo 107, Japan
Tel: (81)-3-403-3471

Cray Research (UK) Ltd.
Seymour House, The Courtyard
Denmark Street, Wokingham
Berkshire, England
Tel: (44)-734-7911-80

CRAY
RESEARCH, INC.

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