A CRAY RESEARCH, INC. PUBLICATION

CRAY CHANNELS

Volume 5, Number 2

FEATURE ARTICLES:

REGULAR COLUMNS:

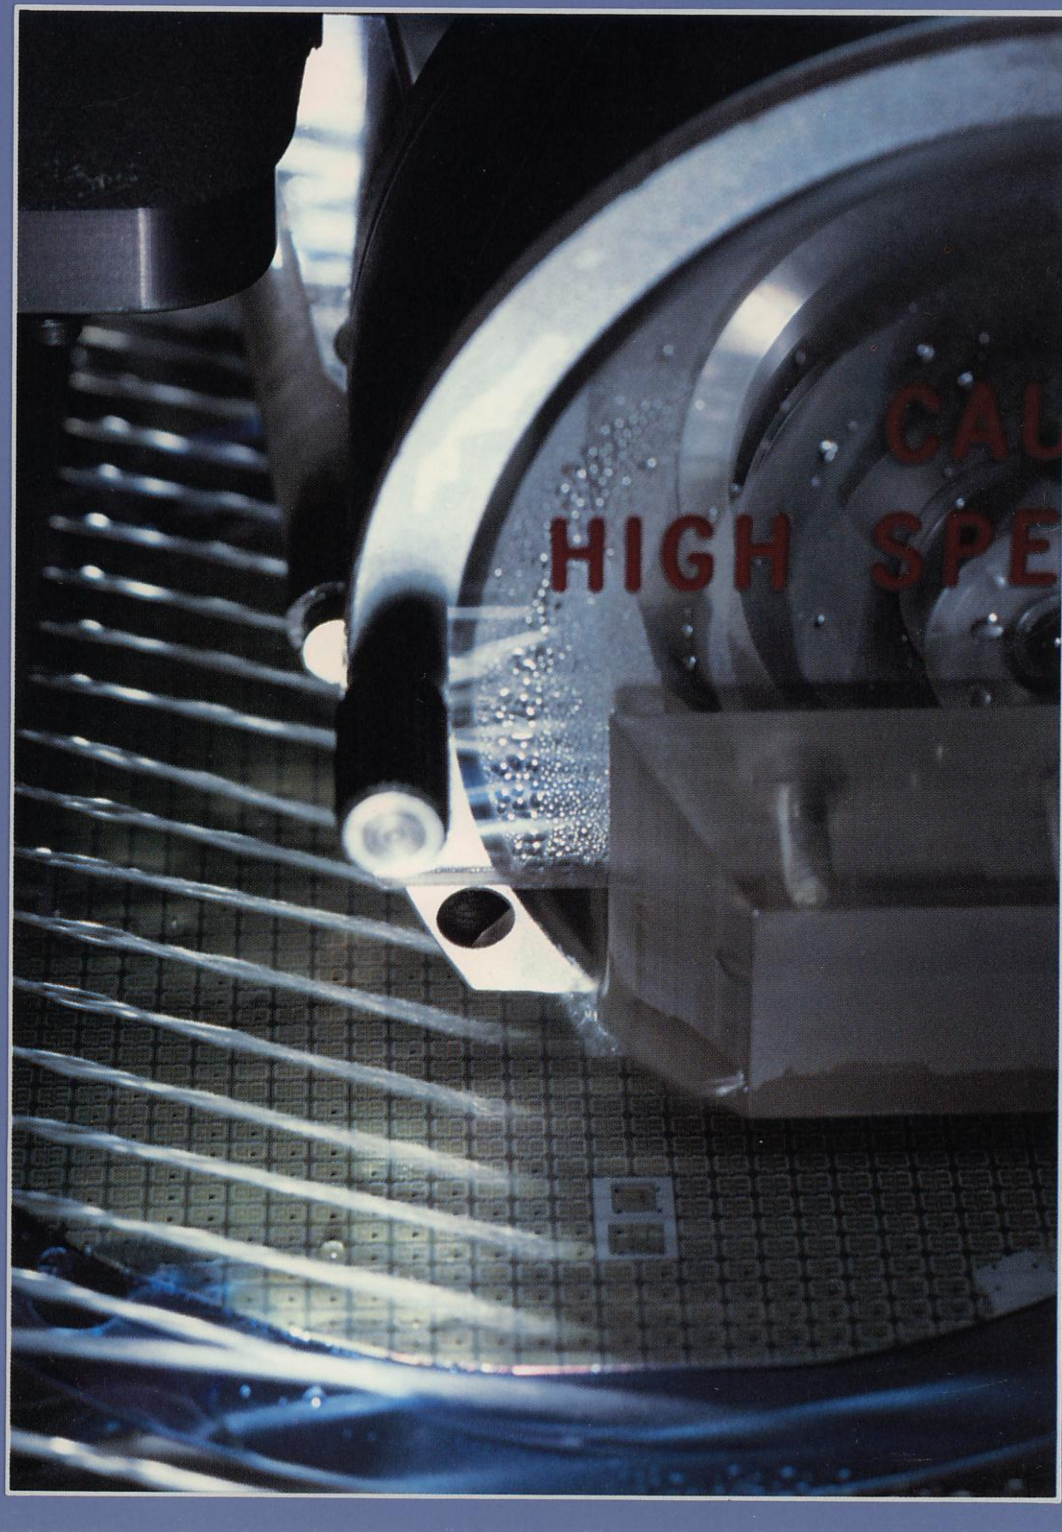

In this issue

Recently, Cray Research announced its ninth order for a CRAY system from a major company involved in petroleum research. Some of Cray Research's petroleum industry customers are ARCO, Chevron, Shell and Phillips. One of the most powerful CRAY computers in the world is a four-million word system installed at EXXON Corporation. Why do so many petroleum companies need supercomputer power? Simply, the increasing challenges of discovering new reservoirs and recovering oil from existing ones require very complex analyses.

This issue of **CRAY CHANNELS** discusses seismic processing and how CRAY power is being applied in the petroleum industry. Our first feature article, beginning on page 2, provides a basic explanation of what seismic analysis is and why it is so computationally demanding. The second article focuses on the computational concerns of the geophysicist. Those features of the CRAY that address the most cumbersome aspects of seismic processing — input/output and computation — are discussed. In addition, universal seismic benchmarks run on the CRAY are described. This article begins on page 6.

Also in this issue, a basic explanation of integrated circuit fabrication is provided in a third feature article. Our Applications Highlight column focuses on a powerful graphics applications package. And as always, our regular columns keep you abreast of product, company and user news.

About the cover

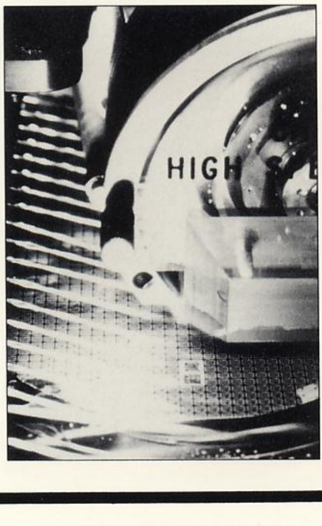

Wafers cleansed during IC fabrication

Cray Research is in the process of setting up an integrated circuit (IC) fabrication facility in Chippewa Falls, Wisconsin. Shown on our cover is one step of the production process. During many stages of IC fabrication silicon wafers are cleansed in a bath of deionized water. Even while being sawed into hundreds of ICs, the wafer and diamond-bladed wafer saw are flooded with the ultra-pure water. In addition to removing contaminants at that stage, the water also helps keep the saw blade cool. Completed wafers are placed in a special water spinner that blows warm nitrogen over them to ensure that they dry without water spots. The contaminated deionized water is recycled for future use.
Photo credit: Joe Gianetti

CRAY CHANNELS

A Cray Research, Inc. publication Volume 5, Number 2

CONTENTS

FEATURE ARTICLES

2

REGULAR COLUMNS

Staff

Editor Tina M. Bonetti
Assistant editor Carol A. Smith
Art/Design Chad Jewett
Kathleen Van Kirk
Jeanne M. Jacobowski

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.

Petroleum exploration takes seismic crews to distant corners of the earth. Inset: Marine acquisition operations.

Seismic Processing Overview

Petroleum exploration involves a lot of money, time and luck. In 1981 alone, 75,000 holes were drilled in the hope of finding crude at a cost of about $3 million each. Of those, only 30% were considered large enough to produce commercial amounts of oil and gas. To alleviate wasted drilling expenditures, explorationists need a better understanding of an area's geology prior to drilling. Seismic prospecting for oil and gas provides that insight. First used for petroleum exploration in the late 1920's, it is the most successful method of petroleum exploration today, accounting for over 90% of all such efforts.

Seismic principles

Reflection seismology is based on analysis of sound vibrations reflected from different rock layers (interfaces). Energy artificially induced into the ground encounters discontinuities between the layers and is partially reflected back to the surface. The returning energy (reflections) is detected and their strengths and reflection times are recorded. By processing the amplitudes, frequency and elapsed time factors of the reflections, geophysicists derive seismic sections of the substrata. Figures 1 and 2 illustrate basic seismic principles.

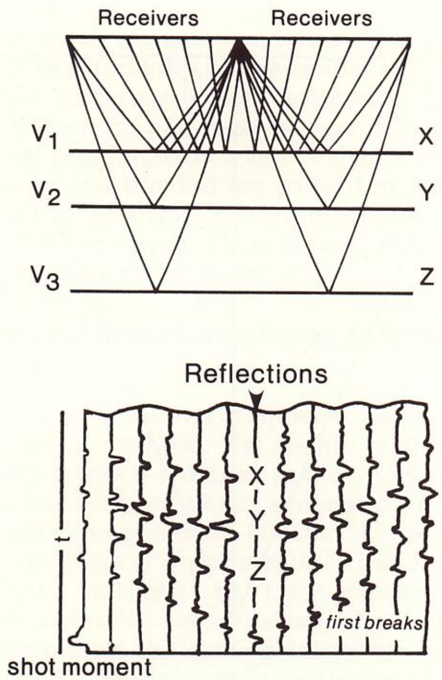

Figures 1 and 2 Basic seismic principles. When there is a seismic impulse at the earth's surface, some of the seismic energy radiates from the shot point. The waves reflect and refract at each boundary layer, with part of the wave penetrating lower and part reflecting at the interface.

The seismic section is not an actual picture of the interfaces, but the reflections. The curvature of the subsurface interfaces, variations in the velocity of propagation through the layers and variations in the thickness of layers distort the actual signals much as one's reflection is distorted by a "fun house" mirror. The degree to which these signals represent the subsurface depends on the quality of the seismic data and processing undertaken. Analysis and interpretation of the data along with other geological information provide clues about an area's petroleum yielding potential.

Data acquisition

To aquire seismic data, controlled artificial energy called shots are induced into the subsurface. Dynamite explosions are very popular. However, more controllable energy devices (sources) are often used.

One type of onshore device is a truck-like vibrator vehicle. These trucks have a heavy metal "base plate" connected to their undersides by a hydraulic cylinder. After lowering the "base plate" to the ground and raising itself onto it to add weight, the truck violently vibrates the earth using controlled hydraulics. Often, these trucks are ganged together in groups of four or even eight to increase the shot size. The most popular device used in offshore exploration is the air gun, which is pulled through the ocean by a specially outfitted ship. The air gun systematically releases shots of highly compressed air, the impulses of which penetrate the ocean floor and create reflections.

Acoustical receivers detect and send the reflections to a recorder where they are amplified, digitized and stored on tape. Normally, 48 to 96 receivers are active for each source input, although new technology is making the use of up to 1024 receivers practical.

Conventional two-dimensional mapping requires that this operation be systematically repeated in a straight line. Typically, 24 shots are fired per mile. A seismic "line", which varies in length from a few hundred yards to tens of miles, is acquired by moving the receivers and source along a straight line over the subsurface. This repetitive operation ensures that reflections from a single reflecting point will be recorded at many receiver locations. The redundant information helps geophysicists sort out primary reflections (those that have reflected directly and not ricocheted onto other layers) from noise and other distortions. This multiple coverage process is illustrated in Figures 3 and 4.

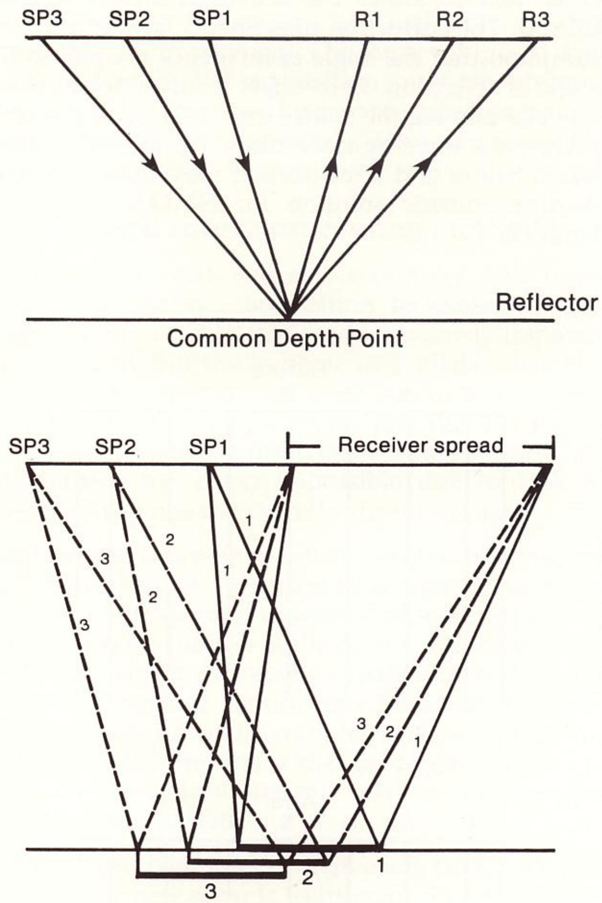

Figures 3 and 4 Multiple coverage principles.

The receivers typically detect reflections from each shot at sample rates of four, two or one milliseconds (ms) over a six-second detection period. The latest technologies in geophysical equipment make even greater resolution possible. Sampling intervals as small as 0.5 and 0.25 ms are now practical. Tremendous quantities of data are acquired for processing. It is not unusual for one billion values to require processing for a single seismic line.

Data processing

In its rawest form, the volume of seismic data is overwhelming. Analysis is also difficult because of the distortions, reverberations and noise in the collected reflections. Data processing attenuates noise and enhances the accuracy of the primary reflections. Average signal velocity is computed for use in other computations. These early operations ready the seismic data for analyses that will help geophysicists glean information about the geology below.

Normal moveout

A major first step is to apply a normal moveout correction (NMO) to the data. Ideally, geophysicists want to analyze traces that would be recorded if the source and receiver occupied the same surface point (zero-offset). But because most traces are not collected that way, the data must be corrected for the distance that separates the source and each receiver (offset). The correction uses Snell's Law and the assumption that the angle of incidence is equal to the angle of reflection. Zero-offset is solved for by theoretically moving the source and receiver to the same point on a trace-by-trace basis. Figure 5 illustrates traces before and after normal moveout correction. A diagrammatic solution for NMO is shown in Figure 6.

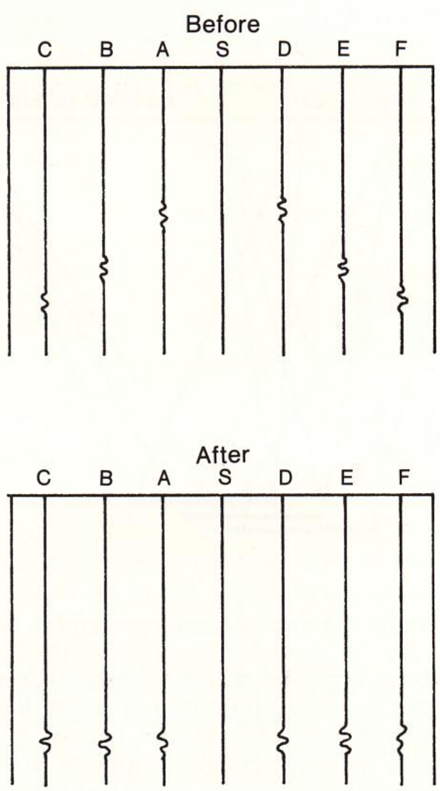

Figure 5 Traces before and after NMO.

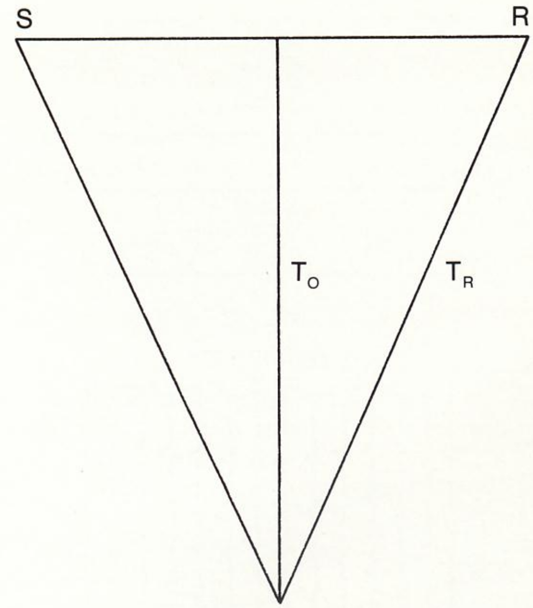

$$T_R = \sqrt{T_0^2 + \left(\frac{X}{V}\right)^2}$$

T_O = Two-way zero offset travel time

T_R = Two-way travel time with source/receiver offset

X = Source/receiver offset

V = Average propagation of velocity

S = Source

R = Receiver

Figure 6 Solution for NMO.

CDP stack

After NMO correction, traces that correspond to the same subsurface reflection point but were recorded from different receivers and shots are summed into a single trace. The summing process is known as common-depth-point (CDP) or common-midpoint (CMP) stacking. The common-midpoint is that point on the surface of the earth halfway between shots and receivers. Stacking enhances the signal-to-noise ratio by suppressing multiples, reinforcing the signal content. Reflections that show a different normal moveout due to refraction, diffraction, etc., relative to primary reflections, are attenuated. The volume of data is significantly reduced (i.e. by factors of 24 to 48) as traces are summed into the stack. This means that more complex algorithms can be used in later processing.

While the stack provides a meaningful display of the subsurface, it is still an artificial representation with limited use. An important shortcoming of stacking is that the results assume the reflecting layers are horizontal, when in reality they dip. In the presence of structural dip, the reflectors will be plotted vertically underneath the common-midpoint between the source and receiver arrays instead of in their true position, which can be determined by the downward and upward ray path of the signal. The shortcoming of stacking is overcome by "moving" the stacked traces to their true spatial positions through the migration process.

Additional processing using migration methods

By solving the scalar-wave equation via one of several migration methods, the depth of the dipping boundary layers is obtained as the traces are shifted to their correct positions in space and time. Migration algorithms generally fall into the categories of Kirchhoff integral or summation, finite difference and finite element or frequency-wavenumber domain methods. These velocity-dependent algorithms can yield the depth solution for the seismic section by converting reflection time to depth. This profile of the depth and inclination of the reflecting layers provides very meaningful information for geophysicists.

Ideally, migration would be applied to all of the data prior to stacking, but because of the complexity of the migration computation and the amount of pre-stacked data, migration is usually relegated to post-stack processing. The availability of computer power offered by the CRAY makes pre-stack migration more practical.

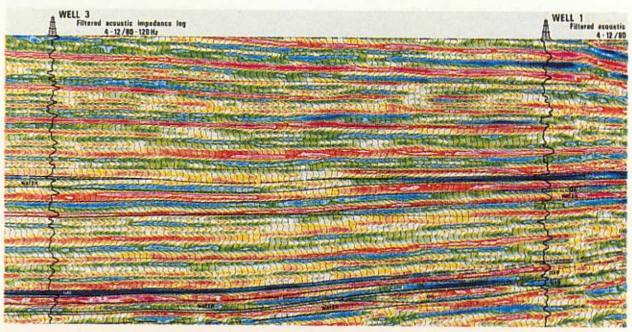

A seismic section displayed.

Interpreting the results

When processing is completed, the seismic section is printed on paper or displayed on a graphics terminal. Geophysicists analyze the section in hopes of spotting reservoirs. If a production and/or geologic history is available, that information may be used to calibrate the seismic section.

Seismic analysis attempts to identify the subsurface geometry where traps may be located. However, it cannot indicate whether and where oil and gas are present. The analyst tries to identify traps, which would appear as large porous formations relative to the surrounding structures on the sections. Anticlines and faults where the overlying formation is impermeable are likely areas for reservoirs. Explorationists rely upon a combination of the geophysicists' analyses, intuition and experience for accurate assessments.

Three-dimensional reflection seismology

In contrast to conventional seismic prospecting using two-dimensional profiles, three-dimensional methods generate seismic data for a volume of the earth. During acquisition, receivers are placed in a large rectangular area, rather than just on a line. Three-dimensional data has distinct advantages over two-dimensional. An important difference between them is that 2-D data is distorted by energy returned from layers located outside the plane of the section. While 2-D sections are restricted to the direction of the seismic line, 3-D sections contain areal (for an area) information.

The discussion of 3-D seismic processing has been academic until very recently. Recent advances in field recording technology are just now making areal acquisition practical. The large volume of data generated with 3-D acquisition only underscores the need for high-performance supercomputers. Traditional scalar computer systems are incapable of addressing true 3-D processing.

Current and future processing methods

As you have read, seismic is complex and time-consuming. Today's acquisition technologies provide volumes of data that were previously impossible to obtain. But the increase in available data is a double-edged sword. At some seismic processing centers, the large computers are just able to keep up with basic processing of this data. More sophisticated analyses are often impossible due to lack of power. This limits the overall value of the data.

But now the business of seismic exploration is at an exciting crossroad. Advances in computing technology and its application to seismic processing provide power to make full use of the data. CRAY computers with their vector processing capabilities and strong scalar performance significantly reduce seismic processing time and make complex analyses practical. What this really means is that geophysicists can extract the most accurate and relevant information possible from the data.

In the future, computer centers using the CRAY will be applying new seismic techniques. Processes such as 3-D analysis and post-stack migration will provide the accuracy needed to identify potential reservoirs. Better selection of reservoirs with the highest oil-yielding potential can save the petroleum companies literally millions of dollars. These are the reasons that CRAY computers are becoming part of this exciting exploration process. □

Acknowledgement

Photographs provided courtesy of Compagnie Generale de Geophysique of France.

The seismic processing problem:

Cray Research's response

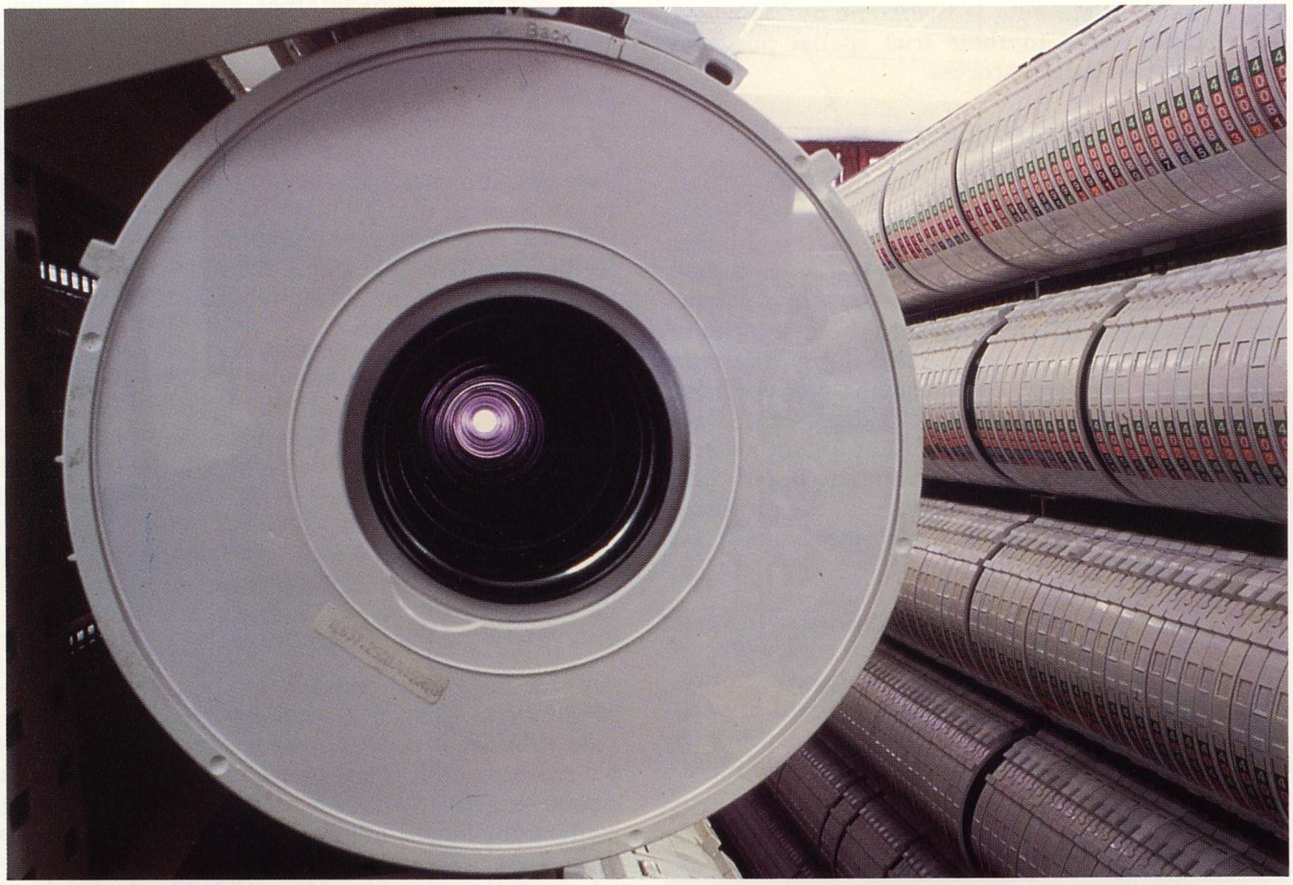

This may be a familiar sight to those who deal with the large libraries of seismic data.

Consider a typical seismic line, 20 miles long with recording stations every 200 feet. Suppose a dynamite shot is fired every 400 feet, for a total of 264 shots. For each shot fired, the reflected sound is recorded at 96 recording stations. From each of the 264 shots, 25,344 traces are produced, each of which contains about 3,000 32-bit words.

Obviously, processing this data will involve considerable computation requiring strong I/O capabilities. What are the reasons for using a CRAY in seismic processing? In order to answer this question, let us look at the requirements of a good seismic data processing (SDP) system.

Requirements of a good SDP system

First, an SDP system must be designed to collect, store and retrieve vast amounts of data for processing. Mathematical processes such as Fast Fourier Transforms (FFTs) and deconvolution processing are performed on very large data streams and the results must be saved.

Second, I/O devices must have sufficient bandwidth to pass data to and from the central processor effectively. Adequate storage devices are required to save permanent data files, and intermediate or temporary storage is necessary for processing.

Third, a computer or processor that performs both scalar and vector operations at very high speeds (millions of floating point operations per second) is required. Since not all code can be vectorized, good scalar processing capability is important.

Fourth, an efficient operating system is needed to schedule jobs and balance I/O with CPU processing.

Finally, and importantly, the user must have quick and easy access to the system. He must be able to interface easily with the system in a familiar language. Thus, standard FORTRAN is desirable, if not a must, for the user. Users on a network must be able to communicate with a large computer center in order to get their jobs done and receive results to enable them to do their analyses.

Here, we will describe characteristics that make CRAY systems ideally suited for seismic processing and will relate some benchmark results demonstrating system capabilities.

The CRAY configuration

Three Cray products, the CRAY-1/S, CRAY-1/M and CRAY X-MP computers, address the major requirements of large-scale processing. The CRAY-1/S, introduced in 1979, today is installed at numerous petroleum company computer centers to address seismic data processing. Its successor, the CRAY-1/M, established a new supercomputer price/performance curve upon introduction last fall. The CRAY X-MP is a multiprocessor system and Cray Research's most powerful computer. Characteristics of these systems that are important in seismic processing are

- Very efficient I/O
- Large disk storage devices
- Fast scalar and vector processor
- An efficient operating system
- An ANSI 78 FORTRAN compiler
- Networking capabilities

In the case of the CRAY X-MP, two processors share 16 or 32 Mbytes of Central Memory and a CRAY-1/M System is expandable from eight to 32 Mbytes of memory. A CRAY system (the X-MP in this case) is shown in Figure 1. It consists of the central processor, the I/O Subsystem (IOS) and the Solid-state Storage Device (SSD). Each CPU has four six-Mbyte channels that operate in parallel for I/O control, giving a bandwidth that allows for simultaneously executing two reads, a write and an I/O in each channel.

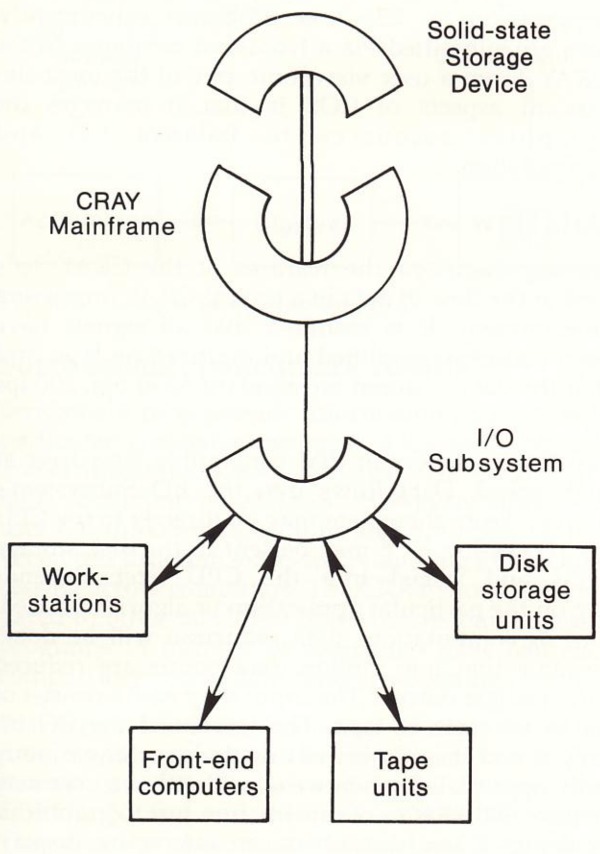

Figure 1 CRAY computer system.

I/O handled by the I/O Subsystem

Handling I/O is just as important as computational speed in SDP. CRAY computer systems can handle large I/O requirements without slowing down the CPU significantly. The I/O Subsystem was designed as an integral part of CRAY mainframes to evenly balance I/O and computation. I/O work is done by the I/O Subsystem rather than the CPU. An important feature of the subsystem is that it directly accesses Central Memory through high-speed channels. Each channel has a maximum transfer rate of 850 Mbits/sec. A large capacity Buffer Memory enables the I/O Subsystem to move data at very high rates without interrupting the CPU.

Communication from front-end computers is also handled by the I/O Subsystem, leaving the CPU free to do the work for which it is intended. A CRAY-1/M, with its I/O Subsystem, can support 16 streams of disk and eight streams of tape simultaneously transferring data at its maximum rate. (Actually, a total of 32 disks and 64 tapes can be connected.)

The Solid-state Storage Device

Complementing CRAY systems by providing additional storage capacity is the Solid-state Storage Device. The SSD provides extremely high-speed mass storage. It is used in the same way as a large disk file, but offers significantly improved performance. The SSD is available in 64, 128 or 256 Mbytes. On the CRAY X-MP, the SSD connects to Central Memory through a high-speed channel, increasing performance by a factor of 50 and 100 for short random and long sequential transfers, respectively. On the CRAY-1/M, the SSD connects to Central Memory via one 100 Mbyte/sec channel, with corresponding improvement factors of 25 and 30.

Disk and tape streaming benchmarks

In a benchmark demonstrating CRAY systems' disk reading capabilities run on a CRAY-1/S, one cylinder was read from each of 16 disks simultaneously. Each cylinder was read at the maximum rate. Twelve Mbytes were read in 0.2 second, demonstrating a transfer rate of 60 Mbytes/sec. During the transfer, the CPU was kept occupied for 80% of the time with a FORTRAN DO-loop. The CPU was able to continue its work without having to wait for data. In the benchmark, data transfer occurred five times faster than the CPU could have performed the delta-T correction algorithm. This illustrates that the transfer rate from disk can keep pace with the CPU. In another benchmark, five seismic tapes were read simultaneously at full tape speed, yielding an observed data transfer rate of 1.1 Mbytes/sec per tape.

CRAY hardware applied to SDP

Let us turn our attention to Cray design philosophy. CRAY architecture tends toward increased parallelism by featuring pipelining, multiple independent functional units, vector processing and multiple processor systems. In addition, Cray computers can be and are configured in networks of other computer systems, terminals, printers, etc. These features increase performance and the overall benefit derived from the use of CRAY computers.

The Central Processing Unit

The CRAY CPU contains functional units and operating registers that are associated with vector, scalar and address processing. Data is transferred from Central Memory via registers to functional units that do the actual arithmetic operations. Several functional units can operate simultaneously so that several operations can be done in parallel. This is an important feature of the CRAY. Code is generated by the compiler so that scalar and vector operations can be done simultaneously. Except for a start-up time, the vector operations are capable of one output result per clock period, (9.5 nsecs on each processor of the X-MP or 12.5 nsecs on the CRAY-1/S and CRAY-1/M).

Parallel vector operations can be obtained by using the resultant output stream from one vector simultaneously as the input to a different vector operation. This process is called "chaining." The chaining of two vector operations together effectively extends the power of the CRAY so that full advantage of the vector operations is obtained.

Using FORTRAN

Seismic exploration is a rapidly evolving science, so many new programs are being developed and old ones modified. To accommodate these changes, a computer used in seismic data processing must have a FORTRAN compiler capable of generating efficient code. Standard FORTRAN DO-loops can thus be translated into machine instructions that use the vector registers efficiently. For example, a delta-T correction routine that interpolates between samples has been written in FORTRAN. It can delta-T correct a 3000 sample trace in 2 ms.

CRAY systems can process short vectors efficiently. Consider the following FORTRAN DO-loop.

```
          DO 100 I = 1, N
100        A(I) = B(I) * (C(I) + D(I))
```

Computation is extremely rapid and performance improves as N increases. Even when N is as small as 10, one-half of full performance is realized. This is especially important in SDP where short vector lengths are common. Central Memory is large enough to hold several hundred traces easily, so traces can be processed in groups.

Suppose the following scalar FORTRAN statements are executed for each trace:

```
      TIME = DIST/VEL
      IF (TIME.GT.TMAX) TIME = TMAX
      SAMPS = TIME/DIGI
```

The scalar instructions may be processed for all traces first, followed by subsequent processing for all traces. Thus, the above code would become:

```
      DO 100  I = I,N
      TIME(I) = DIST(I)/VEL
      IF(TIME(I).GT.TMAX) TIME(I) = TMAX
100   SAMPS(I) = TIME(I)/DIGI
```

The index, I, refers to the I'th trace in the group of N traces. The entire DO-loop, including the IF statement, can be vectorized.

Many computations performed on seismic traces are not amenable to vectorization if the traces are processed one at a time. But as long as the same computation is performed on each trace, vectorization is usually possible by using the above procedure. Statements such as GO TO's cannot be vectorized, of course.

A common operation in seismic processing is

A(I)=B(K(I))

A, B and K are vectors and the computation is carried out for a range of values of I. This operation can be done on a CRAY much faster than on other available computers.

The CRAY Operating System

The CRAY Operating System (COS) manages system resources. It was designed to be a low overhead multi-programming operating system that supports up to 256 user programs concurrently. Jobs are submitted via a front-end computer to the CRAY. From a user viewpoint, one of the most significant aspects of COS is that it manages the computer resources to balance I/O and computation.

Data flow

Having discussed the features of the CRAY, let's look at the flow of data in a typical CRAY processing environment. It is assumed that all signals have been collected, amplified and digitized on tape, and that the data is stored on standard 6250 bpi, 200 ips tape.

Tapes are read on an IBM compatible tape drive at high speed. Data flows into the I/O Subsystem's buffers. From there data may go directly to the CPU for processing or it may be sent to the disk storage units and staged into the CPU later, depending on the particular application or algorithms used. During computation, data reduction will be done. Assume that one million data points are reduced into a single output. The input may easily consist of 40 to 60 reels of tape. The impact of the I/O involved and the number of calculations done is more fully appreciated when we consider that a user may require 600 Mbytes of information just to graphically depict a single result on an interactive display. Data may be stored on a CRAY disk and then sent via high-speed channel to a front-end computer.

Numerous users may be connected to the CRAY via a network or a data net such as a Network Systems Corporation HYPERchannel™. Standard graphics software can be used to display various outputs, or results can be printed for analysis. Networks of computers, interactive terminals and other peripherals have easy access to the CRAY. The user is operating in a familiar environment, and the CRAY has handled millions of words of data efficiently. Figure 2 shows data flow into, through and out of the CRAY.

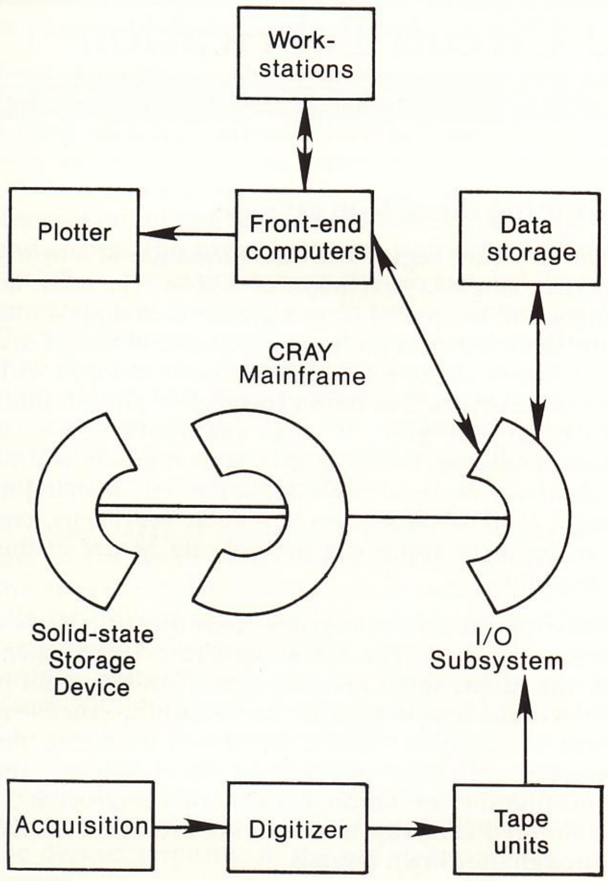

Figure 2 Seismic processing data flow.

Some seismic benchmark results

Benchmark runs provide information on how well a particular computer performs. Our earlier benchmarks illustrated CRAY systems' disk reading and tape streaming capabilities. However, since each user has different algorithms and priorities, no universal test program exists for evaluating seismic processing across computers. The CRAY should be compared to other machines based on customer-specific programs. A number of companies have had such comparisons run to date, with eye-opening results.

Bearing in mind the limitation of a universal test program, Cray has written two seismic test programs in CRAY Assembly Language and FORTRAN, and run them on a CRAY-1/S system. The results of these tests indicate the CRAY's computational speed and demonstrate CRAY performance in executing SDP.

The input for both jobs was a standard SEG Y tape with approximately 19,000 traces, each containing 1,500 samples. The first benchmark consisted of time ramp scaling, deconvolution, trace equalization, NMO, stacking, bandpass filtering and outputting to tape. Only 3 ms of CPU time per trace was required. With two identical copies of this job running simultaneously, 280 total traces/sec were processed. CPU utilization (CPU time/elapsed time) was 80%.

The second benchmark included all the processes of the first. In addition, it performed a disk sort of the input data into CDP order. The input tape was modified so that traces were in shot order. An F-K filter was applied to each subpoint. The addition of the sort and the F-K filter added 2 ms/trace to the CPU time, for a total of 5 ms/trace. The throughput with two identical jobs in the machine was 160 total traces/sec with a CPU utilization of 77%. The sort could have used an SSD or run a number of disk units in parallel which would have reduced the I/O time even further.

Conclusion

Many considerations must go into the design of any computing system, and a seismic processing system is no exception. However, criteria such as fast processing and efficient I/O are redefined with SDP. Meeting those criteria has not been easy. Conventional computing philosophy has not been completely successful. However, Cray computers provide a unique approach to large-scale processing needs that effectively addresses the SDP challenge. For this reason, a number of large petroleum companies have made room for CRAY systems in their computing environments.

A spokesman for SOHIO explained, "Because we are dealing with more complex geological structures and more data than ever before, I/O bandwidth and CPU power were overriding concerns in our identifying a good seismic system." The CRAY at SOHIO was accepted in the second quarter of this year. Already, good results are being realized in preliminary testing of advanced seismic algorithms written for operation on the system.

Early in 1984 a CRAY-1/M will be installed at Phillips Petroleum Company. William Shack, Exploration and Production Technical Services Manager at Phillips, explained, "We have every reason to expect that the CRAY will process our data many times faster than our current systems. This type of capability brings processes like pre-stack migration into the realm of practicality." □

Acknowledgements

This article was prepared with the assistance of several Cray personnel with special thanks extended to Robert T. Haase and Max K. Miller of Cray Research, Inc.

Photograph provided courtesy Western Geophysical.

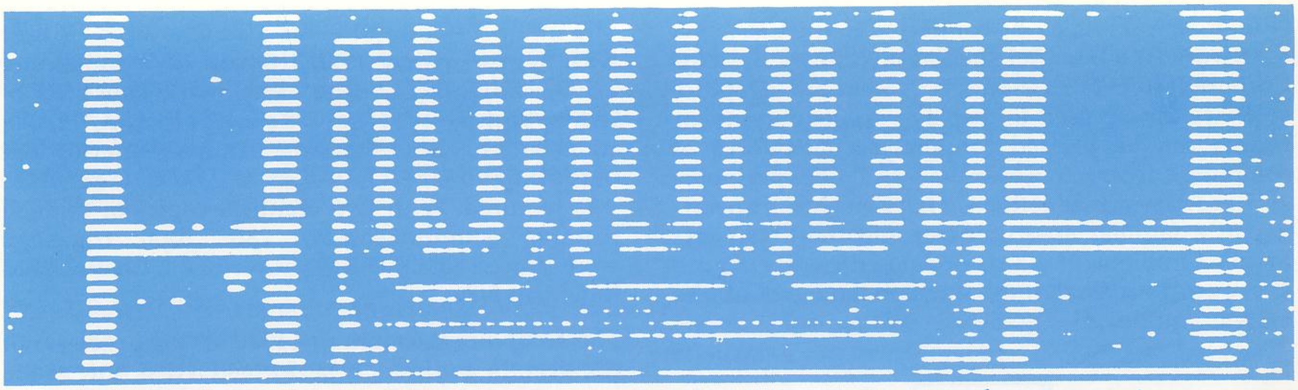

Fundamental Integrated Circuit Fabrication

Availability of integrated circuits on a mass scale has been a major thrust in the advancement of information technologies. Today, semiconductor technology is beginning to lead mainframe design. But the leadtime for integrating new circuit designs into hardware can easily be four to six years. In many cases, leadtimes like that are unacceptable — like in Cray Research's case.

Delbert Eberlein,Chief IC Process Engineer for Cray, explained, "In order to maintain our leading position in supercomputer design, we must shorten the circuit design to implementation cycle. Cray is now beginning to develop new circuits and mainframes in parallel. This will decrease leadtimes by about half, down to about two or three years. We feel very strongly that those who can develop and produce integrated circuits in-house will lead the industry in the introduction of more powerful computers. Cray expects to maintain its position in that arena."

With that, CRAY CHANNELS decided to take a look at the world of semiconductor technology. In this article, we look into integrated circuit fabrication. In our next issue we will focus on the design process and other aspects of the integrated circuit technologies.

The proliferation of integrated circuits may lead one to believe that the production process, if not the technology, is fairly straightforward. Actually, integrated circuit production relies on close interaction of technologies as diverse as solid-state physics, chemical engineering and photography. A silicon wafer (the basic material of a chip) can undergo 150 different processing steps in fabrication. Many of these are chemical cleaning procedures, rinses and making, but all are essential to achieve the final result — integrated circuits. Of the several thousand chips that will be produced from a single four-inch-diameter wafer, industry statistics show that roughly 30-60% of them will be rejected. High precision tools and stringent clean room conditions are a must in this production process where lengths are measured in angstroms and displacement of a few electrons makes a difference.

Forming the silicon crystals

The sequence begins with the formation of a nearly perfect single-crystal ingot of silicon. Typically, an ingot will be created from a perfect seed dipped into and drawn from a molten pool of silicon heated in a crucible to about 3000° C. The silicon is doped with an impurity such as boron (p-type) or phosphorous (n-type) to give it the characteristics of a semiconductor. In doing this, silicon atoms will be selectively replaced with dopant atoms. Keeping in mind that 10^{22} atoms are in a cubic centimeter, one can begin to appreciate the delicate nature of this operation.

Drawing the silicon crystals from the crucible is a precise process. The lattice-structure arrangement of the atoms must be maintained. As the ingot is drawn, the temperature of the remaining crucible is kept constant at 3000°C by slowly reducing the amount of heat applied to it. Likewise, the phosphorous or boron content of the crucible is monitored carefully to ensure its even distribution through the drawn crystals.

Crystal ingots may be drawn to a specified diameter from roughly three-fourths of an inch to six inches. Length may range up to 60 inches, depending on the pulling rate, melt temperature and other externally controlled factors. Then the ingot is sawed into thin wafers with a diamond saw. The wafers are carefully polished to a reflecting finish by mechanically lapping the wafer and then polishing with successively finer grit. A final chemical and mechanical polish leaves the surface virtually free of scratches and imperfections. Many companies that produce integrated circuits purchase the silicon wafers from manufacturers that specialize in this operation.

The fabrication sequence

The cornerstones of the integrated circuit fabrication process are oxidation of silicon, photolithography, impurity diffusion and epitaxial growth. Most of this work is automated and requires stringent clean room conditions.

Initially, a layer of silicon dioxide (SiO_2), called substrate during processing, is grown on the surface of the wafer by heating it to a temperature of about 1100° C and exposing it to oxygen or steam. This formation of oxide, or glass, protects the silicon surface from contamination and prevents diffusion of impurities into the silicon during subsequent processing.

Special photolithographic techniques are used to remove the oxide from regions where transistors are to be placed. The regions from which oxide has been removed now provide "windows" through which impurities can be deposited on the silicon surface.

Next the wafer is placed into a diffusion furnace heated to 1000° C. For a p-type substrate a gas containing an n-type impurity such as arsenic is passed over the surface of the wafer. Successively lower temperatures are introduced to the wafer to prevent destruction of the already completed depositions and etches. The gas decomposes and impure atoms are left on the entire surface of the wafer, but do not penetrate the SiO_2 surface. Where windows exist, the high temperature causes the impurities to diffuse into the silicon, forming the highly conductive n-type layer. This layer is called the buried layer and reduces the series collector resistance of the transistor. After buried-layer diffusion, all oxide is removed from the wafer surface.

A second layer of silicon is now grown on the surface of the wafer by means of epitaxial growth. In this step, the wafer is exposed to an atmosphere containing a silicon compound that decomposes at high temperatures, leaving silicon atoms on the wafer surface. Since these atoms are deposited in accordance with the crystal structure of the substrate, the epitaxial layer thus grown and the substrate are part of the same single-crystal structure. The epitaxial layer is doped, for example, by phosphorous for an n-type epitaxial layer by including small amounts of the desired impurity in the gas stream during the growth.

Successive oxide growths, photolithographic steps and diffusions are carried out to place isolation walls and base and emitter impurity distributions on the wafer. This process is shown in Figure 1. The successive diffusion of isolation walls, base region, and the emitter requires impurity compensation of the p-type base by the n-type emitter diffusion so that the region containing the emitter is converted to a net n-type region.

Finally, contact windows are opened in the oxide layer for connecting to the device. Typically aluminum is evaporated into the contact windows and selectively etched, thus making contact with the silicon itself and forming route lines to interconnect the circuit. One or two additional layers of metal may be deposited — at temperatures hovering around 560°C. A passivation layer of glass is applied over the metal to protect the thin metal lines from scratches and moisture. Selectively, small holes are cut into the passivation layer so that the chip can

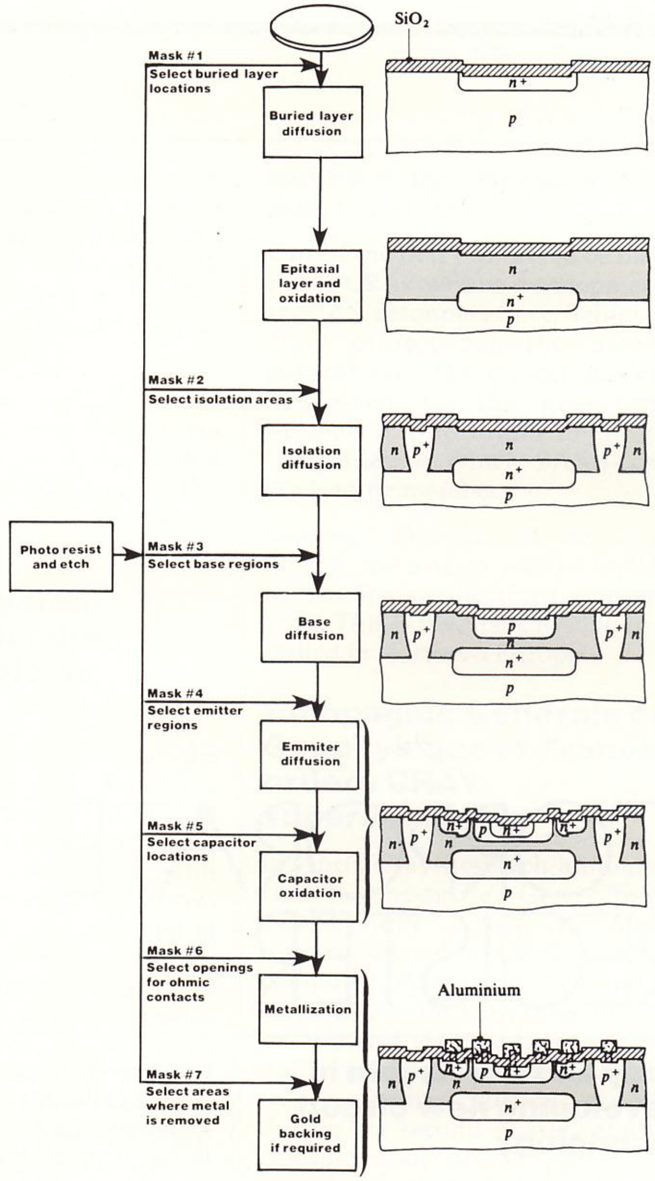

Figure 1 Flow chart of the sequence of processes used in the fabrication of a single-crystal circuit.

later be wirebonded. The wafer is then sawed into small chips, each of which is a single integrated circuit. The chips are then mounted on the case, or header, and wirebonded. The circuit is now tested and readied for use. □

Acknowledgement

This article was prepared with assistance from Delbert Eberlein, Cray Research, Inc.

References

Hamilton, Douglas J., Howard, William G., **Basic Integrated Circuit Engineering**, New York, McGraw-Hill, Inc., 1975. Portions reprinted here in modified form with permission.

Warner, Raymond M., Fordemwalt, James N., **Integrated Circuits Design Principles and Fabrication**, New York, McGraw-Hill, Inc., 1965.

CORPORATE REGISTER

Cray and Harris join in developing new circuit technology

Cray Research and Harris Microwave Semiconductor, Inc., a subsidiary of Harris Corporation, have signed an agreement on a joint program to explore the use of gallium arsenide (GaAs) integrated circuits in supercomputers. The Harris unit specializes in the development and manufacture of GaAs products for communication and information processing applications. The two companies will develop concurrently several types of GaAs circuits using compatible designs and processes.

Commenting on the agreement, Chairman John Rollwagen said, "In the realm of supercomputer architectures, it seems that the speed limitations of silicon are now being reached. The trend there is for greater densities, without significant changes in speed. Consequently, newer technologies — like gallium arsenide — are being explored. GaAs seems to be the technology of the future as far as speed goes, but it will be coupled with other technologies as well, such as CMOS."

Dr. Richard Soshea, vice president and general manager of the Harris unit added, "GaAs integrated circuits can operate at speeds up to five times that of silicon-based equivalents. This advantage is being increasingly recognized in the design of microwave and high-speed signal processing systems. GaAs technology is now sufficiently well-established to encourage its application in very-high-speed computer systems like Cray products."

Cray Research to install CRAY-1 at Ecole Polytechnique in France

Cray Research, Inc. recently announced that an association of French educational and governmental research organizations ordered a CRAY-1 S/1000 computer system. The system has been installed temporarily at CISI, a leading French service bureau that installed its own CRAY-1 in 1981. By late summer, the newly installed CRAY-1 will be moved to a new computer center at L'Ecole Polytechnique, a prestigious engineering school in the Paris area specializing in computers and robotics.

Among the users of the newest CRAY system in France will be the National Ministry for Education; the National Ministry of Transport; which encompasses the Meteorological Weather Bureau; the National Center for Scientific Research (CNRS); the National Office for Aerospace Research (ONERA) and the National Institute of Research and Automation (INRIA).

Cray Research Chairman John Rollwagen said, "We are impressed with the French recognition of the importance of providing supercomputer resources to the research and academic communities. These French organizations have demonstrated their eagerness to have a CRAY-1 system by forming an association for the express purpose of acquiring the system."

Phillips Petroleum Company orders first CRAY-1/M system

During March Cray Research announced that the Phillips Petroleum Company placed an order for the installation of a CRAY-1/M computer system. The two-million word CRAY-1/M system with three I/O Processors will be installed at Phillips' Corporate Information Center in Bartlesville, Oklahoma during the first quarter of 1984.

Phillips is engaged in petroleum exploration and production on a worldwide basis, and petroleum marketing and refining in the United States. The company produces and distributes chemicals in the United States and conducts business in 27 other countries. Since the early 1970's, the company has been developing businesses in other energy fields, including coal, oil shale and geothermal power.

John Rollwagen, Cray Research Chairman, said the order from Phillips was significant for two reasons: "It is the first announced order for the new CRAY-1/M since the product was introduced in September, and it represents further evidence of the acceptance of CRAY-level computing in the petroleum industry."

Cray Research to install two systems in the UK

A CRAY-1 S/1000 has been ordered by Honeywell UK under a contract with the Royal Aircraft Establishment at Farnborough in the United Kingdom. The system, which will be purchased, will be installed at Farnborough during the first quarter of 1984. The RAE is the largest government research establishment in the UK, participating in national and multinational development programs and collaborative research. The CRAY-1 at Farnborough will aid RAE's research in aerospace technology.

In addition, Cray will install a second CRAY computer at a research establishment of the United Kingdom Ministry of Defense. The system will be installed during the fourth quarter of 1983.

CRAY-1/M system selected by leading French petroleum company

Cray recently announced that the Societe Nationale Elf Aquitaine (Production) ordered a CRAY-1/M computer. The two-million word system will be configured with three I/O Processors. The system will be purchased and installed at the customer's facility in Pau, France during the first quarter of 1984, pending export license approval. The Societe Nationale Elf Aquitaine (Production) is France's leading petroleum company and plans to use the computer for reservoir modeling and seismic work.

Cray Research Chairman John Rollwagen said, "This is the first international order we have received for the CRAY-1/M. We're very pleased to provide Elf with the computing power they now require."

Cray to install first system in Scandinavia

SAAB-Scania AB, one of the leading private industrial companies in Sweden, will purchase a CRAY-1A computer system. The company is a major manufacturer of aircraft, automobiles and trucks. The Cray will be used in the development of these products. Among the company's recent product introductions is the 340 commuter aircraft which was introduced in late 1982. The fuel-thrifty aircraft was developed by SAAB with the cooperation of Fairchild Industries.

One of the first projects to be moved to the CRAY will be development of the JAS reconnaissance/fighter aircraft project for the Swedish government. The project was commissioned by the government through a consortium of Swedish companies, of which SAAB-Scania is a leading member.

Pending approval of an export license, the system will be installed in Sweden in the third quarter of 1983. This is the first CRAY to be installed in Northern Europe.

Compagnie Generale de Geophysique of France orders CRAY supercomputer

In April, Cray Research announced that Compagnie Generale de Geophysique (CGG) of Massy, France ordered a CRAY-1 S/2300 computer system. A two-million word CRAY-1/S with three I/O Processors and eight million words of buffer memory will be installed at CGG's facility outside of Paris during the second quarter of 1984, pending export license approval.

CGG is a leading French geophysical company with major operations in France, Canada, the U.S. and the U.K. CGG will use its CRAY-1 for seismic activities. The system is the ninth order received from the petroleum industry, indicating strong acceptance of Cray products in that industry.

SSD attachment now available for CRAY-1A and CRAY-1B systems

Cray Research has announced that it will support the attachment of the Solid-state Storage Device (SSD) to all configurations of CRAY-1A and CRAY-1B computers. The attachment is provided by a field-installed option using two channel pairs to deliver transfer rates of up to 80 Mbytes/second.

APPLICATIONS IN DEPTH

BLUEPACK-3D

BLUEPACK-3D, a geostatistical package, was recently converted for use on the CRAY by GEOMATH, INC. The program is running at Lawrence Livermore National Laboratory where the conversion was done. The primary function of BLUEPACK-3D is Best Linear Unbiased Estimation (kriging) of a spatially correlated variable in three dimensions. This is applicable to the analysis and interpolation of stationary and non-stationary phenomena sampled in 2- or 3-D. BLUEPACK-3D provides users with the best estimates between known data points and the variance of those estimates. Variables can be selected within specified limits and rejected or transformed according to different kinds of criteria. Concurrently, the code validates output with a hypothetical model and indicates those data points that may be suspect based on the model.

The system works interactively building a journal file of parameters that can be processed and updated in time-sharing or batch mode. The program can be prerun in order to build the file of parameters and validate their syntax without executing the calculations. An on-line user guide can be consulted at any time during the execution of the program.

Geostatistics can be used in a number of application areas ranging from petroleum production to meteorology and forestry. BLUEPACK-3D is being used for analysis in petroleum production, hydrosciences, geothermal energy research, coal and shale extraction, 3-D mineral deposit analysis, meteorology and forestry.

Those interested in additional information should contact Dennis Gerken at GEOMATH, INC., 4891 Independence St., Suite 250, Wheat Ridge, CO 80033; telephone: (303) 425-9648, TWX: 910 938 0391 GEO WHET.

MAGNA converted for use on the CRAY

MAGNA, Materially and Geometrically Nonlinear Analysis, is an engineering program that uses the finite element method of analysis. The package is used in a variety of practical applications, primarily in the aerospace industry. MAGNA is particularly useful in impact simulation, damage and residual strength assessment, failure mode analysis and material characterization studies. The system allows one to evaluate complex behavior such as large displacements, "live" loading, material yielding and nonlinear supports with a wide variety of numerical solution options.

MAGNA can also perform linear static and dynamic solutions. Linear static (with multiple load cases), linear transient dynamic and frequency and mode shape calculation are three linear solution options available with the system.

For more information regarding MAGNA contact: Dr. Fred K. Bogner, The University of Dayton, Research Institute, Aerospace Mechanics Division, 300 College Park, Dayton, OH 45469; telephone: (513) 229-3018.

ABAQUS running on the CRAY

ABAQUS is now available for operation on CRAY computer systems. The package is a structural and heat transfer analysis system that includes a large range of material and geometric modeling capabilities. Models may include beams, shells and continuum elements with large rotation – small strain capability. Important material behaviors allowed include plasticity, creep and swelling.

The program includes both user and automatic control of step size in statics and dynamics. Both steady-state and transient heat transfer analyses are possible.

A complete set of plotting options, print controls and versatile restart capabilities are included. Input is free-format, keyworded and uses set definitions for easy cross-reference. Unique controls for simplifying use in nonlinear analysis are possible.

For additional information about ABAQUS contact: Bengt Karlsson, Hibbitt, Karlsson, & Sorensen, Inc., 35 Angell Street, Providence, RI 02906; telephone: (401) 861-0820.

Cray Research to host image processing conference

On July 14-15 Cray Research will host an image processing conference at the Mendota Heights facility. Display device performance in a CRAY computing environment will be demonstrated on Ramtek, Viacom and Evans & Sutherland equipment. Results from SADIE and GIPSY image processing software packages will be shown. The program will also include a variety of talks to be given by experienced image processing users. Those interested in attending the conference should contact: Bill Samayoa, Cray Research, Inc., 1440 Northland Drive, Mendota Heights, MN 55120; telephone: (612) 452-6650.

APPLICATIONS HIGHLIGHT

UNIRAS

Ah, the dilemmas of modern times. Sure, supercomputers can do lots of analyses never dreamed possible. But how about understanding the results? Reams and reams of numbers can be overwhelming without being enlightening. So computer graphics, once a new area of computer technology, has become an important research tool.

Enter UNIRAS, a graphics package recently converted for use on CRAY systems. The package is a powerful tool for users running analyses like seismic, kriging and image processing on the CRAY. UNIRAS is entirely raster-oriented — all vector-to-raster conversion (VRC) is performed by the software computation. This important capability distinguishes UNIRAS from other graphics systems, and saves researchers valuable programming steps and time in preliminary processing and analysis. The system is organized into FORTRAN-callable subroutine libraries for seismic, image processing, mapping and presentation graphics applications. Raster images can be generated quickly at the front-end of data processing, an important feature for these data-intensive applications. Rasterization of line data is especially well-suited to running on CRAY systems because it involves extensive computation.

Vector-based and raster-based comparison

Vector-based packages typically sort vectors to maximize efficiency of the plotting or display device. Adding or removing information requires re-sorting the vector data base. Not only that, before sorted vectors can be sent to a raster output device, they must first be processed through a VRC. For color output, VRCs are largely device-dependent. Color information must be given in a format the device will understand.

UNIRAS, on the other hand, converts input data for raster output and stores the new data in an internal database. In essence, the database is a virtual plot area whose physical properties can be controlled at will. The internally managed raster data base and internal color cross-reference table in UNIRAS avoid the vector sorting and conversion problems. The system handles all conversions for raster and high resolution devices such as ink-jet and electrostatic printers or laser film devices.

The system

Raster device independence is a key feature of UNIRAS, allowing it to produce pictures for any raster device. The system, however, cannot be used with purely vector-oriented devices. UNIRAS takes the pixel size of the actual plotter device into account and creates images that take advantage of the resolution of the device regardless of raster display resolution.

Another important feature gives the user full control of the number of bit planes and the way in which multiple bit planes are used to generate colors and/or intensities. The application program to produce a particular picture is the same for all raster devices; only the UNIRAS drivers differ from device to device. Just one version of the program is needed regardless of the number of displays.

The final raster scan plot has a unique advantage over conventional plotting methods. The raster binary image can be stored on an auxiliary storage medium and accessed for additional plot copies in a fraction of the time required for a vector plot. This raster image format means that processing can be done in a variety of ways. For example, a common background can be logically merged with varying data to produce a set of plots.

The five applications packages contained in UNIRAS — SEISPAK, GEOPAK, BIZPAK, KRIGPAK and GIMAGE — all depend on RASPAK, the kernel package.

RASPAK features

RASPAK provides the tools for operations such as VRC, area shading, point plotting and a variety of text and lettering capabilities. It also contains utility routines required by the application packages. These include the raster processing technique, device control and scaling of data. Plot image storage and a technique for fast area shading of any specified polygon with defined patterns and quadratic surface filling are also included in RASPAK.

Applications packages

SEISPAK is a general purpose subroutine package that applies color graphics to seismic data. The program performs wiggle plotting and raster scanning and can plot negative and positive amplitudes. Each data point along a trace is represented by the raster-scanner as a small color rectangle, which is very cost-effective when producing millions of samples.

GEOPAK is one of the most powerful UNIRAS packages, providing 2- and 3-D mapping capabilities. In addition to generating grid and contour maps in two and three dimensions, highly advanced routines interpolate irregularly formatted data into gridded data. GEOPAK accepts unlimited data input — 3-D contour maps have been processed in color from a data base with more than 1.5 million data points.

With BIZPAK a user can easily produce time consuming graphics such as shaded bar and pie charts, colored backgrounds, clustered bars, thick colored solid lines and 3-D graphs. Other important routines are included and programming effort is reduced with the comprehensive automatic parameter setting system.

GIMAGE can be used in the field of remote sensing, seismic exploration and in mapping applications with large gridded data bases for black-and-white or color image processing. The color and data values in a grid map indicate values assigned to a given "cell". In processing Landsat data, one such "cell" might consist of a one-acre area of the earth's surface. Each acre or cell then becomes a color (or black-and-white) representation of Landsat signets that indicate varying attributes of the ground cover — vegetation, water, and so on. An image composed of many cells showing the distribution of elements across the terrain is useful in land development or environmental impact studies.

KRIGPAK, based on kriging, produces highly reliable geostatistical information. (Geostatistics is the popular name for the applications of modern statistical methods in the disciplines of geology, mining and environmental studies.) With the kriging technique, KRIGPAK employs a stepwise procedure that gives the user full knowledge of the data's spatial structure.

Conclusion

UNIRAS is being used by those engaged in activities as diverse as petroleum exploration, planetary gravity mapping, broadcasting election results and even computer-generated artistry. Running on the CRAY, UNIRAS helps sort out complex information that is best understood in the raster format. For more information about the UNIRAS system in the United States and Canada, contact: American Software Contractors, Inc., 48 Cummings Park, Woburn, MA 01801; telephone: (617) 933-6102. In Europe, contact European Software Contractors A/S, Lyngby Hoved gade 15C, DK-2800, Denmark, telephone: +45 2 93 01 33.

USER NEWS

CRAY BLITZ - world computer chess title contender

How does CRAY BLITZ do it? Does it have a positive mental attitude that just can't be beat? Or is it just nimbler and shrewder than most? Whatever it is, the prodigious program is making its mark. In CRAY CHANNELS Vol. 4 No. 3, we reported that CRAY BLITZ drew the final match against BELLE of Bell Laboratories in the 1982 world chess tournament sponsored by the Association for Computing Machinery. Bob Hyatt, who has been competing with CRAY BLITZ since 1980, and Harry Nelson at Lawrence Livermore Laboratory refined the program for 1982 and are fine-tuning it for the 1983 tournament. We thought we'd find out exactly what they've been up to with our favorite chess program, CRAY BLITZ.

Speed and smarts are key factors in the performance of a computer chess program. CRAY BLITZ has the speed, processing an average of 10,000 nodes/sec. (At last report, its closest competitor, CDC 4.9, was operating at 2,700 nodes/sec.) But Bob and Harry figured that if CRAY BLITZ could run even faster, the extra time could be used to analyze each move with even greater depth. That would give the program an additional edge against BELLE. While BELLE (a computer hard-wired exclusively for chess) can look at about 160,000 nodes/sec, it is not a particularly smart program. Enhancements to CRAY BLITZ's intelligence coupled with greater speed give it a good deal of clout against BELLE. Revectorizing and recoding parts of the program in CRAY Assembly Language (CAL) were important factors in its 1982 performance.

Bob Hyatt explained that Harry Nelson found certain "hot spots" where the program was spending too much time. Vectorization and CAL-coding decreased the amount of time spent in each of those routines by a factor of 10 to 20. Of CRAY BLITZ's 10,000 lines of FORTRAN coding, 2,000 were designated to be coded in assembly language. Those 2,000 FORTRAN lines turned into 8,000 lines of CAL. Care was taken to ensure that there were no surprises when it really counted. Before entering the tournament, the changes underwent rigorous testing for months, virtually 24 hours a day. Overall, the program is now running about five times faster than before.

Some areas where the program was taking a lot of time were in the Move Generator, Exchange Evaluator and Check Analyzer routines. In the Move Generator, for instance, the program was spending a tremendous amount of time generating potential moves via a tree search. In this operation, potential move sequences are identified that will avoid the loss of a player. A tentative move is selected. The Move Generator then switches sides and analyzes the opponent's potential response moves from the new position. This analysis continues until a fixed depth is reached.

The Exchange Evaluator subroutine was also revectorized and coded in CAL. This routine constantly checks to see if a move to a particular square is safe. The Exchange Evaluator had been absorbing 20-30% of the program's total operation time.

Bob went on to relate that the program is constantly worried about whether its king is in check. In normal play it would spend about 25% of its time per move determining the king's safety. And in tactical situations, up to 50% of the allotted time could be spent in the Check Analyzer. Vectorization has also significantly reduced the amount of time spent in this routine.

Intelligence enhancements that have recently been added will augment the program's capabilities in 1983 play. Bob Hyatt explained one of these: "Most computer programs don't recognize the difference between a "good" and "bad" bishop. In the game of chess, a "good" bishop is one whose movements are not impeded by its own pawns. A "bad" bishop is one that is restricted in the moves it can make because its own pawns are on the same color square it is, thus blocking it. If the program can recognize this situation, it can either avoid or correct it. Until now, CRAY BLITZ couldn't do that. By adding a simple check to determine the potential for this situation, problems can now be avoided." Bob went on to say, "To my knowledge, no other chess program has this particular checking capability. It's a nice feature to have, but not as necessary to the game as others. The routine takes up valuable time that most other computer chess programs simply cannot spare."

"CRAY BLITZ also plays an excellent end game," Bob said. "When there are only a few pieces left on the board — like a couple of pawns and a king, it plays a game that is far superior to other programs."

Hyatt is hopeful that CRAY BLITZ will be able to run on the CRAY X-MP in this year's tournament. He explained that the move evaluation time should be cut in half by using the parallel power of the X-MP. "Without doing anything to the program we should see time improvements of about 30% with the faster clock period. But what we are really waiting for is the CRAY-2."

28th Mersenne Prime found

The CRAY computer will keep its place in The Guiness Book of World Records as the discoverer of the world's largest prime number. One Saturday morning several months ago, the CRAY nonchalantly printed out the message: 2^{86243} - 1 is a prime number (containing 25,962 digits). This event was no small cause for celebration, although the CRAY had no way of realizing it. Over 600 hours of CRAY time divided among several systems had been committed to the search for the 28th Mersenne Prime. This marks the second time that a CRAY system has discovered a largest known Mersenne Prime. The 27th was found in 1979.

Many people have been involved in the efforts leading to the discovery of both numbers. Cray personnel, interested parties at several user sites and several CRAY systems have all taken part in the projects. Among them, Harry Nelson of Lawrence Livermore National Laboratory and David Slowinski, a Systems Analyst at Cray Research, have been most interested. Nelson and Slowinski collaborated on writing the program that found the 27th Mersenne Prime, and Slowinski was largely responsible for orchestrating the effort in finding the 28th.

The algorithm used to execute the Lucas-Lehmer test for primality was streamlined in the search for the 28th Mersenne Prime. Dave Slowinski explained, "About 90% of the time is spent squaring 25,000-digit integers. By optimizing squaring procedure in the program, I was able to improve performance by about 40% over the first program."

Slowinski had estimated that it would require 2,000 hours of computer time to hunt for the 28th Mersenne Prime, but with good luck and optimized code it was discovered in only 600 hours. To reduce the problem of finding the number to a workable scale, Slowinski divided the task among several CRAYs around the country, each one looking at some small part of the Mersenne sequence. The systems executed the program only when the computer was idle. Slowinski says that the program is also a useful confidence test for the hardware because it computes a residue check that catches many different kinds of failures.

For additional reading about David Slowinski's work with the 27th and 28th Mersenne Primes see CRAY CHANNELS, Vol. 4 No. 1, *"Searching, for the 27th Mersenne Prime"*, David Slowinski, pp. 15-17, and DISCOVER, Vol. 4 No. 2, February 1983, *"Biggest Prime, Longest Pi"*, Bruce Schechter, pp. 92-93.

A CRAY challenges a TRS-80

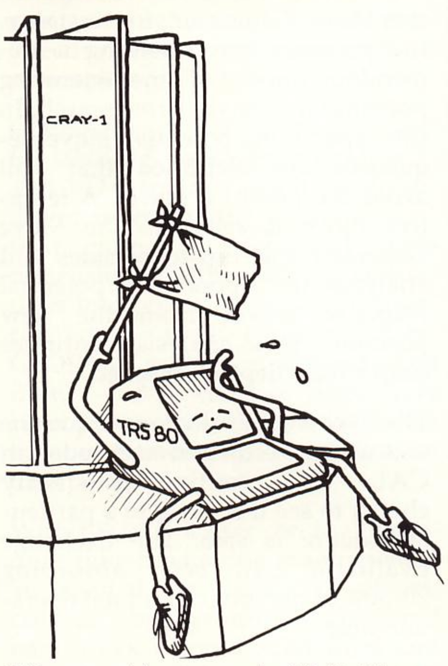

Why would anyone do it? Putting a lil' old TRS-80 just minding its own business up against a CRAY just isn't fair. And yet there are some people in the world who'll try just about anything. A Cray analyst who shall remain nameless is one of them. The gentleman benchmarked the performance of a TRS-80 Model 3 16K computer against the CRAY-1 S/2400, explaining that the unusual action was prompted by his son's school teacher who asked him to talk to the class about computers. He thought it would be fun comparing the performance of the school's microcomputer to the CRAY. He thought it would be something that the kids could relate to. Uh huh.

For his comparison, the analyst executed a 50 x 50 matrix multiply in both machines. Then he compared the time and cost of processing. The results are found below.

Machine	Execution Time (in seconds)	Per Unit Cost
TRS-80 Model 3	3180	$999
CRAY-1 S/2400	0.002	$7,620,000

Performance Ratio: 1/1,590,000
Cost Ratio: 1/7628

It was reported that screams of "Uncle, Uncle" could be heard from the school's computer room where the TRS-80 met its fate. The class was duly impressed with the speed of the CRAY. So was the TRS-80. The CRAY executed the problem approximately 1.6 million times faster than its distant cousin.

But you say, "Nobody expects the TRS-80 to be able to tackle a problem like that." We agree. However, there is a similarity between the systems in that they are both innovative wonders in their class. And the one basis for comparison may be the price/performance ratios of the systems. Based on the results, in order for the two to have equal price/performance ratios, the CRAY-1/S should cost 200 times more than it does, or the TRS-80 should cost 200 times less. While the classroom audience was not interested in that bit of trivia, we thought you might be.

Japanese CRAY users create a variety of images with computer graphics

What connection could there possibly be between Buddhism and a CRAY-1? One we've found is computer graphics. Recently, Seibu Promotion Network Company (SPN) of Japan generated some very special graphics with the help of Cray equipment. The two clients for whom the graphics work was done were a broadcasting company and a publisher. Logos and Buddhist meditation images were to be produced. SPN is a frequent user of CRAY equipment at the Mitsubishi Research Institute (MRI) in Japan.

A representative from SPN explained that their company, with promotion responsibilities for the Seibu Consumer Group in Japan, recognizes information as a very tangible product. SPN is preparing to deal in this commodity. As a first step in that direction, SPN has formed a computer graphics division to develop new visual tools for the expression and visualization of information.

Currently, SPN is polishing their software that is being built into a generalized raytracing package for operation on CRAY computers. Because tracing a picture in high resolution can result in handling tens of millions of vectors and calculating hundreds of millions of intersections, they feel that it is important to have access to CRAY-level computing power. Current development is about 50% in FORTRAN and 50% in the language "C." The latter code is being run on a VAX 750, but SPN anticipates that it will soon be operational on the CRAY. (Many programmers feel that "C" is a more expressive language for their graphics algorithms.)

The photographs below illustrate a few of the 400 frames used in generating the 12 second opening title of "NC9", a popular Japanese news program shown on NHK (Japan Broadcasting Corporation).

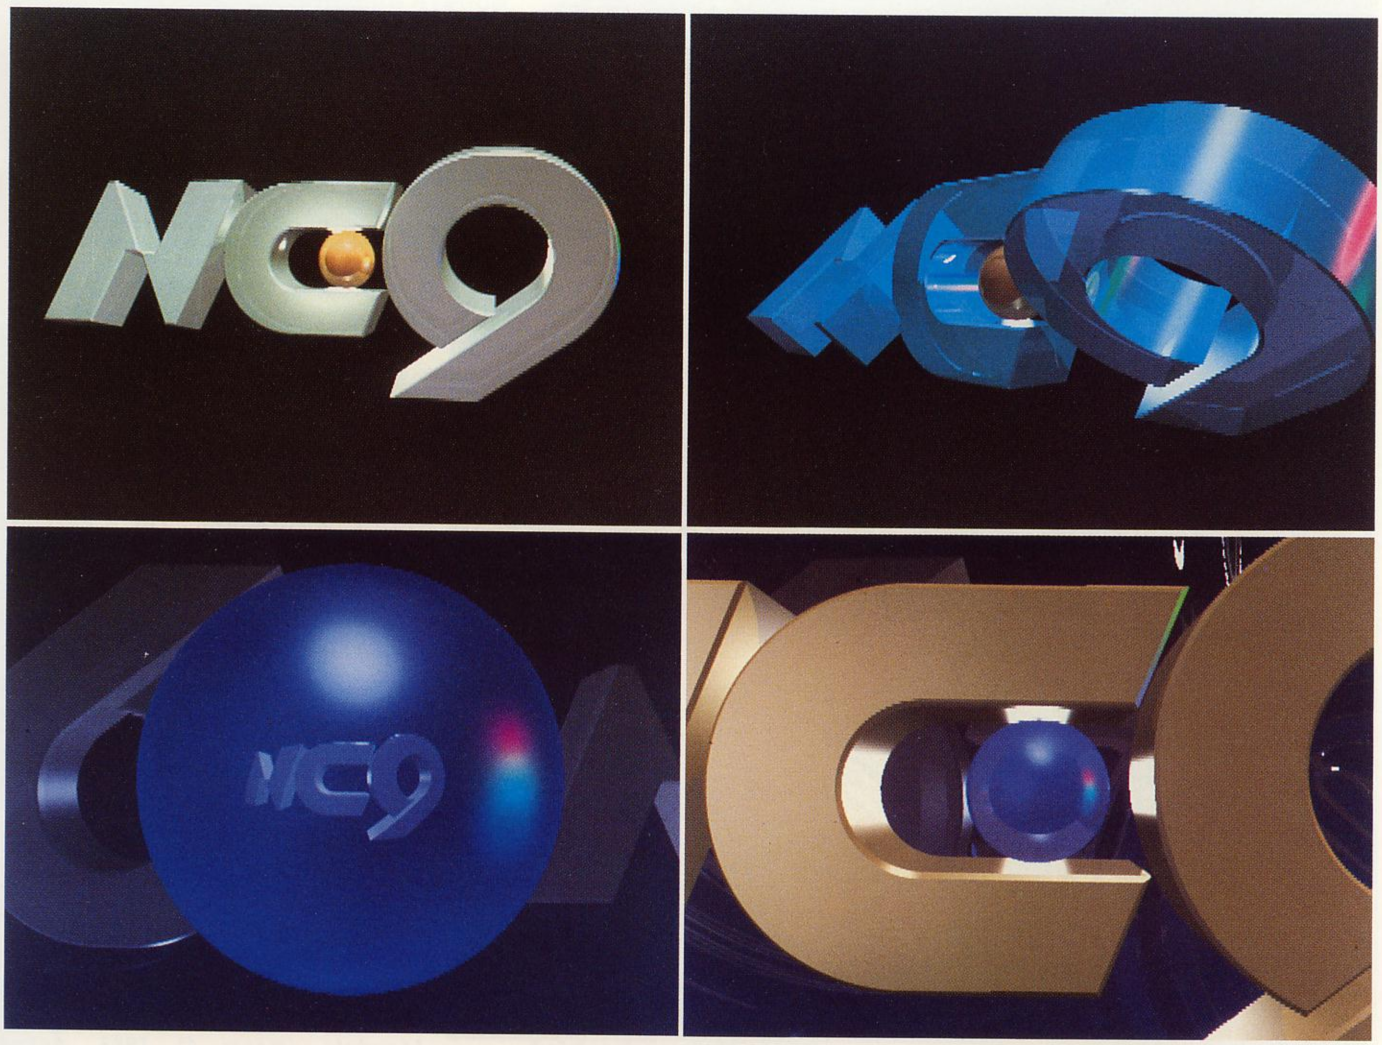

NC9 is a logo graphic recently generated on a CRAY.

USER NEWS

In addition, the group from SPN created beautiful 3-D conceptualizations of the ethereal Buddhist maṇḍala as illustrated in the photographs below. Maṇḍalas belong to the esoteric tradition of Buddhism, the term itself meaning "that which possesses essence". The maṇḍala is a visual representation of the Buddhist view of the world and is an expression of the world of images met with in meditation. In actual meditation practice the maṇḍala is used as an object of visualization. There are various modes of expression, ranging from the abstract to concrete, including those composed of syllables of the Sanskrit alphabet.

The maṇḍala is a Buddhist object of meditation.

A representative from the publishing company for whom the maṇḍala was created commented, "Computer graphics was used in making this maṇḍala because it was felt that it provided the most suitable technique for expressing the continual unfolding of images as experienced during meditation. The 'moving maṇḍala' thus created can truly be said to be an example of 'universal image language.'"

The image was produced in different resolutions on Ramtek equipment. The photos here were taken at 1024 x 1024 resolution. Who would have thought that ancient philosophy and modern technology could complement each other in such a pleasing way?

Acknowledgements

This article is based on text submitted by SPN Company, Ltd.
Maṇḍala images © 1983, Āgama Publications, Tokyo, Japan.

Corporate Addresses

Worldwide Business Centers

Domestic

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

Albuquerque, NM
Chicago, IL
Minneapolis, MN

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

Dallas, TX

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

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

Western Region Sales 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: 01-501-5270

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

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

Cray Research (UK) Ltd.
Seymour House, The Courtyard
Denmark Street, Wokingham
Berkshire, England
Tel: 0734-791180

CRAY RESEARCH, INC.

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