A CRAY RESEARCH, INC. PUBLICATION

CRAY CHANNELS

Volume 5, Number 3

FEATURE ARTICLES:

REGULAR COLUMNS:

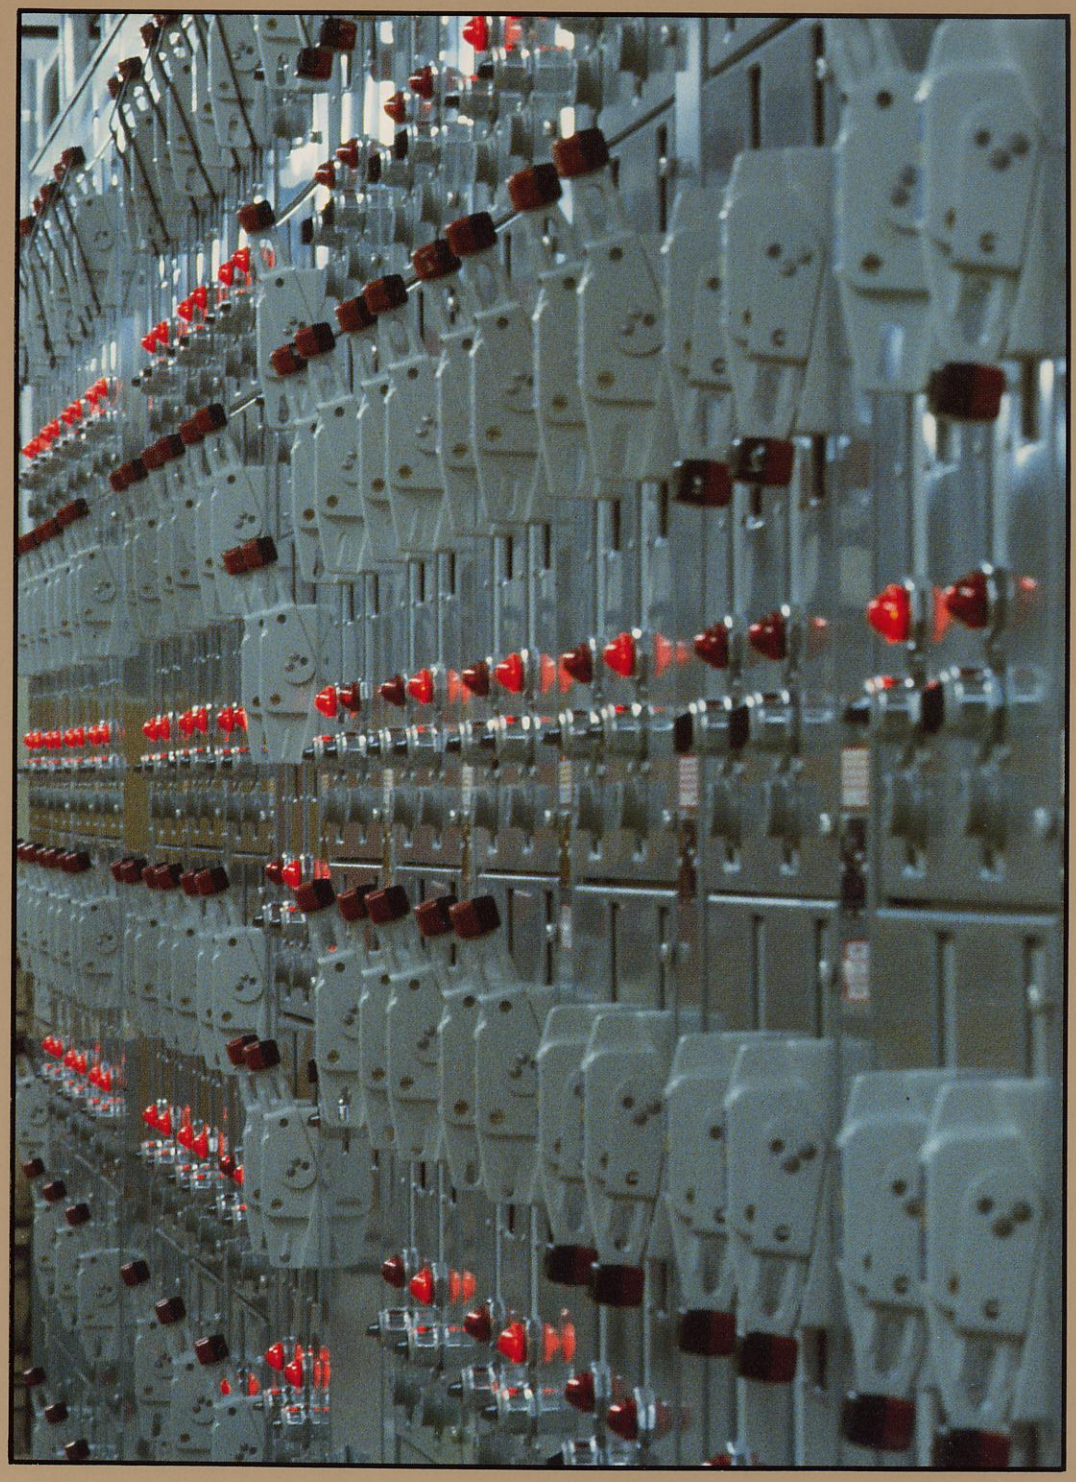

In this issue

Semiconductor technology is now an integral factor in virtually every industry today. From the automotive industry, to the aerospace industry, to medicine, advanced tailored integrated circuits provide enhancements to existing and new products. Certainly, major technical advancements in computer technology will elude scientists without the benefit of more sophisticated circuits and altogether new semiconductor technologies. A slow-down in the introduction of new circuits over the past few years is one bottleneck in the introduction of better products in many industries — including the computer industry.

But the integrated circuit industry itself, is not in a very comfortable position. It now faces rapidly increasing development costs as technology advances. At the same time, profits from the sale of new products are decreasing. Cray products both contribute to and benefit from advances in the field of semiconductor technology. Our computers are important tools that help restore efficiencies to the integrated circuit design cycles.

The two feature articles in this issue of CRAY CHANNELS offer a brief study of the IC industry, the impact that supercomputers can have in that area, and the basic design procedures that preceed any new circuit introduction. In preparing these articles, we developed a keener appreciation for the tremendous progress that engineers have made over the past 15 years and hope that you will too. At the same time, we realize that in order to keep pace with the accomplishments of the past, a more than proportionate amount of resources will have to be directed to future semiconductor development.

We hope that you will enjoy reading our regular columns that include news about Cray Research and you. Take special note of the article in User news describing CRAY users' work in semiconductor research. Work that is being done today at two of our national laboratories may someday affect the shape of computing in the future.

About the cover

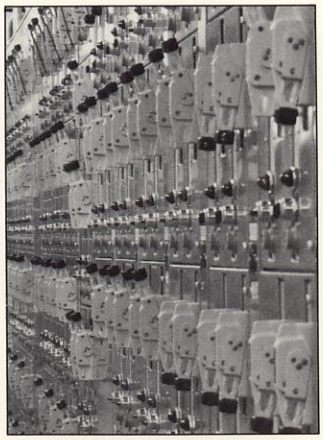

Not your average set of bells and whistles

This imposing wall of switches, buttons, and lights is the starter and breaker panel of the motor control center at Cray's IC fabrication facility in Chippewa Falls, Wisconsin. At the heart of the control center is the computer programmed to monitor and control much of the equipment operation. Gear under the command of the system includes fans, water pumps, air scrubbers, environmental controls, and some of the process equipment. The system is programmed to handle emergency and daily operations. For instance, when the clean room is not in use, the control center's 'unoccupied status' function is activated. In this mode, the air scrubbers and fans operate at a lower level, saving energy and money.

CRAY CHANNELS

A Cray Research, Inc. publication Volume 5, Number 3

CONTENTS

FEATURE ARTICLES

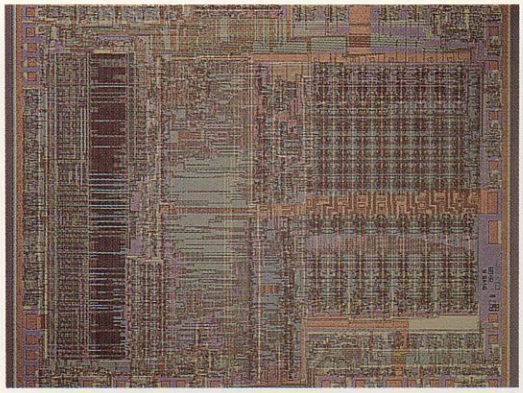

2

REGULAR COLUMNS

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.

Development turnaround: a

Kelly Wild, Cray Research, Inc.

In this article, CRAY CHANNELS looks at the challenges facing the technical, fast-moving world of the semiconductor industry, and how CRAY systems offer a competitive edge.

an important factor in the IC industry

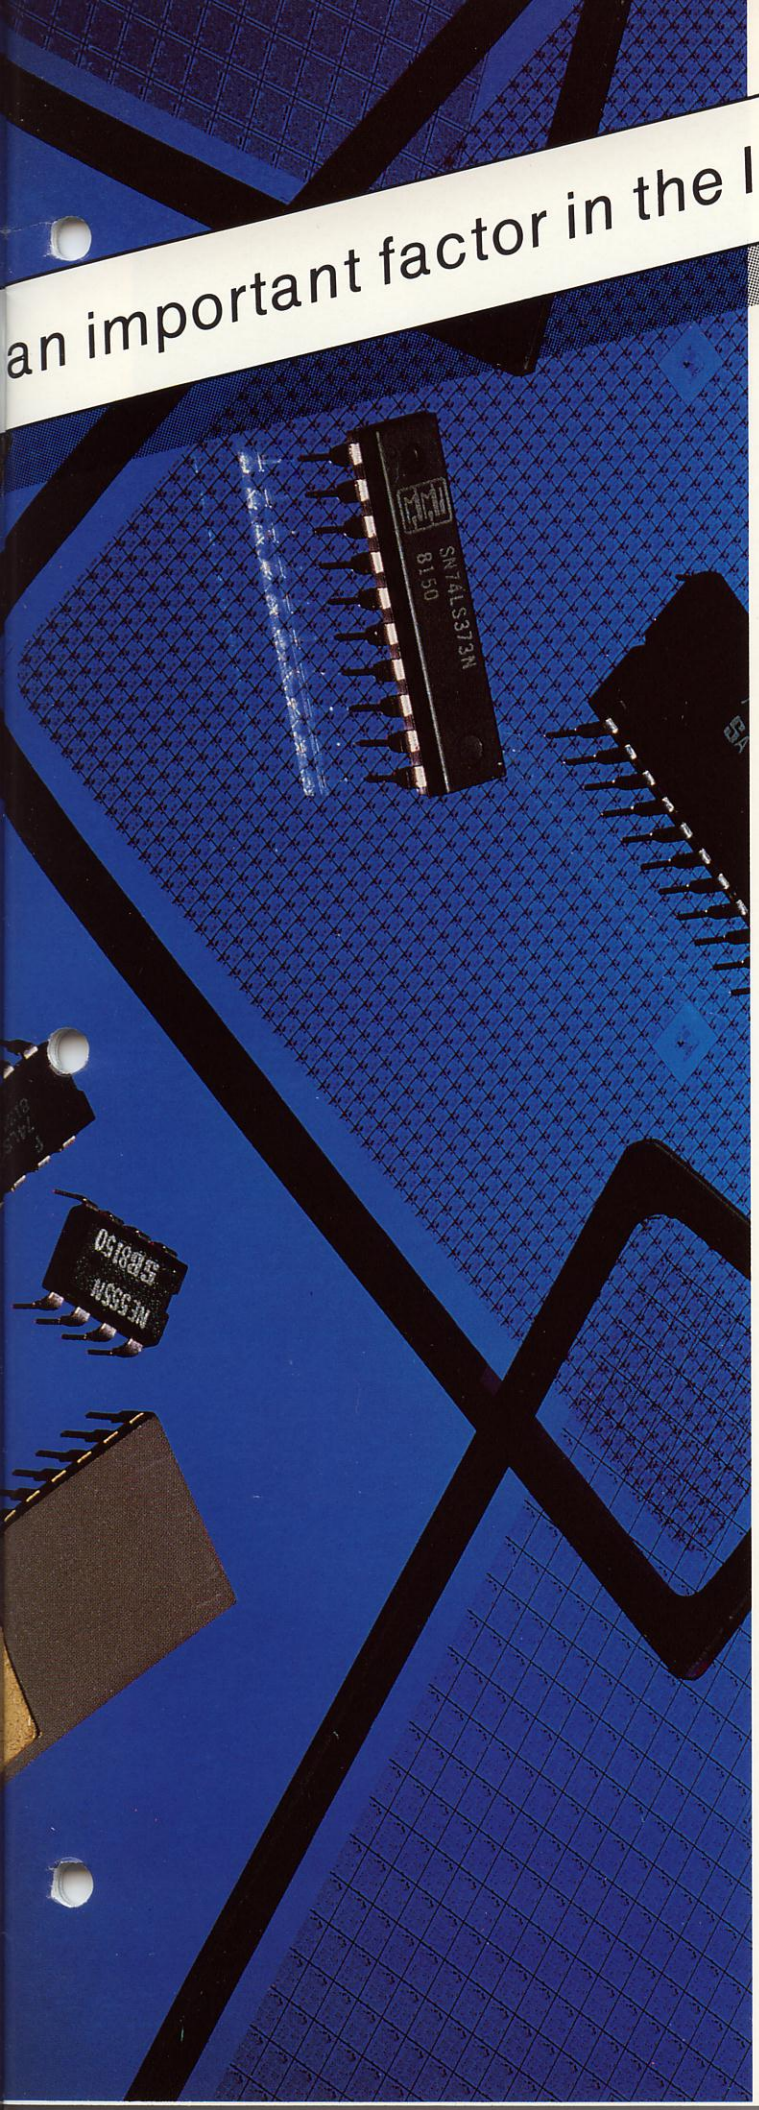

Introduction

Paradoxical but true: as microcircuit development becomes more complex, the end product of microcircuit technology — the computer — becomes essential for further development. No one need be reminded that integrated circuit (IC, chip) technology has launched the second industrial revolution, which promises to bring greater social changes than the first ever did. Machinists of an earlier era built special purpose tools to help them create their end products. Not unlike those machinists, computer designers have also developed special tools that run on the most powerful and sophisticated computers.

Industry background

It was in 1958 that IC technology came into its own. In that year, Jack Kilby at Texas Instruments devised a method of building a complete electron circuit on a single piece of semiconductor. The integrated circuit was born. Then in 1959, Robert Noyce at Fairchild Semiconductor unveiled another technique of forming and interconnecting circuits on slices of silicon. By the 1960's the race was on to pack more and more devices on each chip.

Today, IC technology is no less dynamic than it was two decades ago and the goal is the same: development of more powerful and compact chips. Worldwide competition is fierce, especially between manufacturers from the United States and Japan, the primary IC producing countries in the world today. Ten years from now, industry forecasters expect the semiconductor industry to be one of the largest in the world, with annual sales expected to exceed $90 billion.

For many years, the doubling of IC capabilities occurred annually. However since the development of the 100,000-device circuit in the late 1970's, that pace has slowed. It now takes more time than ever before to develop the new complex ICs. While compute power per chip has increased dramatically, prices for the new products have dropped at the rate of 20 to 30 percent annually.

The rewards for buyers are great as higher levels of integration are achieved. The most obvious benefit is the dramatic decrease in the cost per function as the cost of a chip is amortized over a larger number of functions. Equally important are the increases in performance and overall reliability that accrue as devices shrink in size. By some estimates, densities can increase by a factor of 100 before the fundamental limits of chip technology are reached.

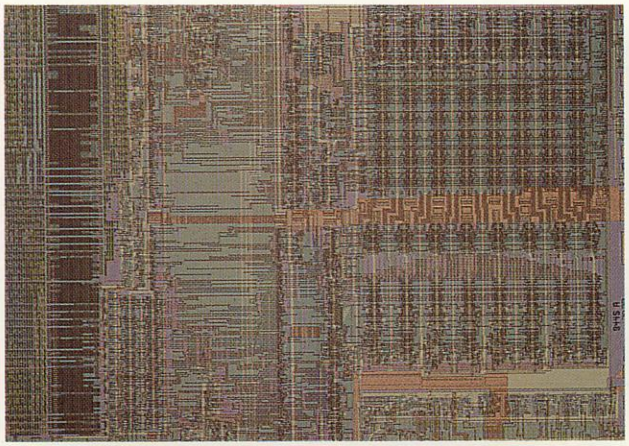

A high-speed injection logic 16-bit microprocessor.

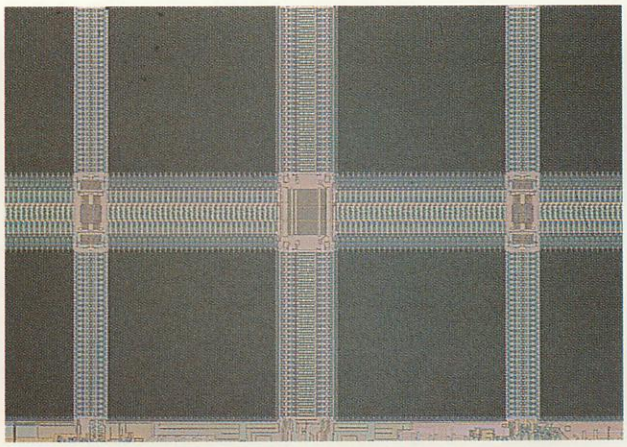

A low power 64K Dynamic Random Access Memory (DRAM) integrated circuit.

As ICs become more complex, the amount of analysis (and subsequently, the computer power) needed for their development grows exponentially. Supercomputer power is becoming a major component in microelectronics development. The large memories and tremendous computing power of CRAY systems provide the circuit designer with the power to develop new, more reliable ICs in less time. Time-consuming simulations that may not have been practical on other systems due to time constraints, can now be run efficiently on the CRAY. Using a CRAY, the engineer may execute many different variables quickly and know all the behavioral characteristics of the chip before the long process of fabrication ever begins. Able to scrutinize every aspect of a chip, an engineer can be more confident of his design and produce a better, more reliable product sooner.

The integrated circuit design problem

Whereas in the 1960's, component dimensions on circuits were about 25 microns wide, today they are typically three microns wide, and manufacturers are now moving to 1.5 micron technology. In the late 60's and early 70's the industry moved from small and medium scale component integration to large scale integration where circuits contained 1,000 to 100,000 devices. Today, very large scale integration (VLSI), where the number of devices on a circuit ranges from 100,000 to ten million devices, is now standard in the semiconductor industry. By the end of the century, researchers expect to introduce ultra large scale integration which will cross over into ten-million plus devices on a single chip.

The challenge is that engineers can barely design state-of-the-art circuits with existing tools. The design problem can be likened to the challenge of designing, mapping, and monitoring the highway system for metropolitan Los Angeles in an area 1/20 the size of a postage stamp.

After mapping all the streets and highways onto that small area, one would have to keep track of where each street led. All possible routes to any given location in the micro city would have to be tracked. In addition, the length of time required to arrive at any location taking any possible route would have to be known. One would be required to monitor all traffic lights in the system and note which ones were red, yellow, or green at a given point in time. Information about which streets had high or low volume traffic, in which direction, and at what times of day would have to be available upon request. It's easy to see that it would require large-scale compute power to structure and understand the dynamic behavior of this intricate system. This problem is not unlike that facing integrated circuit design engineers today.

Now that it is possible to put a complete microprocessor on a chip, the designer must be concerned with a much broader range of problems. Compared to the days when circuit design consisted of a combination of device physics and circuit layout, the scope of the task has expanded to include logic design, machine organization, and even system architecture. A balanced strategy encompassing tools, methodologies, and quick-turnaround fabrication is required in the IC development environment today. In order to devise increasingly more sophisticated circuits, design technologies will have to evolve just as the nature of the IC design task has.

"By 1991 we'll have in the range of 15 million to 20 million devices on each chip, compared with a half-million on today's 256K RAM," predicted J. Jeffrey La Vell, technology strategist at Motorola Inc.'s semiconductor group in Phoenix. "There are simply not enough chip designers in the world (to design new chips at that level of complexity)." A manager at Intel commented, "Given the complexity of these circuits, the number of hours needed for human design alone will make further advancements virtually impossible."

The challenge facing IC manufacturers

The semiconductor industry is truly a business anomaly. In most industries, as development and production costs go up, prices and margins typically rise. Thus, development efforts can only continue if manufacturers are confident of recouping their in-

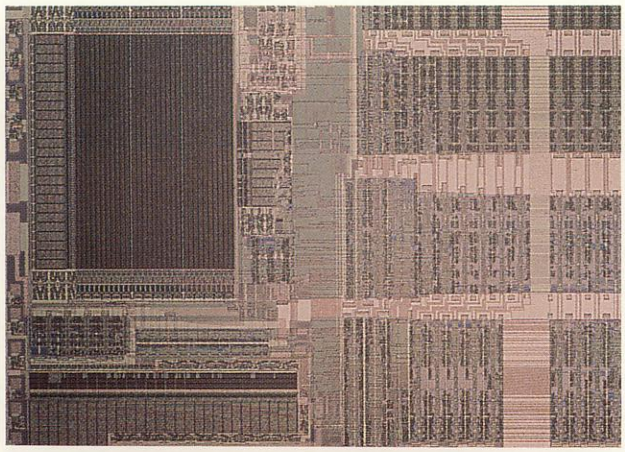

This 16-bit microprocessor has built-in 32 48-bit floating point.

vestments over time. Not so in the semiconductor industry. Research and development spending is exceptionally high (about 10% of sales) and capital equipment costs enormous. To stay in step with technology, most IC manufacturers figure they need to refurbish their plants every two to three years at a cost of $50 to $75 million.

But in the chip business, as these costs have risen, the price of ICs has dropped. The impact of increased density on the price of an IC is dramatic. For example a two-transistor, five-part circuit that cost about ten dollars in 1960 was about ten cents in 1978; it has been said that by 1990 ICs with 20 million transistors will be ten for a penny.

Manufacturing expertise is critical in remaining competitive, but without an advanced product, all the production capability in the world won't guarantee an IC manufacturer a place in the market. Technological competitiveness is just as important as manufacturing capability. The combination of the two are requisite in maintaining a foothold in the industry over the long term. Computers like the CRAY will help IC manufacturers produce more sophisticated reliable products and, at the same time, reduce the time required for their development.

The development cycle as competitive edge

To appreciate the importance of introducing a new product in an aggressive timeframe, consider the hypothetical life cycle for a product as espoused by Merrill W. Brooksby and several of his associates at Hewlett Packard. In their analysis they consider the benefits of quick-turnaround IC processing as illustrated in Figure 1.

The theory rests on the premise that once an idea is formed into a product embodiment, even while still an idea, its potential market is fixed in time. Region 1 designates the time when the idea for the product was conceived, developed, and preparations made for manufacturing build-up. Region 2 is the area of concern. By the time the IC is readily available, demand has peaked and begins to drop sharply as the product matures. At the same time, a new and more cost-effective product may be introduced. The development investment will never have had the chance to be maximized. Not only that, by the time the product comes to market, competitive products will take some share of the total possible sales. In order to maximize return on R & D and manufacturing costs, one would want to bring the product to market earlier in its life cycle.

Figure 2 shows the sales volume that a manufacturer can expect to enjoy when large quantity deliveries begin in 1982 and alternatively, in 1983. Earlier introduction of a new product will result in establishing a strong position in the marketplace. The effect is even more pronounced when the new integrated circuit is the basis of a new or modernized product. The belief in a product's fixed life cycle is a primary driving force for the R & D organization's insistence on shorter product development schedules.[1]

Japanese producers captured the 16K RAM chip market in 1979 when they were able to offer a better product in higher quantities than their U.S. counterparts. One reason (among several) that they were able to do so is that they used large mainframe

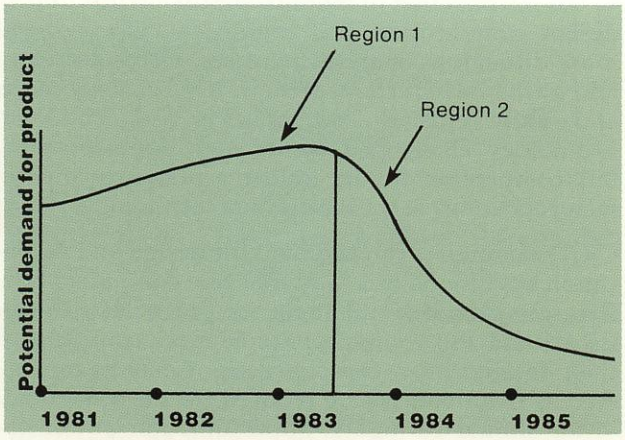

Figure 1. Product life-cycle curve for a hypothetical product. The curve is independent of the product's existence.

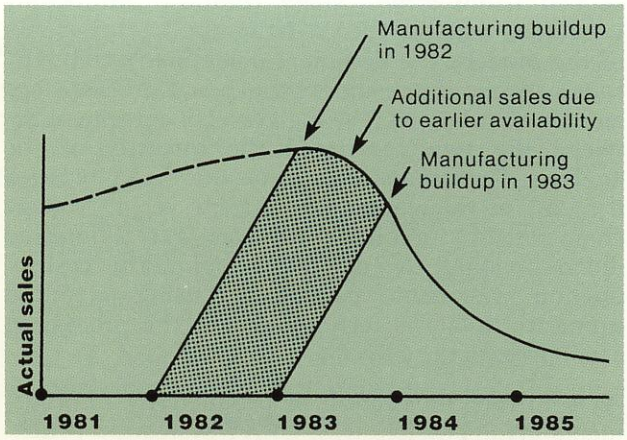

Figure 2. Actual sales achieved if product development is completed at the end of 1983 or one year earlier, at the end of 1982.

computers to design and simulate the complex circuits. When the next generation of computer memories came along in 1980, the Japanese did it again, beating U.S. firms to market with the 64K RAM chip. Japanese producers were able to ride the product life-cycle curve earlier and capture a larger portion of the total potential sales for the product, thus maximizing their return on development and production costs. By 1982, they were producing 70% of the 64K RAMs sold worldwide. Today, CRAY computers installed in Japan and the United States are being used for IC development.

As more advanced circuits continue to be introduced into the market, manufacturers that keep pace in both technical and manufacturing arenas will earn their places in the 1990 marketplace. Supercomputers are now beginning to play a vital role in helping IC manufacturers remain competitive. The boost in compute power allows them to undertake more thorough analyses on more complex designs in a fraction of the time required with conventional systems. CRAY computers can help ease the computing burden while allowing engineers to stay within (or shorten) their design schedules.

Impact of the CRAY in electronic circuit design and simulation

For some time, CRAY computers have been helping integrated circuit research efforts progress at a healthy pace. For instance, Bell Laboratories relied heavily on their CRAY-1/S in the development of the 256K dynamic RAM, one of the largest VLSI devices made to date. The 256K chip promises to be a mainstay in the industry well into the 1980's. Bell Labs simulated a portion of the circuit on its CRAY; the dataset contained 1457 MOS devices. The simulation executed efficiently on the system even though the memory requirements were large. A circuit simulation that would have required 4.4 CPU hours on Bell Labs' Honeywell HIS 6080, was executed on the CRAY in 20 minutes.[2]

In another example, the CRAY-1 played a significant role in the development of 1-μm NMOS technology at Bell Laboratories. Many complex simulations were done with the CRAY. To assure the correctness of the models, the designers made many comparisons between the results of the actual process steps and the simulations. The high level of agreement between the two confirmed the integrity of the analysis. The simulation run on the CRAY reduced the cost of calculating the current versus voltage characteristics for a transistor from several hundred dollars to about two dollars per point. This cost reduction was made possible by using advanced numerical analysis techniques and the processing capabilities of the CRAY-1.[3]

As a further illustration of productivity improvements made possible with CRAY systems, a benchmark was run recently at Cray Research's Mendota Heights, Minnesota facility. A circuit was analyzed with CSPICE running on a CRAY-1/S and SPICE running on a VAX 11/780. (CSPICE is a version of SPICE that has been highly optimized for execution on the CRAY. See page 14 for additional information.) The dynamic RAM circuit was a production-type sense amplifier that had more than 3000 MOS devices. The same dataset was run on both machines. While the benchmark was run on a two-million word CRAY, it could easily have been executed on a one-million word machine. On the CRAY-1 S/2400 the simulation was completed in 34 minutes. The same routine took 4352 minutes (3.02 days) to execute on the VAX 11/780. The performance ratio is 1/128. In many benchmarks comparing CRAY and super minicomputer performance, the CRAY has outperformed the latter system by factors of over 100.

Geared up for the task

Cray products fit well with existing IC engineering workstation environments. The CRAY interfaces with many vendors' systems including IBM and VAX computers. With a standardized data base and netlist extractor, an engineer can quickly lay out his chip and then submit the design to the CRAY for simulation. After simulations are completed and the design finalized, a tape for mask generation can be made on a CAD system in preparation for prototype production.

Many powerful simulation codes have been converted for use on the CRAY. Among some of the major EE design codes already executing on the system are: SPICE, CSPICE, ASPEC, ADVICE, MOTIS, MOTIS-C, TEGAS, NCA/DVS, NEMOS, SUPREM, and SUPREM2. Currently, RSIM/C and SUPREM3 are undergoing conversion. CSPICE, a high performance, highly vectorized code, is only available on the CRAY.

Conclusion

As we progress in the 1980's, the computer has become vital in development of its successors. There is little doubt that without more powerful computers, major advances in microcircuit technology will cease. Already, we are seeing a slow-down in the introduction of more powerful circuits due to design complexities. The engineer's task becomes more difficult as analysis increases, yields decrease, and early introduction of new ICs becomes more and more elusive. Not only is semiconductor technology challenged by its own sophistication, but competition in the industry pressures manufacturers to stay at the forefront of technology.

CRAY computers can augment the design and verification operations in a way that few computers can. Rigorous circuit simulations are practical with the CRAY and allow engineers to be more confident of their designs. More complex designs may be considered as redundancies are reduced. At the same time, yields should be higher. Supercomputer power offered by the CRAY can be a valuable component of integrated circuit design strategy for VLSI and beyond. □

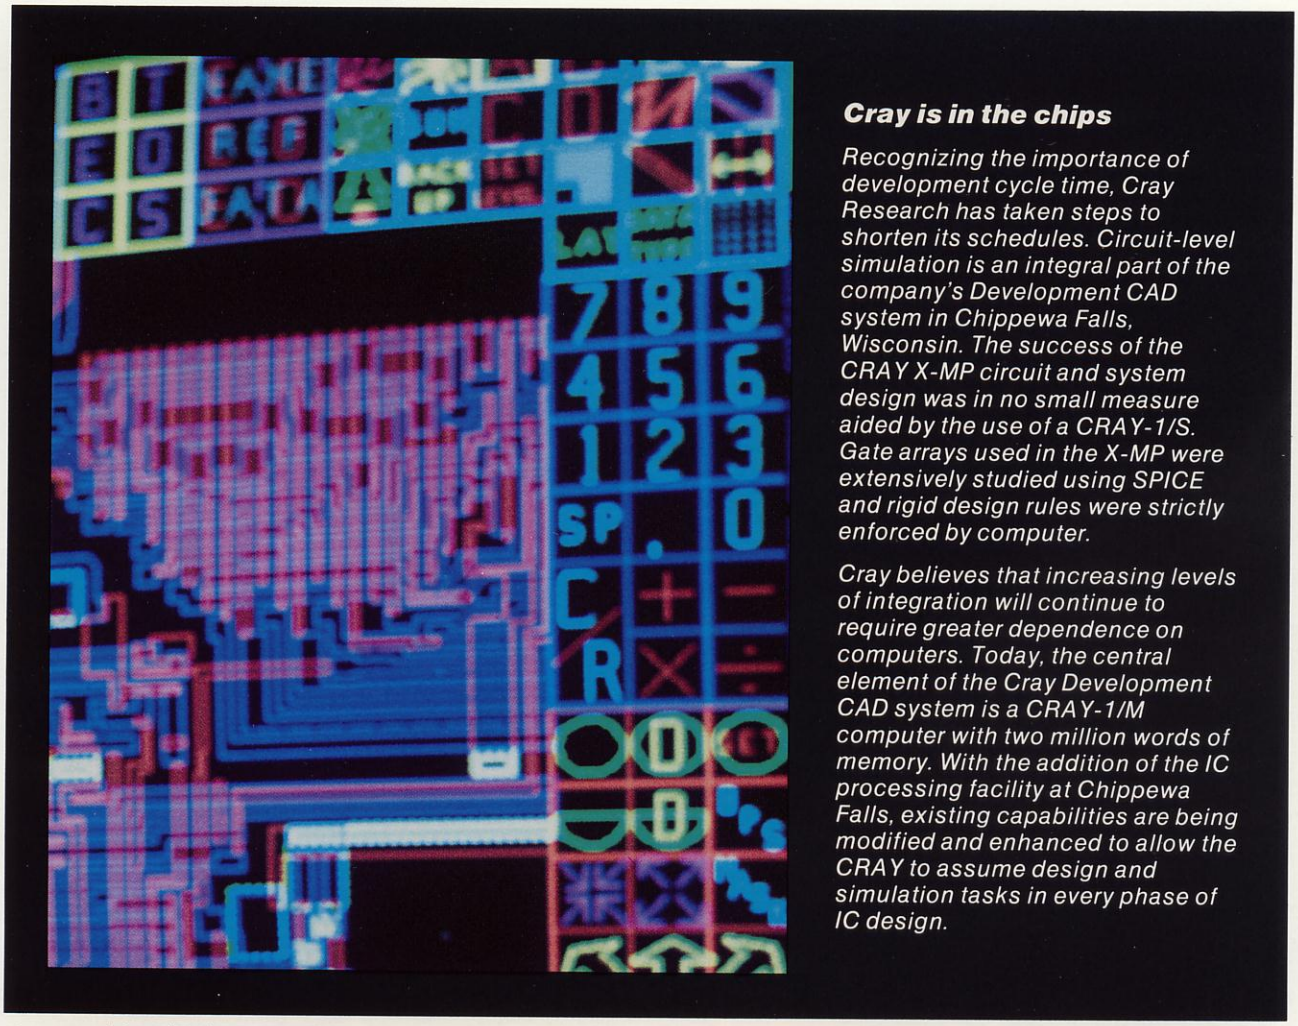

As part of the I/O function, this raster display in the Cray CAD system allows the designer to view a portion of an integrated circuit.

About the author

Kelly Wild received his B.S. in Physics from the University of Colorado and conducted research in high energy physics. Before joining Cray Research earlier this year, he worked at United Information Services, Inc. for four years, where he supported electrical engineering and finite element method codes on the CRAY.

Acknowledgements

Thanks are extended to Fairchild Camera and Instrument Corp. and NCA Corp. for the use of photos on pgs 4 & 5 and 2 & 3, respectively. (Photo credit pgs 2 & 3: Judson Allen.)

References

Boraiko, A. A., "The Chip," *National Geographic*, Vol. 162, No. 4, October 1982, pp 421-441.

Breuer, M. A., Friedman, A. D., Iosupovicz, A., "A Survey of the State of the Art of Design Automation," *Computer*, October 1981, pp 58-75.

"Chip Wars: The Japanese Threat," *Business Week*, May 23, 1983, pp 80-96.

Electronics — Special Commemorative Issue, Vol. 53, No. 9, April 17, 1980.

Garcia, S., Sriram, K. S., "A survey of IC CAD Tools for Design, Layout, and Testing," *VLSI DESIGN*, September/October 1982, pp 68-69.

Hewlett-Packard Journal, Vol. 32, No. 6, June 1981.

Mead, C., Conway, L., *Introduction to VLSI Systems*, Addison-Wesley Publishing Company, Reading, MA, 1980.

Scientific American, Vol. 237, No. 3, September 1977.

Notes

1. M. Brooksby, P. L. Castro, F. L. Hanson, "Benefits of Quick-Turnaround Integrated Circuit Processing," *Hewlett-Packard Journal*, Vol. 32, No. 6, June 1981.
2. L. W. Nagel, "Advice for Circuit Simulation," *IEEE International Symposium on Circuits and Systems*, Houston, Texas, April 28, 1980.
3. M. P. Lepselter, et al., "A Systems Approach to 1-μm NMOS," *Proceedings of the IEEE*, Vol. 71, No. 5, May 1983, pp 640-641.

A view of circuit design principles and procedures

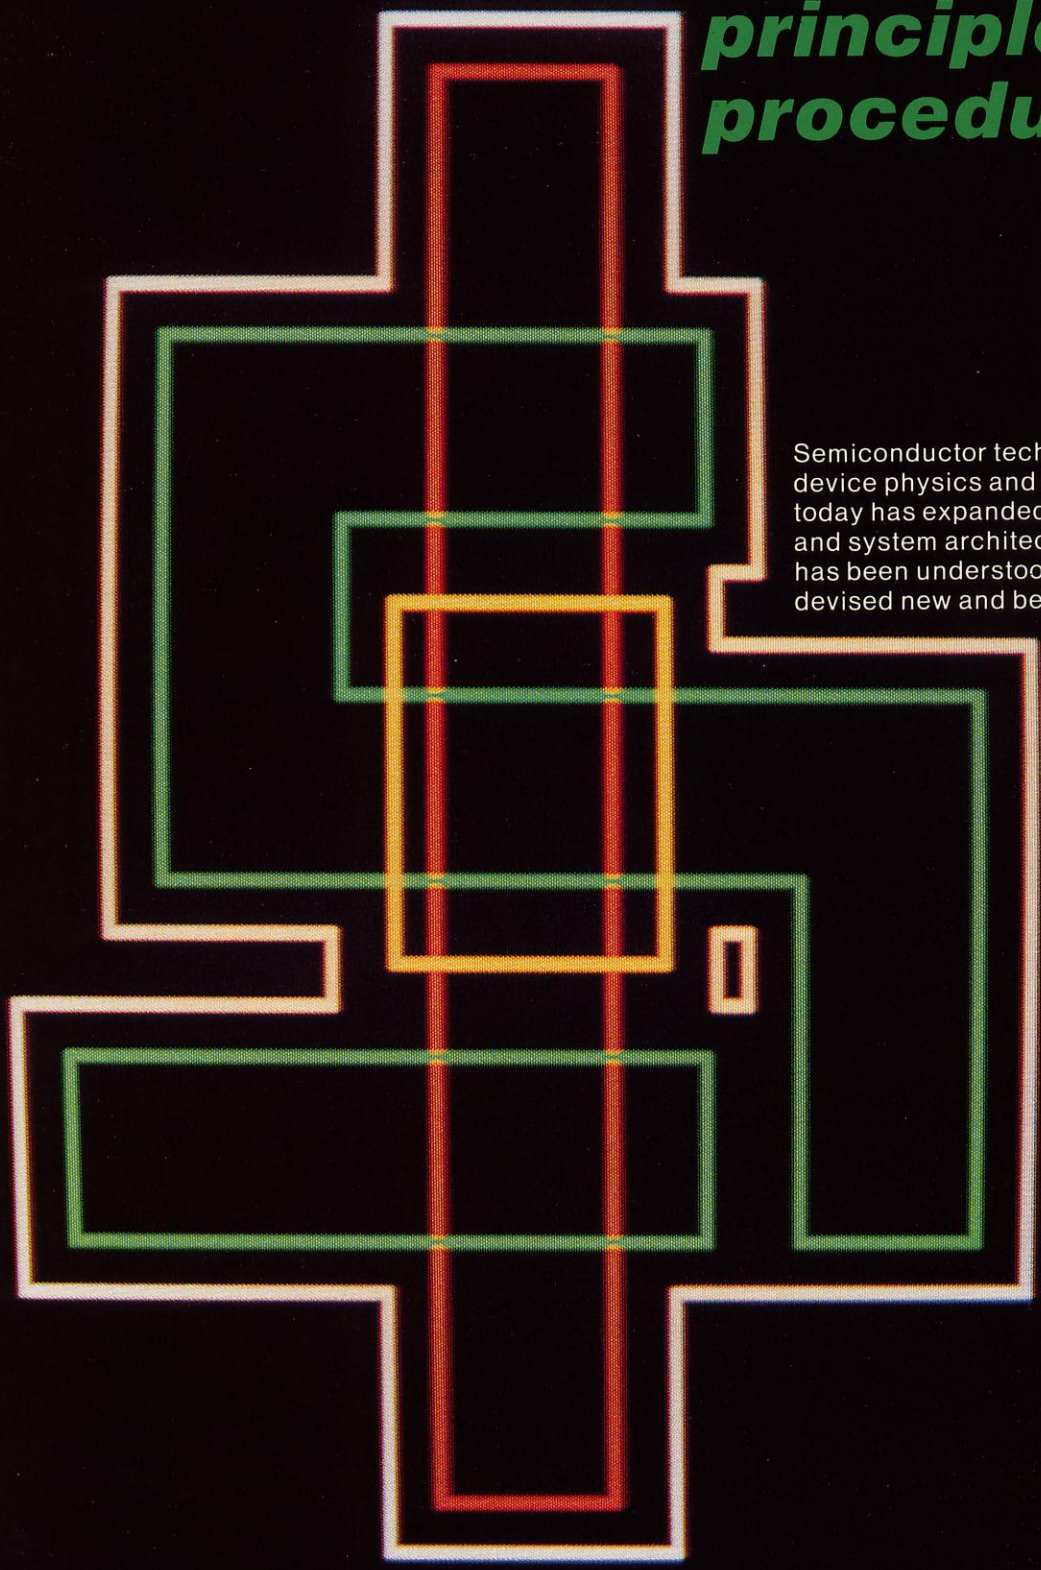

Semiconductor technology — a blend of device physics and circuit design that today has expanded to include logic design and system architecture. Since the physics has been understood, engineers have devised new and better ways of harnessing semiconductor capabilities at a feverish pace. At the same time, sound methodologies to generate and verify circuit designs have been developed. In this article, we take a look at the basic principles behind one semiconductor technology followed by a look at the general design steps from inception to fabrication of a new circuit.

Design rule checking graphics output

Basic principles of n-MOS technology

For our design example, we consider n-MOS semiconductor technology in which integrated systems contain three levels of conducting materials separated by intervening layers of insulating material. Proceeding from top to bottom, the layers are termed metal, polysilicon, and diffusion. Patterns for paths on the three levels and the locations of contact cuts through the insulating material to connect certain points between levels are transferred during the fabrication process from masks similar to photographic negatives. (See "Fundamental Integrated Circuit Fabrication," CRAY CHANNELS Vol. 5 No. 2.)

In this technology, wherever a path on the polysilicon level crosses a path on the diffusion level, a transistor is created. Such a transistor has the characteristics of a simple switch, with a voltage on the polysilicon-level path controlling the flow of current in the diffusion-level path. These transistors interconnected by pattern paths on the three levels form the basic building blocks of integrated circuits.

The technology employs a positive power supply voltage and conducts when a positive signal is applied to the gate. The three elements of an n-MOS transistor are the source, the gate, and the drain. The basic operation performed by the transistor is to use the charge on its gate to control the current between the source and the drain. Transistor operation is controlled by specifying the dimensions of the polysilicon and diffusion materials used in the layout of the transistor. When a positive potential — approximately 2 volts — is applied to the gate, the transistor conducts; when the voltage is reduced to 0, conduction ceases.

Logic gates can be composed to compute arithmetic and logic expressions in which values are represented by voltages. A value 1 (or logical true) is represented by a high voltage, and a value 0 (or logical false) by a low voltage. The inverter is the simplest gate to consider; it emulates the logical function 'not'. When the input to the 'not' function is false (or 0), the output is true (or 1); similarly the input of a true generates an output of false.

An n-MOS representation of an inverter consists of a single transistor connected in series with a resistor. The supply voltage (a positive potential of about 5 volts) is applied to the resistor. The input signal is applied directly to the gate of the transistor, and the output is sensed at the point where the transistor and resistor are joined.

When the input to the inverter is in the low state, the transistor does not conduct. Thus, no current flows from the supply through the resistor and the transistor to the ground point. As a result, there is no voltage drop across the resistor; the output is connected to the supply voltage. When the input is low, the output is high.

If a high signal is now applied to the gate, the transistor conducts. With a large current flowing through the circuit, there is a substantial voltage drop across the resistor. The output of the gate is thus effectively connected through the low resistance of the transistor to the ground; it is in the low-voltage state. This is precisely the behavior required of the 'not' function.

The engineer designs the circuit logic and layout so that these signals are transmitted in a meaningful way in space and time. The job is no mean task. A good circuit design that has the prescribed behavior characteristics is dependent on the cleverness of both man and machine. The process begins with the algorithm describing the basic set of functions to be performed, and evolves into the set of masks used to fabricate the device. Major steps of the circuit design process are illustrated in Figure 1 on page 10.

IC design steps

The first phase of designing a circuit is probably the most engineering and computer-interactive intensive, requiring engineers from a variety of disciplines to work together. Designers consider the functional specifications and produce a logic design. Their basic concerns in creating a new circuit are whether the design will work, perform to specifications, and can be manufactured and tested at reasonable cost. Speed, power, space, and design time can all be traded off against each other, but a wrong decision can lead to a great deal of backtracking, because all portions of a design may be interdependent.

The next major design phase — layout — is more compute-intensive and is the starting point of the CAD environment. All physical attributes of the circuit such as the location of RAM or read in/read out address registers are defined. The most common way that ICs are represented is in geometric form, with each shape associated with some layer on the final chip. Various portions of the allowable delays and power consumption are allocated to parts of the total circuit. This step requires a significant amount of experience — optimizing a design is very complicated.

The tedious and error-prone task of entering the circuit or logic level description into the CAD system has been greatly simplified over the last few years. The data base contains the circuit description generated during the layout phase, and with input provided by the engineer, the circuit's characteristics are compiled as a netlist to be read by different simulators. The netlist is the nodal element list of the network.

At this point we move into the very computationally intensive area of design verification, where CRAY systems find their greatest application. In the data entry steps described above, specialized CAD and/or front-end computers often handle the computing workloads. However, the computers are quickly running out of capacity in executing design verification functions.

Design verification is critical because enormous savings can be realized if design errors are eliminated

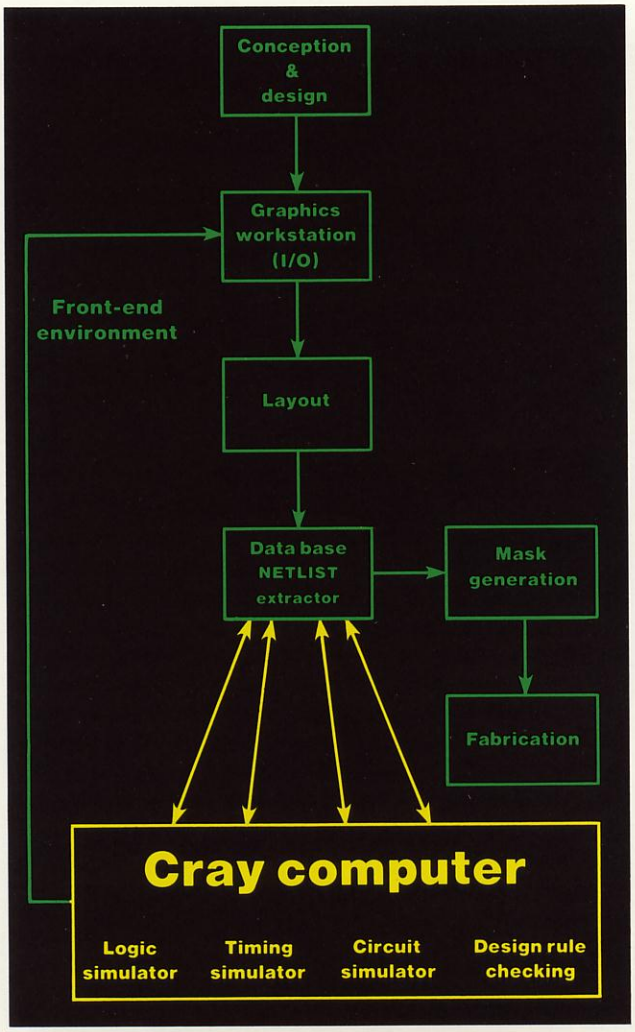

Figure 1. Integrated circuit design flow

in the first pass. Beyond the tremendous cost savings, possible reductions in the time required to introduce a new product can be realized in this phase. CRAY computers can help streamline verification procedures by executing complex simulations. For example, circuit simulation can be run in minutes compared to the hours it would take on conventional large systems. By moving the simulation to the CRAY, more in-depth analysis becomes practical and the work can be done in far less time.

Types of simulators

One of the most CPU-demanding simulators is the circuit simulator. It simulates events that would take place in all the components of a circuit if voltage were applied. It allows the designer to look at speeds, powers, noise margins, best and worst cases, and the effects of process variations by taking a global look at the device. Circuit simulators use modified nodal analysis techniques and numerically solve a set of ordinary differential equations. It generates data relating to the current and voltage behavior at user requested nodes and devices.

The SPICE circuit simulation program developed by the University of California at Berkeley has become a workhorse for the IC industry. A version of SPICE called CSPICE has been enhanced to run on the CRAY. On the CRAY-1/S, some simulations run with CSPICE execute more than 100 times faster than the same dataset simulated on a VAX 11/780 (see page 14 for additional information).

Another major simulator is the logic simulator, in which automatic test-pattern generation and fault simulation programs are run. As circuits have become dense and very sequential, test-pattern generation has become very computation-bound. The logic simulator deals with a network of Boolean gates, true and false gates, 'flip flops', etc. By inputting an initial state of 0's or 1's, the program will flow through the various branches of the network to output final states of 0's, 1's, or intermediate states (not 0, not 1). Voltages are accurately defined, but for MOS devices, timings of the circuit switching may not be.

The timing simulator's primary function is to monitor signal arrivals at different locations. A network of transistors is modeled as resistors; the behavior of the network is accurately described in time, whereas the exact voltage values may not be.

Design rule checking (DRC) is a process that allows the designer to verify the topology of his circuit design. For a circuit to be reliable it must be designed with the prescribed limits of the manufacturing process in mind. Each IC fabrication facility has a set of specifications of minimum spacings and widths an IC process can reliably fabricate. These design rules must be followed, but mistakes may be made because of the complexity of the task. DRC checks for compliance with the design process rules. The user inputs his mask layout from his CAD system, then the simulator checks for various physical layout flaws. Such flaws may include situations like irregular device spacing or an illegal orientation of cells. After DRC has been completed, the circuit design will be readied for fabrication.

Conclusion

By applying the basic principles of device physics, integrated systems are orchestrated to perform complex functions. As we saw with n-MOS technology, different levels of conducting materials are patterned, connected and insulated to produce the device elements that are the functional building blocks of integrated circuits.

Engineers with expertise from different disciplines work together to develop designs that function to specifications. But that is only the first in a series of design synthesis and verification steps that tax both human intellect and compute power. CRAY systems are playing an increasingly important role in the most compute-intensive areas of circuit design. These computers enable engineers to be more efficient and accurate in developing today's integrated circuits. □

Acknowledgement

The photo of DRC output was provided by NCA Corp.

CORPORATE REGISTER

French customers order X-MP and CRAY-1/S computers

This past summer, Cray announced that it received orders from two French organizations. Compagnie Internationale de Services en Informatique (CISI) placed an order for a CRAY X-MP/22 computer system. CISI, the leading computer service bureau in France, has been a heavy user of a CRAY-1/S system installed since 1981. Cray received export approval for the system that will be installed during the first half of 1984.

The second order was from Commissariat a l'Energie Atomique (CEA), a French government agency that will purchase a CRAY-1 S/2200 computer. The system will be the agency's second CRAY and will be used for scientific research.

ECMWF orders CRAY X-MP

Cray Research, Inc. recently announced that it was awarded a contract by the European Centre for Medium Range Weather Forecasts (ECMWF) for the installation of a CRAY X-MP/22 computer system. The two-million word CRAY X-MP with a Solid-state Storage Device will be installed at ECMWF in Shinfield Park, Reading, U.K., during the fourth quarter of 1983.

ECMWF is a cooperative venture of 17 European Member States established in 1975 to provide medium range (up to 10-day) weather forecasts. The tasks of the Centre are to issue daily medium range weather forecasts, to conduct research on improving forecasting techniques, to make available large-scale computing facilities for the Member States, and to provide advanced training for meteorologists.

John A. Rollwagen, Cray Research chairman, commented, "ECMWF is a world leader in medium range weather forecasting. We are pleased that ECMWF has chosen Cray Research for the second time as the supplier of its major computer system." He noted that ECMWF was Cray's first international customer, having installed a CRAY-1A in 1977.

The Director of ECMWF, Dr. Lennart Bengtsson, says that operational weather forecasts for a week or more ahead have been carried out for the Member States each day since 1980. Many of the six-day forecasts now issued by the Centre are of comparable quality to the two- to three-day forecasts available in Europe in 1970 when the establishment of the Centre first was proposed.

The Centre reports it is confident that further improvements will be made because of the greater 'number-crunching' power of the CRAY X-MP. This will allow use of a more detailed atmospheric model for forecasting with a consequent increase in the accuracy of the predictions.

Chevron to lease CRAY X-MP supercomputer

In June, Cray announced that it will install a CRAY X-MP/24 computer

CORPORATE REGISTER

system at the Chevron Oil Field Research Company in La Habra, California, during the first quarter of 1984.

Chevron plans to use the system for research and development in exploration data processing and petroleum reservoir engineering. The oil company has had a CRAY-1 S/2300 installed since early 1982.

Lockheed orders CRAY computer

Lockheed Missiles & Space Co., based in Sunnyvale, California and part of Lockheed Corporation, placed an order for a CRAY-1 S/1000 Computer System in June. The system, which will be leased, is scheduled for installation this September.

The Lockheed operation plans to use the CRAY computer in aerospace applications such as structural analysis, computational aerodynamics, fluid dynamics, simulation, and electronic engineering. R.A. Fuhrman, president of Lockheed Corporation's Missiles and Electronics Group, said that the use of the CRAY computer system "is representative of efforts to provide our engineers the most advanced and cost-effective tools available to achieve technical excellence."

John Rollwagen, chairman of Cray Research, noted that the order "represents the first time an aerospace company in the U.S. will use one of our computers entirely for its own internal design and engineering needs."

Second annual internal Cray Technical Symposium held

In June, Cray Research held its second internal symposium in Minnesuing, Wisconsin. About 50 people from different areas within the company were invited to participate in the conference. Individuals from the company's hardware and software development groups presented information on a broad range of ongoing projects. In addition, a number of analysts and engineers located at customer sites around the world shared insights on practical matters of the CRAY environment.

Holding a molecular model of gallium arsenide, Dr. Freeman explained how supercomputers are needed to handle increasingly complex calculations in the sciences.

Some of the topics covered were: networking software for CRAY systems, consolidation of processor and memory cells in VLSI mainframes, gallium arsenide research, and software advances in the use of the SSD. A panel discussion was held to share information about multi-tasking.

Special presentations were made by several guest speakers, including John Rollwagen, Cray Chairman, Margaret Loftus, Vice President — Software Development, and Les Davis, Executive Vice President. Dr. Arthur Freeman, a noted physics expert from Northwestern University and advocate of large-scale computing, spoke to the group about the importance of concentrated supercomputer development research. Dr. Freeman is a member of a special National Science Foundation committee that has made an appeal to Congress to make funding available for supercomputer proliferation.

Cray releases new software

The company recently announced that new versions of the CRAY-1 Operation System (COS) Version 1.12 and the FORTRAN Compiler (CFT) Version 1.11 with its accompanying libraries were released. Significant features included in these releases follow.

COS 1.12 features include:

- Software support for the CRAY X-MP mainframes, the Solid-state Storage Device, and the 80-Mbyte disk pack on the I/O Subsystem
- Enhanced COS job scheduler task
- Station Call Processor now dynamically allocates and releases station buffers
- FETCH (non-permanent ACQUIRE), which may be used either as a control statement or as a call to a routine in $SYSLIB
- Performance improvements to UPDATE, and addition of the YANK and UNYANK directives
- TMS support via the IBM MVS station
- Permanent dataset privacy, allowing users to specifically grant or deny mass storage dataset access to other users for a variety of access modes
- CFT callable SKIPR, SKIPF, and SKIPD routines
- Support for formatted conversion of IBM-compatible on-line tapes
- COS now permits the user to accept or skip bad data read from tape or disk
- COS security, which enables sites to define for each user privileged functions to be granted or denied
- System-controlled and user-controlled memory management for user jobs' field length
- Single-point and multi-point NSC HYPERchannel® support via the I/O Subsystem
- Ability to reconfigure mass storage devices during and after system startup
- Enhanced error reporting and recovery during startup
- Ability to reconfigure front-end channels connected to the CRAY during system startup

Significant features in CFT 1.11 include:

- 46-bit integer multiply and divide
- MAXBLOCK and OPT parameters added to CFT control statement: MAXBLOCK permits the user to set maximum block size for optimization and vectorization; OPT allows the user to select partial, full, or no optimization of certain IF statements and to enable or disable 46-bit integer multiply and divide
- New compiler directives
- Vectorization of expressions using constant increment integers
- Automatic conversion of appropriate IF statements to MAX or MIN functions, enabling vectorization
- CFT generation of new calling sequence code for subroutine calls following conversion to new calling sequence
- $FTLIB split into four specialized libraries

Customers interested in obtaining these products should follow the same ordering procedures as for all other standard Cray products.

Cray Research announces VM Station availability

Users of VM/SP computer systems now have access to CRAY-1 and CRAY X-MP Computer Systems while retaining the capabilities of the CMS environment. The VM Station Software Service enables IBM-compatible systems running under control of the Virtual Machine/System Product (VM/SP) and Conversational Monitor System (CMS) to be linked with a CRAY computer system running under control of the CRAY Operating System (COS) version 1.11 or later. The IBM-compatible hardware is connected to the mainframe through either a CRAY front-end interface or Network Systems Corporation's HYPERchannel® hardware. The station works with either configuration.

The first release of the VM Station Software Service provides facilities for transferring COS job input and output datasets between the CRAY and VM/SP systems. Further, a VM/SP user will have access to COS job status information and COS interactive facilities. A summary of these features is provided below.

- Dataset Transfer — A CMS-compatible command allows the submission of user job files to the CRAY system or the saving of CMS datasets as permanent CRAY datasets. When running on the CRAY system, a job may stage datasets to and from the VM system using the COS DISPOSE and ACQUIRE control statements. Datasets may be staged from the CRAY system to magnetic tape, the system printer, or a user's virtual reader through either the station virtual machine print or punch. Datasets may also be staged to the CRAY system from magnetic tape or CMS minidisks. The station supports only Enhanced Disk Format (EDF) disks, IBM models 3330, 3350, or 3380 or equivalent. The COS reference to the dataset to be transferred may specify character mode or transparent mode. For character mode, the dataset is translated between EBCDIC and ASCII and between CMS disk format and COS blocked format. For transparent mode, no translation is performed.
- Job status — A user can inquire about the status of jobs running on the CRAY system. Also, one can control his or her own job but not those originating from other users.
- Interactive — A COS interactive session can be initiated for efficient program development.

The VM Station Software Service is available for installation at customer sites beginning September 15, 1983. Those interested in additional information about the station should contact the nearest Cray field office.

APPLICATIONS IN DEPTH

Circuit simulation with CSPICE

Simulation Program with Integrated Circuits Emphasis (SPICE) is a circuit simulator that was originally developed by the University of California, Berkeley. The program is used by an overwhelming majority of electronics companies and has been converted to run on many different computer systems.

Two versions of SPICE, CSPICE and SPICE2, operate on CRAY computer systems and can be obtained from Cray Research. SPICE2 has been available for some time while CSPICE has only recently become available for commercial use.

CSPICE is a highly optimized version of SPICE and has execution speeds of over 100 times faster on CRAY machines than SPICE running on popular super minicomputers. The reason for the program's high performance is that it takes full advantage of the vector and pipelining capabilities of CRAY hardware. CSPICE operates on CRAY-1/M, CRAY-1/S, and CRAY X-MP systems running under COS 1.11 or later.

The software allows an electronic circuit designer to submit an electronic circuit description to a CRAY for processing. It requires that the netlist be generated on a front-end computer system, such as a VAX 11/780. The netlist is preprocessed on the front-end, then presented to the CRAY for solution. Results are brought back to the front-end for plotting or display.

A hierarchical architecture that distributes the functions of the design cycle among processors in a cost-effective manner is important for a good circuit simulator. This is referred to as a loosely-coupled circuit simulator. At the present time, CSPICE is the first and only loosely-coupled circuit simulator. CSPICE makes use of a super-mini and/or an intelligent workstation in order to coexist with the designer's environment. CSPICE allows the user to control the limits of the circuit simulation and interrogate the results of the simulation.

Those interested in additional information about the program should contact: Derek Robb, Cray Research, Inc., 1440 Northland Drive, Mendota Heights, MN 55120; telephone: (612) 452-5560

TEGAS-5 converted for operation on Cray systems

Conversion of TEGAS-5 for the CRAY was recently completed by Cray Research. The TEGAS program is a very powerful logic and timing simulation system that analyzes and verifies digital logic networks. The program provides a network design language, logic and design verification, worst-case timing analysis, testability analysis, automatic test generation, and fault simulation capabilities.

Logic and design verification and worst-case timing analysis are used to study the logical behavior and timing characteristics of digital networks. The modeling of the signal timing propagation is made more complex as the designer progresses through these three design stages. Completed designs can be manufactured without timing problems.

Logic verification is an important capability of TEGAS. Designers can specify complex waveforms in an arbitrary time frame to be applied to the inputs (and/or internal nodes) at simulation time to verify that the logical implementation of a network is functionally correct. Logic verification typically uses unit delay network models to minimize use of computer resources.

To assist in design verification, switching delay information is included in the network under simulation to identify timing problems within the logically correct design. Delay specifications can be included at the primitive-element output pins to reflect loading delays. Separate delay values can be included for rising and falling signal transitions.

Worst-case timing analysis is implemented by modeling minimum and maximum rise and fall delays on primitive elements and their output pins to verify that the design works within the user-specified range of delay possibilities. Worst-case analysis simulation uses special logic states to represent the regions between the minimum and maximum delay values. These are treated as 'unknowns' when they appear on sensitive inputs.

Testability analysis performs analysis of the network with respect to the ease or difficulty of testing at each net for stuck-at-one or stuck-at-zero faults. The algorithms for this program process both combinational and sequential logic.

Automatic test pattern generation is based on network topology and information provided by testability analysis. The generation algorithm is a 14-valued implementation of the path-sensitization D-algorithm. Approved test data can be automatically interfaced to commercial testers.

The TEGAS-5 software automatically creates a fault map of stuck-at-one, stuck-at-zero faults on a set of nets specified by the users. A faulted network model for each fault in the map is simulated. The program compares simulation output of the good network to that of each faulted network and finds discrepancies at network test points to accomplish fault detection.

For more information about TEGAS-5, contact: Comsat General Intregrated Systems, Inc. at 7801 N. Lamar Boulevard, Austin, TX 78752; telephone (512) 451-7938.

EE codes converted to run on CRAY systems

Several EE codes are converted and available for use on CRAY systems. These programs work on a variety of problems involved in integrated circuit design development. For instance, a 2-D modeling program for designing n-channel enhancement-type MOS transistors is available. The code named NEMOS offers the following computational options: drain current for a prescribed set of voltages or for prescribed voltage steps on the gate electrode and drain junction, threshold voltage for prescribed drain and substrate voltages, and drain to source punch-through voltage.

On the other hand, SUPREM (Stanford University PRocess Engineering Models), simulates most typical IC fabrication steps. The package is designed so that these steps can be simulated either individually or sequentially, just as they would occur during the actual fabrication process. The output which is available at the end of each step, consists of the 1-D profiles of all dopants present in the silicon and silicon-dioxide materials. The current version of SUPREM available on the CRAY is SUPREM 2. SUPREM 3 will soon be available from Cray.

Another simulator, RSIM, is a resistive simulator that models transistors as resistors (the logic states of a transistor's terminal nodes determine the value of the resistor). Currently, this system is undergoing conversion for operation on CRAY systems. RSIM was developed in response to a need to alleviate some of the difficulties in abstracting accurate gate-level representations from the electrical networks of bi-directional MOSFETS used in MOS integrated circuits. Networks of transistors and electrical nodes form an R-C tree (R from the transistors, C from the interconnect and gate capacitance); the network's behavior under different inputs is calculated by a selective-trace (event-driven) algorithm.

For additional information about these and other electrical engineering codes converted for use on the CRAY, please contact Kelly Wild, Cray Research, Inc., 1440 Northland Drive, Mendota Heights, MN 55120; telephone: (612) 452-6650.

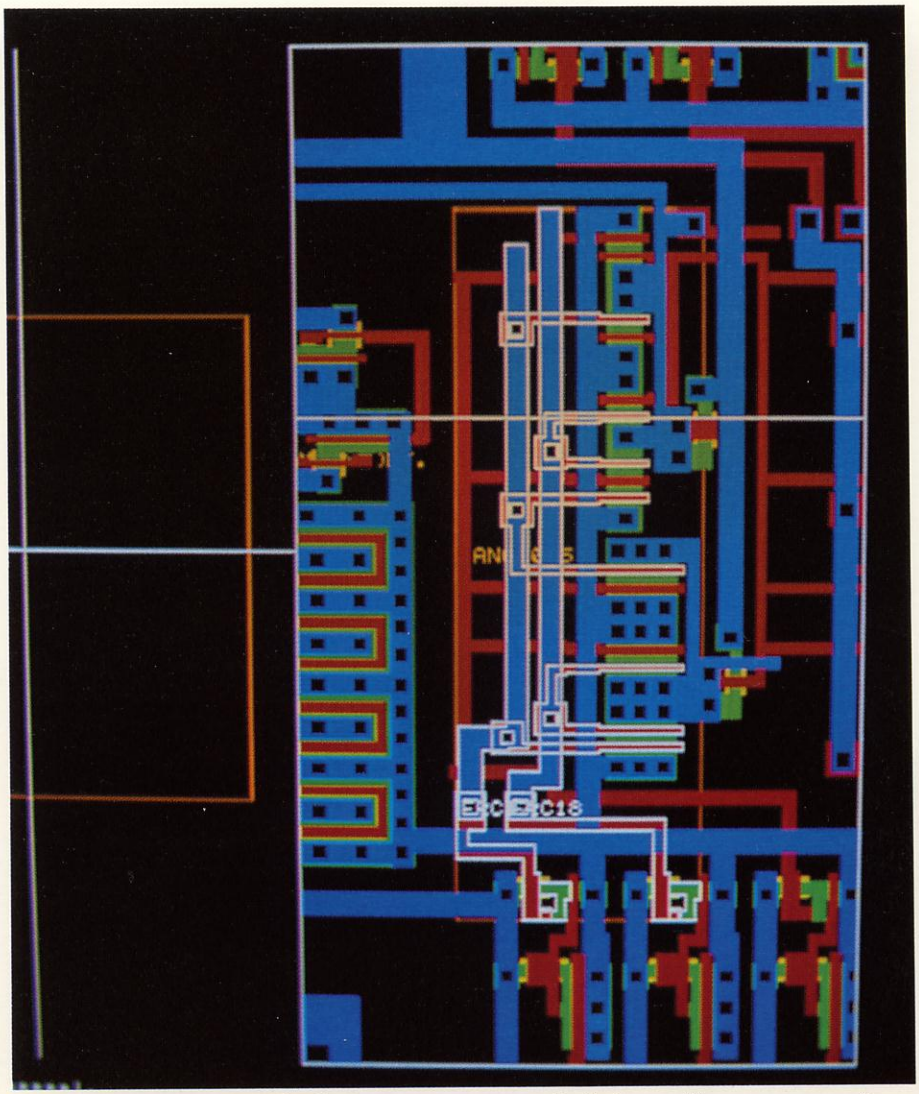

Network Consistency Check compares a network design with a corresponding layout to verify consistency.

NCA/DVS™ converted for use on the CRAY

For some time, NCA's Design Rule Check™ (DRC) program has been available for operation on CRAY systems. DRC is a CAD system for electronic circuits that identifies design rule violations in ICs. The user inputs a circuit description and test request file to specify the design rules to check and instructions used to generate other useful mask layers. The program verifies whether the design rules are met on the input layers and on the generated layers. DRC is generally accepted as an industry standard, and is used by virtually every major semiconductor manufacturer.

NCA, in cooperation with United Information Services recently converted the entire NCA/Design Verification System™ (NCA/DVS) of which DRC is a part, to run on CRAY systems. NCA/DVS is an integrated circuit design verification system developed to reduce design cycle times and ensure design reliability. Currently, the system is operating on the CRAY-1/S system located at UIS.

The system can verify that a circuit design obeys basic electrical and technology rules, ensure that the graphical layout and original logic design are exactly the same from the functional level down to the individual transistor level, verify design rules, and prepare the verified design for the mask-making equipment. The packages that comprise NCA/DVS are Electrical Rules Check™ (ERC™), Master Design Language™ (MDL™), Network Consistency Check™ (NCC™), Electrical Parameters Check™ (EPC™), Mask Data Preparation™ (MDP™), and DRC.

Those interested in obtaining additional information about NCA/DVS should contact: Dennis Elkins, NCA Corporation, 388 Oakmead Parkway, Sunnyvale, CA 94086; telephone: (408) 245-7990 or United Information Services, Inc., P.O. Box 8551, Kansas City, MO 64114; telephone: (913) 341-9161.

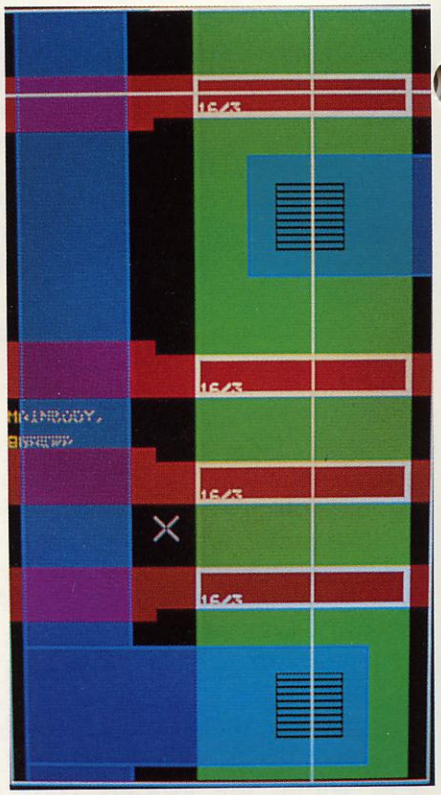

Within NCA/DVS, the Electrical Parameters Check function simulates electrical parameters that affect load and speed by checking factors such as transistor gate areas and gate output capacitances.

USER NEWS

The French connection

Sharing a CRAY across the miles has allowed two French firms to grow into supercomputing one step at a time. After establishing a consortium in 1980, the two organizations imported the first CRAY system to France in 1981. Strange to say, the arrangement has been so successful that the consortium was recently dissolved. The full story is that the two groups' supercomputing needs have grown so rapidly that each will soon have its own CRAY on-site.

Groupement d'Exploitation de Techniques Informatiques Avancees (GETIA) was formed by Electricite de France (EDF) and Compagnie Internationale de Services en Informatiques (CISI). Their CRAY acquisition had an unusual twist — EDF and CISI are located about ten miles apart outside of Paris. This necessitated the establishment of a complex communications network to enable sharing of the CRAY. Communication with the CRAY for remote users has therefore been a key factor in the success of GETIA's arrangement.

The network installed two years ago (illustrated in Figure 1) has been a workhorse for CISI users, giving them the needed accessibility to supercomputer power. It is configured with Network Systems Corporation (NSC) adapters for the CRAY and front-ends, and with NSC link adapters connecting two remote HYPERchannels® through two dedicated high-speed communication lines, called Modulation Impulsions Codees (MIC) lines. The MIC lines, offered between all major cities in France by the French telecommunications company were recently upgraded from 0.5 M bits/sec. to 1 M bits/sec.

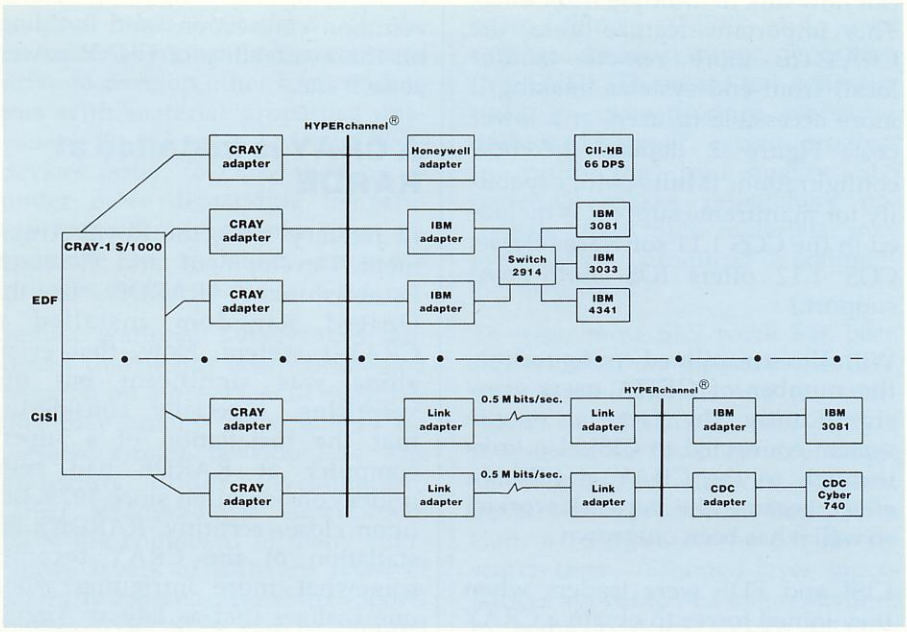

Figure 1. GETIA remote front-end configuration PRIOR to multi-point

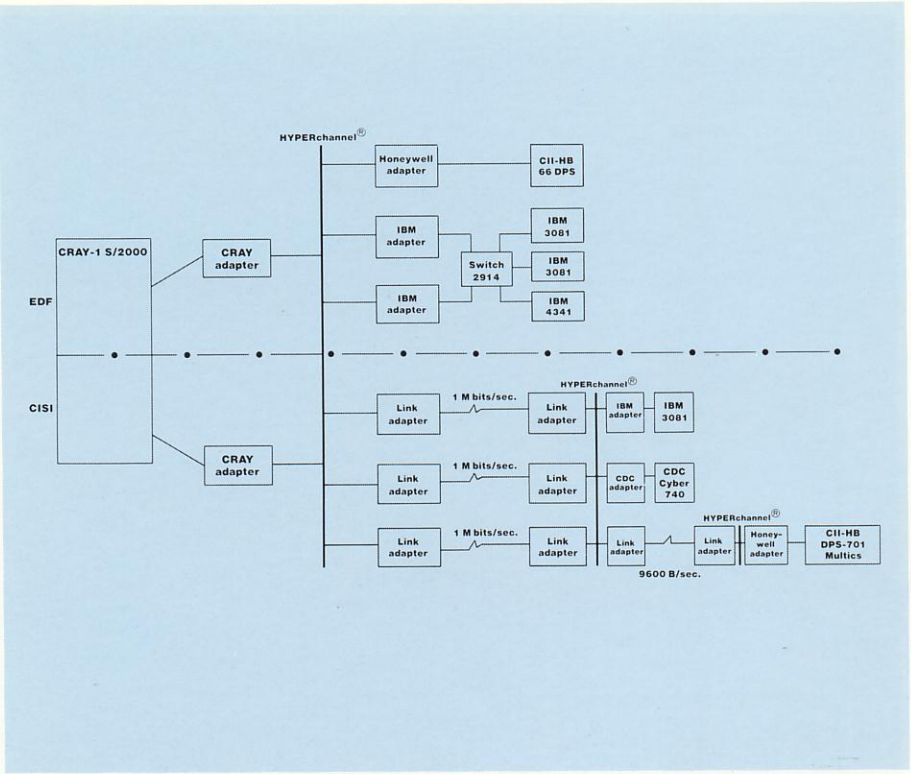

Figure 2. GETIA multi-point configuration

In early 1983, a major enhancement was incorporated into the system. In the past, each CRAY channel connected with an adapter had to be dedicated to a single front-end system. Single CRAY channel ports can now link to multiple front-ends. This important feature links the CRAY to more remote (and/or local) front-end systems, making it more accessible to users at a lower cost. Figure 2 depicts the new configuration. (Multi-point capability for mainframe support is included in the COS 1.11 software release; COS 1.12 offers IOS multi-point support.)

With the streamlined configuration, the number of GETIA users grew significantly. In fact, a remote system connected to CISI also links through to the CRAY at EDF. In effect, because the network worked so well it has been outgrown.

CISI and EDF were leaders when they joined forces to obtain a CRAY in 1980. Their high standards for computing power and capabilities helped evolve multi-point and remote front-end access on the CRAY to where they are today. Now, in 1983 each organization moves into the next phase of supercomputing — severing their common connections and building on the availability of CRAY power at each site.

A CRAY is installed at RARDE

In January 1983, the Royal Armament Development and Research Establishment (RARDE) in the United Kingdom installed a CRAY-1 system. Now, that event alone was significant but not surprising, especially considering that the installation of a supercomputer at RARDE had been under consideration since 1979. But upon closer scrutiny, RARDE's installation of the CRAY becomes somewhat more intriguing when one realizes that as late as August 1982 major decisions regarding the installation had yet to be made. For one thing, funding for a supercomputer purchase was not expected until after April 1984. Furthermore, there was no place to install the computer. An existing computer center was to house a new system, but it required complete refurbishment and preparation for installation. Also, a completely new plant room had to be constructed. Finally, a decision on a front-end system, part of any supercomputer configuration, and its funding had yet to be made. Clearly, things did not look too promising for a January 1983 install date.

Then a breakthrough was achieved late in the summer of 1982 and funds became available to RARDE. However, funding hinged on some demanding requirements: the computer had to be purchased, installed, and running by the end of the first quarter 1983 — a difficult task for the research establishment that had no building in which to put the system, no front-end system, no CRAY, and only a few months to organize the task.

All that not withstanding, RARDE proceeded to call Chris Windridge, Account Manager at Cray, to discuss stepping up its timetable for installation by 15 months. (Right about now, one begins to understand how the British earned their reputation for optimistic determination.) At the end of the conversation, Windridge and Cray UK were left with three questions:

(1) Could a CRAY be installed by March 1983?
(2) How could the computer center be made ready for the CRAY?
(3) How could the purchase and installation of a front-end computer be organized for RARDE?

All three questions were answered by Cray UK's offering to take responsibility for completing the requirements within the timetable. The decision to undertake the project was not made lightly by Cray Research and the UK subsidiary,

out the company was confident that together with RARDE, they had the skills and resources to complete the seemingly impossible task.

Quickly, many RARDE and Cray personnel rolled up their sleeves to get to work. Reservation of the CRAY system aside, efforts were concentrated on the reconstruction of the building. Essentially, Cray UK acted as a third party contractor for refurbishing and re-equipping the computer center and the plant and services. Cray UK also acted as a prime contractor for a VAX 11/780, which was chosen as the front-end by RARDE. The arrangement worked out well for RARDE because there simply wasn't time to place multiple contracts; for the entire installation RARDE was able to work with a single vendor — Cray Research. The RARDE experts, of course, had to be satisfied not only about the specifications and choice of subcontractors, but also that keen pricing was being pursued, despite the tight delivery schedules.

Late in January, the CRAY was installed in the new center. By the middle of February the CRAY was accepted, followed by the VAX acceptance late in the same month. In April of this year the dedication of the new RARDE computer center was made by the Minister for Defence Procurement, Mr. Geoffrey Pattie. Pretty nice handiwork. There is little doubt that RARDE regards itself as having achieved singular value for its money by bringing a major capital investment into effective use in a very short time and offsetting other costs that could have been incurred.

The installation at RARDE is indicative of the kind of working relationship Cray Research develops with its customers. Over time, the company and its people have committed time and energy to fill the special needs of many of our customers. Without the strength and variety of expertise that Cray personnel have, it would have been impossible to work with RARDE on this project. Likewise, the soundness and flexibility of the Cray organization itself allows such efforts to be undertaken. While Cray cannot guarantee that it will always be ready to respond in every unusual situation, the company and its people attempt to do as much as possible for each customer to make each CRAY installation successful.

National labs further IC technology

CRAY users are involved in circuit design and fabrication research on many different levels. Many organizations including computer manufacturers, universities, and our national laboratories, are committing major resources to developing new semiconductor technologies. Recently, two Cray customers announced progress in different areas of the semiconductor realm. The following articles offer some insight into the scope of microelectronics research.

Sandia research holds possibilities for new semiconductor materials

Despite silicon's amazing success as a semiconductor, for years scientists and technologists have 'engineered around' the limitations of the material. There has been a renewed drive to develop other semiconductors with material properties unavailable to silicon in order to make devices faster, to make them work under more demanding temperature conditions, to make them responsive to a wider range of voltages, and so on.

Sandia National Laboratories disclosed that it has been conducting research on a new class of materials that may become important to integrated circuit development. The materials under investigation, called strained-layer superlattices (SLS) are composed of many alternating layers of different types of carefully grown, ultra-thin (>300A) crystals. The SLSs can be tailored precisely to the task to be performed, so scientists are quite certain that they will be valuable in semiconductor technology. The flexibility of design arises because the layers can be chosen independently of the fact that they are mismatched — the distance between the atoms of the material used for a given layer can be different from the corresponding 'atomic spacing' of the material used for the next layer. The layers in SLS materials are so thin (it would take 5000 layers to equal the thickness of a sheet of writing paper) that their atoms align easily by elastic strain during growth, without any of the imperfections between adjacent layers that would otherwise degrade device performance. The imperfections, which plague similar structures with thicker layers, have prevented full use of many semiconductor compounds whose properties are superior to those currently in use. The new SLS semiconductor compounds come from elements in Groups III and V of the Periodic Table of the Elements, for example: gallium arsenide phosphide, indium gallium arsenide, and others.

SLS crystal layers are now being grown at Sandia by means of two specialized processes — molecular beam epitaxy (MBE) and metal organic chemical vapor deposition (MOCVD). These are not generally used to mass-produce electronic devices. However, recent progress in producing high-quality SLS materials makes researchers optimistic that SLS materials may eventually be produced in commercial quantities.

To date, most SLS work has been concentrated on theoretical studies, growth of materials and measurement of electronic and optical properties. So far, only very simple semiconductor devices, diodes, have been produced. Says John Galt, a member of the Sandia research team, "Strained-layer superlattices appear to be a major advance in semiconductor device science. The discoveries are very recent, and

the basic properties of various SLS materials are still being explored so that we can achieve a better idea of their potential for actual device applications. Nevertheless, we are confident that they will prove valuable in the design of improved electronic systems."

Note: Sandia does not claim credit for inventing the superlattice, but has done pioneering work in opening up the electronic and optical properties of superlattices to a broad range of potential applications by varying — with the strained-layer effect — the material properties of the structures.

LLNL laser process 'paints' ICs on silicon

In another vein — fabrication — scientists at Lawrence Livermore National Laboratory (LLNL) are now using pulses of laser light flashing millions of times per second through selected gases to 'paint' integrated circuits directly onto silicon wafers. When fully developed, the revolutionary new process, called 'laser pantography,' and a new computer code for making ICs can reduce the time needed to design, make, and test new circuits from weeks to a few hours. "We hope to have a system painting 1,000 transistors a second by the end of the year," said Dr. Irving Herman, a physicist and co-leader of the LLNL Special Studies Group.

In contrast to the traditional method of making integrated circuits, laser pantography shows potential to be fast, direct and reversible. The process involves rapid reactions using intense green laser light directed onto silicon surfaces with an intensity a billion times greater than noontime sunlight. At times, various gases are introduced into a reaction chamber. The laser is pulsed on and off so fast that reactions occur only in the center of the focused laser spot, about one micron in diameter. This equals the resolution of current integrated circuit technology.

The intense light superheats the silicon surface for such a short time that heating and cooling rates of hundreds of billions of degrees per second are achieved, giving the laser pantography process its great sensitivity and resolution. When the gases are admitted into the chamber, reactions occur on the hot surface that can remove or change the electrical properties of the silicon, or deposit the desired special electrical materials. Dr. Bruce McWilliams, co-leader of the Livermore project and amateur painter in his free time, describes laser pantography as "a set of laser paint brushes which we dip into paintpots of various exotic gases in order to paint on a canvas of silicon."

Scientists expect that laser pantography will one day be part of automated computer systems that will virtually 'reproduce' themselves without the extensive human participation and effort now required.

The Sandia story was prepared based on text appearing in Sandia National Laboratories LAB NEWS, Vol. 35, No. 12.

Sandia's DASSL solves differential-algebraic equations

Sandia National Laboratories has developed a new approach to solving a class of difficult computational problems that commonly arise in science and engineering. The problems are those involving differential equations in combination with nonlinear algebraic equations. These mixtures occur surprisingly often — for instance, in connection with partial differential equations. Most computer codes for solving differential equations are designed for solving systems that must be written in the standard form $y' = f(t,y)$. But often it is not convenient or possible to write systems in this restrictive form. A far more general class of problems can be formulated as systems of differential-algebraic equations — $F(t,y,y') = 0$. Until now, satisfactory techniques for their solution have been lacking.

That is where Linda Petzold of the Applied Mathematics Division at Sandia comes in. She developed the DASSL code which solves nonlinear differential-algebraic systems. "DASSL solves all kinds of problems that are important to Sandia," says Linda. "Problems in flame modeling, ignition processes, chemical vapor deposition processes, solar energy applications and components design are all written as systems of differential-algebraic equations." Before this code, scientists had to spend their time developing special purpose methods to solve these types of problems. DASSL relieves the user of the burden, and at the same time it provides the reliability and diagnostic capabilities that are required.

The mathematical methods used in the code are similar in many respects to those that are used for solving standard differential equations. Therefore, Linda's objective was to extend some of the relatively well understood solution algorithms for ordinary differential equations to a wider class of problems. "To deal with these problems," Linda explained, "we've developed the underlying mathematics, devised new computational algorithms and designed the Differential-Algebraic Solver code. In addition to its usefulness at Sandia, the code has also met with considerable success outside Sandia on problems arising in space shuttle flight dynamics, magma flow in volcanoes and electrical network design."

DASSL has been extensively tested and documented and is running on Sandia CRAY systems under COS. Those interested in obtaining a copy of the DASSL code should contact: Argonne Code Center, Argonne National Laboratory, 9700 South Cass Avenue, Argonne, IL, 60439.

This article was prepared based on text appearing in Sandia National Laboratories LAB NEWS newsletter, Vol. 35 No. 9.

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