EIGHT DOLLARS A CMP PUBLICATION AUGUST 1988

VLSI Systems Design

FOR DESIGNERS OF HIGH-PERFORMANCE SYSTEMS

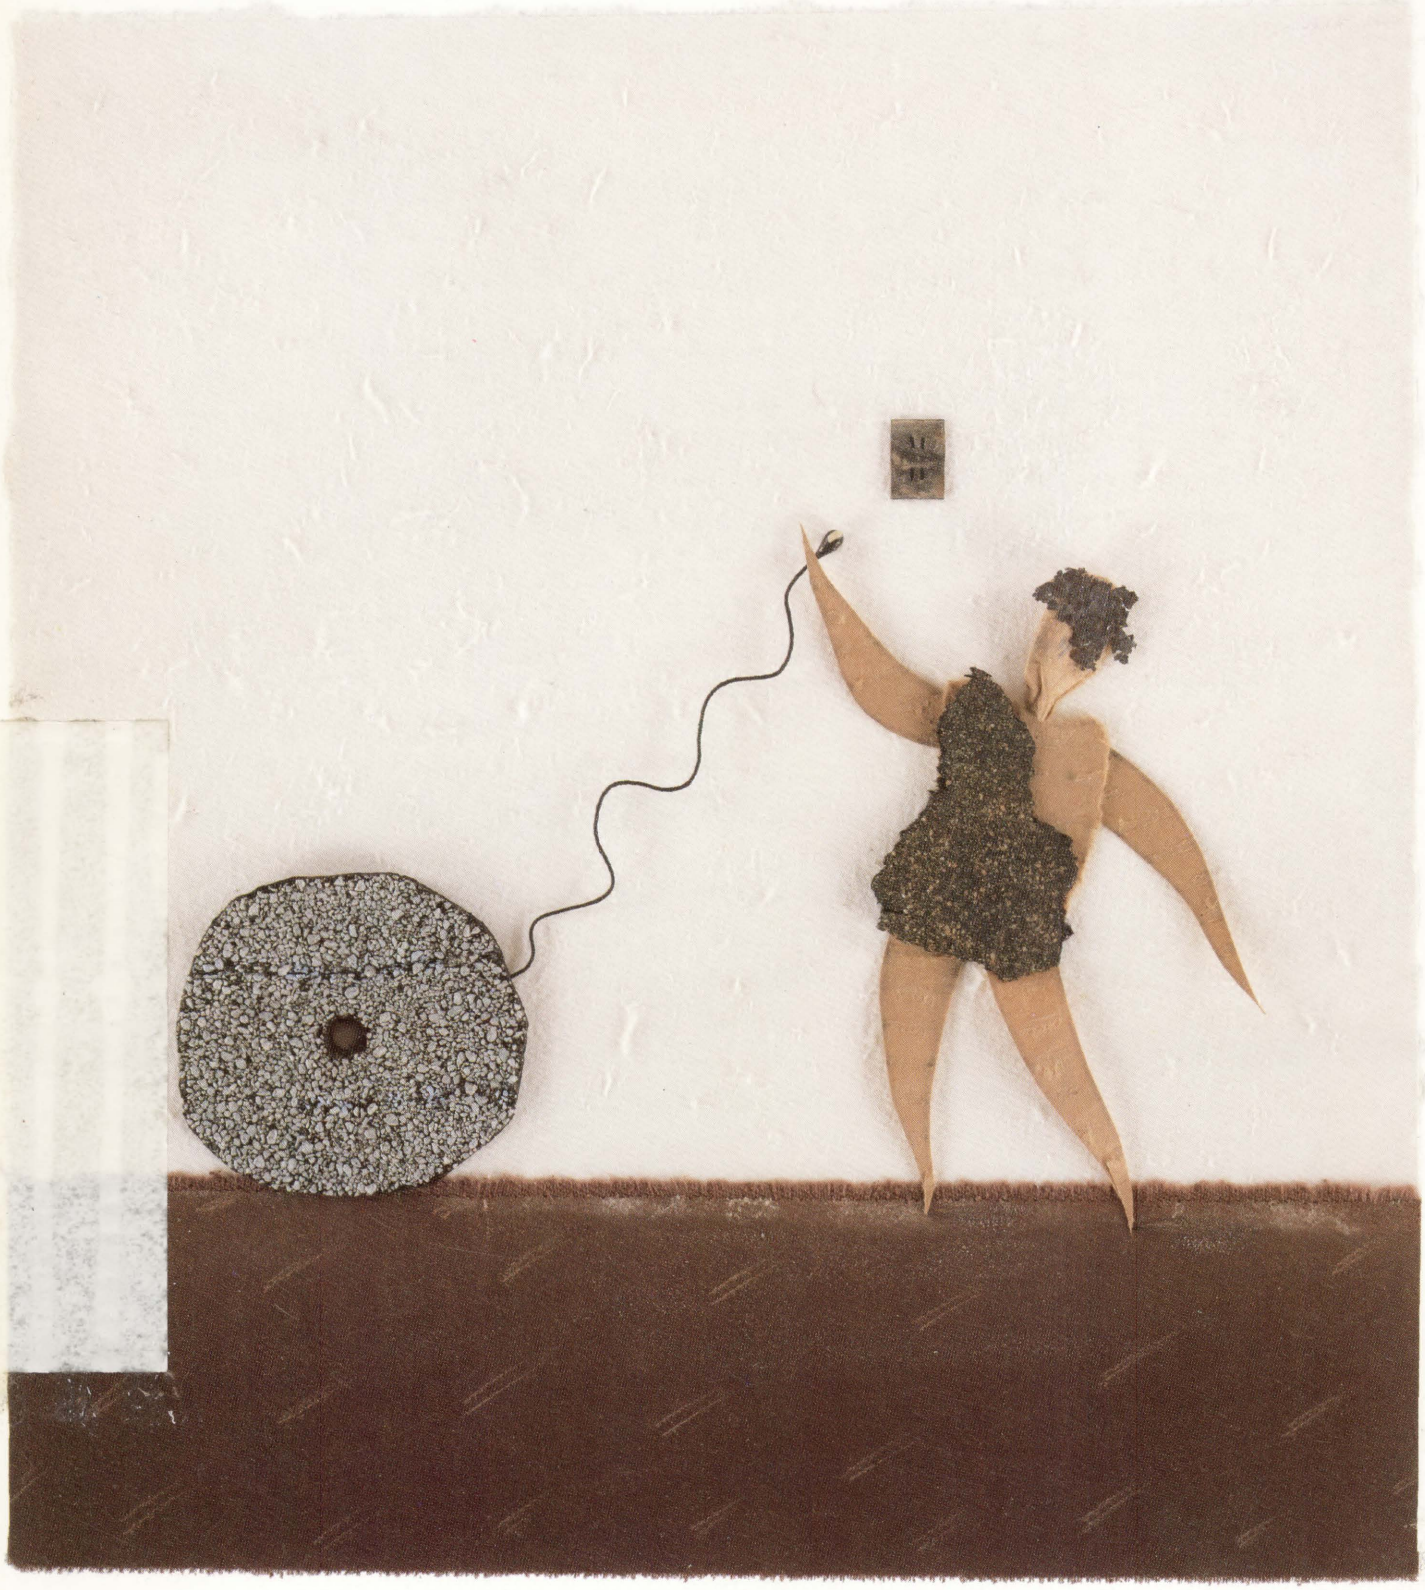

HARDWARE ACCELERATORS BOOST SIMULATION SPEED

rship that works.

In today's highly competitive marketplace, it takes more than technology and tools to meet your ASIC needs. It takes a long-term partner you can count on. One who's ready to help you turn your hot idea into an even hotter new product.

At Fujitsu Microelectronics, we believe that the only way we can achieve our goals is to help you achieve yours. So we've committed our technical, financial and human resources to providing you with the ASIC technology, tools and trust you need to meet your objectives. From design support through volume production of your advanced VLSI ICs. The heart of your new products.

When you shake hands with Fujitsu, you get more than just a business deal. You get the industry's most experienced volume producer of ASIC devices as a working partner.

We become an extension of your ASIC design team, providing you with a nationwide network of field application and technical resource engineers, ready to offer local technical support on both products and tools.

You also get ASIC sales and marketing support to help you smooth out all the administrative wrinkles.

Plus coast-to-coast technical resource centers, fully equipped with easy-to-use design tools. And local training and twenty-four hour design facilities, so you can work whenever inspiration strikes.

All supporting the most advanced process and manufacturing technologies in the industry, including CMOS, ECL, BiCMOS and LSTTL.

All of which is structured for the fast design input-to-prototype turnaround, design security and timely delivery of volume production you need to get to market faster.

It all adds up to a partnership that works. Which, after all, is everything an ASIC partner should be.

FUJITSU

FUJITSU MICROELECTRONICS, INC.

Everything an ASIC partner should be.

3545 North First Street, San Jose, CA 95134-1804 (408) 922-9000

CIRCLE NUMBER 1

You need a partne

When it comes to ASICs, you need more than technology that works.

Change Is Usually a Two-Way Street

Reader Feedback Is One of the Most Important Inputs to Our Publication

In this, our fourth issue since the redesign of VLSI Systems Design, we'd like to pause for a moment and thank our readers for the flood of comments that have been pouring in since May.

Change of any kind tends to make people uncomfortable. So we thought long and hard about the changes we were making to this magazine. By the time we debuted the May issue, we felt sure that what we were doing was right and would much better serve the needs of our audience. But the ultimate test would be whether or not you liked the "new" VLSI Systems Design.

And, it seems you do. In fact—judging from the hundreds of cards we've received with comments ranging from "Bravo!" to "You've made a great publication even better!" to "I love it!" to "Keep up the good work!"—your reaction was overwhelmingly enthusiastic.

To the handful of readers who voiced their concern that we might abandon our original editorial charter, you have our assurance that VLSI Systems Design will continue to provide the quality of editorial you've come to expect from this magazine.

To those readers who liked the first issue, but decided to reserve judgment until you'd seen a few more, we hope you've become a convert by now.

To the majority of readers—and advertisers—who share our enthusiasm and excitement over the changes we've made, our heartfelt thanks. You make the long hours, missed deadlines and hard work our editors put in all worthwhile.

Finally, thanks to everyone who took the time to write. Your feedback is important to us and your comments are always appreciated.

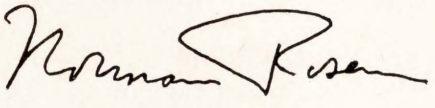

NORMAN ROSEN
PUBLISHER

Contents

ARTICLES

Hardware accelerators are a must if interactive simulation is required by the systems designer

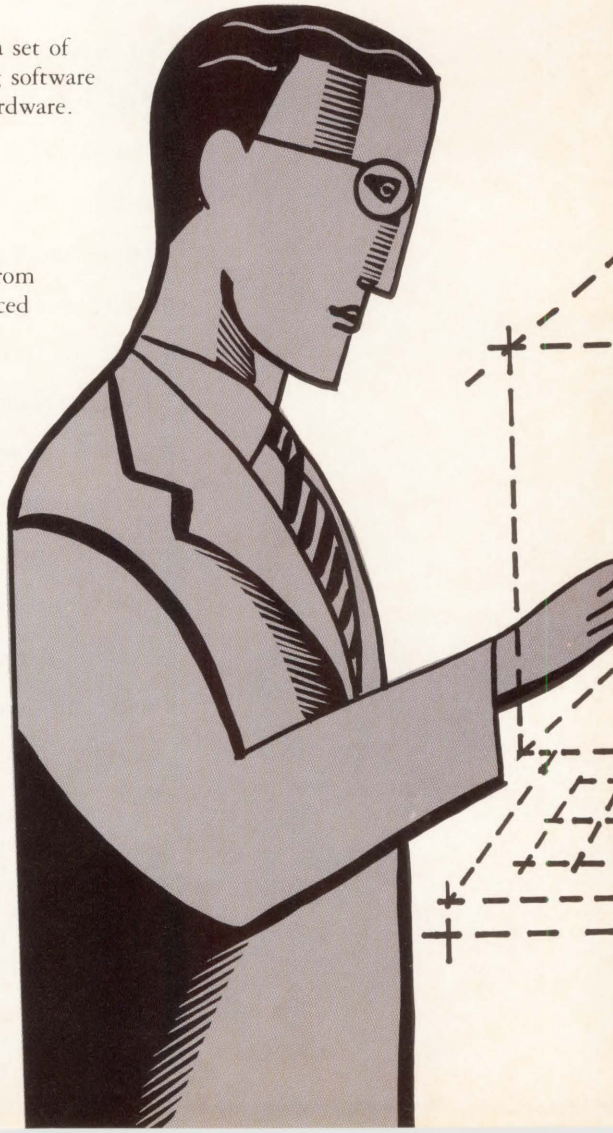

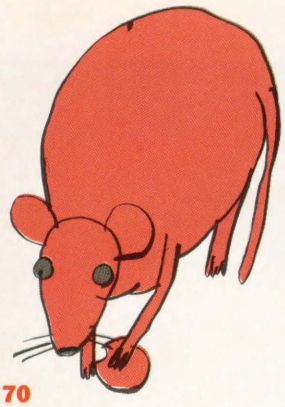

DEPARTMENTS

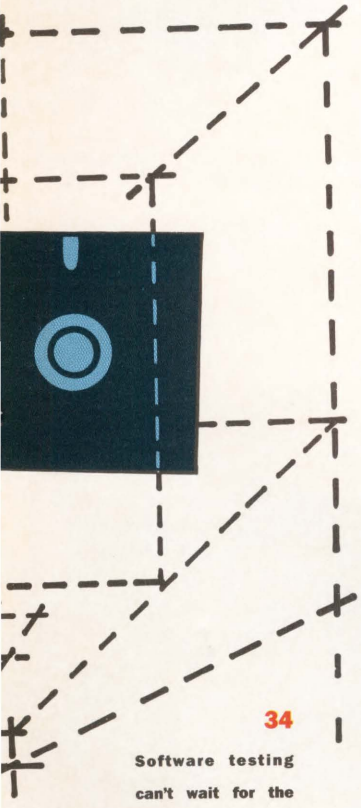

COVER ILLUSTRATION BY ALEXA GRACE

VLSI Systems Design

A CMP Publication

EDITOR/EDITORIAL DIRECTOR
Girish Mhatre

EXECUTIVE EDITOR
Roland C. Wittenberg

MANAGING EDITOR
Mike Robinson

SENIOR EDITOR
Bob Cushman

EDITOR EMERITUS
Roderic Beresford

EDITOR-AT-LARGE
Stan Baker

WESTERN REGIONAL EDITOR
David Smith

DIRECTORIES EDITOR
Michelle A. Losquadro

TECHNICAL ADVISERS
John A. Darringer
Jeffrey T. Deutsch
Edward J. McCluskey
Alan F. Podell
Daniel G. Schweikert
Susan L. Taylor

EDITORIAL PRODUCTION
Patricia L. Gaynor, Sr. Production Editor
Deborah Porretto, Ass't Production Editor

EDITORIAL ART
Sharon Anderson, Art Director
Marie D'Ippolito, Design Director

MANUFACTURING
Marie Myers, Production Manager
James Pizzo, Production Supervisor
Jane Mahoney, Asst. Production Supervisor
Vance Hicks, Coordinator

PUBLISHER
Norm Rosen

VLSI SYSTEMS DESIGN (ISSN 0279-2834) is published monthly with an extra issue in May by CMP Publications, Inc., 600 Community Drive, Manhasset, NY 11030. (516) 562-5000. VLSI SYSTEMS DESIGN is free to qualified subscribers. Subscriptions to others in the US: one year $60.00, two years $95.00; Canada and Mexico: one year $90.00, two years $165.00; Europe, Central and South America: one year $120.00, two years $225.00. Asia, Australia, Israel and Africa: one year $150.00, two years $285.00. Second-class postage paid at Manhasset, NY and additional mailing offices. POSTMASTER: Send address changes to VLSI SYSTEMS DESIGN, Box No. 2060, Manhasset, NY 11030.

CMP ELECTRONICS GROUP
Kenneth D. Cron
Vice President/Group Publisher
Electronic Buyers' News
Electronic Engineering Times
VLSI Systems Design

CMP PUBLICATIONS, INC.
600 Community Drive
Manhasset, New York 11030
(516) 562-5000

Publishers of:
Electronic Buyers' News, Electronic Engineering Times, VLSI Systems Design, Computer Systems News, Computer Reseller News, VARBUSINESS, UNIX Today!, InformationWEEK, CommunicationsWeek, CommunicationsWeek International, Business Travel News, Tour & Travel News, Long Island Monthly, HealthWeek

Michael S. Leeds, President
Pearl Turner, Vice President/Treasurer
Daniel H. Leeds, Vice President

Lilo J. Leeds, Gerard G. Leeds
Co-Chairpersons of the Board

National Semiconductor

Now you can mold the silicon to fit the logic. Not the other way around.

PROGRAM OUR GENERIC ARRAY LOGIC ANYWAY YOU WANT. REPROGRAM ANYTIME YOU WANT

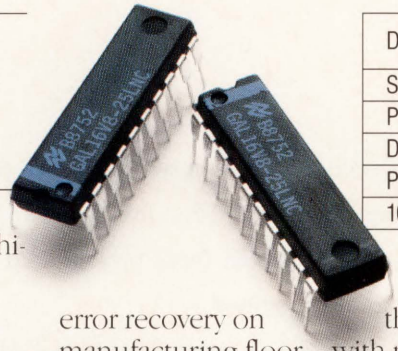

These new GAL® devices emulate the full range of PAL® architectures with 100 percent socket compatibility. They can even be configured to all the in-between architectures — like a 16R1 or 20R7.

You can also change architectures during the design process. And when your design is complete, you can move functions from pin to pin to simplify your board layout. You can even make changes during manufacturing — without huge penalties.

The one-device-does-all approach not only gives you unlimited design freedom, it also reduces parts counts and inventories and decreases your system costs.

REAL TESTABILITY

The first PLDs to use electrically erasable CMOS technology, GAL devices are instantly erasable and reprogrammable. (No time-consuming trips to the UV lamp.) And we guarantee 100 erase/write cycles.

This makes them ideal for prototyping and for pattern changes or error recovery on the manufacturing floor — with no waiting or waste.

More important, it allows us to program and erase each device many times during manufacturing to directly test all characteristics, including AC, DC and functionality. As a result, we can assure you 100 percent programming and functional yields without any board rework.

REAL REDUCTIONS IN POWER

Another advantage of CMOS technology is that these new devices are just as fast as bipolar but consume less than half the power — as little as 45 mA active and 25 mA standby.

And because they consume less power, you'll have fewer cooling problems and a much more reliable system.

REAL SOFTWARE SUPPORT

GAL devices are fully supported by industry-standard development tools, including ABEL,® CUPL,™ and PLAN.™ You don't need any sophisticated or expensive upgrades.

And you don't need any special programming hardware either. Almost all hardware manufacturers support GAL parts on existing models.

So get real design freedom. These new devices are like putty in your hands.

Call National — for information on GALs or any other type of PLD. We're the only manufacturer that offers three different technologies — ECL, TTL, and EECMOS.

For a free copy of our new programmable logic brochure, call your local National sales engineer, authorized distributor, or National Semiconductor, MS23/200, P.O. Box 58090, Santa Clara, CA 95052-8090, 1-800-252-4488, ext. 733.

Device Type	ECL PAL	TTL PAL	EECMOS GAL
Speed	fastest	fast	fast
Power	highest	high	low
Density	low	medium	high
Programmability	OTP	OTP	100%
100% Testability	no	no	yes

National Semiconductor

CIRCLE NUMBER 2

Balancing The Software-Hardware Team

TTL Logic Changed the Rules of the Design Game

The introduction of TTL logic about twenty years ago had a profound effect on the electronics industry. It resulted in the mass adoption of integrated circuits by the systems designer. But, even more profound, it changed the way that engineers designed digital electronic products.

Prior to the introduction of TTL, one of the digital designer's biggest problems was the implementation of his logic design in hardware. Debugging the breadboard was bad enough, but, when the product hit production, the designer's woes became almost unbearable. However, TTL was something different. Almost every new TTL chip had been put through the semiconductor manufacturer's torture tests. If an engineer implemented his logic design in TTL, and followed the manufacturer's design rules, the circuit was almost guaranteed to work.

The smart systems engineer quickly capitalized on this ability to design boards and systems with standard logic blocks. The advent of LSI and VLSI chips only served to accelerate this movement.

The up side was that systems engineers could now concentrate on a higher level of design—using logic blocks of complex functions rather than wasting effort on the detailed design of individual transistor circuits. Innovative logic design (and following close behind, innovative software) became the motto of the successful high-performance systems engineer. Logic and software design became the glamour jobs in engineering.

But there was a down side too. Hardware design lost a lot of its glamour. In the early 1980s, software engineers began to outnumber hardware engineers in many companies. The software teams performed magic (design automation tools, for example), while the hardware suffered.

Today, as faster and higher performance systems become a must for success, hardware design is again becomming critical. Fortunately many of today's leading systems designers have recognized this growing problem and are grabbing every available talented hardware engineer to try to keep the balance in their design team.

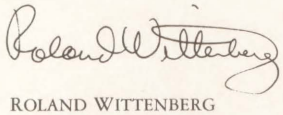

ROLAND WITTENBERG
EXECUTIVE EDITOR

Stop Wasting Power

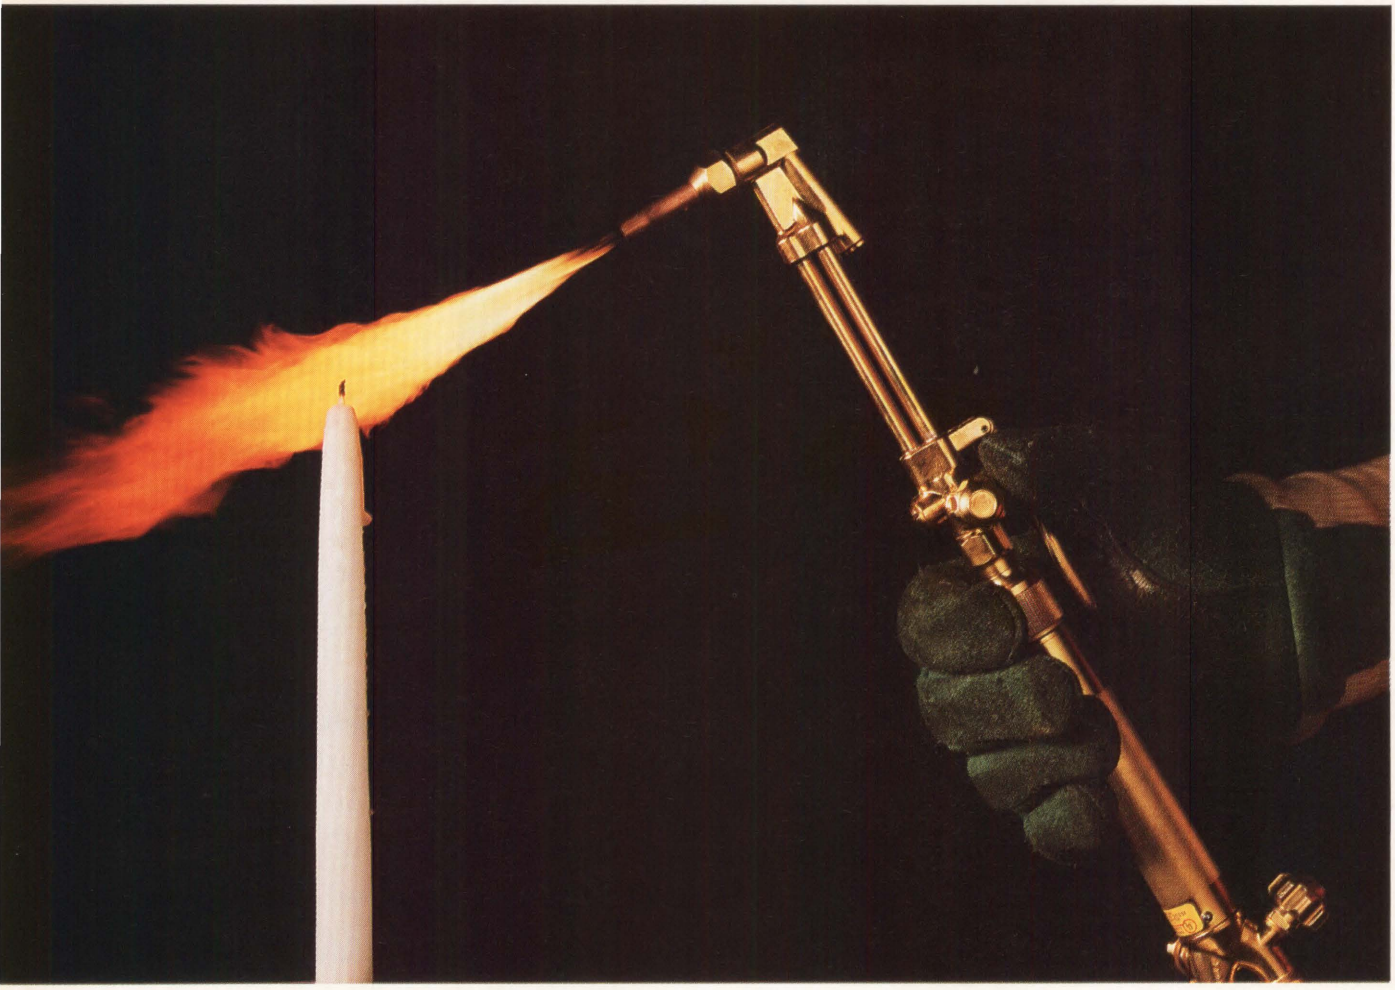

New CMOS array features the lowest power dissipation: 8μW/Gate/MHz

Raytheon's newest CMOS array family, the RL1000, helps you achieve optimum power performance. It offers the lowest power dissipation available—at high densities—without sacrificing speed.

☐ **Lowest Power:** Raytheon's 8μW/gate/MHz CMOS arrays, with 1.0 micron effective channel lengths, have available densities ranging from 5670 to 20,440 gates. With 20 tracks per channel instead of the standard 16, gate utilization is typically 90 percent or better.

☐ **High Speed:** At 0.3 ns unloaded inverter delay, this low-power, high-speed array family is perfect for portable equipment or where power sources are limited.

☐ **Other features:** The RL1000 series, with symmetrical switching delays, operates at 250 MHz flip-flop frequency and is TTL/CMOS compatible.

☐ **Packaging:** All packaging options are available. And Raytheon's design support includes an extensive macrocell library on major workstations.

☐ **Also available in CMOS:**

- Approved VHSIC PHASE I gate array: RVG 1.25 family
- Second source for LSI Logic Inc.'s LL7000 Series: RL7000

Call Raytheon for access to the right CMOS technology. Your success is our first priority.

Raytheon Company
Semiconductor Division
350 Ellis Street
Mountain View, CA 94039-7016
(415) 966-7716

Raytheon

Where quality starts with fundamentals.

CIRCLE NUMBER 3

Calendar

OIS '88

September 7–9
Washington Sheraton
Washington, D.C.

Sponsored by Meckler Corp., the eighth annual Optical Information Systems Conference and Exhibition will focus on write-once and erasable optical storage systems and digital document-image automation. Sessions are planned in such areas as electronic image and document storage systems, erasable optical disk media developments, erasable optical disk drives and systems, evaluating and selecting a WORM subsystem, write-once and erasable optical media manufacturing approaches, operating-system software for write-once optical disk, future trends and new developments, converging optical information technologies, and integrated system development. Additional information may be obtained by contacting Marilyn Reed, OIS '88 Conference Manager, Meckler Corp., 11 Ferry Lane West, Westport, Conn. 06880. (800) 635-5537. ■

7th VLSI and GaAs Packaging Workshop

September 12–14
San Jose, Calif.

This workshop is being sponsored by the IEEE's Components, Hybrids, and Manufacturing Technology Society and the National Bureau of Standards. Topics that will be addressed include package thermal design and interconnection options, GaAs IC packaging, die attachment solutions for large chips, VLSI and wafer-scale package design, new failure mechanisms in VLSI packaging, and VLSI package materials advances. Additional information may be obtained by contacting Paul Wesling, IEEE Council Office, 701 Welch Road, Suite 2205, Palo Alto, Calif. 94304. (415) 327-6622. ■

International Test Conference 1988

September 12–14
Sheraton Washington Hotel
Washington, D.C.

The International Test Conference provides a major forum for an exchange of information about the testing of electronic devices, assemblies, and systems. This year the conference focuses on the test techniques and equipment needed to meet the challenges of new technologies. Technical papers will be presented on such topics as analog devices, yield modeling and process diagnosis, testability analysis, education and training, application-specific devices, microcontrollers and microprocessors, printed circuit boards, wafer-scale assemblies, hardware and software, process and test data management, surface-mount assemblies, computer-aided engineering, quality and reliability, fault modeling and simulation, design for testability, contactless probing, fixturing, and computer-aided test generation. For additional information, contact Doris Thomas, Executive Secretary, International Test Conference, Millbrook Plaza, Suite 104D, P.O. Box 264, Mount Freedom, N.J. 07970. (201) 895-5260. ■

ESSCIRC '88

September 21–23
University of Manchester
Institute of Science
and Technology
Manchester, UK

The aim of the fourteenth European Solid-State Circuits Conference is to provide a European forum for the presentation and discussion of recent advances in solid-state circuits. Topics that will be covered include MOS, bipolar, and GaAs circuits; performance limits of integrated structures, data conversion circuits, circuit design in new technologies, design for testability, integrated filters, VLSI architecture and implementation, and CAD for IC design. For additional information about the conference, contact Dr. Peter J. Hicks, Dept. of Electrical Engineering and Electronics, University of Manchester Institute of Science and Technology (UMIST), P.O. Box 88, Manchester M60 1QD, UK. Phone: 061-236-3311 x2035. ■

1989 IEEE Aerospace Applications Conference

February 5–10, 1989
Breckenridge, Colo.

Sponsored by the IEEE Los Angeles South Bay Section, the Aerospace Applications Conference will offer technical presentations, free format workshops, and informal technical discussions and interactions. Topics of interest within the aerospace field will include system concepts, computer and microcomputer applications, system management, millimeter and microwave technology, communications and telemetry, software methodology, electro-optic applications, VLSI and semiconductor technology, instrumentation and measurement, graphics and display systems, and energy and space applications. Interested authors should submit, by September 6, 1988, two copies of a 500-word summary to Leo A. Mallette, Program Chairman, Hughes Aircraft, Building R-10, A9026, P.O. Box 92919, Los Angeles, Calif. 90009. (213) 334-2909. ■

Illustration by Renee Klein

Simple as...

pushing a button!

DESIGN ENGINEER	ASIC
CSPICE	PLD
HILO	LAYOUT

Intergraph makes electronics CAE easy. We offer interactive graphics solutions that do it all – from product concept and design through implementation – solutions that support hierarchical designs for PCB, ASIC, or hybrid applications. Using a common interface, you can capture schematics and verify both digital and analog designs with graphical waveform editors.

At one workstation...

Push a button and you're doing hierarchical schematic design with Intergraph's Design Engineer.* Push again for analog circuit simulation or digital systems simulation using Intergraph's CSPICE or the industry-standard HILO. Select ASIC, PLD, or LAYOUT for a wide choice of vendor libraries and design tools, supporting a variety of technologies.

Using our 5-MIPS CLIPPER workstations and servers, you can pass data via Ethernet and other industry-standard networks. Intergraph supplies the common database and all the application tools you need to control the entire development and revision process – and to get your products to market faster.

**For a low-cost solution to design entry, ask about Intergraph's Design Engineer PC.*

Intergrating CAE solutions...
Call Intergraph...800-826-3515 or 205-772-2700. CAE really is as simple as pushing a button.

INTERGRAPH

BIT Stream

High-Level Linkage Planned

Silc Technologies Inc. (Burlington, Mass.) and Vantage Analysis Systems Inc. (Fremont, Calif.) announced plans to interface their high-level design automation products. Vantage was the first commercial company to implement the full IEEE 1076 VHSIC Hardware Description Language (VHDL) standard, and Silc is one of the pioneer companies in logic synthesis. With the new interface, ASIC designers will be able to develop their designs using VHDL in the Vantage simulation environment, and then with Silc's Design Synthesis System automatically generate gate-level structural implementations. ■

Mips Latest Computer Delivers 20 Mips

Mips Computer Systems Inc. (Sunnyvale, Calif.) rolled out its latest RISC-based computer system, the M/2000 RISComputer, which provides a full 20 Mips of code crunching power. The new M/2000 is based on the company's recently announced R3000 chip set. It includes the 25 MHz RISC processor and the R3010 floating-point coprocessor plus 128 Kbytes of instruction and data caches.

Benchmarks made on the new computer include 42,000 Dhrystones per second and 14 million double-precision Whetstones per second. The M/2000 includes standard networking capabilities. It comes in 16 and 20 Mips versions, and is priced from $100,000 to $150,000. ■

Rad-Hard ASICs Take Off

Recent announcements from Honeywell Inc. and United Technologies Microelectronics Center Inc. (both of Colorado Springs, Colo.) boost the use of semicustom technologies in applications requiring radiation hardness. First, Honeywell has produced Radiation-Insensitive CMOS (RICMOS) technology that can withstand 3 megarads and still meet performance specifications. This 1.2-micron, bulk CMOS technology will be used to fabricate gate arrays with 4,000 to 15,000 usable gates.

UTMC announced Level S processing and screening for its UTD and UTD-R 1.5-micron arrays. UTMC's rad-hard UTD-R family can withstand a total dose of 1 megarad and meet specified performance. ■

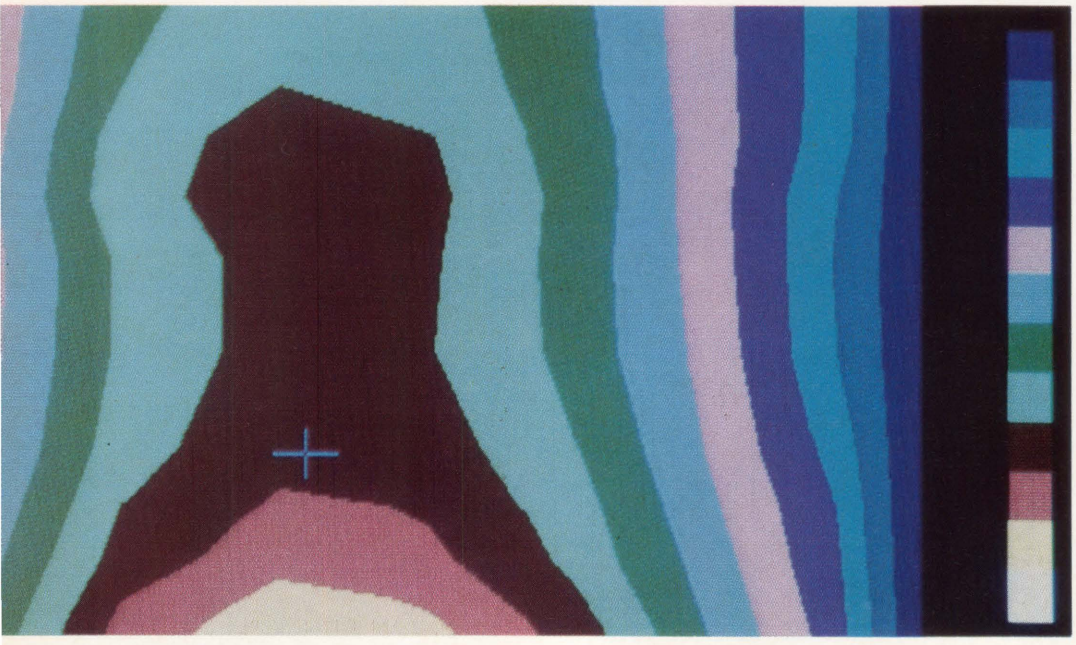

3D Thermal Analysis

Helios Software Engineering Ltd. (Tel-Aviv, Israel) has introduced its Thermax heat transfer CAD systen for the thermal simulation and analysis of printed circuit board assemblies. The new software package—which can run on a VAX system, or an IBM PC XT, AT, or compatible—provides 2D, 2.5D, or 3D color isothermic maps; one for the board and one for the components. The maps allow the user to interpret the temperature topography of the board and the junction temperature of the semiconductor components. Thermax can accomodate two types of conduction cooling, three types of convection cooling and radiation cooling. ■

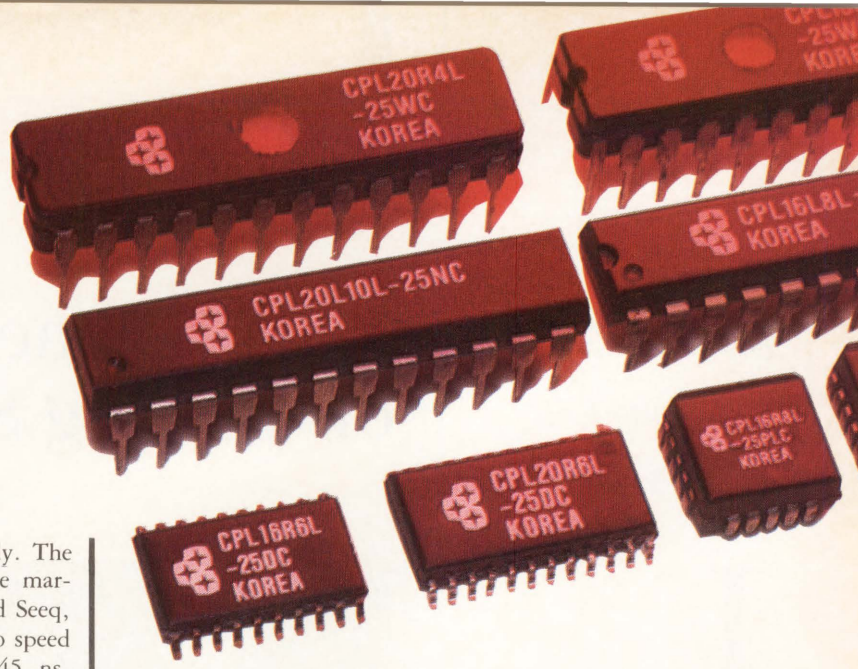

Joint Venture: PAL Architecture Goes CMOS

The first in a family of PAL products fabricated in an electrically erasable CMOS process was rolled out jointly by Advanced Micro Devices Inc. (Sunnyvale, Calif.) and Seeq Technology Inc. (San Jose, Calif.). The new PALC20RA10Z is a low power alternative for AMD's PAL20RA10, a 24-pin bipolar PAL device.

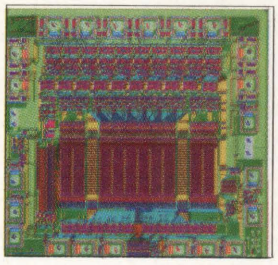

The new 24-pin device features zero standby power (less than 1 mW), and is reprogrammable in plastic packages. To provide maximum design flexibility, each of the 10 registered asynchronous macrocells in the new device can be clocked, set, reset, or bypassed individually. The product, which will be marketed by both AMD and Seeq, will be available in two speed versions, 40 ns and 45 ns. Standby current is specified at 150 μA maximum, while operating current is pegged at 5 mA per MHz.

The new PAL is the first product emanating from the four-year partnership begun two years ago by AMD's Monolithic Memories subsidiary and Seeq. It represents Seeq's entry into the programmable logic market, which is part of the company's strategy to broaden its base in electrically erasable processes beyond memory applications. ■

Complex Analog ASICs Get Personal Attention

Micro Linear Corp. (San Jose, Calif.) and Viewlogic Systems Inc. (Marlboro, Mass.) announced a joint marketing agreement to offer Micro Linear's Analog ASIC Design Library on Viewlogic's Workview systems.

The Workview systems run on IBM standard personal computers and Digital Equipment Corp. VAXs. In addition to providing schematic capture, PSpice simulation, and a waveform analysis tool, Workview systems have complete documentation and networking capabilities.

The library consists of Micro Linear's FB300 and FB3600 1 GHz analog tile families. The database includes symbols, subcircuits, component models, and worst-case design verification utilities. The company claims that

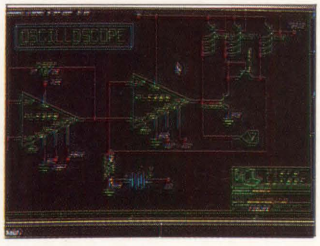

its analog library is unique in that all components are modeled as subcircuits to account for parasitics, in addition to temperature and voltage characteristics. The library also provides "extremely" accurate Spice parameters for the FB300 and FB3600 devices.

The analog library is available for immediate delivery. ■

First PLDs from Korean Manufacturer

With a family of nine CMOS devices, Samsung Semiconductor (San Jose, Calif.) has become the first of the Korean semiconductor manufacturers to market programmable logic devices. The CMOS Programmable Logic (CPL) family uses 1.2-micron, EPROM technology to implement four 20 pin devices and five 24 pin devices that are compatible with PALs. All devices come in 25ns and 35 ns speed grades with current drains of 70 and 35 mA, respectively. Samsung also offers a CPL starter kit with sample devices, a manual, data sheets and the CUPL software development package. The CPL line will be manufactured in Samsung's new CMOS facility in San Jose. ■

Daisy and Asix Forge Design-to-Test Links

Norman Friedmann, president and CEO of Daisy Systems Corp. (Mountain View, Calif.) and Wayne Pittenger, president of Asix Systems Corp. (Fremont, Calif.) inked a pact to exchange technology that will provide a bidirectional interface between Daisy's CAE workstations and Asix's line of ASIC testers.

The bidirectional link will not only transfer CAE simulation vectors to an ASIC tester, but it will also permit the designer to edit the simulation vectors in order to create an optimal test program, and then transfer the modified set of vectors back to the CAE environment, where they can be evaluated to verify that the new vectors provide the desired level of fault coverage. ■

Did Russian Tanks Help Forge Vlach's Saber?

Martin Vlach Converges on Analog Simulation

IN August of 1968, when Martin Vlach was seventeen, his family left their home in Prague for a vacation. That night Russian tanks rolled into Prague to squash the blossoming spirit of democracy in Czechoslovakia. After a childhood of security and predetermined futures, Vlach found himself facing an uncertain future in the alien, free-wheeling West. But now, as creator of the Saber simulator, he has found a foothold, and a purpose, here in his adopted country.

"It seemed that when I was in Prague growing up, there was one course for my life," he explains. "I would go to high school, I would go to the university, I would get, possibly, a research job somewhere in Prague, and I would live out my days in Prague."

The Russian invasion changed all that. Unsure of what future might await them in their home country, the Vlachs turned to the West for a new start. Fortunately, Martin's father, Dr. Jiri Vlach, had spent several years as a visiting professor at the University of Illinois. The family returned to the university, and young Vlach took classes, although not enrolled in a degree program.

The most intractable part of adjusting to a new culture, Vlach says, was the difficulty in expressing himself in a new language. As a result, he felt out of place: "In Czechoslovakia, I was always in the center of things. I stopped being so when I moved to the United States."

When his father became a professor of electrical engineering at Canada's University of Waterloo, Vlach enrolled in the physics program on his father's advice. Still unispired and undirected, he flunked out.

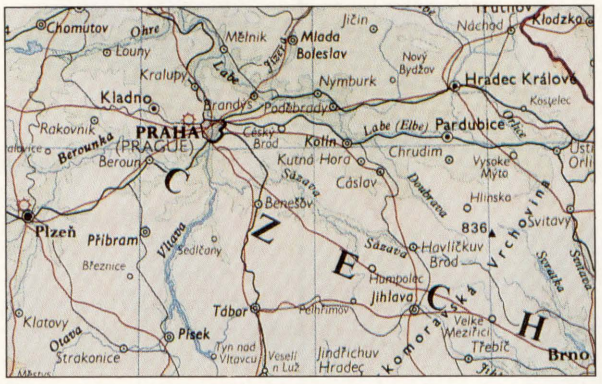

After a summer of pumping gas and working in a restaurant, he returned to Waterloo to try again. He began work as an "apprentice" in a simulation group at the university. Quickly proving his ability, Vlach became the group's focal point, implementing the simulation models and algorithms of professors and graduate students associated with the group. The simulator was WATAND, a public-domain circuit simulator that is still in use.

After a year with the simulator, he was motivated for his mathematics degree. He even began his Ph.D. work, but perhaps more for the university environment than his academic appetite. "I rejected an offer from Bell Northern Research because I didn't want to work" he admits. "University was too much fun."

Two weeks before his Ph.D. proposal was due, however, he was still debating whether to pursue his original interest in modeling or to develop simulation algorithms that he had recently conjured up. In a fortuitous turn of fate, Vlach ended up deciding to pursue his newer ideas about simulation algorithms. The algorithms, a more exclusive and esoteric area of study, satisfied his desire to be back in the middle of things: "I like to be a big fish in a small pond." During his Ph.D. work from 1981 to 1984, he began to write some of the fundamental algorithms that are now a part of Saber.

'I ENJOY THE POSITION OF PEOPLE SEEING THAT I DO SOMETHING WELL'

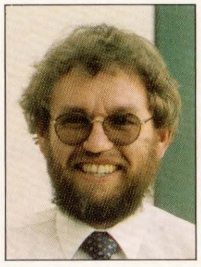

In 1984, Analogy Inc.'s president, Larry Jacob, was looking for academicians and experts to help him create a company to write simulation tools. Jacob was lead to the Dr. Jiri Vlach through his textbooks, and he in turn steered Jacob to his son, Martin. Young Vlach agreed to help Jacob get Analogy off the ground, and in April 1986, the company was formally launched.

"At Analogy, I finally got to do all the things I wanted to do," Vlach says. The technology behind Saber unites two of his Ph.D. proposals: the general-purpose simulation algorithms that he created joined with the flexible description language, MAST, that embodies his emphasis on modeling. Even today, two years after MAST was written, the language is still revealing new powers." Other MAST users continuously find new modeling techniques that even Vlach didn't envision.

For now, he seems content with his position at Analogy and life in Oregon. Perhaps Martin Vlach, through the success of his simulation efforts, again feels that he is a part of things. "I enjoy the position of people seeing that I do something well, and letting me concentrate on it."—David Smith

Feel trapped by one ASIC supplier? Reach for the Super Foundry.

Now you can design complex ASICs to meet your needs, not to fit one supplier's limited capabilities. The Super Foundry™ offers the combined resources of the world's premier CMOS manufacturers. Our process independent design tools and production services give you maximum freedom and control—from concept through delivery.

FREE. ASIC Estimating Kit.

What will it take to do your ASIC? With our free ASIC Estimating Kit, you can analyze design-process trade-offs and explore the performance, cost and scheduling implications of your design in a variety of processes. Take control of your ASIC design and avoid the single-supplier prison. Call the Super Foundry for your free ASIC Estimating Kit: ***1-800-FOR-VLSI ext. 200.***

THE SUPER FOUNDRY

SEATTLE SILICON

3075-112th Ave NE., Bellevue, WA 98004, (206) 828-4422.

CIRCLE NUMBER 5

S·C·F·L

A Solution To LSI-Level GaAs

BILL DAVENPORT AND DONALD LARSON,
TRI QUINT SEMICONDUCTOR INC., BEAVERTON, ORE.

Gigahertz GaAs LSI chips become readily available

Designers of systems with high clock frequencies often turn to GaAs devices for their critical signal paths. Most often, they use discrete devices for applications ranging from supercomputers to cellular phones. TriQuint has developed an alternative for such designers. It's a QLSI standard-cell library that uses Source-Coupled FET Logic (SCFL) circuitry to implement high-speed, LSI-level circuits. Although the use of SCFL introduces some complexities into modeling and layout, TriQuint has hidden most of them so the benefits of SCFL are readily available to the systems designer.

There are several apparent attributes of SCFL for developing a standard-cell library for very fast logic. Its differential current switching, with small voltage swings, is very fast. SCFL has a symmetrical nature (in terms of circuit structure and layout), minimizes signal skew (Vu, 1988), and ensures good fidelity of waveforms. In addition, each logic function generates both senses (true and complement) of every signal, so there is potential for simplified logic designs.

But SCFL's differential current switching also introduces some complexities to cell and circuit design. It uses stacked levels of differential amplifiers, each with its own logic level. All of these logic levels must be available at the output of the cells in order to match the various input levels that the outputs might be called upon to drive. The logic levels also have different prop-

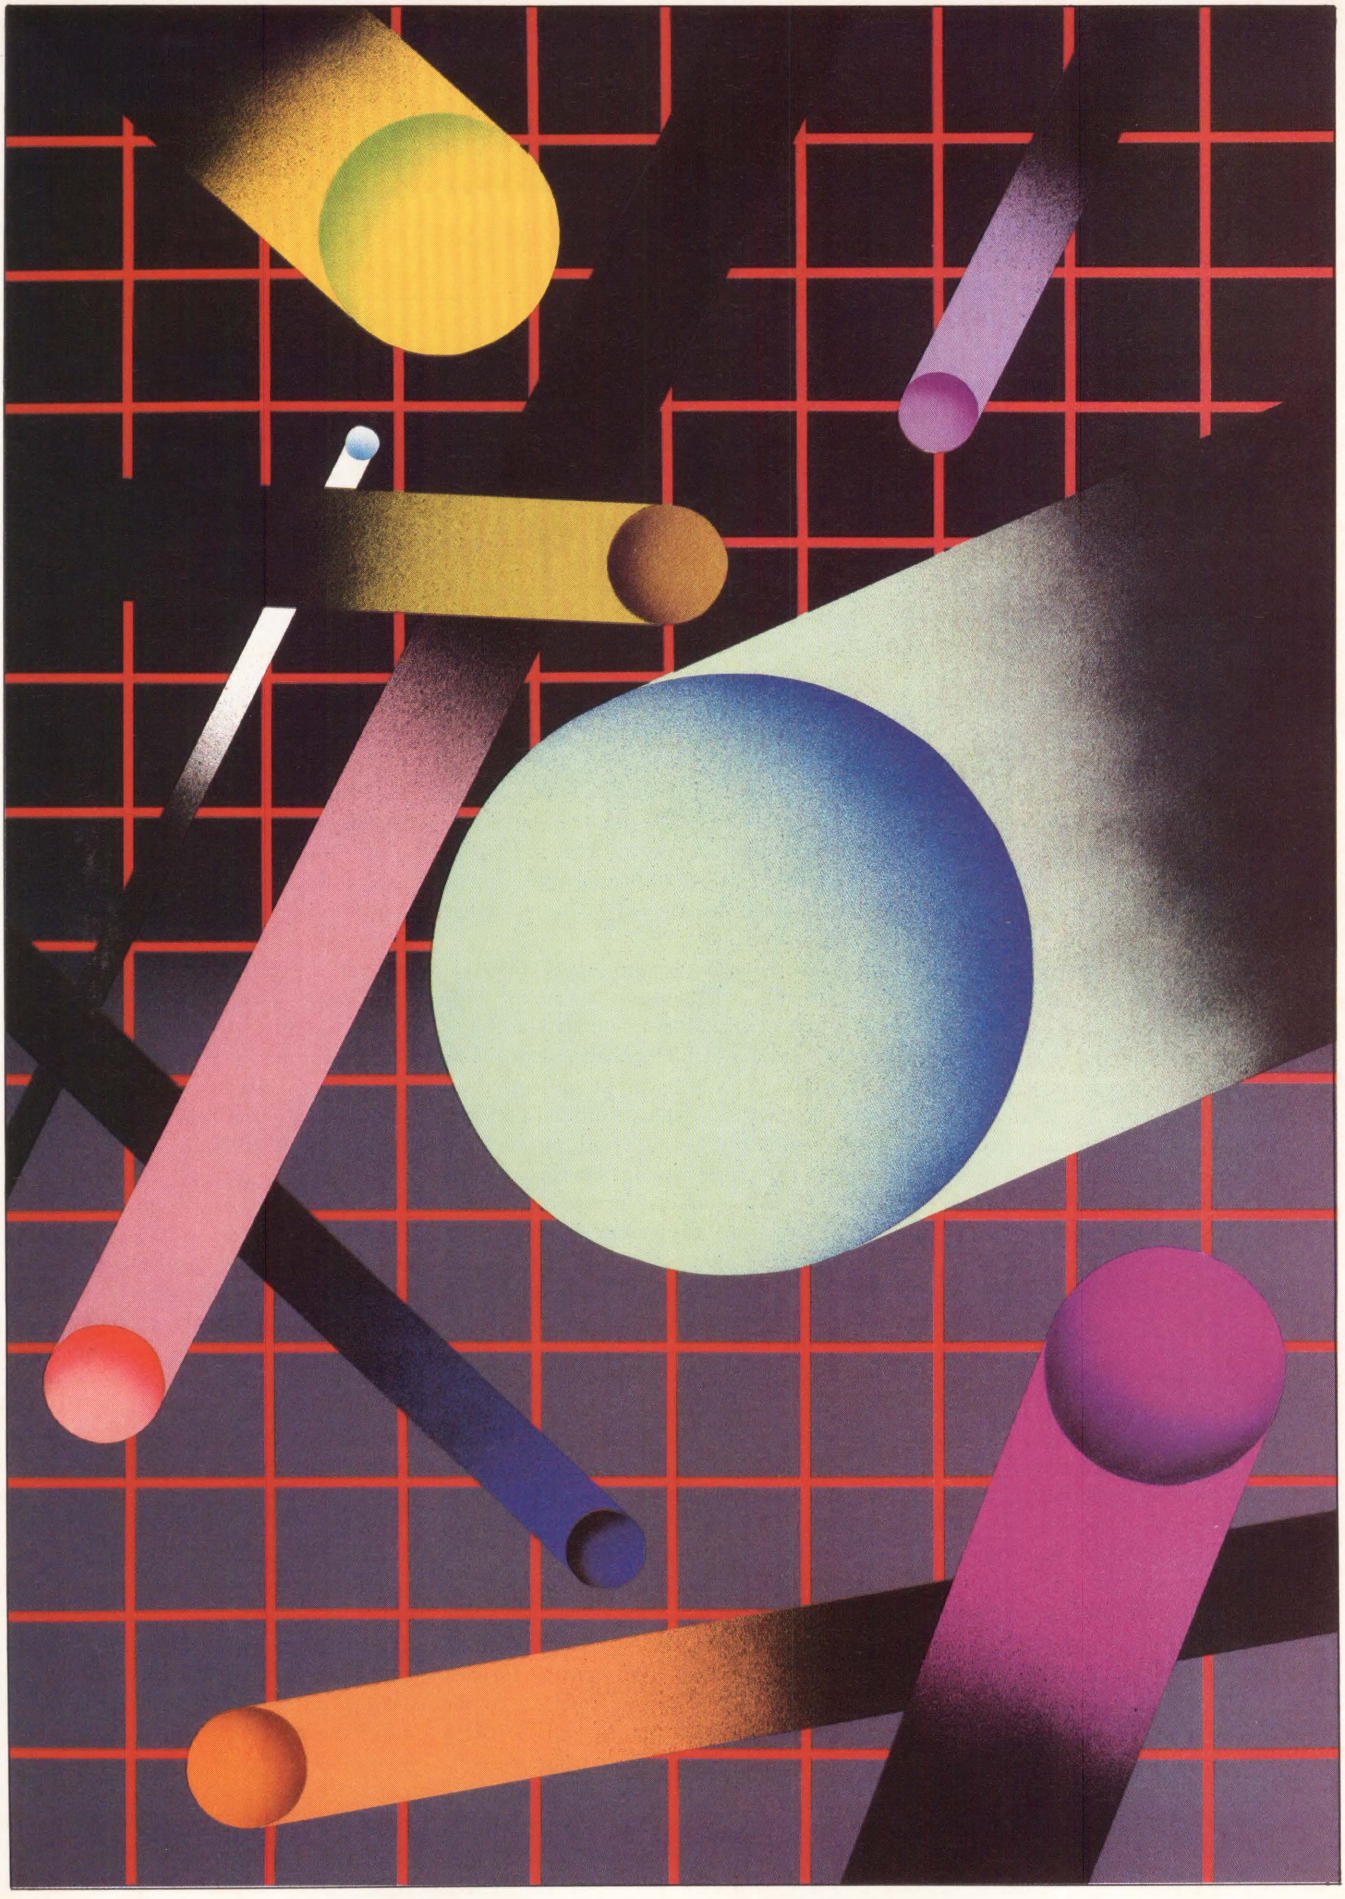

ILLUSTRATION BY GENE GREIF

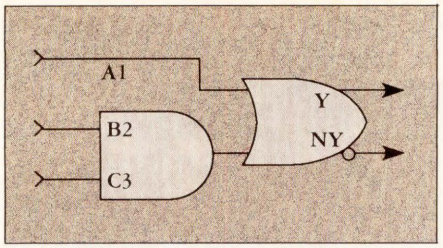

Figure 1. An AND/OR schematic symbol hides the complexity of the SCFL circuitry making design capture and simulation on a workstation straightforward.

CELL CHARACTERISTICS

FAMILY/NODE		DELAY TO Y AND NY (ps)	DELAY (ps/fF)	C_{in}(fF)	POWER (mW)
SM11	A1	65	0.55	25	8
SM11	A1	95	0.55	25	8
SM11	A1	125	0.55	25	8

agation delays. Each logic level—with a true and complemented signal—must be carefully routed to balance and minimize parasitic loads.

In order to incorporate SCFL into a standard cell family, TriQuint tries to address these issues through modeling and automation. With simplified models and specified propagation delays, the designer can use a standard CAE workstation to create and simulate designs. The TriQuint net-lister program, which generates a net-list for routing, also inserts the proper wiring to account for multiple logic levels and differential signals. By following some simple rules, systems designers can use the library to create successful semi-custom GaAs ICs.

■ WHY SCFL?

For years, researchers have investigated SCFL circuits for high-density GaAs chips. SCFL has been used since the early 1980s for devices operating above 5GHz in laboratory processes (Shimano, 1981). Production processes, however, are limited to flip-flop toggle rates between 2 and 3 GHz. Digital applications, requiring good noise margins, usually don't push the maximum limits of the process. However, linear applications take full advantage of the differential circuit.

Noise margin is a principle problem in LSI-level GaAs. In fact, it is the first requirement of any logic family. LSI-level designs exacerbate noise-margin sensitivity because they have longer signal paths and greater on-chip power-supply drops. In addition, the GaAs FET has less gain than its silicon counterparts, and there are no complementary transistors available.

Nevertheless, SCFL has been shown to have an allowable range of threshold voltage, which makes it effective for LSI applications (Katsu, 1982). Its range is actually four times greater than direct-coupled FET logic (DCFL), another common GaAs technology. This advantage stems in part from the differential nature of SCFL, which doesn't require a special reference voltage. The logic-switching level in SCFL is not directly related to FET threshold voltage.

This situation contrasts with single-ended technologies such as pseudo-ECL, which uses an on-board reference voltage. The reference is compared with the input signal, requiring the input signal to swing twice as far as a differential signal for the same current-switching capacity. With the gain of the GaAs MESFET's, which is lower than that of the bipolar transistors in ECL, use of a fixed reference would give away a significant amount of equivalent input swing. This translates into decreased noise margins.

Some other high-speed technologies, such as GaAs HEMT ICs and submicron bipolar technology, depend on more advanced processes, while SCFL remains on production GaAs processes. SCFL is probably the fastest GaAs topology, which is due to the symmetrical design and layout of the logic cells. Each cell uses similar interconnections to load resistors, differential amplifiers, and current sources.

SCFL has another advanatge for GaAs ICs. Its differential signals have twice the logic voltage swing of single-ended technologies like DCFL, so an enhancement/depletion process for SCFL can be optimized for performance. The transistors can be fabricated with higher pinch-off voltages because of its larger swing. Such transistors have higher gain and higher f_T, which translates into better performance.

Power dissipation is another concern for designers of LSI-level chips. Because each SCFL cell uses several current sources, it isn't the lowest power choice. Reliable, lower-power SCFL cells, which can be used in non-critical paths, help ease this problem. In addition, the speed-power product of SCFL cells can be compared with single-ended alternatives, especially for register-intensive designs. SCFL technology can create a flip-flop in only two cells; other technologies require at least four.

SCFL can help make such designs feasible by providing the level of performance required for clocking and control signals. This is done by limiting their use to the implementation of specific critical paths in LSI-level designs. The adoption of a cell-based methodology allows designers to use only the higher-powered cells they need. Potential users of GaAs want large, high-speed LSI designs similar to high-density CMOS and bipolar devices. Many designs, particularly military and portable applications, also need low-power consumption.

Expectations of such users will mature as designers understand the tradeoffs of higher clock rates and the resulting demands on power. If the designer wants the entire design to clock at 2 GHz, he can't use 10,000 gates, but if he needs to integrate 10,000 gates, he must limit the number of gates that run at 2 GHz..

TriQuint's standard-cell library also contains a cell family—zero-diode FET logic, or ZFL —that is denser with lower power needs, but not as fast as SCFL. Circuits out of the critical path can be implemented in the 1-GHz ZFL cells, maintaining maximim density and minimum power consumption. Fortunately, SCFL and ZFL designs can be combined on one chip. This allows medium-density applications (found in many workstations) to contain a portion of 2-GHz circuitry.

TriQuint's ZFL is a single-ended logic family using only one logic level. Its routing requirements, therefore, are much less stringent than SCFL's. Mixing SCFL and ZFL on one chip introduces complexities, such as logic-level shifting and more power buses. ZFL's one logic level is different from those in SCFL. Furthermore, it uses only a -2.5V to -3V supply voltage which, with the -4.5V to -5.5V for the SCFL cells and the +5V needed for CMOS or TTL-level I/O cells, makes as many as three power-supply buses plus a ground bus for the chip. TriQuint routes power buses separately from the signal I/Os to keep this complexity from affecting the use of a commercial router.

■ SCFL CIRCUITRY

TriQuint's QLSI SCFL logic may use as many as three levels of source-coupled differential amplifiers in a current-steering tree. A buffer or inverter cell uses only one level, while a three-input logic gate uses all three. Level one is defined as the top level. Levels two and three are one and two diode drops, respectively, below level one. A three-input AND gate, for example, requires that each set of differential inputs be at a different level. This is easily accomplished because every SCFL logic function

PERFORMANCE

Wafer Foundry Capacity Available.

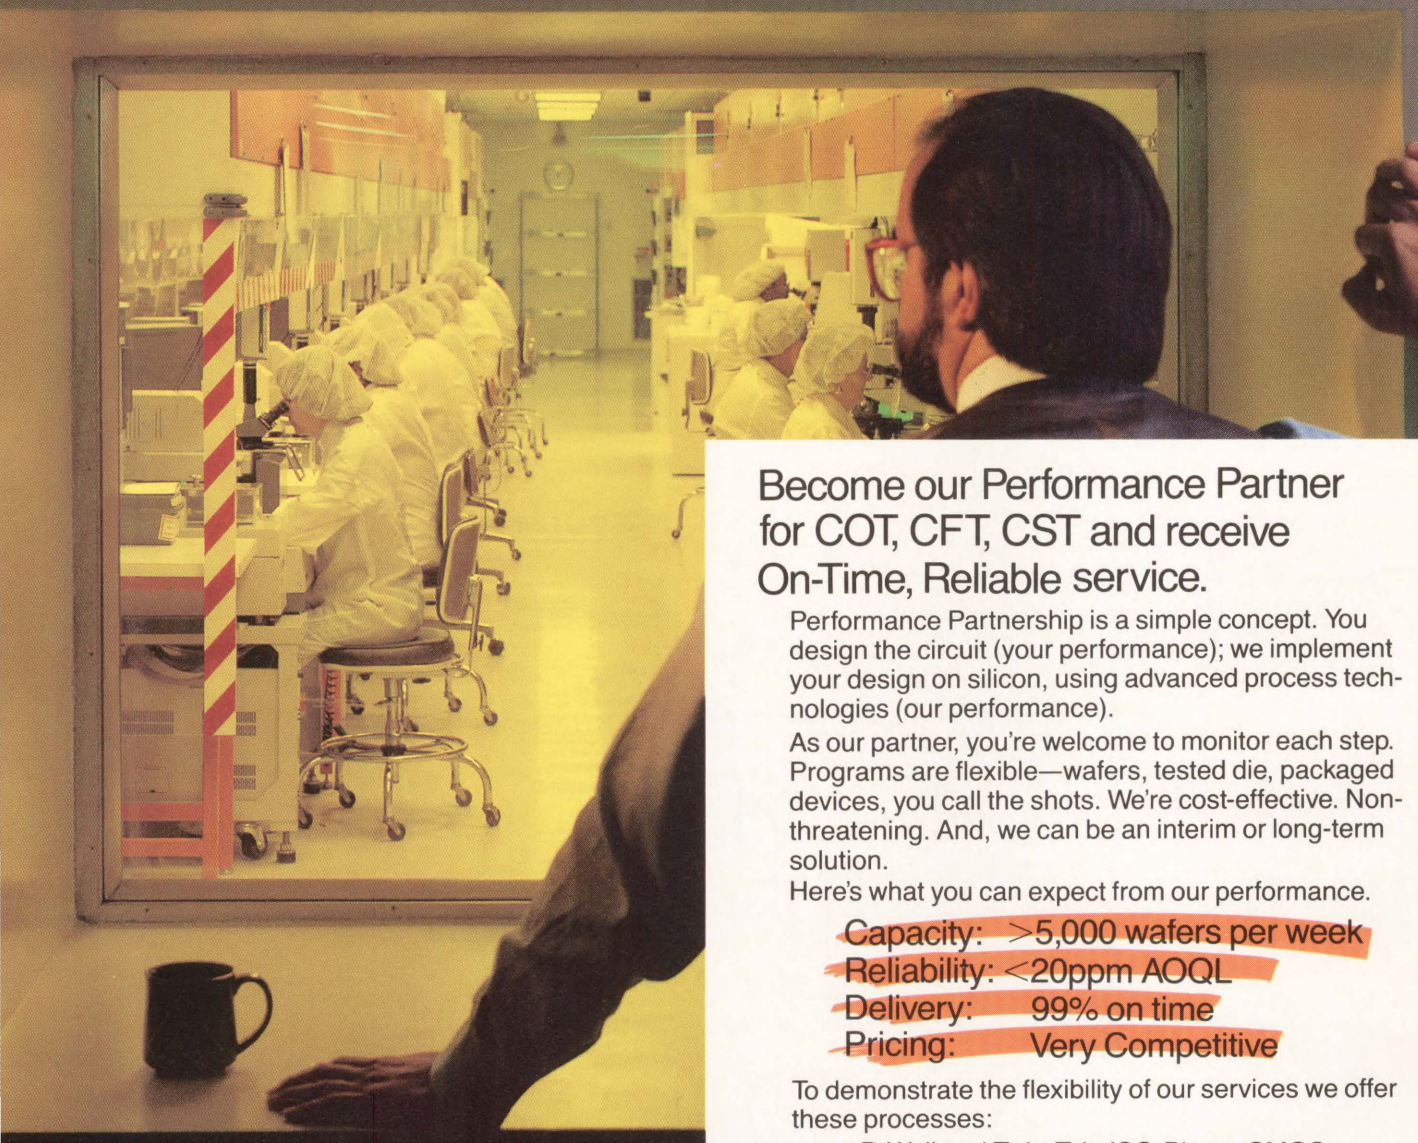

Become our Performance Partner for COT, CFT, CST and receive On-Time, Reliable service.

Performance Partnership is a simple concept. You design the circuit (your performance); we implement your design on silicon, using advanced process technologies (our performance).

As our partner, you're welcome to monitor each step. Programs are flexible—wafers, tested die, packaged devices, you call the shots. We're cost-effective. Non-threatening. And, we can be an interim or long-term solution.

Here's what you can expect from our performance.

Capacity:	>5,000 wafers per week
Reliability:	<20ppm AOQL
Delivery:	99% on time
Pricing:	Very Competitive

To demonstrate the flexibility of our services we offer these processes:

- P-Well and Twin-Tub, ISO-Planar CMOS
- Linear and High-Speed Bipolar
- Silicon Gate NMOS
- Double-Poly Process for Mixed Signal Technology
- Double Metal Process for High Density Applications
- 2 Micron Stepper-based Processing

Whether your need is for commercial or military grade wafers, write today for more information. For faster service, call our President, **Paul Bhasin** at **(602) 921-6011.**

California Micro Devices Corp.
Microcircuits Division
2000 West 14th Street • Tempe, AZ 85281
(602) 921-6000 • FAX (602) 921-6598 • TLX 187202

CIRCLE NUMBER 6

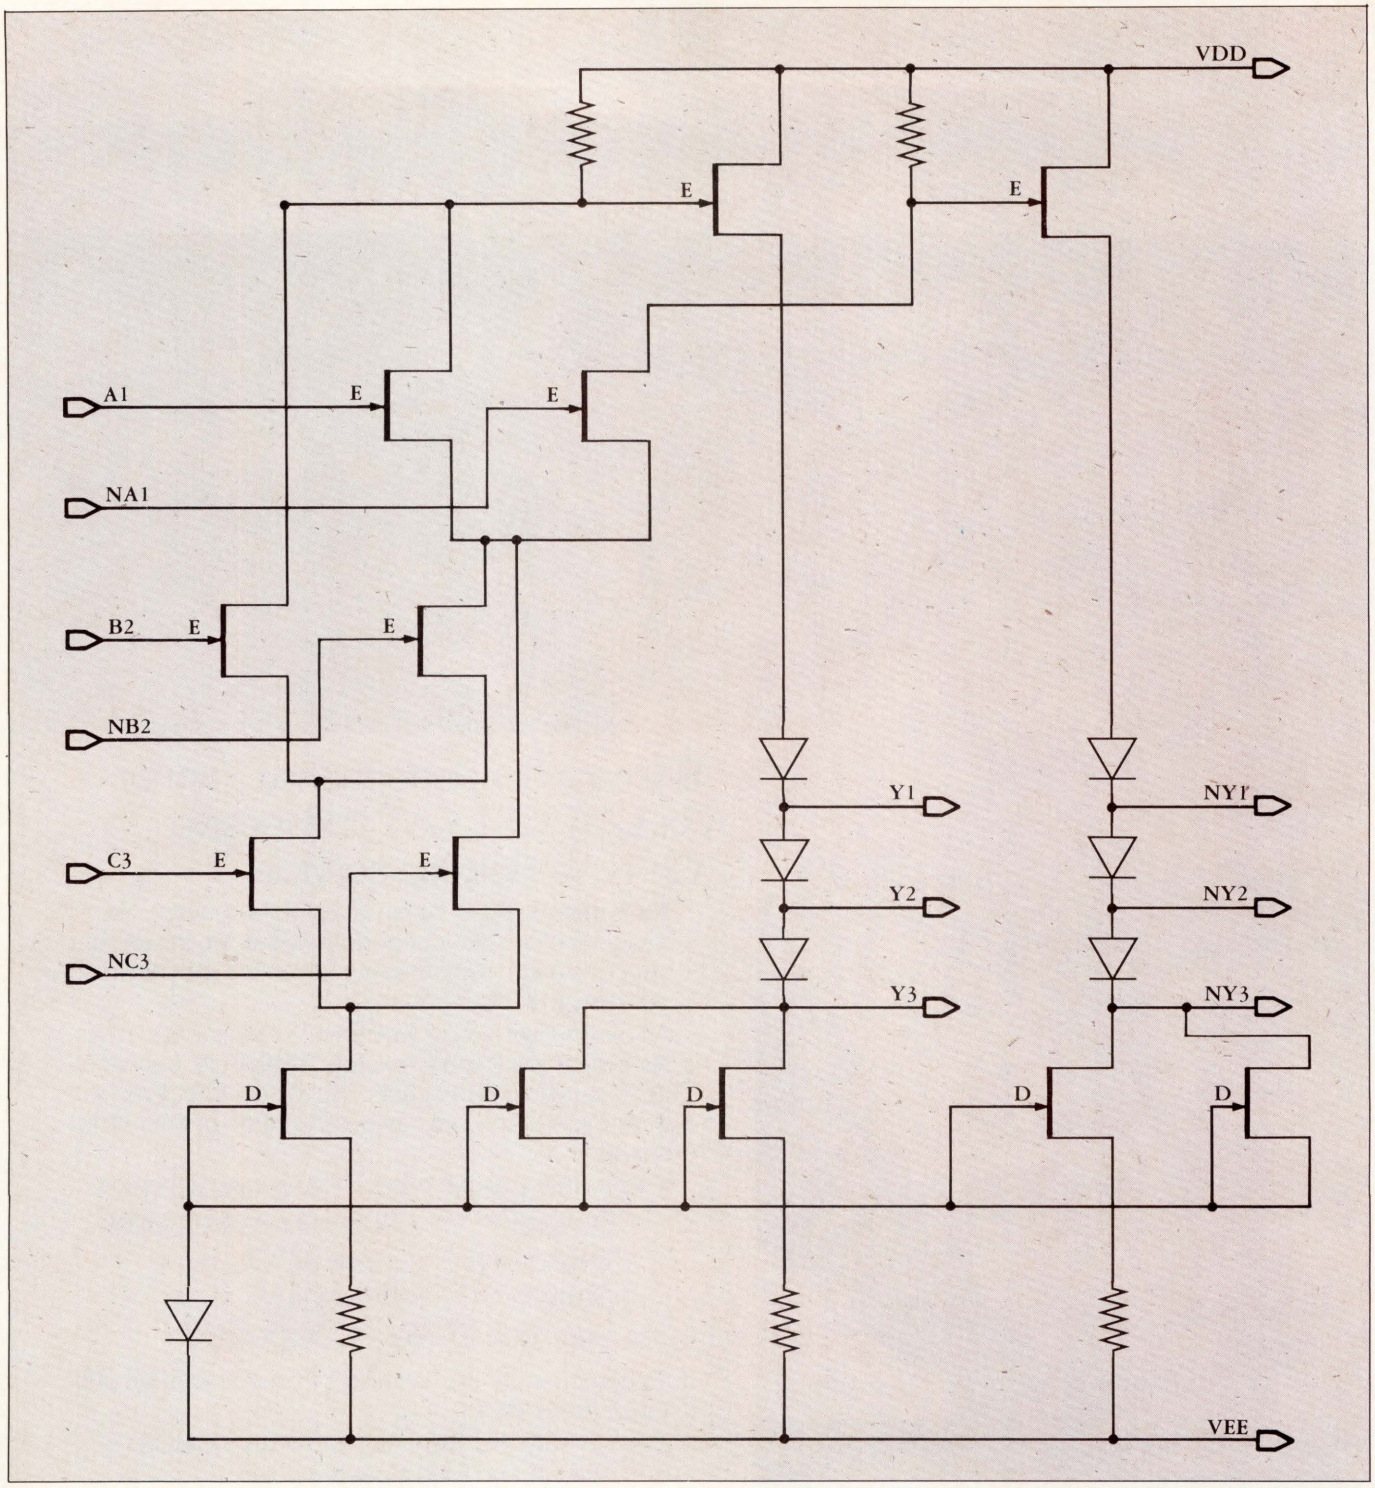

Figure 2. The actual circuitry for the AND/OR logic gate (Figure 1) shows the multiple levels of differential current steering (where the input signals connect) that distinguish SCFL circuitry. In addition, the three levels of differential output signals are shown at the cathodes of the diodes.

generates all three output levels.

Figure 1, for example, shows the schematic symbol for an SCFL AND-OR logic function. Figure 2 shows the corresponding schematic. Each input transistor pair steers one level in the current tree. Outputs of the tree drive complementary source-follower buffers for the true (Y) and complement (NY) sense of the output. The three levels of output signals—Y1 and NY1, Y2 and NY2, and Y3 and NY3—are found at the cathodes of the diodes in the output stacks. The current sources can be seen below the input and output circuits.

The diodes in the output circuits give the differential amplifiers in the current-steering tree sufficient "head room" (voltage swing) to operate in their linear region. (Note that, for GaAs FETs, TriQuint considers "linear" in terms of linear device analysis. Digital designers may consider them in the saturated mode).

Usually an output drives all its loads at one logic level. For example, all D and ND inputs of all latches and flip flops are on level one, while CK and NCK are all on level two. However, a single gate structure can drive multiple inputs, each at a different level. When an output must drive multiple levels, separately routed wires must connect the inputs at each level with the corresponding level on the output. However, the total capacitive load due to routing will be the sum of the routed capacitive loads from each of the output levels.

One benefit of using SCFL is the differential signals. To maintain good signal waveforms and minimal skew, the designer must keep routing capacitance of the two differential signals within 15 percent of each other. Fortunately, their loads (in terms of gate inputs) are the same, and TriQuint's "airbridge metal" keeps interconnection capacitance to roughly 1/7 of the total driven capacitance. This low interconnect method increases performance, and causes very little cross-coupling.

SCFL also has three output taps at the cathodes of the diodes. All three levels of the outputs are available to drive logic circuits. They can be used singularly or all three together to drive different loads. The outputs switch almost simultaneously, with only picoseconds between state changes on different logic levels.

TriQuint has built a library of functions with this logic structure. The functions range from inverters and buffers to three-input exclusive-ORs and set/reset flip-flops. More complex macros are promised.

These libraries run on workstations from Mentor Graphics and Daisy Systems. The schematic symbol for each cell has a single pin for each differential input, and two pins (each representing three levels of output) for each output signal. This simplification dramatically reduces the number of connection points that the designer must consider. For example, a two-input OR gate has 10 physical connections—four for the two differential inputs and six for the three levels of differential output. The model, however, has only four connections that the designer actually uses during schematic capture.

■ DESIGNING With SCFL

Using a more complex circuit structure can present problems to the logic designer accustomed to the simplicity of digital CMOS. The number of voltage levels, variations in delay and differential signals can be daunting. TriQuint tries to simplify the development of SCFL-based designs by hiding some of the complexity in the design software. Models of the macros are set up to appear like ordinary macros. Specific voltage levels are assigned automatically by the netlister, which also takes care of differential signals. Only a few simple rules differentiate the SCFL design process from other technologies.

The different levels of input signals (Figure 2) result in different propagation delays through the logic cells. For each level of current steering, the cell incurs another RC time constant. The gate-source capacitor must be charged before current can appear at load. This delay is roughly 30 picoseconds per logic level. Because the tree works as a delay line, however, all inputs can toggle equally fast.

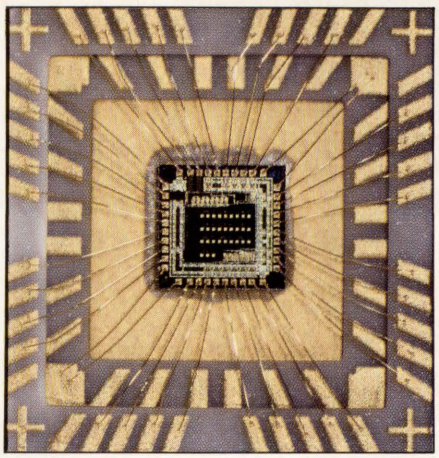

Figure 3. This GaAs IC, containing standard cells for verifying the performance of SCFL cells, operated at clock frequencies as high as 2.5 GHz.

The cells have all been analyzed with Spice, using delays extracted from the characterization. The lumped delays, such as those shown in the Table (for the circuit in Figure 2), are all the designer needs to know about the cells' internal performance. The delays are entered into a logic workstation for simulation. There is no concept of rising and falling signals with differential SCFL. Instead, the lumped propagation delay is measured in terms of delay from input signal zero crossings to output signal zero crossings.

Flip-flops and some logic functions have symmetrical construction, and so have equal delay for either signal sense. NOR gates are physically asymmetrical and exhibit a slight skew (which is higher at the lower signal levels). Reset on a flip-flop requires more design margin, because it also makes the cell asymmetrical. Fortunately, only two or three cells have the skew problem.

Because of this modeling technique, designers can capture schematics and simulate in the same manner as they would for silicon-based circuits. The designer represents differential signals as a single connection on the schematic, and then simulates the circuit on a workstation. The workstation uses a nominal wiring capacitance with each load.

Once the design is simulated, a net-list generation program expands each connection into two differential signals and maintains levels, as explained in more detail in the next section. Even though the connection work is done already, SCFL's idiosyncracies still require some special consideration. Designers should keep in mind the several levels of outputs as well as what level the router will route for critical paths.

The higher-level logic inputs have shorter propagation delays than the lower ones (see Table). When using SCFL AND, NAND, OR and NOR functions, therefore, the 1-level logic signals should be used for data with the most critical timing requirements. The third level of logic should be used when timing requirements are not stringent. The second level of logic should perform well for all but those signals requiring no skew and strict duty cycles.

In addition, signal propagation delay is directly proportional to wiring capacitance. To minimize it, the designer should try to keep all driven signals on the same logic level. Using two or three logic levels can, and usually does, result in more total capacitance. This happens because each must be routed separately, even though they are the same signals, but at different output levels.

Experience shows that about 90 percent of all cells are required to drive only one of the three levels. This is because output signals are usually tied to the same kinds of inputs, particularly in synchronous circuits. For example, a reset signal for a register file will connect to the same port on each register.

In combinational circuitry—for example, address decoding—signals may drive a variety of cells, and are therefore more likely to use two or three output logic levels. The designer can minimize routing capacitance, and reduce routing congestion, by keeping all inputs on a net at the same logic level. As seen in Figure 1, the inputs are marked with logical names (A, B and C) as well as numbers (1, 2 and 3), which represent the logical level. Connecting an output signal exclusively to 1-, 2- or 3-level inputs will yield better performance than using the three input levels randomly. Again, for 90 percent of circuits, this is irrelevant.

The software adds 25fF to account for wiring delays during simulation on Daisy and Mentor workstations. This represents average wiring capacitance to each driven input node. Versions of QLSI libraries for simulators not supported by TriQuint should also include this figure to prevent overly optimistic timing results. Designers should remember that results of simulation don't consider variations in propagation delay due to processing and temperature. They can assess the variation by mutlplying simulated times by 0.85 and by 1.5.

Power consumption is listed for each cell so the designer can determine the overall power consumption of the design. Worst-case power for the SCFL cells can vary from 0.5 to 1.3 times the nominal listed power for the cells, due to process and temperature variations. Power dissi

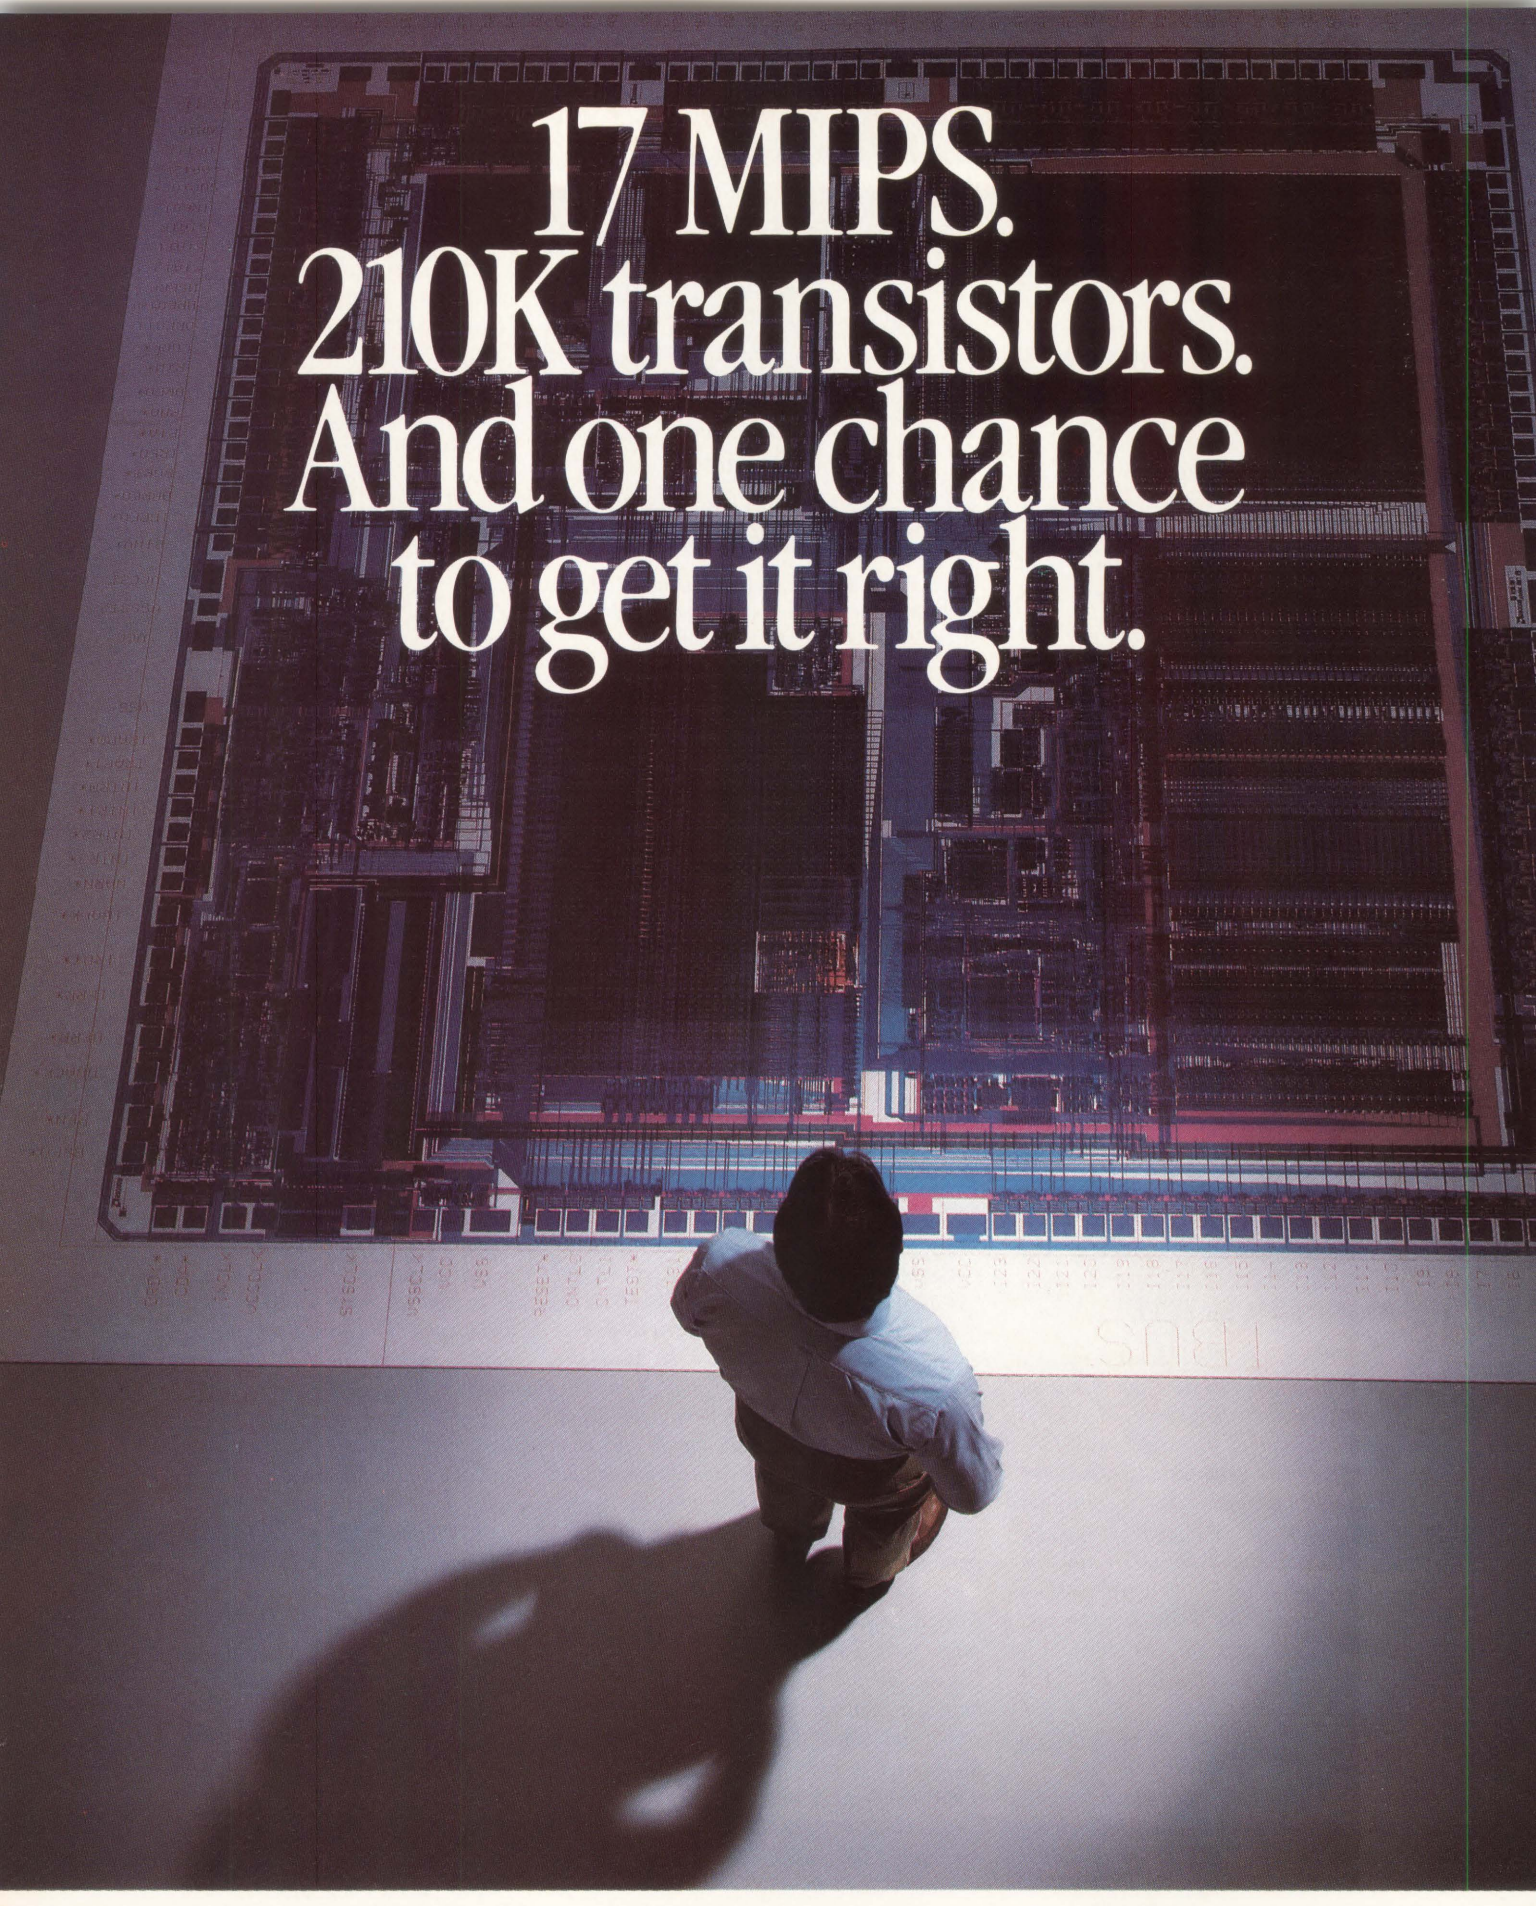
17 MIPS.
210K transistors.
And one chance
to get it right.

To build the world's highest performance RISC microprocessor, AMD turned to Mentor Graphics IC layout tools.

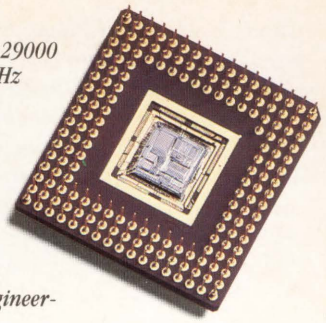

AMD's new RISC-based, 32-bit Am 29000 microprocessor operates at a 25 MHz clock rate with a 40 ns instruction cycle time. It can hit a peak execution rate of 25 MIPS, with a sustained performance of 17 MIPS. Because of its exceptional speed and price-performance ratio, the Am 29000 can be used in a wide range of applications—from embedded controller designs to engineering workstations.

It was a bold, ambitious project: build the fastest 32-bit μP in existence. One delivering a 5X to 7X performance improvement over existing 32-bit processors.

So Advanced Micro Devices turned to Mentor Graphics' Chip Station® to get the Am 29000 into silicon in a single iteration. Why? Because Chip Station provided the most advanced capabilities available, yet could also emulate AMD's existing tools.

Sharpen your competitive edge.

Look inside Chip Station and you'll find features like advanced traversal capabilities that make moves through the hierarchy as simple as point and click. And polygon-based editing that prevents problems like self-intersecting data. Also, programmable stroke recognition, which immediately converts cursor movements into commands. And now, the industry's fastest VLSI color plotting solution.

Complete compatibility with your existing environment.

Chip Station's programmable user interface shortens your learning curve by emulating the commands of your current system. At the same time, Mentor Graphics helps you adapt to the rich feature set of Chip Station and graduate to a truly advanced IC layout methodology.

Chip Station also brings you both the capacity and performance to deal with the coming generation of ULSI designs. Structured Chip Design lets you represent cells symbolically so the organization and management of large designs is greatly simplified. And REMEDI™ expands DRACULA II™ LVS checks to include graphic correlation of errors in both schematics and layouts.

The emerging standard.

Mentor Graphics' IC design and layout tools have already earned a solid reputation for productivity on large, complex VLSI projects. As a result, we're the world's fastest growing supplier of IC layout systems.

It's all part of a vision unique to Mentor Graphics, the leader in electronic design automation. Let us show you where this vision can take you.

Call us toll-free for an overview brochure and the number of your nearest sales office.

1-800-547-7390 (in Oregon call 284-7357).

DRACULA II is a trademark of Cadence.

Mentor Graphics®

Your ideas. Our experience.

pation is also directly proportional to the power supply voltage.

If designers follow these procedures, then they can be confident that SCFL circuits will perform as simulated. To demonstrate, TriQuint fabricated a simple frequency divider consisting of a 2-GHz input cell, divide-by-two, and an output cell (Figure 3). Characterization of the flip-flop showed that its should be able to operate at greater than 2 GHz. Although the simulation indicated operation at 2 GHz, a measure of the conservativeness of simulation was demonstrated by actual operation of the chip at 2.5 GHz.

■ SOFTWARE FOR ROUTING

The netlist generation program automatically resolves the use of multiple, differential signal levels. Schematic nets are doubled to represent the differential signals. TriQuint determines the logic levels by assigning a level attribute to each input on a logic cell. The software traces back from the inputs to find the proper output driving level.

After layout, the design software extracts the capacitance on each signal line for annotation to the simulation. All capacitance for an output signal, regardless of the number of signal levels, is lumped together.

TriQuint built some constraints into the router to help minimize routing capacitance. For example, its GaAs process has two levels of metal: one on the substrate (on a layer of silicon nitride), and another, called airbridge metal, raised above the substrate on posts. The first layer of metal has twice the capacitance as the second layer. This is primarily because air is a better dielectric than silicon nitride. The two routes of a differential signal will have the same amount of capacitance if they contain roughly the same amount of the two routing layers.

The software accomplishes this goal by constraining the orientation of each layer of metal to one direction, with metal "two" orthogonal to metal "one." In point-to-point routing, therefore, the differential signals may take different paths to a load, but the total horizontal and vertical distance—and hence capacitance—should be the same for both lines. In practice, the capacitance of the two lines are typically within 10 percent of each other.

This contraint also prevents corners in metal two. Wherever metal two bends or ends, a post (which is actually a via) must be placed to suspend it. Constraining the metal layers to orthogonal directions ensures that where metal two ends, it must connect to metal one to continue the signal path. If metal two could bend and travel both horizontally and vertically, then placing a post at the corner could create an unwanted connection with metal one underneath the corner.

As an alternative, it's possible to insert routines into the router to determine if a signal path is running under a potential bend in metal two. However, constraining the direction of the metal layers removes the complexity and leaves the insertion of airbridge suspension vias to a post-processing routine.

Power supply distribution also is an important layout consideration. Power supplies must be distributed evenly enough that voltage drops between the power pad and the cells is minimized. In addition, if CMOS or TTL I/O pads are used and if ZFL circuitry is also present, then an SCFL GaAs IC may require two or three power supply voltages on chip.

Triquint removes these concerns for the system designer by building a "power bus carrier" around the completed core to complete all power supply connections satisfactorily. The design is routed with the signal I/O pads placed in their designated locations. Space is appended to the I/O pads to allow for the power bus carrier. The router does not place routing lines in that space, except at locations marked for signal lines. After routing, TriQuint completes the pad ring, inserts the power bus carrier, and completes the internal power routes.

A router, presumably, could route the power supplies as well as the signals. Given the needs of distribution and multiple voltage levels, as well as the need to place power-ring posts all along the power bus carriers, TriQuint feels that it is much simpler to perform this function manually than to try to modify the router's code.

■ SCFL IN SYSTEMS

The use of SCFL GaAs chips can have a major impact on systems design—the technology can implement logic running at gigahertz frequencies. However, it requires a voltage supply between -4.5 and -5.5V, which is readily available if ECL circuits are used in the design. The standard-cell library has I/O cells that receive and drive ECL-, TTL- or CMOS-level signals. To interface with CMOS and TTL, the chip must have access to the +5V bus, while the interface to ECL uses the SCFL power supply.

As the use of GaAs in LSI-level applications grows, the need for higher pin count packages will also grow. TriQuint is extending the number of signal I/Os available in its MLC132 package, which can now support 64 I/O signals. This package conceivably could support over 100 signal I/Os, especially if slower signal transitions could be used.

Slowing down signal waveforms reduces ground bounce associated with changing output states. This technique, therefore, reduces the number of package pins that must be devoted to ground connections. In addition, slower signal transitions cause less crosstalk due to coupling. This allows signal lines to be placed closer together and increases the total number of signal lines.

TriQuint tests completed parts using testers from its parent organization, Tektronix Inc., and Cadic Inc. to verify the functionality of the designs. The chips are also tested at the wafer stage to verify that the process parameters are within specified limits. This test replaces ac testing which, at gigahertz frequencies, is expensive and time-consuming to set up. However, if required by the designer, special arrangements for ac testing are possible. ■

REFERENCES

IDDA, M. et al. 1984. "Analysis of High-speed GaAs Source-Coupled FET Logic Circuits," IEEE Transactions on Microwave Theory and Technology.

KATSU, S. et al. August, 1982. "A GaAs Monolithic Frequency Divider Using Source Coupled FET Logic," IEEE ELectron Device Letters.

LARSON, D. et al. 1986. "A GaAs Counter Family Designed Using Standard Cells," Proceedings of the Custom Integrated Circuits Conference.

SHIMANO, A. et al. February, 1981. "A 4GHz 25mW GaAs IC Using Source Coupled FET Logic," ISSCC Digest of Technical Papers.

VU, Thu T. et al. February, 1988. "The Performance of Source-Coupled FET Logic CIrcuits that Use GaAs MESFETs," IEEE Journal of Solid-State Circuits.

ABOUT THE AUTHORS

BILL DAVENPORT *has been a senior design engineer at TriQuint Semiconductor since 1985. He has also worked at Tektronix, where he was engaged in both system- and IC-level design. He holds a BSEE and MSEE from the University of Wyoming.*

DONALD LARSON *is the digital design manager in TriQuint's digital and linear products group. Before joining TriQuint in 1985, he worked at Tektronix, designing ICs in MOS, bipolar and GaAs technologies. He received his BSEE and MSEE from the University of California, Berkeley.*

Chip, Chip, Array!

For High Speed with Low Power.

AMCC has the chips worth cheering about. When you need the versatility of high speed with low power in a bipolar array, our Q5000 Series Logic Arrays are the answer. They're designed for logic applications requiring speed/power efficiency. And they deliver.

Today's hi-rel commercial and military semicustom applications need high performance and proven reliability. And, our Q5000 Series gives you both—without paying the power penalty.

Our newest bipolar series is comprised of five arrays. All feature 4 levels of speed/power programmable macros and over 600 MHz I/O capability. One comes with 1280 bits of configurable RAM.

Q5000 Series Key Features	
Equivalent Gate Delay:	210–545ps
Flip/Flop Frequency:	>600 MHz
Power Per Gate:	1mW
Speed/Power Product:	0.5pj
Equivalent Gates:	1300–5000
I/O Pads:	76–160
Operating Temperature Range:	−55°C to +125°C

AMCC Bipolar Logic Arrays have been designed with other flexible performance features in mind, too. Mixed ECL/TTL I/O compatibility. Your choice of packaging. Full military screening. AMCC's MacroMatrix® design tools. And, unrivaled customer support.

To talk with an applications engineer about your specific needs, in the U.S., **call toll free (800) 262-8830.** In Europe, call AMCC (U.K.) 44-256-468186. Or write, Applied MicroCircuits Corporation, 6195 Lusk Blvd., San Diego, CA 92121. (619) 450-9333.

A Better Bipolar Array is Here.

AMCC

CIRCLE NUMBER 7

Desktop-Configurable

CHANNELED GATE ARRAYS

AMR MOHSEN, ACTEL CORP., SUNNYVALE, CALIF.

■ DEVICES USE NOVEL INTERCONNECTS, MODULES, AND ARCHITECTURAL ELEMENTS ■

THE IDEAL ASIC would have an unlimited number of gates and offer instant turnaround. A new device family, the Actel ACT 1, has many gates and versatile interconnect resources. It provides the capabilities of a gate array, but the device's user configuration avoids the NRE charges and long turnaround times of gate arrays.

Several new features were developed to provide this combination of benefits, and several are especially significant. They are: the PLICE antifuse programming element optimized for ASIC applications; an innovative, channeled architecture that allows users to take advantage of virtually all of the device's logic modules via automatic layout tools; a logic module that maximizes user flexibility while efficiently using silicon real estate, and a diagnostic feature that permits observation of any node in the device while the device is operating in a system.

To make certain that the maximum number of gates are available to users and that the implementation is as easy as possible, Actel developed the device architecture and automatic place-and-route software in parallel. This approach made it unnecessary to create a device technology that would be difficult for the design automation system to support.

The parallel development program also helped make possible 100 percent automatic placement and routing at device utilization levels between 85 and 95 percent. The first devices to embody this technology are members of the ACT 1 gate array family, and include densities from 1,200 to 6,000 gate-array equivalent gates.

■ A COMPACT ANTIFUSE ALLOWS CUSTOMER CONFIGURABILITY

The PLICE antifuse is the heart of the new device family. It is a small, nonvolatile, two-terminal interconnect element that exhibits low On-resistance when activated. The area occupied by a PLICE antifuse is comparable to that of a via. These characteristics allow the PLICE antifuse to have the same wire-to-wire interconnect capabilities as masked gate arrays. The PLICE antifuse forms a conducting interconnect path, rather than an open circuit, when configured.

There is a crucial difference between the PLICE antifuse and the programming technologies based on EPROM or RAM cells. The latter exhibit too much On-resistance or take up too much space to provide the large number of interconnects needed for a high-density device.

The 2,000-gate ACT1020, for example, contains 186,000 antifuses. A conventional masked array of similar size contains about 90,000 potential metal-to-metal via sites. This multitude of antifuses can be routed like a conventional masked gate array, where the metal interconnects can be laid down as needed. PLICE antifuses can meet the needs of a desktop-configurable, channeled gate array because they can be packed as close as the metal pitch of the CMOS technology used to manufacture the devices.

Specifically, Actel's antifuse consists of a dielectric layer sandwiched between two conducting materials. When voltage is applied across the antifuse, a link is formed connecting its two

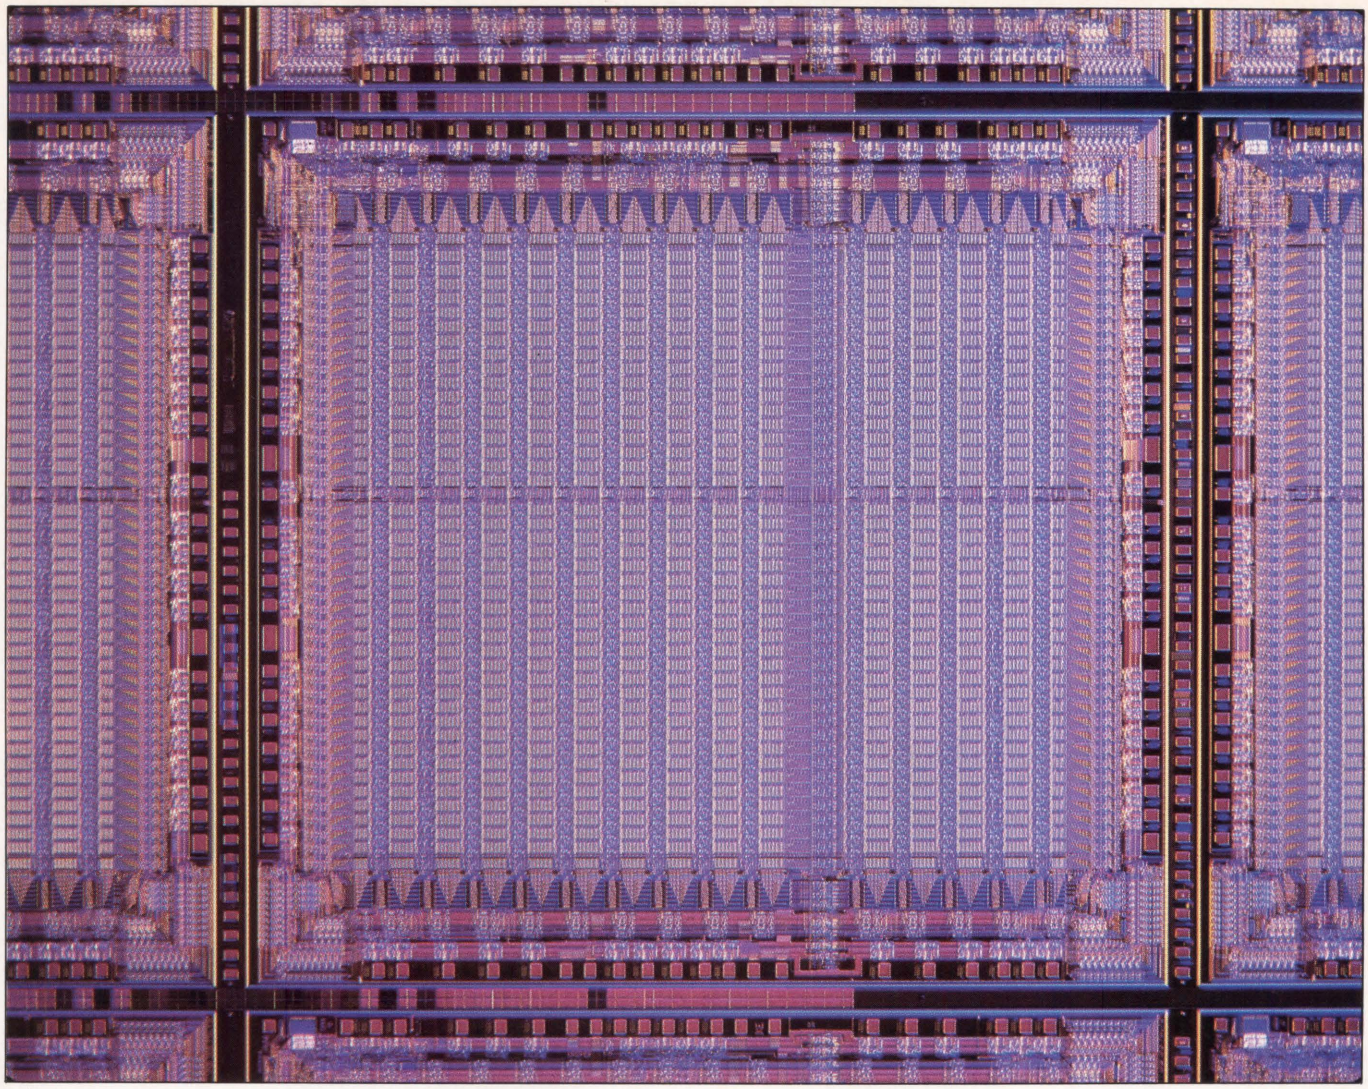

terminals. Most other antifuses reported to date have had at least one of the following drawbacks. They used complicated technologies, needed very high breakdown voltages and currents, were difficult to manufacture, or did not meet reliability requirements of today's ICs.

The PLICE antifuse overcomes these problems. It is manufactured using a typical twin-well, double-layer metal CMOS process (2 micron for the 200- and 2,000-gate devices; 1.2 micron for 3,000- and 6,000-gate densities). The antifuse exhibits an Off-resistance greater than lOOM ohm. Applying 18V across the antifuse configures it permanently to a bidirectional On state, with a resistance less than lK ohm. The configuring pulse's current is less than 10 mA and lasts fewer than 10 ms.

It's important to note that in an array based on antifuse elements, only two to three percent of the total available antifuse elements must be programmed to implement a design. This keeps overall device programming time to a minimum.

In terms of performance, the PLICE antifuse's low On-resistance translates into short wiring delays. This allows configured ACT devices to operate up to 40-MHz in the system. And, because the antifuse is permanent and nonvolatile, its unnecessary to load start-up data. Soft error problems are eliminated as well, because it is nonvolatile. In addition, the antifuse provides higher radiation tolerance than floating-gate storage elements.

■ A Channeled Gate Array Architecture

The overall architecture of an ACT device resembles that of a channeled gate array, in which rows of logic modules alternate with routing channels. The details of the interconnect scheme differ from that of a conventional gate array, however.

Figure 1 illustrates the horizontal and vertical wiring segments used to interconnect the logic modules. A PLICE antifuse is located at each intersection of a horizontal and vertical wiring segment, so that a configured antifuse connects the two wiring segments. Antifuses located between adjacent horizontal wiring segments make it possible to create longer segments by connecting the adjacent segments. The architecture is designed so that usually there are no more than two antifuses in a path, up to a maximum of four antifuses in any given path.

To achieve vertical interconnection tracks, the logic module inputs and outputs extend vertically on dedicated wiring segments. The outputs cross four routing channels. The inputs cross one channel up and onedown. The in-

SOFTWARE PERMITS 100% AUTOMATIC PLACEMENT AND ROUTING

Our ASICs

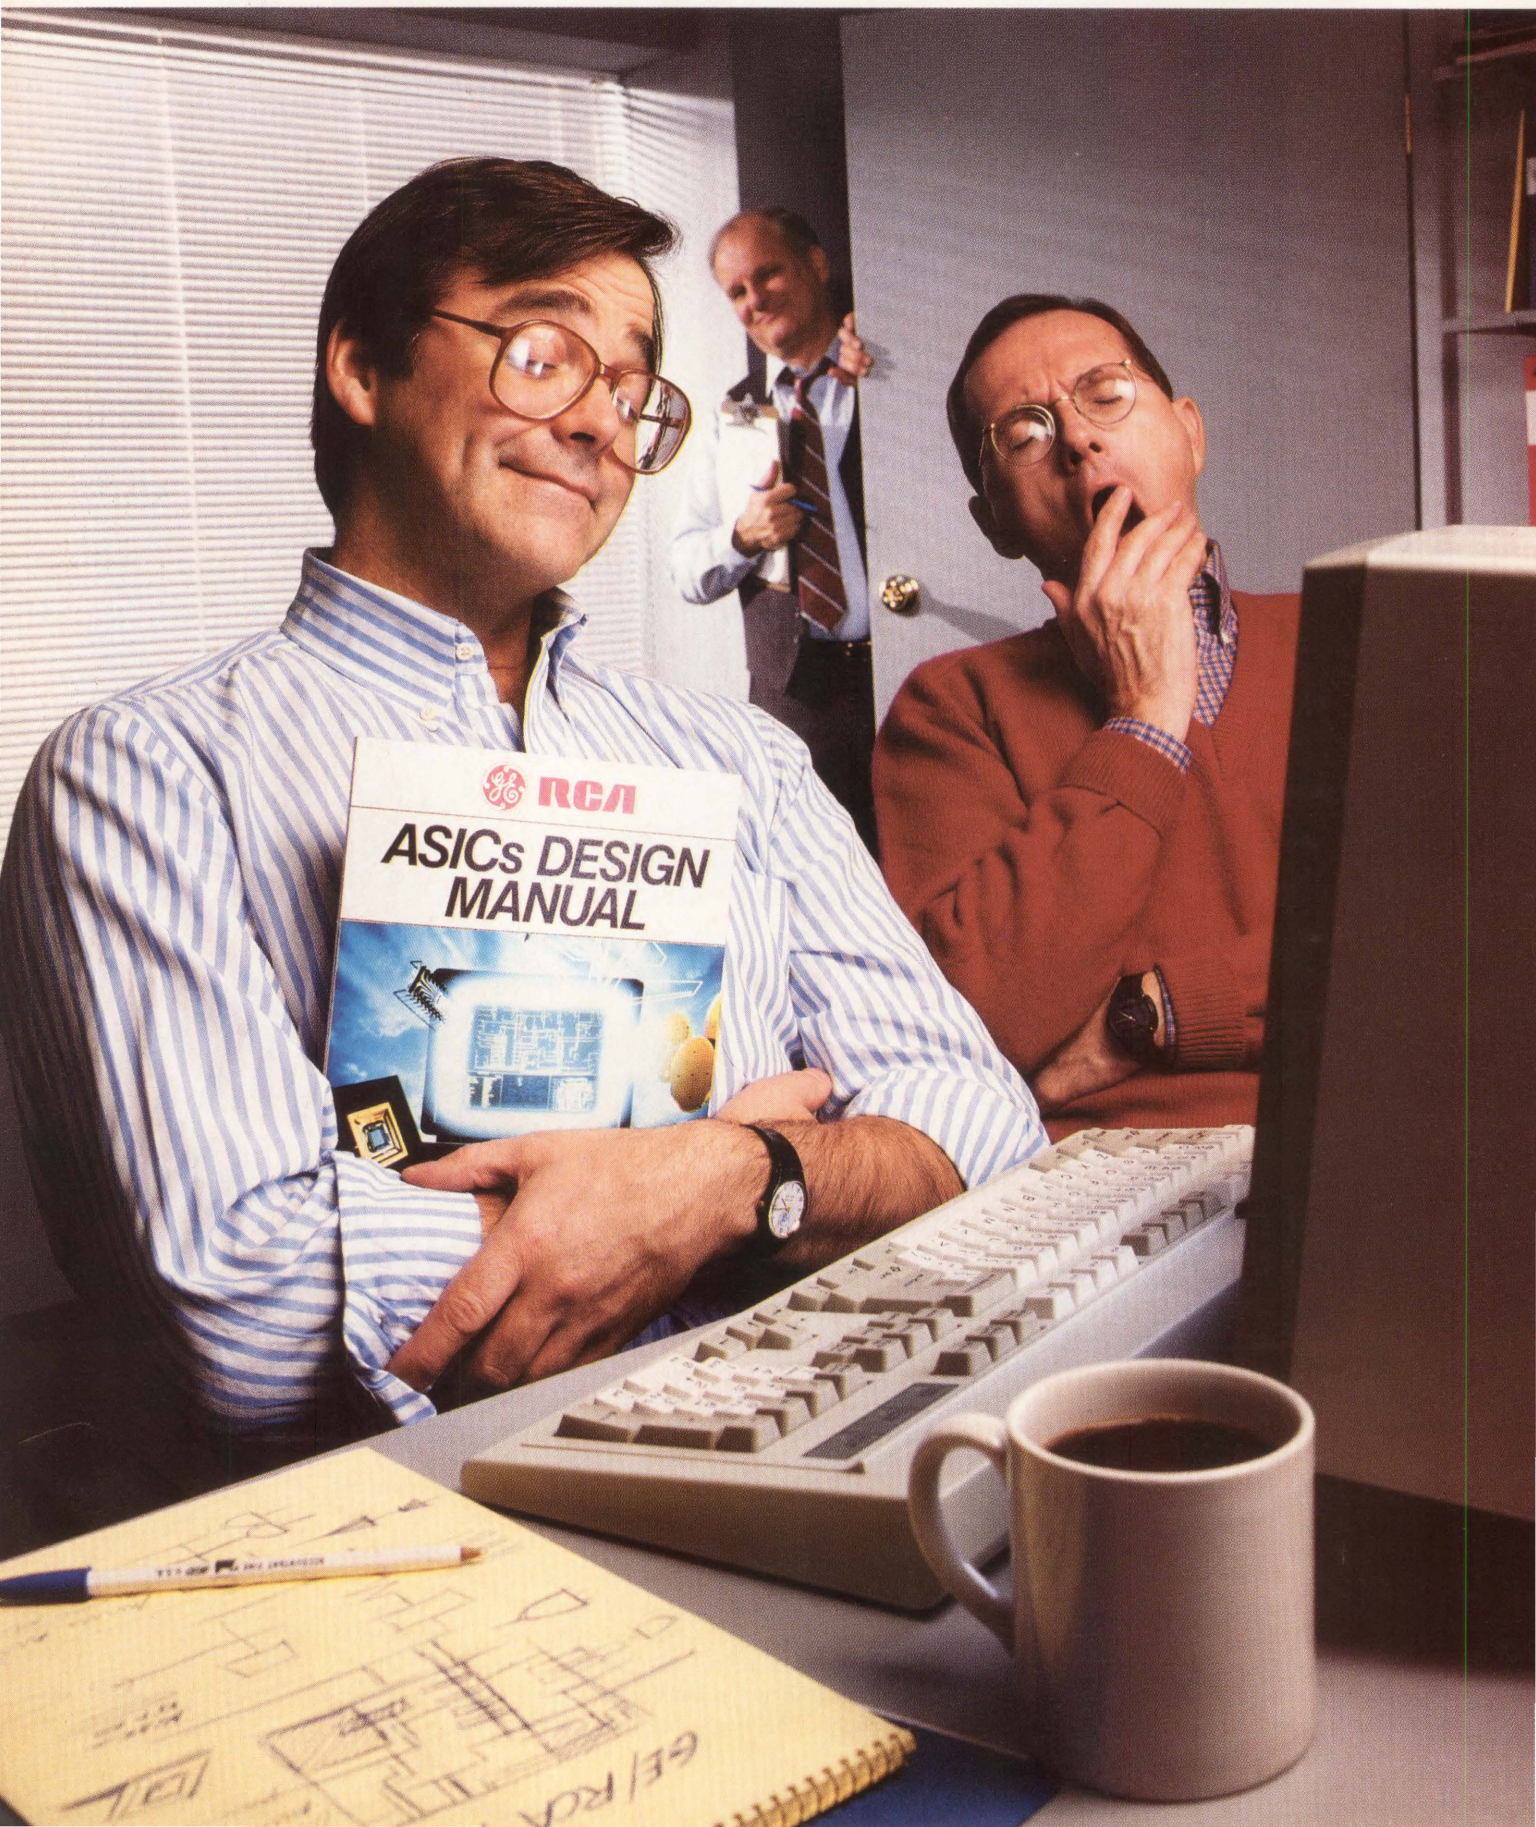

are boring.

They're easy to design. They're ready on time. And first-time success is virtually 100%.

You've heard all about the excitement of ASICs.

They improve performance, lower costs and make many new designs possible.

But, unfortunately, you've probably also heard about one big potential problem: while many ASICs pass the tests specified by the designer, they don't always work in the real world. And that causes excitement you can do without.

How to get first-time success.

It starts with our Design Simulation Software. It's been rated the best in the industry by the people who should know—designers who have used it. Within three days, you can be up to speed, working at any of the major workstations in the industry, creating and revising your ASIC with ease.

The standard cell advantage.

You'll really appreciate the power of our standard cells, which allow you to integrate a whole system, including macros, memories, logic and peripherals, onto a single chip.

We have cells with effective gate length as small as 1.5μ (.9μ coming soon). And double-level metal for higher-density chips that can handle higher clock speeds.

You can choose from a wide range of Supercells, including the leading-edge RS20C51 core micro, RAMs, analog functions, bit-slice processors, HC/HCT logic, Advanced CMOS Logic, and high-voltage cells.

If they aren't enough, we can even generate Supercells to your specs.

And we're also in the forefront of silicon compiler technology. So we can offer you the ability to create designs that are heavily BUS-structured, with your ROMs, RAMs, PLAs and ALUs compiled right into the design.

We also bring you the resources of some very powerful partners, thanks to our alternate-source agreements with VLSI on standard cells; WSI on macrocells and EPROMs; and a joint-development agreement with Siemens and Toshiba on the Advancell® library of small-geometry cells.

Gate arrays, too.

If gate arrays are better for your design, you'll be able to choose from our full line up to 50,000 gates, with effective gate length as small as 1.2μ and sub 1 ns gate delays.

These gate arrays use "continuous gate" technology for up to 75% utilization. They are an alternate source to VLSI Technology arrays.

We also alternate source the LSI Logic 5000 series.

And we have a unique capability in high-rel ASICs, including SOS. Our outstanding production facilities here in the U.S. produce high-quality ASICs in high volume at very low costs.

It almost sounds exciting for something so boring, doesn't it?

For more information, call toll-free today 800-443-7364, ext. 25. Or contact your local GE Solid State sales office or distributor.

In Europe, call: Brussels, (02) 246-21-11; Paris, (1) 39-46-57-99; London, (276) 68-59-11; Milano, (2) 82-291; Munich, (089) 63813-0; Stockholm (08) 793-9500.

USA

GE Solid State

GE/RCA/Intersil Semiconductors
Three great brands. One leading-edge company.

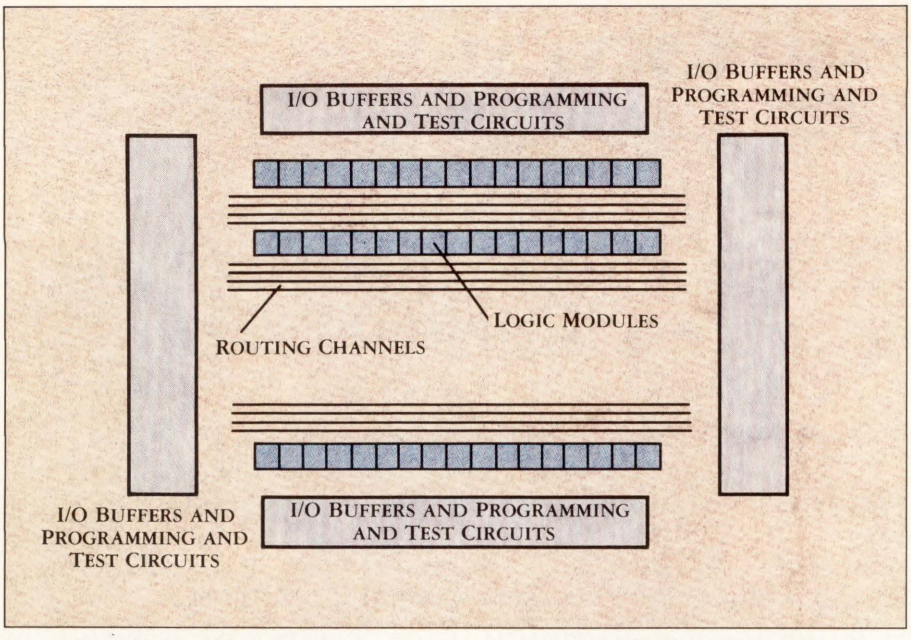

Figure 1. Vertical logic module inputs and outputs connect to other modules or I/O buffers by linking with horizontal wiring tracks. An antifuse at each horizontal and vertical junction allows the segments to be connected electrically.

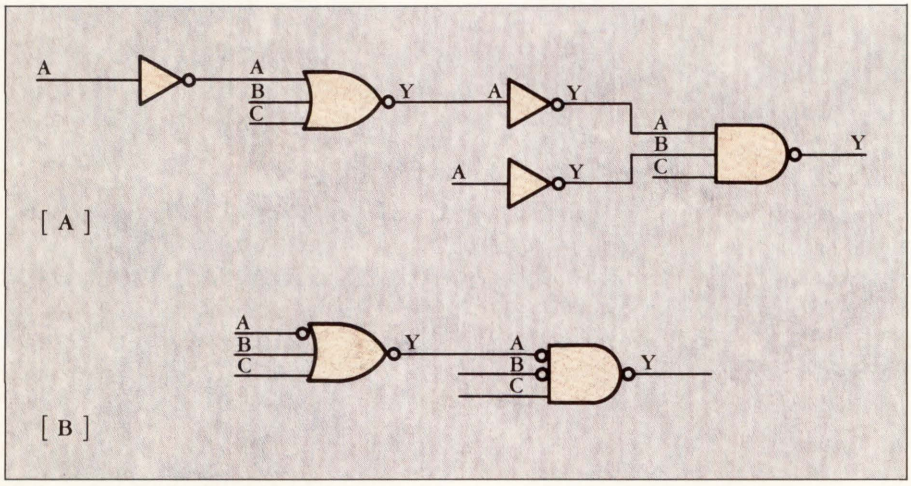

Figure 2. With Actel's negation bubbles, designers pay no speed penalty for inversions, eliminating inverters common in conventional masked gate arrays.

puts are arranged for maximum routing versatility, so that half the inputs extend to the routing channels above the module and half to the channels below. The automatic place-and-route software can choose entry from above or below for the best gate use.

There are also independent vertical segments that can be used for general interconnection along with the independent horizontal segments. A horizontal routing channel contains 25 routing tracks, including general-purpose segments, and power, ground, and a dedicated clock track. Although more clock signals can be routed as needed, the dedicated clock track allows a large fanout with minimal clock skew.

The ACT architecture must accommodate PLICE antifuse programming, as well as interconnect needs. So, along with the basic horizontal and vertical routing segments and I/O buffers, the architecture includes control logic at the periphery of the array. This logic uses the wiring segments along with pass transistors that connect adjacent segments to selectively apply the programming voltage to appropriate antifuses. The pass transistors use shared control lines to minimize wiring. Paths to the logic modules are arranged so they furnish full access to the modules.

An antifuse is programmed by shifting an address into the device and applying a specific programming sequence. The addresses are shifted serially to achieve the desired interconnection pattern. The order is application specific, and is handled automatically by the design software.

■ Versatile Logic Modules Meet A Variety Of Requirements

One key aspect of the ACT device is its configurable logic modules. All logic modules are identical, and have eight possible inputs and one output. The modules, which are similar in capability to conventional gate array hard macros, were designed for the most efficient implemention of both combinatorial and sequential logic. They also take best advantage of the architecture's routing resources. For the first time in a programmable device, no dedicated flip-flops are needed. The user can designate the flip-flops as needed to implement a specific circuit.

Logic modules can implement more than 150 different logic functions. The Actel hard macro library contains functions ranging from simple Nand gates and Exclusive-Or gates, to multilevel logic such as And-Or-Invert, And-Exclusive-Or and Multiplexers. Flip-flops include T, D, master/slave and JK, along with a D flip-flop that has a multiplexed input. Transparent D latches are also available.

The functions mentioned so far consist of hard macros. Because the ACT logic module has about the same complexity as a conventional gate array hard macro, users design with the modules in ways similar to the familiar gate array macro libraries. There is no need to map logic into a more complex module. As with conventional gate array macro libraries, the ACT library includes soft macros such as decoders, comparators, counters, registers, and some commonly used TTL functions.

Because the module includes a versatile method for inverting signals, a variety of "negation bubble" arrangements are possible on gates. In contrast to conventional gate arrays, the input inversions can be implemented as efficiently as noninverting inputs. And the inversions entail no additional propagation delays. Individual inverters and their associated interconnects are rarely necessary (Figure 2).

I/O modules in the ACT devices are also versatile. Each I/O module can be configured as an input, output, or bi-directional by programming the appropriate antifuses. The I/O buffers are TTL compatible, with 4-mA drive capability. In the overall chip architecture, the I/O modules are arranged around all four sides of the logic array.

■ Logic Capacities Compare With Conventional Gate Arrays

A goal of the product architecture was to provide high gate densities over a broad range of applications. A benchmark measurement scheme proved to be a good way to evaluate gate efficiency across application types.

Actel chose four types of logic functions

Cypress Semiconductor introduces the RISC Motorola didn't consider.

The RISC that lets you build faster computers faster.

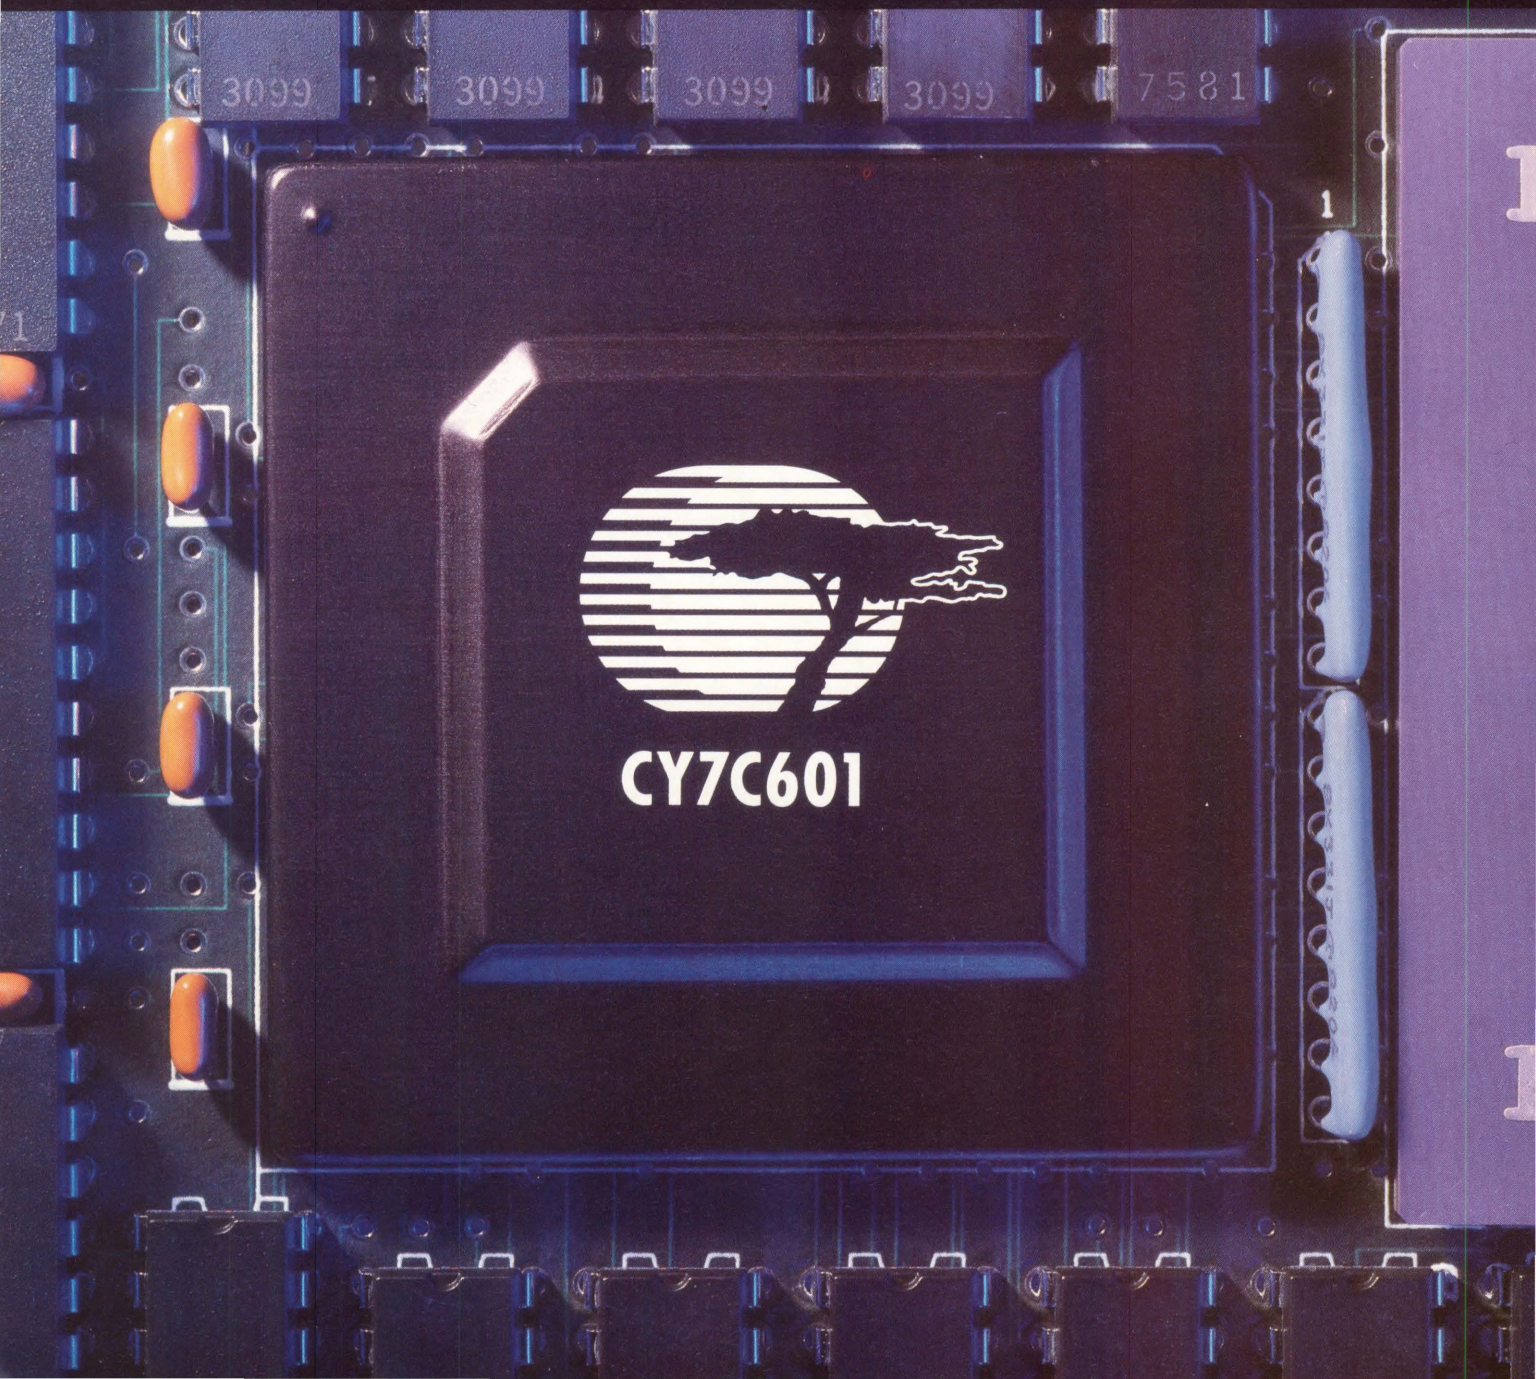

Introducing the RISC 7C600 Family: 20 VAX MIPS, SPARC architecture, and development systems today.

Before you design the next generation of highest performance computers, meet the RISC 7C601 microprocessor from Cypress Semiconductor.

You'll build faster systems because this 20 MIPS RISC chip is available today, running at 33 MHz. It outperforms all others using SPARC™ (Scalable Processor ARChitecture), the fastest RISC architecture, implemented in our fastest 0.8 micron CMOS technology for outstanding performance and cool low power.

A complete chip set.

The fastest microprocessor doesn't stand alone. Besides the RISC 7C601 Integer Unit (IU), you can incorporate our CY7C608 Floating Point Controller (FPC) to interface with a standard floating point unit to perform high-speed floating point arithmetic concurrent with the IU.

Although the IU can function on its own with high speed local memory in a dedicated controller application, for most computer applications our high performance CY7C603 Memory Management Unit (MMU) coupled with the IU and FPC gives you the fastest access to both cache and main memory through the 32-bit address bus and 32-bit data/instruction bus. It also supports the SPARC Reference MMU architecture giving you compatibility with standard UNIX operating systems.

Our CY7C153 32Kx8 Cache RAMs and the CY7C181 Cache TAG RAM maximize your throughput by providing a cache selection capable of running at full speed with a 33 MHz IU.

We also deliver the highest performance SRAM, PROM, Logic, and PLD parts. So you can make the most of those MIPS.

Develop your systems quickly.

You'll build systems faster because the 7C601 is based on SPARC. You have a choice of powerful development systems already running on the target architecture. Plus there's a wide range of UNIX-based languages, tools, and utilities that already run on SPARC. You have more proven development tools, making this RISC easy to integrate into your design.

Design around a proven architecture.

The 7C601's SPARC architecture has already been embraced by companies like AT&T, Sun Microsystems, Unisys, and Xerox. And SPARC systems are already on the market, like the Sun-4. In fact, AT&T has selected SPARC for the first UNIX Application Binary Interface (ABI), which will allow all SPARC-based computers to run the same off-the-shelf UNIX applications. With our RISC 7C600 family you're designing around an accepted architecture with unlimited potential.

Make the fastest decision.

Our free brochure, RISC Factors, will give you more information about the factors affecting RISC microprocessors. Because deciding which microprocessor to use is not a snap decision, but it can be a fast one.

Call the RISC Factors Hotline now for your free copy: 1-800-952-6300. Ask for Department C105.*

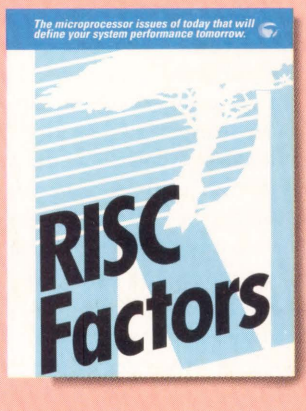

CYPRESS SEMICONDUCTOR

*1-800-387-7599 In Canada, (32) 2-672-2220 In Europe. Cypress Semiconductor, 3901 North First Street, San Jose, CA 95134, Phone: (408) 943-2666, Telex 821032 CYPRESS SNJUD, TWX 910-997-0753. ©1988 Cypress Semiconductor.

UNIX is a registered trademark of AT&T. SPARC is a trademark of Sun Microsystems.

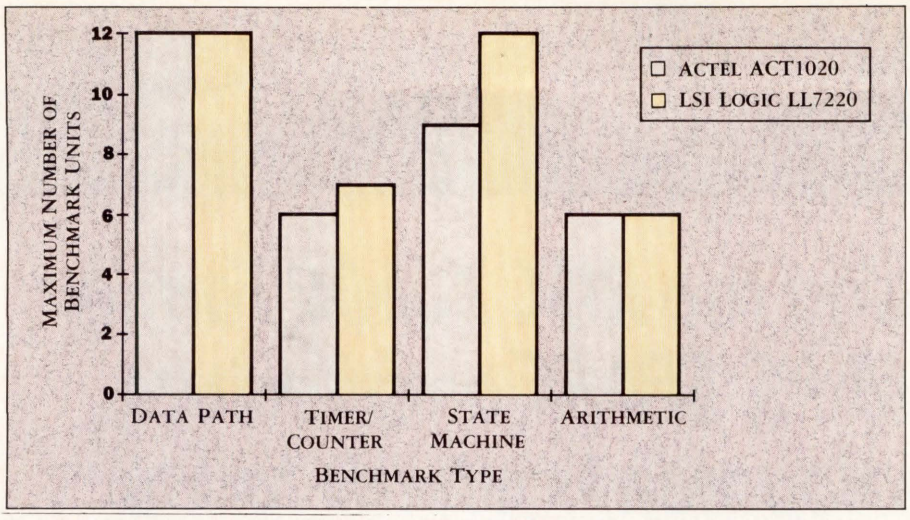

Figure 3. Device benchmarks indicate that the Actel 2,000-gate ACT1020 provides effective gate counts comparable to conventional arrays over a range of application types.

CRITICAL PATHS

RANK	TOTAL	START PIN	FIRST NET	END NET	END PIN
0	24.7	U1:CLK	CQ0	LF0	LDF0:A
1	23.6	U1:CLK	CL0	LDF	LDFF:A
2	23.0	U1:CLK	CQ3	LD3	LDF3:A
3	18.1	U1:CLK	CQ0	LD4	LDF4:A
4					
5					
6					
7					
8					
9					
10					

Figure 4. The ability to see a design's critical paths facilitates designers in optimizing circuits. The Actel timing analyzer provides a report of the slowest timing paths and a ranking of the paths according to speed.

as benchmark circuits to measure ACT devices against application types, and to compare with conventional gate arrays. These four types were a data-path circuit, a counter/timer, a state machine function, and an arithmetic function. Mapping any one of these functions into an ACT 1 device, as many times as the device can accommodate, shows how the ACT 1 compares with a conventional gate array for that same type of logic function.

For example, when this approach is used to compare the ACT1020 against the industry-standard 2,200-gate channeled gate array, it's clear they provide equal capacities for the datapath and arithmetic circuits. However, the conventional gate array rates slightly higher for the timer-/counter and state machine circuits.

Over the range of logic types, the ACT device provides comparable capacities, as illustrated in Figure 3. For example, 12 modules of a 1-of-8 decoder (74138) take only 2.2 percent of the ACT1020's 2,000 gates, while 16 modules of an 8-bit serial-in, parallel-out shift register uses 2.9 percent.

■ Speed Boosted By Short Critical Paths

Speed is also a measure of ASIC performance. The typical flip-flop toggle rate is 70 MHz for the 2-micron ACT devices, which makes possible a system-level speed as high as 40 MHz. Speed also can be evaluated by comparing one ASIC's performance against another's when implementing a benchmark circuit.

In terms of worst-case delay for a 16-bit counter, a conventional gate array's critical path is 11 nets deep, compared with only six nets for the ACT1020. The ACT1020's primary advantage in this case lies in its versatile ability to negate signals, so that extra delays for inverters are not added.

The ACT logic module also allows it to implement multi-level combinatorial logic and multiplexed flip-flops.

When derated for worst-case process, voltage, and temperature effects, the conventional gate array achieves a propagation delay of 54 nanoseconds, compared to 64 nanoseconds for the ACT1020. Despite the wire delays intrinsic to a field-configurable device, the ACT1020 delivers performance close to that of a conventional gate array.

■ Desktop Configuration

The design process for an ACT device is similar to that for a conventional gate array, except that the process is carried out entirely in the customer's facility. Designing with an ACT device begins with schematic capture followed by optional simulation. The designer next runs an automatic place-and-route program, followed by timing analysis. The device is then configured using Actel-supplied Activator hardware, and the design is verified with its diagnostic features. All the design procedures can be carried out on an 80386-based personal computer.

The programs for automatic place and route, timing analysis, and debug are the most interesting aspects of the design and implementation tools. The place-and-route procedure has three phases: constructive placement similar to that of conventional gate arrays; placement optimization using a simulated-annealing-type algorithm, and a unique and detailed routing procedure.

The designer can specify nets as critical paths so that place-and-routing software optimizes these paths. This sequence is completely automatic, and eliminates weeks of tedious manual routing. Placement, routing and programming for Actel's 1,200-gate, desktop-configurable, channeled gate array, for example, takes less than 45 minutes. Conventional masked gate array turnaround is four weeks. More important, ACT devices can be implemented without the cost of NRE charges. If the user must modify the design, then another device can be configured just as quickly and economically.

The timing analyzer, dubbed the Actel Timer, helps designers determine and optimize path delays within a design. The timer provides delay data for the design's slowest paths, and for any other paths specified by the designer, or it can analyze all possible paths at once.

Most important, the timer analyzes a design's critical paths and provides final AC performance specifications. It uses post-routing back-annotation for accurate delay analysis. The timer examines signal paths that start at one set of macro pins and end at another set.

WITH THE WRONG ASIC DESIGN TOOLS, YOU COULD BLOW THE WHOLE PROJECT.

WITH VLSI'S ASIC TOOLS, YOU CAN

VLSI Technology's software is the fastest, surest way to create successful ASIC chips.

Our tools accelerate each and every step of the design process.

And our years of experience insure the success of your finished ASIC chip.

EVERYTHING'S FASTER WHEN YOU HAVE AN ASSISTANT.

Our Design Assistant™ tool partitions your chip and estimates chip size, power consump-

BLOW RIGHT BY THE COMPETITION.

tion, and packaging possibilities. In short, it tells you the best silicon solution to your problem.

Simply enter your design in block diagrams and global interconnect forms. Design Assistant shows you ways to implement it. Gate array, standard cell or cell-based.

You can run dozens of "what if" configurations in a few hours. In no time, you'll have the crucial information you used to wait forever for vendors to supply.

STREAMLINED DESIGN, COMPLIMENTS OF OUR COMPILERS.

Our logic and memory compilers create multipliers, PLAs, or super fast SRAM memory blocks as fast as 8ns.

And we've got a couple of "expert" compilers for cell-based or gate array designs that think exactly the way you do.

When you design a datapath, you think of it as a linear schematic, right?

Well, our Datapath Compiler just happens to use schematics as input. Complex multi-bit datapaths practically pop right out of it.

When you design state machine and other logic blocks, you think in terms of equations. Our State Machine Compiler does, too.

And it even optimizes your equations.

GET IT TOGETHER. FAST.

Ready to place and route your cell-based design? Just call up our Chip Compiler.

It provides floor-planning and auto-routing. It works with our compilers and standard cell libraries to produce an optimized layout.

It even pours standard cells into the gaps between blocks to make sure you use the minimum chip area required.

How fast does it work? Glad you asked.

A company came to us with a layout that took them three months. We did it in two days.

TO BE FAST, YOU HAVE TO BE FLEXIBLE.

Only our Portable Library lets you choose gate array or cell-based, 1.5μ or 2μ, *after* you've completed your logic design.

It can allow you to do that because your library always remains stable. No matter what process you use.

Process obsolescence is now obsolete.

IF YOU NEED TO GET THERE FAST, TAKE THE EXPRESS.

You can buy VLSI's tools in six configurations. From our quick, easy-to-use Logic Express™ to the powerful Design Express.™

If you'd like to find out how quickly you can design successful ASIC chips, give us a call at (800) 872-6753.

Because when you're driving to market this fast, it's good to have insurance.

CONCEPT EXPRESS™:

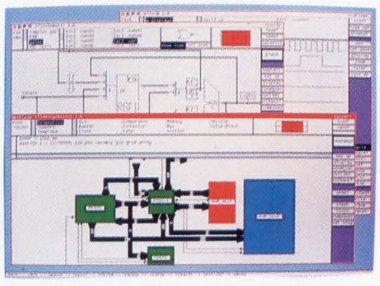

The Concept Express Design System's highly productive logic tools and silicon compilers were used to develop this very-large-scale ASIC. It incorporates a 2901 datapath, RAM, ROM, and over 3,400 gates of random logic.

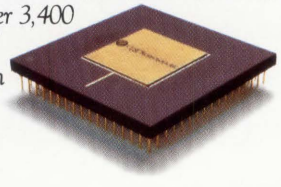

DESIGN EXPRESS™:

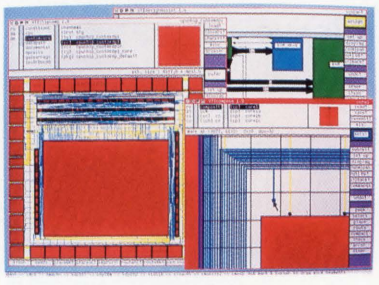

This highly-integrated design combines control logic, a register file, a refresh counter, and five peripheral chips onto a die size of 275x315 mils. The logic design, layout, and verification were completed in only 12 weeks.

SILICON EXPRESS™:

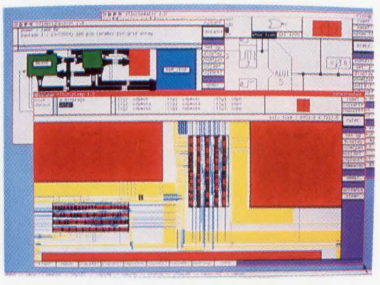

This design integrates all the peripheral chips for an AT computer with six megacells and control logic. Using the Silicon Express Design System, logic and physical designs like these can be implemented in under two man months.

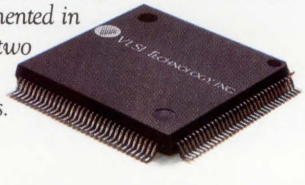

VLSI TECHNOLOGY, INC.

ACT Array Characteristics				
	ACT010 (1988)	ACT1020 (1988)	ACT1130 (1989)	ACT1260 (1989)
Number of Gates	1200	2000	3000	6000
CMOS Process	2μm	2μm	1.2μm	1.2μm
Number of Logic Modules	295	546	823	1400
Maximum Number of J-K Flip-Flops	147	273	411	1400
Maximum Number of Latches	295	546	823	1400
Number of PLICE Antifuses	112,000	186,000	400,000	666,000
Number of Antifuses per Gate	93	93	133	111
Number of I/O's	57	69	112	160

For example, the timer can automatically begin with a starting set consisting of all the clock or gate pins of the design's flip-flops and latches. The corresponding ending set contains all the gated inputs (D, E, multiplexer pins, A, B, and S) of the flip-flops and latches. These are good sets to begin with because most critical paths start and end at clocked (gated) storage elements in synchronous designs. The timer assists in optimizing the design by determining the slowest paths through the circuit, and the pins on which those paths end. This information is displayed as shown in Figure 4.

■ A Completely Observable Circuit

The design verification process can begin when the timing has been analyzed and the ACT device configured. As with other design verification procedures, the designer applies stimuli to the device and analyzes the response. There is a significant difference between the debug process for ACT devices and other ASICs. The ACT devices allow the designer not only to observe the response from I/O pins but to see activity on all internal nodes.

The ability to observe the activity of internal nodes relies on the Actionprobe diagnostics feature.

Under the designer's control, the Actionprobe diagnostics can connect electrically any two points in the device to external pins. This capability gives designers access to any programmed node within the configurable gate array. It permits 100 percent observability of all circuit activity that is normally hidden and unobservable with conventional gate arrays.

Each ACT device contains two such diagnostic circuits, so any two nodes can be observed at a time. The Debugger software assists in addressing the nodes of interest, and any subsequent activity on the addressed nodes appears at two predefined pins. (If users don't need the diagnostic features, these two pins can be used for I/O.) Simply by inputting "Probe A = XX" and "Probe B = YY" commands, where XX and YY are user-assigned net names, signals from internal nodes XX and YY will appear on the two respective pins.

These ports into the internal workings of the device give designers a debug facility similar to that of a logic analyzer on a board-level circuit. This feature is unique to Actel and not available with other conventional mask-programmed gate arrays or programmable devices.

In addition to helping a designer debug a device using artificial stimulus patterns, the Actionprobes assist when the device is operating in an actual system environment. Even when the device is working at full speed, any two internal nodes can be observed in realtime. The Action probe feature shortens design verification and in-system debug, but also reduces the burden of test vector generation by substituting millions of system-generated stimuli.

Other built-in testability features allow testing of an unconfigured ACT device for correct operation of logic modules, I/O buffers, addressing and decoding circuits, vertical and horizontal tracks, and all programming circuits. The integrity of all PLICE antifuses is also checked. This is done even though the one-time-programmable ACT devices are thoroughly tested at the factory, and PLICE antifuse programming is verified in the Activator as programming proceeds.

■ Reliability Comparable to Normal CMOS Levels

Under environmental extremes that accelerate failure mechanisms, the PLICE antifuse in the Off state has shown that, under normal operating conditions, it won't fail in more than 40 years. For the On state, the PLICE antifuse remains conductive under tests and exhibits no intrinsic failure mechanism. The combined PLICE antifuse's contribution to the gate array's hard failure rate is less than 1 FIT. (1 FIT is equivalent to a failure rate of 0.0001%/1,000 hours.)

Product reliability depends not only on circuit elements such as the PLICE antifuse, but also on other features of the manufacturing and assembly process. Product reliability studies indicate that ACT devices attain reliability levels normally encountered with CMOS circuits.

Initial product reliability data on an 89-unit sample burned-in for more than 119,000 device hours indicates no failure or change in AC characteristics. This implys a failure rate of less than 100 FITs. Overall, the ACT devices exhibit about the same reliability as any other CMOS parts.

Channeled gate arrays that can be configured by the customer promise to change the way designers think about ASICs. The ability to implement complex logic functions quickly in the field makes it possible to meet time-to-market demands, reduce the cost of prototyping and design verification, and minimize design cycle risks. ■

About the Authors

DR. AMR MOHSEN, *founder, president and chief executive officer, has more than 20 years of experience in semiconductor processes and device development. Before founding Actel, he was a senior engineering manager in the Technology Development Division at Intel. During his eight year tenure, he was responsible for CMOS DRAM products and CHMOS-III technology developments. He also worked on charge-coupled device development at Bell Laboratories and served as a consultant. Dr. Mohsen holds a Ph. D. in electrical engineering and applied physics from the California Institute of Technology. He has authored more than 40 articles relating to semiconductors and is responsible for inventions covered by 13 patents.*

A Concurrent Hardware And Software Design Environment

ANDREW KERN AND AL BLAZEVICIUS, INTERACT CORP., NEW YORK, N.Y.

According to Dataquest, 70 percent of all 16- and 32-bit microprocessor-based systems in design are embedded systems. Methods for designing these systems, however, have not evolved significantly since the era of 8-bit microprocessors. Hardware and software are developed separately, a hardware prototype is built and tested, instructions are run on the prototype, and then the system is debugged with an in-circuit emulator. If the design works the first time, which rarely happens, it is often rushed into production without undergoing optimization or thorough testing.

The integration of software and hardware at the prototype stage is becoming unacceptable. System developers can't afford to let the software languish untested until the hardware prototypes are ready. Increasing competition requires optimizing both the software and hardware to assure the best performance possible from any system. To address these problems, designers need a common development system for hardware and software.

System Software Testing Can't Wait For Actual Hardware

The System Design Environment (SDE) from INTERACT Corp., a joint venture of LSI Logic Inc. and Advanced Computer Techniques Corp., combines software and hardware engineering tools for the concurrent design of software and hardware—without building prototypes. Using SDE, developers of realtime embedded systems can write and simulate software programs on a Sun-3 or Sun-4 workstation from Sun Microsystems. In addition, they can design and simulate microprocessor-based hardware on the same workstation. Finally, the code written for the hardware-specific, embedded-microprocessor system can be executed and debugged on a hardware simulation of its intended target. No prototypes need be built.

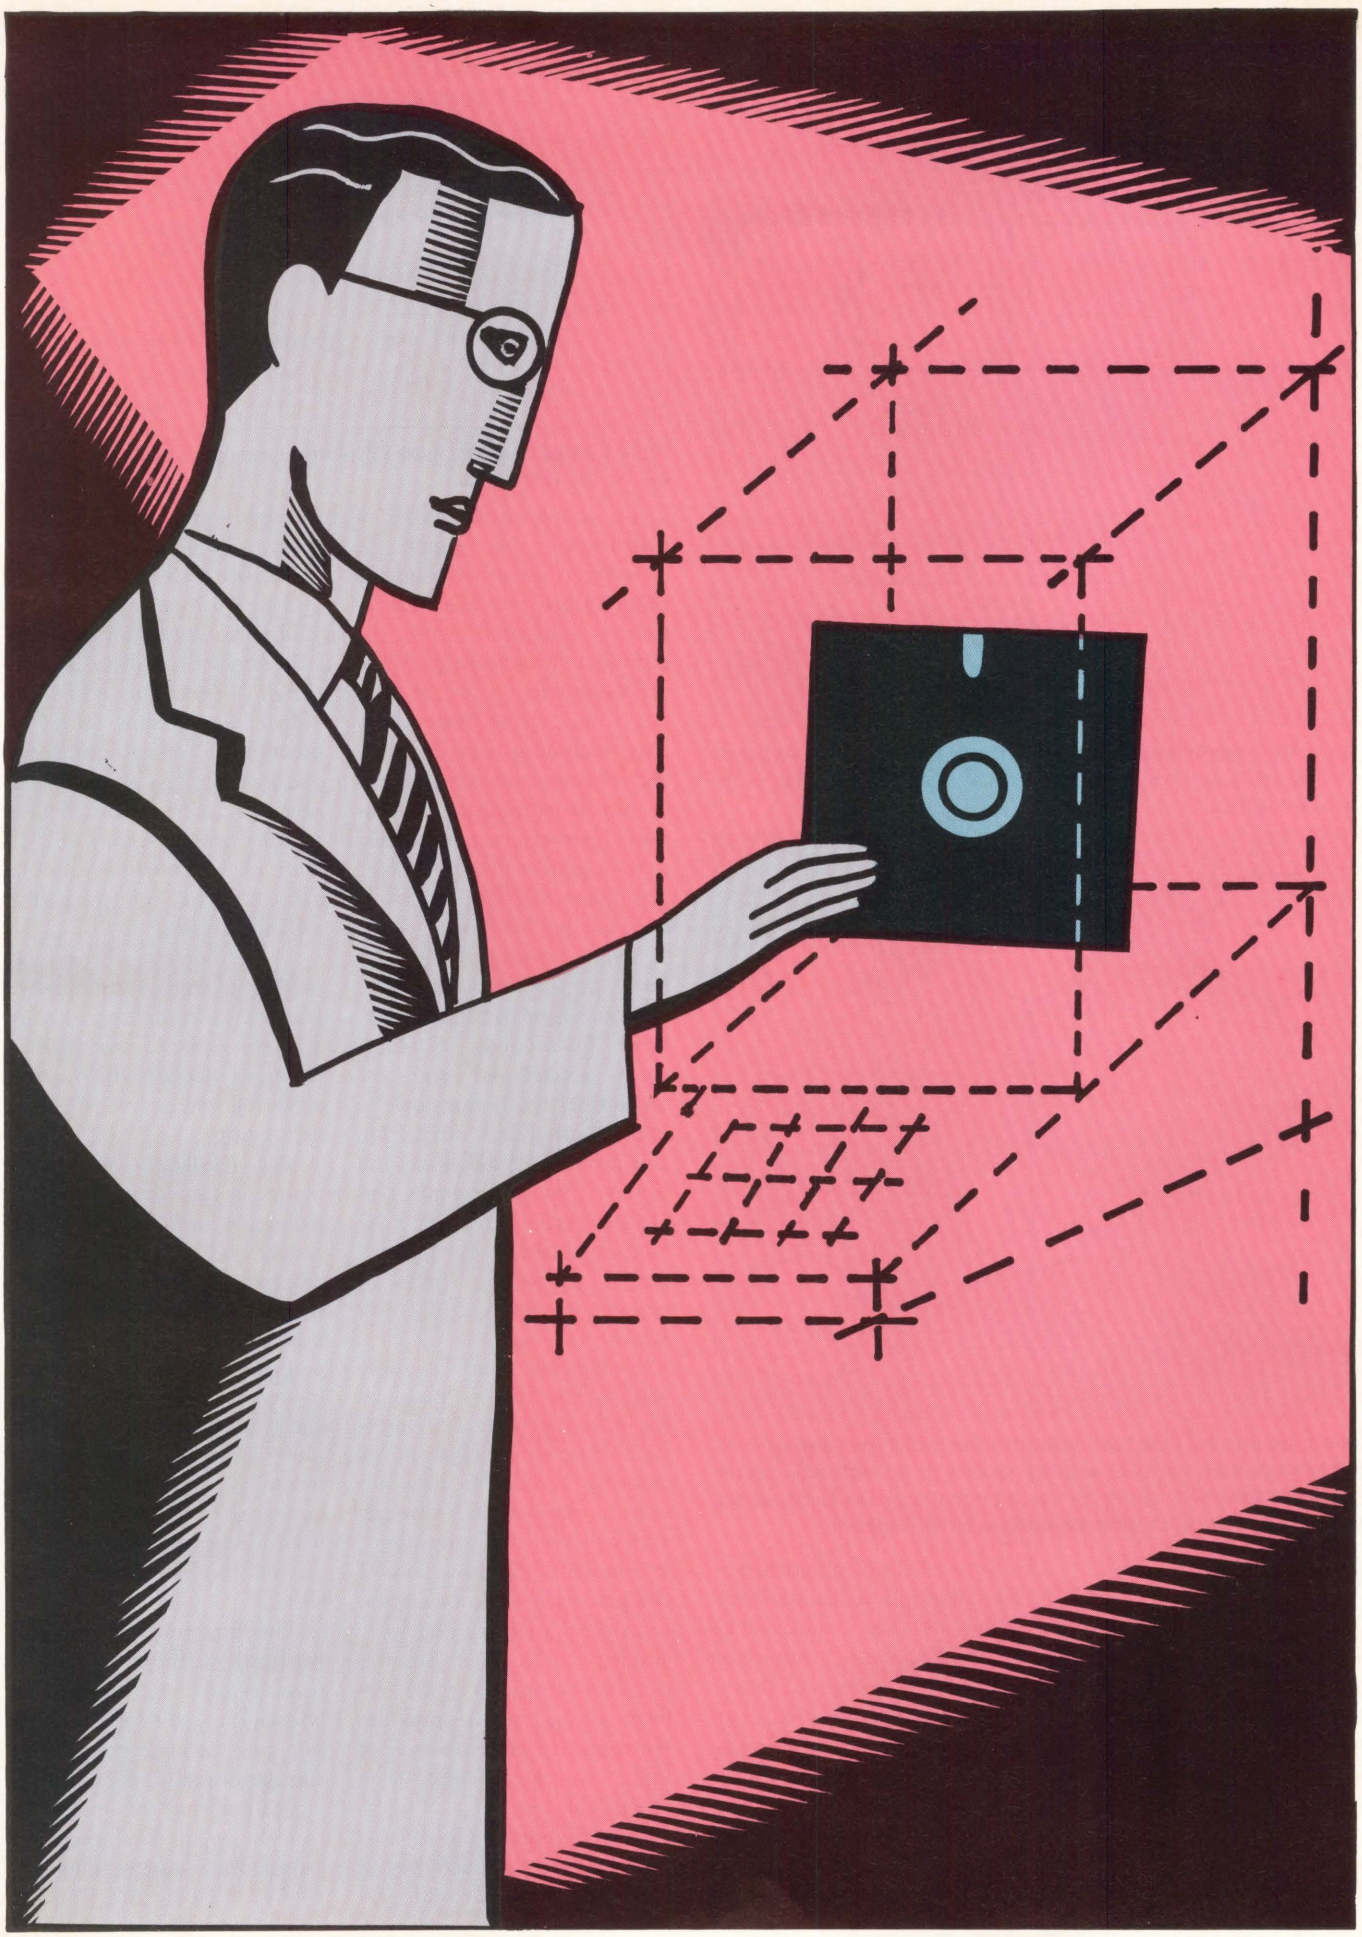

ILLUSTRATION BY ANTHONY RUSSO

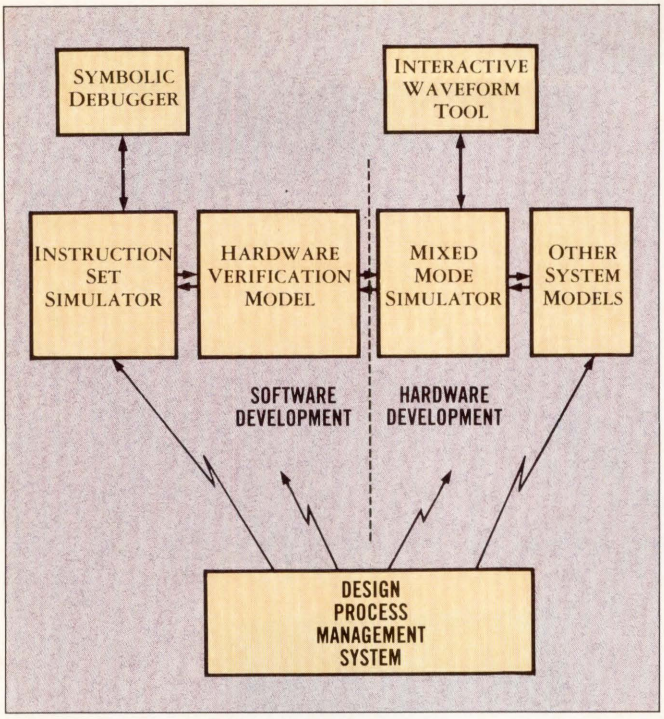

Figure 1. The System Development Environment combines an Instruction-set Simulator (ISIM) with a hardware verification model to allow software to run on the simulated hardware.

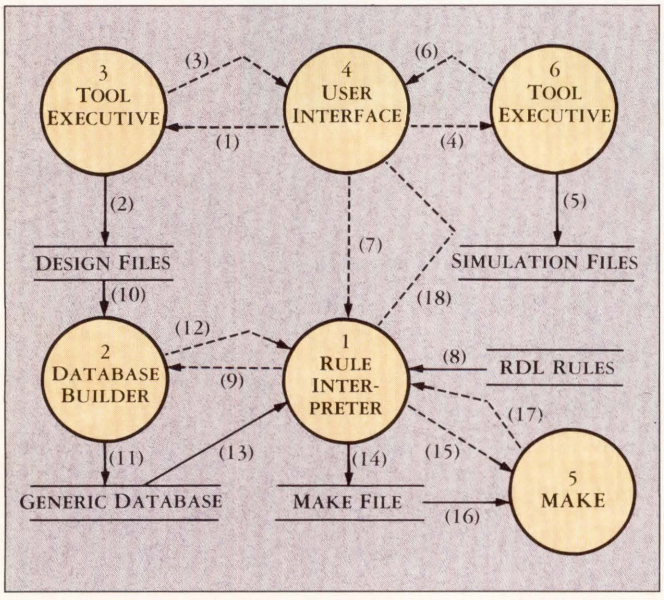

Figure 2. The data flow (solid lines) and control flow (dotted lines) diagram of the Design Process Management System (DPMS) demonstrate the relationship between the user interface, which calls design files and tools through Tool Executives, and the Rule Interpreter, which builds the database for the simulation.

■ Two Types of Simulators

A new perspective on modeling microprocessors is required to unify individual software and hardware simulations into a single simulation. That perspective approaches the problem from a software, rather than hardware, point of view.

Instead of a complete behavioral description of a microprocessor, INTERACT started with an Instruction-set Simulator (ISIM) and then coupled that to a behavioral model that characterizes only a microprocessor's pinouts and bus timing activity (Figure 1).

ISIM is a working model of a given instruction set operating within a processor environment. It represents the set of instructions that would be seen by an assembly-language programmer. ISIM excludes activities associated with the electrical characteristics of the processor. Used alone, it verifies a design's software variables, including data, memory and register contents, and other results of program execution. Instructions are modeled to the extent of their effect on a processor's registers and memory, but not its pins.

Some features that ISIM supports are: vectored interrupts and exception handling; pipelined instructions; floating-point accelerators and co-processors; translation-look aside buffering; instruction and data caches; and clocks and timers. ISIM implements single-step execution during a debugging session by instruction, by various kinds of breakpoints, and by interrupts.

In contrast, the hardware verification model produces the read-and-write and interrupt-acknowledge cycles at the processor's pins, but does not execute instructions. Through a mixed-level logic simulator, it simulates the functions of the processor while connected to external memory and logic circuitry. The logic simulations are controlled by a waveform editor, which provides the means to generate a set of input waveforms and then checks the results.

The combination of the two models fully characterizes a system, both the hardware and software, and executes program code at a speed 1,000 times that of the classical modeling approach. Speed is gained mainly because the instruction-set simulator is optimized to assess only how individual instructions affect a processor's internal registers and memory. As a result, SDE executes more than 5,000 instructions per second on a 1-MIPS platform when running simulations of popular microprocessors.

The alternative, behavioral modeling, tends to run slowly. It executes in logic simulators at a rate of two to four instructions per second, which is too slow for interactive software debugging.

The hardware verification model is the interface to the simulator and produces the proper logic levels for the microprocessor. It translates the instruction-set simulator's inputs and outputs—assembly-language statements—into the hardware world of logic levels and timing signals. For example, a write instruction in ISIM would be translated into the appropriate control signals with waveforms, as illustrated in the microprocessor's data sheets.

The instruction-set and hardware verification simulators can run independently of one another, so that a designer can run through consecutive simulations and debug individual pieces of the design. In this way he can check the set of programmed instructions and the electrical and logical

```
ruleset(ADA->ada_OBJ)
        action("Ada$filemane.ada$ada_options");
end;

ruleset(ada_OBJ->ada_EXE)
        action("link$filename$link_options");
end;

ruleset(ADA_EXE->Ada_Simulation)
        action("SoftICE-Ada-load$filename.ada_EXE$SIoptions");
end;
```

Figure 3. Adding these lines to the SDE rule base integrates a user's ADA compiler into the SDE environment.

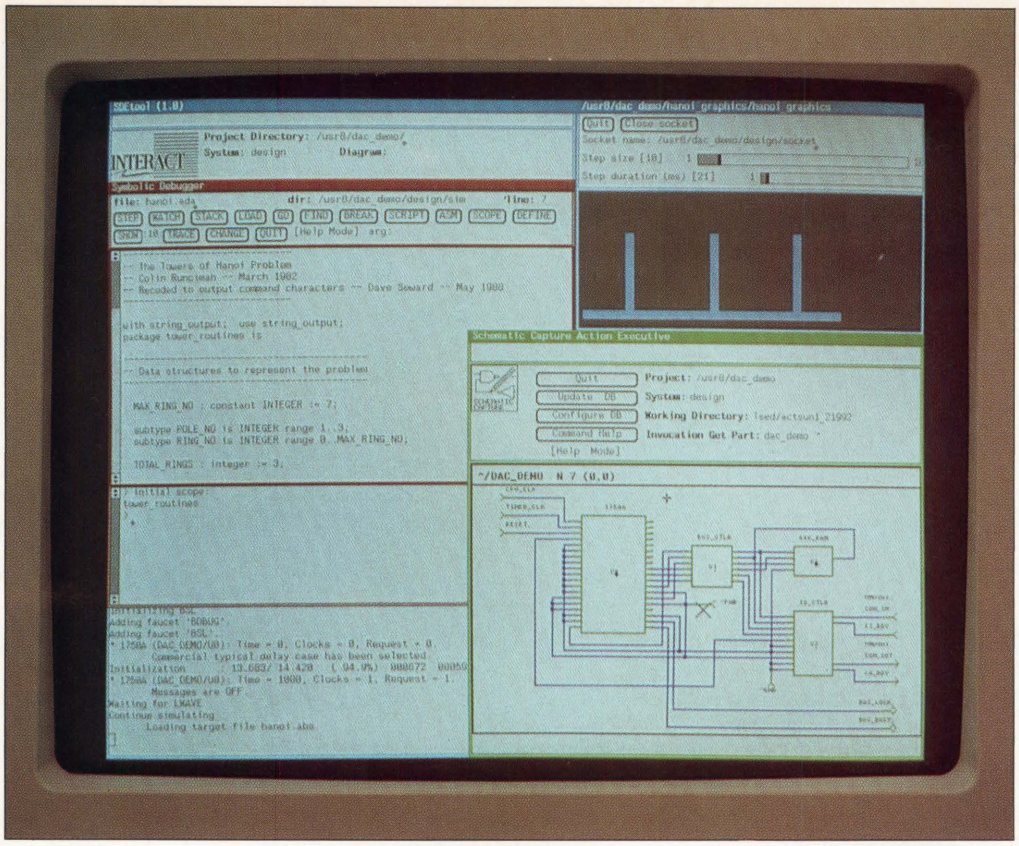

Figure 4. A screen with the system ready for the user to begin a debugging session. The symbolic debugger window shows the "Towers of Hanoi" software program, while the schematic capture tool shows the hardware design. The upper right window indicates that the expected initialization has not taken place.

soundness of the design. This includes operation of memory, I/O, peripheral chips, and interrupt procedures.

The simulators must operate together, so INTERACT had to come up with a way of mapping the values taht are derived during standalone software simulation into a hardware-simulation universe—that is, into logic values. The key to the solution was an algorithm that translates logical events which occur in ISIM to pin (electrical) events in the hardware logic-simulation model. The algorithm also converts bus events in the logic simulator back into information for ISIM. A proprietary communications mechanism synchronizes the flow of data between the two models.

The synchronization scheme results not only in a tremendous improvement in the speed of simulation, but also allows the user to control the extent to which electrical events are simulated. The user can start out with the ISIM only, giving no consideration to electrical events, and run SDE at its fastest speed. Then he can change the configuration so that ISIM passes some information to the hardware simulator, allowing the user to start checking electrical events. The process can continue, while the user controls the degree to which logical events pass to the hardware simulator.

For example, if the microprocessor is running a program out of ROM, the user can direct the software to model the electrical characteristics of only the ROM port. Then the user can start modeling other ports, such as the data bus, on the microprocessor. While looking at the data bus, the user can remove the ROM port model to simplify the simulation. In this way, the user can: change the hardware model's parameters or operating characteristics; and simulate all required electrical events.

To integrate the debugger into the workstation environment, it was enhanced to work like a Sun-type point-and-click debugger, which uses icons and menus instead of line commands. The enhancements were added to the other debugging capabilities such as: stepping through code; tracing code execution at either the source or disassembly level; assembling or disassembling instructions; examination of variables or registers; and dynamic modification of code and variables.

In addition to software breakpoints, a hardware breakpoint can be set by the user based on a time step, a node value or a combination of values. For instance, the user can set the simulator to stop if node x equals a "1" or a set time exceeds 20,000. A series of hardware breakpoints may be set in parallel with a series of software breakpoints. This provides the user with a combined software debugger and hardware waveform analyzer.

■ Database Management

Within SDE, a system manager, called the Design Process Management System (DPMS), connects the user to the tool suite through a uniform graphical interface. DPMS also controls all background tasks, such as invoking appropriate tools, tracking tool and design files, transforming or cleaning up files, and invoking a rules interpreter during simulation.

The DPMS isolates a user from all the intermediate steps typically involved in a CAE development process. The user needn't worry about tedious compilations, linking or "housework."

With DPMS, users also avoid the necessity of learning a new CAE system. They enter the source code and hardware diagrams, draws waveforms, and then pushes a simulate button. That's all the user needs to learn about the system.

SDE's designers provided this seamless framework for integrating tools, automating execution, and tracking data by tying a hierarchical, object-oriented design database to a rule base. The database tracks all design files, while the rule base contains the data-accessing methods.

The database holds the design configuration and allows the user to select simulation parameters. Simulation parameters include: signals to be monitored; model to be active; or combinations of gate and behavioral simulations to be run. The user can modify the interface, rules, or database. He can provide objects, add documentation or graphical files, change components, add processing routines, modify rules, or add other CAE software or hardware tools.

In operation, all intermediate steps are handled by DPMS, which: invokes tools when the user selects an icon; performs the necessary intermediate and final clean-up tasks; and returns control to the user interface. When DPMS receives a command, it calls the generic database functions that access the Design Configuration Database (DCD). The DCD defines the structure and state of a

10 Good Reasons to Get NEC's FREE Gate Array Gateway Kit

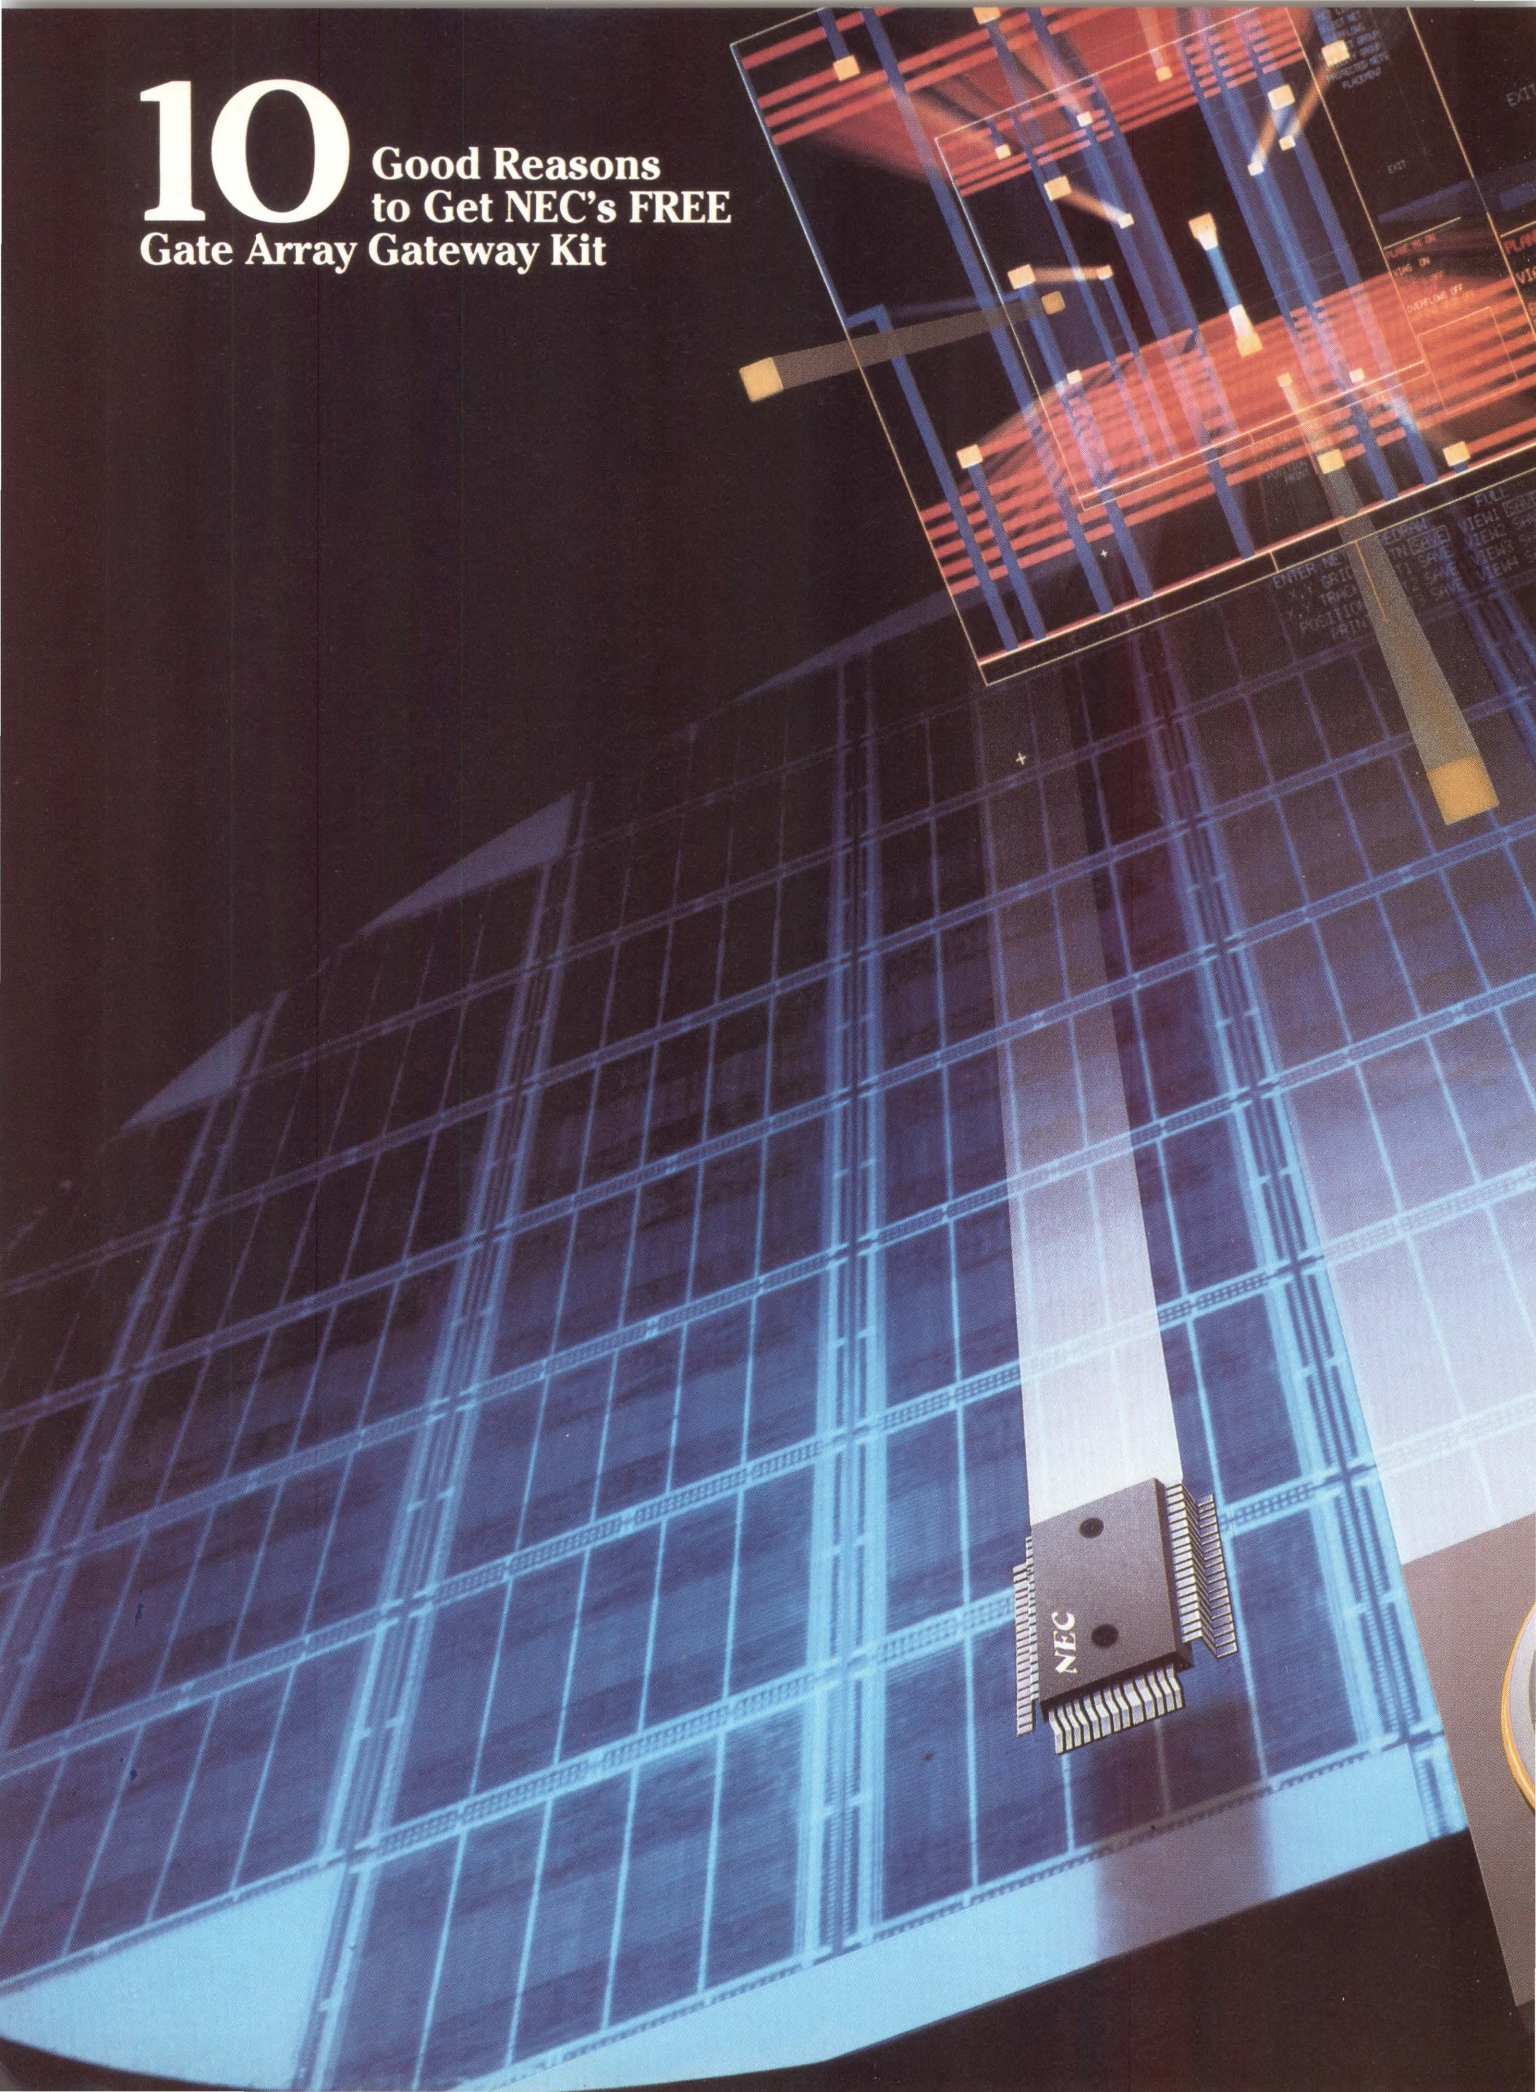

NEC's FREE Design Kit is Your Gate Array Gateway to

1. CMOS, BiCMOS, and ECL technologies.
2. Valid, Sun, Mentor, Daisy, HP 9000, FutureNet, Calma TEGAS V, and HILO compatibility.
3. A powerful block library with more than 170 macros, including analog, peripheral, and megamacro functions.
4. Cell utilization rates of 95% over a range of 800 to 45,000 gates.
5. Better than 95% first-time success with engineering prototypes.
6. The most experienced manufacturing capability.
7. The widest choice of through-hole and surface-mount packages.
8. The most effective way to lower your system costs.
9. Local design support with local FAEs and design centers backed by a satellite network of computers for worldwide communications services.
10. The easiest translation of your ideas into silicon.

NEC Electronics Inc.

401 Ellis Street, P.O. Box 7241
Mountain View, CA 94039

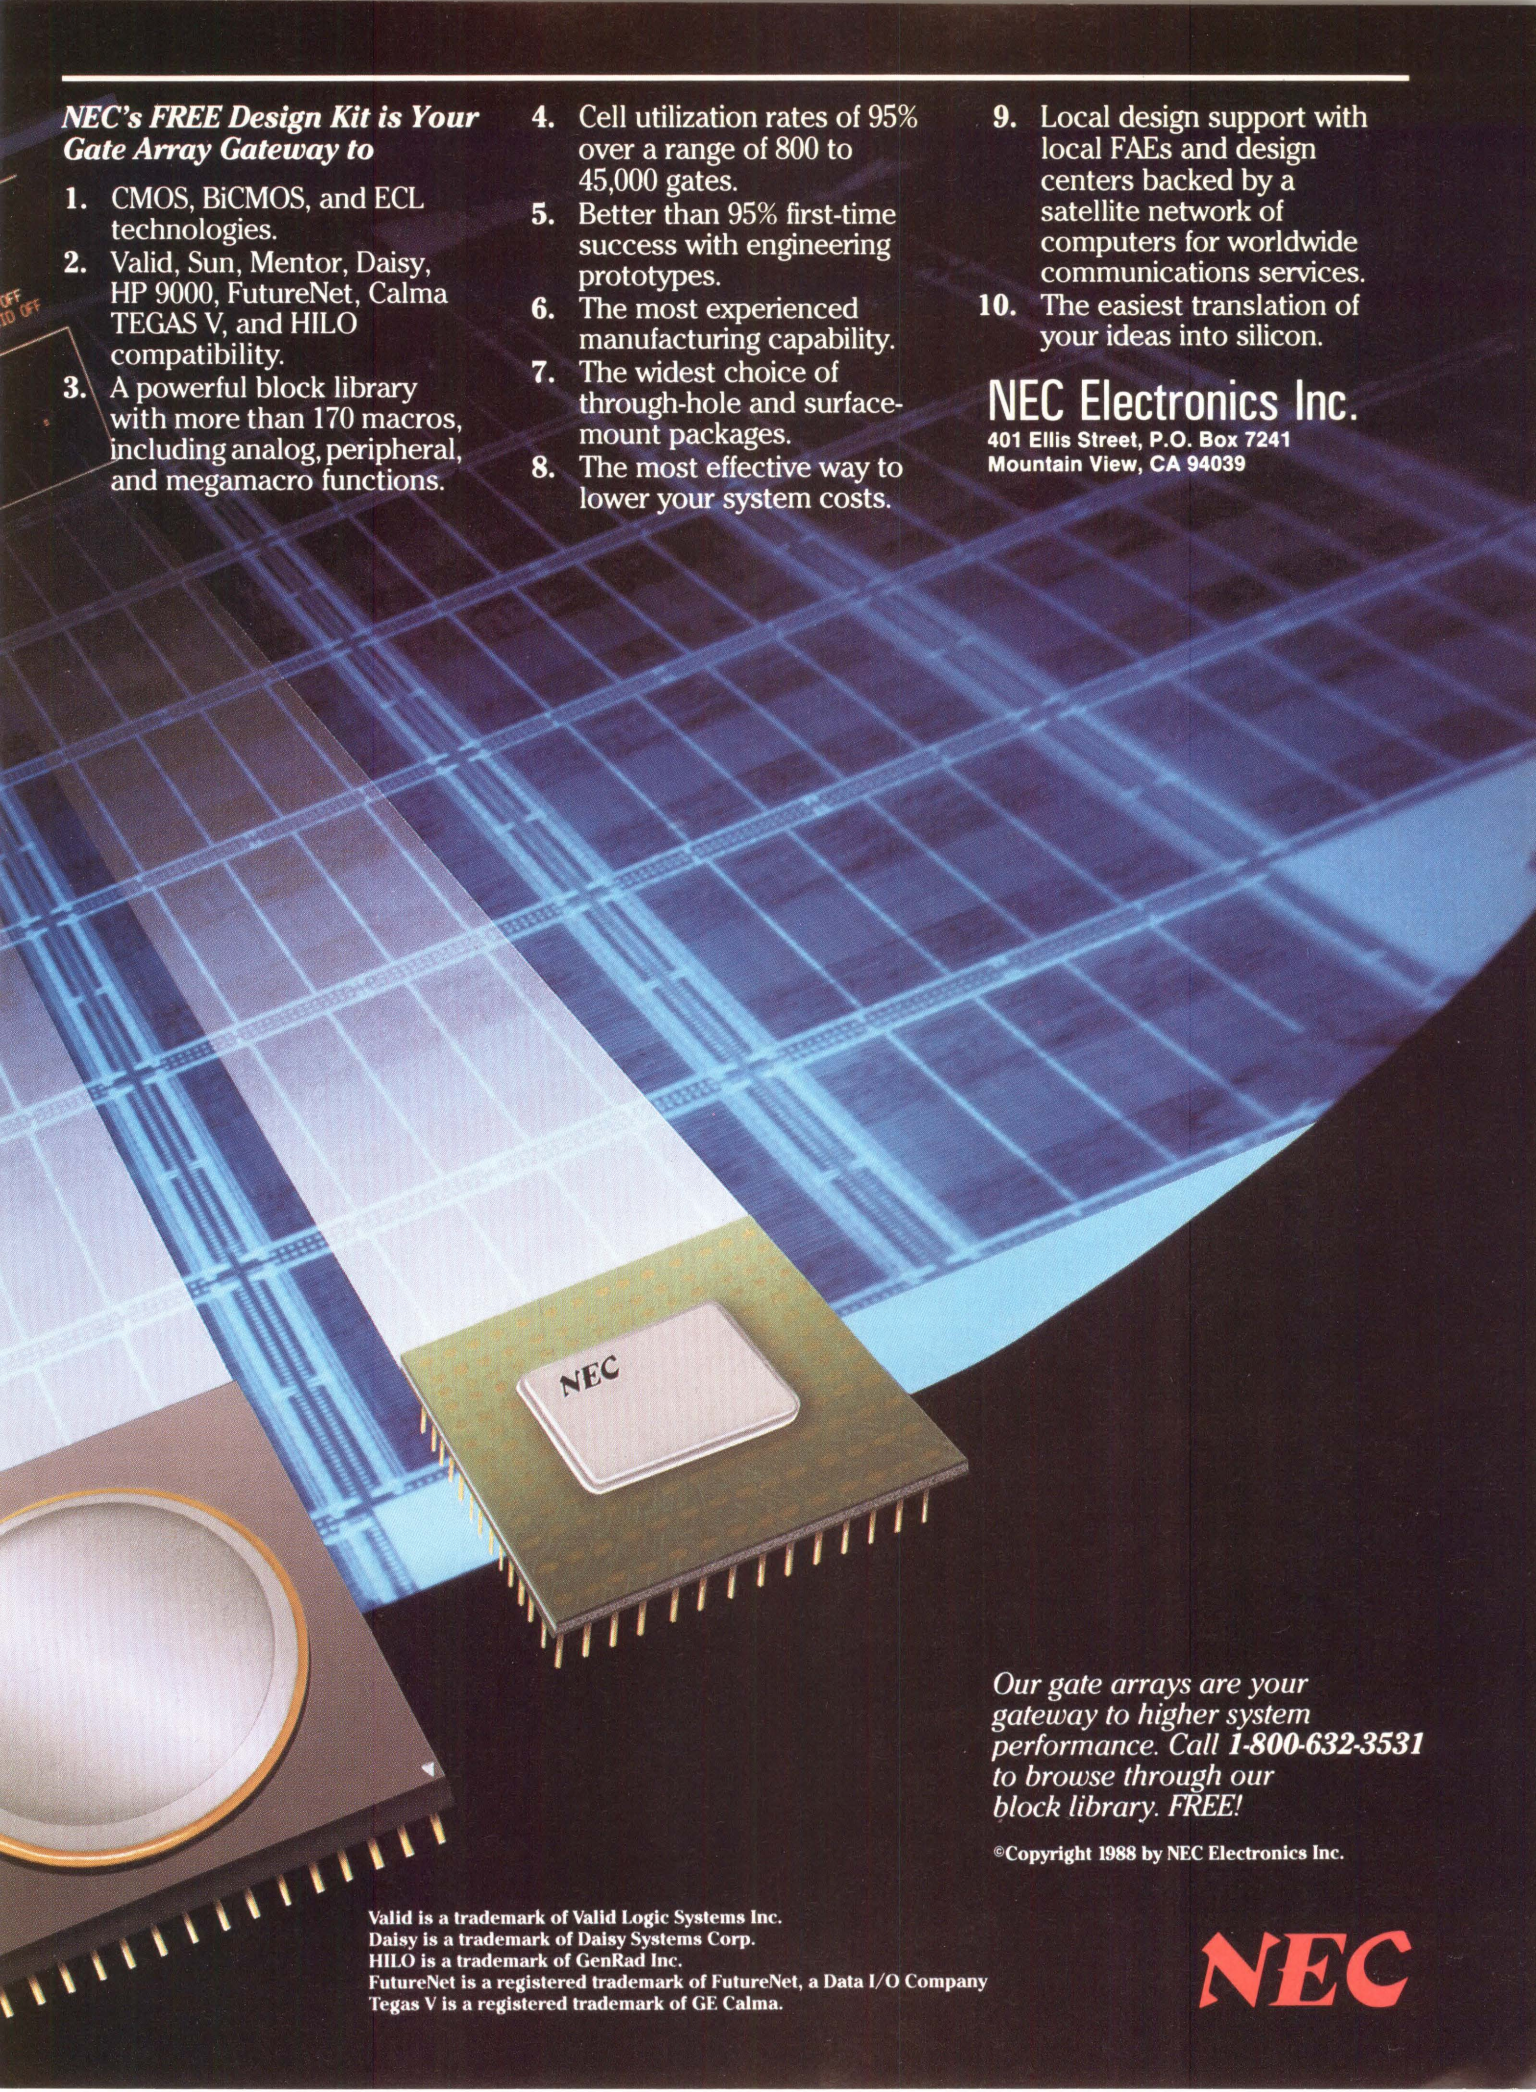

Our gate arrays are your gateway to higher system performance. Call ***1-800-632-3531*** *to browse through our block library. FREE!*

NEC

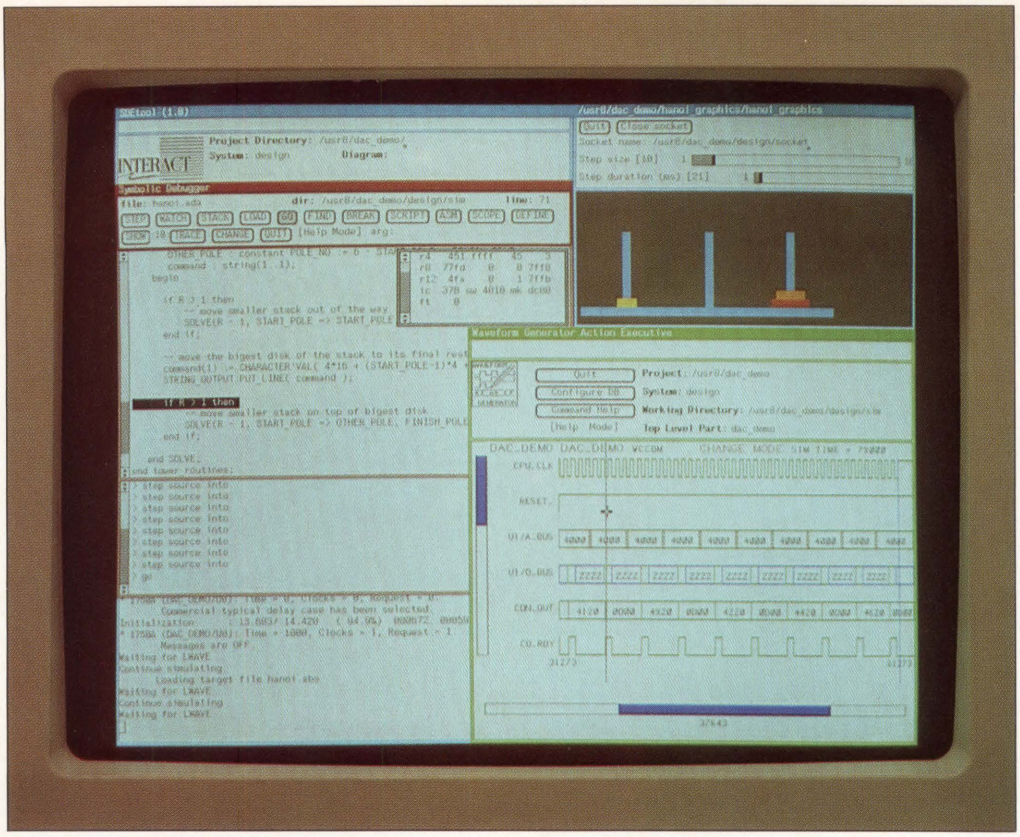

Figure 5. The screen display indicates that a software breakpoint has been reached in a combineed software/hardware simulation.

design; how the components of a design are modeled, as well as pointers to the model of the microprocessor and to the other portions of the design.

The information returned from a call is used to retrieve all files required to invoke a tool. Next DPMS copies files, extracts them from revision-control systems, executes intermediate programs, and invokes the requested tool.

The DPMS interface routines comprise programs that call or add functionality to design tools (Tool Executives) and a rule interpreter, which is called the RDL (Rule Design Language) Interpreter. The RDL Interpreter is invoked when an icon is selected. It handles simulation and other tasks that require file transformations—compiling source codes, for example. The RDL Interpreter accepts two files as input: the RDL rules and the generic database containing the hierarchical representation of the design to be simulated.

The interpreter scans the design hierarchy, and applies rules that match design data stored in the generic database, to find information relevant to the invoked task. The interpreter also invokes the appropriate tools to transform the original design files into the form needed for the task.

Figure 2 demonstrates the operation of the DPMS interface. The user interface calls up a design file (steps 1 through 3), and invokes the simulator (steps 4 through 6). The rule interpreter is invoked (step 7). It uses the RDL rule base to run appropriate database conversion programs to create the design input for the simulator (the generic database). Then the rule interpreter creates a Unix MAKEFILE from the database and the Unix MAKE utility (steps 14 through 17) to begin the simulation.

RDL is written in a textual format that specifies how to manipulate hardware or software files in preparation for simulation. SDE's power is increased by its ability to modify RDL rules to meet individual needs. Assume a user wants to integrate an Ada compiler into SDE, invoke the compiler from the user interface, and simulate software written in Ada. The user would add the code in Figure 3 to the RDL rule file. That code describes how to compile, link, and run the SDE Symbolic Debugger on the Ada source code.

A designer can change the SDE rules to accommodate design flow, but team managers can also install control procedures—they can prevent more than one design change at a time, or force designers to use company standards such as level-sensitive-scan design.

■ SDE's Incarnation

Major tools in the integrated SDE suite include those for front-end software (CASE), hardware design and analysis, simulation and debugging, and the Design Process Management System to control the entire environment. Within the tools are a comprehensive set of compilers, linkers, assemblers, symbolic debuggers, and waveform generators. Typical screen displays during operation are shown in Figures 4 and 5.

SDE is now capable of designing 1750A-based embedded systems on a Sun workstation. There are plans to port SDE to the DEC VAX, the DEC MicroVAX, and Apollo Computer workstations. SDE will include models of the Motorola 68000 and Intel 8086 families, and the MIPS R3000 and Intel 80960 RISC chips. Future enhancements will include the ability to generate and download vectors for automatic test equipment.

SDE is an extensible, open system for the user. A user can tailor it to specific methodologies and needs by modifying rules, adding component libraries, or adding pre- or post-processing routines. The user can take an application-driven system from concept to reality.

The final integrated environment overcomes major inefficiencies in current embedded-systems design practice. They are the inability to integrate the hardware and software elements prior to building prototypes; inability to isolate design errors to the various elements, and inability to predict performance with high accuracy.

In addition, changes can be made on the fly to either the software or hardware. This facilitates the designer's ability to optimize systems. Software and hardware can be traded off as often as necessary to achieve optimization, because all design changes are quickly simulated on the workstation. ■

About the Authors

ANDREW KERN *is a software engineer at InterAct working on the Design Process Management System. Previously, he worked at Standard Microsystems Corp. and Hazeltine Corp. as an IC design engineer. He holds a BSEE from the Polytechnic Institute of New York.*

AL BLAZEVICIUS *is the engineering manager at InterAct for the System Design Environment project. He has previously held management and design engineering positions in CAE, CAD, and custom/semicustom IC design for Hazeltine Corp., LSI Logic Corp., and Hughes Aircraft Corp. He received his BEE and did graduate work at the Cooper Union.*

Hardware Accelerators

VLSI Systems Design Staff

ACCELERATORS CAN MAKE SIMULATION INTERACTIVE

ONLY a few years ago, simulation was something that design engineers liked to talk about. They even bought simulation software packages. But surveys indicated that most design engineers weren't simulating—even if they had a simulator. The most common excuse was that simulation of anything but a very simple design was a batch task—and they couldn't wait for up to two weeks to get their simulation run on the mainframe or mini. In addition, the design engineer "knew" his design was correct—besides, he could correct any errors after testing the prototypes. That was a bad philosophy a few years ago and in today's market it could spell disaster. Most companies can't afford the penalty in time and money for taking a second crack at a design that didn't work the first time. This is particularly true for ICs, but—since many of the printed circuit assemblies coming off today's production lines are high-density boards loaded with SMT components and fabricated in fine-line, multilayer technologies—second tries for board designs have also become unacceptable.

This is where hardware accelerators have come to the rescue. They should have a bright future,

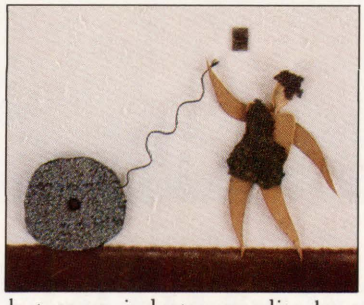

but some industry pundits have predicted that accelerators will have a short lived existance—eventually squeezed out of the market by the growing power of workstations. However, although today's best workstations deliver 10 or even 20 Mips, which is a lot of power for many design automation applications, when it comes to the code crunching requirements for simulating complex ICs and systems, 20 Mips barely scratches the surface.

Most hardware accelerators have code crunching power equivalent to thousands of Mips, and some even an order of magnitude higher. However, since the accelerators are targeted at just one task, they can provide these equivalent Mips at prices that start below $100 per Mips. This compares favorably to workstations that start at several thousand dollars per Mips.

The high Mips capability of hardware accelerators is not a luxury, it's required to make simulation interactive. If an accelerator cuts the simulation time to a few hours or even a day, it can encourage the designer to iterate his design until it's optimized.

When choosing an accelerator be sure that: it can run your simulator software; it supports the type of simulations that you plan (such as logic, fault, concurrent logic and fault, heirarchical, and hardware modelling—at the desired levels); it has the gate capacity to run your most complex designs; and that it has enough power to run your simulations reasonably fast. In the last category, ask the vendor how he defines gate evaluations or events because they don't all use the same basis.

This is no time to test for faults.

If it doesn't fly, it's your fault.

You didn't see it. Neither did your test equipment. But now that the fault is in silicon, it's going to cost you.

In fact, field failures are over 100 times more expensive to fix than faults found in-house.

So why wait?

Complete fault grading at the design stage.

Now you can achieve total fault grading of ASIC, VHSIC, Custom VLSI, and entire system designs.

With a Zycad concurrent fault simulator.

Our simulators are dedicated hardware accelerators that detect faults hundreds of times faster than software

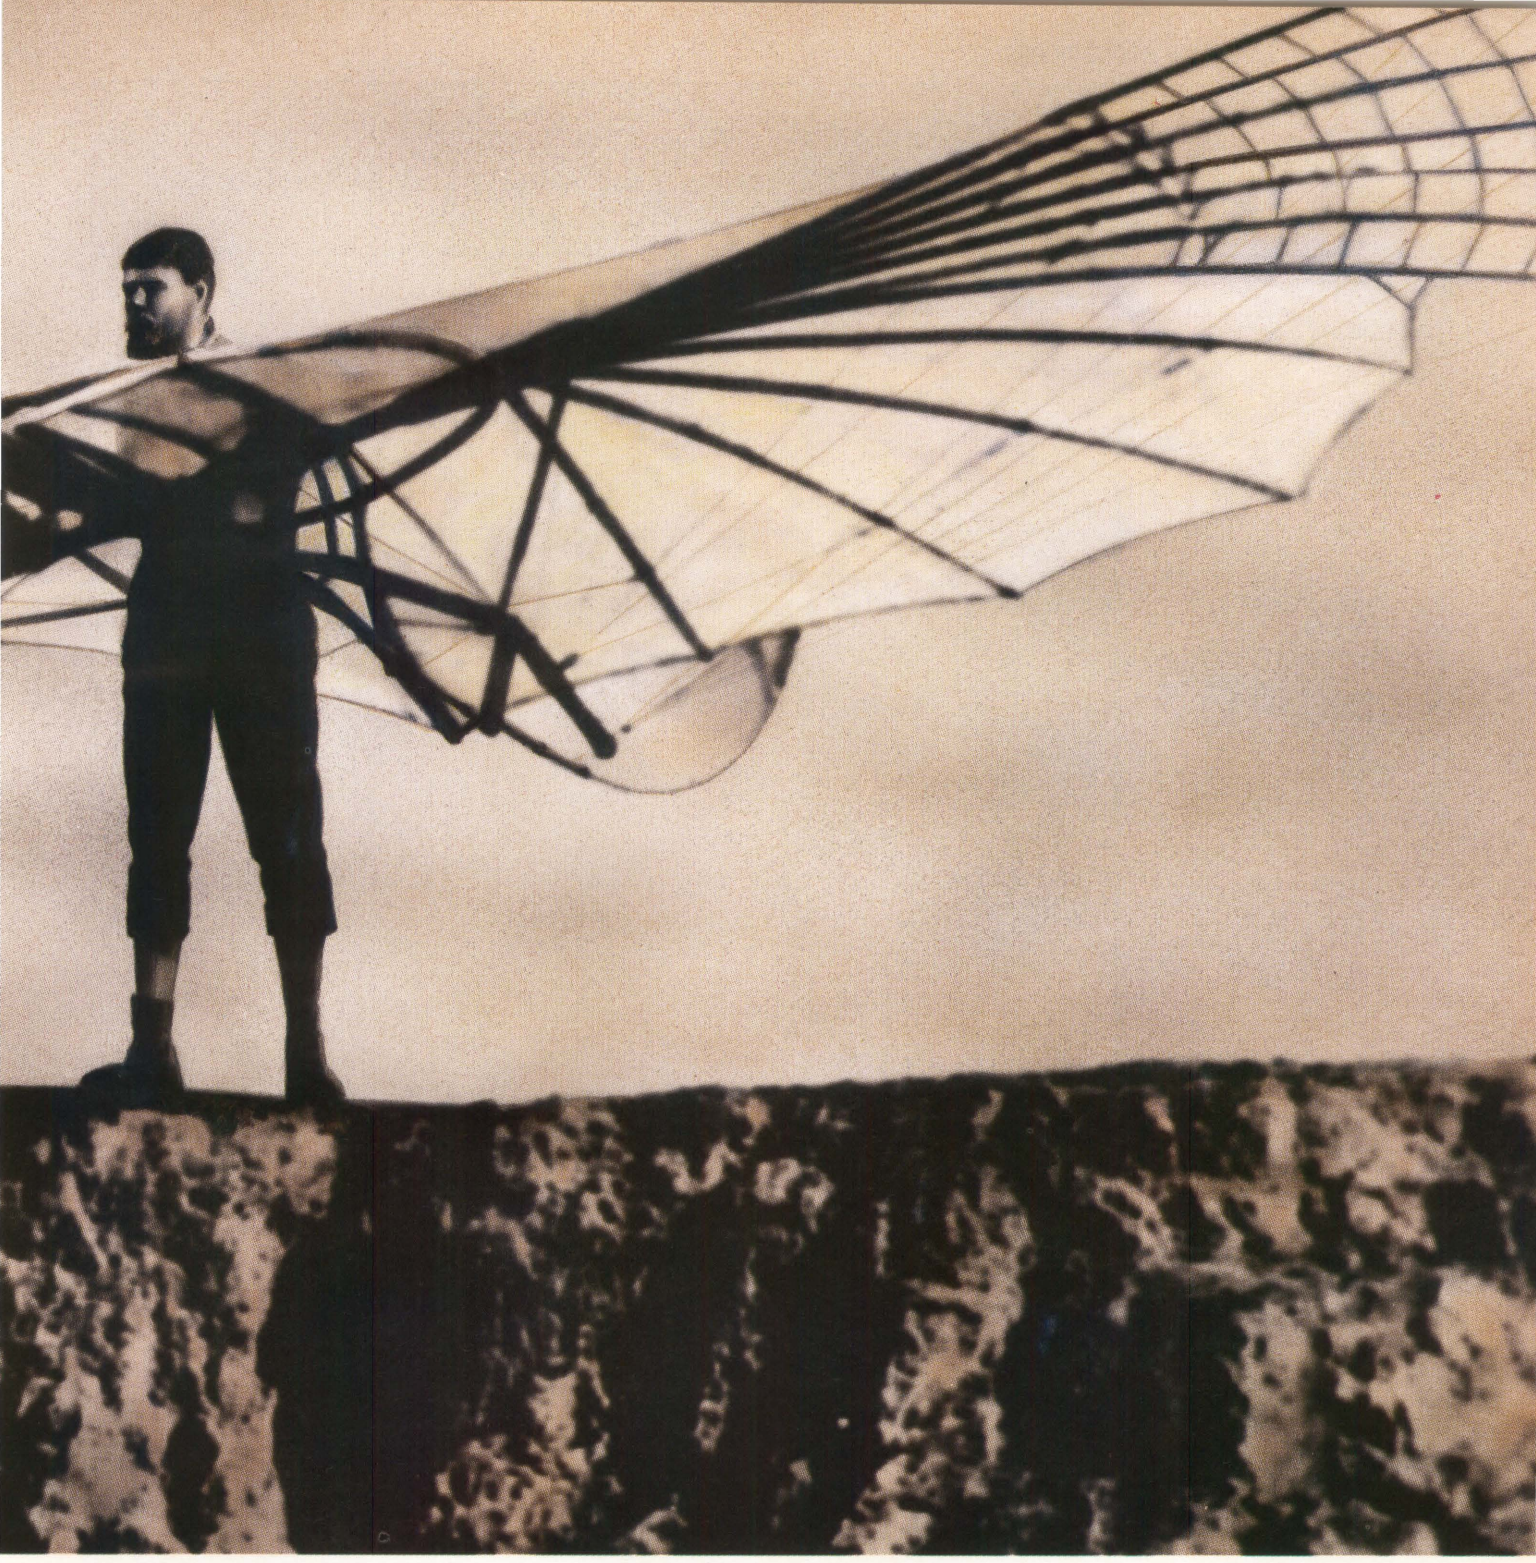

simulators running on mainframes or workstations.

So you can test Full-Custom, ASIC and even system designs with our full line of simulators, including the Fault Evaluator,® the MACH 1000™ and the new Magnum II.™

All our simulators run as direct extensions of the logic simulation environment you're using now. And they can be integrated with your present CAE equipment.

Free fault concept handbook.

To get our free handbook on fault simulation techniques and how you can save money with fault simulation, call 800-631-5040; in Minnesota, 612-490-2500. Or write us at Zycad Corporation, 3900 Northwoods Drive, Suite 200/Dept. FC, Arden Hills, Minn. 55112.

We'll help you fly right the first time.

zycad CORPORATION®

We accelerate your success.

CIRCLE NUMBER 9

DIRECTORY OF HARDWARE ACCELERATORS

VENDOR AND CONTACT	PRODUCT AND PRICE	TYPES OF SIMULATION	HOST REQUIREMENTS
AIDA (TERADYNE INC.) *5155 Old Ironsides* *Santa Clara, Calif. 95051* Peter Hwang Product Marketing Manager (408) 980-5200	**AIDA CO-SIMULATOR ACCELERATORS** $10k–$200k +	Logic and fault (can support concurrently)	Host computer required (not included in price); interfaces available for Apollo DN3000, Apollo DN4000, Apollo DN10000, Apollo DN5xx Sun 3 and Sun 4
DAISY SYSTEMS CORP. *700 E. Middlefield Rd.* *Mountain View, Calif. 94039* Sanjiv Kaul Product Manager (415) 960-6343	**MEGALOGICIAN** $62k	Logic, fault, (can support hierarchical simulation)	Host computer required (not included in price); interfaces available for Sun386i, LOGICIAN 386, Personal LOGICIAN 386, LOGICIAN 286, Personal LOGICIAN 286
	PERSONAL MEGALOGICIAN $24k	Logic, fault, (can support hierarchical simulation)	Host computer required; interfaces available for Sun386i, Personal LOGICIAN 386, Personal LOGICIAN 286
	GIGALOGICIAN $180k–$300k	logic, (can support hierarchical simulation)	Host computer not required
HHB SYSTEMS (CADNETIX CORP.) *1000 Wyckoff Ave.* *Mahwah, N.J. 07430* Ralph Zak Director of Marketing (201) 848-8000	**CATS ACCELERATOR** $240k–$500k	Logic, fault, (can support hierarchical simulation)	Host computer required (not included in price); interfaces available for Sun, Apollo, VAX
IKOS SYSTEMS INC. *145 N. Wolfe Rd.* *Sunnyvale, Calif. 94086* Robert P. Smith Director of Marketing (408) 245-1900	**IKOS SIMULATION SYSTEM** $46.5k +	Logic and fault (can support concurrently); (can support hierarchical simulation)	Host computer required (not included in price); interfaces available for Sun, Apollo, IBM PC/AT, Intergraph
MENTOR GRAPHICS CORP. *8500 SW Creekside Pl.* *Beaverton, Ore. 97005* John Huber Technical Product Manager (503) 626-7000	**COMPUTE ENGINE** $36.9k	Logic, (can support hierarchical simulation)	Host computer required (not included in price); interfaces available for Apollo 3000, Apollo 4000, Apollo 570

MODELING LEVELS	HARDWARE MODELERS ACCOMODATED	SIMULATORS SUPPORTED	GATE CAPACITY	GATE EVALUATION RATE
Gate; functional (store/process in host); behavioral (store/process in host); (RAM and ROM models can be stored/processed in accelerator)	Teradyne D300 (Data Source)	Tegas 5; HILO 3; MDE (NDL); Mentor (DFI); Toshiba (TDL); CASE	64k–8M	5M–10M evaluations per second
Switch; gate; functional; behavioral	Daisy PMX	Daisy DLS	Up to 256k	100k evaluations per second
Switch; gate; functional; behavioral	Daisy PMX	Daisy DLS	Up to 64k	100k evaluations per second
Switch; gate; functional; behavioral	Daisy PMX	Daisy DLS	96k–256k	1.2M–3M evaluations per second
Switch; gate; functional; behavioral	CATS 6000, 8000, 10000, 11000, 12000	Cadat 5.1 and 6.1	64k primitives–256k primitives	300k events per second
Switch; gate; functional	N/A	IKOS simulation system (translator available; translator/compiler is incremental)	Up to 1M	3M–45M evaluations per second
Switch; gate; functional; behavioral	Mentor Graphics hardware modeling library	Mentor Graphics QuickSim, MSPICE, MSIMON	More than 500k primitives; more than 1M gates	100k evaluations per second

DIRECTORY OF HARDWARE ACCELERATORS (continued)

VENDOR AND CONTACT	PRODUCT AND PRICE	TYPES OF SIMULATION	HOST REQUIREMENTS
SIMULOG INC. *3211 Scott Blvd.* *Suite # 104* *Santa Clara, Calif. 95054* Bruce K. Barnard Vice President, World Wide Sales (408) 727-8272	**SUPERSIM** $600k–$12M **MINI SUPERSIM** $35k–$220k	Logic, (can support hierarchical simulation)	Host computer required (not included in price) interfaces available for DEC VAX, IBM mainframe under UTS, Apollo, Sun
VALID LOGIC SYSTEMS INC. *2820 Orchard Pkwy.* *San Jose, Calif. 95129* Don Mazur Marketing Manager (408) 432-9400	**REALFAST** $35k–$100k + **REALMODEL** $70k–$130k +	Logic, (can support hierarchical simulation)	Host computer required (not included in price); interfaces available for Sun 3 and 4, DEC VAX (VME), MicroVAX (VME), Unibus, Qbus; Valid 5320, multibus
ZYCAD CORP. *3500 Northwoods Dr.* *Suite 200* *St. Paul, Minn. 55112* Peggy Kalina Marketing (612) 490-2500 (800) 631-5040	**MACH 1000 LOGIC AND FAULT SIMULATION ACCELERATOR** $85k–$1M	Logic and fault (can support concurrently); (can support hierarchical simulation)	A host is built into the MACH 1000 (included in price); interfaces available for Sun, Apollo, DEC VAX
	LOGIC/FAULT EVALUATOR $150k–$2.1M	Logic and fault (can support concurrently); (can support hierarchical simulation)	Host computer required (not included in price); interfaces available for DEC VAX, IBM, Apollo, Sun, and others
	MAGNUM I & II SERIES ACCELERATOR $38k–$290k	Logic and fault (can support concurrently); (can support hierarchical simulation)	Host computer required (not included in price); interfaces available for DEC VAX, IBM, Apollo, Sun, and others
	SYSTEM DEVELOPMENT ENGINE (SDE) $575k–$3.3M	Logic, (can support hierarchical simulation)	Host computer required (not included in price); interfacea available for DEC VAX, IBM, Apollo, Sun, and others

MODELING LEVELS	HARDWARE MODELERS ACCOMODATED	SIMULATORS SUPPORTED	GATE CAPACITY	GATE EVALUATION RATE
Gate; functional (store/process in host); behavioral (store/process in host)	Proprietary bus available for customers use	Simucad Silos Netlist Input (translators for other input languages available)	8k–123k (Mini SuperSim); 262k–10M gates (SuperSim)	Exhaustive Simulation-Can model 1M gate designs at a simulated speed of 1–10KHz for SuperSim; Mini SuperSim-0.5–10 billion gate evaluations per second
Switch; gate; functional; behavioral (store/process in host)	Valid—Realchip, Networked Realchip, Realchip II, Realmodel	ValidSIM (translator not necessary; translator/compiler is incremental)	32k–1M	200k–500k event per second; 600k–1.5M evaluations per second
Switch; gate; functional (store/process in accelerator and host); behavioral (store/process in host)	HHB CATS	Mentor QuickSim (fully transparent); GenRad HILO 3 (fully transparent); Daisy (fully transparent); VTI; NCR; Silicon Compiler Systems; Intergraph	64k–512k primitives	1.5M–9M evaluations per second
Switch; gate; functional (store/process in accelerator and/or host); behavioral (store/process in host)	No commercial interfaces are currently available, however, the EventLink option is designed to permit proprietary interfacing to hardware modelers	Zycad Endot VHDL, ISP' (translator available; translator is incremental); Zycad ZILOS (translator available; translator is incremental); Mentor Graphics (General), KSIM, QuickFault (translator available; translator is incremental); GenRad HILO (translator available; translator not incremental); Daisy Systems (translator available; translator not incremental); Calma Tegas (translator available; translator is incremental); Silvar-Lisco (available from vendor); ViewLogic Systems (available from vendor); LSI Logic Inc. (available from vendor); AT&T Hasten (available from vendor); SimuCad SILOS (available from vendor)	64k–1.1M	3M–48M evaluations per second
Switch; gate; functional (limited)	No commercial interfaces are currently available, however, the EventLink option is designed to permit proprietary interfacing to hardware modelers	Zycad Endot VHDL, ISP' (translator available; translator is incremental); Zycad ZILOS (translator available; translator is incremental); Mentor Graphics (General), KSIM, QuickFault (translator available; translator is incremental); GenRad HILO (translator available; translator not incremental); Daisy Systems (translator available; translator not incremental); Calma Tegas (translator available; translator is incremental); Silvar-Lisco (available from vendor); ViewLogic Systems (available from vendor); LSI Logic Inc. (available from vendor); AT&T Hasten (available from vendor); SimuCad SILOS (available from vendor)	16k–256k	1.5M–6M evaluations per second
Switch; gate	No commercial interfaces are currently available, however, the EventLink option is designed to permit proprietary interfacing to hardware modelers	Zycad Endot VHDL, ISP' (translator available; translator is incremental); Zycad ZILOS (translator available; translator is incremental); other commercial and proprietary simulators	64k–1.1M	300M–1.5B evaluations per second

HOW WE CREATE AN UNEQUALLED OPPORTUNITY.

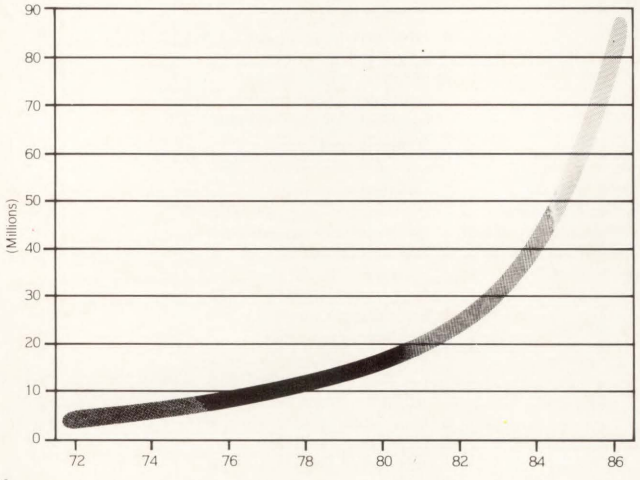

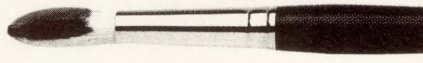

CMP is one of the fastest growing publishers of business newspapers and magazines, with advertising sales up better than 20% a year for each of our fifteen years. But growth isn't an end. Rather, it's our means to unequalled opportunity.

By growing, we create new opportunities. Opportunities for good reporters to advance to responsible editors. Opportunities for staff to become supervisors. Opportunities for sales people to become sales managers. Opportunities for stars to go anywhere their talent takes them. The fact is that you can't hold outstanding people unless you offer them a better opportunity than they can find elsewhere. And you can't offer an unequalled opportunity without consistent growth.

So we plan our growth. We plan to expand the scope and coverage of each of our publications. And we plan to add new publications in a variety of emerging industries. It's not just that we find bigger numbers exciting. What's exciting is seeing good people respond to bigger opportunities. And good people make CMP publications usually the best read in their industries.

CMP Publications, Inc.
600 Community Drive
Manhasset, NY 11030
(516) 562-5000

Electronic Buyers' News
Electronic Engineering Times
VLSI Systems Design
Computer Systems News
Computer Reseller News
InformationWEEK
CommunicationsWeek
Business Travel News
Tour & Travel News

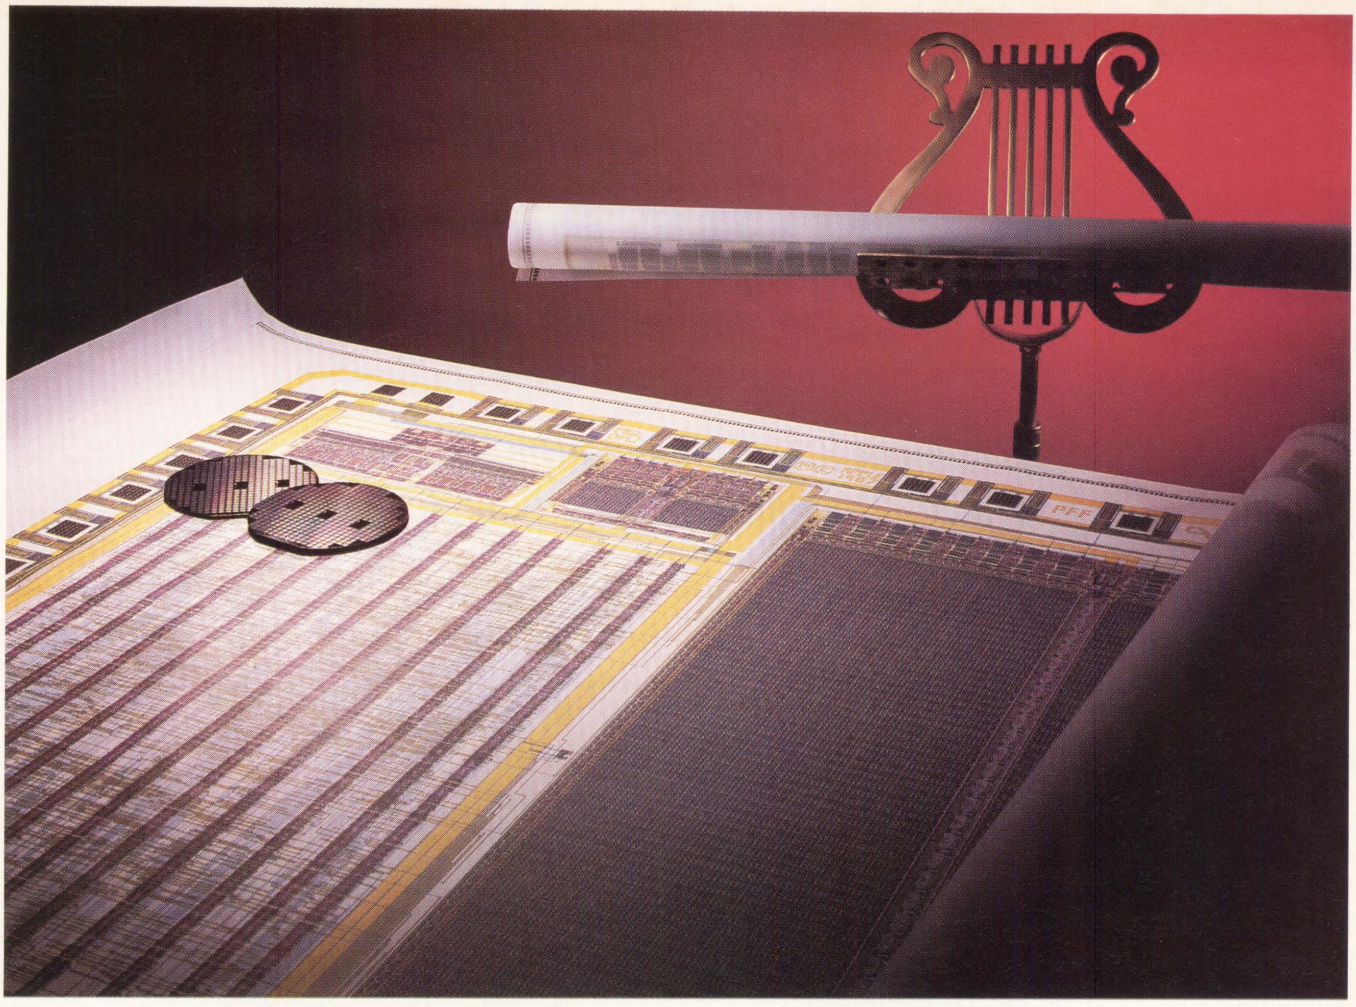

Sheet Music.

We're looking for IC design engineers who see themselves as *composers* of technology. Advanced technology, like the innovative sound chip shown above.

People who can combine creativity and technical vision into exciting breakthrough designs in graphics, sound, parallel processing, IO controllers and system integration.

At Apple Computer, we've had a few hits of our own. In fact, you'd be joining one of the most creative IC design groups anywhere.

The challenge? To create high-volume ASICs using state-of-the-art CAD equipment. The Macintosh™ II contains six such Apple-designed ASICs, including an advanced cell-based sound chip.

Our next generation of computer products will rely even more heavily on innovative IC design. So if you're the IC design genius you think you are, we'd like to talk to you now.

Join us.

This is a rare chance to have a major impact on our products. Or put another way, a chance to have a major impact on the world in general.

For more information, complete and return the attached card. Or send your resume to: APPLE COMPUTER, INC., Human Resources, Dept. JH44V, 20525 Mariani Ave., MS-9C, Cupertino, CA 95014. An equal opportunity employer.

The power to be your best.™

Ardent's Fast Memory

Meet's the Challenge of the Titan Graphics Supercomputer

WM. SPENCER WORLEY III AND WILLIAM S. WORLEY JR.,
ARDENT COMPUTER INC., SUNNYVALE, CALIF.

Scientific and technical applications in areas such as fluid dynamics, structural analysis, geophysical interpretation, molecular design, and astrophysics require significant computing power to develop results and to display those results in comprehensible form. The most effective approach for such applications has been to perform the necessary computations on supercomputers, mainframes, or minisupercomputers, and then to transport the outputs to a workstation that can present the results statically or dynamically as 2D or 3D graphics. The Titan Graphics Supercomputer integrates all of these capabilites into a single system, and was designed to provide the most interactive, productive, and economical platform for such applications.

Titan incorporates high performance vector floating point processors, as well as the fastest available RISC microprocessors (A Mips R2000), because both the applications and the 3D graphics displaying results involve intensive numerical computing. A system with vector processing capabilities requires very high bandwidth to memory, so that the data transfer rate does not limit its computational speed. The memory subsystem (Figure 1) that was developed to meet these needs is described together with some of the technical problems encountered during the design phase and the solutions that we adopted. Achieving the performance goals for this system required a sustainable data bandwidth to memory of 256 megabytes per second. Simulations showed that an effective solution consisted of two 32-bit address, 64-bit data, 16 MHzbuses, one of which is used only for reading vector data from memory, and the other of which is used for all writes and for other reads. These buses are called the R bus (read bus) and S bus (system bus) and are shown in Figure 2.

Bus protocols were designed and validated by simulation of the heavy loads expected during vector processing. These protocols permit each bus to operate at its peak utilization of eight bytes of data transferred every 62.5 nanoseconds. In turn, they place stringent requirements on the design of the memory subsystem. The following discussion describes these requirements, the problems thatemerged, and how the problems were solved.

■ REQUIREMENTS

The memory subsystem must respond to transactions presented simultaneously, each cycle, over both the R bus and the S bus. Transactions are defined for 64-bit double word transfers, 32-bit word transfers, partial word transfers, and synchronization and soft error correction operations. Each data transfer bus consists of a 32-bit address bus and a 64-bit data bus. For a given transaction, the address bus and the data bus are used at different times. For writes, there is a fixed delay between the time the address for a transaction appears on the address bus, and the time the data for the transction appears on the data bus. The response time for reads varies, though it is bounded by a timeout mechanism at the requester. Requests are tagged with a unique request ID so returned data can be routed correctly. All addresses and data have byte parity to ensure high reliability. Requests and request IDs also are checked by parity bits. In addition, handshakes are required to ensure acknowledgement of all transactions. Error signaling is always to the requester when the requester's identity can be determined unambiguously.

Arbitration is single cycle and is rearbitrated every clock cycle. By design, there is no contention on the bus, even with back-to-back transfers. No cycles are wasted for bus hand-off. A request may not be placed on the bus unless the target resource is free. Resource status for the memory subsystem is denoted by backplane signals, designating availability, for each memory interleave. Successive requests over either bus need not occur for adjacent interleaves. The vector units may operate with arbitrary address strides between successive requests.

In addition to responding to simultaneous requests over both the R bus and S bus, the memory subsystem must return data from read requests in the order the requests were issued. This eliminates the need for response tagging, or for enqueu-

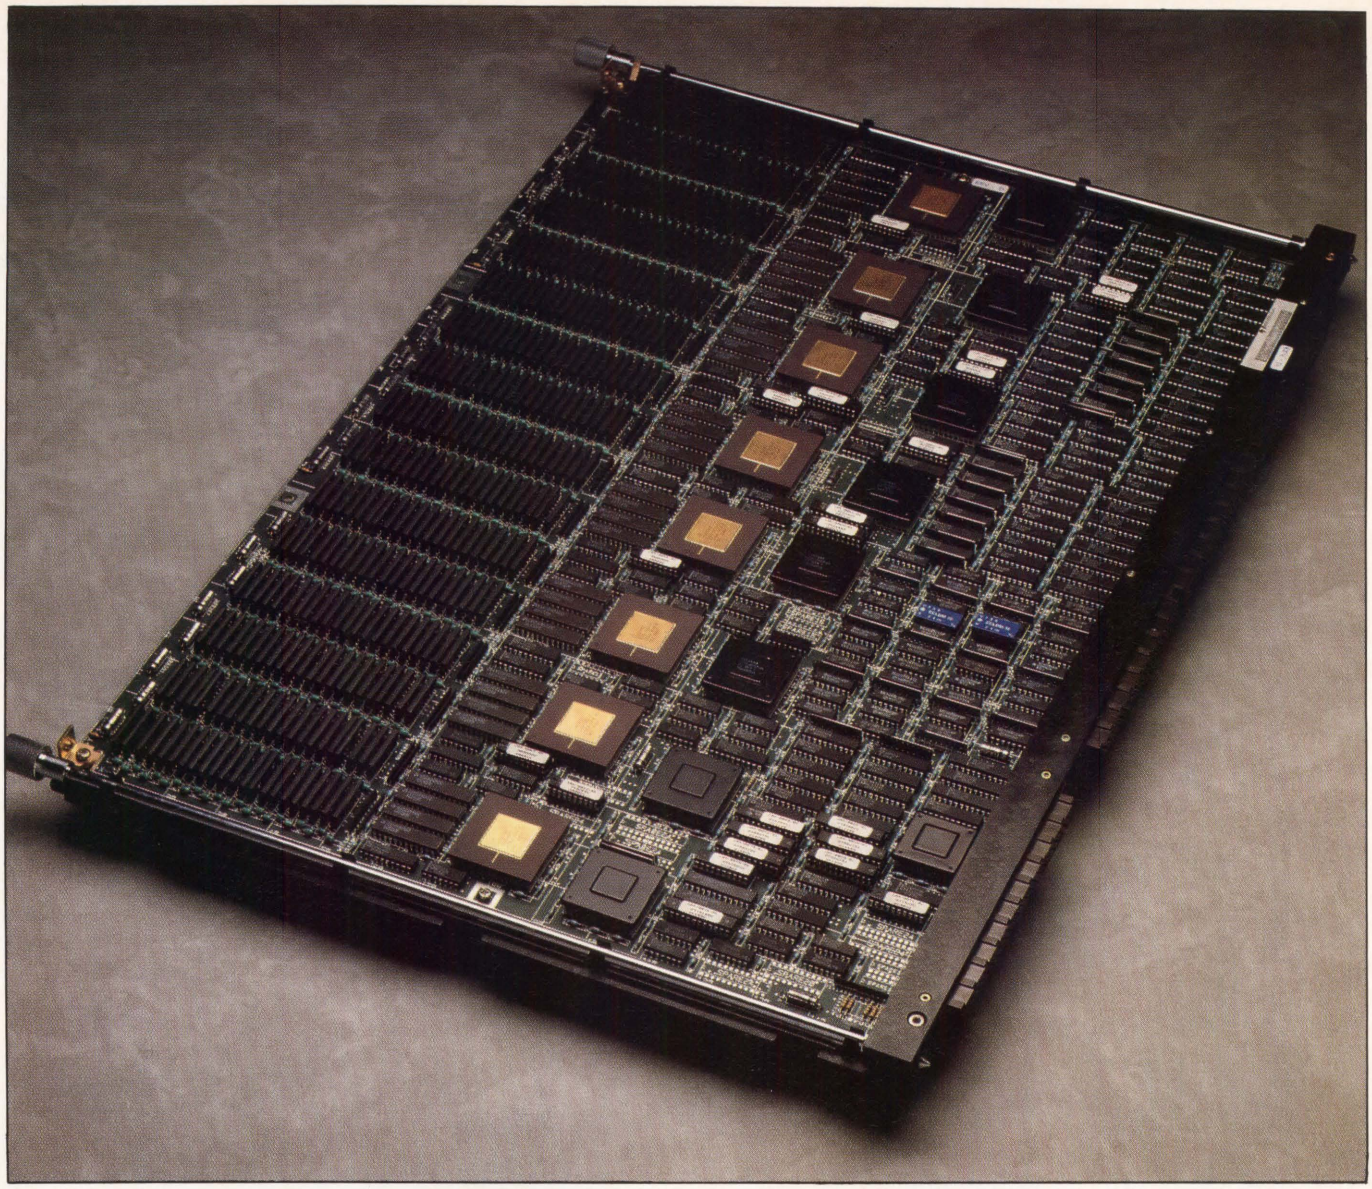

Figure 1. The memory subsystem of Ardent's Titan Graphics Supercomputer was designed to sustain the concurrent transfer of eight bytes per cycle over each of two 32-bit address, 64-bit data, 16 MHzbuses. One of the 32 megabyte memory boards is shown above.

ing returned data. This property must be guaranteed for single board or multiple board memory subsystems. The bus protocols also specify standard status and control registers: to control system configurability; to force errors and ensure testability, and to log soft and hard errors.

As a final requirement, the memory boards accommodate either one megabit or four megabit DRAMs. This permits larger capacity memories once the larger chips are available. For either choice of DRAMs, provision is required for data initialization, refresh control, single error detection and correction, double error detection, and single DRAM failure detection. The latter requirement stems from the use of 256K x 4-bit, or 1,024K x 4-bit, DRAMs. These chips are subject to a chip failure mode which manifests as a 1-bit to 4-bit failure in the data read from a single DRAM chip. Error correcting codes (ECC) used for 1M x 1-bit DRAMs cannot detect this class of failures. A new scheme is required.

■ A New Approach

Requests are presented to the memory subsystem over each of the two system buses every 62.5 nanoseconds. This provides enough bandwidth to support the integer units, vector units, I/O subsystem, and graphics subsystem. Memory is organized into the independent interleaves, on successive double word (8-byte) addresses allowing delivery of data to both the R and S buses every clock.

Bandwidth is obtained by handling successive requests by different interleaves. (As is well known, for N interleaves, where N is a power of two, increasing double word array dimensions to the nearest odd number assures that successive references along any array dimension will cycle through all available interleaves.) Each memory interleave contains one, two, or four megabytes of data for one megabit DRAMs. Or it can have four, eight, or 16 megabytes of data for four megabit DRAMs. Interleaves are organized so any access from the bus is serviced by a single interleave.

In order to minimize the space occupied by an interleave, and maximize the number of interleaves, the memory in each interleave is 32-bits wide. It is possible to access a 64-bit double word in two successive clocks by using static column mode DRAMs. This organization requires each interleave to be addressed on double word boundaries. In a sustained sequence of read operations, the data returned on the bus each cycle consists of bit 0-31 (high order word) for one interleave, together with bits 32-63 (low order word) from the next interleave.

Eight interleaves can be contained on a

Our ASIC line-up is only half the story.

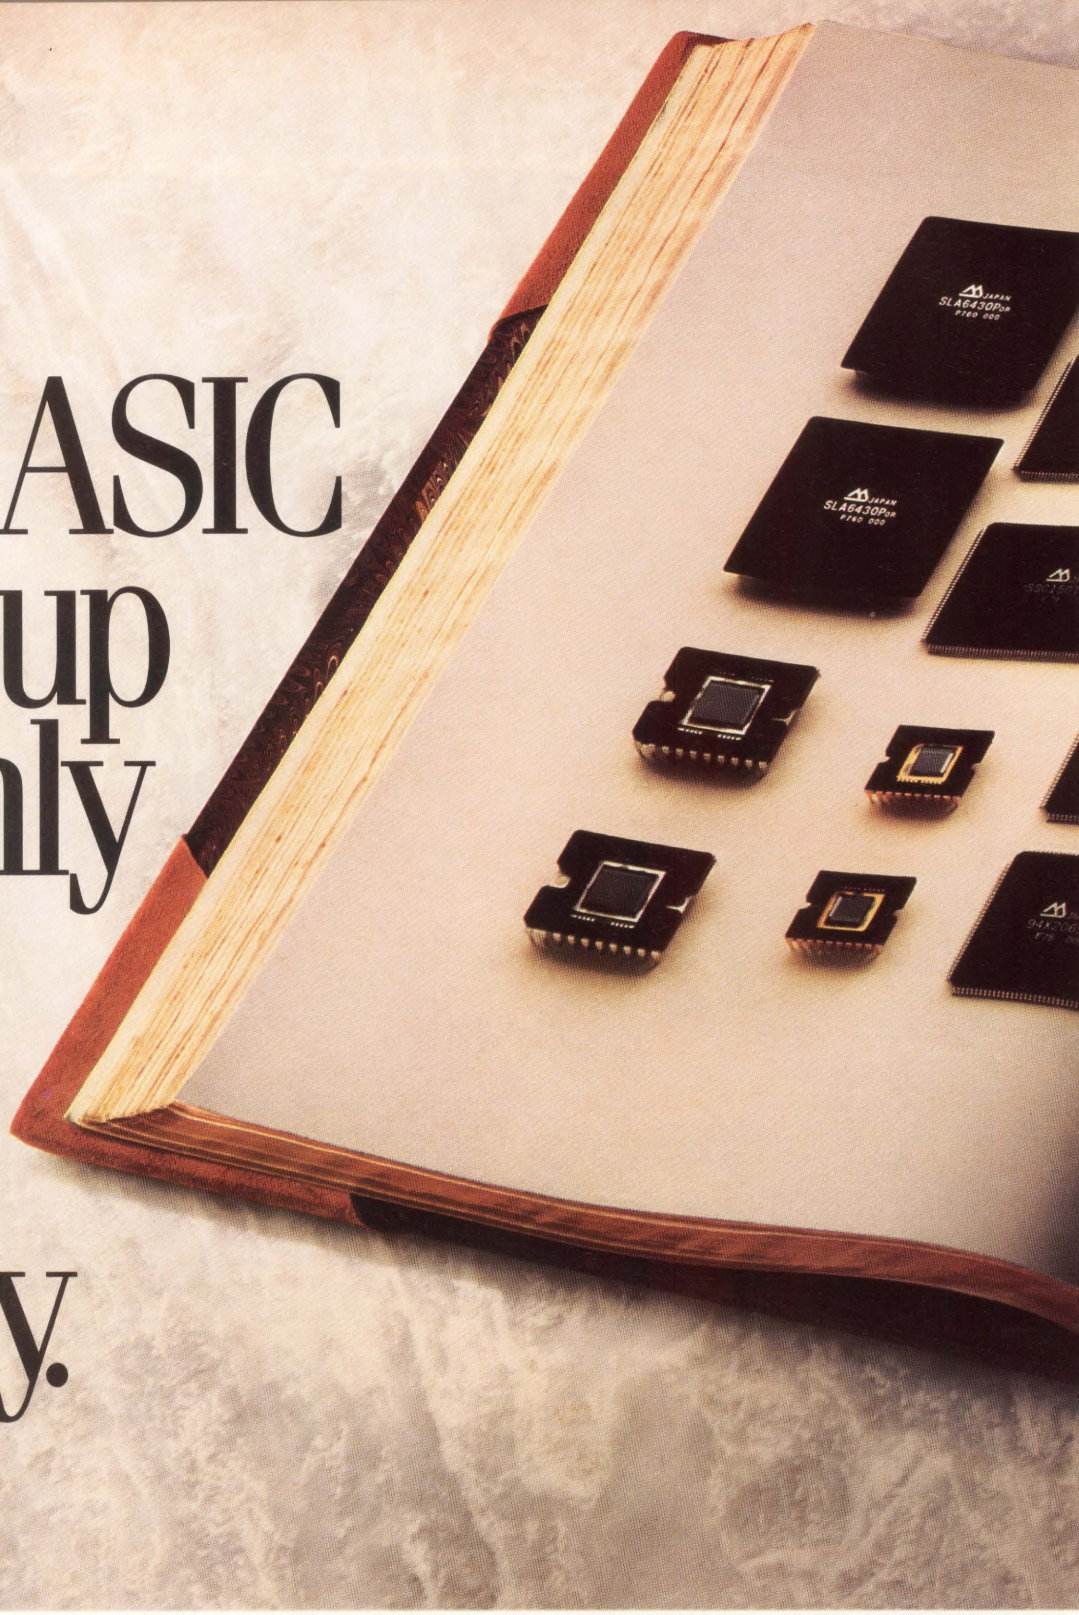

When it comes to delivery of high-quality, reliable ASICs, S-MOS wrote the book.

We did it in collaboration with our manufacturing affiliate, Seiko Epson. With 18 years of CMOS experience, Seiko Epson is one of the world's most advanced CMOS IC manufacturers.

Through Seiko Epson's high-yield manufacturing technology, we ship millions of ASIC units a month, and with a reject rate of less than .0001%. That's our quality story.

Now we've added a new chapter on design. At our advanced R&D design facility, engineers from S-MOS and Seiko Epson are developing new software to simplify circuit design, simulation and the creation of new megacells for our extensive cell library.

Of course, you can still take advantage of our established design tools because S-MOS supports such workstations as Daisy, Mentor, Calma and PC-based systems using FutureNet, OrCAD and ViewLogic.

Our proprietary LADS simulator will speed up the design process. The S-MOS engineering team will support you from concept to production.

If you are looking for an ASIC program that can make your designs best sellers, call us. (408) 922-0200.

CMOS Gate Arrays	CMOS Standard Cells	Compiled Cell Custom
Up to 38,550 available gates ▪ SLA8000 (800 ps)*, 1.2µ drawn, 1.0µ Leff. ▪ SLA7000 (1.0 ns)*, 1.5µ drawn, 1.2µ Leff. ▪ SLA6000 (1.8 ns)*, 2.0µ drawn, 1.5µ Leff. ▪ SLA700B High Drive Output ▪ SLA100L Low Voltage	Complexities to 16K gates** ▪ SSC1000 (1.4 ns)*, 1.8µ drawn, 1.4µ Leff. ▪ Fully migratable from S-MOS gate arrays ▪ RAM and ROM blocks available	*The* Alternative to Full Custom ▪ 1.8µ CMOS process ▪ Can utilize dissimilar cell geometries ▪ 3-button approach to custom design ▪ Currently over 300 fully characterized cells ▪ Fast 14-week implementation time ▪ Timing-driven TANCELL® place-and-route software

TANCELL is a registered trademark of Tangent Systems.
*Typical propagation delay of 2-input NAND gate driving 2 internal loads with 1mm of interconnect.
**Maximum gate utilization depends on amount of interconnect used.

S-MOS SYSTEMS

S-MOS SYSTEMS, INC.
2460 North First Street
San Jose, CA 95131-1002

CIRCLE NUMBER 10

single 16" x 22" board, and a complete memory subsystem comprises from one to four boards. Each memory board, therefore, is eight-way interleaved, and contains either eight, 16, or 32 megabytes using one megabit DRAMs, or 32, 64, or 128 megabytes using four megabit DRAMs. Whenever two memory boards of the same size are installed, they operate in pairs to provide 16-way interleaving. This organization provides memory sizes from eight megabytes to 512 megabytes, with 16-way interleaving for all practical memory sizes greater than eight megabytes.

Figure 1, which is a photograph of a 32 megabyte board, shows the zig-zag inline packages (ZIPs), containing the DRAMs, that occupy the leftmost third of each board. Each interleave contains two rows of ZIPs and is controlled by the adjacent 180-pin, 8,000-gate CMOS gate array. To the right of the interleaves are eight 120-pin, 2,500-gate CMOS gate arrays that constitute a cross-bar switch between the two system buses and the eight interleaves. At the lower right of the board is a 68-pin, 2,000-gate array thatholds the request IDs associated with data returned over the buses.

■ Read And Write Bus Protocols

The R bus and S bus protocols define transactions for reading and writing eight-byte double words, four-byte words, and the one-byte, two-byte, and three-byte subsets of a four-byte word. Transactions also are defined for a read-modify-write synchronization operation and for correction of data containing a soft error. Word transactions to reserved addresses, which can be issued only by the operating system, are used to address status and control registers on each memory board.

There are four status and control registers for each board: for the S bus, R bus, board ID, and LED indicator array. The bus status and control registers include provisions for logging bad address parity, bad data parity, invalid address alignments, invalid transactions, bad request parity, bad request ID parity, and errors detected within interleaves. In addition, the bus status and control registers provide means for identifying the type and capacity of the memory board, setting the number of interleaves and logical ID of the board, invoking diagnostic modes which can force bad parity into DRAMs or onto the bus, and enabling or disabling the entire board.

The board ID register can be read to determine the backplane slot occupied, the board type, and the board revision level. The LED control register permits software to control the OK (green) or fault (red) lights in the LED indicator array attached to each memory board. (Hardware directly controls six other lights in this array: S bus busy, R bus busy, 16-way interleaving, board enabled, ECC enabled, and fatal error.)

In addition to the board status and control registers, two status and control registers are provided within each interleave. These registers provide means for initializing the contents of multiple memory banks with a single write transaction, for logging soft and hard errors, for identifying a failing DRAM chip, and for forcing errors for diagnostic purposes. Forcible errors include: recoverable and non-recoverable data errors, refresh time out, requests inconsistent with arbitration, S bus timeout, and R bus timeout.

Memory is protected by an error correcting code with single bit error detection and correction, and full double bit error detection. Because the DRAMs used are organized four bits wide, this code also detects any single DRAM chip failure.

When an error is detected, the error is logged in the proper interleave status register, and the corrected data is returned on the bus. The corrected result is not immediately rewritten to memory. Instead, the operating system periodically issues a "scrub" transaction to each memory address. This causes any data containing a soft error to be corrected and rewritten. When any transaction requires a read-modify-write cycle, however, the corrected data is rewritten immediately.

Defining an economical error-correcting code with the necessary speed and coverage was one of the difficult design problems .

■ Design Challenges

Many areas of the memory subsystem were difficult to design. Three of the hardest dealt with the physical size and complexity of the board, correct data return ordering across multiple boards, and providing a fast ECC scheme with the required properties. Solutions to each of these problems were made more difficult by very conservative, worst-case design rules.

Although the 16" x 22" board was large, placing all the parts was difficult. This 14-layer board contains 18,000 holes, 4,000 vias, and 11,000 wires. It is populated by 352 DRAMs, eight 180-pin PGAs, eight 120-pin PGAs, one 68-pin PGA, over 300 discrete parts, and close to 1,000 capacitors. (When it's empty, the board is much like a window screen. One can see through it!) As shown in Figure 1, about one third of the area is used for DRAMs; one third is occupied by the crossbar switch, interleave controllers, and DRAM drivers, and one third contains the status registers, clock distribution, address control, return data ordering, error checking, bus latches, and bus drivers and receivers.

Other problems accompanied the large board. Clock distribution was tricky to manage. Backplane clock distribution is done by balanced ECL traces. Clock signals are converted from ECL to TTL by on-board clock drivers. Each driver was limited to eight loads, and the traces were closely controlled. With such care, no more than seven nanoseconds worst-case skew between any two clocks on the board is guaranteed. Losing 7-ns from a total budget of 62.5-ns posed challenges for the remainder of the design.

Dealing with long nets also was difficult. Some traces were as long as 24 inches, which added over 100 picofarads of capacitance to the loading of some drivers. The crossbar switch routes data and a half-dozen other signals between two 64-bit buses and eight 32-bit interleaves. Output traces from the switch vary from three to 20 inches in length.

Great care was exercised when doing timing analyses. Long nets with high capacitance are a problem for CMOS gate arrays, which, when heavily loaded, exhibit large asymmetries between rise and fall times. To meet the necessary timing constraints for data flowing from DRAMs, through an interleave control gate array, through the crossbar switch, and to the bus latches, it was necessary to invert the data path between interleave controller and crossbar switch gate arrays. This ensured that in all cases the timing included the sum of one rise time and one fall time. Otherwise, the worst-case timing would have had to include the larger sum of two fall times.

The use of mixed TTL and CMOS presented signal quality problems. Strong TTL drivers with close to one nanosecond edge speeds exhibited as much as three volts undershoot. Over ten nanoseconds after the incident wave, they would ring to over one volt. Similar problems occured for CMOS drivers, some of which exhibited sub-nanosecond edge speeds when driving TTL inputs. Although TTL inputs tended to clamp the signals better than CMOS inputs, and did a better job of limiting undershoot and ring, resistive termination was still necessary to meet manufacturers' undershoot specifications.

As discussed earlier, the Titan memory subsystem is required, on each system bus, to return all data in exactly the order in which it was requested. This must be

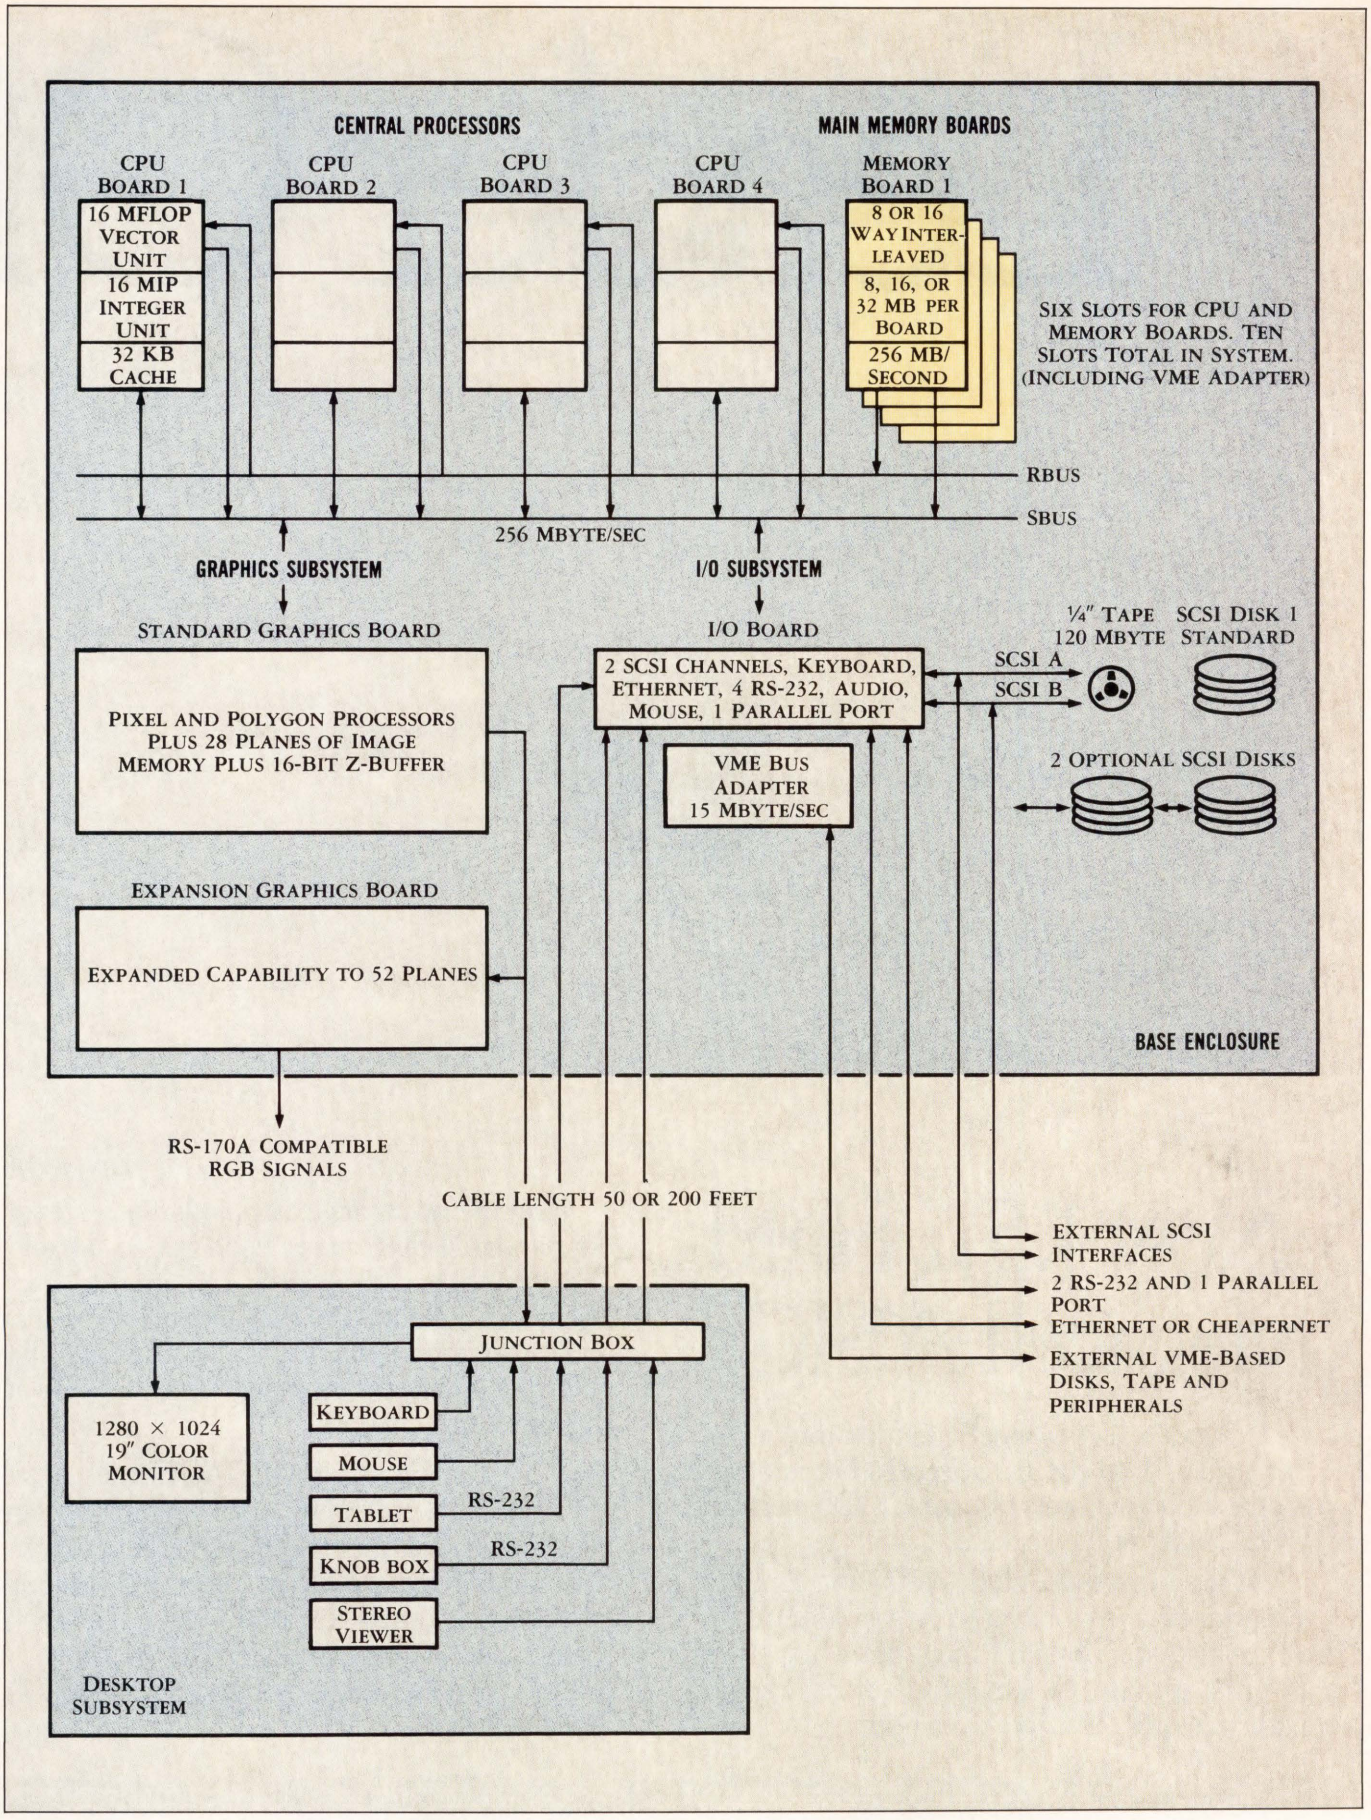

Figure 2. The Titan graphics Supercomputer is a 64-bit multiprocessor system with one to four CPUs, each with concurrent integer and vector floating point processors. The processors, I/O subsystem, and graphics subsystem communicate with the memory subsystem over two synchronous, 16 megahertz buses, each of which operates at 128 megabytes per second.

Leader Of

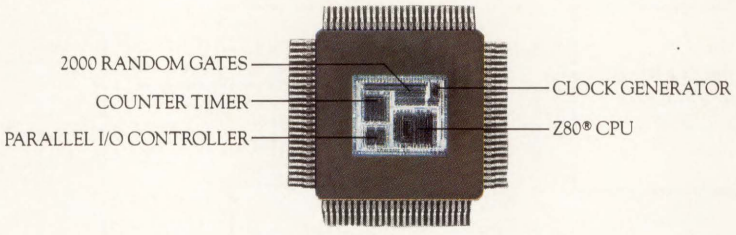

PACKAGE: 100 PFP, 32 MIL CENTERS*

®Z80 is a trademark of Zilog Inc.

Now 1.5μ Hard Megacell ASICs.

Nobody packs as much function into a megacell custom circuit as Toshiba because we've got the broadest library of 1.5μ CMOS megacells. We're Leader of the Packed.

We've been successfully producing complex megacell customs for over four years now. And we've shipped millions of them. So while others are just beginning their megacell efforts, we stand alone in experience and production.

Our megacells are exact mask duplicates of our standard LSI discretes. Each megacell is tested to our standard data sheet specifications. New layout is only required for the random logic section, and total circuit testability is always assured. We offer your application the highest complexity at the lowest risk.

Z80 FAMILY MEGACELLS

Z80. PIO. SIO. CTC. DMA. We have them all. And you can mix and match components with random logic to create integrated solutions for your most complex applications. Like the solutions we've already provided for modems, printers, hand terminals and industrial controls. To name just a few.

AREA SALES OFFICES: CENTRAL AREA, Toshiba America, Inc., (312) 945-1500; **EASTERN AREA,** Toshiba America, Inc., (617) 272-4352; **NORTHWESTERN AREA,** Toshiba America, Inc., (408) 737-9844; **SOUTHWESTERN REGION,** Toshiba America, Inc., (714) 455-2000; **SOUTH CENTRAL REGION,** Toshiba America, Inc., (214) 480-0470; **SOUTHEASTERN REGION,** Toshiba America, Inc., (404) 368-0203; **MAJOR ACCOUNT OFFICE,** Fishkill, New York, Toshiba America, Inc., (914) 896-6500; **BOCA RATON, FLORIDA,** Toshiba America, Inc., (305) 394-3004. **REPRESENTATIVE OFFICES: ALABAMA,** Montgomery Marketing, Inc., (205) 830-0498; **ARIZONA,** Summit Sales, (602) 998-4850; **ARKANSAS,** MIL-REP Associates, (512) 346-6331; **CALIFORNIA** (Northern) Elrepco, Inc., (415) 962-0660; **CALIFORNIA** (L.A. & Orange County) Bager Electronics, Inc., (818) 712-0011, (714) 957-3367, (San Diego County) Eagle Technical Sales, (619) 743-6550; **COLORADO,** Straube Associates Mountain States, Inc., (303) 426-0890; **CONNECTICUT,** Datcom, Inc., (203) 288-7005; **FLORIDA,** Sales Engineering Concepts, (813) 823-6221, (305) 426-4601, (305) 682-4800; **GEORGIA,** Montgomery Marketing, Inc., (404) 447-6124; **IDAHO,** Components West, (509) 922-2412; **ILLINOIS,** Carlson Electronic Sales, (312) 956-8240, R.W. Kunz, (314) 966-4977; **INDIANA,** Leslie M. DeVoe Company, (317) 842-3245; **IOWA,** Carlson Electronics, (319) 377-6341; **KANSAS,** D.L.E. Electronics, (316) 744-1229; **KENTUCKY,** Leslie M. DeVoe Company, (317) 842-3245; **LOUISIANA,** MIL-REP Associates, (713) 444-2557; **MAINE,** Datcom, Inc., (617) 891-4600; **MASSACHUSETTS,** Datcom, Inc., (617) 891-4600; **MICHIGAN,** Action Components Sales, (313) 349-3940; **MINNESOTA,** Electric Component Sales, (612) 933-2594; **MISSISSIPPI,** Montgomery Marketing, Inc., (205) 830-0498; **MISSOURI,** D.L.E. Electronics, (316) 744-1229, R.W. Kunz, (314) 966-4977; **MONTANA,** Components West, (206) 885-5880;

The Packed.

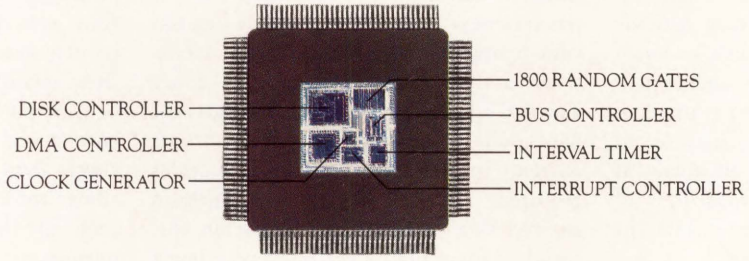

PACKAGE: 144 PFP, 25 MIL CENTERS*

*Die shown larger than actual size.

82Cxx PERIPHERAL MEGACELLS

We can supply all the necessary peripherals you need for PC and compatible environments. To communicate. To control disks. To access memory. To drive buses. To manage interrupts. They're all in our library.

SPECIAL PURPOSE MEGACELLS

CRT controllers. LCD drivers. UARTs and analog circuits. And RAM and ROM. Our special purpose megacells offer these kinds of solutions for your special needs. Solutions not available in other ASIC offerings.

GET A PACKAGE OF INFORMATION

Stop fighting the battle of packing more onto a PC board. Integrate the entire board into one Toshiba Megacell ASIC. And take over as leader of your pack. For complete details contact Toshiba today. Call your Custom IC Product Manager at (714) 832-6300 or a Toshiba Regional Sales Office: NORTHWESTERN: San Jose, CA (408) 244-4070. SOUTHWESTERN: Newport Beach, CA (714) 259-0368. NORTH CENTRAL: Chicago, IL (312) 945-1500. SOUTH CENTRAL: Dallas, TX (214) 480-0470. NORTHEASTERN: Burlington, MA (617) 272-4352. SOUTHEASTERN: Atlanta, GA (404) 368-0203.

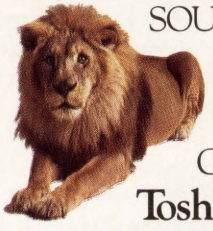

Toshiba. Leader Of The Packed.

TOSHIBA AMERICA, INC.

NEVADA, Elrepco, Inc., (415) 962-0660; **NEBRASKA,** D.L.E. Electronics, (316) 744-1229; **NEW ENGLAND,** Datcom, Inc., (617) 891-4600; **NEW HAMPSHIRE,** Datcom, Inc., (617) 891-4600; **NEW JERSEY,** Nexus-Technology, (201) 947-0151; **NEW MEXICO,** Summit Sales, (602) 998-4850; **NEW YORK,** Nexus Technology, (201) 947-0151; Pi-tronics, (315) 455-7346; **NORTH CAROLINA/SOUTH CAROLINA,** Montgomery Marketing, Inc., (919) 467-6319; **NORTH DAKOTA/SOUTH DAKOTA,** Electric Component Sales, (612) 933-2594; **OHIO,** Steffen & Associates, (216) 461-8333; (419) 884-2313, (513) 293-3145; **OKLAHOMA,** MIL-REP Associates, (214) 644-6731; **OREGON,** Components West, (503) 684-1671; **PENNSYLVANIA,** Nexus Technology, (215) 675-9600, Steffen & Associates, (412) 276-7366; **RHODE ISLAND,** Datcom, Inc., (617) 891-4600; **TENNESSEE,** Montgomery Marketing, Inc., (205) 830-0498; **TEXAS,** MIL-REP Associates, (512) 346-6331, (713) 444-2557, (214) 644-6731; **UTAH,** Straube Associates Mountain States, Inc., (801) 263-2640; **VERMONT,** Datcom, Inc., (617) 891-4600; **WEST VIRGINIA,** Steffen & Associates, (419) 884-2313; **WASHINGTON,** Components West, (206) 885-5880, (509) 922-2412; **WISCONSIN,** Carlson Electronics, (414) 476-2790, Electric Component Sales, (612) 933-2594; **WYOMING,** Straube Associates Mountain States, Inc., (303) 426-0890; **CANADA, BRITISH COLUMBIA,** Components West, (206) 885-5880; **ONTARIO,** Electro Source, Inc., (416) 675-4490, (613) 592-3214; **QUEBEC,** Electro Source, Inc., (514) 630-7486.

CIRCLE NUMBER 11

true regardless of whether the memory subsystem comprises a single board or multiple boards.

Conceptually, each board must contain, for each bus, a FIFO queue of all read requests for all enabled memory boards. Whenever a valid read request to any board occurs, an entry containing the interleave number must be made in every board's FIFO queue. Each entry also must record whether or not the request is for an interleave contained on the same board. The request ID for the transaction is stored in the 68-pin gate array. If, during the very same cycle, the interleave for the request which is first-in-line in the queue is ready to return data, then the data is written to the bus and every board must advance the next-in-line request to the front of the queue.

Though they seem simple, these requirements turned out to be very difficult to implement. During each cycle an entry must be introduced into all queues if a valid read request to any board is presented on the bus. During this same cycle, an entry must be removed from all queues if the first-in-line interleave (regardless of whether it's on the same or a different board) is ready to return data.

TTL memory components could not provide the necessary speed. ECL was required, as well as the necessary conversions between TTL and ECL signal levels. Even with ECL parts, the worst-case address and data setup times for the write into the FIFO queue leave less than a .5-ns margin. On the output side of the queue, the worst-case path (which was from an interleave ready to return data, through a multiplexer selecting the first-in-line interleave, across the backplane to notify all other boards to delete their first-in-line FIFO queue entries, including level conversions) provides less than a 1-ns margin. Even using ECL memory components there was little time to spare.

Titan is a heavily pipelined system. This strongly affects the design of the memory subsystem. It has the potential for up to 16 concurrently active memory transactions, and the entire set of transactions must proceed essentially in lockstep. Under such circumstances, error handling is difficult.

Each interleave control gate array begins its transaction during the first cycle after a memory request has latched onto a bus. In a piplined system, there is no time to delay or stop the request in order to handle an error.

For reads, return data ordering must be preserved, even if an error has occurred. As long as the request ID is known, data must be returned. If there has been an address error, or an unrecoverable data error, data with inverted parity bits in every byte are returned to the requester to signal the error. The data must be returned in the correct order, to the proper requester.

For writes, an error signal is derived that is used to block the write pulse to the DRAMs, and the error is signalled by an interrupt. In both cases, the pipeline advances every cycle.

The ECC scheme also must operate within pipeline constraints. As discussed earlier, the 256K x4-bit DRAMs have a failure mode that cannot be handled by traditional ECC methods. An acceptable ECC scheme not only must solve this problem, it also must not introduce delays into the pipeline. By the time the actual data has been read from the DRAMs, the interleave controller is committed to returning data on the very next cycle. There is no sequential time or any extra cycles in which to decide whether error correction should be done. The ECC scheme had to be fast enough that data could be corrected during every read, without increasing the number of cycles required. Worst-case timing analyses showed that in our implementation this could be met only by three levels or fewer of three-input exclusive ORs.

To meet these requirements, a new ECC scheme was developed. It is mathematically correct, and has been simulated extensively. The scheme detects and corrects any single bit error, whether in data or in check bits; detects any double bit error no matter how the bits are distributed, and detects the total failure of any four-bit-wide chip.

The details are lengthy, but the basic ideas and mathematical constraints are as follows:

Each megabyte is stored in 11 256K x 4-bit chips. Eight of the chips contain data, 32-bits wide, and the other three chips contain check information, 12-bits wide. (This is only one more chip than would be required for a traditional Hamming code over 32 data bits.)

When data is written, three sets of four check bits are generated and stored. Each check bit results from an eight-input parity tree.

For any three check bits, there is either one data bit or no data bit whose value affects all three check bits.

For any four check bits, there is no data bit whose value affects all four check bits.

For any data bit, there are exactly three check bits whose values are affected by the value of the data bit.

For all four data bits on a chip, the intersection of the four sets of check bits (where each set of check bits consists of those check bits whose values are affected by one of the data bits) consists of one distinguished check bit in common.

For all check bits on a chip, there is no data bit whose value affects the values of more than two of the check bits.

Meeting the criteria for the parity trees is sufficient to guarantee that any single bit error can be corrected, and that any double bit error can be detected. In addition to the final two chip constraints mentioned above, data grouping must be such that no chip contains both data and check bits. This ensures detection of a full chip failure.

Except for certain refinements, these constraints determine a practical ECC code. Because these conditions are abstract, the scheme can be more easily understood in the following manner: when data is read, both the data bits and check bits are fed to three sets of four nine-input parity trees (which correspond one-to-one with the three sets of four check bits). If the result of every parity tree is zero, then the data is correct and no error has occurred. If any single data bit fails, then there will be three "zero" bits and a single "one" bit in each of the three sets of parity tree outputs.

There are 64 such combinations. Only 32 can occur as a result of a single data bit error. The particular combination determines which data bit was in error and must be corrected. If a single check bit fails, then the parity tree outputs will consist of all zeros except for a single one bit in one of the three sets. This can be ignored. Any other bit configuration in the three sets of parity tree outputs indicates that an unrecoverable error has occurred.

METHODOLOGY AND RESULTS

Board placement and gate array schematics were done using the VALID S-320 system. Logical simulation was executed on Sun 2/160s and a Sun 3/260 using Gateway Design Automation's VERILOG mixed level simulator. Tools and discrete TTL and ECL component libraries were constructed to permit automatic translation of VALID netlists to VERILOG input. VERILOG logical models of DRAMs and of the interfaces to the rest of the system were added. Finally, a redundant, cross-checking functional model of the refresh logic was incorporated.

A test language was constructed that permitted scripts to be read simulating multiple processors issuing requests simultaneously on both system buses. Scripts were written to exercise all transactions and combinations of transactions, including error forcing and diagnostic fea-

tures. The VERILOG model was also extended to capture signal sequences at gate array pins. These signals were used to derive some of the test vectors for the gate arrays. Once the design had been simulated, the gate array netlists were translated to LSI Logic's LDS format, and "toggle test" coverage simulations were performed using the test vectors.

Once the test vectors were complete and the required test coverage had been reached, the designs were released to LSI Logic Corp. for layout and fabrication. The board used 14 layers, using 100-mil centers, and two lines per channel with controlled impedance. Board layout was performed by Shared Resources, and initial fabrication was done by DICEON Electronics, Inc.

The first pass boards were fully functional. After diagnosing problems due to incompletely seated DRAMs, the boards ran immediately at full speed. Subsequent testing showed that conservative design rules had resulted in good operating margins. In fact, beta-test system memory boards contain first pass gate arrays.

There were, however, several areas that needed improvement. The LED array had been installed backwards, and the signal quality problems discussed earlier required special termination. It was also necessary to improve the timing margins in the cross-bar switch and DRAM control to accommodate slower DRAMs. In addition, there was a logic error in the interleave control gate array that affected only privileged code and could be avoided by the operating system software. Each of these areas has been addressed in continuing engineering work.

For an effort of this complexity—a large, complex board, three gate arrays, tight timing, and mixed CMOS, TTL, and ECL technolgies—successful execution in less than 10 months by a design team of two engineers and two part-time consultants was satisfying. In the future we hope that better tools for logic test generation and coverage, timing analysis, and signal quality assurance will result in even better productivity. ■

ACKNOWLEDGMENTS

The authors wish to acknowledge the contributions to the development of the memory subsystem made by Mr. Jim Weatherford, Dr. Eitan Fenson, and Dr. Balasubramanian Kumar. Weatherford was a key designer of the memory subsystem. He is now with Sun Microsystems. Fenson developed a mathematical proof of correctness of the ECC scheme. He continues to serve as a mathematician-in-residence in addition to his other responsibilities at Ardent. Kumar made important contributions in validating, analyzing, and simulating the gate arrays and the complete memory subsystem. Dr. Kumar is an independent Silicon Valley consultant.

ABOUT THE AUTHORS

Dr. William S. Worley Jr., *vice president and chief scientist at Ardent Computer Corp., has been with the company since June 1986. Prior to this, he spent five years at Hewlett-Packard Co., where he led the development of the Spectrum architecture and system software. Dr. Worley also taught computer science at Illinois Institute of Technology and Cornell University. Inaddition he was involved in the design of systems and storage product architecture at IBM Corp. He holds MSs in physics and information science from the University of Chicago and a Ph.D. in computer science from Cornell.*

Wm. Spencer Worley III, *a senior design engineer at Ardent Computer Corp. and the son of William Worley Jr., has also been at Ardent since June 1986, and prior to that he worked as a design engineer for Hewlett-Packard Co. He holds a BSEE from the University of California at Santa Barbara.*

Apollo's Series 10000

The Tools and Technologies Behind a Personal Supercomputer

PAUL BEMIS, APOLLO COMPUTER INC., CHELMSFORD, MASS.

In early 1985, Apollo engineers began to explore new technologies that would allow them to develop a new class of computer. The goal was to put supercomputer performance under the desk. But when Apollo first began the design of the Series 10000 personal supercomputer, the engineering methodology did not yet exist that would permit the designers to meet performance goals quickly and accurately enough to hit the market window at the right price. The design of Apollo's PRISM (Parallel Reduced Instruction Set Multiprocessing) architecture and the Series 10000 is the story of the development of state-of-the-art CAE tools as much as it is one of advanced hardware and software design. This was the first time that both the hardware and software for a completely new system was to be fully designed and tested on a computing network. The challenge was to design a 20 MHz synchronous computer system that stretched the limits of CAE technology with over a million gates and new diagnostic methods.

■ THE PROBLEM

The nature of the problem that Apollo was undertaking demanded a new kind of development. Engineers recently had designed a CPU board with MSI logic and PALs and had debugged it in the laboratory by trial and error. This approach worked even with earlier versions of gate arrays because a designer could specify an individual ASIC and fix the logic around it by trial and error. But when Apollo engineers looked at the technology required to make the Series 10000 competitive, they realized that the system would be composed of very dense, complex ASICs and RAMs—and little else. This would be a silicon machine.

Laboratory work could show whether the whole system worked or not, but the entire design could not be fundamentally altered at this stage. Silicon of such sophistication is not easily modified by engineering change orders. The product would have to be delayed until the complex gate arrays were redesigned to fix the problems, and sent back to the foundry for another pass. The resultant delay would be unacceptable; for the project to succeed, the system would have to work on the first pass.

It was also decided that the system design and the complete system testing would have to be accomplished on a network that had substantial computing power, using advanced CAE tools that did not yet exist.

In the heart of Apollo's new system, most of the logic is implemented in very high density, 1.5 micron gate array technology. Other than DRAMs and bus drivers, the only other components are address drivers, to boost the power of the lines coming out of the chips. The processor uses no MSI or SSI logic functions.

In 1985, there were no bit-slice devices or standard microprocessors in the desired performance range. The development team had to build its own RISC design with hundreds of thousands of gates. Any mistake would require a redesign and a new chip from the foundry—an extremely costly proposition in terms of time as well as money. Consequently, Apollo set a basic objective to get the design right the first time, and verify it before a single piece of hardware was ever built.

In a traditional test of integrated circuits, the engineer who designed the circuit provides a set of test vectors for a tester to use to compare the actual outputs with the expected results. For ICs up to 10,000 gates, this is an acceptable way to test integrated circuits.

■ TESTABILITY CRUCIAL

As the circuits become more complex, however, the ability to develop test vectors that adequately cover the internal logic deteriorates dramatically. Even with older, less dense ASICs, designers are lucky if test vectors cover 90 percent of the potential logic and fabrication problems on a chip. But with 30,000 to 40,000 gate devices, such as Apollo uses in the Series 10000, it is hard to get even 50 percent coverage.

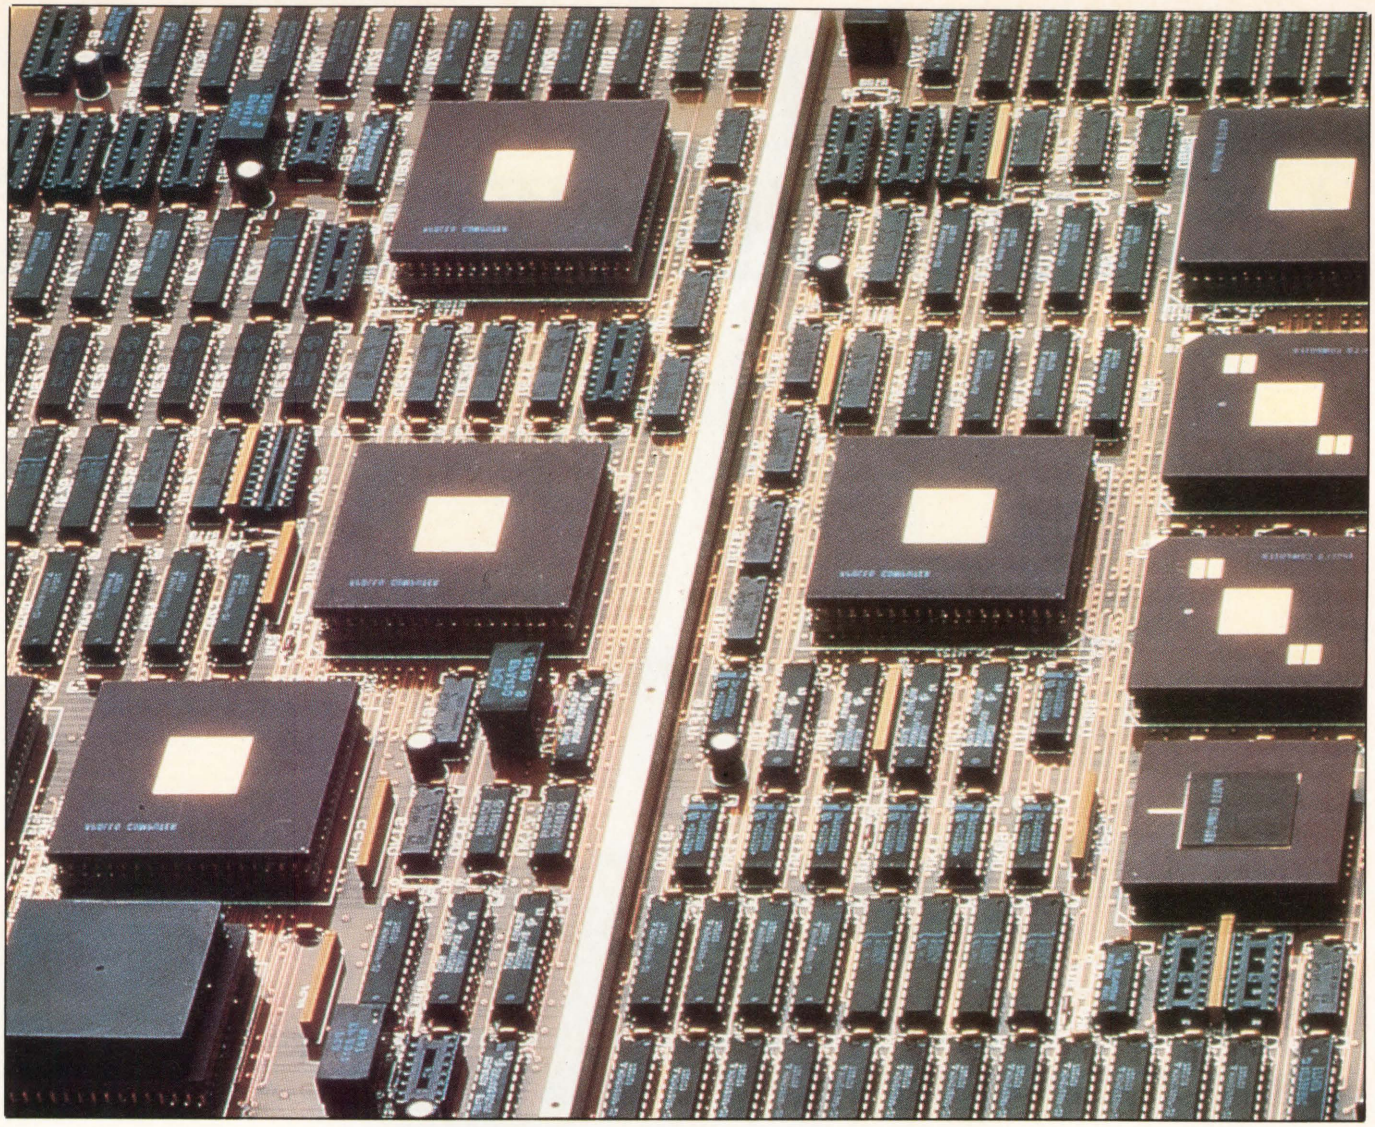

Testability was a crucial issue that had to be resolved by the Series 10000 design team. Testing the four to six thousand gates in older ASIC designs wasn't too difficult. When chip density becomes high enough to bring the control logic on-chip, however, the IC becomes a full sequential state machine with a limited view of the world from the outside pins. Then broadside testing is no longer practical. Consequently, one of the design constraints was the testability of the chips.

It is very difficult to generate correct test patterns for very large gate arrays because in-circuit testing, which is popular for testing printed circuit boards, is not possible. But scan technology quickly generates test patterns and provides very high levels of fault coverage. Scan testing is the only way to accurately test gate arrays of 20,000 or more gates to the gate level.

This method, originally dveloped by Honeywell and IBM, interested Apollo's engineers because of its potential for testing large devices. Although scan testing is not new, the use of discrete components such as microprocessors and bit slice in earlier system designs precluded its use. But the nearly pure ASIC design of the Series 10000 is custom made for this technique. Apollo chose a version similar to IBM's LSSD (Level Sensitive Scan Design).

The decision to solve the problem by implementing full scan testing (for reliability), placed yet another requirement on the tool set and the design team.

The vendors selected had to support scan test circuitry. Integrated CMOS Systems (ICS), the ASIC vendor selected, is not very large or well-known. However, the company had mainframe experience and its product line supported the scan path. ICS is a spinoff from a program that Storage Technology Corp. (STC) ran for a number of years. A group involved in the program left STC, licensed the tools and technology, and started ICS.

Another group from the same program left to form Aida Corp., which designed the next generation of CAE simulation software. Aida provided the timing verification software and hardware accelerators for the project. Mentor Graphics Corp.'s CAE tools were used as the heart of the system because the team needed a solid base of familiar design tools from a company that they could work with closely throughout the program.

Even working with vendors familiar with the problems of this type of design and the design tools necessary to solve the problems, the relationships between vendors' tools and Apollo tools grew very complex and had to be carefully monitored and controlled. Figure 1 shows how complex this relationship was between the Apollo and ICS tools for just one part, albeit an important part, of the entire design process: ASIC design verification and chip release.

The Aida software traces through the logic, determines the state equations for each of the gates and generates exhaustive test coverage for the faults in between register sets. It supports scan testing: its automatic test-pattern generator (ATG) examines the logic between the flip-flop

PHASE LOCK LOOP.
ANALOG/DIGITAL ASIC.

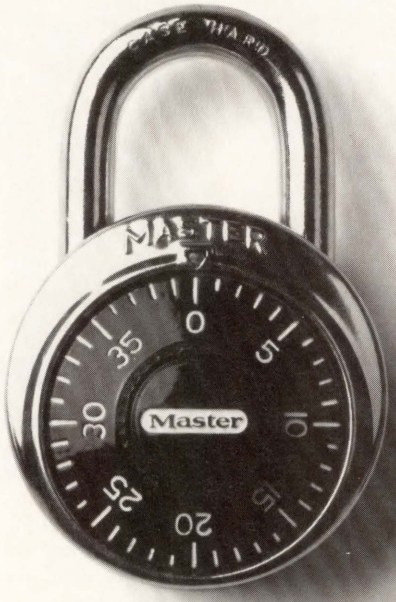

NO ONE ELSE HAS THE COMBINATION.

System level designers can access new levels of integration designing products with a combination of advanced analog/digital circuitry on a single ASIC. Obtaining the smallest form factor, lowest power, and highest performance in data capture applications is now possible.

Consider our 34MHz Phase lock loop subsystem – the PLL34M – an ideal solution for data separators.

The PLL34M is an open loop phase lock subsystem containing a voltage controlled oscillator, digitally controlled charge pump, trimmable reference voltage, and high speed driver I/O circuitry.

Combined with our other analog and digital standard cells, the PLL34M is ideal for increasing system integration of disk drive systems. (See application diagram.)

SIERRA integrates PLL onto a single data separator ASIC.

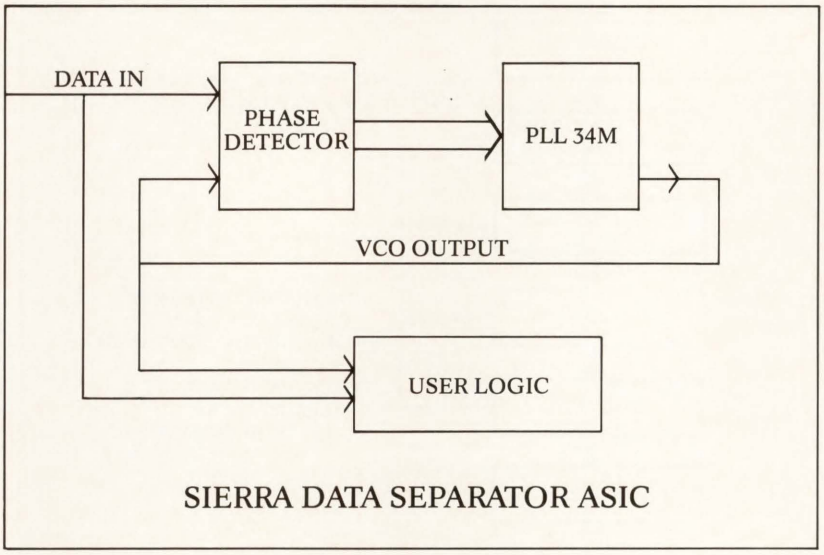

For designers seeking to incorporate analog processing functions, such as amplification, filtering, and data conversion, Sierra offers a wide range of solutions for applications in disk and tape drive systems, local area networks, and servo controllers.

Lock into our CMOS standard cell library which includes over 250 digital standard cells, 50 analog standard cells, and 20 EEPROM standard cells. Our Triple Technology™ gives you the flexibility to combine analog, digital, and E^2 on the same chip.

Sierra reduces packaging size and cost by offering low profile, surface mount packages including SOIC, PLCC, and PQFP Gull wing packages.

CMOS analog standard cells are available now in the same design environment as digital. Importantly, on-site mixed analog/digital simulation is easily achieved through access to our MIXsim™ software tools, which verify design functionality and assure first time silicon success.

Access the unbeatable combination of analog/digital ASIC from Sierra. Lock into the latest data on our CMOS standard cell library and provide your customers with the most advanced products. Just write or call today for our complete library card.

Sierra Semiconductor
Triple Technology™ In CMOS.

2075 North Capitol Avenue, San Jose, California 95132. Telephone (408) 263-9300

CIRCLE NUMBER 14

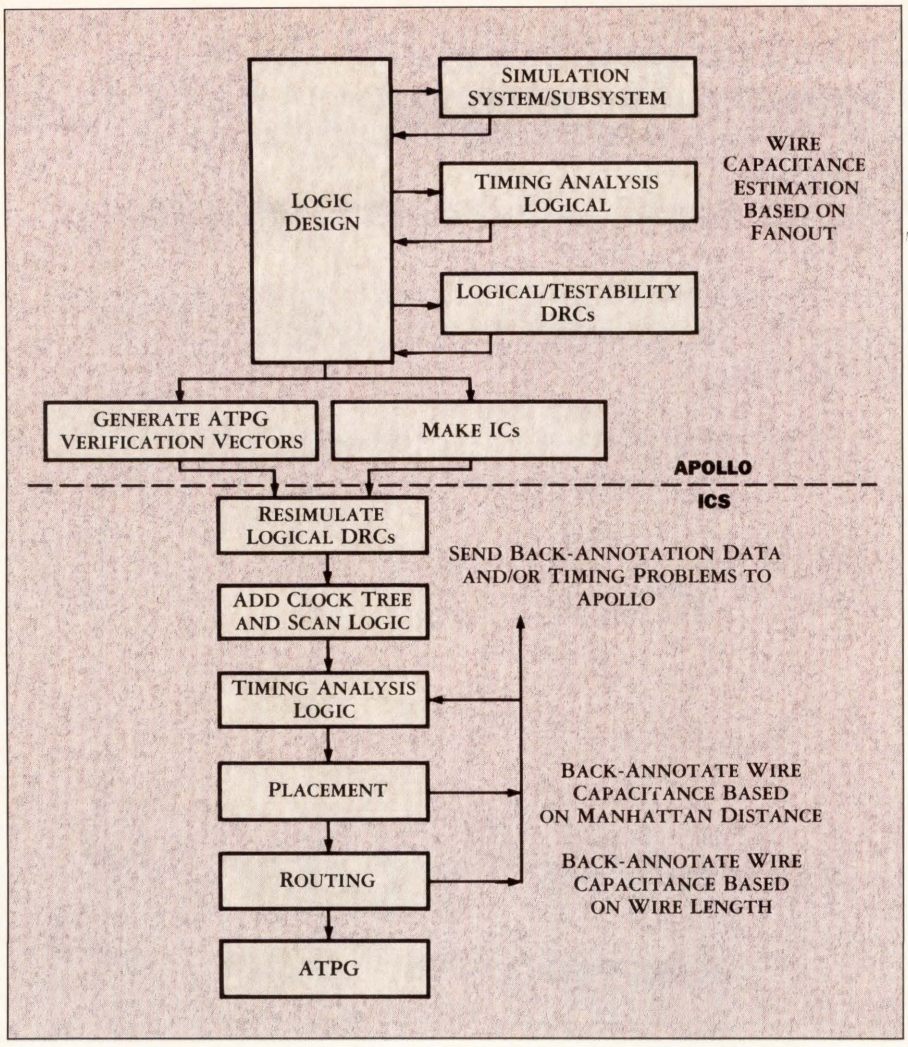

Figure 1. The diagram of the relationship between Apollo Computer and Integrated CMOS Systems design automation tools shows the complexity of just one important part of the entire design process: ASIC design verification and chip release.

sets, and indicates whether all possible faults are covered by the test patterns. The scan method allows diagnostics at the gate level to be performed by junior level operators on the manufacturing floor.

The system designers also needed to perform very accurate timing analysis because performance was critical for the Series 10000. They felt that static, path-sensitive, path-directed, timing analysis was the best method to use for a synchronous system. This method works thoroughly on chip level designs and checks all critical paths. For a full year after the logic was laid out, the team did timing verification with Aida's tools before fabricating any silicon.

■ SCAN TESTING

With the scan path test method, all of the flip-flops in a gate array are hooked together in a shift register to allow test vectors to be serially shifted in or out. This permits easy control of all internal nodes. The scan can be either a boundary or external scan. It essentially lets sequential logic be tested as combinational logic.

Another aspect of this test technology is that a large shift register of all the internal flip-flops can be used as a ring oscillator to calibrate the performance of the chip. CMOS semi-custom devices often have wide variations in timing. With the number of designs in the Series 10000, and the need for them all to work together, it was nccessary that all timing paths, whether obvious or subtle, be accounted for in advance.

The problems can be complicated further by process variations across an individual chip. The circuits at one end of a chip may differ in speed by as much as 20 percent from the circuits at the other end of the chip. In CMOS, logic speed is related to voltage, and voltage drops across a circuit board are related to temperature. Consequently, a designer has to worry about operating environments, including temperature variations, across the circuit board. Scan path was the enabling technology for a new, working high-performance processor.

Scan testing also has its drawbacks. Although it can provide much higher fault coverage of the chips, it requires a synchronous design with no asynchronous circuitry, such as feedback circuits. Scan testing also requires that additional real estate be invested in flip-flops. These flip-flops, designed to deliver high enough performance to minimize their effect on the normal operating logic circuitry, are usually 50 percent larger than normal.

TO KEEP THE DESIGNERS' NEEDS FROM OUTGROWING THEIR WORKSTATIONS, 30 SERVERS WERE ADDED

Scan testing builds confidence in individual chips, but also permits the computer to test itself. The test method impacts the manufacturing and repair center test process. In fact, providing the ability to perform scan testing loads the front end of the entire development cycle. It lengthens the early design stages, but dramatically reduces the time spent on verification. Unfortunately, most standard integrated circuits are not built with this type of self-test. As a result, most systems using these devices cannot incorporate scan testing.

■ NETWORKING

Functionality and timing analysis were the two key activities that tools had to support in this complex, multi-chip design. Apollo engineers specified all individual chips that they planned to design before starting the actual design.

Each designer used a workstation with a Mentor tool set for schematic capture. The designers frequently checked the functionality of chip segments by using Mentor's simulation capabilities. They could, for example, make sure that the ALU worked, and sequence a state machine through its various states. Or they might put the IP and MMU together and then add the caches to see if the design worked.

In the CAD environment, each time an engineer put designs together to see if they worked, extra disk space and computer power were required. With enough pieces in a simulation, a designer can quickly outgrow a workstation. To resolve this problem, Apollo designers networked about 30 server nodes.

These included Apollo's Turbo graphics workstations with floating point accelerators running Aida's Cosim simulation ac

Hardware

The Series lOOOO's 64-bit processor is implemented in 10 gate arrays, two floating point chips, and a few other parts, on a 3 x 6 Eurocard with a set of connectors on one edge. These pins are the one connection from the board to the "X bus", the Series 10000 system bus. The design required a very high bandwidth bus. No commercially available bus could meet the specification, though the Series 10000 supports both AT and VME buses for standard peripherals.

Common physical memory is a globally shared resource of all four processors and I/O. To lock memory locations being used for one processor, there are specific locking mechanisms at the bus protocol level that are unique to locked traffic. Since setting of the lock should not block all normal traffic, the bus protocols know when to do lock references and non-lock references.

Many of the devices used by Apollo are not pure gate arrays. They are single chip combinations of a 22,000-gate array and a complex register file of approximately 8,000 gates, which is not built in a gate array structure. The gate array vendor, Integrated CMOS Systems (ICS), uses VLSI Technology (VTI) and Toshiba as foundries.

The X-bus connector has a set of gate arrays associated with that interface. There is an address chip (bus address unit) and a pair of data chips (bus data units) for the even and odd halves of the data words. The data chips also contain the write-back queue for the caches.

The instruction and data caches, which are more or less symmetric, are fed from the bus interface units. They are controlled by the MMU aboard one of the large gate arrays that contain 30,000 gates, and includes the register file. The bus data units carry the 64-bit data words from the bus to the processor. Data can be loaded into the data cache or into the instruction cache, depending on whether an I-string reference or a D-string reference accessed the memory. The Harvard architecture with separate instructions and data caches is a fairly common way to run one instruction per cycle, a goal of RISC architectures.

Two of the ICs (DITS and DOTS) that interface to the bus address unit are used for multiprocessor support. They are used for duplicate tag store in and out. In a multiprocessor environment, the system must carefully watch the writes to the bus. If the writes collide with the contents of the data cache, then the system must invalidate or update those caches.

A CPU would waste much of the processor-to-cache bandwidth if it had to be checked for every single write command on the X-bus to see if the data was in the cache. So the system makes a duplicate copy of the tag stores of the two caches. If the data is not there, which is common, then the system needn't bother the CPU. A cycle is stolen from the processor to invalidate the cache only if the data is found.

The MMU performs the tag store comparison for the two caches. This is more complicated than it might seem, because the caches are indexed by virtual addresses while their tags are matched by physical addresses. The MMU also contains the logic neeeded to translate virtual addresses to physical addresses.

The MMU is built from the same master slice type of gate array as many of the other chips in the system. This array can be organized differently according to its use (for example, 32 by 32 or 16 by 72).

The register file is used by the gate arrays in various ways. The MMU, for example, uses the register file as part of the primary translation buffer memory while the integer processor uses it as part of the 32 by 32 register set that is an integral part of the architecture. The floating point ALU and multiplier are available from Bipolar Integrated Technologies. The FPC (floating point control), a smaller 13,000 gate array without memory embedded in it, functions as controller for the floating point pipeline. This smaller array is also used for other functions.

The entire system is synchronous. From the integer processor, through the caches and the bus interface, to the memory and back, the system is one massive 18 megahertz pipeline. ■

celerator. The designers could now glue different pieces of the design together and test them in a simulated environment without exceeding the limitations of their "network" system.

As the chips were finished, the designers shifted to Aida simulation. Aida offers a set of software tools and a hardware simulation accelerator that plugs into Apollo workstations and can simulate logic many times faster than a workstation running the software alone. The design is first fed to a tool that generates the appropriate directives and is then loaded into the simulation engine.

Apollo dedicated 17 engines to the program. Each one can handle up to a half million gates.

Although the high-level simulator in the tool set could simulate the gross behavior of the system, the designers did not do a strict hierarchical (top down) design. They kept the high level simulator as a separate entity. The actual gate design was surrounded with functional models of its environment. A cache model, an MMU model, and a number of other models allowed work to proceed in parallel.

By the time the design was finished, the engineers were doing gate level simulations of four processors tied together with the memory and I/O systems. A fully loaded Series 10000, simulating hundreds of thousands of gates, was handled with the help of the hardware acceleration engine. (A simulated system runs much slower than an actual system. In simulation, one 50ns clock cycle occurs every five seconds—or 100 million times slower than realtime!)

Each designer had his own workstation plus file servers on the network. In addition, the team had compute and acceleration servers on the network. For just the core of the hardware design staff, 67 workstations were used; 25 of them were file or compute servers. As the complexity of design increases, so does the requirement for computer resources.

The entire system design, from documentation down to the last gate of the system, used the network. Using a MCAD package, mechanical drawings were viewed on the network as easily as a Mentor schematic. Everyone was also linked through the same network for mail and updates on electrical design, mechanical design, and technical problems. Even documentation was reviewed on line.

Figure 2 shows the complex group of

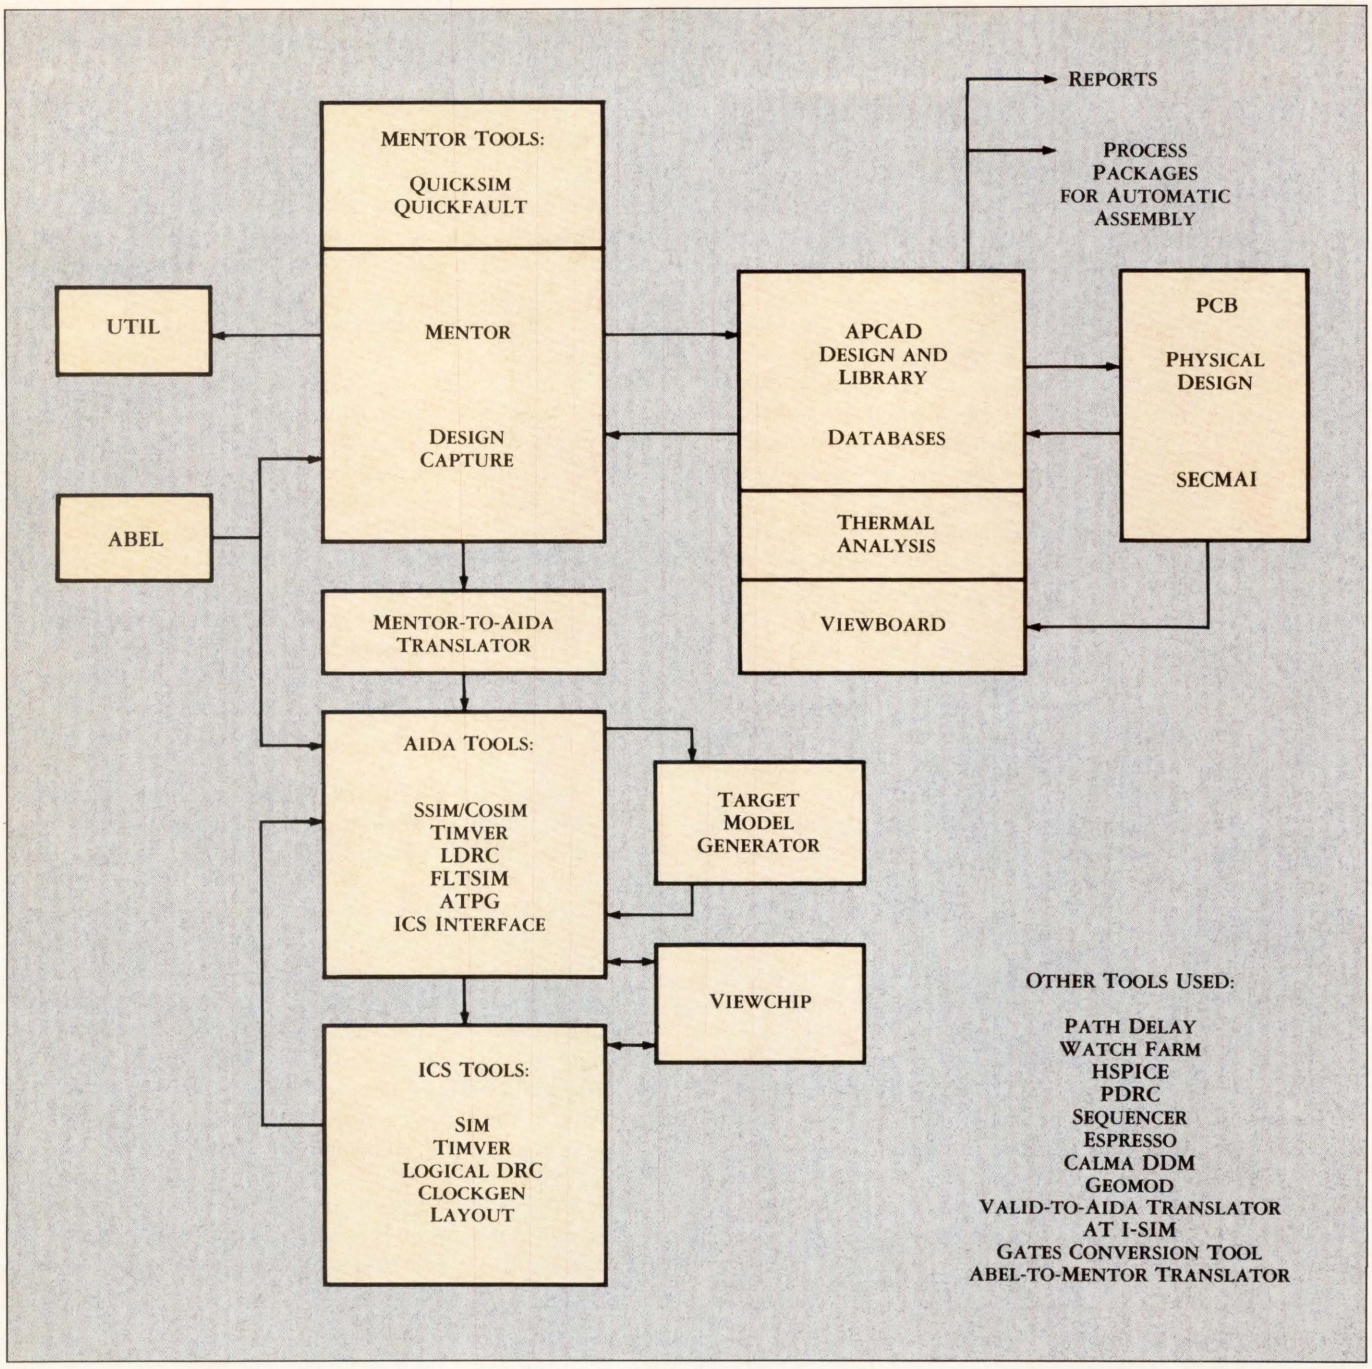

Figure 2. A complex group of design automation tools from different vendors was deemed necessary by Apollo designers in order to successfully design the Series 10000 Personal Supercomputer. The diagram shows the interrelationships of these heterogeneous design tools.

design tools from different vendors that Apollo engineers found necessary to design the Series 10000. The diagram shows the interrelationships of these tools. Additional details on the CAE system used for this project are given in the "Tools" sidebar.

The network also had to deal with the complexity created by numerous programmers developing software for the Series 10000. The CASE group handled that by using Apollo's program, DSEE (Domain Software Engineering Environment). DSEE looked at all of the software modules, identified whether or not they had been changed recently, and maintained a list of all the programs that were required to compile a package.

DSEE also kept track of what had already been compiled, what hadn't been compiled because of changes, and what was out on loan for modification. This horrendous complexity of software was compiled into one runnable object to develop the operating system.

■ BUILDING MODELS

Early in the project the designers built models of the I/O subsystem and memory traffic. They also made a model of the instruction processor to measure the power of the instruction set. These high-level models, written mostly in Pascal, described the hardware data path and ran instructions through it to determine what could be done with a reasonable amount of hardware. The MMU was also modelled to determine associativity versus total size.

At that phase the machine was partitioned into specific pieces. Designers were assigned to each piece, and helped write a very detailed specification. For months, Interleaf was the main design tool as detailed descriptions of each piece of the machine were generated. Next, the designers performed schematic capture and started drawing diagrams using Mentor tools. Each designer worked independent

Tools

The Series 10000 hardware design required some of the most advanced CAE tools available to simulate more than a million gates and make certain that all of the ICs in the entire system would work on the first pass. A team of 20 or 30 hardware designers used these tools on their workstations to design their subsystems. The designers also used them to communicate with each other and pull the pieces of the design together into a complete system model. That model boosted the designers' confidence that the design would work when it was cast in silicon.

Simulation was speeded up by Aida hardware accelerators on the workstations. The accelerator is a small RISC machine that is tailored precisely to the application.

The diagnostic staff was taken away from its usual job, writing test programs for functional hardware, to write test programs for logical hardware. These tests ran on parts of the design as it grew. The engineers added more gates, and made more realistic models with more pieces joined together. Finally, they were able to work with an entire CPU board, memory board, and I/O channel. All four CPUs, memory, an I/O channel and a graphics board eventually were running at gate level. Well over three quarters of a million functioning gates were needed to run the test program.

Another significant advance made by the Series 10000 project was the ability to simulate detailed timing accurately. The chips must be functionally correct, but the timing also must be correct.

The team used a very complete timing analyzer. The designer feeds the net list to the timing analyzer, which examines the flip-flop locations, traces all possible combinations to the destination flip-flops, and indicates how long any given path takes. The designer still has to decide which passes are valid and which are possible but functionally uninteresting. Aida performs timing internally (flip-flop to flip-flop inside the chip) as well as externally (chip to chip). Delays are measured to tens of picoseconds. Consequently, the team is confident that the boards will work without path delay problems.

The tools used (Figure 2) permitted large circuit simulation, static timing analysis, automatic test generation for scan testing and fault simulation. Since Aida's tools were tuned to solving problems that arise in large, CMOS-based synchronous systems, they were a good match for what Apollo engineers were trying to accomplish. Aida had also worked with the IC vendor. There was, therefore, a synergism among the IC vendor, the toolset, and Apollo's design goals for the project.

Mentor's tools were used for design capture, and fed into PCB design and manufacturing. Using Dialogue, Apollo's user interface creation tool, the engineers wrote an interface from Mentor to Aida, and rewrote a primitive timing analysis tool supplied by ICS, the ASIC vendor. Apollo already had a program that could check design rules.

Although gate level simulators are not new, their capacity has been limited to the single chip level. They have not been able to simulate tens of thousands of gates in multiple chips on a board, all working together at the same time. The Aida tool set is the first commercially usable package with that capability. As one of the first customers for the tool set, Apollo engineers quickly learned what worked and what didn't. They worked closely with Aida to improve the tools as they used them.

In the end, Apollo had a working machine and Aida had a very solid tool package. ■

ly until the design was tested to find out whether it all worked together. It was time to simulate the entire system.

The Aida simulation environment also offered timing verification, automatic test pattern generation, logical design rule checking, fault ratings and other facilities that were not needed until later in the design. With these tools the designers could build macrofunctions such as ALUs, shifters, and other major pieces of the data path, and then simulate them to make sure they performed their respective functions. The designers were able to build a complete hierarchy that integrated these pieces into an entire data path that could be cycled from the bottom up.

The high level models described entire subsystems. They provided a clear understanding of the implementation before proceeding with the actual hardware design. The designers could discover whether an approach would work before going to lower levels of the design.

The Aida programming language, ADL/Callows a portion of the gates to be abstracted into the C-like language for high level, abstract models. These models can be mixed with actual gate implementations and run as part of the simulation.

■ Conclusion

The Series 10000 design project started out to create a new class of computer (also see the "Hardware" sidebar) that would combine supercomputer performance with the convenience and versatility of an interactive graphics workstation. To do this required the use of the latest ASIC technology. To use that technology effectively, the design team had to find, develop and apply the most advanced design techniques and tools available. The result, in addition to the design of a new class of computers, was the development of powerful CAE techniques that will make it easier for Apollo to design new systems.

The project proved that system level simulation works and that static timing analysis is the right way to do synchronous design. The project also showed that scan works at both the ASIC and PCB levels.

Close integration of CAE, design teams and ASIC support were necessary for the best results. The networked computing resources provided the considerable power required for the project, and also served as a platform for the design integration. To be competitive, the new system had to be designed and built with dense ASICs—a mandatory process for today's designs. The design process has now been proven, and it works. ■

About the Authors

PAUL BEMIS, *product marketing manager for the DN10000, joined Apollo in 1986. Previously, he worked at Adage Inc. designing and marketing graphics displays. He holds an MBA (Northeastern University) and a BSME/EE (University of New Hampshire).*

Easy ASIC

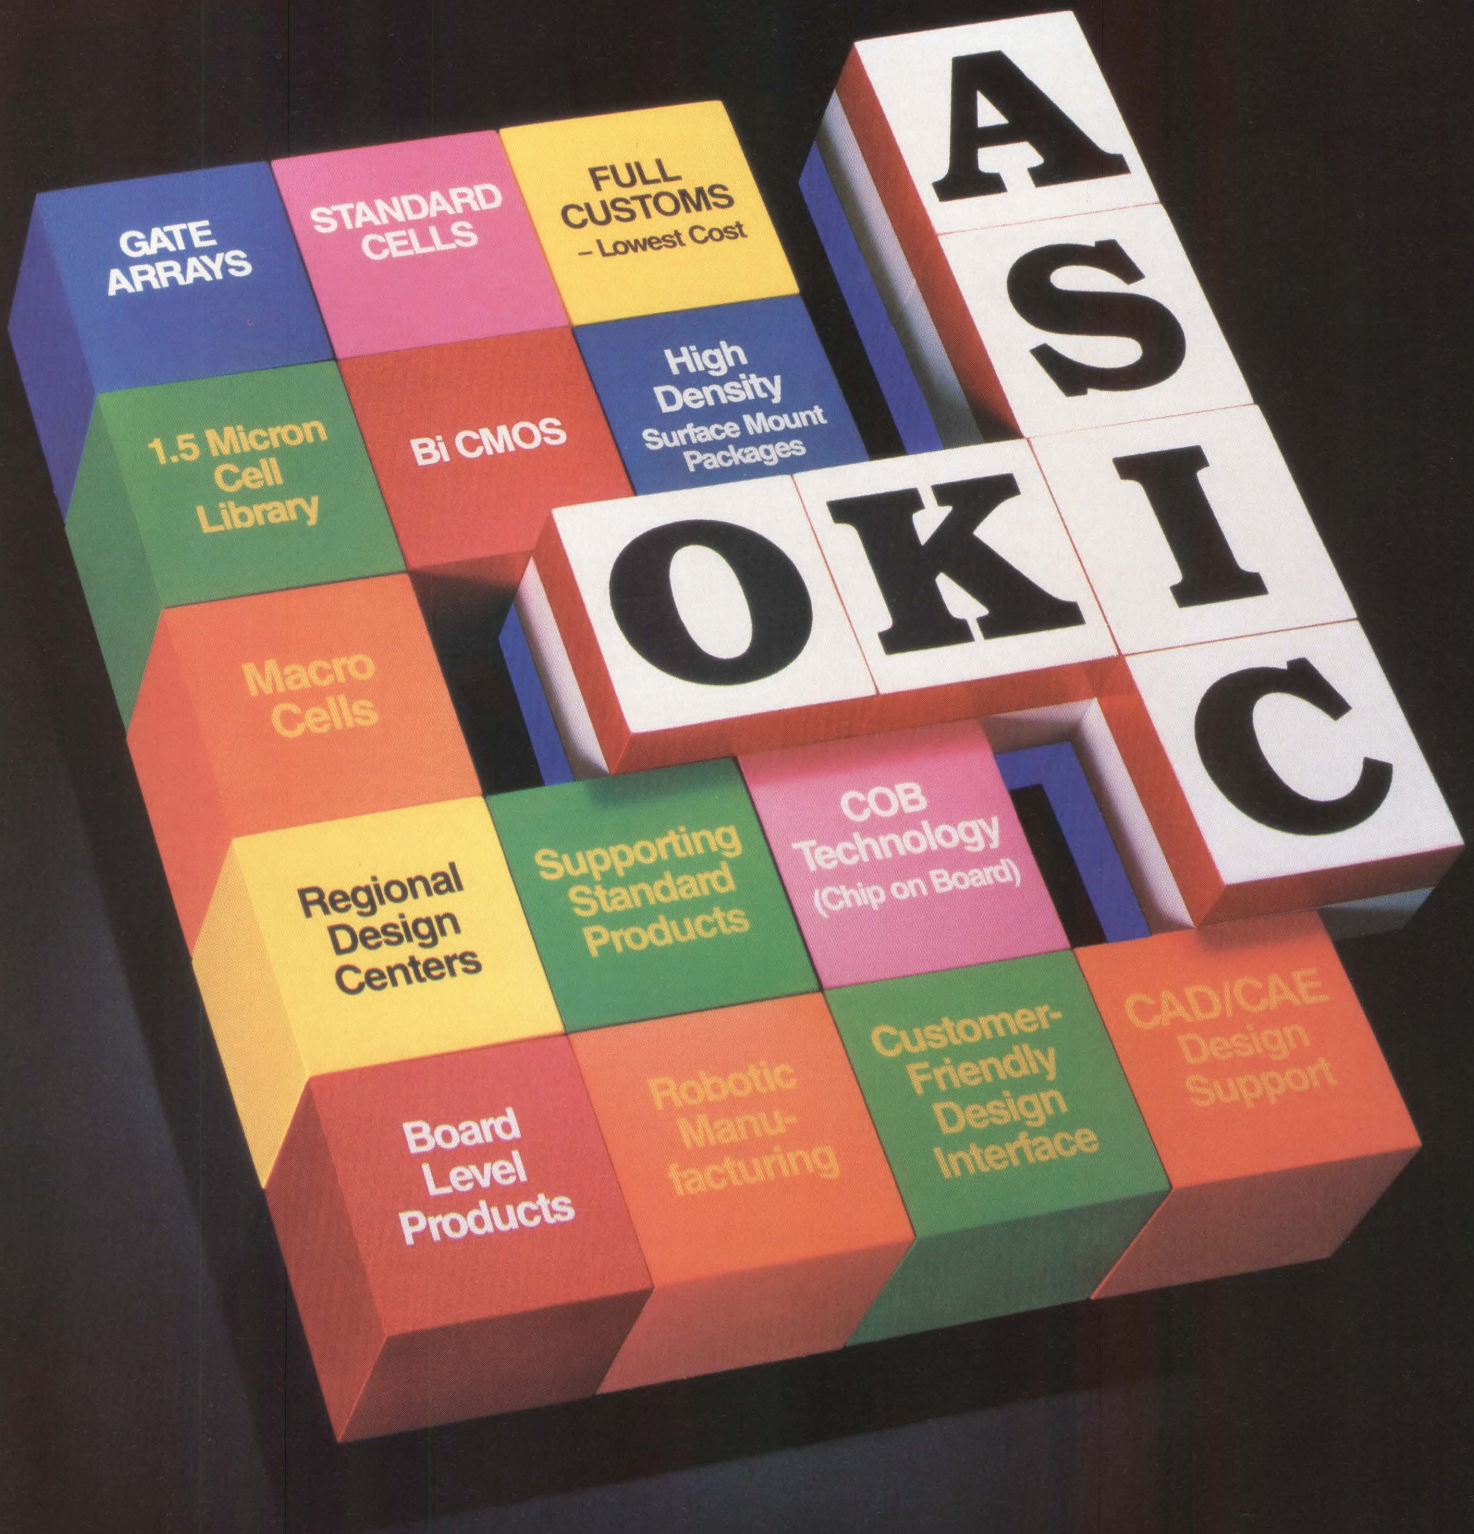
GATE ARRAYS
STANDARD CELLS
FULL CUSTOMS – Lowest Cost
1.5 Micron Cell Library
Bi CMOS
High Density Surface Mount Packages
A
S
O
K
I
C
Macro Cells
Regional Design Centers
Supporting Standard Products
COB Technology (Chip on Board)
Board Level Products
Robotic Manu-facturing
Customer-Friendly Design Interface
CAD/CAE Design Support

Only from OKI:

- Most **complete** ASIC building blocks.
- Most **versatile** design/package options.
- Most **experienced** ASIC technology.

Nobody but nobody puts ASIC technology together like OKI Semiconductor can.

Ease into ASIC with OKI as your close working partner—and you instantly support your VLSI application with the most comprehensive ASIC capabilities on the world market today. Bar none.

From gate array, standard cell and full custom chips to standard components to integration to advanced board level products, OKI alone puts you on the leading edge of ASIC technology and its complete implementation.

OKI: the totally logical choice.

Opt for OKI ASIC, and you open up your options across the board. Only OKI now offers the system designer the unique security and entry ease that only a proven track record in CMOS ASIC problem-solving can provide. This history of performance built up since 1977 has produced the widest range of solid building blocks yet: advanced ASIC products and packaging including surface-mount, backed up with the most flexible cell libraries, CAD/CAE design tools and development aids.

As your working partner, OKI ASIC expertise is available to you at any stage of the development process. We'll help you define system requirements, determine the most cost-effective product solutions and supply complete design software—accessible at your own workstation or through our regional design centers. And then we take it from there: with high volume fabrication, assembly and testing completed in one of the world's most highly robotized manufacturing facilities.

Compare Total ASIC Capabilities	OKI ASIC	Source "B"	Source "C"
Gate Arrays to 10K Gates	●	●	
Standard Cells to 30K Gates	●	●	
Full Customs – Lowest Cost	●		
1.5 Micron Cell Library	●	●	●
Macro Cells	●	●	
Bi CMOS	●		
High Density Surface Mount Packages	●		
Board Level Products	●		
Supporting Standard Products	●		
COB Technology (Chip on Board)	●		
CAD/CAE Design Support	●	●	●
Customer-Friendly Design Interface	●	●	●
Regional Design Centers	●	●	●
Robotic Manufacturing	●		

ASIC Solutions from **OKI:**
You can't beat the logic!

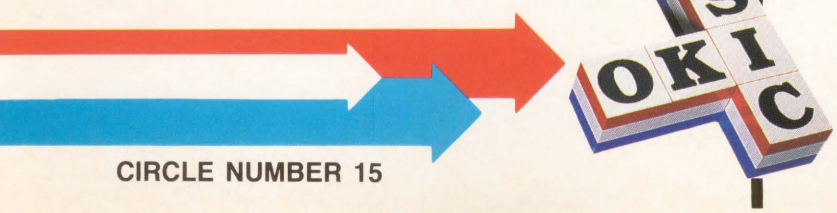

Check out OKI ASIC data:

VLSI 888

Please rush complete technical data/specs on OKI capabilities in:

( ) Gate Arrays
( ) Standard Cells
( ) Full Customs

( ) Please call: we have immediate requirements:

Name/Title _____

Company _____

Attach coupon to business card or letterhead and return to: ASIC Customer Service, **OKI Semiconductor,** 650 North Mary Avenue, Sunnyvale, CA 94086. Tel: (408) 720-1900.

CIRCLE NUMBER 15

Guaranteeing ASIC TESTABILITY

DR. RALPH MARLETT AND STEPHEN R. POLLOCK,
HHB SYSTEMS INC., MAHWAH, NJ

95% fault coverage should be designer's goal (for ASIC'S)

LOGIC DESIGN, layout and fabrication come first when designing with printed circuit boards; then, the test engineers usually take over. With ASICs the situation reverses. The foundry requires both design and test information at the time of design submission. Test considerations have to be moved up further into the design cycle. In this article, we will explore the particular need that ASIC designers must address to achieve high levels of fault coverage, and then present a practical strategy for achieving that level at the lowest possible cost in test generation time and silicon utilization.

The "design, layout, fabrication, and in-circuit-test" strategy for circuit board implementations has worked well for several reasons. First, tests were available for the standard parts loaded onto printed circuit boards. Second, boards could be probed easily to test internal nodes. This is especially important because the complexity of boards often renders functional edge-connector testing impractical.

With ASICs, the same complex logic now resides on a single chip for which the designerhas access only to "edge connectors." The ASIC designer/test engineer is faced with the same kind of test problem, but with far less accessibility. The functional testing problem must be faced because chip complexity will continue to grow.

The rapid-turnaround benefits of ASIC technology are reduced or lost when test engineers write these tests. If the test process is to be relegated to test engineers, and not the ASIC designer, then the test engineers must learn the system's architecture and ASIC's functionality before they can even begin. This can add weeks or even months to a project when it's performed at the end of the

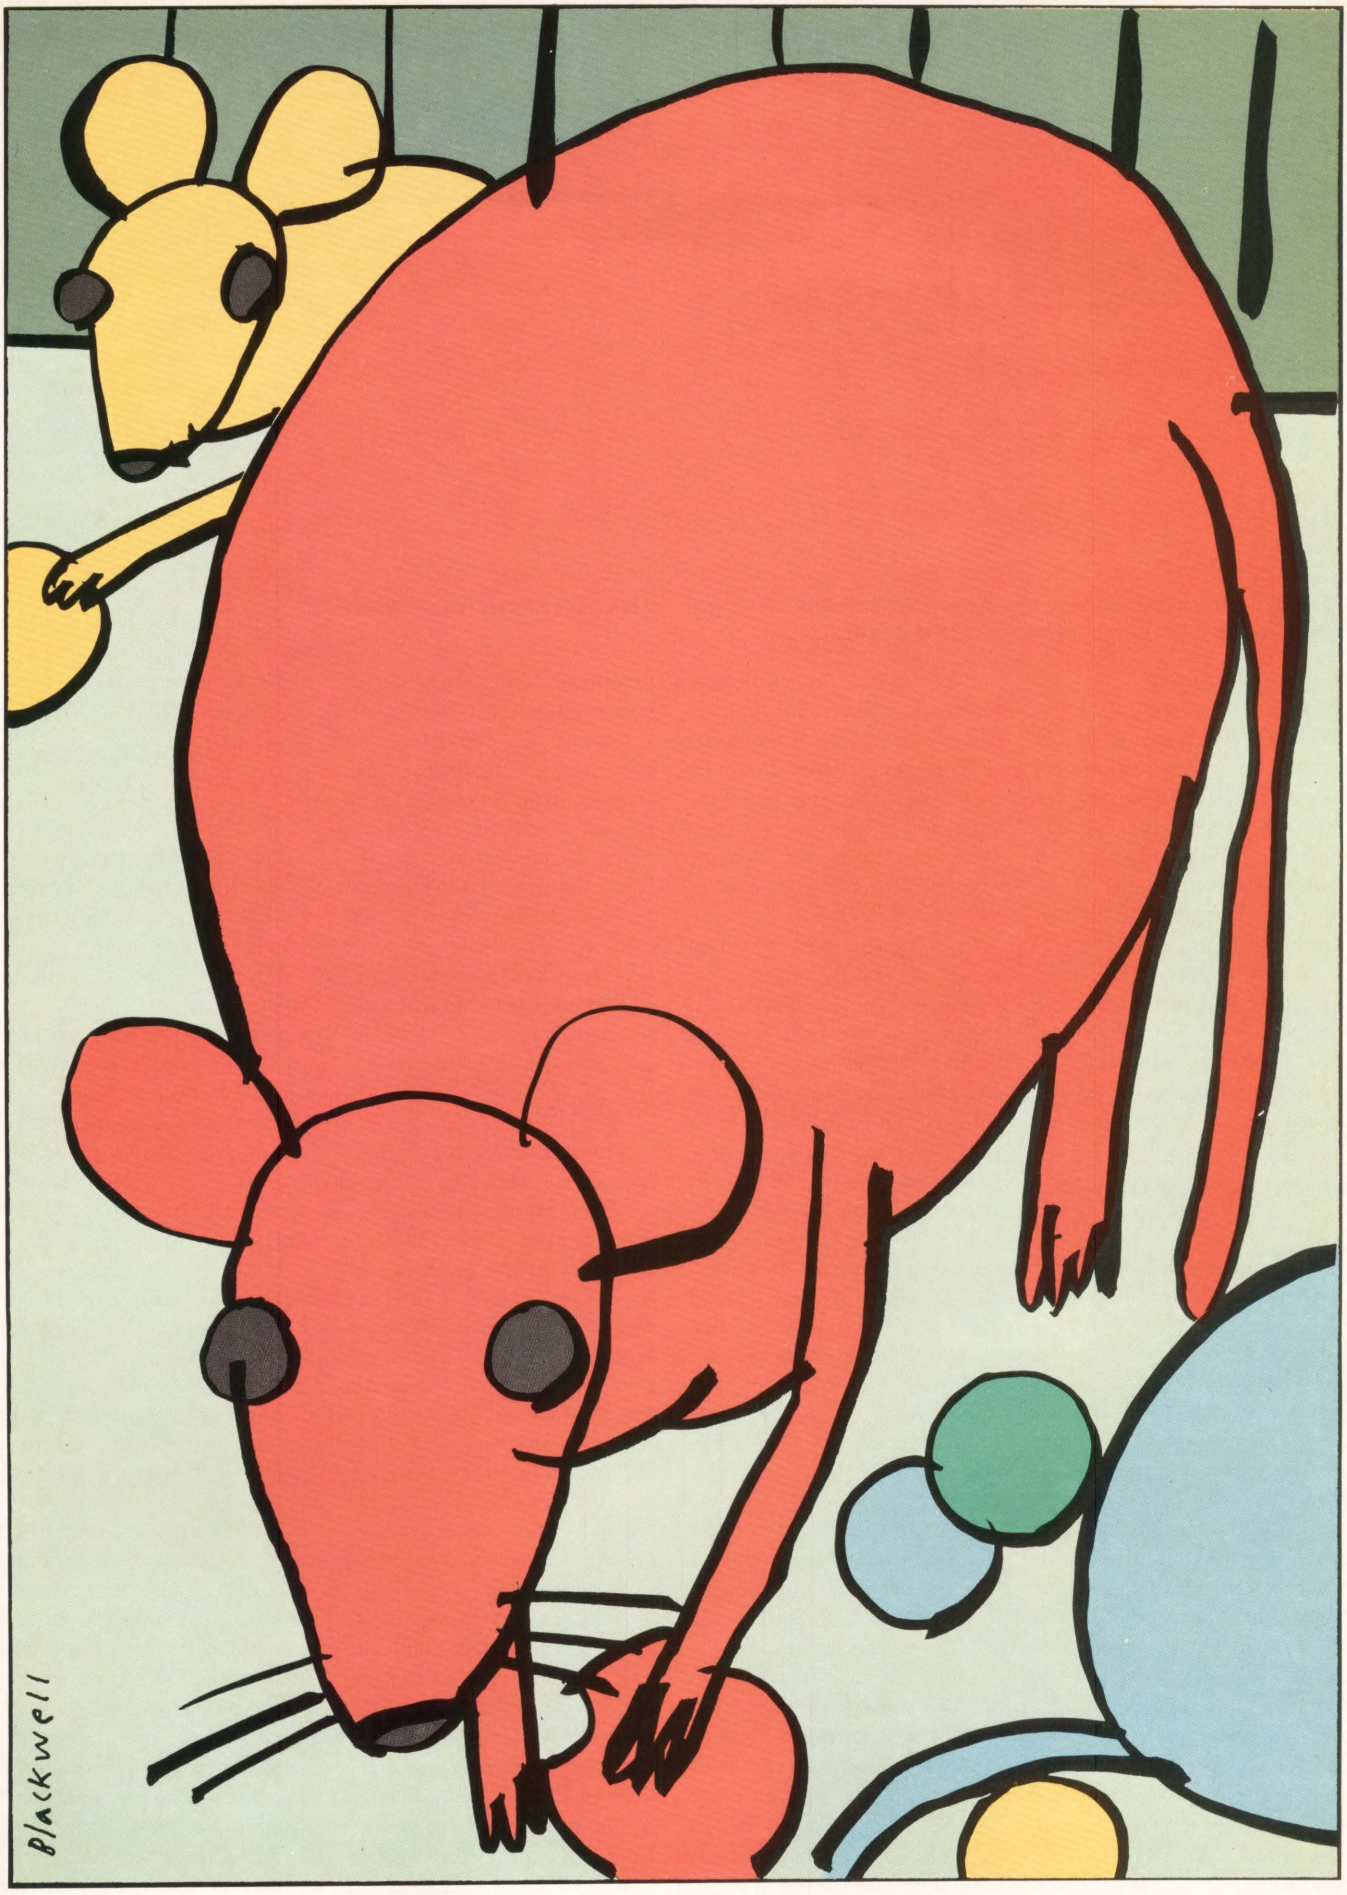

ILLUSTRATION BY PATRICK BLACKWELL

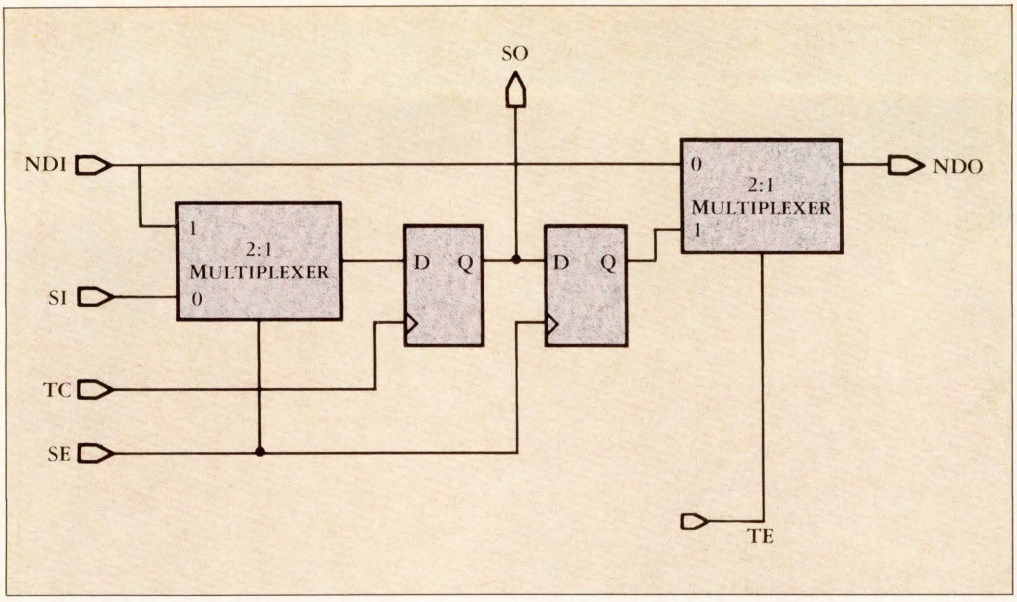

Figure 1. A T-CELL is an improved scan element that, unlike traditional scan elements, can be used on any node (not just storage nodes), and they do not disturb the internal state of the circuit.

of the design cycle.

Designers themselves often skirt the issue by trying to rely on their logic verification vectors, converted in format for ATE. But verification vectors typically provide only 55 to 65 percent fault coverage.

Before designers can be convinced to do what is necessary to improve fault coverage they must first be convinced of the need to achieve at least 95 percent coverage. Bottom-line savings, combined with reasonable test design times must be achieved.

■ Why ASICs Demand Good Test Coverage

Fundamental to quantifying the effect of improving ASIC fault detection on lowering PCB testing costs is the Poisson relationship between yield and average number of defects: $Y = e^{-D}$, where Y is the first pass yield and D is the average number of defects per ASIC.

This Poisson distribution process can be thought of as an ASIC wafer with N die sites and a box containing N defects, so there is an average of one defect per ASIC die. The defects are then randomly "sprinkled" over the wafer.

The result, according to Poisson's law, is that 36 dies will have no faults, 38 will have one fault, 18 will have two faults, six will have three faults and two dies will have four faults. The yield (the percentage of dies with no faults) is 36 percent. This is in agreement with the equation $Y = e^{-D}$, where Y is the first pass yield and D is the average number of defects per ASIC ($Y = .36 = e^{-1}$).

ASIC yields span a broad range depending on process parameters and techniques. If the yield for every die site on a wafer is mapped across the wafer, then it becomes clear that the yield falls off drastically at the wafer's periphery. ASIC parts typically yield 40 percent at best; since only a few thousand parts are typically built the fab line must be shut down and restarted for each run of ASICs. It's too expensive to fine tune the ASIC line to get highly optimized yields.

For example, using Poisson's formula, a 40 percent yield requires that the process be controlled to .91 average defects per die. But, doubling the yield to 80 percent requires that the defect rate be reduced by more than four times to .22 average defect per die.

Once the foundry defects per die site (or in effect, per chip) are determined, multiplying this result by (1 - fault coverage) results in the average number of undetected defects per ASIC at that level of fault coverage. At the typical 40 percent yield for ASICs, for instance, 70 percent fault coverage results in .273 defects per ASIC going undetected. Increasing fault coverage to 98 percent results in .018 defects going undetected for a net defect reduction of .255 defects per ASIC. Parts with an 80 percent yield, however, only benefit by a defect reduction of .062 defects per chip using the same analysis.

■ Actual Dollar Savings

This reduction in escaped defects is especially valuable for dollar savings. Industry estimates range from $100 to $500 as the cost to locate, remove, and replace a defective ASIC on a printed circuit board. If the problem escapes detection before shipment and ends up as a field failure, then the cost goes up by a factor of 10. But more important, if the defect escapes detection at the factory, the manufacturer's quality image and credibility can be damaged.

If we assume an average of $200 per defective ASIC and multiply the average defect reduction per ASIC by this $200 figure, then we can determine the improvement resulting from increasing the test coverage of the 40 percent process yield ASICs from 70 to 98 percent: an average saving of $51 per ASIC. For a system with five ASICs, the saving is $255. The savings per ASIC are shown for other wafer probe yields in the table.

Another way of looking at cost savings is to consider them on an annual basis. Consider a PCB with five ASICs. We can compute annual savings by multiplying the defects reduction per board by the average cost per repair ($200) by the average number of ASICs per board (five) by the number of boards produced per year. For the .255 defect reduction per ASIC already considered, a manufacturer producing 1,000 boards per year can save $233,000 annually.

With potential savings like these, what options are available for improving test coverage?

■ Options for Improving Test Coverage

One obvious alternative is to manually generate more test vectors. There are a number of drawbacks, however.

First, how do we want to expend the designer's resources? Do we want to take away from creative design time? Or do we hire a test engineer who must spend time learning the system and ASIC functionality? The designer would not have to lose that amount of time.

Still, the test engineer is likely to be more experienced writing tests. An engineer who understands the system can probably improve fault coverage better and faster than the designer could. As complexity increases, the problem gets worse.

Because it is so difficult to generate adequate test vectors after design completion, more people in the industry are recognizing that testability is a design issue. We propose two basic strategies:

1) Modify the design for better testability based upon a testability analysis. Use scan discipline throughout, or use a Reduced Intrusion Scan Path.

2) Automatically generate

test vectors for the desired coverage.

Testability can be improved by measuring it using testability analyzers, such as the SCOAP program from Sandia Laboratories, and then modifying the design based on the SCOAP analysis. SCOAP gives a testability figure of merit and lists nodes that are untestable because they cannot be observed or controlled.

A more forceful approach is to use one of the scan disciplines (such as Scan Set, Scan Path or LSSD), which require strict adherence to a set of design rules and the replacement of all registers with scannable latches that are then connected into a scan chain. A third approach is to use built-in self-test techniques that can be used in certain rather qualified situations.

First, let's consider testability analyzers. Testability analyzers frequently produce a figure of merit for the testability of a circuit, but they don't tell you how to write the tests. Virtually every circuit is testable, but some are more difficult to test than others. You may not know how to write the test, but they are testable.

On the positive side, the SCOAP and similar testability analyzers give very rapid results. SCOAP, however, lacks sufficiency. It represents a minimum condition that indicates something is wrong with the circuit in terms of testability, but once it says the circuit is inherently testable, you really don't know if you can write the test or not.

There are legitimate instances in which a designer does not require that a node be observable or controllable. Perhaps the Q-bar of a flip-flop wasn't used. This would appear, correctly, on the SCOAP report as unobservable, and the designer would simply accept that instance as an intended result.

SCOAP can, however, produce misleading results. For example, SCOAP may declare all nodes of a parity detector controllable. But if one tries to test a parity circuit embedded in a larger circuit, one may discover that the only way to test parity detection is to inject that parity data. This may not be possible if the parity circuit comes alive only when a parity error occurs. The parity checker, then, represents a circuit that can't be tested, yet the SCOAP testability analysis says it's testable.

Consider an AND gate with inputs A and A-bar. Under these conditions the AND gate's output can never go statically high. Yet testability analysis indicates that the output node can be controlled both high and low. These are inherent defects in SCOAP-like analyzers. These analyzers may not be sufficient as stand-alone testability tools, but they often are important components of some of the most satisfactory test strategies.

Second, consider the improvement of testability through the use of a scan discipline. Scan disciplines guarantee testability, but at a major cost in silicon area. Twenty percent or more typically is cited as the silicon penalty.

In addition, some designers are reluctant to adopt the constraints of the design rules. These constraints, such as the requirement that the design be fully synchronous with no feedback loops or testability elements in clock lines, are serious limitations. In some cases circuit performance is impaired. The situation is improving, though, with software that automatically synthesizes the scan chains. Still, the designer is hampered by scan restrictions.

New testability techniques, such as those used by HHB Systems' INTELLIGEN, eliminate the need to place scan elements everywhere. Instead, just those nodes required to improve automatic test pattern generation (ATPG) fault coverage are made scannable. This drastically lowers the silicon real-estate penalty.

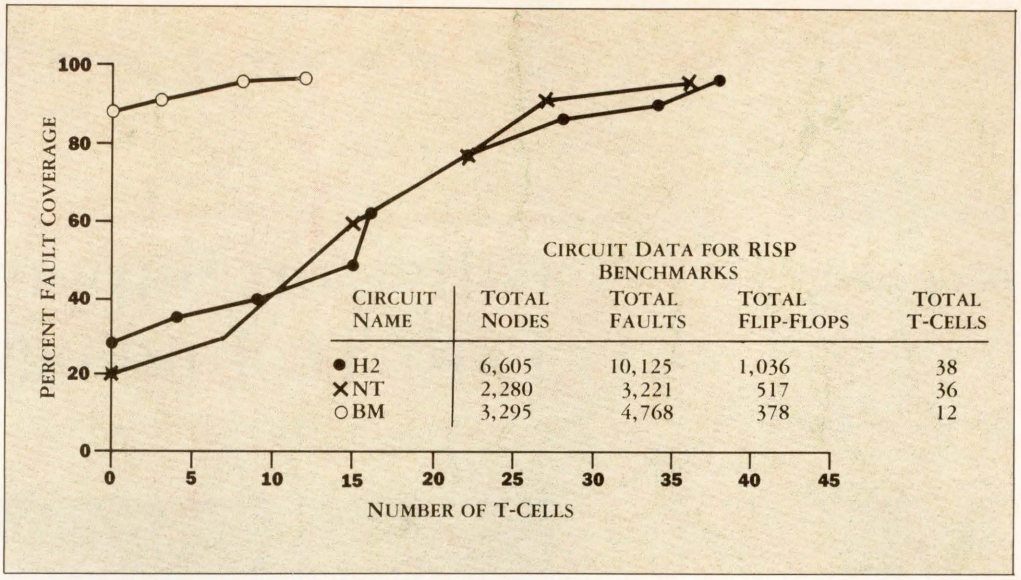

Figure 2. The first-pass results of running INTELLIGEN on the H2, NT, andf BM benchmark circuits yielded fault coverage of 20 to 88 percent with automatically generated test vectors. Adding a relatively small number of reduced intrusion T-CELLS resulted in fault coverages of 98 percent or better.

SAVINGS FROM DEFECT REDUCTION

WAFER PROBE YIELD (PERCENT)	FOUNDRY DEFECTS/ASIC	DEFECTS/ASIC (70% FAULT COVERAGE)	DEFECTS/ASIC (98% FAULT COVERAGE)	DEFECT REDUCTION/ASIC	COST SAVINGS PER ASIC (DOLLARS)
40	.91	.273	.018	.255	51.00
50	.69	.207	.014	.193	38.60
60	.51	.153	.010	.143	28.60
70	.36	.108	.007	.101	20.20
80	.22	.066	.004	.062	12.40

■ EFFECTIVE ATPG FOR SEQUENTIAL CIRCUITS

Most ATPG software requires use of a scan discipline with its comprehensive scan chain, and design and silicon utilization overheads. SystemsINTELLIGENpermit automatic generation of vectors in any design, with or without scan chains, using a backward searching algorithm.

Backward searching algorithms can generate all test

Continued on page 76

Tester Targets Deep Patterns, Fewer Pins, and Lower Test Rates

Gillytron Supports Scan-Path Technology

Many designers resist using design-for-testability (DFT) techniques, such as scan-path design, because the techniques can constrain design practices, increase circuit complexity, and potentially reduce performance. An equally important concern, however, has been the lack of verfication and test equipment that can take full advantage of the benefits of scan-path design. Traditional ATE is focused on providing high test rates at all device pins. Many of today's leading ATE systems provide full tester resources (both high-speed stimuli and high-speed data acquisition) at each device pin. On the other hand, scan-path techniques usually require the driving and monitoring of only a few pins, but with very deep pattern memory and at relatively low data rates.

A new design verification system from Gillytron, Inc., promises to release the full potential of DFT techniques with: an architecture optimized for scan technologies; a novel (and inexpensive) approach to ac testing; and software to help create the tests and analyze the results.

The ScanMaster DV-6005 Design Verifier uses a modular design built around a 68020-based scan module, a 68020-based control module, and one or more functional modules. Each scan module has 23 lines for clocking and control signals, a scan-in data line and a scan-out data line. A "scan generator" feeds scan patterns into the scan path through the scan-in data line and retrieves the results from the scan-out data line. The clocking and control signals drive clocks and control signals at the chip input pins. For a chip designed with a comprehensive scan-path testability approach, these lines are sufficient for a complete test of the device including circuits and nodes that would normally be hidden or difficult to control and observe.

The scan module also contains a frequency counter that is used for the ac testing, programmable power supplies and the parametric testing circuitry. As many as four scan generators can be used in one test system. In addition, the scan memory can be increased up to a maximum of 256M bytes.

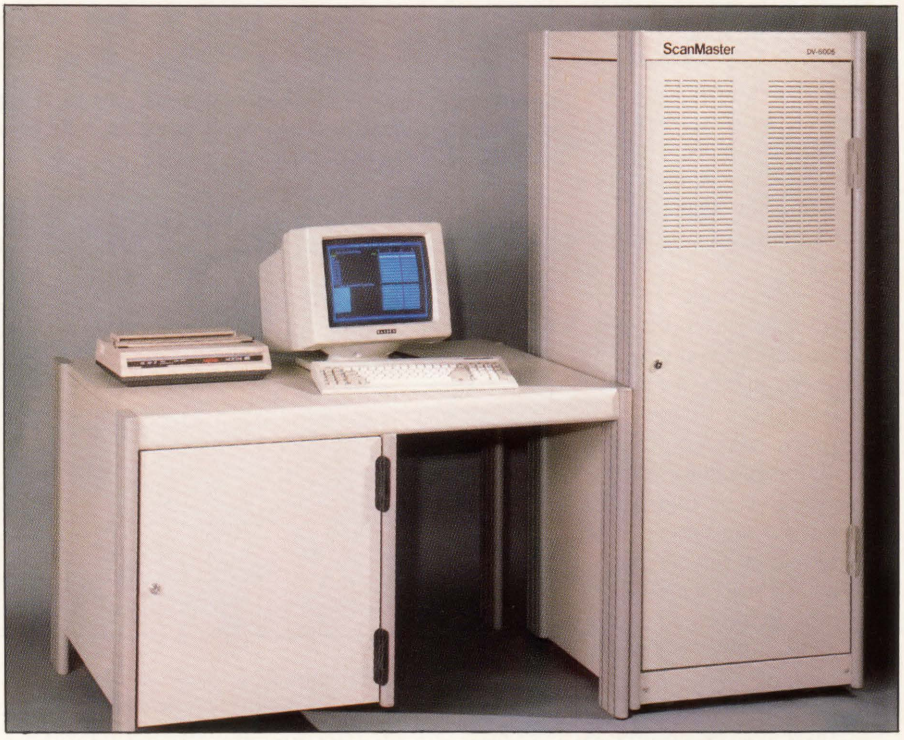

Scan generators in the DV-6005 ScanMaster directly drive scan-path testing in complex ICs and boards.

The functional modules are used during the scan testing to set the chip input values and then sample the chip outputs which are influenced by the scan test vectors. The modules are similar to the pin electronics packages that are used in the more traditional ATE systems in that they contain large numbers of pins (128 per module) and have the ability to apply and capture signals at all pins. The Gillytron functional modules can operate at pattern data rates as high as 5 MHz; they contain a minimum of 16k bits of pattern memory; and they can perform parametric tests of each pin as well.

During a typical test sequence, the scan module serially loads the scan pattern into the device to be tested through the scan-in data line. The functional modules set the desired test conditions at the device inputs, and then the system strobes the device's clock inputs to propagate the effects of the scan pattern through the device via its internal scan-path. Output data from the device's pins is then captured and compared with expected results in the functional modules. The device's states, which are stored in the scan-path are clocked out by the scan module through the scan-out data line for complete analysis. At the same time the resultant data was clocked out, the next scan pattern would be loaded into the chip through the scan-in data line.

According to company claims, the use of scan-path techniques and the DV-6005 tester can result in over 99 percent fault

coverage of most complex prototype ICs. Circuit observability and controllability through the scan-path helps designers find the exact location of any single stuck-at fault condition. In Gillytron's eyes, the extra overhead of 10 to 15 percent of the circuitry devoted to test, is outweighed by the high fault coverage, the availability of scan-based automatic test pattern generation (ATPG) and the reduced cost of the less complex test equipment.

■ SCAN-PATH AC TESTING

The DV-6005 test system uses a fundamentally new approach to ac testing which is called scan-string testing. Using this method, the tester creates a closed loop between the scan-in and scan-out pins which causes the scan path to oscillate. The frequency at which the scan-path ring oscillates depends on the number of the circuits in the scan-path as well as their average propagation delay. The DV-6005 can measure oscillation frequencies up to 100 MHz. As the number of serially-connected scan elements on a chip is increased, it becomes possible, using this oscillation technique, to measure the performance of chips with gates that can operate well in excess of 100 MHz.

Scan-ring ac testing is considered more accurate than critical path testing because the scan-ring testing verifies the performance of a large number of circuits across the chip, rather than one specific signal path. In addition, the chip-wide ac testing path can easily provide an accurate measure of the chip's overall performance when operated over the normal temperature range for the device.

The DV-6005's control module connects to an IBM PC/AT which runs the test development software. Menu-driven programs help the user prepare test programs and generate test patterns. The test program uses a C-like language, proprietary processor commands, and user-defined test macros. Designers can use any ASCII test editor to edit the test programs, which are then compiled with the DV-6005's C compiler on the PC/AT and downloaded to the control module.

Gillytron is working with Gateway Design Automation Corp. (Westford, Mass.) and Aida Corp. (Santa Clara, Calif.) to develop interfaces between their ATPG tools and Gillytron's software. Once in place, Gillytron's software will accept and automatically process test patterns to fit into the DV-6005 architecture. This capability is also being pursued with Daisy Systems Corp. (Mountain View, Calif.).

An "immediate mode" program allows the designer to edit and rerun patterns that are already formatted on the tester. A pattern editor gives the user access to the test in binary format, so he doesn't have to go back to the ATPG program to alter his test patterns. It also facilitates debugging on the tester.

For board testing, which usually have a much larger number of control signals, users can add another scan clock to the scan modules, bringing the total number of clock/control lines to 53 per module. In addition, for production testing, a built-in signature analyzer assists in go/no-go determination of the device or board.

To develop the tester, Gillytron received extensive input from Integrated CMOS Systems (Santa Clara, CA), a company that builds scan-based gate arrays with tens of thousands of gates. ICS uses the tester for wafer, device and board testing.

A basic DV-6005 contains a PC/AT user interface, test software, a control module and a scan module with an interface for 25 test probes. It costs $125,000 and comes with a six-month warranty. Additional pins in function modules average about $700/pin.

Gillytron Inc
San Jose, Calif.
(408) 435-3043

Continued from page 73

patterns to achieve a high fault coverage. Or they can take an existing set of test vectors and improve test coverage by automatically adding test vectors of their own.

INTELLIGEN goes one step further. Its Dynamic Testability Analyzer uses the results (or lack of results) in generating the test to produce a list of the nodes where testability improvement should be applied. This analysis is not SCOAP-like, since the the suggestions for hardware modifications result from trying to generate test patterns, not from a static testability analysis based on circuit structure.

The lists of specific nodes where testability improvement should be implemented makes it unnecessary to broadside the testability issue by placing scan registers throughout the design. Only three to eight percent, rather than 100 percent, of all registers need be scannable.

Test pattern generation can be approached as an iterative process. The designer enters the process either "cold" (that is, with no test patterns) or with a set of his own patterns. The ATPG program suggests nodes requiring testability improvement. The designer implements some or all of them, and repeats automatic test generation. At this point, a new list is generated, and the process is repeated until the desired fault coverage is achieved.

A partial scan chain can be used to improve testability of the nodes identified by the Dynamic Testability Analyzer. We call this new methodology Reduced Intrusion Scan Path (RISP), which generates greater than 98 percent fault coverage on sequential designs, yet removes all the design restrictions of full scan.

RISP uses an improved scan element called a T-CELL (Figure 1). T-CELLS, unlike traditional scan elements, can be used on any node, not just storage nodes, and they do not disturb the internal state of the circuit. T-CELLS introduce only one multiplexer delay. Even without the insertion of T-CELLS, INTELLIGEN automatically generates greater than 95 percent fault coverge for about 70 percent of sequential designs. The addition of T-CELLS handles the remaining 30 percent.

RISP closes in on the desired fault coverage. In addition, about 70 percent of all sequential designs do not require insertion of test elements such as T-CELLS.

The results shown in Figure 2 are for three different circuit designs that did not achieve high fault coverage at first. They illustrate clearly the efficiency of adding T-CELLS. For example, adding just 38 T-CELLS increased fault coverage for the NT circuit to greater than 98 percent from the 20 percent achieved through automatic test generation.

In conclusion, this methodology can automatically generate high fault coverage tests for approximately 70 percent of the designs without any design intervention. For the remaining 30 percent, the system recommends exactly where testability needs improvement to enable automatic generation of high fault coverage tests. ■

ABOUT THE AUTHORS

Dr. Ralph Marlett, *director of ATG products at HHB Systems, joined HHB Systems in 1985. Previously, at Westinghouse, he developed ATG techniques and the BIT architecture for the company's VHSIC chip set. He holds a BSEE and MSEE from Carnegie Institute of Technology, and a PhD in Electrical Engineering from the University of Illinois.*

Stephen R. *Pollock, became director of marketing of HHB Systems' IC Products Division in July 1988 with the acquisition of SIMUCAD by HHB. Prior to joining SIMUCAD in 1987, he managed the Systems and Libraries Marketing Groups at Silicon Compiler Systems. He holds a BSEE from Drexel Institute of Technology.*

MARKET PLACE

FOR IMMEDIATE SALE

VLSI CAD equipment in like-new condition.

Calma: 4280GDSII system with Ethernet, Fast Mask Engine, 300 MB Disk, GSSII/32 workstation (2 ea), 2 MB memory expansion, VRSTC 8242C Color Plotter, programming terminal. Extensive software also available.

Mentor: IDEA Station MC-3.0/DN550, 86 MB Winchester/45 Tape, Imagen Printer Set 8/300, Ethernet Gateway, IDEA Station DM-F-2.0/DN320, 70 MB Winchester, extensive software.

Robert Pozzo
Polo Computer
(617) 999-1436

CONSULTING SERVICES

INCREASE YOUR PROFITS

TAI staff has extensive experience and successful track records in promoting and selling products and firms to the military, automobile, nuclear and industrial markets.

If you are in business or if you are only in the "idea stage", we can help! We will:

- Develop your Business Plan
- Structure your Quality Assurance Program to meet MIL-I-45208A and MIL-Q-9858A
- Develop your Brochure
- Assist in locating and securing financing
- Promote your Products Nationally and Internationally
- Assist in Developing and Implementing Sound Business/Management Practices Throughout your firm

Send information about your type of business and the type of help you need, with a self address envelope, to:

VP, Marketing Dept
Tech for Adva
P.O. Box 171253
Kansas City, KS 66117

CAREER OPPORTUNITIES

DESIGN THE FUTURE.

Lockheed Missiles & Space Company in Sunnyvale, CA has immediate openings for qualified, experienced professionals to join our team designing and developing the most innovative, advanced technologies of the 20th century.

VLSI DESIGNER

If you have a BSEE or BSCS and three or more years' experience in computer aided VLSI circuit design, we need you to develop custom VLSI processors on Mentor Graphics. We have the most updated equipment available and use the latest design tools.

SCIENTIFIC PROGRAMMERS

You will perform design, coding, testing, documentation and maintenance of software and procedures for support of Computer Aided Manufacturing applications, including distributed and computer numerical control.

Six or more years' computer programming and analysis experience and three or more years of FORTRAN experience are required. In addition, you should have extensive knowledge of numerical control software design, development and maintenance. CAD/CAM software experience is highly desirable.

At Lockheed, we offer you challenging projects, excellent compensation and a long list of company benefits, all surrounded by some of the most beautiful landscapes in the country. To make it all yours, send your resume to U. Allen, Lockheed Missiles & Space Company, Professional Staffing, Dept. 576JEUA, P.O. Box 3504, Sunnyvale, CA 94088-3504. We are proud to be an equal opportunity, affirmative action employer. U.S. citizenship is required.

Lockheed
Missiles & Space Company
Innovation
Giving shape to imagination.

World-Class Commitment

Mentor Graphics Corporation is committed to delivering the best design capture, synthesis analysis and simulation tools for all aspects of electronic-based product design. We intend to redefine what constitutes a world-class design environment, covering digital, analog and system design.

We invite engineering professionals like yourself to join us in Beaverton, Oregon, and help set the pace.

You'll help us deliver a system that enhances not only our own applications, but integrates with other tools. A tool set that will save design engineers around the world a lot of frustration from design to test.

We're committed to delivering the best.

We firmly believe our vision of next-generation tools is the answer to a lot of design engineers' dreams. That's why we're committing our resources, our company and our people to this challenge. If you're equally committed, and have a specific view you'd like to share, we invite you to explore these managerial, senior-level and staff opportunities to join us:

Product Development

Engineering Managers

Technical Leaders

Software Development Engineers

- System Simulation
- Logic/Fault Simulation
- Circuit Simulation
- Parallel Logic Simulation
- DFT/BIT/ATG
- Logic Synthesis
- Modeling
- Interactive Graphical User Interface
- User Interface Management Systems
- Design Data Management Systems
- Graphical Design Capture
- System and Architecture Synthesis

Everybody needs a goal.

We'll give you a world-class goal and the resources to meet it. Most importantly, we'll provide a forum for your ideas, a place where your technical opinion is respected.

We offer relocation to the heart of Oregon's "Silicon Forest," which offers unmatched recreational opportunities, affordable housing and a high-quality lifestyle.

Find out more; send your resume to Professional Staffing, Mentor Graphics Corporation, Dept. 2000, 8500 S.W. Creekside Place, Beaverton, Oregon 97005-7191.

We are proud to be an equal opportunity employer committed to affirmative action. We encourage qualified female, handicapped, minority and/or veteran candidates to apply.

Mentor Graphics®

Your ideas. Our experience.

The Intel Influence

Leading the Microcomputer Revolution.

We started out with a goal and a commitment: To lead the way in microcomputer technology. We achieved it by never swerving, never settling for second-rate or second-place. At Intel, our philosophy became second-to-none. We achieved it by giving our professionals the most sophisticated tools, the most creative environment and the allegiance of a highly progressive management. Join the team defining the future computing environment for the office and lab in either Santa Clara, CA or Phoenix, AZ.

ASIC DESIGN CENTER MANAGER

Santa Clara

We are currently seeking a senior ASIC Design Engineer with 8+ years experience supervising at least 2 engineers as well as having designed at least one large (greater than 10,000 gates) standard cell. Experience with basic simulators (Tegas, Hilo, Lasar) in a mainframe CAE workstation platform (Daisy, Mentor, SDA), and working knowledge of digital IC test methodology and tester formats is required. Literacy in computer languages such as C, FORTRAN, PASCAL or Mainsail; experience with place and route; and in-depth knowledge of digital ICs is highly desired.

CAD ENGINEERS

Numerical Analysis CAD Engineers

Requires a strong background in numerical solution of partial equations, semiconductor device simulation and programming.

CAD Application Engineers

Requires 3+ years with VLSI logic design experience and CAD applications experience or interest on CAD. Automatic physical design desired.

Project Managers

Requires at least 3 years management of software or hardware projects. Must have working knowledge of CAD for IC physical design and either direct experience with management of projects utilizing IC CAD.

CAD Graphics Software Engineer

Requires 3 years experience developing graphics systems software. Knowledge of X windows highly desirable. Must have experience designing and implementing general purpose graphics systems.

Graphics Specialist

Responsibilities include developing, testing and supporting design verification procedures and methodologies. One year experience using design verification tools such as DRC, CVS and plotting required. High level language required such as C or Mainsail.

Modeling Engineers

Requires 5+ years in CAE industry. Must have experience in component model development, preferably full functional at behavioral level. Knowledge of hardware description language and familiarity with board/system simulation desired.

Senior Systems Manager

Experience with VAX computers and UNIX is required. NFS is desirable.

SOFTWARE ENGINEERS

Senior Software Engineer

Requires 4-6 years experience in system software design, test and reduction with specific expertise in compiler design, code generation and code optimization. Must have thorough knowledge of C, UNIX, FORTRAN.

Software Engineers
Device Drivers

Requires UNIX device driver and network protocol experience. Experience networking different system architectures is helpful.

System Software Architect

Requires multiprocessor system design, knowledge of UNIX multiprocessor implementation process scheduling, virtual memory, cache coherency, languages, etc.

HARDWARE SYSTEMS DESIGN ENGINEERS

You will use your hardware design experience with PCs and add in boards to develop high performance 80386 system products. These designs will become the design model for high end computer systems. Candidates will be responsible to define the optimal VLSI solution for a PC/AT system.

ALL POSITIONS REQUIRE A BS/MS IN EE OR CS.

We welcome your interest in Intel...please send your resume including location choice to:
Intel Staffing, Dept. M386, P.O. Box 58119,
Santa Clara, CA 95052-8119 or call
Becky at 408-765-1546.

Equal Opportunity Employer M/F/H

intel®

ADVERTISERS' INDEX

VLSI Systems Design

PUBLISHER
NORM ROSEN

NEW YORK
WALTER L. OLSON
District Mgr. (516) 562-5711
KENNETH D. BEACH
District Mgr. (516) 562-5711
DAVID JANOFF
District Manager (516) 562-5846

NEW ENGLAND (617) 244-5333
KAREN TALLARIDA
District Manager

DALLAS (214) 661-5673
CLAIRE FLORA
Senior Accounts Manager

CHICAGO (312) 565-2700
MARY SHUTACK
District Manager

SAN JOSE (408) 252-6191
KATHY MICKELSON
Western Regional Sales Manager
JULIE TAFEL
District Manager
TONY UPHOFF
District Manager
MICHAEL TORCELLINI
District Manager
JONNO WELLS
District Manager

PORTLAND (503) 636-7694
RICHARD CARLISLE
District Manager

LOS ANGELES (213) 473-9641
TODD BRIA
District Manager

NEWPORT BEACH (714) 851-2022
KATHY MICKELSON
Western Regional Sales Mgr (408) 252-6191

INTERNATIONAL ADVERTISING
DANIEL H. LEEDS, Vice President
(516) 562-5000
STEVE DRACE, International Sales Manager
FAX 852-5-861-0668
PacificMedia, Inc.
Fax: Tokyo 252-2780 Telephone: 256-8456
Telex: J25472
Young Media, Inc., Seoul, Korea
Telephone: 756-4819 Fax: (02) 757-5789
Mr. Owen Wang
Ace Marketing, Inc.
Telephone: 751-3636 Telex: 78514142

CLASSIFIED ADVERTISING
NEW YORK (516) 562-5000
SANDRA HERMAN/National Manager
MICHAEL ZERNER, East Coast District Mgr.
NEW ENGLAND N.Y. N. Central District Mgr.
KATHY HEALY, Southeast, South Central Space Sales
SONJA WONG, Advertising Coordinator
FELICIA SABATINI, Graphics
SAN JOSE (408) 252-6191
LAURIE MUNCE, West Coast District Mgr.

DIRECT MARKETING SERVICES
JOANNA BRANDI, Manager
PETER CANDITO, Sales Manager
ARLENE BROWN, List Sales Representative
(516) 562-5000

MARKETING SERVICES
DORIANNE WALTHER-LEE/Director

CUSTOMER SERVICE/DISPLAY ADVERTISING
LYNDIANE HERVEY, Group Manager
LINDA USLANER, Ad Coordinator
(516) 562-5000

PROMOTION
GENEVIEVE HIGGINS

CIRCULATION
GRACE MONAHAN, Director of Operations
NANCY MERRITT, Circulation Director
ROBERT ANDERSEN, Circulation Manager
SUBSCRIPTION SERVICES
(516) 562-5882

RESEARCH
DANIEL R. CAMPBELL, Electronics Group Manager

REPRINTS
BESSIE RHODEN
(516) 562-5000

CMP ELECTRONICS GROUP
KENNETH D. CRON, Vice President/Group Publisher
Electronic Buyers' News
Electronic Engineering Times
VLSI Systems Design

CMP Publications, Inc.
MICHAEL S. LEEDS, President
PEARL TURNER, Vice President/Treasurer
DANIEL H. LEEDS, Vice President

LILO J. LEEDS, GERARD G. LEEDS
Co-Chairpersons of the Board

NANCY GRABOW
Assistant to the President

Most simulation systems tell you as much about your ASIC design as this tells you about the Statue of Liberty.

Get the full picture of your ASIC design's performance. Get the IKOS Simulation System. IKOS gives you more simulation information in less time than any other system. Simulate 15,000 gates with 34,000 vectors in 34 seconds. Or 8,500 gates with 5,100 vectors in 4 seconds. Complete with detailed timing information!

IKOS is a comprehensive ASIC simulation system incorporating a combination of sophisticated software and hardware components that run on the Sun, Apollo, or IBM PC/AT platforms. It provides you with a streamlined, high-speed approach to stimulus entry, logic and fault simulation, analysis, and ASIC library support. And, it integrates easily into your existing design environment.

IKOS gives you a tremendous performance advantage.

You enter stimulus the way you visualize it—by drawing waveforms. Powerful signal modeling features allow you to easily and realistically simulate asynchronous signals and interfaces.

With IKOS' real time logic analyzer you can run simulations interactively, pinpointing signals and conditions of interest. Gone are the tedious hours spent pouring over long output listings.

There's no limit to the number of test cases you can generate and analyze with IKOS. You'll be able to simulate seconds, *even minutes* of real time system operation. And, gain the confidence that your ASIC design works as you intended it to.

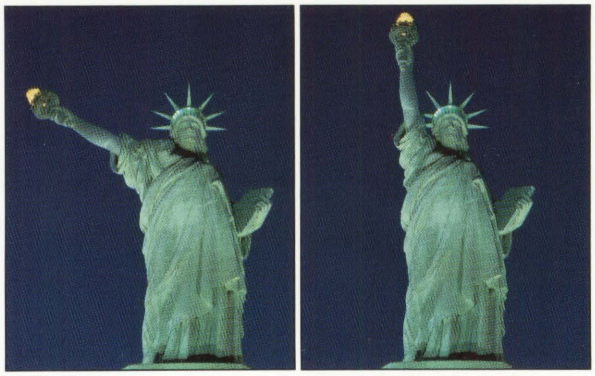

Without the IKOS Simulation System, most ASIC design mistakes show up after production.

IKOS comes complete with all the simulation tools you need—at a low price. Use IKOS to get your ASIC design right *before* committing to manufacturing. It's almost as easy as snapping a picture.

Give us a call.

IKOS

Now available on Sun!

408 • 245 • 1900

145 N. Wolfe Road · Sunnyvale, CA 94086

Sun is a trademark of Sun Microsystems. Apollo is a trademark of Apollo Computer, Inc. IBM PC/AT is a trademark of International Business Machines Corporation.

CIRCLE NUMBER 16

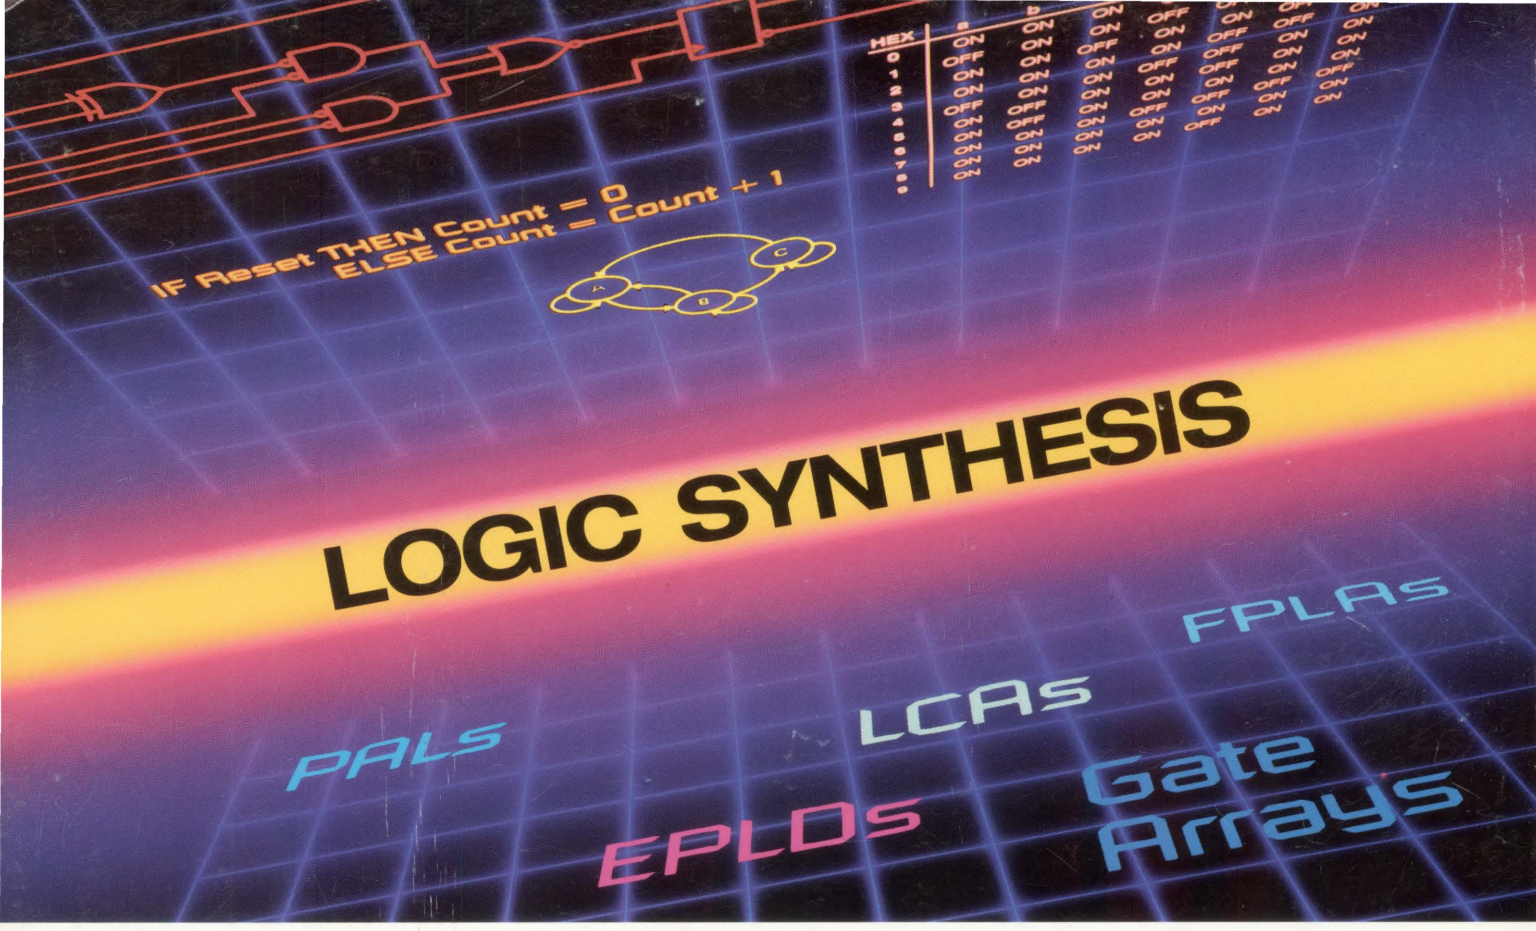

LOGIC SYNTHESIS GIVES YOU MORE DESIGN CHOICES.

FutureDesigner™ gives you more choices than any other design entry software—choices in how you enter your design, in target technologies, and in design output. And only FutureDesigner uses logic synthesis to automatically turn your input choices into your output choices, optimizing and streamlining your design for the technology you select.

CHOOSE THE DESIGN ENTRY METHOD. Only FutureDesigner lets you describe your design in the easiest, fastest, most natural way. You can enter some functions structurally, using DASH schematics. Others can be described behaviorally with any combination of truth tables, state diagrams, or high-level logic equations. Interactive verification and design rule checking help you catch errors up front, as you design.

CHOOSE THE TARGET TECHNOLOGY. FutureDesigner is technology independent. After you've described your design, you can choose any mix of TTLs, PLDs, LCAs, gate arrays, or other ASIC devices for implementation. It's also easy to migrate designs from one technology to another—for example, from TTL to PLD, PLD to LCA, or PLD to gate array.

Choose the platform: FutureDesigner runs on 80386 and 80286 machines, IBM® personal computers, and the Sun-3 Series.

CHOOSE THE OUTPUT FORMAT. With more than 100 DASH-Partners providing a broad range of complementary products and services, FutureDesigner's industry-standard format is accepted virtually everywhere. When you design with FutureDesigner, you'll have more choices in technologies, CAE systems, foundries, and service bureaus.

CHOOSE FUTUREDESIGNER WITH LOGIC SYNTHESIS. With its unique logic synthesis capabilities, FutureDesigner reduces and factors your design, eliminating redundancy and improving efficiency. It optimizes for the particular technology you've selected, making the necessary speed/size trade-offs. Then it generates the schematics, net lists, or JEDEC files for programming PLDs. Automatically.

Call us today for more information. Find out why FutureDesigner is the design entry software of choice.

1-800-247-5700
Ext. 130

Data I/O Corporation 10525 Willows Road N.E., P.O. Box 97046, Redmond, WA 98073-9746, U.S.A. (206) 881-6444/Telex 15-2167
Data I/O Canada 6725 Airport Road, Suite 302, Mississauga, Ontario L4V 1V2 (416) 678-0761
Data I/O Europe World Trade Center, Strawinskylaan 633, 1077 XX Amsterdam, The Netherlands +31 (0)20-6622866/Telex 16616 DATIO NL
Data I/O Japan Sumitomoseimei Higashishinbashi Bldg., 8F, 2-1-7, Higashi-Shinbashi, Minato-ku, Tokyo 105, Japan
(03) 432-6991/Telex 2522685 DATAIO J

FutureNet
A Division of Data I/O Corporation

CIRCLE NUMBER 17