EIGHT DOLLARS A CMP PUBLICATION OCTOBER 1988

VLSI Systems Design

FOR DESIGNERS OF HIGH-PERFORMANCE SYSTEMS

TAKING THE MYSTERY OUT OF LOGIC SYNTHESIS

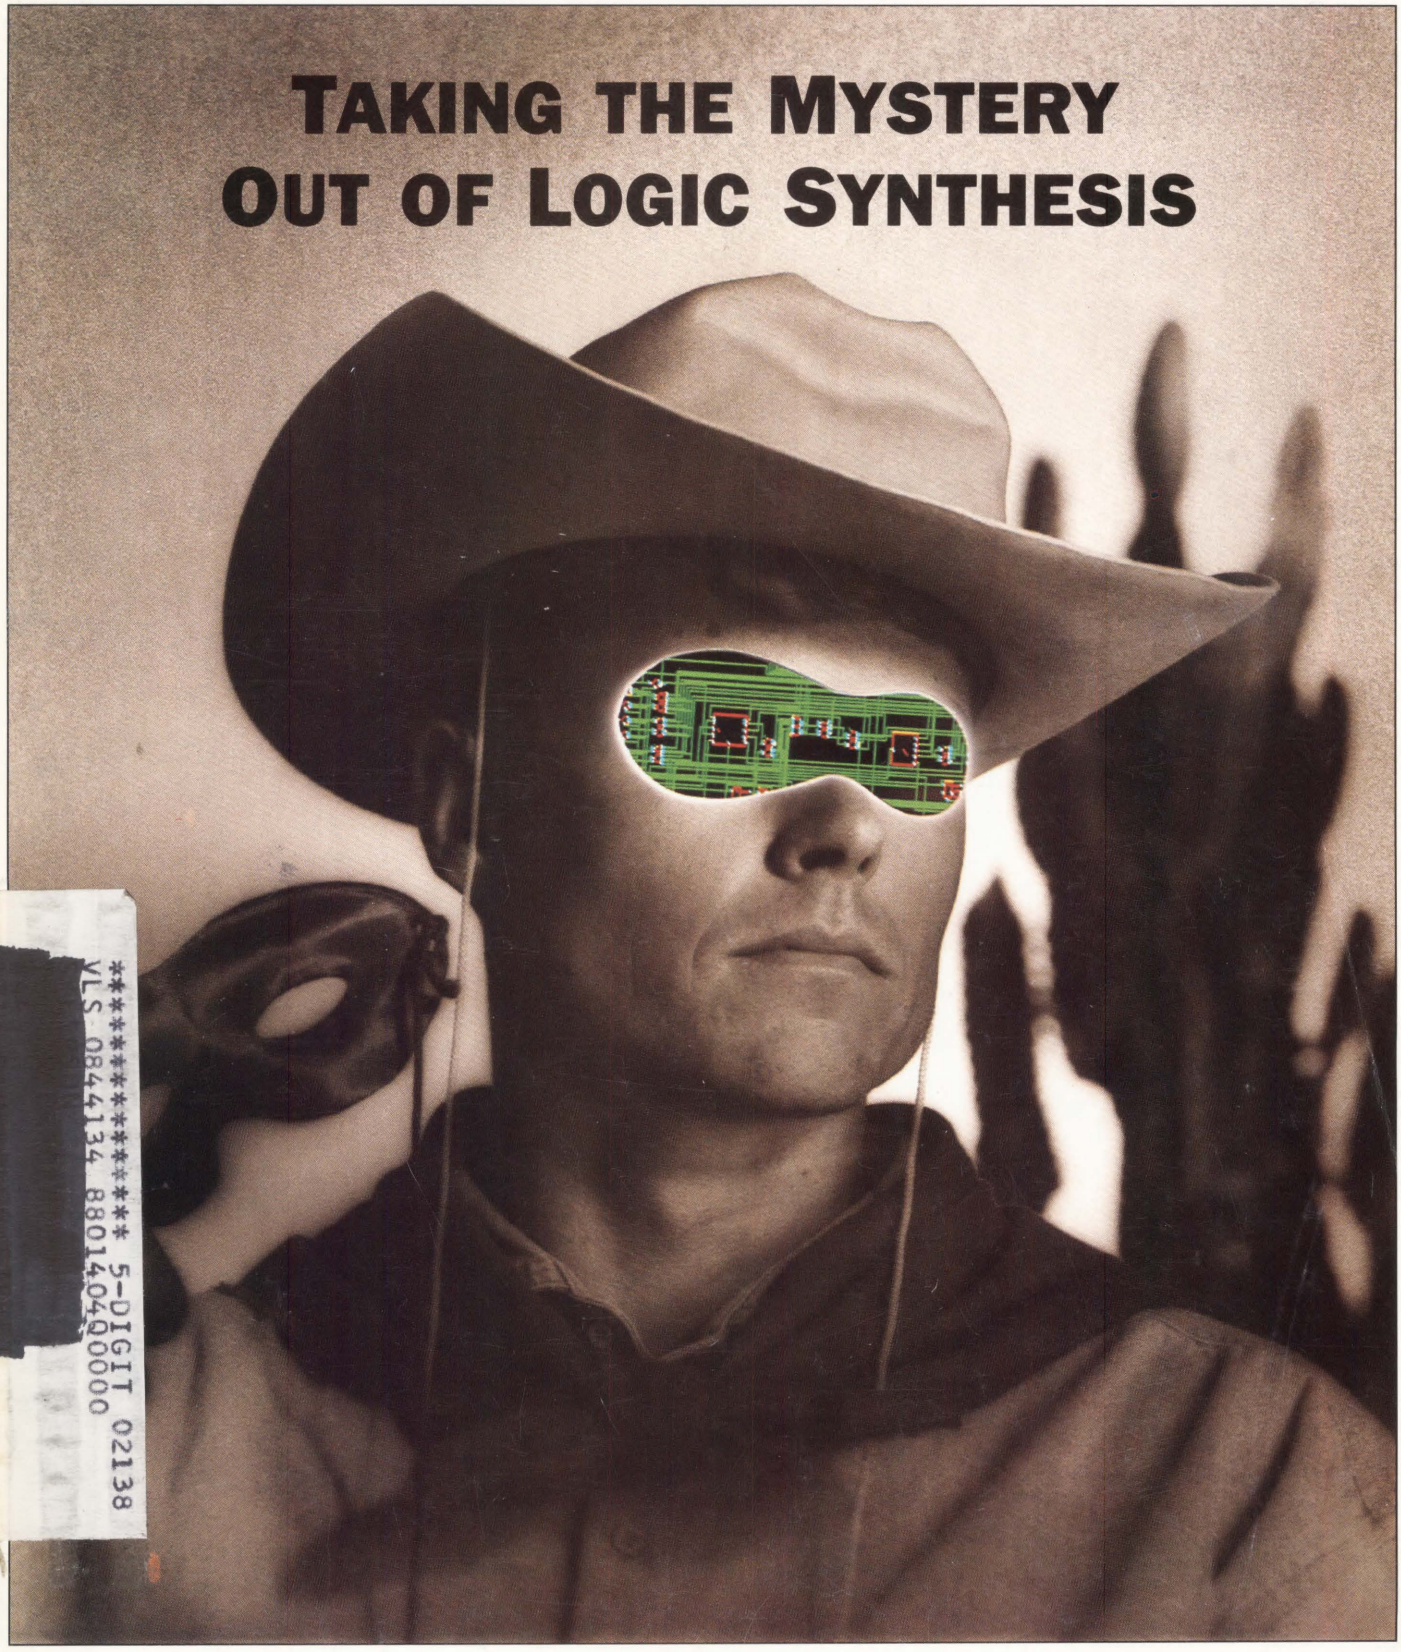

With the right ASIC par

Why shop around at different ASIC suppliers when you can get it all from Fujitsu.

But we'd rather you didn't think of us as just another ASIC supplier. Because, frankly, we don't want to be your ASIC supplier. We want to be your ASIC partner. The only one you'll ever need.

As the industry's most experienced volume producer of ASIC devices, we can provide you with the broadest options in ASIC.

Like standard and compiled cells. Over 70 types of channeled and channeless arrays. And 60 packages. All available now. All from one source.

You're also assured of the most advanced process and manufacturing technologies in the industry. Such as high-speed, silicon-gate 3-layer metal CMOS. Advanced ECL featuring 0.5 and 1.0 micron, 3-layer metal process. And a robust LSTTL bipolar process and design.

We've even taken advantage of our expertise in bipolar and CMOS to develop BiCMOS. A process marriage featuring high-performance ECL and high-density CMOS.

All of which means you can always count on getting the optimum solution to your VLSI needs. Whether you expand to larger gate counts or migrate from arrays to standard cells.

Take packaging, for example. You can choose from PGAs or DIPs. Ceramic or plastic. Through-hole and surface-mount. PLCCs and flatpacks. Even PGAs and flatpacks with heatsinks. All designed, built and backed by Fujitsu.

As your ASIC partner, we expand your design team, providing you with coast-to-coast technical resource centers, fully equipped with easy-to-use design tools.

You also get ASIC sales and marketing support to help you smooth out all the administrative wrinkles. Plus local training and technical support. And twenty-four hour design facilities, so you can work whenever inspiration strikes.

It all adds up to just one name. Fujitsu. The one ASIC partner you need. The only one you'll ever need. Which, after all, is everything an ASIC partner should be.

FUJITSU

FUJITSU MICROELECTRONICS, INC.

Everything an ASIC partner should be.

Fujitsu Microelectronics, Inc. • 3545 North First Street, San Jose, CA 95134-1804 • (408) 922-9000.

CIRCLE NUMBER 1

tner, you only need one.

Contents

ARTICLES

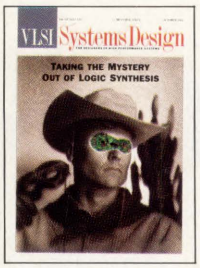

Logic synthesis, the latest design automation buzzword, has different meanings to different CAE/CAD suppliers

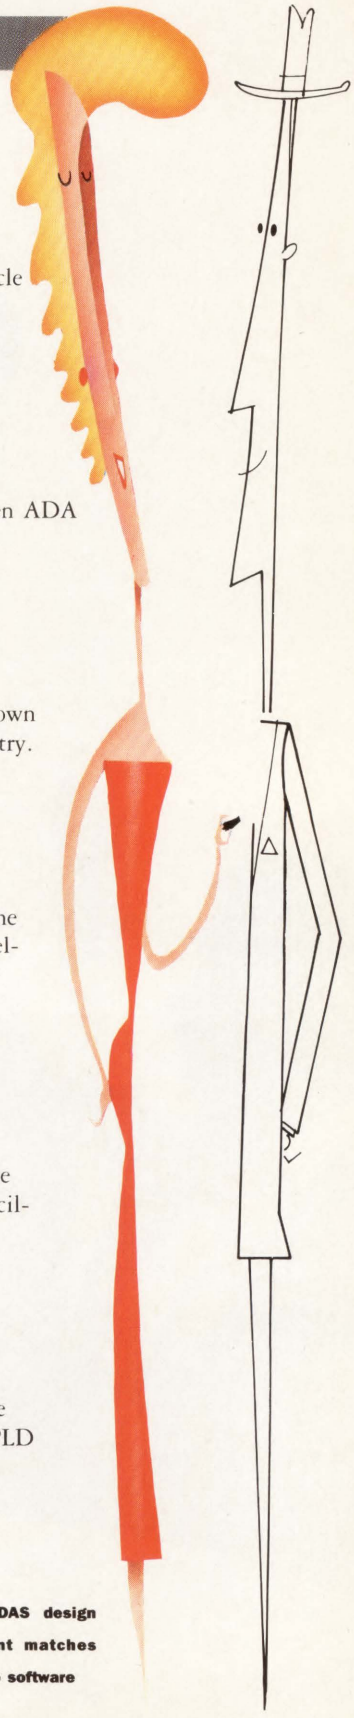

30
The new IDAS design environment matches hardware to software

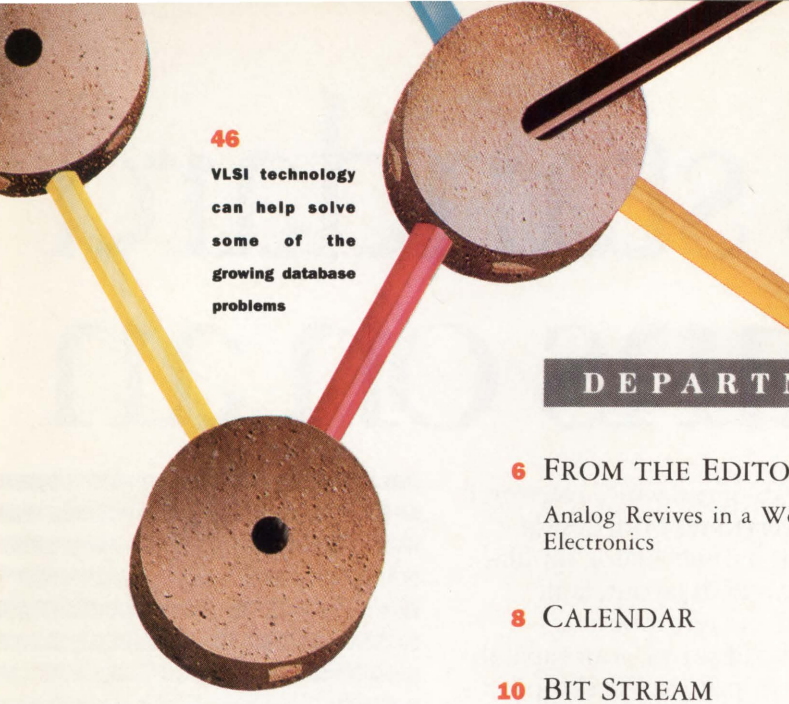

46

VLSI technology can help solve some of the growing database problems

DEPARTMENTS

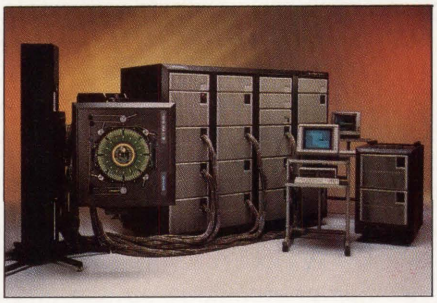

10

New VLSI chip testers support 512 pins with 200 MHz data rates

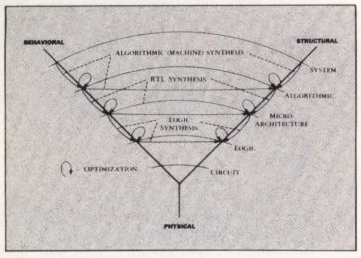

18

Silicon compilation, logic optimization and logic synthesis have many features based on common technological roots

COVER PHOTOGRAPH BY WILLIAM DUKE

VLSI Systems Design
FOR DESIGNERS OF HIGH PERFORMANCE SYSTEMS

A CMP Publication

EDITORIAL DIRECTOR
Robert W. Henkel

EDITOR-IN-CHIEF
Roland C. Wittenberg

MANAGING EDITOR
Mike Robinson

SENIOR EDITOR
Bob Cushman

SOLID STATE EDITOR
Roderic Beresford

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 September by CMP Publications, Inc., 600 Community Drive, Manhasset, NY 11030. (516) 562-5000. VLSI SYSTEMS DESIGN is free to qualified subscribers. One year subscriptions to others; US and Canada $60. Mexico, Europe, Central and South America: $120.00. Asia, Australia, and Africa: $135.00. 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

Daisy's standard now runs on an

Introducing the Advansys Series™ of high performance CAE workstations.

Now the electronic design environment you've dreamed of is here.

Because the most advanced design tools in the world now run on the world's most advanced standard platform.

The Sun 386*i*.™

We call it the Advansys Series. And it encompasses some of the most powerful design tools ever developed. Plus a variety of affordably priced workstations. Including the 20 MHz and 25 MHz Sun 386*i*, as well as Daisy's own LOGICIAN® 386 and the newly enhanced Personal LOGICIAN™ 386. All share a standard system level environment, featuring UNIX,™ advanced X Window System graphics, Sun's NFS™ distributed file system and standard TCP/IP communications.

Now you can get the workstation performance and flexibility you've demanded for your desktop. Including up to 5 MIPS of processing power, high resolution graphics display and an integrated UNIX/DOS environment.

Even better, all these advanced workstations run Daisy's field-proven Advansys software packages. Eight turnkey tool sets that meet all the demands of real world electronic design. Everything from design entry to digital and analog simulation, IC and PCB layout, fault simulation and test tools.

But with Advansys, your capabilities don't stop at your desktop. Because Daisy's unique network computing concept lets you create an affordable team design environment incorporating a wide range of powerful network resources.

Like Daisy's MegaLOGICIAN®—the most widely used simulation accelerator ever created. Or the brand new GigaLOGICIAN™ with up to 30 times greater performance.

For complex system simulations, there's Daisy's PMX,™ the most popular physical modeling system in use today.

You can also link with network servers like the Sun-4™ based XL Server for analog simulation or PCB routing. And Daisy lets you access all this power simply by opening a window on your Advansys desktop workstation. Eliminating file transfers and other time consuming bottlenecks.

To find out more about the new Advansys Series, call Daisy today at 1 (800) 556-1234, ext. 32. In California: 1 (800) 441-2345, ext. 32.

We're raising the standard of excellence for electronic design.

European Headquarters: Paris, France (1) 45 37 00 12.
Regional Offices: England (256) 464061;
West Germany (89) 92-69060; Italy (39) 637251.

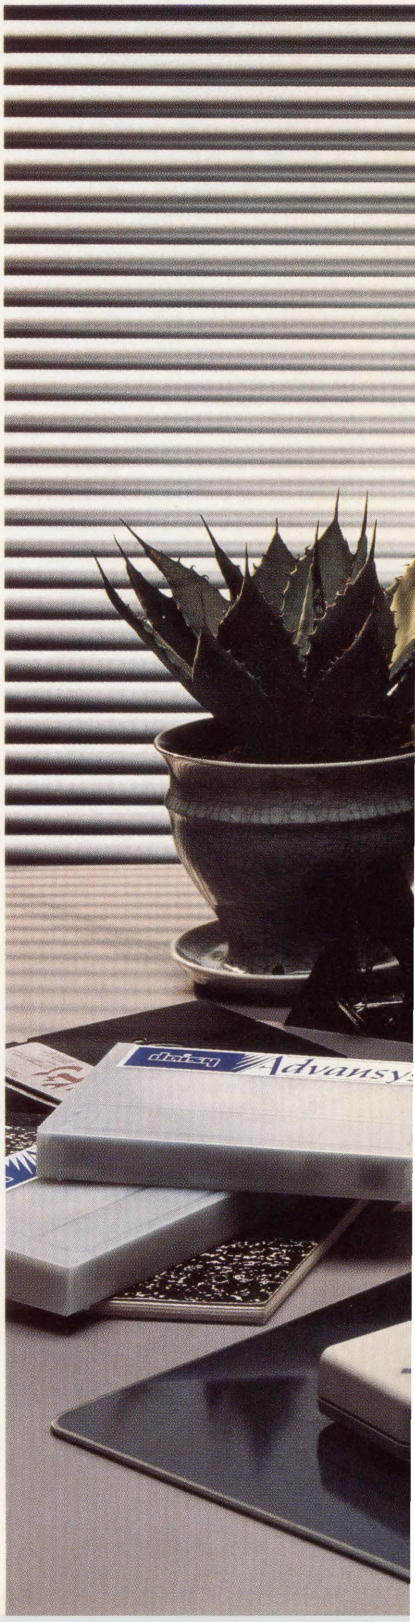

daisy SYSTEMS CORP.

# of excellence
excellent standard.

Analog Revives in a World of Digital Electronics

It seems that the more things change, the more they remain the same

A number of events during the past few years started me thinking about the fortunes of analog engineers and the changing environment that they live in. These included the growth of mixed-signal analog/digital systems, high-speed digital circuits that obsoleted simple logic simulators as far as timing and glitch behavior was concerned, and the increasing number of reader requests for more analog articles.

However, before continuing, I must confess that I started out in the analog world, which might explain my bias. Analog has always served industry and humanity well—from radio and TV to auto-pilots and industrial controls. But when digital electronics became a practical technology, many people started touting digital as the only way to go—anything you could do in analog, you could do better in digital. Well, not quite. The first digital replacements for analog systems were extremely costly, bulky, and slow compared to their analog cousins. However, the foresight of the digital proponents was eventually recognized when reliable, faster, denser, and cheaper digital chips became available. Digital was the winner and the digital engineer was king.

But the tide has turned again. As improvements in VLSI and ASIC technology accelerated, more and more of complex systems were put into the same box, or on the same board, or even on the same chip. It became evident that it was inefficient and costly to keep the analog components and systems separate. So they moved in with the digital domain. The situation was compounded by digital's move into traditionally analog worlds. Added to that was the fact that high-speed digital circuits were beginning to look like RF/microwave circuits, where careful consideration of transmission line effects, stray impedances, matching terminations, and stability were of crucial importance to a successful design.

Suddenly, there seemed to be a dearth of analog engineers, particularly RF/microwave specialists, who thrived on these kinds of problems. This has raised the visibility-level of the analog engineer, who may never be king again, but perhaps a prince consort to the world of digital electronics.

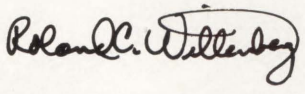

ROLAND WITTENBERG
EDITOR-IN-CHIEF

Access to the Right Technology

Raytheon offers you choices

Whatever your application, Raytheon can open the door to the appropriate ASIC technology. With over 14 years' experience in semicustom design, Raytheon offers you many choices—CMOS, low power ECL and digital bipolar, as well as linear arrays.

☐ **CMOS arrays:** Raytheon's low-power CMOS arrays to 20,440 gates have a 1.0 μm effective channel length and rad hard capability. Both our VHSIC-approved and standard families have the lowest power consumption in the industry. The 2 μm RL7000 array is a second source for LSI's LL7000 series. Standard cell solutions are also available.

☐ **Low-power ECL arrays:** The ECL products available from Raytheon have densities comparable to 1.5 μm CMOS arrays and I/Os compatible with ECL, TTL (10K or 100K), and CMOS. These ECL arrays also deliver the industry's lowest power dissipation per gate. Their superior speed/power performance—<0.1pJ—provides the ASIC designer with state-of-the-art semicustom capabilities.

☐ **Digital bipolar arrays:** Our many years of experience in ISL array design and production give Raytheon a solid foundation for expanding ASIC technology.

☐ **Linear arrays:** Raytheon's bipolar macrocell arrays with 8, 12 or 15 gain blocks and FET input arrays with 4 or 12 gain blocks utilize precision thin-film resistors. They can be configured into a variety of analog functions, including voltage references, switches, comparators, op amps, and VFCs.

Let Raytheon open the door to a world of ASIC choices, with our proven track record in multiple ASIC technologies.

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

Raytheon

Where quality starts with fundamentals.

CIRCLE NUMBER 2

Calendar

1988 Workshop on VLSI Signal Processing

November 2–4
Monterey Plaza Hotel
Monterey, Calif.

Providing a forum for the discussion of new theoretical, and applied developments in signal processing for very large scale ICs is the objective of this workshop. Topics will be discussed in such areas as analog/digital circuits, cell/silicon compilers, design methodologies, testability, algorithms, languages, simulation, functional mapping, architectures, fault tolerance and testing, coding, radar, telecommunications, and data communications. Additional information may be obtained by contacting Paulette E. Powell, University of California, Electronics Research Laboratory, 520 Cory Hall, Berkeley, Calif. 94720. ■

ICCAD-88

November 7-10
Santa Clara Convention Center
Santa Clara, Calif.

This international conference, sponsored by the IEEE Computer Society and the IEEE Circuits and Systems Society, in cooperation with the IEEE Electron Devices Society and the ACM Special Interest Group on Design Automation, is oriented toward the electrical engineering CAD professional. It will feature tutorials, and panel discussions, as well as technical sessions. Session topics will include statistical design techniques, high-level synthesis, timing simulation, channel routing, device simulation and modeling, performance issues for VLSI, circuit verification, row-based CMOS cell generation, automatic test pattern generation, and analog layout. For additional information, contact MP Associates Inc., 7490 Clubhouse Rd., Suite 102, Boulder, Colo. 80301. (303) 530-4562. ■

1988 Government Microcircuit Applications Conference

November 8-10
Riviera Hotel
Las Vegas, Nev.

GOMAC '88 is a government-sponsored conference established primarily to review developments in microelectronics applications for government systems. This year's conference, with the theme, "International Competitiveness: Its Impact on Government Electronics," will feature technical sessions on such topics as radiation effects in electronics, signal processing, VHSIC technology and applications, MIMIC, reliability, digital systems applications, packaging, discontinued parts, systems, and testability. For more information, contact C. Edward Holland Jr., GOMAC-88 Technical Program Chairman, Semiconductor Research Corp., 4501 Alexander Dr., P.O. Box 12053, Research Triangle Park, N.C. 27709. (919) 541-9400. ■

Custom Integrated Circuits Conference

May 15–18, 1989
San Diego, Calif.

CICC '89 is sponsored by the IEEE's Electron Devices Society, and Solid State Circuits Council, and co-sponsored by the IEEE Rochester Section. Its goal is to bring together designers and manufacturers of circuits and systems, and users of custom ICs to discuss new developments and future trends in custom integrated circuits. Papers are invited on such topics as digital signal and data processing applications, fabrication technology, physical design techniques, simulation and modeling, custom and semi-custom circuits, CAD systems and methodologies, custom interfaces and packaging, testing and reliability, and analog circuit techniques. By November 11, 1988, authors should submit 25 copies of a review package which includes a 1-page summary, a 35-word abstract clearly describing the work, and a maximum of 2 pages of figures, drawings and references. Submit review packages to Mrs. Roberta Kaspar, Technical Program Coordinator, CICC '89, 20 Ledgewood Dr., Rochester, N.Y. 14615. (716) 865-7164. ■

26th Design Automation Conference

June 25–29, 1989
Las Vegas Convention Center
Las Vegas, Nev.

DAC '89, which is devoted solely to the field of design automation, will offer tutorials, panel discussions, and technical presentations. Areas of interest include electrical and discrete simulation, timing verification, test validation and testability analysis, floorplanning and placement, global and detailed routing, physical module generation, symbolic layout and compaction, logic synthesis and optimization, behavioral and hardware description languages, design systems and databases, and computer aids for manufacturing. Interested authors should submit, by November 7, 1988, one cover page and 8 copies of a complete manuscript to MP Associates Inc., Att: A. Richard Newton, Program Chairman 26th DAC, 7490 Clubhouse Rd., Suite 102, Boulder, Colo. 80301. For additional information call (303) 530-4333. ■

Illustration by Lorraine Raywood

Saratoga FIFOs.

Our new FIFOs are the world's fastest. Available now in 10, 15, 25, 40 and 50 MHz.

Slow interprocessor communication headaches. You thought you'd tried every buffering remedy in the book to get rid of them.

But here's one you haven't: *raw FIFO speed.* Using Saratoga's new family of BiCMOS FIFOs—the world's first 50-MHz first-in, first-out memories.

Organized as 64 words by-4 and by-5 bits wide, these RAM-based devices deliver performance unmatched in the industry—at 10, 15, 25, 40 and 50 MHz.

Even so, they consume no more power than CMOS FIFOs, while offering high output drive that's TTL compatible. And they can be cascaded to expand in word width and depth. Plus they're available in both commercial and military temperature ranges, in industry-standard pin-outs.

This new generation of FIFOs will soon include 64 by 9 and larger density 512, 1K and 2K by 9 devices. Joining Saratoga's existing lines of high-performance TTL and ECL static RAMs—also among the fastest now available. And all made possible by our proprietary BiCMOS technology—SABIC™—which combines the best of both the bipolar and CMOS worlds.

So if system timing headaches have got you down, take one of our new FIFO buffers. And call us in the morning: (408) 864-0500. Or write: Saratoga Semiconductor, 10500 Ridgeview Court, Cupertino, CA 95014.

Saratoga FIFO Memories	
Clock Frequency	50 MHz (40 MHz military)
Data Access Time	15 nsec
Data Set-up and Hold Time	3 nsec
Bubble-through Time	25 nsec (max)
Power Consumption	385 mW
Output Drive	16 mA

Introducing The First 50-MHz FIFOs. In BiCMOS.

Saratoga Semiconductor

The Technology Leader in BiCMOS

CIRCLE NUMBER 3

BIT Stream

RISC Microprocessor Sales to Top $500 Million

Total revenues for 32-bit RISC chips are predicted to grow from $17 million in 1987 to $505 million in 1992, a compound average growth rate (CAGR) of 96.3 percent—according to a recently published study by The Information Network, a San Francisco-based market research company. This is four times the growth rate (24.7 percent) expected for 32-bit CISC chips, over the same period of time. The study also found that the manufacturers of RISC chips, both 16- and 32-bit implementations, have focused their attention on embedded controller applications, where the increased response speed of RISC-based systems is a must for realtime operation. ■

SPARCs Fly

Motorola has sued the developer of its 88000 chips, Roger Ross, and received a restraining order preventing Ross from using Motorola proprietary information when creating variations of the SPARC architecture.

Competition in the SPARC arena has intensified with Texas Instruments' licensing of-SPARC from Sun Microsystems. TI's now an alternate source for Cypress Semiconductor's SPARC implementations. In addition, startup Solbourne Computer (Longmont, Colo.) announced that it will be developing a 64-bit implementation of SPARC with the help of Matsushita Electric Industrial Co. of Japan. ■

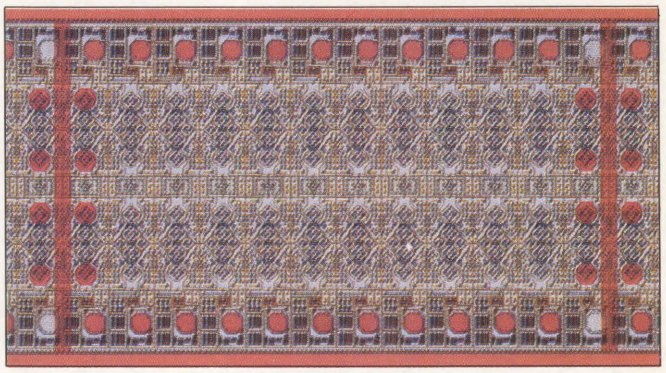

Cut Your Analog/Digital Arrays to Size

Exar Corporation (San Jose, Calif.) has enhanced its Flexar Series of analog/digital arrays to provide programmable die sizes. Exar fixes only the vertical dimension of the die and places bonding pads and I/O circuits at the vertical extremes. No bonding pads interrupt the horizontal rows of active components that result. After the design is laid out across the rows, bonding pads and scribe lines are placed at either end of the layout. Finally, the die are cut to size, achieving close to 100 percent die utilization.

The new Flexar Delta Family of arrays use active components, which are individually programmed with metal interconnects to operate as NPN or PNP transistors. ■

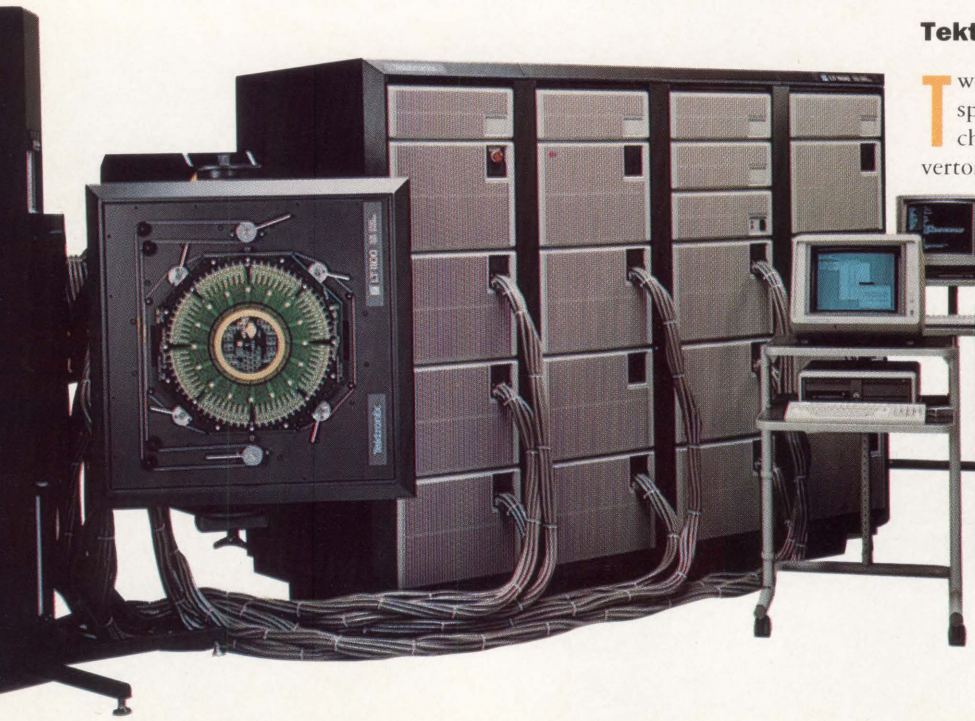

Tektronix Testers Toil at 200 MHz

Two new VLSI test systems—targeted at high speed, high pin count CMOS, ECL, and GaAs chips—were unveiled by Tektronix Inc. (Beaverton, Ore.). Both new testers, the LT-1100 and LT-1100+, support up to 512 pins and deliver pattern rates at 100 and 200 MHz, respectively. The LT-1100+ can deliver 100 MHz stimulus signals without multiplexing, and with an overall placement accuracy of ±275 ps, while the LT-1100 delivers 50 MHz signals with ±500 ps accuracy. Priced from $1.6 million, the new Vista Series of semiconductor testers were designed to compete with the Advantest 3381 and Ando 9037. ■

GaAs for the Masses

To convince designers that GaAs will be economically competitive with silicon, Gazelle Microcircuits Inc. (Santa Clara, Calif.) has forecast that its 10 nsec GA22V10 GaAs PLD will be priced at $13 each in 1990. This price is based on 10,000-piece orders, and represents a 58 percent drop from today's prices, which were reduced 27 percent just last month.

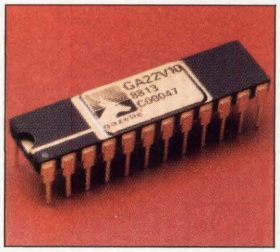

David MacMillan, Gazelle co-founder and vice president of marketing, said, "This lower price reflects the first step on our learning curve for gallium arsenide. While the silicon industry is nearing the end of its learning curve, we are just beginning to accelerate on ours." For example, GaAs requires only half of the critical mask steps as silicon does.

Gazelle is wasting no time making faster and more flexible PLDs as well. The company's new GA22V10-7, quantity priced at $37, has a maximum propagation delay of 7.5 ns—allowing the chip to operate with clock frequencies to 110 MHz.

Gazelle's new GA22VP10-7 is functionally compatible with the silicon 22VP10 from Texas Instruments Inc., but has less delay. The 22VP10 adds internal feedback paths to the 22V10 design to make registered outputs bidirectional. ■

Communications Controller Fits ISDN

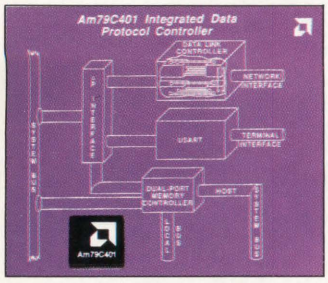

The Am79C401 Integrated Data Protocol Controller (IDPC) from Advanced Micro Devices Inc. (Sunnyvale, Calif.) integrates key components for building communications processors and terminal adapters for HDLC-based packet networks, including ISDN, X.25, the System Network Architecture (SNA), and the Digital Multiplexed Interface (DMI). It is designed to work seamlessly with AMD's current family of ISDN devices.

The IDPC includes a High-level Data Link Controller (HDLC), a USART, and a dual-port memory controller. Comprehensive software development packages from AMD, such as its AmLink3, help designers implement the layer-two functions, according to the seven-layer ISO model of ISDN (with LAPD protocols) and X.25 (LAPB) applications. It comes in a 68-pin PLCC or LCC package for $18 in 100 quantities. A plug-in evaluation board for personal computers is also available from AMD for $1,750. ■

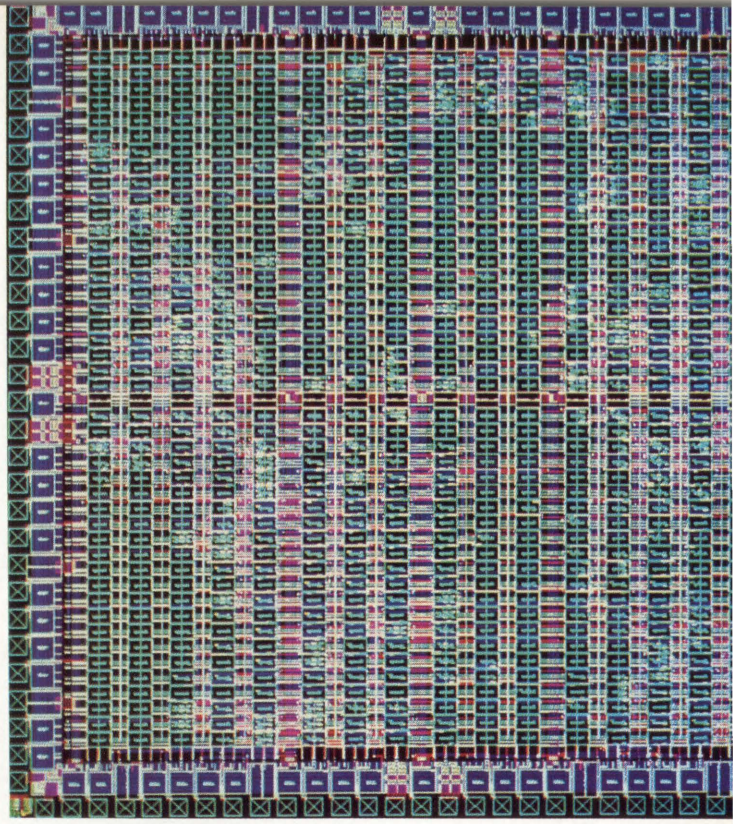

TriQuint Toggles 4,200 GaAs Gates at 100 GHz

Up to 4,200 equivalent gates are available in TriQuint Semiconductor Inc.'s (Beaverton, Ore.) new GaAs gate array family that can operate in applications requiring toggle rates as high as 1 GHz. The three new arrays, which are fabricated in a 1 micron Enhancement/Depletion MESFET process, provide 2,000, 3,000, or 4,200 gates with internal support for 40, 64, or 84 dedicated high-speed I/O pins, respectively. The I/O pins can be programmed to interface with TTL, CMOS, or ECLlogic. ■

Dataprobe Creator Builds 1.2-GHz Test System

Using electro-optical connections to a device under test, the System E/O from Photon Dynamics Inc. (San Jose, Calif.) is able to test devices at speeds of 1.2 GHz, with an overall timing accuracy of 50 psec and resolution of 10 psec. Creator Francois Henley, shown with Photon Dynamics President Dick Barr, was the developer of Mitsui's laser-based Dataprobe engineering characterization test systems. Based on Henley's Ph.D. work, the System E/O uses a Nd:YAG laser and electro-optical sampling techniques to determine output voltage without an electrical connection to the DUT pins. ■

Odds are 50-50 your perfect ASIC is a perfect dud the first time you plug it in.

That's why Mentor Graphics lets you combine ASIC and board circuitry in a single simulation.

Trouble in ASIC paradise.

The big day has arrived.

Your first gate array is back from the foundry. With high expectations, you plug it into your board and power up.

It doesn't work.

Don't feel alone. Over 50% of ASICs aren't operational when first installed in their target system. Even though 95% pass their foundry tests with flying colors.

An immediate solution.

Mentor Graphics shifts these even odds heavily in your favor with our QuickSim™ logic simulator, which lets you simulate both your ASIC and board circuitry in a single run.

With QuickSim, you not only track the internal operations of your ASIC circuitry, but also its transactions with the system at large. If there's a problem, you see precisely where it's located, either inside or outside your ASIC. All in a single, interactive simulation environment, where you can view and graphically "probe" the circuitry created by our NETED™ schematic editor.

Check out our libraries.

Library support is an ideal benchmark to gauge the true worth of an electronic design automation system. The more diverse and plentiful the component modeling libraries, the greater the design capability. It's as simple as that.

By this simple, yet decisive measure, Mentor Graphics brings you unequaled design capability. While other EDA vendors scurry to produce their own ASIC libraries (with little guarantee of accuracy), more ASIC vendors put their libraries on Mentor Graphics workstations than any other. And in most cases, we're the first workstation supported, which means you have the first shot at exploiting new chip technologies.

With Mentor Graphics, you get a breadth of LSI and VLSI component models, both hardware and software based. All of which can be mixed with ASICs in a single simulation that cuts your run time to an absolute minimum.

To be continued.

So much for the present. We're already developing new systems EDA tools that will extend to every dimension of electronic product development. From high-level systems descriptions to CASE. It's what our customers expect. It's what we'll deliver.

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.

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

Sydney, Australia; Phone 612-959-5488 **Mississauga, Ontario;** Phone 416-279-9060 **Nepean, Ontario;** Phone 613-828-7527 **Paris, France;** Phone 33-1-39-46-9604 **Munich, West Germany;** Phone 49-57096-0 **Neu-Isenburg, West Germany;** Phone 49-6102-25092-94 **Hong Kong;** Phone 852-566-5113 **Givatayim 53583, Israel;** Phone 972-777-719 **Milan, Italy;** Phone 39-824-4161 **Asia-Pacific Headquarters, Tokyo, Japan;** Phone 813-505-4800 **Tokyo, Japan;** Phone 813-589-2820 **Osaka, Japan;** Phone 816-308-3731 **Seoul, Korea;** Phone 822-548-6333 **Spanga, Sweden;** Phone 468-750-5540 **Zurich, Switzerland;** Phone 411-302-64-00 **Taipei, Taiwan;** Phone 886-2-776-2032 **Halfweg, Netherlands;** Phone 31-2907-7115 **Singapore;** Phone 65-779-1111 **Bracknell, England;** Phone 44-344-482848 **Livingston Scotland;** Phone 44-506-412222 **Middle East, Far East, Asia, South America;** Phone 503-626-7000 **Helsinki, Finland;** Phone 358-0-45571 **Madrid, Spain;** Phone 34-1-754-3001

Mentor Graphics®

Your ideas. Our Experience.

ATE and EDA challenge an Ocean Engineer

Teradyne's Joe Lassiter Makes Up His Mind Fast

HE was only nine years old when he first made up his mind about his future. "From the minute I saw that movie in the Rialto Theatre in El Darado, Ark., I knew what I wanted to do," said Joe Lassiter, vice president and general manager of the Electronic Design Automation Group at Teradyne Inc. (Boston). The movie was *"Twenty Thousand Leagues Under the Sea,"* a Walt Disney production based on the Jules Verne novel. "There's nothing as magic as the ocean if you live in a town like El Darado," he added.

Joe Lassiter kept his career on a steady nautical course until 1974, when—as an assistant professor of Ocean Engineering at MIT—he again made up his mind quickly, and decided that Teradyne, a leading ATE manufacturer, was the new course he should follow.

His perseverance in keeping to his chosen course was very evident—even when he was 13. By this time his family had moved to Houston, Tex. and Joe had signed up as a deck hand on a research vessel operated by Texas A&M University's Oceanography department. He spent the first of many Summers working in the Gulf of Mexico. But in later years his activities were in more exotic locations like the Red Sea and the Persian Gulf.

Joe continued working at his first love and in 1969 received a BS in Engineering from MIT, quickly followed by an MS and PhD. in Ocean Engineering.

During the time he spent on the Gulf of Mexico, he was fascinated by the reduced wave activity in certain offshore areas when hurricanes slammed into the Gulf. Apparently a muddy bottom helped disperse the waves' energy. This phenomena was still on his mind ten years later and it became the topic of his Ph.D. thesis.

M.P. O'Callahan

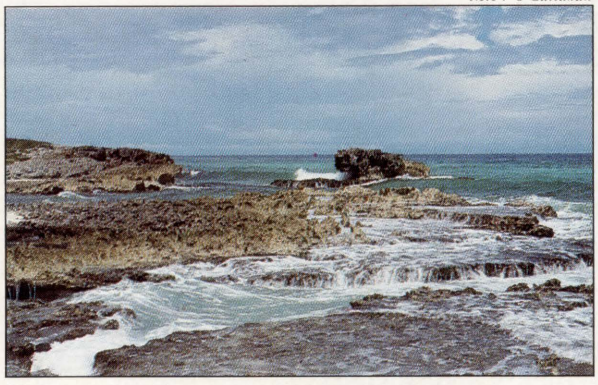

While teaching Ocean Engineering at MIT, he met Alex d'Arbeloff—president of Teradyne Inc. (Boston), an MIT alumnus, and a member of MIT's committee on the future of engineering education. d'Arbeloff convinced him to spend his sabbatical at Teradyne. "It only took me about six or seven weeks to make up my mind that my future was here with Teradyne," Lassiter observed.

While employed in ocean engineering, he made many technical and economic studies of projects such as off-shore drilling platforms. Later at Teradyne, this combination of engineering and economics served him well. The international aspects of his ocean engineering career also helped prepare him for the worldwide role that Teradyne products now play.

At his new company, he found a "tremendous opportunity to do a lot of things that can make a big difference to customers in terms of the value they get out of their investments. That means there are opportunities to try things that are interesting technically, opportunities to figure interesting ways to sell the customer, and opportunities to help the customer apply the product. It's a very rare combination.," said an enthusiastic Lassiter.

'FROM THE MINUTE I SAW THAT MOVIE

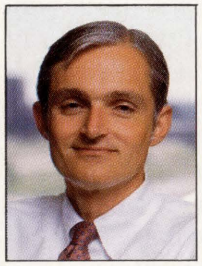

IN EL DARADO I KNEW WHAT I WANTED TO DO.'

A Ph.D. and an academic background didn't inhibit Joe at Teradyne. He quickly rolled up his sleeves and worked as a technician for six months to get acquainted with electronics from the ground up. He soon became product manager of the board test group, where he was involved in board testing, backplane testing, laser trimming, and simulation.

But most of his activity was concentrated in the marketing arena, where he had the chance to bring new products to the market and work with the engineering and sales departments. "I helped in getting products that met the customers' needs and also in getting feedback to the engineers so they could make the right technical trade-offs," he added.

Joe now oversees the company's design automation as well as the board test business. This includes Teradyne's recent acquisitions—Aida, Case Technology and Zehntel—moves that he strongly urged as a board member at Teradyne.

Fortunately, the challenge of his new career excites him enough so that he's never looked back at his ocean engineering days. But, whenever he gets a chance for a vacation, he always takes his wife and two young daughters to a spot near the water—because he still remembers that movie and still "loves to snorkel and fish." ■

—Roland Wittenberg.

Trying to design tomorrow's ASICs with yesterday's tools? Now there's ChipCrafter.

ChipCrafter™ is the integrated design tool that takes your complex CMOS ASIC design out of the stone-age and into the future.

High-level compilers, configurable libraries, process independence, and logic synthesis make ChipCrafter designs efficient. Automatic place and route, buffer sizing, and timing analysis at your Mentor Graphics™ workstation makes them easy, and as dense as hand-packed.

FREE. ASIC Estimating Kit.

What will it take to do your design with our cutting-edge design tool? Our free ASIC Estimating Kit lets you analyze design trade-offs, including performance and cost implications, in a variety of processes. Find out how ChipCrafter and Seattle Silicon chip away at design restrictions to deliver the next generation of ASICs.
Call for your free kit: 1-800-FOR-VLSI ext. 500.

SEATTLE SILICON
Launching the next generation of ASICs.

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

CIRCLE NUMBER 4

Does Logic Synthesis Work For Multi-Chip, High-Gate-Count Systems?

The effectiveness of logic synthesis is determined by more than speed and size

TOM MILLER, VICE PRESIDENT, SALES AND MARKETING,
INTEGRATED CMOS SYSTEMS, INC., SUNNYVALE, CALIF.

AT the 1987 Design Automation Conference, Andy Rappaport of the Technology Research Group Inc. said that logic synthesis represents the new frontier in electronic design automation, opening the way for a quantum increase in system complexity with only a small increase in design complexity. Looking at the crop of programs at this year's Design Automation Conference, the current product offerings don't fulfill that promise for high-density, multi-chip system designs.

The high-performance systems designer is looking for design tools that can significantly reduce the complexity of design while maintaining control of any variables that impact performance. Local optimizations that reduce logic are desirable but do not deliver the full benefits that logic synthesis is capable of. At Integrated CMOS Systems we have been applying logic synthesis approaches to high-performance systems for more than five years. We have identified three key leverage points that logic synthesis must observe to be beneficial in solving system design problems.

Logic synthesis can be focused on minimization of combinatorial logic, sequential logic (finite state machines), and hierarchical structures such as chip interconnects and functional boundaries. Today's logic synthesis tools optimize the combinatorial and sequential logic for each functional blocks within an individual chip.

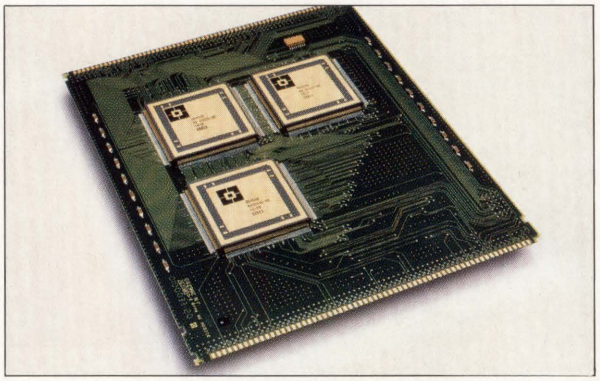

The key leverage points for high-performance, multiple-ASIC systems, however, are found in the optimizations at the boundaries of each functional block.

The first leverage point applies both global and local optimizations interactively to achieve the system design objectives. For this purpose, *user-programmable scripting* lets the user change the order of application of optimization routines and introduce application-specific transformations. User scripts tune the synthesis to the target system, reduce the time to process the logic, and help ensure that the synthesizer uses the same optimum sequence of operations each time it is invoked (an issue called repeatability).

The second leverage point is the importance of *traceability* of the original logic functions and signal names during the logic synthesis transformations. When logic networks take on different forms, engineering changes become a challenge if the original design disappears. In addition, designation of high-level structures such as multiplexers, tri-states and adders can be used for global transformations if the information is not lost during synthesis.

'FOR COMPLEX DESIGNS, LOGIC SYNTHESIS

DEPENDS ALMOST EXCLUSIVELY ON HEURISTICS'

The third leverage point is the *intelligent application* of logic synthesis to prevent synthesis from taking excessively long. Because the speed of processing increases with the number of gates, designers of systems with high gate counts must apply logic synthesis carefully to avoid excessive costs and project delays. For even a few thousand gates, the processing can take several hours. The best approach applies local transformations first, and then explores global optimizations in the remaining problem areas.

Logic synthesis can be an important addition to the designer's toolbox, controlling complexity while allowing exploration of more design options. Automation of combinatorial logic minimization is applied successfully today to the design of PLAs, PLDs and even some low-density gate-array and standard-cell designs.

But for more complex high-level designs, logic synthesis depends almost exclusively on heuristics. Unfortunately, significant development in heuristics still remains before logic synthesis can deliver the capabilities required for multi-chip, high-gate-count system design. ■

TOM MILLER *is the vice president for Marketing and Sales at Integrated CMOS Systems, Inc. (Sunnyvale, Calif.). Previously, he was director of Strategic Marketing at Fairchild Semiconductor, director of VLSI at NCR Corporation, and also a product-line manager at Texas Instruments. He holds a BS in Computer Science from the State University of New York at Stony Brook.*

Previously, the best CAE tools were available on only one platform.

From the leader in desktop CAE solutions comes a new concept in electronic design: you pick the platform.

Viewlogic® was the first to deliver a premium CAE solution on 286/386 PCs through "native mode" technology. Now we've ported the entire Workview® suite of CAE solutions—including VHDL simulation, a mixed analog/digital simulator, and ASIC vendor kits—to Sun and VAX™ workstations. For example, you can run a VHDL circuit simulation of 50,000 gates in minutes on a 386/PC, Sun Workstation® or VAXstation. Standardize on one, or pick a combination and design complex ICs, ASICs, and systems in a multiple platform environment—and retain a common user interface and shared database.

With Workview, you have access to superior CAE solutions on today's three most popular workstations in electronic design. For more information, call: **1-800-CAE-VIEW**; in Massachusetts, 1-508-480-0881.

VIEW*logic*®

desktop CAE and beyond

Sun 3/60
VHDL Simulation
50,000 gates/4,000 vectors
15 min 22 sec

80386 PC
VHDL Simulation
50,000 gates/4,000 vectors
15 min 14 sec

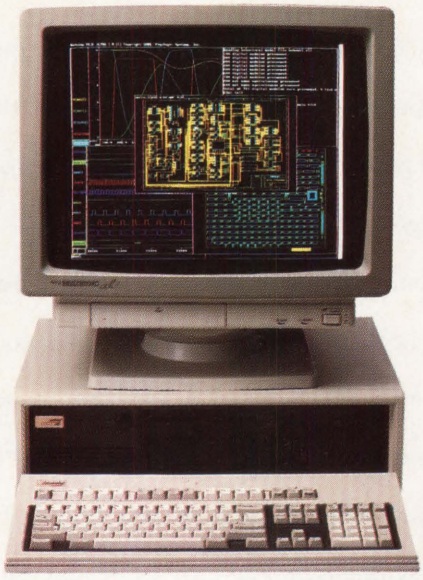

VAXstation 3000
VHDL Simulation
50,000 gates/4,000 vectors
15 min 31 sec

Now, take your pick.

Viewlogic Systems Inc., 313 Boston Post Road West, Marlboro, MA 01752

CIRCLE NUMBER 5

What is Logic Synthesis?

DAVID SMITH, SENIOR EDITOR, VLSI SYSTEMS DESIGN

Synthesis and compilation: mutual subsets?

The generally held definition for the term "synthesis" is the combining of parts or elements to form a whole, or compound. In electronics, synthesis has come to mean the building, by computer program, of a description of all or part of an electronic design—be it a chip, board or system. A logic synthesis tool, therefore, builds a logical description of an electronic design.

The confusion currently involving logic synthesis has several causes. First, the term "logic synthesis" is used popularly to represent a synthesis tool that constructs any type of electronic design, ranging from a PLD fuse map to a register-transfer block diagram. The second source of confusion concerns the term "logic." A gate-array logic design is composed of gates; a PLD logic design may be a fuse map; and a logic diagram of a board design contains 7400-series logic packages. A logic-synthesis program, therefore, can use any of these elements to build a design. Finally, as soon as logic synthesis became a buzzword, many vendors of design automation tools hung out the "logic synthesis" shingle simply by renaming existing products.

Since logic ranges from fuses to packages, and synthesis spans PLA mapping to algorithmic development programs, it's not surprising that logic synthesis is such a potent, and misunderstood, term. To the designers who stand to benefit from this relatively recent addition to design automation, the question is more than, "What is logic synthesis?" The question is, "What does

IS LOGIC SYNTHESIS A BREAKTHROUGH FOR DESIGN PRODUCTIVITY?

ILLUSTRATION BY GARNET HENDERSON

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 6

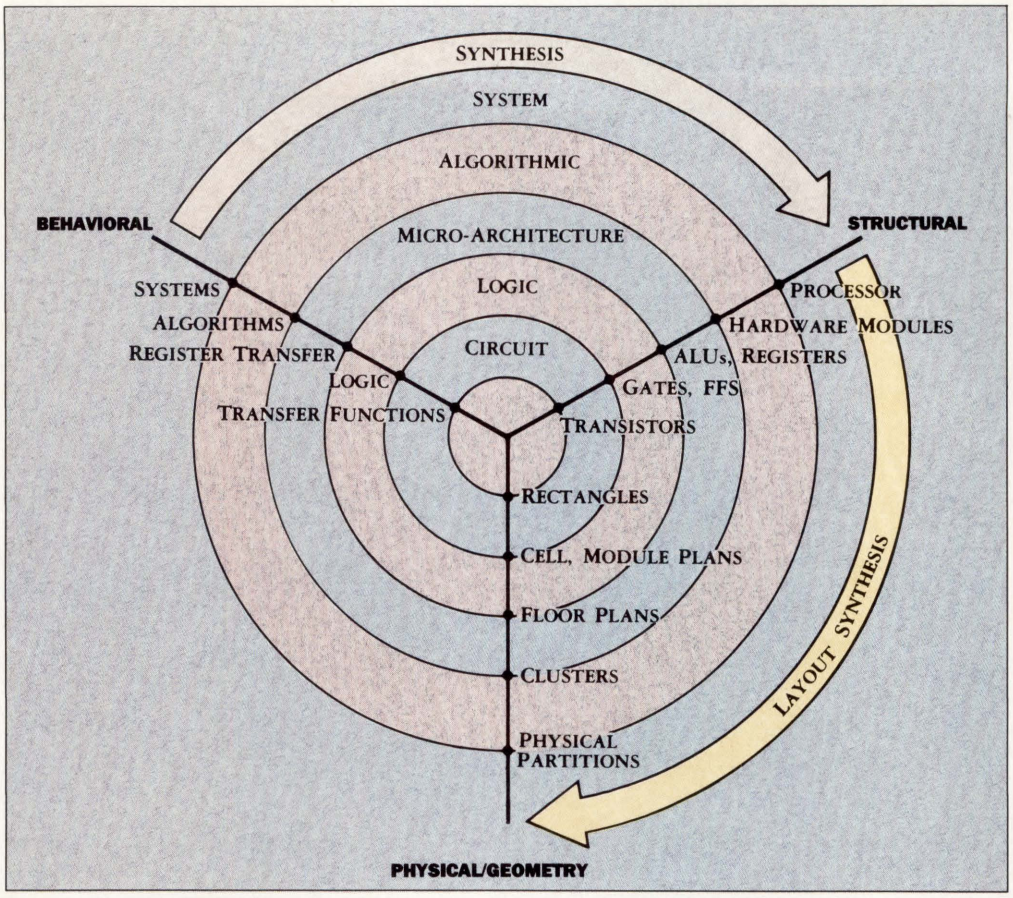

Figure 1. The Y chart proposed by Gajski and Kuhn plots the evolution of an electronic design through three domains of description—behavioral, structural, and physical. Levels of detail are represented by concentric circles. Creating a structural description from a behavioral one is popularly called synthesis or logic synthesis; creating a physical description from a structural one is referred to as layout synthesis.

this program do for me?"

■ MOVING ON THE Y CHART

To understand how logic synthesis is perceived, it's helpful to consult the Y chart (Figure 1) first proposed by Daniel Gajski and Robert Kuhn (Gajski and Kuhn, 1983). It contains three axes representing domains of electronic design description: behavioral, structural, and physical. The concentric circles represent levels of detail, such as micro-architectural, register-transfer level, and logic level. The end-result of any design is its physical implementation, e.g. for an IC, a fabrication database would be expressed in polygons.

The Y chart maps the progress of a design through different types of descriptions (domains on the axes) and progressive levels of detail (circles). For example, if a design is expressed as a set of Boolean equations, it can be translated into a netlist made up of logic gates from a gate array library. The process would be represented by an arrow from the behavioral to the structural domain, starting and ending on the logic circle. The design could then be placed and routed by layout programs, cess taking the design from the structural to the physical domain. "Fracturing" the design database for manufacture move the design to the circuit circle.

The Gajski chart, therefore, plots the evolution of a design through domains of description and levels of detail. It is conceptually useful in discussing logic synthesis because it provides a visualization of the mechanics and purpose of a proposed logic synthesizer. It indicates where the synthesizer starts, what its intermediate steps are, and what its output is.

The term "synthesis," in electronic design, means the building of a structure for a design based on its behavioral description. This can be shown as an arrow from the behavioral to the structural domain on the Y-chart. "Logic synthesis" is often used synonymously with "synthesis" to represent an entire class of tools that builds designs in the structural domain.)

Developers of CAD tools often use the word synthesis in the context of automatic generation of an IC layout, known as layout synthesis. This usage suggests that the modifier "logic" was coined to differentiate synthesis tools that build structural descriptions from those that build layouts.

The "building structure" definition is supported by such industry leaders as the Technology Research Group, Jeff Fox of Silc Technologies, and Gajski himself. The definition seems convenient when attempting to distinguish tools that build structure from those that improve it—optimization programs—and from those that create a physical design from a structural one—layout synthesis (which many perceive as silicon compilation).

Unfortunately, the type of *input* that a logic-synthesis tool accepts is not as cut-and-dried. For example, Professor Giovanni de Micheli at Stanford University's Center for Integrated Systems agrees that the term "synthesis" implies the creation of a structure. He adds, however, that the design input may not necessarily be behavioral: "Synthesis is closely related to structure, [generated] from either a function or another structure".

De Micheli is alluding to the fact that most designs are a combination of behavioral and structural descriptions rather than *pure* behavioral. Art de Geus, of Synopsys Inc., says that adding structural information directs the tools to produce efficient designs. "What we are is an RTL-level down system," says de Geus, referring to the register-transfer-level input (which has behavioral and structural components) that Synopsys' Design Compiler breaks down into a logic-level output.

Structural input is also important when the designer is using certain functions with well-known, dense implementations that logic synthesizers don't recognize. For example, Howard Moscovitz of AT&T explains that "logic synthesis tools can't yet select appropriate DSP-architectures, such as bit-serial, bit-parallel, or serial parallel" (Moscovitz, 1988). Structural statements can be used to preserve the structure when partitioning designs.

According to these experts, logic synthesis generates a structural design description from an input specification containing behavioral and structural information. But if the input is primarily structural, synthesis is when a more detailed structural design results, such as a flip-flop schematic resulting from a register-file input. A design tool is

COUNT PINS
COUNT COSTS
CALL CADIC

...FOR ASIC TESTING

This is the ASIC tester many Fortune 100 companies count on when they want to control costs. Capacities range from 32 to 352 pins, at prices as low as $163,000 for a 256 pin system.

BASIC ASIC

The CADIC STM5200 is straightforward, simple to learn and easy to use. A special utility, which generates test vectors from simulator output, makes testing even easier.

POWER AT THE PIN

All pins are fully bidirectional. A patented SoftWire™ system eliminates fixture building, makes all pins usable and expensive personality cards obsolete.

COST AND EFFECT

ASICs make your products competitive. They make ours competitive too. In fact, they're the secret to the low cost of this 25 MHz dynamic tester.

PIN US DOWN!

Free demo disk. Call 1-800-824-1617, ext. 5200.

CADIC

The Masters of ASIC Testing

SoftWire is a trademark of CADIC, Inc.

CIRCLE NUMBER 7

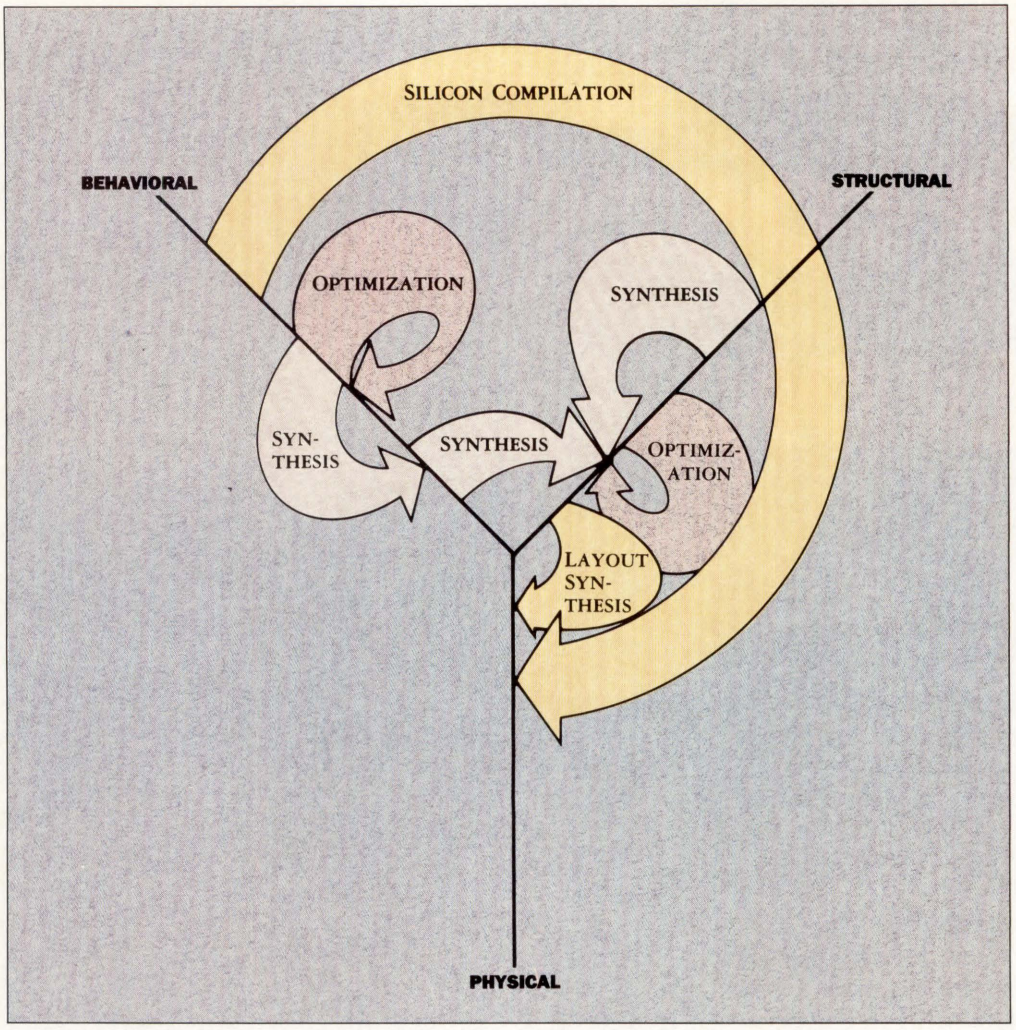

Figure 2. In addition to building structure from behavior, synthesis is also used to describe the creation of more detailed levels of description. If a tool begins and ends with the same type of description and level of detail, it's an optimizer. All synthesis and optimization tools are usually lumped within silicon compilation.

actually performing optimization (Figure 2) if its input and output have the same level of detail.

■ Is It Logical?

The modifier "logic" introduces other nuances to the term "synthesis." Specifically, the circle on the Y chart labeled "logic" represents a level of detail at which semicustom IC users develop design specifications. These users require the logic synthesis tools' *output* in a logic-level (Figure 3) structural form: either a netlist or a schematic. The definition that logic synthesis creates logic-level descriptions is bolstered by an analogy to the term "*silicon* compilation," the automated design of silicon-level circuits.

However, Gajski views logic synthesis as working strictly at the logic level, accepting logic-level *input*, such as Boolean equations, and creating a logic-level structural output. Similarly, Jeff Fox points out that logic synthesis starts with "a low-level behavioral description."

But Fox is trying to differentiate his company's tool, SilcSyn, from those that accept only logic-level inputs. SilcSyn is labeled as a "behavioral-level synthesis tool" because its primary mode of *input* is a behavioral language. Its output, however, is a structural design composed of gate-array or standard-cell *logic* elements. If the output defines the synthesizer, then SilcSyn should be a logic synthesizer.

The Ascyn Synthesizer from Algorithmic Systems Corp (Braintree, Mass.) supports this conclusion. This product accepts a behavioral description similar to that used with SilcSyn (they share the same progenitor: the MacPitts silicon compiler system). Still, Ascyn is presented as a logic synthesizer, creating logic-level netlists for gate arrays, standard-cell ICs, and full-custom ICs.

Users should beware when any other word besides "logic" appears before "synthesis." For instance, "behavioral synthesis" would seem to contradict the idea that synthesis creates structure. According to Fox, the term is used to indicate a logic synthesizer with a behavioral-language input.

On the other hand, de Micheli defines "behavioral synthesis" as *creating* a behavioral description, "making transformations which are independent of the structural model." In this case, the input and output may be of the same type, with synthesis creating a more detailed version.

In a related vein, "structural synthesis" can be viewed as creating structural descriptions that are composed of elements more abstract than logic-level elements. The JRS synthesis tool (see page 30), for example, can build a structural-design description using off-the-shelf ICs. But, the definition remains cloudy. For example, Misha Burich, vice president of engineering at Silicon Compiler Systems (San Jose, Calif.), writes that behavioral synthesis begins with a behavioral description of a system and "results in a structural description of a design, consisting of high-level logic elements, like registers and memories (Burich, 1988)." What, then, is structural synthesis?

These definitions may seem like exercises in semantics, but for the designer's sake, inputs should be distinguished from outputs. Given the analogy to silicon compilation, structural descriptions should issue from structural synthesizers, and behavioral synthesizers don't exist. Designers should ask, "What are the input and the output formats? What levels of detail is the program creating?"

■ Compilation

The third area of confusion results from the fact that many of the capabilities already found in compiler products are associated with logic synthesis. Gajski looks at the entire design process, from behavioral through structural and on to physical domains, as falling under the term "compilation"— with logic synthesis is a subset of silicon compilation.

SCS agrees. It's LogicCompiler is joined by VLSI Technology's Logic Synthesizer and Seattle Silicon's Finesse software in the synthesis-from-compiler product category. According to Burich, logic synthesis "provides an additional

Great things happen when System HILO™ is your simulation choice... *design and test efficiency soar!*

GenRad has solved the common dilemmas of design and test groups...which simulation system to buy for IC design?...which for PCB design?...can the test group use it?...does it all work together?

Now, for the first time, your design and test people can work with a common, integrated simulation tool set, regardless of whether they're working on design verification or test development for integrated circuits or printed circuit boards.

With GenRad's System HILO, design verification and test development for both chips and boards can take place at the same time. And this means optimum working efficiency between your design and test people. The result is tremendous savings in time and money, as well as better designs and more comprehensive test programs.

The key to System HILO is a new, modular architecture. It lets your design and test people solve a broad range of chip and board problems that cannot be addressed with any single tool. It enables design to proceed in parallel with test development. Utilizing System HILO's Test Waveform Language, the test engineer can use the same simulation data as the designer. The test program can then be downloaded and run without translation on GenRad's 275X board test systems. The result is faster, more comprehensive program development with maximum diagnostic effectiveness.

There are many more great aspects about this new simulation system from GenRad. Find out more by calling 1-800-4-GENRAD.

GenRad

The difference in software is the difference in test™

CIRCLE NUMBER 8

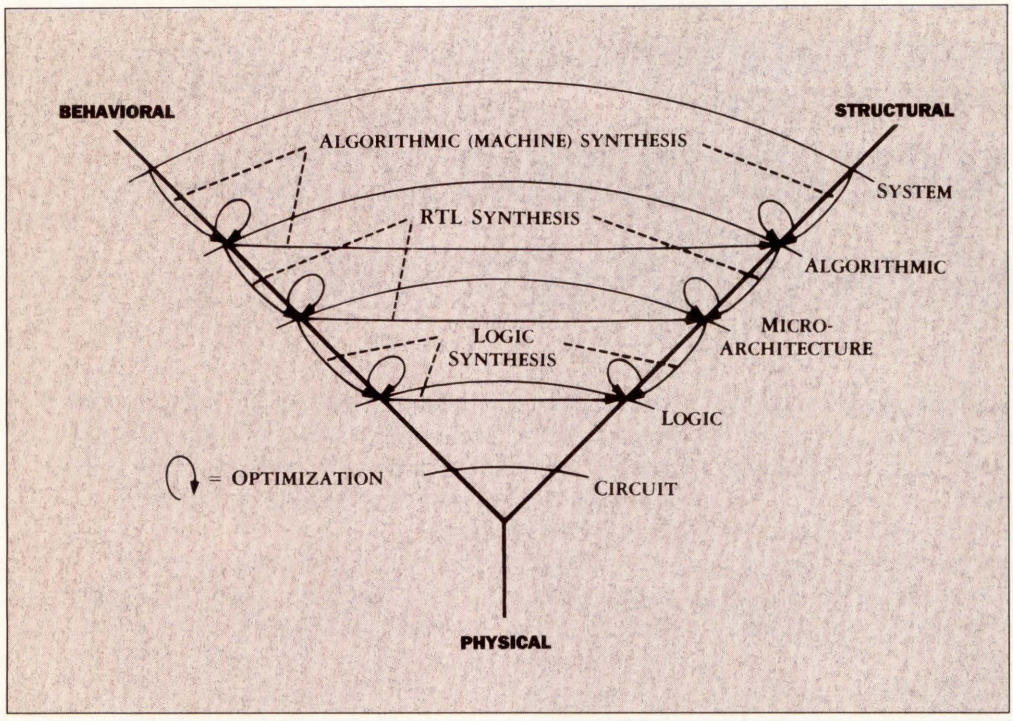

Figure 3. Synthesis tools are named for their output format: Logic synthesis produces a logic-level description. Synthesis tools creating more abstract descriptions should do the same, though this has not always been the case so far.

mechanism for synthesizing layout of parts of a chip or an entire chip."

The input specifications for SCS's logic-synthesis tool consist of RTL languages and finite state-machine descriptions. The structure of the design is largely specified, and the tool replaces a structural description with a more detailed structural description in "the primitives in a particular cell-based technology." The tool creates a logic-level design from a less-detailed structural description. Similarly, the Finesse program accepts equations, state machines and truth tables. The Logic Synthesizer outputs schematics just like the Design Consultant logic synthesizer from Trimeter Technologies Corp (Pittsburg, Pa.). Thus, these subsets of compiler environments may be viewed as legitimate synthesis programs.

A different perspective results from from the work of users closer to front-end, CAE design. Earl Reinkernsmeyer, director of software products for NCR's Microelectronics Division, has worked with Silc to bring SilcSyn into NCR's CAE tool kit. Because contemporary logic synthesizers can produce designs for a variety of custom/semicustom implementations—such as silicon compilation—then synthesizing a logic design can be considered a superset of compilation. "When silicon compilation has a high-level-langauge input, it's a subset of logic synthesis," Reinkernsmeyer says.

The choice of semicustom or custom implementations is often used to distinguish logic synthesis from silicon compilation. Silicon compilation must produce an output for a particular full-custom process. Logic synthesis produces netlists for gate arrays, cell-based ICs, or full-custom ICs (including compiled ones). Alice Parker of the University of Southern California concurs (Parker, 1984), saying that a silicon compiler accepts "logic-level, register-transfer or abstract behavior specification and produces an integrated circuit layout." She adds that they focus on IC layouts rather than synthesis.

De Micheli considers compilation and synthesis "just a question of terminology." SCS's Dan Payne says, "different words come into vogue in the marketing battle." Designers should expect the terms "synthesis" and "compilation" to be used synonymously, because eventually the products which come from both of these processes will offer many of the same features.

■ USING SYNTHESIS

Synthesis and compilation products will eventually form two legs of an overall IC design methodology. A particular design system at AT&T, for example, positions logic synthesis as just one element out of many (such as compilers, generators, and processor cores)in an overall system for designing custom DSP chips (Moscovitz, 1988). What distinguishes logic synthesis from silicon compilation in this system is the type of circuits produced. Synthesis is best for "less-common DSP functions or for random logic functions." Compilers are most effective for for highly regular structures like registers and data paths.

Logic synthesis' most important characteristic is how it fits in with the methods of the designer. According to Nancy Madison of Valid Logic Systems, her customers most often ask how it can eliminate certain steps in the design process. In addition to filling in the blanks regarding input format and output format and the level of detail under consideration, this question can best be answered by understanding how well logic synthesis will be working with the designer's other tools, such as compilers, in his or her overall IC design methodology.

Tom Waugh of Xilinx stresses the importance of productivity in deciding what logic synthesis is. "If a set of tools provides an open system that lets the designer choose the methods of design entry that best express a design, and if the set of tools provides optimization software that frees the designer from worrying about architectural details to concentrate on higher-level design issues, then these tools are logic synthesis tools," he explains. Although convenience and management of design detail may not be the best definition of logic synthesis, they may be the most concrete impressions designers get from logic-synthesis tools. ■

REFERENCES

BURICH, M., 1988, *"The Role of Logic Synthesis in Silicon Compilation," VLSI Systems Design, Semicustom Design Guide, September, 1988, pg 56.*

GAJSKI, D. and KUHN, R., 1983, *"Guest Editor's Introduction: New VLSI Tools," IEEE Computer, December, 1983, pg 11.*

MOSCOVITZ, H., 1988, *"A Hybrid Approach to Application-Specific DSP Design," VLSI Systems Design, Semicustom Design Guide, September, 1988, pg 32.*

PARKER, A., 1984, *"Automated Synthesis of Digital Systems," IEEE Design and Test, November, 1984, pg 75.*

TECHOLOGY RESEARCH GROUP LETTER, *Vol2, no. 11, February, 1987.*

TREVILLYAN, L., 1987, *"An Overview of Logic Synthesis Systems," Proceedings of the Design Automation Conference, June, 1987, pg 166.*

We're eliminating the competition with something everyone else seems to have forgotten you need...

...the maximum performan

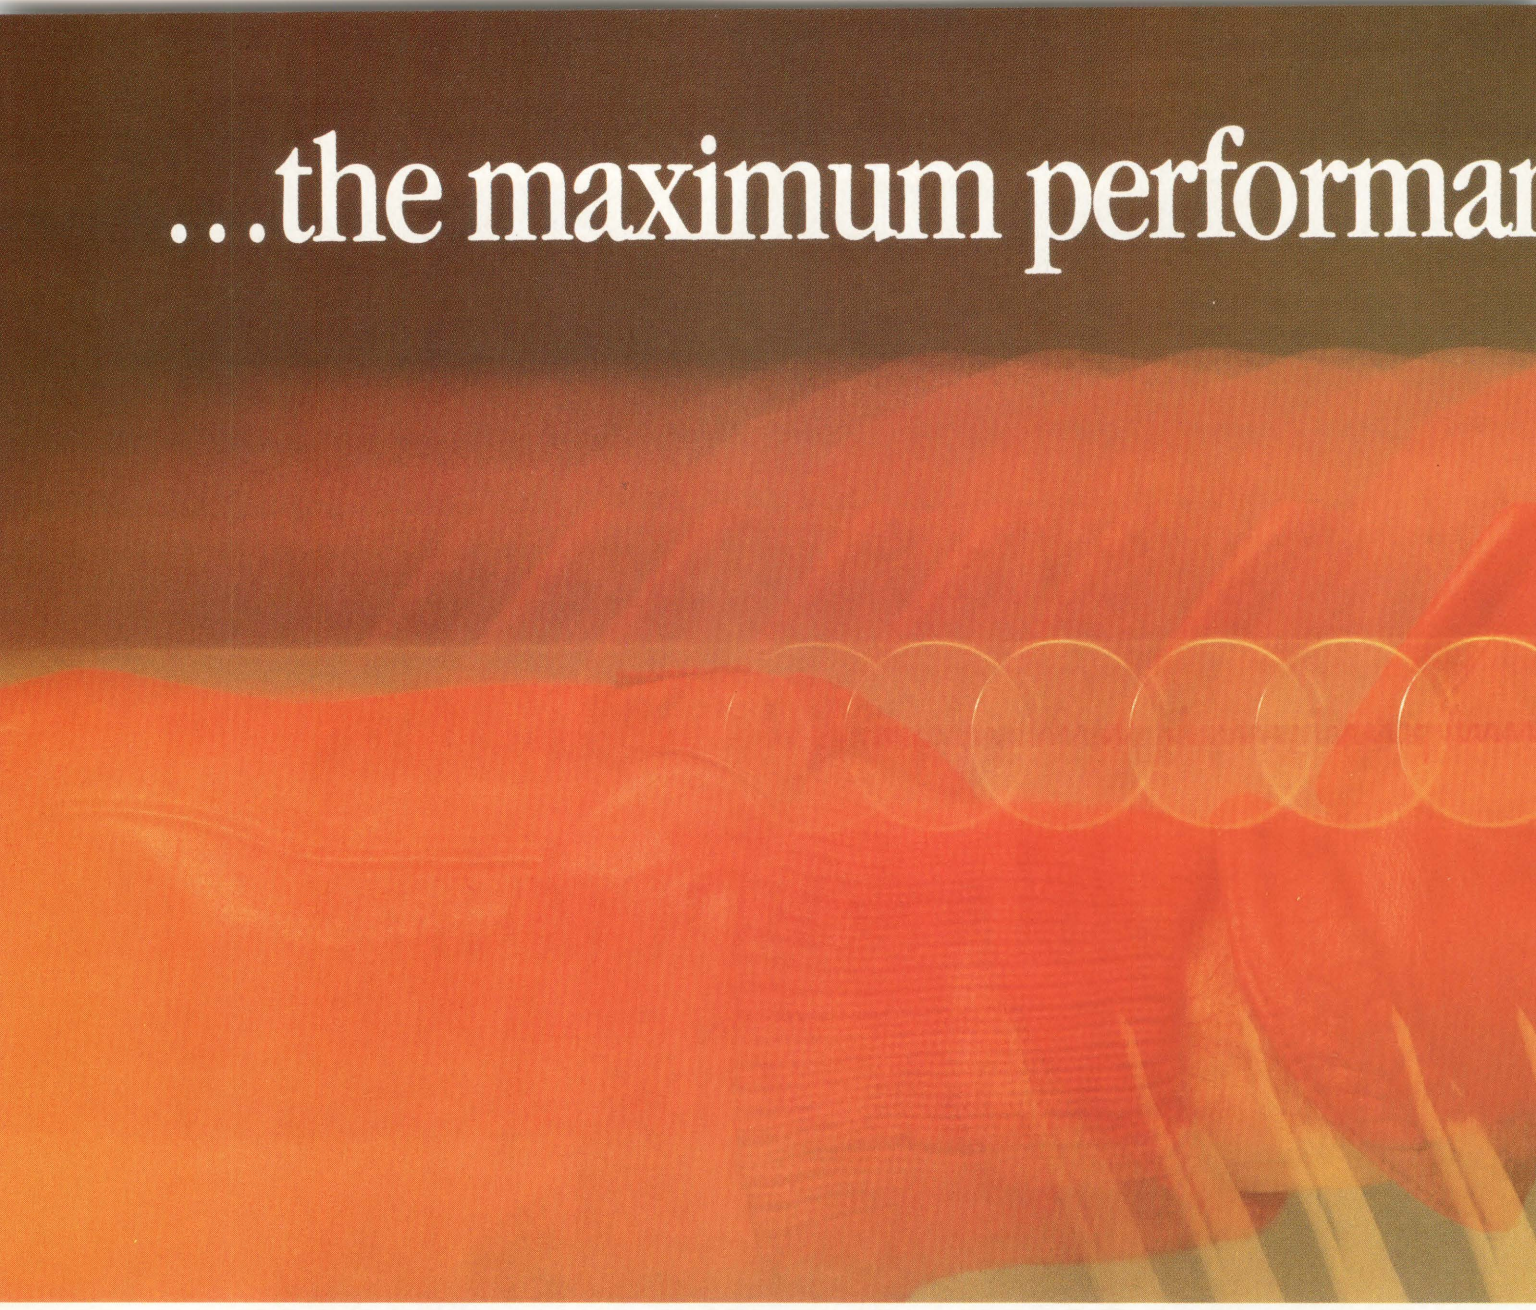

Plessey - Unsurpassed Process Technology

As system design becomes more and more challenging, and product life cycles become increasingly shorter, design flexibility and getting it right the first time have become critical factors in gaining and maintaining that maximum performance edge you've been looking for.

Plessey's investment in advanced process technology is unequaled in the industry. Successive reductions in feature size and continued improvement in process techniques are at the heart of leading-edge Plessey products.

Plessey - The Ultimate in ASIC Technology

Our broad range of ASIC products has grown to the point where we are now able to meet all the needs of ASIC users. We offer a full ASIC product range with a variety of options for digital, analog and mixed analog/digital applications, in gate arrays, standard cells, and full-custom. Advanced, state-of-the-art processes in fine geometry, high-density CMOS, bipolar and ECL technologies give you the highest levels of performance and system integration available today.

Plessey - Unparalleled CAD Support

The Plessey Design System (PDS) is a comprehensive suite of software embracing the design, simulation and implementation of gate arrays, standard cell and compiled ASICs in CMOS and bipolar technologies.

Customers who want to use their own CAD workstations or simulators are accommodated by flexible design interfaces at various stages into PDS.

Plessey - Standard Products And Discrete Components

Plessey's standard product family offers the highest performance product range available in the world today. Capabilities range from CMOS DSP devices operating in excess of 20MHz to the world's most advanced 1.3GHz monolithic log amplifier.

High performance solutions are also offered in radio communications, digital

nce that gives you the edge.

PLESSEY KEY PROCESS TECHNOLOGY

BIPOLAR

DESCRIPTION	Ft	EMITTER WIDTH	METAL LAYERS
Industry standard	400MHz	14μm	1
High voltage	400MHz	20μm	1
High speed linear	4.5GHz	4μm	2
High speed digital	6GHz	3μm	2
Ultra-high speed	14GHz	0.6μm	3

MOS

PROCESS FAMILY	fCLOCK	MINIMUM FEATURE	VSUPPLY
KC Industry standard CMOS	20MHz	4μm	3-10V
JG Double SiGate NMOS	10MHz	6μm	9-18V
VB High speed CMOS	40MHz	2μm	3-5V
VJ Very fast CMOS	50MHz	1.5μm	3-5V
VQ Ultra fast CMOS	75MHz	1.2μm	3-5V
MH/MA SiGate CMOS	30MHz	4μm	3-15V

BIPOLAR (CDI)

PROCESS	EMITTER WIDTH/ FEATURE SIZE	GRID PITCH	MAX. SPEED	MAX. POWER	MIN. POWER
ORIGINAL CDI	5μm				
CDI FAB I	3.75μm	11.5μm	10ns	2.4pJ	1.5pJ
CDI FAB IIa	2.5μm	8μm	4ns	1.2pJ	0.8pJ
Geometry change (utilizing multi-level differential logic-DML)					
CDI FAB IIb	2.5μm	8μm	800ps	0.8pJ	0.54pJ
CDI FAB III	1.5μm	6μm	400ps	0.4pJ	0.27pJ
CDI FAB IV	1.2μm	4.5μm	200ps	0.2pJ	0.14pJ

frequency synthesis, data conversion, telecommunications, data communications and consumer products.

Complementing the standard IC family, Plessey manufactures a complete line of discrete components including FETs, transistors and diodes available in SOT-23 and TO-92 packages.

Plessey - Over Two Decades Of Quality Commitment

For more than 20 years, Plessey Semiconductors has been commited to supplying the latest technology, highest quality, and highest performance semiconductor products in the industry. With our unique combination of CAD support, major advances in process technology, and the most advanced research facility in the world, Plessey Semiconductors is, today, a totally commited leader in the industry.

To learn more on how Plessey can help you achieve the maximum performance that gives you the edge, send for our new comprehensive, full color, 72-page short form brochure, or call Plessey Semiconductors today.

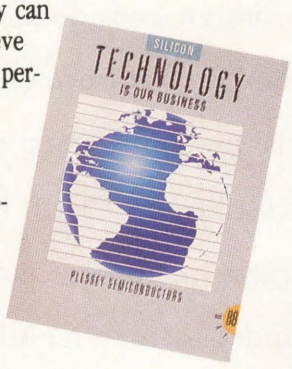

In North America call 1-800-441-5665.
Outside North America call 44-793-726666.

For further information you can write to us at one of the following addresses:

Plessey Semiconductors
1500 Green Hills Road
Scotts Valley, CA 95066
U.S.A.

Plessey Semiconductors Ltd.
Cheney Manor, Swindon
Wiltshire SN2 2QW
United Kingdom

PLESSEY

Number 1 for **ASICs**

CIRCLE NUMBER 9

SYNTHESIZING ADA'S IDEAL MACHINE MATE

DEBORAH W. RUNNER AND ERWIN H. WARSHAWSKY,
JRS RESEARCH LABORATORIES, INC., ORANGE, CALIF.

ADA programs mapped onto VHDL described machines

Embedded computer applications can demand performance that is beyond the current capabilities of standard microprocessor systems. The capabilities of a computer designed for these applications is measured by the execution speed of the actual problem set on the real machine. Design inadequacies are determined by the difference between actual performance and that optimum performance that can result from mapping software algorithms onto hypothetical hardware systems.

For such design considerations, the JRS IDAS tool set automatically derives relevant performance indices about hardware and software designs. It maps Ada programs onto machines described in VHDL, synthesizes machines described in VHDL from specifications expressed as Ada programs, retargets microcode compiler tools from a VHDL machine description, and guides users in rapidly evaluating numerous design alternatives and comparing the alternative approaches objectively and quantitatively.

Embedded computer applications such as signal processing, real time

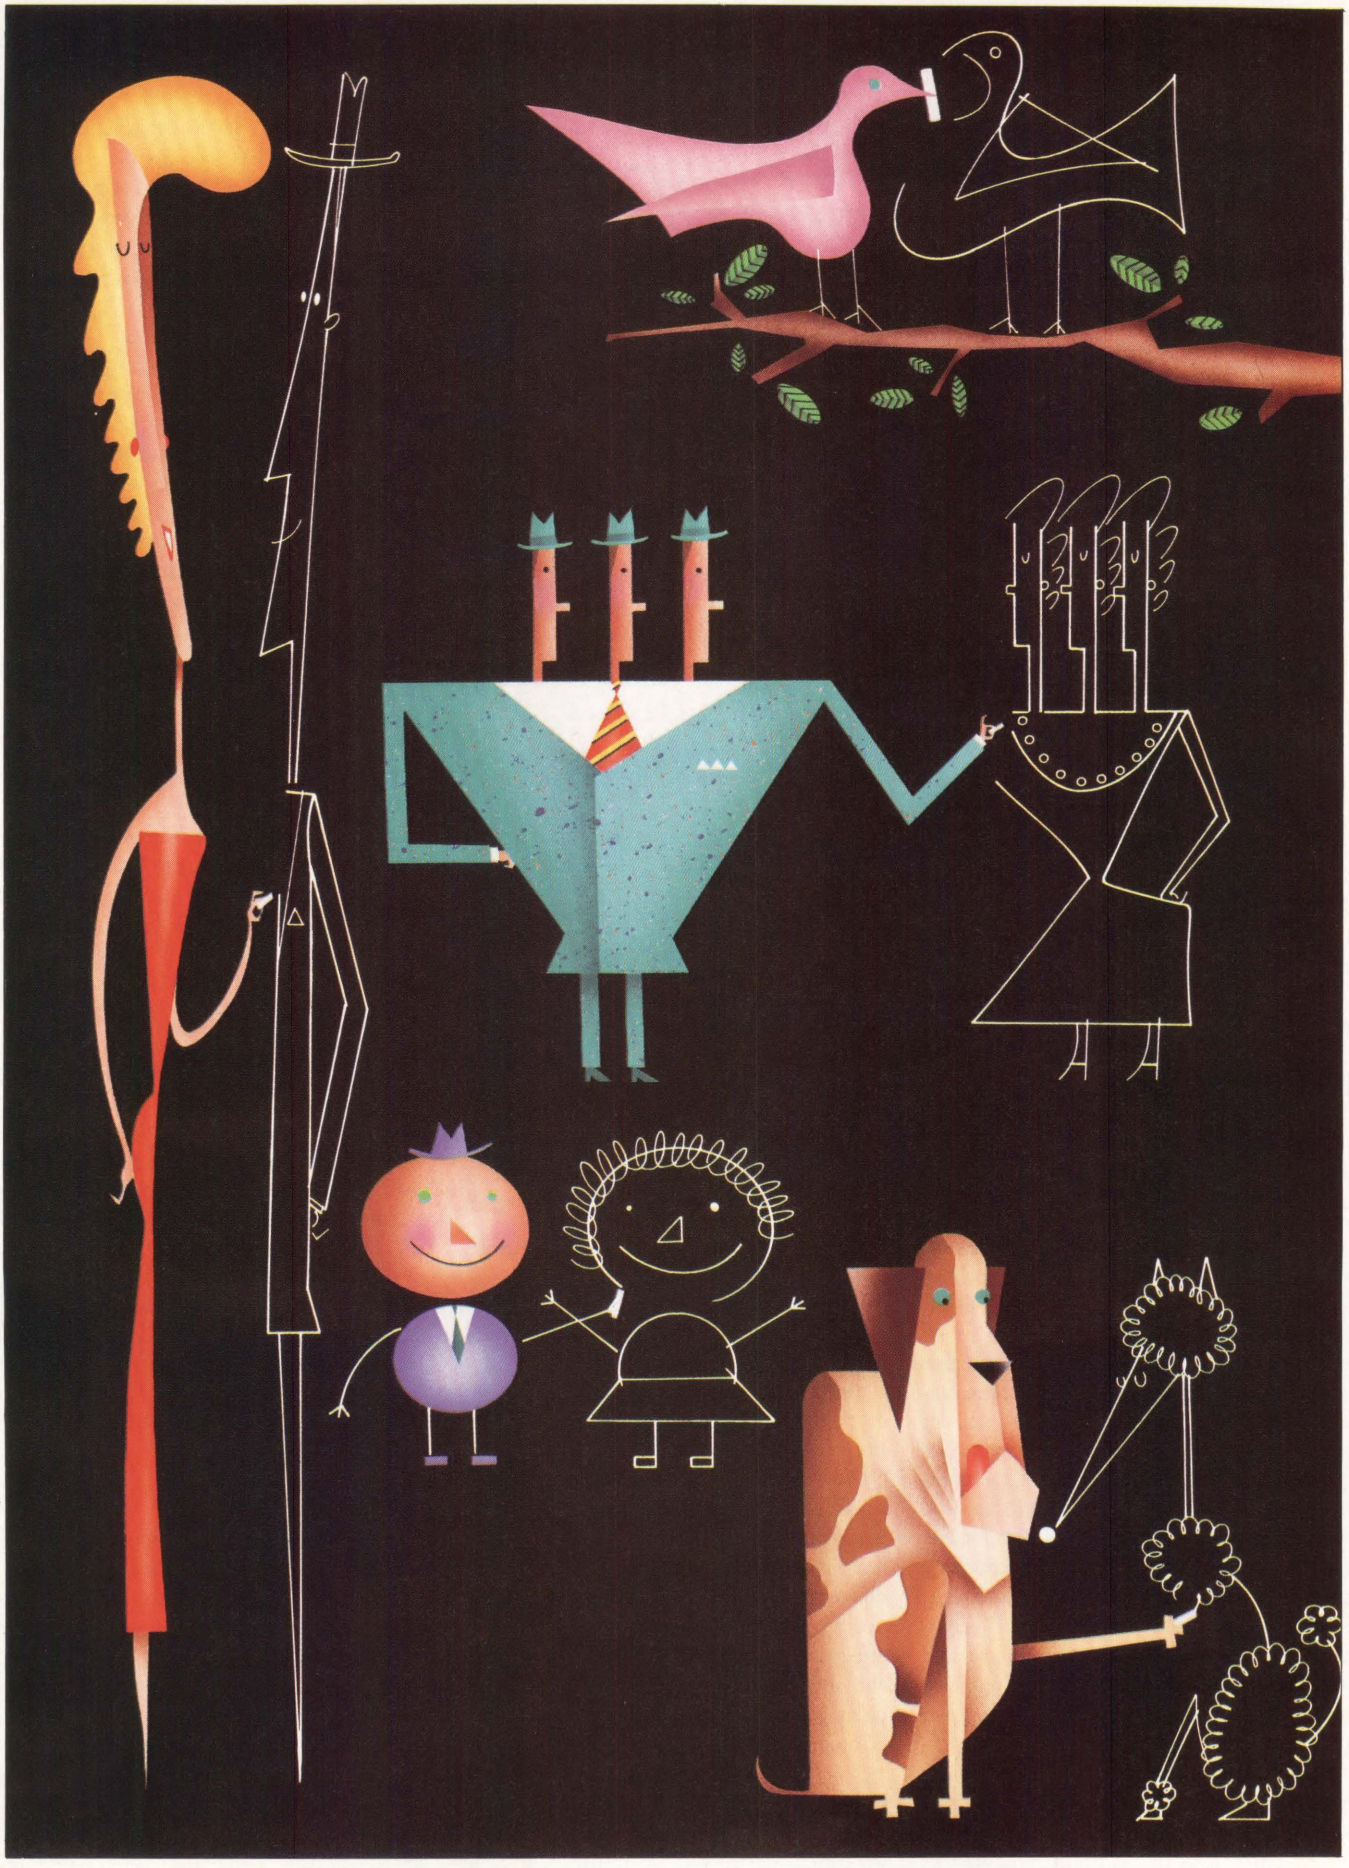

ILLUSTRATION BY TERRY ALLEN

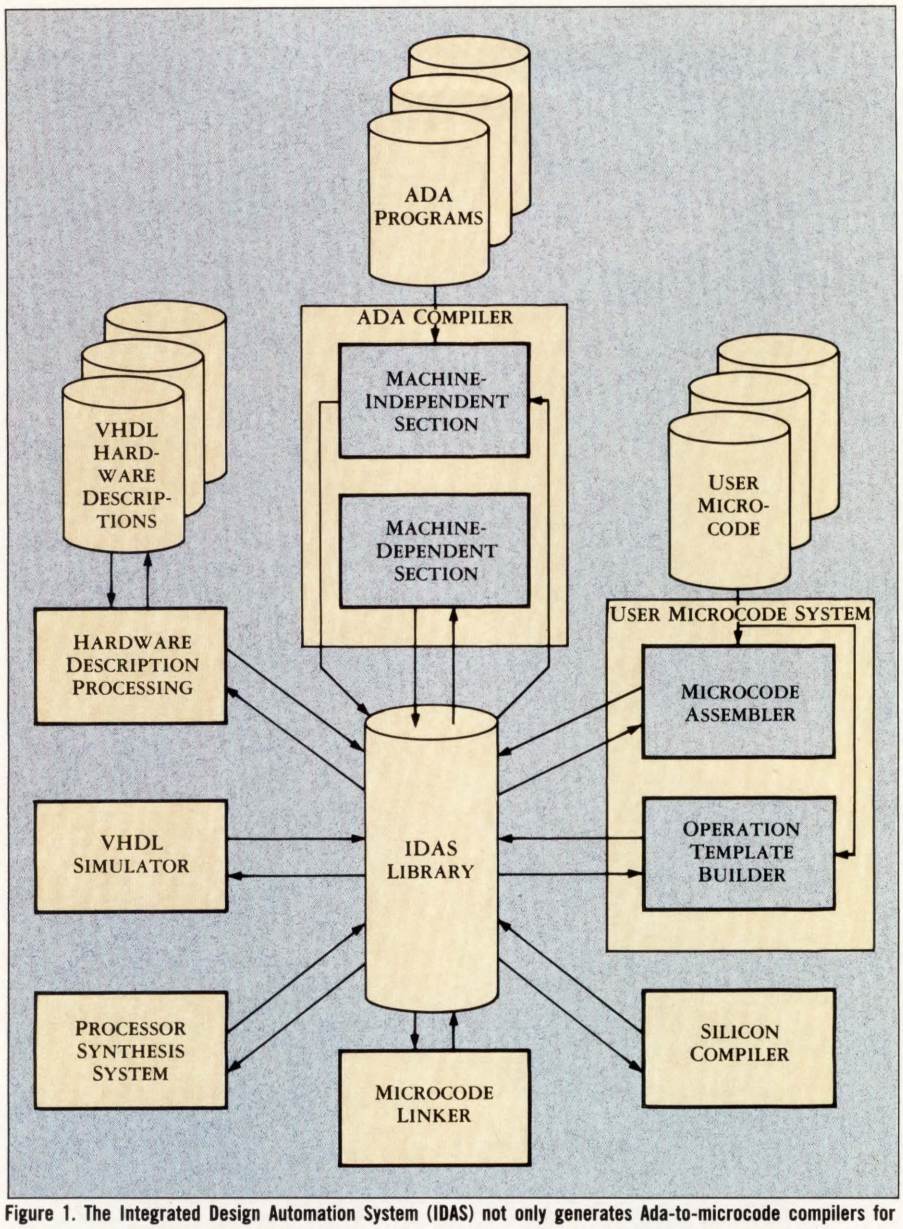

Figure 1. The Integrated Design Automation System (IDAS) not only generates Ada-to-microcode compilers for VHDL-described machines, it can also create a computer design (expressed in VHDL) for a given Ada program.

artificial intelligence, and image processing, have unlimited computing performance requirements—the more performance, the better the solution that is possible or the larger the problem that may be attacked. Furthermore, these embedded computer applications are usually severely constrained by other factors such as power, weight, size, schedule, reliability objectives and cost. The JRS Integrated Design Automation System (IDAS) focuses on solving the design and programming problems inherent in this type of system.

Embedded computer design and application problems are bound problems and are well-suited, therefore, for the application of design automation and design optimization techniques. They have known, specific requirements to solve one distinct problem. Computational requirements are known and may be explicitly characterized by specifications or program benchmarks or both. Non-computational requirements, such as power and cost, may also be known and explicitly specified. Valid solutions to the design problems must simultaneously satisfy all of the design requirements and constraints. Finally, everything is subject to change without notice.

The number of stringent criteria that must be simultaneously met make these high-performance embedded-computer applications extremely difficult to design effectively without tools; in fact, these designs are still a great challenge with an effective automation system. Fortunately, the specificity of the embedded applications provides focus and explicit success criteria to the design process, characteristics of great value in design automation.

Using that focus and conscious of the design criteria, JRS has integrated the major functional blocks in Figure 1 into the JRS IDAS embedded-computer design and programming environment. Machine (computer) and component descriptions expressed in VHDL (VHSIC hardware description language) are stored in the IDAS library, which is an object-oriented database. JRS IDAS has two major functions: generation of an Ada-to-microcode compilation environment for a machine specified in VHDL and creation of a machine expressed in VHDL that executes a user-supplied Ada program.

The Ada compiler block at the top of Figure 1 compiles Ada programs into microcode for a selected machine described in VHDL. User-written microcode can be merged into the automatically generated microcode and linked together for execution on a VHDL simulator for the target machine. The VHDL simulator is automatically generated, based on the VHDL machine description, and is then available for simulations of Ada applications.

IDAS has the capability of synthesizing a machine that is uniquely suited to a given application expressed in an Ada program. The machine is synthesized using library components within the IDAS database. The VHDL description of the synthesized machine may be used to automatically generate an Ada compilation environment and a VHDL simulator. It can also be transferred to a silicon compiler to produce an ASIC design compatible with a VHSIC-class semiconductor fabrication line.

IDAS has several other functions as well to support design and analysis of embedded computers. The user interacts with the synthesizer to manage software/hardware trade-offs, allowing the designer to create and analyze alternative architectures using actual applications code. Feedback from the system allows objective measurement of design trade-offs, removing some of the guesswork. IDAS allows designers to reuse existing code and machine descriptions in new applications, simplifying technology transfer and getting the most use out of existing designs. In short, it makes design more of a science and less of an art.

■ ADA COMPILATION

For input of Ada programs, IDAS provides an Ada compiler capability that generates microcode for an arbitrary horizontal processor; it translates application programs written in Ada directly into microcode that executes at the "microengine level" of a target machine (Figure 2). By so doing, it obtains substantial performance advantages over compilers that compile to target macro-languages or instruction set architectures (ISAs). The ability to use Ada as the application program-

TEXAS INSTRUMENTS REPORTS ON

SYSTEMS LOGIC

IN THE ERA OF

MegaChip TECHNOLOGIES

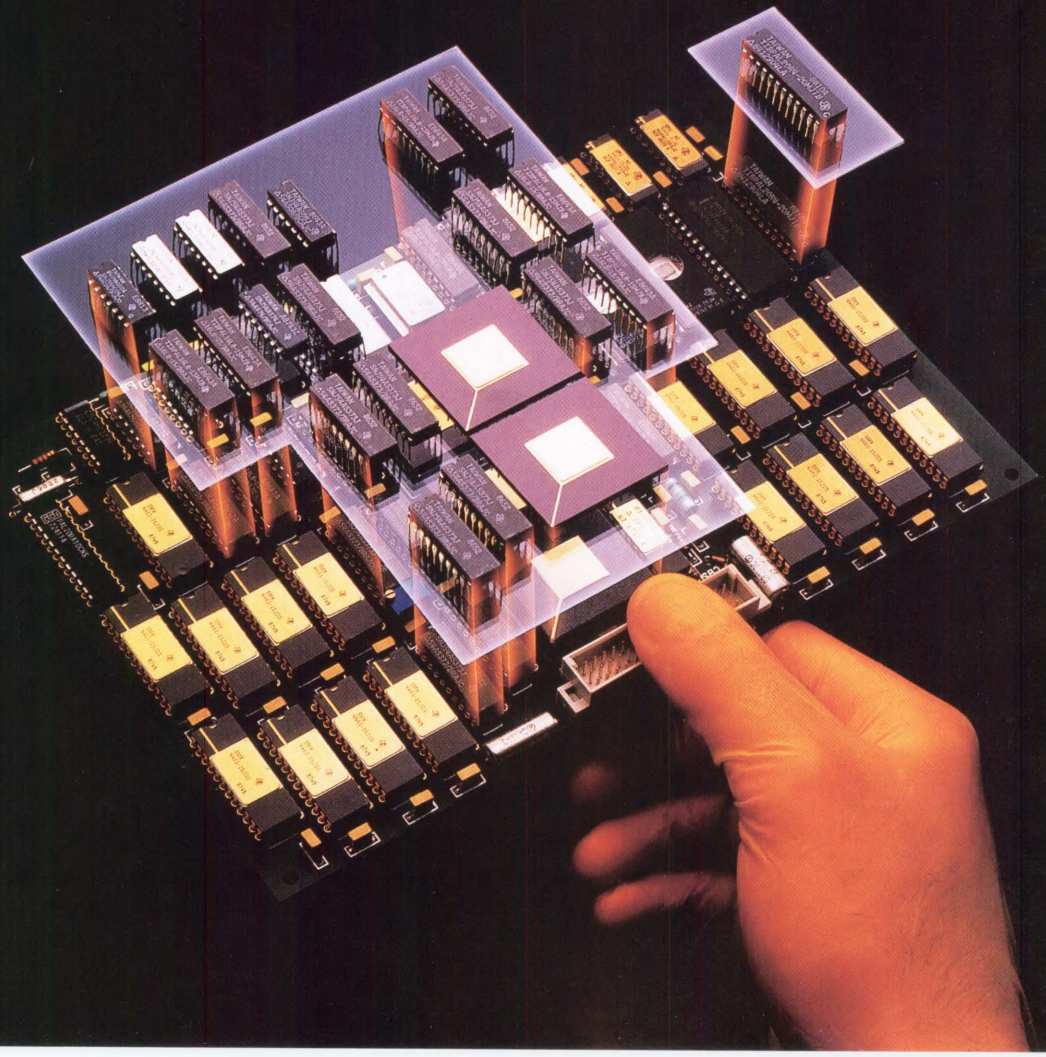

Systems logic in the Era of MegaChip Technologies:

No system should ever be limited by its to help your design perform at its best.

Up to 65% of the components in today's systems are logic. Such a large proportion demands that your logic devices perform on a par with other advanced building blocks—and be chosen with equal care. Systems logic alternatives from Texas Instruments can help you better realize the performance potential of your system design.

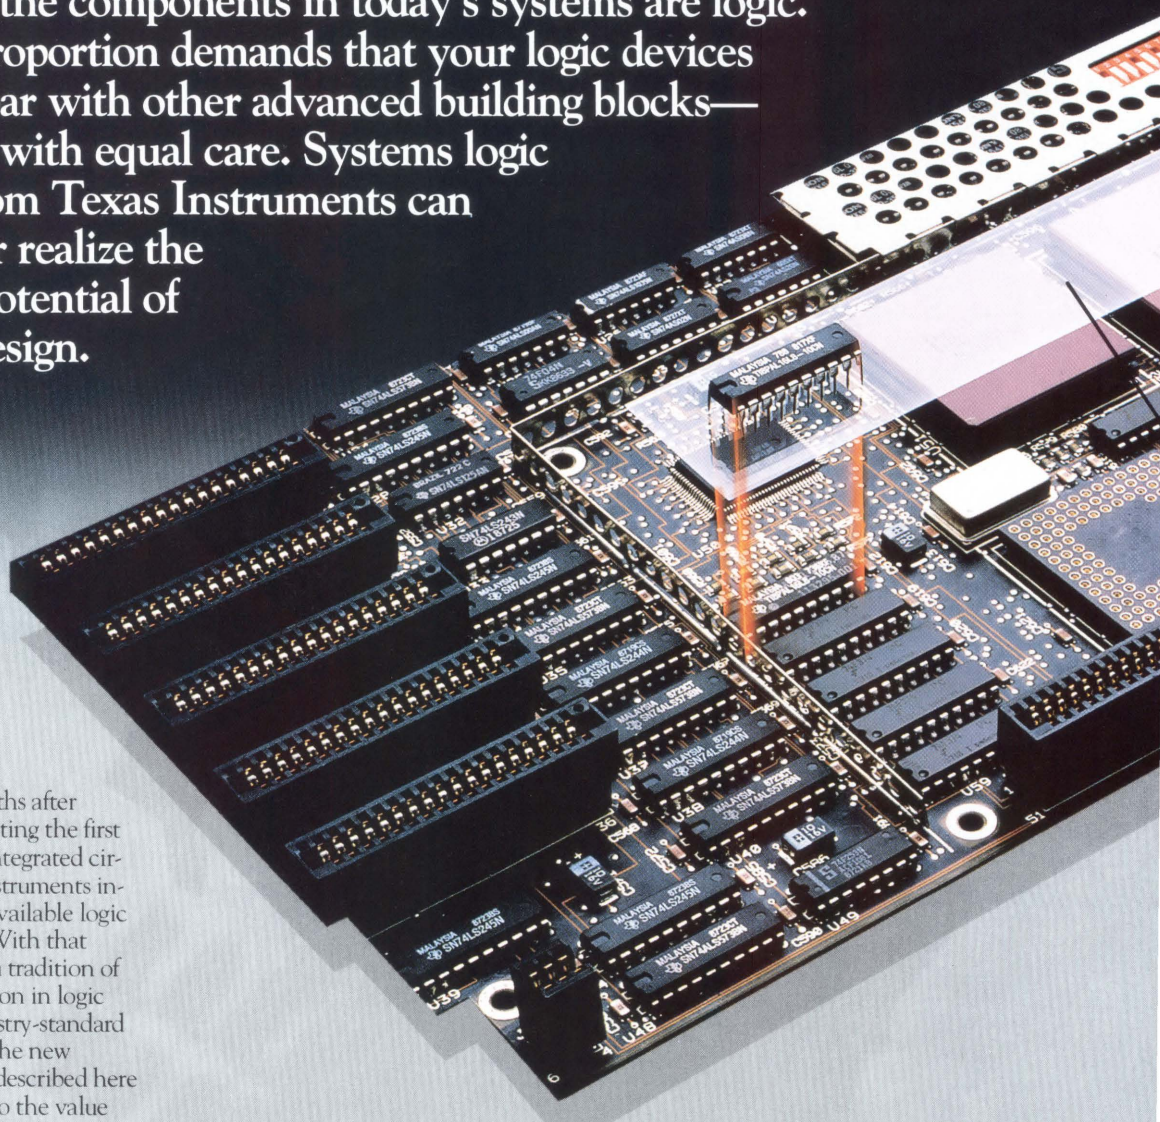

Within months after demonstrating the first working integrated circuit 30 years ago, Texas Instruments introduced a commercially available logic function, an RS flip-flop. With that beginning, TI established a tradition of development and innovation in logic that encompasses the industry-standard SN54/74 Series TTL and the new families of advanced logic described here that can add significantly to the value and performance of your overall system.

For example, for systems that require off-the-shelf flexibility with a degree of customization, TI's Programmable Logic Devices (PLDs) include popular 10-ns PAL® ICs available in high volume. And, to keep pace with today's high-speed microprocessors, TI plans to continue to drive PLD performance to sub-10-ns speeds.

TI's Advanced CMOS Logic (ACL) supports the design goal of high performance combined with low-power operation, while TI's new BiCMOS bus-interface family delivers very high drive current at very low power compared to bipolar circuits.

ON THE COVER: Suspended above the board, provided by Rockwell International, Missile Systems Division, are military versions of TI advanced logic devices.

TI's MegaChip Technologies

Our emphasis on high-density memories is the catalyst for ongoing advances in how we design, process, and manufacture semiconductors and in how we serve our customers. These are our MegaChip™ Technologies, and they are the means by which we can help you and your company get to market faster with better products.

For systems requiring moderate densities and fast prototype cycle times, TI offers a new series of one-micron CMOS gate arrays. When you need higher levels of integration plus increased design flexibility, TI's one-micron CMOS standard cells provide the means for system consolidation.

And for military applications, TI offers a wide choice of high-reliability logic functions.

On the following pages are details of what you can expect from TI's range of logic options:

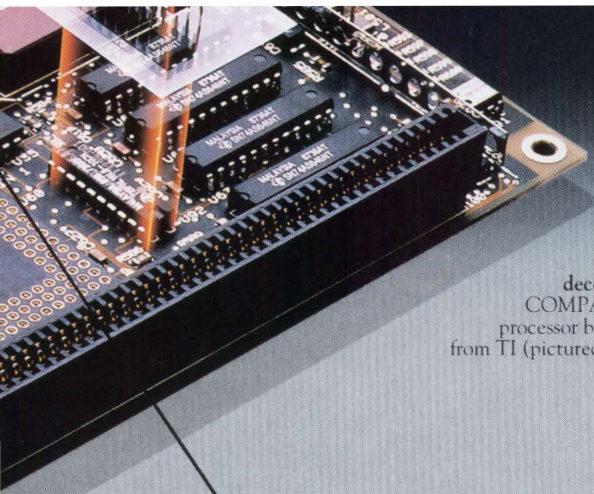

Contributing significantly to fast address decoding in speed-critical paths of the COMPAQ DESKPRO 386/20™ personal computer processor board are two TIBPAL16L8-10 PAL circuits from TI (pictured above a segment of the board).

Speed your system to market with TI's superfast PLDs.

PLDs are a functional alternative to standard logic ICs and gate arrays or standard cells.

Because TI's PLDs are off-the-shelf items you program yourself, you avoid the longer design cycle times of custom ICs and move on to market faster. These PLDs offer very attractive performance advantages. Consider these:

- **TIBPAL16XX-10 PAL ICs** from TI deliver a 10-ns propagation delay and are available in quantity. Clock-to-Q time is 8 ns, and output-registered toggle frequency is 62.5 MHz. IMPACT-X™ technology gives these PAL ICs their superior speed; they are well suited for use with high-speed processors such as the Motorola 68030, the Intel 80386, and RISC-based architectures. The 10-ns performance brings a higher level of integration to speed-critical paths.
- **TI's TIEPAL10H16P8-6 IMPACT™ ECL PAL circuit** delivers even faster operation: 6-ns propagation delay max. You can now streamline conventional ECL designs by consolidating several discrete components into a single custom function.
- **TI's new 7-ns Programmable Address Decoder** is intended to help you squeeze more performance out of memory interface systems. By performing address decoding much faster than conventional PAL architectures—in 7 ns—the TIBPAD16N8-7 allows you to take advantage of the new processors to increase overall system performance.
- **TI's 50-MHz Programmable State Machines (PSMs),** TIB825S105B (16 x 48 x 8) and '167B (14 x 48 x 6), are ideal for use in high-performance computing, memory interface, telecommunications, and graphics. These PSMs may be used to implement custom sequential logic designs such as peripheral I/O controllers and video-blanking controllers.
- **The TIBPAL22VP10-20,** with a 20-ns delay, is 20% faster than the competition's "A" version and much more flexible. A programmable output macrocell allows two extra, exclusive output configurations, for a total of six.
- **TI's TICPAL16XX Series 20-pin CMOS PAL ICs** are the cure for power problems. They operate at virtually zero standby power and are reliable, high-performance replacements for conventional TTL and HCMOS logic. The devices can be erased and reprogrammed repeatedly.

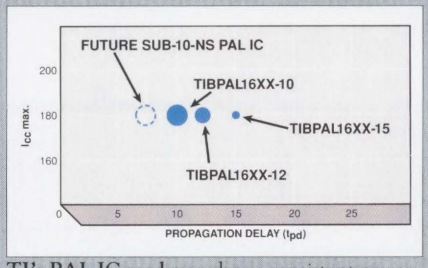

TI's PAL IC road map shows consistent power and consistently higher speeds, with even faster versions on the way.

Turn page for more information

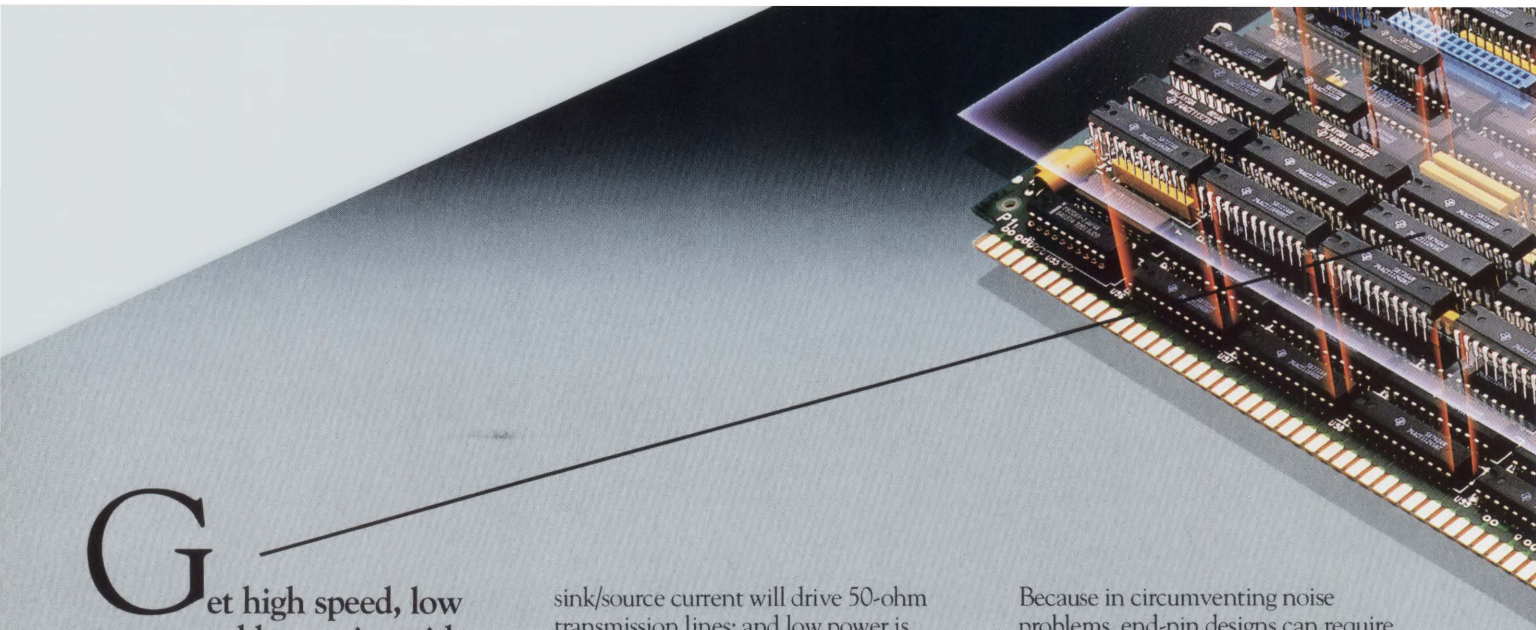

Get high speed, low power, and low noise with TI's broad ACL family.

It's an extensive family that includes gates, flip-flops, latches, registers, drivers, and transceivers. It's a readily available family in DIP and SOIC packages. It's TI's high-performance EPIC™ ACL family, bringing with it an important bonus—major reductions in noise.

Family speed is comparable to advanced bipolar 54/74F; 24 mA of sink/source current will drive 50-ohm transmission lines; and low power is characteristic of TI's EPIC technology. All this with "ground bounce" substantially reduced compared with end-pin ACL. The reasons are innovative packaging and a circuit-design technique called OEC™ (Output Edge Control) which softens the transition states that cause simultaneous switching noise. In fact, EPIC ACL noise levels are typically 10% less than those of bipolar devices.

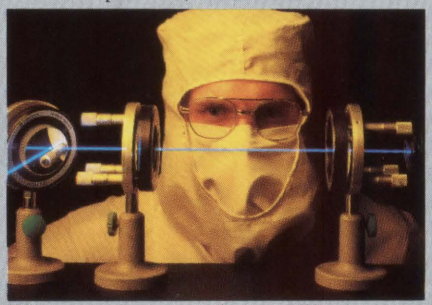

When every nanosecond counts, TI's new high-performance ACL family can help you significantly improve system speed.

The rapidly increasing customer acceptance of TI's ACL family confirms its noise-reduction advantages and its ease of use.

System design advantages

A unique "flow-through architecture" simplifies board design, layout, and troubleshooting. Inputs surround power pins on one side, outputs on the other, and control pins are strategically located at the package ends.

From a systems perspective, TI's arrangement offers the lowest-cost design when compared to end-pin ACL. Because in circumventing noise problems, end-pin designs can require additional components that take up to 32% more board area and slow system performance.

There are 146 functions, in both AC and ACT versions, currently announced in TI's ACL family, including such innovative, highly complex functions as advanced transceivers, line drivers, latches, feedback registers, multiplexers, and counters.

This ACL family, developed in cooperation with and supported by Philips/Signetics, fully meets JEDEC industry-standard No. 20 specifications for Advanced CMOS Logic.

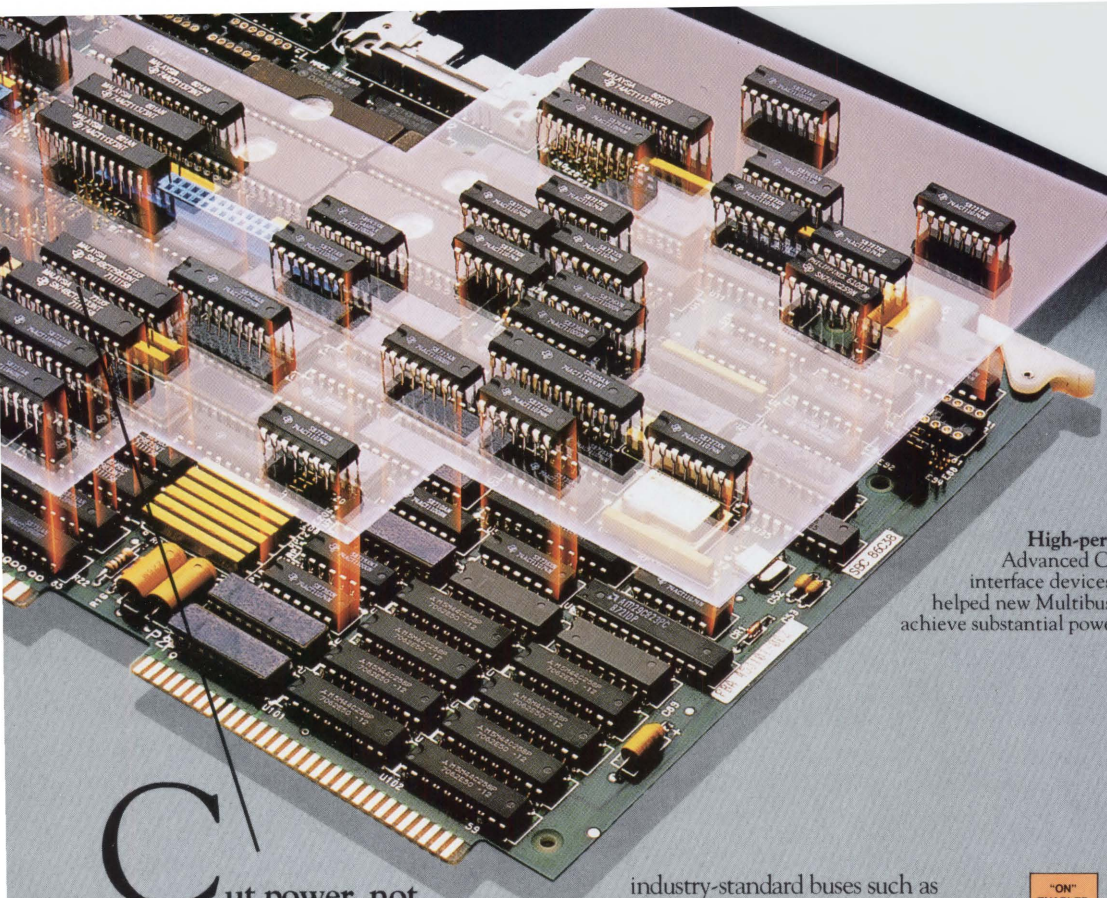

High-performance, low-power EPIC Advanced CMOS Logic and BiCMOS bus-interface devices (suspended above the board) helped new Multibus single-board computer achieve substantial power reductions.

Cut power, not speed or drive, with TI's BiCMOS bus-interface ICs.

This new family is a simple, effective means to reduce system power consumption without compromising advanced performance.

As the BiCMOS name implies, TI combines bipolar IMPACT and CMOS processing to achieve switching speeds comparable to advanced bipolar products and provide the 48/64-mA drive current needed for high-capacitive loads and backplanes. In particular, family members meet the drive requirements of industry-standard buses such as Multibus® and VMEbus™

In addition, TI's BiCMOS devices can reduce disabled currents by 95% and active currents by 50%-80% compared to bipolar equivalents. Result: System IC power savings can be more than 25%.

There are more than 60 functions comprising TI's BiCMOS bus-interface family. Included are 8-, 9-, and 10-bit latches, buffers, drivers, and transceivers—a wide choice that means you can easily find what you need to implement high-performance bus-interface designs.

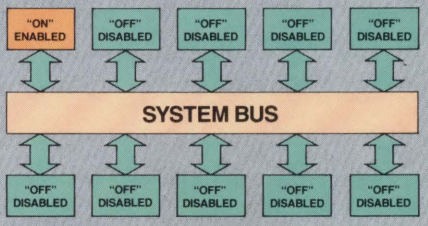

An innovative circuit design in TI's BiCMOS bus-interface logic helps lower disabled currents. This is key to overall power savings because in a typical bus network only one device is enabled at a time.

Turn page for more information

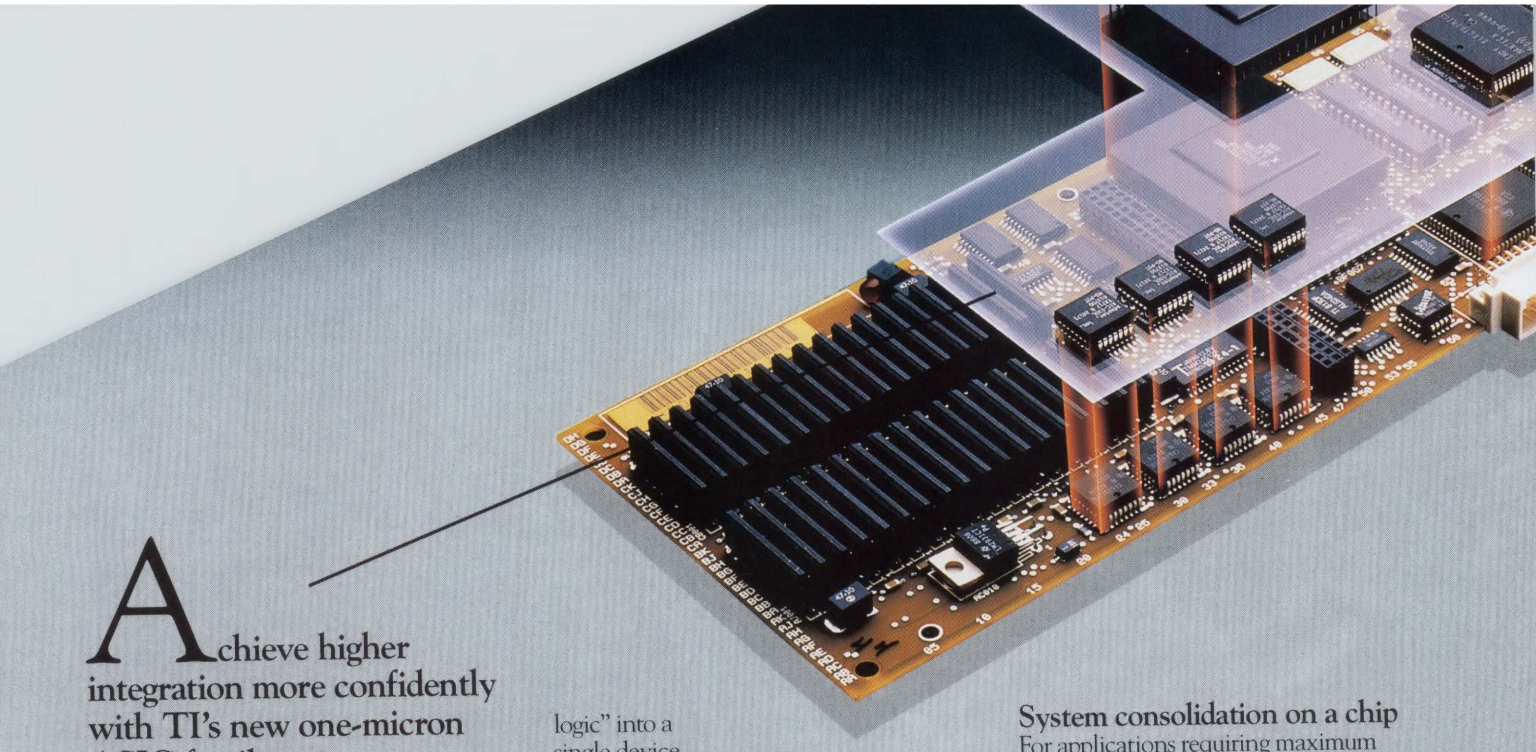

Achieve higher integration more confidently with TI's new one-micron ASIC family.

Now, you can integrate more of your systems logic using TI's new one-micron CMOS ASIC (application-specific integrated circuit) family—the TGC100 Series gate arrays and the TSC500 Series standard cells. Each offers different degrees of design flexibility and system integration. The result is significantly reduced component count which cuts board size and system cost while improving reliability and performance.

And TI is supporting the family with comprehensive kits that help minimize design cost, risk, and time by providing a comfortable, easy-to-use design environment.

Efficient logic consolidation

Using TI's new TGC100 Series gate arrays, you can sweep major chunks of "glue logic" into a single device while realizing fast design and prototype cycle times. Array densities currently range to more than 8K usable gates and 142 bond pads; the Series will be extended to more than 16K usable gates and 216 bond pads in a major production release planned for late 1988. Prototype delivery is typically two to three weeks from approval of postlayout simulation results.

The TGC100 Series Design Kit gives you complete autonomy and control over the design process. It is a comprehensive set of the tools required for successful gate-array design and validation *(see last page for details)*.

Standard packages for the TGC100 Series range from 28-pin DIPs to 84-pin PLCCs, with optional packages up to 144 pins.

System consolidation on a chip

For applications requiring maximum design flexibility and higher levels of integration, TI has disclosed its third-generation standard-cell family, the TSC500 Series.

Complex system designs can be implemented using a growing core of basic SSI/MSI functions, as well as scan cells for testability and MegaModule™ building blocks such as register files. FIFOs, bit-slice family functions, RAM, and ROM are other aids to implementation. Output cells with drive capability up to 64 mA are available.

Package options include conventional through-hole DIPs, surface-mount PLCCs, and plastic quad flatpacks (QFPs) in both JEDEC and EIAJ standards, as well as high-pin-count plastic pin-grid arrays.

Both the TGC100 and TSC500 Series have a typical propagation gate delay of

logic. TI offers advanced logic families

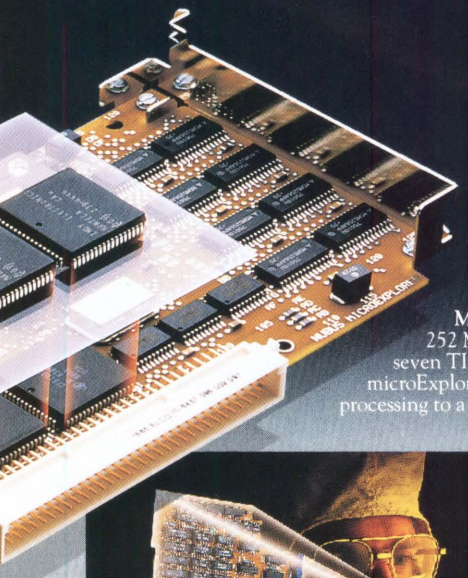

Major logic consolidation, the equivalent of 252 MSI and LSI devices, was possible using the seven TI ASIC functions shown above the microExplorer™ board which brings symbolic processing to a MacIntosh II® desktop computer.

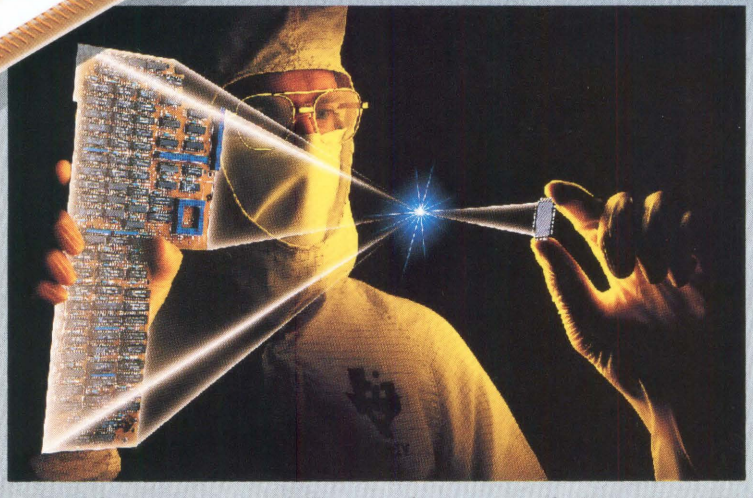

480 ps for a two-input NAND gate with a fanout of three; flip-flop toggle rates range up to 208 MHz. Both series offer output and bidirectional buffers with variable slew-rate control. And both series are fabricated in TI's high-performance EPIC process.

Apply TI's advanced logic to improve the performance of military systems.

Among TI's broad selection of logic devices produced to military requirements is a large PAL family. Propagation delays as fast as 15 ns are available over the military temperature range. The introduction of a 12-ns, 20-pin PAL circuit is planned, as well as military versions of the TIB825S105B and '167B Programmable State Machines.

TI is offering military counterparts selected from its ACL family, as well as 54F functions. Soon to come will be the BiCMOS family of bus-interface functions.

Included among TI's lineup of military ASICs are versions of the one-micron TGC100 Series gate arrays discussed at left, as well as two-micron standard cells.

TI's logic devices are among the more than 800 military functions offered compliant to MIL-STD-883C, Class B. Of this total, TI provides more than 200 to DESC-standard military drawings and is qualified to supply 285 JM38510 Class B devices (QPL 75).

THIRTIETH ANNIVERSARY
1958 1988
INVENTION OF INTEGRATED CIRCUIT

Milestones in Innovation

TI's tradition for milestone innovations extends from the infancy of semiconductor technology into the MegaChip Era. Among the major highlights:

- First commercial silicon transistor (1954)
- First commercially produced transistor radio (1954)
- First integrated circuit (1958)
- First integrated-circuit computer (1961)
- First-hand-held calculator (1967)
- First single-chip microprocessor (1970)
- First single-chip microcomputer (1970)
- First single-chip speech synthesizer (1978)
- First advanced single-chip digital signal processor (1982)
- First video RAM (1984)
- First fully integrated trench memory cell (1985)
- First gallium arsenide (GaAs) LSI on silicon substrate (1986)
- First single-chip Artificial Intelligence microprocessor (1987)

Turn page for more information.

Comprehensive support from TI helps you improve your design performance as you improve system performance.

To enable you to excel in designing the logic portion of your system for maximum performance, TI has compiled or is making available a wide range of design tools and aids:

PLDs: The TI PLD data book (472 pages) contains design and specification data for 78 device types. Four application notes are incorporated as a reference tool. A qualification book is available, and a state-machine design kit is forthcoming.

ACL and BiCMOS Bus Interface: TI's ACL data book (348 pages) contains detailed specifications and applications information on the members of the one-micron ACL family. The ACL designer's handbook (299 pages) spells out the technical issues confronting advanced-logic design engineers and describes methods for handling the issues. A qualification book (358 pages) features extensive reliability and characterization data, die photos, and application derating factors. Customer evaluation capability is enhanced by TI's system evaluation board (available for demonstration through TI field sales offices) and third-party characterization boards.

Data sheets are available on each member of TI's BiCMOS bus-interface family.

ASICs: The TGC100 Series Design Kit gives you the tools needed to successfully complete a gate-array design: A macro library for Daisy or Mentor engineering workstations containing the graphic symbol and functional and simulation models for each macro; a software library of TI-specific software tools that streamline and simplify the design process; a design manual that answers "how to" questions about designing with the TGC100 Series; a two-volume data manual providing detailed specifications for each macro in the TGC100 Series software library; and a software user's manual.

An equally comprehensive design kit for the TSC500 Series is currently in development.

Extensive design support available for TI's systems logic families includes that for the new TGC100 Series gate arrays *(at top)*, Programmable Logic Devices *(at left)*, and Advanced CMOS Logic.

For more information on TI's advanced systems logic ICs and their support tools, complete and return the coupon today. Or write:
Texas Instruments Incorporated
P.O. Box 809066
Dallas, Texas 75380-9066

Texas Instruments Incorporated
P.O. Box 809066
Dallas , Texas 75380-9066

SDV083VD800C

Yes, please send me the following:
RY01 ____ ASIC Information Packet
DZ01 ____ Programmable Logic Device Data Book
CA01 ____ ACL/BiCMOS Information Packet
CB01 ____ BiCMOS Data Sheet Packet

NAME

TITLE

COMPANY

ADDRESS

CITY STATE ZIP

AREA CODE TELEPHONE EXT.

TEXAS INSTRUMENTS

ming language for embedded high-performance computers is a substantial capability in its own right and it is a cornerstone of the system.

The Ada to Microcode Compilation function consists of numerous powerful elements: an Ada compiler with its lexical, syntactical, and semantic analysis and library facilities; an Ada Intermediate Language (AIL) for software representation; a machine-independent optimizer, linker, and simulator; a machine-dependent optimizer, translator, code generator, microcode optimizers and compactors, resource allocator, assembler, linker, and hand microcode insertion facility; and several miscellaneous translators and formatters.

The JRS IDAS uses VHDL as the means for inputting and outputting descriptions of hardware. IDAS invokes the VHDL system developed for the VHSIC Program Office and utilizes its facilities, in particular its simulator generator capability. The VHDL system can be invoked independently of IDAS as well. Currently, IDAS works with the IEEE-1076-standard VHDL system developed by Intermetrics Inc., but it isn't limited to that system.

With VHDL, a user may enter a description of a machine into the system and IDAS will automatically create a simulator for the machine and automatically retarget the Ada compiler system to that machine. In a matter of hours, the system will automatically generate a complete test and evaluation environment for an arbitrary horizontally microprogrammed machine. This is a very powerful prototyping capability. The automatic retargeting capability is the most significant feature of the system because it enables the comprehensive evaluation of alternative designs and the performance of design trade-offs in very short periods of time.

IDAS translates VHDL hardware descriptions into its internal database format and generates VHDL descriptions of completed designs. Documents and program analyzers are provided to aid the user in correctly representing hardware in VHDL for entry into IDAS. The system deals with two levels of hardware description—components and machines. Machines are collections of properly interconnected components and are always digital and synchronous. Components may be existing physical devices, such as the TRW, Texas Instruments, or Honeywell VHSIC Phase I chip sets, or the AMD 29300 chip set; additionally, components may be primitives of a silicon compiler, such as Seattle Silicon's Concorde System. IDAS can synthesize machines constructed with the Concorde primitives and pass the designs to the silicon compiler for chip implementation. The silicon compiler creates an output suitable for implementation of a VHSIC semiconductor production line. The interface between the IDAS and the Concorde uses VHDL, and includes format translators as needed.

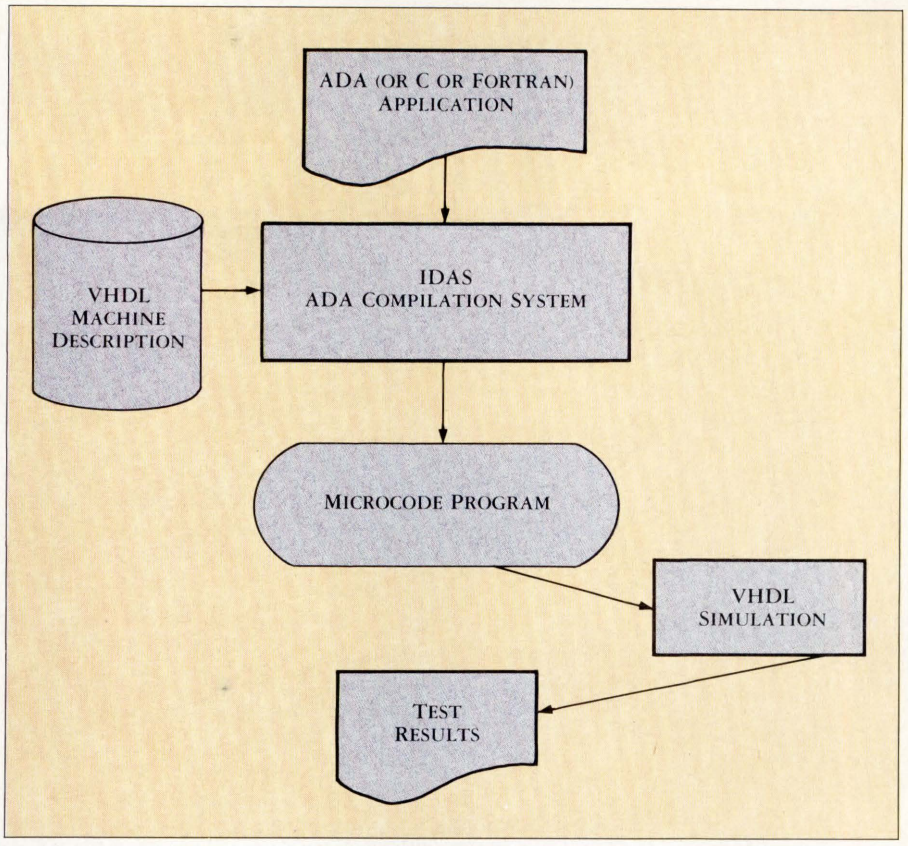

Figure 2. Given a machine description in VHDL, IDAS can create both an Ada-to-microcode compiler and a hardware simulator for evaluating the execution of Ada programs on that machine.

The VHDL -driven automatic retargeting capability is accomplished by a large number of complex programs that perform elaborate analyses and transformations on the input hardware description data. Models of various types, including behavioral, structural, and physical models are built and maintained internally to support the requirements of the code generation, simulation, and synthesis features of the system. Analyzers compute all possible machine paths and generate useful micro-operation sequences. All machine-dependent elements of the Ada compiler and the Ada-to-VHDL synthesis subsystems are retargeted from a common hardware data source.

The use of Ada for embedded, high-performance applications has always been a problem because of the mismatch between the richness and power of the Ada language and the need to make embedded computers simple, to optimize power consumption, weight, size, and reliability. The IDAS system addresses this problem by providing optimizers and translators that tailor the code produced by the Ada compiler to match the behavioral features of the target machines.

The system's tools perform such things as translating real number operations to equivalent integer expansions for machines without floating-point hardware. It can also translate selected pointer operations to the specific functionality available in the address generators of a particular machine. The automatic retargeting of an Ada run-time environment is also included.

The system aims to provide full Ada functional support for any target machine wherever possible so as to minimize reprogramming requirements and maximize software reusability.

■ MACHINE SYNTHESIS

Perhaps the most unique feature of IDAS is the ability to automatically synthesize machines that are "optimal" for the execution of a specified set of Ada programs (Figure 3). It can view a set of Ada programs as specifications for a machine and then synthesize a machine that executes the programs in as small a number of cycles as possible.

The system provides an interactive user interface that allows the user to direct the design process and to impose constraints on parameters of interest. The designer can, for example, limit the number of particular hardware devices in a machine or establish limits on the amount of power

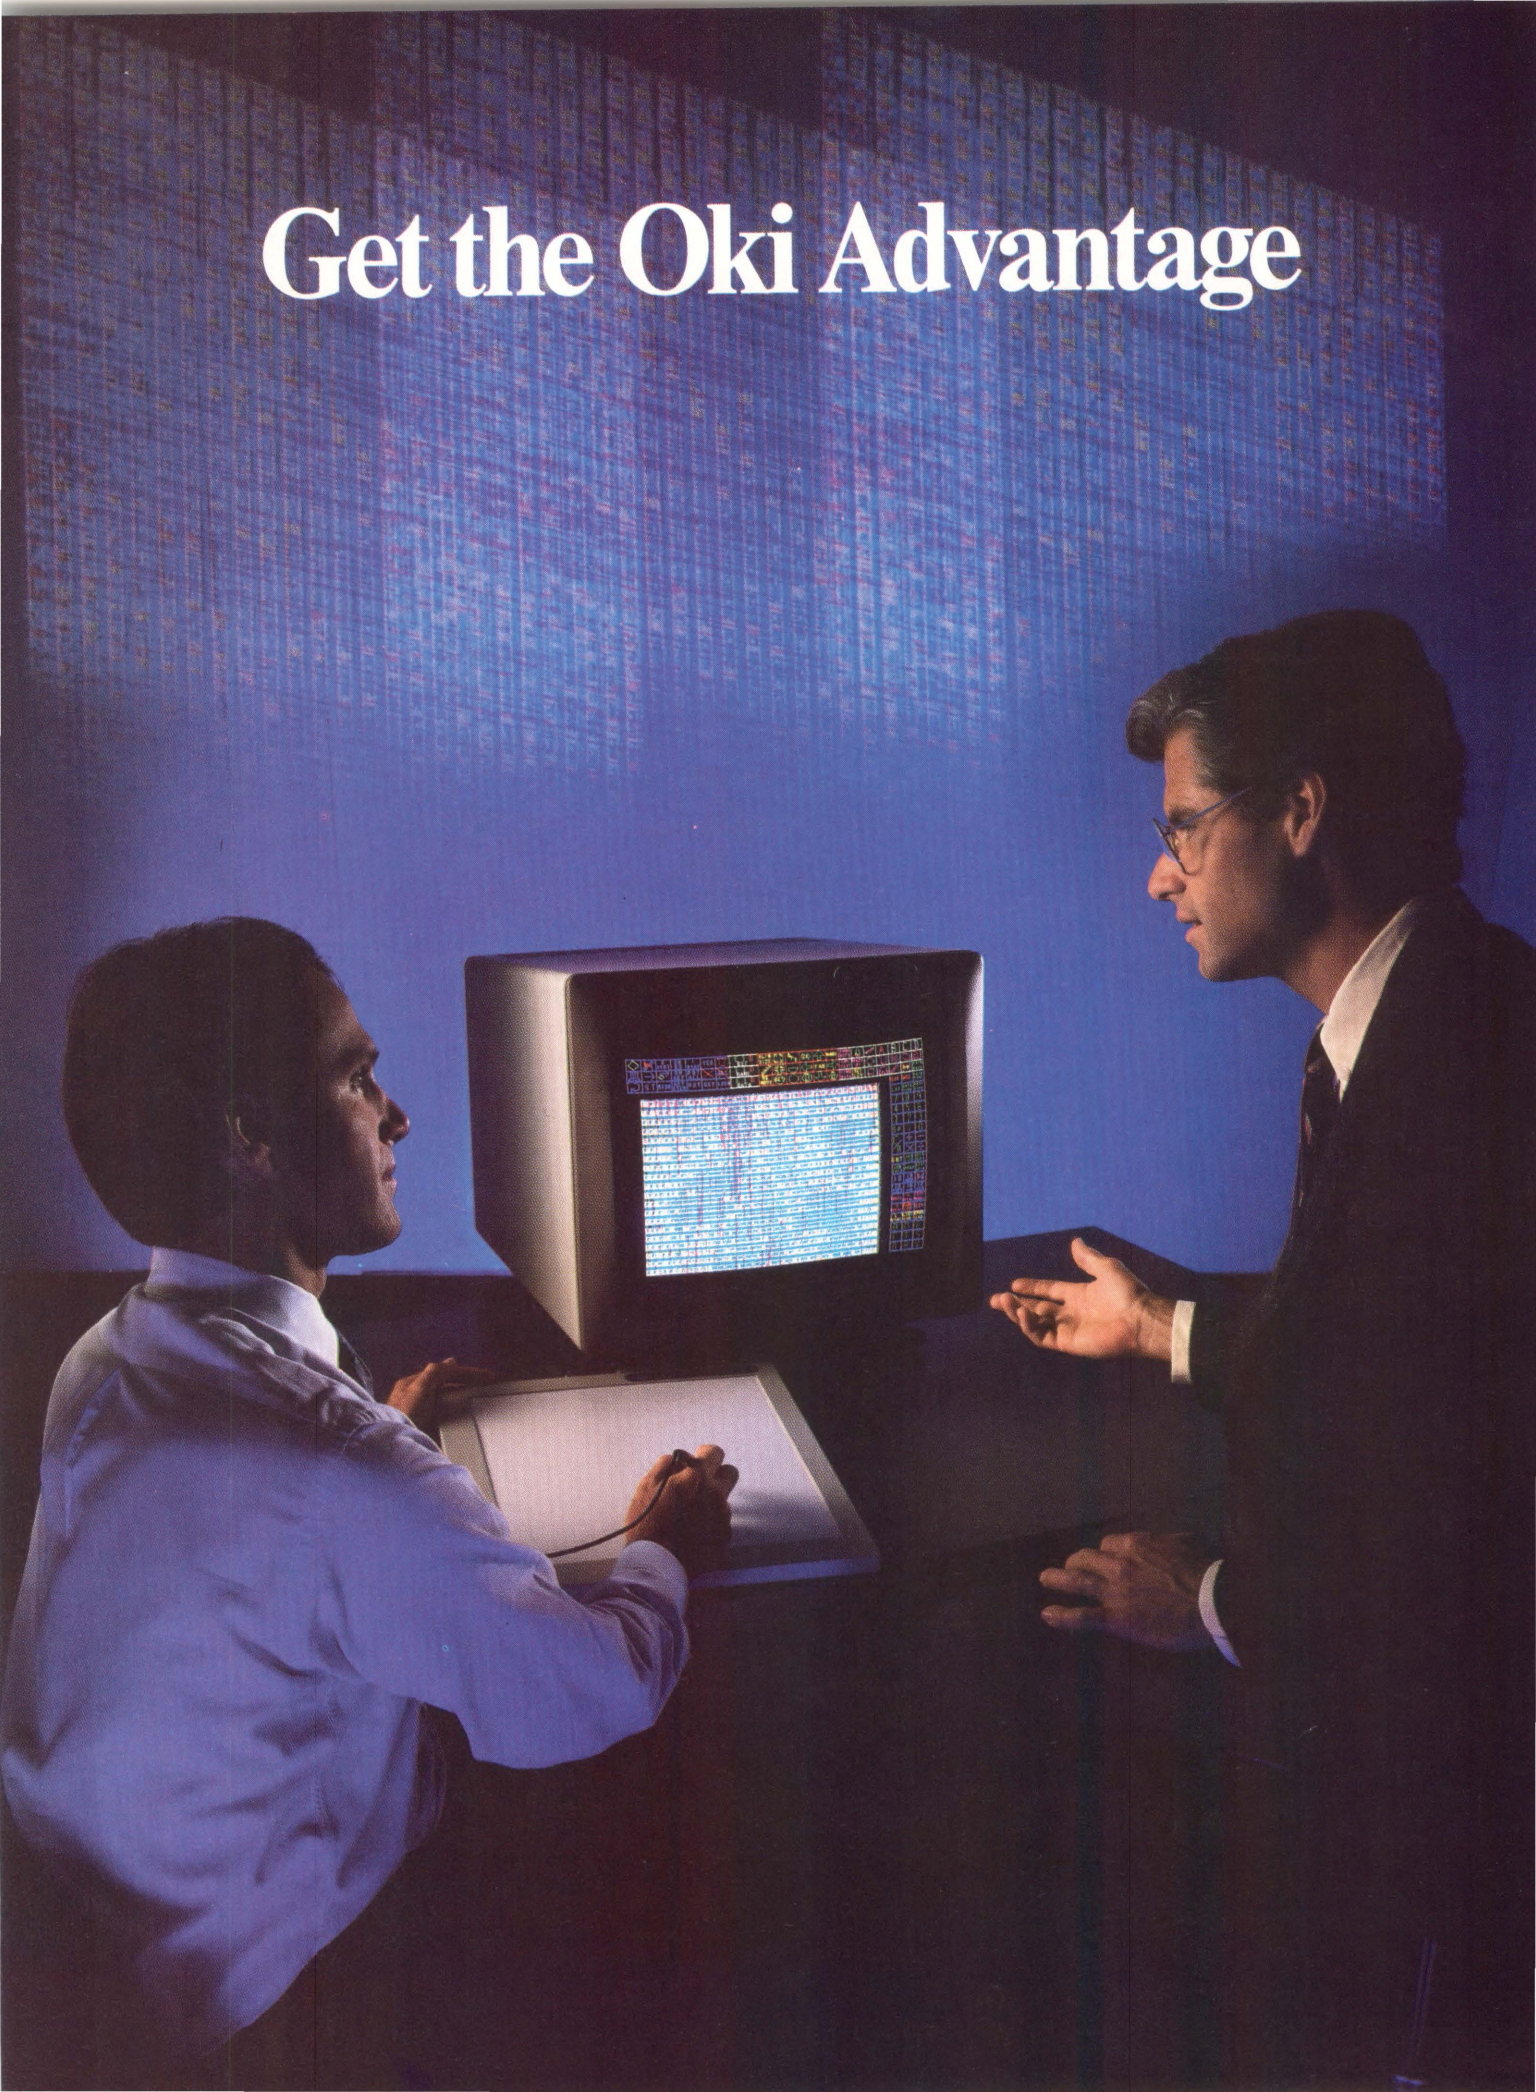
Get the Oki Advantage

Everything you need for ASIC success from one reliable source

An ASIC project is a major commitment of your budget and man hours. Give yourself the advantage of working with a powerful partner. Oki Semiconductor has the experience, resources, and commitment you can rely on to help ensure your ASIC VLSI success.

Advanced ASIC products and technologies

Oki Semiconductor has complete ASIC capabilities, from full custom to semicustom ICs. Our three families of advanced CMOS ASIC products have been designed to meet all of today's high-density, high-speed device requirements.

▶ *Sea-of-gates:*
new sea-of-gates channelless arrays provide an available 100,000 gates and a minimum 40,000 gate circuit logic density of 640 picoseconds.

▶ *Channelled array:*
new 1.2 μ channelled arrays provide speed in the subnanosecond range together with a logic density of up to 30,000 usable gates.

▶ *Standard cell:*
the new 1.2 μ standard cell family offers density up to 60,000 gates and an average speed of 600 picoseconds, plus memory capability of 32K bits RAM and 128K bits ROM.

ATG and logic transparency

With automatic test-point generation built into each of these new products, test programs can be generated in a fraction of the time you'd normally spend—without sacrificing logic or speed. All three product families also use the same cell design library, for logic transparency.

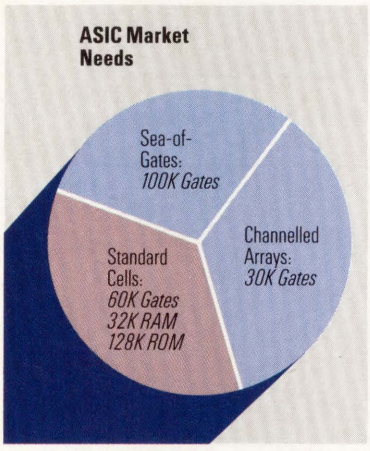

Complete ASIC support

Working with Oki means you can draw on our vast resources and experience to back you up at any stage of the development process. We have one of the finest ASIC teams in the industry to support you. We also provide the state-of-the-art design tools, packaging options, and manufacturing capabilities to successfully implement your project. With so much on the line, give yourself the security of working with Oki—the one source you can rely on for all your ASIC needs.

Check Oki: your complete ASIC resource.

Please send complete technical data/specs on Oki capabilities in:
☐ Sea-of-gates ☐ Channelled arrays
☐ Standard cells ☐ Full customs
☐ Please call. We have immediate requirements.

( )
Phone

Name

Title

Company

Attach coupon to business card or letterhead and return to: ASIC Customer Service, Oki Semiconductor, 785 North Mary Avenue, Sunnyvale, CA 94086. Phone: (408) 720-1900.

OKI
SEMICONDUCTOR

CIRCLE NUMBER 10

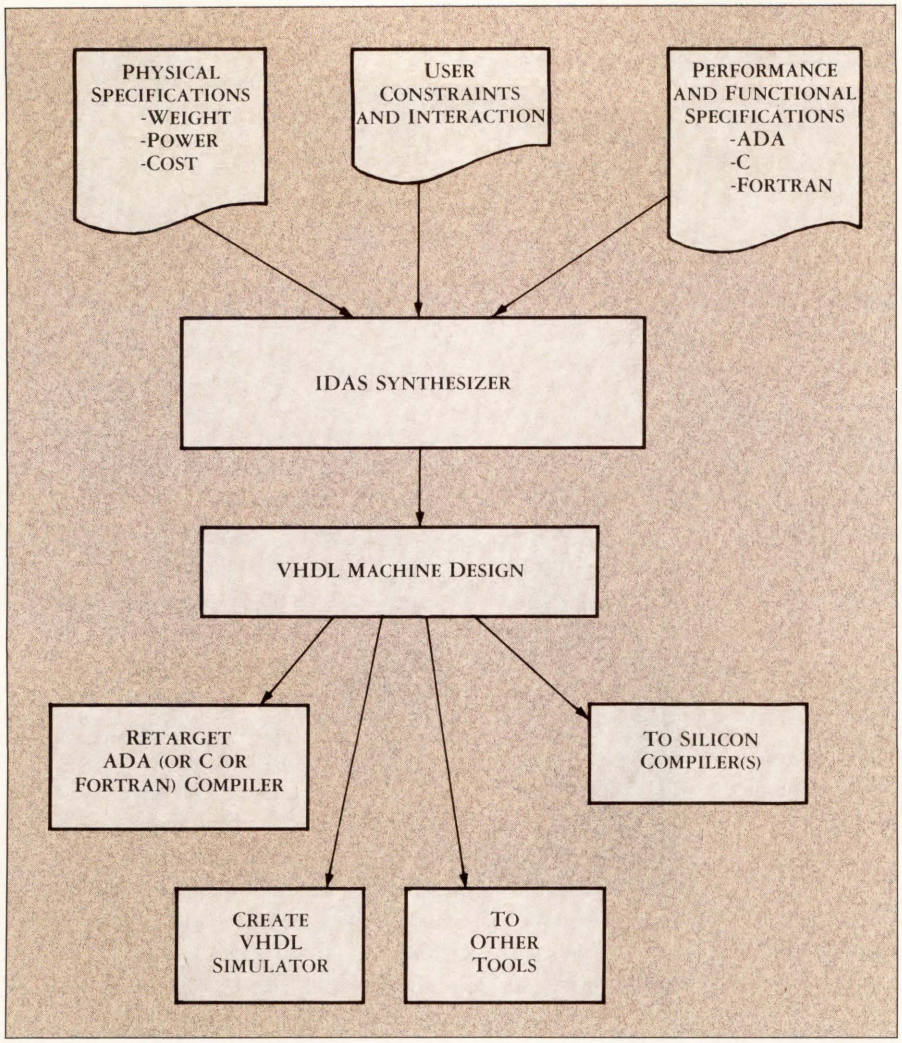

Figure 3. Given an Ada program IDAS can synthesize a machine, expressed in VHDL, that can run that code. After iterating to find the best balance of performance, physical specifications and user constraints, the system generates a VHDL description, a VHDL simulator and a software compiler for the machine.

or weight required by the machine. A user may direct IDAS to implement a new machine with any selected existing component set (e.g., TRW VHSIC or AMD 29300) or with compiler primitives for generation of ASIC devices. The user may make multiple selections and the system will provide a comparative analysis of the results of the choices.

The primary output of the synthesis process is the design of a synchronous, digital, programmable machine. The design is expressed in both the internal system form and in VHDL. This design may then be used to retarget the Ada to microcode compiler and to generate a simulator.

The synthesis facility consists of a large number of complex software elements that perform the optimizations and analyses in conjunction with user interactions. In addition to the capabilities that are utilized in the Ada and VHDL subsystems, the synthesizer contains software tools for manipulating "idealized" logic design components to create an "idealized machine form." The user uses these tools to vary architectural rules and constrain the number of components used for a design. The tools immediately provide performance results and resource utilization statistics. The idealized machine form is extremely valuable to the synthesizer because it is the intermediate machine representation between the software specifications and the hardware component set.

The synthesizer automatically creates an architecture model based on JRS-supplied architectural rule sets together with constraints entered by the designer. The user can manipulate the rule set as well as the constraints to see the effect on benchmark performance and resource utilization. This procedure allows the user to evaluate hardware and software design trade-offs with respect to cost, size, performance, and power dissipation.

As an example of the synthesis process, let us assume that an embedded computer is needed to quickly execute matrix multiplication. The designer will write an Ada benchmark to perform the task or select an existing benchmark from the IDAS library. The Ada code for the multiplication inner loop is shown below:

```
for k in a'range(2) loop
   sum:= sum + a(i,k) * b(k,j);
end loop;
```

The designer will then select a component set from the library with which to build the machine. For purposes of this example, the designer selects the silicon compiler primitives. At a click of the mouse, the synthesizer builds an ideal machine to perform the matrix multiply loop. This machine executes the equivalent compacted microinstructions shown in Figure 4a.

This initial machine has been synthesized using as many silicon compiler components as are necessary to execute the matrix multiply. The resources have not yet been constrained so the synthesizer has selected as many as it needed to give optimum compaction and best performance. As a result, the inner loop of the matrix multiply on this machine architecture takes 1,000 cycles.

The synthesizer reports resource usage to help the designer make better trade-off decisions. This first pass results in the following unconstrained resource usage:

One sequencer is used in 40 percent of the execution cycles;
Six literal fields with percent usage of 1=93 percent, 2=86 percent, 3=40 percent, 4=7 percent, 5=7 percent, and 6<1percent;
Two memories with percent usage of 1=85 percent and 2=78 percent;
Three address generators with percent usage of 1=86 percent, 2=86 percent, and 3<1 percent;
Two register-ALUs with percent usage of 1=80 percent and 2=72 percent.

The designer may then place constraints on the resource usage, invoke the synthesizer to redesign the machine, and examine the effect on the execution speed of the algorithm as well as the code size. This evaluation of software/hardware trade-offs can continue until the designer is satisfied with the constraints placed upon the hardware as well as the benchmark execution. A final synthesized machine might consist of the following:

One sequencer used in 23 percent of the execution cycles;
Two literal fields with percent usage of 1=81 percent and 2=54 percent;
Two memories with percent usage of 1=49 percent and 2=45 percent;
Two address generators with percent usage of 1=69 percent and 2=31 percent;
One register-ALU with percent usage of 1=88 percent.

On this smaller, constrained machine,

PERFORMANCE

16-bit MICROPROCESSOR

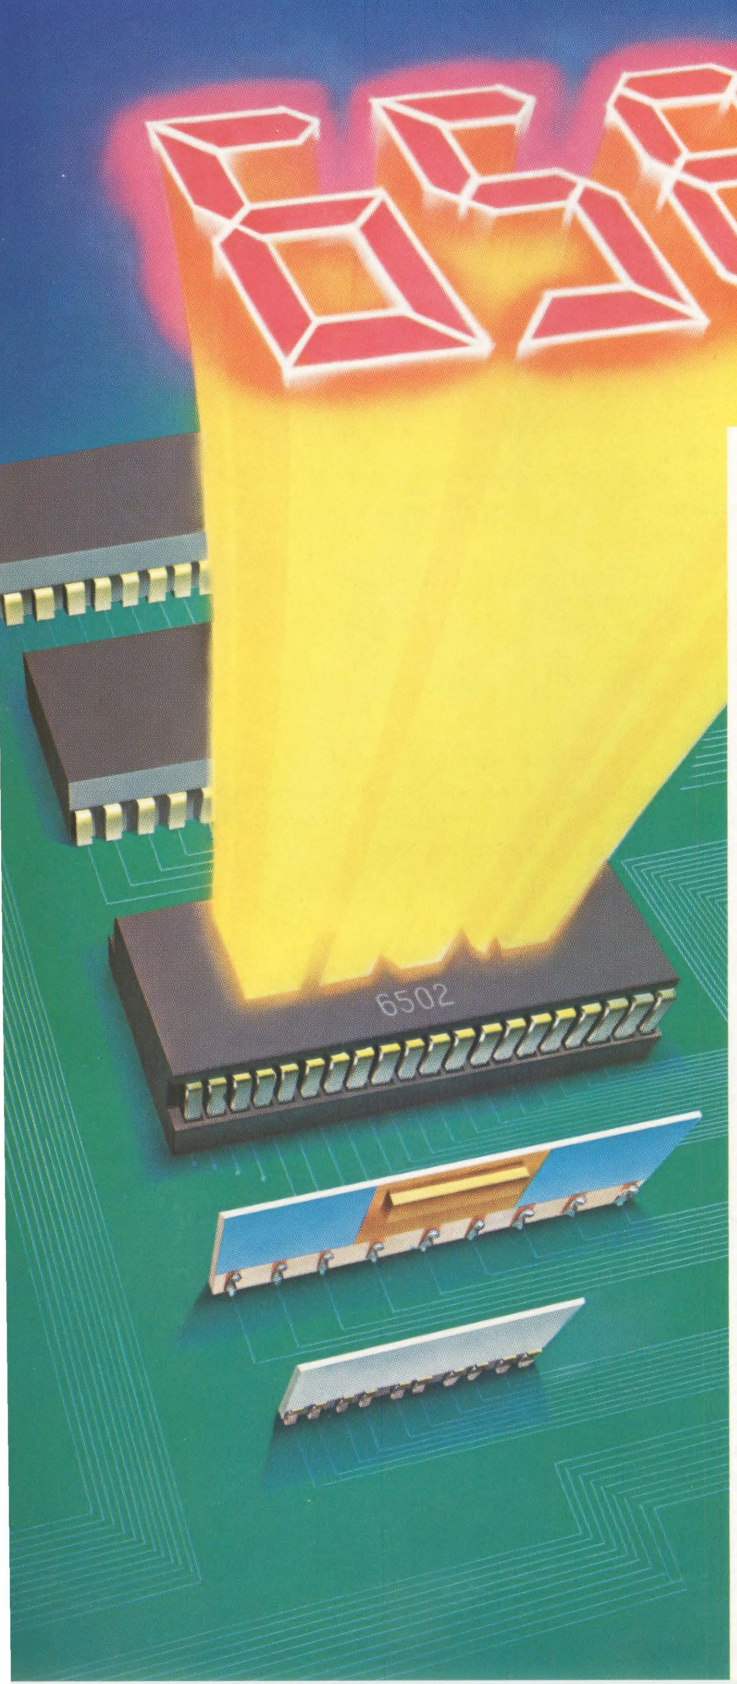

16-bit Performance Boost for 6502 Designs

Upgrade with Ease. With the 16-bit G65SC816 you can design embedded control using what you already know. It is fully software compatible with a performance boost of 16 megabyte addressing, 24 addressing modes, 91 instructions and 255 op codes. All built on a familiar base–easy to use without compromising functionality.

Flexibility. Both G65SCO2 and G65SC816 code can be run by switching from Emulation mode to Native mode through software control. Coprocessors supported through both software and signal pins.

Performance. High performance CMOS for low power consumption, high noise immunity and high speeds.

Compatibility. The world's most popular 8-bit microprocessor family of peripherals is completely compatible and available for immediate delivery. And, if pin-for-pin and software compatibility are key, the G65SC802, with internal 16-bit architecture, is ready for plug-in upgrading.

Performance Products From a Performance Company. We're solid and fully resourced, including microprocessor families, telecom devices and thin film resistor networks. We offer services in wafer fabrication, ASIC design and packaging technologies. You're invited to see our performance first hand.

Call Steve McGrady, Marketing Manager at (602) 921-6526.

California Micro Devices Corp.
Microcircuits Division
2000 West 14th Street • Tempe, AZ 85281
(602) 921-6000 • FAX (602) 921-6298 • TLX 187202

CIRCLE NUMBER 11

VHDL For Synthesis

The VHSIC hardware description language (VHDL) and system was developed under the sponsorship of the Department of Defense to satisfy several compelling needs, particularly those related to modern VLSI technology: documentation of complex VLSI chips, design descriptions at all levels of digital system design which could be simulated and evaluated, multilevel mixed-mode simulation capability, and a common design/description language. The language and system that attempts to satisfy those needs, VHDL, is now IEEE standard 1076.

JRS utilizes, and incorporates into its IDAS, the IEEE -1076 version of the VHDL system provided by Intermetrics Inc. VHDL is used as the language for describing all digital hardware for both input to and output from IDAS.

VHDL is a comprehensive language for describing all levels of digital hardware. Using VHDL, one could provide, for example, documents ranging from architectural diagrams to interconnect lists. Clearly, the information content of each type of document can vary widely. A VHDL architectural description of a machine would bear little resemblance to a VHDL detailed logic description.

VHDL, like English, is a language for expressing things at many levels of detail and complexity. Human or automated readers of VHDL, like those of English, must have an understanding of the level of expression for useful communication to occur.

The JRS IDAS requires a VHDL machine description that provides the information content necessary to retarget its Ada compilation tools, for simulation at a microarchitectural level and to support its hardware synthesis tools. JRS uses VHDL descriptions, which describe the behavioral and structural nature of each component. Behaviors such as add, subtract, and move are defined together with timing information, as shown in the following portion of an ALU description:

```
if clk = '1' and sel = B"011001100001" then
  BEHAVIORS.iadd
   in—port1 = ›in1,
   in—port2 = ›in2,
   out—port = ›iadd16—out—port,
   status = ›iadd16—status,
   Name = ›"iadd16";
   ot‹ = iadd16—out—port after 35ns;
   cc‹ = Status—Mapiadd16—status, CARRY, OVERFLOW,
NEGATIVE, NOT—ZERO
   after 35ns;
end if;
```

VHDL structural description, or net-list, defines instances of components and the interconnections to other components as shown in this next example:

```
ALU: ALU16
   generic map ROW = ›1, COL = ›3, ROTATION = ›90
   portmap
   cc = ›alucc,
   clk = ›phase3,
   in1 = ›muxaoutport,
   in2 = ›muxboutport,
   ot = ›aluot,
   selout = ›alusel
   ;
```

JRS provides a detailed style guide, with examples, that defines its requirements. As standard VHDL documents are defined in terms of their information contents, JRS will redefine its style requirements in terms of standards.

the code size of the matrix multiply inner loop (Figure 4b) has increased to four microinstructions, increasing the execution time from 1,000 cycles to 2,000 cycles. The machine, however, is smaller and may satisfy the power and cost requirements of the designer. To implement this machine in the compiler primitives, IDAS uses the following list of components:

3 ALU16 (one is used as a register-ALU, two are address generators)
1 clock8
1 control-store
5 latch16
2 memory
3 mux21
3 mux41
8 mux81
3 regstor8
1 sequencer

The machine design is converted into VHDL description for which the system creates a full code-generation environment, including a VHDL simulator. Ada applications can be written and simulated on the new target machine in a matter of a few hours. The synthesis capability of IDAS provides an embedded-computer designer with the capability necessary to do true "software first" design.

■ Design Process Instrumentation

A basic problem in the design of high-performance embedded computing systems is evaluating the mapping of a selected set of algorithms onto a selected machine. The design process attempts to synthesize a machine that is particularly well-suited to execute some key algorithms; it also seeks to structure an algorithm to make it execute especially well on a specific machine architecture.

The system provides instrumentation routines to deliver the information that designers may use to improve their designs intelligently and to iterate toward high-quality solutions. Numerous software elements that perform analyses and optimizations help the designer map an algorithm onto a target machine architecture—precisely the problem that the designer is attempting to solve. In performing their functions, the software elements derive information about the relationship between algorithms and the machine, and discover the reasons why an optimum mapping cannot be achieved. This information consists of such things as datapath, control-field and resource conflicts as well as identification of software constructs that are not well-supported in the hardware. This information is invaluable to a designer attempting to optimize a hardware/software design.

The JRS IDAS resides on both the VAX and the IBM PC computers. The VAX, which contains most of the computational and database elements of the system, has a

```
    t4 = IADD16( SUM, t3);
       t7 = IMUL16( t5, t6);
       JNZD ( IS@TMP3, L/72611);
       t1 = Read16( P@LT2, p1);
       t2 = Read16( P@LT1, p2);
       p4 = PADD16( P@LT2, #1);
       p5 = PADD16( P@LT1, #10);
    SUM = IADD16( t4, t7);
       NOP;
       t3 = IMUL16( t1, t2);
       t5 = Read16( p4, P@LT2);
       t6 = Read16( p5, P@LT1);
       P@LT2 = PADD16( p4, #1);
       P@LT1 = PADD16( p5, #10);
[ A ]

    t4 = IADD16( SUM, t3)
    p4 = PADD16( P@LT2, #1);
       t7 = IMUL16( t5, t6)
    JNZD( IS@TMP3, L/72611);
       SUM = IADD16( t4, t7);
       p5 = PADD16( P@LT1, #10);
       t2 = Read16( P@LT1, p2);
       t1 = Read16( P@LT2, p1);
    NOP;
       P@LT1 = PADD16( p5, #10);
       P@LT2 = PADD16( p4, #1);
       t6 = Read16( p5, P@LT1);
       t5 = Read16( p4, P@LT2);
       t3 = IMUL16( t1, t2);
[ B ]
```

Figure 4. The first pass of the synthesizer creates a highly parallel machine that executes in fast, horizontal code (a); limiting the machine resources to reduce size, power consumption, and cost means that the same algorithm requires more clock cycles to execute (b).

VMS environment with the software implemented in Ada and Pascal. The PCs are used primarily as IDAS workstations with the interactive synthesis processes and the system data-entry processes residing there. The PC software is implemented primarily in Prolog and Smalltalk V. ■

REFERENCES

FISHER, J., 1981, *"Trace Scheduling: A Technique for Global Microcode Compaction," IEEE Transactions on Computers, Vol C-30, No 78, pp 478-490, July 1981.*

GIESER, J., 1982, *"On Horizontally Microprogrammed Microarchitecture Description Techniques," IEEE Transcactions on Software Engineering, Vol SE-8, No 5, pp 513-525, September 1982.*

ISODA, S. et al, 1983, *"Global Compaction of Horizontal Microprograms Based on the Generalized Data Dependency Graph," IEEE Transactions on Computers, Vol C-32, No 10, pp 922-932, October 1983.*

LANDSKOV, D. et al, 1980, *"Local Microcode Compaction Techniques," Computing Survey, Vol 12, pp 261-294, September 1980.*

SHERAGA, R. and GIESER, J., 1983, *"Experiments in Automatic Microcode Generation", IEEE Transactions on Computers, Vol C-32, No 6, pp 557-569, June 1983.*

ABOUT THE AUTHORS

DEBORAH W. RUNNER *is the Product Line Manager for JRS Research Laboratories. She was a major contributor to the IDAS system, particularly with the user interface and design-data instrumentation. She has 15 years of technical and marketing experience, including work with the Supercomputer Applications Association, California Computer Consultants, TRW and Data General. She holds a BA in Mathematics from the University of California, Riverside.*

ERWIN H. WARSHAWSKY *is the President of JRS Research Laboratories. He has contributed to the IDAS product concept and system design. He has 30 years of management, product and system design, and marketing experience, including prior associations with Data General, Teledyne Systems, Planning Research and Douglas Aircraft. He holds a BS in Physics from the Illinois Institute of Technology.*

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

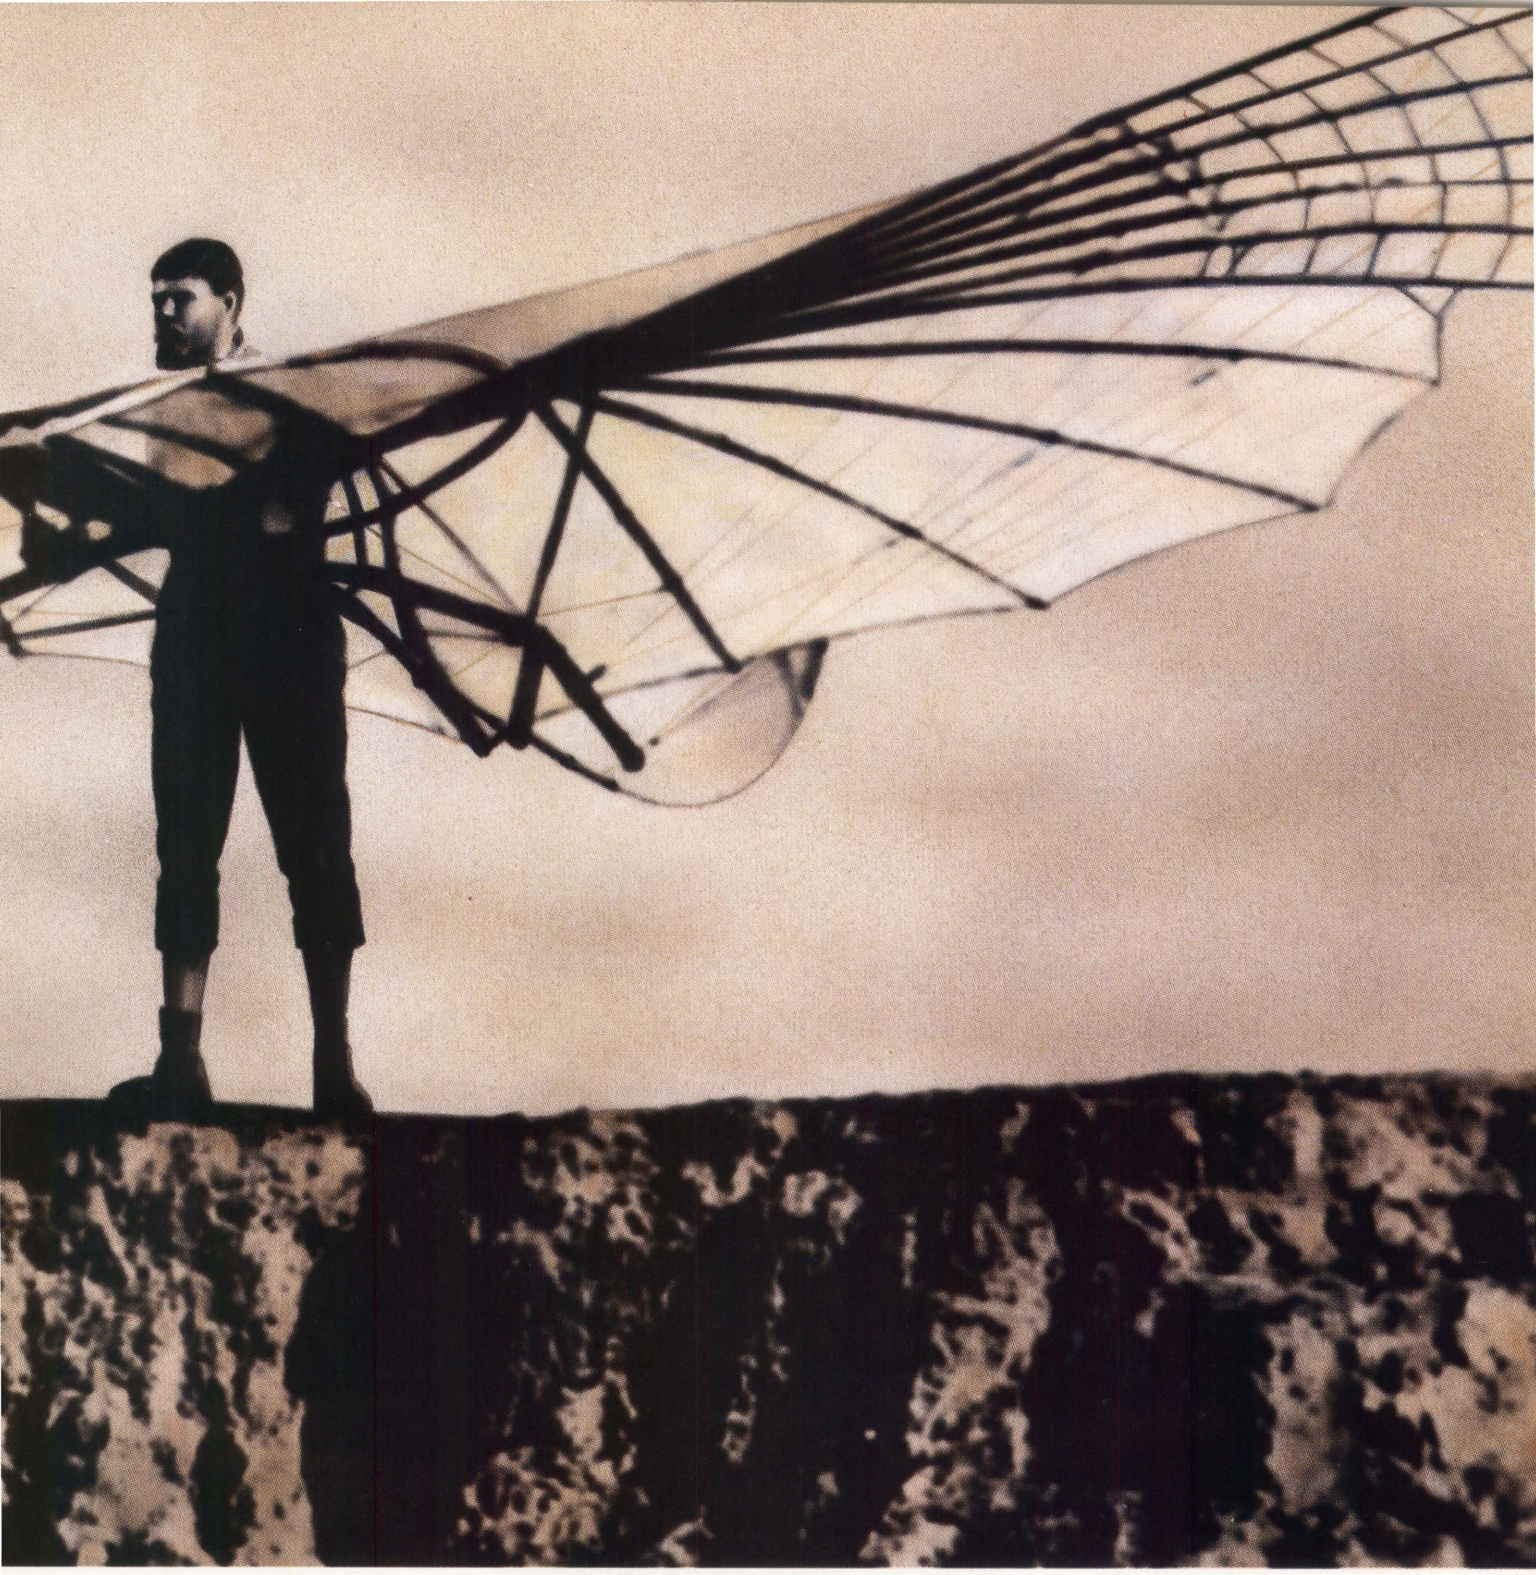

simulators running on mainframes or workstations.

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.

PCB-CAD

A Survey of PCB-CAD Vendors

VLSI Systems Design Staff

Computer-aided layout of printed circuit boards has been around for over twenty years. Many of the originally tools were complex and difficult to use. But fortunately over time the tools have been enhanced and are now not only user-friendly, but several orders of magnitude faster than the early tools. However, although most of today's PCB-CAD tools are interactive, the actual layout, except for a few very rudimentary boards, is still mainly a batch process, which can take hours or even days to complete.

SYSTEMS AVAILABLE FOR PC'S TO MAIN FRAMES

When it comes to percentage completion, there are many definitions. Some vendors, without blinking an eye, promise "100 percent completion" with no any strings attached; others say 100 percent completion for all printed circuit boards that are capable of being 100 percent completed; and others quote typical automatic completion percentages in the high nineties. These latter vendors sometimes suggest stopping the automatic routing when the board has been routed to either a user-selected percentage of completion or when only the user-selected number of nets are still unconnected. The remaining nets can often be completed faster manually than in the automatic mode.

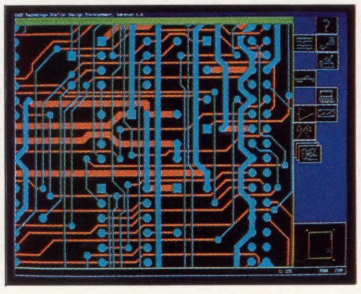

While all of the systems listed have some degree of automatic routing, only about half have automatic placement features. Although many designers like to do the placement manually, the automatic feature can be a big time saver, particularly for complex, high-density boards. Also, the automatic placement function usually provides for user input of special initial placements for specified components.

Since many of the features can vary considerably from one vendor to another, the final selection of a PCB-CAD system, should, if possible, be based on a benchmark of the systems that can handle the user's typical boards, and that are priced within the user's budget.

In addition, all systems in our list have many standard features such as menu-driven interface, ASCII library representation, capability of moving data between all layers, multiple trace-width support, realtime trace stretching, on-line modification of grid size, zoom, pan, consistent grid at all zoom levels, and multi-page design support.

Those readers familiar with our last few PCB-CAD surveys, will see that this year we have included only a short form of the directory—listing the vendors, their product features, and the contact names, addresses and phone num-

Simple as...

pushing a button!

Design Engineer	PCB Engineer
Hybrid Engineer	CAM Engineer
Mechanical	Analysis
Publication	Database

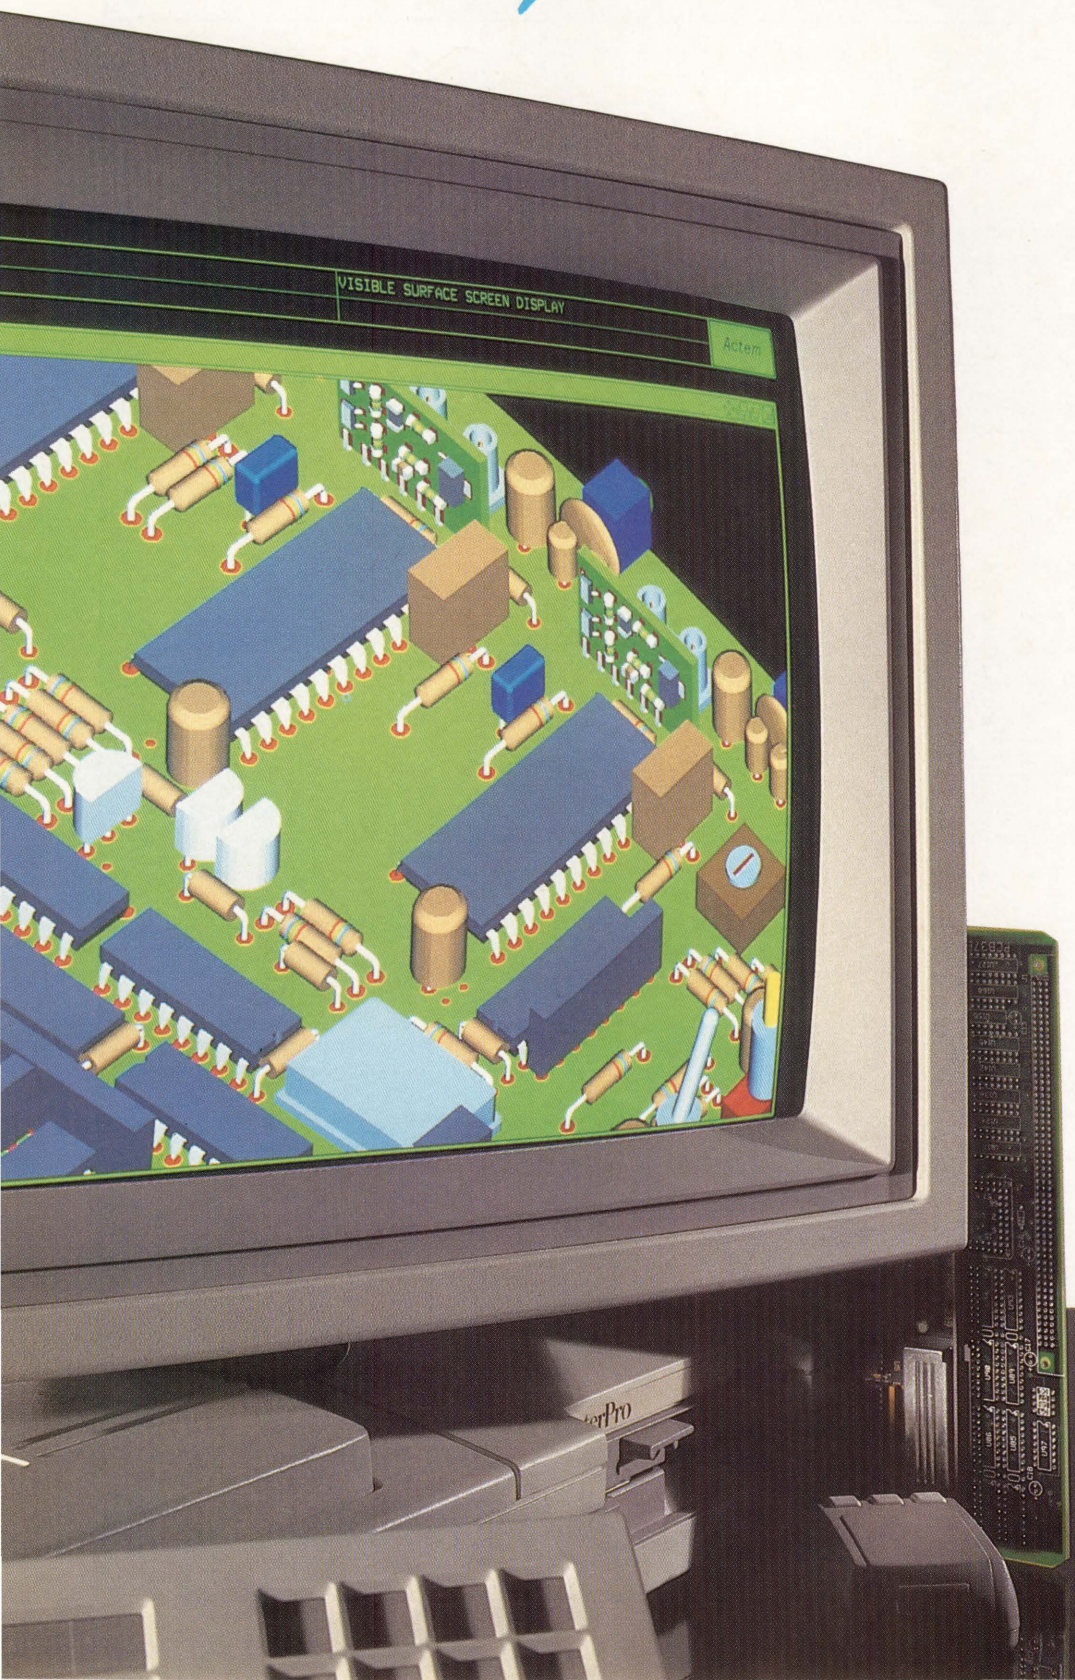

Intergraph makes PCB design and manufacturing easy. Our interactive graphics solutions address a wide variety of design and manufacturing methodologies. And our product interface is designed for both the beginning and the expert user.

From your workstation, simply push a button to...

- ☐ **Place**
- ☐ **Route**
- ☐ **Evaluate**
- ☐ **Package**
- ☐ **Document**
- ☐ **Output**

In short, Intergraph provides everything you need to get your products to market faster.

Our 5-MIPS CLIPPER workstations and servers give you the performance to handle your most difficult boards, while allowing you to pass data via Ethernet and other industry-standard networks.

To see how Intergraph can make your PCB work as simple as pushing a button, call 800-826-3515 or 205-772-2700.

CIRCLE NUMBER 14

INTERGRAPH

	ERC	DRC	ASCII OUT	ASCII IN	DRAW DOTTED LINES	DRAW ANGLED LINES	RUBBERBANDING	PLOWING	ON-LINE PAGE TRAVERSAL	AUTO NET CONNECTION	POINT-TO-POINT WIRING	KEYBOARD MACROS	TEXT ANY SIZE	EDIT SYMBOLIC DATABASE	ANGLED TEXT	AUTO NET NAMING	AUTO PLACEMENT	AUTO PIN NAMING	AUTO INSTANCE NAMING	SNAP CURSOR TO GRID	SMT SUPPORT	LAYERS OF FIXED SIZE	NAME LAYERS	DISPLAY SELECTED LAYERS
ACADEMI SYSTEMS	—	•	•	•	•	—	•	—	•	•	•	•	•	•	—	—	•	•	•	—	•	—	•	•
ACCEL TECHNOLOGIES	•	•	•	•	—	•	•	—	—	—	—	—	•	•	—	—	—	•	•	•	•	•	•	•
APTOS SYSTEMS	•	•	•	•	•	•	•	—	•	•	•	•	•	•	•	•	—	•	•	•	•	—	—	•
AUTOMATED SYSTEMS	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
CADAM	•	•	•	•	•	•	•	—	•	•	•	•	•	—	•	•	•	•	•	•	•	—	•	•
CADISYS	•	•	•	•	•	•	•	—	•	—	•	•	•	•	•	•	—	•	•	•	•	—	•	•
CADNETIX	•	•	•	•	•	•	•	—	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
CAD SOFTWARE	•	•	•	•	•	•	•	—	•	•	•	—	•	•	—	•	—	•	—	•	•	—	•	•
CALAY SYSTEMS	•	•	•	•	•	•	—	•	—	•	•	•	•	•	—	•	•	—	•	•	•	—	•	•
COMPUTAMATION	—	•	•	•	•	•	•	—	•	•	•	•	•	•	•	•	•	•	—	•	•	—	—	•
COMPUTERVISION	•	•	•	•	•	•	•	—	—	—	•	—	•	•	—	•	•	•	—	•	•	—	•	—
CONTROL DATA	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
DAISY SYSTEMS	•	•	•	•	•	•	•	•	•	•	•	•	•	—	•	•	•	•	•	•	•	—	•	•
DOUGLAS	—	—	•	—	•	•	•	—	•	•	—	—	•	•	—	—	—	—	—	•	•	•	•	•
HEWLETT-PACKARD	•	•	•	•	•	•	•	—	•	•	—	—	•	•	—	•	•	•	•	•	•	—	•	•
IBM	—	•	•	•	•	•	•	•	—	•	•	—	•	•	•	—	•	—	•	•	•	—	•	•
INTERACTIVE CAD	•	•	—	•	•	•	•	—	•	•	•	•	•	•	—	•	—	•	•	•	•	—	•	•
INTERGRAPH	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	—	•	•
MENTOR GRAPHICS	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	—	•	•
OMATION	—	•	•	•	—	•	•	—	•	—	•	—	•	•	—	•	—	•	—	•	•	—	•	•
P-CAD SYSTEMS	•	•	•	•	•	—	•	•	•	•	—	•	•	•	•	•	•	•	•	•	•	—	•	•
RACAL-REDAC	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	—	•	•
ROYAL DIGITAL	•	•	•	•	•	•	•	—	•	•	•	•	•	•	•	•	•	•	•	•	•	—	•	—
SCHLUMBERGER	•	•	•	•	•	•	—	—	•	•	•	•	•	•	•	•	—	•	•	•	•	—	•	—
SCI	•	•	•	•	•	•	•	•	•	•	•	•	•	•	—	•	—	•	•	•	•	—	•	•
SHARED RESOURCES	•	—	•	•	•	•	—	—	—	•	•	•	•	•	•	—	—	—	—	—	•	—	•	•
TERADYNE EDA	•	•	•	•	—	•	•	—	•	•	•	•	•	•	—	•	—	•	•	•	•	—	•	•
VALID LOGIC SYSTEMS	•	•	•	•	—	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	—	•	•
VAMP	—	—	•	—	•	•	•	—	—	•	•	•	•	•	—	•	—	•	•	—	•	•	•	•
VISIONICS	•	•	•	—	•	•	•	—	—	•	•	—	•	•	—	•	•	•	—	•	•	—	•	•
WINTEK	•	•	•	•	•	—	—	—	—	—	•	—	•	—	—	•	—	•	—	•	•	•	—	—

Directory Of PCB-CAD Vendor Contacts

ACADEMI SYSTEMS INC.
2418 Armstrong St.
Livermore, Calif. 94539
Cynthia Peddie
Vice President
(415) 449-3294

ACCEL TECHNOLOGIES INC.
7358 Trade St.
San Diego, Calif. 92121
Ray Schnorr
Vice President of Marketing
(619) 695-2000

APTOS SYSTEMS CORP.
5274 Scotts Valley Dr.
Scotts Valley, Calif. 95066
John Roth
President
(408) 438-2199

AUTOMATED SYSTEMS INC.
4105 Sorrento Valley Blvd.
San Diego, Calif. 92121
Andy Chmielinski
Manager, PranceGT Product Marketing
(619) 546-0024

CADAM INC.
1935 N. Buena Vista St.
Burbank, Calif. 91504
Alan Cohen
Electrical Products Marketing
(818) 841-9470

CADISYS CORP.
624 E. Evelyn Ave.
Sunnyvale, Calif. 94086
Paul Ling
Customer Support Manager
(408) 732-1832

CADNETIX CORP.
5757 Central Ave.
Boulder, Colo. 80301
Greg Skomp
Manager of Marketing and Communications
(303) 444-8075

CAD SOFTWARE INC.
Box 1142
Suite 6
Littleton, Mass. 01460
H. Marsh
(617) 486-9521

CALAY SYSTEMS INC.
16842 Von Karman Rd.
Suite 100
Irvine, Calif. 92714
Nancy J. Nykanen
Marcomm Manager
(714) 863-1700

COMPUTAMATION SYSTEMS LTD.
40 Lake St.
Leighton Buzzard
Bedfordshire LU7 8RX
England
J. Yelland
Phone: (0525) 378939

COMPUTERVISION CORP.
100 Crosby Dr.
Bedford, Mass. 01734
Barry Heller
Manager, Electronic Product Marketing
(617) 275-1800

CONTROL DATA CORP.
8100 34th Ave. South
P.O. Box 0
Minneapolis, Minn. 55440
John W. Barnes
(612) 853-8390

DAISY SYSTEMS CORP.
700 Middlefield Rd.
Mountain View, Calif. 94039
Jerry Harvel
Marketing Manager
(415) 960-0123

DOUGLAS ELECTRONICS INC.
718 Marina
San Leandro, Calif. 94577
Dana Dotson
Marketing Manager
(415) 483-8770

HEWLETT-PACKARD CO.
19310 Pruneridge Ave.
Bldg. 49AV
Cupertino, Calif. 95104-9826
Customer Information Center
(800) 752-0900

IBM CORP.
1503 LBJ Freeway
6th Floor
Dallas, Tex. 75234
C.W. Liles
CAD/CAM/CAE Marketing Programs Manager
(214) 406-7518

INTERACTIVE CAD SYSTEMS
2352 Rambo Court
Santa Clara, Calif. 95054
Eddy Ozomaro
President
(408) 970-0852

INTERGRAPH CORP.
1 Madison Industrial Park
Huntsville, Ala. 35807
Randy Anderson
Product Marketing Manager
(205) 772-2000

MENTOR GRAPHICS CORP.
1940 Zanker Rd.
San Jose, Calif. 95112
Jim Behrens
Director, PCB Marketing
(408) 436-1500

OMATION INC.
1210 East Campbell Rd.
Suite 100
Richardson, Tex. 75081
Jackie Williams
Marketing Analyst
(800) 553-9119
(214) 231-5167

PERSONAL CAD SYSTEMS INC. (P-CAD)
1290 Parkmoor Ave.
San Jose, Calif. 95126
Sales Desk
(408) 971-1300

RACAL-REDAC INC.
238 Littleton Rd.
Westford, Mass. 01886
Susan Cook
Marketing Communications Manager
(508) 692-4900

ROYAL DIGITAL SYSTEMS INC.
2855 Kifer Rd.
Santa Clara, Calif. 95050
A.V. Seshadri
Vice President, Sales
(408) 970-0909

SCHLUMBERGER CAD/CAM
4251 Plymouth Rd.
PO Box 986
Ann Arbor, Mich. 48106
Brian F. Barton
Director of Electronics
(313) 995-6000

SCIENTIFIC CALCULATIONS INC.
7796 Victor-Mendon Rd.
Fishers, N.Y. 14453
Douglas W. Spice
Director, Marketing Services
(716) 924-9303

SHARED RESOURCES INC.
3047 Orchard Park
San Jose, Calif. 95134
Robert Wong
Executive Vice President
(408) 434-0444

TERADYNE EDA
2141 Landings Dr.
Mountain View, Calif. 94043
Steve Chidester
PCB Product Manager
(415) 962-1440

VALID LOGIC SYSTEMS INC.
2 Omni Way
Chelmsford, Mass. 01824
Katherine Gambino
Product Marketing Manager
(617) 256-2300

VAMP INC.
6753 Selma Ave.
Los Angeles, Calif. 90028
J. Soluk
Marketing Director
(213) 466-5533

VISIONICS CORP.
343 Gibraltar Dr.
Sunnyvale, Calif. 94089
Alex Wellins
Director of Marketing
(408) 745-1551

WINTEK CORP.
1801 South St.
Lafayette, Ind. 47904
Marcia Borton
CAD Sales
(317) 742-8428

bers. A more complete directory listing product names, prices, and size and accuracy limitations will be published as usual in our annual *User's Design Guide to Design Automation*.

As a result of a number of mergers and dropouts, our annual survey listing of PCB-CAD vendors has shrunk from 36 to 31. Those companies that are no longer in the PCB-CAD business or chose not to be listed in our 1988 survey, include: Bishop Graphics CAD Systems Corp.; Calma Co., a division of General Electric Co. that dropped its PCB-CAD line, and sold its IC-CAD business to Valid Logic Systems Inc.; FutureNet Corp., now a Data I/O division that is concentrating on the CAE end of the market devoted to programmable logic devices; Modula Corp.; Optima Technology Inc.; Pro.Lib Inc.; Tektronix Inc., which sold most of its design automation business to Mentor Graphics Corp.; Vectron Graphic Systems Inc.; and Xerox Corp.

Newcomers to our 1988 list are Cadisys Corp. and Omation Inc. Several old timers still appear on the list, but are listed under their new corporate names. Applicon, which has been a subsidiary of Schlumberger, is now re-labeled Schlumberger CAD/CAM, while Case Technology Inc. now sports the name of its new parent company, as Teradyne EDA.

A V·L·S·I JOIN MODULE

C. R. PETRIE AND A. R. HURSON, PENNSYLVANIA STATE UNIVERSITY, UNIVERSITY PARK, PENN.

Searching a database record (or tuple) calls for massive parallelism. What is desired is a memory so intelligent that an associative search can be performed right in the memory itself.

Rather than incur the von Neumann bottleneck of ordinary computer architectures—the problem that occurs when words must be addressed and brought in one by one to a single central ALU—it would be much better if multiple ALUs could be distributed out in memory so the search could be made where the data resides. It would be ideal if the reference or source data to be searched for could be simultaneously compared directly against all the target data in the memory with no extraneous addressing or data-movement steps.

The VLSI circuit designs described in this article are capable of this, and it was found that the increase in total memory cell size with the addition of local intelligence was only three or four times. This is not an unreasonable trade-off, considering the orders of magnitude of speedup in record searching that can result.

Massive Parallelism and VLSI Chip Tackle Database Problems

The objective of this "associative join chip" (AJC) was to accomplish an associative search of memory records, or tuples, for a data-base join operation. For example, key words in two fragmented data bases files could be matched so that a third composite file could be assembled (see Sidebar).

The architecture of the AJC chip is shown in Figure 1. It can be viewed as an 8,000-word by 40-bits/word memory array with local intelligence logic added at every memory bit position (details given in Figures 2, 3 and 4) so that the comparisons can be made right in the memory. Additional logic in the form of a comparand register (shown at top of Figure 1) and tag bit logic (at right or least

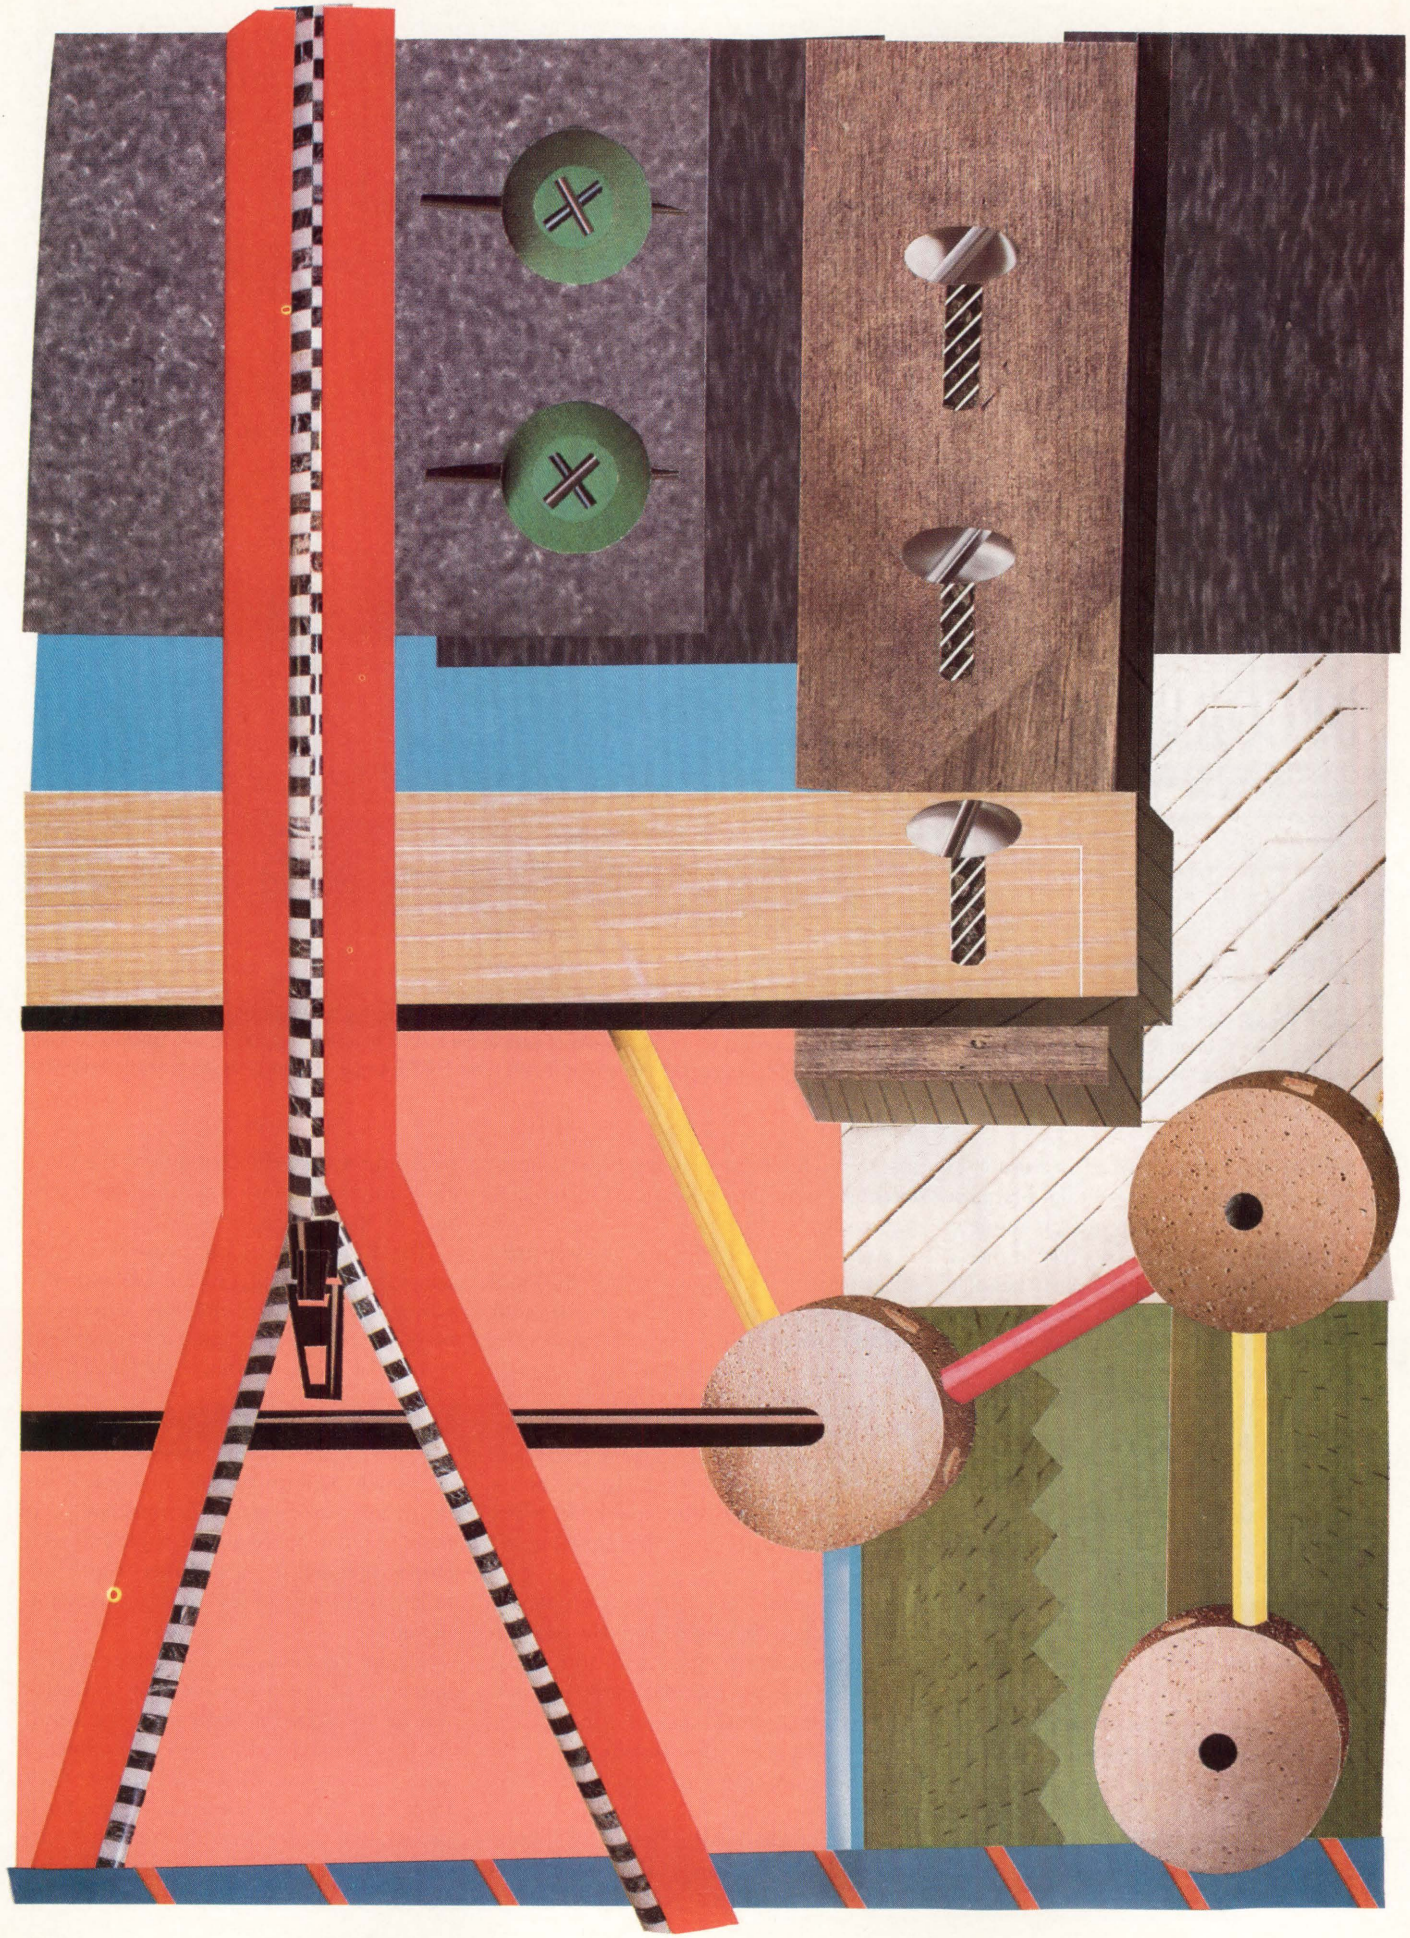

ILLUSTRATION BY RENEE KLEIN

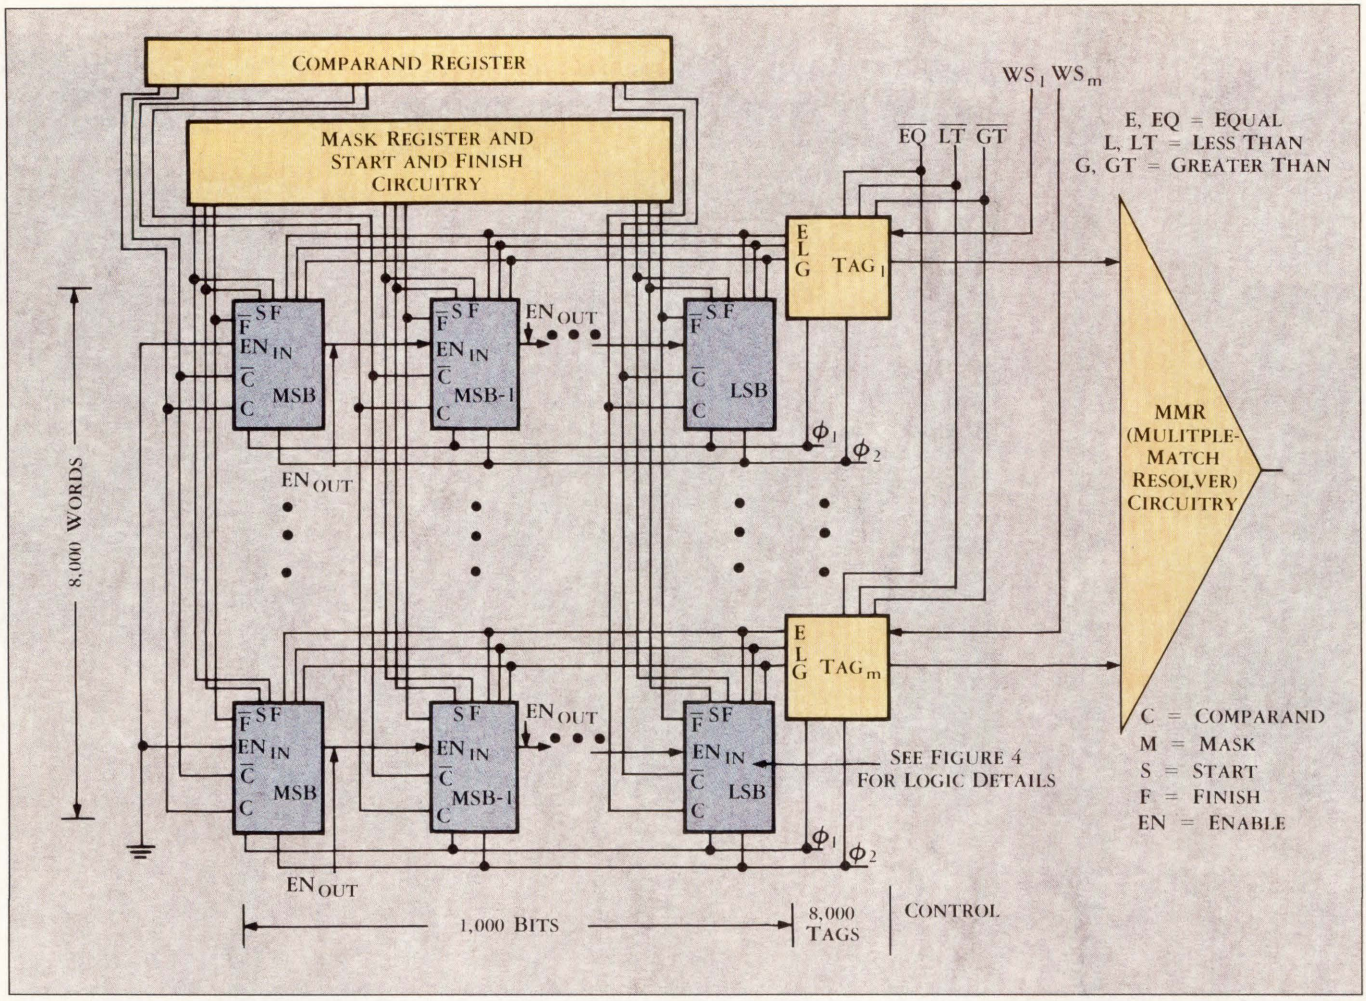

FIGURE 1. The AJC (associative join chip) incorporates 8,000 words, each 40 bits long. The single 40-bit comparand word at the top is simultaneously compared against the 8,000 rows of words below to see which match. The massive parallelism is carried out to the extreme of each of the 8 Mbits of memory, having its own comparison logic. Different ways of doing the comparison will be described in FIGURES 2, 3 and 4; this diagram relates specifically to the FIGURE 4 approach, but is helpful in understanding the purpose of the AJC.

significant end of each word in Figure 1) allows the whole process to take place on the AJC chip. A multiple-match resolver unit (MMR) is shown to the very right of Figure 1. As its name implies, it handles the situation where the tag outputs indicate that more than one word matches the source criteria of the unmasked portion of the comparand.

■ MODIFIED AJC USES SHIFT REGISTER

Figure 2 shows an early version of the AJC that backs off from full parallelism because it uses a shift register. With the shift register, it was only necessary to have one word worth of comparison logic. This compromise allowed it to be more readily put into silicon and used for initial verification of the joining process and the validity of the associative join algorithm. This chip was designed using IBM's MVISA standard-cell CAD system at Pennsylvania State University and then fabricated by IBM. Due to the limitation of the MVISA cell library and chip size, the proposed join system architecture was modified to perform the equi-join operation by connecting multiple chips.

As shown in 2a, this AJC is essentially an associative memory with a comparand register, mask register and a circular shift register bank (CSRB). The CSRB stores the data, and then combinational circuitry compares the unmasked portion of the comparand with a portion of the CSRB.

A two-phase clocking system (BCLK and CCLK) is used. The chip can also be configured in a test mode (ACLK and BCLK) where IBM's level-sensitive scan design (LSSD) becomes active.

The control line INPUT COMPARAND enables data to be loaded into the comparand register while the control line INPUT MASK is used to load the mask register. The LOAD/SEARCH control line distinguishes between the loading and the searching modes of the associative memory.

When in the loading mode, the target tuples are sequentially loaded into the CSRB shift register. And then in the searching mode, the contents of the CSRB are compared against the unmasked portion of the comparand register.

The MISMATCH IN functions as a word select so that data can be individually selected. This allows multiple searches to be performed on the data—for example as would be needed in a join-based on AND/OR.

The MISMATCH OUT signals of all the individual AJCs are sent to a mismatch resolver, which determines the total result for the whole tuple (match or mismatch). If there is a match, the source tuple and each chip's portion of the target tuple are linked together and sent to the destination memory.

Figure 2b shows the flow chart for the sequence of operations the AJC performs during an execution of the join operation. As the contents of the CSRB are compared with the comparand, they are also sent out from the chip. At the end of the BCLK period, the signal from the mismatch circuitry is latched into a D flip-flop. This signal then indicates where the target tu-

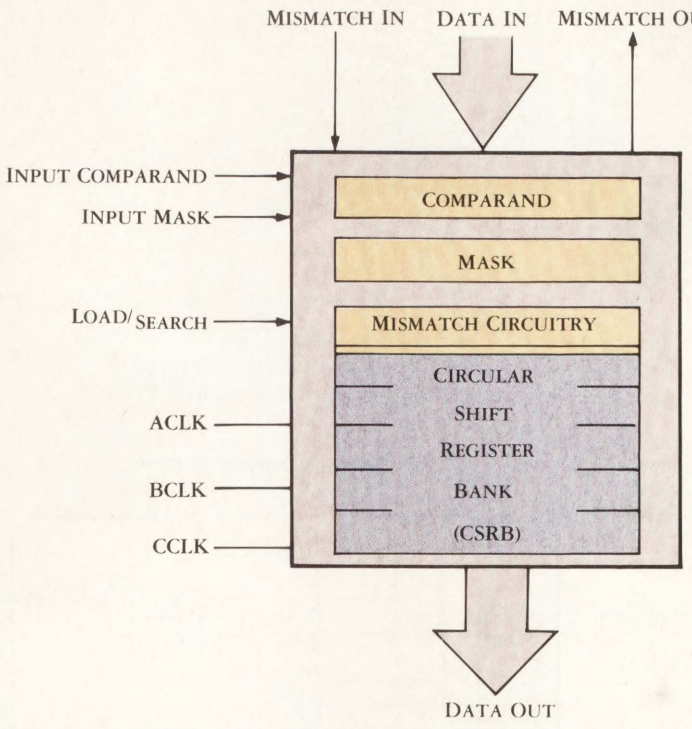

[A] SHIFT-REGISTER VERSION OF ASSOCIATIVE COMPARE CHIP

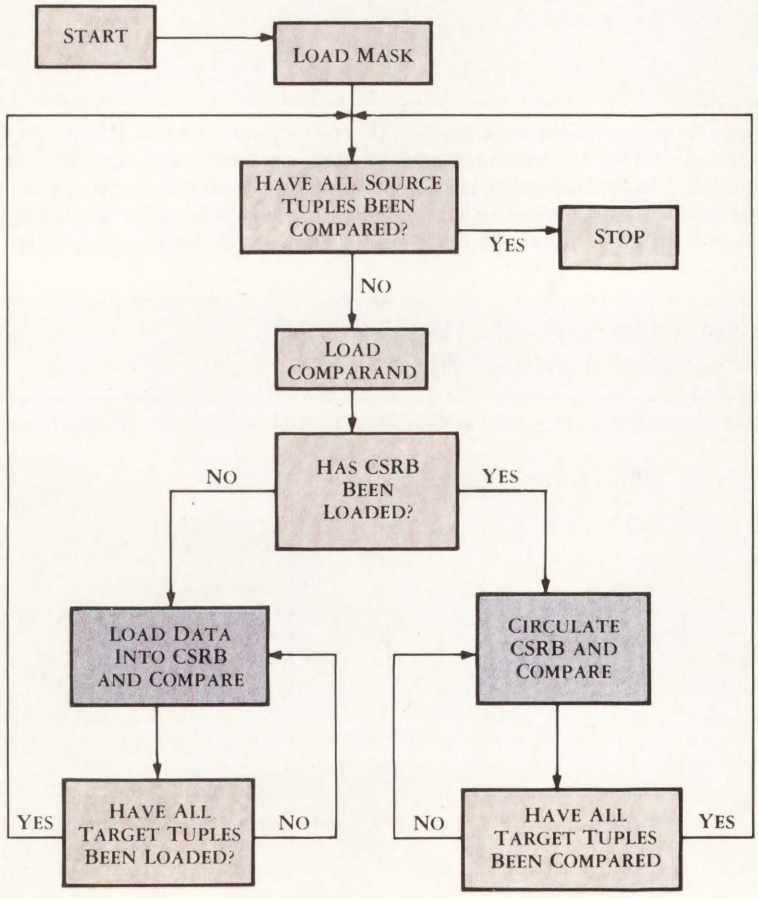

[B] FLOW DIAGRAM FOR SHIFT-REGISTER VERSION

FIGURE 2. This shift register version of the AJC was fabricated in silicon and used for prototype trials to test the validity of the join algorithm. The chip (a) is sequenced (b) so that the data is compared on a word-wide basis.

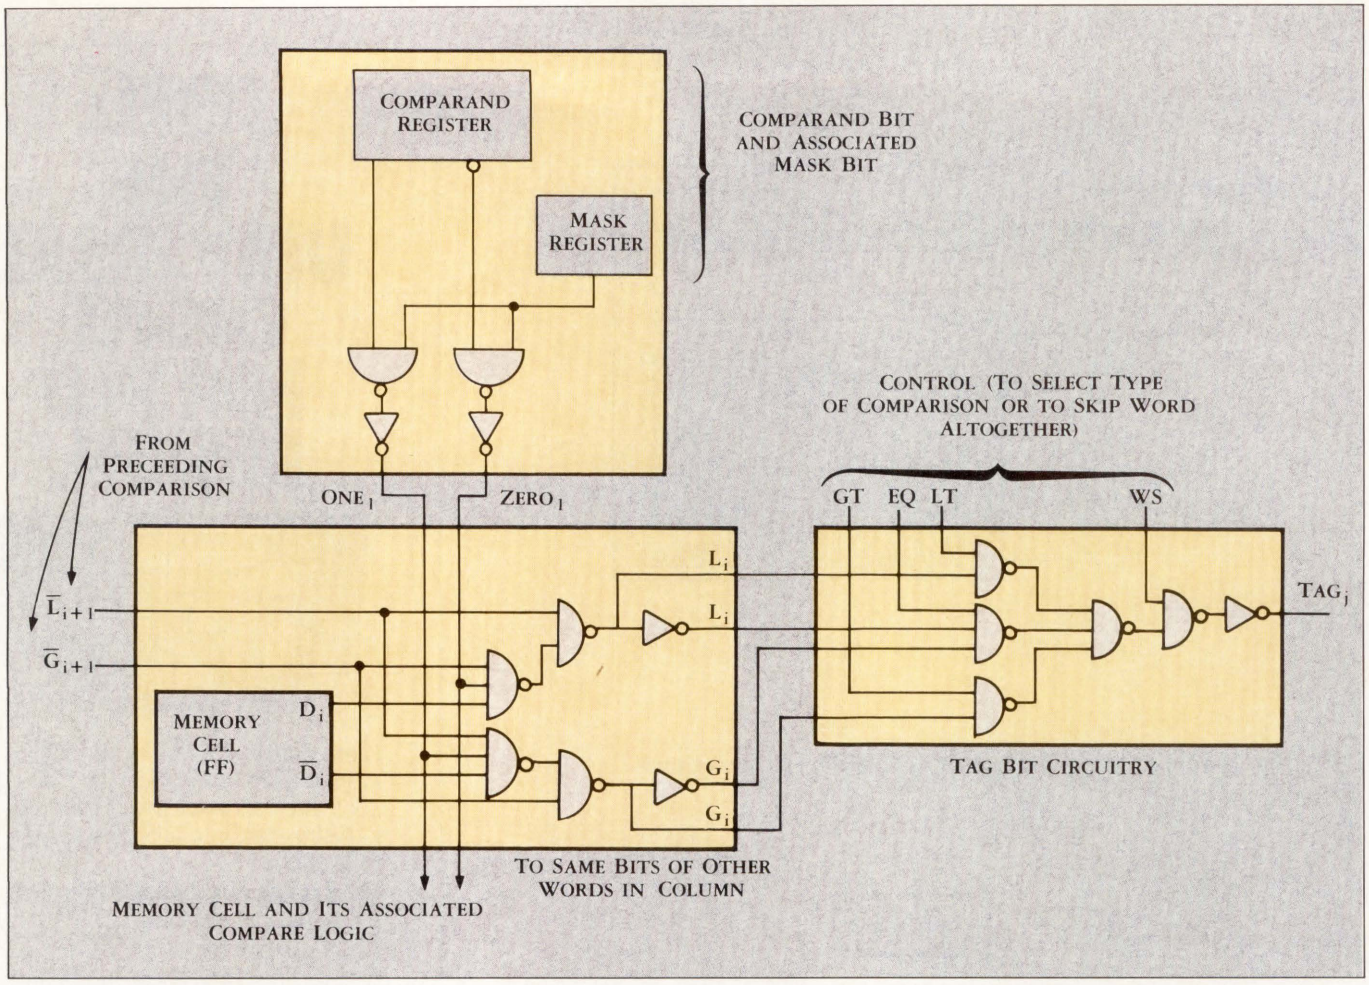

FIGURE 3. CMOS version of the first fully-parallel AJC circuitry provides the most easily understood explanation of how the comparisons are made. It depicts the logic in one of the vertical bit columns of the 8,000 words by 40-bits/word array. Starting at the top, the comparand bit (unless screened out by the mask bit) is sent down vertically to all the corresponding bits of the 8,000 words. Each memory bit has its own logic to compare itself with the reference comparand bit. With this type of circuit, the comparisons must proceed left to right along each word row so that there are decision inputs from the more significant bit on the left and the outputs on the right going to less significant bit. At the very right end of each word, the resulting output from last bit of the word is fed into the Tag bit logic, which outputs the overall decision for that word.

TABLE 1. AREA AND SPEED PROJECTIONS FOR VARIOUS NMOS AND CMOS VERSIONS OF ASSOCIATIVE MEMORY CHIP

CIRCUITRY	ORIGINALLY PROPOSED DESIGN		NAND IMPLEMENTATION OF DESIGN		NOR IMPLEMENTATION OF DESIGN		CMOS IMPLEMENTATION OF DESIGN	
	ACTIVE AREA (λ^2)	SPEED (ΔT)	ACTIVE AREA (λ^2)	SPEED (ΔT)	ACTIVE AREA (λ^2)	SPEED (ΔT)	ACTIVE AREA (λ^2)	SPEED (ΔT)
MASK/ COMPARAND	112	2	112	2	48	1	200	2
MEMORY CELL W/O DATA MEMORY	232	4L	232	3L	232	3L	464	3L
TAG BIT	256	4	256	4	128	3	508	4
TOTALS (EQUATION)	112(L) + 232(L)(m) + 256(m)	4(L) + 4	112(L) + 232(L)(m) + 256(m)	3L + 4	48(L) + 232(L)(m) + 128(m)	3L + 3	200(L) + 464(L)(m) + 508(m)	3L + 4
TOTALS $\lambda = 1\mu m$ L = 1 KBITS m = 8,000 $\Delta T = 1NS$	0.1903 Cm^2	4.100 μSEC.	0.1903 Cm^2	3.076 μSEC.	0.1902 Cm^2	3.075 μSEC.	0.3805 Cm^2	3.076 μSEC.

TABLE 2. AREA AND SPEED PROJECTIONS FOR DOMINO VERSION OF CMOS ASSOCIATIVE MEMORY CHIP

CIRCUITRY	CMOS IMPLEMENTATION OF DESIGN		DOMINO CMOS DESIGN	
	ACTIVE AREA (λ^2)	SPEED (ΔT)	ACTIVE AREA (λ^2)	SPEED (ΔT)
MASK/ COMPARAND	200	2	288	1
MEMORY CELL W/O DATA MEMORY	464	3L	338	$\Phi 1 = 3$ $\Phi 2 = 4A$
TAG BIT	508	4	612	3
TOTALS (EQUATION)	200(L) + 464(L)(m) + 508(m)	3(L) + 4	288(L) + 338(L)(m) + 612(m)	$\Phi 1 = 13$ $\Phi 2 = 14A + 3$
TOTALS W/ $\lambda = 1\mu m$ L = 1 KBITS m = 8,000 A = 8 BYTES ΔT = 1NS	0.380 Cm^2	3.08 μSEC.	0.274 Cm^2	0.90 μSEC.

ple on the output lines of the AJC qualified for connection with the source tuple are stored in the concatenation unit. Therefore, the prototype uses the MISMATCH OUT line to control the write operation (and subsequent address change) of the destination RAM.

When operated at 1 MHz, these AJCs could perform 1 million search operations per second (neglecting the first three cycles used to load the mask, comparand, and first target tuple). It was estimated that the worst-case execution time for a search would be approximately 120 ns, but so far, our laboratory measurements on the prototype have indicated search times in the range of 28 to 50 ns.

■ DEVELOPING FULLY-PARALLEL AJC DESIGNS

After the shift-register version of the AJC verified the validity of our algorithm, we then turned our attention to developing more ambitious, fully-parallel circuits for the associative join chip.

Figure 1 actually applies to our most advanced form of AJC implementation, the precharged domino CMOS circuitry that will be described in Figure 4. An earlier and perhaps easier-to-understand form of fully-parallel AJC circuitry is shown in Figure 3. The logic in Figure 3 depicts what happens in a vertical bit column of the AJC memory array—how the comparison is handled along the word row and how the tag-bit circuitry provides the result for the word.

Starting at the top, the comparand bit, unless screened out by the tag bit, is sent vertically down the bit column, being compared simultaneously to all the memory cells. As can be seen, it doesn't take an impossibly large amount of logic at each memory cell position to make the comparisons.

For the Figure 3 logic, the comparisons must proceed from left to right along the word (from most significant bit to least significant bit). The inputs to the comparison logic from the left take into account the decisions made in the more significant bits and the outputs to the right go on to less significant bits.

At the right end of the word, the output from the least significant bit comparison logic feeds into the tag bit logic. There, the control lines shown determine whether the decision is based on greater than, equal to, or less than. In addition, a word-select (WS) control allows omitting particular words from the result. Finally, the MMR circuitry handles situations where more than one word is shown to match the comparand.

The proposed AJC differs from previous associative memory designs in several ways: First, the design's complexity, size, and power consumption are reduced by implementing the read/write and search operations at word level. Second, the AJC comparisons are based on both greater-than and less-than operations. Also, the AJC makes use of an MMR, which guarantees both the efficiency and the integrity of a search resulting in multiple responses.

The AJC can be loaded up with target words and the matched words can be selectively read out because loads can be done

Our ASICs

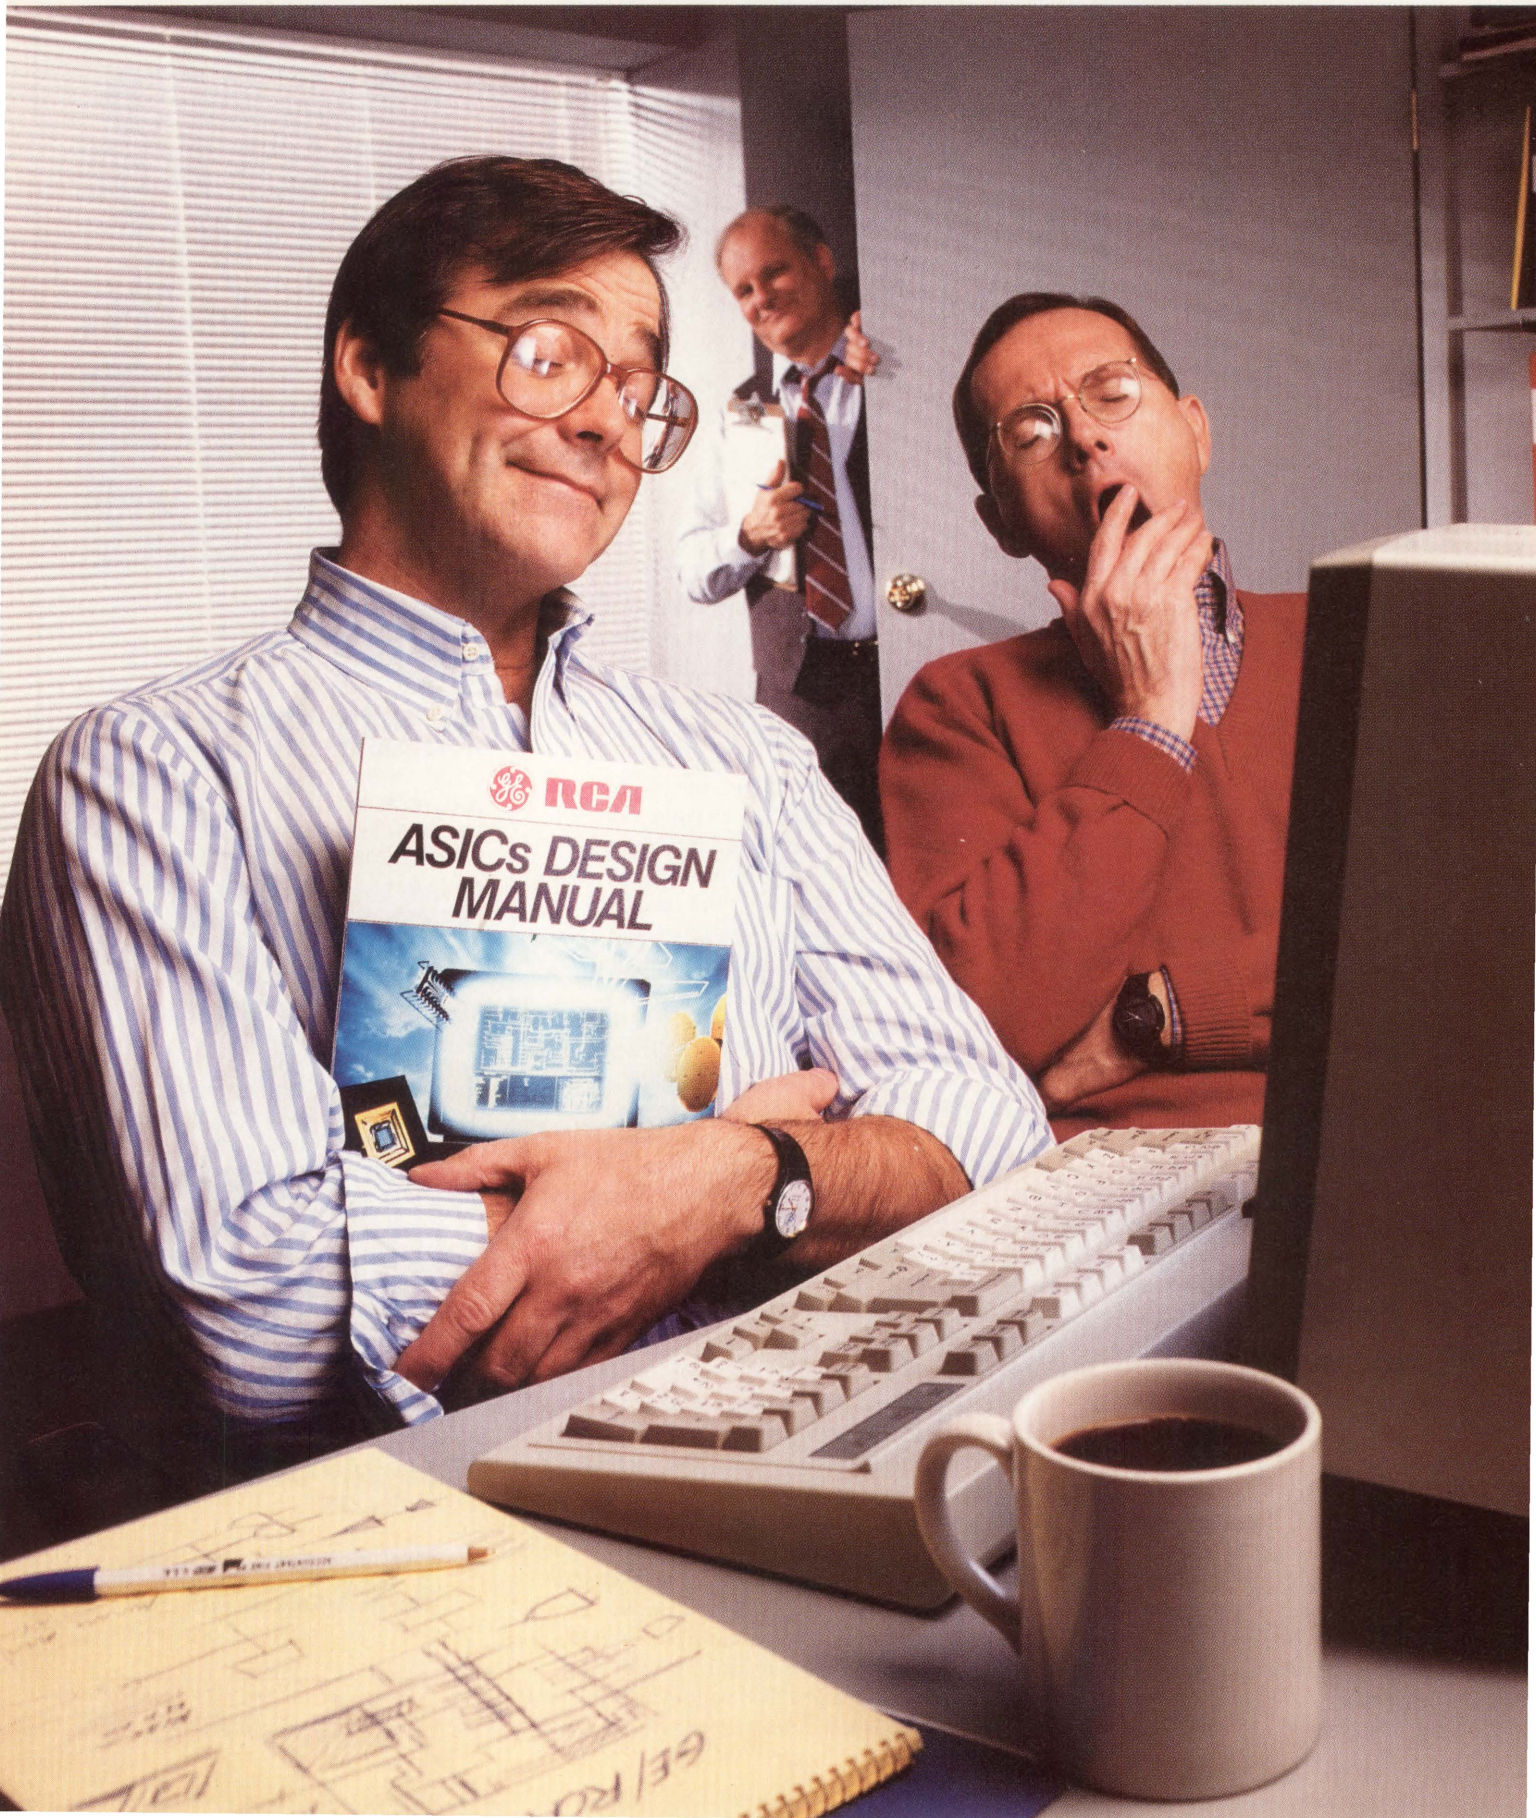

are boring.

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/RCA/INTERSIL
SEMICONDUCTORS

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

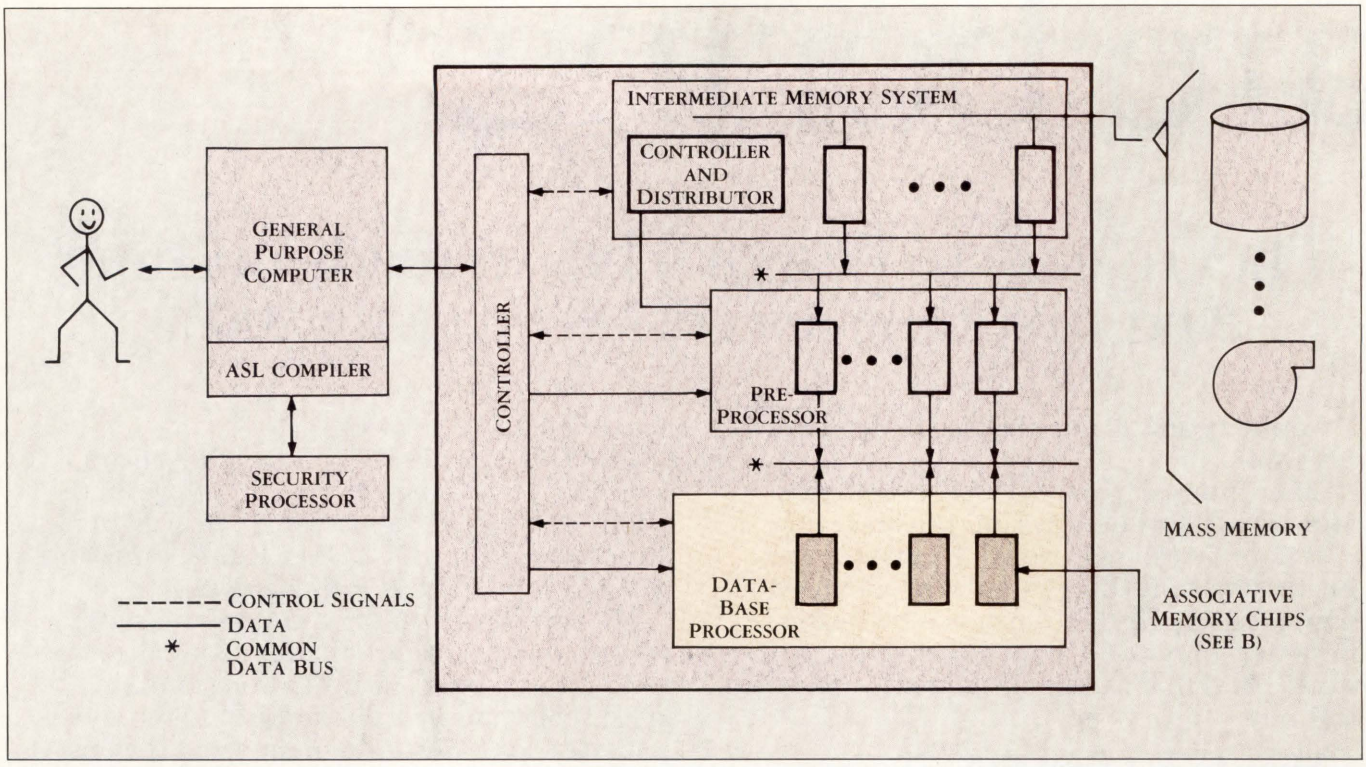

FIGURE A. The proposed database system would use the AJC's (here called associative memories) to accelerate searching, matching and joining operations. They would be used in the data base processor subsystem and preceded with caches and data filters.

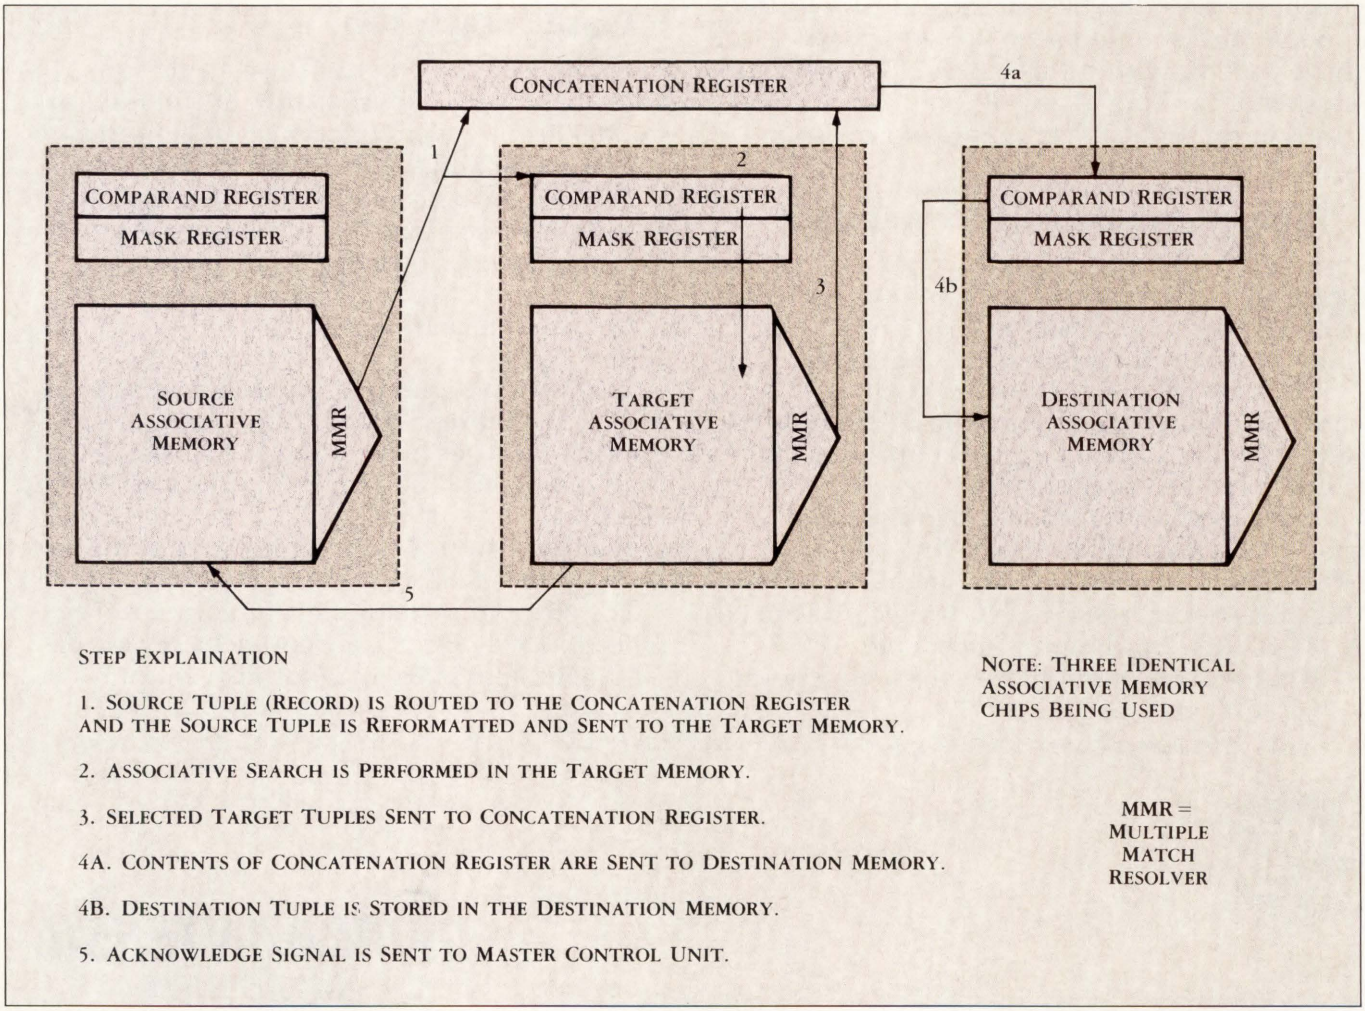

FIGURE B. The AJC chips can be used in sets of three to perform the steps called out for join operations. They are flexible enough that the controller can assigned them to different roles—a flexibility that makes them potentially redundant for fault tolerance.

The Role AJC Chips Play in Rapid Joining of Data Bases

Modern data bases are typically broken up into smaller files for efficient processing. But by the same token there is a constant need to recombine and re-arrange these smaller files into larger, or at least different files.

Assume that a company's employee data base has been broken up into a "name plus social security number" data base for the personnel department and a "salary plus social security number" data base for the payroll department. These two smaller data bases would be constantly used and updated by the two different departments and many changes would occur. Names would get updated with marriages; salaries would get changed with raises. The original file would soon become obsolete. Thus when the company president asks for a list of his employees and their salaries it would be necessary to "join" the two files so that the most recent name would be paired with the most recent salary.

The common, unchanging social security number would be the key for the merge. If the personnel file were used as the source, its records would be held one by one in the Comparand register and each searched against all the payroll records that would be loaded into the AJC memory. The mask would isolate out the social security number. Thus the corresponding names could be joined with the corresponding salaries.

An overall system for using AJC's in this sort of join operation is shown in (a). The AJC's are in the bottom block, Database Processor. Diagram (b) describes how the AJC's are used in a join operation. (Note that three identical AJC's are shown, each performing part of the operation).

The numbers cited in (b) list the five steps of the join. Before the operation is begun, the Database Processor is initialized by the controller (see a). Once preprocessed (by the upper two blocks of the associative search language machine—again see a), the source and target tuples are routed to the Database Processor and stored in their assigned AJC's.

The first step of the join operation is to send a source tuple to the concatenation register where the source tuple is temporarily stored. In addition, a copy of the source tuple has its join attribute aligned and sent to the target AJC where it is stored in the comparand register.

The second step begins once the source tuple is loaded in the target AJC module. The source tuple is compared in an associative fashion simultaneously against all the target tuples that have been loaded.

In the third step, if a match occurs, the selected target tuples are sent to the concatenation unit in a pipeline fashion. If multiple matches occur, the MMR (multiple match resolver) circuitry will control the output of selected tuples.

The fourth step is to route the contents of the concatenation unit to the designated destination AJC.

And finally, in the fifth step, after all of the selected target tuples have been outputted, an acknowledge signal is sent to the source module to prompt the loading of the next tuple, for its processing. This keeps the pipeline full.

Several of the steps can be performed in parallel. In particular steps 1, 4 and 5 can be executed concurrently with step 3. This will further increase the parallelism.

sequentially through the comparand register using the WS lines to address words. Thus the comparand register does double duty, for it serves as an I/O port for the AJC before and after the comparison.

■ VLSI Design Considerations for AJCs

AJC requirements—including modularity, regularity, and simplicity—are well-suited for VLSI implementation, with the possible exception that the area at each memory-bit position may be too large. Certainly, having 8,000,000 identical memory/logic cells in a regular orthogonal pattern is well-suited to CAD tools. The obvious design goals are to keep the memory/logic cells small, the interconnections short, and the execution speed high.

The AJCs can use either static or dynamic memory cells for the comparand, mask, and memory cells. If dynamic, the cells would refresh themselves on the alternate phase of a two-phase clock system (phase two). The detailed analysis regarding this design presented in the second reference provides some size and speed estimates and drives a complexity-performance formula for the basic search operation. The search time is estimated as:

$4\Delta t^*(L) + 4\Delta t,$

where L is the word length and Δt is the average delay time of an inverter circuit.

The cell geometry is estimated at $40\lambda \times 20\lambda$, including the area used for routing among cells. Thus, an associative memory of m words, each of L bits long occupies an area of $40L\lambda * 20m\lambda$. A search time for $m = 512$, $L = 256$ and $\Delta t = 1$ ns is estimated to be approximately 1 μs.

■ VLSI Design of Proposed Chip

Table 1 lists some possible approaches to the design for the chips and gives some computed values for their various figures of merit. The obvious goal was to find a combination that would give small area and high speed.

Our analysis showed that in logic design, for NMOS technology, NOR gates would use less active area than NAND gates. However, since the all-NOR design used more gates than an all-NAND design, the size advantage of the NOR gate could not be exploited.

We then turned our attention to CMOS technology, its well-known advantages of better speed, lower power dissipation, and noise immunity, being attractive to these "intelligent memories." However, the CMOS advantages could not be obtained without a cost. The CMOS logic gates use more transistors than their NMOS counterparts and there is always the danger of "latch up" if proper precautions are not taken.

Nevertheless, as power consumption is a problem for large systems, CMOS is still the preferred technology.

Figure 3 has already shown our basic CMOS logic. Though one cannot directly compare CMOS and NMOS devices in terms of speed and size (because NMOS is ratio-based logic and CMOS is not), it is still possible to generalize about some of the counterbalancing characteristics. Table 1 enumerates our rough first-order estimates

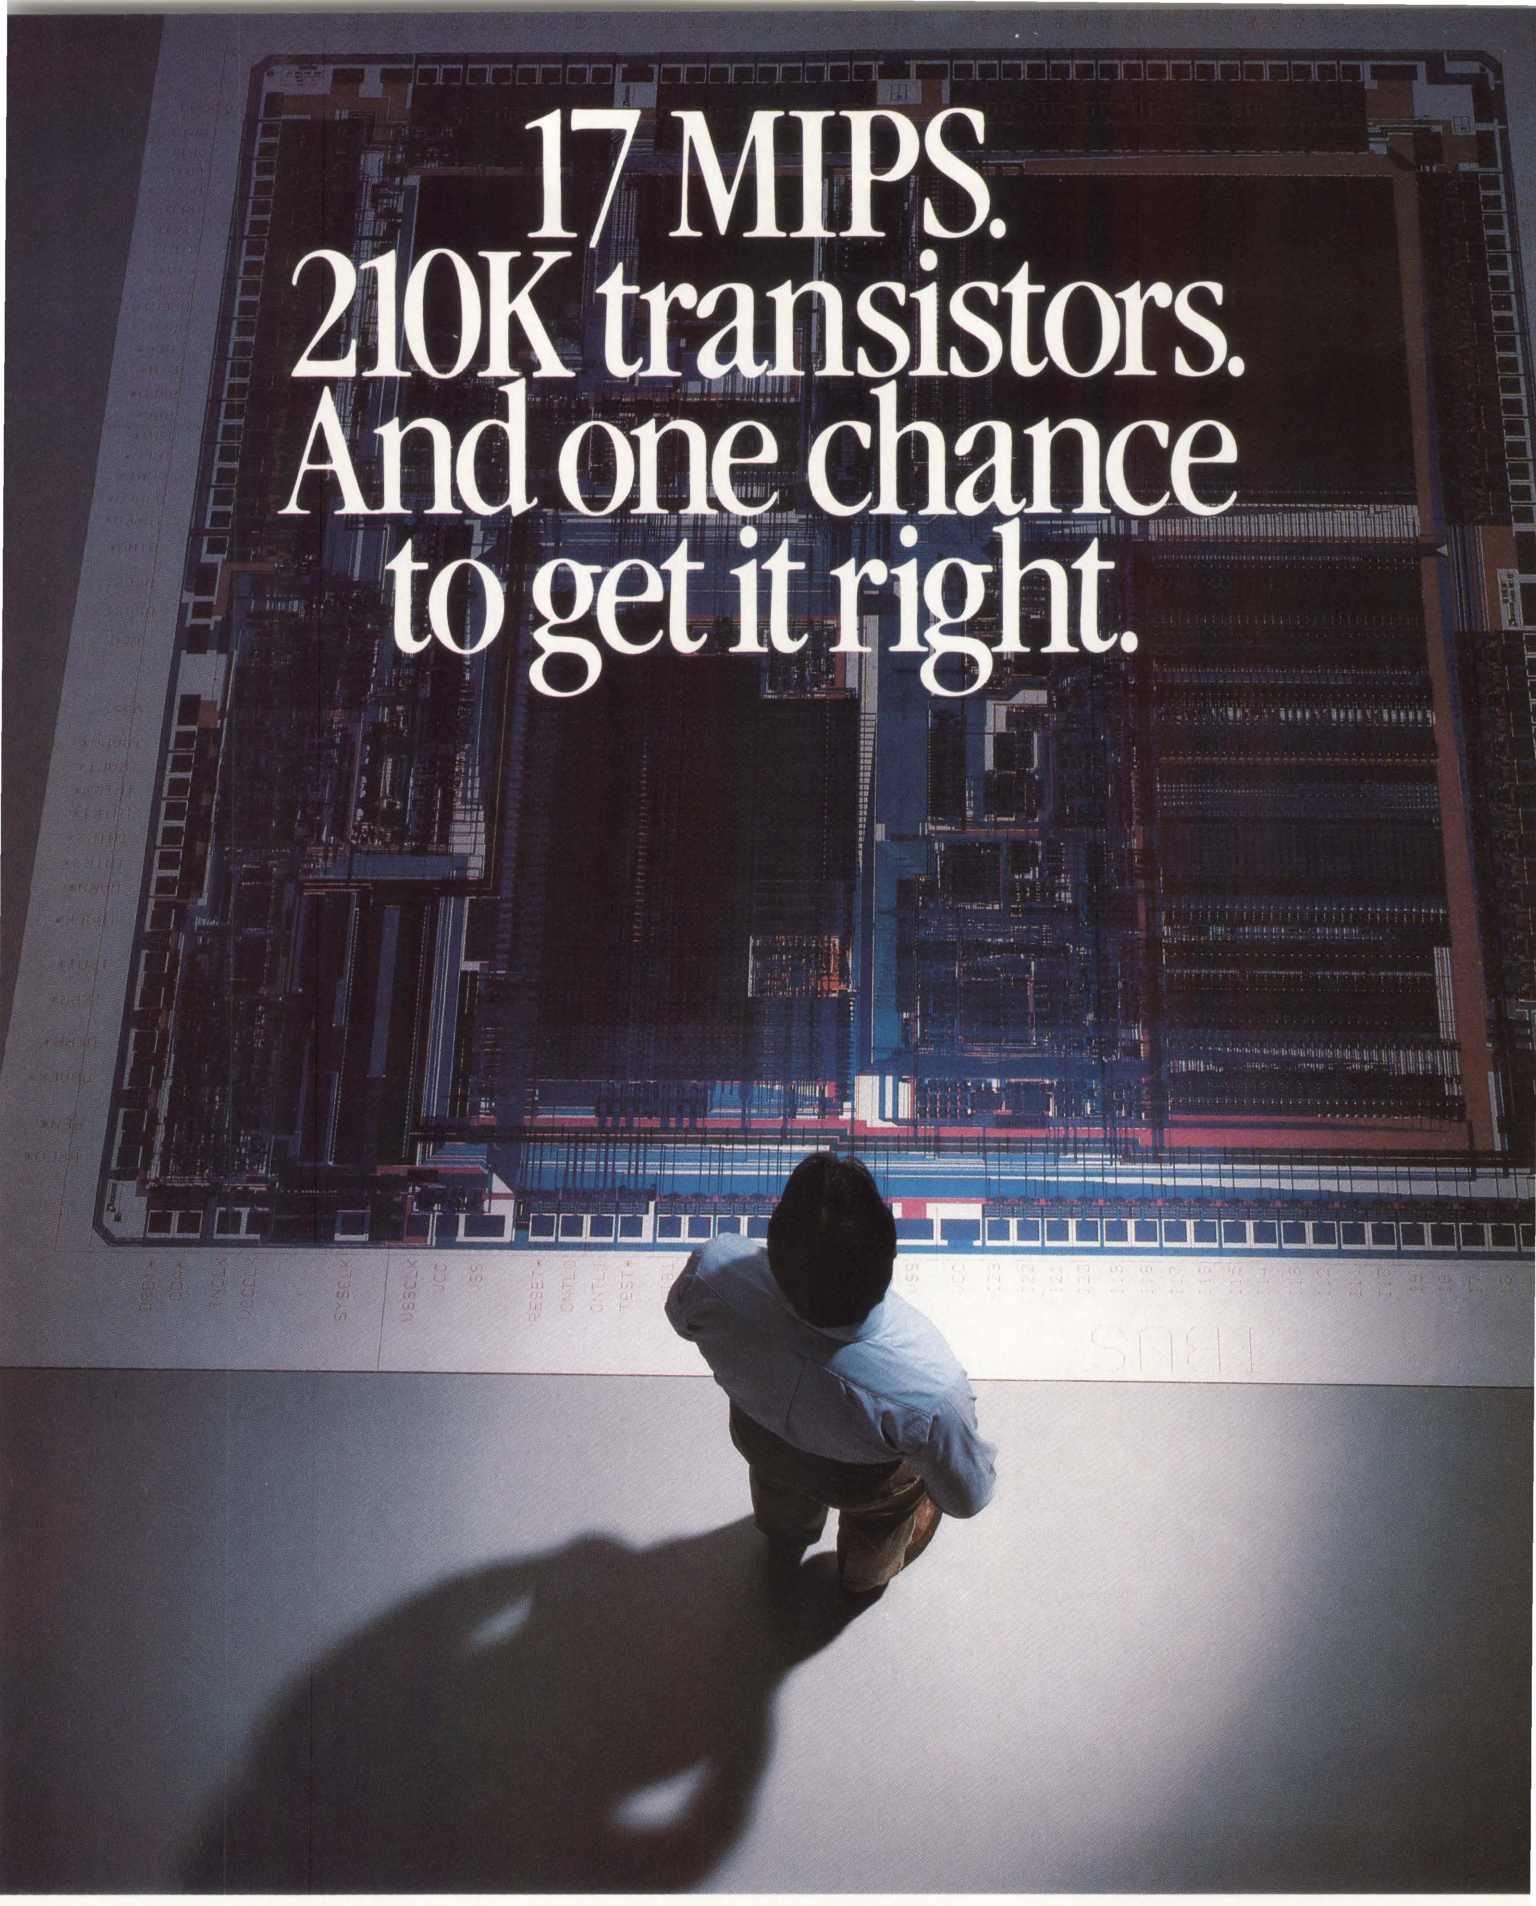
17 MIPS.
210K transistors.
And one chance
to get it right.

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

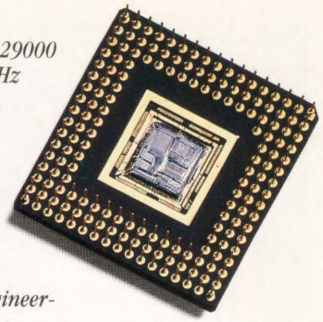

AMD's new RISC-based, 32-bit Am 29000 microprocessor operates at a 25 MHz clock rate with a 40 ns instruction cycle time. It can hit a peak execution rate of 25 MIPS, with a sustained performance of 17 MIPS. Because of its exceptional speed and price-performance ratio, the Am 29000 can be used in a wide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

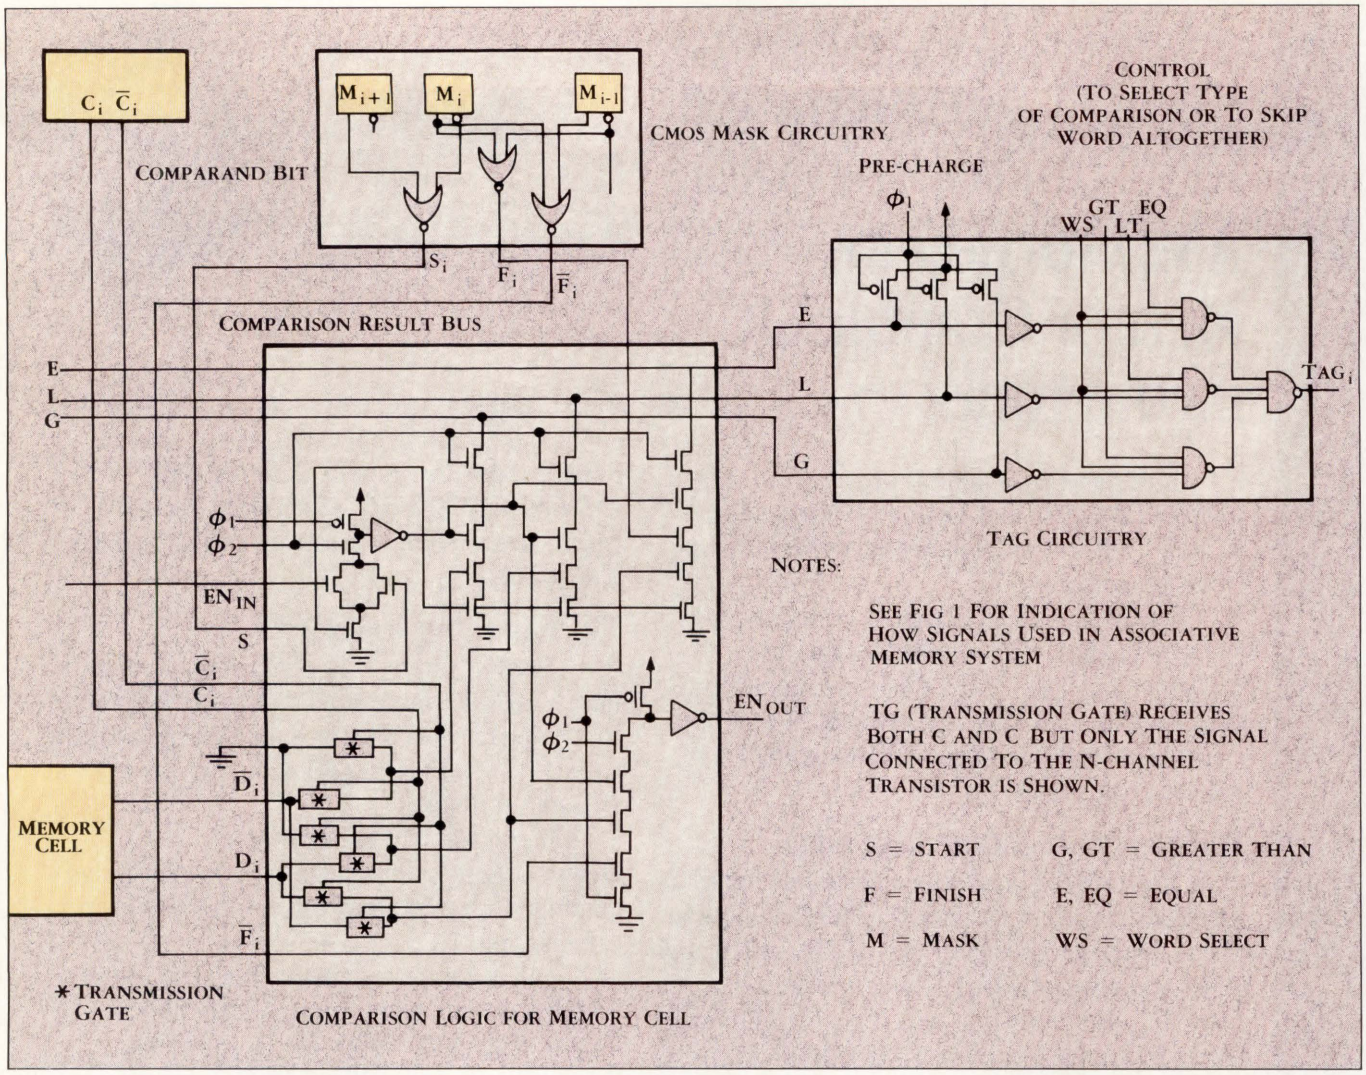

FIGURE 4. Domino CMOS version of fully-parallel logic has common comparison buses for the results and start and finish signals that limit the comparisons to just the attribute fields (instead of to the whole word as in FIGURE 3). It incorporates precharging so that the decision-circuit transistors only have to pull their respective bus lines down. FIGURE 1 showed the total chip architecture for this domino approach.

in terms of the speeds and sizes of the various semiconductor technologies.

■ AJC Chip Enhanced With Precharging

As our initial designs and prototyping (Figures 2 and 3) satisfied our questions about circuit feasibility and algorithm validity, we turned our attention toward speed. In the original design, search time is proportional to the length (L) of the tuple. Therefore, regardless of the length of the search attribute (A), the search time is approximately equal to L × memory cell execution time + tag circuitry execution time.

Figure 4 is a CMOS version of the new design that uses a precharged bus architecture. This architecture minimizes search time because it only searches the length of the search attribute, a much smaller area than the whole tuple. This represents a considerable improvement if the ratio of the attribute length to the overall tuple length is small.

The precharge aspect of the circuit design capitalizes on a precharge phase in which meaningful work is also performed. The results of the comparisons are assessed during the evaluation stage of two-phase clocking.

As can be seen in Figure 4 (and in Figure 1, which relates to Figure 4), this design simultaneously performs its comparison based on equality, greater-than, and less-than by utilizing a separate bus for each. The location of the attribute to be compared is sensed by the mask bit pattern and identified by two control signals, the start bit (S) and the finish bit (F).

During the precharge phase, each word in the memory is simultaneously compared with the comparand-bit pattern. During the evaluation phase, the results of the comparisons are allowed to pull the precharged lines down with the control of a rippled enable signal.

The S signal ensures that the comparison starts at the most significant bit of the search attribute while the F signal ensures that it stops at the end of the attribute. Depending on the type of comparison selected by the greater-than (GT), less-than (LT), and/or equality (EQ) search control lines, a tag bit is set.

Therefore, this design performs the join during every search and is capable of performing multiple-attribute joins by just changing the mask bit pattern. In addition, as in the original fully parallel design (Figure 3), a WS control line is used to enable the comparison of any word in the memory and can be used to perform AND/OR selections.

This design has a more robust distributed control system than the original design, due to the asynchronous operation of the precharged circuitry. Instead of having a single controller in charge of a pool of associated control units, one controller will be assigned to a collection of designated associative memory modules and will control the execution of a join operation.

Continued on Page 84

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

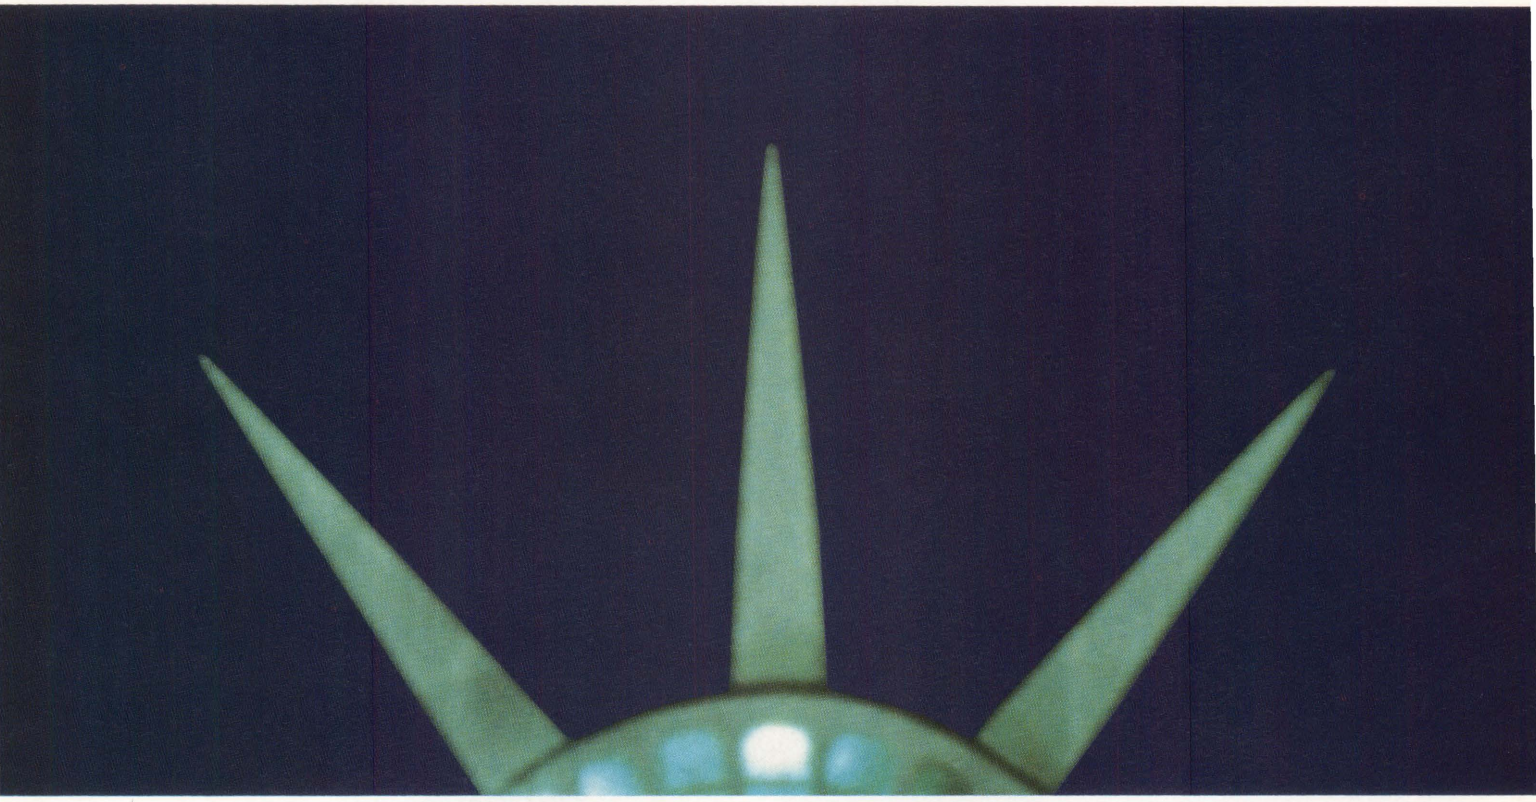

Get the full picture of your ASIC design's performance. Get the IKOS Simulation System. IKOS gives you more simulation information in less time than any other system. Simulate 15,000 gates with 34,000 vectors in 34 seconds. Or 8,500 gates with 5,100 vectors in 4 seconds. Complete with detailed timing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

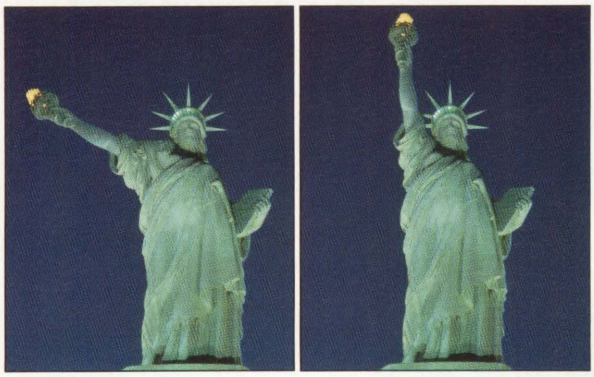

Without the IKOS Simulation System, most ASIC design mistakes show up after production.

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 17

E²/DIGITAL/ANALOG. COMBINATION Y

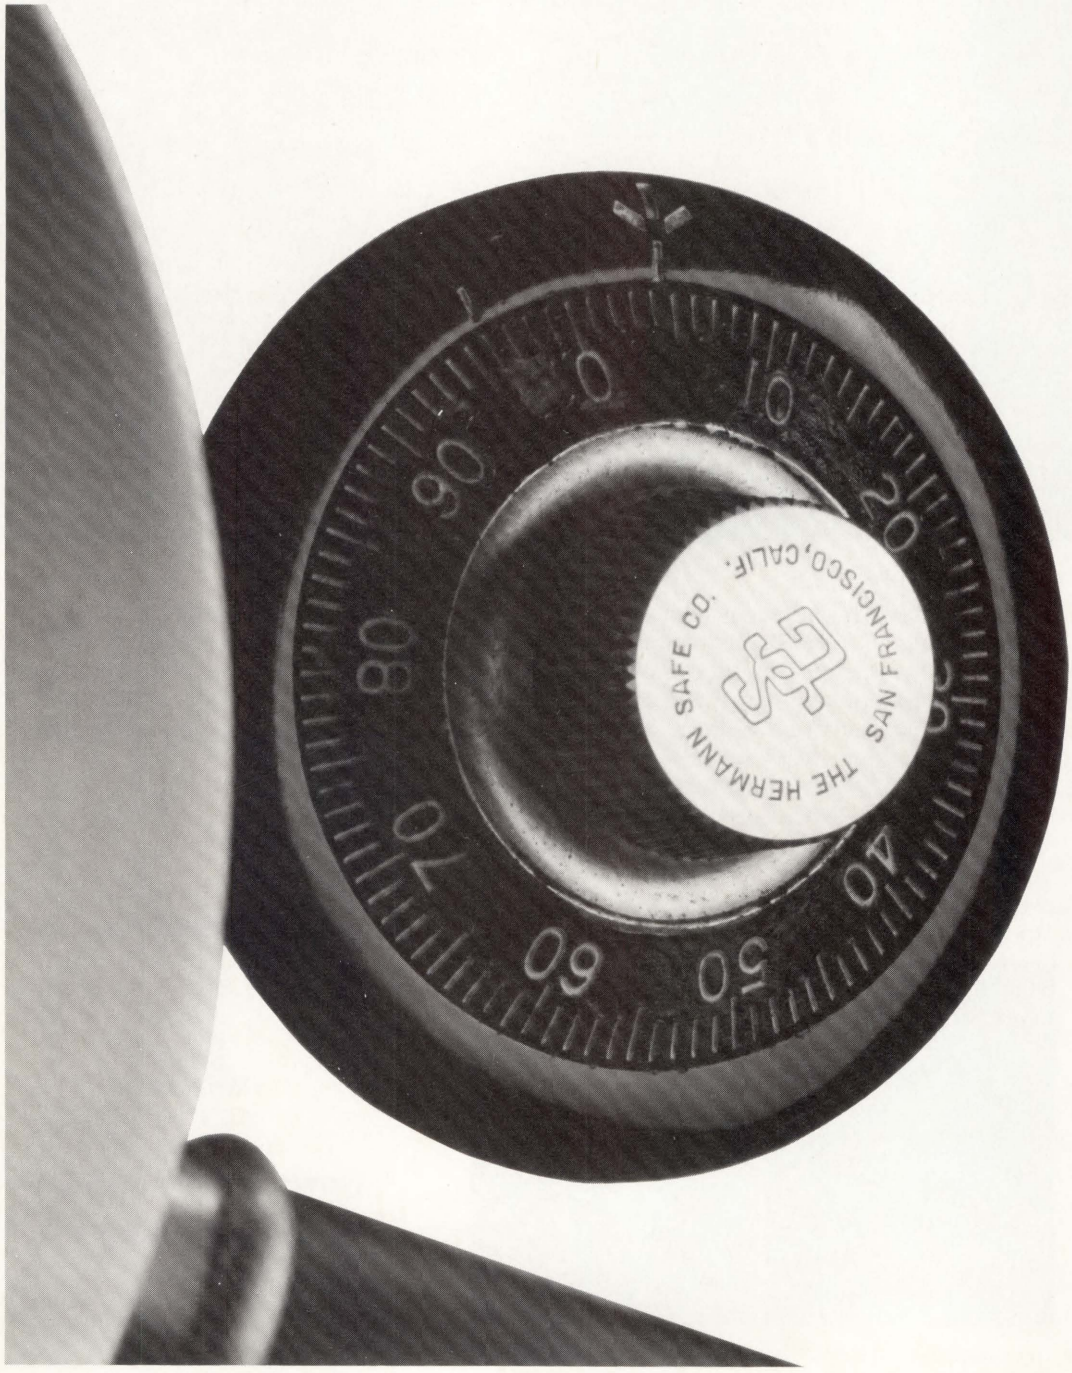

THE TECHNOLOGY OU CAN BANK ON.

To be perfectly honest, we didn't invent the concept that says high integration equals high profit. But as you can see from the application diagram on this page, we definitely perfected it. With our Triple Technology,™ a process that allows you to combine E², digital, and analog functions on the same chip. And, create higher levels of integration than ever before.

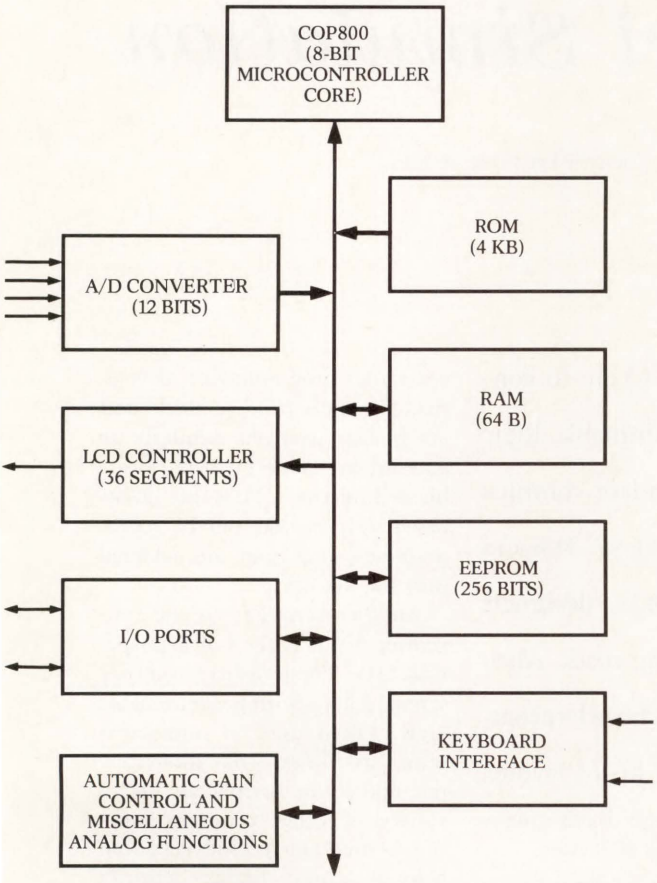

Single chip measurement and control system, integrating several EEPROM, analog, and digital functions.

In this case, our customer's last product was a medical instrument the size of a paperback, with 70 different components. By combining a sophisticated 8-bit controller, RAM, ROM, A/D converter, and 256 bytes of EEPROM on the same chip, we helped them shrink the same instrument to the size of a matchbox. And cut the costs just as dramatically.

As a result, they have a product that sets new standards for the industry. And for their shareholders. And by working closely with their designers, we were able to create this one-chip solution with standard cells from our library. In fact, our customer only had to design about 200 gates of logic using our standard digital cells.

Turnaround time from code to first silicon was only 18 weeks. And because of our development tools and mixed-mode simulation (MIXsim™), the first prototypes worked.

Of course, this is only one example. With 250 digital, 50 analog, and over 20 EEPROM cells in our library, we can create literally thousands of combinations. Including Analog/E², E²/Digital, Analog/Digital, and E²/Analog/Digital. For every application you can imagine. And we can execute them all in high performance CMOS.

So, no matter what you're designing, call or write for our complete library card. And we'll show you a combination you can always bank on. Your ideas and our technology.

Sierra Semiconductor
Triple Technology.™ In CMOS.

CIRCLE NUMBER 18

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

PLDs in Board-Level Simulation

KENT MOFFAT, MENTOR GRAPHICS CORP., BEAVERTON, ORE.

■ SIMULATION IN THE TARGET SYSTEM IS CRITICAL FOR PLDS ■

TODAY'S printed circuit boards designs are likely to contain anywhere from one to over 100 programmable logic devices (PLDs). Designers use PLDs to consolidate complex circuitry that formerly may have required many SSI, MSI and LSI parts. By substituting a single or a few PLDs, designers conserve board space and reduce manufacturing costs. Also, using PLDs can reduce the number of costly board turnarounds because, when errors are detected after board layout, they can be corrected by programming another PLD rather than making changes to the board.

The interaction between PLDs and other devices in the system—such as microprocessors, memory devices and microcontrollers—must be simulated at the board level. This simulation ensures that the optimal PLD architecture has been selected and that the programmed part will function correctly and satisfy all timing constraints when operated in the system. Board-level logic simulation can save development costs and time, especially when compared to the costs incurred by repeated design iteration using some of the more advanced and usually more expensive PLD devices.

■ ARCHITECTURES LEAD TOOLS

Formerly, the term "PLD" could safely be assumed to mean a device that defined a logic function using one or more programmable AND or OR arrays. When designers wanted programmable synchronous parts, registers were added to the array logic. Today, however, the definition of a PLD must be expanded to include a wide range of devices that break with these conventions. To provide greater design flexibility and higher densities, for example, IC manufacturers have developed user-programmable devices that contain both logic blocks and I/O blocks arranged similarly to gate arrays. Some of these "flexible-architecture" PLDs are RAM-based devices that can be reprogrammed after they are soldered into the system.

Another variation on the conventional PLD is the function-specific PLD. These devices combine standard logic with programmable logic. Often used to implement state machines, bus-interfaces, microprocessor peripherals and a variety of other common functions, function-specific PLDs can help to achieve higher densities and speed than devices whose entire function must be designed and programmed.

These general classifications can be used to categorize the more than 100 different PLD architectures that are available today. There are more than two thousand PLDs based on these architectures. They exist in several speed classes and are available from many manufacturers. The increased variety of PLDs now available and their various programming methods pose a significant challenge to system designers who must include

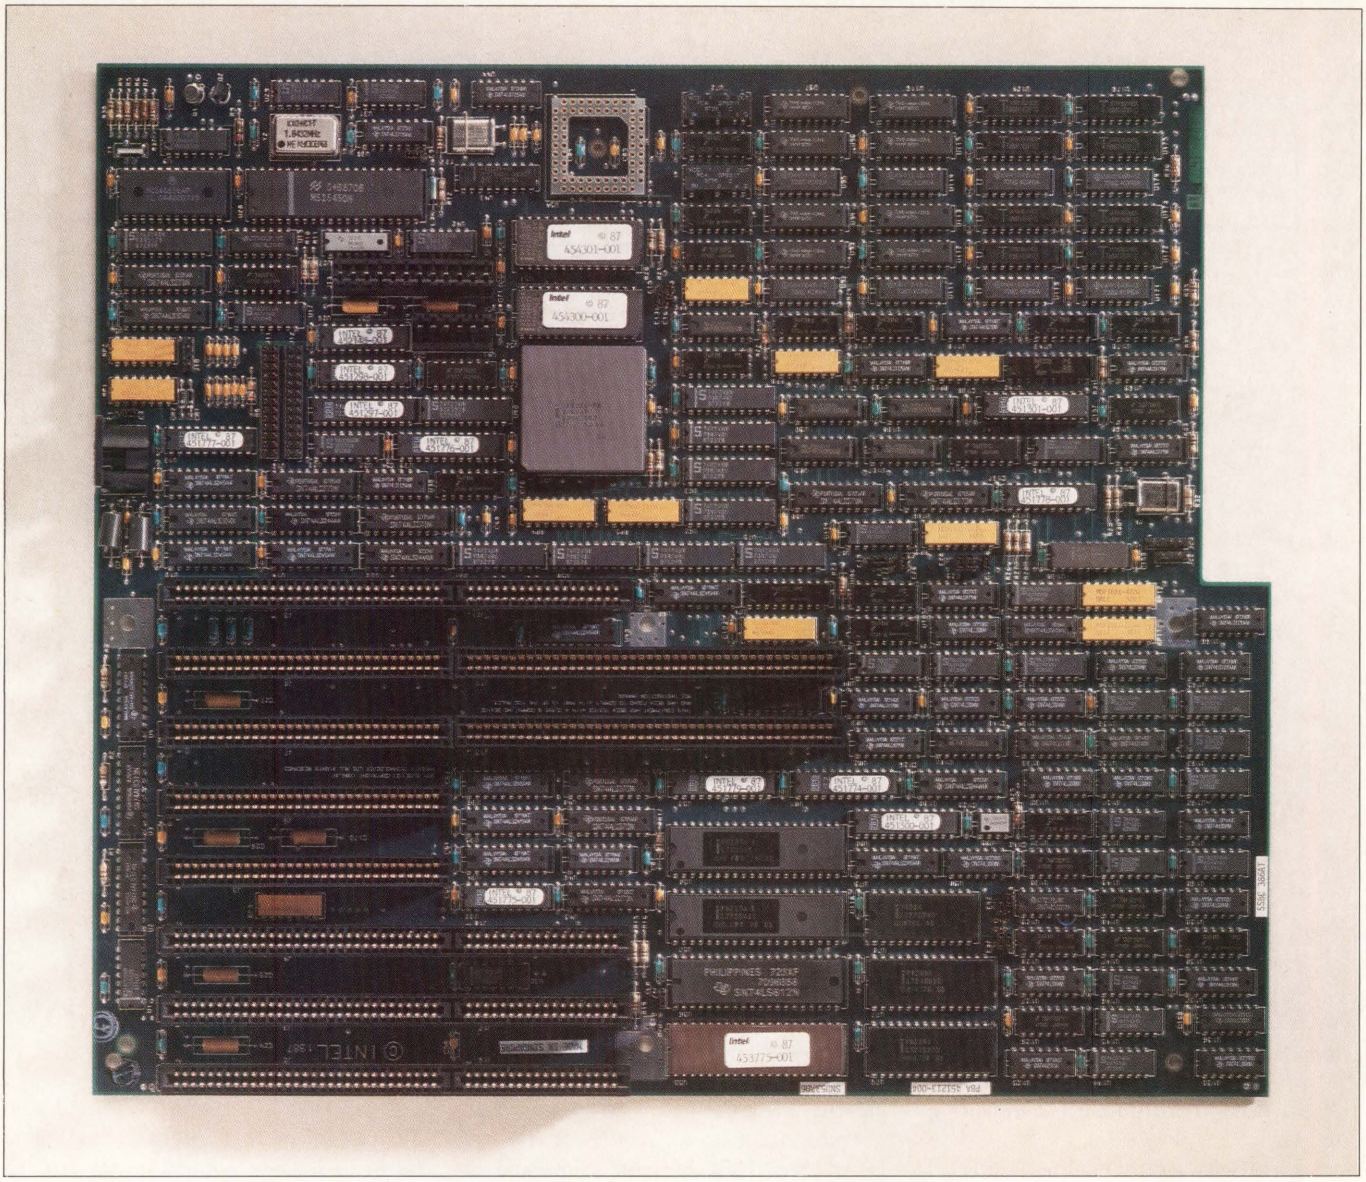

these PLD models in their board-level simulations.

■ Model Categories

These models fall in to two distinct categories: behavioral and structural. The availability of these models depends primarily on the type of PLD tools that support the chosen device.

In the first (and probably still the most common) PLD design method, the designer enters the logic description into a design automation system using a PLD programming tool such as ABEL from Data I/O Corp. (Redmond, Wash.), PALASM from Advanced Micro Devices Inc. (Sunnyvale, Calif.) or PLDesigner from Minc Inc. (Colorado Springs, Colo.). He then directs the tools to create fuse maps from the design which are specified in a standard format defined by the Joint Electron Device Engineering Council (JEDEC).

The newer flexible devices and function-specific architectures depart from conventional programming methods. As a result, they have usually appeared in conjunction with new design tools. Some of these tools rely on proprietary design entry methods that are explicitly adapted to meet the needs of the various architectures. The output from some of these tools is not in JEDEC format.

Data I/O and other programming tool makers now offer tools to functionally simulate the PLD at the device level. Unfortunately, these are not board-level design tools. While they can be used for stand-alone simulations of the device, the models of the PLD designs typically cannot be included in board-level simulations with other parts. Also, proprietary PLD functional simulation tools don't include timing information in their models, and there is no mechanism for adding this data after programming.

■ Behavioral vs Structural Models

To represent PLDs in both device- and board-level simulations, one method used is to merge a JEDEC programming file with a generic behavioral model that represents the unprogrammed device architecture. The generic behavioral model is designed to interpret programming data from a JEDEC file and create an accurate behavioral model of a programmed PLD. This model can then be used in board-level logic simulations, as well as in subsequent analysis steps such as fault

MODELS OF PLDS MUST INCLUDE PROGRAMMING DATA

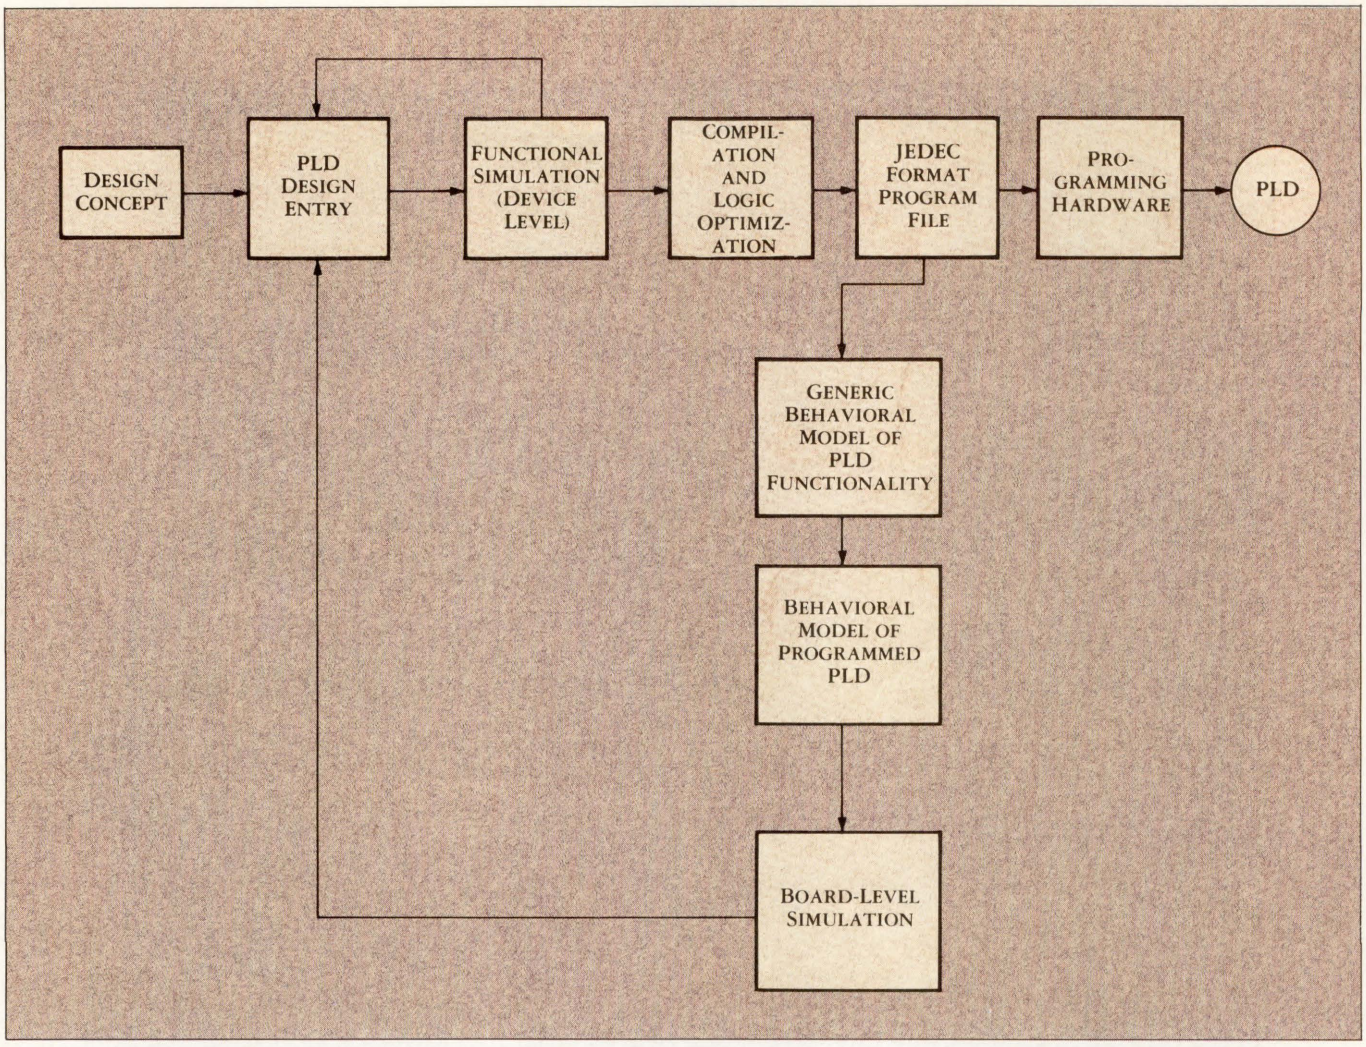

Figure 1. Behavioral PLD models can read the JEDEC-format program file to create a PLD model for board-level simulation.

simulation and critical path analysis (Figure 1).

Such a behavioral modeling strategy is viable as long as the two key components—a programming file in the JEDEC format and the generic behavioral model—are readily available. A wide range of PLD models now exist within electronic design automation (EDA) system libraries, and behavioral models for more complex PLDs can be purchased from modeling firms such as Logic Automation Inc. (Beaverton, Ore.). These commercial PLD models are usually designed for simulating with industry-standard, board-level simulators.

The behavioral modeling strategy is more difficult if either the JEDEC programming file or the generic model do not exist. At the rate new PLDs are being introduced, there will always be devices that haven't been modeled yet. In addition, the JEDEC 3A file format is closely tied to the conventional AND-OR definition of a PLD. Manufacturers of newer PLDs, especially flexible architecture and function-specific PLDs, are devising alternative program file formats that more efficiently reflect their devices' architectures.

Custom behavioral models of PLDs can be written if the necessary time and resources are available. However, PLDs are one of the more difficult parts to model behaviorally because all internal timing structures must be accurately represented. Unlike many standard parts whose behavior can be modeled by simply incorporating pin-to-pin delay values from the device's data sheet, a PLD's internal timing is directly dependant on how the device is programmed. Therefore, the internal delays must be modeled. The task is slightly easier if the device modeled is programmed according to JEDEC standards because the programmer who writes the behavioral model can follow the conventions set by off-the-shelf models. In some cases, however, the modeling burden may be too high a price to pay for the behavioral models and an alternative strategy is required.

■ An Alternative Modeling Strategy

The main alternative strategy utilizes the schematic description of the PLD's design as a model for board-level simulation. To implement this approach, PLD manufacturers include simulation data in the macrocells used to build schematics. The structural models built to define the function of a PLD can then be used to simulate it. In other words, once a PLD schematic model is created, it is used in two separate design streams: one that outputs a programming file, the other that simulates the device (Figure 2).

From the designer's perspective, the structural modeling of a PLD is analogous to modeling the structure of a custom gate array. Design flow for either type of part doesn't vary significantly until fabrication, at which point the PLD is programmed while the gate array is fabricated. Because the design flows are similar, designers like to use the same tools and processes to build both PLDs and gate arrays, as well as to incorporate these devices along with other standard parts in their board designs.

Toward this end, many PLD manufacturers are working with major EDA vendors to make qualified PLD design kits available for industry-standard simulators,

Z I L O G

Z180

Great code compatibility. Terrific performance. Superintegration.™

Zilog's Z80180™ is the CMOS general purpose controller with the high performance and the on-board peripherals that make it clearly the cost-effective, space-saving choice. Whether you're upgrading a Z80 application or designing a totally new system.

Full software compatibility.

You'll be up and running with the Z180 immediately. Because it's 100% object code compatible with Z80/8080.™ You probably already know the code, so you can port right onto the Z180. Not only that, since Zilog originally developed the part jointly with Hitachi, the Z180 is directly compatible with Hitachi's version, the HD64180Z.™

Enhanced performance.

Of course, the Z180's CPU core gives you more power and speed than discrete CPUs. Besides that, there are several new instructions. You also get operating frequencies to 10 MHz. And you have the overall performance advantages of CMOS and Superintegration.™

The important peripherals are on board.

The Z180's high integration results in impressive savings in costs and real estate. The MMU gives you one Mbyte of addressing space. You have 2 DMA channels, 2 UART channels, and 2 16-bit programmable counter-timers. Plus wait-state generators, an interrupt controller, a clock oscillator/generator, and a clocked serial I/O port. All integrated on the Z180 chip.

If this isn't enough to convince you to take a look at what the Z180 can do for your design project, here's a little more to consider. The full complement of development support tools are readily available from industry leaders. And the Z180 comes to you off-the-shelf, backed by Zilog's proven quality and reliability. Find out more about the Z180 or any of Zilog's growing family of Superintegration products. Contact your local Zilog sales office or your authorized distributor today. Zilog, Inc., 210 Hacienda Ave., Campbell, CA 95008, (408) 370-8000.

Right product. Right price. Right away. Zilog

ZILOG SALES OFFICES: CA (408) 370-8120, (714) 838-7800, (818) 707-2160, **CO** (303) 494-2905, **FL** (813) 585-2533, **GA** (404)923-8500, **IL** (312) 517-8080, **MA** (617) 273-4222, **MN** (612) 831-7611, **NJ** (201) 288-5737, **OH** (216) 447-1480, **PA** (215) 653-0230, **TX** (214) 987-9987, **CANADA** Toronto (416) 673-0634, **ENGLAND** Maidenhead (44) (628) 39200, **W. GERMANY** Munich (49) (89) 612-6046, **JAPAN** Tokyo (81) (3) 587-0528, **HONG KONG** Kowloon (852) (3) 723-8979, **TAIWAN** (886) (2) 741-3125, **SINGAPORE** 65-235 7155, **DISTRIBUTORS: U.S.** Anthem Electric, Bell Indus., Hall-Mark Elec., JAN Devices, Inc., Lionex Corp., Schweber Elec., Western Microtech. **CANADA** Future Elec., SEMAD, **LATIN AMERICA** Argentina—Yel.-(1) 46-2211, Brazil—Digibyte (011) 241-3611, Mexico—Semiconductores Profesionales (5) 536-1312.

CIRCLE NUMBER 19

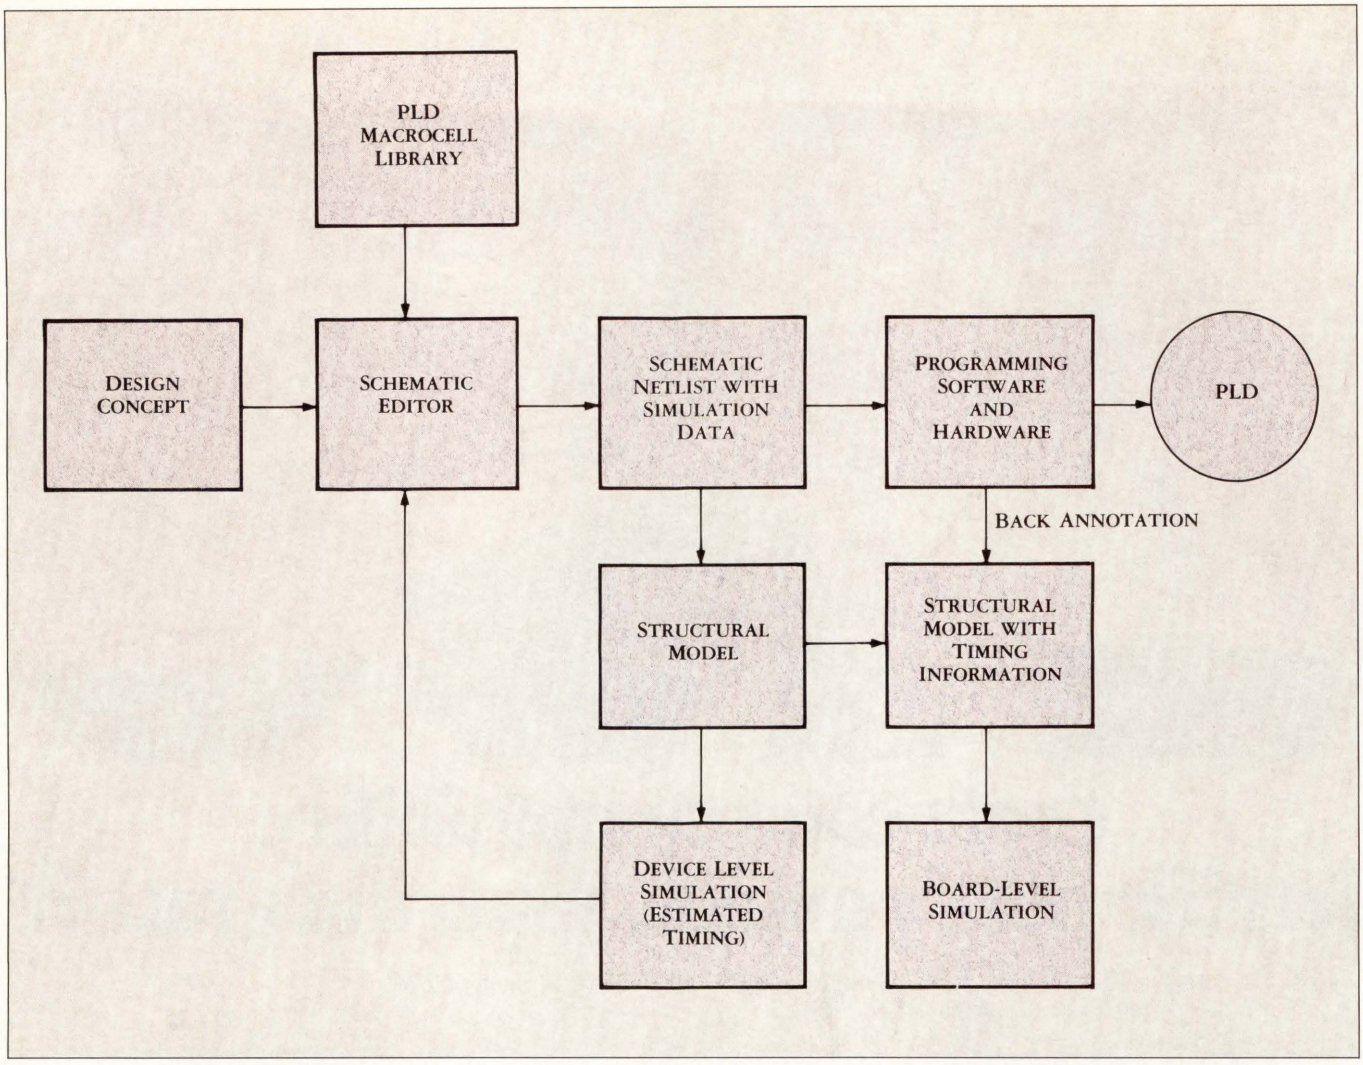

Figure 2. In structural modeling, the netlist that describes the PLD design also serves, when annotated with timing information, as a model for board-level simulation.

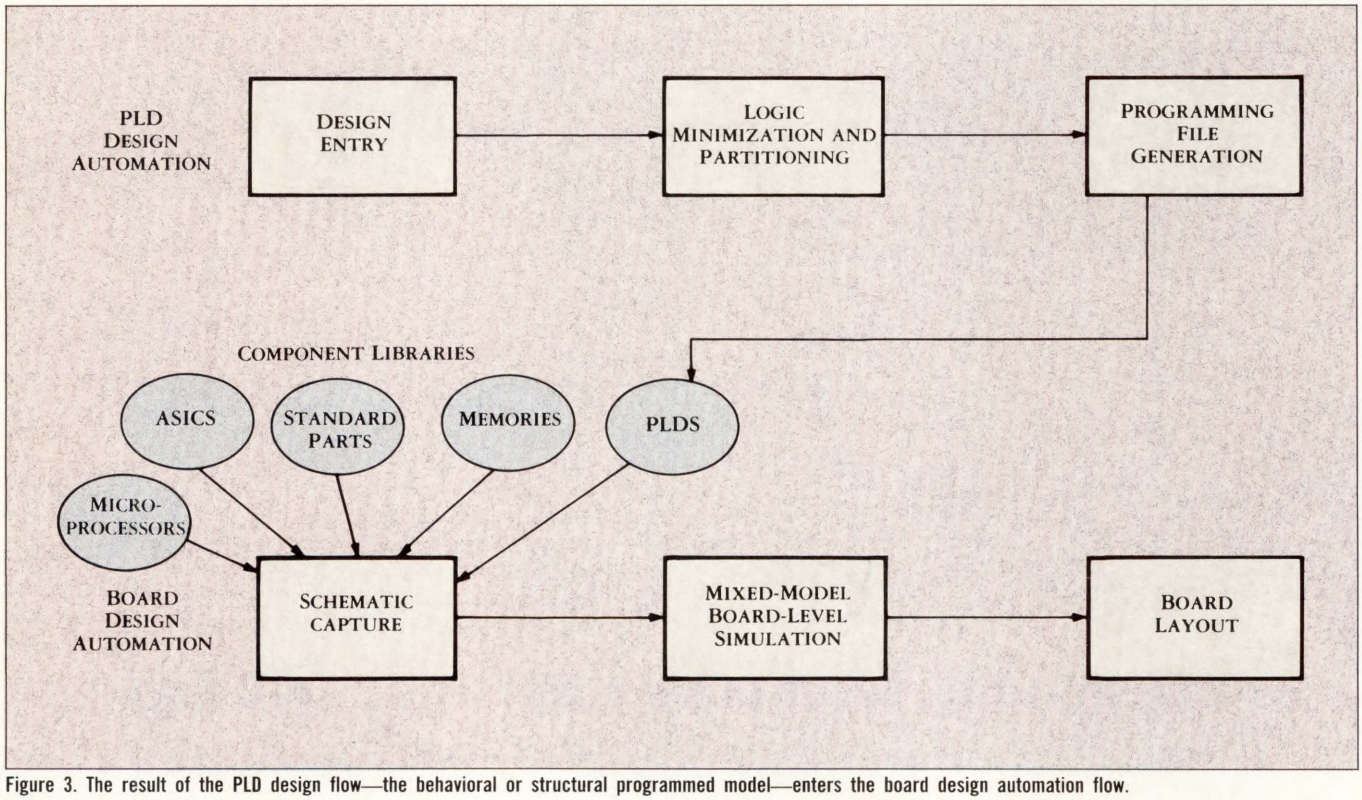

Figure 3. The result of the PLD design flow—the behavioral or structural programmed model—enters the board design automation flow.

much like the ASIC design kits that now exist for custom and semicustom gate array design. These kits include symbols and simulation data for each available macrocell and a "netlister" (software that generates a netlist from a schematic) that translates the completed PLD schematic into a format that can be read by the PLD manufacturer's proprietary programming and layout tools. In this way designers can simulate the PLD in a familiar environment, then output the verified PLD model for programming file creation. Once programming is complete, timing values can be added to the simulation model through the simulator's back-annotatiion capabilities to improve the model's timing accuracy. The updated model can then be used for board-level simulations (Figure 3).

■ TRADE-OFFS

While behavioral modeling of a PLD device is more difficult than structural modeling, there are several advantages to using the higher-level model in a board simulation. For example, behavioral models are more flexible. They represent the PLD's timing and functionality using high-level language constructs and, unlike structural models, they are not tied to the structure of the part. Therefore, they

TIMING ANALYSIS IS CRITICAL TO GAINING INSIGHT INTO HOW A PLD WILL BEHAVE IN A CIRCUIT

can accommodate aids such as error flags that can help debug the system-level simulation at run time (Figure 4).

Another advantage to behavioral modeling of PLDs is that the models can easily accept new programming files. With the structural modeling method, each time a change is made to the PLD programming, the model must be recompiled to create a new structural model that is suitable for board simulation. Then the new timing data must be added through back-annotation, and the board-level simulation must be recompiled, expanded and restarted. In contrast, a behavioral model can usually simply reference the new JEDEC file and reprogram itself without recompilation. In fact, when using a Logic Automation behavioral model with Mentor Graphics' QuickSim simulator, the user doesn't even need to exit the simulator to substitute a programming file.

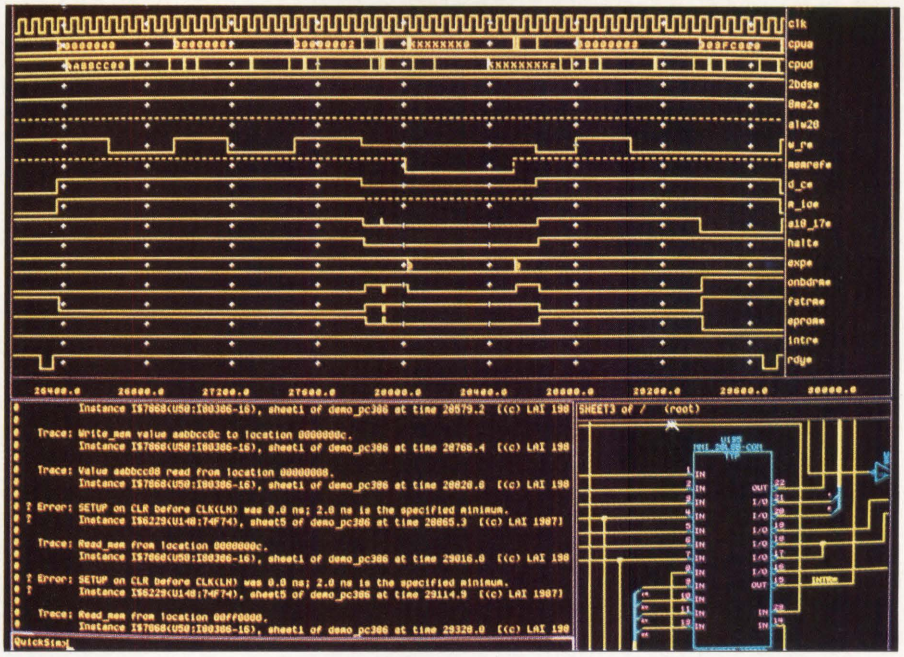

Figure 4. Behavioral models represent the PLD's timing and functionality using high-level language structures, so they can accommodate aids such as error flags that can help debug the system-level simulation at run-time.

The easy substitution of programming files also makes it possible to use the same commercial behavioral model for multiple PLDs in a board design, as long as each instance is based on the same generic architecture and only the programming varies for the individual devices.

To gain any insight into how a PLD will behave in circuit, thorough timing analysis must be performed in addition to logic simulation. Timing can be modeled with either the behavioral and structural modeling methods. The resulting timing accuracy depends on the design software's comprehensiveness.

Behavioral models can incorporate timing information by either using nominal worst-case databook timing values or reading the timing values from a programming file. The latter is typically necessary because the path through internal structures can be changed with programming and, in some cases, nominal values and functions are ambiguous until programming is completed. Off-the-shelf behavioral PLD models, such as those from Logic Automation, usually reference the JEDEC files to acquire programming information of the internal structures before calculating the pin-to-pin timing.

Alternatively, the PLD macrocell libraries used for structural modeling can be implemented with approximated delay values or unit delays, so that in-circuit timing can be estimated before the device is programmed. Then, once programming is complete, the designer can direct the tools to annotate the models with accurate delay data for board-level simulation. For instance, some of the PLD macrocell libraries available for use with Mentor Graphics' QuickSim simulator now provide timing characteristics, and the simulator supports back-annotation.

System designers must also simulate the timing characteristics of interactions between PLDs and other parts on a board. Tools that perform worst-case, min/max timing analysis, such as critical path analyzers (CPAs), are often used for this purpose. CPAs detect timing violations by calculating delays along all paths between clocked elements, and comparing the path delay with clock specifications. Again this type of analysis requires an accurate model of the programmed PLD because basic pin function data—such as whether a pin is a data or a clock pin—is sometimes unknown until after programming.

Fault simulation at the board level also demands that PLD programming data be included in all PLD models. Faults can only be accurately applied and propagated when all internal paths are known. With a PLD this information is not available until the device is programmed. Thus, fault simulation also requires that updated models contain programming data. ■

ABOUT THE AUTHOR

KENT MOFFAT *is the product manager for synthesis tools at Mentor Graphics Corp. He has also served as a technical marketing engineer for both schematic entry and modeling tools. He holds a BA in Physics from Williamette University and a BSEE from Stanford University.*

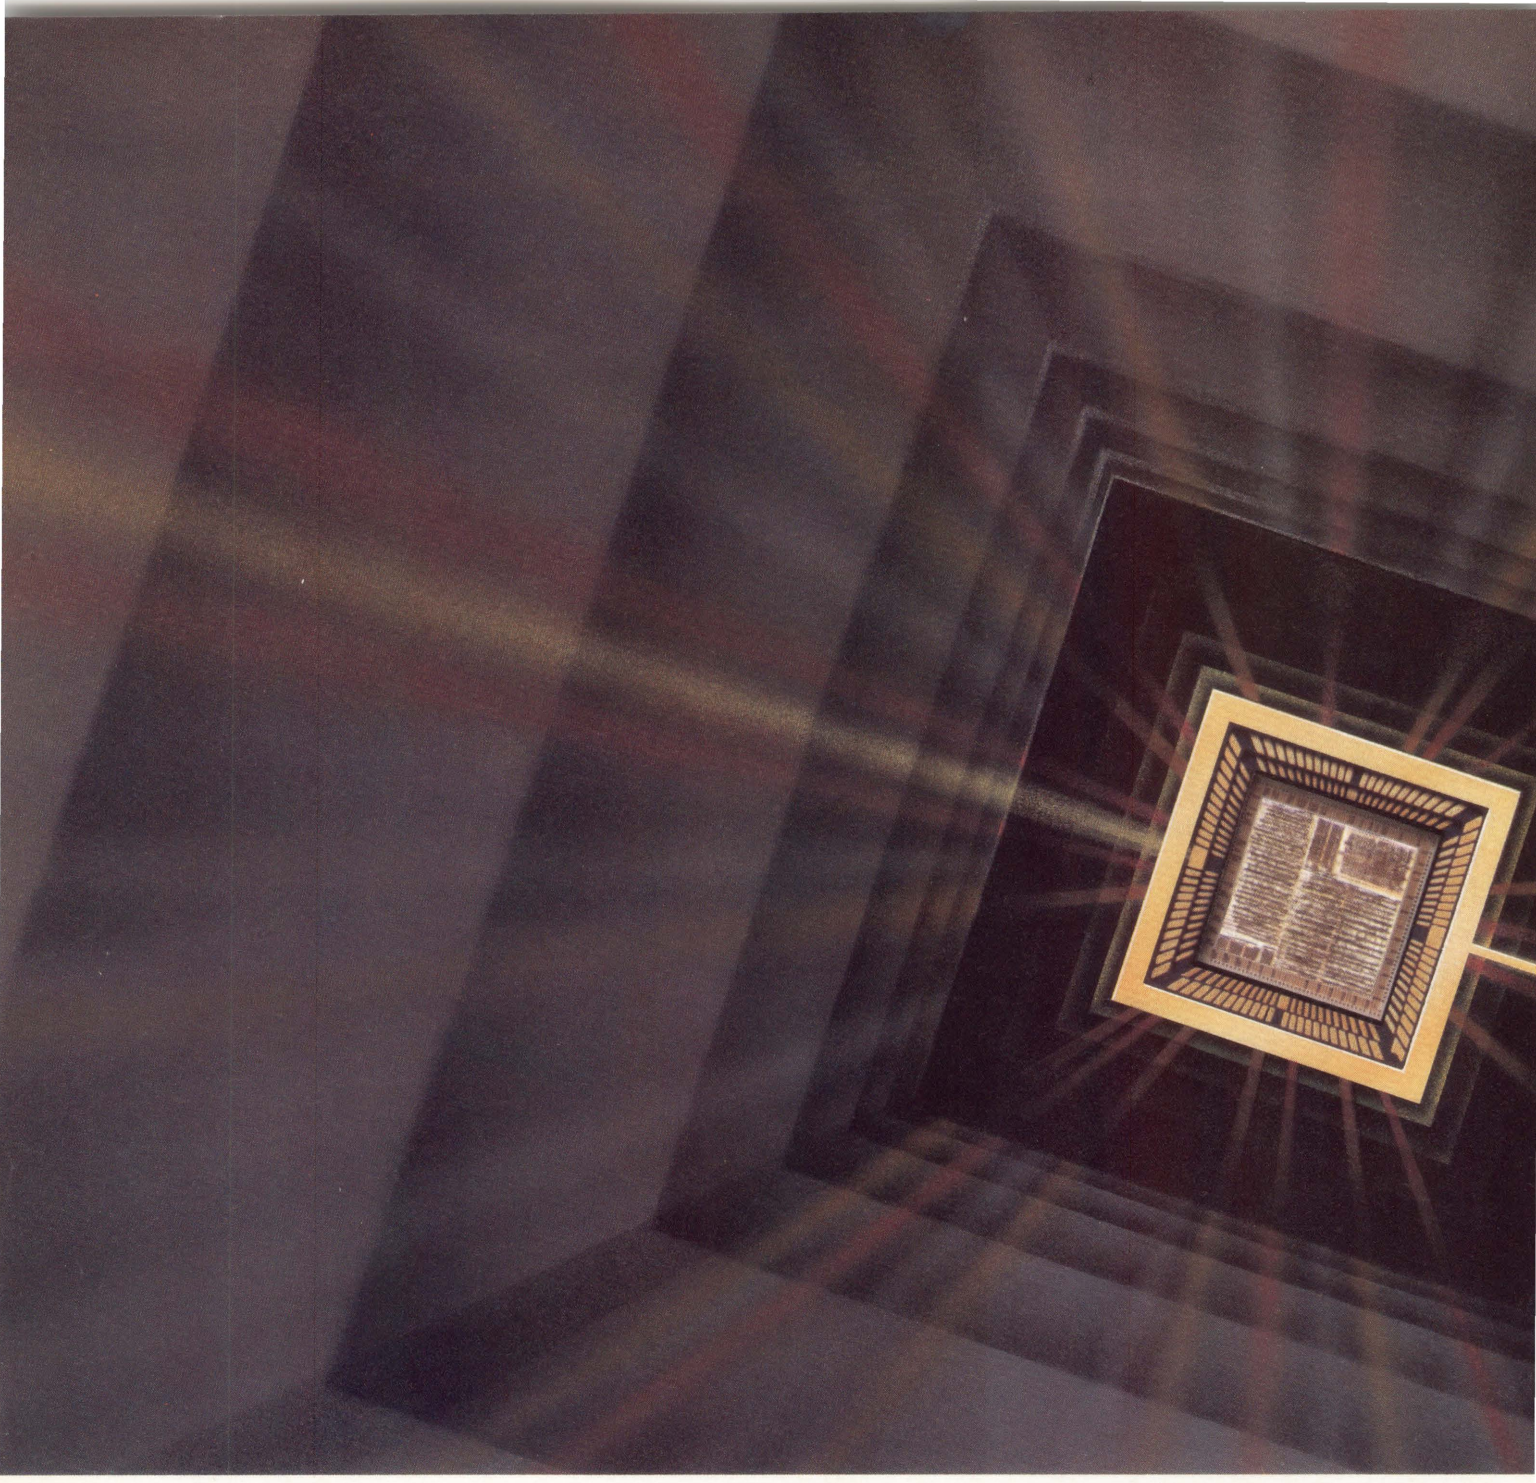

Our 50,000 gate standard cell can reduce your production costs.

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, (503) 629-0818; **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;

Our TC23SC Standard Cell is the most cost-effective solution for high-volume, memory-intensive ASIC applications. Because it uses less silicon. It's as simple as that. Less than an equivalent gate array. And our proven 1.5 micron process provides reliable high density.

Of course, the higher your production volume, the greater your savings. But the TC23SC can even cut costs on low volume applications, because it can simplify your ASIC design task immensely. You don't even have to need all 50,000 gates. The TC23SC Series includes a full range of sizes from 700 to 50,000 equivalent gates. It's simple to choose the optimum standard cell for your design.

Then there's the extensive library lineup. The TC23SC Series library, ADVANCELL™*, is an ultra-high function LSI library. The full product lineup includes 166 primitive cell types, 108 74HC Series macrofunction types, and 67 I/O cell functional types. Macrocells include 2900 Series, CPU peripherals, multiplexers and more. Several RAM and ROM megacells are also available.

BASIC CHARACTERISTICS		
PRODUCT LINEUP	TC22SC	TC23SC
	Macrocells & Macrofunctions	
LIBRARY	74HC Series Compatible Macrofunctions RAM ROM Functional Macros (PLA**, ALU, MPY)	ADVANCELL™* RAM ROM Functional Macros (PLA**, ALU**, MPY, Barrel shifter, adder, FIFO**) 2900 Series Macros*** CPU Peripherals** Analog Macros***
PROCESS TECHNOLOGY	2μm	1.5μm
	HC^2MOS Si-gate double layer metal	
MAXIMUM TOGGLE FREQUENCY	100MHz	150MHz
GATE SPEED (INNER GATE)	1.5ns	1.0ns
SUPPLY VOLTAGE	5V	
GATE COMPLEXITY (MAX.)	10K gates	50K gates
AVAILABILITY	NOW	

* ADVANCELL is a trademark owned by Toshiba or licensed from Siemens or General Electric Co., U.S.A. in certain countries.
** Under development.
*** Development complete, to be available shortly.

The TC23SC Series is fully supported by the VL-CAD system. The system is a CAD design tool developed by Toshiba with full support from logic design to autoplace and route.

The 1.5 micron technology used in the TC23SC system ensures high integration and performance and it provides speeds of 1.0 ns typical gate delay.

To help you even further, we now have a total of five ASIC design centers around the United States for your convenience. For complete information, contact your nearest Toshiba Regional Sales Office: Northwestern: (408) 737-9844, Southwestern: (714) 259-0368, Central: (312) 945-1500, South Central: (214) 480-0470, Eastern: (617) 272-4352, Southeastern: (404) 368-0203.

TOSHIBA

TOSHIBA AMERICA, INC.

1220 Midas Way, Sunnyvale, CA 94086
(408) 733-3223

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.

MAS-88-005

CIRCLE NUMBER 20

SQUEEZING

State Machines into PLDs

DAVID SHARP AND GEORGE BARBEHENN,
HEWLETT-PACKARD CO., SAN DIEGO, CALIF.

Electronics systems are pushed to higher levels of integration by such factors as increasing functionality and manufacturability and decreasing cost and size. Higher integration can be achieved through the use of PALs, EPLDs, and ASICs. PALs are well understood and entail very low risk. They are also relatively expensive, power hungry and offer only a moderate level of integration. ASICs offer high levels of integration and low power consumption but also long lead times, unrefined design tools and higher risk. EPLDs are a compromise. They are reasonably priced, offer moderate to high integration and use familiar, low-risk PAL-type tools. They combine the user programmability of PALs—which makes the design cycle shorter than with custom or semicustom devices—with low power consumption and reprogrammability.

■ PROPER ASSIGNMENT OF SIGNALS IS A KEY FACTOR ■

This article describes some of the challenges encountered when designing large state machines into programmable logic. Many of the ideas presented here, while discussed in the context of EPLDs, are applicable to any state-machine design. The major challenge discussed is the reduction of the state machine's product terms so that it can fit in the EPLD.

For design purposes we had to implement three interlocking state machines (Figure 1), used to control the host interfaces in a computer peripheral, in a single IC. The signals can be categorized into six general groups: clock signals, microprocessor handshaking signals, control register signals, first in/first out (FIFO) memory control signals, parallel-interface control signals and HP-IB-interface control signals.

The number of signals, both input (13) and output (24), was very large. Pin constraints required internal address decoding, shared control signals, and analysis of machine interactions to eliminate unneeded signals. The disparate nature of the signals added to the problem. The microprocessor signals must be fast to avoid degrading system performance, but the standard parallel interface signals are relatively slow, with durations of many microseconds. Many devices were evaluated, and the 5C121 was selected as the best single-chip solution. It was the only PLD with sufficient logic complexity that could also meet performance requirements.

Design steps, such as pin allocation and clock selection, are an iterative process. Apparently insoluble problems may be solved if the premises are re-evaluated and the prescribed methods are applied.

■ CLOCK QUALIFICATION

Our design required three iterations, during which "insoluble" problems were solved. For example, initial designs had far too many states and product terms to fit in the PLD. The use of clock qualifiers and special reduction techniques solved these problems. State diagrams were used to design the state machines. The final design incorporated three Moore-type state machines with a total of 41 states. The output of Moore

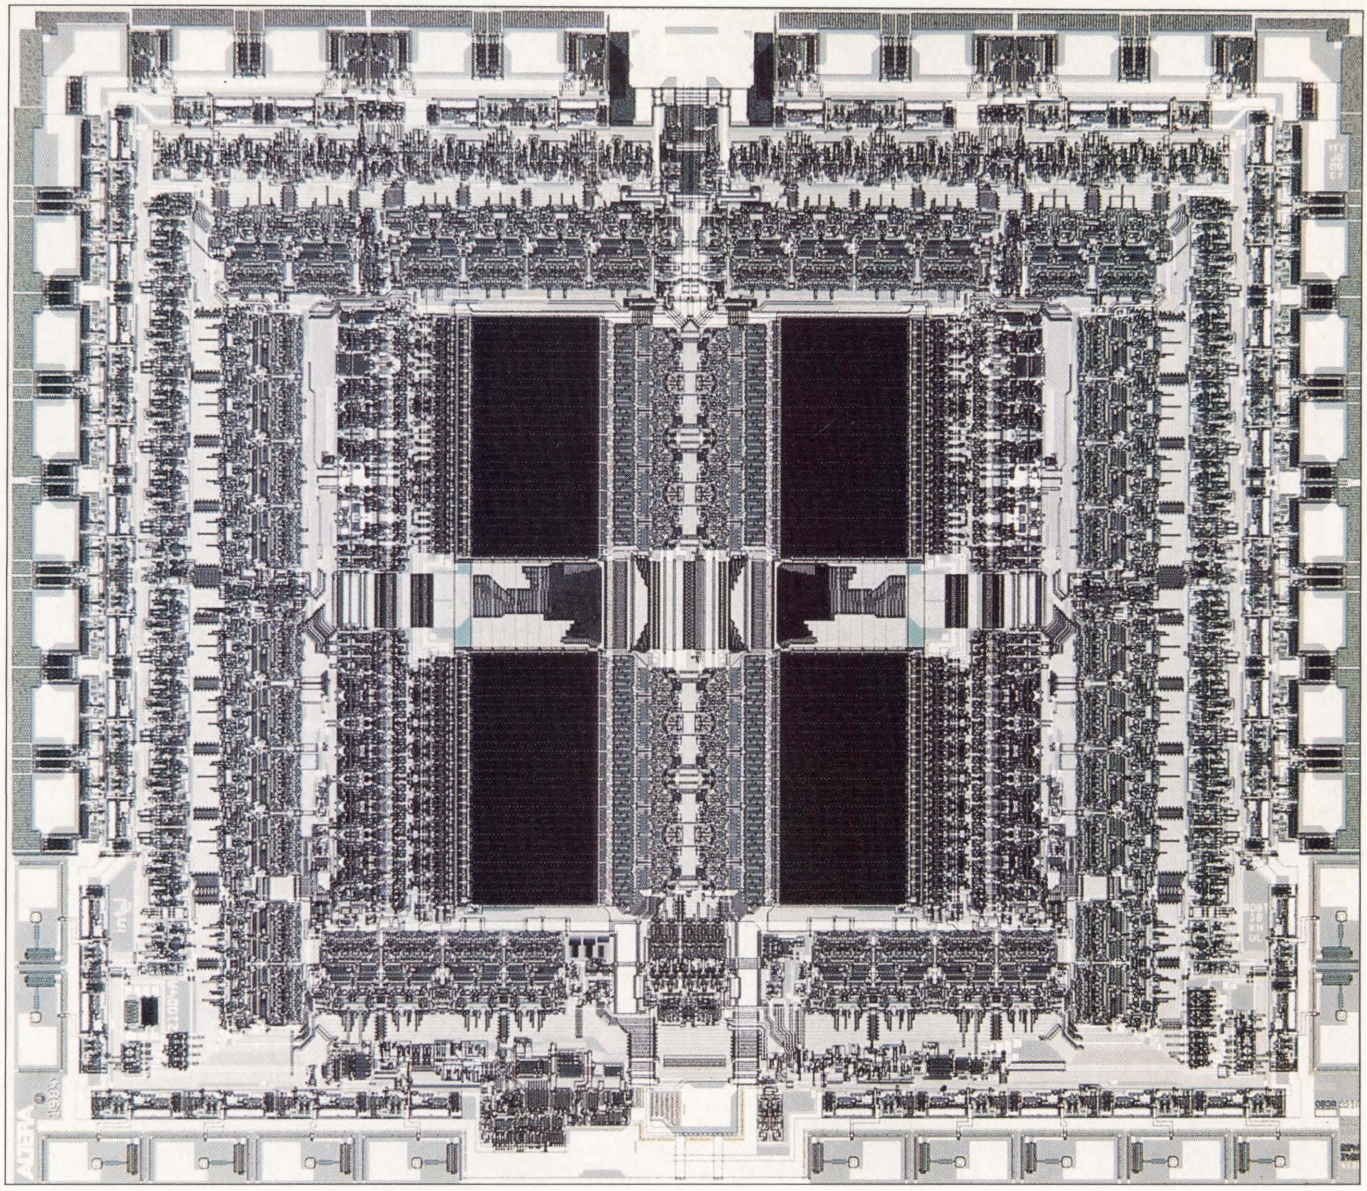

5C121 Photo Courtesy of Intel Corp.

state machines are a function of the present states only and not the current inputs.

After selecting the 5C121 and manipulating the available signals to fit on the device, we found that a practical implementation would require several different clocks. Not only were these extra clock pins not available, but all registers in the 5C121 would be using the same clock input signal. A method was devised by which slower clocks were derived internally from a 16-MHz clock and used to qualify, rather than initiate, the state-machine transitions.

To satisfy timing constraints, a state machine may need to pause in a certain state. One way to implement such a pause is to replace a state with N states whose sequence corresponds to a delay of N-1 clock cycles. This technique rapidly increases the number of states for a given function, and subsequently the required number of product terms. In addition, The number of states necessary for a given time delay also rises linearly with the system clock speed. With this in mind, we built pauses in our current design by using clock qualification of the state transitions.

As an illustration, consider the portion of the state diagram shown in Figure 2. The clock qualifiers in the transition from state A to B force the state to change after the rising edge of the 2-MHz clock. The condition for sequencing to state C is then specified to occur after the 2-MHz clock falling edge. These two elements cause the machine to remain in state B for 250 ns. Note that the transition to B could not be specified as simply "2 MHz" unless it is known that the 2-MHz signal is zero when state A is entered. Thus, the time spent in A is a variable. Depending on the value of the clock signals when A is entered, up to 500 ns may be spent in A. The variable time is accepted in order to achieve a known fixed delay in state B. This method could be applied to successive states to achieve a string of states with delays. Only the first state would have a variable delay.

By properly choosing the transition condition exiting state B, a pause of up to 500 ns can be implemented. The available values are multiples of 62.5 ns, assuming a 16-MHz input clock. The maximum delay is a function of the lowest frequency clock qualifier entering B. For example, if the transition from A to B is "8

LOGIC MINIMIZATION IS ESSENTIAL TO SUCCESS

The RISC that lets you build faster computers faster.

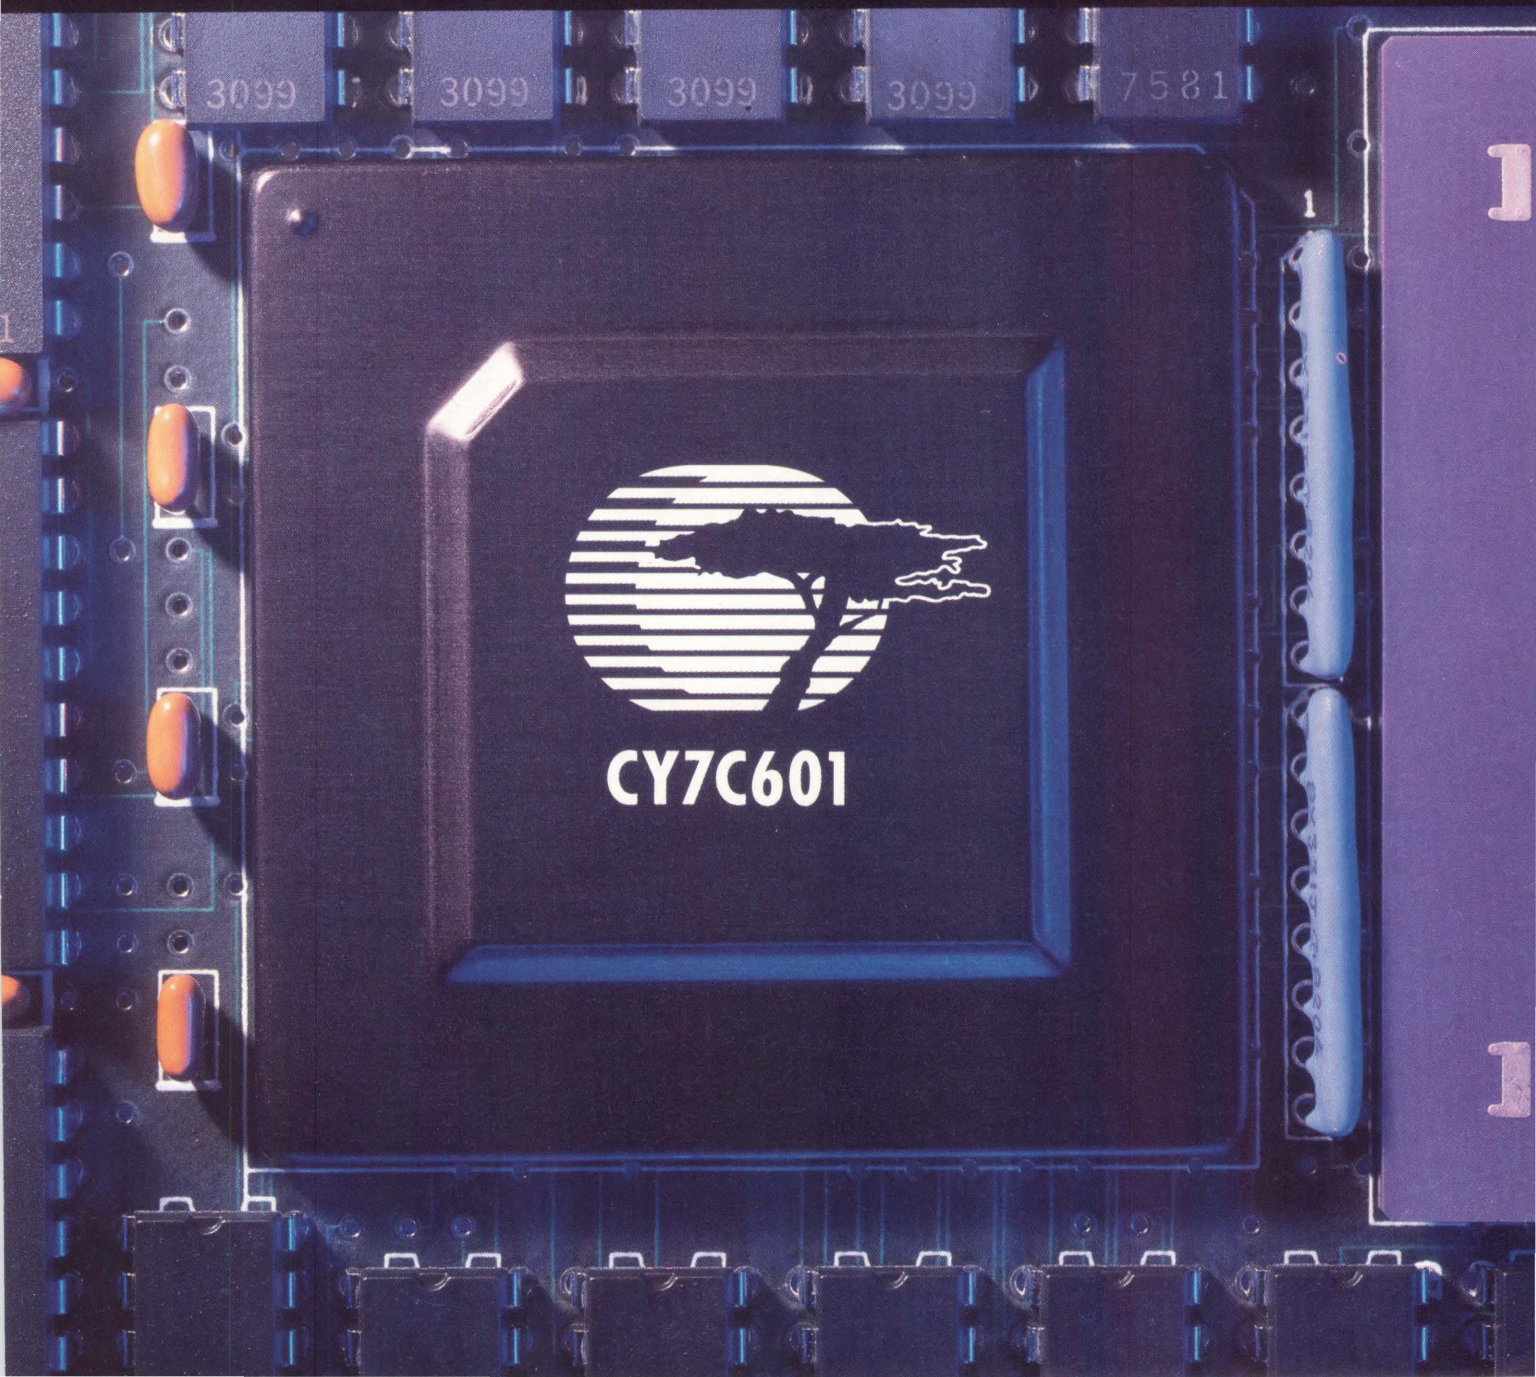

Introducing the RISC 7C600 Family: 20 VAX MIPS, SPARC architecture, and development systems today.

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.*

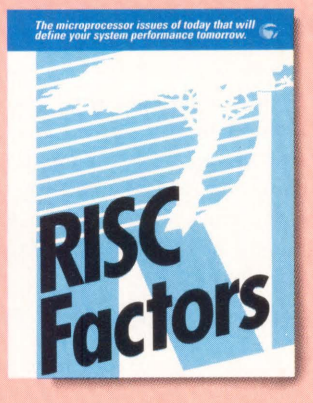

CYPRESS SEMICONDUCTOR

*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.

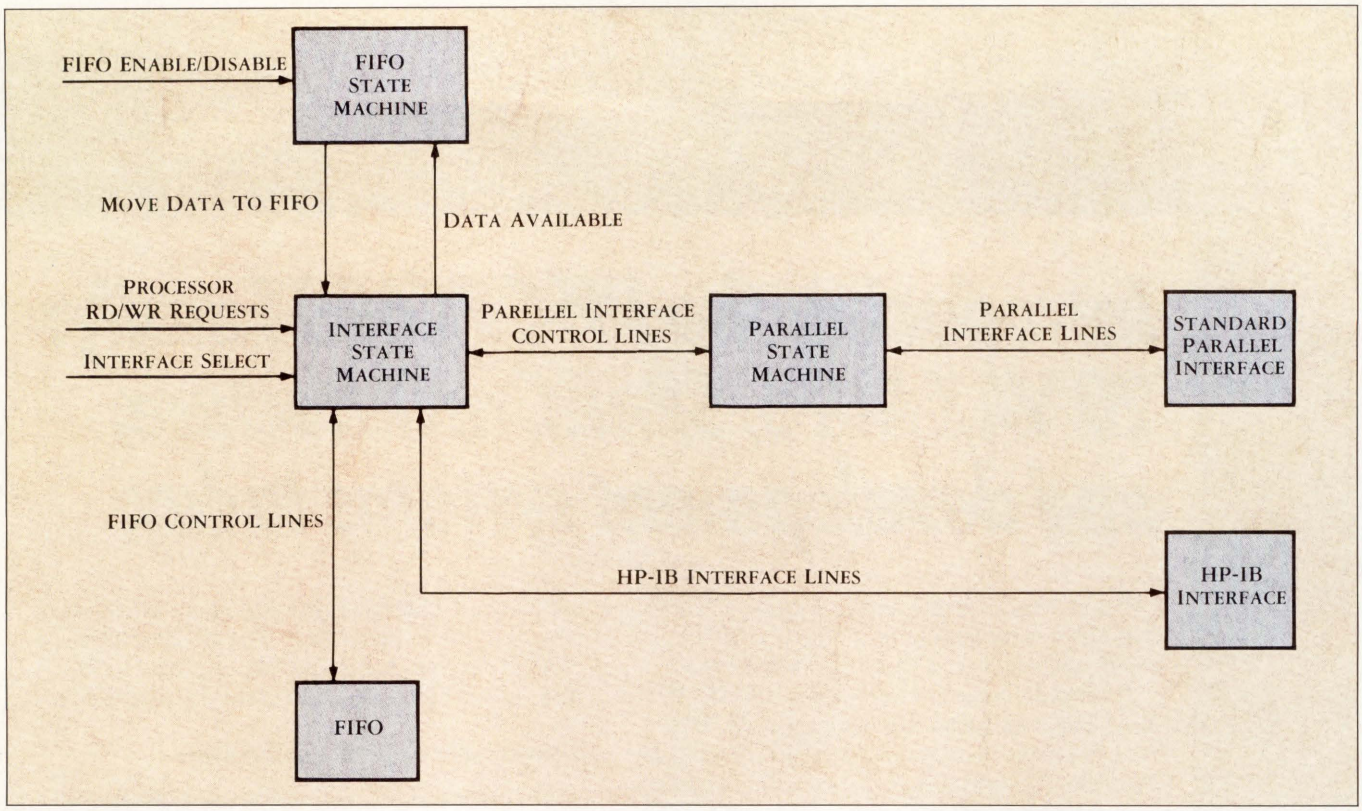

Figure 1. Clock qualification, intelligent state assignment, and product-term reduction enabled these three state machines to fit in a single EPLD.

FLIP-FLOP TYPE	FLIP-FLOP INPUT	ACTIVE HIGH INPUT		ACTIVE LOW INPUT	
		PRESENT OUTPUT	FINAL OUTPUT	PRESENT OUTPUT	FINAL OUTPUT
J-K	J K	0 1	1 0	0 1	0 1
S-R	S R	0 1	1 0	FINAL OUTPUT IS 0 FINAL OUTPUT IS 1	
D	D	FINAL OUTPUT IS 1		FINAL OUTPUT IS 0	
T	T	OUTPUT CHANGES		OUTPUT STAYS THE SAME	

MHz * 4 MHz * 2 MHz * 1 MHz," the largest pause possible would be 1 μs, or one period of the slowest clock. In the present design no extra input pins were available on the 5C121 to bring the clocks onto the chip. Therefore the 8-MHz, 4-MHz, 2-MHz, and 1-MHz clocks were generated internally.

■ Reducing Product Terms

Minimizing the product terms to get complex state machines to fit within a PLD can be a major task. There are many ways to reduce the product terms required for a given state-machine implementation. Software that uses a state diagram description has greatly eased the design process. However, it also can obscure the product-term implications of state-machine manipulations. With D-type flip-flops, specifically, every bit active in a given state, generates a product term for every product term in every transition to the state. Therefore, the designer must understand how changes in the state diagram affect the resulting number of product terms.

It's helpful to note how the relationship between product terms and state diagrams is defined for other types of flip-flops. In the 5C121, as well as many PLDs, registers may use either one or two inputs to control the next output register value (i.e., one for D-type flip-flops and two for J-K flip-flops). Each register actually represents only one bit in the state machine. In general, one product term will be necessary for every flip-flop input that must be active to give the desired next state.

To learn how many product terms will be generated for each register input, bits from the state machine are examined one at a time. The designer must determine the transitions that enter the state and consider each product term in those transitions. If he finds the desired output transition for a given product term (see Table), then a product term will be generated for the state machine.

For example, consider how product terms would be generated for the clock qualification example in Figure 2. Suppose that one specific bit (register) in the state machine takes on values of 1, 0, and 1 for states A, B, and C, respectively. Assume that J-K flip flops are being used with J and K inputs driven by active-high logic. There is one transition condition entering state A —!2MHz + !4MHz + !8MHz (! indicates logic negation). For this transition, the flip-flop output must evolve from 1 to 1. Looking in the Table for the case of a J-K flip-flop with active high inputs, it becomes clear that a product term for the J input will be generated only if the output moves from 0 to 1, and that one is needed for the K input only if it moves from 1 to 0. Since the desired output transition is from 1 to 1, neither input requires a product term.

For state B, one transition condition comes from A, one comes from B, and both have one product term. According to the Table, neither transition adds product terms for input J, but the A to B transition

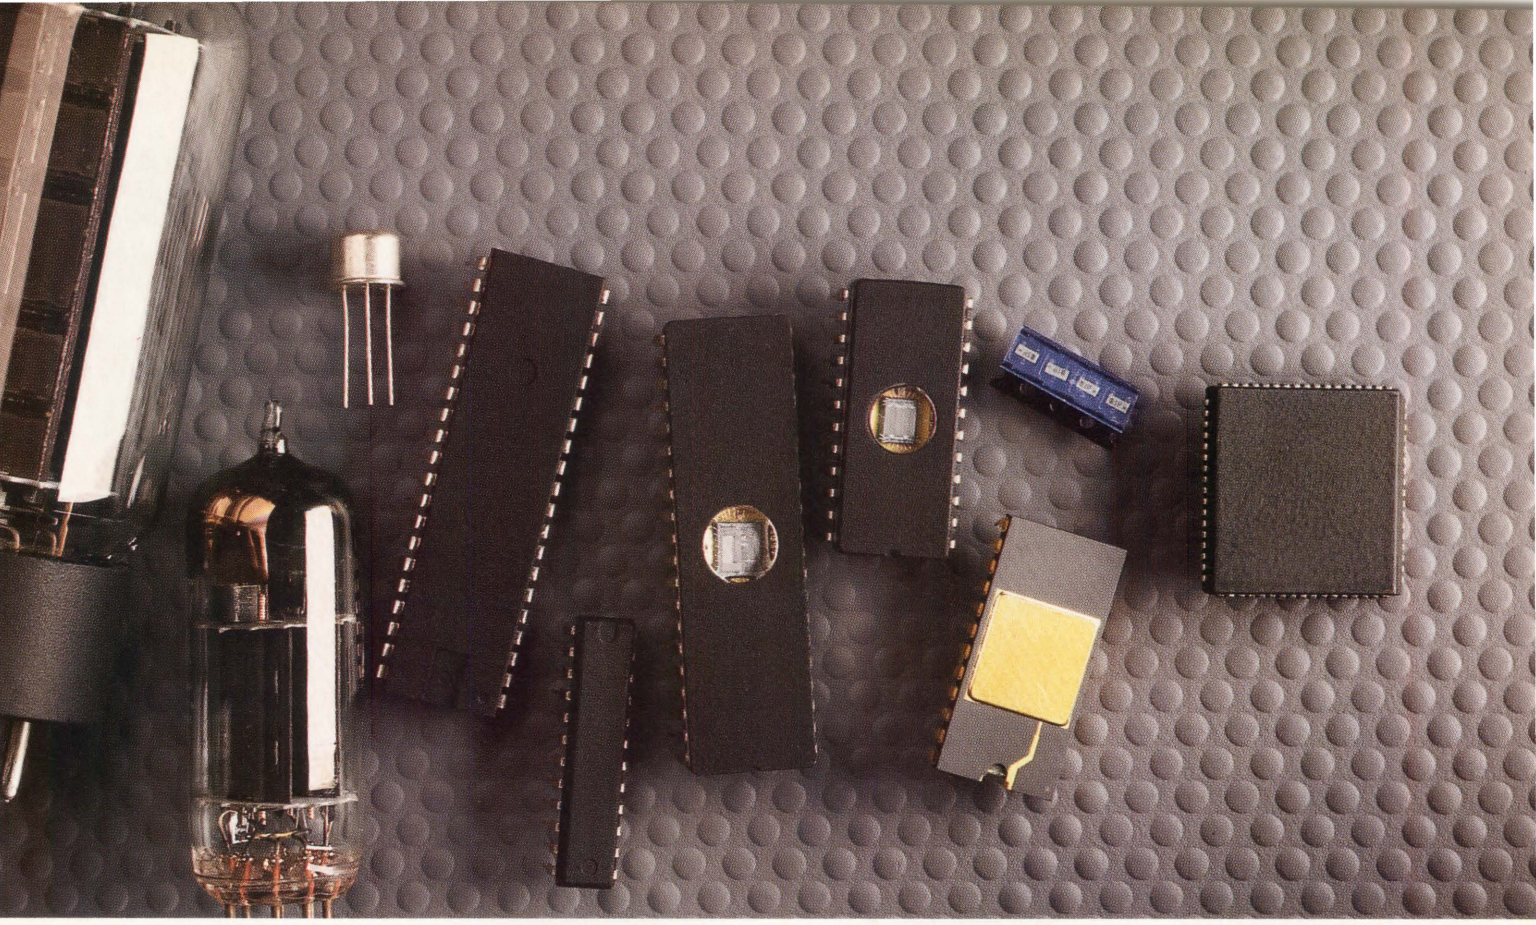

THE UNISITE 40 PROGRAMMER: BECAUSE STATE-OF-THE-ART IS A STATE OF CHANGE.

PROGRAMMING TECHNOLOGY THAT SUPPORTS ADVANCED DESIGNS—TODAY AND TOMORROW. The UniSite™ 40's universal programming technology is the fastest and easiest way to keep up with new devices and packages. Its software-configured pin driver system provides a single site for programming any DIP device up to 40 pins, including PLDs, PROMs, IFLs, FPLAs, EPROMs, EEPROMs and microcontrollers. The same site accommodates the most popular surface-mount packages—PLCCs, LCCs and SOICs.

And now the UniSite 40 is also a gang/set programmer. With the new SetSite™ module, you can program and test as many as eight devices, up to 40 pins each, simultaneously.

INSTANT ACCESS TO NEW DEVICES. The UniSite 40's universal pin driver electronics stores device-specific instructions on a 3½" micro diskette. To update your UniSite 40 with the latest device releases, simply load a new master diskette.

FAST, EASY PROGRAMMING. Menu-oriented operation with step-by-step prompts makes programming simple. Or bypass the menus and zoom directly to specific operations by selecting key commands. Help messages are available whenever you need assistance.

To speed parts selection, the UniSite 40 provides a built-in list of devices. And you can save your most frequently-used programming parameters for instant recall.

DESIGN FREEDOM FOR TOMORROW. When leading-edge designers use the latest devices in their designs, they need the programming freedom only the UniSite 40 provides. Call Data I/O® today and ask about the UniSite 40. Because state-of-the-art never stops changing.

1-800-247-5700
Ext. 648

Data I/O Corporation 10525 Willows Road N.E., P.O. Box 97046, Redmond, WA 98073-9746, U.S.A. (206) 867-6899/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

DATA I/O Corporation

CIRCLE NUMBER 25

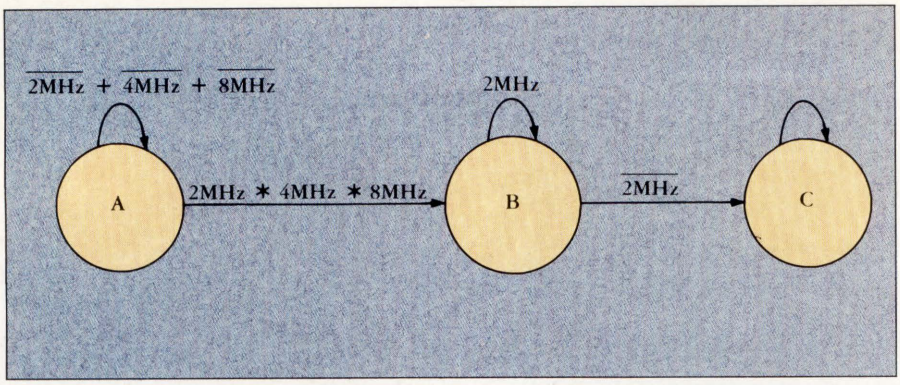

Figure 2. The clock qualifiers in the transition from state A to B force the state to change after the rising edge of the 2-MHz clock. The condition to sequence to state C is then specified as being after the 2-MHz clock falling edge, causing the machine to remain in state B for 250 ns.

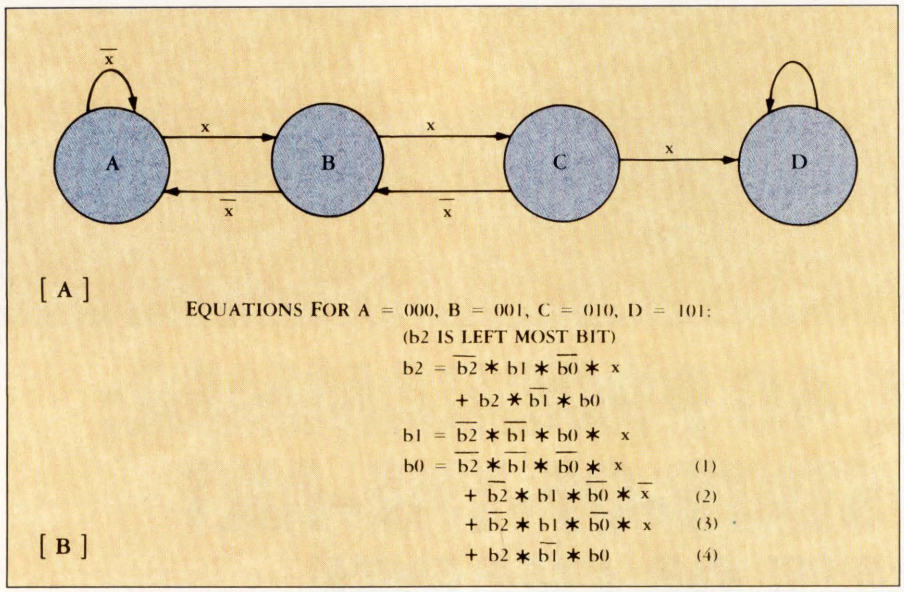

Figure 3. To assign states to this state machine (a), the goal is to keep adjacent states similar and to use the minimum number of active bits for the states. With the given state assignment for A, B, and C (b), the alternatives for D—101, 011, or 100—will all yield the same number of product terms after reduction.

matches the entry for the K input causing the output to change from 0 to 1. One product term, therefore, results for input K due to state B.

Finally, the single transition condition entering state C (from state B) adds one product term for the J flip-flop input and none for the K input. It can be concluded that the J and K flip-flop inputs for this example would each require one product term. Since this approach gives the number of product terms generated from the state diagram before reduction, the final design will usually require less terms.

■ State Assignment

The product-term generation rules for different types of flip-flops show that product terms are a direct result of active bits in the states. To minimize the number of product terms that enter the reduction phase of the design, choose states that have the least number of active bits.

The designer should reduce product terms as much as possible for each single register by selecting states similar to the state preceding it. The reduction can be eased as follows: Given state A , which, depending on the input, goes to one of the states 1, 2, . . ., N, choose the states 1, 2, . . ., N to be as similar to state A as possible. This principle causes many of the factors in the product terms required for the state transitions (such as A to 1, A to 2, . . ., and A to N) to be identical, thereby increasing the possibility of minimization. Note that one of the states 1, 2, . . ., N, may be identical to A, causing the circuit to remain in its current state.

These two guidelines are illustrated by an example. Suppose bit values need to be assigned to the states in Figure 3a. After being reset to state A, this machine detects that input samples have been mostly 1. There are no restrictions on the assignments and D-type flip-flops are used.

If A is the current state, the next state will be A or B, depending on the input X. Therefore, A and B are intentionally chosen to be as similar as possible. If B is the current state, then the next state will be either A or C. Likewise, A and C should be similar. A similar situation holds for B and D. The goal is to satisfy these similarities using a minimum number of active bits for the states.

By assigning A = 000, B = 001, and C = 010, both objectives are accomplished. But D cannot simultaneously satisfy the conditions of the minimum number of active bits, *and* differ from B in only one bit position. The alternatives—101, 011, or 100—will all yield the same number of product terms after reduction.

The unreduced equations if D is 101 are shown in Figure 3b. Examining the terms for the right-most bit b0, term (1) comes from the transition from A to B; (2) from C to B; (3) from C to D; and (4) from D to D. Since both B and D are chosen to include the assignment b0 = 1, terms (2) and (3) would be combined into one in the reduction phase of the design.

The clock qualification (Figure 2) provides another example. The qualifier for the transition from A to B can be chosen as anything from 8MHz to 8MHz * 4MHz * 2MHz * 1MHz—and the qualifier choice doesn't add any more product terms than the unqualified transition. The difference occurs when staying in A—the complement of the condition to go to B. Therefore, the condition to stay in A will range from !8MHz to !8MHz + !4MHz + !2MHz + !1MHz, which contains from one to four product terms. If state A is assigned to have all bits inactive, then any conditions that cause a transition to A will add no extra product terms. The extra terms in the condition to remain in A are therefore not a disadvantage. To minimize the terms generated as a result of the qualification exiting B, the designer should assign the qualification a single clock term. Otherwise more than one product term will be required for the condition to stay in B.

For example, consider two ways to achieve a 250 ns delay. One, is to qualify the first transition with 8MHz * 4MHz * 2MHz and the second with !2MHz. The other is to specify both conditions as 8MHz * 4MHz. If the originating state can be chosen as having fewer active bits (ideally none) than B, then the first scenario will generate fewer additional product terms. If A has only inactive bits, then this scenario generates only one additional product term for the pause. Whenever possible in our state-machine design, A was chosen as the inactive state, and the second qualifier was selected as a single clock.

If a state assignment cannot be found to sufficiently reduce the number of product

META-SOFTWARE

Has Taken A Bold, New Step...

From design to silicon, Meta-Software provides the best circuit simulation tools in the industry. Meta has recently repositioned its product offering to better meet the demanding needs of today's design engineers.

For convenient, one-stop shopping, the innovators at Meta-Software now offer:

HSPICE: The industry's leading analog circuit simulator for integrated and discrete circuit design. HSPICE includes a multi-target optimizer supporting all SPICE and HSPICE models.

RADSPICE: HSPICE plus radiation effects modeling provided by SAIC. Effects include total dose, ionizing photocurrents and neutron radiation.

HSPLOT: Meta's high-resolution interactive graphics post- processor for HSPICE and RADSPICE. HSPLOT provides graphic terminal and hardcopy support for a wide range of display services.

Discrete Device Library (DDL): Includes more than 750 models of discrete components for use with HSPICE. Included are BJT, MOSFET, HEXFET, Diode, JFET, Op Amp. Comparator, A/D converter, D/A converter, Timer and SCR models.

ATEM: Meta's lab test equipment interface program which creates measured data files and initial guesses for optimization features of HSPICE. ATEM provides an easy method for scanning transistor characteristics and selecting devices for full optimization.

MetaTestchip™: A test chip tailored to customer's design rules, providing all structures necessary for complete, automated process and device characterization.

Lab Services: Products and services for discrete device and wafer level characterization.

Circuit PathFinder (CPF): A path timing analysis program, providing full chip analysis at interactive speed. Circuit Rule Checker module locates slow nodes and gates, and a variety of circuit-configuration rule violations.

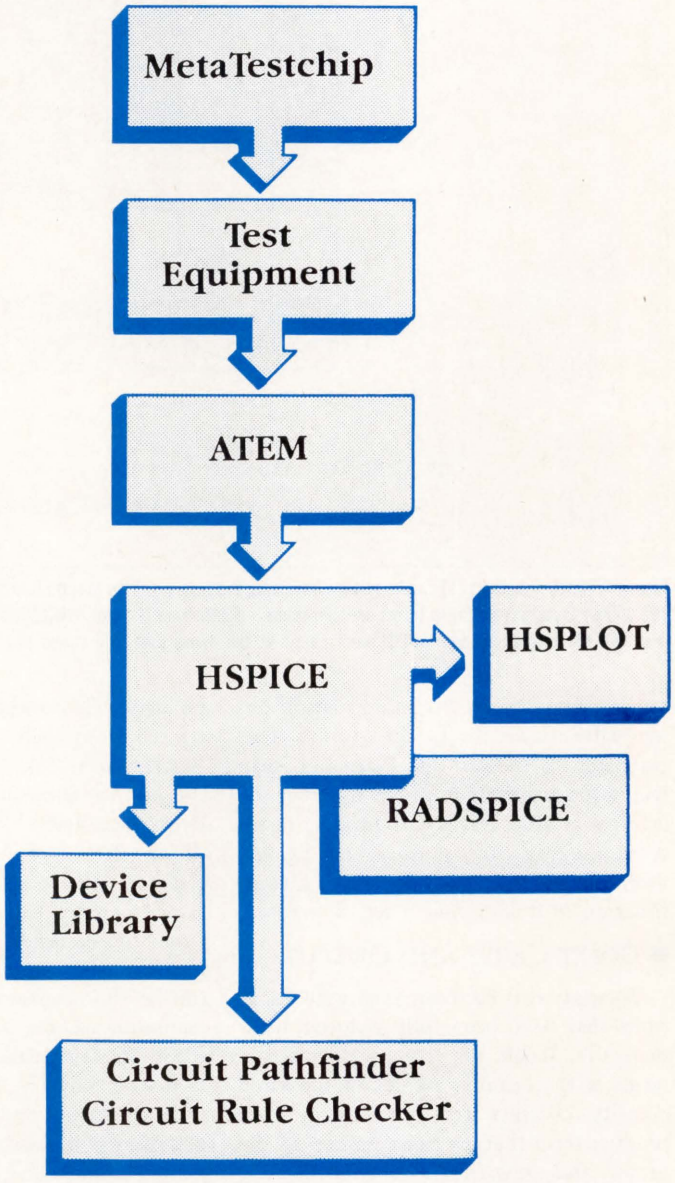

Meta-Software provides all essential support services. For answers to technical questions from experienced engineers, please call Meta's toll-free HOTLINE SUPPORT number: (800) 346-5953

HOTLINE SUPPORT is also available at our main number: (408) 371-5100

Get your circuit simulation and characterization problems resolved quickly and professionally. Call Meta today!

Meta-Software, Inc. • 50 Curtner Avenue, Suite 16 • Campbell, CA 95008
Phone (408) 371-5100 • FAX (408) 371-5638 • TLX 910-350-4928

CIRCLE NUMBER 21

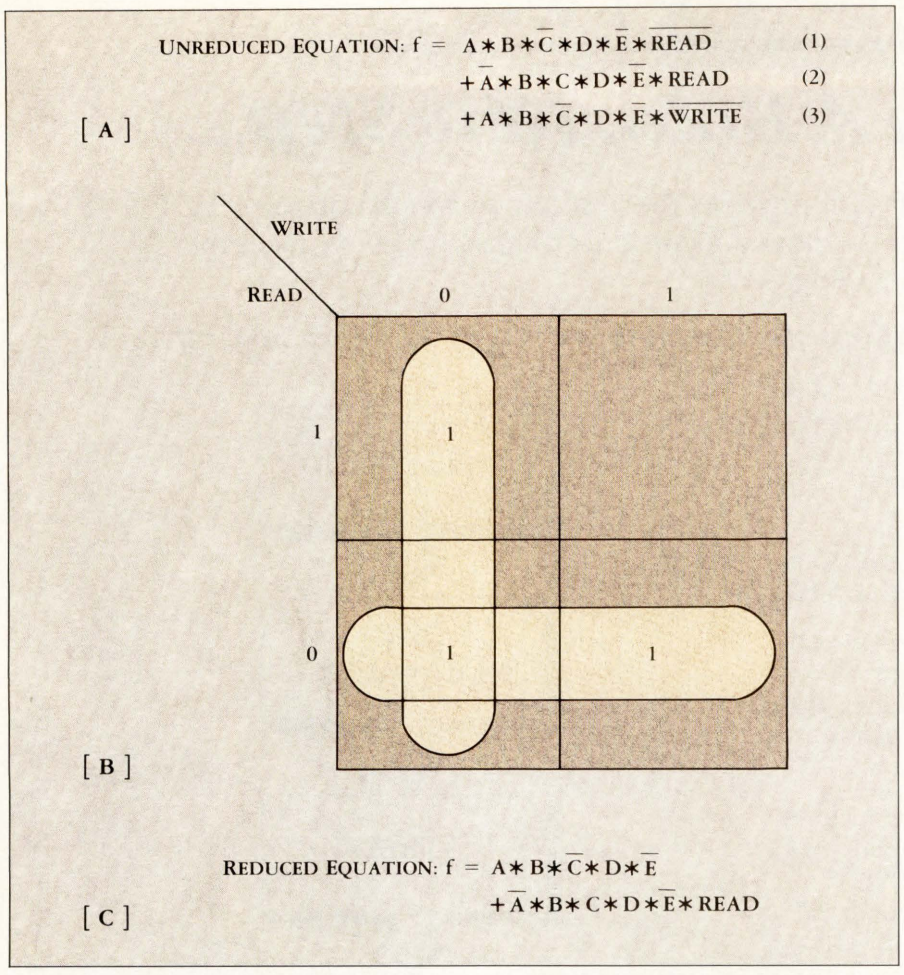

Figure 4. These equations (a) can be reduced by using the terms with no factors that appear with opposite polarity. The unique factors from those terms are extracted—!READ from (1), and !WRITE from (3)—to create a Karnaugh map (b). Knowing READ and WRITE can't occur at the same time, the terms can be combined (c).

terms, then "don't care" information can manually reduce the product terms. This extension of the familiar Karnaugh map technique attempts to group together the original product terms with product terms corresponding to unused conditions, where the term condition denotes a combination of a state and a set of inputs.

■ DON'T CARES AND OUTPUTS

To apply this method, start with equations that have been fully reduced automatically. Look for groups of product terms with common signals of the same polarity. Extract the portions of each product term that are not common to the group, and draw their Karnaugh map. If the terms for reduction can never occur, then the original set of product terms can be combined into one term.

This process is applied to the equation in Figure 4, which has already been reduced as much as possible without using don't care information. READ and WRITE are signals that denote a read and a write to a computer memory, respectively. Starting with Figure 4a, terms (1) and (3) have no factors that appear with opposite polarity. The unique factors are then extracted from each: !READ from (1), and !WRITE from (3). Their Karnaugh map (Figure 4b) shows that to combine these two terms, READ * WRITE must be a condition that never exists—a don't care. Knowing that a read and a write of the same memory can't occur at the same time, the designer can conclude that the terms can be combined (Figure 4c). This can be time-consuming, but it's a useful technique when a complicated output exceeds the available number of product terms by only one or two.

For functions too large to fit in a single PLD macrocell, two outputs may be used in parallel. The equation is split into two halves, with one half becoming the function for one macrocell and the other half forming a second output. The two outputs can then be combined in an external OR gate for the desired function.

This external gate can introduce a glitch in the output (called a static hazard) if the product terms are not allocated carefully. In general, these hazards may be present any time a state is combinatorially decoded to give an output. Suppose two macrocell outputs, X and Y, are externally combined to give the function f = X + Y. Then assume that the macrocell outputs change: X from 0 to 1 and Y from 1 to 0. The output f should remain at 1, but could have a glitch if Y hits 0 before X reaches 1. The same behavior on the inputs to the macrocell flip-flops pose no problem because they are latched after they reach steady state conditions. However, by grouping product terms that are generated from adjacent states into a single macrocell, the output problem can also be avoided.

■ USING THE 5C121

The use of the 5C121 EPLD has ramifications on state-machine design as well. To understand these ramifications, a number of useful architectural facts about the 5C121 can be gleaned from its block diagram in Figure 5. Of the eight groups of macrocells (A-1 through A-3, B-1 through B-3, and two groups of buried registers), each group, except for the buried registers, can be programmed as inputs or any of several registered or combinatorial output types.

All macrocells can communicate with others on its local bus as well as two of the buried registers. The two-way arrows between the local and global busses for the A-3, B-3 and buried register groups indicate that the outputs of these cells are routed to the global bus and are thus available as inputs to all other macrocells. The designer has the opportunity, therefore, to route signals from one side to another without using any I/O pins. For all other groups, the one-way arrows from the global to the local buses signify that the outputs from these cells are not routed to the global bus, and are therefore available as inputs only to macrocells residing on the same local bus.

The ramifications of the global and local buses must be constantly kept in mind. A signal at any given macrocell in the device may not be available to every other cell in the part. All of the registers for each state machine, therefore, were grouped to reside on the same local bus. For example, assume there were two state machines in the application, each with eight bits. Then all eight bits for the first machine would be located in registers on one local bus, and all of the flip-flops for the other state machine would reside on the other local bus. This facilitates the communication between the bits of a single state machine, but still leaves an interconnection problem for those signals that synchronize the different machines.

There are two ways to make a local signal globally accessible, assuming all

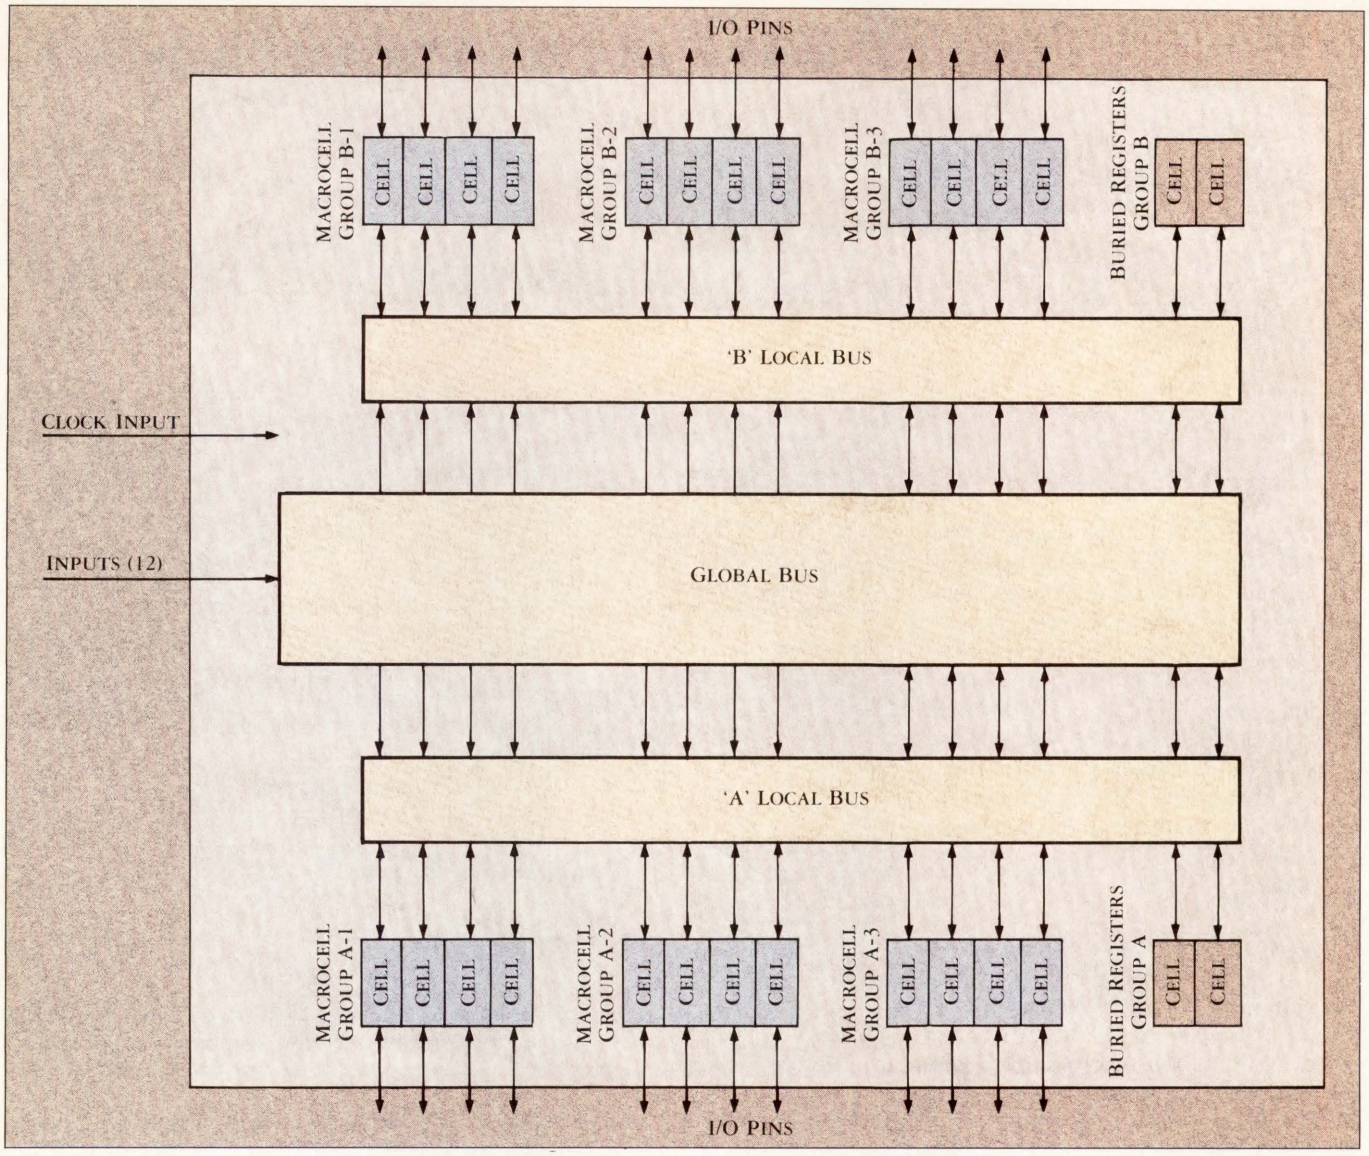

Figure 5. The ramifications of the global and local buses on the 5C121 must be kept in mind. A signal at any given macrocell in the device may not be available to every other cell in the part. In addition, with variable product terms inside of each macrocell group, designating a whole group as inputs requires wasting a high-product-term macrocell.

globally available output pins are used. Either the output can be taken off the chip and brought back in through an input pin, or the signal can be sent through a buried register to the global bus. The first approach uses up an extra input pin, while the second uses up a buried register.

The buried-register technique is facilitated by the connections from the buried registers to both the global and local buses. For example, assume that an output in the A-1 macrocell group needs to be used as an input to a state machine on the bus. To route this signal to the macrocells, it would be sent through a flip-flop in the buried registers group A to the global bus. The buried register can either produce an exact duplicate of the original output, or can use the output as its input to achieve a delayed version for the macrocell, depending on the desired system timing.

The buried register method has an added advantage in the case of state machines with high clock rates. If the output is brought in through an extra input pin, the clock period must be less than the sum of the propagation time from clock to output and the input setup time of the device. For a 65-ns 5C121, this approach reduces the maximum clock frequency from 18 MHz to 12.5 MHz. Note that a fast external flip-flop could be added so that the clock frequency would depend on the clock-to-output delay of the external register instead of that of the EPLD macrocell. By using a separate flip-flop with a 10-ns, for example, the maximum frequency would rise to 17.5 MHz. This alternative would obviously cause an additional clock-cycle delay in receiving the signal on the other local bus, however, and may cause synchroiation difficulties.

For implementing the product terms, the 5C121 also includes a variable-product-term architecture. Each group—A-1, A-2, B-1 and B-2—has four macrocells: one with four product terms, one with six, one with eight and one with ten. Macrocells in the A-3 and B-3 groups provide the largest number of possible product terms by joining two macrocells, one containing twelve terms and one containing four, with four shared product terms. The twelve-term macrocell can accept, therefore, a maximum of sixteen product terms. In that case, the neighboring macrocell has only four product terms, unless the shared product terms are common to the two macrocell outputs.

The 5C121 provides a cost-effective solution to complex state-machine design. There are a number of caveats, however. Although sixteen product terms are available for some outputs, they must be manually fitted. In addition, the device (just like PAL devices) has very long input setup times, a situation requiring resynchronization of all inputs which do not

Continued on page 84

CAE/CAD Pioneer integrates its two IC design systems

Valid Pulls Itself Together

AFTER purchasing Calma Co.'s IC-CAD business from General Electric in June, Valid Logic Systems found itself with two distinct IC design systems to support. Valid has now merged these two lines into one environment and has released some details on how it plans to integrate the two lines into one product.

Valid is also bolstering its CAE environment with the introduction of a new concurrent fault simulator and a schematic generation tool.

■ IC DESIGN

Valid's IC Design System integrates its logic design, chip design and verification software with the IC layout tools from Calma. It supports cell-level design with either Calma's EDS III layout editor, which provides direct access to GDS II libraries, or Valid's original ValidLED (Figure 1). Both tools feed cell designs to the block-level design tool Compose, which assembles blocks and chips automatically or interactively.

Completed blocks and chips are fed through either editor to Valid's verification tools in GDS II stream format. Both flat (DRC) and hierarchical (ValidDRC) verification of physical design rules is possible. For comparison of the completed layout with the original schematics design, the Extract program creates a netlist for comparison with netlists created within the ValidGED schematic editor. Extract also provides physical parasitics for back-annotation to simulators. Completed IC designs can be passed to Calma's CAM tools to create data compatible with plotters and optical or e-beam pattern generators.

Because EDS III and ValidLED serve an established base of designers, Valid will continue to support both as it prepares a single, unified IC design editor to replace them. The unified tools will be called EDS III version 2.0, indicating that that product will be the primary vehicle. According to David Hightower, Director of System Architecture at Valid, it will incorporate the interactive interface now used with Compose, EDS III's programming language, and a new database which will be smaller and faster than the present one. It should be functional next spring and commercially available next fall.

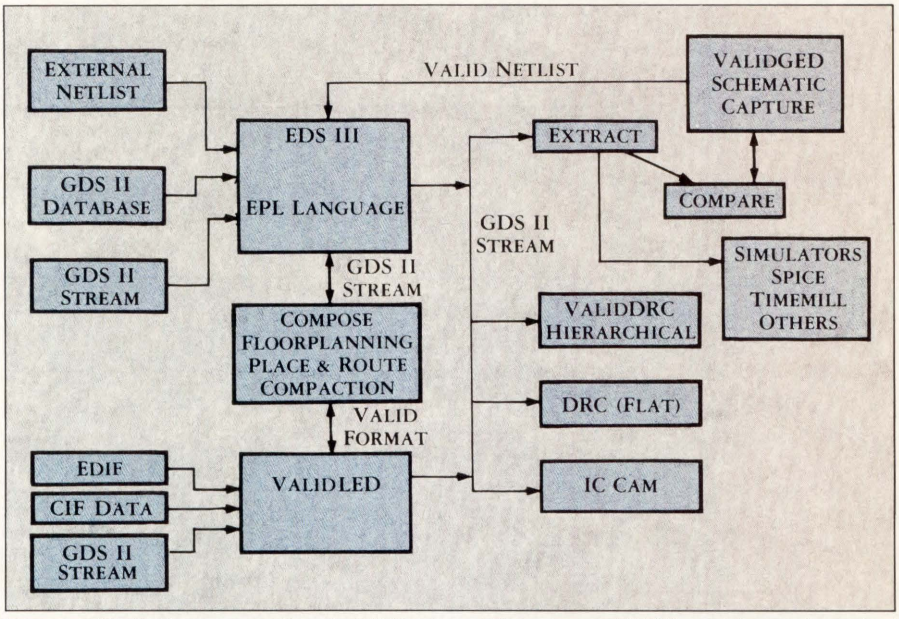

Figure 1. Valid's IC CAD system combines the EDS III editor, a DRC program and IC CAM output from newly-acquired Calma with Valid's Compose layout tool and other verification programs. Today, users can choose between EDS III and Valid's homebrew ValidLED editors; Valid is unifying the editors in upcoming EDS III version 2.0.

■ NEW CAE TOOLS

RapidTEST is the first of Valid's three TestBridge products that link design and test. RapidTEST is based on a proprietary concurrent fault simulation technique that, according to Valid, evaluates 3,000 events per second for each MIPS of platform processing power. It accepts schematics from ValidGED and the libraries and stimulus vectors used with ValidSIM, Valid's logic simulator.

A "hierarchical fault injection facility" gives designers control over the types and locations of simulation faults. It performs automatic fault collapsing to remove redundant faults. The program uses an Automatic Tester Attribute Rules file to ensure that the test can be recreated accurately on the target ATE system.

RapidTEST's fault-coverage reports include fault coverage levels, node toggling levels, undetected fault lists, and a histogram of faults detected per test vector.

The other new CAE tool, Transcribe, converts hierarchical designs, under interactive control, into schematic drawings on a page-by-page basis. Integrated with ValidGED, it expands the design hierarchy to place all schematic components in relative locations according to the original coordinates in the hierarchical description. The user can modify these placements. Next, Transcribe routes all the interconnecting signals to complete the schematics. It can create drawings that conform to drafting standards such as MIL-STD-100C, DOD-D-1000 and ANSI Y14.1 and Y14.2.

The IC Design System is available now for Sun Microsystems Inc. platforms for $59,000. RapidTEST and Transcribe run on Sun and Digital Equipment Corp. workstations. RapidTEST is priced at $17,500, while Transcribe is tagged at $19,900.

Valid Logic Systems Inc.
San Jose, Calif.
(408) 432-9400

More Than 50,000 Good Reasons To Attend **Wescon**/88®

In addition to the most qualified audience you can reach at any show of its kind, Wescon gives you:

- More than 6,000 live product demonstrations.
- More than 1,300 vendors.
- 12,000 factory experts and engineers.
- 36 idea-generating technical sessions.
- 11 incisive and problem-solving tutorials.
- Electronics Theater featuring state-of-the-art developments.

And This Year, Wescon Gives You More:

- New Automated Design Center.
- New Products and Technology Achievement Awards.

Wescon is an exciting, one-of-a-kind opportunity where the engineer and the people who will supply him congregate for an exchange of new ideas and new technologies. It is a place where new ideas are introduced and design problems are solved.

There is no better event to enhance your professional growth and technological awareness than Wescon/88.

Wescon Is More Than An Event, It Is An Investment!

Wescon/88®
Nov. 15-17, 1988
Anaheim Convention Center
Anaheim, California

To pre-register, fill out the form below, or for a Preview Program, call (800) 262-4208 in California or (800) 421-6816 outside California.

Wescon/88®

Pre-Registration Fee – $5.00

6-90140

☐ ZX **IEEE Member**
☐ ZY **ERA Member**
☐ ZZ **Other Affiliation** _____

Please PRINT as you want shown on badge

Name _____ (FIRST OR INITIALS / LAST)
Position • Telephone _____ • _____ (AREA CODE / NUMBER (OPTIONAL))
Company/Organization _____
Address: Bus. or Home _____ (DEPT. or M/S)
City, State, Zip Code _____ (STATE / ZIP CODE)

CIRCLE APPLICABLE BOXES

INDUSTRIAL CLASSIFICATION

BA	Computers
BB	Communications
BC	Military/Aerospace
BD	Automotive
BE	Electronic Components, Devices
BF	Business Equipment
BG	Medical
BH	Consumer/Entertainment
BI	Consulting/Education
BJ	Other _____

JOB FUNCTION (ONE ONLY)

AA	Systems Design
AB	Circuit Design
AC	Packaging, Prototype Design
AD	Production/Fabrication
AE	Research/Development
AF	Purchasing
AG	Engineering Management
AH	Administration; Management
AI	Testing, QC/QA
AJ	Consulting; Education
AK	Other _____

AREAS OF MAIN PRODUCT INTEREST

AL	CAD/CAE Tools; Systems
AM	Computer Software
AN	Electronic Components
AO	Electro-Mechanical Components
AP	Production Equipment, Machinery
AQ	Production Supplies, Materials
AR	Consulting Services
AS	Fabrication Services
AT	Test, Insp., Measurement Inst.
AU	Testing Systems
AV	Other _____

208764

To receive your badge in advance, please complete this form and mail it along with your check to arrive NO LATER THAN October 28, 1988. After October 28 the at-the-door fee applies. If you do not receive your badge prior to Wescon/88, check in at the advance registration counter in the Anaheim Convention Center. **NO ONE UNDER THE AGE OF 16 ADMITTED. No refunds after November 11, 1988.**

MAIL TO: WESCON/88, P.O. BOX 54760, TERMINAL ANNEX, LOS ANGELES, CA 90054-0760.

CIRCLE NUMBER 22

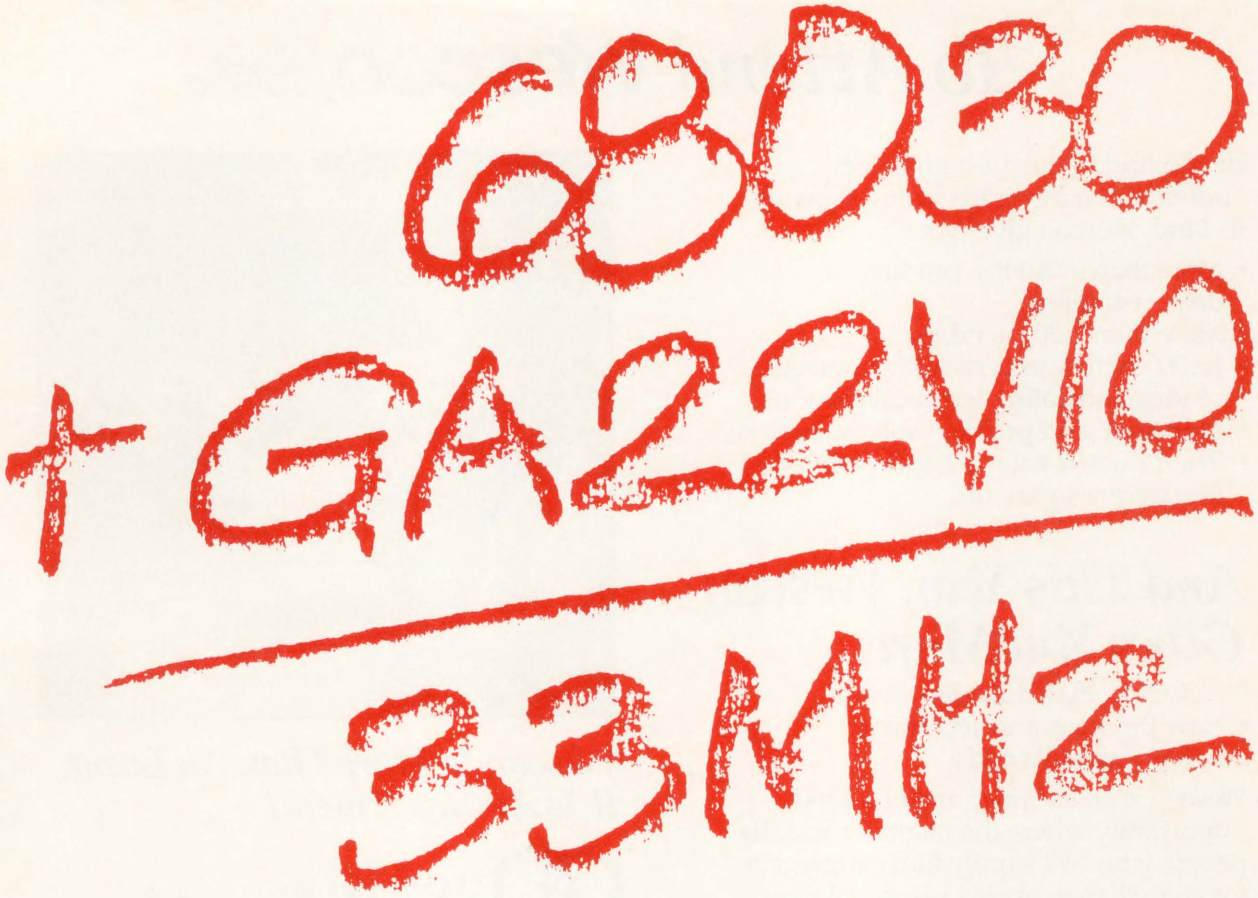

Breakneck Performance at Breakthrough Prices.

Traditionally, the prime obstacle to outrageous performance is an outrageous price.

Meet the GA22V10s from Gazelle Microcircuits.

They're the two fastest, high-density, user-configurable logic chips available. Chips that are up to twice as fast as their silicon equivalents. Chips that eliminate the high price of high performance.

Running at 110 MHz with a propagation delay of just 7.5 ns, the GA22V10-7 liberates the true performance of your system. It's fast enough to let 68030s and RISC processors like the 88000 and SPARC™ run at full speed. Cache control, bus arbitration and wait state bottlenecks simply disappear while workstations and PCs gain a 33% boost in performance.

PART NO.	t_{PD}	t_S	t_{CO}	f_{MAX}
GAZELLE MICROCIRCUITS				
GA22V10-7	7.5 ns	3.0 ns	6.0 ns	110 MHz
GA22V10-10	10.0 ns	3.6 ns	7.5 ns	90 MHz
AMD				
PAL22V10	15.0 ns	10.0 ns	10.0 ns	50 MHz
PAL16R8	7.5 ns	7.0 ns	6.5 ns	74 MHz
TEXAS INSTRUMENTS				
TIBPAL22VP10	20.0 ns	20.0 ns	12.0 ns	37 MHz
CYPRESS SEMICONDUCTOR				
PALC22V10	25.0 ns	15.0 ns	15.0 ns	33 MHz

Gazelle's TTL-compatible GaAs is the reason. Our GA22V10 is 100% pin and function compatible with slower silicon 22V10s. We created 25 new internal circuit techniques to deliver this performance at this price. Which is why you don't see these specs from any other manufacturer. In GaAs or silicon.

All it took was a break from tradition.

gazelle

(408) 982-0900

Call for an information kit. Or send your name and address to Gazelle, Dept. A, 2300 Owen St., Santa Clara, CA 95054.

CIRCLE NUMBER 30

Continued from page 79

arrive within a very narrow time window. An input latch function with one-half as much setup time is provided to offset this drawback. For fast clock rates, however, even these functions may not latch inputs that do not change within a small time window. Although flip-flops other than D type may reduce the number of product terms for an implementation, most software packages that accept state diagram descriptions as input do not support all of the flip-flop types available in the devices. All in all, EPLDs form a good solution to designs that require large state machines on short lead times. ■

ACKNOWLEDGMENTS

Many thanks go to Pedro Vargas of Intel Corporation, who provided valuable applications assistance during design, and Patricia Stoltz of Hewlett-Packard Company, who forwarded helpful comments and suggestions on this manuscript.

REFERENCES

DOLOTTA, T. and MCCLUSKEY E., 1964, *"The Coding of Internal States of Sequential Circuits," IEEE Transactions on Computing, EC-13, October 1964, pp. 549-562.*

ETTINGER, M., 1968, *"The Synthesis of Sequential Circuits," Computer Design, December, 1968, pp. 46-53.*

FLETCHER, W., 1980, *An Engineering Approach to Digital Design, Prentice-Hall, pp. 290-295.*

RHYNE, V. and NOE, P., *"On the Number of Distinct State Assignments for a Sequential Machine," IEEE Transactions on Computing, January, 1977, pp. 73-75.*

TUMBUSH, G. and BRANDEBERRY, J., 1974, *"A State Assignment Technique for Sequential Machines Using J-K Flip-Flops," IEEE Transactions on Computing, January, 1974, pp. 85-86.*

ABOUT THE AUTHORS

DAVID SHARP *is a Development Engineer at Hewlett-Packard* Company*'s San Diego Division. He holds a BSEE from the University of Missouri at Rolla, Mo., and an MSEE from Stanford University. He enjoys music, having played the piano for 20 years.*

GEORGE BARBEHENN *is also a Development Engineer at Hewlett-Packard's San Diego Division. He has also worked as a design engineer, helping to develop computer peripherals and plotters such as the 7600 Series. He holds a BSEE from the University of Maryland at College Park, Md. Barbehenn is bilingual and enjoys scuba diving and soccer.*

Continued on from page 58

The Ultimate Intelligent Memory

The AJC search is maximally fast because the AJC architecture is inherently that of a contents-addressable-memory (CAM). The CAM has for many years been envisioned as the ultimately-desirable form of intelligent memory. Data is located by direct, simultaneous comparisons of the source data in the comparand register against all the target words in the AJC. There are no wasteful, time-consuming addressing and data movement steps as would be needed in a Von Neumann machine that would sequentially bring in each word one at a time for the comparison.

The distributed control system is more complex than the single controller system; however, it decreases the administrative burden of the earlier single-controller design. The precharged design also requires a more complicated clock distribution network. The new design's ability to restrict the search to just the attribute length, however, greatly increases the speed of a search.

Because the transistors only have to discharge a bus and do not have a direct path to ground, both NMOS and CMOS precharged designs can use minimum-sized transistors, which consume less power. This is particularly advantageous for NMOS when compared to traditional depletion-mode NMOS circuits.

In a precharge scheme, the speed of the design is related to the speed of the pull-down devices. Although transistors driven by NMOS pass transistors are normally doubled in size to produce the same speed, in our design we chose to accept the speed reduction to obtain minimum geometry.

Table 2 compares the CMOS precharged design area and speed parameters to the original CMOS design. The final CMOS design tabulated there uses both precharging and a domino technique. The domino CMOS associative memory cell's mask and tag bit are shown in detail in Figure 4 and at full AJC level in Figure 1.

A tabulation of both the active area and speed estimates for the CMOS version of the original design and the domino precharged CMOS design is given in Table 2. Domino CMOS uses approximately 28 percent less active area than the original CMOS design. The NMOS precharged version, however, resulted in a 43 percent reduction in the active cell area.

■ FUTURE PLANS

Our future plans include additional improvements in the associative memory design in an attempt to even further reduce its active area. By combining several gates into a few complex gates, such as in the tag bit circuitry, we can further reduce the area and execution time of precharged configurations.

We also plan to implement the precharged design in a full custom 1-micron AJC (with the help of the NCR Corp.) and to use these to build a prototype system that will include several AJCs in a multiprocessing environment.

Meanwhile, we have been speculating about other, more general applications for these massively parallel architectures. Having a 1-bit ALU associated with every memory cell might prove useful for array and digital signal processing. The performance gains with massive parallelism might more than offset losses because of the elementary 1-bit ALUs. For example, massive parallelism and the ability to do associative searches might compensate for having to do the multiplications necessary in array and signal processing by crude successive additions. ■

REFERENCES

PETRIE, CHARLES R., *"A Database VLSI Associative Join Module," Technical Report, Department of Electrical Engineering, The Pennsylvania State University, University Park, Pa.*

HURSON, ALI R., *"VLSI Time Space Complexity of an Associative Parallel Join Module," Proceedings of the International Conference on Parallel Processing," Chicago, August 19 - 22, 1986, pp. 379-386.*

ABOUT THE AUTHORS

CHARLES R. PETRIE *is a Captain in the US army. He received his MSEE degree from Pennsylvania State University, using the subject of this article as his thesis.*

ALI R. HURSON, *is a Professor of Computer Architecture in the Electrical Engineering Department of the Pennsylvania State University. He recieved his Ph.D. from the University of Central Florida, Orlando Fla., continuing on from University of Iowa, Iowa City, Iowa. He has a special interest in database machines.*

Some people are never satisfied. We need lots of them.

We're running this ad because we have lots of immediate openings at Mentor Graphics.

The people who do well here are looking to break out of the pack. You know the type. Rule changers. Rebels. Leaders. People with a vision and passion so intense it gets everyone else excited.

In fact, at Mentor Graphics, we work in a controlled frenzy. Controlled because there is no uncertainty about the company's stability. Frenzy because we have an exciting vision of the future of design automation.

Our vision sees a world-class systems solution to design tools. A vision where application and integration are equally important.

If you only feel comfortable when challenged, explore these immediate opportunities in the following areas.

Product Development

Engineering Managers
Software Development Engineers
QA Engineers

Marketing

Product Marketing Managers
Technical Marketing Engineers
Technical Writers

As the leader in electronic design automation, Mentor Graphics offers you the opportunity to work in a variety of technologies. You should have 3 or more years of professional experience developing or marketing any of the following technologies:

- **CASE**
- **Design Creation**
- **Simulation**
- **IC Design and Layout**
- **PCB Design and Layout**
- **Electronic Packaging**
- **Design Management and User Interfaces**
- **Computer-Aided Documentation**

Satisfied?

We hope not. Because we want you to learn more about the many benefits to a career with Mentor Graphics.

Among these is a choice of lifestyles offered in either our Beaverton, Oregon or San Jose, California locations. Nestled in the green hills of the Pacific Northwest, the serene surroundings of Beaverton provide a healthy balance to our fast-paced technical environment. Plus, the area boasts a very low cost-of-living. Or you can work in our facility in San Jose, the nation's hub of high tech, where you are as close to skiing and wine tasting as you are to the beautiful Pacific coast.

To apply, please send your resume to Professional Staffing, Mentor Graphics Corporation, Dept. 3819, 8500 S.W. Creekside Place, Beaverton, Oregon 97005-7191.

We are an equal opportunity employer committed to affirmative action.

Mentor Graphics®

Your ideas. Our experience.

ADVERTISERS' INDEX

READER SERVICE NUMBER		PAGE NUMBER
12	Cad Software	39
7	Cadic	23
11	California Micro Devices	37
	Cypress Semiconductor	73
	Daisy Systems Corporation	5
15,16, 25,24,	Data I/O Corporation	49,51 75,CV4
22	Electronic Conventions Management	81
1	Fujitsu	1
30	Gazelle Microcircuits Inc.	82
	GE/SOLID STATE	53
8	GenRad	25
23	Harris Semiconductor	CV3
17	IKOS	59
14	Intergraph Corporation	43
	Mentor Graphics	13,57
21	Meta Software	77
10	OKI Semiconductor	35
9	Plessey Semiconductor	27,29
2	Raytheon	7
3	Saratoga Semiconductor	9
4	Seattle Silicon	15
18	Sierra Semiconductor	61
6	S-MOS Systems Inc.	21
20	Toshiba America, Inc.	69
5	Viewlogic	17
19	Zilog Incorporated	65
13	Zycad	41

This index is provided as an additional service. The publisher does not assume any liability for errors or omissions.

VLSI Systems Design

PUBLISHER
NORM ROSEN (213) 473-9641
NATIONAL SALES MANAGER
JOHN C. COX (516) 562-5710
DIRECTOR OF MARKETING
JAY McSHERRY (516) 562-5780

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

NEW ENGLAND (617) 244-5333
KAREN TALLARIDA, District Mgr.

DALLAS (214) 661-5673
CLAIRE FLORA, Senior Accounts Mgr.

CHICAGO (312) 565-2700
MARY SHUTACK, District Mgr.
TOM NILSEN, District Mgr.

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

PORTLAND (503) 636-7694
RICHARD CARLISLE, District Mgr.

LOS ANGELES (213) 473-9641
TODD BRIA, District Mgr.

NEWPORT BEACH (714) 851-2022
BILL BARRON, District Mgr.

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

CLASSIFIED ADVERTISING
NEW YORK (516) 562-5000
MICHAEL ZERNER, East Coast
NEW ENGLAND N.Y. N. Central, District Mgr.
KATHY HEALY, Southeast, S. Central District Mgr.
SONJA WONG, Advertising Coordinator
FELICIA SABATINI, Graphics
SAN JOSE (408) 252-6191
LAURIE MUNCE, West Coast District Mgr.
DIRECT MARKETING SERVICES
PETER CANDITO, Sales Manager
ARLENE BROWN, List Sales Representative
(516) 562-5000
CUSTOMER SERVICE/DISPLAY ADVERTISING
(516) 562-5000 FAX (516) 365-4829
LYNDIANE HERVEY, Group Mgr.
LINDA USLANER, Ad Coordinator
PROMOTION
GENEVIEVE HIGGINS
CIRCULATION & RESEARCH
RICHARD MATTUCCI, Circulation Mgr.
DANIEL R. CAMPBELL, Electronics Group Res. Mgr.
SUBSCRIPTION SERVICES
(516) 562-5882
REPRINTS
KATHI STEBAN
(516) 562-5873

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

BPA

CMP PUBLICATIONS, INC.
(516) 562-5000
President, MICHAEL S. LEEDS
Vice President/Treasurer, PEARL TURNER
Vice President, DANIEL H. LEEDS
Co-Chairpersons, Board of Directors
GERARD G. LEEDS, LILO J. LEEDS

BUSINESS/OPERATIONS
Director of Operations, GRACE MONAHAN
National Classified Mgr., SANDRA HERMAN
Mktg. Srv. Dir., DORIANNE WALTHER-LEE
Customer Service Manager, GEORGETTE ROSS
Direct Mktg. Services Mgr., JOANNA BRANDI
Director of Manufacturing, STEPHEN J. GRANDE
Production Manager, MARIE MYERS
Systems Manager, HILDE DONISI
Circulation Director, NANCY MERRITT
Research Director, JOHN EATON
Design Director, MARIE D'IPPOLITO
Art Services Supervisor, PAT CARASONE
Assistant to the President, NANCY GRABOW

The Wrong Place To Find Out Your ICs Aren't Rad-Hard Enough.

Harris guarantees performance of its rad-hard circuits. Looking for the best protection for your systems? Come to the most specified source for military rad-hard ICs...Harris.

Our technology credentials span more than two and a half decades of programs, from Minuteman and Polaris to Trident II. Soon SDI.

The only way to stay a leader in rad-hard and hi-rel circuits is to know where technology has been, and where it's going. Our know-how shows through...from more progress in new isolation technologies like Silicon-On-Insulator to continued advances in tapping the inherent speed of GaAs.

Your job's not finished so neither is ours. Count on us for the rad-hard military and commercial range ASICs you need to keep your systems at peak performance.

Our Rad-Hard/Hi-Rel Data Book tells the whole story of our products and processes. To order your copy, call Harris Semiconductor Custom Integrated Circuits Division.

In the U.S.: 1-800-4-HARRIS, Ext. 1908.
In Canada: 1-800-344-2444, Ext. 1908.

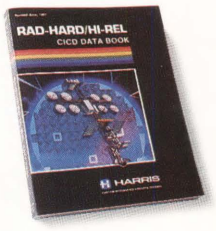

HARRIS
SEMICONDUCTOR SECTOR

CIRCLE NUMBER 23

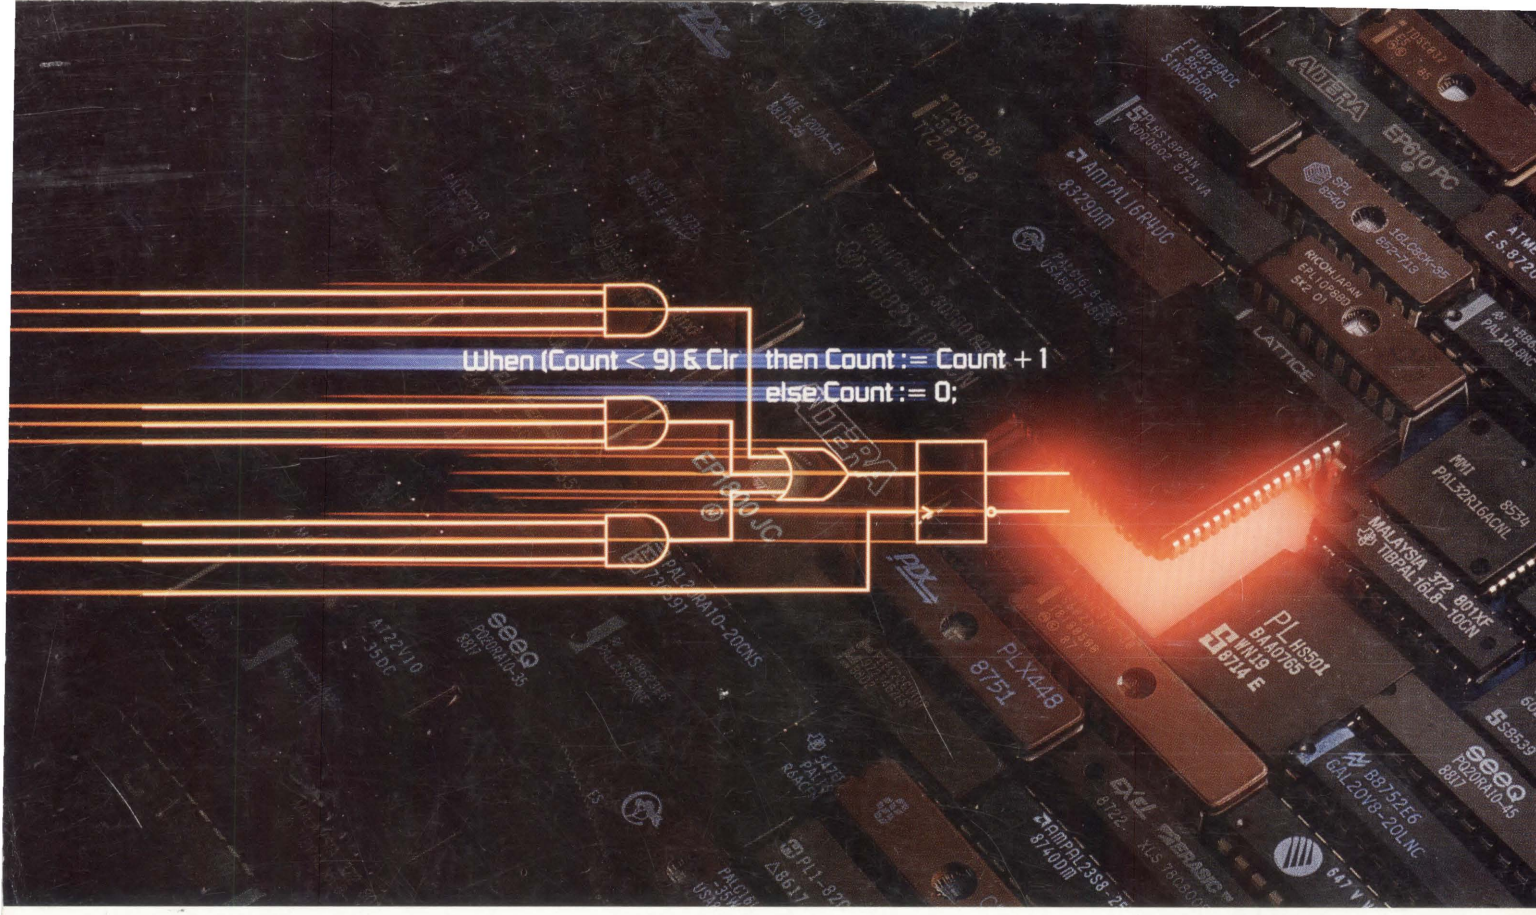

17 MILLION PLDS A MONTH BEGIN WITH DATA I/O.

No company supports as many programmable logic devices as Data I/O®—virtually every PLD on the market. From design, to programming and test, Data I/O is the industry choice for PLD development tools.

UNIVERSAL SUPPORT FOR EVERY PLD MANUFACTURER. We give you the freedom to choose the right PLD for your design—from any manufacturer. You can use our tools to design and program all PLD architectures, including PALs, FPLAs, PLEs and EPLDs, and all technologies from vertical-cell AIM, ECL, and CMOS to bipolar. And you can pick any package—DIP, PLCC, LCC or SOIC.

UNIVERSAL DESIGN TOOLS. ABEL™ 3.0, Data I/O's high-level design language, lets you describe your design with any combination of equations, state diagrams, or truth tables. For multiple PLD design, use the advanced logic synthesis capabilities of GATES™. And for complete device testability,

PLDtest® is our universal tool for automatic fault grading and test vector generation.

UNIVERSAL PROGRAMMING TOOLS. The UniSite™ 40, our most sophisticated programmer, handles all devices—memory, logic, and microcontrollers—in a single site. The modular approach to logic programming is the 29B with a LogicPak™. Or for portable convenience, choose the 60A.

UNIVERSALLY APPROVED BY PLD MANUFACTURERS. With Data I/O, you can depend on the accuracy and reliability of your programmed PLD. We work with semiconductor manufacturers—long before devices are released—to develop, and obtain approvals for, our programming algorithms. So reliable support will always be ready when you need it.

A CLOSER LOOK AT UNIVERSAL SUPPORT. Call us today for more information about our complete line of PLD development tools and qualify for your FREE copy of the ABEL 3.0 demonstration disk. Your PLDs will be off to a better start and a faster finish.

1-800-247-5700
Ext. 341

Data I/O Corporation 10525 Willows Road N.E., P.O. Box 97046, Redmond, WA 98073-9746, U.S.A. (206) 867-6899/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

DATA I/O Corporation

CIRCLE NUMBER 24