EIGHT DOLLARS A CMP PUBLICATION SEPTEMBER 1988

VLSI Systems Design

FOR DESIGNERS OF HIGH-PERFORMANCE SYSTEMS

GaAs & RISC

Drive Microprocessors Even Faster

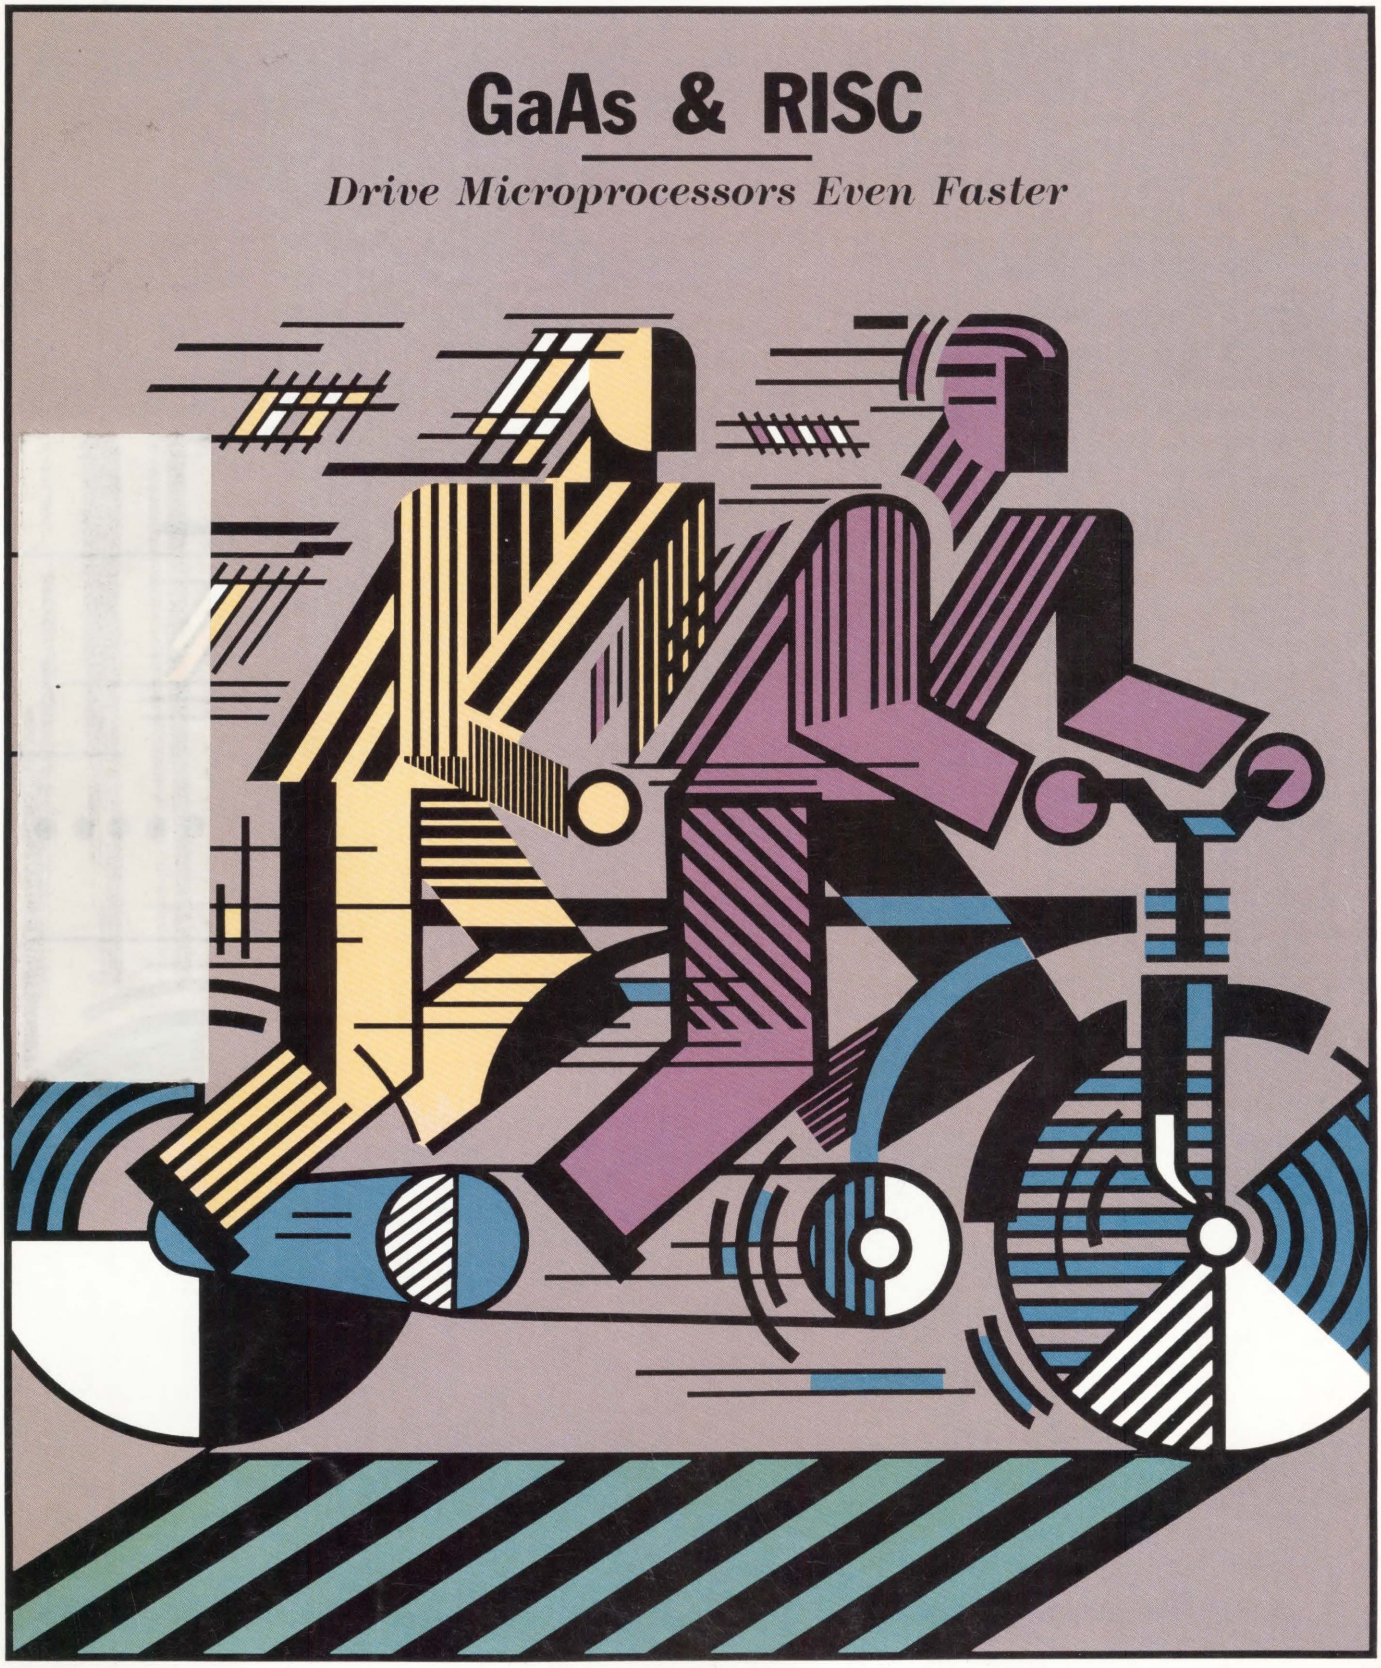

Daisy's standard now runs on an

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.

Contents

ARTICLES

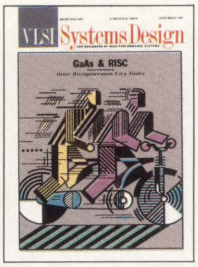

GaAs and RISC combine to offer the promise of new levels of microprocessor performance

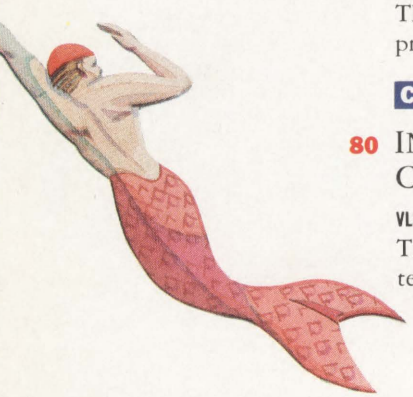

DEPARTMENTS

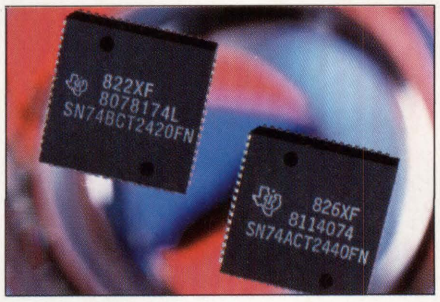

COVER ILLUSTRATION BY JAVIER ROMERO

VLSI Systems Design

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

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

CMP PUBLICATIONS, INC.
600 Community Drive
Manhasset, New York 11030
(516) 562-5000
Publishers of:
Electronic Buyers' News, Electronic Engineering Times, VLSI Systems Design, Computer Systems News, Computer Reseller News, VARBUSINESS, UNIX Today!, InformationWEEK, CommunicationsWeek, CommunicationsWeek International, Business Travel News, Tour & Travel News, Long Island Monthly, HealthWeek

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

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

Our ASIC line-up is only half the story.

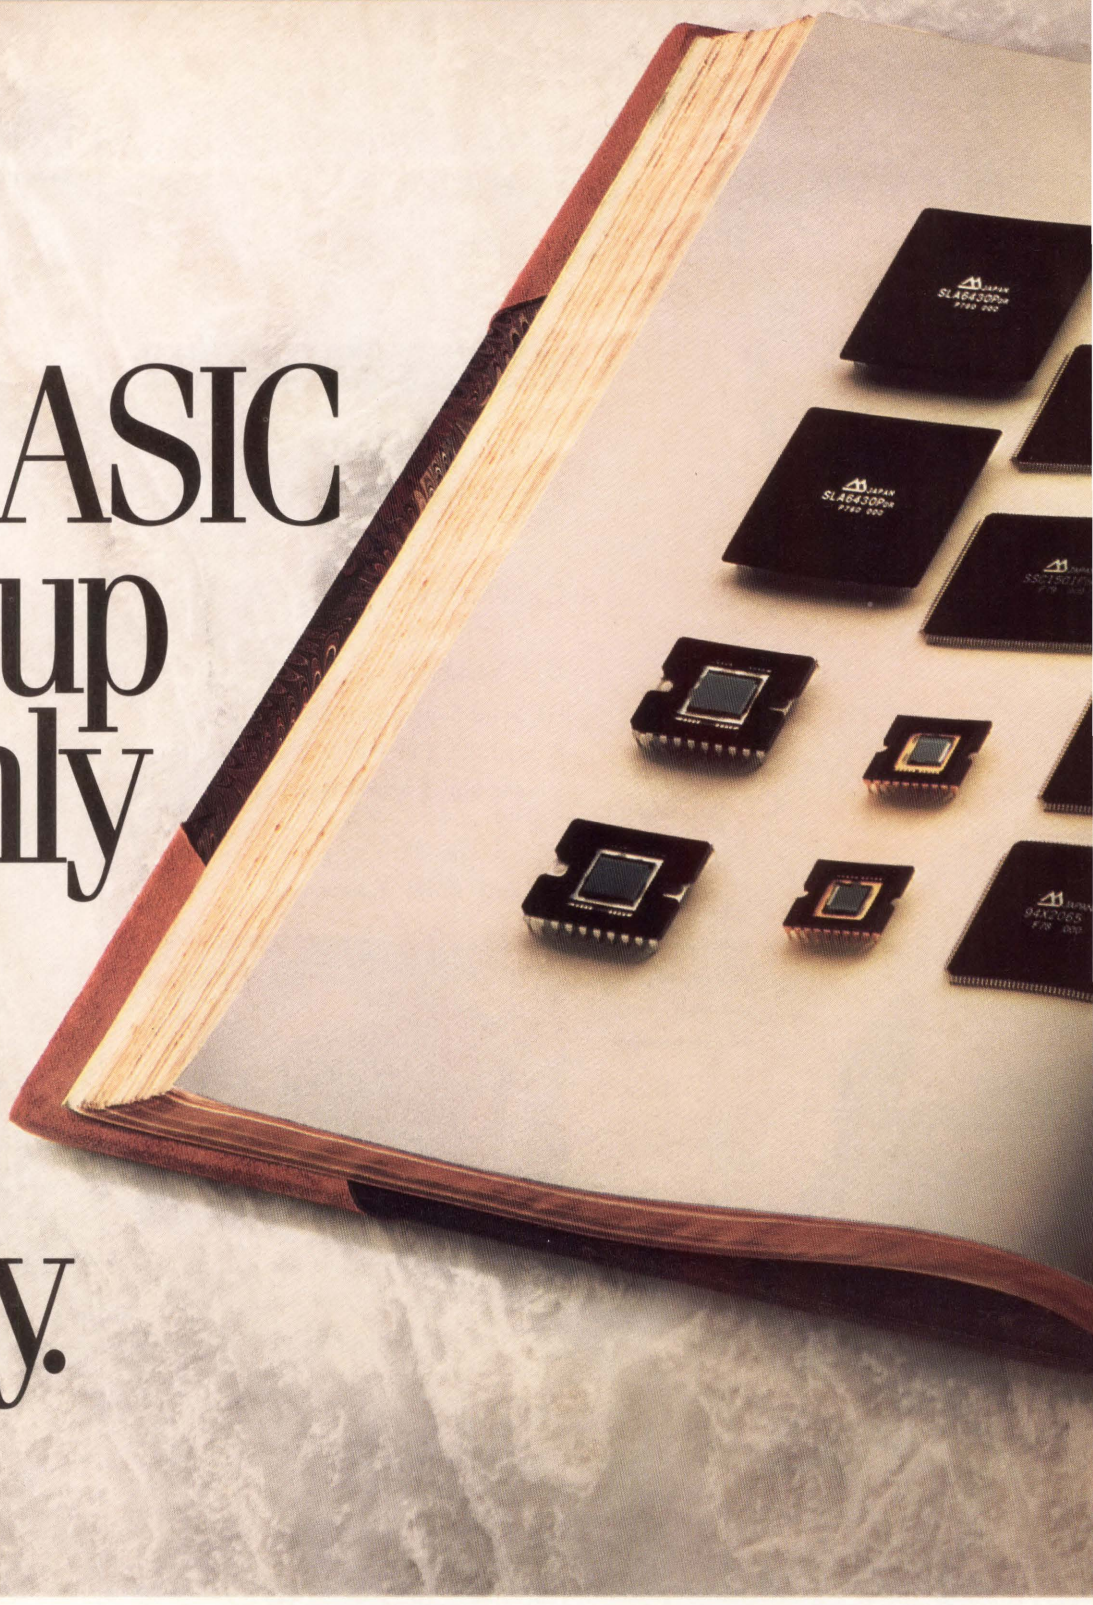

When it comes to delivery of high-quality, reliable ASICs, S-MOS wrote the book.

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 1

Has the ATE Industry Reached a Plateau?

The time is ripe for the ATE industry to take inventory

The consensus of industry experts is that the total 1988 worldwide sales of automatic test equipment (ATE)will be about $2.2 billion—with roughly an even split between chip and board testers. That's about seven percent of the anticipated total worldwide sales of semiconductors this year. Seven percent wouldn't be too bad if it represented the cost of testing. Unfortunately, most semiconductor device testing costs are a lot higher, and some have risen to almost half of the total manufacturing costs. That's when all the programming and test labor, fixturing, and overhead costs are added to the amortization costs of the capital invested in the test equipment.

Is the high cost of testing the reason that some industry pundits are predicting a relatively flat growth curve for the ATE industry? Probably not. It's more likely a result of the uneasiness that the electronics industry feels about the economy after the last prolonged downturn.

Instead of feeling down, this could be a great opportunity for the ATE industry to review its strategies for the future. In the past, with business booming, everybody was too busy to plan ahead. They tended to react to the electronics industry's growing testing problems, rather than work within the industry to develop new ways to avoid testing's many pitfalls.

It's time for the members of the ATE industry to get together and jointly start tackling the problems of testing today's and tomorrow's increasingly complex chips, boards, and systems. With a united front, it would be a lot easier to convince their customers of the benefits of planning—at the earliest stage in the design cycle—for standard test buses, BIST techniques, scan systems, or whatever the ATE industry recommended as a solution for slowing down the accelerating cost of testing.

This could boost the ATE industry into orbit again, and at the same time bring the cost of testing back down to Earth.

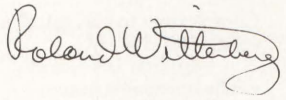

ROLAND WITTENBERG
EDITOR-IN-CHIEF

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

ESSCIRC '88

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

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

ICCD '88

October 3–5
Rye Town Hilton
Port Chester, N.Y.

The annual International Conference on Computer Design is sponsored by the IEEE Computer Society and the IEEE Circuit and System Society in cooperation with the IEEE Electron Devices Society. It will emphasize the interactions between system and memory design, logic circuit design, architecture, software, CAD, testing, physical design, and VLSI technology. The conference will feature technical sessions on topics that include test generation, advanced system interconnect and packaging, microprocessor architecture, and simulation. Additional information may be obtained by contacting Prathima Agrawal, General Chairman, AT&T Bell Laboratories, 600 Mountain Ave., Room 3D-480, Murray Hill, N.J. 07974. (201) 582-6943. ■

2ND SYMPOSIUM ON THE FRONTIERS OF MASSIVELY PARALLEL COMPUTATION

October 10–12
George Mason University
Fairfax, Va.

This symposium will focus on the increasing importance of massively parallel computer systems and data parallel programming techniques. Topics that will be presented include programming languages, architectures, algorithm development, graph theory, image processing, numerical modeling, database management, interconnection networks, hierarchical structures, neural networks, and new technologies. To receive an advance program with registration information, send your name, address, and phone number to Frontiers '88 Symposium, P.O. Box 334, Greenbelt, Md. 20770. ■

5TH INTERNATIONAL ELECTRONIC MANUFACTURING TECHNOLOGY SYMPOSIUM

October 10–12
Lake Buena Vista, Fla.

The Fifth International Electronic Manufacturing Technology Symposium (IEMT) is sponsored by the IEEE's Components, Hybrids and Manufacturing Technology Society. Technical presentations will cover such topics as developments in surface mounting technology equipment, automation applications in manufacturing of semiconductors and electronic products, computer integrated manufacturing applications for electronic products, developments in semiconductor materials and processing, advanced materials and processing applications for packaging interconnections and surface mount, and techniques and examples of designing for automation. For more details, contact Bill Moody, Vice Chairman for Administration, 2529 Eaton Road, Wilmington, Del. 19810. (302) 478-4143. ■

1988 INTERNATIONAL SYMPOSIUM ON MICROELECTRONICS

October 17–19
Washington State
Convention Center
Seattle, Wash.

The 1988 International Symposium on Microelectronics will feature panel discussions and educational tutorials as well as technical presentations on topics including automation, CAD/CAM, cofire tape technology, hybrid design, hybrid microcircuit technology, interconnections, reliability and failure analysis, materials, and multilayers. For more details, contact the International Society for Hybrid Microelectronics, P.O. Box 2698, Reston, Va. 22090. (703) 471-0066. ■

INDUSTRY-UNIVERSITY ADVANCED MATERIALS CONFERENCE

March 6–9, 1989
Embassy Suites Hotel
Denver, Colo.

Sponsored by the Advanced Materials Institute, the Colorado Advanced Technology Institute, and the Materials Research Society,

Continued on page 16

ILLUSTRATION BY RENEE KLEIN

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

BIT Stream

GE Solid State Heads South to Harris

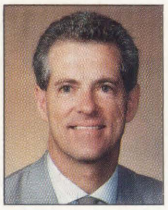

JON E. CORNELL

Harris Corp. (Melbourne, Fla.) and the General Electric Company have agreed on the acquisition of GE Solid State by Harris. Included in the deal are the RCA and Intersil semiconductor products as well as the original GE power and opto-electronic lines. The acquisition will add approximately $550 million in annual sales to the Harris semiconductor operation, which has annual sales of about $300 million. The new 14,000 employee operation will come under the wing of Jon E. Cornell, Harris senior vice president and head of the Semiconductor Sector. ■

News PLDs Cut Power, Metastability

Gould Electronics (Pocatello, Ida.) has produced the first CMOS replacements for the PLS 153 and PLS 173 bipolar PLDs from Signetics. Gould's PEEL 153 and PEEL 173 use one-third as much power as the bipolar devices and are reprogrammable.

Six edge-activated programmable flip-flops in the PAL22IP6 store signals independently of clock signals, thereby simplifying the design of multi-clock systems and reducing the possibility of metastability—up to four-to-one in a VME bus application. Advanced Micro Devices (Sunnyvale, Calif.) calls the device an Interface Protocol Asynchronous Cell (IPAC). ■

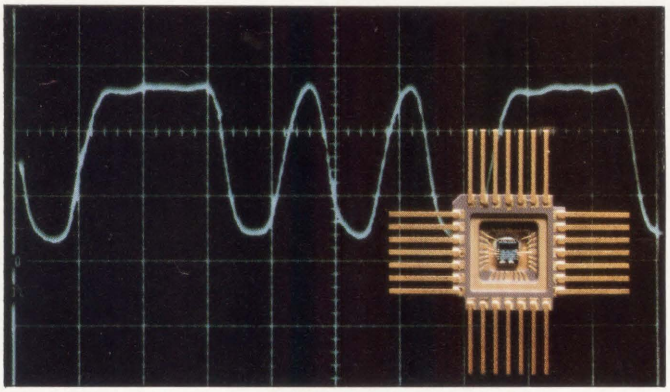

Vitesse Multiplexer Claims Speed Record

The world's fastest four-to-one multiplexer chip, the VS8004, was rolled out by Vitesse Semiconductor Corp. (Camarillo, Calif.). The new IC, which perks along at 2.5 GHz, will be avilable in sample quantities this month. The VS8004 is targeted at the telecommunications industry. The chip's speed will allow its use at the STS-48 level (2.488 Gb/s) of the Synchronous Optical Network (SONET) standard.

The VS8004 is fabricated in the company's proprietary GaAs,self-aligned gate, E/D MESFET process. It has industry ECL inputs and requires only one -5.2V supply. The new IC will use a 28-pin ceramic package. A companion one-to-four demultiplexer is presently under development. ■

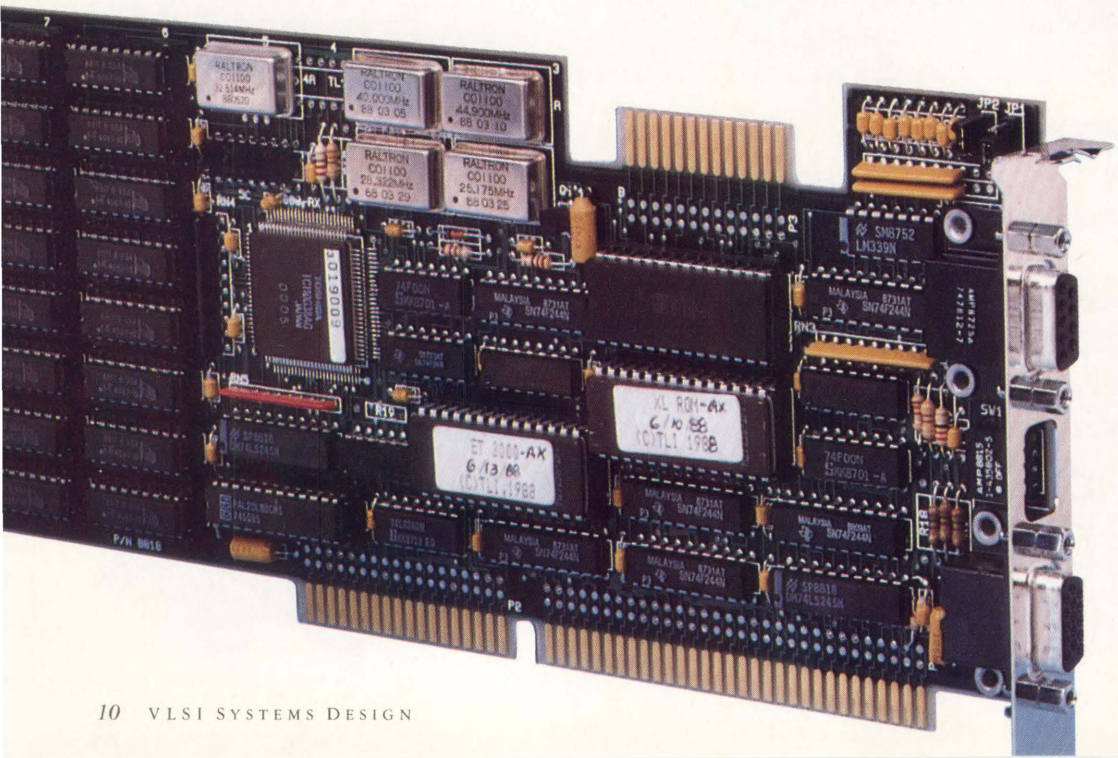

VGA Grows Up

Aydin Controls (Fort Washington, Penn.) announced its Patriot VGA/1024 Graphics board aimed at IBM PC, PC/XT, and compatibles—in addition to the IBM PS/2 Model 30.

Aydin claims that the new video adapter card is 100 percent hardware-registered compatible with the new VGA graphics standard, including CGA, EGA, Hercules, and MDA graphics standards. It incorporates an 8/16-bit data bus, 2/5 CPU/CRTC access, and an extended 1,024 by 768 pixel mode. The VGA/1024, which easily drives Aydin's 20-inch VGA/20 color monitor, is list priced at $695. ■

TI Chips Build Workstations

Texas Instruments (Dallas, Tex.) has developed highly-integrated chips for designers of engineering workstations and computers. First, a set of NuBus interface chips that reduce board area for NuBus interfaces by 66 percent. The 32-bit SN74ACT2440 interface controller and 16-bit SN74ACT2420 address/data transceiver satisfy all requirements for master, slave and combination master/slave interfaces in accordance with IEEE P1196 NuBus specifications.at $24 and $13 respectively, the 2440 provides signalling protocol, including arbitration, bus locking/unlocking, and status bits for cycle control, while the 2420 has three 16-bit ports and a comparator for backplane slot identification.

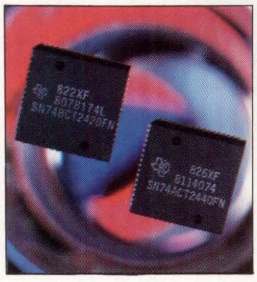

TI also unveiled its second-generation graphics chip, the TMS34020 Graphics Systems Processor (GSP). It will be object-code compatible with the TMS34010, and operate 3 to 20 times faster than its predecessor—largely as a result of its 32-bit address/data bus. Its 512-byte instruction cache and 10-MHz clock provide 10-Mips peak burst rates for small, iterative instruction loops, and it will draw lines at 200-ns per pixel. Samples at $500 are due in the first quarter of 1989. ■

Sun Displays Graphics Choices

Sun Microsystems (Mountain View, Calif.) broadened the range of graphics options for its workstations. Most important to design engineers is the Sun VGA, an EGA/VGA compatible plug-in card for the Sun 386i workstation. The card completes the capability of the workstation to run all engineering tools developed for IBM-compatible personal computers on the 80386-based Sun platform. The $895 card plugs into the system's PC/AT expansion bus and co-exists with the system's frame buffer. As many as four EGA/VGA applications may run at once in a window under the SunOS operating system.

For designers involved in low-end image processing, including animation and publishing, the Sun-4/110TC combines the Sun-4 SPARC microprocessor with a 24-bit, color graphics system which sells for as low as $25,900 (diskless). The desktop system can accommodate as much as 32 Mbytes of main memory (independent of the frame buffer) and 1.3 Gbytes of disk storage; it is rated by Sun at 7-Mips and 0.8-MFlops. ■

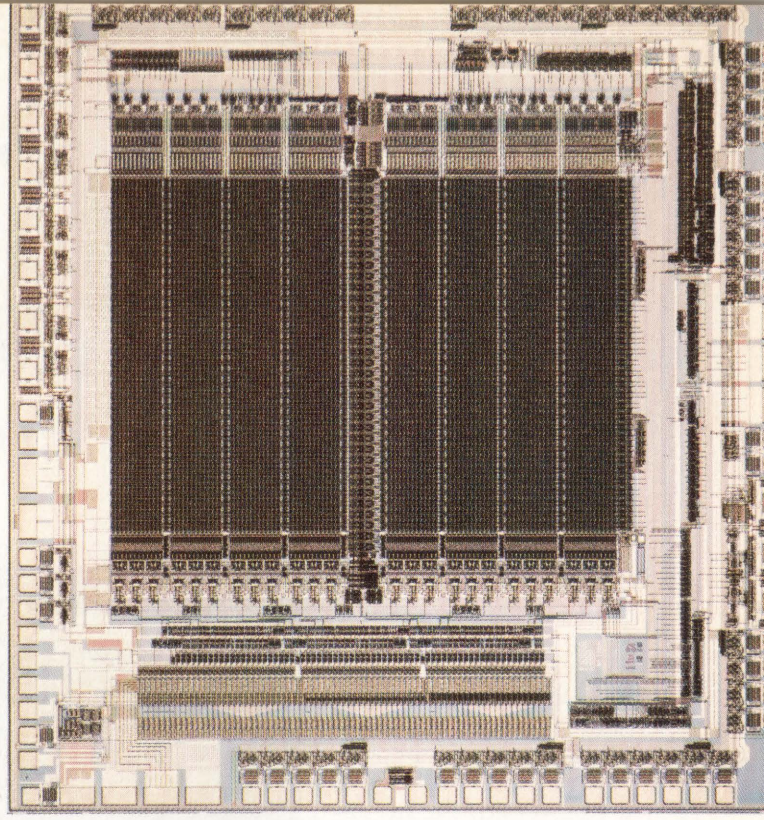

Full Page Display Color Palette Chip

Advanced Micro Devices Inc. (Sunnyvale, Calif.) rolled out its Am81C458 CMOS Color Palette chip, the first in a family of chips that support full page and larger displays. The new chip—which is targeted at image processing, design automation, solid modeling, and animation applications—is offered as an alternate for the Brooktree Bt458 Color Palette chip. AMD's product has also been designed for use in systems that employ the AMD Am95C60 Quad Pixel Dataflow Manager. ■

CV Measurements go DOS

Keithley Instruments Inc. (Cleveland, Ohio) has enhanced the capabilities of its semiconductor capacitance versus voltage measuring system by adding the capability to operate with all DOS-based computer systems. Previous models of the company's Package 82 Simultaneous CV System were only compatible with the Hewlett-Packard 9000 Series 200 and 300 programmable controllers. Included in the system are a Keithley Model 590 CV Analyzer, a Model 595 Quasistatic CV Meter, a ±100V bias source, and calibration sources. Also included are graphics and analysis software for IBM PC/ATs, PS/2s, compatibles, and the Hewlett-Packard Series 200 and 300 controllers.

Helix's creator works at high levels of abstraction

The Creative Urge Fires Up David Coelho

THOSE who witnessed the Fourth-of-July fireworks show in San Leandro may be surprised to learn that the pyrothenic operator behind that spectacle had also created two of silicon valley's high-tech companies. According to David Coelho, licensed pyrotechnic operator and enterpreneur, "I like creating things. That's the driving factor behind everything I do."

It hasn't been easy to fulfill his creative urges. David had to choose between a budding musical career and an aptitude for science and engineering. An active violinist from the age of 2, he was a soloist for his high-school and community orchestras. This self-proclaimed "pretty pragmatic person" hasn't regretted his decision to pursue engineering. Instead, he's focused his creative energy into forming companies with unusual new products.

His diligence in math and computer programming while in high school landed him at Stanford University, where, from his first year, he satisfied his creative needs by pursuing graduate-level research projects. He eventually ended up working for Professor Bill Van Cleemput and, along with two other students, created a program for routing printed circuit boards.

In his senior year (1979) of college, the four of them started Silicon Valley Reseach to create and market a CAD product. David sums up the excitement of that decision by saying, "we had nothing to lose and a lot to gain."

During the start-up phase, David drove himself to the point of becoming ill. Still, he looks back fondly on the creation his first company, saying, "It was a hell of a lot of fun". Unlike many other start-ups in the early 1980's, the company grew out of its own revenues. Luckily, at that time the CAD business was just beginning to "take off like a rocket."

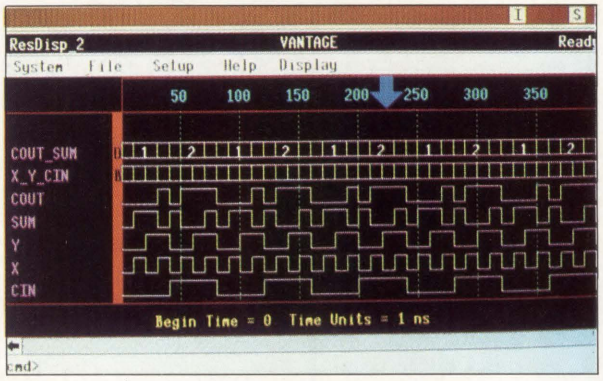

The founders decided they could only be second-tier competitors against the larger, established PCB-CAD vendors. By adapting the software to gate array design, however, they could become leaders in this embryonic market. Their early customers were big companies just starting to experiment with gate arrays, such as Cray Laboratories.

The new company shared an office with a Belgian marketing firm called Leuven Industrial Software Company. It sold standard-cell layout software. Because their product lines were complementary, the two companies merged to form Silvar-Lisco.

David then turned to creating a behavioral simulator based on concepts he had developed at Stanford with Dwight Hill (the two ultimately published a textbook on simulation). His effort resulted in Helix, Silvar Lisco's "architectural-level" simulator.

'I LIKE CREATING THINGS. THAT'S THE

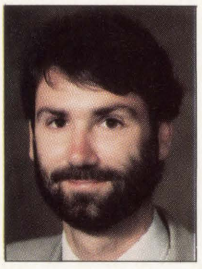

DRIVING FACTOR BEHIND EVERYTHING I DO.'

The Helix approach starts with a good base language (Pascal), and enhances it with features for hardware description modeling. It also adds a user-defined value system for mixed-level simulation capability. This same philosophy was used to define VHDL.

Coelho remained with Silvar-Lisco until April 1986 as manager of simulation products. But, he felt constrained by the larger, more mature Silvar-Lisco, and he was intrigued by an opportunity for VHDL tools that Silvar-Lisco wasn't interested in. He also admits "I'm an enterpreneur at heart".

He left to start Vantage Analysis Systems, which has streamlined software development through heavy use of computer-aided software engineering. Using such tools as Apollo Computer's Dialogue program generator and Metaware's TWS compiler generator, Vantage achieved "a staggering level of productivity," according to David. A team of roughly twelve people took only two years to develop 850,000 lines of code.

He keeps his creative edge by staying close to the technical aspects of Vantage, although he's now active in strategic marketing. He's also kept busy as the chairman of the Vantages' board and as the "leading advocate" for defining new products.

His off-hour interests take him further afield—from setting off fireworks displays to occasional sky diving and scuba diving excursions. It's his creative activities with computers, however, that may light up the sky for users of design automation. ■

—David Smith.

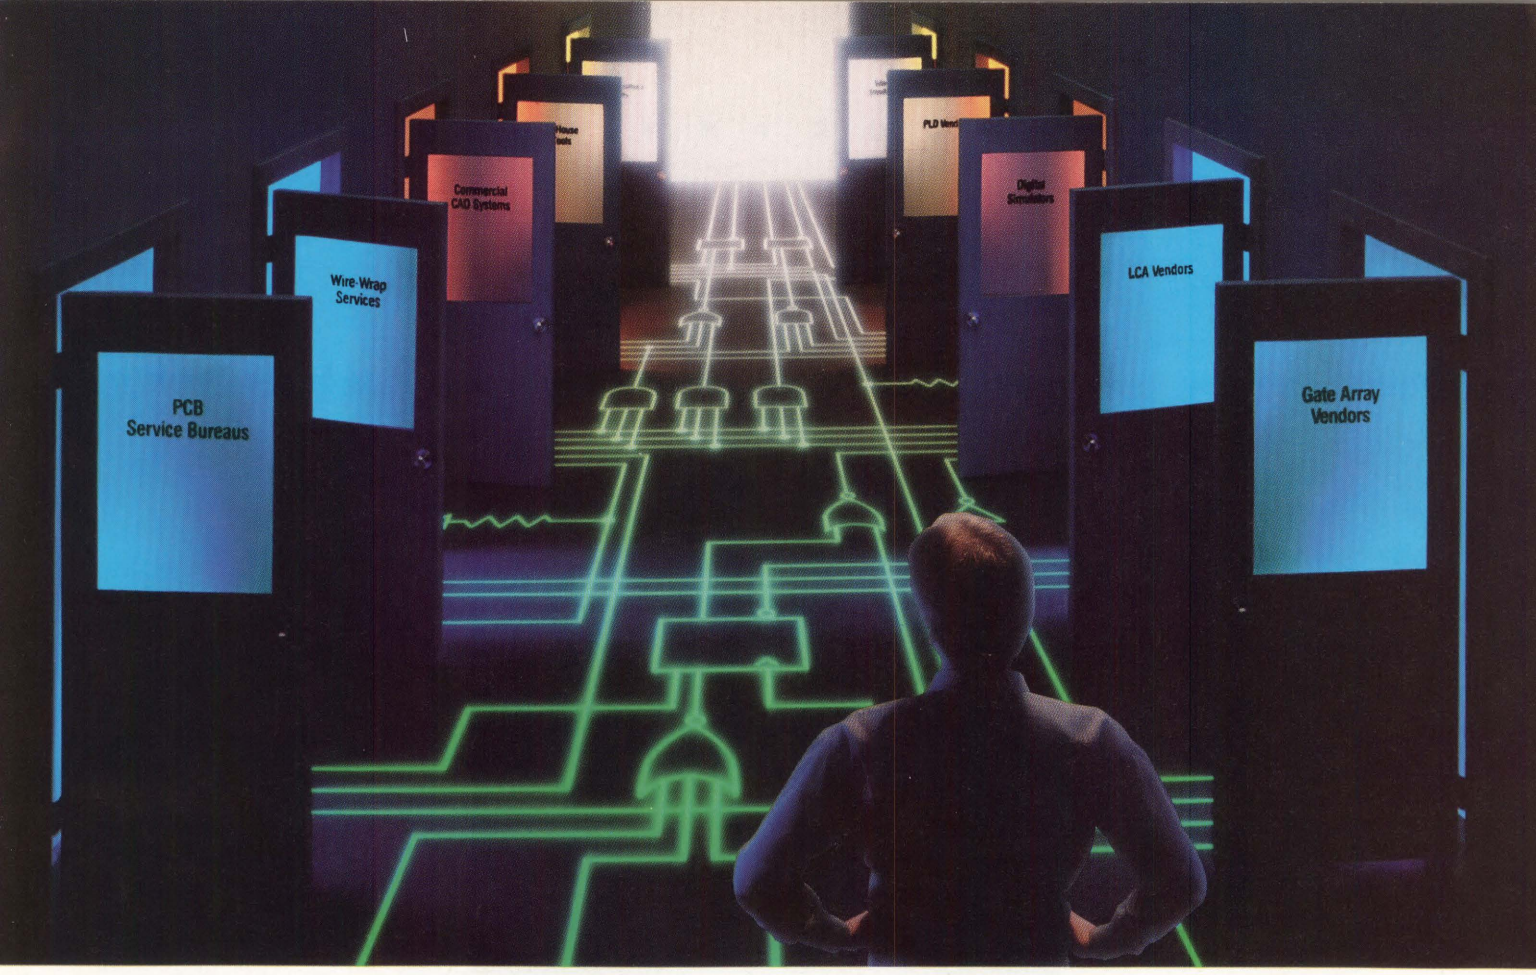

DASH OPENS DOORS TO MORE VENDORS, SERVICES, AND SYSTEMS.

DASH™ Schematic Designer from FutureNet® gives you options no other schematic capture package can. With more than one hundred DASH-Partners providing a broad range of complementary products and services, DASH's industry-standard format is accepted virtually everywhere. So you won't end up with just a schematic. When you design with DASH, you'll have more choices in technologies, CAE systems, foundries, and service bureaus. DASH is a universal front end design tool that has already opened doors for thousands of users worldwide.

AN OPEN DOOR TO ASIC VENDORS. DASH has won the support of dozens of ASIC vendors, including National Semiconductor, Mitsubishi, Motorola, LSI Logic, MMI, Hitachi, and Xilinx. They provide vendor-specific symbol libraries and/or accept DASH netlists for simulation so you can design

DASH runs on 80386 and 80286 machines, IBM® personal computers, and the Sun-3 Series.

ASICs in DASH. Whatever your choice of vendor or technology—PLDs, LCAs, gate arrays, or other semicustom devices—DASH is the schematic entry software of choice.

AN OPEN DOOR TO CAD SYSTEMS. Translators to a wide variety of PCB and simulation systems are available from FutureNet and our DASH-Partners. DASH is the only design entry tool that can be used throughout your company, in multiple CAD environments, regardless of the mix of vendors.

AN OPEN DOOR TO SERVICE BUREAUS. When you don't want to do it all yourself, DASH gives you access to outside services. Service bureaus throughout the world accept the DASH netlist as standard input for their design, wire wrap, and PCB manufacturing services.

FREE DASH-PARTNERS DISKETTE. Take the first step. Call us today for your free copy of the DASH-Partners diskette of vendors, services, and systems that support DASH. You're on the threshold of new design freedom.

1-800-247-5700
Ext. 265

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

FutureNet
A Division of Data I/O Corporation

CIRCLE NUMBER 5

The Demanding Relationship Between ASIC Technology And ATE

The lack of ASIC pinout standards can be costly

RONALD H. LECKE, MEGATEST CORP., SAN JOSE, CALIF.

ASIC market trends have a tremendous impact on the requirements semiconductor manufacturers place on automatic test equipment (ATE). Many of these trends are obvious to most ASIC users, especially technology-dependent trends such as greater IC complexity, density, and pin count. But, commercial aspects of ASIC technology that drive ATE technology may not be apparent to ASIC users—pressures for faster prototype delivery, lower non-recurring engineering charges, and smaller unit volumes.

With ASICs, IC users and manufacturers can no longer treat chips on a part type basis, but must instead look at the whole design-through-shipment process as the product. This process is characterized by many different part types, each with relatively small volume compared to single-type, standard-part processing. To handle the lower volumes, some leading suppliers integate multiple part types on a single wafer using direct write-on-wafer photolithography technology. Having multiple IC designs on each wafer entails lower risk than devoting an entire low-volume run to a single wafer because, with the former approach, one bad wafer does not obliterate an entire run for a design. Unfortunately, it makes wafer-probe testing more complex.

The integration of multiple designs, coupled with higher design density, requires each tester to hold enormous numbers of test patterns. The tests are made even

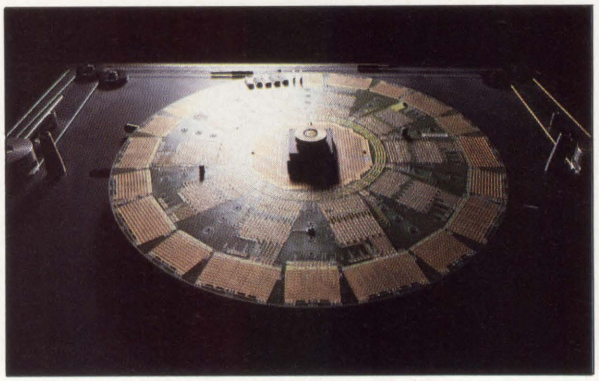

more complex because ASIC wafers depend heavily on wafer monitor structures; the low production volumes don't provide enough feedback on yield parameters. More test monitors are needed per ASIC wafer, and more analysis must be done for each wafer.

Because of the number of designs, test engineers can't pay as much attention to each test program. Test responsibility is placed on the end user or designer. Application of the tests, however, still rests with the ASIC supliers, who need to maintain maximum throughput as well as high test quality. Consequently they need to use larger, more powerful and regretably more expensive VLSI test systems. However, most major players in the ASIC business recognize that *total unit test cost* is more important than minimum capital spending.

There are alternatives to large ATE systems for this problem, namely verification systems and low-cost modular ATE. These approaches, however, compromise other testing requirements, such as overall timing accuracy, flexibility in producing waveforms, and parametric accuracy. ASIC testers must work over a wide range of operating speeds and design complexity, and they must be able to satisfy the highest common denominator. The ASIC industry has not traditionally taken the approach of fully testing devices; for example, many suppliers still use only functional test vectors, and they apply them at low frequencies (typically 1 MHz). Industry-wide drives to higher quality and lower PPM levels, however, make thorough testing, at the intended operating frequency, an eventuality.

'WE MUST LOOK AT THE WHOLE DESIGN-THROUGH-

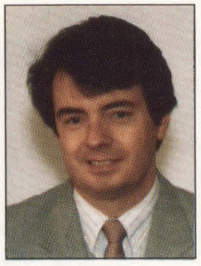

SHIPMENT PROCESS AS THE PRODUCT.'

Another important, but overlooked, ASIC testing problem is the cost of tooling and fixturing. These costs can be very high, especially for cell-based ICs which, unlike gate arrays, have no standard I/O- and power-pad "footprints". Each cell-based chip (including compiled designs) has not only its own test program but also its own probe card for wafer-level test. As ASIC pin counts rise, the cost and manufacturing cycle of the probe cards increases as well, directly impacting prototype cost and delivery cycles.

In addition, package tooling can raise the cost and length of testing. Although some manufacturers have automated the process of fitting die into packages and generating bonding diagrams, few have considered the impact of arbitrary power-pin placement on testing. VLSI test systems can quickly adapt to new I/O pin placement, but test engineers must manually allocate and decouple power connections either on the performance board or at the handler contactor set.

These tooling problems could be alleviated if ASIC manufacturers would trade-off silicon efficien-

Simple as...

pushing a button

DESIGN ENGINEER	ASIC
CSPICE	PLD
HILO	LAYOUT

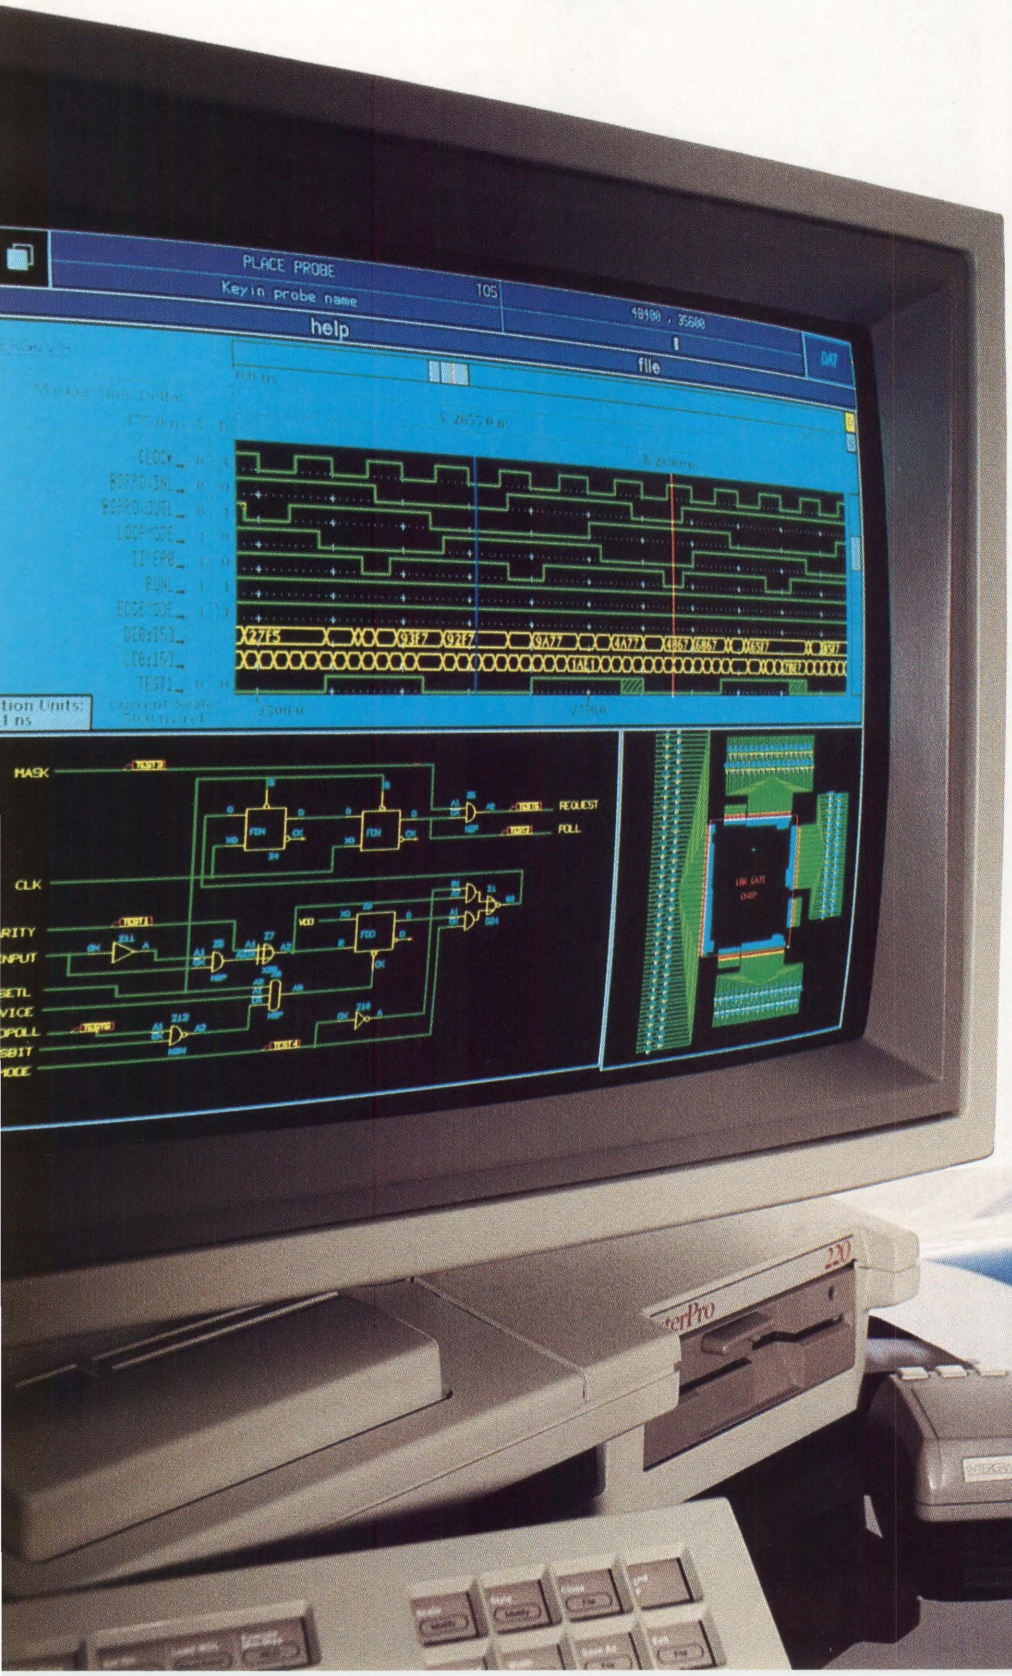

Intergraph makes ASIC design easy. We provide a complete solution for ASIC development to help you create optimized and verified electronic circuits.

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

- ☐ hierarchical schematic design
- ☐ online electrical rules checking
- ☐ leading ASIC model libraries
- ☐ user-generated custom libraries
- ☐ high-speed logic and fault simulation

Resulting verified netlists can be used with a wide selection of physical design tools that Intergraph offers for standard cell, gate array, programmable gate array, and programmable logic device designs.

The ASIC Engineer Series offers schematic symbols, simulation models, netlisters, back-annotation utilities, and simulation vector conversion utilities.

Our 5-MIPS CLIPPER workstations and servers give you the performance to produce advanced ASIC designs with lower part cost and greater architectural flexibility – to get your products to market faster.

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

CIRCLE NUMBER 6

INTERGRAPH

Wafer probing and package fixturing play a major part in total ASIC costs.

cy for *test-fixturing efficiency*. Most suppliers have extensive design-for-testability rules relating to the quality of circuit functionality and performance, but how many have testability design rules controlling the die footprint and packaging parameters? It may be worth sacrificing 10 to 20 percent of the silicon area to gain some footprint standardization to minimize tooling costs for probing, bonding and packaged-IC testing.

ASIC devices are pushing tester technology to satisfy their requirements for more flexibility and lower manufacturing cost. These requirements can be met through close cooperation between ASIC manufacturers and test equipment manufacturers. After all, it is in the best interest of designers of test systems to help the ASIC manufacturer who is not only a customer, but also a supplier of components! ■

RONALD H. LECKE *is Vice President of Engineering for VLSI Test Systems at Megatest Corporation, San Jose, Calif. Prior to joining Megatest, he worked at Signetics Corporation in Sunnyvale, CA, serving as CMOS Product Engineering Manager, Quality and Reliability Manager, and Product Engineering and Test Operations Manager for the Bipolar LSI & Semicustom Division. He had transfered to California from Signetics' facility in Linlithgow, Scotland, where he worked in product engineering, test, yield improvement and customer applications. He received a BS in Electrical and Electronic Engineering from Heriot Watt University in Edinburgh.*

CALENDAR

Continued from page 8

this conference has been designed to review current materials research that is of common interest to academia, industry, and government. Topics will include interfacial phenomena in electronic and non-electronic materials, metal and ceramic matrix composites, high-speed electronic devices, and superconductivity. Authors should submit, by November 1, 1988, a one-page abstract to Dr. Jerome G. Morse, Director, Advanced Materials Institute, Colorado School of Mines, Golden, Colo. 80401. (303) 273-3852. ■

PARLE '89

June 12–16, 1989
Eindhoven, The Netherlands

PARLE '89, the second conference on Parallel Architectures and Languages Europe, is intended to serve as a meeting place for researchers in the fields of theory, design, and applications of parallel computer systems. Papers are solicited on such topics as semantics and models for parallelism; programming environments; memory management, fault tolerance, and real-time aspects; dedicated processors for AI; debugging and monitoring tools; and simulation. By October 21, 1988, interested authors are invited to submit 5 copies of a full paper that does not exceed 6,000 words to Dr. M. Rem, Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands. ■

9TH INTERNATIONAL SYMPOSIUM ON COMPUTER HARDWARE DESCRIPTION LANGUAGES

June 19–21, 1989
Washington, D.C.

The goal of CHDL '89 is to bring together, from both industry and academia, the leading researchers and the best work in description languages. Papers are invited on such topics as optimization techniques, silicon compilation, EDIF, VHDL, design for verification, high-level simulation, advanced software simulation, graphic tools, databases, and industry methodologies. By October 31, 1988, interested authors should submit five copies of a manuscript, not exceeding 20 pages, to Professor Dr. Franz J. Ramming, Program Chairman, Universitat Paderborn, Warburger Str. 100, D-4790 Paderborn, Federal Republic of Germany. ■

Z I L O G

Big 8-bit performance. Tiny 18-pin package. Miniscule price.

Introducing the smallest Z8 microcontroller yet, the high-performance Micro8™ (Z86C08). Just imagine what you'll be able to do with all this capability, packed into a DRAM-sized chip. For what amounts to pocket change.

The Z8 Family: Still setting the standards.

From the first, the Z8 microcomputer has been one of the industry's finest examples of simplicity and elegance. The most sophisticated microcontroller family continues to grow.

Today, there is Z8 support for every phase of your system development, from prototying to full production. Along with devices you know and trust, there are new parts with an increasing number of options and features for your design. Recently, the Z8 is found in such diverse uses as fans, puppy door controllers, induction hot plates, and high performance hard disk controllers, printers, and local area networks (LAN).

Since demand for Z8 products, and their importance to the industry continues to increase, we are developing new members of the Family. The Super8™ is now clearly established as the high-end Z8, and the Micro8 can be expected to have as much an impact on low-end Z8 applications. Among the other new Z8 Family members you should be keeping an eye on are:

- *Z86C91 High-performance CMOS ROMless microcontroller*
- *Z86C10 Low cost 28-pin CMOS, has 22 I/O lines and 4K bytes of on-board ROM*
- *Z86C21 8K ROM Z8, has 32 I/O lines, 2 levels of security*

Mighty strong performance

First of all, the Micro8 features the high-end Z8 architecture. Then there's 128 bytes of RAM, two counter/timers, two single-supply analog comparators, and low power consumption. Not to mention all the advantages of Superintegration™ and CMOS technology.

Mighty powerful protection

The Micro8 may be tiny, but it's as bullet-proof as they come. You get brown-out protect and a watch-dog timer, for instance. You get an operating range of 3-5.5V. And you get CMOS I/O levels and hysteresis for noise protection.

Mighty impressive bit bang for your buck

The Micro8 gives everything you want in an 8-bit microcontroller. In the smallest package you've ever seen. For about a buck and a half. Plus you're working with the familiar software and proven performance of the Z8 Family. And it's all backed by Zilog's solid reputation for quality and reliability.

You really ought to see for yourself just what the mighty Micro8 can mean to your design application. Why wait? Contact your local Zilog sales office or your authorized distributor today. Zilog, Inc., 210 Hacienda Avenue, 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) 885-8080, **MA** (617) 273-4222, **MN** (612) 831-7611, **NJ** (201) 288-3737, **OH** (216) 447-1480, **PA** (215) 653-0230, **TX** (214) 231-9090, **CANADA** Toronto (416) 673-0634, **ENGLAND** Maidenhead (44) (628) 39200, **WEST GERMANY** Munich (49) (89) 612-6046, **JAPAN** Tokyo (81) (3) 587-0528, **HONG KONG** Kowloon (852) (3) 723-8979. Taiwan (886) (2) 731-2420, **U.S. AND CANADA DISTRIBUTORS:** Anthem Electric, Bell Indus., Hall-Mark Elec., JAN Devices, Inc., Lionex Corp., Schweber Elec., Western Microtech. **CANADA** Future Elec., SEMAD.

CIRCLE NUMBER 8

VERSATILE

Broadband Analog IC

JOHN ADDIS, TEKTRONIX, INC.,
BEAVERTON, ORE.

A high-speed bipolar analog integrated circuit is the heart of four plug-in amplifiers used in the Tektronix 11000 series oscilloscopes. This IC, internally designated the M377, forms almost the entire signal path in three of the plug-ins, and a majority of the fourth. With over 700 transistors, it is almost a "plug-in on a chip". A microprocessor and custom logic IC are used to control the M377 and add a sophisticated calibration routine.

ASIC IS A PLUG-IN ON A CHIP

The M377 IC features: gain switching over a 50:1 range in a 1,2,5 sequence of six discrete steps; a continuously variable gain that is proportional to a dc input voltage; multichannel operation with other M377s; three identical outputs that can be independently inverted or enabled in less than 200 ns; a bandwidth of 800 MHz with 420 ps risetime (for gains between 0.4 and 12) and a bandwidth of 320 MHz at gain of 60; a four-pole, 100 MHz bandwidth limiting filter; a four-pole 20 MHz bandwidth limiting filter; a high common mode rejection differential input configuration; and overdrive recovery to within 0.04 percent in 6 ns.

To dissipate its three watts, the 4.32mm (170 mils) X 2.92mm (115 mils) plug-in amplifier chip is mounted on a 1.22 cm (0.480 inch) square thin film ceramic substrate. No bypass capacitors, resistors, or inductors are required, just a conductor pattern and the monolithic IC on a leadless substrate. The entire amplifier is connected to an etched circuit board using the company's proprietary patented elastomeric connector.

■ A PHILOSOPHY CHANGE

The M377's circuitry represents a major departure from earlier wideband oscilloscope amplifiers. Changing industry needs, both for increased precision at high frequencies and for lower assembly costs, dictated a fresh approach. The 1 GHz bandwidth of Tektronix's 7104 oscilloscope (introduced in 1979) was found to be wide enough to satisfy all but a few high frequency needs. However, there were increased requirements for high precision over these bandwidths. The advent of 10- to 16-bit analog to digital and digital to analog converters as well as digital signal processors was part of the driving force behind high-precision, high-frequency oscilloscopes.

The cost of the assembly labor for complex instruments such as oscilloscopes was rapidly becoming too high. Designing more of the complex circuitry into ASICs resulted in decreased labor costs for digital circuit board assemblies, and there was no reason that the same philosophy shouldn't work for analog circuits as well. Board space was another factor which dictated an increased level of integration. A four-channel plug-in was required, but there wasn't enough room on the plug-in circuit board to accept four channels of conventional analog circuitry.

The labor costs for calibration and adjustment were also expensive. This dictated circuit designs that reduced the number of manual adjustments. The entire 11000 oscilloscope calibration is automatic. Only one manual adjustment per channel exists in most of the plug-ins that use

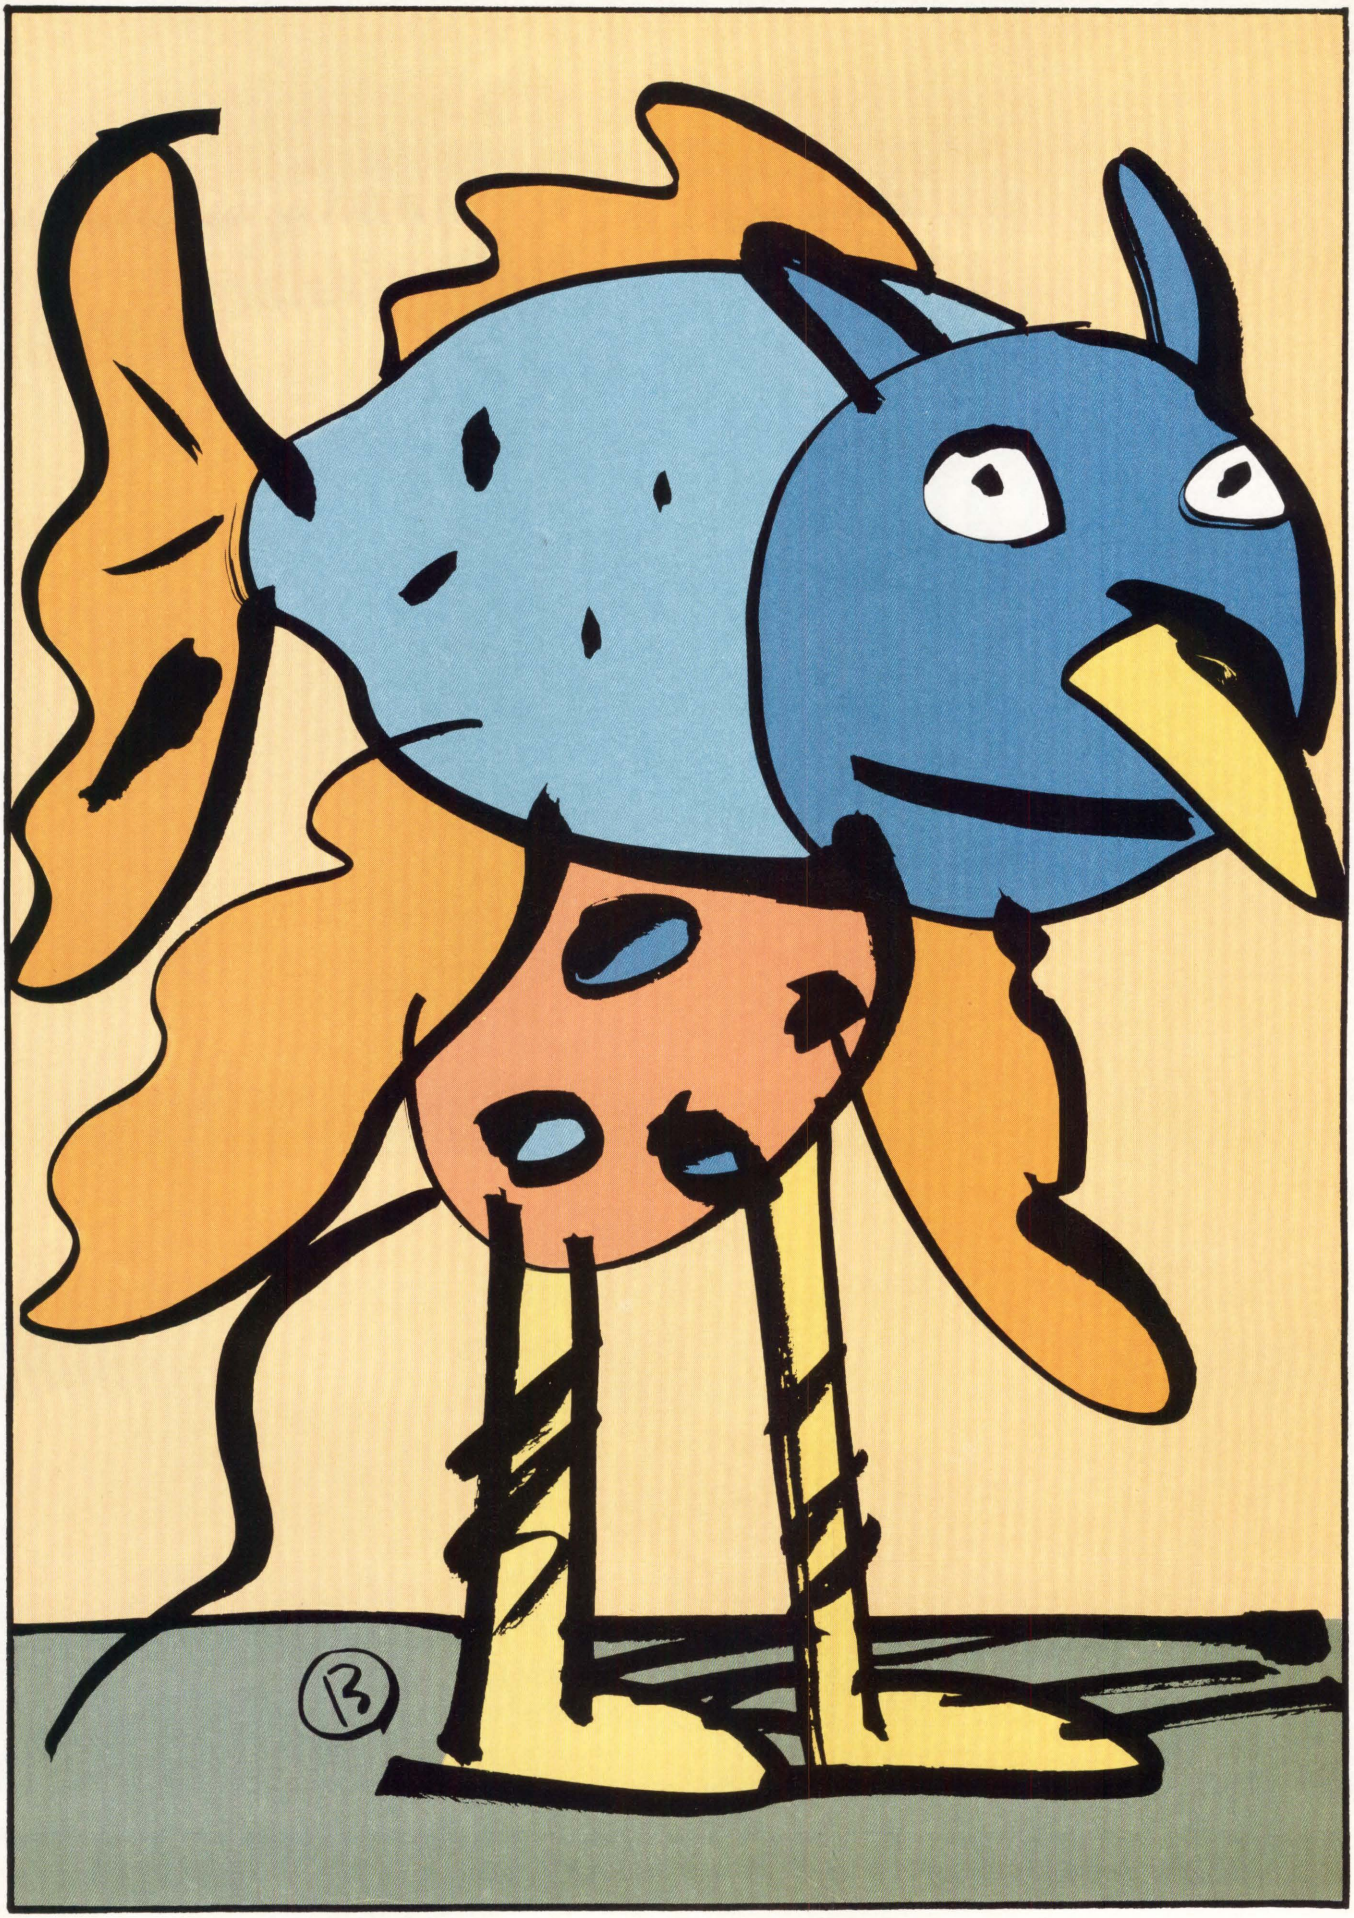

ILLUSTRATION BY PATRICK BLACKWELL

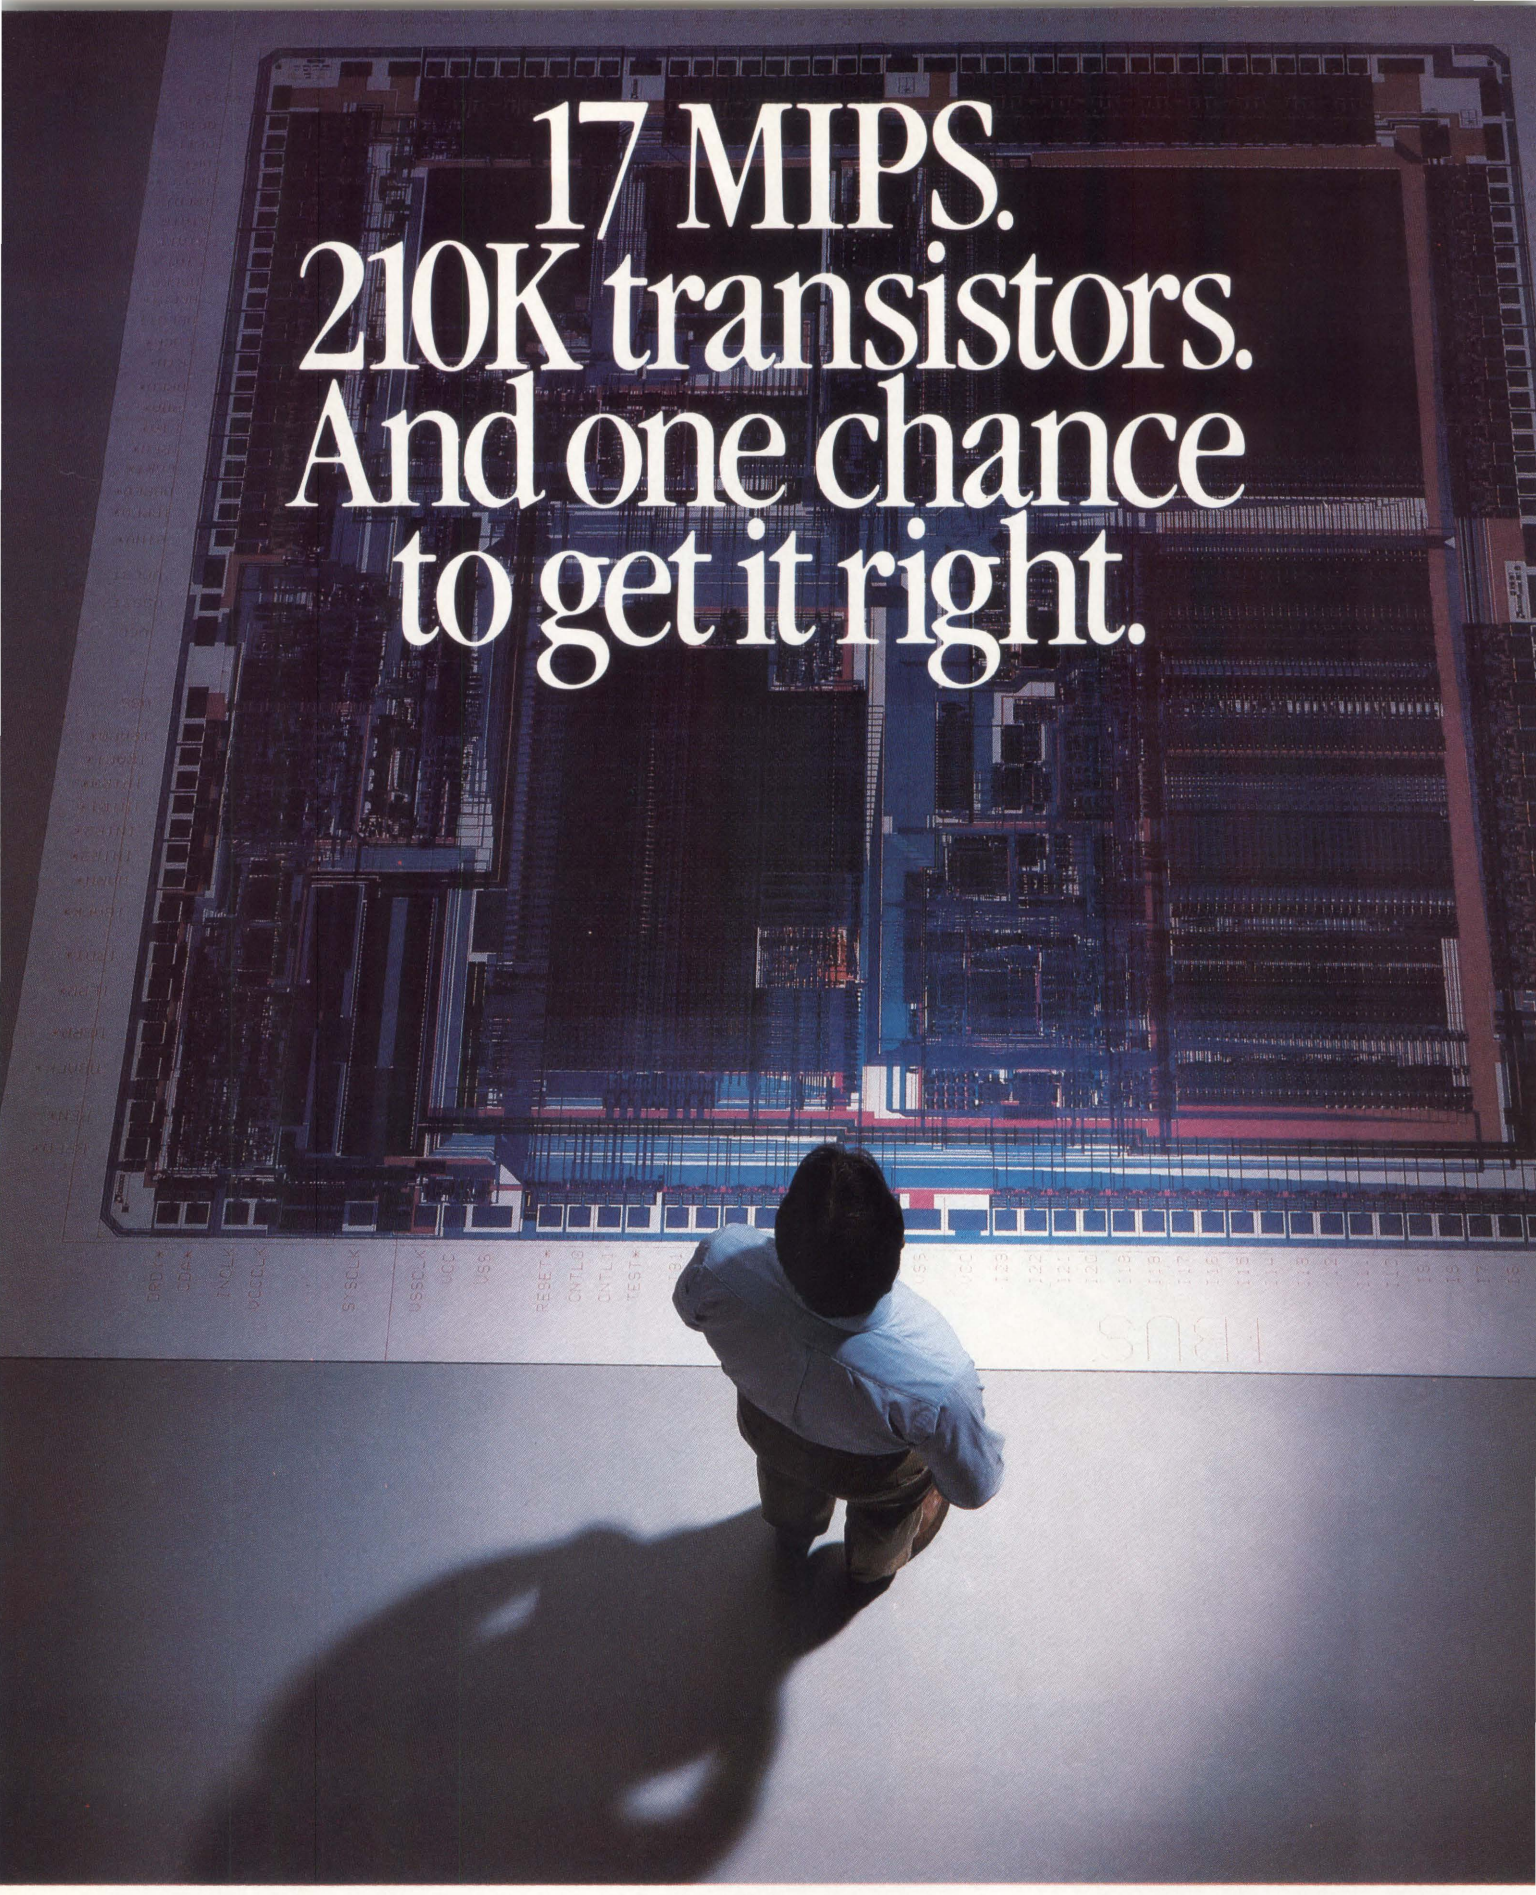
17 MIPS.
210K transistors.
And one chance
to get it right.

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

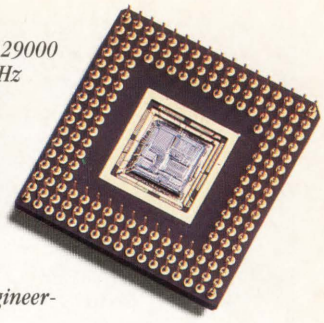

AMD's new RISC-based, 32-bit Am 29000 microprocessor operates at a 25 MHz clock rate with a 40 ns instruction cycle time. It can hit a peak execution rate of 25 MIPS, with a sustained performance of 17 MIPS. Because of its exceptional speed and price-performance ratio, the Am 29000 can be used in a wide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

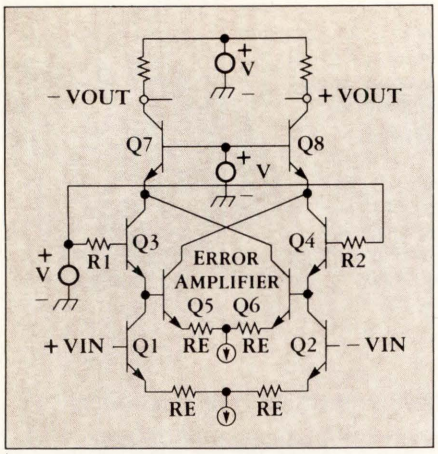

Figure 1. The cascomp (compensated cascode) amplifier was one of the first attempts at broadband precision amplification.

the M377 IC, and that adjustment is for transient response. Even the frequency compensation for the high impedance input attenuators is factory-adjusted by laser trimming under machine control.

■ THE CASCOMP AMPLIFIER

Greater precision and fewer adjustments are not necessarily mutually exclusive goals. Indeed, they are compatible goals if the right circuit design choices are made.

The Tektronix 2465 oscilloscope was the first instrument to use this approach. It uses a circuit we call the cascomp (compensated cascode) shown in Figure 1. If the emitter-coupled differential pair is a first generation amplifier, then the cascomp can be considered a second generation amplifier. It is the first attempt at broadband precision amplification, and an improvement on the simple differential pair.

The cascomp senses the error voltage due to the non-linear (logarithmic) junction characteristics and subtracts the error from the final output. The cascomp does not sense the actual error voltage in Q1 and Q2, but instead senses the nearly identical error voltage generated by Q3 and Q4. Q7 and Q8 are added to provide good error amplifier bandwidth. For the same quiescent current, the cascomp demonstrates increased linearity over the uncompensated circuit.

But there is another source of error plaguing the first generation dc-coupled linear amplifier: thermal effects commonly called "thermals" or "thermal tails". These effects occur because the each transistor's operating point and hence power dissipation changes with signal. The power dissipation change causes the transistor's temperature to change, which in turn changes the base-emitter voltage. Since the base-emitter voltage is in series with the input signal, there is a thermally induced error in the amplifier's transient response. The effect can be substantial (5 to 7 percent per stage in extreme cases of very wideband and therefore power hungry amplifiers), and it is especially troublesome because it results in the multiple time constants characteristic of heat flow through silicon. Discrete amplifiers usually have "thermal balance" resistors (between Q1 and Q3, and Q2 and Q4) which nearly eliminate thermals, but integrated circuit amplifiers cannot provide the large capacitors necessary for bypassing these thermal balance resistors. Fortunately, the effects are usually linear and can be compensated. However, first generation IC amplifiers require a large number of manual adjustments to compensate for signal-related thermal effects. For example, a 7104 with two single channel plug-ins required 32 manual adjustments to correct for thermals.

If the operating points of Q3 and Q4 are chosen carefully, the cascomp cancels thermal errors as well as non-linear errors. By eliminating the need for manual adjustments, the parts costs, board area, and labor costs associated with calibration are all reduced. The calibration labor costs are non-trivial because the adjustments for thermals are among the most difficult and frustrating to make.

The patented (Pat Quinn of Tektronix) cascomp is successfully used in many of the company's oscilloscopes today.

While an improvement over a simple differential pair, the cascomp has some limitations. First, very high gain is not possible from a single stage without loss of all its advantages. Since the error amplifier and the main amplifier have the same gain, non-linearities in the error amplifier become appreciable at high gain (at very high gain, the error amplifier is not able to cancel non-linearities very predictably).

Second of all, the stack of three transistor pairs in the signal path causes the gain to be 1.5 times as sensitive to alpha losses as the standard cascode. As temperature increases, the gain increases as $(alpha)^3$. This amounts to about 225ppm/°C in the cascomp for a typical beta of 80 versus 150ppm/°C for a standard cascode. The standard cascode also has a counteracting temperature dependent gain term in the emitter circuit due to the dynamic emitter resistance which the cascomp has eliminated. This term is about -185ppm/°C for an emitter current of 20mA/side at a junction temperature of 60°C and using a 40 ohm emitter resistor. To compensate for alpha effects, two resistors, R1 and R2, may be added. However, it is usually not possible to make R1 and R2 large enough to compensate for the alpha loss without creating too much high frequency peaking. Thus the standard cascode may have a lower gain temperature coefficient than a cascomp circuit.

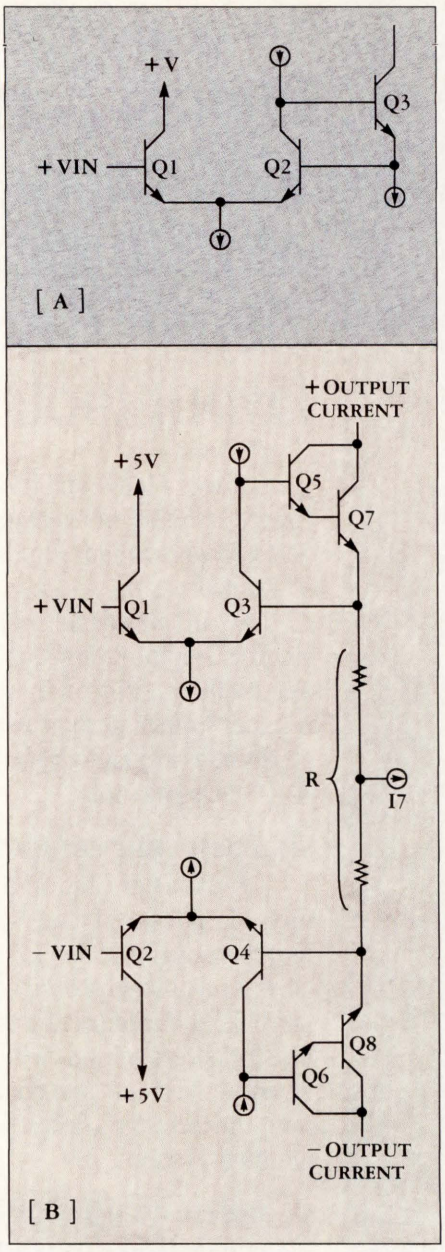

Figure 2. The basic amplifier can be considered as a compound transistor (a), however, changing the output transistor to a Darlington (b) improves the gain stability.

Third, the cascomp has a disadvantage to the IC designer in that the three stacked devices use up more of the available supply voltage. In circuits which require level shifting back down toward the negative supply, the voltage shift required is greater with the cascomp than the cascode.

A fourth disadvantage is in the cascomp's ability to handle overdrive signals. Since the error generating devices (Q3 and Q4) do not see the full input signal, the error correction circuitry and the main amplifier generally have different thermal

histories in overdrive. However, with some extra circuitry, thermals resulting from overdriving the cascomp can be made about as good as a thermally balanced differential pair.

■ THE BASIC AMPLIFIER

The M377 represents a third generation broadband amplifier design. The basic amplifier is shown in Figure 2a. This circuit—whose original design goes back at least to the mid-seventies and is a variation of the LM102 operational amplifier—is a very versatile feedback amplifier. It can be viewed (somewhat imperfectly) as a compound transistor in which a differential pair, Q1 and Q2, compares an input signal with the emitter voltage of an output device, Q3. The compound device has increased gm and beta over Q3 alone when operating at the same current.

When viewed as a compound transistor, it is obvious that the output can be taken from either the emitter or collector of Q3. The analogy with a compound transistor falls apart because the alpha of Q3 is not helped by Q1 and Q2. This flaw can be overcome in several ways. For example, the collector of Q1 could be connected to the emitter of Q3. This not only increases the compound transistor's alpha, but it bootstraps Q1's collector. However, the operating point of Q1 will compromise its F_t. Additionally, at very high frequencies this design has a negative input impedance, which is potentially unstable. The best technique for high frequency designs is to change Q3 into a Darlington as in Figure 2b and accept the lost bootstrapping advantages as the price for stability and a 1 GHz bandwidth.

When two such voltage followers are connected together as in Figure 2b, the configuration is that of an instrumentation amplifier—the basic amplifier stage used in the M377.

A block diagram of the M377 would show only two stages of gain. In between are a level shift, a Gilbert multiplier variable gain control, and a choice of two bandwidth limiting filters or a full bandwidth path. The input stage and the three identical output stages are all instrumentation amplifiers whose inputs are the bases of Q1 and Q2 and whose outputs are the darlington collectors (Figure 2b). Gain for these stages is set by the resistance R between mirror-imaged voltage followers. To a very high degree of precision, the signal current output is equal to the voltage input between the Q1 and Q2 bases divided by R.

■ ELECTRONIC GAIN SWITCHING

The ability to change gain over a wide range in precision steps is important for two reasons in oscilloscope preamplifiers. First, it allows the input attenuator to be simplified to just three settings, unity, ten times attenuation, and one hundred times attenuation. This is significant because the input attenuator, which is built with electro-mechanical relays, is inherently expensive and less reliable than well-designed solid state components. The fewer the components, the more reliable the attenuator will be. Second, the gain changing technique employed should use the best trade-off between gain and noise. For example, if the amplifier needs a maximum sensitivity of 1 mV per division, use of a passive attenuator before the amplifier as the only means of altering the amplification would result in the high gain amplifier's full noise level being displayed at all sensitives. Use of a passive attenuator after amplification increases the dynamic range requirements for the pre-attenuator amplification. Even at 10:1 attenuation, the cost of increased dynamic range can be severe.

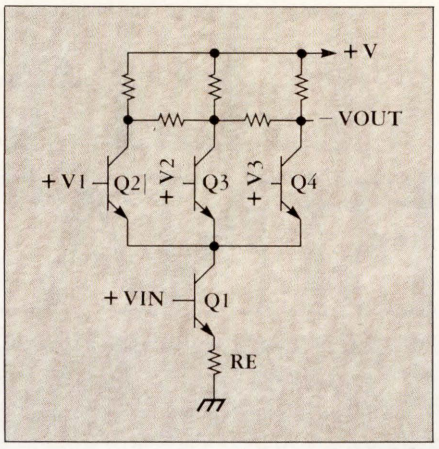

Figure 3. Using a passive constant resistance network provides gain changes with minimum effect on bandwidth.

There are several ways of changing the amount of amplification in an IC. One method is to use a passive constant resistance network as shown in Figure 3. The current in Q1's collector will result in a different Vout depending upon whether Q2, Q3, or Q4 is conducting. This method can have very broad bandwidth and has the advantage of minimal change in bandwidth as a function of the attenuations selected. Used alone, it has the disadvantage of requiring the amplifier preceding it to handle the full dynamic range of the lowest gain setting. The input amplifier cannot simultaneously have high gain and wide dynamic range without prohibitive power dissipation, so it must have low gain to handle the largest expected signals. This requires the post attenuator amplifiers to operate at high gain with correspondingly high output noise.

Some means of controlling the input amplifier's gain offers a way around this dilemma. For example, increasing the input transistor's emitter resistor will allow the amplifier to handle large signals at low gain settings, while decreasing the emitter resistor boosts the gain enough to minimize noise from subsequent stages at high gain settings. The disadvantage of this approach is that the bandwidth will change from one gain setting to another.

In the M377, discrete gain steps are obtained by changing the first stage gain setting resistor (as shown in Figure 4). TTL logic inputs select one of six current sources such as I1 or I2 which force the first stage quiescent current through the appropriate gain setting resistors and associated diodes. (For clarity, only two gain settings are shown in Figure 4). Diodes D1 and D2 or D3 and D4 close the feedback loop. The idle diodes are back biased by the appropriate 100K resistor. The diodes must withstand the brunt of the full input signal swing. Since the diode-connected high frequency transistor will punch through at about 1.5V reverse bias, Schottky barrier diodes, which can withstand reasonable voltages, are used. Two 6.7K nichrome resistors are used in place of active current sources to reduce noise.

In the linear range, thermal effects are quite small in this amplifier (when compared to a simple differential pair) because the large quiescent current and current swings necessary to produce an output signal are handled by Q7 and Q8. Thermals in these devices are inside the feedback loop and are reduced about 250 times by the voltage gain of Q1 and Q3 or Q2 and Q4. The operating points of Q1, Q2, Q3, and Q4 are virtually unaffected by input signals. The voltage change on these devices is small not only because the input signal is small, but the current change is also small because Q5, Q6, Q7, and Q8 have such high current gain. In fact, Q1, Q2, Q3, and Q4 barely change their operating points up to the level at which the signal cuts off either Q7 or Q8. Tight thermal coupling of Q1, Q2, Q3, and Q4 help to keep the remaining linear (as opposed to overdrive) thermals well below the 0.05 percent level. However, once Q7 or Q8 cuts off, the four input devices quickly go into the overdrive range. Taming the resulting overdrive thermals is a unique accomplishment of the M377.

■ FAST OVERDRIVE RECOVERY

Typically, when an amplifier is overdriven, it takes a reasonable amount of time to reestablish its operating point. Unless damage or destruction takes place,

EEPROM/D
COMBINATION F

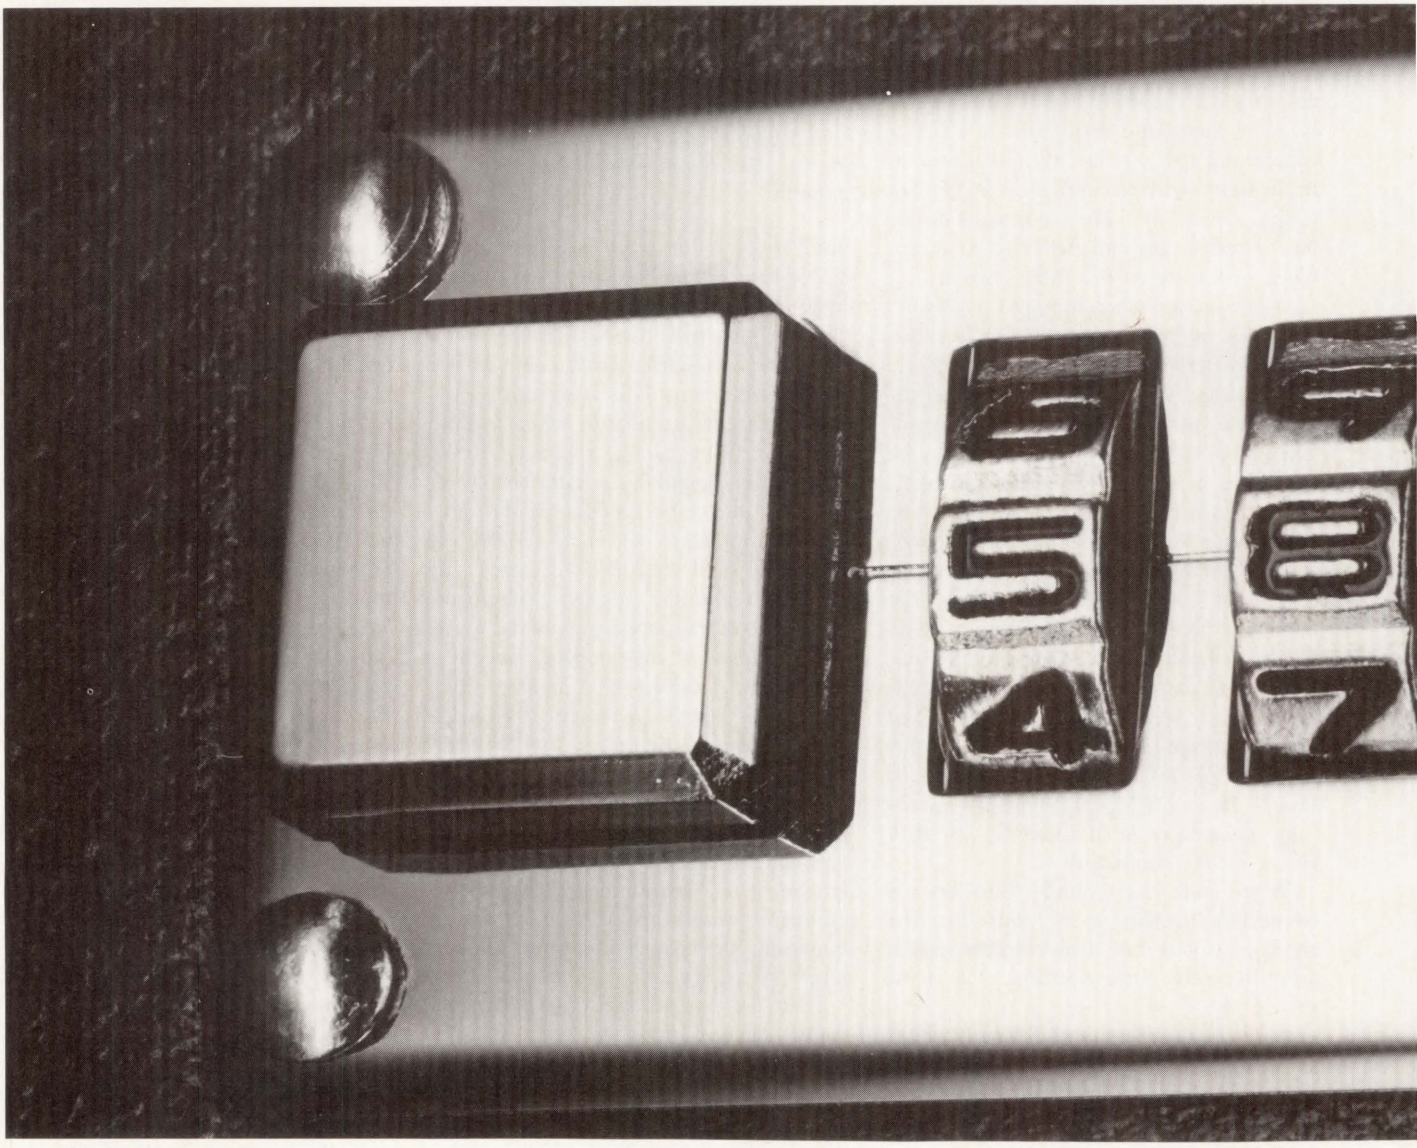

There are literally thousands of combinations of LAN, security, and industrial applications where you have to integrate EEPROM and digital technologies. Now, you can combine them with the confidence and security that they'll work together, since they come from the same proven source — Sierra.

Our CMOS cell library contains over 20 EEPROM cells, which can be combined with our digital cells in thousands of combinations. So you can get exactly the functionality you need.

Just to give you one example, you can design a software switch that replaces DIP switches in setting board/ID configurations. Whether you're designing a mother board, add-in board, or LAN card, this lets you cut your parts count way down. And give your customers an easier product to use.

Of course, with a library of over 300 cells, we don't limit you to EEPROM/digital combinations. We make it just as easy to combine analog and digital, EEPROM

IGITAL. THE OR JOB SECURITY.

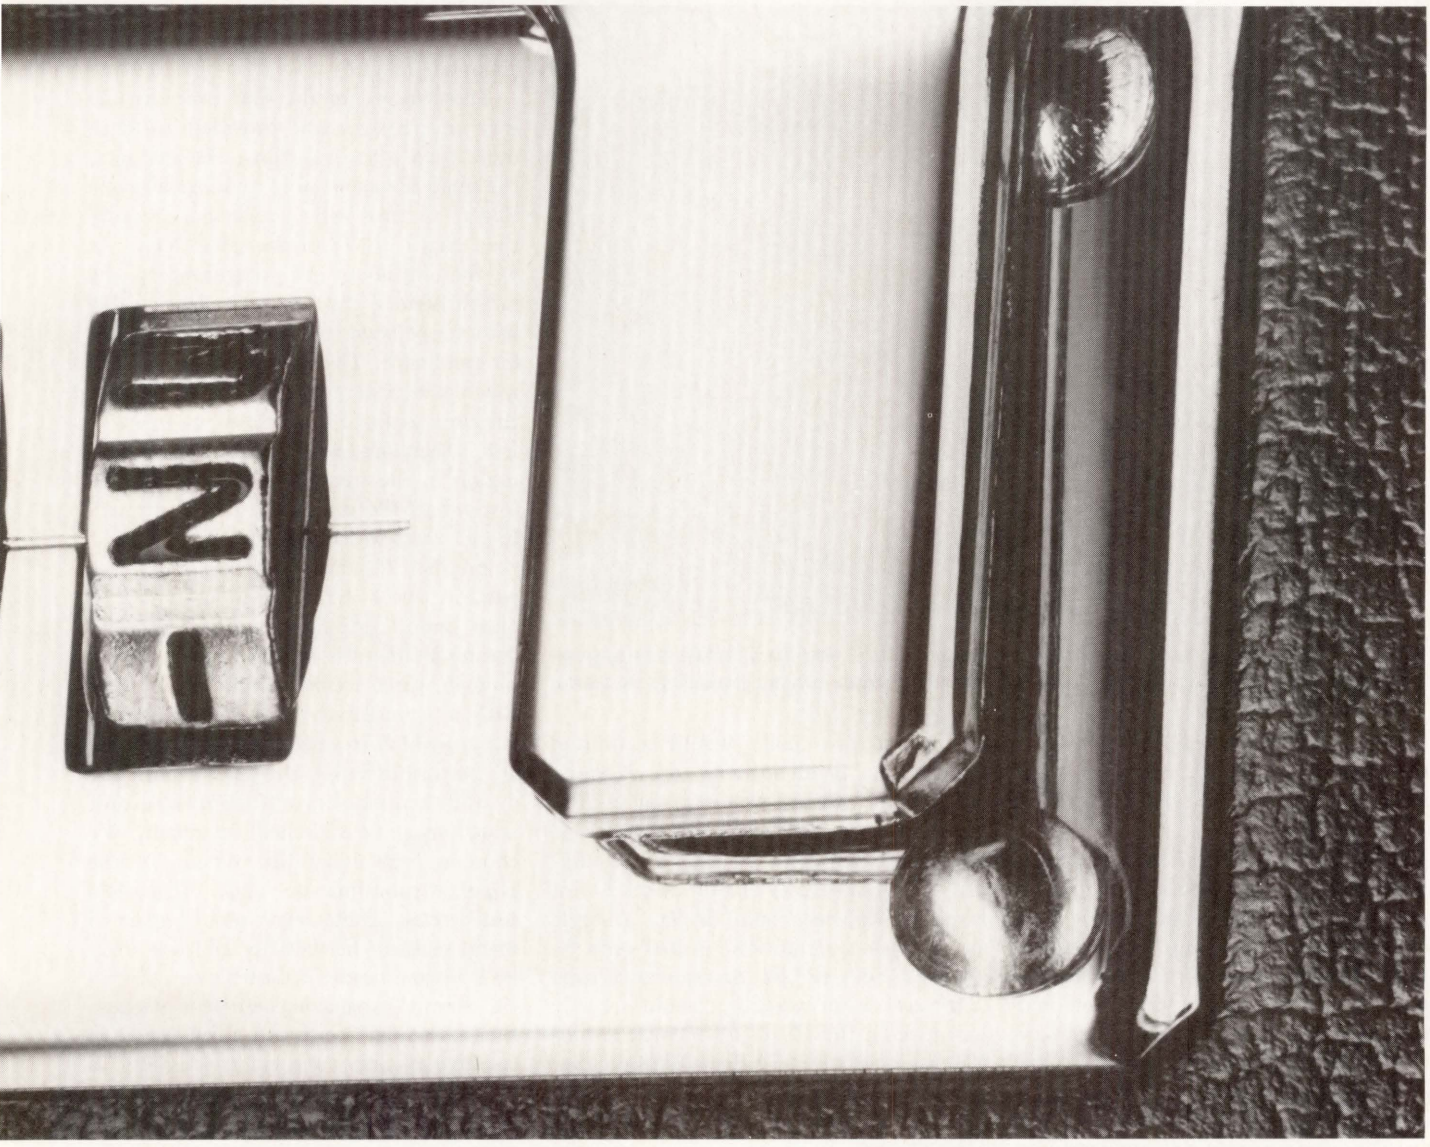

and analog, or all three on the same chip. Which is the whole idea behind our Triple Technology.™

Not only is that a Sierra exclusive, we also give it to you in CMOS. So you can look for reliable, low-power 5V operation. We give you the most efficient packaging, too. You can choose from low profile surface mount packages in SOIC, PLCC, and PQFP gull-wing versions.

As if that weren't enough, we also give you the kind of training and documentation that make ASIC solutions easier to design than you ever thought possible. Just write or call for our complete library card. And we'll show you how to easily secure your next application. All you need is the right combination.

Sierra Semiconductor
Triple Technology.™ In CMOS.

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

CIRCLE NUMBER 9

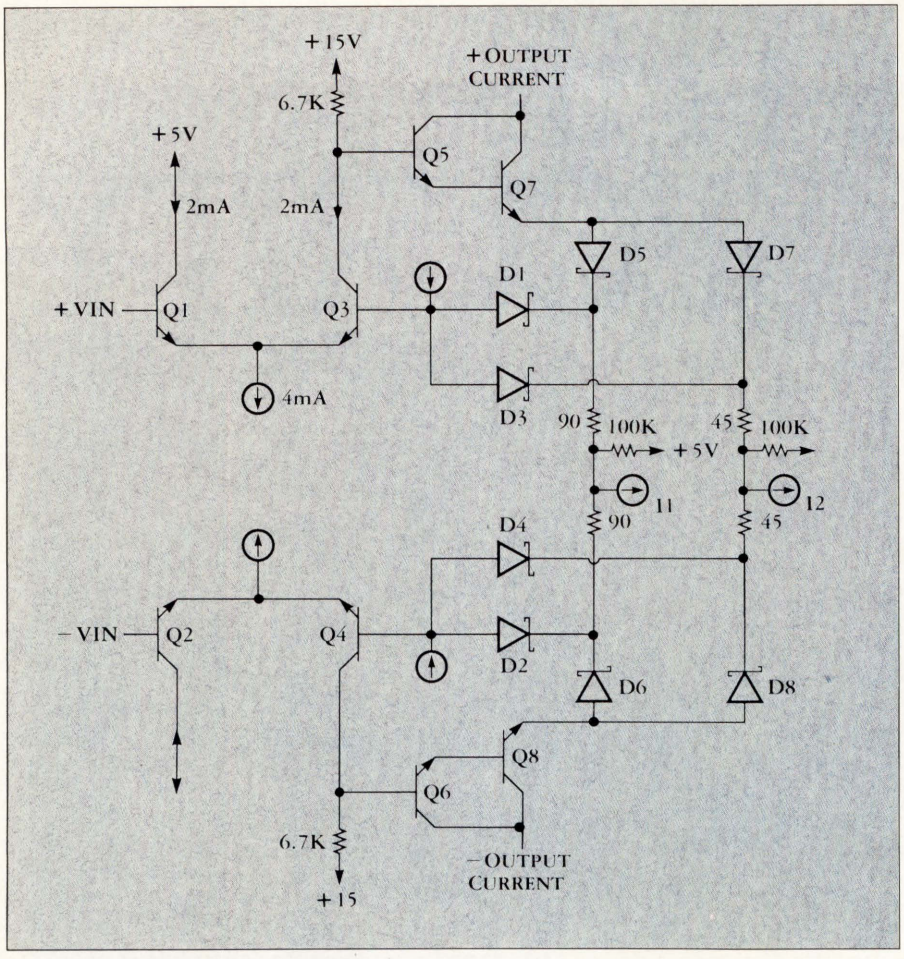

Figure 4. TTL logic inputs select one of six current sources such as I1 or I2 which force the first stage quiescent current through the appropriate gain setting resistors and associated diodes (only two gain settings are shown).

all amplifiers will eventually recover their original operating points. However, some amplifiers are better behaved in overdrive than others.

The overdrive specification of an amplifier depends upon its application. For an oscilloscope, the amount of recovery from overdrive is that percentage of the oscilloscope screen that is readily observable. For example, in observing a waveform on a screen, the recovery from overdrive should be between 0.2 and 2.0 percent of the display range. The advantage of fast overdrive recovery becomes apparent when comparing the equivalent number of bits a digitizer needs to observe the same detail. A 20 Megasample/second digitizer (sampling every 50 ns) would require 12 bits of resolution just to have its LSB (least significant bit) equal 0.02%. But installed in a 10 bit digitizer such as the 11401 oscilloscope, 0.02% of a 2 volt signal takes up 0.4 division at 1 mV/div. Each division is displayed with 100 codes of resolution, so each code then represents 10 μV. Ten microvolts is the equivalent of 17.6 bits of resolution of the 2 volt signal.

A 20 Megasample/second digitizer, unless it can be triggered and operated in an equivalent time mode, would be limited to 10 MHz bandwidth by the Nyquist criterion. However, in a realtime oscilloscope or in an equivalent time digitizer, quick overdrive recovery could display signal details in excess of 300 MHz. A real time oscilloscope could display a single event with a resolution limited only by noise to about 12 bits. An equivalent time digitizer, with averaging, could actually improve on the 10 μV (17.6 bits) resolution for repetitive signals. It is ironic that a digital scope is so effective for repetitive, analog-type signals and that analog scopes are so good a finding the single, isolated event frequently found in digital signals.

The M377 is capable of recovering from overdrive signals of up to 2 volts to within 0.04% in about 6 ns. Recovery to the 0.01% level takes 25 ns. This is more than three orders of magnitude faster recovery to the same level than the company's 7A13, a plug-in designed some years ago specifically for fast recovery. It is about 2 orders of magnitude faster than the settling of good modern operational amplifiers, and about five times faster than the best 12 bit DACs.

We became aware of just how good 0.02 percent recovery in 20 ns was when we tried to make the rest of the plug-in signal path support that kind of performance. For example, a 42 inch length of RG 58 coaxial cable, terminated in 50 ohms at each end still exhibits 0.02 percent skin effect loss 200 ns after a transient even though the total transit time of the cable is only 5 ns! The small diameter cable just 10 inches long used to connect the front panel BNC connector to the M377 inside the 11A52 plug-in exhibits 0.02 percent skin effect loss of its own about 20 ns after a step.

These losses for the first time became easy to measure and provided us with more than a few surprises. For example, the small diameter cable paradoxically settled to its final value much more rapidly than larger RG-58 of the same delay. As we thought about our measurements, the reason became apparent. The small cable has a copperweld (copper plated steel) center conductor. The dc loss of this cable is much greater per foot than the RG 58, and the time required to settle to its final (dc) loss is therefore less. An alternative explanation in the frequency domain: The dc loss equals the skin effect loss at a much higher crossover frequency.

Another discovery was a little more painful. The 11A52 plug-in has only a 50 ohm input impedance. This provides greater bandwidth and lower noise than a plug-in which requires a buffer amplifier with a 1 megohm input impedance. In fact, two M377s, one for each channel, are all the amplification the 11A52 has. The 11A52 then requires a 50 ohm attenuator which must be switched via relays. We chose a high reliability version of the popular subminiature style hermeticly sealed relay. These relays require glass to metal seals for hermeticity and use kovar for leads and header because kovar matches the thermal expansion coefficient of glass. But kovar is twenty to fifty times as resistive as copper and is ferromagnetic as well. Kovar is a great candidate for skin effect loss. Most of the kovar is gold plated, but the small segment inside the glass seal is not because the gold must be plated after the very high temperature glass to metal seal is made. It is in this very short section where serious skin effect losses occur. We found that, although the total path length through 5 relays in the attenuator was only 4 inches, the skin effect error did not die out to the 0.02 percent level for a full microsecond! To further complicate matters, when the relays are driven from a high impedance, such as a passive probe, the skin effect almost disappears. The reason for this is that the skin effect loss is an increase in the

series path resistance at high frequencies. When the attenuator is driven by a high impedance, the series loss is swamped out by the source's resistance.

Our solution was to introduce a network of resistors and capacitors in the attenuator which compensated for the skin effect loss. The whole attenuator has an 9 percent dc loss, but is flat from dc to 50 MHz within about 2 percent. The compensated attenuator is flat to within 0.05 percent at 40 ns, compared to about 2 percent at 40 ns before compensation. The attenuator risetime is just 130 ps.

■ Precision Overdrive Recovery in the IC

Precision overdrive recovery is a subject of some myth and misunderstanding. Most designers think about preventing transistors from saturating when they think of quick overdrive recovery. True, Schottky TTL switches are made faster than TTL by preventing saturation. But amplifiers are not the same as digital circuits, and transistor saturation is not what prevents a transistor amplifier from recovering quickly to within 0.1 percent.

It is thermal effects which dictate the amplifier recovery time at the 0.5 percent level and less. The faster the amplifier, the more power it dissipates, and the greater the potential for thermals. For example, amplifiers in the 7104 individually have a bandwidth of about 2.5 GHz per stage and thermals of about 6 percent per stage.

The M377 handles overdrive by using a Schottky diode to change the feedback loop in the input stage during overdrive. Figure 5 shows the input stage with the added components required for fast overdrive recovery. For simplicity, only one gain setting is shown.

In order to follow how the circuit works, you must realize that only the circuit half with the negative input becomes non-linear. This is true for the simple circuit of Figure 2b as well as the new circuit of Figure 5. For positive input signals, the amplifer takes all the current in the current source I7 and remains linear, but the mirror image circuit becomes non-linear when it gives up all its quiescent current to the positive input side. It is always the side with the more negative input which goes non-linear. Therefore it is only necessary to understand what happens for negative input signals at Q1's base.

The voltage at the cathodes of D3 and D9 will follow +Vin down until D9 and Q7 cut off because the voltage at the junction of D4 and D10 is held by -Vin and its associated amplifier. At that point, D1 conducts because Q5 and Q9 are still conducting due to current source I3. D1 steals the current flowing through D3 and cuts D3 off. D1 closes a new feedback loop consisting of Q3, Q5, Q9, and D1. Without any loading, this changed circuit will follow the input signal down until I5 saturates or Q5 breaks down.

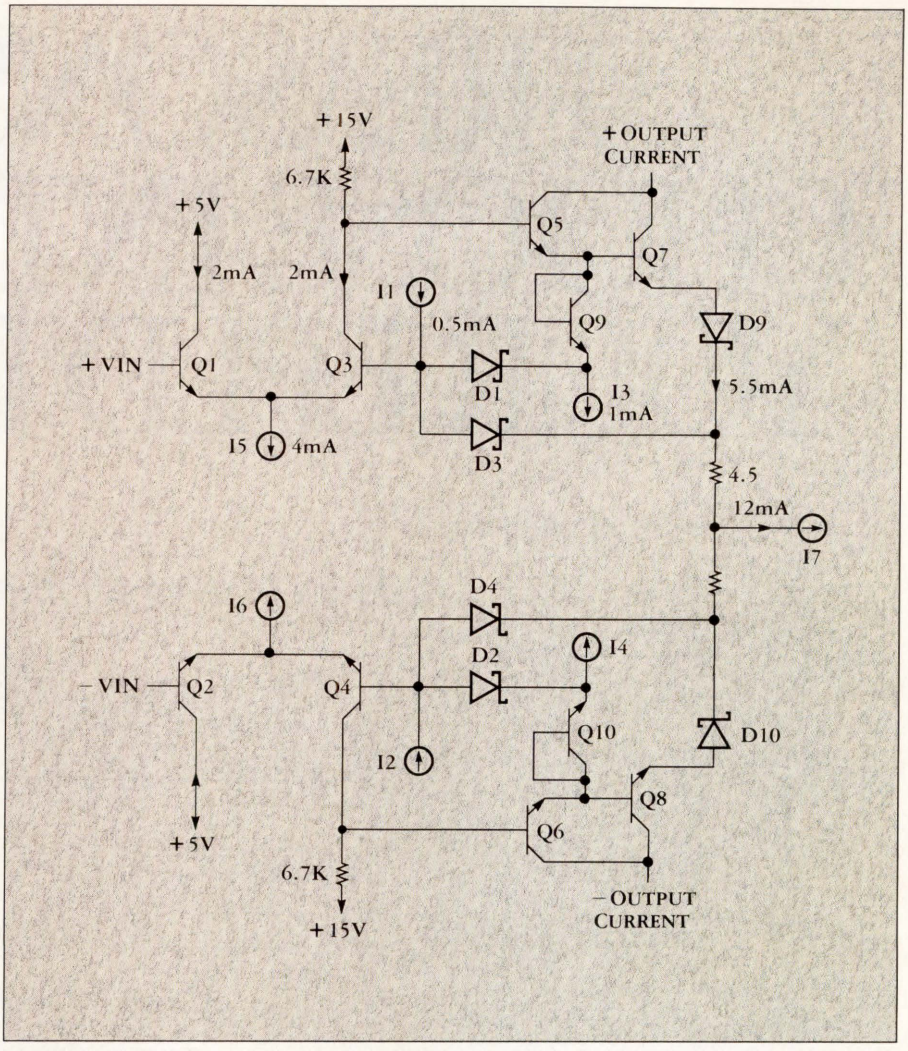

Figure 5. The M377 handles overdrive by using a Schottky diode to change the feedback loop in the input stage during overdrive. For simplicity, only one gain setting is shown.

The linear input range is about 50mV and the gain is very high with the 9 ohm total emitter load. Because of tight thermal coupling, the thermal voltages generated are only 80uV per milliwatt of power difference between Q1 and Q3.

■ Variable Gain Control

Almost every oscilloscope has a continuously variable gain control to allow the user to adjust the size of his trace to fit within the screen area exactly as he chooses. This control is not as frequently used as the coarse (step) control, yet it is a source of several design difficulties. Passive attenuation (use of a potentiometer as a variable voltage divider) works well up to about 50 MHz. The chief advantage of this scheme is low cost. The chief disadvantages are the limited bandwidth and the mechanical constraint tying the potentiometer to the front panel.

Above 50 MHz some form of electronic gain control based upon a Gilbert multiplier is usually used. While the Gilbert multiplier can work well, it is not without its problems. Chief among these are the low bandwidth when compared to other stages in the same IC, the addition of thermals and the generation of a fair amount of noise.

Automatic calibration placed requirements upon the M377 design—requirements which ultimately became useful features. One of these requirements is precision electronic gain control. In order for a microprocessor to calibrate system gain, the gain control element must be predictable and stable. At the same time, this precision allowed calibrated deflection factors in between the coarse step attenuation settings. Thus the deflection factor is always calibrated to better than one percent accuracy, even at 4.52 mV/div for example.

There are several Gilbert multiplier configurations to choose from. The usual

9 REASONS TO CHOOSE INTEL FOR YOUR EPLD DESIGN.

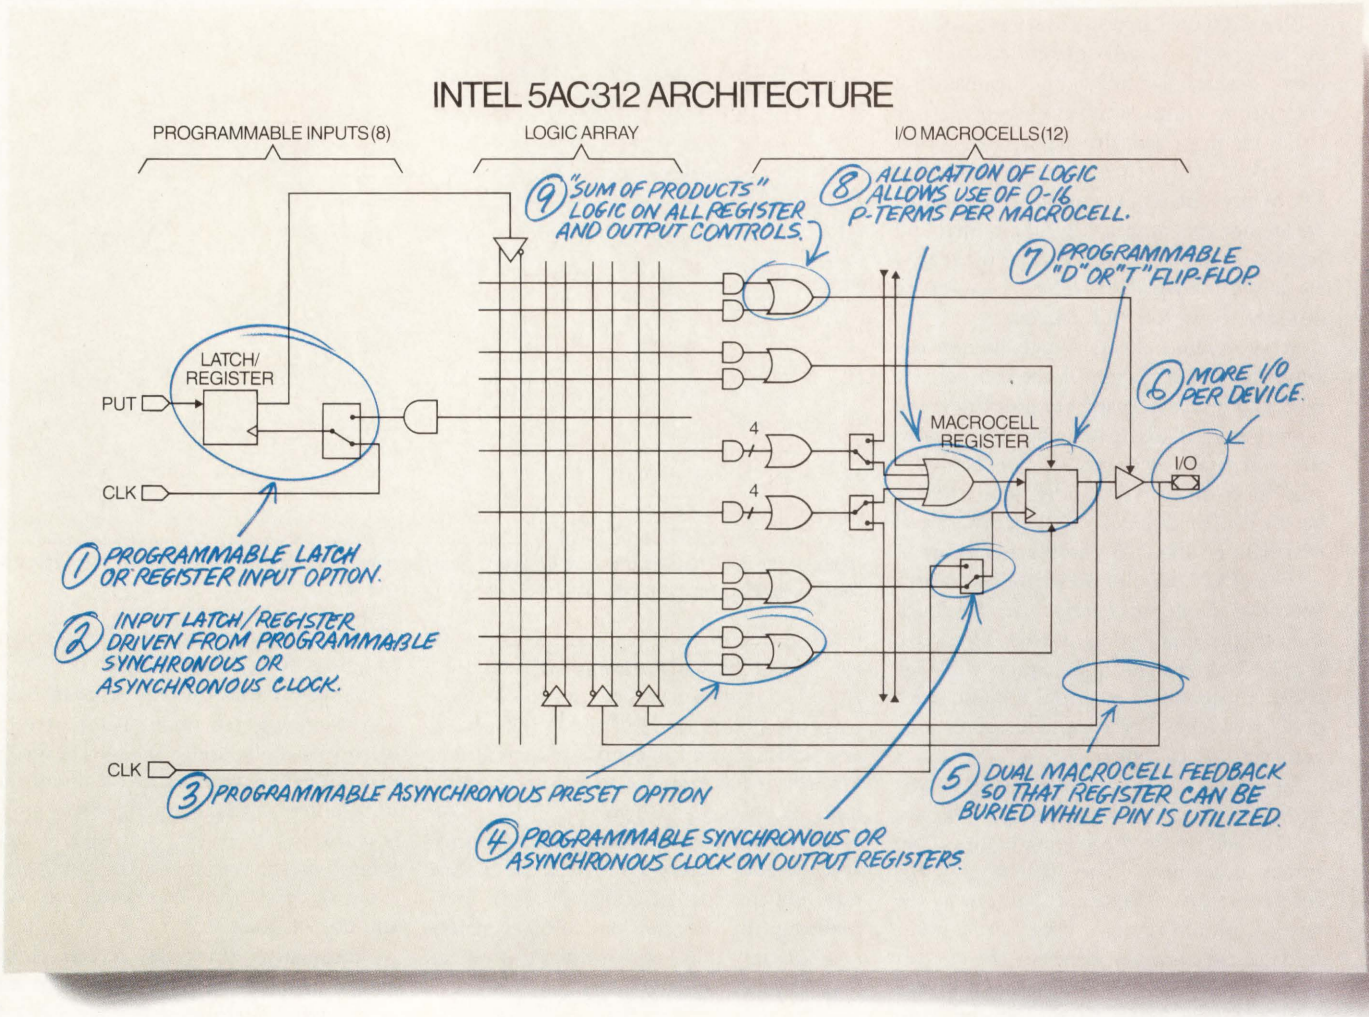

If you're looking for architecturally advanced EPLDs, look no further than Intel's innovative product family.

Start with our 5AC312. Based on Intel's 1.0-micron CHMOS* EPROM technology, this high-performance EPLD clips along at 50MHz with a very low power consumption of only 50mA, 150μA while idling in low power mode. And its enhanced programmable inputs can be configured as latches, allowing the 5AC312 to latch and hold data—a must for complex microprocessor-

AND 12 MORE.

and microcontroller-based systems. Plus you'll save valuable board real estate and improve overall system performance.

In addition, the 5AC312 has an exclusive programmable product term allocation scheme. It improves device utilization up to 33% and boosts performance by allowing implementation of applications requiring complex logic functions.

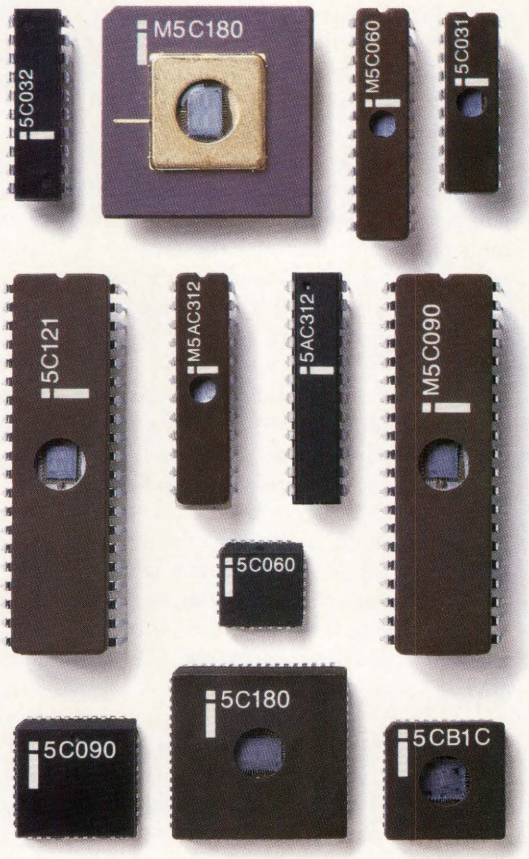

But we didn't stop there. These 5AC312 features are being incorporated into future EPLDs with even higher integration. So your upgrade path is assured.

But, architectural innovation is not the only reason to design with Intel's highly advanced EPLDs.

Even though these EPLDs offer more than six times the logic density per chip than other PLDs, we've dramatically simplified your design process. With our inexpensive and powerful PC-based development tool, iPLDS II, you can use your PC/XT,* PC/AT,* PS/2* or fully compatible system to turn EPLD design concepts into working silicon quickly, automatically and efficiently.

Of course, we provide extensive technical documentation to help you in system design. And, to make it even easier, we've taken service one step further. To ensure that the design process is as smooth as possible, users have access to Intel's toll-free EPLD "Logical Hotline" and a 24-hour electronic bulletin board for any technical assistance they might need. It saves hours of valuable time.

Best of all, the entire Intel EPLD family and all the development support you need are available now.

So call (800) 548-4725 and ask for Literature Dept. #W457. We'll tell you everything you need to know about the most technologically innovative EPLDs around.

Intel's 5AC312.
And family.

intel

*CHMOS is a patented process of Intel Corp. PC/XT, PC/AT, and PS/2 are trademarks of International Business Machines Corporation. ©1988 Intel Corporation.

CIRCLE NUMBER 10

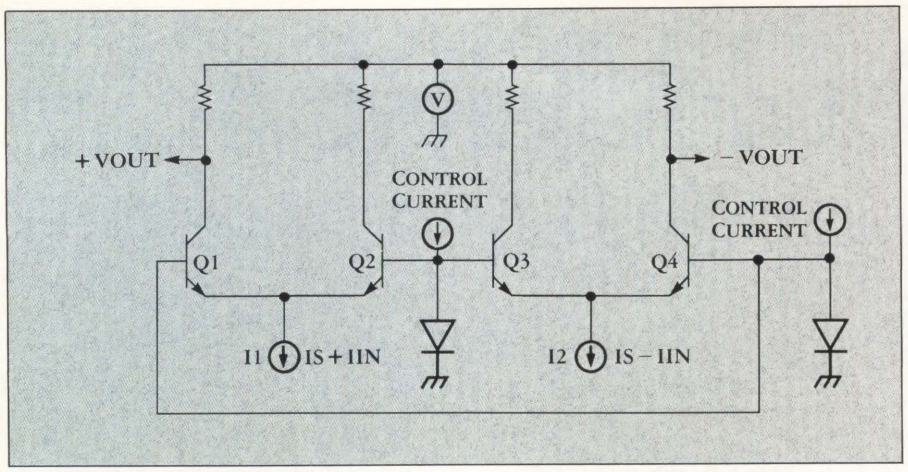

Figure 6. The two-quadrant Gilbert multiplier was chosen for the M377 because of its improved frequency response as a function of gain and its lower noise characterisitics.

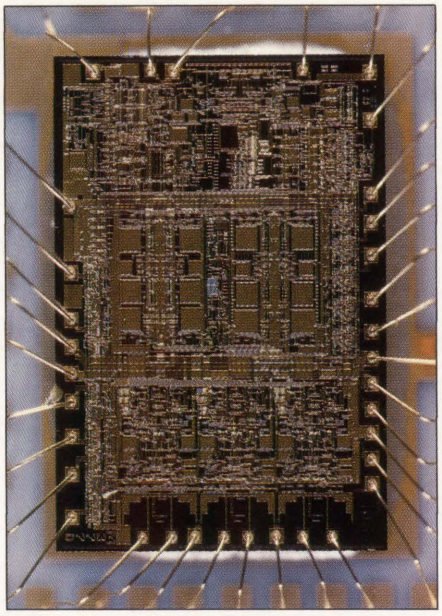

Figure 7. The M377 chip, which contains over 700 transistors, is almost a complete oscilloscope plug-in unit.

four-quadrant multiplier allows the input signal in the form of a current and its complement to be multiplied from +1 through zero to -1. But there's a problem with this configuration. The low frequency signal current splits between emitters in the same ratio as the dc quiescent current does, but at high frequencies the signal splits in accordance with the emitter impedances at that frequency. That impedance is determined primarily by r_b (at the highest frequencies where beta is low). Since the four devices are equal in size, the gain at the highest frequencies tends to be zero. In fact, there are only three gains for which the frequency response is theoretically flat, +1, 0, and -1. By symmetry, it is apparent that +1, 0, and -1 are also the only gains for which thermals are theoretically zero. The M377 multiplier is used at gains between +1 and +0.3, so flat frequency response at a gain of zero is of no benefit.

The four-quadrant multiplier's noise as a function of signal gain is lowest at -1 and +1, but maximum at zero gain. This is because any base resistance generates a thermal noise which is amplified by both emitter coupled pairs and added in the output. When the multiplier's signal gain is zero, both emitter coupled pairs have maximum gain, and therefore maximum noise. The S/N ratio is zero! At signal gains of +1 and -1, the voltage gain for noise generated in base resistors is zero.

Another Gilbert multiplier configuration is of greater use to the M377, and is shown in Figure 6. Here the current in the inner pair of transistors, Q2 and Q3, is wasted while the signal is taken strictly from the outer pair, Q1 and Q4 (in the four-quadrant multiplier, the collectors of Q1 and Q3 are connected together, and similarly for Q2 and Q4). Since currents in the inner pair match the currents in the outer pair when the signal gain is 0.5, the thermals are zero and high frequency current split is perfect by symmetry. Furthermore, half the noise current is thrown away with the unwanted signal current, so the noise is less too! Although the frequency response and thermals are not perfect at gains other than 0, 0.5, and 1, the gain never strays as far from these ideal points as it does in the four-quandrant multiplier. Furthermore, the frequency response is improved because the two-quadrant multiplier (Figure 6) has only the capacitance of one collector at each output instead of two.

A possible disadvantage is that the two-quadrant multiplier will not allow inversion, while the four-quadrant multiplier will. In practice though, there are ways of inverting a signal much more accurately than the Gilbert multiplier. For greater versatility, the M377 employs a separate inverter for each of its three outputs. The signal inversion is accomplished with less than 0.02 percent gain change.

■ Band Gap Reference

The variable gain is controlled by an analog input voltage between -1.0V (zero gain) and +1.0V (full gain). Except at 0V input where the gain is one half that set by the coarse gain control, the gain must be referenced to an absolute dc voltage. This puts the gain at the mercy of the voltage which sets the +1.0V and -1.0V references. For this reason, an internal band gap reference was added to provide independence from power supplied to the chip. With the built-in band gap, change in power supply due to load changes in the instrument does not affect the gain accuracy.

There's another advantage to the use of an internal reference. The internal band gap can have zero output until the +5V supply is at least +3V. All the M377 current sources from the -5V supply are referenced to the band gap. Those current sources cannot turn on in the absence of a +5V supply, and therefore no transistor can go into saturation or cause latch-up if the +5V supply is lost. Of course, the current sources from the -5V supply cannot be on without the -5V supply either, hence there is no power supply sequencing necessary and loss of any supply safely shuts down the chip.

■ Multi-Channel Operation

The M377 took an 80 person-month custom IC design effort, which could only be justified by substantial volume. At Tektronix, that volume required the M377's use in more than one product. The M377 was originally designed for the 11A32, 11A33, 11A34, and 11A52 plug-in amplifiers.

The 11A32 and 11A52 are two channel plug-ins, the 11A33 is a single channel plug-in, and the 11A34 is a 4 channel plug-in. The M377 is a single channel amplifier with three separate outputs, one for display, one for trigger, and an additional channel for auxiliary purposes. Each channel 's output impedance and common mode voltage level remain the same whether an output is enabled or not, so outputs from two or more different M377s can simply be connected in parallel to form a channel switch. Any output may be selected, the outputs can be alternated or chopped. The add mode is accomplished by turning on two channels. Because each output can be inverted in only 200ns, it is even possible to display the sum and difference of two channels with chopped or alternate sweep mode. Since each M377

output is separately controlled, the trigger mode can be different than the display mode.

To accomplish the same gain at the output of each plug-in independent of the number of channels, the M377 is laser trimmed while still in the wafer stage to have an output impedance of 50 ohms, 100 ohms, or 200 ohms per side. The two channel plug-ins use the 100 ohm version and the four channel plug-ins use the 200 ohm part. In this way, the overall plug-in output impedance is 50 ohms per side and the gain is the same independent of the number of channels.

In the case of two channel plug-ins, 100 ohm transmission lines are used to connect together the outputs from two M377s. The combined output is taken midway between the two M377s. This results in a nominally perfect 50 ohm output impedance at all frequencies, a perfect reverse termination.

The four channel plug-in cannot be similarly reverse terminated because it is not possible to construct 200 ohm transmission lines on etched circuit boards. Surprisingly, reflections among the four chips almost cancel, so the small loss of bandwidth which occurs is due primarily to the increased capacitive loading on the output—from the chips themselves. Reverse termination is somewhat compromised by the 100 ohm, 125ps, transmission lines used to connect the 200 ohm outputs together. Fortunately, since the oscilloscope mainframe is terminated in 50 ohms, there is little signal reflected back to the plug-ins.

The ability to laser trim nichrome resistors while the M377 is still in wafer form affords the opportunity to trim more than just the output impedance. Common mode output level is trimmed to zero. Gain at the six discrete steps is trimmed to 1 percent tolerance and two gain settings are trimmed for dc balance as well. The variable gain control (Gilbert multiplier) is trimmed so that full gain and zero gain occur with + 1.0 volts and -1.0 vols at the analog gain control input respectively. In all, eighteen resistors are trimmed and several qualifying tests are performed in about 60 seconds.

Finally, the trimmed wafer is moved to a high speed tester where 126 dc tests are performed, some of them to a .02 percent test limit, in 11 seconds per chip.

The fact that it is very versatile, that its digital inputs are TTL compatible, and that its analog inputs and outputs are differential and centered on ground have made the M377 popular with other designers. The precision analog chip, which contains over 700 transistors (Figure 7), is currently being designed into a number of other Tektronix products. ■

Acknowledgments

The original architecture and much of the circuit design were the author's work, but major contributions were made by Pat Quinn, Gary Polhemus (now with National Semiconductor), and Art Metz. Eight new US patents were issued on circuits in the M377, including the gain switching and overdrive recovery circuits discussed in this article.

About the Author

JOHN ADDIS, *a Principal Engineer in the Portable Instruments Division of Tektronix Inc., joined the company in 1963 after receiving a BSEE from MIT. Prior to this, he worked in Tektronix's Laboratory Instruments Division, where he developed the M377 analog IC. Addis, who holds 12 patents, was also responsible for other analog ICs in the 11000 series oscilloscope plug-ins and the vertical system in Tektronix's 7104/7A29 1-GHz oscilloscope.*

HARDWARE *Modelers*

VLSI Systems Design Staff

As simulation gained acceptance as a necessary part of a rational design strategy, it became quickly apparent that the simulation was only as good as the models that were used in the simulation. Although sparce at first, the libraries offered by the vendors of simulator systems quickly grew to hundreds of standard MSI, LSI, and VLSI chip models. But to the leading edge systems engineer—designing a new high-performance product—the libraries had some serious shortcomings. The models for the hottest new VLSI chips lagged the product introduction by up to two years. There was also no quick and inexpensive way to get a model for a new ASIC.

The main reason for this dearth of models was the large investment in time and money required to develop models of these new complex devices. Semiconductor houses that were developing standard products such as microprocessors or DSP chip sets, were able to justify the investment in these models, but the systems houses that were developing limited run ASICs for military or aerospace applications couldn't afford the cost or delay.

But all was not lost, Dr. L. Curtis Widdoes Jr. of Valid Logic Systems Inc. developed "hardware modeling." This technique allowed the actual device to be tied into the software simulation—performing the same functions as the software model.

There are a few drawbacks to using a hardware model instead of a software model, but there is one big advantage. It is, of course, that the hardware model is often the "only game in town." Without the modeler, there would be no simulation—or require an unacceptable delay in getting the product to market.

In operation, the software simulation sends stimuli to the modeler, which in turn stimulates the proper inputs of the actual device, waits for and captures the appropriate output responses of the device, and then sends the results back to the software simulation.

Some of the limitations of a modeler are: the limited amount of simulation cycles that can be run for devices with dynamic memory components; the maximum speed that stimulus patterns can be delivered to the actual device; and the complexity of the software shells that allow the timing characteristics of the device to be added into the simulation.

Hardware modelers can also be used to verify prototype ICs by plugging the prototype ASIC into the modeler and connecting it into the original software simulation of the ASIC and its accompanying system. A comparison of the simulation with the software model and the hardware model can readily verify the functionality of the "first silicon."

Our survey turned up only six vendors offering hardware modelers, but a new company that is a spinoff of Valid Logic Systems, called Logic Modeling Systems Inc. (San Jose, Calif.), will soon roll out with its first product. It's expected to represent the next generation in performance—particularly since the founder of the company is Curtis Widdoes, the inventor of hardware modeling.

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

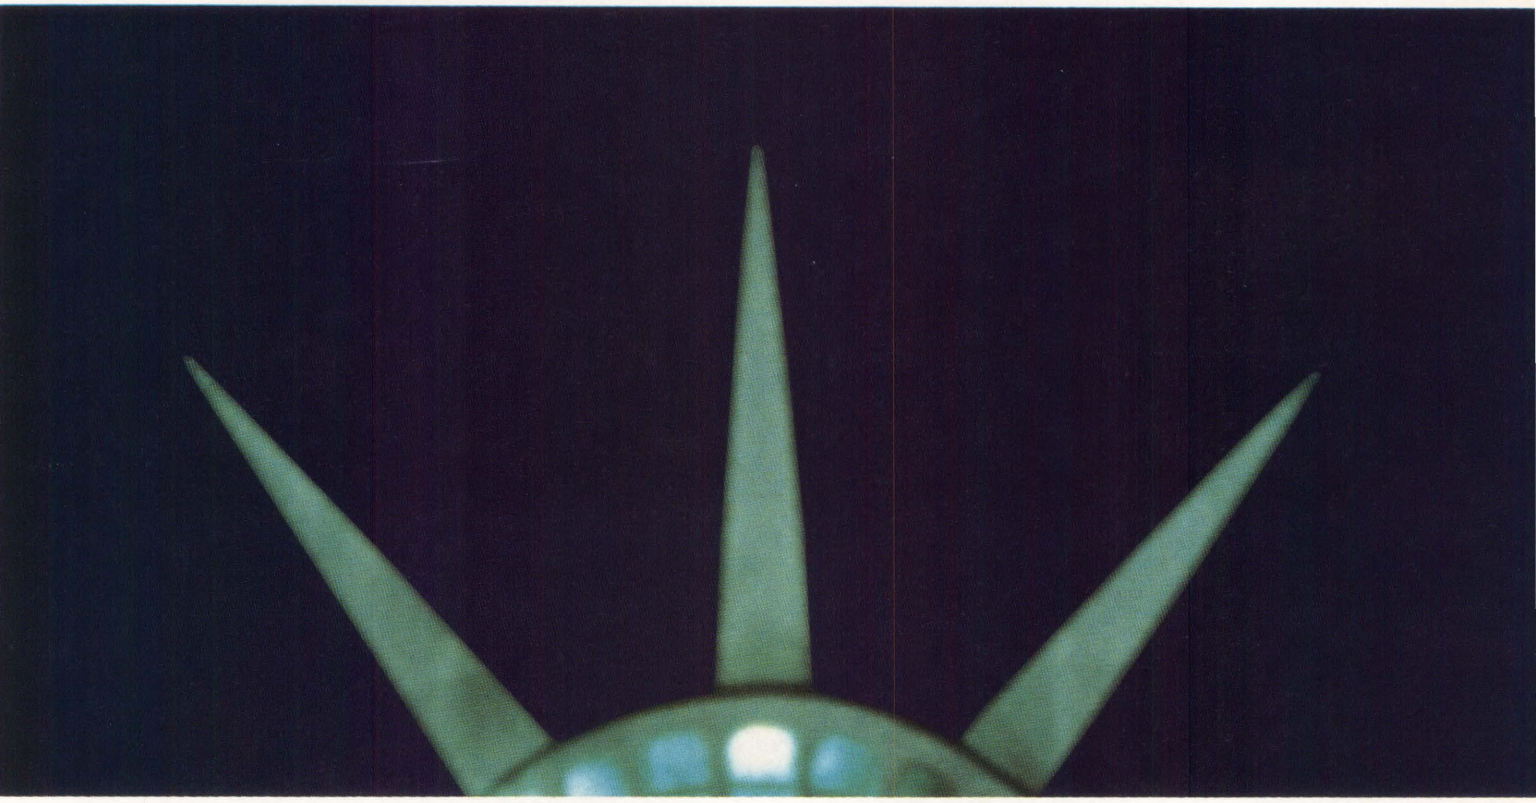

Get the full picture of your ASIC design's performance. Get the IKOS Simulation System. IKOS gives you more simulation information in less time than any other system. Simulate 15,000 gates with 34,000 vectors in 34 seconds. Or 8,500 gates with 5,100 vectors in 4 seconds. Complete with detailed timing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

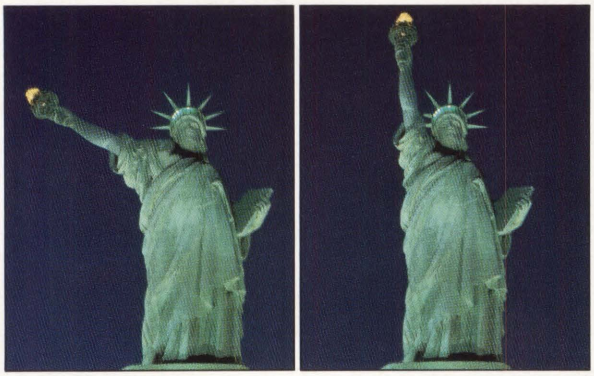

Without the IKOS Simulation System, most ASIC design mistakes show up after production.

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 11

DIRECTORY OF HARDWARE MODELING SYSTEMS

VENDOR	**DAISY SYSTEMS CORP.** *700 E. Middlefield Rd.* *P.O. Box 7006* *Mountain View, Calif.* *94039*	**GENRAD INC.** *Design Automation Products* *300 Baker Ave.* *Concord, Mass. 01742*	**HHB SYSTEMS/CADNETIX** *1000 Wyckoff Ave.* *Mahwah, N.J. 07430*	**MENTOR GRAPHICS CORP.** *8500 S.W.Creekside Pl.* *Beaverton, Ore. 97005*	**TERADYNE INC./EDA GROUP** *321 Harrison Ave.* *Boston, Mass. 02118*	**VALID LOGIC SYSTEMS INC.** *2820 Orchard Pkway.* *San Jose, Calif. 95134*
CONTACT	Manfred Muehter Product Marketing Manager (415) 960-7036	Peter Denyer Marketing Manager (408) 432-1000	Larry Blessman Product Marketing Manager (201) 848-8000	Pat Wolfram Product Marketing Manager (503) 626-7000	Linda Lowe Technical Publicity (617) 482-2700 x3567	Don Mazur Product Marketing Manager (408) 432-9400 x2278
PRODUCT AND PRICE	**PMX SYSTEM** $16.5k for base system	**HICHIP PHYSICAL MODELING SYSTEM** Contact Vendor	**CATS 8000, 10000, 11000, 12000** $80k–$150k	**HARDWARE MODELING LIBRARY (HML)** Starting at $44.8k	**DATASOURCE VLSI MODELING SYSTEM** $48k–$140k	**REALCHIP (RC)** $20k **NETWORKED REALCHIP (NRC)** $40k **REALCHIP II (RCII)** $65k **REALMODEL (RM)** $80k (includes acceleration)
PATTERN DEPTH PER I/O PIN AND NUMBER OF PINS SUPPORTED	(4 different board options) Basic dynamic: 128k, FastBoard: 64k, DeepBoard: 1M; Maximum pins per PMX unit: 780 pins, Maximum PMX units for: Logician, Personal Logician, and Personal MegaLogician = 1, MegaLogician = 5, GigaLogician = 6.	Standard: 64k bits, optional: 256k bits; each HICHIP modeler supports either 8 devices of up to 128 I/O pins, 16 devices of up to 64 I/O pins, or a combination thereof. Each HICHIP modeling system supports a total of 1,024 I/O pins. Either one 128 pin device or two 64 pin devices can access the full complement of memory.	512k; 192 bidirectional inputs, 1,152 outputs.	256k maximum; 2,048 I/O pins per system. With 64 pins/device: depth/pin = 256k, and number of devices/system = 32; with 128 pins/device: depth/pin = 128k, devices/system = 16; and with 256 pins/device: depth/pin = 64k, devices/system = 8.	256k vectors; DATASource provides up to 2,400 pins, and uses dynamic memory allocation to assign its 80 megabit system memory on an as needed basis to hardware models (different models can be allocated different memory depths on the fly). The system can apply full 256k vector depth to a maximum of 160 pins at any one time.	RC, NRC = 4–16k; RCII, RM = 256k–2M. (All devices get maximum pattern memory). 64 pins per device; total pins for RC, NRC = 2,048 total pins (4,096 with expansion); RCII, RM = 1,536pins (3,584 with expansion).
MAXIMUM CLOCK FREQUENCY AND NUMBER OF I/O PINS SUPPORTED	25 MHz (vectors); 376 per unit, total of 1,880 pins with the MegaLogician	64 MHz for on board clock and up to 80 MHz with external clock; 128 I/O pins can be driven with vector pattern rates up to 16 MHz	10 MHz; up to 192 pins per device	16 MHz; 128 pins	16.7 MHz - two formatted clocks can drive a total of 66 million clock edges per second; 160 pins	RC, NRC = 4 million clock edges/second; RCII, RM = 16 million clock edges per second; 64 pins
SIMULATION CYCLES FOR DYNAMIC AND STATIC DEVICES	Dynamic: depends on application, up to 1 million patterns can be presented to modeled devices. Static: unlimited simulation cycles.	The total number of simulation cycles is limited only by the total number of vectors actually passed to the modeler. This is circuit dependent. The modeler does not differentiate static and dynamic chips.	512k clock cycles. Since modeler allows multiple users to interactively access the same model, static models become dynamic.	Dynamic: 256k cycles of the modeled device. Static: unlimited simulation cycles.	256k simulation cycles. Since DATASource is designed to support multiple users and multiple device instances (for concurrent fault simulation), this precludes unlimited simulation cycles.	Dynamic: cycles for modeled device: RC, NRC = 64k RCII, RM = 4M. Static: unlimited simulation cycles.

PIN PATTERN WIDTH AT MAXIMUM I/O RATE	376 pins per PMX unit.	128 pins.	196 pins.	128 pins at 16 MHz, and up to 256 pins at 8 MHz.	160 pins.	64 pins.
DEVICES SUPPORTED SIMULTANEOUSLY AND PIN COUNT LIMIT	Maximum of 3,900 pins can be allocated to any number of devices at the user's discretion. Maximum of 780 static pins or 376 dynamic pins per device.	For up to 64 signal pins, 16 devices per HICHIP box. No practical limit on the number of HICHIP modelers in a given network. For up to 128 signal pins, 8 devices per box. 64 pin devices and 128 pin devices can be intermixed to the limit of the HICHIP configuration. 128 I/O pins per device.	30 devices. 192 input and 1152 output (bidirectional is one input and one output).	32 64-pin devices, 16 128-pin devices, 8 256-pin devices or mixed combinations of the above. 256 I/O pins per device.	60 devices (DATASource supports virtually unlimited device instances for concurrent fault simulation and multiple users). 320 I/O pins per device.	Not specified. RC, NRC = 128 pins per device; and RCII, RM = 256 pins per device.
MAXIMUM PIN COUNT DEVICE SUPPORTED	780 static pins or 376 dynamic pins.	Up to 128 I/O pins.	1,344 pins.	256 I/O pins.	320 I/O pins.	RC, NRC = 2048 pins. RCII, RM = 1536 pins.
DEVICE WIRING SOFTWARE CONFIGURABLE?	Yes.	Except for power, ground, and clock, no custom wiring is generally necessary. Pin-outs of device in modeler are hardwired to edge connector of the daughter card. Software "shell" maps pin list from simulator to actual pins of device in modeler.	No.	Yes.	Yes, DATASource uses a pin map.	Not normally.
ROUNDTRIP EVALUATION TIME	Varies by application.	RS-232 = 120ms (Sun 3-50); TCP/IP Ethernet = 15ms (Sun 3-50). (the time is both host and communications protocol dependent.	Not specified.	20ms including simulation evaluation time, server/driver/network communications, tri-state sensing, vector streaming, component sensing, and simulation update. 12ms for simpler devices.	10–100 milliseconds per MIPS of computer power; time varies according to the complexity of the behavioral shell for the hardware model.	RC = 80ms; NRC and RCII = 120ms (network dependent); RM = 60ms.
TIMING MEASUREMENT CAPABILITY?	Interface signals for synchronizing oscilloscope or logic analyzer are provided.	No.	No.	Yes, 2 ns sample resolution. Typical accuracy = ±4 ns, with absolute worse case accuracy = ±10 ns.	No.	No.
CONCURRENT MULTIPLE USERS? **DEVICES CHANGED WHILE MODELER IS OPERATING?**	No. No.	Yes. No, the HICHIP configuration and model list is read at power-on, so any device introduced later will not be seen.	Yes. Yes.	Yes. Yes, HML notifies and places all simulations on hold when powered down for device insertion/removal. Simulations resume when power is reapplied.	Yes, unlimited. Yes.	Yes for RCII and NRC. No.

DIRECTORY OF HARDWARE MODELING SYSTEMS (continued)

VENDOR	**DAISY SYSTEMS CORP.** *700 E. Middlefield Rd.* *P.O. Box 7006* *Mountain View, Calif.* *94039*	**GENRAD INC.** *Design Automation Products* *300 Baker Ave.* *Concord, Mass. 01742*	**HHB SYSTEMS/CADNETIX** *1000 Wyckoff Ave.* *Mahwah, N.J. 07430*	**MENTOR GRAPHICS CORP.** *8500 S.W.Creekside Pl.* *Beaverton, Ore. 97005*	**TERADYNE INC./EDA GROUP** *321 Harrison Ave.* *Boston, Mass. 02118*	**VALID LOGIC SYSTEMS INC.** *2820 Orchard Pkway.* *San Jose, Calif. 95134*
SPECIAL FEATURES	Software initialization routines can be written for devices that contain internal clock dividers and other devices that present unusual initialization conditions. Also synchronization routines can be written for dynamic devices requiring them. Tri-state sampling of device outputs.	HICHIP modeling language, supports "State Sensitive Modeling" with all features of HILO HDL; Query mode indicates modelers on network and location of modeled devices; also checkpoint/restart; total memory in use; memory pre-allocation; single-user lock; design validation, dynamic timing analysis, and fault simulation modes; and memory dynamically allocated in multi-user and/or multi-chip environments. No tri-state sensing.	Device initialization is handled by software shell. "Hardware Modeling Toolkit" available to aid in generation of software shell. Tri-state sampling of device outputs.	The "config" tool provides detailed graphical and textual information for the automatic (soft) wiring of components. Tri-state sampling of device outputs.	Teradyne's LASAR Version 6 behavioral modeling language is available for writing behavioral shells, which can incorporate: initialization sequences that (a) check that model is initialized, or (b) initialize the model; recognition of up to 15 levels of synchronization matching coded anywhere in the vector stream; information about critical device timing parameters; and rules for propagation of unknown (X) states through the hardware model. Tri-state sampling of device outputs.	Device initialization sequence in software shell. Tri-state sampling of device outputs.
DEVICE SOFTWARE SHELLS AVAILABLE IN MODELER LIBRARY	234 devices.	22 devices.	150 device types, which can repesent multiple speed devices (i.e. 12 MHz, 16 MHz 68020).	80 devices.	60 devices.	More than 140 devices.
SIMULATORS SUPPORTED	Daisy Logic Simulators (DLS, MDLS, GDLS), Daisy Timing Verifier (DTV), and Daisy Fault Simulator (Mega Fault).	HILO-3 and System HILO.	CADAT.	Mentor Graphic's QuickSim (logic), QuickFault (fault), QuickPath (timing).	Teradyne's LASAR Version 6 and AIDA simulators.	ValidSim, Valid FaultSim, Gateway's Verilog and Verilog XL, and Teradyne's LASAR.
HOSTS SUPPORTED	Personal Logician, Logician 386, IBM PC/AT, Accelerators: Personal Mega Logician, Mega Logician, and GigaLogician.	Apollo; VAX (VMS, ULTRIX, BSD 4.2 UNIX); HP9000 Series 300, 500, 800; Ridge; IBM RT-PC; SUN-2; SUN-3; Intergraph Interpro 32, 220; and NEC EWS4800.	Sun, VAX, Apollo, HP9000.	Apollo DN3000, DN4000, DN3500, DN4500, DN10000.	Any VAX/VMS system and Sun workstations.	Valid, Sun, DEC, and IBM PCs.
NETWORKING SUPPORTED	Ethernet: both XNS and TCP/IP pototcols.	RS232, TCP/IP Version 1 or 2, (single gateway).	Ethernet TCP/IP.	Apollo Domain Network.	Any networks in which VAX or Sun systems run.	TCP/IP.

Cypress Semiconductor introduces the RISC Motorola didn't consider.

The RISC that lets you build faster computers faster.

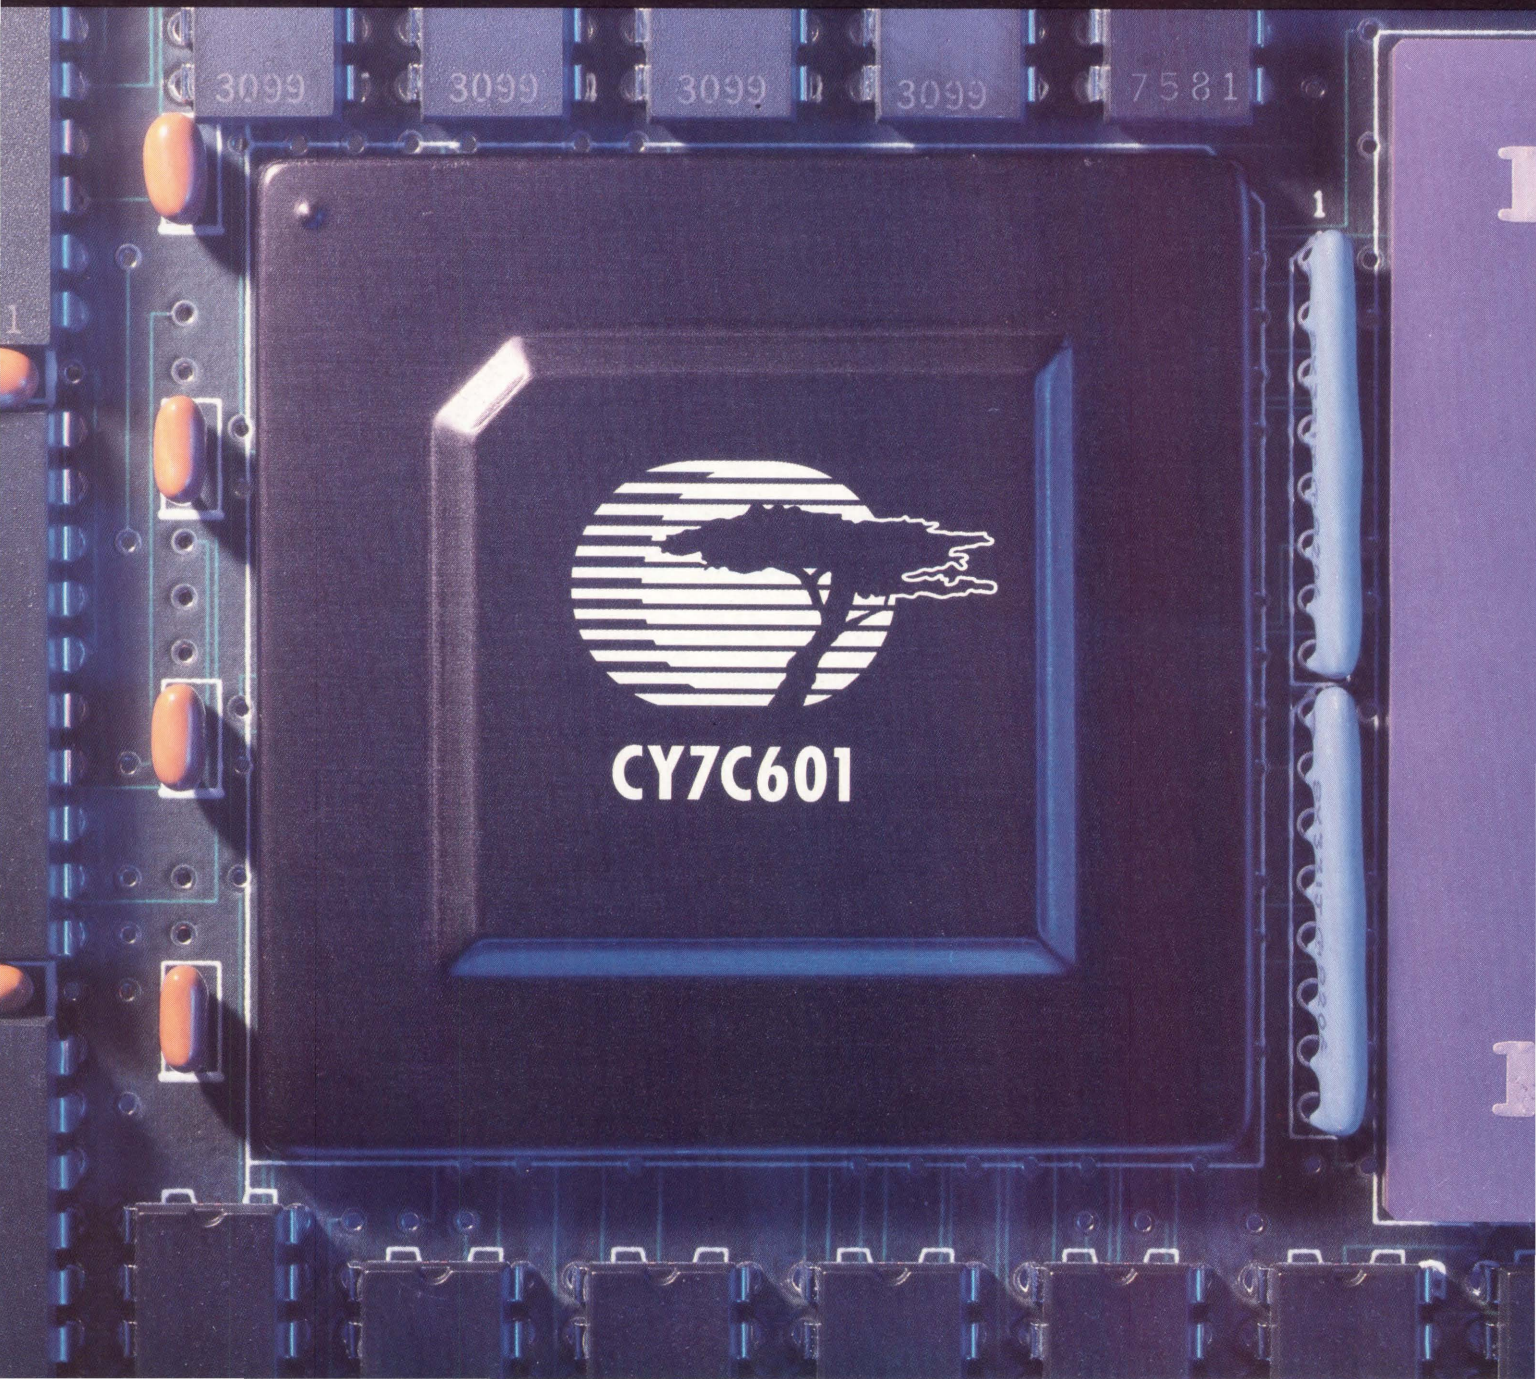

Introducing the RISC 7C600 Family: 20 VAX MIPS, SPARC architecture, and development systems today.

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

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

A complete chip set.

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

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

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

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

Develop your systems quickly.

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

Design around a proven architecture.

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

Make the fastest decision.

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

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

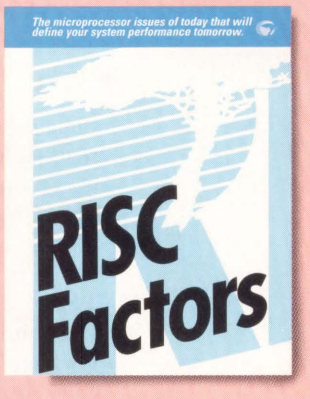

CYPRESS SEMICONDUCTOR

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

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

-A-
VHDL DESIGN
Environment

ERICH MARSCHNER,
CAD LANGUAGE SYSTEMS INC., ROCKVILLE, MD.

VHDL is a newly-adopted IEEE Standard hardware description language that supports structural, behavioral, and dataflow styles of design and documentation for digital systems (IEEE, 1987). Now that the language has become an IEEE Standard, many in the industry are trying to understand exactly how the language will affect their existing design practices. What benefits will VHDL bring to those who use it? How can VHDL be incorporated into existing methodologies? Who will be the suppliers of VHDL-based design tools. How will they work? What can be expected of such tools?

A HIGHER LEVEL OF DESIGN ABSTRACTION IS PROVIDED

■ WHY VHDL?

The need for a language such as VHDL has been felt for many years, in both the government and commercial sectors. Current design practices are tied to proprietary tools with proprietary notations.

For example, one company may use Mentor Graphics' simulator, in which simulation models are defined using BLMs. Another company may be committed to Daisy Systems and DABL models. Others may use Endot (ISP'), Zycad (ZILOS), or Silvar-Lisco (HHDL) simulators and model languages. Schematic and layout data is typically stored using a proprietary database format. The Electronic Design Interchange Format (EDIF) has started to solve the problem of moving schematic and layout data between proprietary databases, but there is still no general solution for simulation models.

The plethora of proprietary design languages used by various DoD contractors has made it difficult for the Department of Defense to reuse parts designed by one contractor in systems developed by another. This situation led to the development of VHDL as part of the VHSIC (Very High Speed Integrated Circuits) program. The DoD has required VHDL in a number of programs. In addition, a recent draft of MIL-STD 454, regulation 64, calls for VHDL documentation of ASIC parts.

A similar problem has plagued many large aerospace, defense, and computer companies. Many of these companies use different design systems in different divisions of the same company, or even in the same division. It is not unusual to find a single design team working with tools provided by different CAE vendors. This diversity is necessary because no one vendor provides the best tool for each application. However, the use of disparate tools requires that design data be translated back and forth between different formats, and such translation often results in data loss or errors.

Here again, EDIF provides a partial solution, but abstract design data simply cannot be handled with the current version of EDIF. VHDL provides an alternative standard form for the communication of high-level designs, particularly simulation models.

Interest in VHDL has been growing steadily. IBM was one of the first companies to implement VHDL for internal use, even before the new IEEE Standard was adopted; the capabilities of this internal implementation were described at the 1987 Design Automation Conference (Saunders, 1987). A number of CAE vendors have announced that they will be support-

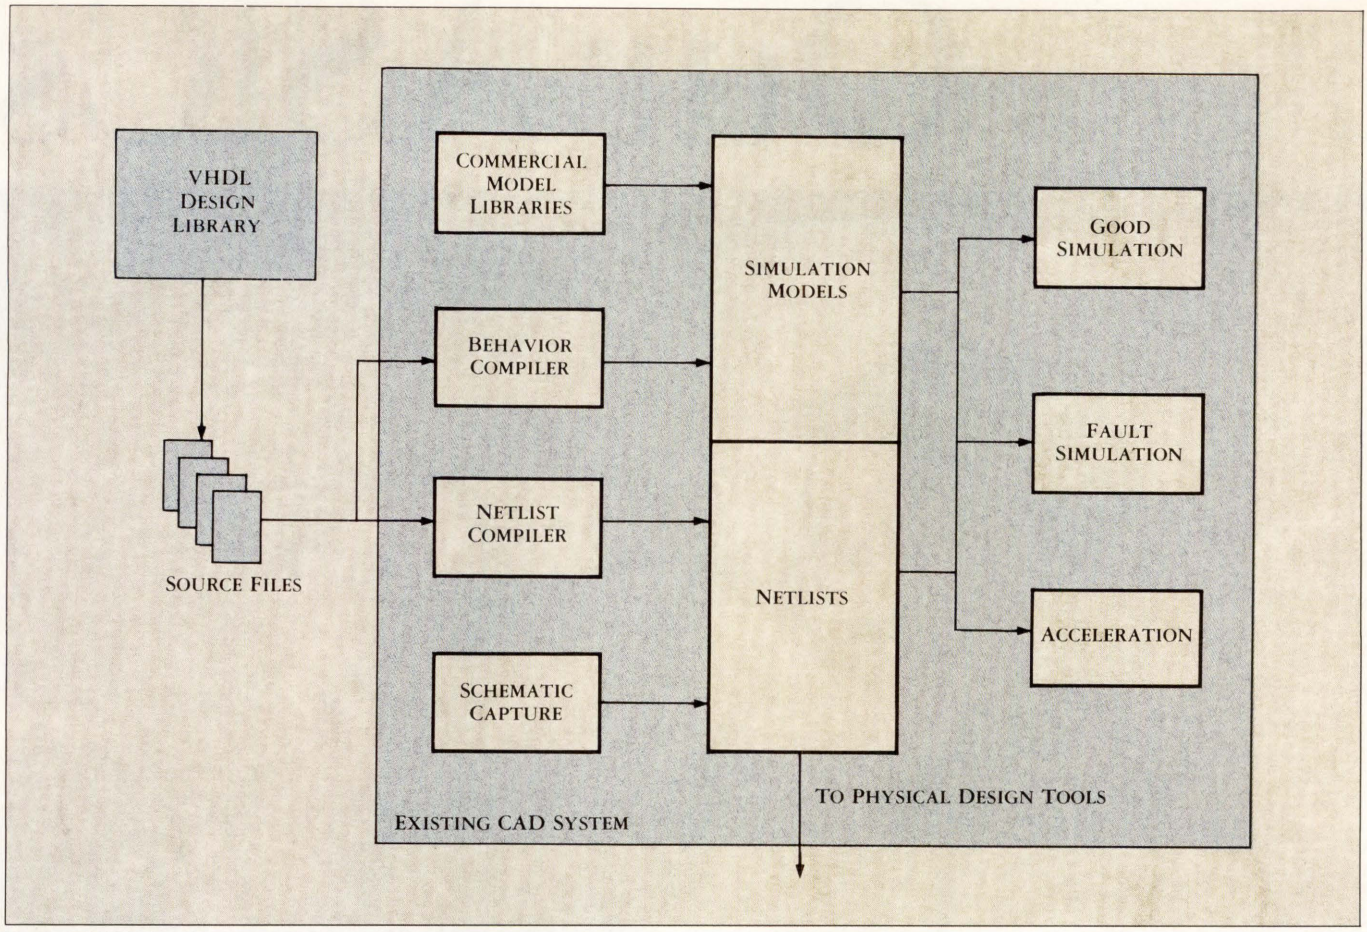

Figure 1. A VHDL interface to an existing system.

ing VHDL. Rather than develop a new, full-scale VHDL simulator, many will elect to provide VHDL interfaces to their existing products, as illustrated in Figure 1.

■ USING VHDL IN HARDWARE DESIGN

What will some of these tools look like? What can VHDL do for you, and how will it affect your existing design methodology? The answer is that VHDL is not a panacea and, in particular, will not replace the low-level design tools in use today.

What it will do, however, is provide a higher level of abstraction for the design of systems, whether such systems reside on a single chip or consist of multiple chips, boards, or complex subsystems. This capability will allow designers to make trade-offs much earlier in the design process, where the cost of making a change is fairly small.

The basic tools required to make use of VHDL are the language processor, typically called a VHDL "analyzer," and a VHDL simulator. A VHDL analyzer functions like a programming language compiler: it reads in a VHDL design unit, verifies that the syntax and static semantics of the design unit are correct and, if so, creates an intermediate-form representation of the design in a design library. This intermediate-form representation, also called a library unit, is then available for use by other tools. In particular, the library unit representing one VHDL design unit can be examined during the subsequent analysis of another design unit that refers to items declared in the first unit.

This allows a large design to be broken up into smaller, more easily managed pieces, without compromising the ability of the VHDL analyzer to verify the interfaces among those pieces. This process is illustrated in Figure 2.

As mentioned, a number of vendors have elected to provide VHDL interfaces to their existing simulators. This approach has several advantages, in particular the availability of simulation model libraries, additional capabilities such as schematic capture and fault simulation, and a large existing user base. Furthermore, until the industry has had a few years to work with VHDL and realize its full potential, the VHDL subset that will most likely be used is the one that corresponds to the features provided by existing simulators. These simulators, with built-in multi-valued logic types and highly optimized algorithms, will run circles around a more general, full-VHDL simulator, in terms of simulation throughput.

VHDL MUST INTEGRATE WITH EXISTING HARDWARE DESCRIPTION LANGUAGES

Such an interface will typically involve one or more VHDL design units that define the data types, functions, and other resources that come pre-defined in the environment of the existing simulator. For example, most simulators are based on a specific multi-valued logic type. The internal algorithms of the simulator are optimized to work with values of that type.

This logic type, and all of the logic operations available for the type, must be defined as part of a VHDL interface. Users of the interface will design with a subset of VHDL that can be easily translated to the notation used by the target

Our ASICs

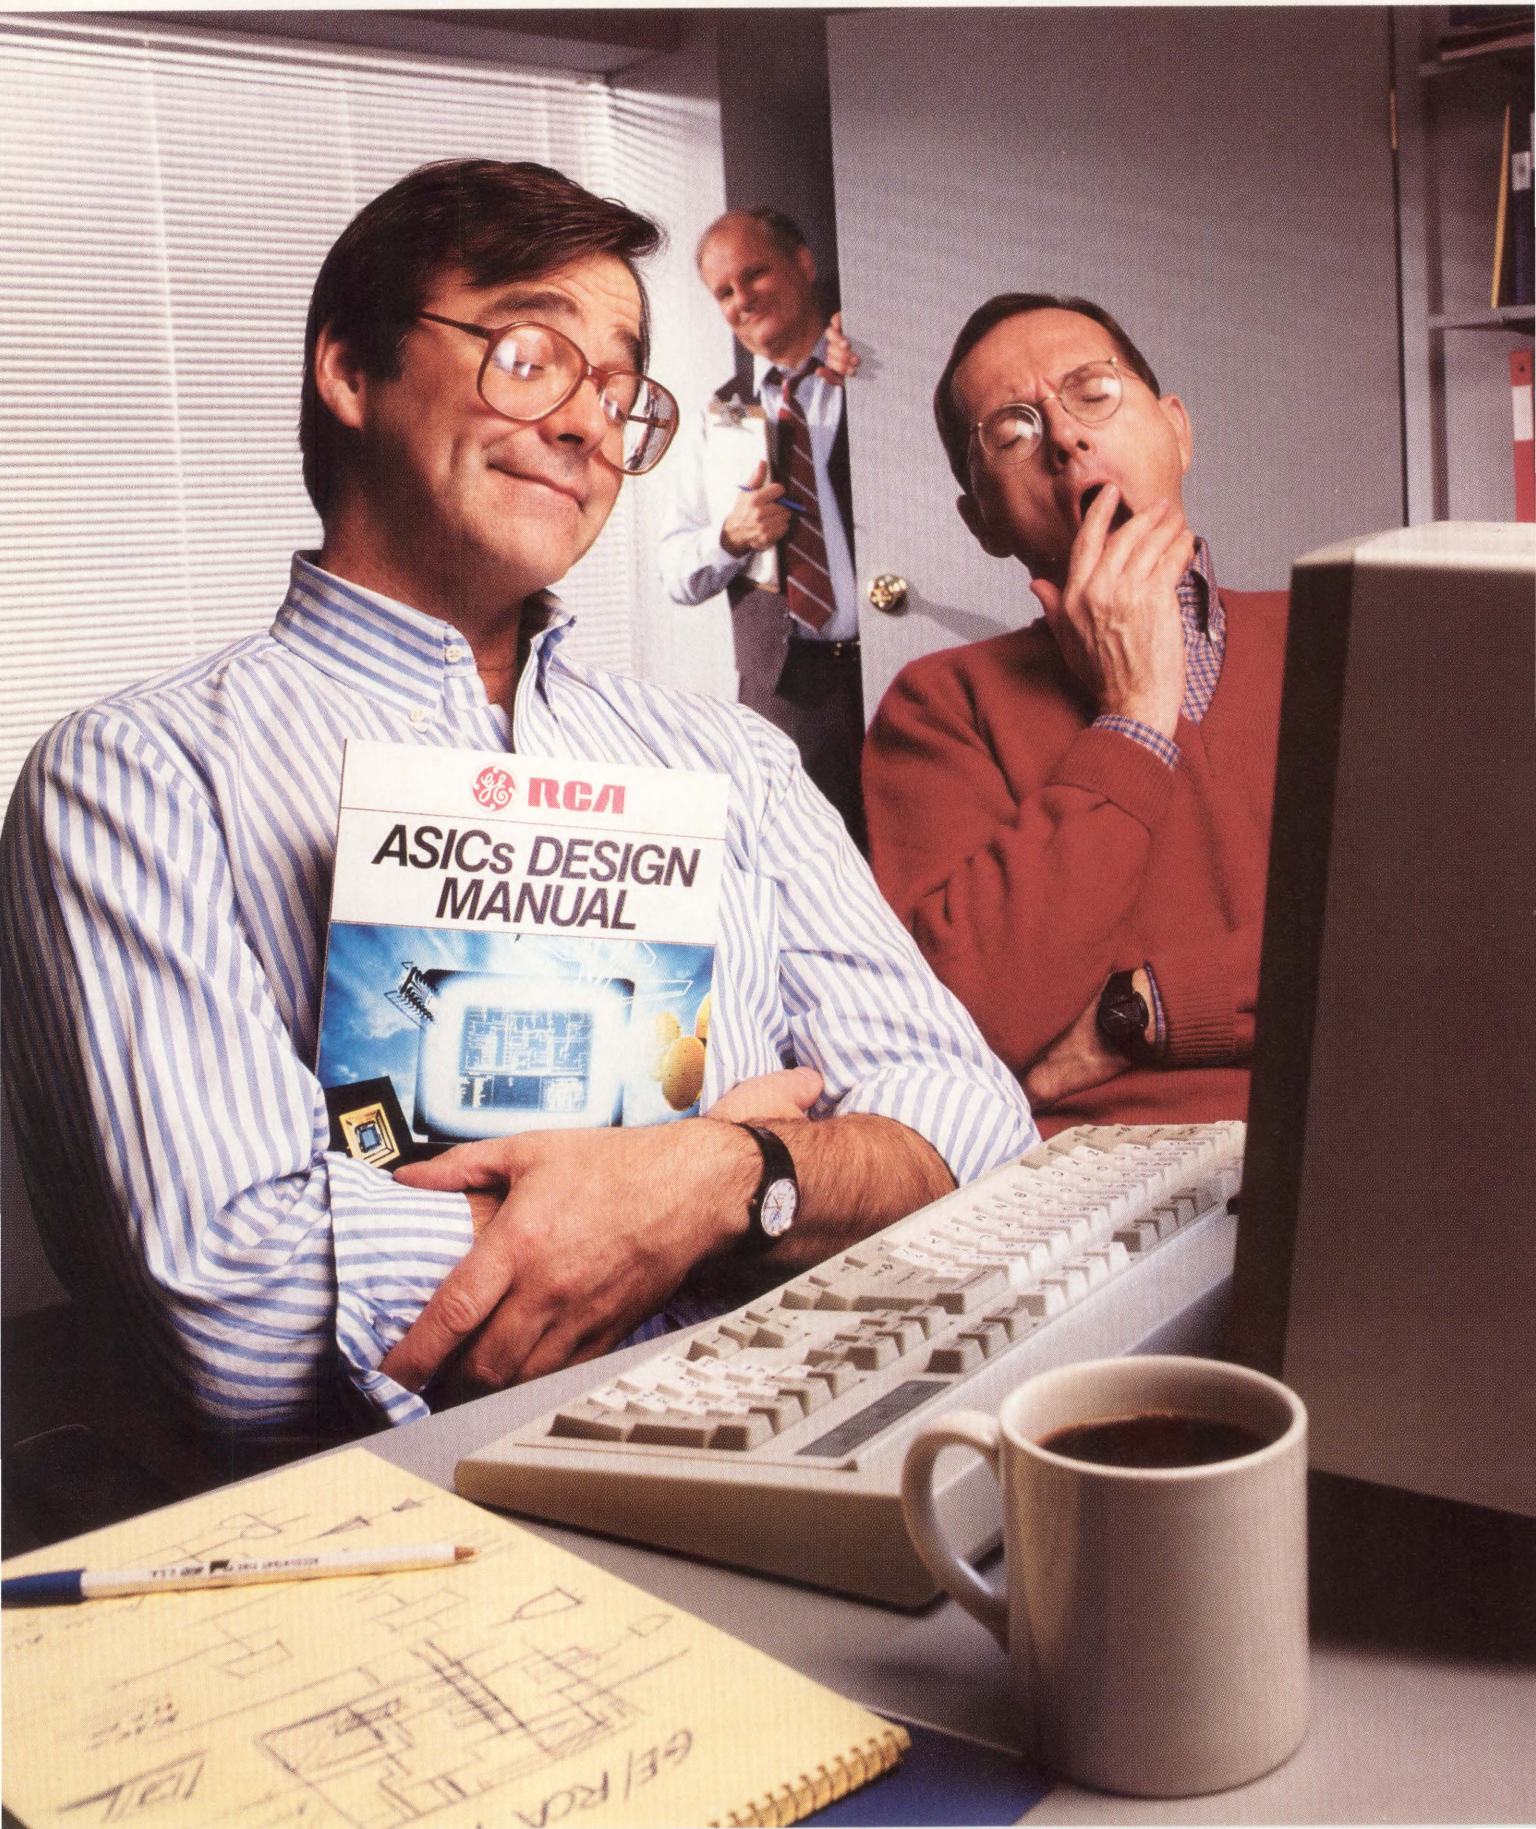

are boring.

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

You've heard all about the excitement of ASICs.

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

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

How to get first-time success.

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

The standard cell advantage.

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

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

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

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

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

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

Gate arrays, too.

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

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

We also alternate source the LSI Logic 5000 series.

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

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

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

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

USA

GE Solid State

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

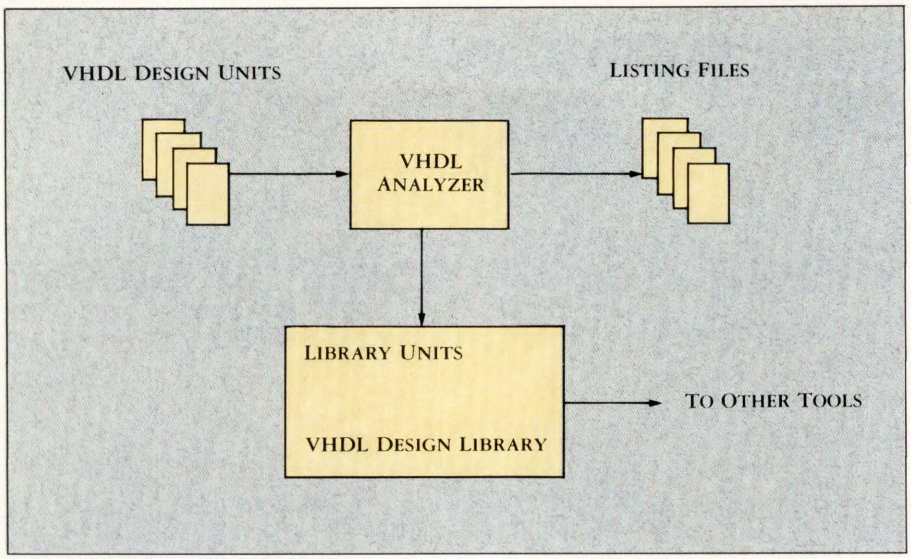

Figure 2. Analyzing VHDL descriptions.

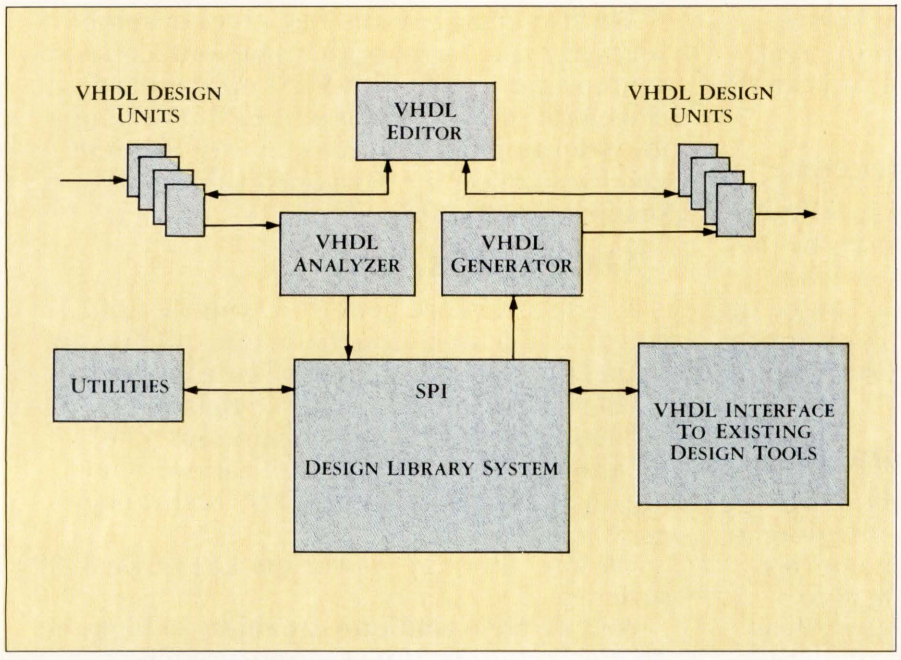

Figure 3. The VHDL tool integration platform.

simulation environment. Depending upon the features of the existing simulator, this may range from a very restricted subset to one that includes many of the features of IEEE Standard VHDL.

■ INTEGRATING VHDL WITH EXISTING TOOLS

In order to be most effective, VHDL must be integrated into a complete design methodology. The ability to describe the behavior of a new device in VHDL and then simulate that behavior is not enough. It must be possible to take the high-level VHDL description and map it into a low-level description that is suitable for use in layout and fabrication. Timing analysis, back-annotation, and gate-level simulation are still required to verify that a device will actually perform as expected. A complete VHDL-based design system ought to have all tools available today as well as a VHDL design and simulation capability.

Schematic Capture. VHDL provides a structural description capability that allows a designer to describe a system in terms of interconnected components. The notation used is essentially a textual, netlist-style of description. However, most designers would prefer to use a schematic capture tool to define a netlist. Although schematic information has no meaning in the VHDL domain, it is certainly possible to build an interface between an existing schematic capture tool and VHDL. One approach would define schematic symbols corresponding to VHDL design entities to be used in developing a netlist via schematic capture. Then it would generate a VHDL netlist description from the schematic. The VHDL netlist could then be analyzed into the design library.

Synthesis. A high-level behavioral model must be mapped into lower levels of design before accurate timing information can be obtained and the device fabricated. VHDL supports this process by providing features at different levels of abstraction: processes for high-level behavioral design, data-flow statements for intermediate levels, and structural (netlist) statements for low-levels. However, the manual transformation of a design from one level to another can be very time-consuming.

At IBM, this process has been automated to a great extent. VHDL is used primarily for developing high-level behavioral models of devices. Once such models simulate correctly, the VHDL code is fed into a synthesis tool that automatically creates an equivalent low-level design, also expressed in VHDL. This system has been in use at IBM for more than a year.

Physical Design. The ultimate goal of all hardware design is to define a network of primitive components that can be fabricated. As a design is refined, the behavioral model might be transformed first into a register-transfer style description in which the control paths and data paths are distinguished. Eventually it will be reduced to a structural description in which all of the functions have been encapsulated within primitive subcomponents.

At this point, a layout tool can be run on the design, and timing analysis tools can be applied to determine the expected delays within the design. Since VHDL does not support physical design, a different collection of tools and notations are required for this phase. Nonetheless it must be possible to bridge the gap from an abstract VHDL description to a physical layout.

Back-Annotation. Behavioral models written in VHDL must be back-annotated with data derived from physical design in order to simulate with precise timing information. Because the final behavioral model will be shipped as documentation along with a fabricated part, the model must be as accurate as possible—it will be used to model the real device in the design of new systems. VHDL behavioral models tend to be fairly abstract, so there is no simple correspondence between the elements of the fabricated part and the statements in the behavioral model. Specialized tools are needed in this area.

Gate-Level Simulation. After a device is designed, but before it is manufactured, it is still a good practice to perform a full gate-level simulation of the device using

Simulating VHDL in a Non-VHDL Environment

A behavioral language such as VHDL is intended to provide descriptions that can be simulated. The language assumes a simulator and the simulator actualizes the language. To implement the language with a given simulator, two mappings are required—one from the objects manipulated by the language (signals and components) to the simulator's data structures and one from the dynamic semantics of the language to the simulator's run-time routines.

Other major issues were considered before deciding on a VHDL interface to GE's proprietary mixed-mode, true-value fault simulator (MIMIC). We wanted to decompose the VHDL-to-MIMIC problem into two subproblems: VHDL to an intermediate form, and then this intermediate form to MIMIC, where the intermediate representation is standardized or widely accepted by the tool builders. We chose to use the intermediate representation provided by CLSI's Design Library System (DLS). Since the cumulative experience with VHDL for model development was still limited, we thought it wise to initially restrict the implementation to a subset—one that includes the behavioral part of VHDL, except for user-defined attributes and data types, such as file and access.

■ Simulating VHDL in the MIMIC Environment

MIMIC supports both true-value and fault simulation of circuits containing mixed components at the switch, gate, and behavioral levels. We expect that future designs will contain imported VHDL macrocells connected to subnetworks at all levels, and that all-VHDL models will be refined into interconnections of components at lower levels of abstraction. If so, behavioral models written in VHDL must coexist in this environment with non-VHDL models at all levels.

Since we wanted to augment the existing MIMIC environment with VHDL capabilities, we decided to incorporate VHDL's data types and semantics into our existing environment rather than building a "pure" VHDL environment. Although our implementation can be used to simulate pure VHDL descriptions, it would be impractical to limit simulation to VHDL models alone and totally ignore existing design practices. For VHDL to be useful, it must be integrated with the existing investment in tools, training, models, and methodologies. Otherwise, its potential never will be realized.

■ Mapping VHDL's Constructs

Mapping VHDL into our existing simulation environment was a twofold task: enhancing the simulator's data structures to support record, enumeration, and real data types; and enhancing the simulation run-time routines to support processes, signal assignments with waveforms, wait-statements, predefined attributes, and so on. Based on these mappings, we built a SPI-to-MIMIC translator that generates an equivalent C behavioral model from the DLS intermediate representation of a VHDL behavioral model, accessed via the SPI. The C behavioral model contains the C version of VHDL statements, with VHDL's dynamic semantics supported by function calls to construct waveforms, to initialize and assign to signals, to suspend and reactivate processes. Each process, for example, is represented as a C switch statement with case alternatives from 0 to N, where N is the number of VHDL wait statements. The bottom of the switch statement reactivates the process by scheduling its reactivation in zero-time. The type and object declarations in VHDL packages, entities, and block statements are passed to the simulator via "include" files also generated from the intermediate representation.

The internal interface between VHDL and non-VHDL descriptions is accomplished via a predefined multivalued enumeration type that includes all of the possible value/strength combinations currently used for gate and switch levels. This same enumeration type is used when a VHDL port is directly connected to a binary primary input and/or output to provide stuck-at fault injection or observation.

Yefim Shor
GE Solid State
Somerville, N.J.

timing data from the physical design phase. Gate-level simulation, while expensive in the course of developing a design, is still a cost-effective mechanism for verifying the end result. A VHDL design environment should make effective use of existing simulation capabilities, which include efficient gate-level software simulators, accelerators, and fault simulators.

■ Developing a Standard VHDL Design Environment

It is clear that the integration of VHDL with existing design tools is necessary if the language is to be useful. But how will such integration occur? Today, "integrated" design systems are typically supplied by each tool vendor. Yet end users still spend a lot of time trying to make tools from different vendors work in harmony. Since VHDL is designed to be a standard medium of communication among tools, it should be possible to use the language itself as a focal point for tool integration, even for tools from different vendors.

What's needed is an open architecture for VHDL-based hardware design tools that allows designers complete access to design data. If such an environment were available, new design tools could be developed by many different vendors, even those without the resources to develop a complete tool set, and all such tools would be automatically integrated via the common architecture. Similarly, existing design tools could be "rehosted" to operate based on the open architecture.

The heart of an open architecture system is the database schema or file formats that it defines for managing data, and the operations it provides on those structures. An IEEE working group is now pursuing the development of a standard format for VHDL design data. This group, the VHDL Intermediate Form Analysis and Standardization Group (VIFASG), meets about every three months to discuss the features and characteristics required to define an intermediate form for VHDL that is capable of supporting the operation of diverse, VHDL-oriented design tools. At present, this group has just completed the definition of requirements for such an intermediate form. They are now examining

One-volt drive gate

Battery-powered products reach new performance heights with our CMOS-4L gate arrays.

For fast answers, call us at:

USA Tel:1-800-632-3531. TWX:910-379-6985.
W. Germany Tel:0211-650302. Telex:8589960.
The Netherlands Tel:040-445-845. Telex:51923.
Sweden Tel:08-753-6020. Telex:13839.
France Tel:1-3946-9617. Telex:699499.
Italy Tel:02-6709108. Telex:315355.

C&C Computers and Communications

arrays.

Your battery-powered system will really take off with the new CMOS-4L gate arrays from NEC. CMOS-4L arrays let you add functions freely. They drastically reduce component count, increase system reliability and require a mere 0.01μA standby current.

Lap-top computers, handheld terminals, game machines – all become smaller, smarter and more reliable when you design-in CMOS-4L gate arrays.

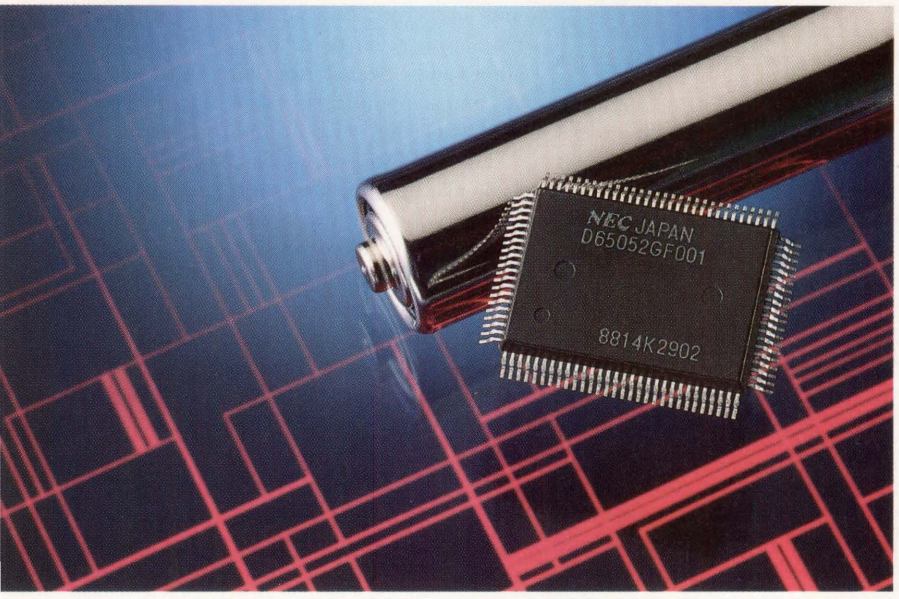

You can choose from six masters with arrays counts ranging from 858 to 5,632 gates.

Features:

- ☐ Low power supply voltage: 1.0 ~ 3.6V
- ☐ Low standby power consumption: 0.01μA typical (V_{DD}=1.5V)
- ☐ High speed: 10ns internal gate delay
- ☐ High drive capability: I_{OL}=3mA

NEC offers the broadest line of CMOS gate arrays in the industry. Call us today for information on our seven families, with 43 masters, including the new CMOS-4L for battery-powered products.

Master	μPD65007	65014	65026	65033	65045	65052
Gate count (2-input NAND)	858	1656	2457	3360	4320	5632
I/O pads	62	82	100	106	120	138
Internal gate delay	10ns (2-input NAND, F/O=3, L=3mm, V_{DD}=1.5V, Ta=25°C)					
Output drive capability	I_{OL}=3mA I_{OH}=1mA (V_{DD}=1.5V, Ta=25°C)					
Power supply voltage	1.0 ~ 3.6V					
Ambient temperature	0 ~ +70°C					
I/O level	CMOS compatible					

UK Tel:0908-691133. Telex:826791.
Hong Kong Tel:3-755-9008. Telex:54561.
Taiwan Tel:02-522-4192. Telex:22372.
Singapore Tel:4819881. Telex:39726.
Australia Tel:03-267-6355. Telex:38343.

NEC

"strawman" intermediate forms.

One such candidate is an intermediate representation of VHDL design data developed by CAD Language Systems Inc. (CLSI). The development of this intermediate form was partially funded by General Dynamics and Unisys in the hope of defining a widely acceptable basis for VHDL tool development. CLSI's intermediate representation has since evolved into a complete tool development environment known as the VHDL Tool Integration Platform, or VTIP. The VTIP includes all of the fundamental front-end software necessary for the development and integration of application-specific VHDL tools.

■ THE VHDL TOOL INTEGRATION PLATFORM

A VHDL analyzer and a VHDL design library system are the primary constituents of the VHDL Tool Integration Platform. In addition, a number of other tools and utilities are provided for building a complete VHDL-based design environment. The major components of VTIP are shown in Figure 3. The VTIP is already in use within a number of companies, both for building new VHDL tools and for building VHDL interfaces to existing tools.

The Design Library System. The foundation of the VHDL Tool Integration Platform is the Design Library System, or DLS, which is the repository for design data. The DLS is built upon the underlying host file system for efficiency, yet it hides the details of any particular host from design tools that use it. As a result, design tools based on the DLS can be written so they will port easily to any host on which the DLS itself is available. The DLS is currently available on VAX/VMS and Sun/UNIX.

The Design Library System supports any number of design libraries, limited only by the capabilities of the host file system. Each library may contain any number of library units, and each unit may be represented with several views. (The concept of a view in the DLS is similar to the concept of views in EDIF.) In fact, The Design Library System is not specific to VHDL design data. Only one of the DLS views, called the Textual View, is designed with that in mind. Other views support non-VHDL data and tools.

A library unit in a DLS library consists of a collection of records that are interconnected to form a directed, attributed graph structure. Each record type is defined in an object-oriented fashion. Each represents a particular kind of design data object, so all of the information about a given design object can be retrieved from a single point in the structure. An inheritance mechanism makes it possible to define variants of more general object types that have all the attributes of the parent object type plus additional, more specific attributes. The equivalent record structures of a few of the more common object types are illustrated in Figure 4.

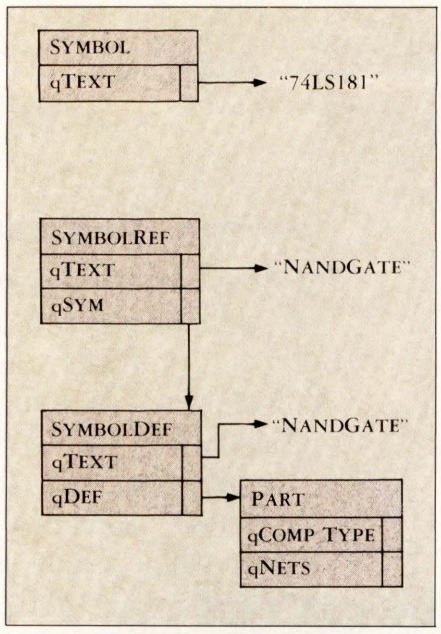

Figure 4. Equivalent record structures of common object types in the DLS.

The Software Procedural Interface. Library units stored within the DLS are accessed by tools via the Software Procedural Interface, or SPI. The SPI is a set of functions that allow a tool to create or modify the contents of a library unit or a library within the Design Library System. All data stored within the DLS is accessible via the SPI, except that certain system-maintained attributes, such as time-stamps, are read-only. The SPI performs extensive error-checking to guarantee that library units constructed or manipulated by tools are internally consistent. Written in C, the SPI can be used by any tool integrated within the VTIP.

The SPI provides both low-level and high-level services to tools. Low-level services include routines for creating, modifying, and retrieving data within a library unit. High-level services include routines for expression manipulation, symbol management, and library management. All of these routines are available via one consistent interface that makes the construction of design tools relatively easy. Since the SPI isolates the tool from the details of how data is stored, a tool-builder can focus on the engineering problem to be solved rather than worry about low-level programming issues.

The most basic routines provided by the SPI allow a tool to set or retrieve the values of attributes of a design object within a library unit. Each attribute of a design object has a value retrieval function with the same name as the attribute. Each attribute (other than system-maintained attributes) also has a value-setting function named Set followed by the name of the attribute. For example, the routines qText and SetqText exist to get and set, respectively, the qText attribute of any Symbol object, including the SymbolDef and SymbolRef variants of object type Symbol (see Figure 4).

Additional SPI routines provide extensive list processing capabilities, including the ability to search for symbols with a given name declared in a given region of text. Constrained lists can be created so that only certain kinds of data can be added to the list. This facility allows a tool-builder to use the SPI to perform various integrity checks within an algorithm. One application of such lists in the DLS allows additional attributes to be associated with any object, similar to the property lists associated with atoms in LISP.

The SPI also provides a comprehensive error-handling mechanism that can be tailored to support the needs of each tool. Tool-specific error codes, messages, and severity levels can be registered with the SPI. The SPI uses this information to determine how to respond when an error is reported by the tool. Automatic logging of error messages and automatic abort on an error with a given severity level may be requested via SPI calls.

The VHDL Analyzer. CLSI's VHDL Analyzer reads a VHDL source file, verifies that each design unit in the file is syntactically valid, and creates one Textual View library unit for each valid VHDL design unit in the source file. The Analyzer also checks that the static semantics of each library unit are correct. Only correct library units are placed into a design library. Before a library unit is put away into the library, all name references within the unit are resolved so that the data within the library unit is unambiguous and can be interpreted easily by downstream tools.

Once the VHDL Analyzer has entered a library unit into a design library, the library unit may be updated via SPI calls by other tools. This provides a mechanism for the direct back-annotation of VHDL library units with data derived from physical design. In particular, it is possible to directly modify the delay characteristics of a design via this facility. Other physical data such as loading characteristics or fan-out may also be added to an existing VHDL library unit.

Although allowing tools to update DLS library units is essential for back-annota-

TEXAS INSTRUMENTS REPORTS ON

SYSTEMS LOGIC

IN THE ERA OF

MegaChip TECHNOLOGIES

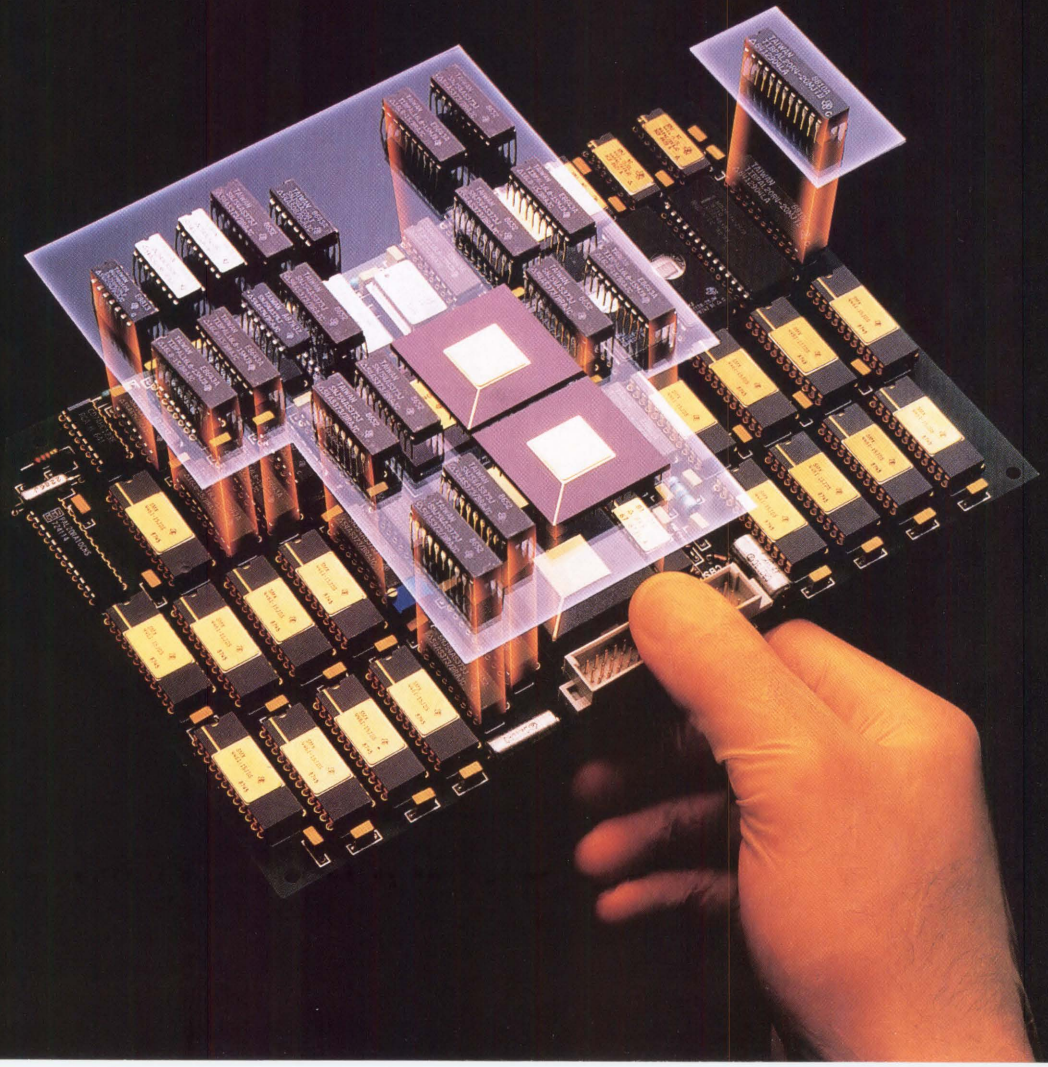

Systems logic in the Era of MegaChip Technologies:

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.

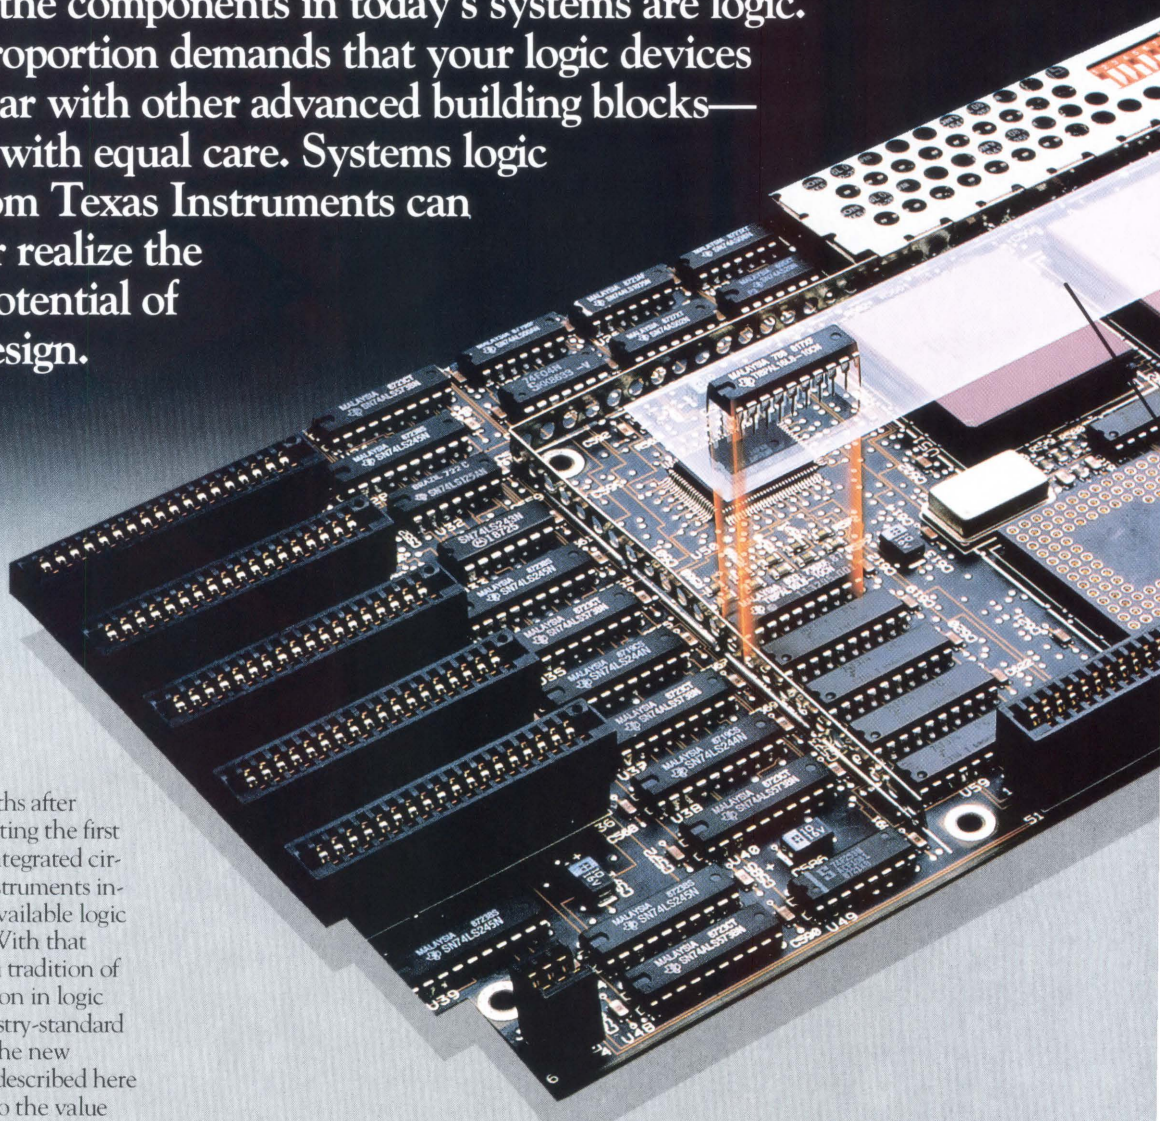

Within months after demonstrating the first working integrated circuit 30 years ago, Texas Instruments introduced a commercially available logic function, an RS flip-flop. With that beginning, TI established a tradition of development and innovation in logic that encompasses the industry-standard SN54/74 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:

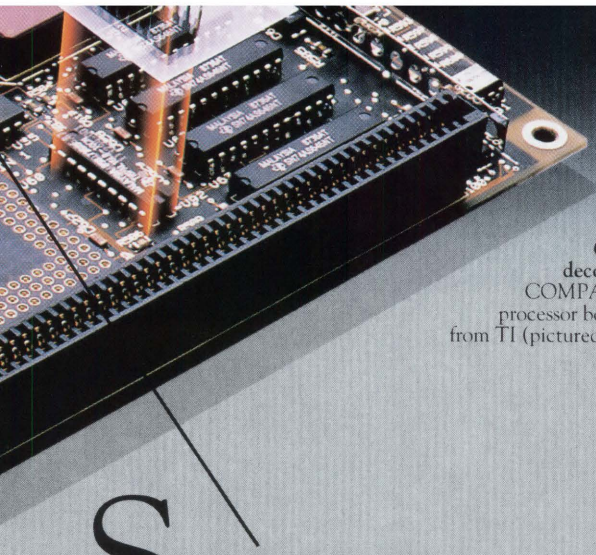

Contributing significantly to fast address decoding in speed-critical paths of the COMPAQ DESKPRO 386/20™ personal computer processor board are two TIBPAL16L8-10 PAL circuits from TI (pictured above a segment of the board).

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.

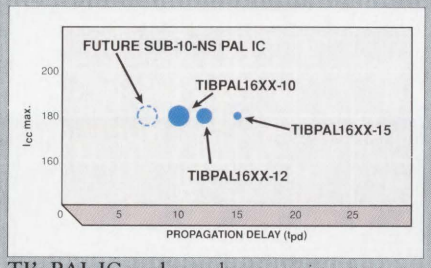

TI's PAL IC road map shows consistent power and consistently higher speeds, with even faster versions on the way.

Turn page for more information

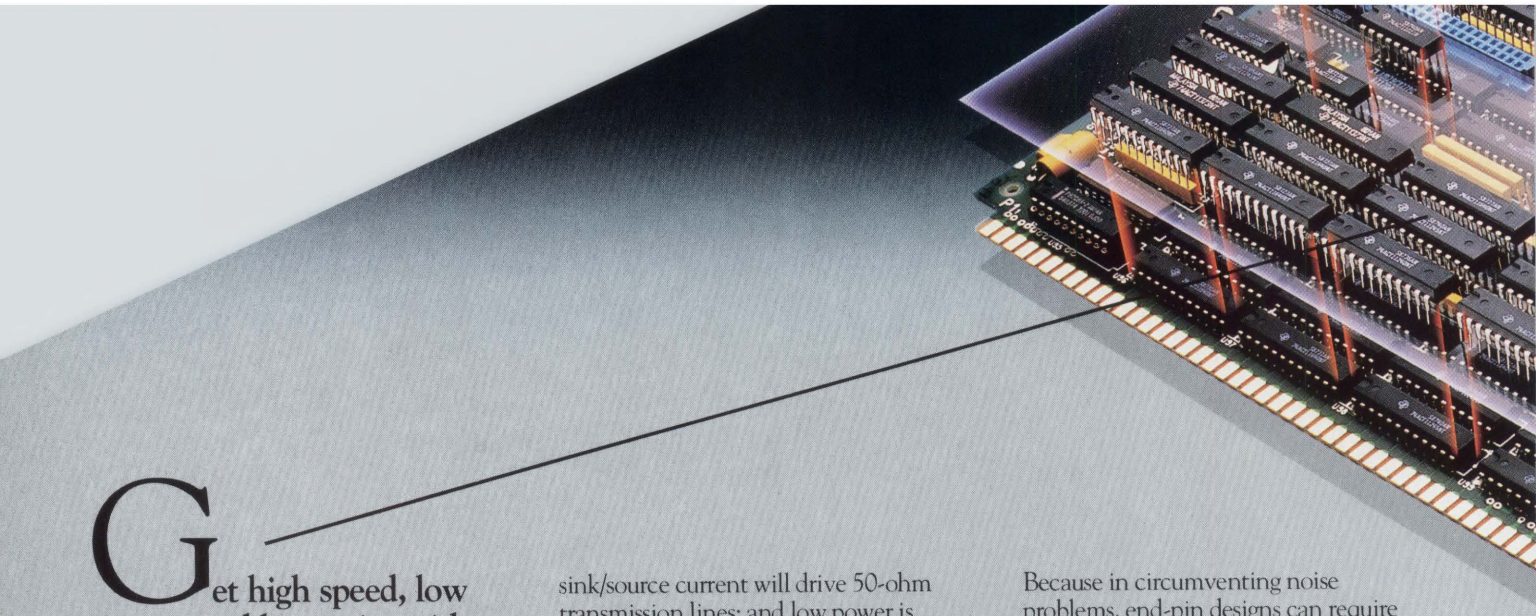

Get high speed, low power, and low noise with TI's broad ACL family.

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.

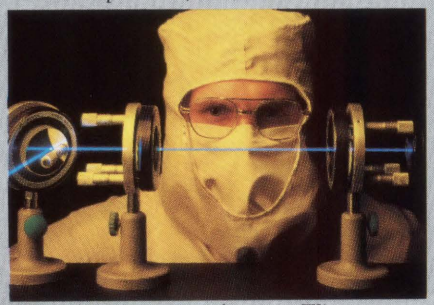

When every nanosecond counts, TI's new high-performance ACL family can help you significantly improve system speed.

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.

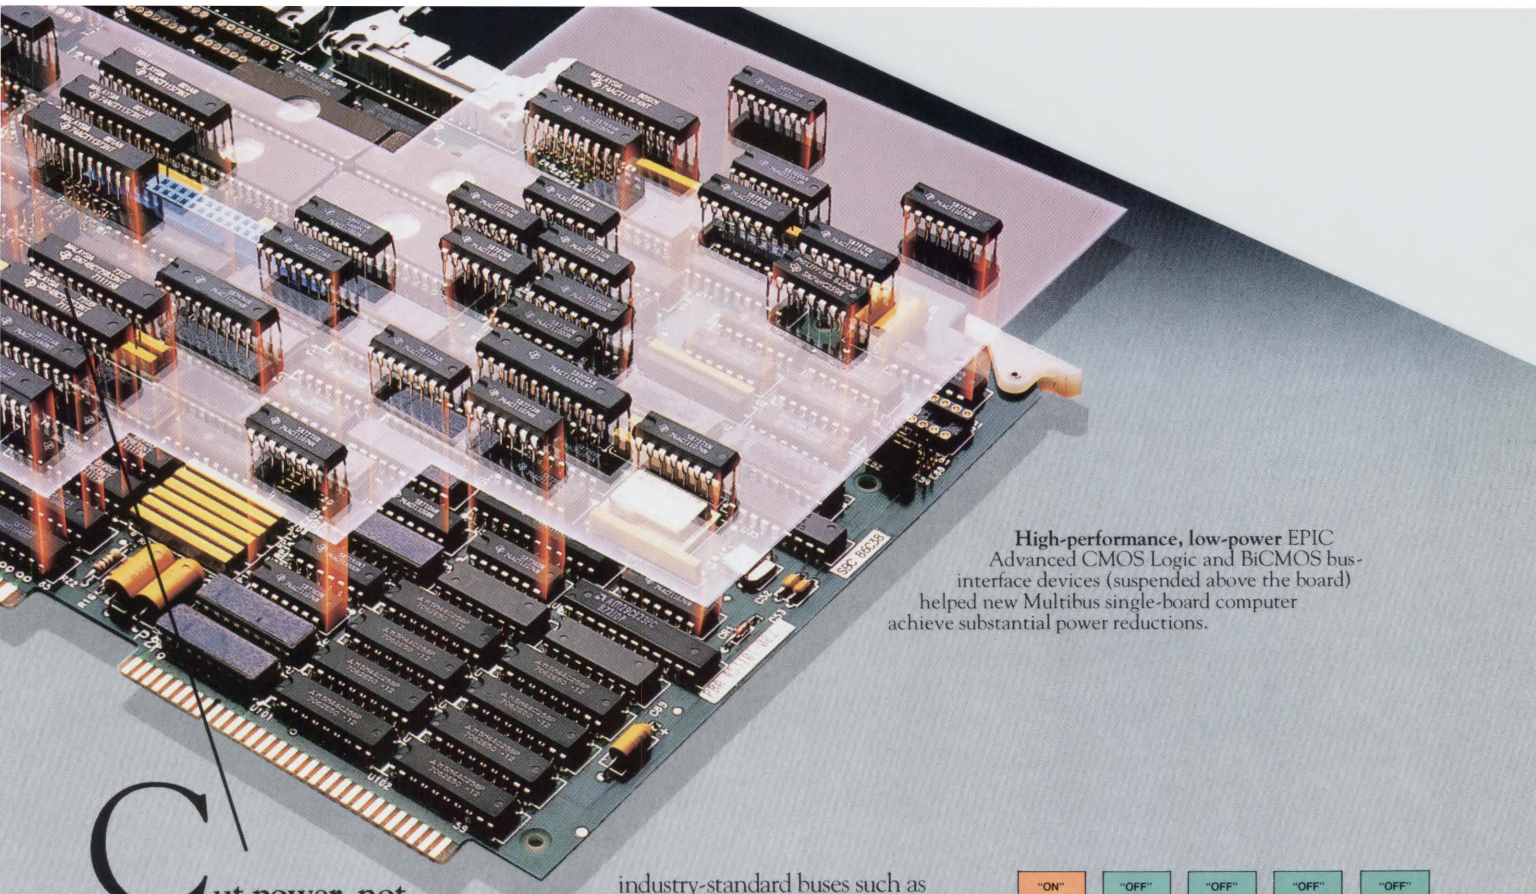

High-performance, low-power EPIC Advanced CMOS Logic and BiCMOS bus-interface devices (suspended above the board) helped new Multibus single-board computer achieve substantial power reductions.

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.

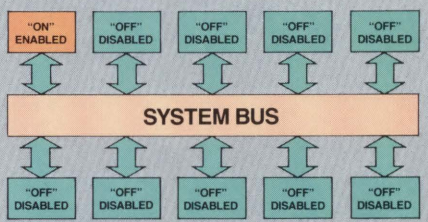

An innovative circuit design in TI's BiCMOS bus-interface logic helps lower disabled currents. This is key to overall power savings because in a typical bus network only one device is enabled at a time.

Turn page for more information

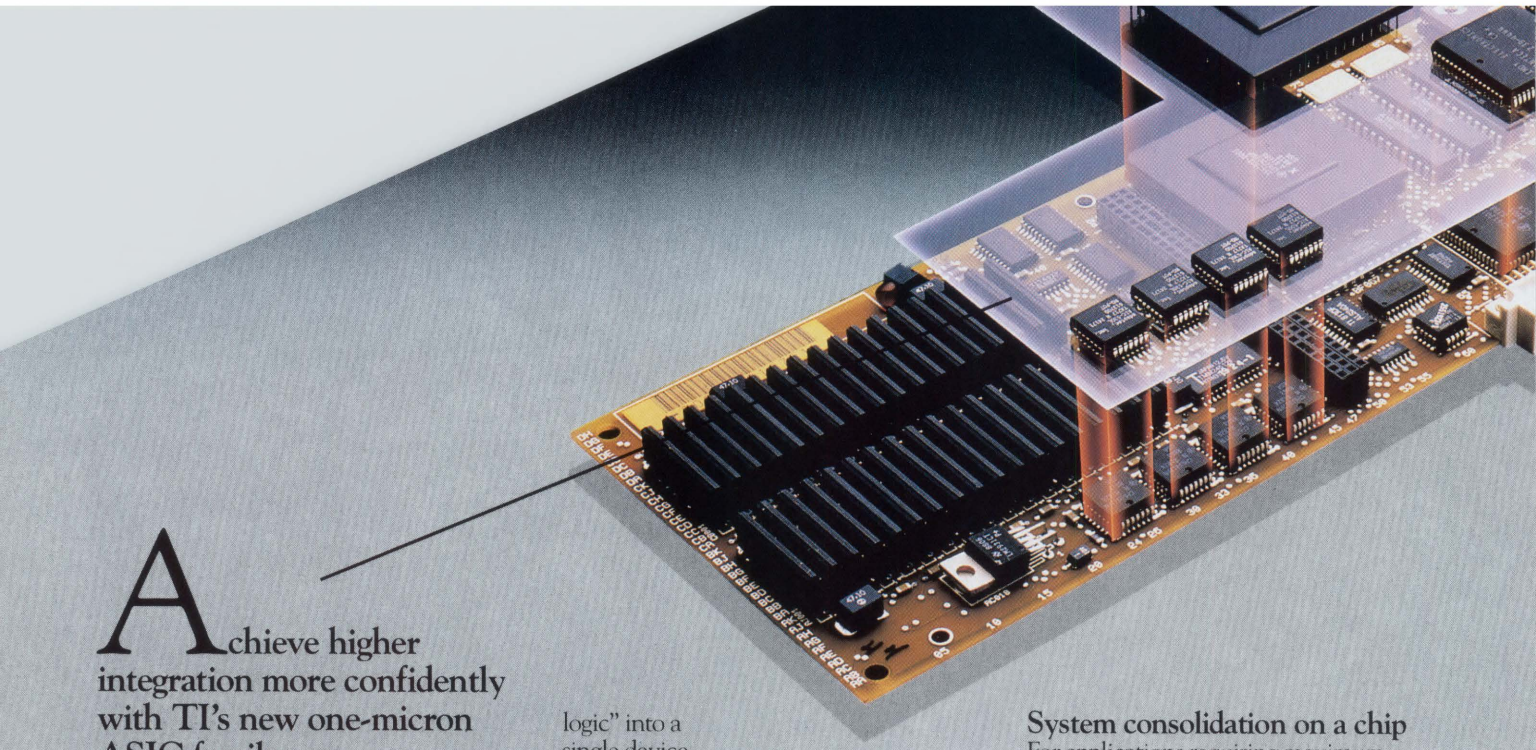

Achieve higher integration more confidently with TI's new one-micron ASIC family.

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

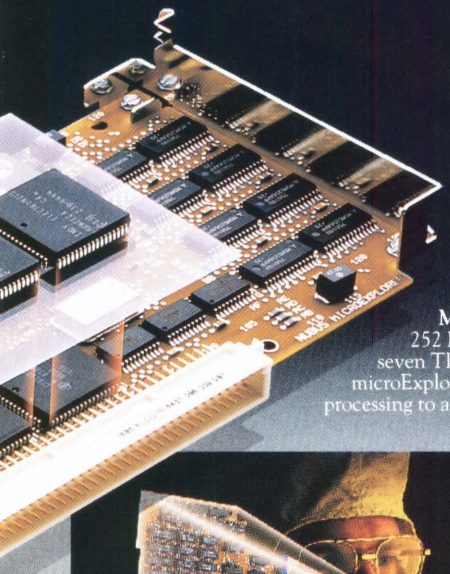

Major logic consolidation, the equivalent of 252 MSI and LSI devices, was possible using the seven TI ASIC functions shown above the microExplorer™ board which brings symbolic processing to a MacIntosh II® desktop computer.

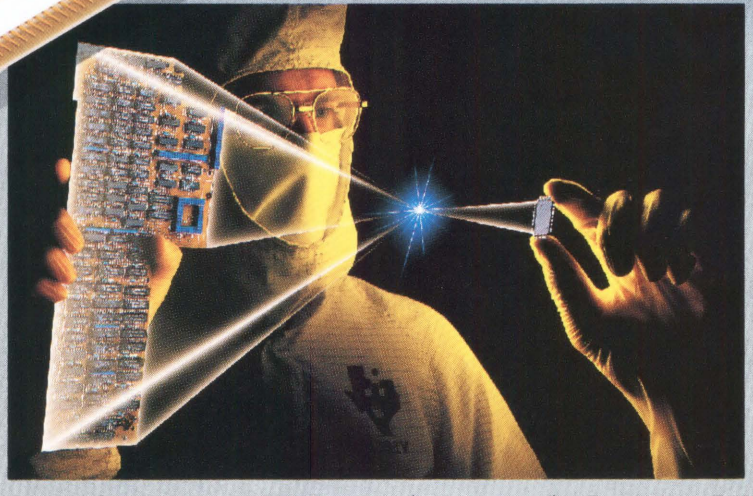

480 ps for a two-input NAND gate with a fanout of three; flip-flop toggle rates range up to 208 MHz. Both series offer output and bidirectional buffers with variable slew-rate control. And both series are fabricated in TI's high-performance EPIC process.

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

SDV$\emptyset$83VD8$\emptyset\emptyset$C

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

Yes, please send me the following:
RY$\emptyset$1 ____ ASIC Information Packet
DZ$\emptyset$1 ____ Programmable Logic Device Data Book
CA$\emptyset$1 ____ ACL/BiCMOS Information Packet
CB$\emptyset$1 ____ BiCMOS Data Sheet Packet

NAME

TITLE

COMPANY

ADDRESS

CITY STATE ZIP

AREA CODE TELEPHONE EXT.

TEXAS INSTRUMENTS

Accelerating VHDL

Simulation of large designs in VHDL requires acceleration, which is now possible through the use of a prototype tool developed by CLSI and Zycad. This tool creates a bridge between CLSI's VHDL Tool Integration Platform and Zycad's ZILOS simulation environment, translating VHDL design data from the VTIP into equivalent ZILOS files.

Translation of VHDL descriptions to ZILOS data files is based on a VHDL package that defines basic data types and primitive components supported by the accelerator. Any VHDL structural design containing instances of only those primitive components can be accelerated by translating it to equivalent ZILOS descriptions. Instantiations of non-primitive components are handled by translating them into ZILOS macro calls and by translating the instantiated design entity into a ZILOS macro definition. Thus hierarchical designs are elaborated via the macro-instantiation process built into ZILOS.

The prototype can also handle some VHDL data-flow (signal-assignment) statements. Zycad accelerators are very good at evaluating concurrent gate-level models—the problem is in synthesizing concurrent gate-level models from signal assignments. The translator generates a network of primitive gates to compute the value of the source; it generates a delay element to assign the resulting value to the target of the signal assignment with the appropriate delay.

The current prototype cannot handle VHDL behavioral descriptions, which consist of sequentially executed statements that form program-like models. Such descriptions could be handled by translating them to ZILOS Behavioral Language, which supports behavioral simulation in conjunction with accelerated simulation of structural descriptions.

Dick Schlotfeldt
Zycad Corp.
St. Paul, Minn.

tion, such updates are potentially dangerous, since a tool could modify one unit in such a way that it invalidates other dependent library units. For example, deleting a component description in one unit will invalidate any other unit that instantiates the component. In order to detect such a situation, the SPI automatically timestamps library units when they are created and again when they are updated, and uniquely labels each node within a unit. When a library unit is opened for processing, the SPI automatically verifies that all dependencies upon other units are valid.

The VHDL Generator. In order to export the results of back-annotation in a human-readable form, it must be possible to transform VHDL library units back into source form. The VHDL Generator performs this task: it produces VHDL source text from any Textual View library unit. Modifications introduced into the library unit are converted into equivalent VHDL text wherever possible. Where such modifications cannot be directly represented as VHDL, comments are inserted into the output source file.

Taken together, the VHDL Analyzer and Generator provide a complete import-export capability for VHDL source code. This is an essential characteristic of a complete VHDL design environment, since the primary function of VHDL is to provide a standard medium of communication among all the parties involved in the design and use of a device.

■ Building VHDL Interfaces to Existing Systems

From the end-user's point of view, the adoption of VHDL involves learning yet another design language. However, it need not involve learning to use a completely new set of design tools. More importantly, it need not involve scrapping millions of dollars of investment in existing systems. Adding a VHDL interface to an existing system is a cost-effective way to plug into the "software bus" represented by VHDL in order to gain a high-bandwidth communication channel with others in the industry.

Still, there is a lot of work involved in building a VHDL interface. Most companies do not yet have many resident VHDL experts, so, education is an issue. Then there is the problem of matching the capabilities of existing tools with the semantics of VHDL, to make sure that VHDL descriptions from elsewhere are correctly interpreted. Since VHDL's features far exceed those of many simulation systems, a subset issue also arises. Finally, the software implementing the interface must be designed and built. All of these issues must be dealt with in order to create a useful VHDL capability.

Several organizations have begun developing VHDL interfaces to existing design tools using the VTIP. For example, GE Solid State has developed an interface to its MIMIC behavioral simulator (see "Simulating VHDL in a Non-VHDL Environment" on page 45). Similarly, CLSI and Zycad have developed a prototype interface to Zycad's ZILOS simulation environment (see "Accelerating VHDL," this page). CLSI is also developing interfaces to other tools, including a prototype interface to Teradyne EDA's AIDA simulator.

As time passes, two trends can be expected. The VHDL-specific simulation systems offered by some vendors will acquire more and more of the capabilities required by the ideal VHDL design environment discussed above. At the same time, the tools for which VHDL interfaces have been developed will be extended to make ever greater use of the capabilities of VHDL.

In the long run, both paths will meet in a fully general, integrated design environment that supports VHDL. In the short run, the interface approach seems to provide the maximum use of existing investment in design tools. ■

References

IEEE Standard 1076-1987, adopted on December 10, 1987. The Standard, which consists of a Language Reference Manual for IEEE Standard VHDL, is available both from the IEEE and from CAD Language Systems Inc.

SAUNDERS, LARRY F. 1987. "The IBM VHDL Design System," 24th Design Automation Conference, Las Vegas, Nev.

About the Author

Erich Marschner *is the technical director of CAD Language Systems Inc., which he cofounded in 1986. He also has held positions at Intermetrics Inc. and at Advanced Computer Techniques (Roslyn, Va.). Erich earned the B.S. in computer science and the B.A. in medieval language and literature from the University of Maryland in 1980.*

ASIC TESTING

WITH HIGH FAULT COVERAGE

THERESA BUTZERIN, ARIF SAMAD, AND ERIC ARCHAMBEAU,
VLSI TECHNOLOGY INC., SAN JOSE, CALIF.

Methodology Converts Functional Block Tests Into Chip-Level Tests

ASICS are characterized by fast turn-around times and relatively low production volumes, so automation is essential at all stages in their design and manufacture. A high degree of automation is particularly important for test program development: an ASIC company can ill afford the time and expense associated with manually crafting an unique test program for every customer's ASIC design.

VLSI Technology has adopted a divide-and-conquer approach to generating test programs for complex ASICs. Each chip is conceptually partitioned, most often into the major functional blocks that make up the chip: RAMs, ROMs, datapaths, state machines, PLAs and megacells. Megacells are the building block equivalents of standard microprocessor peripherals such as the M84C00 CPU and the M84C40 serial I/O controller.

Once a chip has been partitioned, its designer applies a two-step test-generation process. First, he generates or writes vectors that test the overall system function as well as each functional block. Second, these vectors are combined to generate a program that can drive the automatic test equipment to test the fabricated device.

Through a common user interface, VLSI Technology's silicon compilers generate test vectors along with compiled circuit blocks. Megacells are tested using existing functional vectors that have been run through fault simulation to ensure high coverage. The designer provides vectors to test the overall system operation.

ILLUSTRATION BY RENEE KLEIN

```
write(address,value) - write value into addressed location
read(address,value) - read addressed location and check for value

procedure marchingOnesAndZeros;
begin
  for address := 0 upto numOfWords -1 do
    write(address,0);

  for address := 0 upto numOfWords -1 do begin
    read(address,0);
    write(address,1);
    read(address,1);
  end;

  for address := numOfWords - 1 downto 0 do begin
    read(address,1);
    write(address,0);
    read(address,0);
  end;
  repeat the steps above with 0's and 1's interchanged
end;
```

Figure 1. The "marching ones and zeroes" routine checks for opens and shorts, most of the decoder errors and some cell-interaction errors within RAM blocks.

Megacell test programs and the vectors generated by the compilers are described using a tester-independent form called the Vector Intermediate Format (VIF). After vectors have been generated for each functional block, VLSI's test program generation software integrates each set of patterns into a complete chip test and adapts them to one of the available ATE systems.

This article describes both the methodology used to generate tests for individual blocks and the methods used to incorporate them into a complete chip test.

■ VECTOR GENERATION FOR MEMORY

Test programs for RAM blocks must verify the two architectural components of RAMs: a memory array capable of storing information, and address decoders, that determine which cell is being written or read. The most significant faults in RAMs are categorized as follows (Breuer, 1976):

1) Opens and shorts, in which bits in the array cannot be set to desired values.

2) Pattern sensitivity faults cause the contents of a bit in the memory array to be affected by the contents of adjacent bits in the array. This problem is particularly acute with high-density RAMs.

3) Decoder faults lead to errors in addressing memory array cells.

Because RAM tests typically read and write each bit in the array, the number of tester cycles needed to test a RAM can quickly become prohibitive. Manufacturers of RAM chips often have specialized memory testers with built-in hardware that compensates for the number of required cycles. However, ASIC devices containing RAMs are tested on general-purpose testers that do not have special hardware for RAM pattern generation.

A test based on "marching ones and zeroes" has been shown to effectively test memories in a relatively small number of cycles (Breuer, 1976). This test checks for opens and shorts in the array, most decoder errors, and some cell-interaction errors.

As shown in Figure 1, this test addresses memory locations in ascending order, reading a 0 and then writing a 1 at each memory location. If a location has decoder or cell-interaction faults, then a subsequent location will change from a 0 to a 1 and the test will detect this on subsequent read cycles. The ration-ale for marching through the array in descending order of addresses is similar. This procedure results in 14n test cycles for n RAM cell locations.

Compared to RAM testing, ROM testing is relatively straightforward. A ROM is considered functional if the data read from the ROM corresponds to the data specified in the ROM code file by the user.

```
                procedure multiplierTest;
                begin
                    (1) multiply(00...000,00...000);
                    (2) multiply(11...111,11...111);
                    (3) multiply(011...11,011...11);
                    (4) multiply(0101010...;0101010...);
                    (5) multiply(00101...,01010...);
                    (6) multiply(1101010...,1101010...);
                    (7) multiply(1010101...,1101010...);
                    (8) shiftOnes;
                    (9) shiftZeros;
[ A ]           end;

procedure multiply(A,B) applies values A and B to the multiplier inputs
and checks the output for A*B

procedure setLeftMostBitsToOne sets the specified number of bits to 1
and sets the remaining bits to 0
e.g. setLeftMostBitsToOne(A,3) would set A to 0000...00111

procedure shiftOnes;
begin
  for numberOfAbitsToBeSetToOne = 1 to 3 do
    for numberOfBbitsToBeSetToOne = 1 to 3 do begin
      setLeftMostBitsToOne(A,numberOfAbitsToBeSetToOne);
      setLeftMostBitsToOne(B,numberOfBbitsToBeSetToZero);

      for numberOfAshifts := 1 upto lengthOf(A) - 1 do begin
        for numberOfBshifts := 1 upto lengthOf(B) - 1 do begin
          multiply(A,B) ;
          shiftLeft(B)
        end;
      shiftLeft(A)
    end
  end
end;

procedure shiftZeros is the same as shiftOnes with zeros and ones
reversed.

[ B ]
```

Figure 2. The nine steps in a multiplier test (a) result in 99% fault coverage. Steps eight and nine use the shiftOnes and shiftZeroes procedure (b) that exhaustively exercise the adders and multiplexers that make up the multiplier.

■ MULTIPLIERS

Test for multiplier blocks must account for the architecture of the block. For example, one type of multiplier generated by VLSI Technol-ogy's compilers is implemented using banks of adders connected by multiplexers. Two numbers are multiplied by shifting and adding sequences within the adder array.

The multipliers are tested by using the patterns illustrated in Figure 2a. The first seven steps look for global functionality and connectivity. Step eight calls the "shiftOnes" procedure defined in Figure 2b. Each multiplication within the shiftOnes procedure tests a specific portion of the adder array and the multiplexers between those adders.

The bit patterns created by the "SetLeftMostBitsToOne" procedure are shifted through

Great things happen when your design, test and quality management systems are integrated...

Quality goes up because key elements of your products' design, manufacturing and test process begin working in harmony. And this makes for a smoother, more efficient flow from your design department to final test.

Costs go down because your design people understand the needs of your test people and begin designing products that aren't production-test nightmares. This helps you move designs into production and test more quickly. The result is savings in time and money.

You get to market on time because your traditional bottlenecks are easy to overcome. GenRad's TRACS®, for example is a quality management system that helps you fine-tune your manufacturing process to eliminate the causes of faults. It can even help you identify design problems that can stall a new product introduction.

And GenRad can help because we have the newest in design and test simulation systems, a broad line of component and board testers that lets you choose the price performance best for you, and the quality management system that helps you win the fault prevention battle. What's more, we have integrated all of these systems with our powerful software so you won't have problems making them work together.

We have great things in store for your electronics manufacturing process. Call us at 1-800-4-GENRAD.

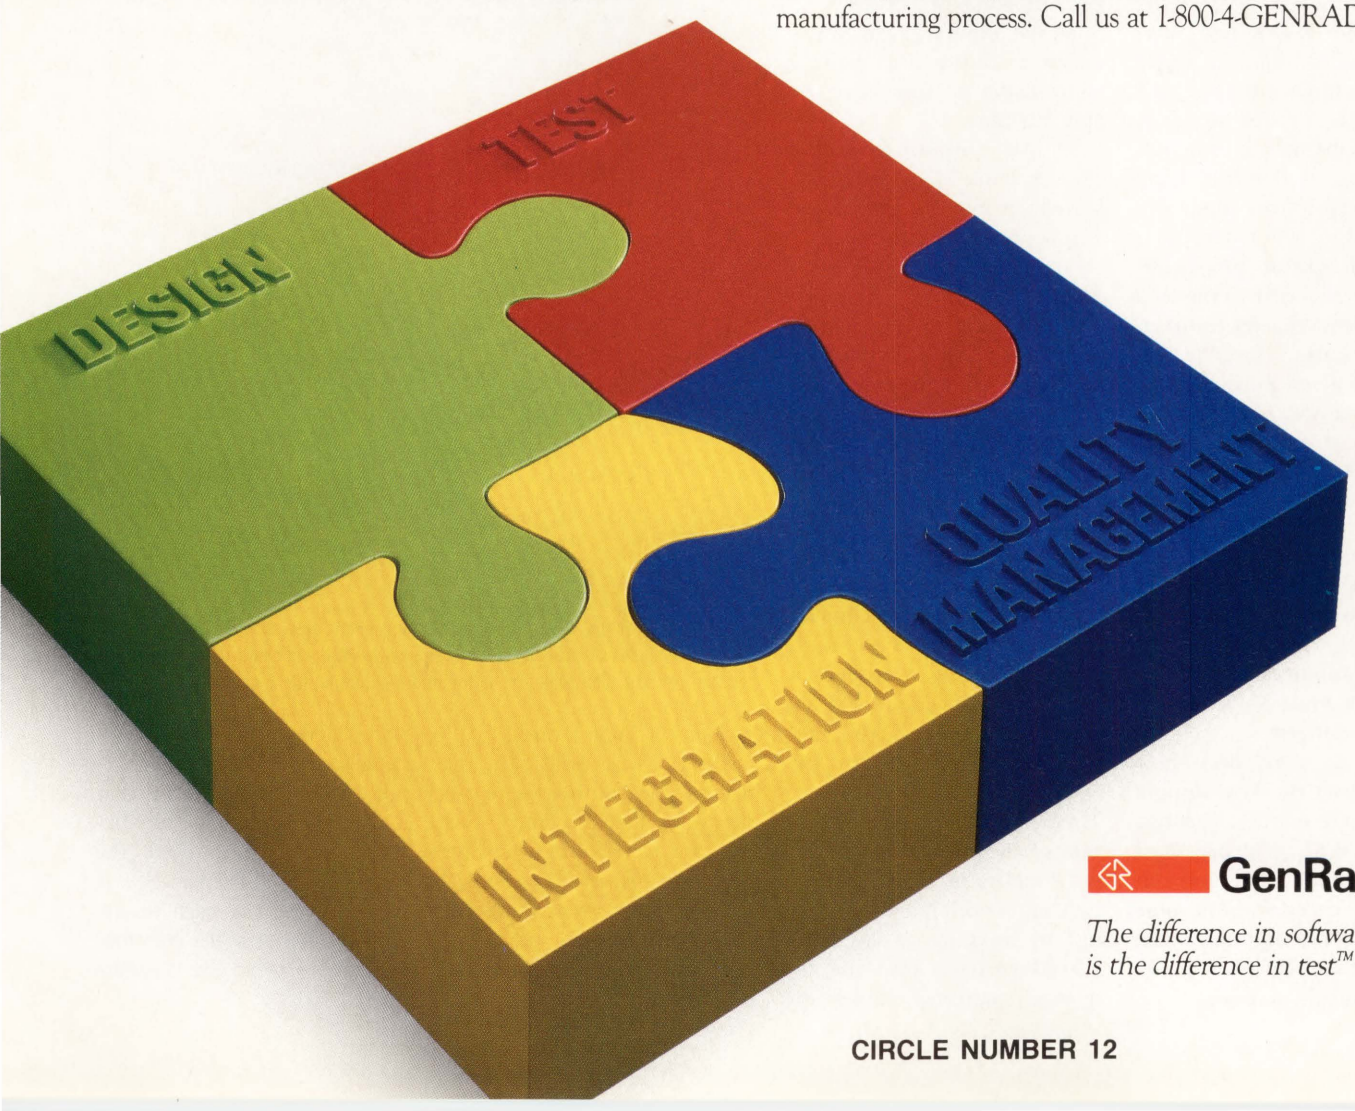

GenRad

The difference in software is the difference in test™

CIRCLE NUMBER 12

the array to test the entire array more exhaustively than the global patterns in Steps one through seven. Steps one through eight result in a fault coverage of about 97 percent, while invoking the analogous "shiftZeroes" procedure (step 9) raises the overall fault coverage to 99 percent.

■ VECTOR GENERATION FOR PLAS

There are several possible kinds of faults in PLAs, such as stuck-at faults, metal-bridging faults and cross-point faults (Abraham, 1986). VLSI Technology's PLA test pattern generation software represents the PLA as a two-level, AND-OR gate structure, a simplification that is reasonably effective for generating tests. It generates tests for all stuck-at faults on the network nodes and may also detect many bridging and cross-point faults.

The algorithm divides the faults in the PLA into three classes: input faults, AND-gate-output faults and PLA-output faults. To test input faults, the input pin under test is set to the desired value. A path is then sensitized through the circuit so that the logic value placed on the input can be viewed at a PLA output.

The algorithm propagates the logic value first to the outputs of the AND gates and then to the outputs of the OR gates. Because a given input pin can feed several AND gates, the algorithm arbitrarily chooses one of these gates (if necessary, the algorithm can backtrack and try a different AND gate). All the inputs of the AND gate except the one under test are then set to their asserted values. Hence, the logic value at the input pin under test determines the logic value at the AND gate output.

The logic value must now propagate to the PLA outputs through one of the OR gates. Since the AND gate may feed several OR gates, the algorithm arbitrarily chooses one of these OR gates. The algorithm can backtrack and choose a different OR gate if necessary.

To propagate the fault effect through the OR gate, the algorithm attempts to set the circuit inputs so that the outputs of all the AND gates feeding the OR gate, except the one being used to propagate the fault effect, are set to zero. This attempt is complicated by the fact that the AND gates may have shared inputs, some of which could be inverted.

Faults on the outputs of AND-gates are tested by first setting the AND-gate output to the logic value needed to detect the fault and then propagating the fault effect through an OR gate. Output faults are tested by setting the outputs to the desired values using the known Boolean equation for each PLA output.

■ VECTOR GENERATION FOR STATE-MACHINE BLOCKS

To generate test vectors for a state-machine, the test-vector-generation software assumes the following properties:

1) All inputs to the state-machine block are directly controllable by signals at the chip I/O pins.

2) All outputs from the state-machine block can be observed during the test at the chip's primary outputs.

3) The state machine has a RESET State.

In addition, the D inputs of all latches of the state machine should be directly observable for optimum fault coverage. Given that these properties are respected, the test sequence exercises all state transitions at least once while testing all outputs in each state.

The test vector generator considers any state machine as two basic parts, a combinational part and a sequential part. For the combinational circuitry, it makes a test based on the PLA generated by the State-Machine Compiler. Then it exercises each transition of the sequential circuitry. Finally, it merges these two sets of vectors into a global test by merging the combinational patterns into the sequential patterns, taking advantage of "don't cares."

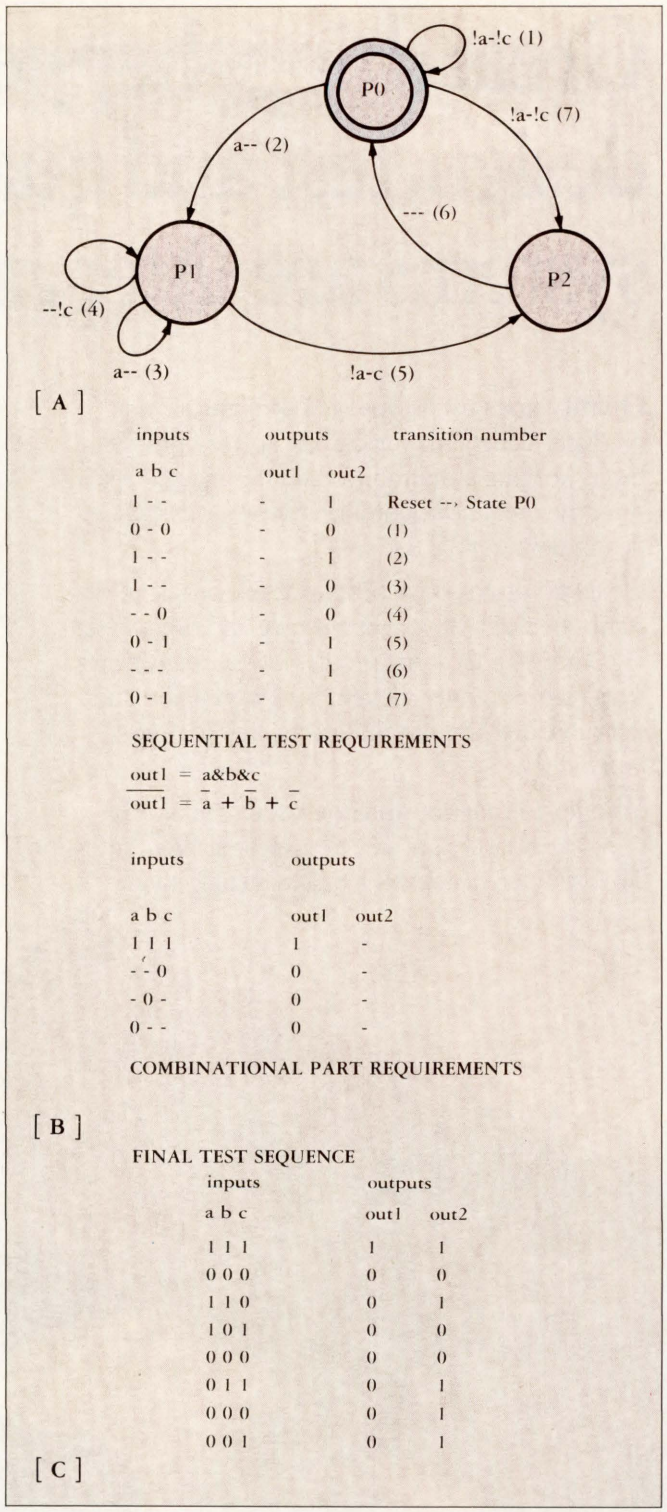

Figure 3. To test the state machine that implements this state diagram (a), the software produces separate tests for its sequential (top) and combinational (bottom) portions (b) and then folds the latter into the former for a complete test (c).

For the sequential part, the Reset command is used to put the state machine into state "P0" for initialization and for getting to any states unreachable from the current state. For the combinational part, the program creates a functional test based on the PLA description generated by the state machine compiler. For each output, each product term that makes the function true becomes a vector to drive the output to "1." The function is then negated and each resulting product term also becomes a vector to drive the PLA output to "0."

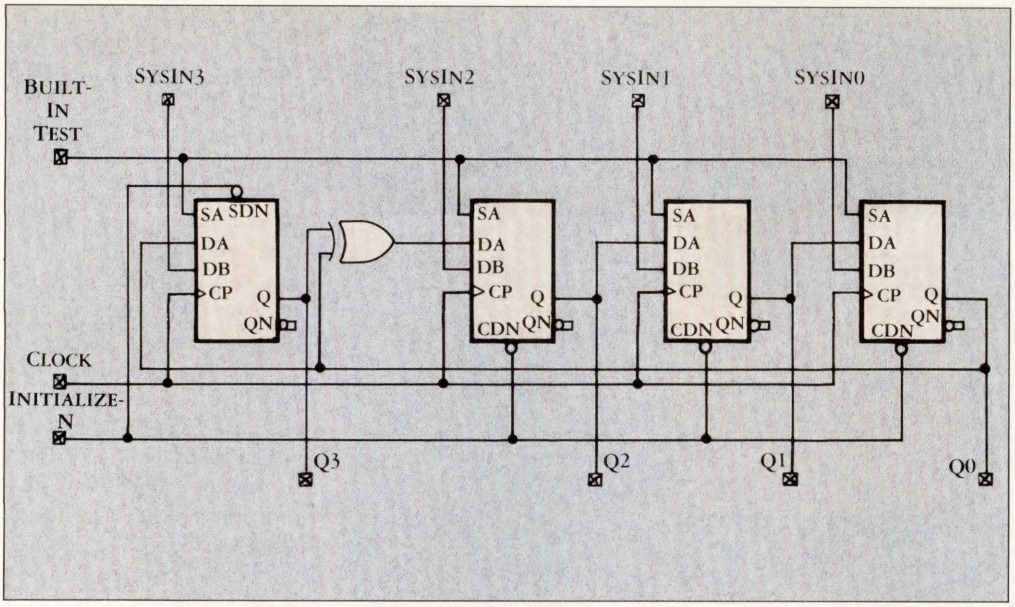

Figure 4. An XOR added to a shift register creates a linear feedback shift register (LFSR), a built-in test circuit that can generate patterns for a block or compile block responses into a "signature" for functional verification.

Finally, these input vector requirements are merged with the sequential ones. This approach does not necessarily insure 100 percent structural testing of the implementation, but it typically provides 95 percent or better fault coverage and requires only a few minutes of CPU time.

The state machine in Figure 3a provides an example for viewing the vector-generation results. Figure 3b shows the sequential (top) and combinational (bottom) vectors for the state machine. Note that most of the vectors contain some unspecified (don't-care) input values. As shown in Figure 3c, the sequential and combinational vectors can be merged to form the final set of vectors. Most of the don't cares are resolved as combinational vectors that are folded into sequential vectors. Those vectors that are not are arbitrarily set to 1 or to 0.

■ Vector Generation for Datapath Blocks

Datapath-type blocks are generated by the datapath compiler. They are typically made of regular logic blocks performing arithmetic or logic transformations on n-bits at a time. These functional blocks are specified in detail by the ASIC designers, and their functionality and expected behavior is well-known. Because of the regularity and simplicity of datapath structures, manually generated test vectors can achieve a high fault coverage, especially if a mechanism is in place to insure controllability and observability of the datapath block's I/O signals at test time. These I/O signals include not only the data buses but also the control logic, whose states will be set with vectors from the state-machine vector generator.

The designer's task to conceive of a test program for a datapath block is simplified by the existence of test modes for each element of the datapath library. In addition, a vector editor simplifies the actual creation of the vectors.

■ Megacells and User-Defined Blocks

In addition to cell compilers, VLSI has a megacell library of full-custom high-level functional blocks. Pre-defined test programs for these blocks are developed by the designer who has implemented the megacell. These programs provide high-coverage vectors -- averaging 90 percent fault coverage -- that are verified through fault simulation.

Users of VLSI's tools may define their own logic blocks and use VLSI's functional block isolation methodology to efficiently test these blocks. This methodology is an effective way of achieving high total fault coverage for a chip design composed of separate functional units that have individual high-fault-coverage vector sets.

Sometimes user-defined functional blocks are actually previous ASIC or board designs that are being consolidated to reduce the cost of the system. In these cases, using existing vectors to test an ASIC functional block can often save the designer weeks of test vector development time.

Automatic test vector generation (ATVG) is often very useful for creating vectors to test other logic blocks, even if the methodology becomes impractical at the chip level. Since VLSI's functional block methodology allows the designers to isolate any area of logic as a functional block, ATVG can be used wherever it proves to be an efficient method of exercising random logic.

■ Built-In Test

Built-in test (BIT) is growing as a partial or complete solution to test generation for ASICs. To incorporate built-in test in a design, it is necessary to have circuitry that can generate test patterns, and other circuitry to compact the circuit's response to these patterns. VLSI Technology has developed a built-in test logic compiler which generates Linear Feedback Shift Registers (LFSRs) for test pattern generation and response compaction (Archambeau, 1988).

An LFSR is a shift register in which certain inputs receive feedback from the last output through an exclusive-OR gate (Figure 4). An "n-stage maximum-length LFSR" is capable of cycling through $2^n - 1$ nonzero distinct states (McCluskey, 1986 and Zierler, 1955).

If the LFSR is used as a pattern generator, there is always at least one specific configuration of XOR gates that can implement a maximum length pseudo-random input generator. This configuration produces an n-bit LFSR whose outputs will cycle through $2^n - 1$ values before repeating. Each individual bit of this LFSR will produce an essentially random input.

The user interface for the LFSR compiler is a window in VLSI Technology's set of interactive tools. The designer edits a template to create a parameter cell describing the BIT circuit to be created. The BIT compiler can generate LFSR modules for both pseudo-random pattern generators (for input to BIT) and signature analyzers (for generating a BIT output). The LFSR module produces a behavioral model and a netlist for simulation as well as an icon for schematic-capture.

A BIT circuit can only be activated by the test equipment during a test mode specified by the designer. This test mode can activate all built-in test logic in parallel, or it can provide decode logic to select subsets of the built-in test circuits. After the test mode has been activated, the built-in test circuit must be initialized to a known state (through a reset signal, for example). Then, the built-in test is executed and the results are examined to determine if its corresponding functional block is correctly functioning.

The test-pattern-generation software discussed earlier re-

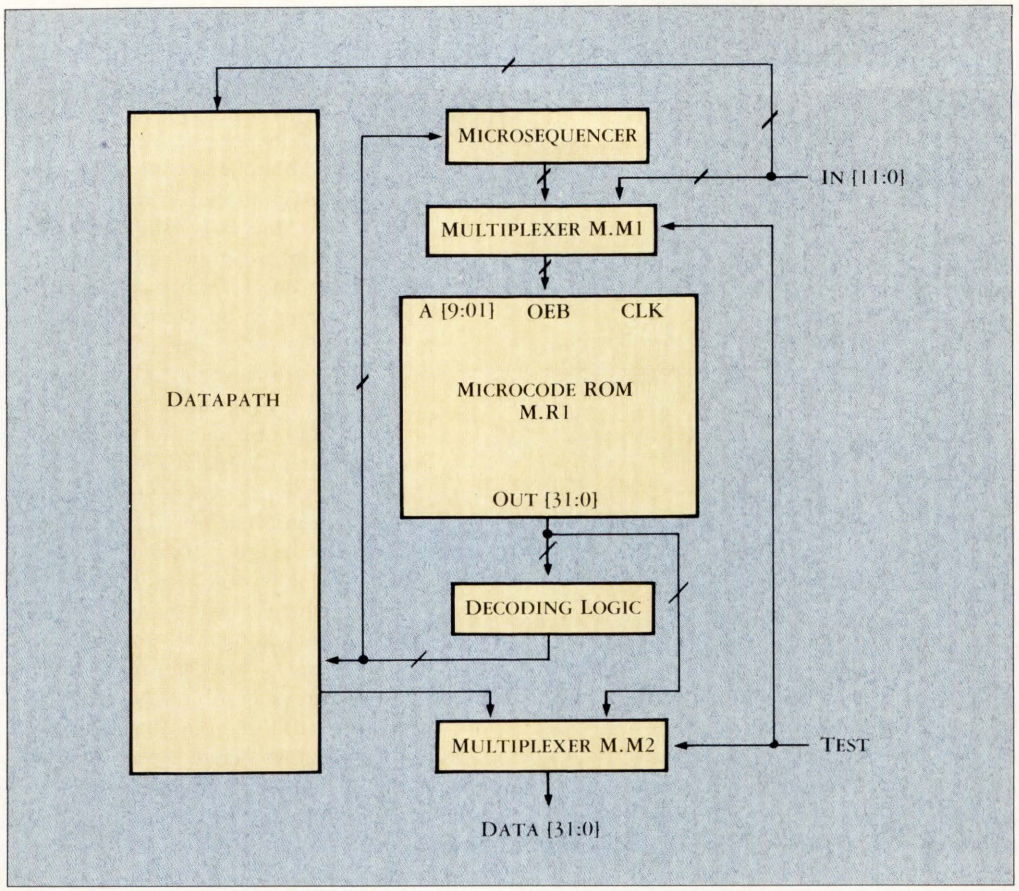

Figure 5. To isolate a ROM for testing, multiplexer M.M1 reroutes the primary inputs IN[11:0] to the ROM inputs and multiplexer M.M2 connects the ROM output signals to primary outputs DATA[31:1].

quires that compiled module I/O signals be accessible from primary I/O signals (chip pins) during a test. In some cases, compiled module I/O signals are connected to chip pins either directly or through simple combinational logic. However, in many cases module I/O signals are embedded deep within a circuit, or no chip pins are available to connect to module I/Os.

■ Functional Block Isolation

In such cases, some form of test logic must be added to the circuit to isolate and access modules. For this purpose, the three most important techniques are multiplexer-based schemes, scan paths and built-in test.

Some of the more important factors that must be considered in choosing a functional-block isolation scheme are area overhead, I/O-pin overhead, delays on critical paths and test time. In addition, provision must be made for hardware to generate a signal to activate the isolation logic. Typically, the values on one or more chip inputs are decoded to provide such a signal.

The simplest form of functional block isolation makes use of multiplexers to substitute primary inputs and outputs for the signals that normally connect to the functional block.

For example, multiplexers M.M1 and M.M2 in Figure 5 connect the primary inputs and outputs, respectively, to the microcode ROM. The biggest advantage of this scheme is that all functional block I/Os can be accessed in parallel, so a whole test vector can be applied in each tester cycle.

The area overhead for multiplexer-based isolation is the sum of the area required for the multiplexers and the area required for interconnect. The area overhead for interconnect varies with the physical location of the functional blocks on the chip and the number of their I/O signals. For blocks with a large number of I/Os, which are embedded deep within a chip, the overhead can be significant. The number of module I/O's may exceed the number of chip I/O pins. Even when there are enough pins, delays due to the multiplexers are introduced on a large number of system paths.

Another commonly used technique to isolate functional blocks makes use of scan paths. A scan path is simply a shift register whose input (the scan-in pin), and output (the scan-out pin), are connected to primary I/O pins. The I/O signals of an embedded module can be made accessible by connecting them to the scan path. Four steps are required to apply a vector to the isolated module:

1) The functional block is put into test mode.

2) The test inputs are shifted into the scan path via the scan-in pin.

3) The test inputs are applied to the block under test and its response is captured in the scan path flip-flops connected to its outputs.

4) The test response in the scan path is shifted out via the scan-out pin.

The area overhead for a scan path can be quite significant if the flip-flops added are used solely for test. However, the inputs and outputs of the block being isolated may already have flip-flops connected to them.

In this case, the area overhead per bit of the scan path equals the difference in size between a regular and scan flip flop (approximately the cost of a multiplexer). The interconnect overhead for a scan path is minimal. Moreover, since only two chip I/Os are used, this technique disturbs fewer unrelated signal paths than a multiplexer scheme. The performance penalty for using scan flip-flops instead of regular flip-flops can be minimized through appropriate design techniques.

The main drawback of using a scan path for isolation is the number of cycles needed to shift a test pattern in and to shift the response of the circuit out. Because input patterns shift in at the same time as the output responses for the previous vector shift out, the number of tester cycles required for a vector equals the greater of the number of input pins and the number of output pins. If the number of functional block I/Os is large, then the time taken to apply a large set of vectors can become unacceptably long.

With the advent of testers with hardware dedicated to scan-path testing, this technique is becoming more practical. However, these testers are not prevalent in the ASIC industry.

The third technique that can be used to isolate a functional block is built-in test. Built-in test can either be used alone or in conjunction with multiplexers and scan paths.

For example, during a ROM test the addresses could be applied from chip I/Os, and the response of the circuit could be captured in a signature regis-

```
$INSTANCE M.R1
$CELLNAME CROM01
$PARAMETERS [PCL]MICROROM
$PINMAP
A[9:0]                    IN[9:0]
OEB                       IN[10]
CLK                       IN[11]
OUT[31:0]                 DATA[31:0]
$TESTMODE
HIGH TEST
$VECTORS [VIF]MICROROM
```

Figure 6. For the ROM isolated in Figure 5, this Test Block Map identifies to the test compilation software the inputs, outputs and test control signal for the ROM operating in test mode.

ter. As with scan paths, the overhead for built-in test circuitry can be reduced if existing circuitry can be used. For the previous example, if the ROM output flip-flops can form part of the LFSR signature analyzer, then less additional hardware will be required.

■ Merging Vector Blocks

Integrating multiple vector sets with unique timing and tester resource allocation requirements into a single test program is a difficult task in itself. A complete test program includes not only functional vector sets but also a complete set of parametric and critical timing checks.

Most testers have proprietary test languages. Modifying the test program to allocate and reallocate test resources can be a difficult problem for the test engineer. VLSI's vector-conversion software automatically combines the vectors for each functional block into a chip test program.

Once a designer has generated all the test vectors, he will have a vector file for each functional block in his design as well as a set of "top level" test vectors that test overall system operation. The timing requirements and data formats for each test set may be unrelated. Each vector set may require an initialization sequence to put the circuit into the required test mode to access each functional block.

All the information necessary to combine all vectors sets into a single test program must be provided by the designer in a Test Block Map (TBM) file. VLSI's test program generation software reads the TBM file along with the complete set of vectors provided for a design and integrates them all into a single test program. Tester resources are automatically reallocated for each set of vectors, and initialization vectors are executed before testing each functional block.

Because the I/O names of the functional block vectors are usually different from the primary I/O names, the TBM file describes the correspondence between the module I/Os and primary I/Os. The TBM file also contains information about how the circuit can be put into test mode. The timing information for each vector set is registered in the VIF vector file.

Figure 6, for example, shows the TBM file for the isolated ROM in Figure 5. The ROM I/Os are "mapped" to the IN and DATA buses. The instance, cellname and parameters of each block must be specified because similiar or identical instances of a cell often exist in a single circuit. The test mode for this ROM involves asserting the chip input TEST.

The TBM file can also be used to describe how a functional block is to be tested using scan or built-in test. To support scan, the TBM file must contain information that identifies scan-in and scan-out and defines the order in which functional block I/Os are linked in the scan chain. The order of functional block I/Os is needed by the vector conversion software to convert the parallel vectors for the functional block into the serial scan vectors.

If a functional block is to be tested using BIT, then the TBM file must specify which chip pin provides a clock signal to the BIT circuitry as well as the number of cycles required for the test. Moreover, the TBM file must also contain the expected result of the test and the chip I/Os on which this expected result can be observed. ■

Acknowledgments

The authors wish to thank Ken Van Egmond, Robert Shur, and Jim Althoff (San Jose, CA) as well as Christian Jay and Chris Kingsley (Sophia-Antipolis, France) for their continuing support throughout the development of this project.

References

ABRAHAM, J.A. 1986. "Fault Modeling in VLSI," in MDULVLSI TestingMDNM, edited by T.W. Williams, North Holland, 1986, pp. 22.

ARCHAMBEAU, E. and VAN EGMOND, K. 1988. "Built-in Test Compiler in an ASIC Environment," to appear in Proceedings of the International Test Conference, Philadelphia, PA, September, 1988.

BREITENWISCHER, T.G. 1987. "Logic Verification and Production Testing of Non-structured Embedded VLSI Blocks," 1987 IEEE Custom Integrated Circuits Conference, pp.62-65.

BREUER, M.A. and FRIEDMAN, A.D. 1976, "Diagnosis and Reliable Design of Digital Systems," Computer Science Press.

MCCLUSKEY, E.J. 1986. "Logic Design Principles: With Emphasis on Testable Semicustom Circuits," Prentice Hall.

SAMAD, M.A. and BUTZERIN, T. 1988 ."A Methodology for the Test of Embedded Compiled Cells," 1988 IEEE Custom Integrated Circuit Conference, Session 16.

WILLIAMS, T.W. 1983. "Design for Testability -- A Survey," Proceedings of the IEEE, Vol. 71, No. 1, January 1983, pp. 98-112.

ZIERLER, N. 1955 "Several Binary Sequence Generators," Technical Rep. 95, Lincoln Laboratories, Massachusettes Institute of Technology, Lexington, MA, September 1955.

About the Authors

ERIC ARCHAMBEAU *is the Manager of Software Development at VLSI Technology, where he has also served as manager of Logic Entry, Workstation, simulation and Test software groups. He has been involved in logic simulation, fault simulation and testability analysis since 1982, first with Matra Design Systems and then with Intercept Microelectronics. He received his Masters, Engineers, and PhD degrees, all in Electrical Engineering, from the University of California at Berkeley, Stanford University and the Institut National Polytechnique de Grenoble in France (respectively).*

THERESA BUTZERIN *is a Staff Member of the Advanced Products Group at VLSI Technology. With VLSI since 1983, she has also served as a Project Leader for test automation software and as a Senior Design Engineer at VLSI's San Jose design center, where she worked on numerous systems-level VLSI designs. She has a BSEE from Stanford University.*

ARIF SAMAD *is a Software Engineer at VLSI involved in the development of the next generation of test-related CAD tools. Prior to joining VLSI in July of 1987, he was a research assistant at Purdue University, where he received his MSEE in 1986. He received his BSEE from Lafayette College in 1984.*

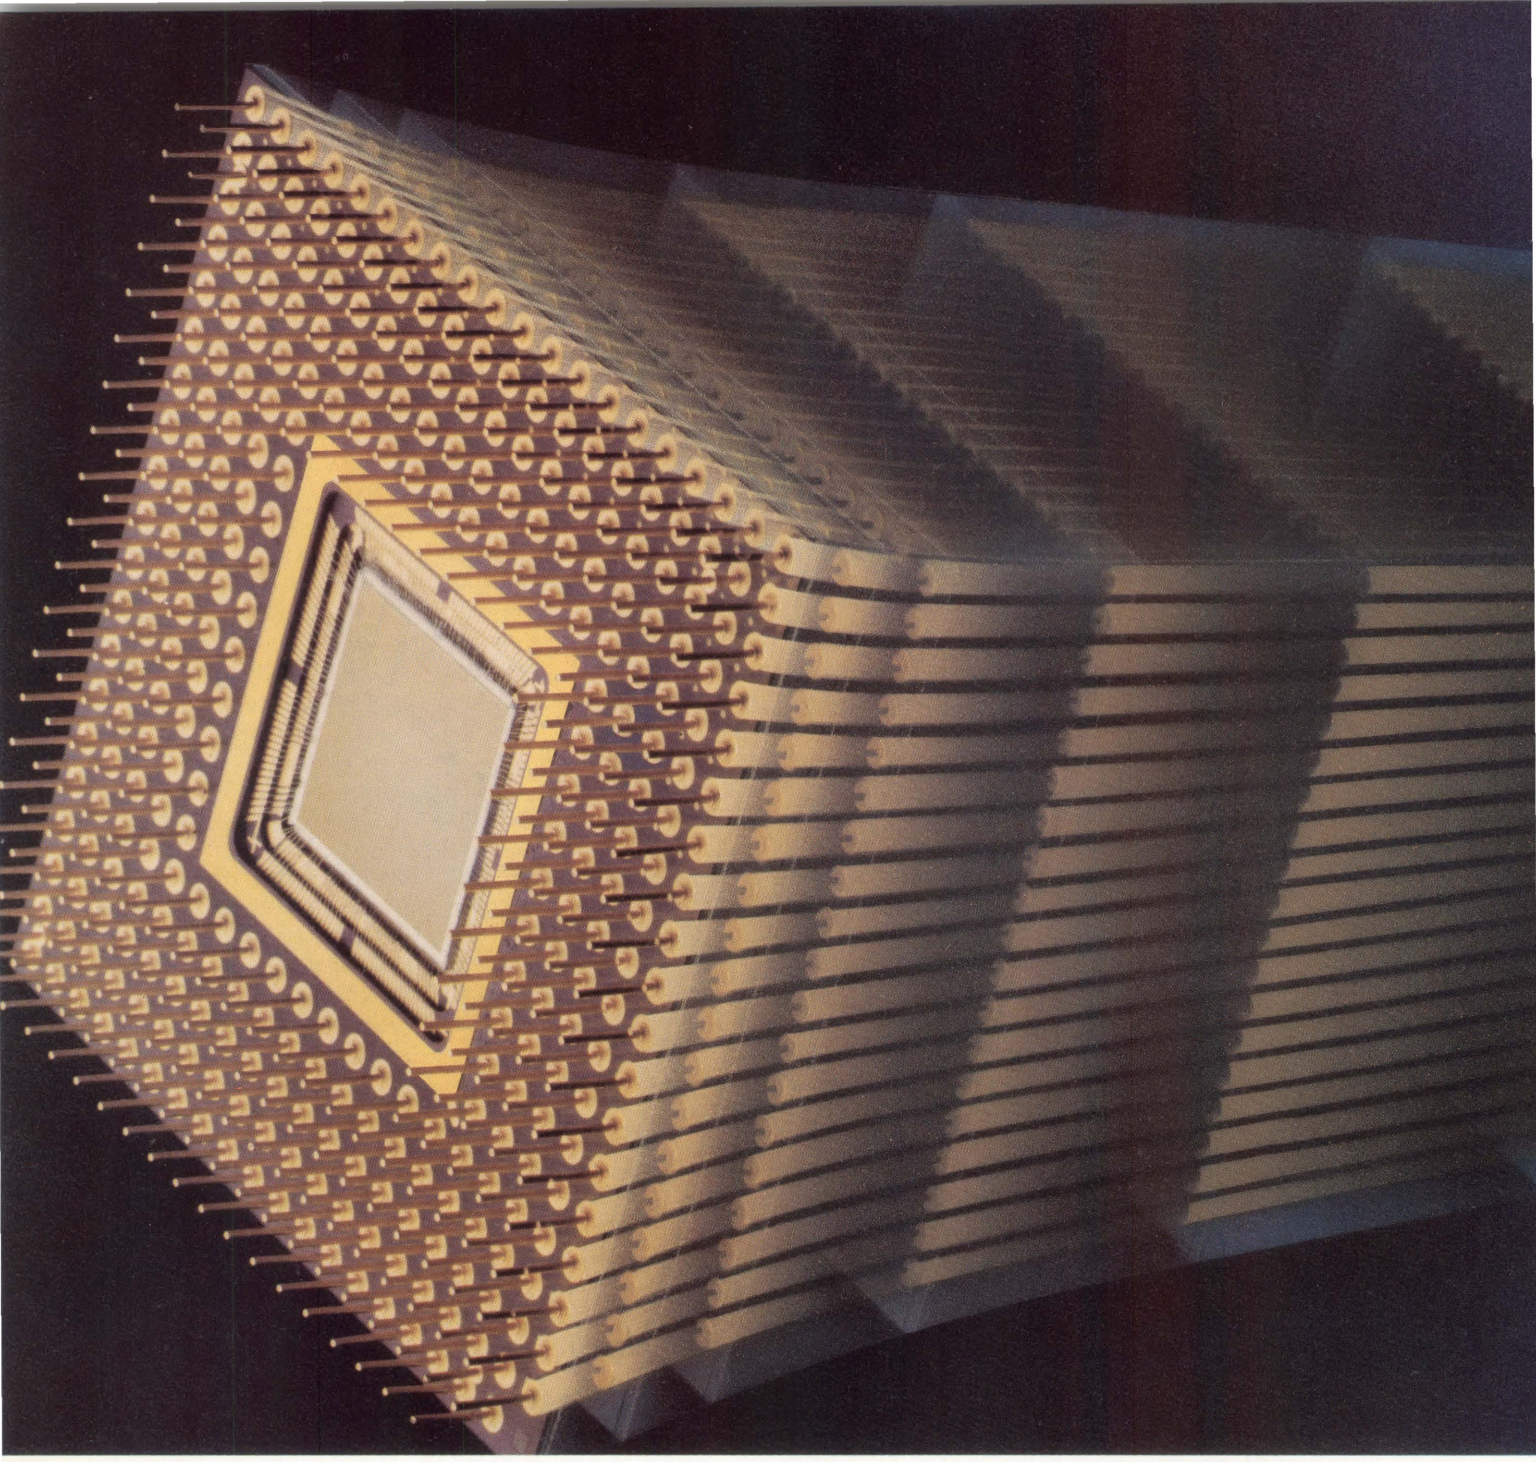

129,000 gates and 400 picosecond to none.

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

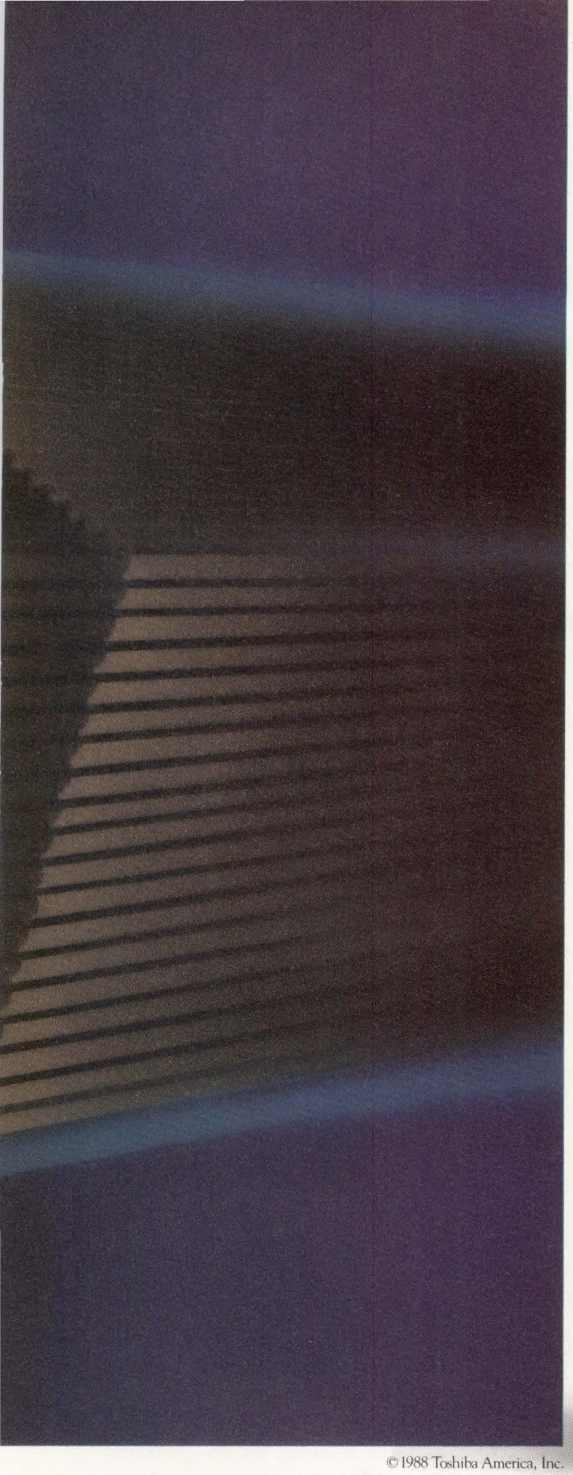

Now Toshiba introduces a new level of performance in gate arrays. It's the TC120G Series, and it features typical gate delays of only 400 picoseconds. That's 35% less than our own industry-leading TC110G Series. The secret is our 1.0 micron CMOS process, and, of course, our proven Sea of Gates non-channelled architecture.

The new TC120G Series is upwardly compatible from the 110G, and supported by compatible CAD tools. In fact, our VLCAD-II System has been designed to permit the easy upgrade conversion of 110G systems to 120G technology.

The TC120G Series can provide the size and speed for consolidating high speed circuits into a single package. It is especially good for high performance applications like mainframe CPU's, minicomputers and telecommunications switching systems.

TOSHIBA. THE POWER IN GATE ARRAYS		
SERIES	TC110G	TC120G
GATES	3,498 to 129,042	37,932 to 129,042
GATE LENGTH	1.5μm	1μm
GATE SPEED	0.6 ns	0.4 ns
PART NUMBERS	14	5
AVAILABILITY	NOW	NOW
All Si-Gate CMOS double layer metal.		

If you don't need 129,000 gates, don't worry. The TC120G Series is available in 5 master array sizes ranging from 37,932 to 129,042 raw gates. So you can design in just the right size for your ASIC application. And now to help you even more, we've opened two new ASIC design centers. That brings the total to five in the United States. 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-88-004

CIRCLE NUMBER 13

MIXED ANALOG/DIGITAL ASICS

standard digital device test techniques are no longer sufficient

KEN DUBROWSKI AND THOMAS WONG, NATIONAL SEMICONDUCTOR CORP., SANTA CLARA, CALIF.

■ DESIGNER MUST SEPARATE FUNCTIONS FOR EFFICIENT TESTING ■

AS CMOS ASIC technology pushes toward higher levels of integration, systems designers have found the need to put analog functions and digital circuitry on the same ASICs. Although the proportion of analog circuitry typically represents less than 20 percent of the die area, mixing analog circuits with digital circuits raises several technical challenges that run the spectrum of design methodology. In addition, design issues such as placement and routing, mixed-mode simulation, test methodology, and equipment requirements become more complex with analog/digital designs.

National Semiconductor's approach to mixed analog/digital ASIC design is based on standard-cell methodology. Using a 2-micron, double-layer-metal CMOS process, National has designed and characterized a library of analog blocks that includes operational amplifiers, comparators, analog switches, multiplexers, voltage references and oscillators.

To simplify the conversion of a printed circuit board to a cell-based IC, the op-amp cells have characterstics that are similar to standard off-the-shelf analog ICs that the designers are very familiar with. For example, the characteristics of the op-amp cell are very similar to those of a standard LM324—when they are both operating from a supply voltage of 5V. In the same way, the comparator cells mimic the behavior of an LM339 IC running at a supply voltage of 5V. Data sheets for the analog cells even look like the standard devices' documentation, with guaranteed minimum and maximum specifications for critical parameters and typical performance graphs for parameter variations as a function of power supply voltage, frequency, and temperature.

Realizing that most system designers prefer to verify their analog designs on a breadboard prior to making the prototype ASICs, National offers an evaluation kit of parts for these purposes. This design approach requires less integrated circuit design experience than other semicustom methods and enables the user to employ a block-level design approach.

■ METHODOLOGY AND EQUIPMENT

Testing mixed analog/digital ASICs is difficult because of the different needs of analog and digital testing. National separates the analog and digital circuits and tests them independently, using the most appropriate tests for each type. This methodology overcomes the present lack of design and test tools for the creation and testing of mixed analog/digital ICs. As new automation tools become available, however, this

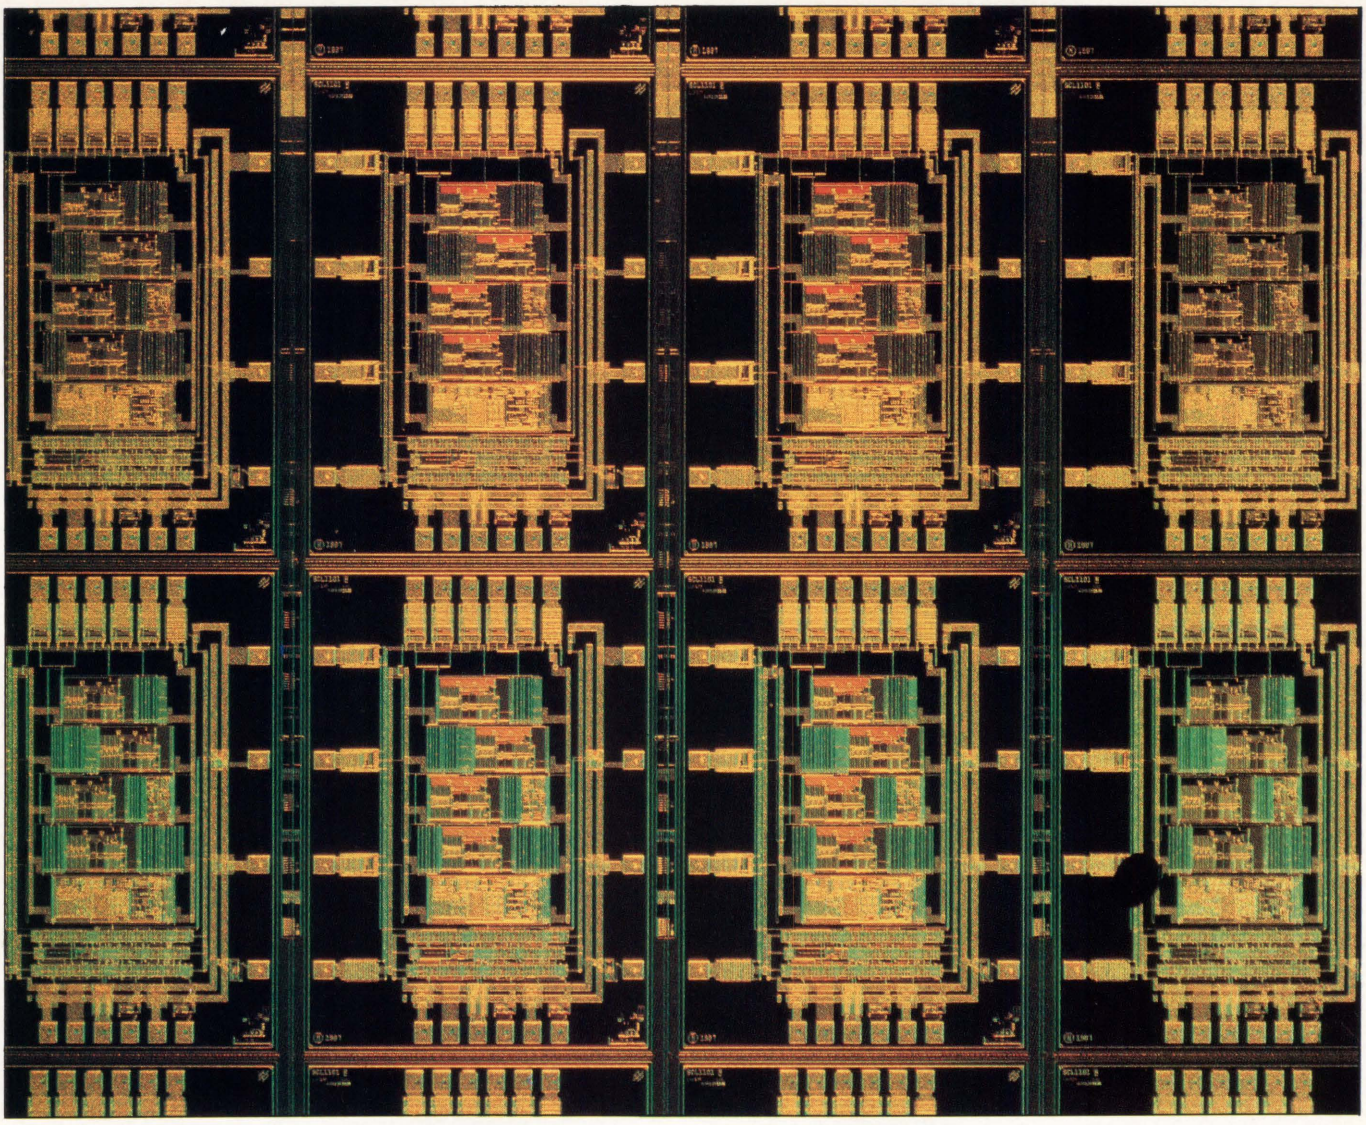

methodology is expected to continue to be effective.

When analog functions are added to a primarily digital ASIC, two fundamental test issues arise: the level of testing and the constraints of the automatic test equipment (ATE) used. The first level of testing is component testing, in which all of the operating parameters of the analog components block are directly verified. This method is used by semiconductor manufacturers to perform tests on standard linear components. For example, with an operational amplifier, parameters such as offset voltage, bias current, input common-mode range and output swing can be determined by using a standard operational amplifier test loop. For mixed analog/digital ASICs, however, a system-level test procedure is usually better because it minimizes the number of test points required and tests the intended application functions instead of checking the parameters of the individual components.

For example, the instrumentation amplifier in Figure 1, a circuit commonly used with transducers such as strain gauges and piezoelectric crystals, contains three operational amplifiers. A system-level specification of the instrumentation amplifier would include closed-loop gain, gain accuracy, dc offset voltage (common mode) and the 3dB bandwidth. Testing the device as a system requires only three test nodes—two inputs and the output. Testing the device at the component level, where the characteristics of each amplifier and resistor are measured, would require more tests and access to nine nodes.

This well-proven method of verifying system performance has been used extensively by hybrid integrated circuit designers. Because of the benefits of using this approach in testing mixed analog and digital ASICs, the system designer must become more intimately involved in testing issues—based on the system specification—than he does in a digital ASIC. Since the techniques for digital ASIC testing are well known and documented, there is no need for further discussion.

ADDING ANALOG FORCES A CHANGE IN TEST PHILOSOPHY

A system used for the testing of mixed analog/digital ASICs requires four basic functions:

1) generating precision analog waveforms;
2) delivery of complex digital patterns at high clock frequencies;
3) digitizing the analog output;
4) capturing the digital output signals.

The selection of test equipment for mixed analog/digital ASICs is complicated by the divergent needs of digital and analog circuits. Digital ASICs are character-

ized by large numbers of I/O signals, a fact that suits them for a high-pin-count digital ATE system such as offered by Sentry and Advantest. Linear circuits, on the other hand, require a level of flexibility and precision—in applying and sampling waveforms—that is generally associated with a linear ATE system such as those available from LTX or Teradyne, for example.

Because the number of digital I/O s and function blocks is much greater than the analog I/Os and blocks, a digital tester modified with the addition of special instrumentation can often provide the best compromise. For example, with a digital tester, a simple sinewave could be generated with an oscillator circuit, or a precision waveform could be mathematically defined and then synthesized through the use of a precision digital-to-analog converter driven by digital signals from the tester. Similarly, the analog outputs could be analyzed either with standard readily available analog instrumentation or digitized for analysis using the capabilities of the digital ATE.

■ Test Methodology

The test policy at National tests 100 percent of the parts for all ac and dc parameters that have specified limits (minimum or maximum). All parts are tested for proper function over the specified operating temperature range, although the testing may occur at only one temperature if performance can be accurately extrapolated to the limits of the operating temperature range.

Verifying the functionality of an ASIC is not the only purpose of a test program. It must also weed out process faults which could cause malfunctions at any time over the lifetime of the device. In addition, the test should also determine that the device will always function properly without regard to the state of the device upon power up or to the state of the device after a power "brown out."

By first testing the functionality of the part, National establishes confidence that the connectivity is correct. The digital signal wires are tested for both high and low states. The analog wires are also tested with either high and low voltage states or sinking and sourcing current states.

This level of testing verifies the connectivity within the circuit is correct to the extent that it provides the required functionality. A more reliable product, however, requires a greater level of fault coverage. This is accomplished by partitioning the chip into smaller parts, and then testing each part exhaustively.

The first partition divides the chip into analog and digital sections through the use of multiplexers (Figure 2). Through these MUXes, the internal nodes at the analog and digital interface are made accessible to the ATE at the package pins. The parasitic effects introduced by this additional circuitry is insignificant at the operating frequencies of the analog and digital circuitry.

The designer uses a logic simulator to create test patterns for the purely digital sections, and verifies the fault coverage of those patterns. Because high fault coverage entails many test vectors, the digital portions should be tested at high frequencies to minimize overall test time. Because the analog test MUXes bypass the comparatively slow analog circuitry, the digital sections can be tested quickly and thoroughly.

The MUXes also isolate the analog circuitry for thorough testing. They provide access to the internal interface nodes between the analog and digital subcircuits. For testing, the access to all analog inputs provides controllability, while access at the package pins to all analog outputs provides observability. The analog block can now be tested independently of the digital circuitry, which avoids lengthy initialization patterns that may otherwise be needed to place the digital circuitry into a known state.

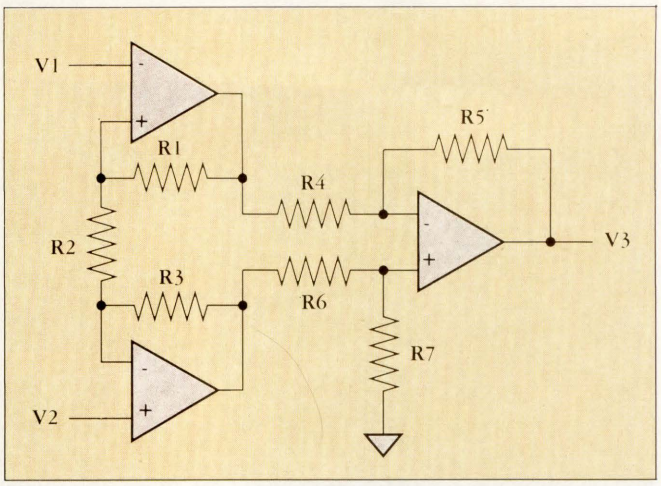

Figure 1. Testing the function of this instrumentation amplifier requires three test points. Testing each component individually would require nine points.

To characterize the operation of the analog circuitry, the designer may want to run the circuitry through several test modes. For this purpose it is inconvenient to use a multiple-channel pattern generator to set the digital logic into a known and desired state. A preferred method would simply tie the inputs from digital circuitry to a given test state.

To minimize the number of I/O pins for this purpose, an alternative strategy uses a counter with a clock input for advancing through test states. As shown in Figure 3, when the test control input is at ground, the counter is reset and the chip operates normally. When the test control input rises above 3.5V, test mode 1 is activated. Test mode 2 is activated by lowering the test control input below 3.5V (while remaining above 1.0V) and raising it above 3.5V again.

This procedure is repeated to advance through the rest of the test modes.

■ Partioning for Testability

The designer may not find it sufficient to have access to only the inputs and outputs of the entire analog subcircuit. To improve fault coverage it may be necessary to partition the analog network further. The designer needs to decide then on how far to partition the design.

The ASIC circuit can be partitioned in three ways: leaving the analog network whole; dividing it into functional blocks; or separating each analog component. If a chip has many analog components—particularly for a chip with an analog signal path passing serially through many devices—the partitioning method has great impact on the test development time, analog "fault coverage," the test cost, and the number of package pins used during the test.

The first partitioning approach tests the analog circuit by accessing only the inputs and outputs of the overall analog subnetworks. Although this approach requires only a single test sequence, the sequence must be unique for each design. Also, the amount of engineering effort required to provide high fault coverage can impact prototype delivery schedules. The risk of error is greater because of the amount of custom test development. In fact, with this approach the designer is limited by the scarcity of people capable of doing a thorough job of writing this test program.

A second approach partitions the analog circuit completely to access each analog component individually. With this access, the designer can test each analog macro to its specified data sheet limits using standard test routines. The

difficulty with this approach lies in the large amount of extra circuitry required for multiplexers to connect the many test points to package pins. The number of test pins required also becomes prohibitive.

National uses a functional-block-level approach to testing the analog circuitry. Most analog designs can be partitioned into standard circuit configurations, such as the summing amplifier, the inverting amplifier, the integrator and the instrumentation amplifier. Each configuration has its own transfer function. Taking into account actual passive component values, standard test routines can be used that completely test the critical factors associated with a particular circuit configuration.

These standard tests are tabulated into National's test tables, such as the one in Figure 4 for a non-inverting amplifier. The tables identify the specific parameters that are measured in the test routine, and the circuit nodes that must be accessible for testing those parameters. If a specified test node is not available to the ATE equipment, then its associated test parameter can't be tested. If the designers are constrained by a small number of test pins, then they can determine the tests which will be passed over as a result of lack of access to particular circuit nodes.

■ ATE Requirements

This "system-level" approach allows standard test routines to be used for testing each standard circuit configuration. It helps automate the development of a comprehensive test program for the entire analog network. Even though access to internal circuit nodes is required, the use of test circuitry, multiplexers and test logic can minimize the number of test pins required. In short, the approach provides for thorough testing, few test pins and some degree of automated test program generation.

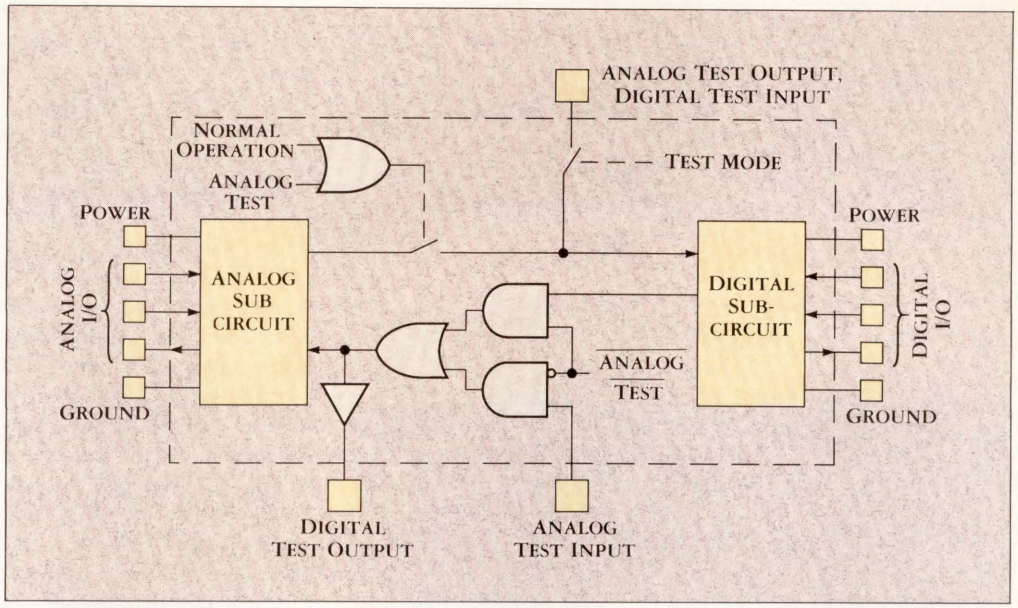

Figure 2. The designer should divide the chip into analog and digital sections with multiplexers to make the internal nodes at the analog and digital interface accessible to the ATE at the package pins.

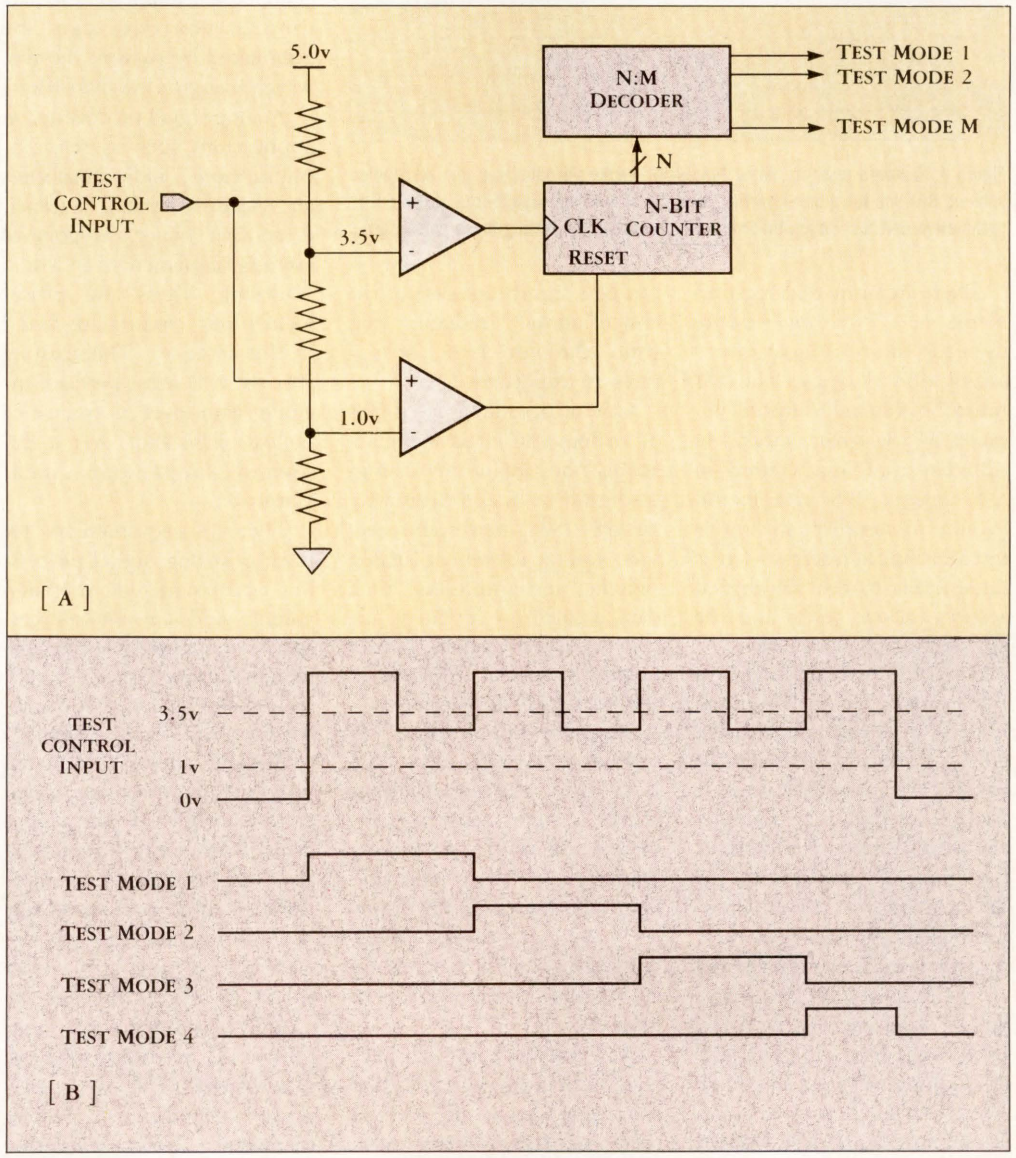

Figure 3. To minimize the number of I/O pins for setting test modes, one test strategy uses a counter with a single clock input for advancing through test states. This circuit advances through test states by toggling the test control input around 3.5V.

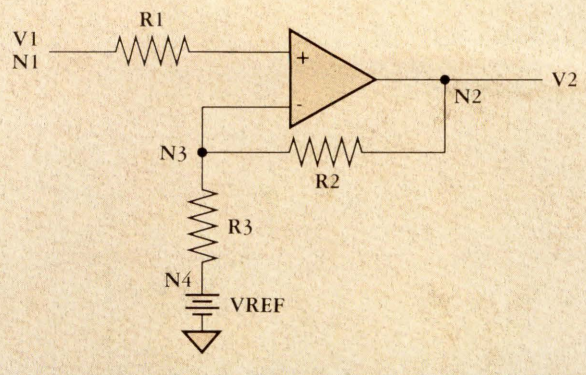

NODES TO BE ACCESSED:
N1, N2, AND N4

$V2 = (R2/R3 + 1)V1 - (R2/R3)VREF$

PARAMETERS TO BE TESTED:	TEST CONDITION
1. CLOSED-LOOP VOLTAGE GAIN	
2. V_{OS} @ V_{REF}	
3. V_{OH}	R_L = 10K OHM
4.A. V_{OL}	R_L = 10K OHM TO GRD
4.B. V_{OL}	R_L = 10K OHM TO VDD
5. SLEW RATE	
6. V_{REF}	

Figure 4. Standard tests for linear functional blocks are tabulated into test tables such as this one for a non-inverting amplifier. The tables identify the specific test parameters and test points that must be accessible for testing those parameters.

There are numerous requirements for ATE systems that are used for the testing of mixed analog/digital chips. The ATE must be general purpose because of the great variety of different functions crafted in ASIC designs. The ATE should be able to complete the entire test in a single insertion—it is undesirable to test the digital section on one tester and the analog section on another.

The digital requirements of the ATE include pin counts as high as 256 pins, high-speed clocking, the ability to handle long test patterns (about 20,000 vectors, for example) and high-speed comparators for sampling the outputs of the device under test (DUT). Chips for the data processing market, for example, need to be tested at their operating frequency, which often exceeds 20 MHz. Even if the operating frequency is lower, the parts should be tested at a rate as high as possible because high levels of fault coverage result in lengthy test programs. Because test time can be a significant part of the cost of an ASIC, reducing that time through high testing rates is important.

One of the major categories of analog ASIC circuits is data acquisition, including analog-to-digital (A/D) and digital-to-analog (D/A) converters, sensors, voltage references, analog switches and amplifiers. Such data-acquisition circuits are tested to verify their transfer functions, such as voltage gain and frequency response.

The ATE requirements for testing data acquisition circuits are precision voltage souces, current sources, voltmeters and ammeters. For example, the present market demands for mixed analog/digital ASICs requires 8-bit resolution. An eight-bit system operating at 5V has a resoltuion of 10 mV.

As a rule of thumb, the ATE system should be able to test at least one magnitude better than the signal being measured, requiring a tester resolution of 1 mV. The data acquisition circuits are generally low frequency, such as 1 hertz for strain guages and up to 10 khz for motor controllers.

Analog test waveforms can be generated on the ATE system by digitizing the waveform as a series of voltage steps. The voltage steps are established by a series of test vectors that drive a D/A converter on the performance board (also called the load board). This situation requires the tester to have a much higher clock rate than the analog signal. For example, to produce a 10 khz sine wave using an 8-bit D/A converter, the tester must generate test vectors at a rate of 2.56 MHz to achieve a total harmonic distortion of less than 0.5 percent. When required, the harmonic content can be attenuated by passing the sine wave through a low-pass filter.

Analog signal processing, in applications such as telecommunications, audio amplifiers and communicatíons systems, is another major category of circuits implemented in ASICs. Multipole filters and phase-locked-loop circuits are found in this category. The analyses that the ATE must perform include gain over a frequency range, phase shift over a frequency range and signal distortion.

The ATE requirements for testing analog signal-processing circuits include precision voltage and current sources, precision voltmeters and ammeters, analog waveform generation and analog-output signal capture and processing.

Analog signal processing tests include measuring signal gain, output slew rate, phase delay from input to output, and signal distortion. In addition, the input waveforms may be continuous, pulsed, or complex. Ideally, the tester should be able to digitize the output waveform (by sampling and recording voltage measurements at evenly spaced time intervals) and process that data to determine whether or not the part is operating correctly without any undesired responses such as short bursts of oscillation.

The ATE system for analog signal processing must be capable of measuring the output signal strength at multiple input frequencies.

For example, a sixth-order low-pass filter with a corner frequency of 1 khz should be tested for a maximum pass-band attentuation of one decibel for input frequencies less than 800 hz. It should also be tested for a minimum stop-band attentuation of 70 dB for input frequenies above 2.5 khz. If the input signal is a one-volt sine wave (two volts peak-to-peak), the tester must be able to accurately measure voltages less than 63 mV peak-to-peak in order to ascertain that the output signal is attentuated by at least 70 dB.

■ ANALOG TESTING ON DIGITAL ATE

The ATE systems used for testing digital ASICs is typically inadequate for testing analog circuitry. With the addition of some hardware, however, many of the shortcomings are overcome. For example, National uses a Sentry ATE system for testing its CMOS digital gate arrays and standard-cell products. This tester can verify data-aquisition analog components when it has some extra hardware on its performance board (Figure 5).

The major limitation of the Sentry is its limited voltage accuracy and its sole precision measurement unit. Testing data-acquisition analog circuits requires the application and measurement of precise voltages. Sentry's Precision Measurement Unit (PMU) is used to apply or measure a voltage or current. Because the PMU is used at the output, the input voltage must be applied from one of the digital power supplies, which have resolutions of ±20 mV plus an accuracy 0.2 percent of the programmed voltage. These specifications can be improved to ±0.3 mV by adding a 16-bit D/A converter to the performance board (Figure 6). The

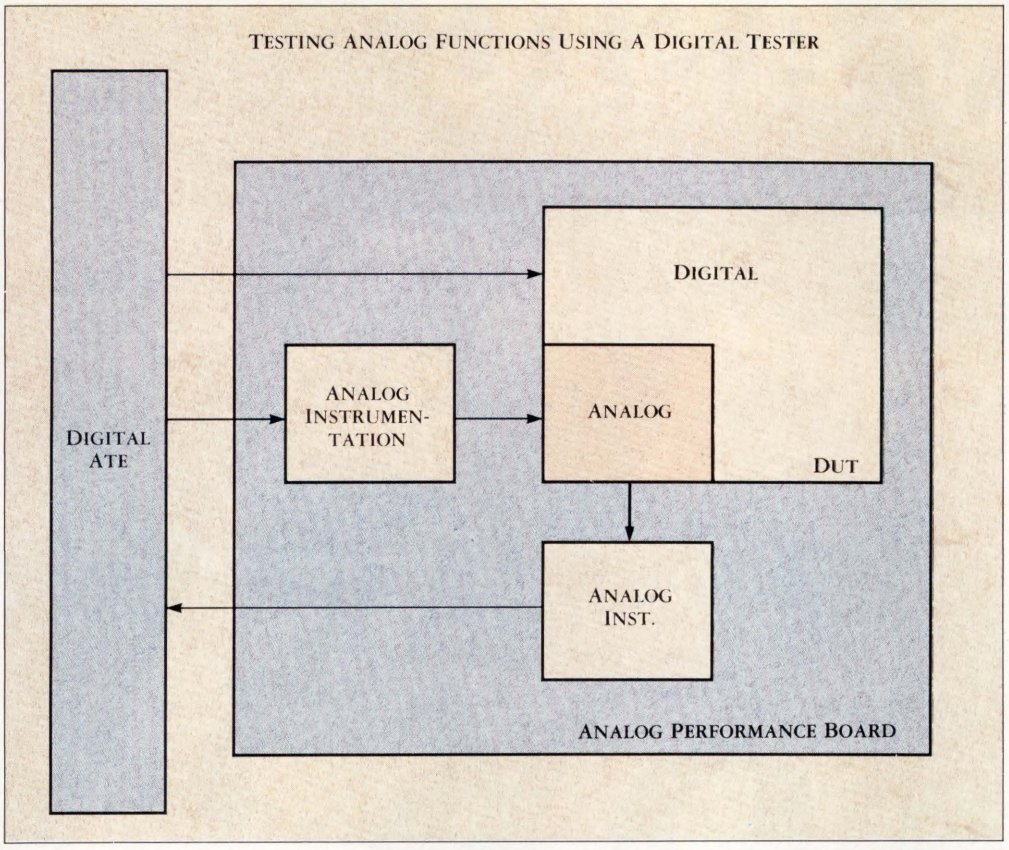

Figure 5. A digital ATE system can verify data-aquisition analog components when it has some additional hardware on its performance board.

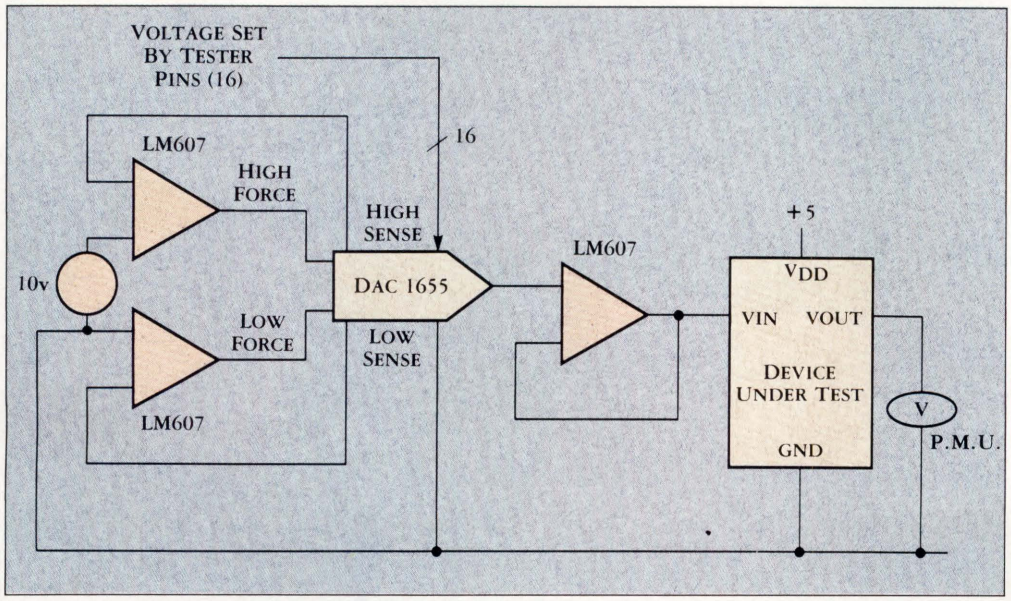

Figure 6. The voltage source accuracy specification can be improved to +-0.3 mV by adding a 16-bit D/A converter to the Sentry tester's performance board.

voltage is set by writing the 16-bit digital word to the D/A coverter, so 16 tester channels are tied up in creating the analog input.

Another method to improve the voltage accuracy is to force higher voltages from the tester and use a voltage divider at the DUT. The voltage-source precision of the tester is a function of the programmed voltage plus a constant accuracy. The accuracy is improved by voltage division.

If the analog input must be 2.5V, for example, driving the input directly from the PMU results in a voltage of 2.5V with ± 0.15 percent accuracy and precision of ± 10 mV. If the PMU voltage is set to 10V and divided by a resistor network, however, the resulting precision is improved to ± 2.5mV. In this case the maximum error voltage is 6.25 mV, much lower than the maximum 13.75 mV when the signal is driven directly. Although this improvement in voltage precision is modest, this method is easier to implement than adding digital-to-analog converters to the performance board.

The Sentry does not have an analog waveform generator. It does, however, have an IEEE-488 General Purpose Instrumentation Bus which allows a wide variety of standard and special analog as well as digital test equipment to be connected to the tester. Using this IEEE bus, the Sentry can automatically control the external instrumentation—thereby providing many more options for ASIC testing.

The Sentry can be used for testing data-acquisition-type analog circuits, but it is not really suitable for testing analog signal-processing circuits that require frequency-domain or distortion testing. Fortunately, at this time more ASICs are designed for data acquisition than for analog signal processing. ■

ABOUT THE AUTHORS

KENNETH DUBOWSKI *is an engineering section head in the ASIC Division of National Semiconductor. He is responsible for the development of analog standard cells. Since he joined National in 1976, he has also worked as a design engineer for the custom MOS/LSI group. He received a BSEE from Illinois Institute of Technology and an MSEE and MBA from Santa Clara University.*

TOM WONG *is the strategic marketing manager for National's ASIC Division, where he is responsible for new product planning of gate arrays and cell-based ICs. An employee of National for eight years, Tom has also served as an applications engineering manager in the hybrid systems products department. He has published numerous papers in areas such as optical communication, medical instrumentation, video amplifiers, high-frequency offline converters, data-acquisition systems and mixed analog/digital design. He holds a BSEE and an MSEE from University of Wisconsin, Madison.*

Easy ASIC

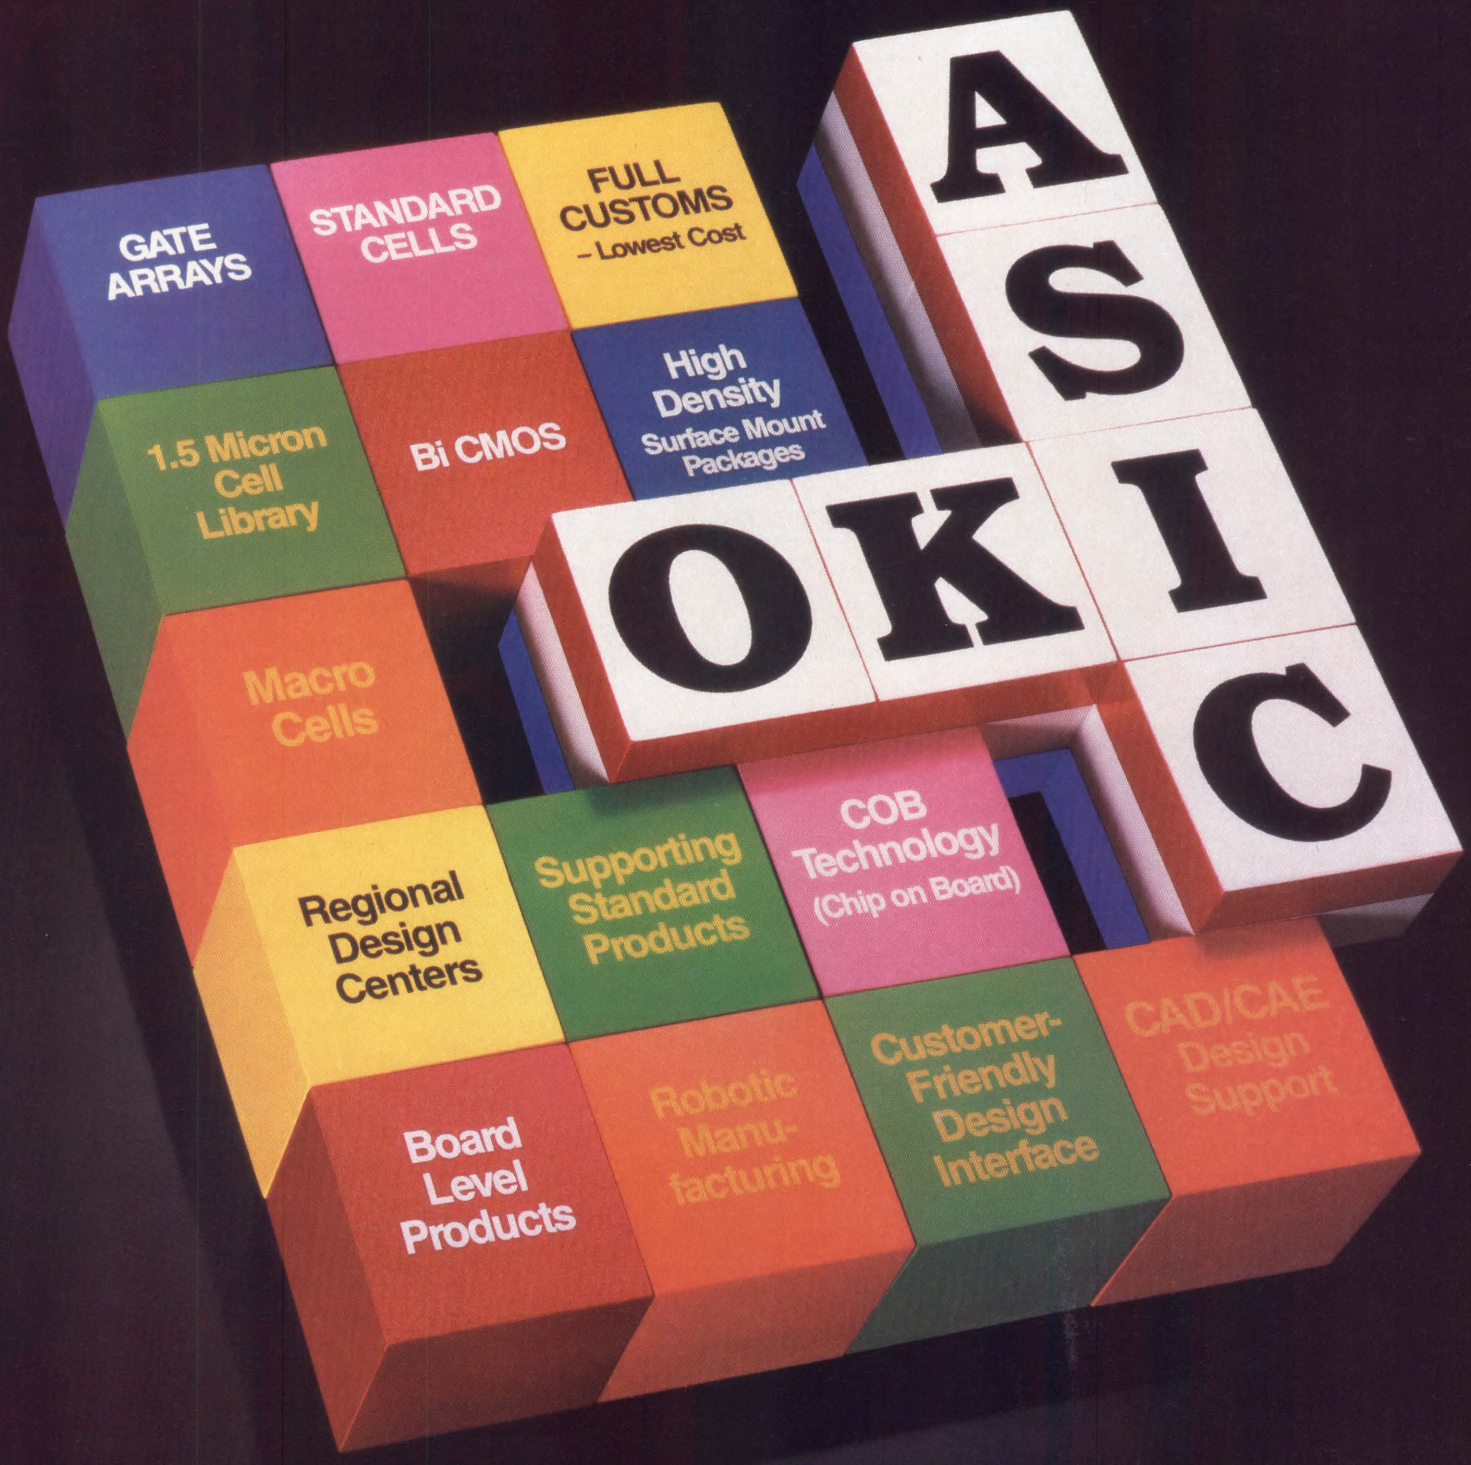
GATE ARRAYS
STANDARD CELLS
FULL CUSTOMS – Lowest Cost
1.5 Micron Cell Library
Bi CMOS
High Density Surface Mount Packages
Macro Cells
Regional Design Centers
Supporting Standard Products
COB Technology (Chip on Board)
Board Level Products
Robotic Manu-facturing
Customer-Friendly Design Interface
CAD/CAE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

Only from OKI:

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

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

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

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

OKI: the totally logical choice.

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

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

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

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

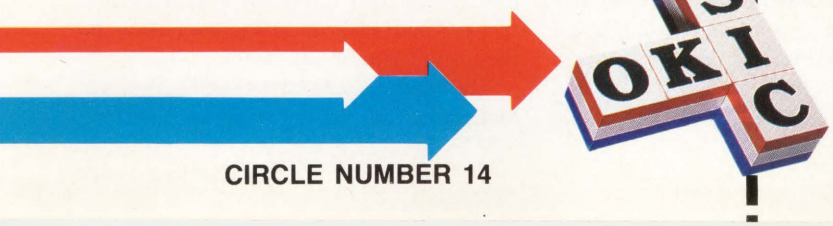

CIRCLE NUMBER 14

Check out OKI ASIC data:

VLSI 988

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

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

( ) Please call: we have immediate requirements:

Name/Title _____

Company _____

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

GaAs Technology *meets* RISC Architectures

BOB CUSHMAN, SENIOR EDITOR

There are two reasons for studying GaAs RISCs: first, its time to re-evaluate both RISC and GaAs—now that there are 32-bit GaAs RISCs in existence with 'first GaAs' running at 60 MHz and 200 MHz promised in two years. Second, the combination of simplicity enforced by the present GaAs technology and the potential for a 5X speed advantage over silicon make GaAs RISCs particularly useful for gaining insight into RISC design tradeoffs. We can only ponder whether the mating of GaAs with RISC will have such a clear speed advantage over currently-dominant silicon CISCs and rapidly-emerging silicon RISCs that it will be the start of new era. Also, is RISC what digital GaAs needs to finally get it past its long drawn-out infanthood?

These GaAs 'super-microprocessors' point up key RISC concepts, especially the use of compile-time software to offset architectural weaknesses. Their-necessarily longer pipelines depend on compiler reorganization of programs to prevent confusion during execution. The pipelines have to be longer as the speed increases to allow time for external memory accesses.

Both the GaAs RISCs we'll be describing have been sponsored by the Defense Advanced Research Agency (DARPA), Arlington, Va. Both are based loosely on a high-level instruction-set description of a 32-bit RISC that is an outgrowth of the original MIPS (microprocessor without interlocked pipeline stages) that evolved out of DARPA research contracts with Prof. John Hennessy of Stanford University. Though DARPA's aim is an ultra-high-performance, radiation-tolerant microprocessor for military programs, the hope is that these advanced prototypes will hasten commercial efforts. DARPA has also let contracts for two silicon versions of the MIPS RISC: to GE's Military Electronics Operation in Syracuse, NY: and to UNISYS Corp. (Minneapolis, Minn).

While the GaAs device developers are working on improvements that are expected to raise operating speeds to DARPA's goal of 200 MHz within the next two years, the silicon counterparts are reported at 40 MHz and headed toward 60 MHz, but with higher relative average throughputs because of higher levels of integration and increased parallelism.

One of the GaAs RISCs is the one we described briefly in our previous article in this series on RISCs (Ref 2). The CPU chip (Figure 1) is a joint project of Texas Instruments (Dallas, Tex.) for GaAs implementation, and Control Data Corp. (Minneapolis, Minn.) for the architecture. The 445 x 415 mil chip contains 12,895 gates. It uses a form of i^2L logic (which TI calls heterojunction integrated injection logic (Hi^2l) and has six pipeline stages. This bipolar type of circuitry tends to consume a lot of power, despite its small 400 to 500 mV logic swings, but TI expects to reduce the power to less than 10 W by the time they reach 200 MHz. If the power seems high, bear it mind that present CMOS RISCs, with their linear rise of power with speed, would scale up to well over 10 W, if they could run at 200 MHz.

The second GaAs CPU (Figure 2) is by McDonnel Douglas Astronautics Corp.(M-

This is the third in a series of articles on RISC architectures. The first two articles appeared in the June (p. 64) and July (p. 60) issues respectively.

ILLUSTRATION BY CHRIS SCHMIDT

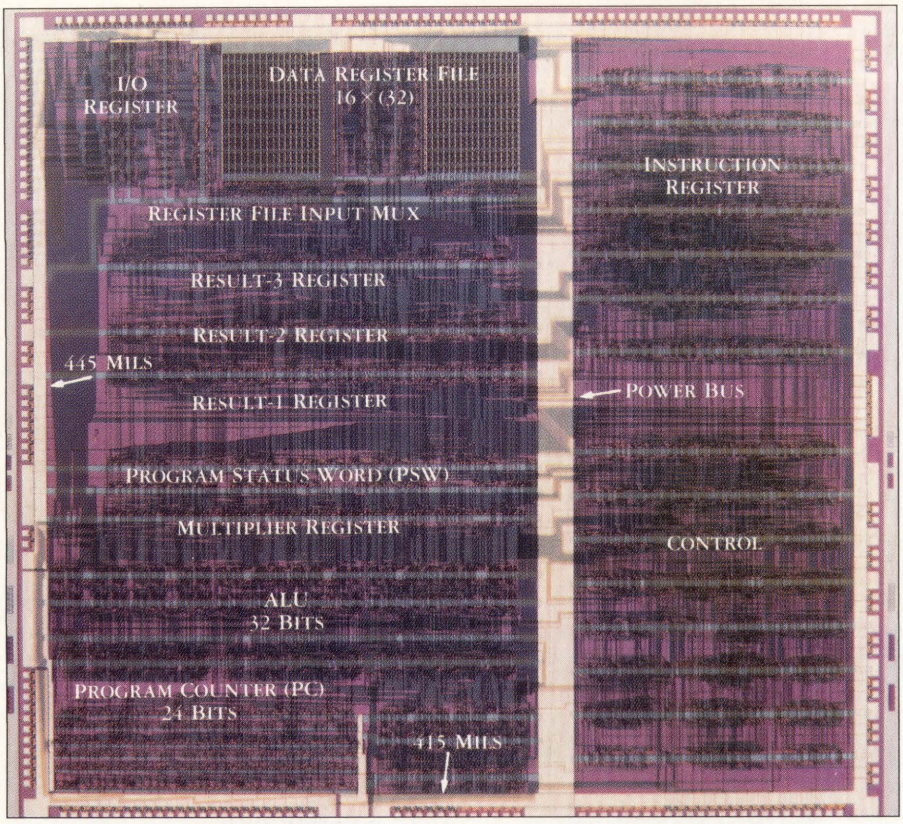

Figure 1. TI chip (a) with its floorplan overlayed. The large control area is partially due to TI's use of a semi-custom approach and will be tightened up for TI's next iteration.

DAC), of Huntington Beach, Calif. The MDAC chip uses the more conventional (for GaAs) JFET logic which is similar to NMOS silicon. The 335 x 437 mil chip contains 23,178 transistors and about 10,000 resistors. It has only 3 pipeline stages in the CPU but depends on external pipelining in the instruction fetching so that its total pipeline length is 5 stages—similar to TI's. It only consumes 3.5 W at 60 MHz—partially because the circuitry has allowed incorporating complementary P-channel devices in the registers without appreciable loss of speed. Neither TI nor MDAC expect their circuit power to go up appreciably as the speed is increased. MDAC projects that this CPU will consume 4 W average and 6 W maximum at 200 MHz.

■ Illustrating the Story

Figures 1, 2, 3 and 4 interrelate the chips to their underlying RISC architectures. The simplicity enforced by the GaAs complexity restrictions makes this possible, almost by visual inspection.

The chip floor plans (Figures 1b and 2b) reveal the layout of the main architectural elements, registers, and logic subsystems. The computer block diagrams (Figures 3a and 4a) show how the architectural elements are tied together by the buses to produce a working system.

Finally Figures 3b and 4b show the timewise pipeline flow through the architecture. This indicates how thesystem cycles through its basic, ever-repeating, fetch-execute behavior—as the system works on the sequence of software instructions delivered by the program. For clarity, brief written descriptions are shown for each pipeline stage.

The TI CPU handles all six stages of its pipeline. Of particular interest are the three RESULTS registers that are needed to space out the pipeline. They permit the two extra M1 and M2 cycles that send the operand address out and get the operand data back.

The MDAC CPU only handles part of its pipeline. It depends upon an external PC that is part of the memory system to fetch its instructions. The external PC is reminiscent of the old 8-bit F-8 uP. But in this case it was added to be able to adjust the depth of the instruction fetch pipeline independently of the CPU.

Comparing the chip areas in Figures 1 and 2 to the architectures diagramed in Figures 3 and 4 should help in understanding the appeal of the RISC concept. The value of RISC for these GaAs chips becomes even more apparent if you compare them to an elaborate 'embellished' RISC like the Motorola 88000 (Figure 5), which has roughly the same chip size. It can be seen that the bare-bones GaAs RISCs take up their whole chips in implementing the RISC CPU core that just takes up a fraction of the Motorola 88000. The 88000 has room to spare for a floating-point unit.

■ Simple Control = Speed

Note in Figure 2 the surprisingly small amount of area devoted to control in the MDAC chip. Most of central area of the MDAC CPU is taken up with data registers, with the control being effected by the narrow region on one side. MDAC says the control uses only 5 percent of the chip's transistors. According to MDAC, the bits in the instruction register (IR) are used very directly to control the ALU and to address the data registers. The hardwired PLA state machine that controls the pipeline sequencing is also very direct and simple because the designers have really adhered to the concept behind the MIPS philosophy—microprocessor without interlocking pipeline stages. They have not included hardware to interlock the pipeline stages, but have left that 'interlocking' up to the compiler. Interlocking refers among other things to the synchronizing of data flow between instructions in different stages of the pipeline. For example, interlocking—at run time by on-chip hardware or at compile time by a reorganizer—to prevent an instruction from going ahead until its operands are available.

TI's chip uses much more area for its control, but that is partially due to TI's use of a semi-custom approach rather than the full-custom approach used by MDAC. Also TI needed more area to control all the pipeline stages.

What Figures 3 and 4 can't fully show is all the circuit short cuts that speed up execution. For example, when an instruction needs data for the ALU that has been received in the result or input register but hasn't been placed into a data-file register, these architectures often provide by-pass paths to gate data directly to the ALU.

■ GaAs Frugality Hurts

On-chip data storage is one part of the architecture where GaAs's limitation on the number of transistors per chip puts GaAs at a distinct disadvantage in implementing RISC. The TI register file contains only 16 registers. The MDAC register file at present contains just 24 registers of which only 17 are general purpose. These are small numbers compared to the 136 registers of the Sun SPARC or 192 of the AMD 29000 silicon RISCs. The effect of this paucity of on-chip data registers is compounded by GaAs's inferior performance when going on or off chip, and the fact that GaAs—because of the same transistor limitations—requires more chips.

PERFORMANCE

16-bit MICROPROCESSOR

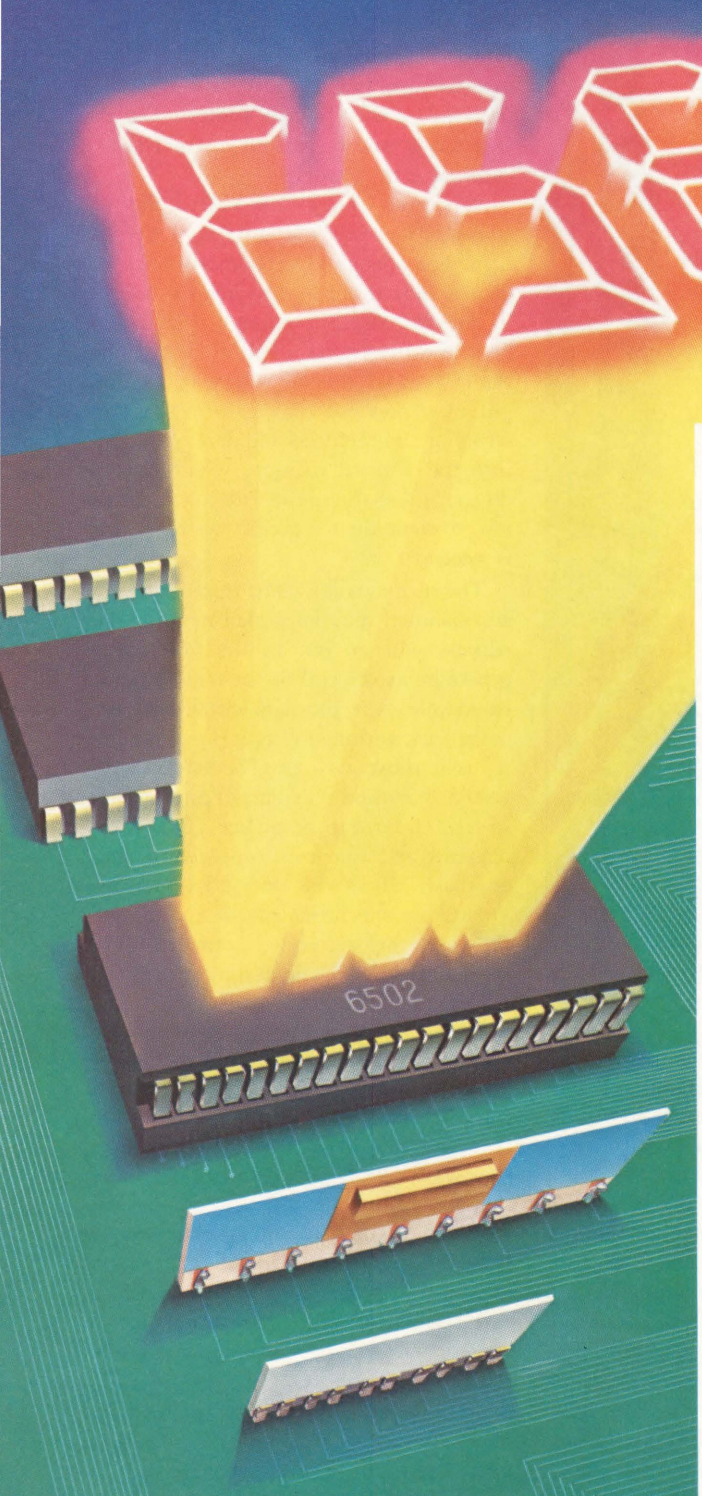

16-bit Performance Boost for 6502 Designs

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 15

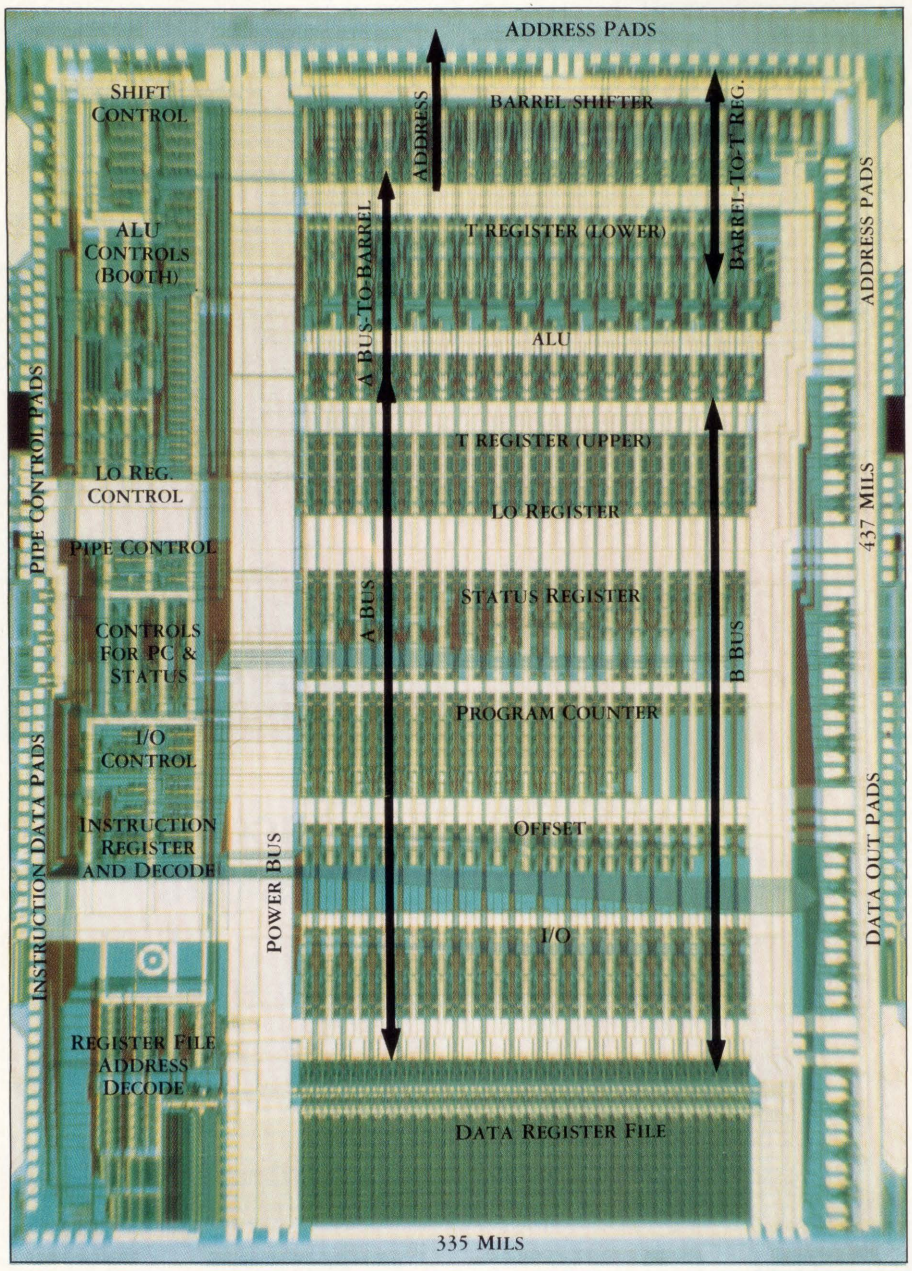

Figure 2. MDAC chip (a) with its floorplan overlayed. The area used for control is remarkably small, just 5 %, partially because some of the control is merged with external memory.

This same transistor count limitation also prevents GaAs chips from having other desirable datapath functions on chip, such as multiplers and floating-point units. However, MDAC has managed to squeeze in a barrel shifter and a 2-bit-per-cycle Booth multiplier assist.

■ Pipeline to Memory

It's hard enough to run a CPU at 200 MHz but it is even more of a challenge to access the memory within the 5 ns clock intervals at that speed. Nevertheless it is fundamental to RISC performance that the memory be accessed for a new instruction each and every clock cycle. The two GaAs RISCs use quite different approaches to accessing the memory.

TI uses full Harvard busing with separate address and data bus pairs for both instructions and data operands. (Fortunately the large areas dictated by GaAs' low device density produce chips with adequate periphery length for the buses). Then TI has the CPU count out extra pipe stages so that there is a stage for sending the instruction address out and a stage for moving instruction in and, similalry, two stages for the operand addresses and data (Figure 3b). These extra stages give 2 to 2.5 ns GaAs memories, such as those made by Vitesse Electronics Corp. (Camarillo, Calif.), and Gigabit Logic Inc. (Newbury Park, Calif.) time to respond. These memories have various combinations of address and data latching so they can play their part in the pipelining. Along these lines, DARPA has sponsored GaAs memories that match the voltage levels of the TI and MDAC GaAs chips.

MDAC implements the extra pipeline stages for the instruction fetch externally. It duplicates the PC as shown in Figure 4b. The external PC is synchronized with the internal PC upon initialization and from then on is incremented on each clock to fetch one instruction after another.

One advantage of this scheme is that only a single address bus is needed. Normally that address bus is used to direct the flow of the data operands to and from external operand memory, but it is also used to transfer branch addresses to the external PC and to pass on instructions to external coprocessors.

Whatever its use, the address bus becomes active in the CPU WRITEBACK pipestage. For operand STORE instructions the data is moved out along with the address in the WRITEBACK pipestage. The data to be stored to external memory comes from the DATA OUT register where it has been put by a previous register-to-register instruction. The address and data are latched into the external memory which can the do the actual storage in a decoupled fashion, to some degree taking whatever time is needed.

The operands for LOAD instructions are also handled in a decoupled manner. The address will go out in the WRITEBACK pipestage at the end of the LOAD instruction cycle, but the data doesn't have to come back within any rigid pipestage limit, as it must with TI. The LOAD data is FIFO'd into the CPU's data-input register, so that LOADs can be picked up as source operands in sequence. Compared to the TI approach where the operand LOADs and STORES are synchronously locked into the pipeline, this decoupled approach allows more flexibility in the design of the external operand memory.

The compiler's re-organizer must see that LOAD commands are issued sufficiently ahead of the need for the data. The compiler may be forced to insert additional instructions to provide enough time for the LOAD data to get into the CPU input register. In that case, the compiler's challenge is to find useful instructions for these delay slots; the last thing that should be in a RISC program is NOPs. But just in case a LOAD doesn't get the data in place in time MDAC designers have relented from the pure MIPS philosophy and provided a 'scoreboarding' bit that implements a hardware interlock to stall the instruction until the input data is available.

■ Non-Sequential Problem

Pipelining schemes work smoothly for straight-line sequential code, but not for

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.

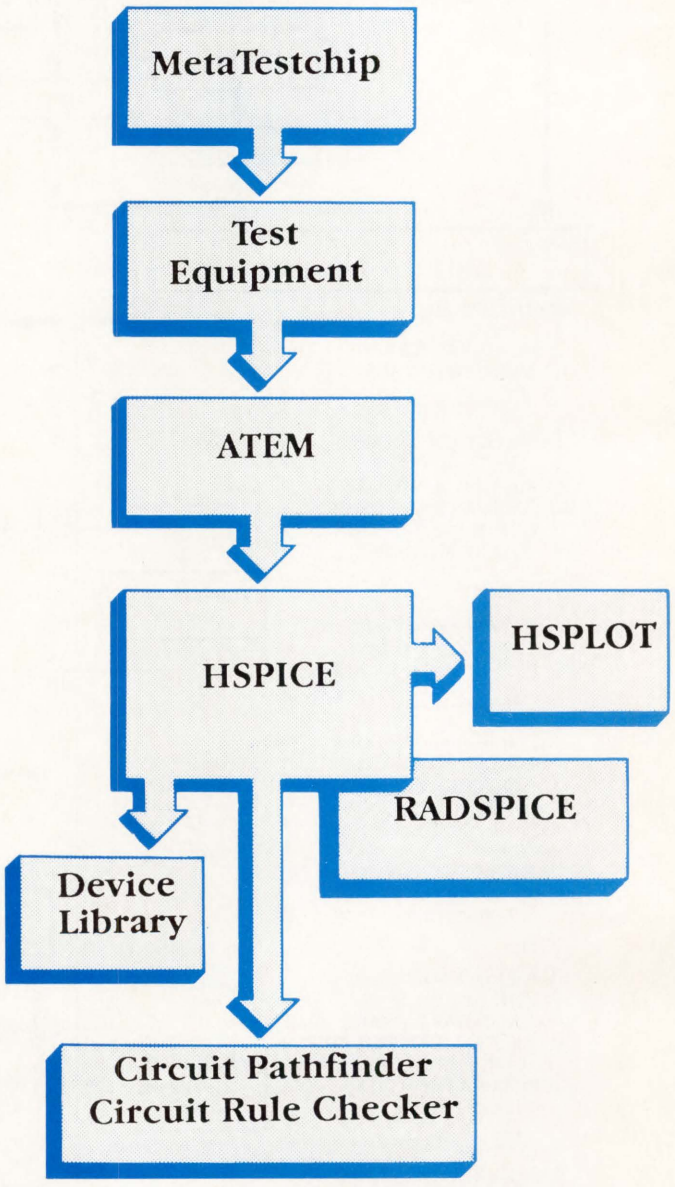

Meta-Software provides all essential support services. For answers to technical questions from experienced engineers, please call Meta's toll-free HOTLINE SUPPORT number: (800) 346-5953

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

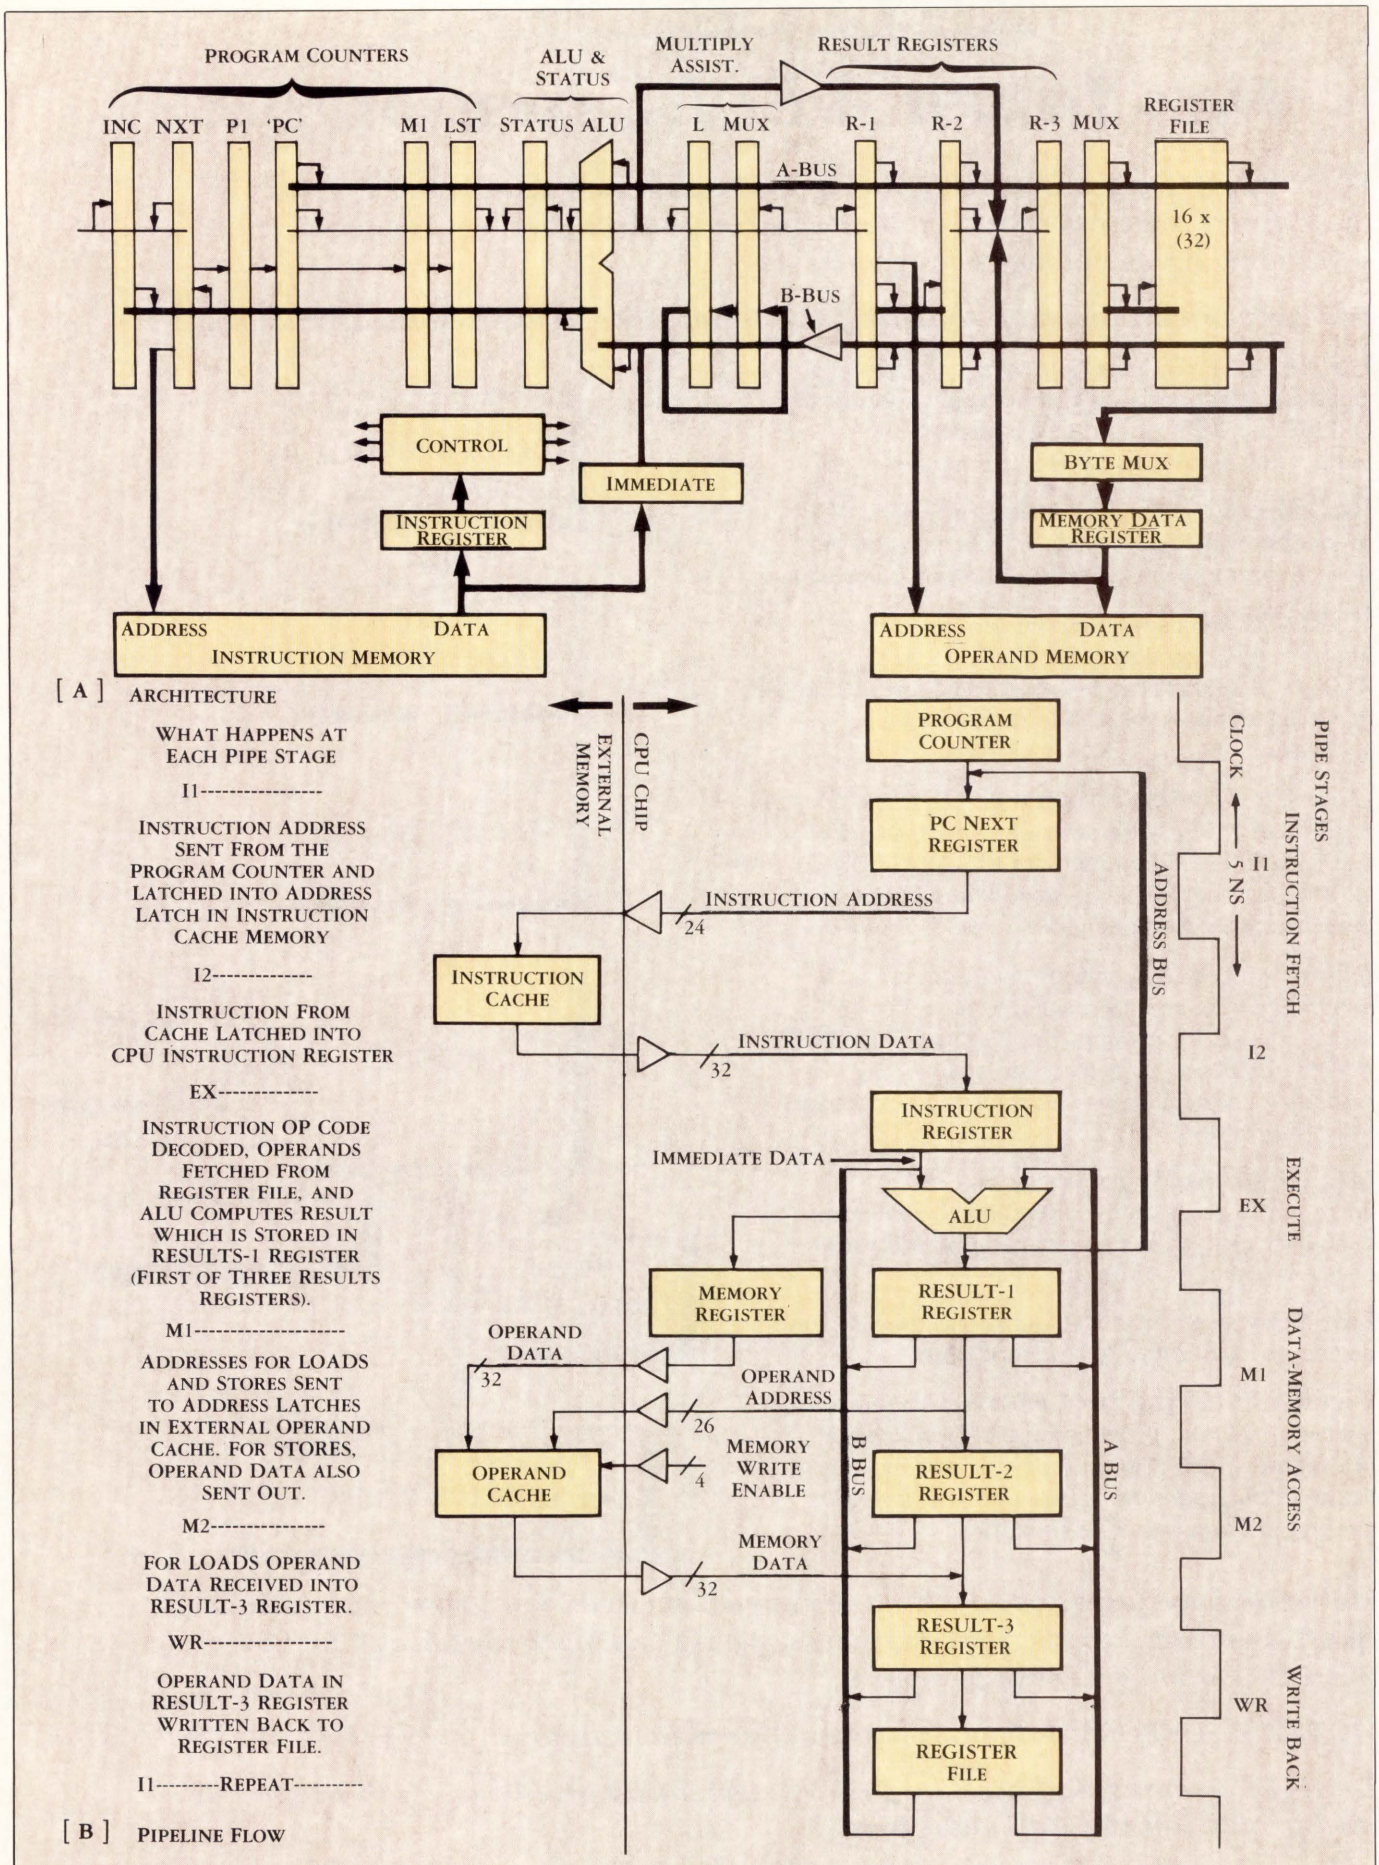

Figure 3. TI architecture (a) and pipeline flow (b). The long pipelines become necessary when interfacing to external memory at high speed.

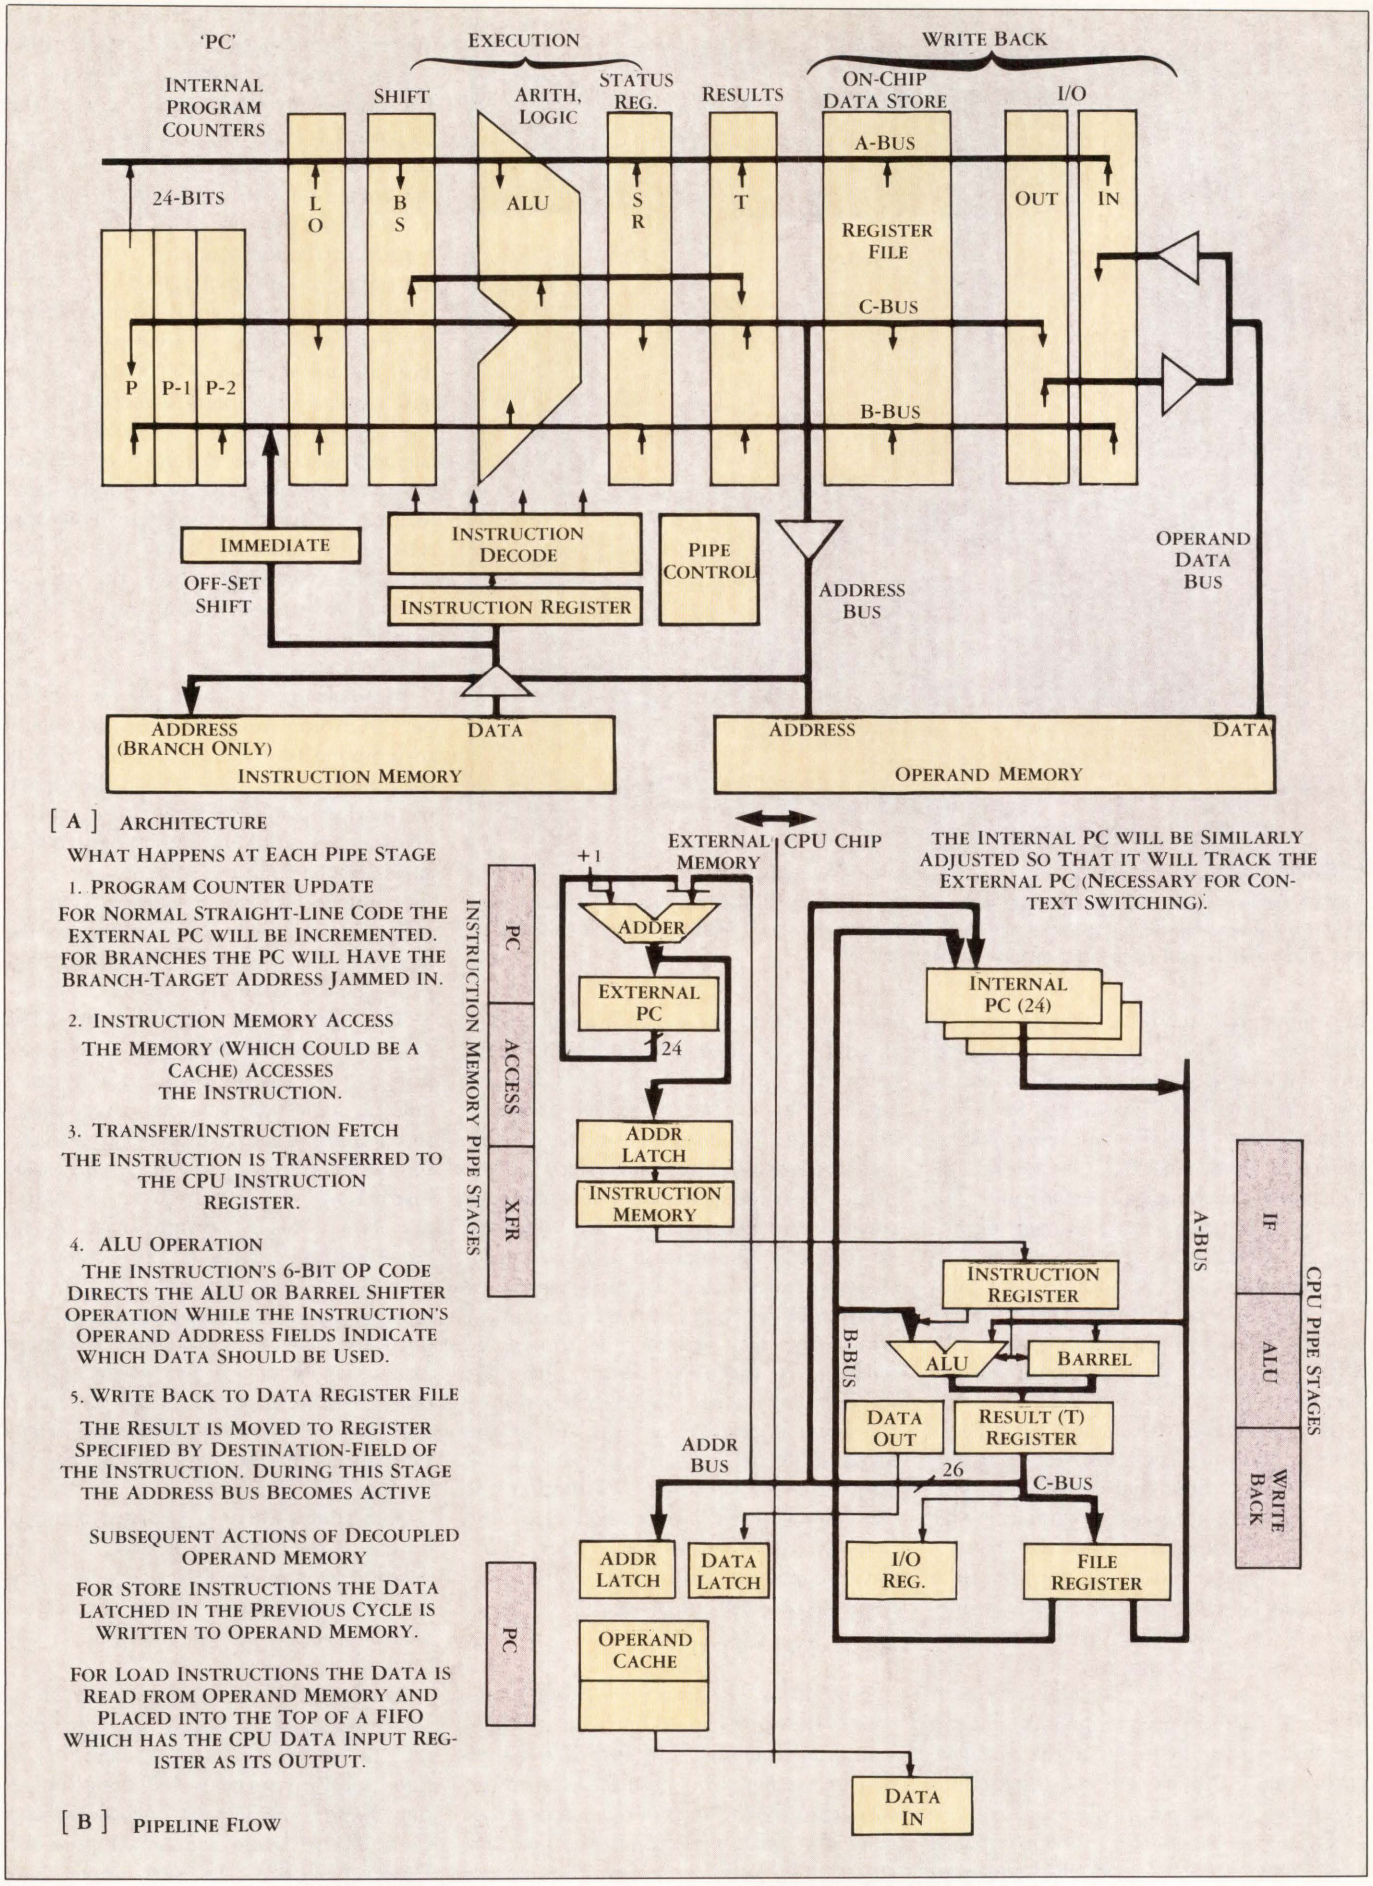

Figure 4. MDAC architecture (a) and pipeline flow (b). The external memories are decoupled from the CPU pipeline for application flexbility.

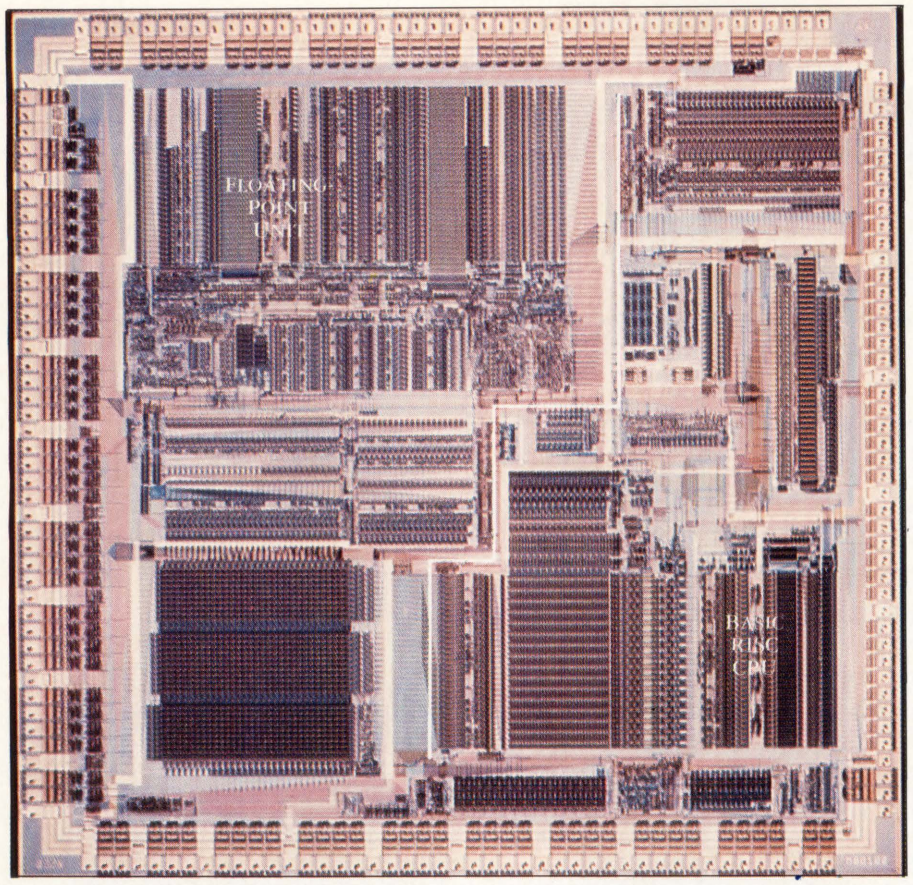

Figure 5. Motorola 88000 RISC has a chip size that is similar to the two GaAs RISCs but because the Motorola part has the higher density permitted by silicon, only about 1/3rd of the chip area is devoted to the basic RISC CPU. The rest of the area is used for a large, on-chip floating-point unit.

non-sequential fetching that occurs in branches. Then the pipeline is disrupted and a poorly designed system can throw away the benefits of pipelining. Such disruptions are why the sustained performance of a RISC (or CISC) CPU averages so much lower than itspeak performance. For example, TI estimates its average performance will be about 130 MIPS rather than the 200 MIPS peak. Similarly, MDAC projects only 125 to 150 MIPS. Let's look at how the breaks in the pipeline that occur with discontinuities in program flow are handled.

A generalized three-stage pipeline is shown with just one instruction going down it at a time (Figure 6a) and with three instructions going down it one after the other (Figure 6b). If just one instruction is clocked down the pipeline at a time, the control is simpler but only one-third of the hardware is used, and, more important for RISC, the rate of instruction execution is just 0.33 per clock cycle. If three instructions are being clocked down the three-stage pipeline then the hardware utilization is 100 percent and instructions execution is one per clock.

Figure 6c shows the problem when a conditional branch breaks the sequential accessing of instructions. In a commonly occurring situation, a conditional branch instruction is used to decide whether a software loop has been executed the desired number of times. The loop iterations are counted by decrementing a register and then the register is tested by the branch instruction to see if the count has reached zero, where the iterations should be stopped. The conditional branch tests the zero status bit and makes known its decision by its choice of the location from which the next instruction should be fetched. If the loop counter has not reached zero and the loop is to be repeated the branch instruction will jam in the location of the first instruction in the loop—the "branch target instruction". If the loop counter has reached zero, the branch decision will let the PC go fetch the next sequential instruction as it normallywould (fall out of loop).

■ Branch Timing

Now if just one instruction were going down the pipeline (Figure 6a), there would be time for the branch decision to deliver the non-sequential location to the PC in time for the fetch of the next instruction. But with pipelining, the next instruction (or perhaps the next two instructions) have already been fetched and are on their way down the pipeline by the time the branch instruction makes its decision.

This problem is detailed in Figure 6c. The dashed "not ready" arrow between the T1 and T2 clock times shows how the branch decision comes too late to address the branch-target instruction in case the decision goes that way. This is because the branch instruction makes its decision during its execution pipestage. This is when it tests the zero status bit, and sends out the branch address, if appropriate. The dash-dot arrow between T2 and T3 shows how a misaligned branch decision could send the program back to the loop after the next instruction beyond the loop had executed, possibly causing complete confusion. The associated program-flow chart relates the problem from the software viewpoint, showing where the various instructions would be with respect to the loop. The diamond symbol is of course the branch instruction.

In pipelined CISC machines the remedy has often been to use hardware interlocks to temporarily delay or stall the fetch of the following instruction until the branch instruction has completed its execution pipestage. But in the most popular RISC approach the compiler inserts a non-critical command for the next instruction. (or as many instructions as is dictated by the length of the pipeline). Then what happens doesn't make any difference.

The catch is that because of the RISC emphasis on performance, taking the easy way out by having the compiler insert safe-but-useless instructions is very undesirable. RISC performance is often determined by the per cent of useful delay instructions. The situation is not unlike that where MDAC might need delays while waiting for a LOAD to complete. Usually in a loop like this, the compiler can juggle or "reorganize" the code sequence so that the loop counter decrement instruction (which is also needed for repeats), will be placed in the delay slot after the branch. In Figure 6c this is indicated symbolically by the reorganized version of the program flow chart.

Sometimes, there is also provision for nullifying the delay instruction in case the decision goes in the direction where it is not needed (though a loop decrement would probably not be harmful in any case).

The longer the pipeline, the more delay slots that must be added, and the more difficult it becomes for the compiler to find useful instructions. TI's longer pipeline (Figure 3b) demands two delay slots versus MDAC's one delay slot (Figure 4b). However TI says its simulations have shown it can find useful instructions for

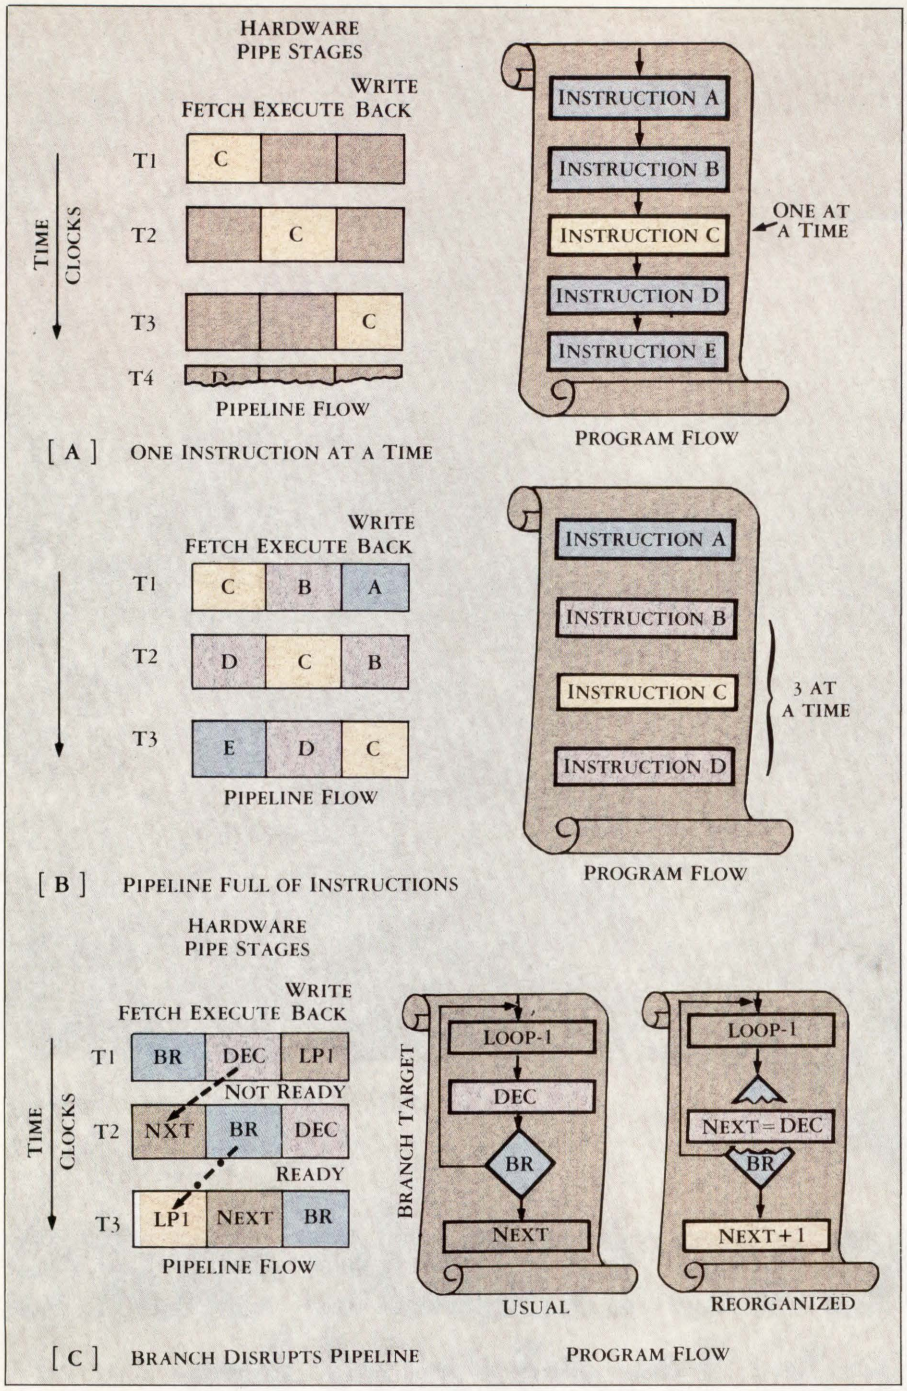

Figure 6. Pipeline benefits and problems are illustrated by a hypothetical 3-stage pipeline. If only one instruction flows down the pipeline at a time (a) then the architecture is realizing only 1/3rd of its potential. If new instructions are being constantly clocked in one after the other (b), then the architecture's MIPS rate is the same as its pipe-advance clock rate, or 3 times higher than in (a). However, if a branch instruction comes along (c), then only 'safe' instructions must be fetched until the branch instruction has executed and decided if the loop is to be repeated.

the additional delay slot in a sufficient percentage of the time to prevent undue degradation of their average MIPS rate. (See Ref. 5 for more on this problem).

Although the MDAC needs only one delay slot its use of an external PC introduces complications. As shown in Figure 4b, there is a path by which the address bus (normally inactive in instruction fetching) becomes active and jams in the branch target instruction address. Not shown is a delay storage register to hold the branch target instruction so it will be properly timed for the pipeline despite the extra memory pipestages.

MDAC is putting these features into an external branch target cache: (similar to that used in AMD 29000). Idea is that first time around a loop the system will have to stop and get the branch target instruction from memory but after the first time around the loop, the branch target instruction will be in the branch target cache (along with some of the following loop instructions) so that the pipeline will stay full despite the break.

Probably the most difficult pipeline break to handle is that which occurs upon interrupt. Both GaAs RISC chips keep track of the addresses of the instructions going down the pipe by passing the PC values used to fetch the instructions down a series of PC's. These additional PC's can be seen in Figures 3 and 4. Note that for the MDAC architecture this is when the otherwise useless on-chip PC's of the CPU become of use. Then upon interrupt the contents of these "history" PC's are saved so that all the instructions in the pipe can be reloaded at the return from interrupt to restart the pipeline from where it left off at interrupt.

MDAC admits that their unusual external PC arrangement means they have to provide additional registers to hold information during interrup. Goal here is fast context switching ■

ACKNOWLEDGMENTS

For the information on which this article was based: David Whitmire, project leader and Steven Sabin, lead design engineer, Texas Instruments and Shaun Whalen, Control Data, and Terrence Rasset, manager of systems architecture and Roger Niederland, manager of processor design, McDonnell Douglas Astronautics and Sherman Karp, consultant, DARPA.

REFERENCES

CUSHMAN, B.,*'RISC Changes the Balance, VLSI Systems Design, June 88, Pg. 64. Part 1 of this series.*

CUSHMAN, B., *'Surveying the RISC Realm,' VLSI Systems Design, July '88, Pg. 60. Part 2 of this series.*

FOX, E., KIEFER, K., VANGEN, R., WHALEN, S., *'Reduced Instruction Set Architecture for a GaAs Microprocessor System,' IEEE Computer, Oct '86. Pg 71. Initial plans for TI project, some of which have changed.*

RASSET, T., NIEDERLAND, R., LANE, J., GEIDEMAN, W., *'A 32-bit RISC implemented in Enhacement Mode JFET GaAs,' IEEE Computer,' Oct '86, Pg. 60. Initial plans for MDAC project, some of which have been modified.*

MILUTINOVIC, V., FURA, D., HELBIG, W., LINN, J., *'Architecture/Compiler Synergism in GaAs Computer Systems,' IEEE Computer, May '87, Pg 72. Examination of the special relationship between bare-bones GaAs RISCs and their compilers. Analysis of possible code-rearrangements that the compiler could effect.*

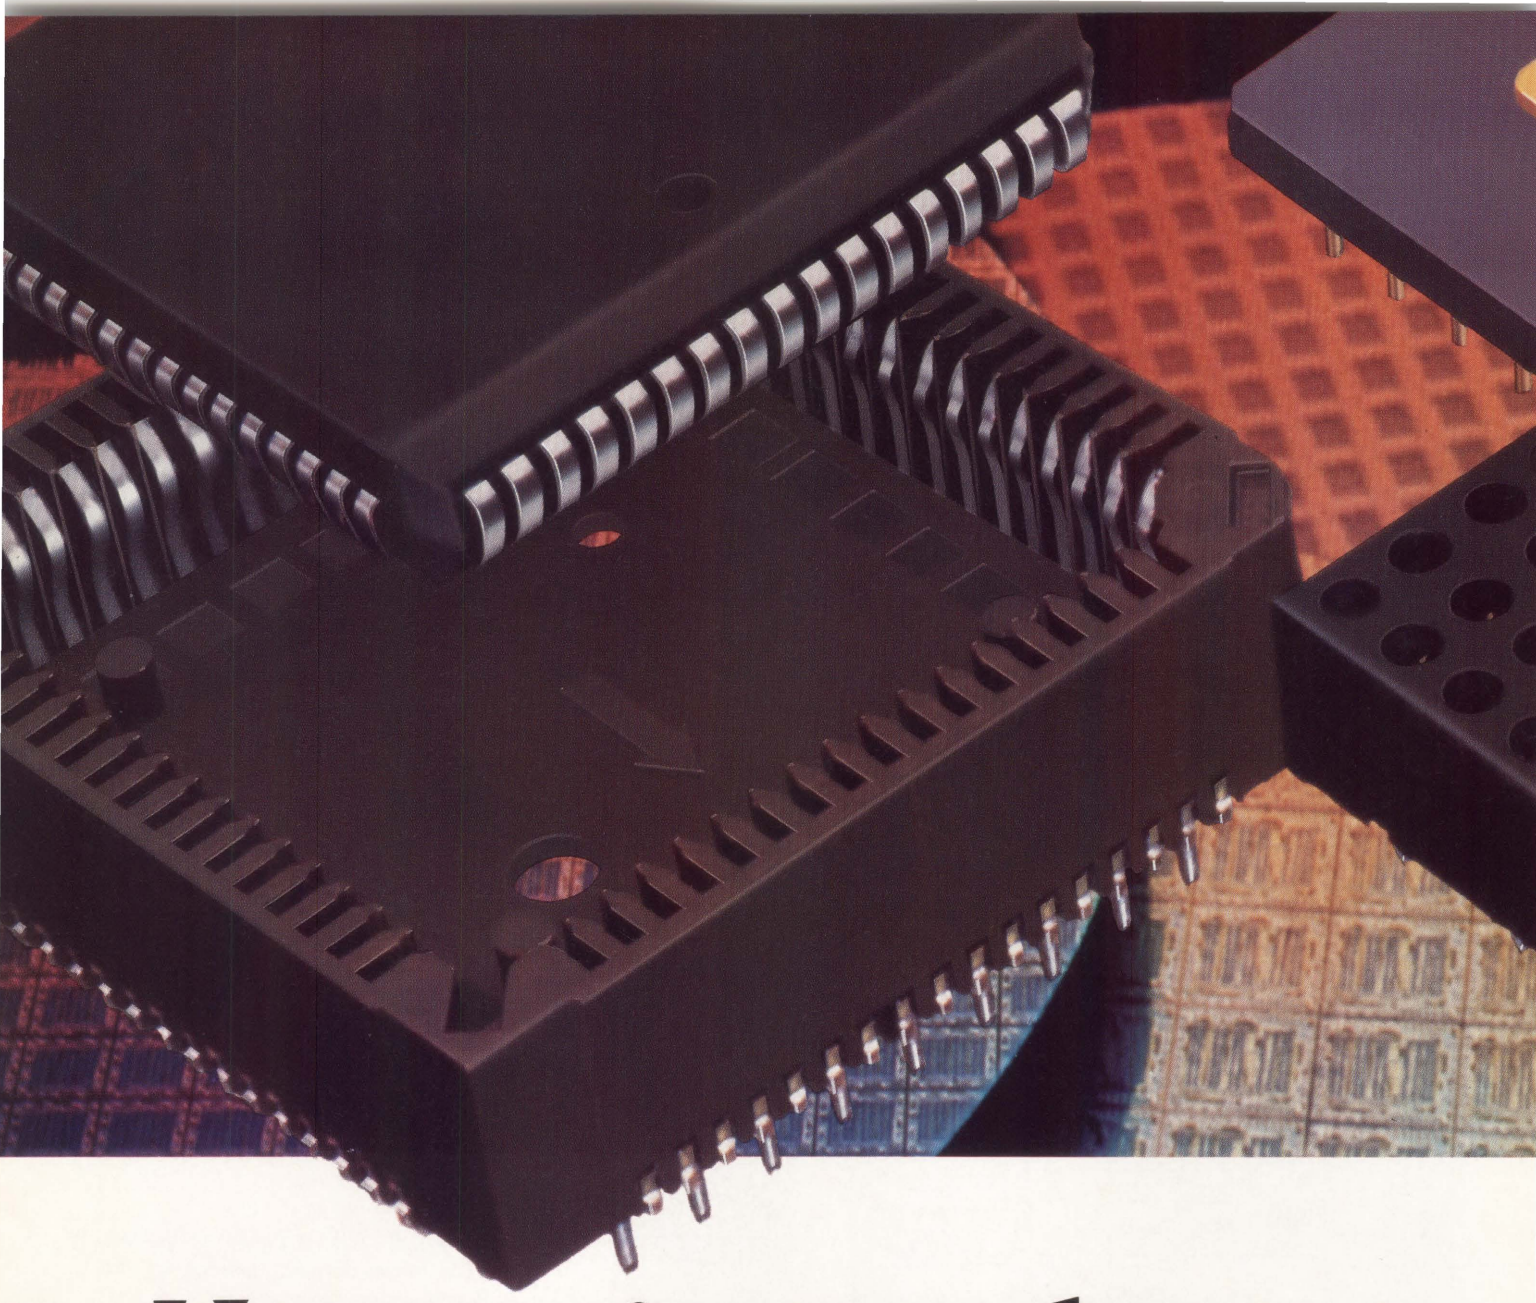

How we've made sure the chip's locked in, but you're not.

ZIF and LIF PGA sockets, plus our minimal profile spring sockets in PGA footprint.

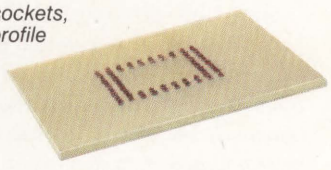

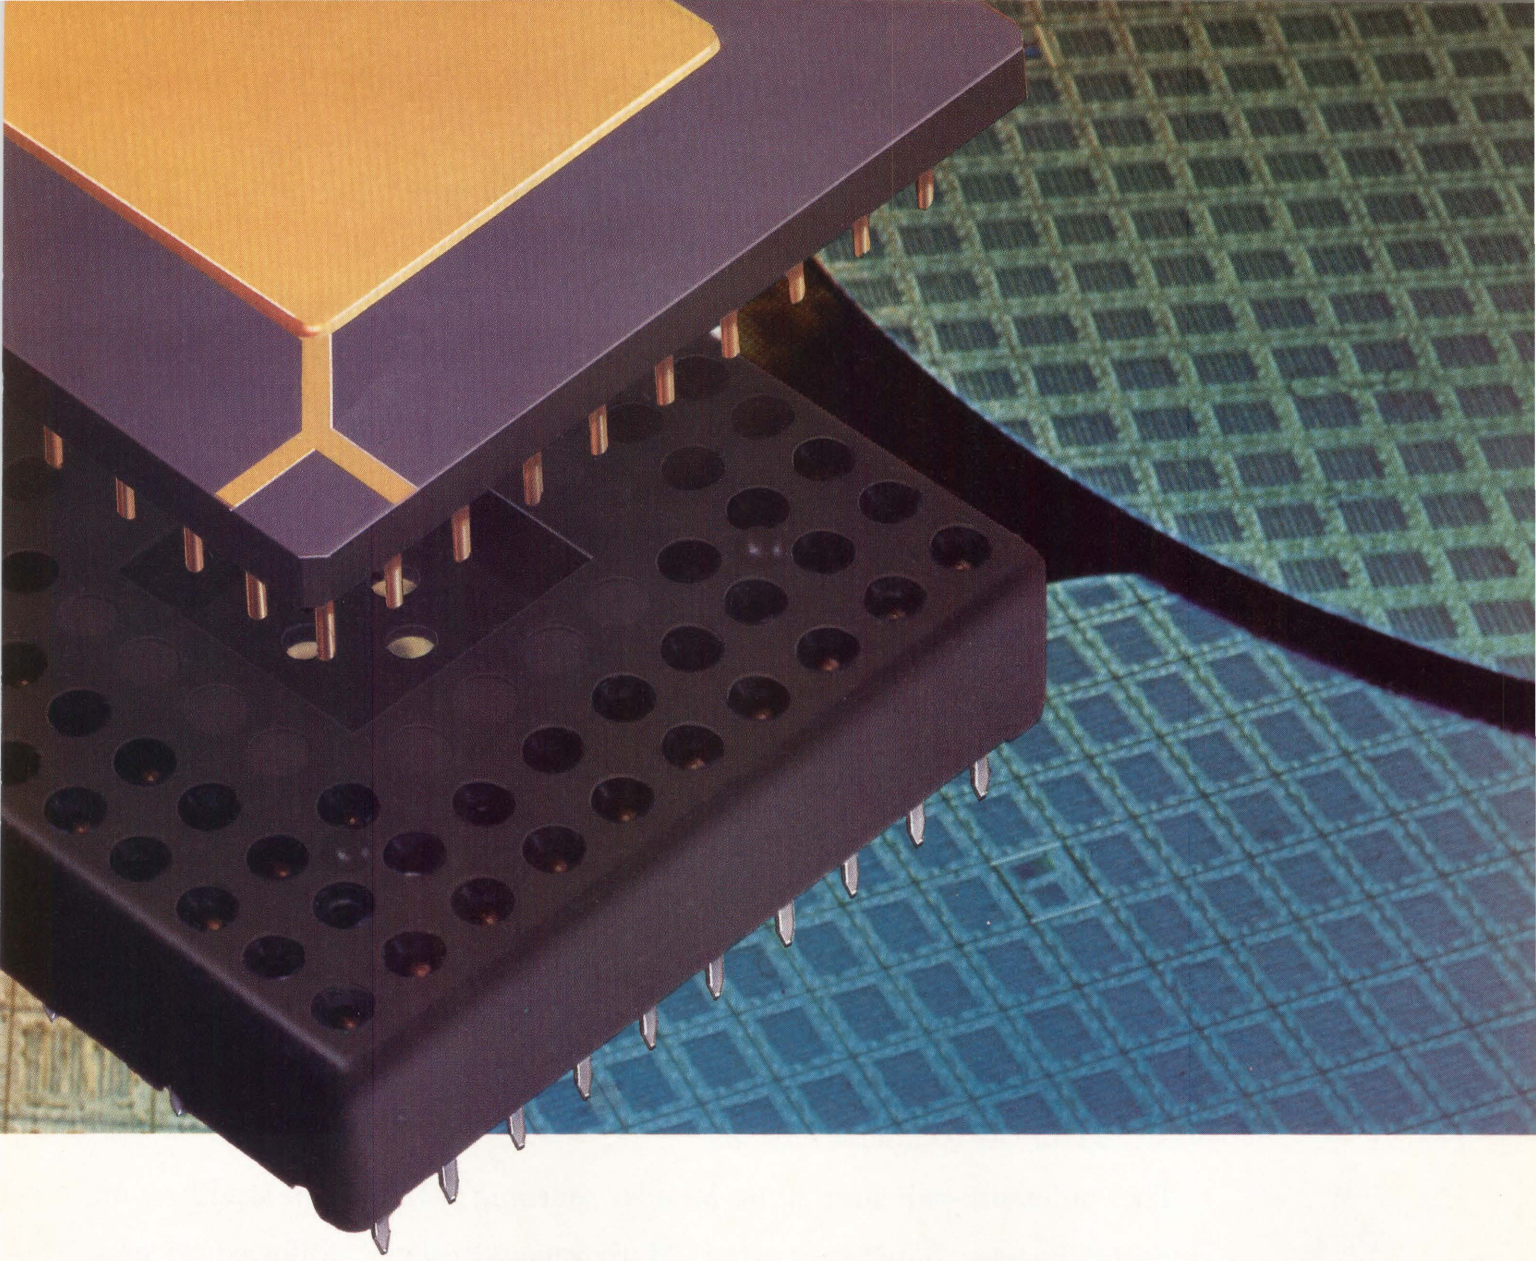

Whichever direction you take in microprocessor technology, AMP makes sure you have the socketing options you need to make it pay off.

Our high pressure tin sockets for plastic leaded packages come with an exclusive Positive Lock retention system that keeps chips secure during handling and shipping. Ceramic chip carrier sockets feature duplex plated contacts and snap-on covers that accommodate heat sinks. Both are available in standard and surface-mount versions.

AMP offers you more: high-speed, surface-mount sockets on .020″ centers. Gold-plated plastic carriers and sockets. For pin grid arrays, sockets in ZIF and LIF styles, plus custom VHSIC capability.

Choose your technology. AMP makes it easy to implement, with full socketing support.

Call 1-800-522-6752 for more information on AMP sockets. AMP Incorporated, Harrisburg, PA 17105-3608.

AMP Interconnecting ideas

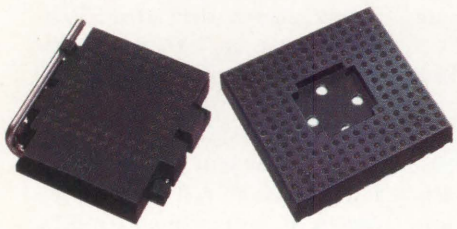

Low-height sockets in all standard sizes (JEDEC A, B, D), with duplex-plated contacts for sure performance. High-speed, surface-mount sockets feature 0.5pF, 1.4nH contact characteristics.

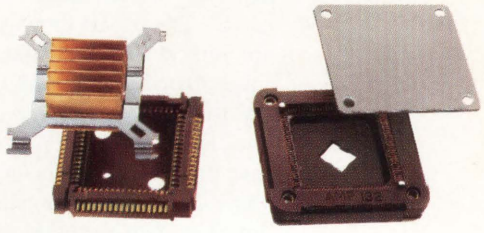

CIRCLE NUMBER 17

International Conference on

COMPUTER DESIGN

VLSI Systems Design Staff

The overall theme of the annual International Conference on Computer Design (ICCD), to be held on October 3-5 in Rye Brook, NY, will be "VLSI in Computers and Processors." A wide variety of topics that stress the interrelationships that exist today in the CAD industry will be covered in 126 papers, presented in 36 technical sessions.

ICCD'88 STRESSES INTERRELATIONSHIPS BETWEEN TECHNOLOGIES

The conference will start off on Monday morning with a keynote address on "General Purpose Supercomputing: Myth versus Reality," followed by four plenary sessions that will present papers addressing the current state of the art for the conference themes. They are: VLSI and Technology (ASIC Technology); Architecture (When will the widespread use of parallel computers become a reality?); Design and Test (Design now, test later?); and CAD (The evolution of the CAD industry from high growth to maturity).

VLSI and Technology: High performance will be on the minds of those who attend the Tuesday session on "High speed VLSI circuits." The papers for this session are typical of most sessions and cover a pot-pouri of topics reflecting the interactions between various technical aspects of the session theme. Topics include high-speed CMOS chips and cells, and a design system for VLSI circuits.

Also one Tuesday is a session on "VLSI Design Integration." Topics include processor design, macrocell-based switching circuits, synthesis of integral designs, and integrating test into the VLSI CAD environment.

Design and Test: This theme starts with a Monday session on "Test Generation," and continues to get attention at a Tuesday panel session on "Proposed Standards in Design and Test." There are also two sessions on Wednesday covering both software and hardware tools used in design for test.

Computer Aided Design: Two Tuesday sessions should be of interest to those involved in IC-CAD. The session on "Cell layout techniques" touches base with the latest technologies in automatic layout, optimization, and compaction. The other session will present three papers on "Silicon Compilation at Philips Research.

Architectures and Algorithms: Monday's session on "Algorithms" covers matrix factorization and computation, in addition to array processors. Microprocessor architectures are also covered on Monday, while architectures for bigger machines and arithmetic algorithms are discussed Wednesday." If you want to put these algorithms to work, then attend Tuesday's session on "VLSI Implementation of Arithmetic Algorithms."

There's lots more sessions covering analog design techniques, advanced system interconnects and packaging , and system applications of the new high temperature superconductors.

INTERNATIONAL CONFERENCE ON COMPUTER DESIGN

RYE TOWN HILTON, RYE BROOK, N.Y., OCTOBER 3–5, 1988

MONDAY 9:00 am	**Introduction and Keynote Address: "General Purpose Supercomputing: Myth vs. Reality,"** by David Kuck, University of Illinois Chair: P. Agrawal, AT&T Bell Labs			
10:30 am	**Technology Plenary** Chair: E. Middlesworth, Hewlett-Packard		**Design and Test Plenary** Chair: S.M. Kang, University of Illinois	
11:30 am	**Architecture Plenary** Chair: W.J. Dally, MIT		**CAD Plenary** Chair: A. Domic, Digital Equipment Corp.	
2:00 pm	**Parallel Processing for VLSI CAD and Testing** Chair: P. Bose, IBM Watson Research Center	**Supercomputing Technology** Chair: W. Maly, Carnegie-Mellon	**Algorithms** Chair: W. Dally, MIT	**Analog Design Techniques** Chair: C. Zukowski, Columbia University
4:00 pm	**Test Generation** Chair: D. Ostapko, IBM Watson Research Center	**Advanced System Interconnect and Packaging** Chair: J. Prince, University of Arizona	**Microprocessor Architecture** Chair: U. Weiser, Intel Corp.	**Simulation** Chair: J. White, MIT
8:00 pm	**Panel: Supercomputers, Challenges and Design, and Applications** Moderator: P. Bose, IBM Research			
TUESDAY 8:30 am	**VLSI Array Testing** Chair: N. Jha, Princeton University	**Silicon Compilation at Philips Research** Chair: C. Niessen, Philips Research	**System Applications of High Tc Superconductors** Chair: B. Ackland, AT&T Bell Labs	**Cell Layout Techniques** Chair: P. Wolff, IBM East Fishkill
10:30 am	**VLSI Design Integration** Chair: S. Kang, University of Illimois	**VLSI Implementation of Arithmetic Algorithms** Chair: V. Oklobdzija, IBM Watson Research Center	**System Support for the NS 32532** Chair: J. Payne, National Semiconductor	**Logic Synthesis and Technology Mapping** Chair: D. Hill, AT&T Bell Labs
2:00 pm	**Error Analysis** Chair: E. Cerny, University of Montreal	**ASIC Technology** Chair: J. Vander Speigel, University of Pennsylvania	**High Performance Computer Systems** Chair: A. Fisher, Carnegie-Mellon University	**Interactive System Design** Chair: D. Boyer, Bell Communication Research
4:00 pm	**Panel: Proposed Standards in Design & Test** Chair: A. Ambler, Brunel University	**High Speed VLSI Circuits** Chair: P. Gelsinger, Intel Corp.	**Microprocessor Design** Chair: B. Shimamoto, IBM Research	**High-Level Synthesis** Chair: R. Camposano, IBM Research
WEDNESDAY 8:30 am	**Design for Test I: Software Tools** Chair: J. Hancock, IBM	**Systolic Arrays** Chair: A. Gupta, Stanford University	**Placement** Chair: C. Sechen, Yale University	
10:30 am	**Design for Test II: Hardware Tools** Chair: A. Albicki, University of Rochester	**Computer Design** Chair: D. Ling, IBM Research	**Arithmetic Algorithms** Chair: M. Ercegovac, UCLA	
2:00 pm	**Vertically Integrated VLSI Design for DSP** Chair: M. Breuer, USC	**Multiprocessing** Chair: A. Agarwal, MIT	**Simulated Annealing** Chair: S. Kirkpatrick, IBM Watson Research Center	
4:00 pm	**Reconfigurable VLSI Processor Arrays** Chair: R. Chandramouli, Sun Microsystems Inc.	**Signal Processing** Chair: T. Knight, Symbolics	**Formal Verification** Chair: P. Prinetto, Politecnico di Torino	

High-level, top-down designs meet logic synthesizer's bottom-up feedback

High-Level Designs Feed Logic Synthesizer

ONE benefit of logic synthesis technology is the automatic generation of a physical design from a high-level specification. Bypassing logic reduction and implementation reduces design time and eliminates errors. Rapid feedback from a real physical implementation gives the designer an accurate estimate of the results of decisions made at a higher level in the design hierarchy. In effect, the designer can perform a top-down design, with the logic synthesis software providing "bottom-up" feedback.

Synopsys Inc., has created an interface to high-level hardware description languages (HDLs) to make this benefit possible. Called HDL Compiler, it is a generic interface, called a "framework" for multiple languages, with ports to specific HDLs.

The first HDL port accepts designs expressed in the Verilog language from Gateway Design Automation (Lowell, Mass.) The product is a result of a Verilog licensing agreement signed with Gateway in May of this year. The company plans to add a subsequent port for behavioral-level descriptions expressed in VHDL, although no availability date has been announced.

The HDL compiler acts as another source of input to Synopsys' Design Compiler (see Figure). The Design Compiler, introduced at this year's Design Automation Conference (see *VLSI Systems Design,* June 1988, p. 12) accepts designs in terms of netlists, PLA functions, and equations. These are the types of outputs that the HDL Compiler prepares from HDL descriptions.

■ Synthesis Policy

To work with the HDL compiler, a high-level language must adapt to a series of guidelines. These guidelines, called a synthesis policy, cover the types of statements can be made within the language. In short, the synthesis policy forces the designer to work within an "RTL-level style of design" according to Bob Dahlberg, Synopsys' marketing manager. This style is a mixture of structural and functional elements. The structural statements describe the RTL architecture and the combinatorial statements describe the logic functions within that architecture.

"This style is how engineers are already designing today," according to Dahlberg. To illustrate his point, he refers to a one-year-old design from Sun Microsystems Inc. that was expressed in 3,000 lines of Verilog. This design was successfully passed through the HDL compiler, with Synopsys having to modify only one type of construct—a reference to a parameterized cell. The compilation took 30 minutes on a Sun 4/280, and the resultant design required only 83 percent of the area of Sun's manually-implemented design.

Future HDL's will include VHDL which, according to Synopsys, is an order of magnitude more complex, but not an order of magnitude better, than Verilog. In evaluating which HDLs deserve ports, Synopsys is looking for simulators that produce quality results, support mixed-level simulation, and are being used by a substantial number of ASIC designers. The second criterion underscores the fact that both the high-level description and synthesized circuits should be verified by the same tool to ensure consistency.

Synthesis of a design from an HDL description remains in the province of IC design, since the designs created by available synthesis tools are unlikely to be implemented with standard VLSI-level parts. Instead, the designer will start with complex parts (such as microprocessors) and create the "glue logic" to complete the system design.

Like the rest of Synopsys tools, the HDL compiler runs within Mentor Graphics Corp. environment. It's written in C, runs under Unix, and runs on Apollo Computer Inc.'s DN3000 and DN4000 workstations and Sun's 3 and 4 series and the Sun 386i. Prices starts at $17,500.

Synopsys Inc.
Mountain View, Calif.
(415) 962-5000

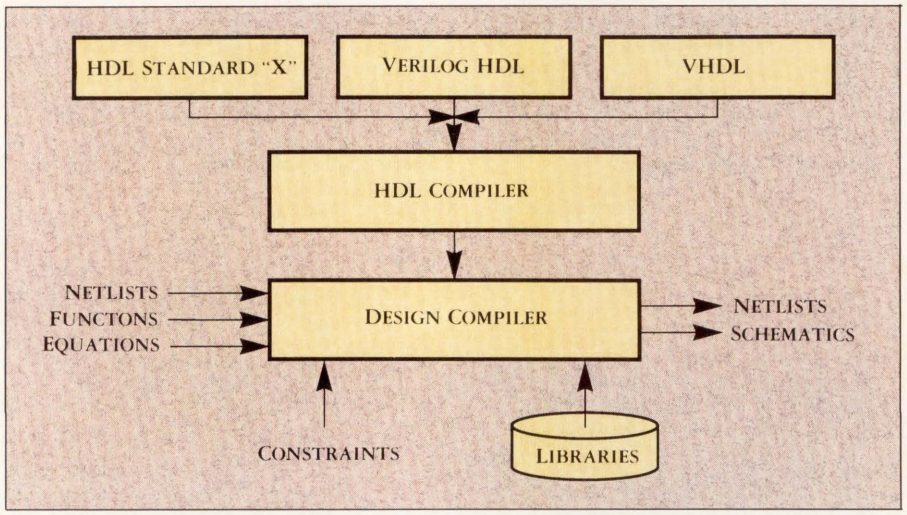

The HDL Design Complier generates input for Synopsys' Design Compiler logic synthesis tool from a high-level description expressed in a high-level hardware description language (HDL). The first HDL supported is Verilog; later, other high-level languages, including VHDL, will also be supported.

A preview of the PCB layout eases the decisions of the systems designer

Greasing the Path to Approved PCB Layouts

TOO often board-level designs go through several iterations as the system and board designers negotiate a compromise between system requirements and space on the PCB board. With many systems tied to specific form factors, such as VME bus- or PC/AT-size cards, physical constraints can have as much impact on system design as throughput specifications.

Most design environments have a distinct interface between the tools and database used in the CAE environment and those in the CAD environment. The former environment consists of simulators, netlists and logical models; the latter contains placement and routing programs, coordinates, and packaged components. Moving between environments is largely a manual process so resolving the conflict between the logic requirements and board-area constraints is costly and time consuming.

Mentor Graphics Corp. attempts to span this impass with its Protoview option to its front-end IDEA Series of CAE systems. More a new concept than a new product, Protoview borrows two modules from Mentor's Board Station PCB design system and puts them in the CAE environment at the control of the system designer. The package module assigns logical entities to physical packages, allowing the engineer to get an accurate estimate of the area needed to lay out the design on a board. The layout module gives the designer automatic and interactive placement capabilities so he can pre-place, finish, and modify the board placement.

■ EASES DESIGN TRANSITION

Focused on the transition from schematics to physical packaging, Protoview lets the designer specify the placement of critical components. He can partition his design and automatically or interactively place components according to design rules, device properties, and connectivity. It provides an esitmate of total area and shows board areas with high congestion. It can place components on both sides of a board. Forward annotation of reference designator, pin number, part number, geometry and board location is automatic.

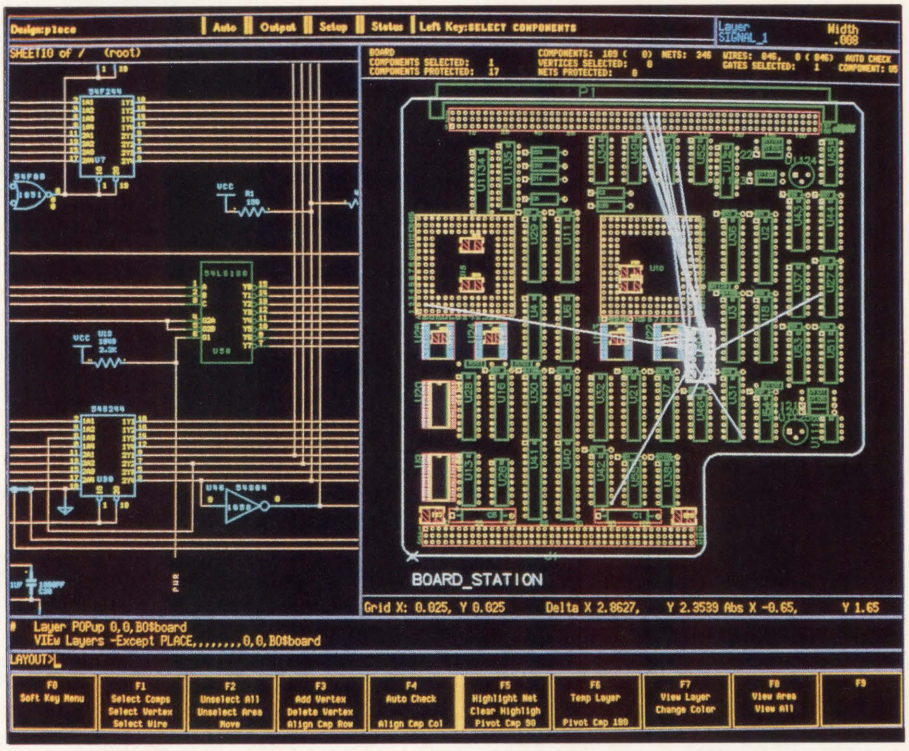

To review a PCB layout after routing, the Protoview option can display a schematic and its corresponding layout simultaneously. A logic element selected on the schematic (green) automatically appears highlighted (white) on the layout, along with all its interconnections.

The split screen capability (see Figure) allows the designer to view schematics simultaneously with placed or completed board designs. A designer can select a device on the schematic and the software will highlight the corresponding physical package, simplifying the modification of the board design.

The designer can also invoke Mentor's other analysis tools, such as thermal analysis, from the Protoview screen. These tools, on the same workstation or across a network, let the designer analyze his placement before committing it to routing on a BoardStation.

In this way, layout and manufacturing concerns can be addressed much earlier in the design cycle, reducing the numnber of iterations. William E. McEachnie, manager of Electronic Packaging Computer Aided Design at Martin Marietta Electronics & Missiles Group, is familiar with Protoview. He believes it will "reduce by about 25 percent the design cycle time" from design conception to an assembled and tested product.

The product rolls out in October, with a special price of $8,000 until December 31, when it jumps to $14,900.

Mentor Graphics Corporation
San Jose, CA
(408) 436-1500

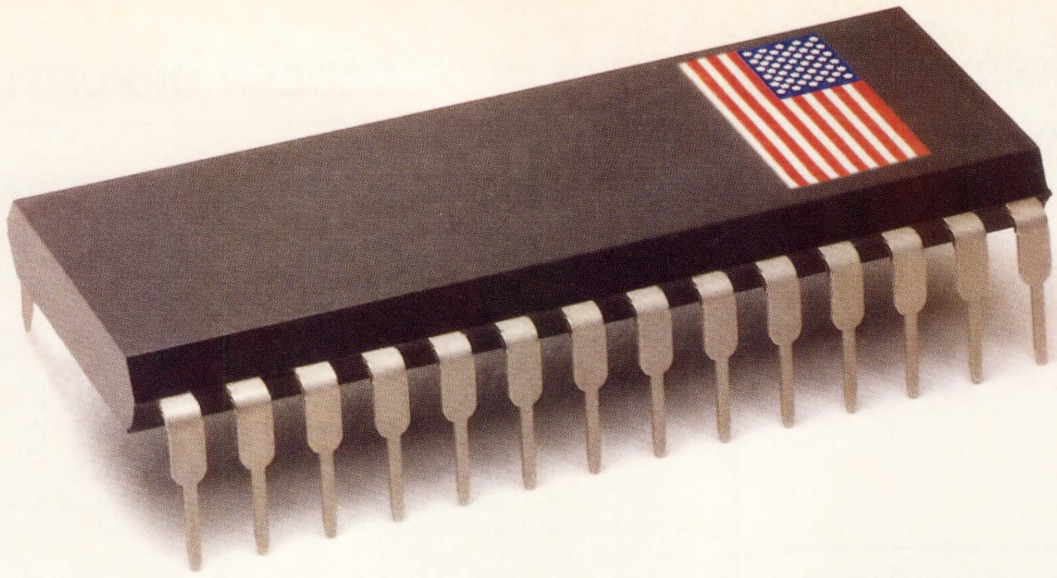

American history?

It doesn't have to be. The American semiconductor chip need not slip into the Smithsonian as merely a faded icon of bygone U.S. ingenuity and technical leadership.

History is what you make it.

And SEMATECH, a new developmental manufacturing collaboration between government and industry, plans to return the U.S. to a preeminent position in semiconductor manufacturing.

The biggest companies in the business are with us, such as IBM, Hewlett-Packard and AT&T.

Now we need you. SEMATECH is hiring the top minds in the industry to join our 5-year, billion-dollar-plus commitment. Here at our advanced laboratories in Austin, Texas, you'll join a skilled staff in the development of new manufacturing techniques.

With your help, we can put the American semiconductor back on the upswing. Before it goes down in history.

Assembly

Openings are available for Manufacturing and Ceramic Engineers.

Test

Several openings exist for Test Engineers to work with IBM and AT&T during technology transfers.

Total Quality Deployment

We're in need of a Vendor QA Manager, an Analytical Chemist, and Engineers/Technicians in the areas of: Failure Analysis - Electrical/Physical Failure Analysis - SEM/EDX - Electron/Ion Beam Analytical Instrumentation - Chemical Analysis & Instrumentation - Gas Analysis - Reliability Engineering & Stress.

Manufacturing Systems Development

Seek individuals qualified to work in the positions and areas of: Manager, System Transfer & Sustaining - Manager, Business Systems - Personnel Computer Support - Engineering Systems Statistician - Senior Analyst, Data Base and Data Architecture - Analyst, Specification & Control - Network Services Specialist - Manager, Automation & Robotics - Case and End-user Tools Specialist - Computer Operations Supervisor.

Strategic Planning Manager and Competitive Analysis Engineer

Requires extensive experience in semiconductor analysis, reverse engineering, planning and setting strategies. Advanced degrees required.

Program Managers

Planning, production control, budget, scheduling, and manufacturing and engineering coordination for program modules of all aspects of semiconductor manufacturing.

Technicians

Lithography Development Techs to assist engineers installing and qualifying photolithography equipment. Contamination Techs to help measure purity of process chemicals, gasses, and ultrapure water. Etch Techs to help install and start-up develomental/pilot line and supporting equipment.

Circuit Design Engineers

Will work on manufacturing demonstration vehicle which is primarily concerned with manufacturing interface design problems.

Development Specialists

Seek a variety of individuals with specialized experience in the areas of: Standards & Metrology - Optical/Electrical Properties - Physical Properties - Auger Spectroscopy - X-Ray Photoelectron Spectroscopy - Secondary Ion Mass Spectrometry - Rutherford Backscattering - Scanning Electron Microscopy - Transmission Electron Microscopy.

Facilities

Openings available for: Fab Support Manager - Facilities Planning Manager - Manufacturing Engineer - Facilities Engineer - Contamination Control Engineer - Planning Engineer - Facilities Contamination Control Manager.

SEMATECH offers competitive salaries and an advanced manufacturing environment in the attractive, growing city of Austin. For immediate consideration, send resume to: SEMATECH, 2706 Montopolis Drive, Austin, TX 78741. Equal Opportunity Employer.

SEMATECH

INNOVATION FOR AMERICA'S FUTURE

ADVERTISERS' INDEX

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
ROBERT ANDERSEN, 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
Human Resources/Admin. Dir., R. POCSIK
National Classified Mgr., SANDRA HERMAN
Mktg. Srv. Dir., DORIANNE WALTHER-LEE
Customer Service Manager, GEORGETTE ROSS
Direct Mktg. Services Mgr., JOANNA BRANDI
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

ARRAY FOR BiCMOS!

180 MHz with low power.

It's cause for celebration. AMCC extends its lead as the high performance/low power semi-custom leader with three exciting, new BiCMOS logic arrays that optimize performance where today's designs need it most. In throughput (up to three times faster than 1.5µ CMOS).

Today, system designers look at speed, power and density. For good reasons. As CMOS gate arrays become larger and faster, designers can't meet their critical paths due to fanout and interconnect delay. As Bipolar arrays become larger and faster, power consumption becomes unmanageable. So AMCC designed a BiCMOS logic array family that merges the advantages of CMOS's low power and higher densities with the high speed and drive capability of advanced Bipolar technology. Without the disadvantages of either.

Our new Q14000 BiCMOS arrays fill the speed/power/density gap between Bipolar and CMOS arrays. With high speed. Low power dissipation. And, mixed ECL/TTL I/O compatibility, (something CMOS arrays can't offer).

For more information on our new BiCMOS logic arrays, in the U.S., call toll free (800) 262-8830. In Europe, call AMCC (U.K.) 44-256-468186. Or, contact us about obtaining one of our useful evaluation kits. Applied MicroCircuits Corporation, 6195 Lusk Blvd., San Diego, CA 92121. (619) 450-9333.

Q14000 SERIES

	Q2100B	Q9100B	Q14000B†
Equivalent Gates	2160	9072	13440
Gate Delay* (ns)	.7	.7	.7
Maximum I/O Frequency (MHz)	180	180	180
Utilization	95%	95%	95%
Power Dissipation (W)	1.8	4.0	4.4
I/O	80	160	226
Temperature Range	COM, MIL	COM, MIL	COM, MIL

*(2 loads, 2 mm of metal) †Available soon

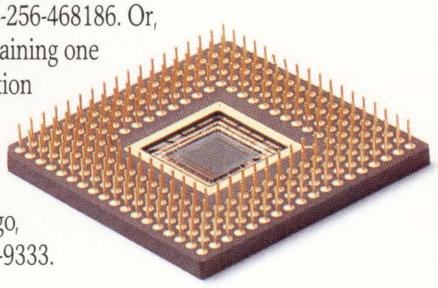

A Better BiCMOS Array is Here.

AMCC

CIRCLE NUMBER 16

How do you justify going with the runner-up in PCB layout?

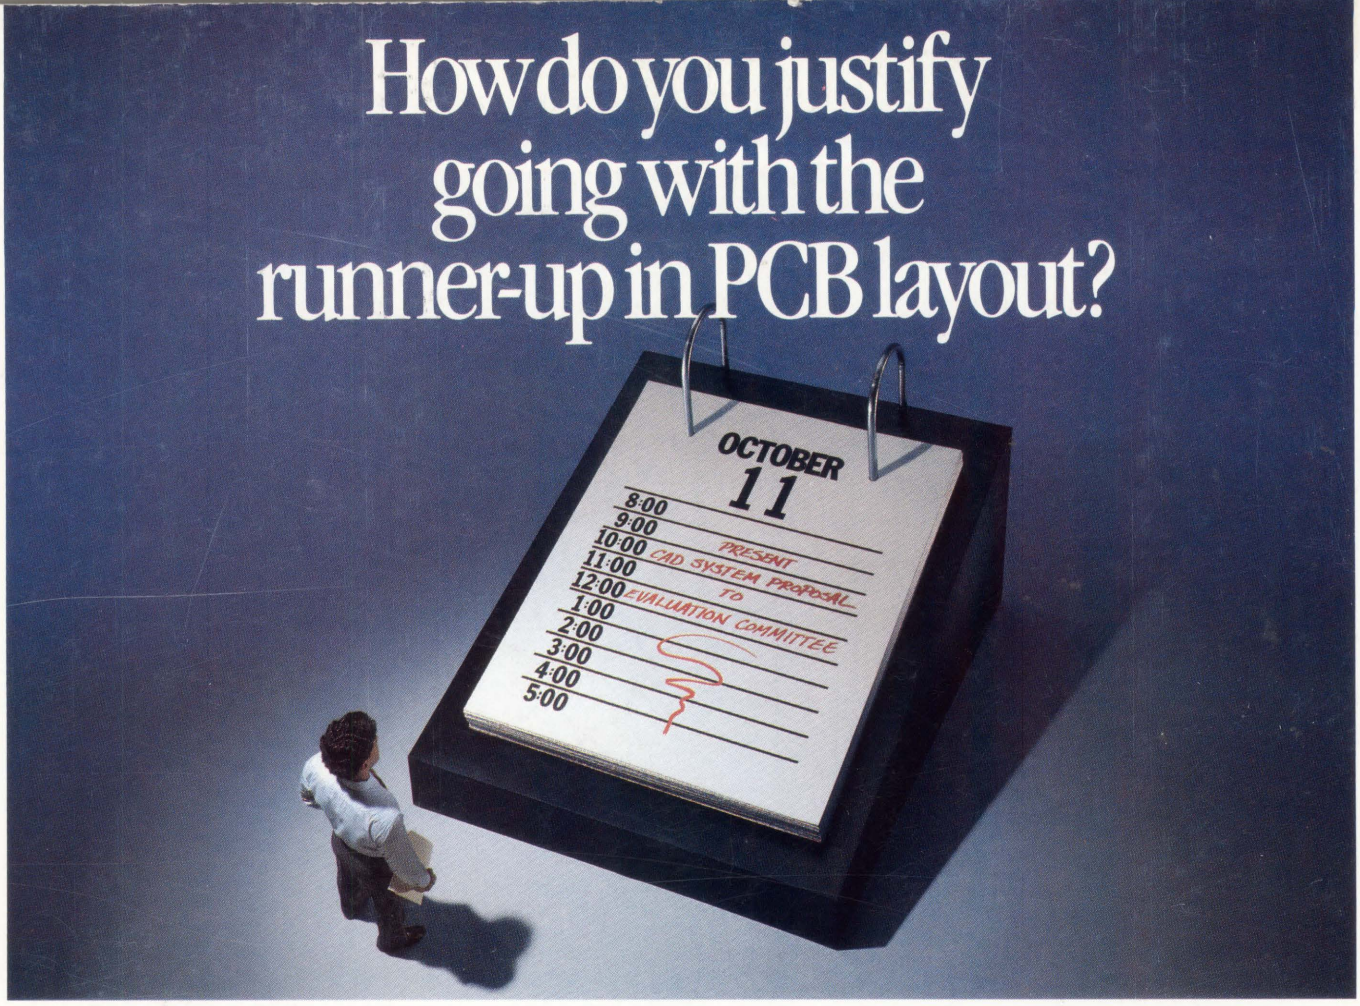

You don't. Especially when Mentor Graphics' Board Station® takes the lead by winning 15 out of 15 recent benchmarks.

One meeting from now, you're going to commit to a new PCB CAD system. Not an easy decision. You need speed, reliability and a guaranteed growth path, all in a single product.

So where to turn? Simple. To the sales and performance leader in PCB layout, Mentor Graphics.

More CAD managers are purchasing Board Station than any other PCB layout system. And with good reason. Board Station's technical performance is second to none. Good enough, in fact, to recently win 15 benchmarks in a row against the former industry leaders[1].

Unequaled PCB design productivity.

How did we achieve a position of leadership in PCB design automation? By reacting promptly to the critical issues facing CAD departments everywhere.

Like the need to mix interactive and automatic operation in a way that boosts quality without sacrificing speed. And the demand for rule-based layout tools which produce correct designs in a single iteration. And the importance of a neutral ASCII database that readily adapts to heterogeneous computing environments. And the added flexibility achieved through macro programming.

A universal solution.

Besides unmatched performance in PCB CAD, Board Station participates directly in the entire Mentor Graphics Electronic Design Automation (EDA) environment. Which means you can refer to on-screen schematics during layout, or back-annotate net lengths to design files, or export layouts to mechanical packaging, or track ECOs through system-wide documentation tools.

To be continued.

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).

[1]Based on total PCB design turnaround.

Mentor Graphics®

Asia Pacific Headquarters, Tokyo, Japan: Phone: 81-3-505-4800, Telex: 2427612 **Middle East, Far East, Asia, South America, Mentor Graphics, International Department, Beaverton, Oregon, U.S.A.:** Phone: 503-626-7000, Telex: 160577 **Mentor Graphics Corporation, European Headquarters, Velizy, France:** Phone: 33-1-39-46-9604, Telex: 696805 **Mentor Graphics Corporation, North American Headquarters, Beaverton, Oregon:** Phone: 503-626-7000, Telex: 160577.

Your ideas. Our experience.