EIGHT DOLLARS A CMP PUBLICATION MAY 1988

VLSI Systems Design

FOR DESIGNERS OF HIGH-PERFORMANCE SYSTEMS

HIGH-PERFORMANCE DESIGNS VAULT TO A NEW ORBIT

DO YOU HAVE WHAT IT TAKES TO MAKE IT BIG IN ASICS?

Not long ago, designing ASICs wasn't even part of your job. Now it's the part everyone's counting on.

That's where Daisy comes in.

Daisy CAE tools are used by more ASIC designers than any other CAE workstations.

Because from schematic creation through post-layout simulation, Daisy has what it takes to keep ASICs on time and on budget.

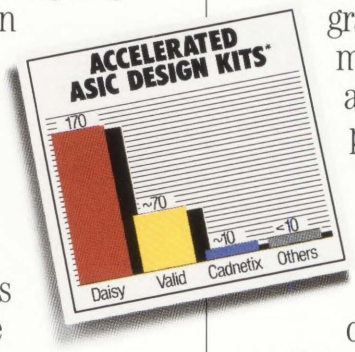

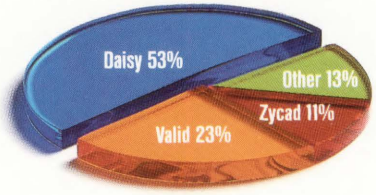

Simulation accelerator market share. Source: Prime Data, 1985 and 1986 unit shipments.

For example, our MegaLOGICIAN™ simulation accelerator is seamlessly integrated with the schematic, so you can locate and correct design problems interactively. That means faster debugging and more time to improve the quality of your design. Plus no other accelerator is as well supported, with more than 170 design kits supplied by 70 different vendors. So you can build productivity instead of libraries.

Which may explain why more MegaLOGICIANs are in use today than all other accelerators combined.

Speaking of combining, you

*Based on minimum 10x performance improvement compared to 32-bit workstations. Source for design kit estimates: VLSI Systems Design's Semicustom Design Guide, 1987.

can share a MegaLOGICIAN with a network of our 386-based desktop workstations, for a high-powered low cost ASIC design environment.

And that's just the beginning.

With our library of more than 4,500 system-level components, you can include your ASIC in complete "real world" system simulations to ensure that your designs will be ready for production, instead of revision.

All of which makes Daisy today's choice for no-sweat ASIC success.

But what about tomorrow?

Gate counts are on the rise. If your tools run out of steam at 5,000 gates, so could your future.

No problem.

Our ASIC design tools glide through 20,000-gate designs without even breathing hard.

In fact, new design kits already support arrays of over 100,000 gates.

So you'll never have to worry about hitting a dead-end.

But don't take our word for it, listen to what Rockwell and other industry leaders have to say. For a free copy of "Making It Big In ASICs" call Daisy at 1 (800) 556-1234, Ext. 32. In California, 1 (800) 441-2345, Ext. 32.

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.

Advanced technologies and tools are lofting high-performance systems to rarefied levels.

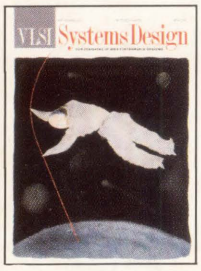

Contents

ARTICLES

86
ASIC costs

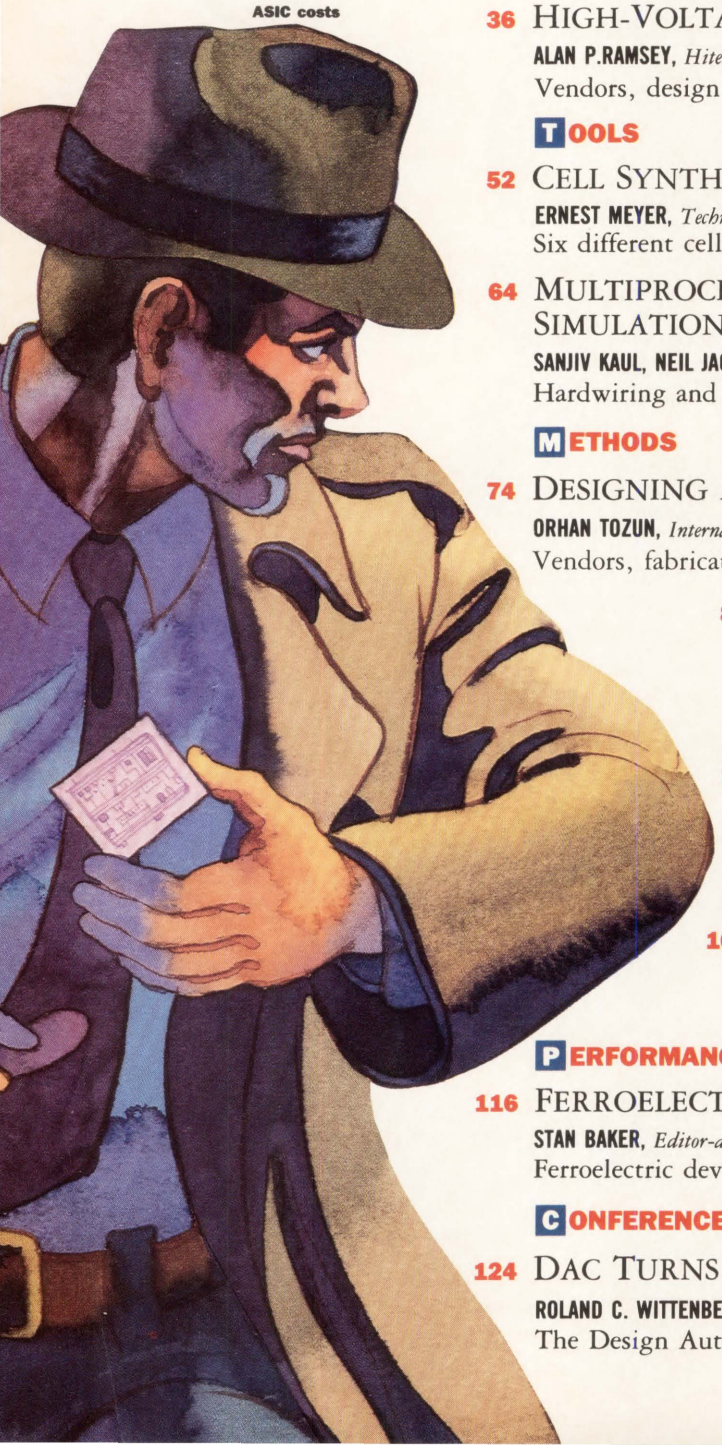

DEPARTMENTS

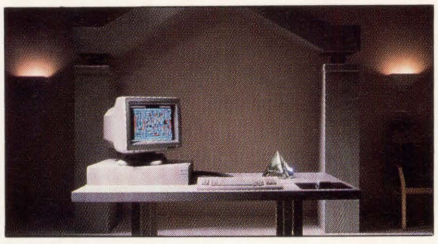

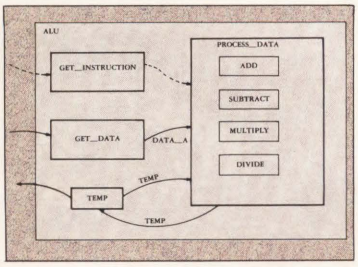

COVER ILLUSTRATION BY ALEXA GRACE

VLSI Systems Design

EDITOR/EDITORIAL DIRECTOR
Girish Mhatre

EXECUTIVE EDITOR
Roland C. Wittenberg

MANAGING EDITOR
Mike Robinson

EDITOR EMERITUS
Roderic Beresford

EDITOR-AT-LARGE
Stan Baker

WESTERN REGIONAL EDITOR
David Smith

DIRECTORIES EDITOR
Michelle A. Losquadro

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

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

EDITORIAL ART
Terry Koppel, Art Director
Rosemary Intrieri, Designer
Marie D'Ippolito, Design Director

MANUFACTURING
David Getlen, Dir. of Manufacturing
Marie Myers, Production Manager
James Pizzo, 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) 365-4600. 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, 600 Community Drive, 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.

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 2

High Performance: A New Orbit

The focus is on technology levers for boosting system performance

This magazine was conceived as a forum for discussion of a technology that had only recently emerged from the laboratory and was finding its way into the hands of a few enterprising engineers. The technology was the semicustom, or application-specific, IC and the motive force behind it was the commercial availability of automated design tools.

The technology was eagerly embraced by a special breed of pioneers driven by the insatiable need to make their systems do "more"—much more, not just a little more. These pioneers shared a special kinship with those who were the first to embrace the microprocessor and, many generations ago, with those who were the first to endorse the integrated circuit. All these pioneers firmly grasped technology developments as levers to boost system performance—as tools for increasing system speed, throughput, functional density, and reliability, and for getting more bang for the buck.

Performance pioneers share a number of traits: Not content with evolution, they prefer discontinuities; they look to leaping into the next orbit. They thrive on being at the vanguard of technology revolutions that roll back performance frontiers.

Today the chips and tools we discussed in the early years seem almost primitive. Some of the concerns seem almost trivial. But to those pioneering souls who explored uncharted territory and in the process staked their professional reputations—and often their careers—this technology was a major gamble.

The semicustom IC continues to evolve on many frontiers, and design tools continue to become more automated. But they are no longer the exclusive purview of a select few.

What then is the next leap for the restless performance pioneer?

We think the next orbit is to view the semicustom IC within the context of the system. In other words, the designer of high-performance systems has to expand his locus of interest to include system concerns, among them architectures, the partitioning of hardware and software, and packaging. He also must be vigilant to the evolution of new technologies. With that in mind, we have recast VLSI SYSTEMS DESIGN.

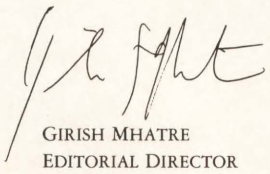

GIRISH MHATRE
EDITORIAL DIRECTOR

There Will Still Be a Few Uses for Conventional ECL ASICs.

Cold facts: now the highest-density ECL logic array runs at a cool 1/10 the gate power of competing devices.

Raytheon's ASIC design expertise and proprietary technology make conventional ECL arrays too hot to handle. The superior performance of the new CGA70E18 and CGA40E12: the ECL logic array family with the highest density and the lowest power requirement now available.

☐ **Superior performance:** 300 pS delay and 300 μW (typical gate) power dissipation deliver the industry's lowest speed-power product: <0.1 pJ. Toggle frequency 1.2 GHz (typical).

☐ **Highest density:**
CGA70E18 — 12540 equivalent gates
CGA40E12 — 8001 equivalent gates

☐ **Lowest power:** Industry's smallest bipolar transistors result in power dissipation that is a fraction of conventional ECL at comparable propagation delays. Typical chip power dissipation of 3W to 5W.

☐ **Et cetera:** Interface TTL, ECL (10K, 10KH, 100K), ETL. Customer access to proven, fully integrated CAD system. Commercial and military operating ranges.

Call Raytheon for access to the right ECL technology. We're not blowing any smoke, and neither should your system's performance.

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

Access to the right technology

Raytheon

CIRCLE NUMBER 3

Calendar

CICC '88

May 16–19, 1988
Rochester Riverside Convention Center/Holiday Inn–Genesee Plaza
Rochester, N.Y.

CICC '88 is sponsored by the IEEE Electron Devices Society and the IEEE Solid-State Circuits Council and cosponsored by the IEEE Rochester Section. Its goal is to bring together designers and manufacturers of circuits and systems and users of custom ICs to discuss new developments and future trends in custom chips. Session topics will include application-specific memories, circuit and logic simulation, drivers and interfaces, layout analysis and generation, analog circuit techniques, silicon compilers, testers and testability, user-programmable logic devices, packaging and systems interconnection, and reliability. For further information, contact Mrs. Roberta Kaspar, Executive Secretary, CICC '88, 20 Ledgewood Drive, Rochester, N.Y. 14615. (716) 865-7164. ■

15th Annual Symposium on Computer Architecture

May 30–June 2, 1988
Ilikai Hotel
Honolulu, Hawaii

This annual symposium is sponsored by the Computer Society of the IEEE and the Association for Computing Machinery. It will feature presentations that include language-oriented architectures, distributed and parallel architectures, memory systems, performance evaluation and measurement, advanced devices, architectures for transaction-based systems, interconnection networks, the impact of VLSI on architecture, novel computing techniques, operating-systems–oriented architectures, and tools and methods for architecture design and description. For details, contact H.J. Siegel, General Chair, Supercomputing Research Center, 4380 Forbes Blvd., Lanham, Md. 20706. ■

1988 IEEE International Symposium on Circuits and Systems

June 7–9, 1988
Helsinki University of Technology
Espoo, Finland

This international symposium will cover all aspects of the theory, design, and applications of circuits and systems. Session topics will include power electronics and circuits, VLSI design and applications, large-scale networks, computer-aided design, modeling and simulation, digital signal processing, switched-capacitor networks, communication circuits, graph theory, fault analysis, circuit layout, filter theory, active and digital filters, distributed networks, and computer and microwave networks. For additional information, contact Prof. Yrjo Neuvo, General Chairman, Tampere University of Technology, P.O. Box 527, SF-33101 Tampere, Finland. (358) 31-162698. ■

Design Automation Conference '88

June 12–15, 1988
Anaheim Convention Center
Anaheim, Calif.

DAC '88, sponsored by the IEEE Computer Society and the Association for Computing Machinery, is devoted solely to the field of design automation. This year's conference will offer tutorials, panel discussions, and technical presentations. General session topics will include design for testability, VHDL in use, floorplanning and area estimation, timing verification, parallel simulation, channel and global routing, engineering information databases, high-level synthesis, application-specific simulation, placement algorithms, layout compaction, register-transfer-level synthesis, logic synthesis and optimization, data structures for integrated design systems, physical design verification, ideas in circuit verification and simulaton, ideas in system generation, ideas in testing, fault simulation, and micro-architecture synthesis. For additional information, contact Pat Pistilli, MP Associates Inc., 7490 Clubhouse Road, Suite 102, Boulder, Colo. 80301. (303) 530-4333. ■

International Workshop on VLSI for Artificial Intelligence

July 20–22, 1988
University of Oxford
Oxford, England

This workshop will provide a forum where AI experts and system and VLSI designers can come together to discuss trends in AI applications and their computational requirements, VLSI implementations, and computer architectures. Topics to be discussed will include alternative technologies, functional-language architectures, knowledge-oriented machines, logic-oriented archictectures, rule-based engines, and fifth-generation computers. For further information about the workshop, contact Dr. Jose G. Delgado-Frias or Dr. Will R. Moore, Dept. of Engineering Science, University of Oxford, Parks Road, Oxford OX1 3PJ, England, U.K. Phone: (0865) 273188. ■

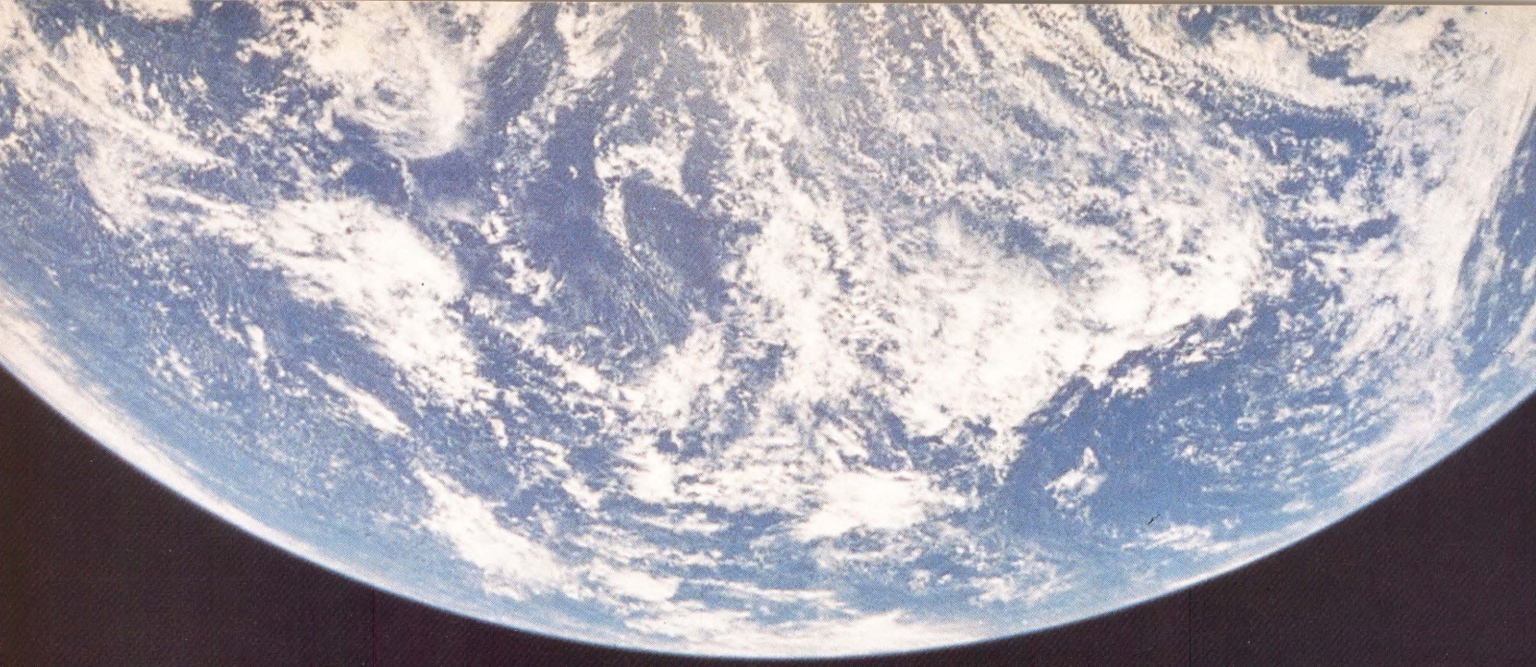

At 100 MHz, the IMS Logic Master XL creates a new world standard for ASIC verification.

But speed isn't everything.

You might think it sufficient that we're introducing the highest performance ASIC verification system available anywhere.

And that our new Logic Master XL is the first system with data rates of 100 MHz.

But we didn't think that was enough.

Instead, we concentrated on the kind of features that turn superior performance into superb productivity.

Like advanced driver electronics which produce symmetrical waveform edges for measurement accuracy.

And complete per-pin control. So you can handle complex timing configurations with ease. And mixed technology devices.

Plus automatic wiring for rapid, error-free setup, and superb system skew for greater measurement confidence.

Best of all, the Logic Master XL comes from Integrated Measurement Systems, the international leader in ASIC verification.

Contact us for more information.

IMS

Integrated Measurement Systems, Inc.
Corporate Headquarters
9525 S.W. Gemini Drive
Beaverton, Oregon 97005
TEL 503 626-7117

CIRCLE NUMBER 4

Photo courtesy of NASA

U.S. Regional Offices: Boston, MA 617-366-3676; Irvine, CA 714-768-5808; Portland, OR 503-626-7117. **International Offices:** Belgium 02-219-2453; Canada 416-890-2010; France 01-4534-7535; Hong Kong 5-701181; Japan 0423-33-8111; Sweden 753-38465; Switzerland 830-5040; Taiwan 02-501-0055; The Netherlands 01-40-582911; United Kingdom 06-284-76741; West Germany 5722-2030.

BIT Stream

AT&T Appointment

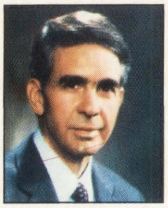

ROBERT ALLEN

ROBERT E. ALLEN was elected chairman of AT&T following the death of AT&T's former chairman James E. Olson. Allen was previously president and chief executive officer. But even though Olson had only been in the chairman's slot for 20 months, Allen will have a big pair of shoes to fill. Olson, as the first and only chairman of the "new" AT&T, guided the company out of its chaotic breakup to its present position as a formidable competitor in the semiconductor, fiber optics, computer, and telecommunication industries. The board awaits Allen's recommendations for the president and CEO slots. ■

AMD Unveils Universal 64-Bit FPU

Advanced Micro Devices Inc. (Sunnyvale, Calif.) rolled out its latest floating-point processor, the Am-29C327, a double-precision CMOS processor that supports all of the industry-standard floating-point formats: IEEE-754 standard; DEC's D, F, and G formats; and IBM's standard. The chip can operate in a flow-through or a pipelined mode and delivers double-precision accuracy at a 10-MHz rate. It can execute over 70 floating-point instructions and can perform both arithmetic and logic operations on 32- and 64-bit integers. The processor is priced at $595 in 100-unit quantities and comes in a 169-lead PGA. ■

Sun Shines on AT&T's Unix

Open Look is the name that AT&T and Sun Microsystems Inc. (Mountain View, Calif.) have given to the new user-friendly interface for the Unix operating system. Sun's president, Scott McNealy, and AT&T Data Systems Group president, Vittorio Cassoni, are very optimistic about the future of Open Look. The interface was developed by Sun for AT&T and is licensed from Xerox Corp., which developed the concept of a heuristic, graphical, menu-driven interface at its Palo Alto Research Center approximately 20 years ago. It is the same technology that Apple Computer adopted for its Macintosh line of user-friendly computers. Open Look will operate under X Window and NeWS. ■

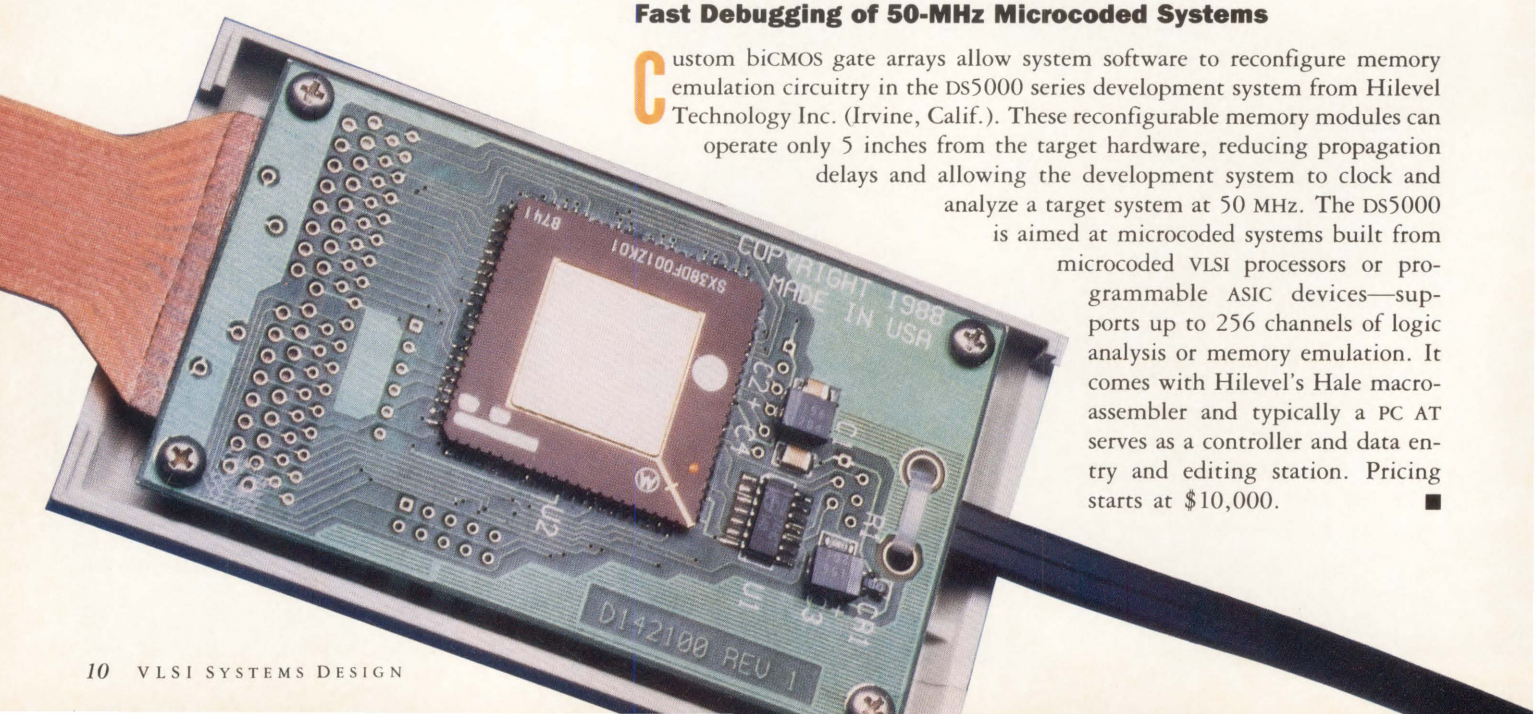

Fast Debugging of 50-MHz Microcoded Systems

Custom biCMOS gate arrays allow system software to reconfigure memory emulation circuitry in the DS5000 series development system from Hilevel Technology Inc. (Irvine, Calif.). These reconfigurable memory modules can operate only 5 inches from the target hardware, reducing propagation delays and allowing the development system to clock and analyze a target system at 50 MHz. The DS5000 is aimed at microcoded systems built from microcoded VLSI processors or programmable ASIC devices—supports up to 256 channels of logic analysis or memory emulation. It comes with Hilevel's Hale macro-assembler and typically a PC AT serves as a controller and data entry and editing station. Pricing starts at $10,000. ■

14 MIPS from NMOS Chips

At the heart of Hewlett-Packard Co.'s two newest midrange computers in the company's HP 9000 Series 800 family is a single-chip NMOS III microprocessor that drives the two-board computers at 14 times the integer performance of a Digital Equipment Corp. VAX-11/780. The Models 835S and 835SE provide a standard 66.7-ns instruction cycle time and a 128-Kbyte cache memory. The 835S can support up to 30 users, and the 835SE can handle as many as 78 users.

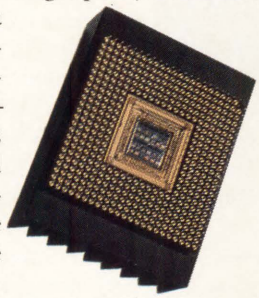

The 32-bit RISC-based HP Precision Architecture includes 48-bit virtual addressing, which allows it to support up to 112 megabytes of RAM and 16 disk drives, storing 9.1 gigabytes of data.

Both models feature a new floating-point coprocessor that delivers 2.02 double-precision MFLOPS when measured with a Linpack benchmark. Hewlett-Packard says that the 835S and 835SE provide 75% more MFLOPS than the Sun-4/260 at a comparable price. The new platforms are object-code–compatible with the other members of the Series 800 and source-code–compatible with the Series 300. An 835 configuration starts at $45,000. ■

DSPs Cut High-Speed Modem Costs

Although the use of higher levels of integration has stemmed the rising costs of many of today's designs based on complex VLSI chips, there are still cases where multiple chips can be the economical best bet. Texas Instruments Inc. has made a study of the total component costs for modems operating at data rates of 300 to 19,200 bits/s. The study indicated that as the rates pass 1,200 bits/s, the cost of dedicated modem ICs increases almost exponentially, while a combination of a DSP and an analog interface chip increases only linearly with data rates—offering a more economical solution. The TI study used a DSP chip together with its TCM29C18 single-chip analog interface for the 1,200–to–2,400-bits/s range. The interface chip includes companding A/D and D/A converters, together with input and output filters. The converters have 8-bit resolution, but as a result of the companding transfer characteristics of these devices, the dynamic range is extended to 12 bits. ■

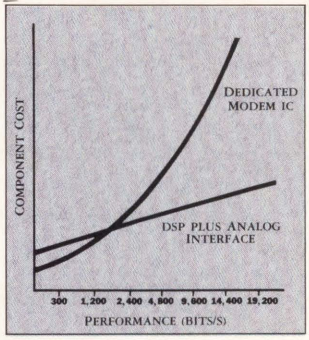

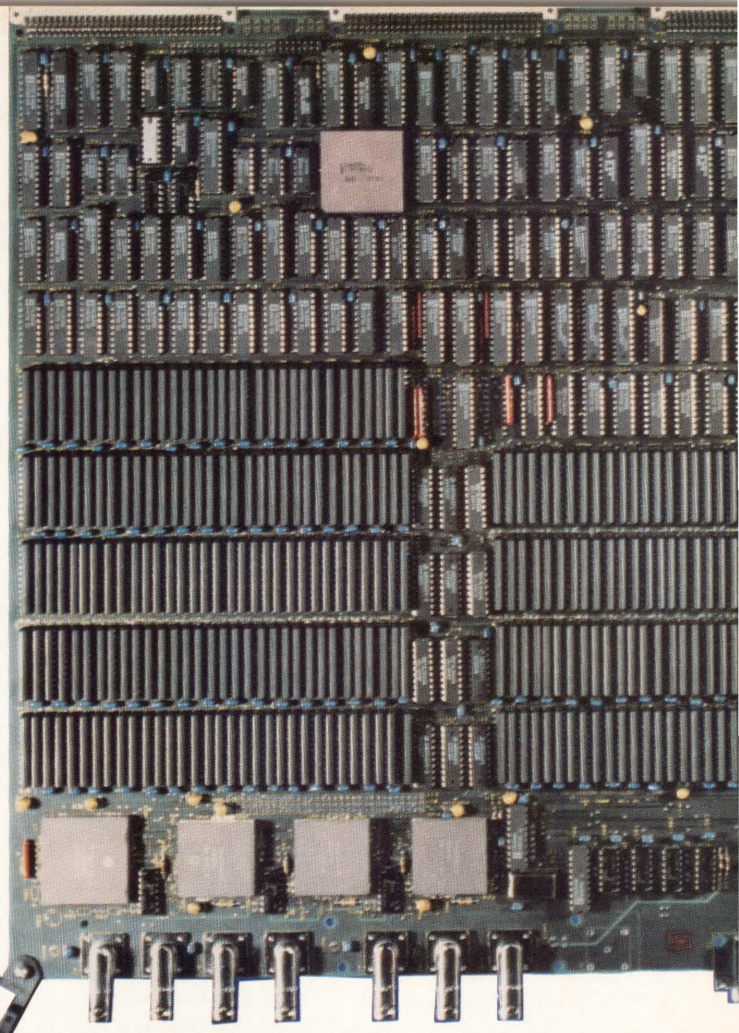

All the Colors under the Sun

Univision Technologies Inc. (Burlington, Mass.) hopes to unleash all the colors of the sun, or more literally, Sun Microsystems' Sun-3 family of workstations, with its recently unveiled UDC-3400 series of high-resolution color display controllers that plug directly into a Sun standard VMEbus slot. The plug-in boards can deliver color image displays having a resolution of up to 1,280 × 1,024 pixels at 34 bits per pixel. They are priced from $3,995 to $8,995. ■

Wilf Speaks Out on RISC Chips

Commenting on Motorola's recent unveiling of its 88000 microprocessor chip set, Wilfred J. Corrigan, chairman of LSI Logic Corp., said, "With Motorola finally in the RISC fray, system designers now have a clear vision of all the significant architectural approaches open to them in the RISC arena. We believe that most of the major designs will be selected from a three-entree menu: MIPS Computer's R2000 and R3000, Sun's SPARC, and the 88000. It's going to be a battle between these three, and later RISC approaches will not be much of a factor." ■

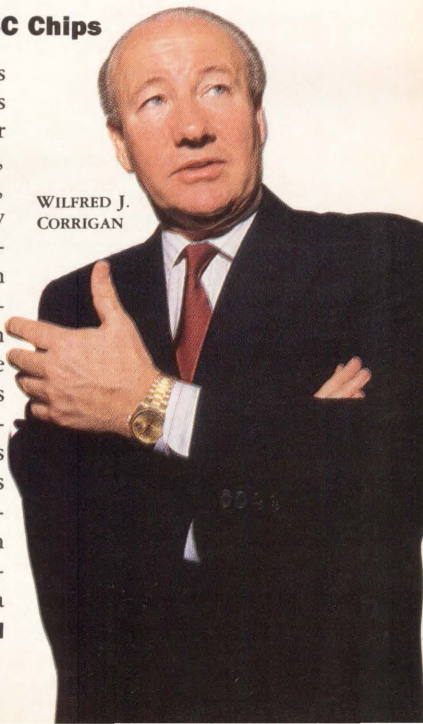

WILFRED J. CORRIGAN

News Analysis

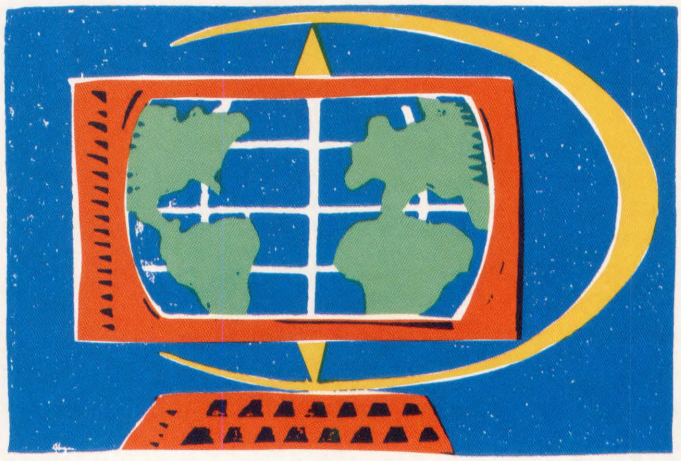

The Other Paths of RISC

New or improved RISC microprocessors have been much in the news recently, and flowing from that core of events are the many announcements by companies joining hands to cooperatively develop bit-level operating-system interfaces, design tools, compilers, and hardware. However, almost all these newsy developments are focused on putting RISC technologies into central processors, the general-purpose high-profile darlings of the computer industry.

Nevertheless, there are also other important routes that RISC is already taking, paths of development that will involve many more engineers and types of systems than will the CPU push. One that has become active is embedded RISC processing. Embedded RISC microcontrollers are beginning to appear, and next will come application-specific systems on a chip built from cores in standard-cell libraries and silicon compilers.

RISC has become a mainstream movement to improve computing power and reliability and keep costs and design times down. It is a whole-systems movement, not just a hardware trend. And systems aimed at different applications have very different needs in terms hardware, software, and support.

The primary difference between RISC processors made for CPUs and those made for embedded processing and control is that the former are optimized for performance; whereas the latter are optimized for specific application targets, in speed, function, and cost. Embedded intelligence here must be much cheaper than in CPUs.

The computer systems companies that have been the most successful in bringing leading-edge performance to workstations are those driving the acceptance and growth of general-purpose RISC computing. Their system environments are becoming dominant. In this arena, the semiconductor companies are not in the driver's seat, although Motorola is trying to be a force with its new 88000 RISC family (see the article, p. 24, and Product Showcase, p. 128).

But highly important application areas for RISC approaches exist outside the general-purpose computing arena—and there the semiconductor companies can be a primary force. In such categories as military, industrial, and consumer, the applications and customers are more fragmented and require different methods of bringing IC technology to their systems. These are the areas in need of strong support by IC producers, and that is where new RISC products are aimed by Texas Instruments, Advanced Micro Devices, Intel, VLSI Technology, and now Sanyo.

Embedded processors are like CPUs in that they use large programs, compilers, large chunks of memory, and often have complex memory hierarchies. They mirror CPUs but are hidden from users and programmed for specific uses for their lifetimes. In contrast, 32-bit RISC microcontrollers are an extension of the traditional microcontroller market that is still dominated by 4- and 8-bit processors in low-cost consumer items. They are I/O-intensive and event-driven, manipulate lots of data at the bit level, and carry much of their memory on chip.

The users of embedded processors come from a computer development environment and are used to software tools. The controller engineers come primarily from industrial backgrounds, are familiar with hardware, and are often optimizing assembly language code.

Embedded processing and control are also important targets for the 32-bit CISC processors of Motorola, National Semiconductor, Zilog, and others. These CISC devices have sudden competition from the RISC products and will get more. Motorola will try to relieve some of the pressure put on by other companies' RISC offerings by aiming its 88000 at higher-end embedded opportunities.

AMD last month introduced its 29000 microprocessor chip set, which has found its earliest design acceptance in controlling graphics and high-speed networks such as those using the Fiber Distributed Data Interface (FDDI). A few weeks earlier, Intel introduced its 80960 architecture and the first two processors in that family (see Product Showcase, p. 128), although purists won't accept them as RISC units because they use some microcode. A primary target for Intel is automotive applications. Both the 29000 and the 80960 are general-purpose architectures and hardware sets for embedded intelligence, with AMD's aimed at the highest-performance end and Intel balancing its features at the midrange of performance.

VLSI Technology just last month enlisted the partnership of Sanyo to develop versions of its Acorn RISC processor for single-chip microcomputer applications. Its chip is by far the smallest on the market, allowing it soon to be a core processor in a library, available for custom designs done by users. VLSI Technology's product now has lower performance than Intel's, but it will sell at volume prices of about $50.

Many of the traditional embedded microcontroller de-

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

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

FREE. ASIC Estimating Kit.

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

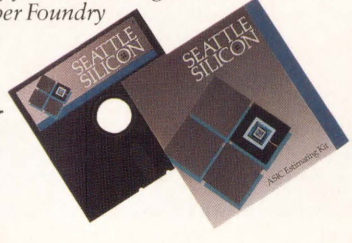

THE SUPER FOUNDRY

SEATTLE SILICON

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

CIRCLE NUMBER 5

signers are avidly searching for more performance. Their more complex systems have been in industrial applications with high prices and low volumes. They need more and more performance but expect it without increased costs, by leveraging the rapidly growing price/performance ratio of dense CMOS ICs. RISC offers the hardware simplicity for such economy.

The high-volume controller applications have been in very low cost consumer items, but these and consumer electronics products in general are demanding more and more computing power. Here, continuing low costs are critical to achieve high volume. Hence, the interest in RISC hardware.

Good examples of the high-performance, low-cost trend at the 32-bit level are application-specific graphics and network microprocessor control chips that Texas Instruments has pioneered, starting before many engineers had heard of RISC. The company is now explaining that its speed-enhancing techniques are the same as found in the new wave of RISC processors.

TI came in the back door of the RISC trend, with embedded microprocessors. Other semiconductor makers have found the RISC path open to them is also embedded, but they do not take TI's application-specific processor approach. They are challenging TI to a general-purpose versus application-specific control battle. Meanwhile. VLSI Technology and Sanyo will try to get around the fray and score with single-chip RISC control devices. ■

Scaling Down Analog-Digital Cell Libraries

When a technology shrink is announced for a standard-cell library, designers can make certain assumptions about the digital cells in that library. The cells will become smaller, roughly by the square of the ratio of the new design rules to the old ones. The cells' operating frequencies will rise, roughly by the ratio of the new and old design rules. Of course, density and frequency depend on other factors besides effective channel length, but such first-order estimates generally hold true.

When a library with analog cells migrates to a new process, however, such density and performance scaling cannot be assumed. Analog circuits are more sensitive to layout, parasitics, and the operating characteristics of their transistors. Although a smaller transistor is faster in digital or analog cells, it is also more susceptible to noise, and its transfer characteristics may be more difficult to control. As a result, designers of cell-based ICs that have both analog and digital circuitry must examine the new operating characteristics of the cells they use.

Recent announcements of products from International Microelectronic Products Inc. (San Jose, Calif.) are revealing some of the trends defining the design of cell-based semicustom ICs that have both analog and digital circuitry. IMP now offers digital and analog cells in 3-μm, 2-μm, and most recently, 1.2-μm technologies. As the company sees it, the spectrum of mixed analog-digital applications for semicustom chips needs a variety of processes (see the figure). The 1.2-μm process yields the highest-performance digital cells and those analog cells with the highest bandwidth, but high-resolution analog circuitry needs the 3-μm analog cells because their operating characteristics are more easily controlled.

IMP's new analog and digital cell libraries, ACL 1.2 and DCL 1.2, respectively, enable designers to build chips with analog and digital functions. The digital cells can give the designers twice the functional density of designs implemented in 2-μm technology. The smaller design rules also give the digital section a 40-MHz clock frequency, up from 25 MHz for 2-μm designs. These results reflect the direct effects of scaling down digital circuits.

The analog cells in the 1.2-μm process have a wider bandwidth than those in the 2-μm libraries. The unity-gain bandwidth rises from 175 to 250 MHz, enabling the technology to implement a 120-MHz video amplifier, for example. On the other hand, smaller design rules constrict the cells' operating voltages. The 1.2-μm cells will operate off 0 to 5 V, whereas their 3-μm counterparts can accommodate signals between −5 and +5 V. The wider signal range of the larger geometries makes them better for applications requiring high precision, such as biomedical systems.

NCR Corp. (Fort Collins, Colo.) has released details on its 2-μm and 1.5-μm analog cell families that show how analog performance may not keep pace with digital as technology sizes shrink. Some of their shared characteristics are characterization over the military temperature range, a power-down mode in some cells, operation down to 3 V, 0-V common-mode range, and kit parts for every cell for breadboarding. Both libraries include

Continued on page 129

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

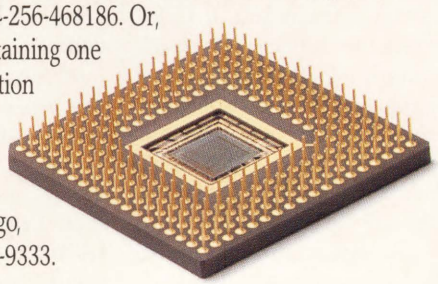

A Better BiCMOS Array is Here.

AMCC

CIRCLE NUMBER 7

Design Automation Rates High Marks at HP

HP's Parzybok Keeps His Balance

BALANCE HAS always been a key ingredient in Bill Parzybok's professional life, and in his private life as well. Couple that with a man who has a tremendous amount of drive and enthusiasm for his job and it helps to explain the fast track record that Parzybok has had at Hewlett-Packard, where he is now a corporate vice president and general manager of the Engineering and Manufacturing Systems Group.

His college days were a harbinger of the future. While keeping a busy schedule as a full-time electrical engineering student at Colorado State University, he was also president of both Sigma Chi and the local chapter of the IEEE, in addition to playing on the varsity soccer team. To keep his balance and accomplish all the goals he had set for himself, he recalled, "I had to learn how to manage my time better. When you're forced to prioritize your time, you tend to be more efficient."

Parzybok earned his bachelor's in electrical engineering and a master's in management at Colorado State. In between, he spent the summer of 1967 working at HP "doing circuit simulation on an IBM computer." He joined HP full time in 1968, after completing his master's.

Although he was interested in math and science and loved to build electronic kits and experiments as a teen-ager, at HP he "got into marketing fairly early on with some important new products." But he didn't abandon his engineering roots when he started as a customer service engineer at HP's Loveland, Colo., operation.

"Actually, I moved my desk to the lab and sat next to the people who invented HPIB. I worked with a guy named Jerry Nelson, and we built frequency synthesizers and network analyzers," Parzybok said. "I enjoyed the circuit design, but I really got my kicks working with customers and showing them how to use HP instrumentation to solve some of their real problems."

His enthusiasm and drive have served him well. Just five years after joining HP, he became marketing manager of the Loveland Instrument Division. Three years later he became general manager of the Electronic Measurements Group and by 1981 Parzybok had most of the instrument division under his wing.

HP was founded almost 50 years ago in the test and measurement business. But with the birth of the electronic design automation industry, many of the engineer's tools began migrating from the laboratory workbench to the engineering workstation.

As a result, HP formed the Design Systems Group in 1984 to lead the company into design automation. Parzybok was chosen as the group's first general manager and is now vice president of the group, which has been reorganized as the Engineering and Manufacturing Systems Group.

'I THINK TO BE A REAL CONTRIBUTOR

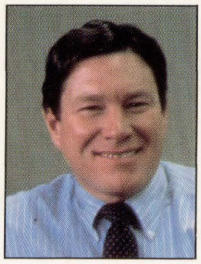

YOU NEED SOME SORT OF BALANCE'

Although design automation is a fast-growing, multibillion-dollar business, it's not an easy market to penetrate. While many competitors floundered, Parzybok moved his group from essentially nowhere in 1984 to one of the top 10 leaders last year.

Being responsible for more than 2,200 employees and multiple operations located throughout the world keeps Parzybok on the move. How does he feel about this heavy work schedule? "My professional life is very important to me," he said, "and it takes the majority of my time. But my personal life is also very important to me, and a job that's all-consuming is not healthy. I think to be a real contributor you need some sort of a balance."

One place that he balances his busy professional life is in the nearby Rocky Mountains. Whenever he gets a chance, Parzybok likes to relax by mountain climbing in the summer or skiing in the winter. He's also involved with woodworking and photography.

But just as with his business career, Parzybok takes his leisure-time activities very seriously. He's already climbed 30 of the more than 50 peaks in Colorado that are 14,000 feet or higher.

"I think you can get consumed by one part of your life or other, and that if you're not balanced, you end up making bad decisions," he continued. Judging by the results, it appears that Bill Parzybok has kept his balance. ■

—ROLAND C. WITTENBERG

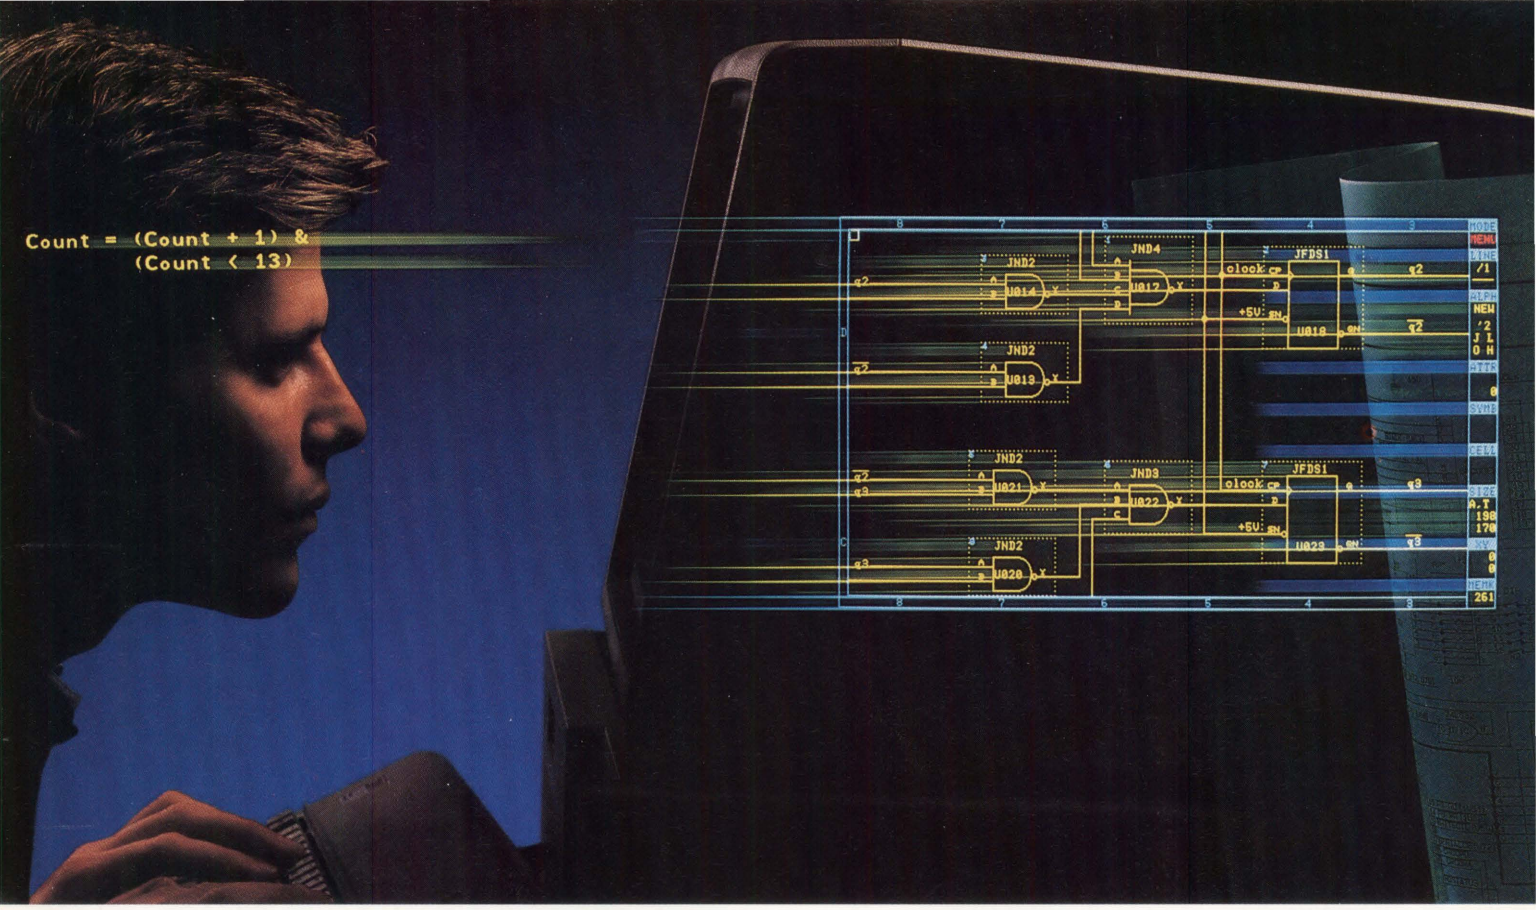

AUTOMATIC SCHEMATIC.

FUTUREDESIGNER: DRAW LESS. DESIGN MORE. Introducing FutureDesigner™ — the only advanced design entry workstation that lets you describe your circuit in compact, high-level terms and create more complex designs faster. FutureDesigner's flexible, new techniques encourage creativity and experimentation, helping you produce innovative products quickly and more accurately.

MULTIPLE DESIGN ENTRY MODES FOR SPEED AND FLEXIBILITY. Describe your circuit with any combination of structural and behavioral representations. Use schematics to enter the structural portions of the design, such as data paths in a memory array. For portions easier to describe behaviorally, like sequencers or decoders, simply enter equations, truth tables or state diagrams using on-screen input forms.

ADVANCED DESIGN VERIFICATION HELPS YOU GET IT RIGHT THE FIRST TIME. For the behavioral portions of your design, use FutureDesigner as a "what if" tool to try different design approaches. Immediately verify that your circuit works as you intended. For the structural portions, design check tools detect and help you correct connectivity and other common design errors. Together these features significantly shorten the design iteration cycle.

LOGIC SYNTHESIS CONVERTS YOUR EQUATIONS INTO SCHEMATICS. Once you've entered equations, state diagrams or truth tables, FutureDesigner's logic synthesizer eliminates redundant circuitry and optimizes your design for size/speed trade-offs. FutureDesigner is the only design entry workstation that will then automatically produce the correct schematics and integrate them with the total structural design.

MORE CHOICES IN TECHNOLOGIES, VENDORS AND SYSTEMS. FutureDesigner is technology independent. Choose the most convenient mix of TTL, PLDs, gate arrays or other ASICs from a wide range of semiconductor manufacturers. You can easily migrate from one technology to another without redesign.

FutureDesigner output is an industry standard, widely accepted by engineering service bureaus and semiconductor vendors. You'll also have access to both FutureNet® and other CAD systems for simulation and PCB layout.

Call us today and learn how a FutureDesigner workstation gives you the flexibility and accuracy to design innovative products faster.

1-800-247-5700
Dept. 123

Data I/O Corporation 10525 Willows Road N.E., P.O. Box 97046, Redmond, WA 98073-9746, U.S.A. (206) 881-6444/Telex 15-2167
FutureNet 9310 Topanga Canyon Boulevard, Chatsworth, CA 91311-7528 (818) 700-0691/Telex 910-494-2681
Data I/O Canada 6725 Airport Road, Suite 302, Mississauga, Ontario L4V 1V2 (416) 678-0761/06968133
Data I/O Europe World Trade Center, Strawinskylaan 633, 1077 XX Amsterdam, The Netherlands (20) 622866/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 Data I/O Company

CIRCLE NUMBER 8

Build your gate array on a solid foundation.

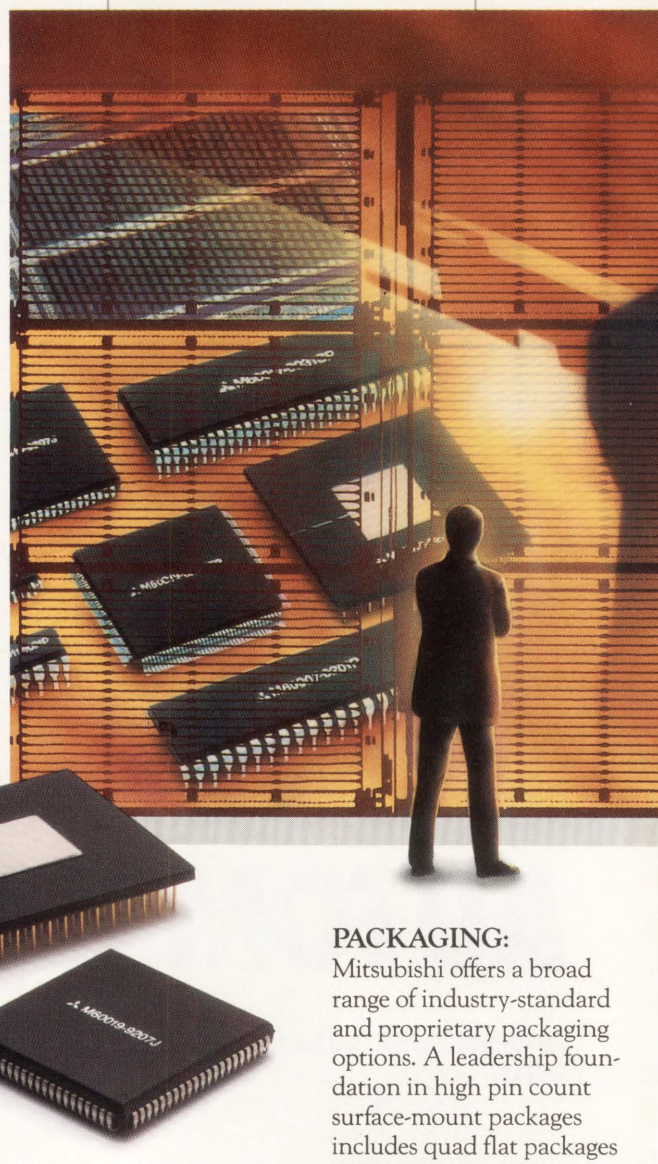

You want to implement your ASIC design quickly. With technology that meets your requirements. You need quality products and dependable delivery to maintain a competitive edge. In short, you expect a partnership built on a solid foundation. One you can count on for future designs as well.

And, that's precisely what you get with Mitsubishi. A solid partnership with a technology base that includes in-house wafer processing, advanced packaging, demonstrated production capabilities and long-term commitment to R&D.

Mitsubishi offers a broad range of gate array product capabilities, from 200 to 20,000 useable gates. And, continuing development work at Mitsubishi's research laboratories is leading to sub-micron CMOS and even more advanced ASIC materials and processes. Unique variable track masterslice (VTM), 1.3μm gate arrays provide an ideal architecture for mixed random logic and memory designs.

There's a broad range of packaging options: DIP and shrink DIP, SOP, PGA, as well as high pin count quad flat packages (QFP) and plastic leaded chip carriers (PLCC).

Mitsubishi is committed to supporting your ASIC needs with CAD/CAE design tools to help you develop designs fast. For maximum flexibility, you can design on Mitsubishi's workstations or your own. Remote access to Mitsubishi's proprietary mainframe CAD system allows design verification on any terminal with telephone access.

Mitsubishi's design centers, located in Sunnyvale, CA and Durham, NC, are networked with Regional Technical Centers and a gate array engineering support staff ready to assist in all phases of design.

Mitsubishi Electronics America, Inc.
The solid foundation for all your ASIC needs.

PACKAGING:
Mitsubishi offers a broad range of industry-standard and proprietary packaging options. A leadership foundation in high pin count surface-mount packages includes quad flat packages (QFP), now available with up to 160 leads. For through-hole technology, Mitsubishi provides a cost-effective alternative to ceramic PGAs by mounting QFPs on a PC board adaptor (MPGA).

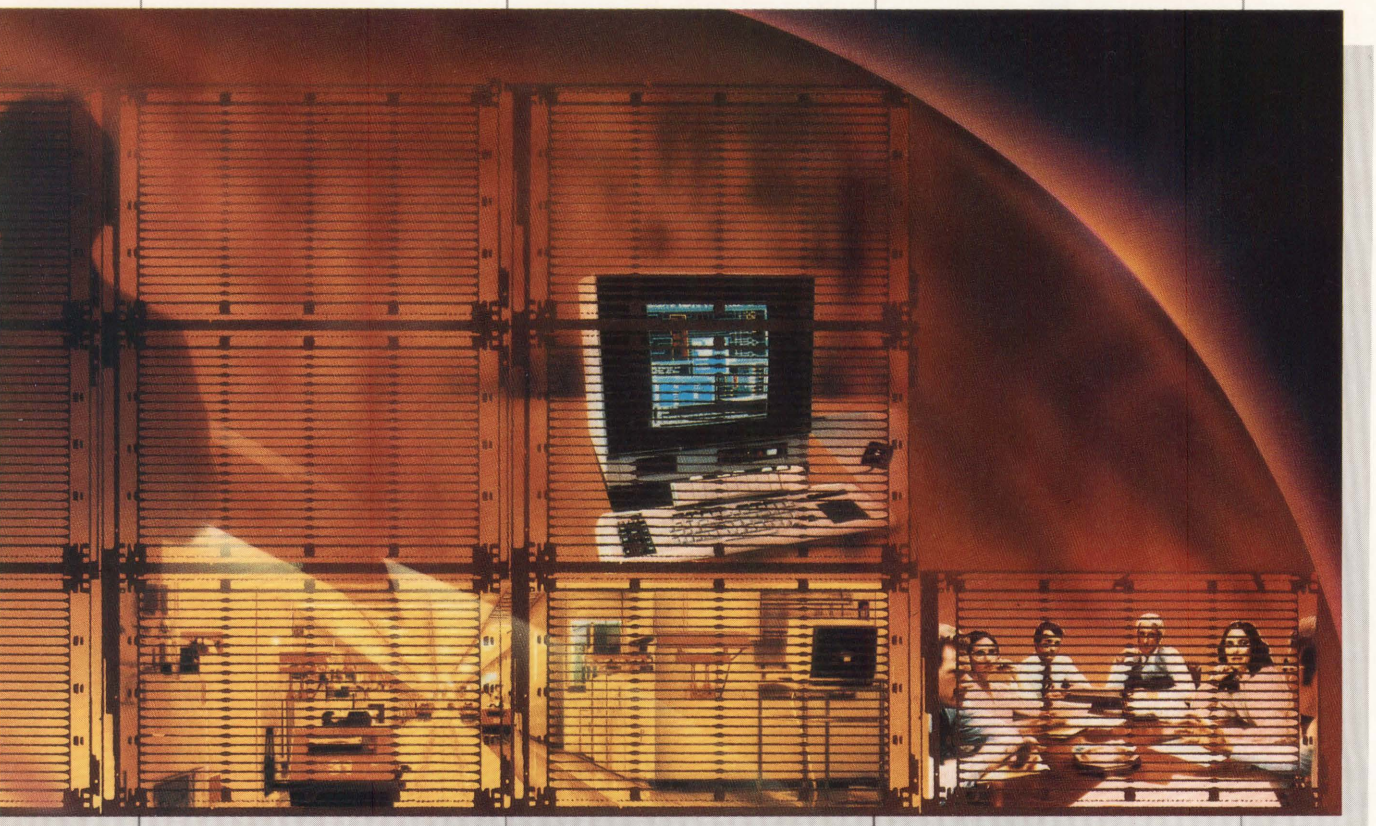

PROCESS:
Low power, 1.3μm, double metal CMOS technology is available today, with submicron levels on the way. Mitsubishi's patented* gate isolation structure provides 10% to 20% faster performance, with a 15% to 25% higher gate density than conventional oxide isolation.
*U.S. Patent No. 4,562,453

MANUFACTURING COMMITMENT:
Mitsubishi is committed to continued advancements in ASIC technology, quality and production capacity. Typical time from design to prototype is three to five weeks. Prototype approval to production is eight to ten weeks.

DESIGN SUPPORT:
At your workstation or Mitsubishi Design Centers, use industry-standard, advanced CAE tools with Mitsubishi's mainframe CAD system. Or, let Mitsubishi integrate your design.

SERVICE:
Mitsubishi offers design and applications support at centers in Sunnyvale, CA; Durham, NC and Regional Technical Centers.

Quality through commitment. **MITSUBISHI ELECTRONICS**

For further information, call or write Mitsubishi Electronics America, Inc., Semiconductor Division, 1050 E. Arques Avenue, Sunnyvale, CA 94086, (408) 730-5900.

U.S./CANADIAN REGIONAL SALES OFFICES AND TECHNICAL CENTERS:
NW: Sunnyvale, CA (408) 730-5900, **SW:** Torrance, CA (213) 515-3993, **N:** Minnetonka, MN (612) 938-7779, **NC:** Mt. Prospect, IL (312) 298-9223, **SC:** Carrollton, TX (214) 484-1919, **NE:** Woburn, MA (617) 938-1220, **MA:** Hackensack, NJ (201) 488-1001, **SA:** Norcross, GA (404) 662-0813, **SE:** Boca Raton, FL (305) 487-7747, **Canada:** St. Laurent, Quebec (514) 337-6046.

U.S./CANADIAN AUTHORIZED SALES REPS: AL: Beacon Elect. (205) 881-5031, **AR:** OM Assoc. (214) 690-6746, **AZ:** SMS & Assoc. (602) 998-0831, **CA:** Pathfinder Elect. (619) 578-2555, QCI (408) 432-1070, SC Cubed (805) 496-7307, SC Cubed (714) 731-9206, **CO:** Simpson Assoc. (303) 794-8381, **CT:** Comp Rep Assoc. (203) 269-1145, **DE:** Trinkle Sales (609) 795-4200, **FL:** Beacon Elect. (305) 997-5740, (813) 796-2378, (305) 332-1940, **GA:** Beacon Elect. (404) 256-9640, Beacon Adv. Comp. (404) 662-8190, **IA:** Mid-Tec Assoc. (314) 275-8666, **ID:** ES/Chase (503) 292-8840, **IL:** Phase II Mktg. (312) 303-5092, **IN:** Carter, McCormick & Pierce (317) 244-1685, **KS:** Mid-Tec Assoc. (913) 541-0505, **KY:** Makin Assoc. (513) 871-2424, **LA:** OM Assoc. (214) 690-6746, **MA:** Comp Rep Assoc. (617) 329-3454, **MD:** Trinkle Sales (609) 795-4200, **ME:** Comp Rep Assoc. (617) 329-3454, **MI:** Carter, McCormick & Pierce (313) 477-7700, **MN:** Gibb Elect. (612) 935-4600, **MO:** Mid-Tec Assoc. (314) 275-8666, **MT:** Simpson Assoc. (801) 566-3691, **NC/SC:** Beacon Elect. (919) 787-0330, **ND/SD:** Gibb Elect. (612) 935-4600, **NE:** Mid-Tec Assoc. (913) 541-0505, **NH:** Comp Rep Assoc. (617) 329-3454, **NJ:** Trinkle Sales (609) 795-4200, Win-Cor Elect. (516) 627-9474, **NM:** SMS & Assoc. (602) 998-0831, **NV** (Northern & Central): QCI (408) 432-1070, **NV** (Southern): SMS & Assoc. (602) 998-0831, **NY:** Win-Cor Elect. (516) 627-9474, Lamtec (315) 637-3738, **OH:** Makin & Assoc. (513) 871-2424, (614) 848-5424, (216) 248-7370, **OK:** OM Assoc. (214) 690-6746, **OR:** ES/Chase (503) 292-8840, **PA:** Trinkle Sales (609) 795-4200, **PUERTO RICO:** Beacon Elect. (809) 728-5040, **RI:** Comp Rep Assoc. (617) 329-3454, **TN:** Beacon Elect. (404) 256-9640, (205) 881-5031, **TX:** OM Assoc. (512) 388-1151, (713) 789-4426, (214) 690-6746, SMS & Assoc. (El Paso) (602) 998-0831, **UT:** Simpson & Assoc. (801) 566-3691, **VT:** Comp Rep Assoc. (617) 329-3454, **VA:** Trinkle Sales (609) 795-4200, **WA:** ES/Chase (206) 823-9535, **WI:** Gibb Elect. (612) 935-4600, Phase II Mktg (312) 303-5902, **WV:** Trinkle Sales (609) 795-4200, **WY:** Simpson Assoc. (801) 566-3691, **CANADA:** Tech Rep Elect. (416) 890-2903, (613) 225-9186, (514) 337-6046, (604) 254-2004, (604) 432-1788

CIRCLE NUMBER 9

CAE in Japan: Misadventures and Opportunities

Japan Has Been Slower to Buy into CAE

ISADORE TAUB KATZ, DATAQUEST INC.

OVER THE past 25 years, the Japanese have established themselves as leaders in the electronics industry, and to paraphrase a current TV advertisement, the Japanese companies gained this leadership position the old fashioned way: they earned it. There were no magical software programs that made their designs better. Instead, they developed systematic approaches to designing, manufacturing, and delivering high-quality products to customers on schedule and at the right price. Long before CAD or CAM was ever invented, the Japanese had created organizations that knew how to bring products to market successfully.

It is interesting to observe how Japanese companies have approached the usage of computer-aided engineering (and I mean just CAE and not computer-aided design). Whereas the American, and then the European, electronics companies were quick to adopt CAE, the Japanese have been much slower to buy into these technologies, despite the promise of major productivity gains in the design process.

Far Eastern investment in CAE has lagged behind both Europe and North America. While North American spending soared to over $500 million in the space of five years, 1981–1986, and European companies' spending climbed to over $200 million in the same period, Far Eastern spending has remained under $200 million (see the graph). If one takes into account the fact that part of the increase in Far Eastern CAE revenues reflect the appreciation of the Japanese yen, then the lag in the Japanese CAE market becomes even more apparent.

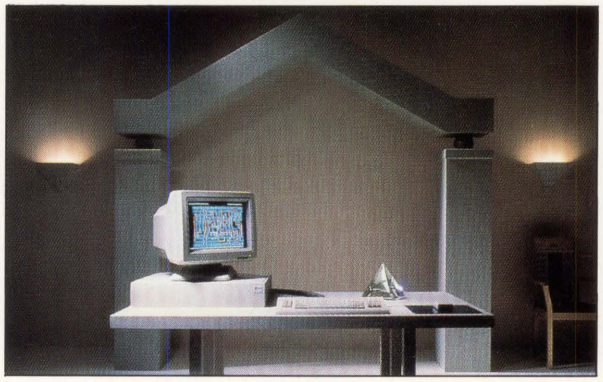

These figures are, however, somewhat deceptive. The Japanese semiconductor industry has made a significant commitment to CAE, even though other industry segments, particularly consumer electronics, have not. Also, the computer and system manufacturers are now beginning to make investments in engineering tools. The Japanese CAE market does appear finally to be taking off.

■ WHY THE WAIT?

There seem to be several reasons why most Japanese companies, excepting semiconductor manufacturers, have been slow to join the CAE revolution.

- Technology and competitive pressures in semiconductors
- The role and behavior of ASIC users in Japan
- The responsiveness and behavior of CAE vendors

Taken together, these three forces begin to explain why many Japanese companies have held off major purchases of CAE tools.

'THE MAJORITY OF ASIC DESIGN WORK

IS STILL DONE BY THE ASIC SUPPLIER'

First, the Japanese semiconductor industry, as is true worldwide, was driven into the use of CAE tools by technology and competitive pressures. In the case of VLSI chips, design complexity forced the adoption of software-based engineering. For semicustom chips, the sheer design volume and need to service the customers pushed Japanese semiconductor suppliers into CAE. The Japanese semiconductor suppliers invested in internally developed tools, such as logic simulators or customized versions of SPICE; established joint ventures with foreign companies with software expertise (for example, Toshiba with LSI Logic), and invested in commercially available tools from CAE suppliers like Daisy, Mentor, and Valid.

It is important to note that the movement to CAE by the semiconductor companies also pushed the U.S. and European system houses to the use of CAE workstations, especially for ASIC design. In fact, the majority of the early purchases of CAE stations were for gate array development, and even today close to 40% of workstation-based CAE sales are for ASIC design. Indeed, many of the initial acquisitions of CAE by the Japanese semiconductor companies were explicitly to support their U.S. and European ASIC customers, not internal development.

This fact points to the second major reason for the slow adoption of CAE in Japan. The majority of the ASIC design work in the Japanese market is still done by the ASIC supplier, not by the end user. That is a result partly of the fact that many of the chips are com-

The state-of-the-art in gate array design is a Silvar-Lisco product. Again!

Introducing AVANT GARDS.™ For VLSI sea-of-gates arrays. With 150,000 gates on a chip. And more to come.

Until now, the state-of-the-art in gate array design systems was Silvar-Lisco's GARDS.™ But not anymore.

AVANT GARDS,™ Advanced Gate Array Design System can place and route over 150,000 gates on a single chip using sea-of-gates and three-layer designs. With gate utilization as high as 90%. Adhering to your timing specifications. In less than a week. And supported by the most popular foundries.

That's the *new* state-of-the-art in gate array design. And it's also a Silvar-Lisco product.

But AVANT GARDS™ isn't alone. Silvar-Lisco's cell-based and full custom chip layout software are state-of-the-art as well. And they get better all the time because we improve them all the time. All of them fully integrated with Silvar-Lisco's full line of CAE software that takes you from design capture through mask data preparation. All production proven in thousands of designs.

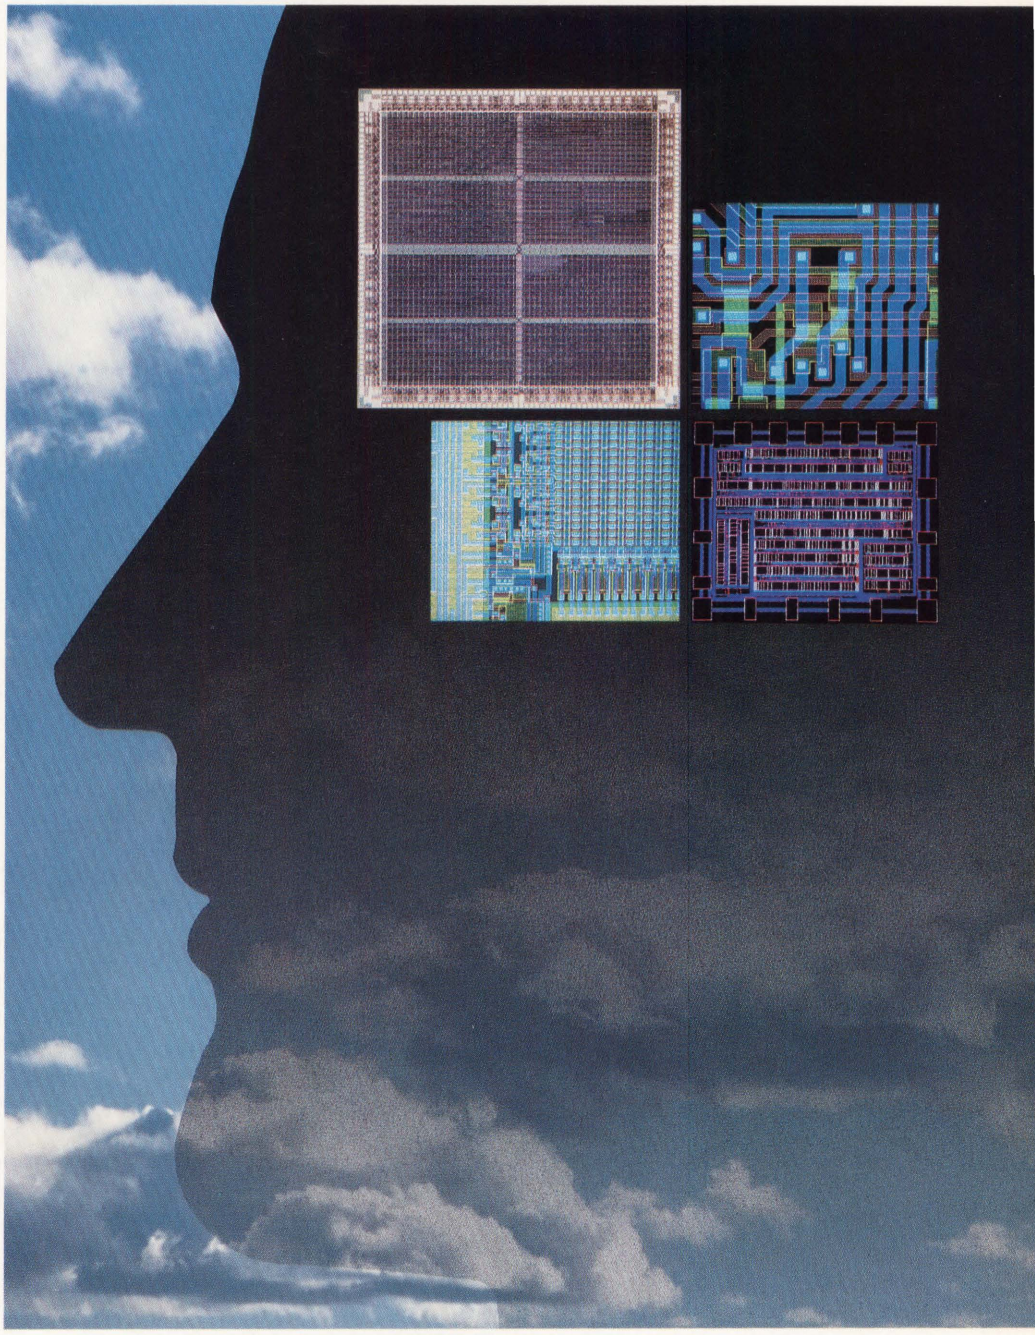

Silvar-Lisco Offices

Corporate Headquarters
1080 Marsh Road
Menlo Park, CA 94025
Telephone: (415) 324-0700

European Headquarters
Brussels, Belgium
Telephone: 32-16-200016

Far East Headquarters
Nihon Silvar-Lisco
Tokyo, Japan
Telephone: 81-3-449-5831

SILVAR-LISCO

CAE Software products for electronic design

Call or write to get detailed literature about the IC layout products that have made Silvar-Lisco the world's leading CAE software company.

CIRCLE NUMBER 10

plex, full-custom ICs with large analog content for the consumer market, and partly of the service traditions in the Japanese market. As a result, Japanese system houses have never been pressured into purchasing CAE workstations as were their U.S. and European counterparts. While American and European companies eagerly turned to workstations as a means of controlling the design process and costs, Japanese companies were closely serviced by teams of application engineers from their ASIC suppliers. Hence, whereas ASICs served as an introduction to CAE for many U.S. and European companies, such that they are now applying engineering tools to a much wider variety of design problems, potential Japanese users have not had the benefit of this focused start in design automation.

Finally, the third major factor was the behavior of the CAE suppliers themselves. Although there are several major Japanese vendors in the CAD (IC and PCB layout) markets, such as Zuken and Seiko, the majority of the CAE companies have been new, U.S.-based companies. Not only have these companies, collectively, experienced significant ups and downs in a very short span of time, but individually many of the suppliers went through excruciating gyrations with their distributors. A few suppliers, such as Mentor Graphics and, now, Valid Logic, have set up strong direct-sales organizations, but this move has been the exception and not the rule. In a marketplace where a company's reputation, service, and support are just as important as product functionality, this sort of behavior does not lend itself to success.

Moreover, the U.S. CAE suppliers have not made a great deal of progress toward "Japanizing" their products. Few suppliers have gone to the extra effort of delivering kanji documentation, let alone kanji versions of their software. In addition, rather than trying to integrate their tools into the Japanese methods of developing products, the CAE vendors have tried to push their own design methodologies onto the user. Yet the Japanese vendors have successful and competitive methods for bringing products to market, and as the saying goes, "if it ain't broke, don't fix it."

To the extent that engineering tools had been used in the Japanese market, they were mainframe-based and closely linked to their design and manufacturing capabilities. Forcing isolated, workstation-based tools into these environments was bound to meet resistance.

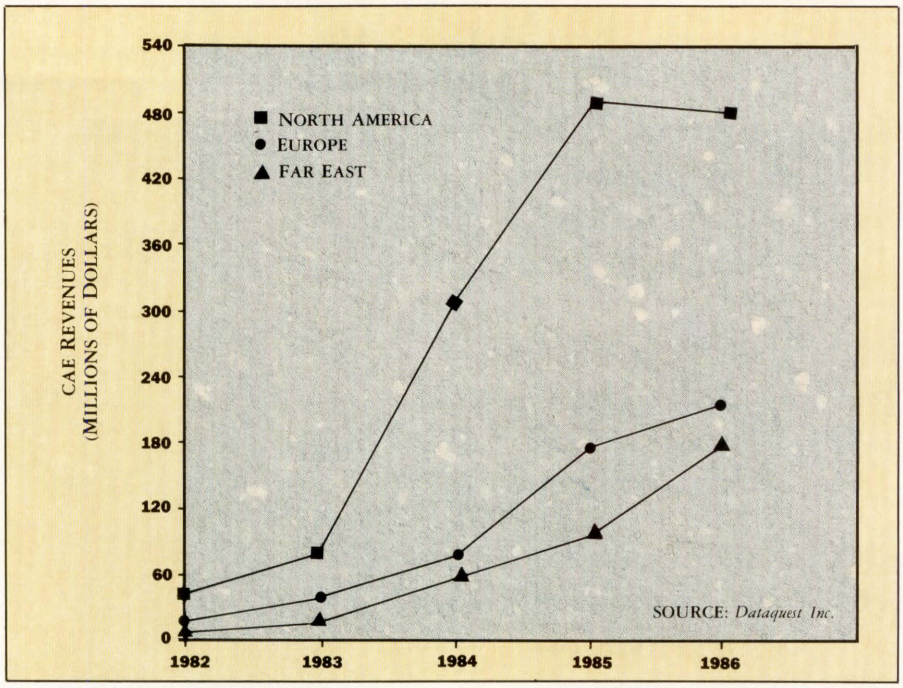

■ WHAT HAS CHANGED

Nonetheless, the Japanese market is one of the fastest-growing regions for CAE today. Dataquest's current forecasts are that Japanese CAE revenue growth should remain between 10% and 20% for the next two to three years. By 1992, we believe that the Japanese CAE market could begin to rival the U.S. market in size. This growth is coming both from continued purchases by the semiconductor companies and, as preliminary results from a Dataquest survey of Japanese EDA users indicates, much wider penetration among the Japanese system houses. (The complete analysis of this survey is to be completed in the near future.)

The initial barriers that kept Japanese companies from purchasing CAE in the past seem to be breaking down. First, and perhaps most importantly, are changes among the CAE suppliers. These changes include the original U.S. CAE suppliers, key distributors, and new entrants to the market from the traditional electronics companies—U.S. and Japanese—like Fujitsu and Hewlett-Packard. All these organizations are paying increasing attention to the service, support, and reliability issues that are essential to selling within the Japanese market. We are even starting to see the first kanji documentation and products that are the key to long-term market growth.

The second major change is the Japanese workstation revolution and the shift from centralized to distributed computing. From 1986 to 1987, the volume of technical workstation shipments in Japan rose from 4,500 units to 16,328, an increase of over 260%. Moreover, over 30% of these stations were made by Japanese vendors such as Fujitsu, Hitachi, NEC, and Sony. Falling hardware prices and the rapid diffusion of workstations into Japanese companies have made the introduction of typical CAE products much simpler as they increasingly fit within the computing environment.

The final shift concerns the Japanese electronics companies themselves. It appears that rising design complexity at the ASIC and system level, competitive pressures, and growing familiarity with CAE capabilities and suppliers are leading the Japanese system houses to invest in CAE. These companies are assuming increasing responsibility for ASIC design from the semiconductor supplier through CAE tools and applying their new expertise to larger and larger problems. CAE has begun to arrive on the desk of the average Japanese engineer.

What all that means is that there is a vast opportunity in the Japanese CAE market. CAE workstation penetration today in Japan is at most 15% of the total available market, and the recent spread of engineering tools and workstations to a much wider audience shows that the market is starting to take off. Nonetheless, the Japanese market is very complex, as is CAE itself, and will take time to reach its full potential. It will be up to the CAE suppliers to deliver products and services that fit with the requirements of the Japanese customers and not vice versa. ■

ISADORE TAUB KATZ *is an industry analyst at Dataquest, responsible for managing and directing research in the electronic CAD/CAM market. Before joining Dataquest, he was a product marketing manager at Daisy Systems.*

In the race to market, GenRad's HILO gives the smooth hand-off you need from your ASIC design to your board design.

Getting your new product to market faster than your competition means added profits for a longer time. And that's what HILO® is all about.

The HILO Universal Logic Simulator helps you design quality into your ASICs fast. And you'll find many leading ASIC foundries using HILO in their in-house simulation.

This means your ASIC design via HILO is the ideal connection with the foundry's system. The result is faster production of high-quality custom and semi-custom devices.

And since HILO can feed forward to your board design, your HILO-designed devices move quickly onto your HILO-designed pc boards where they are complemented by the GenRad HICHIP™ physical device modeler.

Think about it. You'll get quality designed devices and boards in less time. And your finished product will get to market faster. Call 1-800-4-GenRad to find out more about HILO.

GenRad

CIRCLE NUMBER 11

See us at DAC Booth 236

Supercomputing on Chip

CARL DOBBS AND PAUL REED, MOTOROLA INC., AUSTIN, TEXAS

TOMMY NG, SILICON COMPILER SYSTEMS CORP., SAN JOSE, CALIF.

GENESIS OF THE 88000

A new 32-bit microprocessor architecture from Motorola, the M88000 was designed in only 20 months to meet several stringent design goals. An extensible architecture and instruction set that also balanced CPU power and memory throughput led the designers to follow a supercomputer model similar to Control Data Corp.'s 7600 computer.

The centerpiece of the architecture is multiple pipelined function units that execute independently and concurrently. These units execute out of a register set with hardware-monitored interlocks called scoreboarding. In addition, separate data and instruction paths to memory and combination cache-and-MMU chips completed the initial design goals. Silicon compilation tools satisfied the short design time requirements by leveraging the designers' engineering expertise.

■ DESIGN REQUIREMENTS

Motorola began work on a new microprocessor architecture to meet some challenging system goals. Designed with the aim of being the basic processor of future workstations and high-end embedded systems, the architecture had to support Unix and provide a direct path for high-level language software developed on existing machines. It also had to support multiple simultaneous processes. More importantly, to provide higher power for superworkstations and superminicomputers, the architecture had to accommodate multiprocessor designs. Multiprocessor support meant that different levels of throughput could be achieved with the same basic system design but different numbers of CPUs. Additionally, to support one-clock-cycle memory throughput, separate data and instruction paths and elegant cache algorithms were necessary.

All these system requirements had to be put into silicon under the pressures of a short design cycle, aggravated by developments occurring at other companies.

The designers decided that partitioning a supercomputer architecture into a set of chips was the best method to achieve the design requirements. They borrowed from existing supercomputer designs several concepts for high throughput. The existence of multiple pipelined function units (processors), as well as independent data and instruction ports communicating with cache chips having on-board memory management, illustrates the 88000 family's use of fine-grain parallelism. In addition, pipelining techniques allow the processing units to deliver results every clock cycle. To support concurrent processors with pipeline structures, a flexible register file and hardware scoreboarding mechanism incorporate data-flow concepts. These concepts allow the computer and software to use the multiple processing units efficiently. The initial results of the design, the first members of the 88000 family, consist of a CPU chip and a cache-MMU chip (Figure 1).

■ DESIGN TIME REQUIREMENTS

As mentioned, when the development of the 88000 family began in 1985, increased competition and economic condi-

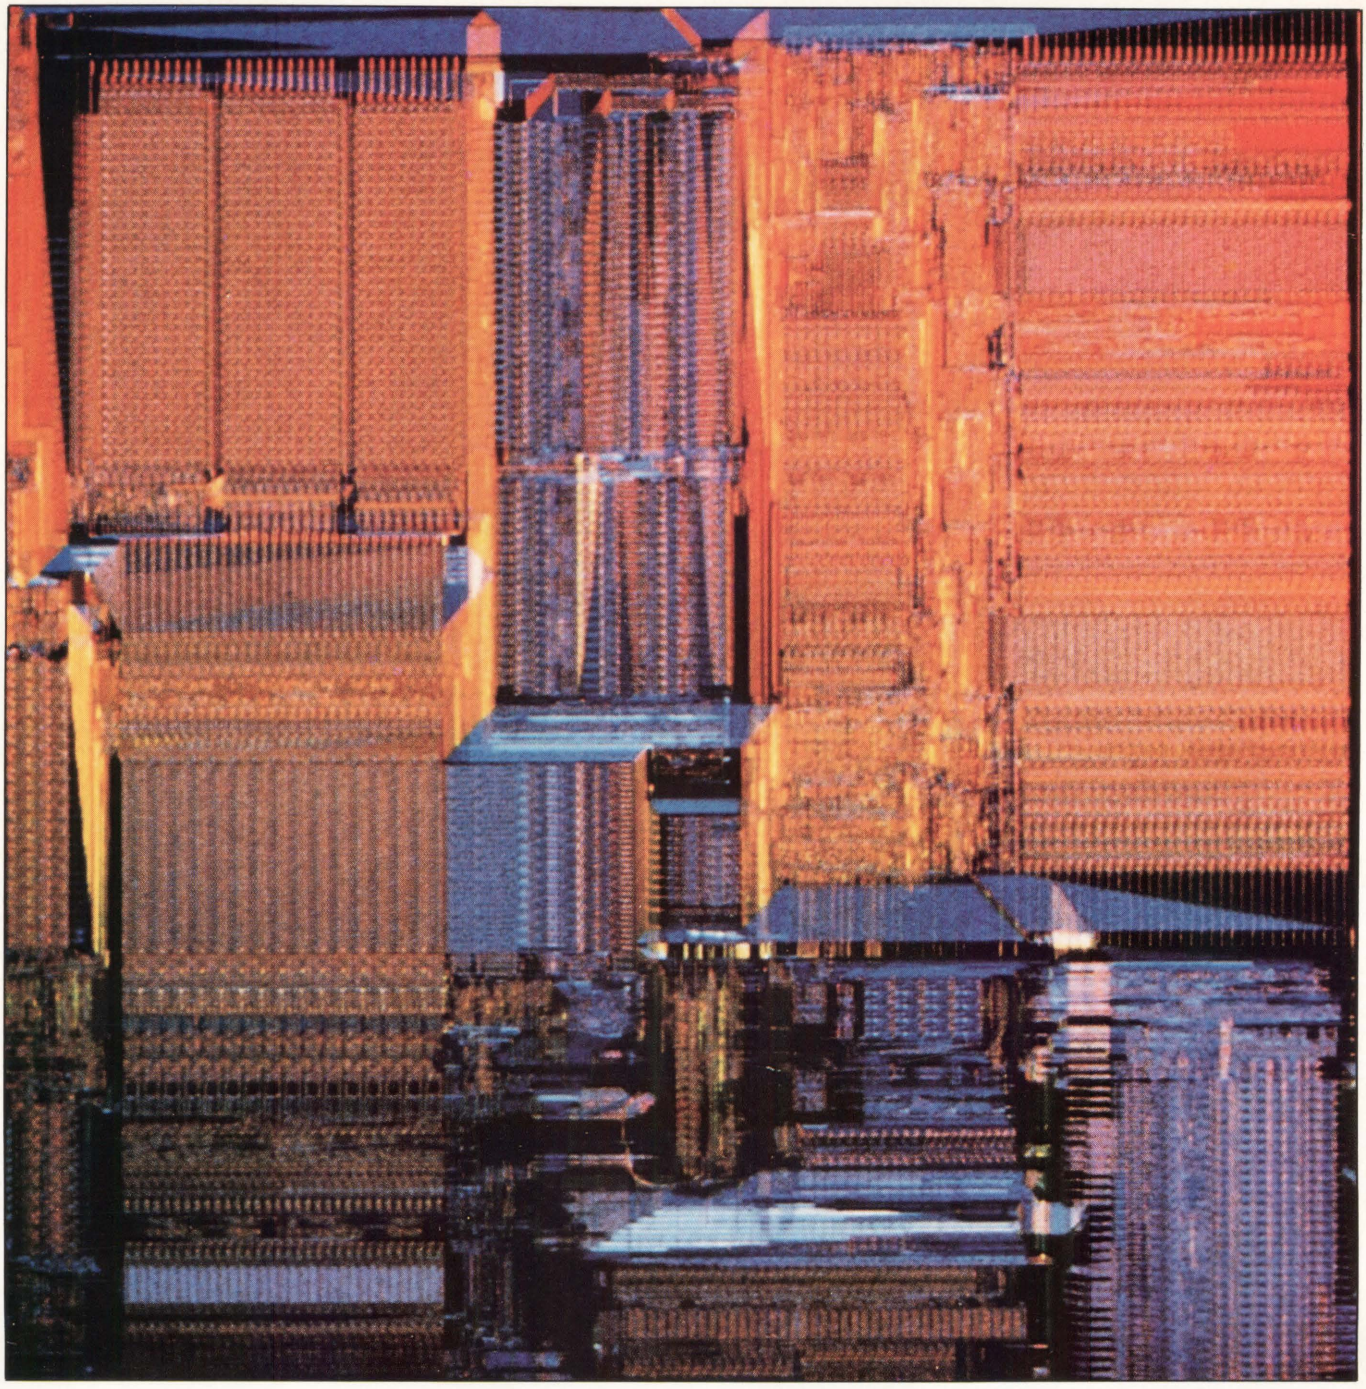

tions constrained the development program. Not only was the design team few in number, but also the design cycle had to be compressed to get the chip in production as soon as possible. In addition, "flexibility" had to be designed into the product so that future derivatives could be developed quickly. Even so, the design had be compact for manufacturability and performance.

These constraints produced a need for a highly automated design system, a turnkey one that needed little overhead support. Also, because of Motorola's extensive experience in semiconductor design and manufacturing, the system had to be configurable so that company engineers could have complete control over the design.

Motorola chose the GDT system from Silicon Compiler Systems (SCS) because of its integration into one environment of all the critical design tools: functional modeling, circuit design, physical layout, simulation, and verification. The integration allowed the one system to carry the project from early arhitectural decisions to final layout. Also, the integrated mixed-mode simulation and procedural layout capabilities of GDT allowed Motorola to use an iterative, rather than a serial, design process for the 88000 family. In addition, Motorola was able to customize the tools to fit its needs by defining a design methodology, integrating transistor models, and connecting existing CAD tools to the GDT environment.

The basic micro-architecture of the processor was defined in a top-down approach. Using M, the behavioral modeling language in GDT, the team developed a functional model of the chip. In most respects, this model reflected the hierarchy of the circuitry. It provided an "executable specification" that was compatible with other aspects of the design, thus avoiding the need for a breadboard and saving engineering resources for other design tasks.

The circuitry of the 88000 chips was designed in a bottom-up approach in parallel with the top-down design. The de-

BUILD KILLER CELLS.

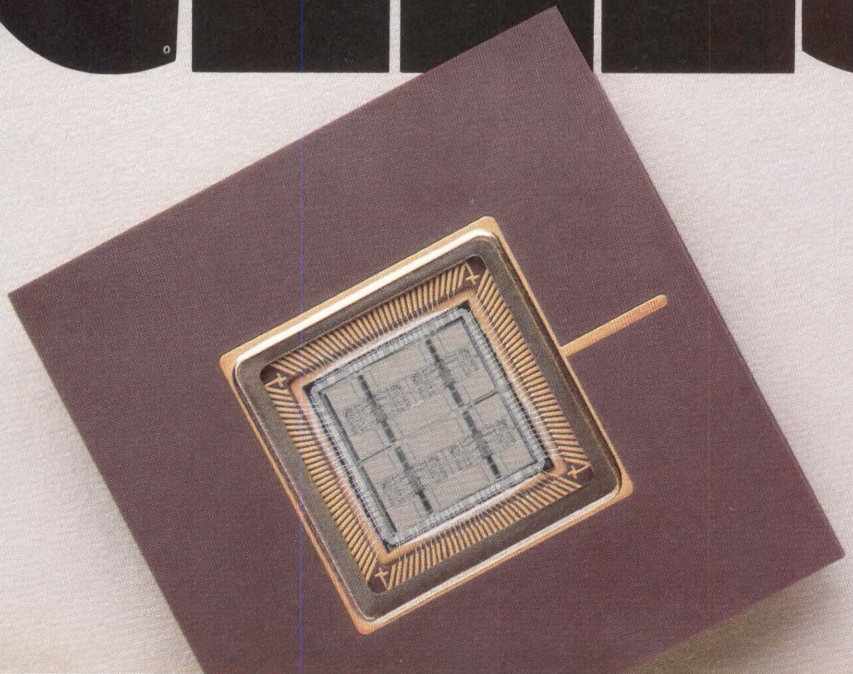

This is the stuff reputations are built on.

A cell design with 100,000 gates. More than 1000 library elements. 2MB of memory. And sub-micron HCMOS technology.

You can search high and low, but there's only one place in the world you can find a chip of such staggering complexity.

LSI Logic.

Designs like this are precisely why we're the #1 domestic ASIC supplier with more than 4,000 working cell and array-based designs in the field.

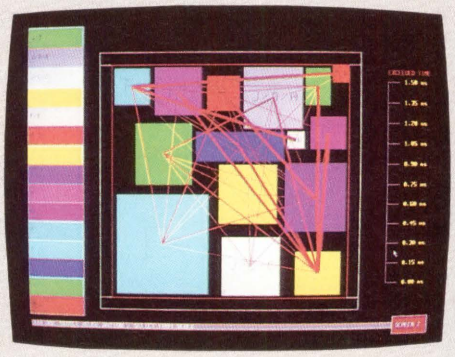

High complexity Cell-Based ASICs require advanced tools like the MDE chip floorplanner to optimize delays and verify performance prior to layout.

And why our Modular Design Environment (MDE)™ software is the best foundation for building killer cells.

MDE is the industry's most advanced design software for ASICs. It arms you with the capability to build today's most sophisticated cell-based designs. Easily. And with the smallest possible die size.

Besides MDE, LSI Logic delivers more than 400 LSI and VLSI building blocks, including 16 and 32-bit microprocessor cells and industry-standard processor and peripheral functions; the longest list of SSI and MSI functions; and memory compilers to develop exactly the RAM or ROM your chip requires.

And if your application calls for lower densities, that's covered, too. Because LSI Logic has more cost-effective cell-based solutions than you can imagine.

No matter what kind of Cell-Based ASIC you build, LSI Logic will deliver a fully-tested prototype in as little as 4 weeks.

All in the production quantities you want, thanks to our advanced worldwide wafer fabrication, assembly, and test facilities.

So find out more about LSI Logic's Cell-Based ASICs by calling the sales office nearest you.

After all, we can help you make a killing.

LSI LOGIC

LSI Logic Sales Offices and Design Resource Centers: Scottsdale, AZ 602-951-4560, Milpitas, CA 408-433-8000, San Jose, CA 408-248-5100, Irvine, CA 714-553-5600, Sherman Oaks, CA 818-906-0333, Denver, CO 303-756-8800, Altamonte Springs, FL 305-339-2242, Boca Raton, FL 305-395-6200, Norcross, GA 404-448-4898, Chicago, IL 312-773-0111, Bethesda, MD 301-897-5800, Waltham, MA 617-890-0161, Ann Arbor, MI 313-769-0175, Minneapolis, MN 612-921-8300, Bridgewater, NJ 201-722-7522, Poughkeepsie, NY 914-454-6593, Raleigh, NC 919-783-8833, Worthington, OH 614-438-2644, Beaverton, OR 503-644-6697, Trevose, PA 215-638-3010, Austin, TX 512-338-2140, Dallas, TX 214-788-2966, Bellevue, WA 206-822-4384, Calgary, Alta 403-262-9292, Paris, France 33-1-46212525, Israel 972-3-403741/6, Milan, Italy 39-39-651575, Tokyo, Japan 81-3-589-2711, Seoul, Korea 82-2-785-1693, Nidau/Biel, Switzerland 032-515441, Bracknell, United Kingdom 44-344-426544, Munich, West Germany 49-89-926903-0.

CIRCLE NUMBER 12

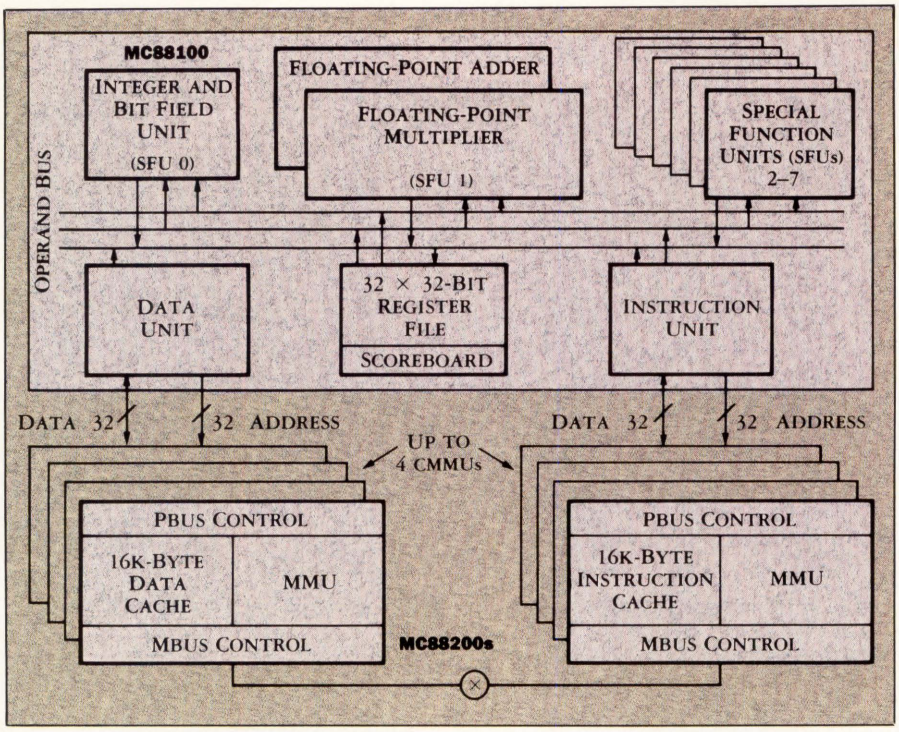

Figure 1. The 88000 "supercomputer" architecture, composed of independent, pipelined function units, takes form in the MC88100 CPU and the MC88200 cache–and–memory management unit (CMMU).

signers created the layout for the lowest-level leaf cells by hand using Led, the graphical editor in GDT. The designers used L, the GDT programming language, to create compiler modules that wired together the leaf cells into blocks in increasing levels of hierarchy. This combination of graphical and procedural circuit design resulted in a chip density comparable to that of handcrafted chips, and it also retained the flexibility of an automated design approach.

■ Register-Driven Processors

The heart of this multiprocessing architecture is the register file's "operand bus" and the hardware scoreboarding. The scoreboard (Figure 2) synchronizes operand usage by the multiple function units. The instruction register (IR) has the register operands in the same fields of every instruction, allowing register decoding to be done in parallel with instruction decoding. The register decoder fetches the operands and also checks the scoreboard bits for those operands. If any of the S1, S2, or D (destination) operands has a scoreboard bit set, execution of the instruction cannot begin. The scoreboard bit of a destination register is set when an instruction is started and cleared when the result is written back into the register.

The function units are clustered on the operand bus with equal access to the register file, as in a supercomputer design. The scoreboarding not only makes possible a type of data-flow architecture—by ensuring that instructions are processed only when the operands contain valid data—it also relieves the compilers of having to worry about the use of registers. Software can run on the 88100 microprocessor without any regard to register usage. Properly written compilers, however, can make better use of the pipelines than the hardware scoreboarding by itself. The scoreboarding mechanisms will also let future members of the 88000 family, which may have different levels of pipelining, be binarily compatible with the original.

Using Led, the designers were able to lay out both the register file and a 32 × 32-bit combinatorial multiplier in a matter of days. The layout helped the designers determine the area requirements and performance characteristics for these basic functions, factors that guided decisions about the structure and number of registers in the processor and about the design of the instruction set. Similarly, Lsim, SCS's mixed-mode analog-and-digital simulator, was used to determine pipeline delays and, consequently, to make decisions about the number of pipeline stages and their length. Many different regions were simulated at the same time using different levels of abstraction.

With the register file and scoreboarding providing the backbone, the processor uses four processing circuits called function units. These basic functional units are the integer and bit field unit, the data memory unit, the floating-point multiplier, and the floating-point adder. The two floating-point units are grouped as a single special function unit (SFU)—the two physical circuits reside at one logical address—so that they can be considered one unit logically, although they execute concurrently. Similarly, the integer unit and the data memory unit occupy another SFU logical address.

The integer and floating-point units can be considered coequals; they have equal access to the register file and scoreboarding function. Because of tighter coupling to the register file and instruction decoder, this access yields more balanced throughput than those systems in which a floating-point processor resides on the microprocessor bus. Also, coequal access lowers the costs in machine cycles of switching types of processes. For example, an operation to remove the exponent from a floating-point number can use the ALU in the integer unit in the next machine cycle, because the number is in a register common to both processors. All software containing a mix of integer and floating-point instructions should run faster with coequal processing units.

The lower overhead to initiate an instruction begins floating-point operations in only one half a cycle; if the floating-point processor were on the system's memory bus, the floating-point instruction would take three cycles to get to the processor.

The functional units and register schemes were defined to support software developed on existing architectures. This capability depends basically on maintaining the same representation of data and objects.

In addition to the integer and floating-point units, six more function units can be connected to the operand bus at SFU logical addresses. The hooks were placed in the architecture to make it extensible. Future enhancements —SFUs designed either by Motorola or by customers—also can have equal access to the register file. These SFUs can operate concurrently with the other SFUs and can be pipelined. The SFU bus specification uses a very general prototcol to permit the incorporation of arbitrary extensions.

To support the addition of SFUs, the architecture reserves 512 unique op codes for each of the six unused SFU addresses. One precise exception vector, one imprecise exception vector, and 64 control registers are allotted to each SFU. The SFU can use these registers in many ways, even as a FIFO. Bits in the control register space are allocated to enabling and disabling SFUs, thus supporting exception-handling facilities in those units.

The function units were also designed from the ground up. By wiring together several leaf cells, the designers created primitive functional blocks such as an adder and a barrel shifter. Generators then wired these blocks together to form larger blocks, until such blocks as the floating-point adder pipeline and the integer execution unit were complete.

■ BUILDING THROUGHPUT

To keep the processor executing smoothly, a pipelined data memory unit and an instruction unit provide data and instructions to the CPU at a rate of one each per clock. The data unit is bidirectional; it serves as a pipeline for new data to the register file and also loads data from the register file back into the data cache. Because it operates in a pipelined manner, it can accommodate cache misses without stalling the rest of the machine. It can be considered part of a processing pipeline because it operates independently while other instructions are processed.

The instruction unit both fetches instructions and performs most of the instruction decoding. It assigns processes to a particular function unit, which will then execute them. It also sends control signals to the register file to coordinate instruction execution and the scoreboard circuitry with the operands needed for the instructions.

Under the control of the hardware scoreboarding, the MC88100 can execute optimized and unoptimized code. Consider, for instance, the inner loop of a fast Fourier transform. Uncompiled, it executes sequentially, requiring 53 clock cycles. When compiled with a scheduler algorithm, the instructions overlap significantly because the scheduler understands data dependencies and the pipelined operation of the functional units. In Figure 3, the inner loop has been compiled with the scheduler algorithm (release 1.0) and executes in only 27 clock cycles.

Critical to the data and instruction pipelines is the "Pbus" between the processor and the CMMUs. It is synchronized through phase-locked loop circuits between the clocks on the chips so that the clocks operate synchronously. By coordinating the clock signals on the chips, the Pbus needs no handshaking cycles. Functionally, it seems that there is one clock for both chips, so that the chip boundaries are invisible. The synchronous bus between the parts, in effect, becomes another part of the overall pipeline.

After designing the data memory unit and instruction unit in the same manner as the previous blocks, the entire chip was assembled from its major blocks by using generators. Specialized L programs routed the global signals and various buses from one major block to the next.

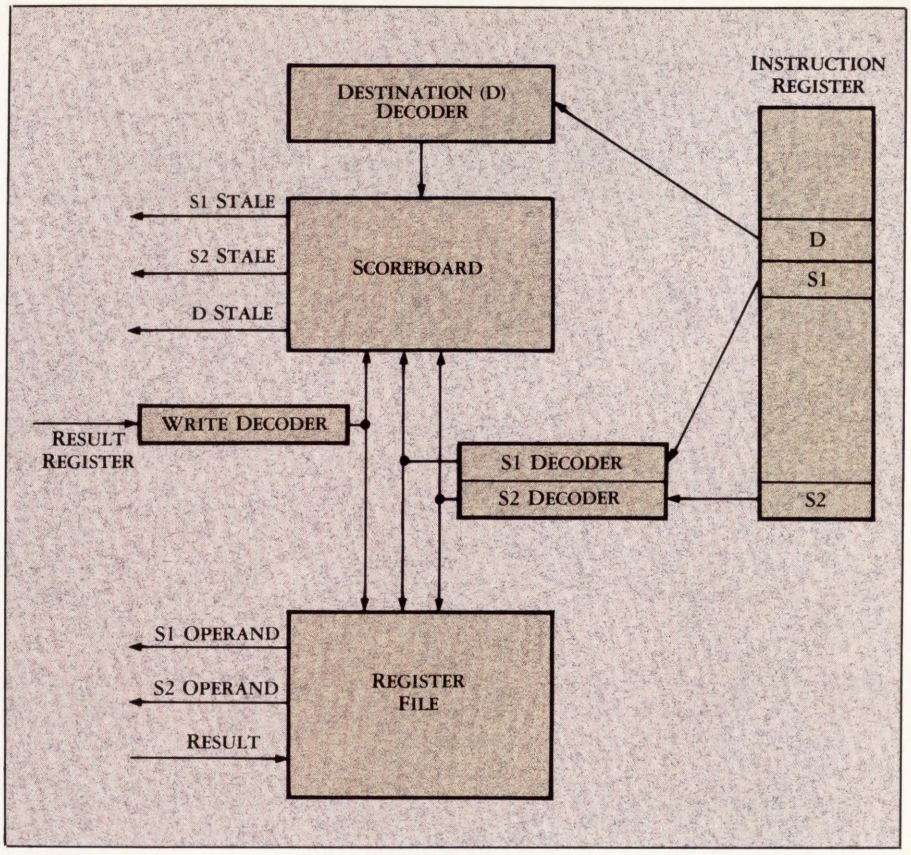

Figure 2. The register scoreboard mechanism concurrently fetches operands and scoreboard bits for returning results. The sequencer can thus control concurrency using the scoreboard as a synchronization mechanism.

The full chip was verified in several different ways. The functionality and timing of the major blocks of the circuitry had been verified as part of the design process. The entire chip was simulated at the switch level to check functionality, which was automatically verified by comparing the results of two levels of simulations of the chip, one constructed of the switch-level integer unit and the functional model of the floating-point unit and the other a pure functional model.

Timing verification was done by replacing circuit modules with equivalent functional models. These functional models had timing information built into them so that they could accurately replace the circuitry. As a final functional check, the entire integer section was simulated using Adept, an advanced circuit-level simulator included in Lsim that is as accurate as SPICE but runs in a fraction of the time. This 75,000-transistor full-circuit-level simulation executed in 30 clock cycles and took some 30 hours to run on a 68020-based workstation.

Finally, the entire chip database was translated into the GDS II format and design-rule-checked using Motorola's batch checker. The translated database was then sent to the mask shop.

■ CACHE-MMU DESIGN

Any high-performance processor requires high-performance memory mechanisms to keep the processing hardware operating as close to the peak execution rate as possible. Without caches, memory accesses become a severe bottleneck to throughput. Incorporating caches presents difficult design choices based on whether the cache stores logical or physical addresses. Placing the MMU on the processor chip may speed address translation, but it slows addresses on their way to the cache, and that slowup can require the use of wait states or an expensive cache design. Storing logical addresses in the cache, on the other hand, makes it difficult to maintain cache coherency.

The designers decided therefore to place the MMU with the cache in one, separate chip (Figure 4), with one or more of these chips for the data port and one or more for instruction port. Each access to the cache –MMU (CMMU) chip is processed in one clock period, providing the data or a miss signal back to the processor each machine cycle. This concurrency offers the simplicity of a physical cache with the speed of a logical cache by overlapping address translation with cache look-up.

When a logical address is presented to the CMMU, it decodes bits 11 through 2,

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

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

Trouble in ASIC paradise.

The big day has arrived.

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

It doesn't work.

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

An immediate solution.

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

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

Check out our libraries.

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

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

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

To be continued.

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

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

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

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

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

Mentor Graphics®

Your ideas. Our Experience.

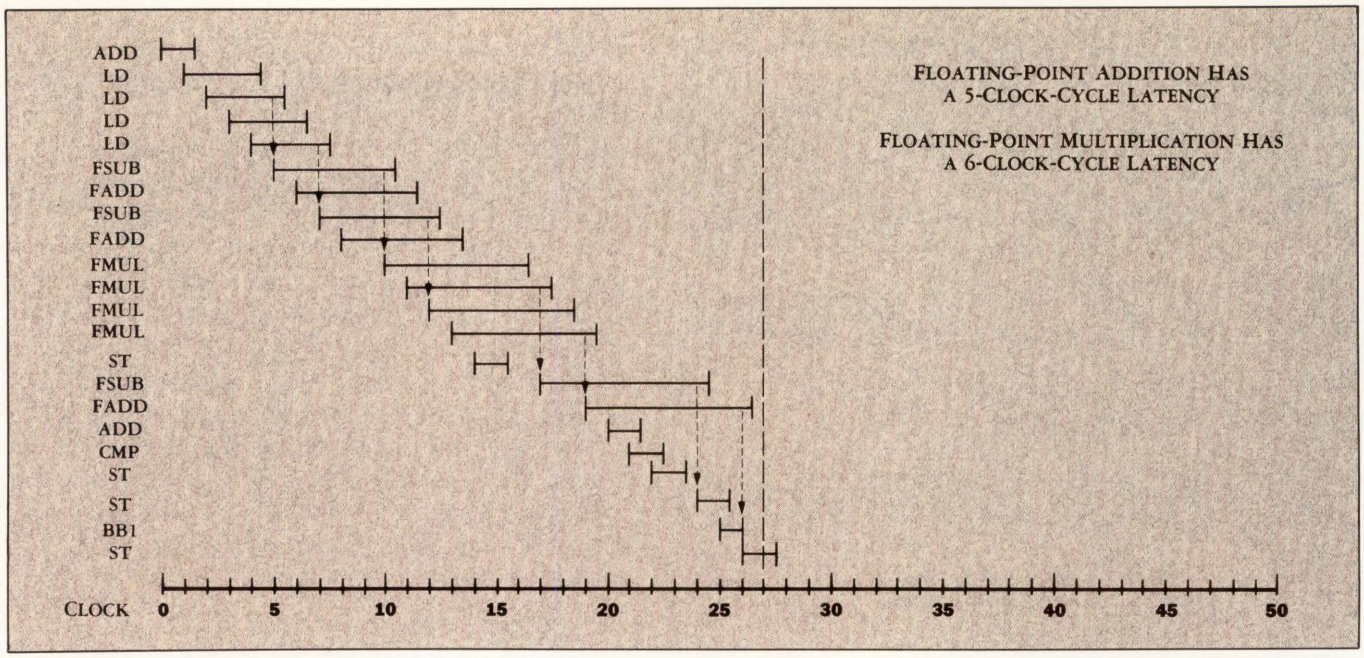

Figure 3. Unscheduled, the execution of an FFT single-precision inner loop takes 53 clock cycles; a compiler can cut that figure down to only 27 clock cycles through intelligent scheduling of the operations. The arrows show where data from one operation is fed directly to another, according to direction from the scoreboard mechanism.

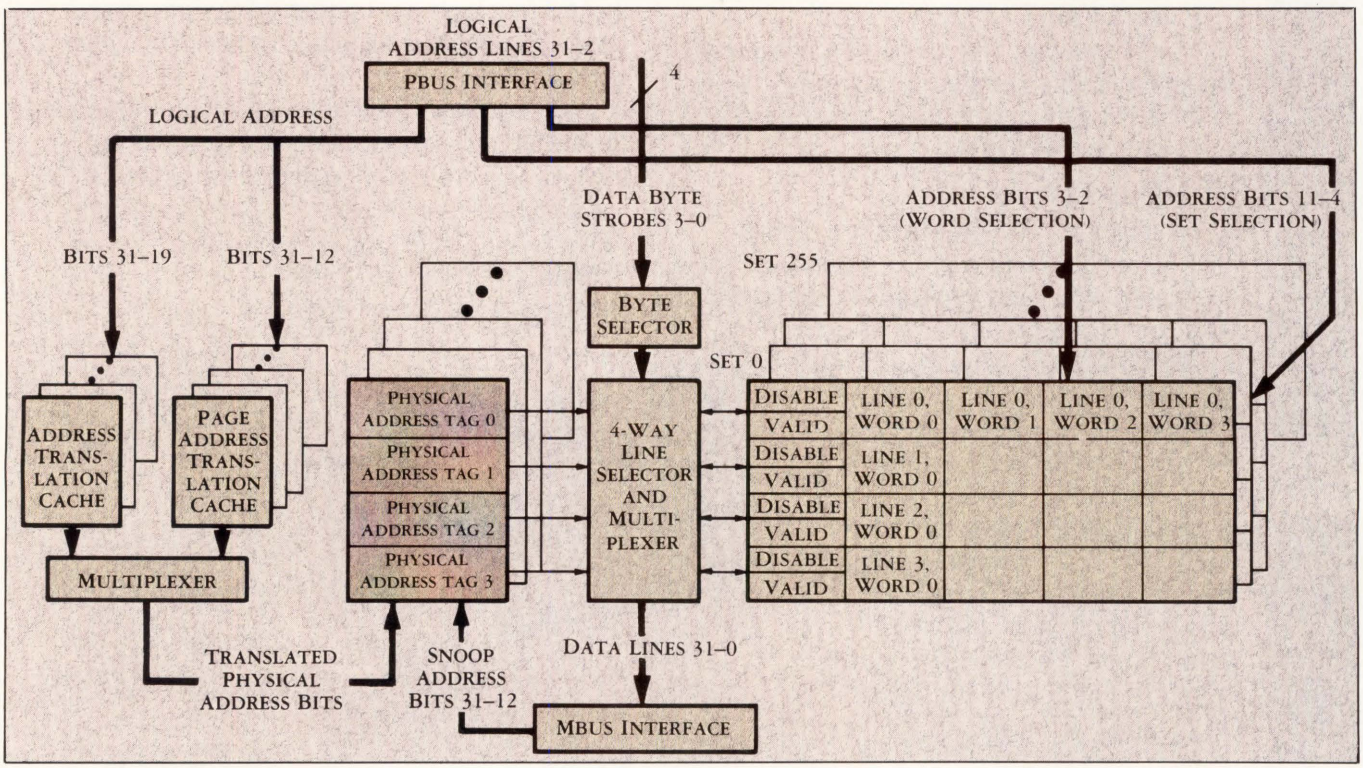

Figure 4. The cache-MMU (CMMU) performs address translation (on the left) and cache access (on the right) simultaneously, producing a result in a single machine cycle.

which select four words in the cache. Simultaneously, bits 31 through 12 access block and page address translation caches (ATCs), producing a physical address. This address is compared with four physical address tags corresponding to the four selected words in the cache. When a match is found, the matched word passes through the multiplexer and on to the processor. When no match is found, the Mbus interface initiates a memory read cycle on the Mbus (a synchronous bus between the CMMUs and main memory) to fetch the requested data or instruction.

Because the CMMU holds a physical cache, it can watch the memory bus and maintain its own coherency. Whenever it sees an address on the bus that matches an address corresponding to data it is storing, it takes steps to maintain coherency.

The cache is four-way set-associative, an organization proven effective in preventing thrashing that can occur in direct-mapped caches. Most designers have difficulty in designing any type of associative caching scheme, so Motorola provided the complete solution. The 16-Kbyte static RAM block was the largest cache that could fit in the chip design.

Because the GDT system captures designs and design experience, Motorola could leverage its MC88100 design experience when designing the MC88200 CMMU.

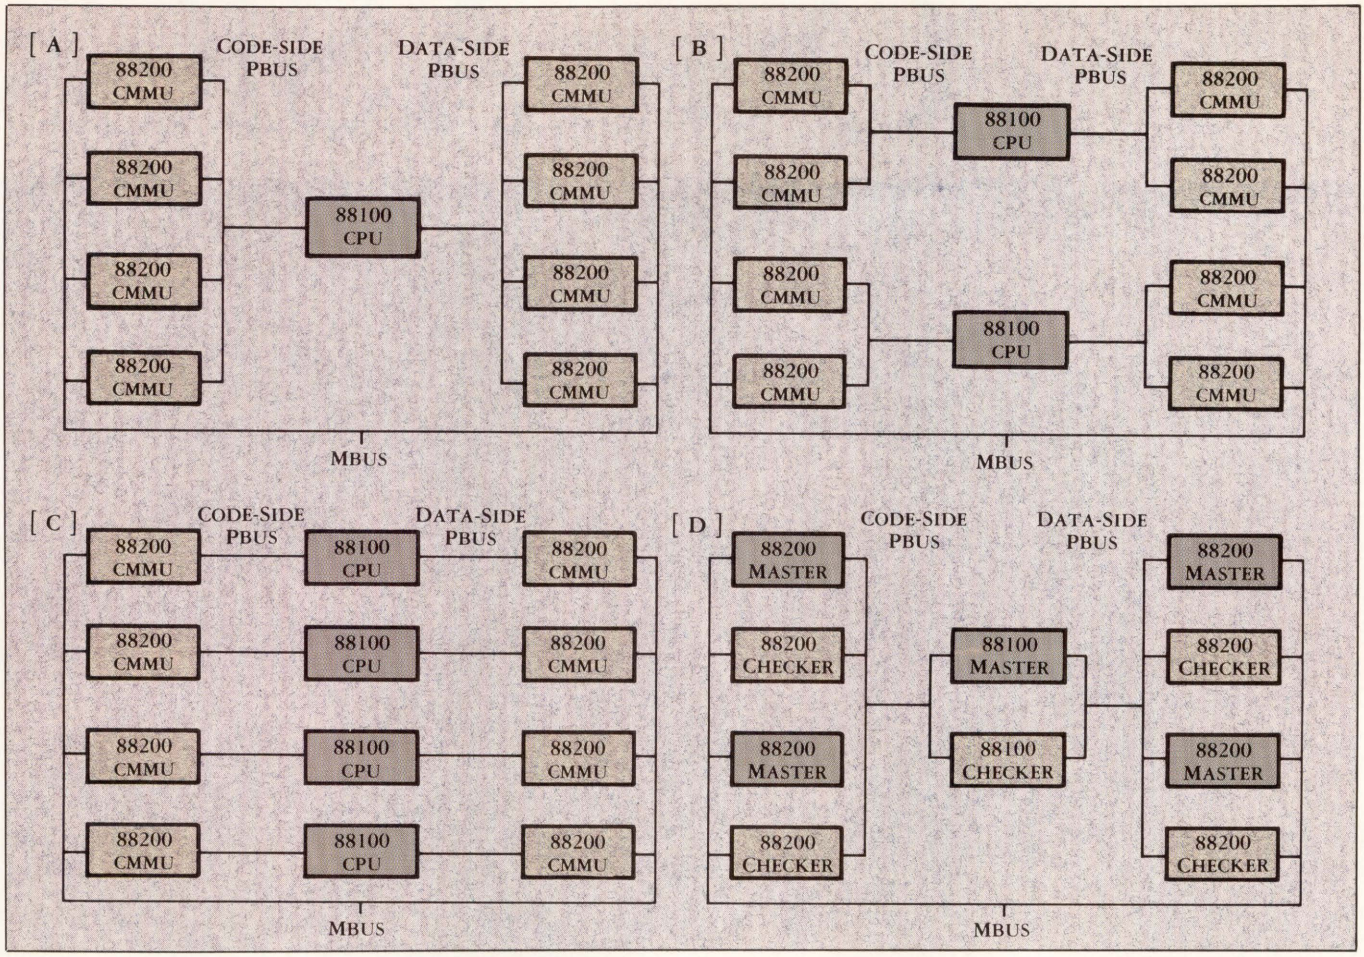

Figure 5. Four systems built from the MC88100 and MC88200 chip set: single-CPU (A), dual-CPU (B), quad-CPU (C), and fault-detecting dual-CPU (D). In the fault-detecting system, one CPU is a master and one a checker; if the outputs differ, error pins are asserted within one-half clock cycle.

The specifications for the CMMU called for the same clocking scheme and a similar output scheme as those of the 88100. Although the CMMU was designed in a p-well process to create a stable static RAM cell, the n-well 88100 designs could be reused because the design database is independent of the design rules.

To meet design completion requirements, the 88200 was jointly developed by Motorola and SCS. Motorola designers created the memory arrays and control circuitry, achieving 20-MHz, no-wait-state performance from the CMMU. SCS designers implemented the control logic—including the memory translation logic, data cache/CAM access and replacement logic, and hit/miss processing logic—and bus interfaces, and they integrated the major blocks into a complete chip.

The functional modeling language simplified the coordination of the two design teams (who were 2,000 miles apart). Motorola wrote functional models of the memory blocks that SCS used to verify the functionality of the entire chip design. The models also helped the designers verify the control circuitry that controls 131,072 bits of the cache memory array, and they provided a convenient way of exploring architectural trade-offs. The result of any changes to the architecture or micro-architecture could be accurately determined by changing the models. Also, the models provided complete documentation of the functionality of the chip's various sections. The 88200 was completed about the same time as the 88100, requiring only 11 months of design time.

Like the microprocessor, the CMMU has implementation flexibility. The two bus interfaces—one with the processor, one with the system bus—may be replaced with interfaces for other processors and system buses. In addition, the size and configuration of the cache and the memory management tables can be customized in future versions of the architecture to fine-tune it for specific types of systems.

The resulting chip set, operating at 20 MHz, can execute instructions at a sustained rate of 17 VAX-equivalent MIPS. Floating-point operations are executed at a sustained rate of 7 MFLOPS. The 88200 CMMU offers the combined functions of both a 16-Kbyte cache and a segmented, demand-paged memory management system. ■

About the Authors

CARL DOBBS *is the circuit design manager for the MC88100 processor at Motorola's Advanced Microprocessor Operations. He has six years' experience in CMOS circuit design, having worked in ASIC design for Motorola and previously for Texas Instruments before joining the microprocessor products group in 1985. Carl received his BSEE and MSEE degrees with high honors from the University of Arkansas, with emphasis on device physics, fabrication and circuit design.*

PAUL REED *is the circuit design manager for the MC88200 Cache-MMU. He has 10 years' CMOS circuit design experience and was involved with both EPROM and SRAM memory design at Texas Instruments and Motorola before joining the Microprocessor Products Group in 1986. He earned a BSEE from the University of Missouri at Columbia.*

TOMMY NG *has been an IC designer at Silicon Compiler Systems since 1984 and is currently director of IC design. From 1980 to 1984, he was a member of the technical staff at Bell Laboratories and was involved in VLSI chip design. Tommy received his MSEE in 1980 from the University of California at Berkeley.*

Our ASICs

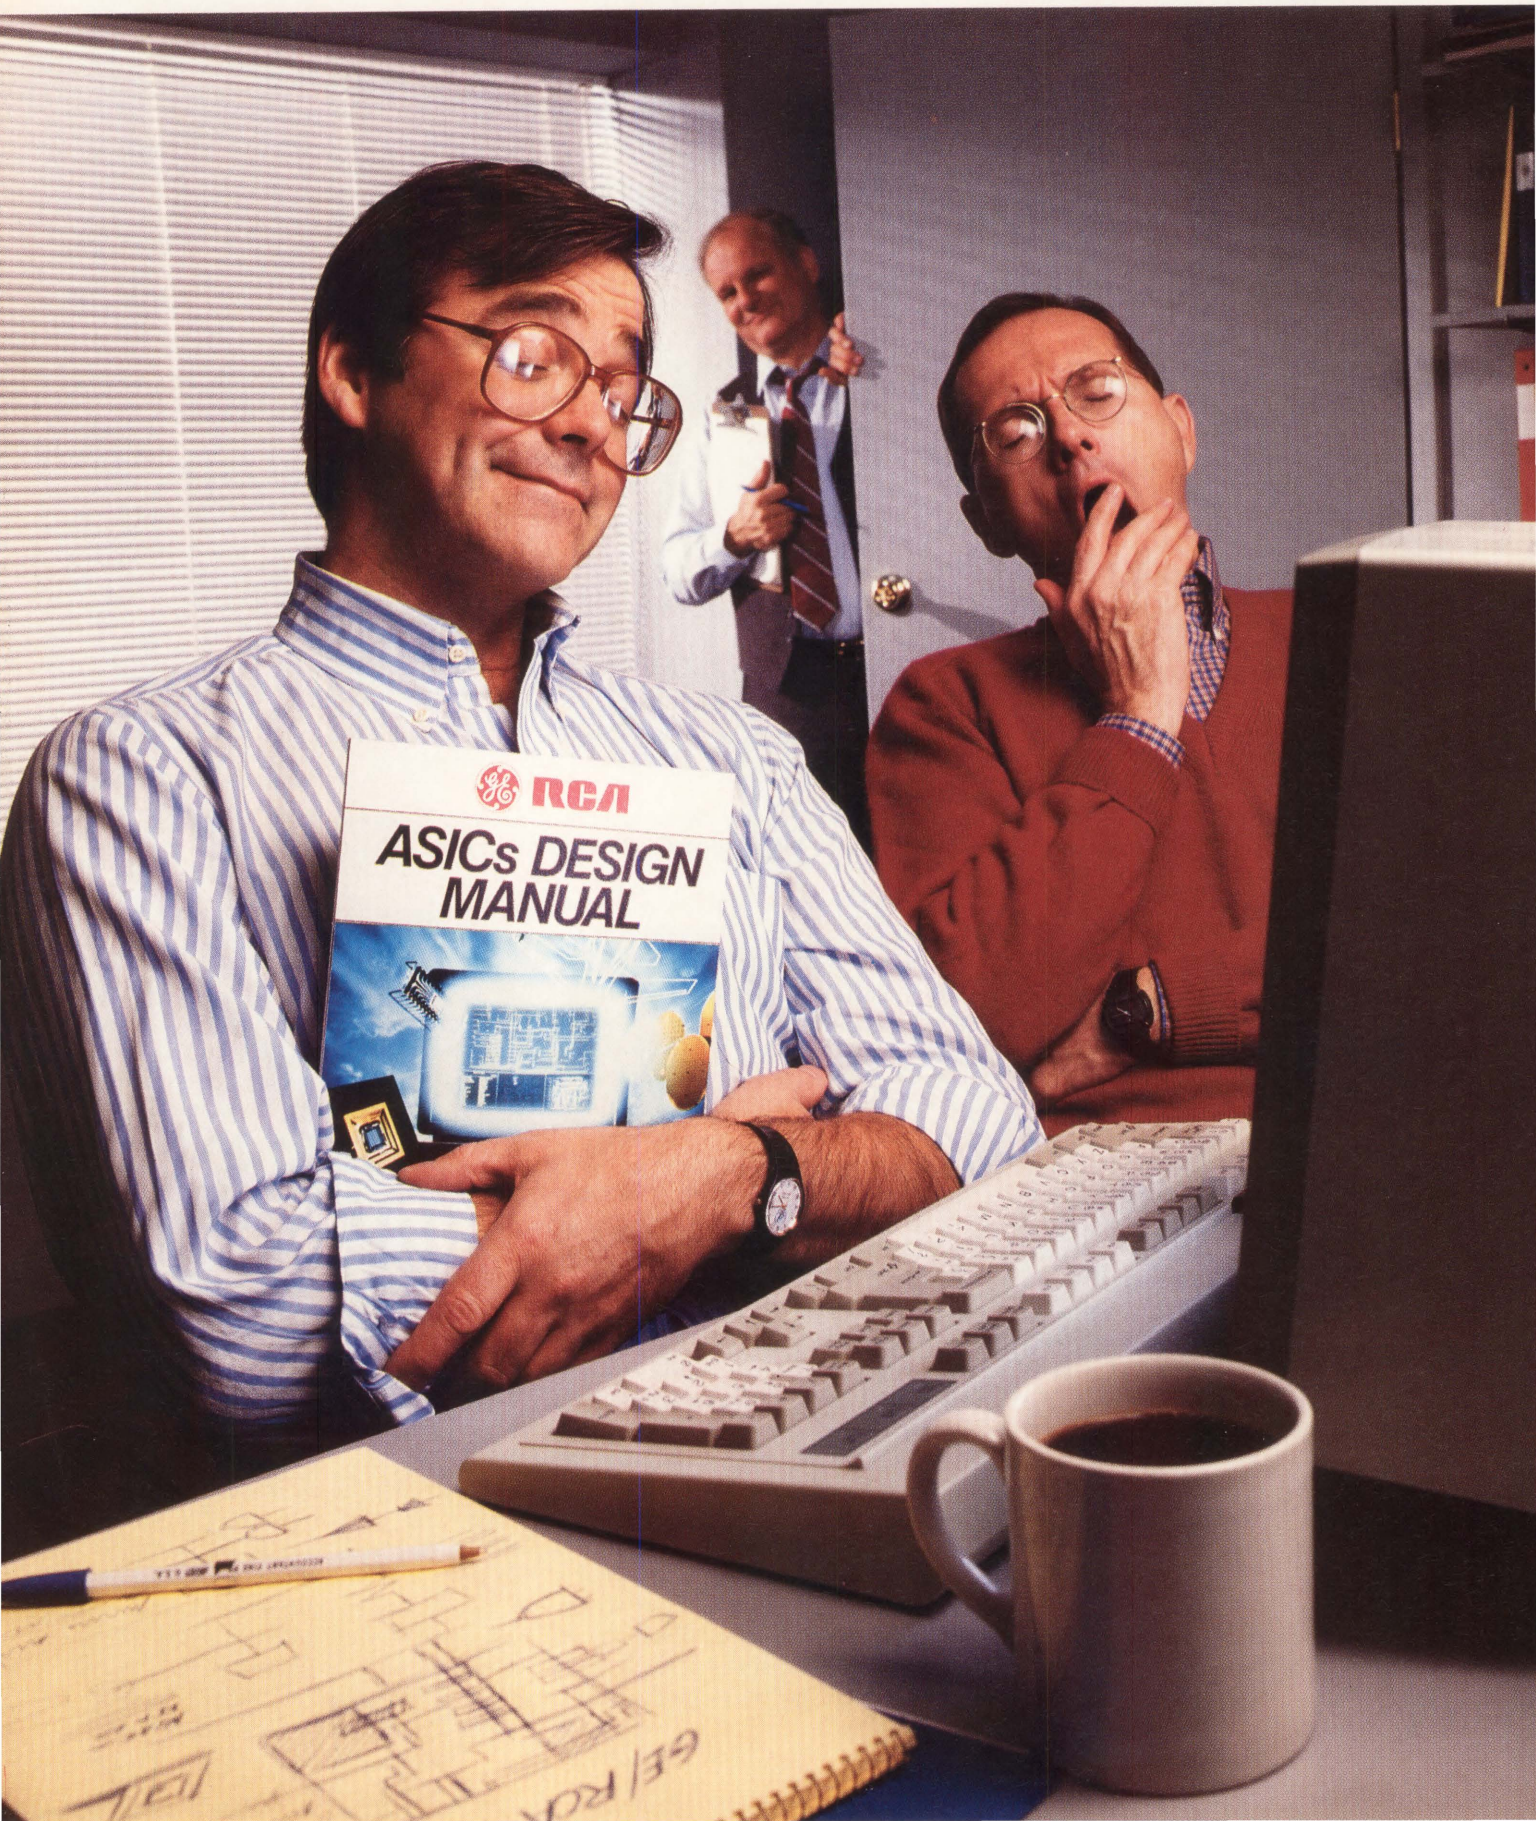

are boring.

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.

GE Solid State

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

*Trademark of General Electric Company, U.S.A., not connected with the English Company of a similar name.

HIGH-VOLTAGE

Analog Semicustom ICs

ALAN P. RAMSEY, HITEK CONSULTING, AUBURN, NY

■ ANALOG SEMICUSTOM TECHNOLOGY HAS DEVELOPED MORE CAPABILITIES FOR SUPPORTING HIGH-VOLTAGE SIGNALS ■

As analog semicustom technology pervades the design of systems, it has developed more capabilities for supporting high-voltage signals: higher breakdown voltages, higher allowable current levels, better heat dissipation, overtemperature protection, and higher gain-bandwidth products. To implement analog circuitry, building-block circuits, usually gain cells, are implemented within analog arrays, and predesigned analog cells, such as op amps, make up macro libraries for cell-based ICs. SOIC surface-mount packages are also available, and some companies are trying to develop packages that have better thermal conductivity than the generic SOIC package. In addition, computer-aided design and design support are expanding, particularly in the area of sophisticated but low-cost software for industry-standard workstations like the IBM PC and compatible personal computers.

These developments are closely related and vital to extending the applications for analog semicustom ICs. Progress with high-voltage implementations in the last year may underline the broad promise for the future of analog semicustom technology. High-voltage ASIC designs are shrinking the size and weight, although not necessarily the cost, of a growing list of electronic systems. These semicustom chips, in both analog and analog-digital forms, are functioning as PWM (pulse-width modulation) controllers in switching power supplies; motor-speed controllers, timer controllers, and ground-fault detectors in appliances, machine tools, and electrical machinery; high-level drivers and receivers in computer peripherals; line-interface drivers for telecommunications, especially modems; and servo amplifiers, actuators, and transducer buffers in industrial control.

A well-designed high-voltage analog semicustom IC, in fact, is capable (with the help of an external component or two for tuning, large capacitances, and inductance) of performing all the functions of a typical linear circuit board designed to work at, for example, a dynamic operating range of ±16 V.

■ VENDOR SELECTION

Most high-voltage analog semicustom devices are fabricated in bipolar technology, and there are considerable differences among bipolar silicon foundries. The physical structures and resulting performance of active and passive components can vary as widely as the chips' architectures. The first step for any designer considering high-voltage semicustom ICs is to narrow down the field of foundries to those who have the combination of devices most likely to successfully integrate his design. Most foundries have specific strengths that tend to imply a special fitness for implementing particular types of applications.

For the highest-rated transistors on analog semicustom ICs, Micrel puts DMOS transistors on its MDP8020 chip, along with bipolar and CMOS circuits (Figure 1). In some configurations, this chip can drive signal levels of 200 V peak to peak (±100 V). Each DMOS transistor can drive signals ranging from 20 to 100 V, with current levels as high as 200 mA. On the

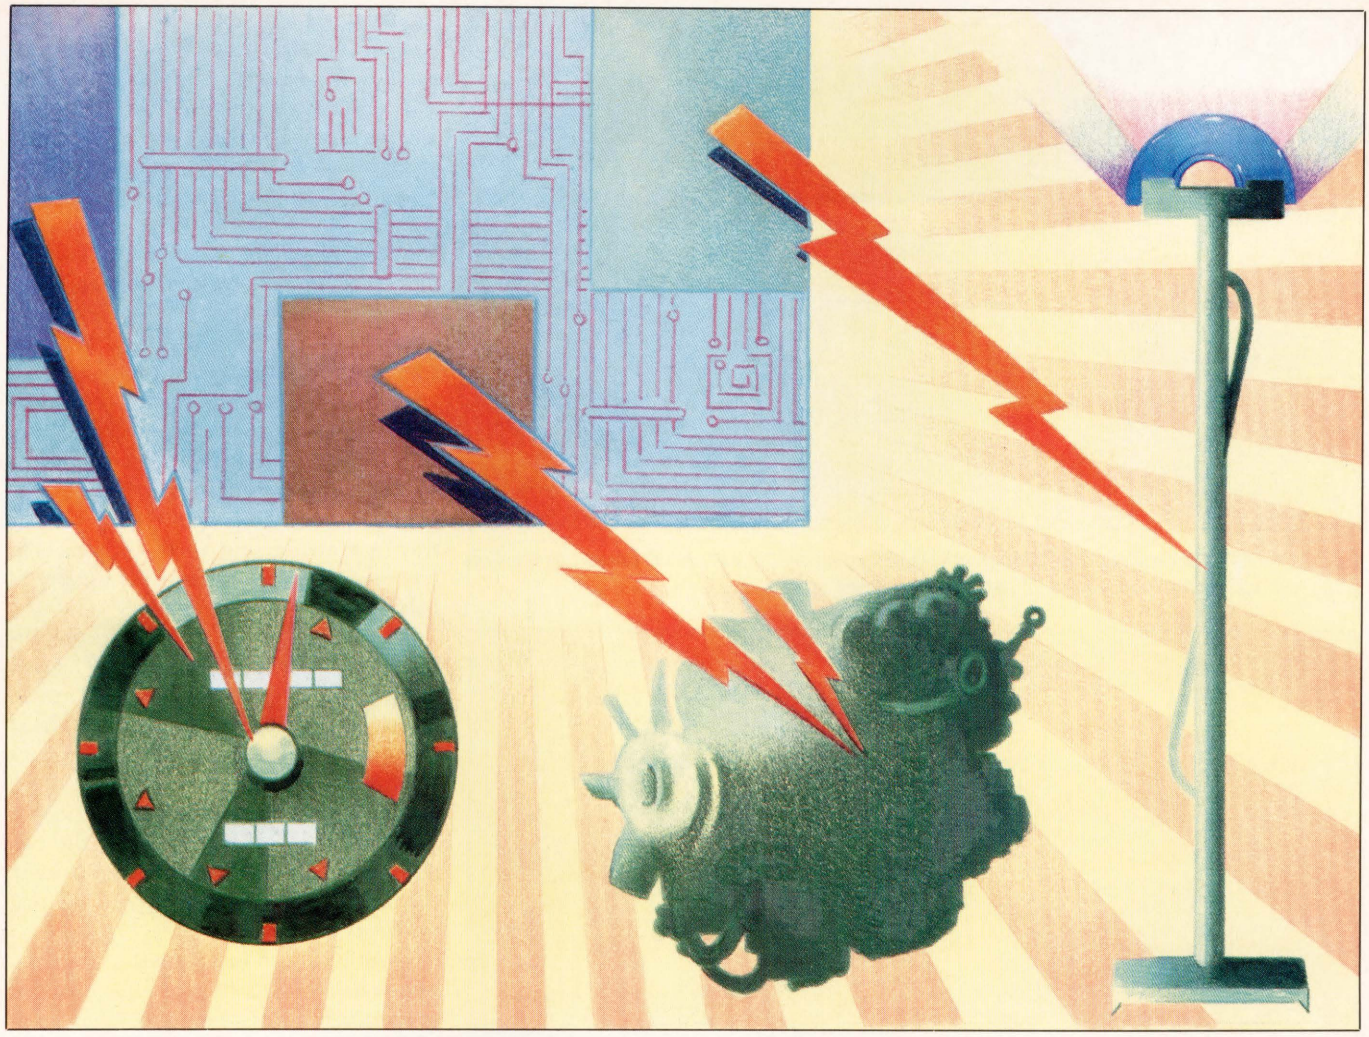

same chip, a single power supply, that delivers from 5 to 15 V, powers the CMOS digital and bipolar analog circuitry that the designer customizes.

AT&T's complementary bipolar integrated circuit (CBIC) technology fabricates fully complementary vertical npn and pnp transistors on the same chip. Most integrated processes use lateral pnp's, which exhibit inferior performance. There are ways of designing around the inferior pnp devices, but the conversion of discrete designs should be simpler with complementary transistors.

If a circuit design requires very close tolerance resistances, or resistors with a high breakdown voltage, the designer should look for semicustom ICs with thin-film resistors. These resistors have better relative tolerance and durability than resistors fabricated in either polysilicon or bulk silicon. Raytheon, for example, uses sputtered SiCr resistor arrays on all its high-voltage analog chips. Tektronix offers NiCr resistors as an option.

Plessey Semiconductors, having acquired Ferranti Interdesign and its Macrochip family in March, offers a compound active device called a "monistor," in addition to three sizes of npn transistors (including vertical npn devices). Because monistors are included in each array cell on the high-voltage MV series chips, efficient device utilization and relatively easier layout should be possible. Plessey offers a variety of design tools, including libraries of macrocells, seven different versions of SPICE models and proprietary programs on standard workstations.

Robust IC design backs up the Genesis 8000 semicustom array from Cherry Semiconductor, making it a good choice for high-current designs. By attaching a large lead frame on a standard-size package, the company, which also designs, fabricates, and markets standard ICs, has developed one of the first surface-mount packages that supports relatively high power dissipation.

Given the variety in foundry specialties, it is difficult for the designer to evaluate all the foundries in depth. Once the decision is made, however, the foundry's strengths should make the rest of the semicustom design experience easier.

■ High-Voltage Design Concerns

Designing a high-voltage analog circuit for semicustom integration is not very different from designing a low-voltage IC. The designer, however, must be more careful of such typical design characteristics as layout parasitics, transients, high currents, thermal gradients, and noise.

Parasitics. Stray leakage currents often occur at "cross-unders"—where one signal crosses another. These currents can be controlled by using cross-unders sparingly and by following the foundry's recommendations con-

HIGH-VOLTAGE ASIC DESIGNS ARE SHRINKING THE SIZE AND WEIGHT OF A GROWING LIST OF ELECTRONIC SYSTEMS

ILLUSTRATION CHERI DORR

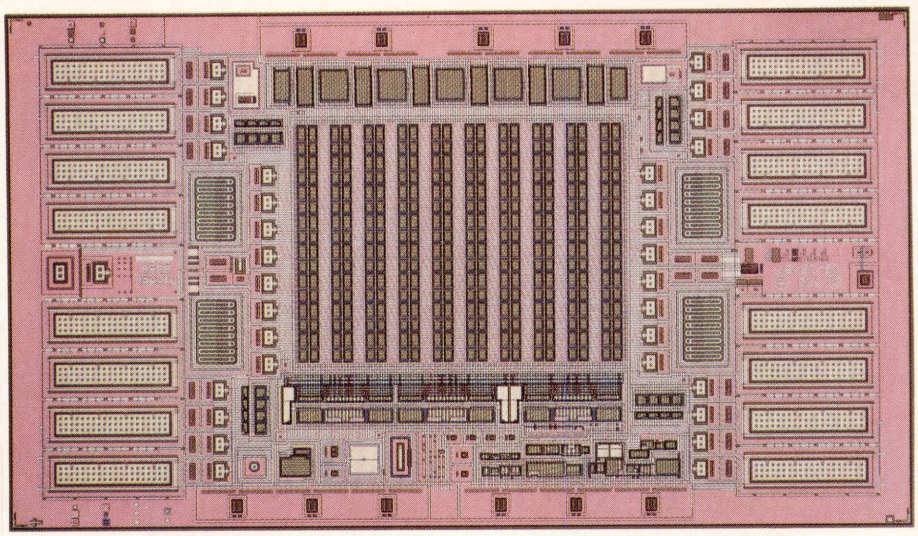

Figure 1. The MDP8020 combines 100-V DMOS power FETs with analog and digital arrays and components to provide a chip that can support a wide range of custom power control applications.

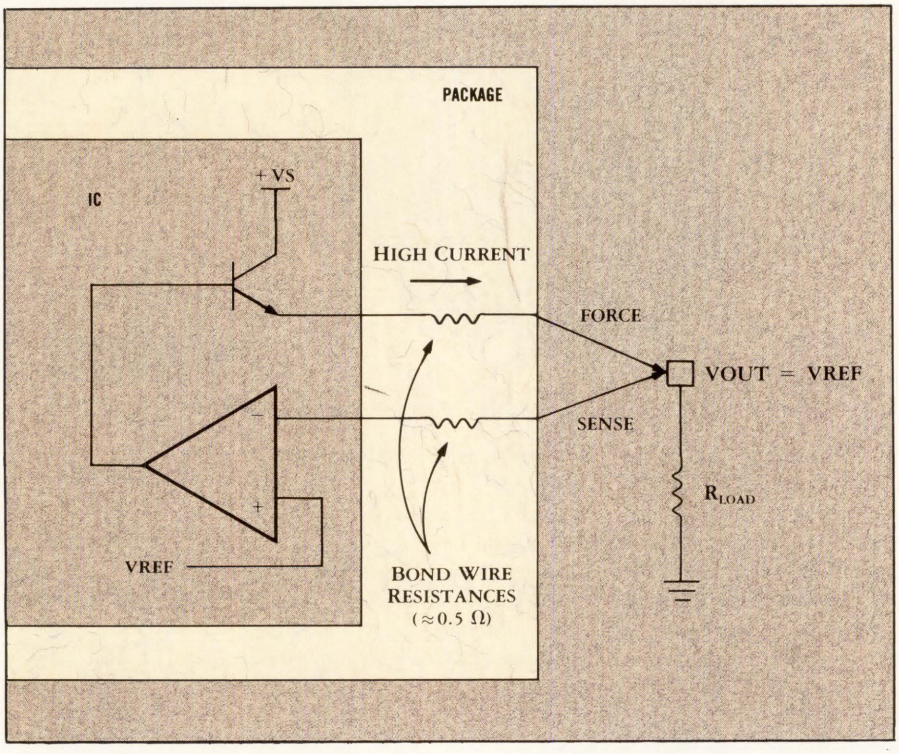

Figure 2. Separate force and sense lines can be used to reduce the effect of common bus voltage drops.

cerning which terminals of an on-chip element to use.

Cross-unders also introduce stray capacitance between the signal line and the substrate. Although the magnitude of such capacitance is only a picofarad or two, it adds up as a signal crosses a chip. In the presence of the large voltage swings of a high-voltage IC, stray capacitance can reduce operating frequencies significantly.

Transients. High-voltage spikes that exceed the breakdown voltage of the chips' transistors should be suppressed at the power supply. Although some ICs may withstand transients as high as 500 V for durations of 1 μs, such stresses are likely to reduce circuit reliability over the long term.

High currents. Extra care must be taken in the routing of high-current paths. Although the metal used will have a low resistivity (roughly 30 Ω/sq), even a small resistance can cause a measurable voltage drop with sufficient current density. In sensitive input-signal paths, the voltage drops can cause unwanted offset voltages and ground loops.

At least one silicon foundry (Raytheon) will widen high-current traces whenever practical if the circuit designer identifies them on his schematic. Widening a trace, however, may not completely solve a voltage drop problem, especially if bond wire resistance is significant. One approach to bond wire resistance uses separate force and sense lines (Figure 2). A similar scheme might be considered for eliminating ground loops by bringing separate ground traces together at a common point.

Thermal gradients. A thermally induced difference in base-emitter voltage (V_{BE}) among on-chip transistors can upset the balance of a current mirror or a bandgap reference. Because of thermal gradients, in general it's good practice to group all power elements toward one edge of a chip and all heat-sensitive elements (differential pairs, for instance) toward the opposite edge. Some chip architectures prevent this segregation, but these chips usually provide some type of thermal symmetry or balance in the circuit layout.

Ideally, sensitive matched pairs of transistors should be equidistant from the heat source. If not, they should be laid out in formations that balance the effects of thermal gradients approaching from any direction. For example, the designer can lay out resistors in a square and connect opposite sides of the square in series; the two resulting resistor chains will experience the same degradation from thermal gradients.

Even with such precautions, compensation voltages may still need to be included in thermally sensitive designs. Compensation techniques are most effective with designs that already have some thermal balance.

Noise. The interface between digital logic and analog circuits should be designed to avoid noise injection from high-voltage rails or from transmission lines. A bandgap regulator and buffer circuitry are needed in typical ASIC designs.

Other "common sense" practices need to be followed during the layout of a design that contains not only high voltages and high currents, but also a mixture of low-level analog and digital control elements. Outputs and noninverting amplifier inputs, for example, should not be placed on adjacent bonding pads to minimize positive feedback that could be introduced by the coupling between adjacent leads. If some I/O pads are unused, it is preferable to spacing signals evenly around a chip is preferable to bunching them together. Also, low-noise, high-gain inputs should be located away from bonding pads that have short-transient waveforms to prevent noise injection, and a ground or supply pad should, if possible, be placed between the pads in this case.

Provided that the chip and its elements can withstand the potentials, it is reason-

Chip, Chip, Array!

For High Speed with Low Power.

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

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

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

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

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

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

A Better Bipolar Array is Here.

AMCC

CIRCLE NUMBER 13

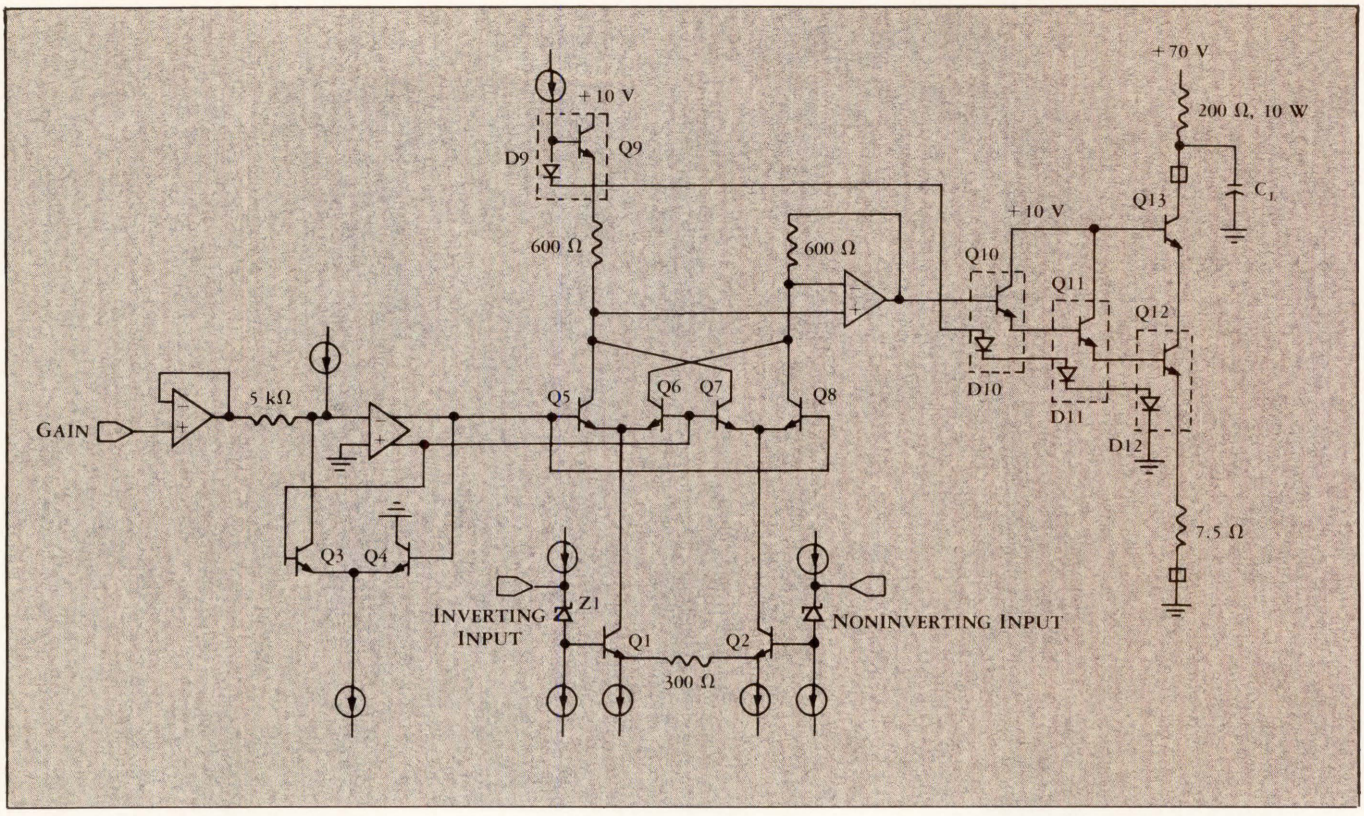

Figure 3. A 175-MHz video display driver chip drives a full 50-V transition into a 6-pF load in just 2 ns. It takes advantage of cascoding techniques to extend its peak-to-peak voltage-handling capability.

ably easy to extend the operating range of a semicustom IC. Normally, the maximum operating voltage is defined by the transistor's collector-emitter breakdown voltage (BV_{CEO}). However, cascoding the output of one transistor into the input of another (as seen in Figure 3) multiplies the maximum operating voltage by the number of cascoded transistors. The new operating limit becomes the breakdown voltage between the collectors and the substrate.

■ HIGH-VOLTAGE APPLICATIONS

A look at some recent applications of analog semicustom ICs suggests their virtuosity. All of the chips mentioned, as well as others, are summarized in the directory beginning on p. 44.

A high-voltage chip from Exar provides a solution to a common telecommunications application: a compact telephone-line controller that can connect directly to the line. The direct connection is achieved through the chip's npn transistors, which have a breakdown voltage of 75 V. Both the switching power supply and a regulator have been integrated. The resulting chip contains a temperature-compensated voltage reference, an oscillator, an error amplifier, a comparator, the output drive transistor, the flyback diode, and short-circuit protection. This integrated version of the application has low-enough power requirements to operate off the low voltage and negligible current available from the phone lines.

Another application combines high-voltage signals with broad-bandwidth response. The video display driver in Figure 3 operates at 175 MHz, driving a full 50-V transition into a 6-pF load in only 2 ns. This circuit contains 147 devices and consists essentially of an input buffer, a variable-gain stage, a gain control circuit, a differential-to-single-ended amplifier, and a cascode output stage.

The variable gain stage uses a pair of transistors (Q_3 and Q_4) to implement a signal-gain stage. The collector circuits of this pair are controlled by a feedback loop, and the difference in V_{BE} is applied to the gain stage to set the gain. This structure cancels the nonlinearity and drift contributed by the transistors' ohmic resistances and beta variations.

Thermal distortion of less than 1% is achieved by using four temperature-sensing diodes (D_9–D_{12}) connected in series; these diodes produce a voltage that compensates for the change in base-emitter voltage as the emitter followers and the common-emitter stage are heated from the common-base stage.

Some external impedance-matching components (not shown) are used to achieve the maximum bandwidth into the capacitive load. The IC's package, a modified 24-pin DIP with an integral copper heat sink, dissipates 7 W and exhibits a thermal resistance of 6°C/W.

This high-voltage, high-current chip from Tektronix, designated QuickChip 3, currently supports signal levels that are only as high as 34 V when used as a semicustom device; higher voltage levels will be accommodated through some custom techniques until the chip's double-layer metallization has been rated for a full 65 V.

Raytheon's RLA-120 forms the basis of a semicustom high-voltage differential amplifier by using its op-amp cells and sputtered thin-film resistors. Four channels of the differential amplifier shown in Figure 4 were required for a power-supply monitoring application. As long as the slew rate of the op amp cells is not exceeded, the inputs will withstand high voltages, with a maximum differential voltage of 200 V.

Using just two resistor ratios helps to minimize the gain error, thereby avoiding the accumulation of errors from more extensive resistor networks found in other implementations. Dielectrically isolated thin-film resistors, with breakdown voltage ratings of between 100 and 1,000 V, allow the inclusion of on-chip scaling networks whose values track during temperature changes. On-chip resistor networks also reduce the overall component count.

Another IC from Raytheon's RLA series serves as a bandgap voltage reference in a missile system. It uses "zener zapping" instead of the more common laser trimming to control the ratio of current densities in core transistors. Zener diodes, created using the emitter-base junction of npn transistors, are selectively short-circuited by passing a short pulse of high current through them. Thin-film resistors complement this technique because they can withstand the 60-V surge needed to blow the diode junction cleanly. A designer wishing to zap zeners must reserve bonding pads (in this case, nodes A, B, and C) for wafer probes to access the zener-trim networks—one pad for each bit of adjustment, in addition to a "return" pad, which is usually the power supply or ground.

According to Raytheon Semiconductor, the circuit (Figure 5) stability and wide-range operations that IC designers expect from a bandgap reference. Its use has tightened the output voltage manufacturing tolerance to ±1% from a typical 3% to 5%. Levels of current for transistor Q_1, which controls the reference voltage, can be selected by zapping the combination of zener diodes Z_1, Z_2, and Z_3.

The largest smorgasbord of macrocells, circuit components, and protection structures is found on Micrel's MDP8020 (see Figure 1 again). One popular application of these resources is motor control, achieved by connecting the winding of the motor as a load on an H bridge. Alternatively, this bridge can just as easily connect to a high-frequency transformer (100 to 400 kHz) in a switching-mode power supply.

Three half-H bridges form a standard delta or wye three-phase motor driver circuit. Each transistor in the bridge, an n-channel DMOS FET with ratings of 100 V, 200 mA, and 10 Ω, can be connected in parallel with one or three other DMOS FETs for added current capability. Having 16 FETs, 16 level shifters, and 16 drivers on the chip makes it possible for designers to realize a considerable number of drive applications. On-chip functions are combined with a CMOS gate array and high- and low-voltage linear elements, all of which implement the user's proprietary control circuits.

Furthermore, Micrel has added protection against overheating to the MPD8020. Short-circuit current limiting and a bandgap voltage reference form the first line of defense. An overtemperature detection circuit shuts down the chip when its temperature reaches 150°C; with built-in hysteresis, it can restore functionality when the chip gets back down to 85°C. For any given design, these trip points are mask-programmable within the bandgap reference, whose output goes to the detecting circuit. ■

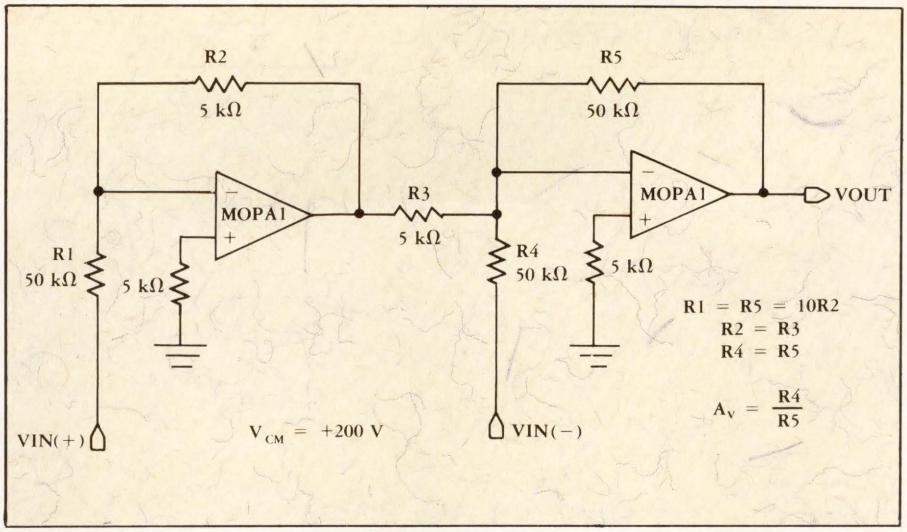

Figure 5. A 200-V differential amplifier is realized by using on-chip sputtered thin-film high-voltage resistor dividers to reduce the amount of common-mode voltage seen by the op amps.

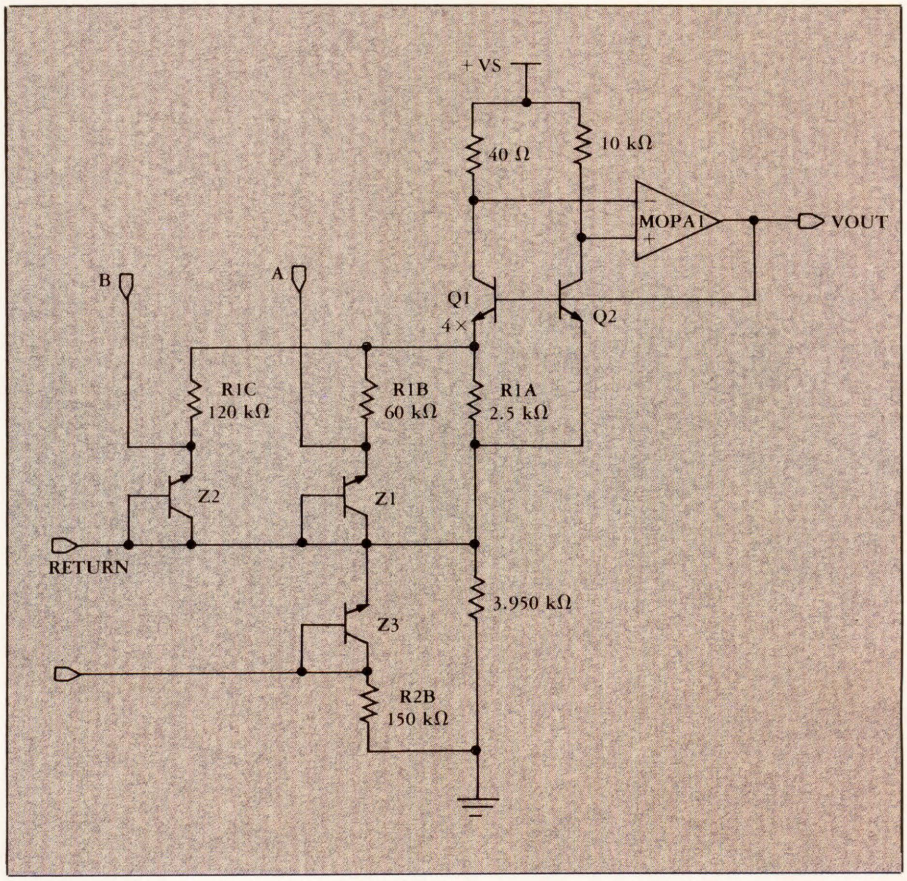

Figure 6. Selectively "zapping zeners" in this voltage reference circuit during manufacture has resulted in a tightening of the output voltage tolerance.

About the Author

Alan P. Ramsey *is a free-lance writer and communications consultant to client companies in the high-tech industry. He has worked in the industry since the early 1960s and in various writing and communications management positions for corporations and their agencies, including GE's Aerospace Division, Calspan Corp. and General Dynamics, before starting his own consulting service. Ramsey holds a British diploma in business from Kilburn Polytechnic Institute (London) and an AAS in electronics from Lindsey Hopkins (Miami, Fla.).*

National Semiconductor

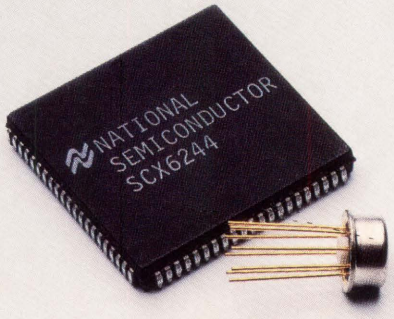

Digital and analog functions in standard packages

Similar functions in a standard cell available for packaging

We've mixed analog with digital to help you make it small in ASIC.

NATIONAL INTRODUCES OP AMP, COMPARATOR, ANALOG SWITCH, AND BAND GAP REFERENCE CELLS DESIGNED TO CUT YOUR ANALOG REAL ESTATE DOWN TO SIZE

Now you can bring the finest in analog performance right into the heart of your digital cell-based design.

Because now, the people who wrote the book on linear offer you industry-standard functions in analog cells.

With all the quality and reliability you've come to expect from National. Plus all the design and high-yield manufacturing expertise it takes to ensure you get the cells you need when you need them.

EIGHT CELLS YOU CAN BUILD ON

Our cell library now includes three general-purpose op amp cells, based on our popular LM324, with gain bandwidths ranging from .75 MHz to 3 MHz.

We've also developed the A021 and A022 Comparators, which feature two different input common mode ranges. Both cells are similar in performance to our industry-standard LM339.

Our A241 Analog Switch is designed to handle signal voltages equal to, or less than, VDD while exhibiting a lower on-resistance than most switches in our AH50XX, LM, or LF families.

As you'd expect, we've also added a 2.5V Voltage Reference, plus resistors. And all our analog cells operate at the standard 5V.

NATIONAL'S CELL LIBRARY INCLUDES:

Op Amps
- Wide input common mode (0.3V to VDD – 1.5V)
- Up to 3 MHz gain bandwidth product

Analog Comparators
- High input common mode range (1.5V to VDD – 0.5V)
- Low input common mode range (0.3V to VDD – 1.5V)
- Less than 70nS response time

Voltage Reference
- 2.5V Band Gap Reference
- Tempco under 100 ppm/°C

Analog Switch
- Analog signal range 0V to VDD
- Low on-resistance 100 ohms

All are available in military, commercial, and industrial temperature ranges and in a variety of high-density surface-mount packages.

PROVEN CELLS YOU CAN WORK WITH

We're already working with digital ASIC designers who need basic analog functions without making basic analog trade-offs in size, weight and performance.

One major automotive subsystem manufacturer, for example, is currently testing our op amp and comparator cells in a revolutionary anti-skid braking system under severe environmental conditions.

Wherever there's a need for mixed digital and analog solutions that improve efficiency and reliability by eliminating bulky, power-hungry multi-chip configurations, there's a need for analog cells from National.

FUTURE CELLS YOU CAN COUNT ON

All our analog cells are fabricated in our proprietary M²CMOS process. Which guarantees you three important competitive advantages: outstanding low-power performance, noncopyable functional implementation, *and* a steady stream of other proven National analog products into our cell library.

The result? Your decision to go with National's analog solutions will continue to pay off long after the decision is made.

CELL EXPERTS YOU CAN TALK TO

We can show you how a little analog functionality can make a big difference in your digital ASIC design. For details, call our Applications Hotline at (408) 721-6247, or write:

National Semiconductor
ASIC Solutions
M/S 23-200
P.O. Box 58090
Santa Clara, CA 95052-8090

National Semiconductor

DIRECTORY OF HIGH-VOLTAGE ANALOG SEMICUSTOM

COMPANY	PRODUCT (TECHNOLOGY)	ACTIVE COMPONENTS					PADS		RESISTORS	
		n (<0.1 A)	p (<0.1 A)	n (>0.1 A)	p (>0.1 A)	Digital gates	I/O	Power only	Diffused (#)	Impl. (#)
AT&T MICROELECTRONICS *1 Oak Way* *Berkeley Heights, N.J. 07922* S.K. Smith Product Planner (215) 439-7531	ALA 400 series (complementary bipolar)	66–100	66–100	2	2	0	38–44	0	0	232–656 (100 Ω–10 kΩ)
CHERRY SEMICONDUCTOR *2000 South County Trail* *East Greenwich, R.I. 02818* Gary Nielsen Marketing Manager (401) 885-3600	Genesis 8000 (bipolar)	42	26	18	8	0	23	0	427 (508 kΩ total)	0
DATA LINEAR *491 Fairview Way* *Milpitas, Calif. 95035* Valentino Liva ASIC Manager (408) 945-9080, x466	DL104 family of standard-cell arrays (bipolar, dielectrically isolated)	104–650	104–650	—	—	—	14–34	0	312–1930	0
	SP1104 family ofmacrocell arrays (bipolar, dielectrically isolated; rad-hard)	216	216	6	6	—	24	0	800 kΩ uncommitted	NiCr
ELECTRONIC TECHNOLOGY CORP. *525 E. 2nd St.* *Ames, Iowa 50010* Liz Partrick National Sales Manager (515) 233-6360	ETC-X (bipolar)	208	82	4	0	0	40	0	1672	0
EXAR CORP. *750 Palomar Ave.* *Sunnyvale, Calif. 94086* Surjit Nijjer Product Marketing Manager (408) 732-7970	X-100 Masterchip (bipolar)	N.s.	N.s.	4	N.s.	0	18	0	N.s. (615 kΩ total)	N.s.
INTEGRATED CIRCUIT SYSTEMS INC. *1012 W. 9th Ave.* *King of Prussia, Pa. 19406* Bruce J. Rogers National Sales Manager (215) 666-1900	3-μm DLM Si-gate CMOS	N.a.	N.a.	N.a.	N.a.	>100	N.a.	N.a.	N.a.	N.a.

CAPACITORS			MACRO-CELLS	NPN OR NMOS PERFORMANCE		MIN. BV_{CEO} (V)	SUPPLY VOLTAGE (V)	TEMP. RANGE (°C)	PACKAGES	NOTES
Fixed (#)	Prog. (#)	Range (pF)		f_T (MHz)	h_{FE}					
0	7–14	1–7	0	350	40–250	33	2.5–33	N.s.	Plastic and ceramic DIP, PLCC, LCC, plastic SOJ, SOIC—narrow and wide body	Complementary vertical npn and pnp; 2 JFET tiles on ALA-400
0	0	—	Programmable bandgap reference (1)	N.s.	N.s.	>50	1–50	N.s.	Standard epoxy and ceramic, DIP, flat pack, transparent molded	Two dedicated 28-V zeners
260 (1.6 pF)	104–650	0.25–2.5	—	1,000	150	>35	35	−55 to +125	Cerdip, side-brazed, LCC, QLCC, others on request	—
14	14	5, 8, 18	Bandgap references (2 mA), op amp (30 MHz, 72 dB, 5 mA; 80 MHz, 110 dB, 20 mA), comparator (25 ns, 0.4 mA), transimpedance amplifier (400 V/μs, 100 mA)	—	—	35	1.5, 2.5, 5	−55 to +125	Kit part	Macrocells can be lifted for system layout
8	0	0–10	—	400	158	75	N.s.	N.s.	Plastic and ceramic DIP, SOIC, LCC, PLCC	—
—	—	—	—	N.s	N.s.	75	N.s.	N.s.	N.s.	—
N.a.	N.a.	1–100	Configurable op amp, R-2R, A/D, D/A, comparator, bandgap reference, balanced modulators, power drivers, switches, current sources	N.s.	N.s.	>35	35	−55 to +125	Plastic and ceramic DIP, to 84 pins; PGA	Cells assembled from library on industry-standard workstations; fabricated to order

Here are the keys to faster design of military ASICs...

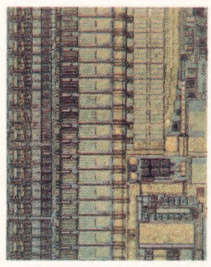
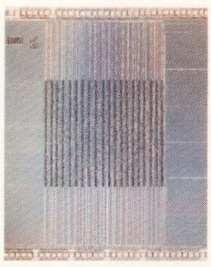
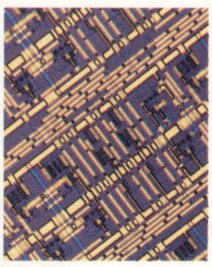
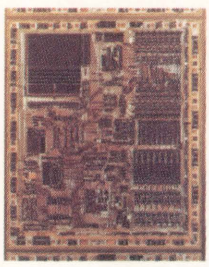
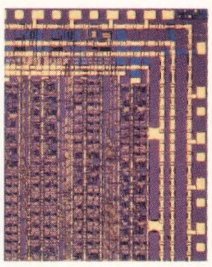
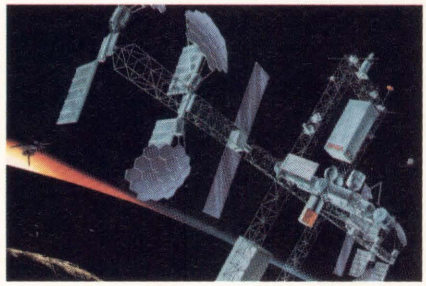

One workstation, one company, for all your military/space designs.

Remember when semicustom IC designing was impractical because your technology options were too limited, your CAD tools too new and turnaround too slow?

Times have changed! Now Harris makes designing military/space ASICs — in the broadest range of technologies — easier than ever. Because you use the same workstation and design tools for all.

Libraries of solutions: Our GaAs DIGI-II library offers nearly 50 fully characterized cells that are ECL, TTL, CMOS or GaAs compatible. Or...in silicon our advanced HSC 2.0-micron rad-hard CMOS library provides cells and macros you can intermix. Resident libraries also exist for CMOS and bipolar, analog, and digital plus analog/digital mixed technologies.

Hardware and software: Our open-system software CAE/CAD toolset fully integrates with Harris/Masscomp workstations, and meshes easily with other UNIX-based platforms. Daisy and Mentor support is also available.

On to manufacturing: You handle the front-end; leave the back-end to Harris. We carry your ASIC design through to mask, then manufacturing — with the screenings you want and most popular packages: metal flat-pack, side-brazed DIP, chip carriers, and more.

Guaranteed survivors: Now you can create ASICs that tap the high speed of GaAs, or the low power of CMOS, and are guaranteed to survive military and space environments. And you can do it faster.

For more on military/space silicon and GaAs ASICs, call for our Rad-Hard/Hi-Rel Data Book. In U.S. phone 1-800-4-HARRIS, or in Canada: 1-800-344-2444; for GaAs ask for Ext. 1525, for CMOS ask for Ext. 1925.

IN MILITARY ASICs,
THE NAME IS

HARRIS

Harris Semiconductor: Analog - CMOS Digital
Gallium Arsenide - Semicustom - Custom

HARRIS

CIRCLE NUMBER 15

DIRECTORY OF HIGH-VOLTAGE ANALOG SEMICUSTOM

COMPANY	PRODUCT (TECHNOLOGY)	ACTIVE COMPONENTS					PADS		RESISTORS	
		n (<0.1 A)	p (<0.1 A)	n (>0.1 A)	p (>0.1 A)	Digital gates	I/O	Power only	Diffused (#)	Impl. (#)
MICREL INC. *1235 Midas Way* *Sunnyvale, Calif.* *94086* Marvin Vander Kooi Director of CMOS/DMOS (408) 245-2500	MPD8020 (CMOS/DMOS/bipolar)	See Notes	See Notes	See Macrocells	See Macrocells	200	78	40	(1–250 kΩ)	(1–750 kΩ)
PLESSEY SEMICONDUCTORS *1500 Green Hills Rd.* *Scotts Valley, Calif.* *95066* Richard Padovani Strategic Marketing (408) 438-2900	MV series (bipolar)	60–208	2–5	0	0	0	16–36	0	100–428	68–280 (1.153–4.469 MΩ total)
POLYCORE ELECTRONICS INC. *1107 Tourmaline Dr.* *Newbury Park, Calif.* *91320* S.K. Leong Vice President (805) 499-6777	Maxi-chip MX-21 (bipolar)	6	3	2	0	0	10	0	24	1 (54 kΩ total)
RAYTHEON COMPANY *Semiconductor Division* *350 Ellis St.* *Mountain View, Calif.* *94039* Bruce D. Moore Linear Applications (415) 966-7757	RLA series (bipolar)	97–146	75–85	3–4	0	0	24–44	2	0 (see Notes)	0 (see Notes)
TEKTRONIX INC. *P.O. Box 500* *Beaverton, Ore. 97077* David Bernel Marketing Manager (800) 835-9433, x100	QuickChip 4 (bipolar)	294	174	6	0	User-selectable	66	6	0 (1.38 MΩ total)	N.s.
	QuickChip 3 (bipolar)	198	48	12	0	0	28	0	0 (1 MΩ total)	926
	QuickChip 2 (bipolar)	150–524	82–240	3–12	0	0	24–70	0	0 (2.2 MΩ total)	N.s.

CAPACITORS			MACRO-CELLS	NPN OR NMOS PERFORMANCE		MIN. BV_{CEO} (V)	SUPPLY VOLTAGE (V)	TEMP. RANGE (°C)	PACKAGES	NOTES
Fixed (#)	Prog. (#)	Range (pF)		f_T (MHz)	h_{FE}					
(40 pF total)			16 floating 100-V, 200-mA DMOS FETs; 16 100-V p- and n-channel level shifters; 12 TTL/CMOS I/O buffers; 16 logic predrivers; 3 configurable gain cells; unity-gain buffer; bandgap reference; programmable master bias; voltage doubler; low-voltage pass regulator; high- and low-level current mirrors	N.s.	N.s.	100	20–100	Commercial, industrial, military	Commercial, military, power	Unapplied n- and p-channel transistors from gate array can be used for general-purpose analog circuits; dedicated zeners, various voltages
4–8	0	5	0 (see Notes)	350	100–400	40	Up to 40	Commercial, industrial, military	Variety, from plastic DIP to SOIC	Extensive library of linear macro designs (4/20-mA linear array cells each contain the small npn's, 22–88 monistors, and resistors to support library macros)
0	0	0	1-A npn Darlington (1)	500	150	80	Up to 80	N.s.	Plastic and ceramic DIP, LCC, flat pack, unpackaged dice	Dedicated 1-A diode, 80-V BV dedicated 5.8-V zener
8–16 (for amplifier compensation only)	0	8.5	Gain cells (8–15) configurable as op amp or comparator (some as optional input amplifier)	400	250	32	3–32	−55 to +125	Plastic and ceramic DIP, LCC, PLCC	Thin-film sputtered SiCr resistor arrays standard
16	16	0–3	0	5500	85	32	<32	−55 to +125	Plastic and ceramic DIP, quad, PLCC, hybrid	Thin-film NiCr resistors with laser trimming available
16	12	0–3	0	2500	90	34 (see Notes)	4–34	−55 to +125	Plastic and ceramic DIP, quad, PLCC, hybrid	Thin-film NiCr resistors with laser trimming available; double-layer metallization awaits approval for 65-V breakdown rating
20–72	20–72	0–3	0	5500	85	32	<32	−55 to +125	Plastic and ceramic DIP, quad, PLCC, hybrid	Thin-film NiCr resistors with laser trimming available

"AUTOMATIC ASIC TES
THAT'S WHAT WE
AND WE NEED IT AT EVERY S

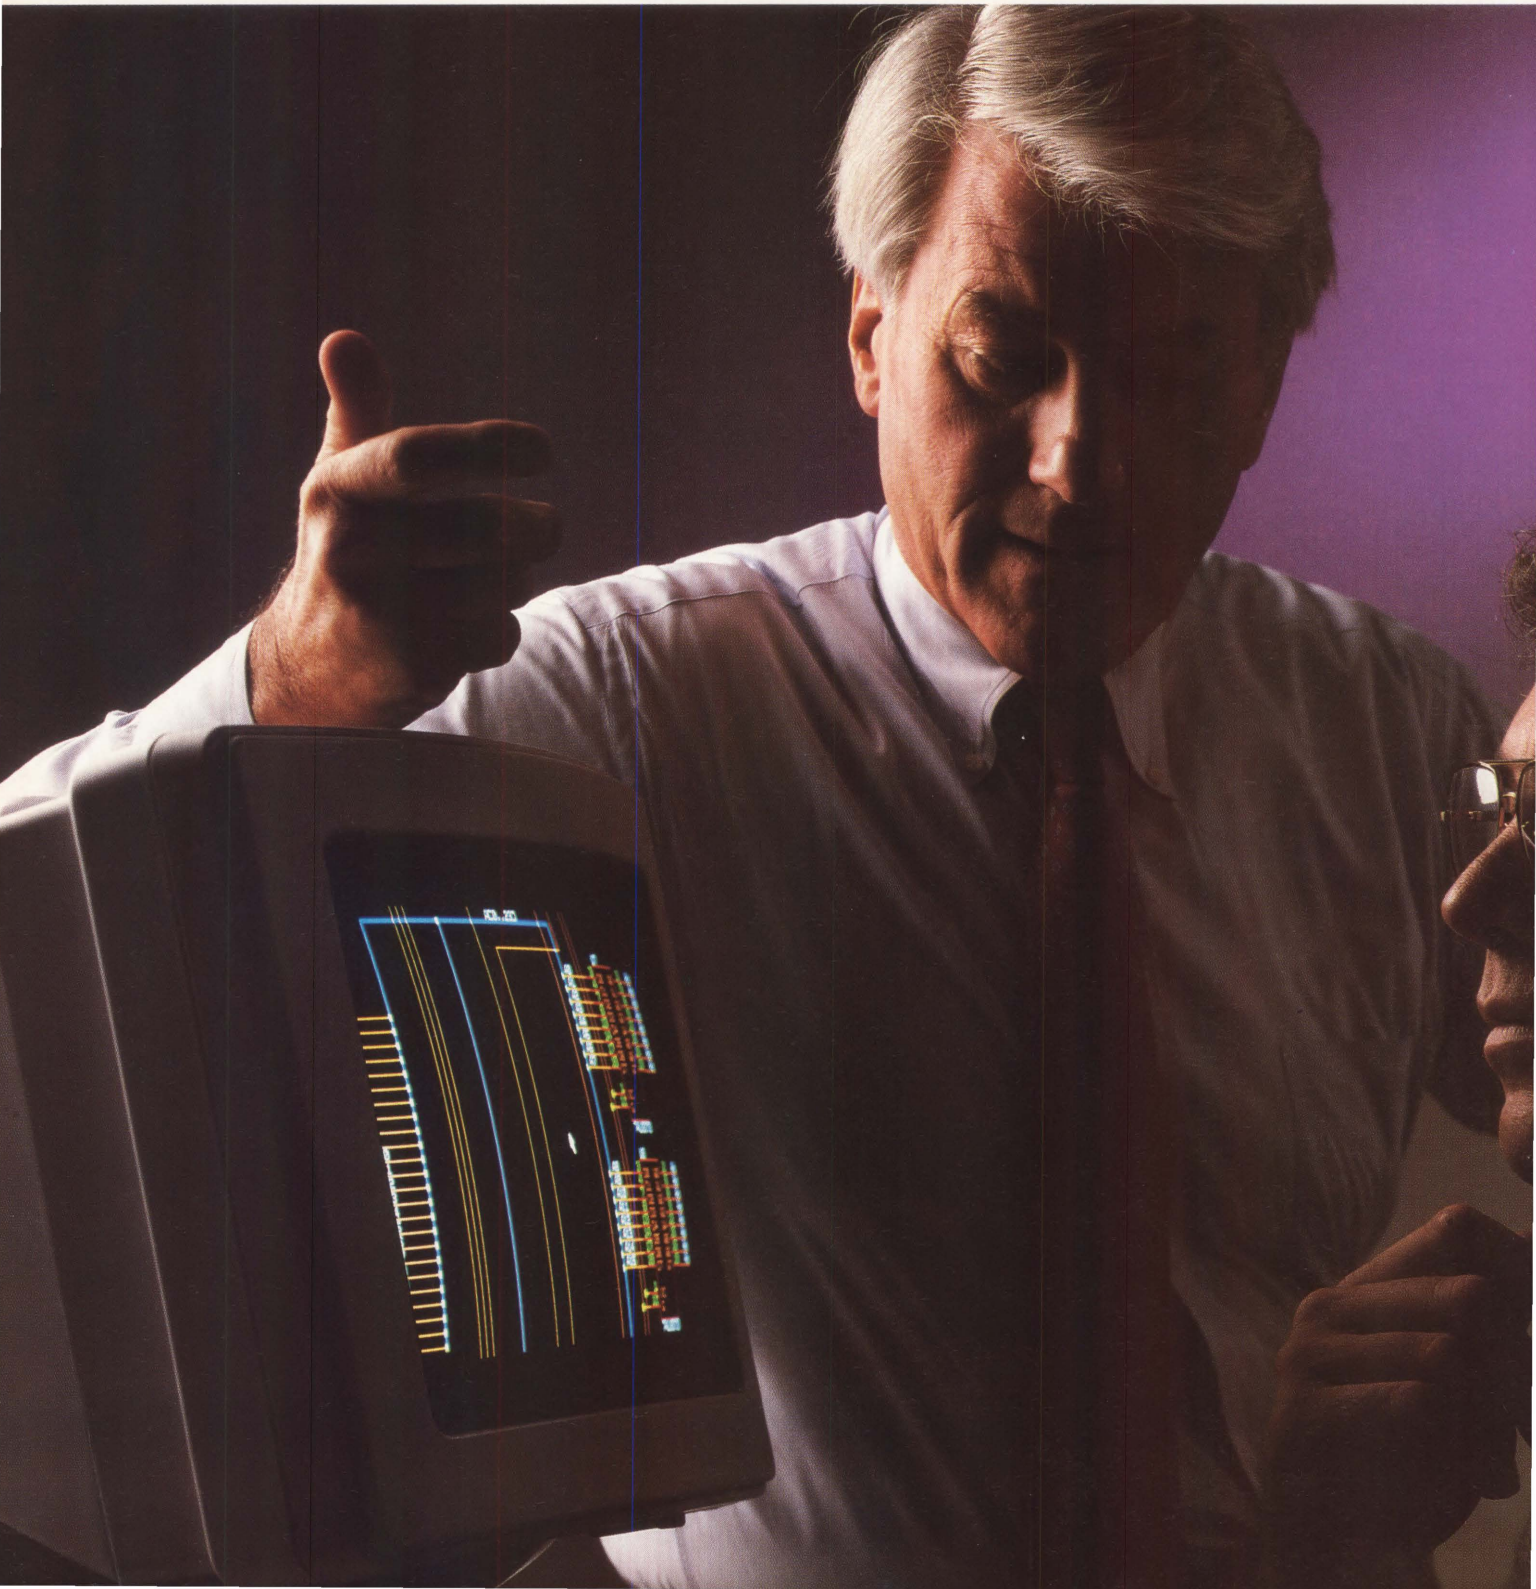

T PROGRAMING.
REALLY NEED.
TAGE IN THE PROCESS."

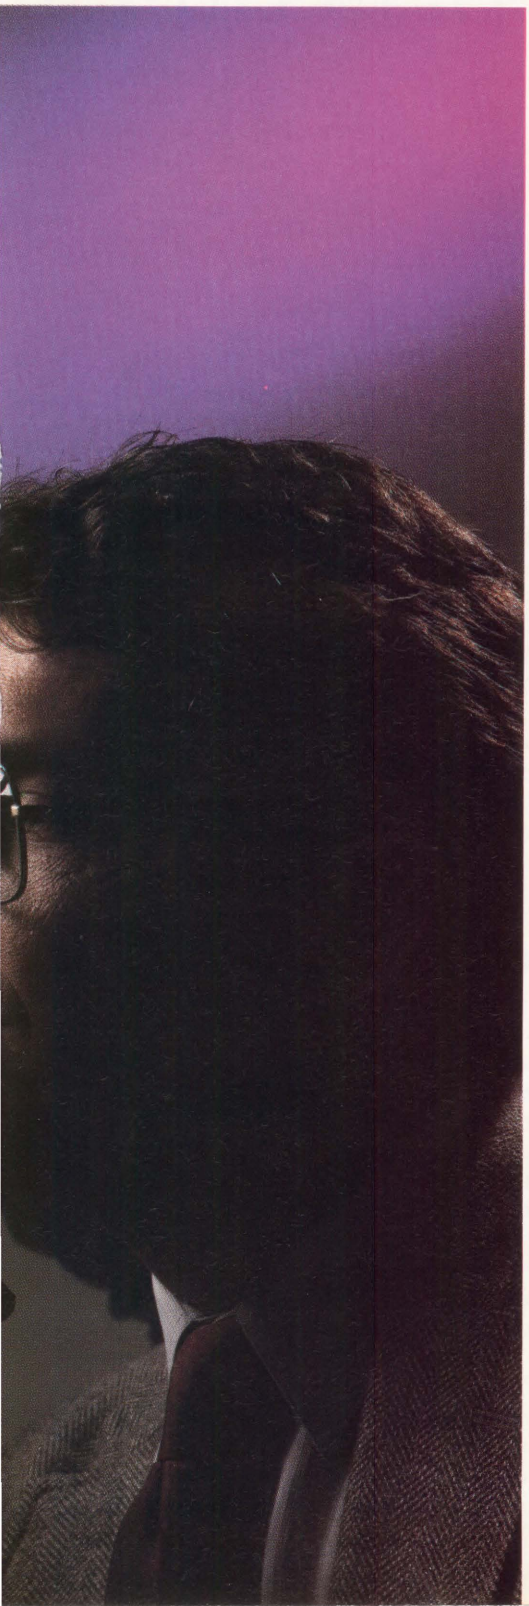

WE COULDN'T AGREE MORE.

Testing ASIC devices creates a whole new set of problems. Developing test programs, characterizing, verifying and debugging are, at best, unwelcome and time-consuming for ASIC designers. ASIC vendors can test the silicon, but not the custom functionality of the ASIC prototype. Traditional test approaches and traditional ATE merely compound the problem.

At ASIX Systems our focus has always been exclusively on ASICs. From the start we took an entirely different approach to solving the unique ASIC test problems. We saw immediately that adapting existing ATE to try to fit to needs of ASICs didn't make sense. Designing a totally new, focused ASIC test system did. Not only did programs need to be automated, they needed to be developed from the design data base and menu-driven, so changes would be simple to make. And the test system itself had to be easy to use, designed for its particular environment, and a cost-effective alternative to the huge, expensive, complicated ATE.

TEST SOLUTIONS FOR THE WHOLE ASIC COMMUNITY.

Our unique perspective allowed us to understand that the ASIC world is not so much Design Engineers, Test Engineers and Quality Engineers working independently. It's more a "community" of specialists whose tasks are intrinsically linked. So we made sure that we could provide another crucial element. Communication. In order to capture the vital time-to-market edge, what ASIC end users and vendors really need is the opportunity to use the same test programs and the same tester. That's what gives both environments a common frame of reference and the chance to generate test programs automatically from the design data base. That's the ASIX-1 family of test systems.

ASIX-1: ASIC TEST SYSTEMS THAT MAKE SENSE.

We don't have the room here to tell you everything the ASIX-1 family has to offer. But here are a few things to think about: automatic, menu-guided programming; data base management; ATE architecture and flexibility at an affordable cost; 256 true I/O pins; "zero footprint"; fully integrated PMU; automatic calibration; simple fixturing; no cabling; high MTBF. That's enough. If you're testing ASICs you know you have to see for yourself what the ASIX-1 can do. And the sooner the better. ASIX Systems Corporation • 47338 Fremont Blvd • Fremont, CA 94538.

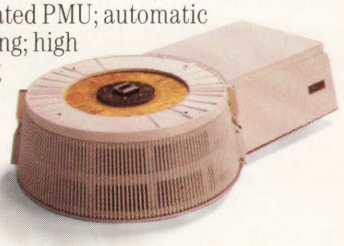

CALL: 1-800-FOR-ASIX

ASIX SYSTEMS

REAL ASIC TEST SOLUTIONS.
FROM THE REAL ASIC TEST COMPANY.

CIRCLE NUMBER 16

CELL SYNTHESIS *in action*

CELL synthesis tools produce transistor-level layouts of complex logic functions automatically, directly from a net list or a schematic. They are technology-independent and permit symbolic editing. Further, they make use of existing algorithms for leaf cell design, placement, routing, and compaction.

The cell synthesis tools can generate optimized rows of p- and n-channel transistors and permit the user to specify the location of ports and power rails. As such, cell synthesis tools are especially suited to the design of MSI-scale cells that contain little or no regular structure. If the methodology is used intelligently, the resultant designs can be superior to those possible with the orthodox standard-cell approach. Cell synthesis also can be used to design VLSI-scale chip layouts, although it will be some time before the technique can actually supplant standard-megacell design entirely—preliminary designs of several-thousand-gate complexity are still in the embryo stages at some of the cell synthesis tool vendors. On the other hand, it can be an ideal complement to a standard-cell or silicon compilation system.

We have asked six vendors to use their system to automatically generate IC layouts for the same circuit and design rules. We will examine the results qualitatively, describing the complex algorithms at work. We will also compare the synthesized cells with a standard-cell layout for the same circuit. For the task assigned, the cell synthesis tools produced layouts that were one half to three quarters the size of a standard-cell layout.

*Now a free-lance technical writer.

Cell synthesis is not to be confused with logic synthesis, silicon compilation, or module generation. Logic synthesis refers to the generation of an optimized net list from an initial hardware description (de Geus and Cohen, 1985). Layout is not part of logic synthesis, whereas cell synthesis is directly concerned with a geometric layout. Of course, it's useful to perform logic optimization before embarking on cell synthesis, and we may reasonably expect integrated logic and cell synthesis systems to appear in the near future.

Silicon compilation, on the other hand, refers to a methodology rather than a process. Originally, a silicon compiler was conceived as a "program with the translational attributes of a compiler, yet whose object is an integrated circuit layout, rather than another software program" (Peskin, 1982). The term *silicon compilation* has come to refer to a complete design system, differentiated from standard CAE and CAD systems by its use of layout compilation as well as library elements.

Module generation is another capability that is offered by many silicon compilers. *Module generation* refers to the process of abutting a small number of layout blocks for functions that contain many repetitive elements, such as multipliers, PLAs, and memory blocks (Meyer, 1987). Module generation works at the block level; cell synthesis, in contrast, generates a transistor-level layout that can subsequently be used as a block.

In fact, a designer could use cell synthesis to design the building blocks for later use by a module generator. Therefore in an integrated design system, logic synthesis could precede cell synthesis, and cell synthesis in turn could precede module generation. All three processes could be embedded in a silicon compilation system.

■ TRANSISTOR-LEVEL LAYOUT

In order to better gauge the actual workings of the current commercially available cell synthesis systems, we asked a number of vendors of cell synthesis systems to generate a CMOS layout for a transistor-level circuit containing 60-odd transistors. The final circuit presented to the vendors, shown in Figure 1, is a version of a circuit that provides synchronous square-wave division by an odd integer (Hsieh, 1985). This version produces a single pulse every three cycles.

The particular circuit was chosen for the following reasons. First, it contains transmission gates, which limit the use of shared diffusion regions, as employed by simple stack-based synthesis algorithms. Second, the circuit contains a number of

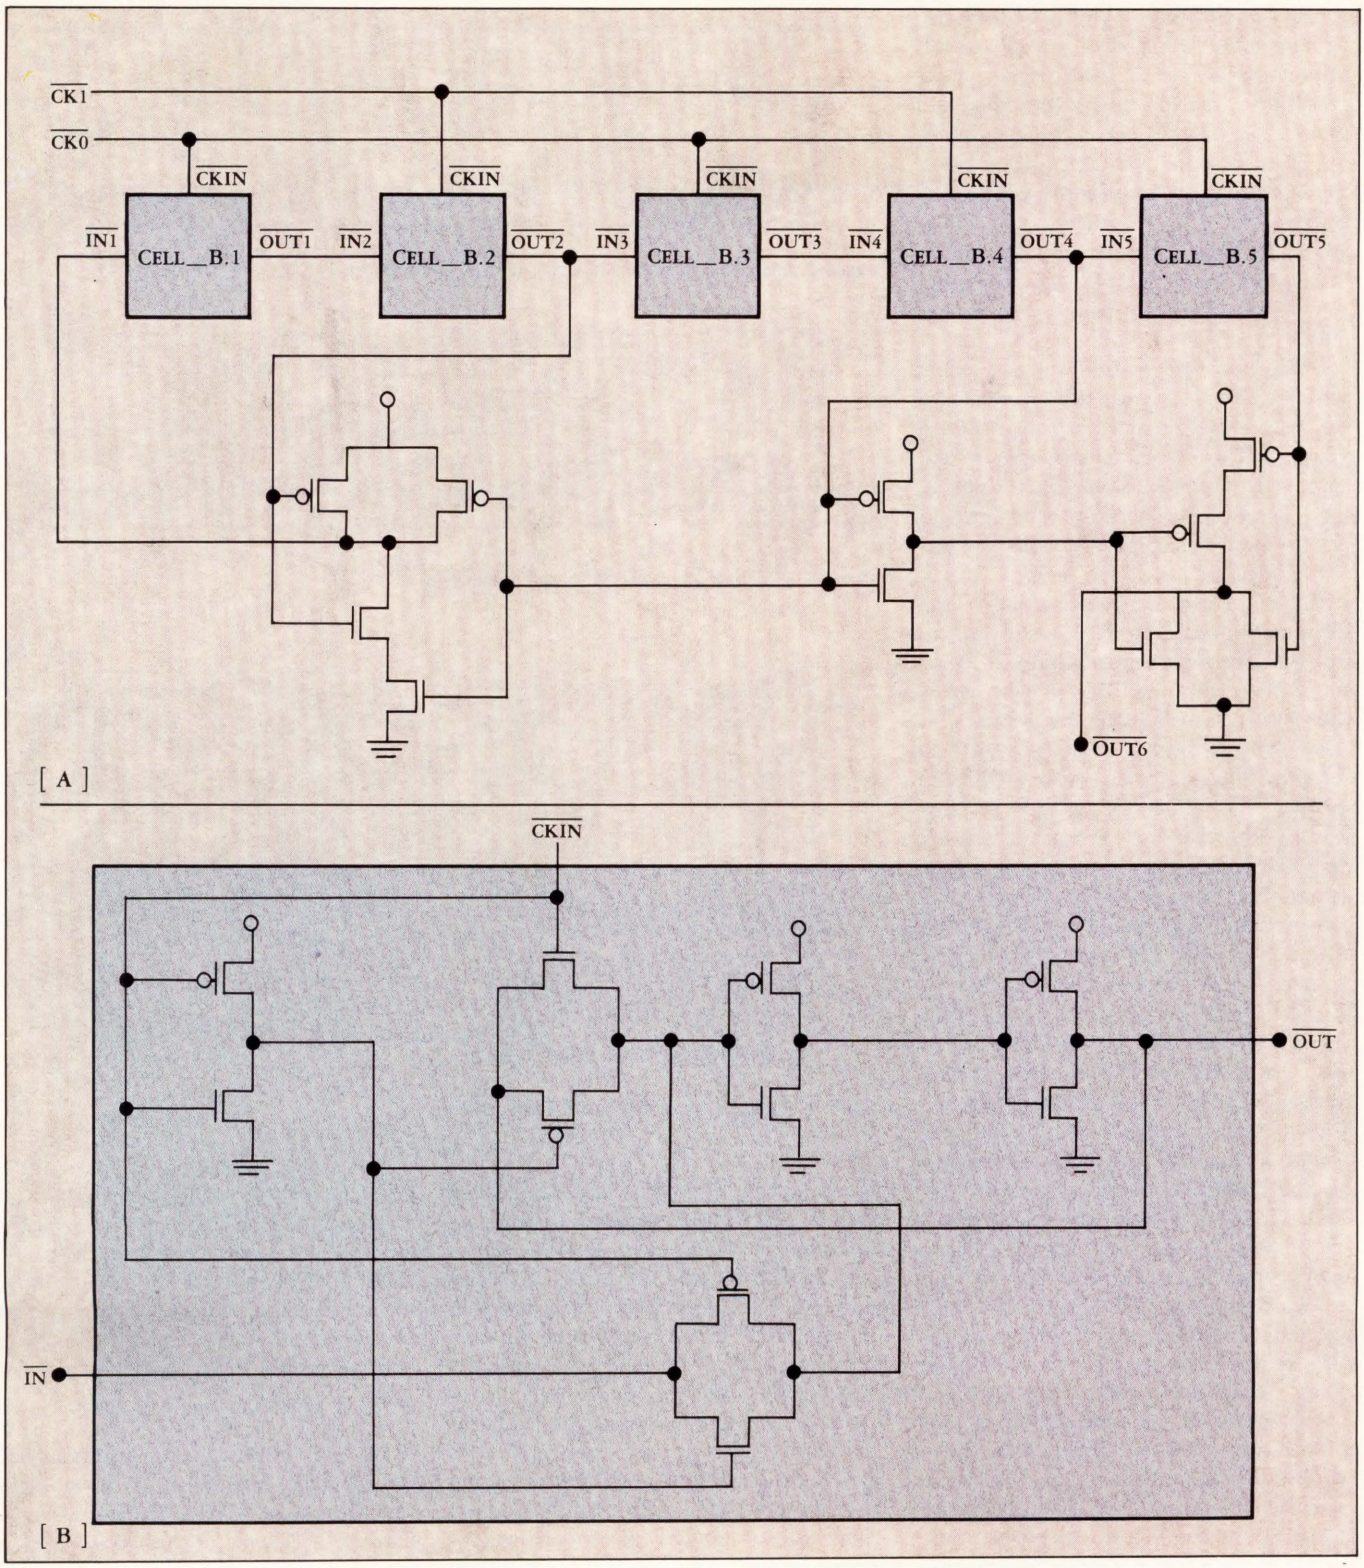

Figure 1. The test circuit for cell synthesis (A), built up using the same subblock (B) five times, contains about 60 transistors.

repetitive elements with some long wiring paths between them. As a consequence, the circuit contains some hierarchy, which simplifies design entry, and some irregularity, which leads to some signal paths that are difficult to route. Third, the circuit has only two inputs and one output, which simplifies simulation. Finally, the circuit is large enough to make manual layout difficult but small enough to permit the completion of cell synthesis in less than a day. In fact, several vendors reported that the design entry took less than half an hour, with process porting taking approximately four times as long. The actual synthesis time is typically about two minutes for a circuit of this size.

Six vendors readily agreed to participate in the design rally: Andrew Tickle Associates, AT&T, Caeco, Emerald Design Systems, Integrated Silicon Systems (ISS), and Silicon Compiler Systems Corp. (SCS). Motorola produced the layout generated by an SCS system.

AT&T generated three different layouts for the circuit, one of which is shown in Figure 2. The other two layouts use folded

and snaked diffusion regions and were two to four times larger than a standard-cell implementation of the same design. AT&T is still investigating the use of folding and snaking, since it is theoretically possible to produce more compact layouts with these techniques. The algorithms, however, are expected to be much more complicated. Incidentally, none of the systems from AT&T are for sale, but we appreciate the company's taking part in this design rally.

Caeco has a schematic front end for its system. The cell synthesizer works with the rest of Caeco's tools, which include a geometry editor, a schematic editor, and automatic layout software, plus a link to SPICE. Caeco's software runs on Sun and Apollo platforms. Figure 3 shows the layout from Caeco.

Emerald Design Systems has ported the software from Andrew Tickle Associates to workstations from Silicon Graphics, Sun, and Apollo. Andrew Tickle, who wrote the Cellgen system, sells the software on an 80386 platform. Cellgen is also marketed by Silvar Lisco. Andrew Tickle Associates and Emerald Design Systems worked together to produce a layout, shown in Figure 4, for the example circuit using a Silicon Graphics platform.

ISS, a small company with a 386-based IC design system, has a tool suite that is similar to the Tickle software in that it includes a pad ring generator and a GDS II driver. However, ISS provides a geometry editor, whereas the Tickle software uses symbolic editing. ISS provided us with the line plot shown in Figure 5. The layout is produced on a dot-matrix printer, using a vectorized output translator from the HP-GL driver supplied with the ISS software.

SCS offers a cell synthesis module embedded within its Genesil silicon compilation environment. The module is the only cell synthesis system that produced a layout with multiple p- and n-type transistor rows. The resulting cell is larger than it would otherwise be; however, it is also thinner and taller, which may sometimes be desirable. Motorola employed the Genesil system to generate the layout shown in Figure 6. (Motorola also generated a layout for its own 2-μm process, to demonstrate the process portability.)

VLSI Technology Inc. also generated two standard-cell layouts for this circuit, which served as a rough comparison for the cell-synthesized layouts. A single-row and a double-row standard-cell layout are shown in Figure 7. The company's 2-μm rules were used, which are reported to be very similar to the Orbit design rules given to the cell synthesizers.

Although the various layouts are not to

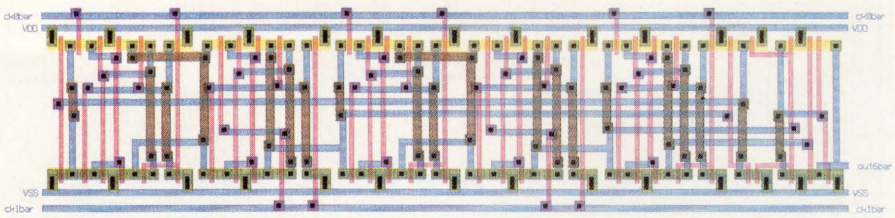

Figure 2. One of the three layouts produced by AT&T's software.

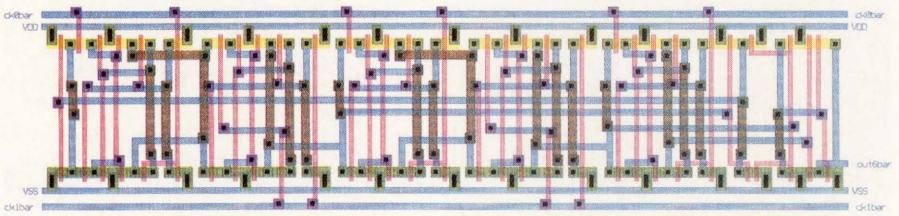

Figure 3. The layout generated by Caeco's software.

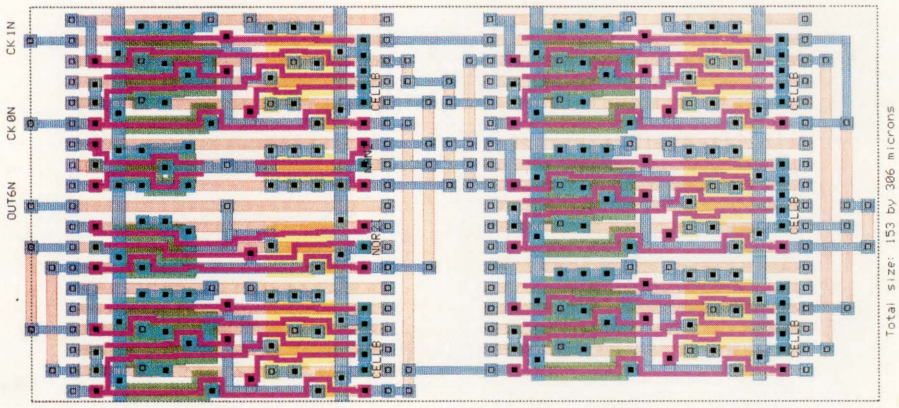

Figure 4. The layout generated by the Tickle software.

the same absolute scale, their size relative to the individual transistors in the layout can provide a rough estimate of the size of the different layouts submitted. It is important, however, to keep in mind that these results do not indicate the absolute performance capabilities of the various design systems. In other words, this circuit example is not intended to be a benchmark of these systems. Each system is suited to different design demands, and the size and type of circuit is very likely to moderate strongly the quality of the generated layouts. The differences are due to algorithmic variations between the different systems. For example, some allow the creation of multiple p- and n-channel transistor rows. Some place all the routing between the pand the n transistor rows, and some route outside as well as between the diffusion regions. In addition, the systems are endowed with a variety of different bells and whistles, and they support differing degrees of designer interaction. Also, the systems support different types of user and simulation interfaces. As a result, cell size alone cannot be used as an absolute criterion for evaluation. The needs of any particular design may in fact make the use of a number of these systems for one design a necessary methodology for optimal results.

■ STACK-BASED SYNTHESIS

In cell synthesis, the transistors are usually placed in two rows: p-type transistors at the top and n-type transistors at the bottom. This structure optimizes the sharing of diffusion regions by devices with common sources and drains. The algorithm used to perform this task is loosely known as a p-and-n-transistor row generator. The software used to perform a complete cell layout is sometimes called a stack-based synthesis system.

Some systems can generate more than one stack of p and n transistors, in which case the rows of p- and n-type devices typically alternate (forming p-n-p, n-p-n, or p-n-p-n structures). The layout pro-

Capture Your Solutions

With SCHEMA PCB LAYOUT and SCHEMA AUTOROUTER

Introducing The SCHEMA Family Of Integrated CAE Solutions.

Omation, the originators of SCHEMA I and II, announces two new solutions in the SCHEMA family of integrated software for the professional PCB design engineer — *SCHEMA PCB LAYOUT* and *SCHEMA AUTOROUTER*.

Now, you can enter a schematic using SCHEMA II, then process it with SCHEMA-PCB and SCHEMA-ROUTE into a finished printed circuit card — *from start to finish in one easy step*.

SCHEMA-PCB and SCHEMA-ROUTE operate on an IBM personal computer or compatible, yet they offer all the features of design work-stations costing $50,000 or more *for a fraction of the cost*:

- Boards of over 400 equivalent ICs, up to 32″ × 32″, and up to 30 layers
- Resolution on placement and tracking of 1 mil
- Autoplacement
- Selectable track width and pad size/shape (250 increments of 1 mil each)
- Integrated schematic input and back annotation
- Entry of engineering changes at the schematic level
- Autorouting with 4 separate algorithms for accuracy
- Complete manufacturing kit output
- 'Air Gap' and connectivity design rule checking

As always, each member of the SCHEMA family of products comes with a *30-Day Money Back Guarantee*.

Capture your solutions by calling us Toll Free: **1-800-553-9119** (in Texas, 214-231-5167). Ask us about our SCHEMA Family Demo Evaluation Kit.

OMATION

1210 E. Campbell Rd., Suite 100
Richardson, Texas 75081

CIRCLE NUMBER 17

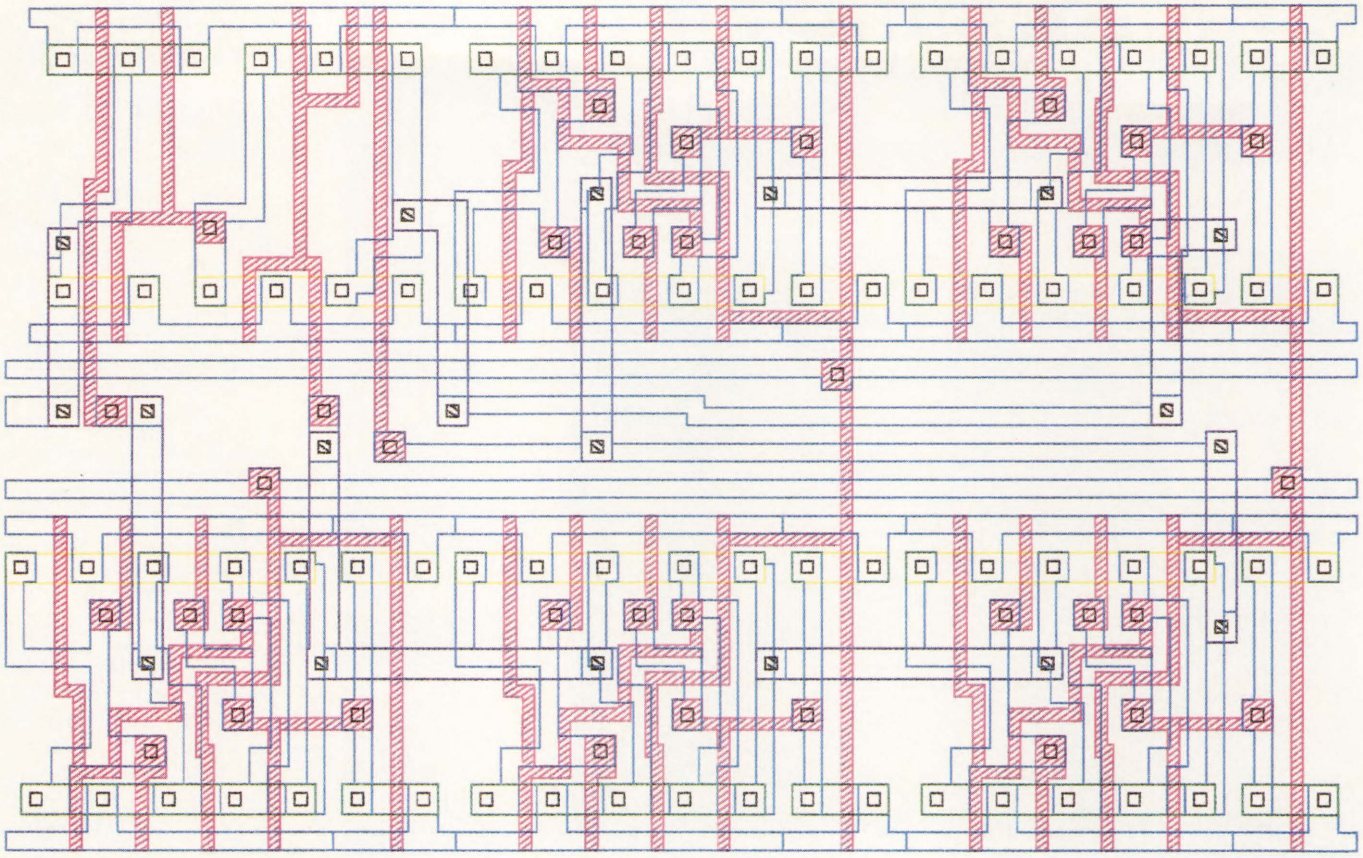

Figure 5. The layout produced by the ISS software.

duced by Silicon Compiler Systems uses three pairs of p-n rows; the number of rows used for the layout can be adjusted by the user. The Tickle software will permit multiple row pairs in a release scheduled for this quarter; Caeco is also planning to provide this capacity.

■ Leaf Cell Layout

Most of the systems use leaf cell generators for the SSI gates within the design. Leaf cell generation is actually a complex problem, as there are many ways even a small number of transistors can be organized. As a result, some of the systems use a library of primitives for the basic SSI functions. The Tickle software, for example, can be used to make an SSI gate; the gate can then be used in a hierarchical manner for larger designs.

■ Placement and Routing

Most systems permit the user to define transistor placement manually. A number of cell synthesizers also contain automatic row generation algorithms. The current level of experience with cell synthesis has restricted the effectiveness of the available placement algorithms, however.

Some systems use a constructive (or force-directed) placement algorithm, which looks at the circuit as a web of springs and lets the transistors fall into places where the strain on any strand in the web is minimized. The user can often adjust the "strength" of particular strands, so that critical paths can be minimized. Force-directed placement is fast but usually inefficient, as the assignment of strength is a nontrivial problem. The Tickle software reportedly supplies a force-directed algorithm, although Tickle advises manual placement in most cases.

Other algorithm designers prefer a min-cut algorithm over force direction. For example, Caeco uses a min-cut approach. The min-cut algorithm cuts the circuit into two clusters such that there is a minimum number of interconnections between the clusters. Each cluster is then placed on one side of the center line. The process is repeated on each cluster, forming additional, but smaller and smaller, clusters. The process continues down to some preselected level, such as the primitive or transistor level. Although this procedure sounds simple, the performance is intimately tied to the way nets with more than two nodes are divided up. Both of these approaches produce placements that can be iteratively improved by swapping pairs of clusters.

After the transistors are placed, their terminals are interconnected by means of a single- or multiple-pass routing algorithm. Most systems use a "greedy" channel routing algorithm on the first pass. In this pass, as many horizontal wiring channels as needed are inserted between or around the transistor rows. Most systems permit the placement of wiring channels above the p-type and below the n-type devices, as well as between the transistor rows. A few systems permit the user to select whether channels can be used both between and outside the transistor rows.

If a multiple-pass router is used with channel routing on the first pass, the routing can be subsequently optimized with a rip-up and reroute algorithm. Routes that cause congestion are removed first, and any freed-up channels are removed by closing up the empty horizontal rows. A Lee (maze) algorithm is then employed to reconnect the unconnected terminals. Here, the layout is treated as a three-dimensional labyrinth, with features as obstacles and the empty areas in the layers as the maze.

Further routing phases can be added. In fact, a six-phase router is probably not atypical. For example, a preliminary power route is sometimes used to minimize the length of power and ground wires. These wires often need to be wider than the signal wires, and of necessity reduce the available chip area. A channel route for the nets that are easier to complete, fol-

NEW From META-SOFTWARE

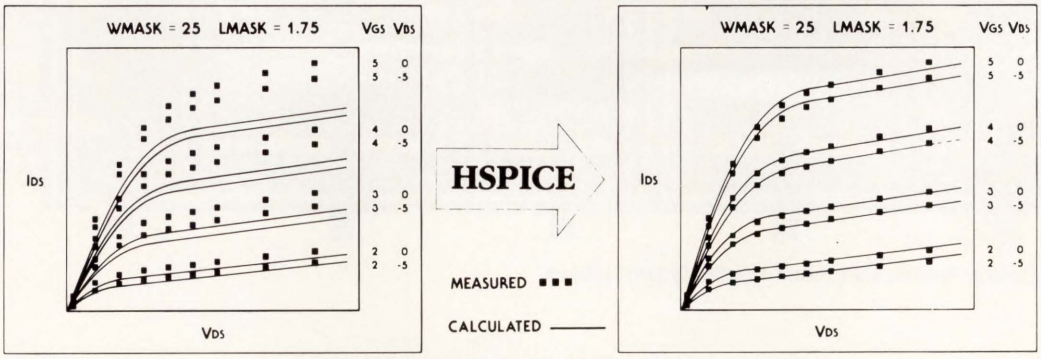

Model Parameter Optimization

Optimizing HSPICE

Meta-Software announces Optimizing HSPICE, incorporating full optimization into the HSPICE circuit simulator.

HSPICE is now a multi-target optimizer that supports all SPICE and HSPICE models. Optimization is included with HSPICE as a new feature, and is available to all HSPICE customers on Software Maintenance at no extra charge.

HSPICE is an integrated solution, optimizing not only DC currents for models, but also capacitance for AC analysis and transient parameters for transient analysis.

HSPICE effectively replaces the functionality of SUXES-10 with full multi-target optimization capabilities. No pre- or post-processing is required.

HSPICE is the result of more than ten years of research in both optimizing algorithms and in user interface. The optimizing function has been integrated into the core of HSPICE, resulting in optimum efficiency. Optimizing HSPICE results will always agree with HSPICE circuit simulation.

Special features of Optimizing HSPICE include:

- Incremental optimization technique, DC, AC and transient optimization
- Uses HSPICE language format
- Model, device, subcircuit and circuit level parameters may all be optimized
- Optimizing Results Targets include Device Currents, Capacitance, Power, Time Delays, Unity Gain Frequency and S Parameters

Meta-Software also offers an extensive Discrete Device Library, HSPLOT graphics post-processor, ATEM process characterization system, Discrete ATEM for characterizing BJTs, MOSFETs, JFETS, HEXFETs and diodes, MetaTestchip™, and the Circuit PathFinder path timing analysis tool.

Other new features of HSPICE:

- Monte Carlo Analysis
- Pole Zero Analysis
- S Parameter Output
- Mixed Domain Analysis
- Instantaneous and RMS power
- Individual element temperature
- Measure Statement
- Multi-Gamma Model for MOS 6 Level
- Small Signal Network Analysis
- ALTER Statement
- Improved BSIM Model
- Data Statement

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 18

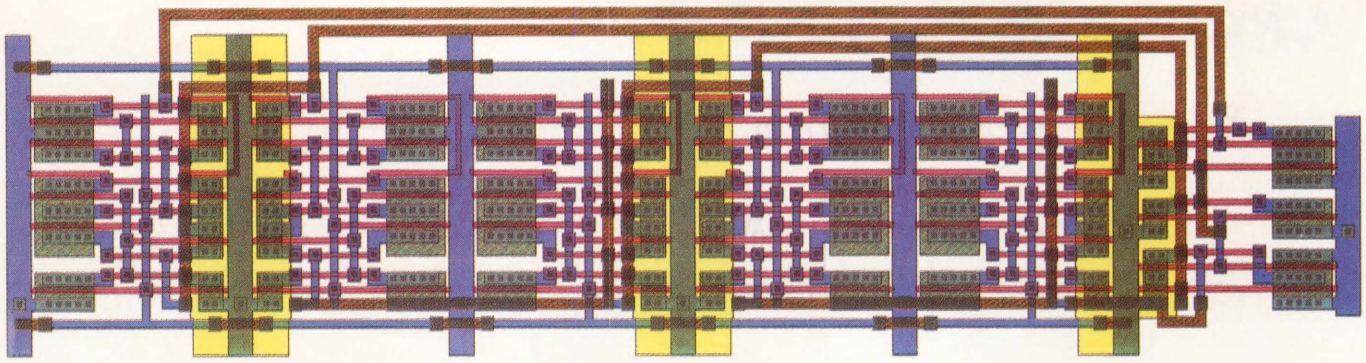

Figure 6. A layout produced by Motorola using Silicon Compiler Systems' software.

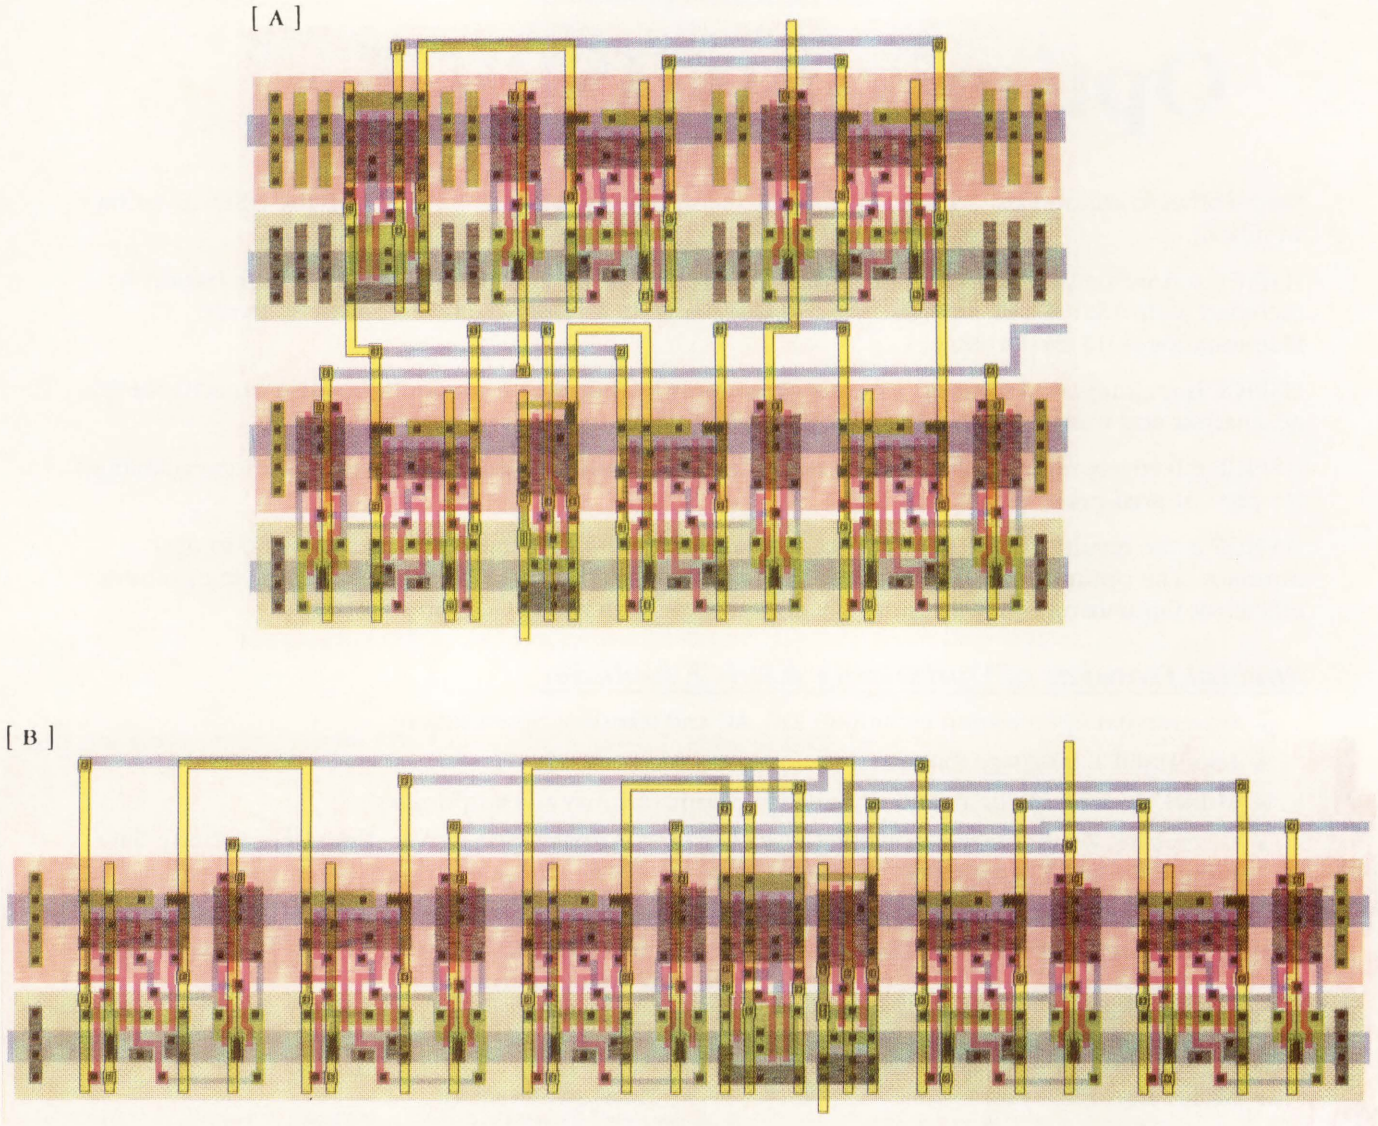

Figure 7. A single-row (A) and a double-row (B) standard-cell layout for the test circuit, performed by VLSI Technology.

lowed by maze, rip-up, and rerouting of the more difficult nets, is a typical succession of routing phases.

It is fairly easy for a user to change the exact structure of any router by changing the cost factors assigned to the various routing possibilities. However, when layout tools are purchased, they usually have pre-assigned "default" costs for each phase of each algorithm. Typically the costs assigned to rip-up, backtracking, and contact insertion are more refined in older tools; many older tools that have been in operation at numerous sites over long periods of time are likely to have had their algorithms continuously fine-tuned and upgraded. Each algorithmic variation is also more suited to different design tasks. For example, the AT&T system tends to create airy routes, which are actually desirable for very large circuit designs. The air

A close look at digital testers will give you a new Outlook for design verification.

The DAS 9200 from Tektronix is a high performance digital test system. But for even higher performance, take a close look at the T-100 from Outlook Technology.

Both products have a maximum recording clock of 2 GHz. But only the T-100 uses intelligent sampling for precise data recording with 100 ps resolution. That's up to five times the resolution of the 9200... the difference between just seeing what happened and finding out why.

Both instruments find timing problems, but only the T-100 can trigger on and track down setup and hold time violations to save countless hours of searching for logic problems.

Both products can be used in automated setups to test boards and chips at high speed and high resolution. But the T-100 can perform up to ten times more tests per hour. And it can act as a 250 MHz pattern generator (stimulus), 250 MHz logic recorder (response), or both.

The T-100 also comes with a friendly human interface, including LogicProbe, a new utility program that makes setup and use faster and easier than ever before.

For a new outlook on digital testing, look into the T-100 family, with prices starting at just $15,000. Contact Outlook Technology, Inc., 200 E. Hacienda Ave., Campbell, CA (408) 374-2990.

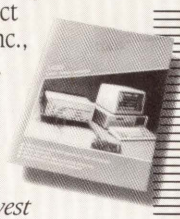

Discover the newest measurement techniques in our 12-page brochure.

OUTLOOK

CIRCLE NUMBER 19

allows more "through routing" between each row.

The Caeco and SCS routing software, on the other hand, create extremely tight layouts, making these tools suitable for the design of MSI and LSI standard cells. The layouts also appear complex and difficult to improve. The Tickle software, on the other hand, creates a fairly simple layout; and the ISS software creates a layout that is routed together manually.

■ Expansion And Compaction

When the routing is completed, the cell synthesis system has produced a symbolic layout: the primitive cells are represented as nothing more than nodes, and the wires have no assigned widths. Symbolic editing refers to the editing of a text or graphical description of the symbolic layout. All the cell synthesis systems provide for symbolic editing of routing; some provide for symbolic editing of the leaf cells.

During expansion, the symbolic description is converted into a physical layout. To perform this simple task, the software needs to know the minimum dimensions permitted for the target technology, as well as the power rail widths, transistor sizes, and so on, that the user wants. The information is entered by the user in a process file. When the layout is expanded, a layout file is created by merging the information in the process file and the symbolic description. Since the physical layout is generated from the symbolic description and the design rules specified in the process file, the resulting layout is correct by construction.

The process of expansion can create a layout with empty spaces, but compaction can be used to compress an airy layout. Most compactors are two-phase and one-dimensional: they squeeze everything in one direction, almost accordion style, and then they squeeze everything in a direction perpendicular to the first direction. Often it is better to compact different parts of a circuit separately, as compaction in one direction is likely to create obstacles for subsequent compaction in the perpendicular direction. Effective compaction thus requires some experience, and systems suppliers sometimes do not recommend it. However, Caeco and SCS use compaction, as an inherent part of cell generation. With other systems, a GDS II (or "Calma stream") format description of the physical layout can be passed into a commercial compactor such as Ecad's.

■ Manual Editing of Physical Layout

All the cell synthesis systems produce reasonably readable layout files and permit the user to edit a layout file with a word processor. The layout file can also be processed to produce a graphical image of the physical layout. Some designers prefer to edit graphical images rather than text files. Since all CAD systems use a graphics editor to construct the physical image, there is always at least one graphics editor available. However, the text in the original layout file does not then match the edited graphical image, and back annotation of the layout file from an altered graphics image is not necessarily a feature. Instead, it is necessary to regenerate the entire layout from the graphic image, which can be a lengthy process. Of course, designers who want to edit the physical layout will probably prefer to use the CAD system they are familiar with; virtually all current IC CAD systems support the Calma-stream physical description format for input; and all the cell synthesis systems supply a Calma-stream output, which can therefore be easily used to move physical layout descriptions between a cell synthesis system and the designer's favorite layout editor.

The final layout is then reformatted in Calma stream and "fractured" into constituent polygons for mask fabrication. However, the physical description may no longer be correct by construction after the expanded layout is manually edited. Designers may therefore prefer to edit the symbolic layout rather than the physical layout. Alternatively, the edited geometric circuit can be run through a design rule checking (DRC) software package (such as Dracula from Ecad) to insure that that there are no design rule violations.

■ Advantages And Disadvantages

Cell synthesis has one major disadvantage compared with standard-cell design: its use results in ICs that contain a large number of elements that have not been precharacterized. Performance is therefore potentially less predictable than with standard-cell design. Consequently, it is advisable to use cell synthesis only with processes that have been well characterized.

On the other hand, cell synthesis does seem to offer many advantages. Cell synthesis can be used to make MSI standard cells when the needed cell is not available, resulting in more compact blocks than would result from the use of many SSI cells. It can be used as an integral part of a standard-cell design system, as well as as a stand-alone tool.

Because the layout is performed at the transistor level, the location of ports and power rails is adjustable. Ports can be moved to reduce congestion in the external wiring channels. Critical paths can be optimized by reducing the relevant wire lengths within the cells. Also, since the layout is at the primitive level, transistors can be sized to a user's specifications. Sizing can eliminate the need for buffer cells and also allow bus drivers to be accommodated within logic cells. If a design needs a NAND gate with a three-state output, the cell is quickly synthesized, as is any other unusual leaf cell configuration.

Cell synthesis can be used to construct many or all of the blocks needed for a particular IC design. The aspect ratio of the created blocks is adjustable, especially if multiple pairs of p and n transistor rows are permitted. Then, after a preliminary layout, the aspect ratios of all the synthesized cells can be adjusted so that the available space is efficiently utilized.

The use of many logic synthesis blocks permits the designer to intermix functions between different building blocks in an IC. If there are long structures, such as NAND trees or complex control functions, the circuits can be spread out to reduce delays.

All the systems are producing MSI layouts that can rival handcrafted design. Design houses that are generating standard-cell libraries can therefore use cell synthesis as a productivity tool. After initial synthesis, the generated layout can be edited, resulting in cells that can actually be better than full-custom cells, especially when the cells contain more than about two dozen transistors. The resulting designs will always be produced faster than if the entire design were constructed manually. ■

Acknowledgments

The author wishes to thank all those who helped in the creation of this article, without whose cooperation the project would never have been possible.

References

De Geus, A.J., and W. Cohen. August 1985. "A Rule-Based Expert System for Optimizing Combinational Logic," *IEEE Design and Test of Computers*.

Hsieh, M.M. March 1985. "Synchronous Square-Wave Division by an Odd Integer," *VLSI Design*.

Meyer, E.L. March 1987. "On Module Generation," *VLSI Systems Design*.

Peskin, A.R. 1982. "Towards a Silicon Compiler," *The 1982 Custom Integrated Circuits Conference*..

NOW YOUR RIGHT HAND CAN KNOW WHAT YOUR LEFT IS DOING.

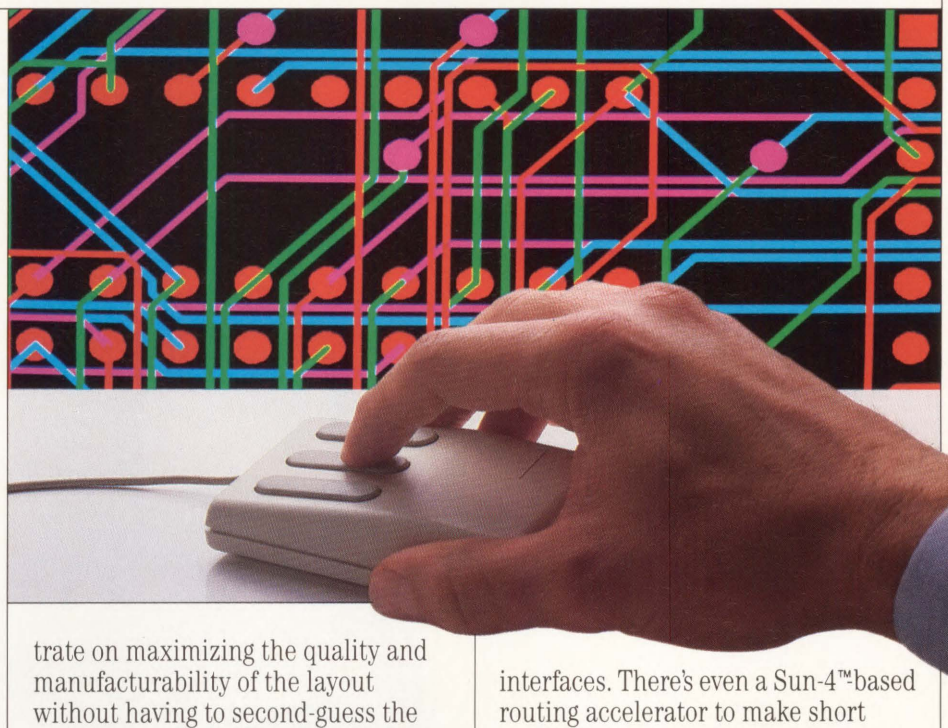

Ever seem like your CAE and CAD people are playing for different teams? Chances are it's because your design systems can't communicate critical information from the engineer to the layout designer. So instead of a smooth handoff, you get hand-to-hand combat.

But now, there's a system that streamlines the way CAE and CAD teams work together.

It's Daisy's BOARDMASTER.™ The first automated system that plays by the rules of real world system design.

With its rules-driven PCB design environment, BOARDMASTER gives engineers the flexibility to specify key design rules in the schematic. Rules for signal priority. Ordering and termination for ECL nets. Package types and power definition. Pre-packaging and pre-placement priorities. Pin and gate swapping. And many other important design considerations.

This critical information becomes part of the design database. And is passed directly to BOARDMASTER's powerful set of PCB layout tools.

With these rules guiding the process, layout designers can concentrate on maximizing the quality and manufacturability of the layout without having to second-guess the engineer's real intentions.

BOARDMASTER comes with the most advanced set of layout tools available anywhere. Like 100% autorouting. Full support for SMD, ECL, analog and ultra fine line designs. Plus a variety of manufacturing interfaces. There's even a Sun-4™-based routing accelerator to make short work of your most complex layouts.

BOARDMASTER even takes the frustration out of design changes. Because it updates only the parts of the database that need updating. Which keeps ECO from becoming a four-letter word.

So if you'd like to get a grip on better board design, put BOARDMASTER to the test in your next project. And give your entire team a hand.

For a demonstration or more information, call Daisy at:
1 (800) 556-1234, Ext. 32.
In California: 1 (800) 441-2345, Ext. 32.

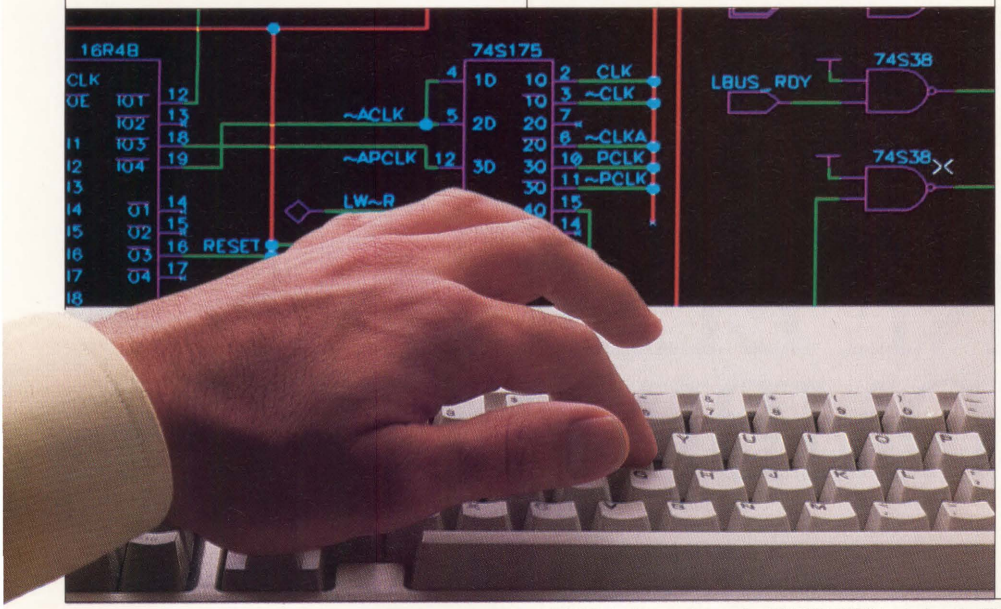

European Headquarters:
Paris, France (1) 45 37 00 12.
Regional Offices:
England (256) 464061;
West Germany (89) 92-69060;
Italy (39) 637251.

daisy SYSTEMS CORP.

PHASE LOCK LOOP.
ANALOG/DIGITAL ASIC.

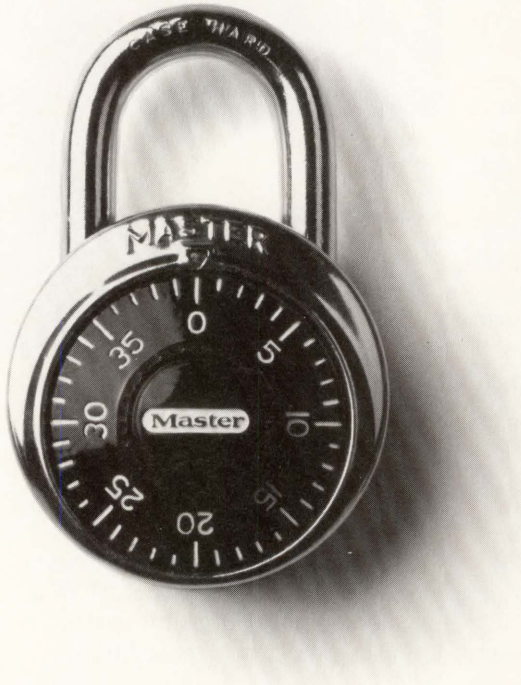

NO ONE ELSE HAS THE COMBINATION.

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

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

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

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

SIERRA integrates PLL onto a single data separator ASIC.

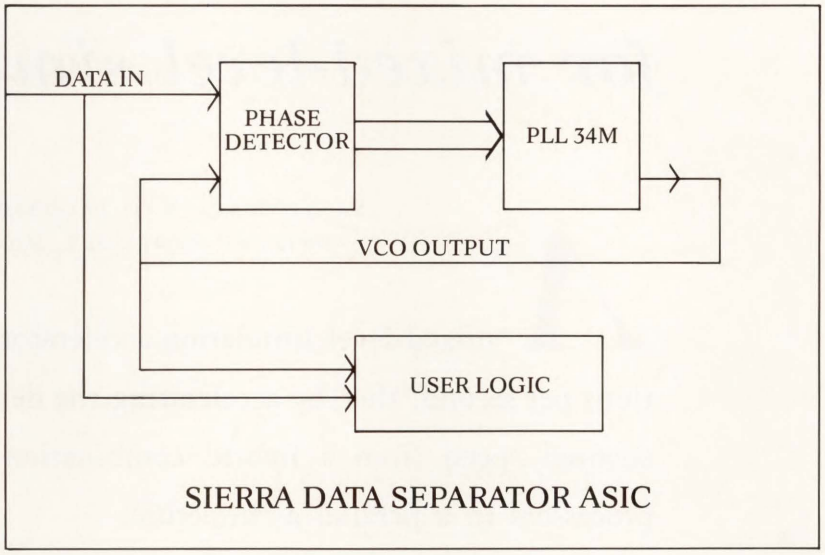

For designers seeking to incorporate analog processing functions, such as amplification, filtering, and data conversion, Sierra offers a wide range of solutions for applications in disk and tape drive systems, local area networks, and servo controllers.

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

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

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

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

Sierra Semiconductor
Triple Technology™ In CMOS.

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

CIRCLE NUMBER 20

MULTIPROCESSOR ACCELERATOR

for mixed-level simulation

SANJIV KAUL, NEIL JACOBSON, AND ISRAEL LIVNAT
DAISY SYSTEMS CORP., MOUNTAIN VIEW, CALIF.

A mixed-level simulation accelerator can perform millions of evaluations per second, thereby accelerating the debugging loop. It gets its unprecedented speed from a hybrid combination of hardwired and microcoded processors in a parallel architecture.

HYBRID APPROACH COMBINES HARDWIRED AND MICROCODED PROCESSORS IN A PARALLEL ARCHITECTURE

The accelerator responds to growing demands for two capabilities in simulation. First, engineers want very fast mixed-level simulation (switch- and gate-level through functional, physical, and behavioral) with the capacity to support large system designs with multiple ASICs. Second, they want fast debugging because they now spend half their design time on that task.

Existing hardwired accelerators cut simulation time with fast switch- and gate-level simulation, but they lack high-performance behavioral modeling (if they offer any behavioral modeling at all), and their debugging environment tends to be poor. In contrast, mainframe-based simulation packages have good behavioral modeling, but their gate-level simulations are at least an order of magnitude slower than what's available on comparably priced hardwired accelerators.

These trade-offs point to the need for mixed-level simulation accelerators, which are fast becoming a universal requirement. Additionally, there's a growing trend for ASICs to include RAM, ROM, and PLAs. These and other complex structures are modeled most efficiently as high-level primitives. Also, system-level designs include components ranging from TTL logic and PLDs all the way to the latest microprocessors.

Further, modeling VLSI components at the gate level can require man-years. Therefore it's usually necessary either to use physical modeling, where the actual component is used to model itself, or to buy a behavioral model.

Another reason for mixed-level simulation stems from the most common cause of late projects—the time required to fine-tune the design to meet all the specifications. By using a behavioral language to model architectural scenarios, the design team can quickly explore and evaluate various design choices, finally homing in on a design that best meets the design goals and specifications. Writing a specification as a behavioral model and then simulating it provides an unambiguous definition. As each designer in a team implements his parts of a design, he can compare his implemented block with its specifications.

■ TWO BASIC ARCHITECTURES

There are two basic architectures for accelerators—hardwired parallel processors and microcoded parallel processors.

Accelerators that can simulate

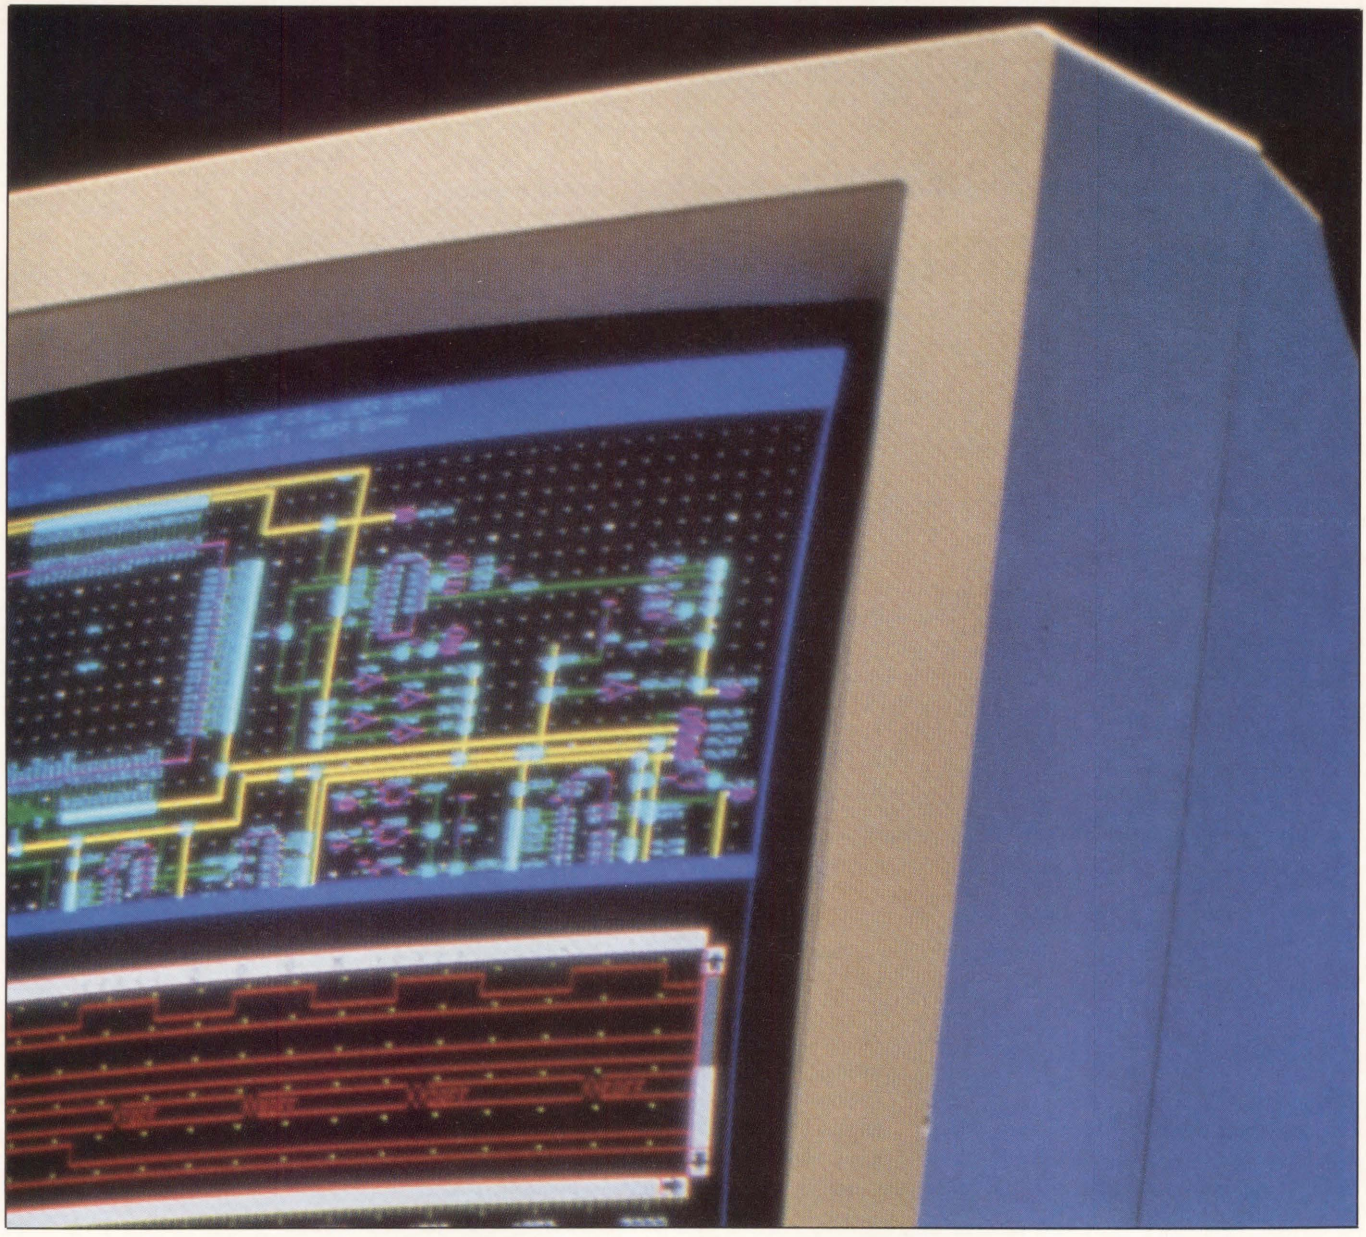

designs larger than 100,000 gates have hardwired parallel architectures. They implement the evaluation algorithms for a predefined set of three or four input primitives in silicon. All models used in a particular design are then translated into these primitives. The types of primitives directly modeled vary from accelerator to accelerator but almost never get more complex than flip-flops.

Hardwired processors can simulate a fixed number of primitives. Capacity and performance are increased by adding processor cards to the accelerator. This expandability enables users to configure the system to their current needs, then add processors as necessary. The processor speed ranges from hundreds of thousands to millions of evaluations per second.

■ HARDWIRED LIMITATIONS

Hardwired accelerators, doing switch- and gate-level simulation, have been targeted at IC design. But they are unacceptable for mixed-level simulation for the following reasons:

- Hardwired accelerators do not model memories—like ROM and RAM—and PLAs directly. That is a serious limitation because most system-level and large ASIC designs have memory components. One can work around the memory limitation either by modeling the memory on the host processor or by invoking a special module containing physical memory every time the memory model needs to be evaluated. The second method is cumbersome, and both methods give unsatisfactory speed.
- Hardwired accelerators have a further limitation. Either they don't support behavioral models at all, or they support them only on a host processor connected through a relatively slow bus, rather than on the accelerator processor cards themselves. The offloaded processing requires that simulations stop and wait for a response every time a behavioral model is evaluated, and that slashes speed.
- Further, these accelerators lack physical modeling support. It is important to provide not only physical modeling, but also a direct high-speed connection between the accelerator and the physical modeler. Otherwise,

ABOUT 80% TO 90% OF A TYPICAL DESIGN IS BUILT OF LOW-LEVEL PRIMITIVES

The Workload Distribution Algorithm

To get maximum performance from an accelerator complicated by two types of processors, Daisy developed an algorithm to distribute a design to be simulated across multiple processors.

The objectives of the algorithm are threefold:

- Primitive assignments must be made across multiples of the same type of processor and divided among the hardwired and software (microcoded) processors (HWP and SWP). Simple primitives are placed on the HWP, and complex primitives on the SWP.
- The parallelism must be optimized without causing excessive interprocessor communication.
- The total processing time of the distribution algorithm should not adversely affect the circuit load time.

Experimentation has shown that the speed advantage of the HWPs relative to the SWPs is such that the assignment of the HWP primitives is not as critical as that of the SWP primitives.

The distribution of the SWP primitives is more critical. Two heuristics are used to attain maximum parallelism during the evaluation phase. First, primitives with very long processing times (like RAMs, ROMs, PLAs, behavioral models, and physical models) are assigned to processors in a round-robin fashion. Second, connectivity information is used to provide parallelism. For example, multiple fan-outs of single outputs of primitives are distributed in separate processors, causing fan-out evaluations to occur in parallel. The effectiveness of this algorithm is, of course, directly related to the number of processors available.

The increased parallelism must not come at the expense of excessive communication. To offset this possibility, primitives are clustered in processors before the round-robin assignment begins. If a gate output fans out to three gates, for example, each of the three gates will be distributed to separate processors. This distribution ensures that all the gates will be evaluated simultaneously during the next evaluation cycle. The amount of clustering depends on the primitive type. Experimentation has shown that, for almost every benchmark and for almost every primitive type, clusters of one primitive are optimal. This finding suggests that interprocessor communication is not the bottleneck.

Distributions are saved in a separate file and then reused if the simulation database has not been changed. Overall, the distribution algorithm assigns primitives quickly and effectively.

speed will be limited to the point of defeating the purpose of the accelerator.

- These accelerators require that high-level building blocks like counters, multiplexers, and eight-input NAND gates be decomposed into three- or four-input gate primitives. The designer must work from this new intermediate design composed of low-level primitives in net-list format, not from a schematic. Because he must trace all errors back through the net list, the tracking down of design errors requires a good deal of time and can significantly lengthen the analysis portion of the debugging cycle.

■ Microcoded Architecture Benefits

Daisy's first-generation accelerator, the MegaLogician, uses a microcoded architecture that supports mixed-level logic and fault simulation. The MegaLogician is a microcoded data-flow machine consisting of three processors running separate portions of the simulation algorithm. The microcoded architecture lends itself to porting and to upgrading of the simulation algorithms. As a result, the MegaLogician makes it easy to write software to evaluate complex primitives.

This ability to model complex primitives and to run mixed-level simulations and microcode simulation algorithms comes at a price, however. One cannot use a microcoded architecture to design an accelerator capable of millions of evaluations per second. There are two reasons. First, the architecture is parallel only within the evaluation cycles, but the evaluation cycles themselves are processed sequentially. Second, the accelerator's performance is limited by memory access times. Even with the fastest memories available, it would not be possible to design a system capable of millions of evaluations per second.

■ A Hybrid Architecture Solution

About 80% to 90% of a typical design is built of low-level primitives. The remainder of the design involves high-level primitives like behavioral models; large Boolean expressions; ROMs, RAMs, or PLAs (or combinations of such blocks); and physical models. Because no processor could provide the speed required, we decided to develop a hybrid architecture (Wong and Franklin, 1987).

The new mixed-level simulation accelerator, the GigaLogician, combines hardwired and microcoded approaches in a parallel architecture (Figure 1), as stated, with communications over a high-speed token ring. Low-level primitives are modeled by a hardwired processor (HWP), and higher-level primitives, like long Boolean expressions, behavioral models, and bidirectional transfer gates, are evaluated in a very fast microcoded "software" processor (SWP). The SWP interfaces directly with as many as five physical modeler (Physical Modeling Extension, or PMX) boards, which fit into the GigaLogician chassis.

The hardwired processor is a nine-stage pipelined implementation of Daisy's simulation algorithm, which can simulate up to 64,000 low-level primitives at a peak of 1.75 million evaluations per second and an average of 1 million. It supports the following primitives:

- Simple gates (five inputs or less)
- Three-state gates (five inputs or less)
- Flip-flops (D and JK) and latches
- Unidirectional transfer gates
- Delays

The software processor can simulate 16,000 high-level primitives at 100,000 evaluations per second. It can simulate all primitives used by Daisy's logic simulator, but when used with the HWP, it supports:

- Logic elements and expressions with greater than five inputs
- RAM, ROM, and PLAs
- DABL (Daisy's behavioral language)
- Bus-contention-and-resolution gates (wired ORs)
- Physical modeler control

A GigaLogician must have at least one

Or...

The choice is really quite simple. Dozens of different PLDs expensively stocked to meet every need. Or the generic GAL® family of programmable high-performance logic devices.

Reprogrammable CMOS.

Behind this family is our Generic Array Logic (GAL) architecture. First developed and produced by Lattice, this architecture lets one GAL device replace all common 20 or 24 pin PAL® devices.

Fabricated with Lattice's proprietary E²CMOS™ technology, GAL devices give you bipolar speed and lower power. Half-power parts are available with maximum gate delays of 15ns. Quarter-power delays are 25ns. Both reduce heat without loss of performance.

Gone are discarded fuse-link parts caused by misprogramming. Gone is the need for expensive window packaging. Gone are long UVerase cycles. Instead, you simply plug your GAL device back into the programmer and erase and reprogram in less than a second. Design changes are simple and easy, too.

Inventory reductions.

From a manufacturing and management perspective, the GAL family offers even more. Now PAL and other PLD inventories can be dramatically reduced. At the same time, design alternatives multiply.

Because GAL devices are reprogrammable, they are also reuseable. And last-minute design changes mean simple reprogramming rather than costly replacement.

100% tested.

Unlike other PLDs, GAL devices are 100% tested for optimum system quality. There is simply no need to overstock in anticipation of a high failure rate.

So forget about that pile of "Eithers." And learn more about the Lattice "Or." To find out how GAL devices fit all of your PLD needs, ask for a free copy of the Lattice GAL Data Book today.

Lattice Semiconductor Corporation™

5555 N.E. Moore Ct. • Hillsboro, OR. 97124

The GAL Revolution Starts Here℠

CIRCLE NUMBER 21

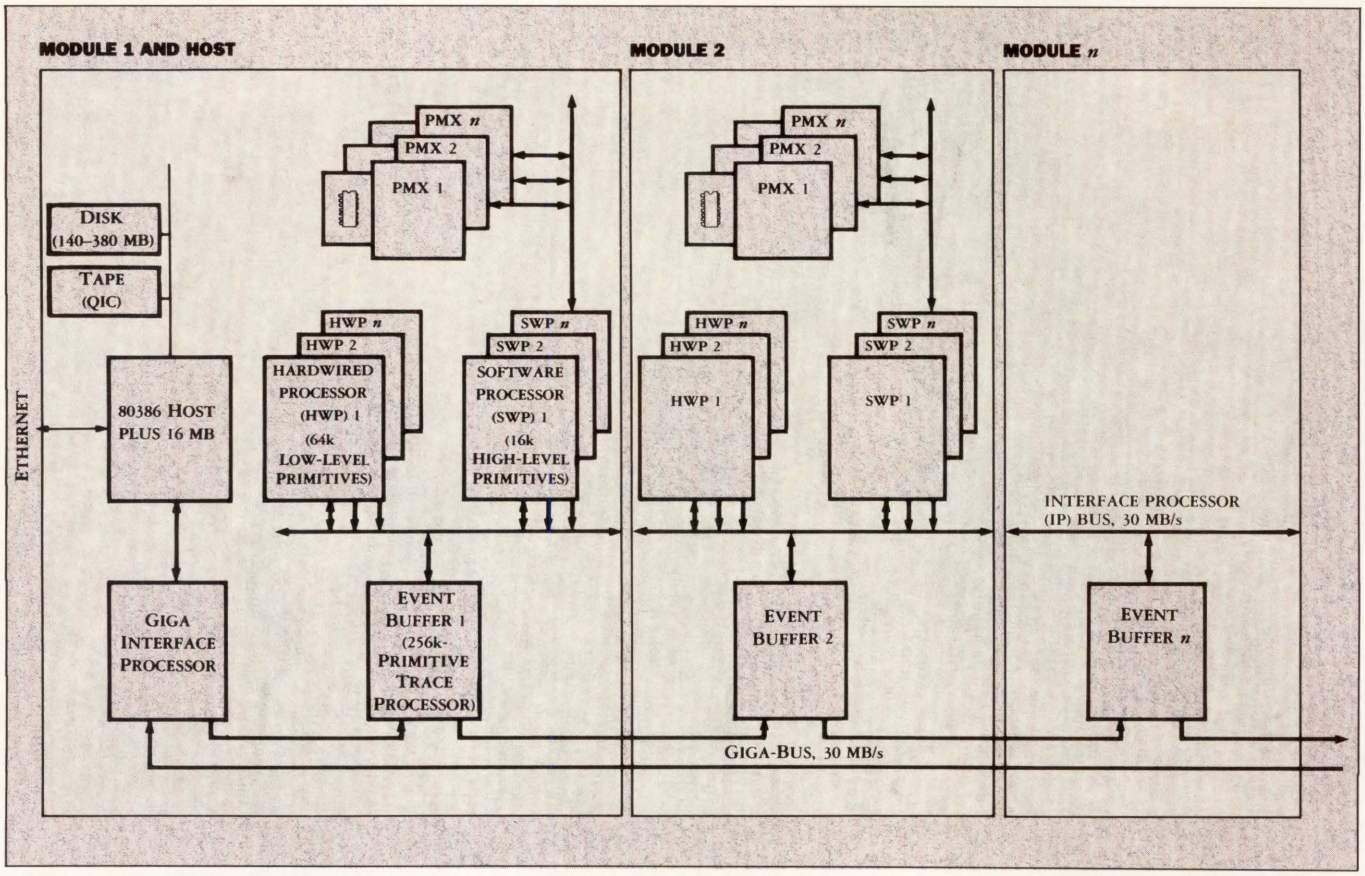

Figure 1. In the GigaLogician simulation accelerator, hardwired processors simulate at the gate and switch level and microded ("software") processors simulate higher-level structures like behavioral models.

SWP per module. The rest of the processors can be any combination of HWPs and SWPs. Each module can have a maximum of 11 boards. The GigaLogician is limited to 64 modules.

In addition to the hardware and software simulation processors, there is an 80386-based host processor to handle Ethernet communications, processor workload distribution, and system utilities. Because the GigaLogician is a network resource, it has its own disk—140 megabytes today—and its own 12-megabyte RAM.

■ Maximizing Simulation Speed

The major challenges in designing the GigaLogician were workload distribution across the processors and fast interprocessor communication.

In any parallel architecture, the workload distribution between processors must be carefully balanced for optimum performance. The problem was more severe with the GigaLogician because there were two types of processors. One algorithm was selected that provides efficient parallelism (see "The Workload Distribution Algorithm," p. 66).

The communications system between the processors minimizes bus loading (Hirose et al., 1987). Three main factors help ensure that interprocessor communications do not become the bottleneck: First, each processor has all the memory and data structures needed to evaluate the primitives on board as well as to evaluate data on states, fan-ins and fan-outs, event scheduling, and other design-related matters. Second, the distribution algorithm balances the workload between the processors. Third, an integral token ring transfers data at 30 megabytes per second.

The GigaLogician interface processor (GIP) is a DMA board that acts as the master controller of the token ring linking all the modules. It can put data onto the ring at the full 30 megabytes per second.

A separate bus provides communications between processors within a module. This bus is controlled by the event buffer processor (EVB), which also operates at 30 megabytes per second.

The buses are processor-independent for two reasons. First of all, this independence allows customers to install any combination of the two processors in their systems. Second, it paves the way for an upgrade path for further improvements in the individual processors or for switching to completely new processors. This upgrade path is extremely important because it enables users to take advantage of rapid improvements in computational horsepower.

■ Toward Faster Debugging

Traditionally, accelerators have been judged only on the merits of raw simulation. Because half the design time is spent on debugging, with simulation time being only a small part of that, users are looking beyond this single measure (Kaul, 1988). The name of the simulation game is fast debugging turnaround, which leads to shorter design time.

In addition to providing fast simulation, the GigaLogician enhances two key features of the Daisy Logic Simulator (DLS): high-speed capture of all node waveform history and fast incremental recompilation.

More specifically, an important benefit of DLS is its ability to open schematics while simulating and to interactively track down problems by probing signals on the schematic while displaying the full history of the simulation for that node. Design problems tend to propagate from nodes that are not specified in the list of nodes to be tracked. The result is that engineers must run simulations iteratively to track down problems. DLS's trace-back capability helps the user in two ways. First, it reduces the overall number of simulations required, because the user does not need to

Introducing OrCAD Verification and Simulation Tools

Logic Simulation

IS ONLY $995

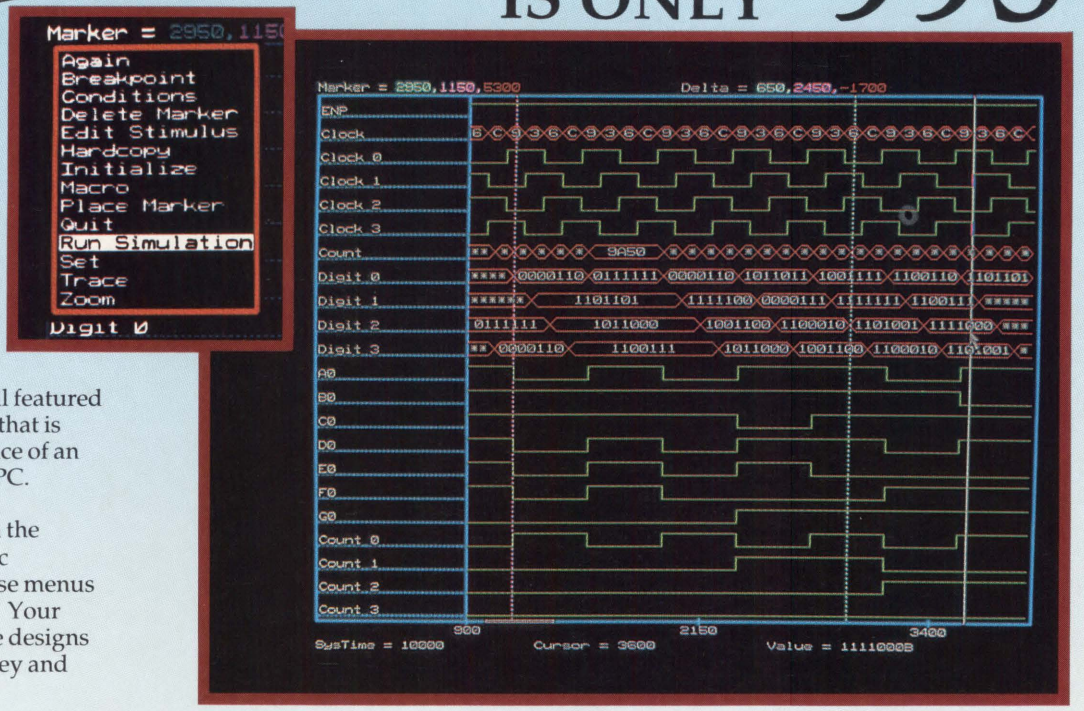

Easy to use menu driven commands speed and simplify design verification.

OrCAD / VST

The Next Logical Step

Introducing OrCAD / VST, a full featured verification and simulation tool that is designed to place the performance of an expensive workstation on your PC.

OrCAD / VST is integrated with the popular OrCAD /SDT schematic capture package, with easy-to-use menus and powerful keyboard macros. Your valuable time is used to simulate designs without investing the time, money and resources to built prototypes.

With unsurpassed performance on a PC, you'll discover that nothing comes close to OrCAD's features and price. Benchmark the specifications for yourself. We guarantee satisfaction, or your money back!

- Event driven, 12-state functional simulator
- Exceeds 10,000 events/sec. on an 8 MHz AT without additional hardware
- Over 14,000 gate capacity
- Logic analyzer display format. Virtual screen displays 50 channels
- 10 breakpoints can be set as AND/OR condition of up to 16 signals
- User selectable minimum and maximum delays
- Input stimulus is easily defined with an integrated pop-up editor
- Includes component models of TTL, ECL, CMOS, Memory devices, and easy to use utility for creating custom models
- Of course, OrCAD's excellent support: technical staff to answer questions, 1 year of free product updates, and a trained sales and support network

Call or write today for our FREE Demo Disk and brochure.

OrCAD™ Systems Corporation

1049 S.W. Baseline St. Suite 500
Hillsboro, Oregon 97123
(503) 640-5007

Contact your local OrCAD Representative for further information

1. WA, OR, MT, ID, AK
   Seltech, Inc.
   206-746-7970
2. N. CA, Reno NV
   Elcor Associates, Inc.
   408-980-8868
3. So. CA
   Advanced Digital Group
   714-897-0319
4. Las Vegas, NV, UT, AZ, NM, CO
   Tusar Corporation
   602-998-3688
5. ND, SD, MN, W. WI
   Comstrand, Inc.
   612-788-9234
6. NE, KS, IA, MO
   Walker Engineering, Inc.
   913-888-0089
7. TX, OK, AR, LA
   Abcor, Inc.
   713-486-9251
8. MI, E. WI, IL
   Cad Design Systems, Inc.
   312-882-0114
9. IN, OH, KY, WV, W. PA
   Frank J. Campisano, Inc.
   513-574-7111
10. TN, NC, SC
    Tingen Technical Sales
    919-878-4440
11. FL
    High Tech Support
    813-920-7564
12. DE, VA, MD, DC
    MGM Visuals
    703-352-3919
13. MS, AL, GA
    Electro-Cadd
    404-446-7523
14. E. PA, NJ, NY
    Beta Lambda, Inc.
    201-446-1100
15. CT, RI, MA, VT, NH, ME
    DGA Associates, Inc.
    617-935-3001
16. BC, AB, SK, MB
    Interworld Electronics & Computer Industries, Ltd.
    604-984-4171
17. ON, PQ
    Electralert, Ltd.
    416-475-6730

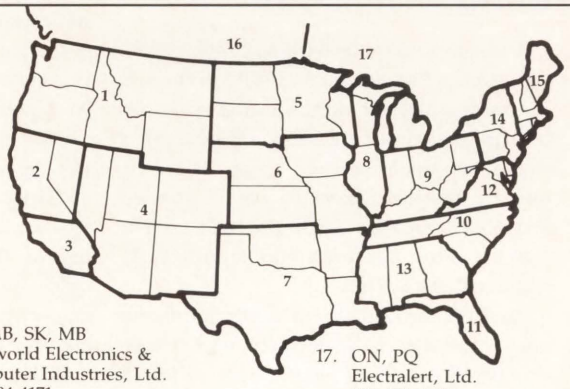

OR 8746

CIRCLE NUMBER 22

prespecify nodes he wants to observe. Second, because the user doesn't have to wait for another simulation, his train of thought is not interrupted and he can concentrate on tracking down bugs.

The event buffer processor solves the problem of recording trace information during simulation while maintaining high simulation speed. Without the EVB, storing the full history of all nodes in the design would make the simulation disk-bound. That binding is the reason most accelerators allow users to trace only a limited number of signals.

■ Tracing Data with No Slowdown

Several features help the EVB collect trace data with minimal impact on simulation speed. First, the EVB has 3 megabytes of RAM available. Second, it uses memory wraparound and split-memory techniques. As the RAM gets filled, the oldest data are stored in a cache file on the 80386. This file is buffered and downloaded onto the hard disk as a background task. The user has the option of filing his simulation data on the disk in a wraparound mode or expanding the file size.

Engineers can select from various modes: selected trace signal, full history, and mixed (selected full trace and window tracing for all remaining nodes, with no disk bandwidth limitations). The 80386-based host and its peripherals are isolated from the simulation processors during simulation, leaving the host to provide only control functions, mass storage services, and LAN communications.

Large designs running on accelerators normally take a long time to compile (Mott and Hall, 1985). To address this problem, we have developed a new modeling system that speeds recompilation by an order of magnitude. The modeling system is based on modular software compilation, in which a change to a subroutine requires recompilation of only that subroutine.

■ How Sublinks Help

Users designate individual blocks at the top level of the design as sublinks (Figure 2). Submodules of these sublinks can also be designated as sublinks. Whenever a change is made on any page, only the relevant sublinks have to be recompiled and then relinked with the rest of the system. This approach can significantly speed recompilation.

Another advantage of sublinks shows up in multiple-ASIC designs, which pose two problems. First, ASICs from different vendors often use common cell names. Consequently, name-clashing problems frequently come up when designs are compiled, forcing users to change names manually. The separate compilation of each sublink allows conflicting names to coexist. Second, multiple-ASIC designs often use dissimilar technologies (for example, 1-μm and 2-μm CMOS or vendor A mixed with vendor B). Therefore they have different equations for back-annotated delays from layout. With sublinks, each ASIC can be placed on an individual sublink that can support specific delay information pertaining to that ASIC's technology.

Using sublinks, a GigaLogician beta site was able to simulate a system consisting of 10 ASICs using different technologies and from different vendors without any of the problems associated with name clashing or back annotation of delays.

The new accelerator works in the Daisy environment (Figure 3), which consists of PC-based platforms, 80386-based workstations, and computation servers (Sun-4's). Designers can enter schematics and compile designs on the lower-cost workstations. They can even simulate individual modules or gate arrays on the workstations alone or, depending on the size, with the aid of a networked accelerator.

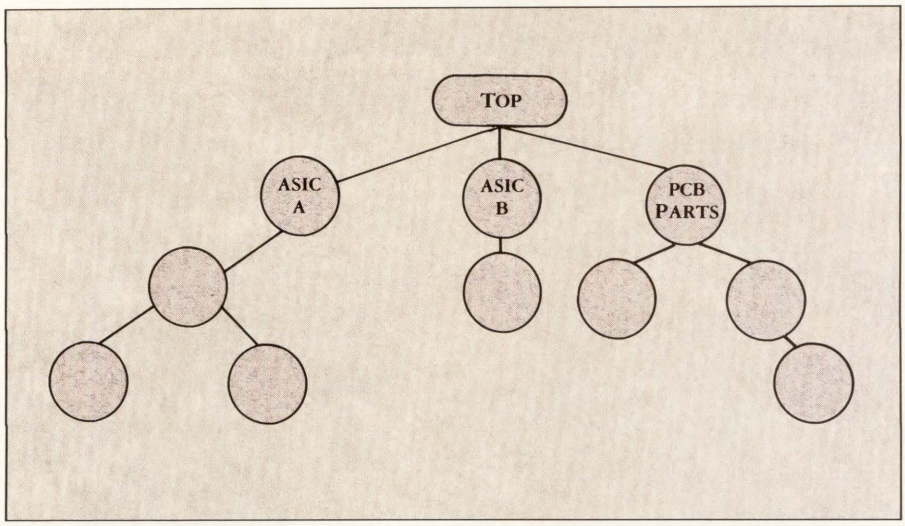

Figure 2. A sublink structure makes for fast recompilation because only affected sublinks need be recompiled. Sublinks also support multiple ASIC technologies in a single simulation.

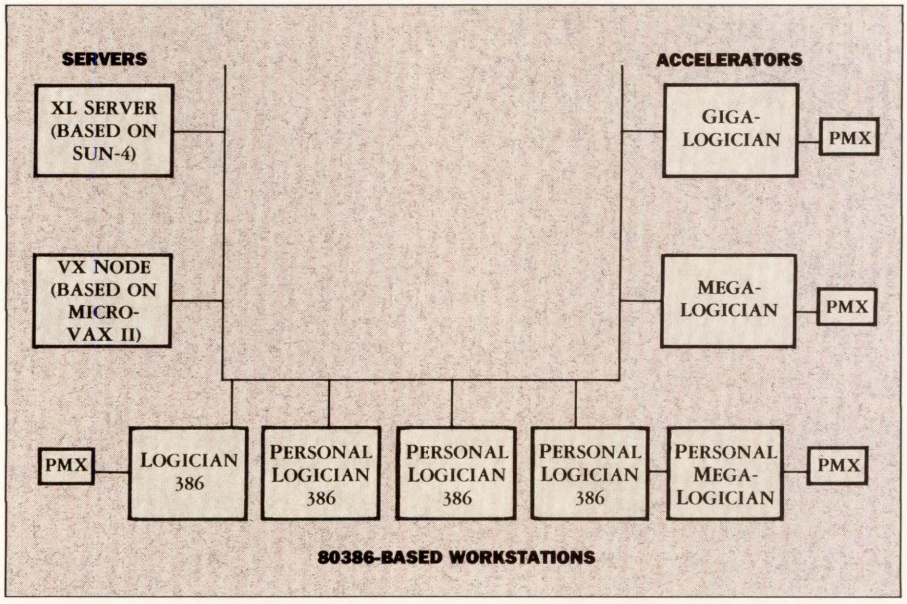

Figure 3. Distributed accelerators and servers can be accessed from 80386-compatible workstations over Ethernet.

■ Merging to System-Level Simulations

Once the individual modules have been verified, the designer can perform system-level simulations on the GigaLogician by designating each ASIC as a sublink, which can then be merged on the new accelerator.

A typical 32-bit workstation can simulate designs with up to 4,000 gates, beyond which simulation runs become unacceptable. An accelerator like the MegaLogician can handle designs in the 20,000-to-50,000-gate range. However,

for designs larger than 100,000 gates, a different class of performance is required.

Unfortunately, technical and scientific computers are optimized for general-purpose numerical processing, whereas digital simulation requires integer processing. The new mixed-signal accelerator, which is targeted at the simulation process, can perform this type of task 10 to 20 times faster than even the most powerful general-purpose machines, such as a Cray X-MP.

To put this performance in perspective, a simple calculation shows that simulating one second of operation of a 10-MHz, 100,000-primitive system requires 10 billion evaluations, assuming 1% activity. The most powerful general purpose computer would take 20 to 30 hours, while accelerators like the GigaLogician could do it in an hour or two.

The GigaLogician also shares the user interface and library database of the Daisy Logic Simulator and the MegaLogician.

■ The Bottom Line

Because simulation accelerators are specifically designed to run digital simulations, they are rated in terms of evaluations per second, not MIPS. A rule of thumb beginning to emerge is that the most efficient portable mixed-level simulators deliver 5,000 evaluations per second on the average per MIPS (in a design consisting of 20% behavioral primitives and 80% gate-level primitives). A GigaLogician configuration with two SWPs and one HWP can deliver an average of 1.2 million evaluations per second, the equivalent of 240 MIPS.

Recompilation speed is equally impressive, with improvements over the MegaLogician ranging from 5 to more than 50 times. One beta site recompiled a design with 240,000 primitives in 12 minutes, compared with 690 minutes for a "cold" compilation. ■

References

WONG, K., AND M.A. FRANKLIN. 1987. "Performance Analysis and Design of a Logic Simulation Machine," *Proceedings of the 14th Annual International Symposium on Computer Architecture.*

HIROSE, F., ET AL. NOVEMBER 1987. "Simulation Processor SP," *ICCAD-87 Digest of Technical Papers.*

KAUL, S. FEBRUARY 8, 1988. "Accelerators: Beyond Raw Performance," *Electronic Engineering Times.*

MOTT, G., AND R.B. HALL. OCTOBER 1985. "The Utility of Hardware Accelerators in the Design Environment," *VLSI Systems Design.*

About the Authors

Sanjiv Kaul *is the product manager for accelerators at Daisy Systems. He joined Daisy in 1984 as a technical specialist for test products. Prior to joining Daisy, Kaul was a design engineer for Develco. He holds a BSc from the University of Delhi and a BSEE from the University of Maryland. He is currently enrolled in an MSEM program at Santa Clara University.*

Neil Jacobson *has been a senior software engineer since joining Daisy Systems in 1986. He has been involved in the design and development of the GigaLogician software at the application, system, and microcode levels. Jacobson holds a BA Sc in engineering science from the University of Toronto and an MSEE from the Technion, the Israel Institute of Technology, Haifa.*

Israel Livnat *is the director of hardware engineering at Daisy Systems. He has been responsible for Daisy's simulation acceleration and physical modeling products. Before joining Daisy, he was head of the systems electronics department for the Israeli Aircraft Industries. Livnat holds BSc and MSc degrees in electronics engineering from Technion, the Israel Institute of Technology, Haifa.*

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

CICC HOSPITALITY SUITE

MAY 17, 3-9 PM
HOLIDAY INN, GENESSE PLAZA
SUITE 304

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

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

NEC Electronics Inc.

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

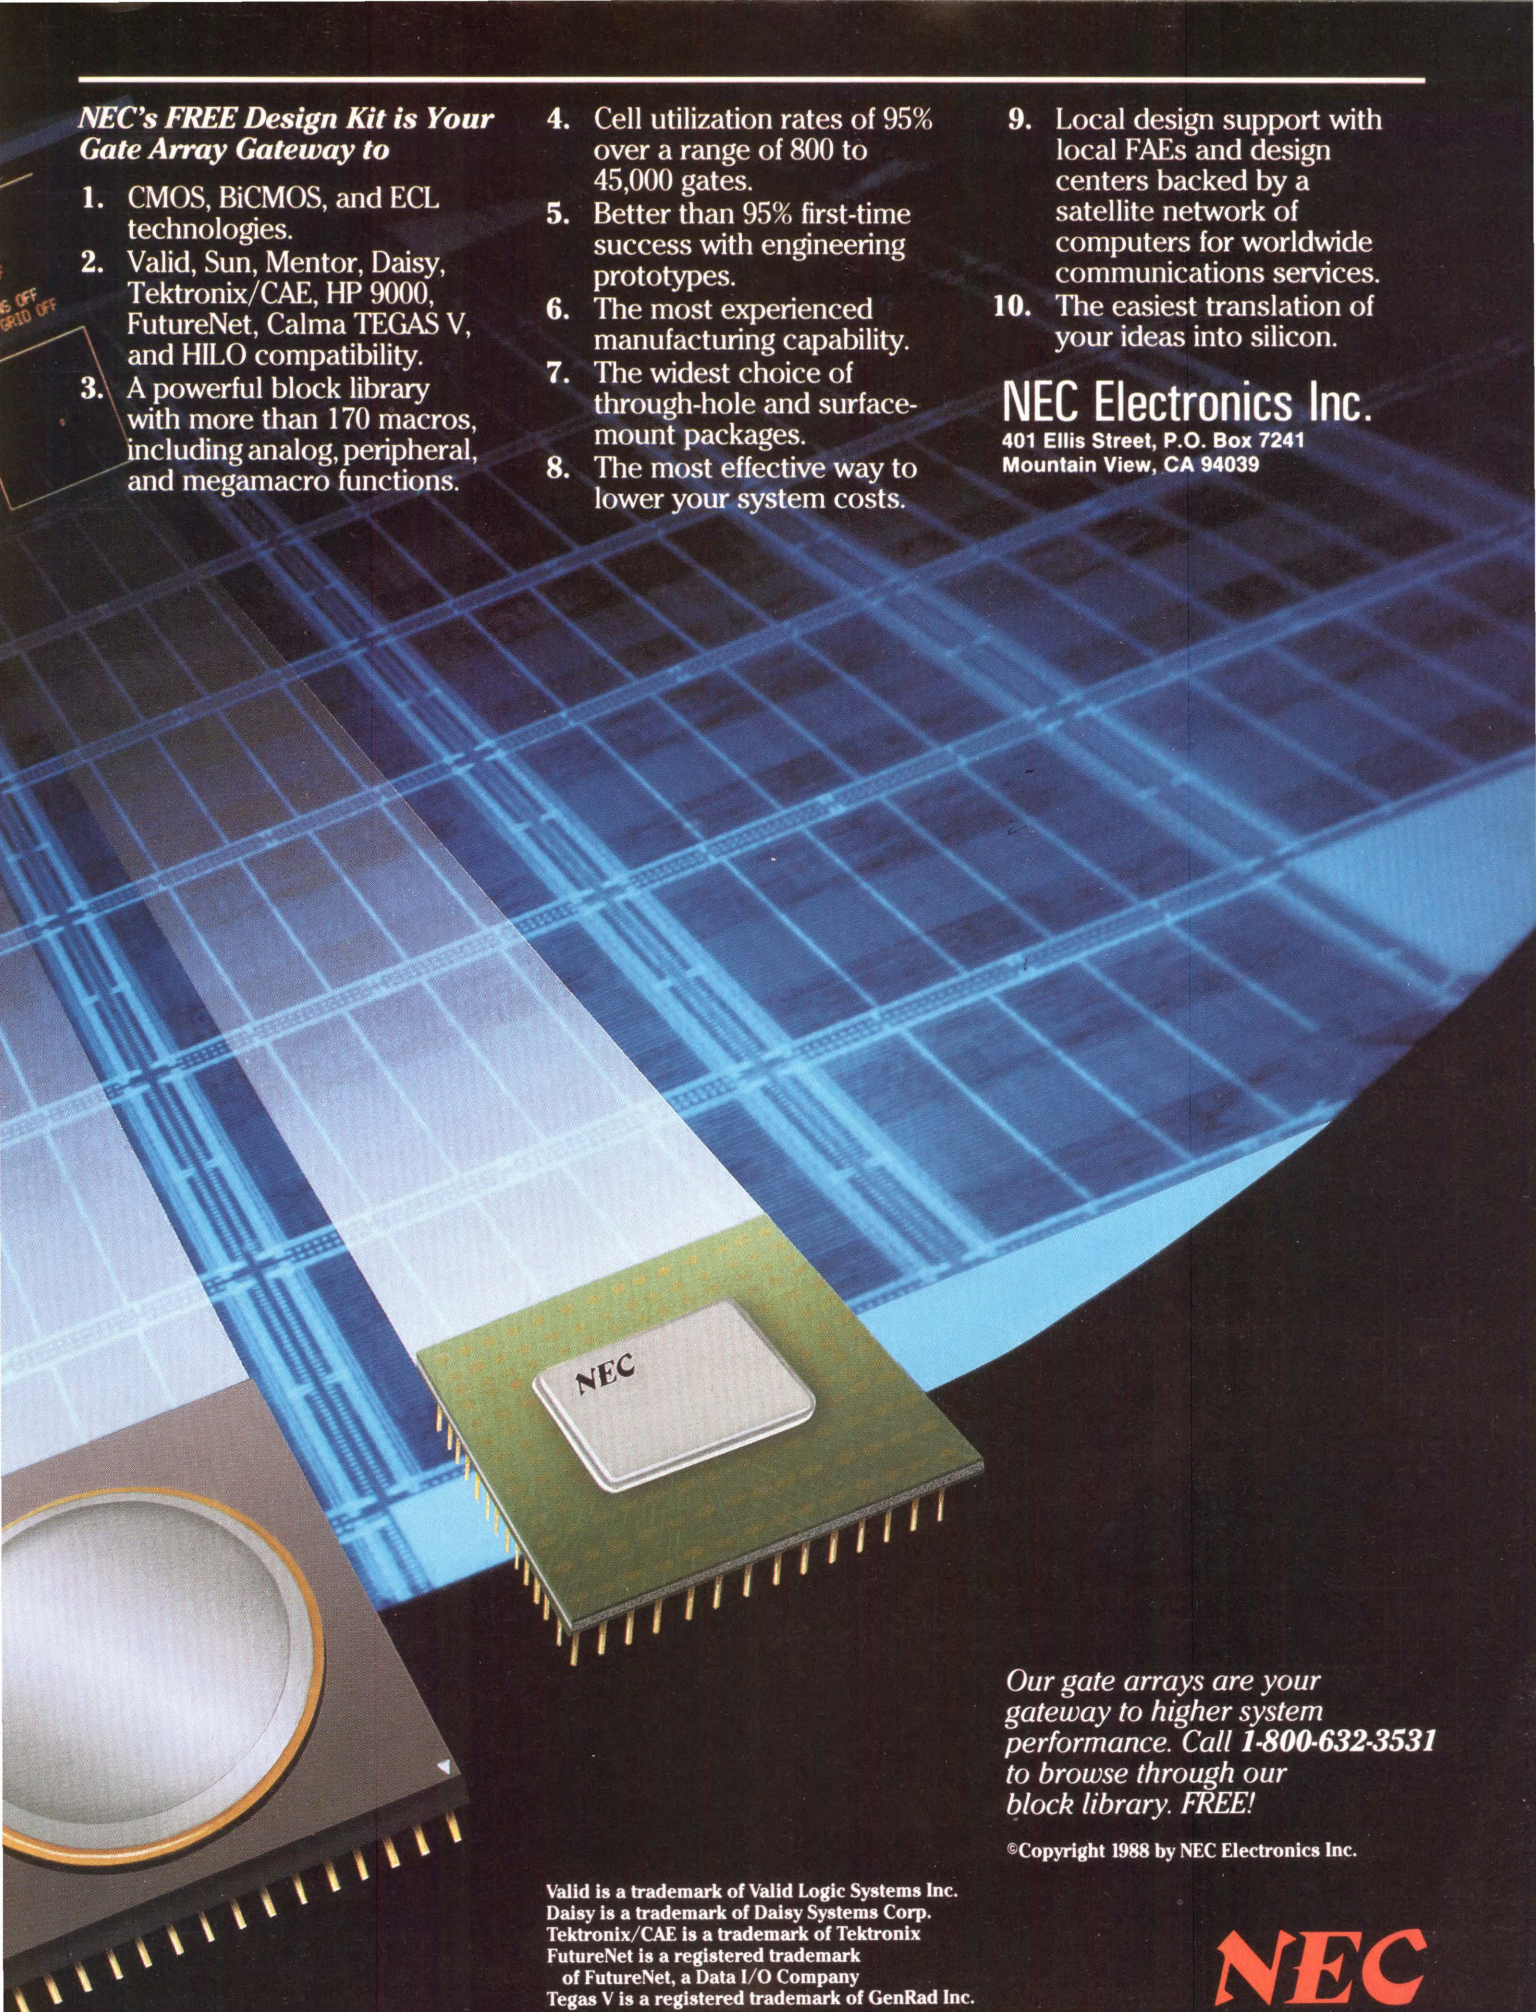

Our gate arrays are your gateway to higher system performance. Call ***1-800-632-3531*** *to browse through our block library. FREE!*

Valid is a trademark of Valid Logic Systems Inc.
Daisy is a trademark of Daisy Systems Corp.
Tektronix/CAE is a trademark of Tektronix
FutureNet is a registered trademark of FutureNet, a Data I/O Company
Tegas V is a registered trademark of GenRad Inc.

NEC

CIRCLE NUMBER 24

DESIGNING ASICS FOR High Reliability

A Little Common Sense Goes a Long Way

ORHAN TOZUN,
INTERNATIONAL MICROCIRCUITS INC.,
SANTA CLARA, CALIF.
ROBERT E. PERRY, CUBIC DEFENSE
SYSTEMS, SAN DIEGO, CALIF.

The design of high-reliability gate arrays not only requires selecting a hi-rel ASIC vendor but also requires considering several circuit-level, logic design, test vector, and simulation issues. The designer wishing to optimize a design for reliability must be aware of the following: current density levels of the internal array, peak current levels of the chip, power mesh current distribution, the nature of IC processing faults, and the relationship of testability level to reliability. In addition, the designer must be able to evaluate his layout for optimized reliability and be able to generate high-reliability test vectors.

Reliability Requires a Lot More Than Just a Reliable Vendor and Fabrication Process

With those issues in mind, this article discusses the design process of VLSI high-reliability chips. Specifically, the article discusses the fundamentals of designing for reliability, logic block design techniques and partitioning for reliability, layout and packaging for reliability, and test vector generation and simulation for achieving field-reliable array products.

As VLSI technology matures, it is gaining acceptance as a medium that can be used for high-reliability applications. However, designing high-reliability arrays is not just a simple matter of standard high-reliability environmental screening and testing. It also involves a fundamental knowledge of, and experience with, the basics of IC design. Imperative to the design of reliable CMOS ASICs is the design flow process and the design methodology that are used, in addition to the standard high-reliability tests.

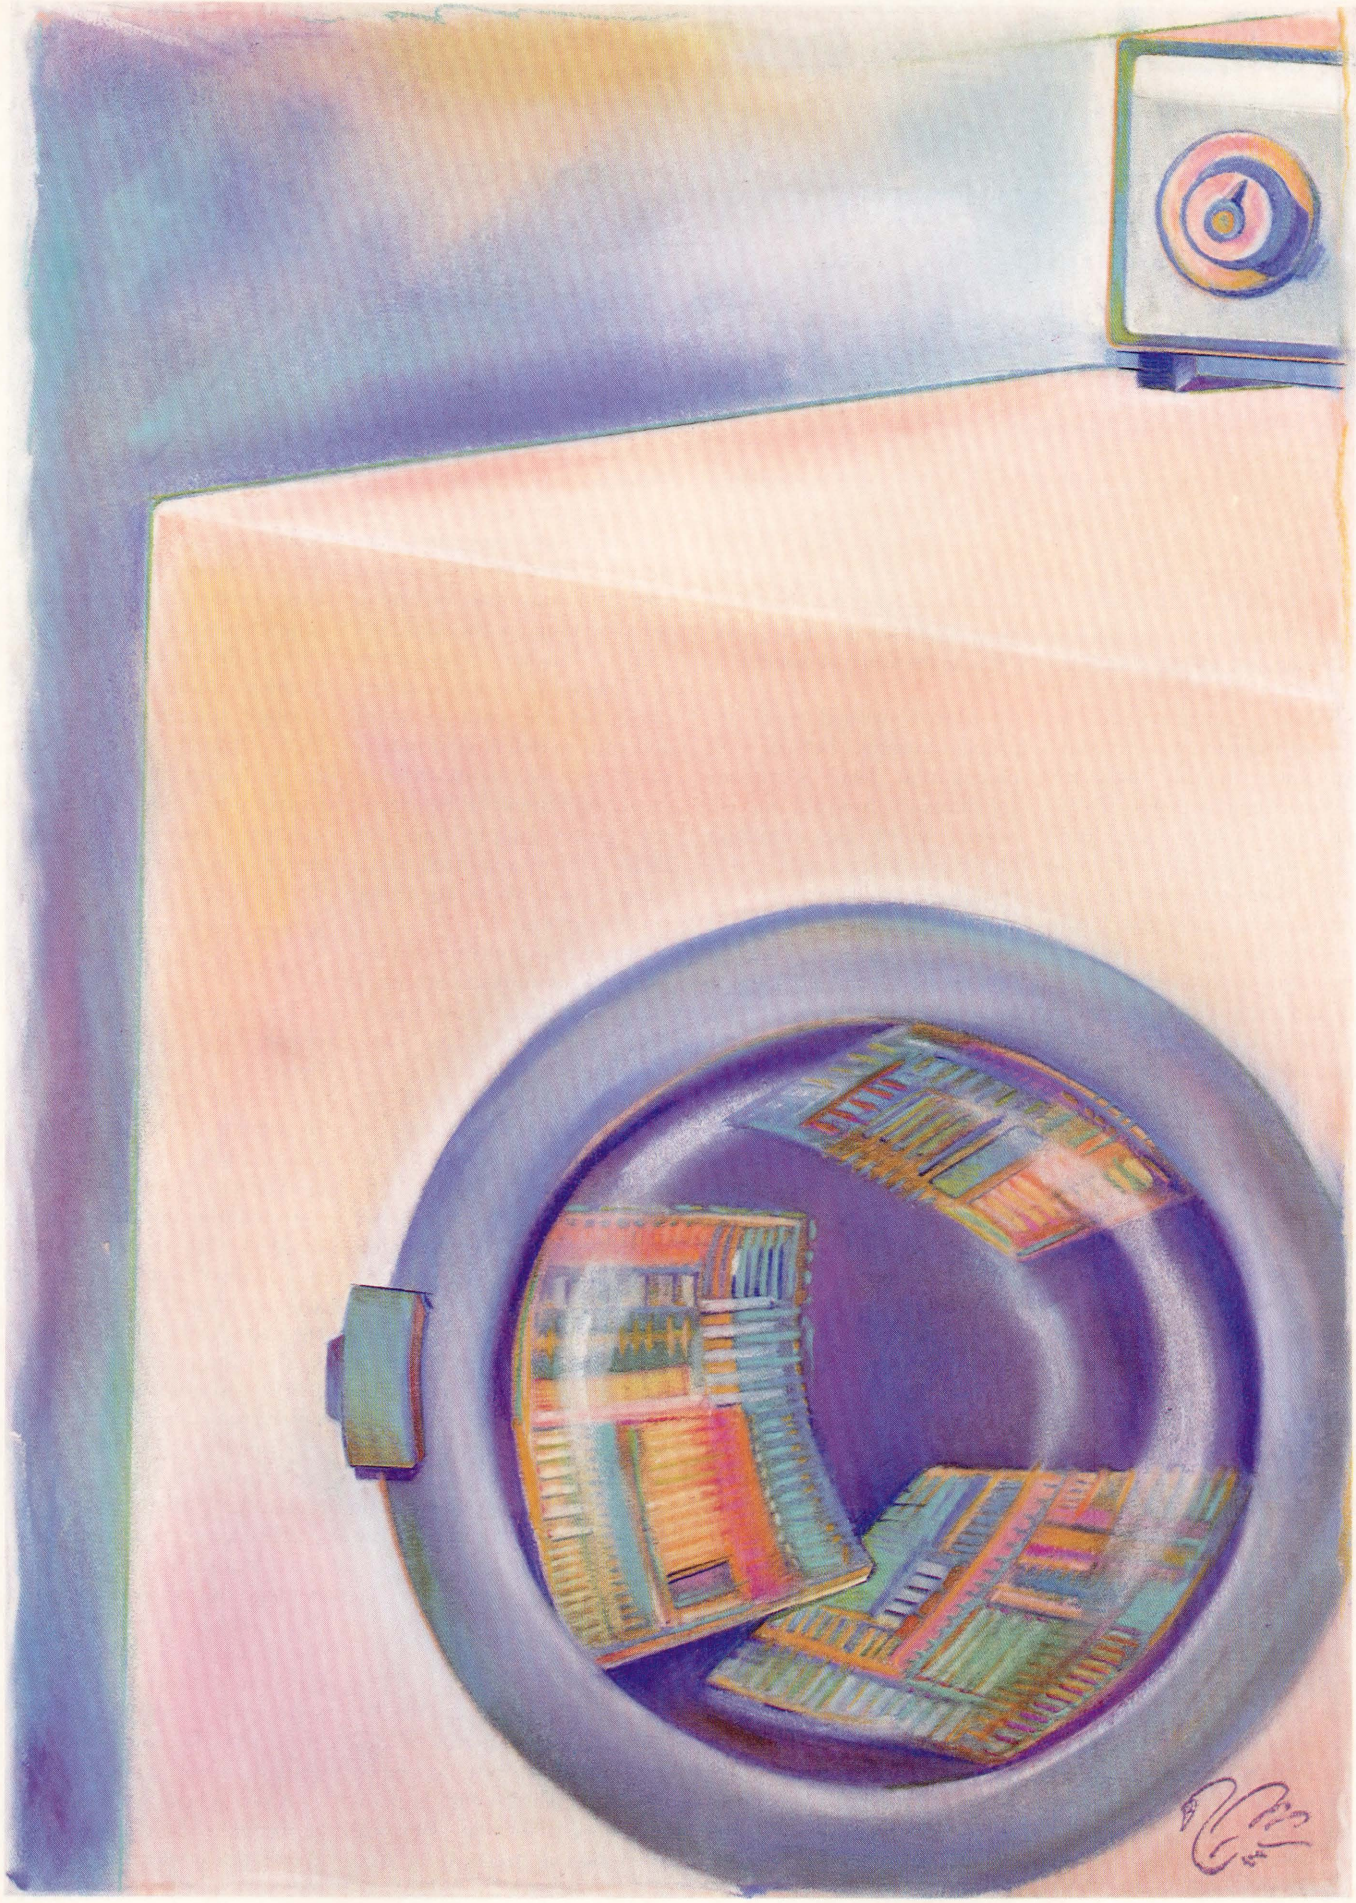

ILLUSTRATION SARA SWAN

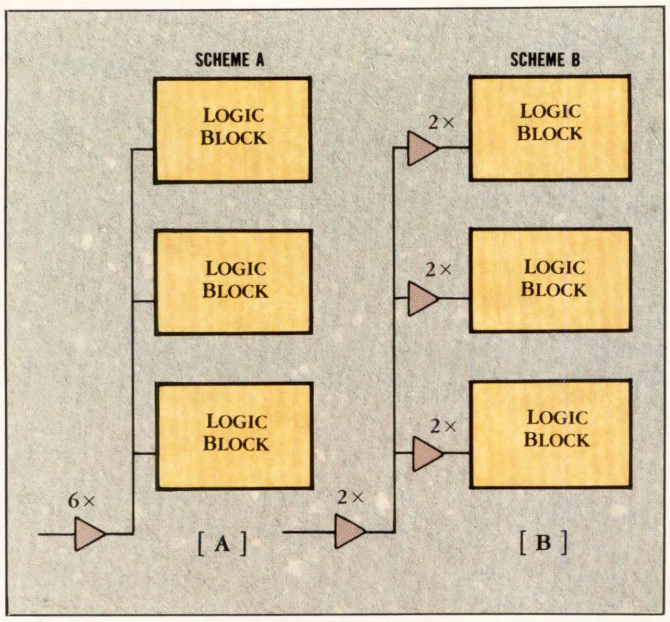

Figure 1. Scheme A uses one high-drive buffer in parallel with three logic blocks. Using lower-drive buffers, as in Scheme B, reduces the danger of overdrive.

Many techniques can be used to improve ASIC reliability. These include those for minimizing processing faults; reducing the amplitude, duration, and frequency of transients; improving the circuit's immunity to transients; minimizing the number of chip areas with high current density; reducing current density and peak currents; providing adequate current drive; and designing circuitry for ease in testing.

The adequacy of current drive depends on the power supply. It must be able to deliver peak currents large enough for worst-case conditions. For CMOS designs, these occur during transitions. The largest currents are drawn during I/O switching and during switching of buffers that drive large capacitive loads. Thus it is important to specify a power supply, not for average currents but for peak currents.

■ The Right Current

Inadequate current drive can allow noise to ride on a long rising or falling signal edge. The slow rise and fall could lead to incorrect timing. Many CMOS simulators do not compute rise and fall times; they just calculate propagation delays—the time from midsupply to midsupply. The simulator would indicate proper timing because the propagation delay is correct, yet the overly long rise and fall times, which the simulator might not detect, could lead to trouble.

Inadequate drive is not the only source of poor circuit functioning or inadequate speed. As an example, an open circuit in a line leading to one of two inverters paralleled to provide greater drive current, and thus greater speed disables that inverter. It is therefore wise to use test vectors that detect opens at operating speed (Banerjee and Abraham, 1984; Courtois and Baschiera, 1984).

One approach to coping with excessive rise and fall times is to use a Schmitt buffer to sharpen the pulse. Another approach is to use buffer trees instead of high-drive buffers. In Figure 1, for example, the circuit in scheme A uses six paralleled buffers (a high-drive 6× buffer) to drive three logic blocks. The blocks here would be far more susceptible to the effects of overdrive than would be those in scheme B. There, a lower-drive buffer (2×) drives three other lower-drive buffers to deliver the same current to each logic block at approximately the same delay without the danger of overdrive.

One must avoid overdrive as well. Overdriving a node can produce high peak-current levels and thus large voltage transients.

■ Cutting Transients

Voltage transients can be reduced by using multiple power supply lines and separate power lines for the internal power distribution mesh and for the I/O power bus ring. But these approaches increase chip size and therefore cost.

Transients can also be reduced by partitioning the chip to minimize concurrent I/O switching, and they can be reduced further by designing a circuit to minimize concurrent I/O switching in the same direction. Further, extra power and ground pins can be added to shorten the path to nodes and thus to reduce lead inductance, and one can use separate power supplies or at least separate power supply lines for those sections requiring large currents and those requiring small currents.

Noisy environments are particularly harmful to sequential circuits. A noise pulse invading a counter circuit, for example, can set the circuit into a state that is not part of the designed-for state sequence. A BCD counter might continuously cycle through its six unused states (10 to 15), instead of resetting to 0 after it reaches a count of 9.

Undesired states can be prevented by a self-correcting circuit or separate error correction circuitry to reset the circuit to a state within its proper state sequence. Alternatively, at lower cost, one can design circuitry to alert an operator, who can then reset the circuit.

■ Current-Distribution Problems

A chip may appear to be perfectly safe in terms of the total current it draws, yet be in serious danger because of poor current distribution. Current density is a function of frequency (in a CMOS design), as well as of current drive level and capacitive load.

Generally 1 to 2 mA rms is a safe rating for signal interconnects and 50 to 100 mA rms is safe for the power distribution network in a 2-μm CMOS design. One should be careful to maintain these current ratings, particularly with high-drive buffers operating above 20 MHz and with situations involving many I/O switches operating at greater than about 10 MHz.

There are several approaches to minimizing current density on a chip. One is the classical power ring structure. Here, the periphery of the die has a wide (low-inductance) interconnect bus to supply the large current drawn by I/O devices. Smaller currents are supplied through thinner lines to the logic circuitry. The currents are distributed more evenly by partitioning the circuits in quadrants and supplying different quadrants through different V_{DD} and V_{SS} ports.

■ The Value of Partitioning

Partitioning can be used to reduce drive current requirements by using scheme B instead of scheme A in Figure 2. In scheme A the divide-by-5 counter must drive the package capacitance of another chip as well as the capacitance of the divide-by-500 counter, whereas the divide-by-5 counter in scheme B must merely drive the capacitance of the divide-by-500 counter. In scheme B, the second chip is driven at a much lower frequency, so that the frequency of the voltage transients is reduced.

Further, a circuit can be partitioned to make it easier to test. As an example, scheme A is easier to test than B because one could more readily access the divide-by-5, which might be more failure-prone as it switches at a higher frequency than the divide-by-500 and thus dissipates more heat.

Also, scheme A would require fewer test vectors, as one could test the divide-by-500 and subsequent circuitry with

PROTECT YOUR ASIC INVESTMENT

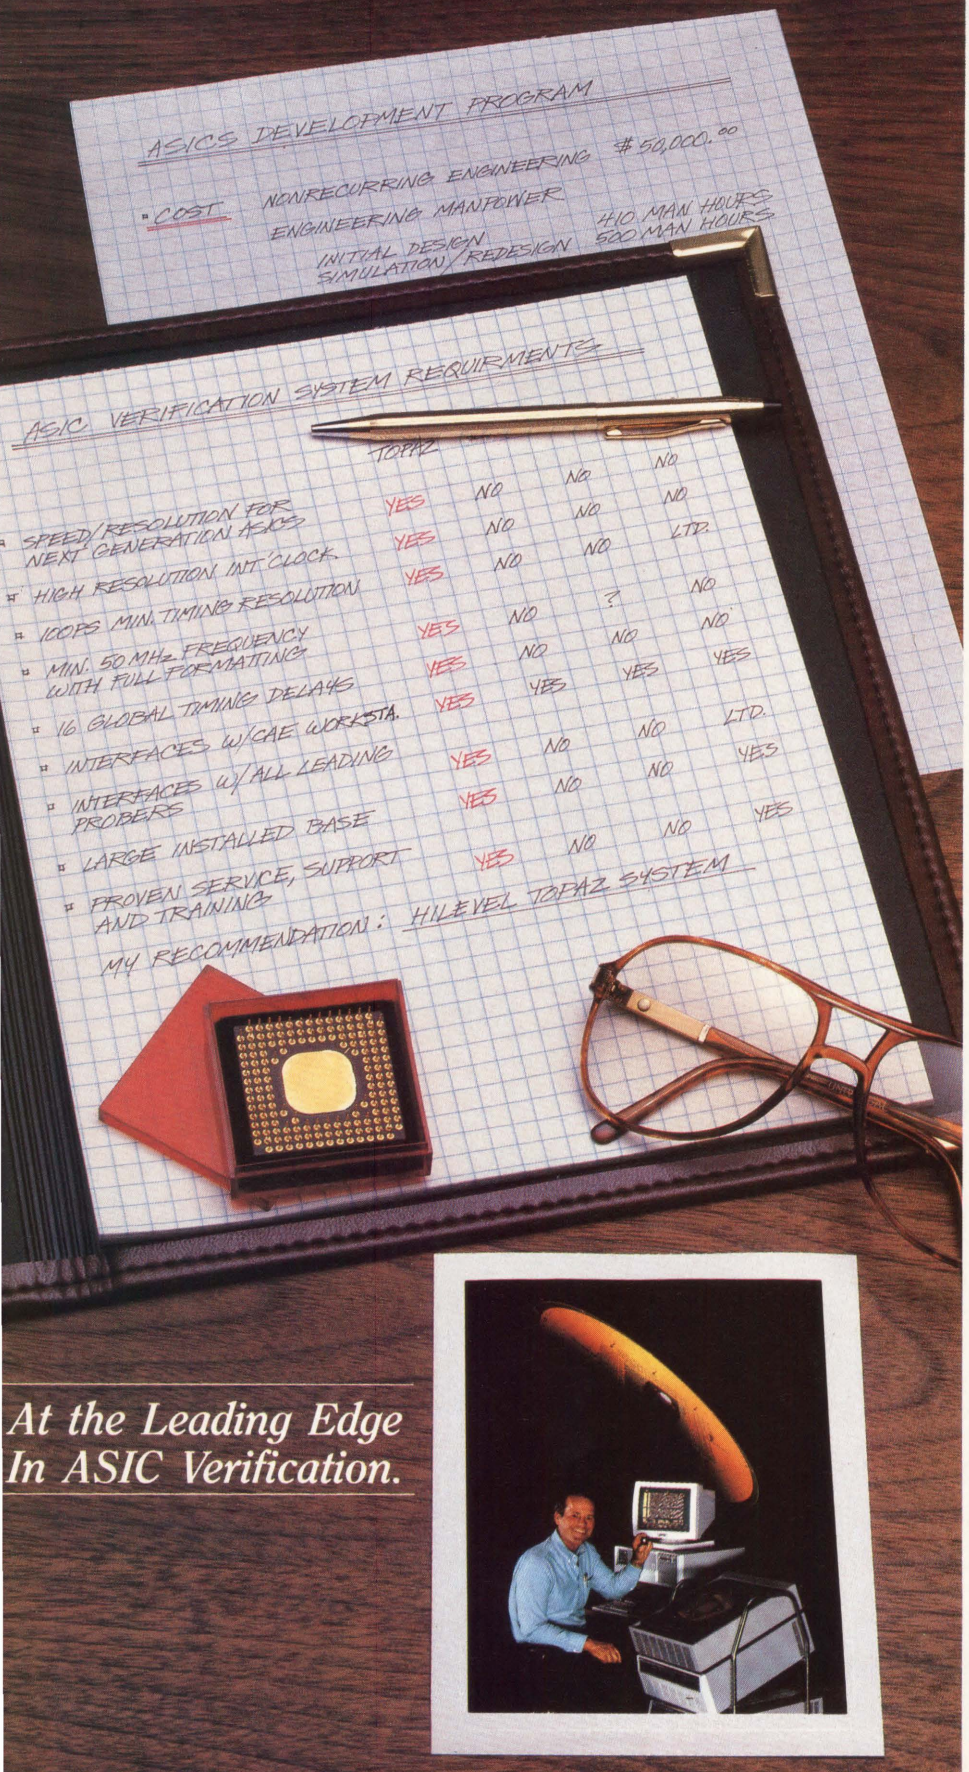

At the Leading Edge In ASIC Verification.

Insure Against Schedule Slippage And Unexpected Costs.

Committing to an Application Specific Integrated Circuit may be the only way to stay ahead of competition, but it has its risks. Discovering an undetected design error after production begins can be devastating. Topaz Series Design Verification Systems offer features which assure your schedules and costs will stay on target.

TEST IT LIKE YOU'LL USE IT.

Test your ASIC prototype at operating speeds—with full data formatting and no multiplexing—all the way to 50 MHz. It gives you confidence your design is right *before you take it to market.*

FAST DEBUG SHORTENS DESIGN CYCLE.

Error mapping, real-time compare, and automatic characterization give Topaz the best diagnostic and analysis capability in the industry. Chances are that Topaz will have your design errors diagnosed before other methods have even indicated a failure.

BEST RETURN ON INVESTMENT.

Topaz has the power to test all of today's logic families, and those you'll use tomorrow. This allows you to amortize your investment over several years; not just today's project . . . and you can take that difference to the bank!

HILEVEL—A SOUND DECISION.

You'll be working with *the* technical leader in this demanding field, with over eight years dedicated experience and a proven record of customer support. With a world-wide installed base and a host of customer testimonials, Topaz gives you confidence in today's designs and reassurance for tomorrow's challenges.

Call today.

HILEVEL
TECHNOLOGY, INC.

31 Technology Drive, Irvine, CA 92718
Phone: (714) 727-2100
DIAL TOLL FREE 1-800-HILEVEL
(In California 1-800-541-ASIC)

CIRCLE NUMBER 25 SEND LITERATURE
CIRCLE NUMBER 26 NEED DEMO

VISIT US AT ELECTRO/88
BOOTH NO. 2034 – WORLD TRADE CTR.

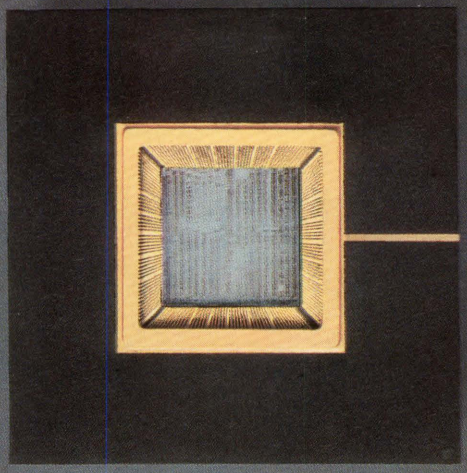

50,000 Gates At 0.7 Nanoseconds.

Toshiba's process technology has made us the world's leading supplier of CMOS devices. From this experience comes the capability to build a 50,000 gate ASIC.

And from Toshiba's manufacturing capabilities comes the volume production to meet your smallest or greatest needs.

These new Toshiba "Sea of Gates" Compacted Arrays™ provide high packing density for more effective silicon utilization and ultra-high speed. And they are available in 5 base arrays with usable gates from 15,000 to 50,000.

This new family is compatible with Toshiba's existing channelled array families, using compatible macrocell libraries and CAD software.

If you are working on advanced products – or working to advance existing products, talk to Toshiba at one of our conveniently located Design Centers. We're leading the way. Toshiba. The power in gate arrays.

TOSHIBA. THE POWER IN GATE ARRAYS.

	3 MICRON CHANNELLED ARRAY	2 MICRON CHANNELLED ARRAY	1.5 MICRON CHANNELLED ARRAY	SEA OF GATES COMPACTED ARRAY™
SERIES	TC15G	TC17G	TC19G	TC110G
GATES	880 to 6,000	540 to 10,000	3,200 to 10,000	2,100 to 50,000
DESIGN RULE	3μm	2μm	1.5μm	1.5μm
GATE SPEED	2.5ns	1.5ns	1.0ns	0.7ns
PART NO's	6	9	5	5
AVAILABILITY	NOW	NOW	NOW	NOW

All Si-Gate CMOS double layer metal. Compacted Array is a trademark of LSI Logic Corporation.

AREA SALES OFFICES:EASTERN AREA, Toshiba America, Inc., (617) 272-4352; **NORTHWESTERN AREA**, Toshiba America, Inc., (408) 244-4070; **CENTRAL REGION**, Toshiba America, Inc., (214) 480-0470; **REPRESENTATIVE OFFICES: ALABAMA**, Montgomery Marketing, Inc., (205) 830-0498; **ARIZONA**, Summit Sales, (602) 998-4850; **ARKANSAS**, MIL-REP Associates, (512) 459-8602; **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; **DISTRICT OF COLUMBIA**, ArboTek (301) 825-0775; **FLORIDA**, Sales Engineering Concepts, (305) 834-7656, (305) 426-4601, (813) 963-1042; **GEORGIA**, Montgomery Marketing, Inc., (404) 447-6124; **IDAHO**, Components West, (509) 255-6224; **ILLINOIS**, Carlson Electronic Sales, (312) 956-8240, R.E.P.S., (215) 529-5711; **INDIANA**, Leslie M. DeVoe Company, (317) 842-3245; **IOWA**, C.H. Horn, (319) 393-8703; **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; **MARYLAND**, ArboTek, (301) 825-0775; **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,

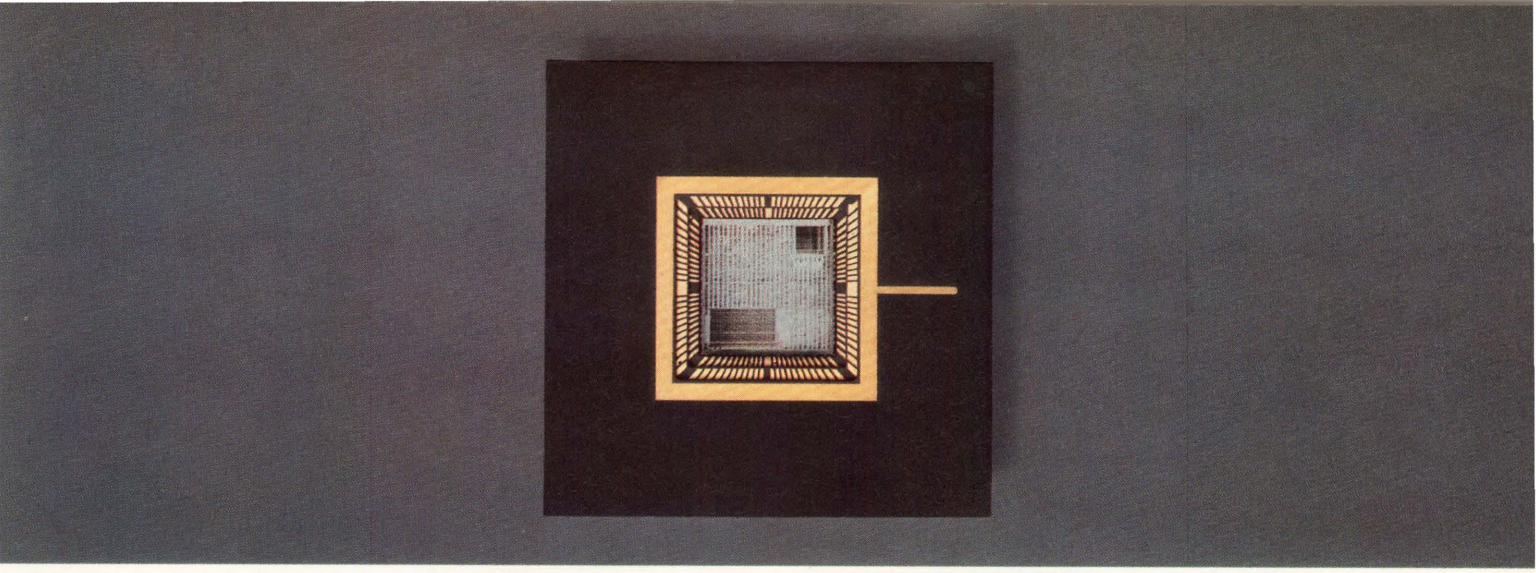

RAM & ROM Memory In CMOS ASICs.

Now add custom-built memory to your standard cell designs. Add ROM in various word and bit combinations up to 16 kbits. Add RAM in word and bit configurations up to 4 kbits.

Based on Toshiba's proprietary 2-micron, double-metal CMOS process, this new series of devices is compatible with existing and future Toshiba standard cell and gate array families.

These advanced devices integrate up to 10,000 gate equivalents with typical speeds of 1.5 nanoseconds for a loaded 2-input/2-output NAND gate.

Look to Toshiba for the newest advances in custom LSIs and for the manufacturing strength that means unmatched production capability.

For complete details, call your Regional Sales Manager at one of Toshiba's Regional Sales Offices:

NORTHWESTERN: San Jose, CA (408) 244-4070.
SOUTHWESTERN: Newport Beach, CA (714) 752-0373.
NORTH CENTRAL: Chicago, IL (312) 945-1500.
SOUTH CENTRAL: Dallas, TX (214) 480-0470.
NORTHEASTERN: Burlington, MA (617) 272-4352.
SOUTH-EASTERN: Atlanta, GA (404) 493-4401.

TOSHIBA. THE POWER IN STANDARD CELLS.		
SERIES	TC21SC	TC22SC
GATES	MAX. 10K	MAX. 10K
CAD	VLSI CAD-I	VLSI CAD-IA
DESIGN RULE	2.0 μm	2.0 μm
SPEED	1.5 ns	1.5 ns
MACRO FUNCTIONS	TC17G Gate Array MACRO Cell	TC21SC MACROs, plus: RAM 4K ROM 16K Functional MACROs 74 Series
AVAILABILITY	NOW	NOW
All Si-Gate CMOS double layer metal.		

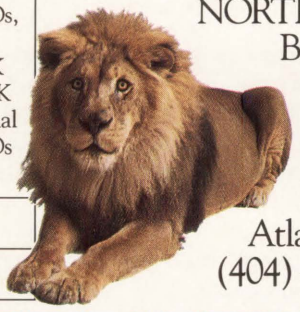

TOSHIBA AMERICA, INC.

(316) 744-1229; **MONTANA**, Components West, (206) 885-5880; **NEVADA**, Elrepco, Inc., (415) 962-0660; **NEW ENGLAND**, Datcom, Inc., (617) 891-4600; **NEW HAMPSHIRE**, Datcom, Inc., (617) 891-4600; **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) 295-6549, Steffen & Associates, (412) 276-7366; **RHODE ISLAND**, Datcom, Inc., (617) 891-4600; **TENNESSEE**, Montgomery Marketing, Inc., (615) 984-7080; **TEXAS**, MIL-REP Associates, (512) 459-8602, (713) 444-2557, (214) 644-6731; **UTAH**, Straube Associates Mountain States, Inc., (801) 263-2640; **VERMONT**, Datcom, Inc., (617) 891-4600; **VIRGINIA**, ArboTek, (301) 825-0775; **WEST VIRGINIA**, Steffen & Associates, (419) 884-2313; **WASHINGTON**, Components West, (206) 885-5880, (509) 255-6224; **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) 726-1452.

CIRCLE NUMBER 27

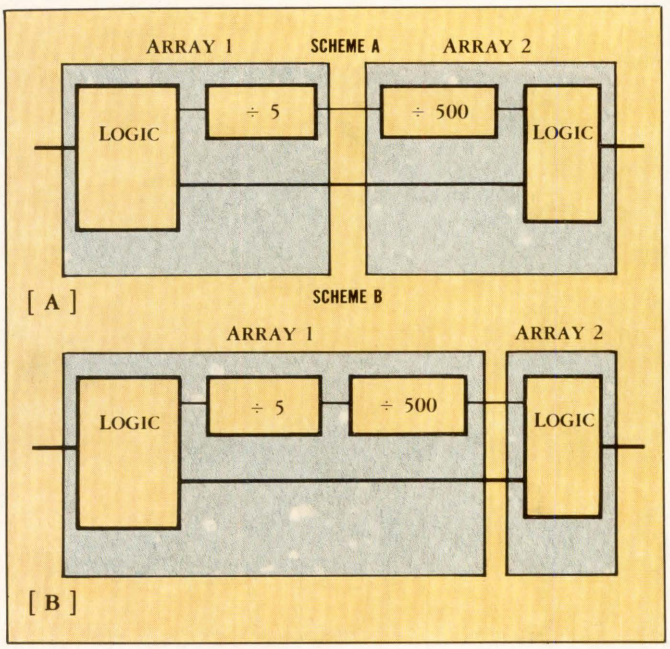

Figure 2. Partitioning can reduce drive current requirements. In scheme A the divide-by-5 counter must drive the package capacitance of another chip as well as the capacitance of the divide-by-500 counter; the divide-by-5 counter in scheme B must merely drive the capacitance of the divide-by-500 counter.

500 test vectors, rather than using 2,500 in front of array 1.

■ A Good Test Strategy

A good test strategy is also key. It is wise, for example, to develop a test procedure in which critical parts are tested first. This approach saves the trouble of making full tests on parts that would fail early with a critical test. Generally, I/O dc parametric tests are conducted first, because failure of the I/O will cause the entire chip to fail. Quiescent current testing is a good choice next, since high levels of quiescent current indicate internal shorts.

One can apply fault grade vectors next, using first those that are likely to detect most faults. Finally, one can use timing, power, current, and noise vectors.

To be most effective, one should use test vectors that will stress sensitive parts of the chip and a test sequence that can reveal likely faults early. As an example, one would make early tests in areas of high gate density.

Additionally, one would want to be sure that tests provide information. For example, instead of switching all I/Os simultaneously, which could result in chip failure without much added knowledge, one could switch one I/O, then two I/Os, and so on, until the specifications are reached or the chip fails.

One can of course test noise-sensitive circuits by injecting noise. Current surges should be injected into circuits that are likely to fail in the presence of such surges.

■ Some Testers Are Too Slow

Some testers do not provide high-speed switching. Since high-speed switching makes for high current densities, these testers do not stress nodes to their current density limits. It would therefore be wise to apply ATE signals at V_{in}(high) and V_{in}(low) levels to stress those nodes, the power mesh, and the power ring with high current densities. Another approach would be to apply midsupply levels (2.5 V) at ATE speeds. These tests place the internal gates in their linear region, which leaves them partially on (and hot) all the time.

One can determine the worst-case current requirement from simulation data. The worst case shows up in listing files as a transition from all 0's to all 1's or the reverse. There are postprocessing programs that can examine the listing files for such worst-case transitions and call them out, making it unnecessary for a user to go through the listing files visually. Listing files can also locate areas of excessive current density and can calculate the density.

Other tests should be used to evaluate the effects of varied propagation delays that can result from processing problems. Because the characteristics of the p-channel transistors are not identical with those of the n-channel transistors, there can be four different transition times. One can have fast transitions from high to low, slow transitions in both directions, and fast transitions in one direction and slow in the other. Therefore it is necessary to provide vectors for all four possibilities and to test with these vectors under temperature and voltage extremes.

■ Testing the Tests

Some testers do not provide the variety of stimuli available from CAE simulators. It is therefore necessary to make sure that the simulator's test vectors are compatible with the ATE used. If the vectors are converted for use with the ATE, one should be certain that no timing errors are introduced.

A basic difference between a simulator stimulus and an ATE stimulus is that the latter is synchronized to a clock and has a fixed pulse width. The simulator's test vectors do not need to be synchronized and do not need a fixed pulse width.

It should be obvious here that a reliable design depends on high levels of testability, and that normally calls for extra pins on the chip and extra logic within the chip. One might include, for example, multiplexers or scan path circuitry to detect hard-to-access faults. These requirements involve trade-offs of reliability and testability against chip size and complexity.

After all the test results have been reviewed, one can move on to placement and routing. Once a tentative layout is prepared, one should evaluate its effects on performance, testability, and the reliability of the chip.

Remember that a dense gate array has shorter interconnect lengths and shorter distances between interconnects, thus lending itself to higher speeds. However, a dense array might be more susceptible to shorts and to noise pickup from nearby high-current interconnects. On the other hand, a low-density array requires higher currents to achieve the same speed. So there are trade-offs among chip size (and therefore cost), speed, and reliability.

One should naturally select place-and-route software that will provide maximum gate usage, distribute power as evenly as possible, and minimize noise coupling.

After a design is completed and put into production, one can estimate the success from a reliability viewpoint by determining the failure rate. There are several approaches. One of them, the widely used Wadsack model, is expressed as:

$$\text{Field Failure Rate} = (1-Y)(1-F)$$

where Y is the yield of good dice to total dice on a wafer and F is the level of fault grading—the ratio of the number of detected faults to the number of possible faults. Thus, if the yield is 70% and the fault grade is 70%, the expected field failure rate is 0.09, or 9%. ■

References

BANERJEE, P., AND J. A. ABRAHAM. AUGUST 1984. "Characterization and Testing of Physical Failures in MOS Logic Circuits," *IEEE Design and Test of Computers.*

COURTOIS, B., AND D. BASCHIERA. OCTOBER 1984. "Testing CMOS: A Challenge," *VLSI Design.*

ONE PROCESS...

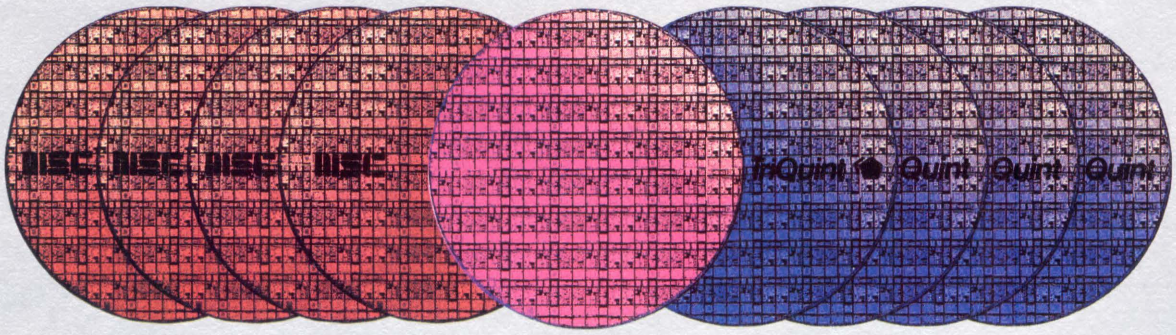

...TWO SOURCES

Announcing the ultimate solution to GaAs IC second sourcing.

Now, for the first time, you can get identical GaAs ICs...from separate sources!

As a result of an extensive, year-long collaboration, MSC and TriQuint have established uniform GaAs processes—used at both locations—that yield devices with the same characteristics and performance.

It's a unique partnership that gives you unprecedented convenience and confidence...and gives exciting new meaning to the words "second source." Because now you can use two dependable suppliers as if we were one. In fact, since you'll be using the same design rules, you can supply the same GDS II™ tapes or masks to either foundry!

Extensive testing proved the results. Let us prove it to you.

Call us and get the details straight from the sources mouth.

MSC and TriQuint.
Two sources are better than one.

msc
MICROWAVE SEMICONDUCTOR CORP.
A Siemens Company
100 School House Rd., Somerset, N.J. 08873
(201) 563-6300

TriQuint SEMICONDUCTOR
A TEKTRONIX COMPANY
Group 700, P.O. Box 4935, Beaverton, OR 97076
(503) 644-3535

GDS II is a trademark of GE/Calma.

CIRCLE NUMBER 28

Enjoy the "Crisp" Sights and Sounds of High Performance and Reliability with DSP-Technology

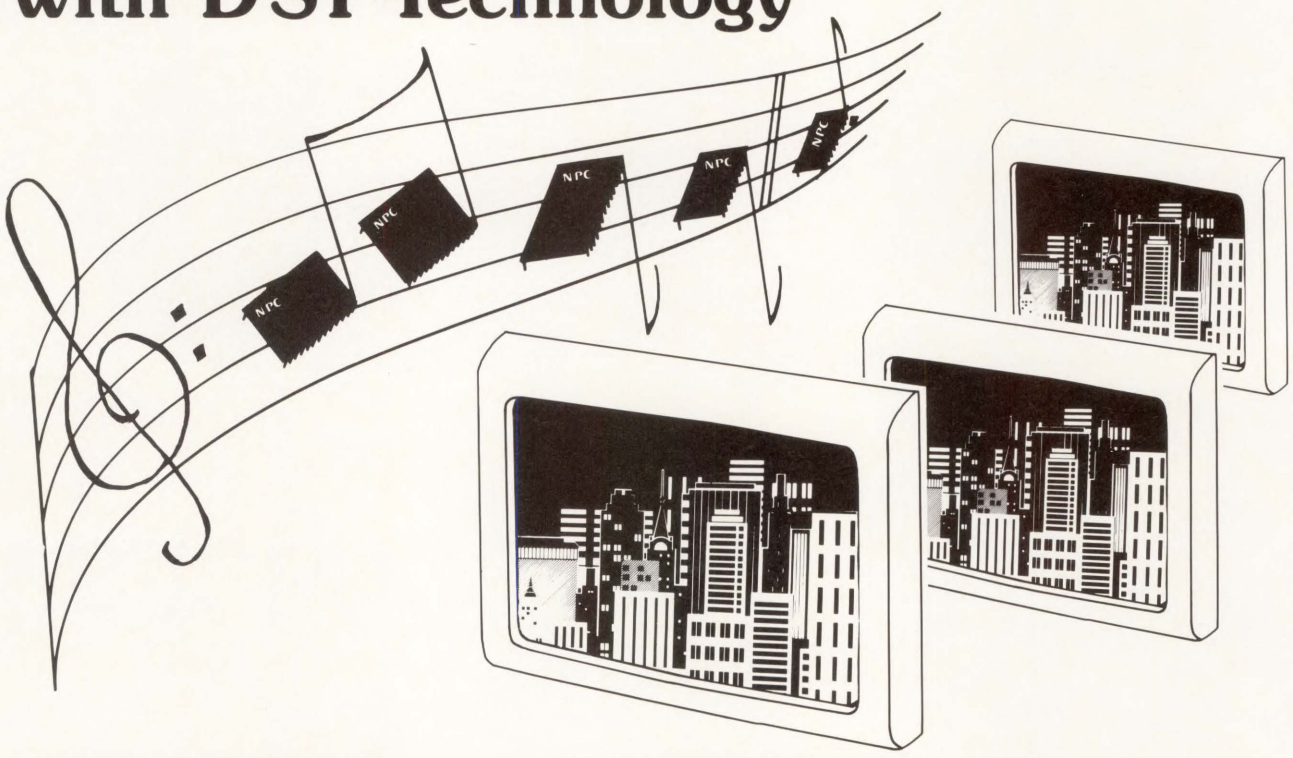

DIGITAL SIGNAL PROCESSING

SM5831F
◄ Digital Video Filter ►
4 to 8 tap variable FIR
f clk =15 MHz
Package: 64PIN FPP

SM5828
◄ Video Shift Resistor ►
8 bit word, 1 to 128 variable step
f clk = 20 MHz
Package: 24PIN DIP

SM5805
◄ PCM Audio Digital Filter ►
121th order filter × 2 ch.
Package: 28PIN DIP

SM5808
◄ 8 × 8 bit Multiplier ►
t mac = 45 nS
Package: 48PIN DIP

SM5810
◄ 16 × 16 bit Multiplier ►
t mac = 65 nS
Package: 64PIN DIP/68PIN PGA

SPECIAL GRAPHICS & AUDIO

SM6100
◄ 8 bit A/D Converter ►
Conversion time 2.1 s

No S/H required
P-bus compatible
Package: 20PIN DIP

PLL2001
◄ PLL Frequency Synthesizer ►
f in = up to 200 MHz
Package: 16PIN DIP

SC6433
◄ B/W TV Camera Sync. Generator ►
NTSC or CCIR
Packge: 22PIN DIP

SM8530B
◄ IEC Bus Interface Controller ►
Meets IEEE-488 (GP-IB)
Package: 40PIN DIP

SM5804
◄ Compact Disc Digital Filter ►
FIR (80 stage + 15 stage)
Package: 60PIN FPP

SM5130
◄PLL IC for Cordless Telephone ►
46 - 49 MHz
Package: 16/18PIN/SOP

For Technical Information, and Orders

NPC

NIPPON PRECISION CIRCUITS LTD.
USA, Canada, Australia and Asia Operations
c/o Seponix Corporation
2151 O'Toole Ave., Suite L
San Jose, CA 95131 U.S.A.

TEL: (408) 922-0133
FAX: (408) 922-0137
TLX: 9102405802/SEPONIX USA
TOLL FREE: **1-800-237-4590**

◄ Seeking New Sales Reps. in USA & Canada, Pacific & Asia. ►

CIRCLE NUMBER 29

Custom ICs receive their just reward

Seponix "Full Custom IC Designs" don't believe in Band-Aids.

Are you satisfied with your "Gate Arrays, Semi-Custom ICs?"

Our experienced full custom IC design team integrates your ideas onto a silicon chip. Our extensive application and technology covers:

- Digital & Analog on chip ICs
- Telecommunication ICs
- Data acquisition ICs
- High voltage driver/ peripherals ICs

Our latest VLSI process technology will guarantee high reliability and low/low cost in:

- CMOS
- DMOS
- Bipolar
- ECL
- Others

For Information:
Please write and send your specifications and requirements for your custom IC to:

Attn: Engineering Department
Design Center

SEPONIX CORPORATION
2151 O'Toole Ave., Suite L
San Jose, CA 95131 USA

TEL: (408) 922-0133
FAX: (408) 922-0137
TLX: 9102405802/SEPONIX USA
TOLL FREE: **1-800-237-4590**

"Are you looking for higher yield in your VLSI production?"

Your solution may be as simple as obtaining the New Generation Beads Blaster.

- ► Discover now the secret of Japan's top VLSI production high yield & high quality VLSI products through the use of Seponix's New Beads Blast Systems for semiconductor production equipments (Ion Implantation, CVD, Vacuum Evaporator Sputtering, Wafer Processing and others).
- ► These new Dry Blaster Systems have been designed for the ultra cleaning purpose of semiconductor manufacturing equipment without any mess during impurity contamination and during cleaning processing. Most major Japanese VLSI manufacturing companies are now enjoying high yield production of VLSI products due to detailed emphasis on the use of cleaning systems to insure very high yield.
- ► Seponix's new Beads Blast System can be used in either gray areas or clean room atmospheres.
- ► The new Beads Blast Systems will guarantee you with low maintenance requirements.

General features shown in some of our standard beads blast systems:

—Variable workpiece sizes
—Auto transfer mechanism
—Applicable in class 1000 cleanroom
—Toxic and poisonous material handling possible
—Easy disposal of waste abrades

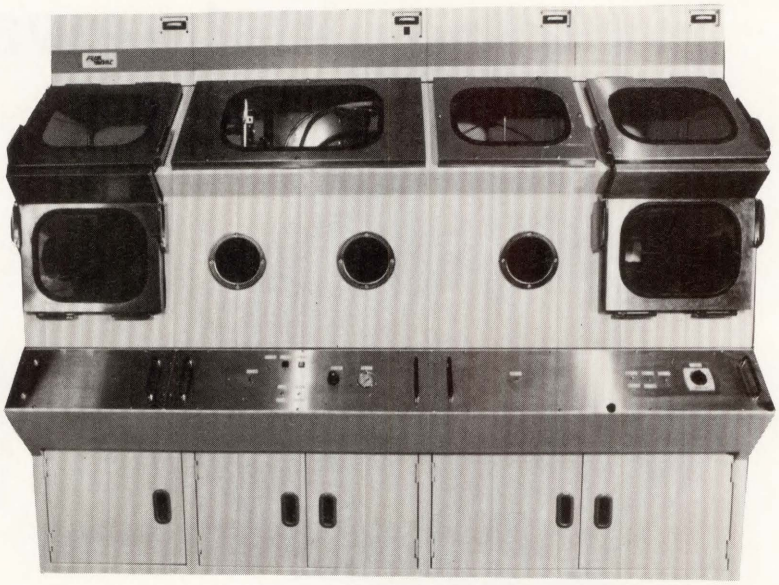

Model: System 2300

—Some of our systems have been developed so that the operator may disassemble, blast clean, air blow, sonic clean and re-build the work piece all in one unit.

* Standard Clean Beads Blast Systems offered by Seponix.

(Prices subject to change without notice)

MODEL TYPE	MAIN SPECIFICATION	APPLICATION PURPOSE	** STANDARD UNIT PRICE PRICE IN U.S. DOLLARS
FV-10	Main blast room plus one gate isolated room (GIR) with bench	General cleaning for small tools: Ion Impla., C.V.D., sputtering, jigs and parts	$ 37,500.00
FV-20	Main blast room plus one (GIR)	General cleaning for small tools: Ion Impla., C.V.D., sputtering, jigs and parts	$ 48,000.00
FV-30	Main blast room plus two (GIR)	General cleaning for small tools: Ion Impla., C.V.D., sputtering, jigs and parts	$ 75,000.00
FV-40	Main blast room with two (GIR) plus ultra sonic wafer cleaning room	General cleaning for small tools: Ion Impla., C.V.D., sputtering, jigs and parts	$ 105,000.00
TR-100	Main blast room plus two (GIR)	General cleaning for small tools: Ion Impla., C.V.D., sputtering, jigs and parts / cleanroom application	$ 105,000.00
MS-20CR	Main blast room plus one (GIR) with dust collector systems for clean room use	C.V.D. parts cleaning for big plate of 600 mm diameter in clean room	$ 96,000.00
System 2300	Main blast room with two (GIR) plus ultra sonic water cleaning room for cleanroom use with Class B	Ion Impla., C.V.D., electrodes, sputtering parts for cleaning in clean process of class 1000 VLSI production	$ 142,000.00
CS-800	Main blast room plus one (GIR) with clean dust collector hepa filter	Cleanroom use for general cleaning for small tools: Ion Impla, C.V.D., sputtering, jigs and parts/bench type	$ 60,000.00

* Seponix will also accept custom design beads blast systems and any modification of standard systems. (There will be an additional cost from the standard price.)

** All **standard unit prices** includes the initial installation fee and operational training fee at the customer's site in the continental U.S.A. and Canada. Customers located outside these areas will be quoted F.O.B. San Jose, California U.S.A.

Note: Special blast beads powder can be supplied upon request.

All beads blast systems are included with a one year limited warranty from manufacturing defects.

For Technical Information and Orders

SEPONIX CORPORATION
2151 O'Toole Ave., Suite L
San Jose, CA 95131 USA

TEL: (408) 922-0133
FAX: (408) 922-0137
TLX: 9102405802/SEPONIX USA
TOLL FREE: **1-800-237-4590**

◄ Seeking New Sales Reps. in USA & Canada, Pacific & Asia. ►

CIRCLE NUMBER 31

Why high performance designers are so excited about the new PLD 7C330 State Machine:

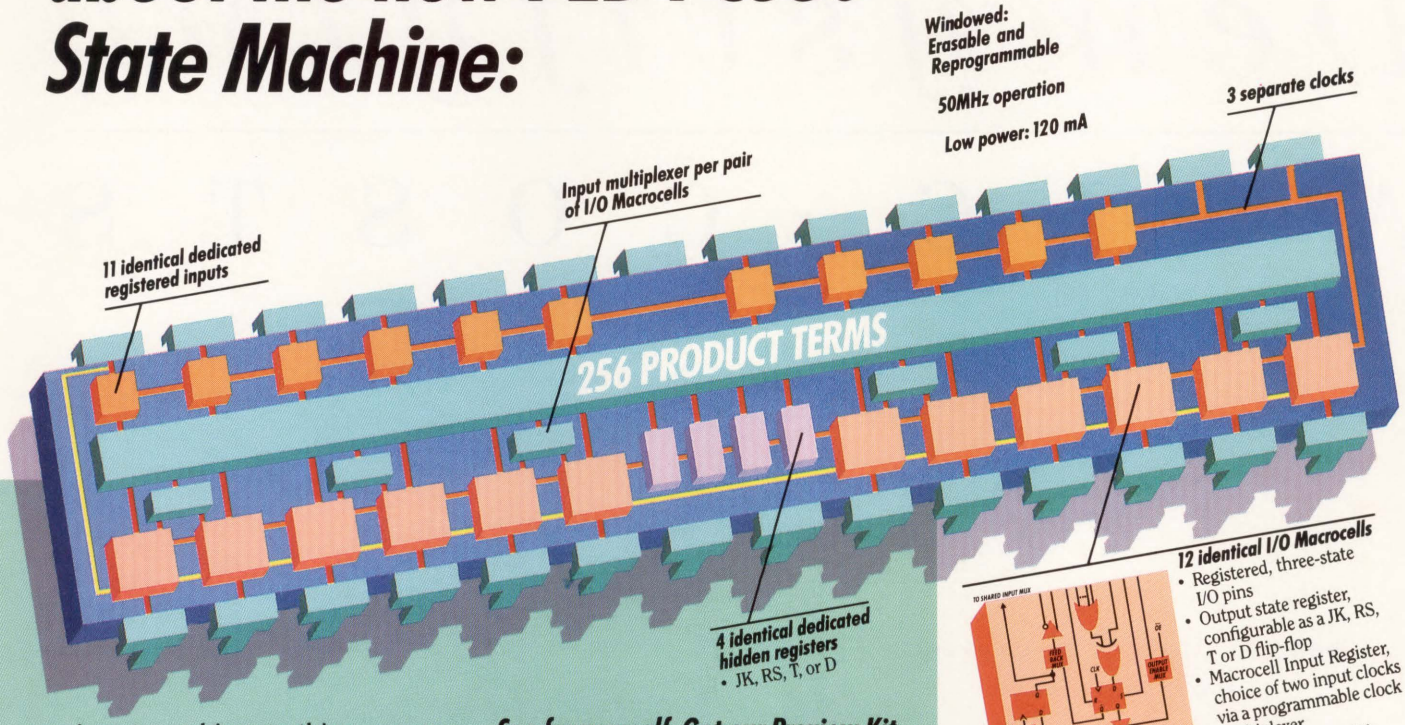

As state machines go, this one goes the fastest. With the highest functional density available.

A system that lets you design state machines that can execute control sequences at a full 50MHz without even breathing hard.

High performance designers are excited by the extensive features, led by the richness of 256 product terms.

By this part's ability to deliver twice the performance and density of previous generation PLDs, with about half the power.

By the ease of programming, using popular programming languages and machines.

And by the ability to design the highest performance state machines, with 1000 to 1500 gate functionality, using reprogrammable PLD technology.

See for yourself. Get our Preview Kit.

Our Preview Kit, with demo floppy and technical paper, even includes application ideas, like our step-by-step design example showing you how to create a 16 x 4 programmable multiplexer using a single CY7C330.

You'll also get the databook that has the information you need on this great part, and on ALL our high performance parts.

CMOS high speed SRAM.
CMOS high speed PROM.
CMOS high speed PLD.
CMOS high speed Logic.

PLD 7C330 Preview Kit and DataBook. Yours free, for a toll-free call.

DataBook Hotline:
1-800-952-6300
Ask for Dept. C103
1-800-423-4440 In CA
Ask for Dept. C103
(32) 2-672-2220 In Europe
(416) 475-3922 In Canada

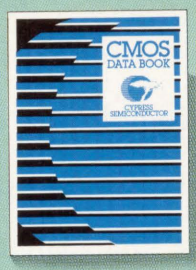

CYPRESS SEMICONDUCTOR

Cypress Semiconductor, 3901 North First Street, San Jose, CA 95134, Phone: (408) 943-2666, Telex: 821032 CYPRESS SNJ UD, TWX: 910-997-0753. ©1988 Cypress Semiconductor.

CIRCLE NUMBER 32

Demystifying ASIC·COSTS

HOWARD K. DICKEN, DM DATA INC.,
SCOTTSDALE, ARIZ.

Estimating the cost of an ASIC needn't be a mystery. It is determined by several well-known factors. By understanding the relationships of these factors, designers can plan their ASIC development to achieve the lowest ASIC cost and therefore a lower system cost.

The cost of any IC, including ASICs, is a function of the wafer-processing cost, the yields, the assembly and package costs, and the amortized tooling costs—that is, the factory cost. The selling price is the factory cost marked up by a factor of 1.8 to 2 to cover overhead, R&D, and profits (Figure 1).

IT DOESN'T HAVE TO BE A CLOAK-AND-DAGGER AFFAIR

The details of these costs and the relationships to other factors, such as the process technology and the circuit complexity, are of critical importance to the ASIC designer and user. Optimum system partitioning and total system cost are directly influenced by the cost of designing and manufacturing ASICs.

The discussion here is based on an extensive computer cost model of the semiconductor design and manufacturing sequence. (The model software is available in several forms, ranging

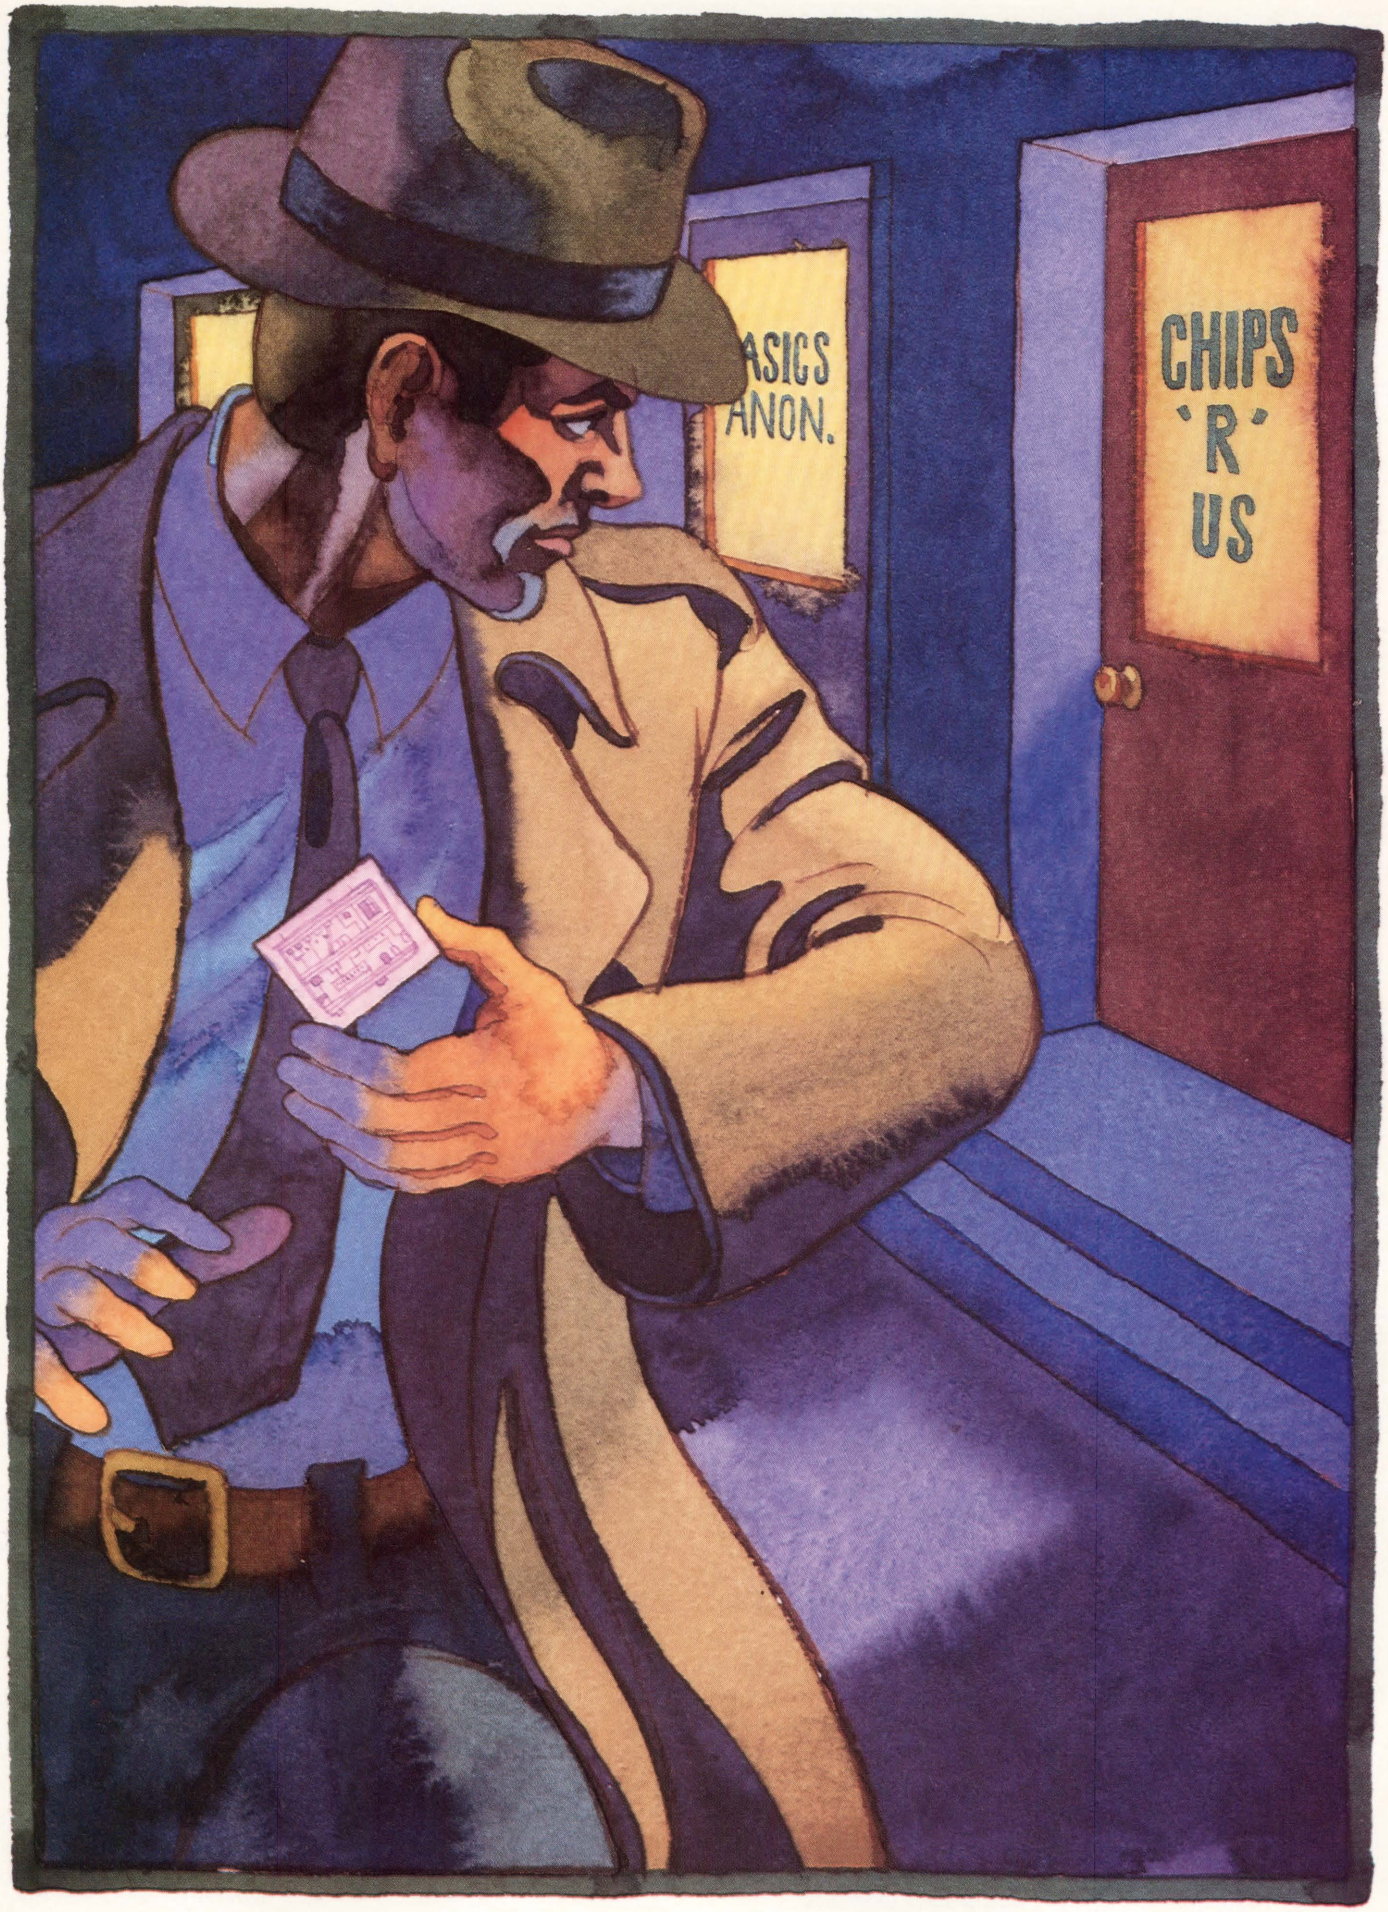

ILLUSTRATION GARNET HENDERSON

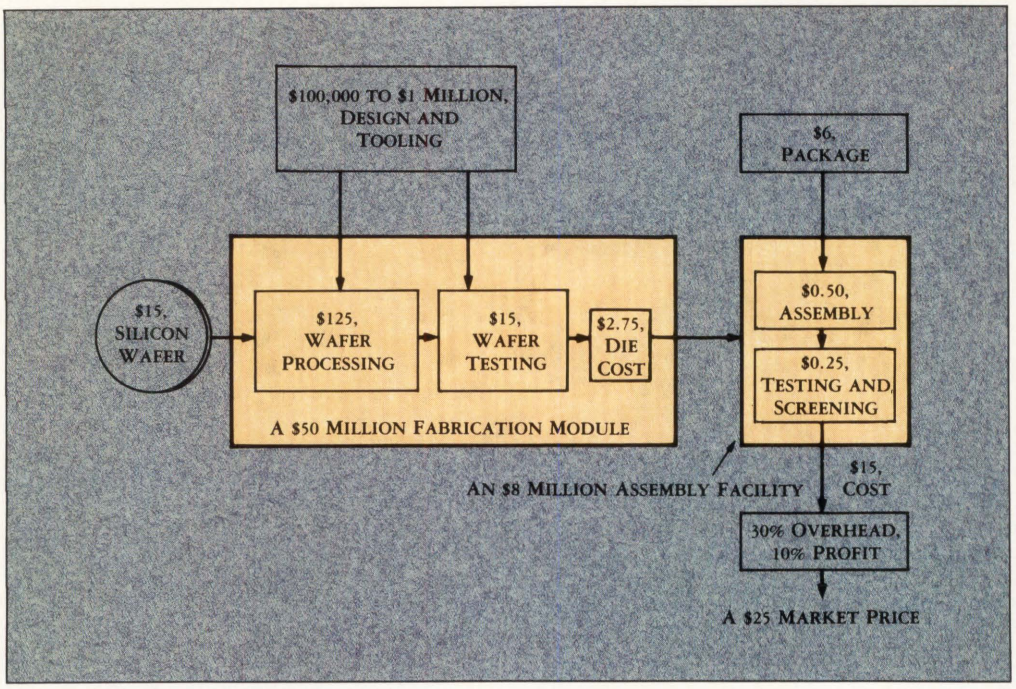

Figure 1. Costs associated with the VLSI manufacturing process.

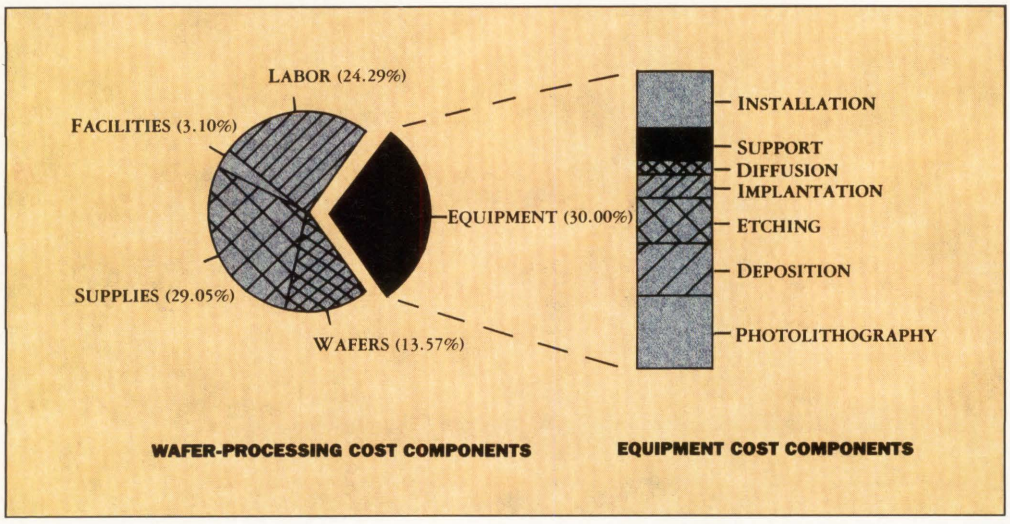

Figure 2. Contributions to the cost of wafer fabrication.

from a preprogrammed pocket computer to more detailed PC-based interactive programs.)

■ WAFER-PROCESSING FACTORS

What should you look for in a cost-effective semiconductor vendor? What are the critical factors that determine the cost and ultimately the market price of any semiconductor?

As with any product, the primary factor is production volume. For a given part type, the typical semiconductor vendor will follow a classical learning curve for production costs for several years. The slope is typically in the range of 70% to 80%. In other words, for each doubling of production volume, the cost will decrease by 20% to 30%.

The user should therefore avoid the "lonesome" technologies and parts that are not in the mainstream of industry production, unless his application justifies their use. Not only is there a major cost decrease for established technologies and high-volume parts, but also the reliability of the parts, whether standard or semicustom, improves.

In addition to the volume, wafer processing includes a series of technical factors that can have a major effect on the final market prices. Here are the key vendor-related factors and their ranges of impact on the final market price of an IC:

- Wafer diameter, 15% to 20%
- Photolithography resolution, up to 50%
- Maturity of the processing technology, up to 300%
- Size of the facility, up to 200%
- Geographical location and labor rates, up to 20%
- Effectiveness of the clean rooms and handling, up to 200%

For a modern 6-inch, 1-μm CMOS fabrication facility primarily producing VLSI circuits for commercial use, the largest cost factors are equipment, supplies, labor, and testing. Other factors, like the cost of the silicon wafer, have consistently ranged from 5% to 6% of the sales price of the product. The facility cost—the depreciated value of the building itself, including the clean rooms—runs about 3%.

The costs of a processed wafer depend primarily on the equipment, particularly for photolithography (Figure 2). For example, even with a fully loaded facility and efficient use of the equipment, over 3 cents of every sales dollar goes to photomasking equipment depreciation. This ratio assumes that steppers with prices of $750,000 to $1 million are used to achieve this process resolution. If X-ray or E-beam photolithography equipment will be required for processing below under 0.5 μm, this ratio will increase by a factor of 3 to 4 because of the higher equipment cost and lower throughputs.

One added cost factor is the installation. Typically, the cost of installation (plus setup and adjustment), plumbing, and supplying electrical connections is 15% to 20% of the equipment value. This added cost is normally capitalized and added to the depreciation costs per wafer.

Supplies are another major cost factor for wafer processing. With the requirements of submicron processing, the relative costs of the acids, source materials, and de-ionized water increase. Experiments have shown that some 30% of the wafer defects come from the chemicals and the water. The industry is continually improving the grade of the materials, but at an increasing cost.

Clean rooms are another area that is causing an increase in the relative cost of supplies and materials. Employee ma-

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

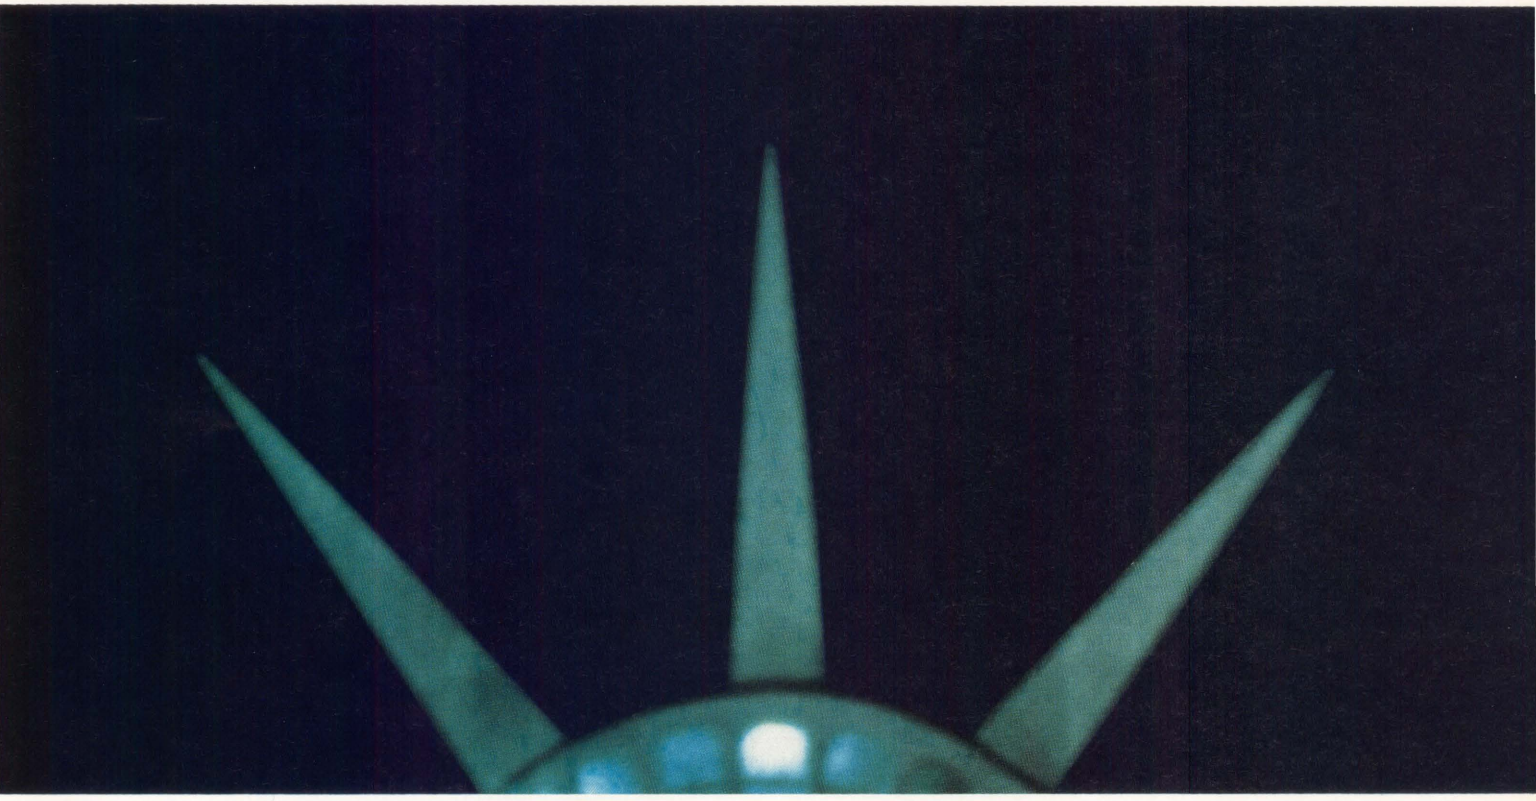

Get the full picture of your ASIC design's performance. Get the IKOS Simulation System. IKOS gives you more simulation information in less time than any other system. Simulate 15,000 gates with 34,000 vectors in 34 seconds. Or 8,500 gates with 5,100 vectors in 4 seconds. Complete with detailed timing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

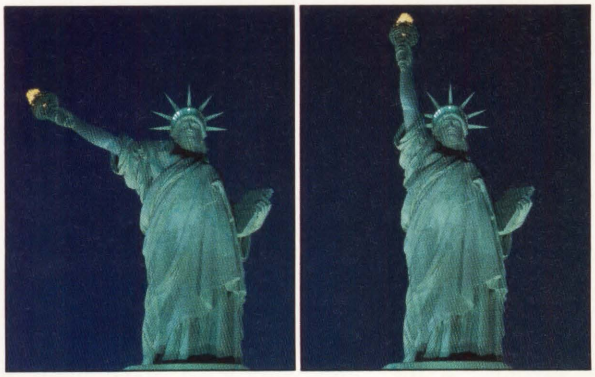

Without the IKOS Simulation System, most ASIC design mistakes show up after production.

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.

DAC Booth #565

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 33

TABLE 1. CALCULATING THE COST OF A GATE ARRAY DIE	
YIELDED WAFER COST, 125-mm DIAMETER	$218
WAFER PROBE COST	$41
TESTED WAFER COST (ADD FIRST TWO ITEMS)	$259
TOTAL DICE PER WAFER	270
YIELD FACTORS ■ DEFECT YIELD ■ MASK YIELD ■ PROCESS MATURITY ■ DESIGN MATURITY	 49% 91% 93% 95%
TOTAL PROBE YIELD	39%
NUMBER OF GOOD DICE (MULTIPLY TOTAL DICE PER WAFER BY TOTAL PROBE YIELD)	105
COST OF EACH GOOD DIE (DIVIDE TESTED WAFER COST BY NUMBER OF GOOD DICE)	$2.46

TABLE 2. CALCULATING TOTAL ASIC COST, INCLUDING PACKAGING AND TESTING	
DIE COST (TABLE 1)	$2.46
PACKAGE COST	$7.50
ASSEMBLY COST	$0.47
TOTAL PACKAGED COST (ADD FIRST THREE ITEMS)	$10.43
ASSEMBLY YIELD	89%
YIELDED ASSEMBLED COST (DIVIDE TOTAL PACKAGED COST BY ASSEMBLY YIELD)	$11.72
FINAL TEST COST	$0.26
FINAL TEST YIELD	89%
FACTORY COST (ADD YIELDED ASSEMBLY COST AND FINAL TEST COST AND DIVIDE BY FINAL TEST YIELD)	$13.46

terials, such as gowns, gloves, booties, face masks, and caps —whether disposable or reusable—are now a major cost factor in daily operations.

Larger wafer diameters can mean lower-cost parts, assuming that the vendor has had enough experience with the new line. In theory, switching from a 5- to a 6-inch line will reduce the manufacturing cost by over 15%. However, these savings are possible only after a one- to two-year process improvement cycle to obtain equivalent yields.

Where the user has a choice of technologies, as with any ASIC design approach, the process resolution used can be a significant cost factor. At any time, there is an optimum range of process resolution. For typical commercial and industrial requirements today, this value is in the range of 1.5 μm. The relation to cost is simple. For example, if the user specifies a gate array design with an old 3-μm set of design rules, the resultant die area is excessively large. (Note that a larger die not only results in fewer parts per wafer, but also increases the likelihood that each die will have a defect). Thus the cost could increase by as much as 50%. On the other hand, specifying a very high resolution process increases the cost because of lower process yields and higher equipment investments.

Unless a vendor is forward-pricing (setting costs in line with the price of the product when it matures), the price of products from a new facility or a new process recipe can be as much as three times higher than its mature selling price two years later. Very few vendors can debug a new process and achieve typical yields in less than 12 to 18 months. Some new facilities have taken more than two years to reach full-production economies of scale. This experience curve also applies to other changes, like a move of the equipment or larger wafer diameter.

How large should your vendor's facility be? If we assume a normal facility loading factor of 85% to 90%, a facility should have a capacity of better than 3,000 wafer starts per week to obtain the lowest-cost production. In smaller facilities, the mix of equipment, including the need for spares, is less efficient. This inefficiency, coupled with fewer wafers per run, contributes to a higher cost. Consequently, locations and local labor rates are not the most sensitive cost factors. However, the large overhead portion, which includes R&D engineering, does bring the total labor content up to almost 45%.

One very cost-sensitive factor is the "effective defect density" of the vendor's facility, a value not readily available to most users. This factor is a measure of the number of defects (such as accumulated dust particles) per square centimeter of processed wafer that will prevent the circuit from working. Usually, there is no relation between the claimed class of the clean room and the resultant defect density.

The value of this factor, which can be considered a relative rating, can be determined for a given vendor if both the die area and the resultant wafer probe yield is known. The typical value of effective defect density for most U.S. companies is in the range of 1.5 to 2 defects per square centimeter, whereas some Japanese companies are operating in the range of 1 to 1.5. This difference could translate into a price difference of as much as 2 to 1.

■ FABRICATION COSTS

The yield at wafer probe, and thus the cost of each good die, is determined primarily by the die area and the effective defect density. If a fixed level of defects and a fixed location is assumed, all dice on any one wafer can't work, because a defect will always be present in the defined area. As the chip is designed smaller, more parts will "miss" the defects and the yield percentage will increase. This basic relation is defined by several yield-versus-area models, where the yield is typically an exponential function of the area.

The relation of the number of defects to the die area and the resultant yield can be represented by equations drawn from the models. The most basic representation is Murphy's model, where the yield is represented as follows:

$$Y = [(1 - e^{-AD})/AD]^2$$

where A is the die area and D is the number of defects per square centimeter.

Yields can of course be improved by reducing the average number of defects. The extent of the reduction in defects, however, is limited by the investment required in clean rooms and handling equipment.

Of course, too, this basic yield curve will eventually be modified downward to account for the effects of the smaller line widths of the higher-resolution processes. With 1-μm

XILINX PROGRAMMABLE GATE ARRAYS from HAMILTON/AVNET

PERFECT DESIGNS ALL THE WAY DOWN THE LINE

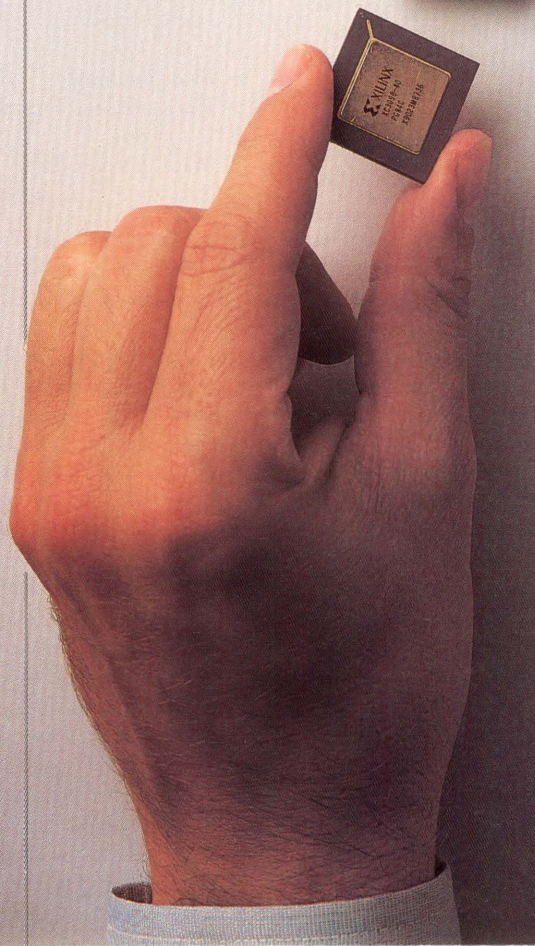

FIRST USER-PROGRAMMABLE GATE ARRAYS AVAILABLE THROUGH 57 LOCATIONS

Create complex logic designs with the Xilinx family of Logic Cell* Arrays, the only user-programmable gate arrays. You get the performance and densities you expect from a conventional gate array with none of the risk. That could add-up to substantial savings in time and money.

Hamilton/Avnet's experienced team of field application engineers are available to demonstrate the versatility and convenience of using Xilinx LCA*s. And if you want help with your design, our design center engineers are always on hand.

Xilinx's patented architecture is responsible for gate counts that range from 1200 gates in our XC2064 to 1800 in the XC2018. While the 3000 series, with enhanced architecture, features the 2000 gate XC3020.

But this is only the beginning. The 9000 gate XC3090 is soon to come, offering enough density for just about any logic application.

With speeds that run with the fastest performance microprocessors available, VLSI densities and 100% tested parts, the Xilinx Programmable Gate Array is the ideal logic device. And it's available now from Hamilton/Avnet.

So, if you want the very best in high performance, high density programmable logic, and the strongest technical support group around, get on the line with your local Hamilton/Avnet representative today.

For more information or to find the location nearest you, call toll free: 1-800-HAM-ASIC. (In Canada, 1-800-387-6879; 1-800-387-6849 in Ontario and Quebec.)

*Logic Cell and LCA are trademarks of XILINX, INC.

Hamilton Avnet ELECTRONICS AN AVNET COMPANY

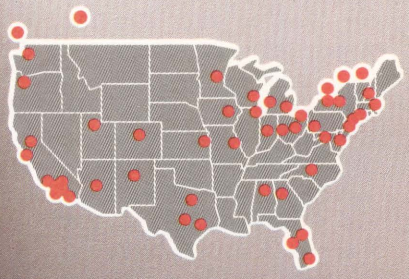

CIRCLE NUMBER 34

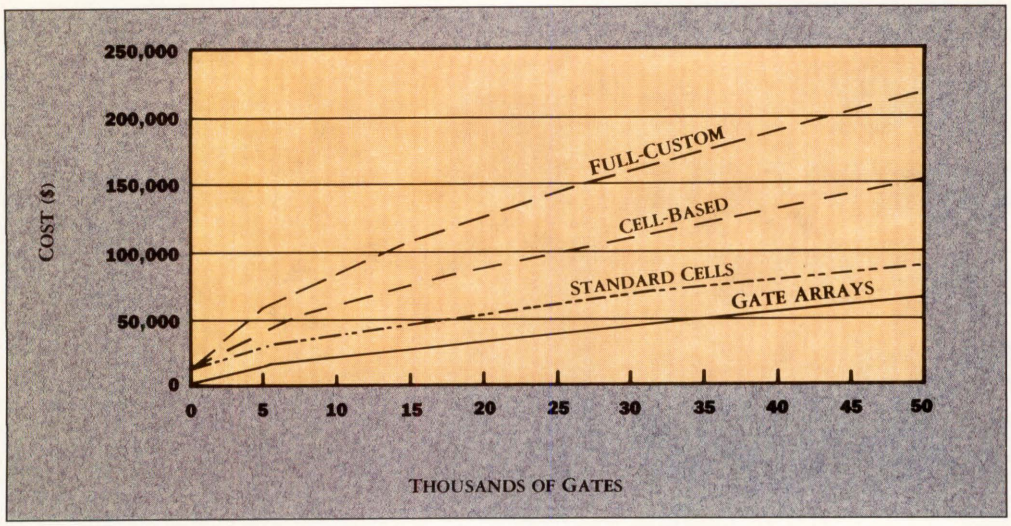

Figure 3. Tooling costs for custom and semicustom ICs.

rules, for example, a dust particle of a given size will have a greater chance of causing a defect that reduces the yield.

The total assembly operation for plastic DIPs, including the cost of the package parts (lead frame, molding compound, wires, and such) is approximately 5% of the market price. In contrast, the onshore contribution of wafer fabrication and testing is approximately 50%.

■ The Economics of ASICs

One of the hidden costs incurred when implementing a semicustom or custom IC program is the time and labor involved in the planning stages, before the design and layout phases. This cost is typically a function of the complexity of the system and the percentage of the system incorporated into the IC. Typically, the planning cost is a function of the complexity as expressed by the number of gates. These costs include the steps of both system definition and initial logic design.

The following items are typically included as part of the planning phase:

- Partitioning the system
- Preparing the logic block diagrams
- Eliminating redundancy and logic optimization
- Analyzing the logic in terms of gate count
- Determining the speed/power requirements
- Determining the package/pin-out requirements
- Selecting the technology
- Determining the transistor count per gate
- Determining the custom design approach
- Estimating the die area
- Determining the potential vendor and an alternate source

When the designer has the opportunity to choose the complexity of his ASIC, he should aim for a level that balances the yield, and hence cost, of his chip with the total system cost. Given this balancing, the optimum wafer probe yield is typically about 50%. In other words, the added cost of this modest yield, which results from high complexity, is compensated for by a reduction in system packaging cost. Based on this tradeoff, typical industry values for the optimum yield for industrial and computer applications is 40% to 60%; for consumer products, the yield should be over 60%; and for military and aerospace, the optimum yield is 10% to 20%. Yield consistently below 5% indicates processing or design problems and the wrong system partitioning.

With the increased use of ASICs, this optimum partitioning is now available for a wider range of product designs. Many system manufacturers now analyze the system partitioning in detail to determine the optimum functional complexity and consequently the die size and yield of each section of the system.

The major factor in determining both the total program cost and the optimum IC design approach is the time and labor required for the tooling phase. Typically, this phases consists of the following steps:

1. Logic checking, or simulation
2. Floorplanning
3. Placement and routing
4. Digitizing
5. Mask generation
6. Test tape development

The time and the costs (Figure 3) are a function of both the circuit complexity as expressed by the number of gates and the design approach. Full-custom design requires that every mask layer be designed and checked, whereas, in terms of masks, gate arrays require only that the interconnections be routed and checked.

The accuracy of the tooling and the chance of success for the first layout have improved, thanks to the application of various CAE/CAD checking programs that perform logic simulation, check the routing, and verify both the design rules and the masks. In fact, the success rate on the simpler gate arrays is usually good—between 70% and 100%. Larger arrays, of course, and other approaches, like standard cells and full-custom design, have lower success rates.

After the tooling is completed and the mask set is obtained, the prototype run is next. This typically consists of one wafer run (10 to 30 wafers) to process the new circuit. After initial wafer probing, 50 to 100 chips are packaged for final testing and evaluation.

If the devices meet all specifications, and if there are no changes or errors in the design or specifications, with a good yield this operation is the last step before the quantity production runs. This step, then, would also represent the last of the nonrecurring tooling charges. Typical costs for this step are $3,000 to $5,000.

However, if there is a design or logic error, or if the design or one or more specs have changed, an iteration phase is necessary. This phase includes the costs of diagnosis (in the case of an error), redesign, one or more new masks, and a second wafer-processing and assembly run. The iteration cost depends on both the complexity and the design approach.

The total time required to the first tested circuit depends primarily on the design approach. With gate arrays, only two to three mask levels must be modified and processed; therefore it is possible to obtain customized devices in less than eight weeks. At the other extreme, it can take over a year to design and evaluate complex, full-custom circuits. Additionally, with increasing complexity, the chances of requiring an iteration increases.

The basic problems in the use of an ASIC approach are:

- Correctness of the logic design
- Length of time to the first working product
- Test programs
- Design-related reliability problems, like latch-up

Most vendors and users will agree that the biggest risk fac-

HOW TO SUCCEED IN ASIC PROTOTYPE VERIFICATION. FIRST TIME. EVERY TIME.

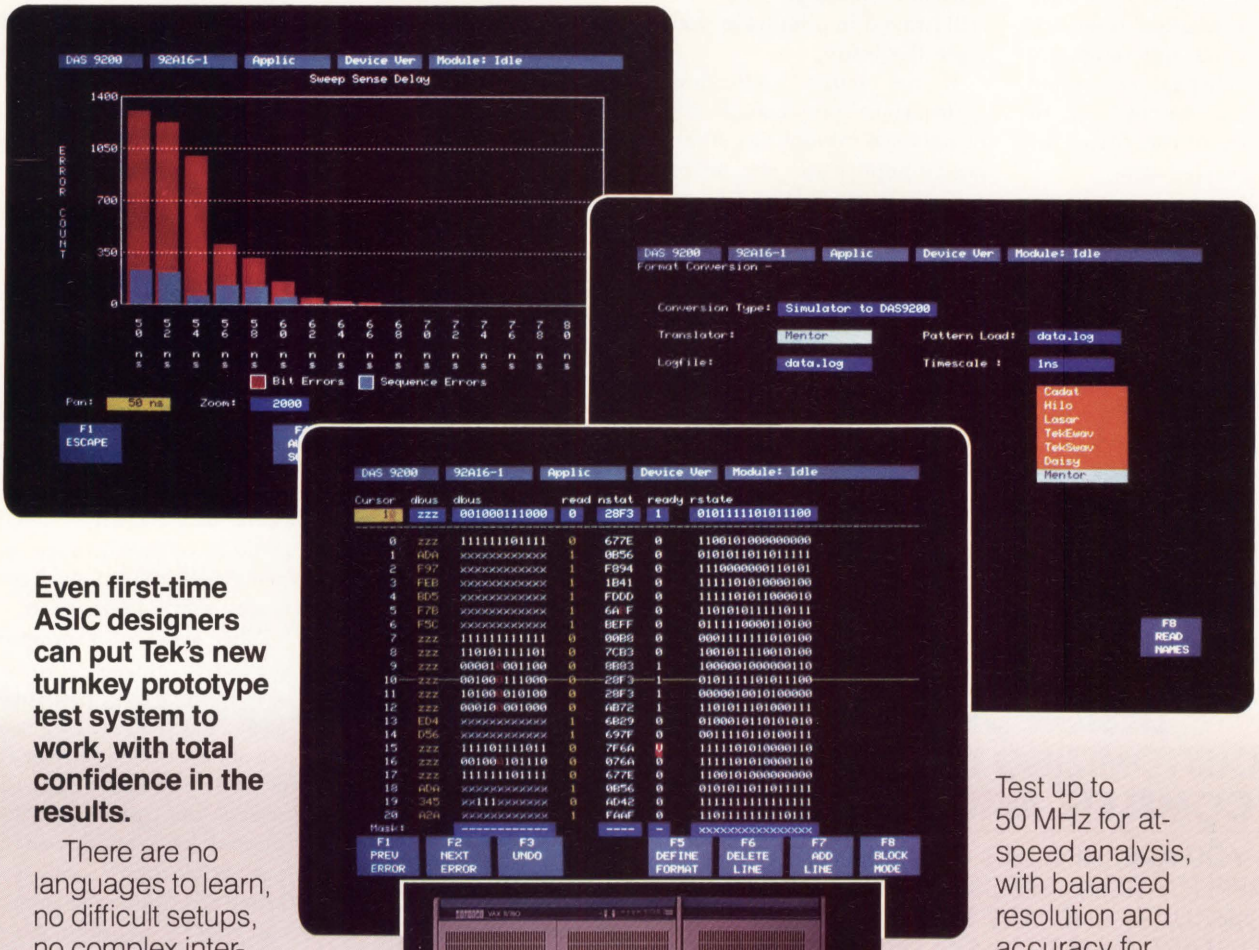

Even first-time ASIC designers can put Tek's new turnkey prototype test system to work, with total confidence in the results.

There are no languages to learn, no difficult setups, no complex interpretations, no incomplete answers.

You simply advance from one menu to the next. Use the same test vectors you developed during design. Debug with highly flexible, interactive pattern editing and timing characterization routines.

Test vector data from your logic simulator is translated and checked for tester compatibility automatically.

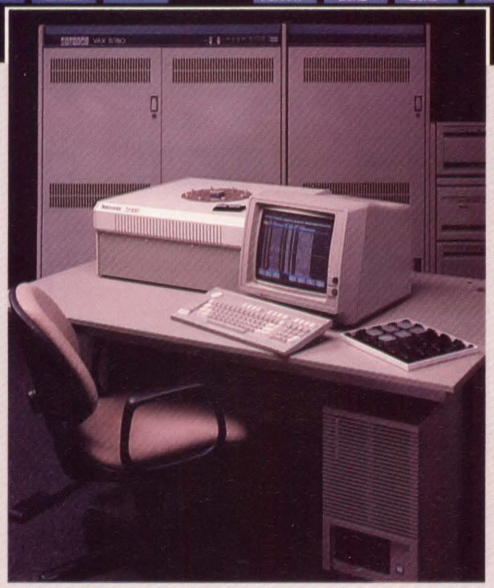

Test up to 50 MHz for at-speed analysis, with balanced resolution and accuracy for unprecedented signal measurement quality. The test fixture adapts to your DUT quickly.

The system can be easily reconfigured to help with many other complex problems—including integrating your ASIC onto the circuit board.

To learn more about Tek's DAS9200 Personal ASIC Verification System, call: **1-800-245-2036**
In Oregon, 231-1220

Tektronix®
COMMITTED TO EXCELLENCE

CIRCLE NUMBER 35

tor is correctly processing the wrong logic design. Another issue, as noted, is a change in the logic or specifications.

When an error occurs, typically most of the correction time is spent diagnosing the problem and redesigning the circuit or components. Diagnosis and redesign times (the latter whether due to an error or to a change) are, again, typically a function of the complexity of the circuit and the design approach.

The minimum production iteration time is the relative fixed times required for one or more new masks, a new run of wafers, and the assembly of the initial samples for testing. Typically, this time is six to eight weeks under normal industry conditions.

Time is still a critical factor in ASIC design. Many companies are using gate arrays to meet certain product market windows. However, in some cases even the typical turnaround time of 8 to 12 weeks is too long. This market need has created a new class of semicustom products called fast-turnaround chips, which typically rely on laser patterning to obtain the custom metallization. The claimed times have been as low as 24 hours. These approaches, although fast, are still limited in complexity and logic flexibility.

If the system's electrical performance specifications can be obtained through any of the design approaches, the choice then depends on the interacting factors of available time, degree of risk, total number of units required for the life of the program, and minimum program cost.

For any volume requirement, there is a minimum total program cost, which is the minimum of the sum of the amortized tooling and prototype costs and the cost of each circuit. For example, gate arrays have the lowest tooling costs but the highest unit cost. Thus, for a low-volume requirement, this combination would involve the lowest cost.

The lowest-cost approach, however, is overridden by the time and risk factors. For example, if a working circuit is needed in 15 weeks, it is unlikely that the choice could be full-custom, even if this approach provided the lowest program cost. These areas can be defined more accurately by making some of the cost modeling calculations shown later.

■ An Example Gate Array

What are the manufacturing costs and typical market prices associated with an average gate array? The typical array today uses a 2-μm (1.5-μm effective channel length) silicon-gate CMOS dual-metal 12-mask technology on a 5-inch wafer. This technology allows performance of up to 50-MHz clock rates and delays of 1.2 ns. The arrays are designed to operate from a 5-V supply and dissipate 20 μW/gate/MHz. Assuming an average-sized array with 3,700 gates (2,800 usable), the die size is 250 × 250 mils (62,500 mil^2).

The worksheet shown in Table 1 demonstrates how the calculation proceeds. Given a manufacturer's facility and manufacturing costs for a 5-inch wafer, you can estimate the wafer cost. Total wafer cost includes the cost of wafer probe testing. A 5-inch wafer holds about 270 dice measuring 250 × 250 mils.

Next, the total of 270 dice is multiplied by yield factors that statistically determine total yield. These include defect yield from processing; mask yield from aligning masks and handling wafers; process maturity, which reflects processing control; and design maturity, which indicates the likelihood of remaining design bugs. The product of all these yield factors is what we expect the wafer probe testing to reveal. From the numbers in Table 1, the total probe yield is 39% of 270, or 105 good dice. The cost of each good die, therefore, is \$259/105 = \$2.46.

Assume this device is placed in a 68-pin, hermetically sealed, ceramic surface-mount package. Assembly itself can create defective parts, from handling the dice and defective bond wires, so that assembly yield must be considered. Once packaged, the parts go through a more exhaustive test than at the wafer probe station. Assuming 89% yield at these two steps, the total assembled and tested cost would be about \$13.46 each. Given a markup of 1.8 to 2, the typical market price would then be in the range of \$25 to \$30 each, not including the nonrecurring tooling (NRE) costs. The tooling costs, including test tapes and masks, would be about \$17,500, which would result in a quoted price of \$30,000 to \$35,000 for the NRE portion. For a part of this complexity, the initial samples could be delivered (under normal market conditions) 8 to 10 weeks after authorization.

This article has examined the factors influencing the cost of IC manufacture, particularly ASICs, and presented an example of how to estimate the cost of a semicustom . As discussed, the cost of an ASIC is affected by a number of factors. Production volume is the prime determinant of an ASIC's cost, but by no means is it the only influence. The choice of a particular vendor imposes a number of production-related technical constraints, such as process technology maturity and wafer size. A vendor's manufacturing procedures also dictate efficiency of the entire process. Then, there are elements under closer control of the designer. Careful system partitioning, for example, can play an important role. In general, careful planning prior to design and layout is key. ■

About the Author

Howard K. Dicken *is the president of DM Data, a high-technology consulting firm. He earned his BSEE at Iowa State College in 1956 and his MSEE at Iowa State University in 1961.*

MILITARY 256K RACE

EDI—First from the Start

Since 1983, when EDI introduced the first Military 64K SRAM, the company has been first in the race to market with the latest generation, highest density, CMOS MIL-STD 883 devices.

With a new family of fast 256K Monolithics, DESC qualified devices, and a new JEDEC standard Megabit Module, EDI is maintaining its' leadership position with the widest array of Military CMOS Static RAMs available today.

EDI's 256K Gold Medalists

256Kx1 35, 45, 55ns
64Kx4 35, 45, 55ns
32Kx8 45, 55, 60ns

And, all are available in a variety of ceramic, JEDEC standard DIP and LCC Packages.

EDI
The future...today.

Electronic Designs Incorporated
Corporate Headquarters: 42 South Street, Hopkinton, MA 01748 USA (617) 435-2341, TFX: (617) 4356302, TLX: 948004
Electronic Designs Europe Ltd.,
Shelley House, The Avenue, Lightwater, Surrey GU18 5RF United Kingdom, 0276 72637, TFX: 0276 73748, TLX: 858325

MODELING COMPLEX SYSTEMS

a graphic approach to CASE using state diagrams

ROBERT A. TIERNEY, I-LOGIX INC., BURLINGTON, MASS.

BEHAVIORAL models are often used by VLSI system designers to represent portions of a system during the design process. Created using a high-level language like C or Pascal, behavioral models are compiled and linked to the CAE logic simulator to allow the detailed portion of the system design to interact with the behavioral model. Because behavioral models can be extremely complex, outside contractors are often relied on for the development of the models. The process, however implemented, is always time-consuming, expensive, and error-prone.

CASE SYSTEMS CAN REPRESENT THE FLOW OF DATA AND CONTROL JUST AS CAE SYSTEMS SHOW CIRCUITS

Advances in the emerging field of computer-aided software engineering (CASE) may provide a more efficient approach to behavioral modeling. It is now possible to graphically represent an accurate behavioral model and link it with an electronic design.

■ WHAT IS CASE?

Like CAE tools, CASE tools use sophisticated graphics workstations. They provide a means for the user to represent the flow of data and control, much the way CAE workstations show chips and circuits. Some of the more sophisticated CASE tools have the ability to represent real-time systems that must work in interrupt-driven environments. Tools of this caliber incorporate timing characteristics, which are vital to a model of a large system.

One such tool is Statemate, a workstation-based system engineering tool for the specification, simulation, and design of complex, real-time systems (Harel, 1987). Unlike most CASE tools, this new tool can represent not only the functionality and structure of complex systems but also their behavior (Figure 1). With it, designers can create a logical rapid prototype of their system and then execute it to verify the system's performance.

The system to be designed and analyzed is described with Statemate using three interrelated graphical languages: activity-charts, which are reminiscent of hierarchical data-flow diagrams; statecharts, which extend state transition diagrams; and module-charts, which represent architecture.

The purpose of the activity-chart is to define what processes are to take place in the model. This "view" of the model shows where control and information comes from and goes to. It uses various types of boxes and arrows

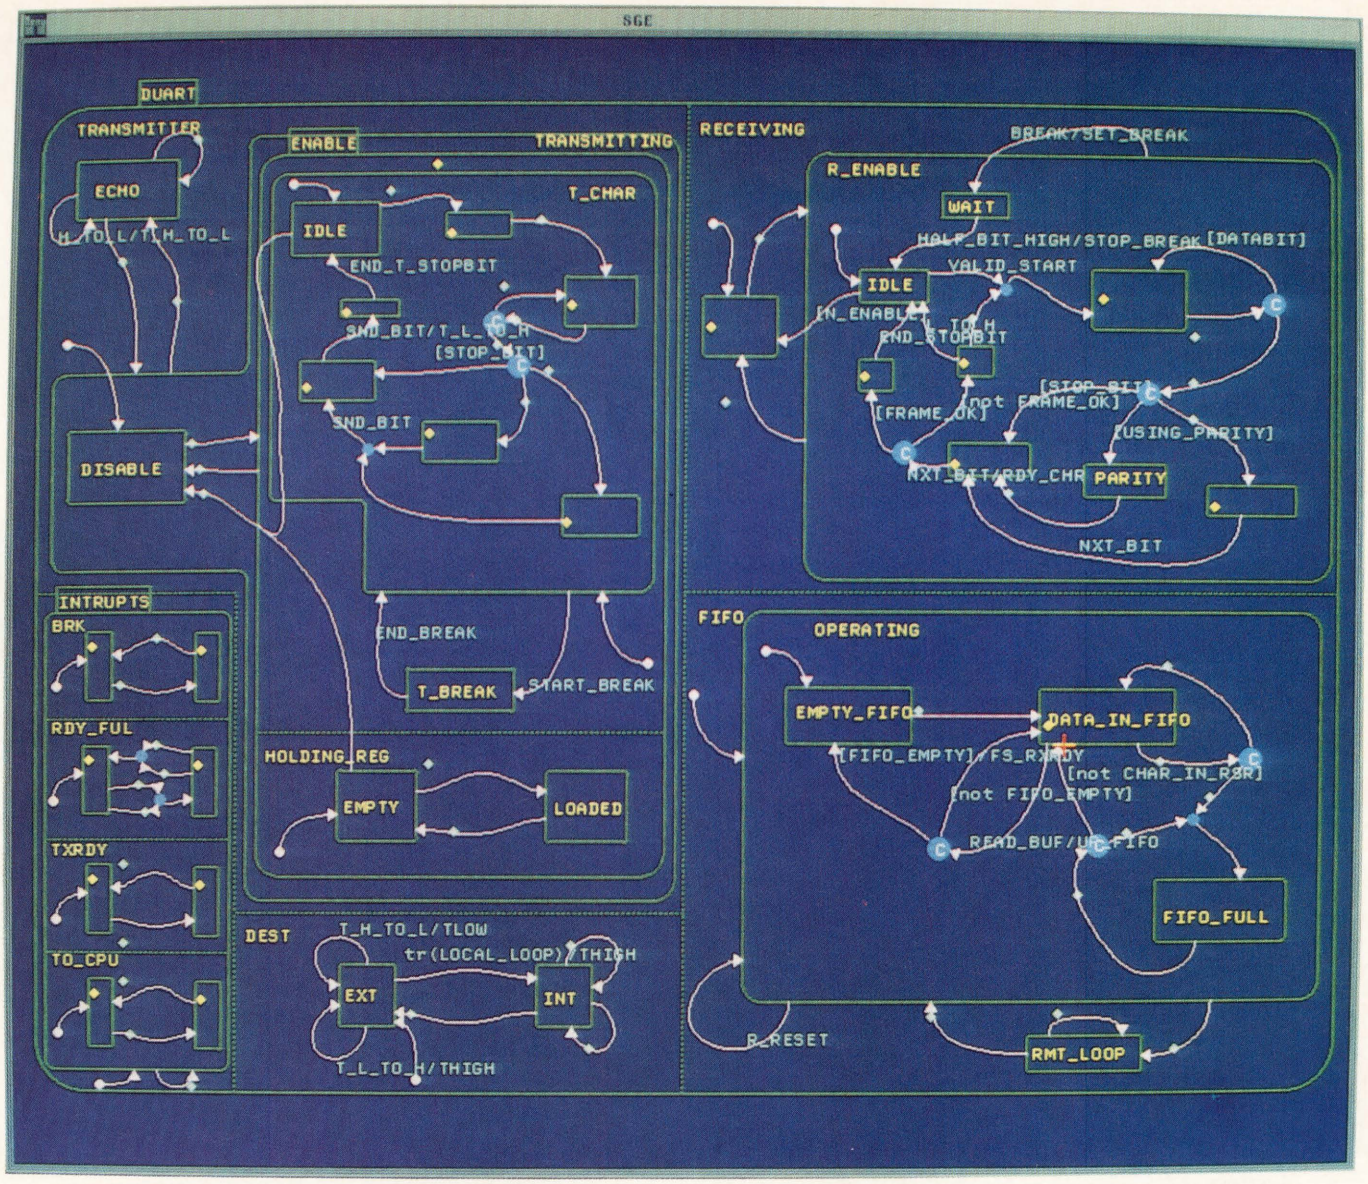

Figure 1. Statemate enables designers to graphically express complex VLSI devices at the behavioral level. The screen display shows a statechart of the behavioral specification for Signetics' SCN2681 dual UART chip.

(Figure 2); each type has a specific meaning with respect to the model's functionality.

The large box represents the high-level activity being modeled. Within this box are sets of activities (boxes) that are used to define the model. The surrounding dotted boxes represent the external or environmental activities that the model will interact with. In the example that we will consider shortly, the dotted boxes would represent the logical interface between the model and the rest of the system.

To interconnect activities, which are represented by boxes, we make use of data and control flows, represented by solid and dashed arrows, respectively. A data flow represents data items that flow between activities. A control flow represents activation commands (such as *start, stop, suspend, resume*) or activity status feedback.

In a complex model it may not be possible or even practical to represent all the activities on one level. For this reason, Statemate makes use of hierarchical descriptions. Hierarchical descriptions enable a user to define a group of high-level activities (functions) and then create lower-level activities as he or she continues the modeling process.

Once the multilevel activities and data and control flows have been defined, we have in effect a visual mapping of where our model will accept, process, and then pass on data.

To use an example, let's design a very simple model for an ALU. This device may have data (information) and control signals (instruction set) as its I/O, and it may be further defined as a group of lower activities that do the following:

- get instruction
- get data
- process data
  - add
  - subtract
  - multiply
  - divide
  - etc.
- output data

There is also a data store called "TEMP."

Each of these activities represents a function for the device (see

AN ACCURATE BEHAVIORAL MODEL CAN NOW BE GRAPHICALLY REPRESENTED AND LINKED WITH AN ELECTRONIC DESIGN

(Forward Thinking, Re

It takes more than desk-changing for two companies to merge. They need to have a shared vision of the world and the future.

With AMD and MMI, the calling was clear. To produce advanced programmable logic devices for a wide range of needs. To lower system costs. To develop higher performance designs. And to get your products to market as soon as possible.

We're delivering on this promise of the future in myriad ways. With a product line that doesn't stop halfway. With the first 10ns parts for designs where speed's important. In CMOS, we have the broadest line of parts, including the industry standard, C22V10. And for high density, no one can beat our programmable gate arrays. We've got the best selection of PAL® parts around. (And no one sells more to the military.)

We're bringing new meaning to the term "product availability." Our combined manufacturing capacities mean we have more PAL chips coming out of our pipeline. Ready to be put into yours.

And more Field Application Engineers give you more answers to get

sponsive, Meticulous)^2

projects up and running.

Because it's impossible to create great designs without great software, we also offer PALASM® software, the industry standard. We're investing heavily to keep PALASM as current as our hardware.

With the best network of distributors, we can now offer a whole new concept in distributor support: TestPro Centers. Now you can speed through programming and testing at authorized centers.

Just bring your logic pattern to your closest TestPro Center. They'll program and they'll test the parts to factory quality standards. Relieving you of the burden of test vector generation.

And we'll always be here with immediate and expert help in getting your design moving. Just call AMD at (800) 222-9323.

Or write Advanced Micro Devices, Inc., 901 Thompson Place, P.O. Box 3453, Sunnyvale, CA 94088.

One try and you'll be sharing our vision, too.

Advanced Micro Devices

Monolithic Memories MMI

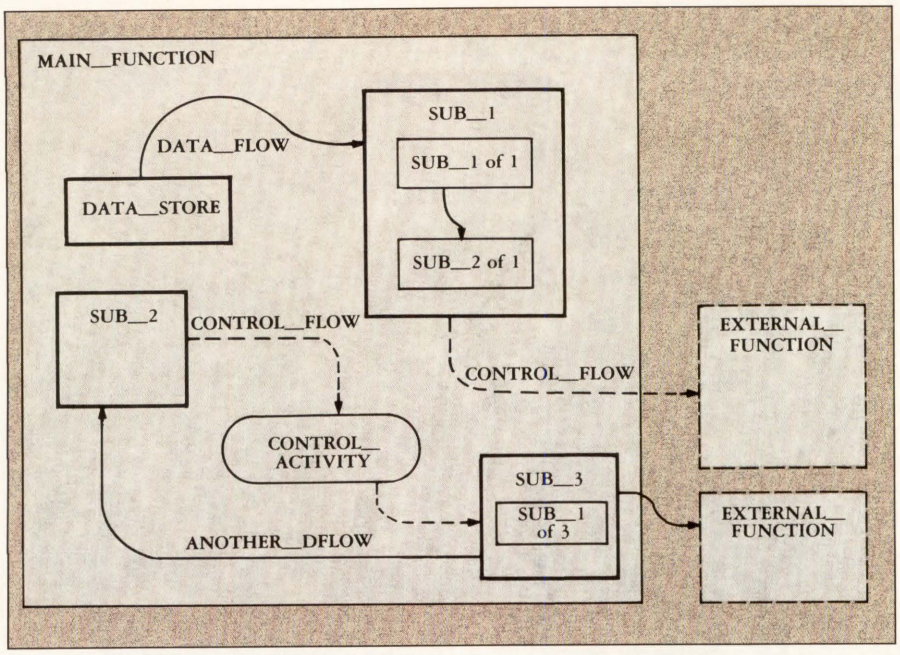

Figure 2. A sample activity-chart demonstrates the data and control flow. Each of various types of boxes and arrows has a specific meaning with respect to the functionality of the device being modeled.

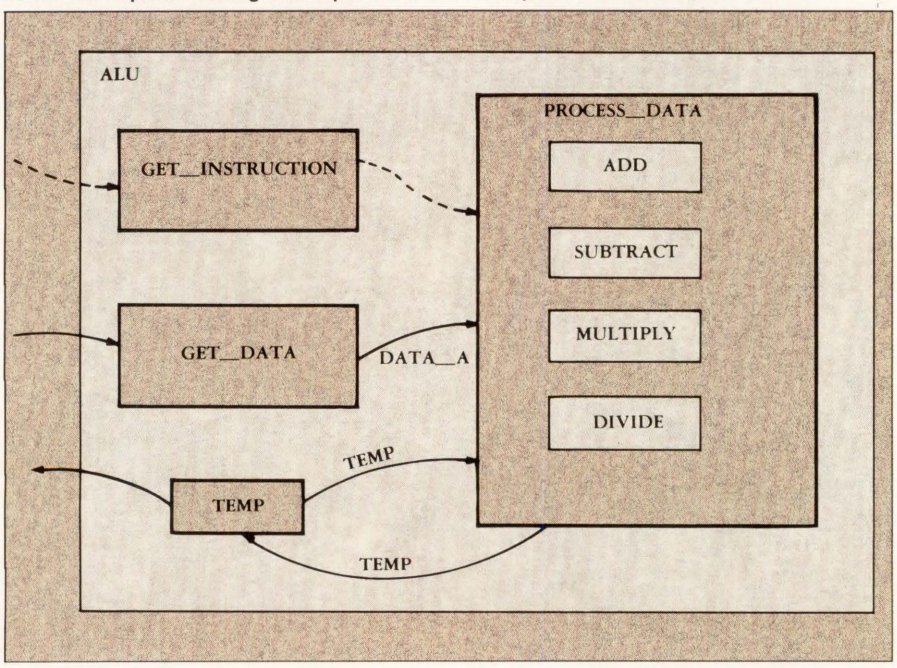

Figure 3. In this activity-chart for a simple ALU, the arrows show where data and control are generated and accepted.

Figure 3). The arrows show where data and control are generated and accepted. Since our model must interact with the rest of our system, it makes sense to model these additional parts as a set of external activities. The exchange of information and control between external and internal activities is also represented by such arrows.

■ Defining the Functions

The specific functions that this ALU performs are defined by an instruction signal (control) from the outside world (external activity). The data that is processed comes from the information flow.

With this activity-chart, we are able to define what functions are to take place in our ALU. These functions have information flow between themselves and the outside world. Control is provided as well and represents the instruction set for the processor. We now have a multilevel description of functions and a mapping of the information interface of this device.

This description is not enough to define a behavioral model. The activity-chart does not provide how, when, and why all this information and control are to be processed.

It is the statechart view that allows us to define the behavior of these activities. We then combine the two charts in such a way that the statechart controls the activity-chart. In this way, we can develop a complex model of a device at any level of detail. At each level we may define a set of necessary activities (processes) and the corresponding statecharts that define their behavior.

Statecharts are a derivative of finite state machines (Harel et al., 1988). They represent the modes, or states, that a system may be in. For example, a light can be in the on state or the off state; it cannot be in both states at the same time.

By including transition lines, we provide a mechanism to change from one state to the next. The transition lines may have a corresponding level describing the discrete event that causes the model to change its current configuration. These triggering events can also be combined with conditions or other events to form compound labels.

The definitions of these labels follow specific syntactic and semantic rules in order for the system to be well specificed behaviorally. For example:

make_true!(x)	;the value of x is set to TRUE
read(y)	;y is read from a data store
/stop (adder)	;the process "adder" is halted
GO	;the event go
if GO then ABC endif	;if the event go occurs, then generate the event ABC

With the use of statecharts we can define processes that are concurrent. This capability is useful when, for example, the device has two processes that are simultaneously active. This notion is implemented with the use of an "and line" to create two statecharts that communicate and therefore have some level of control over each other. To the CAE user, the use of communications between states, termed "broadcasting," would be most like a handshaking protocol between two communications devices.

Figure 4 gives an example of a statechart with concurrency. The transition from state D to state C occurs after the trigger "TM(EN(D),3)/GO_B." This label can be interpreted as, "a time-out occurs three clock units after the event of having entered the state D; at this time, also take the transition to C and broadcast the event 'GO_B." The effect here is the transition from state A to state B.

The use of broadcasting is very impor-

tant with respect to a complex model design. It provides a controlling mechanism between concurrent components.

■ SIMULATING THE CIRCUIT

The specification described visually with the Statemate design system is rigorous and formal. It "understands" the visual descriptions perfectly, to the point of being able to analyze them for crucial dynamic properties, carry out rigorous animated executions and simulations of the circuit or system described, and create running code automatically. In the world of CASE, these features are invaluable when it comes to the quality and reliability of the final outcome of a complex system. For the world of electronic design, one ramification of these features is that the Statemate description is ideally suited for use as a behavioral model.

With activity-charts and statecharts, designers have the necessary tools to model the functions and behavior of a complex VLSI device. This model can be executed and dynamically analyzed. To verify that the model represents the device as intended, a test suite would be created. The test suite would consist of stimuli in the form of data and events from the external activities represented in the model. If specified correctly, the model should process these environmental interrupts and data changes and produce predictable results. If not, the user can edit the charts and labels to correct the problems and then rerun the simulation.

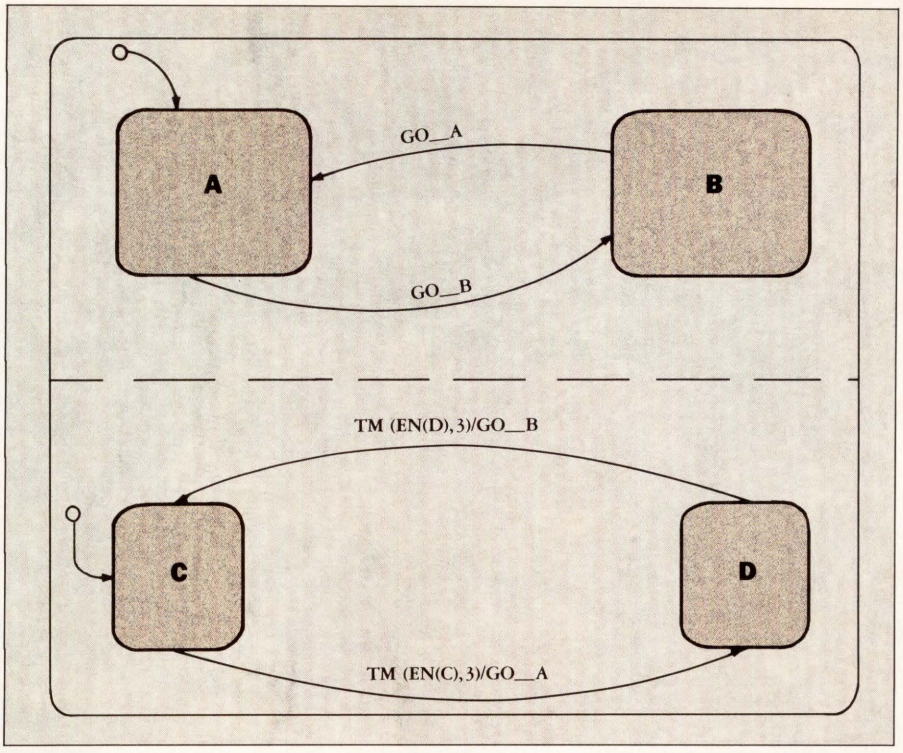

Figure 4. This statechart shows an example of concurrency. It uses "broadcasting" to provide a controlling mechanism between concurrent components.

Up-Front CASE

Though less than 1% of the dollars invested in computer-aided software engineering (CASE) in 1987 was for front-end tools, this is expected to change rapidly as system engineers and their managers realize that the cost of a design error grows almost exponentially as it moves down the design cycle from abstract idea to finished product. Front-end tools allow testing of the design at the early stages of design.

R.C. Wittenberg

Simulation with the Statemate environment can be performed in three ways.

The first is interactively, allowing the user to focus on specific areas of interest. Changing data, causing events to occur, scheduling actions, and so on, all can be done from the keyboard or through menu selection.

The second means of simulation makes use of a batch file. Interactively entering a large amount of data and control becomes tedious and reduces the amount of time available to analyze the model. The system allows the user to write a simulation control program (SCP) to drive the simulation. Through the SCP, the user is also able to make use of more sophisticated actions like defining breakpoints, reading from and writing to files, and generating random numbers.

The third way to simulate is with the design system's newest module (not yet released), which translates the Statemate model into prototype Ada code for execution in the target environment. In principle, the next logical step would be to link that code with the logic simulator to provide a complete simulation of the system under development.

Currently, behavioral models run as external programs linked to the logic simulator. When a predetermined breakpoint is reached, the logic simulator will halt, passing data and control to the behavioral model, which processes them and then passes new data back to the logic simulator, allowing it to continue on with the simulation. By creating an environment that permits the logic simulation portion to generate the external control and data for the Statemate model, a user could simulate a complete system in much the same way.

We have seen that with the new design environment, a behavioral model can be created graphically, as opposed to the creation of traditional behavioral models, which require an extensive programming background. Moreover, the system can be used to generate a model at any level of complexity necessary. In addition, the graphical environment reduces the time required to create a behavioral model. The use of Statemate would also reduce the risk of inacuracies entailed with present approaches. ■

REFERENCES

HAREL, D., ET AL. 1988. "STATEMATE: A Working Environment for the Development of Complex Reactive Systems," *Proceedings of the 10th IEEE International Conference on Software Engineering.*

HAREL, D. . "Statecharts: A Visual Formalism for Complex Systems," *Science of Computer Programming,* no. 8 (North-Holland, Elsevier Science Publishers).

ABOUT THE AUTHOR

Robert A. Tierney *is an applications engineer with i-Logix. He has had over seven years' experience in VLSI logic simulation. Before joining i-Logix, he was an applications engineer for logic simulation accelerators at Zycad Corp. He also studied electrical engineering for four years at the University of Southern California.*

Easy ASIC

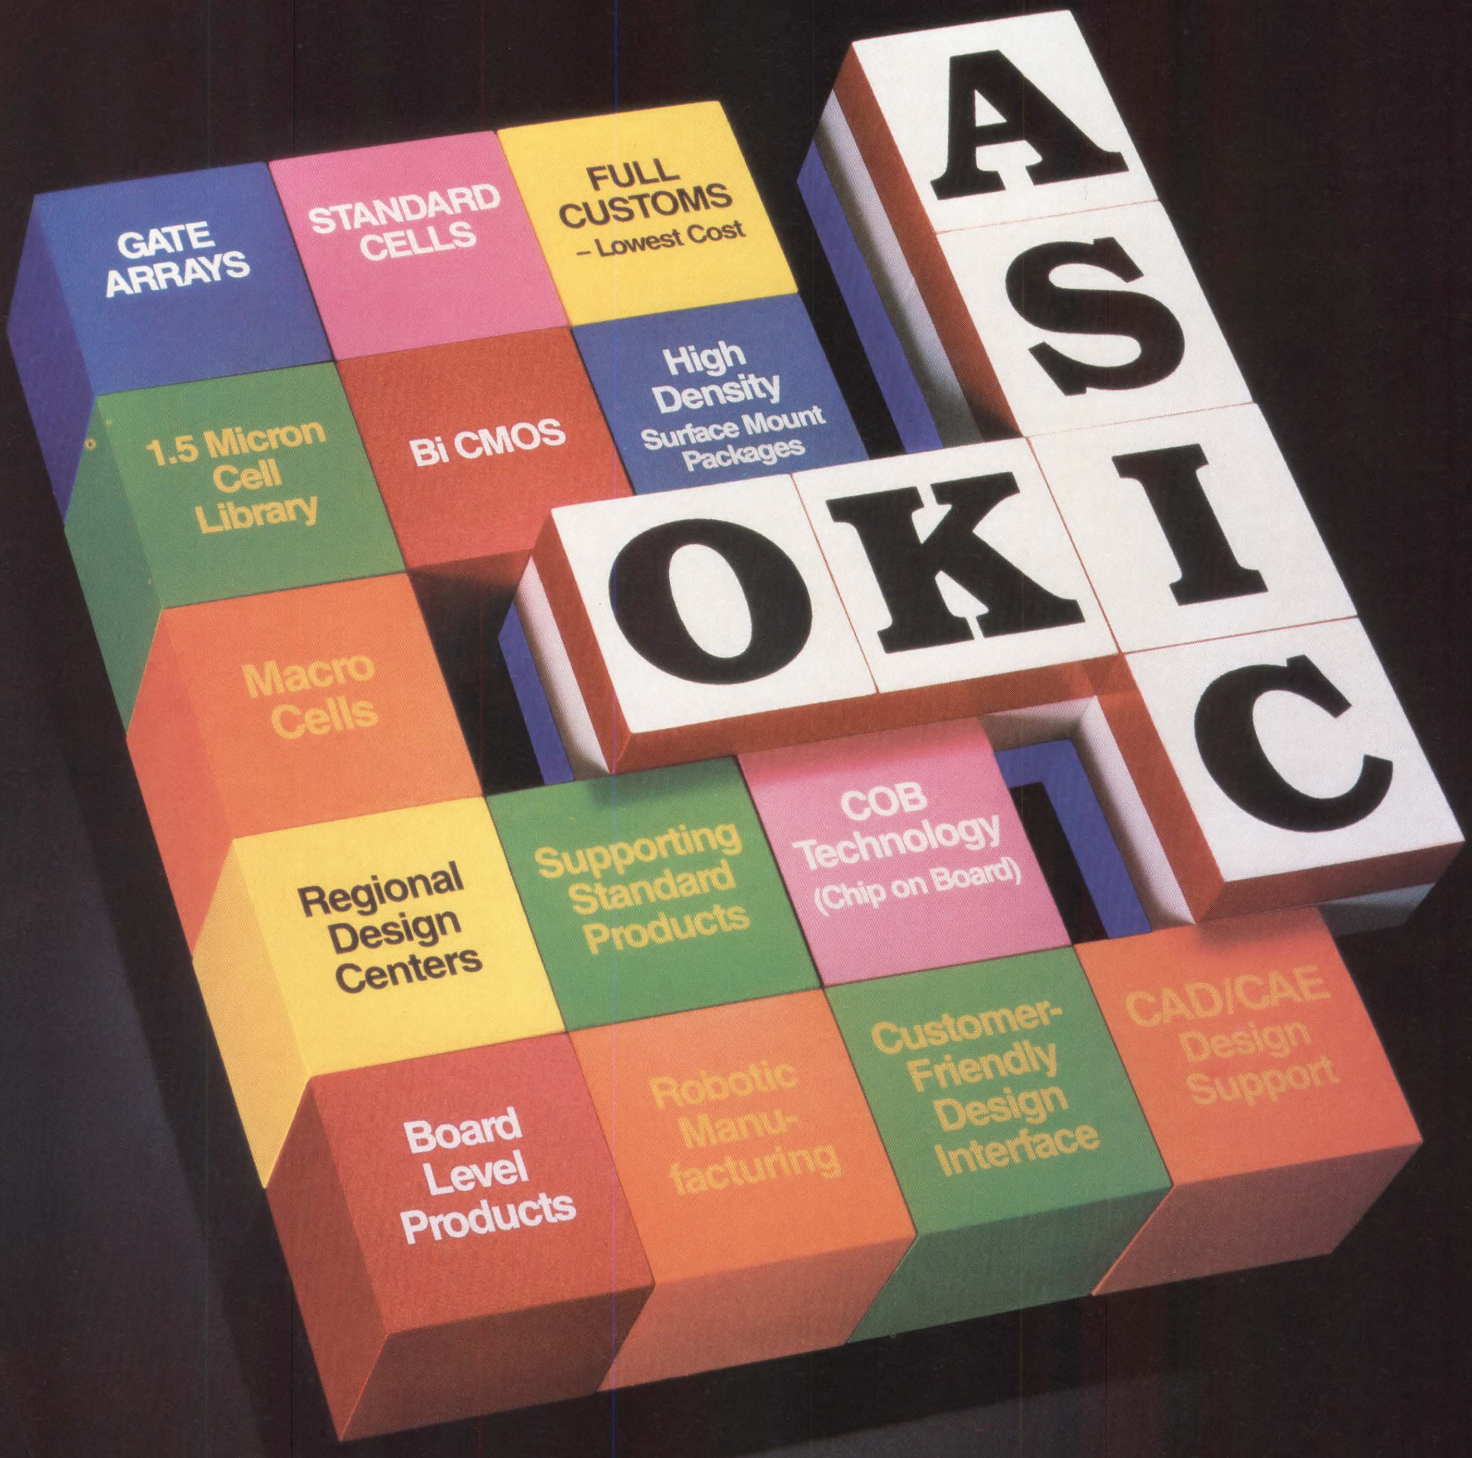
GATE ARRAYS
STANDARD CELLS
FULL CUSTOMS – Lowest Cost
1.5 Micron Cell Library
Bi CMOS
High Density Surface Mount Packages
Macro Cells
OKI
ASIC
Regional Design Centers
Supporting Standard Products
COB Technology (Chip on Board)
Board Level Products
Robotic Manu-facturing
Customer-Friendly Design Interface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

Only from OKI:

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

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

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

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

OKI: the totally logical choice.

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

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

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

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

Check out OKI ASIC data:

VLSI 1588

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

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

( ) Please call: we have immediate requirements:

Name/Title _____

Company _____

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

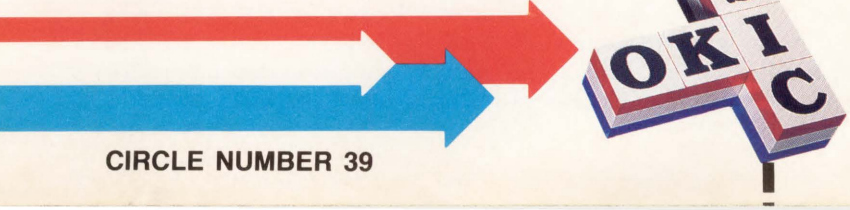

CIRCLE NUMBER 39

Convex's C Series

Top-Down Simulation of a Minisupercomputer

HAROLD DOZIER, CONVEX COMPUTER CORP. RICHARDSON, TEXAS

THE DESIGN CYCLE of the recently announced Convex C series of supercomputers was significantly shortened by extensive use of top-down simulation. The entire CPU and memory system were simulated and extensive diagnostic tests were executed on the simulation before building any hardware. The new processor was designed using a mixture of discrete 100K ECL logic, ECL static RAMs, dynamic RAMs, 3,000-to-10,000-gate ECL gate arrays, and 20,000-gate CMOS gate arrays. A total of 14 unique gate array types were designed. The simulation of the new gate array designs was virtually a fallout of the overall simulation strategy.

This new CPU design is the basis for five new Convex machines: the C210, with one CPU; the C220, C230, and C240, with two, three, and four CPUs, respectively, tightly coupled in a multiprocessor environment; and the C130, a lower-cost, lower-performance uniprocessor system. For this article, we will call the newly designed CPU the C series design.

The top-down design approach represented a significant departure from the approach taken for the first Convex machine, the C1. In that design, only the ASIC devices received any significant logic simulation, with verification of the rest of the system being done by first constructing a wrapped-wire version. The approach taken appeared to be the only practical approach available to us in 1983. Our computer resources were limited to a single VAX-11/780, which was only available to the hardware design team 12 hours per day (the software team had it the other 12 hours).

There was no true behavioral-level simulator available. One could write extended functions in a high-level language and link them to a gate-level simulator, but that wasn't the simulation environment necessary for the top-down design of a supercomputer. Everyone recognized the necessity of thoroughly simulating the ASICs, but we intentionally limited the use of semicustom chips to discrete arithmetic functions, which contained virtually no control logic. That being the case, we were able to simulate each ASIC as a simple stand-alone function.

The ASICs used in the C1 design were 8,000-gate CMOS arrays, and we were capable of performing about 100 simulation vectors per hour with the SCALD simulator running on the VAX. At that rate, it would have been impossible to simulate any significant portion of the system, given the speed of the simulator and the VAX time available. In contrast, using the Endot N.2 simulator at the behavioral level running on the C1 provided us with a simulation capability that was about 1,000 times faster than what we had for the C1 design, thus making a thorough top-down simulation strategy possible.

■ WHY TOP-DOWN

The top-down approach was motivated by several factors. We felt that it was a better way to design a complex computer system and that it could significantly shorten the design cycle. The latter assumption proved to be correct. Whereas the time from the start of the project to producing the first hardware was about the same as for the C1, the many months of debugging required to take the C1 design from a wrapped-wire prototype to a hard-tooled production machine were reduced to a few weeks in the C series design. Also, wire wrapping would not have been suitable for the C series machine, which is primarily an ECL design. A by-product of the top-down approach was that it allowed each of the system-level designers and ASIC designers to work more independently and simulataneously—saving additional time.

■ THE IMPACT OF SYSTEM-LEVEL SIMULATION ON ASIC DEVICES

The C series design called for a significant amount of control logic to be implemented in ASICs, and it was important to simulate these devices in their system environment. We now feel that ASICs can be designed that will produce the expected functionality and performance on the first pass. This success is virtually ensured by

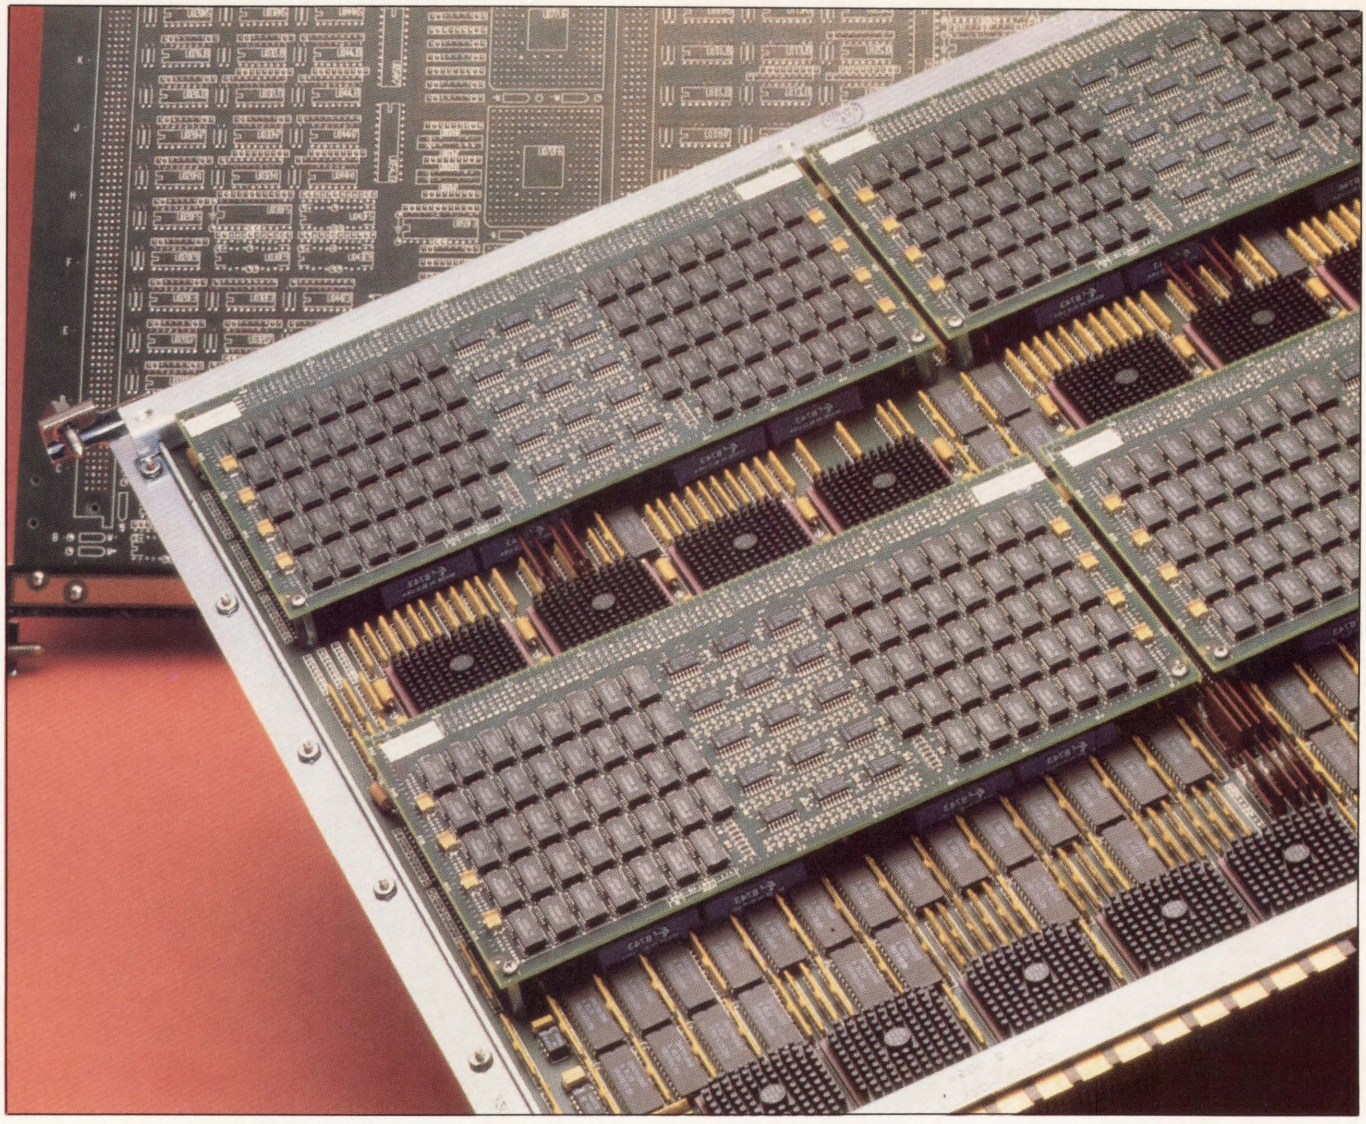

the ASIC vendors using extensive up-front simulation, capacitance extraction, back annotation, and resimulation—even if the designers don't go through this procedure before releasing their net lists.

Our biggest concern in designing the ASICs was that the *designed* functionality was not the functionality actually required for proper operation in the system as a whole. Making sure that the designed functionality is, in fact, the *required* functionality, is best accomplished by simulating the entire system, including the ASICs. Further, performing this simulation at the behavioral level allows one to determine the required functionality of the ASICs and all the other system components prior to gate-level design. In addition, the simulation runs significantly faster, which in turn allows more extensive simulation per fixed CPU resource.

■ The Tools Selected

The N.2 simulator was selected, primarily because of its strong behavioral-level capabilities, since behavioral was the bulk of the simulation task. Although the N.2 simulator doesn't employ built-in gate-level primitives, gate-level simulation was also done with it by using small behavioral models to handle the basic gate-level logic functions. All simulations were done with zero delays, and a separate static path analysis timing verifier was used to verify proper timing of the entire system.

We used Valid Logic Systems' workstations for schematic capture. One of the primary benefits of the Valid system is that the database for a design is available to the user in several different forms, and each form is documented and in ASCII text of one variety or another. This variety and the accompanying documentation make it rather easy to write utility programs to translate Valid data into another format or, for that matter, to translate other forms of design data into the appropriate Valid format.

To aid in gate-level simulation, we wrote a program to translate the Valid database into a format acceptable to the N.2 simulator. The .2 simulator does not, however, perform fault grading, and it was felt that fault grading of the final test patterns for the ASICs was an extremely important task. The C series design required 14 new gate array designs, plus one previous array design. However, most of the arrays are used several times throughout the design, adding up to more than 100 ASICs total in a single-processor system. Consequently, avoiding any possibility of "test escapes" is extremely important. Fault grading was performed with the Silos simulator from SimuCad running on the C1.

The tools and programs we used are summarized in the table.

■ Behavioral-Level Simulation

The logic simulation was constructed hierarchically and went through considerable evolution during the course of the project. The initial behavioral models that were written modeled the operation of the various parts of the machine. The partition-

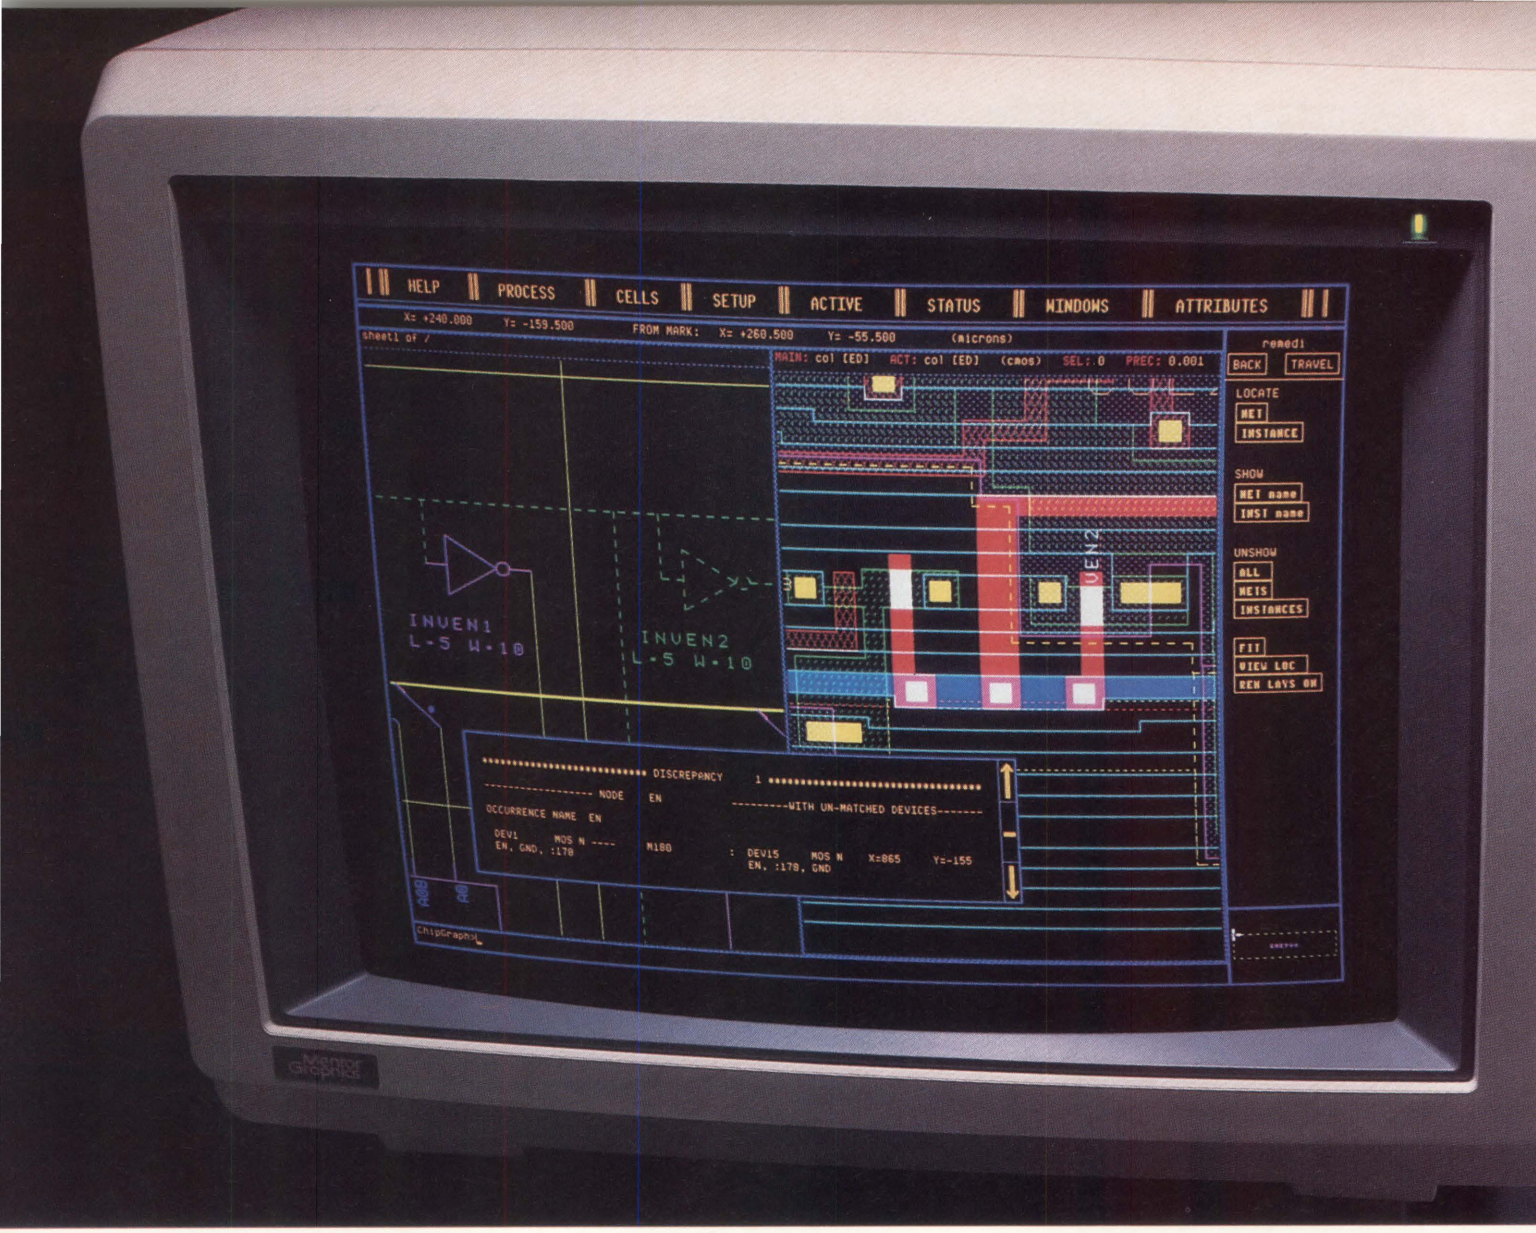

Only Mentor Graphics maps symbols to silicon.

The bigger the IC design, the bigger the problem: you're trying to locate a discrepancy between the schematic and your IC layout, but all you have is an ASCII error report. It's like driving all over a strange city to find an address — without a map.

That's why Mentor Graphics created REMEDI,™ a graphical interface that helps debug complex layouts. REMEDI works with Dracula II,™ taking the leading layout verification package's layout-versus-schematic checks a step further. LVS errors detected with Dracula II can be quickly pinpointed on both the layout and schematic using REMEDI's interactive graphical correlation capabilities.

And, because REMEDI is part of ChipGraph,™ the powerful Mentor Graphics full custom IC layout editor, as soon as you find an error you can fix it. There's no need to move back and forth between the layout editor and the debugging tool to correct the design database.

Today's complex VLSI designs demand a layout tool that lets you create and navigate efficiently through a maze of mask data. So ChipGraph provides flexible geometry editing and fast cell-based layout tools. Beyond this,

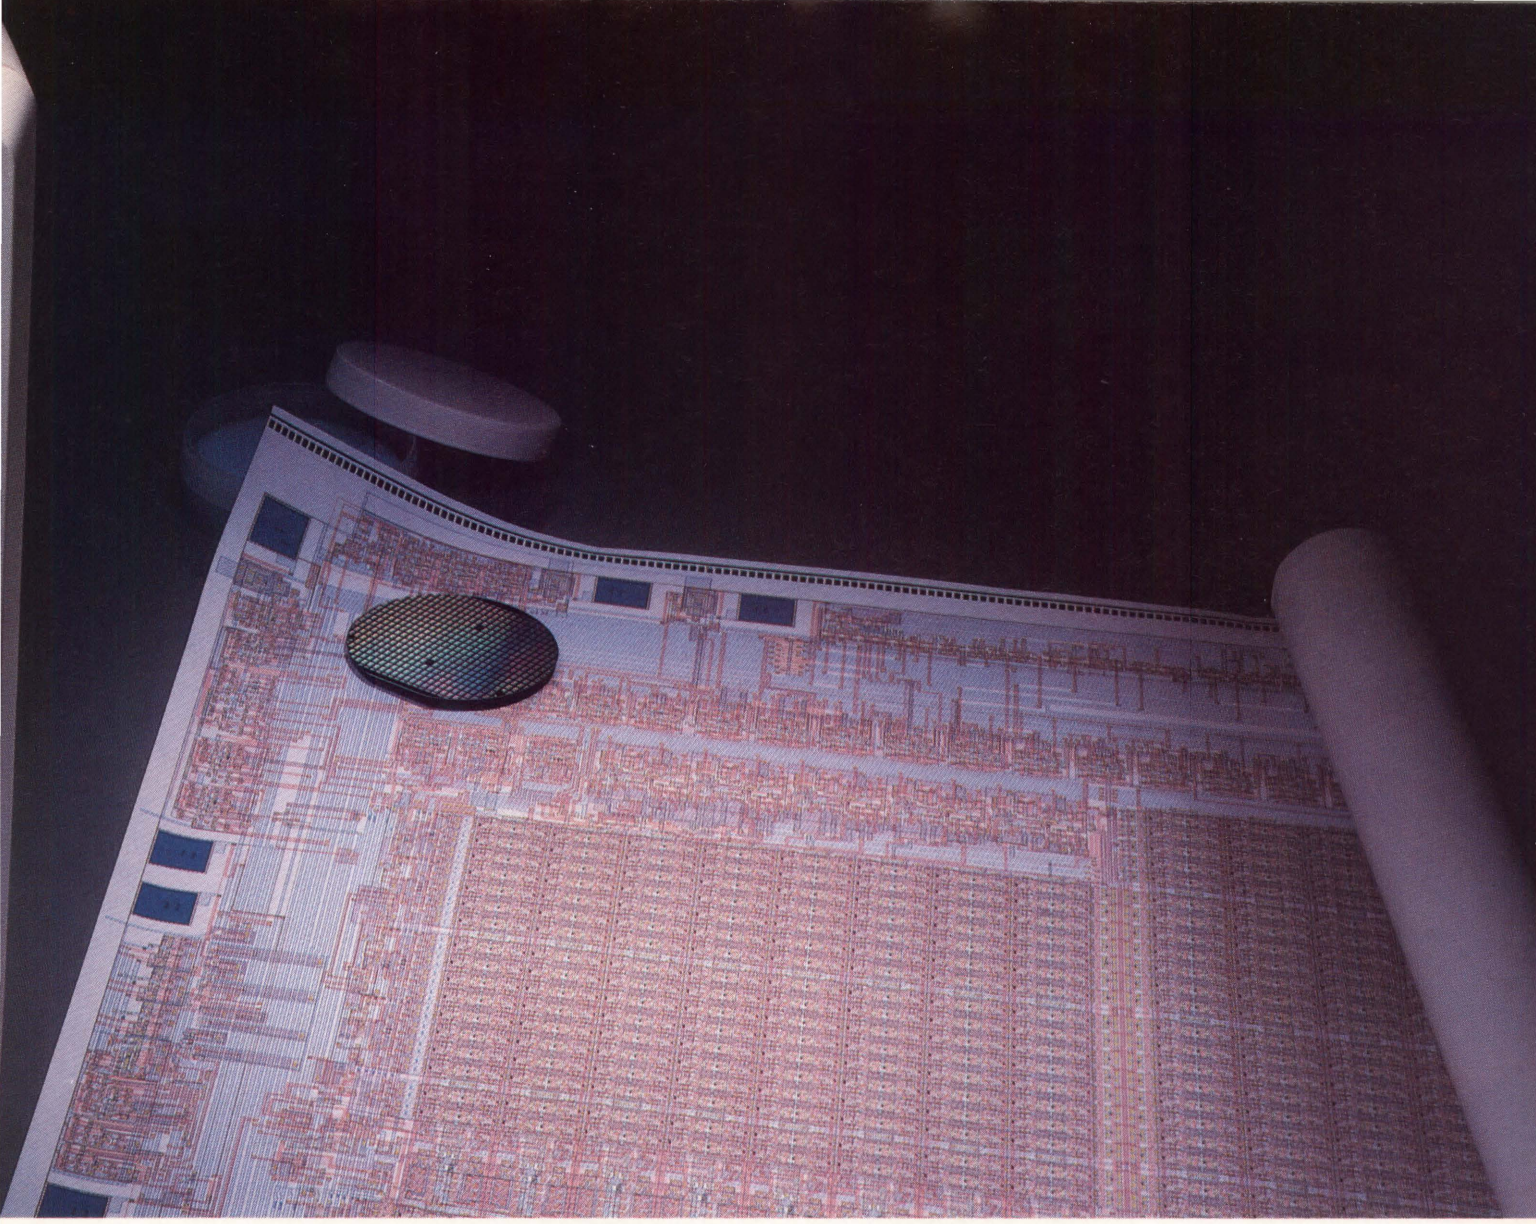

Structured Chip Design (SCD), a hierarchical approach to physical layout, removes much of the unnecessary and confusing mask data, while retaining correct functionality and interfaces.

With SCD, you can work with a simplified representation of the cell when making higher level decisions. And you can easily move between SCD and more familiar tools, with no loss of productivity. The result? The tightest possible layout, created quickly and accurately.

ChipGraph also lets you partition a large design over any number of workstations while maintaining version control, through a network-wide shared database.

And there's no need to worry about losing old data when switching to a new tool. ChipGraph offers full data compatibility with your existing design files.

Best of all, ChipGraph is not just an idea — it's a working reality. Designers are already using ChipGraph to lay out 32-bit microprocessors and multi-megabit RAMS.

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

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

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

Sydney, Australia; Phone 02-959-5488 **Mississauga, Ontario;** Phone 416-279-9060 **Napean, Ontario;** Phone 613-828-7527 **Paris, France;** Phone 0145-60-5151 **Munich, West Germany;** Phone 089-57096-0 **Wiesbaden, West Germany;** Phone 06121-371021 **Hong Kong;** Phone 0566-5113 **Givatayim 53583, Israel;** Phone 03-777-719 **Milan, Italy;** Phone 02-824-4161 **Asia-Pacific Headquarters, Tokyo, Japan;** Phone 03-505-4800 **Tokyo, Japan;** Phone 03-589-2820 **Osaka, Japan;** Phone 06-308-3731 **Seoul, Korea;** Phone 02-548-6333 **Spanga, Sweden;** Phone 08-750-5540 **Zurich, Switzerland;** Phone 01-302-64-00 **Taipei, Taiwan;** Phone 02-7762032 or 02-7762033 **Halfweg, Netherlands;** Phone 02907-7115 **Singapore;** Phone 0779-1111 **Bracknell, England;** Phone 0344-482828 **Livingston, Scotland;** Phone 0506-41222 **Middle East, Far East, Asia, South America;** Phone 503-626-7000

Mentor Graphics®

Dracula II is a trademark of ECAD, Inc.

PRIMARY TOOLS IN THE CONVEX DESIGN AUTOMATION SYSTEM		
TOOL OR PROGRAM	SUPPLIER	FUNCTION
N.2	ENDOT	BEHAVIORAL, GATE-LEVEL, AND MIXED BEHAVIORAL AND GATE-LEVEL LOGIC SIMULATION
SILOS	SIMUCAD	FAULT GRADING OF ASIC TEST PROGRAMS
GED	VALID LOGIC SYSTEMS	SCHEMATIC CAPTURE, MANUAL ASIC AND PCB PLACEMENT
VALIDTIME	VALID LOGIC SYSTEMS	TIMING ANALYSIS
MDT PACKAGE	PCK TECHNOLOGIES	MULTIWIRE PRINTED WIRING BOARD ROUTING (KERNEL VECTORIZED ROUTING ALGORITHM WRITTEN BY CONVEX)
ALLEGRO	VALID LOGIC SYSTEMS	ETCHED PCB ROUTING
UTILITIES	INTERNALLY WRITTEN	NUMEROUS, FOR DATABASE TRANSLATION AND DISPLAY

ing of the system into separate models was along the expected physical boundaries.

The initial simulations were of a more architectural nature, to prove that all the necessary building blocks existed, to identify any possible bottlenecks in the system, and to prove that there were no bus conflicts on the major system buses. In some cases, functions that were to be implemented in gate arrays were identified early on, and the initial set of behavioral models followed this physical partitioning of the implementation.

The initial system simulation is represented pictorially in Figure 1. The topology files shown in this and the subsequent figure are basically a net list that describes the "hookup" of the various behavioral models.

Obviously, some consideration of how the logic would ultimately be designed at the gate level went into the initial set of behavioral models. But having completed the early behavioral simulation, the designers then concentrated their efforts on the implementation of the logic, and throughout this effort, there were significant changes. For example, certain pieces of logic that were initially assumed to fit with other pieces on a given printed circuit board were shifted around as better estimates of package counts were made. As we considered early timing issues, it became obvious that certain events would have to occur at the "half-clock" edge, as opposed to the major cycle boundary, and that certain other pieces of logic may require a clock that led or lagged the system clock.

Additionally, many of the initial behavioral models represented a combination of discrete logic and gate arrays. Though we may have assumed that gate arrays would be used to implement part of this logic, frequently we didn't know exactly how many different gate array designs or how many copies of a single design would be used for the final implementation.

Through the course of doing cell count estimates, taking a harder look at gate array pin counts, and making a first-pass analysis of some critical paths, and so on, the definition of gate arrays began take shape. As a result, the number of behavioral modules in the system simulation was increased, and the simulation evolved into a slightly different form. As an example, Figure 2 shows a more finalized form of the vector processor portion of the simulation.

■ PROVING FUNCTIONALITY

The simulation became the primary vehicle for proving the correct functionality of the entire design. Diagnostic engineers developed a complete set of tests, predominantly scan-based, to verify the entire system. They also wrote a diagnostic scan interface to the simulator which emulated the same interface that would be available on the console of the actual machine. Not only did these diagnostics serve to verify the proper functionality of the simulated machine, but also the same diagnostics were used on the actual hardware as well. Whereas the diagnostics served as the primary verification tests for the simulated system, the simulation served as a debugging vehicle for the diagnostics, resulting in a set of proven diagnostic tests available before the hardware was ready.

As various parts of the logic were proven to operate correctly on various tests, the simulated cycle-by-cycle results were saved, along with the tests themselves and their proper method of execution. These then formed the regression tests to be re-executed each time the logic was changed. For example, one set of diagnostics may have simulated correctly, but later another test may have failed and some change to the behavior may have been required to make that test pass. Having made such a change, all of the previous tests would be run again to prove that the most recent change had no effect on operations that were already working properly.

■ BEHAVIORAL MODELS AID IN ASIC DEVICE SPECIFICATION

A by-product of having each ASIC specified with a separate behavioral model in the system simulation is that the behavioral models can become good functional specifications for the ASICs. They do not necessarily replace the need for a written functional specification of the devices, since reading and interpreting a behavioral model requires a good familiarity with the simulation language. Certain other people associated with the project, like the board-level test engineers, quality assurance engineers, and manufacturing technicians, all may need to gain at least a working familiarity with the function of an ASIC, and this familiarity is probably better accomplished with traditional block diagrams and textual description of the operation of the device.

However, given the complexity or the functions performed by an ASIC, it can be virtually impossible to write a functional description that, in testing, describes exactly what the device will do under every possible circumstance, input data pattern, and so on. However, should any question arise about what a device will do in a particular situation, the system designer can look at the behavioral model to determine the device's behavior and make fixes where necessary.

Similarly, the ASIC designer can look at the behavioral model of the surrounding system to gain insight into what the system requires of the device. Obviously gate-level logic schematics provide a similar vehicle, but it is desirable that the system designer and the ASIC designer be able to proceed with their gate-level designs fairly independently of one another, without the design work of one person holding up another. Of course, the same person may be doing the design of an ASIC and the surrounding external logic (or other ASICs). Still, behavioral models can serve the same purpose, and the designer

Where Are You Going in the Nineties?

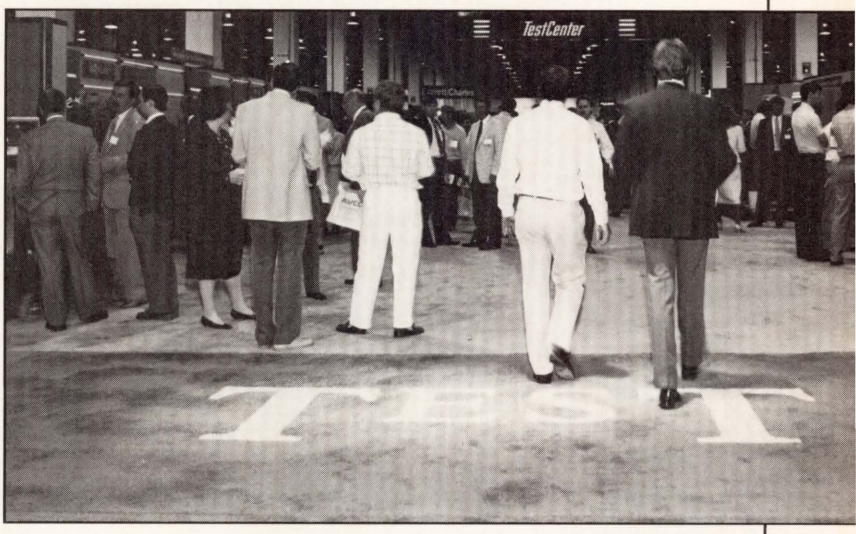

Develop your test strategies for today and tomorrow.

It's time to start planning now for a new decade in electronic product development. Whether you're in design, manufacturing or service, you constantly confront new problems in quality, testability, manufacturability, serviceability and more. In every area, test technology plays an important part in the solution.

Focusing on integrated test strategies at the chip, board and systems level, the ATE & Instrumentation Conference provides you with solutions to your problems. What you see and learn will be invaluable in your product plans *and* professional growth.

Plan now to attend. Join industry leaders for seminars on the newest directions in Design Test, Manufacturing Test, and Field Test. Join leading test companies for hands-on demonstrations of their products and equipment. Join your colleagues in engineering and management at the only conference dedicated to test solutions.

Join us at the ATE & Instrumentation Conference East on June 7–9, 1988 at the World Trade Center, Boston, MA.

Sponsored by *Electronics Test, Circuits Manufacturing* and *EOS/ESD Technology* magazines.

ATE & Instrumentation Conference East

For Test & Instrumentation Users

For conference details and registration forms, contact Al Dean at MG Expositions Group, 1050 Commonwealth Avenue, Boston, MA 02215, telephone (617) 232-3976, (800) 223-7126.

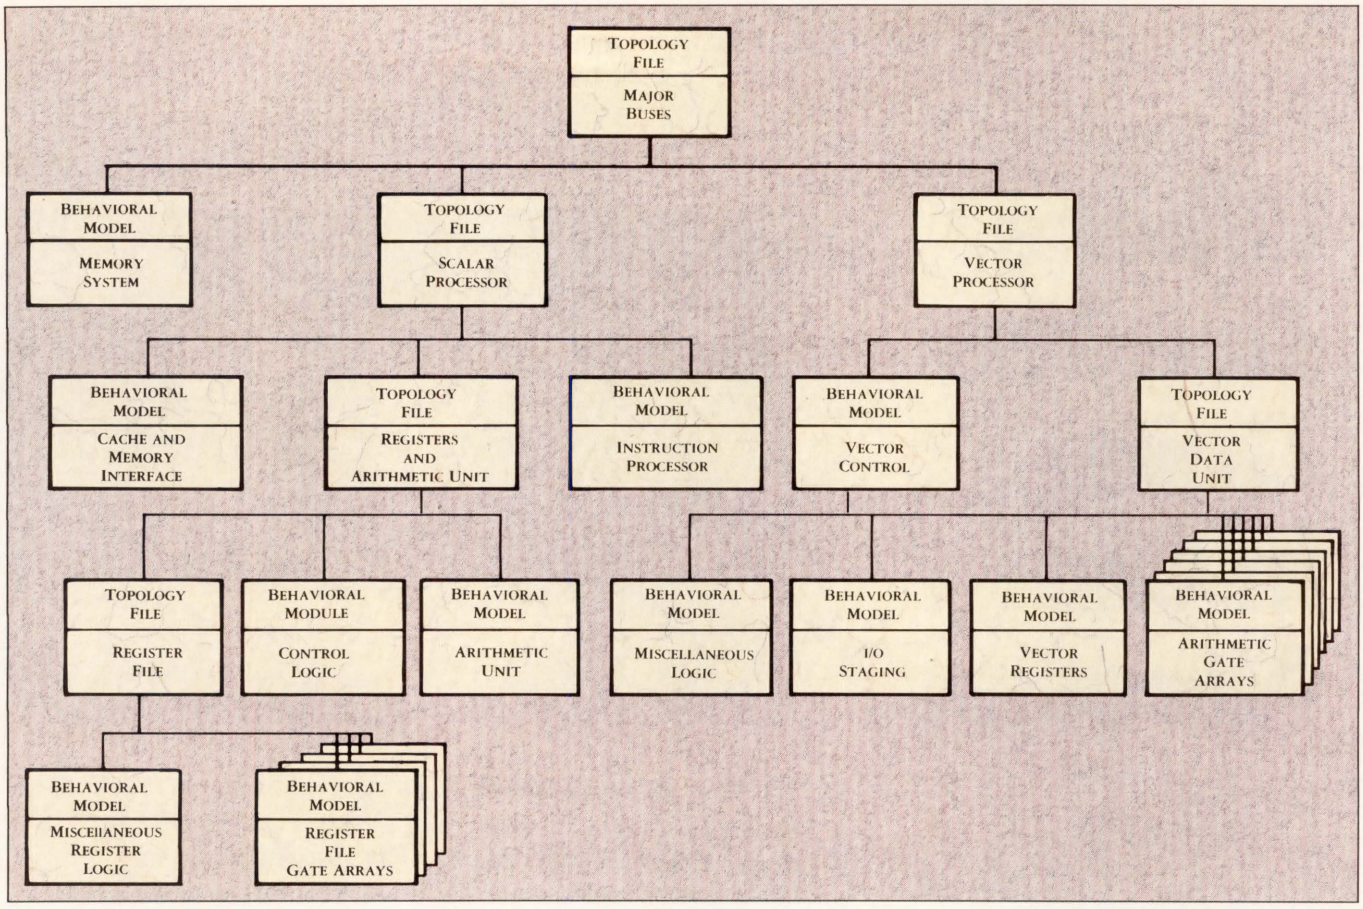

Figure 1. In this pictorial representation of the initial system simulation, the topology files are basically a net list that describes the "hookup" of the various behavioral models.

of any piece of logic can refer to the behavioral model for that logic in much the same way that a flow chart can be used in developing a piece of software.

■ Gate-Level Logic Design and Verification

The task of gate-level logic design became one of creating logic that provided the same behavior as did the higher-level models. The gate-level design of any part of the logic could easily be done by an engineer, independent of the work being done by others, using the behavioral model(s) as the blueprint. Obviously there were still some trade-offs to be made for the sake of ac performance or for logic simplification or parts count reduction, and when such changes were made, the behavioral models would also be correspondingly changed.

In certain cases, a simulation at a level lower than the entire system simulation proved valuable. As a simple example, testing an ADD instruction from the system level would require loading registers; performing a register-to-register addition; and to make the result visible outside the bounds of the CPU, writing the register containing the result back to main memory. Of the several machine cycles involved, the one cycle in which the addition actually takes place is the only one of interest to the person designing the adder. Therefore creating a separate simulation driver that focused on the operations of interest proved to be an efficient method for detailed verification of certain pieces of the logic.

Verification of the gate-level logic required only that the designer prove that the gate-level logic produced the exact same behavior as the behavioral models. This proof was generally accomplished in one of two ways. One method was to replace the behavorial model of the logic with the gate-level expansion in the full system simulation or in a simulation made using a separate simulation driver as described above. Within the N.2 simulator, the interface with the behavioral models is identical to the interface for a gate-level expansion of that same logic.

It is therefore very easy to form a simulation that replaces any behavioral model or set of behavioral models with the gate-level equivalents. Regression tests could then be run to prove that the operation of the system or subsystem was identical to what had previously been simulated using the behavioral model.

Another method was to run a simulation using the behavioral model and trace all inputs and outputs of that model throughout the course of the simulation. Then the same inputs could be applied to the gate-level logic and the outputs could be compared with those previously traced to ensure proper gate-level functionality.

The second method required more overhead than the first to extract data and reformat it for application during the gate-level simulation; however, it allowed the gate-level simulation to proceed much more rapidly. The trade-off between the two generally revolved around the complexity of the initial simulation—was the initial simulation a full system or was it a subsystem? Also involved in making this trade-off was the number of machine cycles of gate-level simulation that one wanted to perform. If the gates in question could be simulated in a subsystem, or if relatively few cycles were involved, it was easier and more efficient to just replace the correct behavioral model with gates.

■ Verification of ASIC Devices

As previously indicated, behavioral models should be partitioned along physical boundaries. This partitioning is especially important when it comes to the

PUTTING YOU AT THE CUTTING EDGE OF THE COMPUTER AGE

Yours free for 15 days.

Take full advantage of expert systems!

PUTTING EXPERT SYSTEMS INTO PRACTICE

By Robert G. Bowerman and David E. Glover

$39.95, 384 pages, illus.

Get the strategies and tactics you need for putting expert systems into practice in the most practical, cost-effective, and profit-making way! This real-world guide clearly covers the seven phases of a recommended expert system development cycle, demonstrating how a project develops from conception to successful completion.

To order, circle # 132 on reader service card.

AI—it's time to start using it!

ARTIFICIAL INTELLIGENCE
An Applications-Oriented Approach

By Daniel Schutzer

$39.95, 224 pages, 60 illus.

Artificial intelligence (AI) takes a quantum leap forward into the real world with this first guide to current applications of expert systems, machine perception, and natural language processing. The author—an AI system expert with Citibank—shows how to apply AI to solve a wide range of problems in business, industry, science, medicine, and defense.

To order, circle # 134 on reader service card.

The practical programming guide you've been waiting for

LISP
The Language of Artificial Intelligence

By A.A. Berk

$28.95, 168 pages, illus.

Professionals familiar with traditional programming languages will welcome this highly practical guide to LISP, the up-and-coming language of the new artificial intelligence.

To order, circle # 135 on reader service card.

Detailed coverage—from microchips to full-scale computer systems

HANDBOOK OF ADVANCED SEMICONDUCTOR TECHNOLOGY AND COMPUTER SYSTEMS

Edited by Guy Rabbat

$64.95, 960 pages

For the first time, one source details computer-system development by starting at the microchip level. It covers everything from semiconductor processing methods to very large scale integration (VLSI), computer-aided design of computers, and computer systems architecture. This source provides information on VLSI routing techniques, cell-based VLSI design, and VLSI design verification and logic simulation. Software design techniques complete this chip-to-final-system reference.

To order, circle # 126 on reader service card.

New !
A self-teaching guide...

MODULA 2 FOR MICROCOMPUTER SYSTEMS

By J.E. Cooling

$23.95, 208 pages

Here is a complete guide to Modula 2—one of the latest high level computer languages designed specifically to enhance software reliability, quality, portability, maintainability, and cost control.

To order, circle # 136 on reader service card.

A convenient source for fast, accurate information

ENCYCLOPEDIA OF COMPUTER SCIENCE AND ENGINEERING

Edited by Anthony Ralston

$93.95, 1678 pages, illus.

Drawing on the expertise of more than 300 respected authorities, this one-volume reference answers a wide range of questions about computer science, engineering, and related fields. Offers over 550 informative entries plus over 700 explanatory illustrations, tables, and charts. Useful appendices focus on abbreviations, acronyms, high-level languages, research journals, and university departments.

To order, circle # 133 on reader service card.

A truly comprehensive soldering handbook—

SOLDERING HANDBOOK FOR PRINTED CIRCUITS AND SURFACE MOUNTING

By Howard H. Manko

$52.95, 370 pages, illus.

All the information you need on soldering printed circuits is in this special handbook. It thoroughly covers all aspects of printed circuits—from design, material selection, processes, equipment, inspection, and quality, touchup, and repair to line management and cost studies. You'll learn about the new technology of surface mounting on to printed circuits that is fast becoming standard in the field.

To order, circle # 121 on reader service card.

State-of-the-art computer security techniques

BUILDING A SECURE COMPUTER SYSTEM

By Morrie Gasser

$34.95, 288 pages, illus.

A must for the practicing computer professional who wants to understand and implement technical solutions to computer security problems. It is a major reference source for all security issues that are likely to come up in the day-to-day work of all computer professionals.

To order, circle #125 on reader service card.

TO ORDER FOR 15-DAY FREE EXAMINATION: Simply circle the appropriate number(s) on the Reader Service Card at the back of this magazine. Your book(s) will be sent to you for your 15-day examination. If you are satisfied, keep the book(s) and pay the purchase price plus postage and handling. Otherwise return the book(s) by the end of the 15-day period and owe nothing. You may also order by writing directly to *VLSI*.

To order, use the card provided or call toll-free **1-800-645-6278** and ask for Kathie Conlon.

B 4460

VNR VAN NOSTRAND REINHOLD

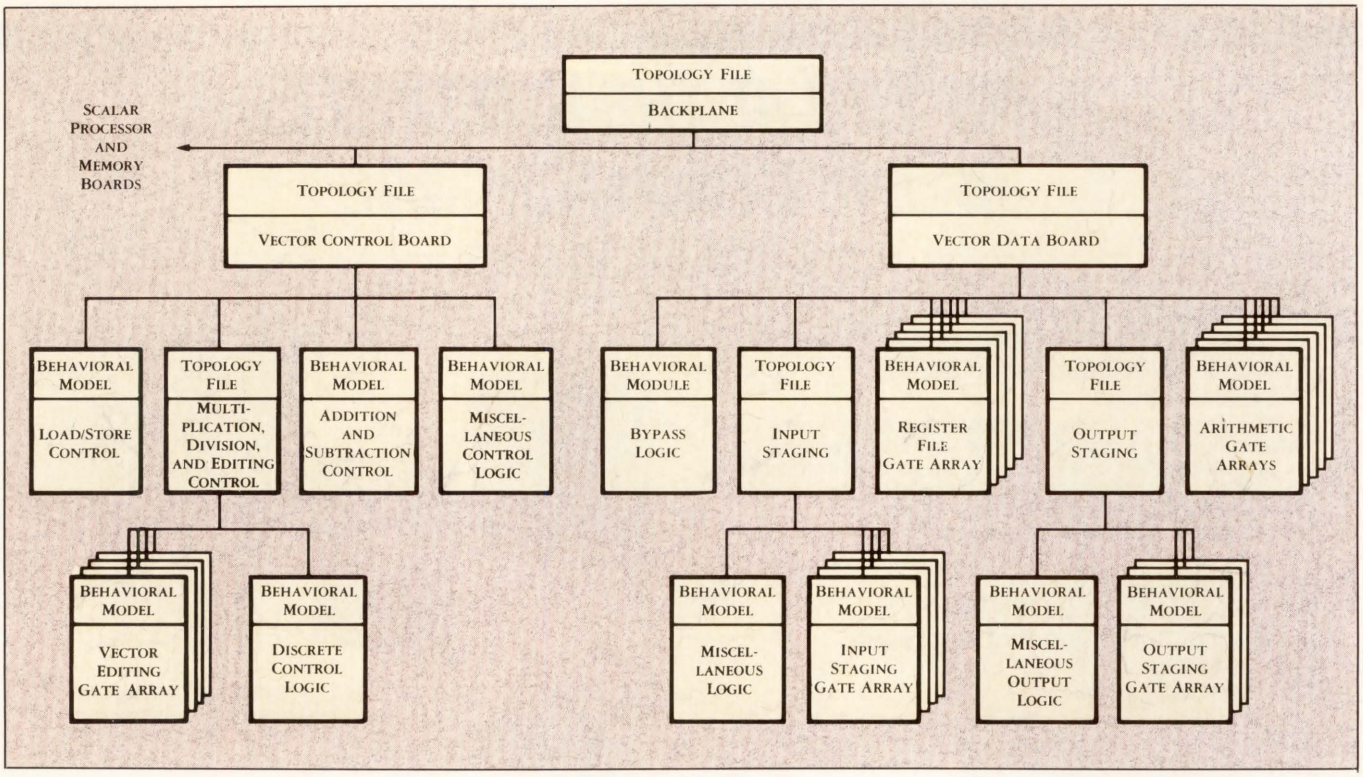

Figure 2. As the simulation progressed, the number of behavioral modules increased, and the simulation evolved into the form shown above, which represents a more finalized form of the vector processor portion of the simulation.

verification of the ASICs. If a single behavioral model is used to represent several ASICs or to represent a mixture of ASICs and discrete logic, the signals that represent the pins of the ASIC may not exist as such in the system simulation and instead may just be buried in the internal state of the behavioral model. These buried pins make tracing of the inputs and outputs of the device much more difficult and preclude the ability to simply replace a behavioral model in the system simulation with the gate-level logic of a single gate array.

Both gate-level verification methods described for verification of the board-level logic were also used for the ASICs, but most of the verification was done by tracing the inputs and outputs of the behavioral model and then independently simulating the gates.

One reason for this approach was that during a single behavioral-level simulation, the inputs and outputs of several ASICs could be traced at one time and then each device could be individually simulated using that data. A second reason was that the ASIC designers needed a method of independently simulating the ASIC using test vectors as the input and expected output so that they could verify the final set of tests that would be provided to the ASIC vendor.

As a result of this requirement, we used a set of utility simulations that employed traced data as inputs and expected outputs. One utility was used to read test vectors in the vendor's format and automatically build a simulation driver to apply those patterns to either a gate- or a behavioral-level model of the ASIC and to cause the outputs of the model to be traced. Fujitsu was the vendor chosen to supply both the ECL and the CMOS ASICs. A second utility was used to compare the traced outputs with the expected outputs from the Fujitsu-format test vectors and provide a formatted table of differences. With this utility in place, a third one was written to allow the trace file from a system-level simulation to be translated into the Fujitsu test vector format.

Although the test programs developed for many of the ASICs contained a number of vectors that were derived from system-level simulation, the main purpose of the tests were to prove that each device had been properly manufactured. Therefore fault coverage and verification of an ASIC's ac performance were the primary concerns in the set of patterns that were delivered to the ASIC vendor. Many of these patterns were generated by hand, and others were generated using intelligent random-number techniques rather than relying strictly on patterns derived from system-level simulation.

It appears that some users of ASICs don't make this distinction. Many times, ASIC designers concentrate on running simulations to prove that the device does produce the correct behavior that the system requires, then they use those same patterns as the manufacturing test of the device, even though the patterns may do a relatively poor job of detecting many of the possible manufacturing defects or verifying the true ac performance of the device. Conversely, some ASIC designers seem to concentrate almost entirely on fault coverage and do not perform a thorough enough job of verifying that the design exhibits the exact behavior required by the system. Obviously, proper system behavior tests and a high-quality manufacturing test are both very important, though neglecting to put enough emphasis on system behavior has the most severe impact. Such neglect can lead to a costly redesign of the device and may bring the rest of the project to a halt. This most important aspect of ASIC design came virtually as a fallout of all of the extensive system-level simulation done on the C series project.

It might be interesting to note that, at Convex, the ASIC designers perform all phases of the design, including the generation of the manufacturing tests. Fault grading of the manufacturing tests was performed with the Silos simulator. As we did for the N.2 simulator, we wrote utility programs to convert the Valid database into a Silos format and to convert Fujitsu format test vectors into the format required by the Silos simulator.

With the ever increasing density of

25th ACM/IEEE DESIGN AUTOMATION CONFERENCE®

JOIN THE 25th ANNIVERSARY CELEBRATION

For 24 years the Design Automation Conference has been the meeting place for electronic CAD/CAM Engineers

A MUST FOR CAD/CAM ENGINEERS AND DEVELOPERS

Always 1st in the latest CAD/CAM Technological Innovations

NOW IN THE 25th ANNIVERSARY YEAR, DAC WILL OFFER:

- Over 130 papers, tutorials and workshops.
- Over 110 vendors of CAD/CAM hardware and software will exhibit their products.
- Over 50 vendor-technical presentations on Sunday afternoon June 12th - Many announcing new products.
- Full day of paid tutorials on Thursday June 16, 1988.

DYNAMIC APPLICATIONS AND CONCEPTS

Anaheim Convention Center, Anaheim, CA

June 12-15, 1988

Advance Registration ends **May 14, 1988**

For More Information, Call **1-303-530-4333**

acm SIGDA®

COMPUTER SOCIETY OF IEEE® DATC

HOTEL RESERVATION FORM
25th Design Automation Conference
June 12-15, 1988

Hotel reservations will be handled by the Anaheim Visitors & Convention Bureau. All reservations will be handled on a first-come first-served basis. The following hotels are available:

HOTEL	RATE SINGLE	RATE DOUBLE	HOTEL	RATE SINGLE	RATE DOUBLE
_ Convention Center Inn	$48.00	$48.00	_ Inn at the Park	$78.00	$78.00
_ Disneyland	$90.00	$90.00	_ Jolly Roger	$68.00	$68.00
_ Emerald	$88.00	$88.00	_ Marriott	$95.00	$105.00
_ Grand	$78.00	$78.00	_ Paul Bunyan	$55.00	$59.00
_ Hilton	$98.00	$112.00	_ Quality Inn	$78.00	$78.00
_ Holiday Inn	$68.00	$68.00	_ Raffles	$53.00	$58.00
_ Hyatt Regency Alicante	$82.00	$95.00	_ Sheraton	$73.00	$83.00
_ Hyatt Anaheim	$81.00	$81.00			

PLEASE MAKE FOUR CHOICES AND NUMBER THEM

RESERVATIONS SENT TO THE DAC OFFICE WILL BE DISCARDED

The Anaheim Visitors & Convention Bureau will attempt to make your reservations as indicated by the choice of hotel you specify. If the hotels you choose are not available, the Bureau will make your reservation at an available Conference Hotel. Complimentary busing will be provided from all Conference Hotels listed above the Anaheim Convention Center, when needed.

INDICATE HOTEL PREFERENCES ABOVE AND SEND COMPLETED FORM TO:

DAC HOUSING BUREAU c/o ANAHEIM VISITORS & CONVENTION BUREAU
P.O. BOX 4270
ANAHEIM, CALIFORNIA 92803
(714) 999-8939 - FOR INFORMATION ONLY - NO TELEPHONE RESERVATIONS!

FIRST NAME_____ LAST NAME_____
COMPANY_____
MAIL STOP_____
STREET ADDRESS_____
CITY_____ STATE_____ ZIP_____
ARRIVAL DATE_____ TIME_____ AM_____ PM
DEPARTURE DATE_____ TIME_____ AM_____ PM

A deposit of $75.00 is required by check or credit card. Make check payable to DAC HOUSING BUREAU.
Your reservation will be confirmed by both the DAC Housing Bureau and the hotel to which you are assigned.

CREDIT CARD INFORMATION

☐ VISA
☐ AMERICAN EXPRESS Number_____
☐ MASTERCARD Expiration Date_____

All changes must be made directly with the DAC Housing Bureau listed above, 72 hours before your arrival date or you will lose your deposit.

ADVANCE REGISTRATION FORM
25th Design Automation Conference

To register, mail this form or a facsimile to:

Mr. P. O. Pistilli
25th DESIGN AUTOMATION CONFERENCE
7490 CLUBHOUSE RD. #102
BOULDER, CO 80301

Telephone Registrations WILL NOT BE ACCEPTED. NO CREDIT CARDS.
For Information Only, Call (303) 530-4333.

PLEASE TYPE OR PRINT CLEARLY. Your name and company will be used to print your identification badge.

FIRST NAME _____ LAST NAME _____
COMPANY _____
MAIL STOP _____
STREET ADDRESS _____
CITY _____ STATE _____ ZIP ____
COUNTRY _____
TELEPHONE NUMBER (___)_____
*MEMBER # ACM/IEEE _____

CONFERENCE REGISTRATION FEES

(circle one)	Advance Registration Before May 14, 1988	Registration at Conference
Member ACM or IEEE	U.S. $120	$150
Non-Member	U.S. $150	$190
# Students	N/A	$75
*One Day Only	N/A	$70

#Good for one day of attendance for technical program and exhibits. Price does not include social events or proceedings.

*This discount is for individual memberships only! This does not apply to company memberships. If this individual number is not included, non-member fees will be charged.
For membership information, see page _____.

TUTORIAL REGISTRATION FEES, JUNE 16,1988

You must register for at least one day of the Conference to attend tutorials, but if you plan to get a one day only registration and attend a tutorial, you must register at conference.

MEMBER $140.00 U.S. Dollars (plus conf.
NON-MEMBER $180.00 U.S. Dollars reg. fee)

_____ Physical Design Automation
_____ Using EDIF to Describe Electronic Design Data
_____ Parallel Programming for CAD Applications
_____ An Introduction to VHDL

To determine your total fee, find your conference registration amount in the table above, and add $140 or $180 (whichever is applicable).

AMOUNT ENCLOSED $ _____

ASICs, we are often asked how we will be able to deal with the design and simulation of new or next-generation devices. Our response is simple. To properly design ASICs, one really needs a CAE system that is capable of simulating the entire system, not just an individual ASIC device. Any CAE system that can handle the top-down simulation of an entire supercomputer is certainly capable of dealing with ASICs of any size. The actual size and complexity of each ASIC really becomes a simple issue of partitioning.

■ Board-Level Testing Benefits

Each printed circuit board was also treated like an ASIC. A simple utility was written to allow signal traces from the simulation to be translated into the format of a circuit board functional tester. As the first boards were completed, we were able to independently debug these boards on the board tester without waiting for the entire board set to be completed and without having the uncertainty of which of several boards might contain the error when a problem was uncovered. With the extensive simulation done before fabrication of the first hardware, the number of problems actually encountered were few, but this method of testing each board against simulated results served to further reduce the time required to produce a fully functional machine.

■ Supporting the Design System

We have not discussed several other aspects of the CAE/CAD system used for the C series design, such as timing verification, gate array macrocell placement, and printed circuit board placement and routing. Yet they were also partially responsible for the success of the project. The full electrical design automation system in place at Convex is a mixture of products from several vendors, supplemented with tools and utilities created internally.

Although it may be desirable to have a design automation system supplied by a single vendor that can take a design from the architectural phase all the way through the generation of PCBs and ASICs, the single system that would exactly fit our needs did not exist when we started the C series development, and it does not appear to exist today. That does not imply that there are no systems on the market that attempt to cover this scope. Design automation vendors are making progress, but our view is that each of these systems has some weakness in one area or another.

Many vendors do not have the behavioral-level simulation power of the N.2 simulator; some do not have the ac timing verification capabilities of the SCALD timing verifier; others do not have the PCB placement and routing capabilities that we require; and so on. Also, many of the commercial systems restrict their use to a single platform. They may offer a choice of platforms, but once the platform is chosen, one is expected to perform all operations from interactive schematic capture to computationally intensive fault simulation on it. We have found it more desirable to turn the interactive tasks over to the workstation, while relying on the general purpose supercomputer for computationally intensive tasks.

Our experience has been that if one is willing to expend some effort in internal support, a better-fitting system can be created by picking the best items available from several different vendors. However, we should quantify the level of "internal support" that went into the series project. The entire CAE/CAD effort was supported through about half of the project by a single dedicated CAE engineer, and a CAD engineer was added later as the focus switched from simulation to printed circuit board routing. Also, we benefited greatly from the efforts of a computer systems manager who kept the design team up and running without fear of losing any of the valuable design data.

Now, these were (and still are) very intelligent and hardworking individuals, and one can argue that they did more work than their actual numbers would tend to indicate; but the point is that this entire design system did not require the support of an army of engineers. Throughout the design, the efforts of these people were supplemented by those of some of the knowledgeable and willing members of the system-level and ASIC design teams. In many cases, these designers would find a case where the system could be improved, and thus their "manual" effort minimized, by some operating procedure or utility program. They would then simply write the program or put the procedure in place. There is no doubt that the designers themselves are generally in the best position to recognize where they can most benefit from additional automation and to know exactly what form it should take. They also have a much higher sense of urgency when it comes to putting some design aid in place for their project than any central CAE/CAD organization could possibly have.

■ Summary

The success of the C series design is an excellent example of the benefits a thorough top-down design and simulation strategy. A design automation system that allowed us to take this approach was put in place from a mixture of commercially available products, along with some internally generated tools. Using workstations for the interactive tasks and a supercomputer for computationally intensive tasks allowed us to perform a large set of system-level simulations. Though this approach requires the investment of additional time before beginning the actual logic design, it can significantly reduce the total time to market and can yield a much more producible and error-free design. Most of the design problems were found and fixed on paper rather than in hardware. In the case of the C series effort, all of the semicustom chips and many of the printed circuit boards were shipped to the first customer on revision A. ■

About the Author

HAROLD DOZIER, *vice president of advanced development at Convex Computer Corp., leads the company's VLSI activities and is responsible for new CPU projects. Before joining Convex, as its first employee in 1982, Dozier was manager of single-chip microcomputer products at Mostek Corp. He received his BSEE and MSEE from Texas A&M University and holds several patents in semiconductors and IC digital logic design.*

The complex world of microprocessor development just got simpler.

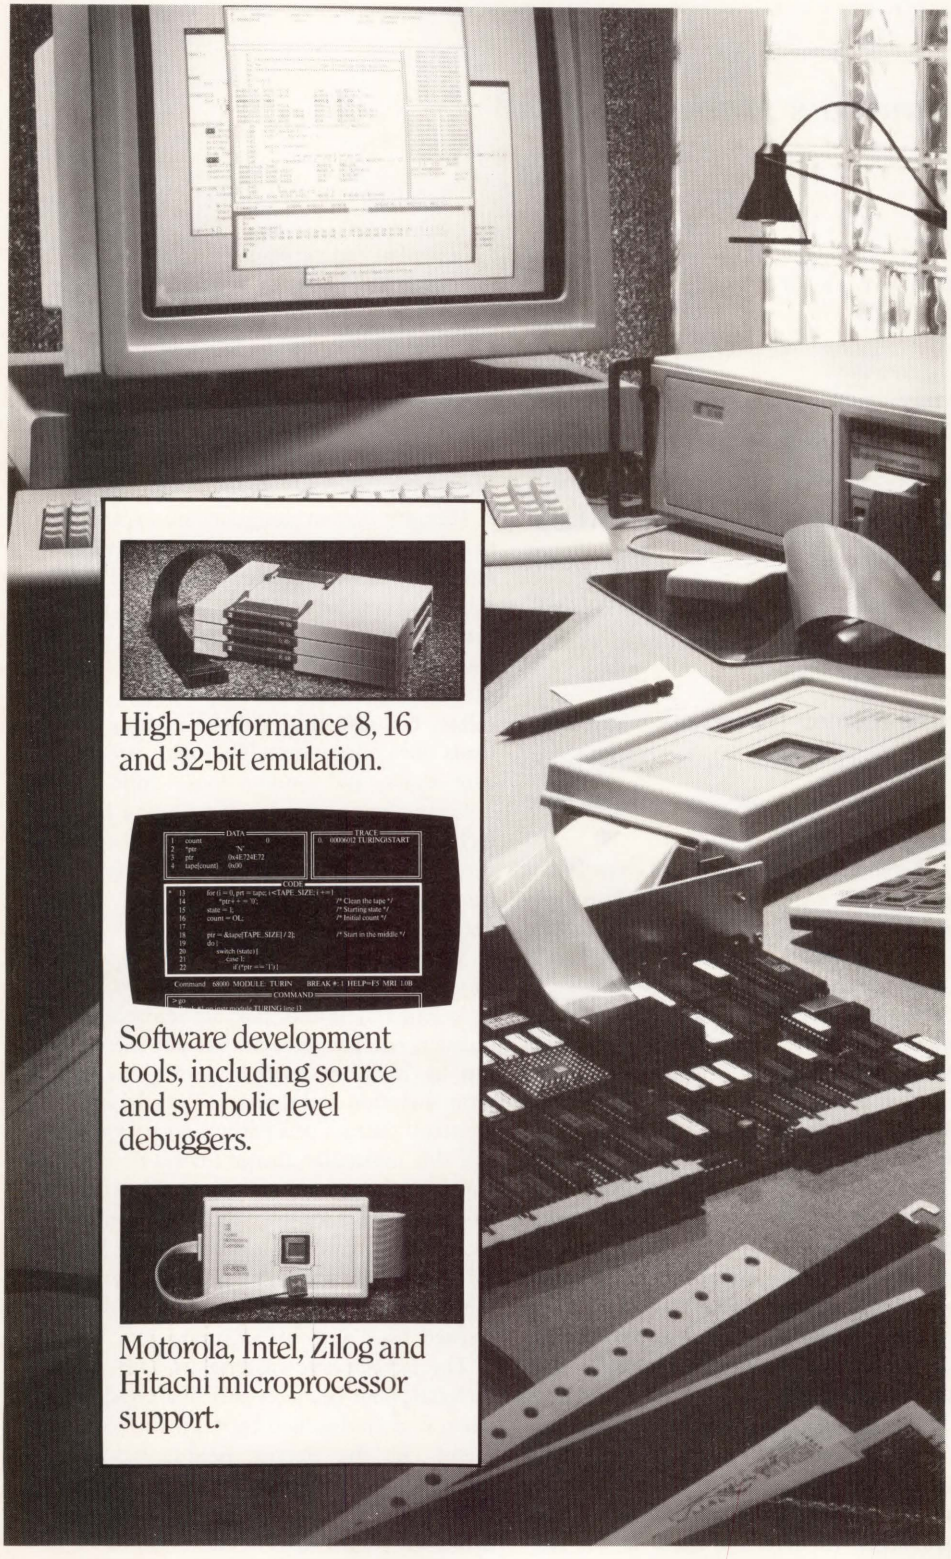

High-performance 8, 16 and 32-bit emulation.

Software development tools, including source and symbolic level debuggers.

Motorola, Intel, Zilog and Hitachi microprocessor support.

As the challenges of developing and debugging embedded microprocessor designs becomes more and more complex, you need the best development tools possible.

And nobody offers you higher reliability, better performance or a quicker way to develop and debug your 8, 16 or 32-bit systems than Applied Microsystems.

Our emulators, for example, can be controlled from the host computer you work with (VAX, Sun, Apollo, IBM PC or compatible). Your target system will run exactly as if its microprocessor were in place. And you'll get a clear picture of your design and how it interfaces with interrupts, clocks and the flow of software.

Our emulators are designed to work with easy-to-use VALIDATE software to give you powerful source level or symbolic debugging capabilities. You also get our powerful Event Monitor System, a hierarchical, real-time breakpoint and triggering system.

The list of benefits goes on and on; but the end result is a faster, easier and more accurate approach to developing, debugging and integrating your design.

To find out more, write Applied Microsystems Corporation, P.O. Box 97002, Redmond, Washington, USA 98073-9702. Or call (800) 426-3925, in WA call (206) 882-2000.

In Europe contact Applied Microsystems Corporation Ltd., Chiltern Court, High Street, Wendover, Aylesbury, Bucks, HP22 6EP, United Kingdom. Telephone 44-(0)-296-625462.
AMC-230

Applied Microsystems Corporation

CIRCLE NUMBER 41

Ferroelectric Chips

Are Memories Made of This?

STAN BAKER, EDITOR-AT-LARGE

THE CONTINUING quest for new memory technologies has led researchers back to the study of a well-known natural phenomenon, ferroelectricity. According to its proponents, ferroelectricity reapplied could provide the ideal nonvolatile memory. If current developments are extrapolated, ferroelectric devices could seriously challenge all other types of semiconductor memory devices; they promise essentially unlimited nonvolatility without degrading function, performance, density, or cost.

The developers of this new entry in the memory sweepstakes are bent on paring the ever growing number of memory device types. They hope to replace all currently used memory devices with one general-purpose nonvolatile ferroelectric RAM. If it can prove itself in the memory arena, ferroelectric technology will alter system architectures and software. Beyond that lie even more exciting possibilities: it has the potential to radically change logic devices.

■ UNUSUAL PROPERTIES

The new memory devices leverage an unusual electrical property of ferroelectric materials. These materials can permanently store charge as the orientation of ions in their lattice structures rather than as electrons in traditional capacitors, in which the charge can leak rapidly away.

Ferroelectric research began after 1921, when it was discovered that the polarization of Rochelle salt could be reversed by the application of an electric field. In later decades many more materials, in different chemical classes, were found to exhibit similar effects. In some cases the effect was a problem because scientists were trying to find better dielectrics, especially during World War II, when a substitute for mica was being sought.

A general understanding of ferroelectric phenomena was slowly developed over several decades. But in the 1960s a unified theory of the ferroelectric effect finally began to emerge. This advance gave rise to better metrology techniques and experimental methods, and after 1970 much of the work began to focus on feroelectric films and their place in electronic devices. This work led directly to the founding of Ramtron Corp. of Colorado Springs, Colo., and Krysalis Corp. of Albuquerque, N.M., to exploit the electronic possibilities of ferroelectricity. The prime result was the birth of ferroelectric materials with adequate reliability, good response speed, and operating voltages and dimensions that make them compatible with semiconductors.

The new patented technology has been coupled lately with silicon semiconductor structures. However, while the history and data of the materials research is readily available, the details of many of the electronic-related patents and their supporting data is not, since most of this information is classified as proprietary data by the companies involved.

Memory system designers are expected to apply these new nonvolatile devices rapidly over the next few years, and the champions of this technology are boasting that not only will ferroelectric devices compete with and replace many DRAMs, SRAMs, EEPROMs, EPROMs, PROMs, and ROMs, but also they will be making inroads into programmable logic devices, gate arrays, and random logic chips.

Certainly, that's a lot of bragging for a technology that hasn't made it to the field yet, but laboratory studies and results from prototype devices support the optimism.

The ferroelectric capacitor is inherently a bistable device. The ferroelectric material contains molecules that are ionized and unbalanced. The unbalance causes them to line up with one another, leaning in the same direction and therefore causing a net electrical charge, or a natural polarization; and this molecular charge rotates to align with an impressed electric field. The effect has also been used to electrically control the optical polarization of such materials. But the bistability comes from the hysteresis characteristic produced as the ions are twisted back and forth (Figure 1).

The ferroelectric material is used as a dielectric that becomes part of a capacitor when it is sandwiched between two metal plates. As the electric field across the capacitor is switched in direction, the ions follow the electrical field. However, a major difference from a conventional ca-

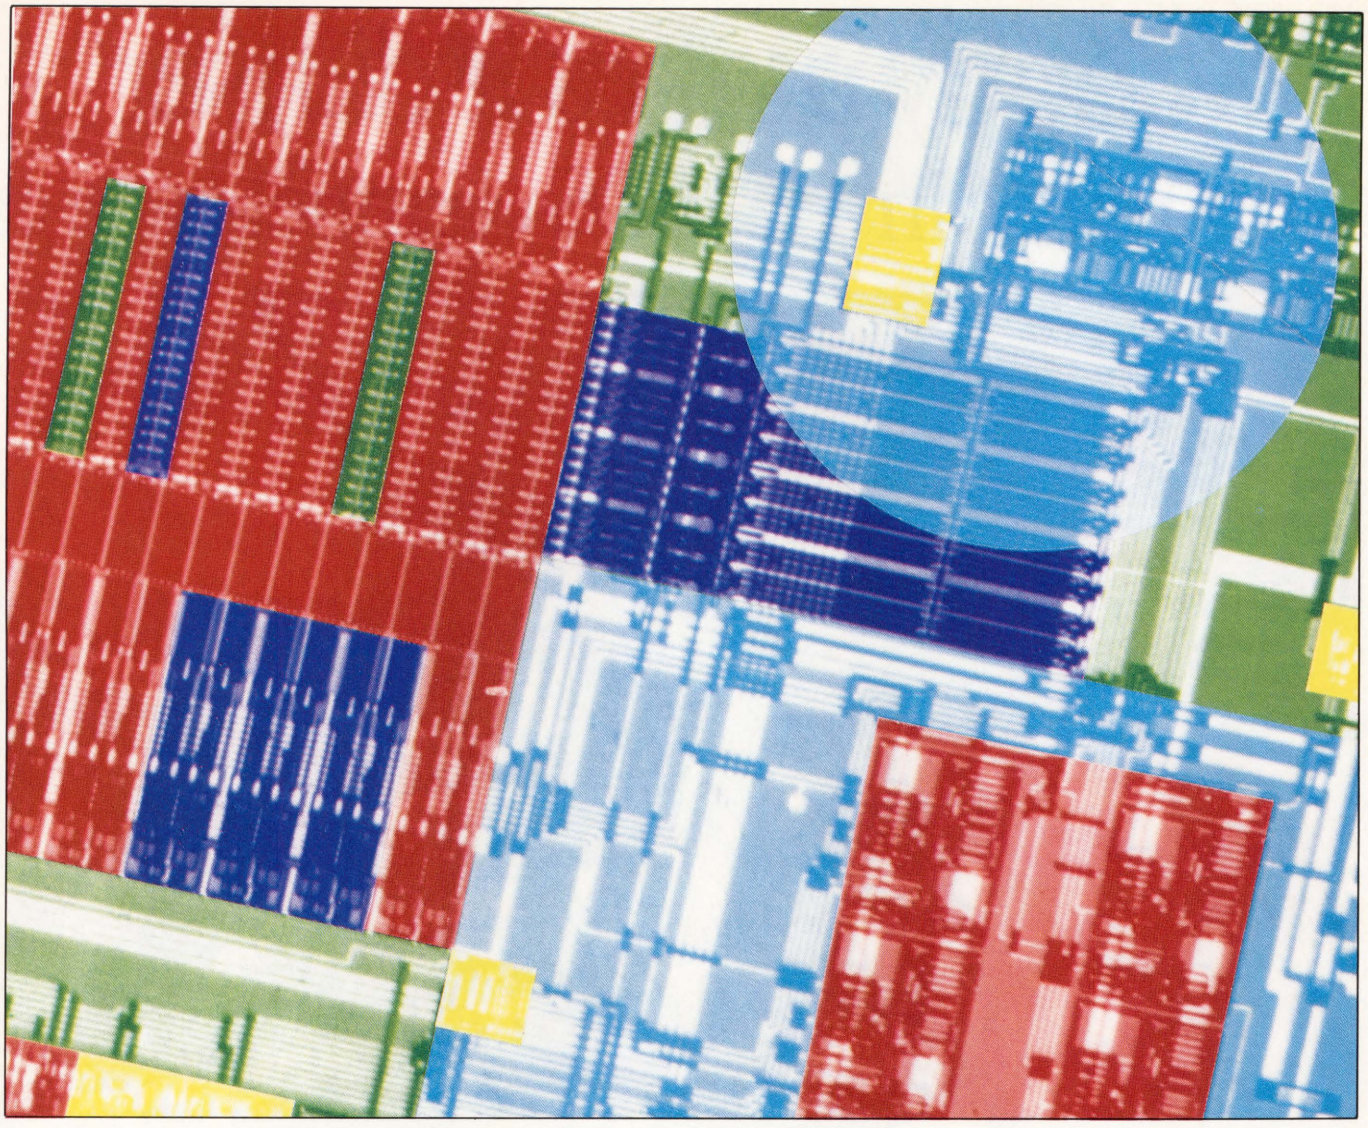

pacitor is that this charge orientation remains even when there is no voltage across the capacitor. In other words, what is involved is not the storing of a charge, but an alignment of existing charges; the phenomenon is analogous to magnetic poles aligning in magnetic storage materials.

The ferroelectric effect creates stable, reliable nonvolatile electrical charges. Its researchers project that the materials under development will give more than 100 years of nonvolatility and operate satisfactorily for more than 10^{15} cycles. At present, ferroelectric RAMs have achieved endurance ratings of at least 10^{11} cycles, whereas EEPROMs are achieving endurance ratings of 10^5 cycles at best. In addition, the new RAMs have good radiation resistance.

Experts from Ramtron and Krysalis say the ferroelectric capacitors will change about 2:1 over the full military temperature range, whereas the capacitors now used in DRAMs will exhibit even larger changes. Of course, there are other variations that can affect all types of memory capacitors. For example, manufacturing defects can degrade the capacitors to a point that it becomes difficult to accurately and reliably sense the output signals from the memory cells.

Additionally, to balance out the common-mode changes, the outputs must be sensed by differential amplifiers. As in DRAMs, dummy cells at the ends of rows or columns are used to drive a reference side of each amplifier. This is only one of many techniques that will be borrowed from familiar traditional memory chip designs.

The borrowing points out the underlying simplicity of what is happening. A new material with a rather simple phenomenon is being incorporated into traditional device structures. Integrating the new phenomenon has been worked out at the laboratory level. Putting it into general practice involves merging it with many familiar techniques.

Of course, the new devices can function as a nonvolatile memory only if the "states" can be detected reliably. Detection is not difficult, but the act of detection does alter the state. If a voltage pulse is applied to the plates with a polarity that reverses the charge alignment, the voltage source will have to supply a minimum amount of charge in order to effect this change. But if the polarity of the voltage pulse is such that it doesn't change the alignment, much less charge is required. This difference in charge is detected across a standard capacitor placed in series with the "ferrocap." In actual practice, the bit lines in a memory array form the capacitors collecting the charges to be sensed (see "Current Devices and Designs," p. 122).

Detecting the states of an array of ferrocaps is accomplished by driving them all to the same state and noting which ones change. However, as noted, the lost states must be restored if the device is used as a memory. The technique now used resets the lost states concurrently with the readout of data by the sense amplifiers. The resetting is internal to the chip and is invisible to the memory user.

Nobody makes it denser

45,000 gates in a CMOS channel-type gate

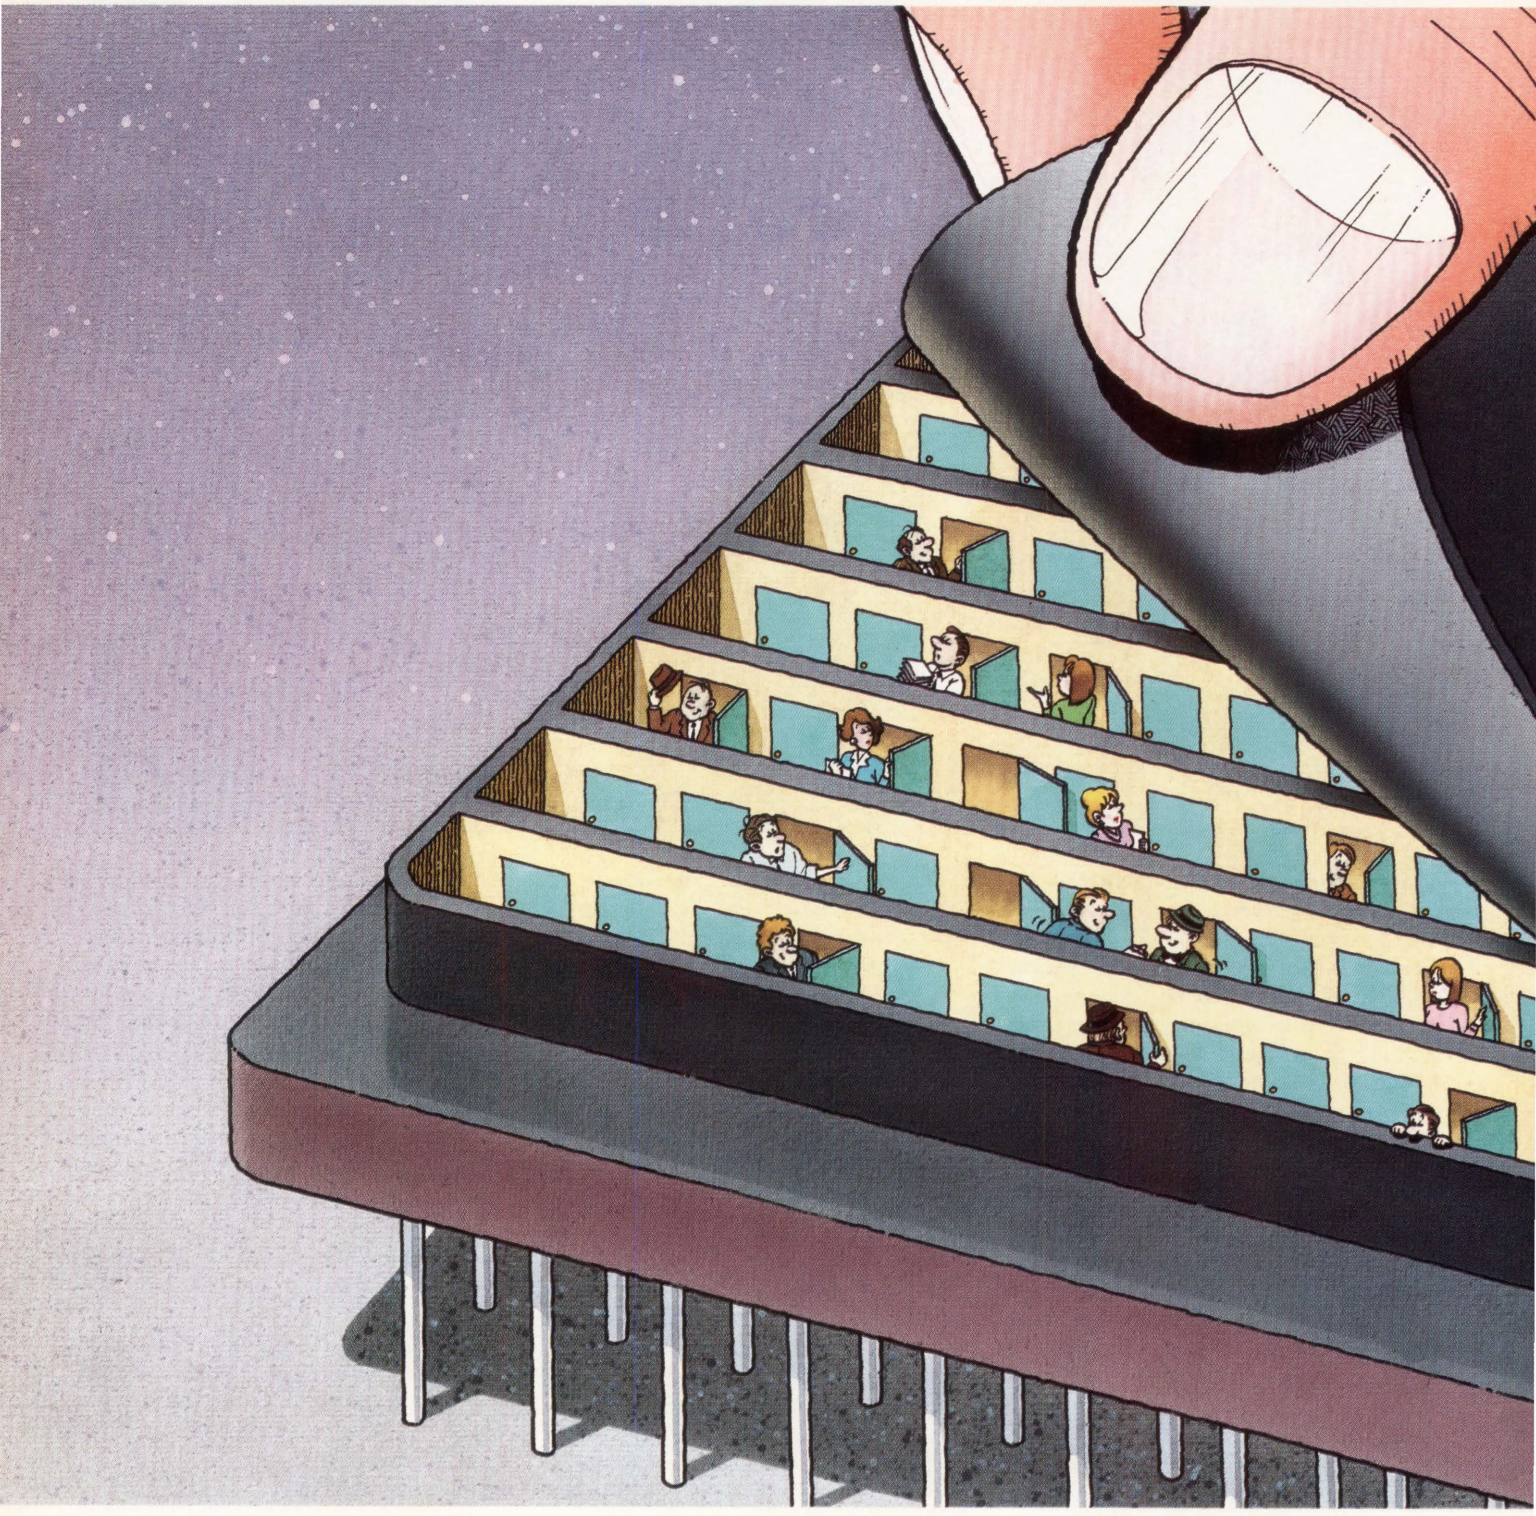

C&C Computers and Communications

array

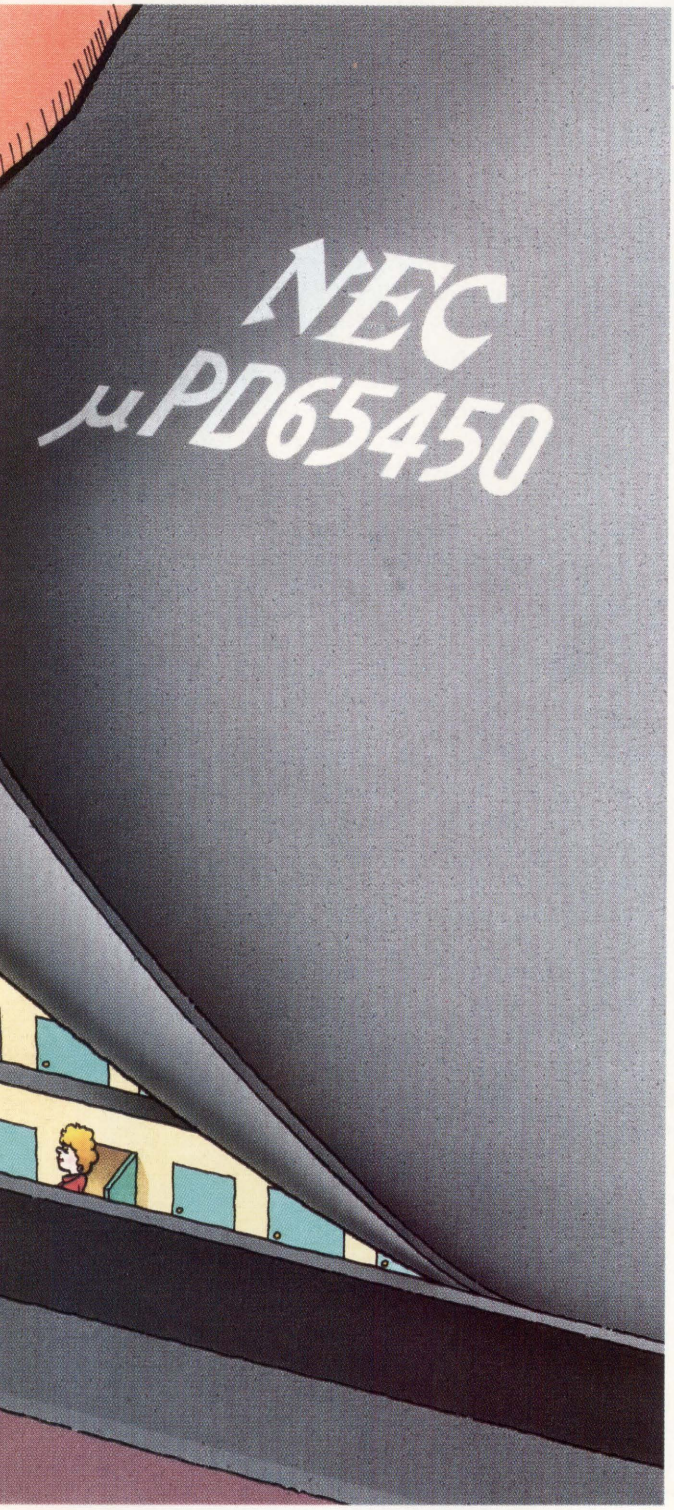

High integration and high speed—you get it all from NEC's new CMOS-5 and -5A channel-type gate array families. Gate counts start at 2000 and range all the way up to 45000, the world record for density in a channel-type CMOS device. Gate utilization is over 90% (μPD65450).

Speed rivals ECL performance. Internal gate delays are 1.0ns for a standard gate; 0.65ns for a power gate (F/O = 3, L = 3mm).

CMOS-5 and -5A families protect your system with high resistance to latchup—over 1A (typ). That's more than enough for security in most applications. Other features include:

- ☐ High drive capability: I_{OL} = 3—18mA.
- ☐ 5K-ohm pull-up resistor-incorporated buffers; ready for TTL bus lines.

NEC now offers you the broadest CMOS gate array lineup—six families, with 37 masters, including the world's densest channel-type gate array. For more information, call NEC today.

Family	Master	Gate Count	I/O Pads
CMOS-5	μPD65025	2016	71
	032	3366	81
	044	4440	112
	051	5292	120
	061	6348	132
	071	7500	144
	082	8748	152
	103	10800	164
	140	14256	188
	180	18144	224
	240	24000	256
CMOS-5A	μPD65300	30600	256
	450	45012	256

For further information, please contact:
USA Tel:1-800-632-3531. In California: Tel:1-800-632-3532. TWX:910-379-6985. W. Germany Tel: 0211-650302. Telex:8589960. The Netherlands Tel:040-445-845. Telex:51923. Sweden Tel:08-732-8200. Telex:13839. France Tel:1-3946-9617. Telex:699499. Italy Tel:02-6709108. Telex:315355. 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

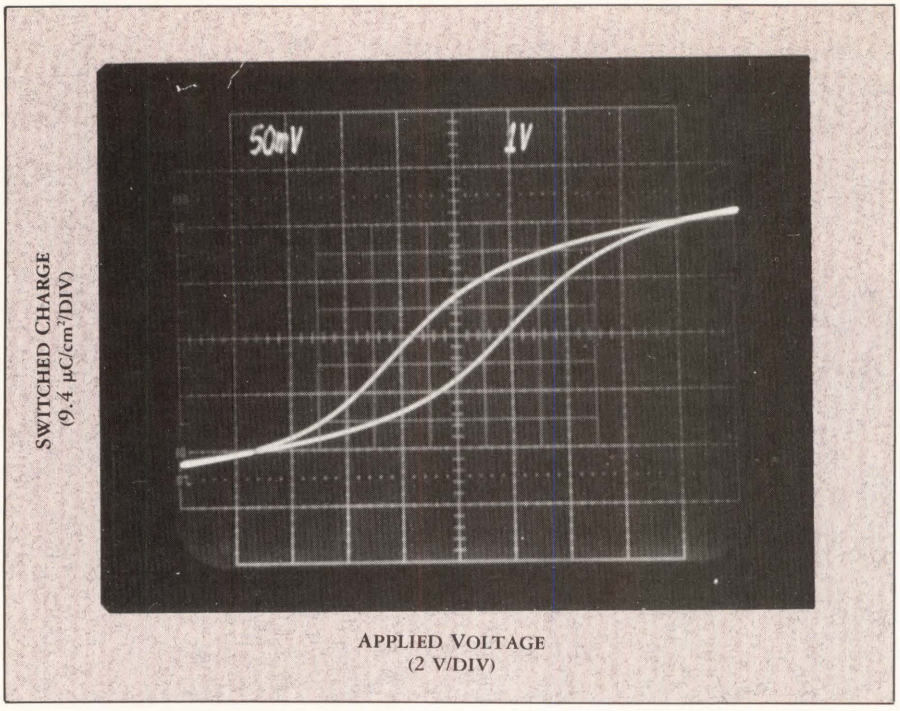

Figure 1. The oscilloscope trace shows the hysteresis loop of a ferroelectric capacitor. The two stored charge "states" are the two intercepts of the loop and the zero voltage vertical axis.

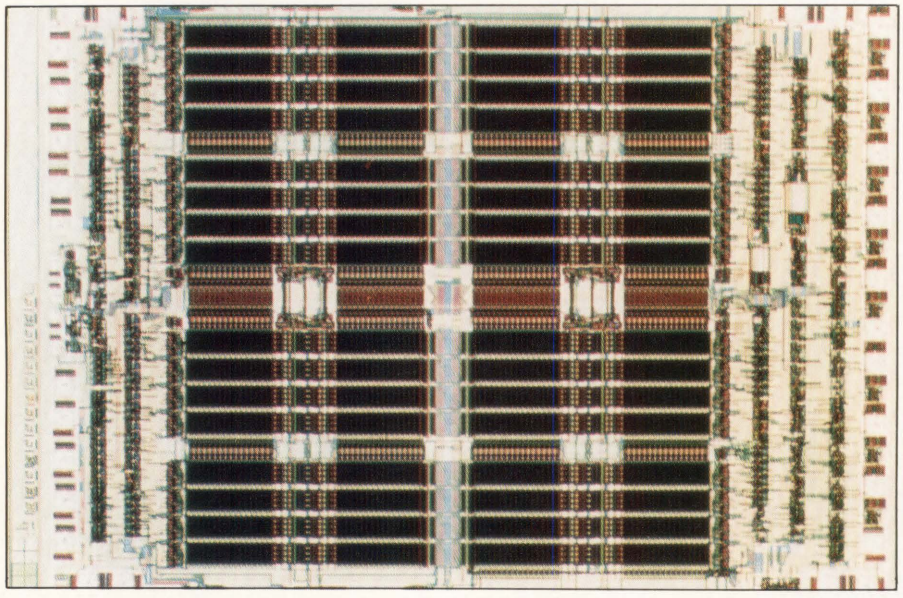

Figure 2. Krysalis has seen first silicon on its second device, a 16K UniRAM, which, like the 256-bit version, employs the ferrocapacitors actively in the nonvolatile memory.

To enter data into the memory, the state of one or more ferrocaps is set simply by impressing a voltage that is high enough to saturate the charge alignment in the desired direction. When data are read from the ferrocaps, the amplifier sensing the voltage accross the bit line capacitors drives the ferrocaps with positive feedback and forces them to return to the states that existed before sensing.

This basic theory seems simple enough. However, many pitfalls had to be overcome before making practical devices. Materials had to be developed that produced reliable readout signals—particularly with extremely small capacitor plate sizes, typically a few tens of square microns. The electrical characteristics must be predictable and stable. Also, many ferroelectric materials exhibit a "wearing" effect whereby the difference in readout from the two states gradually approaches zero. Additional material requirements include operation over the full military temperature range and the ability to be easily layered into the processing of silicon ICs.

■ USABLE DEVICES

In the last few years researchers at Ramtron and Krysalis have solved enough of the problems to field the first usable devices.

Most of the ferroelectric materials investigated were in the Perovskite family, selected for their crystalline structure. Ramtron evaluated a number of these materials, including potassium nitrate (KNO_3), bismuth titanate ($Bi_4Ti_3O_{12}$), lead germanite ($Pb_5Ge_3O_{11}$), and various lead zirconate titanate (PZT) compounds. It has had the most success with the PZT compounds, and its present memory chips use this group of compounds. Krysalis has had the best results with a "modified" lead titanate ($PbTiO_3$).

Krysalis made a prototype device with 512 bits and has recently fabricated a 16K device (Figure 2). It is also designing 256K devices. The company made an alliance with National Semiconductor Corp. in January under which National will fabricate Krysalis's designs and be a partner in development and marketing. Krysalis estimates the sampling of its 16K chip will begin late this summer.

Ramtron, which is making its silicon prototypes at the University of Colorado, has licensed the use its new technology with gallium arsenide devices to a young Australian semiconductor company, Ramex Ltd. Ramtron has built a 256-bit device (Figure 3) and is also designing 16K and 256K devices. However, its 256-bit device used the new ferroelectric technology to act as a shadow RAM. The ferrocaps act as a backup for a standard SRAM when the power is turned off, and they force the SRAM cells into the proper state when power is reapplied (see again "Current Devices and Designs," p. 122).

■ LOGICAL APPLICATIONS

The first goal for Ramtron and Krysalis

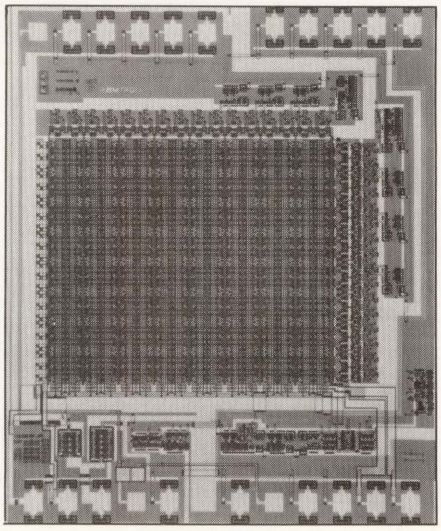

Figure 3. Ramtron's first device is the 256-bit FM801 FRAM, which uses the ferrocaps in a shadow RAM arrangement for power-off data storage.

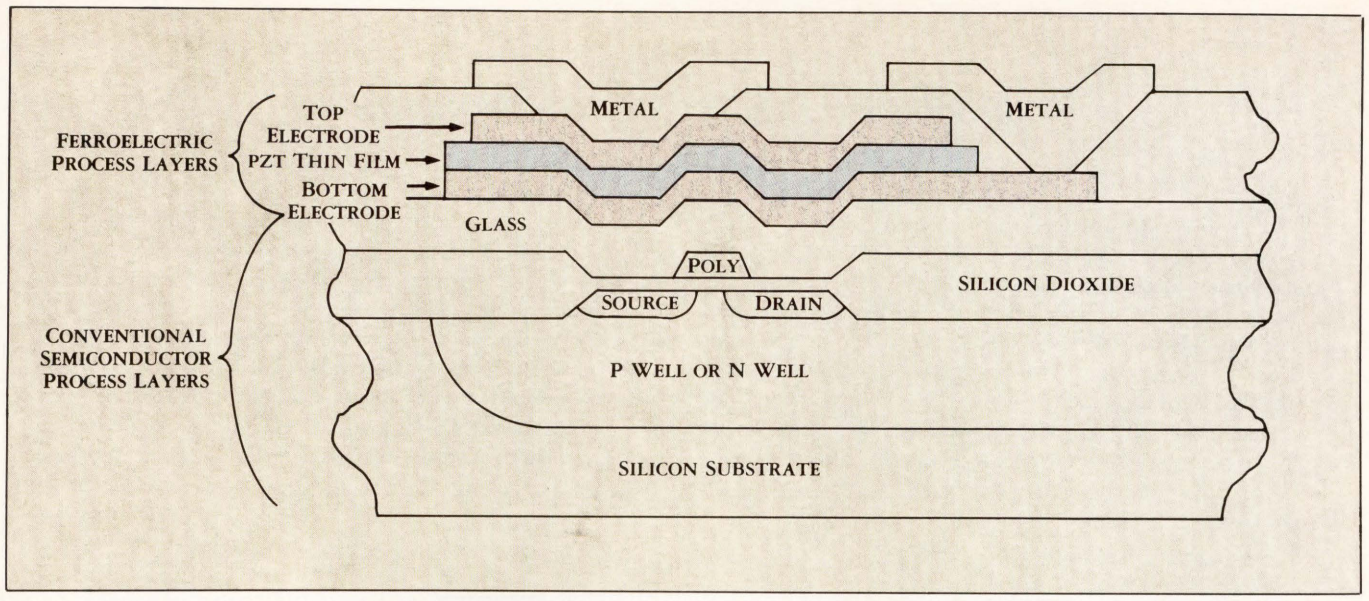

Figure 4. Ramtron uses a three-step process to layer thin-film lead zirconate titanate (PZT) ferroelectric capacitors on the surface of a CMOS SRAM structure.

is to bring ferroelectric devices to the memory market. But success there will surely be followed by applications in the logic arena. Both companies are researching all the possibilities, but they won't supply any details. They suggest that they don't have anything to talk about now, since their resources are focused on memories.

The most apparent way to use ferroelectric materials in logic is for the gate oxide in MOS transistors. The alignment of ions will interact with the flow of charge in the adjacent channel. The alignment could be switched, as in a ferroelectric memory, altering the current flow in the transistor channel and causing two states of bias. This action could lead to flip-flops with only one transistor and a quantum simplification of logic structures.

The logic possibilities, in both semicustom ICs and standard random logic parts, can be imagined, but neither company is willing to project what might happen.

■ Impact on Architectures

The long-range impact of ferroelectric memories on architectures is more evident than future logic based on the ferroelectric phenomenon. Today, the many types of memories support the complex hierarchies used in memory subsystems in the computer world, even down to desktop units. There are cache memories, main memories, read-only memories, dual-port memories, hard disks, floppy disks, and many more. If many of these can be replaced by one type of low-cost memory device, it would greatly simplify many aspects of computer systems. The hierarchy requires a considerable overhead of hardware and software and consequently increased price. The complexity also reduces system speed and reliability. Simplification thus has many architectural implications—for system software as well as hardware.

Having low-cost, high-density nonvolatile RAMs available would likely cause code that is now booted to become resident and capable of being updated without exchanging parts. Because of a different mix of cost, density, function, and performance, the handling of code would be reorganized.

There are no architectural implications from the 16K designs expected to be ready as samples by the end of this summer. But the 256K designs will offer the possibility of making significant changes in systems—primarily from power savings, simplification of the many types of memory subsystems into one, and elimination of the overhead circuitry required by DRAMs or EEPROMs.

■ Circuit Densities

A primary cost factor is the number of memory cells that can be crammed into a square centimeter. DRAMs of course have the smallest cell area. If other memory device cells are referred to a DRAM cell, with its area normalized as 1, we get the following approximate relations:

- DRAM = 1
- EPROM = 1
- Ferroelectric RAM = 2 (current designs)
- SRAM = 3–4 (depending on the number of transistors used)
- EEPROM = 2 (floating gate)

This picture, however, is changing, and will therefore force changes in the memory hierarchies, because cell size is related so strongly to cost. DRAMs are trying to race

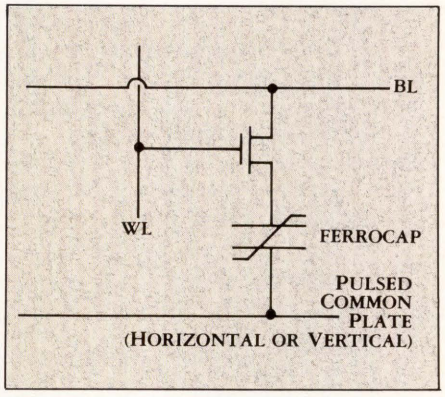

Figure 5. Ramtron is developing this simple one-transistor RAM cell, and Krysalis is working on a similar design.

further ahead in density, but shrinking the dimensions can make the capacitance too small for sufficient alpha radiation insensitivity or adequate sensing of signals. DRAM designers have been working around this problem by resorting to trench capacitors and by stacking transistors over and within the capacitors. Unfortunately, shrinking dimensions also have increased the complexity to the point that it's rapidly undermining the economic viability of the DRAM.

The nonvolatile structures show promise of closing the gap with DRAMs in terms of density. EPROMs are essentially there already, and stacked EEPROM structures are on the horizon that promise to do the same. The next round of ferroelectric RAM designs are likewise expected by their developers to close the gap.

Ferrocaps have 100 times more capacitance per unit area than those used in DRAMs, which are made with oxide films. A ferrocap in a cell can be hidden under the area of a contact, whereas DRAM designers are going to extremes of processing

Current Devices and Designs

The two companies working on ferroelectric memories have initially approachd the design of the memory cell from different directions. In Ramtron's ferroelectric RAM, the ferrocaps act like "shadow" memories. They are not used during normal read/write cycles; instead, they set the SRAM cells to the correct states when they are powered up. The ferrocaps also retain the data when the power is turned off. However, in the current designs at Krysalis, the ferrocaps are cycled during reading and writing. Whereas Ramtron reads and writes to the SRAM, using the ferroelectric cells as backup nonvolatile memory, the Krysalis chip actually reads, writes, and stores the data in the ferroelectric cells.

Ramtron's technique allows it to exploit the full speed of the SRAM cells, since the ferrocaps are not cycled through reversals during the read/write operation. Therefore the cycle time of the FRAM is essentially that of the CMOS SRAM, which for the comany's first chip is specified at 35 ns.

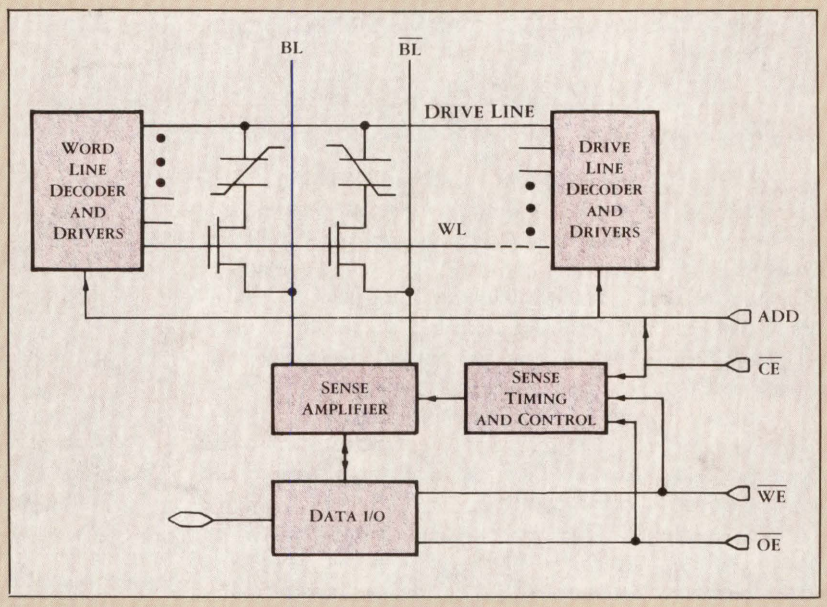

Figure A. In Krysalis's double-capacitor cell, the ferrocaps are used as the nonvolatile memory.

The present Krysalis devices, in which the ferrocaps are actively cycled in normal read/write operation, are slower, with 200-ns cycle times and 100-ns access times. The company's first chip is a 512-bit experimental version of a ferroelectric RAM. It is an externally controlled device with on-chip test circuitry and a separate timing interface that has proven useful in the company's process and timing development programs. Krysalis is also currently designing an improved 16K version and has seen first silicon (see Figure 2, main text). The devices are trademarked as UniRAMs.

Even though the Krysalis cell is slow today, ferrocaps are not the culprit, and the company expects future designs to cycle much faster. The ionic displacement mechanism is certainly slower than the response of electrons in normal dielectrics. But ferrocaps have been tested at frequencies well into the gigaherz range, and no one has yet experimentally determined the switching limits of the materials. According to researchers at Krysalis, experiments so far have indicated that their cycle times can be as low as 10 ns or less.

Krysalis is working on a 256K design with a cycle time goal of 100 ns or lower. Ramtron is also designing a 256K device, but it projects that the chip will achieve a cycle time of 55 ns.

In Krysalis's UniRAM cell (Figure A), the word line controls two pass transistors that connect the ferrocaps to the bit lines and the sense amplifier. The drive line pulses the opposite side of the ferrocaps during both the read and write cycles.

When writing into the cell, the sense amplifier is set to the desired state and used to drive the bit lines to opposite levels, that is, one to ground and the other to the value of the drive line voltage. When the high drive line voltage is impressed across the ferrocap that is tied to the grounded bit line, it writes a 0 state into that capacitor. But when the drive line drops to ground at the end of the drive pulse, the opposite capacitor has a 1 written into it from the high bit line voltage.

The read operation is such that when the cells are read, the

complexity, such as digging deep narrow trenches, to keep the surface areas of their capacitors small. Added to this is the fact that today's complexity of ferrocap processing will soon be reduced.

For example, Ramtron is now adding three mask steps on top of an SRAM structure (Figure 4) to make its FRAM (a Ramtron trademark). There is no increase in cell size because the ferrocaps are added on top of the SRAM cells. In the near future, though, the processing will be done in only one step: the double metallizations already required in the SRAM will be used for the plates of the ferrocaps, eliminating two metal steps. (The three steps are necessary now because the ferrocaps are being added to an already existing SRAM cell design.)

It therefore appears that the ferroelectric memory should match DRAMs in size as well as in complexity of manufacturing. Consequently, ferroelectric chips, with speed competitive with that of DRAMs, are potential threats not only to DRAMs, but to EPROMs, EEPROMs, and SRAMs as well.

■ SINGLE-TRANSISTOR CELLS

To close the density gap with DRAMs, ferroelectric memories must be made with cells that have only one pass transistor. That is the aim of both Ramtron and Krysalis for their 256K device, and they are taking essentially the same tack to reach their goal (see Figure 5 and "Current Devices and Designs," above). These memory cells will be smaller than planar DRAM cells made with the same design rules but a little larger than DRAM cells that use stacked techniques. Here, the difficulties will be to manitain the long endurance achieved in the four-transistor cells and still attain good speed.

But the single-transistor cell will also have an important power advantage over

ferrocap charged opposite to the reading voltage is reversed and the data is temporarily lost from the ferrocap. But during this "destructive" readout part of the cycle, the bit lines are disconnected from the sense amplifier. These "floating" bit lines collect charges during the readout. As the readout is completed, the regenerative sense amplifier is reconnected to the bit lines, and the polarity across the bit line capacitors is such that regenerative amplifier drives the bit lines and ferrocaps to the states they originally had at readout. The typical difference in voltage developed across the two bit line capacitors for this type of cell is approximately 1 V. The sense amplifier simultaneously provides an output to the I/O circuitry. The process of restoring the read data to the memory cells is invisible to the user and is designed to occur in parallel with the output gating of the read data to the I/O ports.

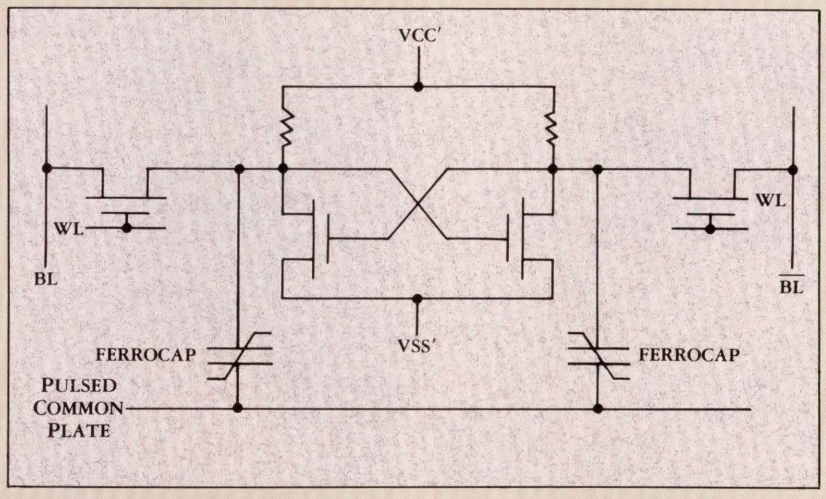

Figure B. Ramtron employs a four-transistor static RAM cell in which two ferrocaps act as a "shadow" nonvolatile memory.

The Ramtron cell operates in a completely different mode. The ferrocaps act as a backup to a standard SRAM (Figure B) and are only asked to store data when the power is off. Of course, they must also force the SRAM to read and store the ferrocap data in the SRAM cells when the power is restored. During normal operation, the bit and word lines are used only for normal reading and writing to the flip-flops in the SRAM, and the pulsed common plate (PCP) of the memory cells remains tied to ground.

If we assume that the V_{CC} supply for the SRAM is +5 V and V_{SS} is tied to ground, then the upper plates of the ferrocaps will only see +5-V pulses during normal read and write operations, when the PCP is grounded. That means that both ferrocaps will usually end up polarized with +5 V on the upper plates, since the upper plates of the ferrocaps are constantly bombarded with +5-V pulses. When the user's external circuitry senses that the power is going off, it must pulse the PCP line to +5 V. The ferrocap that is tied to a 1 state will see its voltage drop to zero (since both plates of the capacitor are now at +5 V). This ferrocap won't change its state, but the ferrocap tied to the 0 state will receive an effective −5-V pulse (since the PCP is tied to +5 V and the upper plate is grounded to zero). This 5 V of opposite polarity is sufficient to drive the ferrocap to the opposite polarity. Thus the ferrocaps effectively "read" and store the data that was in each SRAM cell at power-off.

One benefit of this design is that the ferrocaps are reverse-cycled only at power-up and power-down, which extends their effective endurance. Ramtron reports that it is currently achieving endurance figures of at least 100 billion read/write cycles and expects improvement by several orders of magnitude in the next few years, while Krysalis is specifying the endurance of its 16K UniRAM as "unlimited."

For designs of 256K and higher, Krysalis is planning on using a single-transistor cell design that can take advantage of the finer geometry design rules. The single-transistor cell is essentially one half of the present Krysalis cell, but with the sense amplifier shared by all the stages along the bit line and the other input of the sense amplifier attached to a dummy capacitor, for common-mode rejection. One drawback of this design is that there will be less differential signal provided to the amplifier, resulting in lower noise immunity. ■

SRAMs, as well as a density advantage. If the speed can equal that of SRAMs, the power and density improvements will be highly significant.

The new technology will first be put into devices that plug directly into existing sockets. Once accepted in these forms, it will forge its own way, altering the systems it is used in and becoming an essential part of new types of systems.

For example, Ramtron intends that ferroelectric RAMs will first start to replace EEPROMs and battery-backed SRAMs. For its part, Krysalsis says that aiming at existing nonvolatile sockets. In programmable logic devices, they will likely be incorporated into chips such as Logic Cell arrays, the programmable gate arrays made by Xilinx, and soon also to be offered by Advanced Micro Devices. These PLDs hold their logic patterns in an associated off-chip nonvolatile SRAM, which can be replaced by the ferroelectric RAMs.

At Ramtron, marketeers are projecting that by the early 1990s:

- One-transistor cells will be used.
- 1.5-μm rules will be reduced to 0.8-μm rules.
- The added three mask steps of its four-transistor device will become only one.
- The cell size of 100 μm^2 will drop to 10 μm^2.
- Ferroelectric devices will reach 1 and 4 megabits in size.
- Cycle times will be cut to 25 ns.

Ramtron is also very optimistic about ferroelectric RAMs becoming competitive with other technologies within a few years. Although the initial cost per bit is some 30% higher than that of EEPROMs and roughly 30 times that of SRAMs and DRAMs, the company is projecting that the cost per bit will drop below that of EEPROMs by 1990 and equal those of DRAMs by the mid-1990s. ■

D·A·C *turns* SILVER

Still crazy after all these years

The Design Automation Conference will celebrate its silver anniversary next month in Anaheim, Calif. From a humble beginning in Atlantic City, N.J., it has become the premier event for the multibillion-dollar design automation industry.

Early in 1964, Pat Pistilli, today's DAC conference manager, promoted a SHARE/Design Automation Workshop. It was a success; a "couple of hundred" engineers and designers showed up to hear 20 papers presented at the workshop, which became an annual event. In 1969 the Association for Computing Machinery became a sponsor, followed shortly by the Computer Society of the IEEE.

It was 19 years before the conference held its first exhibition, at which a handful of companies showed their wares. But this year more than 120 vendors will be touting their latest offerings.

In addition to this year's bumper crop of 124 papers, covered in 34 sessions, there will be two tutorials, five panel sessions, and a special commemorative session that gives a historical view of CAE/CAD.

The tutorials, presented Monday afternoon, will be "Automating the Design of Electronic Packaging" and "Opportunities in Computer-Integrated Manufacturing." The former will discuss the effects that different manufacturing and packaging methods have on the final performance of an overall device, including mechanical and thermal problems. The latter will identify some of the problems encountered in CIM today that can be addressed with the tools and know-how of the design automation community.

ILLUSTRATION GENE GRIEF

Four of the five panel sessions will target CAE/CAD tool environments, the integration of heterogeneous tools, and simulation of a complete system that includes everything from simple glue logic to complex ASICs, memories, microprocessors, and DSPs. The last panel will discuss the new cell generation tools and their future role.

As in previous conferences, approximately 60 of the DAC exhibitors will hold brief product previews during a five-hour marathon session starting at 2:00 PM on Sunday.

In addition, on Thursday, following the conference, four paid full-day tutorials, aimed at the design automation professional, will be offered: "Physical Design Automation," "Parallel Programming for CAD Applications," "Using EDIF to Describe Electronic Design Data," and "An Introduction to VHDL." The tutorials, which are limited to registrants of DAC, will be held in the Anaheim Marriott Hotel.

■ Sessions and More Sessions

As simulation continues to grow in popularity, the speed at which the simulation can take place is getting increasing attention. Three sessions are devoted to methods of boosting simulation speeds.

On Tuesday morning, "Parallel Simulation" looks into using the inherent parallelism typical in logic circuit operation as a means of speeding up the process. The very next session, "Mega-Simulation Accelerators," analyzes the results of using both large dedicated computing resources and extensive software support. The third session, the next morning, will be devoted to incremental techniques that reduce the total simulation time by reusing previously computed results.

The cost of testing has become one of the major contributors to the total manufacturing cost of ICs, and a number of the technical sessions will cover test-related issues. On Monday afternoon, a session will be devoted to design for testability. In addition, other sessions will cover fault simulation, timing and connectivity verification, and automatic test generation, including some new methods for checking sequential circuits without the full penalties of traditional LSSD methods. Hardware and physical design verification will also have their own sessions.

The latest progress in logic synthesis will be under the spotlight at five sessions to be held on Tuesday and Wednesday. Micro-architecture, register-transfer-level, and high-level synthesis will be covered, as well as logic optimization techniques.

The design automation environment also will be explored, in several sessions and panels. Three will be devoted to databases, two to hardware design languages, and four to present and future design automation systems and environments.

Ten sessions will be targeted at the latest techniques in physical design, but only a few will be devoted to the design of printed circuit boards. The majority of papers will be aimed at ICs: new and improved simulated annealing and sea-of-gates placement techniques and other improved placement, routing, and compaction methods.

—Roland C. Wittenberg

25th ACM/IEEE

DESIGN AUTOMATION CONFERENCE

FOR ADDITIONAL INFORMATION:

MP Associates Inc.
7490 Clubhouse Road, Suite 102
Boulder, Colo. 80301
(303) 530-4333

ANAHEIM CONVENTION CENTER, ANAHEIM, CALIF., JUNE 12-16, 1988

TIME	ARENA	CALIFORNIA B	CALIFORNIA A	PACIFIC A	ANAHEIM
MONDAY 8:00-9:30 am	Opening Remarks, Awards, and Keynote Address	—	—	—	—
10:00-11:30 am	**Twenty-Five Years of Electronic Design Automation** Chair: A. Richard Newton	—	—	—	—
1:30-3:30 pm	—	**Design for Testability** Chair: Vishwani Agrawal	**VHDL in Use** Chair: Lt. John Hines	**Floorplanning and Area Estimation** Chair: Kenneth Roberts	Tutorial: **Automating the Design of Electronic Packaging** Barry Whalen (1:30-2:30)
					Tutorial: **Opportunities in Computer-Integrated Manufacturing** David Hodges (2:30-3:30)
4:00-5:30 pm	—	**Automatic Test Generation** Chair: Randal Bryant	**Hardware Design Languages and Environments** Chair: Mario Barbacci	**Physical Design Acceleration** Chair: Jonathan Rose	Panel: **Future Computing Environments for DA** Chair: Andrew Rappaport
TUESDAY 8:00-9:30 am	—	**Timing Verification** Chair: James Kleckner	**Parallel Simulation** Chair: Alberto S.-Vincentelli	**Channel and Global Routing** Chair: Manfred Wiesel	Panel: **Behavioral Modeling for System Design** Chair: Tom Blank
10:00-11:30 am	—	**Applying Formal Verification to Hardware Design** Chair: John Darringer	**Mega-Simulation Accelerators** Chair: Rob Smith	**Xerox Parc Design Automation System** Chair: Bryan Preas	Panel: **What Is a Design Automation Framework, Anyway?** Chair: Wayne Wolfe
1:30-3:30 pm	—	**Engineering Information Databases** Chair: Randy Katz	**Application-Specific Simulation** Chair: Dale Hocevar	**Placement Algorithms** Chair: Ulrich Lauther	**High-Level Synthesis** Chair: Daniel Gajski
4:00-5:30 pm	—	**Distributed Databases for the Support of Design Teams** Chair: Steve Lusky	**DA for Analog Circuits** Chair: Chin-Fu Chen	**Layout Compaction** Chair: Alfred Dunlop	**Register-Transfer-Level Synthesis** Chair: Aart de Geus
WEDNESDAY 8:00-9:30 am	—	**Logic Synthesis and Optimization** Chair: Robert Brayton	**Transistor-Level Layout** Chair: Dwight Hill	**Physical Design Verification** Chair: Lou Scheffer	Panel: **CAD Tool Needs for System Designers** Chair: Randal Bryant
10:00-11:30 am	—	**1988 High-Level Synthesis Workshop Overview** Chair: Ewald Detjens	**Incremental Techniques in Logic Simulation** Chair: Mark Horowitz	**Gate Extraction and Connectivity Verification** Chair: Carl Ebeling	Panel: **Will Cell Generation Displace Standard Cells?** Chair: Alfred Dunlop
1:30 -3:30 pm	—	**Microarchitecture Synthesis** Chair: Ted Kowalski	**Routing-Related Subjects** Chair: Michael Burstein	**Ideas in Testing** Chair: Jacob Abraham	**Data Structures for Integrated Design Systems** Chair: Arnold Goldfein
4:00-5:30 pm	—	**Ideas in System Generation** Chair: Osamu Karatsu	**Ideas in Placement and Routing** Chair: Jiri Soukup	**Fault Simulation** Chair: Ernst Ulrich	**Ideas in Circuit Verification and Simulation** Chair: Resve Saleh

PHYSICAL DESIGN

SIMULATION, VERIFICATION, AND TESTING

SYNTHESIS

DATABASES, LANGUAGES, AND DESIGN ENVIRONMENTS

CMP PUBLICATIONS

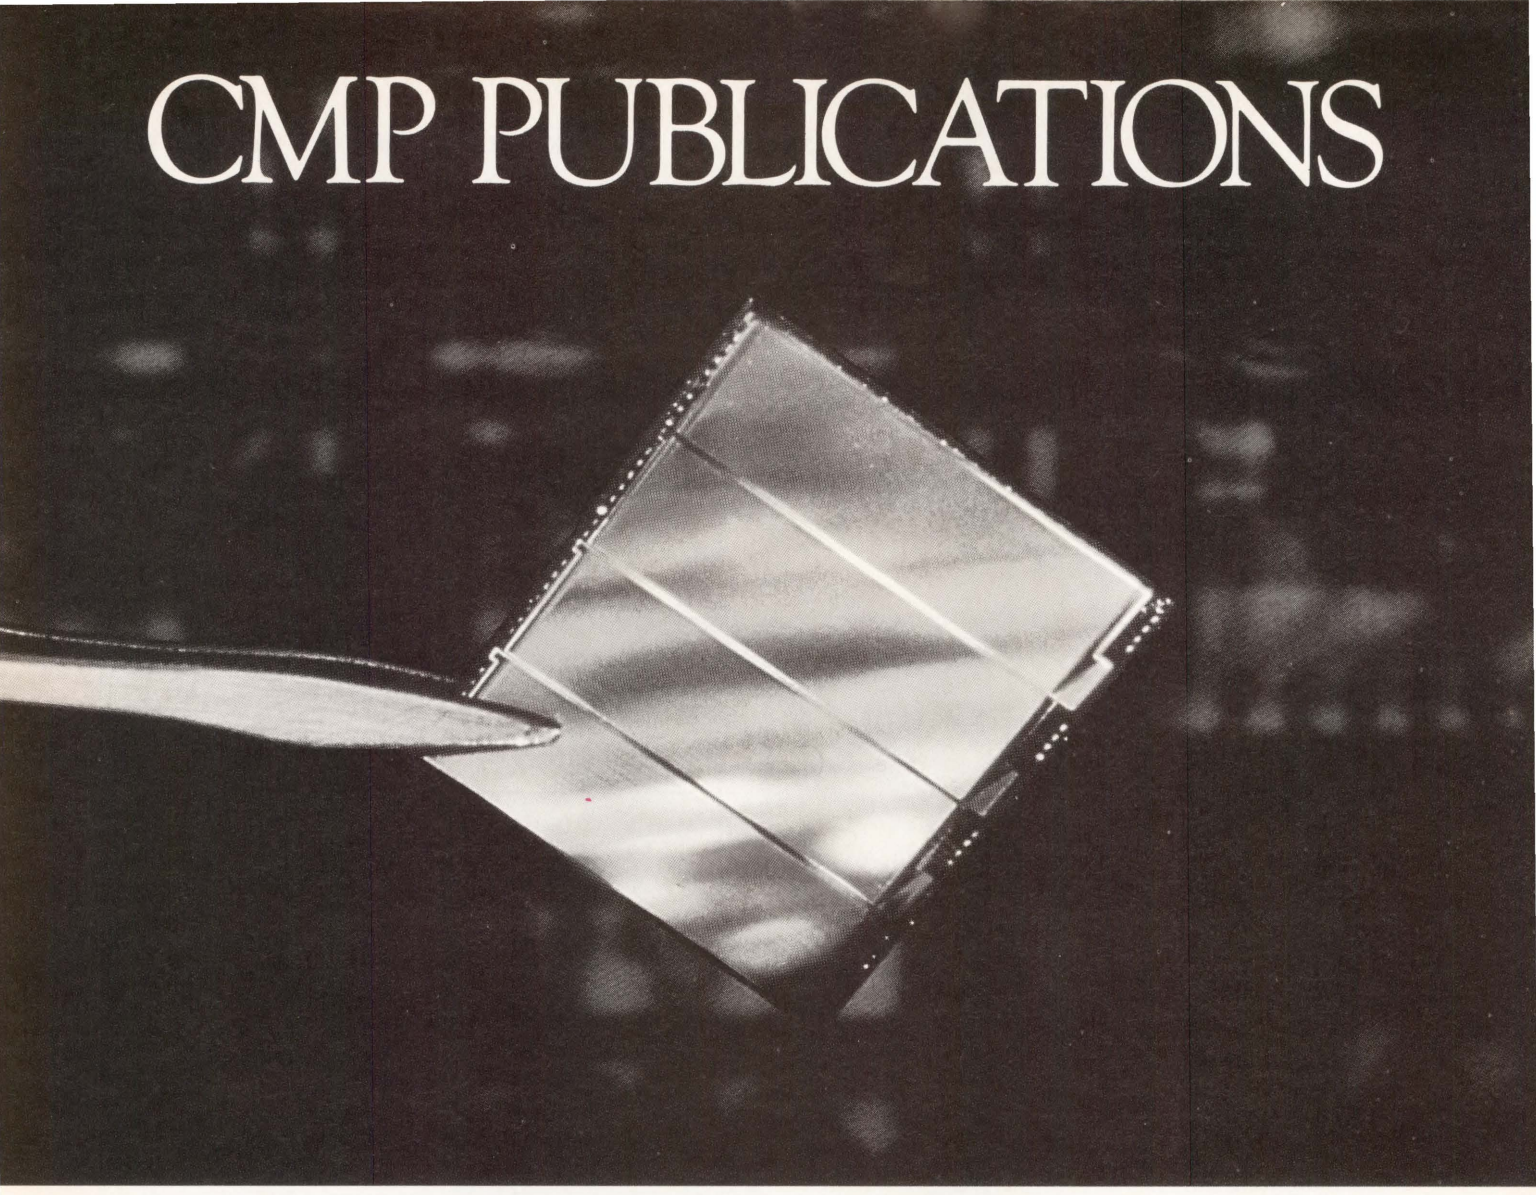

IN ELECTRONICS

The electronics industry has brought greater change to the way we think design, build and communicate than any other industry. CMP publications serve its information needs.

Electronic Engineering Times is the leader in delivering business and technology news to engineers and technical managers. *Electronic Buyers' News* is the number one publication reporting on business and procurement trends important to purchasing and other corporate managers. *VLSI Systems Design* serves the innovators in the design and use of very-large-scale integration technology. *Electronic Buyers' Handbook* is a reference for buyers who need product information.

READERSHIP GETS RESULTS

All CMP publications build readership in dynamic industries through editorial excellence. Innovative circulation programs reach each industry's key decision makers. Research identifies their information needs. And outstanding editorial delivers the news, analysis and perspective that meet those needs.

The result is exceptional readership. And readership gets results for our advertisers.

Find out how CMP can get results for your advertising. Contact CMP Publications, Inc., 600 Community Drive, Manhasset, New York 11030. (516) 365-4600.

CMP publishes *Electronic Buyers' News, Electronic Engineering Times, VLSI Systems Design, Electronic Buyers' Handbook, Computer Systems News, Computer Reseller News, VARBUSINESS, Business Solutions, CommunicationsWeek, InformationWEEK, Business Travel News, Tour & Travel News, Manufacturing Week, HealthWeek.*

New RISC Designs for Computers and Embedded Systems

Intel and Motorola Toss Their Chips into the Ring

INTEL AND Motorola have taken the wraps off their new RISC architectures. The Intel 80960 family is designed for embedded control applications, putting it in competition with the Am29000 from Advanced Micro Devices Inc. (Sunnyvale, Calif.) and the VL86C010 from VLSI Technology Inc. (San Jose, Calif.). The Motorola M88000 products are more applicable for workstations and minicomputers, like products from MIPS Computer Systems Inc. (Sunnyvale), Sun Microsystems Inc. (Mountain View, Calif.), Intergraph Corp. (Huntsville, Ala.), and their semiconductor partners. The similarities in the new architectures highlight some of distinguishing features of RISC processors; their differences underline the split in RISC approaches—those designed to meet the needs of high-speed multiprocessing computers versus those aimed at low-cost, flexible embedded controllers.

Both the 80960 and the 88000 (see article, p. 24) are built around a large set of registers, one characteristic trait of RISC processors. Both have a selection of processing units that plug onto a bus though which they access the register file, and scoreboarding techniques control the use of registers by the processing units. Scoreboarding tracks the accumulation and dispersion of register data as it is manipulated by the processing units. In effect, it implements a type of data-flow architecture, in which instructions are regularly assigned to processing units but execute only when all of the data is available.

The 80960's integer unit is rated by the company at 7.5 VAX-equivalent MIPS, and its floating-point processor is rated at 4M Whetstones/s, both at 20 MHz. Both are integrated on the 80960KB (see figure), which runs overall at 20 MHz, is available now in a 132-lead PGA, and costs $390 (100-piece quantities). The floating-point unit executes 32-, 64-, or 80-bit arithmetic functions according to the IEEE-754 standard. In comparison, the 20-MIPS R3000/R3100 chip set from MIPS Computer Systems and its semiconductor partners costs $895 per chip.

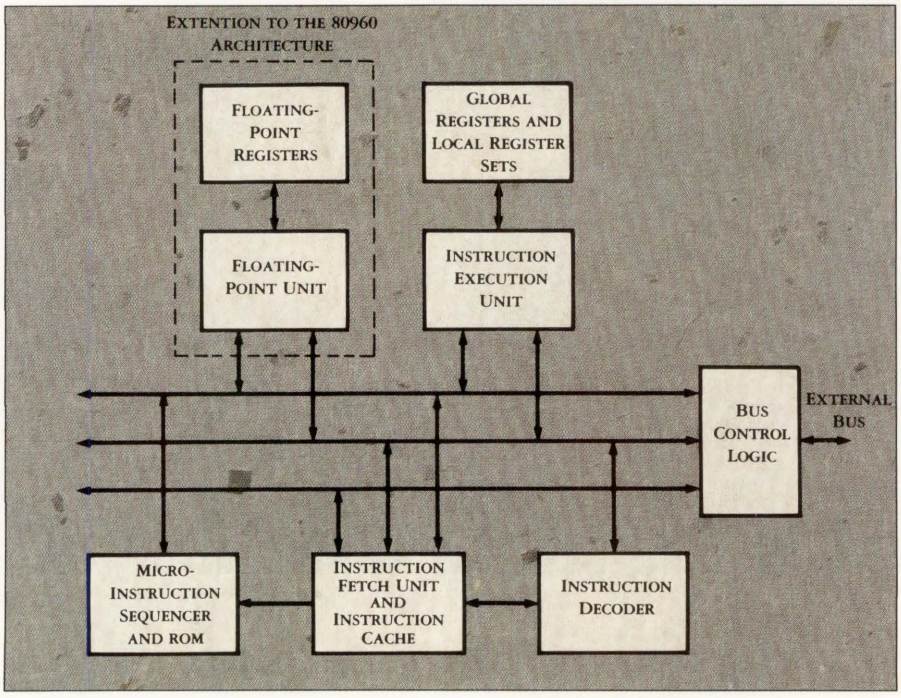

The 80960KB architecture contains concurrently executing processors and functional blocks.

Within both architectures, concurrently operating circuitry keeps the registers full and the instructions executing. The high degree of parallel operation within the architecture is a trait borrowed from supercomputer architectures—for example, as the authors of the 88000 article point out, Motorola borrowed design concepts from the Control Data 7600 supercomputer. The parallelism helps the chips toward another goal of RISC design: single-cycle execution of all instructions. In fact, with multiple processing units on the register bus, such architectures could occasionally execute more than one instruction during a clock cycle.

Like the 80960KB, the MC881000 CPU includes a floating-point processor. In contrast, floating-point coprocessors are used for the R3000, Sun's SPARC chip set, and the 29000 chips. The use of coprocessors does not imply inferior performance, however; MIPS' R3010 is rated at 7 MFLOPS for single-precision arithmetic and at 4 MFLOPS for double-precision.

Both the 88000 and the 80960 were designed with modular techniques, making it easy for the companies to create new versions of their architectures—what Intel calls "application-specific processors" (a term also used by Texas Instruments, but very differently), a concept critical for embedded controllers. In a departure from the pre-announcement of many 32-bit processors, both companies had functional silicon of the processors at the time of announcement. As a result, these products

should be relatively mature and bug-free. In addition, having silicon makes it easier to complete software development systems and tools. Starter development kits based on a personal computer are already available from Intel and a hardware model of the chip is available from Mentor Graphics Corp. (Beaverton, Ore.).

Now for the differences in the two architectures, which, as stated, underline the disparate needs of central processors and embedded controllers.

The 80960 is designed to work with low-cost memory, a feature critical in such products as printers and graphics systems. It has a single multiplexed bus, making board design simpler than with dual-bus RISC processors. It has a version with no floating-point unit, the 80960KA, and a version for embedded military systems, the 80960MC. It contains a 512-byte instruction cache but no data cache, relying instead on its bank of 32 registers—16 local and 16 global. In addition, a register cache of three sets of 16 registers backs up the local registers, lowering the overhead of switching tasks.

The military version contains all the circuitry on the 80960KB plus a memory management unit. It also supports the design of fault-tolerant systems through the use of a separate Bus Extension Unit, the M82965. The latter pairs two 80960MCs operating in parallel, comparing operations each machine cycle. When a discrepancy is detected, a fault is raised and a recovery procedure is initiated. Both the 80960MC and the M82965 will be qualified to MIL-STD-883C in the first quarter of 1989. The 80960MC is priced at $2,400 in 100-unit quantities; the M82965 runs $1,700.

The 88000 architecture has the features needed for central processing units in workstations and computers. Separate data and instruction buses give maximum throughput. A separate cache-and-MMU chip, the MC88200 CMMU, keeps maximum throughput on each bus and serves as the basis for cache coherency schemes. The MC88100 CPU chip has the capability to integrate as many as five additional processing units on the register bus, and it can perform up to 11 operations concurrently. In addition, the 88000 series has an expandable, integrated cache architecture that not only supports multiprocessing, but also fault tolerance.

The 88100 sells for $495 and the 88200 for $795, also in 100-unit quantities. ■

Intel Corp.
Chandler, Ariz.
(800) 548-4725

Motorola Inc.
Austin, Texas
(512) 440-2839

NEWS ANALYSIS

continued from page 14

general-purpose op amps, comparators, voltage references, switches, matched-resistor building blocks, capacitor cells, and analog input and output cells.

As the cells move from 2-μm to 1.5-μm technology, a few of the specifications improve noticeably: op-amp gain-bandwidth products change from 2.4 to 2.8 MHz, and response time for ECL interface circuits, which NCR specifies as analog cells, drops from 6.0 to 4.5 ns. Other than that, the 2.0-μm specifications are very similar to those of the 1.5-μm cells. Part of the reason is that transistors in analog cells often don't use the minimum channel length. The name of the game for analog circuits is control over the operating characteristics, so the analog transistors are often two or more times larger than the minimum channel length.

Plans for new 2-μm cells include an 8-bit, 5-MHz analog-to-digital converter; a 10-bit dual-slope A/D converter; an 8-bit, 2-MHz digital-to-analog converter; a rail-to-rail common-mode op amp; a low-current op amp; a higher-speed comparator; and a VCO with output frequencies as high as 30 MHz. Such figures show how a larger technology can make use of greater control over transistor dimensions to make higher-quality cells.

IMP's 1.2-μm cell libraries contain predesigned functions that use a cell architecture of fixed-height, variable-width blocks. The digital library contains 46 SSI- and MSI-level core cells and 24 I/O-pad cells. It is complemented by compilers for RAM, ROM, and PLA blocks. It is available now; the analog module of the processing technology, which adds a second layer of polysilicon to the processing for stable capacitor structures, and the analog cells will be available in the fourth quarter of this year. This process contains both lateral and vertical pnp transistors for low-noise analog inputs and high-drive output circuits, respectively. The process is an n-well process, allowing IMP to add EEPROM cells in the future.

The 1.2-μm digital cells from IMP don't seem so speedy when compared with new digital cells in NCR's 1.5-μm VS1500 cell family. The new cells, which have a toggle rate of 140 MHz and a system clock rate of 80 MHz, include eight high-performance flip-flops and latches. A representative cell, a D flip-flop, has a setup and hold time of 0.7 ns, a minimum clock pulse width of 1.0 ns, and a clock-to-Q specification of 0.9 ns typical. This performance is 70% better than that of existing cells in the 1.5-μm family. The speed increase for all new cells is between 30% and 100% over the original cells.

In addition, NCR has added new I/O-pad cells with drive capabilities of up to 48 mA. Schmitt triggers, oscillators, 4-bit adders, and power-on/reset cells have been added too. New cells with restricted slew rates help minimize voltage pulses from rapid signal transients. With these new cells, the performance is 50% better than that of 2.0-μm cells; and the density of digital circuits increases about 28%. Designs from the 2-μm cells can be migrated to the 1.5-μm technology directly, requiring only resimulation. ■

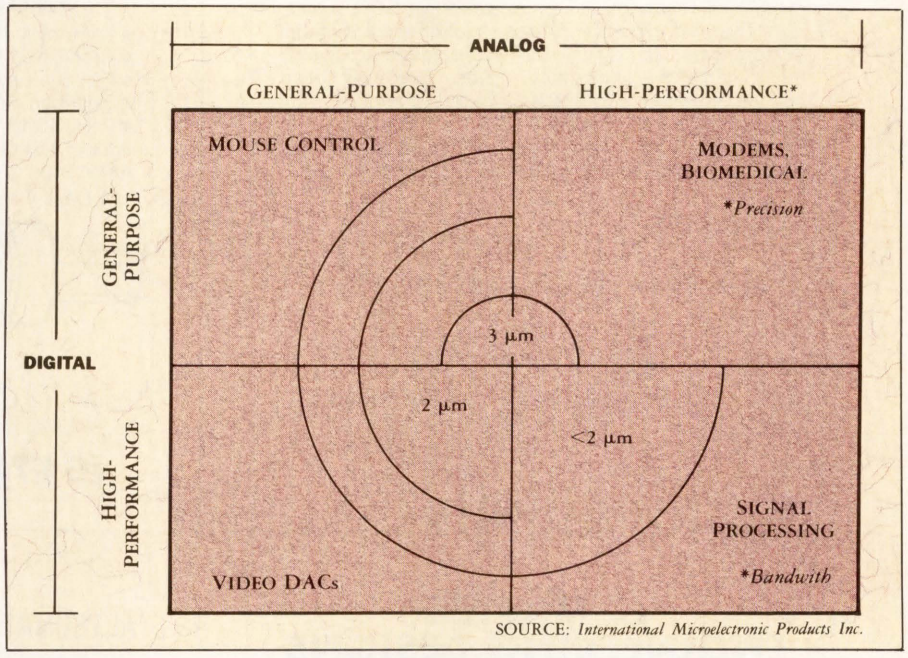

Different applications, be they analog or digital, call for different process technologies.

DESIGNER'S MARKETPLACE

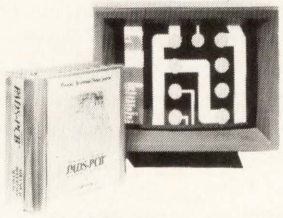

PADS-PCB™
PCB CAD
Layout System
Basic Package **$975**

New -- Version 2.1

It's all there -- 3000+ part library, routing, autoplacement, 1 mil resolution, 30 layers, checking, ECO, reporting. The post processor has pen plotter and dot matrix printerdrivers. For a little extra you can add a powerful autorouter and a photoplotter driver. Even with these options the whole package is just a little over $2K. And the system runs on cheap MS-DOS machines.

We believe you'll like it if you try it so we're offering a **FREE** evaluation kit with 4 disks and a 66 page manual.

Call for your **FREE** evaluation kit -- today.

Distributed by
Zyrel, Inc.
(408)433-0488 In California
(800)553-8838 Outside California

PADS-PCB is a registered trademark of CAD Software, Inc

CIRCLE NUMBER 46

AZURE COAX SCOPE

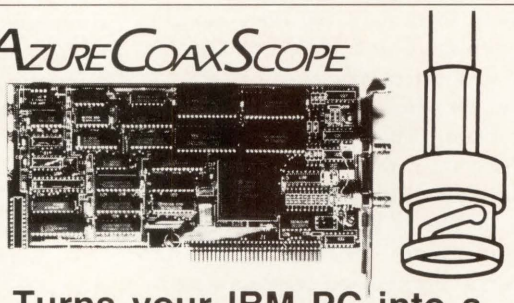

Turns your IBM PC into a 3270 coax line monitor

- Completely Passive Monitor
- Records up to 16 MB of Data
- Function Keys, On-line Help
- Extension Search Features
- Polls Compressed at Capture
- µsec Accurate Times Stamps
- Hardcopy of all Displays
- 3299 Multiplexer Option
- DFT Protocol Option
- SNA Protocol Option

14 Days Free Trial Evaluation

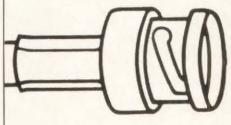

Priced at $2995.

AZURE
TECHNOLOGIES
38 Pond Street, Franklin, Massachusetts 02038

(617)-520-3800

CIRCLE NUMBER 47

IC MASTER ENGINEERING DESIGNS BEGIN HERE

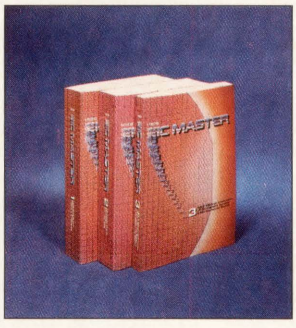

All new in 1988. The complete guide to integrated circuits. Now in three volumes covering all commercially available ICs including 10,000 new devices and related products. Has complete alternate source directory listing over 92,000 replacement devices, master selection guides, ASICs, part number guides and more. If you don't have the new 1988 IC MASTER you are wasting time and money. Regular price (US) $130 plus $10 shipping and handling.

IC MASTER
645 Stewart Ave.
Garden City, NY 11530
516-227-1300 FAX 516-227-1901

CIRCLE NUMBER 48

IC MASTER ALTERNATE SOURCE DIRECTORY Software

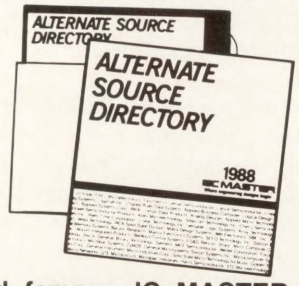

The world famous IC MASTER Alternate Source Directory is now available in floppy disk form. The program allows extensive modification to suit individual needs. The database includes exact pin for pin replacements for over 100,000 ICs. It is organized by original manufacturer and shows page references in the IC MASTER. Everything you need to find the alternate source you need...fast! For IBM XT/AT and compatibles. US $149.50.

Hearst Business Communications
645 Stewart Ave.
Garden City, NY 11530
516-227-1300 FAX 516-227-1901

CIRCLE NUMBER 49

DESIGNER'S MARKETPLACE

ANALOG CIRCUIT ANALYSIS
ECA-2

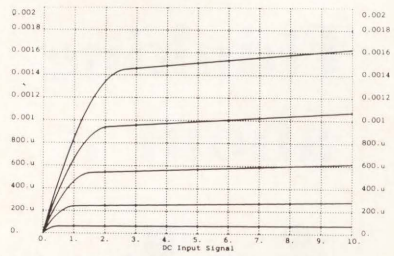

ECA-2 is a high performance, low cost analog circuit simulator. **ECA-2** 2.31 includes more models, increased graphics capabilities, increased AC analysis speed, and expanded documentation. Now available for the Macintosh.

ECA-2 Offers

- AC, DC, Transient
- Fourier, Temperature
- Worst-case, Monte-Carlo
- Full nonlinear simulator
- Interactive or batch
- SPICE compatible models
- Twice as fast as SPICE
- Over 500 nodes
- Detailed 525 page manual
- Sine, Pulse, PWL, SFFM, and Exponential generators
- Money back guarantee

ECA-2 IBM PC or Mac $675
ECA-2 Apollo $3000

Call 313-663-8810 For FREE DEMO disk

Tatum Labs, Inc.
1478 Mark Twain Court, Ann Arbor, MI 48103

CIRCLE NUMBER 50

AUTOMATED DEVICE CHARACTERIZATION and MODELING

If you are concerned with **IC design or with semiconductor manufacturing process monitor and control,** or if you are involved with **power supply design or board-level analog design,** you are also concerned with CAD/CAE Tools and Test Systems for...

- Semiconductor Parametric Test
- Computer-Aided Model Parameter Extraction
- Commercial Analog Component SPICE Model Libraries for Board-Level Simulation & Design

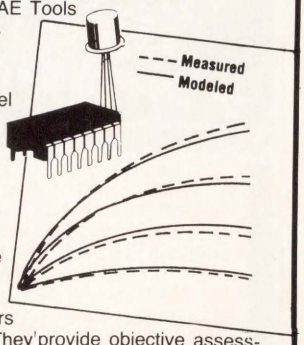

MOSEC reports reduce risk and minimize time spent by engineers and managers who need such tools. They provide objective assessments of the available commercial solutions with:

- Comparative in-depth reviews of commercial products from HP, Keithley, Lomac, Silvaco, Acotech, Daisy, Analog Design Tools, AB Assoc., Microsim, Intusoft, Meta-Software, and many other vendos.
- Discussion of methodologies
- Industry trends
- In-house presentations available.
- Inquiry privileges
- Benchmark Results
- Satisfaction quaranteed.

MOSEC

(415) 726-3039
591 Terrace Avenue
Half Moon Bay, CA 94019

CIRCLE NUMBER 51

DESIGNER'S MARKETPLACE

REACH OVER 35,000 SYSTEMS DESIGN ENGINEERS IN DESIGNER MARKETPLACE

IT'S AS EASY AS A PHONE CALL TO ANGELA ANDERSON OR NELDA BRANCH

800-645-6278 OR (IN NEW YORK) 516-365-4600

UNIVERSAL LOGIC PROGRAMMER

New! BP from MICROSYSTEMS $798

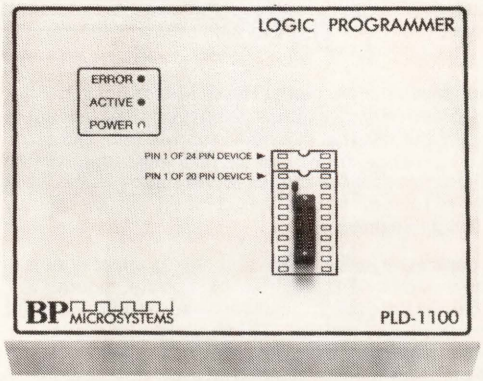

- PROGRAMS, READS, DUPLICATES, TESTS AND SECURES HUNDREDS OF 20- AND 24-PIN LOGIC DEVICES
- MENU DRIVEN OPERATION IS EASY TO LEARN AND QUICK TO OPERATE
- 23 UNIVERSAL PIN DRIVERS WITH INDEPENDENT DAC, ADC & SLEW FUNCTIONS PROGRAM ALMOST ANY LOGIC DEVICE
- CONNECTS TO ANY IBM COMPATIBLE VIA PARALLEL PRINTER PORT
- EDITS FUSE DATA & TEST VECTORS WITH FULL SCREEN EDITOR
- AUTOMATICALLY TESTS WITH VECTORS & SECURES AFTER PROGRAMMING
- SUPPORTS ALL POPULAR PLD DEVELOPEMENT SOFTWARE
- ONLINE HELP FUNCTION
- SELF CALIBRATING
- JEDEC FILE INPUT & OUTPUT
- GOLD TEXTOOL ZIF IC SOCKET
- EPROM PROGRAMMERS ALSO
- ONE YEAR WARRANTY
- SAME DAY SHIPMENT
- 30 DAY MONEY BACK GUARANTEE
- UPDATABLE VIA FLOPPY
- TOLL-FREE TECH SUPPORT

CALL FOR FREE DEMO DISK OR INFO 800/225-2102

10681 HADDINGTON #190 HOUSTON, TX 77043
713/461-9430 TLX: 1561477 FAX: 713/461-7413

CIRCLE NUMBER 52

The Intel Influence

It's Excitement. It's Challenge. It's Opportunity. It's Pride.

We're on the move. Lots of growth. Exciting projects and breakthrough technologies in systems and components. In short, there's never been a better time to be at Intel.

We have the influence and the resources that can make a profound impact on your career. If you thrive on hands-on responsibility and would like a chance to excel, look into the following opportunities.

CAD Group or Workstation Managers

Requires hands-on experience in CAD/CAE industry with at least 3 years in a management capacity. MS in CS or EE required.

CAD Engineers

Requires 3+ years experience developing or supporting CAD tools utilized by Design Engineers for ASICs, microprocessors, microcontrollers, peripherals, memories or other devices. In-depth knowledge of simulation, timing analysis, modeling, APR, widows or graphics experience is highly desirable. MS/BS in CS or EE required.

CAD Software Networking Engineers

Requires 1+ years of UNIX or TCP-IP networking experience.

CAD Technical Writers/Course Developers

Support in-house CAD tools utilized by engineering community for all levels of design. Some positions require the ability to write technical documentation and/or the ability to train our users on new tools developed in-house. All positions require 3+ years experience with a BS in CS or EE.

Product Marketing Engineers/ Systems Marketing Managers

Requires a BS/MS in EE with an MBA preferred. Individuals will be considered for marketing positions in the 80486, microcontroller, graphics, systems, telecommunications and ASIC areas. Both tactical and strategic positions are available.

Design Engineers

Microprocessors/Memories/Graphics/ ASICs/Microcontrollers/Systems

Requires 3+ years experience in logic or circuit design. Experience with CAD tools such as circuit/logic simulators and layout verifiers for designing CMOS desired. Some positions require experience in microcode development/ logic design. Assembly language (286/386), computer architecture, memory management techniques, pipelined implementation, microsimulators, 80x86 based architecture, functional design, RTL modeling, EGA or PC/AT systems design experience required along with a MS/BS in CS or EE.

We welcome your interest in Intel...please send your resume to the location of your choice.

Arizona: Intel Staffing, Dept. 4115, 5000 W. Chandler Blvd., Chandler, AZ 85226

California: Intel Staffing, Dept. 4115, 3535 Garrett Drive, Santa Clara, CA 95050, or call Rebecca at 408-765-1546 if you do not have a resume and would like to be considered for one of these sites

California: Intel Staffing, Dept. 4115, 1900 Prairie City Road, Folsom, CA 95630

Equal Opportunity Employer M/F/H

intel®

SEMICONDUCTOR ENGINEERING

Digital has it now.

Blue Chip Opportunities

Blue Chip opportunities are not just on the stock market.

You'll find them in Digital's Semiconductor Engineering group.

You'll be working for one of the most successful groups in a remarkably successful company.

Apply your skills to VLSI design using emerging semiconductor processes, design tools and packaging technologies. Build advanced COS microprocessors for multiprocessor configurations.

Develop advanced CPU and peripheral megacells for use in VAX*-based systems on a chip, working with advanced processes and CAD tools to define design methods and deliver subsystem building blocks for final silicon system products.

We're looking for innovative, enthusiastic and visionary designers at the consulting level (more than 10 years' experience) and principal engineering level (more than 5 years.)

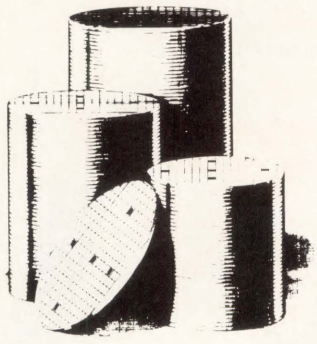

Semiconductor and CAD Design

PRODUCT MANAGER

Maintain a product management organization in support of the Semicondutor Engineering Group. Duties include writing of high level strategic business plans, End of Life responsibility, support of MOS Marketing base plan and customer interface. Requires at least 5 years' product management experience in the semiconductor industry.

VLSI Design

MODULE DESIGNER

Drive a high performance CPU module into production. Integrate custom VLSI CPU chips and semicustom interface chips to deliver a complete CPU module. Experience in high speed circuit design, transmission line concepts, as well as reliability and manufacturability will all contribute to the introduction of a revolutionary product. Experience delivering high speed module designs through manufacturing required. Requires a BS/MSEE and at least 5 years' experience.

Microprocessor Development

MICROPROCESSOR ARCHITECT

Design for a complex VLSI processor chip set. Responsibilities include intitial top level chip partitioning, chip microarchtecture, microcode development, performance modeling, behavioral modeling and specification development of the system interface and the memory subsystem(s). Requires a BSEE/CS and 5 years' VLSI processor design experience.

VERIFICATION/TESTABILITY ARCHITECT

Develop total verificaiton strategy for design, finihed chips and complete system(s). Implement this strategy with and through the design team for a very complex fourth-generation microprocessor system. Requires a BS/MSEE and at least 5 years' experience.

RELIABILITY/ RECOVERABILITY SPECIALIST

Develop models of, and a strategy for, chip and system error detection and correction for a fourth-generation microprocessor system. Analyze data about chip and system failures, develop a chip and system error detection and recovery strategy and implement this strategy with and through the design team. Requires a BS/MSEE and at least 5 years' experience.

For the preceding positions, please contact: Gary V. Schipani, Department 0401 8825, Digital Equipment Corporation, 77 Reed Road, HL02-2/K12, Huson, MA 01749-2895.

SBO Engineering

SENIOR ENGINEERING MANAGER

Responsible for technical strategy and operations to manage vendor, technology and device selection, specification, qualification and production support for standard VLSI business. Requires BS in engineering or computer science and MBA with at least 10 years' technical and business operations experience in computer/ semiconductor industry.

PRODUCT ENGINEERING SUPERVISOR

Provide technical leadership of a group responsible for the technical aspects of non-volatile memory devices (PROM, EPROM, EEPROM). Interface with IC vendors, design engineering, and manufacturing. Requires BS in physics or equivalent and at least 5 years' semiconductor technology with supervisory experience.

For the preceding positons, please contact: Jerry Temple, Department 0401 8825, Digital Equipment Corporation, 111 Locke Drive, Box 1006, Marlboro, MA 01752-9106.

To cash in on these blue chip opportunities with a $9.4-billion computer company, please contact the appropriate person listed above.

*Trademark of Digital Equipment Corporation.

We are an affirmative action employer.

digital™

World-Class Commitment

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

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

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

We're committed to delivering the best in simulation.

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

Product Marketing Managers
Software Development
Technical Leaders
Software Development Engineers

- System Simulation
- Logic/Fault Simulation
- Circuit Simulation
- DFT/BIT/ATG
- Logic Synthesis
- Hardware Description Languages

Everybody needs a goal.

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

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

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

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

Mentor Graphics®

ADVERTISERS' INDEX

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

VLSI Systems Design

PUBLISHER
NORM ROSEN
DIRECTOR OF MARKETING
TIM L. TOBECK

DISPLAY ADVERTISING
JONNO WELLS, District Manager
1054 Saratoga-Sunnyvale Road
San Jose, CA 95129
(408) 252-6191
KAREN TALLARIDA
7 Wells Avenue, Newton, MA 02159
(617) 244-5333
WALTER OLSON
DAVID JANOFF
600 Community Drive, Manhasset, NY 11030
(516) 365-4600
CLAIRE FLORA
5151 Beltline Road, Dallas, TX 75240
(214) 661-5673
WILLIAM HIGGENS, III
MARY SHUTACK
225 N. Michigan Ave, Chicago, IL 60601
(312) 565-5943
RICHARD CARLISLE
4500 SW Kruse Way, Lake Oswego, OR 97035
(503) 636-7694
KENNETH D. BEACH
209 Arlene Court, Newark, DE 19702
(302) 368-3116

INTERNATIONAL ADVERTISING
DANIEL H. LEEDS, Vice President
(516) 365-4600
STEVE DRACE, International Sales Manager
FAX 852-5-861-0668
PacificMedia, Inc.
Fax: Tokyo 252-2780 Telephone: 256-8456
Telex: J25472
Young Media, Inc., Seoul, Korea
Telephone: 756-4819 Fax: (02) 757-5789
Mr. Owen Wang
Ace Marketing, Inc.
Telephone: 751-3636 Telex: 78514142

CLASSIFIED ADVERTISING
NEW YORK (516) 365-4600
SANDRA HERMAN/National Manager
MICHAEL ZERNER, East Coast District Mgr.
NEW ENGLAND N.Y. N. Central District Mgr.
KATHY HEALY, Southeast, South Central Space Sales
SONJA WONG, Adverstising Coordinator
FELICIA SABATINI, Graphics
SAN JOSE (408) 252-6191
LAURIE MUNCE, West Coast District Mgr.
DIRECT MARKETING SERV.
JOANNA BRANDI, Manager
PETER CANDITO, Sales Manager
ARLENE BROWN, List Sales Representative
(516) 365-4600
MARKETING SERVICES
DORIANNE WALTHER-LEE/Director
CUSTOMER SERVICE/DISPLAY ADVERTISING
LYNDIANE HERVEY, Group Manager
TRACY RENK, Ad Coordinator
(516) 365-4600
PROMOTION
GENEVIEVE HIGGINS
CIRCULATION
ROBERT ANDERSEN, Circulation Manager
RESEARCH
DANIEL R. CAMPBELL, Electronics Group Manager

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

CMP PUBLICATIONS, INC.
MICHAEL S. LEEDS, President
PEARL TURNER, Vice President/Treasurer
DANIEL H. LEEDS, Vice President

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

NANCY GRABOW
Assistant to the President

BPA

Metal gate CMOS yesterday. Silicon-On-Insulator tomorrow.

Our non-stop development of rad-hard technology is guaranteed.

American ICs for American strength...that's been our commitment since 1965. In the Polaris program. In Poseidon. In B-1B, Trident, MX/Peacekeeper, SICBM, SDI, and more. Today, in fact, we're the number-one supplier of military rad-hard integrated circuits.

Some IC companies have uncertain futures. We don't. We'll be here. American-owned and American-run. Providing products and product accountability throughout the lifetime of your systems.

And developing tough new products to keep your systems competitive...like our HS-6504RH rad-hard 4K RAM, that's produced in our MIL-M-38510, Class S certified facility.

We wrote the book on your future needs. Harris' Rad-Hard/Hi-Rel Data Book: hundreds of pages of information on what's here now and what's coming next. Ask for a copy.

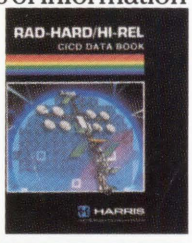

Phone Harris Semiconductor Custom Integrated Circuits Division. In the U.S.: 1-800-4-HARRIS, Ext. 1908, or (305) 724-7418. In Canada: 1-800-344-2444, Ext. 1908.

IN RAD-HARD/HI-REL, THE NAME IS HARRIS

Harris Semiconductor: Analog - CMOS Digital
Gallium Arsenide - Semicustom - Custom

HARRIS

"So Harris' rad-hard data book is Class-A?"

"Yeah, and their rad-hard IC facilities are MIL-M-38510, Class S certified."

CIRCLE NUMBER 42