EIGHT DOLLARS A CMP PUBLICATION DECEMBER 1988

VLSI Systems Design

FOR DESIGNERS OF HIGH-PERFORMANCE SYSTEMS

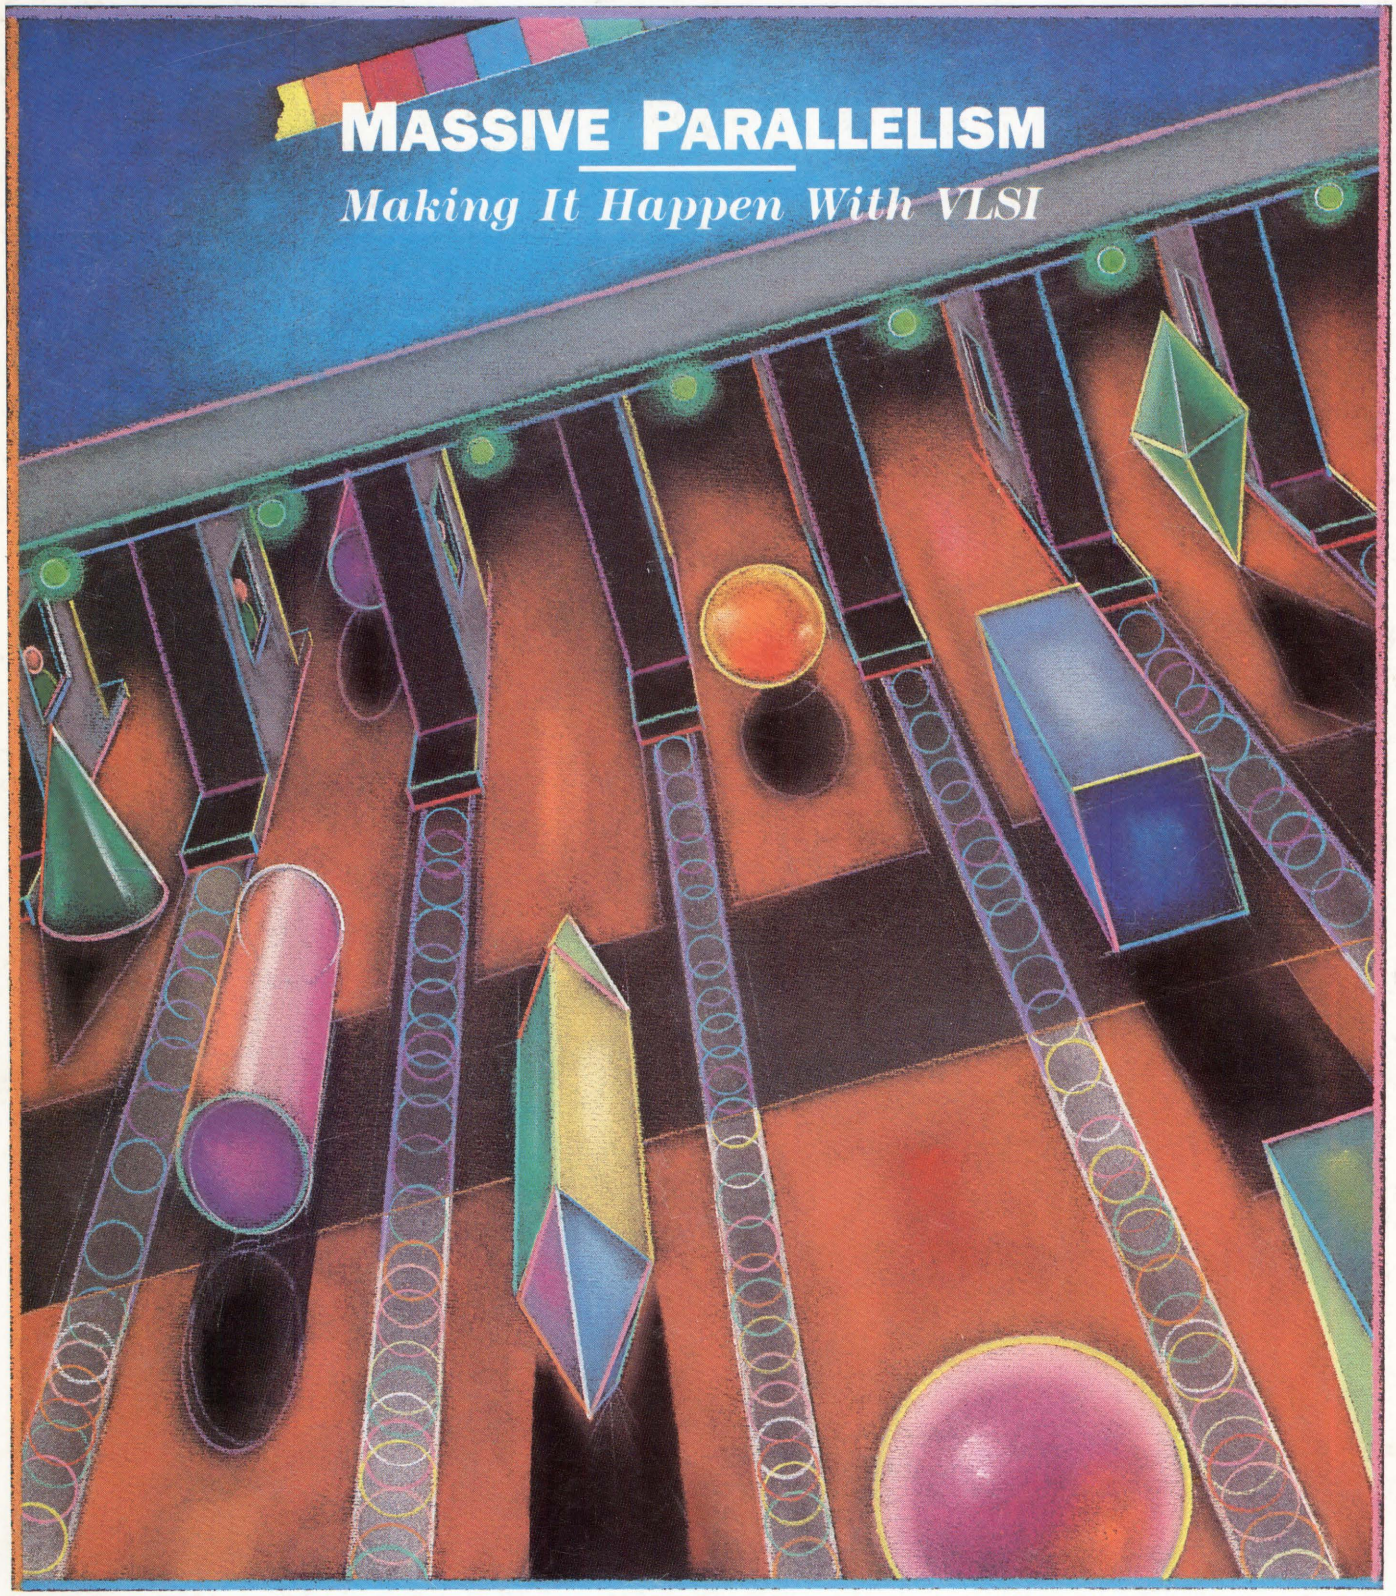

MASSIVE PARALLELISM

Making It Happen With VLSI

Only an ASIC partner lets

you choose your own tools.

It's no secret that the faster you can design your ASICs, the faster you can get your products to market. Which is why Fujitsu lets you decide which design tools to use. Then we provide you with all the on-site support you need to breeze through not only product development but also product verification.

As your ASIC partner, we've got what it takes to help you turn a good idea into a great new product.

If you choose a popular engineering workstation, for example, we'll support you all the way. With libraries, convertors, test vector editor and a powerful, menu-driven file-management program called FAME (Fujitsu ASIC Management Environment) that boosts productivity by 300%.

If, on the other hand, you're using computers running third-party software, we'll support you there too. Our support goes even further with ViewCAD.™ A complete set of easy-to-use tools that gives you a choice. Use ViewCAD for on-site verification or as a complete, stand-alone ASIC engineering workstation.

ViewCAD supports your entire design cycle. From schematic capture and logic design rule checking to test data entry, simulation, analysis and data conversion. All running on your favorite industry-standard UNIX®-based platforms. And all from Fujitsu.

ViewCAD puts years of ASIC experience under your belt. Experience that has led to more than 10,000 successful designs. And made Fujitsu the world's leading ASIC supplier.

ViewCAD's design methodology helps you shorten your product development cycle. Because it catches more design errors earlier in the process. Its powerful design capabilities and X Windows™ software lets you easily handle the overwhelming complexities of high-density ICs. Even our 100,000 gates and beyond.

With ViewCAD, you get the benefit of a powerful simulator that provides you with virtual mainframe compatibility. Including a simulation-to-silicon correlation of over 99.9%.

What's more, you can integrate ViewCAD with third-party tools for interactive design verification.

Which ensures the integrity of your design and reduces design verification and modification time. So your circuit is closer to silicon than ever before. And you're closer to market.

But supporting your design tool decision is only part of the story. As your ASIC partner, we expand your design team, providing you with fully staffed and equipped technical resource centers, coast to coast.

Each able to train your engineers and provide twenty-four hour design suites, so you can work whenever inspiration strikes. You also get ASIC sales and marketing support to help you smooth out all the administrative wrinkles.

So no matter what your decision, you can count on Fujitsu to support it. Which, after all, is everything an ASIC partner should be.

FUJITSU

FUJITSU MICROELECTRONICS, INC.

Integrated Circuits Division

3545 North First St., San Jose, CA 95134-1804. (800) 642-7616

Everything an ASIC partner should be.

CIRCLE NUMBER 1

Contents

ARTICLES

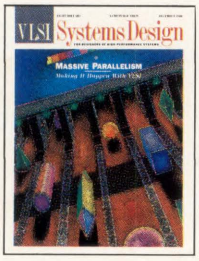

Massive parallelism can be a key to higher performance, and VLSI chip technology can facilitate the path to success

DEPARTMENTS

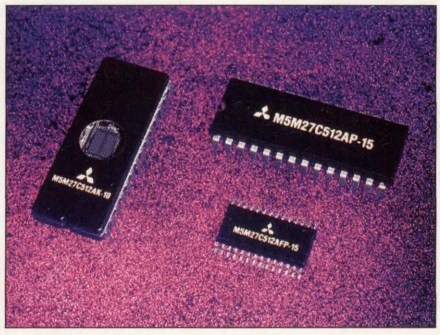

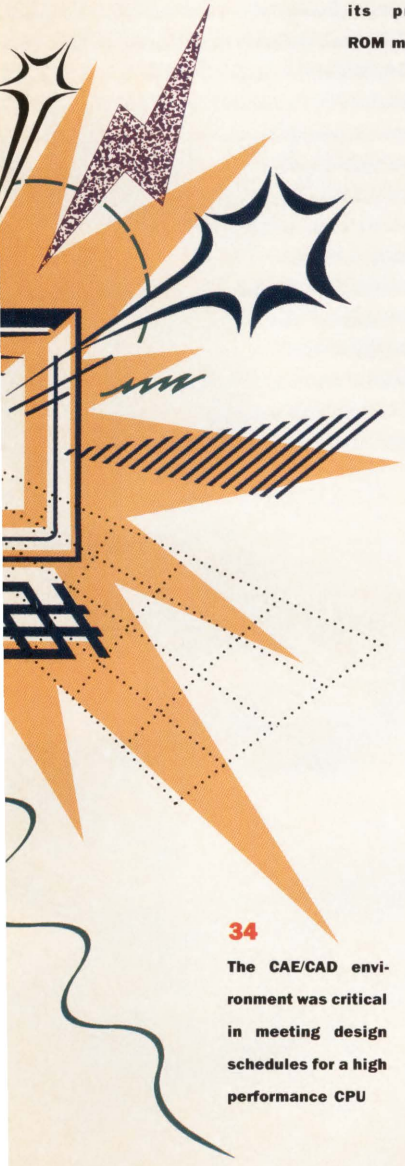

COVER ILLUSTRATION BY ANDRZEJ DUDZINSKI

VLSI Systems Design

FOR DESIGNERS OF HIGH PERFORMANCE SYSTEMS

A CMP Publication

VLSI SYSTEMS DESIGN (ISSN 0279-2834) is published monthly with an extra issue in September by CMP Publications, Inc., 600 Community Drive, Manhasset, NY 11030. (516) 562-5000. VLSI SYSTEMS DESIGN is free to qualified subscribers. One year subscriptions to others; US and Canada $60. Mexico, Europe, Central and South America: $120.00. Asia, Australia, and Africa: $135.00. Second-class postage paid at Manhasset, NY and additional mailing offices. POSTMASTER: Send address changes to VLSI SYSTEMS DESIGN, Box No. 2060, Manhasset, NY 11030.

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

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

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

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

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

Daisy's standard now runs on an

Introducing the Advansys Series™ of high performance CAE workstations.

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

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

The Sun 386*i*.™

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

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

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

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

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

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

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

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

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

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

daisy SYSTEMS CORP.

# of excellence
excellent standard.

ASICs: The Migration to Mainstream Accelerates

ASICs will be the foundation for tomorrow's high performance products

Only five years ago, when ASICs were "the new kid on the block", one of the system engineer's biggest challenges was to convert his/her board design to one or more ASICs. The tools available for designing ASICs were relatively crude and required that the user have a lot of IC design expertise in order to operate them efficiently. The design-to-prototype cycle was long, first pass success was a big gamble, and there was a strong possibility of missing an important market window.

But the rewards were great. Converting a board full of standard logic chips to a few ASICs produced systems with shorter data paths and higher system speeds. The uniformity and reliability of the integrated circuits were much higher than most printed circuit board assemblies. The custom circuits were more gate-efficient than designs using standard TTL. The physical size of the systems were considerably smaller. In addition, the overall cost was usually lower. However, since the high performance of these products depended directly on the ASICs, the systems engineers tended to concentrate their efforts on the ASICs (and their design) rather than on the system (and its design).

Today, ASICs are even more important. They can be the key to keeping a company's products competitive. Now, however, the systems engineer's biggest challenge isn't the ASIC and its design, but rather in how ASICs can best be leveraged to deliver even higher performance systems and products. They also still look at ASICs as a way to get a proprietary "leg up"on the competition. Perhaps the biggest testimonial to this growing trend can be seen in the literature describing today's latest and hottest new products. Many times, more ink is devoted to the number of ASICs and their gate counts than to the system features and specifications.

It may sound like ASICs are used mainly as a promotional gimmick. Maybe it's true in a few rare cases, but ASICs have proven their worth in many product applications that have ranged from medical (pacemakers) to musical (CD players). ASICs have accelerated into the mainstream where they will continue to deliver the performance and functionality that is a must for the next generation of high performance systems and products.

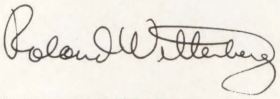

ROLAND WITTENBERG
EDITOR-IN-CHIEF

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

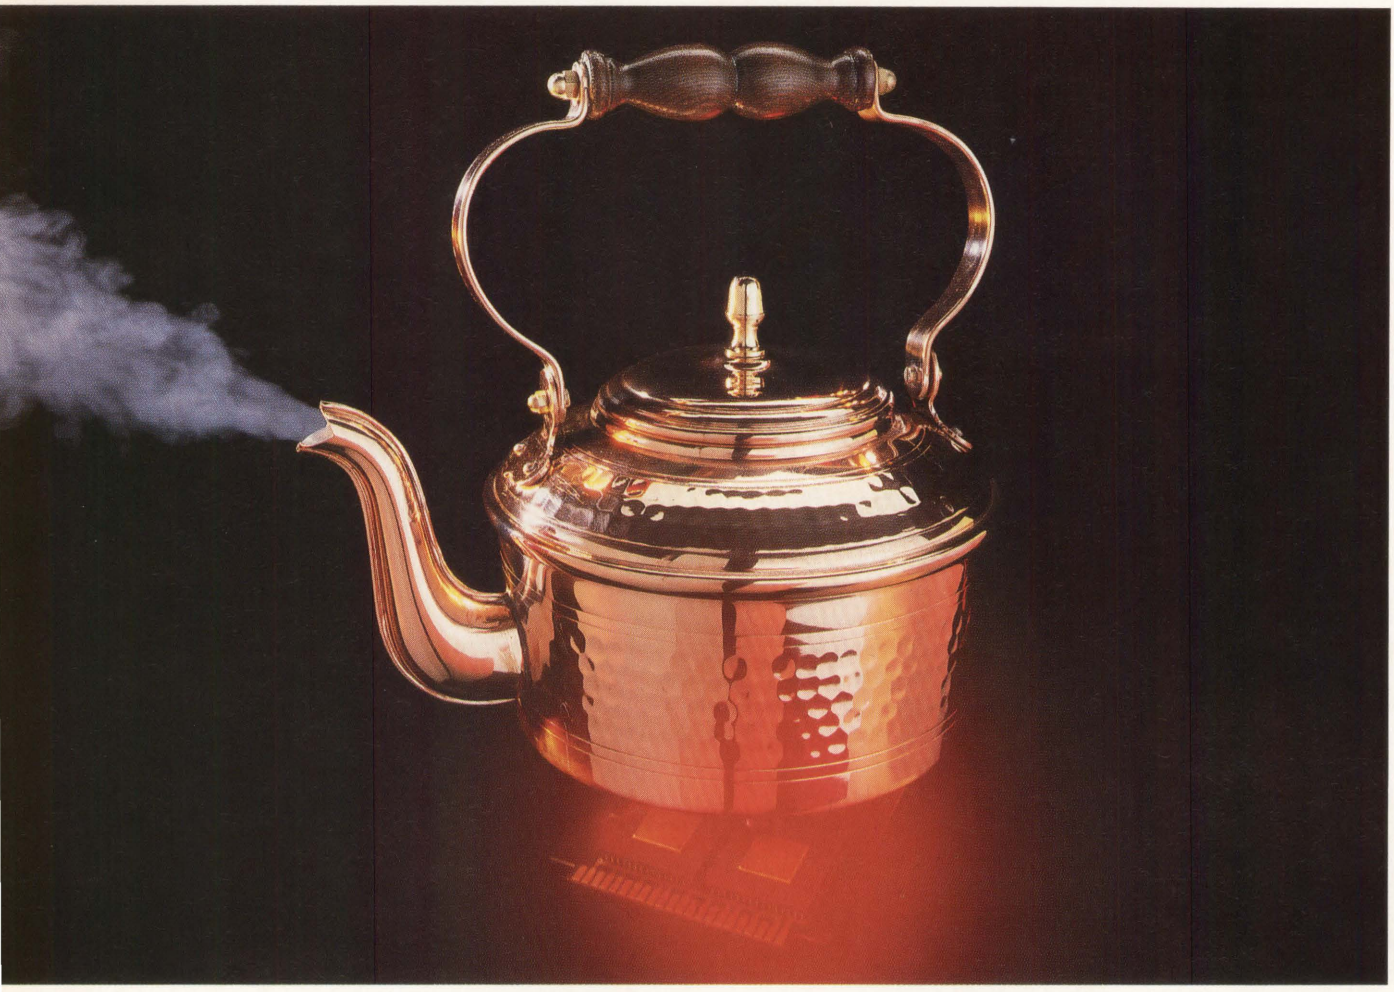

Cold facts: now the highest-density ECL logic array runs at a cool 1/10 the gate power of competing devices.

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 2

Calendar

IEEE International Conference on Wafer-Scale Integration

January 3-5, 1989
Fairmont Hotel
San Francisco, Calif.

This conference will present a balanced program of all the aspects of monolithic wafer-scale integration, including theory, technology, applications, and products. The program will feature contributed papers, poster presentations, and panel discussions and will cover topics such as WSI reliability, yield modeling, wafer-scale CAD systems, packaging, power/ground distribution, signal and image processors, and wafer-scale memory. For further information, contact Patty Patterson, TRW Defense Systems Group, 1 Space Park (R2/2076), Redondo Beach, Calif. 90278. (213) 812-0788. ■

IEEE 1989 Aerospace Applications Conference

February 12–17, 1989
Breckenridge, Colo.

Sponsored by the South Bay Harbor Section of the IEEE, the emphasis of this conference will be on applications, present and future. Sessions will cover such topics as system concepts, computer and microcomputer applications, system management, millimeter and microwave technology, communications and telemetry, software methodology, electro-optic applications, VLSI and semiconductor technology, instrumentation and measurement, graphics and display systems, energy and space applications, aerospace manufacturing and testing, and small aperature terminals. Additional information about the conference may be obtained by contacting Harvey Endler, Registration Chairman, 15137 Gilmore St., Van Nuys, Calif. 91411. ■

1989 International Symposium on VLSI Technology, Systems, and Applications

May 17–19, 1989
Taipei, Taiwan, R.O.C

Papers are being solicited for presentation at the 1989 International Symposium on VLSI Technology, Systems, and Applications. Topics will include modeling and simulation, materials and processing; logic, memory, and analog ICs; design tools, custom VLSI and gate arrays; personal computers, microprocessors, fault tolerance, design for testability, CAD/CAM, automation and robotics, workstations, signal and image processing, software and expert systems, and computer peripherals. Interested authors should submit, by January 8, 1989, a 35–50-word abstract and 20 copies of a 300–600-word summary with supporting figures to: Dr. John Y. Chen, Technical Program Chairman, Boeing Electronics, High Technology Center, P.O. Box 24969, MS 7J-56, Seattle, Wash. 98124. ■

International Test Conference 1989

August 29–31, 1989
The Sheraton Washington Hotel
Washington, D.C.

Sponsored by the IEEE's Computer Society and Philadelphia Section, the ITC provides a major forum for the exchange of information about the testing of electronic devices, assemblies, and systems. This year's conference focuses on innovative test techniques and equipment needed to meet the challenges of the future. Technical presentation topics will include built-in-self-test, computer-aided engineering, design for testability, design verification, fault modeling and simulation, memory devices, microcontrollers and microprocessors, printed circuit boards, surface mount assemblies, system test, wafer-scale assemblies, quality and reliability, standards, and test economics. Authors are invited to submit, by January 16, 1989, a 35-word abstract, and either a 500-word summary or a full manuscript to Ray Mercer, Program Chair, International Test Conference, Millbrook Plaza, Suite 104D, P.O. Box 264, Mount Freedom, N.J. 07970. For more details, call Doris Thomas, at (201) 895-5260. ■

International Conference on Semiconductor and Integrated Circuit Technology

October 22–28, 1989
Beijing, China

Designed to provide an international forum on semiconductor and integrated circuit technology, this conference will cover such topics as amorphous silicon, bipolar technology, CAD, CMOS technology, dielectrics, electrical characterization, fab safety and maintenance, IC design, interconnect technology, multilevel interconnect, packaging, process characterization, rapid thermal processing, and reliability/yield. By January 9, 1989, interested authors should submit a 300-word abstract detailing the work to be presented. To submit abstracts or to obtain additional information contact Linda Reid, Continuing Education in Engineering, University Extension, University of California, 2223 Fulton St., Berkeley, Calif. 94720. ■

Illustration by Lorraine Raywood

GAZELLE'S NEW GA22V10 LOGIC DEVICES

110 MHz @ 34¢/MHz

Breakneck Performance at Breakthrough Prices

Anyone designing a 25 MHz to 40 MHz microcomputer is familiar with the three trade-offs of support logic: speed, affordability and availability.

Until now, nobody's ever solved all three at once.

The problem is that support logic has to run twice as fast as the CPU. Or the whole system slows down by a full third. Or more.

So a 33 MHz CPU needs a 66 MHz 22V10. But that doesn't exist in silicon. And CPUs just keep getting faster.

Performance	GA22V10-7	GA22V10-10
t_{PD}	7.5 ns	10.0 ns
t_s	3.0 ns	3.6 ns
t_{co}	6.0 ns	7.5 ns
f_{MAX}	110 MHz	90 MHz
Volume Price	$37.00	$31.00

Problem solved: Gazelle's new 110 MHz GA22V10s. At great prices. In quantity.

The result?

Swifter 68030s. Extraordinary 80386s. Accelerated 88000s. Full-speed SPARC.™ Red-hot RISC. All priced competitively.

The reason?

Gazelle's GA22V10s are TTL-compatible GaAs—100% pin and function compatible with silicon. At just 34¢/MHz, they deliver the price advantage of silicon 22V10s. But at 110 MHz, they're more than twice as fast. And fast enough to outrun the fastest CPUs.

But inexpensive enough to outsmart the toughest competition.

gazelle

(408) 982-0900

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

CIRCLE NUMBER 3

BIT Stream

Retargetable C Compiler and Assembler

Step Engineering has augmented its Metastep microprogram language and added a front-end C compiler, allowing designers to create C compilers for custom architectures. The C compiler performs machine-independent optimizations such as reuse of expressions in the body of the code, strength reduction, common subexpression elimination, and loop unrolling. The output is then consumed by the Metastep language system, which the designer configures for his architecture by using augmented Metastep macros. So, designers can program custom systems using C, Meta-step assembly-level language, and microcoding. The Metastep Microprogram C Compiler runs on MS-DOS computers ($4,995), Sun Microsystem workstations (from $9,995), and VAX computers (from $19,995). ■

Software Development Support for DSP Chips

A recent announcement by Texas Instruments Inc. (Dallas, Texas) brings extensive software support, at a level normally associated with microprocessors, to TI's 32032 digital signal processor. The SPOX real-time operating system from Spectron Microsystems (Santa Barbara, Calif.) simplifies software development for 32032-based systems and, by providing a 32032-resident operating environment, eliminates the need for a general-purpose CPU beside the 32032. SPOX is included in TI's XDS-1000 development package ($16,000), which includes an emulator, a C compiler/assembler/linker, and a development board. ■

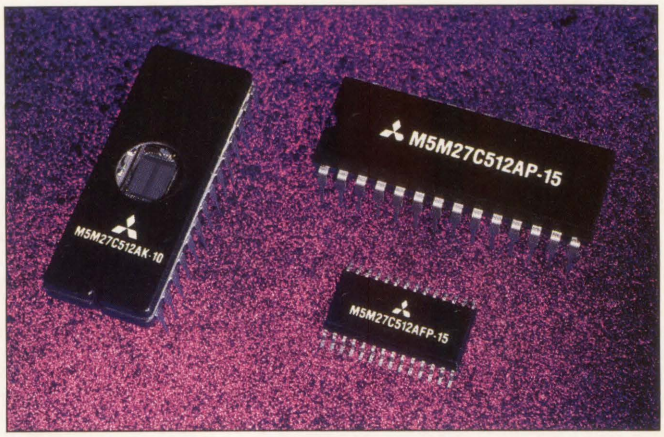

More Memory from Mitsubishi

A new 512-kbit programmable ROM memory chip has been added to the low-power CMOS memory line offered by the Semiconductor Division of Mitsubishi Electronics America Inc. The high-speed, 100-ns 512-kbit UV EPROM is available in a 28-pin, 600-mil Cerdip package (M5M27C512AK-10) that is tagged at $22.60 in 100-unit quantities. The chip, which is targeted at embedded-control microprocessor memory applications, is organized 64 K x 8 bits. It features TTL-compatible inputs and outputs in both the read and program modes. Since the configuration is interchangeable with NMOS 512-kbit EPROMs, it allows easy upgrading from 64-kbit, 128-kbit and 256-kbit EPROMs. ■

One-Board 68030-Based Computer for VME Systems

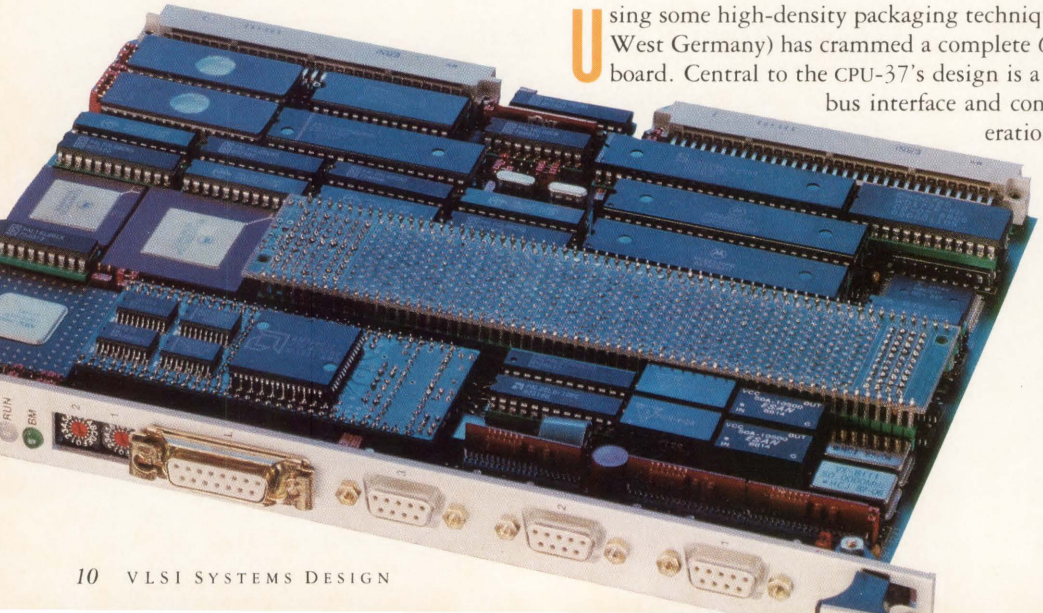

Using some high-density packaging techniques, Force Computers Inc. (Ottobrun, West Germany) has crammed a complete 68030 computer onto a single VMEbus board. Central to the CPU-37's design is a 135-pin gate array that provides VME bus interface and control functions, including DSACK generation, bus error generation, system reset, bus clock, and all on-board control logic. Given this chip and some daughterboards, the CPU-37 can accommodate a 16.7-MHz or 25-MHz 68030 and 68882, up to 4-Mbytes of RAM, three serial ports, an SCSI interface, a floppy-drive controller, and an Ethernet transceiver interface. Precision surface-mount technology is the shoehorn for putting all this capability on one board. A bare-bones board is priced at $3,990, and the fully configured computer costs $5,890. ■

Update Pampers PADS-PCB

CAD Software Inc. (Littleton, Mass.) has rolled out the latest release of its printed-circuit-board layout system, PADS-PCB Version 3.0. The new release has added several features to the previous version, such as: enhanced design rule checking; an improved surface-mount design system; networking capabilitiesthat allow users to share a common library; and added support for metric units. Version 3.0 also features "group operations" that allow users to define groups of components, including connections and routes. These groups can then be handled as a single component and moved, rotated, copied, stepped and repeated, placed on the reverse side of the board, and even stored on a disk for future use.

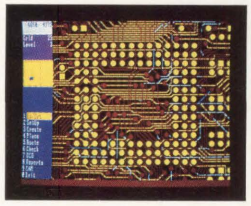

Two new options were also announced for the new release of PADS-PCB. These options include an autoplacement package and a graphics package that supports high-resolution graphics cards from Number 9 Computer Inc. (Cambridge, Mass.). The PADS-PCB Version 3.0 is priced at $975, while the autoplacement option (PADS-PLACE) and high-resolution option (PADS-HI-RES) run $350 and $495 respectively. ■

Analyzers Scale New Lows

Hewlett-Packard Co. has unveiled two new scalar network analyzers. The HP8757E is priced from $7,500, while the HP8757C, which includes additional features and a color display, is tagged at $9,000.

The 8757E includes: three detector inputs with choice of ac or dc detection; a 76-dB dynamic range with ac detection; two display channels; an internal plotter/printer buffer that allows hardcopy output simultaneously with testing; and fully annotated displays including trace cursors and min/max search functions.

The 8757C has all the features and performance of the 8757E, but it also provides: a four channel color display; a limit line test capability that provides on-screen pass/fail indication; disk interface capa-

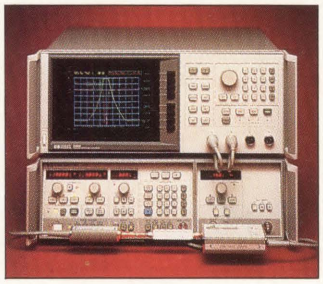

bility that allows external storing and recalling of both test setups and data; adaptive normalization for calibrated measurements on narrow sweep ranges after a wideband calibration; and up to 1,600 measurement points per sweep. ■

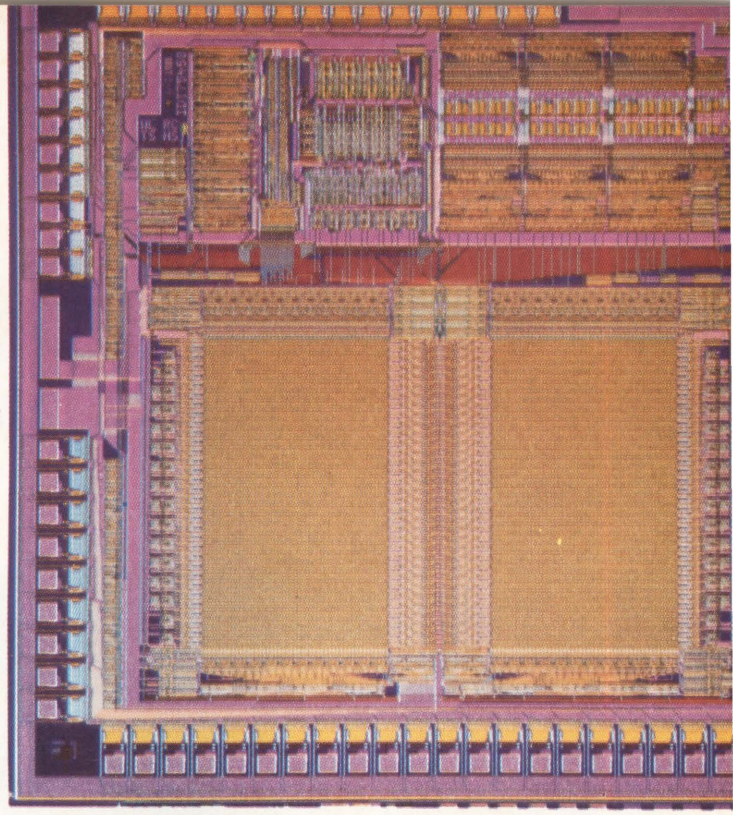

Cranking up Triple Pallette DACs

By developing a triple pallette DAC that runs at 165 MHz, Integrated Device Technology (Santa Clara, Calif.) has removed one of the constraints that kept graphics manufacturers from exceeding the 1,200- X 1,000-pixel resolution standard. The higher clock frequency can refresh screens that have resolution as high as 1,600- X 1,200-pixels. The IDT75C458's higher speed can also allow systems to refresh at 70 Hz rather than 60 Hz, a step that minimizes flickering on the display. Although pin compatible with the BT458 pallette DAC, the part consumes 1 W, half the specification of competing devices. Its accuracy is rated at ½ LSB. The 165-MHz version is tagged at $213; lower-cost versions are available that run at 122 MHz, 110 MHz, and 80 MHz. ■

Harris Buy of GE Semiconductor Sealed

The payment of $206 million before the end of this month by Harris Corp. (Melbourne, Fla.) marks the final scene of the General Electric semiconductor saga. The deal also includes Intersil and most of the RCA solid state business that GE picked up last year. The signing of the agreement was announced by John T. Hartley, chairman and CEO of Harris. The merging of these operations will make Harris the nation's sixth largest semiconductor vendor. Jon Cornell, who headed Harris' Semiconductor Sector before the acquisition, has been given the nod to run the expanded operation. ■

HARTLEY

Compute-bound problems become I/O-bound problems

Justin Rattner Seeks New Horizons

LONG before Justin Rattner started Intel Scientific Computers, he had to choose which college to attend. Unlike other high school students in Los Angeles, he eschewed the nearby California schools and reached out across the country to wintry upstate New York, where resides Cornell University.

Justin was bitten by the computer bug during a summer job at Scientific Data Systems. So he strayed from the traditional EE course study at Cornell to take computer science courses, which were then taught by the Liberal Arts school. Faculty at the EE school advised him that no one would hire him if he wasted so much time in computing science. Justin prevailed, however—he landed one of the coveted entry positions with Hewlett-Packard Co. in 1972.

While programming in HP's minicomputer lab, Justin became intrigued by the new type of ICs coming from Intel Corp.: microprocessors. Attracted to this new method for system design, he jumped to Intel in 1973. Once again he found himself in a novel environment, because HP had been a mature company while he describes Intel as still being "a brash upstart." As part of the small group supporting the 8080 family, he contributed to all portions of the microprocessor support. "Just being involved with the microprocessor was the exciting thing," he remembers.

He was in the right place at the right time when, in 1975, Intel

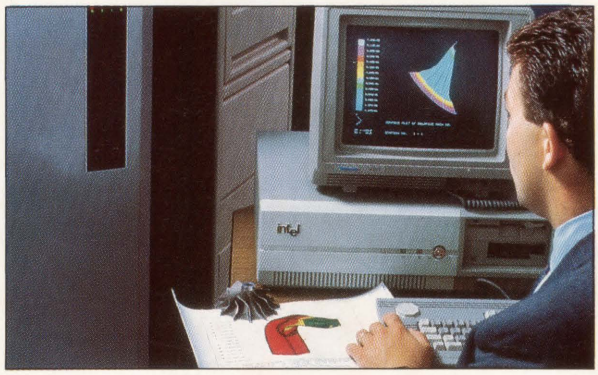

decided to stretch toward new possibilities in integration. The company saw that, with the high levels of integration possible in its scaled NMOS processes, 100,000-transistor chips incorporating entire systems were conceivable. And Justin was chosen to lead a group of fewer than 10 people in an effort that culminated in Intel's 432. Although advanced in design, its low throughput killed it in the marketplace.

It was apparent that Intel would have to spend a great deal of money on the 432 to make it commercially successful. Another system-level project, codenamed Gemini, got the nod instead. The Gemini project pulled Justin on board to oversee the architecture specification of a new computer system that, ironically, incorporated many of the concepts that the 432-team had addressed. The result of the project: the 80960 microprocessor and the BiiN computers.

After specifying the Gemini design, Justin took a sabbatical, in 1983, to plot his next move. He wanted to identify the new horizons for microprocessors, now that the 80386 and 80960 were on the boards. As he explored areas of possible promise, he kept bumping into parallel processing. He became particularly impressed with a project at CalTech that combined 8086s in a parallel architecture. Starting with that idea, Justin founded Intel Scientific Computers (ISC) which, in just 13 months, built a 128-node, 80286-based Intel Parallel Super Computer (iPSC1).

'USING MANY SMALLER DRIVES IS AN

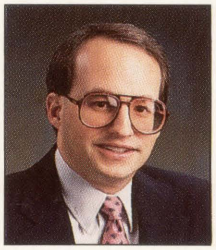

IDEA THAT WILL BECOME STANDARD'

By not waiting for the 32-bit processors to appear, Justin says, "the iPSC1 broke the no parallel computers, no parallel software cycle." Without parallel hardware, who could develop the parallel software that would spur the use of parallel hardware? The second-generation iPSC2, powered by the 80386, has pushed the power of parallel computers to ranges of performance associated with the Cray 1.

After getting ISC off the ground, Justin turned over the managerial reigns to concentrate on technology development. Now, as director of technology (and the fourth Intel Fellow), he is addressing the I/O bandwidth problems of supercomputing. "A supercomputer is a machine that transforms a compute-bound problem into an I/O-bound problem," he points out. ISC recently introduced parallel I/O concepts, built around banks of low-cost disk drives rather than fewer, larger drives; this should loosen the I/O bounds. He feels that using many smaller drives is "an idea that will become a standard architectural feature of future computers."

"It's inexorable that microprocessors will eventually have the power of . . . supercomputers," he says. "I'm excited about what parallel processing systems built with them can do."

—David Smith

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

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

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

FREE. ASIC Estimating Kit.

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

SEATTLE SILICON
Launching the next generation of ASICs.

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

CIRCLE NUMBER 4

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.

Semiconductor foundries have dropped the ASIC ball

Next-Generation ASIC Suppliers Must Have Tool And Service Tech Base

MANY large semiconductor firms are failing to master a new set of skills required for success with application-specific ICs. A few companies have already backed out of the business, others deemphasized it. That leaves the door open to a new generation of ASIC suppliers with keystone technology in process-independent tools and a responsive service infrastructure.

The main contribution of merchant semiconductor suppliers historically has been manufacturing muscle. In such a business, the company's energy is directed inward—process development and fabrication operations typically absorb 80 to 90 percent of management's energy. While beginning with a different emphasis, the first-generation ASIC suppliers drifted to a very similar business model based on a captive manufacturing capability.

But for most ASIC users, especially those who deal with highly complex designs, the needs have become very different. These users find themselves limited by design tools that cannot handle complex designs; by production services that are geared toward four-year project schedules and 1-million-unit production volumes; and by designs that lock them into a single manufacturing source.

Some of the first-generation ASIC suppliers have tried to overcome the captive fabrication business model. But the first-generation ASIC suppliers are still trapped by captive foundries and are looking more and more like the traditional chip makers. They are forced to develop new processes and continue huge investments. Once the investment is made, they frantically redesign libraries for each new process. As a result, alliances disappear almost as rapidly as they are announced. The foundries' design tools inevitably steer people toward the foundries' own process. And production minimums grow higher to make sure the facility is full. Recently, some first-generation ASIC suppliers announced their intention to focus support on a limited number of high-volume customers.

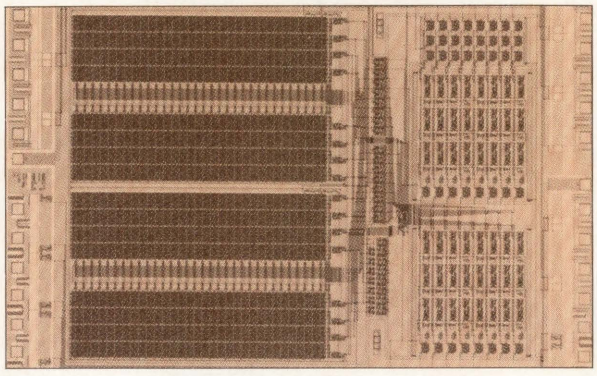

'FIRST-GENERATION ASIC SUPPLIERS

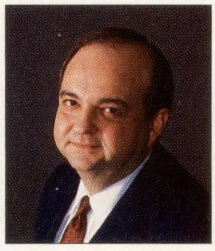

ARE STILL TRAPPED BY THEIR CAPTIVE FOUNDRIES'

It is logical, then, that market pressures are creating a new generation of ASIC suppliers with a business based on a different model. This model offers access to manufacturing advances rather than captive ownership of them.

The technology base of these companies begins in integrated design tools that provide a link between the designer and manufacturing in a foundry-independent environment. This tool-oriented technology is essential for two reasons.

First, only when tool design is made a fundamental priority can the tools provide the level of design capability needed to create high-quality, high-complexity designs on schedule. For the designer of high-end ASICs, therefore, design tools should include physical as well as logical design. They should also provide a path to testable circuits and integrate a combination of approaches, including standard cells, logic synthesis, and compilers.

Second, by creating tools that can retarget designs for many processes, suppliers can finally break the dominance of the captive fabrication facility. That means letting designers choose objectively from among several processes—a prospect that is impossible when the ASIC supplier is concerned with keeping the fab full.

Balancing the emphasis on effective tools, the next-generation ASIC suppliers are also being built around a process-independent production service component.

They offer project-oriented support that welcomes lower production volumes and can help a designer get working devices in any one of several manufacturing processes. A service infrastructure geared toward process-independent support means that every design has an easy path to multiple sources. ■

AMAURY PIEDRA *is president and CEO of Seattle Silicon Corp. (Bellevue, Wash.). Previously, he was vice president of Strategic Alliances at Fairchild Semiconductor Corp., head of Fairchild's linear division, and managed Zilog's international operations.*

PERFORMANCE
means 15ns wide word SCRAMs

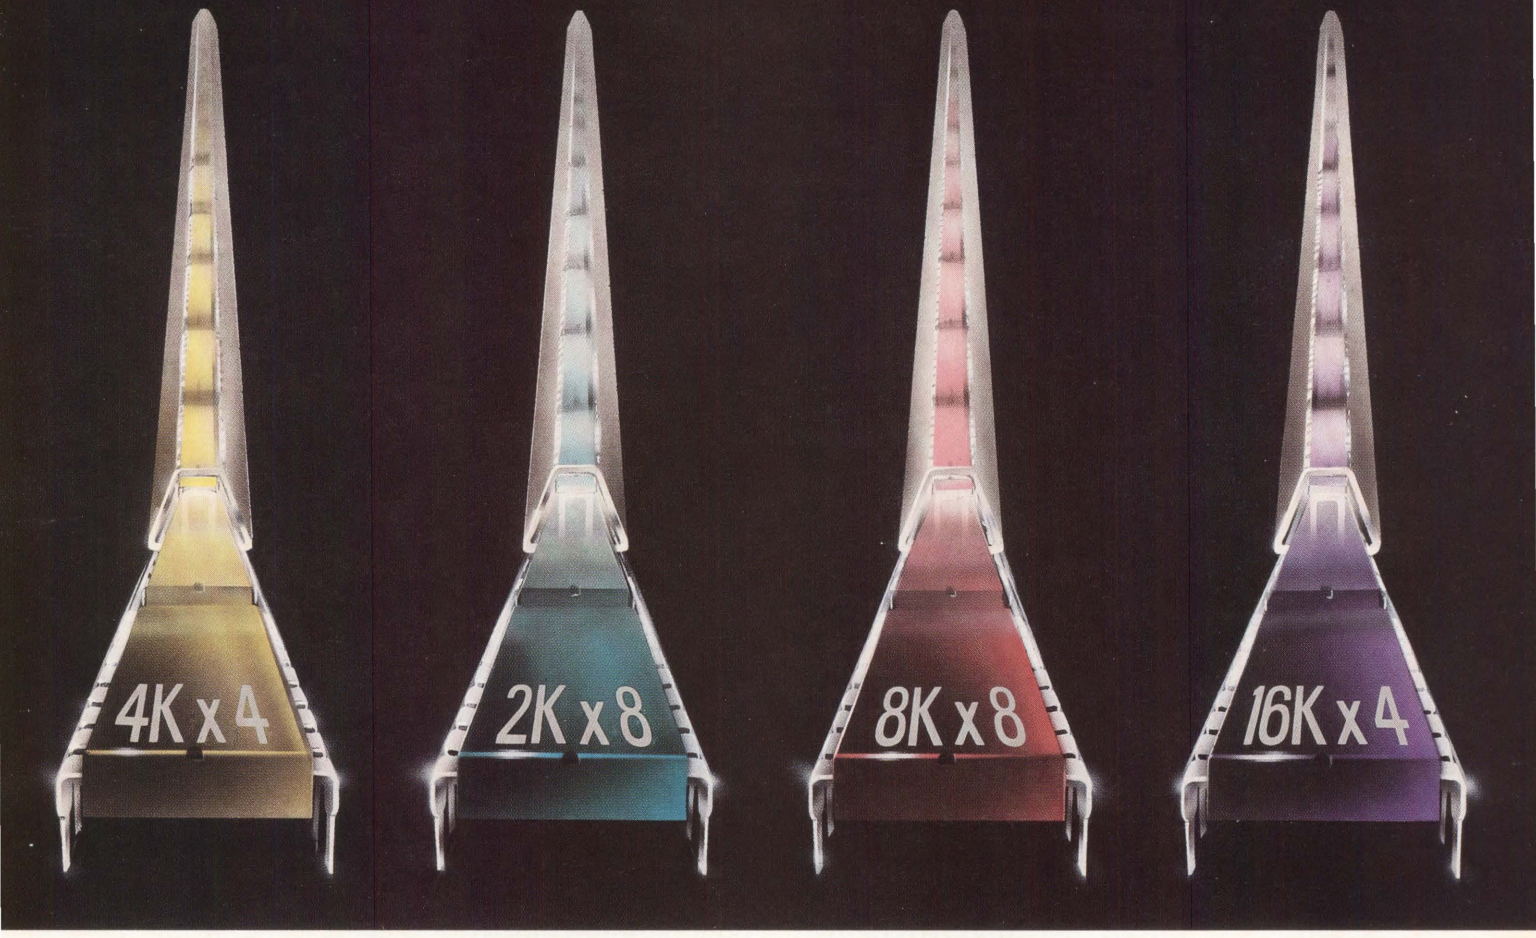

15ns

SCRAMs are Static CMOS Random Access Memories from Performance Semiconductor. At 15ns address access time these 64K's and 16K's are the world's fastest. SCRAMs are manufactured in Performance's six inch class 1 fabrication facility using PACE II 0.7 micron gate length technology which has set the standard for memory speed.

There is immediate availability of 15ns 64K and 16K bit SCRAMs compatible with JEDEC standard pinouts. Also available from stock are 17, 20 and 25ns speed versions.

15NS SCRAM PRODUCT GUIDE

PART	CONFIG.	SPEED	AVAIL.
P4C164	8K x 8	15ns	NOW
P4C188	16K x 4	15ns	NOW
P4C198	16K x 4	15ns	NOW
P4C198A	16K x 4	15ns	NOW
P4C1982	16K x 4	15ns	NOW
P4C1981	16K x 4	15ns	NOW
P4C168	4K x 4	15ns	NOW
P4C169	4K x 4	15ns	NOW
P4C170	4K x 4	15ns	NOW
P4C1682	4K x 4	15ns	NOW
P4C1681	4K x 4	15ns	NOW
P4C116	2K x 8	15ns	NOW

FAST, COOL, & AFFORDABLE

For further information or to order 15ns SCRAMs call or write:
Performance Semiconductor
610 E. Weddell Drive,
Sunnyvale, CA 94089
Telephone: 408 734-9000

CIRCLE NUMBER 5

PERFORMANCE SEMICONDUCTOR CORPORATION

MATRIX CRUNCHING WITH MASSIVE PARALLELISM

BOB CUSHMAN, SENIOR EDITOR

Massively parallel architectures are just about the only way to get past the Von Neumann bottleneck. That's the thinking of many computer experts, not the least of them Gordon Bell, known as the father of Digital Equipment Corp.'s VAX line. The bad news is that many algorithms can't be parallelized. But, fortunately, the matrix operations that are at the heart of so many modern scientific and engineering analyses are inherently able to be parallelized. This route to increased parallelism is now being explored in several ways. Roughly speaking, they can be classified by their level of ambitiousness:

NEW PARALLEL ARCHITECTURES BOOST PERFORMANCE

- Modest attempts that increase the word width of a conventional, Von Neumann, sequential computer and add more execution units;
- moderate attempts that operate Von Neumann machines in loosely coupled, multiprocessor networks communicating via nearest-neighbor or hypercube message passing; and
- all-out extremes that involve non-Von Neumann massively parallel architectures.

The non-Von Neumann massively parallel architectures are of special interest and warrant a closer look. In particular, those that start with memory and try to mix in ALU functions in a fine-grain manner. The current progress towards economical VLSI chips with hundreds of thousands of transistors is making some of these schemes more practical.

One example of this integrated memory-and-ALU approach comes from Steven Morton, who has founded Oxford Computer, in Oxford, Conn., to promote his family of "intelligent memories." Morton's methodology, such as the partitioning of the matrix multiplication algorithm, offers some insights into how advances in VLSI might be exploited for the rapid, economical execution of matrix math. They provide a helpful foundation for understanding some of the other, competing approaches to massive parallelism. These competing approaches, which cover both matrix and non-numerical parallelizable functions, will be covered in future articles. (Meanwhile, for some historical background on massive parallelism, see References at the end of this article.)

Figure 1 shows a basic version of Morton's proposed architecture. It's a configuration offered for handling large matrices. The architecture uses multiple convolution-type intelligent memory chips plus a few other auxiliary and support chips.

The intelligent memory chips have two SRAMs (or DRAMs) that hold the M and V matrix values that are to be multiplied together for the [M] x [V] = [R] operation. Below these M and V memories are a sufficient number of logic devices to perform the multiply-accumulates that gener-

ate the partial products of the matrix multiply.

The bottom of the figure shows a less complex chip that Morton uses to sum up the partial products from the intelligent memories and produce the result elements for the R matrix. Not shown is a control chip that directs the operations of these chips, so that, to the outside world, complex matrix operations are being performed at a high level. This is similar to what happens in mathematics software packages such as Matlab, from MathWorks Inc., in South Natick, Mass. The control chip could be a microprogram sequencer or a controller-type microprocessor.

Multiple copies of the intelligent memory chips are needed because each chip only stores one bit of the M-matrix words. Therefore, there have to be as many intelligent-memory chips as there are bits of precision in the M-matrix values. Meanwhile, the V-matrix values are repeated redundantly in each chip. What Figure 1 does not show clearly is that the intelligent-memory chips must be clocked through as many cycles as there are bits of V-word precision.

The bitwise spreading out of M words and the timewise cycling of V words are the result of two design objectives. One was to use the full width available internally in conventional memory structures. Typically, the actual memory structure is square and the wide words obtained from the square arrays are "wastefully" narrowed by addressing to produce much slimmer output words.

For example, 64-kbit memories, such as those used in the chips in Figure 1, would have 256 × 256-bit structures. These would have a wide 256-bit word coming off internally. Normally, because of I/O-pin limitations, as little as 1 bit of this 256-bit-wide internal word might be used. But because the ALUs are brought onto the chip, the full parallelism of the 256-bit-wide word can be used.

Second, like others in massive parallelism, Morton required the simplicity of bit-serial operation. Only with bit-serial operation could the 256 multipliers be crammed into the chips in Figure 1. With bit-serial operation, the multipliers need only be 2-input AND gates (although for pattern matching, Morton also includes EXCLUSIVE-OR gates). When smaller matrices are involved, Morton trades up to more ALU parallelism, as will be seen in Figure 7.

MARRIAGE OF MEMORY WITH ALU's HOLDS GREAT PROMISE

■ TEXTBOOK MATH PROVIDES CHIP'S ROOTS

Figure 2a shows the classical mathematical notation for two matrices, M and V, being multiplied to produce a resulting matrix R. This is straight out of a textbook. For simplification purposes, we assume there is just the first column of the V matrix, which of

The traditional approach to IC design.

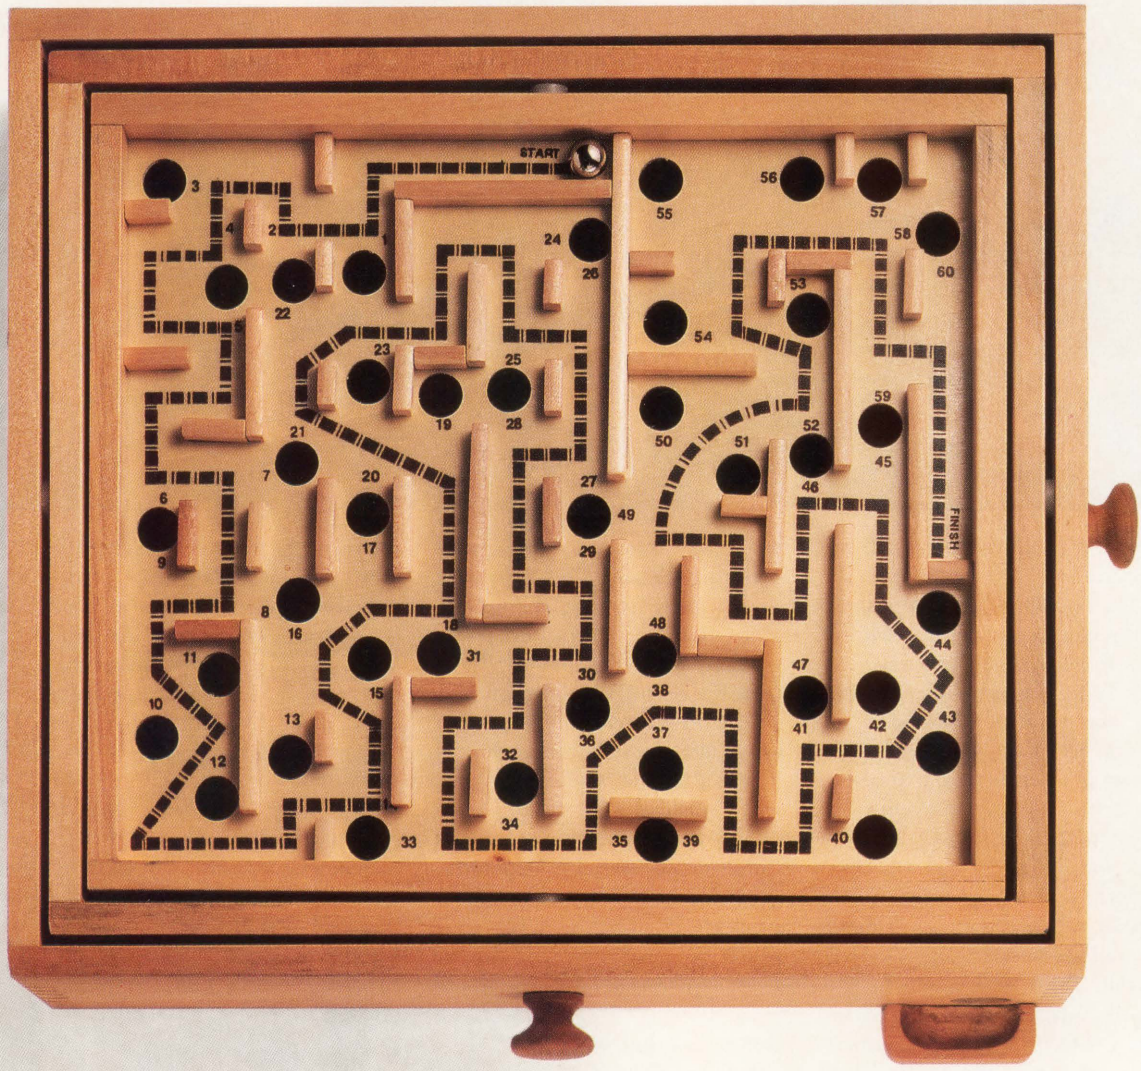

The Cadence approach.

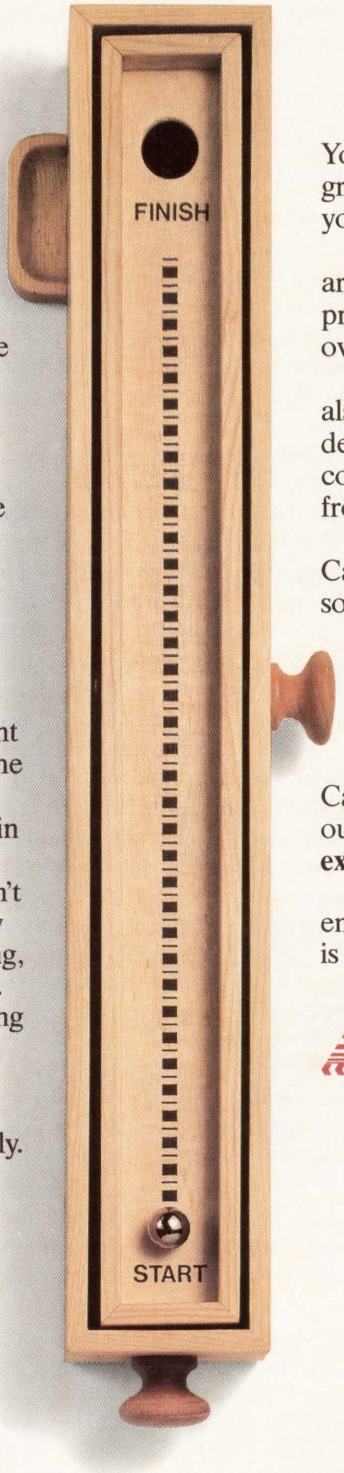

How you play the game determines if you win or lose.

And there's plenty at stake. Your design. Your product. Maybe even your company.

Traditional IC design is full of pitfalls and blind alleys. While there are plenty of good tools available, none of them really work well together.

There is a better approach.

Design Framework™ architecture from Cadence. The first integrated design environment to support the entire IC design process. A system that lets you go from start to finish in one smooth, direct path.

Not a "shell," Design Framework architecture is a unified environment where all design tools share the same user interface and design database. So important details never get lost in transit. Or garbled in translation.

But Design Framework tools don't just passively coexist. They actively cooperate. As your design rolls along, you see the impact of every change. In real time. Catching and correcting errors as they're made. Eliminating the need to go back and start over.

The bottom line—you finish designs faster and more economically. You get to market sooner. And put greater distance between you and your competition.

In fact, Design Framework architecture can boost your design productivity five times or more over the traditional approach.

Design Framework architecture also fits easily into your existing design environment. You can even couple tools you developed or bought from other vendors.

And it's all brought to you by Cadence. The IC design automation software tool leader. We'll be happy to tell you more. Write or call for a copy of our IC Design Game Plan: Cadence Design Systems, Inc., 555 River Oaks Parkway, San Jose, CA 95134, inside California: **1-800-672-3470, ext. 866,** outside California: **1-800-538-8157, ext. 866.**

Because in today's competitive environment playing by the old rules is a losing game.

CĀDENCE

Leadership by Design

CIRCLE NUMBER 6

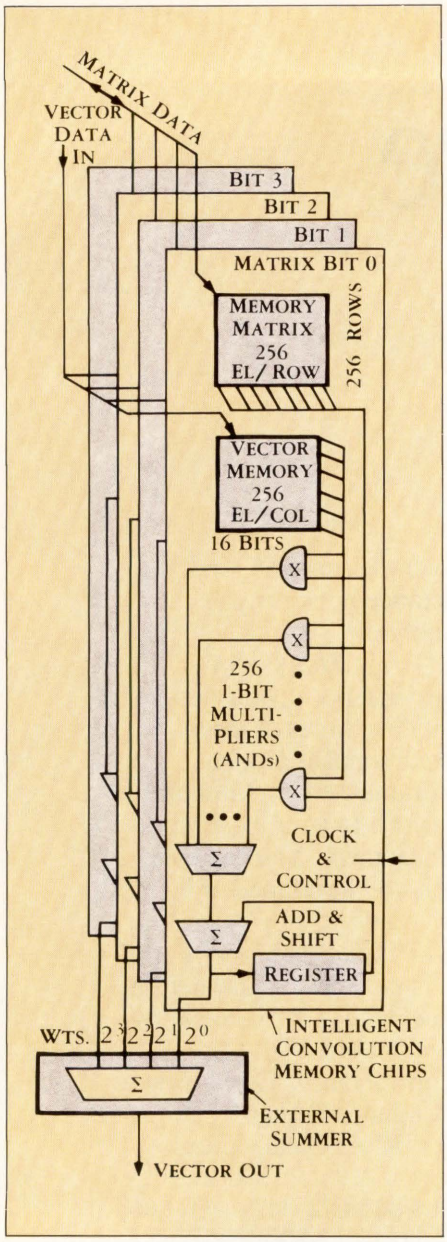

Figure 1. One form of Oxford's 'Intelligent-Memory' chips for massive parallelism in multiplying whole matrices.

$$\begin{matrix} & M & x & V & = & R \end{matrix}$$

$$\begin{bmatrix} M_{1,1} & M_{1,2} & M_{1,3} \\ M_{2,1} & M_{2,2} & M_{2,3} \\ M_{3,1} & M_{3,2} & M_{3,3} \end{bmatrix} \times \begin{bmatrix} V_{1,1} \\ V_{2,1} \\ V_{3,1} \end{bmatrix} = \begin{bmatrix} R_{1,1} \\ R_{2,1} \\ R_{3,1} \end{bmatrix}$$

[A] TEXTBOOK NOTATION FOR VECTOR MULTIPLICATION

$$R_{1,1} = M_{1,1}\,V_{1,1} + M_{1,2}\,V_{2,1} + M_{1,3}\,V_{3,1}$$
$$R_{2,1} = M_{2,1}\,V_{1,1} + M_{2,2}\,V_{2,1} + M_{2,3}\,V_{3,1}$$
$$R_{3,1} = M_{3,1}\,V_{1,1} + M_{3,2}\,V_{2,1} + M_{3,3}\,V_{3,1}$$

[B] TEXTBOOK EQUATIONS FOR MATRIX MULTIPLICATION (WORD LEVEL)

Figure 2. Review of matrix multiplication: (a) is the notation for the overall operation; (b) shows how the R result elements are obtained (at word level).

course makes V a column vector. However, what is shown can be extended to full multicolumn arrays or matrices.

Figure 2b is a reminder of how the elements of the resulting matrix R (like V, just a column vector) are computed. The computation is the ubiquitous sum-of-products that's so widely found in all computer programs for engineering and science (a fact that makes some of these massively parallel approaches all the more universally interesting).

Figure 3 gets down to the bit-level details of the mathematics that must be performed by the chips. It shows how the first result element of Figure 2b would be calculated in longhand with paper and pencil. To the left we start to sketch in the connections between the longhand operation and Morton's architecture. Note that we have chosen 4 bits of precision.

■ Mapping Math Into VLSI

Figure 4 continues the mapping of the Figure 2b computations into Morton's chips. Because the precision is 4 bits, four of Morton's intelligent convolution memories would be needed. The M memories would hold the 3 × 3 M matrix with the precision bits strung out over the four chips as shown.

The V memory would hold the V column vector, with that vector loaded redundantly into each of the four chips (each of the chips has a complete copy of V).

Three multiplier ANDs would be provided below on each chip. Three would be needed to match the 3 × 3 size of the matrix, so that the three multiplications indicated in Figure 3 could be moved along in parallel. The summers and accumulator registers at the bottom of the chips (refer back to Figure 1 for best details) would add up the products from the different longhand computations as the operation progressed. A key premise of Morton's architecture is that it makes no difference how or when the products are summed as long as they all do get summed. Of course it's necessary that there be built-in shifting to take care of bit-position scaling.

Since we are assuming that this is a dedicated application, only the capacity to handle the example is provided. Obviously, for flexibility in more general applications, the matrix would be sized to handle the largest problem of interest. Then smaller problems would only partially fill the memories and fewer precision cycles might be used.

Let us summarize how Morton performs the longhand computations of Figure 3 in terms of the Figure 4 diagram.

1. Each row of partial products is handled by a different chip, according to the precision bit of the M-matrix element.

2. For the computation to move along each of the partial-product rows, it takes successive cycles in time, in which the chip's given M bit is multiplied against all the V bits.

3. The sums of the like rows are continuously accumulated as they are generated by the adders and registers at the bottom of the three chips (refer back to Figure 1 for details). By the end of the computation, the external summer chip will have accumulated the total sum that will represent an element for the result matrix (vector) R.

This process is quite confusing to follow unless you meticulously and methodically keep track of all the matrix subscripts. At the same time you should be constantly checking that what is going on inside the chips is what *should* be going on as indicated by the Figure 3 mathematics.

The software (or firmware, depending on how you view it) flow-diagram of Figure 5 will help you see what happens as the chips do a matrix multiplication.

■ Program Flow Helps Explain Steps

It's easier to follow what happens if you unravel the loops of the program flow given in Figure 5, working from the inner loop outwards.

The inner loop counts out the bits of V-element precision. Each chip's given bit of M-element precision is being swept across all the bits of V precision. Note that the inner loop operation is parallel with respect to bits of M elements, but in series with respect to bits of V elements.

In our example, there would be four traverses of the inner loop, one for each of the 4 bits of V-vector-element precision. At the end of these traverses, a complete answer—a result element for R—will be in the external summer's output register.

The middle loop iterates through all the rows of the M matrix and the outer loop does any additional V columns—if V is a

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

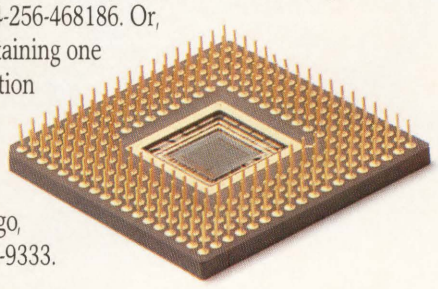

A Better BiCMOS Array is Here.

AMCC

CIRCLE NUMBER 7

You can't catch a rabbit who refuses to rest.

Recently, some ASIC competitors have entered the ECL race and tried to tortoise their way past Motorola's rabbit. Unfortunately for them, they're hoping for a lazy rabbit, while through three generations of ECL arrays the one they're chasing has refused to rest.

So how can they catch an energetic rabbit? They can't.

Performance that won't quit.

Motorola's ECL arrays are the fastest track to total system performance. With speeds of up to 1200 MHz and typical gate delays of 100 picoseconds they're the ultimate in bipolar performance. No longer will designers have to curtail their imaginations to use parts with limited specs.

Motorola ECL arrays not only give you speed and up to 12,402 gates of logic power, they give you the versatility to use them. Programmable speed-power levels put designers in control of both macro and drive currents to allow peak performance where timing is critical and the ability to trim power consumption where speed is less crucial. Three-level series gating lets you multiplex latch inputs without the cost of additional cell usage and delay.

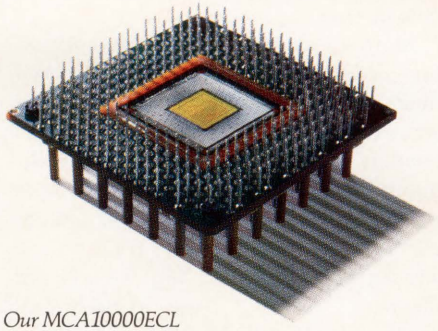

Our MCA10000ECL array is available in a ceramic 235 PGA with 180 I/Os, a polyimide/glass 289 PGA with 256 I/Os, and a 360 lead TAB tape.

Runs faster, cooler and longer.

The MOSAIC III* process used in our third generation ECL arrays utilizes an innovative "edge defined" technique to achieve submicron features without submicron lithography. The end product is very fast—with typical gate delays of 100 picoseconds and metal delays down to 40 ps for a fan out of 2.

An integral heat sink package assures thermal efficiency for reliability and ease of system design. Using impinged air, our packages have a junction-to-ambient thermal resistance of only 2.0°C/W with a forced air velocity of 500 LFPM. This allows high performance arrays to dissipate over 20 watts in an air cooled environment.

Advanced support keeps you race-ready.

To simplify your designs and maximize your potentials, Motorola offers the most comprehensive library available. No longer will you be forced to build system-level functions from a limited selection of library macros and have to accept the loss of performance and density that comes with it. Our library features over 225 functions, three-level series gating, and expandable macros, all of which yield higher performance and better logic density. The library is supported by popular engineering workstations plus a dedicated mainframe which puts even physical layout under your control.

The race doesn't stop.

We know that Motorola has to live up to some pretty tough expectations, after all, we wrote the book on ECL. So we've added a host of innovations designed to make sure your products begin competitive and stay that way. Innovations like tape bonding and STECL outputs.

Multilayer TAB technology provides a controlled impedance environment to minimize interconnect delays. And programmable Series Terminated ECL (STECL) outputs simplify multichip applications.

One-on-one design-in help.

Motorola's ECL arrays provide the most advanced features and the most progressive designs, without ever compromising reliability and the always-important cost/performance ratio. Give us a call for more information. Call toll-free any weekday, 8:00 a.m. to 4:30 p.m., M.S.T.

1-800-521-6274

If the call can't answer all your questions, we'll have a local applications engineer contact you. For published technical data, just complete and return the coupon below.

We're on your design-in team.

MOTOROLA

*MOSAIC is a trademark of Motorola.

To: Motorola Semiconductor Products, Inc.
P.O. Box 20912, Phoenix, AZ 85036
Please send me more information on Motorola's semi-custom ECL arrays.

373VLSI120088

Name _____

Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Call me (_____) _____

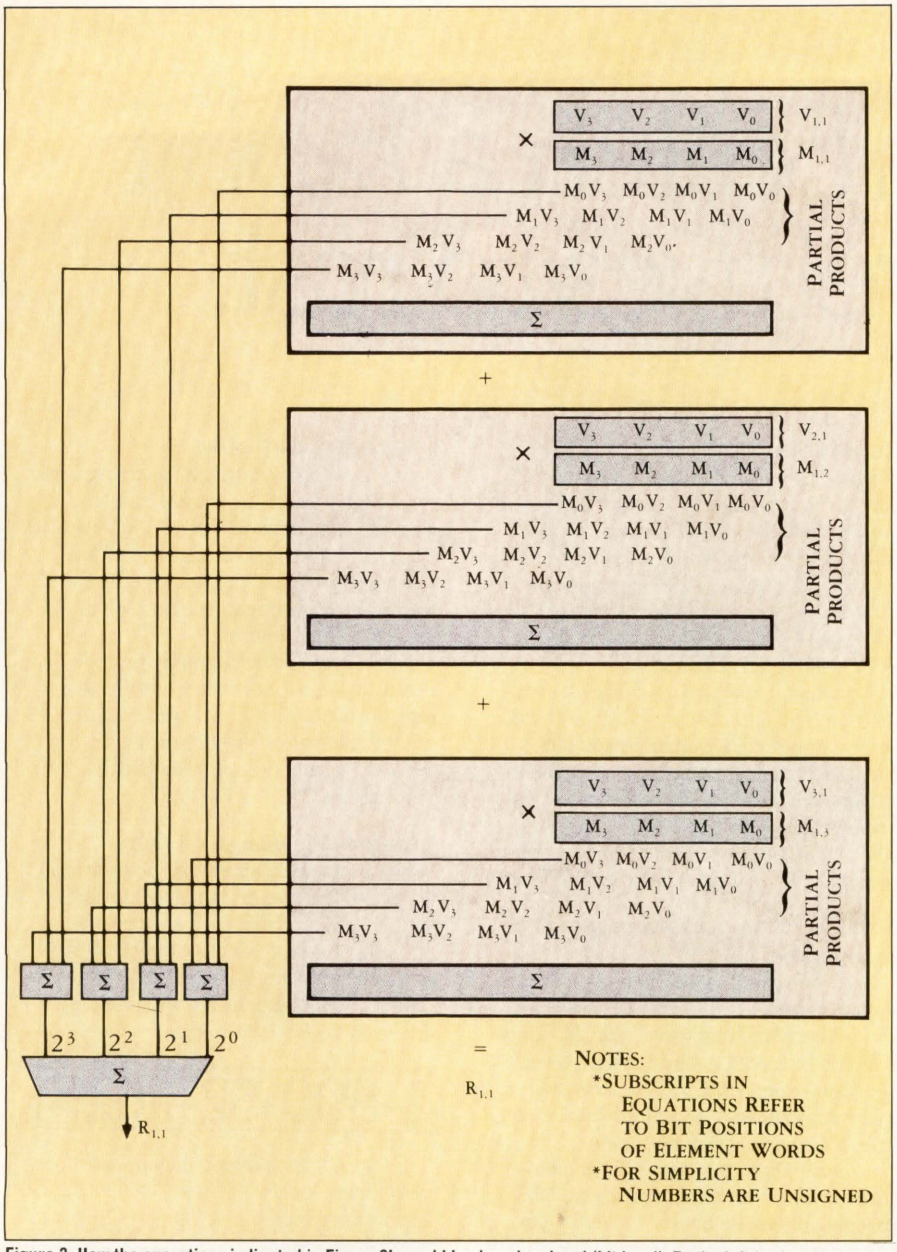

Figure 3. How the operations indicated in Figure 2b would be done longhand (bit level). To the left is sketched in how this would be mapped into Morton's chips.

full matrix rather than just a single column vector, as in our example.

■ Performance Gain vs. Sequential Architecture

What, if any, is the performance gain with this mostly parallel architecture? For our small 3 × 3 matrix times a 3 × 1 vector with just 4-bit precision, there is not much advantage. If the precision were raised to 16 or 32 bits there would be even less advantage. Some of the modern DSP microprocessors that can handle 16- or 32-bit multiply-accumulates in one cycle would outperform the Morton architecture. But there are two aspects of real-world matrix operations that favor massively parallel architectures such as Morton's. First of all, as implied by the values in Figure 1, many matrix applications call for much larger matrices than our 3 × 3 example. Even if the end user's problem doesn't start out large, it is one of the appealing characteristics of matrix formulations for real-world problems that it is often a trivial matter to scale up the problem for more accuracy or resolution; matrix formulations, in fact, encourage problem growth.

Second, many algorithms require that the matrix be used repetitively. An architecture like Morton's that keeps everything on a chip during computations obviates having to waste time continually transferring the matrix from memory to the processor. This is especially true of algorithms in which the data is iteratively updated—for example, when adaptive feedback or network learning is involved.

Morton says his architecture allows maintaining performance when problems get bigger, in contrast to sequential machines in which the performance progressively bogs down as problem size grows. By now you should appreciate how these massively parallel architectures exploit the parallelism inherent in the matrix computations, especially that part of the partial product formation that is sometimes called the "dot" product. This allows multiple groups of chips to work in parallel to simultaneously handle very large matrices. Morton also points out that he has different configurations of the architecture that can increase the performance when handling smaller matrices, such as the graphics chip described in Figure 7.

The middle and outer loops build up the number of reuses of the architecture. The middle loop iterates through all the M rows and is traversed three times in our example. The outer loop iterates through all the columns of the V matrix and is traversed only once in our example (but it is shown to remind that V can be a full multicolumn matrix).

Morton says that these outer loops do not affect the execution time, because they are absorbed in internal pipelining. They only add delay in getting the first result out, as happens with all pipelines.

■ Avoiding Memory-Load/Store Bottleneck

All too often array processors lose so much time in getting their memories loaded up before the computation that their overall performance falls far short of what was promised by their internal operation. Morton estimates that the use of conventional structures as the foundation for his memories and ALU functions on VLSI chips will permit the chips to avoid being slowed down from the memory load and unload times.

Morton's strategy is to use two addressing modes for his memories: the special mode for doing the matrix operations, as we have shown, and the conventional memory mode. Thus, in addition to the highly parallel internal accesses that send the data to the on-chip ALUs, his designs use conventional memory reading and writing to go on and off chip. To keep our diagrams simple, these conventional read/write addressing and data paths are only given the merest suggestion in our diagrams.

Because of the dual-port nature of the memories envisioned by Morton, in many applications the loading of one or both of the memories can go on concurrently with the basic massively parallel matrix operations. For example, data can be written in or read out by the host microprocessor

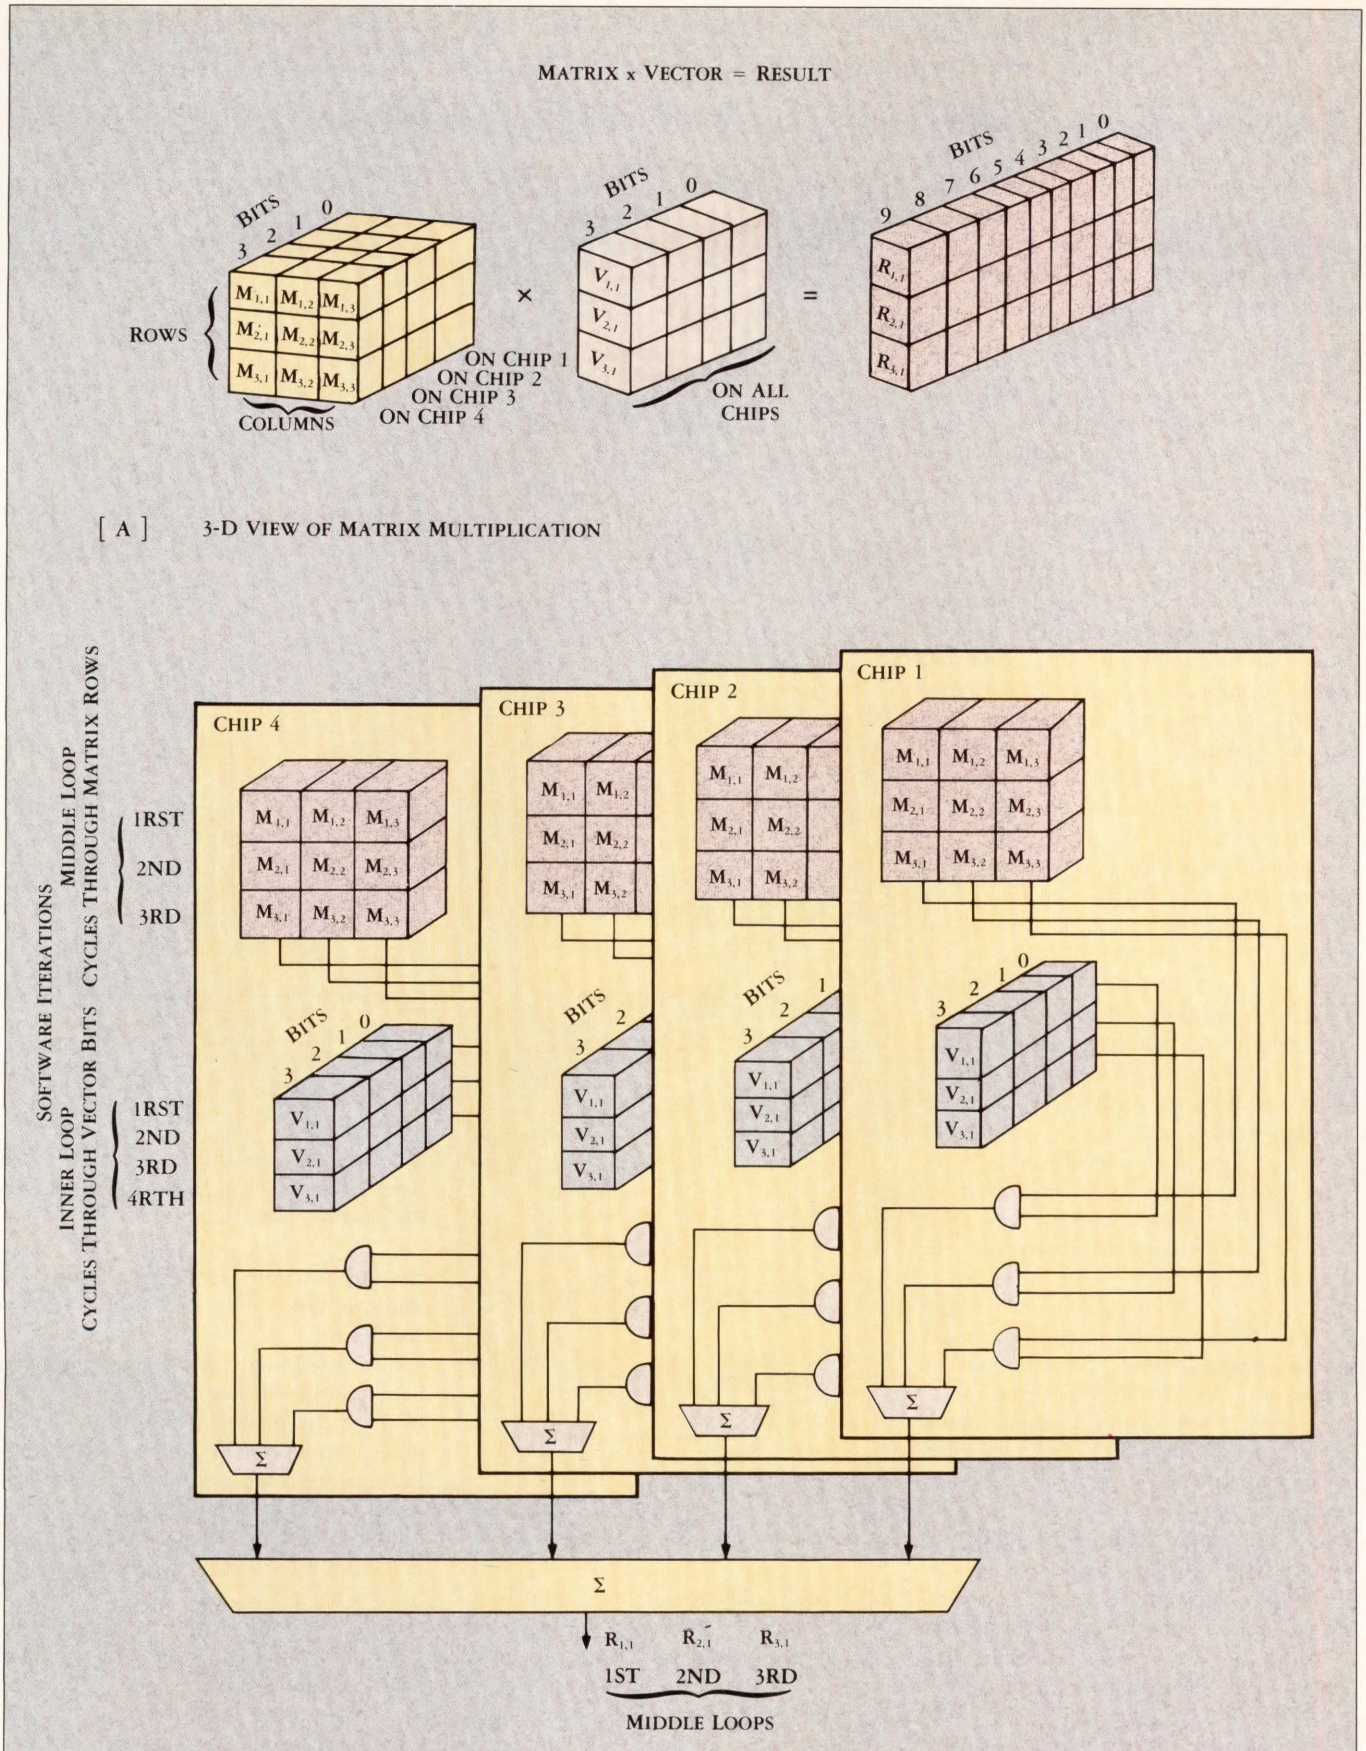

Figure 4. The multiplication and summing operations of the textbook example of Figures 2 and 3 are fully mapped onto the Oxford chips. The four chips must be put through four timewise cycles.

or via DMA. It would be somewhat like what is done in video RAMs, where the host microprocessor can be updating display information as the CRT is constantly refreshed. This would be valuable for selectively upgrading or modifying coefficients for neural "learning" networks and adaptive DSP filters.

Morton says the design's ability to randomly read and write into any memory location is a distinct advantage over those massively parallel systems in which data can only be shifted in one bit at a time from the chip edges and then moved one bit at a time from neighboring cell to neighboring cell.

■ Configurations For Different Applications

As with most other massively parallel schemes, Morton trades off generality for hardware optimization for particular classes of application algorithms. Consequently, Morton finds it necessary to recast his concepts into different forms for various end uses. Initial variations were aimed at pattern recognition, 2-D FFTs and 3-D graphics.

Figure 6 shows the concept's use for a DSP FIR filter. The FIR filter coefficients or weights would be preloaded into the M memory while the sampled-data values from the signal stream would be shifted into the V memory. Morton's chips have internal features that permit this shifting. In effect, these techniques allow him to maintain the dense structure of a regular random-address memory yet provide the shifting. Though the shifting adds to the chip's complexity, it is actually worth it, Morton says, because of the many other applications in which this type of shifting is essential.

The outputs from the chips would not be considered elements of an R vector per se, but as signal samples. From the matrix viewpoint, the filtering operation might be considered as taking an input vector that matched the length of the FIR filter coefficients and producing a similar length output vector. The difference from ordinary static matrix operation would be that the vector's contents would change because after each output element was computed, a new element would be shifted into the V vector, pushing the oldest element off. The new one-element-shifted vector would be used to compute the next output element. Thus, the input and output vectors would represent windows slid over the input and output signals.

The software flow diagram for the FIR filter would be similar to that shown in Figure 5. The iterations of the inner loop would attend to cranking out the bitwise-

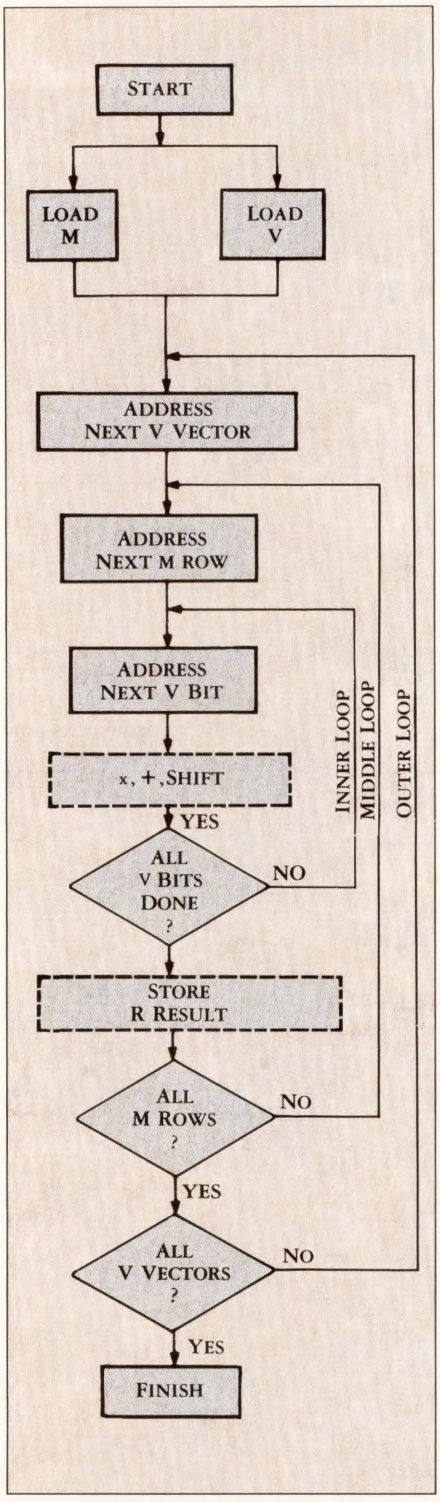

Figure 5. Software flow for the Figure 4 hardware.

serial multiplications and accumulations that generate the output samples. The middle loop would take in the new input samples. In this case there would be no decision at the end of the middle loop; it would just run continuously as it fed upon the endless samples of the signal data stream. Normally there wouldn't be an outer loop, except in situations where the applications designer wanted to change the filter coefficients adaptively. Then a loop could be added that would observe

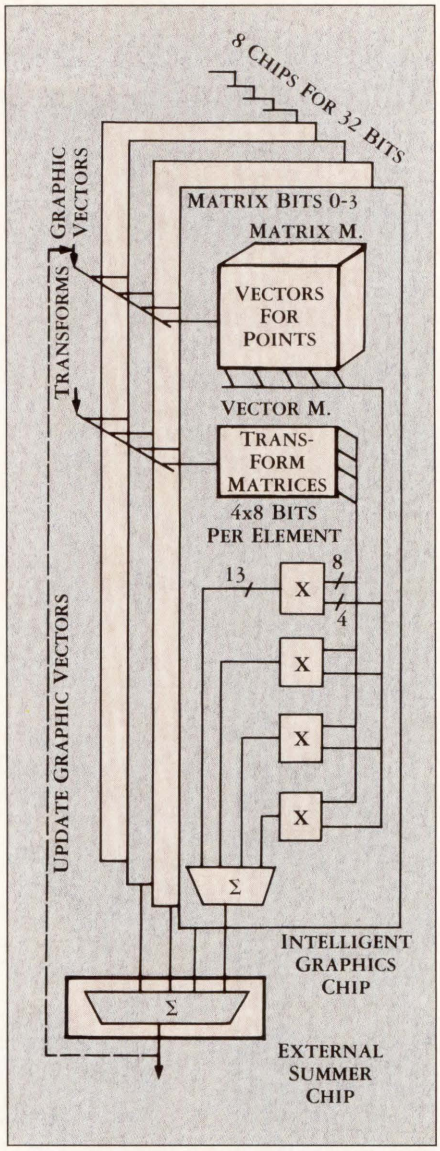

Figure 7. Variation of the Oxford chip for graphics transformations (and FFTs).

the output stream, and if that didn't meet some reference criterion, change the FIR coefficients.

The sampling period and thus the bandwidth of the FIR would depend mainly upon the precision or the number of inner-loop reprises. For the 40-MHz clock Morton uses for his performance estimates, the filter would take 25 ns for each bit of precision of the vector. For the 8-bit precision often used in video, this would be 200 ns or 5 MHz (assuming the signal samples are inputted in a transparent, pipelined manner). Thus Morton's approach, while superior to the Von Neumann sequential DSPs, still can't match some of the dedicated-hardware parallel-bit schemes for FIRs such as have been offered by TRW, Zoran, and Inmos. He says he could double his speed (to 100 ns per pixel) by doubling up on the number of chips and dividing them into two sets. One set would handle bits 0-3 of the

SEMI-CUSTOM SMART 100V ICs IN 45 DAYS.

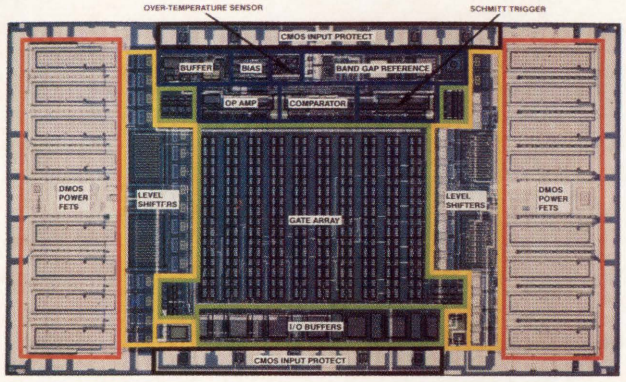

The MPD8020 uses Mixed CMOS/DMOS/Bipolar Technology to provide the user true monolithic smart power management.

Only 45 days after you give us the layout of your breadboard, built from our Kits #1 and #2, we'll give you perfectly tailored smart 100V ICs. What's more, you'll get them for a fraction of the cost of custom circuits.

Design Power ICs Faster.

Micrel's new MPD8020 CMOS/DMOS Semi-custom High Voltage Array combines CMOS analog circuits, TTL/CMOS compatible high speed CMOS logic, and high voltage DMOS power drive circuits on one monolithic IC.

A Great Library of Parts on Every Array.

16 fully floating 100V, 200 mA N-channel DMOS FETs. • 3 op amp/comparator/Schmitt trigger programmable macro cells and numerous array op amps. • 1 unity gain analog output buffer. • 1 band-gap reference. • 1 overtemperature sensor. • 16 high voltage CMOS level shifters • 200 CMOS gates in an uncommitted array. • 12 TTL/CMOS I/O buffers. • 16 medium current sink pre-drivers. • 4 internal high and low voltage power supplies. • Zeners, resistors, capacitors and more.

Two Development Kits.

For breadboarding your semi-custom MPD8020, Micrel offers two development kits to demonstrate the operation of key SSI and MSI circuits. Each is housed in a 40 pin DIP. Kit #1 ($20) provides 11 commonly used analog circuits. Kit #2 ($15) carries 8 digital circuits to check speed and digital timing characteristics.

After the customer has used the development kits to determine the interconnect pattern, Micrel turns each IC into a proprietary smart 100V ASIC for about one-sixth the cost of a custom IC. As your needs grow, we can quickly turn your semi-custom chip into a full custom chip for even greater savings.

The MPD8020 can be packaged in a 44 PIN JEDEC PLCC with an integral heat sink, 16 to 48 PIN DIPs, or single in-line power package. Packaged units available to MIL STD 883C. Dice are also available for hybrid manufacturers.

A single +5 Volt to +15 Volt supply powers the logic and analog circuitry. High voltage portions operate at +20 to +100 volts. The chip can also derive the +15 volt supply from a 24V, 48V, or 100V high voltage supply. An internal voltage pump can be used to drive the high side gates of the power N-channel DMOS FETs at 15 volts above the +100 volt supply for rail-to-rail high voltage switching.

Wide Range of Applications.

Use the MPD8020 in switching regulators, motor control, relay and solenoid drivers, smart switch with bus decode, smart lamp drivers, automotive switching, printer solenoid drivers, and high voltage display drivers.

Build Safety and Reliability Into Your Product.

You specify which safety features you want built-in such as overtemperature, overcurrent, short circuit, and overvoltage protection. The circuit can then take immediate action whenever any of these faults are detected and send a status signal back on your microprocessor data base. You can design safety in at the outset.

For more information, fill out the coupon below or contact Marvin Vander Kooi, Micrel, Inc., 1235 Midas Way, Sunnyvale, CA 94086.
Phone (408) 245-2500. FAX (408) 245-4175.

MICREL
THE INTELLIGENT POWER COMPANY

SEND THE DETAILS ON INTELLIGENT POWER.

NAME _____ TITLE _____

COMPANY _____ PHONE _____

ADDRESS _____

APPLICATION _____

ESTIMATED YEARLY VOLUME. _____

MICREL, INC., 1235 MIDAS WAY, SUNNYVALE, CA 94086

CIRCLE NUMBER 29

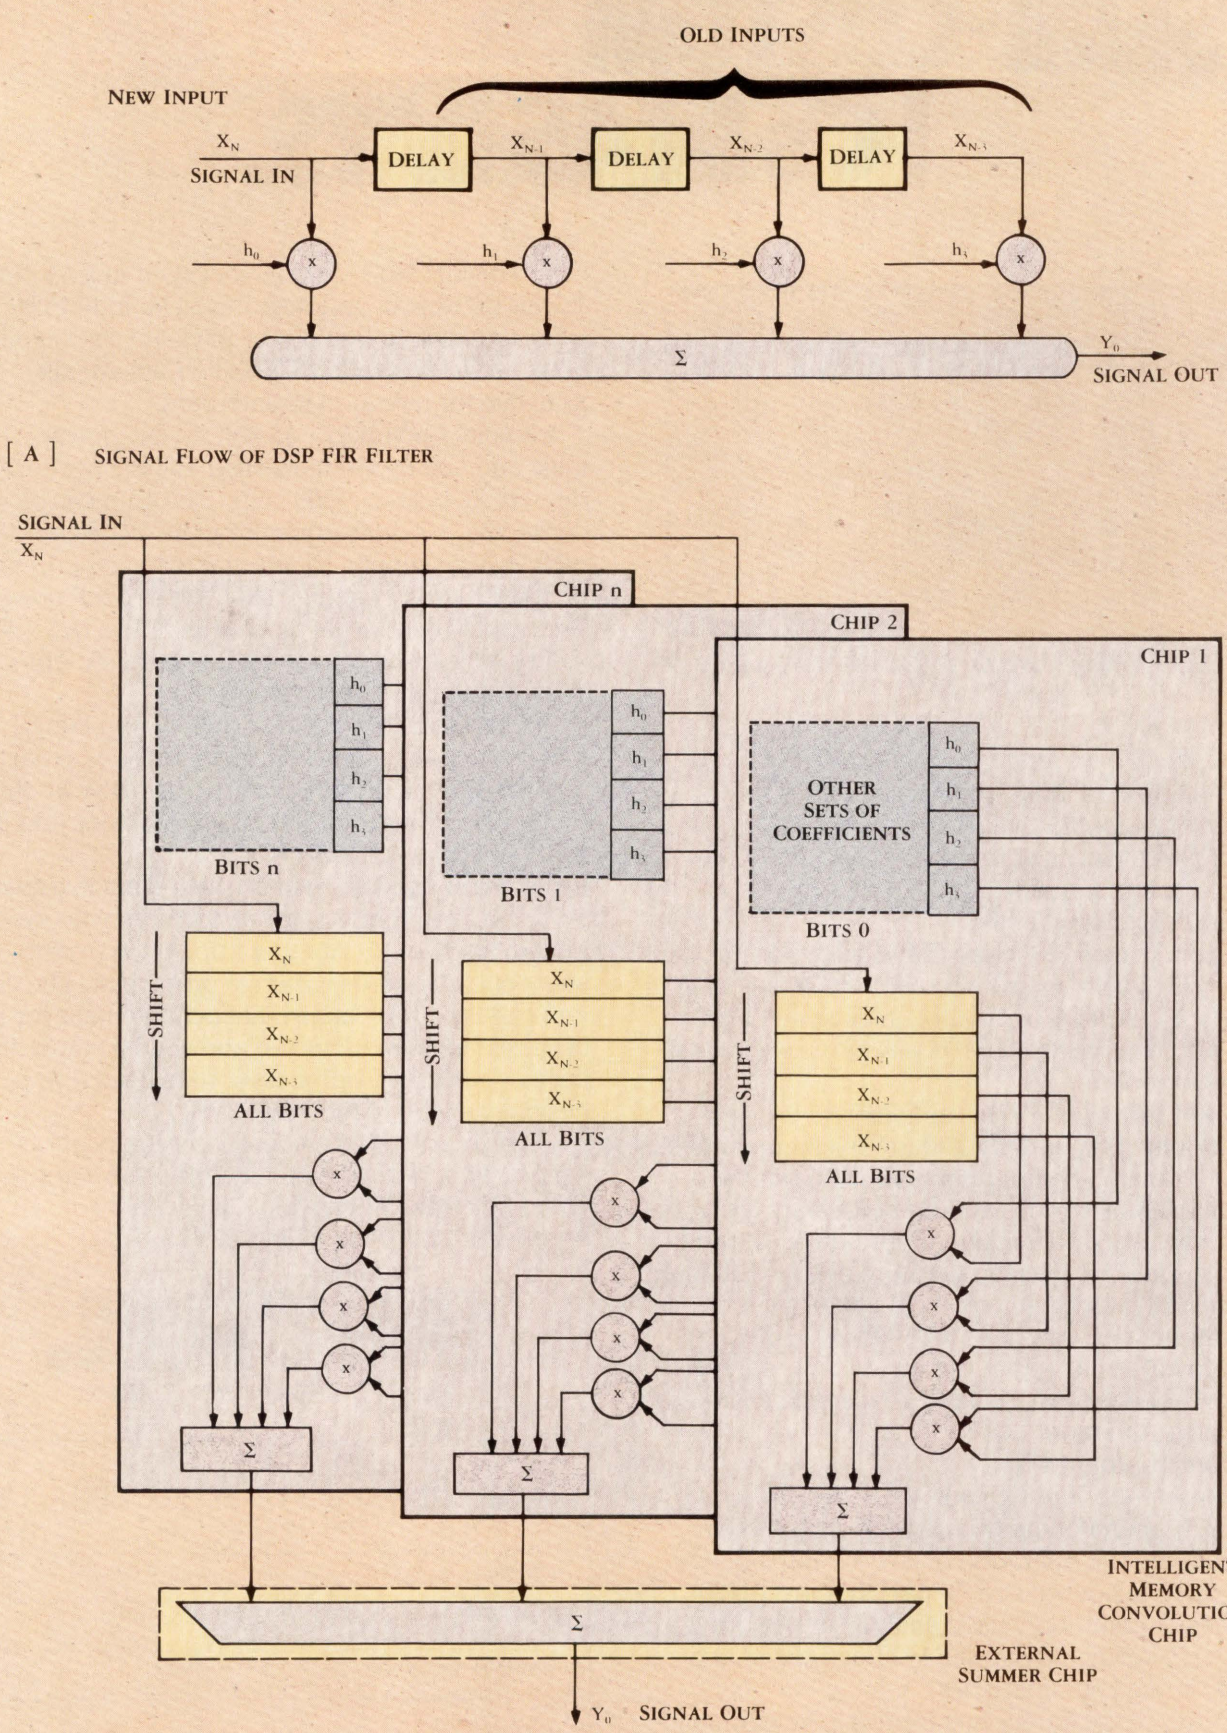

Figure 6. How a FIR filter structure might be done with the Oxford chips. The filter coefficients would be in M memory while the sampled-data signal values would be shifted into V memory.

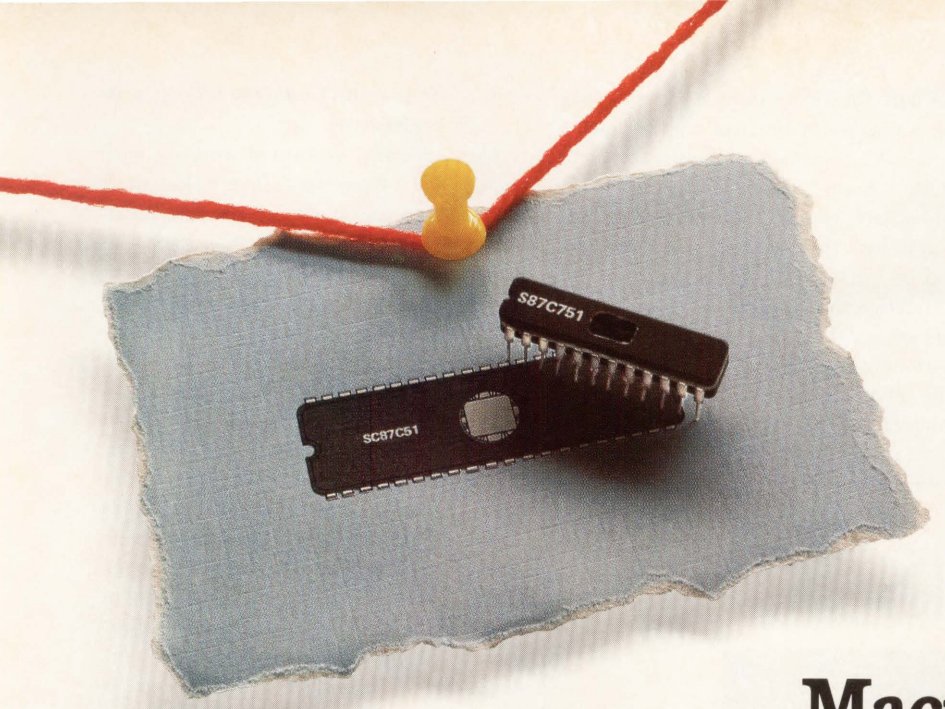

Macro performance in a micro size

Signetics 83/87C751 microcontroller—the performance of an 80C51 packed into a 24-pin skinny DIP or SMD 28-pin PLCC.

Designed around the 80C51 architecture, our new 83/87C751s let you use your existing instruction set and code to gain the enormous benefits of a smaller package at a smaller price.

A lot of guts in a little package. Its 16 MHz operation means you don't have to sacrifice performance. You'll appreciate the convenience of EPROMs in UV or OTP. Coupled with the flexibility of an I^2C serial bus port. And the efficiency of an 8-bit architecture backed by a complete Signetics development system.

Think of it. Increase the power of your 4-bit applications. Or replace your logic blocks. Or lighten your handheld designs. All while decreasing system costs.

Call us for a free '751 Microcontroller Information Packet at (800) 227-1817, ext 991I. For surface mount requirements and military product availability, contact your local Signetics sales office.

One standard. 0 defects.

Signetics

a division of North American Philips Corporation

PHILIPS

PHILIPS

CIRCLE NUMBER 8

vector and the other set would handle bits 4-7.

This illustrates, again, that it takes the larger matrix-crunching operations for the bitwise-serial massively parallel schemes to outperform competing architectures. In general, Morton says, the advantage of his "vector-slice" architecture is its relatively greater flexibility. For example, he points out that the Figure 6 structure would be justified where multiple banks of FIR coefficients had to be in place on chip for instantaneous switching between filters.

Perhaps the best example for showing where Morton's architecture might excel is the graphics processor of Figure 7. Here the task is to do the rapid transformations of 3-D objects that are expected of modern CAD systems. In this case, the M matrix would be loaded with all the vectors that would define the hundreds of thousands of points of a 3-D object; for example, an aircraft or molecular model.

The V memory would be loaded with the standard 4×4 transformation matrices used in graphics to manipulate the 3-D object. These could rotate the object, scale it, add perspective, etc.

Here, because Morton is working with smaller 4×4 matrices, he forsakes strict bit serialism and trades up to more parallelism. This parallelism uses $4 \times 8 = 13$ multiplier units, for which he now has the chip area since he is using fewer of them.

Morton says that his actual designs for the graphics applications are far more complex. CAD workstation designers now want not only to give designers the ability to rapidly manipulate complex objects, but also to provide realistic coloring and shading when the object is illuminated by various light sources. Morton says his chips can encompass the nonlinear operations, such as reciprocals and square roots, by employing Taylor series approximations. Taylor series are convenient sum-of-products computations, so they can be handled by Morton's architecture as vector-times-vector operations.

The Figure 7 configuration is also ideal for doing complex 2-D FFTs at video rates, Morton asserts. For FFTs, the vector memory would hold successive sets of FFT butterfly coefficients. Meanwhile the matrix memory would be acting as a double buffer, with one butterfly being updated for each layer of FFT.

The advocates of massive parallelism talk in large numbers. Morton is no exception; he sees systems such as those in Figures 1 and 7 as being packaged first in SIPs, then on boards, to generate billions of MIPS at board level (trillions at system level).

In our next article in this series, we'll cover some of the other competing approaches to massive parallelism. ■

ACKNOWLEDGMENTS

John Little, MathWorks Inc., South Natick, Mass; John Poulton, University of North Carolina, Chapel Hill, N.C.; Howard Sussman, NEC, Natick, Mass.; Roger Thorpe, Active Memory Technology Inc., Irvine, Calif.

REFERENCES

KOWALIK, J. S. (editor), "Parallel MIMD Computation: HEP Supercomputer and its Applications," 1985, MIT Press, Cambridge, Mass.

POTTER, J. L. (editor), "The Massively Parallel Processor," 1985, MIT Press, Cambridge, Mass.

REED, D. A., and FUJIMOTO, R. M. (editors), "Multicomputer Networks: Message-Based Parallel Processing," 1987, MIT Press, Cambridge, Mass.

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

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

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

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

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

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

GenRad

The difference in software is the difference in test™

CIRCLE NUMBER 10

THE PEGASUS CPU

JAMES J. BOHANNON, ELXSI CORP., SAN JOSE, CALIF.

The design automation environment was critical in accomplishing our time-to-market goals

With the heat turned up on a new CPU project, the Pegasus 6460, Elxsi Inc. sought the best available approach for using CAE software. The design automation phase was critical to increasing the 6400 series' performance by a significant factor. Another project objective was shortening the time taken to develop a working prototype from when the logic design was completed on paper. By using CAE software, we intended to reduce the bring-up time to three months from the full year earlier machines had taken. The CAE software would also allow us to verify the correctness and performance of the intricate operation of the multistage pipelined CPU (see sidebar) before building a prototype.

The system designers minimized the use of exotic or leading-edge semiconductor technologies in the new CPU to reduce project risks. Instead, they adopted a more complex architecture to accomplish the system performance goals. The CPU would be built with ECL RAMs—15 different ECL gate-arrays designs, and discrete ECL parts totaling about 250,000 gates spread across two very large boards. The CAE tools had to be fast for managing such a large design. Most CAE libraries were already available for design simulation and timing analysis. However, because the design methodology

Elxsi's 6400: Nearly 100 Percent Efficient

The Pegasus "Superframe," the latest addition to the Elxsi System 6400 family of computers, was designed to compete head-on with the most powerful IBM and Digital Equipment Corp. platforms. The system, which can have up to 12 parallel processors, has a modular architecture that provides almost 100 percent processing efficiency regardless of the number of parallel CPUs installed. For example, while the individual Pegasus CPU delivers 25 VAX Mips and 10 Mflops (100 × 100 Linpack), a fully configured 6400 with 10 Pegasus CPUs is designed to deliver 250 Mips and 100 Mflops.

A major design requirement for the new Pegasus 6460 CPU was that the 6460 modules could be easily plugged into any existing System 6400 computers using 6410 or 6420 CPUs. This upgrade, which is fully object-code compatible, will increase the power of a system using 6420s by a factor of more than three—based on the guaranteed 25 Mips minimum rating for the 6460. At press time, the prototype 6460 boards were executing code and easily surpassing the 25 Mips requirement. It is anticipated by Elxsi that benchmarks to be run this month will show the 6460 delivering in excess of 40 Mips per CPU.

The heart of the System 6400 is a tightly coupled, message-based, and bus-oriented multiprocessing system. The system can be configured with up to a dozen 6410 or 6420 CPUs (or 10 6460s), four I/O processors, and 2 Gbytes of memory. All units have access to the common Gigabus (Figure A), and a synchronous bus that clocks data and instructions at a 320 Mbyte/s rate.

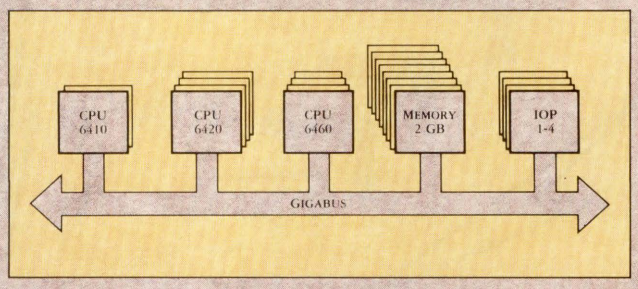

The new 6460 CPU uses a highly bypassed, five-stage pipeline that is optimized for fast branching. The CPU includes dual floating-point units and multiple register files. It's fabricated with 15 different ECL gate arrays, comprising more than 250,000 gates. Each CPU has an on-board 1-Mbyte cache memory, evenly divided between instruction and data functions. For applications such as real-time processing, the system features user-controllable partitioning that provides the ability to reserve either one-eighth, one-fourth, or one-half of the cache memory. In addition, the cache memory has write-back, write-through, and write-around modes.

One interesting result of the new CPU design experience that proved the worth of the CAE system used at Elxsi was that the more powerful 6460 CPU required only two boards—compared with the three required for the lower performance 6420.

—*R.C.W.*

adopted was not purely synchronous, traditional tools for doing timing analysis would not work.

There was also an emphasis on making new CAE tools match the paradigm of the CAE tools already in use at Elxsi. That emphasis was sought because the CPU design project had already begun and was under a tight schedule. This approach would also reduce the learning curve for the designers. And because of the tight schedule, there was not enough time to hire many programmers to put together a customized design automation system from scratch. Since no single third-party CAE software vendor could fulfill all of our needs, we concentrated on selecting and integrating third-party software with our tools.

■ THIRD-PARTY SOFTWARE

Third-party software packages had to satisfy several criteria before we would select them. Most important was functionality and performance, which we typically evaluated with our own benchmarks, a process that turned out to take more effort than we expected. The level of service and support provided was also important because of our limited resources in debugging problems. Next came purchase and maintenance costs. The quality of documentation was considered the least important criterion.

After selection, a tool had to be integrated into our design environment. Integration was required at several levels.

First, we needed to write netlist translators and other "bridging" programs (Table 1) between schematic capture software, our physical design database, simulation tools, timing analysis tools, and the netlist formats required by our gate-array vendor.

Second, the user interface needed to be customized for all of these programs. The designer was familiar with analyzing his design using the signal and gate names that appear on his logic drawings. Unfortunately, each of the tools supplied by different software vendors has its own naming conventions for signals and gates, as well as unique ways of handling buses, sized parts, and other logic-drawing contructs. Rather than require the designers to track perhaps four different names for the same signal, we decided to build extra features into our netlist translators to generate information that would allow front-end software to translate names back and forth between different netlist notations while the user interacts with the other tools.

Third, none of the tools (Table 2) we selected did a good job tracking design changes during parallel design activity. Not only did we have to track different versions of logic, but we also had to keep test vectors for each of the gate arrays and diagnostics for the CPU itself up to date. We still don't have a good way of handling this problem.

■ PHYSICAL DESIGN

Elxsi already had a physical design system in place before

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

Plessey - Unsurpassed Process Technology

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

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

Plessey - The Ultimate in ASIC Technology

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

Plessey - Unparalleled CAD Support

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

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

Plessey - Standard Products And Discrete Components

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

High performance solutions are also offered in radio communications, digital

ıce that gives you the edge.

PLESSEY KEY PROCESS TECHNOLOGY

BIPOLAR

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

MOS

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

BIPOLAR (CDI)

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

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

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

Plessey - Over Two Decades Of Quality Commitment

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

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

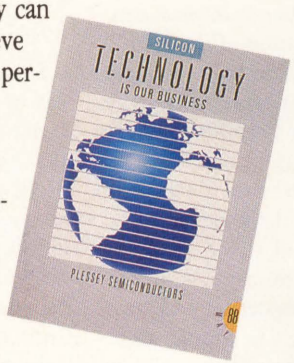

In North America call 1-800-441-5665.
Outside North America call 44-793-726666.

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

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

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

PLESSEY

Number 1 for ASICs

CIRCLE NUMBER 11

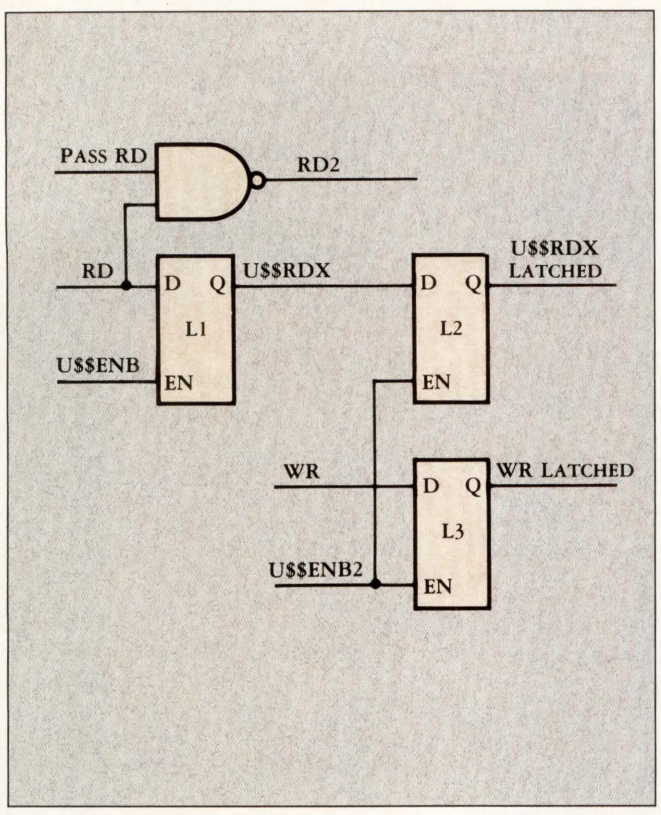

Figure 1. An example of fault-list back-annotation. Signals with a U$$ prefix have undetected faults. The program Drawfaults adds the U$$ prefix to the signal names in the Valid Logic schematics (in GED) for signals with untested faults.

the Pegasus CPU project began. Valid Logic workstations were used to enter logic diagrams schematically, and Valid software was used to package the design. The Valid Logic graphics editor (GED) was also used to manually enter placement of components on the two boards. Automatic placement was not used because the designers felt that they could optimize the timing at the board level by hand better than they could using automatic placement tools. Both the logic drawings and the placement drawings were interpreted by software written at Elxsi to create a board-level physical design database, called DAD (design automation database). All information required to specify the board is entered through the drawings as properties of gates and signals in the logic drawings and components on the placement drawing.

DAD is the source of information for an automatic board router supplied by Shared Resources. After routing is complete, the routing information is read back into DAD. Also, information for board fabrication by outside vendors is produced by the Shared Resources software. DAD is used to generate files for automatic insertion of components into bare boards after they are fabricated. DAD is also used for consultation by test and manufacturing people to track down problems with a freshly assembled board. Since board fabrication is expensive, rework of a board is done when changes are required rather than rerouting and refabrication of the board. The rework instructions are generated by DAD after the changes have been made to the physical design database.

A new use of the DAD physical design database was the generatopn of delay information for all the signals on the board for simulation and timing analysis. A simple algorithm was derived to calculate delays based on the assumption of correctly terminated ECL wiring on the board. A delay file (containing minimum and maximum delays for the entire board) is generated and can be loaded by analysis tools. Rise and fall delays were not considered because there is typically not much difference between the two—the ECL wiring behaves like transmission lines and it is mostly the propagation delay that matters.

A new tool was created to read the logic drawings and a placement drawing for the gate arrays. This tool generates netlist files for input to Motorola's design automation system. It also catches many of the errors a designer might make before the design is transferred to Motorola's computer, which helped to reduce the time necessary to getthe gate-array logic right. A cross-reference is generated between our signal and gate names (based on the Valid SCALD netlist language) and Motorola's signal and gate names (based on the netlist language). We also wrote a program to read back placement and delay information from Motorola. Often, the designer does not completely hand-place a gate array, but only places gates on critical paths. Sometimes, pin assignments are made by hand on the placement drawing as well. The placement information from Motorola is annotated back onto the placement drawing for review by the designer. The pin placement information is also fed into the physical design database for board-level routing to the gate arrays.

■ Logic Design And Analysis

Something new for Elxsi was the extensive use of logic design and analysis tools during the design of a board set. We looked into several different kinds of tools. Logic-level simulation would be used to verify correct functionality of the machine. Behavioral-level simulation would be required to supply large missing pieces of the design until the gates for them could be designed. Timing verification would be used to verify the timing correctness of the design. Automatic test generation would be used to generate test patterns for the gate arrays. A fault-grading tool would tell us if our test patterns were sufficient, whether they should be expanded, or if a better strategy for testing a particular gate array was required. We also needed programs to compare simulation results between our simulator and Motorola's version of the LOGCAP simulator. Of course, we also needed libraries for all of these tools.

■ Logic Simulation

There are many suppliers of logic simulators. It was difficult to find time to evaluate many of them, so we studied the problem and developed a list of criteria that the simulator would have to meet.

First, the simulator had to be fast. We wanted to be able to simulate up to 10,000 cycles of the entire CPU within several hours, or overnight in the worst case. This would enable us to simulate already existing actual hardware diagnostics and snapshots of our existing CPUs (which run the same instruction set) to debug the new machine. We determined that generating test patterns especially for debugging the machine was a difficult, time-consuming alternative.

Second, it was necessary to be able to simulate the entire machine all at once so that hardware diagnostics and actual processes could be simulated. This meant being able to simulate about 250,000 gates and a little more than 8 Mbytes of main memory and caches.

Third, we decided this had to be done mainly at the gate-level to ensure correct operation of the machine. It was not clear at first whether we wanted to do simulation with actual delays or whether unit-delay simulation would suffice. It turned out that some logic in the machine depends on minimum delays, and we were forced to use reasonably realistic delays for that part of the logic.

Z I L O G

Z180

Great code compatibility. Terrific performance. Superintegration.™

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

Full software compatibility.

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

Enhanced performance.

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

The important peripherals are on board.

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

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

Right product. Right price. Right away. **Zilog**

ZILOG SALES OFFICES: CA (408) 370-8120, (714) 838-7800, (818) 707-2160, **CO** (303) 494-2905, **FL** (813) 585-2533, **GA** (404)923-8500, **IL** (312) 517-8080, **MA** (617) 273-4222, **MN** (612) 831-7611, **NJ** (201) 288-3737, **OH** (216) 447-1480, **PA** (215) 653-0230, **TX** (214) 987-9987. **CANADA** Toronto (416) 673-0634, **ENGLAND** Maidenhead (44) (628) 39200, **W. GERMANY** Munich (49) (89) 612-6046, **JAPAN** Tokyo (81) (3) 587-0528, **HONG KONG** Kowloon (852) (3) 723-8979. **TAIWAN** (886) (2) 741-3125, **SINGAPORE** 65-235 7155, **DISTRIBUTORS: U.S.** Anthem Electric, Bell Indus., Hall-Mark Elec., JAN Devices, Inc., Lionex Corp., Schweber Elec., Western Microtech. **CANADA** Future Elec., SEMAD, **LATIN AMERICA** Argentina—Yel.-(1) 46-2211, Brazil—Digibyte (011) 241-3611, Mexico—Semiconductores Profesionales (5) 536-1312.

CIRCLE NUMBER 12

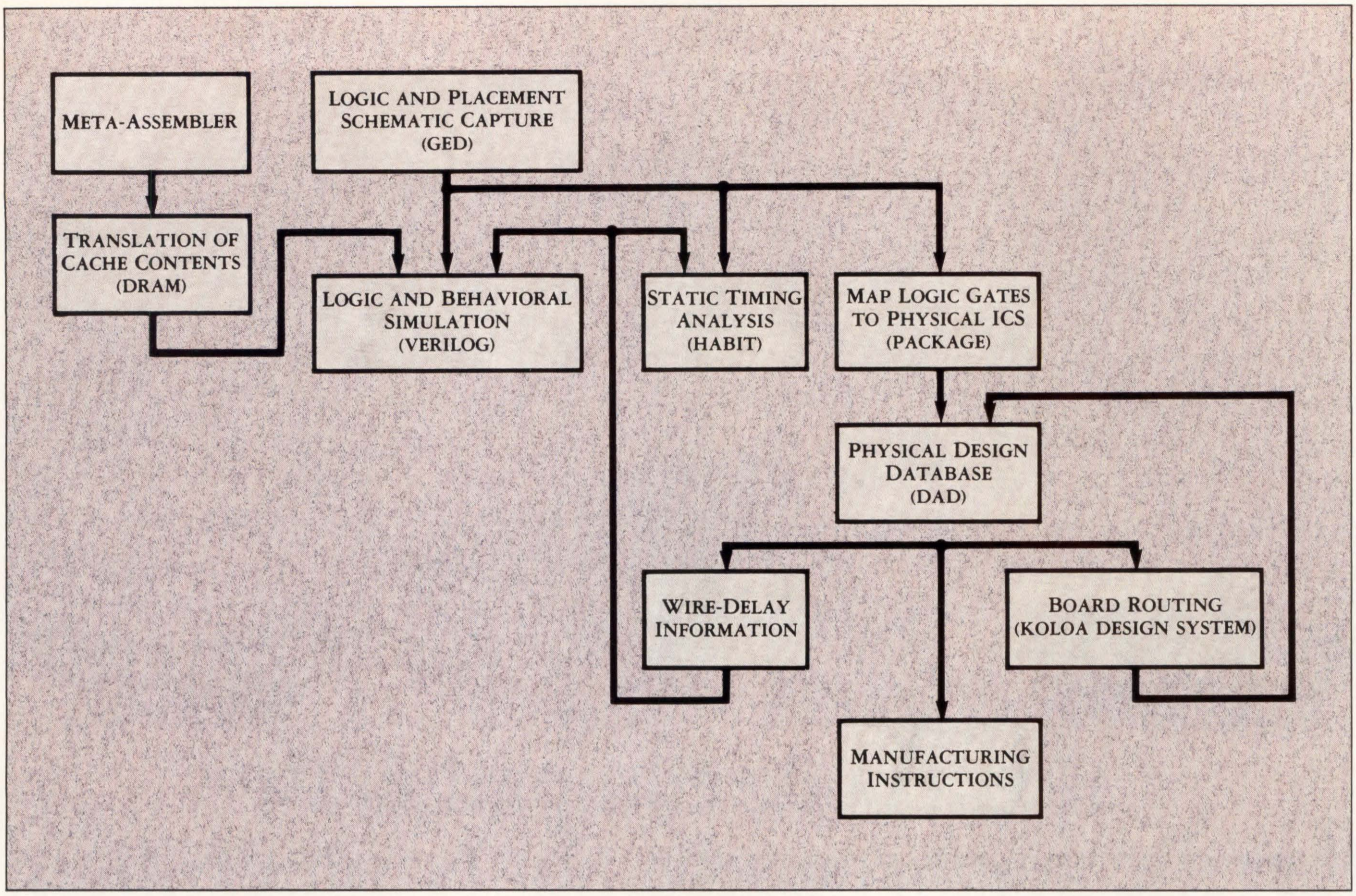

Figure 2. An overview of the board-level design flow used at Elxsi in the design of the 6460 CPU.

Fourth, we needed mixed-behavioral and gate-level simulation capability. The caches and main memory could only be done this way. Also, the clock-generation logic (which is the last logic to get implemented in gates) needs to be done first in behavioral models.

Finally, the ability to do breakpointing and patching was important to us because these imply that the simulator is interactive.

Breakpointing is the ability to stop simulation interactively (such as when a diagnostic begins to fail) and take a look at what happened. Patching is the ability to then change the behavior of the circuit by breaking connections and adding logic, then continue simulating from that point on. The design-loop time to go back to the logic drawings, make a change, recompile the logic, reload the simulator, and get back to the point of failure would be on the order of hours.

■ SIMULATORS EVALUATED

We looked at both software logic simulators and hardware simulation accelerators. The hardware simulation accelerators were all less flexible than the software logic simulators (though it appears much development is in progress to correct this deficiency). None of the hardware simulators supported patching. Most didn't support interactive simulation and breakpointing.

The Valid software simulator was the best choice if it could meet our criteria because we were already using Valid Logic workstations to enter our designs. Integration problems would thus be reduced. However, the speed of the Valid Logic simulator was too slow. A benchmark of an earlier version of our design ran at about six seconds per cycle of the machine, assuming the simulation wasn't paging too heavily (our workstations did not have enough physical memory to run even that fast). It does support breakpointing and patching.

The Valid Realfast hardware simulation accelerator does run fast enough (an estimated 0.5 seconds per cycle of our machine). It also supports breakpointing, but not patching. Though it is capable of behavioral simulation, there is a performance penalty that would result in a longer time per simulation cycle. It is also inconvenient for multiple concurrent users, which is a problem for us, since we planned to be running as many as six simulations of the entire machine simultaneously.

Gateway Design Automation's Verilog software simulator was estimated to run at a rate of 0.25 to 0.5 seconds per cycle of our machine. We could only guess at the effect on performance that reading and writing several megabytes of RAM during simulation could have, since the design wasn't done yet, but this looked promising. Verilog handles both breakpointing and patching and has a very good behavioral language. It also has a C programming interface, which allowed us to put our own user interface on it for doing graphics waveform displays and for letting the user use his own signal and gate names. A Valid-to-Verilog netlist translator was available from a consultant, and we would get the source code for it.

■ THE VERILOG SIMULATOR

We chose the Verilog logic simulator, but it took a lot of work to integrate it into our environment. The netlist translator we got from the consultant was not general enough to handle all of the constructs we made use of in the SCALD netlist language, and so we had to rewrite the translator. Also, we added a feature to generate a comprehensive cross-reference between SCALD names and Verilog names. One complication that arose

Saratoga FIFOs.

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

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

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

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

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

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

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

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

Saratoga Semiconductor

The Technology Leader in BiCMOS

Introducing The First 50-MHz FIFOs. In BiCMOS.

CIRCLE NUMBER 13

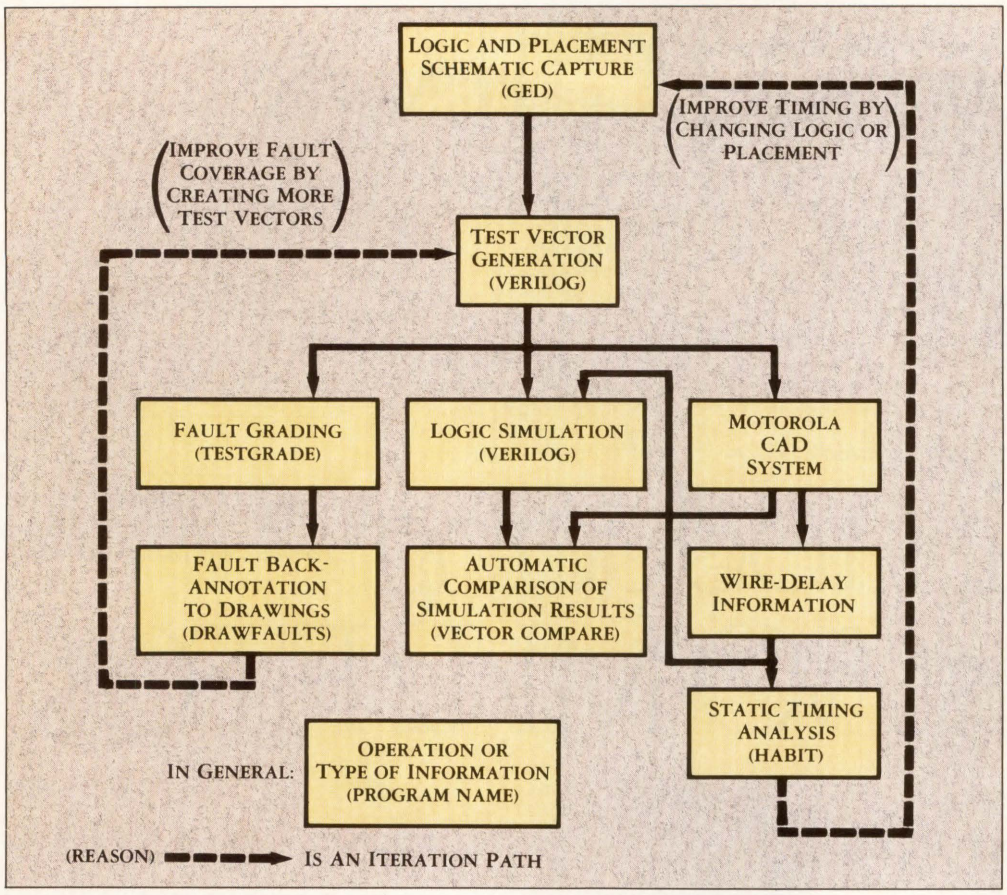

Figure 3. The gate-array design flow for the Pegasus design team is shown above.

was the need for "pin-order" libraries to allow the translator to translate between SCALD connections, which are made by pin name, and Verilog connections, which are made by pin number.

Other complications arose when we tried to speed up the rate of simulation. The Verilog simulator can simulate a model in two different ways: accelerated and non-accelerated. The difference in performance can be a factor of ten, so we spent much time finding out why gates were being simulated using the non-accelerated algorithm and then tried to fix the problem. This problem appears to be unique to Verilog. Once we learned enough about this issue, however, we indeed did get good performance and didn't have any further trouble tracking down this kind of problem. One of the primary causes of the non-accelerated gate problem was the heavy use of RAMs in our simulation (which generated a lot of non-accelerated events because they were implemented with behavioral models). Finally, we achieved an actual simulation rate of about one second per cycle, not including the time to initialize simulation (i.e. the time to read in the netlist and load the RAM contents).

We ran Verilog under the BSD Unix operating system running on our Elxsi 6420. This worked out well for us because our machine had eight CPUs plugged into it at once, giving us plenty of throughput for the six simultaneous simulations of the new CPU design. It also had a gigabyte of shared real memory so that the simulations (which each required over 120 Mbytes of virtual memory) would not page. Using the Elxsi for simulation, however, meant that we couldn't use the graphics interface supplied by Gateway Design Automation. This would only work if we ran the software on a Sun workstation that didn't have enough throughput or real memory. What we finally did was to build our own graphics interface on the Sun that communicated over the Ethernet to the Elxsi machine during simulation. This interface was designed to also understand SCALD signal and gate names for the design being simulated. The interface program reads information (generated during netlist translation) to translate between Verilog and Valid (SCALD) signal and part instance names. It then automatically translates Valid names embedded in Verilog syntax to Verilog names before passing the input to Verilog itself. This was a great advantage to the designers working with Valid schematics during simulation. Other custom features, such as special commands to load, examine, and modify cache RAM contents by field name, made design verification easier. (A given RAM field may be spread across several physical RAM locations. To make this feature work, the designer creates and maintains a cross-reference file between physical RAMs and logical field names that the interface program understands).

We created extensions to the netlist translator to allow it to read delay files generated for both the gate arrays and the board-level signals. It has the ability to combine delay information created separately for each gate array and for each of the two boards into a single comprehensive delay simulation. It is much easier to read the waveforms displayed by the graphics interface when the signals change at approximately the right places in the cycle.

Was all the work worth it? Yes! In fact, the simulation process has caught many errors in the logic (which perhaps would not otherwise have been discovered until a prototype was built). By running actual diagnostics from previous CPU designs, the designers were able to verify the correctness of the instruction set being implemented. This is advantageous because of the simulator's ability to show the value of any signal at any time with ease, using the names on the logic drawings themselves.

■ STATIC TIMING VERIFICATION

A static timing verification program allows the designer to analyze the complete timing correctness of a design, without requiring him to generate test vectors for all the possible conditions. Some of the timing problems that may occur include set-up or hold violations on the inputs to flip-flops or latches, glitches on clock lines, and excessive delay along a critical path.

One reason that timing simulation cannot catch all these problems is that the number of combinations of signal values in the circuit is very large. It would be difficult to try all the combinations to make sure that one of them doesn't result in a timing problem. Another reason is that some race conditions may not show up with a simulation that uses only

PERFORMANCE

16-bit MICROPROCESSOR

16-bit Performance Boost for 6502 Designs

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

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

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

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

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

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

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

CIRCLE NUMBER 14

TABLE 1. PRIMARY TOOLS IN THE ELXSI CAD/CAE SYSTEM

FUNCTION	TOOL(S)	SUPPLIER
DESIGN ENTRY (SCHEMATIC CAPTURE OF LOGIC AND PLACEMENT	GED	VALID LOGIC
GATE-LEVEL AND BEHAVIORIAL SIMULATION	VERILOG VERSHELL,VERGRAPH	GATEWAY DESIGN AUTOMATION WRITTEN INTERNALLY
FAULT GRADING	TESTGRADE DRAWFAULTS	GATEWAY DESIGN AUTOMATION WRITTEN INTERNALLY
STATIC TIMING ANALYSIS	HABIT	WRITTEN INTERNALLY
TEST VECTOR GENERATION	VERILOG + MANUAL EFFORT	GATEWAY DESIGN AUTOMATION
PACKAGING	PACKAGE	VALID LOGIC
BOARD ROUTING	KOLOA DESIGN SYSTEM (MANY PROGRAMS)	SHARED RESOURCES
BOARD-DELAY AND ECL TERMINATION CALCULATION	GENDELAYSANDTERMS	WRITTEN INTERNALLY
GENERATION AND MAINTENANCE OF MANUFACTURING INSTRUCTIONS	DAD (MANY PROGRAMS)	WRITTEN INTERNALLY

maximum (or minimum) delays—it may only show up with the right combination of delays through different components on the board.

Static timing verification gets around these problems by doing a special type of design analysis and displaying the results for the user. Some timing-verification programs are interactive and allow the user to set up a state in the circuit (by forcing some signals to a logic high and others to a logic low) and then requesting the critical path delay between two points.

Unfortunately, all commercially available tools that perform this function place severe restrictions on the design methodology, that is, the design must be synchronous. Elxsi's Pegasus CPU design does not fit within that clocking paradigm, and existing programs did not fit our needs. While some timing-verification tools allow limited gating of clocks (Valid Logic's timing verifier allows gating of clocks with AND and OR gates), the Pegasus design uses a much more complicated clock generation scheme. The solution was to write our own timing verification tool with the required extra capabilities.

We started with an existing Pascal program (this same code was the predecessor to the timing verifier supplied by Valid Logic). It took about a month to get familiar with the old program, and then we began to enhance the algorithms to understand our clocking paradigm. We modified it to read the output of the Valid compiler directly, and to read delay files generated for the signals in the gate arrays and at the board level. The program was also transformed from a batch-oriented program to an interactive one which allows the user to ask for information about critical path timing through the circuit from point to point. It also provides a good interface for browsing the logic and examining delays. The user may interactively force the timing behavior of inputs or other signals and cause the tool (called Habit) to re-evaluate the circuit.

This program proved to be very valuable in analyzing the performance of the gate array designs and making changes to them to speed their performance. The program only takes about a minute to run for a 2,500-gate array. Re-analysis of the circuit takes only about 30 seconds after a change is made in the input timing specification for the gate array. It takes about 55 minutes to load the entire CPU and another 15 minutes to read delay information in. The last phase of the logic design was to verify correct timing behavior at the board level with this tool.

■ AUTOMATIC TEST GENERATION

We wanted an automatic way to generate test patterns for the gate arrays since there were 15 different arrays. The only tool on the commercial market we could find was the Testscan automatic test generator supplied by Gateway Design Automation. Testscan assumes a scan-based design, but our design is not scan-based. So our idea was to utilize the scan-in pattern and scan-out pattern suggested by Testscan for each test. We loaded our circuit with the state it specified using our own means (though we don't use scan design, it is generally very easy to load the design with any arbitrary state and then to read it back out again).

Unfortunately, even though Testscan is supplied by the same vendor that supplies the Verilog simulator, the two programs require different libraries! Moreover, the netlist format for Testscan is much more limited than the format for Verilog. It doesn't support buses (i.e. vectored signals) and it does not allow long names. Also, the program does a check to make sure it can understand the circuit by analyzing how all the clocks are hooked up. Unfortunately, most of our gate arrays violated these rules and we needed to delete logic from the netlist input to the program to get it to generate tests for the rest of the gate array.

We decided not to use Testscan because of all these problems. It was then that we found a manual test-generation method aided by programs written in the Verilog behavioral language that proved to be effective and easier to use. Verilog provides random signal-value generators that we used to provide inputs to data paths. We then specified the correct clocking for the circuit and that combination of data and clock inputs does a good job of testing the gate array.

■ FAULT GRADING

We wanted to know how much fault coverage was needed to test the gate arrays and how much coverage we were actually getting. To determine how much was required, we went back to old ECL gate array

TABLE 2. TRANSLATORS IN THE ELXSI CAD/CAE SYSTEM

NAME	TYPE OF DATA	FROM	TO
VCMP (1) DAD PROGRAMS	NETLIST + DELAYS NETLIST + PLACEMENT	GED GED	VERILOG DAD PHYSICAL DESIGN DATABASE
PLACETOFIXFILE + GLOGCAP (2)	NETLIST + GATE-ARRAY PLACEMENT	GED	MOTOROLA (LOGCAP)
VECTORTOMOTO	TEST VECTORS	VERILOG	MOTOROLA (LOGCAP)
DELAY	GATE-ARRAY DELAYS	MOTOROLA	VCMP + HABIT
DRAM	RAM (CACHE) CONTENTS FOR SIMULATION	INTERNAL META-ASSEMBLER	VERILOG

NOTES: (1) VCMP WAS SUPPLIED TO US BY A CONSULTANT AND WE EXTENDED IT SUBSTANTIALLY
(2) GLOCAP IS SUPPLIED BY VALID LOGIC

ALL OTHER TOOLS WERE WRITTEN INTERNALLY.

designs and checked the fault coverage of the tests used for them. We correlated that coverage with the actual failure rate of the gate-arrays during system testing of boards that they had been plugged into.

We then had to choose a fault-grader. LOGCAP had been used in-house to do fault-grading on the old gate arrays. However, it only allowed faults to be simulated on the outputs of gates and it ran extremely slowly. A 2,500-gate array with about 500 test patterns would take overnight to run. Some of the gate arrays had over 2,000 test patterns. The Testgrade fault-simulation program allowed coverage of all stuck-at faults in our models and ran much faster, completing the 500 test patterns in less than one-fourth the time of the LOGCAP fault grader (this meant that we could run up to three fault-grading sessions in a 24-hour period for the same array).

To make it easier to generate additional patterns for covering faults that were missed, we wrote a program (Figure 1), back-annotating the logic drawings. Markings in the program showed which signals had faults that were missed. Often, the person generating test patterns could look at the drawings and see the pattern of what wasn't covered and thus could more easily figure out what to do. This was something that couldn't be done by looking at a list of the names of uncovered faults.

■ CONCLUSIONS

Integration of design automation tools *is* a big deal. It requires the dedication of people, time, and money. However, the payback can also be large. We discovered that integration cannot be solved by simply buying all the required tools from a single vendor. No single vendor meets all the needs of a complex design methodology such as ours. Many of the needed tools are not available from any vendor.

Two constraints work against each other during the integration of a design automation environment. On one hand, there's usually not enough time to hire CAE people to build tools internally, implying the need to buy from tool vendors whenever possible. Also, building all the tools internally can be expensive for a small company.

On the other hand, the use of vendor-supplied tools has problems. They are not supported as well as internally written software would be, as measured in terms of time from bug report to bug fix and also in getting the functionality provided to match the needs—including timely enhancements. What's more, the use of vendor-supplied tools can also lead to a more poorly integrated system.

Some vendors made it easier to integrate their software with other software. This was true of Valid Logic and Gateway Design Automation. Both made extensive use of user-readable file formats that can be read and generated by tools written internally.

We also discovered that debugging newly integrated tools places an extra demand on logic designers. However, these tools are now in place now and there are many people experienced in their use. The final board-level and gate array design flows adopted at Elxsi are shown in Figures 2 and 3 respectively. ■

ACKNOWLEDGMENTS

The author would like to thank Nick Whyte for his helpful comments during the writing of this article.

ABOUT THE AUTHORS

JAMES BOHANNON, *who joined Elxsi Corp. (San Jose, Calif.) in 1987, is the manager of its CAD Software Development group. Before joining Elxsi, he was a member of the CAD software development group at LSI Logic Corp. He received his BS in mathematics and computer science from the University of California at Los Angeles in 1982, and is currently studying for an MBA.*

SOME THINGS JUST AREN'T WORTH THE RISC.

$174*
AMD RISC Chip
16 MHz, 12 MIPS

$495**
Motorola RISC Chip
20 MHz, 17 MIPS

$179*
LSI Logic SPARC Chip
16 MHz, 12 MIPS

$174*
Intel RISC Chip
20 MHz, 7.5 MIPS

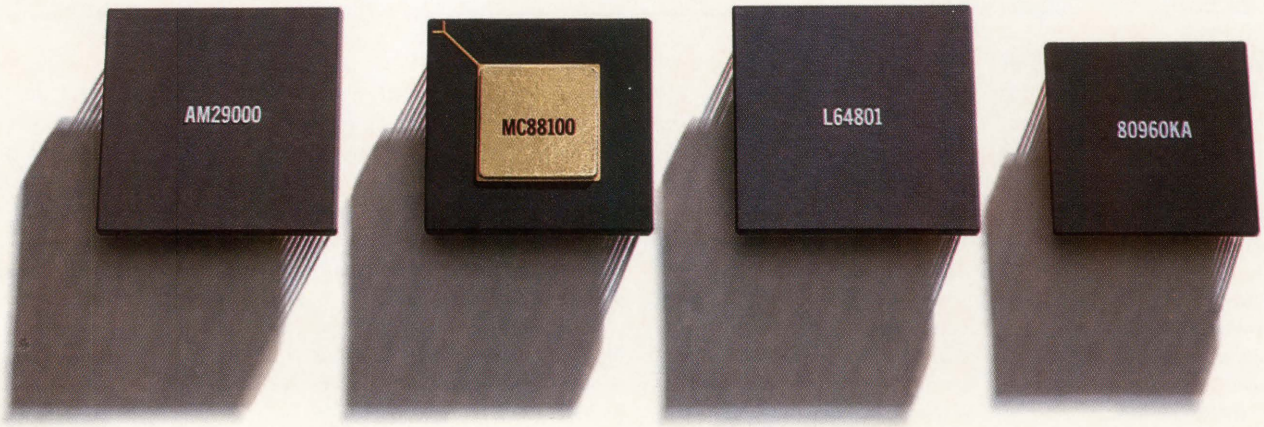

Any chip above would do a good job as an embedded controller in a disc drive, laser printer, or workstation.

It's just that one can do it for a lot less money. The VL86C010 32-bit RISC chip (ARM) from VLSI.

It gives you the performance you need right now, in volume, for about $5 per MIPS. (The nearest competitor is $10 per MIPS. If and when the product is available.)

YOU NOT ONLY SAVE CASH, YOU SAVE CACHE.

Because our chip is not only less expensive, but easier to use than most RISC chips, you save a bundle on your entire system.

You don't have to pay for fancy memory schemes like cache or full static memory. Our RISC chip only needs 80ns DRAMs to run at a full 12 MHz.

You save on real estate, too. You can design an intelligent controller with 4 Mbytes of DRAM that will fit on a standard XT expansion card.

YOU'LL THANK US FOR OUR SUPPORT.

Nobody else in the industry offers the

AND SOME THINGS ARE.

$50*
VLSI RISC Chip
12 MHz, 6.33 MIPS

peripheral support we do. Nobody.

Ours is the only RISC chip with off-the-shelf peripheral support for memory, video, audio, and other I/O functions.

Don't want to mess with microcode? That's okay. ANSI C, FORTRAN 77, and an Assembler/Linker will do very nicely, thank you.

And we're working on development support tools like industry standard real-time executives and in-circuit emulators.

But we saved the best for last. Because VLSI is a leader in ASIC and always has been, we're now using our RISC as an ASIC core.

That means you can build your very own RISC system quickly and easily using VLSI design tools.

Call 1-800-872-6753 and ask for our free VL86C010 brochure. Or write to us at 8375 South River Parkway, Tempe, AZ 85284.

And reduce your RISC.

*Prices based on: EDN, June 9, 1988; Electronics, April 14, 1988; Microprocessor Report, May and September 1988; Quantities of 100. **Price based on sample quantities.

VLSI TECHNOLOGY, INC.

CIRCLE NUMBER 15

A SINGLE CHIP MODEM FRONT END

RAOUF HALIM AND DANNY SHAMLOU, HAYES MICROCOMPUTER PRODUCTS INC., NORCROSS, GA.

CLOSE VENDOR USER RELATIONSHIPS ARE A MUST

One of the keys to successful in-house design of mixed analog/digital ASICs by system houses is a close liaison with a silicon vendor offering state-of-the-art computer-aided engineering tools. Vendor-supplied libraries that accurately model the vendor's process in addition to supporting the required design styles also play a key role.

Another key is a high level of custom chip design expertise on the part of the system house (tool users). However, as mixed analog/digital tools advance, the required level of "chip" expertise will decrease. Examples of such next-generation tools are analog module generators, general easy-to-use synthesis tools, cell compilers, and high-level and mixed-mode simulators.

Hayes Microcomputer Products Inc. has successfully designed a single-chip analog front end (AFE) for international 2400/1200/300-bps modems. It did so by using a CAE system (Testa and Mittal) and library cells internally developed and supplied by Silicon Systems Inc. (SSI). Design time from concept to first-pass successful silicon (Figure 1) was eight months.

The AFE chip incorporates 80 poles of high-performance analog switched-capacitor filtering, a mixed analog/digital tone-generator block, tone-detection circuits, modulator, gain stages, an 8-bit analog-digital converter, and digital microprocessor and DSP interfaces. The chip is approximately 62,000 square mils in a 3-micron double-poly CMOS process. It was designed primarily for modem performance improvement, particularly in the V.23 1,200-bps mode.

■ CAE System Overview

The CAE system from SSI was made available to Hayes in an arrangement whereby Hayes operated as a remote design center to SSI. The CAE system (Figure 2) is an integrated design toolbox developed specifically for design of mixed analog/digital chips. It accommodates design strategies ranging from full-custom to standard cells and predesigned building blocks, in both CMOS and bipolar technologies. The system runs on Apollo Computer Inc. workstations.

The front end (schematic capture) for the electrical design phase consists of a Mentor Graphics Corp. workstation with SSI-developed tools. These tools are comprised of custom menus, entry macros, checking utilities, and component libraries.

The available component libraries are: custom components (including primitives such as capacitors and transistors); standard cells with fixed-height layouts; and predesigned building blocks such as filter stages and A/D converters.

After completion of schematic entry, the performance and functionality of the design are evaluated using simulation and analysis tools to verify that it meets all requirements. SSI-developed hierarchical netlisters automatically generate design netlists and user-defined stimuli from the single schematic database. Formats are compatible with the various third-party simulators, such as HSPICE, SWITCAP, and SILOS.

Chip floor-planning, layout, and checking are entirely carried out by SSI, including layout-to-schematic verification (circuit trace). Given SSI's experience in physical design of full-custom mixed analog/digital chips, this was determined to be the optimum approach to the "back-end" IC design phase.

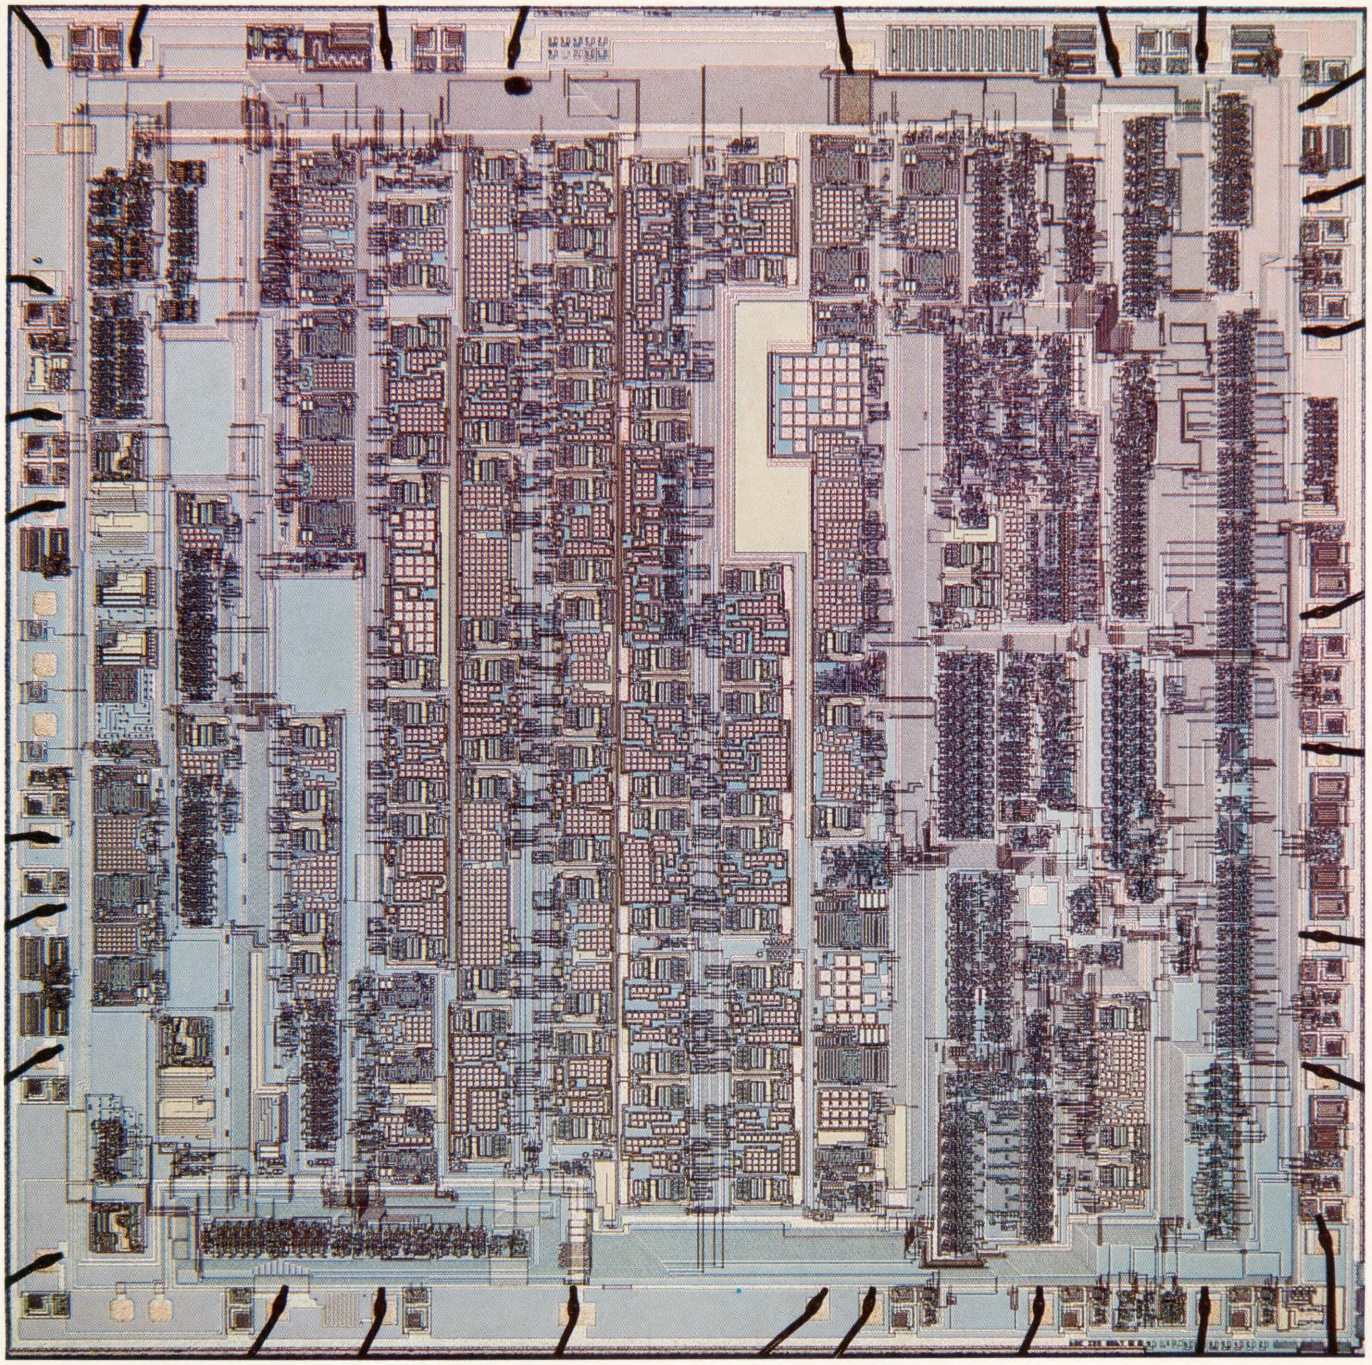

Figure 1. A die photograph of the modem analog front end.

■ ARCHITECTURE

The architecture of the 2,400/1,200/300-bps modem is shown in Figure 3. In this architecture, the AFE chip functions as a peripheral to the microcontroller, performing all the front-end analog-signal conditioning and processing for the modem. All data transfer and signaling to and from the AFE is digital.

The microcontroller performs data-terminal-equipment to modem-command interpretation, data buffering, scrambling and descrambling, as well as call setup and programming of the AFE chip in the appropriate configuration. It communicates with the AFE chip via a 4-bit multiplexed address/data bus. All AFE control and status information is stored in 11 4-bit registers that are accessible to the processor.

The DSP demodulates the signal, adaptively equalizes it, then decodes it according to the constellation of the selected mode. It communicates with the AFE via a bidirectional serial port.

The modem is designed for international applications and meets the requirements of most foreign countries. It implements four full-duplex modem communication standards (a "QUAD " modem), as recommended by the CCITT :

■ V.22 bis: 2,400 bps using QAM modulation and compatible with V.22 mode;

■ V.22: 1,200 bps with fallback to 600 bps using DPSK modulation;

■ V.23: asymmetrical channel frequency division with up to 600/1,200 bps in the main channel, and up to 75 bps in the back channel using FSK modulation. V.23 is used primarily for videotext applications; and

■ V.21: up to 300 bps using FSK modulation.

The AFE architecture is shown in Figure 4. It makes extensive use of macrocells and building blocks developed by SSI for other chips, including a V.22 bis AFE chip

More gates. More More memory.

NEC's 60,000-gate standard cell gives you maximum design flexibility.

C&C Computers and Communications

For fast answers, call us at:

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

speed.

Maximize your design options with NEC's newest standard cell. The 60,000-gate SC-5 is a complete system on silicon. Fabricated with a 1.2μ CMOS process, it gives you higher integration, greater gate speed, and the largest memory on the market. Memory access and drive capability are tops in the industry.

Check the specs item by item. You'll see that the SC-5 edges the competition in crucial performance parameters, allowing you to make bold new departures in design.

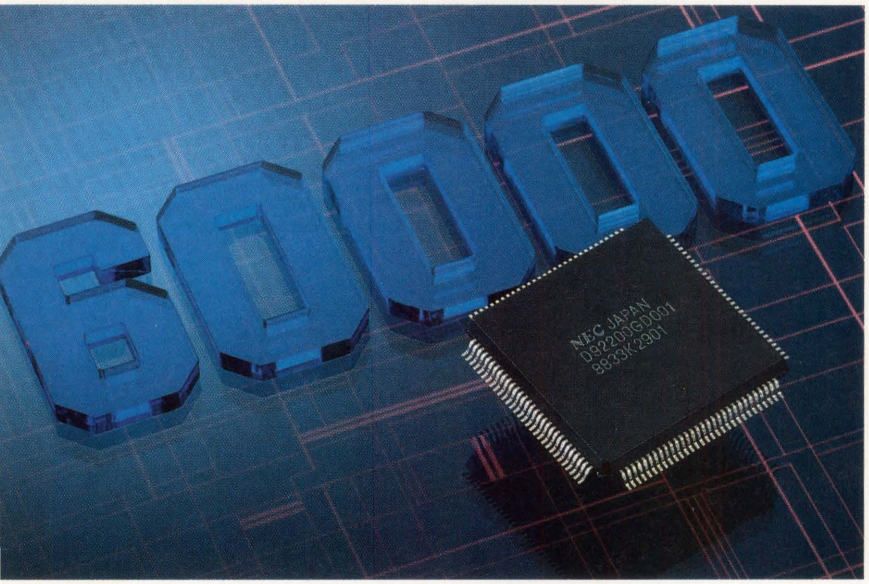

- ☐ Integration: Up to 60,000 gates (2-input NAND).
- ☐ Internal gate delay: 1.0ns (F/O=3, l=3mm).
- ☐ Memory access*: 22ns RAM; 15ns ROM. ** for 256 x 8 organization.*
- ☐ Max memory size: 128K-bit RAM; 1M-bit ROM.
- ☐ Output drive capability: Up to 48mA.
- ☐ Packages: DIP, QFP, PLCC, PPGA, and PGA.

NEC offers users of the SC-5 the benefit of an extensive macro library.

- ☐ Functional blocks: Functionally compatible with our CMOS-5 gate arrays.
- ☐ Soft macros*: 74 LS series-compatible macros.
- ☐ Mega macros*: CPU peripherals including multipliers, adders and FIFO.
- ☐ Analog macros*: A/D and D/A converters, comparators, OP amps, analog switches and more.
- ☐ I/O blocks: Compatible with our CMOS-5 gate arrays, including CMOS/TTL-level I/O, and pull up/down resistors.

** under development.*

The purpose of a standard cell is to help designers achieve enhanced performance with reduced effort and expense. To find out how the SC-5 increases your design flexibility, call NEC today.

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

NEC

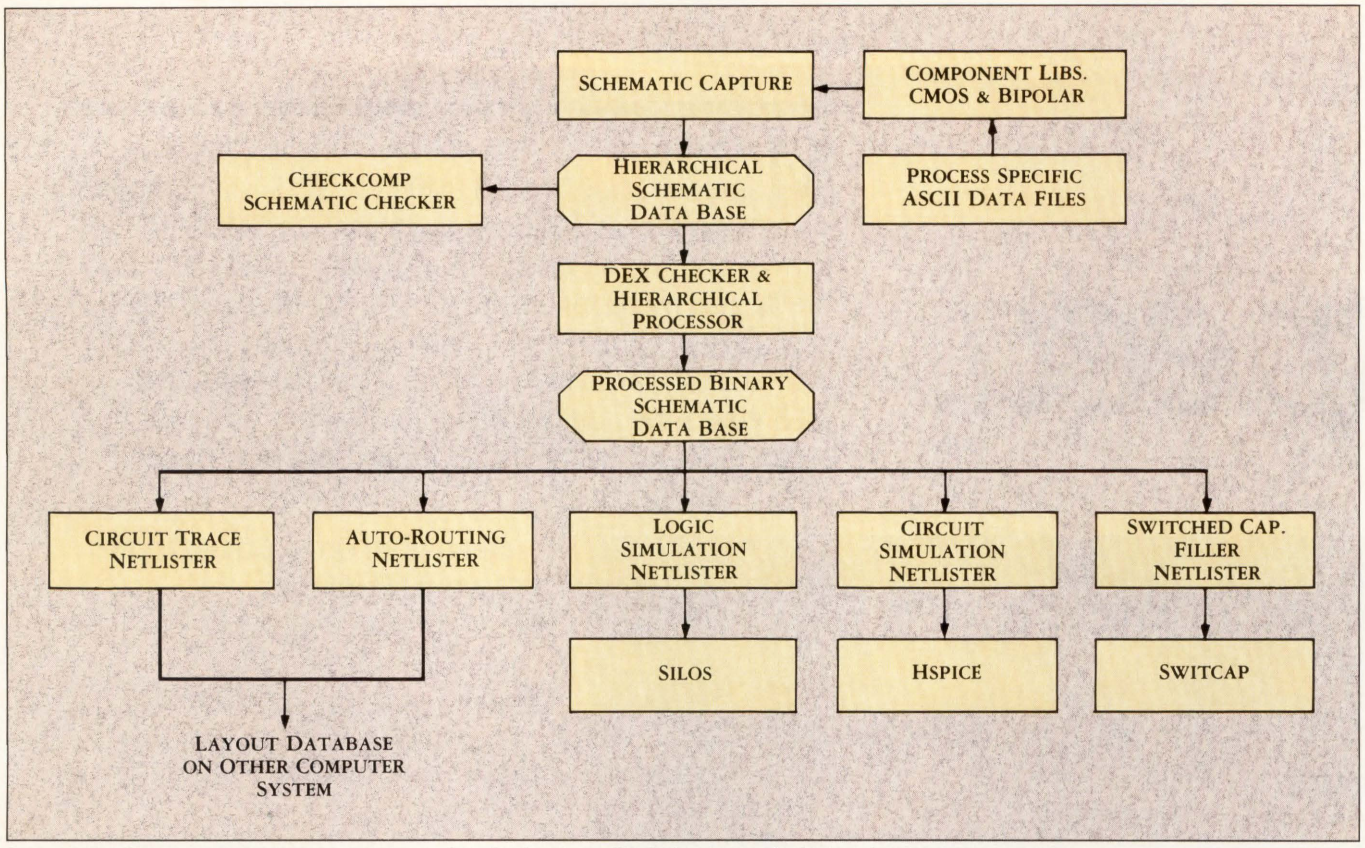

Figure 2. The CAE system is an integrated design toolbox developed specifically for design of mixed analog/digital chips.

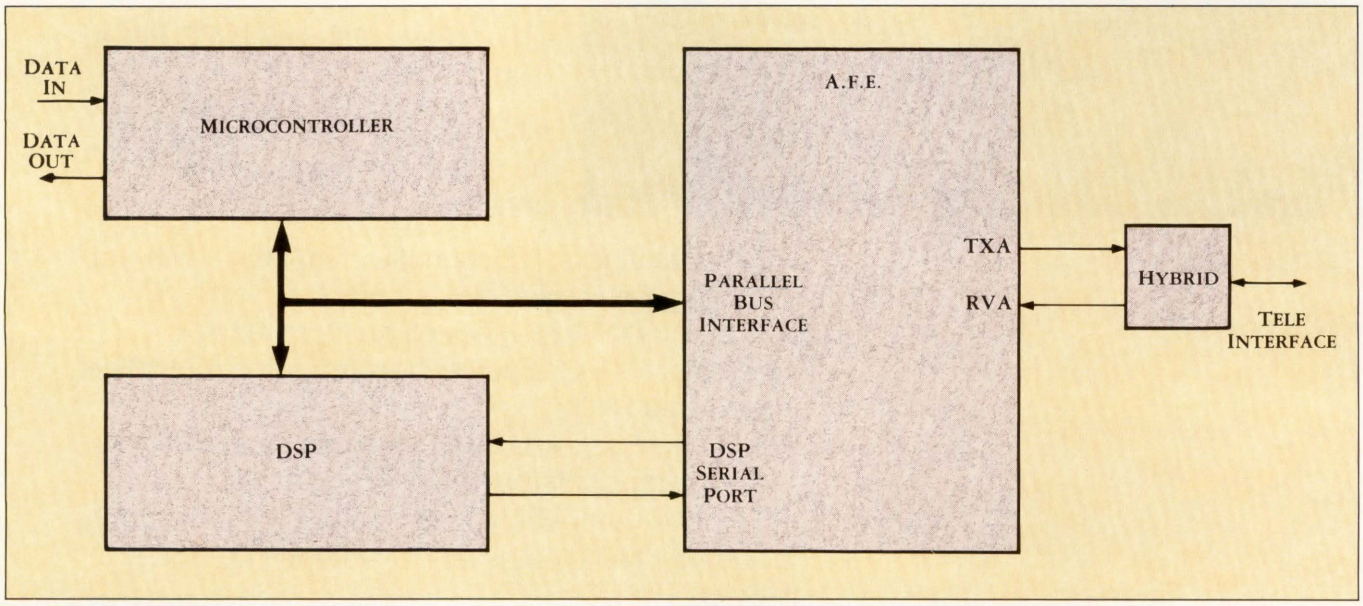

Figure 3. The architecture of the 2,400/1,200/300-bps modem; the AFE funcitons as a peripheral to the microcontroller.

(Hurst et al). On the transmit side, the AFE encodes and modulates the data from the controller in QAM and DPSK modes. In FSK modes the tone generator block acts as the modulator, as well as generating guard, DTMF, and answer tones. The passband signal is next shaped by an anti-aliasing filter followed by the transmit bandpass filter for the selected mode. Guard tones are next summed with the signal (if enabled), followed by a switched capacitor programmable attenuator to set the transmit signal power. An output-smoothing filter eliminates high-frequency images generated by the switched capacitor filters, completing the transmit signal processing.

On the receive side, the signal first passes through an anti-aliasing filter with selectable boost (0/6/12 dB), followed by the main receive bandpass filter (BPF) for the mode selected. The BPF provides opposite band rejection, as well as compromise levels of amplitude and phase equalization based on an average phone line. The signal is next passed through a second anti-aliasing filter, followed by a programmable switched-capacitor gain stage controlled by the DSP. The signal is then converted to an 8-bit twos-complement representation by an algorithmic A/D, and the data is passed to the DSP.

Carrier, answer tone, and call-progress tone detectors are also included in the receive path. These are energy detectors with current states that are stored in status

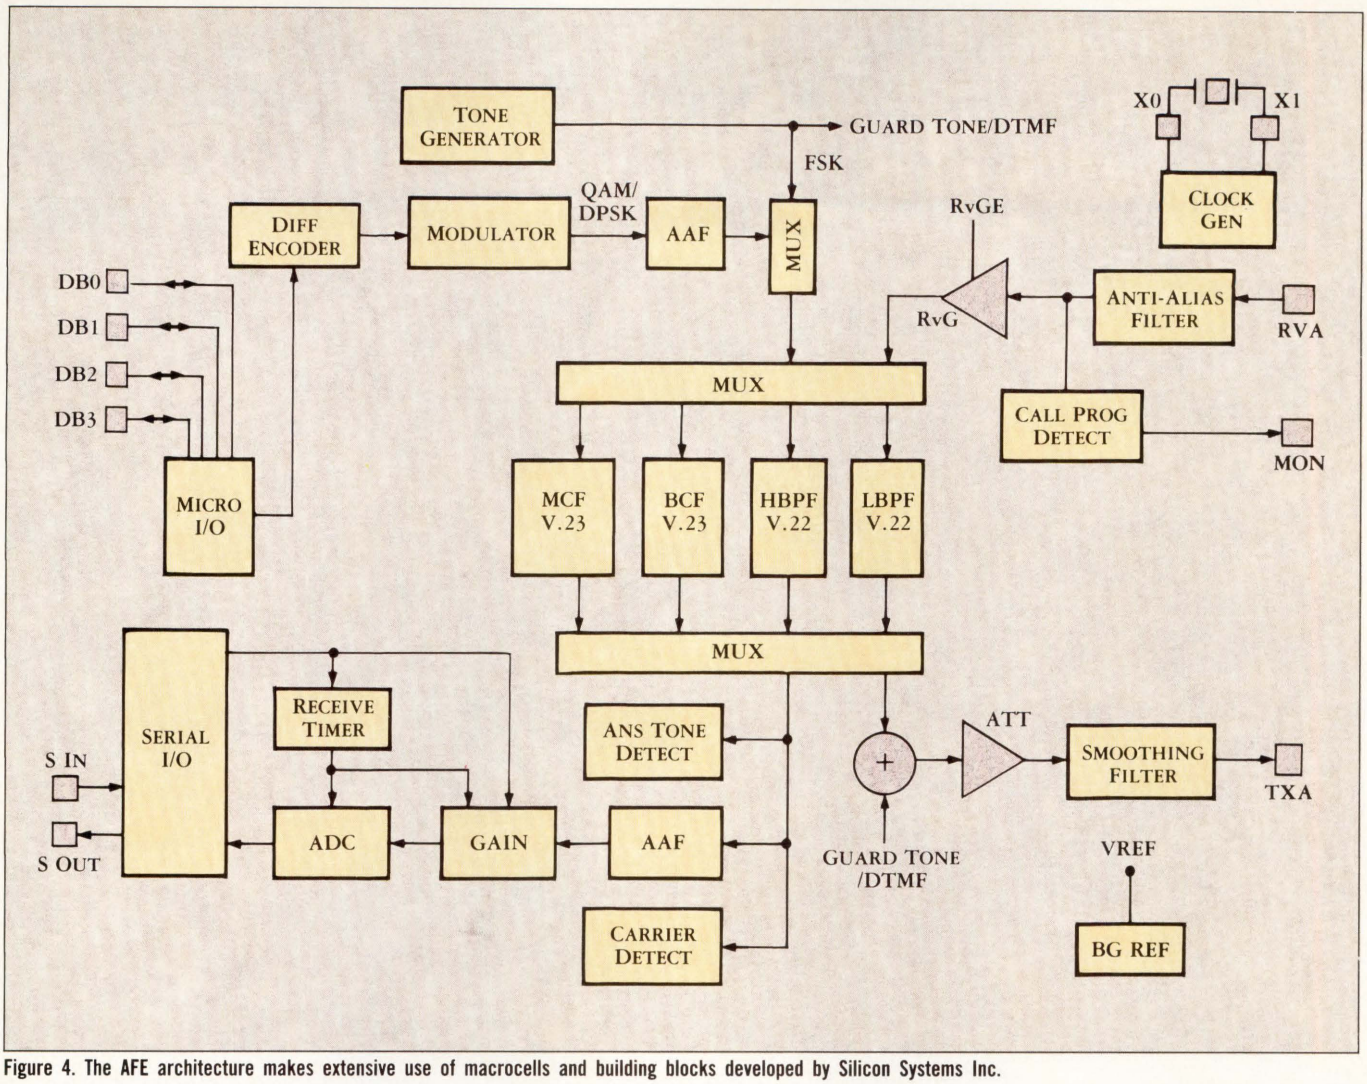

Figure 4. The AFE architecture makes extensive use of macrocells and building blocks developed by Silicon Systems Inc.

registers and can be read by the microcontroller. In addition, a programmable timer is included for clock recovery, as well as a programmable four-level gain stage for off-chip audio monitoring of call-progress tones.

Band-split filters for V.22/V.22 bis were incorporated as predesigned building blocks. Filtering for V.21 mode is achieved by a 25 percent reduction in the V.22 filter clock rate in order to pass V.21 FSK tones.

The majority of the design effort concentrated on new high-order switched capacitor bandpass filters for V.23 mode, more fully described in the next section.

■ Filter Design Methodology

The V.23 filters are composed of two bandpass filters, one for the main channel and the other for back-channel filtering. Their responses are shown in Figure 5. One is connected on the transmit (tone generator) side and the other on the receive side depending on AFE answer/originate setup.

The filters' primary function is to reject the opposite band transmitted signal leaking through the hybrid. They also handle channel anti-aliasing and reduce the noise bandwidth. In addition, the main channel filter incorporates a fixed delay equalizer, as well as compensation for channel gain roll-off. Also, the filters should introduce minimum amplitude or group delay distortion.

To further complicate the filter designs, worst-case signal-to-filter noise at the filter output must be at least 20 dB. For a signal level of -39 dBm (which is boosted from -45 dBm by a hybrid gain of 6 dB), this implies filter noise levels of roughly 800 microvolts.

Due to the conflict between the high order dictated by filter response requirements and low noise required for receiver dynamic range, much care was taken in the design of the filter transfer functions, as well as in their synthesis and implementation, to ensure that they met all specifications. The filter design methodology is shown in Figure 6.

The filter transfer functions were synthesized from pass and stopband specifications as cascaded second order sections using FILSYN.

The main channel filter is a 12th-order bandpass synthesized as the cascade of a 7th-order high-pass and a 5th-order low-pass filter. The 12th-order magnitude section is followed by a 4th-order all-pass delay equalizer that compensates for both filter and channel group delay variation over the pass band. The main channel filter has two modes of operation: a 1200-bps and a 600-bps (half-speed) mode. In the 1200 mode, the center frequency is 1700 Hz with a bandwidth of 1400 Hz, while in its half-speed mode it is centered on 1500 Hz with a bandwidth of 1000 Hz. The back-channel filter is a fixed 8th-order bandpass. Its center frequency is 420 Hz with a bandwidth of 140 Hz. This ensures symmetry around the different V.23 mark and space frequencies, with bandwidth limiting for optimum receiver performance.

In the top-level partitioning of the V.23 filters, the different modes of V.23 main channel filtering (half-speed and amplitude equalization) are achieved by capacitor programming of the low-pass section of the main channel filter. The delay

Get the Oki Advantage

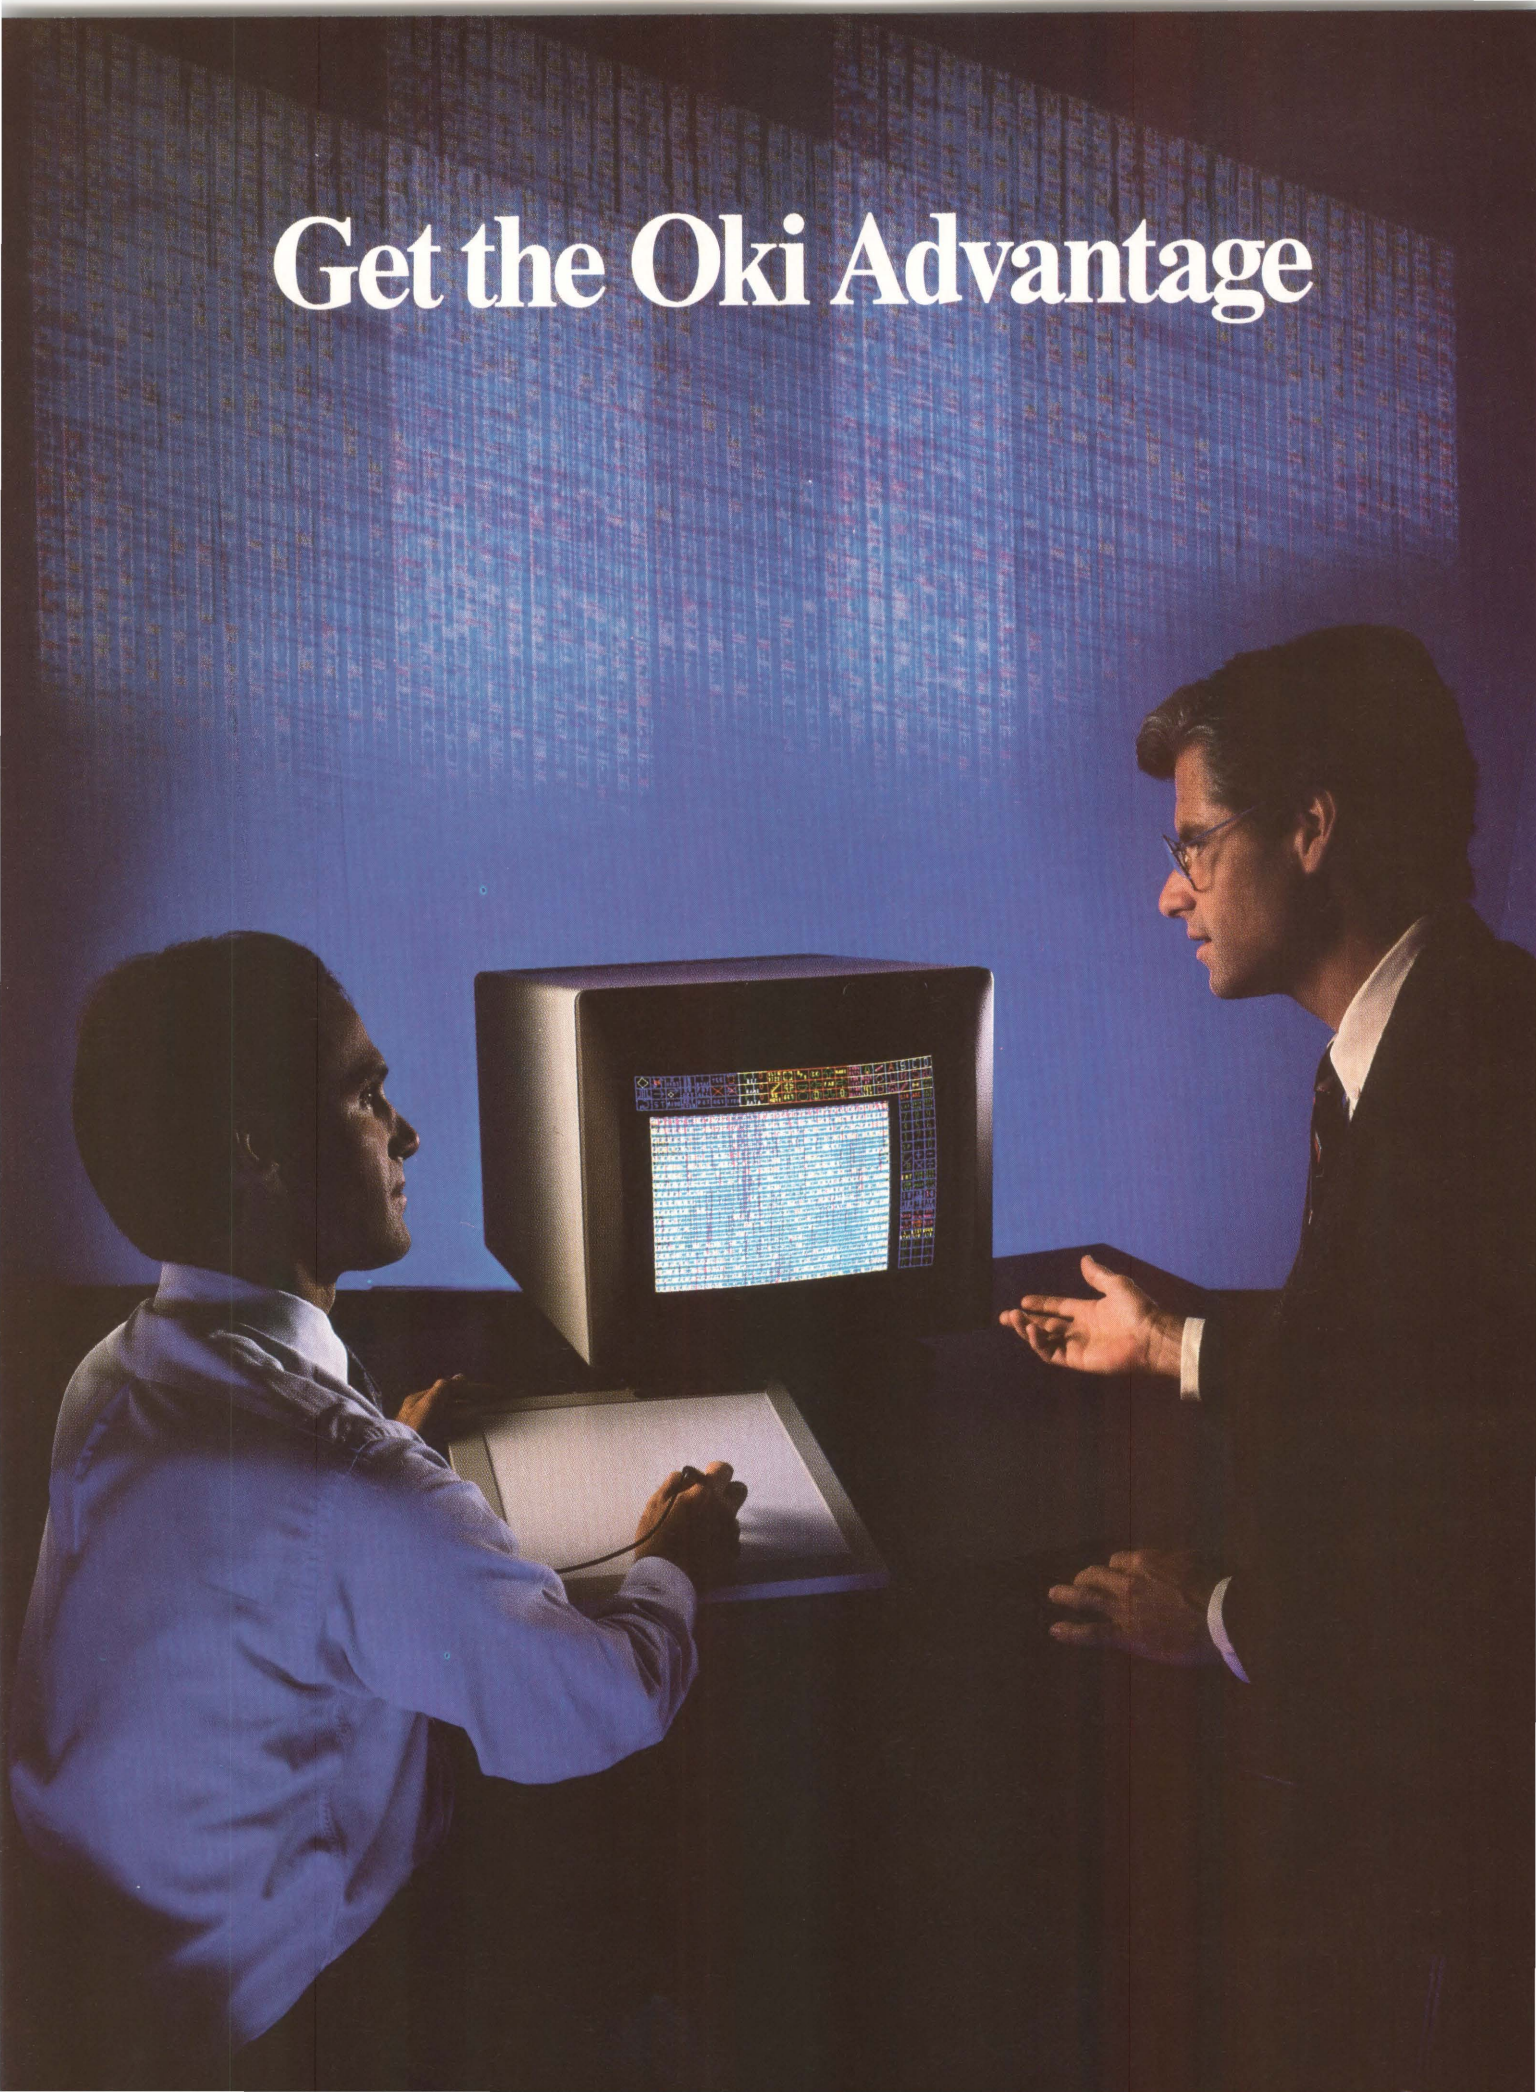

Everything you need for ASIC success from one reliable source

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

Advanced ASIC products and technologies

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

- *Sea-of-gates:*

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

- *Channelled array:*

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

- *Standard cell:*

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

ATG and logic transparency

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

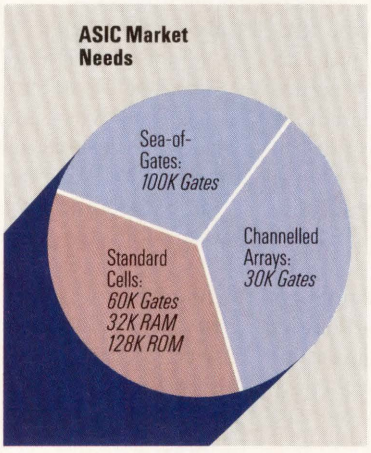

Complete ASIC support

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

Check Oki: your complete ASIC resource.

VLSI 12/88

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

( )
Phone

Name

Title

Company

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

OKI SEMICONDUCTOR

CIRCLE NUMBER 16

equalizer section can be independently bypassed. All programming of the main channel filter is under control of the processor via control register bits. The back channel filter is a fixed structure requiring no programming.

The filters are realized as cascaded biquads of the Fleischer-Laker topology (Fleischer and Laker). Capacitor values for each biquad were computed from its z-domain transfer function, then scaled for the signal swing of the internal operational amplifier node using a BQSYNTH program. Once that was done, the capacitors were then normalized to a unity capacitor of 0.42 pF (30 × 30 microns). Afterwards, the dimensions of all of the fractional pieces were computed in order to establish close to constant area-to-perimeter ratios using the BQCAP program. A clock rate of 38.4 KHz was determined to be acceptable in order to provide a good trade-off between the over-all capacitance and the amount of noise folded by the switched capacitor stages.

Standard CMOS operational amplifiers and switches from the component libraries were used. Array sizes and dimensions for precision capacitors (poly/poly) were specified as capacitor properties on the schematic.

For each filter, a large number of different biquad orderings, in addtion to pole-zero combinations, were synthesized. For each synthesis, HSPICE was used to estimate the filter noise floor. This was accomplished by means of a program (BQ2SPICE) that automatically generates an equivalent resistive-capacitive filter as a noise model in SPICE-compatible syntax. This model also incorporates effects of 1/f, kT/C, thermal, and folded noise in the switched capacitor filter (Fischer). This process was iterated to minimize noise while maintaining acceptable capacitor ratios for efficient silicon area utilization. Another program (BQMONTE) was used to perform Monte Carlo simulation of effects of random capacitor variation on filter gain and phase response.

In addition to noise, a large variety of simulations were run to verify filter functionality and performance. SWITCAP was employed in order to simulate filter gain and phase response, signal swing at each op amp's output, dc offset gains, transient response, and the effect of low op-amp gain. HSPICE was used to simulate filter gain response, op amp offsets, and worst-case settling time on a single clock phase.

■ MODIFIED BLOCKS

While the majority of macro-blocks from SSI's library were incorporated in the AFE "as-is," several blocks needed modification to support V.23 mode, as well as to comply with foreign Post, Telephone, and Telegram (PTT) requirements.

TABLE 1. AFE MEASURED CHARACTERISTICS

TECHNOLOGY	3 MICRON 2P CMOS
POWER SUPPLY	+/- 5 V, 350 mW
BPF NOISE:	
V.23 MAIN	‹ 450 MICROVOLTS
V.23 BACK	‹ 280 MICROVOLTS
V.22/V.21	‹ 700 MICROVOLTS
BPF DISTORTION	› 60 dB S/THD
A/D CONVERTER	8 BITS +/-0.5 LSB

The tone generator block was modified to generate V.23 mark/space tones by picking new counts from a programmable counter. The receive analog-to-digital timer was modified to guarantee that there was an additional sampling rate specifically for receiving in V.23.

Another area of modification was the SC gain stage. For reasons of blocking dc offsets, the gain stage incorporates two first-order high-pass stages, one at its input and one at its output.

Since dial tones for some countries can be as low as 150 Hz, these high-pass stages were redesigned to move their composite corner down to approximately 200 Hz. For those countries, dial tone detection is implemented by bypassing the receive band pass filter and performing the filtering and detecting of the signal entirely in the DSP.

Since PTT restrictions on out-of-band energy emissions are tighter than FCC guidelines,new SC third-order anti-aliasing and smoothing filters were designed. These provide 55-dB suppression of the first spectral image generated by the 38.4-KHz SC clock.

■ RESULTS AND CONCLUSIONS

The AFE chip was tested both in a complete modem system and stand-alone to verify its functionality and performance. The AFE was fully functional and exceeded all performance requirements on the first silicon pass. Total development time was eight months, evenly distributed between definition/specification, electrical circuit design, physical design, and fabrication phases. Table 1 summarizes the AFE measured characteristics.

The recipe for successful design of mixed analog/digital ICs by a systems house relies heavily on the vendor-supplied CAE tools. In addition, it is important that the cell libraries are well documented and supported. There should also be an effective partitioning of responsibilities from device definition through electrical and physical design. Some level of software problems should be anticipated as well on the user side. In many cases, the vendor's tools may be exercised in new areas for the new design that uncover bugs. This makes vendor tool training and support an important issue. Finally, both sides must work closely together during

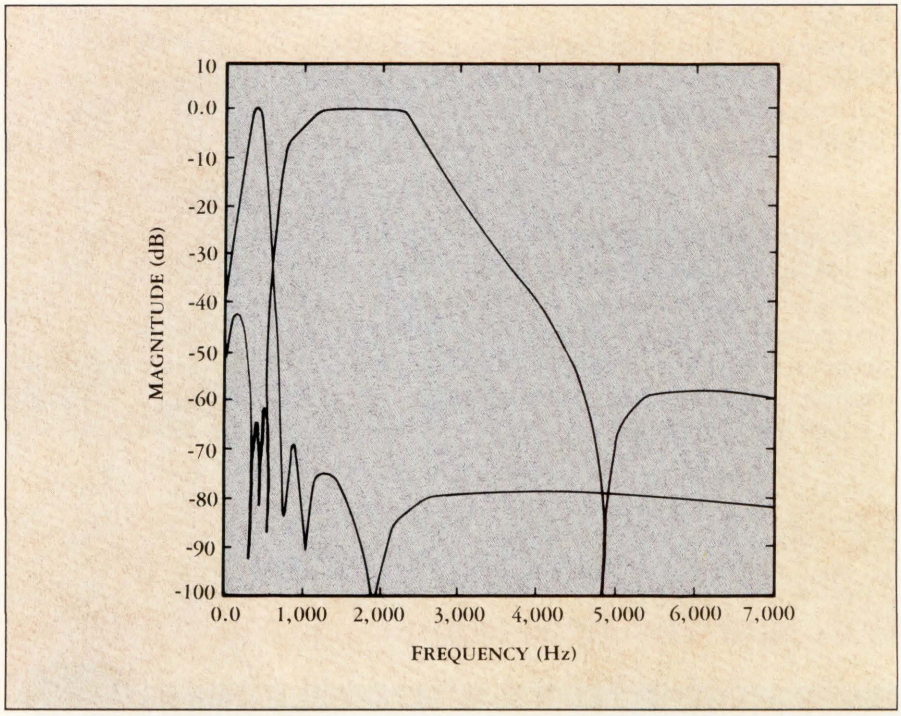

Figure 5. The V.23 filters are composed of two bandpass filters, one for the main channel and the other for back channel filtering. Their responses are as shown above.

DON'T GET CAUGHT IN THE ASIC SPEED TRAP

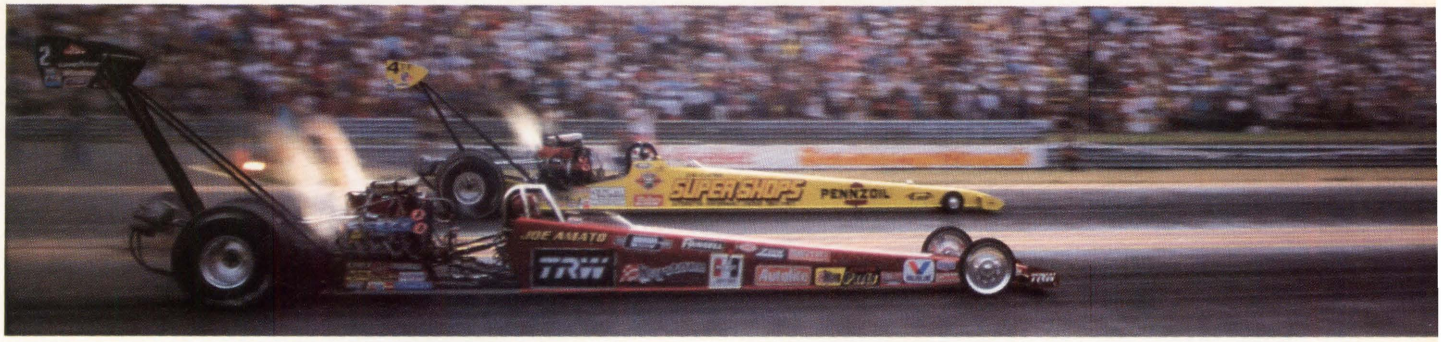

Winning in Design Verification Demands Highest Speed And Lowest Elapsed Time. Topaz-V Gives You Both.

There's a significant difference between test rate and testing time. That difference could add hours, days or even weeks to your VLSI design verification effort. Don't be trapped by only comparing the obvious.

The Topaz-V family of VLSI Design Verification Systems redefines the meaning of speed—concentrating on how fast it can work for you and your ASIC verification requirement. Not simply how fast it can run.

These fifth generation systems incorporate the revolutionary, new virtual vector memory (V^2M) architecture. It allows large vector sets to be downloaded, edited, and actively maintained on line. 'What if' analysis takes only seconds; simply change the set-up or vector and press 'start.' Time-consuming recompiles, vector re-downloading, and use of multiple vector pages are completely eliminated, slashing characterization time. Coupled with the system's 110 MHz test rate, 1.5 V/ns slew rate, 544 I/O pins, and 0.5 Gigabit per second communication interface, it will characterize your most complex device.

The Superior Solution to Prototype Testing, Failure Analysis, Quality Assurance, and Low Volume ASIC, VHSIC and VLSI Production

Specify Topaz-V. Thoughtfully engineered to satisfy your requirements for shorter test time, higher speed and greater accuracy.

Don't be caught in the ASIC speed trap. Look beyond the obvious. Redefine what speed and time really mean to you, and to your next new product's time to market.

Demand a demonstration of the Topaz-V systems; designed to meet your needs for high performance, accuracy and throughput into the decade of the Nineties. Then decide for yourself.

HILEVEL
TECHNOLOGY, INC.

At The Leading Edge of ASIC Verification

31 Technology Drive
Irvine, CA 92718
(714) 727-2100 • FAX 714-727-2101

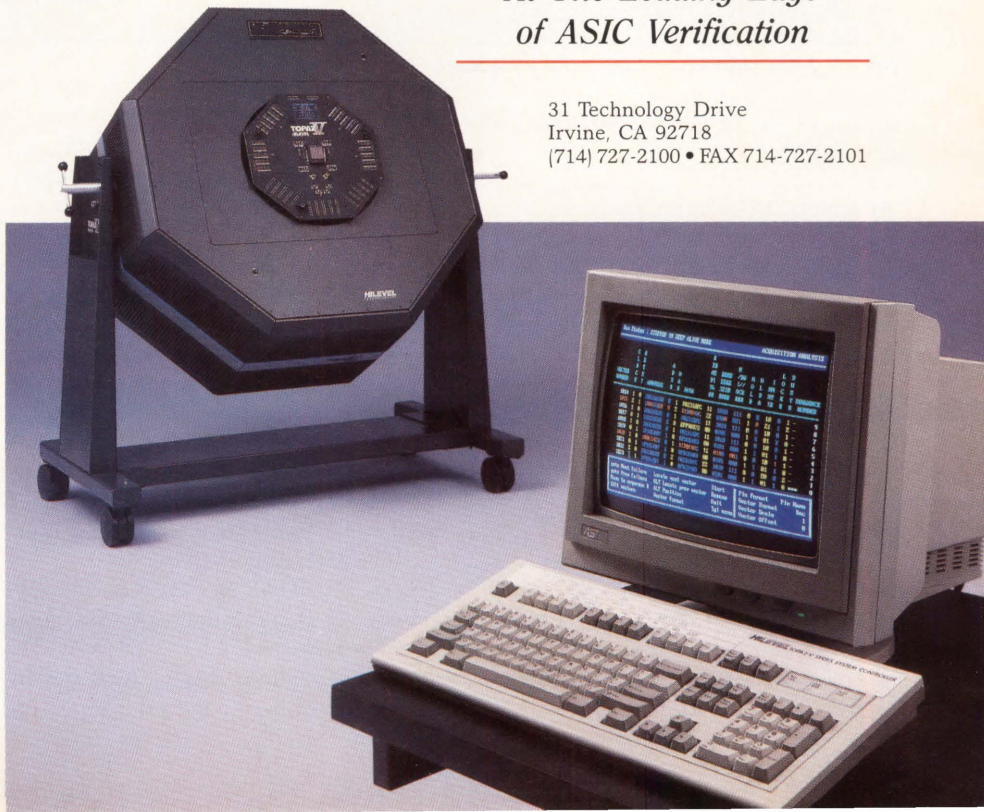

DIAL TOLL FREE:
1-800-HILEVEL
In California, 1-800-541-ASIC

Circle No. 17 for Literature
Circle No. 18 for Demonstration

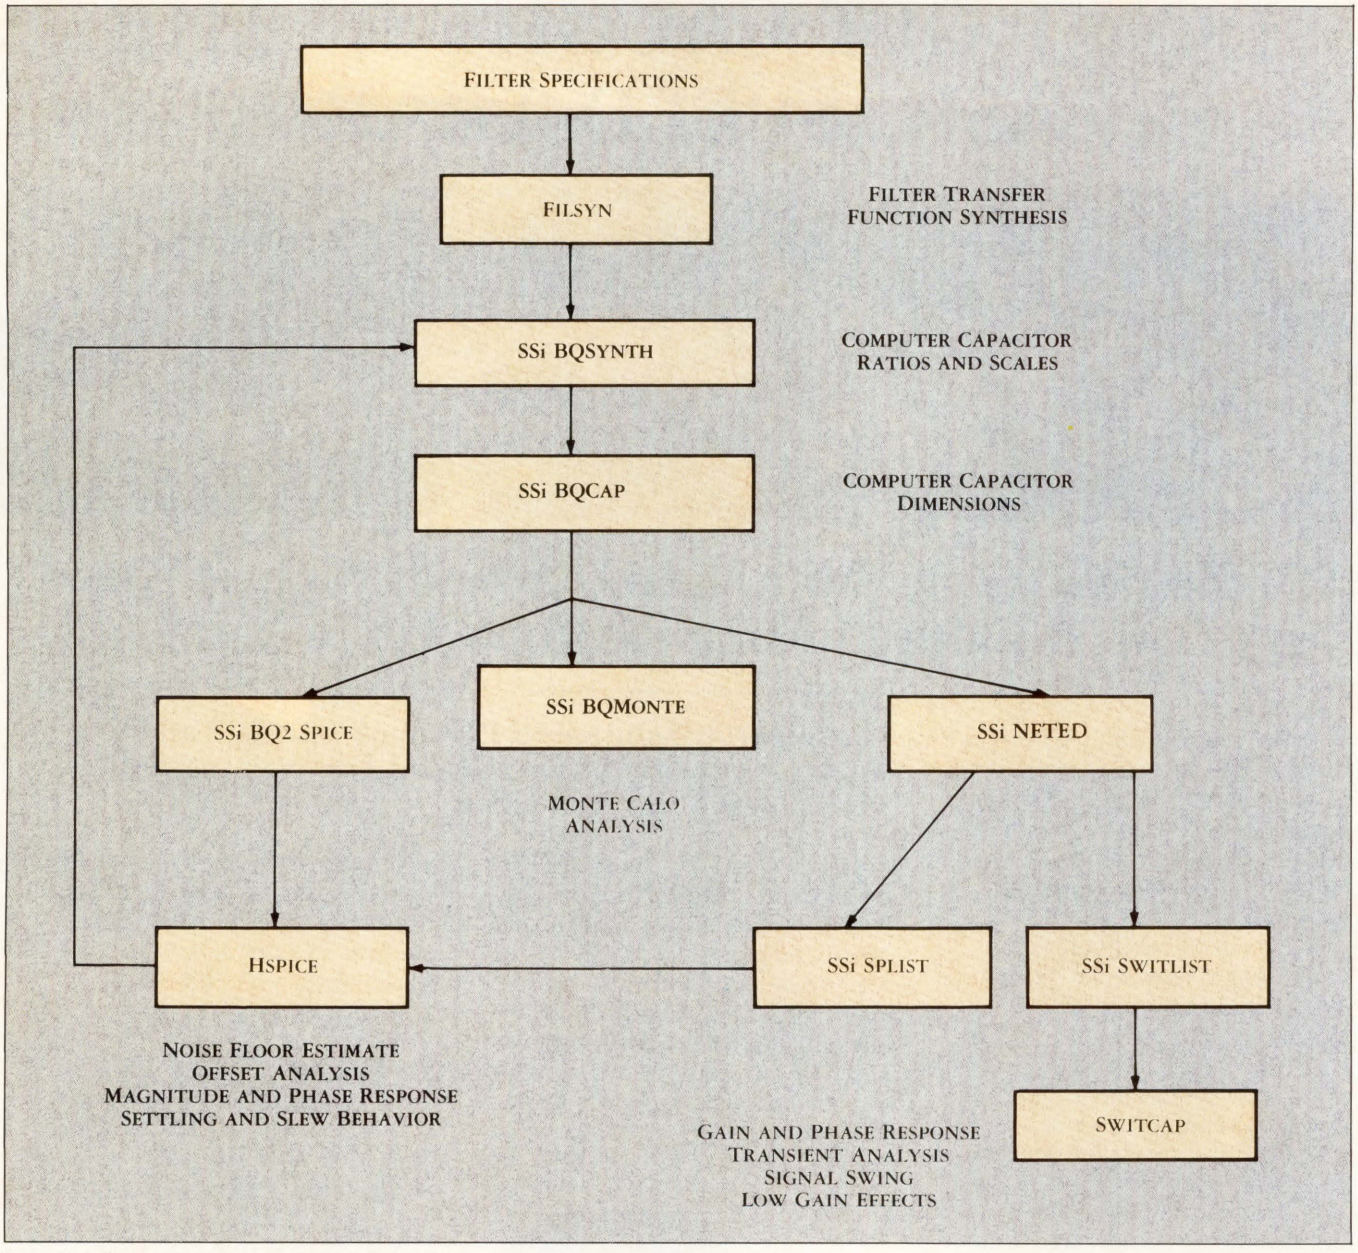

Figure 6. A flow chart for the filter design methodology.

design to ensure good testability of the device. ■

ACKNOWLEDGMENTS

Authors are indebted to Dr. David Rife for his valuable contributions to system definition issues; Matt Easley for help in defining filter requirements; and Fred Bunn and Mike Pugh for assistance in chip testing.

Our appreciation, is also extended to Tom Glad, Jeff Illgner, and Lou Testa of SSI for invaluable help with design and tool issues.

REFERENCES

Data Communication over the Telephone Network, Recommendation of the V Series, vol VIII, CCITT, Geneva, Switzerland, 1985.

FISCHER, J.H. "Noise Sources and Calculation Techniques for Switched Capacitor Filters," IEEE J. Solid-State Circuits, vol. SC-17, pp. 742-752, August 1982.

FLEISCHER, P.E., and K.R. LAKER, "A Family of Active Switched Capacitor Biquad Building Blocks," Bell Systems Technical Journal, vol. 58, no. 10, December 1979.

HURST, P.J., T.J. GLAD, J.J. ILLGNER, and G.F. LANDSBURG, "An Analog Front End for V.22bis Modems," IEEE J. Solid-State Circuits, vol. SC-23, pp. 978-986, August 1988.

TESTA, L., and M. MITTAL, "An Integrated CAE Design Automation System for Mixed Analog and Digital ASIC Design," in Procedures of ATE and Instrumentation Conference, West, pp. 287-299, January 1988.

ABOUT THE AUTHORS

RAOUF HALIM *joined the Advanced Technology Department of Hayes Microcomputer Products Inc. in 1985, where he is currently a principal engineer working on architecture, design and implementation of CMOS mixed analog/digital telecom ICs. He obtained a BSEE from the University of Alexandria, Egypt, in 1982, and an MSEE from Georgia Institute of Technology, in 1985.*

DANNY SHAMLOU *joined the Advanced Technology Department of Hayes Microcomputer Products Inc. in 1986, where he is now a design engineer working on analog IC design. He obtained a BSEE from the University of Tennessee, Knoxville, in 1985, and an MSEE from Georgia Institute of Technology in 1988.*

Microwave Triggering to 18 GHz.

we never stop asking "What if..."

Yesterday.

Today.

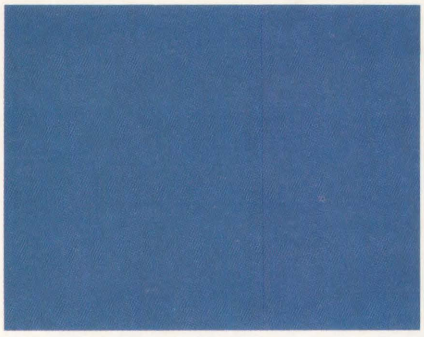

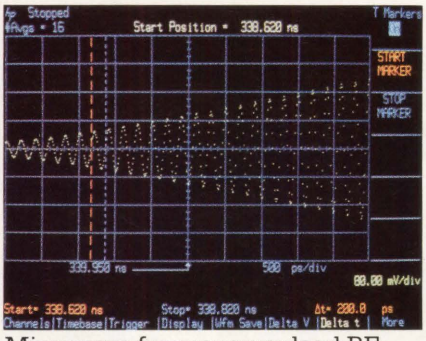

Microwave frequency pulsed RF

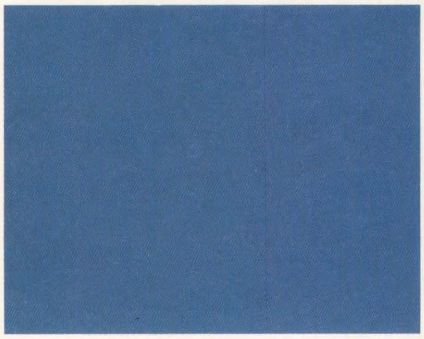

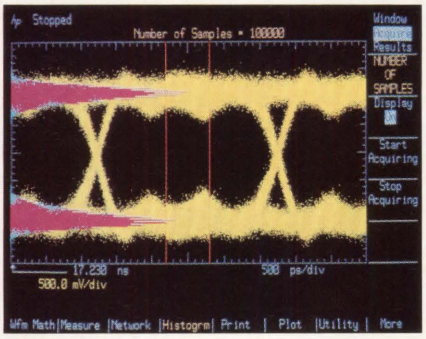

Multi-gigabit eye diagrams

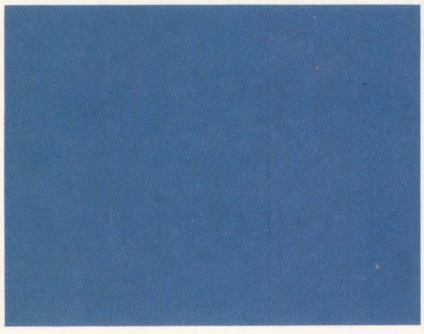

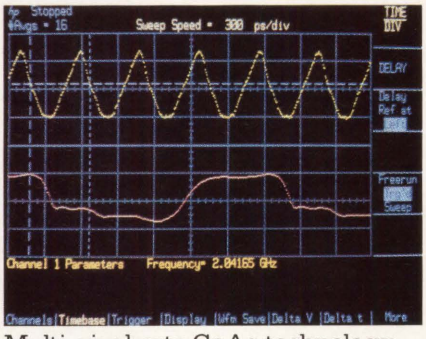

Multi-gigahertz GaAs technology

Now, trigger by level and slope at microwave frequencies with the HP 54120 Digitizing Oscilloscope.

You can now get images and make measurements never before possible with a benchtop oscilloscope. Now, the standard HP 54120 triggers from dc to 2.5 GHz—a five-time increase. When that's not enough, the new HP 54118A Trigger frees you from using countdown synchronizers to trigger at even higher microwave frequencies. For the first time, true event triggering is possible up to 18 GHz, extending measurement capability beyond CW to include pulsed RF, pulses and non-sinusoidal inputs.

You can trigger on the carrier in radar pulses or on CW signals containing a large noise component. Applications include testing and debugging of non-linear devices, lightwave communications systems, radar systems and components.

The HP 54118A's dual tunnel diode circuit also provides better dc stability and minimizes diode switching spikes. The trigger event will be unique and remain stable in the scope's display, even if the input signal has frequency drift or large deviation FM modulation.

Call HP today. 1-800-752-0900, Ext. 146C.

To get complete specifications and application information before you order, ask for our free brochure on the HP 54120 and HP 54118A. See for yourself that simple, stable triggering at microwave frequencies is a reality. Only from HP.

HEWLETT PACKARD

Recreating the Communications Environment

BILL JENNINGSCHECK AND MIKE FERGUSON,
LEVEL ONE COMMUNICATIONS INC., FOLSOM, CALIF.

Accurate simulation and testing require a proper system environment

Communications systems are getting more complex as telephone lines, once reserved primarily for voice transmission, become the main arteries for transmitting data and graphics information as well. These arteries depend on digital transceivers—complex devices that encode, shape, and transmit digital data over a transmission medium, and also receive and decode data coming in from the medium. Highly varied environments confront these communications systems, which must transmit digital information using analog transmission lines. Combining digital and analog circuitry within a transceiver presents design, simulation, and test difficulties. One hurdle is recreating the system's environments for accurate simulation and testing.

A modular design technology, in which system-specific transceivers are built from standard device cores, offers some advantages. Designers at Level One Communications Inc. have developed such a technology; it allows the system interface to be tailored for optimizing system performance. Also available are techniques that produce customer-specific transceivers. This way, designers can define the transmission technology, line code, bit rate, and other parameters for use in their transceivers.

The foundation of the Level One design methodology is a specialized set of modeling, simulation,

ILLUSTRATION BY JEAN TUTTLE

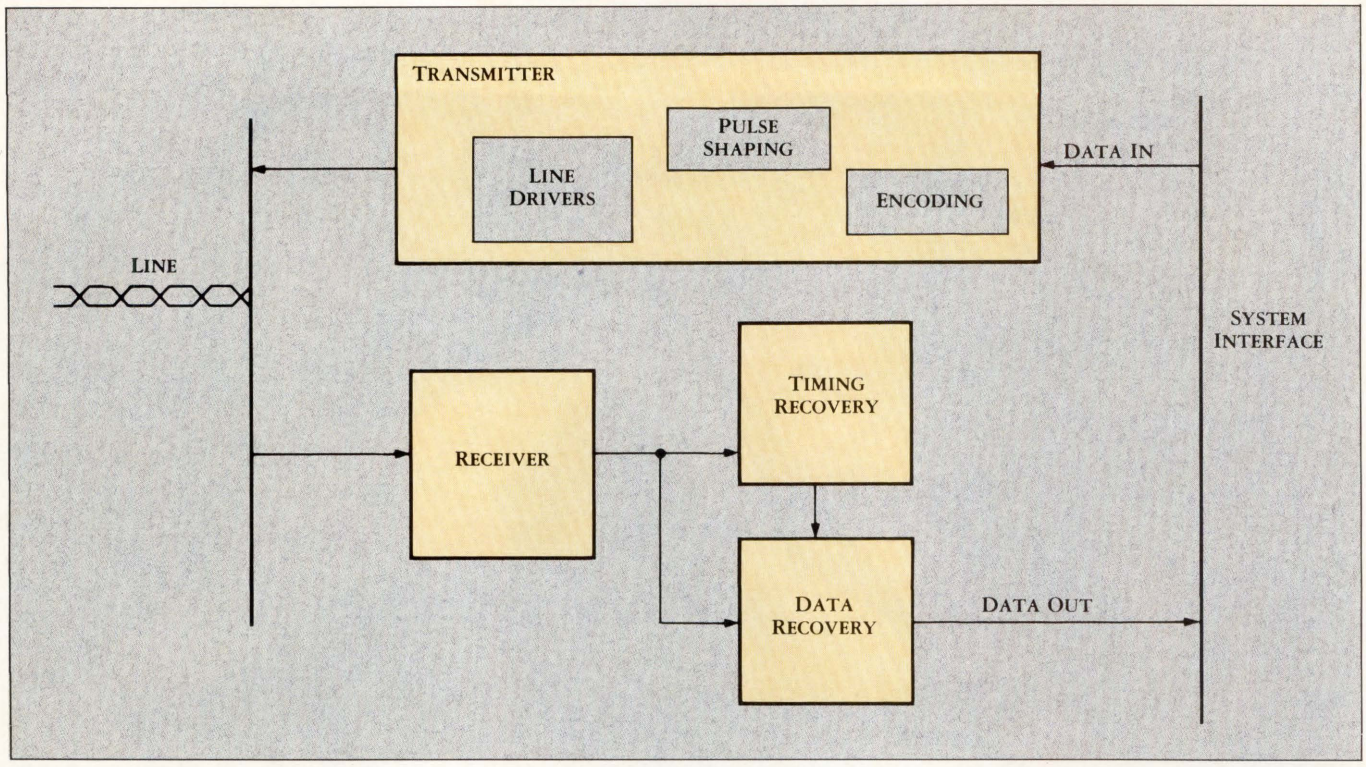

Figure 1. A data transceiver is a complex assembly of analog and digital functions: the transmitter function requires pulse shaping, conversion to line code, and matching the appropriate interface impedance. The receiver function performs equalization, noise filtering, reconversion of the line code to digital data, echo cancellation, and error detection and correction.

and development tools, supported by an extensive library of analog functional blocks and digital standard cells and functional blocks. These tools allow design and simulation to incorporate the effects of environment on design trade-offs. In addition, Level One has developed an approach to transceiver testing that allows testing of mixed analog and digital transmission circuits over a variety of network types and configurations. The test program is derived directly from the circuit and network simulation processes. That assures that the final product will be tested for the most critical parameters under conditions that closely resemble the actual system environment.

This design approach is meant to overcome a wide array of problems that challenge designers striving to implement transceiver functions on chips. The transceivers must operate over readily available, low-quality twisted-pair cable which, with more than 600 million installed lines, represents the major portion of the global telephone network. Digital transceivers designed to operate in this environment require a flexible design technology.

In the first place, supporting high-speed digital communications on a medium designed for analog signals introduces severe problems in signal attenuation, signal dispersion, and noise. These effects must be recreated during simulation and test to verify the circuit design.

Second, a transceiver is a complex combination of analog and digital functions (Figure 1). This, too, imposes special problems on design and test. The difficulty of integrating analog and digital functions on the same chip, and the additional problem of testing these functions simultaneously, often results in a more costly multichip solution unless specialized expertise can be brought to bear. In relatively complex systems, for example, the transmitter function requires pulse shaping, conversion to line code, and matching the appropriate interface impedance. The receiver function also performs several functions: some form of equalization to counteract transmission-line impairment of the signal, noise filtering, reconversion of the line code to digital data, echo cancellation, and error detection and correction. The timing recovery function must extract the clock from the received signal for synchronization with the transmitter at the far end (or with the system backplane, if that is the master timer source). Depending on the application, any one of these functions can be implemented in digital or analog circuitry.

Transmission networks, furthermore, implement a wide variety of operating environments, formats, and interface standards, any or all of which can differ from system to system. They include 80-to-144-kbit/s, time-compression multiplex (TCM) systems, 1.544-Mbit/s T1 pulse-code modulation (PCM) networks, and the 10-Mbit/s Ethernet. Therefore, no single transceiver solution can satisfy all applications.

The choice of the transceiver, transmission technique, line code, and other parameters is based on the system requirements and the network over which it must operate. These could include PBXs, which handle digitized voice and data from digital phones, terminals, or file servers; local area networks (LANs) for communicating from computers to the file server or between desktop communications equipment; high-speed T1 lines from the public network with signals that must be converted and distributed on twisted-pair wires; and central offices that transmit to customers' homes over telephone wire. This last will become even more important in the emerging ISDN environment.

■ SPECIALIZED TOOLS

To accommodate this wide spectrum of specifications on a timely basis, VLSI designers need a flexible design technology to make the appropriate design trade-offs for optimum performance within the system's anticipated environment. After verification and fabrication, the design must be tested with a program that, by recreating the environmental conditions, will assure the desired performance and reliability.

Level One's design flow begins with the

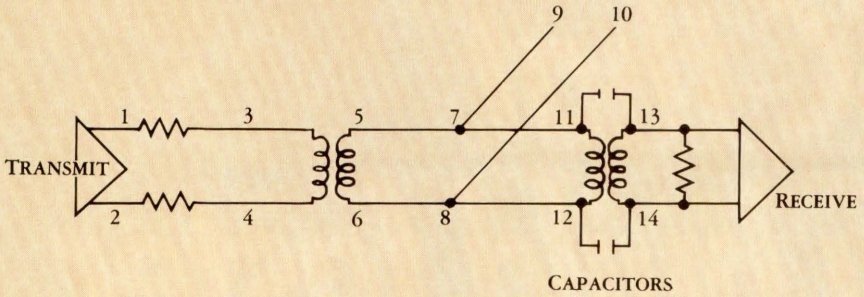

[A]

SAMPLE LOOP WITH BRIDGED-TAP

```
Rx1     1  3   200.0
Rx2     2  4   200.0
Xxmt    5  6   3   4    5.0M 2.0 0.9989  ! 5.0hM 2:1 xfrm
T1      5  6   7   8  1,200.0  24ga070    ! 24ga cable @ 70F
T2      7  8   9  10  2,050.0  26ga120    ! 26ga cable @ 120F
T3      7  8  11  12  9,730.0  24gaDIW    ! 24ga, 25pr DIW cable
Xrcv   11 12  13  14    5.0M 2.0 0.9989  ! 5.0hM 2:1 xfrm
Cww1   11 12  25p
Cww2   12 14  25p
rrcv   13 14  400.0
```

[B]

Figure 2. For transmission-line simulation, the designer constructs a model of his transmission environment (a) and extracts an LxNET topology file to provide as input to the line simulator (b).

system's specifications. The degree of difficulty in establishing the specifications depends on a variety of factors, such as whether the system's line interface is standard or nonstandard, as well as the topology of the system. For example, a transceiver designed for the ISDN U-interface is a rather complex chip that conforms to established international standards. On the other hand, a transceiver for a PBX manufacturer need not meet any standards except his own, since he supplies both the telephone and the system. In the latter case, the designer has control over the parameters that are of major concern for his system.

As important as transceiver functionality, however, is the description of the system environment in which the transceiver will operate. Parameters of the communications medium—such as the type of cable, the length or operating range of the loop, gauge changes in the cable, worst-case parameter tolerances on line interface elements, the presence of bridged taps (unterminated stubs of cable in parallel with the main loop), and noise characteristics—are critical in establishing the final product's performance criteria. Establishing and meeting specifications in a relatively short time depends on simulation tools that can incorporate the effects of the system environment.

■ SAVE THE ENVIRONMENT

The specifications are the starting point for system design. At the heart of this process is Level One's proprietary simulation tool, called LxWave, which combines behavioral simulation of the system with automated development of environmental parameters for the simulation. The designer uses LxWave to develop models of his system and its environment and to simulate those models to verify the system design. LxWave resides on a Sun 4/280 computer that functions as a server and is accessed from a number of Sun workstations connected to the server via Ethernet.

TOGETHER, LXNET AND LXSYS ACCURATELY PREDICT SYSTEM-LEVEL PERFORMANCE

LxWave consists of two modeling tools: a transmission line simulator (LxNet) and a system simulator (LxSys). Using LxNet, the designer constructs a model of his transmission line topology by drawing from a library of transmission line parameters. The parameter values in the transmission line parameter library can be assembled from one of three sources: published data from cable manufacturers; derivation from cable construction characteristics; or from laboratory measurements on actual cable. The laboratory technique is preferable, as it allows the most accurate modeling of the line in its expected operating environment.

In the lab, the impulse responses of cables are logged using a Hewlett-Packard Co. HP3577A network analyzer connected over an IEEE-488 bus to a Macintosh II computer. The Mac II runs instrument controller software that controls the response logging and collects the responses. These responses form models of the cables for the LxNet line-model library. Using this approach, Level One has developed line models for wire gauges ranging from 19 to 26, at environment temperatures as high as 120° F. These include models of single and multipair bundles of between one and 25 pairs. The line models are accurate from 5 to 30 MHz, depending on the application. The models are used by the simulator to calculate the response of the cable over a range of topologies.

Given a line model, the designer constructs a "topology" file that describes the transmission environment (loop). This file and the line model form the input to a simulation program in LxNet that calculates the response of the loop anywhere in its topology. The file can include transmission lines, transformers, inductors, capacitors, resistors, op-amp models, and other elements to accurately reflect the network operating conditions. The topology file is similar to a SPICE deck, consisting of numbered nodes and various components that describe the transmission environment. Figure 2a shows a sample transmission topology, with a transmitter driving a loop containing two transformers and a bridged tap; Figure 2b shows the corresponding line simulator input deck.

Once the network response has been obtained from LxNet, the designer can simulate the behavior of a transceiver design with LxSys, the other component of LxWave. LxSys is a proprietary system simulator that accepts a behavioral design of the transceiver constructed from software "modules" that are analogous to standard cells. These modules correspond to actual circuitry in the design library; users can construct modules for unrepresented functions by describing the function in Fortran (according to a user manual). The functional blocks within the behavioral design can be as large and abstract or as small and detailed as the designer desires.

In the simulation, the designer can include non-ideal behavior, such as jitter and noise, analog offsets and nonlineari-

The RISC that lets you build faster computers faster.

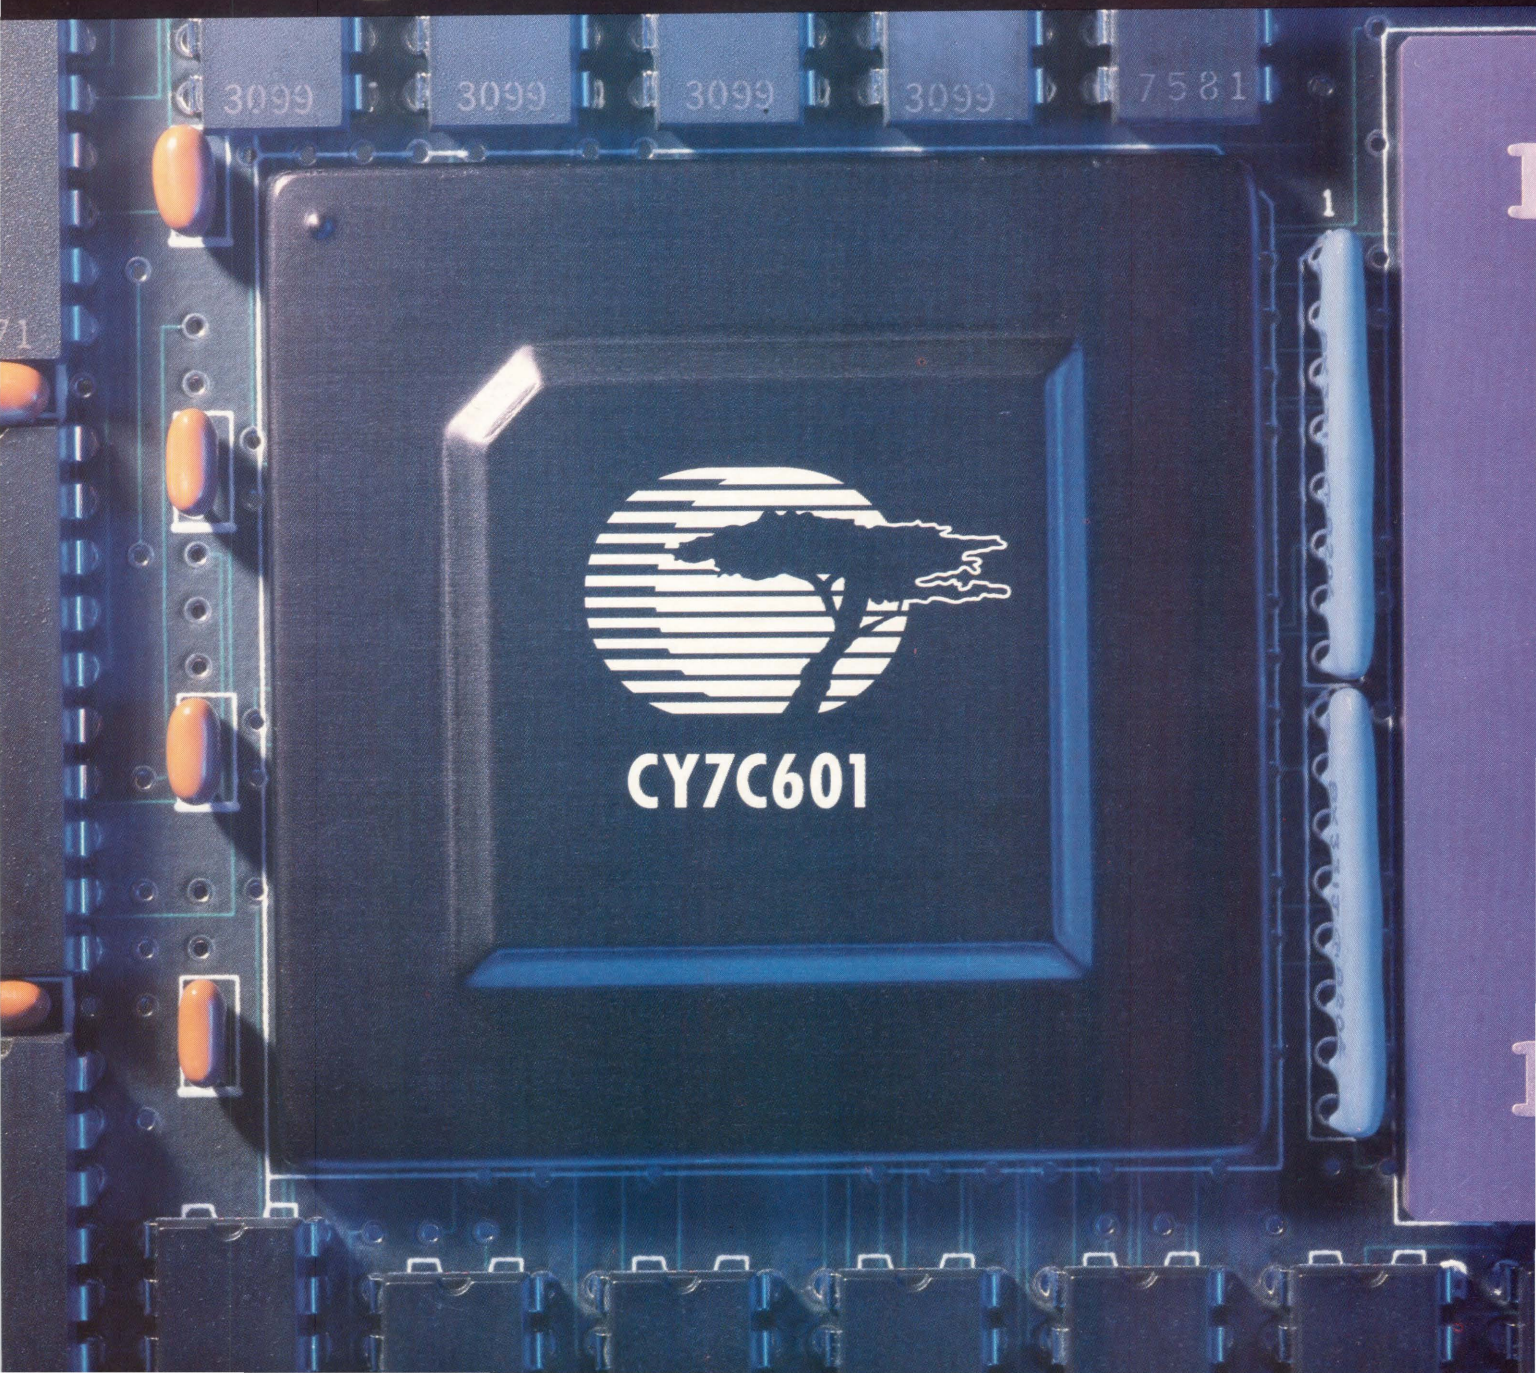

Introducing the RISC 7C600 Family: 20 VAX MIPS, SPARC architecture, and development systems today.

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

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

A complete chip set.

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

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

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

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

Develop your systems quickly.

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

Design around a proven architecture.

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

Make the fastest decision.

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

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

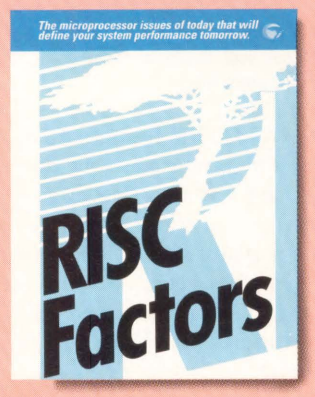

CYPRESS SEMICONDUCTOR

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

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

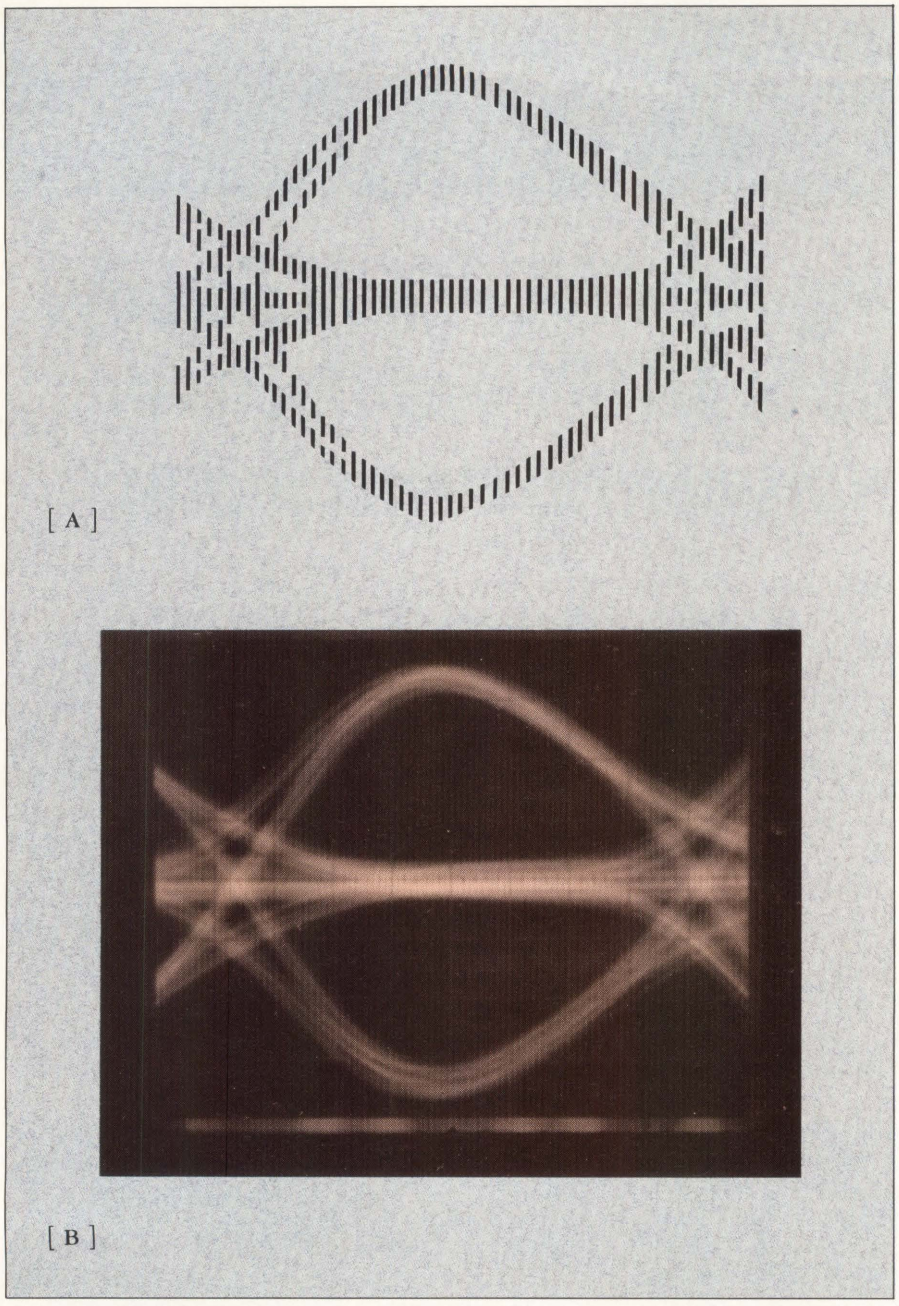

Figure 3. The eye diagram, generated by superimposing successive periods of the received pulse train, shows the amplitude and clarity of received signals. The simulator results (a) show close agreement with actual line measurements (b).

ties, bandwidth limitation, or quantization noise produced in the digital-to-analog conversion. Descriptions of the behaviors are written into the modules. The simulator also includes programs that can define various types of noise spectra, giving the user appropriate feedback so he can eliminate sensitivity to a specific type of noise.

Together, LxNet and LxSys accurately predict system-level performance of actual transmission networks by simulating the system design together with an accurate model of the network's signal response. Combined with a timing recovery module, the simulation can predict noise performance, recovered clock jitter, and other transceiver performance parameters. Modifications to achieve optimum performance can be checked rapidly for convergence of the timing recovery scheme or other adaptive circuitry such as equalization, echo cancellation, and filtering. A function optimizer in LxSys, for example, helps the designer find an optimum set of receiver characteristics to reduce noise for a given set of variables. The optimizer automates the process of exploring ranges for those variables through iterations of the simulator.

A valuable product of the LxWave tool is the "eye diagram" (Figure 3), which is generated by superimposing successive periods of the received pulse train. By combining positive pulses, negative pulses, and steady-state (no-pulse) bit periods, the eye diagram provides a convenient method by which to observe the amplitude and clarity of received signals. Its "openness" indicates the amplitude of the pulses, which, in turn, determines the relative noise immunity of the "slicing" that converts the analog received signal into digital data. The clarity of the eye diagram indicates the amount of intersymbol interference caused by pulses overlapping in succeeding bit periods. Intersymbol interference appears by a thickening up and curving of the center line or even of the top and bottom curved lines.

Figure 3a is an LxSys-simulated eye diagram of a transceiver with an adaptive equalizer, operating over 6,000 feet of a 24-gauge AT&T DIW 25-pair bundle. Figure 3b is an oscilloscope photograph of the eye diagram generated by the actual transceiver chip, operating under the same conditions as were simulated. This example shows close agreement (less than 10% difference) between simulation and actual performance for this example.

The behavioral simulator is efficient as well as accurate. Because transceiver block modules are behavioral descriptions rather than transistors or gates, changes can be made rather quickly to find an optimum system design. Without such modeling tools, verifying a transceiver design would require breadboarding with cable spools and hardwired impedance simulators.

Another unique advantage of this system is that both analog and digital signals are sampled and stored in a file for use later in design and test-generation cycles. These files can easily be converted to a piecewise linear source for use with HSPICE or to the tabular I/O format that drives the logic simulator. What's more, these files can also be translated to a format that can be downloaded to Level One's tester for recreating both analog and digital signals during IC test.

Once the assembled functional block models meet the desired specifications in the LxWave simulation, the next step is to implement those functions in actual circuits. The blocks used in LxWave have corresponding cells for standard-cell circuit design, although custom cells can be created for functions not included in the library. Using the cells, a circuit schematic is captured with the GED tools from Valid Logic Systems.

A standard verification and layout process follows. Analog circuits are simulated using HSPICE and digital portions with Valid's SIMULATE logic simulator. The results of these simulations, like those of LxWave, are available for conversion to

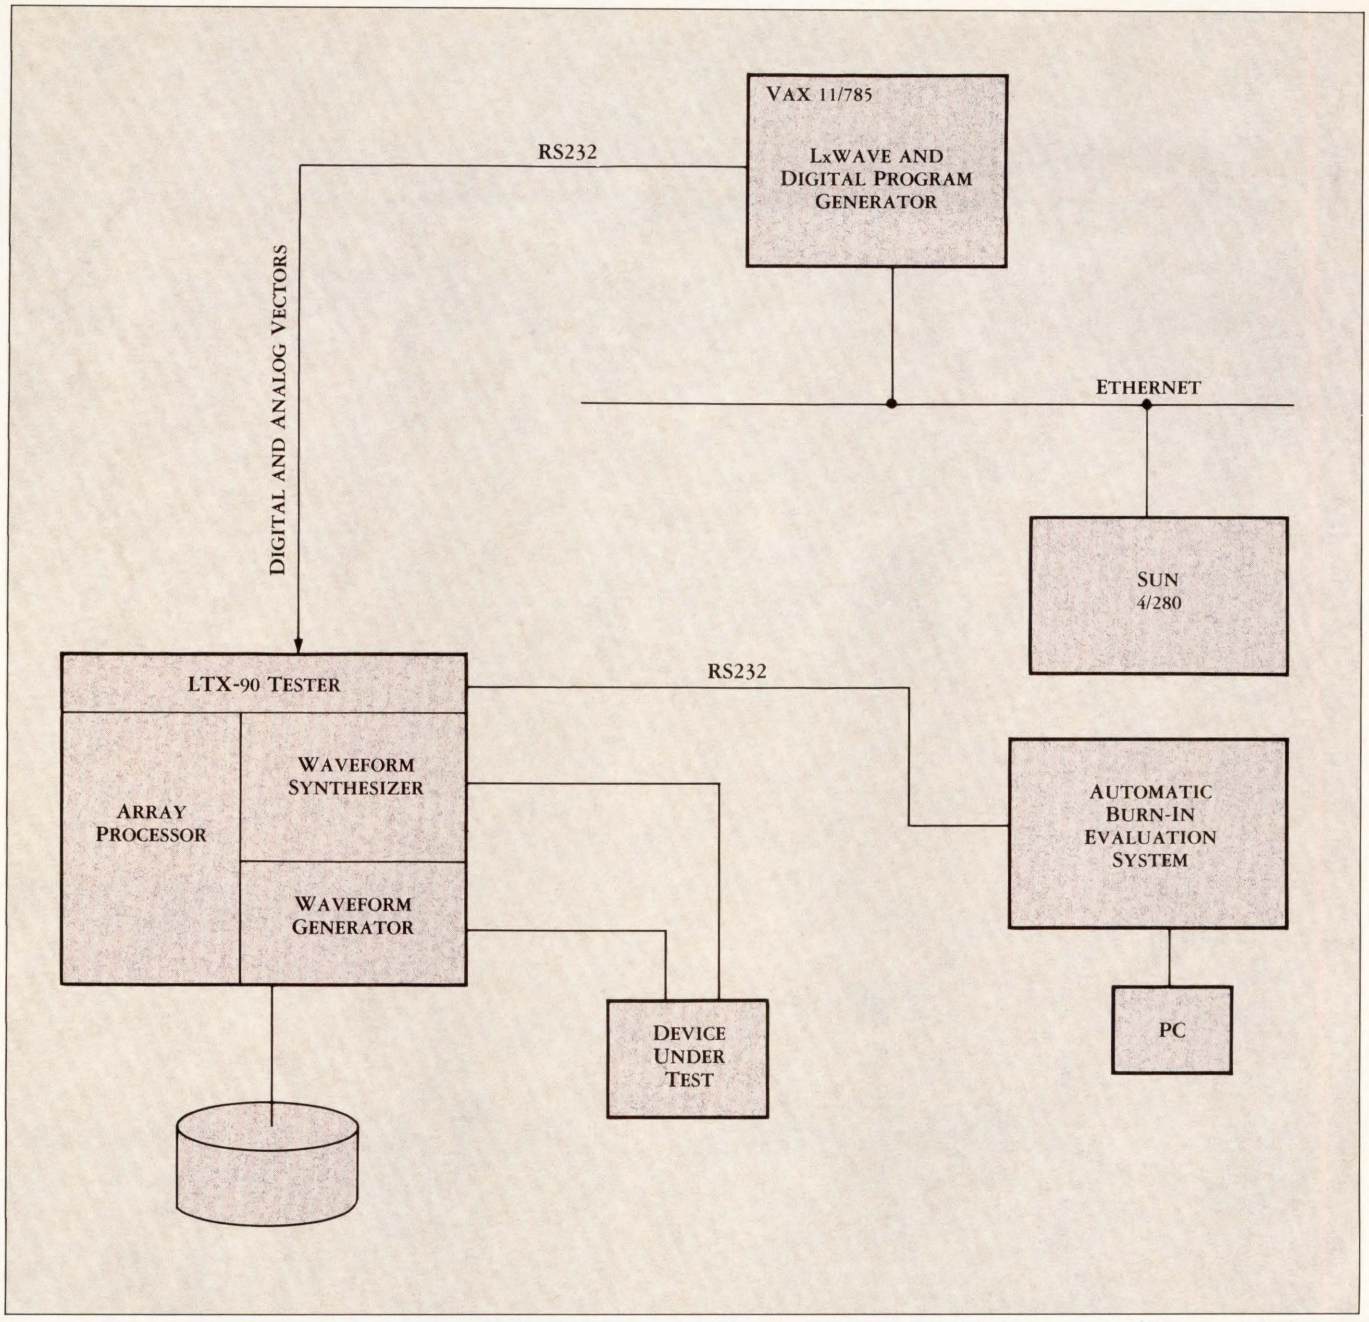

Figure 4. In this design environment, design and initial simulation are done on Sun workstations, while the analog and digital vectors produced by LxWave in the VAX are converted into test vectors and downloaded to the LTX tester.

test programs. The verified design is laid out, standard-cell style, with the help of proprietary techniques for power-line buffering and stabilization of the chip's substrate and well potentials. In addition, the use of "tweener" cells during placement creates routing channels within cell rows for a tight layout.

■ TRANSCEIVER TESTING

When the first silicon is returned from the IC foundry, it is subjected to the all-important step of design verification. By automating a large part of the production test program generation process, the time was reduced to less than four weeks from a typical six-month cycle. The method is based on the ability to derive analog waveforms and digital test vectors from the output of the LxWave simulation and the logic simulator. These waveforms and vectors can be transmitted to the test system for incorporation into a test (Figure 4).

The files that form the basis of the test vectors are generated throughout the entire design process. During the later stages of architectural definition, the performance of the transceiver is evaluated under various loop configurations and noise environments. The results of these simulations provide a means of generating analog test waveforms and developing the specification for the transmitted pulse. The files, containing the sampled analog waveforms, are post-processed on the VAX and converted to a format compatible with Level One's tester, the TS90 from LTX Corp. The conversion software matches the magnitude and time period of the sampled waveforms with the driving capability and cycle time of the LTX tester's WS801 waveform synthesizer. Once downloaded from the VAX, the waveforms are reconstructed by the WS801 and applied to the device under test.

The characteristics of the transmitted pulse can also be determined from system simulation data. In some cases this specification may be a standard. In others, however, it is determined by evaluating the transmitter's output for various cable and loop configurations.

From this data a specification for a properly working transmitter can be de-

FROM COPPER TO FIBER, COVER THE FREQUENCY RANGE FROM 700 MHz TO 8 GHz...

AND BEYOND.

8 GHz

5 GHz

3 GHz

PATTERN GENERATION FOR SYSTEM EVALUATION AND ANALYSIS HAS NEVER BEEN EASIER OR MORE COMPLETE.

Anritsu has a complete range of pattern generators and receivers that cover 1 MHz to 8 GHz and beyond, manual or full GPIB control.

Virtually any type of pattern can be generated including pseudorandom, variable mark ratios and programmable patterns up to 16 kps. Pattern and clock output levels, as well as offset voltage, can be user controlled. In addition, delay can be set on all the units to permit flexible interfaces.

Anritsu helps you every step of the way whatever your test requirement. From high speed digital ICs to high speed optical devices...from electrical to optical transmission systems.

So take the first step and contact Anritsu today at 1-800-255-7234 or in NJ 201-337-1111. Anritsu America, Inc. 15 Thornton Rd., Oakland, NJ 07436, Fax 201-337-1111.

Anritsu

CIRCLE NUMBER 28

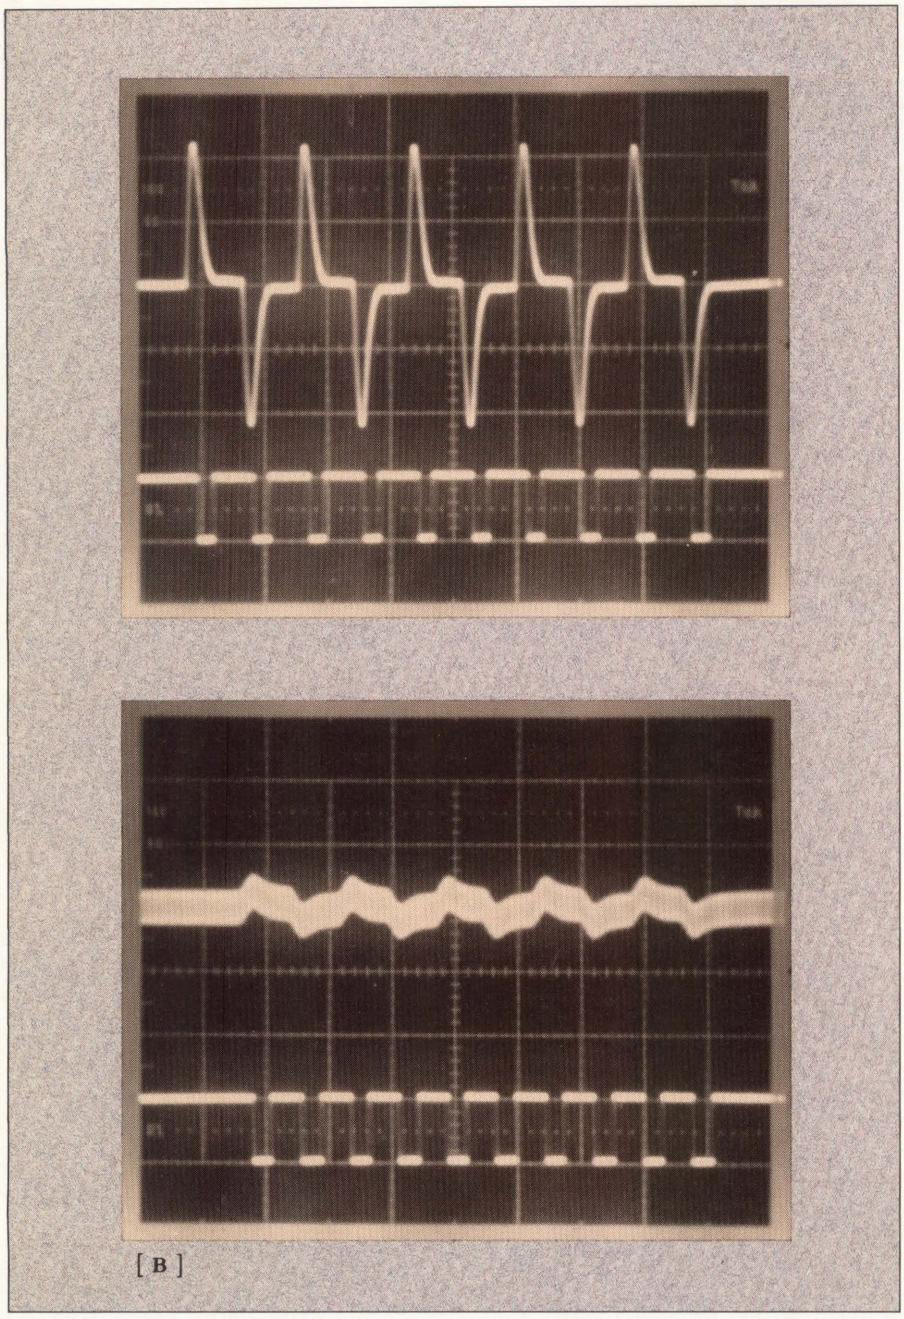

Figure 5. During testing, a transceiver first receives a "zero-length line" pattern for a functional check of the major blocks on the chip (a). A subsequent "worst-case" test presents a waveform as it would appear at the end of the "worst-case" loop, with degraded waveforms and, in this case, a 2.048-MHz noise component (b). As seen in the digital waveforms, this tranceiver is still able to discern bit patterns correctly.

termined. The specification is used by the test engineer to develop a "template" that defines the maximum and minimum slope and amplitude boundaries of the transmitted pulse. The tester can then sample the transmitted pulse using the waveform digitizer (WD801) and compare the sampled data points to the boundaries described by this software template.

When the design simulation and fault-grading are complete, the digital vectors defining the behavior of the final chip at its pins are also post-processed on the VAX and downloaded to the LTX. In this step, digital events are mapped to tester cycles. The test engineer then determines how to align the digital and analog waveforms to exercise the transceiver accurately. He also adds parametric tests to the pin responses to verify individual pin characteristics. The final test program, therefore, has three components: analog test vectors from system simulation; digital vectors from system, logic, and circuit simulation; and pin parametric tests.

Figure 5 shows a set of two patterns as typically run on a transceiver under test. The first (Figure 5a) is a "zero-length line" pattern, in which the tester is working as if the driver and receiver are directly tied to one another. This pattern is basically a functional check to see if all the major blocks on the chip are working properly.

Another pattern would be a marginal, or "worst-case," test in which the receiver would be tested using a simulated waveform as it would appear at the end of the "worst-case" loop (of a specified length, gauge, and inductance). In Figure 5b, a worst-case test includes not only degraded input signals but also a 2-MHz noise component. Without the availability of the waveform synthesizer and the processed simulation patterns, the test engineer would have had to build a jury-rigged setup using a reel of cable or an equivalent artificial line.

As mentioned previously, the design engineer has determined (from the simulation with LxNet) what the pulse amplitude, width and slope of the driver's output waveform must be, so that ones and zeros can be accurately discriminated at the end of the specified transmission line after undergoing the anticipated degradation. The amplitude data becomes a template that is downloaded into the tester's memory. The driver's output is then compared with the template in the tester's array processor to determine whether or not the output falls within the template's limits.

Bit error rate (BER) is an important specification, but it takes a great deal of testing time to accumulate the data. This is because typical BERs range from 1 bit in 10^7 to 1 bit in 10^9. To enhance quality assurance, Level One engineers designed "smart" burn-in equipment with the capability of testing BER during burn-in. The equipment contains an array of "mini-testers," which energize the transceiver with appropriate analog and digital vectors derived from the test program and log the BER over the test period. This final step ensures that the IC will work as well in its target communications system as it did in the IC tester. ■

About the Authors

BILL JENNINGSCHECK *is engineering manager at Level One Communications Inc. in Folsom, Calif. JenningsCheck joined Level One from Xicor, where he was a section head for standard product development. Prior to that, JenningsCheck has worked as a program manager for Intel's Telecommunications Operation Department. JenningsCheck received his BSEE from Lehigh University.*

MIKE FERGUSON *is director of operations at Level One. Before joining Level One, Ferguson was Southwest-area applications manager at LTX Corporation, with previous experience at Fairchild and Teradyne. Ferguson received his BSEE and BSCS degrees from the University of Florida.*

Can you solve this problem with a one-clock telecom tester?

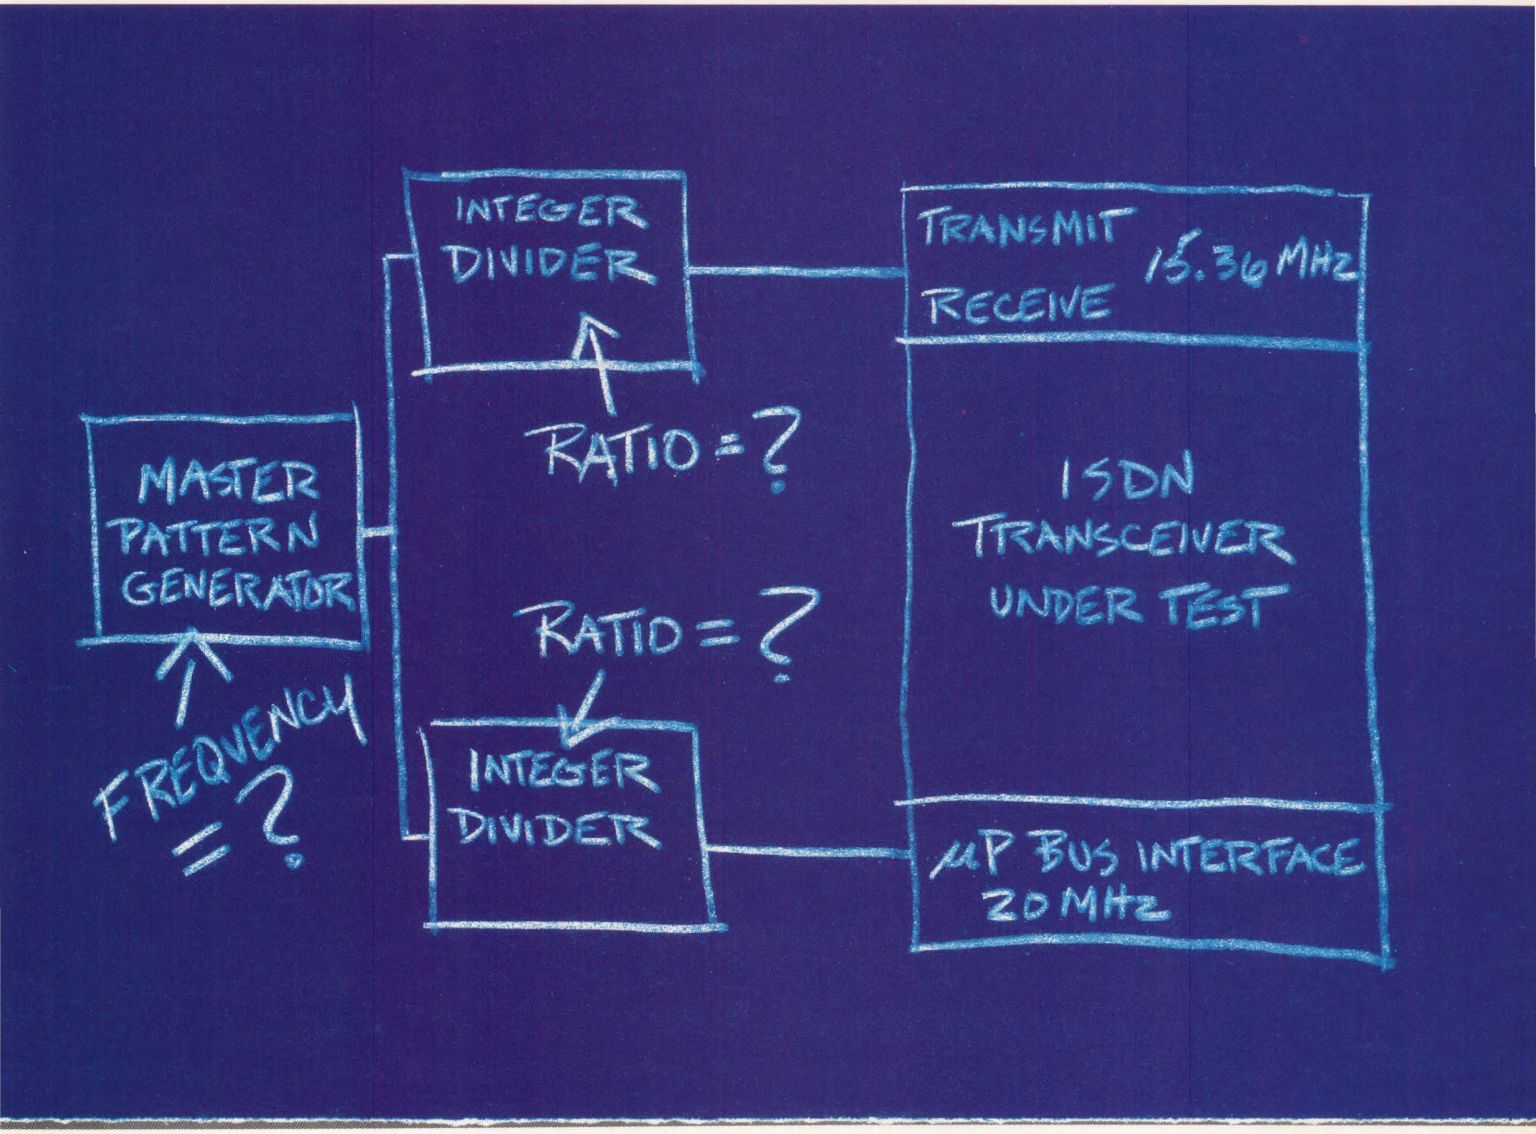

This ISDN transceiver requires synchronized transmit/receive data at 15.36 MHz *and* a bus interface pattern at 20 MHz. Given the tester's single master clock, what frequency do you program? What integer ratios do the dividers need?

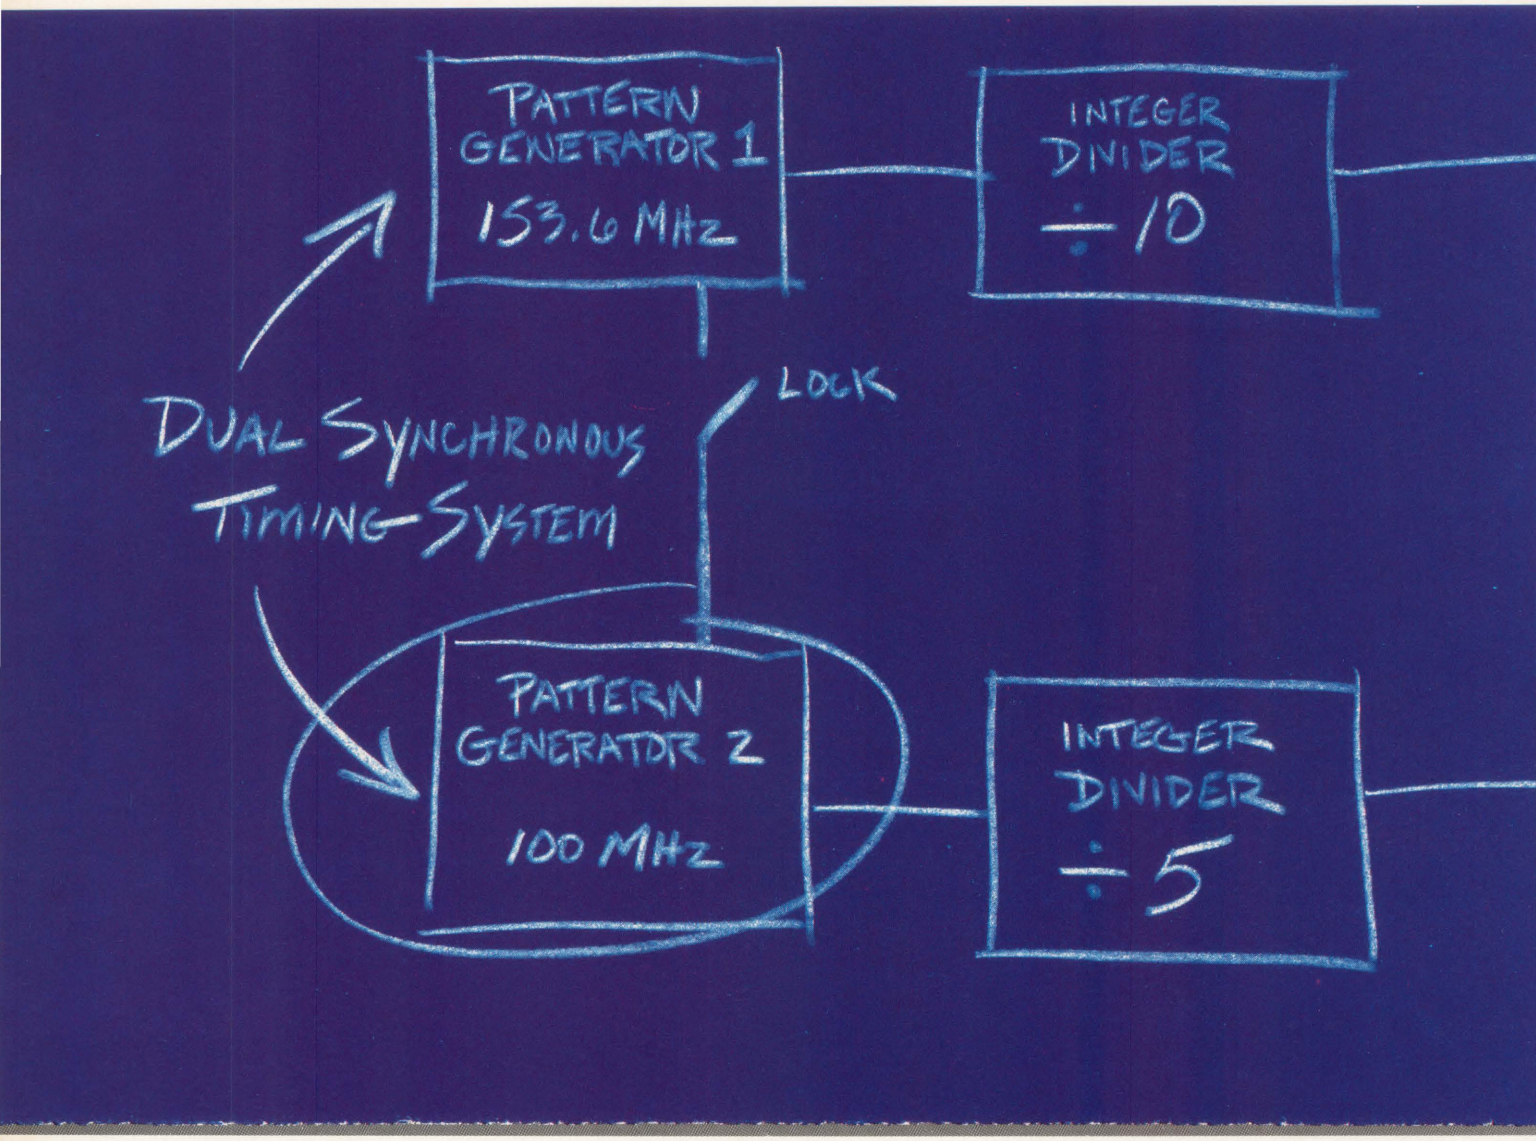

If you settle for anything less than Synchromaster, the new resource-per-pin mixed-signal test system from LTX, you're going to wind up one clock short. In order to produce test frequencies of 15.36 MHz and 20 MHz with a single-clock test system, you would have to choose integer ratios of 125 and 96 and set the pattern generator at 1.92 GHz — well beyond the typical 100-200 MHz range.

Only Synchromaster can produce the exact combination of frequencies this IC requires. Its Dual Synchronous timing system uses *two* high-resolution generators, so patterns and waveforms aren't constrained to just ratios of a single frequency, as they are in single-generator testers.

We've been through exercises like this many times, for many years. In fact, no one else comes close to LTX in mixed-signal testing experience. That's why the Dual Synchronous timing system is a design fundamental. And that's why Synchromaster

Need more time? Get the tester with two clocks.

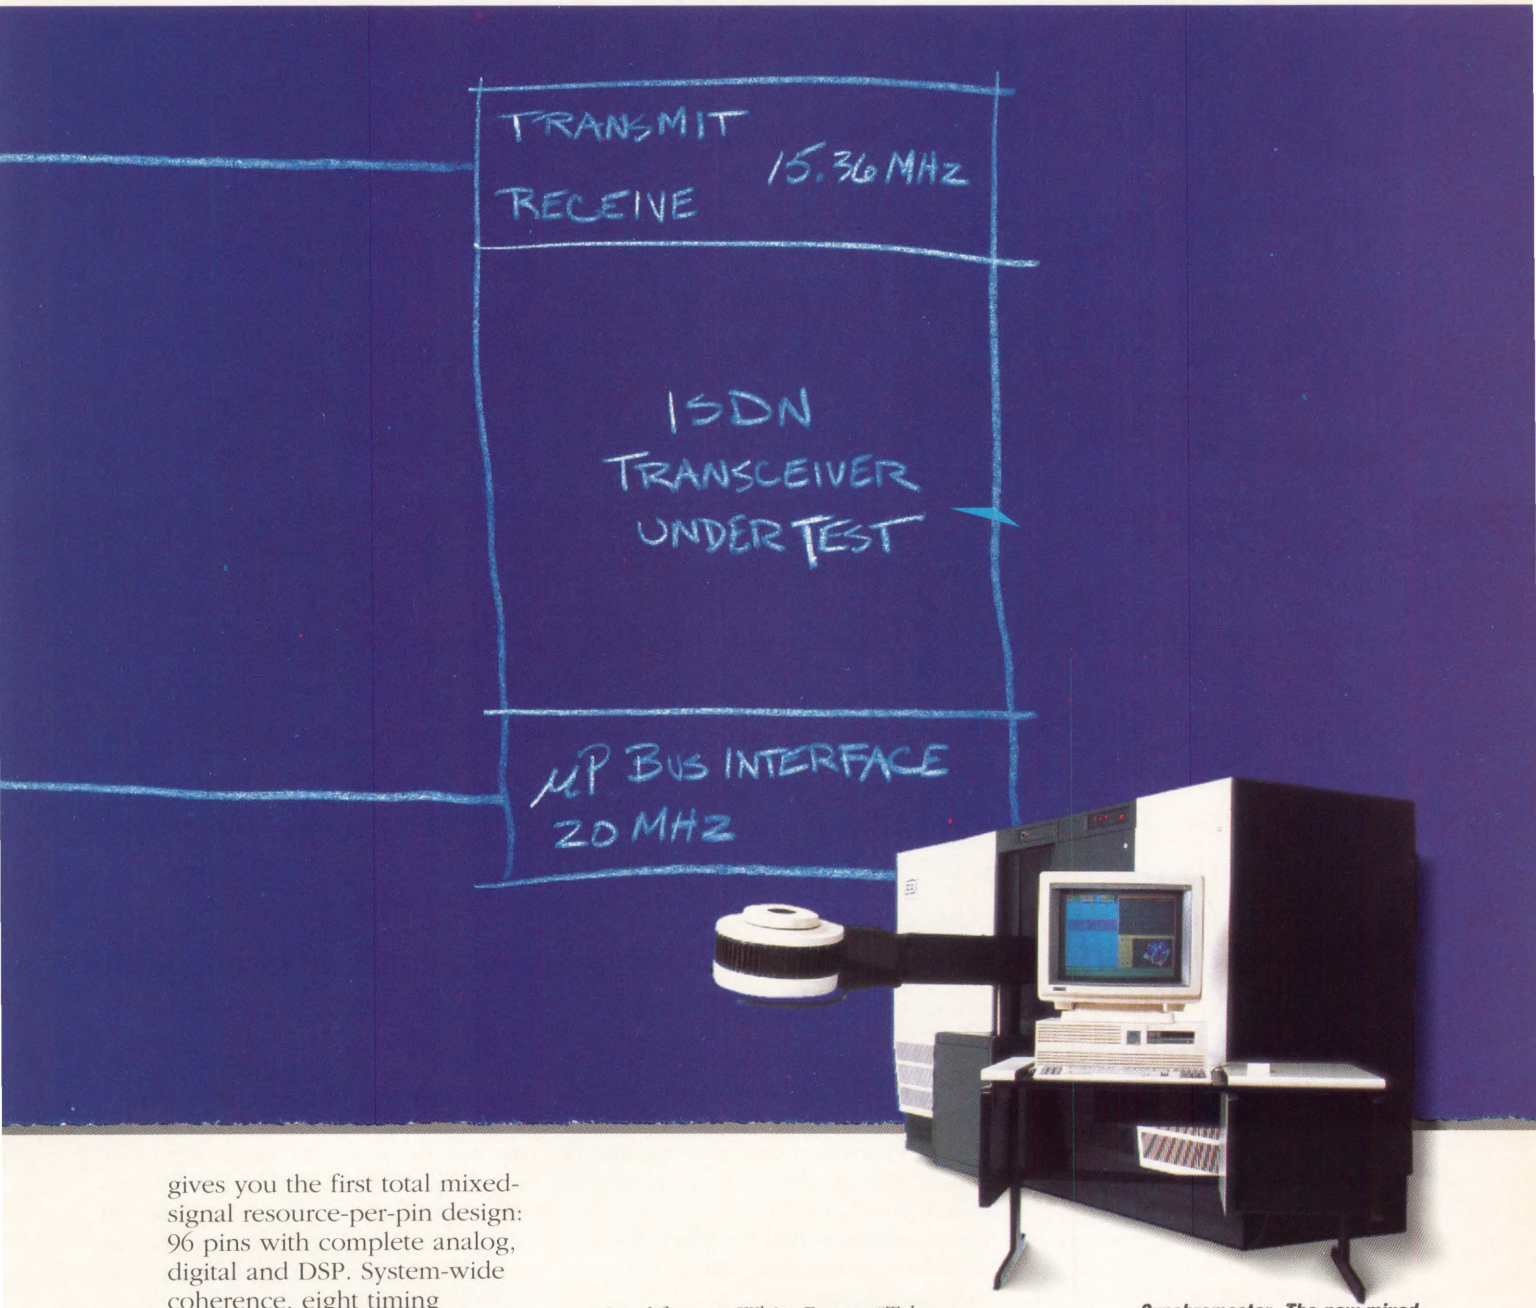

gives you the first total mixed-signal resource-per-pin design: 96 pins with complete analog, digital and DSP. System-wide coherence, eight timing markers per pin, on-the-fly time set switching, and pattern rates up to 80 MHz. *And* techniques for solving the problems of jitter, pulse mask testing and other telecom measurements ... without compromise.

Send for our White Paper: "Telecom Testing Without Hangups," for the details on Synchromaster's Dual Synchronous timing system. Better yet, use the telephone. Call Carol Everett at (617) 461 1000.

Synchromaster. The new mixed-signal tester from LTX.

It's why time has run out on one-clock testers.

LTX ***Responds.***

CIRCLE NUMBER 20

DIGITAL SIGNAL PROCESSOR IC's

Top down and bottom up design techniques are combined

LUIS BONET AND TIM A. WILLIAMS,
MOTOROLA INC., AUSTIN, TEXAS

With the arrival of VLSI design tools, engineers can now design efficient and cost-effective digital signal processing ICs for specific applications. The efficiency of these DSP machines is improved by reducing the number of instructions and/or modifying the architecture of more general-purpose designs. As in RISC architectures, application-specific DSP designs provide single-cycle instructions tailored to the algorithms to be executed. Such simple single-cycle instructions reduce the machine cycle time, increasing the design's efficiency.

The concepts and procedures involved in designing an application-specific DSP machine are discussed in this article. Augmenting the concepts will be a case study of a design recently completed by the authors. This machine was designed to execute the algorithms necessary to implement the ANSI and CCITT standard algorithms for adaptive differential pulse code modulation (ADPCM), G.721 and T1.303.

The advantages of application-specific DSP designs include their ability to perform algorithms that will not fit into general-purpose machines because of memory or speed limitations. Another advantage is that a higher degree of concurrency can be integrated into an application-specific DSP than with the general-purpose DSP, because of the limited number and types of communication paths available within the general-purpose DSP devices.

The major advantage of application-specific DSPs is their highly flexible architecture. With this flexibility, the designer has the freedom to create an architecture that is optimized in terms of performance and cost. Special data paths or processor elements can be added to increase performance. Concurrent processors can also be added, such as address ALUs, to perform critical operations. Pre-ALU functions, i.e. functions that modify the data before it is presented to the ALU, also enhance the machine's power.

Using a general-purpose DSP in situations where the algorithms are changing rapidly or several sets of algorithms are performed on the same machine is generally considered the only approach. But there are several ways to design application-specific DSPs in such environments. Partitioning the tasks to be performed will, in many cases, allow the application-specific DSP to be used in systems where several sets of algorithms are required. Classification of different algorithms is discussed in this article to aid the reader in selecting similar sets of algorithms to be performed on a single application-specific DSP.

The case study involves the design of a machine to execute the algorithms necessary to perform the ANSI and CCITT standard algorithms for 24- and 32-kbit per second ADPCM. A good review of these algorithms is presented by Daumer et al. (1984). The process of encoding and decoding involves adaptive filters that predict the next PCM data word based on the past PCM data words. The adaptive filter contains six zeros and two poles. Data within the algorithm is represented in many formats. A-Law or μ-Law PCM format is used for the input and output data. Floating-point data formats are used for intermediate values in the algorithm, i.e. the state of the filter. Both sign-magnitude and two's complement formats are used for fixed point numbers within the algorithm. These range in precision from

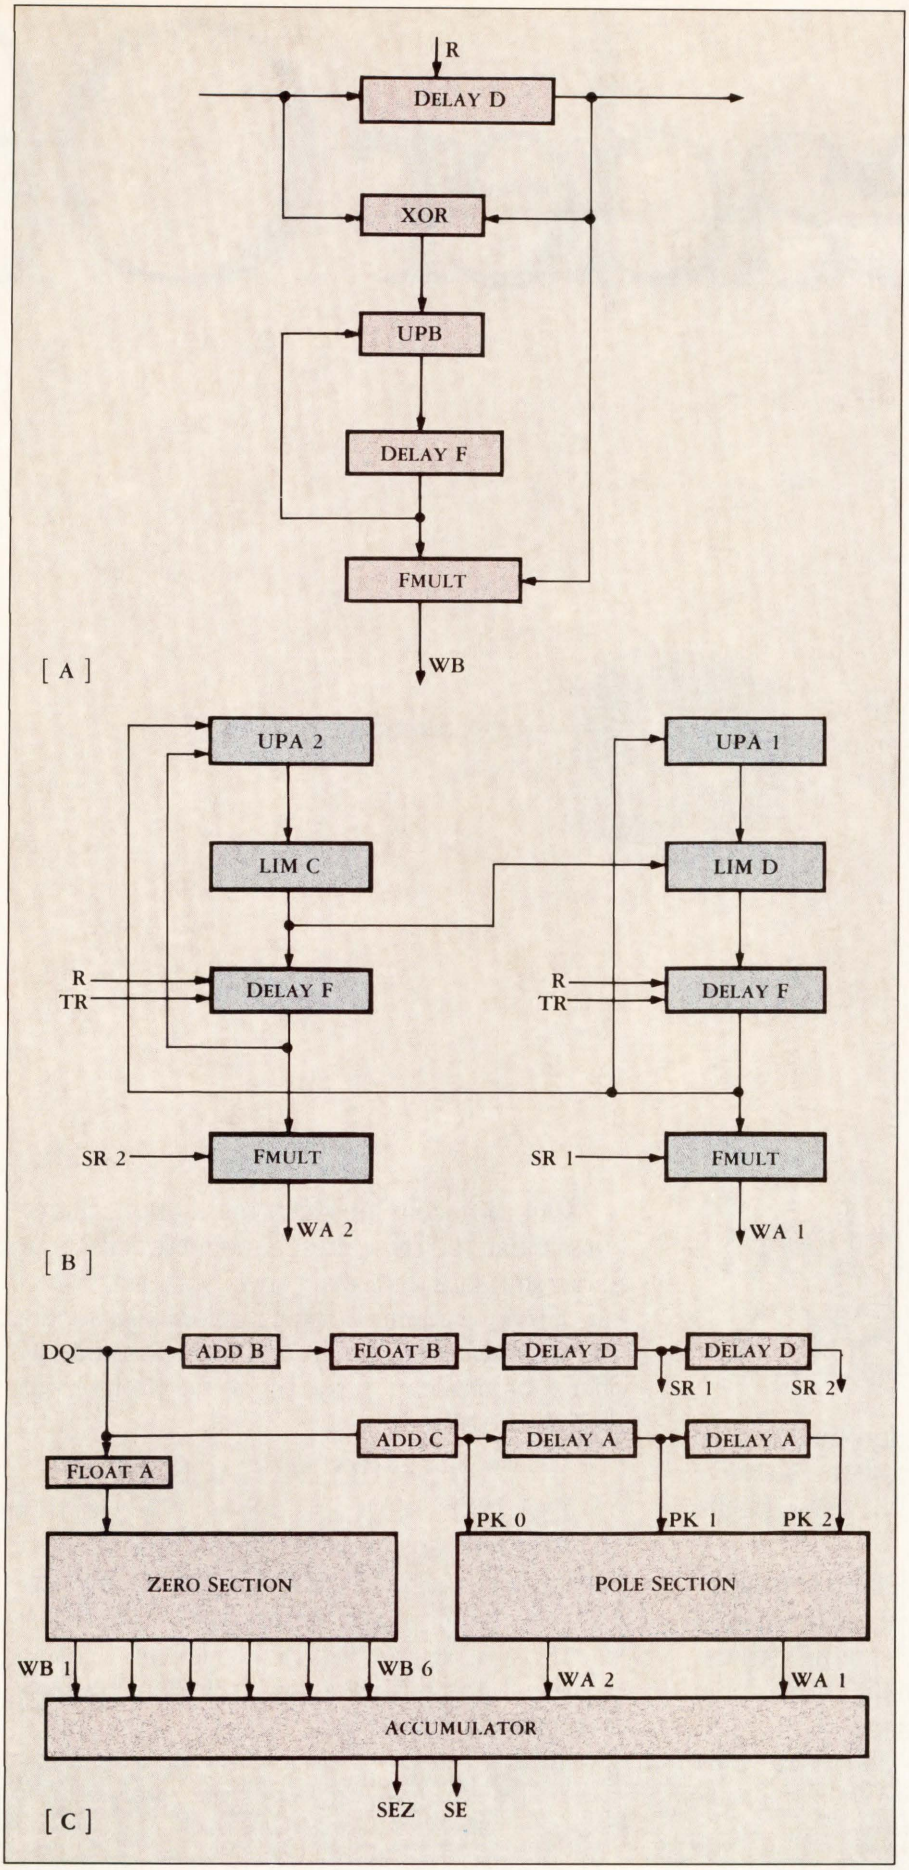

Figure 1. Filter section of the ADPCM chip. The six zeros of the transfer function are implemented with the identical sections shown in (a); along with the pole section (b) these data paths feed the accumulator as in (c).

6 bits to 19 bits. The efficient manipulation of these different number systems is a challenge to the designer.

The following sections describe the sequence of operations used in designing an application-specific DSP. This design process takes a top-down approach initially, then jumps to a bottom-up design at the point where the initial coding of the algorithms is done.

■ Overall Goals and Constraints

The case study's design goals were set from a knowledge of both the end users' desires and the limitations of the VLSI processes. The main goal was to have a single chip implementation of both standards.

The second design goal was to allow full duplex operation. This forced the requirement on the processor that it be able to operate with asynchronous channels. Because the I/O was designed with interrupts and an interrupt scheduler, we gained the benefit of being able to process either two encodings, two decodings, or an encode/-decode pair in a single sample period.

A third goal was that the chip have a small number of pins. This is possible with the serial protocol selected for the I/O channels and the fact that the part is not user-programmable.

Selecting a technology in which to achieve the stated goals is the next step. General-purpose machines have advantages where quick system prototypes are required. However, for production volume equipment, the application-specific DSP usually has the cost advantage.

Partitioning of the tasks to be performed will, in many cases, allow the application-specific DSP to be used even in systems where several sets of algorithms are required. For example, in algorithms where transform operations need to be performed in conjunction with other operations, a DSP specialized for transforms can be utilized effectively in conjunction with another processor.

Application-specific DSPs are useful in systolic architectures, since the designer has a great deal of freedom to design the data and communications flow. For example, multiple data words of various precisions could enter a certain node of the systolic array. If the node were a general-purpose DSP, the designer would have to buffer the data or arrange the transfer of the data at staggered times since general-purpose DSPs have limited bus availability. The result may be an undesirable latency. If the node were an application-specific DSP, the data paths and precisions could be incorporated in the node, thus

Leader Of

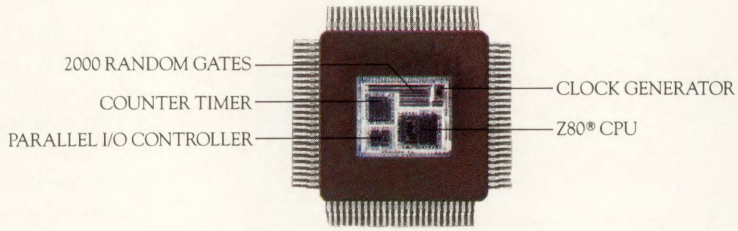

PACKAGE: 100 PFP, 32 MIL CENTERS*

®Z80 is a trademark of Zilog Inc.

Now 1.5 μ Hard Megacell ASICs.

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

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

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

Z80 FAMILY MEGACELLS

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

AREA SALES OFFICES; CENTRAL AREA, Toshiba America, Inc., (312) 945 – 1500; **EASTERN AREA,** Toshiba America, Inc., (617) 272 – 4352; **NORTHWESTERN AREA,** Toshiba America, Inc., (408) 737 – 9844; **SOUTHWESTERN REGION,** Toshiba America, Inc., (714) 259 – 0368; **SOUTH CENTRAL REGION,** Toshiba America, Inc., (214) 480 – 0470; **SOUTHEASTERN REGION,** Toshiba America, Inc., (404) 368 – 0203; **MAJOR ACCOUNT OFFICE, FISHKILL, NEW YORK,** Toshiba America, Inc., (914) 896 – 6500; **MAJOR ACCOUNT OFFICE, BOCA RATON, FLORIDA,** Toshiba America, Inc., (305) 394 – 3004. REPRESENTATIVE OFFICES: **ALABAMA,** Montgomery Marketing, Inc., (205) 830 – 0498; **ARIZONA,** Summit Sales, (602) 998 – 4850; **ARKANSAS,** MIL-REP Associates, (512) 346 – 6331; **CALIFORNIA** (Northern) Elrepco, Inc., (415) 962 – 0660; **CALIFORNIA** (L.A. & Orange County) Bager Electronics, Inc., (818) 712 – 0011, (714) 957 – 3367, (San Diego County) Eagle Technical Sales, (619) 743 – 6550; **COLORADO,** Straube Associates Mountain States, Inc., (303) 426 – 0890; **CONNECTICUT,** Datcom, Inc., (203) 288 – 7005; **FLORIDA,** Sales Engineering Concepts, (813) 823 – 6221, (305) 426 – 4601, (305) 682 – 4800; **GEORGIA,** Montgomery Marketing, Inc., (404) 447 – 6124; **IDAHO,** Components West, (509) 922 – 2412; **ILLINOIS,** Carlson Electronic Sales, (312) 956 – 8240, R.W. Kunz, (314) 966 – 4977; **INDIANA,** Leslie M. DeVoe Company, (317) 842 – 3245; **IOWA, Carlson Electronics, (319) 377 – 6341; KANSAS,** D.L.E. Electronics, (316) 744 – 1229; **KENTUCKY,** Leslie M. DeVoe Company, (317) 842 – 3245; **LOUISIANA,** MIL-REP Associates, (713) 444 – 2557; **MAINE,** Datcom, Inc., (617) 891 – 4600; **MASSACHUSETTS,** Datcom, Inc., (617) 891 – 4600; **MICHIGAN,** Action Components Sales, (313) 349 – 3940; **MINNESOTA,** Electric Component Sales,

The Packed.

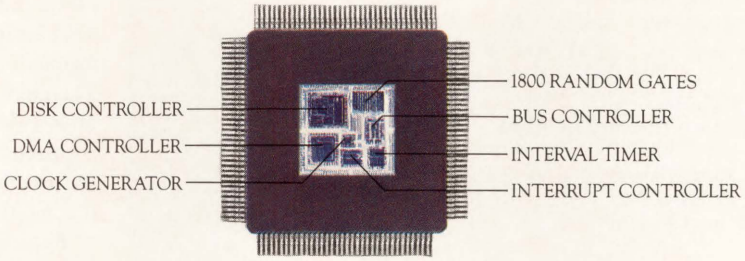

PACKAGE: 144 PFP, 25 MIL CENTERS*

*Die shown larger than actual size.

82Cxx PERIPHERAL MEGACELLS

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

SPECIAL PURPOSE MEGACELLS

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

GET A PACKAGE OF INFORMATION

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

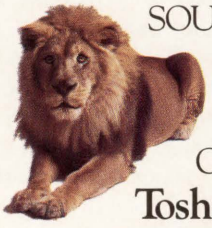

Toshiba. Leader Of The Packed.

TOSHIBA AMERICA, INC.

(612) 933 – 2594; **MISSISSIPPI,** Montgomery Marketing, Inc., (205) 830 – 0498; **MISSOURI,** D.L.E. Electronics, (316) 744 – 1229, R.W. Kunz, (314) 966 – 4977; **MONTANA,** Components West, (206) 885 – 5880; **NEBRASKA,** D.L.E. Electronics, (316) 744 – 1229; **NEVDA,** Elrepco, Inc., (415) 962 – 0660; **NEW ENGLAND,** Datcom, Inc., (617) 891 – 4660; **NEW HAMPSHIRE,** Datcom, Inc., (617) 891 – 4600; **NEW JERSEY,** Nexus Technology, (201) 947 – 0151; **NEW MEXICO,** Summit Sales, (602) 998 – 4850; **NEW YORK,** Nexus Technology, (201) 947 – 0151; Pi-tronics, (315) 455 – 7346; NORTH CAROLINA/SOUTH CAROLINA, Montgomery Marketing, Inc., (919) 467 – 6319; NORTH DAKOTA/SOUTH DAKOTA, Electric Component Sales, (612) 933 – 2594; OHIO, Steffen & Associates, (216) 461 – 8333, **(419) 884 – 2313, (513) 293 – 3145; OKLAHOMA,** MIL-REP Associates, (214) 644 – 6731; **OREGON,** Components West, (503) 684 – 1671; **PENNSYLVNIA,** Nexus Technology, (215) 675 – 9600, Steffen & Associates, (412) 276 – 7366; **RHODE ISLAND,** Datcom, Inc., (617) 891 – 4600; **TENNESSEE,** Montgomery Marketing, Inc., (205) 830 – 0498; **TEXAS,** MIL-REP Associates, (512) 346 – 6331, (713) 444 – 2557, (214) 644 – 6731; **UTAH,** Straube Associates Mountain States, Inc., (801) 263 – 2640; **VERMONT,** Datcom, Inc., (617) 891 – 4600; **WEST VIRGINIA,** Steffen & Associates, (419) 884 – 2313; **WASHINGTON,** Components West, (206) 885 – 5880, (509) 922 – 2412; **WISCONSIN,** Carlson Electronics, (414) 476 – 2790, Electric Component Sales, (612) 933 – 2594; **WYOMING,** Straube Associates Mountain States, Inc., (303) 426 – 0890; **CANADA, BRITISH COLUMBIA,** Components West, (206) 885 – 5880; **ONTARIO,** Electro Source, Inc., (416) 675 – 4490, (613) 592 – 3214; **QUEBEC,** Electro Source, Inc., (514) 630 – 7486.

MAS-88-001

CIRCLE NUMBER 21

```
Inputs:      {An or Bn, SRn or DQn
Outputs:     {WAn or WBn

AnS =          An >> 15
AnMAG =        {An >> 2                    if AnS = 0
               {(16384-(An>>2)) & 8191     if AnS = 1

AnEXP =        {13,      if 4096 <= AnMAG
               {12,      if 2048 <= AnMAG <= 4095
               {.
               {.
               {2,       if 2 <= AnMAG <= 3
               {1,       if AnMAG = 1
               {0,       if AnMAG = 0

AnMANT =       {1 << 5 if AnMAG = 0
               {(AnMAG << 6) >> AnEXP, otherwise

SRnS =         SRn >> 10
SRnEXP =       (SRn >> 6) & 15
SRnMANT =      SRn & 63

WAnS =         SRnS ** AnS
WAnEXP =       SRnEXP + AnEXP
WAnMANT =      ((SRnMANT * AnMANT) + 48) >> 4

WAnMAG =       {(WAnMANT << 7) >> (26 - WAnEXP)      if WAnEXP <= 26
               {(WAnMANT << 7) << (WAnEXP - 26) & 32767      if WAnEXP > 26

WAn =          {WAnMAG                       if WAnS = 0
               {(65536 - WAnMAG) & 65535     if WAnS = 1

*              - Multiplication
**             - Exclusive OR
>>, <<         - Shift right, shift left
&              - Logical AND
```

Figure 2. The FMULT routine listed here reflects the ADPCM standard for floating-point multiplication and format

allowing a higher speed implementation.

The use of standard cells or gate arrays for application-specific DSPs imposes architectural constraints on the designer, particularly regarding on-chip memory. The typical microcoded ASDSP has a wide microcode control word. However, the typical gate array or standard cell design has only limited capacity for ROM. One method to avoid these ROM limitations is to implement several basic micro routines in on-chip ROM. Sequences of these micro routines can be executed from outside the chip. Another technique is to remove the ROM from the chip entirely and allow direct control of the architecture from the outside. This approach runs into a pin limitation if taken to extremes and also is speed limited by the off-chip access time of the memories. An instruction cache also could be built, allowing high speed execution of a few routines on chip.

Implementing data RAM on a gate array or standard cell design has similar problems. The storage and organization of the data RAM often has a dramatic effect on the throughput of the machine. Gate array and standard cell products are not efficient when multiple arrays of memory or unusual widths of memory are required. Accessing long arrays of off-chip data RAM tends to slow the chip down. A small array of data RAM on chip may be utilized in a cache scheme in conjunction with a cache controller to increase the speed of execution.

Gate array and standard cell design technology are useful for application-specific DSPs if certain problem areas are avoided. This may involve nonconventional architectures. Raw speed of the process is not the only criteria when selecting a technology to implement ASDSPs with gate arrays or standard cells. Often a unique architecture can reduce the throughput requirements of the design dramatically.

■ CLASSIFICATION OF THE ALGORITHMS

At this point, the systems designer's task is to transform a high-level language implementation of the algorithm developed using general-purpose computers into a form that can be implemented in the DSP. This transformation typically involves the design of both hardware and software. The transformation cannot be performed in a single step since there is no architecture or instruction set defined. Instead, iterations are made to the architecture and the microcoding until the design goals are met. The first task is to decompose the algorithms into general classes of operations, for example, transforms, FIR filters, IIR filters, lattice filters, and ladder filters. Once the operations are classified, the designer can get a clear idea of what the chip's structure must be.

A block diagram of the case study adaptive filter is shown in Figure 1. The adaptive filters predict the next PCM data word based on the past PCM data words. The adaptive filter contains six zeros and two poles, which can be identified by the variables WB1-WB6 and WA1, WA2 as they enter the accumulator from the FMULT blocks. The delay blocks that hold the variables DQ1-DQ6 retain the state of the zeros of the filter.

If we analyze the algorithms, we find that it takes 50 RAM locations to store the state of the algorithm from one sample period to the next. At this point we can check what bandwidth is required for data movement. You can assume, for the worst case, that each arithmetic operation has two inputs and one output and that each RAM location will have one operation performed on it during every sample period. Therefore, in a single-bus architecture, 50×3 data moves per sample period are required. This estimate can tell you if the architecture must be a multibus design or if a single bus will suffice.

Typically the multiply-accumulate time is the worst-case cycle in a DSP, so it is an important elementary operation. In the case study, these are the FMULT blocks, which receive both floating-point and two's-complement fixed-point inputs. The outputs are two's complement numbers. Format conversion therefore must take place within the FMULT block.

Format conversions are therefore important elementary operations. Other types of elementary operations in this machine are bit shifting, table look-ups, and arithmetic functions. As each type of elementary operation is identified, a general timing estimate can be revised to give a clearer picture of the required architecture and processor speed.

I/O requirements should also be considered. If the DSP has to accommodate asynchronous inputs, as in this case, time must be allocated to cover the scheduling of the calculations. Data to the encoder/decoder arrives at 125-μs intervals. The ADPCM algorithm is sample- rather than block-oriented so each data point can be processed independently. If the data needs to be blocked, as in a Fast Fourier Transform, temporary storage must be used.

By now, the memory structure should be approximately known, the important elementary operations identified, the I/O form and timing understood, and a rough timing estimate completed.

■ HIGH LEVEL REPRESENTATION

Expressing the algorithm using a high-

level representation enables the designer to identify the requirements for the structures responsible for the arithmetic operations, data movement, memory space, I/O operations, and control sequencing. In the case study, this step was already included in the standards. As an example of the level of detail present in the ADPCM standards, the FMULT routine is included as Figure 2.

Careful study of the algorithms shows the following characteristics:

- The majority of the arithmetic operations consist of addition and subtraction of operands with different bit lengths. Logical operations such as AND, OR, or XOR are also needed.
- Operands involved in either arithmetic or logical operations may be left or right shifted by 1 to 15 bits before use. Sign extension is sometimes specified when shifting to the right.
- Bit-test capabilities are required to perform conditional branching based upon the state of a bit in a computed value.
- Few but lengthy and cumbersome conversions from floating point to sign magnitude notation and vice-versa are required. The algorithm frequently requires the computation of the absolute value or signum of an operand.
- The algorithm requires a significant number of constants for decoding tables, quantizer limits, and masking operations.
- Only a small number of variables need to be stored in RAM. Only rarely is a RAM value used more than once during a single sample period. Most of the variables have lengths of 16 bits or less with the exception of one that is 19 bits.

Performing an analysis of the algorithm after writing a high-level description helps identify the required hardware structures. For the case study, these include specifications like the type and size of the arithmetic logic unit required to realize the data manipulation, the data storage size and width for both fixed and variable parameters, and the type of control structures that manage the sequence of operations.

A data flow representation of the algorithm can help to identify the order in which the functions specified by the high level representation must be executed. The proper execution of a data flow diagram on a sequential machine can be determined by first identifying all of the delay blocks or data sources. A new sample period begins with the information leaving those blocks. Follow the flow of information through the data flow diagram until an input to a delay block or an output is encountered. This terminates the flow in that branch of the diagram. The sample period ends when no more information can flow.

■ Minimum Architecture

The next design step is to define a minimum architecture—the architecture "barely capable" of executing all the operations required by the algorithm. The term "barely capable" applies, because the minimum architecture will probably not be powerful enough to execute the entire algorithm within the actual time constraints. However, by doing the initial coding under minimum assumptions, the designer will have a better idea of what is essential to accomplish the goals. A smaller final architecture results compared to starting with a deluxe version and eliminating resources.

A microprogrammable architecture suits many DSP applications, due to the low cost, high degree of flexibility, and speed of development. A microprogrammable architecture allows the designer to make efficient hardware versus software implementation tradeoffs.

Figure 3 is a block diagram illustrating the minimum architecture judged necessary to implement the ADPCM algorithms. The control structure is not shown because initially only the arithmetic components and data paths are of interest.

Because of the large number of multibit shift operations required in the algorithm, a barrel shifter is placed between the 3:1 multiplexer and the B input of the ALU as part of the minimum architecture. The barrel shifter increases the throughput of the machine and reduces the number of lines of microcode. Most of the shift operations require zero filling rather than sign extension, so the shifter implements zero fill directly.

Two lines out of the barrel shifter are made available for use in conditional branching. Having two lines that can be tested in a single microcode operation gives a significant improvement both in speed of execution and in reducing the number of instructions that it takes to implement a binary tree decoding scheme.

Because the chip is designed to perform both the encoding and decoding algorithm in a single sample period, it must have enough RAM to accommodate the variables corresponding to both functions. The RAM is split into two identical pages with the current page determined by a signal generated in the control unit.

■ Instruction Definition

Instruction definition and execution time are intimately related because the designer cannot define instructions that cannot be executed within the prescribed cycle time. To decide on a target cycle time, the designer sets a goal for the number of instructions that will be required to implement the algorithm. Dividing the sample period by the estimated number of instructions gives the instruction cycle time. The ADPCM algorithm specifies a single sample period of 125 μs (8 kHz sampling rate). We estimated that it would take at least 1,250 instructions to accomplish the required tasks, resulting in an instruction cycle time of 100 ns.

Knowing the instruction cycle time and the speed limitations of the IC process, the designer judges which operations can be performed in the available time. Based on our minimum architecture we imposed the following rules and constraints on the instructions:

- Arithmetic/logical operations can only be performed using the X, Y, or the ACC registers as operands. RAM or ROM values cannot be used directly, rather they must be brought by means of a data move operation into a temporary register, X or Y.
- Data can be moved in parallel with arithmetic/logical operations. The key rule here is that if an operand is used in arithmetic operations it cannot be specified as a destination in data move operations. The only register that can be specified as a source and destination in a single arithmetic/logical operation is the ACC register. This overlapping or pipelining improves the machine's throughput and efficiency but requires widening of the microcode word.
- The ROM LATCH cannot be used as a destination and at the same time be used to address the ROM.

OFTEN A SINGLE CLASS OF INSTRUCTION WILL BE THE LIMITING FACTOR

■ Initial Coding

The initial coding of the high-level representation of the algorithm is a critical step in the design process. At this point the design jumps from a top-down approach to a bottom-up approach. Initial coding begins with a decomposition of the

Call (800) 227-1817, ext. 9781

Breakthrough

Programmable Macro Logic. An architectural innovation that breaks through the barrier of AND/OR design constraints. This unique single NAND array architecture simplifies design and eliminates constraints by delivering total utilization of on-chip resources. Each NAND gate connects with any other NAND gate—there are no interconnect restrictions.

A programmable "instant gate array" with no "gate-a-risk." The high-speed (t_{PD} = 22ns) PLHS501 Programmable Random Logic unit is the first of our powerful new family of Programmable Macro Logic (PML) devices.

It blows away your gate array development time. And with it your NRE, inventory problems and quality concerns. It's programmable or reprogrammable within hours—not weeks. And delivered on schedule, fully tested.

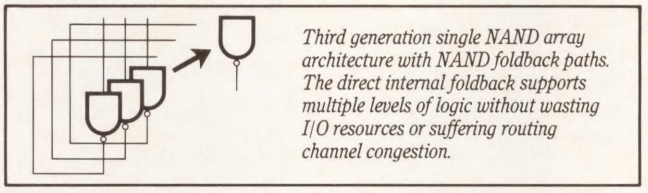

Third generation single NAND array architecture with NAND foldback paths. The direct internal foldback supports multiple levels of logic without wasting I/O resources or suffering routing channel congestion.

Single-chip space saver. The PLHS501 provides 1300 effective gates—a complete solution in a single 52-pin PLCC package.

Software that simplifies. All Signetics PML devices are supported by our powerful AMAZE software that makes complex designs easy by simplifying logic entry, simulation and device programming.

We've got the guts! That's right, we have those essential PLD devices you need to improve total system performance. PAL®-types, PLAs, Logic Sequencers—and our new PML products such as the PLHS501, that simplify complex design problems via innovative architecture supported by powerful software.

Break your bottleneck! Call Signetics at (800) 227-1817, ext. 978I, for a PLHS501 Development Kit, including a free sample.

One standard. 0 defects.

Signetics

a division of North American Philips Corporation

PHILIPS

PHILIPS

CIRCLE NUMBER 22

high-level operations into register-level operations that are executable by the minimum architecture. The goals in this first step are minimum execution time and minimum number of operations. The goal of this step is to implement a sequence of operations that execute the data flow graph in the minimum number of cycles with smooth transitions between major sections of the code.

The transformation from a high-level representation to a format that is specific to the architecture is not a unique process, because of the commutative and associative properties of the arithmetic operations in the algorithm. As a result, the designer has great flexibility in defining the ordering of the operations. This flexibility is restricted by the architecture and the IC process technology. Only after an initial coding of the algorithm on an elementary architecture does the designer have enough information to make reasonable trade-offs regarding the number and placement of the hardware resources.

Although a high degree of flexibility exists in the ordering of some routines, there are also some strict rules relating to the delay routines commonly found in signal-processing algorithms. A given sample delay routine cannot be executed until all other routines have used the variable provided by this delay routine from the past sample period.

A consideration in the optimum ordering of routines is the transition from one routine to another. A perfect transition occurs when a given routine leaves most of the parameter values required by the next routine stored in the temporary registers. This way when the next routine starts it has its input values readily available and does not have to waste cycles obtaining those values from memory.

■ EVALUATION OF THE ARCHITECTURE

After completing the initial coding, the total number of operations required to implement the algorithm is calculated by summing all of the cycles in the critical paths. This value is checked against the value used for computing the instruction cycle time. Several modifications can be made if this total substantially exceeds the maximum.

The first one should be to consider reducing the instruction cycle time, which will increase the number of operations that the architecture can execute in a single sample interval. Often a single class of instruction will be limiting the instruction cycle time. For example, in the case study, executing arithmetic operations on values in the data memories limited the instruction cycle time. If the original instruction definition was just too aggressive, the only alternative is to augment the architecture with hardware components.

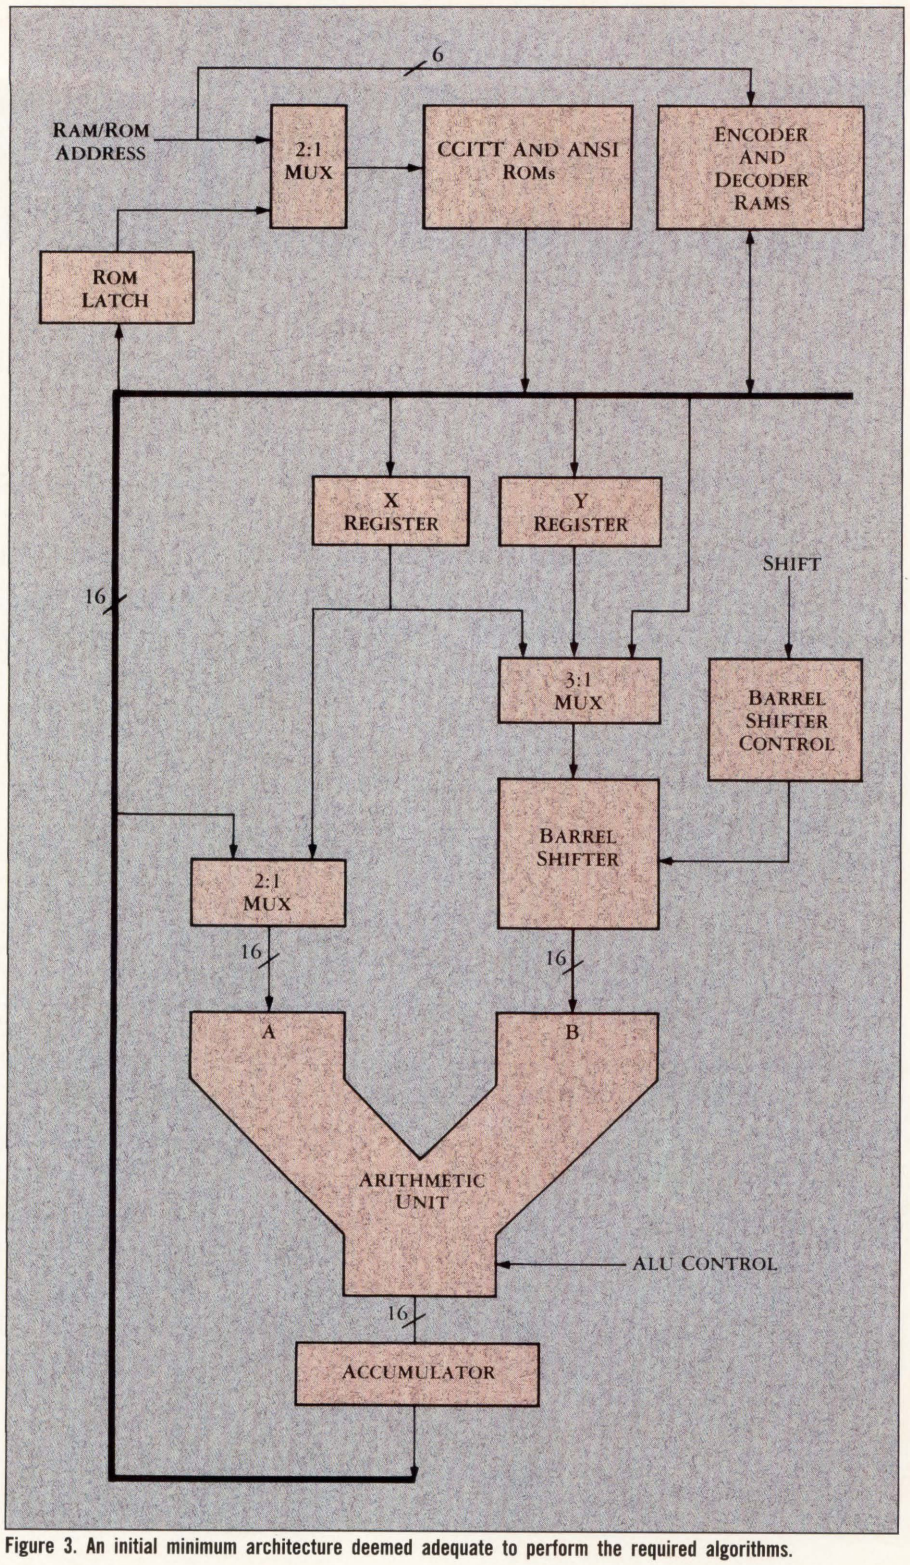

Figure 3. An initial minimum architecture deemed adequate to perform the required algorithms.

In our case, the minimum architecture greatly exceeded the maximum number of operations permitted to execute the algorithm in real time. We generated the following list of ideas to improve the design—improvements which were not readily identifiable when the minimum architecture was originally defined.

1. The FMULT routine, which is used 16 times per sample period, was a logical place to start. Three operations within FMULT were consuming the majority of the cycles. These were the multiplication of the two 6-bit mantissas, the floating-point conversion of the An parameter, and the sign-magnitude conversion of the WAn value.

2. Because memory operands must be

transferred to a temporary register before performing an arithmetic/logical operation, an extra cycle occurs each time a constant is used. Also, the small number of temporary registers leads to increased bus traffic when constants have to be used in a computation. Further, ROM is not being used efficiently because a large number of the constants stored in it are used only once during the algorithm. We discovered that the majority of these constants can be derived by shifting the values $FFFF or $8000 to the right or to the left with the barrel shifter. The addition of two temporary locations containing these values would result in a dramatic reduction in the size of the ROM, plus an improvement in the execution time.

3. Some routines are executed more than once and each time they are executed they operate over a different set of values, suggesting that the overall control logic should be able to call a given routine several times with different pointers to data memory.

4. Because of the similarities in the ordering of the encoding and decoding functions, it would be extremely efficient if both sequences were to be combined into a single sequence of operations. In order to accommodate this feature the control logic must be able to allow conditional branching based on a flag that determines whether the encoder or decoder function is being performed.

5. The algorithm contains many decision points during which no arithmetic operation or data transfer is being performed. During such cycles, most of the resources of the architecture are not being used. In some circumstances, it is possible to perform computations while a branch condition is being tested.

■ Migration from Software to Hardware

The minimum architecture was defined under the assumption that the majority of the functions required by the algorithm could be implemented in software. However, at this point in the design we must augment the available hardware resources. Migrating functions from software to hardware is an iterative process in which additions are made and then evaluated in terms of performance improvement and cost.

Based on the evaluation of the initial coding, the following enhancements were implemented:

Constant ROM and temporary registers. The two constant registers were connected such that their values could be modified by the barrel shifter and used in conjunction with the temporary registers (X, Y, ACC or exponent register) in a single cycle. The use of these registers reduced the cycle count of certain operations but did not bring the overall cycle count of the algorithm under its maximum limit.

FMULT conversion to hardware. Optimization of this routine involves three hardware improvements. The first modification is to implement, in hardware, the conversion from sign-magnitude to floating-point notation. This required a priority encoder, exponent decode logic and register, and a shift control PLA.

The next idea involves the enhancement of the multiply operation with a software/hardware hybrid implementation. The 6×6 and 13×7 multiply operations required can be coded into 4 instructions in a modified Booth's algorithm. The additional hardware consists of an additional temporary register (multiplier register) and an expanded barrel shifter control PLA.

The third optimization to the FMULT routine simplifies the conversion of the WAn variable from floating-point to fixed-point representation. This involves shifting the WAnMANT variable based on the value of the WAnEXP variable.

An SC register was added to hold WAnEXP while the shift control PLA decodes its value to determine the amount and direction of the shift.

■ Control Structure

Up to this point in the design process little emphasis has been given to the DSP's control structure. Rauscher and Adams (1980) give examples of several types of microprogrammable control structures. In the case study, we selected a split control structure because several of the functions required by the algorithm are repeated several times (Bonet and Williams, 1987).

The microcode control store is only 27 bits wide, due to multiple use of the bits (in different contexts). A direct coding of all of the control lines required by the architecture would cause the microcode control store to be extremely wide. The cost of reducing the width of the microcode control store is some further decoding of the fields as they are presented from the micro latch. The post micro-latch decoding slows the machine's cycle time. It also reduces the flexibility of the architecture to respond to algorithm changes. Direct control of the control lines is most advantageous when the algorithms change rapidly. However, if the instructions and algorithms are somewhat set, this decoding can drastically reduce the size of the microcode control store, which is often the largest feature on the chip. Another advantage of reducing the width of the microcode control store occurs when the micro latch can be accessed from outside the chip. In this case the width directly translates to pins on the chip.

■ Conclusion

Throughout this article we have stressed that a major advantage in the design of application-specific DSPs is that the architecture is flexible. Due to this flexibility the designer has the freedom to create an architecture that is optimized in terms of performance and cost. He also has the freedom to "overdesign" to the extreme that the solution is no longer cost-effective. Our intention was to provide the designer with a methodology that can yield an optimum application-specific DSP IC: one that fully implements the algorithm requirements at minimum cost. ■

References

American National Standard T1/87-301. 1985. *32 kbits/s Adaptive Differential Pulse Code Modulation (ADPCM)*, ANSI T1Y1 Committee.

BONET, L., and T.A. WILLIAMS. 1987. "A Split Control Store VLSI for 32 kbps ADPCM Transcoding," *International Conference on Acoustics, Speech, and Signal Processing.*

CCITT Recommendation G.721. June 1984. *32 Kbits/s Adaptive Differential Pulse Code Modulation (ADPCM)*, CCITT Study Group XVIII.

DAUMER, W.R., X.C. MAITRE, P. MERMELSTEIN, and I. TOKIZAWA. 1984. "Overview of the ADPCM Coding Algorithm," *IEEE GLOBECOM '84, Conference Record 23.1.*

RAUSCHER, T.G., and P.M. ADAMS. January 1980. "Microprogramming: A Tutorial and Survey of Recent Developments," *IEEE Transactions on Computers.*

About the Authors

LUIS BONET *received the BSEE degree from the University of Puerto Rico in 1979 and the MSEE from the Georgia Institute of Technology in 1980. Since 1981, he has worked at Motorola's Semiconductor Products Sector, in the areas of electronic tuning and speech synthesis ICs, and digital signal processing architectures.*

TIM A. WILLIAMS *joined Motorola Communications Sector in 1976 after attending Michigan Technological University. He transferred to the Semiconductor Products Sector in 1979 and completed the MSEE and PhD degrees at the University of Texas at Austin in 1982 and 1985, respectively. Tim designs custom DSP architectures.*

E²/DIGITAL/ANALOG
COMBINATION Y

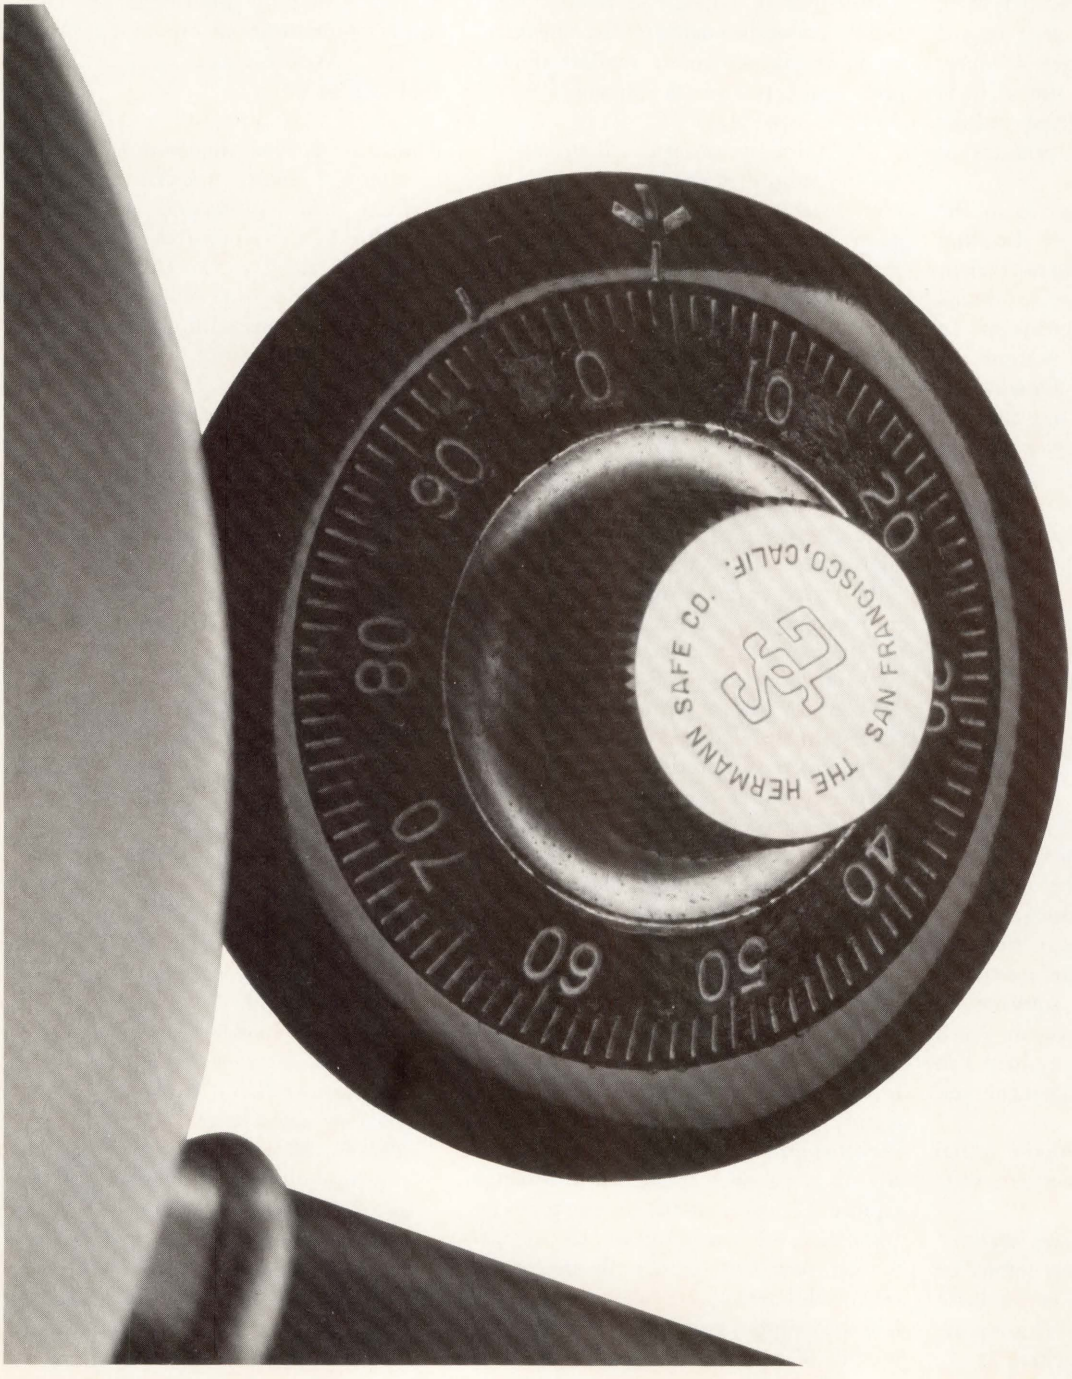

THE TECHNOLOGY OU CAN BANK ON.

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

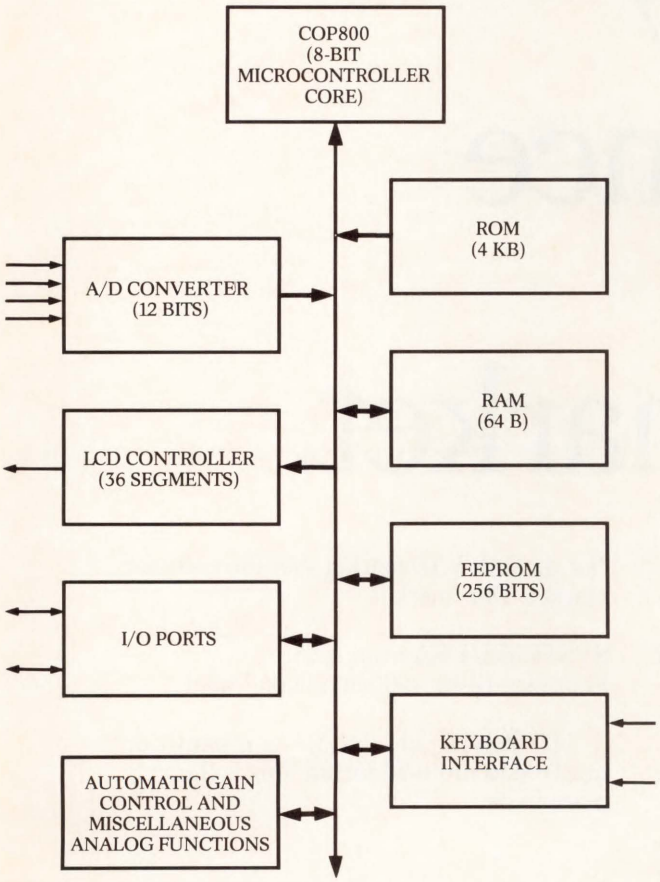

Single chip measurement and control system, integrating several EEPROM, analog, and digital functions.

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

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

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

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

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

Sierra Semiconductor
Triple Technology.™ In CMOS.

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

CIRCLE NUMBER 23

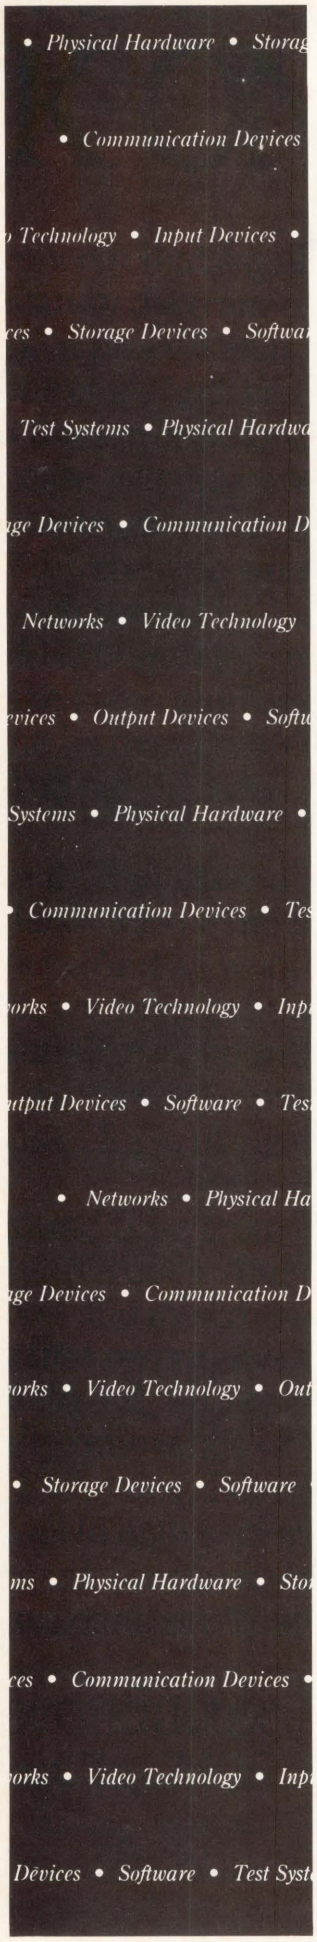

Finally.
A major industry conference for the OEM market.

The expanding computer technology market has created a big gap between COMDEX and WESCON.

SYSTEMS/USA, the first technical conference and trade show focused exclusively on the OEM systems and subsystems market, fills that gap.

It's exclusively for the OEM market. *Not* the electronics components market. *Not* theVAR-VAD market. *Not* the end user market. *Your* market.

SYSTEMS/USA from AEA.
February 13-15, 1990 in Silicon Valley.

If you're an OEM company or involved in product design, this is the only place to be.

SYSTEMS/USA... Sponsored by the **American Electronics Association**

☐ **Yes!** I'm interested in exhibiting at SYSTEMS/USA the first trade show designed exclusively for the OEM market. Send details.

☐ **Yes!** I'm interested in attending SYSTEMS/USA. Send details.

Name _____ Title _____

Company _____ Phone _____

Address _____ City/State/Zip _____

Return to:
SYSTEMS/USA, American Electronics Association, 5201 Great America Parkway, Santa Clara, CA 95054

Design-automation software written for IBM PCs keeps getting better

PC-Based Design Auto Hits The Hot Spots

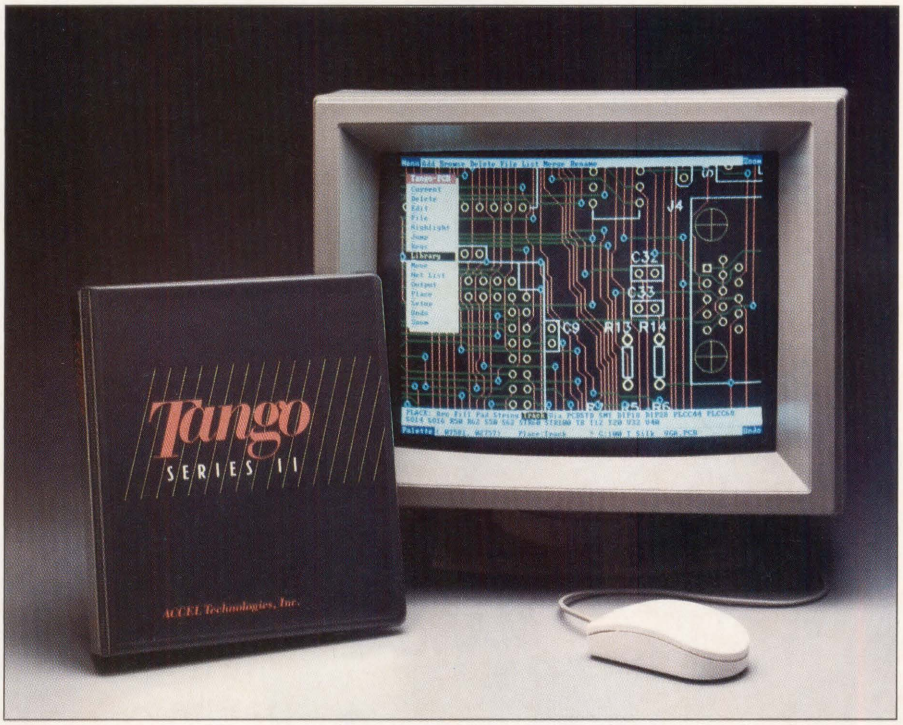

Accel's enhanced Tango-PCB and Tango-Route packages increase the workspace and number of layers

PERSONAL computer-based CAE is getting more adept at almost every point in the design process, from PLD design to board placement and routing. At the front end, designers of ASICs can buy Custom Silicon's Design Kit and create an ASIC design that can be implemented in one of the three IC technologies currently supported. Dubbed vendor-independent ASIC design, the kit contains generic primitives that map to libraries of gate arrays from Motorola, NCR, and Seiko.

After designing in the generic primitives, the designer compiles the design into a target library. The compiled design gives the designer a gate count, while vendor-specific technology files supply timing information for performance comparison. In this way, the designer can compare implementations in the different technologies. At $9,950, the Design Kit includes the generic library, a dozen special software tools, and a choice of two vendor technology files. The kit works with the Workview family from Viewlogic Systems Inc. (Marlboro, Mass.).

Designers using PLDs have two new tools from OrCAD Systems Corp. OrCAD/PLD is used to design devices in conjunction with the OrCAD schematic-capture package or in a stand-alone mode, while OrCAD/MOD adds to the company's simulator the capability to model and simulate board designs containing PLDs.

OrCAD/PLD allows designers to enter PLD designs in several forms: a state-machine language, truth tables, Boolean equations, a schematic netlist, or as a series of indexed equations. OrCAD/PLD then extracts the information from the PLD, reduces the logic, and generates a JEDEC fuse map. In turn, OrCAD/MOD reads a JEDEC file and creates a simulation model of the programmed PLD. The designer can then run a simulation with timing analysis on his PLD-based board design. OrCAD/MOD contains more than 200 timing models of PLDs, and users can develop new models. OrCAD/PLD and OrCAD/MOD cost $495.

For the rest of the board design, both analog and digital engineers have new tools to choose from. First, Personal CAD Systems Inc. (P-CAD) has released a new version of its Master Designer II PCB design software that supports high-speed graphics boards from vendors such as Nth Graphics and Autotasc. The new version also offers features for analog board design. Designers can identify on the schematic critical paths and component groupings that must be maintained during the layout process. Master Designer II products start at $8,495.

Finally, Accel Technologies has introduced the Series II versions of its Tango-PCB and Tango-Route software for PCB design. Accel has doubled the number of board layers to 19, increased the size of the workspace by 50 percent, and replaced the basic grid of nine fixed sizes with a flexible system of three grids (visible, snap, and relative). Track widths can now range from one to 255 mils in width, and the variety of hole and pad shapes and sizes has increased. All of the programs have been bundled together to create a more cohesive environment. Tango-PCB Series II costs $595; Tango-Route Series II costs $495.

Accel Technologies Inc.
San Diego, Calif.
(619) 695-2000

Custom Silicon Inc.
Lowell, Mass.
(508) 454-4600

OrCAD Systems Inc.
Hillsboro, Ore.
(503) 640-9488

Personal CAD Systems Inc.
San Jose, Calif.
(408) 971-1300

Our ASICs

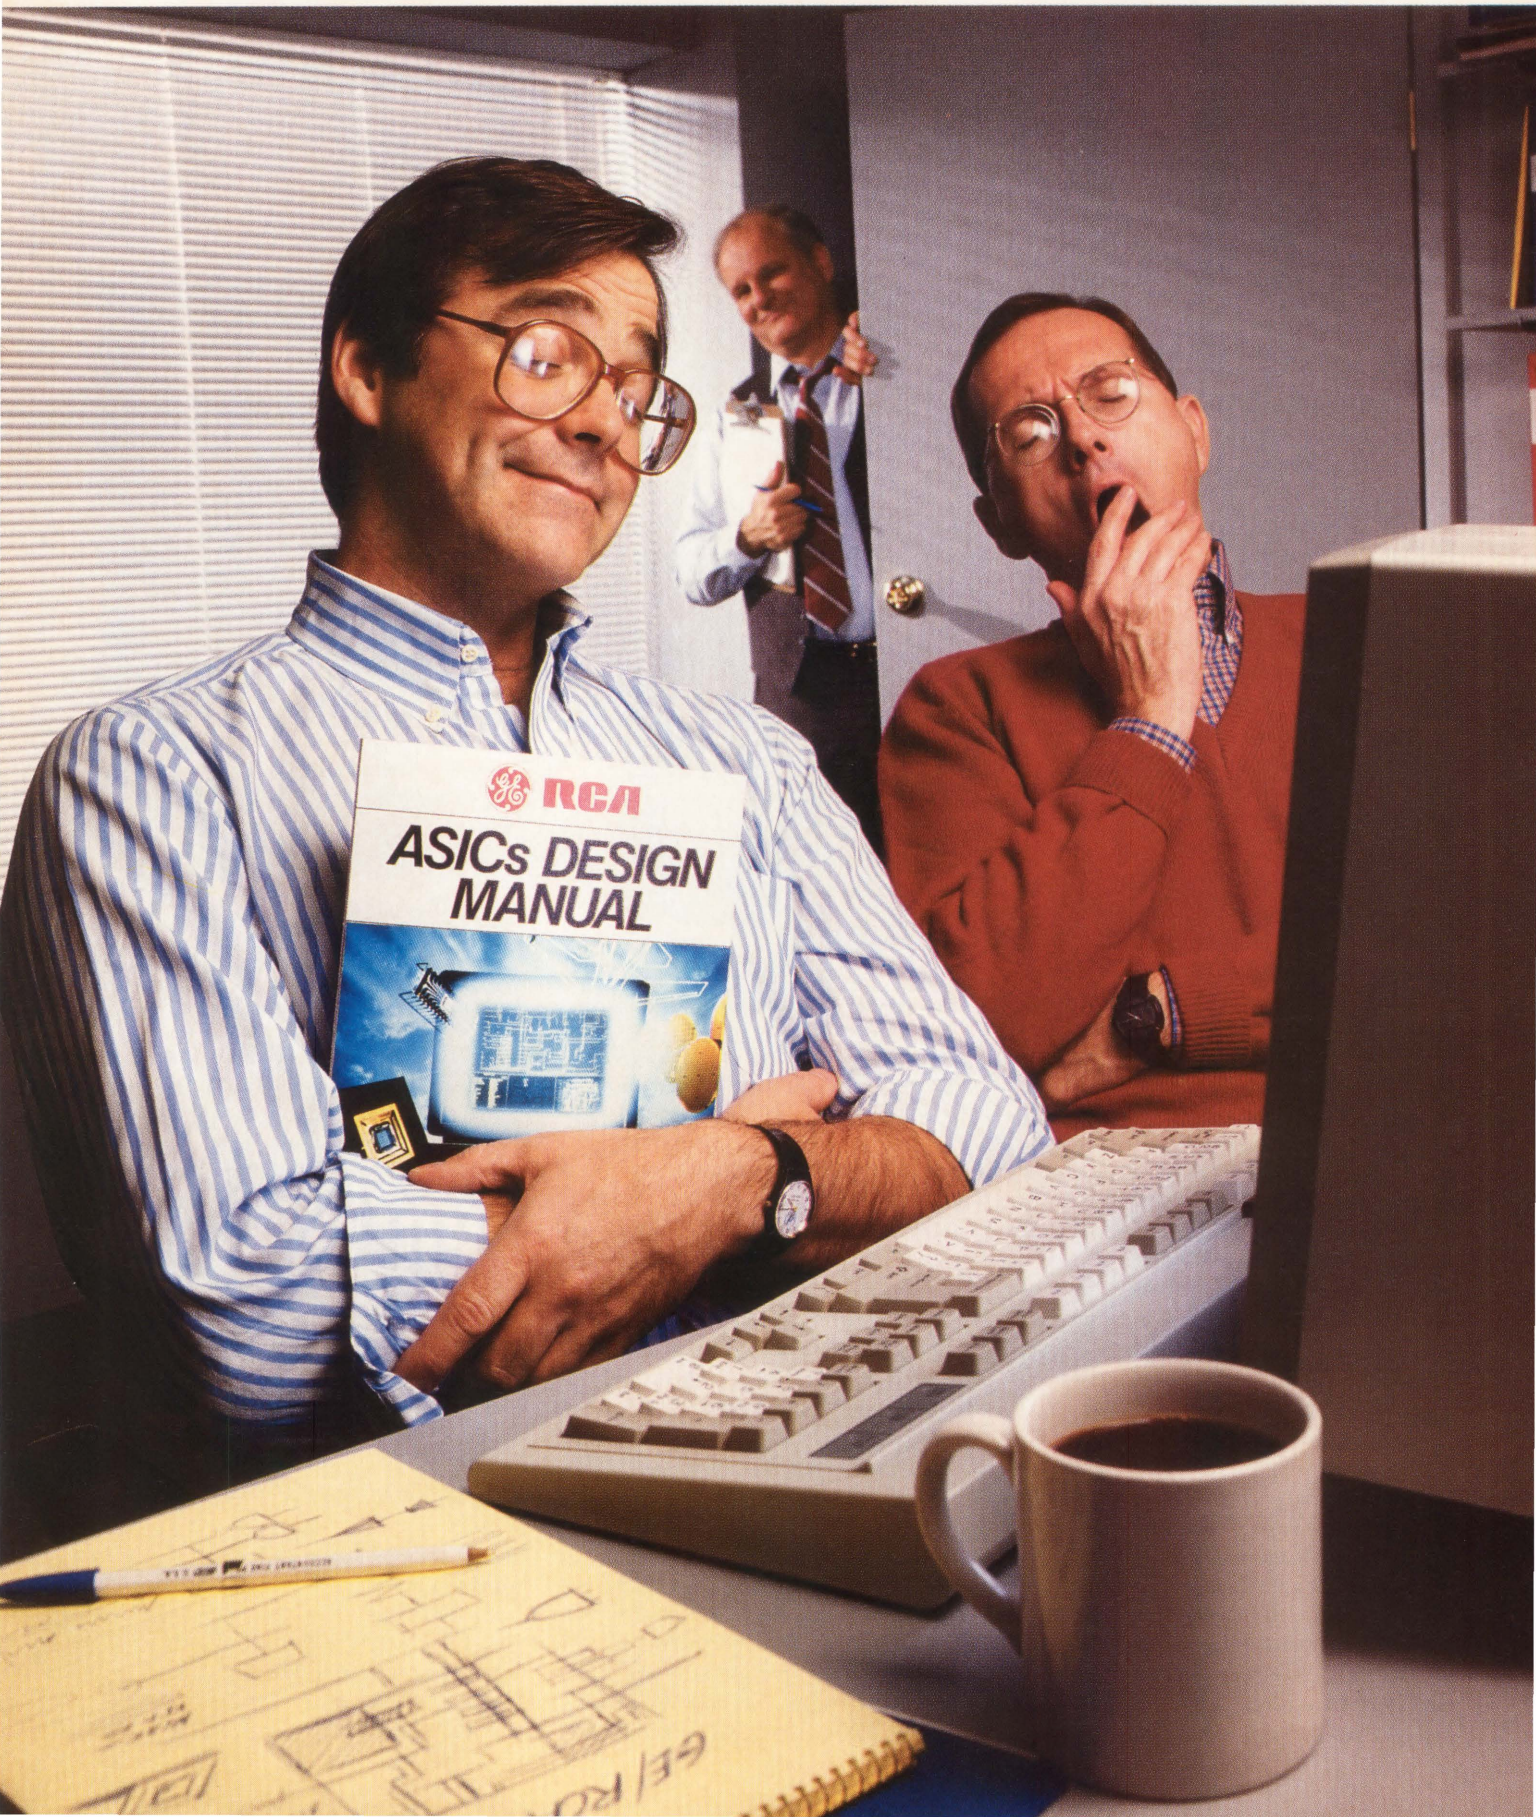

are boring.

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

You've heard all about the excitement of ASICs.

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

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

How to get first-time success.

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

The standard cell advantage.

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

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

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

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

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

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

Gate arrays, too.

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

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

We also alternate source the LSI Logic 5000 series.

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

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

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

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

USA

GE/RCA/INTERSIL
SEMICONDUCTORS

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

Advanced processes enhance gate delays and increase chip densities

One-Micron ASICs Surpass 100,000 Gates

Industry leaders LSI Logic Corp. and VLSI Technology Inc. have driven ASIC integration past the 100,000-gate mark. Taking advantage of processes with 1-μm drawn channel lengths, these companies have pushed cell-based IC densities to 200,000 and 150,000 gates, respectively. VLSI's technology has made possible new gate array dies with 243,000 gates, which, with the company's claim of 30 to 40 percent utilization, can implement designs as large as 97,000 gates. Density limits on the cell-based products are determined by the size of the cells and the largest die with which companies feel comfortable manufacturing.

Migrating to advanced processes enhances the gate delays as well. The 1-μm gates in LSI's LCB007 shrink to 0.7-μm during processing, resulting in a typical gate delay of 450 ps for a two-input gate driving two loads. LSI claims that internal toggle rates are as high as 250 MHz.

The 1-μm gates in VLSI's VSC300 shrink to 0.85 μm during processing, but are specified at a typical gate delay of only 350 ns. While these products seem clearly faster than LSI's, the actual performance depends on the structure and interconnect of the application.

LSI's chips are fabricated with three layers of metal, and therefore may introduce lower capacitance on signal lines than an equivalent implementation in VLSI's two-layer-metal process. VLSI plans to rectify this imbalance by introducing triple-layer-metal processing by the middle of next year. Also, the higher density claimed for the LCB007 may, in part, be the result of the greater density of triple-layer-metal layout. LSI has gained experience in the triple-layer metal through its LCA100K gate arrays.

The LCB007 can also implement a great deal of RAM and ROM storage on chip, up to limits of 144,000 and 1 million bits, respectively. Through LSI's Modular Design Environment (MDE), libraries of standard designs are available. The MDE also contains memory compilers and logic synthesis tools for creating large circuit blocks. LSI's LPACE chip-planning software allows the designer to determine performance characteristics prior to layout. The MDE should be ready now to support LCB007 designs. In addition, LSI has added a 391-pin ceramic PGA to accommodate the pin-intensive designs that will undoubtedly result from such high integration.

VLSI has migrated both its gate arrays and cell-based products to one micron at the same time. Both product lines use VLSI's "bent-gate" architecture (introduced with the VGT200 family) and offer the same level of performance. The VGT300 gate arrays come in seven die sizes, starting at 28,090 gates.

Designs for both the VGT300 and VSC300 can be created with VLSI's Portable Libraries, so named because designs created with its compilers and libraries of cells can be implemented in either gate arrays or cell-based ICs. The cell-based designs, furthermore, can be implemented in CMOS, radiation-hardened CMOS (through an agreement with Harris Semiconductor) or GaAs (through a relationship with Vitesse Semiconductor). The libraries contain memory and data-path compilers, a logic synthesis tool, a RISC processor and other standard designs.

The products will be initially produced in San Jose. Starting next year, the 1-μm processing will be moving to VLSI's new San Antonio facility, which has been designed to produce cell-based ASIC designs in only two weeks.

TABLE 1: NEW 1-micron ASIC PRODUCTS

PRODUCT	LCB007 STANDARD CELLS	VCT300 GATE ARRAYS	VSC300 STANDARD CELLS
COMPANY	LSI LOGIC	VLSI TECHNOLOGY	
MAXIMUM DESIGN SIZE (GATES)	200,000	97,000	150,000
TYPICAL GATE DELAY (PS) (FAN IN = FAN OUT = 2)	450	350	350
MAXIMUM LEADS IN PACKAGES	391	299*	299*
TYPICAL I/O DELAY (NS)	3.0 (50 PF LOAD)	3.5 (25 PF LOAD)	
DESIGN TOOL HIGHLIGHTS	MDE, LPACE	PORTABLE LIBRARY	
FIRST CUSTOMER PROTOTYPES	MID-1989	Q2, 1989	

*HIGHER PIN COUNTS PROMISED FOR 1989.

LSI Logic Corporation
Milpitas, Calif.
(408) 433-4091

VLSI Technology Inc.
San Jose, Calif.
(408) 434-3000

META-SOFTWARE

MAKES THE CONNECTION!

ARE YOU USING . . .

- *Mentor Graphics' IDEA Series*™? Meta's Mentor Server version of HSPICE interfaces directly into the IDEA MSPICE environment.
- *Cadence EDGE™ Design Framework™ System*? Meta's HSPICE accepts full hierarchical netlisting and generates WAVES output.
- *EDA's Electronic Design Management System*? EDMS provides an open framework for electronic design activity incorporating HSPICE.
- *CAECO Schematic*™? HSPICE interfaces directly with CAECO's full-function hierarchical schematic editor.
- *Teradyne/Case Stellar Schematic Capture System*? Teradyne/Case supplies a fully functional CAE package interfacing with HSPICE on standard system configurations.
- *Performance CAD's Circuit PathFinder*? CPF extracts HSPICE netlists of critical paths from large circuits.
- *Analog Design Tools' Analog Workbench*? The Workbench version of HSPICE runs in ANALOG's design and simulation environment, providing access to advanced analysis tools.
- *Interactive Solutions Limited's MINNIE*? Meta's HSPICE interfaces with ISL's interactive graphical circuit design system.
- *IBM VM/CMS*? Meta-Software's HSPLOT high-resolution interactive graphics post-processor drives all devices supported by IBM's GDDM.
- *VIEWlogic® Workview*™? Workview covers the IC, ASIC and PCB engineer's total workday needs, including integrated circuit simulation using HSPICE.
- HSPICE accepts a standard SPICE netlist, making it compatible with most electronic design tools.
- Interfaces currently under development include the IBM Circuit Board Design System (CBDS), mixed-mode analog/digital simulation and more.

NO MATTER HOW COMPLEX THE PROBLEM, META OFFERS THE CIRCUIT SIMULATION SOLUTION!

Software evaluations are available at no charge. For detailed information on Meta-Software products, please contact us!

META-SOFTWARE

Meta-Software, Inc. • 50 Curtner Avenue, Suite 16 • Campbell, CA 95008
Phone (408) 371-5100 • FAX (408) 371-5638 • TLX 910-350-4928
Toll Free (800) 346-5953

CIRCLE NUMBER 19

SAFECON

The Software and Firmware Engineering Conference

WHY EXHIBIT?

Chip manufacturers and vendors of software-development tools (compilers, assemblers, CASE systems, software components, and in-circuit emulation equipment) will find **SAFECON** the ideal forum for meeting quality attendees in an unhurried atmosphere. You'll meet developers specifically interested in software/firmware programming requirements related to high-end microprocessors.

These professional software engineers are producing the next generations of products and are looking for tools to make their jobs easier. Your one-on-one contact with key individuals who can specify your products is an opportunity that will pay off!

WHY ATTEND?

As a software engineer who selects software and tools, you'll get an in-depth education in the latest programming techniques and tools, and tips for using them in your daily work. You'll receive the information you need to make an intelligent selection of processors, support chips, software, and development tools for your application. No matter whether the end product you write code for is a peripheral device or consumer product, you will find **SAFECON** an ideal refresher on the latest real-time programming methods and the application of the latest controller technology.

SEMINAR SESSIONS/TOPICS: Understanding the New Chips (Intel, Motorola, National, AMD, etc...) including Bus Structures, Instructions Sets, Memory Management, Interrupt Structure; Leveraging Processor Performance: Registers, Burst-Made Read/Write Operations to Reduce Band Width; Maintaining Software Quality: Change Control, Testing; Multi-Processing: Peer Processors, and more...

SAFECON

The Software and Firmware Engineering Conference

Writing Code for High-End Microprocessors

February 7 - 9, 1989

Fairmont Hotel, San Jose, CA

WHY SAFECON?

Intense competition has made design turnaround time a major factor in the success or failure of a product. At the same time, the performance of the newest chips depends more and more on efficient firmware. And as endproducts become more sophisticated, writing this firmware gets harder and takes longer.

SAFECON looks at the tools and techniques best suited to the rapid development of high-quality software and firmware. **SAFECON** is the first forum to address these firmware development issues and provide this essential, highly-focused applications knowledge.

For more information on reserving booth space & registering to attend

SAFECON

The Software and Firmware Engineering Conference

call, write, or FAX

Multidynamics, Inc., 13762 Newport, Ste. 204,

Tustin, CA 92680

Phone: **(714) 669-1201**, FAX: **(714) 669-9105,**

MCI-ID: MULTIDYNAMICS

CIRCLE NUMBER 24

INNOVATIONS IN TEST

See them at ATE & Instrumentation Conference West

January 23–26, 1989, Disneyland Hotel, Anaheim, CA

ATE & Instrumentation West is the only conference dedicated to all aspects of test and instrumentation. While other horizontal conferences have some test here and some test there, ATE West has dedicated state-of-the-art conferences and an exhibit hall filled with the products and vendors you want to see.

Applications-oriented technical sessions cover test in design, manufacturing and field support including the latest on Design to Test Integration, In-Circuit Test, SCAN Design, Built-In-Test, VXIbus Instrumentation, Designing for SMT, Testing SMT Assemblies, CAD Power and much more! In fact, there are four days of tutorials, workshops and sessions for you to choose from.

See hands-on demonstrations of IC Functional Test Systems, CAD Translators, ATE based on SCAN Techniques, Artificial Intelligence Software for Board and IC Test, Test Fixtures for SMT PCB's, VXIbus Instrument Systems and much, much more in over 240 booths with instruments, systems and technology for the broad spectrum of test.

Plan now to attend! Call today for a complete show and conference preview.

Registrar, MG Expositions Group,
(800) 223-7126, (617) 232-3976, (between 9am and 5pm EST)

Produced by: MG Expositions Group,
1050 Commonwealth Avenue, Boston, MA 02215

Sponsored by: *Electronics Test, Circuits Manufacturing, EOS/ESD Technology* and *Computer/Electronic Service News* magazines

Endorsed by: The American Society of Test Engineers (ASTE)

ATE & Instrumentation Conference West

CIRCLE NUMBER 25

Advertisers' Index

Reader Service #		Page Number
28	Anritsu	71
7	Applied Micro Circuits Corp.	23
6	Cadence	21
14	California Micro Devices	45
	Cypress Semiconductor	67
	Daisy Systems Corp.	5
27	Data I/O Corporation	CV4
1	Fujitsu	1
3	Gazelle Microcircuits	9
	GE/Solid State	91
10	Genrad	33
	Hewlett-Packard	61
17,18	Hilevel Technology	59
20	LTX Corporation	73,75
	Mentor Graphics	15
19	Meta Software	93
25	MG Expositions	96
29	Micrel	29
	Motorola Inc.	25
	NEC Corporation	53
16	OKI Semiconductor	57
5	Performance Semiconductor	17
11	Plessey Semiconductor	37,39
2	Raytheon	7
24	Safecon	94
13	Saratoga Semiconductor	43
4	Seattle Silicon	13
23	Sierra Semiconductor	87
8,22	Signetics	31,83
26	S-MOS Systems Inc.	CV3
	Systems/U.S.A.	88
21	Toshiba	79
9	USC/MOSIS	32
15	VLSI Technology	49
12	Zilog	41

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

VLSI Systems Design

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

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

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

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

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

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

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

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

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

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

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

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

BPA

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

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

Trust S-MOS. Our ASICs won't leave you out in the woods.

Instead, we'll help you along the path to higher productivity.

Through high-volume, high-yield technology, our manufacturing affiliate Seiko Epson Corp. produces millions of ASIC devices each month.

S-MOS backs up that production with a dependable design program that provides back annotation simulation and fault grades every chip to help your designs succeed.

To keep costs low, there are no CPU simulation charges.

Our full line of ASICs are migratable from gate arrays into standard cells and beyond to our Compiled Cell Custom cell-based designs.

Our ASIC solutions span from 513 to 38,550 gates with technologies down to 1.2 micron (drawn).

To save you time, we can use your existing arrays as future building blocks.

Most ASIC products are available in plastic quad flat packs, pin grid arrays, plastic leaded chip carriers, small outline packages and plastic dual-in-line packages.

So if you're looking for an ASIC program that will get you out of the woods, call us.

(408) 922-0200.

S-MOS SYSTEMS

S-MOS Systems, Inc.
2460 North First Street
San Jose, CA 95131-1002

CIRCLE NUMBER 26

LOGIC SYNTHESIS GIVES YOU MORE DESIGN CHOICES.

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

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

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

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

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

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

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

1-800-247-5700
Ext. 138

Data I/O Corporation 10525 Willows Road N.E., P.O. Box 97046, Redmond, WA 98073-9746, U.S.A. (206) 867-6899/Telex 15-2167
Data I/O Canada 6725 Airport Road, Suite 302, Mississauga, Ontario L4V 1V2 (416) 678-0761
Data I/O Europe World Trade Center, Strawinskylaan 633, 1077 XX Amsterdam, The Netherlands +31 (0)20-6622866/Telex 16616 DATIO NL
Data I/O Japan Sumitomoseimei Higashishinbashi Bldg., 8F, 2-1-7, Higashi-Shinbashi, Minato-ku, Tokyo 105, Japan
(03) 432-6991/Telex 2522685 DATAIO J

DATA I/O Corporation

CIRCLE NUMBER 27