A CMP Publication
EIGHT DOLLARS

VLSI SYSTEMS DESIGN

Tony Zingale on tool integration

Concurrency in analog-digital simulation

Power distribution in sea-of-gates arrays

FEBRUARY 1988

The Magazine for Systems Design Using VLSI Technology

Can VLSI chips help designers tame graphics and imaging systems?

THE POINT OF NO RETURN.

When you put a Fujitsu ASIC to work, you can rest assured it will work the way it should. And keep on doing its job for a very long time to come. In fact, when you look at our performance record over the years, you'll be hard pressed to find any field failures at all.

This is no empty promise. Product reliability has been a way of life for us for more than 15 years.

That's why we always take a conservative approach to the design process. Giving you realistic worst case specs that no production device will exceed.

Guaranteeing a minimum 90% utilization of all gates. And giving you a simulation-to-production correlation of 99%.

It's also why we control every step of the production process. From design to wafer fab to assembly and final test, including 100% AC testing at frequency. So nothing is left to chance.

To us, reliability in the field is everything. And when you remember we've taken over 8,000 ASIC devices from design through mass production, you can see that we'll give you a level of confidence no one else can offer.

So count on parts that have longer life expectancies.

Call our Hot Line today at (800) 556-1234, Ext. 82; in California (800) 441-2345. Look into ASICs you can send out the door. Never to return again.

FUJITSU
FUJITSU
MICROELECTRONICS, INC.
Technology That Works.

CIRCLE NUMBER 1

VLSI SYSTEMS DESIGN (ISSN 0279-2834) is published monthly with an extra issue in May by CMP Publications, Inc., 600 Community Drive, Manhasset, NY 11030. (516) 365-4600. VLSI SYSTEMS DESIGN is free to qualified subscribers. Subscriptions to others in the US: one year $60.00, two years $95.00; Canada and Mexico: one year $90.00, two years $165.00; Europe, Central and South America: one year $120.00, two years $225.00. Asia, Australia, Israel and Africa: one year $150.00, two years $285.00. Second-class (Requester) postage paid at Manhasset, NY and additional mailing offices. POSTMASTER: Send address changes to VLSI SYSTEMS DESIGN, 600 Community Drive, Manhasset, NY 11030.

VLSI SYSTEMS DESIGN was founded to explore, expand, and define the interrelations between very-large-scale integrated circuits (VLSI) and computer architecture, design strategies, costs, and aids, as well as the electronics industry as a whole. VLSI SYSTEMS DESIGN is unique in that it is written by and for the participants in this dynamic field. VLSI SYSTEMS DESIGN intends to be the communication focus of a new VLSI design community, encourage its development, and help define its directions.

Over 41,000 copies of this issue printed.

VLSI SYSTEMS DESIGN

Volume IX, No. 2, February 1988

The Magazine for Systems Design Using VLSI Technology

Cover

It hasn't been easy for VLSI devices to get a rope around the ornery demands of graphics and imaging systems. Evolving standards and seemingly endless customer demand for more throughput have made it difficult for IC vendors and systems designers alike to decide which functions should be set in silicon. Still, graphics processors and floating-point chips look as if they'll remain in the saddle. Computer graphics by Susan Felter, University of Santa Clara, CA.

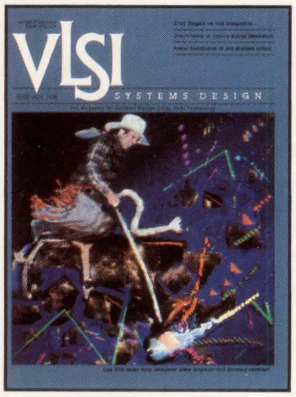

Departments

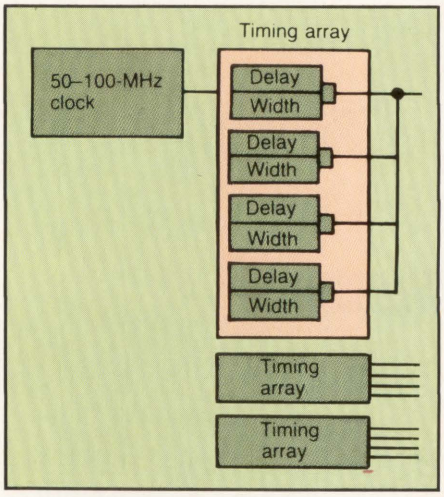

Page 64

Articles

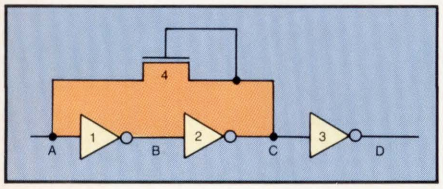

Page 32

16 Aluminum Screen Gives Superior Power Distribution in a Sea-of-Gates Array

Richard Blumberg and Charles Waggoner, *SGS-Thomson Microelectronics*

A 1.2-μm sea-of-gates array with 128,000 raw gates uses an augmented aluminum "screen of power" and a novel controlled-skew-rate I/O design to produce nearly foolproof power busing.

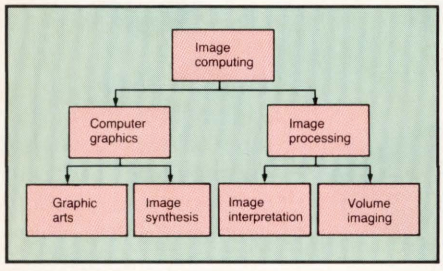

Page 22

22 The Integration of Graphics and Imaging Systems

David Smith, *Western Regional Editor*

Four architectural issues dominate the design of graphics and imaging systems and their constituent VLSI components: the choice of programmable or hardwired functions, variable formats and resolutions of data, the use of building-block ICs or one integrated chip, and the use of parallel processing and pipelines to enhance throughput.

32 A Smart System That Compiles RTL Models from Schematics

Arie Brish, Ronen Keinan, and Yiftach Ravid
Motorola Semiconductor Israel Ltd.

For accelerating simulation, register-transfer-level modeling offers some unique advantages. This article presents a set of model reduction and compaction rules that take a detailed network model and produce an RTL model. The method speeds simulation by an order of magnitude.

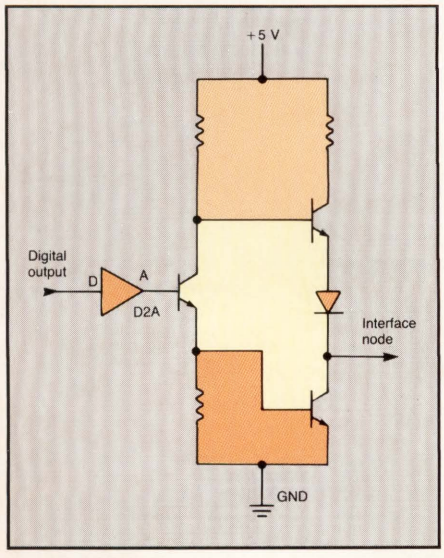

Page 38

38 Coupling a Digital Logic Simulator and an Analog Circuit Simulator

Tedd Corman, *Viewlogic Systems Inc., Marlboro, MA*
Michael U. Wimbrow, *MicroSim Corp., Laguna Hills, CA*

To simulate circuits with analog and digital circuitry, an interface process passes state and timing information between a digital logic simulator and a circuit simulator. Concurrent execution of the simulators and user-configurable interface models let designers simulate mixed analog-digital circuits, even those with feedback loops.

50 1988 Survey of Logic Simulators

VLSI Systems Design Staff

Our annual survey, based on information supplied by each primary vendor of a logic simulation product, highlights the salient features of these critical tools.

Build your gate array on a solid foundation.

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

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

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

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

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

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

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

PACKAGING:
Mitsubishi offers a broad range of industry-standard and proprietary packaging options. A leadership foundation in high pin count surface-mount packages includes quad flat packages (QFP), now available with 160 leads and soon over 200. Mitsubishi provides a cost-effective alternative to PGAs by mounting QFPs on a PC board adaptor.

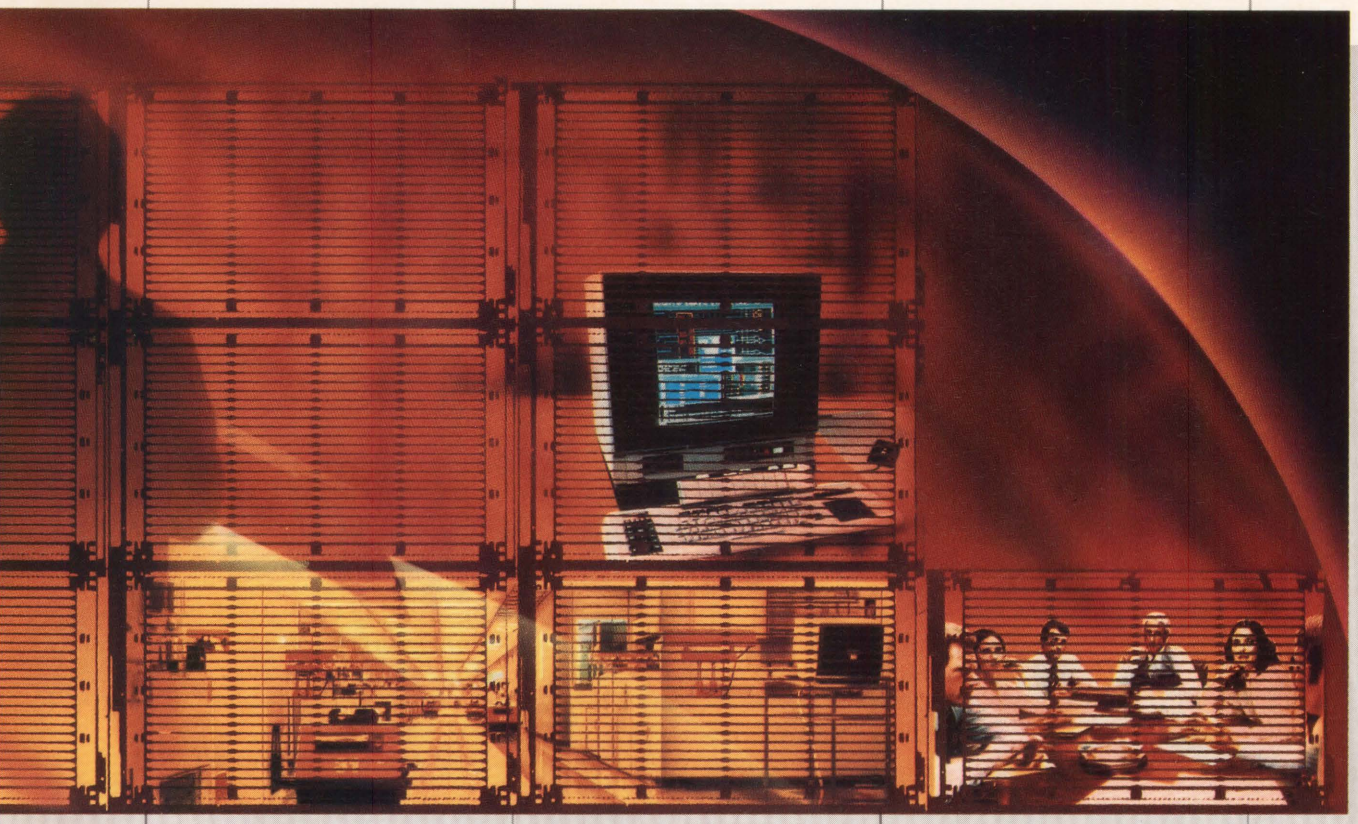

PROCESS:
Low power, 1.3μm, double metal CMOS technology is available today, with sub-micron levels on the way. Mitsubishi's patented* gate isolation structure provides 10% to 20% faster performance, with a 15% to 25% higher gate density than conventional oxide isolation.
*U.S. Patent No. 4,562,453

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

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

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

Quality through commitment. MITSUBISHI ELECTRONICS

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

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

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

CIRCLE NUMBER 2

Breaking Down the Gates

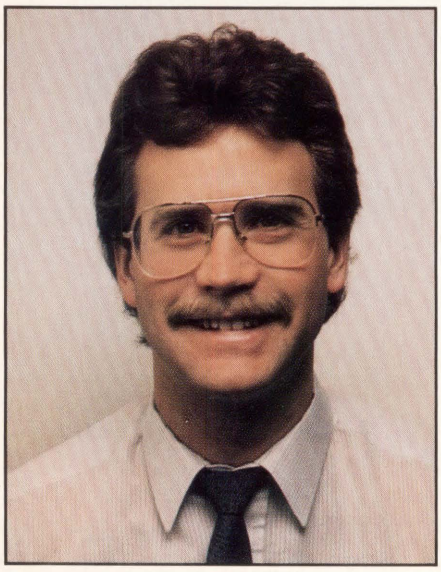

It's getting more difficult for engineers to categorize and compare the various alternatives available for implementing a circuit in silicon. If you design a custom controller chip, for example, you may be able to implement it using a full-custom chip, a standard-cell IC, a "sea-of-gates" gate array, a conventional gate array, a logic cell array, a field-programmable logic sequencer, a registered PAL, an EPLD, an EEPLD, or any number of products with catchy names.

How does an engineer evaluate each type of device for its suitability for his design? He can compare the number of inputs and outputs on the devices with those on his design. He can compare manufacturers' specifications for maximum operating frequency with his design goals. But how can he evaluate the function density of a device whose specific function is not yet defined?

Ideally, the designer should be able to use the product and gauge its potential first-hand. Unfortunately, there are too many products and not enough time to evaluate all the possibilities.

Certain terms do exist to specify the utility of ASICs. The term *gates* is used most often to describe density. Unfortunately, a gate in a VLSI device may be four transistors, the ECL equivalent of a CMOS NAND gate, or the total number of transistors on the chip divided by 4, 5, or 6, depending on the vendor and technology.

Last month in our Industry Insights column, Andy Haines of Actel Corp. argued that a benchmark for density is necessary. He suggested four benchmark functions: a datapath block, a state machine, an arithmetic block, and a counter/timer circuit. IC vendors would use these functions to describe the capacity of their user-specified ICs.

Benchmarks are a convenient way of comparing the capabilities of products. A benchmark should be a generic example of the tasks that the products are designed to perform. Its generic nature is equitable to all products. A designer looking at benchmark results is supposed to get an unbiased evaluation of the product, a true "apples-to-apples" comparison. When properly delineated, such benchmarks can yield more explicit information than a vague term like *gates*. For example, Andy's benchmark includes a counter/timer circuit. If one ASIC can implement four 8-bit counters and a designer has three 8-bit counters in his design, he can be reasonably certain that the ASIC can satisfy his need for counters. On the other hand, he can't compare the number of gates in his three counters with the number of gates in the ASIC and be certain that the ASIC has any counter resources.

The generic nature of a benchmark is also a liability. The benchmark may not exercise special features of a product that may enhance the use of a product for particular applications. By the time a set of benchmarks expands to incorporate all relevant details, it may be unwieldy. In addition, the designer may believe that a benchmark is the final word on density when in fact it is only another yardstick.

Although we acknowledge the problem, we're not completely sold on Andy's solution. Please write or call us with your opinions and, if you agree with Andy, your suggestions for density benchmark circuits.

David Smith
Western Regional Editor

The One Week Analog ASIC.

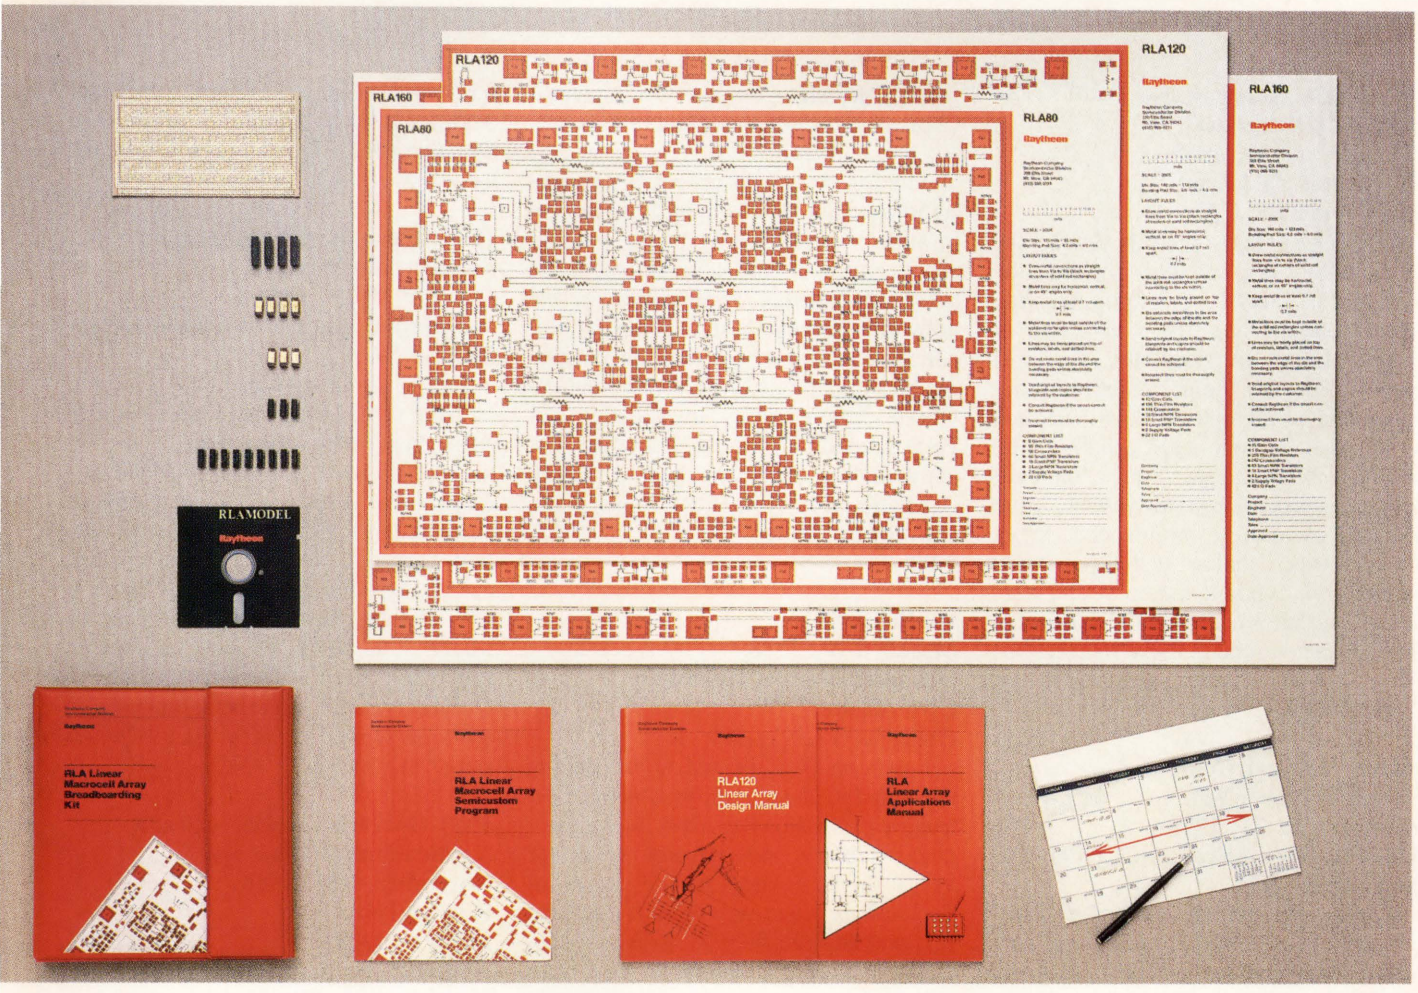

The leading technology . . . the best support . . . go!

The leading analog ASIC family: RLA80, RLA120, and RLA160 user-configurable macrocell arrays. On-chip thin film SiCr resistors and dual-layer metal for ease of interconnect routing and maximum array utilization and performance. Wide supply voltage range — ±1V to ±16V. And simplified design procedures that take days, not weeks.

☐ **Configurable macrocells:** 8, 12, or 15 gain blocks in any combination as general purpose op amps, open collector output comparators, or as input amplifiers with ground sensing function for single-supply systems. RLA160 has a preconfigured on-board adjustable ±30 ppm voltage reference.

☐ **Thin-film resistors:** on-chip SiCr resistors exhibit 1% matching and temperature drift characteristics comparable to discrete film resistors. High performance and high values (1.25 kΩ to 150 kΩ) handle a wide variety of applications.

☐ **Design support:** The RLA Breadboarding Kit contains complete design and applications documentation, 200x plots, a pre-drilled printed circuit board, and 23 ICs. Friendly RLAModel software includes menu-driven user interface program for PSPICE[1], SPICE[2] model library for RLA series, documentation files, and full screen editor. Not to mention the attention you get from applications engineers who have 15 years experience meeting custom and semi-custom requirements.

Call Raytheon for access to RLA program information and details on the RLA Macrocell Array Breadboarding Kit. We promise you a week that will go down in history.

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

[1]PSPICE is a trademark of MicroSim Corporation.
[2]SPICE developed by University of California.

Access to the right technology

Raytheon

CIRCLE NUMBER 3

1988 IEEE International Solid-State Circuits Conference
February 17-19, 1988
San Francisco Hilton
San Francisco, CA

The ISSCC, sponsored this year by the IEEE Solid-State Circuits Council, the IEEE San Francisco Section and Bay Area Council, and the University of Pennsylvania, provides the foremost global forum for the presentation of advances in solid-state circuits. Session topics will include high-speed digital circuit techniques, gate arrays, integrated signal-processing subsystems, nonvolatile memories, high-speed logic, static RAMs, analog techniques, A/D conversion, dynamic memory, and ASIC design and interface circuits. For further information, contact Lewis Winner, 301 Almeria Ave., Coral Gables, FL 33134. (305) 446-8193.

Fourth Annual Computer Graphics New York
February 22-24, 1988
Jacob Javits Convention Center
New York, NY

The Fourth Annual Computer Graphics New York is an exhibition and conference devoted to the latest computer graphics technology. Topics to be discussed will range from advanced turnkey graphics systems to low-cost PC-based solutions. Conference panels will cover applications in business presentations, CAD/CAM/CAE, desktop publishing, biomedical applications, AV and multimedia presentations, aerospace, and television graphics and electronic effects. For additional information, contact David J. Small, Exhibition Marketing and Management Inc., 8300 Greensboro Dr., Ste. 1110, McLean, VA 22102. (703) 893-4545.

Compcon Spring 88
February 29-March 4, 1988
Cathedral Hill Hotel
San Francisco, CA

Sponsored by the IEEE Computer Society and the Institute of Electrical and Electronics Engineers Inc., Compcon Spring 88 will be a broad-based conference for computing professionals. Compcon will begin with a day of tutorials entitled Software Reuse Update, VHDL Tutorial, Computer Architecture Choices, and Software Quality Control. As the conference progresses, it will feature topics such as processors, workstations, artificial intelligence, high-performance computer systems, parallel processing, software engineering, computer-aided design, legal issues, neural networks, and databases. Compcon will conclude with a day of tutorials entitled An Introduction to Artificial Neural Networks, Microchannel Interfacing, and Fault Tolerant Distributed Software. For additional information, contact Hasan AlKhatib, Dept. of EECS, Santa Clara University, Santa Clara, CA 95053. (408) 554-4485.

Southcon/88
March 8-10, 1988
Orange County Convention/ Civic Center
Orlando, FL

Southcon is dedicated to facilitating the transfer of technology among electronics professionals. This seventh annual exhibition and convention will use as its theme "The Leading Edges of Technology" and will feature technical sessions and tutorials, as well as a special purchasing conference. Sessions will include Advances in High Performance Silicon Cache Controllers, Productivity and Computer Applications, Medical Electronics, Programmable Logic Based Sequencers, and PC Building Blocks for Industrial Applications. For more information, call Alexes Razevich at (213) 772-2965.

2nd Annual ESD/SMI Expert Systems Conference and Exposition
April 26-28, 1988
Hyatt Regency
Dearborn, MI

This conference will assemble users with existing needs and applications, suppliers offering expert-system solutions, and researchers who are improving and expanding the capabilities of expert systems. Topics will include analysis, configuration, control factors, design, diagnosis, documentation, engineering, enhancement, integration into mainstream, operations, real-time systems, production, simulation, scheduling, scoping, system development, training, user interfaces, and validation. Additional information may be obtained by calling the Expert Systems Planning Committee at (313) 832-5400.

1988 International Symposium on Microelectronics
October 17-19, 1988
Washington State Convention Center
Seattle, WA

Papers are being solicited for presentation at the International Symposium on Microelectronics, sponsored by the International Society for Hybrid Microelectronics. Topics of interest include automation, hybrid design, interconnections, sensors, fiber optics, microwave, CAD/CAM, and manufacturing technology. Authors should submit abstracts by March 1 to ISHM '88 Call for Papers, P.O. Box 2698, Reston, VA 22090. Questions concerning the submission of abstracts should be directed to Charles Bauer at (503) 627-4908.

1988 Government Microcircuit Applications Conference
November 8-10, 1988
Riviera Hotel
Las Vegas, NV

GOMAC is a government-sponsored conference established primarily to review developments in microelectronics applications for government systems. This year's conference, with the theme, "International Competitiveness: Its Impact on Government Electronics," is soliciting papers for presentation on topics such as VHSIC and MIMIC insertion, computer systems, fault-tolerant systems, ASICs, CAD/CAM/CAE, and signal processing. In addition, solutions to the unique requirements of government microelectronic systems in the following are being solicited: avionics, discontinued parts, quality assurance, and radiation hardness. By March 16, authors are requested to submit double-spaced summaries not exceeding 2 pages to Palisades Institute for Research Services, Attn: Jay Morreale (G-88), 201 Varick St., Rm. 1140, New York, NY 10014. For more information, contact C. Edward Holland Jr., GOMAC-88 Technical Program Chairman, Semiconductor Research Corp., Research Triangle Park, NC 27709. (919) 541-9400. □

Sea-of-Gates Arrays

Oki Semiconductor Inc. (Sunnyvale, CA) and Gould Inc.'s Semiconductor Division (Santa Clara, CA) join the ranks of vendors of sea-of-gates gate arrays. Oki has adopted an architecture reminiscent of that from the now-defunct Integrated Logic Systems Inc. The arrays contain rows of dedicated flip-flops separated by 11 rows of sites for logic gates. Oki claims that its architecture heralds a "third generation" of sea-of-gates arrays whose special structures, including upcoming transmission gates built from polysilicide, create semicustom ICs as dense as those built from standard cells. The new arrays are built with 1.5-μm CMOS technology and can implement designs as large as 50,000 gates.

The GC series from Gould features nine arrays with gate capacities as high as 40,000 gates. The typical utilization is around 35%, so that designs as large as 14,000 gates can be implemented. Utilization is expected to double later this year as new design software becomes available. The parts are built with 1.2-μm drawn channel lengths, two layers of metal, and 100-MHz operation of D flip-flops.

32-Bit Architecture from Japan

Hitachi Ltd., Fujitsu Ltd., and Mitsubishi Electric Corp. have implemented a 32-bit microprocessor and three peripheral chips that support Japan's bid for a new 32-bit architecture. *T*he *R*eal-Time *O*perating-System *N*ucleus (TRON) was developed by Ken Sakamura at Tokyo University to span a wide range of computing machines, from personal computers to supercomputers. It is regarded as the three companies' challenge to Motorola's and Intel's dominance of the 32-bit microprocessor market.

The microprocessor and support chips are called the GMICRO/200 family. The microprocessor is rated at 4 Whetstone MIPS and has a 32-bit, 4-GB physical address space. The most interesting peripheral is the DMA controller, which, with four DMA channels, handles data transfers between two buses and has a maximum transfer speed of 27 MB/s. A cache-tag RAM and an interrupt controller round out the chip set, which operates at 20 MHz and is fabricated in 1.2-μm CMOS. The chips are scheduled to be complete (and, presumably, available for sampling) this summer.

Retargetable CASE Tools

Retargetable software development tools from Quantitative Technology Corp. (Beaverton, OR) can be customized by designers to produce code for custom hardware architectures. The Software Foundry comprises a compiler, an optimizer, a simulator, a symbolic debugger, and a combination of assembler, linker, and formatter.

The assembler-linker-formatter, available now, is customized by the designer, using a proprietary hardware description language. Once the designer has specified his architecture, the assembler produces object code for conversion into a program module by the linker and formatter for the target architecture.

The optimizer (available in May) compacts the microcode, taking into account parallel and pipelined operations specified in the designer's input of his hardware design and using a method called loop folding. It operates both automatically and interactively.

The rest of the tool suite will be available in September. It will cost $50,000 for a workstation and $90,000 for a VAX 8600 or 8800; it runs under the UNIX and VMS operating systems. Individual tools can be purchased separately.

Design Tool Kit

Sun Microsystems Inc. (Mountain View, CA) is releasing two more workstations based on its RISC processor, SPARC. The Sun-4/110 and -4/150 are similar to the Sun-3/110 and -3/150 models, with a SPARC CPU board replacing the 68020-based CPU board of the earlier workstations. The new machines are rated at about 7 MIPS and start at $18,900. A Sun-4/110 with a floating-point accelerator, 327 megabytes of disk storage, tape backup, and a 19-inch monochrome monitor costs $30,200. The Sun-4/150, which has additional slots for graphics options, will be available this summer; its introduction is delayed because of difficulty in getting memory devices.

With ACCESS, a software-resource controller program from Valid Logic Systems Inc. (San Jose, CA), a user checks application software out from a network server and returns it when he is done. In contrast, most design software requires the user to be working on the network node that contains the software. ACCESS makes design software a shared network resource; a complete range of design software can be available to all users on a network without buying a version for each workstation. The company estimates that for an environment with 35 workstations, ACCESS can save $125,000.

Aldec Co. (Newbury Park, CA) has reduced the price on its IBM PC-based SUSIE simulator to just $295 and has removed its software protection. New options include modeling of PLDs from JEDEC fuse maps ($795), a library of 160 types of RAM devices ($795), and a model compiler that converts Boolean descriptions of a device into assembly language models ($195). The simulator accepts designs created with the schematic capture tools from OrCAD.

SIMCOMPARE, from Logic Automation Inc. (Beaverton, OR), compares the output from different simulators and flags significant differences, according to parameters defined by the user. The results are formatted, classified, and presented as either functional or timing differences. The first license (for Apollo workstations) costs $8000; subsequent licenses cost $2000.

33-MFLOPS Chip Set

A CMOS floating-point chip set from Integrated Device Technology Inc. (Santa Clara, CA), the IDT721264 multiplier and IDT721265 ALU, can execute 32-bit operations at 33 MFLOPS and 64-bit operations at 25 MFLOPS. It conforms to IEEE Standard 754 version 1.0 and is pin-compatible with Weitek's WTL 1264/1265 chip set.

The chips have a pipelined architecture that operates with 30-ns clock periods. Flow-through latency is 180 ns for 32-bit arithmetic and 270 ns for 64-bit arithmetic. Single- and double-precision operations run at 16.7 MFLOPS in the ALU; single- and double-precision multiplication occurs at 16.7 and 8.3 MFLOPS, respectively. The chips come in 144-pin PGAs and cost $406 each (100). □

Our ASIC Program targets all your design needs.

In gate arrays, standard cells or Compiled Cell Custom, you'll always score with S-MOS.

Complete design tools.

S-MOS SYSTEMS supports a large and growing number of workstations, including Daisy, Mentor, Calma and PC-based systems using FutureNet, Orcad and Viewlogic. Our proprietary LADS simulator is one of the fastest in the world, and our engineering support will keep you ahead of the competition.

Advanced manufacturing.

Our manufacturing affiliate, Seiko Epson, is one of the world's most advanced CMOS IC manufacturers, with 18 years of CMOS experience.

At our highly automated Class 10 facility in Japan, Seiko Epson can produce 40,000 4" wafers, 20,000 5" wafers and 20,000 6" wafers. Each month.

That's a total equivalent of about 23 million units a month.

High quality and total reliability.

With a product reject rate of less than .0001%, S-MOS and Seiko Epson deliver the ASICs you can depend on.

For a fully integrated ASIC program, call us. **(408) 922-0200.**

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

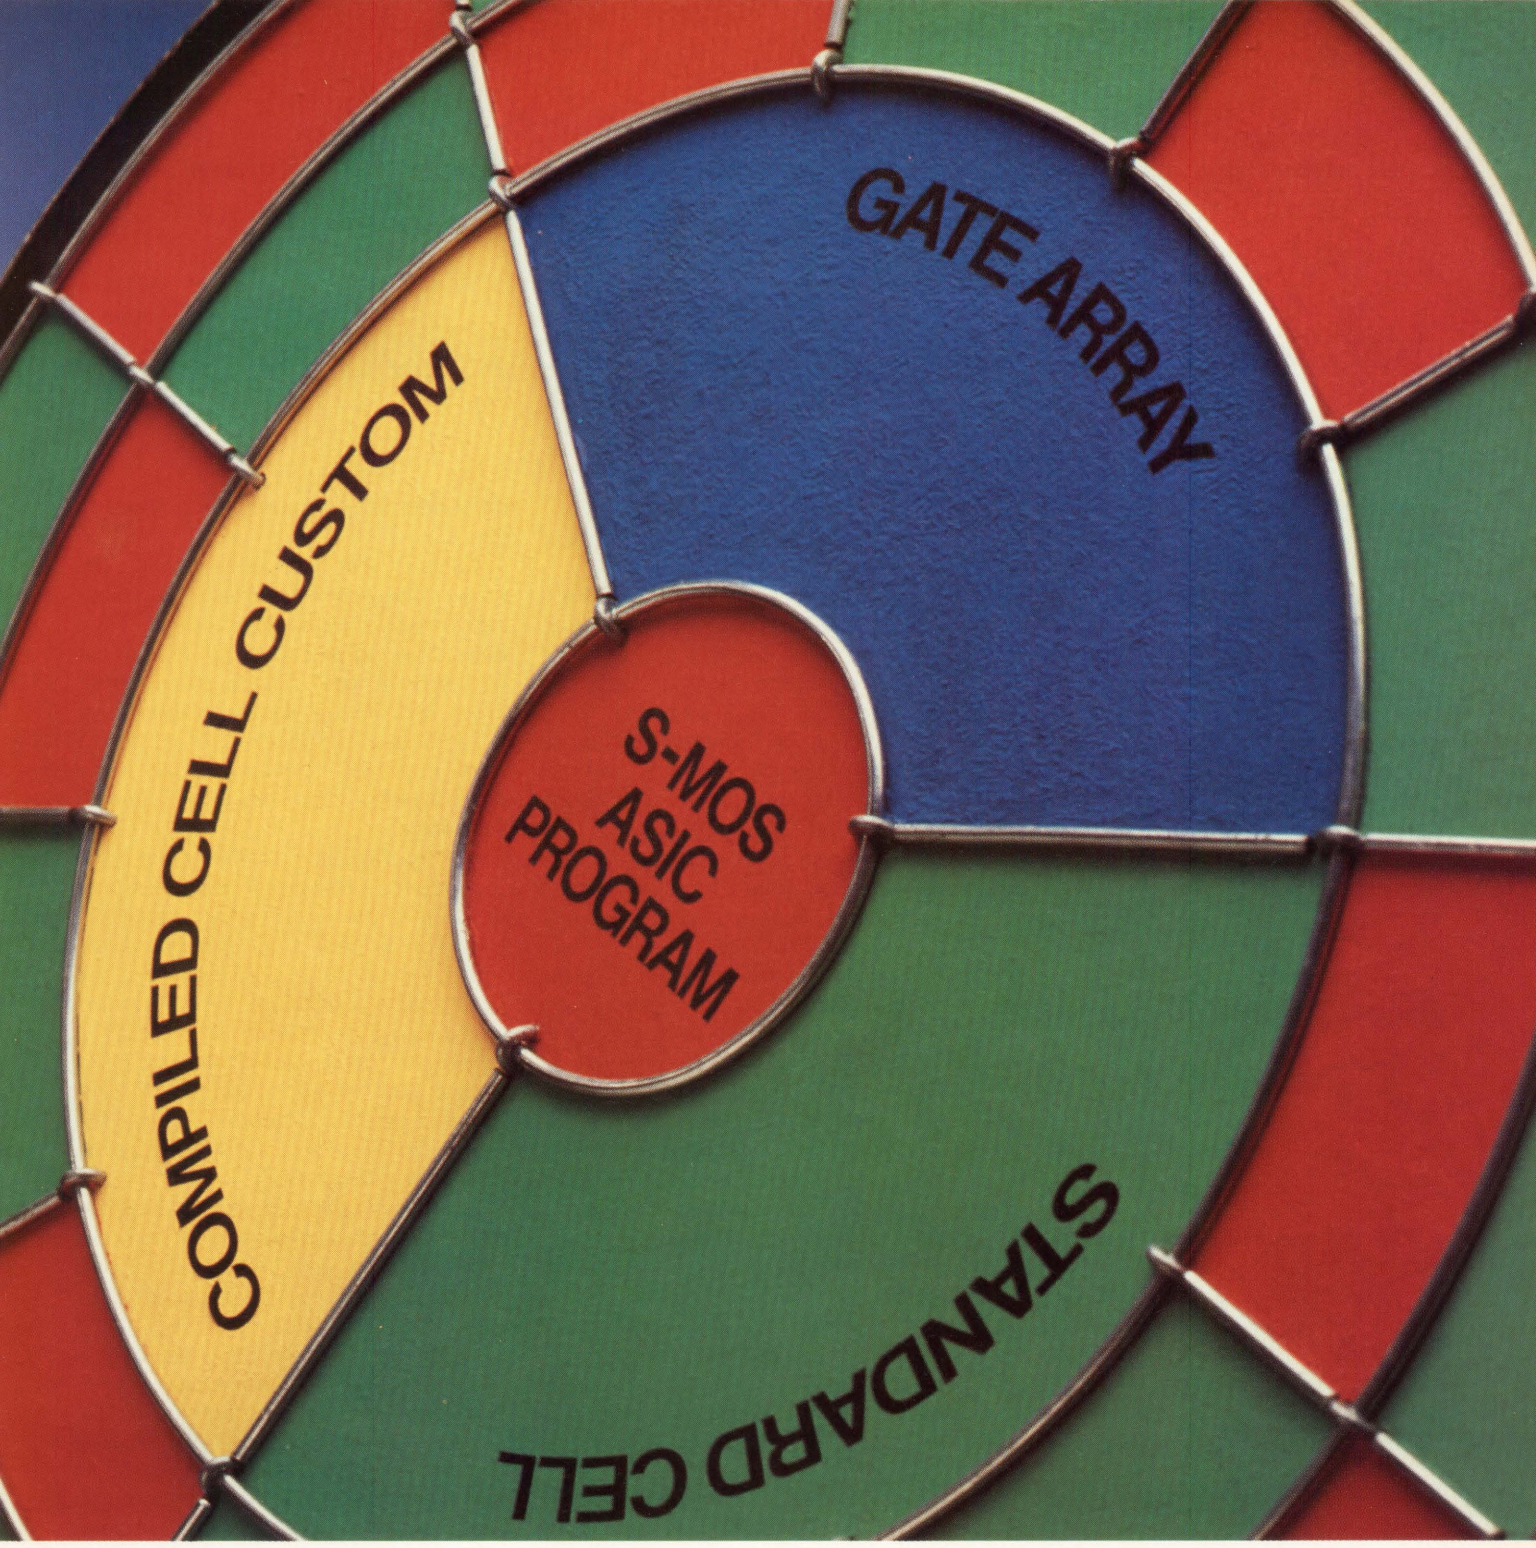

CMOS Gate Arrays	CMOS Standard Cells	Compiled Cell Custom
Complexities to 16K gates • SLA8000 (500ps)* 1.2μ drawn, 0.95μ Leff • SLA7000 (640ps)* 1.5μ drawn, 1.2μ Leff • SLA6000 (1.3ns)* 2.0μ drawn, 1.7μ Leff • SLA5000 (2.2ns)* 3.0μ drawn, 2.3μ Leff	Complexities to 16K gates • SSC1000 (720ps)* 1.8μ drawn, 1.5μ Leff • Fully migratable from S-MOS gate arrays • RAM and ROM blocks available	*The* Alternative to Full Custom • 1.5μ CMOS Process • Can utilize dissimilar cell geometries • 3-button approach to custom design • Currently over 300 fully characterized cells • Fast 14-week implementation time • Timing-driven TANCELL® Place-and-Route Software

*Typical propagation delay of 2-input NAND gate driving 1 internal load including 100μ interconnect.
®TANCELL is a registered trademark of Tangent Systems.

SMOS SYSTEMS
A Seiko Epson Venture

CIRCLE NUMBER 4

John Eurich Brings Engineers Together

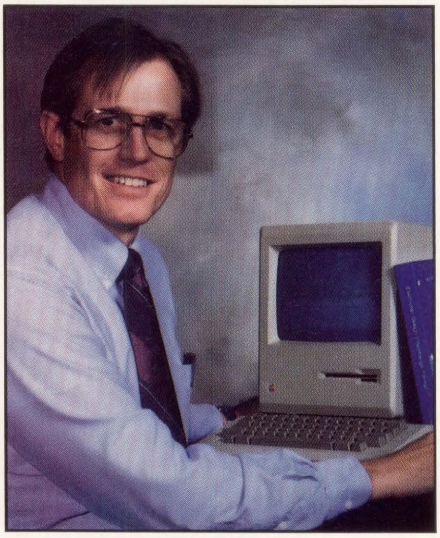

John Eurich has the ability to create consensus from a group of disparate ideas. In a room of people, brainstorming, he can often "put together the pieces" and build an understanding.

"I was never an R&D soloist," he says; "I've always been a team player." His career has been a string of assignments to pull together people and ideas to get the job done. His new job—president of DataXpress, a three-man company developing translation tools for engineering databases—builds on that string. DataXpress is a focal point for engineers needing to pull together design environments.

John himself is a focal point of ideas, a trait that has made him an effective manager. Even off the job, he reads textbooks instead of, say, novels for relaxation. Afraid he may miss important insights by not using his time wisely, he believes that "you never know what will trigger an idea."

John began to put together the pieces of his specialty after his education at San Diego State University and in the Navy. As a "grunt programmer" at Xerox in El Segundo, CA, he designed interfaces between the company's fledgling CAE system and its manufacturing floor.

"I lived the problem" of database translation, he says. "It was hell."

For example, an interface between a CAE system and a piece of autoinsertion equipment must provide the physical height of every component, as well as its coordinates on the board. As John learned to create the interfaces that made such data interchange possible, he rose to become the manager of the group developing VLSI CAD tools.

In subsequent posts at Calma and then Daisy Systems, John continued to build design teams and manage technology development. While at Daisy, he was pulled into the committee developing the then embryonic Electronic Design Interchange Format, EDIF. Not surprisingly, he ended up heading the committee.

Upon leaving Daisy, John turned his hand to serving as a consultant about the use of EDIF as a translator for design environments. After almost a year, "I got so hooked on the whole problem that I started a company to continue the effort," he says. He founded DataXpress in early 1986.

"We want to own data translation," he boasts, although sometimes it seems that EDIF owns John. "I got caught up in data translation because of my work in EDIF. As long as there seems to be a path to take, I'll be taking it."

John explains that there's more to establishing EDIF than just coming up with the present version, EDIF 2.0. Expecting this version to leap into immediate use is "like handing someone a dictionary and telling him to learn English." To develop an EDIF translator, you have to have a good working knowledge of the standard.

DataXpress's flagship product is the EDI (Engineering Data Information) translator development environment. This environment includes an EDIF reader (parser), an EDIF writer, a syntax checker, a semantics checker, and the EDI database, which has its own procedural interface for retrieving and storing data.

Users of EDI need only develop the readers and writers from their own database to the intermediate format. They don't have to worry about reading or writing EDIF files, creating data structures, and moving data into and out of the data structures (because of the procedural interface). "Let us worry about the more esoteric details of EDIF syntax," John says.

An intermediate format may seem redundant. When an interface designer has to map one data format to another, however, he won't get data out in the right format and sequence that he needs. As a result, interface designers spend most of their time creating a parser and determining what the temporary data structures will be. Designing the best data structures to solve the mapping problem is, according to John, "gut-wrenching."

He defends the EDIF standard against those who still think it is insufficient for translating CAE data. "EDIF can transfer symbols and logic models just fine," he says.

He admits that there's a need for a standard for symbols, but whatever standard results, EDIF can support it. He suggests that one CAE company should offer its symbols as a standard, telling its customers to insist to other CAE companies that they support the same symbols. A consensus-building snowball effect, similar to that which doomed copy-protected software, should result.

"There is an EDIF effort because the industry has a need to transfer information," John reiterates. Engineers need to share design data with other engineers, design teams, and even engineering disciplines. Defining the methods and formats to make the interchange easy and accurate won't be easy, but John Eurich is positive, even reassuring: "I know there's a solution; we just have to get together and find it." *—David Smith*

WHAT YOU SIMULATE IS WHAT YOU GET

Be sure product performance matches simulation

Get Optimal Accuracy in Circuit Simulation with HSPICE and ATEM

Accurate model equations and algorithms provided by **HSPICE**

Model Parameters derived from Silicon by **ATEM**

Advanced equations and algorithms have made **HSPICE** the logical choice for accuracy in simulation throughout the industry. Meta-Software now complements **HSPICE** with the **ATEM** system for generation of optimized model parameters.

ATEM is based on the **SUXES-10** parameter optimizer. The **MetaTestChip** provides all the data needed by SUXES-10 for complete HSPICE models with full temperature compensation. ATEM subsystems automatically take the necessary data from the MetaTestChip by driving HP and Tektronix laboratory equipment.

With ATEM, a complete process characterization can be completed in a single day.

ATEM provides complete and optimized **CMOS, BIPOLAR** and **GaAs** model parameters. If you are using **BiCMOS** and need accuracy in simulation, ATEM and HSPICE are the ideal answer.

CRITICAL PATH TIMING ANALYSIS: For path timing analysis of very large scale digital MOS designs, consider **Circuit PathFinder (CPF).** CPF has been successfully used to analyze circuits of over 500,000 transistors. CPF uses transistor level simulation based on HSPICE models for accuracy.

CPF provides an easier way to determine changes in critical path timings introduced by the layout process.

NEW HSPICE FEATURES

- Discrete Component Models based on lab measurement:
  - * A sample library is included free with HSPICE
  - * A large library is available for a fee
- Measure Statement: simplifies Cell Characterization
- BSIM, CASMOS and Fluke MOS models
- Statz GaAs model
- Additional temperature compensation for Bipolar and MOS models
- Geometry-based MOS model selector
- Wire model
- Pole-Zero analysis
- Automatic Op Amp Macro generator
- Switches: Voltage controlled resistors
- Ideal Inverters
- "FAST" option for greater speed
- Interfaces to Tektronix and Mentor workstations

In-house evaluation copies of HSPICE are available to qualified customers. Contact Meta for details.

Meta-Software offers laboratory services for process characterization, based on the MetaTestChip, and for discrete device modelling.

Meta-Software, Inc.

50 Curtner Ave., Suite 16 • Campbell, CA 95008 • Fax: (408) 371-5638 • (408) 371-5100

CIRCLE NUMBER 16

Tools vs. Data: The Debate for '88

Tony Zingale, EDA Systems Inc., Santa Clara, CA

Tool users and vendors tend to focus on tool functionality—and rightly so! Functionality is the heart and soul of a good design tool. Productivity, efficiency, and migration also are fundamental issues—although perhaps secondary to functionality. However, while these four aspects of a tool can be improved, the data generated by the software is really more important to users. Isn't the data where the real value of the design is held? How is the data generated by the design process tracked and management? What tools focus on that dimension of design automation?

Data management issues have been festering in the design automation community for several years, for good reason. The Technology Research Group reports that the installed value of design data in 1987 was more than 10 times the value of the installed base of tools. Unfortunately, tool vendors are only just realizing the significance that the design data holds to users.

The value of the tools versus the value of the data the tools generate is best understood by the replacement cost of the two. To replace a tool, a user can purchase another copy of the software, with a minimal amount of down time. On the other hand, replacing the design data, if it is possible, requires an incredible time sink. In most cases, a company cannot afford the time.

Let's consider an example. A company buys a schematic entry package with its own user interface and a proprietary database from one vendor. It then buys a logic simulator with a textual user interface and a high-performance application database from another vendor. To exchange data, the buyer uses netlist interfaces or database translators, both of which are time-consuming to develop and even more costly to support and maintain. Companies usually focus on making the tools work instead of on maximizing the productivity of the entire process.

Some tool manufacturers have acknowledged that their tools will be used with tools from other vendors. Even so, users cannot be assured that a single supplier will provide the best tool for a future project. And even when one vendor has all the right tools, the design data is not necessarily well managed within the vendor environment. Users still have to distinguish between *data compatibility* and *data management*.

To achieve a system with real data management, a different approach is required. Some sort of electronic design automation "framework" is needed to support the infusion of tools from any source, even in-house proprietary tools, and manage them in the context of the entire design process. The framework should house design management tools that track the intricate relationships between the data, tools, projects, and users. It should be flexible enough to allow change. It should accommodate the reuse of portions of a project completed with a tool set from a previous project. Finally, it should be able to run on the user's choice of platforms and conform to the user's specifications.

Who best supplies such a framework? There are four potential sources: a design tool supplier, a computer manufacturer, an internal development team, or a third party who supplies only the framework and nothing else.

Design tool suppliers and computer manufacturers inevitably force a migration path on their users, or at least give a certain direction to them. Their main interest is in protecting the software or hardware purchases made through them, not to protect a user's data investment.

Some users therefore develop a solution internally. Many users can recount the nightmares associated with developing such a framework. It is a mammoth effort and a major time and resource drain. In most cases, the framework must be reworked for each project, because the problem is dynamic whereas the internally developed solution is static. The user's expertise is in designing, not in making the tools more manageable.

Independent framework suppliers, who supply only framework tools, have no vested interest in any particular tool or platform. It is their charter to provide an overall "open" framework that will let users plug in and manage the data on any tool they want. It is their charter to accommodate new and future tools as they become available and to manage those tools independent of platform, tool vendor, and database structure.

If progress is to be made, the data must be managed, the tools must be managed, and the overall design process must be managed. Framework suppliers are paving the way for that to happen in 1988. □

Tony Zingale joined EDA Systems as director of business development in September 1986. Prior to that he spent three years in various management positions at Daisy Systems. Before that he worked for three years at Intel.

SILVAR-LISCO Performance Redefined.

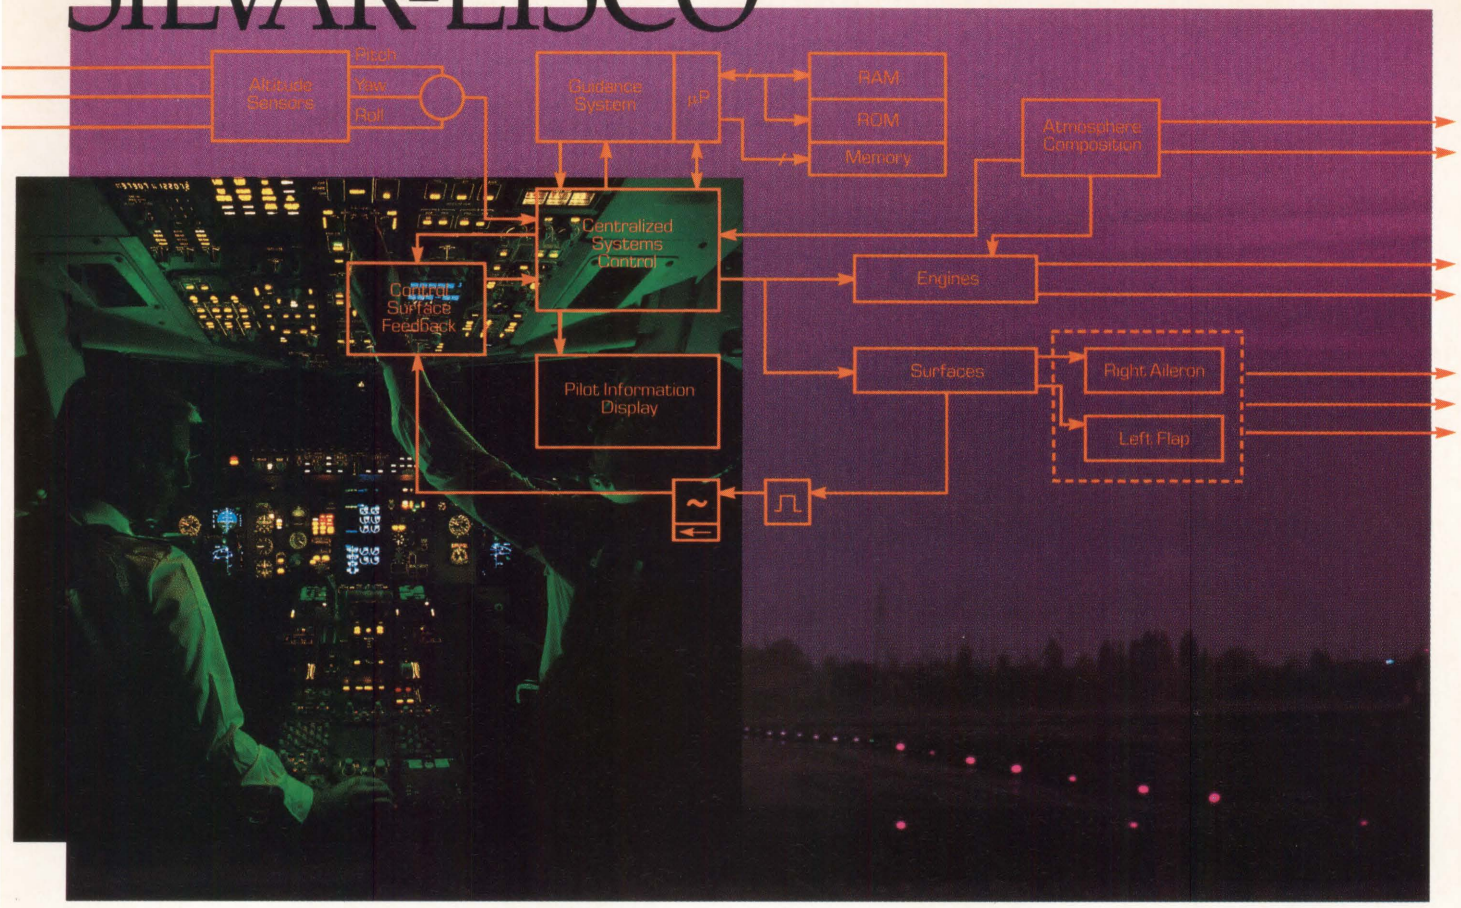

For complex designs, HELIX™ is the "what if" simulator!

When designing large systems, printed circuit boards or VLSI/VHSIC chips, simulation becomes an invaluable aid. Logic, switch level and circuit simulators are excellent tools for validating design implementations. Silvar-Lisco's Helix Behavioral Simulator takes you one giant step further. In addition to design validation, system architects now can optimize the design itself through analysis of various implementation alternatives.

With Helix you can define and test your system concepts first. If the concept works, you move down to the register level. Then the gate level. Multi-level throughout. If it doesn't work, you can rethink the concept and probe Helix with additional design alternatives. In the end, you will get the best design. The first time.

Helix is a vital component of Silvar-Lisco's Architectural Design Series. Incorporating design capture, simulation and design libraries, the Architectural Design Series gives you the competitive edge.

Be a design winner! Contact Silvar-Lisco today!

SILVAR-LISCO OFFICES

Corporate Headquarters
1080 Marsh Rd.
Menlo Park, CA 94025
TEL: (415) 324-0700
WATS: 1 (800) 624-9978
TWX: 910-373-2056
FAX: (415) 327-0142.

European Headquarters
Leuven, Belgium
TEL: 32-16-200016
TWX: 221218
FAX: 32-16-236076

Far East
Nihon Silvar-Lisco
Tokyo, Japan
TEL: 81-3-449-5831
FAX: 81-3-449-4040

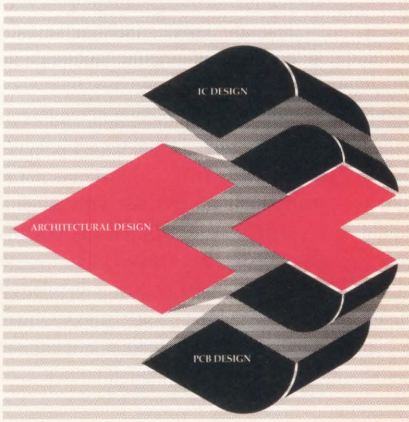

Architectural Design Series

SILVAR-LISCO

CAE Software products for electronic design.

CIRCLE NUMBER 5

Aluminum Screen Gives Superior Power Distribution in a Sea-of-Gates Array

Richard Blumberg and Charles Waggoner, SGS-Thomson Microelectronics, Carrollton, TX

Poor power distribution schemes can cripple even the best sea-of-gates designs. They can result in difficult-to-test designs with poor yields and poor reliability. A good power distribution structure and methodology are far more likely to yield reliable first-pass silicon.

In designing a 1.2-μm sea-of-gates array with 128,000 raw gates, SGS-Thomson Microelectronics used an augmented aluminum "screen of power" and a novel controlled-skew-rate I/O design to produce nearly bulletproof power busing.

Poor power distribution in a sea-of-gates array is not well recognized. Poor distribution on the chip results in noise at the system level. Unfortunately, there are no tools to help predict these problems before testing. The chips have unpredictable yields and the boards have unpredictable yields. When the customer gets the part, it sometimes works in his system and it sometimes doesn't.

One reason for the paucity of analytical tools is that each application places different demands on the distribution system. For example, the number of simultaneously switching internal circuits and drivers will vary with circuit topology and data flow. Therefore certain assumptions must be made. When an ASIC manufacturer designs the master slice of the array, he designs it for general use, with no specific application in mind, yet he has to provide a technology that any customer can use or, perhaps, abuse.

Power distribution problems are often first encountered during testing. That is when one finds that it becomes necessary to overdrive the inputs to make the ASIC work, and overdriving calls for broadening the drive specifications. The first clue that noise is being injected is that, when the I/Os or internal circuits are switching, the switching-current change causes the ground to rise, and V_{CC} to collapse, reducing the voltage across the receiver, internal latches, and flip-flops.

Coping with Noise

Because the noise generated depends on the system-level environment, power distribution problems at the chip level are difficult to analyze. It is hard to determine, for example, how much inductance and decoupling the system will provide. One cannot know in advance if the part will be socketed or not, how well system power will be distributed, and how many drivers switch simultaneously.

Some designers are well aware of the problems of noise. Therefore they provide thick copper planes and well-decoupled power supplies. Others bring power in on a "yellow" wire and never think of decoupling capacitors. Thus even if there were an analytical tool, because of the vast differences in operating environments the array planner would never be completely sure what assumptions to make to use it. Further, the finished die may be housed in any of several packages, all with different inductances.

In sum, combining a highly inductive external environment with fast 1.2-μm technology is bound to create transient switching problems. With on-chip delays of 300 ps, off-chip delays of 5 ns or less, and a voltage swing of 5 V, $C \times dV/dt$ produces a large current swing that must be accounted for.

Power distribution problems are aggravated by faster switching speeds and greater component density. Further, larger gate counts make it more likely that more internal circuits and I/Os will switch simultaneously. Larger arrays typically house circuits with wider data, address, and control buses—that is, 4- and 8-bit buses are giving way to 16- and 32-bit buses—again making simultaneous switching more likely. In short, high-performance, high-density sea-of-gates arrays are far less forgiving of power distribution problems than were their 2-μm-and-up forbears.

Power distribution in a double-layer metal ASIC is best accomplished on the second level, where the metallization is thicker and the resistance is the minimum. One leading sea-of-gates array distributes the power on the first level of metal. Then it automatically routes power on the second level, using wiring areas remaining after signal routing is completed. Unfortunately, as the array density (and hence the signal interconnection density) increases, fewer routable areas are available for power distribution. This approach clearly takes power distribution in the wrong direction.

Guessing about Abuse

A better approach is to make some assumptions about how the user will abuse the technology and to plan a conservative power distribution architecture and methodology that take those assumptions into account. It then becomes a question of making the right assumptions.

The architecture of the ISB12000 arrays shows the use of

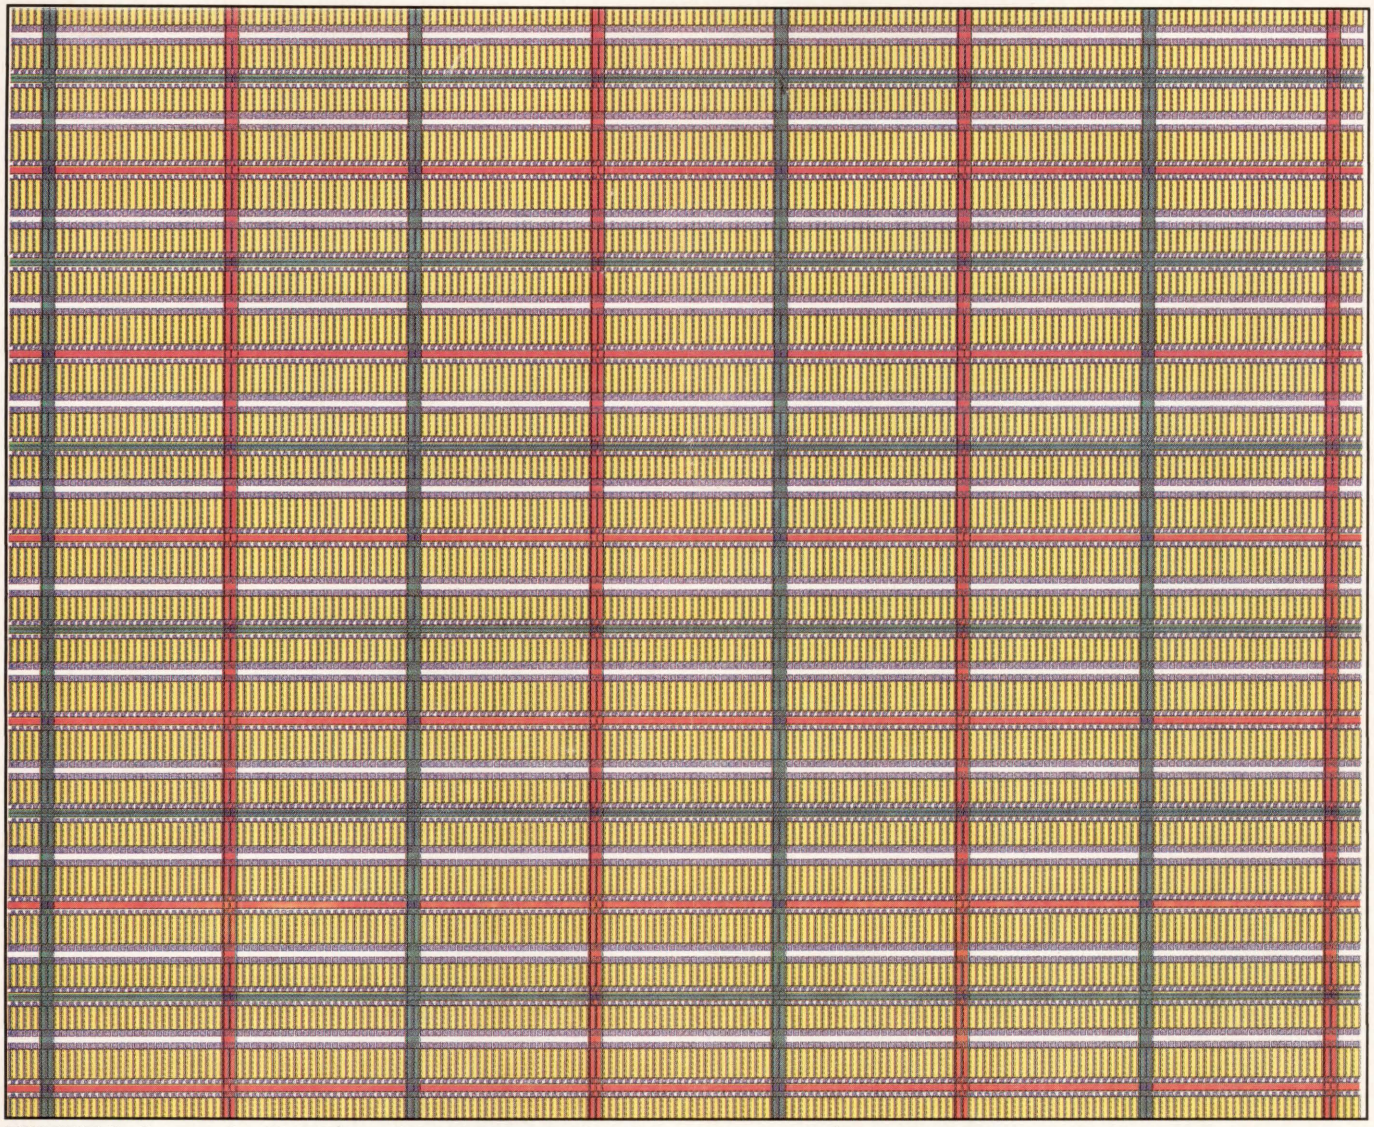

FIGURE 1. Power is distributed in this sea-of-gates array by a fine aluminum grid, then redistributed at every 25 columns of cells on metal layer two and between every row of cells on metal layer one.

reasonable assumptions. Power is distributed on a microlevel between cells by a fine aluminum power mesh. Figure 1 shows that power is distributed every 25 columns of cells on layer two and between every row of cells on layer one.

Instead of stacking cells, the cells are mirrored one on top of another. This arrangement allows for sharing power or ground buses between adjacent rows of cells, increasing the width of each bus to 4.6 μm and eliminating the need for the space between adjacent buses in stacked configurations. Wherever possible, this bus is widened further by the metallization associated with each macrocell.

But why is power distributed every 25 columns on level two? Once you decide to distribute power in fixed buses on level two (already a novelty), you must make some assumptions about the right periodicity of the power busing.

Three assumptions were made for the ISB12000. First, 40% of the raw gates will be used and a 10% safety margin left. Thus half the rows of gates are depopulated to allow for wiring interconnects. Second, the macro is assumed to be a 25-cell-wide inverter in a clock driver. Power is redistributed every 25 cells and a clock driver is used as a model because it gives the worst-case switching and transient currents. Third, the clock driver is assumed to switch rapidly, at 25 MHz.

A SPICE analysis showed that a single redistributed bus was sufficient to drive the clock driver while keeping the voltage drop and electromigration to acceptable levels. One 25-cell-wide inverter/clock driver equals the bus periodicity of one redistributed second-level 10.8-μm-wide bus.

All the power used by the internal cell matrix originates at the power and ground pads at the chip periphery. Though the current feeding each row or each twenty-fifth column may be manageable on a microlevel, the currents are additive between buses. Without additional bus augmentation, by the time the currents approach the chip edges they will be excessive, so an additional set of buses is needed.

If current is forced to flow through small vias and narrow lands, electromigration will cause an outflow of aluminum ions, which will hurt reliability over time. Augmenting the

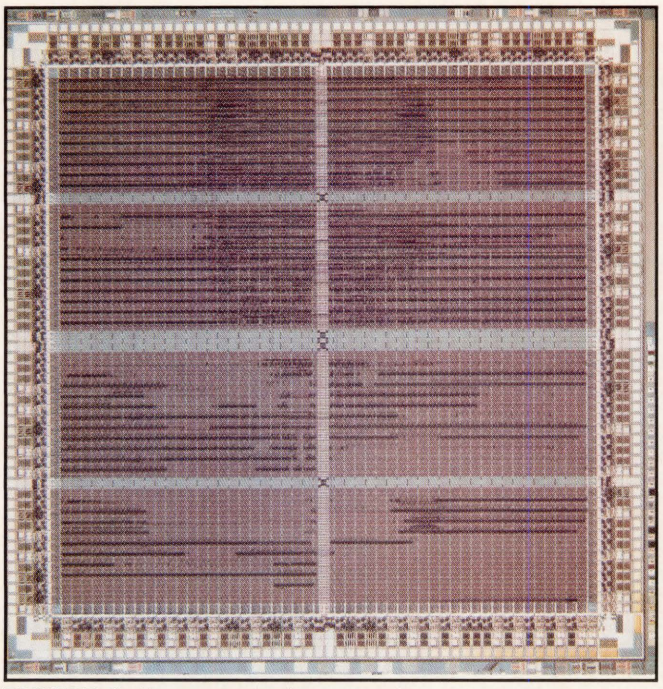

FIGURE 2. Augmented power buses appear here as three wide horizontal buses and one vertical bus.

power distribution buses helps prevents electromigration. A titanium barrier between the glass dielectric and the aluminum metallization provides further resistance to electromigration. It also prevents hillock formation between the silicon and the aluminum metallization.

To implement an augmented "coarser" mesh (one with larger pitch and wider conductors) requires the depopulation of both rows and columns. Typically, on the first level, two to four rows are depopulated every 20 to 40 rows to allow for a 150- to 300-μm-wide bus. On the second level, 50 to 150 columns are depopulated in the center of the master slice to allow for a 300- to 900-μm-wide bus. In Figure 2, 50 columns were depopulated in the center of the master slice.

This power distribution approach is particularly conservative. It minimizes design constraints, even for those who choose to operate the design out to the 3-σ operating points. The scheme also allows operation under military environments without degrading performance.

There is more to the system than just the fine and coarse power grids that overlay the array. There is the effect on power distribution to the I/O cells and the matter of supplying power to the output buffers without affecting the array matrix or off-chip receiver circuits.

Six Power Rings

In the entire power distribution system, there are actually six power rings that surround the array core (see Figure 3). All six are distributed on second-level metal. The two outer rings supply power to the p-channel and the n-channel output transistors of the I/O drivers and are completely independent of the other power rings. The next four rings supply power to the internal matrix, and the predriver and receiver areas of the I/Os. The large output transistors can turn on simultaneously and not affect the internal matrix and receivers. The fine

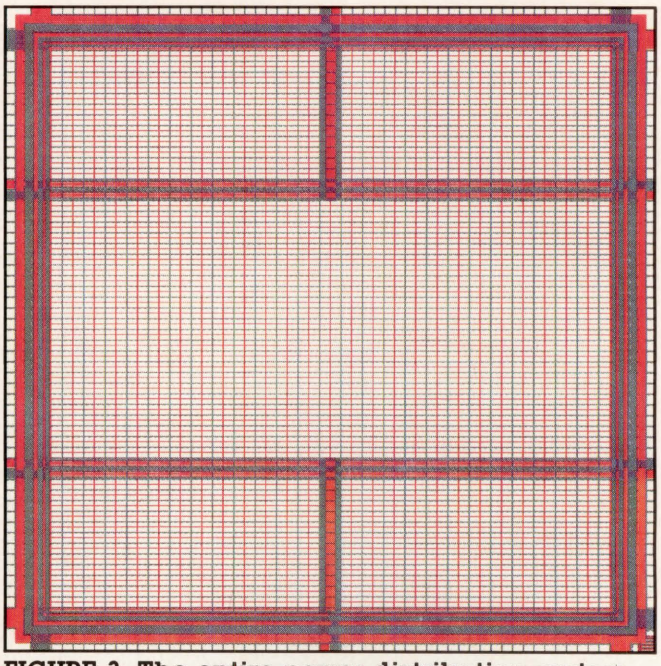

FIGURE 3. The entire power distribution system of the ISB12000 sea of gates.

aluminum grid and the augmented power buses tie into the inner four rings. The augmented buses are then tied directly to the internal power supply chip pads.

Before proceeding further, one should evaluate the real-estate penalty of an adequate power distribution system. Just how much silicon is sacrificed?

The best measure is the number of square mils per wirable gate. The ISB12000 achieves an average of 7 square mils per gate, an extremely attractive number using 1.2-μm design rules. Perhaps 6.5 square mils per gate might have been achieved with less attention to power distribution. But the additional area is a small price for predictable results.

A desire to automatically generate master slices of various sizes had a major effect on the design of every macro cell. In particular, alignment of the macros with the power bus structure became an overriding concern, since the location of the cell contacts relative to fixed power buses could present a serious placement and routing problem.

The need to align cells with power buses is obvious. Another need was to be able to generate 10 arrays with 8000 to 128,000 raw gates without requiring a unique power distribution architecture for each array. The fixed grid and the array's symmetrical topology made it possible to use the Calma GDS II system's Graphic Programming Language to generate each master slice by simply defining the number of I/O cells on each side. The program automatically stepped the I/Os, the internal matrix, and the power distribution. The advantages of symmetrical geometry go far beyond the ability to rapidly generate master slices. It also simplifies coding for the place-and-route chip images; it makes the power distribution uniform over all cells, and hence it makes design rule checks (DRCs) constant between master slices. There is nothing unique in any one master slice.

But what about cells that just happen to fall under a power bus? How do you connect to them if their connections are blocked? The second-level-metal power buses over the active

internal matrix areas are two cells wide. Every internal macrocell is at least three cells wide. To ensure that every cell's input and output are accessible, even when the cell falls directly under a power bus, ''antennas'' are provided for each connection, as shown in Figure 4. One antenna carries at least three ''hit'' points, so that at least one hit point always sticks out. In fact, with macros, such as D flip-flops, that tend to use all the *x*-direction wires, there may be seven or eight different hit points per antenna for the A input, B input, Q output, D input, Reset, and so on. No matter where the cell falls relative to a second-level-metal power bus, the wiring program will always be able to make connection with every pin. In addition, the low-resistance polysilicide links that form the transistor gates facilitate interconnecting the macrocell metallization without using any second-level-metal routing channels.

The wiring program assumes that one hit point on each antenna is always a ''protected'' pin—a keep-out area during routing until the pin is connected. If the protected pin happens to fall directly under a power bus, the placement-and-routing software (Silvar-Lisco's Advanced GARDS) reassigns the protected pin to one of the unblocked pins.

Coping with Instantaneous Power

The fine and coarse aluminum power meshes and inner power rings distribute power to the inner array core, but we must also attend to the instantaneous power demanded by the output transistors, based on their slew rate. The ISB12000 has optional 2-mA, 4-mA, and 8-mA drivers without paralleling I/O cells. The 600-μm-wide drivers are capable of hefty slew rates of better than 200 mA/ns. If all the bits in a 16-bit or a 32-bit bus are firing simultaneously, an intolerably large current swing will result. This current swing can cause a dangerous noise spike. The best solution is to provide a slew-rate control circuit in the I/O cell that limits the slew rate to approximately 20 mA/ns.

Using a slew-rate-controlled I/O also makes the off-chip delay almost independent of load. Whether you drive 20 pF or 100 pF, you get practically the same delay, because the slew-rate control determines the rate at which the capacitance is charged or discharged. For example, if the slew rate is constant at 20 mA/ns, the driver delay is 6 ns ±1 ns.

Besides slew-rate control in the I/O cells, the I/O pads have two clamp diodes electrically connected, one from the pad to V_{CC} and one from the pad to ground. These clamp diodes prevent reflections from open-ended transmission lines from doubling and causing false switching and sharper waveforms. The driver circuit also contains short-circuit protection that limits the current through the output buffers during a low-impedance fault.

Spare the Power and Spoil the Package

If attention is not paid to power distribution on chip, special attention will have to be paid at the package level to minimize package inductance, perhaps using ground planes in a custom-tooled package with fixed power and ground. As a direct benefit of careful attention to powering at the chip level, open-tooled packages can be purchased from off-the-shelf suppliers. One can easily satisfy an application requiring specific power pin assignments, and it is easy to provide additional power pins for bus-oriented designs with high-power output buffers.

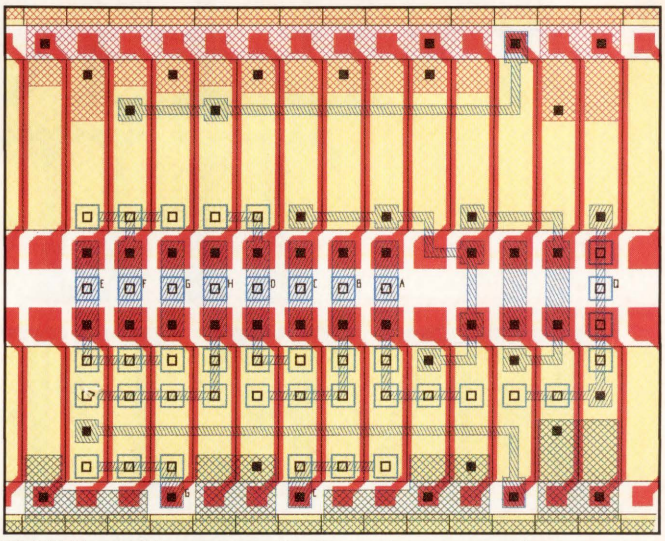

FIGURE 4. "Antennas" at each connection ensure accessibility to each cell's input and output even if the cell falls directly under a two-cell-wide power bus.

At about $15,000 in tooling costs per package, using custom-tooled packaging is expensive. With open-tooled packages, there is no packaging NRE and there is approximately a 50% saving in the per-piece price. Further, lead times for custom packages can extend to 26 weeks, whereas open-tooled parts may be available overnight.

Cutting Delays, Skewing

Good power distribution on chip also can reduce delay tolerances and clock skew, which can result from voltage distribution variations of as much as 0.5 V around the chip. Latch-up characteristics improve, too, for several reasons: First, the slew-rate control cuts noise. Second, one can tie down the substrate, the p wells, and the n wells so that the chip itself is less susceptible to noise-induced latch-up. In fact, one can push or pull more than 750 mA through any signal I/O without inducing latch-up. □

About the Authors

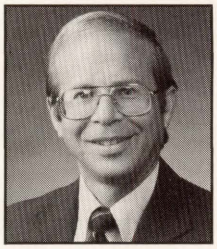

Richard Blumberg, director of advanced array development, joined SGS-Thomson Microelectronics (formerly Thomson Components–Mostek Corp.) at its inception in November 1985. He was previously with UTC Mostek, and before that with IBM, where he managed the design and development of semicustom bipolar products. He earned a BSEE from George Washington University in 1968.

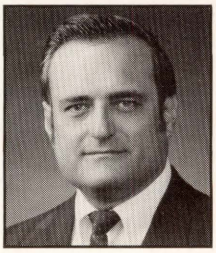

Charles Waggoner, senior design engineer of advanced array development at SGS-Thomson Microelectronics, joined predecessor UTC Mostek in June 1983. He has been involved in CMOS gate arrays and, from 1968 to 1983 at IBM, in the design and development of semicustom bipolar products.

ONE PROTOTYPE,
THREE PROT
FIVE PROTOTYPE
SEVEN PROTO

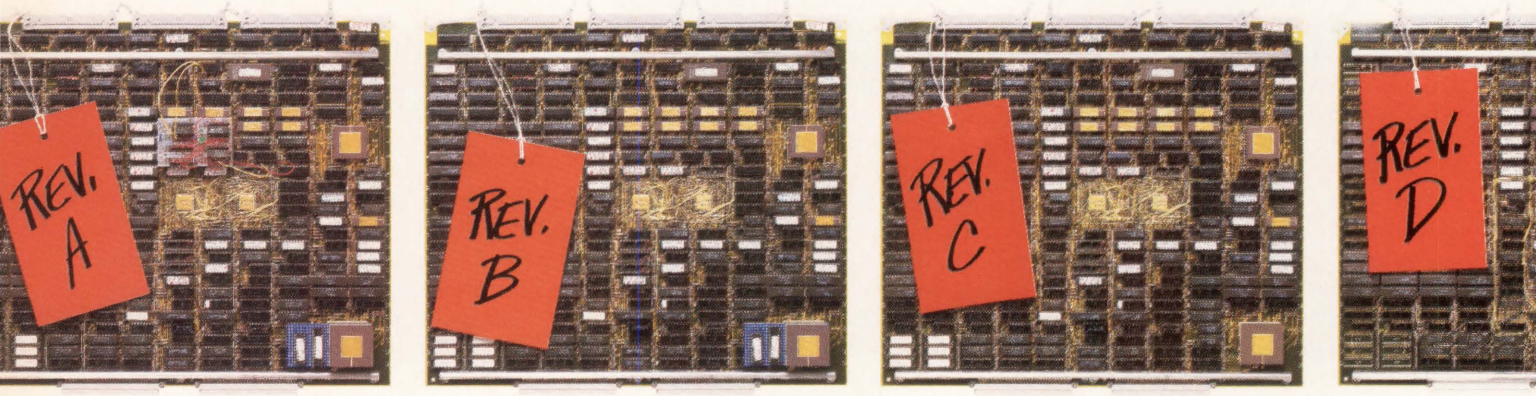

Teradyne's new CAE tools help you get to market faster by

Designing VLSI circuit boards that work the first time isn't child's play. You can't do it without fast, accurate feedback on design and test quality.

Now count on Teradyne to help. With our new family of CAE products, you'll uncover flaws before they're cast in hardware. So you'll be able to jump from initial design directly to final prototype. And from design to manufacturing test.

Start with a firm design foundation.

The process starts with DATAView our new design entry system. Running it on a standard AT-class PC, you can control the whole design process right at your desk. From schematic entry and waveform analysis on the PC to simulation and fault simulation on more powerful networked computers. All with the same mouse-and-menus, multiple-window interface.

You move from schematic capture to simulation effortlessly. DATAView's links with our LASAR Version 6 simulation system save hours of model compilation time. Incremental compiling lets you revise a design and resimulate in minutes.

Make your design work together before it comes together.

LASAR Version 6 is the only sure way to avoid multiple trips to artwork as well as costly rework steps in manufacturing. Because its simulation accuracy is unmatched

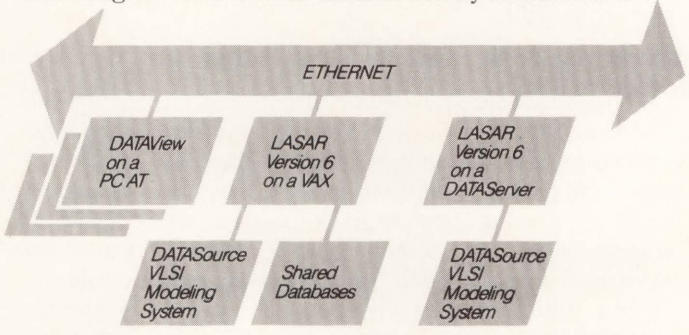

Teradyne networks advanced CAE tools for convenience *and* performance.

TWO PROTOTYPE,
OTYPE FOUR.
' SIX PROTOTYPE,
,TYPE, MORE.

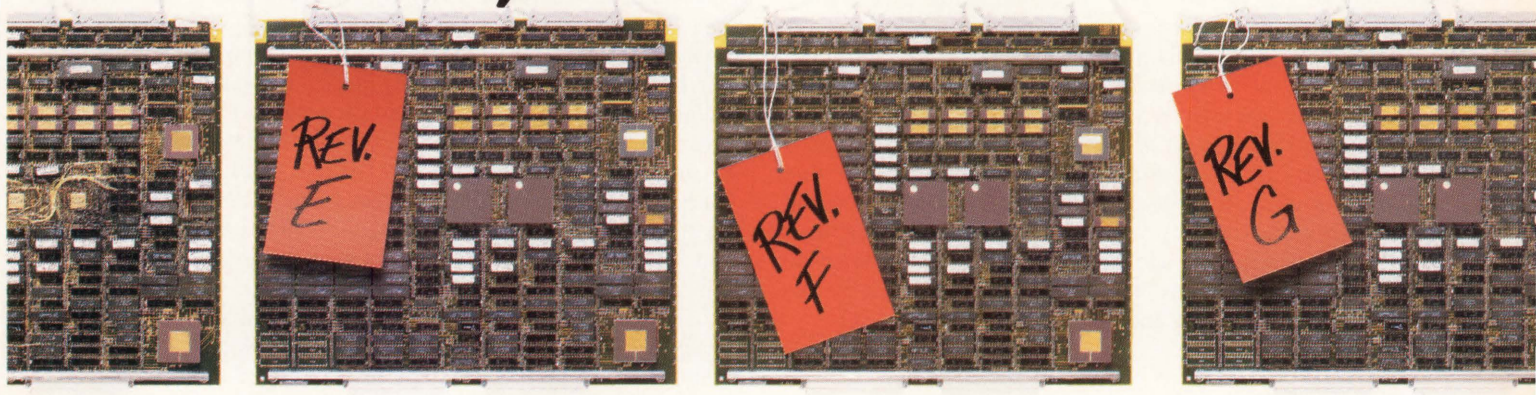

giving you a working prototype the first time.

for analyzing PC board designs prior to prototyping. For instance, LASAR takes full device timing specifications into account for true worst-case timing analysis. And it eliminates shared timing ambiguity in reconverging signals. Both of which mean LASAR finds real design errors reliably.

Finally, you can use LASAR fault simulation. It will uncover testability problems and untested circuit functions before it's too late.

Simulation without trial and tribulation.

LASAR also takes care of the board modeling problem. It includes models of over 4000 devices. And the most popular gate array libraries.

For new VLSI parts, our DATASource hardware modeling system uses actual devices instead of software models. A single system supports multiple users and concurrent fault simulation. With extraordinarily fast response.

LASAR and DATASource work with any VAX. For simulations up to 25 times faster, you can use our new parallel/multiprocessing host, DATAServer. It will give you quick results, with full LASAR precision.

Teradyne makes it easy.

If you want to avoid multiple prototypes and get to market faster with better products, Teradyne's new CAE system is the way. Why not call Daryl Layzer at (617) 482-2700, ext. 2808 for more information. Or write Teradyne, 321 Harrison Avenue, Boston, MA 02118.

VAX is a trademark of Digital Equipment Corp.
AT is a trademark of International Business Machines Corp.

TERADYNE

We measure quality.

CIRCLE NUMBER 6

The Integration of Graphics and Imaging Systems

David Smith, Western Regional Editor

Workstations are now the predominant form of computing power for technical and business professionals. They run application programs previously found only on supermincomputers and so put more powerful software in the hands of more people.

These applications require greater display quality and interactivity than the products previously running on the desktop. More information is presented at once, such as three-dimensional images, graphs, charts, and colors. Each of the new workstations has to have the image-rendering power once found only on terminals tied to a supermini.

These two reasons explain why all computers are becoming commodity items. Why put 10 MIPS into a workstation when you can't view the results in a meaningful, productive display? In addition, the cost of all other computer subsystems is coming down fast. A 1.5-MIPS workstation costs $5000; five years ago it cost 10 times as much. Graphics and imaging systems are under the same pressures to deliver higher functionality at a lower price. However, the use of VLSI devices to reduce cost and improve throughput is hampered by the peculiarities of graphics and imaging architectures.

Image Computing

The terms *graphics* and *imaging* are sometimes used interchangeably. A convenient classification comes from Pixar Corp. (San Rafael, CA), which lumps both terms under *image computing* (see Figure 1). In Pixar's scheme, under computing graphics, two-dimensional systems are used for graphics arts, including CAD and publishing. *Image synthesis,* the 3D leg of computer graphics, includes the modeling of 3D images from mathematical models.

Under *image processing,* 2D image interpretation includes map making. Volume imaging, the 3D portion of image processing, includes medical imaging, in which 2D sections of a body are constructed into 3D representations. In this view, therefore, image processing differs from graphics in that the basic data is an image built of pixels, whereas graphics systems receive data as models. As shown in Figure 2b, the memory-processor interface is more closely linked in imaging systems. Figure 2a, the graphics pipeline, shows how the memory of graphics models, the display list, is just the beginning of a graphics pipeline. It feeds a succession of processing steps, many of which can be implemented as VLSI devices.

VLSI devices for the display end of the graphics pipeline are available because the drawing of pixel data on a given-resolution screen is a well-defined process. For example, memory devices that translate a pixel description into a color—*color lookup tables*—are combined with digital-to-analog converters to form complex analog-digital output devices.

VLSI devices also are used to control access to the memory that contains the pixel data—the *frame buffer*—from both the host CPU and the video output circuits. Such *display controllers*—the Hitachi 6845, for example, used in the IBM PC XT—can often perform rudimentary logical operations on the pixel data.

More advanced VLSI display controllers draw and manipulate pixel data, such as drawing lines, and may also fill in patterns and move images around the screen. These *graphics processors* are available from Advanced Micro Devices Inc. (Sunnyvale, CA), Hitachi Ltd. (Tokyo), Intel Corp. (Santa Clara, CA), National Semiconductor Corp. (Santa Clara), and Texas Instruments Inc. (Houston, TX). They perform integer and logical calculations while drawing the two-dimensional images.

To create the data that describes the 2D image, *geometry processors* accept mathematical models of objects and commands to alter the position and orientation of those objects. They calculate the position of the objects and the effects of lighting and clipping. Then they pass 2D coordinate and color information to a display controller for rendering in the frame buffer. Graphics processors must perform a variety of floating-point algorithms, such as 4×4 transformations. They are implemented, therefore, with general-purpose floating-point processors, like the 3332 from Weitek Corp. (Sunnyvale), and controlled by microcode that implements the designer's specific algorithms.

The retrieval of mathematical models and graphics commands is managed by the *display list processor.* The graphics system's host CPU, running the application software, usually constructs a list of objects called the display list. The display list processor receives the commands to be executed, retrieves the objects that the commands are to be executed on, and schedules execution of the rest of the pipeline. These processors are usually built from general-purpose chips like the 68000.

The graphics pipeline is sequential: The display list processor fetches models and commands, the geometry processor calculates model coordinates and attributes, the display controller constructs an image in the display buffer, and the video circuits display the frame buffer's contents on a screen. Image

processing, although often considered inherent in graphics, is largely more cyclical than sequential.

The raw data for image processing is usually in pixel form, a large collection of 3D coordinates and attributes, instead of mathematical models. Image-processing systems (see Figure 2b again) therefore fetch image data, perform operations on them such as filtering and convolution, and place them back in image memory. Although, like graphics systems, image-processing engines may move and scale object data, their more important function is the extraction of information from the image, rather than the extraction of an image from information (models).

Although the data flow and types of operations are different in image processing and graphics, the factors influencing the architecture of the systems are similar: The stability and portability of standard operations and data types influence the use of programmable or dedicated devices as processing engines. The abundance of resolution requirements, both within algorithms and in displaying data, create different datapath widths for different systems and for different parts of the same system. The need to tailor processing power and functionality affects the integration of circuits in either complete processors or building blocks. Finally, throughput and expandability requirements can force designers into creating parallel or highly pipelined architectures.

Programmable vs Hardwired

It has been difficult for manufacturers of standard VLSI components to create products for graphics and imaging systems because of the variety of formats and throughputs of these systems. For example, PC-based graphics systems may have to respond to commands in Extended Graphics Adapter (EGA), Hercules monochrome, Video Graphics Array (VGA), and Computer Graphics Interface (CGI) formats.

Higher-end workstations, if they support any standard software at all, can conform to Programmers' Hierarchical Interactive Graphics Standard (PHIGS) or Graphical Kernel Standard (GKS) commands for 2D graphics and PHIGS+ or GKS-3D for 3D applications. "If PHIGS or PHIGS+ becomes widely accepted," postulates John Dalrymple, a principal scientist with Tektronix Inc.'s Graphic Workstation Division (Wilsonville, OR), "then there exists some opportunity for semiconductor manufacturers" to make standard VLSI devices for systems using those standards. Data and command formats for the most advanced graphics and imaging systems, however, can change with each successive model, as manufacturers try to expand the flexibility and throughput of their systems.

For these reasons, VLSI graphics processors implement a small set of primitive functions that are copied in many graphics systems. Texas Instruments' TMS 34010, perceived as the most programmable such chip, is a general-purpose microprocessor with graphics functions added—specifically block pixel moves and line drawing. Because it is a general-purpose part, its performance is somewhat lower than that of other graphics processors, drawing 700,000 random vectors/s versus 2.5 million vectors/s for the more specialized Intel 82786.

More-specific graphics ICs, like Intel's, AMD's QPDM, and Hitachi's 63484, have instructions that implement more graphics functions, such as circle and arc drawing. If a graphics operation maps directly onto these instruction sets, these chips are faster at those operations than the more general-purpose, programmable parts.

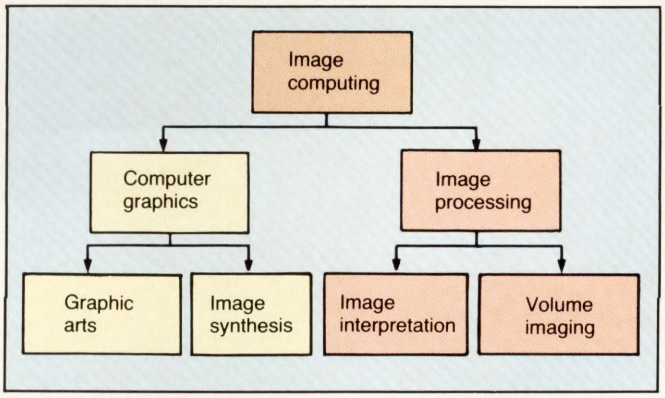

FIGURE 1. Image computing encompasses computing graphics and imaging processing, as conceived by Pixar Corp.

The Intel 82786, for example, implements windowing functions in hardware. Pointer registers within the device enable it to keep track of the different areas of memory that contain the contents of windows displayed on the screen. If the application supports these pointers (as does Autocad, for example), then hardware windowing is more efficient than using block transfers of window data to build a screen image.

However, the more advanced a system is, the more likely that it contains specialized graphics functions such as shading. Such functions do not map directly into the instruction sets of the specialized ICs, therefore requiring the host CPU to build the image directly. Moreover, many such functions are the ones that make a product useful and unique. "Another problem with standard VLSI graphics chips," adds Gene Chao, president of Metheus Corp. (Hillsboro, OR), "is no product differentiation from your competitor."

Even graphics system manufacturers have difficulty defining the exact instructions to implement in VLSI devices. For example, Metheus set out to implement its graphics pipeline, previously built with bit-slice components, in gate arrays for its UGA graphics card. Because of the need to remain compatible with earlier products while implementing new features, the company designed a general-purpose processor, in line with the philosophy of the TI part. The processor executes routines from microcode, performing calculations corresponding to the geometry processor in the graphics pipeline and handing off commands and vector coordinates to one or more custom datapath chips.

Similarly, in creating the new 4330 series of terminals, designers at Tektronix had to maintain compatibility with the older 4230 series. They used a general-purpose processor (the 59032 bit-slice processor from Waferscale Integration Inc., Fremont, CA) to processes the set of 4230-compatible graphics commands. They then used several gate arrays to implement specific functions that operate under control of the processor.

The algorithms used in imaging systems are as complex and mutable as those graphics. Most pioneering work in imaging therefore occurs on general-purpose mainframe and supercomputers. The wide range of imaging applications, including medical imaging, animation, and geographic survey, prevents the integration of much of the mathematics.

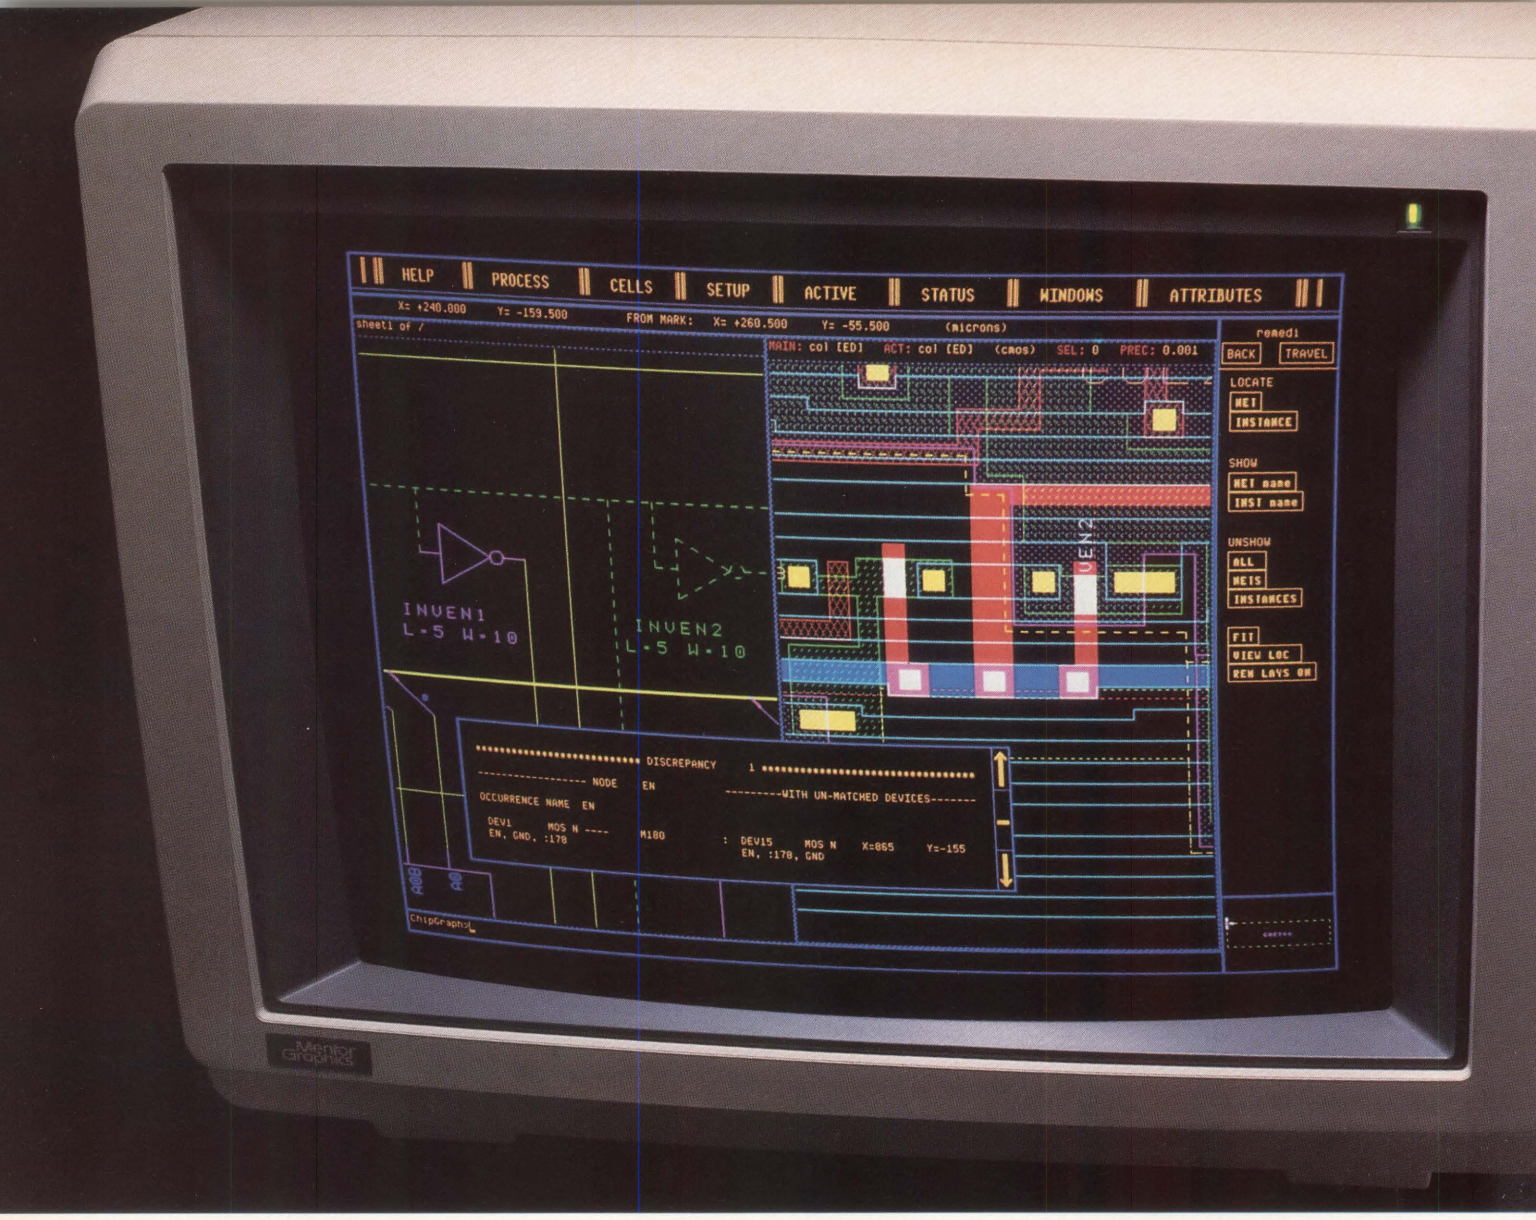

Only Mentor Graphics maps symbols to silicon.

The bigger the IC design, the bigger the problem: you're trying to locate a discrepancy between the schematic and your IC layout, but all you have is an ASCII error report. It's like driving all over a strange city to find an address — without a map.

That's why Mentor Graphics created REMEDI,™ a graphical interface that helps debug complex layouts. REMEDI works with Dracula II,™ taking the leading layout verification package's layout-versus-schematic checks a step further. LVS errors detected with Dracula II can be quickly pinpointed on both the layout and schematic using REMEDI's interactive graphical correlation capabilities.

And, because REMEDI is part of ChipGraph,™ the powerful Mentor Graphics full custom IC layout editor, as soon as you find an error you can fix it. There's no need to move back and forth between the layout editor and the debugging tool to correct the design database.

Today's complex VLSI designs demand a layout tool that lets you create and navigate efficiently through a maze of mask data. So ChipGraph provides flexible geometry editing and fast cell-based layout tools. Beyond this,

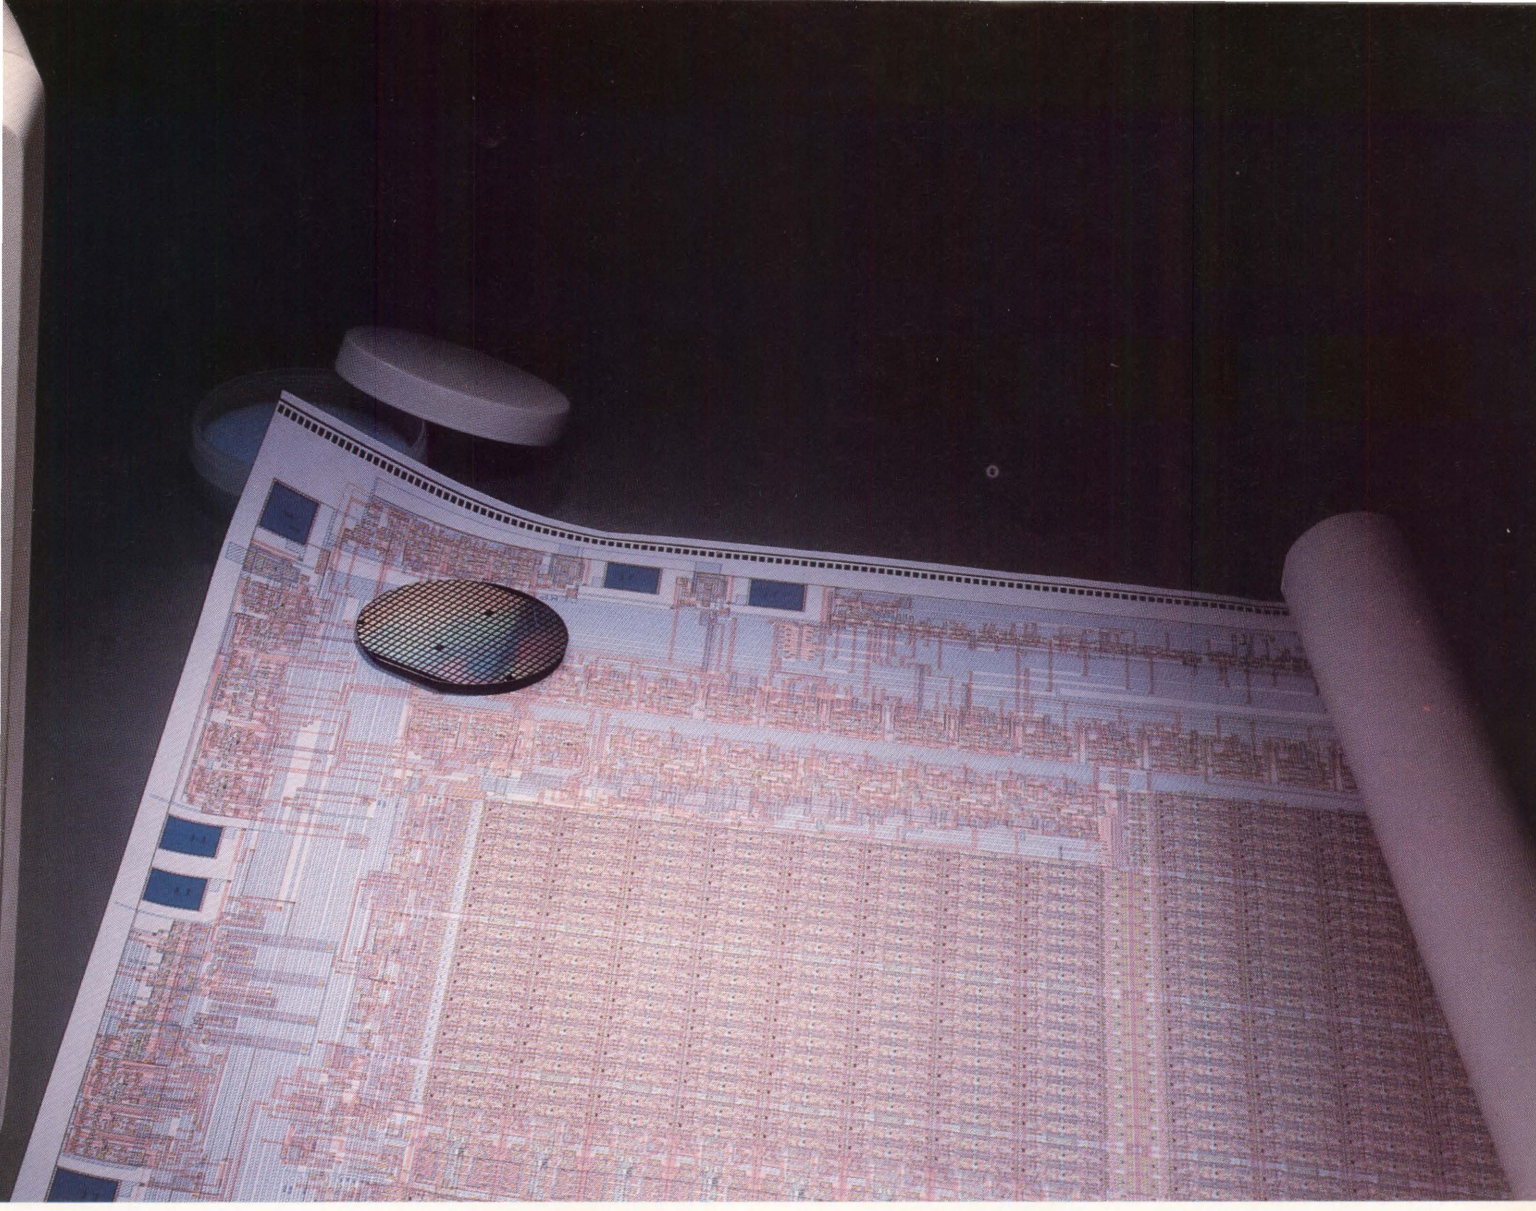

Structured Chip Design (SCD), a hierarchical approach to physical layout, removes much of the unnecessary and confusing mask data, while retaining correct functionality and interfaces.

With SCD, you can work with a simplified representation of the cell when making higher level decisions. And you can easily move between SCD and more familiar tools, with no loss of productivity. The result? The tightest possible layout, created quickly and accurately.

ChipGraph also lets you partition a large design over any number of workstations while maintaining version control, through a network-wide shared database.

And there's no need to worry about losing old data when switching to a new tool. ChipGraph offers full data compatibility with your existing design files.

Best of all, ChipGraph is not just an idea — it's a working reality. Designers are already using ChipGraph to lay out 32-bit microprocessors and multi-megabit RAMS.

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

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

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

Sydney, Australia; Phone 02-959-5488 **Mississauga, Ontario;** Phone 416-279-9060 **Napean, Ontario;** Phone 613-828-7527 **Paris, France;** Phone 0145-60-5151 **Munich, West Germany;** Phone 089-57096-0 **Wiesbaden, West Germany;** Phone 06121-371021 **Hong Kong;** Phone 0566-5113 **Givatayim 53583, Israel;** Phone 03-777-719 **Milan, Italy;** Phone 02-824-4161 **Asia-Pacific Headquarters, Tokyo, Japan;** Phone 03-505-4800 **Tokyo, Japan;** Phone 03-589-2820 **Osaka, Japan;** Phone 06-308-3731 **Seoul, Korea;** Phone 02-548-6333 **Spanga, Sweden;** Phone 08-750-5540 **Zurich, Switzerland;** Phone 01-302-64-00 **Taipei, Taiwan;** Phone 02-7762032 or 02-7762033 **Halfweg, Netherlands;** Phone 02907-7115 **Singapore;** Phone 0779-1111 **Bracknell, England;** Phone 0344-482828 **Livingston, Scotland;** Phone 0506-41222 **Middle East, Far East, Asia, South America;** Phone 503-626-7000

Mentor Graphics®

Dracula II is a trademark of ECAD, Inc.

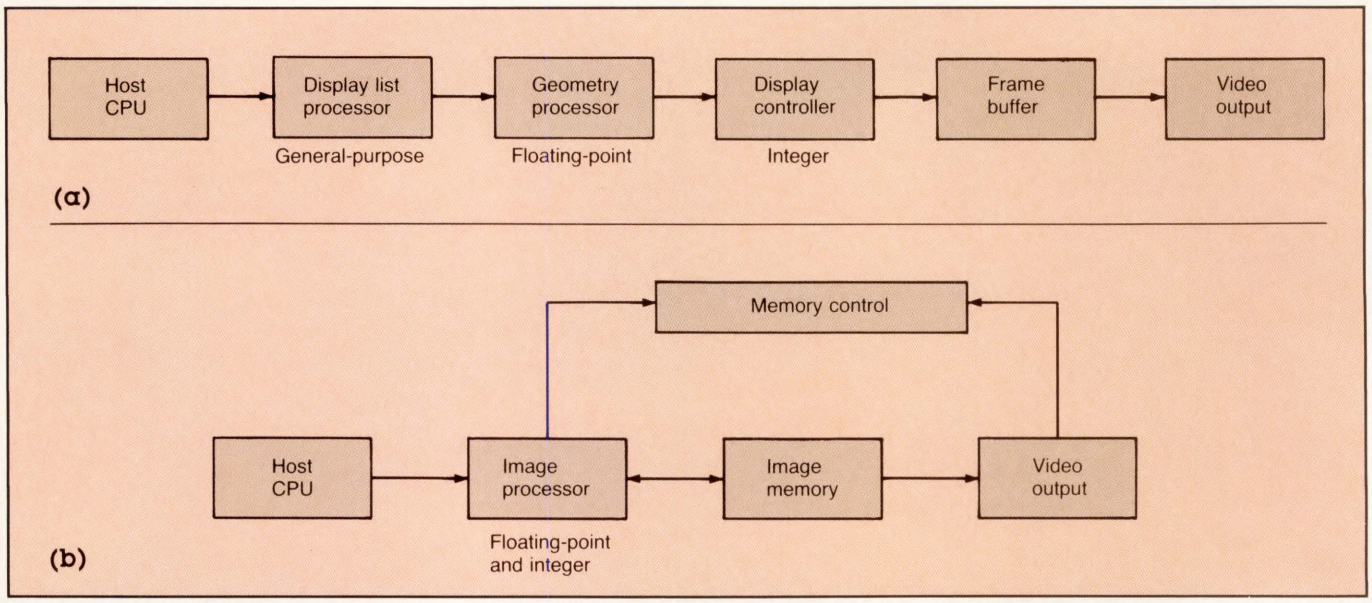

FIGURE 2. The sequential graphics pipeline (a) and a cyclical imaging system (b).

Consider the architecture of one high-end imaging machine, the channel processor, or CHAP (see Figure 3), within Pixar's Image Computer. Says the company's director of product engineering, Jim Wilson, "For a complete solution, you need a computer. The CHAP is more like a minicomputer" than a graphics engine. He estimates that the CHAP typically performs 200 times faster than a VAX-11/780. He points out that for such systems, designers typically trade off integration for performance.

All algorithms for the CHAP reside in a microcoded control circuit (not shown) that directs the execution of four sets of general-purpose ALUs and multipliers. Two 48-bit vectors buses carry four 16-bit words of pixel data (red, green, blue, and a unit called alpha) each memory cycle. One ALU-multiplier pair is available for each word. The crossbar switches betweeen the scratchpad memory can direct any of the words to any of the four ALU-multiplier pairs. In addition, a routing network can extract a single 16-bit piece from the vector bus to the Sbus (scalar bus), where scalar arithmetic can be performed.

The hardware can execute a great variety of graphics and imaging algorithms, processing any combinations of the data in the local scratchpad memory. The use of four independent arithmetic circuits allows all four pixel elements to be processed simultaneously. In addition, other hardware, including other CHAPs, can be attached to the YAPbus (*y*et *a*nother *P*ixar bus) for further parallel processing. As many as three CHAPs can exist in a single system, sharing access to the image memory over the Pbus (Pixar bus). Wilson says that a fairly linear speedup is attained as CHAPs are added.

The use of general-purpose parts lets the designer run a variety of algorithms, even switching them on the fly. For example, when moving an image on a screen, the Image Computer can reduce the sampling bandwidth of the hardware on the image, allowing the computer to keep up with the dislocation of the image about the screen. Real-time response is an important, and difficult, feature for designers of imaging systems to implement. The demands of creating realistic images "is an infinite sink on computer power," according to Adam Levinthal, one of the architects of Pixar's Image Computer.

Data Format and Resolution

The unusual pixel organization of the Pixar machine highlights another aspect of graphics and imaging systems: the variation in the format of the image data. The organization of the frame buffer itself can be either sequential (planar) or parallel (packed-pixel). The frame buffer may have from 1 to 96 memory planes, each containing a bit for each pixel. As data are processed, they may change from floating-point to integer format to design-specific formats, with intermediate results at higher resolution along the way.

For a planar frame buffer, the bits in each plane are stored in the same area in memory. A memory access for 16 bits, for example, accesses 1 bit for each of 16 pixels in the plane. In packed-pixel memories, all of the bits for a pixel are stored and accessed linearly, so that a 16-bit access may produce a two 8-bit pixels. Since drawing a line may change all of the bits in the line's pixels, packed-pixel memories are better for line-drawing commands. Operations such as masking may affect only a few planes of the frame buffer and are therefore more efficiently carried out in planar memories.

Advances in memory devices may help to combine the attributes of the two types of organization. NEC Corp. (Tokyo) has introduced a triple-port DRAM (the μPD42232) that can support both types of access modes, even (with an appropriate supporting controller) switching between the two on the fly. The chip contains two RAM ports similar to those found on dual-port DRAMs. A third, serial port, the pixel access port, can be used for accessing a bit corresponding to an individual pixel. By enabling eight μPD42232s that are connected in parallel, a system can perform a packed-pixel access through the chip's pixel access port. Planar accesses proceed through the conventional parallel port.

National Semiconductor's Advanced Graphics Chip Set, based on a programmable approach like that taken by the TMS 34010, are the only graphics chips that could operate in both a planar and a packed-pixel memory configuration. Optimized

The Gate Array WorkSystem™ Makes Layout As Easy As Pushing A Button.

4 IN A SERIES

Tektronix Aided Engineering

Now logic designers can control the entire physical layout of gate array designs. From a single schematic entry environment. Just by pushing a button.

Developed by Tektronix as part of Tektronix Aided Engineering, the Gate Array WorkSystem eliminates the need for IC layout expertise. Because it gives you everything you need to quickly develop ASIC vendor-certified layouts.

And since you're controlling the layout from the schematic, you can tune your design using iterations of simulation and automatic layout to achieve your performance requirements.

The Gate Array WorkSystem creates a unique, performance-driven design environment integrating Designer's Database Schematic Capture (DDSC™) and industry-standard HILO®-3 logic simulation with MERLYN-G™ automatic physical layout.

The system also introduces vendor-certified TurnChip® ASIC Layout Modules for knowledge-based, automatic control of MERLYN-G layout of specific array families.

Layout so automated you can place and route a 5000-gate array 100% automatically. Just by pushing a button. With results so accurate that your layout is ASIC vendor-endorsed.

Using TurnChip modules, you can generate ASIC vendor-certified layout designs, then send them directly to the ASIC vendor. Which cuts your design time, lowers your costs and delivers complete control of your sensitive design data.

It's all part of Tektronix Aided Engineering. Integrated WorkSystems that take you beyond traditional CAE solutions. And into prototype verification, software development and testing, systems integration, mechanical design and manufacturing. All running on industry-standard platforms from Apollo® and DEC.™

Best of all, it's from Tektronix. The name you've always trusted to get the engineering job done. So you're assured of worldwide service, support and training.

If you'd like to take control of physical layout, contact your local Tektronix, CAE Systems Division, sales office. Or call 800/547-1512. Tektronix, CAE Systems Division, P.O. Box 4600, Beaverton, OR 97076-4600.

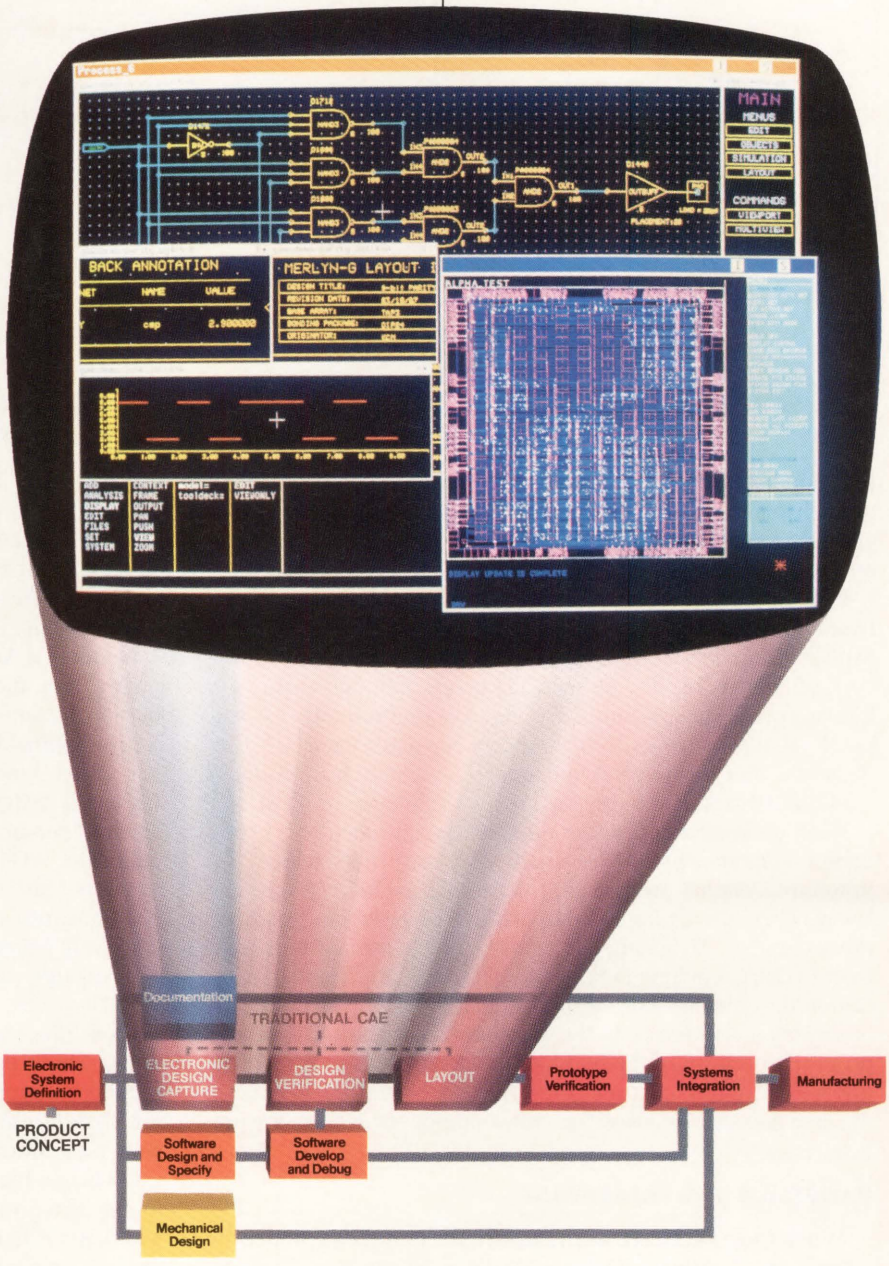

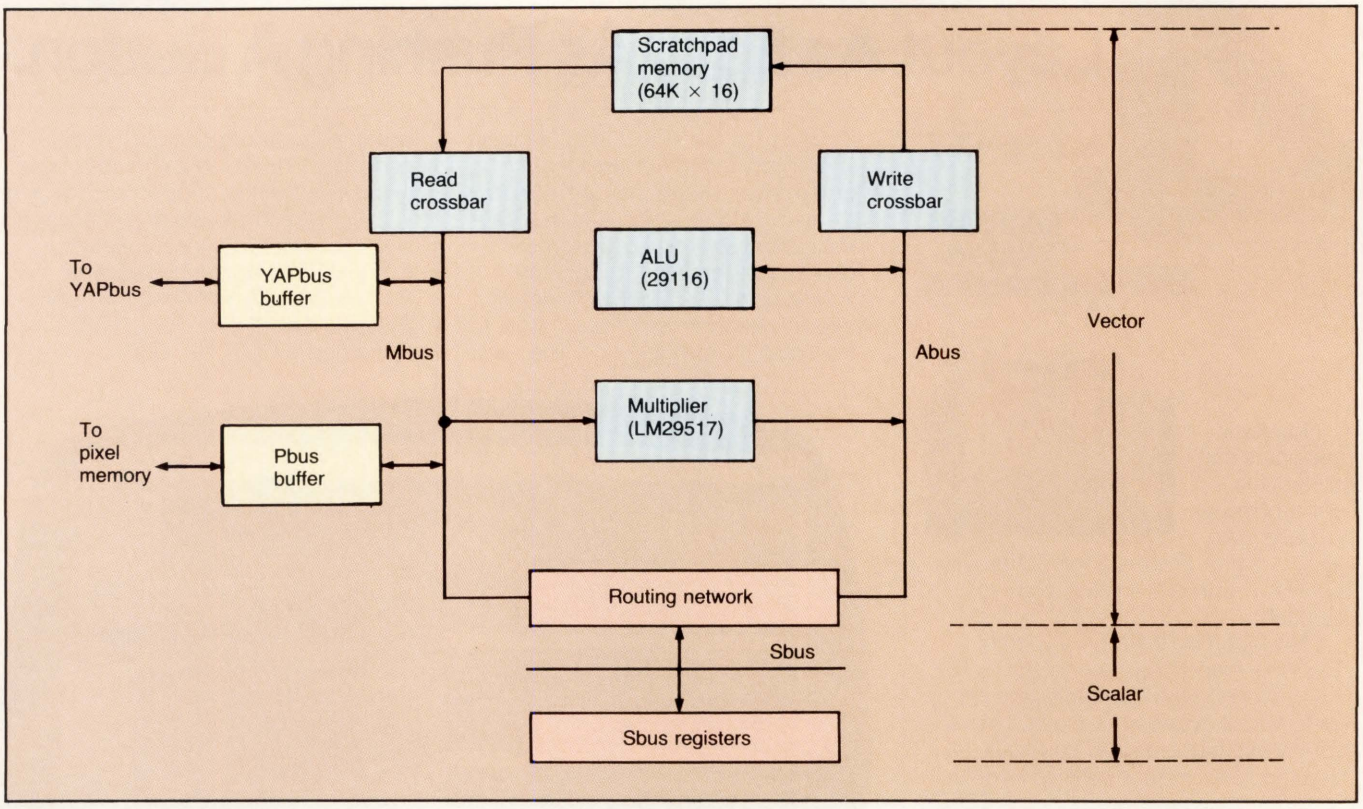

FIGURE 3. The CHAP, or channel processor, in Pixar's Image Computer.

for planar operation, it assigns a BITBLT processing unit (BPU) to each plane of the system. For packed-pixel accesses, the graphics processor would direct the BPUs to access a single bit, with all BPUs acting in parallel.

Before they are stored, image data may metamorphose between formats. Within the raster subsystem of the Iris GT from Silicon Graphics Computer Systems Inc. (Mountain View, CA), 96 bits are used to define each pixel, although a mere 74 bits per pixel are spit out to the video output circuitry. Also, pixel calculations use 27-bit integers for pixel values. "The datapath was designed to do exactly what it has to do," explains Bharat Patel, the company's manager of VLSI.

With datapaths dependent on specific architectures, it becomes clearer why off-the-shelf VLSI may not suffice. In graphics systems from Sun Microsystems Inc. (Mountain View, CA), for example, the shading of polygons requires processing of 32-bit integer values, which expand to 64-bit intermediate products to maintain accuracy. Jeff Buchanon, a senior designer at Sun, explains that integer processing is becoming a bottleneck in the company's processing systems, primarily because of the width of the datapath. Sun is moving toward designing custom chips for this bottleneck, because the gate arrays it has used for some component designs are not fast enough.

Partitioning vs Integration

When functions have been identified for implementation in silicon, the designers faces the decision of partitioning those functions. Some IC vendors and graphics companies have put all the basic functions on a single chip; others have taken a "Buliding-block" approach, putting logically distinct functions on separate chips. Implementing an entire pipeline on one chip can increase throughput and reduce system cost, provided that the chip's architecture matches the system designer's needs. By implementing sections of the graphics pipeline in different chips, on the other hand, the designer can use any or all of the chips as needed, adding extra circuitry between chips as required.

National's Advanced Graphics Chip Set (AGCS) is the best example of the building-block approach. It comprises the DP8500 Raster Graphics Processor (which, like the TMS 34010, is programmable), the BITBLT Processing Unit, clock generator chips, video shift registers, and video RAM controllers. Although designed to work seamlessly together, the parts are generic enough to fit in a variety of systems. The new Tektronix 4330 series, for example, uses the video shift registers, but no other AGCS devices, to translate parallel pixel bits into a serial stream for the system's d/a converters.

Certain operations within graphics and imaging systems are well enough defined to allow implementation in silicon. In the 4330 series, Tektronix used gate arrays instead of general-purpose processors to implement some pixel-processing functions. These operations were well enough defined to allow such as solution to be possible. System designers are designing custom chips for point problems in imaging systems as well (see Shugard, 1987, for example).

LSI Logic Inc. (Milpitas, CA) has a range of chips dedicated to specific problems in the processing of images. Its L64240 Multi-Bit Filter can be configured as a 2D filter for image-processing applications such as edge enhancement and detection. It accepts data directly from the company's L64210/11 video shift registers, which reformat raster-scanned video into a 2D data stream. The L64220 Rank-Value Filter can be used as a median filter, producing the median value from a sampling

of 2D data. For pattern matching and noise removal, the L64230 works on 2D images as large as 32×32 pixels.

The architecture of such point solutions, however, can still limit their use. The Vector Signal Processors from Zoran Corp. (Santa Clara, CA), for example, can impement some imaging functions (Eidson, 1987), but they are "too rigidly defined" for the breadth of image processing tasks, according to Pixar's Wilson. To begin with, the processor is built to process vectors of data, an approach unsuitable to the organization of Pixar's image memory. More applicable are general-purpose floating-point arithmetic chips that can be adapted to a greater variety of algorithms. "The Weitek 3332 is not a graphics chip. It's just a well-thought-out floating-point processor," explains one user of the part, Kurt Akeley, principal scientist at Silicon Graphics.

Parallel and Pipelined Circuits

"The challenge of the designer is to balance the throughput through the graphics pipeline," said Sun and Cates in discussing high-performance graphics (1987). For those bottlenecks that are too malleable for implementation in silicon, the use of parallelism and pipelining can improve performance while allowing the hardware to remain programmable (Budge, 1988). Parallel processing can be traced from the chips through the subsystems and even to entire systems.

The tradeoff for Metheus's programmable graphics processor was size, not speed, says Chao. The chip is relatively large, as a reult of the internal parallelism controlled by a 64-bit op code. Within Intel's graphics coprocessor, the four main blocks—a graphics processor, a display processor, a bus interface unit, and a memory controller—all operate in parallel. Within AMD's QPDM, four datapath ALUs operate independently, each on its own bit plane.

At the subsystem level, the use of generic processors in such systems is not inconsistent with high performance, if enough capability exists to run systems in parallel. Simple architectures can be fast if their functional and memory units are closely coupled, if the functional units are general-purpose, and if there are no constraints on the ability to program the functional units in parallel. In fact, for complex graphics problems such as ray tracing (where reflection off objects in the scene is factored into the color of each pixel of an object), the ideal system may have one general-purpose processor for each pixel.

The architecture of Silicon Graphics' Iris machine (see Figure 4) demonstrates the use of parallel processing and pipelining to increase the throughput of designs. Its geometry subsystem consists of five autonomous floating-point engines, based on Weitek's 3332 floating-point unit, connected in series. A graphics object arriving on the VMEbus (as a group of 3D coordinates and attributes) is subjected to, in order, coordinate translation, lighting calculations, object clipping to fit into the viewer's perspective, perspective generation, and conversion into screen coordinates. Each stage operates independently of the others, as does the whole geometry engine from subsequent circuitry.

The scan conversion, raster, and display subsystems demonstrate parallelism. Within the scan conversion subsystem, a polygon processor and a slope calculator break down the polygons into trapezoids and assign slopes to the edges of the trapezoids. The seven edge processors use this vertex and

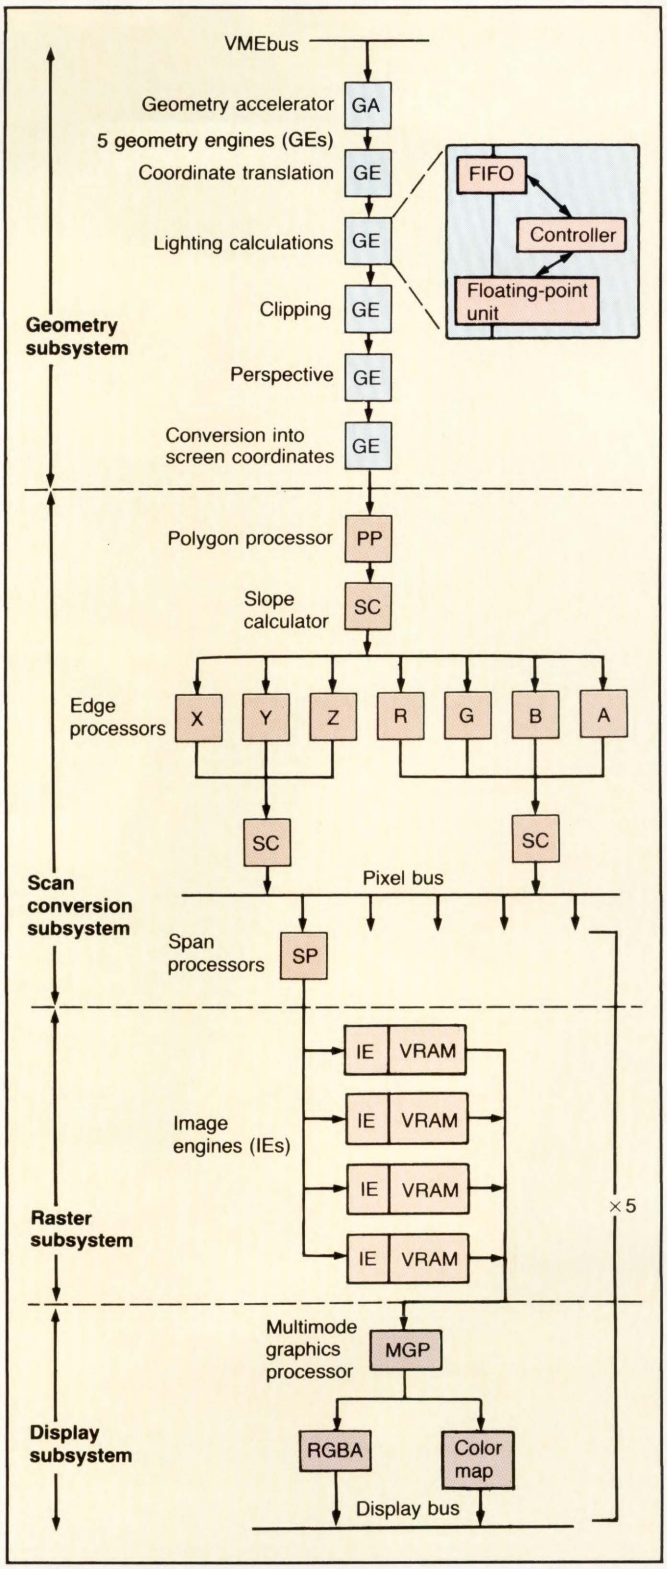

FIGURE 4. Parallel processing and pipelines in Silicon Graphics' Iris GT graphics system.

slope information to compute the coordinates and color of pixels between the vertices. Then five span processors compute the vertices and position of the pixels between the edges of the polygons.

The raster subsystem is fed with the color and coordinates

Continued on page 47

Easy ASIC

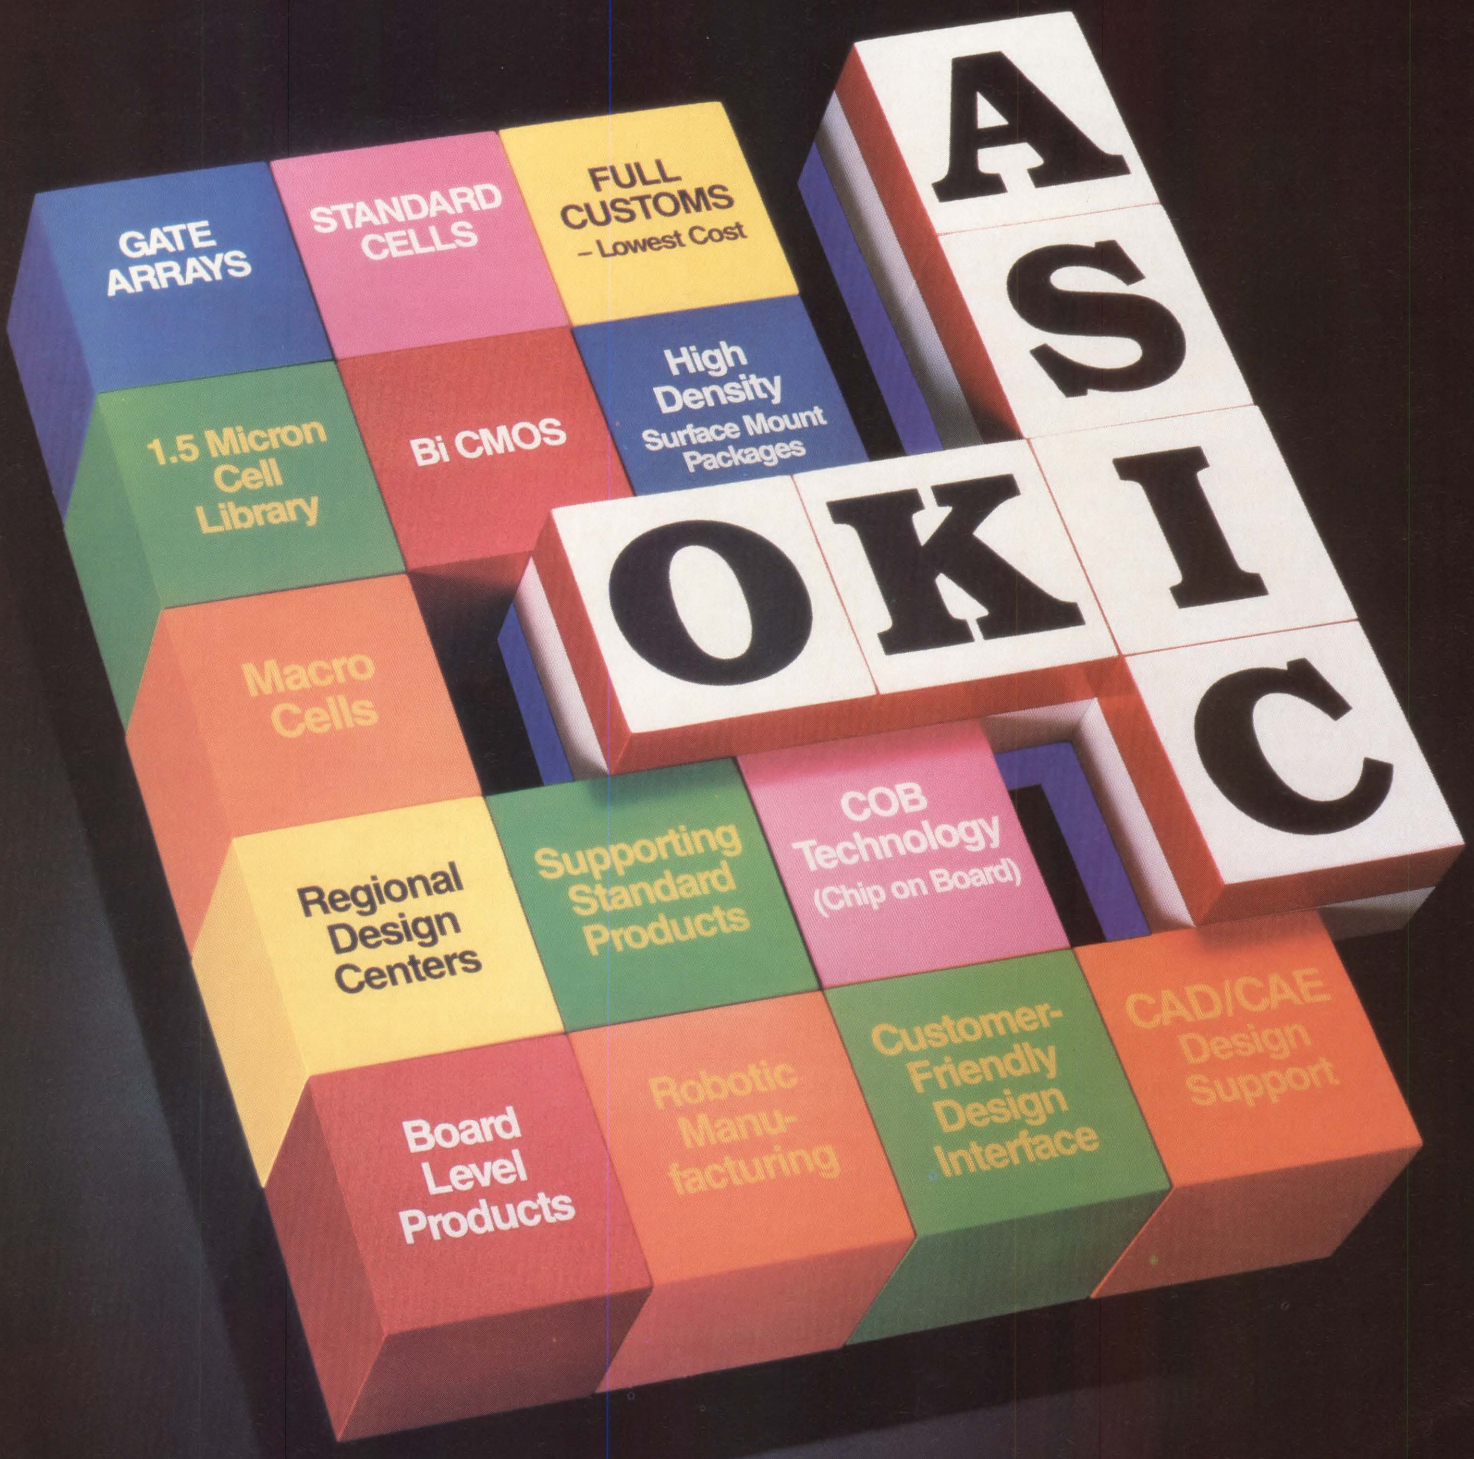
GATE ARRAYS
STANDARD CELLS
FULL CUSTOMS – Lowest Cost
1.5 Micron Cell Library
Bi CMOS
High Density Surface Mount Packages
Macro Cells
Regional Design Centers
Supporting Standard Products
COB Technology (Chip on Board)
Board Level Products
Robotic Manu-facturing
Customer-Friendly Design Interface
CAD/CAE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

Only from OKI:

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

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

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

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

OKI: the totally logical choice.

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

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

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

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

Check out OKI ASIC data:

VLSI 2/88

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

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

( ) Please call: we have immediate requirements:

Name/Title _____

Company _____

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

CIRCLE NUMBER 8

A Smart System That Compiles RTL Models from Schematics

Arie Brish, Ronen Keinan, and Yiftach Ravid, Motorola Semiconductor Israel Ltd., Ramat-Gan, Israel

The ongoing goal of cutting VLSI design turnaround time is approached from two complementary directions. From one side, engineers seek to improve design synthesis productivity; from the other, turnaround time is reduced by maximizing test and debugging productivity. Simulation forms a major bottleneck in the latter of these two approaches.

During the "LSI age," circuit and logic simulators were the most popular simulation tools among digital IC designers. Unfortunately, circuit and logic simulators are too slow and heavy to simulate a whole chip when the transistor count is in the five-digit range.

The register-transfer-level (RTL) simulator is one tool that helps improve testing and debugging productivity. Its contribution to the effort is essentially through two features: its capability of high-level hardware description and its speed.

This paper describes acceleration and modification techniques for extracting RTL simulation models directly from a logic-level schematic. A set of network compaction rules allows a logic-level description of a circuit (automatically extracted from the workstation database) to be compacted into an RTL description. This process yields benefits in simulation performance while preserving the reliability and validity of a schematics-driven logic network description. The network compaction process includes a number of separate elements, including switch gathering, complex logic gate gathering, complex gate simplification, elimination of irrelevant or redundant elements, and macro call gathering to register-level description. Test results indicate that a simulation runs between 2 and 30 times faster with a fully reduced and gathered RTL model.

Different Simulation Acceleration Techniques

In an effort to improve simulation throughput, developers have essentially taken four different approaches: hardware modeling, simulation algorithm improvement, hardware acceleration, and higher-level software modeling. Before discussing RTL modeling in specific, it is worthwhile to consider the merits and problems presented by alternative simulation acceleration methodologies.

Simulation through breadboarding (Pucket, 1983), or hardware modeling, is the fastest approach. Sadly, it is difficult to maintain. An interface between the simulator and the breadboard requires complex and closely knit interaction, and engineering changes must be rewired on the board. In addition, timing must still be calculated by software with hardware models: The hardware itself can provide only functional modeling.

Techniques that directly speed up the logic simulator itself have been presented in the past (Cernei and Gescai, 1985; Miyoshi et al., 1985). However, the remedies possible taking this route for simulation speed problems are limited, owing to the inherently detailed nature of the code needed for accurate logic simulation.

The third approach is to use special-purpose hardware accelerators (Blank, 1984). Accelerators are currently a costly solution. Since they are usually single-task machines, investment in a hardware accelerator is not often justified.

RTL simulation on a general-purpose computer is an alternative route by which logic simulation may be accelerated. In an RTL simulator, the design is described in high-level terms, such as registers, flip-flops, memories, and counters. Operations on the high-level elements and data transfers among the elements are performed by the RTL simulator (Evanczuk, 1986). Control logic is described with Boolean equations and state-machine behavioral descriptions. The language is usually very similar to a high-level programming language like PASCAL or C, with some relevant modifications.

Some Advantages of RTL Models

The high-level description in an RTL simulator affords the designer a number of benefits not always offered by other simulation acceleration methodologies.

An overwhelming advantage of RTL simulation is its speed. An RTL simulation can be approximately two orders of magnitude faster than a logic simulation of the same circuit.

As a direct result of the higher speed, productivity is increased. With RTL simulation, the designer is treated to a fast response. The use of a high-level description improves coding productivity, because of the lack of "noise" from detailed modeling. Debugging productivity is also increased: The designer can debug concepts early in the cycle, before the detailed design starts.

The high speed indirectly results in an increase of capacity, because larger circuits can now be simulated in a reasonable amount of time. In addition, the smaller design description database increases the capacity of the simulator; therefore, larger circuits or systems may be simulated.

Besides increasing speed, productivity, and capacity, RTL models have other advantages. For example, RTL modeling can be used for fault simulation: We have achieved a fault coverage of 98.45% with RTL modeling. This characteristic

deserves a separate discussion, which will not appear here.

Further, an RTL listing may be used as a master specification document. In many design groups, an RTL description is used primarily for circuit specification (Hollander, 1983; Rimkus et al., 1983); in addition, a detailed design can then be verified automatically, either by comparing simulation results (Doshi et al., 1984; Tham et al., 1984) or by static correctness proof (Barrow, 1984).

Moreover, maintenance of the models is much easier with RTL models than with hardware breadboarding; design management is simplified because there is a reduced amount of code; and multi-user multitask computers can be employed for simulation runs—dedicated, single-user, single-task machinery is needed for hardware acceleration.

Some Limitations

Unfortunately, the generation of suitable models is slower for RTL network descriptions than for either hardware modeling or hardware acceleration. RTL models are usually hand-coded, because the circuit description is at a high level; and manual network coding is risky, because the resulting model may not reflect the real circuit behavior. An automatic schematic-driven simulation netlist for logic simulation is more reliable; such a route causes slow simulation, however, owing to the detailed level of the generated model.

RTL simulations have several other limitations, including potential unreliability. Inaccuracies are usually referred to as "miss accuracies." An RTL simulation's accuracy problems can be due to a lack of charge-sharing support, fewer logic levels, or limited or no contention handling.

Charge-sharing effects are usually not supported by RTL simulators. A floating level will in many cases be treated as any other logic level, thus causing wrong propagation of floating levels. Problems can occur when a small capacitor is charging a large capacitor or when a small capacitor is causing contention with a driving gate.

RTL simulators usually support a limited number of logic levels. A limited number of levels can cause some misleading results, especially in those that have two logic levels only.

Adequate contention handling is lacking, too, because of the lack of strengths, so that wrong handling of contention cases may occur. In multipass RTL simulators, the strength is sometimes determined indirectly by ranking the equations; in such cases, the equation executed first is the weakest. The engineer should take this limit into account while coding his circuit into the RTL simulator. In RTL simulators that support unknown states, an UNKNOWN level will occur in contention cases. In those that do not support unknown states, a contention may cause nonconvergence or may give a wrong result.

RTL Model Generation Techniques

One way to increase simulation speed without losing accuracy is to combine hand-coded RTL descriptions with schematic-driven netlists (Doshi et al., 1984; Tham et al., 1984; Rimkus et al., 1983; Lathrop and Kirk, 1985; Adler, 1986). This approach is often called mixed-mode simulation (which can be confusing because the term is also used to refer to mixed analog-digital simulation). Unfortunately, the hand-coded high-level circuit description may not match the original design, because of human coding errors. Such a case can hide bugs, as simulation on a wrongly hand-coded circuit

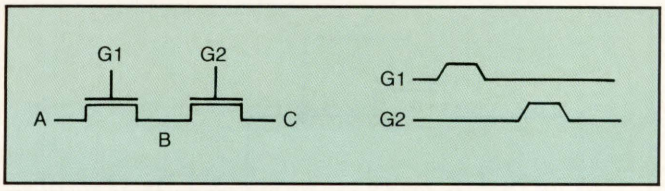

FIGURE 1. Transfer gates can cause problems.

description may accidentally give correct results although a bug exists in the schematic.

A second technique arises from the automatic design synthesis feature of silicon compiler systems (Rimkus et al., 1983; Thomas and Nestor, 1983; Thomas, 1986). The reliability of hand-coded RTL descriptions is higher with the compilation approach because the lower-level design is automatically synthesized by the systems. Thus speed is improved and accuracy is maintained.

The third technique is "smart" schematic-to-RTL translation. Automatic schematic-driven RTL coding is possible. However, a straightforward approach may eliminate major RTL modeling advantages, because of the detailed code generated and the inaccuracies that may occur.

Automatic Translation

With most VLSI designs, a significant amount of the detailed design is hand-drawn on a workstation screen. It is therefore desirable to extract RTL descriptions automatically from the schematics.

In logic, switch-level, or circuit simulation, an automatic netlist generation from the workstation is a straightforward task. In RTL simulation, the translation is not as easy. The translation program should take into account the limitations mentioned above and any other limitations of the specific RTL language. In addition, a one-to-one translation will generate a very detailed code, thus diminishing the major advantage of RTL modeling.

In the following sections, we describe some particular techniques for synthesizing an efficient high-level RTL code from a detailed circuit schematic while avoiding simulator limitations that may cause erroneous results. The RTL code generation involves several code reduction and compaction steps: transfer gate gathering, equation reordering, redundant component drop, network simplification, and macro cell gathering.

Gathering Transfer Gates

Transfer gates must be gathered in order to avoid charge-sharing mishandling as described above. Let us consider the network in Figure 1. In a straightforward translation, the RTL equations would be:

```
WHEN G1 DO B = A ; WHEN G2 DO C = B        (1)
```

Notice that with the G1 and G2 timing as described in Figure 1, C may erroneously change its value: In the real world, B's capacitance may not be large enough to change C's value. The translator would therefore translate this case to the following:

```
WHEN G1·G2 DO C = A                        (2)
```

Reordering the Equations

The next phase is equation reorganization. In this step,

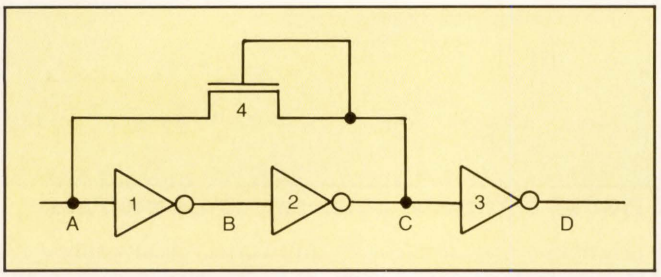

FIGURE 2. Redundant circuitry can be removed.

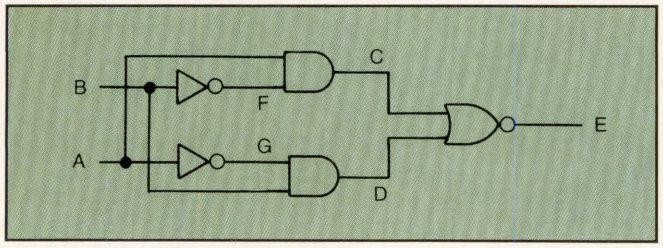

FIGURE 3. A candidate for simplification.

equations are rearranged according to their mutual dependencies. By reordering the equations we achieve both performance improvement and limited strength handling. Owing to our multipass relaxation algorithm dependency, ordered equations need fewer loops to stabilize. For example, the equations in (3) will take longer to settle down than the equations in (4):

$$A = B \;:\; B = C \;:\; C = D \tag{3}$$
$$C = D \;:\; B = C \;:\; A = B \tag{4}$$

By reordering the equations, we can account for limited strength handling. We take advantage of the fact that later equations override earlier equations. We put the weakest equation first while avoiding contention later.

Removing Irrelevant and Redundant Elements

In the next step we remove from the RTL description any element that does not add information to the RTL level, such as positive feedback and even inverter trains. Let us consider the example in Figure 2. In a straightforward translation, the RTL code would be:

$$B = \overline{A} : C = \overline{B} : AB = C : D = \overline{C} \tag{5}$$

After the depletion feedback and inverters 1 and 2 are dropped, the RTL code would be:

$$D = \overline{A} \tag{6}$$

Note that by doing so we also provide false contention problems at node A.

Simplifying the Network

In some cases predefined gate combinations are reduced to more simple forms. For example, a boostrap circuit can be simplified into an AND gate and a complex XOR combination can be simplified into a single XOR gate. Let us consider the circuit example in Figure 3. In a straightforward translation, the RTL code would contain five operations:

$$F\overline{B} : G = \overline{A} : C = A \cdot F \tag{7}$$
$$D = G \cdot B : E = \overline{(C + D)}$$

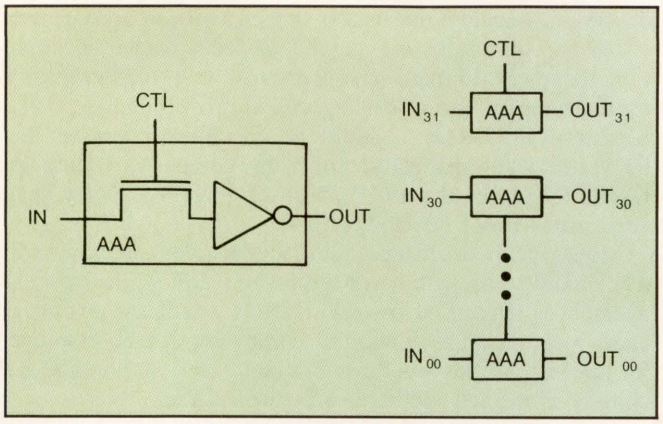

FIGURE 4. Repeatedly used macros can be gathered for code compaction and efficiency.

whereas the simplified RTL equation contains only one operation, as shown in equation (8):

$$E = \overline{(A + B)} \tag{8}$$

Gathering the Macro Calls

The cherry in our cake is macro call gathering. In this step, the circuit description level is raised from signals and single-bit macro calls to buses and multiple-bit macro calls. Up to this step, a 32-bit register will be translated to 32 calls of a single-bit latch macro; after this step, the register is reduced to one call of a 32 bit register macro. The speed improvement here is due to the implementation of buses as integers, as well as signals, in our simulator. Therefore the time needed for an AND operation between two signals is equal to the time needed to do an AND operation between two buses. In other words, the time it takes to do an AND operation between two 32-bit buses, bit by bit, is 32 times longer than the time it takes to do one operation on the whole bus. Without macro call gathering, the RTL list would be:

```
begin (* expansion of macro AAA *)
*E00113 AAA (CTL, IN31, OUT31)
*E00114 AAA (CTL, IN30, OUT30)
*E00115 AAA (CTL, IN29, OUT29)
*E00144 AAA (CTL, IN01, OUT01)
*E00145 AAA (CTL, IN00, OUT00)
end (* of AAA's expansion *)
```

With the macros gathered, the RTL list is:

```
begin (* expansion of macro AAA *)
*E00113 AAA (CTL, bus (IN31, IN00),
a bus (OUT31,...,OUT00) end
```

The single macro AAA shown in Figure 4 is defined as:

$$\text{WHEN CTL} : \text{OUT} = \overline{\text{IN}} : \text{END WHEN}$$

where IN and OUT are declared to be signals. The same macro definition will work for the gathered case with the difference that IN and OUT are declared as 32-bit buses.

We have expanded this idea from vector gathering to include matrix gathering. Any $n \times m$ calls of macro AAA with m data bits and n control signals will be gathered into one macro call with an n-bit control bus and an m-bit data bus. Let us consider the previous example and expand it to the case in

Tool	Execution time[1]	Code size[1]
Switch level	74.60	27.15
Detailed RTL simulation	9.00	26.31
Reduced RTL, data macros gathered	3.44	7.02
Reduced RTL, all macros gathered	1.00	1.00
High-level hand-coded RTL	0.11	0.23
[1] Relative to reduced RTL with all macros gathered		

Table 1. Performance of RTL models.

which macro AAA is expanded with different controls, CTL1 . . . CTLn. The equation for this further gathering is:

$$OUT = (OUT \cdot (\overline{CTL}) + (\overline{IN}) \cdot CTL) \quad (9)$$

Notice that register AAA is now duplicated n times with different control CTLi ($i = 1 \ldots n$) for each register. Also notice that, with this additional gathering, we reduce a register file matrix to a single macro call.

Macro gatherings significantly reduce the number of nodes in the model, thus contributing significantly to fault simulation throughput.

Results

We ran several tests on the reduced RTL code achieved by the translation program. The results show an improvement that varies from 2 to 30 times. The worst results are achieved in highly random designs; the best results are achieved for highly structured, highly repetitive designs.

For the test, we took one example with 10,000 transistor blocks, containing approximately 70% sequential logic (registers, counters, shift registers, and such), 20% PLAs, and 10% random logic. We generated a simulation model in five different representations. The lowest level was the switch level; the next level was an automatically generated "detailed" (logic-level) RTL; then a reduced RTL with macro gathering on data bits only; followed by a reduced RTL with macro gathering on data bits and control signals; and finally a hand-coded high-level RTL description for the same circuit. The relative results are summarized in Figure 5. Implementation of these minimization techniques, along with automatic translation, has clearly yielded an order of magnitude improvement in simulation throughput. □

Acknowledgments

We would like to thank the MC56000 design team at Motorola Semiconductor Israel. Being the first heavy users of this tool, they made a great contribution to its debugging. Thanks also to Yehuda Schaik, who defined and coded the network simplification rules.

References

Adler, D. 1986. "SIMMOS—A Logic/Timing Simulator for Digital CMOS VLSI," *23rd Design Automation Conference*, Las Vegas, NV.

Barrow, H.G. July 1984. "Proving the Correctness of Digital Hardware Designs," *VLSI Design*.

Blank, T. August 1984. "A Survey of Hardware Accelerators Used in Computer Aided Design," *IEEE Design & Test*.

Cerny, E., and J. Gescei. October 1985. "Simulation of MOS Circuits by Decision Diagrams," *IEEE Transactions on CAD*.

Doshi, M.H., R.B Sullivan, and D.M. Schuler. 1984. "THEMIS Logic Simulator—A Mix Mode, Multi-Level, Hierarchical, Interactive Digital Circuit Simulator," *21st Design Automation Conference*, Albuquerque, NM.

Evanczuk, S. January 1985. "Getting Out of the Gate: High-Level Modeling in Circuit Design," *VLSI Design*.

Hollander, Y. September 1983. "Using an RTL Simulator to Simplify VLSI Design," *VLSI Design*.

Lathrop, R.H., and R.S. Kirk. 1985. "An Extensible Object-Oriented Mixed-Mode Functional Simulation System," *22nd Design Automation Conference*, Las Vegas, NV.

Miyoshi, M., Y. Kazama, O. Tada, Y. Nagura, and N. Amano. 1985. "Speed Up Techniques of Logic Simulation," *22nd Design Automation Conference*, Las Vegas, NV.

Pucket, R.R. March 1983. "Standard Schottky TTL ICs Emulate MOS Logic Designs," *Electronic Design*.

Rimkus, C.J., M.R. Wayne, D.D. Cheng, and F.J. Magistro. 1983. "Structured Design Verification: Function and Timing," *20th Design Automation Conference*, Miami, FL.

Tham, K., R. Wiloner, and D. Wimp. 1984. "Functional Design Verification by Multi-Level Simulation," *21st Design Automation Conference*, Albuquerque, NM.

Thomas, D. February 1986. "RTL Simulation Makes a Comeback for Complex VLSI," *Computer Design*.

Thomas, D.E., and J. A. Nestor. July 1983. "Defining and Implementing a Multilevel Design Representation with Simulation Applications," *IEEE Transactions on CAD*.

About the Authors

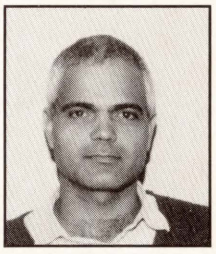

Arie Brish has been department manager of engineering support at Motorola Semiconductor Israel since 1982. From 1978 to 1982 he worked for National Semiconductor as a design engineer on the 32000 program, both in Israel and in Santa Clara, CA. Brish received a BSEE degree from Ben Gurion University in 1978 and an MBA degree from Tel Aviv University in 1984. His interests include VLSI design productivity, functionality verification, and operation management.

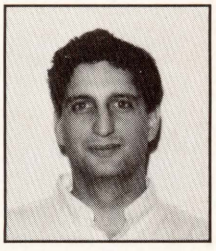

Ronen Keinan has been a CAD software engineer at Motorola Semiconductor Israel since 1984 and is currently leading its layout automation activities. The recipient of a BSc in mathematics and computer science from Tel Aviv University in 1984, he is working on his MSc in computer science at the same school. His interests include logic simulation, compilation, and layout automation.

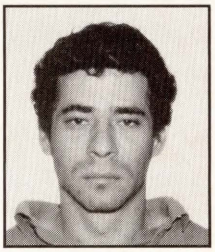

Iftach Ravid has been a CAD software engineer at Motorola Semiconductor Israel since 1985. He received his BSc in mathematics and computer science from Tel Aviv University in 1986 and is currently working on his MSc in computer science at the same school. Iftach's interests include logic simulation, compilation, random algorithms, and combinatorics.

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

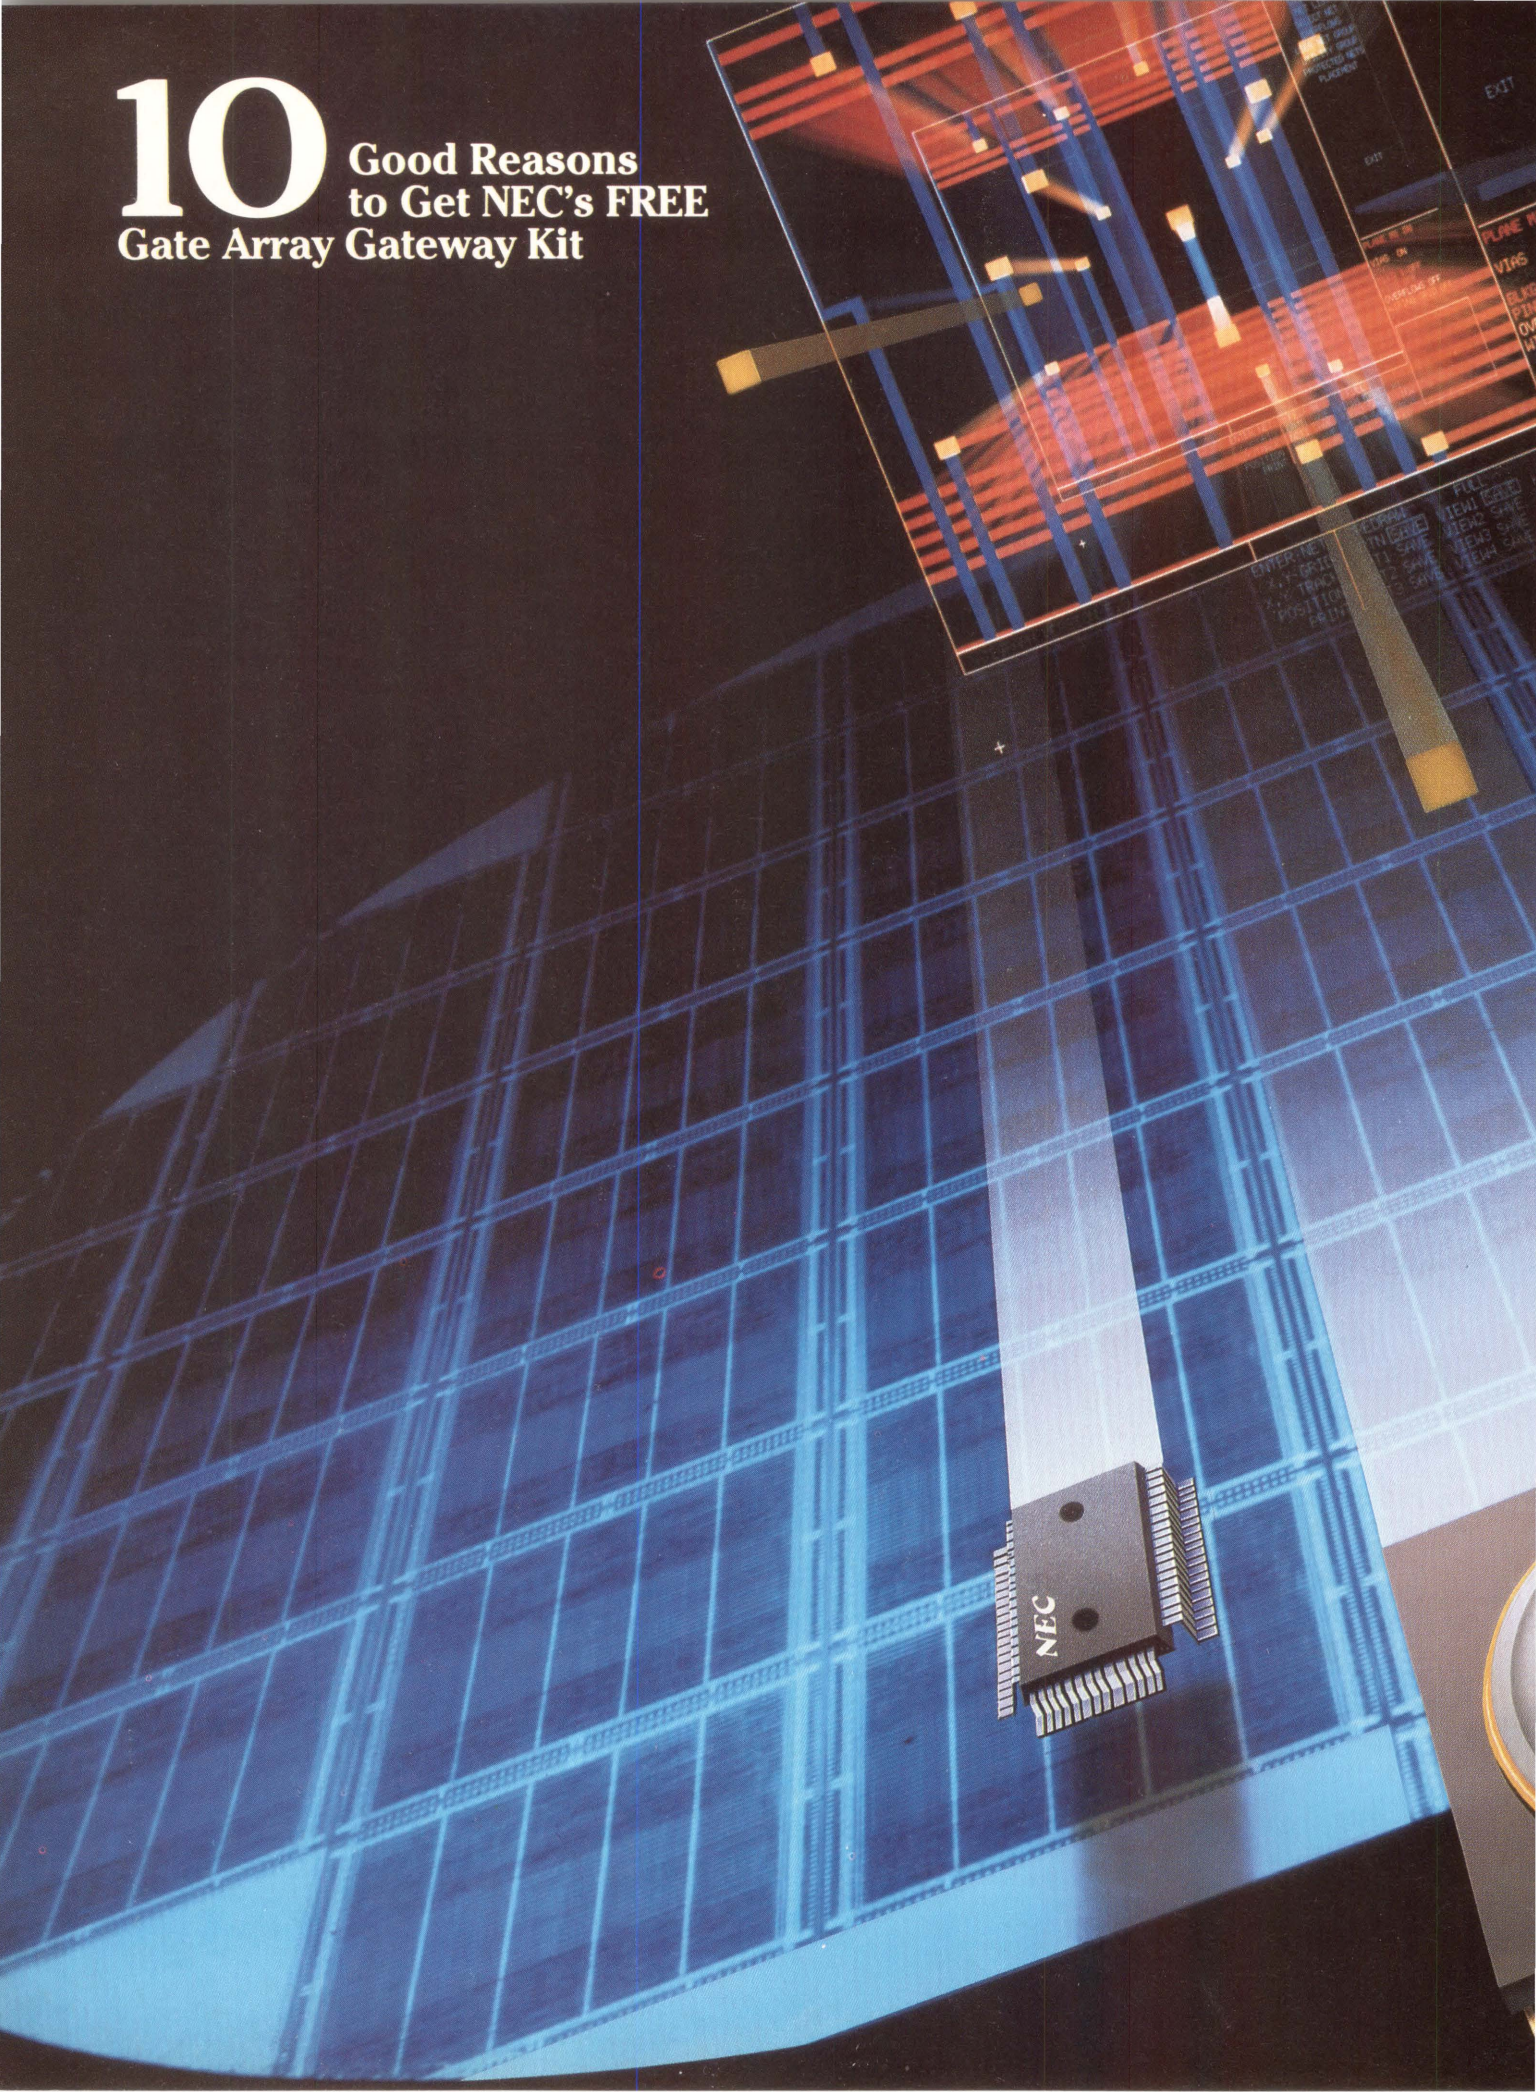

NEC's FREE Design Kit is Your Gate Array Gateway to

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

NEC Electronics Inc.

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

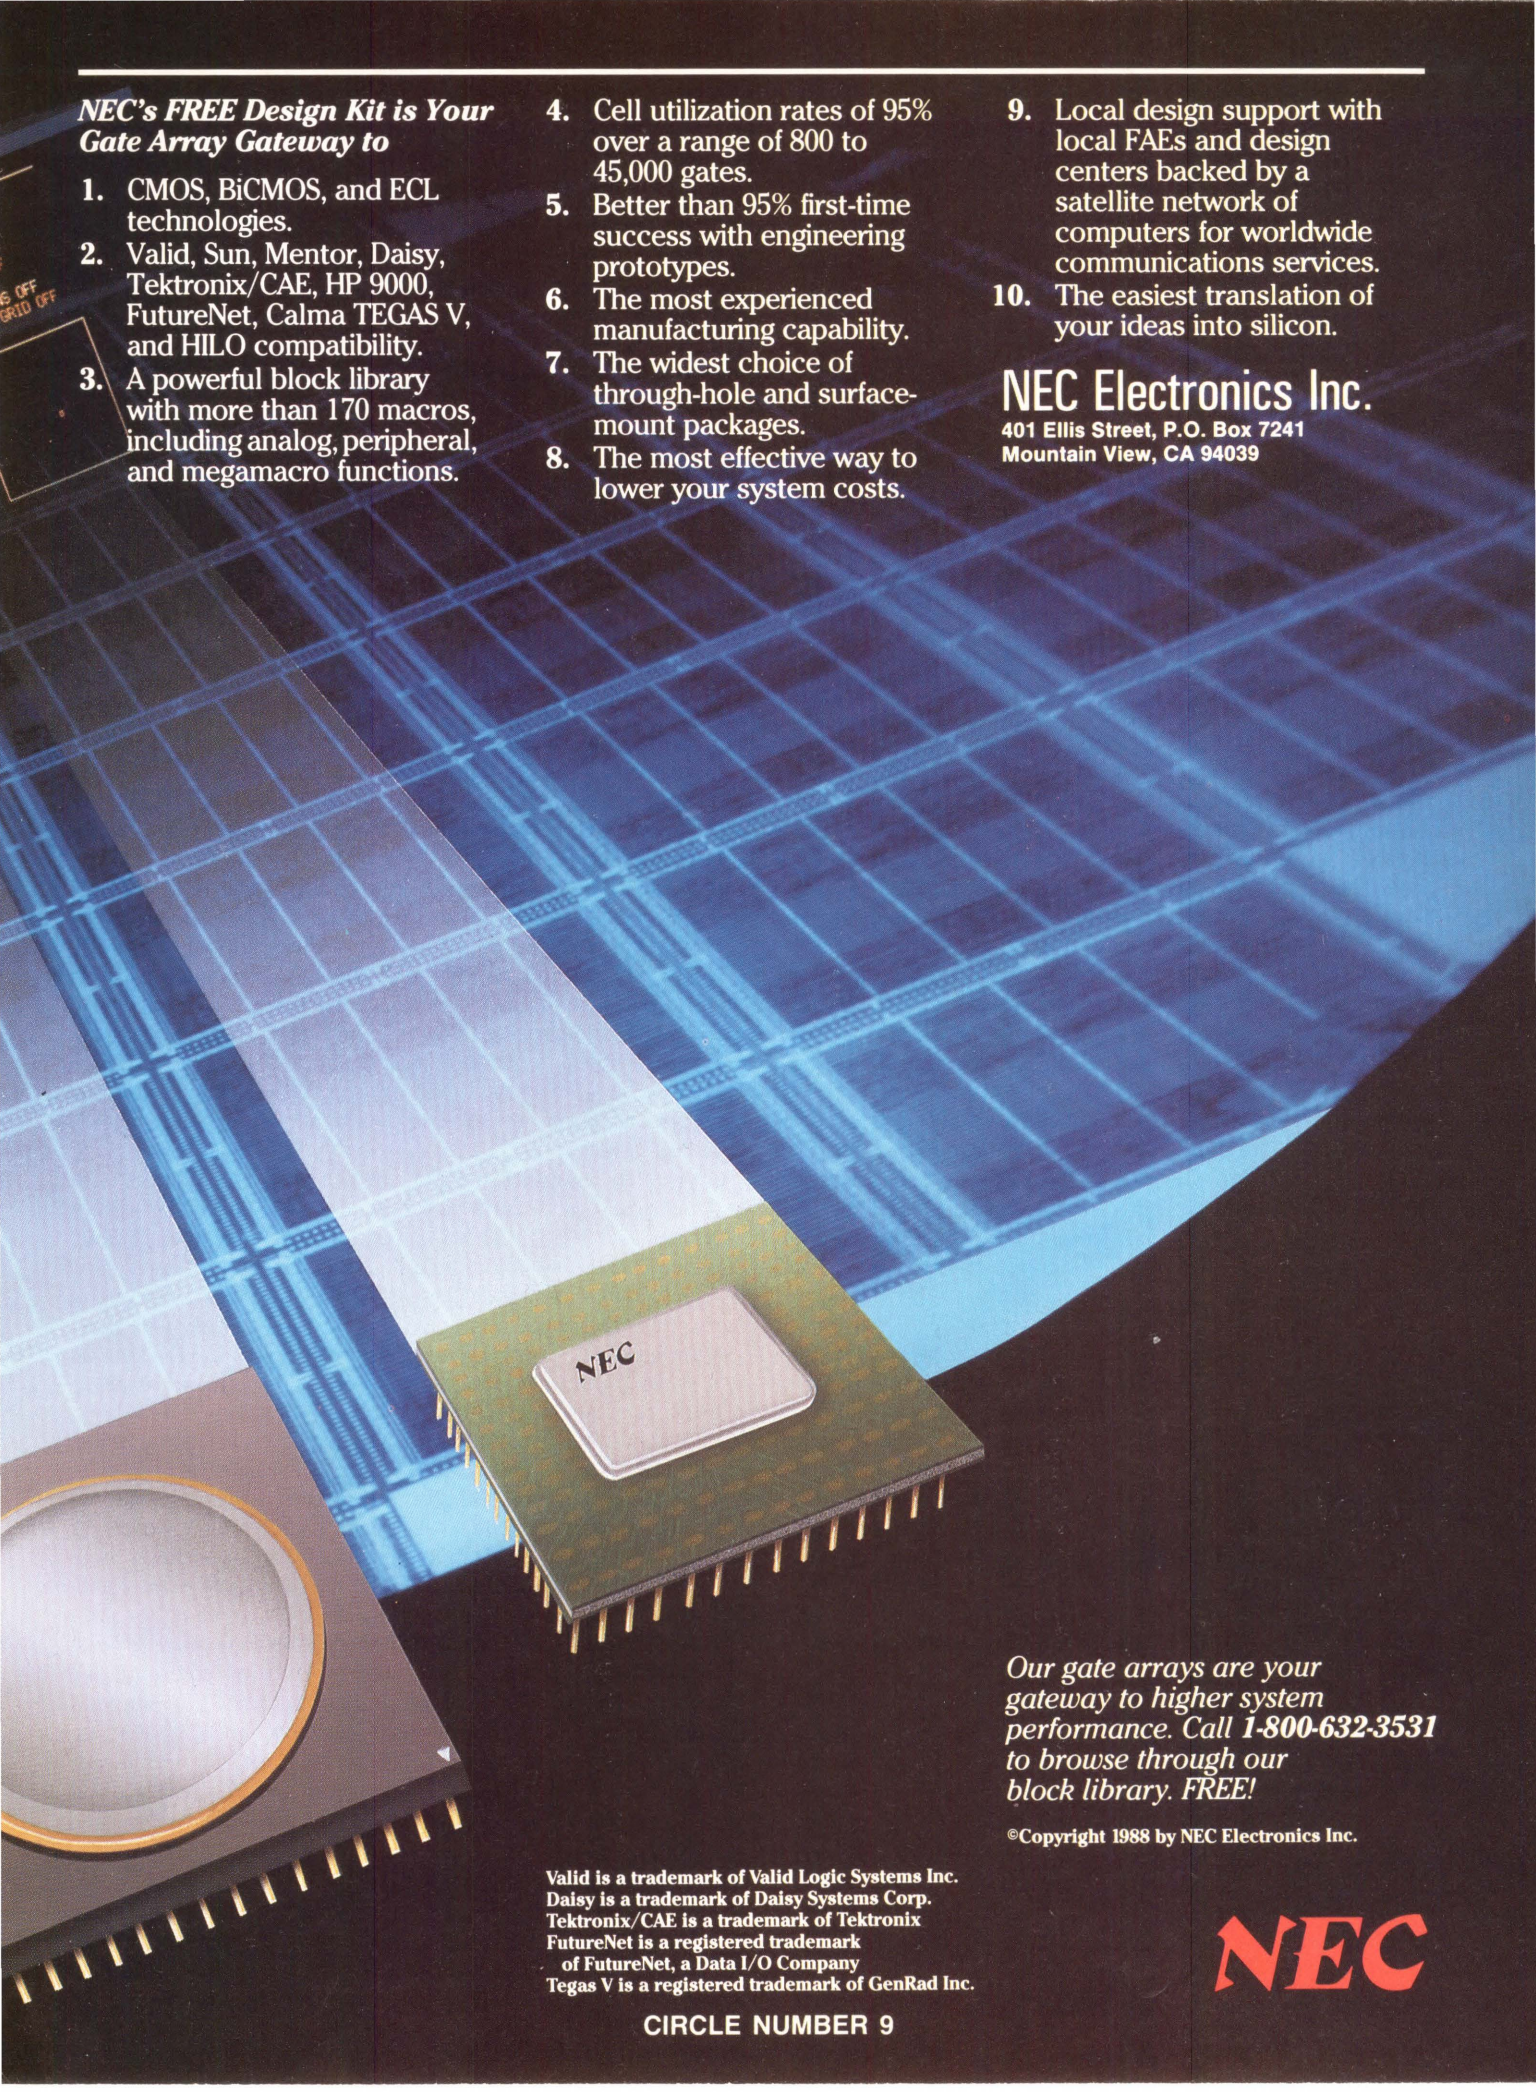

Our gate arrays are your gateway to higher system performance. Call ***1-800-632-3531*** *to browse through our block library. FREE!*

NEC

CIRCLE NUMBER 9

Coupling a Digital Logic Simulator and an Analog Circuit Simulator

Tedd Corman, Viewlogic Systems Inc., Marlboro, MA
Michael U. Wimbrow, MicroSim Corp., Laguna Hills, CA

As systems integrate analog and digital functions, separate simulation of the analog and digital sections no longer provides an adequate level of verification. Entire designs must be simulated as an integrated system because the designer must verify both the interaction between analog and digital subsystems and the functionality within them. Mixed-mode solutions—simulators that can model and simulate analog and digital circuitry—must provide the accuracy required for critical analog circuitry, as well as the performance of a logic simulator necessary for the verification of large electronic systems.

Along with mixed analog and digital systems, digital-only and analog-only designs can benefit from a mixed-mode solution. Mixed-mode simulation can provide more accurate analog models for digital devices in critical signal paths in all-digital systems, and large analog-only designs can be simulated faster when portions of a design can be characterized in terms of digital logic models instead of more complex analog device models.

Historical Perspective

Some early mixed-mode simulators (Newton, 1978; Chawla et al., 1975) attempted to merge analog and digital simulation algorithms into a single simulaton environment. Although many of these simulators provided some performance increases over existing analog-only simulators, most did not demonstrate the desired performance improvements, because of architectural problems resulting from combining two different simulation algorithms. Many of these simulators evaluated the entire circuit at time points when either analog or digital portions changed value, resulting in overall performance close to that of the slower analog simulator. In addition, partitioning designs into separate analog and digital sections in the manner required by the simulator proved to be difficult. Additional code had to be added to handle the different analog and digital simulator primitive evaluation routines.

Other systems (Friedman, 1987) were built upon existing analog or digital simulators, adding enhancements to support modeling of the missing counterpart. Some of these systems used behavioral models to characterize analog components in a digital simulator; others added special digital models to an analog simulator that utilized simpler, faster model evaluation algorithms. Athough sometimes effective in simulating predominately analog or digital circuits, neither approach provides an effective solution for a large class of mixed analog-digital systems.

Special models must be written to simulate mixed systems. These models are usually difficult to develop because they tend to be very different from the rest of the simulator models. For example, the models may have to be modeled in a programming language. Because the models are retrofitted into an existing analog or digital simulator, they tend to suffer from accuracy or performance problems. Finally, although capable of simulating digital portions of a system faster than analog portions, these simulators cannot simulate digital circuitry nearly as fast as existing digital simulators.

Another approach, unified mixed-mode simulation, was developed based upon entirely new simulation algorithms (O'Rouke, 1988). These algorithms provide for flexible, generalized modeling techniques, supporting not only analog and digital circuit modeling capabilities, but also modeling of mechanical, chemical, and other systems. Although these simulators offer the clear advantage of broadening the simulation capabilities, they do suffer from some significant drawbacks. Because they are not based upon existing technology, they cannot take advantage of the wide base of existing model libraries. New models have to be developed using new, unfamiliar modeling techniques. Moreover, once a model has been built, it is difficult to determine its accuracy, or even the accuracy of the simulation algorithm itself, because these simulators have a brief history compared with more established simulators such as SPICE.

Some recently developed "glued" mixed-mode simulators (Sarin et al., 1987)) use two loosely coupled digital and analog simulators. Glued approaches have the advantage of providing distinct, special-purpose analog and digital simulation algorithms. Because of their loosely coupled, sequential operation, however, they cannot effectively simulate systems that have tight feedback between analog and digital circuitry.

VIEWSIM/AD Architecture

VIEWSIM/AD (VIEWSIM Analog Digital Simulator) is an

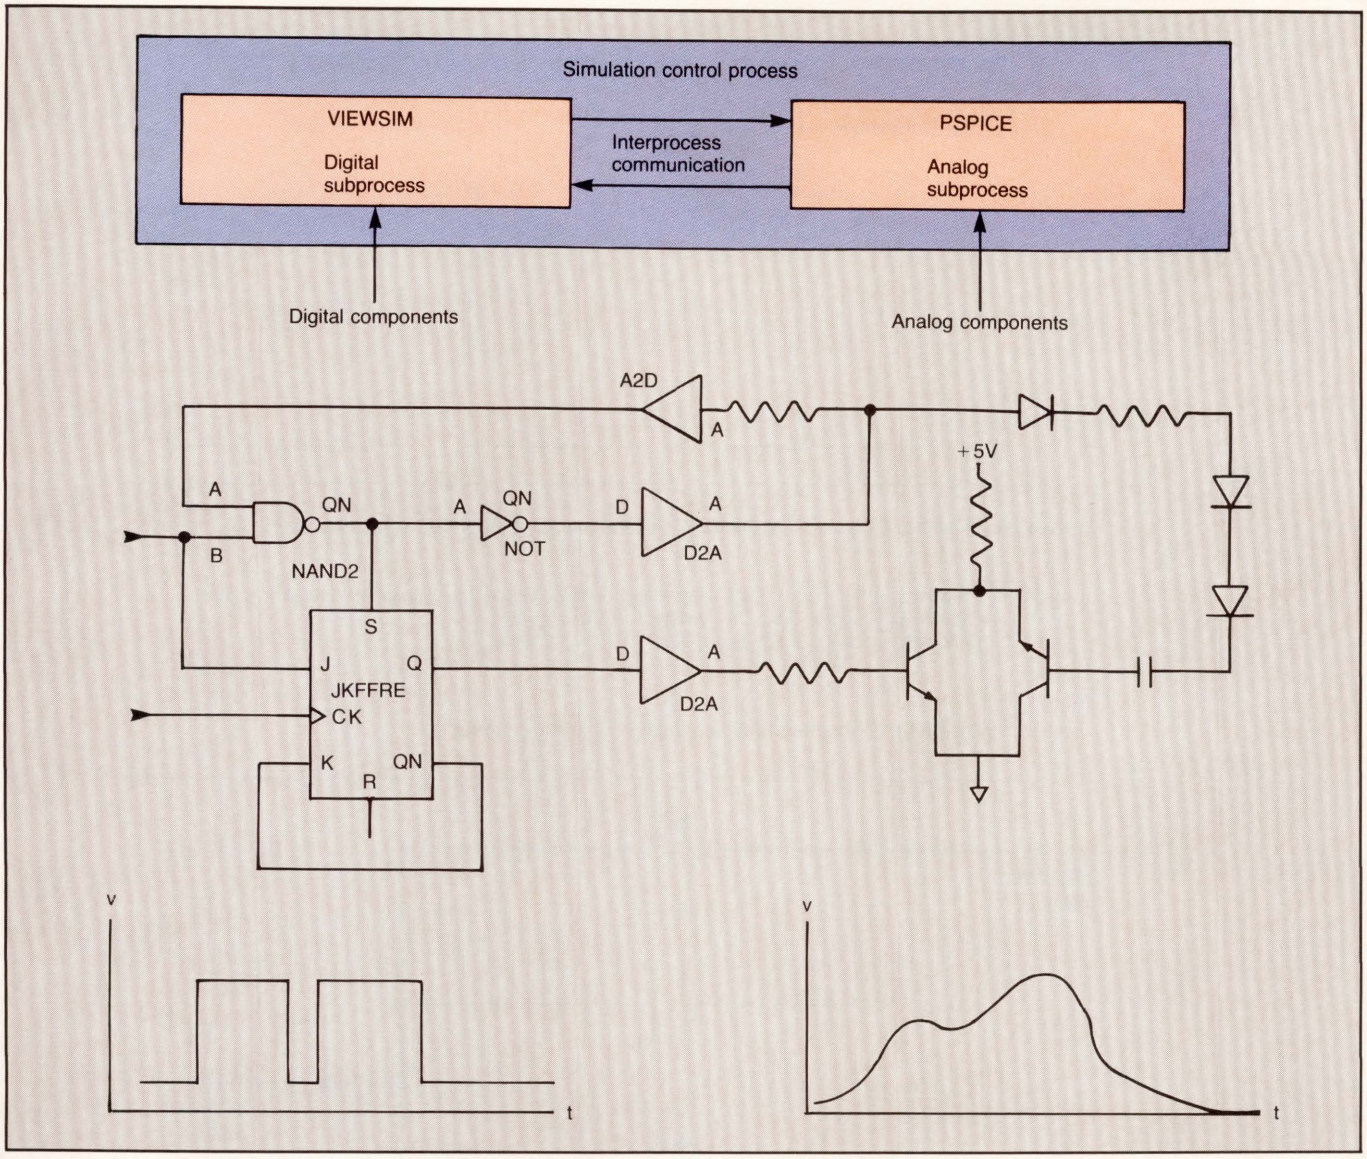

Figure 1. Architecture of the VIEWSIM/AD analog-digital simulator.

integrated, mixed-mode simulator capable of simulating mixed analog and digital systems. It is based on VIEWSIM, our 28-state logic simulator, and PSPICE, a Berkeley SPICE2-based analog circuit simulator from MicroSim Corp. The VIEWSIM/AD simulator comprises three parts: the analog simulator, the digital simulator, and the VIEWSIM/AD simulation control process (see Figure 1).

Because multitasking operating system capabilities are exploited, the digital and analog simulators run as separate subprocesses under the control of the VIEWSIM/AD simulation control process. The control process is responsible for initialization and synchronization of the PSPICE and VIEWSIM simulation subprocesses. Synchronization of the simulators is performed by a proprietary algorithm that has access to information about the PSPICE internal time step and the VIEWSIM simulation event wheel. The algorithm allows VIEWSIM to schedule logic events in its regular fashion and PSPICE to calculate circuit values in its regular fashion. PSPICE time step sizes are adjusted only when a signal that affects both digital and analog circuitry changes state. If the amount of activity across the boundaries between analog and digital circuits is small in comparison to the activity in the rest of the circuit, performance is significantly improved over simulating a circuit in PSPICE alone.

State information is passed between the two simulation subprocesses by interprocess communication. When a VIEWSIM "analog output" changes state, a node identifier is sent to PSPICE, along with the new state and time. PSPICE's matching "digital input" device maps the digital logic state into corresponding analog circuit values. PSPICE also has a "digital output" device, which monitors the voltage on the node to which it is connected. When the voltage crosses one of the specified thresholds, PSPICE passes the corresponding logic state value to the VIEWSIM "analog input" device, which applies the state change to the connected digital node. The size and frequency of the information transfer is minimal because only a node identifier, a simulation time, and a digital state value are transferred when mixed nodes actually change their digital state value.

This approach has several advantages over other mixed

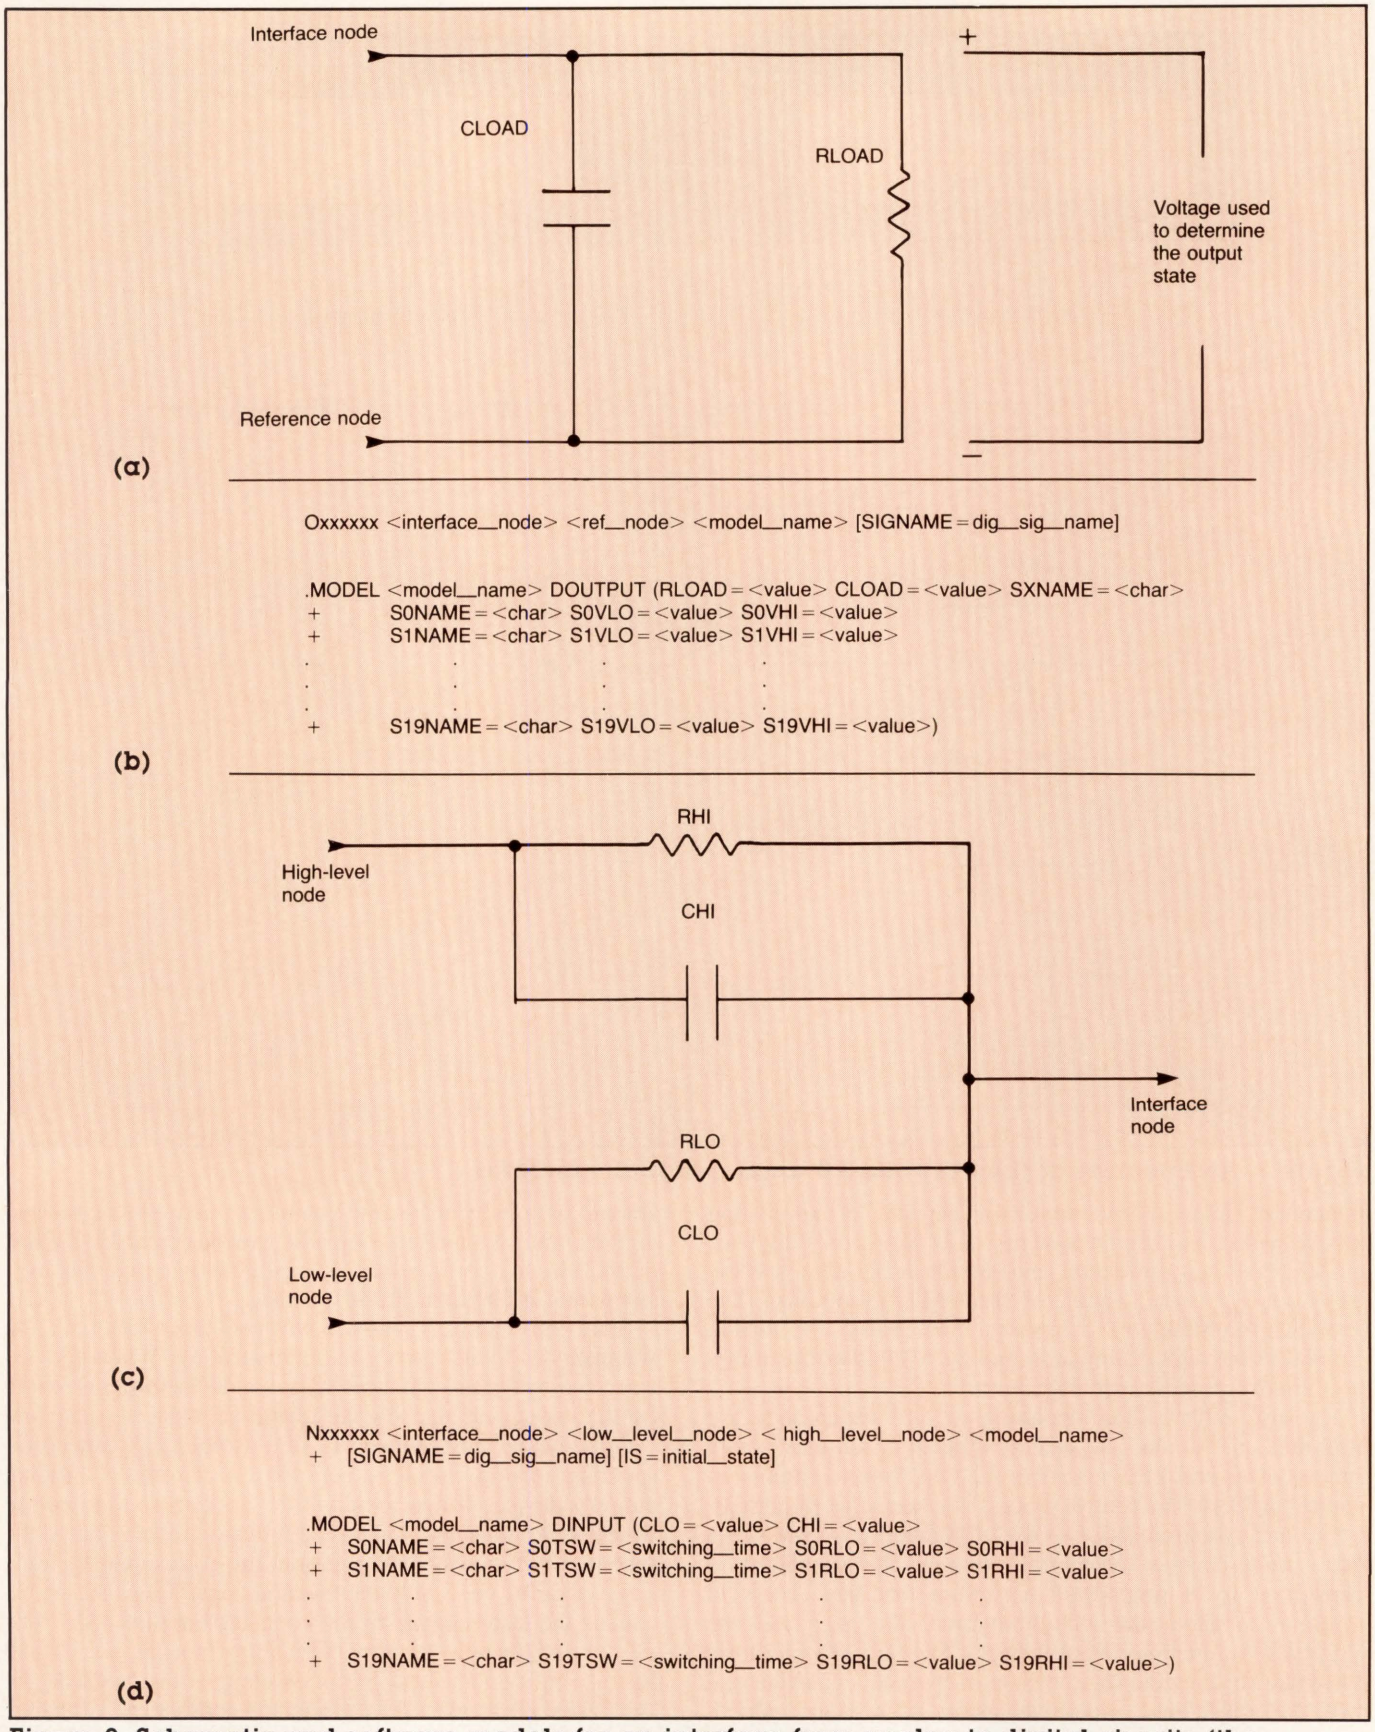

Figure 2. Schematic and software models for an interface from analog to digital circuits (the "digital output device," a, b) and from digital to analog circuits (the "digital input device," c, d).

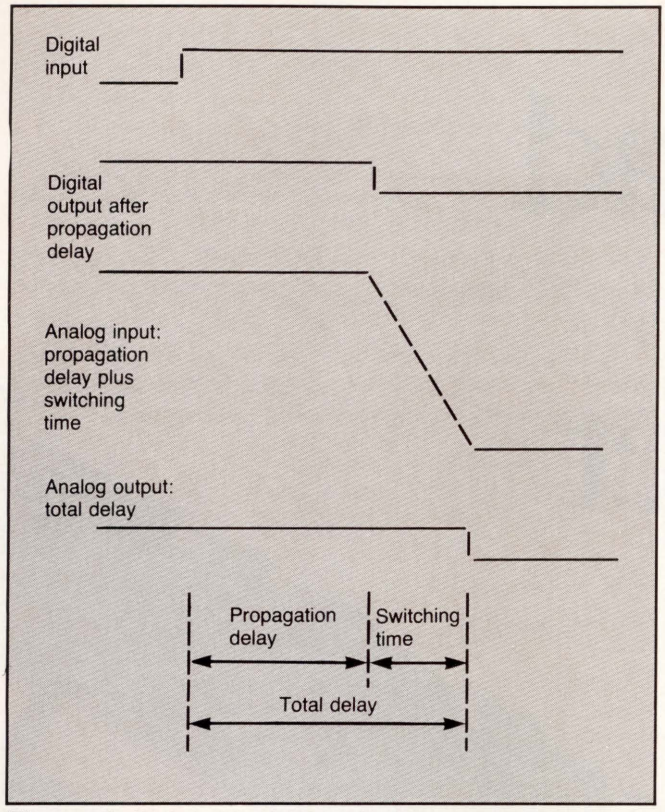

Figure 3. Delay modeling between analog and digital simulators.

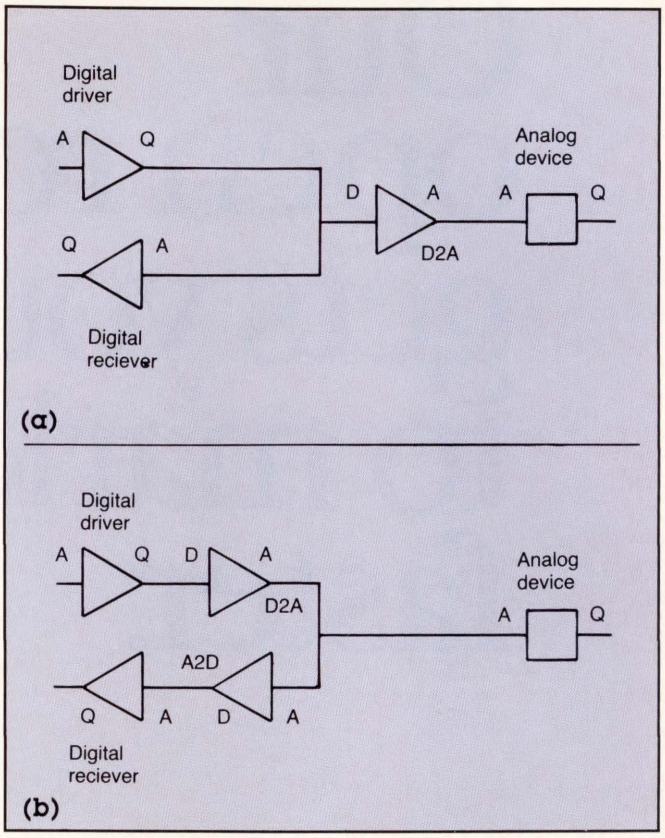

Figure 4. Modeling a bus interface with a small analog load (a) and a significant analog load (b).

analog-digital simulators: Because VIEWSIM/AD is based on existing simulation technology, it can use existing digital and analog model libraries. The basic simulation algorithms of PSPICE and VIEWSIM are unmodified, so that VIEWSIM/AD makes few performance compromises to simulate mixed analog-digital designs. Arbitrarily interconnected mixed designs can be simulated, including systems with feedback and even those with a mix of analog and digital components within a feedback loop.

Furthermore, the system partitions the design into analog and digital subcircuits. Many other mixed-mode approaches require the user to partition his design; such analog/digital partitioning considerations impose design constraints. For example, SPLICE (Newton, 1978) cannot simulate hierarchical designs and MOTIS (Chen, 1984) places restrictions on how a designer can interconnect devices.

Modeling the Analog-Digital Interface

Two special types of PSPICE devices model the interface between the logic simulator and the circuit simulator: a digital-to-analog (d/a) interface and an analog-to-digital (a/d) interface. PSPICE model statements control the characteristics of the interface devices. By varying these PSPICE models, the user can create interfaces for any logic family. In fact, a single circuit can contain several interface models, each representing an interface for a different logic family.

The user can denote d/a and a/d interface connections explicitly by placing vendor-supplied d/a and a/d interface components on his schematics. Alternatively, he can elect to have the VIEWSIM/AD's netlister (described below) automatically insert the necessary interface models. PSPICE interface model names are specified by adding an attribute to an interface component or, if interface components are not used, by adding an attribute to the net connecting analog and digital components. This choice allows the designer to use different technology models at different interface points within a design.

The a/d device in PSPICE is called the "digital output" model (see Figure 2a). The digital output is modeled as a fixed resistor and a fixed capacitor connected between the interface node and a reference node (usually ground). The resistor and capacitor model the loading presented by the logic circuit's inputs to the analog circuit. (The loading may be modeled more accurately if desired, as discussed later.)

The voltage across the digital output device determines the state of the logic simulator for the interface node. The PSPICE model for the interface device specifies a voltage range associated with each VIEWSIM logic state. A digital output model signals a state change to VIEWSIM when its node voltage crosses a specified voltage threshold. Because voltage ranges can overlap, the user can build hysteresis into the digital output model.

The VIEWSIM/AD netlister, the program that produces a netlist from the user's schematic design, produces a device statement for each analog-to-digital interface node. The statement specifies PSPICE interface and reference nodes (logic ground), the name of the model that describes this digital output device and the name of the digital signal in VIEWSIM controlled by this analog-to-digital interface. The user can insert as many digital output devices (and hence analog-to-

Our open road gets you to market faster.

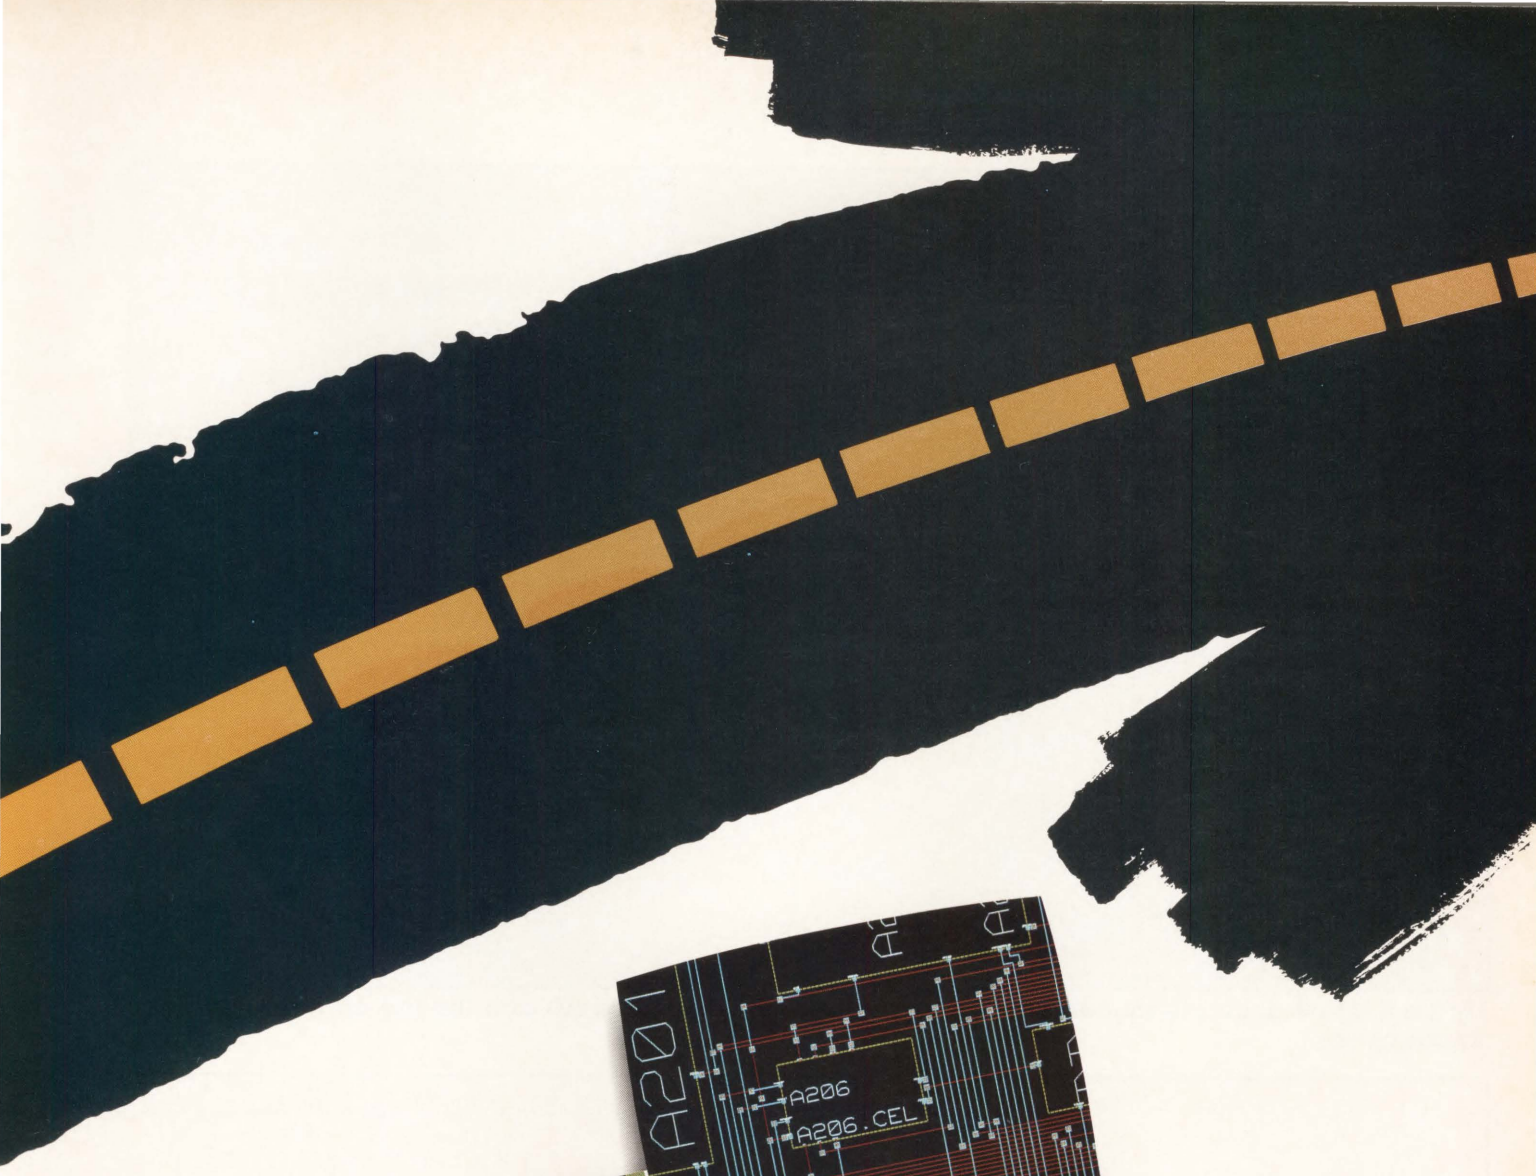

Our open road architecture means you can choose the IC design tools and hardware you want, when you want them, knowing that each addition contributes synergistically to the power of your IC design system.

Our SYMBAD™ family of integrated layout tools can take you from netlist to maximum density layout from five to ten times faster than with conventional tools. Design rule independent results mean that creation of new layouts for changed processes is almost instantaneous. A new layout language captures your methodologies, converting design expenses into assets. Our industry standard DRACULA™ family, interactively integrated with SYMBAD, ensures that your silicon will work with good yield, the first time you make it. DRACULA's established hierarchical techniques verify your full layout from five to ten times faster than with conventional methods.

Take the fast lane to market and enjoy our open road: an open database and language for total control, an open chance to use multiple artwork and netlist sources, the open opportunity to work with all the popular hardware choices (DEC, Apollo, Sun, IBM), and the peace of mind of dealing with an industry leader. Contact us today: ECAD, 2455 Augustine Drive, Santa Clara, CA 95054. (408) 727-0264.

Western U.S. (408) 727-0264; Southwestern U.S. (714) 752-8451; Mid-America (214) 869-0033; Eastern U.S. (201) 494-8881; Boston (617) 448-3543; United Kingdom 44-932-568544; Japan 81-3-989-7881; France 33-1-64460531; West Germany 49-89-1782047; Asia 852-5-660136.

CIRCLE NUMBER 10

ECAD

Optimizing time and space

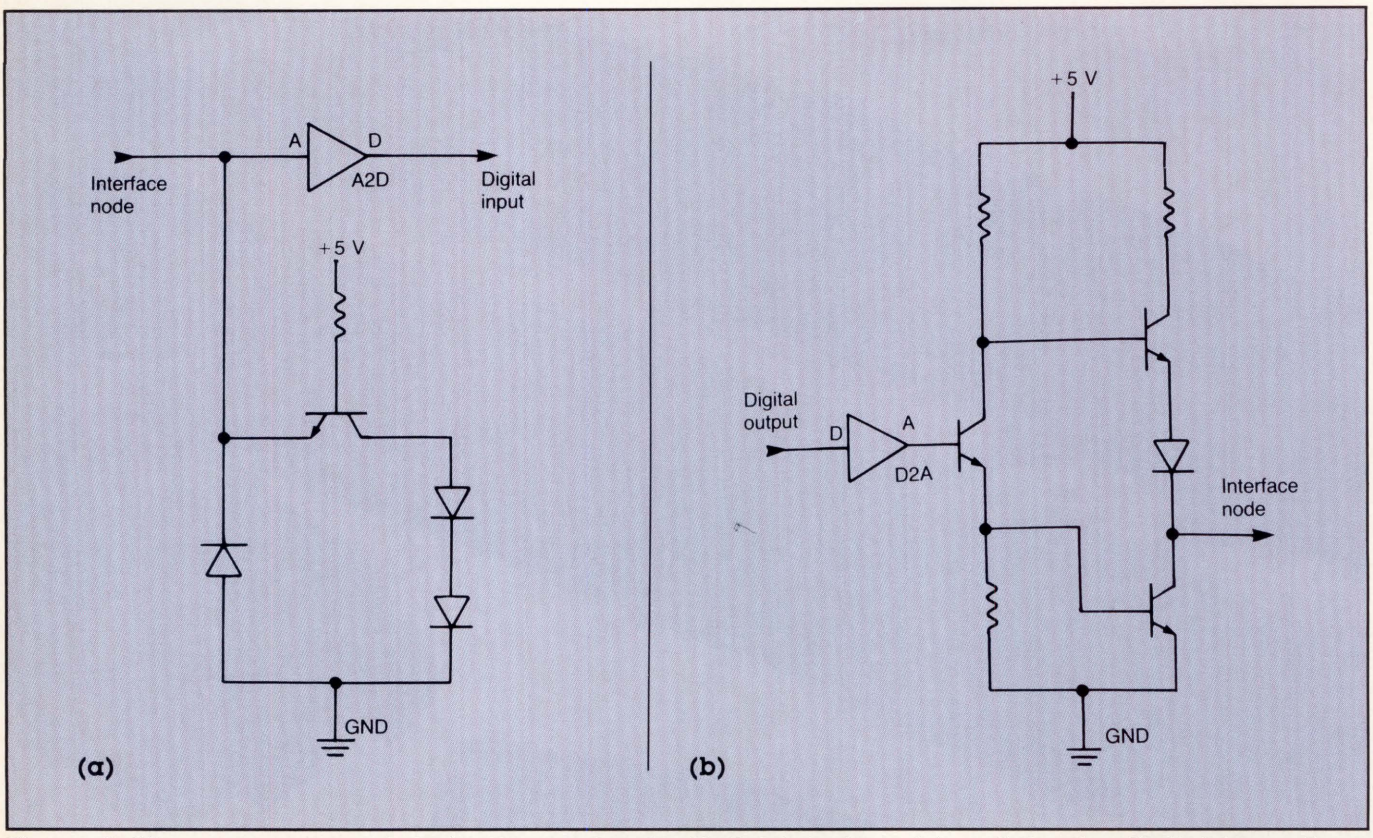

Figure 5. More accurate models for the analog-to-digital interface (a) and the the digital-to-analog interface (b).

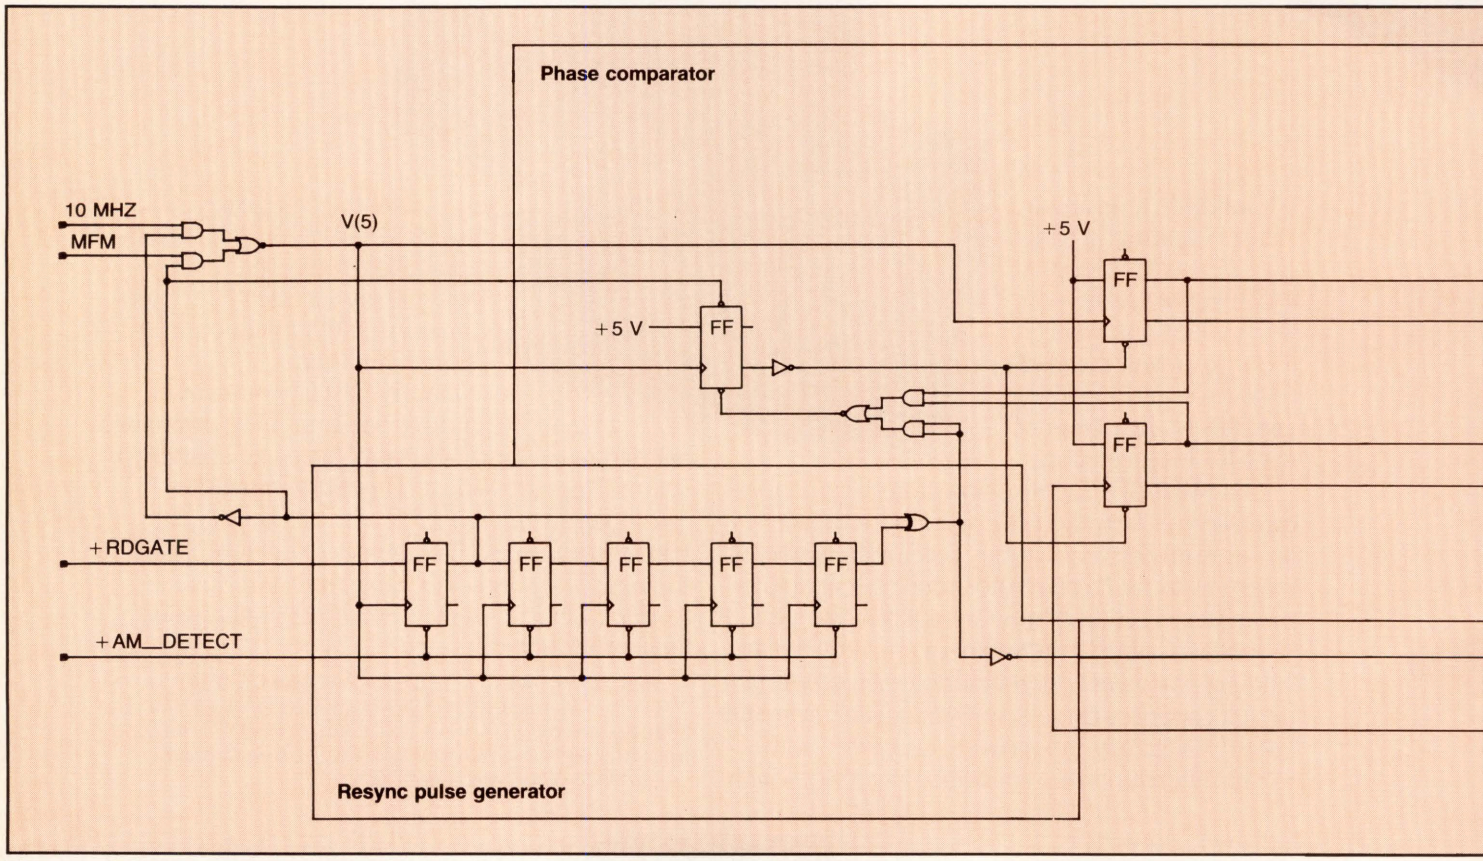

Figure 6. Data separator example for analog-digital simlation.

digital interfaces) as required by the system being simulated.

The model statement defines the characteristics of the digital output devices that refer to it. It defines the value of the load resistor and capacitor (parameters RLOAD and CLOAD) and the voltage range for each logic state (parameters SnVLO and SnVHI, where *n* identifies the paricular state). The SnNAME parameter identifies the VIEWSIM logic state of 1, 0, X, or Z. The user defines as many different model statements as he needs to define the different types of a/d interface characteristics in the system being simulated.

Usually, the user does not create device or model statements, because the netlister creates device statements and the models are part of supplied component libraries. However, a description of the parameters demonstrates the flexiblity of the interface model and the ease with which a user can change models to meet specific needs.

There are no delays associated with the digital output device. PSPICE maintains the current state in VIEWSIM for each digital output device at each VIEWSIM time point. (PSPICE only sends VIEWSIM the state of a digital output on changes.) However, some apparent delay can be caused by the fact that VIEWSIM represents time as multiples of a discrete time step and PSPICE uses real numbers to represent time. When a digital output device causes a change of state in a VIEWSIM device, the change is not seen until the next multiple of the VIEWSIM time step.

The Digital-to-Analog Model

The digital-to-analog interface device, called the digital input model (see Figure 2b), is modeled as a three-terminal device. One terminal connects to the positive digital power supply, another to the negative digital power supply (usually ground), and the third to the d/a interface node in PSPICE. The digital input is modeled as a time-varying resistor between the positive digital power supply node and the interface node and another time-varying resistor between the negative digital power supply node and the interface node. The values of the resistors, which depend on the state of the interface node in VIEWSIM, model both the voltage level and driving strength of the logic output. Both resistors have optional fixed-value capacitors in parallel with them (parameters CHI and CLO) that model loading effects.

When the digital signal state changes, the resistors change from their current values to values specified for the new state. The voltage at the interface node will change to a new steady-state value, based on new resistor values. This change is continuous, taking place over the time specified by the SnPSW parameters (switching time); it thus provides a smooth transition of both the driving voltage and the driving strength (resistance).

The device statement identifies three terminal connections, the model to be used, the name of the VIEWSIM net that controls the device state, and an optional initial state value. The user can specify the initial state so that the analog simulator can use an input state other than X (the VIEWSIM unknown logic state) for its bias point calculation. (VIEWSIM initializes all nodes to X when it starts.

The model statement defines resistor values (SnRLO and SnRHI) and switching time (SnTSW) to be associated with each VIEWSIM state (SnNAME). Again, the user does not normally

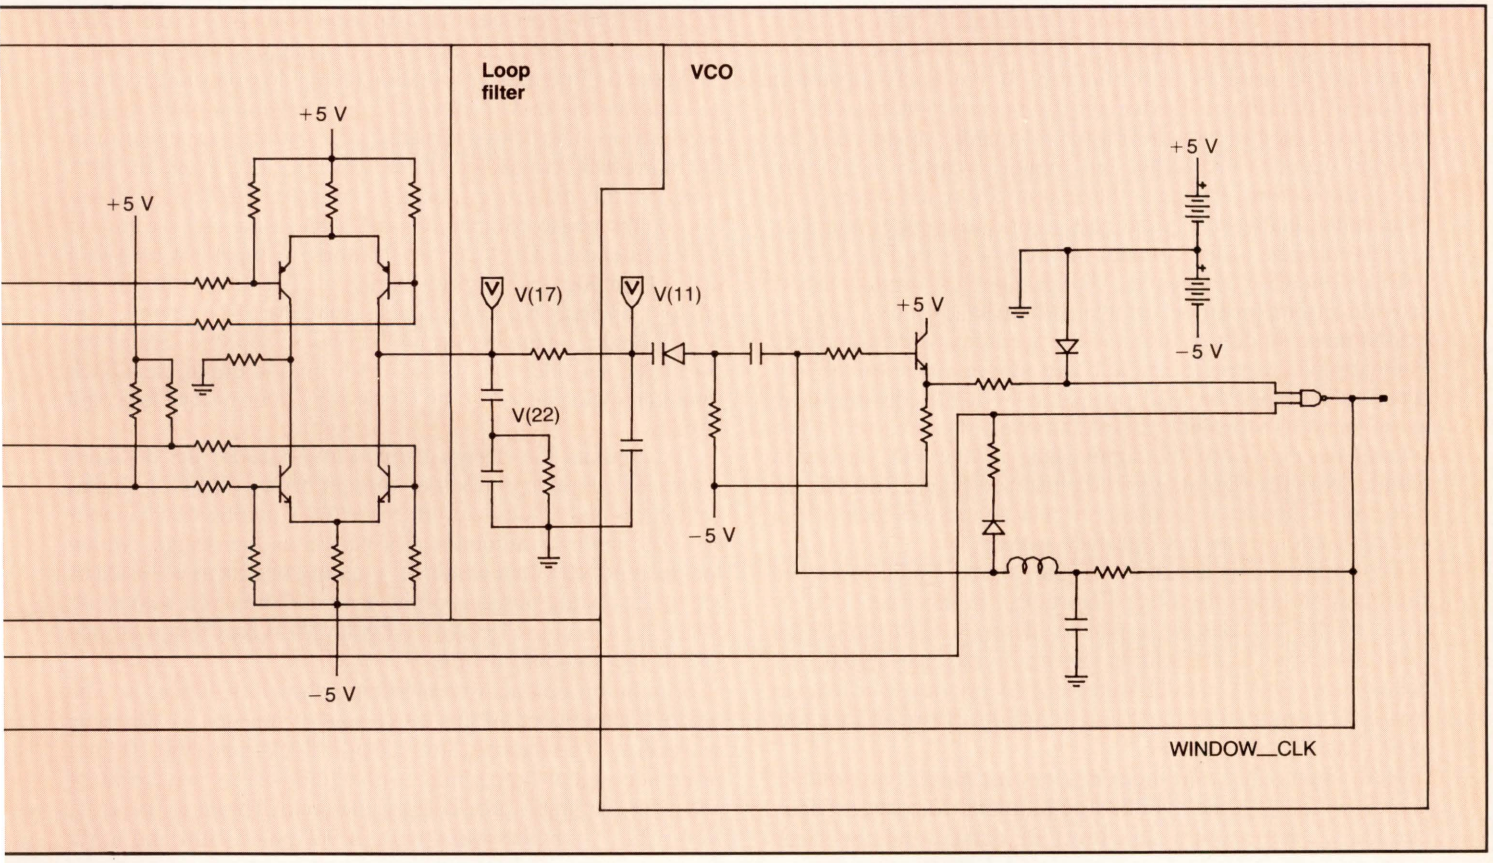

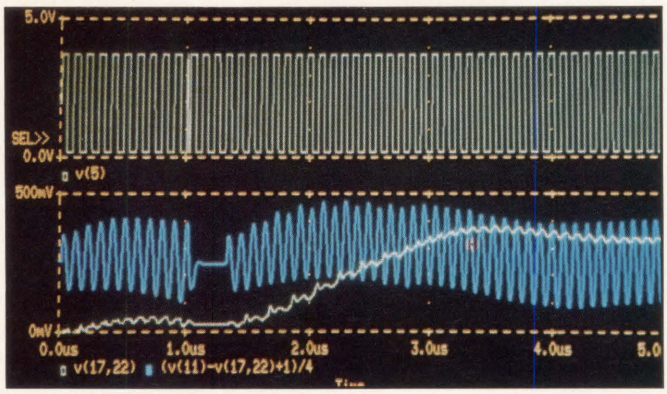

Figure 7. Output of the simulation of the data separator example.

need to read or write the PSPICE device or model statements. This description shows how the user can control the details of the interface.

Interface Considerations

When modeling digital component delays across the d/a interface, the analog and digital models must cooperate to break a signal delay into a propagation delay and a switching time. VIEWSIM models the propagation delay for components that drive digital-analog interfaces, and PSPICE models the switching time. Figure 3 illustrates the relationship between the the propagation delay of a digital component and the total delay, combining propagation delay and switching time, of an analog component.

This delay modeling consideration is especially significant when a digital output signal is connected to both a digital input and an analog component. One of two types of interfaces can be used in this situation. If the delay (loading) effects of the analog device on digital components can be neglected, then a simple digital-to-analog model can be inserted between the digital and analog circuitry (see Figure 4a). If the loading effects of the analog device are significant, then both analog-to-digital and digital-to-analog interface models must be used to model the interface (see Figure 4b).

Basic interface devices (the digital input and digital output) provide a simple model of the a/d interface. It is easy to extend these models to improve the accuracy of the interface. For example, a single resistor and single capacitor do not accurately model the load presented by a TTL device to an anlog input. To provide a better model, the user can place several devices from the input stage of the TTL device in parallel with the digital output device (see Figure 5a). He can place these devices in a subcircuit that he can use each time a critical analog-to-digital interface occurs.

To improve the accuracy of the digital input model, a similar approach is used (see Figure 5b). In this case, a driver stage of the TTL device is placed in series with the digital input. Because a TTL output driver inverts the signal, we change the PSPICE model statement to provide a low-level signal for a high logic state, and a high-level signal for a low logic state.

Both approaches add more devices for PSPICE to simulate and therefore degrade overall simulation performance. Many circuits do not require this level of detail, and the basic interface models in the digital input and digital output devices provide sufficiently accurate and faster simulation results.

Using VIEWSIM/AD

The VIEWSIM/AD simulation environment consists of five parts: schematic capture, stimulus creation, netlist extraction, simulation, and waveform analysis. The designer captures his schematic on low-cost, PC-based workstations using VIEWDRAW, Viewlogic's schematic capture package. Designers can freely interconnect analog and digital library components on their design schematics.

The user creates design stimuli in a conventional manner using ASCII SPICE decks and VIEWSIM command files. He can also create analog stimuli by adding parameterized signal-source symbols to a schematic. Optionally, he can create digital stimuli graphically using Viewlogic's VIEWWAVE waveform processor, or he can enter stimuli interactively during simulation.

Once schematic capture is complete, netlist extraction is performed using the VIEWSIM/AD netlister. The netlister analyzes the design topology and automatically partitions the design into analog and digital sections. As output, the netlister generates two files, a standard-format SPICE input deck and a VIEWSIM netlist description file.

When invoked, the simulator reads in the files created by the netlister and spawns the PSPICE and VIEWSIM subprocesses. The user is then placed into the VIEWSIM interactive simulation environment. He is now free to execute interactive simulation commands, execute simulation command files, generate simulation output files, and perform other design tasks. After a simulation session is completed, the user can analyze both digital and analog output waveforms by using the waveform processor. He can save or modify these waveforms and use them as input to subsequent simulation runs.

Figure 6 illustrates a typical mixed analog-digital system—a combination disk controller, data separator, and phase-locked loop. The design is patterned after the data separator in the hard-disk controller of an IBM PC XT. Four sections compose this data separator: the voltage-controlled oscillator (VCO), the phase comparator, the loop filter, and the resync pulse generator. The VCO and the phase comparator contain a mix of analog and digital circuitry. The loop filter is all analog, whereas the resync pulse generator is all digital. The VCO includes an analog-digital feedback loop that creates the oscillator, in addition to the main analog-digital feedback loop that locks the VCO onto clock transitions of the MFM data stream.

To simulate the design using VIEWSIM/AD, the designer represents digital portions of the circuit as VIEWSIM digital primitives and analog circuitry as PSPICE analog device models. The netlister automatically partitions the design, which the designer can then simulate, yielding results as accurate as PSPICE in analog portions of the design. Compared with using SPICE for the entire design, VIEWSIM/AD takes less than one fourth the simulation time.

Simulation output waveforms are shown in Figure 7. The analog component of the VCO signal appears as the blue waveform. The data separator output locks onto the 10-MHz reference signal (appearing in the light green waveform of V[5]) for the first 1 μs. At 1 μs, the control input to the VCO, V(17, 22), is stable. After 1 μs, a 90° phase shift and a 5% frequency change are introduced by gating the MFM signal

onto node 5, representing the "training sequence" of a new disk sector. At the same time, the +AM_DETECT signal is asserted to produce a resync pulse. The resync pulse stops the VCO for four cycles, restarting it in phase with the MFM signal. As shown by the control input to the VCO, the oscillation frequency changes during the period from about 1.4 μs to about 3 μs. Small VCO input changes continue until about 4.2 μs, after which phase errors are small. □

References

Bloom, M. 1987. "Mixed-Mode Simulators Bridge the Gap between Analog and Digital Design," *Computer Design,* January 1987.

Chawla, B.R., et al. 1975. "MOTIS—A MOS Timing Simulator," *IEEE Transactions on Circuits and Systems,* December 1975.

Chen, C.F., et al. 1984. "The Second Generation MOTIS Mixed-Mode Simulator," *21st Design Automation Conference,* Albuquerque, NM.

Friedman, M. 1987. "A Swifter Way to Simulate Analog and Digital ICs," *Electronic Products,* September 1987.

MicroSim Corp. 1987. *PSpice User Manual,* MicroSim Corp., Laguna Hills, CA.

Nagel, L.W. 1975. "SPICE2: A Computer Program to Simulate Semiconductor Circuits," Memorandum No. M520, University of California, Berkeley, CA.

Newton, A.R. 1978. "The Simulation of Large-Scale Integrated Circuits," Memorandum No. UCB/ERL M78/52, University of California, Berkeley, CA.

O'Rouke, B. 1988. "Behavioral Modeling of Digital Devices in an Analog Simulation Environment," *VLSI Systems Design,* January 1988.

Sarin, H., et al. 1987. "Simulator Environment Handles Mixed Designs," *Computer Design,* January 1987.

Viewlogic Systems Inc. 1987. *Workview Manual Set, V 2.1,* vol. 3, Viewlogic Systems Inc., Marlboro, MA.

About the Authors

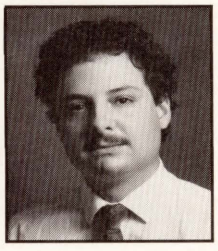

Tedd Corman, a Viewlogic employee since 1985, is a senior engineer within the digital simulation group and principal engineer on the VIEWSIM/AD project. Previously, he was a member of the technical staff at ITT Avionics in Nutley, NJ, where his primary responsibility was the evaluation and integration of CAE tools. He holds a BS in computer science from Rensselaer Polytechnic Institute, Troy, NY.

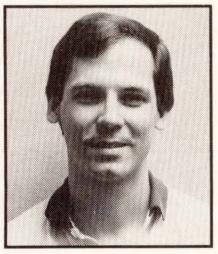

Michael U. Wimbrow, a software engineer at MicroSim, is responsible for mixed analog-digital simulation projects. He developed the techniques used to synchronize PSPICE with an event-driven logic simulator. Prior to joining MicroSim, he spent 10 years at Silicon Systems Inc., where he designed bipolar and MOS ICs and headed the company's CAD group. He received a BSEE from the California Institute of Technology in Pasadena.

Continued from page 29

of all of the pixels in the object to be displayed. Five parallel paths, each with four image engines, operate simultanaeously on a distributed frame memory. In effect, 20 independent processors update the display memory simultaneously, according to the position, color, and transparency of the new object. Similarly, five multimode graphics processors determine the color modes (which depend on the windowing environment) for the pixels streaming from the raster subsystem.

The new Tektronix series has a similar pipeline, with a transform engine based on the ADSP3221 from Analog Devices Inc. (Norwood, MA), which is analogous to the geometry engine. Tektronix claims that for many applications, using one processor instead of a pipeline can be as fast or faster. For example, for moving within a representation of a 2D drawing, few vectors are clipped, making the transform engine very efficient. Similarly, if no transformations are necessary at all for a command, the command proceeds more quickly down the drawing processor. Such commands, however, can hardly be considered the bottlenecks to real-time display.

Raster Technologies Inc. (North Billerica, MA) places the floating-point engines in its GX4000 series of workstations in parallel (Torborg, 1987), instead of in series as the Iris architecture does. The company promotes a parallel architecture for front-end processing that is more adaptable to different algorithms, particularly for complex algorithms like surface tesselation and lighting calculations. Through queue management, the algorithmic tasks are evenly distributed among the GX4000's four processors, optimizing the throughput regardless of the sequentiality of a given algorithm. Because the system offloads drawing operations to a drawing processor, it makes the processing time of the commands independent of the number of pixels that are modified. Such uniformity of command processing enhances the efficiency of the parallel processing.

At the systems level, Pixar has integrated three of the boards on the Image Computer into one with gate arrays. Using an analog multiplexer board, two systems can work in parallel, a feat conducive to animation (one system manipulates the background;the other creates the foreground or main characters). Such parallelism multiplies the parallelism among and within CHAPs.

Additional parallelism results if the memory devices themselves can perform operations on internal data. For example, NEC's μPD42232, discussed earlier, has a logical operation unit that performs operations on raster data (raster ops). Through this circuitry, logic operations and BITBLT transfers can occur on chip. □

References

Budge, D. 1988."Designing High-Speed VLSI: The Hardware-Software Partnership," *VLSI Systems Design,* January 1988.

Eidson, S., et al. 1987. "Image Processing with Vector Processors," *Wescon/87,* San Francisco, CA.

Shugard, D., et al. 1987. "A CMOS VLSI Chip for Real-Time Centroid Calculation," *IEEE 1987 Custom Integrated Circuits Conference,* Portland, OR.

Sun, E., and J. Cates. 1987. "High Performance 3-D Graphics Processing," *Wescon/87,* San Francisco, CA.

Torborg, J. 1987. "A Parallel Processor Architecture for Graphics Arithmetic Operations," *Siggraph '87,* Anaheim, CA.

National Semiconductor

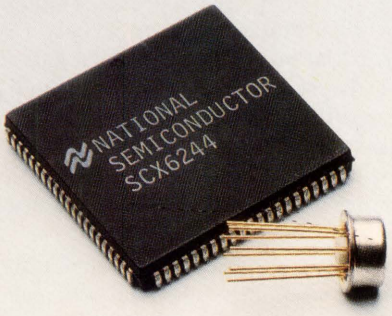

Digital and analog functions in standard packages

Similar functions in a standard cell available for packaging

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

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

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

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

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

EIGHT CELLS YOU CAN BUILD ON

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

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

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

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

NATIONAL'S CELL LIBRARY INCLUDES:

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

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

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

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

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

PROVEN CELLS YOU CAN WORK WITH

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

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

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

FUTURE CELLS YOU CAN COUNT ON

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

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

CELL EXPERTS YOU CAN TALK TO

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

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

National Semiconductor

1988 Survey of Logic Simulators

VLSI Systems Design Staff

Updating our 1987 survey of logic simulators, we are surprised to find a relatively stable market. Five companies have departed from the survey: Aptos, Calma, Clarity, Data General, and Rune Software. In turn, five companies are now reporting on products not previously surveyed: ECAD, Gateway Design Automation, OrCAD, Roche Systems, and Silvar Lisco. Pricewise, things are much the same, although Teradyne has almost halved the price of LASAR and Silvar Lisco has doubled the price of HELIX. On the platform front, AIDA and Valid now support the Sun; Daisy has added the Personal Logician 386 to its line; and the CAD Group now supports Apollo, Alliant, and Cray and Cyber supercomputers. In addition, VLSI Technology offers a link to the Mach 1000.

At the front end, Cadnetix now supplies graphical waveform entry, and Matra has added OrCAD schematic entry to its system. Although there have been some general alterations to the simulators' basic capabilities, we particularly note that Silicon Compiler Systems (which now offers Silicon Design Labs' LSIM timing simulator) reports that LSIM supports 4096 states instead of 64. AIDA and Simulog, the two compilation-based simulators in the survey, no longer report on switch-level simulation capabilities.

Significant developments have occurred in the modeling arena, with Daisy adding 2500 models to last year's library of 1500 parts, Mentor offering 1400 behavioral models and support of Logic Automation models, and Silvar Lisco expanding its library to 4000 parts.

ATE interfaces remain sparse. We do note that Mentor and Valid have significantly expanded their ATE interfaces. Also, Viewlogic now employs Test Systems Strategies Inc.'s services for tester interfaces.

Directory Definition

This survey does not reproduce the complete listing of CAE vendors that remarket these simulators under different user interfaces—that listing may be found in the 1988 *User's Guide to Design Automation* and the June 1988 issue. A particular simulator may appear more than once here, but only in those cases where the vendor provides a simulator-specific hardware accelerator (or has significantly modified the original simulator's kernel).

Besides the self-explanatory listings of hosts, pricing, input method, fault simulation, and ATE interfaces, the directory lists information about features that bear further explanation. In the column "Simulation level," an **S**, **G**, or **B** indicates that the corresponding product supports simulation at the switch, gate, or behavioral level, respectively. In this directory, behavioral-level simulation reflects the ability to describe a design (or component) in abstract terms using some form of hardware description language.

For features, a bullet (●) indicates availability, whereas an open circle (○) indicates its absence. In the column "Operating modes," the directory describes the simulator's ability to support incremental compilation. Following a simulation run, a user may modify a portion of a design, recompile only that redesigned portion, and simulate the complete (but modified) design without needing to recompile the entire design. Using a simulator with a save/restart capability, after he has compiled and loaded a design, a user may save the current simulation image in a host data file for later resumption of the simulation.

Image patching indicates the ability to modify, or "patch," the executable image and resimulate without a need to recompile or reload the design into the simulator. In some cases, a simulator may not support complete patching, but it may permit a more restricted variation in which an engineer may break selected network connections or force a node to a specific value. Alternatively, an interactive simulator may inherently provide this and the other listed operating modes, because even though the simulator requires netlist compilation, it may work directly from the design database.

The column "Signal states" lists the number of supported logic levels (such as 0, 1, and unknown) and signal strengths (such as driving, resistive, and high-impedance).

Under the heading "Timing support," the directory provides information about various critical timing-support features of logic simulators. Here, the "Propagation delay" column describes whether the simulator supports a different signal propagation delay (through devices or wire), depending on the direction in which the logic level is changing.

Similarly, the column "Transition times" lists the specific ability of a simulator to support skews of different durations, depending on which direction the logic level is changing.

The simulator's ability to support minimum, maximum, and nominal timing is noted in a separate column.

The final column in this group describes the ability to provide warning of setup, hold, and pulse width timing violations for gate-level and behavioral simulations.

The directory also lists the number of unique elements supported as simulation primitives (and thus directly recognized by the simulator). Separate columns indicate the number of models available through libraries, as well as the availability of source code—defined here as a complete description expressed with one or more of the circuit description methods listed for the product. In many cases, the vendor provides a user interface that includes extensive model development support.

Directory of Logic Simulators

Vendor and contact	Simulator and cost	Host	Simulation level	Operating modes: Incr. compile	Operating modes: Save/restart	Operating modes: Image patch	Signal states: Logic levels	Signal states: Signal strengths	Timing support: Prop. delays	Timing support: Transition times	Timing support: Min/max/nom	Timing support: Violations	Primitives and models: No. of primitives	Primitives and models: No. of models	Primitives and models: Source provided	Circuit description method	Fault simulation	ATE interfaces
AIDA Corp. 5155 Old Ironsides Dr. Santa Clara, CA 95054 Georgia Marszalek Director of Marketing Communications (408) 748-8571	**AIDA Verification System** $50k	Any Apollo Domain workstation; Sun-3	G B	● ●	● ●	○ ○	4 4	4 4	● ●	● ●	●[1] ●[1]	● ●	>41 >41	●[2] ●[2]	○ ○	Mentor Graphics' NETED; TEGAS netlists; AIDA Schematic Design Editor; C-language behavioral descriptions; AIDA PG-language functional descriptions; cosimulator instruction-set operators	AIDA Fault Simulator	None
The CAD Group Inc. 3911 Portola Dr. Santa Cruz, CA 95062 Vinnie Apicella Director of Marketing (408) 475-5800	**SALT** $3.5k (IBM PC XT, AT); $15k (workstations); $50k (VAX); $100k+ (mainframes and supercomputers)	IBM PC XT, AT; VAX (UNIX, VMS); Ridge; Sun; Apollo; Alliant; Cray; Cyber	S G B	● ● ●	● ● ●	○[3] ○[3] ○[3]	6 6 6	4 4 4	● ● ●	● ● ●	● ● ●	N.a. ●[4] ●[4]	5 96 100+	0 2000 100	● ● ○	Case, CAECO, OrCAD schematic entry; network description language; register-transfer language	Probablistic fault grading	Sentry series
Cadnetix Corp. 5757 Central Ave. Boulder, CO 80301 Greg Skomp Manager, Marketing Communications (303) 444-8075	**CDX** $15k (simulator); $4k (waveform editor for Sun); $2.9k (waveform editor for IBM PC AT)	CDX-9510, 9610 (Sun-based workstations); CDX-70000 (Configurable Analysis Engine with digital and accelerated digital simulation); CDX-56010SP (Sun-based workstation with digital simulation and CAD)	S G[6]	● ●	○ ○	○ ○	3[5] 3[5]	4 4	● ●	● ●	● ●	N.a. ●	25 110	N.a. 629[7]	● ●	Cadnetix hierarchical schematic entry	None	GenRad 2270; Zehntel; Factron
Control Data Corp. 8100 34th Ave. South Minneapolis, MN 55440 Robert Biggs Marketing Manager (612) 853-3117	**MIDAS** $35k–$740k	Cyber 176, Cyber 800 series	G B	● ●	● ●	● ●	4 4	8 8	● ●	● ●	● ●	● ●	N.a. 0	262 ●[8]	○ N.a.	Graphic schematic entry	Parallel detection of stuck-at faults; Zycad FE	Sentry; Teradyne; GenRad; Takeda-Riken; neutral tester format for gate level

N.a. = not applicable.
1. Plus manufacturing tolerances.
2. Number varies with supported manufacturer's library.
3. Not recommended.
4. Special timing models are available to automatically monitor timing violations.
5. 21-state simulator supports 15 internal levels for unknown.
6. Also (hierarchical) functional.
7. Plus 1258 functional models.
8. User-defined.

The GE/RCA ASIC cell library:

where best-sellers begin.

With one of the largest standard cell and gate array libraries, and a first-time success rate of virtually 100%, GE Solid State has all the tools you need to design leading-edge products.

Our library of hundreds of verified cells and functions is one of the broadest in the industry.

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

You can choose from a wide range of Supercells, including core micros, RAMs, analog functions and bit-slice processors. Right now, we're working on new editions, including Advanced CMOS Logic, EPROMs and high-voltage cells. Additionally, we can generate custom cells to meet your requirements.

So you can easily reduce an entire system onto a single ASIC chip, and get the ASIC benefits of smaller power supplies, lower system cost and higher reliability.

Plus a full library of gate arrays.

If your design calls for gate arrays, you can select from a broad line with up to 13,000 gates. With alternate sourcing on the LSI/Toshiba series.

Our gate arrays have effective gate lengths as small as 1.2μ. And they're compatible with the industry standard.

Writing made easy.

We've made it easy for you to write your best-selling design.

Our designer-friendly software gets you up to speed in as little as the three days it takes to complete our training course.

And our software supports the major workstations in the industry, including Daisy,™ Mentor Graphics,™ Valid,™ P-CAD® and FutureNet.®

Through the use of Twins, our transparent workstation interface system, you can complete or revise your design at practically any workstation.

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.

The best editors in the business.

We know how important accurate performance modeling is to you. That's why we've developed the industry's most advanced simulation and parasitic extraction programs, which make it easy to create designs that work. In fact, our first-time success rate is virtually 100%.

And you can count on our ASICs specialists for technical support from design through production. One other thing you can count on with GE Solid State: there are no hidden charges in addition to your NRE.

On the best-seller list for years.

We've been the CMOS company since the day we invented the technology. We've produced one best-selling CMOS circuit after another. And built the industry's broadest line of High-Rel CMOS parts, including super radiation-hardened Silicon-on-Sapphire.

With all this behind you, think how easy it could be to design a best-seller.

For more information, call toll-free 800-443-7364, ext. 14. Or contact your local GE Solid State sales office or distributor.

In Europe, call: Brussels, (2) 246-21-11; Paris, (1) 39-46-57-99; London, 0276-685911; Milano, (2) 82-291; Munich, (89) 63813-0.

GE Solid State

GE/RCA/Intersil Semiconductors
These three leading brands are now one leading-edge company.
Together, we have the resources—and the commitment—
to help you conquer new worlds.

Vendor and contact	Simulator and cost	Host	Simulation level	Operating modes: Incr. compile	Operating modes: Save/restart	Operating modes: Image patch	Signal states: Logic levels	Signal states: Signal strengths	Timing support: Prop. delays	Timing support: Transition times	Timing support: Min/max/nom	Timing support: Violations	Primitives and models: No. of primitives	Primitives and models: No. of models	Primitives and models: Source provided	Circuit description method	Fault simulation	ATE interfaces
Daisy Systems Corp. 700 Middlefield Rd. Mountain View, CA 94039-7006 Nancy Morrison Public Relations Manager (415) 960-6591	**DLS, Daisy Logic Simulator; DTV, Daisy Timing Verifier; MDLS accelerated logic simulator** $10k–$15k	Personal Logician 386; Compaq 386; IBM PC AT; Logician 386; MegaLogician hardware accelerator	S G B	● ● ●	● ● ●	●[9] ●[9] ●[9]	3[10] 3[10] 3[10]	4 4 4	● ● ●	● ● ●	● ● ●	N.a. ● ●	2 30 44	●[11] 4500 ●[11]	N.a. ● N.a.	Schematic editor; Daisy behavioral language; EDIF netlist	Mega-FAULT, hardware-accel-erated concurrent fault simulation	Sentry; GenRad GR160, GR180; Factron 700; IMS; general-purpose ATE interface utility package
ECAD Inc. 2455 Augustine Drive Santa Clara, CA 95054 John Miles Marketing Manager (408) 727-0264	**ELLA** $30k–$150k	VAX (VMS, ULTRIX); Sun; Apollo	G B	● ●	● ●	○ ○	●[11] ●[11]	●[11] ●[11]	● ●	● ●	● ●	● ●	●[11] ●[11]	●[11] ●[11]	N.a. N.a.	ELLA language	None	None
Endot Inc. 11001 Cedar Ave. Cleveland, OH 44106 Michael Radovich Public Relations Manager (206) 881-6444	**N.2 Systems Architects Workbench** $30k–$250k	VAX; IBM; Apollo; Sun	G B	● ●	● ●	● ●	4 4	2 2	● ●	● ●	● ●	● ●	●[11] ●[11]	●[11] ●[11]	N.a. N.a.	ISP′ HDL; N.2 network description language; VHDL output from ISP′	None	None
Gateway Design Automation Corp. PO Box 573 6 Lyberty Way Westford, MA 01886 Pete Johnson Marketing Manager (617) 692-9400	**VERILOG/ VERILOG-XL** $25k–$35k (workstations); $50k–$200k (mainframes)	VAX; Apollo; Sun; IBM 43xx, 30xx; Masscomp; Elxsi; Silicon Graphics	S G B	○ ○ ○	● ● ●	● ● ●	6 6 4	8 8 N.a.	● ● ●	● ● ●	○ ○ ○	N.a. ●[13] ●[13]	12[12] 12[12] N.a.	0 6500 20	N.a. ● ●	VERILOG language; translators from Mentor, Daisy, Valid, CAECO schematics	Interface to TestGrade	TSSI interfaces
GenRad Inc. 510 Cottonwood Dr. Milpitas, CA 95035 Peter Denyer Product Marketing Manager (408) 432-1000	**HILO-3** $14k (Opus-516); $103.5k (VAX-11/780)	VAX (UNIX, ULTRIX, VMS); IBM (VM/CMS); Apollo; Sun; HP9000; Opus; Ridge; Harris HCX-7; IBM PC RT	S G B	● ● ●	● ● ●	●[14] ● ○	3 3 3	3[15] 3[15] 1	● ● ●	● ● N.a.	● ● ●	N.a. ● ●	4 22[18] N.a.	●[16] ●[16] ●[16]	●[17] ●[17] ●[17]	Schematic entry; HILO HDL	Parallel value list	GenRad GR16, GR18, 227X series, GR160, GR180, GR125; TSSI interfaces

N.a. = not applicable.
9. Timing delays; users can interactively set node values.
10. 6 in DTV.
11. User-defined.
12. User can also define combinational and sequential primitives.
13. User can define assertions to test for any desired condition.
14. Nontopological changes only.
15. 15-state simulator also uses additional levels to represent stored-charge behavior.
16. Total of about 6000 modeled at switch, gate, or behavioral levels as appropriate, not including gate array or standard-cell libraries.
17. Some gate array and standard-cell libraries are supplied without source code at the manufacturer's request.
18. Includes 8 interconnect (wire) elements.

Vendor and contact	Simulator and cost	Host	Simulation level	Operating modes: Incr. compile	Operating modes: Save/restart	Operating modes: Image patch	Signal states: Logic levels	Signal states: Signal strengths	Timing support: Prop. delays	Timing support: Transition times	Timing support: Min/max/nom	Timing support: Violations	Primitives and models: No. of primitives	Primitives and models: No. of models	Primitives and models: Source provided	Circuit description method	Fault simulation	ATE interfaces
HHB Systems 1000 Wyckoff Ave. Mahwah, NJ 07430 Larry Blessman Product Marketing Manager (201) 848-8000	**CADAT 6** $3k (IBM PC, PC-DOS); $20k (workstations); $80k (VAX-11/780)	VAX (VMS, UNIX), MicroVAX; Sun; Apollo; IBM PC (PC-DOS and UNIX coprocessor); Data General; Sumitomo U-Station; HP-9000; Masscomp	S G B	● ● ●	● ● ●	○ ○ ○	3 3 3	4[19] 4[19] 4[19]	● ● ●	● ● ●	● ● ●	N.a. ● ●	27 31[20] ○[11]	200 2500 100[21]	● ● ●	CADAT circuit description language; behavioral modeling language (BML)	Concurrent functional; also used in CATS accelera-tor and CATS hardware modeler	Eaton Model 800; Factron and Sentry testers through CADDIF; GenRad; Teradyne L200
Ikos Systems Inc. 145 North Wolfe Rd. Sunnyvale, CA 94086 Mick Westhoff Vice President, Marketing and Sales (408) 245-1900	**Ikos 800/1900 hardware accelerator** $40k	IBM PC AT; Apollo Series 3000, 4000, DN570; Sun-3	S G	● ●	● ●	● ●	3 3	8 8	● ●	● ●	● ●	N.a. ●	256 256	●[22] ●[22]	● ●	Schematic capture through supported workstations	Parallel	None
LSI Logic Inc. 1551 McCarthy Blvd. Milpitas, CA 95035 Van Lewing Director of Software Marketing (408) 433-4546	**LSIM single-chip logic simulator; MSIM multi-chip simulator; ZSIM accelerated simulator; FSIM fault simulator; BSIM behavioral simulator** $80k (Sun-3)	IBM (VM/CMS); VAX, MicroVAX (VMS); Sun-3; Apollo Series 3000, DN570, DN580 (DOMAIN IX)	G B[24]	● ●	● ●	○ ○	3 3	4 4	○ ○	● ●	● ●	● ●	●[23] ●[23]	●[23] ●[23]	●[23] ●[23]	NDL (netlist description language); LSED graphic schematic editor	Through FSIM, serial (Zycad LE) and concurrent (Zycad FE)	Sentry; Ando
Matra Design Semiconductor 2895 Northwestern Pkwy. Santa Clara, CA 95051 Pradip Madan Vice President, Marketing and Sales (408) 986-9000	**ARCIS logic simulator; WAVE waveform analyzer; ARCOP testabili-ty analyzer** $945	IBM PC, XT, AT, 386, and compatibles	G	●	●	○	6	4	●	○	●	●	10	150	○	OrCAD-compatible graphic schematic entry and block description; netlist entry; PLD Boolean equation entry	Concurrent	Sentry; IMS

N.a. = not applicable.
19: Plus 12 indeterminate states, for a total of 21 states.
20. Plus 43 functional models.
21. Plus 65 hardware behavioral, plus Engineering Information Systems (Santa Clara, CA) and Quadtree.
22. Number varies with supported manufacturer's library.
23. Depends on particular technology among the several technologies supported.
24. BSIM only.

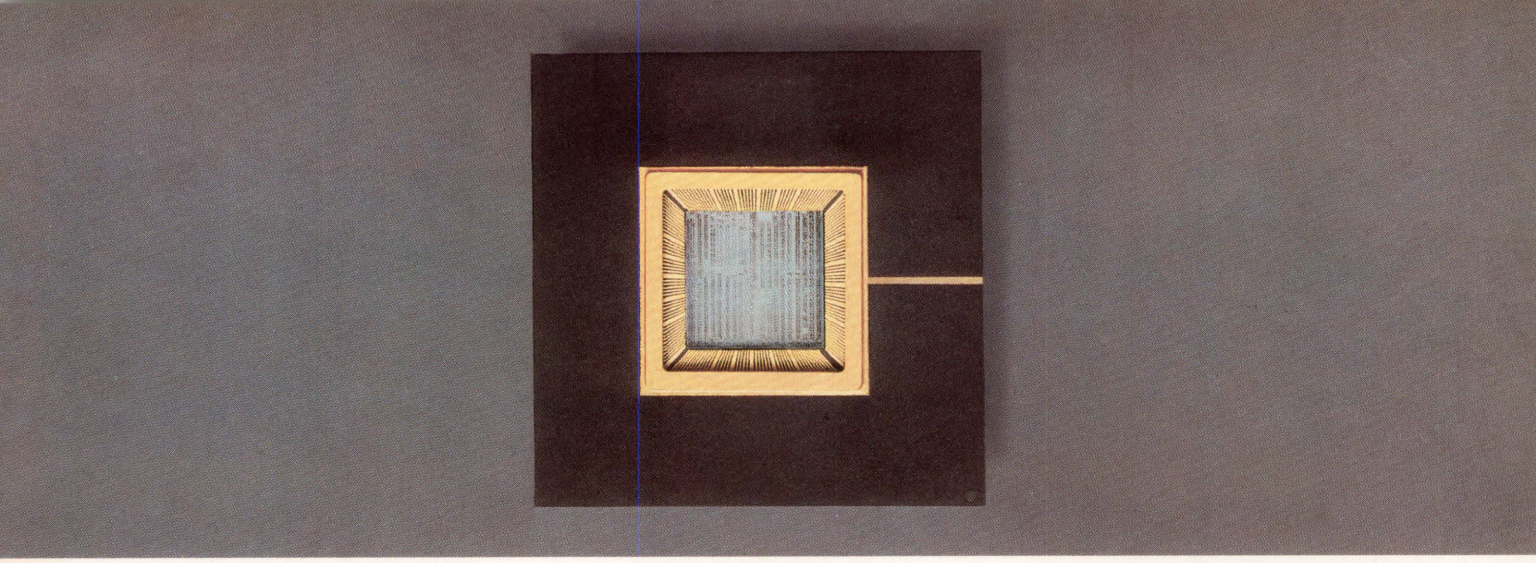

50,000 Gates At 0.7 Nanoseconds.

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

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

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

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

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

TOSHIBA. THE POWER IN GATE ARRAYS.				
	3 MICRON CHANNELLED ARRAY	2 MICRON CHANNELLED ARRAY	1.5 MICRON CHANNELLED ARRAY	SEA OF GATES COMPACTED ARRAY™
SERIES	TC15G	TC17G	TC19G	TC110G
GATES	880 to 6,000	540 to 10,000	3,200 to 10,000	2,100 to 50,000
DESIGN RULE	3μm	2μm	1.5μm	1.5μm
GATE SPEED	2.5ns	1.5ns	1.0ns	0.7ns
PART NO's	6	9	5	5
AVAILABILITY	NOW	NOW	NOW	NOW
All Si-Gate CMOS double layer metal. Compacted Array is a trademark of LSI Logic Corporation.				

AREA SALES OFFICES: CENTRAL AREA, Toshiba America, Inc., (312) 945-1500; **EASTERN AREA,** Toshiba America, Inc., (617) 272-4352; **NORTHWESTERN AREA,** Toshiba America, Inc., (408) 244-4070; **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, POUGHKEEPSIE, NEW YORK,** Toshiba America, Inc., (914) 462-5710; **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, (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,** C.H. Horn, (319) 393-8703; **KANSAS,** D.L.E. Electronics, (316) 744-1229; **KENTUCKY,** Leslie M. DeVoe Company, (317) 842-3245; **LOUISIANA,** MIL-REP Associates, (713) 444-2557; **MAINE,** Datcom, Inc., (617) 891-4600; **MASSACHUSETTS,** Datcom, Inc., (617) 891-4600; **MICHIGAN,** Action Components Sales, (313) 349-3940;

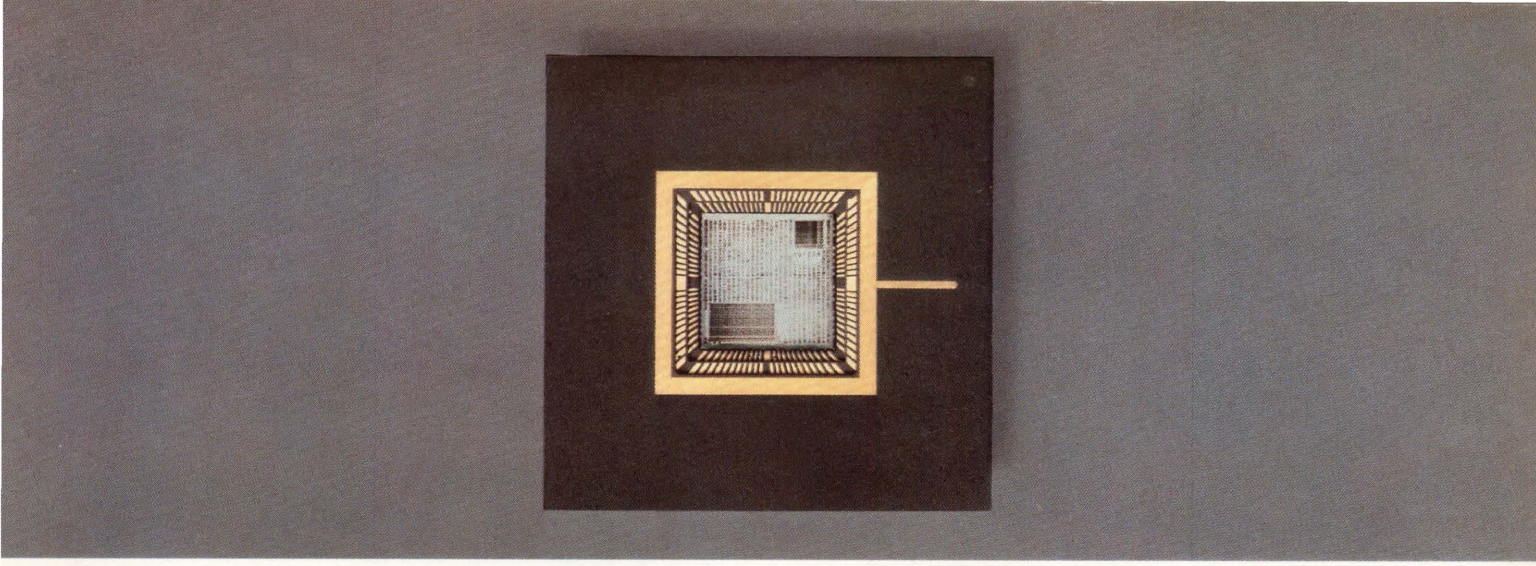

RAM & ROM Memory In CMOS ASICs.

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

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

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

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

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

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

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

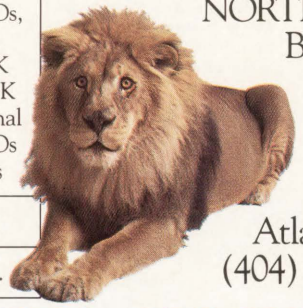

TOSHIBA AMERICA, INC.

MINNESOTA, Electric Component Sales, (612) 933-2594; **MISSISSIPPI, Montgomery Marketing, Inc., (205) 830-0498; MISSOURI,** D.L.E. Electronics, (316) 744-1229; **MONTANA,** Components West, (206) 885-5880; **NEVADA,** Elrepco, Inc., (415) 962-0660; **NEBRASKA,** D.L.E. Electronics, (316) 744-1229; **NEW ENGLAND,** Datcom, Inc., (617) 891-4600; **NEW HAMPSHIRE,** Datcom, Inc., (617) 891-4600; **NEW JERSEY,** Nexus-Technology, (201) 947-0151; **NEW MEXICO,** Summit Sales, (602) 998-4850; **NEW YORK,** Nexus Technology, (201) 947-0151; Pi-tronics, (315) 455-7346; **NORTH CAROLINA/SOUTH CAROLINA,** Montgomery Marketing, Inc., (919) 467-6319; **NORTH DAKOTA/SOUTH DAKOTA,** Electric Component Sales, (612) 933-2594; **OHIO,** Steffen & Associates, (216) 461-8333; (419) 884-2313, (513) 293-3145; **OKLAHOMA,** MIL-REP Associates, (214) 644-6731; **OREGON,** Components West, (503) 684-1671; **PENNSYLVANIA,** Nexus Technology, (215) 675-9600, Steffen & Associates, (412) 276-7366; **RHODE ISLAND,** Datcom, Inc., (617) 891-4600; **TENNESSEE,** Montgomery Marketing, Inc., (205) 830-0498; **TEXAS,** MIL-REP Associates, (512) 346-6331, (713) 444-2557, (214) 644-6731; **UTAH,** Straube Associates Mountain States, Inc., (801) 263-2640; **VERMONT,** Datcom, Inc., (617) 891-4600; **WEST VIRGINIA,** Steffen & Associates, (419) 884-2313; **WASHINGTON,** Components West, (206) 885-5880, (509) 922-2412; **WISCONSIN,** Carlson Electronics, (414) 476-2790, Electric Component Sales, (612) 933-2594; **WYOMING,** Straube Associates Mountain States, Inc., (303) 426-0890; **CANADA, BRITISH COLUMBIA,** Components West, (206) 885-5880; **ONTARIO,** Electro Source, Inc., (416) 675-4490, (613) 726-1452.

CIRCLE NUMBER 12

Directory of Logic Simulators (continued)

Vendor and contact	Simulator and cost	Host	Simulation level	Operating modes: Incr. compile	Operating modes: Save/restart	Operating modes: Image patch	Signal states: Logic levels	Signal states: Signal strengths	Timing support: Prop. delays	Timing support: Transition times	Timing support: Min/max/nom	Timing support: Violations	Primitives and models: No. of primitives	Primitives and models: No. of models	Primitives and models: Source provided	Circuit description method	Fault simulation	ATE interfaces
Mentor Graphics Corp. 8500 SW Creekside Place Beaverton, OR 97005 Mohan Nair Marketing Manager (503) 626-7000	**QuickSim; QuickFault; Compute Engine hardware accelerator** $33.9k (workstation)	Apollo; Mentor Graphics' Compute Engine	S G B	● ● ●	● ● ●	● ● ●	3 3 3	4 4 4	● ● ●	● ● ●	● ● ●	N.a. ● ●	13 16 ●[25]	0 300 1400[26]	N.a. ● ●	Schematic entry; Pascal-like language; C-like language	Concurrent; accelerated concurrent through local-area network	Sentry; GenRad; Takeda-Riken; Tektronix; Teradyne; Marconi; Factron; SGS; HP; IMS; Advantest; Trillium; Hi-level
OrCAD Systems Corp. 1049 SW Baseline St., Suite 500 Hillsboro, OR 97123 Ken Seymour Vice President (503) 640-5007	**OrCAD/VST** $995	IBM PC, AT, XT, PS/2, and compatibles	G	○	●[27]	○	3	4	●	●	●	●	35	N.a.	●[28]	Schematic entry from OrCAD/SDT III	None	None
Personal CAD Systems Inc. 1290 Parkmoor Ave. San Jose, CA 95126 Tracy Kahl Product Manager (408) 971-1300	**PC-LOGS-II** $4k (includes schematic capture, net-list extraction, and other utilities); $2.5k as upgrade for existing P-CAD owners	IBM PC XT, AT, and compatibles; IBM 5550; TI Professional; NEC	S G B	● ● ●	● ● ●	● ● ●	3 3 3	4 4 4	● ● ●	● ● ●	○ ○ ○	N.a. ○ ○	3 27 25	●[22] ●[22] ●[22]	● ● ●	Graphic entry; PC-Model HDL	None	IMS
Roche Systems Corp. 1705 N. Rankin St. Appleton, WI 54911 Les Roche President (414) 733-6077	**DSIM logic simulator** $995–$2.5k (IBM PC); $15k (Apollo)	IBM P(XT, AT; Apollo	S G	○ ○	○ ○	○ ○	3 3	5 5	● ●	● ●	○ ○	● ●	4 14	●[29] ●[29]	● ●	DSIM macro language supporting user-defined nested modules; primitive net-list	Concurrent (3Q88)	None
Silicon Compiler Systems Corp. 2045 Hamilton Ave. San Jose, CA 95125 Peter Odryna Product Marketing Manager (201) 580-0102	**LSIM mixed-mode analog and digital simulator** $49.5k; $88k (with GDT silicon compiler tools)	Sun; Apollo; DEC VAXstation II, GPX	S[30] G B	● ● ●	● ● ●	● ● ●	3[31] 3 3	4096 4096 4096	● ● ●	● ● ●	● ● ●	N.a. ● ●	10[32] 28[32] ●[32]	32[32] 28[32] 36[32]	● ● ●	Netlist; layout; schematic; M modeling language; SPICE; EDIF	Serial probabilistic	Sentry; Takeda-Riken; IMS

N.a. = not applicable.
25. User-defined.
26. Includes VLSI models from Logic Automation Inc.; hardware modeling; 2100 compiled code models.
27. May 1988.
28. No charge for source.
29. User-defined; CMOS semicustom library at switch-level, 82.
30. Also supports circuit level.
31. At circuit level, up to 2^{31} discrete voltage levels.
32. User-extensible through the M modeling language.

Vendor and contact	Simulator and cost	Host	Simulation level	Operating modes: Incr. compile	Operating modes: Save/restart	Operating modes: Image patch	Signal states: Logic levels	Signal states: Signal strengths	Timing support: Prop. delays	Timing support: Transition times	Timing support: Min/max/nom	Timing support: Violations	Primitives and models: No. of primitives	Primitives and models: No. of models	Primitives and models: Source provided	Circuit description method	Fault simulation	ATE interfaces
Silvar-Lisco 1080 Marsh Rd. Menlo Park, CA 94025 Herman Beke Vice President of Marketing (415) 324-0700	**HELIX system design** $40k–$80k+	DEC VAXstation GPX, VAXstation 3xxx, 7xx, 8xxx (VMS); Sun-3 (UNIX); Apollo (AEGIS); IBM 43xx, 9370, 30xx (VM/CMS)	G	●[33]	●	●	●[34]	●[34]	●	○	●	●	500[35]	4000[35]	○	CASS graphic schematic entry; SDL alphanumeric network language; HHDL hardware description language	None	TSSI interfaces
			B	●[33]	●	●	●[34]	●[34]	●	○	●	●	500[35]	4000[35]	○			
Simucad Inc. 1040 Marsh Rd. Suite 200 Menlo Park, CA 94025 Jan Schuppert Manager, Marketing Administration (415) 321-2350	**SILOS** $10k–$20k (workstations); $45k–$100k (mainframes)	VAX (VMS, UNIX); IBM (MVS, CMS); Data General MV series; Apollo; Elxsi; Sun; MicroVAX II; CCI	S	●[36]	●	○	3	4[37]	●	○	●	N.a.	11	N.a.	○	SILOS netlist description	Concurrent	Sentry
			G	●[36]	●	○	3	4[37]	●	○	●	●	28	N.a.	○			
			B	○	●	○	3	4	●	○	●	●	N.a.	64	●			
	P/C-SILOS $995–$5k (1.5k–16k gates)	IBM PC XT, AT, and compatibles	S	●[38]	●	○	3	4[37]	●	○	●	N.a.	11	N.a.	○	SILOS netlist description	None	None
			G	●[38]	●	○	3	4[37]	●	○	●	●	28	N.a.	○			
			B	○	○	○	3	4	○	○	○	○	N.a.	8	○			
Simulog Inc. 2336H Walsh Ave. Bldg. G Santa Clara, CA 95051 Bruce Barnard Vice President of Sales (408) 727-8272	**SuperSim** $250k (Model 02) through $2.7M (Model 40)	VAX; IBM mainframes	G	●	●	●	3	1	○	○	○	○	14	●[39]	○	SILOS netlists	Serial or multiple copies in parallel	None
Spectrum Software 1021 S. Wolfe Rd. Sunnyvale, CA 94086 Karen Burchfiel Customer Service Representative (408) 738-4387	**Micrologic II** $895 (includes editors)	IBM PC XT, AT, and compatibles	G	●	●	●	3	1	○	○	●	○	N.a.	300	●	Graphic schematic entry	None	None
Teradyne Inc. 321 Harrison Ave. Boston, MA 02118 Daryl Layzer Marketing Services Manager (617) 482-2700, ext. 2808	**LASAR Version 6** $15k (VAXstation 2000); $175k (VAX 8800)	VAXstation 2000 through 8800 (VMS); Teradyne Dataserver	S	●	●	○	15[40]	●[40]	●	●	●	N.a.	5	0	N.a.	Wirelist; netlist converters; JEDEC file converter; Teradyne Modeling Language (TML); LABEL behavioral language	Concurrent; supports all simulation levels (S, G, B, H); fault selection options, including statistical sampling	Teradyne; Computer Automation; GenRad; HP; Schlumberger; TSSI interfaces; ASCII output format
			G	●	●	○	15[40]	●[40]	●	●	●	●	9	4000+[41]	●[42]			
			B	●	●	○	15[40]	●[40]	●	●	●	●	17	9[43]	●			

N.a. = not applicable.
33. Incremental compilation of HDL models into model database only.
34. User-defined; also simulated analog.
35. Only 24 models with mixed analog-digital simulation.
36. Input stimulus and ROM/RAM tables only.
37. 4 additional states also supported.
38. ROM/RAM tables only.
39. Supports SILOS libraries.
40. Actively driven signals may be assigned min/max drive strengths from up to 256 explicit values from 1 nA to 1 A for 1, 0.
41. Plus hardware models.
42. Not for LSI/VLSI nor for gate array libraries where requested by the vendor.
43. Behavioral models also available from third parties, including Logic Automation.

Vendor and contact	Simulator and cost	Host	Simulation level	Operating modes: Incr. compile	Operating modes: Save/restart	Operating modes: Image patch	Signal states: Logic levels	Signal states: Signal strengths	Timing support: Prop. delays	Timing support: Transition times	Timing support: Min/max/nom	Timing support: Violations	Primitives and models: No. of primitives	Primitives and models: No. of models	Primitives and models: Source provided	Circuit description method	Fault simulation	ATE interfaces
Valid Logic Systems Inc. 2820 Orchard Pkwy. San Jose, CA 95134 Ben Tang Product Marketing Manager (408) 432-9400	**ValidSIM** $28.8k (software only)	Sun-3; VAXstation II; personal computers; Valid SCALDsystems; VAX	S[44]	●	●	○	3	5[45]	●	●	●	N.a.	2	●[46]	○	ValidGED graphics editor; Concorde silicon compiler interface; PALASM, CUPL, ABEL PAL generators; device plus timing data for hardware models; SCALD III language	Concurrent through LASAR 6	Through LASAR 6: Teradyne L200, L100, J941; GenRad 1790, 2225, 2235; HP DTS70; TSSI interfaces to GenRad, HP, Sentry, Trillium, IMS, STS
			G	●	●	○	3	5	●	●	●	●	50	2000[46]	●			
			B	●	●	○	3	5	●	●	●	●	●[47]	●[46]	N.a.			
Vamp Inc. 6753 Selma Ave. Los Angeles, CA 90028 John Soluk Marketing Manager (213) 466-5533	**McCAD D/SIM** $395	Apple Macintosh 512K; Mac Plus	S	○	●	○	3	2	○	○	●	○	4	●[47]	N.a.	Netlist; parts list; schematic entry	None	None
			G	○	●	○	3	2	○	○	●	○	8	●[47]	N.a.			
Viewlogic Systems Inc. 275 Boston Post Rd. West Marlboro, MA 07152 Sri Sriram Director of Marketing (617) 480-0881	**VIEWSIM** $3k (IBM PC); $10k–$30k (VAX)	IBM PC and compatibles; VAX; NCR Tower; Sun	S	●	●	○	4[48]	3	●	○	●	●	5	3000	●	Schematic entry; PAL equations; C behavioral modeling, VHDL behavioral modeling	None	TSSI interfaces
			G	●	●	○	4[48]	3	●	○	●	●	68	3000	●			
			B	●	●	○	4[48]	3	●	○	●	●	●[49]	3000	●			
VLSI Technology Inc. 1109 McKay Dr. San Jose, CA 95131 Bill Murray Software Marketing Manager (408) 434-7660	**VTIsim** $25k	Apollo; VAX, MicroVAX; HP 9000 Model 300, 350; Elxsi; Ridge; Sun-3; Mach 1000 link	S	●	●	●[50]	●[51]	●[52]	●	●	●	N.a.	●[53]	●[54]	○	Graphics schematic entry (equation input, VTI-model modeling language); silicon compiler input (state machine, datapath, memory, ALU compilers)	Concurrent through Mach 1000 hardware accelerator; parallel through HILO-3	Sentry 10, 20; STS 256; MegaOne
			G	●	●	●[50]	3	1	●	●	●	●	27	27	○			
			B	●	●	●[50]	3	4	●	●	●	●	●[55]	1000	○			
The Western Design Center Inc. 2166 E. Brown Rd. Mesa, AZ 85203 William D. Mensch, Jr. President (602) 962-4545	**LOGIC** $1995 (IBM PC AT); $9995 (VAX)	Apple II GS; IBM PC AT; Prime; VAX	S	○	○	○	3	2	●	○	○	N.a.	3	●[56]	○	Graphic schematic entry; hand coding	None	Generated vectors can be edited
			G	○	○	○	3	2	●	○	○	○	10	●[56]	○			
			B	○	○	○	3	2	●	○	○	○	3[57]	●[56]	○			

N.a. = not applicable.
44. MOS pass transistor.
45. Hard, soft, mem, HiZ, and undetermined.
46. Plus 70 RealChip and RealModel models and 3000 (primitives) with LASAR 6.
47. User-defined.
48. 28-state simulator also supports 21 unknown states.
49. No hard upper limit
50. Restricted.
51. Continuum (voltage).
52. Continuum (resistance).
53. 2 CMOS (continuum of sites), 4 HMOS.
54. 2 CMOS, 4 HMOS.
55. Expandable.
56. Macros.
57. Plus macros.

Vendor and contact	Simulator and cost	Host	Simulation level	Operating modes: Incr. compile	Operating modes: Save/restart	Operating modes: Image patch	Signal states: Logic levels	Signal states: Signal strengths	Timing support: Prop. delays	Timing support: Transition times	Timing support: Min/max/nom	Timing support: Violations	Primitives and models: No. of primitives	Primitives and models: No. of models	Primitives and models: Source provided	Circuit description method	Fault simulation	ATE interfaces
XCAT Inc. 2855 Anthony Lane Minneapolis, MN 55418 Charlie Loegering Vice President (612) 789-2050	**MX/MXT hardware accelerators** $39k–$120k	IBM PC XT, AT, and compatibles; Apollo Series 3000; DEC VAXmate	S	●	●	●	3	4	●	●	○	N.a.	N.a.	N.a.	●	Translators for Mentor, FutureNet, TEGAS, EDIF, HILO, SILOS	Parallel	Through supported vendor
			G	●	●	●	4	4	●	●	○	●	128	350	●			
Xerox Corp. 2945 Oakmead Village Court Santa Clara, CA 95051 Petros Xides CAE Product Division Manager (408) 982-8132	**Expert Logic Simulator** $14k (unit quantities; includes HDL)	Xerox 8000 series, Xerox 6085 Professional Workstation	S	●	●	○	3	3	●	●	●	N.a.	10	50	●	Graphic schematic entry; XDDL hardware description language; hierarchies of graphics, HDL	None	None
			G	●	●	○	3	3	●	●	●	●	10	100	●			
			B	●	●	○	3	3	○	○	●	●	20	400	●			
Zycad Corp. 1315 Red Fox Road St. Paul, MN 55112 **Silicon Solutions Corp.** 1380 Willow Rd. Menlo Park, CA 94025 Peggy Kalina Marketing Communications (612) 490-2500	**Logic Evaluator; Expeditor; Magnum** $38k–$2M	VAX, MicroVAX; Apollo; Prime; IBM; Sun	S	●	●	●	3	4	●	●	○	N.a.	10	1400	●	Zycad Intermediate Format; ZILOS; TEGAS TDL; Mentor; Tektronix CAE Systems; Daisy; HILO (GenRad); SILOS; Silvar Lisco; CADAT (HHB Systems)	Serial; concurrent with fault evaluator upgrade	In development
			G[58]	●	●	●	3	4[59]	●	●	○	○	31[60]	1400	●			
			B															
	System Development Engine $1.7M–$2.5M	VAX; IBM mainframes and compatibles	S[61]	●	●	●	3	4	●[62]	●	○	N.a	10	N.a.	N.a.	Zycad Intermediate Format; ZILOS; TEGAS TDL; Mentor; Daisy; HILO (GenRad); SILOS; CADAT (HHB Systems)	None	None
			G	●	●	●	3	4	●[62]	●[62]	○	○	31	1400	●			
	Mach 1000 $95k–$1M	Stand-alone processor; Ethernet server	S	●	●	●	4	8[63]	●	●	●	N.a.	28	●[64]	N.a.	Binary formats; hierarchical EDIF format; netlists of various simulators; Mentor; Daisy	Accelerated concurrent	Sentry
			G	●	●	●	4	8[63]	●	●	●	●	192	●[64]	N.a.			
			B	●	●	●	4	8[63]	○	●	●	●	●[64]	●[64]	N.a.			

N.a. = not applicable.
58. Also supports functional level.
59. 3 at functional level.
60. 19 functional primitives.
61. Unidirectional models only.
62. Zero or unit delay only.
63. Uses 136 intervals (combinations of states and strengths) to avoid order-dependent calculations and to resolve contentions.
64. Depends on supported software simulator.

This baby is gifted!

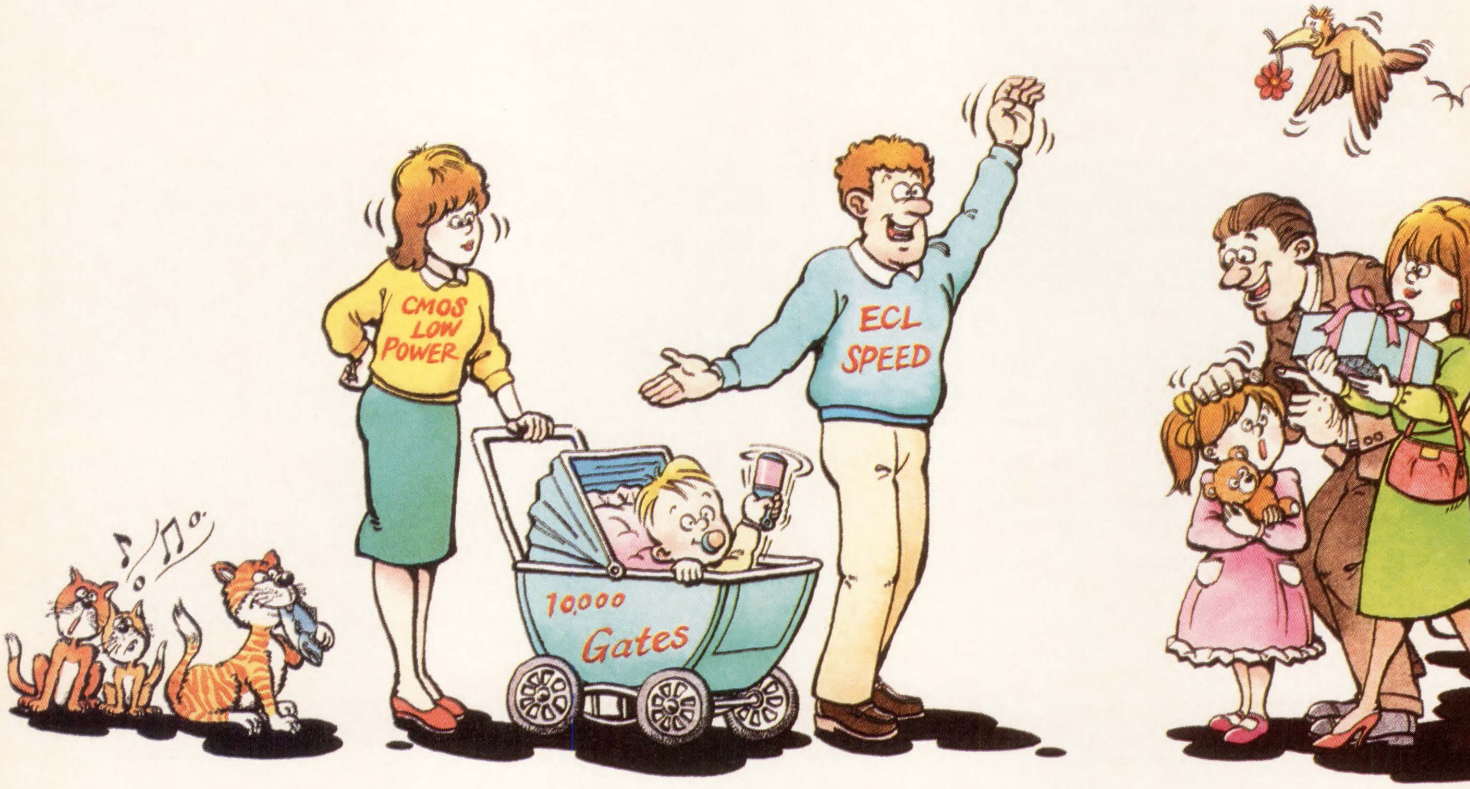

The new generation BiCMOS gate array with 10,000 gates, ECL speed and CMOS power economy.

Mating ECL speed with CMOS power savings, NEC created the gifted family of BiCMOS gate arrays. Now the next generation arrives.

Announcing BiCMOS-4A. It gives you more gates: 10,000 and 6,000. Double the drive capability: 24mA/output buffer. And more I/O pads: 228 for the 10,000-gate device; 180 for the 6,000-gate version.

Of course, BiCMOS-4A inherits all the advantages of the BiCMOS-4 family. Including high

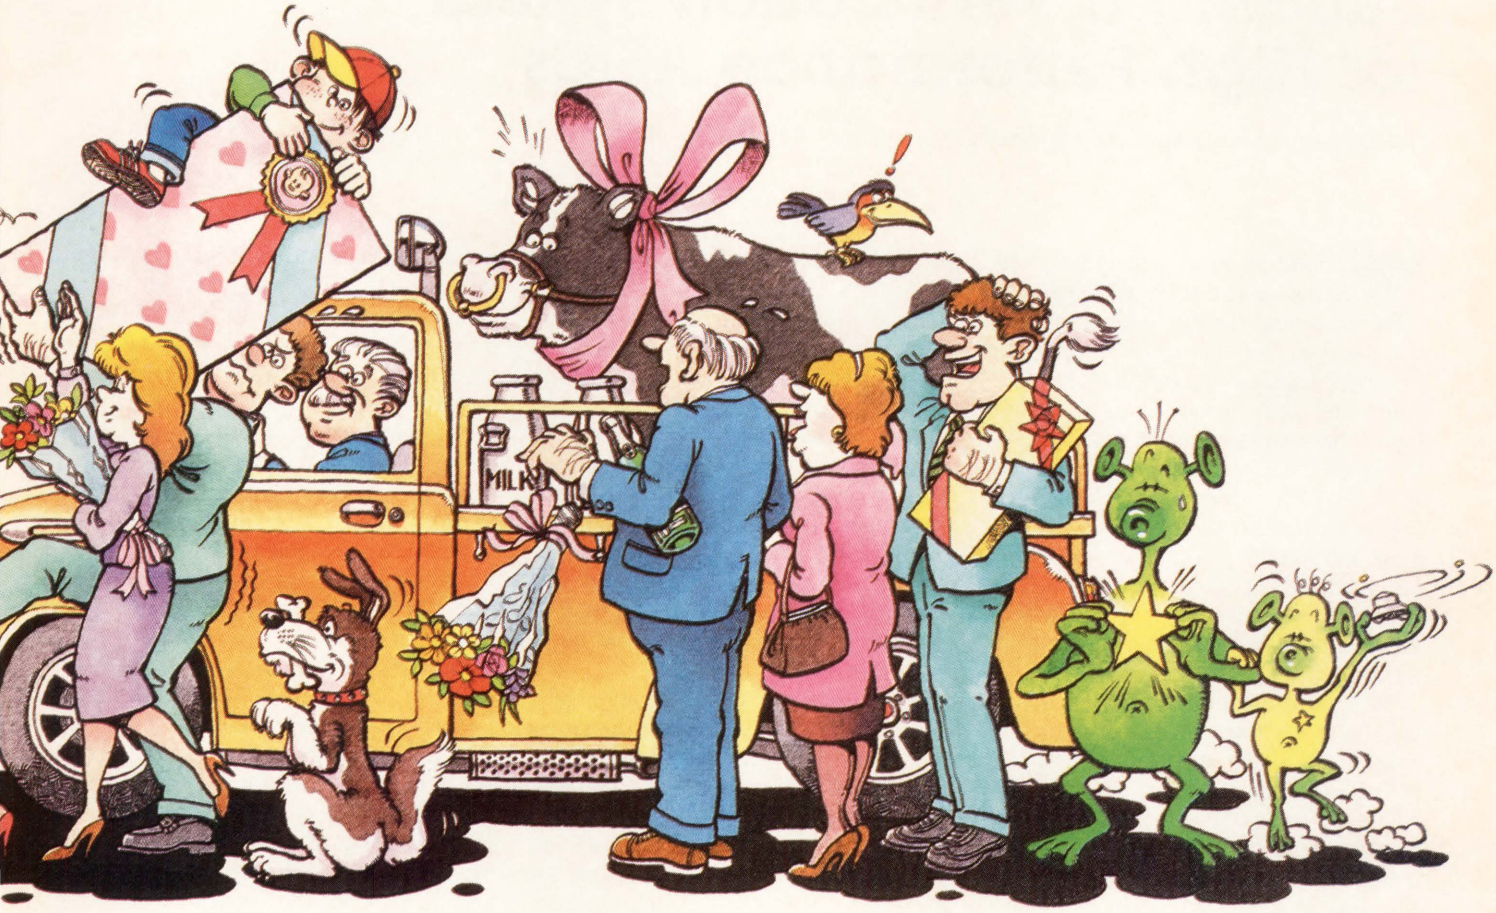

speed: 0.8ns/gate. Low power consumption: 18μw/gate. And freedom from latch-up problems.

BiCMOS-4A gate arrays come in a wide variety of package types. Including Flat, PLCC, PPGA and PGAs with diverse pin counts. 280-pin PPGA and PGA packages are available for the 10,000-gate device.

Users can draw on a library of 89 macros and 103 I/O blocks. We support interface with seven popular engineering workstations.

To find out how our gifted gate arrays can brighten your product's future, call NEC today.

Device	μPD67100	67060	67030	67020	67010	67001
Gate (2-input NAND)	10,348	6,372	3,140	2,248	1,124	624
I/O pads	228	180	140	120	84	64
Internal gate delay	0.8ns (F/O = 3, L = 3mm)					
Input-buffer delay	1.2ns (F/O = 3, L = 3mm)					
Output-buffer delay	3.0ns (C_L = 15pF)					
Output current (I_{OL})	24mA		12mA			
I/O compatibility	Input: CMOS, TTL. Output: TTL.					
Power supply voltage	5V ± 0.5V					

For further information, please contact: USA Tel:1-800-632-3531. In California: Tel:1-800-632-3532. TWX:910-379-6985. W. Germany Tel: 0211-650302. Telex: 8589960. The Netherlands Tel:040-445-845. Telex:51923. Sweden Tel:08-732-8200. Telex:13839. France Tel:1-3946-9617. Telex:699499. Italy Tel:02-6709108. Telex:315355. UK Tel:0908-691133. Telex:826791. Hong Kong Tel:3-755-9008. Telex:54561. Taiwan Tel: 02-522-4192. Telex:22372. Singapore Tel:4819881. Telex: 39726. Australia Tel:03-267-6355 Telex:38343.

NEC

Building a Verification System for High-Performance ASICs

Jim Graydon, Integrated Measurement Systems Inc., Beaverton, OR

To build an ASIC prototype verification system with 100-MHz test rates and flexible user operation, the designers of the Logic Master XL followed a strategy of shrinking key circuitry onto semicustom ICs and employing a distributed-resource architecture. Designed to verify fast, complex devices such as VHSIC phase II components, the Logic Master XL provides edge placement and sampling resolution of 100 ps; 125-ps cycle period resolution at a 100-MHz test frequency; ±1-ns system skew; and automatic, wireless signal-pin fixturing.

The semicustom integrated circuits that lie at the heart of the XL hardware consist of a linear array design and two ECL digital gate array designs. The linear design integrates pin driver and comparator resources, with one array required for each input/output channel. The system timing resources are generated using three identical timing arrays, each array providing four timing channels (eight signal edges). The I/O-channel data formatting and memory multiplexing are done using another ECL array—the formatter-multiplexer, or format/mux (F/M) array. One F/M array is used for every two of the 224 (maximum) bidirectional channels in the system.

Distributed Architecture

At the outset of the development effort, the XL design team decided that a "per-pin" architecture—where each pin has a complete set of test circuitry—was the best way to ensure ease of use. Also, the ability to make arbitrary pin assignments would reduce or eliminate the amount of custom wiring needed to fixture the device under test. Fixturing was an important consideration, given that timing errors related to inefficient fixturing only get worse at higher test frequencies.

Some production ATE systems provide per-pin control by duplicating all resources behind each channel. Though such a tester-per-pin architecture was deemed too expensive for a prototype verification system, the design team felt that they could provide most of the same advantages through a distributed-resource approach. Following this approach, timing and voltage resources were concentrated on a timing module and then distributed through the backplane (see Figure 1).

Force, Inhibit, Expect, Acquire Data, and Mask Data for all pins are contained on the data module, along with the control circuits that select timing and voltage levels from the distributed resources on a per-pin basis. Driver/receiver arrays also reside on this board, with one linear array available for each I/O channel in the configuration.

A separate control module contains the system processor that provides the user interface with the hardware. This module also contains a high-speed sequencer that controls the data program flow.

The Timing Module

The timing and voltage resources distributed to individual pins are incorporated within the system's timing module. The module provides the reference voltages for the system's four driver voltage limits (called driver rails), four threshold rails, and four comparator threshold rails. Programmable in 10-mV increments, the driver rails enable the user to generate ECL, TTL, and CMOS levels all at the same time while leaving a fourth set of rails for voltage margin testing at each pin.

The three digital gate arrays contained within the module each provide four timing sets with delay and width signal edges, for a total of 24 edges. These timing channels, coupled with the data module's formatting capabilities, ensure that the system can generate the complex timing relationships typically required to verify the performance of high-density ASICs.

The timing board also contains the system clock generators, which provide the system's 125-ps cycle time resolution, for test cycles ranging from 10 ns to over 6 μs.

On the data module, the timing channels are combined with test data and applied to the device under test (DUT). This module contains the 5 bits of memory at each pin that control Force, Inhibit, Expect, Mask Data, and Acquire Data for each test cycle. It also provides the user with full bidirectional channel capability, even allowing Force and Acquire bits on the same channel during one cycle.

The Format/Mux Array

The workhorse of the data module is the format/mux array. As its name implies, it serves two functions. First, it provides a 4:1 multiplexer for the data and control bits for each pin, so that 25-ns static RAMs can be used for the memory while still achieving a 10-ns cycle time at the DUT. Second, it formats the forced data, samples the acquired data, and makes real-time comparisons of expected and acquired data. The array also performs more complex functions, such as selecting the timing for each individual channel and providing calibration delays for the 100-ps resolution force and acquire timing. One F/M array supports two I/O channels.

With 6,000 gates each, digital gate arrays made by NEC were selected for

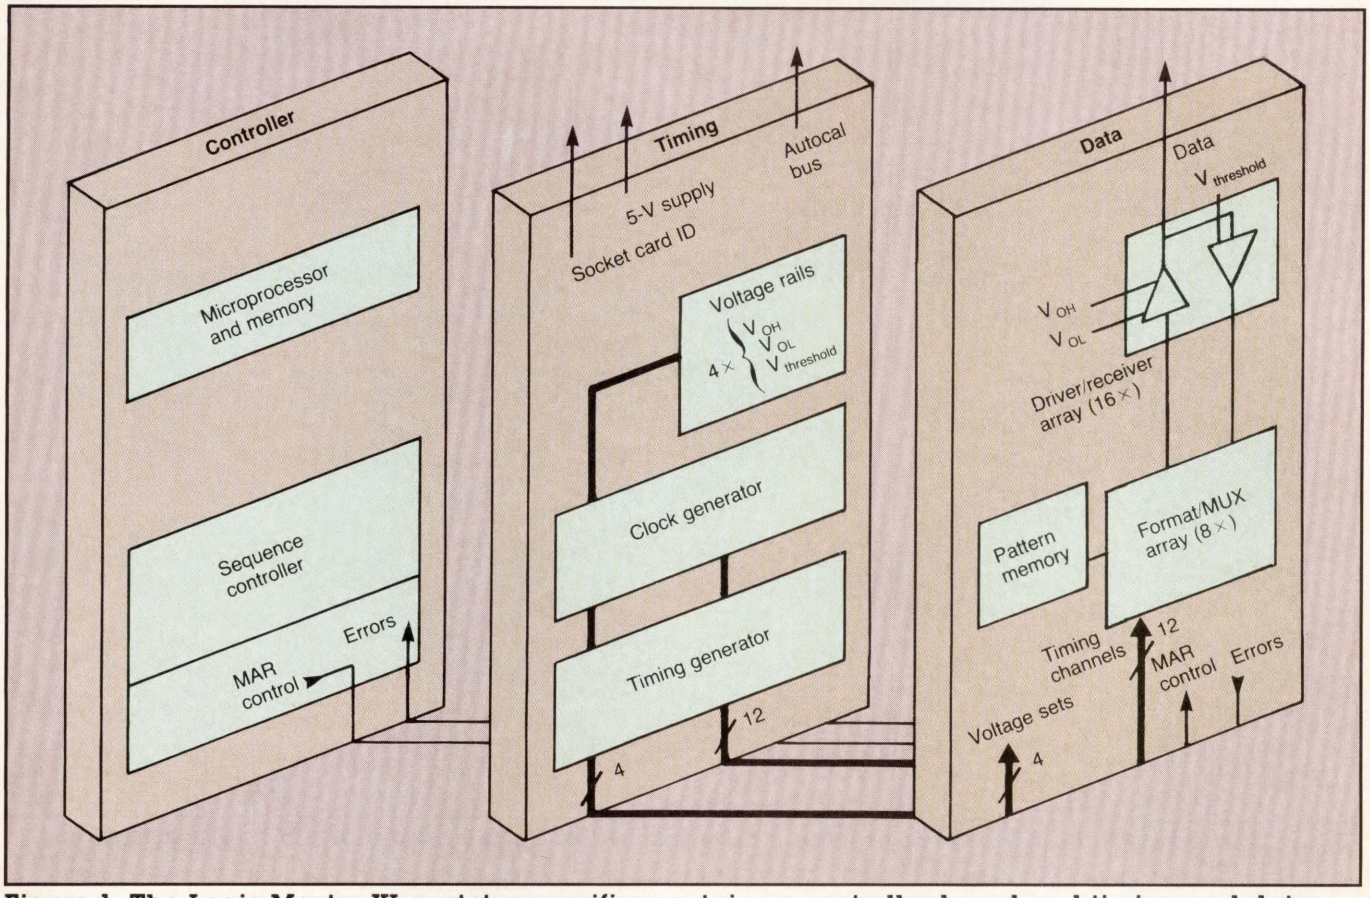

Figure 1. The Logic Master XL prototype verifier contains a controller board and timing and data boards.

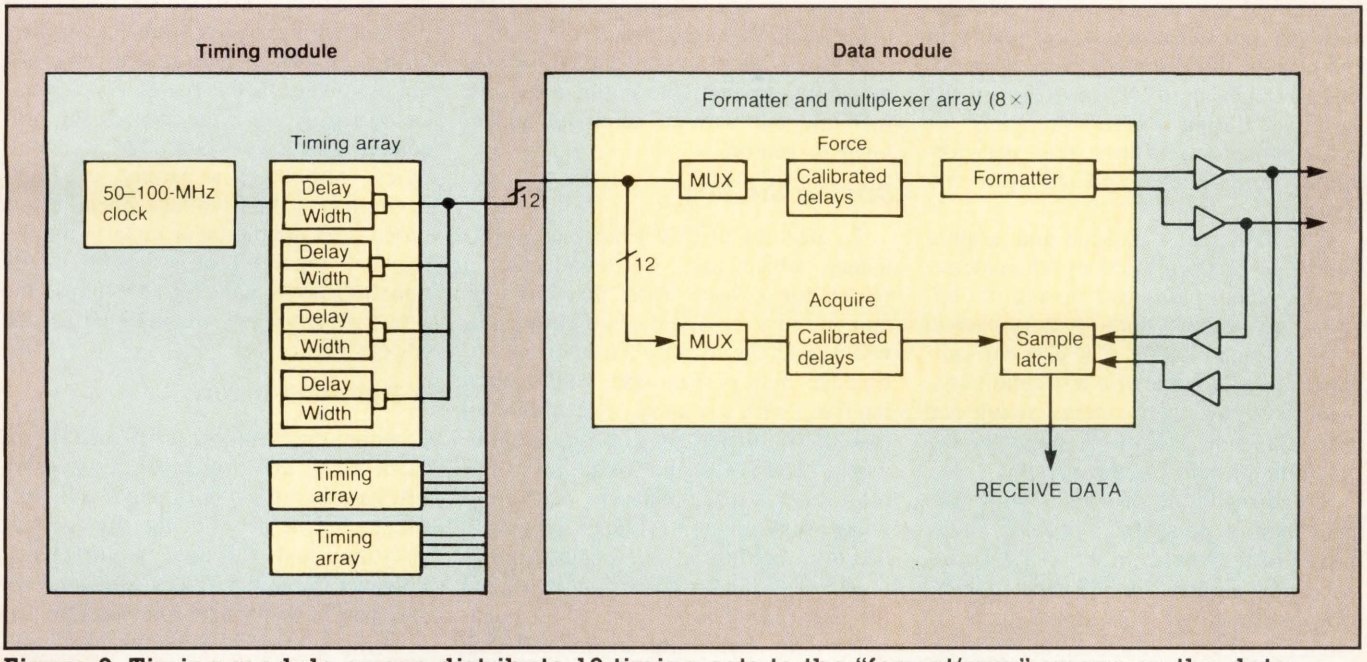

Figure 2. Timing module arrays distribute 12 timing sets to the "format/mux" arrays on the data card.

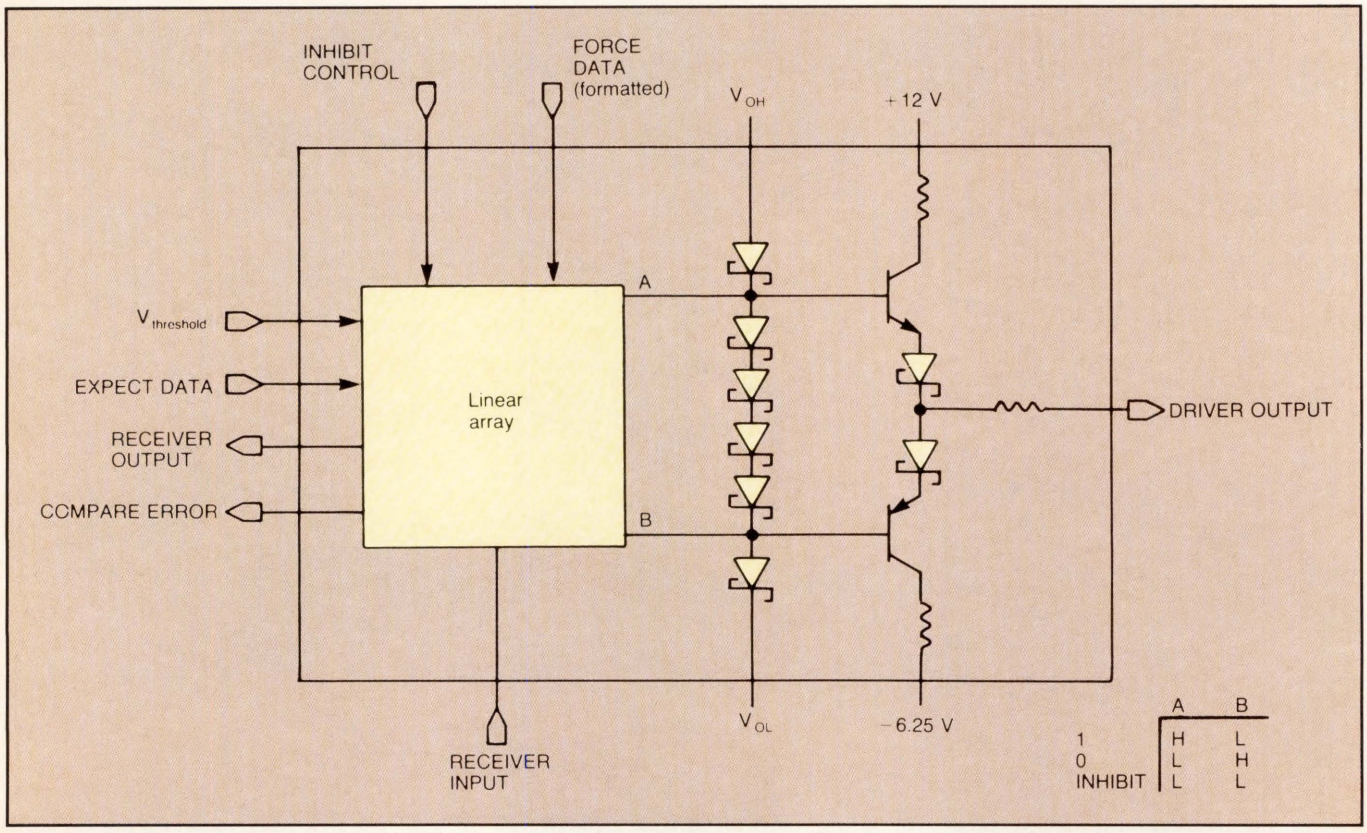

Figure 3. Discrete Schottky diodes provide signal resolution for the output of the linear arrays.

the F/M array, as well as for the timing generator, because they were among the largest semicustom ECL devices manufactured at the time. Available real estate was put to good use, especially in the case of the F/M array, where the XL design team used up to 97% of the array's building control structures for force and acquire timing sets, data formats, and mask and comparison circuitry.

Calibration of the force and acquire timing is accomplished by characterizing the delay paths and storing the appropriate values into battery-backed RAM on each module. The user can calibrate the XL for specific test conditions, such as a particular voltage swing and input threshold. Figure 2 shows the basic timing chain for the system.

The distribution of signals with 100-ps resolution presented a problem in the layout of the boards in the XL. All critical timing signals were distributed using 50-Ω differential ECL signal paths to reduce signal noise and coupling problems. Differential signal paths must have two complementary component signals that are the same length and that are terminated correctly at the end of the signal trace. The designers had no automatic tools that could route the critical signals, so they routed many signals manually. In fact, the entire data card was laid out without automatic placement or routing.

Driving the DUT

To meet the XL driver/receiver performance objectives, we conducted a search for a linear array that provided npn and pnp transistors with gain-bandwidth products greater than 2 GHz. Also, because high-performance ASICs often interface with a variety of technologies that have different voltage requirements—TTL, ECL, and CMOS, for example—the XL would have to generate a wide programmable voltage swing. In addition, because of the special requirements of ECL ICs, the designers wanted the linear array to support a four-quadrant driver that could source and sink current when driving either a logic 1 or a logic 0.

After making inquiries at many foundries, we discovered that the fastest npn transistors provided by most complementary processes were just 1 GHz, with even slower pnp devices. We had almost given up the search for a semicustom array and were considering full-custom processes when AT&T Technologies introduced a family of high-performance, high-voltage linear arrays based on a double-layer metal, complementary bipolar process. With 4-GHz npn transistors and 2.5-GHz pnps, the AT&T ALA200 UHF was selected for the XL design.

A Hybrid Solution

Although the linear array met almost all of the XL driver/receiver design criteria, its process did not include Schottky diodes on the die. For the driver's fast rise-time and fall-time requirements, however, Schottky diodes were a necessary part of the design. Our solution was to build an external Schottky-based output stage. This stage and the array (see

Figure 3) were then incorporated on a single hybrid substrate. With this configuration, the linear array provides tightly controlled current sources that can drive controlled parasitic capacitances linearly, ensuring signal fidelity.

Implementing this hybrid solution did present a challenge to the designers and foundry. SPICE simulation models provided with the array did not accurately model transistor quasi-saturation performance or predict potential device breakdowns. Quasi-saturation is the operating point of a transistor when V_{CE} is roughly 1 V—at which point the operating characteristics of the device are making the transition from linear operation to that of saturation. Because performance could not be modeled in this operating area, the part had to be redesigned to avoid quasi-saturation.

Automatic Fixturing

The high-quality signals from the array would do little good if the fixturing degraded their waveforms before they arrived at the DUT. To that end, we wanted prewired fixturing implemented with controlled-impedance traces. If users did not have the flexibility of configuring any channel for input, output, or bidirectional signals, the only solution would be to create a custom fixture board with connections hard-wired for each DUT.

Fortunately, with the XL's per-pin resource assignment capabilities, standard fixturing can be used and pin assignments can be made in software. By eliminating the need to wire the DUT to specific pins on the tester, "automatic fixturing" minimizes errors and fixturing time. The DUT socket solders into a fixture board that connects to the driver and receiver through matched-length, 50-Ω traces. Users can simply switch fixture boards to accommodate different package types, such as DIPs, pin-grid arrays, and leadless chip carriers. Custom fixturing for other packages and wafer probing are also available.

From the timing generator to the fixturing, building a prototype verification system that pushed current measurement speeds and that provided per-pin control over system resources required a major reworking of former system architectures. The use of ASIC technology allowed the XL team to design a system that met these challenges while maintaining the advantages of relative low cost, compact size, and ease of use compared with production ATE. Fortunately, we found ASIC vendors whose semicustom logic met our design needs. Otherwise, we would have had to implement a full-custom solution, which would have been more difficult, been more expensive, and taken longer to implement. □

Mechanical Design and Thermal Analysis in a CAE Environment

Mentor Graphics has taken a significant step into the future by developing a workstation for the design of electronic packaging that is linked to the company's Board Station PCB design and layout system. The Package Station, which is indicative of a growing industry belief that more engineering tasks should be integrated with the electronic designer's traditional tool environment, includes three-dimensional geometric modeling, drafting, thermal analysis, and interface software for designers of subsystem and system enclosures, PCB card cages, and IC packages.

Using this system, designers construct a representation of the enclosure or package. To assist in drafting, the system uses a concept called a "work plane." Rather than define three-dimensional coordinates, shapes are defined in two dimensions and automatically placed with reference to the work plane.

The work plane can be at any angle to the plane on which the shapes are to be situated. Thus it can be made parallel to the workstation screen so that the designer can "peer down" on the image, as he would on a drafting table. With its usual finesse, Mentor also provides user-definable, multiple-window views. The angle of each view is adjustable, and , impressively, the image in each view is continously updated.

Designers can create composite parts by grouping model items; these parts can then be repeated, moved, and rotated as a unit. Several styles of text, numbers, and arrowheads include support of the Department of Defense standards. The designer can also add additional properties, like material cost and supplier, to the models.

For thermal analysis, the system extracts a description of the circuit board

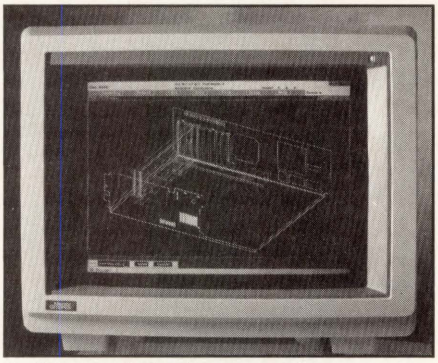

A 3D model of an enclosure on the Package Station.

and the placement of its components from Mentor's PCB layout software. Temperature attributes can be assigned from a resident parts library, containing such thermally related properties as resistivity, conductivity, power dissipation. Boundary conditions and heat transfer coefficients can be defined by the user. Designers can also define their own models.

The system then generates, in a matter of minutes, a mesh representation of the board for finite-element analysis. Traditionally, mesh generation has been a time-consuming phase, and its automation by Mentor distinguishes this system from those offered by CADAM and Valid.

The tools can provide conduction, convection, and radiation analyses. However, the analysis is two-dimensional only: Separate analyses are required to examine lateral and vertical thermal gradients. Similarly, gradation of thermal coefficients is linear only: The software assumes a uniform coolant flow in one direction over one surface. Consequently, the designer must manually adjust the coefficients for all the components if the airflow in a fan-cooled container is irregular because of varying heights or shapes in the cabinet or if there are to be other localized heat sources near to parts under analysis. Nevertheless, the analysis does not just assume monotonic heat dispersion from all parts of the surface, which is somewhat better than many other commercially available systems.

The thermal analysis tools can detect potential thermal failures caused by board and component overheating, predict junction temperatures, and estimate the effects of heat dissipation within IC packages. Analysis results can be displayed as an isoline map (showing relative criticality), as a color-coded temperature profile, and as a tabular report. Temperatures can appear as colors, or the system can display the absolute temperatures of packages.

The analysis tools do not account for dynamically related thermal effects. Some other systems use the gate toggle frequencies, extracted from simulation, to predict heat generation more accurately. Mentor's "snapshot" approach also does not account automatically for differences between active and standby power consumption. Instead, a part's power dissipation rating can be edited by the user to reflect these differences. Of course, it may be questionable whether more elaborate analysis is really worth the effort—it may actually be *easier* to edit part attributes, and added dynamic analysis may not be much more accurate any way.

Package Station includes an Apollo 3000 workstation with 4 megabytes of RAM, a 19-in. monitor, software, and documentation for $55,000. It will be available in the second quarter.

Mentor Graphics Corp.
Beaverton, OR 97005
(503) 626-7000

Advertisers' Index

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

S E M I C O N D U C T O R E N G I N E E R I N G

Digital has it now.

Blue Chip Opportunities

Blue Chip opportunities are not just on the stock market.

You'll find them in Digital's Semiconductor Engineering group.

You'll be working for one of the most successful groups in a remarkably successful company.

Apply your skills to VLSI design using emerging semiconductor processes, design tools and packaging technologies. Build advanced CMOS microprocessors for multiprocessor configurations.

Develop advanced CPU and peripheral megacells for use in VAX*-based systems on a chip, working with advanced processes and CAD tools to define design methods and deliver subsystem building blocks for final silicon system products.

We're looking for innovative, enthusiastic and visionary designers at the consulting level (more than 10 years' experience) and principal engineering level (more than 5 years).

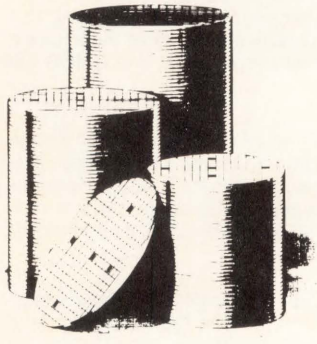

Semiconductor and CAD Design

PRODUCT MANAGER

Maintain a product management organization in support of the Semiconductor Engineering Group. Duties include writing of high level strategic business plans, End of Life responsibility, support of MOS Marketing base plan and customer interface. Requires at least 5 years' product management experience in the Semiconductor industry.

VLSI Design

CIRCUIT DESIGN CONSULTANT

Provide technical leadership innovation and tackle the challenging circuit design issues in the rapidly growing internal semi-custom design business. Requires at least 6 years' experience in transistor level CMOS circuit design, working knowledge of device physics and process technologies, and a good understanding of logic design, VLSI layout and associated CAD tools.

MEGACELL ARCHITECT

Lead a VLSI design team as the architect and project leader for high performance megacells that will form the basis for tomorrow's one chip systems. You should be expert in many of the following areas: computer systems architecture, VLSI microarchitecture definition and behavioral modeling, CMOS logic design and verification, CMOS transistor level design, layout planning and review and test generation. Requires at least 5 years' experience in VLSI microprocessor or peripheral chip design.

CAD Development

CAD SIMULATION MANAGER

Manage a project that will provide state-of-the-art tool suites to the Digital community. You will need a BSEE and 10 years' industry experience, 5 years of which should be in simulation (electrical). Some VAX and logic simulation background, and at least 3 years' group management.

CAD SENIOR ENGINEER

Develop automatic layout tools for Integrated Circuits Design. The design process includes standard cell, placement and routing, and chip-level assembly.

You will need a BSEE and 5 + years' hands-on development experience.

Microprocessor Development

INTERCONNECT DESIGN CIRCUIT DESIGNER

Opportunity to work in a CPU Development group involved in interconnect of high speed VLSI CMOS processor chips. Direct analysis, simulation and design of high speed buses and interconnect at package, module and backplane level. Education in and experience with transmission line theory and high frequency design is important. Industry experience in TTL or ECL design.

MICROPROCESSOR ARCHITECT

Design for a complex VLSI processor chip set. Responsibilities include initial top-level chip partitioning, chip microarchitecture, microcode development, performance modeling, behavioral modeling and specification development of the system interface and the memory subsystem(s). Requires a BSEE/CS and 5 years' VLSI processor design experience.

VERIFICATION/ TESTABILITY ARCHITECT

Develop total verification strategy for design, finished chips and complete system(s). Implement this strategy with and through the design team for a very complex fourth-generation microprocessor system. Requires a BS/MSEE and at least 5 years' experience.

RELIABILITY/ RECOVERABILITY SPECIALIST

Develop models of, and a strategy for, chip and system error detection and correction for a fourth-generation microprocessor system. Analyze data about chip and system failures, develop a chip and system error detection and recovery strategy and implement this strategy with and through the design team. Requires a BS/MSEE and at least 5 years' experience.

To cash in on these blue chip opportunities with a $9.4-billion computer company, contact: **Gary V. Schipani, Employment Department 0201 8825, Digital Equipment Corporation, 77 Reed Road, HLO2-2K12, Hudson, MA 01749-2895.**

*Trademark of Digital Equipment Corporation.

We are an affirmative action employer.

digital™

Keeping an eye on freedom.

Sometimes you cannot watch what is most important to you. Our basic freedoms are felt, not seen.

So keeping an eye on freedom becomes a question of technology. At Hughes Electro-Optical & Data Systems Group, we're helping our nation "see" with many types of sensors, lasers and data processing systems.

We currently have opportunities for:

- Technical Writers
- E-O System Engineers
- Optical Physicists
- Image/Signal Processing Engineers
- DP System Engineers
- Program Managers
- Focal Plane Electronics Engineers
- Component/Material Process Engineers
- ATE Engineers
- Quality Assurance Engineers
- Design Engineers
- Reliability Engineers
- Passive Device Specialists
- Gate Array Specialists
- Semiconductor Specialist
- Tribologist
- Metal Finishing Specialist
- Contamination Control Engineers
- Elastomer Specialist
- Specification Engineer
- Printed Wiring Board Process Engineer
- IC Test Engineer
- Environmental Engineers
- System Test Engineers
- E-O System Analyst
- E-O Sensor Mission Analyst
- Signal Processing System Engineers
- VLSI Design Engineers
- A to D Conversion Engineers
- Univac Scientific/Business Programmers
- Univac RDMS DataBase Programmers
- S/W Configuration Manager
- Application S/W Testing-Maintenance Engineers
- Contracts Administrators
- Contracts Negotiators
- Subcontract Negotiators
- Developmental Engineers
- Composite Engineers
- Material & Process Engineers
- Polymer Applications Engineers
- Configuration Management Administrators
- Property Administrators
- Test Engineers
- Reliability Engineers
- Design Engineers
- Program Managers
- Buyers
- Financial Analysts
- Facility Engineers
- Accountants
- Computer Operators
- Program Controls Administrators
- Environmental Engineers
- Teleprocessing Analysts

No training positions are available. Experience on special access programs is desirable.

We offer an attractive salary and an outstanding benefits package, including tax-deferred savings; medical, dental and vision care coverage.

Please send your resume to: Hughes Aircraft Company, Electro-Optical & Data Systems Group, Professional Employment Dept. ISLV-288, P.O. Box 913, El Segundo, CA 90245.

Proof of U.S. Citizenship Required. Equal Opportunity Employer.

Creativity America depends on.

HUGHES
AIRCRAFT COMPANY
ELECTRO-OPTICAL & DATA SYSTEMS GROUP

The Intel Influence

It's Excitement. It's Challenge.
It's Opportunity. It's Pride.

We're on the move. Lots of growth. Exciting projects and breakthrough technologies in systems and components. In short, there's never been a better time to be at Intel.

We have the influence and the resources that can make a profound impact on your career. If you thrive on hands-on responsibility and would like a chance to excel, look into the following opportunities:

Design Engineers

ASIC Designers

BS/MSEE with knowledge of analog and digital integrated design techniques (Bipolar and CMOS). Working knowledge of device modeling, logic design, transmission line theory and high frequency design.

Component Designers

BS/MSEE or CS with 2+ years experience in chip design, CMOS circuit and logic. Expertise in VMS, UNIX and CMS operating systems with Mainsail. Familiarity with SPICE.

VLSI Designer

MSEE or CS. Requires supervision of an engineering development group including microprocessor architecture, logic design, circuit design and layout of a portion of 32-bit microprocessors. Floating point or memory management experience desired.

CAD Engineers/ Managers

BS/MSEE or CS with 2+ years experience in the development of CAD tools. Positions currently available in simulation and modeling, APR, workstation development, functional design and physical design. Experience with Daisy and Mentor workstations a plus.

Software Network Engineer

BS/MSCS with 1+ years experience with TCP/IP and UNIX.

Marketing Engineers

BSEE or CS with 3+ years experience in Technical Marketing, Product Marketing or Customer Marketing. Product experience with ASICs, microprocessors, minicomputers or mainframes required.

We welcome your interest in Intel... please send your resume to the location of your choice.

Arizona: Intel Staffing, Dept. 2759, 5000 W. Chandler Blvd., Chandler, AZ 85226, or call Mike Foster at 602-961-5728

California: Intel Staffing, Dept. 2759, 3535 Garrett Drive, Santa Clara, CA 95050, or call Rebecca Canary at 408-765-1546

California: Intel Staffing, Dept. 2759, 1900 Prairie City Road, Folsom, CA 95630, or call Monica Lorain at 916-351-6476

Oregon: Intel Staffing, Dept. 2759, 5200 NE Elam Young Parkway, Hillsboro, OR 97124, or call Denise Velaski at 503-696-4597

Equal Opportunity Employer M/F/H

intel®

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

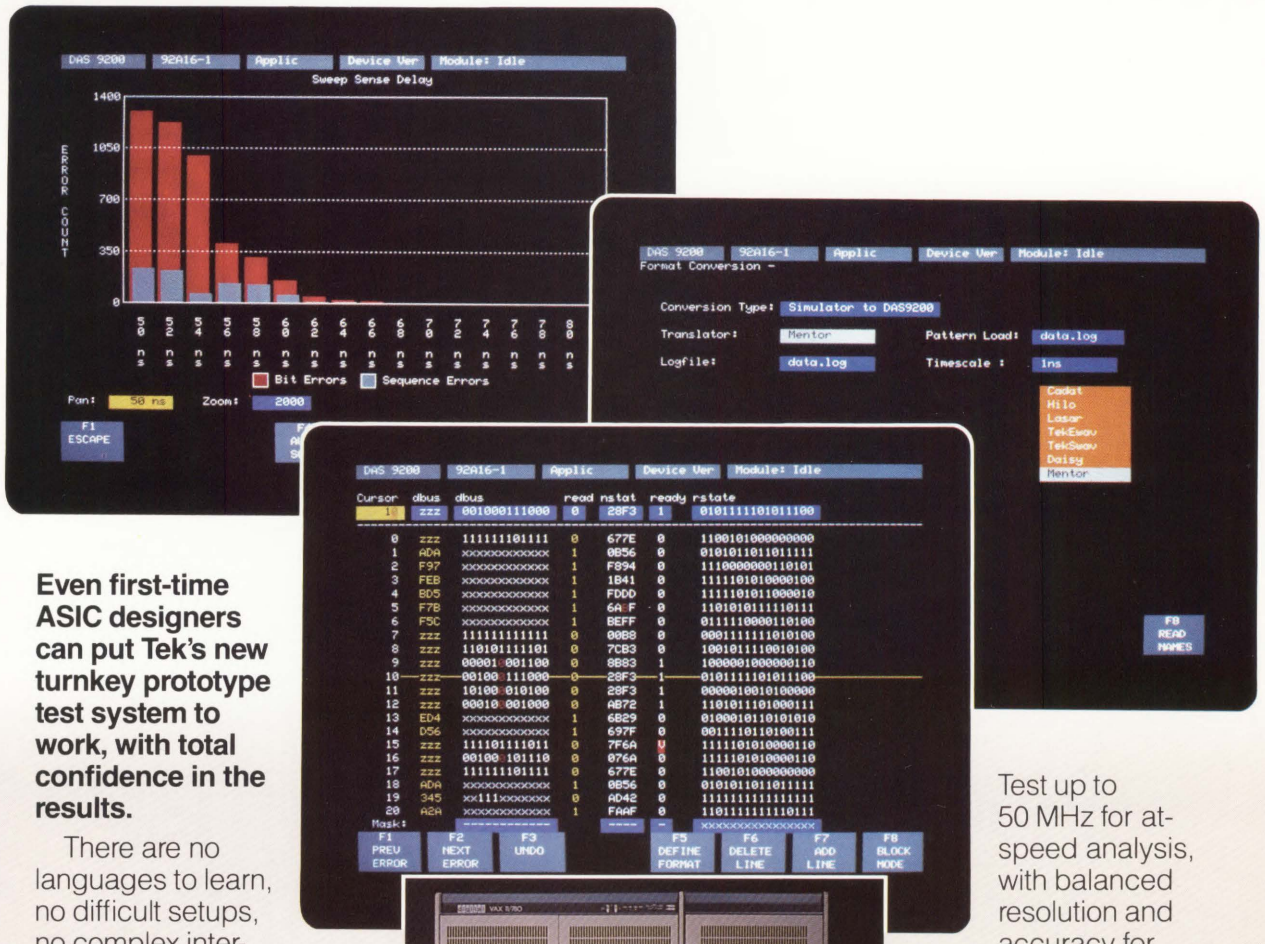

Even first-time ASIC designers can put Tek's new turnkey prototype test system to work, with total confidence in the results.

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

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

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

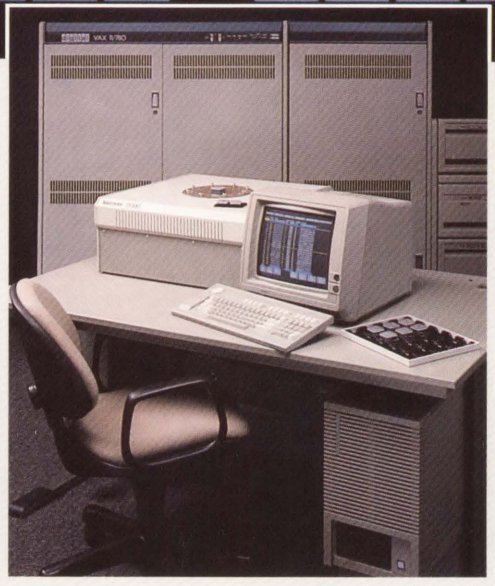

Test up to 50 MHz for at-speed analysis, with balanced resolution and accuracy for unprecedented signal measurement quality. The test fixture adapts to your DUT quickly.

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

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

Tektronix®
COMMITTED TO EXCELLENCE

CIRCLE NUMBER 14

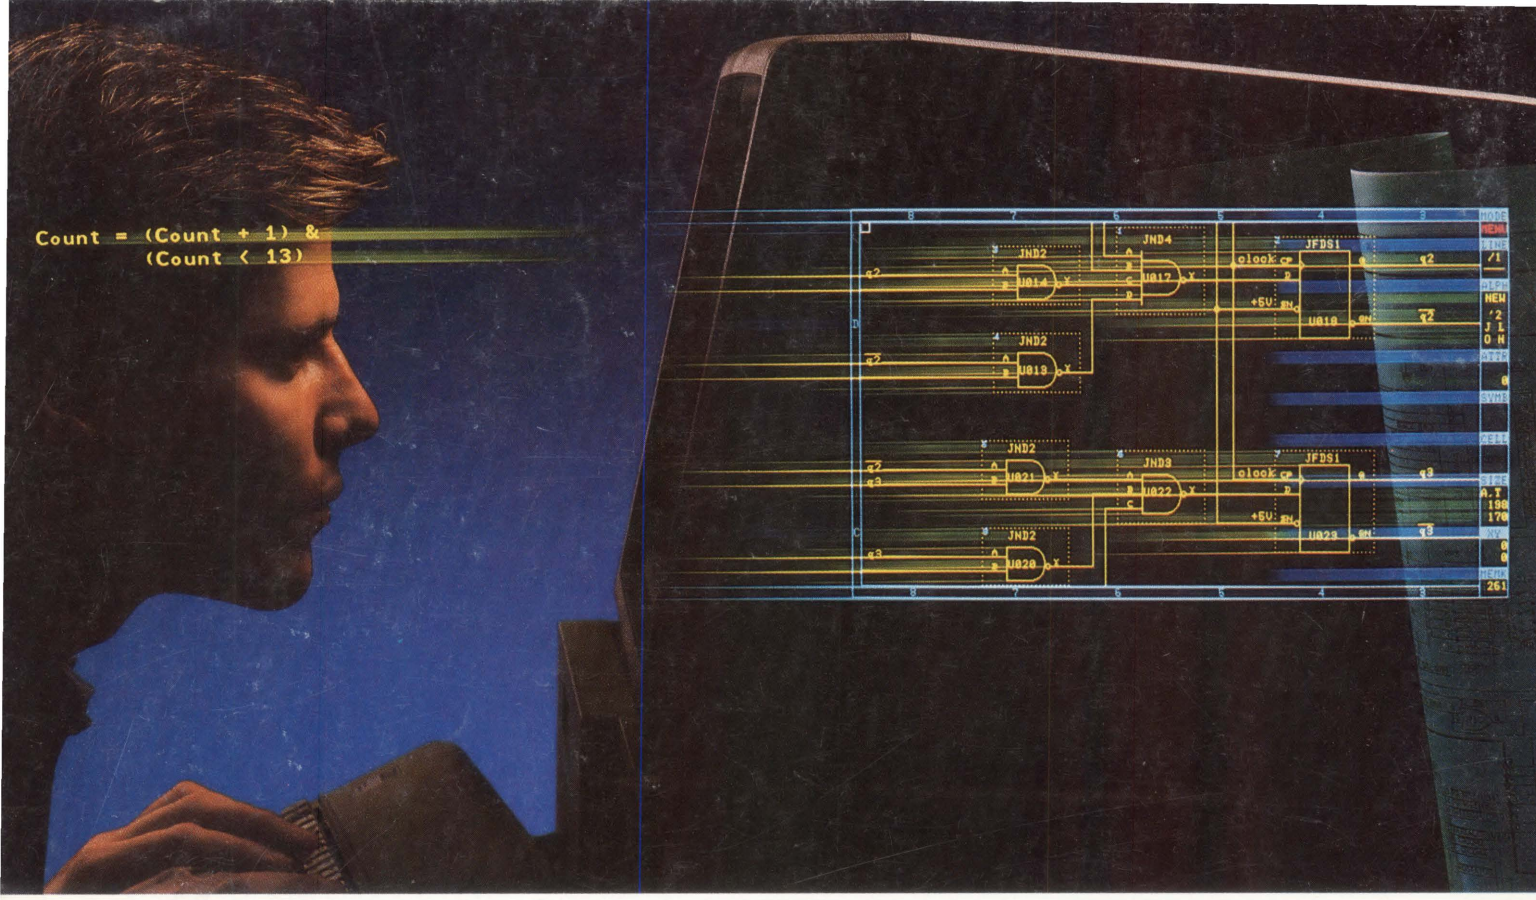

AUTOMATIC SCHEMATIC.

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

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

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

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

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

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

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

1-800-247-5700
Dept. 107

Data I/O Corporation 10525 Willows Road N.E., P.O. Box 97046, Redmond, WA 98073-9746, U.S.A. (206) 881-6444/Telex 15-2167
FutureNet 9310 Topanga Canyon Boulevard, Chatsworth, CA 91311-7528 (818) 700-0691/Telex 910-494-2681
Data I/O Canada 6725 Airport Road, Suite 302, Mississauga, Ontario L4V 1V2 (416) 676-0761/06968133
Data I/O Europe World Trade Center, Strawinskylaan 633, 1077 XX Amsterdam, The Netherlands (20) 622866/Telex 16616 DATIO NL
Data I/O Japan Sumitomoseimei Higashishinbashi Bldg., 8F, 2-1-7, Higashi-Shinbashi, Minato-ku, Tokyo 105, Japan (03) 432-6991/Telex 2522685 DATAIO J

FutureNet
A Data I/O Company

CIRCLE NUMBER 15