A PENNWELL PUBLICATION

COMPUTER DESIGN®

JUNE 1, 1989

THE FIRST MAGAZINE OF SYSTEM DESIGN, DEVELOPMENT AND INTEGRATION

Negotiating the obstacles to building military computers

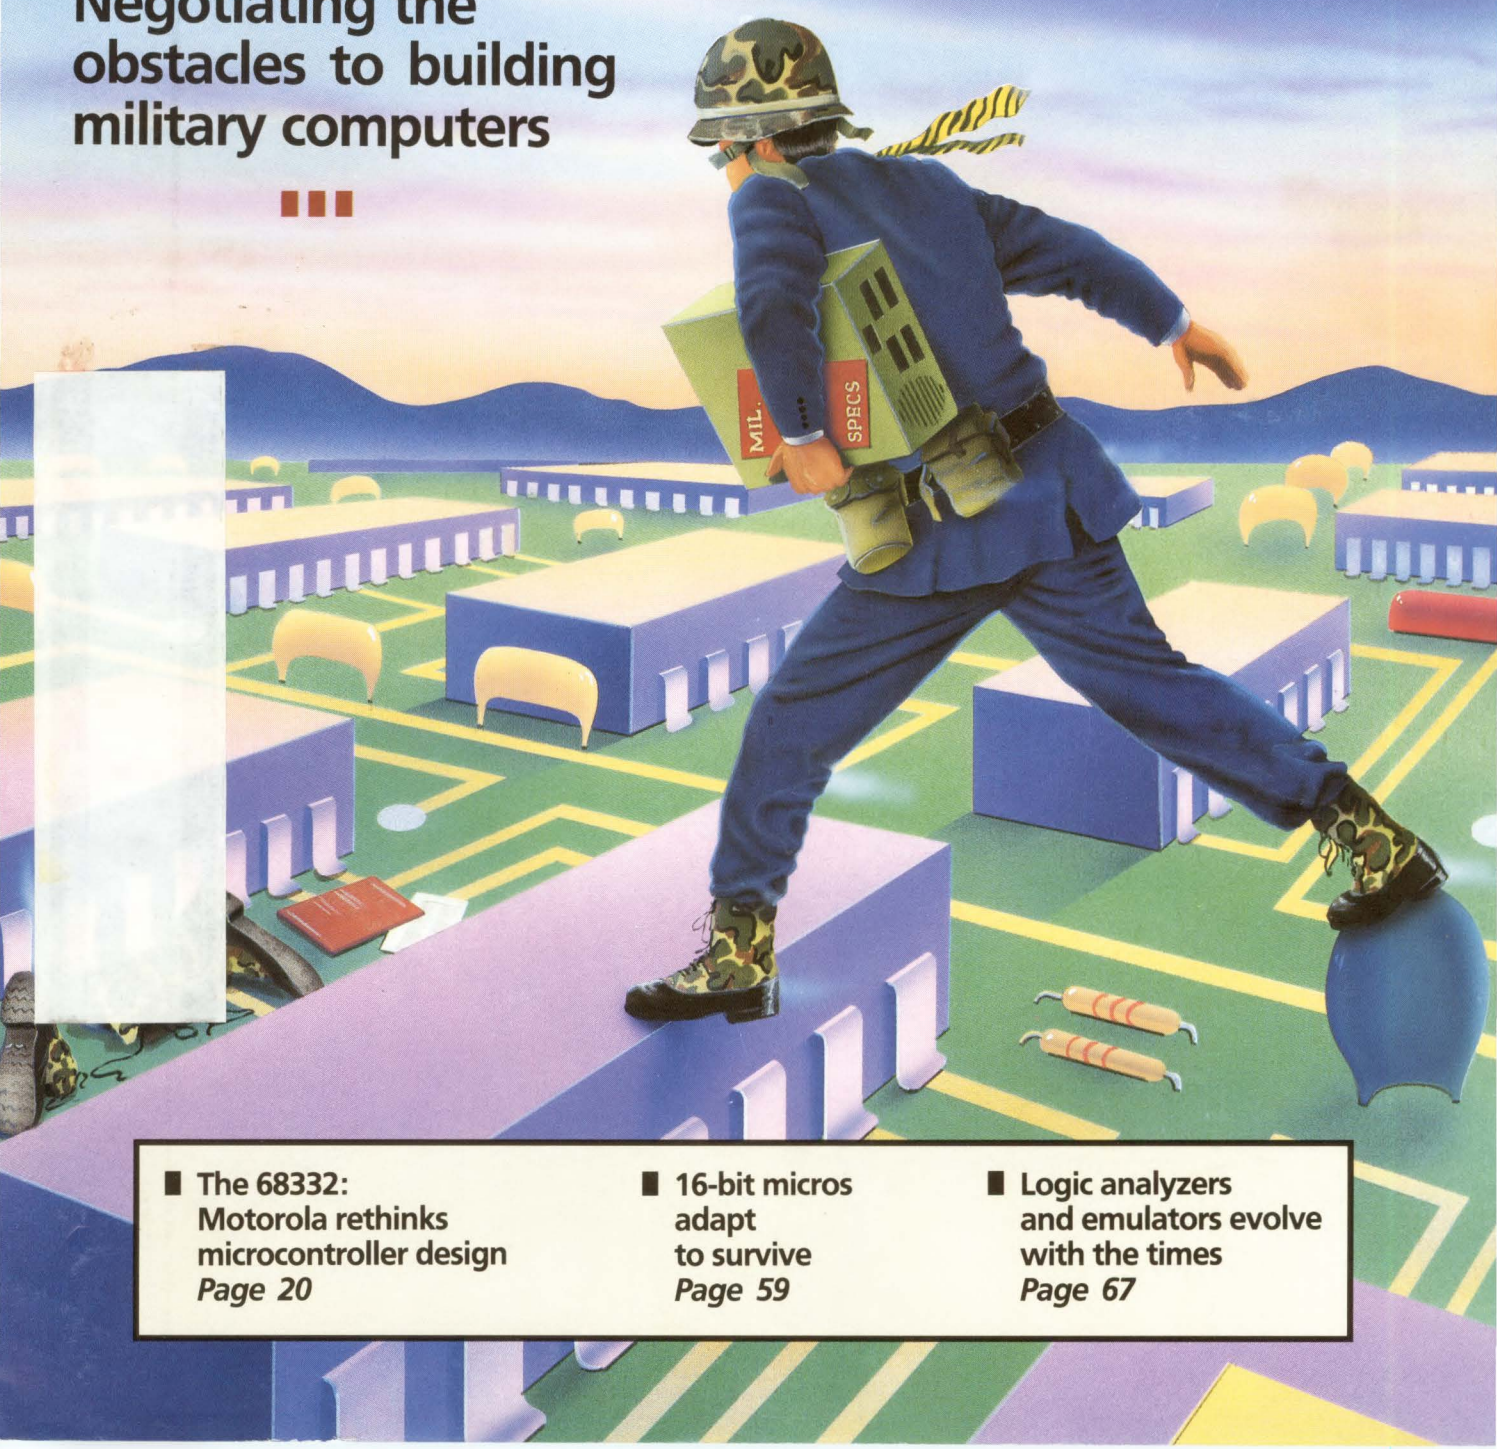

- The 68332: Motorola rethinks microcontroller design *Page 20*
- 16-bit micros adapt to survive *Page 59*
- Logic analyzers and emulators evolve with the times *Page 67*

WANT TO TALK SCSI? CALL CIPRICO AT 1-800-SCSI-NOW.

Why talk SCSI with Ciprico? To start with, we're the only vendor with a complete line of high-performance SCSI host bus adapters for Multibus® I, VMEbus, and Multibus II. Each board was designed to optimize performance with its system bus. And consider our experience. Ciprico has over 50,000 boards installed worldwide. Our design expertise provides you with the highest possible performance at the lowest possible price.

But that's not all. Ciprico has *respondability*. We have the largest staff of customer support engineers in the industry to help you resolve any integration problems. Plus software drivers for all major operating systems.

Our SCSI adapters are all based on an 80186 microprocessor, providing a command queuing, pass-through SCSI command software interface. Each board supports SCSI disconnect/reconnect and asynchronous and synchronous transfer rates of 2MB/s and 5MB/s respectively. And a floppy port is optional for cost effective system design.

So if you're currently designing a system based on Multibus I, VMEbus, or Multibus II, give us a call to talk SCSI.

RIMFIRE 3510
SCSI Adapter for
VMEbus

RIMFIRE 1500
SCSI Adapter for
Multibus® I

RIMFIRE 2500
SCSI Adapter for
Multibus® II

CIPRICO

2955 Xenium Lane
Plymouth, MN 55441

CIPRICO LISTENS. AND RESPONDS.

Multibus is a registered trademark of Intel Corporation.

CIRCLE NO. 26

MaxChassis.

A New VME Enclosure with Benefits That Really Stack Up:

40 Slot Capacity. Big systems in less space. In its maximum configuration, MaxChassis has *two* 20-slot VME card cages and two power supplies—two complete 20-slot VME systems in a single package only 14″ high, 26″ deep. One takes cards from the front, the other from the rear.

Match your needs exactly. There are 12, 20, and 32-slot systems, too. Tabletop or rack mount. The 12 and 20 slot models are only 20″ deep for VME or Multibus II, but you can get triple height (9U) cages and card depths to 400mm. There are three power supply choices from 400 to 800 Watts.

Ready to play. All MaxChassis come fully wired, complete with high-performance VME or Multibus backplanes. Just add your cards.

Disk drive modularity. The 12 and 32-slot models have slide-in peripheral mounting assemblies that hold up to four half-height drives. No brackets to design, no harnessing to make.

RFI tight. There's a front cover panel that holds RFI in (and out) so your system will be looked on approvingly by the FCC. The cover panel covers all those cables and connectors on the card fronts, so everyone else will look on your system approvingly, too.

So before you stack this magazine, phone for our MaxChassis facts right now. And ask for our complete catalog of VME, Multibus and Multibus II system packaging.

We'll FAX you the facts.

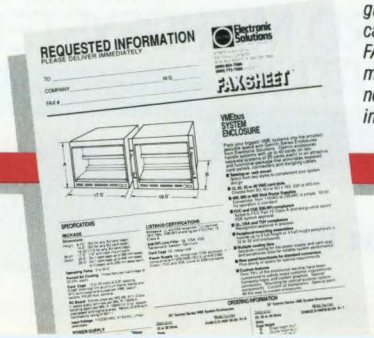

Want the latest data in a hurry? Nothing is faster than Electronic Solutions' new "FAX the FACTS" program. If you have a FAX machine, just call our "800" number, give us your FAX number and type of FAX machine, and the information you need from us. We'll FAX it to you immediately.

ZERO **Electronic Solutions**
UNIT OF ZERO CORPORATION

6790 Flanders Drive, San Diego, CA 92121 · (619) 452-9333 Telex II(TWX): 910-335-1169

Call Toll Free: (800)854-7086 In Calif: **(800)772-7086**

CIRCLE NO. 1

SMC®'s adaptable SuperCells™ get you over design humps on the road to mass storage.

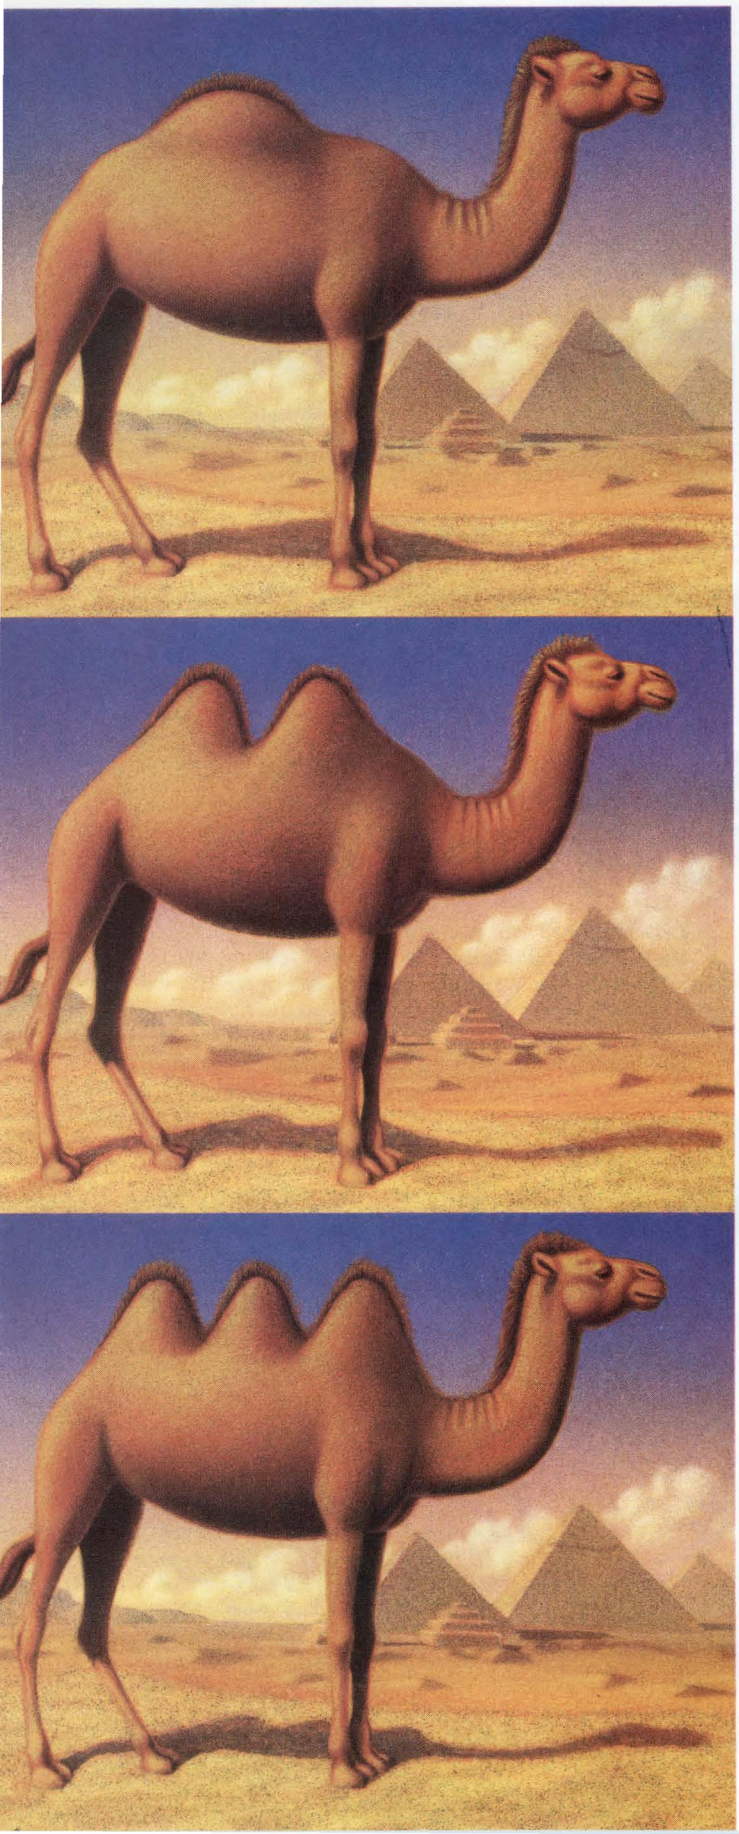

SMC is the perfect partner to help you over the long haul towards more efficient control of your mass storage devices. That's because SMC's powerful ASIC tools let you use our flexible VLSI storage products either off the shelf, or as the basis for your own unique design.

A perfect example: the MSD95C01/02 Storage Micro-Controller. It's made up of SuperCells which may be easily modified or replaced for your specific application. Based on a programmable micro-sequencer which allows you to choose any data format along with a selectable ECC, variable microprocessor interface, buffer control and MFM or RLL encoder/decoders, this chip provides the flexibility to control data flow to a hard disk, floppy disk, optical disk or magnetic tape.

Interfacing to Micro Channel™? Our MCI94C18 supports two DMA channels with an on-chip arbitration, timer and much more. And if your Micro Channel destination is ESDI, our MSD95C10 MIDAS (Micro Channel-to-Disk Access) is the perfect companion to the MSD95C01/02. And our MSD95C00 SCSI controller provides just what you need to interface your storage device to SCSI.

Working with floppy disk? The FDC37C65 is the right part . . . or the right SuperCell . . . for you. Built around the industry standard 765A floppy disk controller, the FDC37C65 adds SMC's patented digital data separator (analog data separator also available!), as well as the write precompensation generator and glue logic required for floppy disk sub-systems.

Let SMC get you over design humps at prices that won't break the camel's back.

SMC® is a registered trademark and SuperCell™, Midas™ are trademarks of Standard Microsystems Corporation.
Micro Channel™ is a trademark of International Business Machines Corporation.

Yes! I'd like to know more. CD/C

☐ Please send me information on _____

☐ Please send me your most recent Short Form Product Catalog.

☐ My need is immediate. Please have a salesperson call.

NAME _____

TITLE _____

COMPANY _____

ADDRESS _____

CITY _____ STATE _____ ZIP _____

TELEPHONE _____

STANDARD MICROSYSTEMS CORPORATION

35 Marcus Boulevard, Hauppauge, NY 11788 516-273-3100
Attn: IC Marketing

SMC—The Better Breed Of ASICs.

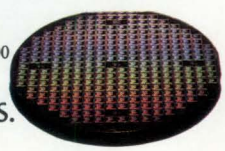

CIRCLE NO. 2

VOL. 28, NO. 11

COMPUTER DESIGN

JUNE 1, 1989

THE FIRST MAGAZINE OF SYSTEM DESIGN, DEVELOPMENT AND INTEGRATION

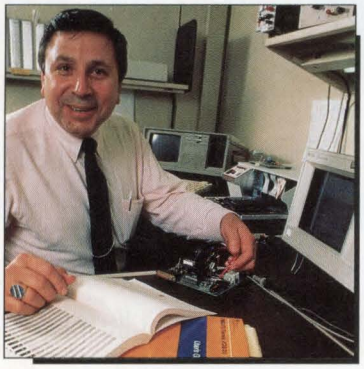

Page 26

TECHNOLOGY UPDATES

Integrated Circuits

Motorola rethinks the fundamentals for one-chip 32-bit microcontroller 20

Vendors eye flash EPROM for role in one-chip microcomputers 26

Workstation chip set bridges Micro Channel to RISC and CISC CPUs.....27

Hardwired language processors enjoy renewed attention 32

Software

Real-time development aids extend beyond kernels 36

Graphics/audio software simplifies man-machine interface in real time.....42

Design and Development Tools

Pipes, hubs and frameworks vie for tool integration 46

Major System Components

Coming soon: flexible film in large, flat displays 49

Pseudo-grayscale scheme generates 16 colors on active-matrix display 54

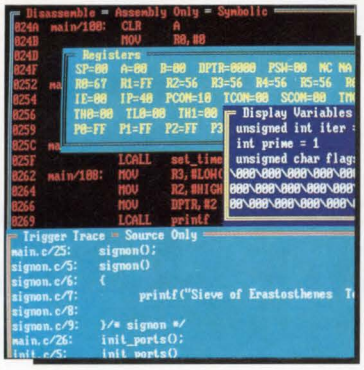

Page 67

TECHNOLOGY AND DESIGN FEATURES

16-bit micros adapt to survive

Trapped between single-chip microcontrollers and 32-bit microprocessors, 16-bit CPUs are fighting for survival. To succeed, they must learn their enemies' tricks 59

Logic analyzers and emulators evolve with the times

Racing to keep up with the demands of today's breed of high-performance systems, logic analyzers and emulators take on new roles and capabilities 67

COVER STORY

Negotiating the obstacles to building military computers

Designers of military systems are doing some fancy footwork to maneuver around unsupportive suppliers and stringent DOD demands. 78

DEPARTMENTS

News Briefs 8
Editorial.....14
Advertisers Index 120

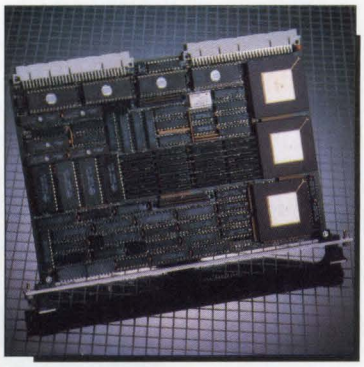

Page 96

NEW PRODUCT HIGHLIGHTS

Integrated Circuits

AMD offering joins second generation of disk-controller chips 93

Chip set offers IBM-compatible, Micro Channel-ESDI disk controllers 94

Computers and Subsystems

VME display modules build workstation graphics.....96

Data-acquisition speed approaches NuBus' theoretical limit 97

Analog I/O boards deliver high performance on PS/2 100

Image processor provides fast link to other boards 102

Intelligent VMEbus board controls up to six stepper axes 105

Design and Development Tools

Simulation systems combine behavioral and gate-/switch-level tasks.....108

Logic synthesizer optimizes IC designs for area and speed 110

Single simulator handles mixed analog/digital chip designs 112

Major System Components

Module brings removable storage to the VMEbus 115

Controller and TCP/IP software connect G-64 systems to Ethernet.....116

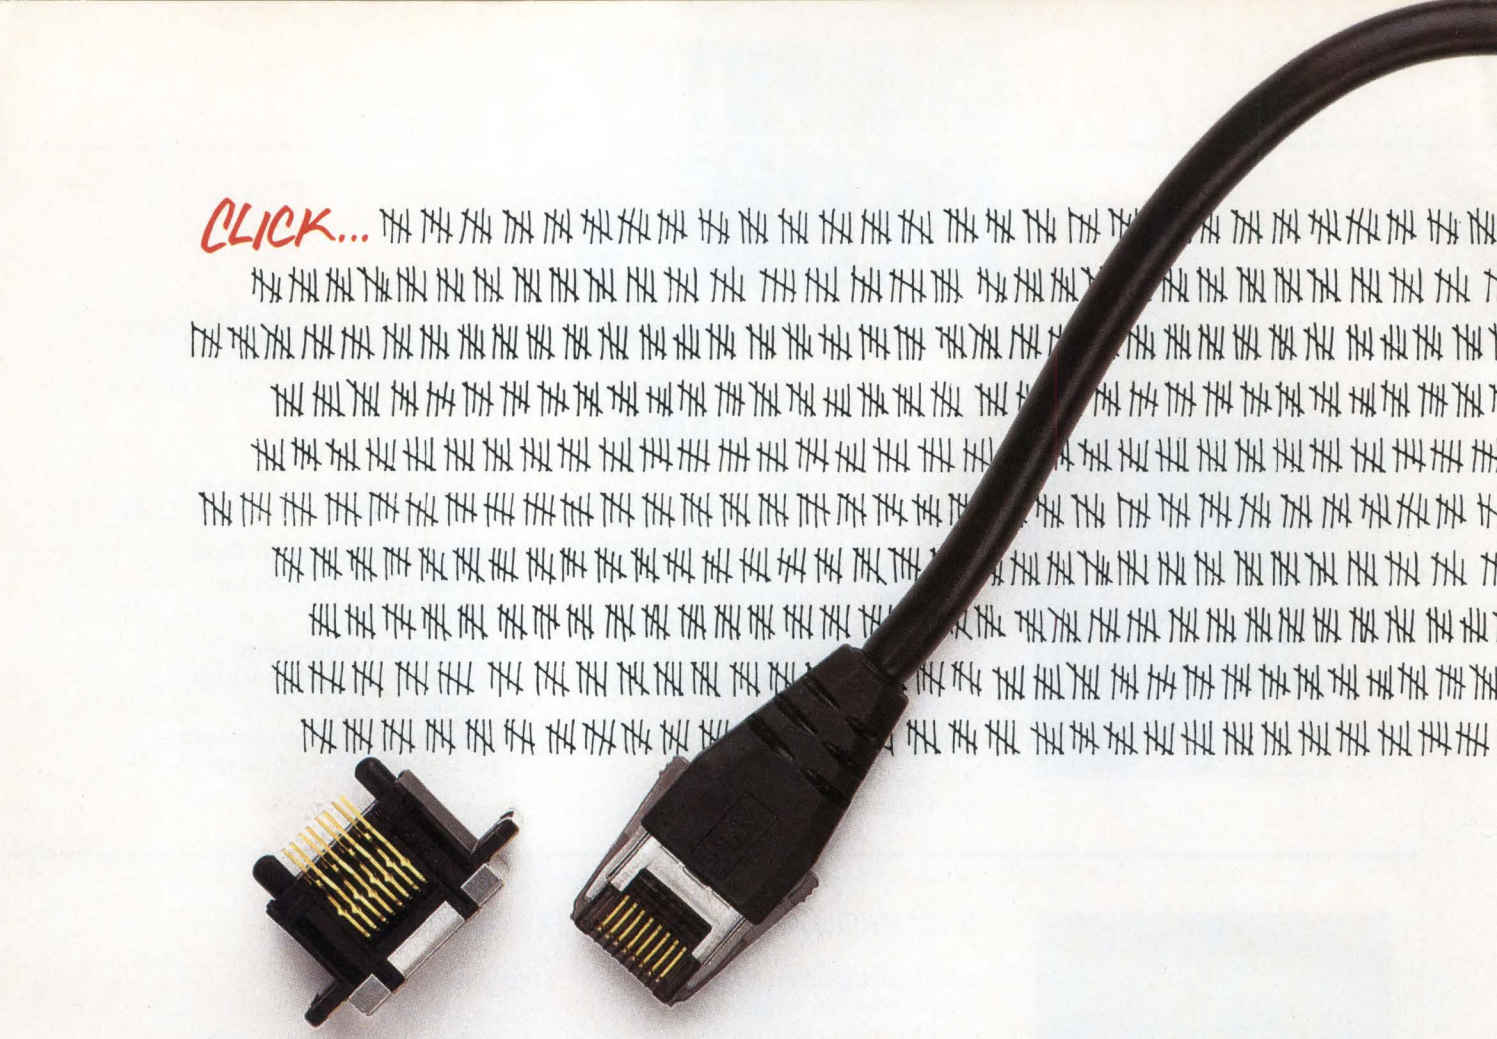

Great ideas tend

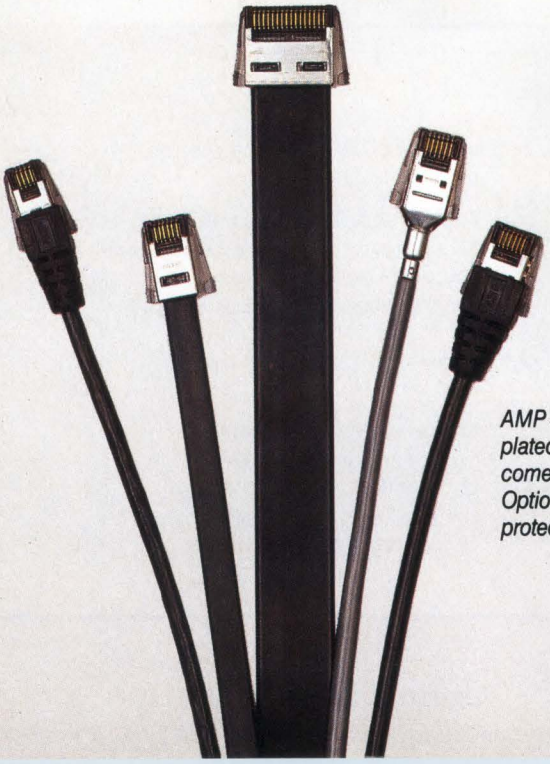

They're simple. They're shielded. They go in with a solid click, and release with a gentle squeeze. Keying and polarization are implicit.

And they're very, *very* durable. In fact, our Shielded Data Link (SDL) connectors are rated at 3000 cycles. They owe their toughness to design and engineering, and attention to such details as selective gold plating. Dependability is the essence of friendliness.

But good-looking and hard-working start well before the user sees this connector. The highly effective shield is designed in, so our SDL connector comes as a no-hassle, one-

AMP SDL connectors—selectively gold plated for reliability and durability—come in 4, 6, 8, and 16 positions. Optional boot for round cable version protects against electro-static discharge.

Offset flange receptacle makes good use of real estate.

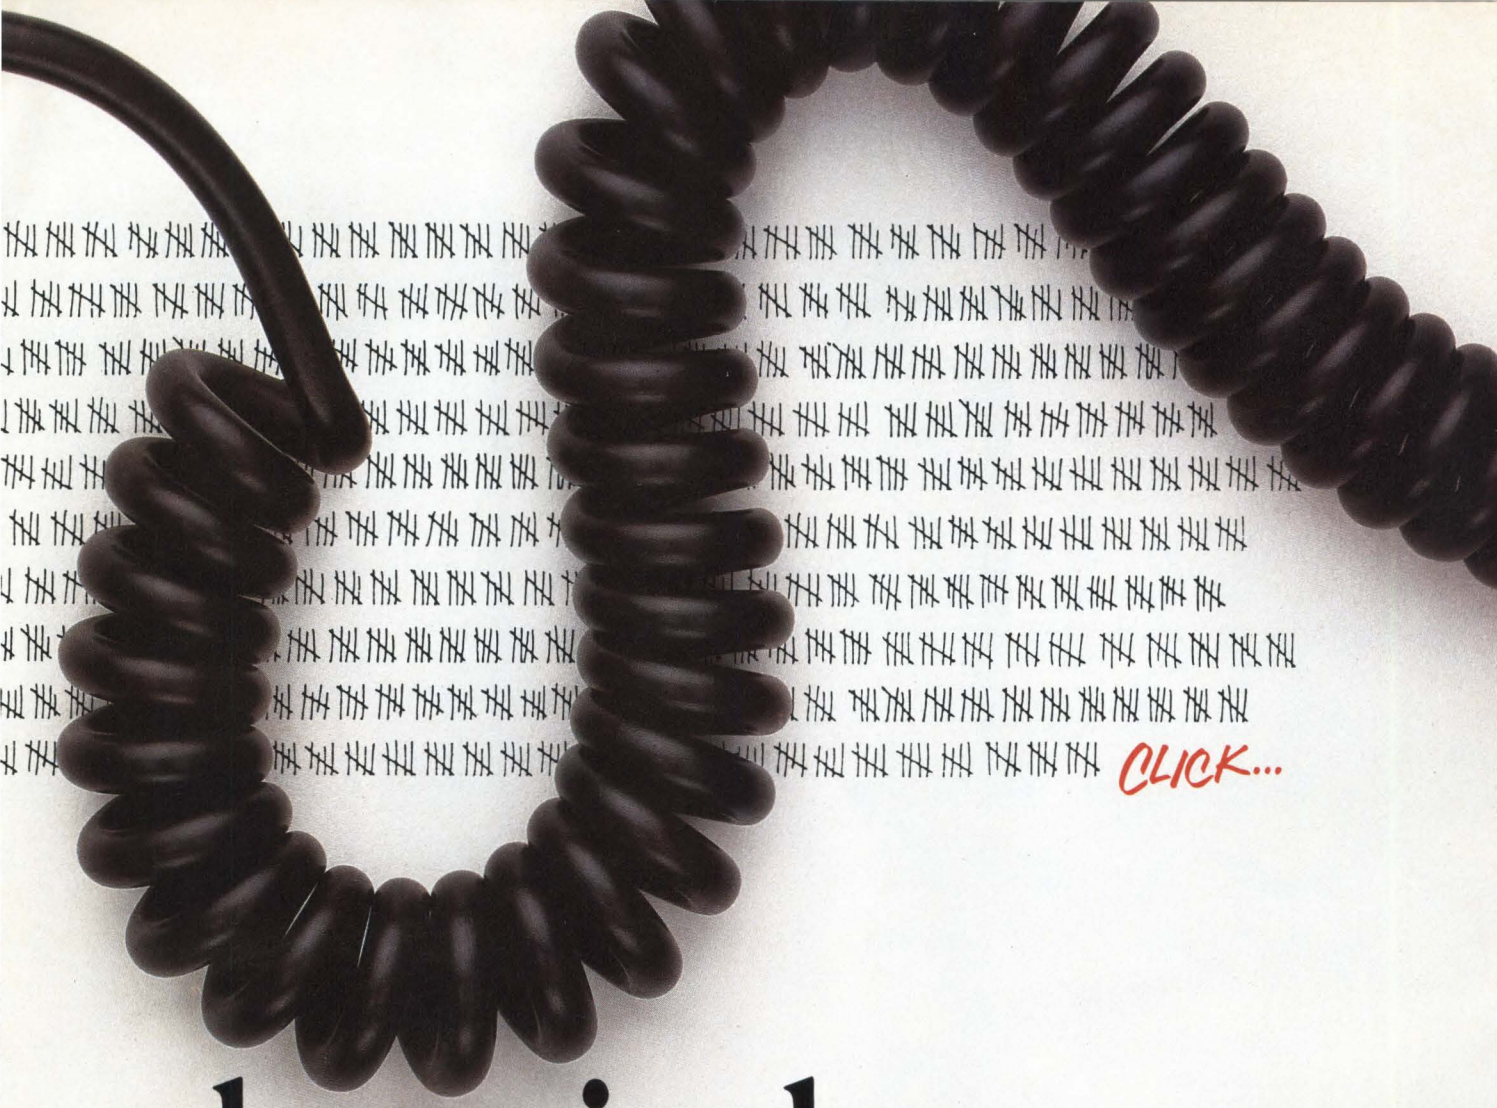

to hang in there.

piece assembly. It terminates to flat shielded cable in a single step—conductors and shield, plus two-point strain relief. Round cable takes only two steps.

AMP SDL connectors are a cost-effective alternative to crimp-snap type products, whether you do-it-yourself or order custom cables from us in flat, round, or coiled styles. Either way, you benefit from the same ease of application, the same well-thought-out design. And you walk away with the reliability and durability you need.

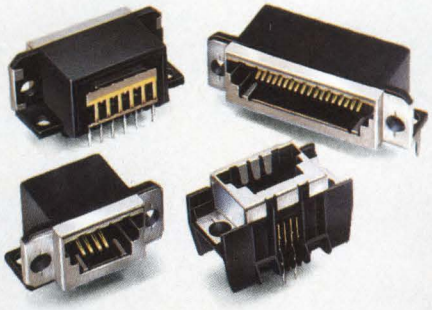

Receptacle options: top and side entry, and shunted side entry versions for network applications.

For technical information and literature on SDL connectors, call 1-800-522-6752. AMP Incorporated, Harrisburg, PA 17105-3608.

AMP Interconnecting ideas

CIRCLE NO. 3

"Did you know that new 1Mb Video RAMs increase throughput

ALABAMA, Marshall Electronics Group, (205) 881-9235, Milgray Electronics, Inc., (404) 393-9666, Reptron Electronics, (404) 446-1300; **ARKANSAS**, Marshall Electronics Group, (214) 233-5200, Milgray Electronics, Inc., (214) 248-1603, Sterling Electronics-Dallas, (214) 243-1600; **ARIZONA**, Marshall Electronics Group, (602) 496-0290, Sterling Electronics-Phoenix, (602) 268-2121; **CALIFORNIA**, Image Electronics, (714) 259-0900, (818) 407-8850, Marshall Electronics Group, (818) 407-0101, (714) 458-5395, (408) 942-4600, (916) 635-9700, (619) 578-9600, Merit Electronics, Inc., (408) 434-0800, Western Microtechnology, (619) 453-8430, (714) 637-0200, (408) 725-1660, (818) 707-0377; **COLORADO**, Marshall Electronics Group, (303) 451-8383, Sterling Electronics-Denver, (303) 792-3939; **CONNECTICUT**, Cronin Electronics, Inc., (203) 265-3134, Marshall Electronics Group, (203) 265-3822, Sterling Electronics-Wallingford, (203) 265-9535; **DELAWARE**, General Components, Inc., (609) 768-6767, Marshall Electronics Group, (301) 840-9450, Milgray/Delaware Valley, Inc., (609) 983-5010, (800) 257-7808, (800) 257-7111, Sterling Electronics-Sterling, (703) 450-2373; **DISTRICT OF COLUMBIA**, Marshall Electronics Group, (301) 840-9450, Milgray/Washington, Inc., (301) 621-8169, (800) 638-6656, Sterling Electronics-Sterling, (703) 450-2373; **FLORIDA**, Marshall Electronics Group, (813) 573-1399, (305) 977-4880, (305) 767-8585, Reptron Electronics-Tampa, (813) 855-4656, (305) 735-1112; **GEORGIA**, Marshall Electronics Group, (404) 923-5750, Milgray Electronics, Inc., (404) 446-9777, Reptron, (404) 446-1300; **IDAHO**, Marshall Electronics Group, (801) 485-1551, Western Microtechnology, (503) 629-2082, Milgray Electronics, Inc., (801) 272-4999; **ILLINOIS**, Goold Electronics, (312) 860-7171, Marshall Electronics Group, (312) 490-0155, (314) 291-8554, Milgray Electronics, Inc., (913) 236-8800, (312) 350-0490; **INDIANA**, Goold Electronics, (312) 593-3220, Marshall Electronics Group, (317) 297-0483, Milgray/Chicago, Inc., (312) 350-0490; **IOWA**, Marshall Electronics Group, (612) 559-2211, Milgray Electronics, Inc., (913) 236-8800; **KANSAS**, Marshall Electronics Group, (913) 492-3121, Milgray Electronics, Inc., (913) 236-8800, Sterling Electronics-Kansas, (913) 492-5406; **KENTUCKY**, Marshall Electronics Group, (513) 898-4480, Milgray/Cleveland, Inc., (216) 447-1520, (800) 321-0006; **LOUISIANA**, Marshall Electronics Group, (214) 233-5200, Milgray Electronics, Inc., (214) 248-1603, Sterling Electronics-Dallas, (214) 243-1600; **MAINE**, Cronin Electronics, Inc., (617) 449-5000, Marshall Electronics Group, (617) 658-0810, Sterling Electronics-Boston, (617) 938-6200, Western Microtechnology, (617) 273-2800; **MARYLAND**, Marshall Electronics Group, (301) 622-1118, Milgray/Washington, Inc., (301) 621-8169, (800) 638-6656, Sterling Electronics-Sterling, (703) 450-2373; **MASSACHUSETTS**, Cronin Electronics, Inc., (617) 449-5000, Marshall Electronics Group, (617) 658-0810, Sterling Electronics-Boston, (617) 938-6200, Western Microtechnology, (617) 273-2800; **MICHIGAN**, Calder Electronics, (616) 698-7400, Marshall Electronics Group, (313) 525-5850, Reptron Electronics, Inc., (313) 525-2700; **MINNESOTA**, Marshall Electronics Group, (612) 559-2211, Reptron Electronics, (612) 938-0000, Sterling Electronics, (612) 831-2666; **MISSISSIPPI**, Marshall Electronics Group, (205) 881-9235, Milgray Electronics, Inc., (404) 393-9666; **MISSOURI**, Marshall Electronics Group, (913) 492-3121, (314) 291-8554, Milgray Electronics, Inc., (913) 236-8800, Sterling Electronics/Kansas, (913) 236-5589; **NEBRASKA**, Marshall Electronics Group, (913) 492-3121,

by over 100%?"

"1Mb Video RAM?"

"Yeah. With the fancy new Flash Write, Block Write and Split Transfer functions to improve performance."

"Not bad."

"And greater density for high resolution color graphics."

"Come on. We need something we can use today."

"Toshiba's got 'em. In quantity."

First the good news. Toshiba, the leader in 1Mb DRAMs, introduces a new high density 1Mb Video RAM for today's high resolution color graphic applications.

Now the great news. It's available in volume production, giving you access to the quantities you need. Today. With capabilities that'll take you into tomorrow.

Toshiba's high density, 256K x 4bit and 128K x 8bit Video RAM supports next generation video board or graphic terminal applications by allowing you to store, manipulate and transmit more information than ever before.

A multi-port feature combining RAM and SAM gives you easy and fast access to data. Bit Masking is achieved through the Write Per Bit feature.

	ORGANIZATION		SPECIAL FEATURES				PACKAGE	
	X4	X8	Fast Page Mode	Flash Write	Block Write	Split Transfer	SOJ	ZIP
TC524256A	X		X				X	X
TC524258A	X		X	X	X	X	X	X
TC528126A		X	X				X	
TC528128A		X	X	X	X	X	X	

With this feature, selected planes of data can be modified using just a Write cycle instead of a Read-Modify-Write cycle.

Use the Block Write feature to write four times faster and simplify high speed window clear/fill. Use the Flash Write feature to accomplish high speed clear of selected planes without using the SAM register.

To speed things along even more, RAM read/write cycle times have been improved through fast page mode. A split SAM register provides the best method to simplify SAM to RAM and RAM to SAM transfer timing in the system.

TC524256A/8A & TC528126A/8A	–10	–12
RAS ACCESS TIME (MAX)	100NS	120NS
COLUMN ADDRESS ACCESS TIME (MAX)	30NS	35NS
FAST PAGE MODE CYCLE TIME (MIN)	60NS	70NS
SERIAL ACCESS TIME (MAX)	25NS	35NS
SERIAL CYCLE TIME (MIN)	30NS	40NS

Want more good news? Toshiba's service is every bit as good as their new Video RAM. With the people and the design specification information you need to get the job done. Call for your complete data sheet today. (1-800-888-0848 ext. 517).

In Touch with Tomorrow

TOSHIBA

TOSHIBA AMERICA ELECTRONIC COMPONENTS, INC.

© 1989 Toshiba America, Inc.

MST-88-009

Milgray Electronics, Inc., (913) 236-8800, Sterling Electronics-Kansas, (913) 236-5589; **NEVADA**, Marshall Electronics Group, (916) 635-9700, (602) 496-0290, Sterling Electronics, (602) 268-2121, Western Microtechnology, (408) 725-1660; **NEW HAMPSHIRE**, Cronin Electronics, Inc., (617) 449-5000, Marshall Electronics Group, (617) 658-0810, Sterling Electronics-Boston, (617) 938-6200, Western Microtechnology, (617) 273-2800; **NEW JERSEY**, General Components, Inc., (609) 768-6767, Marshall Electronics Group, (201) 882-0320, (609) 234-9100, Milgray Electronics, Inc., (516) 420-9800, (609) 983-5010, (800) 257-7808, (800) 257-7111, Sterling Electronics-So. Plainfield, (201) 769-7000; **NEW MEXICO**, Marshall Electronics Group, (602) 496-0290, Sterling Electronics-Albuquerque, (505) 884-1900; **NEW YORK**, Marshall Electronics Group, (516) 273-2424, (607) 798-1611, (716) 235-7620, Milgray Electronics, Inc., (516) 420-9800, (716) 235-0830, Rome Electronics, (315) 337-5400; **NORTH/SOUTH CAROLINA**, Marshall Electronics Group, (919) 878-9882, Milgray Electronics, Inc., (404) 393-9666, Resco Raleigh, (919) 781-5700, (704) 588-1530; **NORTH/SOUTH DAKOTA**, Marshall Electronics Group, (612) 559-2211, Reptron Electronics, (612) 938-0000; **OHIO**, Marshall Electronics Group, (513) 898-4480, (216) 248-1788, (614) 891-7580, Milgray/Cleveland, Inc., (216) 447-1520, (800) 321-0006, Reptron Electronics, (614) 436-6675, (216) 349-1415; **OKLAHOMA**, Marshall Electronics Group, (214) 233-5200, Milgray Electronics, Inc., (214) 248-1603, Sterling Electronics-Tulsa, (918) 663-2410; **OREGON**, Marshall Electronics Group, (503) 644-5050, Western Microtechnology, (503) 629-2082; **PENNSYLVANIA**, General Components, Inc., (609) 768-6767, Marshall Electronics Group, (609) 234-9100, (412) 963-0441, Milgray/Cleveland, Inc., (216) 447-1520, (800) 321-0006, Milgray/Delaware Valley, Inc., (609) 983-5010, (800) 257-7808, (800) 257-7111; **RHODE ISLAND**, Cronin Electronics, Inc., (617) 449-5000, Marshall Electronics Group, (617) 658-0810, Sterling Electronics-Boston, (617) 938-6200, Western Microtechnology, (617) 273-2800; **TENNESSEE**, Marshall Electronics Group, (205) 881-9235, Milgray Electronics, Inc., (404) 393-9666, Resco Raleigh, (919) 781-5700; **TEXAS**, Marshall Electronics Group, (214) 233-5200, (915) 593-0706, (713) 895-9200, (512) 837-1991, Milgray Electronics, Inc., (214) 248-1603, Sterling Electronics, (713) 623-6600, (512) 836-1341, (214) 243-1600; **UTAH**, Marshall Electronics Group, (801) 485-1551, Sterling Electronics-Denver, (303) 792-3939, Milgray Electronics, (801) 272-4999; **VERMONT**, Cronin Electronics, Inc., (617) 449-5000, Marshall Electronics Group, (617) 658-0810, Sterling Electronics-Boston, (617) 938-6200, Western Microtechnology, (617) 273-2800; **VIRGINIA**, Marshall Electronics Group, (301) 840-9450, Milgray/Washington, Inc., (301) 621-8169, (800) 638-6656, Resco Raleigh, (919) 781-5700, Sterling Electronics, (703) 742-8400, (804) 226-2190; **WASHINGTON**, Marshall Electronics Group, (206) 486-5747, Western Microtechnology, (206) 881-6737; **WISCONSIN**, Gould Electronics, (312) 593-3220, Marshall Electronics Group, (414) 797-8400, Milgray/Chicago, Inc., (312) 350-0490, Reptron Electronics, (612) 938-0000; **CANADA**, Marshall Electronics Group, (416) 674-2161, ITT Multicomponents, (604) 291-8866, (403) 451-4001, (403) 253-8575, (306) 933-2888, (204) 786-8401, (416) 736-1144, (613) 226-7406, (514) 335-7679, (506) 857-8011, (902) 465-2350, Space Electronics Sales Corp., (514) 697-8676, (416) 636-8814, (613) 596-5340, (604) 294-1166.

NEWS BRIEFS NEWS BRIEFS NEWS BRIEFS NEWS BRIEF

ASIC vendors focus on testability

A growing number of ASIC vendors are talking about testability as a competitive issue rather than as merely a checklist item. Following the recent announcement of an agreement between LSI Logic (Milpitas, CA) and Crosscheck (San Jose, CA) to introduce a new testing technique into LSI Logic's products, other vendors are revealing their own strategies.

Texas Instruments (Dallas, TX), for example, unveiled 14 cells to implement IEEE P1149.1 Joint Testability Action Group testing procedures on TI cell-based products. The cells include an access port controller and the necessary switchable buffers to let sections of a circuit, or of a chip, be isolated and tested.

A different approach was announced by VLSI Technology (San Jose, CA), which rolled out a pair of compilers for creating self-test circuitry on that company's ASICs. Based on a linear-feedback shift register (LFSR) technique related to signature analysis, the company's compilers can generate pseudorandom patterns, accumulate a series of circuit outputs and compress them into a unique pattern, and then test received patterns against expected ones. One compiler generates the LFSR circuitry, and a second compiler uses the same technology to generate specific test circuits for RAM cells. In addition, the company announced a Test Assistant tool that creates test programs and suggests modifications to circuitry to improve testability.

The appearance of testability as a competitive issue could be a win-win situation for ASIC vendors and customers, giving the former an arena for competition in an increasingly jelly-bean market, and the latter a needed assist in maintaining product quality. —*Ron Wilson*

Touch screen uses CRT beam

A joint research project by IBM's Watson Research Center (Yorktown Heights, NY) and the National University of Singapore will explore a new touch-input technology that uses the CRT display beam rather than LEDs for user input. In a paper presented at the recent Society for Information Display conference, Isay Chang, manager of IBM's Applications Software Institute, discussed the new technology and compared it with existing touch-screen designs.

Chang and students from the university developed the prototype, which uses the overscan regions on a CRT screen to create light bars that are reflected by prisms to detectors at the corners of the display. When the beam is broken by the user, the sensors signal a controller to take the appropriate action. Future designs may locate the light bars anywhere on the screen, with the sensors positioned inside the CRT for compact designs. —*Michael Donlin*

TI cleans up on DRAM patent issues

Several years ago, Texas Instruments (Dallas, TX) successfully forced a number of Japanese and other offshore companies to pay royalties for use of TI's basic dynamic RAM technology, which TI had patented. Now, the largest domestic maker of DRAMs, Micron Technology (Boise, ID), has agreed to settle out of court and pay TI more than $38 million up front as well as ongoing royalties.

One reason that Micron was forced to cough up so much was that TI had also filed suit charging infringement on video RAM technology. While the out-of-court settlement has caused a shakeup among Micron's executives, it's not yet clear what effects the action will have on overall DRAM prices and availability. —*Tom Williams*

Motorola joins design-automation fray

In what could almost be described as a beta test of the CAD Framework Initiative—a user/vendor group trying to get different CAD tools to work together—Motorola's ASIC Division (Chandler, AZ) recently introduced its own line of design automation tools.

Called the Motorola-OACS (Open Architecture CAD System), Motorola's offering is built mostly around third-party design tools from companies such as Gateway Design Automation, EDA Systems, Mentor Graphics, Cadence and Valid Logic Systems. Motorola's own contributions include the Mustang ATPG (automatic test program generation) and Trailblazer digital timing analysis system.

Motorola's close ties to the CAD Framework Initiative (Motorola's Andy Graham is the president of CFI) had some influence on the direction that the company took in developing its design environment. In addition to integrating tools from a variety of companies, Motorola-OACS uses a framework approach to tie the tools together, makes extensive use of EDIF (Electronic Design Interchange Format) netlists, and runs on both Apollo and Sun workstations.

To reinforce its position as a supplier of design automation tools, Motorola will exhibit at the Design Automation Conference for the first time this year. —*Bill Harding*

First mixed-mode fax shown for ISDN

The ability to transmit composite documents that consist of bit-map images and ASCII-coded data via facsimile has been demonstrated by Ricoh (West Caldwell, NJ). The demonstration at the International Communications Association exposition in Dallas used a new Ricoh CCITT Group 4, Class 3 mixed-mode fax adapter. Group 4 fax will be the fax mode for the integrated services digital network (ISDN). With the mixed-mode capability, users can not only transmit hard-copy documents, but can also send images and text from one computer to another, where the receiver can alter the image and edit text.

Today's LANs use the Group 4, Class 1 image-only protocol, which is much more efficient than the Group 3 fax used on the current public phone system. —*Tom Williams*

New memories undercut EPROM cost

Advanced Micro Devices (Sunnyvale, CA) is offering a new type of CMOS memory that can perform the same task as EPROMs at up to

(continued on page 10)

Who gave Sun a license to speed?

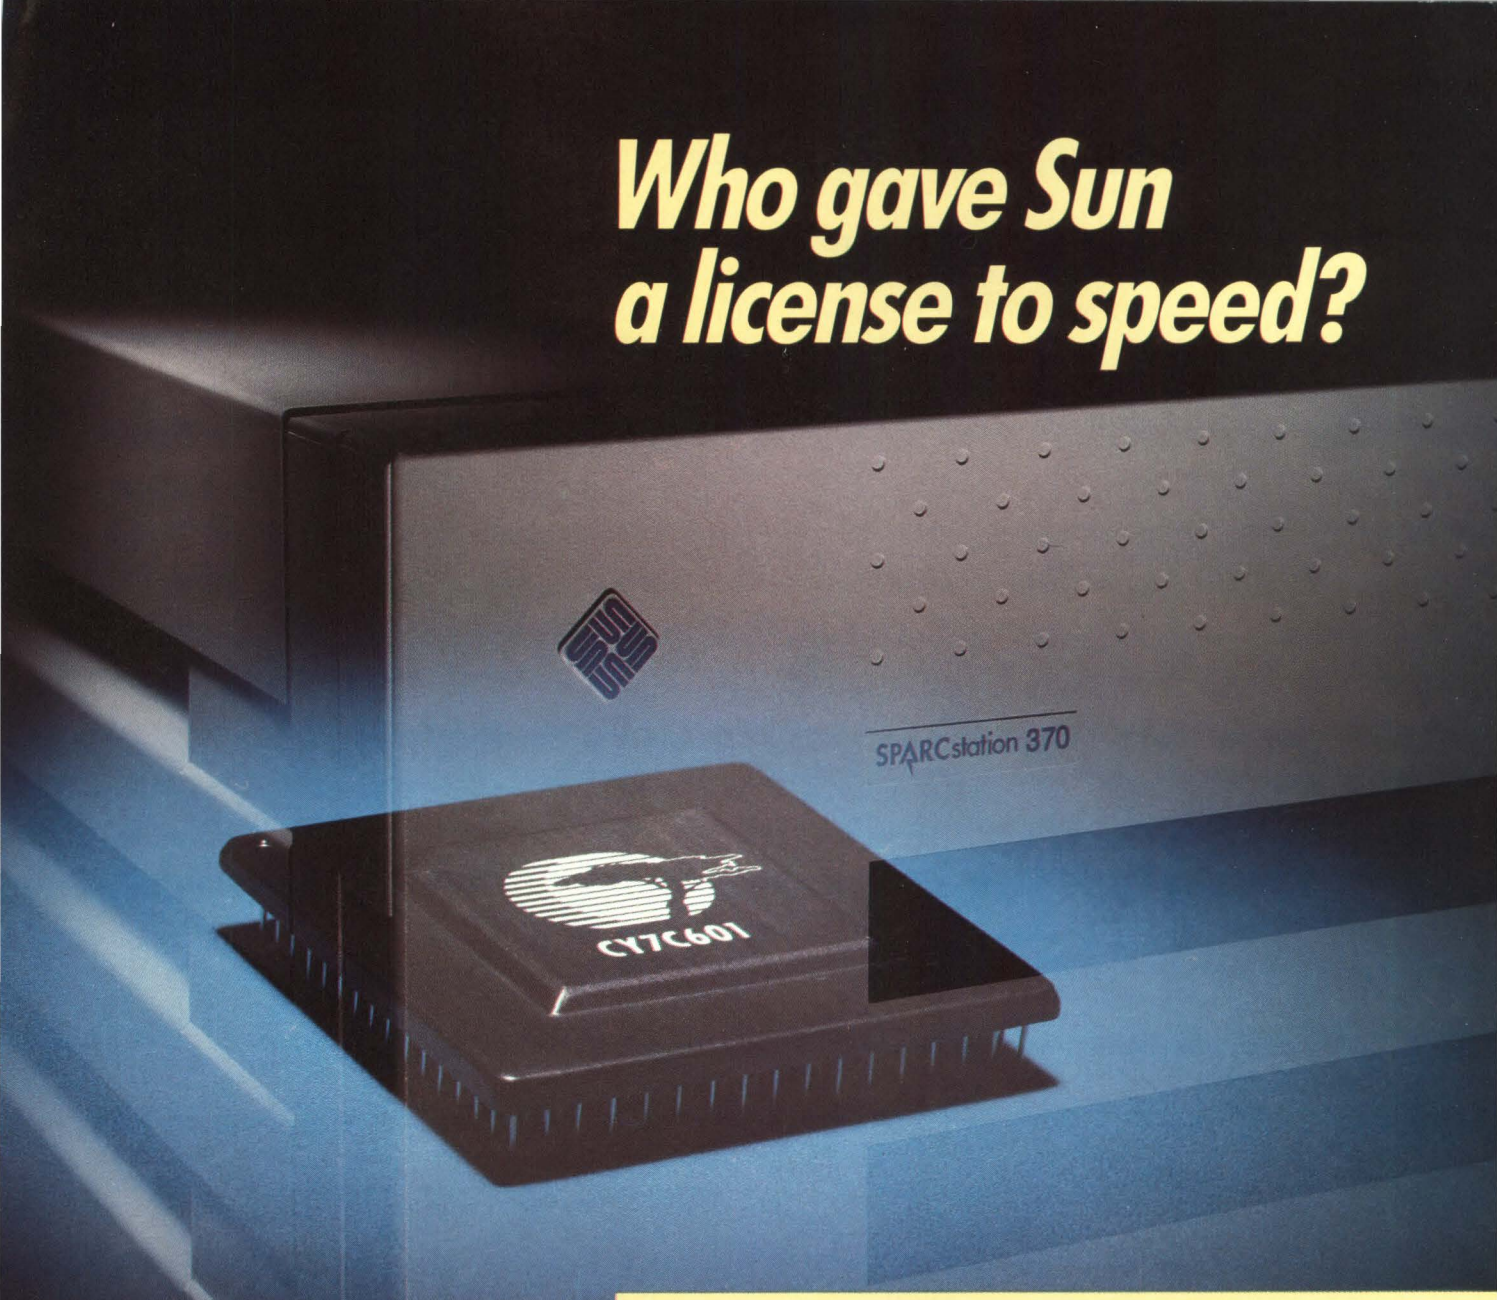

Hot Rod RISC.

Sun's new SPARCstations 330, 370 and 390 deliver 16 MIPS.

The engine? Our 7C601 RISC CPU, running at 25MHz.

An engine that is also available at 33MHz, for up to 24 MIPS. It is part of a complete RISC chipset that includes memory management, multiprocessing, floating point, and cache, all tailored to give you high performance at very competitive prices.

Plus access to the broadest library of applications, software, utilities, languages, graphics, and productivity software available under the Sun.

Cypress Hot Rod RISC 7C601. Another blast from our latest data book. Call for your copy.

Data Book Hotline:
1-800-952-6300
Ask for Dept.
C11A*

CYPRESS SEMICONDUCTOR

*1-(800) 387-7599 in Canada. (32)2-672-2220 in Europe. ©1989 Cypress Semiconductor, 3901 North First Street, San Jose, CA 95134 Phone: (408)943-2600, Telex: 821032 CYPRESS SNJ UD, TWX: 910-997-0753.

NEWS BRIEFS NEWS BRIEFS NEWS BRIEFS NEWS BRIEF

(continued from page 8)

25 percent lower cost. The new devices, called ExpressROMs, are plastic-packaged EPROMs that are programmed and tested in the factory to customers' specifications.

"ExpressROM devices are programmed as part of each wafer's normal test process, so they can be cost-effective at relatively low volumes," says Steve Grossman, director of marketing for nonvolatile memories at AMD.

The new parts, available in as little as four weeks from receipt of a customer's code verification, come in the same densities and pinouts as do AMD's line of standard EPROMs.

—John Mayer

LSI Logic expands Co-Design Environment with link to Mentor

In the latest move in its aggressive Co-Design Environment (CDE) program, LSI Logic (San Jose, CA) has forged a link with Mentor Graphics (Beaverton, OR). This is the first agreement between LSI Logic and a major, full-range electronic design automation vendor under the CDE program, which was established to integrate third-party design tools with LSI Logic's design tools.

The result of the agreement with Mentor is a new LSI Logic design tool, called CDE/MG, that runs on an Apollo workstation and links elements of LSI Logic's Modular Design Environment with Mentor's Idea Series of design tools running on Apollo workstations.

Part of the link includes LSI Logic ASIC libraries for Mentor's Quicksim, which will let Mentor users design ASICs on their Mentor workstations and simulate them at both the chip and system levels. LSI Logic's LSIM simulator (not to be confused with Silicon Compiler Systems' Lsim simulator) is then used for final ASIC simulation prior to final sign-off of the design. While both Quicksim and LSI Logic's LSIM should produce the same results, users must use LSIM for final sign-off prior to fabrication.

—Bill Harding

Multiprocessing board boosts STD Bus systems

A high-speed V40-based multiprocessing single-board computer to be introduced this month by Ziatech (San Luis Obispo, CA) may help maintain the position of STD Bus in the industrial-application marketplace. The 8-MHz ZT 8832 I/O Control Processor (ICP) can be used either as a stand-alone SBC or as an intelligent control coprocessor in an STD Bus system.

The boost in speed is provided through the sharing of 32 kbytes of memory so that the main processor can communicate with other memories directly instead of over the STD Bus. The shared memory is located on the ICP and is mapped into the memory space of the main system.

The ICP contains one main processor and several intelligent control processors and communicates with the master STD Bus processor. Because a single master processor controls the STD Bus backplane, there's no slowdown from competition by multiple master processors for bus access. As many ICPs as necessary can be used, with each ICP mapped into a 32-kbyte addressing space. As an option, up to seven ICPs can be mapped into a common memory space.

—Sydney F. Shapiro

Software tool manages behavioral data

Platform independence will be a common cry at the Design Automation Conference in Las Vegas later this month. Contributing to that trend will be a new tool from TSSI (Beaverton, OR) that should simplify management of behavioral data in the design-to-test environment.

TSSI's Wavemaker is a platform-, simulator- and tester-independent software package that creates and preserves the behavioral/functional data generated in simulation so that it can be used throughout the product development process.

One advantage Wavemaker will bring designers is a simpler way to compare different simulation runs. As an extension to TSSI's TDS system, Wavemaker provides four graphical editors for the TDS Waveform database. The editors allow interactive creation and modification of signal definitions, device timing, data patterns and full composite waveform displays. The interfaces can be displayed in different windows on the workstation screen.

—John Mayer

First patent granted for hot-electron transistor

Researchers at AT&T Bell Laboratories (Holmdel, NJ) have received the first patent issued for hot-electron, or ballistic electron, transistor technology. The basis for the patent is a device the Bell Labs team claims to have used to achieve 165-GHz operation at room temperature.

The patent was granted to A.F.J. Levi, member of the technical staff at Bell Labs, for the underlying physics of a bipolar transistor consisting of thin layers of indium and gallium arsenide. Other research teams have also been pursuing hot-electron technology, but most other devices have operated only at cryogenic temperatures.

—Ron Wilson

Perpendicular recording pushes data densities

The rigid-disk industry's race to pack more information onto a magnetic disk will make perpendicular recording a key technology in data storage markets, according to Ed Zschau, president and CEO of Censtor (San Jose, CA), a supplier of perpendicular recording heads and media. By storing data vertically rather than across the disk's surface (as in conventional recording), perpendicular recording increases the amount of data that can be written onto a disk.

Zschau reported to members of the International Disk Equipment and Materials Association that disk drives are available that reduce the flying height of the read/write heads to only 4 to 6 μin., compared to 8 μin. for conventional heads. Using this technology, the disk drives feature recording densities of up to 110 Mbits/in.^{2}, with future products holding 200 Mbits/in.^{2} using heads that operate in actual contact with the disk's surface.

Censtor, however, has been making predictions such as this for at least the past three years and has yet to produce a product.

—Michael Donlin

T A D P O L E

Why the TP32V is the best single board computer

• The Design •

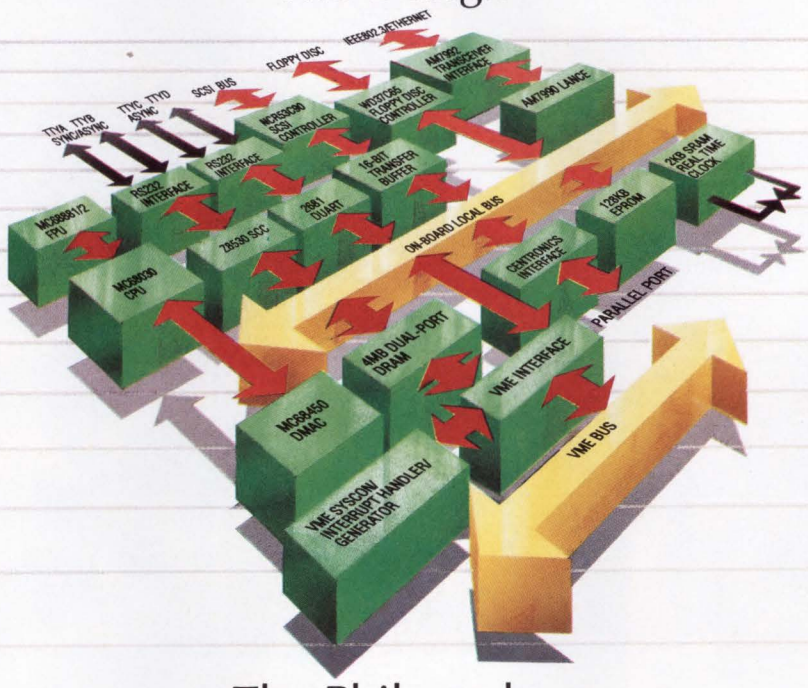

• The Philosophy •

For optimum system performance from a single IEE1014 VME board, the TP32V with its 16-33 Mhz MC 68030 has a tightly coupled processor memory design, 4Mb/sec sustained DMA driven SCSI transfers, on-board I/O and ethernet and advanced dram arbitration techniques.

• The Specification •

- MC68030 16-33MHz
- MC68450 4-channel DMA controller
- 4Mb multi-ported nibble-mode DRAM
- AMD Lance IEEE 802.3 Ethernet with DMA
- Z8530 SCC giving two DMA-driven RS232 sync/ asynchronous ports and two further RS232 asynchrous ports
- NCR 53C90 DMA-driven synchronous or asynchronous SCSI interface • Floppy disk controller
- Full VME Rev C.1 IEEE 1014 interface
- 64-512Kb EPROM • Battery-backed RTC/SRAM
- Full debug monitor • Optional MC68881/2 FPU
- TP-IX/68K version of UNIX V.3.1* • NFS, RFS,TCP/IP

• The Evidence •

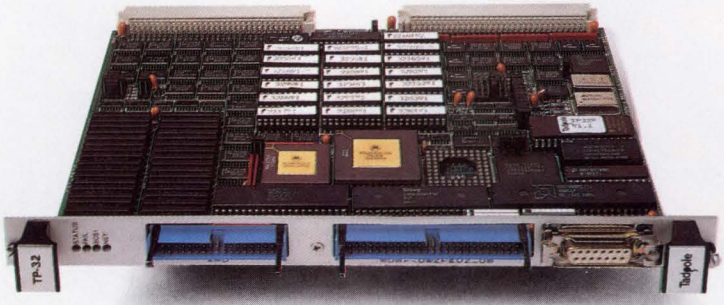

Tadpole Technology

the driving force in 32-bit design

Titan House, Castle Park,Cambridge CB3 0AY, UK
Fax: 0223 460727 **Tel: 0223 461000**

1601 Trapelo Road, Waltham, Massachusetts, 02154, USA
Fax: 1-617 890 7573 **Tel: 1-617 890 8898**
Or outside Massachusetts call 1-800-232 6656

2157 O'Toole Ave., Suite F, San Jose, California, 95131, USA
Fax: 1-408 435 8482 **Tel: 1-408 435 8223**

*Unix is a trademark of AT & T Ethernet is a trademark of the Xerox Corporation VHTX is a trademark of Ready Systems *TP-IX V.3.1 is derived from Unix V.3.1*

CIRCLE NO. 6

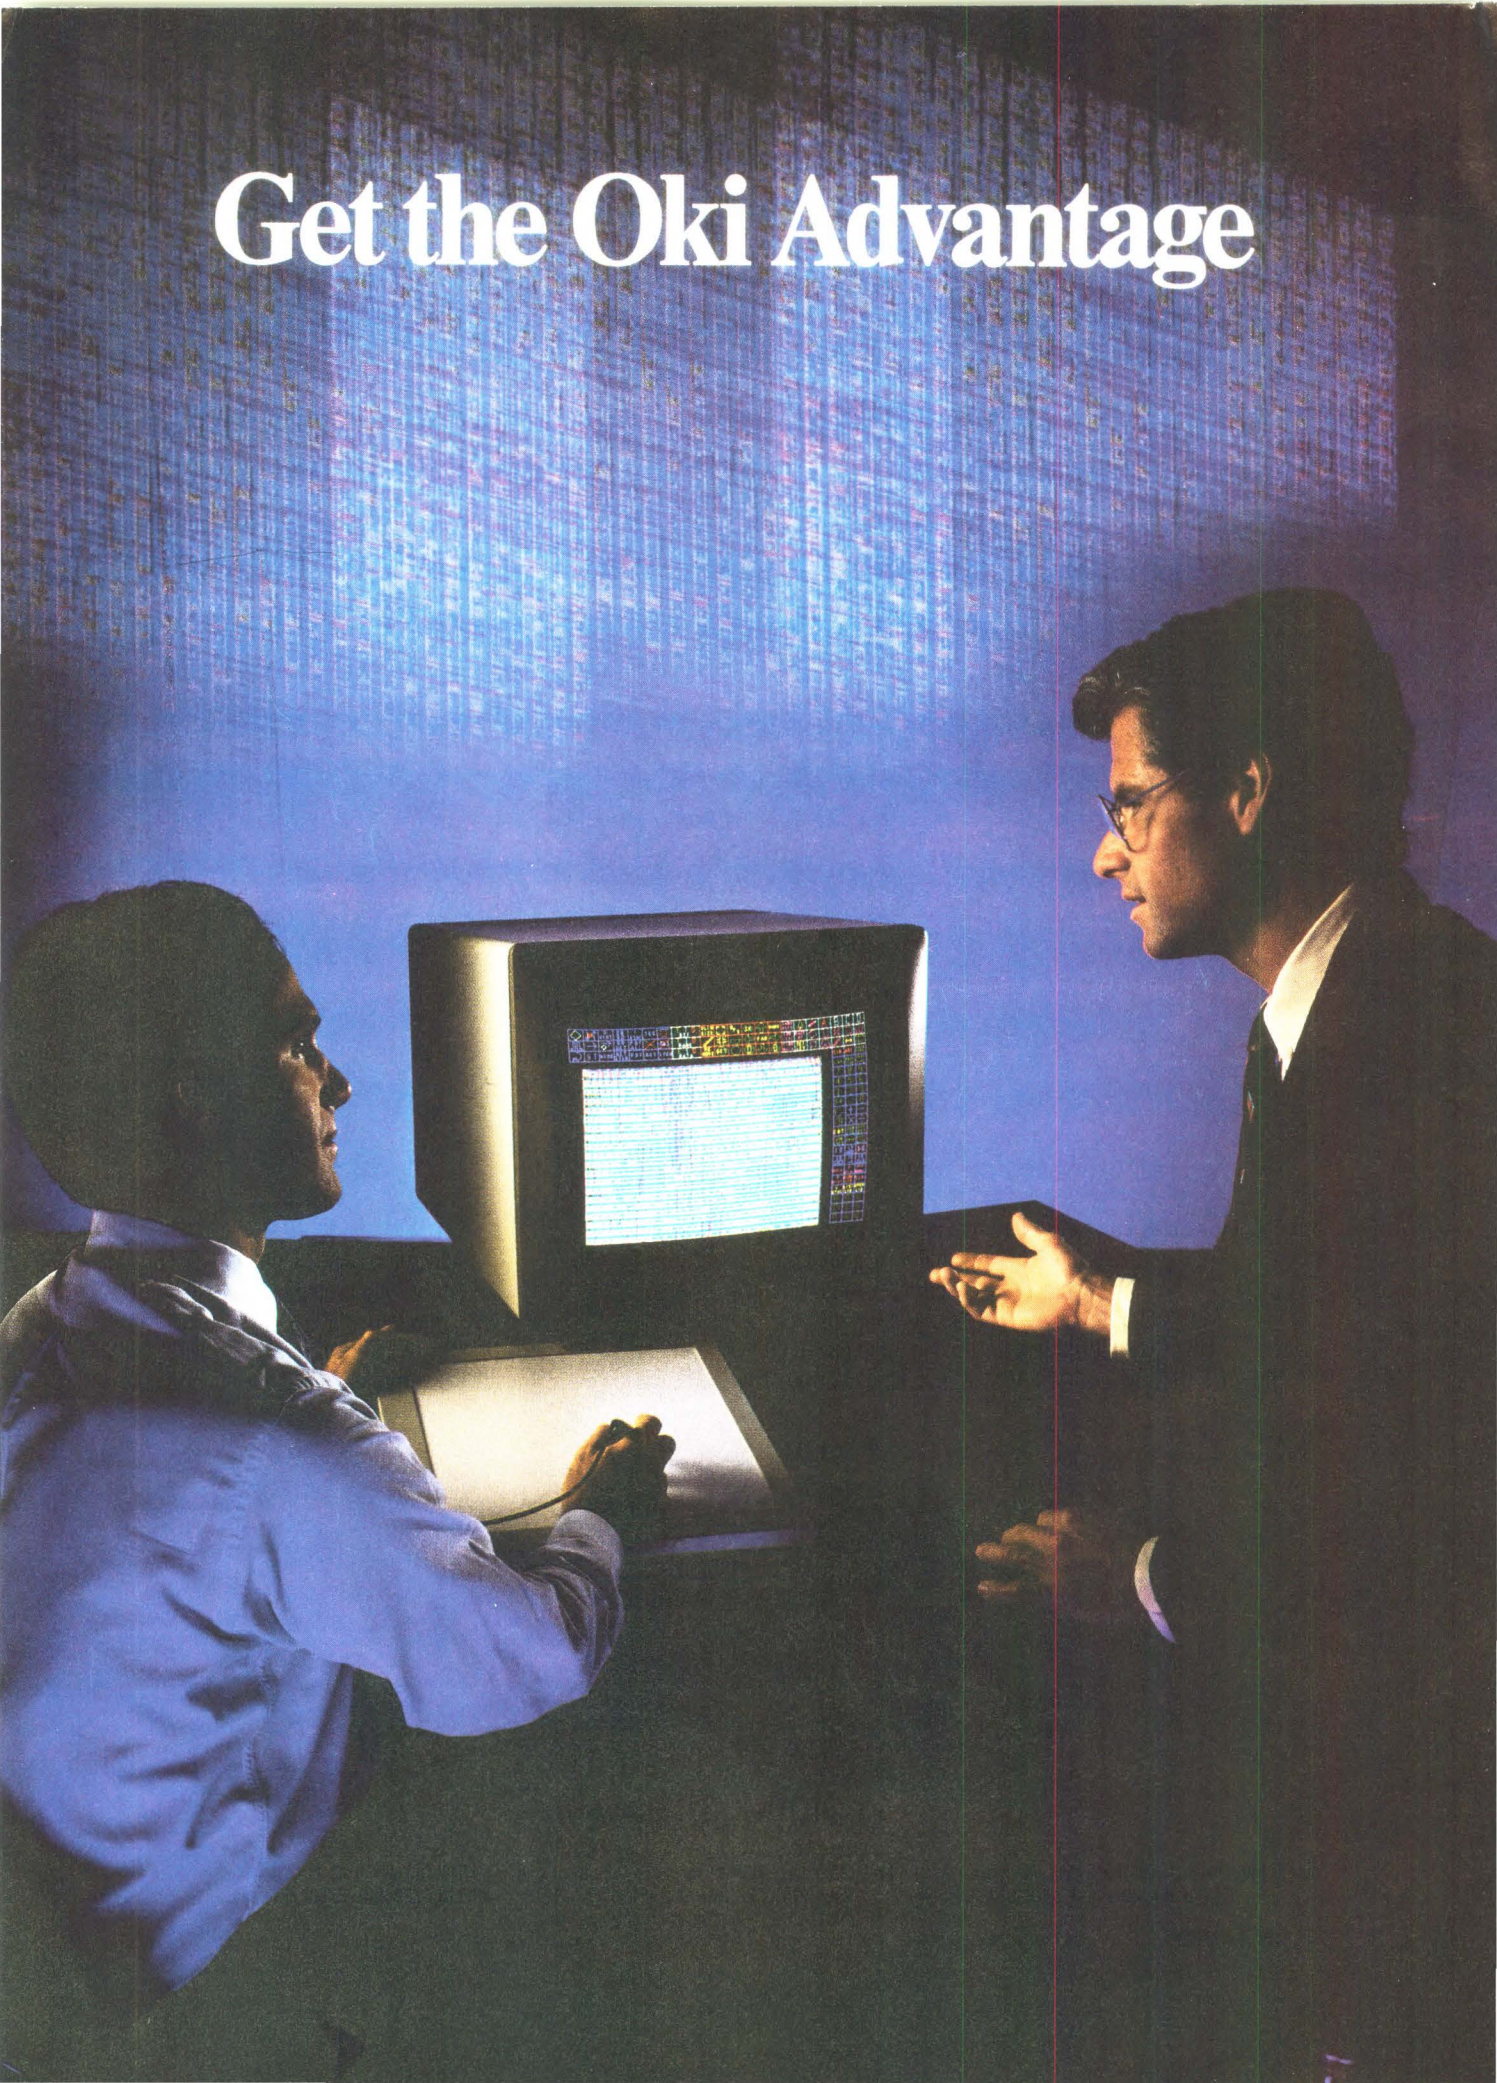
Get the Oki Advantage

Everything you need for ASIC success from one reliable source

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

Advanced ASIC products and technologies

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

▶ *Sea-of-gates:*
Flexible-Cell-Array® channelless arrays deliver up to 40,000 useable logic gates from an available 100,000. Combined with high-density memory and macros similar to the 82xx products, they offer typical delays of < 640 ps.

▶ *Channeled array:*
new 1.2-μm (drawn) channeled arrays deliver logic densities from 1,600 to 30,000 gates and sub-ns delays with output drive from 2 mA to 24 mA.

▶ *Standard cell:*
the new 1.2-μm (drawn) standard cell family offers densities to 60,000 gates, delays of 600 ps, memory capability of 32K bits RAM and 128K bits ROM and macros similar to the 82xx products.

ATG and logic transparency

With automatic test generation built in, Oki guarantees fault coverage > 95%, with no loss of performance. In addition, all three product families share the same cell library for ease of migration.

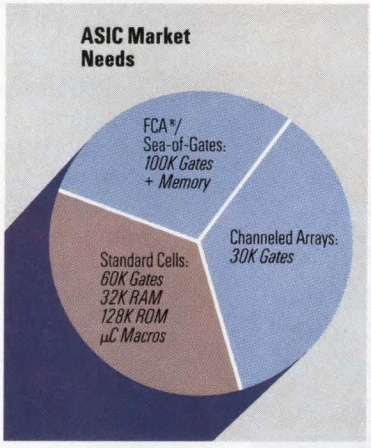

Complete ASIC support

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

Check Oki: your complete ASIC resource. CD0689

Please send complete technical data/specs on Oki capabilities in:

☐ Sea-of-gates ☐ Channeled arrays
☐ Standard cells ☐ Full customs

☐ Please call. We have immediate requirements.

Name ( ) Business Phone

Title

Company

Street Address

City State Zip Code

Mail coupon to:
Marketing Communications, Oki Semiconductor, 785 N. Mary Avenue, Sunnyvale, CA 94086. Or call (408) 720-1900.

OKI SEMICONDUCTOR

So. Calif.
(Irvine)
714/752-1843

(Sherman Oaks)
818/990-3394

No. Calif.
(San Jose)
408/244-9666

Boca Raton, FL
407/394-6099

Atlanta, GA
(Norcross)
404/448-7111

Chicago, IL
(Rolling Meadows)
312/870-1400

Boston, MA
(Stoneham)
617/279-0293

Detroit, MI
(Livonia)
313/464-7200

New York, NY
(Poughkeepsie)
914/473-8044

Philadelphia, PA
(Horsham)
215/674-9511

Dallas, TX
(Richardson)
214/690-6868

CIRCLE NO. 7

Reflections from Munich

European engineering managers seem to be as concerned about Asian competitors as their American counterparts are.

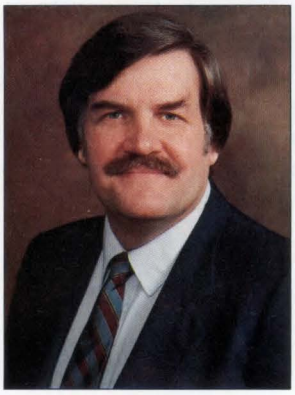

John C. Miklosz
Associate Publisher/
Editor-in-Chief

Spring has been a big tire-kicking season for *Computer Design*'s editors and for some of our readers. In just two weeks, we had Electro, Comdex and NCGA in New York, Chicago and Philadelphia, respectively. Even though these shows ranged from hot to cold in terms of what they offered for both our editors and our readers, I eagerly looked forward to still one more show—EuroBus/89 in early May, in Munich.

Unfortunately, EuroBus/89 was a disappointment because of low turnout—both of exhibitors and of engineers. That's not surprising, however, because this staging was the first, and the format of a combined conference and exhibit is apparently unfamiliar to European engineers and designers. I had been looking forward to meeting with a large number of European board manufacturers in the two days of the show, but as that didn't work out, I had an opportunity to attend some of the conference sessions and chat with the attendees. Munich in the spring also went a long way in making up for my disappointment with EuroBus/89, and I scribbled most of this editorial over a couple of large glasses of beer in one of Munich's famous beer gardens.

With the atmosphere of Munich so different from that of any city in the United States, it was easy to start thinking about some of the similarities and differences between Europe—or Germany at least—and the United States. Here are a few observations:

- **About trade shows and conferences** European designers are no different from U.S. designers in what they expect from conference papers. They want to increase their understanding of design problems and learn some solutions; they don't want sales pitches.
- **About vendors** Designers everywhere, but perhaps especially those in Europe, complain about the speed at which vendors move when they're implementing or conforming to standards, or introducing improved products, especially ones that involve a change in architecture or new standards. It's hard for designers to reconcile their needs with the vendors' desires to extend the life of existing architectures and product lines, and to maintain market position and customer base. It may be even harder on European designers because this market still isn't given a high priority by most U.S. manufacturers.
- **About competition** European engineering managers seem to be as concerned about Asian competitors as their American counterparts are. The Europeans, or at least the Germans, are having just as much trouble finding dedicated, hard-working workers as we are. The nine-to-five syndrome is affecting them as badly as it's affecting us. And where are they turning to find those dedicated workers? You've probably guessed it: Asia and the Middle East.
- **About life in general** Everybody, it seems, listens to American rock and country & western (or at least all the Munich taxi drivers do). The news that smoking is bad for your health hasn't reached Germany. Nobody drives at 55. Nobody worries about drinking too much beer. Burp.

Check DY-4 for the widest choice in VME commercial, ruggedized and Mil-Spec NDI

DY-4 offers the widest variety of Non-Developmental Items (NDIs) right across the board: commercial NDIs for developmental projects; ruggedized NDIs; and Mil-Spec NDIs fully compliant with MIL-E-4158, 5400, and 16400 specs.

All DY-4 NDIs are manufactured in accordance with AQAP-1 and include advanced features like built-in test and system-level diagnostics.

For the most and the best NDIs, call DY-4 — one of the world's leading suppliers of board-level products for commercial and military applications.

California: (408) 377-9822
Los Angeles: (714) 549-2559
Washington: (301) 964-5554
Massachusetts: (508) 692-9308
Canada: (613) 596-9911
Denmark: (06) 963824
Germany: (089) 903 9206

DY-4 SYSTEMS INC.

CIRCLE NO. 8

DY-4 NDI

PRODUCT	COMMERCIAL	RUGGEDIZED	MIL-SPEC
Single Board Computers (680x0/88000)	✔	✔	✔
Military / Communications Interfaces			
MIL-STD-1553 B	✔	✔	✔
NTDS (Serial and Parallel)	✔		
Digital-to-Synchro	✔	✔	
Digital-to-Resolver	✔	✔	
Synchro-to-Digital	✔	✔	
Resolver-to-Digital	✔	✔	
ARINC 429 / 561	✔	✔	
Memory Modules			
SRAM	✔	✔	✔
EPROM	✔	✔	✔
EEPROM	✔	✔	✔
DRAM	✔	✔	✔
System Controllers	✔	✔	✔
Serial I / O	✔	✔	✔
Intelligent Serial I / O	✔	✔	✔
Parallel I / O	✔	✔	✔
Analog-to-Digital	✔	✔	✔
Digital-to-Analog	✔	✔	✔
Graphics Engines			
Medium Performance (34010)	✔	✔	✔
High Performance (68020 / QPDMs)	✔	✔	
Very High Performance (RISC / QPDMs)	✔		
IEEE-802.3 LAN	✔	✔	
ANSI X3T9.5 LAN (FDDI)	✔		
SCSI	✔	✔	✔
SASI	✔	✔	
9-Track Tape Controller	✔		
Relational Database Processor	✔		
Chassis Systems	✔	✔	✔

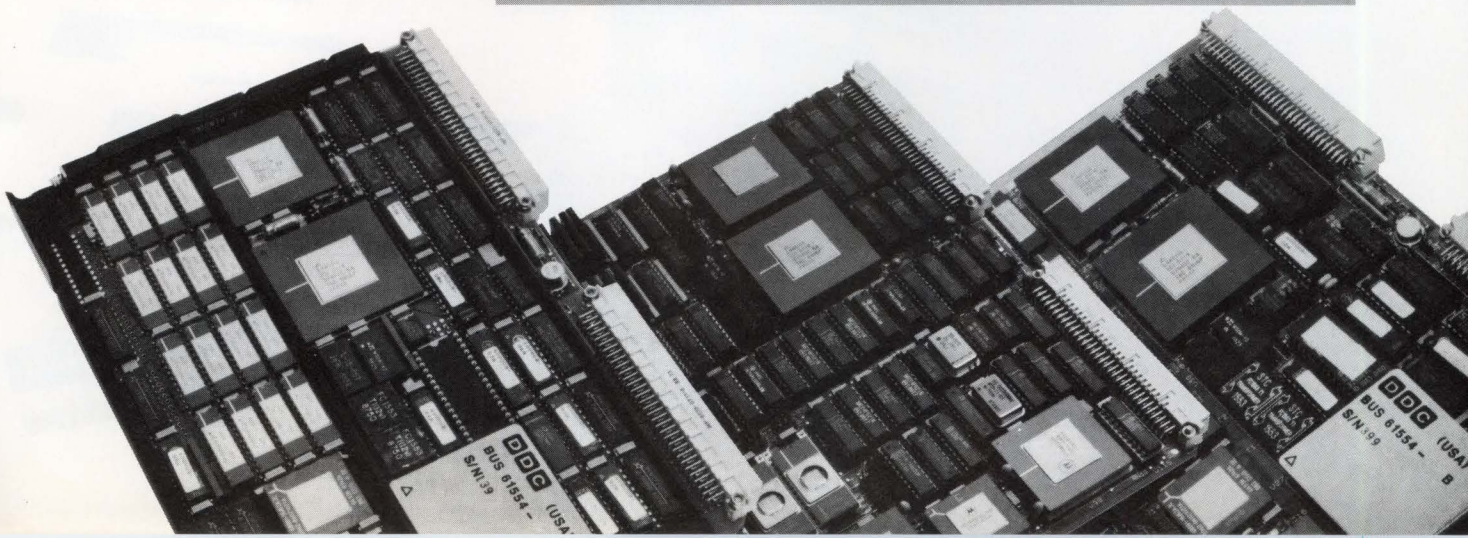

TEXAS INSTRUMENTS REPORTS ON HIGH-PERFORMANCE IN THE ERA OF MEGACHIP™ TECHNOLOGIES

Speed, power, drive, noise — system logic design can be a long series of trade-offs. To put the odds in your favor, get our *free* BiCMOS and ACL Performance Data-Packs.

Boosting system performance and cutting power consumption is a common design goal, but getting there needn't involve as many trade-offs as it used to. Because now, TI's BiCMOS and Advanced CMOS Logic (ACL) families combine the speed and drive you need with low-power operation and enhanced signal integrity. We want to prove it to you with our free Performance Data-Packs.

On your screen, see how our BiCMOS chips stack up

Our BiCMOS technology calculation software tool (BCT Calc) makes performance evaluation simple. Just place the disk into your TI or IBM®-compatible PC with Lotus® and enter the performance parameters you want to check out. You'll see exactly how TI's high-performance BiCMOS logic stacks up against other available bipolar and CMOS devices. Call or use the return card today.

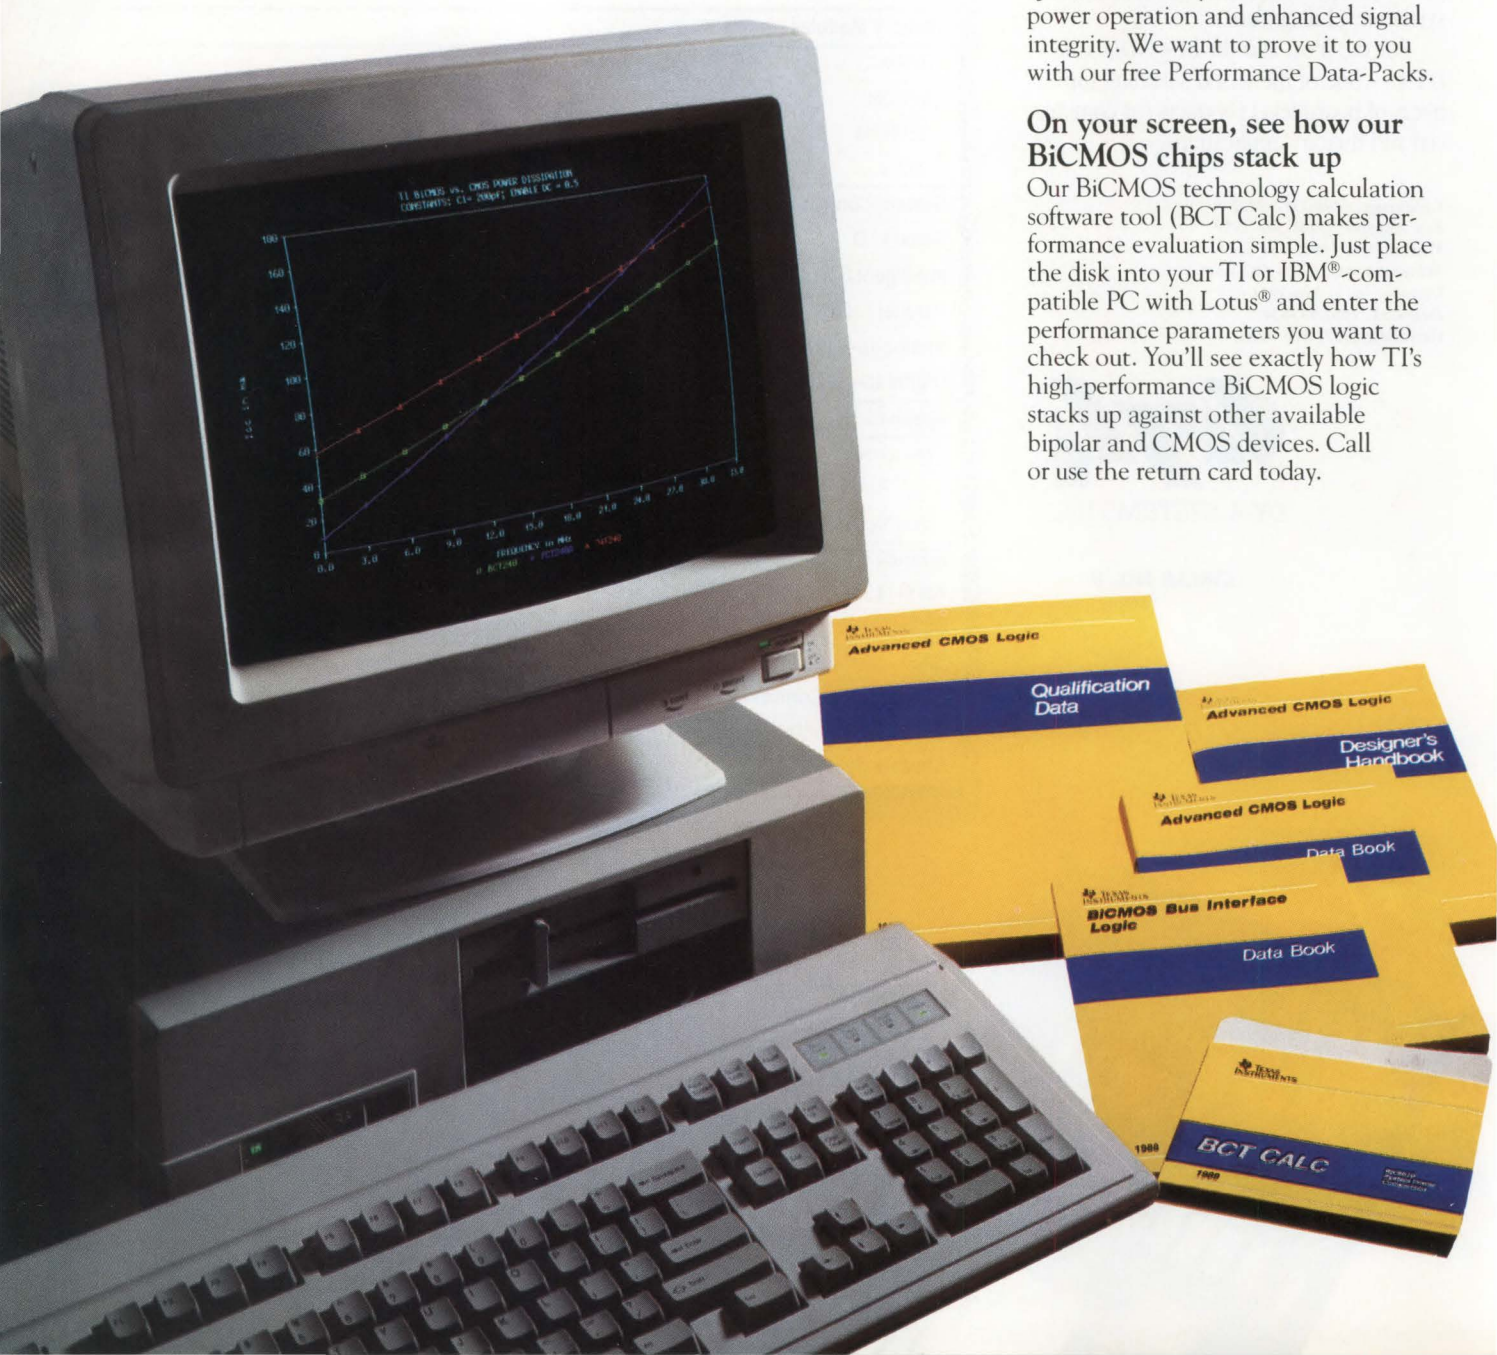

LOW-POWER LOGIC

TI's BiCMOS logic: Use it to cut power without compromising performance

By combining bipolar and CMOS technologies, TI BiCMOS gives you the best of everything.

Switching speeds are comparable to advanced bipolar devices. You get the 48/64-mA drive current needed for high-capacitive loads and backplanes. This meets the drive requirements of industry-standard buses such as Multibus® and VMEbus™.

Yet TI's BiCMOS family can reduce disabled currents as much as 95% and active currents as much as 50% compared to bipolar equivalents. As a result, total system power savings can be more than 25%.

TTL I/O provide enhanced signal integrity, power-down backplane interface, and backplane-level standards.

There are more than 60 members in the BiCMOS family, including 8-, 9-, and 10-bit latches, buffers, drivers, and transceivers—whatever you need to make your bus-interface design a high-performance reality.

Again, check it all out for yourself with our BiCMOS Performance Data-Pack.

TI's EPIC ACL logic family: Use it to keep performance up and squelch "ground bounce"

This is the first advanced CMOS logic family designed to improve system performance while reducing the simultaneous switching noise or "ground bounce."

Performance is comparable to advanced bipolar devices, supporting frequencies up to 125 MHz and providing the 24-mA output capable of driving 50-ohm transmission lines. It's a perfect match for core and local bus interface.

The family is fabricated using TI's EPIC™ one-micron CMOS process, so power consumption is the lowest of any CMOS family, essentially zero standby.

Noise levels are as much as 50% lower compared to conventional advanced CMOS logic. That's because TI's family has two distinctions going for it: Innovative packaging and a circuit-design technique called OEC™ (Output Edge Control) that smooths the sharp edges of the output waveform which are a major cause of simultaneous switching noise.

A unique pinout arrangement creates a "flow-through" architecture. This simplifies board design and can save up to 32% in board space because you don't need passive components to control noise as you usually do with end-pin devices.

TI's more-than-300-member family includes versions compatible with CMOS and TTL voltage levels. Simulation models are available for both BiCMOS and ACL devices.

Call 1-800-232-3200, Inquiry 3004, for your free Performance Data-Packs

A call or a card will start your BiCMOS Performance Data-Pack on its way to you. Included are our performance calculation software and our BiCMOS data book and product bulletin.

Our ACL family Performance Data-Pack includes:

- TI *ACL Designer's Handbook* that describes methods for handling the technical issues facing advanced-logic design engineers.

- TI *ACL Data Book* that contains detailed specifications and applications information on family members.
- TI *ACL Qualification Data* that provides reliability and characterization data, die photos, and application derating factors.

It's all free and it's all yours by making a call or by completing the return card, today.

TEXAS INSTRUMENTS

™ MegaChip, EPIC, and OEC are trademarks of Texas Instruments Incorporated.
VMEbus is a trademark of Motorola, Inc.
® IBM is a registered trademark of International Business Machines Corporation. Lotus is a registered trademark of Lotus Development Corporation. Multibus is a registered trademark of Intel Corporation.
© 1989 TI 08-8451

CIRCLE NO. 9

One–Stop Debugging

UDB From OASYS

✓ **Multi-Window**

✓ **Multi-Processor**

✓ 68000/20
✓ 8086/286
✓ 88000
✓ Sun
✓ VAX/VMS
✓ Motorola VME Delta

- Multi-Window Symbolic Source-Level Cross-Debugger
- Can Debug Multiple Processors at the same time--even a 68020 and an 80286 simultaneously
- 68000/20, 8086/286 and 88000 Support
- Complete Symbolic Disassembly
- Breakpoints, Macros, Full Debug Log
- User-Configurable User Interface
- Target Execution Monitor/Remote Server Supplied
- ICE Support
- Provides Multi-Window Environment even on Character-oriented ANSI terminals

Available for:
* *VAX/VMS*
* *Sun-3*
* *Motorola Delta VME Series*

Works with:
* *OASYS Green Hills C for 680x0, 88000*
* *OASYS Microsoft Cross C*

Call OASYS today about UDB and our other cross-development tools (compilers, assemblers, editors, real-time OS, etc.)

Oasys

230 Second Avenue, Waltham, MA 02154 **(617) 890-7889**

Trademarks acknowledged to: CaseTools, DEC, Motorola, Sun, Microsoft

CIRCLE NO. 65

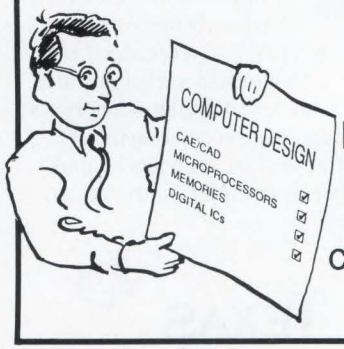

THANK YOU!

More and more of you have been rating this magazine Number One in our readership studies. Thanks for the vote of confidence, and we'll continue to live up to your expectations.

COMPUTER DESIGN

THE FIRST MAGAZINE OF SYSTEM DESIGN, DEVELOPMENT AND INTEGRATION

PUBLISHER
David L. Allen (508) 392-2111

ASSOCIATE PUBLISHER/EDITOR-IN-CHIEF
John C. Miklosz (508) 392-2114

TECHNOLOGY EDITORS
Ron Wilson, *Integrated Circuits* (503) 224-9396
Warren Andrews, *Integrated Circuits* (201) 825-1317
Bill Harding, *Design & Development* (408) 729-1005
David Lieberman, *Computers & Subsystems* (508) 392-2117
Tom Williams, *Software, Graphics & Imaging* (408) 335-5514
John H. Mayer, *Computers & Subsystems* (508) 392-2118
Sydney F. Shapiro, *Data Acquisition & Control* (508) 392-2116

CONTRIBUTING EDITORS
Ernest Meyer, *Design & Development* (408) 378-6241
Howard Falk, *Software* (201) 783-4050
John Bond, *Computers & ICs* (508) 443-2310

RESEARCH/SPECIAL PROJECTS MANAGER
Sydney F. Shapiro (508) 392-2116

SECTION/SPECIAL PROJECTS EDITORS
Ron Wilson, *Technology Updates* (503) 224-9396
Michael G. Donlin, *New Product Highlights* (508) 392-2123
Patti Kenney, *Buying Guides* (508) 392-2124

MANAGING EDITOR: Susan R. Nykamp

GRAPHICS/PRODUCTION MANAGER: Laura C. Smith

ASSISTANT EDITOR: Kerstin R. Feldman

COPY EDITORS: Barbara-Ann Scofidio, Anne Holm

EDITORIAL ASSISTANTS: Claire Coupal, Claire Ellis

ART DIRECTOR: Jan Horner
ILLUSTRATOR: Fred Lomas

PRODUCTION DIRECTOR: Jan M. Lopez
AD TRAFFIC MANAGER: Adrienne Adler
COVER ART: Sergio Roffo

CIRCULATION DIRECTOR: Robert P. Dromgoole

PennWell PUBLISHING COMPANY

Advanced Technology Group
L. John Ford, Senior Vice-President
Leslie P. Cypret, Vice-President, Administration
One Technology Park Drive
P.O. Box 990
Westford, MA 01886
Tel: (508) 692-0700
Fax: (508) 692-0525

Postmaster: Send change of address form 3579 to COMPUTER DESIGN, Circulation Department, Box 3466, Tulsa, OK 74101

VxWorks. More than Ready for Real-Time.

VxWorks Believers	VxWorks Advantages
Amdahl *Supercomputing*	Seamless partnership with UNIX
Bendix *Aerospace testing*	High-speed optimized real-time kernel
Boeing *Flight simulation*	Support for CISC and RISC architectures
Carnegie Mellon *Robotics control*	Ethernet, TCP/IP, sockets, RPC, NFS
Honeywell, Inc. *Systems research*	Easily configurable multiprocessor support
Hughes Aircraft *Flight simulation*	Host and target hardware independent
Jet Propulsion Labs *Robotics and teleoperation*	Tailorable (kernel to complete operating system)
Litton Industries *Image processing*	Fully re-entrant and sharable code and library
Martin Marietta *Land vehicle navigation*	Over 600 callable utility functions
MIT– Haystack Observatory *Telescope control*	Fully integrated network file system
National Bureau of Standards *Robotics control*	Source-level and symbolic debugging tools
Pacific Bell *Network research*	Dynamic and incremental linking and loading
RCA Corporation *Telerobitic control*	Support for high-level languages – C and Ada
Rockwell *Space shuttle telemetry*	Unlimited on-line support
Stanford University *Computing research and development*	Responsive and engineering-oriented company*
Tektronix *Instrumentation control*	rlogin, telnet, ftp, rsh

**And the only one with an in-house pinball arcade!*

Application developers designing sophisticated systems are seeking a real-time operating system that is fast, completely integrated and supported by a rich UNIX development environment.

Only one real-time operating system satisfies these requirements.

VxWorks.

But don't take our word for it. Evaluate VxWorks. Rigorously. Unmercifully. Compare it with products from Ready Systems and other competitors. Contrast it to your in-house operating system design. Then make up your own mind.

That's what the VxWorks Believers did. They compared and contrasted. They considered the performance versus effort, time and money trade-offs. Then they all selected VxWorks. Except one firm that opted to do it in-house. A year later they called back...

for VxWorks.

Get more than ready for a revolution in real-time. Call us at 1-800-545-WIND.

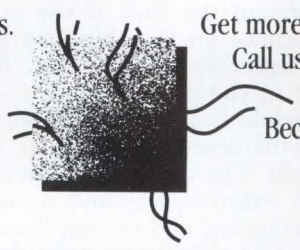

Become a Believer.

Wind River Systems, Inc.
1351 Ocean Avenue
Emeryville, CA 94608

CIRCLE NO. 10

INTEGRATED CIRCUITS

Motorola rethinks the fundamentals for one-chip 32-bit microcontroller

Ron Wilson, Senior Editor

Responding to the needs of the engine-control experts in Detroit, Motorola (Austin, TX) has developed a single-chip microcontroller with a full 32-bit CPU core: the 68332. How successful the new part will be in roles outside the engine compartment will be determined at least in part by the innovative architectural approach the chip designers devised.

A state-of-the-art engine controller, like sophisticated control systems in many industrial and aerospace applications, must compute its control responses on the fly. There are too many variables, and the strategy is too complex, to rely on the table-look-up methods used in earlier, simpler black boxes. This requirement for computation on the fly virtually mandates a fast 32-bit CPU core, just to keep up with the calculations.

But a 32-bit core by itself isn't enough. The control environment is full of event-driven interruptions, such as the arrival of a data sample or the closing of a switch, that need a minimal amount of processor attention. If all of these minor events present themselves as interrupts to the 32-bit core, most of the big CPU's clock cycles are going to be spent on context switching. "If you give customers a powerful CPU and then tie it up doing interrupts, you haven't really given them anything," warns Pat La Violette, Motorola microcontroller system design manager.

To solve this problem, board-level system designers use a multiprocessing approach: they surround the 32-bit processor with 8- or 16-bit microcontroller chips that can handle whole I/O transactions independently, thus reducing the number of interrupts and noncomputational instructions forced onto the big CPU.

This strategy works well for military avionics and laser printers, but not for the maniacally cost-conscious auto industry. Motorola had to find a way to provide an architecture that held a 32-bit CPU core, protected that core from event-driven interrupts and fit everything on a minimum number of chips.

The solution was a modified 68020 CPU core, connected via an on-chip bus to a series of memory modules and specially-designed intelligent I/O modules. "The CPU core executes the majority of 68020 codes, but leaves out some of the things like bit-field instructions that aren't all that useful for control computations," explains Brian Wilkie, advanced MCU operations manager. "Instead, we put in some new instructions to handle interpolated table-look-up operations and a low-power shutdown mode."

In the 68332, the core will have on-chip access to 2 kbytes of static RAM. Future products in the family will add on-chip ROM. Access to additional memory, and to 68HC11-compatible peripheral devices, comes through the first of the major I/O modules on the chip, the External Integration Module (EIM). Connected to the core through the intermodule bus, the EIM is virtually a motherboard chip set in its own right. The module creates an 8-/16-bit external bus and a variety of debugging and testing features, including a four-wire serial test bus.

But this is more than an ordinary bus controller. The EIM contains a bank of sophisticated chip-select logic that not only generates 12 programmable chip-select signals, but also controls bus width and bus timing for each chip-select signal. To give some idea of the module's flexibility, it's possible to have both a 16-bit version of the 68030 two-cycle memory bus and an 8-bit 68HC11-style bus on the same pins, with the nature of the bus on a particular cycle determined by the bus address. This combination of features gives the 32-bit CPU core fast access to external memory and flexible access to slower peripheral devices, without the CPU overhead of synthesized bus cycles.

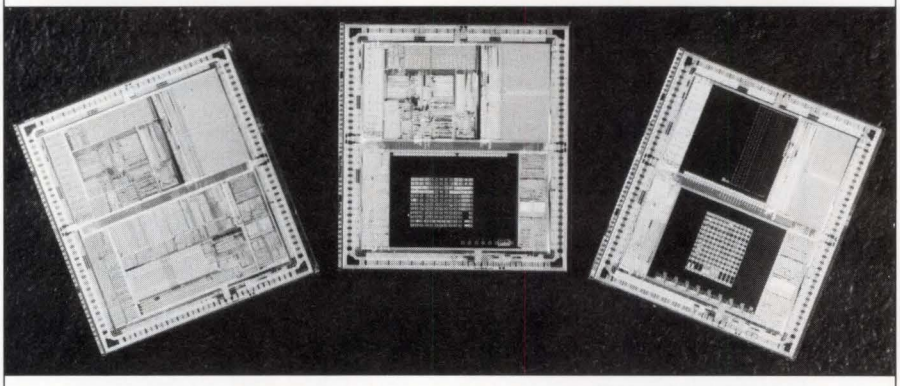

Modularity permitted Motorola to break the 68332 design into three phases of development. At left, the chip contains only the static RAM, the system integration module, the queued serial module and the chip selects. The time processing unit is then added. With the addition of the CPU, the chip is complete (right). The modular approach shortened Motorola's design cycle and gave the company the ability to introduce new versions of the chip quickly.

Intelligent I/O modules

Bus management is only one of the vital issues in a controller design. In systems that control electromechanical devices, counter/timer circuitry is extremely important. And for connection to off-chip, low-speed peripherals or remote systems, serial interfaces are necessary. So the 68332 designers have made counter/timer and serial interface modules with considerable local intelligence.

Many of the functions performed with timers in an electromechanical system are in fact not simple timing operations—in applications such as pulse-width modulation or stepper motor control, the actual output from the chip represents a complex

16-Bit CMOS Microcontrollers. More Functions and More Memory On-Chip. Now. From Mitsubishi.

32K bytes EPROM or ROM. 2K bytes RAM. Two UARTs, with on-board baud rate generator. 8-bit, 8-channel A-D converter. Eight multi-function 16-bit timers that operate in five modes (timer, event-counter, one-shot, pulse-period and pulse-width), and pulse-width modulation. 12-bit watchdog timer. Interrupt controller. 68 independently programmable 1/0 ports. All on-chip.

Plus, linear access for 16 megabytes external memory. Powerful instruction set optimized for high-performance operation. High efficiency hardware multiply and divide instructions. And, configurable for 8- or 16-bit external operation. All in low power, 1.3 μm silicon-gate CMOS technology. All part of a broad 16-bit MCU product line.

IBM®-PC based development and software support tools. Plus, masked ROM, ROM-less and OTP ROM versions and prototyping adapters.

All available now from Mitsubishi. Driven by your applications, we put more on-chip and give you more choices, so you can reduce system component count, save board space, lower overall design costs and get products to market fast. For more information call now. (408) 730-5900 Ext. 2106. 1050 E. Arques Avenue, Sunnyvale, CA 94086.

IBM is a registered trademark of International Business Machines Corporation.

MITSUBISHI ELECTRONICS

Applications-Driven Microcontrollers

CIRCLE NO. 11

It's Fast

200,000 A/D samples per second.

You can capture fast analog signals using Acromag's new VME board in your system . . .only 5 microseconds per channel.

Versatile Inputs.

16 channels differential; or, you can get 32 channels single-ended, selected by a jumper on the VME board.

Analog inputs are jumper selectable to:
±5 volts or ±10 volts or 0 to +10 volts

Gain adjustments are software programmable to cover full-scale voltage inputs from ±0.625 to ±10 volts in nine ranges. Each input channel can have a different gain, if desired.

Flexible Trigger Controls.

Input data scanning can be continuous, or trigger controlled from the internal clock or from your CPU's command. You can set up block mode triggering, so the board takes a programmed number of samples, then stops.

Channels can be scanned in any order, or repeated under programmed control.

Input data scan rates can be from 200,000 readings per second to any desired slower data rate with software programming.

Dual Port RAM Stores the Samples.

The new Acromag A/D board's ability to store up to 64,000 readings means your CPU is relieved to do other tasks while the Acromag A/D board captures your data. This unburdens your CPU from having to wait for each A/D conversion value.

Low Cost.

The cost is about $50 per channel for these great features and performance.

No. of analog inputs	16 Differential 32 Single ended
Ranges	± 5V, ± 10V, 0-10V
A/D resolution	12 Bit
Conversion time	5μ sec. or 10μ sec.

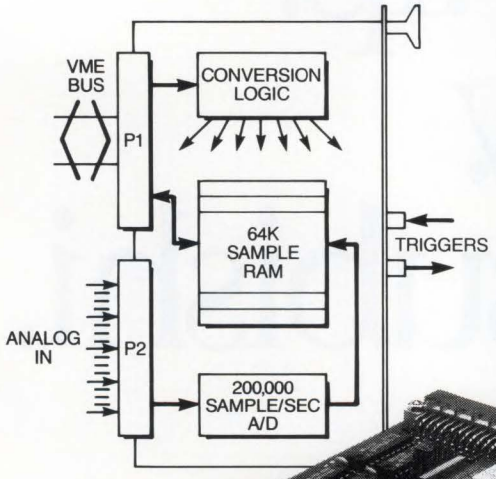

Want More Data?

For full details about our FAST A/D VME boards, ask for your copy of our new technical bulletin 20-004.0.

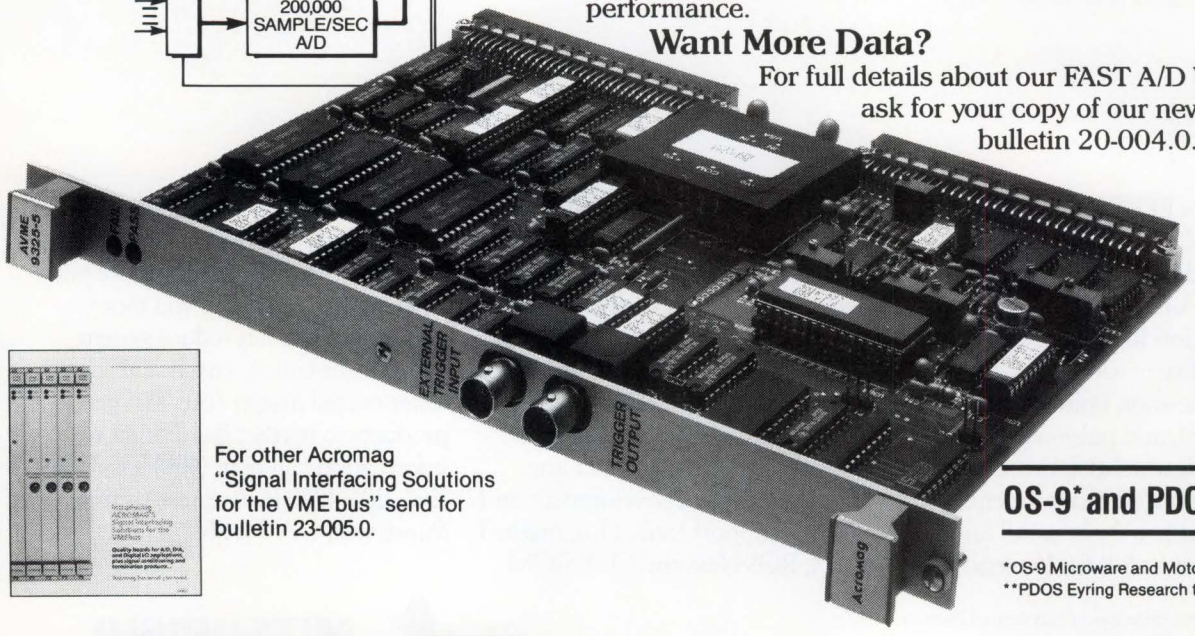

For other Acromag "Signal Interfacing Solutions for the VME bus" send for bulletin 23-005.0.

OS-9* and PDOS** Drivers

*OS-9 Microware and Motorola trademarks.
**PDOS Eyring Research trademark.

Acromag®
LEADER IN SIGNAL CONDITIONING

30765 Wixom Road • Wixom, MI 48096 • 313-624-1541 • Fax 313-624-9234

CIRCLE NO. 12

INTEGRATED CIRCUITS

interaction between two or more timers under control of the CPU. Since Motorola's goal was to remove the CPU from such low-level transactions, the team embedded a programmable processor in the timer module on the 68332, essentially creating a timing microcontroller within the on-chip I/O module.

The Time Processing Unit (TPU), as the module is called, includes 16 free-running counters with capture and compare registers. The counters interface not to the core CPU but to a local timer microengine, which can link counters together, schedule execution of microroutines and respond to counter outputs. "Some of the microcode we provide in the TPU is quite sophisticated," claims Wilkie. "We have a routine that can accelerate, cruise and decelerate a stepper motor, for example, all without interrupting the CPU."

The chip's serial communications module shows a similar level of sophistication. The unit consists of two sections: a Serial Communications Interface (SCI) and a Queued Serial Peripheral Interface (QSPI). The SCI is an ordinary universal asynchronous receiver-transmitter with a rate generator, intended primarily for serial links to remote devices such as engine diagnostic equipment. The QSPI, however, represents new thinking on an increasingly important aspect of controller design: high-speed serial I/O buses.

The QSPI creates an external serial bus with clock, data-in and data-out lines, plus four chip-select signals. Up to 16 devices can be placed on the bus under direct control of the QSPI. The bus control tasks are handled by a local processor, which not only manages serial/parallel conversion and bus control signals, but also maintains a set of data queues in a small local RAM.

The result is that devices on the serial bus appear, to the 32-bit CPU, as locations in the local RAM. Like the other I/O modules, the QSPI reduces external parts count and the impact of I/O activity on the central CPU core.

A futuristic approach

By starting with a set of application requirements, rather than with a preconceived notion of microcontroller architecture, the 68332 designers came up with a unique approach. The chip combines a CPU with the computational clout to handle sophisticated control strategies, the specialized peripherals necessary for electromechanical interfaces, and a means of integrating entire control systems on a high-speed serial bus. In addition, the silicon provides for communications, an efficient external memory interface and elaborate diagnostic features.

But the 68332 team has done all this without trashing the throughput of the CPU core. In every instance, the architecture uses intelligent I/O to protect the CPU from the low-level, high-frequency interrupts. In this respect, the 68332 stands at the leading edge of microcontroller design. The most recent 16-bit controllers have been moving in this direction, but none has taken such a thorough approach.

All indications are that the evolution will continue. The 68332 itself, with its independent modules linked through a silicon bus, is clearly just the beginning of a family that will expand in both hardware and firmware. Not only can Motorola easily add modules to the on-chip bus, but since the existing I/O modules are virtually self-contained microcontrollers, the company can modify the firmware within the modules to implement new functions.

As transistor counts move inexorably higher, it seems likely that single-chip products will reflect the multiprocessing solutions already proven in the world of board-level controllers. That may be the long-term message of the 68332 design.□

For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate number on the Reader Inquiry Card.

Motorola
(512) 891-2034 ..... Circle 247

CIRCLE NO. 13

The old road to IC design.

The Cadence expressway.

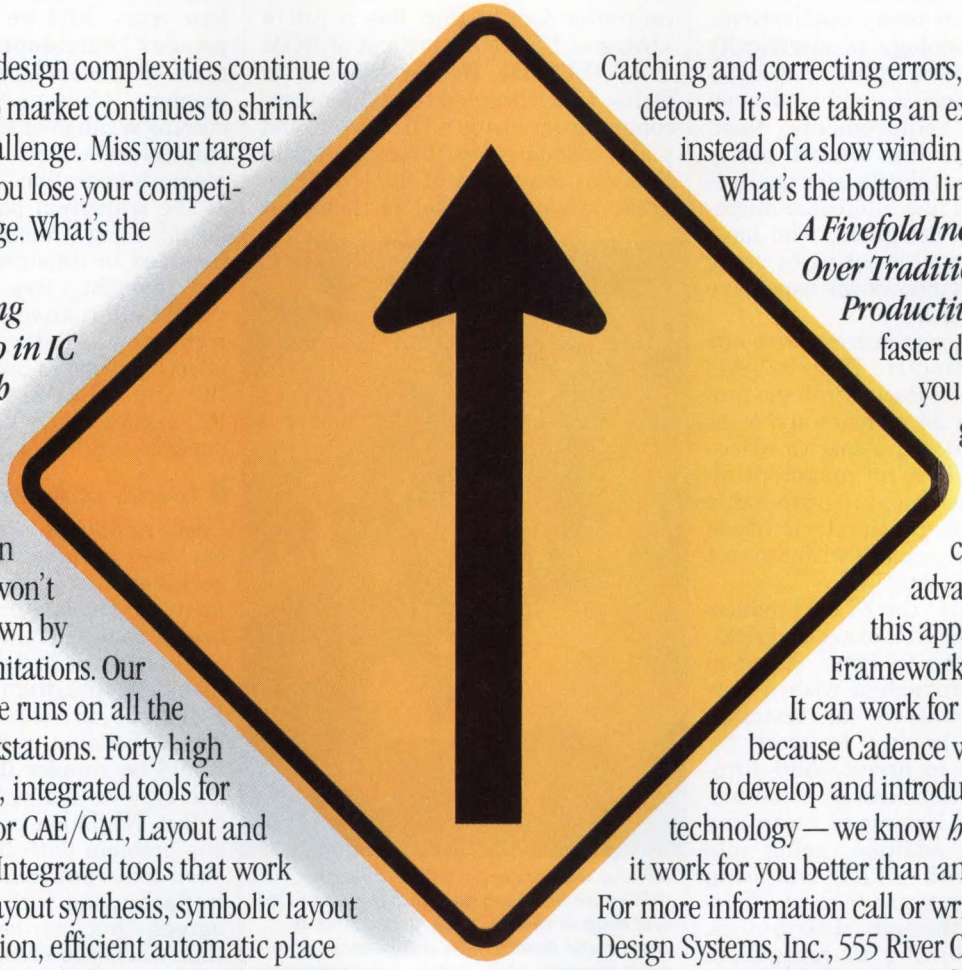

While IC design complexities continue to grow, time to market continues to shrink. It's a real challenge. Miss your target window — you lose your competitive advantage. What's the solution?

A Working Partnership in IC Design With Cadence With Cadence as your IC design partner you won't be slowed down by hardware limitations. Our ICDA software runs on all the leading workstations. Forty high performance, integrated tools for ICDA. Tools for CAE/CAT, Layout and Verification. Integrated tools that work together — layout synthesis, symbolic layout and compaction, efficient automatic place and route tools including gate arrays, module generation. Forty integrated ICDA tools that share a common user interface and design database. As your design moves forward, nothing is lost in transit. You see the impact of every change. In real time. Catching and correcting errors, eliminating detours. It's like taking an expressway instead of a slow winding road.

What's the bottom line?

A Fivefold Increase Over Traditional Design Productivity With faster design times, you hit your target windows and maintain your competitive advantage. We call this approach Design Framework™ architecture. It can work for you. And because Cadence was the first to develop and introduce framework technology — we know *how* to make it work for you better than anyone else.

For more information call or write: Cadence Design Systems, Inc., 555 River Oaks Parkway, San Jose, CA 95134. Inside California **1-800-672-3470, ext. 866.** Outside California **1-800-538-8157, ext. 866.**

CADENCE

CIRCLE NO. 14

INTEGRATED CIRCUITS

Vendors eye flash EPROM for role in one-chip microcomputers

Ron Wilson, Senior Editor

Flash memory has proved itself a serious alternative to UV-EPROM in many applications. The new technology is electrically erasable in blocks, rather than byte-by-byte. In density, cost, and even power-supply requirements, flash memory is moving toward parity with older nonvolatile storage media. This trend has interested microcontroller chip designers, who have long sought an alternative to using little quartz windows on expensive ceramic packages.

"We've had great success with our Zero Turn Around Time (ZTAT) parts," says Ken Pope, strategic program manager for memory and mass storage at Hitachi America (Brisbane, CA). "They're microcontrollers with one-time-programmable EPROM, so you get the flexibility of programmable memory with the cost of plastic packaging."

Customers use the ZTAT products to improve time-to-market and flexibility, according to Pope. A group can go into production with a new design immediately, for instance, rather than having to wait for masked parts to arrive. And firmware modifications can be accommodated without having to scrap any inventory. In addition, going into production with ZTAT lets even high-volume customers meet short-term needs through distributors, rather than having to renegotiate shipping schedules with the factory.

"Given all these advantages, taking the next step to flash memory on the chip would make a lot of sense," Pope says. Other vendors with flash processes see the situation similarly. Mike Polen, microcontroller marketing manager at Texas Instruments (Stafford, TX), for example, sees advantages to using flash parts in the TI line, too. "In the long run, we see the cost difference between flash and masked ROM dropping, maybe to as little as a 70 to 80 percent premium," Polen predicts. "At that price, you would break even on a flash microcontroller the first time there was a code change."

Polen knows whereof he speaks. "We provide a part in the 370 microcontroller family that has a full 4 kbytes of EEPROM instead of ROM or UV-EPROM. We developed it for Delco, which wanted to seal a microcomputer up inside a CD player, and make code changes through a serial line that comes out of the box. It's a great development tool for them, but the cost is too high for production.

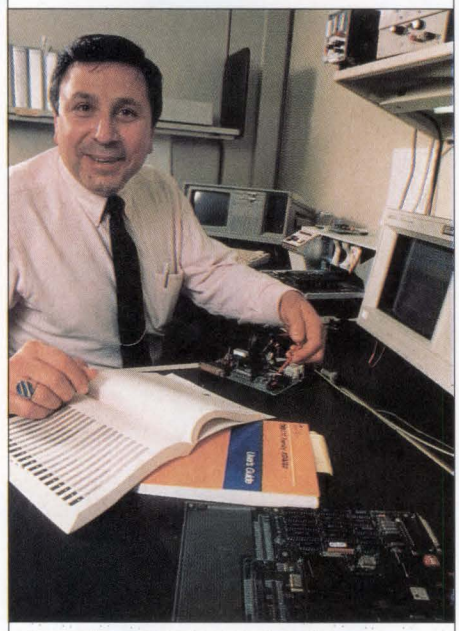

Several vendors expect that the cost difference between flash and masked ROM will drop in the next few years. And as costs come down, flash will become dominant, predicts Mike Polen, microcontroller marketing manager at Texas Instruments. "I think the economics will eventually get to the point where customers will never do masked ROM versions of their parts in the first place," he says.

"More important," Polen continues, "is that the EE cell is so large. We can't get more than 4 kbytes of EEPROM on a chip. Moving to flash will let us expand that memory considerably."

As costs come down, flash will become dominant, Polen believes. "I think the economics will eventually get to the point where customers will never do masked ROM versions of their parts in the first place."

Vinod Mahendroo, operations manager for 16-bit embedded microprocessors at Intel (Chandler, AZ), has similarly concrete ideas about the value of a flash microcontroller. "We expect our flash devices to come near PROM costs within three to four years. And we have a range of popular microcontroller architectures that would benefit from flash memory. Consequently, we're discussing a range of single- and multichip solutions with our Folsom (CA) memory operation."

The transition point may still be far off. Zilog (Campbell, CA), deeply involved in consumer electronics at the moment, sees all of the EEPROM alternatives as too expensive right now. "Today, the most cost-effective way to maintain data on the chip is to put a capacitor beside it," argues Gary Pacey, director of superintegration programs.

Toward a flat memory architecture

Cost arguments aside, some customers see architectural reasons for seeking out a nonvolatile main memory for microcontrollers. They argue that current microcontroller architectures contain a variety of memory types, from masked ROM to EEPROM to static RAM, each of which behaves differently. Things would be simpler if there were just one type of memory on the chip.

But TI's Polen points out an obvious problem with flash in this role. "You could put some data structures in flash along with the code," Polen agrees, "but you have to remember that when you erase a block of flash memory, you lose everything in the block." So a flash-based microcontroller would have to have at least enough SRAM to hold a block of data while the flash memory was being erased. "And if you're going to have that much SRAM on the chip, you might as well put on enough for the rest of the data, too," adds Pope.

So the block-erase nature of flash EPROM limits its use as data memory. This limitation has caught the attention of another vendor with a different technology, ferroelectric RAM developer Ramtron (Colorado Springs, CO).

"If you look at a typical compact microcontroller system," says David

Bondurant, director of marketing, "you see a hierarchy of memory types. For instance, a smart card may have masked ROM, UV-EPROM, EEPROM and battery-backed SRAM all in one system. Ferroelectric RAM can have the operating characteristics of all these devices, so it could condense the whole hierarchy into one memory."

At this stage in its development, the economics of ferroelectric RAM (or FRAM, as Ramtron calls it) for such use remain to be proved. "But we've already established that the process can be no more complex than EE," Bondurant claims.

To get to a prototype FRAM microcontroller, Ramtron will have to go through a number of steps. "First, we need to qualify the FRAMs we're building in conjunction with ITT," says Bondurant. "Then, we need to integrate our PZT capacitor into the source electrode of a transistor, to form a one-capacitor, one-transistor nonvolatile memory cell. After that, we can form an alliance with an established microcontroller vendor to move the FRAM on-chip."

With the economics of flash memory making the technology more attractive, and the possibility of ferroelectric memory on the more distant horizon, it seems a good bet that we'll soon see major changes in microcontroller program store. Surprisingly, given the advantages offered, most of the enthusiasm today seems to come from the vendors rather than from customers.

"So far, the idea has come mostly from the memory guys. We haven't heard much about flash from customers," says Pope. But since flash microcontrollers offer their benefits only in return for some changes in design style and production habits, the subject might bear some consideration by customers, too. □

For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate number on the Reader Inquiry Card.

Hitachi America
(415) 244-7191 ..... **Circle 221**
Intel
(800) 548-4725 ..... **Circle 222**
Ramtron
(719) 594-4455 ..... **Circle 223**
Texas Instruments
(713) 274-2370 ..... **Circle 224**
Zilog
(408) 370-8015 ..... **Circle 225**

Workstation chip set bridges Micro Channel to RISC and CISC CPUs

Warren Andrews, Senior Editor

The recent rash of new RISC and CISC microprocessors—culminating in an announcement by Intel (Santa Clara, CA) of its 80860 and 80486—ensures a new generation of personal computer and workstation performance. And while the high-end workstations will find these processors riding on VMEbus, Multibus II, proprietary buses and maybe even the Next-Generation Architecture—Futurebus+—the low end will undoubtedly jump on the PC buses: AT bus, EISA, Micro Channel and NuBus.

Though it hasn't been determined which bus will be the 32-bit successor to the ubiquitous 16-bit AT bus, there's a strong indication that IBM's Micro Channel is rapidly gaining ground on the yet-to-be delivered EISA bus. As far as NuBus, it has yet to catch on outside the Apple environment. But Micro Channel has both the integrity and performance to emerge as the victor, says Charles Smoot, general manager and executive vice-president of Bull Micral of America (Roseville, MN), a division of Honeywell Bull.

"Micro Channel provides an I/O environment that is sufficiently disciplined for distributed computing and has high-enough performance to meet workstation-level needs," says Smoot. Satisfying PC applications under DOS and OS/2 as well as those under the Unix environment called for by entry-level workstations, however, is difficult because the requirements are somewhat incompatible.

Bull Micral has addressed these design problems and developed the Slik-9000 (Short-Line Interface Kernel) family of chip sets. The Slik family lets OEMs assemble either Intel instruction-set processors or RISC machines on Micro Channel using basically the same chip set.

With the Slik chip set and firmware, OEMs can build Micro Channel workstations that can support both the Intel 80486 and another RISC processor. "Our first target RISC machine is the MIPS R3000 because of its deliverable performance," says Smoot. Bull plans to market the product as software, rather than as chip sets, however, and will let potential OEMs customize the chip set to specific applications and processors to differentiate their products.

The end products are basically 20-Mips-and-up workstations that gain the leverage of the Micro Channel environment. One advantage in addressing that market is that IBM hasn't traditionally been an effective force in the workstation business, thus leaving the Micro Channel workstation business wide open. "The RT hasn't been as strongly received as the Sun- or Apollo-class workstation products," says Smoot.

A different approach

The cost involved in supporting Unix products lies not so much in the CPU and the RAM, but more often in the entire I/O environment. This environment includes packages, power supplies, add-in boards, coprocessors and so forth, all of which need to be qualified. The Micro Channel approach solves all of those support problems with a cadre of manufacturers producing everything from packages to coprocessor add-in cards. This reduces the workstation costs to the same level as those for PCs.

Realistically, an ideal strategy for potential microcomputer/workstation vendors would be one that permits supporting two very similar products in the same Micro Channel architecture—one with an 80486 and the other with a RISC processor. The add-on cards and the rest of the system remain the same. The target application is really what determines whether a customer leans to Intel instruction sets and 80486s, or a predominantly Unix environment where the "bang-for-buck" leader is likely to be a RISC product.

In the Unix environment, development and support costs can run very high. Bull Micral's approach significantly reduces these costs by supporting both the 80486 and the

INTEGRATED CIRCUITS

R3000 architectures, says Smoot. The chip sets are essentially the same, though there will be some bond-out differences between the memory subsystems for the two processors.

"Our implementation of the Micro Channel architecture is very different from that of IBM and other Micro Channel developers," says Smoot. "IBM designs its CPUs by taking the local bus of the processor and building around that and then interfacing to the I/O channel. Each time a processor is changed, all the other subsystems have to be significantly modified to make the new product. IBM's transition from the 286 to 386 processor, for example, was really quite traumatic, judging from the system components that had to be changed."

"One advantage of working off someone else's bus standard," continues Smoot, "is that we were able to build our chip set around the Micro Channel architecture alone, with a simple and independent interface to the microprocessor/memory pair. That's why we were able to build a single chip set that supports the 286, the 386 and the 386SX."

Other vendors have adopted the IBM-style architecture, which hangs the bus on the local CPU pins instead of building the system around the Micro Channel as Bull has done. The transition to go to higher performance processors is thus directly related to whether one has to develop a unique interface to the bus for each different processor used.

The key to Bull's chip set approach is that the interface is adapted directly to the Micro Channel bus itself with a 10-MHz heartbeat—it sits in the asynchronous end of the system, which doesn't change much regardless of the processor used. "What we're really doing is building new pairs of CPU interfaces and memory controllers for the different processors that work within the system, says Smoot.

Instead of working from scratch, as most designers do whenever they contemplate switching to a new processor, "we've wrapped a ribbon around the Micro Channel architecture and its specification and put it in five out of the six parts in our chip set," says Smoot. The interface bridge part is consistent for the CPU and memory regardless of which CPU is used, and it interfaces to the Micro Channel environment, which supports the DMA, peripheral I/O, local I/O devices and video.

Not for PCs only

Although IBM has promoted Micro Channel as an I/O channel for personal computers, the benefits of Micro Channel are really yet to be delivered to the PC marketplace. The real needs for Micro Channel in the business environment haven't emerged; OS/2 hasn't really arrived, and the applications haven't caught up with the available performance.

Smoot says that isn't true in the workstation side, though. "We've been fighting a losing battle to move PCs into workstation niches in the AT environment. We couldn't do it because the integrity of the system wasn't good enough, and the performance features just weren't there—particularly relating to DMA," he says. EISA may help the performance issue somewhat in providing faster DMA and 32-bit buses, but it won't help the integrity side.

With the Micro Channel and with IBM's rigid demands for conformance to specifications, Micro Channel is emerging as a major standard. But, paradoxically, there are fewer benefits in the PC business than in the workstation business, says Smoot. The workstation market demands increasingly higher performance, and there's a greater need for specialized coprocessors than in the general-purpose business and office environment.

"We think there are two good opportunities there: one for OS/2 workstations, which tend to have a business and office flavor and strong support in software, and the other for those applications where Unix is viable," says Smoot. "In the Unix world, the same workstation will be able to deliver 1.5 to 2 times the Mips performance for the same dollar with a RISC processor as opposed to an Intel system processor." □

THE SHORT-LINE INTERFACE KERNEL

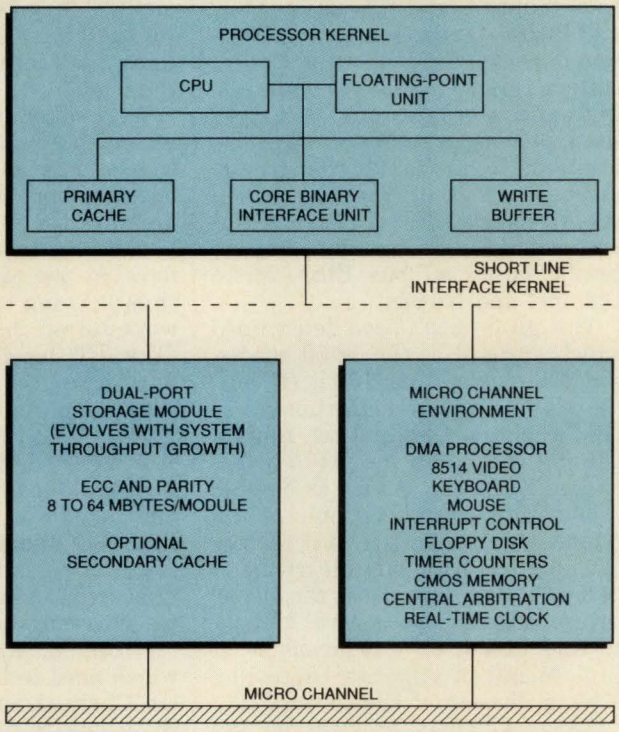

Bull Micral's Short-Line Interface Kernel (Slik) provides a smooth interface from IBM's Micro Channel to a variety of different RISC and CISC processors. The basic Slik function comprises a synchronous CPU interface, high-performance clock generation and distribution, a 32-/64-bit data bus, a 32-bit address bus, byte parity carried throughout the system, optional ECC/EDC, a multiport memory controller and 128-bit burst memory access. Slik also features similar timing to IEEE 488 and efficient mapping of MIPS R3000 bus.

For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate number on the Reader Inquiry Card.

Bull Micral of America
(612) 633-5660 ..... **Circle 259**
Intel
(408) 765-1843 ..... **Circle 260**

Run the new InterTools™ compiler. It's like winning the hundred yard dash by ten yards.

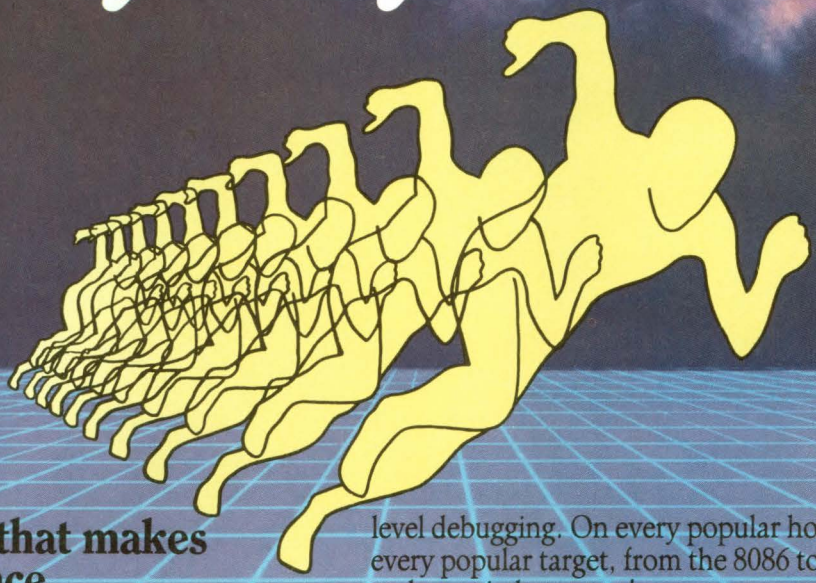

The 10% edge that makes a huge difference.

If you're in the race to develop software for embedded systems, lots of little advantages can make a big difference.
With InterTools, that's exactly what you get.
Take our new InterTools C Cross Compiler for the 68000. It gives you 10% to 20% more compact code. So your application runs faster–up to 27% faster–than code from the most popular compiler!*

Work in the fast lane.

And that's just for starters. Because InterTools is a full set of software tools designed to help you win in every phase of embedded microprocessor development. From compiling and assembling, to source-level debugging. On every popular host. And for every popular target, from the 8086 to the AM29000–and most in between.†

Get complete support.

Best of all, InterTools is from Intermetrics, the only major company that designs and builds every one of its toolkit components to work together as a fully integrated system. InterTools is a total solution, from a team of experts committed to supporting you every step of the way.

Put us to the test!

Just call 1-800-356-3594‡ right now. We'll send you a free InterTools demo kit–and help you get your next C cross development project running ahead of the pack.

The big difference in embedded systems development.

*And we have the benchmarks and data on file to prove it!

†• 68000/008/010/020/030; 68HC11/00/05/09; Z80/64180
• 8086/186/286; 8051; V20/25/30/33/40/50
• AM29000; V60/70; VAX; PDP-11; and many more.

‡In MA (617) 661-0072.

© 1989 Intermetrics, Inc. All rights reserved.
InterTools is a trademark of Intermetrics, Inc.

Intermetrics

733 Concord Avenue, Cambridge, MA 02138
FAX (617) 868-2843

CIRCLE NO. 15

Using the new HP 64000 AxCASE speeds

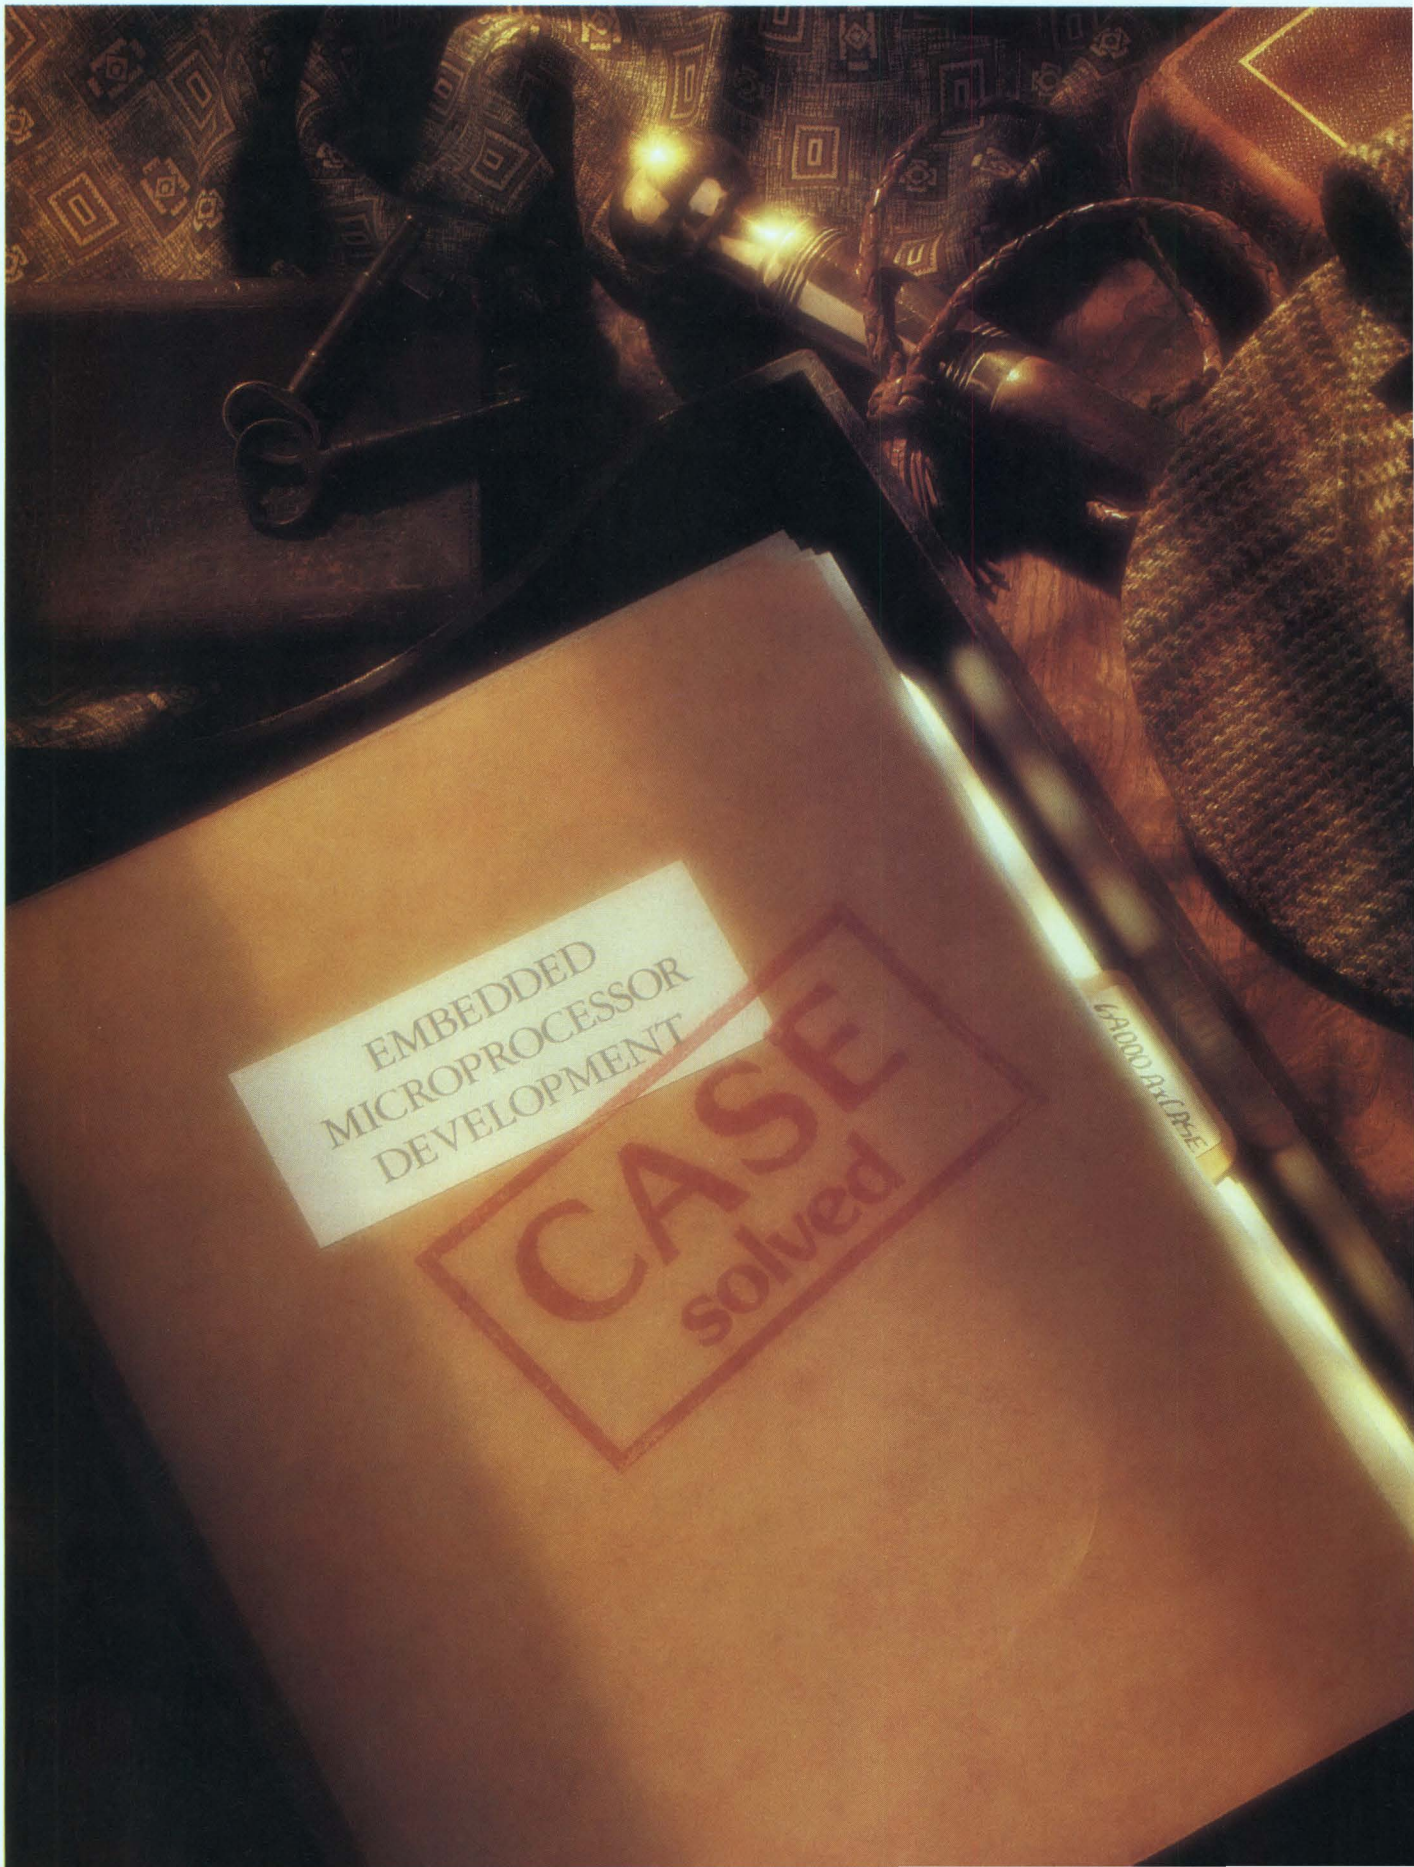

©1989 Hewlett-Packard Co. DS15911/CD

µP development to its logical conclusion.

HP's AxCASE solutions create a complete environment to help you solve embedded software problems fast. And, these days, speed is critical. Because time-to-market can have a big impact on the profitability of your next product. When another month spent testing and debugging represents another million or two in lost revenue, it's easy to see why management takes more than a passing interest in development schedules.

Powerful new tools are part of the HP 64000 AxCASE package. They begin with HP Teamwork and continue with AxLS Cross Compilers that deliver optimized code. Debugging and system integration is accelerated with the new Advanced Cross Debugging Tools (AxDB). You start debugging immediately on the new debugger simulator. As you move to real-time verification, switch to the emulator debugger. Both tools use the same interface and access the same symbols. Software and hardware engineers can use the tools interchangeably.

Your designs will flow transparently from one development state to the next through construction, testing, system integration, verification, delivery and maintenance. Moreover, this is an open system. It's compatible with your design environment. And backed by HP's worldwide support network.

Learn more about our HP 64000 AxCASE solutions. Because developing without them can be a criminal waste of time. Call **1-800-752-0900** today. Ask for **Ext. 786S** and we'll send the HP 64000 AxCASE Solutions Kit. Better still, call to see it firsthand —after all, it's available now, so put it to work.

HEWLETT PACKARD

CIRCLE NO. 16

INTEGRATED CIRCUITS

Hardwired language processors enjoy renewed attention

Ernest Meyer, Contributing Editor

Hardwiring high-level language (HLL) instructions into a processor is one way to improve the performance of artificial intelligence, graphics and real-time applications. In the past, semiconductor technology limited the degree of functional integration too severely to permit the efficient hardwiring of HLLs. As a result, the design technique was rarely used outside of universities and hardwired HLL processors largely disappeared from mainstream research and development. Recently, however, higher degrees of IC integration have resurrected hardwired HLL processors from their too-early demise.

The promise of higher performance incited the hardwiring of AI languages, such as Lisp and Prolog. Both Lisp and Prolog run comparatively slowly on von Neumann architectures, which aren't optimized for the list processing and concurrent statement evaluation that these languages use in force.

Lisp is nevertheless difficult to embody in hardware because the manipulation of long lists requires extensive on-chip memory. Prolog, on the other hand, is well suited to hardware acceleration. A typical Prolog logical operation requires at least 10 operations on a von Neumann machine, and these operations can be performed in one clock cycle by dedicated hardware.

Hardwired HLLs are also very predictable, a trait that's somewhat surprising. Since a hardwired HLL processor executes the same instructions as those written by the software developer, engineers can estimate how long a procedure will take to run while they're writing the code. This predictability has particular benefits for real-time software development, and hardwired HLLs are finding more adherents in the embedded-computing marketplace than expected by their pioneers who mostly focused on the more grandiose objectives of AI.

One group has taken a novel approach to hardwiring general-purpose HLLs. Pierluigi Civera and colleagues at the Politecnico di Torino (Torino, Italy) recently constructed a Prolog coprocessor that works in conjunction with a standard microprocessor to accelerate the most compute-intensive subset of common Prolog operations. The concept of using a language-specific coprocessor is relatively new. A Prolog coprocessor may preface an era where users can insert other coprocessors, besides floating-point accelerators, into a standard coprocessing socket in their workstations.

Most languages put too many layers between the software engineer and the final product, according to William Holloway, president of VME. Compared to Forth, these other languages are like shock absorbers. "With real-time systems, you want to find the potholes in the road," he says. "You don't want shock absorbers. Forth puts you in touch with the road."

Some manufacturers have already seen the benefit of designing coprocessors for application-specific HLLs, particularly in embedded processors for specific subsets of graphics-oriented functions. Tadpole Technology (Waltham, MA), for example, is readying a windowing engine for accelerating viewbox manipulations in the X Window System. And last year, Weitek (Sunnyvale, CA) introduced a coprocessing subsystem for accelerating Postscript, the standard Laser printer format language from Adobe Systems (Mountain View, CA).

"It's natural for leading manufacturers to seek out particular vertical niches and position products for the specific requirements of the application," says Steve Roe, graphics marketing manager at Weitek. "The real constraint on building such systems is no longer the processor's performance—the issue is the available bandwidth to memory."

Even so, it may be possible to design a standard coprocessing interface that could support much more than the acceleration of floating-point arithmetic. If a standard interface can be procured that provides sufficient bandwidth to language coprocessors, the low demand may not obstruct the commercialization of other elegant hardwired HLL chips, such as the Soar Smalltalk chip developed by William Patterson and colleagues at the University of California at Berkeley.

Nevertheless, since there's still no standard interface that offers sufficient bandwidth to tightly couple a language coprocessor with a standard CPU, language coprocessors must be embedded in large and expensive dedicated subsystems. The development system for the Prolog coprocessor, for example, resides on a separate board with its own instruction and data store. Civera's group has clearly seen the possibility of tighter coupling between a language coprocessor and the system CPU chosen, however.

The Prolog coprocessor uses a simple Harvard architecture, with a separate 32-bit bus for instructions and a 28-bit bus for addresses. As such, it would seem possible to place the coprocessor at a direct-mapped location using 4 bits of a 32-bit address space, and still permit memory addressing on the other 28 bits. "We executed great care in the construction of the bus interfaces because interaction with memory may become the bottleneck," says Civera.

Predictability and performance

In contrast with the hardware issues confronting AI coprocessor chips, hardwired HLL processors have met with slow acceptance, largely due to

BRINGING WORKSTATION 3-D GRAPHICS TO THE PERSONAL COMPUTER.

NTH 3D ENGINE.

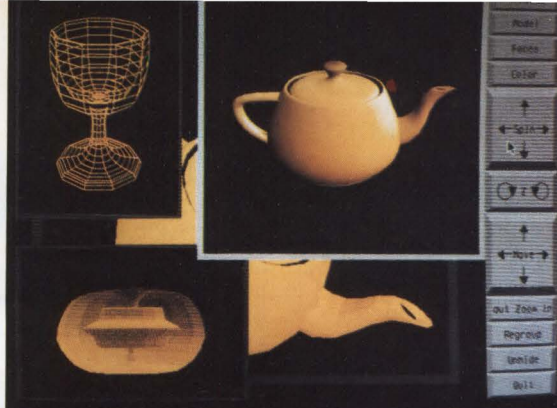

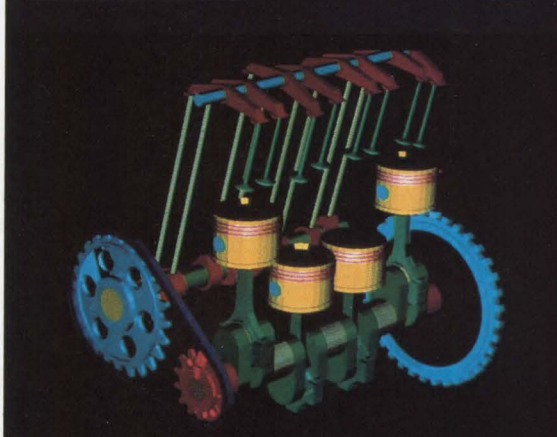

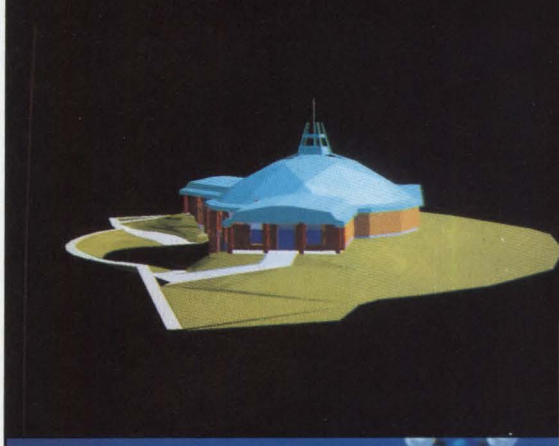

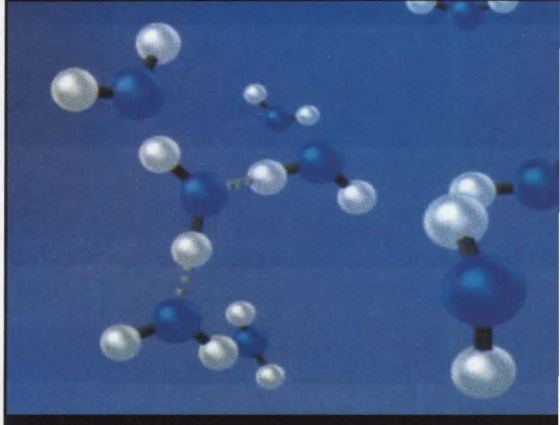

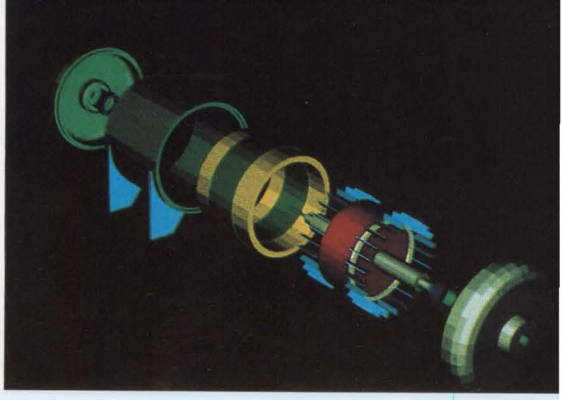

Interactive 3-D Performance on a PC

- Fast 10 MIPS, 3 MFLOPS Graphics Processor on the Board
- 2nd Pipelined Processor Coupled with Custom ASIC for Fast Rendering
- Resolutions from 640 x 480 up to 1280 x 1024
- 256 Colors from a Palette of 16.8 Million
- Standard On-Board 3-D Display List and Optional Z-Buffer
- On-Board Hardware Floating Point Math

Realistic Display of Complex Objects and Data

- Flat Shading at a Continuous 5000 polygons/sec
- Gouraud Shading at a Continuous 3000 polygons/sec
- 3-D Vectors Transformed at a Continuous 40,000/sec
- 2-D Vectors Transformed at a Continuous 50,000/sec

Complete with Powerful HOOPS Graphics Software Tool Kit - *On the Board*

- Full DOS 640k Available for Application Software
- Object-Oriented with Hierarchical Segmented Object-Oriented Database
- Hidden Surface and Line Removal, Smooth and Flat Shading Routines, Multiple Light Sources, Transparency, Built-in Windowing and Picking, Event Manager, and a lot more
- Bindings for "C", FORTRAN and Pascal - Easy to Interface

Prices Start at $6995

For More Information, Call 800-624-7552

CAD Users
Circle 17 on Inquiry Card

OEMS/Software Developers
Circle 68 on Inquiry Card

Nth 3D ENGINE™ DISPLAY CONTROLLER

Nth Graphics 1807-C West Braker Ln. Austin, Texas 78758 512-832-1944 (collect in TX)

Nth 3D Engine is a trademark of Nth Graphics, Ltd. HOOPS is a trademark of Ithaca Software.
Models courtesy of Circuit Studios, Inc. and Versacad Corporation.

INTEGRATED CIRCUITS

A PROLOG COPROCESSOR

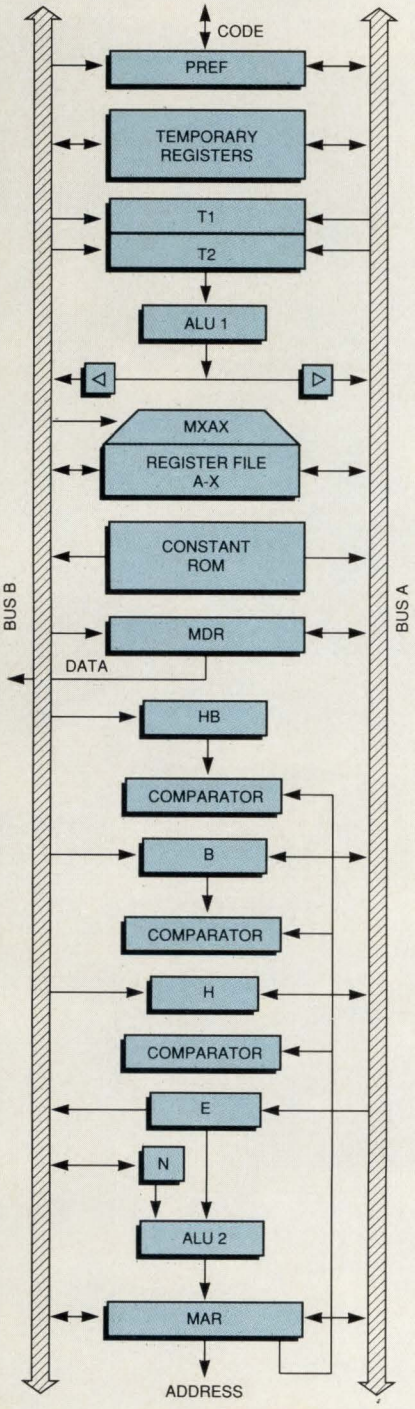

The Prolog coprocessor from the Politecnico di Torino works in conjunction with a standard CPU to accelerate frequently used compute-intensive Prolog operations. The Harvard architecture uses registers that have been transformed into up/down counters to decentralize increment, decrement and zero-test operations.

their lack of development support. This year, however, code and operating system development tools are appearing in force for hardwired HLL controllers, so the chips are likely to become accepted design solutions in the next few years.

Forth is the dominant language in the embedded hardwired HLL court, partly since Forth was originally designed for real-time control. Moreover, Forth is particularly well suited to direct hardware embodiment and offers high performance for relatively low clock speeds.

Unlike virtually all other languages, which are register-based, Forth is stack-based. As a result, context switching is more efficient in Forth than in languages such as C. Changing to a new task in Forth entails little more than changing the stack pointers. In parallel with more efficient multitasking, the stack-based architecture minimizes the operating system required for an embedded application to a kernel structure that typically occupies little more than 4 kbytes of memory.

Similarly, Forth code tends to be more compact than that of other languages, such as C. Federal Express (Colorado Springs, CO), which uses Forth in its "trackers," the handheld paging units carried by its couriers, found Forth's small kernel and code size particularly desirable.

"We were designing a desktop unit, smaller than a telephone, that secretaries can use to call a courier by hitting just one button," says Peggy Mills, a project manager at the Federal Express Cosmos division. "We started with C, but external developers couldn't even fit the code in 128 kbytes, let alone the operating system. We could fit the whole thing in a 64-kbyte ROM with Forth. Also, code development was quick, and the code is far more modular than the C version."

Federal Express used Polyforth from Forth (Manhattan Beach, CA) for code development. Federal Express isn't yet using hardwired Forth in its products, however, primarily because of cost. Instead, Federal Express crosscompiles the code onto the 6303 microcontroller from Hitachi (San Jose, CA) for the final system.

"We've looked at the hardwired Forth chip from Harris, the RTX processor," says Gene Farrar, project manager at Federal Express for the courier-carried tracker units. The RTX chip from Harris Semiconductor (Melbourne, FL) was originally developed by Novix in a gate-array implementation. Harris moved the design into a standard cell, which has simplified the addition of other functions to the Forth core, such as multipliers and counters.

"We were impressed by the speed of the RTX," says Farrar. "Even though it's more expensive than Hitachi's 6303, it offers more flexibility and performance in the long run, so we may use it in the future."

Performance may not seem to be a primary issue in a handheld tracker unit, but a faster processor yields corollary benefits. "We don't need the speed, but we are using battery-operated devices," says Farrar. "Since the chip operates faster, we can power it down sooner and increase the battery life of the system."

The RTX processor is one of two hardwired Forth processors in the marketplace. The other is the SC32 chip from Silicon Composers (Palo Alto, CA). The SC32 is a new 32-bit Forth chip independently developed by Silicon Composers, which also sells the RTX processor.

Using hardwired Forth

A Forth developer that has made use of hardwired Forth is VME (Milpitas, CA). "We replaced a discrete standard-part implementation of a VMEbus DMA controller with Harris' RTX Forth chip, and performance more than quadrupled," says William Holloway, VME president.

Holloway asserts that the chip more than paid for itself in reduced component count and drastically reduced development time. "Since our board needed fewer chips, we were able to fit some extra functions on the board," he says. "We were also able to add data-flow analysis utilities because there was still spare processing power. And if we had needed to revise the design, which we didn't, we only would have had to reprogram a PROM."

In Holloway's view, the combination of Forth and a hardwired processor that runs Forth instructions is ideal for embedded applications. Other programming languages place too many layers between the software engineer and the final product,

obfuscating the actual purpose of the design and increasing the risk of error. "With real-time systems, you want to find the potholes in the road. You don't want shock absorbers," says Holloway. "C is like a shock absorber. Forth puts you in touch with the road."

Growing Forth support

Initial acceptance of hardwired Forth was slow mostly because there was limited software support. That situation is changing. Both Harris and Silicon Composers supply similar development systems for the RTX chip. Forth and VME also supply development boards and systems for the RTX processor, which add customized user interfaces to the interpreter, compiler and debugger from Harris. And Silicon Composers has ported its SC/Forth Development System to the SC32 processor.

For the RTX architecture, the SC/Forth Development System used to offer superior optimization capabilities over the Harris offering. The RTX chip contains four internal buses, all of which can operate concurrently; an optimizing compiler can therefore find neighboring Forth words that can be parallelized and combine them for higher performance and greater code compaction. Harris is now releasing an optimizing compiler that optimizes for this parallelism.

In addition, both Harris and Silicon Composers are offering multitasking operating systems for the RTX chip. The Harris system allows for any number of tasks in either timeslice or non-timeslice modes, and includes the standard features for task scheduling and priority resource management. The Silicon Composers System allows timeslicing with up to 15 levels of task priority: PAUSE statements can be used to force a context switch if one slice is idling.

Harris is also releasing C and Prolog compilers for the RTX. The ANSI-C compiler will be a standalone package with a debugging system that's like Microsoft's Codeview: the C compiler's intermediate language is a Forth-like assembly language that's then translated into RTX code. Developers will use inline Forth for time-critical sections or link in separate Forth modules.

Silicon Composers is also planning a C compiler for the SC32. "Unlike the RTX chip, the top four registers in the SC32 are directly addressable, making the SC32 suitable for general-purpose programming," says George Nicol, president of Silicon Composers. "Caching an off-chip stack removes limits on stack size, making it possible to perform Lisp operations with the SC32."

Harris will counter the SC32 with a new version of the current-generation RTX 2000. "The RTX 4000 will support a lot of C characteristics, such as stack frame addressing," says Mike Mellen, staff engineer in marketing development at Harris. "Developers can use C when performance isn't crucial, yet they won't need to drop into assembly for the critical portions."

To measure how long a routine takes in the Harris development system, the user can run a routine that executes from the PC and uses the PC's internal clock to time the interval between breakpoints with an accuracy of about 20 ms. For more accurate timing, both the Harris and Silicon Composers software provide postoptimization disassembly listings; the user merely counts the lines of code, virtually all of which execute in one 100-ns clock cycle.

The ease by which the run-time of individual processors can be measured on the hardwired Forth chips, coupled with C and operating-system development support, elevates hardwired HLL out of academia into the world of practicality. At the same time, the price of the hardwired Forth chips is dropping to more reasonable levels. □

For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate number on the Reader Inquiry Card.

Adobe Systems
(800) 833-6687 ..... Circle 232
Federal Express Cosmos Div
(800) 238-5355 ..... Circle 233
Forth
(213) 372-8493 ..... Circle 234
Harris Semiconductor
(407) 724-7000 ..... Circle 235
Silicon Composers
(415) 322-8763 ..... Circle 236
Tadpole Technology
(800) 232-6656 ..... Circle 237
VME
(408) 946-3833 ..... Circle 238
Weitek
(408) 738-8400 ..... Circle 239

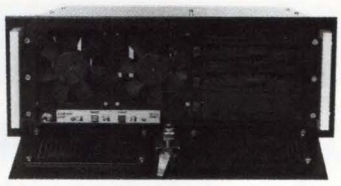

CRM710
Rack Mounted
/AT Compatible System

This ultra rugged 7 inch high, 19 inch rackmount /AT compatible features all welded construction, 10 slot low noise backplane, filtered air system, 3 peripheral slots, internal 3½ inch hard drive, 200 watt power supply, with a lockable front control panel.

PC BUS
SINGLE BOARD /XT /AT
COMPUTER SYSTEMS
Plugs into Passive Backplane

CAT930
/AT 80386 "engine only" board running at 16 or 20MHz with up to 8MB DIP RAM, 80387 socket, keyboard, reset, speaker port and BIOS SHADOWING. An expansion card, CEX706, is available.

CAT900
/AT 80286 computer running at 8, 16, or 20MHz with up to 64KB EPROM memory, 2MB of DIP RAM, 2 serial, 1 parallel port, mouse, speaker, reset port and BIOS SHADOWING.

CXT100
/XT 8088 computer running at 10 MHz, 64KB DRAM, IDE hard disk interface, floppy controller, two serial, one parallel port, seven channels of DMA. Optional 8087 coprocessor is available.

CPC010
/XT 8088 half card size computer (3.9" x 6.2") running at 10MHz with 256K RAM memory, 4 DMA channels, 8 levels of interrupts, keyboard, speaker and reset port. Optional 8087 coprocessor is available.

PACKAGING SUPPORT
4, 6, 10, and 14 Position Card Cages, Backplanes • Industrial & Rack Mount Chassis • Fans • Plug-In Power Supplies • Low Profile Enclosures • Small Footprint Enclosures

Diversified Technology
1-800-443-2667
RIDGELAND, MS 39158 **601-856-4121**
/XT/AT are tradenames of IBM Corp. E-3

CIRCLE NO. 18

SOFTWARE

Real-time development aids extend beyond kernels

Tom Williams, Senior Editor

When it comes to designing embedded real-time multitasking software systems, kernels are no longer enough. This problem becomes greater as real-time control applications increasingly break out of the confines of their "black box" environment and become what Arthur Allen, president of Metasphere (Berkeley, CA), and others call open real-time systems.

Several software companies are starting to address the broader needs of open real-time systems. Help comes with integrated products offering functions beyond a kernel and kernel-oriented development aids to include I/O drivers, user interfaces and file managers.

Deeply embedded real-time systems are the classical black boxes that may be quite complex internally, but they have a limited range of well-defined possible inputs from the outside and are expected to perform sharply defined functions. Open real-time systems, on the other hand, must maintain exact real-time control over their applications but also interact with human operators and mass-storage systems. They must deal with data in unpredictable sizes and formats.

Real-time executives, or kernels, such as VRTX by Ready Systems (Sunnyvale, CA) or PDOS by Eyring Research (Provo, UT) are at the heart of real-time control. A large open multitasking control application, however, may involve some 75,000 lines of C code, points out Yoav Agmon, president of Realtime Performance (Sunnyvale, CA). Of that, the kernel itself represents only 2,000 to 5,000 lines. Writing the rest of the I/O routines and user interface code as well as the application has until now been left up to the system designer.

Realtime Performance has developed a set of software modules, called RPCore, to simplify the development of real-time control applications, graphical user interfaces and generic I/O drivers, all of which can be adapted to a wide variety of specific needs.

At least two other companies are supplying extended development aids to real-time kernels. Metasphere has developed a set of extensions to the Ready Systems VRTX executive that help VRTX work in open real-time applications and connect to user interfaces and storage systems. In addition, Digital Research (Monterey, CA) supplies a real-time operating system called FlexOS that works with a graphics user interface called XGEM for use in real-time control applications. XGEM can also be used to extend a kernel such as VRTX.

Generic software components

The RPCore system is a set of routines and data structures divided into functional groups called managers. These include a graphics manager, screen and window managers, a memory manager, an I/O manager, an alarm manager, drivers and a real-time kernel. "The managers represent a generic set of software components that can be tailored to a diverse field of many different machines with many options," says Agmon. "But we had to ask ourselves what was common among all those different machines."

As a result, RPCore recognizes three basic core areas that all machines share. All require a real-time user interface and the software to control it. Likewise, all must have many different kinds of different I/O and the ability to adapt and integrate that I/O so the machine will have easy and uniform access to data from sources as divers as robot arms or radiation counters. Finally, all machines need a real-time kernel with multitasking services.

The user interface replaces the old control panel with its switches, meters and knobs. That panel is now a CRT, often with a touch screen. The CRT display may have windows, or at least a hierarchy of different dis-

RPCORE'S SOFTWARE MODEL

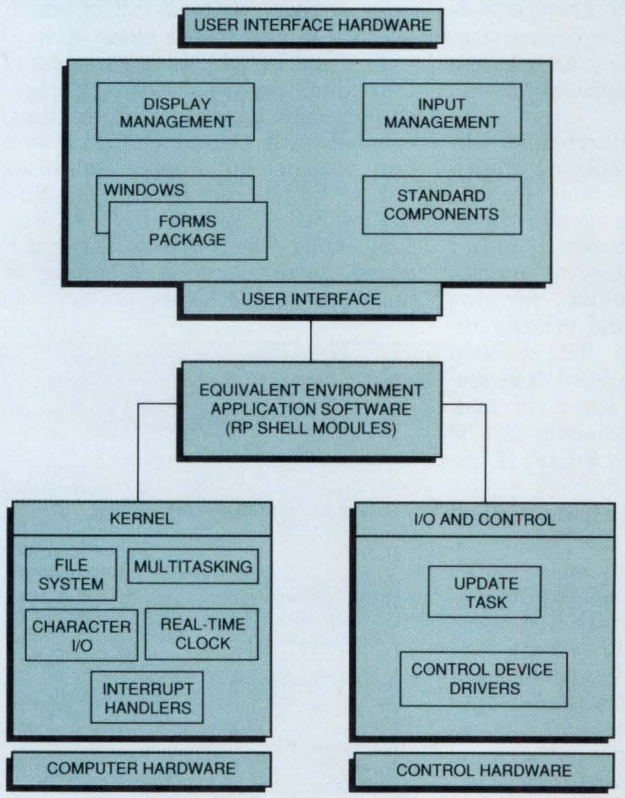

RPCore from Realtime Performance supplies standard software components to control three core elements of embedded real-time systems: the user interface hardware, the computer hardware and the I/O control hardware. The components can be adapted to work with real-time kernels supplied by other vendors, and the I/O components are supplied for popular standard buses including VME, Multibus and STD bus. In addition, RPShell supplies examples and aids for developing the actual application.

plays that appear when needed.

Still, the graphics requirements for real-time control are different from those of the office. The user has no need to move and resize windows, but windows and the items in them—such as status displays and alarms—must respond very quickly, even though they aren't the most time-critical parts of the system. There must thus be an easy way to link graphics entities on the screen with control and interrupt service routines and to set priorities for serving graphics needs.

In the RPCore system, the user interface (as well as the disk manager) runs on an MS-DOS-compatible platform and is connected to one or more I/O controllers via a communications link such as RS-232, RS-485 or Arcnet. The I/O controllers contain microprocessors running the real-time kernel, the application control software and the I/O drivers.

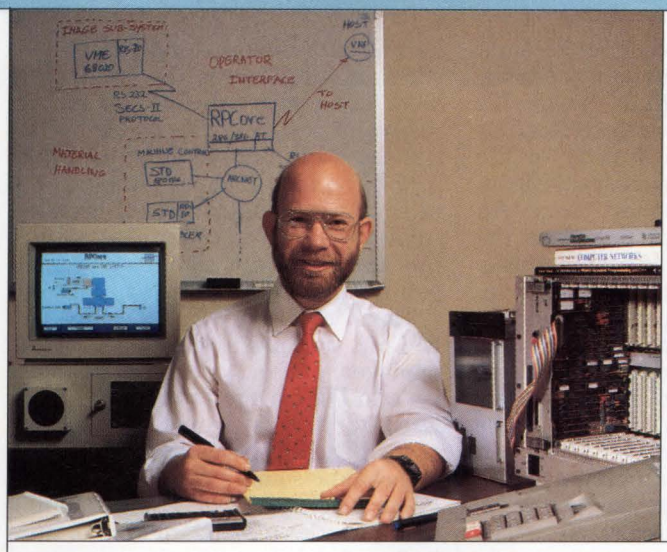

Realtime Performance's set of software modules, called RPCore, eases the development of real-time control applications, graphical user interfaces and generic I/O drivers. The software components can be tailored to many different machines, according to Yoav Agmon, president of the company.

In most cases, communication flows between the user interface and the control application via the user interface data base manager. In special cases, though, links can be set up between the user interface—essentially, graphics and touch-screen or mouse routines, and system code at the manager, device driver or interrupt service routine level.

Mission Ready – Military VME from DY-4

Providing the best in performance, reliability and cost-effectiveness, DY-4 introduces its military VME (DMV) family of fully MIL-Spec VME boards and chassis.

When you're looking for the best in military VME, consider these facts about DY-4's DMV family:

- designed to meet MIL-E-4158, MIL-E-5400 and MIL-E-16400 specifications,
- conduction-cooled for survival in the harshest environments,
- full range of functionality: SBCs (680X0), serial and parallel I/O, A/D, 1553B, SCSI and ATR chassis,
- Ada™ and RTOS support packages
- military-grade, off-the-shelf boards with commercial equivalents to reduce system development costs

™ Ada is a trademark of the United States Department of Defense

Call us with your system requirements. Now you can benefit from cost savings using industry-standard VME, without compromising military specifications.

California: (408) 377-9822
Los Angeles: (714) 549-2559
Massachusetts: (508) 692-9308
Canada: (613) 596-9911
Denmark: 45 6 963624
Germany: 49 89 3171180

dy4®
DY-4 SYSTEMS INC.

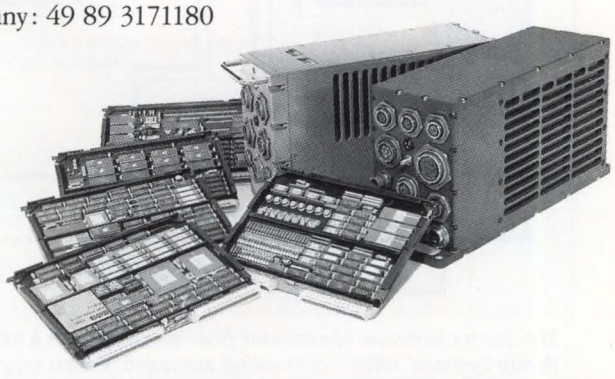

CIRCLE NO. 19

SOFTWARE

Separating the user interface on a DOS platform isn't the only possibility, however. Digital Research's XGEM graphics user interface, for example, interfaces directly to the FlexOS real-time operating system via a shared run-time library. An application manager linked to the library keeps a list of running tasks and makes sure that XGEM responds with the proper displays associated with those tasks.

The point of the direct connection isn't so much speed as cost. A few milliseconds is sufficient time for the user interface to respond because it depends on human interaction, which can't be expected to be extremely fast or predictable. "The graphical interactions aren't usually all that intense in process control systems—things like updating a trend plot, responding to a button being pushed," notes Metasphere's Allen. "You can certainly do that in a fraction of a second in a multitasking real-time system."

Allen is examining different graphical interfaces for possible use with his Vextra product, which is itself a set of extensions to the VRTX kernel, but would prefer to integrate such an interface at the interrupt level as a server process. The main reason is cost. A DOS platform to run the user interface might be fine in a $1 million piece of equipment, he says, but in an instrument even a small increment in component count can translate into significant competitive pricing issues.

The focus points for I/O services in RPCore, for example, are the I/O manager and a small number of generic I/O drivers that can be tailored to a wide range of specific needs. The I/O manager organizes and channels data from the system's I/O to its client tasks. The I/O manager provides a standard system interface for drivers. It also provides for simple simulation of I/O routines. The interface software can thus be exercised in the development stages without being connected to the physical hardware of the machine.

Real-time kernel still important

All this talk of the other software components needed by real-time control systems isn't meant to downplay the importance of the real-time kernel, also called the multitasking executive. RPCore supplies a real-time kernel along with memory-management functions that can be linked into the application software.

But Agmon stresses that Realtime Performance isn't primarily a kernel supplier. "We recognize that there are some very common services and if our components need to run with other kernels, we're very happy to do it," he says, stressing the point that the kernel is a small portion of the total code needed by an embedded real-time system.

Metasphere's Vextra focuses on adding robustness to the Ready Systems VRTX real-time executive. VRTX works very well alone in deeply embedded systems. Vextra adds memory management, and process creation/termination services that can handle exceptions and unpredictable conditions coming from the outside world. Parent processes can spawn and terminate child processes, for instance. But a process can raise a shield against termination if it's in the middle of an operation that could corrupt the system if terminated early. Even if completion yields unwanted results or if an illegal operation takes place (division by zero, for example) the system leaves "clean" errors that can be handled without killing the system.

Vextra's extensions to the VRTX kernel try to promote an open kernel architecture, according to Allen. "You want to be able to assemble an operating system environment that may include graphics, file and I/O managers, network management, and so forth, out of building blocks," he says. He adds that while Vextra isn't the end-all, it does provide a foundation to encourage designers to develop robust system software. The next step is to attach graphics and I/O as building blocks as well.

The concept of building blocks for real-time embedded control is starting to spill over into the actual application programming. Allen says Metasphere is planning on providing a source-code skeleton that can plug into the Vextra services. "If you adhere to our rules, you can develop application-specific components quickly," he says. This skeleton would be source-code examples that system developers could copy, modify or use as examples in producing

VEXTRA ARCHITECTURE

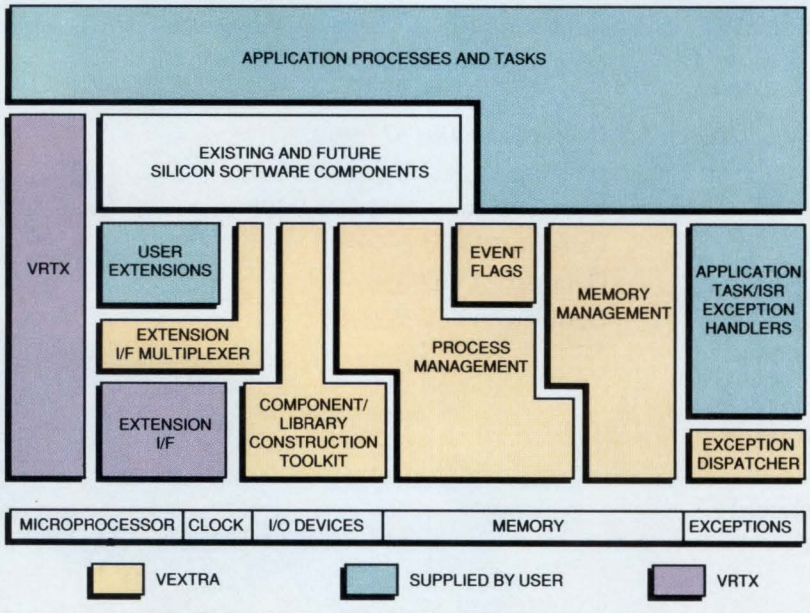

The Vextra software component from Metasphere is a ROM-based set of extensions to Ready Systems' VRTX multitasking executive. Vextra supplies 33 additional system calls to VRTX for memory management, process creation and termination, and exception handling. Additional components, such as a real-time graphics kernel or I/O routines, can interface easily to the open Vextra architecture.

FLUKE AND PHILIPS – THE GLOBAL ALLIANCE IN TEST & MEASUREMENT

FLUKE®

PHILIPS

PHILIPS

Finally! The one thing you've always wanted in a low-cost logic analyzer: No restraint.

The Philips PM 3655. It's the first *true* 100 MHz logic analyzer for under $10,000! That means it comes with a feature no other low-cost analyzer in the world can offer you: Freedom to use 100% of its capability 100% of the time.

UNRESTRAINED PERFORMANCE

- *High-speed timing and state analysis:* Get full 100 MHz synchronous and asynchronous sampling on up to 96 channels. And never again sacrifice width for speed.
- *Deep memory:* 2K memory depth per channel lets you measure complex data flows—with no loss of width, speed or glitch capturing potential.
- *GLITCH mode:* 5 ns glitch capture at 100 MHz on all channels. Again, without forfeiting speed, width or depth.
- *Flexible triggering plus Trace:* Powerful, multi-level triggerword sequence, range recognition and four clock qualifiers help you find and display exactly the data you're looking for.
- *Expandability:* Add 24-channel modules any time for 24, 48, 72 or 96 channel configurations.
- *Broad support:* Choose from a wide range of disassemblers and dedicated personality modules.

SIMPLE OPERATION PLUS PC

- *Immediate familiarity:* All this capability and so easy to use—menu-driven operation; extensive user-defined labeling; and standard PC-style keyboard interface.
- *Plus full MS-DOS-based PC:* Adds versatility and value without compromising logic analysis. On-line disk storage for measurement data and instrument set-ups.

Expansion slots for extra memory, modem or other PC options.

UNRESTRAINED SUPPORT

Count on a one-year warranty and all the application and service assistance you'll ever need. From Fluke—the people who believe that extraordinary technology deserves extraordinary support.

TEST THE DIFFERENCE

Call Fluke today at **800-44-FLUKE ext. 77**. And look into a logic analyzer that unleashes your potential instead of tying your hands.

John Fluke Mfg. Co., Inc., P.O. Box C9090, M/S 250C, Everett, WA. 98206
U.S.: 206-356-5400 CANADA: 416-890-7600
OTHER COUNTRIES: 206-356-5500

© Copyright 1988 John Fluke Mfg. Co., Inc.
All rights reserved. Ad No. 0581-P3655

PM 3655 • 100 MHz • LOGIC ANALYZER

FLUKE®

CIRCLE NO. 20

How P-CAD Put A Smile On Over 10,000 PCB Designers

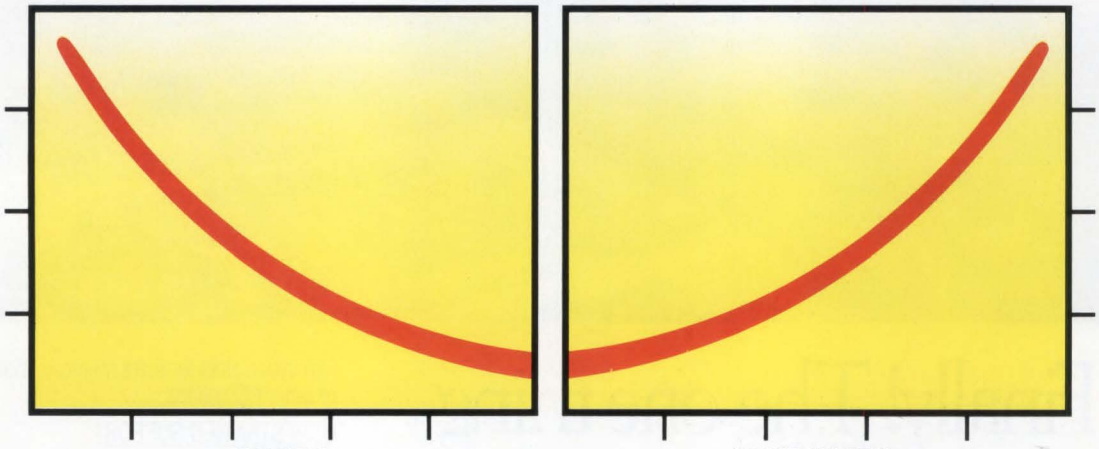

Lower costs and higher productivity add profits to your bottom line.

Work with P-CAD® and you'll be smiling every time you finish a design or a fiscal year.

Because we know successful PCB design requires more than software alone.

That's why we offer the industry's most complete PCB design solution. Including training. Superior tech support. User groups. Bug reports. Plus, the best maintenance, upgrades, user guides, and reference documentation you'll find.

All elements that help reduce redesign, project delays, rework, and manufacturing costs. But even all that wouldn't mean much if our system didn't deliver. And it does!

P-CAD software is modular, with a consistent system command structure. Learn one and you know how to use all our modules. That means greater productivity and lower training costs.

P-CAD also works with all current technologies—SMT, analog, high speed, and fine line. So designers who use P-CAD will always be the most competitive.

You get a complete, upward migration path, too. As your requirements grow, you can take your P-CAD investment with you. No throw-away software. No obsolescence.

Architecture? It's "open" for easy integration with CADAM,® AutoCAD,® Ventura,™ Interleaf,™ and other CAD environments. So your options are open now, and in the future.

Finally, with over 10,000 users in 37 countries, we have the largest installed base in PCB design. So when you're looking for experienced P-CAD designers, there's no need to look far. They're everywhere.

Let us show you how increasing PCB productivity can enhance your profits. Call now for details and a free copy of our **PCB CAD Buyer's Guide:** 800-523-5207, in California, 800-628-8748.

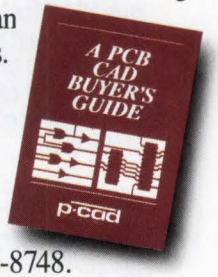

What we have to say will make you smile. P-CAD, 1290 Parkmoor Avenue, San Jose, CA 95126.

p-cad®
PERSONAL CAD SYSTEMS INC.

P-CAD is a registered trademark of Personal CAD Systems, Inc. CADAM is a registered trademark of CADAM, Inc. AutoCAD is a registered trademark of Autodesk, Inc. Ventura is a registered trademark of Xerox Corporation. Interleaf is a trademark of Interleaf, Inc.

CIRCLE NO. 21

SOFTWARE

an actual application.

Realtime Performance appears to already have a generic application example available in a companion product called RPShell. RPShell is a running application that uses the RPCore managers to provide a basic, generic control system. It consists of a hypothetical machine with valves, sensors, gas bottles and chambers, alarm handling and so forth.

The graphics user interface consists of five screens related to this fictional machine: Status, Process, Direct I/O, Plot and Data Log screens. The user can duplicate or modify each of these screens and its associated C source code to produce a real system, or may use the generic user interface as an example for independent code development.

Networking control systems

Deeply embedded real-time control applications in black boxes aren't normally written with local area network communications in mind. They simply concentrate on the control problem. The addition of a user interface, however, makes networking real-time systems a fairly straightforward problem. A user interface running under DOS can easily take advantage of DOS-based LAN interfaces, which are available in great variety, as long as time-critical issues are respected.

When the user interface is running on a platform other than that of the control application and is based on a multitasking operating system, more opportunities arise to use the system's data by feeding it into spreadsheets, graphing packages or files running on other systems on the LAN. Realtime Performance is developing an OS/2 platform that will be able to connect to Ethernet and other LANs.

As long as the primacy of the real-time control function is kept uppermost, many other solutions may be possible. "If you're time-sharing the display with intense process control activities, you'll have to provide sufficient processor bandwidth to accommodate the overall activities," cautions Metasphere's Allen. Those overall activities are bound to expand as designers discover new potentials. As Realtime Performance's Agmon notes, "Our customers have been educated very fast." □

For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate number on the Reader Inquiry Card.

Digital Research
(408) 649-3896 ..... **Circle 240**
Eyring Research
(801) 375-2434, ext. 360 ..... **Circle 241**
Metasphere
(415) 540-7015 ..... **Circle 242**
Ready Systems
(408) 522-7334 ..... **Circle 243**
Realtime Performance
(408) 245-6537 ..... **Circle 244**

Puzzled Over GOSIP?

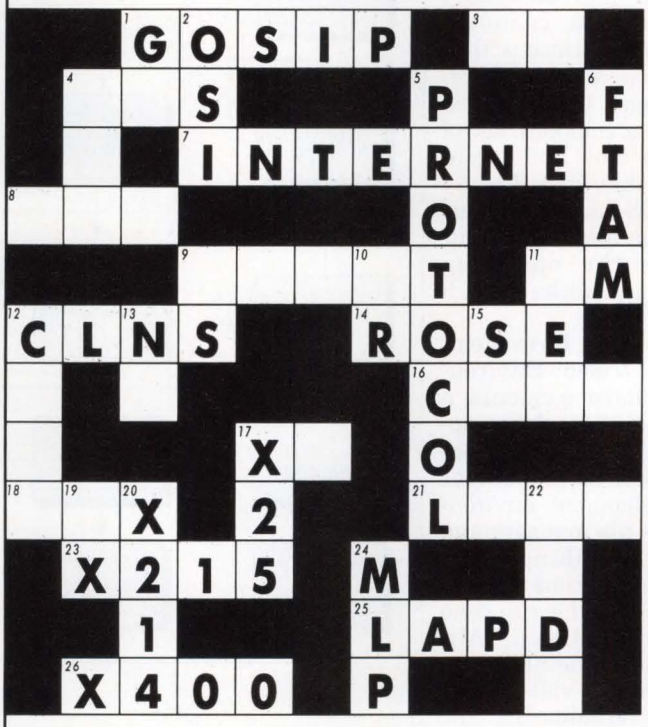

Across
3. . . . Class 0/2, 0/2/4
4. A Session subset
8. Session subset required by FTAM
9. Int'l stds. body
11. Mantra
16. Main net function
17. PAD std.
18. M or F
21. CCITT T.71

Down
1. Gigacycle (abv.)
4. Session subset required by X.400
10. Translation (abv.)
11. Major computer co.
12. Provided by X.25
13. Chance of selling non-GOSIP system to gov't after 8/15/90
19. Former
22. Link async devices to X.25 nets

Not sure how you're going to have a GOSIP-compliant system ready for the federal market by August 15, 1990? We have a four letter solution to that problem: TITN. We provide a GOSIP protocol stack that ports like a dream so you can get your product out the door ASAP.

TITN has eight years of experience in the OSI source code business. Our customers get the benefit of 125 man years of software development and testing at a cost effective price.

We have CCITT 1988 X.400 with a Message Store and User Agent, FTAM, ACSE, ROSE, Presentation, Session, Transport, IP and X.25 source code modules. Our products' rich CCITT and ISO standards implementations give our customers plenty of product development options. Join over 80 of the best names in the computer industry. For the solution to your OSI product development puzzle, call TITN today.

HQ/Western Region Sales
24301 Southland Drive
Suite 200
Hayward, CA 94545
Tel: (415) 785-5970
Fax: (415) 785-1149

TITN

Eastern Region Sales
1300 North 17th Street
Suite 1530
Arlington, VA 22209
Tel: (703) 528-2662
Fax: (703) 528-0388

CIRCLE NO. 22

SOFTWARE

Graphics/audio software simplifies man-machine interface in real time

Warren Andrews, Senior Editor

Real-time computer applications, particularly those in industrial and process control, have traditionally been burdened with an awkward, at best, user interface. Nonexistent development tools and a lack of standard run-time support have prevented designers from improving such an interface. Now, a development and run-time software package lets designers write and integrate a complete graphics and audio interface with a leading real-time operating system.

Microware Systems (Des Moines, IA) has developed a complete set of presentation editor tools and a run-time environment that links directly to the company's OS-9 real-time kernel. Designed to operate on any platform supporting the operating system, the package works with a variety of I/O devices.

Microware's approach, Rave (Real-time Audio and Video Environment), has three major elements: a presentation editor, a graphics support library and a graphics file manager. The presentation editor provides the development environment; the graphics file manager and support library combine with Microware's OS-9 real-time operating system to provide the run-time environment. In addition, Microware has provided the ability to control audio as well as video in the system, therefore eliminating the classic bell, buzzer and beeper approach.

Rave runs on top of OS-9 and takes advantage of the system's programming features. It will run on almost any platform supporting OS-9, including most of the popular 68XXX-based VME platforms. Audio and video support is provided by a number of popular devices readily available on standard platforms. In addition, Rave doesn't require a separate PC or Unix host. OS-9 supports editing, compilation and debugging as well as full Rave development and execution directly on a target system.

The appeal of such a graphics system is enhanced by the widespread acceptance of newer input or pointing devices, most notably the touch screen, though alternate pointing devices such as the track ball and the light pen are also widely used. The combination of real-time response with an interactive interface medium not only simplifies operator control, but in many cases reduces the risk of error.

Without writing any native code, designers can use Rave to create a custom audio/visual user interface that's fully integrated with a real-time application. Objects can be created by using components already in

THE RAVE SOFTWARE PACKAGE

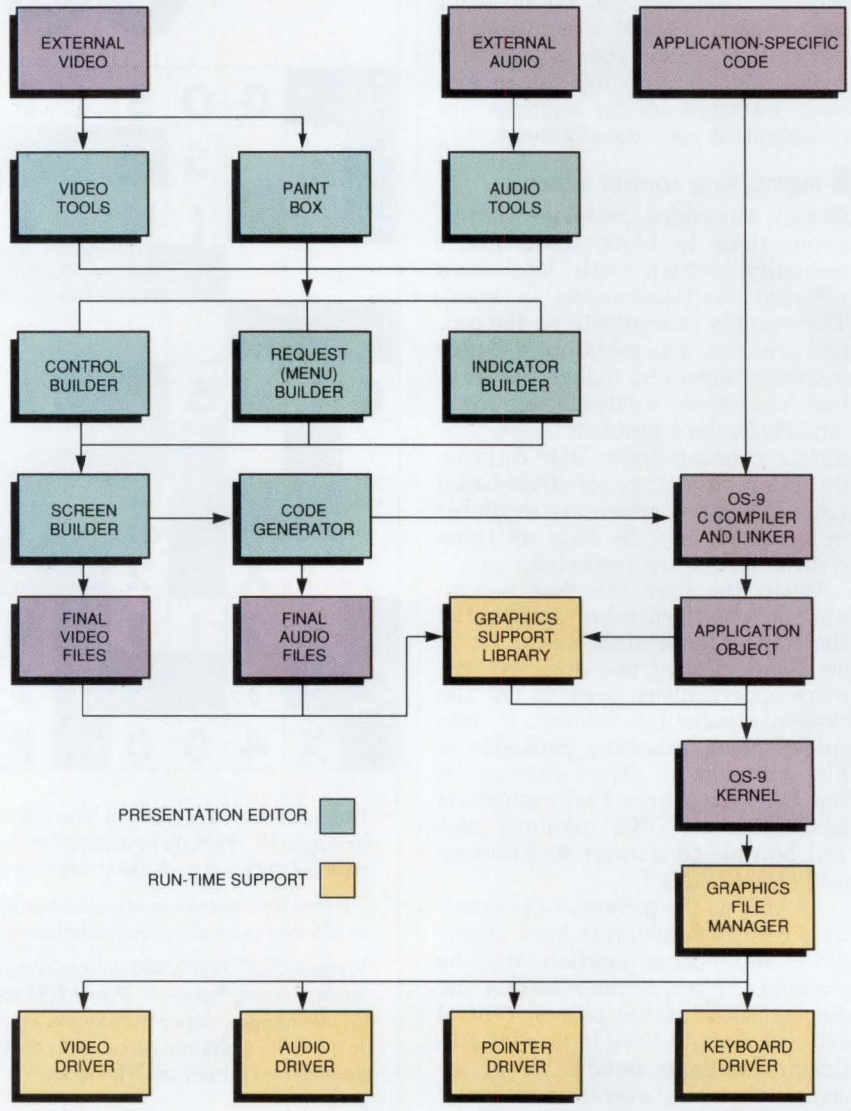

Microware's Rave package consists of three basic functions: the presentation editor, a graphics support library and a graphics file manager. Run-time support is provided with the graphics support library and the file manager combined with the various device drivers. The presentation editor comprises the audio and video interface as well as the drawing tools and code generator.

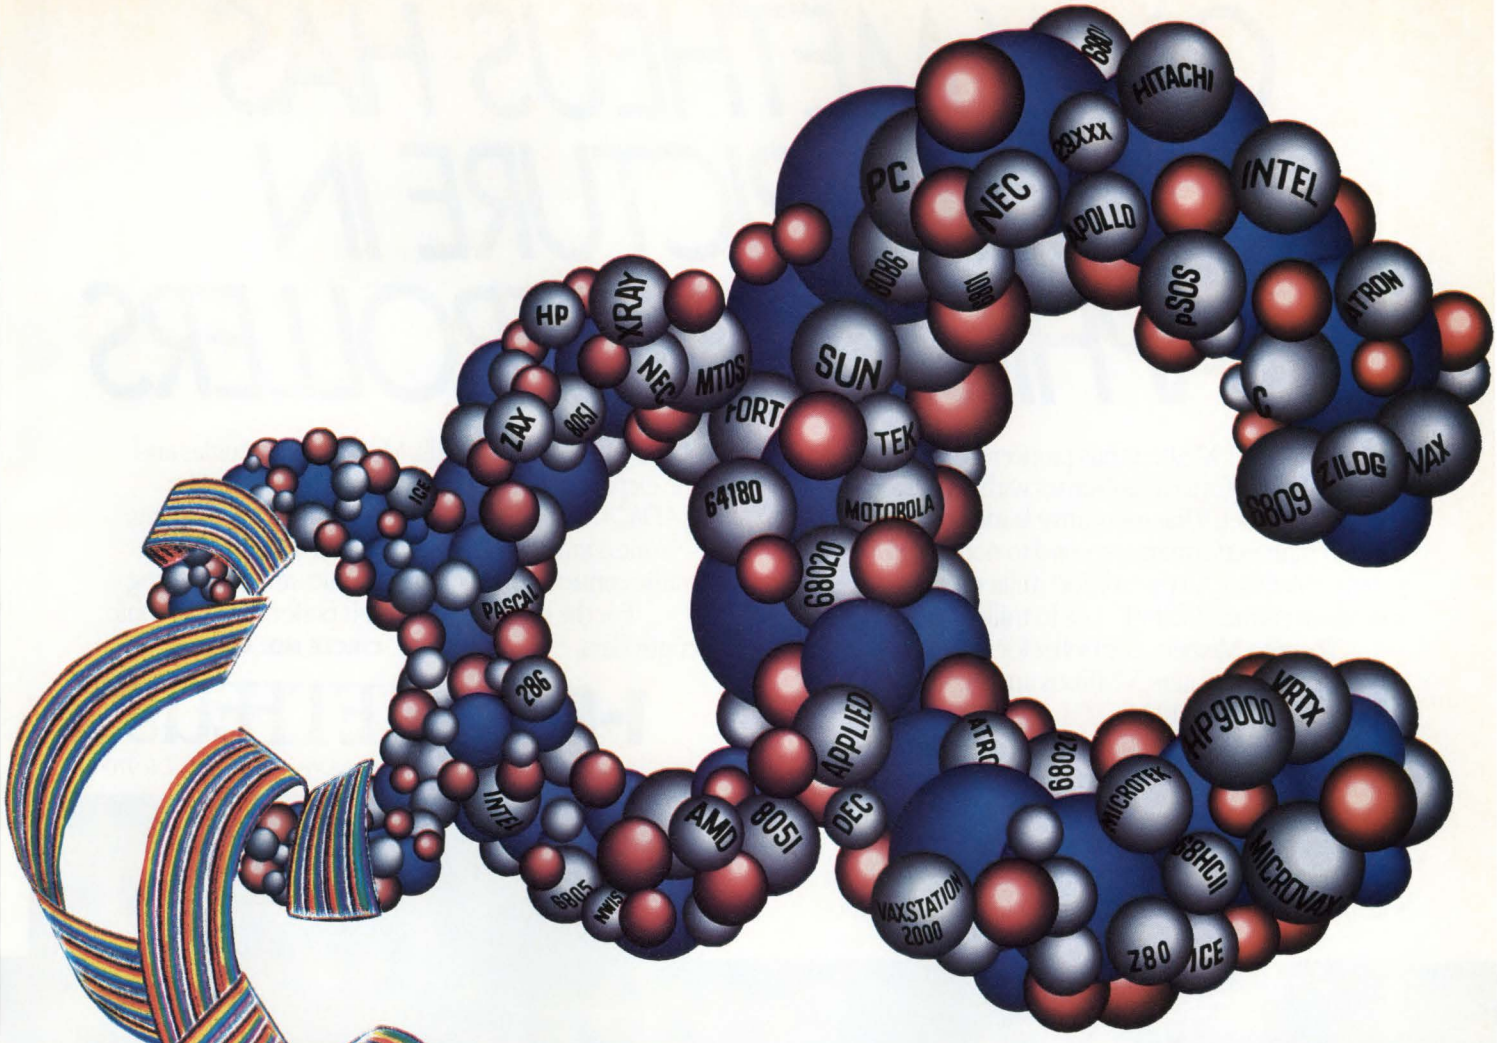

MICROTEC CROSS-COMPILERS DEBUGGERS ASSEMBLERS: THE DNA OF SUCCESS

Microtec Research has done a little genetic engineering on the cross-software development industry. Sure, our quality and customer support have become the industry cross-development standard. And what other company has the confidence to actually publish all fixes and ECOs for all products, for everyone to see? But our real legacy to you, our customer, is our commitment to providing *the same* cross-development software from more host processors to more target systems. From pSOS to VRTX to MS-DOS to UNIX to VMS. From assembler to C to Pascal to Fortran. From Motorola to Intel to Hitachi to AMD to Zilog. From IBM to DEC to HP to Apollo to Sun. From crossassemblers to cross compilers. From source debuggers running on hosts that are simulating target systems to the advanced in-circuit emulators from HP, TEK, Atron, Applied Microsystems, Sophia, ZAX and Taiwan's Microtek. All of which means you can spend your valuable time creating new silicon wizardry, rather than learning new cross-development user interfaces. Call Microtec at 800-950-5554 today (408-980-1300 in California). So you can do some genetic engineering on a competitor's physiology tomorrow.

MICROTEC® RESEARCH

2350 Mission College Blvd. • Santa Clara, CA 95054 • 800-950-5554 or 408-980-1300 in California

© 1987 by Microtec Research, Inc. Microtec® is a registered trademark of Microtec. The other companies mentioned own numerous registered trademarks.
Adv. by TRBA

CIRCLE NO. 23

ONLY METHEUS HAS THE BIG PICTURE IN GRAPHICS CONTROLLERS

Since 1981 Metheus has pioneered graphics controllers for PCs to mainframes with its proven ASIC graphics chip set. That innovative leadership has paid off with high performance second to none. Nothing compares to Metheus' fill rates of 160 million pixels per second and block moves (PIXBLT™) of 16 million per second.

There's a Metheus controller for almost any platform: system level computers, VMEbus and PC-Bus ATs. Resolutions include 1024 x 768, 1280 x 1024 and an amazing 2048 x 2048. Configurations are available from 4 to 32 bit planes. Powerful features include display list processing, double buffering, crossbar switch, anti-aliasing, overlays and DMA bus master control.

Metheus performance, quality and service have satisfied thousands of users worldwide. Users include "USA Today", McDonald's, McDonnell-Douglas and the Department of Defense. For applications like CAD/CAM/CAE, satellite imaging, medical imaging, seismic control and mapping, flight simulation, air traffic control, presentation graphics and publishing.

For the big picture on high performance graphic controllers, call Metheus. **CIRCLE NO. 66**

1-800-METHEUS™

OGC Science Park 1600 NW Compton Drive Beaverton, OR 97006

Metheus graphics controllers support PCs, VMEbus and system level platforms.

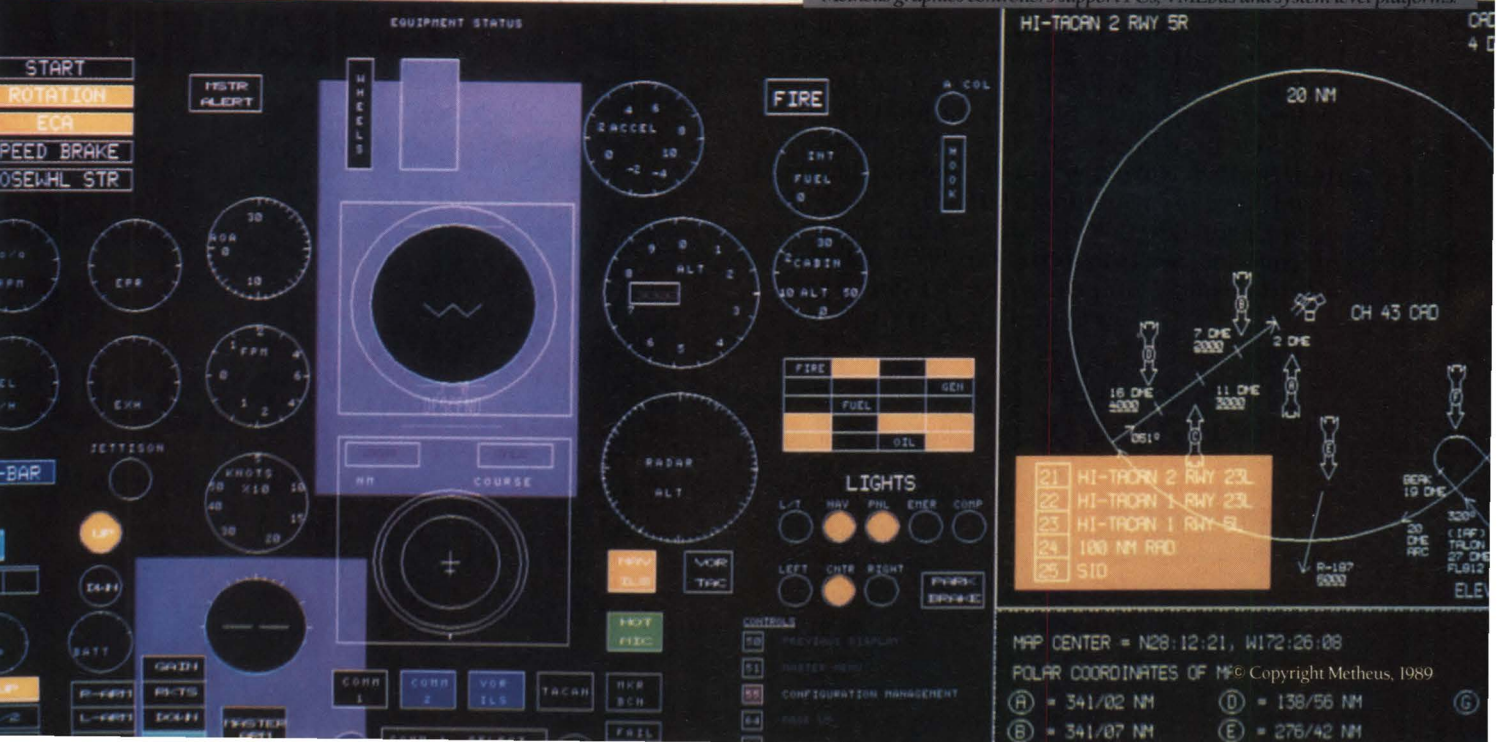

SOFTWARE

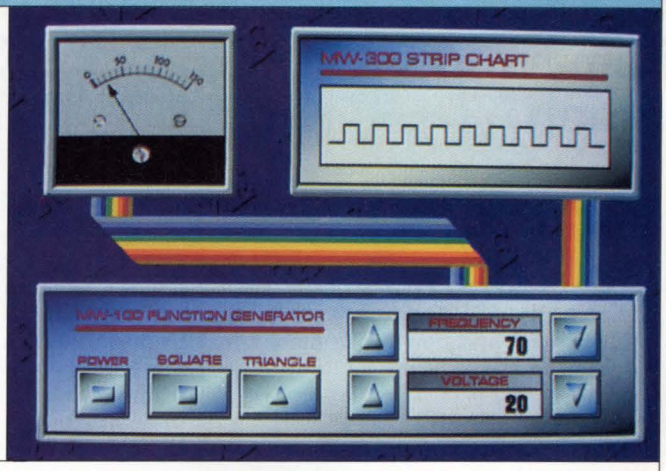

Screen images of an analog meter or a strip-chart recorder are typical of Rave's library. When used with a touch-screen monitor, controls such as those in the MW-100 function generator can be implemented.

the graphics library, or by combining library primitives.

Digitized images from a video camera can be imported directly into the program. Images—created from the library, or digitized and imported—can be modified and refined using Rave's paintbox package. Alternately, Rave includes a utility that lets drawings made and edited under other PC programs be imported.

The audio tools let designers capture, edit and play back audio information that can be input from any external source (microphone, cassette tape, and so forth). The finished audio can be combined with the video image to complete the combined interface. In addition, Rave provides a simple interface for linking programs or functions written in C if the designer wants to supplement Rave's capabilities with hand-written functions.

The graphics file manager provides the audio, video and input drivers needed to support the run-time user interface. The input drivers support a keyboard as well as any other pointing device such as a mouse, touch pad, touch screen or any other device that supplies x-y coordinate information.

In addition to supporting a monitor, the video driver provides the drawing and block-copy primitives used to build the more complex functions used in the graphics support library and the presentation editor. Drawing capabilities include lines, rectangles, polygons, circles, ellipses, rectangles and rounded corners as well as text.

The graphics support library builds on the graphics file manager to create more complex concepts for a specific application. Library elements include controls, indicators and menus. Controls are objects on the display that mimic the behavior of the output devices. Both controls and indicators may be implemented as computer-generated objects or by digitizing an actual device image.

Simulating the equipment

While much is being made of the "look and feel" of PC interfaces, Rave's object is not to have the look and feel of a computer, but rather to simulate the look and feel of the equipment being used. While the audio and video can be used to provide both verbal and visual user warnings or instructions, they can also combine to provide the visual and auditory feedback.

As an example, consider an operator turning a knob. "The visual display can show the knob rotating; the auditory response could indicate clicks as the knob passes detent points," says Steve Johnson, Microware product marketing manager.

As industrial and process control functions continue to automate factories, higher level real-time computer systems are slowly taking over from traditional programmable logic controllers. The availability of an easy-to-implement, interactive graphics and audio interface capability integrated into a real-time system provides the next step in this transition process. □

For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate number on the Reader Inquiry Card.

Microware Systems
(515) 224-1929 ..... **Circle 231**

ARE VME TIMING ERRORS LURKING IN YOUR DESIGN?

FIND OUT BEFORE OTHERS DO.

The VME Bus Anomaly Trigger (VBAT) is a massively parallel trigger board which automatically recognizes violations of the VME specification in real time.

Plug it into a spare slot in your system, and it will find design errors in all boards by watching every bus cycle, during actual operation.

Each timing violation lights an LED and generates a trigger output in less than 80 ns, which will trigger your logic analyzer, to give you an immediate picture of the bad bus activity.

Try one, and be confident.

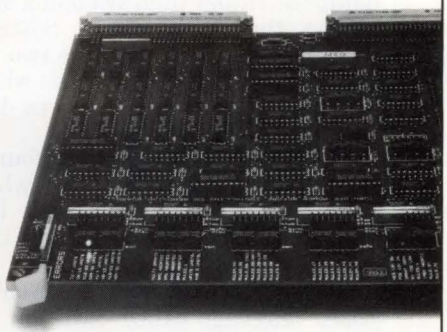

ULTRAVIEW

Ultraview Corporation
475 Yampa Way
Fremont, CA 94539
(415) 657-9501
FAX: (415) 657-0927

CIRCLE NO. 24

DESIGN AND DEVELOPMENT TOOLS

Pipes, hubs and frameworks vie for tool integration

Ernest Meyer, Contributing Editor

The fervent activity in developing object-oriented languages, structured query languages and fourth-generation languages in new markets is paralleled by accelerating movement toward tool integration in design automation. In particular, the traditional pipe model, in which data is translated directly from one tool format to another, diminishes in usefulness as tools and design systems continue to evolve and multiply. As a result, a number of CAD vendors are supplying more sophisticated design-tool management systems.

Some, such as CAD Logic Systems Inc (Rockville, MD) and Engineering DataXpress (San Jose, CA), are replacing the simple data pipe with a hub. In this approach, all data passes through some intermediary exchange format rather than being translated directly.

Other vendors, such as Silicon Compiler Systems (San Jose, CA) and Cadence Design Systems (San Jose, CA), are instead providing frameworks, which link design tools together by using one integrated database—often object-oriented—for all design data. In this approach, all tools access and process the same database, rather than swapping data among themselves. EDA Systems (Santa Clara, CA) combines two approaches in its Framework, which links tools together but moves data through pipes.

Vendors are also supplying fourth-generation languages, with which users can add their own tools into the system by writing scripts that link the tools together or into the database. Cadence supplies a Lisp-based language called Skill, for example. And EDA Systems supplies a macro-generated graphical utility for controlling data and process flow at an abstract level.

Meanwhile, many CAE vendors have joined forces to provide a standard for tool integration, extending all the way from database management to cursor shape definition. The group, known as the CAD Framework Initiative (CFI), won't have any usable answers for at least two years, however. So the issue then comes down to deciding what to do until then, and how much money to spend on the interim solution.

When a pipe isn't enough

It may seem strange to spend any money at all on tool integration when standard operating systems such as Unix, and even DOS, include all the necessary facilities for data format translation and automatic tool invocation. Bourne shell scripts are easy to write, after all, and Unix is free. Even so, rolling your own data translators and process-control batch files leads to three problems: redundancy, elitism and transience.

The work involved in writing a simple translator—a one-way data pipe, for example—becomes redundant when there are more than three tools to integrate. "If you're interfacing four tools to each other, each one needs three translators, for a total of 12," says Paul Menchini, director of research and development at CAD Logic Systems Inc (Durham, NC). "If you instead move all the data through an intermediary format, a central 'data hub,' each tool needs only two translators, one into and one out of the hub. That's a total of eight translators with four tools on a hub—four less than are necessary with pipes."

Obviously, a hub becomes more advantageous as the number of tools increases. Unified databases, as used in tightly integrated frameworks, offer the same advantage.

On the other hand, using a hub or framework does mean learning one more data format description. Also, frameworks and hubs are designed to accommodate all types of data, so their lexicon is much larger and more complex than for any one particular application.

Redundancy by itself is therefore not enough of a reason to do away with pipes. "The real problem with pipes might be that they're too easy to write," says Bob Carver, design framework product manager at Cadence. "We keep coming across companies where every engineer has written his own script interfaces, and no one knows what anyone else is really doing." In fact, this problem of elitism actually magnifies redundancy, as new engineers end up rewriting those same interfaces over and over again.

Perhaps the biggest advantage of hubs and frameworks is their durability. Not only do they reduce redundancy and elitism, but they also provide a more standard data description format, which is more durable over time, as a side effect of

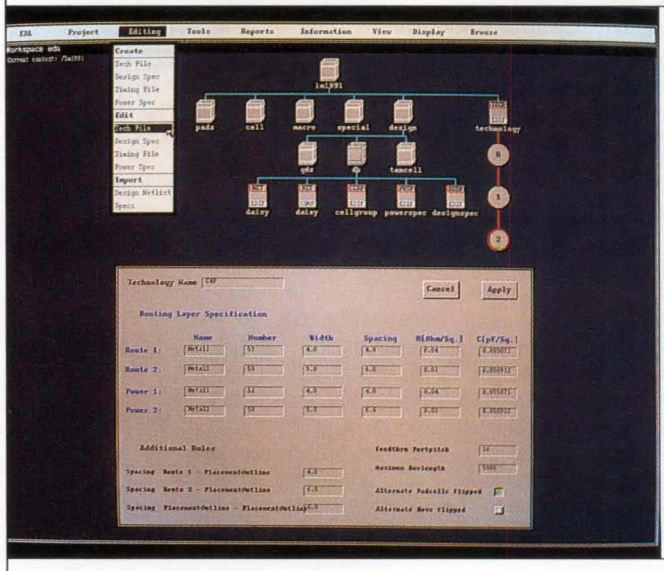

Siemens is using integration tools from EDA Systems as a design-movement system for a design-automation toolset. In the window are design files active on a particular project, together with a diagram that shows how the tools are linked by pipes.

The Fastest Way To Catch A Bus.

VME, VSB, Multibus or Micro Channel.

Have your new boards and systems up and running in record time. Off-the-shelf bus interface chips from PLX Technology get you to market faster and save a bundle in the process.

Pour on the power with bus control chips driving up to 64mA. Choose from a variety of master controllers, slave controllers, and other devices for immediate, space saving solutions. All of our chips incorporate on-board Schmitt trigger inputs and metastable hardened circuitry.

PLX packs everything into slim, 300-mil by 24-pin DIPs, or 28-pin LCC packages.

Suddenly, valuable board space is opened up, leaving you loads of extra room to implement that breakthrough design. Best of all, you'll finish way ahead of your competition.

Get on the fast track with bus interface chips from PLX Technology. Call us toll free 1-800-759-3753. Or write PLX Technology, 625 Clyde Ave., Mountain View, CA 94043.

PLX TECHNOLOGY

CIRCLE NO. 25

DESIGN AND DEVELOPMENT TOOLS

their implementation. Each tool by itself is too transient to ensure that its own data format will be usable in the future.

The transience of each tool's own database is accentuated by the continuing mergers and product consolidations in the design automation market. "Our customers are asking us to archive the versions of the tools we use, as well as the data," says Brent Wilkins, semicustom marketing manager at Plessey Semiconductors (Scotts Valley, CA). "That just indicates how cautious our customers are about maintaining reusability in their designs."

For this reason, both CAD Logic Systems Inc (CLSI) and Engineering DataXpress chose industry-standard descriptions as the foundation for their intermediate format: CLSI uses VHDL, and Engineering DataXpress uses EDIF. These standard description formats are established and therefore likely to survive the waves of time.

Tighter integration

Other vendors choose to integrate data sharing more tightly than is provided by a central hub intermediate format. "A hub is meant to be better because you can get from one place to another relatively easily," says Carver of Cadence. "The problem with a hub approach is you're always going from one place to another, and you can't be in both places at the same time. Customers really want dynamic interaction between the two tools."

Cadence and Silicon Compiler Systems have both adopted the "tightly integrated" database approach, where all tools and users share one central database. Tightly integrated databases aren't ideal for all situations, however.

"Translating data into another format for a loosely integrated tool is useful when the job that the tool is performing is large and mostly noninteractive," says Carver. "Layout verification is a particularly good example." Cadence thus combines a tightly integrated set of about 30 tools with a dozen or so loosely integrated tools. Data is passed to and from the loosely integrated tools and reformatted as necessary, when the performance trade-off is attractive.

Integration is a steady and continual process that shouldn't be held off until the CFI has drafted its standards, according to Cadence. "We start by putting a framework in place for some tools, and then we gradually expand it," says Carver. "When the CFI standards are complete, our framework will be made compatible with them."

"We ask for control of design entry, for schematics and full-custom layout," continues Carver. "If we control design entry, it's very easy to apportion data out to the tools that need it."

With the Cadence approach, design entry is very tightly integrated into the data framework and can thus be used as an all-purpose control tool. "Just plugging your favorite editor into a system is giving up too much of the power that a framework approach offers," says Carver.

"The real problem with pipes might be that they're too easy to write."

—Bob Carver, Cadence Design Systems

EDA Systems has a different approach. "Large design houses are beset with burgeoning archives," says Isadore Katz, director of business development at EDA. "When you put together a framework, you're assuming responsibility for managing, networking and archiving design data. Our approach is to make sure that our customers have to try really hard to lose a design file or crash the system halfway through a design process."

Siemens (Munich, West Germany) is using the EDA toolset to integrate its future CAD system, AFCIA. "We're looking at making a CAD system for full-custom IC design," says Norbert Kraft, project manager of design systems at Siemens. "For that purpose, we were looking at the market to get state-of-the-art CAD tools. The main thing to do after getting the tools is to give them a common user interface, and provide a common user access to a database." Siemens felt EDA Systems met those needs the best.

"Another very important part of the tool integration task is to provide features for project management," says Kraft. "That means keeping track of all design activity. And that's what the EDA Systems approach does best."

Making a choice

The results of each integration vendor's tactics are roughly similar, yet each offers a different emphasis. The "hub" vendors, CLSI and Engineering DataXpress, are really providing for the movement of data between unintegrated tools, and for a long-term storage format. Within that approach, CLSI's use of VHDL makes the company more centered on simulation, whereas Engineering DataXpress's use of EDIF makes that company more suited for IC layout data.

Cadence and Silicon Compiler Systems are perhaps the most similar in philosophy. Both companies provide tightly integrated tools with hooks for translating data to more loosely integrated tools. Cadence, however, also provides "cosmetic" utilities for standardizing the user interface, whereas Silicon Compiler Systems, with the more recent offering, instead provides remote procedure call and similar lower-level interfacing techniques.

EDA Systems is probably the most unusual of the bunch. Its so-called "Framework" approach really provides a high-level language for writing plain-vanilla pipes. EDA Systems is ahead of all the other vendors in its file- and data-management schemes, however, which even provide for some amount of project management. □

For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate number on the Reader Inquiry Card.

Cadence Design Systems
(408) 954-7494 ..... **Circle 261**
CAD Logic Systems Inc
(301) 424-9445 ..... **Circle 262**
EDA Systems
(408) 986-9585 ..... **Circle 263**
Engineering DataXpress
(408) 243-8786 ..... **Circle 264**
Plessey Semiconductors
(408) 438-2900, ext. 2040 ..... **Circle 265**
Silicon Compiler Systems
(408) 371-2900 ..... **Circle 266**

EVERYBODY'S TALKING ABOUT THE FUTURE OF EDA.

WE'RE CR

Talk is fine.
But the competition's breathing down your neck. Designs are getting more complex. And your market windows are shrinking.

It's time to stop talking about EDA's future—and do something about it.

That's why DAZIX was created.

Formed by the merger of Daisy and Cadnetix, DAZIX is the only brand new EDA company with over 10,000 satisfied users. And a record of innovation no one else can touch.

Right now, we're creating

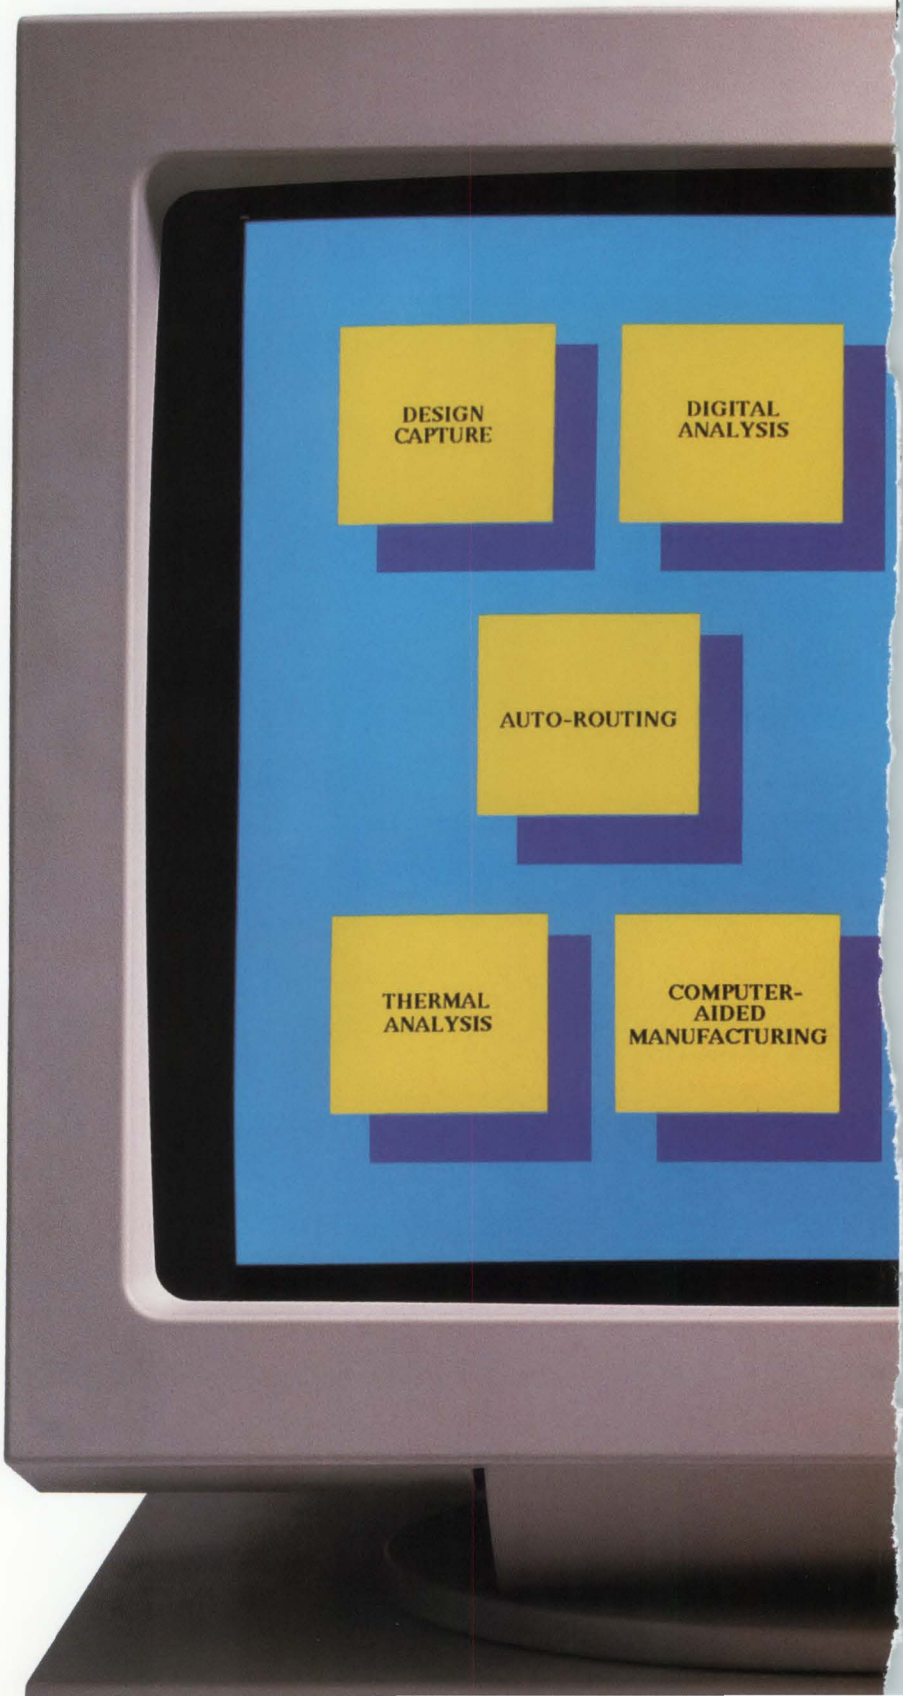

EATING IT.

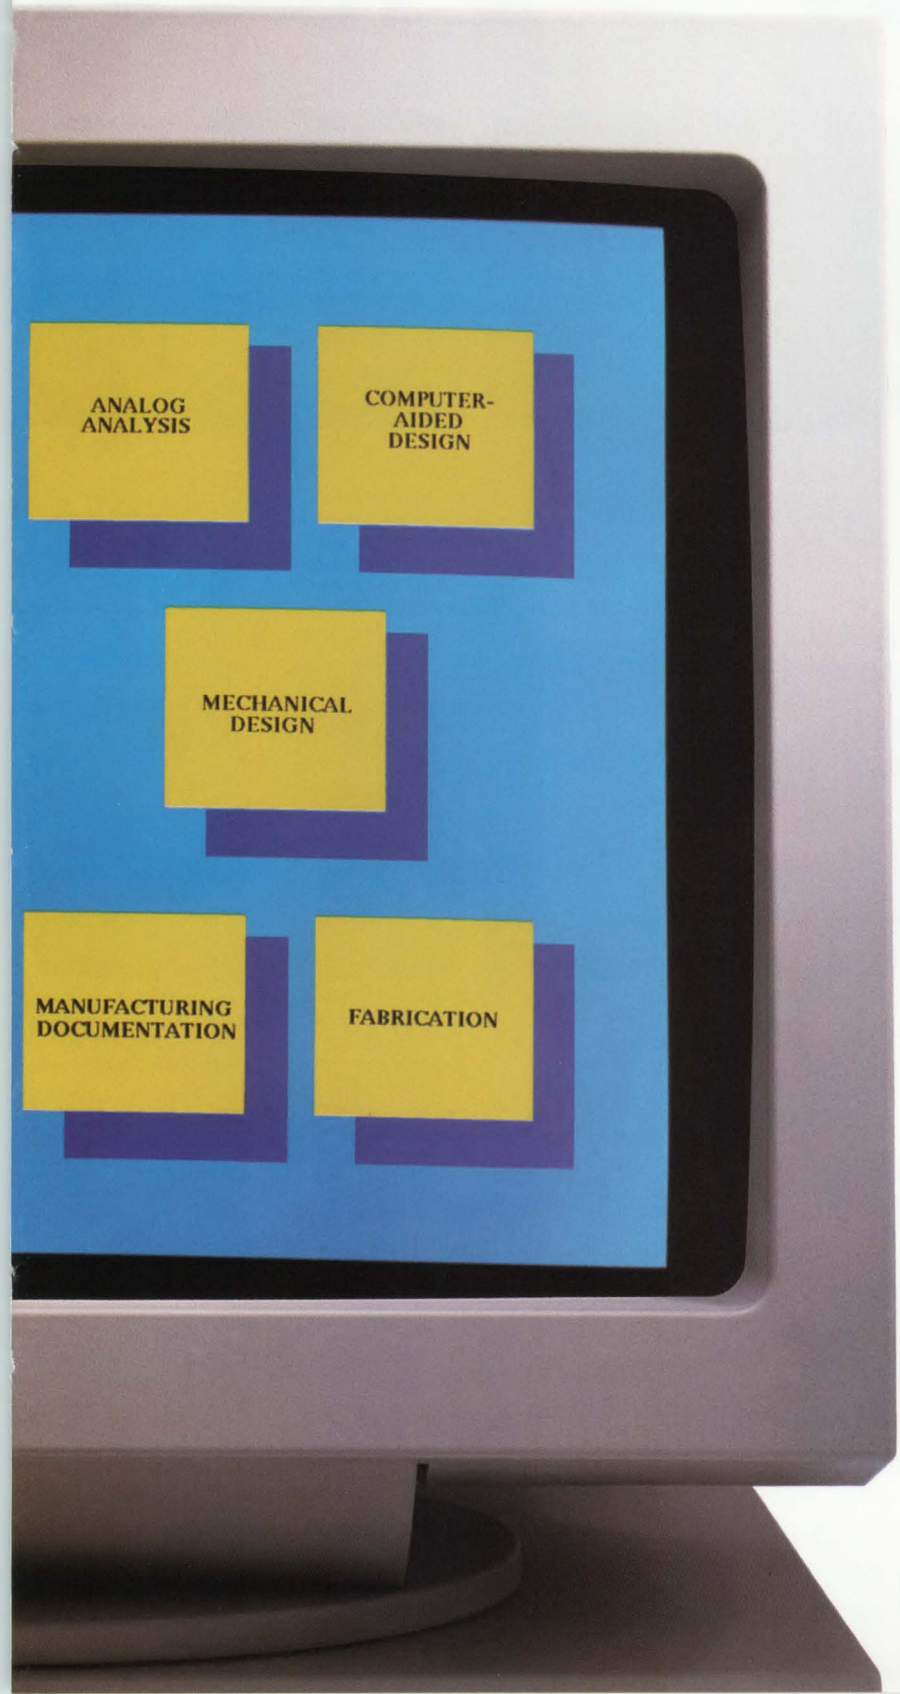

the next generation of EDA. And making it what it should have been all along. A complete set of the industry's richest, most powerful design tools. The best graphic interface. Open systems and industry standards. And the most comprehensive support in the industry.

For designers, DAZIX EDA will be a shot of adrenaline to their productivity. For product managers, it'll be the catalyst that drives design and manufacturing. And for your company, it'll provide the competitive edge to carve out a larger market share.

STARTING WITH A THE WORLD'S MOST

Design Capture

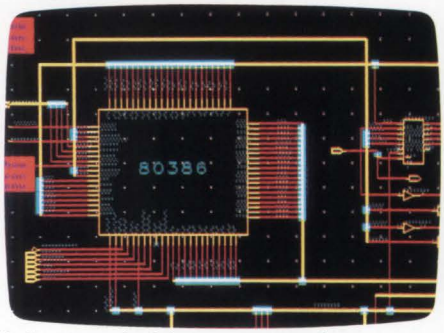

Industry standard design capture packages are available for Sun® products and the IBM® PC AT.® Integrated design analysis tools help assure a successful design. DAZIX ACE schematic editor is tightly coupled with simulation and PCB design tools to improve your design productivity.

Digital Analysis

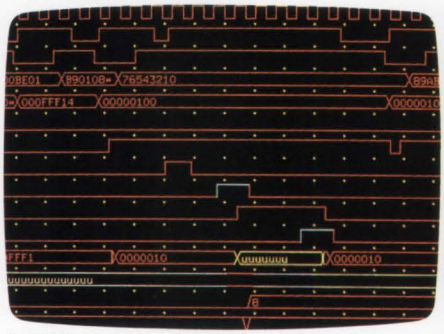

DAZIX mixed-level digital design environment supports over 170 ASIC design kits, and over 4500 popular standard components. Comprehensive VHDL design capabilities satisfy IEEE standards and DOD requirements. DAZIX logic, fault and worst-case timing simulations ensure your design's performance.

Analog Analysis

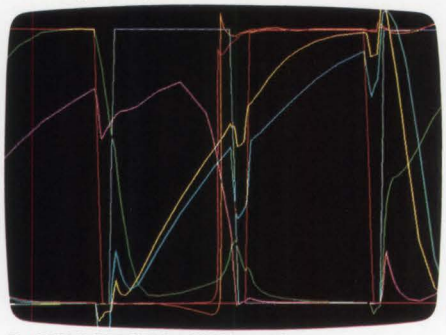

DAZIX DSPICE/PACSIM and Saber Analysis engine share the VLAB interface and draw from a common library. Over 1800 components are currently available in our Analog Library, including special purpose devices such as transformers and pulse width modulators.

Thermal Analysis

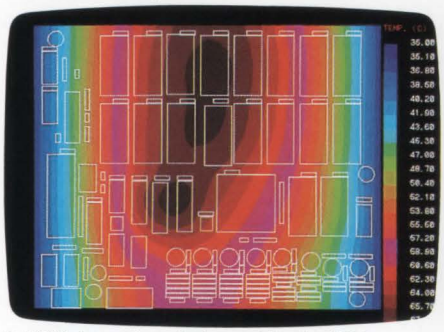

DAZIX thermal analysis provides the EDA industry's most in-depth, leading edge 2-D and 3-D thermal analysis solutions.

Computer-Aided Manufacturing

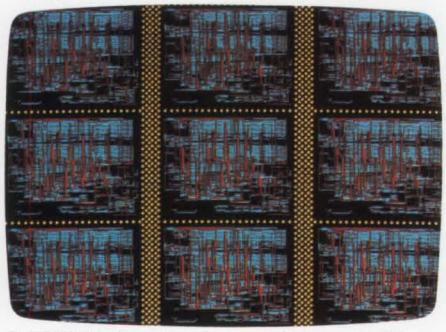

DAZIX CAM workstation contains a powerful set of features to panelize and prepare PCB designs for all manufacturing processes—photo tooling, fabrication, assembly and test to meet diverse data requirements.

We now have the richest set of EDA tools ever assembled in one place. Bar none.

Instead of the usual mind-numbing reiterations, you create, simulate, debug and manufacture uncompromised by design bottlenecks. So your designs move from concept to customer with extraordinary speed and control.

To get your concept down quickly and accurately, we have the industry's premier Design Capture tools. Powerful enough to handle the most complex ideas.

We provide Digital and Analog Analysis tools with remote compilation and analysis capability. Including our innovative Saber Simulator for board and system-level design, DSPICE/PACSIM for ASIC and discrete designs, our proprietary VLAB interface and our well-stocked Analog Component Library.

To test your ideas, we have the industry's most popular physical modeler. It lets you plug in different parts to create "what

COMPLETE SET OF POWERFUL TOOLS.

Computer-Aided Design

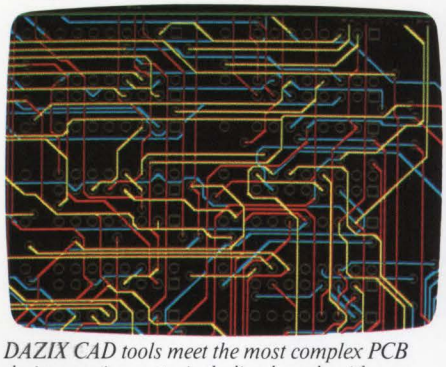

DAZIX CAD tools meet the most complex PCB design requirements, including boards with over 1000 EIC. DAZIX CAD workstations accommodate a wide variety of boards, from two-sided surface mount designs, to those that combine surface mount, pin grid arrays and standard DIP packaging.

Auto-Routing

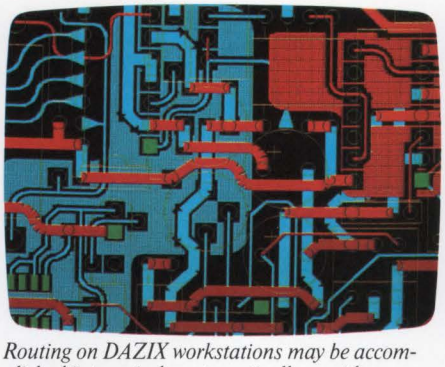

Routing on DAZIX workstations may be accomplished interactively, automatically or with a combination of the two. Routing may also be performed over the network on the DAZIX Route Engine 3,™ which performs true multilayer routing to 100% completion.

Mechanical Design

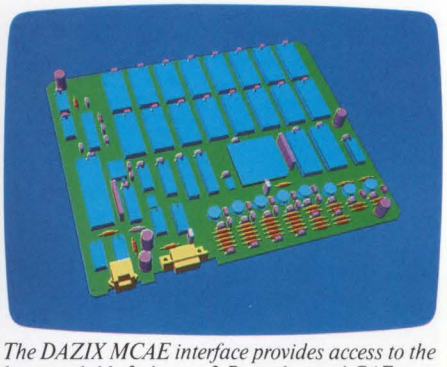

The DAZIX MCAE interface provides access to the best available 3rd party 3-D mechanical CAE software. This allows electrical packaging, thermal and vibration analysis to be performed at both board and system level.

Manufacturing Documentation

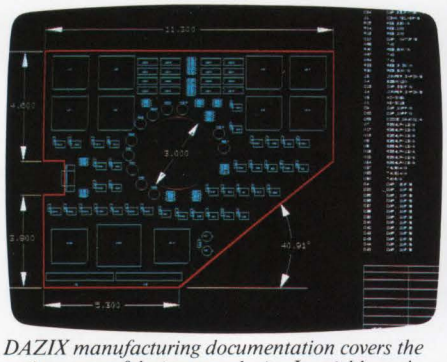

DAZIX manufacturing documentation covers the entire range of design complexity. It quickly and easily handles process flow through complex feature sets for DOD documentation requirements.

Fabrication

DAZIX offers the industry's most flexible fabrication workstation tools. Tools that quickly and easily handle your most complex panels and difficult PCB designs.

if" scenarios in seconds. So you can modify and perfect your design on-the-fly.

And for complex sequential ASIC designs, there's Automatic Test Program Generation, for the most complete fault coverage. When it's time to put it all together, you hit the boards running with the industry's best 100% Automatic Routing tools. Featuring accelerated graphics, comprehensive SMT and ECL tools and "design for manufacture." Your assurance that your boards can be easily built at the lowest possible cost.

To physically package everything into a system with maximum efficiency, we interface with leading-edge Mechanical Design tools. And to make sure your designs don't lose their cool, we have sophisticated Thermal Analysis.

And when you're finished, DAZIX EDA translates it all over to manufacturing, without losing a bit. So what you get off the assembly line is exactly what you had in mind.

WITH A COMMON UNCOMMO

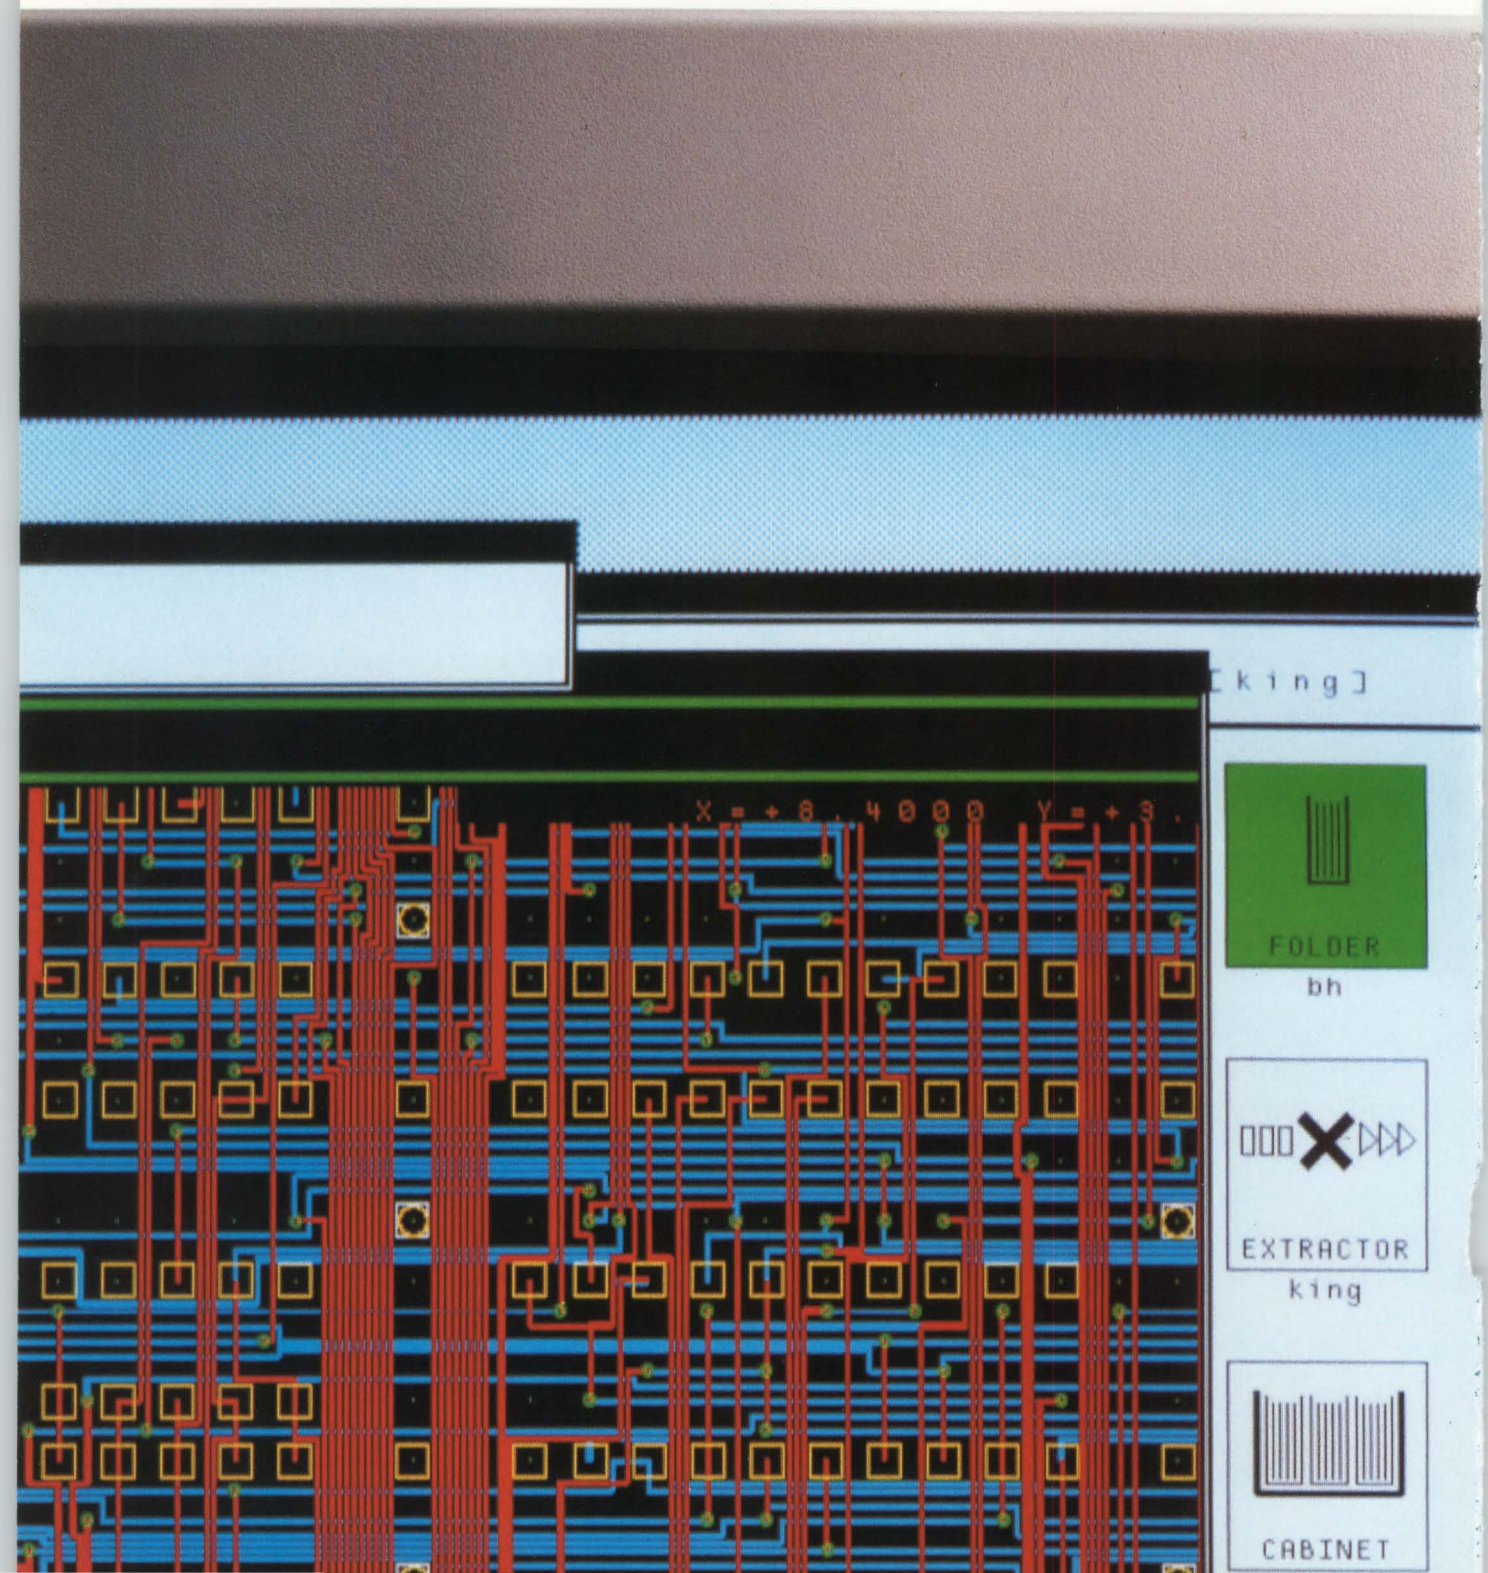

INTERFACE THAT'S NLY SIMPLE.

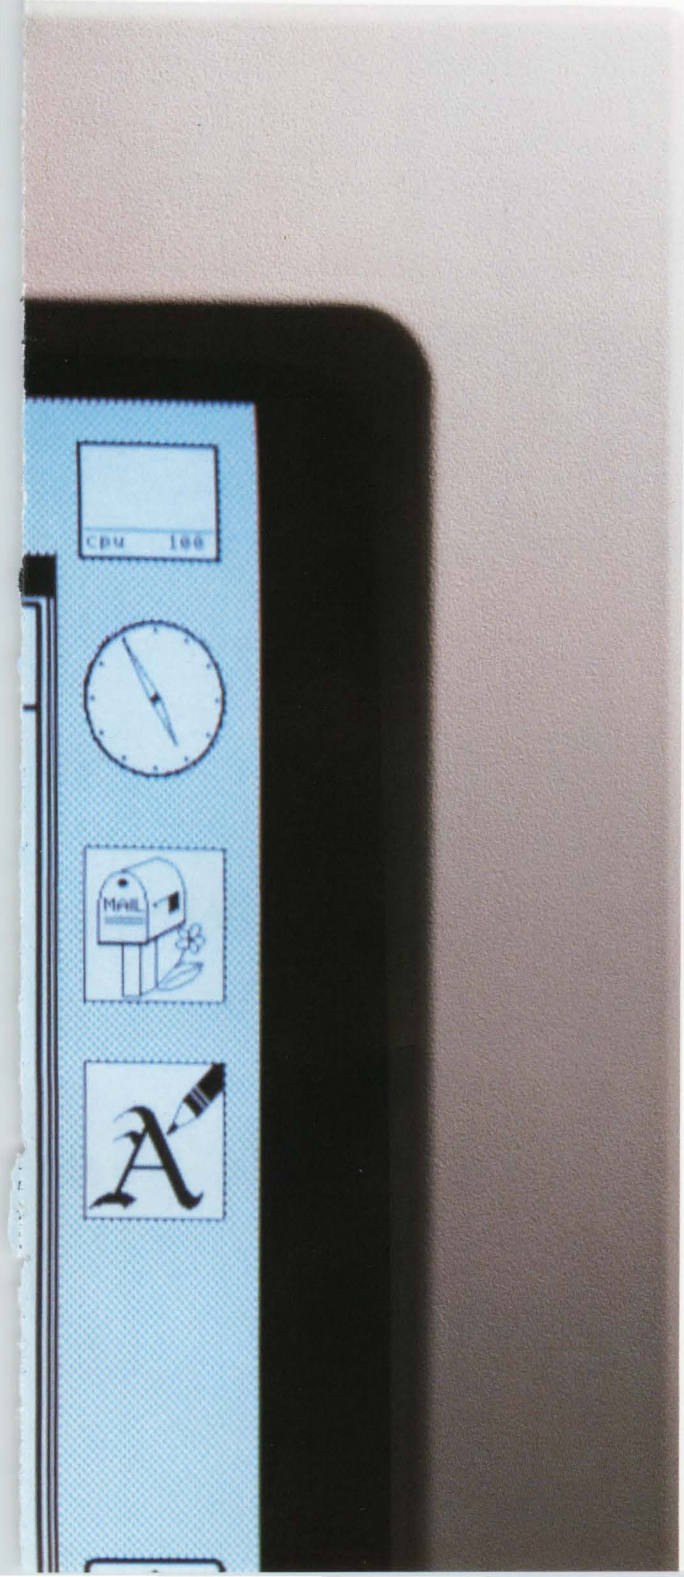

You'll see a visible difference in DAZIX EDA. It's the industry's best user interface.

It uses graphic icons, function keys and a mouse. Just point and click.

With just one uniform interface to learn, there's minimal training. So your design teams are up and running from nearly the first day they sit down.

And our interface is not only uncommonly simple, it opens windows to the most powerful and complete design environment in the industry.

An environment that quickly and accurately takes your design data from engineering to design to manufacturing.

Best of all, it's all yours today. And it's only the beginning.

USING AN O
BUILT TO THE HIG

Our design tools play all your favorites.
Including UNIX, DOS, Ethernet, IEEE 802.3, TCP/IP, NFS and the X-Window system.

With RISC/SPARC architecture you get unparalleled system performance and easy migration to the next generation of processing power.

To guarantee an open database, with protection of your design data, DAZIX EDA works with EDIF, IGES and IPC data exchange standards. So nothing gets lost in porting your design from tool to tool.

We're also open to today's most open platforms. Soon to include the entire Sun family. And that leaves you free to integrate a variety of supporting hardware and third-party software. And have it all work together. Economically. Because we let you cost effectively match the equipment to the job.

All this openness also lets you easily

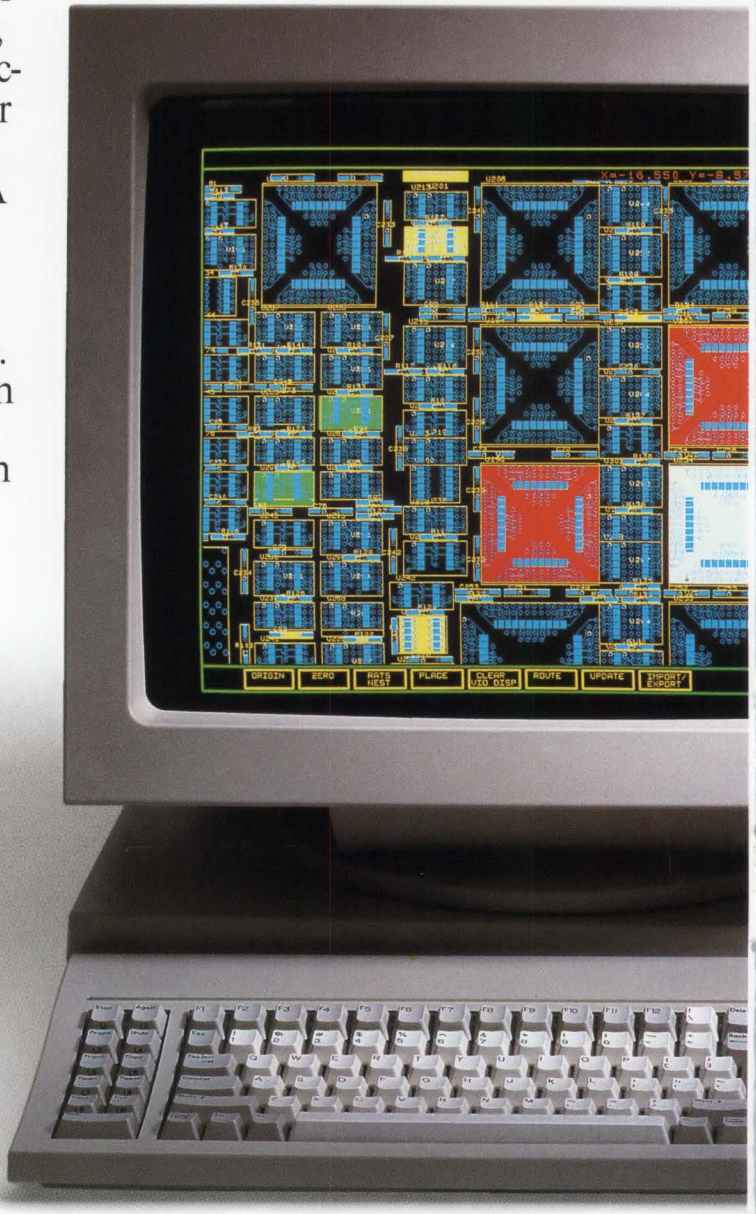

PEN SYSTEM HEST STANDARDS.

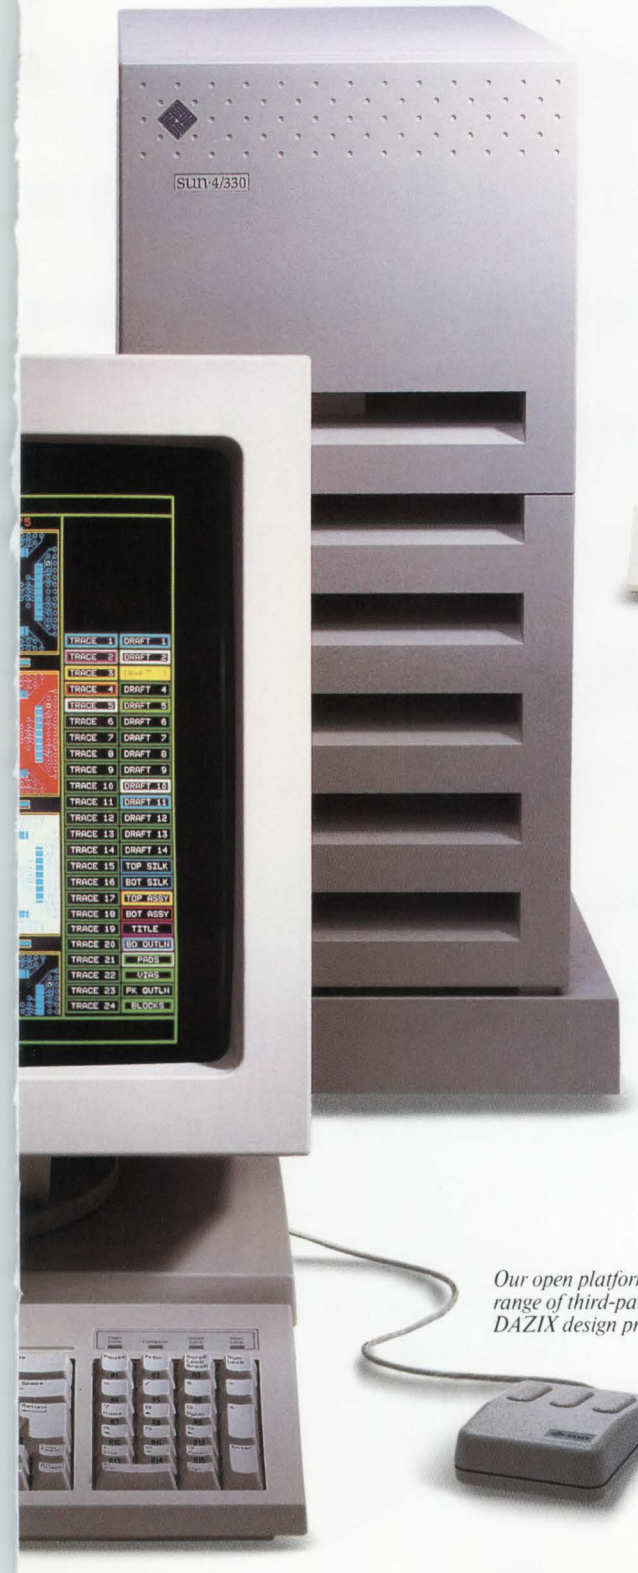

partition projects among your designers. And still have everybody work together, sharing the same libraries and hardware.

With DAZIX, industry standards come standard. So you get the most flexible, powerful solution to all your EDA problems.

So design teams move along in lock step. With fewer errors and iterations. Giving you the best possible chance to finish on time and within budget.

Our open platform approach lets you run a wide range of third-party programs to complement the DAZIX design process.

BACKED BY THE G
SYSTEM KNOWLEDG

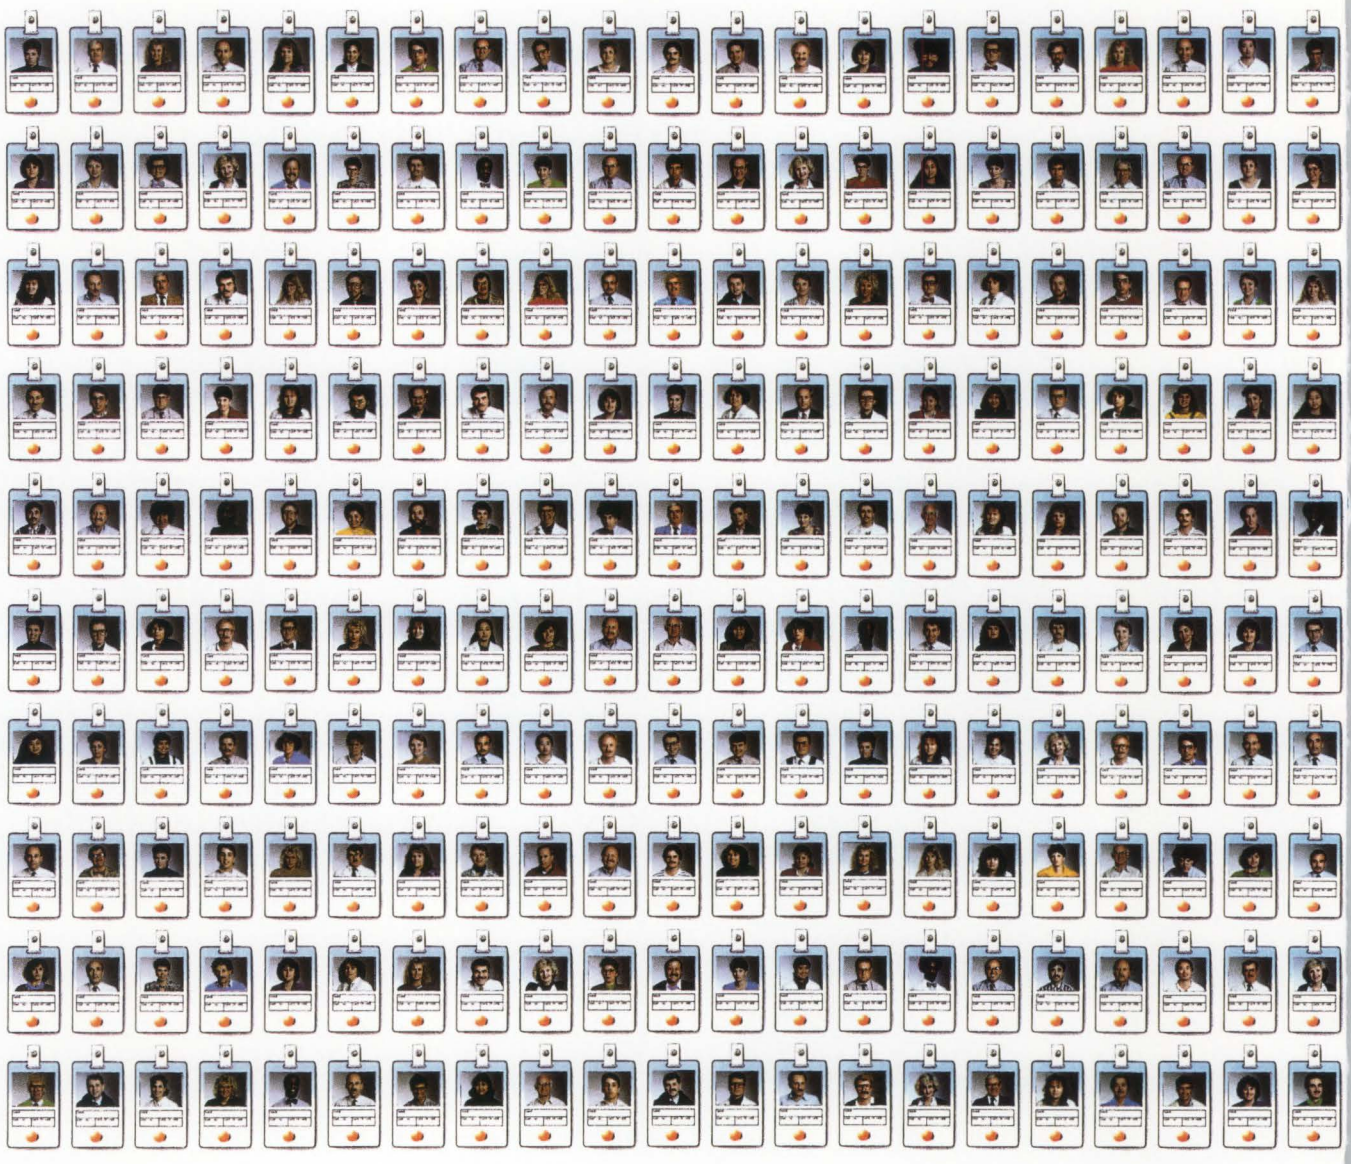

Some of our best design tools are the experience and talent of the industry's largest R&D staff. All 435 of them. All focused on EDA. And our sales and support staff in 44 offices in 17 countries worldwide.

All dedicated to helping you integrate even more capabilities with DAZIX EDA.

REATEST BODY OF
E EVER ASSEMBLED.

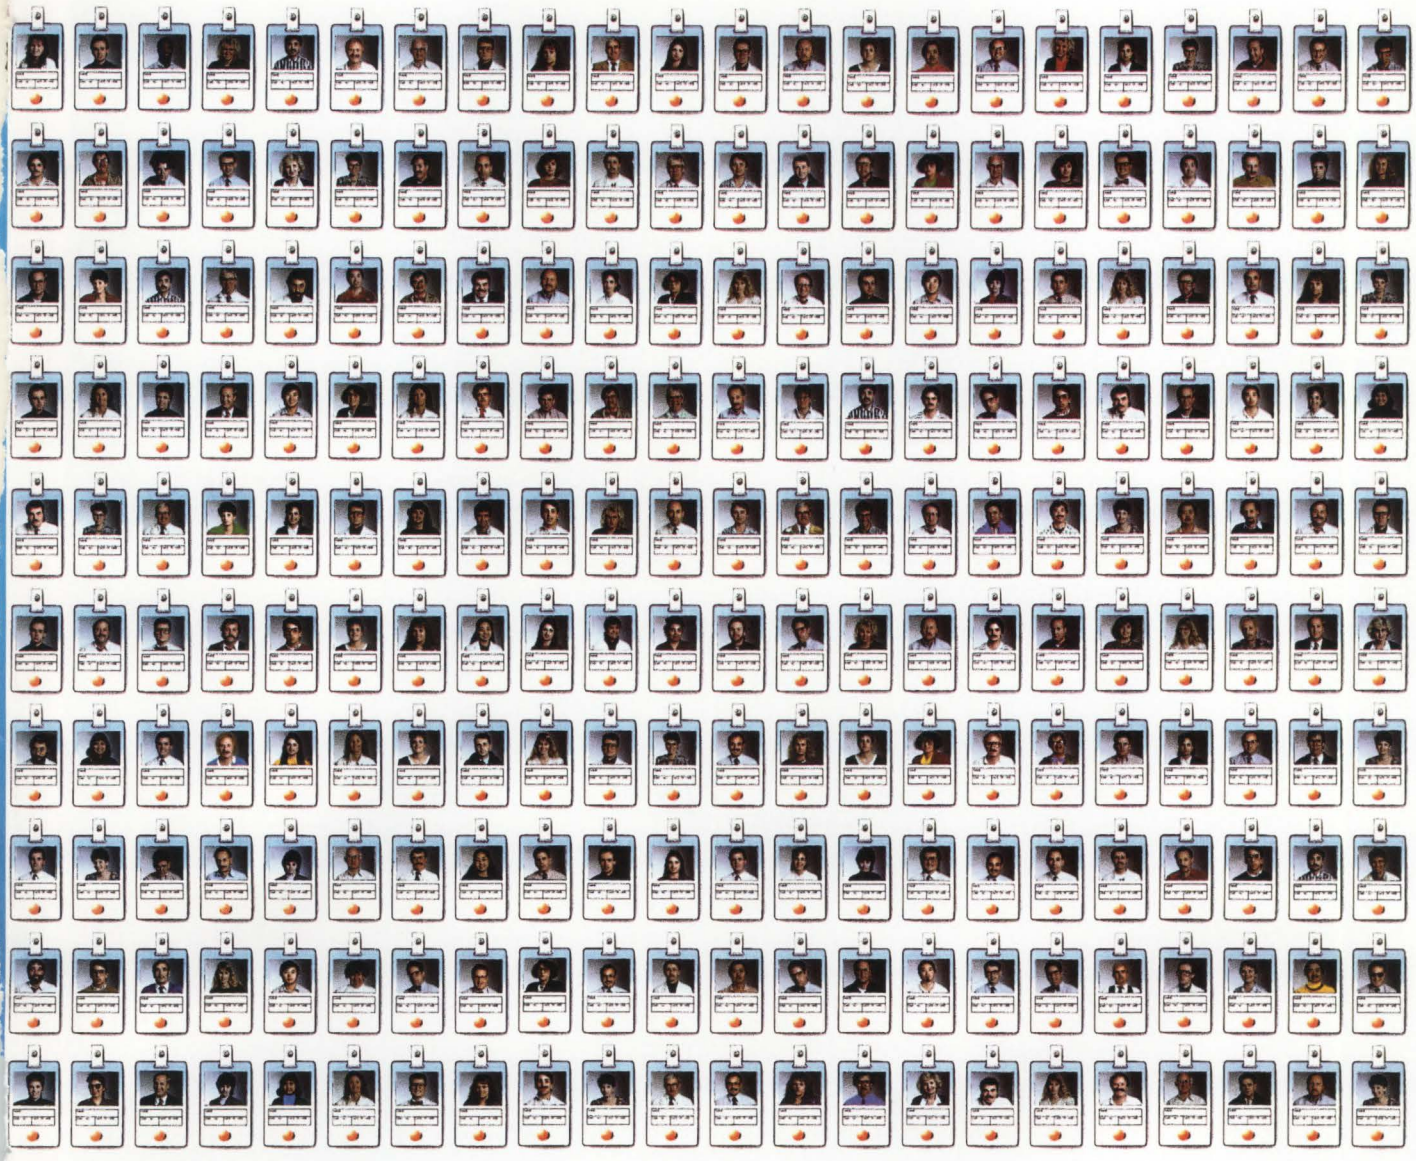

In fact, we're working to bring more advanced technologies on line faster than anyone else in the business. And we're backing that effort with the largest R&D budget in EDA. Proof we're putting our money where our mouth is.

IT'S TIME TO CREATE YOUR FUTURE. TODAY.

Ready or not, the future's here. And it's a moving target. So you're going to need the best tools to deal with it.

We have them. Now.

No matter what tools you're using, we can make the transition smooth and painless. Not that you'll need much help. Because DAZIX EDA is flexible and open, you can implement it as you see fit. In increments that match your needs and budget.

With one of the largest installed bases of EDA in the world, we have a global view of your design needs.

Knowledge we reinvest in increasing the capabilities of our technology. So we'll always offer you the best environment for growing ideas.

See for yourself. Write or call DAZIX, 200 Parkside Drive, San Fernando, CA 91340, 1-800-556-1234 ext. 32; in California, 1-800-441-2345 ext. 32.

It's time to turn talk into action.

DAZIX™

DAISY/CADNETIX INC.

© 1989, Daisy/Cadnetix Inc.
Sun, IBM, AT&T, UNIX, X-Windows, Ethernet, EDIF, IGES, IPC, TCP/IP and NFS are registered trademarks of their respective manufacturers.

Coming soon: flexible film in large, flat displays

Tom Williams, Senior Editor

The push for high-definition television (HDTV) is opening possibilities for new display technologies that will spill over into computer graphics as well. Since HDTV will be digital and therefore computer-based, there's a great deal of overlap in the display requirements between computer and video applications.

In particular, there are two emerging technologies that are remarkable not only because they're both flat-panel technologies, but also because both are implemented on flexible polymer film. Both could thus be used in very large displays, as well as in normal-sized television and computer monitors.

The first technology, under development by Staplevision (East Brunswick, NJ), uses a series of electroluminescent polymer materials to produce colored dots for an RGB color display. The second technology, called Polyvision, is being developed for commercialization by Alpine Polyvision, a division of The Alpine Group (Hackensack, NJ). Polyvision uses an electrochemical process to control the transmission or reflection of light through a proprietary film material that can be applied to a glass or plastic substrate.

Staplevision's pursuit of its display technology is a direct result of a mandate from the Department of Defense's Defense Advanced Research Projects Administration (DARPA), which is the only major U.S. backer of HDTV. But John Stapleton, president of Staplevision, says his first instinct was to pursue a color computer graphics technology before tackling the full-color video problem, which uses analog color levels rather than distinct digital signals. Pressure from Congress, however, made Staplevision go for the consumer product first.

TFTs control color triads

The Staplevision technology uses an active matrix of thin-film transistors (TFTs) to control triads of red, green and blue color dots deposited on a flexible polymer film. The dots are made up of proprietary electroluminescent material. There are three types of material for the three colors.

On the front of the film—between the viewer and the colored dots—is a transparent tin-indium-oxide electrode, which forms one electrode of a storage capacitor. This tin-indium-oxide material is separated with gaps so that it runs in vertical strips down the screen.

Each strip connects a vertical column of dots of the same color. Dots are arranged across the screen in an alternating sequence of red, green and blue. Behind the dots on the film are the other electrodes of the storage capacitors, the TFTs and their matrix of connections. Driver ICs on the edge of the film connect to the matrix conductors.

The bugbear for manufacturing large active-matrix displays has always been the failure rate in the TFTs and the storage capacitors. Just a few bad pixels can ruin a display panel. Staplevision has approached the problem with a fault-tolerance scheme that includes, but isn't limited to, straight redundancy. Each three-dot triad constitutes a pixel, and each pixel has eight capacitors and eight TFTs associated with it—that's the redundant part.

Fault tolerance comes in with the way the transistors are arranged. Two pairs of the transistors are connected in parallel and two pairs in series. If a transistor in one of the series pairs shorts, its mate can thus still handle the transistor's function.

If one of the transistors in a parallel pair is open, the other can still take over. While an open transistor in a series pair will knock out that pair, there's still a series pair left. The opposite also applies if a transistor in a parallel pair shorts. The probability that enough faults will occur at one pixel to disable it is quite remote.

In a computer graphics applica-

STAPLEVISION TECHNOLOGY

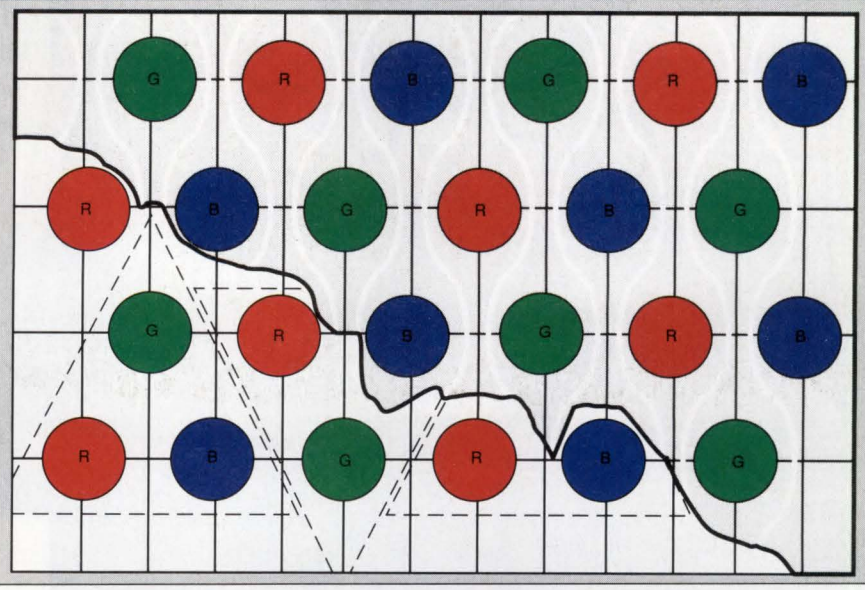

The Staplevision display technology uses an active matrix of thin-film transistors to drive triads of red, green and blue dots, which are electroluminescent material deposited on a polymer film. Each triad of dots represents a pixel. Above the matrix layer of criss-crossed conductors (the transistors aren't shown) is a layer of color dots. The layer above that is a transparent electrode layer of tin-indium-oxide. The wavy lines represent gaps in the electrode material to electrically connect dots of the same color in vertical columns. In a graphics application, the brightness of each dot is determined by the voltage applied to the transparent electrode when that triad's transistors are turned on.

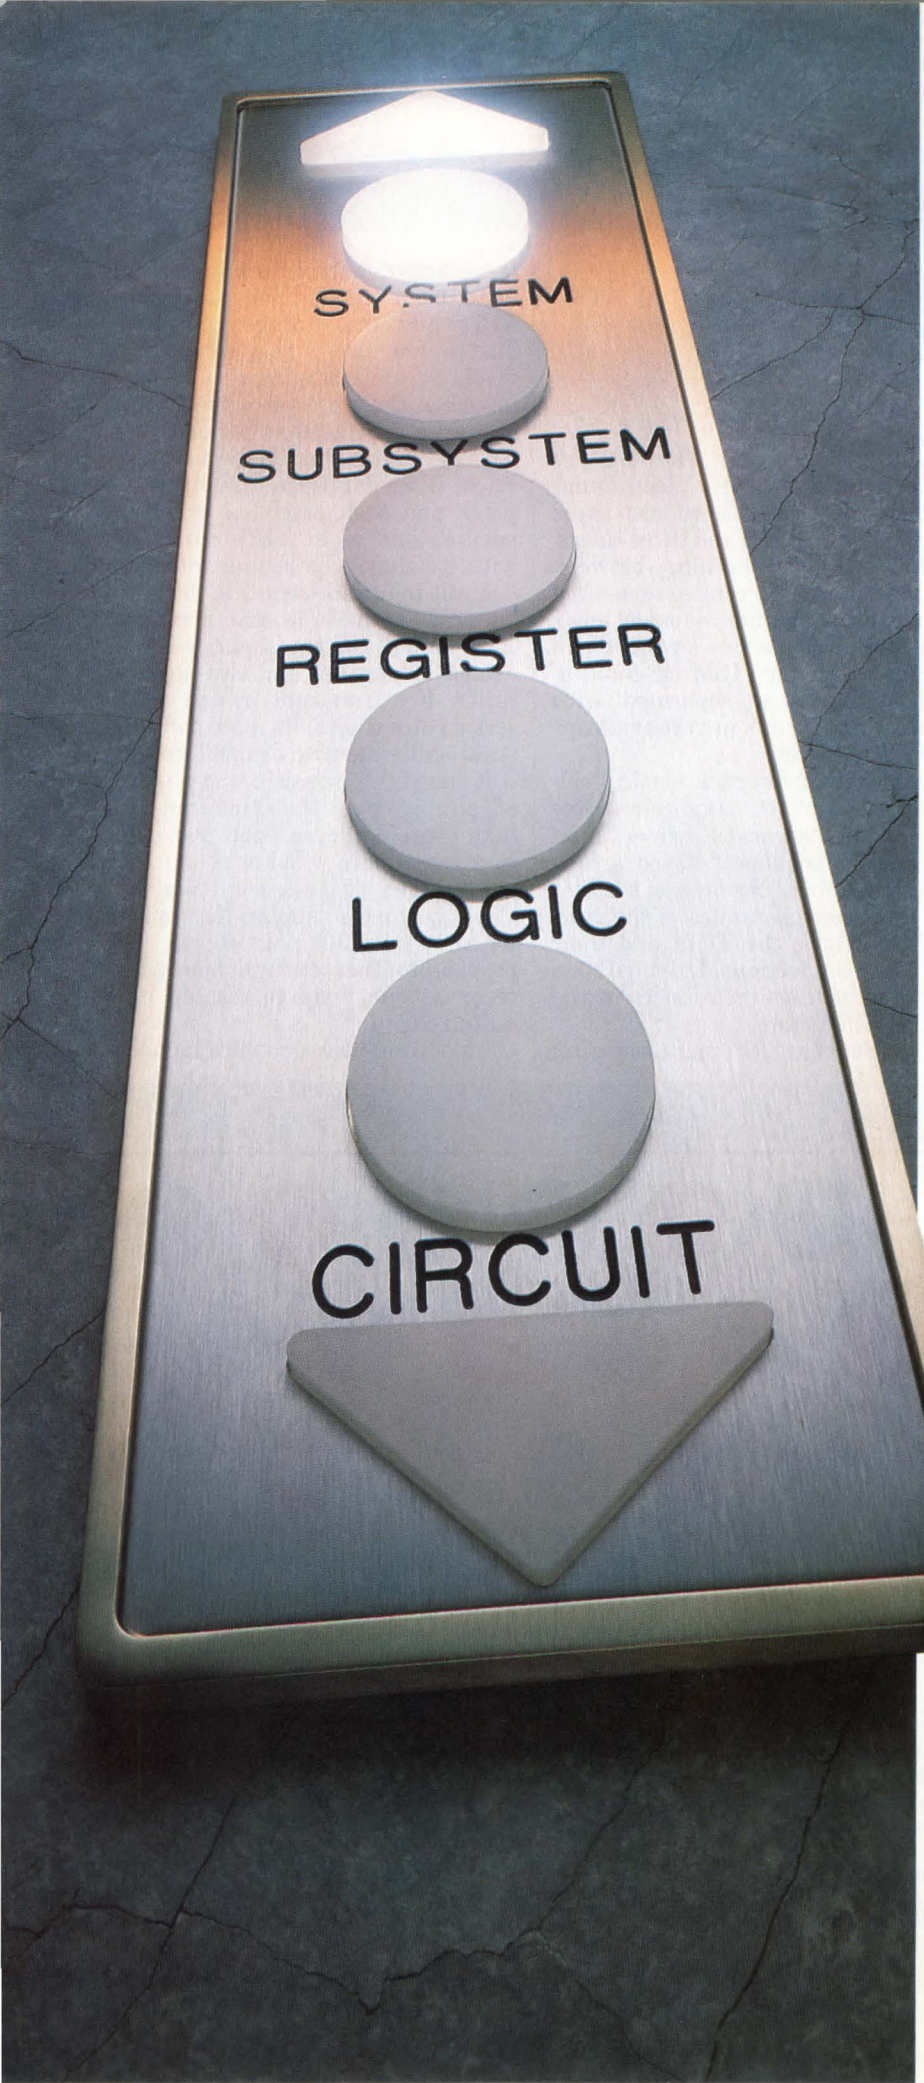

Elevate Your Thinking.

The power and flexibility of VHDL integrated into a complete EDA environment. We call it System-1076.™ It's a comprehensive architectural design and analysis tool suite. And there's nothing else like it.

System-1076 eliminates the barriers between architecture and logic design. In a single automated environment, designers conceive and analyze a system from the highest possible level to its detailed implementation.

System-1076 features the industry-standard VHDL for high-level design description. And since hardware designers know that a picture often says more than a thousand words, we developed an interactive graphical interface.

System-1076 is more than a clever HDL, however. Much more. VHDL models execute directly in Mentor Graphics' QuickSim™ logic simulator, and QuickGrade™ and QuickFault™ fault simulators. No translations required. There's also a powerful source-level debugger to quickly spot architectural design flaws.

System-1076. System design from top to bottom.

For more information about System-1076, call toll-free **1-800-547-7390.**

Mentor Graphics®

Your ideas. Our experience.

MAJOR SYSTEM COMPONENTS

tion, all the dots in a triad share the eight transistors. As in a shadow-mask CRT, different colors result from varying the relative brightness of the dots in a triad. In the Staplevision system, one can vary the brightness by successively turning on the transistors of each row of triads while applying, in synch, a voltage level to the vertical transparent electrode for each color. The electrode voltage determines the brightness of each dot in a pixel.

Full-color video is harder to do, says Stapleton, because in addition to controlling the transparent electrode voltage, one must also modulate the current in the TFTs. Here the triad dots can't share all transistors, and fault tolerance is harder to maintain. The video comes in as a source signal to the TFTs, which are field-effect transistors (FETs).

"The video comes through an analog shift register to drive the source of each FET in succession across a line," Stapleton says. "Since the gain of transistors isn't exactly uniform, we have to correct for it through feedback." In addition, the full video implementation must correct for between-dot crosstalk that isn't critical in computer graphics applications.

Still, Stapleton feels confident that commercialization of the technology is imminent. He predicts a full-size prototype within 18 months and production within 30 months.

Staplevision's production method is unique, to say the least, for a display technology. Displays will be manufactured by vapor-depositing layers of material (transparent electrodes, TFTs, conductors, electroluminescent polymer dots, and various insulation layers) onto a continuous roll of film 62 in. wide.

"That one dimension limits the size of the screen," says Stapleton. "Theoretically, it can be as long as you wish. The other dimension is limited by the technology of the deposition equipment."

Very high contrast

The Staplevision technology emulates a color CRT by illuminating triads of colored dots. In contrast, the Polyvision technology uses the optical response of a proprietary material that changes opacity when a voltage is applied. The system is made up of various formulations of this electrochemical material, which is applied in a layer to a glass or plastic substrate, and of a transparent conductive layer. The transparent layer, which acts as one electrode, is tin-indium-oxide. The Polyvision material reacts very rapidly to a low applied voltage and can produce a very high contrast.

Polyvision has a sharp response curve over a very low voltage range, about 1.1 V. Current, however, drives the display, and current demand depends on the pixel size.

A Polyvision display can be either transmissive or reflective. Depending on the type of material used in the electrochemical film, the off state can be either transparent or opaque white. The fully on state is saturated black. Gray scales can be achieved, however, by varying the level or the duty cycle of the applied voltage.

Polyvision is a very broad technology that has opened up a new area in electrochemistry, according to Bragi Schut, executive vice-president at The Alpine Group. "There are many different Polyvision formulas that respond to different requirements, such as color or opacity," he says. In addition to computer displays and HDTV, The Alpine Group is looking into variable opacity architectural glass and large displays such as sports scoreboards.

Better than LCDs

The variable transmissive characteristic of Polyvision technology shows some superficial similarity to liquid crystal displays (LCDs), but Polyvision has some distinct advantages. For one thing, it's a solid-state material, so there isn't the complex sealing problem associated with liquid crystal material. And because it doesn't need polarizers, Polyvision doesn't have the restricted viewing angles that plague LCDs.

Polyvision also appears to have a much faster response time than non-active-matrix LCDs, which run around 100 to 150 ms. "We believe video response rates are well within the technology's capability, and we have demonstrated video response rates," says Schut. The response rate is measured as the time between turning the material on and then off again by applying a reverse voltage, thus changing the display.

Free Intel Language

Buy an ICE-196KB/HX Emulator now and receive your choice of one of the following Intel 80C196 languages.
ASM • PL/M • C
A $750 value! Absolutely Free!
To order, or for more information, call:
800-874-6835
Offer expires June 30, 1989.

© 1989 Intel Corp
ICE™ is a trademark of Intel Corp.

POWER TOOL.

ICE™-196KB/HX. Intel's all-new, high-power emulator, offers greatly expanded capabilities and new debugging power. Like advanced break/trace. Expanded mappable memory. Precise 12 MHz real-time debugging. And powerful Intel symbolics.

ICE-196KB/HX. It's the new Power Tool for 80C196 designers.

intel

Power Tools for Power Users.

CIRCLE NO. 27

OPEN SYSTEMS::OPEN TOOLS™

MAXIMUM POWER FOR REAL-TIME DEVELOPMENT

Whether over the backplane or network, Heurikon has your VME and Multibus targets. VxWorks • VADSWorks • OS-9 • VRTX • Ada • Unix V.3 • Network solutions • Just a few of the Open Tools™ Heurikon has to speed your application to market. Call now for real-time response.

1·800·356·9602

HEURIKON CORP

MICROCOMPUTERS FOR INDUSTRY

VxWorks is a trademark of Wind River Systems, Inc. VADSWorks is a trademark of Wind River Systems, Inc. and Verdix Corporation. VRTX is a trademark of Ready Systems, Inc. OS-9 is a trademark of Microware Systems, Inc. Unix is a trademark of AT&T Bell Laboratories, Inc. Multibus is a trademark of Intel Corporation. Open Systems::Open Tools is a trademark of Heurikon Corporation. Open Tools is a trademark of Heurikon Corporation.

CIRCLE NO. 28

MAJOR SYSTEM COMPONENTS

The need to apply a reverse voltage to change the material back to its previous state can be used to advantage in displays with high information content. The material has memory; when the voltage is removed, the material tends to remain in the state to which it was set. "Depending on the material's formula, memory can be manipulated to last a very short time or to last virtually indefinitely," says Schut.

This memory characteristic, a form of "persistence," would eliminate the refresh problem in nonactive-matrix LCD displays. In nonactive-matrix LCDs, a pixel begins to fade when the voltage is removed during multiplexing. This aspect limits the number of lines that can be energized in sequence before it becomes necessary to go back and re-energize the first line. With Polyvision, the pixel would remain in its state until reverse voltage was applied; thus it wouldn't need to be refreshed.

In a high-information-content display, such as for a laptop computer, the inherently simple cell structure of Polyvision technology leads to a very simple addressing scheme. The cell structure is basically a layer of glass or plastic coated with a tin-indium-oxide transparent electrode. This is overlayed with the Polyvision material and a counter-electrode layer with conductors. The back is then sealed with an opaque or transparent protective coating. "In matrix addressing, half the current comes down one axis (row) and half down the other (column), and only at the point where both meet is there a reaction," says Schut.

On the issue of color, Schut notes that it's possible to use a transparent Polyvision material with a backlighting scheme and filters much the same way as is done with LCDs. The technology does have some inherent luminance, which might be sufficient in some color applications. Alpine is also investigating a side-lighting scheme in which light would come in from the side through the edges of the front glass panel and reflect off the color areas of the screen. Schut also stresses that while basic research into the technology is complete, commercial products available won't be for several years.

Flat-panel displays on flexible film-like material will lead to computer applications well beyond laptop and desktop systems. The possibility of wraparound simulation displays has already been mentioned. And no doubt many people would like be the first to walk into a meeting, pull down a large screen from a roller in the front of the room, and produce a full-color display without a projector. □

In Polyvision display technology, a layer of tin-indium-oxide is deposited on a glass or plastic substrate to act as a transparent electrode. The proprietary Polyvision material is layered on top of that, and then a counter-electrode layer is applied. The counter-electrode layer is topped by a sealant layer.

POLYVISION CELL STRUCTURE

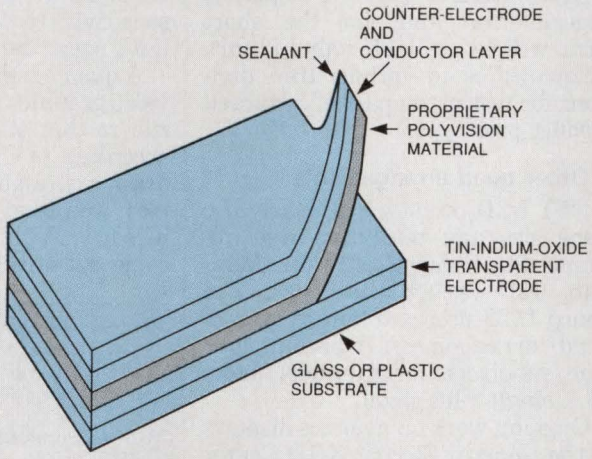

For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate number on the Reader Inquiry Card.

Staplevision
(201) 238-5314 ..... **Circle 229**
The Alpine Group
(201) 343-7600 ..... **Circle 230**

LOW PRICE PC ICE.™
$3,000 Each. Available Now.
80196
There *is* a low-cost 80C196 emulator that lets you debug your software *before* target hardware is available. One that supports in-target debug in real-time. Provides full symbolic support. And displays your source code during debug. All from your PC/AT host. All for $3,000.
And best of all, it's a Genuine Intel® emulator. The ICE™-196KB/PC Emulator. intel
Economy Without Compromise.
At this price, no 80C196 programmer should be without one. To order, or for more information, call:
800-874-6835
© 1989 Intel Corp
ICE™ is a trademark of Intel Corp.
CIRCLE NO. 29

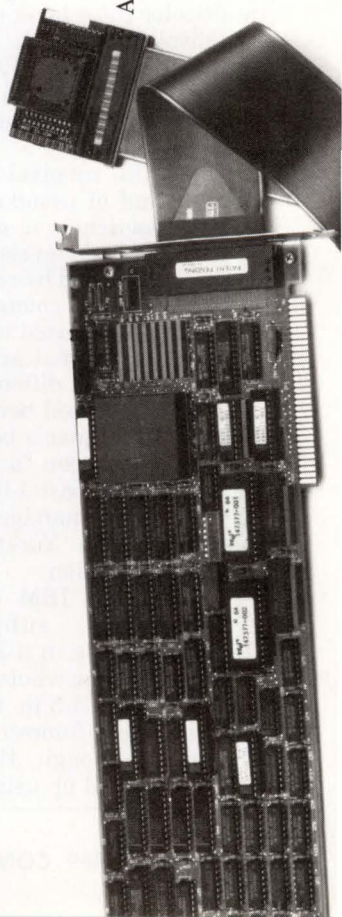

MAJOR SYSTEM COMPONENTS

Pseudo-grayscale scheme generates 16 colors on active-matrix display

David Lieberman, Senior Editor

A prototype display demonstrated at last month's Society for Information Display (SID) conference represents the largest size and highest resolution yet achieved for an active-matrix panel. Packing about 1.6 million subpixels and driven by the same number of amorphous-silicon TFT (thin-film transistor) switches, the 14.26-in.-diagonal color LCD is the result of a two-year joint development project of IBM Japan (Yamato, Japan), IBM's Research Division (Yorktown Heights, NY), and Toshiba's R&D Center (Kawasaki, Japan) and Electron Device and Engineering Laboratory (Yokohama, Japan).

In addition to its size and resolution, the display achieves double the color capability of other displays of this type by taking a novel approach to color generation. As is common in color LCDs, each pixel in the IBM display is made up of multiple subpixels that are colored by means of a multicolor filter layer on the front of the display sandwich. But IBM isn't using the familiar RGB triad; instead, it uses a quad arrangement that adds a white subpixel to the RGB triad.

The white subpixel gives the display a kind of pseudo-grayscale capability: switching it on and off gives two brightness levels to whatever color is generated by combinations of RGB. Up to 16 colors can thus be generated, compared to eight with a non-grayscale triad arrangement.

The perceived difference between, say, a red subpixel being on and red and white subpixels being on is the difference between "a pure, saturated red and a pastel-like red," says Denis Arvay, manager of information for IBM's Yorktown Heights Research Division.

Overall, the IBM display packs $1,440 \times 1,100$ subpixels and 720×550 pixels in a $288- \times 220$-mm active area. The whole display sandwich is a mere 1.5 in. thick and uses a novel type of fluorescent backlight. "Strangely enough, the fluorescent tube we wound up using is the same type used in tropical fish tanks because it has the right degree of brightness and the right color balance," says Arvay.

Arvay stresses that the display is experimental and, for the short term, will be available only in limited quantities to various IBM divisions for testing as part of advanced product prototypes.

Other quad arrangements

A TFT LCD prototype of nearly the same size and resolution was announced by Sharp Electronics (Mahwah, NJ) earlier this year. The Sharp LCD uses a different type of quad arrangement wherein four blue subpixels, for example, make up a single blue pixel.

Ongoing work on avionics displays at the General Electric R&D Center (Schenectady, NY) also makes use of a quad arrangement. The GE quad configures a pair of green subpixels—placed kitty-corner in the $2- \times 2$-subpixel matrix—with one blue and one red subpixel, explains Thomas Credelle, manager of optical electronic applications programs at the GE Center.

While the IBM quad doubles color capability by means of two brightness levels, the Sharp quad provides redundancy and the GE quad boosts perceived resolution because, says Credelle, the human eye has greater sensitivity to the green portion of the light spectrum.

A quad subpixel arrangement provides graphics capability that's superior to that of a triad arrangement, according to Credelle. "If you try to draw a straight line in a triad display, you tend to get a zigzag line," he says. "With a quad, straight lines come out straight." □

For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate number on the Reader Inquiry Card.

General Electric R&D Center
(518) 387-5000 ..... **Circle 226**
IBM
(914) 945-3000, ext. 3037 ..... **Circle 227**
Sharp Electronics
(201) 529-8200, ext. 8419 ..... **Circle 228**

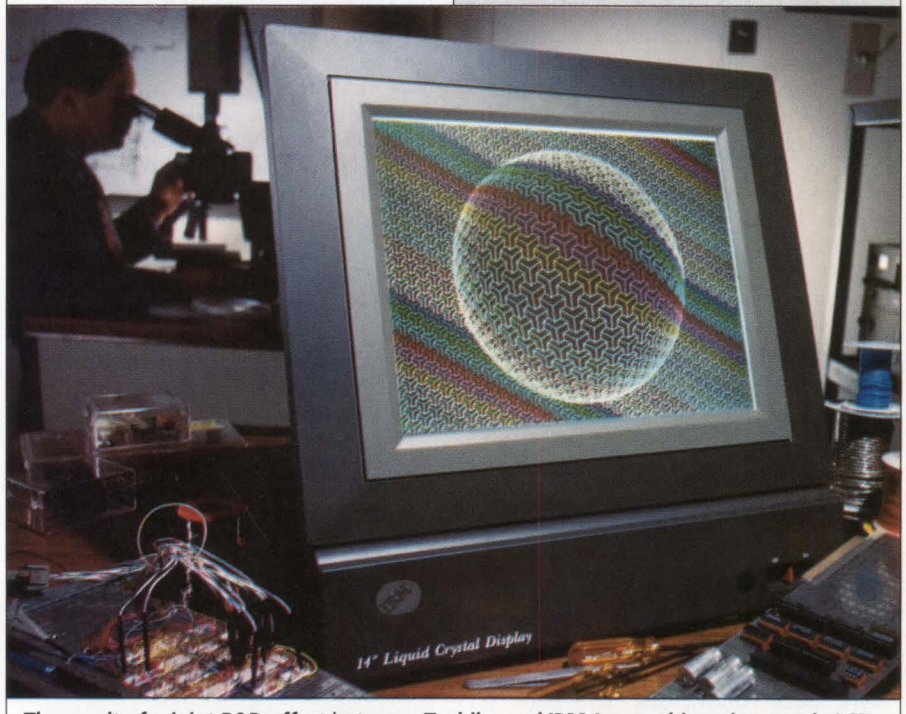

The result of a joint R&D effort between Toshiba and IBM Japan, this active-matrix LCD doubles the color palette for displays of this type. A white subpixel gives two brightness levels to any color created through a combination of red, green and blue subpixels, letting 16 colors be displayed simultaneously.

Here's the train of thought behind our new low priced VME board.

Introducing CPU 33. A high speed, economical VME computer.

Here's your ticket to high performance at an extraordinarily low price.

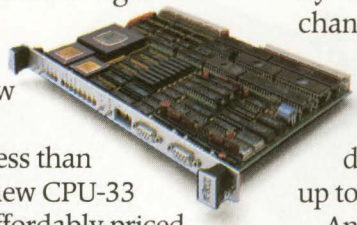

In fact, at less than $2500,* our new CPU-33 is the most affordably priced, performance-driven CPU board available.

And it's available at just the performance level you need. Including 16.7 MHz versions, with or without a turbocharging 68882 co-processor. Or a 25 MHz version with 68882, for those faster-than-a-bullet applications.

The CPU-33 starts with a 68030 processor that can easily deliver over 4 MIPS. And that's only the beginning.

The CPU-33 also features one megabyte of shared DRAM memory. Local and global accesses are interleaved to provide faster, more efficient management of VMEbus traffic.

Throughput is further enhanced by two message broadcasting channels, which allow for simultaneous data transmission to more than one CPU. Up to 256 messages can be defined and transmitted to up to 20 boards.

And if that still isn't enough, eight mailbox interrupts also support interprocessor communications. So you have lots of flexibility in designing multiprocessing systems.

CPU-33 also incorporates VMEPROM,™ our real-time multitasking kernel with debug monitor. It's designed to take full advantage of FORCE architecture. Best of all, it comes at no extra charge.

And to get your project to market sooner, FORCE provides a total system solution. Which includes third party software along with FORCE integration support. Including VxWorks, pSOS, VRTX32, OS-9 and PDOS. You also get extensive documentation. Plus access to FORCE's regional support staff, who are always ready to answer your application or system design questions.

For more information on the CPU-33, call FORCE at 800 BEST-VME. Or fax us at (408) 374-1146 for an immediate response.

If you want high performance at a low price, CPU-33 will put you on the right track.

FORCE COMPUTERS®

VME at its best.

FORCE COMPUTERS INC.
3165 Winchester Blvd., Campbell, CA 95008-6557
Phone (408) 370-6300, Fax (408) 374-1146

FORCE COMPUTERS GmbH
Prof.-Messerschmitt-Strasse 1, D-8014 Neubiberg/Munich
Telefon (0 89) 6 08 14-0, Telex 524 190 forc-d
Telefax (0 89) 6 09 77 93

*U.S. price. © 1989 FORCE COMPUTERS INC.
VxWorks is a trademark of Wind River Systems, Inc. pSOS is a trademark of Software Components Group. VRTX32 is a trademark of Ready Systems, Inc. OS-9 is a trademark of Microwave Systems. PDOS and VMEPROM are trademarks of Eyring Research Institute, Inc.

CIRCLE NO. 30

National Semiconductor

Is this the way your 16-bit controller company does business?

Call National.

THE HPC FAMILY. NOTHING ELSE GIVES YOU THE SAME DEGREE OF FLEXIBILITY AND CHOICE.

If you've been confronted with a "take it or leave it—that's all you need" attitude on the part of your 16-bit controller vendor, we suggest you take a new look at National's HPC™ family. Because for today's complex designs in information control, including printers, faxes, scanners, data storage and communications, you need a choice of optimized solutions that fulfill all your requirements.

Specifically, you probably need more in the way of intelligence, connectivity, interface and data handling. And much less in overall system costs. The HPC family more than measures up to these considerations. Let's look at just how well.

MAKE THE INTELLIGENT CHOICE.

Multiple functions mean multiple benefits. Including much faster context switching and interrupt response, due to our compact 54-instruction set. And higher system performance, because of our fast 30MHz clock rate, 67-ns instruction cycle, and 16-bit memory mapped architecture. And our ANSI-standard C compiler is just one way the HPC family can lower engineering costs, and speed your time to market.

HPC PRODUCT FAMILY SUMMARY

Part #	16-bit Timers	UPI	I/O	Memory ROM	Memory RAM	Features
HPC16003*	8	Yes	32	0	256	4 ICRs
HPC16004	8	Yes	32	0	512	4 ICRs
HPC16064	8	Yes	52	16K	512	4 ICRs
HPC16083*	8	Yes	52	8K	256	4 ICRs
HPC16104	8	Yes	32	0	512	8 CH A/D
HPC16164	8	Yes	52	16K	512	8 CH A/D
HPC16400	4	No	52	0	256	2 HDLC & 4 DMA
HPC16083MH	8	Yes	52	8K UV	256	UV Emulator

Standard features: Watchdog, Synchronous Serial Peripheral Interface, Uniform Memory Address Space, UART, 32X16-bit divide, 16X16-bit multiply, and available as standard cell.

DEVELOPMENT TOOLS:
PC-based development system
Relocatable C compiler, linker and symbolic debugger

ICRs = Input Capture Registers
HDLC = High-Level Data Link Control
**MIL-STD 883C*

Another is packaging, and specifically, our unique feature called TapePak.® As the latest generation in VLSI packaging, TapePak gives you a wide variety of industry-standard, high-density, high-leadcount options. When you put it all together, then throw in eight timers and up to 11 addressing modes, you can easily see why members of the HPC family are considered smart cookies indeed.

OPT FOR INTERFACE AND CONNECTIVITY.

The HPC family has more on-chip integration, and increased I/O. And that can mean a lot more capacity to manage system functions. It can also mean a lot less, when it comes to reducing board space and chip count. In addition, our Universal Peripheral Interface (UPI) allows you to partition your system, by using the HPC as a peripheral to a host processor for high-end applications.

PUT THROUGHPUT AT THE TOP OF THE LIST.

Thanks to our on-and-off chip uniform memory mapped peripherals, all data manipulation can be accomplished on any memory location, regardless of whether it's part of an external device, or on the HPC itself. Which can be very important in ISDN, SCSI and LAN applications.

ASK ABOUT OUR NEW COMMUNICATIONS CONTROLLER.

For data communications, and especially for those applications requiring protocol conversion, we offer the newest member of our HPC family: the HPC16400. It's a complete solution, supported by a full library of application-specific software. With two full-duplex HDLC channels, driven by 4-channel DMA, plus one full-duplex programmable UART channel for rate adaption, the HPC16400 is completely optimized for TE, TA, and line-card applications.

CALL AND FIND OUT MORE.

To get a free brochure on our full HPC family, call us today at 800-825-5805, ext. 100. Once you discover how we do business, you'll agree that when it comes to 16-bit controllers, there's only one company of choice: National.

National Semiconductor

© 1989 National Semiconductor Corporation

HPC is a trademark and TapePak is a registered trademark of National Semiconductor Corporation.

CIRCLE NO. 31

◁ EXCELAN ◁▷ DOS ◁▷ EXCELAN ◁▷ UNIX ◁▷ EXCELAN ◁▷ VMS ◁▷ EXCELAN ◁▷ MAC ◁▷ EXCELAN ◁▷ XENIX ◁▷ EXCELAN ◁▷ OS/2 ◁▷ EXCELAN ▷

◁ XENIX ◁▷ EXCELAN ◁▷ OS/2 ◁▷ EXCELAN ◁▷ DOS ◁▷ EXCELAN ◁▷ UNIX ◁▷ EXCELAN ◁▷ VMS ◁▷ EXCELAN ◁▷ MAC ◁▷ EXCELAN ◁▷ XENIX ◁▷ EXCELAN ◁▷ OS/2 ◁▷ EXCELAN ◁▷ DOS ◁▷ EXCELAN ▷

◁ DOS ◁▷ EXCELAN ◁▷ UNIX ◁▷ EXCELAN ◁▷ VMS ◁▷ EXCELAN ◁▷ MAC ◁▷ EXCELAN ◁▷ XENIX ◁▷ EXCELAN ◁▷ OS/2 ◁▷ EXCELAN ◁▷ DOS ◁▷ EXCELAN ◁▷ UNIX ◁▷ EXCELAN ◁▷ VMS ◁▷ EXCELAN ◁▷ MAC ▷

WHEN YOU INSTALL THE NETWORK AT THE WORLD ENTERPRISE NETWORKING SHOW, YOU'D BETTER MAKE SURE IT WORKS.

When Deere Tech Services of Moline, Ill., took part in installing the demonstration network at the 1988 Enterprise Networking Show, they knew that for the event's five days the network had to be faultless. So they used an Excelan LANalyzer® to set up, troubleshoot and monitor it.

Deere Tech Services knew that few companies have as much experience in connecting heterogeneous systems as Excelan.

And that Excelan's LANalyzer would be up to the job.

The Baltimore show's network was based on the Manufacturing Automation Protocol (MAP) running on broadband and the Technical and Office Protocol (TOP) running on Ethernet.

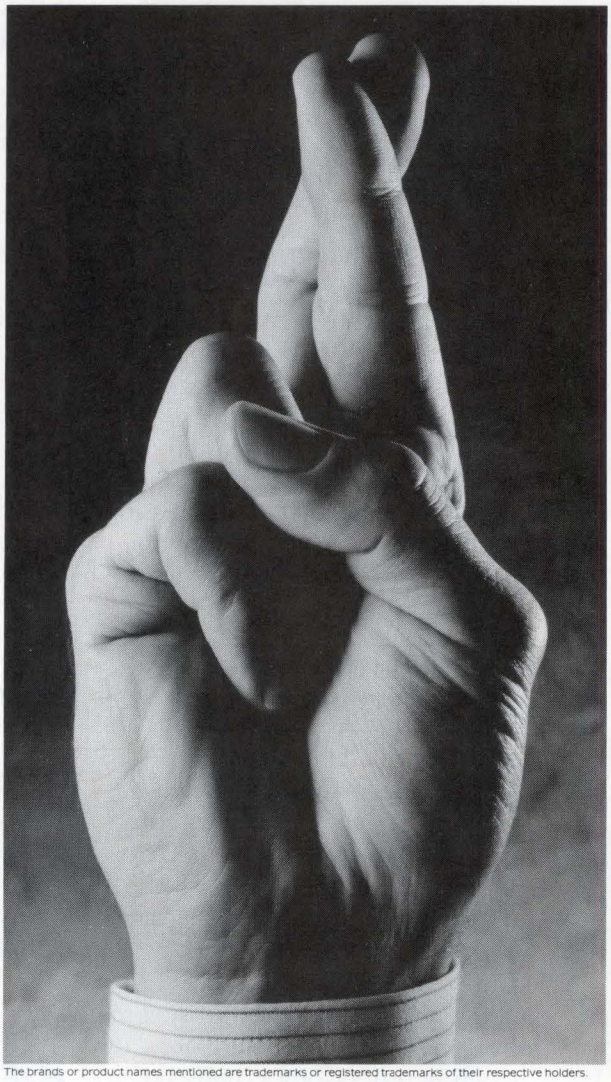

The brands or product names mentioned are trademarks or registered trademarks of their respective holders.

UP AND RUNNING.

It had to support large numbers of systems running several different protocols, all needing to be up and running in days.

But because LANalyzer, with it's mouse and menu driven interface, is extremely easy to use and provides high performance under heavy traffic, the network was ready in time for the show.

All of this was accomplished, because LANalyzer provides such features as simultaneous data capture and display, flexible multichannel filtering, and summary real-time monitoring, along with the ability to decode network protocols including TCP/IP, DECnet, AppleTalk and XNS.

PERFORMANCE NOT PROMISES.

The end result was that for the full five days of the show the network performed faultlessly with no downtime.

It's no wonder that many LANalyzer customers are repeat purchasers. Excelan's LANalyzer provides a user friendly, cost effective method of network analysis. And it's available as a stand alone unit or as a kit which can be installed in any IBM PC/AT or compatible.

And best of all, LANalyzer is made by the company that networks networks. Excelan.

To find out how the LANalyzer can solve your network problems, and to receive a free demo diskette, call 1-800-243-8526. Or write to us at 2180 Fortune Drive, San Jose, California 95131.

◂EXCELAN▸

WE NETWORK NETWORKS.

◁ EXCELAN ◁▷ VMS ◁▷ EXCELAN ◁▷ MAC ◁▷ EXCELAN ◁▷ XENIX ◁▷ EXCELAN ◁▷ OS/2 ◁▷ EXCELAN ◁▷ DOS ◁▷ EXCELAN ◁▷ UNIX ◁▷ EXCELAN ▷

CIRCLE NO. 32

16-bit micros adapt to survive

Trapped between single-chip microcontrollers and 32-bit microprocessors, 16-bit CPUs are fighting for survival. To succeed, they must learn their enemies' tricks.

Ron Wilson, Senior Editor

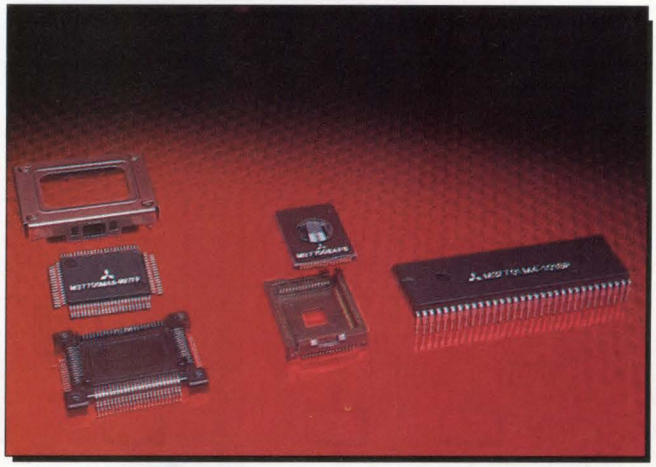

More and more designers are choosing 16-bit CPUs for their combination of speed and compactness. The Mitsubishi M37700 microcontroller family is one of the most recent choices, combining the sophisticated instruction set and computing speed of a 16-bit microprocessor with the memory and peripheral content of a board-level computer, but in a single chip.

Pity the poor 16-bit microprocessors. Once the rulers of personal computing and embedded control applications, 16-bit CPU chips are now besieged on two fronts—8-bit single-chip microcontrollers offer higher integration at lower cost, while 32-bit microprocessors promise greater speed and upward compatibility.

Caught between these antagonists, 16-bit processors must use a number of niche strategies to survive. Some chips are clinging to a particular body of code, running a select group of programs faster than their 32-bit competitors. Other CPU architectures are beginning to look like single-chip microcontrollers, cultivating on-chip peripheral devices. And new microcontroller products with 16-bit data paths are appearing, offering both integration and bandwidth in a single chip.

Using these strategies, both the familiar 16-bit architectures and the new microcontrollers are likely to survive, even prosper. But in the process, many of the ground rules of 16-bit design will change, creating both problems and opportunities for the design team.

The most evident struggle between 16-bit microprocessors and their competitors is in the personal computer market. Here, the 8086 and 80286 architectures from Intel (Santa Clara, CA) are still going strong because a large body of shrink-wrapped software keeps them from being driven off by more modern 32-bit systems. But surprisingly, the 8086 and 286 are holding on even against Intel's own 32-bit offerings, the 386 and 16-bit-bus 386SX.

"The 286 isn't disappearing—it's still the workhorse CPU," says Ahmet Alpdemir, systems product line manager at Chips and Technologies (San Jose, CA). One reason is raw speed. Fast 286 PCs with 20- or 25-MHz CPUs and page-interleaved memory schemes can execute 16-bit PC software—including OS/2 applications—faster than 386SX systems can. And the 286's 16-bit external bus gives it a big system-cost advantage over the 32-bit 386.

The viability of the 286 architecture in personal computing is reflected in continued engineering development. Vendors have pushed the speed of the CPUs to 20 and even 25 MHz. And vendors such as Chips continue to provide motherboard chip sets specifically for the 286, tracking the CPU clock frequency as it grows.

The 16-bit lock on low-end PCs has led a number of industry watchers to speculate about the possibility of a PC-specific single-chip computer. The 286 die, with its 16-bit ALU, registers and data paths, is considerably smaller than a 386SX die. The 8086 die is even smaller. Consequently,

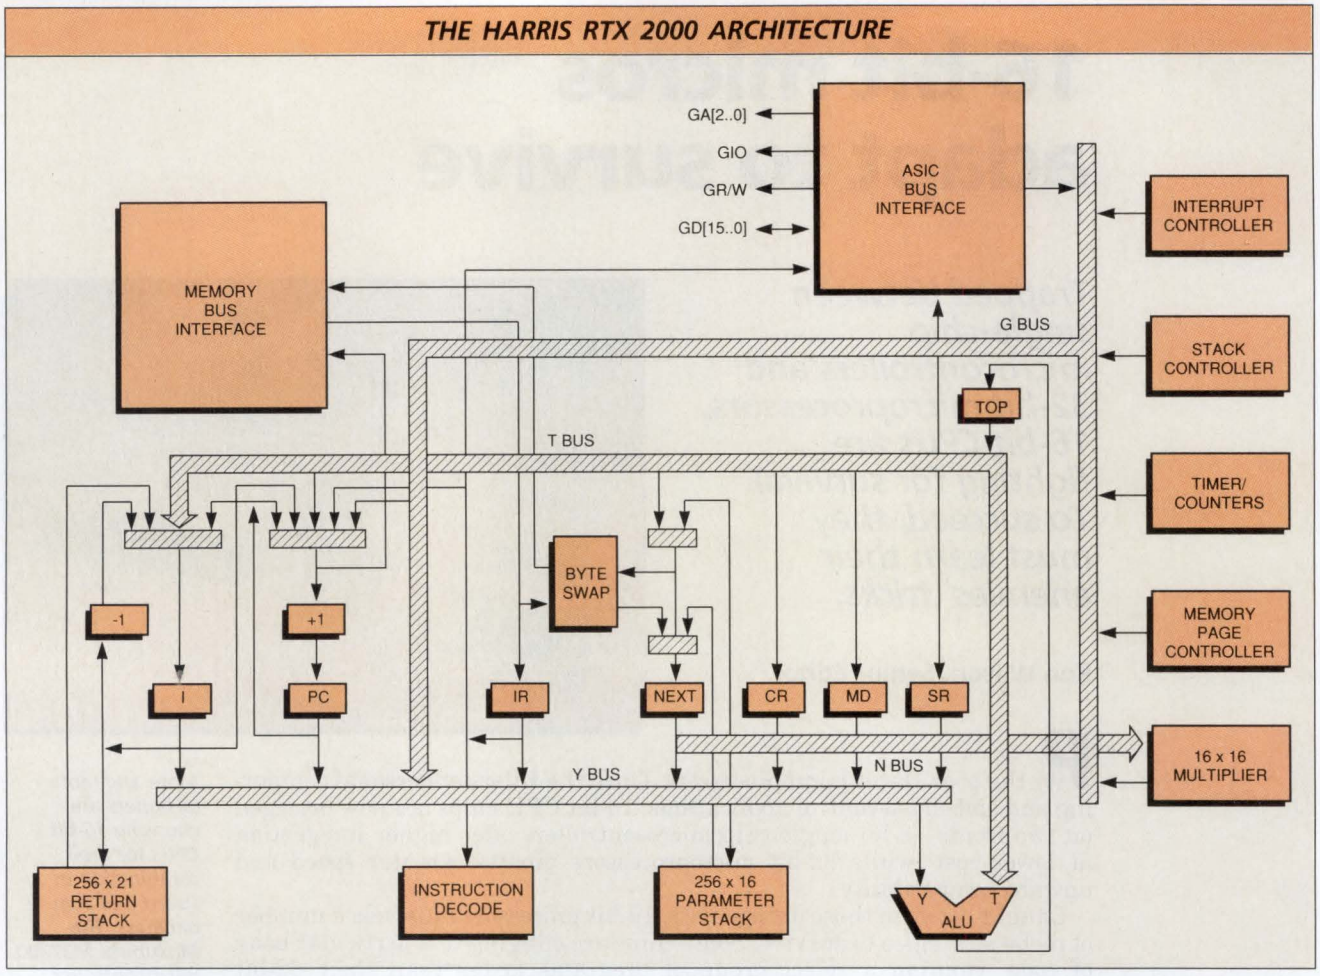

Because 16-bit microprocessors are often used in applications requiring very high bandwidth, the Harris RTX 2000 provides the ASIC Bus, an extremely fast parallel connection that bypasses the chip's memory bus interface and connects directly to the arithmetic unit.

there might be enough room for all of, or most of, a motherboard chip set to be integrated onto the CPU chip, which would result in a one-logic-chip PC.

In fact, some of the members of the NEC V Series are examples of just such an attempt to increase the integration of the 8086 architecture. The V25 and V35 CPUs combine an 8086-superset CPU core with memory and a variety of peripheral devices such as serial ports, parallel ports and timers. The processors are intended not as single-chip PC motherboards but as microcontrollers with PC software compatibility.

"We're seeing microcontrollers replacing microprocessors in embedded applications," says Prabhat Dubey, business operations manager for NEC's microprocessor strategic business unit. "Users compare the performance of the 8086 to the speed of our V30, but then they compare it to the V35 microcontroller. They find that in many applications, the V35 not only increases integration but also is faster for operations that can stay on-chip. This has led to V25 and V35 applications in telephone line cards, fax machines, disk controllers and many other areas."

Tightly integrated subsystems

The ability of 16-bit CPU cores to form tightly integrated, yet fast, subsystems is being explored by a number of vendors. Vadem (San Jose, CA), for example, has developed a peripheral chip set for the NEC V40. "Using the chip set, you can build a complete controller from the V40, nine logic chips and memory," says Vadem chief operating officer Paul Rosenfeld. "The Vadem chips adapt the peripheral devices on the V40 so that they can be used either as strictly PC-compatible peripherals or as general-purpose devices in the control system." Vadem has recently announced a new chip set that extends support to the new 10-MHz version of the V40.

Harris Semiconductor (Melbourne, FL) has taken a rather different approach to increasing the integration of its Forth-executing RTX processor chip. The RTX processors include timers and interrupt controllers, as well as high-speed stack memory, on the CPU chip. But to incorporate more sophisticated peripheral devices, the RTX uses a unique type of dedicated I/O bus.

"We observed that customers were attempting to include very high-bandwidth peripheral devices, such as dedicated CPUs and image-processing chips, in their embedded systems," says Harris senior scientist Christopher Malinowski. "But it's very hard to integrate such devices into a system without a dedicated I/O bus: memory-mapped I/O schemes are just too inefficient. So we developed an I/O bus that bypasses the memory bus circuitry and connects directly to the ALU. This permits the user to access data from I/O devices as quickly as he could get data from the top of the stack. Since

OS/2 makes 286 the best choice for cost/performance

There are situations in which software requirements can give a 16-bit microprocessor a performance advantage over a nominally faster 32-bit chip. The high-end personal computer market is an example, and IBM's OS/2 strategy is the reason.

OS/2 was created to overcome a number of the deficiencies in DOS. For instance, OS/2 is a multitasking system, using protected mode to accommodate multiple active jobs on a single CPU. It's able to use address ranges larger than DOS's 640-kbyte limit. The new operating system also provides for integrated networking and window-managed graphics, facilities that were never native to DOS. Yet OS/2 lets users make a smooth transition from their existing hardware and applications to the new, more powerful system.

One facet of this transition strategy is that OS/2 users can run existing DOS programs, in real mode, on OS/2-equipped machines. Another important facet is that OS/2 and its applications are written in 16-bit, 80286-compatible code. The operating system doesn't use the additional 32-bit instructions or memory-management features of the 80386 architecture. This decision by IBM lets OS/2 run on a wider variety of machines, covering a wider range of prices. But it also has an important performance implication.

The 386 is no faster than a 286 when running 16-bit software at the same clock speed. The primary reason is that the 286 executes more 16-bit instructions in fewer clock cycles than does the 386 or 80386SX. Out of 190 existing 16-bit instructions, the 286 is faster on 74 instructions, the 386 is faster on 50 instructions and the processors are the same on 66 instructions.

OS/386-compatible PCs

In fact, today's applications, as well as those being developed, won't take advantage of the 386 until an OS/386-specific version of the operating system is available sometime in 1990, with OS/386 applications becoming available sometime in 1991 or 1992. But there is no guarantee that today's 386—or especially the 386SX—will be configured to run new 386 32-bit software four years down the line. After all, the first 286-based PCs running at 6 MHz with 640 kbytes of memory are hardly suitable for running 16-bit OS/2 now. The same scenario will be likely for today's 386 PC if it's intended to run OS/386 several years from now.

Processor speed crucial

What is important for the OS/2 operating system, then, isn't whether it's run on a 286 or 386, but rather the speed of the processor. Multitasking is the bulk of the processor's work, actually accomplishing several things at once by dividing the computer's time into 'time slices' lasting a fraction of a second. These time slices are handled so fast that it appears as if programs are run simultaneously. Since the processor is actually doing all the tasks at separate intervals, the faster the processor the quicker the multiple tasks will be completed.

The 386 is no faster than a 286 when running 16-bit software at the same clock speed.

An adequately equipped 286 system running at a minimum of 12 to 16 MHz with Video Graphics Array graphics is a cost-effective OS/2 foundation. The 16-, 20-, and 25-MHz 286 systems of today have the 16-Mbyte memory access capability and protected mode for the multiple applications required of OS/2.

Furthermore, the 286 is an economical choice for the system designer. The 386 is obviously a higher priced component (due to its die size, packaging, and complex processing), and systems built around this device require 32-bit peripherals, thus boosting design cost. A 386-based system costs an average of 35 percent more than an equally fast 286-based system.

Consider, for example, several OS/2 benchmarks that demonstrate how a 286 system compares to a 386 and 386SX system. In these benchmarks, the 286 and 386SX systems are Everex Step PCs, and the 386 is an IBM System 80. The 386 PC uses page-mode memory access for 0.8 average wait states with 80-ns dynamic RAMs. Both the 286 and 386SX run at zero wait states with 60-ns DRAMs.

In a benchmark based on the R:Base database program, the 286 PC outperformed the 386 PC by 4 percent and the 386SX by 24 percent. In a similar benchmark based on the Paradox database program, the 286 was 5 percent faster than the 386 and 18 percent faster than the 386SX.

And in a third benchmark based on IBM Display Write 4.2, the 286 was 7 percent faster than the 386 PC and 8 percent faster than the 386SX.

Benchmark performance is obviously influenced by both the processor and the memory interface. In the PC systems used, the memory interface was relatively equal (0.8 wait states on the 386, and 0 wait states on the 286 and 386SX machine). Thus, the performance difference measured between the 286 and 386SX was primarily a result of the different processors with the 286 performing faster than the 386SX.

A 0.8-wait-state 386 system has about a 9 percent performance degradation over a zero-wait-state 286 system. Taking this into account, the 286 and 386 performed essentially the same.

The 286 has a lot of life left both for DOS and OS/2 operating systems, and will continue the trend toward faster clock speeds. According to Dataquest (San Jose, CA), the 286 will increase its current market share of IBM and compatibles from 30 percent to 33 percent by 1992 and will become the entry-level PC, replacing the 8086/8088 based-machines. Following a stable path to OS/2, the 286 is the most effective platform in the search for optimal cost/performance.

Linda Bishop, BSEE, product marketing engineer, Personal Computer Products division, Advanced Micro Devices

the RTX processors can overlap activity on all five of their internal buses, the I/O operation can be done in parallel with other other tasks."

A role in multiprocessing

The Harris designers are responding to a strong trend in embedded systems design: multiprocessor systems. Increasingly, system architects are partitioning tasks into a number of major pieces and assigning an appropriate CPU to each piece. This trend is creating another important niche for 16-bit microprocessors—as computational or data-intensive nodes in multiprocessor systems.

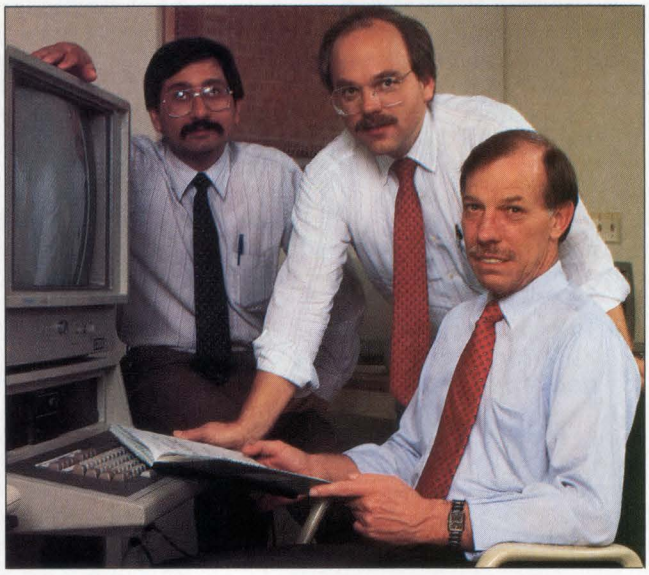

Rao Gobburu (left), Richard Sessions and Bernie Kute of National Semiconductor have watched their 16-bit microcontroller, the HPC, become a mainstay of the embedded computing market. Now, use of the HPC as an ASIC core is growing too, spreading from the auto industry into communications and peripheral controller applications, according to Sessions.

One early adopter of multiprocessor architectures was the telecommunications industry. The huge number of small tasks and the need for on-sight expandability made multiprocessing a natural solution. Pete Kusulas, systems development department head at AT&T (Middletown, NJ), says, "We use a range of 8-, 16- and 32-bit processors in the Definity product line. Typically, we will use 8-bit microprocessors as controllers on circuit packs, and we will use either 16- or 32-bit microprocessors where there is a need for more computing performance, such as in protocol handling."

Unlike the PC market, where customer preference can force designers to use a 32-bit CPU, the telecommunications industry still has the luxury of making an engineering, rather than a promotional, choice. "We don't see any wholesale migration to 32-bit CPUs," Kusulas says. "We will continue selecting processors on the basis of application needs. Right now, 16-bit chips are perfectly adequate in many areas—in some cases, they're actually faster than 32-bit chips."

AT&T field introduction group supervisor Frank Wyatt concurs. "We've looked at 32-bit processors for applications such as terminating protocols, where you need a combination of high bandwidth and large buffer size. But if the address range you have to cover is less than about 32 kbytes, it's very hard to beat the cost and performance of 16-bit chips."

When multiprocessor system designers upgrade from 16-bit microprocessors, they as often as not move to a more highly integrated 16-bit processor. Wyatt cites a case where AT&T had an 8086 CPU terminating an X.25 protocol, and replaced it with an 80186. This tendency to upgrade is putting two pressures on vendors: first, to provide more highly integrated—and more application-specific—16-bit chips; and second, to maintain software compatibility as they improve their product lines. "Software compatibility in new releases is very important to us," says Kusulas. "We don't have revolutions between new product releases—we evolve."

Wyatt adds, "The ideal is instruction-set compatibility, but we also use design methodology to help with portability. Tools such as machine-independent high-level languages and operating software can make the transition easier."

So far, chip vendors have shown more aptitude for integrating hardware than for maintaining software compatibility. Except for single-chip products that are direct descendents of 16-bit microprocessors, the more integrated parts usually have a different instruction set. But the new parts are beginning to reflect, through their on-chip peripherals, the multiprocessing environments in which they are increasingly going to be used.

Many of the more recent 16-bit integrated parts have some explicit hardware for interprocessor connection. National Semiconductor (Santa Clara, CA), for example, includes a universal peripheral interface on its HPC microcontroller. "The interface lets us map the HPC into the memory space of a host microprocessor," says Rao Gobburu, National marketing manager for configurable microcontrollers. "There is no functional equivalent for this capability in 8-bit processors."

Provision for memory-mapped interfaces is a natural outgrowth of the reason 16-bit parts are being used: the need for bandwidth. In a laser printer, for instance, there is a mixture of heavily computational graphics tasks, high-bandwidth communications tasks and event-driven control tasks. "By moving the communications and control tasks onto a $10 microcontroller," claims Bernie Kute, National director of ASIC strategic products, "you can get a 10-times improvement in the speed of the computational engine, because you've prevented context switches."

"And when you have such a large number of tasks, some of which require high bandwidth, a 16-bit controller is ideal," adds Gobburu. So the HPC architects provide a high-bandwidth external interface, to make sure the interprocessor connection doesn't become a bottleneck.

Fighting for compactness

Multiprocessor architectures are great when there is enough space in the system for several chips. But in the most compact systems, such as 3½-in. disk drives, there may not be room for a second processor. So all of the tasks, whether computational, data-moving or event-driven, fall on a single CPU. In this domain, single-chip 16-bit microcontrollers are establishing a niche of their own.

Frank Hannon, director of controller engineering development at Maxtor (San Jose, CA), has observed the move to single-chip controllers. "In our 5¼-in. line, we use an 8031 for interface protocol and an 8096 for

ASIC technology shapes 16-bit embedded control microsystems

At the 16-bit level, embedded microsystems are largely involved in information control. Hence, both their CPU cores and associated hardware differ markedly in function, architecture and topology from those of more familiar data processing-oriented microprocessors.

Microprocessors' strengths are their general-purpose data-processing capabilities, system software, user interfaces and application software base. But mainstream embedded controllers march to different music. Here, the required functionality and performance must be delivered in timely, as well as cost- and space-effective, solutions. Hence, microcontrollers are benchmarked not only against program execution speed but also against code density, number and effectiveness of I/O channels, and the number and utility of *relevant* hardware functions. Moreover, embedded control microsystems often reside in severe environments, such as in vehicle engine compartments.

Because of microcontrollers' application in, or as, embedded microsystems, the only relevant way to benchmark them is in situo. Obviously, a unit designed from scratch could exclude all but those functions required for the application at hand. And a customized microcontroller's performance could be tailored to match throughput and/or control requirements on a task-by-task basis. The architecture, hardware and instruction set—and even the package—could also be custom-tailored. Unfortunately, the cost, reliability, testability and time-to-market penalties of designing such a system would be too great for most mainstream applications.

A great degree of microcontroller customizing, however, is practical. Thanks to the marriage of advanced CMOS processing and standard cell-based application-specific IC technologies, fast, compact, low-power and production-worthy, configurable 16-bit microcontrollers can be put together from a library of standard cells. Several types of standard cells can be used, most commonly CPU core megacells, MSI/LSI functional logic macrocells, expandable memory cells, gate-level random logic cells and linear circuit cells.

Of course, the cells must be members of the same library. This means that they must conform to a universal set of design rules and that all library elements must be fabricated on the same process. Obviously, these cells must also be compatible with the same design automation system.

At National Semiconductor, these requirements are automatically met because all LSI and VLSI devices are built on the same M^2 CMOS core process, on the same design automation system.

And testability isn't ignored: each core megacell includes additional logic to ensure access to the core through port pins. This allows the test vectors for the core megacell to be identical to those for the standard part. In addition, test vectors follow each standard cell into the library. Thus, testability of both the CPU-core megacell and macro-function combinations is assured by National's design-verification team.

Moreover, when modules are added to the production-proven core process, analog, memory and even bipolar elements can be created as needed. To further ensure silicon efficiency, core CPU megacells and functional logic macrocells use a topologically compatible rectangular structure, where the width is fixed and heights vary. Of course, core CPUs for National's families of CMOS 8-bit and 16-bit microcontrollers each mate with their own dedicated set of topologically matching macrocells.

Standard vs. custom-made

Though size, power dissipation and performance influence a system designer's choice between standard and customized microcontroller solutions, cost and availability are usually the most important considerations. Obviously, standard off-the-shelf microcontrollers are available immediately. Standard product families combining features such as analog-to-digital converters, universal asynchronous receivers/transmitters, timers and comparators are also the most cost-effective alternatives. The optimum microcontrollers for most embedded control applications, however, would usually differ from their family counterparts in pin count, instruction set and on-chip hardware.

The degree to which the differences between a fully customized microcontroller and its commercial counterpart affect the size, cost and performance of an entire microsystem must be decided case by case. As a rule of thumb, for 16-bit machines, wedding standard-cell functions to a core CPU via design automation involves from $40,000 to $150,000 in nonrecurring engineering (NRE) charges. Hand packing a microcontroller chip raises the NRE ante to the $750,000 area. Hand packing usually saves over 25 percent of the microcontroller die area.

Mainstream CMOS 16-bit microcontrollers cost from $10 to $25. Hence, from the standpoint of cost alone, a volume of 100,000 could make it practical to use a microcontroller based on a core CPU and macrocell functions, and customized via design automation tools. At volumes of 1 million, silicon area savings would offset the higher NRE charges in a typical hand-customized 16-bit microcontroller.

Microcontrollers customized via design automation are available faster than those that are hand-customized. Prototype turnaround times are only 6 to 10 weeks, compared to the 6 to 18 months to deliver the denser hand-packed versions.

In the face of ever-narrowing market introduction windows, National suggests the following three-phase migration to the smallest, highest-performing and most cost-efficient microcontroller die:

- Phase 1—Enter market and secure share with a catalog core CPU microcontroller and external components such as gate arrays and/or standard LSI, MSI and SSI functions.
- Phase 2—Bring the most cost- and performance-critical circuitry on-chip via design-automated macrocell functions.
- Phase 3—As volume increases, improve performance and reduce cost and size with the fully customized version.

In all cases, it's important to cast the chip in a shrinkable technology. Shrinkability assures that the chip can migrate in size and performance as process technology advances.

Rao Gobburu, MSEE, marketing manager, Configurable Microcontrollers, and
Bernie Kute, BSEE, director, ASIC Strategic Products, National Semiconductor

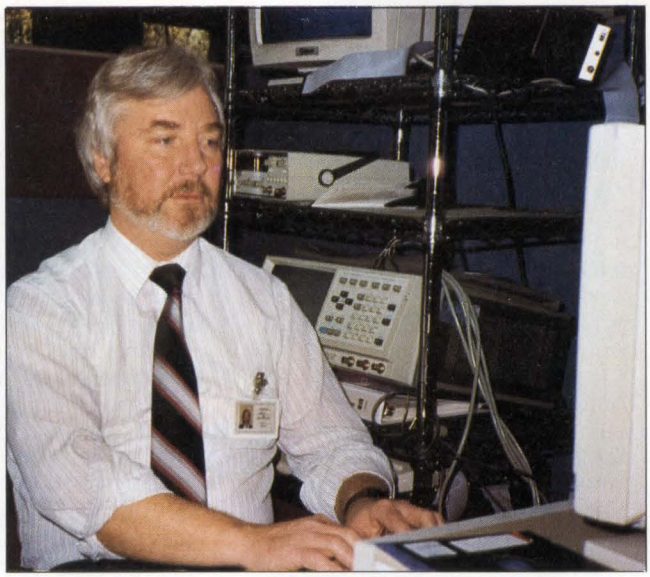

When the space constraints of a 3½-in. disk embedded controller demanded a single-processor solution, Frank Hannon, director of controller engineering development at Maxtor, turned to a 16-bit processor for its speed and multitasking capability. The part had to absorb tasks assigned to two separate processors in the company's 5¼-in. controllers.

servo control," he says. "But on the 3½-in. product, space is at a premium. We use just a single HPC. So both the protocol and the servo tasks land on one processor."

Vendors are acutely aware of this trend. "Disk drives need computing power, but they're also very cost sensitive," says Ike Saeed, product engineering manager MCU at Mitsubishi Electronics (Sunnyvale, CA). "You may have one chip controlling the spindle motor through a pulse-width-modulation output, controlling the head stepper through a timer output and managing the flow of data, too."

The need for elaborate I/O devices in such a system suggests the use of a sophisticated single-chip microcontroller. But the need to address large buffers, handle a bus protocol and execute a high-resolution positioning algorithm for a moving disk head mandates a 16-bit CPU core. Mitsubishi's most recent part, the M37700, is built around a 16-bit core that provides hardware multiply and divide, bit/byte manipulation and a variety of addressing modes. Outside the core, the chip integrates two universal asynchronous receivers/transmitters (UARTs); an 8-bit, 8-channel A-D converter; eight multifunction timers; 2 kbytes of RAM; and 32 kbytes of EPROM.

As such parts gain the considerable computing muscle required to undertake the tasks of an embedded disk controller, the distinction between microprocessors and microcontrollers starts to blur. "Traditionally, microprocessor applications were geared to data processing, and microcontroller applications to device control," Saeed observes. "But in the 16-bit families the gap is narrowing. The 7700s can access up to 16 Mbytes of memory already. And the next generation will see richer instruction sets and 20-MHz CPU clocks. A CPU core such as that could replace a microprocessor, even in data-processing applications."

Remarkably enough, embedded disk controllers aren't the most compact application for 16-bit microcontrollers. That role goes to smart motor controllers. "In inverter applications," says Lee Davidson, product marketing engineer at Intel (Chandler, AZ), "a 16-bit microcontroller will generate ac signals for a motor, controlling magnitude and phase in real time. With this kind of intelligent control, you can reduce peak loading, cut back on noise and, frequently, use a smaller motor with smaller starter windings."

Davidson explains that an inverter application requires the microcontroller to monitor analog quantities such as armature position, torque and temperature. Then the CPU core calculates optimal pulses to drive a particular motor winding.

"The Japanese, with their need for very efficient appliances and air conditioners, have recognized the role of microprocessors in power management for a long time," says Tony Sica, Intel ASIC division marketing manager for microcomputers. "Now, U.S. companies are beginning to understand that point too—microcontrollers aren't just for running the user interface anymore."

These applications require not only low cost and high computing speed but also some very specific peripheral hardware. In the Intel chips, for instance, the CPU can determine the rise time, width and fall time of a pulse, and then send the data to a mailbox. The pulse-width modulation (PWM) output peripheral will generate the pulse without further CPU intervention. And the A-D converter inputs on the parts are getting more sophisticated as well. Because the 8097BH, for example, includes an on-chip sample-and-hold, it can exploit its 22-μs conversion time without requiring an external device to hold the signal at dc for 22 μs.

It's clear that 16-bit processors have established themselves in a variety of niches—as the mainstay CPUs in personal computing, or as computational or data-intensive nodes in multiprocessor systems. In each of these niches, 16-bit chip vendors are using integration as a survival tool. But as the vendors innovate, they strain against the system designers' needs for continuity in instruction sets and development tools.

Each vendor, and each design manager, must decide when new instructions and more-specialized on-chip devices justify beginning a design with minimal tools and having to write new code. That choice, as much as any architectural decision, may determine the staying power of 16-bit processors. □

For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate number on the Reader Inquiry Card.

Advanced Micro Devices
(512) 385-8542 ..... **Circle 248**
AT&T
1-800-247-1212 ext. 190 ..... **Circle 249**
Chips and Technologies
(408) 434-0600 ext. 2563 ..... **Circle 250**
Harris Semiconductor
(407) 729-4629 ..... **Circle 251**
Intel
(602) 554-2388 ..... **Circle 252**
Intel
(408) 765-8080 ..... **Circle 253**
Maxtor
(408) 432-1700 ..... **Circle 254**
Mitsubishi Electronics
(408) 730-5900 ..... **Circle 255**
National Semiconductor
(408) 721-5185 ..... **Circle 256**
NEC
(415) 965-6045 ..... **Circle 257**
Vadem
(408) 943-9301 ..... **Circle 258**

68020 EMULATION NOW AT 25MHz.

68020-25MHz

68000-16.7MHz

80286-12.5MHz

80C186/C188-16MHz

Applied Microsystems has raised speed limits on the best real-time, transparent emulation you can get.

So now emulating faster moving targets is getting easier.

But the powerful edge you get with Applied Microsystems is built on more than speed. Applied Microsystems provides the most comprehensive solution in the industry. You get an easy user interface, full-width trace, multi-condition break points, cross triggering, source level debugging and 2 Mbytes of RAM overlay.

Applied Microsystems emulators are controlled from the host you normally work on. Your target system runs exactly as if its microprocessor were in place. And you have the option of SCSI, the industry standard high-speed bus that transfers data at up to 1.5 Mbytes per second.

Nobody offers more performance or a quicker way to develop and integrate everything from microcontroller designs to 8-, 16- and 32-bit systems. You can also expect the back-up you need from our international network of applications engineers who provide the support to meet your project objectives.

For a demonstration, write Applied Microsystems Corporation, P.O. Box 97002, Redmond, Washington USA 98073-9702. Or call (800) 426-3925. In Washington, call (206) 882-2000.

Applied Microsystems Corporation

In Europe contact Applied Microsystems Corporation Ltd., Chiltern Court, High Street, Wendover, Aylesbury, Bucks, HP22 6EP, United Kingdom. Telephone 44-(0)-296-625462. FAX 44-(0)-296-623460.
In Japan, contact Applied Microsystems Japan Ltd., Nihon Seimei, Nishi-Gotanda Building, 7-24-5 Nishi-Gotanda, Shinagawa-KU, Tokyo T141, Japan. Telephone 03-493-0770. FAX 03-493-7270.

AMC 247

CIRCLE NO. 33

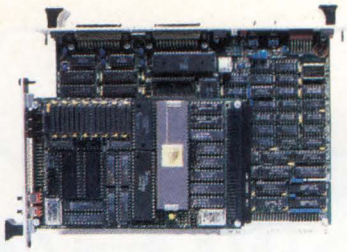

Processor Modules

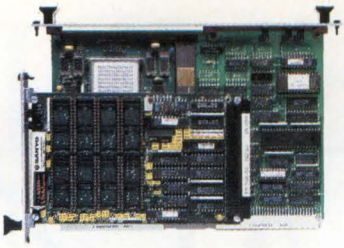

Memory Modules

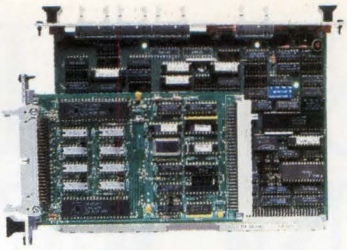

Peripheral Modules

NEW! XVME-682 VMEbus PC/AT Processor Module

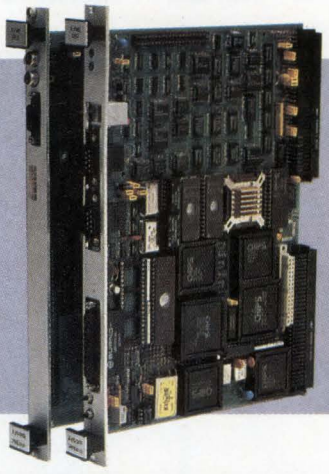

- Fully IBM PC/AT Compatible
- 10 MHz 80286 Processor
- 1-4 Mbytes of Dual-Ported RAM
- EGA/CGA Graphics Controller
- Hard Disk/Floppy Disk Controller

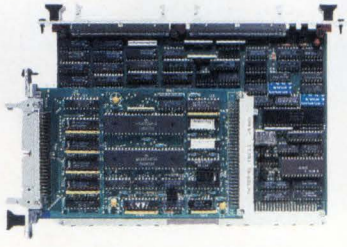

Digital I/O Modules

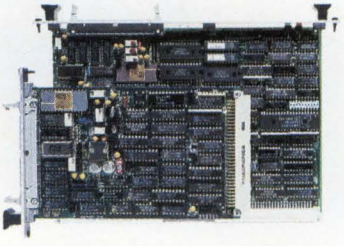

Analog I/O Modules

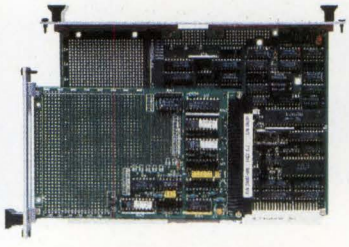

Prototyping Modules

OUR FAMILY HAS GROWN!

And now provides you with complete VMEbus solutions.

Several years ago, we started small in industrial VMEbus I/O modules. Our sales grew. And our customers asked for more.

So our product line grew, too. Today you can select from our complete family of VMEbus boards. In single-high or double-high. Including complete development systems. Backed by the best support and warranty in the business.

For a lot more information on our growing family of VMEbus modules, write or call us today.

xycom®

The Hardhat Computer People.
750 North Maple Rd., Saline, MI 48176
(313) 429-4971 TWX 810-223-8153

VME Products			Software Support
Processor Products	ESDI Controller	3U and 6U Development Systems	OS-9*
Memory Products	BITBUS Controller	VME Terminal Products	PDOS*
Digital I/O	PAMUX Controller	Servo Controller Module	PSOS*
Analog I/O	1553 Interface Module		UNIX V/68K*
SCSI Controller	Thermocouple Input Module		SUN Workstations*
Winchester/Floppy Controller			VERSAdos*
			MTOS*

*Trademarks are acknowledged to: Microware Systems Corp. (OS-9); Eyring Research Institute, Inc. (PDOS); Motorola, Inc. (VERSAdos); V/68K Motorola, Inc. (UNIX); Software Components Corp. (PSOS); SUN Microsystems, Inc. (SUN Workstations); Industrial Programming, Inc. (MTOS).

CIRCLE NO. 34

Logic analyzers and emulators evolve with the times

Racing to keep up with the demands of today's breed of high-performance systems, logic analyzers and emulators take on new roles and capabilities.

Bill Harding
Senior Editor

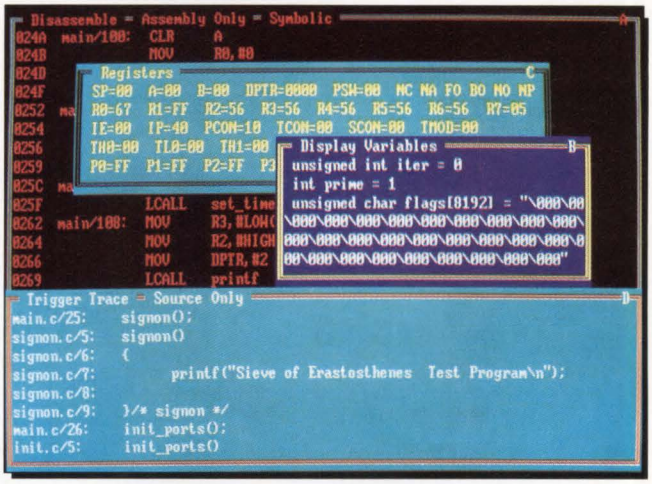

Logic analyzers and in-circuit emulators help designers trace applications software execution in a target system. Many emulators, when coupled with a host computer, allow designers to trace high-level source code as well as assembly language. This source-code debugger from Huntsville Microsystems displays both assembly language code and the high-level language source code that generated the assembly instructions.

Faster, wider, and deeper—those are the watchwords in today's logic analyzer market. To meet the demands of RISC-based, 32-bit and multiprocessor designs, logic analyzers must be able to monitor signals at higher clock rates, must provide more channels to monitor bigger address and data buses, and must have more memory to display larger segments of monitored waveforms.

For in-circuit emulators, the watchwords are faster and smarter. Like logic analyzers, emulators must be faster to keep up with the higher speeds of the new breed of microprocessor chips. But unlike logic analyzers, which are passive monitors, emulators exert some control over a target system. The challenge to emulator manufacturers is to endow emulators with the control capabilities needed to debug complex applications programs and to integrate complex systems.

In the high-end, most-competitive segment of the logic analyzer market, vendors have turned to modularity, which lets users configure their systems without purchasing any more power than they need. And many companies let users trade channel width for sampling speed by halving the number of channels in order to double the sampling speed for timing analysis. Prices for these systems start at about $20,000 and can range to about $200,000 in larger configurations.

One of the most popular high-end logic analyzers is the DAS 9200 from Tektronix (Beaverton, OR). It was one of the first systems to allow the user to trade channel count for measurement speed. In its widest configuration, the DAS 9200 can have 540 channels at 20 MHz, with a 32-kbyte acquisition memory. If 20 MHz isn't fast enough, the DAS 9200 can be configured to acquire 432 channels at 200 MHz with 4 kbytes of memory, or 160 channels at 2 GHz with 8 kbytes of memory. It can also be set up with 1,008 stimulus channels at 50 MHz and 8 kbytes of memory.

The CLAS 4000 Configurable Logic Analysis System from the Gould Design and Test System Division (Cupertino, CA) can be configured with up to four 96-channel logic analyzers in a single chassis, all sharing a Macintosh II for control and display. The four logic analyzers can function independently, or can be combined into a single 384-channel instrument. Time stamping within each module lets displays from different modules be correlated for display.

LOGIC ANALYZERS

Each 96-channel logic analyzer module in the CLAS 4000 configuration supports a 200-MHz synchronous (state) data-capture rate and a 1-GHz asynchronous (timing) data-capture rate. As in the case of the Tektronix DAS 9200, the highest rates require that fewer channels be used. The 200-MHz rate supports 24 channels, 100 MHz supports 48 channels, and 50 MHz supports all 96 channels. Acquisition memory is 4 kbytes per channel.

The KLA/2, a high-end logic analyzer from Kontron Electronics (Mountain, View, CA), is scheduled for release this month. The KLA/2 can be configured with either one or two 96-channel, 100-MHz modules with 4 kbytes of acquisition memory. For higher acquisition speed, the 96-channel 100-MHz modules may be configured as 48-channel 200-MHz modules.

Also available from the company are 16-channel 1-GHz modules with 10 kbytes of acquisition memory, which may be configured as eight-channel 2-GHz modules. Kontron uses a plug-in, industrial-grade, AT-compatible computer for control and display.

What to look for

An examination of system specifications and brochures may help narrow the field for designers who are looking for a logic analyzer, but the specs won't provide all of the information that's necessary to make an intelligent choice. There's really only one criterion that must be satisfied when choosing an instrument, according to Ellis Goldberg, Gould's product marketing manager for logic analyzers.

"You need to be sure that the instrument will make the measurement that you want to make," says Goldberg. "Many of our customers present us with their toughest measurement problems before making a decision. That's an excellent way to evaluate a logic analyzer, particularly a high-end instrument."

Donald McCook, Kontron president, adds another caution. "The whole instrument must be designed to do the measurement job," he says. "In the past, you could switch probes between instruments, but that approach may not work with today's high-performance instruments. A probe that works well on a 20-MHz logic analyzer, for example, may be intrusive and produce erroneous results if it's used with a 100-MHz instrument."

In fact, probes are at the center of many changes in logic analyzers. Surface-mount technology makes it possible to pack tremendous power on a single board, but those tiny packages with their inaccessible pins make it difficult to connect a logic analyzer.

Most logic analyzer manufacturers provide special probes and assemblies that ease the job, but in some cases the designer may have to consider the problem before the board is built. For both design debugging and for production test, designers should bring critical signals to points where they can be ob-

Low-end and mid-range analyzers

Not everyone needs a Tektronix DAS 9200 or Gould CLAS 4000. Many, perhaps most, logic analyzer applications are handled quite adequately by less powerful and less expensive units. The mid-range and low end of the market offer a rich choice of products from which to choose.

As in the high end of the market, modularity is prevalent in the mid-range as well. The 1241 from Tektronix (Beaverton, OR), for example, is a self-contained, 26-lb unit that can be configured with up to 72 channels in nine-channel increments, has up to 2 kbytes of acquisition memory, and has sampling rates up to 100 MHz.

The Philips PM 3570, supported in the United States by John Fluke Mfg (Everett, WA), separates its state and timing channels into two different groups. The timing channels may be configured as 32 channels at 100 MHz, 16 channels at 200 MHz, or eight channels at 400 MHz. The 3570 provides up to 83 state channels and 32 timing channels for simultaneous acquisition of time-correlated data flow and high-speed hardware signals.

Another model, the PM 3655, is a 100-MHz, 96-channel instrument with its own built-in MS-DOS computer. Like the PM 3570, it's a modular instrument configurable in 24-channel increments from 24 to 96 channels. Acquisition memory is 2 kbytes per channel.

The HP 16500A from Hewlett-Packard's electronic design division (Fort Collins, CO) has five card slots that can be configured in a number of ways by plugging in cards containing instruments such as logic analyzers, digitizing oscilloscopes, 1-GHz timing modules, and pattern generators. Each logic analyzer module is equivalent to the 80-channel HP 1650A, thus providing up to 400 channels if all five slots are filled with logic analyzer modules. The 1-GHz timing module contains 16 channels per card, with a maximum of 80 channels per system. The pattern generator contains 48 channels per card, and the digitizing oscilloscope has two channels per card.

The PLA/2 mid-range unit from Kontron Electronics (Mountain View, CA) is based on a new custom chip, the GPDA (general-purpose data acquisition). The PLA/2 is a modular, 50-MHz state and 100-MHz timing analyzer that may be configured with up to 192 channels in 48-channel increments. Acquisition memory is 4 kbytes per channel. By halving the number of channels, timing analysis can be increased to 200 MHz.

Prices for most mid-range systems run from less than $10,000 in minimum configurations to more than $30,000 in maximum configurations.

The low end

Low-end logic analyzers are the tools of choice for engineers who are building systems using smaller 8- and 16-bit microprocessors running at under 20 MHz. Prices for these systems start at under $5,000, and is well under $10,000 for maximum configurations.

HP's two low-end offerings are the 32-channel HP 1651A and the 80-channel HP 1650A. Both the HP 1651A and the HP 1650A support state analysis at 25 MHz across all channels, timing analysis at 100 MHz across all channels, and a 1-kbyte capture memory per channel. Both are self-contained transportable units, weighing in at 22 lb each.

Tektronix has extended analyzer modularity into the low end of the market with its model 1230, a self-contained, 16-channel, 25-MHz instrument that can be expanded in 16-channel increments to a maximum of 64 channels. Acquisition memory is 2 kbytes behind each channel. Each 16-channel module works from a separate time base, allowing different parts of a circuit to be observed using a local timing reference for each different part.

Embedded systems designers have already used **CrossCode C** *in over 413 different applications.*

CrossCode C comes with four powerful tools to help you program your 68000-based ROMable applications

From C source to final object, each tool takes you one step closer to your finished ROMable design

CrossCode C is designed specifically to help you write ROMable code for all members of the Motorola 68000 family. Four powerful tools take you from C source to object code:

1. COMPILER: To get truly ROMable code, you have to start with a truly ROMable compiler. Here are three **CrossCode C** features that you won't find in any ordinary C compiler:

- Compiler output code is split into five independent memory sections that you can assign into ROM or RAM as you please.
- You can optimize the code for your application because *you* control the sizes of data types. For example, you can optimize for speed by using two byte *ints*, or get maximum versatility by using four byte *ints*.
- You can easily write assembly language routines that call C functions and vice versa, because the compiler uses simple, well documented parameter passing conventions.

2. ASSEMBLER: CrossCode C comes with a Motorola-style assembler that has all the features that assembly language programmers require. In fact, you could write your whole application with it:

- The assembler features an advanced macro language, conditional assembly, "include" files, and an unlimited size symbol table.
- Detailed cross references show you where you've defined and referenced your symbols.
- After a link, you can actually convert your "relocatable" assembler listings into "absolute" listings that contain absolute addresses and fully linked object code.

3. LINKER: The **CrossCode C** linker is designed to handle truly huge loads. There are no limits on the number of symbols in your load or on the size of your output file. And you can always count on full 32 bit target addressability, because the linker operates comfortably in the highest ranges of the 68030's address space.

4. DOWNLOADER: CrossCode C comes with a *downloader* that puts you in touch with all EPROM programmers and emulators. It can convert your load into Motorola S-Records, Intel Hex, Tek Hex, Extended Tek Hex, and Data I/O ASCII Hex. You can also produce a binary image and convert that image into any format you might want. In all formats, bytes can be split into EPROMs for an 8, 16, or 32 bit data bus.

Why Wait

Once you start using **CrossCode C**, you may just wonder how you ever got the job done before! It's available under MS-DOS for just $1595, and it runs on all IBM PCs and compatibles (640K memory and hard disk are required). Also available under UNIX, XENIX, and VMS.

CALL TODAY for more information:

1-800-448-7733
(ask for extension **5002**)

Outside the United States, please dial

PHONE: **1-312-971-8170**
FAX: **1-312-971-8513**

SOFTWARE DEVELOPMENT SYSTEMS, INC.
DEPARTMENT 52
4248 BELLE AIRE LANE
DOWNERS GROVE, ILLINOIS 60515 USA

CrossCode™ is a trademark of SOFTWARE DEVELOPMENT SYSTEMS, INC. MS-DOS® is a registered trademark of Microsoft. UNIX® is a registered trademark of AT&T. XENIX® is a registered trademark of Microsoft.

WORLD CLASS QUALITY

Now you can get world famous Wren™ disk drives, the fastest, most reliable 5¼" drives, direct from your local distributor.

Designed for Original Equipment Manufacturers (OEM), Wren drives are used in many of the most popular PC's and workstations. They set the high-performance standards for 5¼" drives.

With capacities from 50 to 766 MBs and industry standard ST506, SCSI, ESDI and PC-AT interfaces, there is a Wren solution for your needs.

Quality you can count on.
Imprimis is the world's leader in 5¼" high-performance, high capacity disk drives. The complete line of Wren drives feature field proven MTBF in excess of their 40,000 hr. specification.

PC-AT compatible
Wren PC-AT models offer capacities from 65 to 380 MB in a half-height form factor. A true 1:1 interleave provides faster data access. Wren provides the ultimate price/performance for AT-compatible microcomputer users.

Easy IBM PC, XT and AT installation
Several models come in "kits." Kits contain everything you'll need to install the drive, plus easy-to-follow instructions. Installation takes less than 30 minutes. A diskette-based software routine prepares the drive for use. Kits are also available for Apple Macintosh and Novel LAN too.

Factory certified RLL
All RLL models are certified at the factory using RLL controllers so you can be sure they work on your system.

Model	Capacity (Mbytes) (UF)/(F)	Interface	Avg. Seek (ms)	Transfe Rate (MHz)
WREN VI	766/676	*SCSI	15.5	15
	766/676	ESDI	15.5	15
WREN VI H/H	383/338	*SCSI	15.0	18-20
	383/338	ESDI	16.0	15
	383/338	AT	16.0	15
	274/242	AT	16.0	15
	164/145	AT	16.0	15
WREN V	702/613	*SCSI	16.5	12-16
	385/339	*SCSI	10.7	15-16
	442/390	ESDI	16.0	10
	383/338	ESDI	14.5	10
WREN V H/H	209/183	*SCSI	18.0	9-15
WREN IV	376/330	*SCSI	17.5	10-15
	350/307	*SCSI	16.5	9-15
WREN III	182/160	*SCSI	16.5	10
	182/160	ESDI	16.5	10
WREN III H/H	106/94	*SCSI	18.0	10
	106/94	ESDI	18.0	10
WREN II	135/115	RLL	28.0	7.5
	96/80	ST506	28.0	5
	86/71	ST506	28.0	5
	85/71	ST506	28.0	5
	86/71	ESDI	28.0	5
WREN II H/H	81/74	AT	28.0	7.5
	74/65	AT	28.0	7.5
	77/65	RLL	28.0	7.5
	51/41	ST506	28.0	5

H/H – Half Height Models. **(UF)** – Unformatted Capacity. **(F)** – Formatted Capacity at 512 Bytes/Sector. All WREN Disk Drives have 40,000 hours MTBF. *Apple Mac compatible interface available.

LOCAL AVAIL-ABILITY.

Get the high-performance edge... Wren disk drives from Imprimis.

Authorized Imprimis Distributor Sales Offices

Alabama
Huntsville
Arrow 205-837-6955
Anthem 205-883-0829

Arizona
Phoenix
Arrow 602-437-0750
Tempe
Anthem 602-966-6600

California
Chatsworth
Anthem 818-700-1000
Irvine
Anthem 714-768-4444
Los Angeles
Arrow 818-701-7500
Hayward
Arrow 415-489-5371
Orange County
Arrow 714-838-5422
Rocklin
Anthem 916-624-9744
San Diego
Arrow 619-565-4800
Anthem 619-453-9005
San Francisco
Arrow 408-745-6600
San Jose
Anthem 408-453-1200
Tustin
Arrow 714-838-1233

Colorado
Denver
Arrow 303-790-4444
Anthem 303-790-4500

Connecticut
Meriden
Anthem 203-237-2282
Wallingford
Arrow 203-265-7741

Florida
Clearwater
Anthem 813-797-2900
Ft. Lauderdale
Arrow 305-429-8200
Orlando
Arrow 407-323-0252

Georgia
Atlanta
Arrow 404-449-8252
Norcross
Anthem 404-381-0866
Arrow 404-449-8252

Illinois
Chicago
Arrow 312-250-0500
Elk Grove
Anthem 312-640-6066

Indiana
Indianapolis
Arrow 317-243-9353

Iowa
Cedar Rapids
Arrow 319-395-7230

Kansas
Kansas City
Arrow 913-541-9542
Lenexa
Anthem 913-599-1528

Maryland
Baltimore
Arrow 301-995-6002
Columbia
Anthem 301-995-6640

Massachusetts
Boston
Arrow 508-935-5134
Wilmington
Anthem 508-657-5170
Arrow 617-658-7100

Michigan
Detroit
Arrow 313-971-8220
Grand Rapids
Arrow 616-243-0912

Minnesota
Eden Prairie
Anthem 612-944-5454
Minneapolis
Arrow 612-830-1800

Missouri
St. Louis
Arrow 314-567-6888

New Hampshire
Manchester
Arrow 603-668-6968

New Jersey
Fairfield
Anthem 201-227-7960
Marlton
Arrow 609-596-8000
Parsippany
Arrow 201-538-0900

New Mexico
Albuquerque
Arrow 505-243-4566

New York
Hauppauge
Anthem 516-273-1660
Long Island
Arrow 516-231-1000
Melville
Arrow 516-391-1847
Rochester
Arrow 716-427-0300

North Carolina
Raleigh
Arrow 919-876-3132
Winston-Salem
Arrow 919-725-8711

Ohio
Cleveland
Arrow 216-248-3990
Columbus
Arrow 614-436-0928
Dayton
Arrow 513-435-5563
Worthington
Anthem 614-888-8707

Oklahoma
Tulsa
Arrow 918-252-7537

Oregon
Beaverton
Anthem 503-643-1114
Portland
Arrow 503-645-6456

Pennsylvania
Horsham
Anthem 215-443-5150
Philadelphia
Arrow 215-928-1800
Pittsburgh
Arrow 412-856-7000

Texas
Austin
Arrow 512-835-4180
Dallas
Arrow 214-380-6464
Houston
Arrow 713-530-4700
Richardson
Anthem 214-238-7100

Utah
Salt Lake City
Anthem 801-973-8555
Arrow 801-973-6913

Washington
Redmond
Anthem 206-881-0850
Seattle
Arrow 206-575-4420

Wisconsin
Milwaukee
Arrow 414-792-0150

Puerto Rico
Hato Rey
Arrow 809-764-4384

Canada
Montreal
Arrow 514-735-5511
Ottawa
Arrow 613-226-6903
Quebec
Arrow 418-687-4231
Toronto
Arrow 416-672-7769
Vancouver
Arrow 604-291-2986

IMPRIMIS
A Subsidiary of Control Data

CIRCLE NO. 35

served. And that means bringing the signals to points where probes can be attached.

Connecting several hundred probes to a system can be a problem in itself. Once the probes are connected, it's difficult to move them to change measurements. There are some logic analyzers that provide a way around that problem, particularly in systems where selecting higher sampling rates leaves some probes inactive.

One such instrument is the Gould CLAS 4000. When it's working at 200 MHz, only 24 channels on each module may be active. The user, however, can connect all 96 probes, and then select the active probes according to the measurements that will be performed. Using an Apple Macintosh II as its controller, the CLAS 4000 lets users switch active channel probes from the Macintosh keyboard, thus reducing the need to open the target chassis and move probes.

■ Emulation moves to software

Like logic analyzers, emulators used to be hardware debugging tools, but that's no longer the case today. "Emulation used to be hardware-related because of the control that it provides over a chip," says Craig Fisk, OSD development tools product line manager at Intel (Hillsboro, OR). "But with the higher software content in today's designs, emulation is more a software development and integration tool than it is a hardware tool."

Software has grown in complexity over the years, and emulators have evolved to support that complexity. In particular, emulators today help debug applications software that is written in high-level languages as well as that written in assembly language.

"Because of their current role in software development, emulators have new, more sophisticated capabilities that are geared to complex embedded system development," says Dave Camp, product marketing manager at Hewlett-Packard's electronic design division (Colorado Springs, CO). "Both emulators and logic analyzers provide a disassembly capability that displays assembler mnemonics instead of 1s and 0s when monitoring a microprocessor bus, but with new higher-level debugging tools, you can trace high-level code execution as well."

Tracing the execution of a high-

level language requires access to the source code so that machine-language instruction sequences that are monitored by the emulator can be compared with the high-level source code.

The debugging tool, in this case an emulator and its host computer, can't regenerate high-level language source code the way that a disassembler can regenerate assembly language code. Instead, the debugging tool interleaves source statements with assembly-level code, thus determining the source of executable instructions.

Increasing speeds and performance in microprocessor technology have put emulation technology to the test. Higher target-system execution speeds mean that signal paths that run from the microprocessor socket to the emulation processor must be as short as possible, or emulation will encounter problems with the laws of physics.

The Gould CLAS 4000 can hold up to four 96-channel logic analyzer models that can function as four separate logic analyzers or as a single 384-channel instrument. An Apple Macintosh II provides control and output display.

Because of the physical limitations involved in signal path lengths that are imposed by emulation, one may question the continued viability of emulation. It seems that there must be a physical limit to emulation technology, because there must be some point at which it's no longer possible to plug an emulator into a system and run the system at its full operating speeds. But that's a problem to be dealt with in the future, many argue.

"We may get to a wall somewhere in the future, but we're not there yet," says Intel's Fisk. "Because we know more about how to design and build emulators today, going from 12 to 25 MHz was less painful than going from 6 to 10 MHz was a few years ago."

The key to keeping up with advancing technology in instrumentation is to use that technology. "New microprocessors, including RISC, are using technologies that are available to emulator makers, too," says Richard Jensen, vice-president of product development at Applied Microsystems (Redmond, WA). "But the emulator makers must be willing to build their own VLSI products to shorten signal paths and reduce bulk."

Higher clock speeds act almost like a magnet, drawing the emulator's emulation pod closer and closer to the target system's microprocessor socket to shorten signal paths. In some emulators, the emulation pod sits directly on the target system's microprocessor socket.

The emulation pod itself looks much different from how it looked a year or two ago because more functionality resides on the end of the probe. Hewlett-Packard calls it an "active probe design."

Even with emulation technology approaching a physical speed limit, technology won't be the limiting factor for emulation, according to Camp. "As microprocessor speeds go up, it becomes more costly to develop the emulation technology to support them," Camp says. "Cost will limit emulation's applicability more than technology will."

Proprietary or universal?

Logic analyzers tend to fall in neat categories, but that's not the case with emulators. For one thing, there are a lot more players in the emulator game. Most of the giants are represented, including Hewlett-Packard and Tektronix among the general-purpose instrumentation manufacturers.

A few companies, such as Applied Microsystems and Huntsville Microsystems (Huntsville, AL), concentrate their efforts on producing a broad line of emulator products. Still others, usually smaller companies, offer one or two emulators. Finally, some microprocessor manufacturers themselves offer emulation tools to support their own product lines, which is a source of debate in the industry.

One of the most difficult technical problems for emulation manufacturers is figuring out what's going on inside a microprocessor chip. With its on-board cache memory, a microprocessor may execute long sequences of instructions with little or no information on the chip's pins to tell an emulator what's happening inside the chip.

Because they have access to the microprocessor design and manufacturing processes, microprocessor manufacturers have an advantage over general instrument manufacturers when it comes to designing and building emulators that can look inside a microprocessor. At Intel, for example, design support tools make up an important part of the company's total product line, and the company manufactures special versions of its microprocessors just for emulation.

These special chips, called bond-out versions, are standard microprocessor chips with some additional pads that bring key signals out to external pins where the emulator can monitor or manipulate them. Bond-out chips allow Intel to provide emulation control that's not available when a standard microprocessor is used.

Intel's objective is to offer development support for a new microprocessor within three months of the chip's introduction. In-circuit emulators that are available from the company include the ICE-386, the ICE-376, the ICE-390 and ICE units for just about all other Intel microprocessors. Intel also provides source-level debuggers for assembly language and PL/M source code, plus Ada source-language debuggers for the 80960.

HP's Camp admits that bond-out chips can be an advantage in designing emulators, but he says that careful evaluation of the signals that do come out of a microprocessor, particularly the status signals, gives sufficient information about what's going on inside. He also contends that instrumentation manufacturers have advantages that aren't available to the chip builders.

"Using production chips isn't that much of a handicap, and our prod-

Specifications: the real design-verification problem

Thorough design verification of any complex system is an extremely challenging task that typically consumes more time and effort than any other phase of a development project. Recent improvements in simulations, simulation accelerators, timing and circuit analyzers, logic analyzers and the like have all helped ease the design-verification process. But they haven't solved the entire problem.

Despite the available verification tools and the enormous effort invested in verification of application-specific IC designs, more than half of all new ASIC designs fail to operate in the end system the first time. This is a severe indictment of today's design-verification tools and their use. The reason for the problem is clear. Today's tools focus on the elimination of errors introduced by the designer during the implementation process and do very little to verify that the design will do what it's intended to do.

Indeed, the implementation of the design itself is usually not the cause of design failures. What is at fault, and must be verified, is the design specification. Most design-verification failures result from correctly implementing an incorrect or insufficient specification. If the specification isn't perfect, if it doesn't cover every nuance of what the product must and must not do in every case, the product will be similarly flawed. Design verification against an imperfect specification can, at best, assure that we have implemented a flawed specification perfectly. This is worse than garbage in–garbage out: it's garbage in–perfect and expensive garbage out.

The real design-verification problem is twofold: first, a correct design specification must be created from which the implementation detail will be derived; and second, the implementation must match the specification. This is a staggering task for a complex design. How is a sufficient specification actually created? Hundreds of pages of specifications will rarely be error-free, self-consistent, unambiguous and complete. It's simply impractical, if not impossible, to create a sufficient specification, except in a few rare cases, such as when implementing a well-known algorithm. Furthermore, even when a perfect specification is created, today's design-verification technology doesn't provide adequate assurance that the implementation will match the specification.

How can designs work correctly the first time when it can't be known with enough precision what the design is supposed to do? They can't. The reality is that a prototype must be built, used in its actual application and modified until it does what it's supposed to do. The ultimate arbiter of correctness is always the end application or end user. Only in rare cases, such as when a semiconductor company builds a new chip for a totally new application, can the specification be changed to match the implementation.

Prototype key to verification

There are several recent design and verification technologies that can ease, but don't entirely solve, the design-verification problem. Synthesis tools offer the promise of efficient correct-by-construction implementation, but can only be as good as the specification from which they synthesize. High-level languages are a better, more precise, method to write a specification than the English language, but offer no assurance that what is written is what the end application requires.

Behavioral simulation of high-level specifications offers the opportunity to run test vectors on the design representation before a low-level detailed implementation. But how does one create vectors that can be guaranteed to test all the conditions that must be met for correct operation of the design?

To achieve complete design verification, and to do so efficiently, a running hardware prototype must be constructed as rapidly as possible. The prototype may be hardware that's specially constructed for design verification before committing to manufacturing, or it may be the first manufactured unit. Verification isn't complete, however, until confirmation that the hardware correctly performs in the end system or application.

The high cost and long lead time to develop ASIC prototypes has driven many designers to try complete design verification with simulation tools before prototype hardware. These attempts have been costly and haven't solved the problem because simulation is simply too slow. Simulating microseconds of real time can't begin to represent the range of activities that occur when a system executes operating system software or complex application software.

What has been lacking from the suite of design-verification tools is the ability to build a design prototype without the cost and time of hardware kludges, and without waiting for the first silicon, which can't be observed internally or modified by the designer. A new design-verification technology that fills these needs is now available.

Enabled by new reprogrammable gate arrays and advances in automatic place-and-route software, the RPM Emulation System from Quickturn Systems takes the design netlist and automatically implements a prototype that can operate in real time. In-Circuit Interface cables let the emulated design plug into the target system, similar to microprocessor in-circuit emulators. The result is the target system that can be run through all of its paces, thereby assuring complete design integrity before investing in silicon.

Because it's now fast and easy to get to a running system, it's no longer necessary to invest enormous energy trying to verify a design by creating huge sets of test vectors while trying to assure that the design matches a specification that itself is probably flawed. The right combination of good implementation-verification tools and hardware emulation will let the end application be the arbiter of correct functionality, while eliminating the impossible task of producing perfect specifications and perfect test programs. The result is a breakthrough in design efficiency and quality.

Phillip A. Kaufman, MSEE, president, Quickturn Systems

ucts have features and basic debugging capabilities similar to those offered by the chip manufacturers," says Camp. "We can use other technology, such as our logic-analysis capabilities, for emulation; and we can do it across our entire emulation product line, not just for the products from a single vendor."

HP covers a much broader range of microprocessors with its HP 64000-UX family of emulators. HP's tools also target embedded systems design, but support more than 40 microprocessors, including those from Intel, Motorola (Austin, TX), National Semiconductor (Santa Clara, CA) and Advanced Micro Devices (Austin, TX). HP also offers emulators for DSP chips from Texas Instruments (Dallas, TX).

Part of HP's strategy is to support partnerships with microprocessor vendors, which are receptive to working with independent tool makers in addition to, or instead of, making their own design support tools, according to Camp. "We're working with National Semiconductor to provide support for National's 32532 chip," he says. HP has a similar relationship with Advanced Micro Devices for the 29000 RISC processor.

Tektronix takes a different tack. The company provides a narrowly focused software-development environment for embedded systems applications based on the Motorola 68000 family and on MIL-STD 1750A microprocessors.

"We stay away from emulators as generic tools and instead offer integrated solutions," says Mike Mihalik, marketing and planning manager for Tektronix microprocessor development products. "We find that our customers—who are involved in defense, aerospace and communications—need tools that the whole design team can use."

"We find that our customers need tools that the whole design team can use."

—Mike Mihalik, Tektronix

The Textronix solution consists of an emulator, a high-level language compiler (C, Ada or Pascal) and a debugging tool called TekDB. Tektronix emulators are connected to a workstation that's then networked with other design team members' workstations via Ethernet.

TekDB provides the network control program and a source-level debugger. The debugger lets the software designer execute code using the emulator and debug it using the source code debugger rather than having to debug in assembly or machine language.

Applied Microsystems is another company that supports a broad range of microprocessors, from 8-bit microprocessors such as the Zilog Z80 to 32-bit microprocessors such as the Motorola 68030. Applied Microsystems' flagship product, the EL 3200, supports the Motorola 68030 at 33 MHz and is designed to work at up to 60 MHz.

Huntsville Microsystems offers about six basic emulators, with adapter boards that let those six support another six microprocessors. At the high end, Huntsville Microsystems supports the Motorola 68020 running at either 16 or 25 MHz. Also available is a source-level debugger that supports C, PL/M and Pascal applications.

Kontron Electronics takes the approach of combining its KSE5 32-bit emulator with a 112-channel logic analyzer in the same chassis. The KSE5 supports the Intel 80386 and the Motorola 68020, both at 25 MHz, and Kontron is developing modules for AMD's 29000 and the Motorola 68030. A source-level debugger that works with the KSE5 supports source code written in C and Pascal.

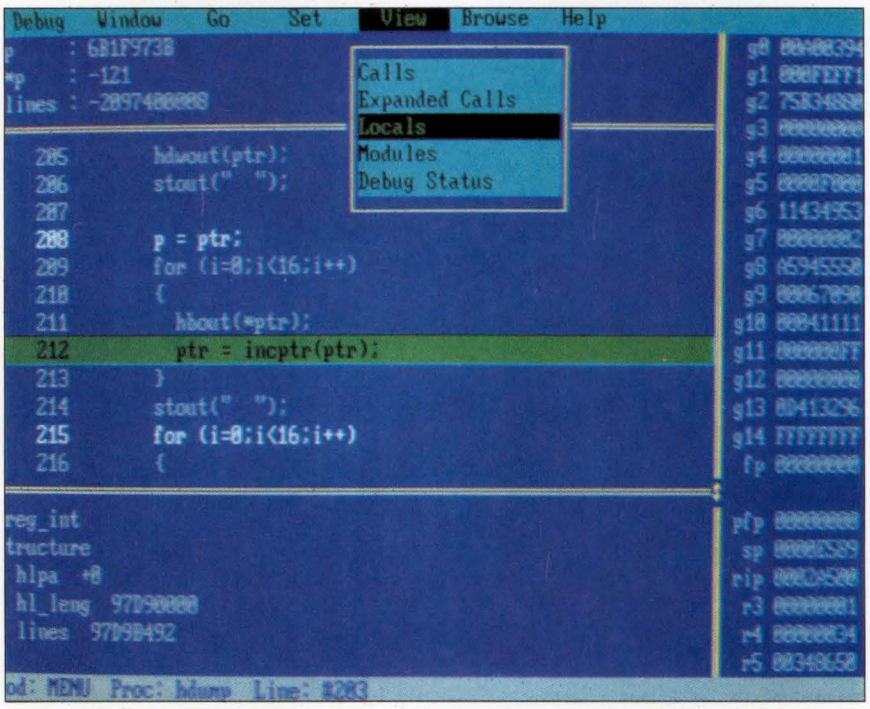

Microprocessor manufacturers often provide development tools for their own microprocessors. Intel, for example, provides various support tools for its RISC-based 80960 microprocessor, including the ICE-960KB in-circuit emulator, Ada and C compilers, a code execution board that plugs into an IBM PC AT, and the source-code debugger pictured here.

Niche microprocessors and beyond

While microprocessor emulation has moved to the software side of the house, emulation technology has expanded to encompass devices other than microprocessors. It's not too large a step from microprocessor emulation to DSP chip emulation, but it's quite a large step when it comes to emulating an application-specific IC in a target system—particularly when the ASIC isn't yet fabricated.

With its RPM Emulation System, Quickturn Systems (Mountain View, CA) applied emulation technology to the design of ASICs. RPM lets designers write their ASIC designs into programmable gate arrays that are located in the RPM system, and then test the ASIC design in a target system via an emulation cable. While this approach is limited to speeds of 10 MHz or so, it's orders of magnitude faster than a simulator, letting designers perform much more thorough verifications of their

DESIGN CHALLENGE: COMPUTER SECURITY

Our challenge: To develop a high security interconnect system for Digital Equipment Corporation.

When DEC faced the problem of designing interconnections for a top-security government computer system, they turned to Switchcraft for help.

A smart move. Switchcraft has a standing inventory of over 5,000 standard, miniature, and sub-miniature electronic components known for their rugged dependability and high tech performance. Chances are we would have a design ideal for DEC. Plus, with a complete model shop, the latest CAD/CAM capabilities, and fully automated manufacturing facilities, Switchcraft had the innovative engineering services package needed to do the job right.

Our Solution: By customizing the design of an existing miniature connector, Switchcraft engineers created the security-sensitive Tini Q-G connector and cable assembly built to pass the Tempest Test, the highest security requirement in the U.S. Fully shielded and filtered, the cable/connector assembly prevents any pick up of external signals and prohibits EMI or RFI emissions. Compact construction increases circuit density while positive latchlock system assures high integrity, vibration-resistant mating, and a positive ground connection between housings.

The next time your designs pose a challenge, call Switchcraft. We've got the answers.

Switchcraft

A **Raytheon** Company

5555 N. Elston Ave.
Chicago, IL 60630
(312) 792-2700

CIRCLE NO. 37

For instant response to your questions or for additional information, FAX us at (312) 792-2129, or call (312) 792-2700.

ASIC designs prior to going to fabrication.

RPM is a significant new development in emulation in that it supports something other than a microprocessor. What's more, it supports hardware design rather than software design—another break with the current emulation position in the design cycle.

RPM could well point the way to other revolutions in the emulation industry, but evolution rather than revolution is more likely. Such evolution will come about because a new microprocessor design will seldom be a single design in the future.

Automated design tools make it a lot easier to create different versions of a basic chip, and microprocessor manufacturers are doing just that. Intel's 80960, for example, targets embedded system designs. In addition, Intel created detuned versions of the 80386, the 80376, for embedded systems designs.

Others, such as TI, have developed DSP chips and special-purpose graphics chips. With many design tools available to develop special versions of microprocessors, we can expect more niche processors in the future, and we can expect support tools to support them.

Emulation's role with conventional microprocessors will continue to evolve, predicts Jensen of Applied Microsystems. Structured design methodology will be imposed on software designers, just as it's imposed on hardware designers today. The purpose will be optimization and verification of software, and emulation is the only technology that has the tools to support such design methodologies. □

WHY DO YOU REGULARLY READ COMPUTER DESIGN?

"Because COMPUTER DESIGN gives the overall systems picture... it doesn't just emphasize one component."

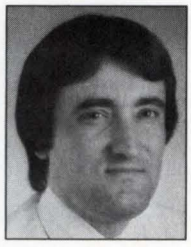

A. Karagiannis
Program Engineer
SUNDSTRAND CORP.

Mr. Karagiannis is a Program Engineer at SUNDSTRAND CORP., and has been a reader of COMPUTER DESIGN for 4 years. His engineering management responsibilities include system design, development and integration.

As a system design engineer, we'd like to know why you read and prefer COMPUTER DESIGN. Drop us a line at One Technology Park Drive, Westford, MA. 01886.

COMPUTER DESIGN®

THE FIRST MAGAZINE OF SYSTEM DESIGN, DEVELOPMENT AND INTEGRATION

For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate number on the Reader Inquiry Card.

Applied Microsystems (206) 882-5257	**Circle 267**
Advanced Micro Devices (512) 462-5257	**Circle 268**
Gould (408) 864-7749	**Circle 269**
Hewlett-Packard (719) 590-5823	**Circle 270**
Huntsville Microsystems (205) 881-6005	**Circle 271**
Intel (800) 874-6835	**Circle 272**
John Fluke Mfg (206) 356-5294	**Circle 273**
Kontron Electronics (415) 965-7020	**Circle 274**
National Semiconductor (408) 721-5000	**Circle 275**
Quickturn Systems (415) 967-3300, ext. 11	**Circle 276**
Tektronix (503) 629-1690	**Circle 277**
Texas Instruments (214) 995-3172	**Circle 278**

Our new 256Kx1 and 64Kx4 Mighty Memory CMOS SRAMs are ready to take on your heavy-weight advanced architecture design projects.

Perfect for high performance applications such as Military 1750, DSP, ALU, and other 32 bit micro designs, our Mighty Memories are available now in DIP and LCC. Both devices feature speeds to 35ns in standard and low power versions, MIL-STD-883 Class B screened.

EDI is also listed as an approved source on 256Kx1 and 64Kx4 DESC-SMD drawings.

High density. High speed. Low power. The new Mighty Memory 256K Static RAMs from EDI, the high performance military Static RAM leader.

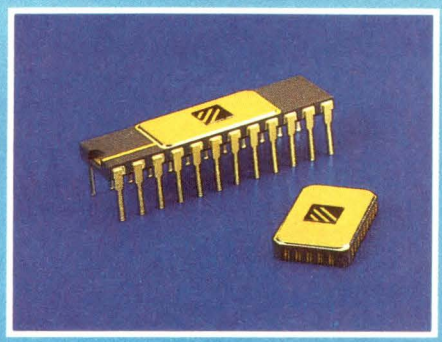

EDI's fast new 256K CMOS Static RAMs, available now in 256Kx1 and 64Kx4 with JEDEC standard pinouts, DIP and LCC.

CIRCLE NO. 38

EDI

The future...today.

ELECTRONIC DESIGNS INC. *42 South Street, Hopkinton, MA 01748, USA, (508) 435-2341, TFX: (508) 4356302, TELEX 948004*
ELECTRONIC DESIGNS EUROPE. *Shelley House, The Avenue, Lightwater, Surrey GU18 5RF, United Kingdom, Phone: 0276 72637 TELEX 858325*

Negotiating the obstacles to building military computers

David Lieberman
Senior Editor

Designers of military systems are doing some fancy footwork to maneuver around unsupportive suppliers and stringent DOD demands.

It's safe to say that the American public isn't willing to foot the bill for $700 hammers or $500 toilet seats anymore, no matter what the military wants to do with them. The military has recognized its desperate need for the economies of commercial standards, and this recognition is creating enormous opportunity for vendors of commercial bus boards, chasses, software kernels, languages, compilers and all the other system-level components involved in the open-architecture design arena.

For commercial OEMs and integrators, however, especially those making their first foray into military design, contracting or subcontracting for military projects is fraught with peril. The military designer is caught between what some claim to be unsupportive suppliers on the one side and, on the other side, what all agree to be the most demanding of customers. Tales of disaster are quite common, especially among those without the depth of commitment (and pockets) to go into the military market for the long haul.

Wanted: off-the-shelf commercial

From what was perhaps the most proprietary of computing environments, the Department of Defense has done a complete about-face and now endorses the procurement of commercial equipment and off-the-shelf parts whenever possible. "That's the implicit statement that the government is making more and more," says John Simpson, manager of the parallel system bus analyzer group at Planning Research Corp (Bellevue, NE), "and it's making key decisions based on the availability of hardware and software standards to support its requirements."

"Most military markets in the NATO world are moving rapidly away from proprietary black boxes into open system architectures," says Gary Dool, president of DY-4 Systems (Nepean, Ontario), "and they're doing it to reduce costs, shorten schedules and limit the risk of implementing new, complex systems."

Multibus I is very strong in military markets these days, especially among those in the Intel processor camp, but by all accounts the VMEbus is quickly becoming the military's bus of choice. "VME has been formally adopted by the German and French armies, and it's the de facto standard for U.S. vetronics applications," says Dool.

Simpson points out, though, that many contractors are bullish on Multibus II. "DOD contractors and projects are currently reponsible for roughly 60 percent of all domestic module-level sales," he says. And Digital Equipment Corp's Unibus and Q-bus enjoy widespread, though diminishing, support. "Q-bus is widely used in the intelligence community, in large part because of the good security certification of the VMS operating system," says Dean Tipa, principal systems engineer at Sanders Associates (Hudson, NH). "In many instances, the decision on whose system to use is based on who's got the best compiler for Ada, and the best today is for DEC's VAX systems."

According to Harry Campbell, project engineer at the Autonetics Marine Systems Division of Rockwell International (Anaheim, CA), the military's use of standard bus products is in transition. "We've seen our NTDS [Naval Tactical Data Systems] interface boards being purchased for the laboratory development environment for some time to take advantage of software, tools, storage capabilities and processing power of commercial equipment," he says, "but it's only recently that we're seeing purchases for use in deployable systems."

Campbell's group is a long-time supporter of the Q-bus, more recently of the IBM PC buses and, just recently, of VME.

Paul Graham, former vice-president of marketing and currently a consultant for Titan Sesco (Chatsworth, CA), described the appeal of commercial standards in a paper given at Buscon '89/West ("Bridging the Past, Present and Future in Military Microcomputer Design"). "The adoption of a well-known commercial architecture," he wrote, "accesses a broad base of engineering knowledge in all areas of the electronics community. This substantially reduces the developer's and end user's learning curve and associated costs. If the supplier has been willing to develop a well-thought-out line of products, usable over a wide range of applications, development costs are incurred only once, and production procurement economies can be passed along to subsequent customers without a further charge for NRE."

Variety of hardware grades

There are a variety of hardware grades that suit one military purpose or another: COTS (commercial off-the-shelf), ROTS (rugged), MOTS

(military), and GOTS (government), and there are different varieties and combinations of these.

"Before the Packard Commission, if you wanted to use ruggedized instead of mil-spec hardware, you had to justify it like crazy," explains Tipa. "Now you look at ruggedized first and have to justify any recommendations for full mil-spec. That's smart; you don't end up with $500 toilet seats in airplanes that way. If you require something to be mil-spec, it's going to cost $500 because of all the paperwork involved and all the tests it has to go through to prove it's mil-spec."

"I have no problem with $700 hammers," says Doug Patterson, manager of military products marketing at Radstone Technology (Montvale, NJ). "The government made the problem. With what they put the vendors through, they're not exactly your standard hammers anymore, and I'm surprised they don't cost even more."

"I have no problem with $700 hammers," says Doug Patterson, manager of military products marketing at Radstone Technology (Montvale, NJ). "The government made the problem. With what they put the vendors through, they're not exactly your standard hammers anymore, and I'm surprised they don't cost even more."

The epic amount of paperwork generated for military design work is legendary. To put some concrete figures on it, Simpson reports that Planning Research Corp (PRC) delivered 16,000 pages on a recent three-year Air Force contract and has delivered "well over 40,000 pages" on other projects. However, as Tipa points out, "That's not something the military saddled itself with. It's Congress that requires the documentation."

But the Packard Commission's push for off-the-shelf procurement has been misinterpreted to mean that full mil-spec equipment should be procured only as a last resort, says Graham. Again, from his paper: "C^3 [Command, Control and Communications] systems, often enclosed in shirt-sleeves environments, are logical candidates for COTS gear, but the deployed tactical systems required for equipment parity with the Eastern Bloc are not. A recently observed trend in our industry is a return to true military gear, after field failures with equipment of lesser environmental design integrity and inherent reliability have become more common."

Commercial pitfalls

The appeal of using commercial off-the-shelf products is great, however, most particularly for companies spending their own money to develop a military system. Then, too, commercial-grade products exist in far greater abundance than do ROTS and MOTS products. There's an insidious pitfall awaiting those looking to cut costs by designing with off-the-shelf commercial grade products, though. When the government buys into a solution based on a standard, Simpson explains, they don't require the contractor or integrating organization to do a proof of implementation on the underlying hardware base. "What they really want to see is proof of implementation to design specs with respect to the application," says Simpson.

The decision to prove a concept on commercial hardware, however, has come back to plague a number of contractors and subcontractors, and reports are rife of integrators who run into serious problems and delays when it comes time to implement their concept in a mil-spec system. "If a prime's subs need to prove a quick point without the long-term goal in mind, they generally put a bunch of commercial hardware together, then throw in hundreds or thousands of engineer hours in doing a first pass at system software just to show feasibility as part of the first phase of a full-scale engineering development," says Patterson.

"At the end of the day," he continues, "they have a system that performs the functions they want and satisfies the customer. 'Fine,' the military will say, 'now mil-spec it.' That throws a big crowbar into the middle of a deliverable system, and what we're seeing time and again, then, is a mad scamper."

Packaging habits that are common in the commercial environment are completely unacceptable in some military environments, for example, and the designer who has come to count on angular-mounted memories and daughter boards in the commercial arena will find himself at a sudden loss for real estate. Then, too, he may have designed with a commercial IC that has no military equivalent, so he's forced to redesign the circuit from scratch, most likely requiring multiple chips. The end result can be a 15-board system that, when military design constraints are factored in, requires 25 boards.

The ramifications of designing with commercial parts ripples throughout a system architecture. "Someone may have designed with a 33-MHz processor and a kernel that are running flat out at 0 to 50 or 70° C," says Patterson. "Now, make it military and put it in a −55 to 100° C environment and you'll get degraded performance, and you probably were able to get only a military version of the processor at 16 MHz, anyway. So, let's say you're now doing your real-time analysis at half speed. You'll probably have to put in another processor to distribute the load, which adds another card, maybe makes the box bigger, and creates more heat to be dissipated which, if not done right, raises internal temperature, pushes down mean time between failures, and drastically drops the reliability of the system."

Going plug-compatible

The use of "commercial-equivalent," rather than just "commercial," hardware for development provides what Campbell calls "the upward path without pain." A commercial equivalent of a mil-spec board that takes all the constraints of military design

What the military needs

The requirements of military computer systems are significantly more demanding than those of commercial systems and far more expensive to meet and maintain. Doing business with the Department of Defense forces the bus-based computer supplier to adopt a number of practices and procedures that are completely alien to the commercial electronics community.

Much of the difference in supplying computers to the military revolves around the long-term nature of DOD programs. Because these programs take such a long time to mature, and because the deployed systems remain in the field for so long, a supplier must be prepared to retain links with the past while planning for the future.

The supplier must be prepared to deliver and maintain products for at least 10 years, which often means maintaining an end-of-life supply of obsoleted components at its own expense. It also means actively maintaining documentation and logistics support for products that are reordered on a regular basis, no matter what their age.

For a number of reasons—reduced budgets, recent procurement scandals, and so forth—there's increasing attention in the military on purchasing economies as well as a new emphasis on upgradeable systems. For the DOD to improve its procurement practices and minimize its acquisition costs, it must buy only enough performance to meet today's requirements. Equally important, the systems purchased today must be able to accommodate future upgrades that boost performance while allowing as much of the original system as possible to be retained. New technology and higher performance cost money. The use of a 32-bit microcomputer to solve a requirements problem that could be handled with an 8- or 16-bit microcomputer will preclude the purchase of necessary quantities within a restrained program budget.

The DOD must also reduce the amount of nonrecurring engineering costs expended to develop custom solutions to program requirements, since every dollar spent on NRE is a dollar that's unavailable for production hardware—thus, the new emphasis on non-developmental items (NDI). The Packard Commission correctly recommended that off-the-shelf hardware be used whenever possible.

Since procurement costs go down significantly when large quantities of the same item are bought, the military has a need for hardware commonality. Using the same gear on multiple programs offers not only initial economies at procurement time, but also reduced logistic support costs, which can often exceed the cost of the hardware. New designs should, of course, retain bus compatibility.

Hardware commonality obligates the military computer supplier to maintain a consistent packaging approach with a long life span. And perhaps most obviously, software continuity across time is critical.

In fact, the military has a crying need for economical software development. For every $4 spent on embedded computers in defense systems, $3 are spent on software. This makes it imperative that software-development costs be contained, and that software be reusable as long as possible, and transportable as well. The imposition of Ada is meant to address the transportability issue.

Producing military computers isn't, of course, simply a matter of buying wide-temperature-range parts and putting them on stiffened circuit card assemblies. To comply with mil specs, every aspect of a supplier's design and manufacturing practices and the entire system of management controls must comply with the appropriate mil specs. It's common to underestimate the investment in time, money and resources to position a company as a true military computer supplier. It's also common to overlook the difficulties of trying to build commercial, rugged and mil-spec equipment all under the same roof.

To satisfy today's requirements, provide growth for the future, and provide continuing support for older programs, military suppliers must maintain a broad line of processors with a wide spread of cost and performance over a span of many years. This single requirement uniquely characterizes military suppliers and separates them from their commercial counterparts. Products or packaging approaches that a commercial supplier would discontinue without a second thought must be maintained by a military supplier to maintain the breadth of product necessary to provide the military with the most cost-effective solutions possible.

Paul Graham, BSME, consultant to Titan Sesco

into account lets the integrator "move all the way from development to deployment with complete hardware, software and mechanical compatibility," he says. The availability of boards in COTS, ROTS and MOTS grades from a single vendor makes an easy implementation, especially on the software side, more likely.

Patterson reports that many integrators have yet to understand and appreciate the difference between commercial and commercial-equivalent hardware. "But every once in awhile," he says, "we run into a Rockwell or Northrup or Hughes saying, 'Let's think long-term and use commercial equivalents of mil-spec hardware so when it comes time to change over, all we have to do is order mil-spec hardware and not change one spec of software.'"

Patterson reports another not-so-merry mixup occurring among contractors and subs involving lack of coordination among the different groups within a company that are working on the same military project. "You'll have the mechanical group, electrical group, digital group, maybe analog group, and software group working in tandem toward the end goal of the program, but each is expecting something different," he says. "For instance, you'll have mechanical people thinking about the space they've got to work with, say, in a tank turret, and that the box may have to have a curved corner in the back. Meanwhile, the digital people are thinking about getting as many functions as possi-

The military spends a great deal of money on computerized training equipment (such as this Multibus I-based flight simulator), which doesn't require mil-spec type environmental tolerances.

Frasca International

ble on a square board, the systems people are thinking about needing X functions on Y number of boards, and nobody has the long-term view in sight."

Serving the contractor better

According to Simpson, the vendors of hardware and software modules who will find success in serving the military market will be those who make the up-front investment to simplify the contractor's and subcontractor's job. Today, however, these vendors lack a system perspective, he says. "Especially in the military, where a premium is placed on reliability and interoperability, it's essential that a systems-level focus be exploited from start to finish."

Taking on a systems perspective, as Simpson sees it, would consist partially of providing the hooks and handles to simplify software integration. Military programs, Simpson laments, always seem to call for "something different" from what off-the-shelf software packages support, and frequently require the incorporation of code generated in a different development environment, most likely for an earlier program.

If suppliers of off-the-shelf operating systems, compilers, assemblers, linkers, and so forth would provide the hooks and handles with their products to bridge different operating environments, Simpson contends, contractors would literally flock to their doors.

"Any time that I can buy something and charge it to the government as another direct cost, it'll typically be billed at a much, much lower rate than the labor cost we'd have to charge the government if we have to build it ourselves," he says. "To stay competitive, we'd almost always prefer buying over building. Then when we deliver labor hours, it's to enhance, value add or extend the capabilities of a commercially available product. Any time you can design engineer hours out of a project, you've saved the government money and made yourself more competitive. And that goes across the full spectrum of engineering disciplines, from planning to documentation. If the software and hardware vendors would make an investment to serve the military market, they'd position their products to be bid and procured more frequently."

Simpson also believes that the absence of off-the-shelf hooks and handles tends to undermine the "standard" nature of standards. "The military has certain standard requirements that we try to meet with standard parts, but when all is said and done, systems wind up being vendor-unique," he says. "The module vendors need to recognize, for example, that the military has well-defined interface standards of long standing, though they're not in the commercial mainstream. If there were enough support in the vendor community for servicing these military interfaces, then I could go to a compiler vendor and get fully validated MIL-STD-1553 or NTDS extensions. Now, however, there's a Rockwell 1553, for example, a TRW 1553 and a Ford Aerospace 1553 since there's no one well-accepted, universally recognized 1553 that's a standard part of software products, so it has to be documented every time we use it."

Nonstandard standards

The ideal situation in military computing, offering the greatest vendor opportunities and military economies, would consist of commonality and consistency of procurement across the whole range of military systems. The different needs of different parts and branches of the military, the different computing histories of different branches of the service, plus the overabundance of procurement agencies, each with its own motivations and eccentricities, makes this a pipe dream. As Simpson reports, two programs may even require two completely different sets of documentation for what is basically the same machine. "Sometimes you have to rewrite documents just to keep in sync with program methodology," he says.

It's not always clear, therefore, what standard approach will meet with the greatest success across the widest range of military programs. And even when a certain standard is unequivocally dictated, it's possible to get exceptions. Nevertheless, according to Tom Radi, president of Software Systems Design (Claremont, CA), it's becoming more and more difficult to convince the DOD to waive the use of Ada.

It's been almost six years since the DOD made Ada the standard for mission-critical software. "It's taken all that time to develop usable compilers that support the policy and a management infrastructure in government that can effectively deal with the complex implementation issues of Ada," Graham says. Ada has,

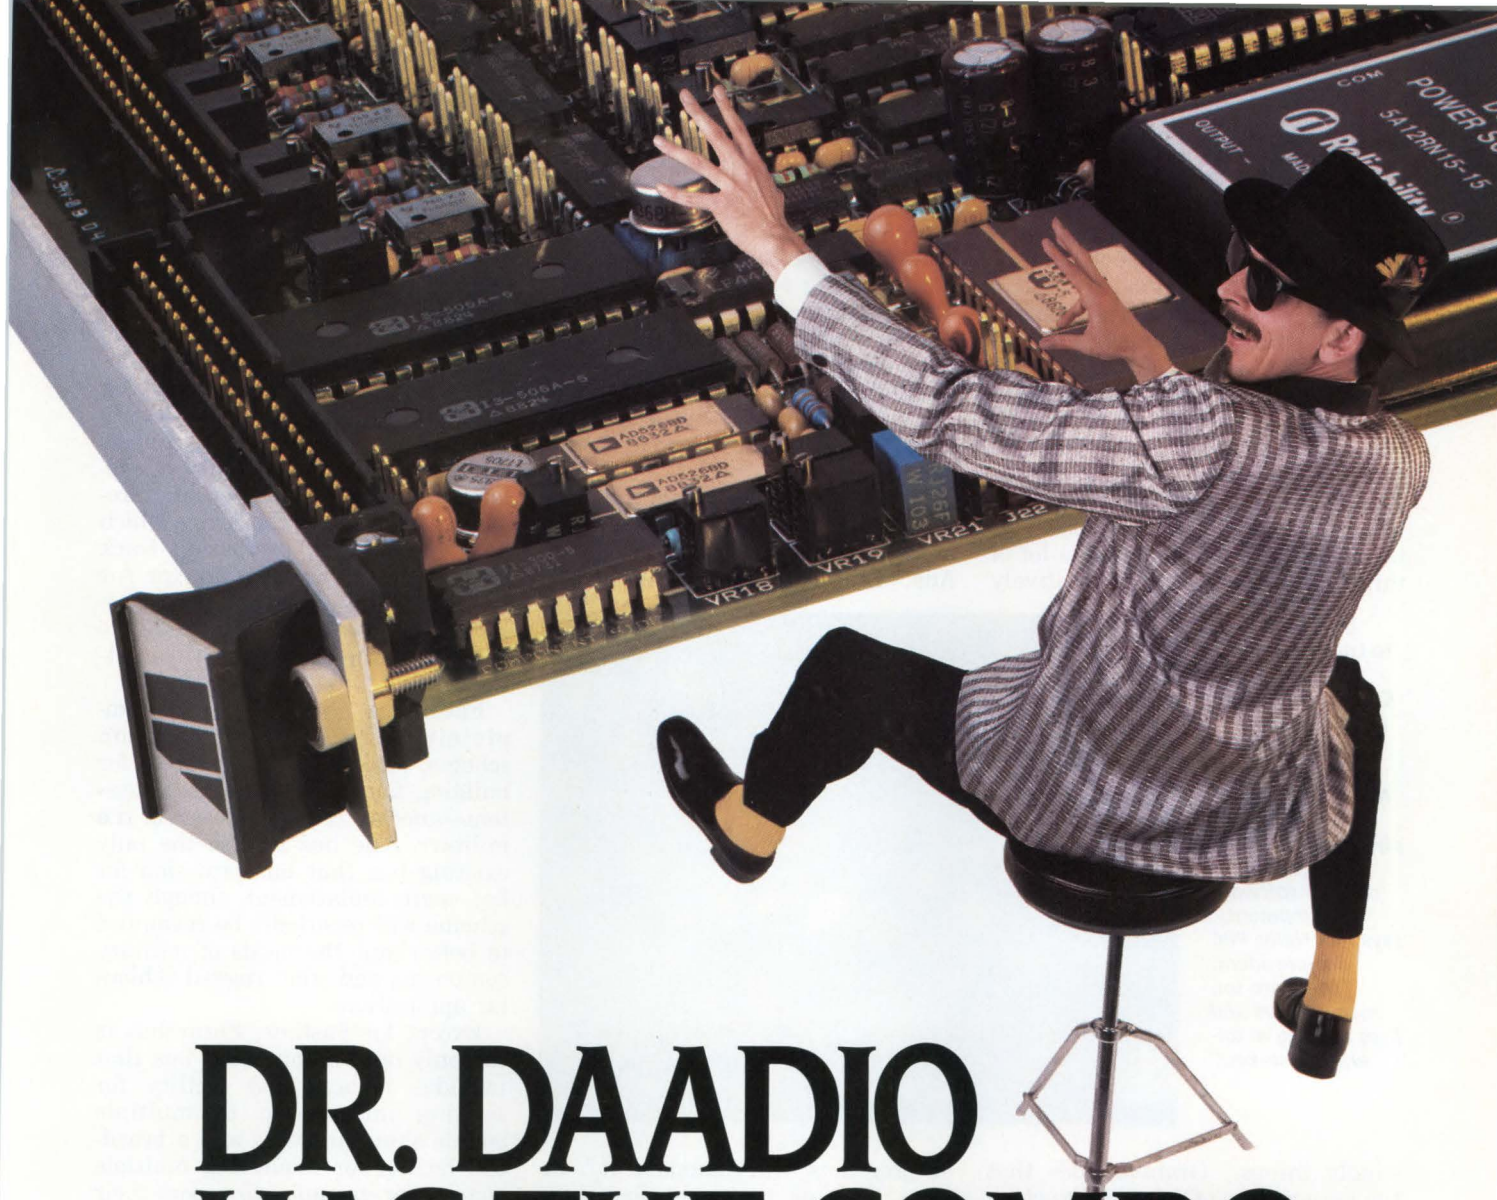

DR. DAADIO HAS THE COMBO

to cure your data acquisition blues!

D/A, A/D and PIO requirements in the same VMEbus application mean three separate boards, right? Not any more! The MD-DAADIO from MATRIX Corporation gives you everything you need in one combination board—and saves you space and money!

125 KHz A/D Rate!
The MD-DAADIO knocks your socks off with a unique pipelined A/D conversion method and ten software selectable gains. Also, the MD-DAADIO features external triggering for synchronous applications.

12-Bit, 8-channel D/A.
The MD-DAADIO gives you analog output any way you want it—you decide the voltage range and the mode. The MD-DAADIO even has current output capability and external scaling.

48 Lines Parallel I/O.
The MD-DAADIO offers six, eight-bit ports, each software configurable for data directionality. In addition, the first port is capable of interrupting the VMEbus on change-of-state or byte recognition.

The Combination Cure.
Why risk the health of your VME data acquisition solution? The MD- DAADIO gives you everything you need in one combination board and gives it to you at less cost.

Call now...and ask for the Doctor!

MATRIX CORPORATION

1203 New Hope Road
Raleigh, NC 27610
Phone: (919) 833-2000
FAX: (919) 833-2550

CIRCLE NO. 39

indeed, come a long way.

By all accounts, the early Ada compilers were inefficient, due primarily to the limitations of the language itself. "Up until recently, compiler validation quite frankly didn't mean all that much with regard to the real world," Graham says. "You could pass all the validation tests and yet find out that what you had was an unwieldy, cumbersome solution to a problem."

"Ada doesn't handle certain things like message queues well," says Patterson, "and it tends to take a lot of target code to handle some relatively simple things." Graham adds that the very nature of Ada, its structure and the way it services interrrupts and handles multitasking are quite inefficient for the real-time environments where much of military computing lies.

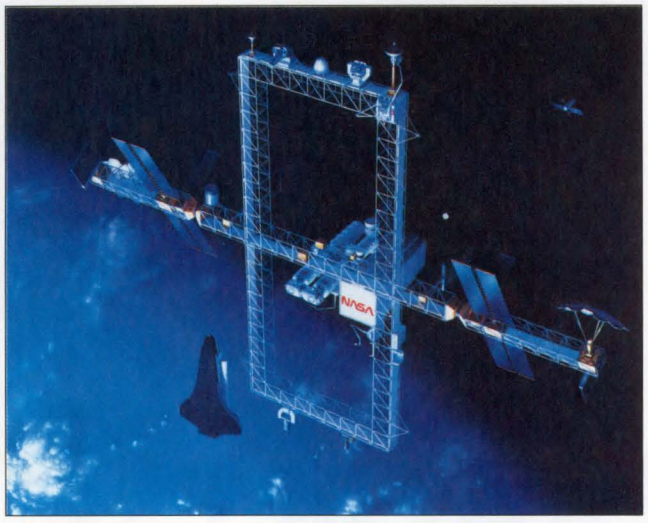

Under contract to Lockheed, Planning Research Corp is developing an Ada software-development environment for NASA's space station. "We're applying a disciplined shop floor mentality to software development," says Jack Nulty, PRC vice-president. "There are too many dollars and lives at stake to tolerate mistakes."

Consequently, the software vendors have over time developed their own macro library extensions that, while making Ada operate more efficiently, raise the specter of incompatibility. Although Patterson doesn't see these extensions as causing major problems, the purists do. "At some point," he says, "the Ada Joint Program Office might draw the line."

Futurebus finds its niche

Although the U.S. Navy has mandated Futurebus as its bus of choice for mission-critical computing in the 1990s, this doesn't mean that enormous markets for Futurebus system-level components will appear overnight. A number of the bus's architectural features that the Navy is keen on are actually uncompleted, though the IEEE 896.1 version of the spec, completed in 1987, contains all the necessary hooks and handles to implement them. Then, too, the military isn't going to swap out all of its computing equipment overnight.

"Any changes in naval computing will be evolutionary because of the huge installed base and tremendous investment in computers and other equipment that will be around for some years to come," predicts Campbell.

DY-4's Dool suspects that the military's use of Futurebus will evolve in much the same manner as its use of Ada. "Initially, only certain classes of programs will mandate it," he says. "But as the technology matures, more money is put into development, and the primes become more educated, it will be mandated on a wider selection of programs."

Many advantages

The appeal of Futurebus, the so-called futureproof bus, is multifold. Using backplane transceiver logic, which was designed specifically for driving a backplane bus, it solves the transmission-line problems that limit the speed of buses based on TTL. Futurebus doesn't specify any particular setup and hold times but, rather, bus transactions take place at the speeds the participating modules are capable of.

Furthermore, Futurebus sports what are commonly called "squeaky clean" protocols that are completely independent of the processor family or technology used to implement the system. "Other buses impose complicated byte ordering and justification rules on the bus interface hardware," says Paul Borrill, Futurebus committee chairman and staff scientist at National Semiconductor (Santa Clara, CA). "With its simplified byte-lane paths, Futurebus doesn't make any assumptions about how bytes ought to be broken up or rearranged."

The bus also contains extensive monitoring, diagnostic and error-detection facilities. "We have parity on absolutely everything," says Borrill. These features are extremely important for the maintainability of systems—a critical military requirement, especially for the Army, which has far less of a computing track record than either the Navy or Air Force. "Never overestimate the resources of the guy who's going to wind up fixing the thing," says Patterson.

Futurebus also features a completely distributed arbitration scheme, which is a real boon for building fault tolerance into a system—another critical concern of the military. The bus is also the only existing bus that has provision for hot spare replacement, though the scheme will reportedly be revamped to better suit the needs of military computing and other rugged vehicular applications.

Except for Fastbus, Futurebus is the only open-architecture bus that includes a broadcast facility for sending information to multiple boards simultaneously and a broadcall facility for requesting multiple boards, for example, to report their current status or the results of local diagnostic tests. It also provides comprehensive support for implementing a variety of caching methodologies within a single shared-memory multiprocessing system while maintaining cache coherency.

The bus also supports "third-party transactions," which permit, among other things, a master with the most current version of data in its cache to update other boards that have an earlier, stale copy of the data in their caches.

The basis of the Futurebus caching scheme is a cache state model commonly called the MOESI model, which is built around three primary attributes of cached data: it may be valid or invalid, shared or not shared, and owned or not owned, with ownership meaning responsibility for providing the latest version of cached data to another board that requests it from the system. Various combinations of these three attributes cede five cache states: M for modified (owned, exclusive), O for

Radstone supplies Norden Systems (US) and Marconi Command and Control Systems (UK) with bulk storage memory units developed for the MLRS fire control computer. On time. On budget.

MIL-Spec. VMEbus

The Sure-Fire Solution To Cost and Lead-Time Problems In Military Computer Applications

Now there is a new generation of Mil-Spec computer hardware. For systems integrators with tactical computer applications who need a sure-fire solution fast, Radstone has the answer — a complete range of field-proven Mil-Spec, 32-bit VME boards.

All you need to build tomorrow's military systems today, available off-the-shelf, 68020 processors, memory and I/O modules, Multiprocessor capability. On board diagnostics. Real-time Ada and C software support. A choice of Mil-qualified system enclosures.

This total capability is designed for projects where lower development costs, shorter and more predictable timescales are important. Developed by Radstone, one of the world's largest manufacturers of VME boards with leading edge design resources to match. Built by Radstone, in a facility approved to NATO AQAP1 standards. Supported by Radstone, with 30 year's experience of military systems applications.

Call us now for your copy of the new PMV 68 MIL-STD-VME systems guide. Let us show you how we put the sure-fire solution into action.

RADSTONE TECHNOLOGY

Radstone Technology Corporation
20 Craig Road, Montvale, NJ 07645

Call Toll-Free: (800) 368-2738
Eastern Region: (201) 391-2700
Western Region: (408) 727-4795

CIRCLE NO. 40

Built around three basic cache attributes, the Futurebus MOESI state model provides the foundation for developing complex shared-memory multiprocessing systems without fear of cache incoherency. Various combinations of the three attributes cede five cache states: M for modified (owned, exclusive), O for owned (owned, shared), E for exclusive (not shared, not owned), S for shared (shared, not owned), and I (invalid). If data in one board's cache is designated as owned, for instance, and another board requests the same data from main memory, the first board must intervene in the transaction to give the second board up-to-date data.

MOESI STATE MODEL

Characteristic	M	O	E	S	I
Data valid in shared memory?	?	?	yes	?	?
Data valid in other caches?	no	?	no	?	?
Cache required to intervene for memory?	yes	yes	no	no	no
Cache allowed to intervene for memory?	yes	yes	yes	no	no
Data valid in this cache?	yes	yes	yes	yes	?

owned (owned, shared), E for exclusive (not shared, not owned), S for shared (shared, not owned), and I (invalid).

Futurebus is also unique in its use of a tagged architecture: that is, it's really a 33-bit bus with an extra bit on every long word that's usable in a variety of ways. The extra bit is useful, for example, in Lisp architectures to define different categories of data structures and in security-intensive systems to flag data that should be made available only to certain users. By providing tags in an extra bit in this fashion, Borrill says, Futurebus maintains a "clean" software architecture.

"Futurebus was designed explicitly to be extensible," says Borrill. "The working group thought of 101 different ways the bus could be extended in the future, and incorporated the facilities for the bus to migrate gracefully without obsoleting older boards."

Layered bus protocols

The 896.1 spec is, in fact, only one layer of the evolving Futurebus architecture: the Logical Layer, which

Personal computers invade the military

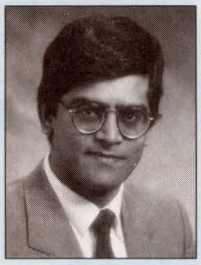

Two major trends are converging to propel personal computers into new applications for the military: increasing budgetary pressures from Congress and increasingly fast computing power for relatively low costs. Add to that the high comfort level today's graduating engineers have with PCs, and you have a formula for rapid acceptance of versatile, off-the-shelf commercial products in an area that was once solidly locked into costly, full-custom, long-lead-time gear.

In one area, commercial PCs are lowering the cost of military system development and emulation. Virtually all design engineers working for the government, prime contractors and their secondary vendors either have their own desktop PCs or have access to them, and most of those PCs have one or more vacant card slots available. By dropping in a low-cost military emulator or coprocessor board, it's possible to eliminate the high cost and computing overkill of, for example, an AN/UYK-44 military computer, a MIL-STD-1750A avionics computer, or a Navy Tactical Data Systems (NTDS) node processor.

In systems-development labs, there's no need for rad-hard gear when the job can be done by a low-cost platform that drives, receives and interprets MIL-STD signals. Engineers at the Naval Surface Warfare Center, for example, recently used $5,000 worth of PC equipment to simulate and stimulate an AN/SLQ-17 Electronic Counter Measures system, which would otherwise have cost 10 to 20 times as much. The PC is used to help develop ECM software, simulate radar bogeys and train operators for real-world scenarios.

In addition to lowering costs, off-the-shelf PCs also help speed systems development and delivery. The procurement cycle for a $100,000 MIL-STD computer can take six to nine months or more, and it might take three months just to get a purchase order to the vendor. A personal computer such as the Zenith Z-248, on the other hand, costs just $3,500, and is already approved for off-the-shelf military delivery.

Useful for diagnostics

PCs are also invading the military in diagnostics, training and actual post-deployment service. There are many noncritical applications on ships and land-based installations where commercial products are acceptable for use, even though they are neither rad-hard nor rated for shock, vibration or the full mil-spec temperature range.

Commercial operating ranges are acceptable, for example, for equipment used to repair parallel-bus shipboard and embedded avionics systems. Defective computing nodes can be diagnosed by simply disconnecting and replacing them with PCs fitted with drop-in emulators and coprocessor boards. PCs are also being used for weapon systems training at land-based test sites, as well as aboard aircraft carriers deployed in the fleet.

Other trends contribute to the relaxation of stringent rad-hard and environmentally rigorous military specs. In areas where it would be impossible for people to work the equipment, such as in a combat information center above 70° C or underwater, it doesn't make much sense for the military to spend premium dollars for, say, keyboards that can work at 125° C.

All of these factors point the direction for personal computers to continue to expand their successful invasion of the military.

Rahim Sabadia, president, Sabtech Industries

consists of the low-level physical protocols and the basic logical protocols. The 896.2, called the System Layer, has been in definition for about three years and deals with higher-level architectural details such as a message-passing model and cache-coherency methodologies. Borrill foresees that three more layers of Futurebus will be developed, though the effort has yet to gain IEEE approval.

The 893.3 Physical Layer of Futurebus is intended to detail the bus's electrical specifications, board sizes, connector types, and a new hot spare replacement scheme; the 896.4 layer will be a system-configuration guide covering such issues as maintainability, fault tolerance and data security; and 896.5 will be a bridge document that describes the relationships between other buses (VMEbus and Multibus II) and Futurebus for both mixed-bus and migratory designs.

It's likely that early Futurebus systems will retain another bus as an I/O structure to take advantage of already available hardware.

Now that the Futurebus is destined to make the transition from a more or less academic bus structure to a commercial reality, the 1987 896.1 document is being modified and extended to better suit the bus for the real world. Among the extensions, a pipelined transfer mode will be specified to make the bus faster, and data width extensions to 64 and 128 bits will be specified. Borrill reports that the committee will

Documentation: let's share the burden

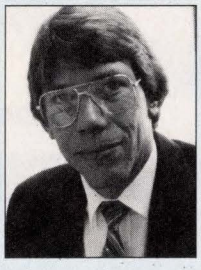

Few vendors of commercial single-board computer products realize just how much time and effort military contractors invest to research and represent the vendors' technology to the government. Nor do most realize the expense of researching, analyzing, evaluating, specifying, testing and documenting an embedded system design based on more than two or three vendors' products.

DOD contractors, especially integrators, spend a great deal of time acquiring various degrees of interoperability data to analyze and evaluate boards as system-level components. Most vendors' hardware and software documentation does a reasonably good job of describing how to utilize the local resources embedded in the product. But very little vendor documentation, if any at all, provides system-level design information for the product.

Therein lie many opportunities: for commercial board vendors to carve out a new market niche; for integrators of military systems to become more cost-competitive; and for the government to use its procurement dollars more effectively. Time is money, and usable technical information is the first step in all product design wins.

Vendors with strategic marketing initiatives aligned with the DOD may find they can create new products by merely repackaging their existing internal technical information. By revamping their documentation under the auspices of DOD-STD-2167A—the Defense System Software Development Standard for Mission-Critical Computer Resources (MCCR)—they can document valuable system-level integration characteristics of their products using a single publication format that will be applicable to the vast military marketplace.

Historically, the DOD has bought a great deal of custom documentation in conjunction with the development of new computer systems, especially in the case of systems developed using a life cycle-oriented, cradle-to-grave methodology. A study recently conducted by a prominent industry trade group of defense contractors indicates that its members typically spend at least 40 percent of their software-development contract budgets for researching, drafting, reviewing, revising, and baselining various specifications and reports to fulfill life-cycle document requirements prescribed by the government.

With DOD-STD-2167A, the 40 percent budgetary impact for life-cycle documentation inside military software-development contracts is expected to increase. Documentation requirements inside current and future MCCR procurements will shrink the dollars available to buy equipment. In some cases, the ability of the system-development team to field the necessary documentation may become a gating factor in product recommendations to the government.

Vendors of computer-aided software engineering tools may be the first group to actively realize opportunity from this dilemma. Some CASE vendors are starting to offer 2167 document generators as an extension to their products.

Vendors will that find the 2167A standard establishes uniform requirements for software development that are applicable throughout the system life cycle (see "Defining the software life cycle," *Computer Design*, Jan. 1, 1989, p 72). The latest revision defines a common development methodology and a suite of 17 standard data item descriptions (DIDs).

Although most of these DIDs are software-oriented, they do require that system hardware components be documented. In general, all aspects of the system are documented using a structured top-down approach to provide a means for establishing, evaluating, and maintaining quality in the hardware, software, and documentation delivered to the government.

Clearly, many DOD contractors would gladly invest in high-quality, system-level product information to defer some of the expense of acquiring it themselves, especially if this information were in a format that would allow it to be directly incorporated into their contracts. Vendors seeking to position system products based on open bus standards in military contracts should seriously consider stepping up to a 2167A library that documents how the system was developed or, at least, should offer 2167A documentation describing each system service that can be solicited through value-added integration.

Those interested in obtaining 2167A guidelines should contact Commander, Space and Naval Warfare Systems Command, ATTN: SPAWAR - 3212, Department of Defense, Washington, DC 20363-5100.

John Simpson, senior systems programmer, Planning Research Corp

Can you find the Ap

Anyone looking for the best Electronic Design Automation (EDA) solutions would undoubtedly consider Mentor Graphics, Racal-Redac, Cadence or VLSI Technology. Four companies whose superior applications and powerful capabilities place them squarely at the forefront of their markets.

Likewise, when these four industry leaders went looking for the best hardware on which to develop and run their software, they chose Apollo. Discovering that our Series 3500™ Series 4500™ and Series 10000™ workstations provided ideal platforms for everything from IC layout to full system design.

What appealed to them were the same things that appeal to the thousands of engineers who've come to depend on our machines. They liked our open network-

ollo in this picture?

ing and advanced UNIX® operating system. Qualities that make it easy for developers to work together productively, allowing them to access all the information and processing power needed to get a job done.

They appreciated a complete family of compatible workstations that starts as low as $5490. Assuring solutions to applications as familiar as simulation or as formidable as microwave design.

And they admired system administration features so efficient that even a network with hundreds of users can be managed by just one person.

As you can see, if you're looking for the perfect EDA workstation, you have lots of options. And fortunately all of them are built by Apollo.

apollo

CIRCLE NO. 51

For more information, call 1-800-323-1846 (in MA: 1-800-847-1011) or write Apollo Computer Inc., 270 Billerica Road, Chelmsford, MA 01824.

Series 3500, Series 4500 and Series 10000 are trademarks of Apollo Computer Inc. UNIX is a registered trademark of AT&T.

Military leads the way in software development

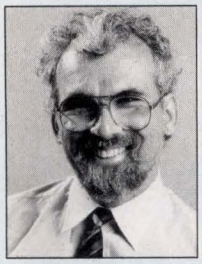

The defense industry is the strongest proponent of state-of-the-art software-development techniques. That's no surprise in light of the criticality of military systems and the disastrous consequences of failure. The military not only has a huge and growing investment in complex embedded software systems, but also has long experience in maintaining them. The maintenance and update of this large and complex software investment isn't a trivial task.

The success of a military project depends on the successful combination of routines written by many hands into an embedded software system that must be integrated with equally complex hardware. Through all this, development costs and delivery schedules must be tightly controlled.

The best way to control costs and development time in complex systems is to ensure that the design conforms to specification. This presumes that one is able to regularly check the design against the spec during development, detecting discrepancies as soon as possible. To do this, the progress of the design must be easily tracked and well documented. Should the specification change, it's important that design changes in all phases of development be monitored and traced.

Another way to save time in program development is to reuse as much code as possible by employing a structured, object-oriented approach. The promise of reusable software lies in the fact that 30 to 70 percent of most software programs contain code that is very similar to code that has been written before, perhaps dozens of times. If developers can eliminate this duplication of effort, enormous gains in productivity are possible. Reusable software also brings with it the benefits of maturity and robustness.

Standardization has provided a solution for part of the software-development problem. Most military projects now mandate Ada as the programming language. Ada is especially well suited to object-oriented design and, hence, code reusability. Standards have been established for the language, and commercially available compilers are tested for compliance. With a single language in use, the issue of incompatibility in software development and maintenance will all but disappear.

Through its acceptance and verification procedures, the military has effectively standardized the development process, tying it to a software life-cycle model that is now being used in most computer-aided software engineering (CASE) tools. The model dictates the resolution of the design problem to a certain level of detail at each phase of the life cycle and demands specific documentation to support it. Each phase of development becomes self-documenting, ensuring a paper trace back to the original specification. The combination of a standard development process and a standard software language expedites the development process as information becomes available across all development phases.

The use of object-oriented design to optimize code reusability is spreading throughout the commercial, as well as military, environment. We see greater and greater integration of CASE and software-development tools in conformity with the software life-cycle model adopted by the military. The newest generation of CASE products can provide integrated support over several life-cycle phases. A single CASE product may, for example, support software requirements specification, software architectural design and detailed design, and may provide ties between these phases for ease and consistency in moving between design stages.

Customer concerns about the reliability of software and its ability to be successfully integrated into target hardware makes it important to reduce design risk. This is especially true in embedded systems, where hardware/software integration may be particularly difficult. Knowledge of the performance characteristics of the target system can now be comprehended in the early phases of software development and modeled early in the design process so that the integration of final hardware and software bears no unwelcome surprises.

David Kalinsky, director of software engineering technologies, Ready Systems

dump the bus's single address mode, which doesn't suit the operation of current microprocessors and unnecessarily complicates the bus protocol, as well as the reflection operation in the cache-coherency protocols which, Borrill says, causes difficulties in implementing the spec.

Future systems

Perhaps one of the greatest potential benefits of making the move into military design work is the possibility of spinoffs into the commercial world. Though it can be an extremely expensive proposition to develop and support a military computer, vendors typically retain ownership of the fruits of their labors.

"Many commercial companies have been able to leverage the development they've done for the military," says Tipa. "The work with Ada, for instance, has given them a leg up on other development, and the same will be true of Futurebus. There will be Futurebus systems that don't have to be full mil-spec, and the government will look for people to develop products that can be sold beyond the government, thereby spreading the base and cost and reducing the per-piece cost to government. I think a lot more things have been spinoffs of military development and the space program than people realize." □

For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate number on the Reader Inquiry Card.

DY-4
(613) 596-9911 ..... **Circle 279**
National Semiconductor
(408) 721-5170 ..... **Circle 280**
Planning Research Corp
(402) 293-3900 ..... **Circle 281**
Radstone Technology
(201) 391-2700 ..... **Circle 282**
Ready Systems
(214) 661-9526 ..... **Circle 283**
Rockwell International
(714) 762-8111, ext 1476 ..... **Circle 284**
Sabtech Industries
(714) 524-3299 ..... **Circle 285**
Sanders Associates
(603) 885-4321 ..... **Circle 286**
Software Systems Design
(714) 625-6147 ..... **Circle 287**
Titan Sesco
(818) 709-7100, ext 117 ..... **Circle 288**

G300 CVC • PROGRAMMABLE VIDEO TIMING GENERATOR • COLOR LOOK-UP TABLE

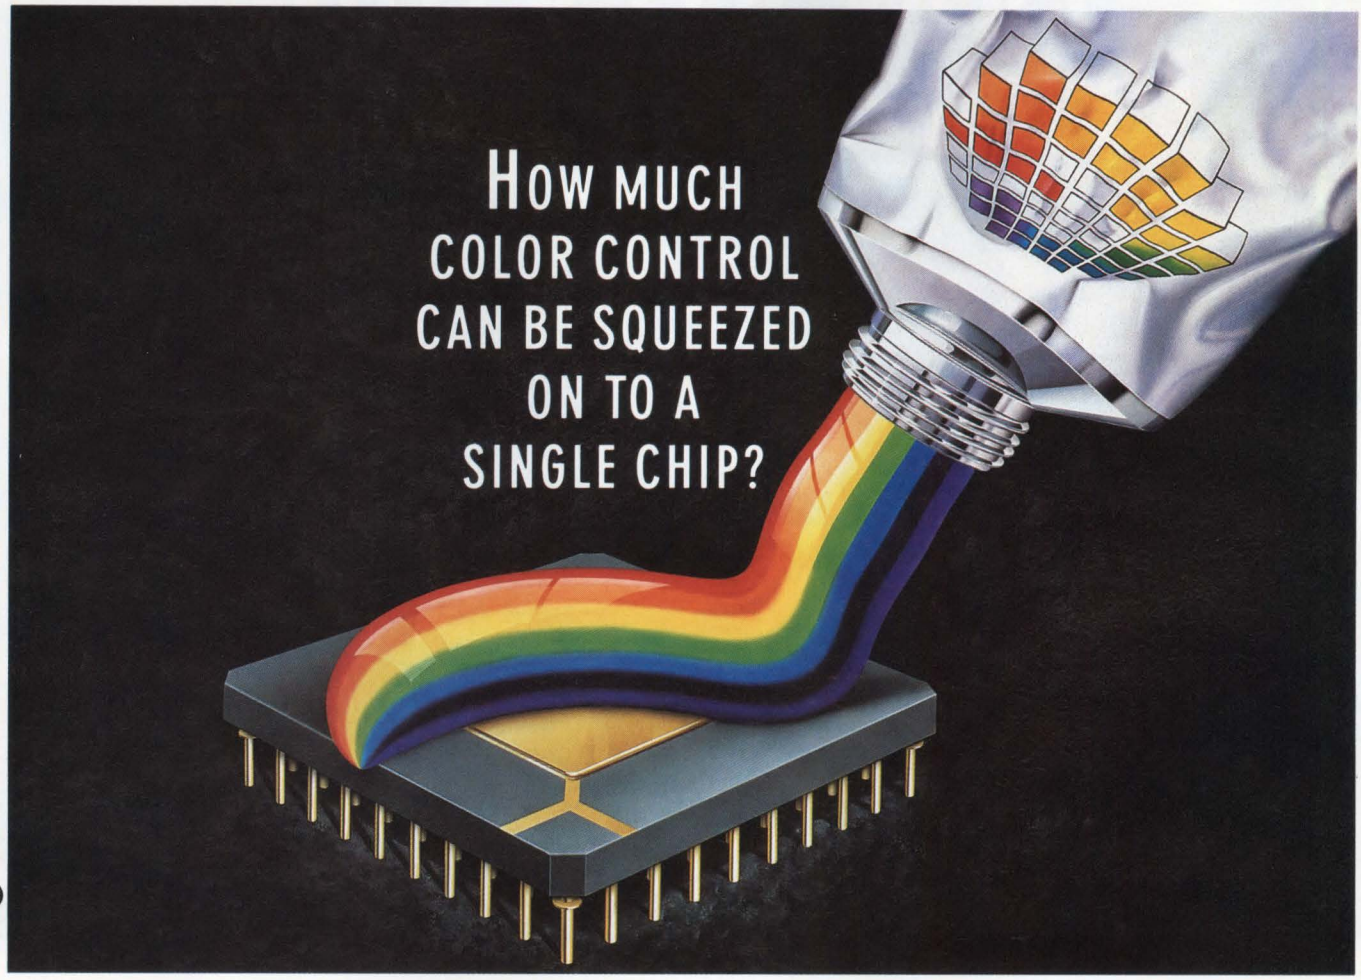

INMOS ● and IMS are trademarks of the INMOS Group of Companies.

FULL BITMAP MANAGEMENT • TRIPLE 8-BIT DAC • PHASE-LOCKED LOOP • G300 CVC

If the chip in question is the new IMSG300 Color Video Controller (CVC) from INMOS, the answer is quite simply, "all of it".

Through high integration, INMOS has combined total graphics control and data conversion onto a single 84 pin device. This includes a color look-up table, programmable video timing generator, full bitmap management, triple 8-bit video DACs and a phase-locked loop.

Because the IMSG300 is programmable, it is easily interfaced to any processor or monitor that may be used in your system. The IMSG300 promotes the technique of using software drawing algorithms as opposed to designing them in hardware. This dramatically reduces design-in time and allows for easy upgradability.

The IMSG300 also provides an optional 24-bit pixel mode for full color as well as pseudo color, an on-chip phase-locked loop for generating all high frequency clocks, and a 32-bit multiplexed pixel port for achieving pixel rates up to 110MHz.

Only the INMOS G300 CVC gives you all the options you need to tailor your systems exactly the way you want them – both now and in the future.

The INMOS G300 CVC single chip solution. Find out more today.

inmos

INMOS, PO Box 16000, Colorado Springs, Colorado 80935. Telephone: (719) 630 4000

Baltimore (301) 995 6952 • New York (914) 897 2422 • Denver (303) 368 0561 • Santa Clara (408) 727 7771
Atlanta (404) 242 7444 • Dallas (214) 490 9522 • Minneapolis (612) 932 7121 • Boston (617) 229 2550 • Orange County (714) 957 6018

CIRCLE NO. 41

How To Picture Printer Technology.

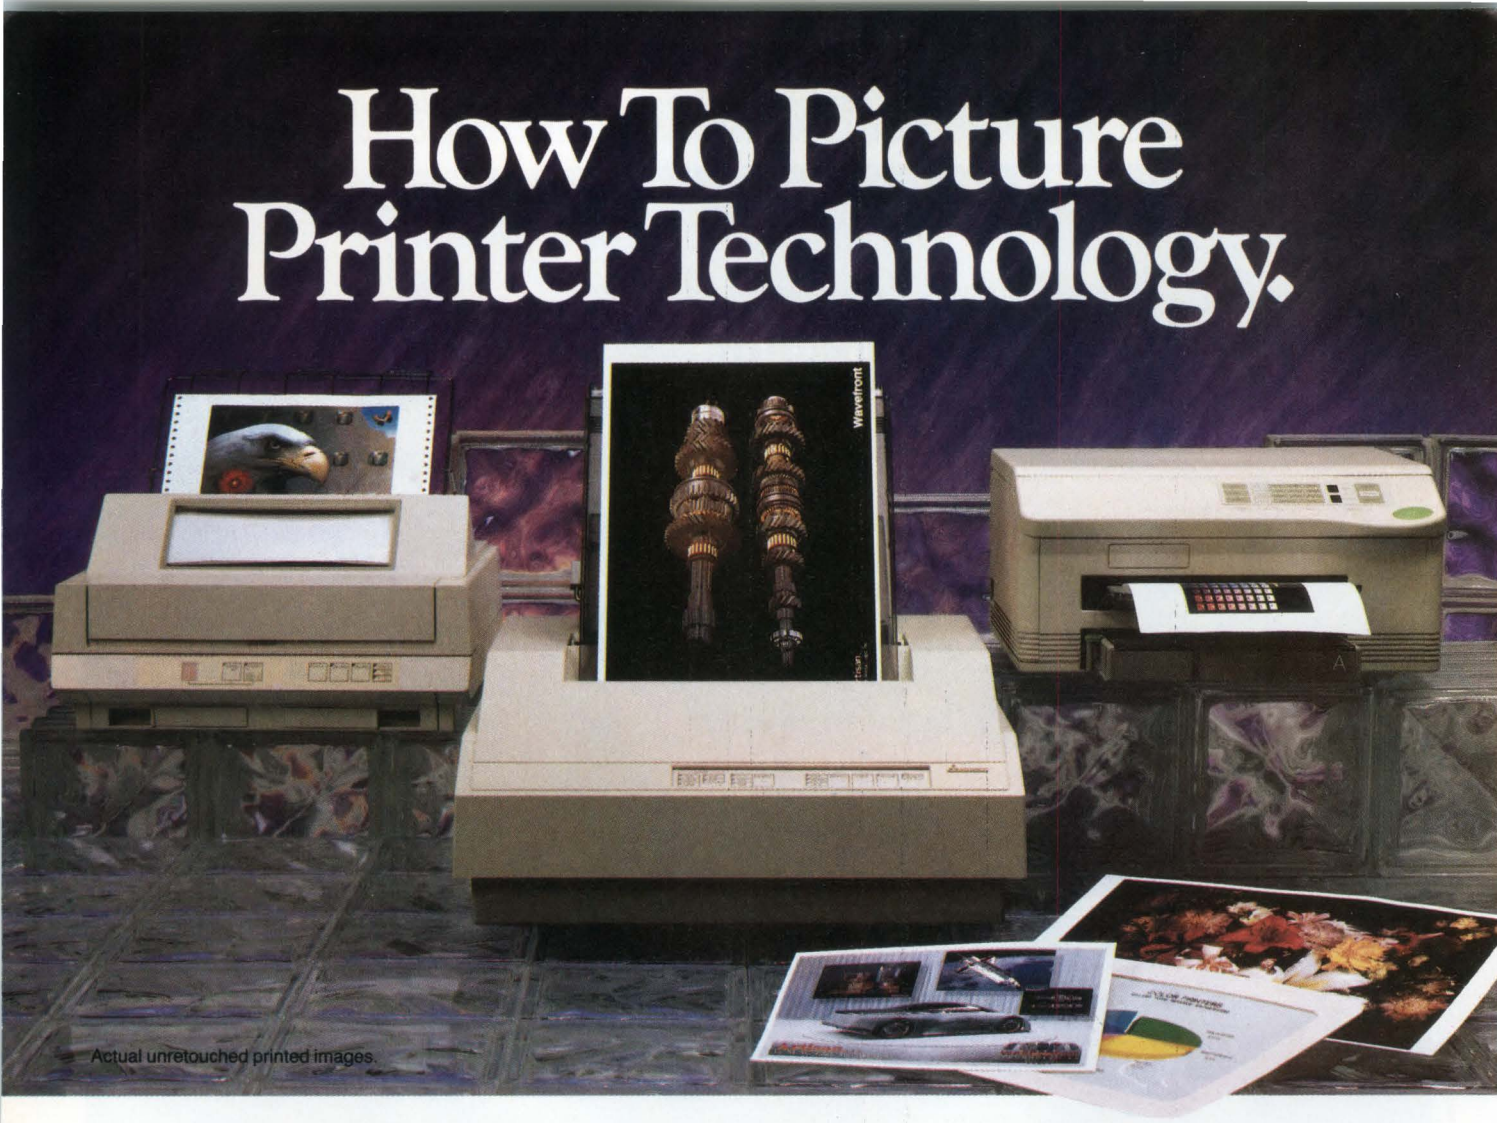

Actual unretouched printed images.

The broadest line of high-quality color thermal transfer printers.

When it comes to thermal transfer printer technology, Mitsubishi Electronics offers one of the most colorful and profitable product lines in the industry.

Applications such as CAD/CAM, presentation graphics, solids modeling and graphic arts demand high-quality color printing. And Mitsubishi® is there to meet that demand with the industry's broadest line of dependable, color thermal transfer printers.

Fast, full-color output.

Mitsubishi's color thermal transfer printers produce bright, bold images in an array of colors. Using transparency film or paper, resolution is exceptionally sharp and printouts are fast: 60 or 80 seconds in full color.

Innovations that sell.

Mitsubishi printers offer other features beyond great color reproduction and speed. The G330-70, with a built-in video interface, is compatible with many popular graphics boards. This allows creation of a permanent hard copy of an image viewed on a computer monitor. As the originator of 300 dpi print resolution in color thermal transfer technology with the G650, Mitsubishi continues to set the performance standards by which others are compared. The high resolution G650 prints on large B-size cut sheet paper or transparency film, as well as on A-size.

Largest OEM supplier of color thermal transfer printers.

As the leading OEM supplier, Mitsubishi Electronics has the capacity to support you with proven color thermal transfer printer technology. Mitsubishi's design and production expertise has earned us a reputation for creating products with unparalleled reliability and long life.

Mitsubishi Model	Resolution (DPI)	Print Substrate	Speed (sec./page)	Interface
G330-70	150	Cut sheet paper: A, A4 size Transparencies: A, A4 size	80 (color) 30 (monochrome)	Video: 12.5 MHz to 27 MHz bandwidth
G500	240	Continuous forms: A size Transparencies: A size	60 (color) 25 (monochrome)	8-bit Centronics® parallel
G650	300	Cut sheet paper: A, B size Transparencies: A, B size	60 (color) 25 (monochrome)	8-bit Centronics parallel

Performance Peripherals.

If you would like to picture more profitable printer sales, call or write Mitsubishi Electronics America, Inc., Computer Peripherals Division, 991 Knox Street, Torrance, CA 90502, (213) 217-5732.

Sales Offices:
Carrollton, TX (214) 241-5300
Minnetonka, MN (612) 938-7779
Mt. Prospect, IL (312) 298-9223
Norcross, GA (404) 368-4845
Piscataway, NJ (201) 981-1001
Sunnyvale, CA (408) 730-5900
Torrance, CA (213) 217-5732
Woburn, MA (617) 938-1220

© 1988 Mitsubishi Electronics America, Inc.
Mitsubishi is a registered trademark of Mitsubishi Electric Corp., Tokyo.
Centronics is a registered trademark of Ekco Group, Inc., formerly Centronics Corp.

Images printed with permission from the following companies *(trademarked software package name follows company name)*: Truevision, Inc. *(Targa, TIPS)*; Media Logic, Inc. *(Artisan)*; Zenographics, Inc. *(Mirage)*; Software Publishing Corporation *(Harvard* Graphics); CAP International. Monitor image produced with permission from Autodesk, Inc. *(AutoShade)*.

MITSUBISHI ELECTRONICS

CIRCLE NO. 42

INTEGRATED CIRCUITS

AMD offering joins second-generation disk controller chips

Second-generation disk controller chips—ICs that can perform most of the functions of a board-level disk controller—are becoming a new battleground for silicon vendors. The recent three-chip entry from Advanced Micro Devices (Sunnyvale, CA) follows the pattern of earlier products by combining the sequencer, disk data paths and buffer management onto a single die. But the offering has a few surprises, such as enhanced performance and specific provision for the needs of optical media.

The three new chips—a magnetic disk controller, an optical disk controller and a companion Reed-Solomon error detection and correction chip—are intended for use in embedded controllers. The new level of integration reduces the chip count for a SCSI disk controller from the previous boardful of components to about a half-dozen packages. This reduction is good news for desktop PC builders, who must cope with shrinking footprints. It's vital in laptop computers, because the disk drives can be as small as 2½ in. and because the system has no room for a separate controller board.

But at the same time that customers are demanding smaller disk controllers, system architects are demanding more sophisticated controllers. To keep up with 32-bit personal computer and workstation CPUs, disk subsystems are resorting to track buffering, disk caching, reduced-latency operations and other tricks that require both large memories and considerable intelligence in the controller. So the vendor offering a second-generation product can't just find a popular board-level controller and shrink it. The new product must be ready to face the demands of SCSI disk designs that resemble supermini mass-storage systems in their complexity.

Tailoring for speed

The AMD 95C95 magnetic disk controller takes on both of these size and sophistication challenges. Like other second-generation products, the chip combines a formatter, a microsequencer, a microcontroller interface, and buffer management and error detection/correction circuitry into a single package.

To ensure flexibility, the chip offers a writable-control-store-based formatter, permitting coverage of ESDI, SMD, ST506 and custom disk formats through RAM pattern changes. The microcontroller interface supports both Intel- and Motorola-style microcontrollers. The chip will connect to just about any type of 8- or 16-bit bus interface, including AMD's new 33C93A SCSI chip.

While the 95C95 can be used to implement a low-parts-count, low-cost controller, many of the chip's advantages become apparent only in high-performance designs. Here, the chip offers low latency, very high throughput, sophisticated buffer management with overlapped transfers, and the ability to attach huge buffer memories.

The maximum data rate of the chip is 32 Mbits/s, well in excess of the needs of any existing disk. Perhaps more important than the raw speed are the reduced command latency and the overlapped operations. The sequencer, for example, can start a multisector read operation at any sector; it doesn't have to begin with the first sector in the sequence. By reading the sectors as they go by (at one-to-one interleaving), and then assembling them in their proper order in buffer memory, the controller saves an average of one half-rotation of latency.

The on-chip, eighth-order Reed-Solomon EDC circuit checks each data block while the following block is read, so checked sectors are available in buffer memory with a delay of only one sector. The controller overlaps transfers into buffer memory from the disk with transfers out of buffer memory to the bus interface, again increasing throughput. Numerous clever addressing modes are available through the sequencer, giving both the bus interface and the microcontroller easy access to sectors in buffer memory.

The memory itself can be extremely large, to accommodate track buffering and aggressive caching schemes. To make these large memories practical, AMD has put an on-chip dynamic RAM controller into the 95C95. The on-chip controller and an external PAL can handle megabyte-sized buffers built from 1-

THE 95C96 OPTICAL DISK CONTROLLER

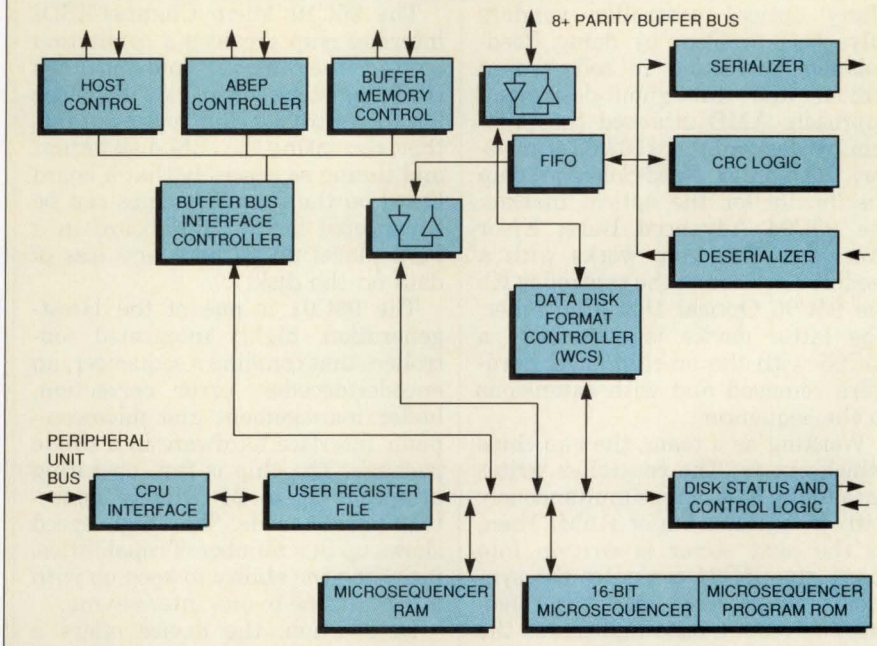

INTEGRATED CIRCUITS

Mbit DRAMs. To minimize the chance of data loss in these large arrays, both the controller and AMD's SCSI chip work with parity memory, using a 9-bit data path.

Looking toward optical media

Since many second-generation controllers can be programmed to handle the formats used in optical disks, AMD has also made provision for opticals. But AMD has done more.

The problem is one of error rates—while a good magnetic disk will have one error in 10^7 bits, opticals are lucky to manage an error in 10^5. Many optical controller vendors solve this problem by doing Reed-Solomon correction in software—a tedious and throughput-destroying approach. AMD attacked the problem by designing a 110,000-transistor, 16th-order Reed-Solomon chip specifically for the optical market: the 95C94 Advanced Burst Error Processor. This chip works with a modified version of the controller IC: the 95C96 Optical Disk Controller. The latter device is essentially a 95C95 with the on-chip EDC hardware removed and with extensions to the sequencer.

Working as a team, the two chips attack errors. The controller writes data to the EDC chip simultaneously with writes into buffer RAM. Then, as the next sector is written into RAM, the 95C94 computes the syndromes for the current sector, calculates correction data and places the correction vectors in a queue.

The 95C96 sequencer then checks the error status. If errors are detected, the sequencer takes the correction vectors off the queue and uses them to repair the errors in buffer memory. As in the case of the simpler on-chip circuitry of the magnetic disk controller, the error detecting and correcting is overlapped with the reading of the next sector.

Samples of all three parts are available, and the prices in 1,000-unit orders are $66 for the 95C95, $65 for the 95C96 and $53 for the 95C94 Reed-Solomon chip.

—*Ron Wilson*

Advanced Micro Devices
901 Thomson Pl
Sunnyvale, CA 94088
(408) 732-2400

Circle number 482

Chip set offers IBM-compatible, Micro Channel-ESDI disk controllers

As IBM shows more signs of actually offering licenses to clone the PS/2 architecture, more and more copies of the PS/2's proprietary IBM chips are showing up in other people's catalogs. The most recent cloning job is on IBM's Micro Channel-ESDI disk controller. Standard Microsystems (Hauppauge, NY) now offers a pair of chips that can be used to reproduce a compatible version of the IBM-proprietary board.

The 95C10 Micro Channel-ESDI interface chip provides a connection between the controller and bus that's IBM register-compatible. But the 95C01 controller chip goes even further, recreating the IBM disk format and timing so precisely that a board based on the Standard chips can be substituted for an IBM board in a PS/2 Model 80 without any loss of data on the disk.

The 95C01 is one of the latest-generation, highly integrated controllers that combine a sequencer, an encoder/decoder, error correction, buffer management and microcomputer interface hardware in a single package. The chip is fast, operating at speeds up to 24 MHz in nonreturn-to-zero mode. This high speed shows up in a number of capabilities, including the ability to keep up with a disk at one-to-one interleaving.

In addition, the device offers a number of programmable options. If the customer wants to change timing, sector formatting, read gating or any of the other innumerable nasty details of disk operation, loading different microcode is all that's necessary. Standard claims that this capability lets the controller operate with a wide range of devices, from floppy drives to ESDI disks, ST506 disks, SMD disks, optical disks and even ¼-in. tape drives.

The error-handling needs of most of these devices can be met by on-chip hardware. The 95C01 provides both IBM-compatible, CRC-16/ECC-32 circuitry and a separate, programmable, 8- to 64-bit ECC circuit.

The 95C01 interfaces to a local microcontroller through a port capable of Z8, 8051 or 80188 connection. This port gives the microcontroller access to the disk controller internals and to a buffer memory.

Buffer memory can be built from up to 64 kbytes of standard static RAM. Access to the memory is managed by a three-channel DMA controller inside the 95C01. With this controller, the disk data circuitry, the 95C10 Micro Channel interface, and the microcontroller chip all have access to the buffer on a priority basis. In addition, the 95C01 circuitry manages the SRAM as a ring buffer, maintaining data pointers, keeping track of data validity, and timing data transfers between the disk and bus through the SRAM.

A Micro Channel connection

The other half of the Standard chip pair is an interface chip to go between the 95C01 controller and the Micro Channel bus. The 95C10 is intended specifically for this role, but the company claims that it can be used just as well with other disk controllers.

As do most Micro Channel chips these days, the 95C10 provides register-level IBM emulation, so there should be no adjustments necessary to the standard PS/2 BIOS. The chip also supports IBM's Programmable Option Select strategy, which includes programmable DMA burst length and fairness mode. Using a local 26-byte FIFO, the 95C10 can gather up DMA bursts and pass them through the Micro Channel at up to 5 Mbytes/s, helping to conserve bandwidth on the increasingly crowded buses of high-end PCs.

On the controller side, the chip has an 8-bit, multiplexed address/data connection for a local microcontroller, and a switchable 8- or 16-bit interface to a ring buffer. This latter connection to the buffer forms the normal read/write data path for the interface, although the chip can be configured to permit DMA transfers between the Micro Channel and the microcontroller's external memory.

—*Ron Wilson*

Standard Microsystems
35 Marcus Blvd
Hauppauge, NY 11788
(516) 273-3100

Circle number 483

FAX'n Easy

Designing Group 3 FAX capabilities into a laptop or other portable device has just been made faster and easier. All it takes is the new YM7109, a true single-chip 9600bps FAX MODEM, and FAXSIM™ (FAX Software Integration Module) by Yamaha LSI.

It's in there.

With the YM7109, the solution to your Group 3 FAX design is suddenly no bigger than a 40-pin DIP or 64-pin QFP. And, FAXSIM can cut weeks, even months off your design time. Because FAXSIM allows you to sidestep thick manuals and detailed T.4 or T.30 documentation. There's no need to interpret chip specs and study the idiosyncrasies of communication protocol or training sequences because we make the source code (developed in Borland Turbo C) available. FAXSIM gives you what you need to beat the competition — a head start.

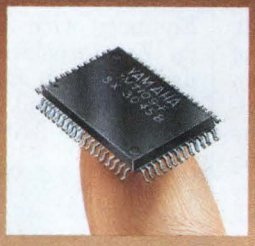

It takes so little power.

The YM7109 typically consumes only 200mW ± 25mW, so batteries will last longer. And, the YM7109 needs only one +5V supply. DTMF detect for remote control or security, flag pattern detect, selectable cable equalization, receive signal monitoring, programmable receive level adjust and auto equalization are all in the chip. Of course, the YM7109 meets CCITT V.29, V.27ter, V.21 channel 2, V.23 backward channel, V.21 and Bell 103 standards. It's compatible with industry-standard chip sets and has been tested on both domestic and international telephone lines.

Get the FAX. Just call 1-800-543-7457

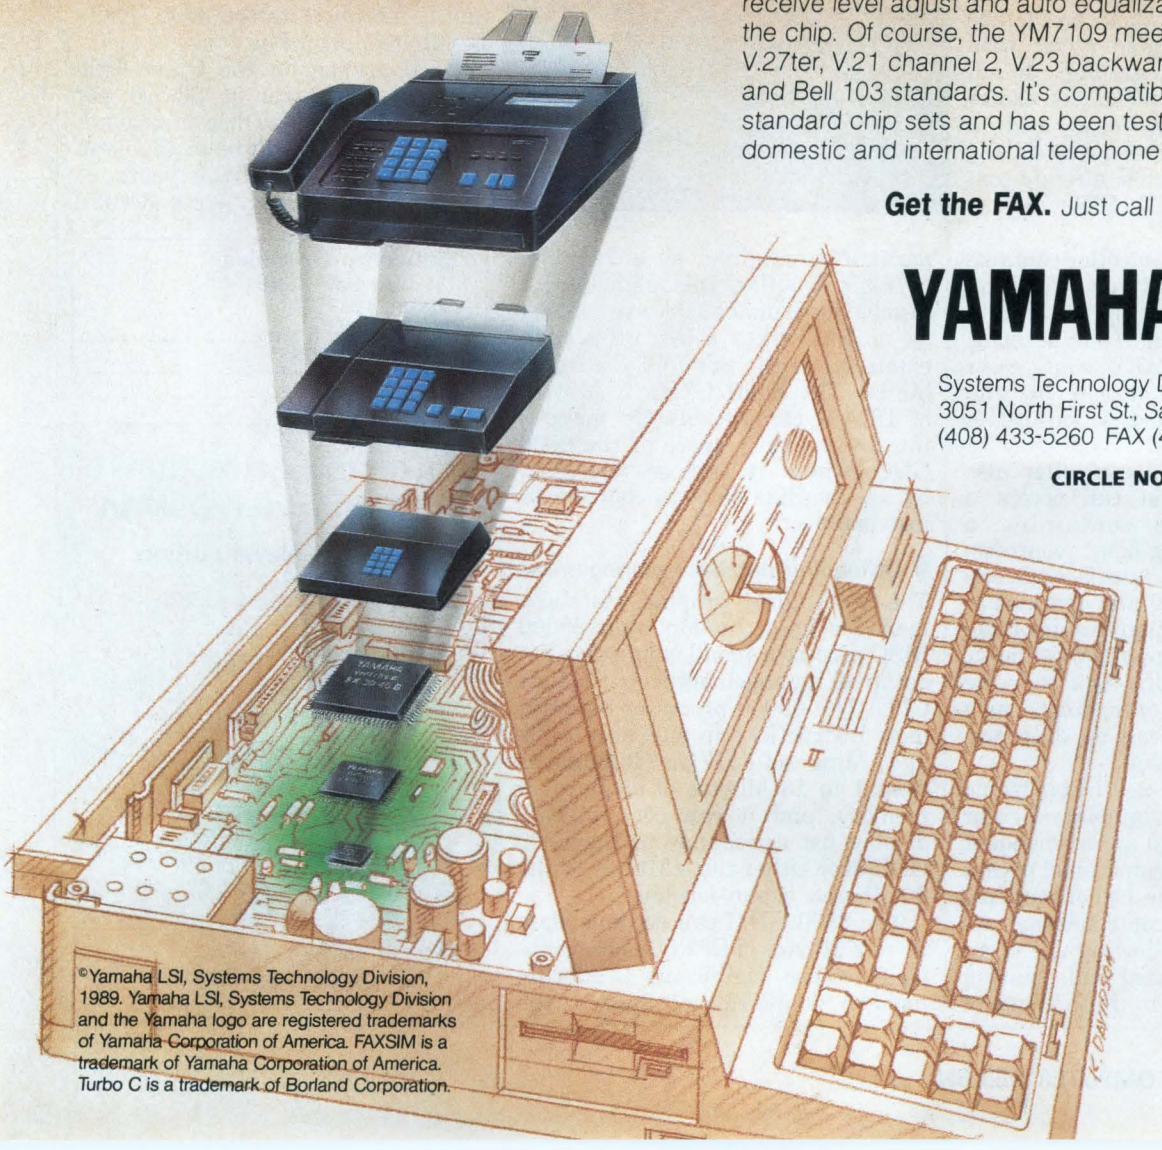

YAMAHA® LSI

Systems Technology Division
3051 North First St., San Jose, CA 95134
(408) 433-5260 FAX (408) 433-5230

CIRCLE NO. 43

©Yamaha LSI, Systems Technology Division, 1989. Yamaha LSI, Systems Technology Division and the Yamaha logo are registered trademarks of Yamaha Corporation of America. FAXSIM is a trademark of Yamaha Corporation of America. Turbo C is a trademark of Borland Corporation.

COMPUTERS AND SUBSYSTEMS

VME display modules build workstation graphics

A family of modular graphics display boards for VME-based workstations is one of the first product lines to use the new TMS34020 graphics processor and the companion TMS34082 floating-point coprocessor from Texas Instruments (Dallas, TX). The Omni 8800 GDS family from Omnicomp also makes use of the new 88000 RISC processor from Motorola (Austin, TX). Although initially designed for the VMEbus, the 8800 GDS family is rapidly being implemented for the PC AT bus as well, according to Anthony Masroff, Omnicomp president.

The Omni family of display boards includes the 8600 GDC display controller, the 8800 GDM graphics database manager and the 8200 FGB frame grabber. The latter board digitizes images so that they can then be used and manipulated along with computer-generated graphics. In addition, there's a frame buffer extension board that can add a 4k-×4k-×12-bit pixel array to the 2k-×1k-×12-bit buffer already contained on the 8600 GDC display controller.

The 8600 GDC controller contains a 10-Mips TMS34020 graphics processor that has a 1-Mbyte program memory as well as a 512-byte cache. Since the TMS34020 is a processor in its own right, it can be programmed to support a variety of graphics standards.

The 8600 display controller also has connectors that can accept a daughter module containing a TMS34082 floating-point coprocessor. The TMS34082 interfaces directly to the TMS34020 and operates at 40 MFlops. In addition to supporting the normal floating-point arithmetic operations, the TMS34082 supports graphics-oriented operations, such as splines, and it can be microprogrammed by the user.

The 8600 GDC's 2-×1-kbit frame buffer has four overlay planes, eight underlay planes and an independent cursor plane. Additional pixel planes can be added by the use of multiple 8600 GDCs. The controller has its own serial ports, allowing use of an independent keyboard and pointing device. In addition to the 1 Mbyte of program dynamic RAM, there's 1 Mbyte of EPROM for on-board graphics firmware.

The 8400 GZB Z-buffer board works as a companion module to the 8600 GDC in systems that perform three-dimensional solids modeling. Equipped with 8 Mbytes of DRAM for Z interpolation, the 8400 GZB uses a custom gate array to perform hidden line and surface removal.

While the 8400 GZB's memory is

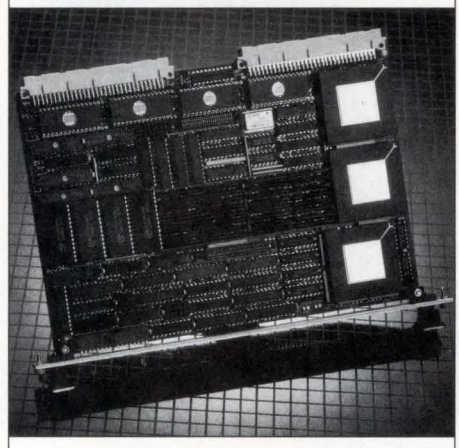

normally configured as a 2k-×2k-×16-bit Z-buffer, the memory can also be used, under software control, as a 2k-×4k-×8-bit off-screen extension to the 8600 GDC's frame buffer. The 8400 GZB's 8 Mbytes of DRAM can be directly mapped into the memory space of the 8600 GDC, where it can be addressed by x-y coordinates from coprocessor instructions.

Offloading database management

The 8800 GDM graphics database manager uses the Motorola 88000 RISC processor to offload display list database management from the host CPU. The 88000 processor executes most instructions in one clock cycle at 17 Mips and 7 MFlops. It manages from 1 to 16 Mbytes of display list memory, and it can communicate display list commands to the 8600 GDC over either the VMEbus or the 20-Mbyte/s expansion bus.

The 8800 GDM can accommodate up to 2 Mbytes of EPROM for graphics firmware. Omnicomp is offering its own Omni*Kernel System, the ANSII standard Graphics Kernel System, and the extensions to the programmer's hierarchical interface to graphics systems, which has solids-modeling features.

In addition, Omnicomp will be offering a firmware version of the Dynamic Object-oriented Rendering Environment (DORE) developed by Ardent Computer (Sunnyvale, CA). DORE lets the user interactively select speed of rendering versus detail (wire-frame, flat-facet, Gouraud or Phong shading, or full surface-texture mapping).

The 8200 FGB frame grabber board digitizes standard video images in either 756-×485-×8-bit (NTSC) or 768-×575-×8-bit (PAL) format. It can store up to four digitized images and can perform operations such as frame averaging, subtraction or addition on them. The 8200 FGB supports four independent overlay plans for text or graphics, and it interfaces to the 8600 GDC's 20-Mbyte/s pixel bus.

OEM pricing for the Omni 8800 GDS family starts at $3,000 per board, with quantity discounts available. The boards will be available in the second quarter of this year.

—Tom Williams

Omnicomp Graphics
1734 West Belt North
Houston, TX 77043
(713) 464-2990

Circle number 478

Coming attractions in *Computer Design*

June 12 News Edition

- Designers' Buying Guide to CAE/CAD tools
- Product Focus Report on modem ICs
- Product Focus Report on memory boards

July 1 Magazine Edition

- Cover Story on multiprocessing
- Technology Report on programmable logic devices
- Technology Report on imaging subsystems

Here Are Over 2000 Different VME I/O Boards

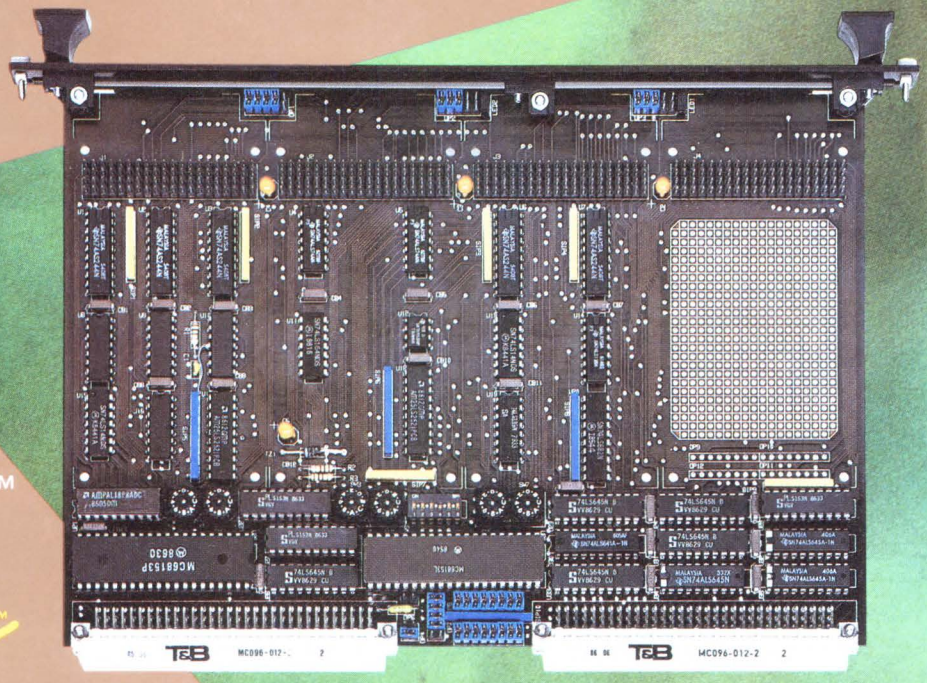

OB68K/VIO™ with Omnimodules™

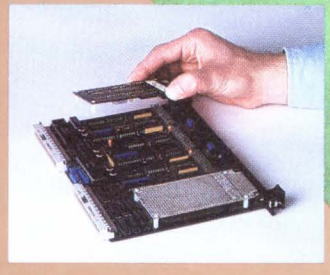

You can use our OB68K/VIO VME I/O board, with over 2000 different standard configurations, to fit your needs. Its flexibility comes from our OMNIMODULE modular I/O. All the I/O functions for the OB68K/VIO are contained on plug-in OMNIMODULE's. Up to four OMNIMODULE's can be plugged into one OB68K/VIO. With our five currently available OMNIMODULE's you can custom tailor our OB68K/VIO in over 2000 different configurations!

The OMNIMODULE's currently available include:

- Kluge prototyping module
- (2) RS232C asynchronous serial ports.
- (2) RS232C synch/asynch serial ports.
- (2) RS422 asynchronous serial ports.
- (20) lines of parallel I/O.

We will be coming out with new OMNIMODULE's in the future such as a SCSI controller, IEEE 488 controller, (2) RS422 synch/asynch ports, etc. If you have some very unique and specialized I/O requirements, you can build it on our kluge module. Our OMNIMODULE specification is available to you for just that purpose. If you will need OEM quantities, we can produce your design as a standard OMNIMODULE.

Each OMNIMODULE has one interrupt standard with a second interrupt optional. Even fully loaded with OMNIMODULE's, the OB68K/VIO will fit into just one card slot. You can choose to bring out up to four of the OMNIMODULE I/O ports to the front panel, or you can bring out up to 3 of the OMNIMODULE I/O ports to the P2 connector.

To obtain a free data sheet or to learn about our latest OMNIMODULE contact our Marketing Manager, Peter Czuchra at 1-800-638-5022 or (312) 231-6880 in Illinois.

Send $10.00 for a detailed technical manual or a copy of the OMNIMODULE specification.

OMNIBYTE

OMNIBYTE CORPORATION
245 W. Roosevelt Road
West Chicago, IL 60185-3790
Fax No. 312-231-7042
CALL TOLL FREE
1-800-638-5022
In IL (312) 231-6880
After Nov. 11, 1989, call (708) 231-6880

A Look At Today . . . A Vision Of Tomorrow.

CIRCLE NO. 120

A VME AND MULTIBUS* ASSORTMENT THAT WOULD FILL OVER 15 HOPPER CARS.

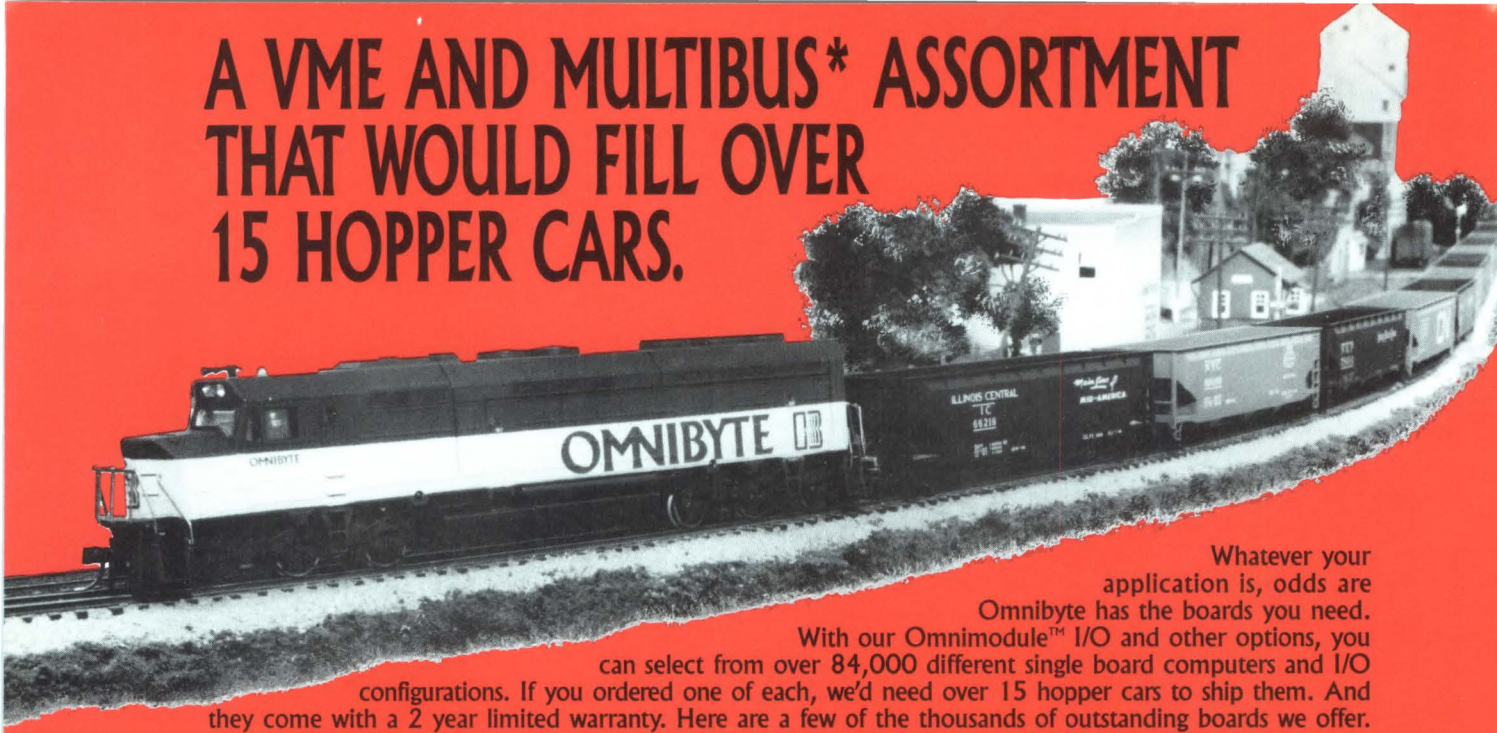

Whatever your application is, odds are Omnibyte has the boards you need. With our Omnimodule™ I/O and other options, you can select from over 84,000 different single board computers and I/O configurations. If you ordered one of each, we'd need over 15 hopper cars to ship them. And they come with a 2 year limited warranty. Here are a few of the thousands of outstanding boards we offer.

If you don't see the board you need in this ad or would like more information, call our Marketing Manager, Pete Czuchra today. He'll put you on the right track to finding the boards that fit your needs.

OB68K/VSBC20™ VME SINGLE BOARD COMPUTER

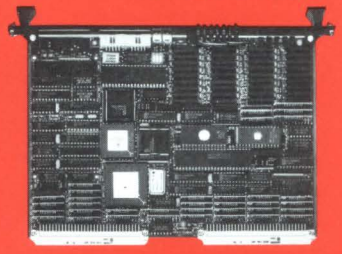

- 68020 12.5 – 33MHz CPU
- 1-4 MB of dual-access, zero-wait-state DRAM with parity
- 68881/2 (optional)
- (2) 32-pin ROM sockets
- (2) RS232C serial ports
- (2) 8-bit parallel ports
- (1) OMNIMODULE socket for a wide variety of I/O (ie. 2 serial ports, 20 parallel lines)
- 4 level bus arbiter (optional)

THE OB68K/VIO™ VME UNIVERSAL I/O BOARD

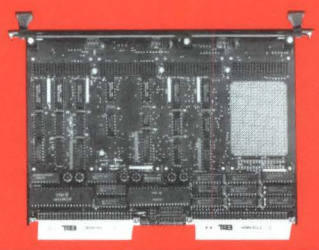

- (4) Omnimodule I/O sockets for a wide variety of I/O (ie. 8 serial ports, 80 parallel lines)
- One (1) interrupt per Omnimodule, two (2) optional

OB68K/VSBC1™ VME SINGLE BOARD COMPUTER

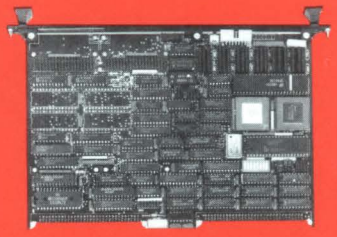

- 68000/10 12.5-16MHz CPU
- 512KB of dual-access, zero-wait-state DRAM with parity
- (4) 28-pin ROM sockets
- (3) 16-bit counter/timers
- (2) Omnimodule™ I/O sockets for a wide variety of I/O (ie. 4 serial ports, 40 parallel lines)
- DMA controller (optional)
- VME bus interrupter (optional)
- Optional 4 level bus arbiter

OB68K1A™ MULTIBUS SINGLE BOARD COMPUTER

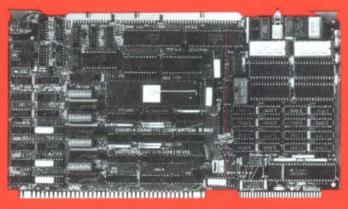

- 10MHz 68000 CPU
- 128K/512K of zero-wait-state dual ported RAM
- Up to 192K bytes of EPROM
- (2) RS-232C serial ports
- (2) 16-bit parallel ports
- A triple 16-bit timer/counter
- (7) Prioritized-vectored interrupts

OB68K/VME1™ VME SINGLE BOARD COMPUTER

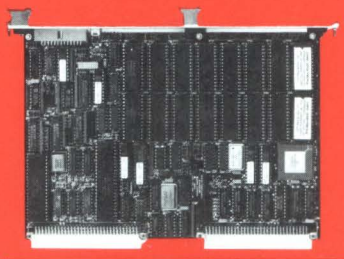

- 12.5MHz 68000 CPU
- (8) pairs of 28-pin sockets for RAM and ROM (up to 448K RAM or 896K ROM)
- (2) RS-232C serial ports
- (2) 8-bit parallel I/O ports
- System controller

OB68K230™ MULTIBUS PARALLEL I/O BOARD

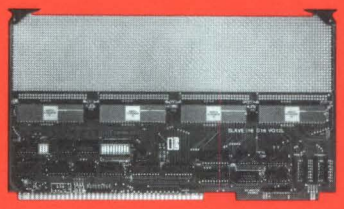

- 96 bits of software definable parallel I/O
- (4) 68230 PI/T chips
- (4) 24-bit timers
- 35 sq. in. of prototyping area

OMNIBYTE™

OMNIBYTE CORPORATION
245 W. Roosevelt Road
West Chicago, IL 60185-3790
Fax No. 312-231-7042
CALL TOLL FREE 1-800-638-5022
In IL (312) 231-6880
After Nov. 11, 1989, call (708) 231-6880

Photo shot at Midwest Railroad Modelers, Batavia, IL. *Multibus and ISBX are trademarks of Intel Corp.

CIRCLE NO 119

Data-acquisition speed approaches NuBus' theoretical limit

Ever since the introduction of 32-bit personal computers, PCs have been replacing many of the CAD/CAM workstations and dedicated instrumentation controllers used during the era of 8-bit PCs. Even the 32-bit PCs had limitations, but high-performance direct memory access (DMA) schemes and other improvements on these machines have overcome many of the original problems. IBM's Micro Channel, for example, has an achievable bandwidth of 20 Mbytes/s, and the NuBus used in the Apple Macintosh II has an achievable limit of 37½ Mbytes/s. But actual speeds only approach these theoretical transfer rates.

One step toward attaining the maximum speed available on the Macintosh II is provided by the NB-DMA2800 block-mode DMA interface board from National Instruments. "Although DMA is generally necessary to obtain continuous sampling above 1 Msample/s on a PC, the Macintosh II doesn't contain a DMA controller, and the system memory on its motherboard won't transfer data in block mode," says Richard House, product manager for National's data acquisition division. The NB-DMA2800 overcomes this problem by continuously transferring data to block mode memory at 29 Mbytes/s, and transferring data to Macintosh II memory at up to 2.4 Mbytes/s.

"One DMA board can sustain several of the recently introduced NB-A2000 high-speed analog input boards," says House, "even when those boards are digitizing at their maximum rate of 1 Msample/s. The DMA board, therefore, can continuously transfer data in bursts up to the NuBus' 37½-Mbyte/s upper bandwidth limit."

DMA board allows direct transfer

I/O data acquired by plug-in data-acquisition boards can be transferred directly to memory by the DMA board rather than by the Macintosh II coprocessor. This frees the processor to perform other program tasks. The DMA board's NuBus-to-GPIB interface enables the Macintosh II to function as a high-performance GPIB (IEEE 488) instrument controller for sophisticated test, measurement, monitoring and control system applications, and to transfer data between Macintosh II computers and GPIB-compatible instruments.

Features of this DMA board include eight DMA channels, eight counter/timer channels and eight interrupt channels. Data can be transferred to and from the full 4-Gbyte NuBus address space. Like other boards in the company's NB series, the DMA board contains a Real-Time System Integration (RTSI) bus interface.

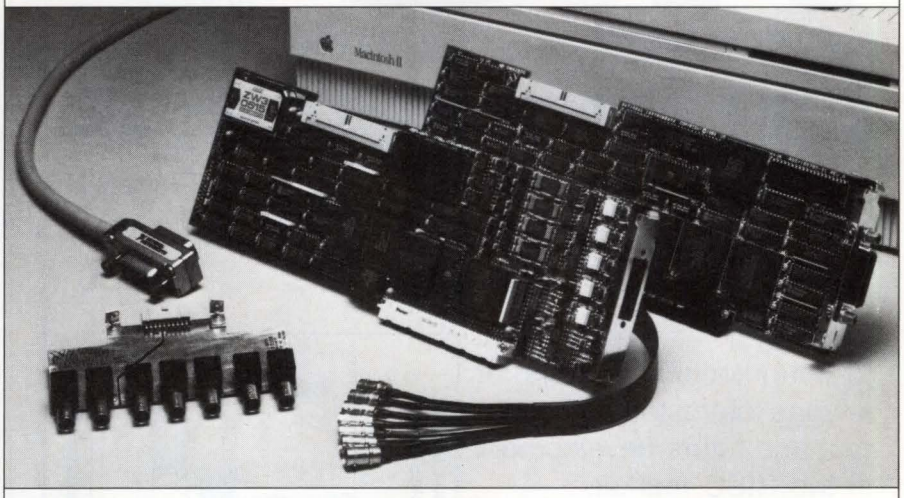

Each of the companion Macintosh II NB-A2000 high-speed analog input boards contains a precision 12-bit analog-to-digital converter and analog circuitry that's capable of sampling waveforms at up to 1 Msample/s. Four independent circuits simultaneously sample multiple analog signals.

Acquired data can be transferred from one of these boards directly to memory by either the Macintosh II processor or the DMA board, each of which contains a RTSI bus to carry DMA requests and acknowledge signals. The RTSI bus provides a data-acquisition system with the ability to synchronize the operation of multiple boards by routing digital trigger and timing signals between boards to coordinate analog, digital and timing operations. "As a result, the only limitation imposed on the duration of the user's sampling process is that of the computer's memory," claims National's House.

Software support includes Labview, a graphical programming system for developing scientific and engineering applications on the Macintosh II. This system acquires data from the NB-A2000, processes the data and displays the results with graphical presentation tools. An NB Labdriver Virtual Instrument Library, a special version of NB Labdriver, provides the user with additional programming language support, and is intended for use with Labview.

An NB-DMA2800 board with NI-488 Macintosh driver software is priced at $1,595 and an NB-A2000 board is priced at $2,995. The data-acquisition operations of both of these boards are controlled by NB Labdriver software, which is priced at $295.

A coax adapter board for the DMA board, with seven connectors, will be available for $225, and an alternative 1-m coax adapter cable will cost $175. The NB Labdriver Virtual Instrument Library will be available to all Labview users at no charge. Labview is listed at $1,995 and is available now. The other items are scheduled for August deliveries.

—Sydney F. Shapiro

National Instruments
12109 Technology Blvd
Austin, TX 78727
(512) 250-9119

Circle number 477

COMPUTERS AND SUBSYSTEMS

Development board brings 25-MHz i486 to Multibus II

Two Multibus II boards from Intel provide OEMs with immediate access to the high end of Intel's architecture. Based on the 25-MHz 80486 and the 33-MHz 80386 microprocessors, respectively, the iSBC 486/125 and iSBC 386/133 boards retain software compatibility with earlier 80386-architecture products.

The two boards are fully compatible, so OEMs can first design and ship systems using iSBC 386/133 boards and then upgrade those systems with iSBC 486/125 boards.

The iSBC 486/125 is designed around the new 32-bit, 1.2 million-transistor 80486 microprocessor, making it the most powerful CPU board Intel has ever provided for the Multibus II architecture. The board has 8 Mbytes of memory and features a two-way interleaved design that uses fast-page dynamic RAMs to optimize memory access.

A configurable memory

The iSBC 386/133 can be configured with up to 16 Mbytes of memory and will support up to 64 Mbytes of memory when 4-Mbit DRAMs become available. The board contains a 33-MHz 387DX coprocessor, an 82258 advanced direct memory access coprocessor, an iLBX II interface, two serial ports and sites for two 32-pin EPROMs.

Both boards have connectors for the new iSBC CSM/002 Multibus II central services module, which performs all central services required by Multibus II. The module also provides a battery-backed, time-of-day clock and a periodic alarm function. If the user installs the module on either board and uses that board in slot zero of a Multibus II system, a slot is freed for another use.

Both boards support the iRMX II real-time operating system, the iRMK real-time kernel and the Unix System V/386 operating system.

The iSBC 486/125 will initially be shipped as the iSBC 486/125DU development unit, with production units available in early 1990. The iSBC 486/125DU has a list price of $13,995. An iSBC 386/133 with 1 Mbyte of memory has a single-quantity price of $6,995. Three other versions are available—with 2, 4 and 8 Mbytes of memory at an additional $700/Mbyte. The iSBC CSM/002 sells for $295. *—Michael Donlin*

Intel
3065 Bowers Ave
Santa Clara, CA 95052
(800) 548-4725

Circle number 484

RISC

IMPACT ON THE COMPUTER MARKET

A 204-Page Comprehensive Market Report Analyzing the Applications and Competitive Trends for RISC Technology $1250

This report was researched for manufacturers, users, or anyone interested in the current and future development in the RISC market.

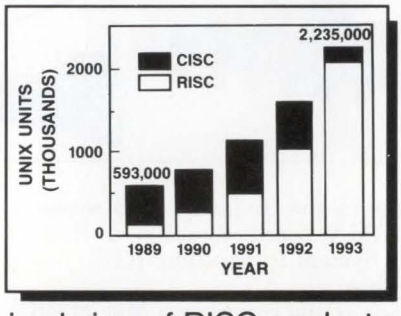

The report analyzes the current market for RISC chips, boards, and systems. It also provides a clear, organized view of RISC products offered by semiconductor manufacturers, computer system manufacturers, and OEMs.

A detailed forecast of the market by price class and a discussion of the key issues of concern to both users and manufacturers are also provided.

Order your copy NOW.
Call 408-996-7416.

Electronic Trend Publications
12930 Saratoga Ave., Suite D1
Saratoga, CA 95070
Tel: 408-996-7416
Fax: 408-996-7871

Application-Specific Knowledge A-S-K ETP Call 408 996-7416

CIRCLE NO. 45

YOU'RE INVITED TO JOIN ARROW'S BOOK-OF-THE-YEAR CLUB.

ABSOLUTELY FREE.

Arrow has just published system solution guides to give you the competitive edge in choosing among today's computer and peripheral products. And these books are available free from Arrow. Each *Systems Product Guide* is targeted to a specific audience: for original equipment manufacturers, value-added resellers, and commercial end-users.

In addition to product highlights, each edition includes selection guides, technical overviews, and a glossary of terms to provide detailed information and in-depth descriptions of computer and peripheral products from the industry leaders.

When it comes to helping you make fast, informed purchasing decisions, Arrow wrote the book. For your free copy of the *Arrow Systems Product Guide* that suits your needs, complete and return the attached coupon today.

*There is a difference in distribution. Arrow.*SM

ARROW ELECTRONICS, INC.

CIRCLE NO. 46

ALABAMA
Huntsville (205) 837-6955

ARIZONA
Phoenix (602) 437-0750

CALIFORNIA
Los Angeles (818) 701-7500
Orange County (714) 838-5422
San Diego (619) 565-4800
San Francisco (408) 745-6600

COLORADO
Denver (303) 790-4444

CONNECTICUT
Wallingford (203) 265-7741

FLORIDA
Ft. Lauderdale (305) 429-8200
Orlando (407) 333-9300

GEORGIA
Atlanta (404) 449-8252

ILLINOIS
Chicago (312) 250-0500

INDIANA
Indianapolis (317) 243-9353

IOWA
Cedar Rapids (319) 395-7230

KANSAS
Kansas City (913) 541-9542

MARYLAND
Baltimore (301) 995-6002

MASSACHUSETTS
Boston (508) 658-0900

MICHIGAN
Detroit (313) 462-2290
Grand Rapids (616) 243-0912

MINNESOTA
Minneapolis (612) 830-1800

MISSOURI
St. Louis (314) 567-6888

NEW HAMPSHIRE
Manchester (603) 668-6968

NEW JERSEY
Marlton (609) 596-8000
Parsippany (201) 538-0900

NEW MEXICO
Albuquerque (505) 243-4566

NEW YORK
Metro New York (516) 231-1000
Rochester (716) 427-0300

NORTH CAROLINA
Raleigh (919) 876-3132
Winston-Salem (919) 725-8711

OHIO
Cleveland (216) 248-3990
Dayton (513) 435-5563

OKLAHOMA
Tulsa (918) 252-7537

OREGON
Portland (503) 645-6456

PENNSYLVANIA
Philadelphia (215) 928-1800
Pittsburgh (412) 856-7000

TEXAS
Austin (512) 835-4180
Dallas (214) 380-6464
Houston (713) 530-4700

UTAH
Salt Lake City (801) 973-6913

WASHINGTON
Seattle (206) 575-4420

WISCONSIN
Milwaukee (414) 792-0150

PUERTO RICO
Hato Rey (809) 764-4384

CANADA
Montreal (514) 735-5511
Ottawa (613) 226-6903
Quebec (418) 871-7500
Toronto (416) 670-7769
Vancouver (604) 291-2986

International Export Group
TLX: 4758034

ARROW ELECTRONICS, INC.
Arrow/Kierulff Electronics Group
25 Hub Drive, Melville, NY 11747
Attn: Chris Johnson, Dept. N

☐ Please send me a free copy of the *1989 Arrow Systems Product Guide**
☐ For original equipment manufacturers. ☐ For value-added resellers.
☐ For commercial end-users.

☐ Please have a salesperson contact me.

Name _____
Phone _____
Title _____
Company _____
Address _____
City _____ State _____ Zip _____

*Quantities are limited.

COMPUTERS AND SUBSYSTEMS

Analog I/O boards deliver high performance on PS/2

Capitalizing on its extensive experience in developing data-acquisition boards for IBM PC/AT/XT and PS/2 computers or compatibles, Data Translation has added two Micro Channel boards to its PS/2 line of I/O products. Both boards—the DT2914 analog input board and the DT2916 analog output board—meet the requirements of advanced measurement and control and of scientific and industrial applications. They're supported for real-time data-acquisition applications by PS/LAB, a subroutine library software package.

The DT2914 and the DT2916 are new designs optimized for PS/2 Models 50, 60, 70 and 80. Each fits into an expansion slot in the PS/2 and contains low-power circuit elements to accommodate the limited power allotted to PS/2 peripherals.

The DT2914 provides 20-kHz throughput on 16 single-ended channels and samples inputs with 12-bit analog resolution to an accuracy of ±0.03 percent of full-scale range. Inputs can be sampled either one at a time or sequentially, through a programmable, on-board pacer clock that automatically initiates analog-to-digital data conversions.

For environments where continuous polling isn't practical, the board can generate an interrupt at the end of a data conversion and automatically signal the CPU that the conversion has been completed. Eight digital I/O lines on the board can be configured in four-line ports for input or output. Connecting the board to signal-conditioning panels provides protection for both the board and the computer and reduces signal noise. Software-programmable gains accommodate inputs as small as 625 mV with full resolution.

The DT2916 contains four independent 30-kHz digital-to-analog converters that can provide either single or simultaneous output. Twelve-bit analog resolution provides an accuracy of ±0.02 percent of full scale. Voltage or industrial current loop outputs from 4 to 20 mA are selectable on each channel.

Subroutine library

PS/LAB supports all functions of the Micro Channel data-acquisition boards on the PS/2, as well as the direct memory access transfers and the continuous-performance capability of selected boards to assure high-speed transfer without memory gaps.

PS/LAB permits control of all board operations from user-written custom programs, eliminating the need to master complex driver and I/O programming on the PS/2. Subroutines sample analog inputs, drive analog outputs and perform digital I/O operations in real time.

Other subroutines sample and process sensor data—for example, when making direct temperature measurements using thermocouples, PS/LAB subroutines sample data, linearize the data according to built-in tables and return temperature values directly in degrees. An error-processing system checks program and board operations for illegal conditions that could result in lost or erroneous data.

User-written routines supported under PS/LAB can be called from Microsoft C, Pascal, Fortran, QuickC and QuickBasic, or from Borland Turbo C and Turbo Pascal.

Both boards, as well as PS/LAB, are available immediately. The DT 2914 analog input board plus free Acquire software costs $595, the DT2916 analog output board costs $650 and the PS/LAB subroutine library software package costs $199.

—Sydney F. Shapiro

Data Translation
100 Locke Dr
Marlboro, MA 01752
(508) 481-3700

Circle number 476

VME I/O THAT GROWS WITH YOU.

Start with just a few digital or analog I/O channels and grow to over a 1000. Burr-Brown's affordable, expandable MPV940 Series brings intelligent VMEbus I/O interfacing to process control, robotics, simulators, data acquisition and more. The MPV940 controller serves as a stand-alone computer, as a mono-master VME system or as a Master/Slave in a multi-processor configuration. RS232 or VMEbus control and software loading. Compatible with Burr-Brown's PCI-20000 termination panels. Call or write for MPV940 facts, plus new full-line catalog: Burr-Brown Corporation, 1141 W. Grant Road, MS#131, Tucson, AZ 85705. (602) 746-1111.

BURR-BROWN®

©1988 5800-8155

MPV940 10 MHz 68000 CPU Controller board can be fitted with two Burr-Brown analog or digital multichannel I/O modules and still occupy one VMEbus slot.

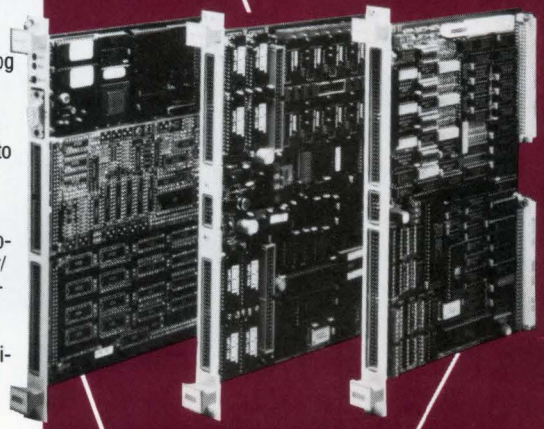

MPV941 Analog I/O expansion board provides 128 input channels & 8 output channels.

MPV942 Digital I/O expansion board provides 32 I/O lines plus 32 input lines with change of state monitor.

CIRCLE NO. 47

An inside look at the Bus/Board Market

- Where are Multibus I and II headed?
- Where are the PC/AT Bus and Micro Channel headed?
- Will STD design-ins increase or decrease in the next two years?
- How important are speed, clocking, message passing, software standards, microprocessor type, etc. in a design team's choice of a bus?
- What's the future for Futurebus?
- Are you in step with other system designers, or are you making a bus choice that will lead to a dead end?

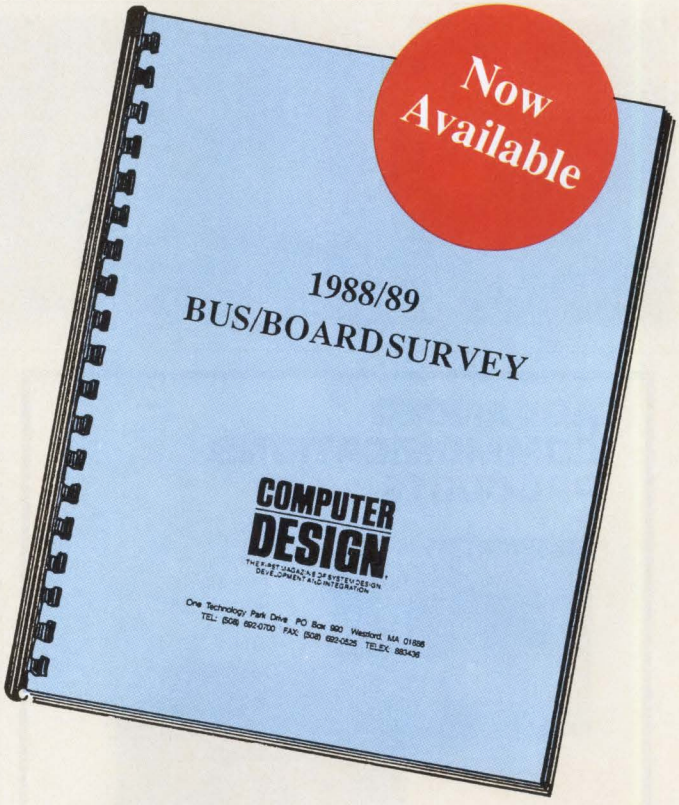

To learn the answers to these questions and a whole lot more, ***Computer Design*** conducted an in-depth survey of bus/board users and solicited comments and observations from nearly 20 manufacturers of board-level products. We've packed the survey results and extensive analyses into a **100-page** report that can be yours for only **$95**. Order your copy today. The competition is moving too fast to delay.

Please send me a copy of ***Computer Design's*** 1988/89 Bus/Board survey today.

Name_____

Company_____

Address_____

City_____ State____ Zip_____

My check for $95 is enclosed.____

Please charge my: American Express____ Master Charge____ Visa____

Account No._____ Exp. Date_____

Name of Cardholder (please print)_____

Signature_____

Mail to: ***Computer Design,*** One Technology Park Drive, Westford, MA 01886

COMPUTERS AND SUBSYSTEMS

Image processor provides fast link to other boards

A high-speed image-processing board from Toshiba America combines PC AT compatibility, a microprogrammable digital signal processor (DSP) chip, and a dedicated image bus that allows direct connection to boards from other manufacturers. The IP 9506 PC AT joins a family of software-compatible VME- and Multibus-based imaging boards, all of which use Toshiba's T9506 DSP chip.

The 32-bit T9506, which runs at 10 MHz on the Multibus and VME boards, is clocked at 5 MHz on the IP 9506 PC AT version. The T9506 has up to three stages of pipelining, and can access memory at up to 60 Mbytes/s. At the 5-MHz speed, the processor can perform a 1,024-point, 32-bit complex Fast Fourier Transform (FFT) in 4 ms. Image rotation, enlargement or reduction takes only 200 ns per 32-bit pixel.

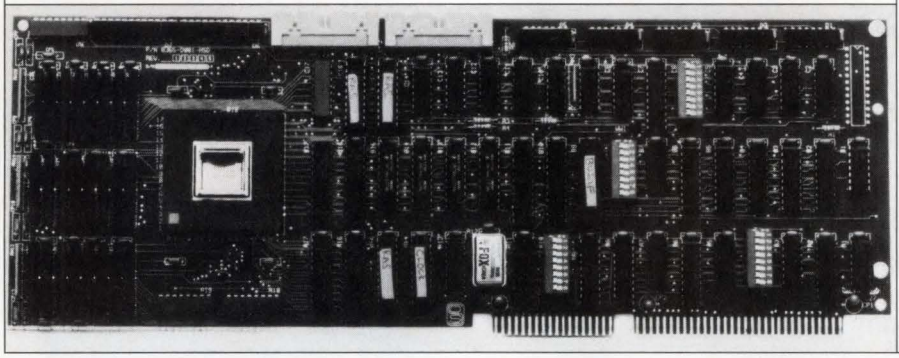

On-board memory consists of 3 Mbytes of dynamic RAM configured as 512 ×512 ×32 bits in three channels. The T9506 can access all three channels in parallel at a rate of 5 M 32-bit words/s for each channel. In addition, multiple processor boards can be configured in parallel for even faster processing and access to even larger pixel arrays.

The IP 9506 PC AT has a dedicated image bus that's compatible with the DT-Connect bus developed by Data Translation (Lowell, MA). The DT-Connect bus can perform 16-bit transfers at 1.67 MHz, and allows 8-,

(continued on page 105)

ADVANCED COMMUNICATIONS PRODUCTS

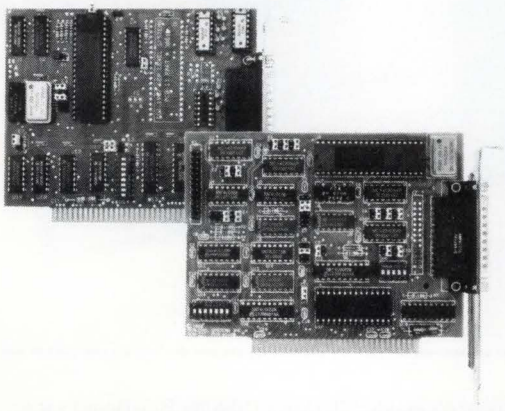

The Sealevel ACB-II and ACB-III Advanced Communications Boards provide SYNC/ASYNC high speed communications capabilities to PCs. Features include:

- SDLC/HDLC/ASYNC Protocols
- RS-232 and RS-422 Interfaces
- One or two channels with DMA
- Software Toolkit Provided

SEALEVEL COMMUNICATIONS & I/O

SEALEVEL SYSTEMS INC.
PO BOX 1808
EASLEY, SC 29641
(803) 855-1581

CIRCLE NO. 49

VME SYSTEM RACKS

EUROCARD/VME
AC SWITCHING
45 - 100 WATT

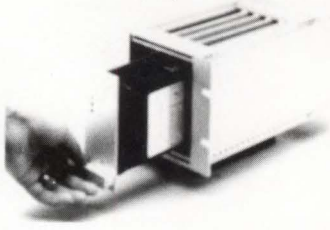

FEATURES

- 3U Packages to include 6, 8, and 10TE
- Meets VDE Safety Requirements
- Power Fail Signal
- Remote Sensing

TOTAL POWER INTERNATIONAL, INC.
418 Bridge Street, Lowell, MA 01850 Telephone (508) 453-7272
FAX (508) 453-7395 TLX 948617

Send for Free Catalog

CIRCLE NO. 50

ARE YOU THE TYPE TO SETTLE FOR HALF RIGHT?

COMPUTER DESIGN reaches more Design Engineering Managers and Senior Design Engineers than any other industry publication. COMPUTER DESIGN's readers are the primary buyers ... the OEM systems people who decide on the technology ... guide the design ... choose the vendors ... and direct how it all goes together. *COMPUTER DESIGN's subscribers are 100% system design and integration qualified.*

If someone tells you that any other design publication reaches as many of *these* key managers ... at best, they're only *half* right.

Only a little more than half of COMPUTER DESIGN's readers receive the leading "how-to" publication; less than half receive the leading "news" tabloid. And as many as 79% do not receive some of the other magazines in the field.

PUBLICATION	PERCENT CD SUBSCRIBERS RECEIVING
ED	52.1%
EDN	51.4
EET	45.1
M-M Systems	39.7
EP	36.2
ESD	32.3
Electronics	30.5
VLSI System Design	21.1

Source: 1988 COMPUTER DESIGN Media Evaluation Survey conducted by The Market Information Center, Marlboro, MA.

So, if your goal is to reach more qualified design managers and project leaders, don't pay full price and settle for *half.* Choose COMPUTER DESIGN and know you're right. All right. All the time.

Why is COMPUTER DESIGN able to attract more leaders as readers? The answer is WHY! Some publications tell their readers "how-to," and others provide industry news. Only COMPUTER DESIGN zeroes-in on WHY. WHY one technology is more suitable than another. WHY to build or buy. WHY to choose one cost/performance tradeoff over another. And more.

Find out WHY this unique editorial makes advertising in COMPUTER DESIGN the better buy — every time.

Go ahead.

COMPUTER DESIGN®

One Technology Park Drive
P.O. Box 990
Westford, MA 01886
(508) 692-0700

THE FIRST MAGAZINE OF SYSTEM DESIGN, DEVELOPMENT AND INTEGRATION

Published by PennWell Publishing

COMPUTER DESIGN®

THE FIRST MAGAZINE OF SYSTEM DESIGN, DEVELOPMENT AND INTEGRATION

Continues to Deliver *more* Buyers of Board-Level Products than the Competition

COMPUTER DESIGN is committed to providing its 100,000+ design engineering managers and senior design engineers with the leading-edge information on significant technologies, products and design options that help them turn their ideas into competitive reality. COMPUTER DESIGN targets areas of critical importance to systems designers, such as bus/board technology and products, with solutions-oriented editorial. That's why COMPUTER DESIGN continues to reach more people who specify, purchase or influence the purchase of board-level products than the competition.

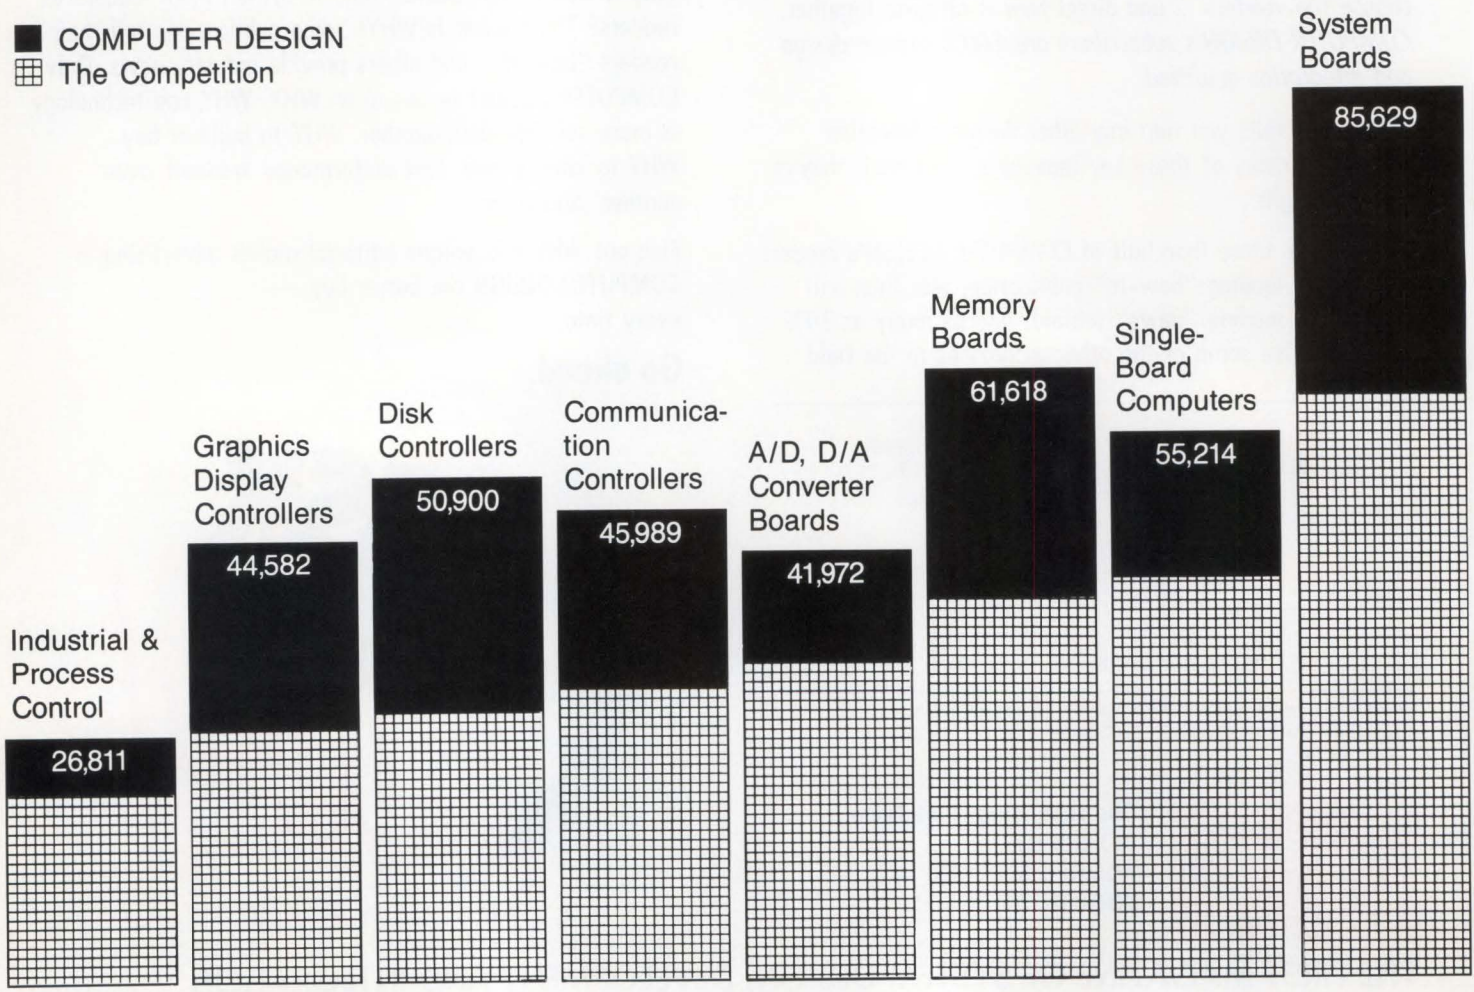

SOURCE: June 1988 BPA statement.

CIRCLE NO. 54

16- or 32-bit operations via the T9506 and I/O port.

The incorporation of a DT-Connect-compatible image bus makes the IP 9506 PC AT easy to use with third-party boards, such as the AT-compatible frame-grabber and image-buffer boards supplied by Data Translation. The DT-Connect bus differentiates the AT version of the processor board from the VME and Multibus versions, which have proprietary image buses.

Software support

Software for the IP 9506 PC AT will include T9506 microcode routines, a subroutine library and an interactive image-processing package. In addition, a microcode development system will soon be available, letting users customize the routine as well as write their own microcode. Software developed on the AT version will be fully compatible with the VME and Multibus versions.

An agreement with Media Cybernetics (Silver Spring, MD) will result in an integrated software suite consisting of that company's Image-Pro interactive processing package and the Image-Pro Developers' Toolkit. The Toolkit is a library of 200 callable routines that include thresholding, edge detecting and spatial filtering. The interactive Image-Pro is a menu-driven, general-purpose image-processing package that includes many operations contained in the Toolkit. In addition, Media Cybernetics will be incorporating an interactive FFT module specifically coded for the T9506.

The Toshiba IP 9506 PC AT image processor is available in limited quantities now, with production quantities scheduled for July. It will be priced at $5,995 in single quantities. The Image-Pro package from Media Cybernetics is due in the third quarter of 1989. The menu-driven interactive package is expected to be priced at $2,000, the Developers' Toolkit at $1,000, and the interactive FFT module at $1,500. —*Tom Williams*

Toshiba America
OEM Division
9740 Irvine Blvd
Irvine, CA 92718
(714) 583-3000

Circle number 479

Intelligent VMEbus board controls up to six stepper axes

An intelligent motor controller for the VMEbus, the VMEX from Oregon Micro Systems allows control of up to six axes in one slot of a CPU. The ability to control six axes breaks new ground for a VMEbus product, according to Wayne Hunter, president of the company. "We had a lot of customer interest in a VME product, so we decided to put our efforts there. Our nearest competitor has a single-axis board, so this is a substantial step forward for sophisticated motor control applications."

Using a proprietary design technique and a 68000 miocroprocessor, the board provides control for stepping, linear or servo motors and delivers pulse rates of up to 524,000 steps/s with a position range of 134 million pulses for each axis. All of the axes are controlled through five ports for control, status feedback, data, commands and the VMEbus interrupt vector.

Each axis has a separate command queue, letting the host computer transfer a command string and then proceed with other tasks while the VMEX manages the motion process. The computer can be interrupted at any point in the command stream to coordinate the motion process with other activities. Each axis can perform unrelated moves or can be coordinated with the others as required by the application.

High-resolution microstepping

The VMEX supports high-resolution microstepping of 50,000 steps/revolution with a standard 200-step/revolution stepping motor by developing the high pulse rates required for these applications. This high resolution lets the motor run smoothly at all speeds and minimizes low-speed torque loss due to mechanical resonance effects. The velocity streaming mode allows arbitrary move contouring under control of the VMEbus computer.

Velocity contouring with circular and linear interpolation on two axes provides constant tool control for machining applications. The VMEX will interpret an input sequence and build a queue of moves to later generate a continuous path at constant velocity. Overtravel limit switches and home sensor inputs for each axis are provided for ease of system implementation.

Simple ASCII commands can be programmed from any high-level language, including Basic, Pascal and C. Additional user-definable I/O

(continued on page 106)

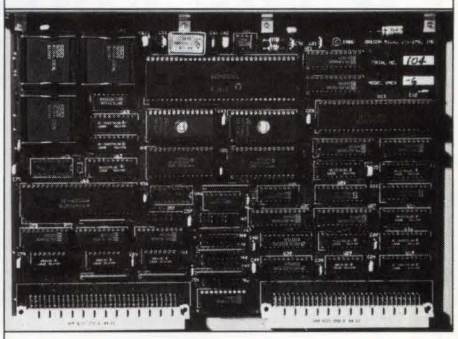

Magnesium For Computer Components

Magnesium is one-third lighter than aluminum, has an exceptional strength-to-weight ratio, and is easily machined to meet your custom requirements.

Versus die cast components there's no comparison: Extruded shapes produce better parts, allowing better machinability, due to less porosity in their structure and startup costs are significantly less.

For computer components, extruded magnesium is the ideal solution.

KAISER ALUMINUM
SPECIALTY PRODUCTS
7311 East 41st Street
Tulsa, Oklahoma 74145-4505
(918) 627-0100 • FAX (918) 664-1645

CIRCLE NO. 52

COMPUTERS AND SUBSYSTEMS

lines can be used to monitor or initiate other events and are under control of the host computer. Complex move sequences, time delays, status checks and control of other external events can be programmed through the VMEbus interface.

Incremental encoder feedback is available as an option on two axes for applications that require precise position feedback or correction. The encoder option can correct for position errors, monitor for slip or stall, or allow tracking of one motor with another.

The company's motion controls are programmed with double ASCII commands through an extensive command structure. When combined into character strings, these commands can be used to create sophisticated motion profiles. The board hosts a 200-command and -parameter buffer for each axis and a command loop counter that allows multiple executions of any command string without host intervention. Each axis can nest loops in the command sequence up to four levels deep. A flag can be passed to the host on the completion of a sequence or at any intermediate point in the command stream.

Acceleration to an optimum speed is provided by the VMEX, followed by constant velocity pulses and controlled deceleration to a stop. Actual position can be simultaneously monitored by use of the encoder feedback option. The VMEX creates a velocity profile by calculating the optimum velocity 1,024 times/s, providing a smooth acceleration. This calculation is used to control a variable frequency pulse train derived from a crystal oscillator, resulting in accurate pulse rates. Linear as well as parabolic and cosine velocity ramps can be generated to fit a variety of system requirements.

The VMEX is available now in two-, four- or six-axis configurations. The two-axis configuration is priced at $1,145. —*Michael Donlin*

Oregon Micro Systems
1800 NW 169th Pl
Beaverton, OR 97006
(503) 644-4999

Circle number 480

IF YOU WRITE IN C, THEN YOU SHOULD DEBUG IN C

The ultimate in time savings is obtained when you debug your code in the same language it was written. Code development is accelerated as constant program printouts are no longer necessary. All displays of your program, including the real-time trace buffer, are in the form you specify, with options for Source only, Source and assembly or assembly only. Use your favorite C or PL/M compiler with our emulation system and SourceGate™ to enhance productivity of your engineering department. If you are working with different microprocessors, SourceGate provides the same interface for each, so learning curves are almost nonexistent when switching between projects or processors.

HMI enhances this software capability with the most advanced line of in-circuit emulators on the market today. SourceGate runs on all IBM PC family computers, Sun Workstations and many Unix systems.

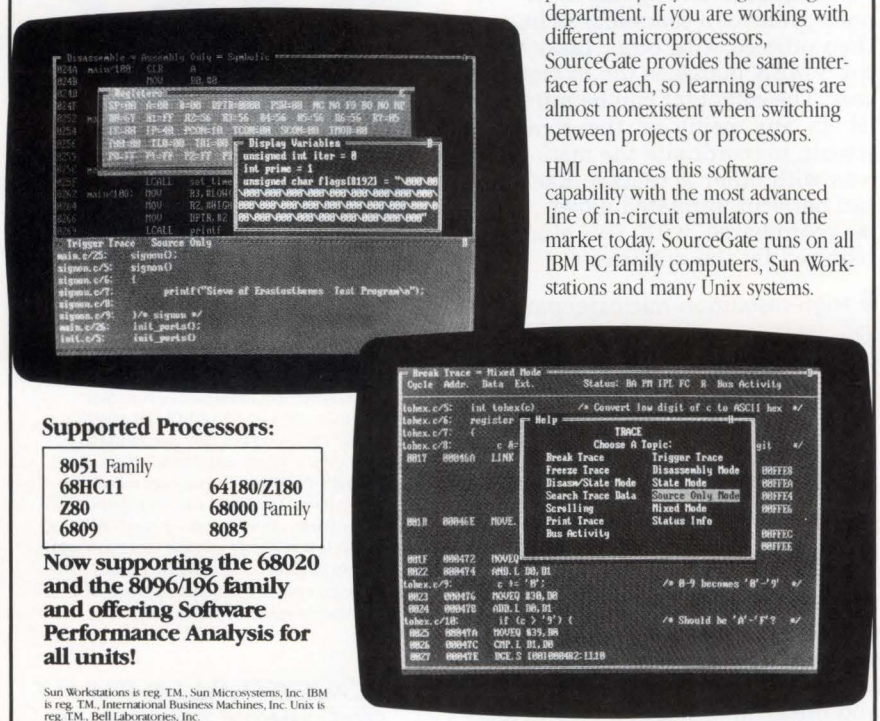

Supported Processors:

8051 Family	
68HC11	**64180/Z180**
Z80	**68000** Family
6809	**8085**

Now supporting the 68020 and the 8096/196 family and offering Software Performance Analysis for all units!

Sun Workstations is reg. TM., Sun Microsystems, Inc. IBM is reg. TM., International Business Machines, Inc. Unix is reg. TM., Bell Laboratories, Inc.

HMI

For complete details, contact:
Huntsville Microsystems, Inc.
P.O. Box 12415, 4040 South Memorial Parkway, Huntsville, AL 35802
(205) 881-6005 TWX: 510-600-8258 FAX: 205-882-6701

CIRCLE NO. 53

Trademark Information

UNIX is a registered trademark of AT&T Bell Laboratories.

PAL is a registered trademark of Advanced Micro Devices, Inc.

SMART-POWER is a registered trademark of Nartron Corp.

DEFINITY is a trademark of AT&T.

RPM is a trademark of Quickturn Systems, Inc.

This space contributed as a public service.

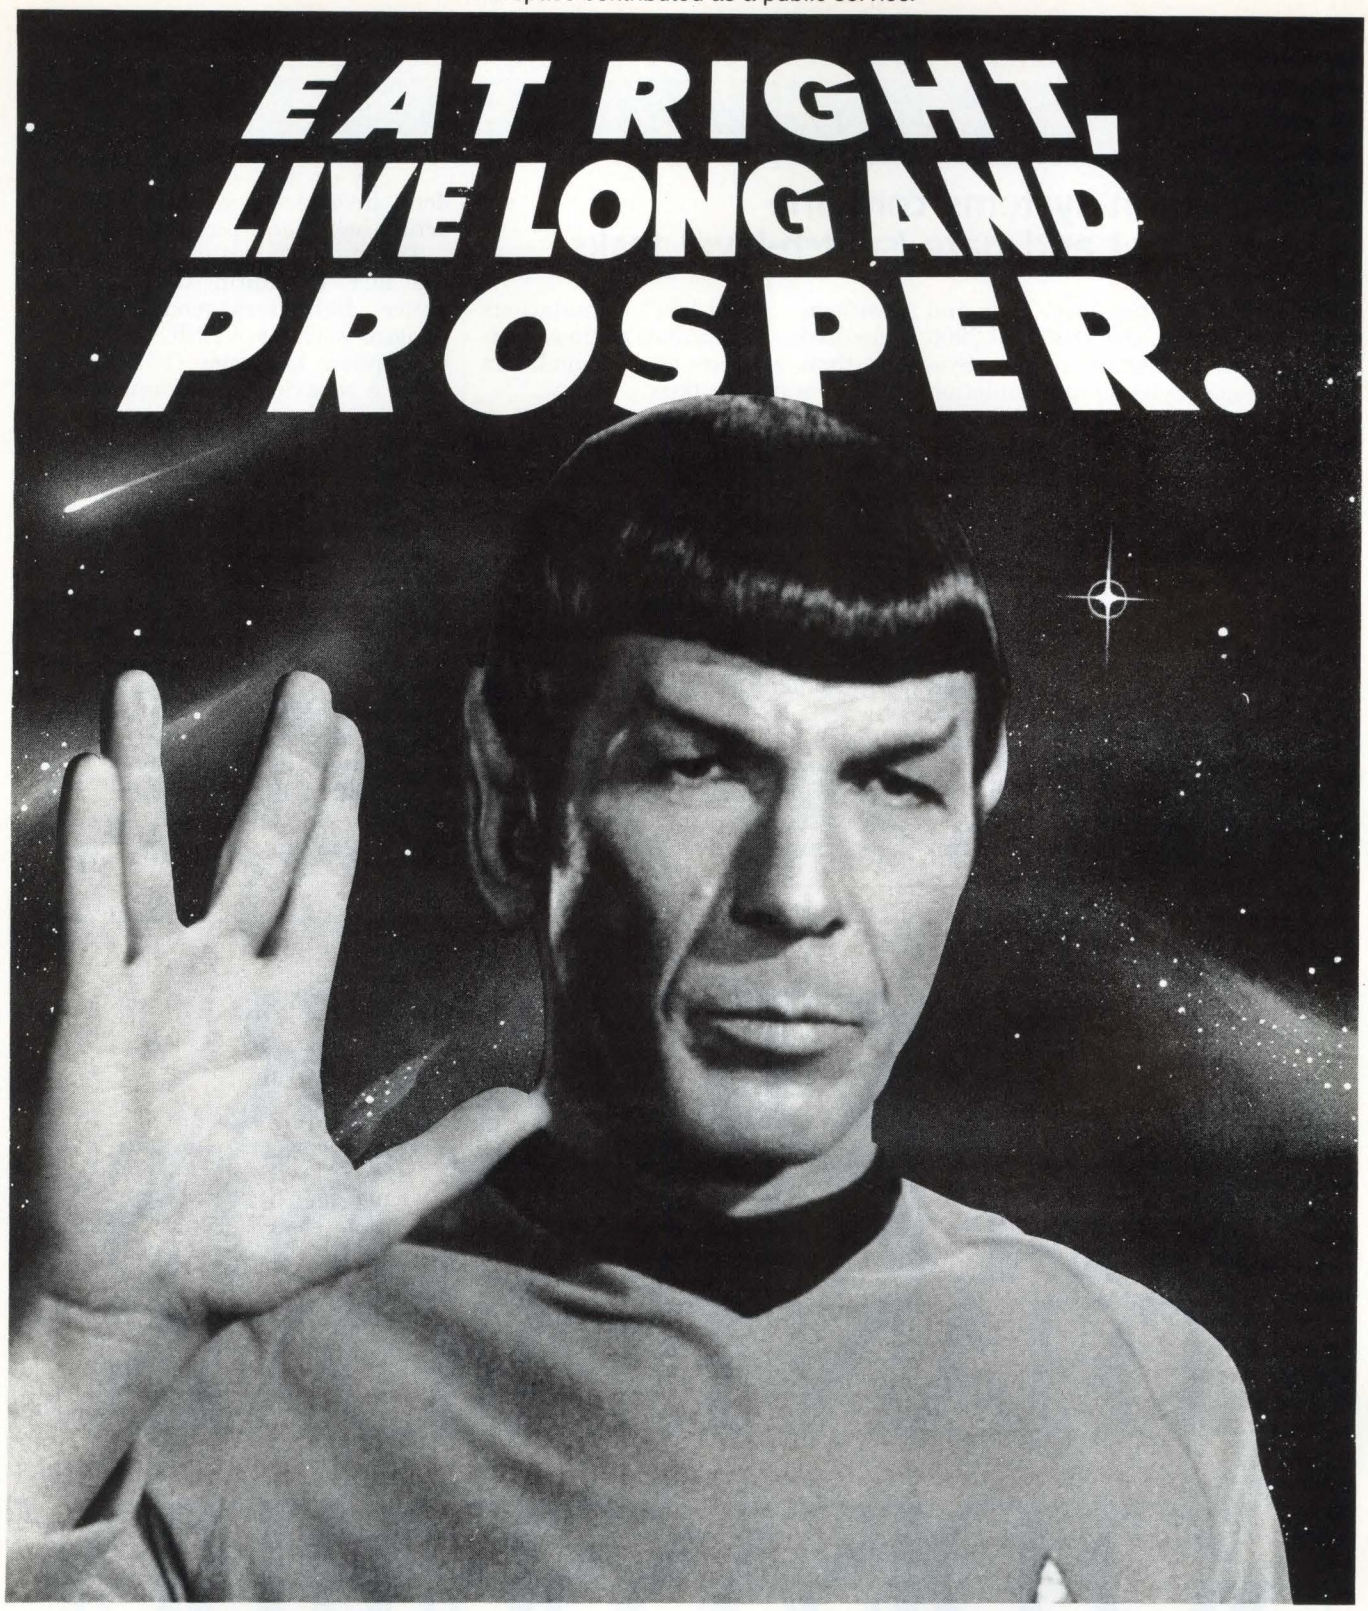

EATING RIGHT IS HIGHLY LOGICAL.

Recommendations:
Eat high-fiber foods, such as fruits, vegetables, and whole grain products.

Eat fewer high-fat foods. Maintain normal body weight. And live long and prosper.

AMERICAN CANCER SOCIETY®

CALL THE AMERICAN CANCER SOCIETY AT 1-800-ACS-2345 FOR FREE NUTRITION INFORMATION.

STAR TREK ® & © 1989 PARAMOUNT PICTURES CORPORATION ALL RIGHTS RESERVED. THE AMERICAN CANCER SOCIETY AUTHORIZED USER.

DESIGN AND DEVELOPMENT TOOLS

Simulation systems combine behavioral and gate-/switch-level tasks

Designed for system-level logic and fault simulation, the 2800 and 2900 Simulation Systems from Ikos Systems couple behavioral simulation with hardware-based gate- and switch-level simulation. The products mark the first time that timing data used in high-performance ASICs can be incorporated in behavioral models and simulated concurrently with gates, according to Ikos. Both tools run on Apollo DN3000 and DN4000 workstations and on Sun 3 and 4 workstations.

Normally, in software-based simulators, mixed-level simulation time is dominated by gate- and switch-level processes. In the 2800/2900 systems, the gate- and switch-level simulation executes in high-speed hardware. As a result, the systems perform mixed-level simulations hundreds to thousands of times faster than software-based simulators do, according to Ikos.

The 2800 system has a maximum gate-level capacity of 320,000 gates, while the 2900 has a 1.2 million-gate limit. The behavioral capacity for both systems is essentially unlimited, since the behavioral models run on the host workstations.

At the heart of both systems is the behavioral coprocessor (BCP), which coordinates communication between the general-purpose computer and the gate-level simulation hardware. The BCP contains a 32-bit RISC machine with 32 Mbytes of real (not virtual) memory. The BCP executes the kernel of a behavioral simulator and runs all overhead tasks that don't involve the execution of the actual behavioral processes.

The most important function of the BCP is multiported time advance. Behavioral-modeling languages, such as VHSIC hardware description language (VHDL), require a behavioral-level time queue, while a hardware simulator requires a gate-level time queue. These time queues must be synchronized at simulation run time.

Without hardware assistance, the gate-level simulation must interrupt the behavioral level at every time advance to determine whether behavioral processing is required. Because the two processes run concurrently, Ikos was able to eliminate the resulting system overhead in its 2800/2900 systems without sacrificing timing accuracy.

The gate-level time-advance logic supports a look-ahead approach where time can advance in "chunks"

20X LOWER UPDATE COSTS THAN EPROMS.

NO DISASSEMBLY.

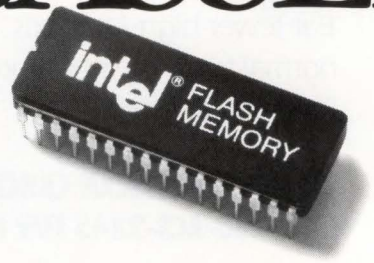

to the next event. A multiport time-control logic function communicates with the time-advance logic on each of the gate-level processor boards and with that of the BCP. This multiport time-control logic function receives the time to the next event from all sources, and coordinates the time advance of all the evaluator boards and of the BCP.

The behavioral environment is implemented with C++. Since VHDL can be translated into C++, Ikos is developing a VHDL front-end for both systems.

The new simulators will be available in the fourth quarter of 1989. Complete base systems start at $95,000. —*Michael Donlin*

Ikos Systems
145 N Wolfe Rd
Sunnyvale, CA 94086
(408) 245-1900

Circle number 490

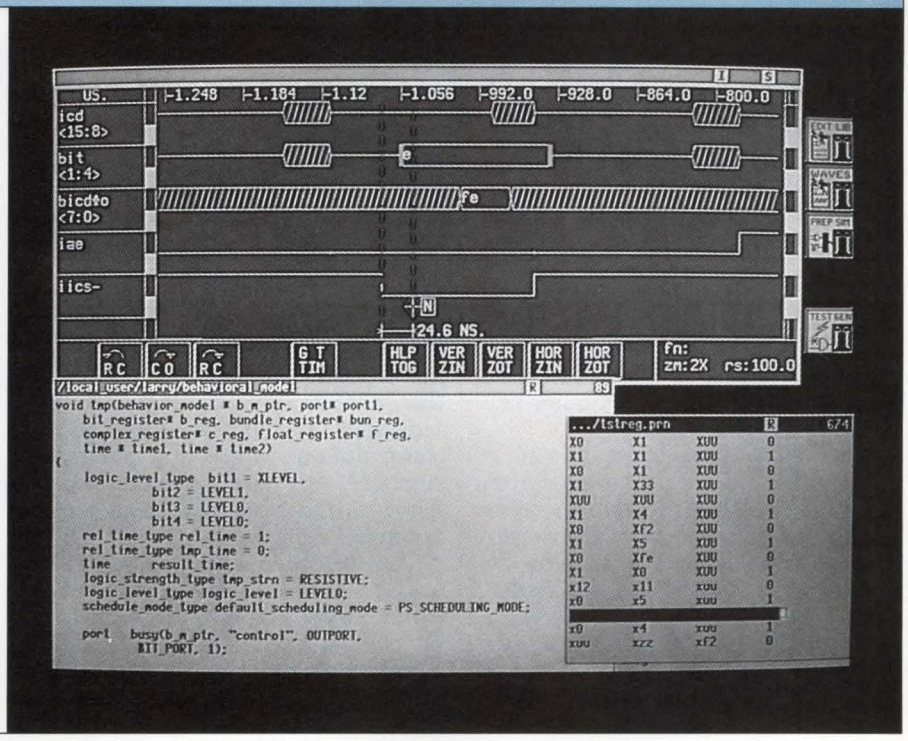

100X THE SPEED OF DISK/DRAM. NO MOTOR.

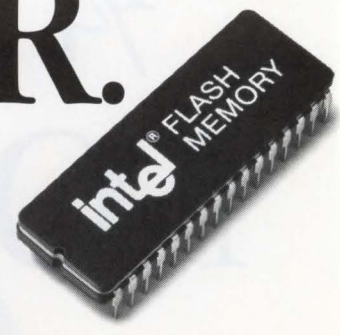

DESIGN AND DEVELOPMENT TOOLS

Logic synthesizer optimizes IC designs for area and speed

The Design Compiler from Synopsys is a set of logic-synthesis tools that targets both full-custom IC and application-specific IC designs. Version 1.1 of the Design Compiler typically produces optimized IC designs that, with no increase in die size, are 10 percent faster than those produced by the original Design Compiler. Version 1.1 also provides more performance enhancements.

The Synopsys package consists of a Design Compiler that optimizes designs for either speed or area, an HDL (hardware description language) Compiler that translates designs from a higher-level description into a gate-level description, and a Library Compiler that may be used to translate existing ASIC libraries into Synopsys format. A new Audit Report feature generates a set of design statistics that can be used to calculate the final cost of a design, and to perform a statistical design-rule check on the gate-level implementation.

Design synthesis

The Design Compiler is an optimizer that accepts gate-level design descriptions and engineering design constraints as inputs and generates smaller, faster netlists as outputs. It may also generate a schematic of the optimized netlist.

The optimizer is complemented by a timing verification system. The timing verifier calculates four timing components to determine timing performance. Version 1.1 highlights all gates in the critical timing path when users view a schematic of the optimized design.

In addition, by allowing the repartitioning of circuit blocks, Version 1.1 lets designers selectively optimize segments of an IC design. Blocks can then be individually synthesized through selective optimization strategies.

Benchmarks conducted on real designs that had been manually optimized show an average improvement of 31 percent when optimizing for speed, and an average improvement of about 24 percent when optimizing for area. Most benchmarks show significant reductions in area, even when designs are optimized for speed. The resulting average is 10 percent faster, with no increase in die size.

Gate-level description inputs to the optimizer may come from the HDL Compiler or from other EDA systems. The HDL Compiler accepts

4X THE DENSITY OF AN EEPROM. ¼ THE COST.

NO WEAROUT.

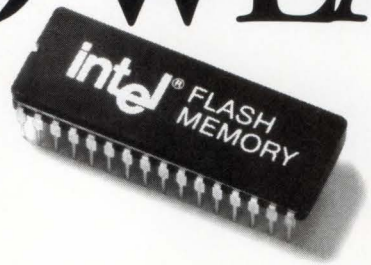

designs in Gateway Design Automation's Verilog HDL and translates them for the optimizer. The HDL Compiler, however, is actually a framework that can support most HDLs.

All Synopsys code is written in C and is designed to run in a Unix X-Window environment. Platform support includes Apollo, Digital and Sun workstations.

Version 1.1 will be automatically distributed to licensed users of the Synopsys Design Compiler. Through recent corporate agreements, the Design Compiler will be available through LSI Logic (San Jose, CA) and Cadence Design Systems (San Jose, CA). *—Bill Harding*

Synopsys
1500 Salado Ave
Mountain View, CA 94043
(415) 962-5000

Circle number 489

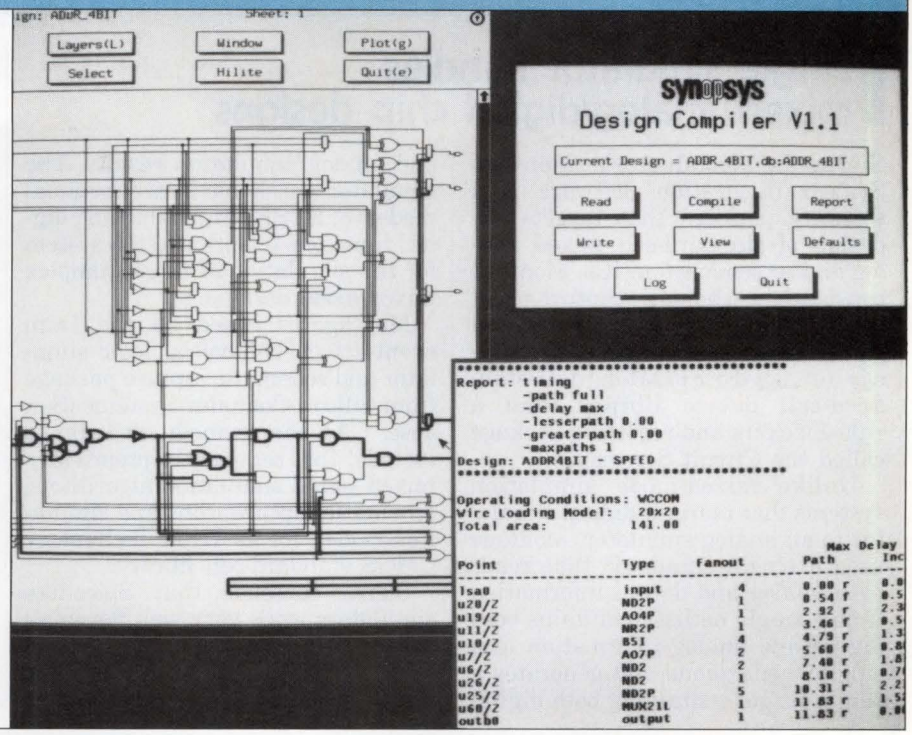

4X THE DENSITY OF AN SRAM. ½ THE COST. NO BATTERY.

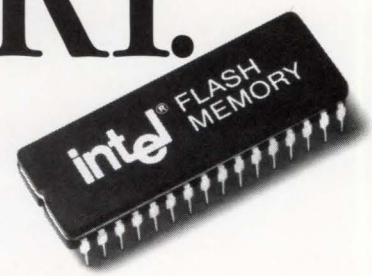

DESIGN AND DEVELOPMENT TOOLS

Single simulator handles mixed analog/digital chip designs

Sierra Semiconductor's Montage System Integration Software is a software package that targets the design of standard-cell mixed analog/digital semicustom ICs. Montage consists of a schematic capture package, an analog/digital behavioral simulator, Sierra's Triple Technology (analog-EEPROM-digital) standard-cell device library, and a rule-checking and reporting package called the Circuit Screener.

Unlike mixed-mode simulation systems that marry a digital simulator to an analog simulator, Montage uses a single simulator that reads both analog and digital information from a single netlist, maintains both digital and analog information in a single event queue and generates a single output containing both digital and analog simulation results. The simulator uses behavioral functional models of EEPROM, analog and digital functions to optimize the system for the simulation of large, complex mixed-mode designs.

Montage is based on the Lsim event-driven behavioral logic simulator and schematic capture package from Silicon Compiler Systems (San Jose, CA). Sierra modified Lsim to include the company's proprietary mixed-signal simulation algorithms, schematic-capture icons and simulation models for its Triple Technology CMOS standard-cell library.

Sierra contends that Spice-like simulators work very well for small designs or for the development of cells for libraries, but that they aren't practical for large, complex designs. Behavioral analog simulators run at least 1,000 times faster than Spice derivatives, so they're more efficient at simulating large designs.

The speed difference between the two is illustrated by a Sierra benchmark, where simulating one conversion in a 10-bit analog-to-digital converter using a Spice simulator took 6 CPU hours on a 4-Mips Elexi computer. Using Montage, the same simulation on the same computer took less than 11 s.

Spice-verified behavioral models

To ensure that analog simulations using Montage produce results as accurate as those produced by Spice-based simulators, Sierra uses a Spice-based analog simulator to develop its behavioral analog libraries. In the initial release, the Montage Triple Technology 1.5-μ cell library contains 245 digital cells, 25 analog

NOW YOU CAN HAVE IT ALL. IN A FLASH.

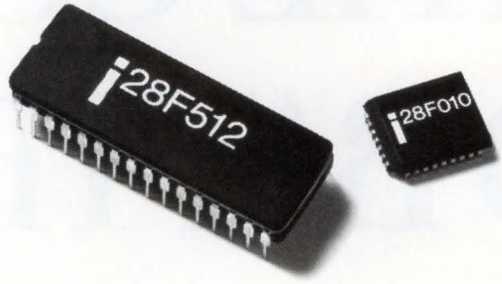

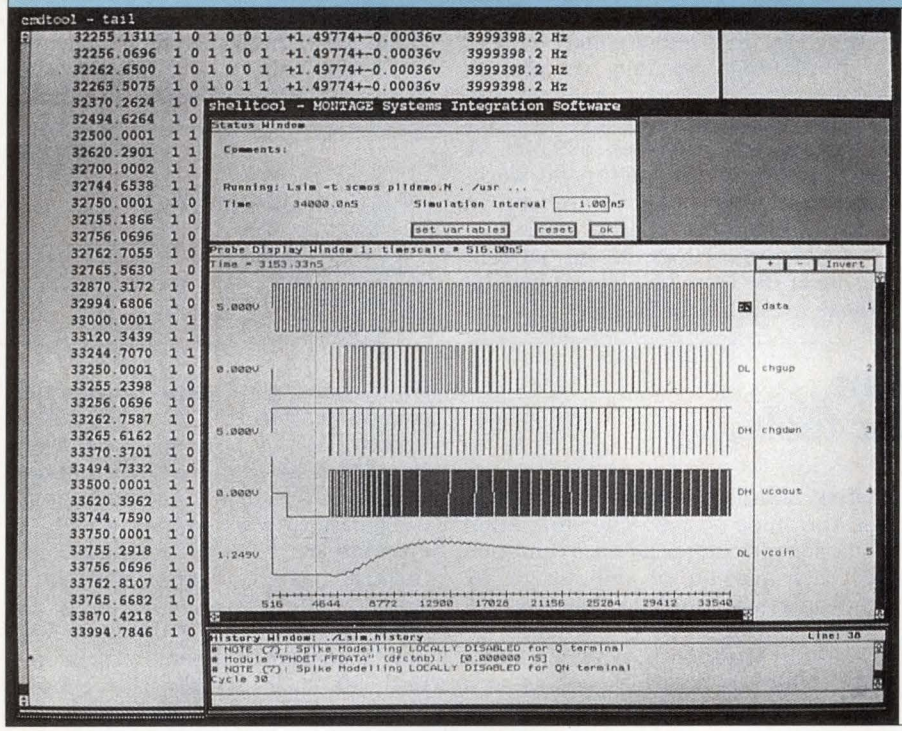

cells and nine EEPROM cells, as well as RAM and ROM compilers. Sierra's confidence in its simulation results is evidenced by the company's guarantee of first-time silicon success for customers using Montage to design mixed-mode semicustom ICs at their own facilities.

Working with error-checking algorithms developed by Sierra, the Montage Circuit Screener analyzes and validates mixed-mode designs. It performs Triple Technology error checking, including connectivity, V_{pp} net checking in EEPROM designs, and loading. In addition to providing error checks, the Circuit Screener supplies users with detailed design reports that include the power consumption, interconnectivity, estimated die size and the total number of gates in the design.

Sierra's Montage uses an open architecture that can be configured to

(continued on page 114)

Intel's new Flash Memories. At first, they might sound too good to be true.

But the more you know, the better they get. Because flash gives you the

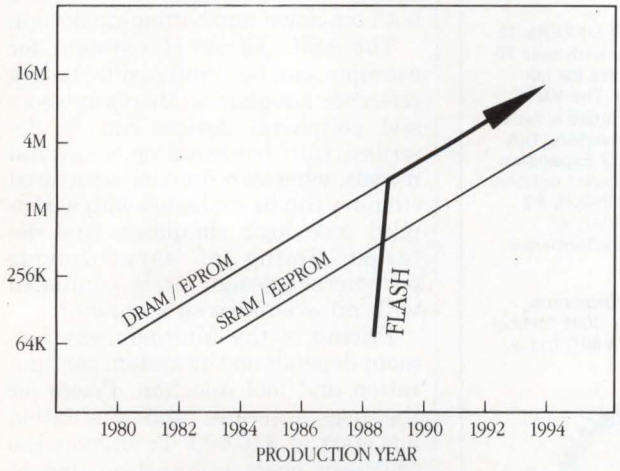

Flash Memories have overtaken EEPROMs and battery-backed SRAMs. They also give you the density of EPROMs, with rewrite capability, or the density of DRAMs with nonvolatility.

advantages of mainstream memory technologies. Without any of the drawbacks.

In one chip, you get high density. Read/write capability. Random access. And nonvolatility. All at the lowest system cost.

In fact, Intel's new 28F010 and 28F512 Flash Memories offer so much that they give you a powerful new alternative to EEPROMs. EPROMs. Battery-backed SRAMs. And disk/DRAM combinations. Whether you compare them on performance. On price. Or on reliability.

No matter how you look at it, this is a memory solution you should learn more about. Before your competition does. So call (800) 548-4725, Lit. Dept. #6A14. And we'll send you everything you need to know. In a flash.

intel®

©1989 Intel Corporation.

CIRCLE NO. 57

DESIGN AND DEVELOPMENT TOOLS

support model libraries from other semiconductor vendors, so customers aren't limited to designing ICs for fabrication by Sierra. Because Montage is based on Silicon Compiler Systems' simulator and schematic capture package, users may expand their systems with any of Silicon's modules that work with Lsim. Montage runs on standard Unix-based hardware platforms such as those offered by Sun and Apollo.

A Montage Systems Integration Software package that runs on Sun Microsystems workstations is available and is priced at $29,500. A package that runs on Apollo workstations is scheduled for release in the third quarter of this year.

A Montage upgrade for current users of the Lsim behavioral simulator is available for $19,950. The upgrade includes the Triple Technology icons and simulation models, the Montage upgrade of Lsim, the Montage Screener and documentation.

—Bill Harding

Sierra Semiconductor
2075 N Capitol Ave
San Jose, CA 95132
(408) 263-9300

Circle number 488

Simulation environment integrates multivendor tools

The Multisim simulation environment from Teradyne EDA allows the integration of electronic design automation tools from multiple vendors, as well as various tools from Teradyne.

Multisim supports mixed-mode, mixed-level simulation and lets different simulators work concurrently on the same design. The simulation bus permits independent evaluation of each portion of the design by multiple simulators operating on different platforms.

Key to Multisim's operation is the Multisim bus, which functions as a simulation backplane to provide the communication link between simulators. The bus protocols have been broadened so that they can handle data from practically any simulator.

The user can access the simulation environment through a common interface based on the Vanguard graphics framework, or through another interface of the user's choice. Any interface is capable of controlling all tools within the chosen framework; and each provides a complete symbolic toolset to link major application tools with a system-independent language.

Verification of complex designs

With Multisim, the user can comprehensively verify the design of systems composed of different technologies created with different design methodologies. It supports both top-down and bottom-up design.

The ASIC portion of a system, for example, can be simulated by native reference simulators. Microprocessors and peripheral devices can be described with hardware or behavioral models, while synchronous structured circuitry can be evaluated with a compiled code logic simulator. And detailed timing of asynchronous portions of a design can be evaluated with an event-driven simulator.

Pricing of the Multisim environment depends on the system configuration and tool selection. Prices for the user interface and simulation bus start at $10,000. To support the Multisim open architecture and to encourage standardization, Teradyne EDA is publishing interface bus specifications. *—Michael Donlin*

Teradyne EDA
5155 Old Ironsides Dr
Santa Clara, CA 95054
(408) 748-7661

Circle number 491

VMEbus ANALOG I/O SOLUTIONS

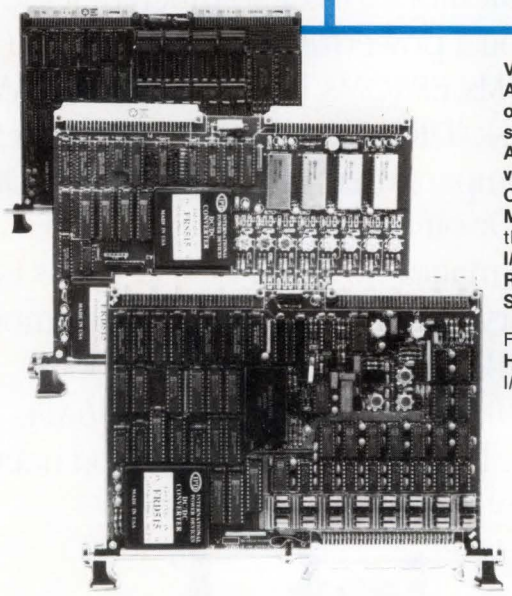

VMIC CURRENTLY OFFERS 13 Analog I/O boards with over 70 off the shelf options for I/O solution flexibility. The VMIC Analog family includes a wide variety of A/D Converters, D/A Converters and A/D Expansion Multiplexers and board options that include Built-In-Test, P2 I/O, and a wide variety of Resolutions and Performance Specifications.

For your High Performance, High Density, Low Cost Analog I/O solutions call VMIC today!

2-YEAR WARRANTY

VMIC VME TM

VME MICROSYSTEMS INTERNATIONAL CORPORATION
12090 South Memorial Parkway
Huntsville, Alabama 35803

1-800-322-3616 Ext. 689 (205) 880-0444 FAX (205) 882-0859

CIRCLE NO. 58

Module with two disk types brings removable storage to the VMEbus

A combined hard- and floppy-disk subsystem that plugs directly into standard VMEbus backplanes, the Mass Storage Module (MSM) from Force Computers can achieve up to 84 Mbytes of formatted Winchester capacity, and 720 kbytes of formatted floppy capacity. The "plug and play" subsystem takes only two backplane slots, and fits into the 6U double-high VMEbus form factor.

There are two models of the MSM, allowing selection of Winchester capacities of 42 Mbytes (MSM-42) and 84 Mbytes (MSM-84). The Winchester is based on the small computer systems interface (SCSI), while the floppy has an SA460 interface. Users whose applications need more capacity than a single MSM can provide may cascade up to seven modules for over 500 Mbytes of storage.

Both MSM models are compatible with the manufacturer's CPU-23, -26, -37 and -30 32-bit single-board computers, as well as with the company's ISCSI-1 intelligent SCSI controller board. All power, control and data transfers take place on the P2 VMEbus connector of the card cage backplane.

The ISCSI-1 includes firmware for control of the drives, but most other operating systems, and many kernels, function with the MSM's interface without requiring special low-level drivers. The only requirements for compatibility with other VMEbus vendors' products are the interface electronics and software that handles the SCSI and SA460 specifications.

The 3½-in. Winchester drives used in the MSM offer an average access time of 19 ms, a synchronous data-transfer rate of 2 Mbytes/s, and an asynchronous data-transfer rate of 4 Mbytes/s. The floppy drive uses standard 3½-in. media and has a transfer rate of 250 kbytes/s.

The user can install the MSM in any VMEbus slot other than slot one, though end-slot installation is recommended. The MSM is designed to slide in and out of the system easily, and features an automatic shipping lock for safe removal and transport.

The MSM-42 is priced at $1,990, and the MSM-84 is $2,490.

—Michael Donlin

Force Computers
3165 Winchester Blvd
Campbell, CA 95008
(408) 374-1146

Circle number 494

COMING IN NEXT MONTH'S
COMPUTER DESIGN MAGAZINE:

CAREER CONNECTIONS

A new advertising section showcasing:

- ☐ Employment Opportunities
- ☐ Outplacement Services
- ☐ Positions Wanted

FOR INFORMATION ON OUR SPECIAL INTRODUCTORY ADVERTISING RATES

Call Nancy Olbers
Recruitment Advertising Manager
(800) 225-0556
(508) 392-2184

August 1 issue "Career Connections"
Closes: July 3

MAJOR SYSTEM COMPONENTS

Controller and TCP/IP software connect G-64 systems to Ethernet

Designed to let G-64-based systems communicate with existing computer resources, the Evlan-11 board from Gespac brings Ethernet communications hardware together with Transmission Control Protocol/Internet Protocol (TCP/IP) support for the OS-9 real-time operating system. The board allows connection of a G-64-based system to any IEEE 802.3 local area network (LAN) through the transceiver, and it connects to either Cheapernet or Ethernet cabling on broadband or baseband networks.

With 64 kbytes of local memory for buffering, the board handles both the physical and the data link layers to free the host processor from communications overhead. Full generation and interpretation of the packet preamble, plus arbitration, are also handled by the on-board processor.

The Evlan-11 is supported by software that makes TCP/IP available for systems running OS-9. File Transfer Protocol and Telnet services are available to the programmer for both file transfer and remote terminal capabilities. A C-compatible 4.2 BSD socket library is available as a single OS-9 data module, or it can be included in a standard Internet database.

Because the Evlan-11 can access TCP/IP, a G-64 system on a factory floor can transfer data to a process engineer's Unix station, to a software engineering system running OS-9, or to board-level systems controlling processes on the factory floor. Production information can be sent to a VAX/VMS running a scheduling program, for instance, and then the VAX/VMS can transfer the information to an IBM mainframe running VM/CMS. All of these communications channels can be dynamically selected on the same cable, and can signal event occurrences in real time.

The board makes effective 10-Mbit/s communication possible between G-64 systems and other popular systems for applications in process control, automatic test and data acquisition.

The Evlan-11 is priced at $1,195 and is available from stock. Single quantity price for the OS-9 TCP/IP driver is $600. —*Michael Donlin*

Gespac
50 W Hoover Ave
Mesa, AZ 85210
(602) 962-5559

Circle number 495

Contact-free soldering and de-soldering

Of SMD, DIP and pin grid components and connectors in a matter of seconds with the Leister-Labor "S"hot-air tool. Electronic control of temperature and air supply. Over 400 special nozzles available.

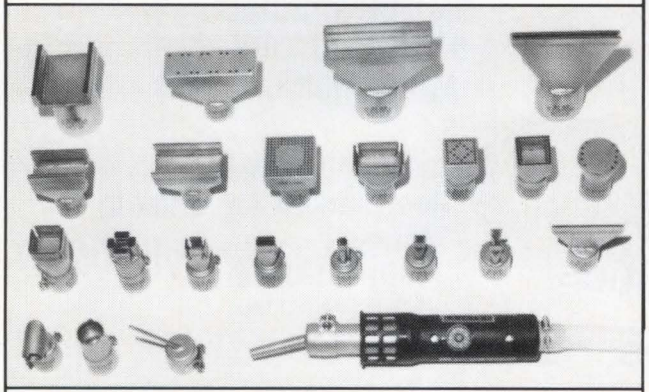

Ask for free brochure UW 154

Brian R. White Co., Inc., 313 Henry Station Road
Ukiah, CA 95482 phone: (707) 462-9795
Farmingdale, NJ 07727 phone: (201) 938-2700

CIRCLE NO. 60

COMPUTER ENGINEER

NIAGARA MOHAWK POWER CORPORATION, an investor-owned electric and gas utility with over 10,000 employees, has a rewarding and challenging position available in our Nuclear Corporate Headquarters located in SYRACUSE, NEW YORK.

The individual we seek must possess a BS Degree in Computer Science or Computer/Electrical Engineering. Background should include knowledge of computer hardware/software and data acquisition systems, combined with a minimum of 5 years experience.

Primary responsibilities will include design, specification, procurement and software verification and validation for nuclear power plant process computer systems, in compliance with industry established guidelines. Experience with PC's, DEC and Honeywell computer systems, computer networking and nuclear computer applications is a plus.

Interested candidates should submit resumes and salary history/requirements, in confidence, to:
MARY KAY MANNS
Supervisor, Employee Relations-Nuclear

NIAGARA MOHAWK POWER CORPORATION
301 Plainfield Road,
Syracuse, New York 13212

Equal Opportunity/Affirmative Action Employer M/F/H

NIAGARA MOHAWK

CIRCLE NO. 61

WIRELESS COMMUNICATION IS NOT A NEW IDEA

Smoldering blankets used to perform amazing feats of communication across long distances. Today, the equipment has altered, but the idea of wireless technology lives on.

Motorola Cellular has utilized the wealth of historical knowledge as a base for technological development. We are seeking engineers with technical know-how, enthusiasm and an inquisitive nature to maintain a legacy of innovation. We have exceptional opportunities for:

SOFTWARE ENGINEERS: Minimum of two years experience in any one of the following: UNIX*, C and Assembly languages; AI/Expert Systems; ISDN; Advanced Communication Networks; Object-oriented Design using C++; Common Channel Signaling System No. 7; Advanced Data Communication Protocols; CASE Environment; Encryption; Packet Switching (X.25).

COMPUTER ENGINEERS: Two or more years experience in any of the following: DSP, ISDN; Advanced Data Communication Networks; Advanced Data Communication Protocols; ASIC Design; CASE/CAE Environment, Encryption, Digital Modulation; Speech Coding; Advanced Error Correction Coding, Packet Switching (X.25), Digital Mobile Communications, Channel Equalizers.

If you are an engineering entrepreneur who wants to help make history rather than merely be a part of it, we need some fresh, creative thinkers to help us with an idea that's as old as the hills.

Send your resume to: **Professional Staffing, Dept. #CD13, Motorola, Inc., Cellular Group, 1501 West Shure Drive, Arlington Heights, IL 60004, or FAX your resume to: (312) 632-5717.** Equal opportunity/affirmative action employer.

MOTOROLA INC.
Cellular Group

Advanced Electronics for a More Productive World.

*A trademark of Bell Laboratories.

CIRCLE NO. 62

SYSTEM SHOWCASE

THE UNIVERSAL PROGRAMMER FOR YOUR PERSONAL COMPUTER

SPRINT™ PLUS

A complete software driven E/EE/PROM, PLD Programmer and development tool in a single PC expansion board!

- Handles all popular NMOS, CMOS, and Bipolar technologies in 20, 24, and 28 pin packages. (32+ pin packages available.)
- Pin-programmable for future device support. Simple, low-cost disk updates.
- Includes Assembler, Unassembler, pin exerciser, screen editor, logic translator, etc.

Call for **FREE** evaluation disk!

adams·macdonald ENTERPRISES INC

800 Airport Road, Monterey, California 93940, Tel (408) 373-3622, Telex 882141, Fax (408) 373-3622, 800-777-1202

CIRCLE NO. 176

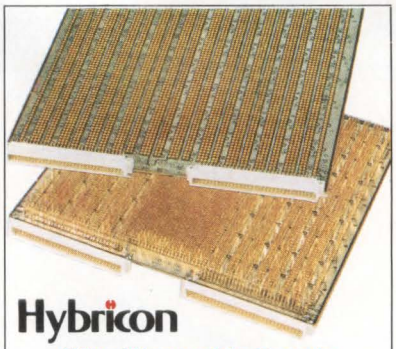

Hybricon

Wire-Wrappable Boards
Multibus II and VME
also SUN VME

ALL SIZES FROM 3U x 160mm TO 9U x 400mm WITH OR WITHOUT PGA - DESIGNED FOR HIGH PERFORMANCE U.H. DENSITY PATTERN ALLOWS OPTIMUM COMPONENT PLACEMENT. SEND FOR DETAILED INFORMATION.

FACTORY (508) 772-5422 — WESTERN OPERATIONS (602) 921-1824

CIRCLE NO. 177

E/EPROM & MICRO PROGRAMMER

$895

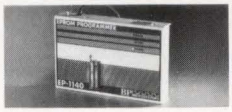

- EP-1140 includes: software, cable, user's manual, 2 free software update coupons, toll-free technical support, one-year warranty & a unconditional 30-day money back guarantee
- Programs 24-, 28-, 32- & 40-pin E/EPROMs
- Supports 874X & 875X series microcontrollers
- Connects to a standard parallel port
- Hardware to support future parts to 8 Megabit
- Software updates on floppy disk
- 32-pin model, EP-1132, available for $695

The Engineer's Programmer™

CALL TODAY 800-225-2102

BP MICROSYSTEMS

10681 Haddington, #190, Houston, TX 77043
713/461-9430 FAX 713/461-7413

CIRCLE NO. 178

IT'S EASY TO LOCATE ROM SOFTWARE SUPPORT

PC-LOCATE: Produce ROM-able code from your "exe" files. PC-LOCATE assigns physical addresses to the re-locatable image based on user inputs. PC-LOCATE supports the entire Intel processor family including the 8086, 8088, 80186, 80188 and 80286.

PC-PROMPAK: A PROM/ROM expansion board for IBM and IBM-compatible computers. PC-PROMPAK provides up to 384 Kbytes of expansion memory and supports most 28-pin JEDEC compatible devices.

PC-ROMDRIVE: Create a "Diskless PC" that can include MS-DOS and your application program. Program execution can be automatically invoked through the use of an "autoexec.bat" file.

ALDIA SYSTEMS INC

P.O. BOX 37634 PHOENIX, AZ. 85069
(602) 866-1786

CIRCLE NO. 179

NEW ZX-386/16 MBI SINGLE BOARD COMPUTER

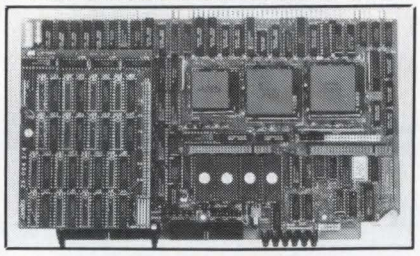

FEATURING:
- On-board 53C90 SCSI
- 8 DMA Channels
- Two 16-bit wide, fully compatible SBX connectors
- Optional "Tru-32" thirty-two bit bus extension

For a free catalog call 1-800-4MB-2DAY
(In California call 1-415-828-3000)

ZENDEX

6700 Sierra Lane, Dublin, CA 94668

CIRCLE NO. 180

Schematic Capture, PCB Layout, Auto Router — $695.00

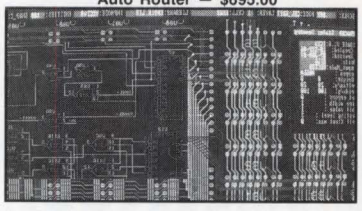

Other Companies talk about Sophistication, but we deliver it. Our full featured CAD software sets industry standard in performance and flexibility with features such as:
- •User definable Menus, macros and user interface
- •User definable grid, pad and trace sizes (1 mil increments)
- •Optimized Rat's Nest
- •DRC with Back annotation
- •Superior Analog design (Irregular filled polygons etc.)
- •Rich editing commands etc etc

Compare our CAD software specification with those so-called sophisticated systems, and if you are that interested, we are that good! Prices start as little as $695.00.
- •30 day money back guarantee
- •Full function Evaluation kit with full user manual $75.00 (Creditable towards purchase of regular system)

Interactive CAD Systems
2352 Rambo Court
Santa Clara, CA 95054
(408) 970-0852

CIRCLE NO. 181

9-Track Tape Subsystem for the IBM PC/XT/AT

Now you can exchange data files between your IBM PC and any mainframe or minicomputer using IBM compatible 1600 or 6250 BPI 9-Track tape. System can also be used for disk backup. Transfer rate is up to 4 magabytes per minute on PCs and compatibles. Subsystems include 7" or 10½" streaming tape drive, tape coupler card and DOS compatible software. For more information, call us today!

QUALSTAR®

9621 Irondale Ave., Chatsworth, CA 91311
Telephone: (818) 882-5822

CIRCLE NO. 182

48 CHANNEL 25 MHz LOGIC ANALYZER $1595.00 + POD PRICE

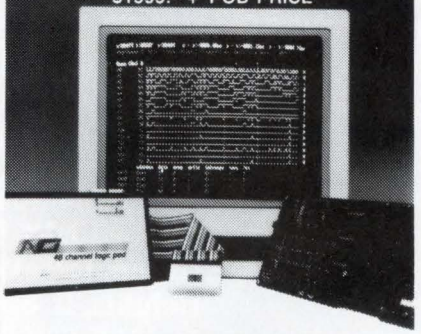

- 48 Chnnls @25MHz x 4K word deep/12 Chnnls @100MHz
- 16 Trigger words/16 Level Trigger Sequence
- Storage and recall of traces/setups to disk
- Disassemblers available for: 68000, 8088, 8086, 80286, 6801, 6811, Z80, 8085, 6502, 6809, 6303, 8031

NCI ☐ 6438 UNIVERSITY DRIVE, HUNTSVILLE, AL 35806
(205) 837-6667

CIRCLE NO. 183

FREE UV ERASER WITH DATA I/O'S 280 SET PROGRAMMER.

For $1595 the 280 gives you:
- 8 sockets for gang/set programming,
- Support for EPROMs and EEPROMs up to 512K,
- Full after-sale support from Data I/O®.

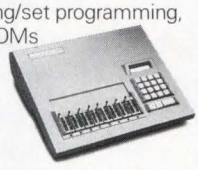

Order a 280 today and get a FREE UV eraser—a $40 value. The Datarase II erases up to 4 EPROMs in 3 minutes.

1-800-247-5700
Ext. 846

DATA I/O Corporation

CIRCLE NO. 184

IEEE 488 Solutions

- Hardware & software interfaces for PC, AT, 386, PS/2, Macintosh, SUN, HP & DEC
- IEEE converters to SCSI, RS-232, RS-422, modem, Centronics, digital I/O & analog I/O
- IEEE extenders, expanders & buffers
- IEEE drivers for DOS, UNIX®, Lotus 1-2-3, Symphony & Quattro.

Call or send for your FREE Technical Guide

Demo disks and application notes available

IOtech

(216) 439-4091
Telex 6502820864 • Fax (216) 439-4093
IOtech, Inc. • 25971 Cannon Road
Cleveland, Ohio 44146

CIRCLE NO. 185

MIL-STD-1553

MIL-STD-1553 Interface Board For IBM-PC/XT/AT or Compatibles

Dual redundant bus controller, remote terminal and monitor with powerful easy-to-use software. **From only $2095**

Call now for data on MIL-STD-1553 and ARINC 429 for PC/XT/AT, VME and VXI

Ballard Technology 1216 N.W. 75th St., Seattle, WA 98117 (206) 782-8704

CIRCLE NO. 186

REAL WORLD I/O
For PC/XT/ATs

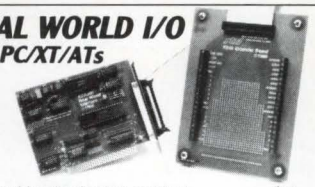

DG24 • 24 line digital I/O; 10 MHz 8255. **$95**
AD500 • 8 channel 12-bit (plus sign) integrating A/D; prog. gains of 1, 10 & 100; 7 digital I/O lines. **$239**
AD100 • Single channel version of AD500; 10 digital I/O lines. Same programmable gains. 700 meg input Z. **$149**
AD1000 • 8 channel 12-bit A/D; 25 uS; sample & hold; 3 5 MHz timer/counters; 24 digital I/O lines. **$295**
ADA300 • 8 channel 8-bit 25 uS A/D; single D/A sample & hold; 24 digital I/O lines. **$239**
AD200 • 4 channel 12-bit 125 uS A/D; 3 5 MHz timer/ counters; 24 digital I/O lines. **$239**
DA600 • Fast settling dual bipolar D/A. **$179**
PD200 • Prototype board w/ address decoder; manual **$99**

All boards include BASIC, Pascal, C, and Forth drivers. 30 day return; 1 year warranty. Call for "Real World Interfacing" application notes.

Real Time Devices, Inc.
P.O. Box 906 State College, PA 16804
(814) 234-8087

CIRCLE NO. 187

ULTIBOARD COMPUTER AIDED PCB DESIGN

The ULTImate PCB layout package featuring:

- Full SMT support
- 32 layer support with blind and buried vias
- Curved traces
- Real-Time Design Rule Check
- Powerful placement aids
- Trace Shove and Reroute-While-Move
- Autoroute by window, component, or net
- Backannotation to OrCAD, DASH, Schema, ViewLogic

ULTIMATE TECHNOLOGY LTD. **USA/Canada:** NorCross Ga. USA • tel. (404) 242-0125/fax 6382 **Europe:** Naarden Holland • tel. (+31) 2150-44424/fax 43345

ASK FOR YOUR FREE DEMO DISK

See us at DAC Las Vegas, booth 1907

CIRCLE NO. 188

Don't Get Zapped!

High inrush current can **destroy** your sensitive VAX CPUs and peripherals **in less time than it takes to flip a switch.**

THE SOLUTION?
Power up with ***Z-LINE*** **TPC 115-10 MTD™ the smallest** power distribution and control system available.
POWER UP WITH — — —

Z-LINE

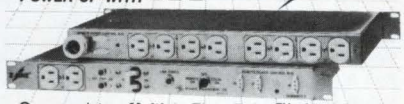

Our proprietary **Multiple Time Delay™ circuitry sequences** your power-up to **protect** your systems from the spikes and surges, EMI & RFI, that **destroy your hardware and erase your data.** And our remote on/off and emergency shutdown **gives the power control back to you.**

All Pulizzi Engineering MTD™ **controllers are compatible with DEC and UPS systems. PRICES FROM $436 TO $305**

DON'T WAIT UNTIL IT HAPPENS, CALL TODAY!
PULIZZI ENGINEERING INC.
3260 S. Susan Street, Santa Ana, CA 92704-6865
(714) 540-4229 FAX (714) 641-9062

CIRCLE NO. 189

AutoKey,
400 Macro Keys

- Plugs inline with your PC Keyboard.
- Uses no system memory.
- Requires no software.
- Has its own microcomputer.
- Eliminates complicated memory resident macro programs such as SuperKey and ProKey.
- Introductory prices $99.00

Mextel Corp. 159 Beeline Rd.
Bensenville, Illinois 60106
Phone (312) 595-4146.

M/C VISA AMEX

Superkey and Prokey are trademarks of their respective manufacturers

CIRCLE NO. 190

PROFESSIONAL OPPORTUNITIES

CIRCULATION 100,000 +

Target your recruitment message to system design engineers — with no waste circulation.

If you need engineers and engineering managers who are skilled in the design of computer or microprocessor systems, **COMPUTER DESIGN** *should be the number one choice for your recruitment advertising.*

Published Twice Monthly — Closing Dates Every Two Weeks

ADVERTISERS INDEX

ADVERTISER	PAGE NO.	CIRCLE NO.
Acromag Inc.	22	12
Adams-MacDonald Int'l.	118	176
Aldia Systems Inc.	118	179
AMP	4,5	3
Apollo Computer Inc.	88, 89	51
Applied Microsystems	65	33
	64A, 64B	
*Arrow Electronics	99	46
Ballard Technology	119	186
†Bayer AG	98, 99	44
BP Microsystems	118	178
Brian R. White Co. Inc.	116	60
Cadence Design Systems	24, 25	14
Ciprico	Cover 2	26
Cypress Semiconductor Corp.	9	
Daisy/Cadnetix	48A-48L	
Data I/O	118	184
Diversified Technology Inc.	23	13
	35	18
DY-4 Systems Inc.	15	8
	37	19
EDI	77	38
Electronic Solutions	1	1
*Electronic Trend Publ.	98	45
Emulex Corp.	Cover 4	64
Excelan, Inc.	58	32
Force Computers	55	30
Heurikon Corporation	52	28
Hewlett Packard	30, 31	16
Huntsville Microsystems	106	53
Hybricon Corp.	118	177
Imprimis Technology	70, 71	35
Inmos	91	41
Intel	51	27
	53	29
	108-113	57
Intelligent Instrumentation	100	47
Interactive CAD Systems	118	181
Intermetrics	29	15
IOTech	119	185
John Fluke Manufacturing Co.	39	20
Kaiser Aluminum	105	52
Matrix Corp.	83	39
Mentor Graphics Corp.	50	
Metheus Corp.	44	66
Mextel Corporation	119	190
Microtec Research	43	23
Mitsubishi Electronics America	21	11
	92	42
Motorola, Inc. Cellular Group	117	62
National Semiconductor Corp.	56, 57	31
NCI	118	183
Niagara Mohawk Power Corp.	116	61
NTH Graphics	33	17, 68
Oasys	18	65
OKI Semiconductor	12, 13	7
Omnibyte Corp.	96A, 96B	119,120
Personal CAD Systems	40	21
PLX Technology	47	25
Pulizzi Engineering Inc.	119	189
Qualstar Corp.	118	182
Radstone Technology	85	40
Real Time Devices	119	187
Sealevel Systems Inc.	102	49
Software Development Systems	69	
Standard Microsystems	2	2
Switchcraft, Inc.	75	37
Systech	Cover 3	63
Tadpole Technology, Inc.	11	6
Texas Instruments	16, 17	9
	16A, 16B	
TITN	41	22
Toshiba America, Inc.	6,7	4
Total Power	102	50
Ultimate Technology	119	188
Ultraview	45	24
VME Microsystems Int'l. Co.	114	58
Wind River	19	10
Xycom	66	–34
Yamaha Systems Technology	95	43
Zendex	118	180

*Domestic Only Issues †International Only Issues

The Advertisers Index is published as a service. The publisher does not assume liability for errors or omission.

COMPUTER DESIGN

THE FIRST MAGAZINE OF SYSTEM DESIGN, DEVELOPMENT AND INTEGRATION

Executive Office: One Technology Park Drive, P.O. Box 990, Westford, MA 01886
(508) 692-0700 FAX:(508) 692-0525 Telex: 883436

Publisher, David L. Allen (508) 392-2111

Associate Publisher, John C. Miklosz (508) 392-2114
Marketing Services Manager, Mary M. Gregory (508) 392-2110
Administrative Assistant, Wendy Ross (508) 392-2112
Ad Traffic Manager, Adrienne Adler (508) 392-2198

Postcards/System Showcase, Shirley Athorn (508) 392-2183
Recruitment, Nancy Olbers (508) 692-0700
Circulation Director, Bob Dromgoole (508) 392-2173
List Rental, Bruce McCarthy (508) 392-2176
Circulation Manager, George Andrew P.O. Box 3466, Tulsa, OK 74101 (918) 831-9443

SALES OFFICES

New England/Upstate New York/ Florida/Eastern Ontario/Quebec
Tim Pritchard
One Technology Park Drive
P.O. Box 990
Westford, MA 01886
Tel: (508) 392-2217
Fax: (508) 692-0525; Telex 883436

New York/New Jersey/ Pennsylvania/Delaware/ Maryland/North-South Carolina/ Georgia
Neil Versen
Park 80 West, Plaza Two
Saddle Brook, NJ 07662
Tel: (201) 845-0800
Fax: (201) 845-6275

Upper Midwest/Western Ontario/ Manitoba/Saskatchewan
Robert D. Wentz
9501 West Devon, Suite 203
Rosemont, IL 60018
Tel: (312) 518-9191
Fax: (312) 698-2675

South Central
Eric Jeter
19627 I-45 North
Suite 220
Spring, TX 77388
Tel: (713) 353-0309; (713) 288-8963

Southern California/Southwest
Tom Boris, Greg Cruse
2232 SE Bristol, Suite 109
Santa Ana Heights, CA 92707
Tel: (714) 756-0681
Fax: (714) 756-0621

Northern California/Nevada
Tom Boris, John Sly, Bill Cooper
1000 Elwell Court, Suite 234
Palo Alto, CA 94303
Tel: (415) 965-4334
Fax: (415) 965-0255

Oregon
Tom Boris
2232 SE Bristol, Suite 109
Santa Ana Heights, CA 92707
Tel: (714) 756-0681
Fax: (714) 756-0621

Washington
John Sly
1000 Elwell Court, Suite 234
Palo Alto, CA 94303
Tel: (415) 965-4334
Fax: (415) 965-0255

U.K./Scandinavia
David Round
69 Imperial Way
Croydon,
Surrey CRO 4RR, England
Tel: 01 686 7655 Telex: 938420
Fax: 01 688 2134

France/Belgium/S. Switzerland/ Spain/The Netherlands
Daniel R. Bernard
247, Rue Saint Jacques
75005 Paris France
Tel: (1) 43 54 55 35
Telex: 214235F MISIVZ
Attn: Missitex: PENNWELL
Fax (1) 47 07 59 01

W. Germany/Austria/ N. Switzerland/ Eastern Europe
Johann Bylek
Verlagsbuero Johann Bylek
Stockaeckerring 63
D-8011 Kirchheim/Muenchen
Federal Republic of Germany
Tel: 089 903 88 06
Telex: 529355 vbb d

Italy
Luigi Rancati
Rancati Advertising
Milano San Felice Torre 5
20090 Segrate Italy
Tel: 2 7531445
Telex: 328601 RANCAD I
Fax: 02 7532354

Japan
Sumio Oka
International Media Representatives, Ltd.
2-29 Toranomon 1-chome
Minato-ku, Tokyo 105 Japan
Tel: 03-502-0656
Telex: J22633 MEDIAREP
Fax: 03-591-2530

Southeast Asia
Anne Goh-Taylor
Seavex, Ltd.
400 Orchard Rd.
10-01, Orchard Towers
Singapore 0923
Republic of Singapore
Tel: 734-9790
Telex: 35539 SEAVEX RS
Fax: 732-5129

Jay G. Seo
Yong-Jin Park
Doobee International, Ltd.
Center Building (Byulgwan)
1-11 Jeong-dong, Choong-ku
CPO Box 4557
Seoul, Korea
Tel: 776-2096
Telex K27117 DOOBEES
Fax: 755-9860

COMPUTER DESIGN (ISSN-0010-4566) is published twice monthly, by PennWell Publishing Company, 1421 S. Sheridan, Tulsa, OK 74112. Second class postage paid at Tulsa, OK 74112 and additional mailing offices. Editorial offices are located at One Technology Park Drive, P.O. Box 990, Westford, MA 01886. Annual subscription price: $70 in the U.S.A., $110 elsewhere. Call (918) 831-9401 for subscription information. Microfilm copies of COMPUTER DESIGN may be purchased from University Microfilms, a Xerox Company, 300 North Zeeb Rd., Ann Arbor, MI 48106. POSTMASTER: Send change of address form to COMPUTER DESIGN, Circulation Department, PO Box 3466, Tulsa, OK 74101.

© 1989 COMPUTER DESIGN by PennWell Publishing Company. All rights reserved. No material may be reprinted without permission from the publisher. Officers of PennWell Publishing Company: Philip C. Lauinger, Jr., Chairman and Chief Executive; Joseph A. Wolking, President; John Ford, Senior Vice-President; Carl J. Lawrence, Senior Vice-President; John Maney, Vice-President/Finance; Leslie P. Cypret, Vice-President/Administration; Steve Zimmerman, Vice-President/ Corporate Services.

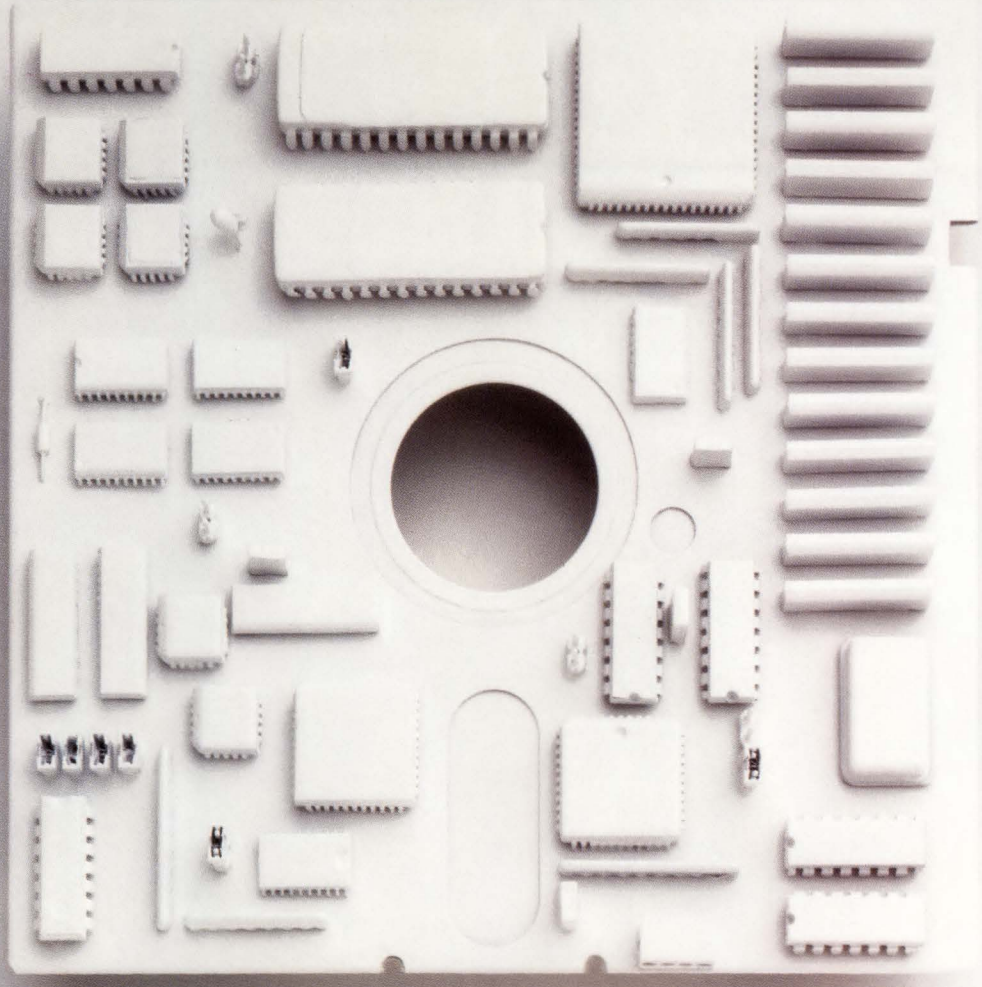

IF YOU'RE SPENDING TOO MUCH TIME ON SNA DESIGN, YOU'RE LEAVING OUT ONESTEP.™

Most likely, if you need SNA compatibility, you'll buy the software source code from one company and hardware from another. Then you'll spend several months desperately working out the integration.

Or worse, you'll design it from scratch.

Systech knows you've got better things to do with your time. So we've created a way to cut down that lengthy process. With OneStep.™

OneStep is a complete SNA/DCP board package. Which makes it the quickest and easiest off-the-shelf solution for IBM communication. In fact, with our help, the OneStep can be up-and-running in just five days.

To make this unique integration work, we teamed up with the most reputable and experienced company in IBM compatibility around—Systems Strategies, Inc. With their superb emulation software and Systech's modular design, OneStep gives you the wide-area links you need for VME and Multibus systems. What's more, the OneStep takes a load off your CPU. That way, you have more power and improved system performance overall.

Systems Strategies, Inc.'s extensive variety of popular protocols includes SNA 3270, RJE and LU6.2. BSC 3270 and RJE. Plus, X.25 and QLLC. All thoroughly tested and successfully ported to Systech's hardware.

Along with easy installation, the OneStep will keep your up-front costs down. And its modularity makes it easy to maintain and upgrade.

Most importantly, with the experience and expertise of Systech and Systems Strategies, Inc., you'll be in step with the latest software and hardware developments. And you'll always be OneStep ahead of the competition.

For a demonstration or more information, call Systech at (619) 453-8970. Or write: Systech Corp., 6465 Nancy Ridge Dr., San Diego, CA 92121.

SYSTECH

WE OPEN UP A LOT OF POSSIBILITIES.

OneStep is a trademark of Systech Corporation.

CIRCLE NO. 63

EMULEX ALREADY HAS THE ONE THING VME SYSTEMS NEED FOR SMOOTH SAILING...

MORE

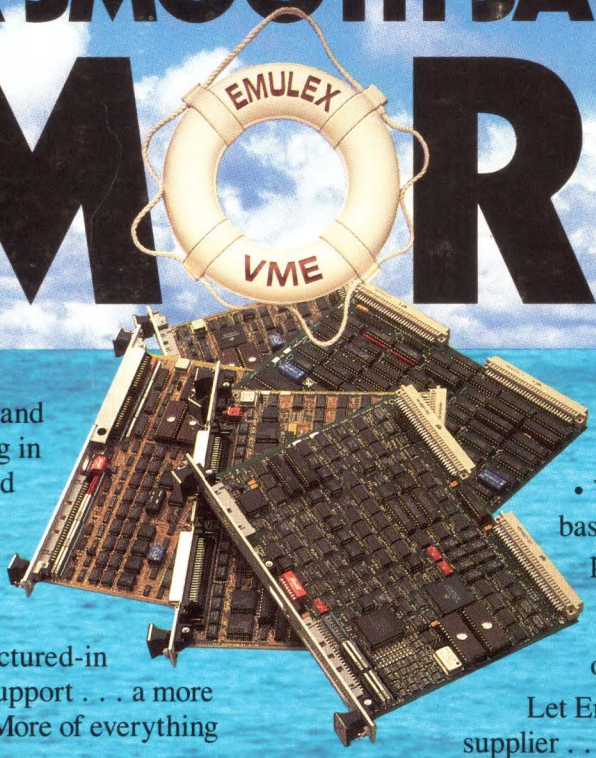

To assure system designers and integrators of smooth sailing in the technologically advanced VME marketplace, Emulex has more of what's needed most. More designed-in quality . . . more VLSI expertise . . . more manufactured-in reliability . . . more OEM support . . . a more dependable supply source. More of everything that's important.

And our "MORE" philosophy applies to every member of our new family of products for the VMEbus.

Introducing:

- VT01 – a 9-track Pertec tape formatter controlling up to eight 1/2-inch tape drives. Tape densities can range from 800 to 6250bpi, tape speeds to 200ips.
- VM21 – a disk controller that can handle four ESDI disk drives with the highest transfer rates available – up to 20Mbps.
- VM31 – a disk controller for four SMD-E disk drives with transfer rates up to 3.0MB/sec and no penalty on performance.
- VS02 – a communications controller designed around a proprietary I/O processor that can handle 16 asynchronous lines at 38.4Kbps simultaneously.
- VH01 – a SCSI host adapter based upon Emulex's own SCSI processor chip that supports synchronous transfer rates of 4.8MB/sec and asynchronous rates of 3.0MB/sec.

Let Emulex be more than your OEM supplier . . . think of us as your OEM partner. That means you can count on us for product quality and reliability, and for stability, capability and integrity. From initial design to final audit, Emulex is dedicated to meeting your needs . . . and MORE.

Call toll-free (800) EMULEX-3, (714) 662-5600 in California.

EMULEX

VMEbus CONFORMANCE CERTIFIED • VME Laboratories

I want to know MORE about VME products.

Name: _____

Company: _____

Address: _____

City: _____ State: _____

Zip: _____ Phone: _____

Emulex Corporation, 3545 Harbor Blvd., Costa Mesa, CA 92626

Regional Offices: Anaheim, CA (714) 385-1685 Atlanta, GA (404) 587-3610 Burlington, MA (617) 229-8880
Chicago, IL (312) 490-0050 Dublin, CA (415) 829-1170 Teaneck, NJ (201) 836-3717 Washington, DC (703) 264-0670
International Offices: Bracknell 344-484234 Munich 89-3608020 North Sydney 2-957-1669 Paris 134-65-9191 Tokyo 3-234-8951 Toronto (416) 673-1211

©1988 Emulex Corporation

CIRCLE NO. 64