A PENNWELL PUBLICATION APRIL 1992

COMPUTER DESIGN

Technology and Design Directions®

FOR ELECTRONIC ENGINEERS & ENGINEERING MANAGERS

QuickTurn's Kaufman on: Prototyping alternatives

FPGAs race for the gold in product development

Is RISC or DSP best for your application?

Analyzers fine-tune high-performance buses

Special Report on Future Computing:

Fuzzy Logic is anything but fuzzy

Don't miss **SBus News** *in this issue*

MAYBE YOU JUST NEED BETTER DRIVERS.

Stuck in a SCSI data-transfer bottleneck?

A first-class SCSI-2 VME board *could* get you out of that jam... *if* getting that board up and running doesn't leave your whole engineering team in gridlock.

That's why, when you buy a Rimfire SCSI-2 VME board from Ciprico, you don't just get a product. You get a partner.

As your partner, we'll help you write the drivers you need, and provide on-site installation. We'll even work with your engineering team to customize hardware or software for your special application needs.

And you get Rimfire's unmatched performance, quality and features. Like the new Rimfire 3870 and 3590 – the first VME boards with independent two-channel SCSI-2 architecture.

First-class hardware *and* first-class support and service. *That's* how you beat bottlenecks. And *that's* what Ciprico is all about.

TURN TO CIPRICO WHEN YOU'RE IN A SCSI JAM.

CALL 1-800-SCSI-NOW
Ciprico, Inc.
2955 Xenium Lane
Minneapolis, MN 55441
Fax: (612) 559-8799

CALL (44) 635 873666
Ciprico (Europe) Ltd.
7 Clerewater Place, Lower Way
Newbury, Berkshire RG13 4RF England
Fax: (44) 635 871996

CIRCLE NO. 1

CIPRICO

New

10-Layer VME Backplane!

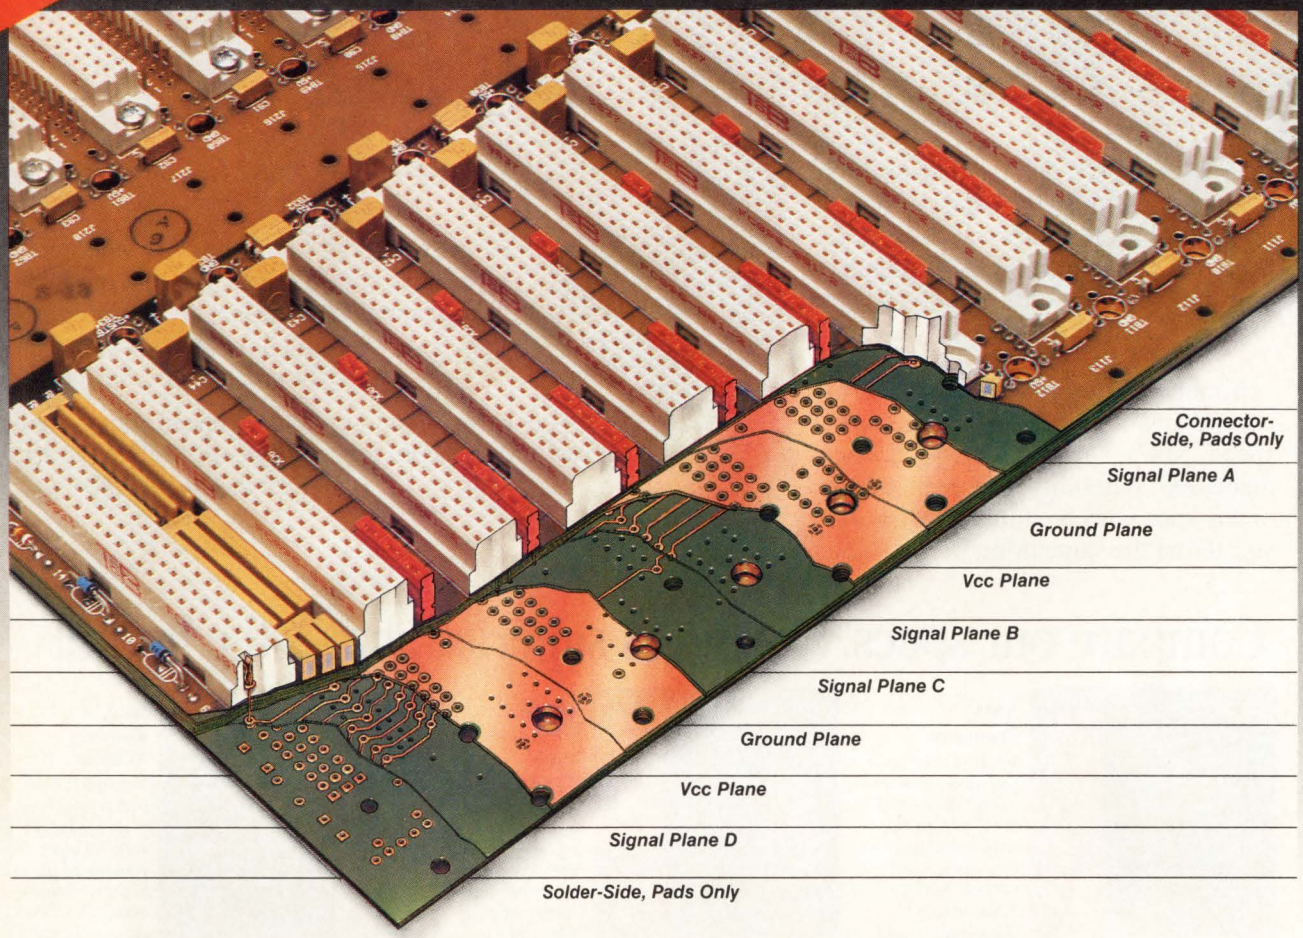

There's Logic To Every Layer!

Look inside the first truly reliable, glitch-free VME backplane and you'll see the logic behind each detail of its design.

Start with the signal layers—four of them—with the most "noise sensitive" lines isolated from the others. The spacing and placement of the layers in this stripline (with embedded microstrip) design was calculated to enhance the backplanes' high frequency bypass characteristics.

It's the quietest backplane ever! Sharply reduced crosstalk is achieved by the increased separation of signal lines—both within and between the layers. Superior decoupling comes from the placement of both ceramic and tantalum capacitors between each slot.

Low inductance, low impedance power distribution is insured by the close proximity of the dual power & ground planes and the strategic placement of power connectors for each voltage across the entire length of the backplane.

There's more logic on the outside. Schottky diode terminations. Interrupt and Bus Grant jumpers accessible from either side. Wire wrap pins with shrouds on the end slots of the J1 and every slot of the J2. Ground pins in between J2 slots to provide I/O cabling with interstitial grounds. And much more.

Choose any size from 2 to 21 slots in separate J1 or J2 or monolithic J1/J2 configurations. Call us today for the inside (and outside) story on the highest performance VME backplanes ever made. It's the logical thing to do.

We'll FAX you the facts.
Want the latest data in a hurry? Nothing is faster than Electronic Solutions' "FAX the FACTS" program. Call our "800" number and give us 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 FAX: 619-452-9494

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

CIRCLE NO. 2

© Copyright 1990, Electronic Solutions

Why Settle for ½ an '040 Board?

You've chosen the '040 because you need maximum performance in your VME system. But look carefully, because other Single Board Computers may only give you only half of what you expected from the '040.

Compare Synergy's SV430 performance to any other SBC. Compare bus speed, MIPs, support, flexibility, documentation, reliability, I/O intelligence or any spec you can think of. We think you'll find the same thing we did—the SV430 outperforms every other SBC on the market by as much as 150%.

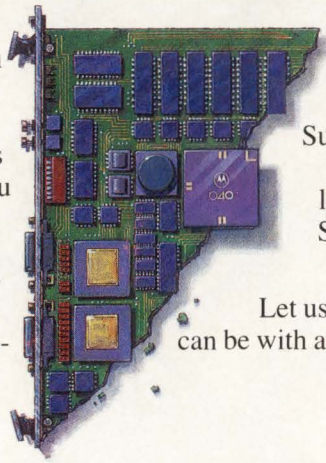

Surprisingly, this kind of quality won't cost you any extra, because Synergy products lead in another important area—value. At Synergy, you don't have to pay a premium price for premium performance.

Let us show you just how far ahead your system can be with a Synergy processor board. Call us today, and get the *whole* '040 story.

Compare our specs. Synergy is superior across the board!

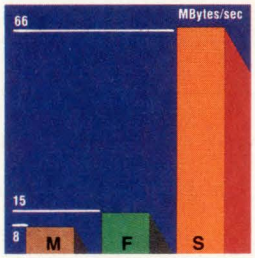

VME Transfers

VME64 doubles bus performance to 66 MB/s—and the SV430 is the only '040 board that has it. But we don't need VME64 to win this comparison. Even normal 32-bit transfers race at 33 MB/s. That's 200% faster than Force or Motorola.

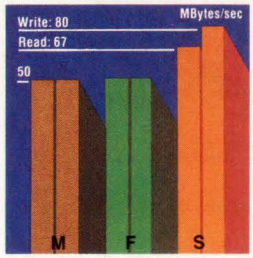

DRAM Burst Rates

A 25 MHz '040 is capable of accessing memory at 80 MB/s. The closer you are to this maximum, the more '040 performance you're gaining. SV430 bursts are 26% faster than Force and Motorola.

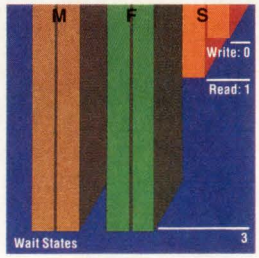

DRAM Random Accesses

Non-burst '040 performance is measured in wait states. Fewer wait states mean higher performance. The SV430 is not only 66% faster than Force or Motorola, it supports twice the on-board memory—32 MB.

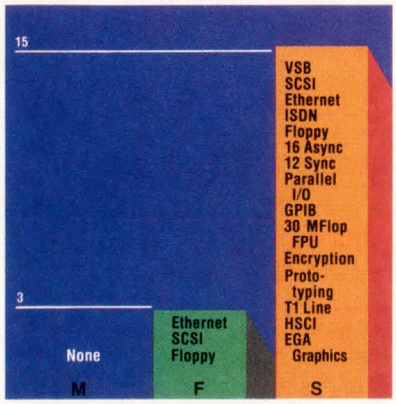

I/O Modules

Synergy's EZ-Bus modules are compatible with our entire line of SBCs. This means Synergy's current line of 12 intelligent I/O modules are immediately available for the SV430—today. No other vendor comes close for selection, functionality or availability.

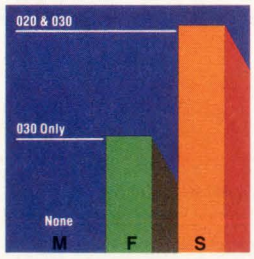

'020/'030 Compatibility

Software compatibility between Synergy SBCs means users have simple upgrades to the SV430 from our '020 and '030 SBCs. Force offers compatibility only from the '030 level, and Motorola offers "upward migration"—a polite phrase that means rewriting your code.

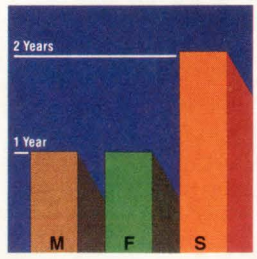

Product Warranty

Synergy backs the reliability of its SBCs with a two year standard warranty. Force and Motorola only offer you one.

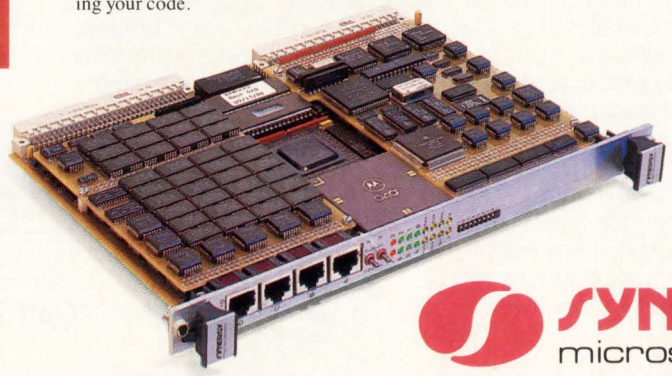

Data from Motorola MVME165 data sheet dated 2/90, and Force CPU-40 data sheet A1 Rev. 1. DRAM measurements shown are with parity. VMEbus transfers are to a 60ns slave.

VME64 is a trademark of Performance Technologies, Inc.

SYNERGY microsystems

Synergy Microsystems, Inc., 179 Calle Magdalena, Encinitas, CA 92024 **(619) 753-2191** FAX: 619-753-0903

INSTANT DATA ACCESS (IDA) DIAL (617) 494-8338 DOCUMENT NO. 1042

CIRCLE NO. 3

A PENNWELL PUBLICATION VOL 31, NO 4 / APRIL 1992

COMPUTER DESIGN

Technology and Design Directions

FOR ELECTRONIC ENGINEERS & ENGINEERING MANAGERS

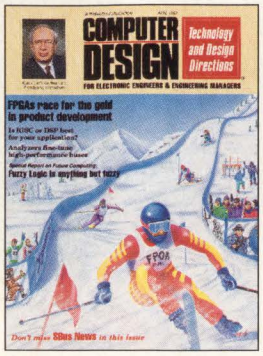

This month's cover story explains why FPGAs are edging out gate arrays and ASICs when it comes to product development. Will they get the design wins? **88**

Illustration by Mike Gardner

Calendar **16**

Editorial **22**

Letters, Covers & Other Diversions **153**

Advertiser Index.....**159**

Special Supplement

SBus News

Technology and Market Update for SBus Users and OEMs

Stiff competition for hosts and boards

SBus chases VME, ISA, EISA

See Page 129A

NEWS BRIEFS

Xilinx files lawsuit against Actel for FPGA architecture . . . Fuzzy logic to be available as ASIC core . . . Mentor sells its emulation technology . . . CFI plots a credible course . . . Trouble in paradise? . . . PC processor products proliferate . . . Army vehicles go VME . . . Is it goodbye CISC, hello RISC in laser printers? . . . Thanks for the memory . . . Multimedia for the one and only **12**

TECHNOLOGY DIRECTIONS

CAE/CAD Tools

Newest VHDL tools encourage conformity to standard **30**

Integrated Circuits

DRAM/SRAM combo chip cures cache headaches **34**

ASICs & ASIC Design Tools

Choices expand to incorporate design for testability **40**

Computers & Subsystems

Digital's Sun killer also rises.....**46**

Embedded PC bus standard gains ground.....**50**

Multimedia—here today, not tomorrow **56**

NEW PRODUCT DEVELOPMENTS

Computers & Subsystems

Two Sparc-based CPUs unveiled **134**

Ethernet multiplexer takes over network processing **139**

Software & Development Tools

Graphic tool shows multiple views of real-time behavior **136**

ASICs & ASIC Design Tools

High-density PLD runs 80 MHz, programmable in-system.....**140**

Xilinx 10,000-gate FPGA supports 32-bit systems.....**141**

CAE/CAD Tools

Scicards revamped for high-speed digital design.....**143**

CAE suite spans platforms, entry methods and technologies.....**144**

Integrated Circuits

Chip set optimizes 486 caching **145**

MORE ON PAGE 5

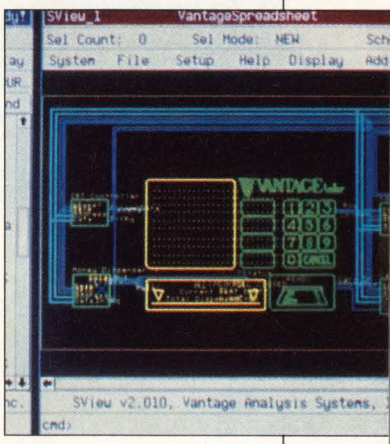

Page 30

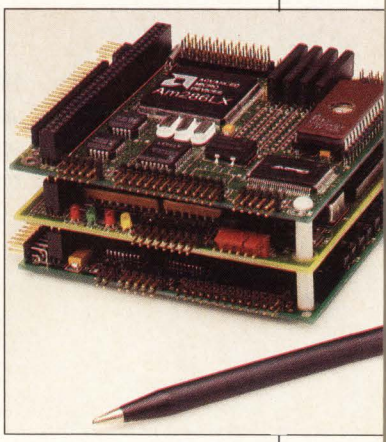

Page 50

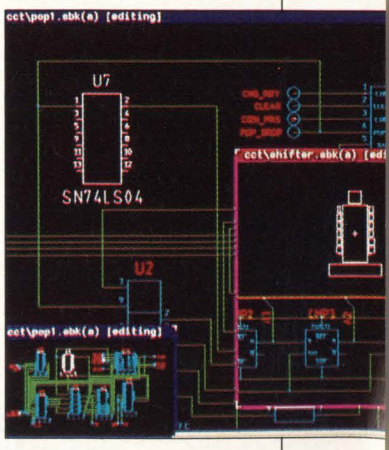

Page 144

Real-Time Software Performance Analysis and Test Coverage

HMI's Performance Analysis Card (PAC) provides real-time software performance analysis and real-time software test coverage for all HMI-200 series in-circuit emulators. This option operates completely transparent to the system under test and collects its data in real-time to establish a true profile of the software execution.

Features:

- Hardware implemented
- Up to eight modules can be defined
- Histograms for each module are displayed
- Minimum, Maximum and Average time duration for each module displayed
- Coverage mode displays which pieces of the code did and did not execute
- Trace data has a time stamp

Benefits:

- More efficient code produces higher performance products for your company
- Better tested code eliminates bugs generated from untested code and creates higher quality software for your company
- The Emulator and Performance Analysis together shorten the design cycle time allowing your company to have its window of opportunity in the marketplace

Available Emulators:

68000	68302	6809/6809E	68HC001	8085
68008	68301/303	68EC020	8051 Family	64180/Z180
68010	68330/333	68EC030	DS5000	Z80
68020	68331/332	68HC11 including	8096/80196 Family	68HC16 Family
68030	68340	F1 and D3		

Now supporting 68040 Series

Write or call for further information and *free* demo disk.

HMI

HUNTSVILLE MICROSYSTEMS, INC.
3322 South Memorial Parkway, Huntsville, AL 35801
(205) 881-6005

CIRCLE NO. 4

A PENNWELL PUBLICATION VOL 31, NO 4 / APRIL 1992

COMPUTER DESIGN

Technology and Design Directions

FOR ELECTRONIC ENGINEERS & ENGINEERING MANAGERS

CONTINUED FROM PAGE 3

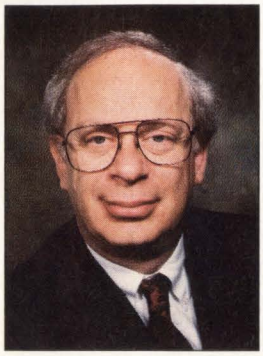

QuickTurn's Kaufman on: Prototyping alternatives

Every design process ends in a prototype, whether it's called that or not. Kaufman evaluates first articles, scratch-built prototypes and CAP. **27**

COMING NEXT MONTH

Analog vs. digital approaches to signal processing

Intelligent power control

PCB design and layout

RISC processors

Instrumentation

TECHNOLOGY & DESIGN REPORTS

Is RISC or DSP best for your application?

RISC processors have already kicked CISC out of more than a few embedded sockets. Now, DSPs are trying to prove they can be just as useful in embedded applications. —*Dave Wilson*.....**65**

Analyzers fine-tune high-performance buses

As standard buses increase in number and complexity, bus analyzers have progressed from simple logic analyzers to specialized, protocol-specific machines. —*Warren Andrews*.....**75**

COVER STORY

FPGAs race for the gold in product development

Gate arrays and standard-cell ASICs are still the big winners when it comes to product development, but FPGAs are closing in with competitive benefits such as design flexibility, reduced cost and quicker time-to-market. —*Barbara Tuck* **88**

SPECIAL REPORT ON FUTURE COMPUTING

Fuzzy Logic is anything but fuzzy

Fuzzy logic is real, and its use is gaining popularity in the United States. Widely used in Japan, it can deliver results that are simpler to achieve, more precise and more reliable than conventional approaches to computing and control. —*Tom Williams*.....**113**

DESIGN STRATEGIES

SDI testbed relies on standard hardware to rein in costs

Sparta designs a battle management system for the US Air Force's SDI. —*Bob Margolin*.....**107**

PRODUCT FOCUS

C/C++ compilers bring faster code crunching to PC platforms

The beauty of the C programming language is that it combines the ease of use of a high-level language with the ability to manipulate bits. —*Jeffrey Child* **129**

COLUMNS

Mixed-Signal Design —*Stephan Ohr*

Where speed replaces precision **147**

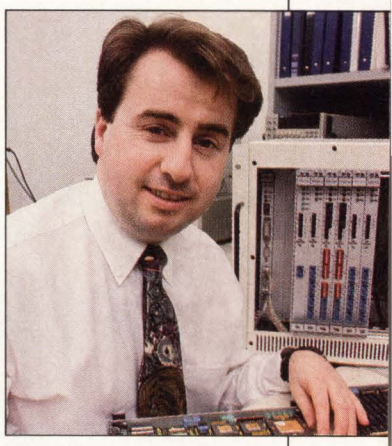

Page 65

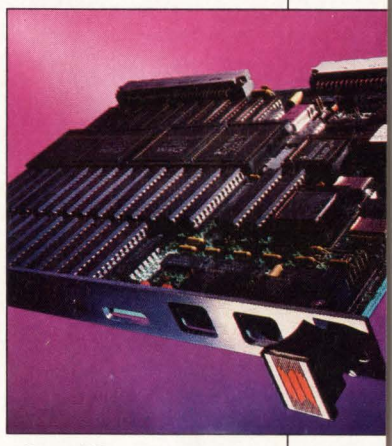

Page 75

Page 113

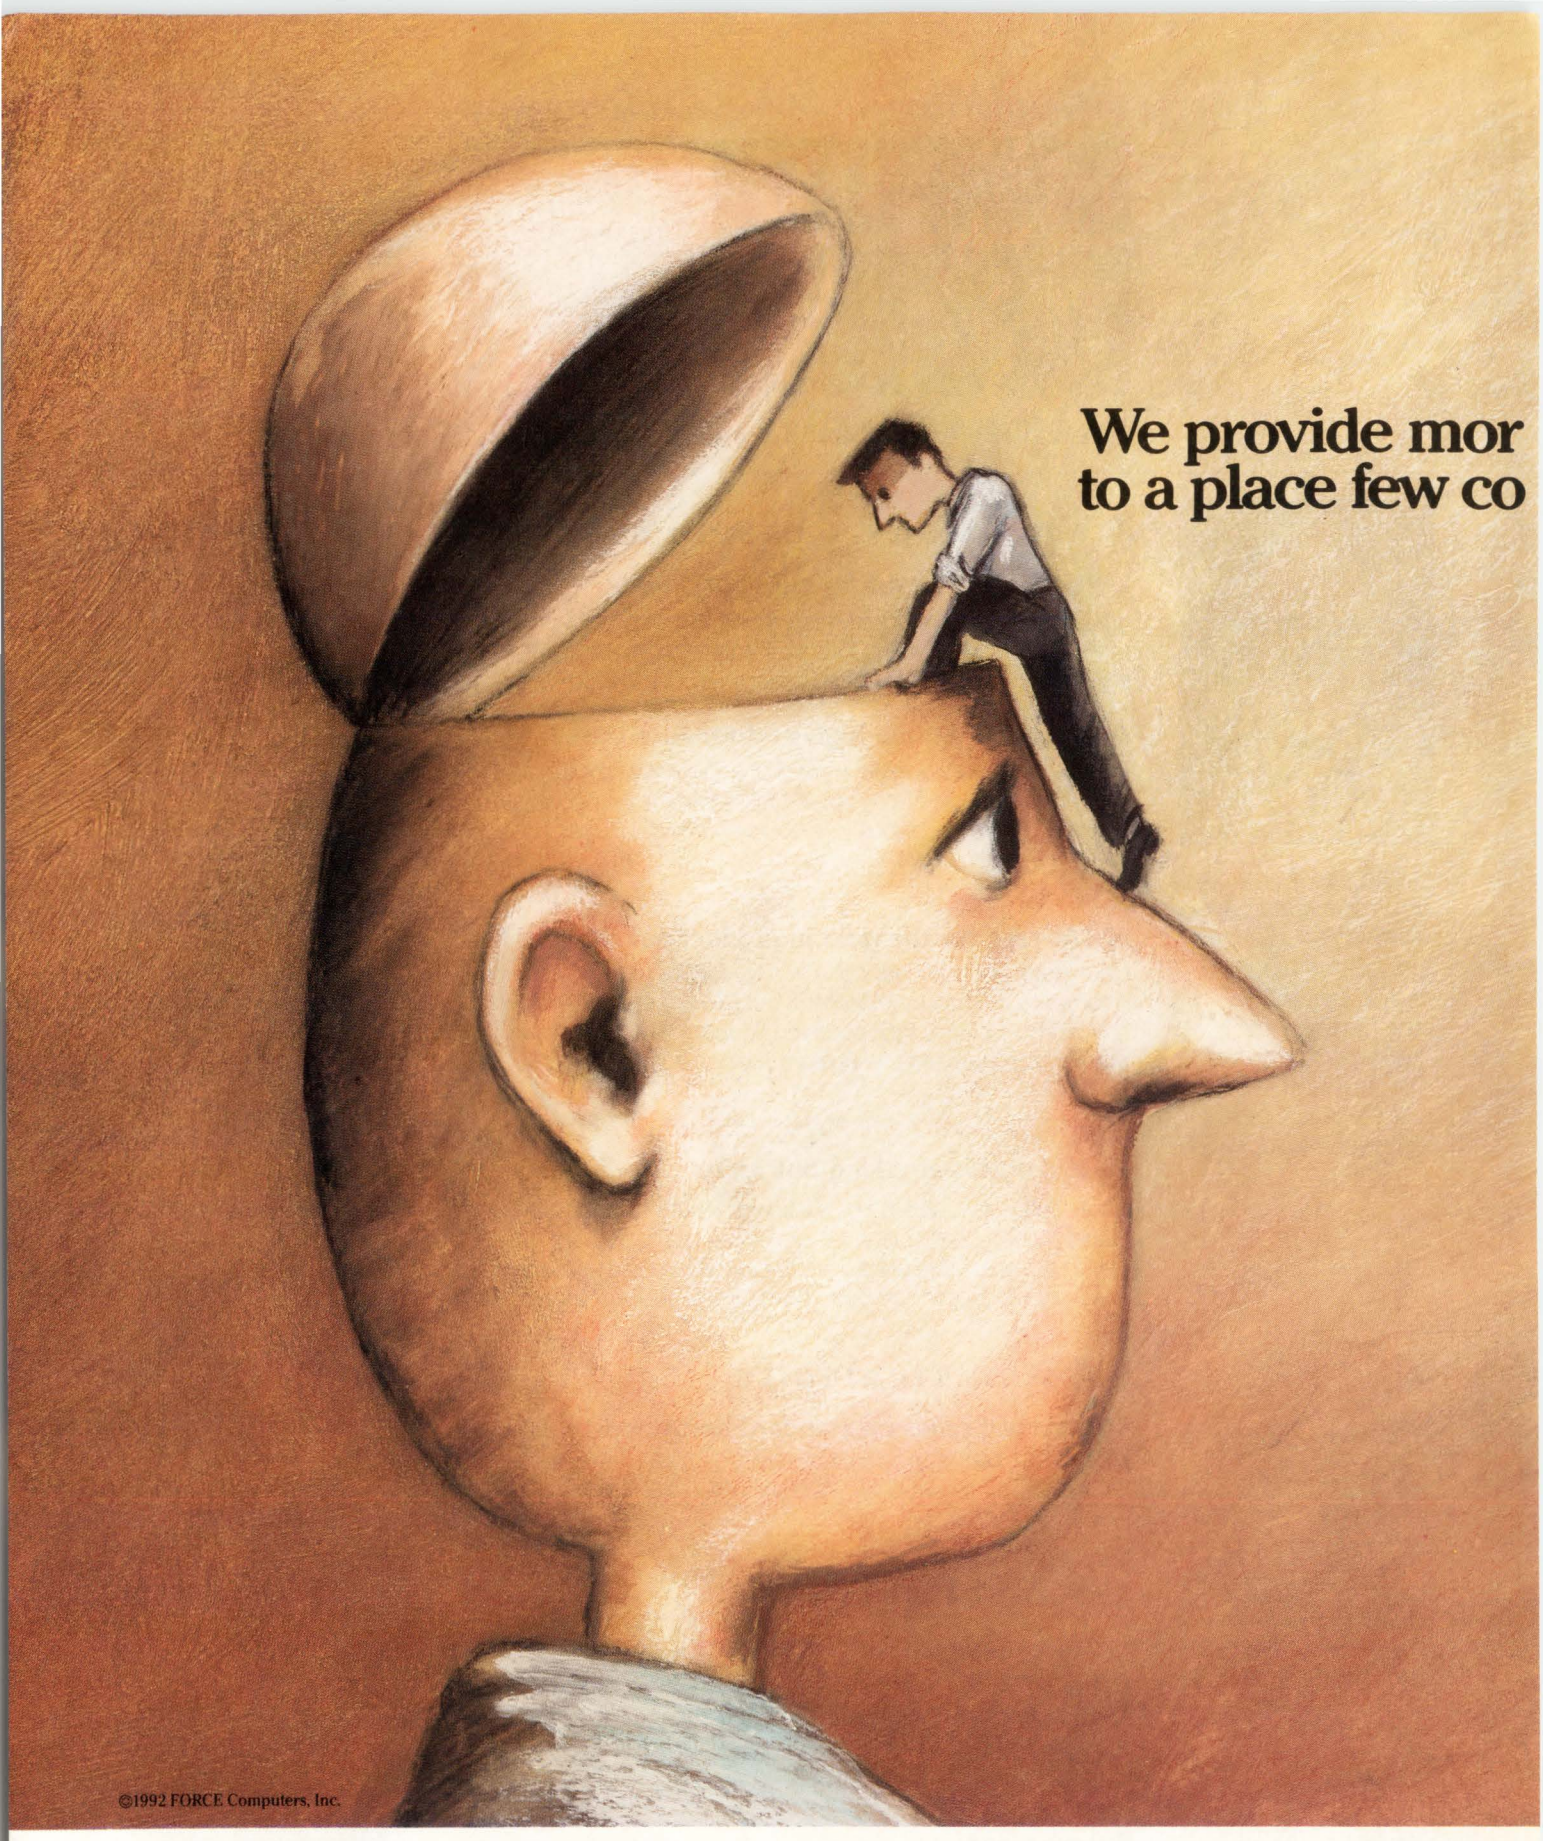
We provide mor
to a place few co
©1992 FORCE Computers, Inc.

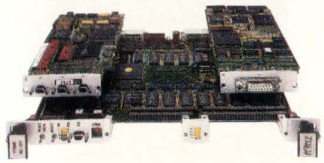
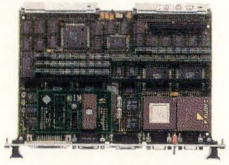

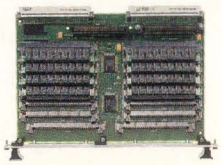

NEC JAPAN
FGA-002A

e solutions because we've been mpanies have dared to venture.

When it comes to embedded systems design, nobody offers better solutions than FORCE.

But, we have to admit, we had a little help. From you.

At FORCE, we get inside the heads of our customers to learn about their requirements.

The result is the world's most advanced CISC and SPARC™ based VME single board computers, including the latest SPARC CPU-2E-based boards.

What else did you have in mind? How about bus extensions – like SBus and our own FLXibus.™ Plus SunOS™ and the widest variety of real-time operating systems and kernels.

And when your designs need to evolve to the next generation, count on FORCE for the best in VME64/Plus™ or Futurebus/Plus.™ Because better products come from better knowledge.

For a partner that can get your application up and running fast, call 800-237-8863, ext. 10. In Europe, call 49.89.608-14-0.

And thanks for keeping an open mind.

FLXibus, VME64/Plus and Futurebus/Plus are trademarks of FORCE Computers, Inc. All other brands or products are trademarks of their respective holders.

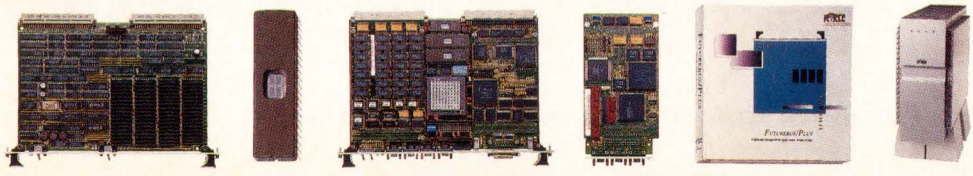

FORCE COMPUTERS®

3165 Winchester Blvd., Campbell, CA 95008-6557
Prof.-Messerschmitt-Str. 1 W-8014 Neubiberg/München

CIRCLE NO. 5

$4,995

Sun SPARCstation ELC

Save $5. Gain 14 MIPS.

*U.S. list price for grayscale display with diskless, 8MB of memory. SPECmarks: HP Apollo, 34; Sun, 20.3. MFLOPS: HP Apollo, 8.4; Sun, 3.

©1992 Hewlett-Packard Company ADCSWG018

$4,990

HP Apollo Series 700 Model 705

With the new HP Apollo Series 700 Model 705 RISC workstation, you'll also gain an additional 13.7 SPECmarks.* And 5.4 MFLOPS.*

So, while the difference in price is just a pleasant surprise, the difference in performance is astonishing.

And that $5 savings won't short you on anything. We've designed the Series 700 to optimize X-window performance. What's more, our Model 705 has a 19" gray-scale display (Sun only offers a 17" monochrome display on the ELC). And, with over 1,800 applications already available, you're ready to get up and running fast.

The open architecture of our workstations also lets them work smoothly in a multivendor environment. And, if you should need our help, HP's worldwide support is at your service.

So call **1-800-637-7740, Ext. 2797.** We'll show you a very fast way to save money.

HEWLETT PACKARD

CIRCLE NO. 6

Considering

Multimedia or Interactive Systems ?

Consider **MICROBOARDS**

Authoring systems for:
○ **DOS, Mac, UNIX, OS-9, CD RTOS**
○ **Player Hardware**

Red Book
Yellow Book
Green Book
Orange Book

Standard and Custom OEM designs

MICROBOARDS, Inc.

MBi America
Tel: (612) 448-9800
Fax: (612) 448-9806

MBi Japan
Tel: 0474 (37) 9811
Fax: 0474 (37) 9821

CIRCLE NO. 9

200 MFLOPS OF POWER!

- Single or dual i860 processors
- True color frame buffer option
- DT Connect, ITI VISION*bus*, SCSI I/O options
- Up to 64 MB memory
- AT or VME

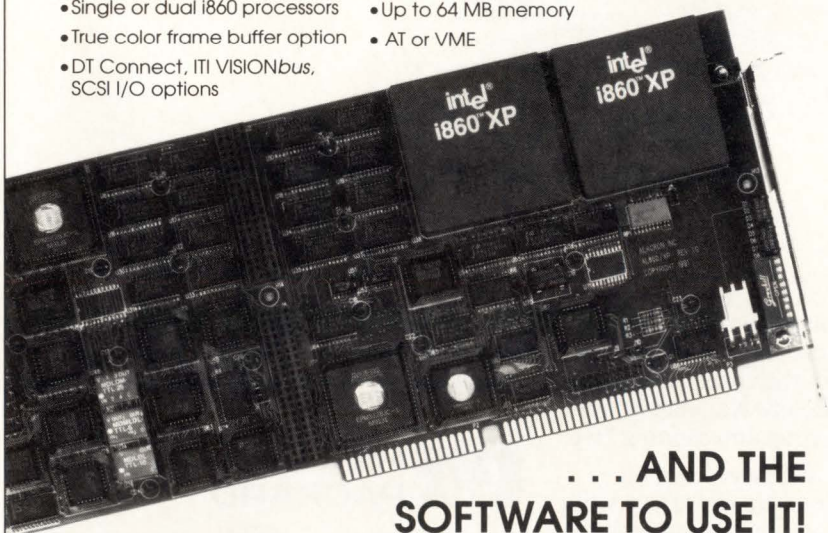

. . . AND THE SOFTWARE TO USE IT!

ALACRON

Suite 204, 71 Spitbrook Rd.
Nashua, NH 03060
Phone: (603) 891-2750
Fax: (603) 891-2745

- For DOS, Unix 386, Sun OS, VxWorks
- i860 Native Unix Sys V.4 with X Windows 11.4
- Fortran and C compilers
- Scientific and imaging libraries with over 250 routines
- 256^2 2DFFT ≤ 55MS
- 512^2 3x3 convolution ≤ 88MS

Trademarks are property of their respective holders.

CIRCLE NO. 10

COMPUTER DESIGN *Technology and Design Directions*

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

SENIOR EDITORS
Warren Andrews, *Computers & Subsystems*
(508) 283-2102
Barbara Tuck, *ASICs & ASIC Design Tools*
(516) 473-1661
Michael G. Donlin, *CAE/CAD Tools*
(508) 392-2123
Tom Williams, *Software & Development Tools*
(408) 335-5514
Dave Wilson, *Integrated Circuits*
(508) 392-2176

SECTION EDITOR
Jeffrey Child, *New Product Highlights*
(508) 392-2126

CONTRIBUTING EDITOR
Stephan Ohr, *Analog & Miixed-Signal*
(908) 232-1380

MANAGING EDITOR:
Arlyn S. Powell, Jr.

CHIEF COPY EDITOR:
Annette M. Staron-Wilson

ASSISTANT EDITOR:
Lisa A. Coleman

PRESENTATION MANAGER:
Richard Sarno

EDITORIAL ASSISTANT:
Claire Ellis

ART DIRECTOR: Jan Horner

ILLUSTRATOR: Manuel King

AD TRAFFIC MANAGER:
Kelly Rice, (508) 392-2198

CIRCULATION DIRECTOR:
Robert P. Dromgoole, (918) 832-9213

CIRCULATION MANAGER:
Paul Westervelt, (918) 832-9287

REPRINTS:
June Bozarth, (918) 835-3161

PUBLISHER
David L. Allen, (508) 392-2111

ASSOCIATE PUBLISHER/ NATIONAL SALES MANAGER:
Tim L. Tobeck, (508) 392-2116

PennWell PUBLISHING COMPANY

Advanced Technology Group
Dr. Morris R. Levitt, 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

BPA ABP

FOR THOSE WHO THINK ALL ATTACHED PROCESSORS ARE ALIKE, PROCESS THIS.

Your choice of processors can make or break your system design. Performance in = performance out. So start with the leader in high performance real-time computing solutions. Mercury's MC860VS. The only attached processor that offers up to 32 Intel i860s in no more than 8 VME slots. So you get 80 Mflops to 2.5 Gflops to handle the most demanding applications in defense signal processing and medical imaging. And all Mercury products can be configured in workstations and chassis systems to deliver scalable performance at a scalable price. So if you're building or buying a high performance computing solution, take advantage of Mercury's unparalleled investment in standards, innovative hardware, and software development environment. And get the most complete, flexible, high performance computing solutions available. If you've been thinking about multiprocessing solutions, here's some important information to process:

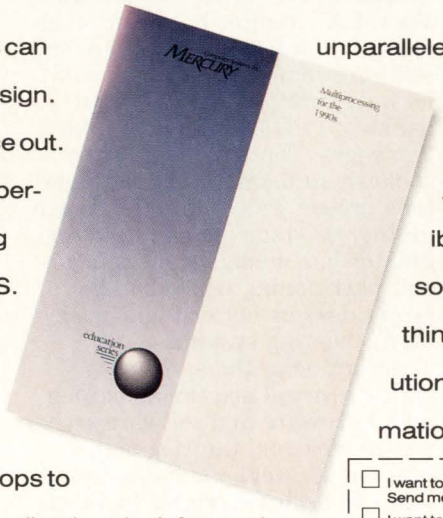

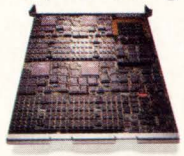

☐ I want to learn more.
Send me your white paper "Multiprocessing for the 1990s."
☐ I want to subscribe to Mercury's ongoing Education Series.
☐ Please send me more information on the MC860VS right away.

Name _____
Title _____
Company _____
Address _____
City/State/Zip _____
Phone _____

Please send to
600 Suffolk Street
Dept. K.
Lowell, MA 01854
Fax (508) 458-9580
(508) 458-3100

Computer Systems Inc
MERCURY

CD4/92

Mercury's MC860VS. The Ultimate Performance Machine.

INSTANT DATA ACCESS (IDA) DIAL (617) 494-8338 DOCUMENT NO. 1020

CIRCLE NO. 11

Xilinx files lawsuit against Actel for FPGA architecture

With FPGA competition getting hotter every day, it's not surprising that the first lawsuit with regard to patent infringement has been filed by inventor and market leader Xilinx (San Jose, CA). The company has filed suit in the Federal District Court for the Northern District of California against Actel (Sunnyvale, CA) for infringement of a patent covering the architecture used in FPGAs. "This is the first time we've had to enforce our patent rights," commented Xilinx president Bernard Vonderschmitt. "It's vitally important to the company and its investors that we enforce these patents vigorously to protect our intellectual property," he added.

Legal counsel for Xilinx, Robert Hinckley, reports that the lawsuit covers the basic structures used in Actel's FPGAs. Xilinx informed Actel of its claim of patent infringement within a year of the 1989 issuance of that patent, according to Hinckley. Discussions have taken place between Xilinx and Actel, but the two companies have not managed to reach an out-of-court solution.

Actel maintains that its ACT product architectures don't infringe the Xilinx patent. "It's the opinion of two independent counsels that Actel will not be held to infringe any valid claim of the Xilinx patent," said John East, president and CEO. "Indeed, Actel intends to vigorously defend against this action."

Asked if Xilinx will file suit against other FPGA makers, Hinckley says, "We haven't focused on others yet. The focus has been on Actel." *—Barbara Tuck*

Fuzzy logic to be available as ASIC core

Togai Infralogic (Irvine, CA), which currently markets software development tools, its own specialized FP110 fuzzy logic processor and accelerator boards based on the FP110, has announced that it will be licensing its fuzzy processor technology in the form of processor cores for use in ASICs to a wide variety of interested semiconductor manufacturers and OEMs. OEMs would license specific core designs to be provided by Togai. Togai foresees such cores, which can be custom designed for different price/performance requirements, being built into custom ASICs, hybrid conventional/fuzzy processors and fuzzy coprocessors.

At the same time, Togai announced an agreement with Siemens AG (Munich, Germany) to jointly develop fuzzy logic and applications for the European market. While no licensing agreement involving the core technology was announced with Siemens, the agreement will pool Togai's expertise in fuzzy logic hardware and software with Siemens' market presence and skill in developing electrical and electronic applications. *—Tom Williams*

Mentor sells its emulation technology

Mentor Graphics (Wilsonville, OR) advanced its plan to become a software-only vendor with the sale of its hardware emulation technology to Quickturn Systems (Mountain View, CA). Mentor has been developing the project—a hardware and software combination that would let users prototype ASICs with banks of FPGAs—for more than two years. Though a Quickturn spokesman pointed out that Mentor's project was still in the "rough prototype" stage, he admitted that Mentor has made advances in circuit partitioning which handle larger designs better than Quickturn's existing systems.

According to the agreement, Quickturn will acquire all of Mentor's hardware and software architecture designs, working prototypes of Mentor's emulation system and exclusive rights to all existing and pending hardware emulation patents for an undisclosed amount of cash and stock. The sale also means that Mentor will now endorse emulation technology, something it has been reluctant to do while it had a competitive product under development.

Quickturn's strengthened position comes at a time when it's facing its first head-on competition in the hardware emulation market. In March, start-up PiE Design Systems (Sunnyvale, CA) introduced a logic emulation product line for VLSI-based system design verification that's targeted directly at potential Quickturn customers. In the face of this competition, it will be interesting to watch whether Quickturn will be able to repeat in 1992 the 100 percent sales growth it experienced last year, when Intergraph, Harris, Kodak, National Semiconductor, NCR Germany, and Bull Italy embraced the Quickturn computer-aided prototyping verification methodology.

—Mike Donlin, Barbara Tuck

CFI plots a credible course

The CAD Framework Initiative (Austin, TX) recently announced a multirelease program at its annual business meeting that president Andy Graham described as, "...a credible plan to deliver on our promises in a pivotal year." CFI's new technical director, Nick English, held out hope that some of the products based on the 1.0 specification released in December might be ready in time for this year's Design Automation Conference.

The 1.0 release includes the capability to manipulate and exchange netlist data, while future releases of the spec will permit the exchange of design representation data and intertool communication. Release 2.0, targeted for December 1993, will support multiple views of a design and will include limited support for design management. *—Mike Donlin*

Trouble in paradise?

Although supporters have been racing to get all the latest advances incorporated into Revision D of the VME (1014) specification, there are some conservative voices calling to put on the brakes a little to make sure everything is going to work. The VME64 part of the specification has been well exercised and at least three interface ICs—and over a dozen board-level products—incorporate the protocols.

It seemed as if source-synchronous block transfer had just about slipped through with little objec-

Continued on page 14

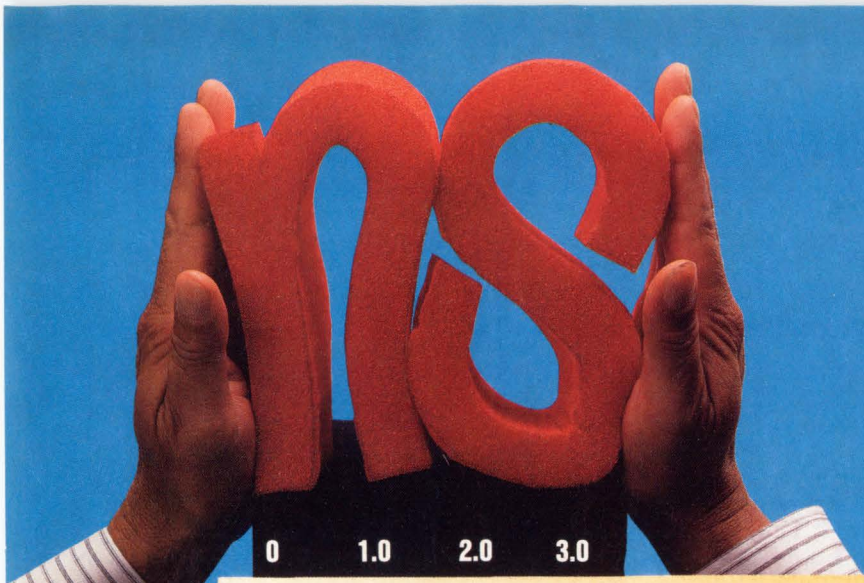

COUNT ON IDT

When Every Nanosecond Counts

Squeeze critical nanoseconds from your high-speed logic interface with the fastest FCT logic available. IDT's FCT-CT family offers speeds that are 50% faster than standard FCT or FAST logic families — as fast as 3.4ns (typical)!

The Perfect System Solution

As a system designer, you need the perfect combination of:

1. ***Fastest speed***
2. ***Low ground bounce***
3. ***Low power consumption***

FCT-CT logic has true TTL compatibility for ease of design. The reduced output swings and controlled output edge rate circuitry ensure low system noise generation. No other technology offers higher speeds or lower power consumption.

The FCT-CT family is completely pin- and function-compatible with FCT logic, and is available today in all standard packaging.

FUNCTION	PROPAGATION DELAY (Max)	OUTPUT ENABLE (Max)	OUTPUT DISABLE (Max)
Buffers	4.1ns	5.8ns	5.2ns
Transceivers	4.1ns	5.8ns	4.8ns
Registers	5.2ns	5.5ns	5.0ns
Latches	4.2ns	5.5ns	5.0ns

Free Logic Design Kit

Call our toll-free hotline today and ask for Kit Code 3061 to get a *1991 High-Speed CMOS Logic Design Guide* and *free FCT-CT logic samples.*

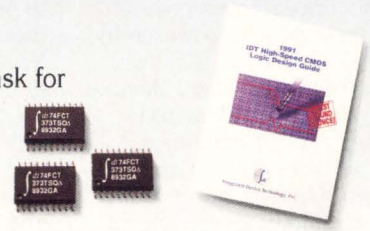

12ns 256K SRAMS

Fastest cache solutions for RISC and CISC CPUs. 36+ ultra-high-speed sub-micron SRAMs for 33MHz processing & beyond are in the **SRAM Data Book**.

35mips RISC CHIPS AND MODULES

R3000A for the most mips at *any* MHz; R3051 for CPU, cache, & buffers on one chip. Modules, eval. boards & software complete the family. See them in the **RISC Data Book**.

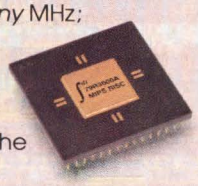

HIGHEST-PERFORMANCE MEMORIES

Fast FIFOs, dense dual-ports, BiCEMOS ECL, & memory modules. 120+ FIFOs & multi-port memories, 5ns ECL, & multi-chip modules are in the **Specialized Memories Data Book**.

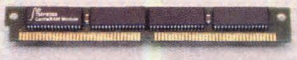

Call today for your new IDT data books with complete technical specifications and application information.

(800) 345-7015 • FAX: 408-492-8674

Integrated Device Technology, Inc.

The IDT logo, CEMOS, BiCEMOS, and R3051 are trademarks of Integrated Device Technology, Inc.

CIRCLE NO. 7

Continued from page 12

tion until at least one "fatal flaw" was discovered in the potential SSBLT definition. Dieter Nattkemper, Openbus products manager for Newbridge Microsystems (Kanata, Ontario), reports that the SSBLT specification relies on triggering a positive edge of one of the VME signals. "This isn't reliable," he says, "and if allowed to be included in the specification, [it] would have been a major problem."

While the fatal flaw probably has a number of possible "fixes" which will undoubtedly be put forth at future sessions, it's alerted the committee that it may not be wise to rush into any specification changes without thorough examination and testing. If some fatal, overlooked flaw were to be appended to the specification, it would not only be an embarrassment, but it could undermine the credibility of the bus. *—Warren Andrews*

PC processor products proliferate

Advanced Micro Devices (Austin, TX) is readying a slew of Am486 processors as a follow-up to its highly popular Am386. First members of the family will include 25-, 33-, and 50-MHz versions of the Am486DX device in PGA packages. And 25-MHz versions of the Am486SX devices will appear in PGA and surface-mount packages. A 25-MHz Am486DX will be available in surface-mount at a 20 percent premium over the 25-MHz Intel 486SX—that's the one with the disabled on-chip coprocessor.

On the 386 front, AMD hasn't stood still either. The company announced a 33-MHz version of the Am386DX/DXL device, as well as a 40-MHz version in a plastic package.

On the subject of Intel clones, don't forget Integrated Information Technology (Santa Clara, CA) either. That company just debuted a 33-MHz coprocessor for the 386SX. Dubbed the 3C87SX, it's just waiting for an Am386SX to work with. If you need a lower-power part, consider a family of surface-mount coprocessors from Cyrix Corporation (Richardson, TX). The 87LC and 87DLC coprocessors offer design engineers a 2.7-V alternative to Intel's offerings in a small 14 × 14-mm, 80-pin quad-flat pack. *—Dave Wilson*

Army vehicles go VME

After almost two years of jockeying, the U.S. Army has decided to go with a standard VMEbus form factor for its vehicle electronics. What has been at issue is SAVA (Standard Army Vetronics [vehicle electronics] Architecture). Initially it was proposed by the U.S. Army TACOM (Tank Automotive Command), based in Warren, MI, that it be based on a standard bus, and bids went out to a number of contractors to develop the necessary systems. The four bid winners, General Electric, Texas Instruments, FMC Corporation, and General Dynamics, formed AVTA (Army Vetronics Technical Association) to finalize the program.

AVTA elected to use the standard VME electrical interface, but chose a slightly larger board—9.5 in. in place of the standard 9.2 in.—and a slightly different connector—four rows of pins instead of three. This format, which had taken on the SAVA name, persisted for almost two years. Now, looking to take advantage of the greater economies of scale available with standard form factor VMEbus cards, the Army has elected to change the mechanical configuration to conform to the IEEE-P1101.2 conduction-cooled, double-Eurocard standard.

—Warren Andrews

Is it goodbye CISC, hello RISC in laser printers?

Underscoring the wide price/performance range of RISC technology, a flurry of laser printer announcements in March revealed a slew of design wins for the 29k family. Dispelling the myth that RISC technology is reserved for high-end systems only, Apple Computer (Cupertino, CA) has announced an entry-level PostScript laser printer powered by the Am29005, the 32-bit RISC processor from Advanced Micro Devices (Sunnyvale, CA). The Personal Laser Writer NTR, introduced March 23, sells for $2,100 and is the first Apple laser printer based on RISC technology. According to AMD, the new printer uses the Am29005's processing power to print documents up to six times faster than the Personal Laser Writer NT, Apple's 68000-based laser printer.

Meanwhile, the 29k RISC architecture continues to score wins at the high end. GCC Technologies (Boston, CA) introduced a Phoenix PostScript color laser printer based on the Am29000. Also following the RISC route, Japanese firm Minolta announced an Am29000-based laser printer at last month's Cebit show in Germany. *—Jeffrey Child*

Thanks for the memory

Working in conjunction with a research team from Purdue University, scientists at Cree Research (Durham, NC) have demonstrated a silicon carbide capacitor, the major building block of a nonvolatile RAM based in silicon carbide. According to Cree, the SiC RAM cell will be able to store data indefinitely without refresh. The longer storage time is a function of the SiC's energy bandgap, which is 2.5 times larger than that of silicon. (Storage time increases exponentially with increases in bandgap.) As a result, leakage current in SiC has been measured to be at least that of 0.00001 typical silicon DRAMs. This means that the Cree device can store data for years instead of minutes. Another benefit is that the Cree device can be written to indefinitely without data fatigue. Nonvolatile silicon memories today use floating-gate technology which can lead to the device wearing out after a large number of write/erase cycles. *—Dave Wilson*

Multimedia for the one and only

Redmond-based software giant Microsoft said that it intends to make a version of its Multimedia Viewer authoring system available to enable multimedia title development for Sony's portable CD-ROM XA (extended architecture) player, to be released this fall. Viewer for the Sony player is based on Microsoft Viewer for Windows. Viewer will let software developers create titles which combine text, audio and video. *—Dave Wilson*

If you don't think our real-time systems have taken off, consider some of the programs we've landed.

You already know the significance of the Alsys name when it comes to quality Ada. But you may not have realized that over half of our business is in—and over half of all our resources are devoted to—real-time applications. Right now Boeing, Lockheed, McDonnell Douglas, NASA, AIRBUS, European Space Agency and many others rely on us for solutions for mission and safety-critical applications. Alsys offers one of the broadest ranges of compiler products and developer tools, plus unparalleled expertise and guidance every step of the way.

If you need to reduce risk, deliver projects on time, and do both within budget, the first step is easy: send in the coupon, or call 617-270-0030. We'll rush you a handbook detailing some powerful reasons why you should consider enlisting our support. Alsys. We're serious about real-time.

Please rush me the Alsys handbook detailing your real-time capabilities.

(name) (title)

(company) (phone no.)

(address)

Alsys, Inc., 67 South Bedford Street, Burlington, MA 01803-5152, phone: 617-270-0030 fax: 617-270-6882

Alsys maintains major operations in the U.S., Sweden, France, England, Germany and Japan.

alsys
A Thomson–CSF company

CD

CIRCLE NO. 8

AMX™

REAL-TIME MULTITASKING KERNEL

8086/88, 80x86/88 80386

68000/10/20 Z80, 64180, 8080/85

- Fast, reliable operation
- Compact and ROMable
- PC peripheral support
- DOS file access
- C language support
- Preemptive scheduler
- Time slicing available
- Configuration Builder
- Complete documentation
- Intertask messages
- Message exchanges
- Dynamic operations
  - task create/delete
  - task priorities
  - memory allocation
- Event Manager
- Semaphore Manager
- List Manager

No Royalties
Source Code Included

INSTANT DATA ACCESS (IDA) DIAL (617) 494-8338 DOCUMENT NO. 1036

Manual only $85 US
AMX 86 $3000 US
(Shipping/handling extra)

Call for prices for other processors.

KADAK Products Ltd.
206-1847 West Broadway
Vancouver, B.C., Canada
V6J 1Y5
Telephone: (604) 734-2796
Fax: (604) 734-8114

CIRCLE NO. 16

TOUCH SCREEN PROJECT?

We stock touch screens, decoding electronics and software for immediate delivery.

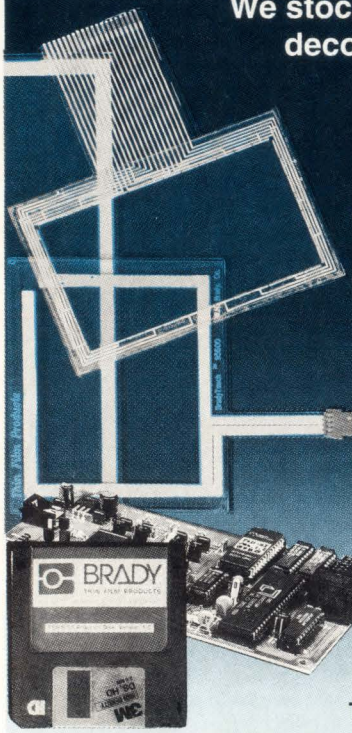

- Analog and matrix touch screens.
- Decoding electronics and software (enables screen to function as a mouse).
- High gloss and anti-glare surface finishes.
- Custom manufactured screens also available.

BRADY®

W.H. BRADY CO.
THIN FILM PRODUCTS

8225 W. Parkland Ct. • P.O. Box 571 • Milwaukee, WI 53201
Phone: 414-355-8300 • Fax: 414-354-0453

Copyright 1991 W.H. Brady Co. All rights reserved.

CIRCLE NO. 17

CALENDAR

CONFERENCES

April 13 - 16
IDEA '92

San Jose Convention Center, San Jose, CA. The International Design Engineering & ASIC Exposition will focus on the latest ASIC, FPGA and EDA technologies in a program of 36 seminars and 13 training sessions. One day will be devoted to programmable logic with two training courses about introducing and evaluating FPGA architectures. Other issues addressed by the conference include: design automation, application-specific ICs, programmable logic, and system development. Contact: IDEA '92, 480 San Antonio Rd, Suite 245, Mountain View, CA 94040, (800) 848-4332 or (415) 949-2742, Fax (415) 949-3635. **Circle 366**

April 13 - 16
SMTCON with Modern Circuit Making

Meadowlands Hilton Hotel and the Meadowlands Convention Center, Secaucus, NJ. The Surface Mount Technology Conference and Exposition and its new companion event, Modern Circuit Making, will cover design, fabrication, assembly, and testing of electronics using state-of-the-art packaging technologies—surface mount, thin and thick film circuitry and multichip modules. Contact: IC Management, 900 N Franklin St, Suite 700, Chicago, IL 60610, (312) 944-3434, Fax (312) 944-5289. **Circle 367**

April 26 - 29
Joint Mountain-Western Regional Technical Symposium

Tropicana Hotel and Nellis Air Force Base, Las Vegas, NV. The symposium theme is "The threat after Desert Storm," for discussion of the aftermath of the Gulf War and the future of electronic warfare and related fields. In addition to technical papers and exhibits, this symposium will feature an aircraft static display with the F-117 Stealth fighter. This symposium is sponsored by the Association of Old Crows (AOC) that supports the defense efforts of the United States and its allies to advance the understanding and state-of-the-art of electronic defense. Contact: AOC Convention Department, 1000 N Payne St, Alexandria, VA 22314-1696, (703) 549-1600. **Circle 368**

April 27 - May 1
XWorld

Sheraton New York, Manhattan, NY. XWorld is the largest conference and exhibition devoted to the X Window system. Over 40 technical and managerial sessions will be offered as well as exhibits. Tutorials will include: "Using the X Toolkit," "Widget Writing," "Debugging X," "System Administration," "C++ Programming," and "Windows for X Programmers." Contact: SIGS Publications Group, 588 Broadway, Suite 604, New York, NY 10012, (212) 274-0640, Fax (212) 274-0646. **Circle 369**

Continued on page 18

FLUKE AND PHILIPS – THE GLOBAL ALLIANCE IN TEST & MEASUREMENT

FLUKE®

PHILIPS

PHILIPS

Test results:

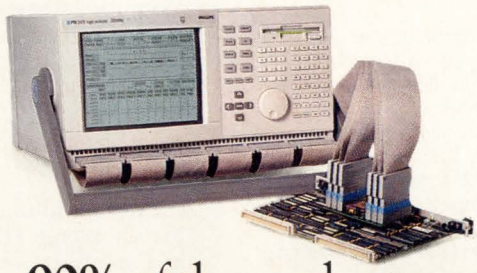

90% of those who try a Philips Logic Analyzer from Fluke buy one.

100% get a free DMM*.

Our logic analyzers sell themselves. All we have to do is get one in your hands. To make sure you do, we're giving you a Fluke DMM*, whether you buy our analyzer or the competition's. (See attached card for complete details).

Only the Philips PM 3580 family of logic analyzers give you *true* dual state and timing on up to 96 channels - simultaneously. All accessible with one probe and one keystroke. Which means no more dual probing or reconfiguration between state and timing. Or no probes at all if you use our boundary-scan test option!

*The top-of-the-line Fluke 12 in our newest DMM family. It combines a smart set of troubleshooting features in a new design that's exceptionally fast and simple to operate — with one hand. It's yours after our 30 minute demo, no matter whose logic analyzer you purchase.

All our analyzers feature 50 MHz state and up to 200 MHz timing speeds. As well as integrated state and timing triggering for fast debug of complex hardware and software problems. Plus broad µp support like Intel®'s i486; i386; 80286; 80186/88 families. The MCS-96, 8051, and i960 families. And the Motorola 68040 to 6800, 68HC11, 68332/1, 68302, 68340, 56001, AMD®'s AM 29030, and TI's 320Cxx family.

The PM 3580 family of logic analyzers is priced from $4495 to $11,450 - about half the cost of comparable analyzers. What's more you can have them up and running in only 30 minutes.

Find out why the PM 3580 family of logic analyzers were the only ones cited for excellence and innovation by *Electronic Design, EDN, Embedded Systems, Electronic Products*, and *R&D* magazines. Take the Fluke Challenge. The odds are 100% you'll be totally impressed.

For literature, our video or a demonstration, call **1-800-44-FLUKE**.

John Fluke Mfg. Co., Inc., P.O. Box 9090, M/S 250C, Everett, WA 98206-9090. U.S. (206) 356-5400. Canada (416) 890-7600. Other countries: (206) 356-5500.
©1992. All rights reserved. Registered T.M. of Advanced Micro-Devices and Intel Corp. Ad No. 00178.

FAST ANSWERS

FLUKE®

CIRCLE NO. 14

DCP

POWER

DCP means the fastest, most flexible communications available for PCs.

For more information on our DCP-286 family of mega-performance communications controllers featuring pick-your-protocol (SNA, Bisync, X.25, frame relay), pick your interface (RS-232, RS-422, X.21, parallel), pick your bus (ISA, EISA, Micro Channel), wide-open 64Kbps and T1 support, full 1MB on-board memory, and flat-out fastest 80286 (and 80386!) design, there's only one call to make.

1-800-EMULEX-3.

EMULEX Intelligent Communications for PCs

CIRCLE NO. 19

80386SX / 80386DX / 80486SX / 80486DX

INDUSTRIAL SINGLE BOARD COMPUTER

TWO YEAR WARRANTY

All TME products come with Two Serial ports, One Bi-directional Parallel port, Floppy & IDE Hard Disk Interfaces.

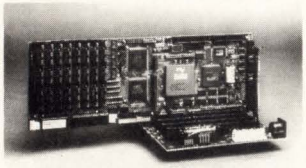

HQP486
* 20/25/33/50Mhz, 80486SX/DX
* Up to 48Mbytes on-board memory
* Optional 256K write back cache memory

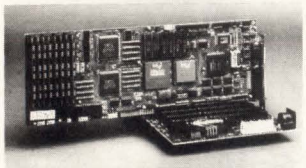

HQP386C
* 25/33/40Mhz, 80386DX
* Up to 32Mbytes on-board memory,
* 128K write back cache memory

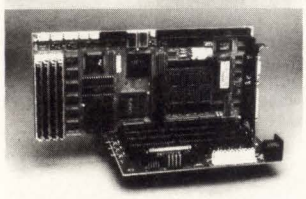

HQP386SXC
* 16/20/25Mhz 80386SX
* Up to 16Mbytes on-board memory
* 64K cache memory
* Up to 1.5Mbytes FLASH EPROM on-board

TME Toronto MicroElectronics Inc.
1848 Bonhill Rd. # 1 Mississauga, Ontario, L5T-1C4 CANADA
Tel: 416 564-4833 • Fax:416 564-4768

CIRCLE NO. 20

CALENDAR

CONFERENCES

Continued from page 16

April 28 - 30
AFCEA Europe Spring Symposium and Exposition

Hotel Melia Castilla, Madrid, Spain. This conference on simulation and high technology C^3I will explore how to solve a broad range of C^3I operational, training and system design problems. The major session topics include: "Simulation and modeling in support of C^3I system design and optimization," with subsessions on future NATO communications networks and more. Contact: J. Spargo & Associates, 4400 Fair Lakes Ct, Fairfax, VA 22033, (703) 631-6200, Fax (703) 818-9177. **Circle 370**

May 4 - 7
EDS '92

Las Vegas Hilton, Las Vegas, NV. The Electronic Distribution Show and Conference emphasizes the needs of manufacturers, distributors and sales representatives of electronic components, test equipment and accessories. Contact: The Electronic Industry Show Corporation, 222 S Riverside Plaza, Suite 270, Chicago, IL 60606, (312) 648-1140, Fax (312) 648-4282. **Circle 371**

May 5 - 6
The Canadian High Technology Show & Conference

Lansdowne Park, Ottawa, Canada. Nearly 400 exhibitors from Canada, the United States and Europe will display products in components and microelectronics, instrumentation, production and packaging equipment. The conference program will feature four workshops and seven technical papers. Contact: Suzanne Stevens, Reed-Macgregor Exhibitions, 800 Denison St, Unit 7, Markham, Ontario L3R 5M9, (416) 479-3939, Fax (416) 479-5144. **Circle 372**

May 11 - 14
Test & Design Expo

Garden State Convention Center, Somerset, NJ. The Test & Design Expo is the first show to revolve around the concurrent engineering process. Test equipment and design products will be showcased side by side on the exhibit floor. The conference will feature sessions, tutorials and minicourses that are structured for test and design professionals, emphasizing concurrent engineering, boundary scan and VXI/systems integration. Contact: Miller Freeman Expositions, 1050 Commonwealth Ave, Boston, MA 02215-1135, (800) 223-7126, Fax (617) 730-5708. **Circle 373**

Would you like your event listed here?

Computer Design *can include a calendar announcement for your upcoming conference or seminar as long as it's received at least three months prior to the date of the event. Be sure to include a specific location, a brief description of the event and a telephone, fax number and address of a contact person. Please address calendar items to: Lisa Coleman, assistant editor.*

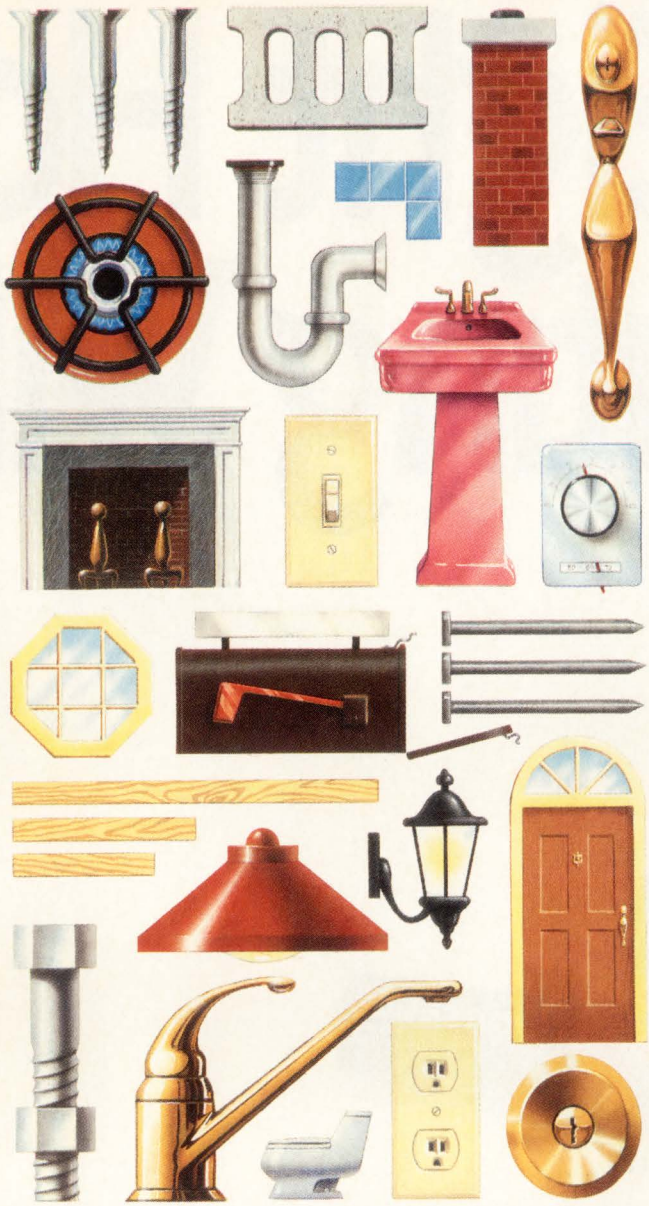

You Can Try Building Your Own Real-Time Operating System,

Building your own operating system for a real-time application can cost you up to a year in expensive programming time. And lost business opportunities. In a fast-moving market, it simply doesn't make sense to build. Not when VxWorks™ from Wind River Systems lets you move right in.

VxWorks gives you all the components and tools you need to start developing applications immediately. Which cuts your costs dramatically. And because VxWorks offers a more full-featured development platform than any other off-the-shelf operating system, it gets your product to market faster. It's the one true turn-key solution, fully compatible with industry standards right out of the box.

Hosts	Targets
Sun	68K
HP	SPARC
DEC	MIPS
IBM	i960

VxWorks supports multiple platforms.

Or You Can Feel Right At Home With VxWorks.

Wind River Systems pioneered off-the-shelf high-speed UNIX® networking for real-time applications. And VxWorks still leads with the most complete, robust networking available. The lean, efficient **wind**™ kernel gives it unsurpassed multitasking speed and functionality.

So before you invest the time and money trying to build your own real-time operating system from the ground up, find out more about VxWorks. And feel right at home from day one. 1-800-677-1586.

WIND RIVER SYSTEMS

Real-Time Solutions for The Real World

© 1992 Wind River Systems, Inc. Wind River Systems, VxWorks, **wind** and the Wind River Systems logo are trademarks of Wind River Systems, Inc. UNIX is a registered trademark of AT&T.

CIRCLE NO. 17

Soon, Eight Ho Computing Will

(actual size)

AMD Introduces The World's First 386 Microprocessor With 3-Volt Technology.

Two standard dry-cell batteries. There's really nothing special about them. Aside from the fact that they can run a powerful, portable 386 computer for a full eight hours. Provided, of course, that portable is built around a low-voltage Am386™ microprocessor.

The 25MHz DXLV and the 25MHz SXLV are available in PQFP packaging.

Thanks to the low-voltage Am386 microprocessors, laptop, palmtop and notebook computer designs will become smaller, lighter, and more powerful than ever before.

With battery life of up to eight hours or more. That's a full day's worth of 386 performance—the per-

901 Thompson Place, P.O. Box 3453, Sunnyvale, CA 94088 © 1991 Advanced Micro Devices, Inc. "We're Not Your Competition" is a service mark and Am386 is a trademark of Advanced Micro Devices, Inc.

urs Of Portable Look Like This.

formance you need to run sophisticated applications like Windows™ 3.0.

And rest assured, the low-voltage Am386 microprocessors are proven compatible and comply fully with JEDEC standards for low-power, 3-volt computing. We can even supply you with the 3-volt EPROMs your systems will need. Other 3-volt system logic is also readily available.

For more information on the low-voltage Am386 microprocessors call AMD today at **1-800-222-9323**. You'll never look at dry-cell batteries the same way again.

Advanced Micro Devices

"We're Not Your Competition."℠

All brand or product names mentioned are trademarks or registered trademarks of their respective holders. DURACELL is a registered trademark of Duracell International, Inc.

CIRCLE NO. 18

Bringing you another dimension

***Our goal has been to make designer and user comments, observations and insights an integral part of* COMPUTER DESIGN.**

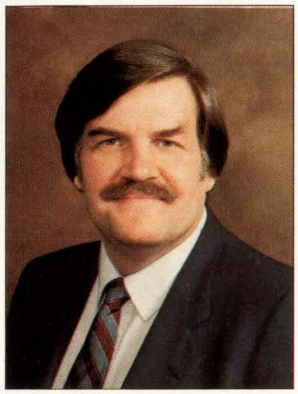

John C. Miklosz
Associate Publisher/
Editor-in-Chief

Like all of you working on a challenging project, we get involved with what we're doing on *COMPUTER DESIGN*. So involved that we take it for granted that you'll notice when we add something new or make some other change to the editorial package you receive every month.

One of the major changes we've made is in how we cover developments in technology, design and new products. Traditionally, editors of electronics/computer design and product publications have gotten their information exclusively from vendors (not exactly unbiased sources) and industry analysts (who focus on business issues more than technology, design or product issues).

Designers and design managers—like yourselves—have an understanding of applications and design problems that goes far beyond that of the vendors who are offering the ASIC, IC, board-level, CAE/CAD, and software solutions you're using for your designs. I should point out that we're not alone in realizing this. The vendors themselves realize it, and more and more are forming closer relationships with the users of their products. It's called "partnering." Our goal, too, has been to "partner" and to make designer and user comments, observations and insights an integral part of *COMPUTER DESIGN*. Look at this issue, and especially at the Design Strategies feature (page 107), and you'll see what we mean.

Trying to go beyond specifications, feature sets and nitty-gritty design details and to provide comparative information, perspective, analysis, and insights that, hopefully, help you along the way toward evaluating a design approach or a new product is anything but easy. Probing vendors, designers and users to go beyond the obvious specs and "how it works" detail is hard and time-consuming, and editors, like everyone else, usually want to take the path of least resistance. To keep our editors on the straight and narrow, so to speak, we laid down the rule that anything they wrote in the Technology Directions section up front in *COMPUTER DESIGN* had to be more than a single-vendor, single-product article. It had to have some mix of comparison, perspective and insight from vendors *and* designers or users. The "traditional" single-vendor, single-product features were relegated to the New Product Developments section—at the back of the magazine. When it was brought to our attention that most other design publications were putting the kinds of products that we were putting in the back of our magazine at the *front* of theirs, we realized that we might be downplaying what we do just a little bit.

Rather than change the flavor of our up-front technology section, we decided to redesign the table of contents. Now, new developments in technology *and* new product developments appear on the first page. All of our features are grouped together on the second page. We think that this grouping—essentially, news and features—now gives proper emphasis to our coverage of the newest developments in technology and products, as well as to our focus on the analysis of ongoing developments and trends.

RADISYS GIVES VME THE INSIDE ADVANTAGE

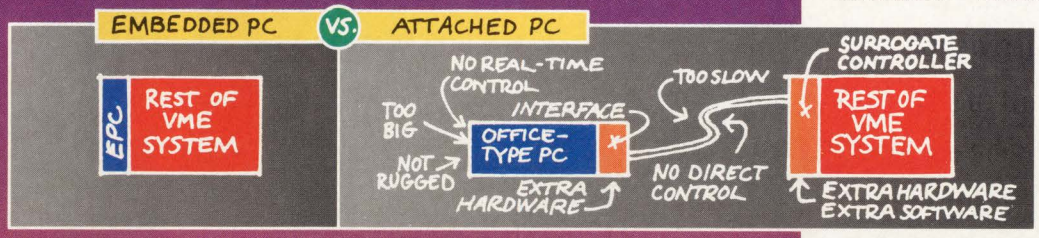

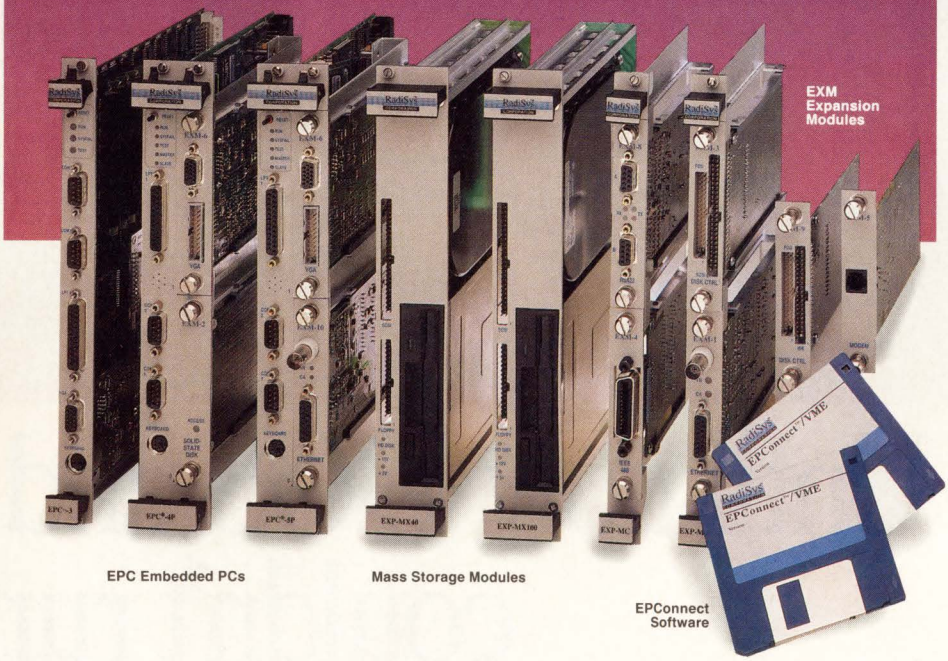

Integrating a PC with your VME system is a smart move. The "PC advantage" provides a superior human interface and access to the PC's huge base of system, application and development software.

The PC Advantage belongs inside your VME system. Not attached to it. By embedding a PC inside your VME card cage, instead of attaching it externally, you break through the inherent communications bottleneck that constricts system performance. You also eliminate the superfluous hardware and software needed to attach two system architectures.

Only RadiSys EPC® Embedded PCs completely integrate the strengths of PC and VME. An EPC, with its exclusive EPConnect™ Software, is the only 386- or 486-based, PC-compatible computer with software that integrates the VMEbus into the DOS, Windows, UNIX and OS/2 environments. EPCs give your VME systems:

- **Highest system performance** from the real-time responsiveness of the direct 32-bit interface between the 386 or 486 and the VMEbus.
- **Improved system packaging** in 1/10th the volume, with integral VME ruggedness, and no bus link baggage.

And EPCs cost you less. EPC-based systems avoid the costly pitfalls of attached PC systems. No extra interfaces, cables, surrogate controllers, or the software to make them work.

Give your VME systems the EPC advantage. Call (800) 950-0044. We'll send all the details. No strings attached.

EPC MODEL	EPC-1 (shipping since Aug '88)	EPC-3 (shipping since Aug '89)	EPC-4 (shipping since Mar '90)	EPC-5 (shipping since Oct '90)
Processor Modules:				
CPU	80386	80386SX	80386	80486
CPU Clock	16 or 20MHz	16MHz	25 MHz	25 or 33 MHz
DRAM	1 or 4 MBytes	1, 2 or 4 MBytes	4, 8 or 16MBytes	4, 8 or 16 MBytes
Graphics	EGA (640 x 350)	VGA (800 x 600)	VGA (800 x 600)	VGA (800 x 600)
Mass Storage Modules:				
Hard Disk Capacity	40 MBytes		40, 100 or 200 MBytes	
Floppy Drive Size/Cap.	3.5" / 1.44 MBytes		3.5" / 1.44 MBytes	
Expansion Capabilities:				
PC Add-in Cards	Yes		Yes	
EXMbus Expansion	N/A	EXM Expansion Modules:		
		EXM-1 Ethernet	EXM-5 Modem	EXM-9 IDE/Floppy Ctrl.
		EXM-2 Solid State Disk	EXM-6 VGA Graphics	EXM-10 Ethernet
		EXM-3 SCSI/Floppy Ctrl.	EXM-7 RS232 Serial I/O	EXM-11 Timer/Counter
		EXM-4 IEEE 488	EXM-8 RS422 Serial I/O	EXM-12 Prototyping Card
Software Support:	EPConnect development, run-time, and multiprocessing software package for DOS, Windows, UNIX,and OS/2			

Copyright © 1989 RadiSys Corporation. All rights reserved. RadiSys and EPC are registered trademarks and EPConnect is a trademark of RadiSys Corporation. 386 and 486 are trademarks of Intel Corporation. OS/2 is a trademark of International Business Machines Corporation. UNIX is a trademark of American Telephone and Telegraph.

The Intel Inside logo is a trademark of Intel Corporation.

RadiSys® CORPORATION

THE INSIDE ADVANTAGE

RadiSys Corporation
19545 NW Von Neumann Dr.
Beaverton, OR 97006 USA
(800) 950-0044
(503) 690-1229
Fax (503) 690-1228

CIRCLE NO. 19

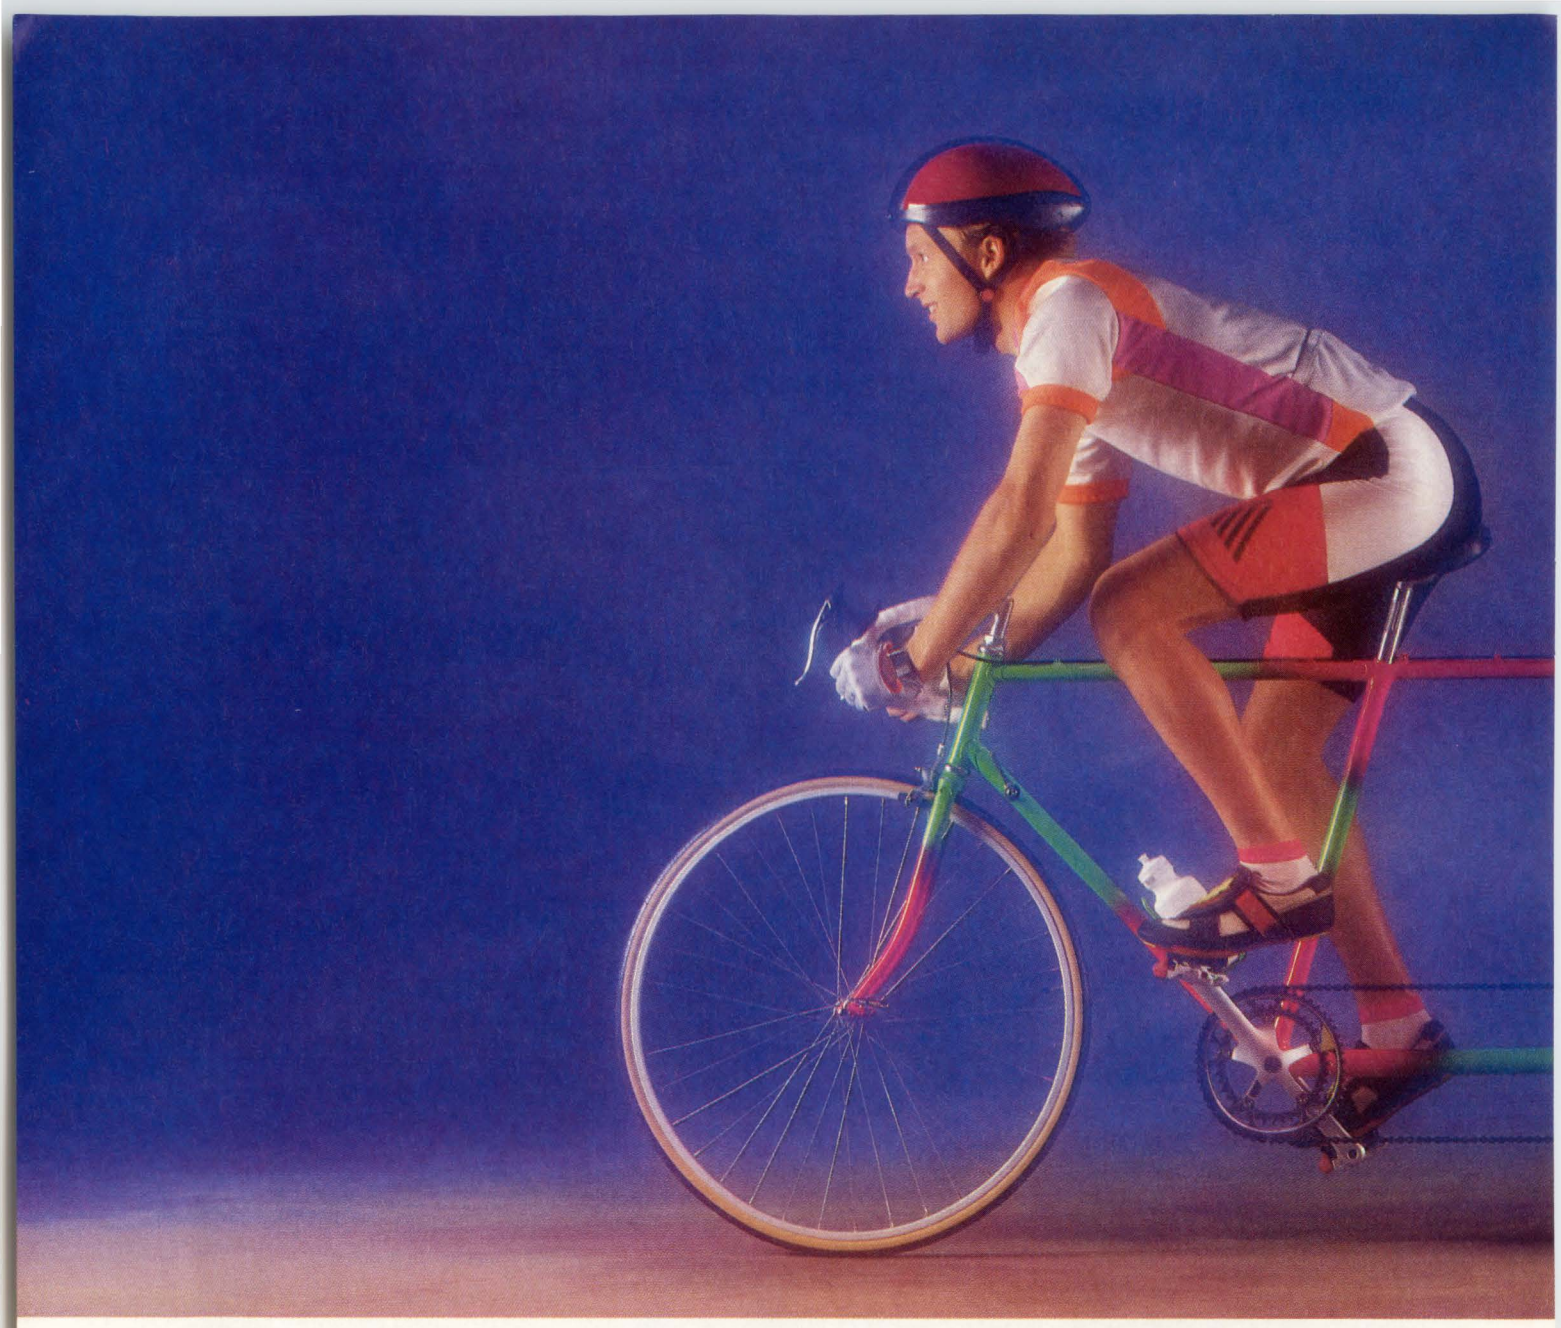

Until now, DOS and couldn't run

Introducing iRMX® for Windows: Real-Time Windows™ for PCs.

Running DOS or Windows* and real-time on a PC used to create some rather formidable challenges. After all, you could only run one at a time.

Fortunately, there's a whole new direction in real-time processing. It's called iRMX® for Windows. And it's a proven environment that lets DOS and iRMX software work together to generate powerful applications.

This reliable operating system takes full advantage of the protected-mode features of the popular Intel386™ architecture. So it opens new avenues for cost-effective solutions.

Of course, iRMX for Windows also opens some powerful options for software

AWARD

© 1991 Intel Corporation. iRMX is a registered trademark and Intel386 and Real-Time Windows are trademarks of Intel Corporation. *Windows is a trademark of Microsoft Corporation.
iRMX for Windows is an *EDN Innovation of the Year Award* winner for 1991.

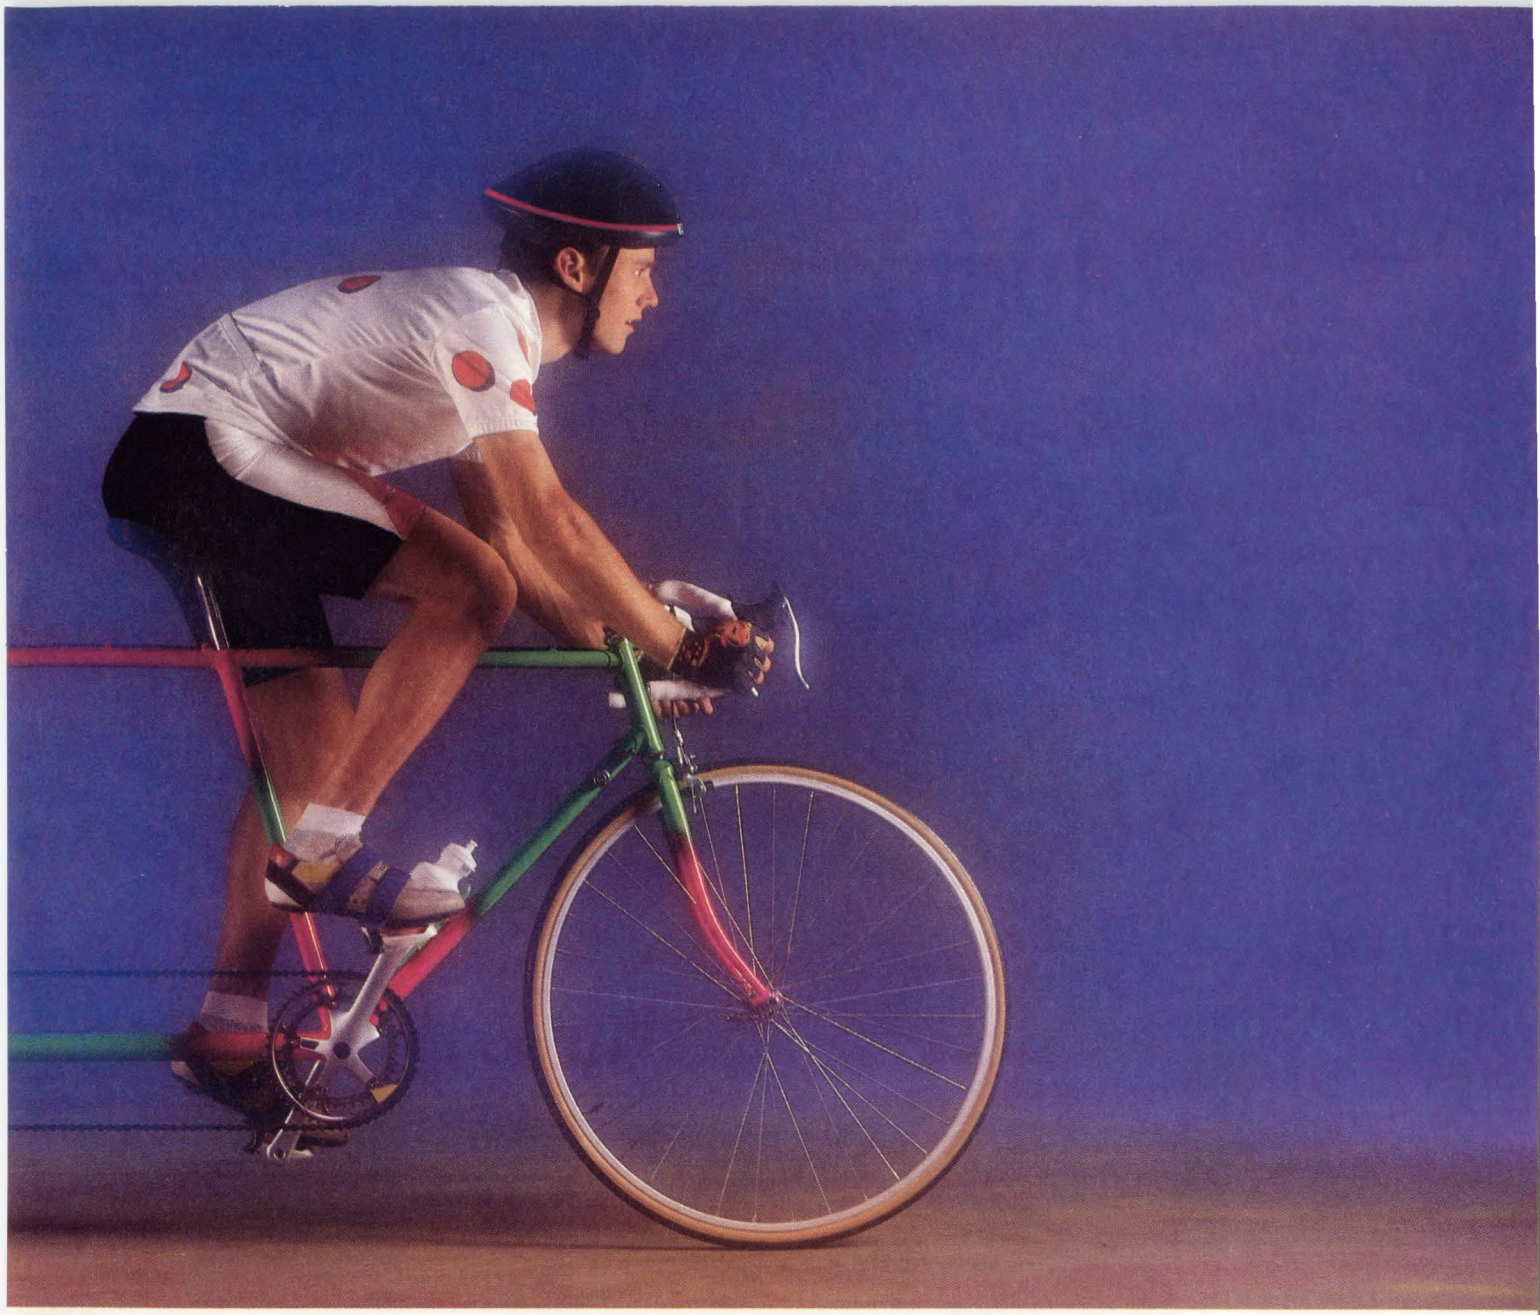

real-time applications concurrently.

engineers. Like bringing the huge installed base of DOS applications, tools, and the popular Windows environment to real-time development. What's more, as a developer, you'll receive professional support from Intel engineers.

All of which is important when you're racing to bring your application to market.

So make the shift to iRMX for Windows.

Call 1-800-GET-iRMX and ask for Lit. Packet #1A. You'll receive our Real-Time Response Kit. And give your next application some serious momentum.

intel®

The Computer Inside.™

TADPOLE CUSTOM DESIGN
FOR SINGLE BOARD COMPUTERS

Don't Compromise Your Next Design.

It's your product. Your market window.

So why compromise by settling for less than what you really need in a single-board computer?

When off-the-shelf boards don't meet all of your specifications, there's no reason to accept costly or inefficient workarounds. There's a better alternative: Tadpole Technology's Custom Design Services.

"Tadpole is extremely responsive, very thorough and, most importantly, played a key role in our development program."
— Mr. Tom LeBlanc, Principal Systems Engineer Polaroid Medical Imaging Systems

The OEM's Partner

Since 1985 we've helped OEMs like you build better products — and meet critical market opportunities — by designing and building custom computing engines that meet *exact* specs.

Our customers include some of the world's largest companies, who have turned to us with applications ranging from process control and flight simulation to application-specific graphics and airborne communications.

Built To Your Specs

CISC or RISC. VME or Multibus II. Real-time or UNIX®. Whatever your requirements, Tadpole has the expertise to design industry-standard, single-board computers that fit your application needs. We complement our design services with state-of-the-art semiconductor packaging techniques and advanced surface-mount assembly.

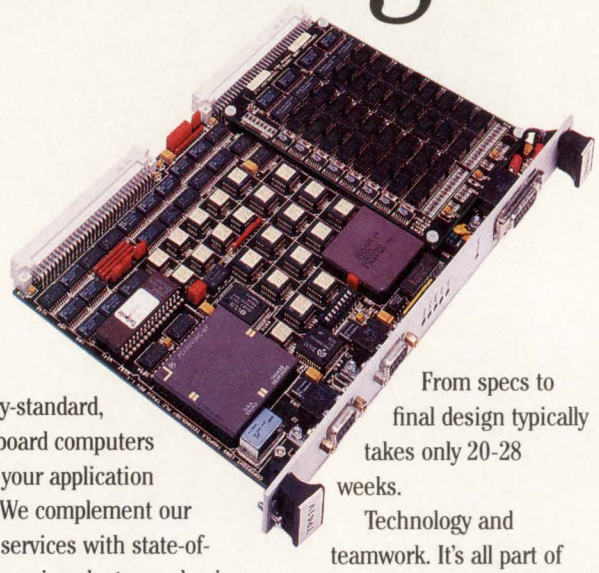

Integrated software is a major component of Tadpole custom design. We have one of the world's most experienced UNIX software porting teams, with 13 ports completed. We also support real-time operating systems.

The Team Approach

When you work with Tadpole, the Custom Design team serves as an extension of your company. Our team becomes *your* team, helping you optimize the use of your internal resources. This reduces your investment in product development and improves your time to market.

From specs to final design typically takes only 20-28 weeks.

Technology and teamwork. It's all part of the Tadpole approach to custom-built, single-board computers. And it all adds up to one thing: computing without compromise.

Should you settle for anything less?

For information on Tadpole's Custom Design Services, call in the U.S.:

800-232-6656

Cambridge, England
0223 423030

Paris, France
(1) 60 86 27 92

Tadpole Technology offers an extensive range of single-board computers and graphics processors, in addition to Custom Design Services.

UNIX is a registered trademark of AT&T.

Computing Without Compromise

T A D P O L E

CIRCLE NO. 20

Phillip A. Kaufman on: Prototyping alternatives

The two highest priorities of most product development teams are delivering the product to market as soon as possible and delivering a quality product that meets the customer's needs. The economic penalties for shipping late and missing a market window have been well documented many times. The penalty for delivering a flawed product is even higher. Pick up any trade paper and you'll find articles about how various companies lost their market lead or missed a major opportunity because they delivered late or had to recall a product for a "minor" problem.

Despite major improvements in design tools over the past five years, product development cycles are neither as short as they need to be nor as predictable as they should be. The principal reason for long, unpredictable design cycles lies not in the task of creating a design, but in the enormous time and effort required to perform comprehensive design verification and product integration. Products today are so complex, must meet such strenuous compatibility requirements and incorporate so much software that we're hard pressed to know how to perform adequate verification, much less accomplish the task in a reasonable amount of time and with a reasonable amount of effort.

Debug early

It's quite clear that the inevitable specification errors and ambiguities, as well as design implementation bugs, must be ferreted out as early in the development cycle as possible. As the bar chart on the next page shows, the cost of finding and correcting a problem increases dramatically as we move through the development cycle and into production. The solution to long and unpredictable development schedules, therefore, must be based on achieving the maximum possible amount of design verification and system integration as early in the development cycle as possible—that is, we must achieve concurrent design verification.

Today's complex products can't be adequately verified, much less carried through the full system integration phase, with simulation alone. A product must be fully exercised and perform all of its intended functions, including running all software and interfacing to its external environment, before we can say that the development task is done. The only solution to this problem is a full prototype of the product.

Every product prototyped

Every new product is prototyped. In many cases the prototype is the first article, using the first silicon for its ASICs. In some cases, a separate prototype is constructed well before an actual first article of the product exists. In either case, the first time that power is applied to a physical instance of the product and an attempt is made to validate all of its functionality, a prototype debugging process begins.

Whether this process is called design verification or system integration or quality assurance—whether or not the project plan formally includes a prototyping phase—prototyping always takes place. Problems are always discovered during prototyping and system integration, and design changes are always needed. To achieve rapid time-to-market, we must not only move the prototype phase to a point as early in the development cycle as possible, but we must also make the construction and debugging of prototypes as inexpensive and productive as possible.

A major challenge to any design project is to decide how and when to prototype. This is particularly crit-

Phillip A. Kaufman is president and CEO of QuickTurn Systems in Mountain View, CA.

ical for any product that incorporates ASICs or custom ICs. The choice of prototype technique, and when and how the prototype will be used, will have a substantial impact on the cost of development, time-to-completion and quality of results. There are several alternatives for the creation of a prototype. The pros and cons of each are examined below.

The major prototyping alternatives are:

- First article: Build the first real product as soon as possible and use it to validate the design.
- Scratch-built: Build up a prototype from standard parts and programmable logic well before a real first article exists.
- Computer-aided prototyping (CAP): Very early automatic implementation of designs that do not yet exist as hardware (for example, ASICs) combined with any actual hardware that already exists.

Goals of prototyping

Each of these alternatives has pros and cons. To evaluate them it's necessary to address the goals, objectives and constraints that are involved. The number one design verification goal must always be to find as many problems as possible as early as possible. To accomplish this, the prototype must be available very early in the design cycle.

The second goal is to make using the prototype as productive as possible. Productive debugging means assuring that when a problem is discovered it's easy to isolate the cause and correct the prototype implementation. To accomplish this one must create a prototype that has very high controllability and observability. The one absolute truth about prototypes is that bugs will be found and changes will have to be made. If it takes days, or weeks, or even months to effect a change in a prototype, it will be very difficult to uncover multiple bugs in any reasonable amount of time.

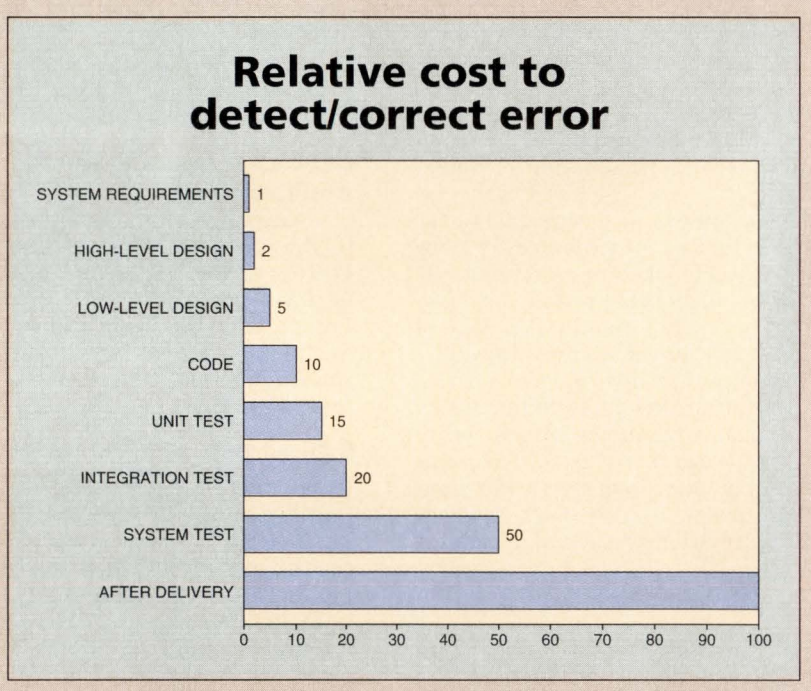

Another key goal of prototyping is, of course, to minimize costs. Prototype costs come in two parts—the cost of creating the prototype and the cost of using it. Prototype creation costs are relatively straightforward. They consist of the actual hardware and software that must be procured to construct the prototype, plus the manpower costs to design and build it. The cost of using a prototype is somewhat harder to quantify. The most obvious piece is the cost of making a change to the prototype when a bug must be corrected. The difficult part is the time and dollar cost of low debugging productivity.

Another attribute of prototyping is its verification span—aspects of the design that can be verified with each type of prototype, such as functionality, timing, electrical, and mechanical. There's a definite trade-off between how much can be verified and when it can be done. A full verification of a product requires having the product in its final form, a first article. By definition, a first article requires waiting until very late in the development cycle. On the other hand, early prototypes which may allow full verification of some critical aspects, such as a functionality, will likely not be able to have the full form fit and speed of the final product. If these early prototypes enable finding many critical problems, however, they can result in a substantial improvement in schedule. The table on the next page summarizes various attributes versus type of prototype.

First article prototypes

The most common form of prototype today is the first article. It's so common that many projects don't even consider the first article as a prototype. Rather, evaluating the first article is considered to be a system integration or quality assurance phase. It's not uncommon for development projects to schedule six to nine months after a complete system is available for system integration. For mainframe developments, it's not uncommon to schedule up to two years of system integration on the "bring-up-floor."

Using a first article to do design verification is a very inefficient activity. A first article has two major limitations; it can't be available until all design is finished and all parts are fabricated and assembled, and debugging and correcting errors with a first article are extremely difficult. In addition, use of a first article means that we're beginning full design verification just when manufacturing should be ramping up volume production. Any problem found not only causes a long recycle time for redesign and reimplementation, but has tremendous impact on manufacturing and product introduction.

The limitations of first articles are particularly severe when ASICs or custom ICs are involved. The expense and time to spin new chips drives the development team to attempt exhaustive design verifica-

Key prototype attributes versus type of prototype			
	First article	**Scratch-built**	**Computer-aided**
Availability	Very late	Early	Very early
Problem isolation	Very difficult	Difficult	Easy
Ease of changes	Very hard	Hard	Easy
Construction costs Material or tools Labor	 High Low	 Low High	 High Low
Verification span Functionality Speed Timing Electrical Mechanical	 Yes Full Yes Yes Yes	 Yes Very fast No Partial Partial	 Yes Fast No Partial Partial

tion using simulation before tape-out. While this approach finds some problems earlier, the simulation effort tends to grow unbounded. How much simulation is enough? Simulation can never replace full system integration and there will likely still be many problems discovered with the first article.

The importance of getting to the first article rapidly is reflected in the way that most projects are managed. The most critical schedule milestone is very often tape-out. Simulation is continued until we run out of time and then we rush to get chips. This approach, however, is backward from what should be happening. In the ideal situation, all verification and system integration take place prior to tape-out, and the time from tape-out to production shipments is minimized.

Scratch-built prototypes

A real hardware prototype is a very valuable design verification tool. It can potentially be available quite early in the design cycle, and it can be quite productive. A scratch-built prototype is very effective if it can be built with a small number of components and without the prototype development itself becoming a major project. The major down side of scratch-built prototypes is the difficulty of making design changes. Where the design is small and doesn't contain any ASICs or custom chips of any significant complexity, a scratch-built prototype is often the best solution.

When there's much logic, the practicality of a scratch-built prototype diminishes rapidly. Although programmable components help significantly, their use in scratch-built prototypes creates many problems. Converting a design from a medium- or high-complexity ASIC to programmable components is a major effort. The primitives don't match well, and the number of programmable components required is usually far higher then one would estimate based upon the specification for today's programmable components. The result can easily become a design project that's larger than the actual product development project.

Computer-aided prototyping

Computer-aided prototyping (CAP) is a relatively new technology. The essence of CAP is to automate the conversion of ASIC and custom-IC designs into an implementation using reprogrammable logic and to incorporate into the implementation debugging tools and incremental capabilities that make finding and correcting design problems very easy. CAP overcomes the problems of late availability and rigidity associated with first articles. It also overcomes the enormous effort required to implement complex chip designs into programmable parts manually. CAP is the only way to move very quickly from high-level design to functioning hardware.

There are actually two different ways to utilize CAP for system prototyping and debugging. With CAP capabilities now available for up to one million gates, entire systems can be implemented. The standard components in the system can be plugged into the CAP prototype in a manner analogous to the use of hardware modelers with simulation. Alternatively, CAP can be used just for the ASIC and custom-IC parts of the system and can be combined with scratch-built or first-article prototypes of the rest of the system.

The use of CAP does entail some limitations. As with any design tool, a capital investment is required. Speed, while three to seven orders of magnitude faster than simulation, will be somewhat slower than using first articles, because of the use of many reprogrammable components. Nonetheless, effectiveness and payback on investment have been demonstrated to be very high by many successful customers.

In every hardware development there's a point at which prototyping comes into play as the final design verification and system integration vehicle. With the advent of programmable logic and CAP, designers now have many more choices about how and when to prototype. Made properly, these decisions can make a dramatic difference in time-to-market and product quality.

CAE/CAD TOOLS

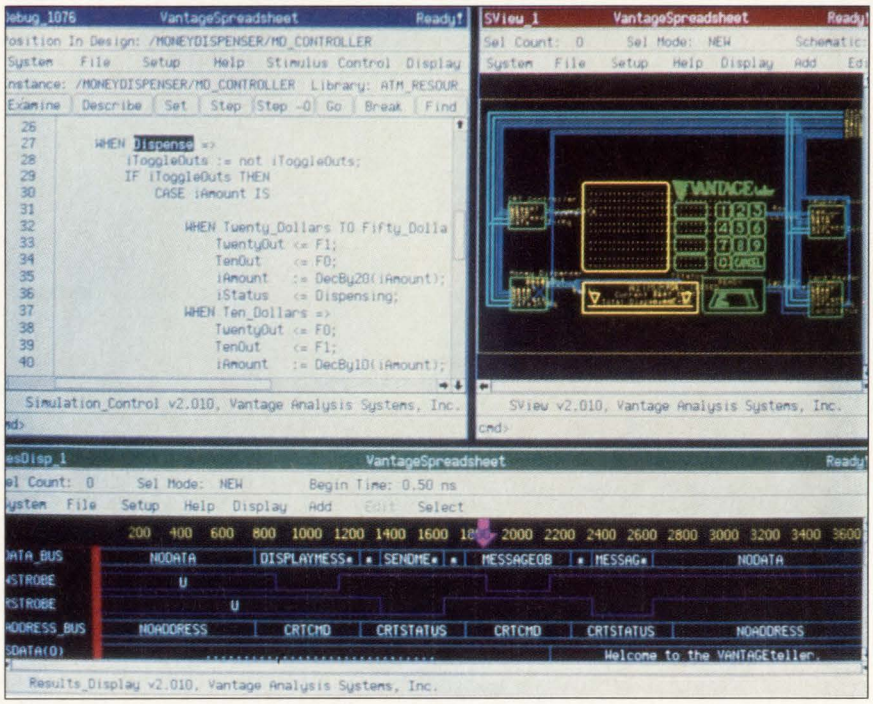

VHDL simulations, unlike traditional computer programs, can have numerous VHDL models running in parallel. This photograph of Vantage's Debug 1076 shows a behavioral model of an automatic teller machine, the VHDL code and simulation waveforms.

Newest VHDL tools encourage conformity to standard

Mike Donlin, Senior Editor

Following in the footsteps of their counterparts in the ASIC world, system designers are turning to hardware description languages to simulate their designs. VHDL, in particular, is emerging as a de facto standard which lets you simulate everything from multichip modules (MCMs) to PCBs and complete systems.

But if VHDL is to enjoy continued success, designers must master the fine art of writing accurate models, the linchpins in the top-down design paradigm. These VHDL models must include a broad range of characteristics to provide simulation results that are accurate, efficient and, most important, synthesizable into actual hardware. Educating hardware engineers to write software code that fulfills all of these requirements remains a challenge.

"You have to take a lot of things into account when you're assembling a team to write VHDL code," says William D. Billowitch, president and CEO of the VHDL Technology Group (Allentown, PA). "If you don't establish standards for model development, then everyone will go off in their own direction and you'll end up with spaghetti code that's impossible to maintain. Remember, a good deal of the risks of systems integration are painfully realized too late in the design cycle. By standardizing procedures, you can avoid many of the pitfalls caused by semantic differences in modeling techniques, as well as interface differences caused by a diversity of modeling methodologies."

Sticking to the standard

To try to keep designers on this straight and narrow path of adhering to standard design practices, CAE vendors provide modeling software packages that are compliant with the IEEE-1076 VHDL standard. The VHDL Technology Group, for example, has released a package called Std Developers Kit, which contains building-block routines for VHDL model development. These routines let you stay within 1076 guidelines while developing models from the macrocell to system level. Parameters such as min/typ/max timing, I/O routines and arithmetic operations can be incorporated into models within guidelines that are designed to keep design teams working from a common set of standards.

Though it would seem that such regimentation is unnecessary if all members of a modeling team are writing to an IEEE standard, the opposite is true. "In an ideal world, all simulators that boast VHDL compatibility would work with models that adhere to the standard," says Billowitch. "But in reality, some simulators are more compliant than others. There's a level of compatibility that's essential or model performance degenerates quickly. If your simulator is 90 percent 1076-compatible, you can be pretty sure that models written with our Developers Kit will work. Adhering to standards also gives you a better chance that your models will be portable to simulators from different vendors."

Room for diversity

The main problem in writing coherent VHDL models is adhering to the parts of the standard that any one vendor chooses to embrace. There's general agreement among leading simulator vendors about which portion of the VHDL standard is essential for accurate simulation, but the remaining 10 percent can cause problems, particularly when models are written for simulation rather than synthesis. "If you're dealing with simulators that support the whole standard, it's easier to write models that will simulate accurately and synthesize efficiently," says Steve Carlson, manager of methodology consulting at Synopsys (Mountain View, CA). "Unfortunately, there aren't many simulators that do. Remember, VHDL was written for the whole design process, and that goes beyond the basic functional description of what you want your design to do. There are constructs for error checking, for example, which you don't necessarily

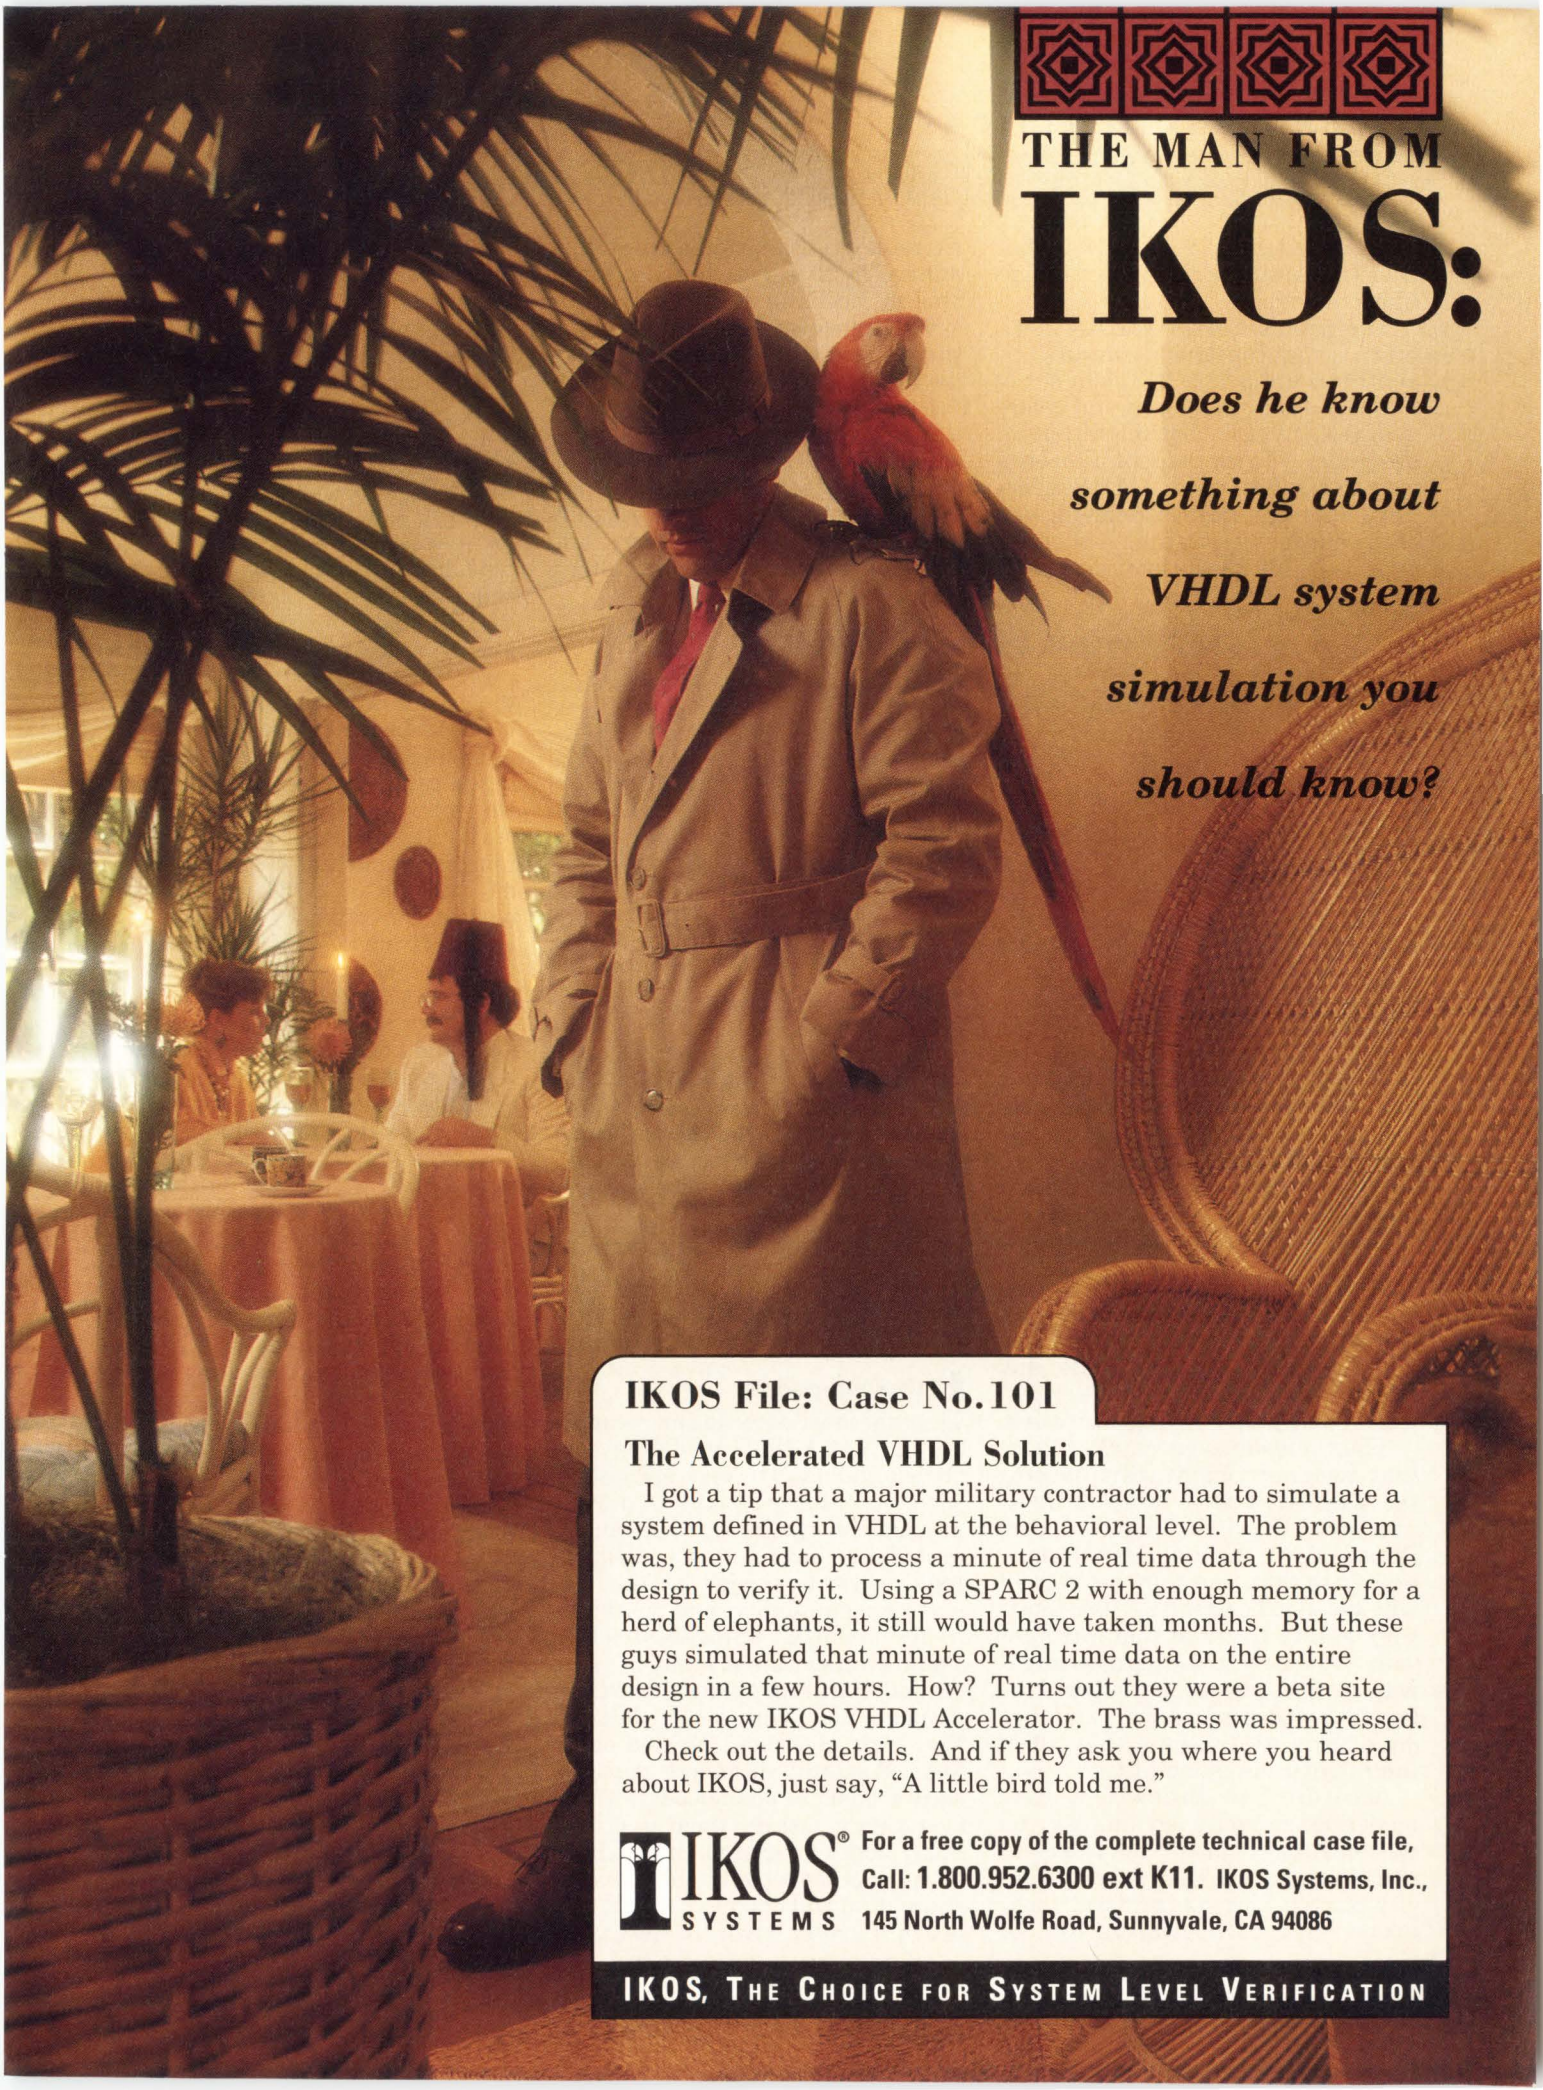
THE MAN FROM
IKOS:
Does he know something about VHDL system simulation you should know?
IKOS File: Case No. 101
The Accelerated VHDL Solution
I got a tip that a major military contractor had to simulate a system defined in VHDL at the behavioral level. The problem was, they had to process a minute of real time data through the design to verify it. Using a SPARC 2 with enough memory for a herd of elephants, it still would have taken months. But these guys simulated that minute of real time data on the entire design in a few hours. How? Turns out they were a beta site for the new IKOS VHDL Accelerator. The brass was impressed.
Check out the details. And if they ask you where you heard about IKOS, just say, "A little bird told me."
IKOS® SYSTEMS
For a free copy of the complete technical case file, Call: 1.800.952.6300 ext K11. IKOS Systems, Inc., 145 North Wolfe Road, Sunnyvale, CA 94086
IKOS, THE CHOICE FOR SYSTEM LEVEL VERIFICATION

CAE/CAD TOOLS

want included in your hardware but that you want included in your simulation model."

The problem with VHDL's flexibility is that constructs that can't be synthesized have to be taken out of a model after simulation and before synthesis, a task that might put off designers who are reluctant to write code in the first place. CAE vendors have different ways of relieving you of the burden of rewriting models for synthesis, but they sometimes must exact a performance penalty for convenience. Some synthesis tools have a preprocessor which warns you when code isn't synthesizable, and others simply give error messages, while some stop running altogether.

"The consequence of all this is that designers might start limiting the ways that they write models," cautions Billowitch. "Some vendors throttle down their simulation capabilities so that a model written for it will be synthesizable. That limits your ability to write constructs for your test bench. Designers might want to include parsing algorithms for test pattern generation, and they don't want restrictions that affect their synthesis capabilities."

According to Billowitch, good modeling software should let you concentrate on model development and design rather than development of tools that let you perform those tasks. The Std Developers Kit, for example, contains guidelines that adhere to 1076 standards to assist you in building timing into your VHDL models. Because pin-to-pin timing is critical in simulating today's submicron technologies, models must accommodate a diversity of timing representations at the macrocell and board levels. The timing package in the Developers Kit provides a mechanism for you to specify burned-in actual timing data and lets you supply your own timing data. The tool also includes constructs that let you describe environmental factors affecting timing.

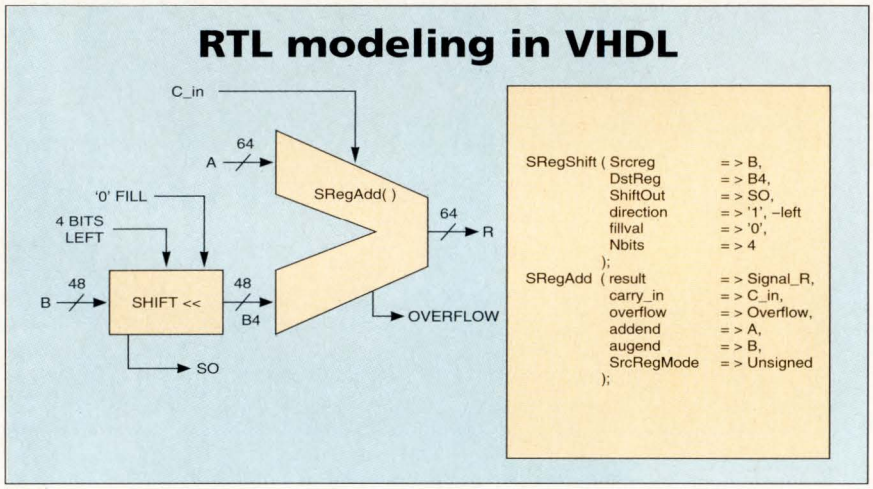

The VHDL Consulting Group provides over 200 routines in its Std Developers Kit for modeling digital architectures at the register-transfer-level (RTL). The drawing on the left represents a microarchitectural view of a datapath with models running concurrently. The code on the right, generated by the kit, shows the circuit's VHDL implementation; the user invoked the code with simple procedure calls.

Tailoring models

Though most simulation vendors agree that VHDL modeling constructs shouldn't be severely limited, some hold that it's better to provide a seamless synthesis environment, even if it means limiting modeling or requires you to provide different versions for simulation and synthesis. "This issue of being restrictive sounds worse than it really is," says Roland Mattison, product marketing manager for synthesis products at Viewlogic Systems (Marlborough, MA). "Usually, designers partition their models to accommodate whatever they're asking them to do. An ASIC model, for example, might have a top-level model with one block for simulation and one block for the test bench. Further down the design cycle, there might be 10 blocks which include VHDL and schematics. A designer who's savvy enough to use full VHDL knows when to use which constructs."

Some simulation vendors provide debug capabilities to help you isolate problems during simulation. The Debug 1076 tool from Vantage Analysis Systems (Fremont, CA) lets you set breakpoints for subprocesses, subprograms, source-code line numbers, or time points generated by the main simulation. VHDL model source code is displayed as a simulation run in Vantage's SpreadSheet to let you track the execution of the simulation and debug a design as needed.

Though programs such as Vantage's could be used to debug models, many designers use commercially available models from companies such as Logic Automation (Beaverton, OR) or hardware models generated by systems such as those from Logic Modeling (Milpitas, CA). For those designers, there's no need to debug models, so most debug work centers around glue logic and basic design constructs. "I would contend that most designers don't synthesize models most of the time anyway," says John Willey, vice-president of marketing for Vantage. "ASIC vendors or model vendors don't want you messing around with their models or synthesizing them into something else. If people do want to synthesize their models, though, they have to be careful. First of all, you have to be aware that the model you're writing is synthesizable or it's just not going to work. Second, you have to write the model for the particular synthesis tool that you're going to use. Some synthesis tools are better suited for some applications than others." ■

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

Logic Automation
(503) 690-6900 ..... **Circle 240**
Logic Modeling
(408) 957-5200 ..... **Circle 241**
Synopsys
(415) 962-5000 ..... **Circle 242**
VHDL Technology Group
(215) 266-9791 ..... **Circle 243**
Vantage Analysis Systems
(510) 659-0901 ..... **Circle 244**
Viewlogic Systems
(508) 480-0881 ..... **Circle 245**

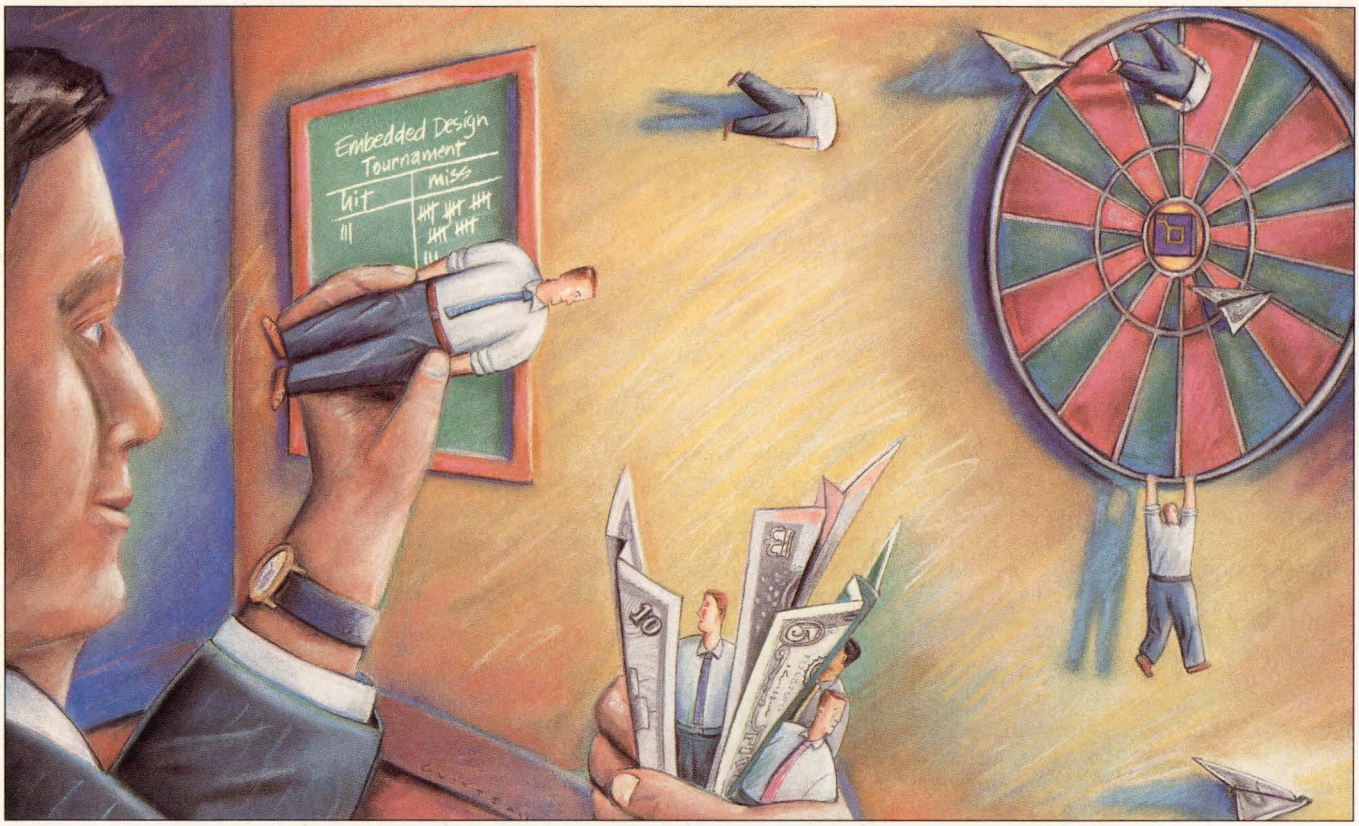

Maybe you're throwing the wrong things at your development problems.

You're looking at a fairly common approach to solving sticky embedded design problems.

Hurling more money and engineers may have worked on smaller-scale projects. But it often falls short of the mark with today's more complex 16- and 32-bit designs.

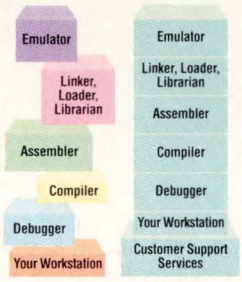

Mismatched, nonintegrated development tools can throw off a design team's efforts in no time.

At Applied Microsystems, we're helping software and hardware teams develop better embedded systems more efficiently with high-performance, integrated development systems.

With our tools and support, your project's cost is more manageable. Your design team is more productive. And the end product is more of what you envisioned from the beginning.

This may all sound too good to be true, but Applied Microsystems has the solutions and experience to make it real. Witness the fact that we've set up over 15,000 development systems worldwide.

What to look for in a development systems company. And what to look out for.

To learn more, call us for details or free literature at 1-800-343-3659 (in WA, 206-882-2000).

We'd like to throw a few solutions your way.

Applied Microsystems Corporation

Where it all comes together.

In the U.S. and Canada, call 1-800-343-3659 (in Washington, 206-882-2000). Or contact Applied Microsystems Corporation, P.O. Box 97002, Redmond, WA 98073-9702 USA. **For the name of your nearest distributor in Europe, call 44-(0)-296-625462.** Europe Fax 44-296-623460. Or contact Applied Microsystems Corporation, Ltd., AMC House, South Street, Wendover, Aylesbury, Bucks, HP22 6EP, England. **In Japan, call 03-493-0770.** Japan Fax 03-493-7270. Or contact Applied Microsystems Japan, Ltd., Nihon Seimei, Nishi-Gotanda Building, 7-24-5 Nishi-Gotanda, Shinagawa-KU, Tokyo T141, Japan. ©1991 Applied Microsystems Corporation. All rights reserved. Other names indicated by ™ and ® are registered trademarks of their respective holders. AMC-42.

CIRCLE NO. 21

INTEGRATED CIRCUITS

DRAM/SRAM combo chip cures cache headaches

Dave Wilson, Senior Editor

Traditional microcomputer designs sport a hierarchy of solid-state memory. First, there's usually a small amount of fast memory on the processor itself. This is complemented by both static cache memory and slower DRAM memory residing on the system motherboard. Breaking with this tradition is a new memory architecture from Mitsubishi (Sunnyvale, CA).

Dubbed the CDRAM (cached DRAM), it integrates a DRAM and an SRAM on the same chip. "For PCs and X Terminals, there's enough bandwidth to support second-level cache, main memory and video frame buffers," says Chris Tennant, a Mitsubishi applications engineer.

No buffers needed

The device comprises a 1 Mbit × 4 DRAM and a 4-kbit × 4 SRAM. Capable of 100-MHz performance, it can be coupled directly to the CPU, eliminating the need for buffers. The synchronous clock design of the CDRAM also qualifies the device as the industry's first synchronous DRAM.

The CDRAM transfers a 16-line block (16 × 4 bits) between the array and the cache in one cycle on a cache miss. This high bandwidth eliminates the bottleneck in transfer time between cache and DRAM. The large number of internal bus lines between the SRAM and the DRAM permits a large block size, so the device, without incurring any miss penalty, can achieve a higher hit rate than that of a separate DRAM and SRAM.

The CDRAM implements a fast copy-back scheme that results in a miss cycle access time approximately ⅓ that of a conventional copy-back method. Two internal data transfer buffers (DTBs) facilitate the scheme. The missed data not found in the SRAM cache is latched to one DTB, while the expected data is transferred simultaneously from the DRAM to the SRAM through the amplifier of the second data transfer buffer. This makes the expected data available in one DRAM access time.

Others in the wings?

If the CDRAM is so hot, are other manufacturers racing to build parts just like it? David Chapman, applications manager for fast SRAMs (FSRAMs) at Motorola (Austin, TX), says that his company doesn't have such a product in its plans at present, but he doesn't discount the possibility that it may at some point. He admits that the part "looks like a neat idea for some applications."

Such applications would be in systems with a limited number of bus masters. That's because one of the main chores an SRAM cache performs is isolating the system bus from the processor so that the bus can be used by the DMA or I/O controllers. "If you use a CDRAM," says Chapman, "you lose that capability. The CDRAMs would be used for the main system storage. So the processor would spend all its time on the processor/main memory bus instead of isolated from it."

On the other hand, the part may be useful in embedded applications where there aren't multiple bus masters vying for time on the system bus—but designers may still need to build a system around a high-performance processor, implying a high bandwidth between the processor and memory.

Ajit Deora, a vice-president of engineering at Sand Microelectronics (San Jose, CA) and a designer of high-performance 50-MHz i486 PC motherboards, thinks that—at least conceptually—Mitsubishi's CDRAM looks attractive from a price/performance perspective. "There are many system configurations that are possible with the device compared with the standard DRAMs and SRAMs that are available today," he says.

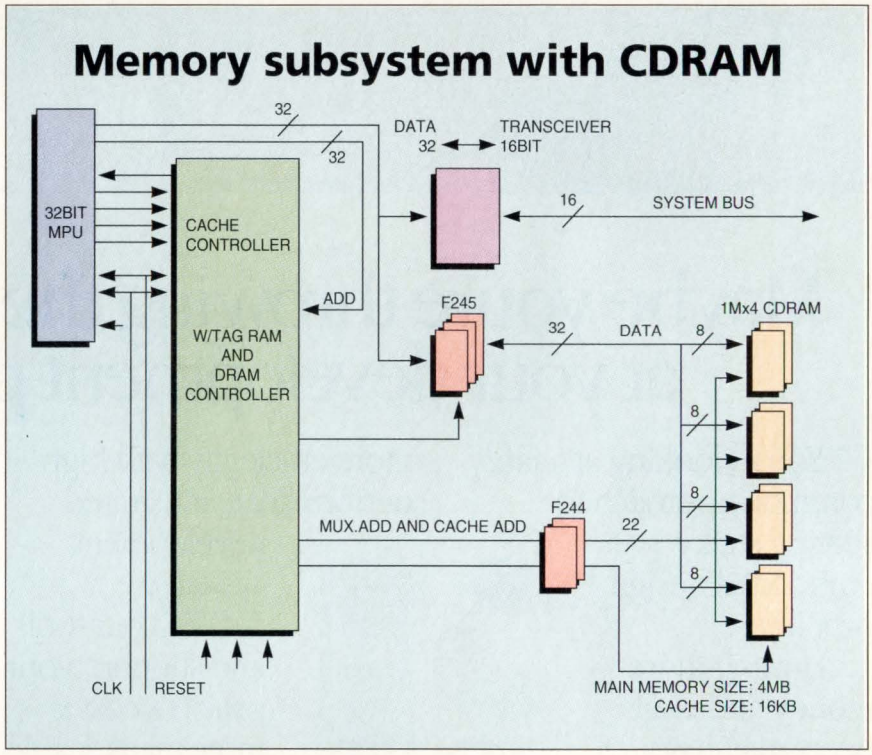

A conventional memory subsystem consists basically of the microprocessor unit (MPU), cache controller and tag RAMs. They interact with a separate SRAM cache and DRAM controller (not shown). A memory subsystem using the CDRAM requires no external SRAM for cache and internalizes the DRAM controller.

In Search Of The Lost Generation In Desktop Video

When outputting to tape, the Bt858 avoids image degradation and maintains the superior video quality as represented here.

When outputting video from desktop to tape, something gets lost in the translation. That something is **image quality** — which is why the resulting picture resembles a second generation copy of the original material.

Image degradation can be avoided, however, by outputting with an encoder at least one level higher than the tapedeck. Since most desktop video users rely on corporate-level video equipment with a 4 MHz bandwidth and a SNR of 43-50 dB — such as S-VHS — a studio grade encoder with a 5 MHz bandwidth and a SNR of 50-62 dB is necessary to maintain corporate-level quality.

Such an encoder is now available with Brooktree's Bt858 — the industry's **first all-digital, studio-grade NTSC/PAL encoder** — which maintains image quality at a level identical to the original material.

Utilizing high-quality digital filters to minimize color artifacts, the Bt858 eliminates hue and saturation errors that generally accompany encoding and achieves greater color accuracy with more precise sub-carrier generation. In addition, it employs rounding algorithms rather than truncation to minimize cumulative processing errors and uses controlled edges to maintain rise and fall times within NTSC/PAL specifications.

A videotape, **"The Ins and Outs of Video Out,"** discusses various aspects of the Bt858 in greater detail. Call Brooktree at 1-800-VIDEO IC for your free copy.

CIRCLE NO. 22

INSTANT DATA ACCESS (IDA) DIAL (617) 494-8338 DOCUMENT NO. 1026

Why Wait For ISA XGA?

Use Brooktree's Bt484 RAMDAC with a VRAM GUI accelerator to enhance Windows capabilities *and* pick up many of the same features found in the **XGA standard.** You get 8-24 bit color flexibility and display size up to 1024 x 768 — all for about **30 percent less** than the price of an XGA chip set.

INSTANT DATA ACCESS (IDA) DIAL (617) 494-8338
DOCUMENT NO. 1029

CIRCLE NO. 23

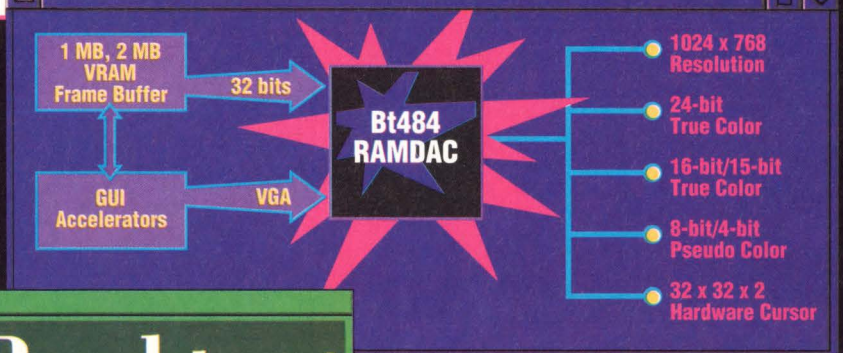

Brooktree®

Need technical information?

CALL 1-800-VIDEO IC

Brooktree Corporation
9950 Barnes Canyon Road
San Diego, CA 92121
(619) 452-7580
FAX (619) 452-7294

All trademarks and registered trademarks are owned by their respective companies. ©1992 Brooktree Corporation.

40 MB/SEC
DUAL-PORT VME/VSB HIGH - DENSITY MEMORY

from

Chrislin Industries, Inc.

Your

TOP QUALITY Memory Supplier for OVER 16 YEARS

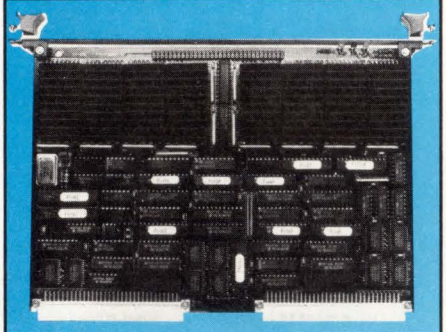

CI-VME40

- FAST ACCESS TIMES
  20ns READ/WRITE BLOCK CYCLE
- FAST CYCLE TIMES
  83/62NS READ/WRITE BLOCK CYCLE
- VERSATILE CONFIGURATIONS
  4, 8, 16, 32 or 64MB in one 6U slot
- VERSATILE ADDRESSING
  Addressable in 24 or 32 bit through 4 Gigabytes
  Memory start and end addresses selectable on 256kb boundaries
  VME/VSB configured independently
- RELIABLE
  VME Revision C.1 Compatibility
- LIFETIME WARRANTY

OTHER QUALITY VMEbus MEMORIES AVAILABLE ARE:

THE CI-VMEmory
LOW-COST VMEbus BYTE PARITY MEMORY with 4, 8, or 16MB in one 6U VMEbus slot

THE CI-VSB EDC
Dual-port VMEbus/VSB memory with Error Detection and Correction, single-bit error detection/correction, double-bit error detection, 4MB up to 64MB in one VMEbus/VSB slot

Chrislin Industries, Inc.
31312 Via Colinas, Suite 108
Westlake Village, CA 91362
TEL: (818) 991-2254 FAX: (818) 991-3490
CALL TOLL FREE: (800) 468-0736

CIRCLE NO. 24

INTEGRATED CIRCUITS

Cost/performance comparison

Performance	SRAM & DRAM	CDRAM
Cache size	128 kbytes	16 kbytes
DRAM	4 Mbytes	4 Mbytes
Block size	4 bytes	64 bytes
Tag size	32 kbits x 8	256 kbits x 8
Miss rate	4%	4%
Board area	10,500 mm^2	7,500 mm^2
Avg. access time	17.6 ns	17.6 ns
Price	$193	$124

A cost/performance comparison of the CDRAM with the SRAM and DRAM configurations shows that, in comparable systems, the CDRAM has faster average access time, uses less board space and is less expensive.

"For PC applications, where the memory performance is the biggest issue, it's very attractive."

Nevertheless, Deora agrees with Motorola's Chapman that the device is unlikely to be used in file-server applications. "In a file server where there's a lot of DMA, it would be a problem unless the device was dual-ported—one port to the CPU and one to the DMA device," he says. On the other hand, the device may be fast enough in such an application to let the bus interface be multiplexed between bus masters without paying a performance penalty.

Both price and size, of course, are important design considerations for designers such as Sand's Deora to consider. The Mitsubishi device is offered in a 300-mil, 44-pin TSOP (thin small-outline package) with an 0.8 mm lead pitch. That's approximately seven percent longer than the standard Mitsubishi 4-Mbit DRAM. The increase in package size is more than offset, according to the company, by the number of devices needed in a system. Samples of the devices will be available in the second quarter, with production in the third quarter. Depending on cache access speed, the device will be priced at 10 to 20 percent more than the prevailing price for 4-Mbit DRAMs. In 100-piece quantities, the price will be $15 each for 20-ns parts with 80-ns DRAM arrays, $15.50 for 15-ns parts with 75-ns arrays and $16.20 for the 10-ns parts with 70-ns arrays. ■

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

Mitsubishi Electronic Device Group
(408) 730-5900 ..... Circle 254
Motorola
(512) 928-6000 ..... Circle 255

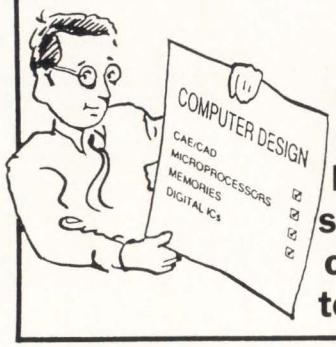

THANK YOU!

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

Use our low power, low voltage memory and buy more time.

If you design battery operated systems such as laptop and notebook computers, Micron's low power, low voltage memory components can buy you more time.

Micron's extensive line of leading edge, low power, low voltage DRAMs and SRAMs are designed to give you optimum 3.3 volt operation in battery powered systems, extending battery life and system operating time. Since the parts run cooler, they also increase system reliability. And we offer the latest packaging technologies such as TSOP and PQFP.

So call Micron today at 208-368-3900. And find out how to buy more time.

Micron. Technology that works for you.

Part Number	Memory Configuration	Availability
3.3 Volt, Low Power, Extended Refresh DRAMs		
MT4C4001J VL	1 Meg x 4	3Q92
MT4LC4001 S*	1 Meg x 4	4Q92
MT4LC4001 L	1 Meg x 4	4Q92
MT4C4256 VL	256K x 4	Now
5 Volt, Low Power, Extended Refresh DRAMs		
MT4C1004J L	4 Meg x 1	Now
MT4C4001J L	1 Meg x 4	Now
MT4C8512 L	512K x 8	3Q92
MT4C16256 L	256K x 16 DW[1]	3Q92
MT4C16257 L	256K x 16 DC[2]	3Q92
MT4C1024 L	1 Meg x 1	Now
MT4C4256 L	256K x 4	Now
MT4C1664 L	64K x 16 FPM[3]	Now
MT4C1670 L	64K x 16 SC[4]	Now

Part Number		Memory Configuration	Availability
3.3 Volt, Low Power Specialty SRAMs			
MT5LC2818	Latched	16K x 18	Now
MT58LC1618	Synchronous	16K x 18	Now
MT5LC2516	Latched	16K x 16	Now
MT58LC1616	Synchronous	16K x 16	Now
5 Volt, Low Power, Low Voltage Data Retention SRAMs			
MT5C1001 LP		1 Meg x 1	Now
MT5C1005 LP		256K x 4	Now
MT5C1008 LP		128K x 8	Now
MT5C2561 LP		256K x 1	Now
MT5C2564 LP		64K x 4	Now
MT5C2565 LP		64K x 4 $\overline{OE}$[5]	Now
MT5C2568 LP		32K x 8	Now

*Self Refresh [1]DW– Dual Write Enable [2]DC– Dual $\overline{CAS}$ [3]FPM– Fast Page Mode [4]SC– Static Column [5]$\overline{OE}$– Output Enable

MICRON TECHNOLOGY, INC.

2805 E. Columbia Rd., Boise, ID 83706 (208) 368-3900
Customer Comment Lines: U.S. 800-932-4992; Intl: 01-208-368-3410

© 1992 Micron Technology, Inc.

CIRCLE NO. 25

Build on our Digital faster ASICs faster.

System Arrays™ for

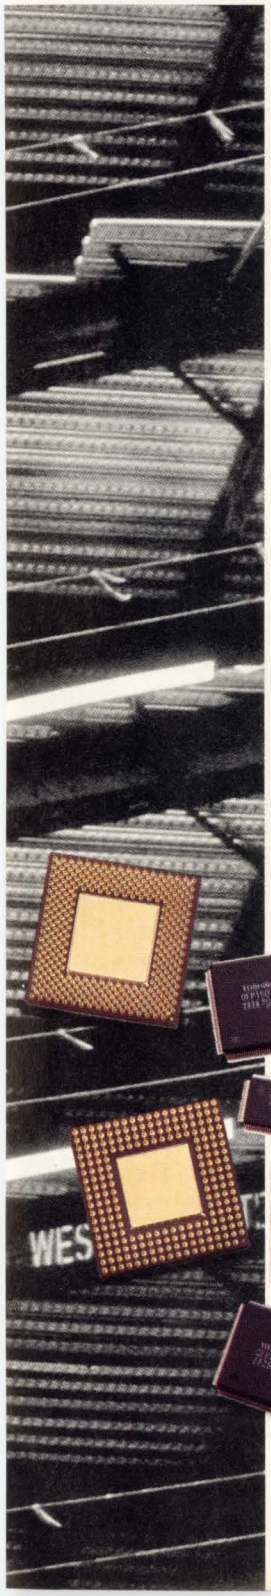

Optimize for speed, testability, functionality or high pin/gate ratios. You're the architect.

Toshiba's Digital System Arrays (DSAs) provide the superstructure system designers need to optimize ASICs for market and application.

Which Digital System Array is right for you?

Optimize for performance and testability.
Our new TC165G/E, architected by Vertex and rendered in Toshiba's proven 0.8 micron CMOS technology, show the power of DSA.

These new ASICs offer the highest usable gate count in the industry—up to 250K. Gate delays of <260 picoseconds combined with Vertex's proven performance-enhancing techniques and design methodologies optimize these arrays for high clock rates. And automatic design-for-testability achieves >99% fault coverage with no performance penalty.

TC165G gate arrays are complemented by the TC165E embedded array series, offering a wider variety of functionality. With embedded functions, the TC165 series reaches new heights in system integration.

Optimize without compromise.
Toshiba's TC160G/E provide the optimum combination of features and cost effectiveness for most applications.

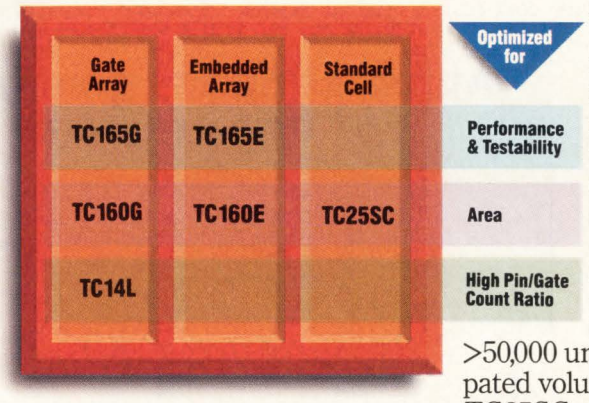

You can't optimize without options. Choose the Toshiba solution that's best for you.

Employing sea-of-gates architecture, the family offers up to 210K usable gates at 300 picoseconds.

Toshiba mega-functions and super integration libraries make the embedded version a high density, silicon efficient choice.

Optimize for efficiency or for I/O.
Efficient gate utilization your top priority? >50,000 units/month your anticipated volume? Look into our TC25SC standard cells. They put an area-optimized 0.8 cell library at your command.

If you've ever had to select a gate array with more gates than you needed just to get enough I/O pads, our TC14L is for you. It provides the highest available pin-to-gate ratio, offering up to 20K gates with reduced pad spacing.

Optimized Design Environment.
Our Digital System Arrays are supported by a designer-friendly design environment.

We support all popular libraries, including Verilog, Synopsys, IKOS, Dazix, Mentor, ViewLogic and Valid.

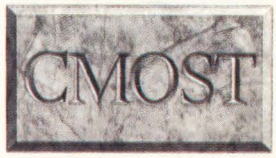

CMOST is the cornerstone of Toshiba's Unified Device Architecture.

The sky is the limit with Toshiba's DSA family. *Please call 1-800-321-1718 for technical literature.*

In Touch with Tomorrow

TOSHIBA

TOSHIBA AMERICA ELECTRONIC COMPONENTS, INC.

© 1992 Toshiba America Electronic Components, Inc.
Brand names and/or product names mentioned herein may be trademarks or registered trademarks of their respective companies.

MAS-91-065

CIRCLE NO. 26

Choices expand to incorporate design for testability

Barbara Tuck, Senior Editor

When looking for high-fault-coverage test programs, you have to decide if you're willing to support full scan, or if a back-end approach will be less intrusive.

Mentor Graphics (Wilsonville, OR) and CheckLogic Systems (Santa Clara, CA) have just entered into a product development and remarketing arrangement covering design-for-testability and back-end tools, both of which will begin shipping to Mentor's beta customers next month. A Mentor-developed tool for scan-based/combinational automatic test pattern generation (ATPG), to be used with Mentor's AutoLogic synthesis tools or standard netlists, will compete against test synthesis tools such as Test Compiler from Synopsys (Mountain View, CA) and SilcSyn II test synthesis integrated with the Intelligen ATPG engine by Racal-Redac (Mahwah, NJ), which has begun shipping to beta sites.

A sequential ATPG tool developed by CheckLogic will compete against back-end tools, such as the Test Design Expert (TDX) from ExperTest (Mountain View, CA) and the recently announced test generation tool from Sunrise Test Systems (Sunnyvale, CA). Under the terms of the agreement, CheckLogic will further develop and support Mentor's scan-based ATPG as well as its own sequential ATPG products.

Tied in with synthesis

Designers interested in full-scan methodology can use the scan-based ATPG tool with Mentor's AutoLogic, within which scan insertion can take place. Mentor claims that, for circuits that contain in excess of 100,000 gates, its scan-based ATPG tool can create high-quality test programs with fault coverage higher than 99.9 percent. Though full scan risks increasing silicon area, lengthening the design cycle and extending critical paths, Mentor's sales and marketing development manager for the simulation, synthesis and test division, Steve Eichenlaub, says that "A certain subset of AutoLogic users will justify full scan because of the quality results it yields."

AutoLogic users receiving beta copies of Mentor's scan-based ATPG will not be working with a fully integrated toolset, however. "The integration will be fair; a seamless handoff of information will take place," explains Eichenlaub. By the time production copies of the ATPG software are available, integration will have taken place at the database level. "In addition to a tie-in to schematic regeneration, we're developing heuristic guidance to ATPG within the AutoLogic environment so that the software can take advantage of the testability inherent to the circuitry," he says.

Mentor claims that today's scan tools are more amenable to design features that have in the past been considered incompatible with scan. For example, you're no longer required to strip portions of the design, such as the clock and scan-path circuitry, to create a circuit that is compatible with the scan tools. The clock and scan-path circuitry may now be as complex as desired, as long as memory elements are accessible by a predefined procedure. On top of that, those adopting full scan can lean on the high-impedance state of tristate drivers for fault detection.

Partial-scan option

Mentor users not wishing to deal with test issues up front in the design process can, once designs have been implemented, rely on CheckLogic's sequential ATPG tool for fault coverage. Sequential ATPG users will have the option of identifying critical flip-flops in the circuitry and creating partial-scan chains to improve fault coverage. Asked what sort of fault coverage sequential ATPG users will get, Eichenlaub says, "It depends on how much patience they have." Among factors affecting fault coverage are the sequential depth of the design and the amount of CPU time you want to devote to the task.

Though the potential market for partial scan is larger than that for full scan, Eichenlaub predicts Men-

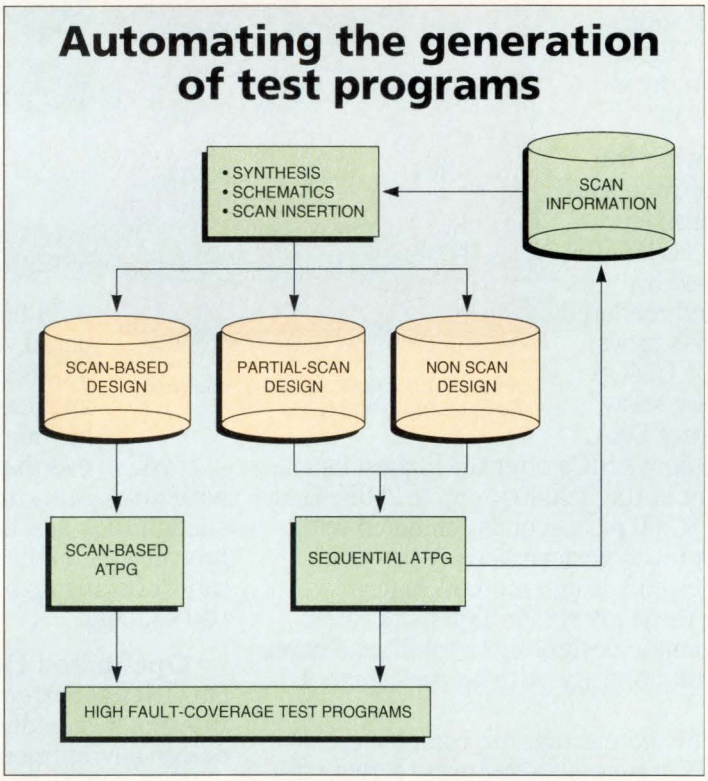

As ASIC complexity increases, system designers are increasingly having to choose between the risks associated with an up-front, design-for-testability strategy based on full scan, or of adopting an iterative back-end strategy using a sequential ATPG tool and perhaps partial scan, but running the risk of getting less-than-optimum fault coverage.

THE ONLY 5-VOLT MEGABIT FLASH. NO WAITING.

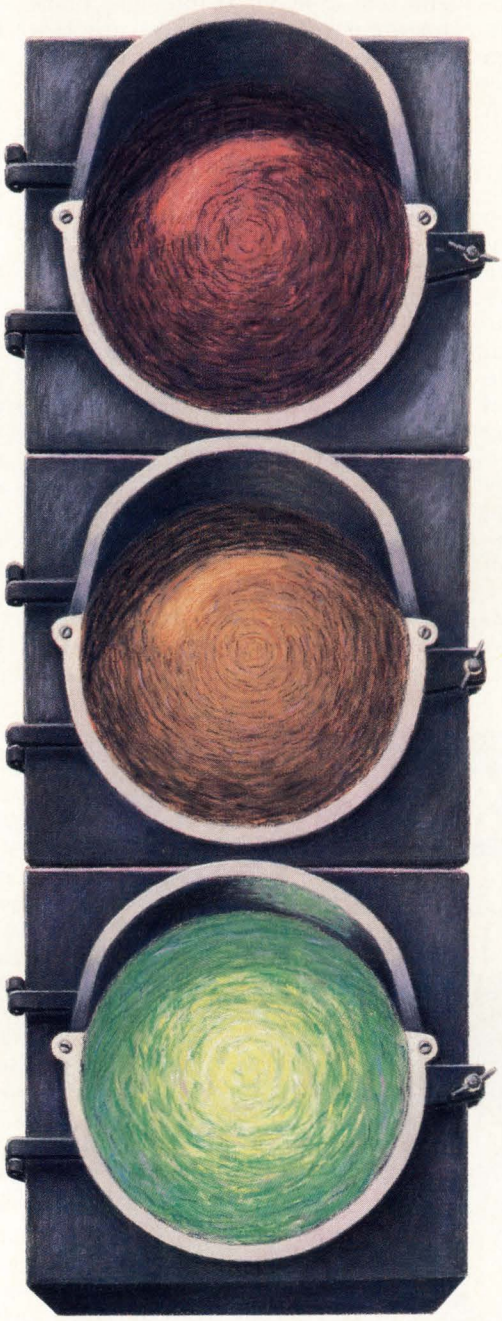

Atmel's 5-Volt-only, one-megabit Flash is available. Production quantities. No waiting.

That's 5 Volts to read and 5 Volts to write. No costly 12-Volt converter or regulator circuits to clutter your system. And, there's more:

PERFORMANCE:

- The AT29C010 is fast—90 nanoseconds commercial, 120 nanoseconds military.

HASSLE FREE:

- The AT29C010 has one thousand 128-byte sectors, and each can be individually written or erased.
- The AT29C010 has the easiest programming algorithm in town. It's self timed and has automatic erase, so you don't have to erase before writing.

So, if you like blazing access times, need sector write and want to save space, call, fax or write us about your application and we'll get you a sample of the Atmel 5-Volt-only Flash.

No Waiting.

INSTANT DATA ACCESS (IDA) DIAL (617) 494-8338 DOCUMENT NO. 1045

ATMEL®

The people who make the difference.

ATMEL CORPORATION
2125 O'Nel Drive
San Jose, CA 95131

Tel. 1-800-292-8635
Tel. (408) 441-0311
FAX (408) 436-4200

CIRCLE NO. 27

ASICs & ASIC DESIGN TOOLS

tor users will be about equally interested in the sequential and scan-based ATPG tools. "Market segmentation will depend on types of design methodologies, mindsets and functionality-versus-cost decisions. A different balance will apply to different users."

Racal-Redac's SilcSyn II/Intelligen test synthesis product, which has lagged behind the synthesis capability of SilcSyn II by a year, offers both partial- and full-scan methodologies. Racal-Redac has taken the guts of its Intelligen sequential ATPG tool and modified the algorithms to work with combinational as well as sequential logic, according to Jeff Fox, vice-president of engineering for ASIC tools. The company claims that SilcSyn II/Intelligen, driven by both netlists and behavioral descriptions, puts control of sequential depth into your hands. You have the option of using multiple or gated clocks, and can also insert JTAG boundary scan. A front-end tool, SilcSyn II/Intelligen adds testability elements before any circuit optimization takes place.

Back-end tools

The TestGen partial-scan sequential ATPG tool from Sunrise Test Systems, on the other hand, is used after a preliminary chip design is completed, avoiding interference with the design process itself. TestGen supports combinational and sequential logic, synchronous and asynchronous circuits and ASICs and full-custom ICs of both low and high complexity, including those with embedded functions. After a fault simulator analyzes the circuit design, you can take advantage of test synthesis to selectively insert scan points to make circuits more controllable and observable. You can exclude performance-critical logic blocks, paths or gates from scan insertion, to avoid any speed penalty.

In contrast to some ATPG products that force you to write a special high-level description of the chip design, TestGen's vector generator works from a standard design netlist to produce test programs automatically. Up until now, users of the Test Design Expert from ExperTest have needed a behavioral description of their chips at the register transfer level to test their sequential designs. Though ExperTest will announce this month that a netlist-driven automatic test equipment (ATE) solution has been bundled into TDX, Kevin Hall, vice-president of marketing, says the company is still recommending that designers use a behavioral model for best results. "Because a behavioral model makes it easier to navigate through the chip," he explains, "fault coverage is higher and vector counts lower."

Moreover, with a behavioral model you don't have to change the test vectors when you change your structure.

HEAT SINKS for INTEL 80486 and i860XR MICROPROCESSORS

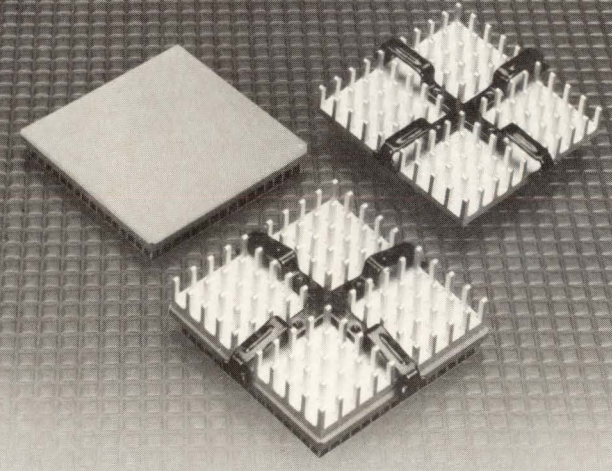

EG&G WAKEFIELD ENGINEERING
60 AUDUBON ROAD, WAKEFIELD, MA 01880
TEL: (617) 245-5900 • FAX: (617) 246-0874

Designed for use with the Intel 80486DX and 80486SX microprocessors and the Intel i860XR RISC processor, packaged in 168-pin ceramic PGAs, the EG&G Wakefield 669 Series Heat Sink/Clip Assembly offers a cost-effective heat dissipation solution for today's high-speed microprocessors. This assembly provides the highest clamping force available with a nylon-coated stainless steel clip, for the most efficient interface heat transfer and to meet system shock and drop test requirements.

Our omnidirectional heat sink offers optimized heat dissipation and ease of application; the symmetrical clip is suitable for high volume installation with the EG&G Wakefield 162-IT installation tool.

Heat dissipation with the 669 Series Heat Sink/Clip Assembly is optimized for PC, workstation, and server applications with low airflows (e.g., 50-200LFM). Pressure drop is minimized in multiple-processor applications.

Intel, 80486, i860XR, and 80386 are trademarks of Intel Corporation.

Call our Application Engineering Department today at (617) 245-5900 for information about the EG&G Wakefield Engineering 669 Series (patent pending) and other heat sinks for the Intel 80386 and 80486 microprocessor family. Also ask for information about the EG&G Wakefield DELTEM™ family of heat sinks for cooling high-speed cache RAM used with 50 MHz microprocessors.

CIRCLE NO. 28

So Much Real-Time Power. Such A Small Package.

Pick up these ROMs and you're in touch with the world's most powerful real-time system software solution—OS-9®. But OS-9 is so much more than just a powerful real-time kernel. Check out this *Critical Features Checklist* to find out how OS-9 can give you "Real Time In No Time."

CRITICAL FEATURES CHECKLIST

Real-Time Operating System Features	Robust Development Environment Features	Sophisticated I/O Features
Compact (28K), high-performance real-time kernel for demanding applications Uses UNIX process and I/O models Multi-user, multi-tasking, pre-emptive scheduler Modular architecture User-installable system calls Interprocess communication facilities: • semaphores • pipes • signals • events	UNIX-hosted development (UniBridge) PC-DOS-hosted development (PCBridge) Complete development capabilities: • highly-optimizing ANSI C compiler, assembler/linker • C source level and system level debuggers • PVCS source code control system • advanced shell interface (MShell)	Hard and flexible disk support, SCSI Common Command Set Tape support WORM support Networking: • Ethernet (IEEE 802.3) • NFS Version 2 • ARCNET (SMC Data Point) Graphics: • X Window System V11R4 • OSF/Motif Version 1.1.1 • RAVE for real-time graphics and multimedia

Also, ask us about OS-9000® for 386/486 and RISC processors.

Grab the Power

So get your hands on these ROMs that have been proven in thousands of specialized and demanding applications. Call Microware® today to put OS-9 to work for you.

1-800-475-9000

Call Microware Today!

In California,
call (408) 980-0201

microware®

MICROWARE SYSTEMS CORPORATION

1900 N.W. 114th Street
Des Moines, Iowa 50325-7077

MORE CHOICES • MORE OPTIONS • TOTAL SUPPORT

CIRCLE NO. 29

Run PC Software on a Hardhat Computer

NEW!
XVME-688 80386SX CPU
$1350 in OEM Quantities*
*100 Units

With Xycom's large family of PC/AT products for the VMEbus, you can run the most popular PC software packages, as well as take advantage of the multiprocessing capability and increased reliability designed into the VMEbus architecture. Complete your system with a full line of Xycom hardhat hardened VMEbus 68000 family CPUs, I/O, peripheral controllers and support modules.

Xycom offers one-stop shopping for all the VMEbus hardware and software you'll need in any automation, control or data acquisition application.

FREE

Xycom's complete family of PC/AT hardware includes: 80386SX™, 80486SX™ and 80486™ PC/AT processors. Call today for a detailed catalog of all Xycom VMEbus systems.

1-800-AT XYCOM

xycom

THE HARDHAT COMPUTERS

750 North Maple Road, Saline, Michigan 48176
(313) 429-4971 FAX:(313) 429-1010
In USA call: 1-800-289-9266 (1-800-AT-XYCOM)
In Canada call: (416) 825-0281
In Europe call: +44 604 790 767
In Asia call: (65) 251-0161

Genesis is a registered trademark of Iconics Corporation. The Norton Utilities is a registered trademark of Symantec Corporation. Microsoft Windows and Microsoft Excel are registered trademarks of Microsoft Corporation. Harvard Graphics® is a registered trademark of Software Publishing Corporation. dBASE and dBASE IV are registered trademarks of Ashton-Tate Corporation. Lotus and 1-2-3 are registered trademarks of Lotus Development Corporation. WordPerfect® is a registered trademark of WordPerfect Corporation. PC Tools is a registered trademark of Central Point Software, Inc. 386 and 486 are trademarks of Intel Corporation.

CIRCLE NO. 30

"You can just change the model, resynthesize the design and resynthesize the test vectors," Hall reports.

Synthesis is key to ASIC testability. He adds, "Our customer base is committed to top-down design and synthesis." Because TDX is compatible with synthesis tools from Synopsys, Viewlogic (Marlboro, MA) and Cadence Design Systems (San Jose, CA), those of you who design with these tools don't need special models for TDX. You can use the same models to drive test vectors that you use for synthesis.

What about floor planning?

In addition to selecting testability techniques, you'll also have to decide whether to use third-party tools or silicon vendors' test tools which have access to floor planning and layout information. Farzad Zarrinfar, marketing manager for LSI Logic (Milpitas, CA), says that an effective test synthesis tool should use layout and floor planning information. "This is the only realistic way to optimize the design for area and speed with the most compact test vectors," claims Zarrinfar.

"Unfortunately, due to lack of floor planning and routing information, third-party test synthesis tools synthesize the scan logic using just the circuit netlist. This approach can result in a nonoptimized ASIC," he adds. LSI Logic's year-old Test Builder test synthesis tool uses floor planning information to determine whether the clock skew between two scan elements will cause an error in clocking data and for automatic insertion of lock-up latch.

Zarrinfar states that LSI Logic has over 50 scan elements in its library, giving you lots of freedom to design. "Third-party tools aren't the most efficient way to go, since they support only a subset of scan elements," he says. But Mentor's Eichenlaub says that if silicon vendors such as LSI Logic choose to release only a subset of scan cells to third-party vendors, "It'll stick in the craw of users who want to base foundry choices on considerations of turnaround time and price, not functionality in the library." As users gain experience with ASIC design, more and more of them are looking for foundry independence. ■

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

CheckLogic Systems		Racal-Redac	
(408) 988-3846	Circle 260	(508) 692-4000	Circle 264
ExperTest		Sunrise Test Systems	
(415) 965-2000	Circle 261	(408) 739-4000	Circle 265
LSI Logic		Synopsys	
(408) 433-8000	Circle 262	(415) 962-5000	Circle 266
Mentor Graphics			
(503) 685-7000	Circle 263		

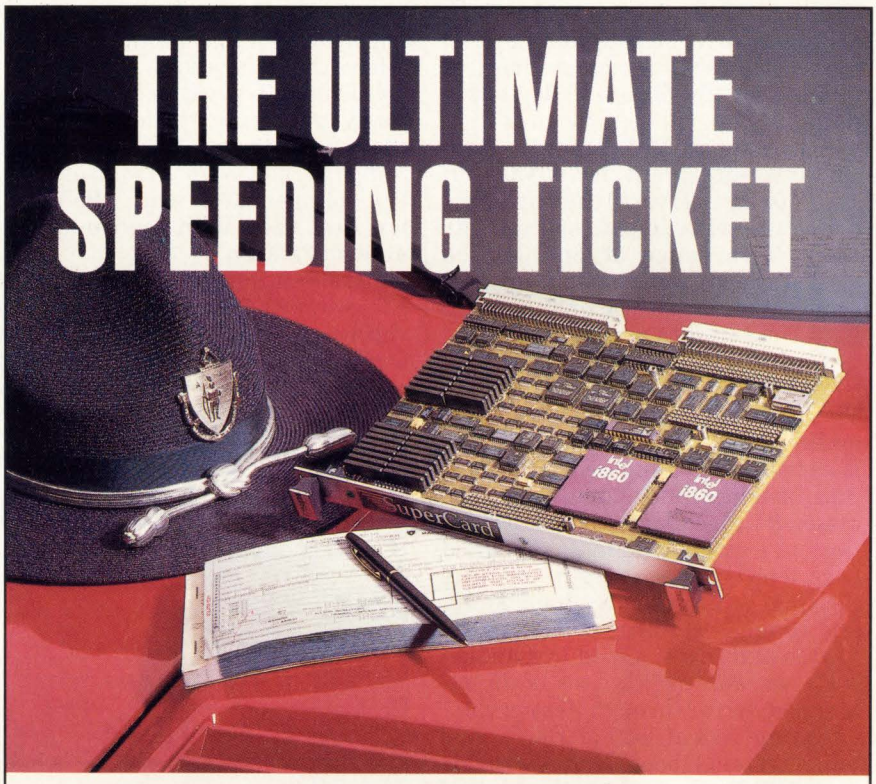

SuperKit: When Development Time is as Critical as Compute Time

SuperCard i860-based vector processors are superfast, more than 2.5 GigaFLOPS of raw multiprocessor power, but then you'd expect that. The real news is SuperKit, a robust software toolkit which cuts development time to the bone and gets you up and running superfast.

Fortran, C and ADA compilers... the industry's largest scientific subroutine libraries... diagnostics aplenty... a pSOS+ real-time kernel running in a UNIX-like environment... CASE tools... and much more.

And talk about versatility... from 6U VME boards to fully packaged multiprocessor systems to MIL spec units... large memories... peripheral I/O mezzanine boards... SuperCard has it all.

Find out how you can benefit from SuperCard and SuperKit solutions just right for you. Call or write CSPI, 40 Linnell Circle, Billerica, MA 01821. **1-800-325-3110**. Or 617-272-6020. Fax: 508-663-0150.

SUPERCARD CSPI

CIRCLE NO. 31

COMPUTERS & SUBSYSTEMS

Digital's Sun killer also rises

Dave Wilson, Senior Editor

Representatives of Sun Microsystems (Mountain View, CA) have derided the new 64-bit Alpha processor architecture from Digital Equipment Corporation (Hudson, MA) as "just perfect for a minicomputer designer," but it seems that DEC has some surprises in store for the Silicon Valley-based upstart.

According to Bill Demmer, Digital's vice-president for VAX/VMS systems and servers, the Alpha architecture will compete successfully with Sun's range of Sparc-based workstations, not to mention personal computers based on the 80X86 from Intel (Santa Clara, CA).

Make way for the 21064

The first incarnation of the Alpha architecture, the 64-bit 21064 itself, is a classic RISC superscalar, superpipelined machine. A dual-issue processor, it's able to launch two instructions in a given cycle to its pipelined functional units: integer operate, floating-point operate, load store, and branch. The chip comes with 8 kbytes of direct-mapped instruction cache as well as 8 kbytes of direct-mapped, write-through data cache.

Taking a leaf out of Sun's own "open systems" book, Digital intends to sell the processor at chip, board and system levels. Furthermore, in one for the books, Digital has said that it will license its operating systems, including DEC OSF/1 and VMS, its compilers and its layered software products.

The 21064 is far in advance of its competitors, perhaps by as much as a year, although it isn't the world's first 64-bit microprocessor. That honor goes to the MIPS Computer Systems (Sunnyvale, CA) R4000 design, announced a few months ago. Nevertheless, Digital's entry, built using a 0.75 micron VLSI technology at Digital's facilities in Hudson, MA and South Queensferry, Scotland, is the world's fastest 64-bit chip.

It's easy to see why a Digital system running applications on Windows NT from Microsoft (Redmond, WA) might gain the upper hand over an Intel-based design in the PC market. Even the forthcoming 100-Mips 586 (P5) and R4000 processors are going to have a hard time competing with the 300-Mips Digital component. In addition to its speed advantage, Digital is at least six months ahead of its competition in supplying a low-voltage processor. While others are just talking about manufacturing 3.3-V parts, Digital is already producing them—the 21064 operates from a 3.3-V supply. The company has been producing such low-voltage parts since 1990, and has used them with previous generations of silicon.

Needless to say, it's a significant achievement to reduce a processor's voltage and keep its speed up at the same time. One of the things that makes Digital's process extremely fast is the very low-resistance, low-capacitance third-layer metal used to build the chips. Usually only gate array vendors boast of triple-level metal, and there it's used to maximize the routability of the gate array. In contrast, Digital's third-level metal is intended solely for power supply and clock distribution, as opposed to signal interconnect. "You can't get to 200 MHz without the third-layer metallization," says Dirk Meyer, principal engineer with the semiconductor engineering group and one of the co-architects of the microprocessor.

Digital's circuit designers worked extremely hard to wring the last little bit of performance out of the process technology, too. They built very refined circuits using a lot of simulation—synthesis was used only where it didn't get in the way of cycle time. "In contrast to another design style where they use synthesis everywhere, we use it at very localized spots and absolutely don't

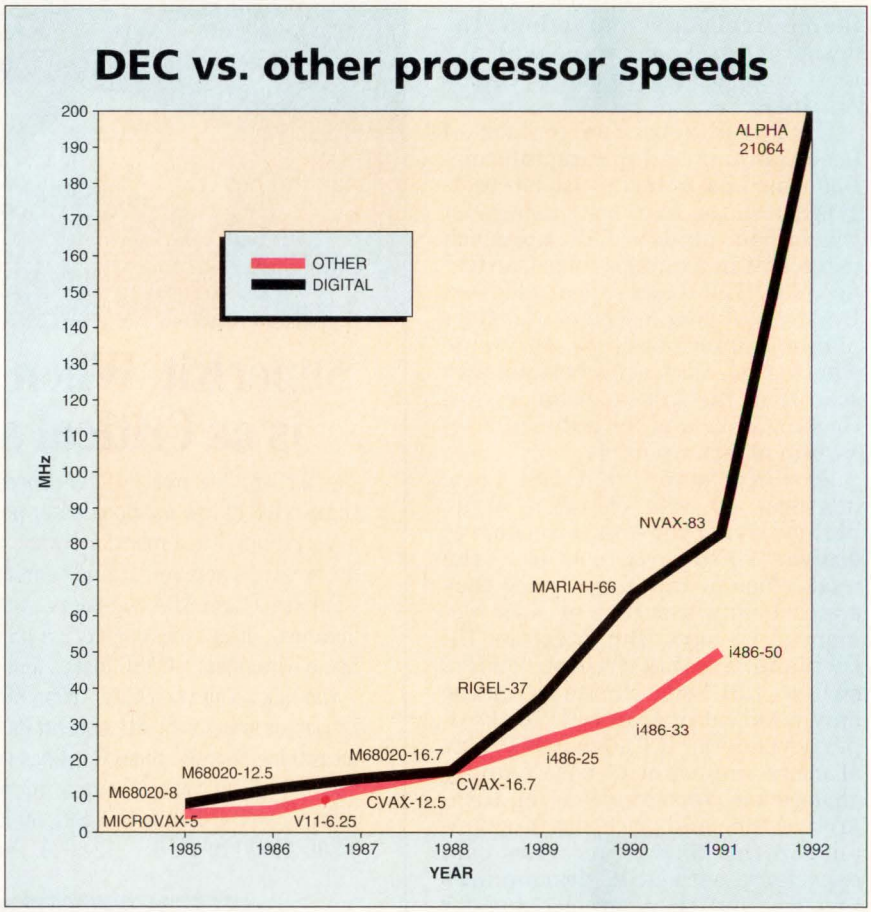

Digital's processors have outrun the industry's since 1988, but the recently announced Alpha is something else again—offering a fourfold speed advantage over the fastest competitor.

INTRODUCING THE TP-IGC6V
TADPOLE'S DEDICATED X WINDOWS SERVER FOR VME

What To Look For In High-Performance X

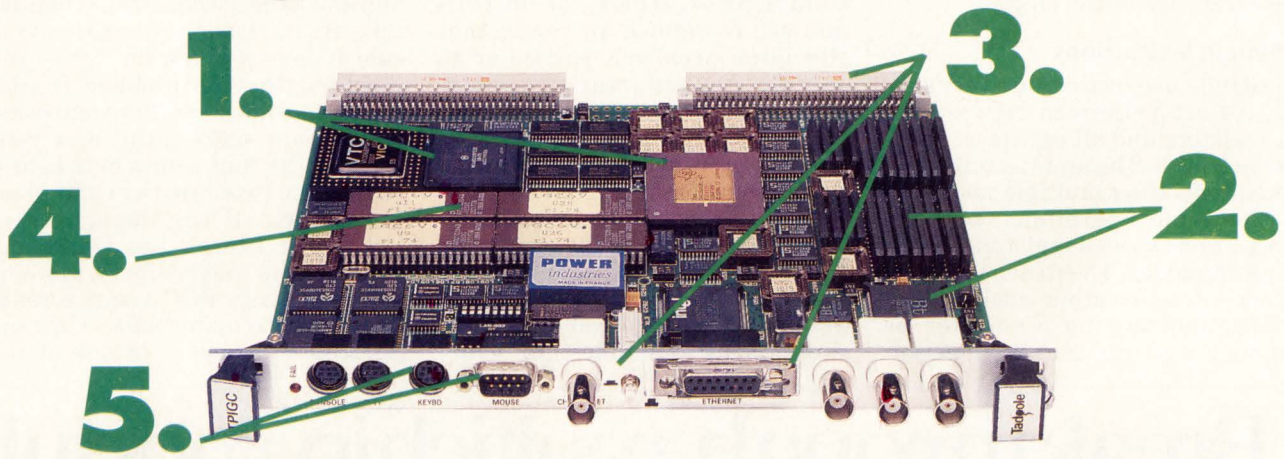

Looking for maximum performance in VME-based X Windows applications?

Your search is over.

Tadpole Technology's new TP-IGC6V Intelligent Graphics Controller gives you more of everything for powerful VMEbus or stand-alone X Windows performance.

1. More Processing Power

While most X servers rely on a single processor, the IGC uses two. A 40 MHz TMS34020 handles the X server code and graphics, while a 20 MHz MC68020 takes care of the TCP network communications and user interface. By dividing computing tasks and off-loading X Window server routines from the host, our dual-processor design significantly improves overall system performance.

2. High-Resolution Performance

When your application demands high-speed, high-resolution color graphics, the IGC delivers. It runs 37,000 X Stones with 1280 x 1024 non-interlaced resolution — while displaying 256 colors from a palette of 16.7 million.

3. More Configuration Flexibility

No other VME-based X server offers you the configuration flexibility of the IGC. It can communicate with X clients over the VME backplane, through Ethernet™ or both — making it ideal for a wide range of VMEbus X applications. You can even install several IGCs in the same VME system.

The IGC is compatible with Tadpole's extensive line of high-performance RISC and CISC single-board computers running VxWorks™ or TP-IX, Tadpole's implementation of UNIX® System V.3.2.

4. More On-Board Software

On-board EPROM contains the complete X11R4 server code, which is optimized for the IGC's dual processors. The IGC also offers the convenience of built-in VT100/200 terminal emulation, diagnostics and a user-friendly firmware monitor.

5. Just Plug And Play

The IGC comes complete with a mouse and keyboard interface. Just add power. We've taken care of everything else.

So what should you look for in high-performance X? You're looking at it: Tadpole's new TP-IGC6V.

For information in the U.S. call:

800-232-6656

Cambridge, England
0223 423030

Paris, France
(1) 60 86 27 92

All trademarks belong to their respective companies.

Computing Without Compromise

T A D P O L E

CIRCLE NO. 32

COMPUTERS & SUBSYSTEMS

use it if it's going to get in our way," says Ed McLellan, the microprocessor's other co-architect.

When a designer defines the instruction set of a microprocessor based on any architecture, he or she tries to ensure that cycle time limitations aren't imposed on the design. The Alpha architecture was designed to prevent circuit implementers from building hardware structures that would get in the way of the cycle time of the chip.

Simple instructions

The Alpha instructions are very simple. All are performed between 64-bit registers and all registers are 64 bits in length. The architecture facilitates pipelining multiple instances of the same operations, because there are no special registers and no condition codes. The instructions interact with each other solely by one instruction writing a register or memory and another instruction reading from the same place. That makes it easy for Digital to build implementations that issue multiple instructions every CPU cycle.

Alpha also makes it easy to maintain binary compatibility and full speed across multiple implementations. For example, there are no implementation-specific pipeline timing hazards, no load delay slots and no branch delay slots. This is in contrast to other architectures, where the instruction set is defined in a way that alleviates problems particular to specific implementation generations. Delayed branch mechanisms, for example, have proved popular on other RISC machines, popping up on the Hewlett-Packard PA, MIPS and Sparc architectures. All three use the concept of delayed branches or branch delay slots. "That's a concept that works well in reasonably well-pipelined, single-issue machines. In many cases it can bury the branch latency. The problem comes when you go to deeper pipelines or wider issue machines. Now the branch latency is much bigger and delayed branching is no longer appropriate," McLellan concludes.

Another important differentiating characteristic of the Alpha is the division of hardware and software support for byte- and word-length operations. Says Digital's Meyer, "Every other RISC architecture in its load/store instruction subset has support for byte- and word-length. The problem is that you need some circuit structure between your data cache and your register file that—for a byte load—swings the byte into the right place and zeroes off the upper half of the word. Stores are worse because cache and memory systems which support ECC require read-modify-write operations. That necessitates hardware in one of the

Breakthrough multichip modules

critical paths of every chip that I'm aware of."

With the Alpha architecture, this problem has been overcome by only providing loads and stores of longword and quadword—32-bit and 64-bit memory operands. Byte and word operations are still supported in hardware but use register file instead of memory operands. This simplifies and removes bottlenecks from cache and memory system designs.

Changes in system design

While it's one thing to design a processor, it's quite another to design a system to go with it. Even designers accustomed to building 50-MHz subsystems around the 486, with its 12-ns SRAMs, might be a tad rattled dealing with a RISC processor roaring along at 150 MHz. "We wanted to put this chip in desktop workstations as well as full-blown multiprocessor servers," says Meyer. "To do that we didn't want to design a memory bus interface that was designed for one of those applications at the expense of the other."

The solution Digital designers came up with lets you run the external cache memory subsystem and the other system logic at some frequency less than the clock frequency of the processor.

If, for example, you were to build a system based around a 200-MHz processor, you might choose a fast 100-MHz second-level cache subsystem interfaced to a somewhat slower 40-MHz PC peripheral chip set. The subsystem is configured by software upon initialization. Although the first chip offering includes no support devices for second-level cache management, it's likely that Digital will support its processor with such devices in the future. The processor itself doesn't force a particular cache coherency methodology—that's left to the designer.

Perhaps the best thing about Digital's chip is that you don't have to wait to use it. Packaged in a 431-pin PGA and dissipating around 23 W, evaluation quantities of the device are available now. The chip might be priced at a hefty $3,375 in 1- to 100-unit quantities, but the price drops rapidly to $1,600 for over-1,000-unit quantities. ■

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 Equipment Corporation
(800) 332-2717 ..... Circle 281
Hewlett-Packard
(408) 725-8900 ..... Circle 282
Intel
(800) 548-4725 ..... Circle 283
Microsoft
(206) 882-8080 ..... Circle 284
MIPS Computer Systems
(408) 720-1700 ..... Circle 285
Sun Microsystems
(415) 960-1300 ..... Circle 286

for breakneck speeds.

That's AT&T "Customerizing."

If beyond 50MHz performance is where you're heading in workstations, AT&T's multilayer Multichip Module (MCM) solution is the most reliable, viable way to go.

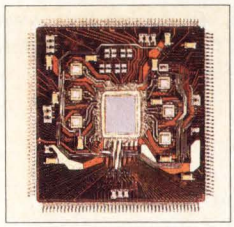

POLYHIC packaging, developed by AT&T Bell Laboratories, combines copper thin film with our patented polymer. It delivers one of the industry's lowest dielectric constants (2.8) to minimize access times between CPU and memory devices. And low-loss POLYHIC technology helps prevent timing problems on critical paths.

Every AT&T MCM is comprehensively tested to meet critical function requirements. Designs work right the first time because Bell Labs' design expertise assures that your crucial high performance requirements are met. And POLYHIC MCM modularity helps you get system designs up and running faster.

Should the going get rough, AT&T engineers are on stand-by to advise you. Analyze your circuits. Even recommend a solution to take you to the next generation of high-end workstation performance. That's "Customerizing."

For more about the POLYHIC power of AT&T's MCMs, call AT&T Microelectronics at 1 800 372-2447, ext. 900.

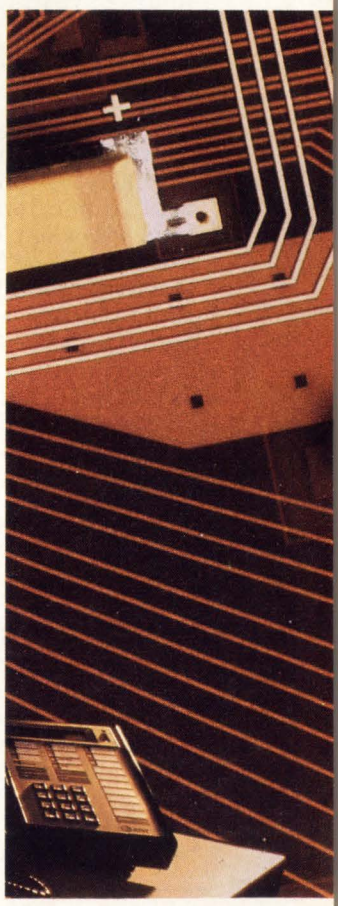

AT&T
Microelectronics

CIRCLE NO. 33

COMPUTERS & SUBSYSTEMS

Embedded PC bus standard gains ground

Warren Andrews, Senior Editor

Although the ISA bus probably has the largest following of any standard bus, its success in embedded applications and the industrial environment has been less than scintillating. Now, a new PC bus, PC/104, will cater to just such embedded applications. Its key features are a very small form factor and a "stack-pack" connector architecture. Already, a cadre of more than 20 manufacturers has adopted the fledgling specification. They've even formed a consortium and have started to march the specification through the mine field of IEEE certification as an extension to the IEEE-P996 PC bus standard.

Traditionally, bus architectures have adopted PC bus protocols and software standards and implemented them on other pinouts. PC-DOS-compatible systems have emerged on STD Bus, G-64/96, Multibus I and II, and even VME and VXI, to mention a few—and it's likely to continue in that direction. Take Intel, for example. "We're focusing all our Multibus II activity toward our PC-compatible environment," says Richard Binns, Multibus marketing manager. That's evidenced in the company's latest Multibus II release and the direction it's taking with iRMX for Windows.

The PC/104 Consortium, the group charged with specifying the new bus, sees the highly simplified approach taken in its very small form factor—only 3.6 × 3.8 in.—as being most desirable, providing a low-cost, low-power solution. "And should a designer or OEM want to include some PC-DOS component on another bus such as VME, Multibus or even Futurebus+, the PC/104 bus can easily be used as a mezzanine or daughterboard," says Rick Lehrbaum, vice-president of strategic development at Ampro Computers (Sunnyvale, CA) and chairman of the consortium.

At least one VME vendor, Xycom (Saline, MI), uses a PC/104 interface for a VME card. "The PC/104 may soon become the first fully standardized VME mezzanine bus," says Lehrbaum. But while he sees the mezzanine-bus approach as one use of PC/104, he believes the small form factor, low power dissipation, stack-through connector technology, and low cost will make the bus suitable for a far greater number of embedded applications.

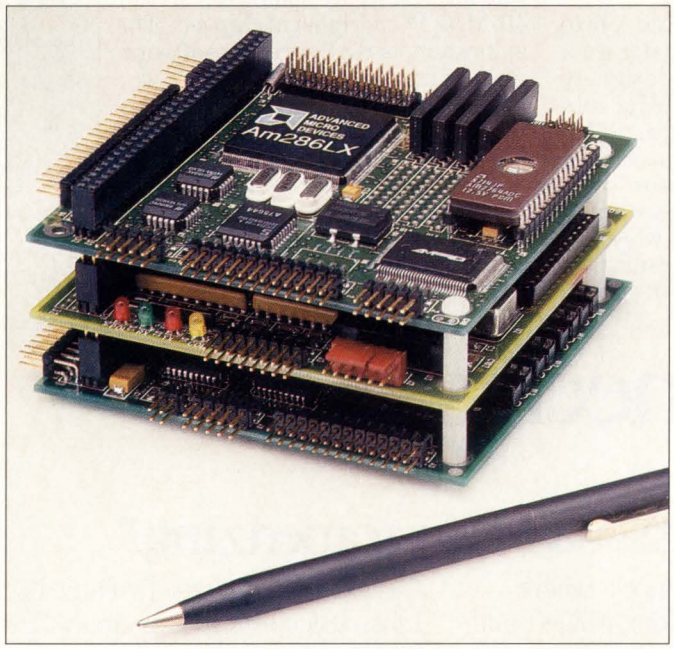

PC/104 is an emerging bus derived from the basic ISA architecture. It features a compact form factor (3.6 × 3.8 in.), low power dissipation, low cost, and a unique stack-through connector design. Pictured is a stack representing a full-featured PC.

Several options available

Though the standard has well-defined mechanical and electrical definitions, there are a handful of options available to designers. First, you can utilize a standard 8-bit PC/XT version which uses a single 64-pin connector (P1). This connector can be set up mechanically to be either the stack-through or non-stack-through type. The non-stack-through connection is used in applications where the board mates with another PC/104 board or a motherboard, and where vertical height must be kept to a minimum. The connectors exit the circuit board at the edge and are parallel to its plane.

In applications where the full 16 bits of the PC/AT bus are required, a 40-pin P2 connector is appended to the board next to the P1 connector. Presently, only a handful of fully developed PC/AT modules are available, but it's expected that as the technology catches on, more and more OEMs will be demanding the transfer capability of the wider bus, in addition to the attendant power of the faster 286, 386 and 486 processors. Ampro already offers a 286-based module, while other manufacturers in the consortium provide mating modules with such capabilities as high-resolution graphics.

In addition, vendors subscribing to the PC/104 approach already provide a broad variety of add-on boards. IOtech (Cleveland, OH), for example, offers a 488.2 interface module. This gives OEMs the capability to quickly and almost painlessly add IEEE-488.2 capability to an embedded application. Other examples include a relay driver board from DDI (Beverly, MA), a display controller from Reflection Technology (Waltham, MA), analog data acquisition boards from Automation Instruments (Carrollton, TX) and Diamond Systems (Palo Alto, CA), and a variety of analog I/O and digital control products from Real Time Devices (State College, PA).

Ampro also offers a processor core module it calls CoreModule in both XT (Hitachi V20) and 286 versions. The company has more than a dozen other so-called MiniModules, including Super VGA, VGA, CGA, and

NEW & NOW

ICEBOX™ VMEbus System Enclosure

The Case For The Best Case.

Would you believe close to half the spurious board level problems in your system are enclosure-related? Odds are, they are. The new Ironics ICEBOX roots out tricky system problems — ***before they start.***

ICEBOX answers cooling demands that leave muffin fans gasping for air, because our 550 CFM Blower/Diffuser/Plenum (patents pending) ensure vigorous, even cooling no matter how dense your slot loading — ***with no "dead" spots.***

ICEBOX's 16 gauge steel cabinet, swing-aside power supply (patents pending), and uncluttered interior layout give you a tough, accessible case with plenty of room to mount ***all the I/O your system requires.***

ICEBOX's premium electronics produce vast power reserves, while limiting electrical noise. Each component is selected to reduce headaches and improve reliability, because we know ***you'd rather tune your application — not the box you've put it in.***

Call our VME Product Hotline today — **800-334-4812.**

IRONICS Incorporated

IRONICS Incorporated
798 Cascadilla Street
Ithaca, New York 14850 USA

607.277.4060
Fax 607.272.5787
Telex 705742

INSTANT DATA ACCESS (IDA) DIAL (617) 494-8338 DOCUMENT NO. 1046

ICEBOX is a trademark of Ironics Incorporated.

CIRCLE NO. 35

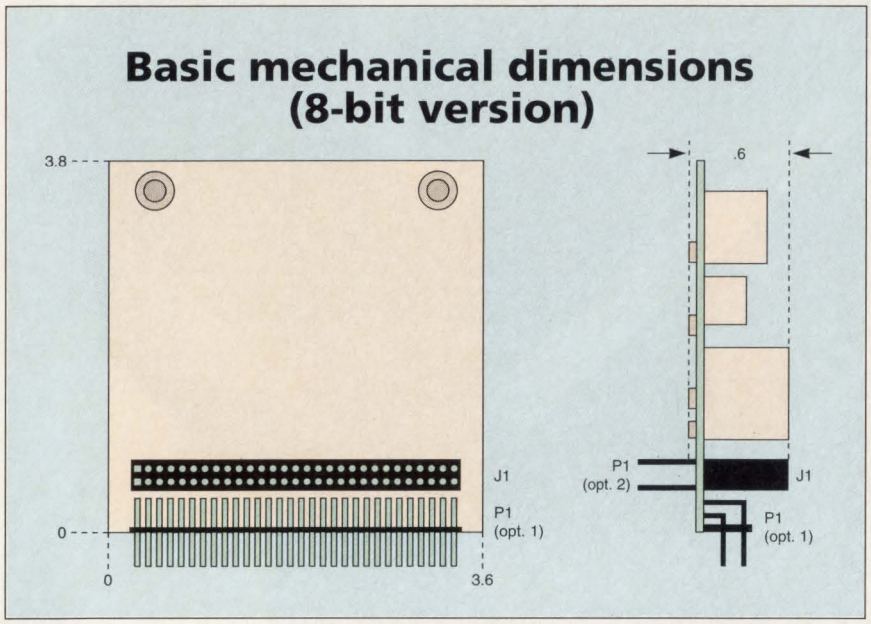

The PC/104 provides several options in its connector approach. The connector can offer the new stack-through arrangement or, for applications requiring a minimum of space, the connections can come off the edge of the card, parallel to the plane of the circuit board. The specification also provides for a P2 connector (not shown), which resides next to the P1 connector when full 16-bit PC/AT capability is required.

LCD display cards, Ethernet and Arcnet networking cards, modem cards, and FAX cards. This is not to mention Ampro's variety of serial, parallel, hard-, and floppy-disk drive cards, including a SCSI interface.

Changes may be unwelcome

With vendors already lined up behind PC/104, the application to IEEE for an addition to the ISA standard is more a sign of solidarity than a requirement for growth. In addition, the large number of vendors supporting the current version of PC/104 and the quantity of products on the market would make any change in the specification very unwelcome, so it's likely to be mandated by the IEEE specification process unchanged—especially if proposed changes were to require substantial alteration of the mechanical configuration or bus protocols.

The need for a small form-factor PC-bus alternative—particularly one that can stand up to industrial environments—has been well documented over the past several years. Traditional PC-bus approaches have been woefully inadequate in such hostile environments for a variety of reasons. These range from poorly specified components to heat build-up to too-long circuit traces picking up radiated energy in environments with electrical noise. The PC/104 should go a long way toward mitigating these problems. ■

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

Ampro Computers
(408) 522-2100 ..... Circle 380
Automation Instruments
(214) 418-6827 ..... Circle 381
DDI
(508) 927-7976 ..... Circle 382
Diamond Systems
(415) 813-1100 ..... Circle 383
Intel
(800) 548-4725 ..... Circle 384
IOtech
(216) 439-4091 ..... Circle 385
PC/104 Consortium
(408) 245-9348 ..... Circle 386
Real Time Devices
(814) 234-8087 ..... Circle 387
Reflection Technology
(617) 890-5905 ..... Circle 388
Xycom
(313) 429-4971 ..... Circle 389

VMEbus 486DX -33MHZ

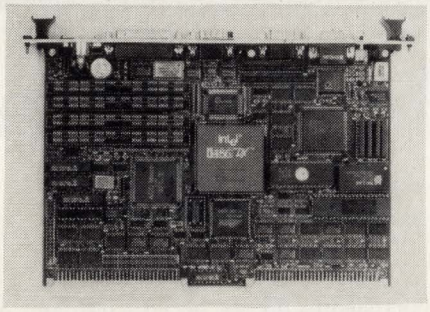

- New generation **single slot** 486DX -25 & 33MHZ -100% PC/AT compatible (A32/D32)
- No platforms, no piggy-backs, just clean streamlined **low-cost** one-side SMD technology
- Up to 16M of on-board DRAM
- 1024x768 SuperVGA w/512K Video DRAM
- 2 Serial Ports , 1 Parallel Port and Real Time Clock w/Battery
- AT style keyboard Controller and Connector
- Floppy and Hard Drive IDE Controllers on-board -4 ways to connect to Disk Drives
- Three-way Byte/Word Swapping Hardware for interfacing to 68000 type CPUs
  Access VMEbus through Real Mode Window or Protected Mode
- Custom Gate Array technology provides
  -Interrupter, Interrupt Handler, System Controller, all SW configured
- Reset Switch and Speaker
- Choice of **SCSI** or PC/AT Bus on P2

INTERGON INC.
3000 Wilcrest
Houston, Texas 77042
1-800 - 68VME-AT

Tel: (713) 780-7003 Fax: (713) 974-6059

CIRCLE NO. 36

AWESOME

TEK-AT3
386DX HALF-SIZE PC/AT

We think this is the best way to describe the new TEK-AT3 all CMOS 386DX single board computer. And rightly so. Because the TEK-AT3 comes standard with a powerful 33 Mhz CPU, high performance Cache memory, and a capacity to support up to 16 Megabytes of DRAM.

Plus, you get all the bells and whistles you've come to expect from conventional PCs: hard disk interface, floppy disk controller, serial and parallel ports.

And like all Teknor computer cards, the TEK-AT3 benefits from the latest innovative technologies. Options like 1 Meg of onboard programmable Flash EPROM, 1 Meg of battery-backed SRAM, and the ability to boot from Flash EPROM disk offer possibilities never before conceived on a 7 x 4.7" PC/AT form factor.

In fact, when you consider all the advantages the TEK-AT3 will bring to your system designs... you'll think it's awesome, too!

Call us today at **1-800-387-4222** for your free literature package.

QUALITY 2 YEAR WARRANTY

TEKNOR
MICROSYSTEMS INC.

P.O. Box 455, Sainte-Therese (Quebec), Canada J7E 4J8 • Tel: (514) 437-5682 • Fax: (514) 437-8053

INSTANT DATA ACCESS (IDA) DIAL (617) 494-8338 DOCUMENT NO. 1028

CIRCLE NO. 37

Now you can afford to

Presenting a very small development in Ethernet.*

Chipsets that are matched to your system and your budget. In fact, they cost you as little as 5 square inches. Which, by the way, is less total real estate than any competitive solution. But sizable reductions don't stop with board space, because we're also reducing the price up to 30 percent.

Needless to say, true plug-and-play simplicity requires an intelligent network interface. So our new high-integration 82503 Dual Serial Transceiver goes beyond IEEE 802.3 to include automatic port selection, polarity switching and a jumperless interface to AUI or TPE.

For unmatched desktop performance, we offer

© 1992 Intel Corporation. Intel386 and Intel486 are trademarks of Intel Corporation. *Ethernet is a registered trademark of Xerox and Diamond is a registered trademark of Diamond Brands, Inc.

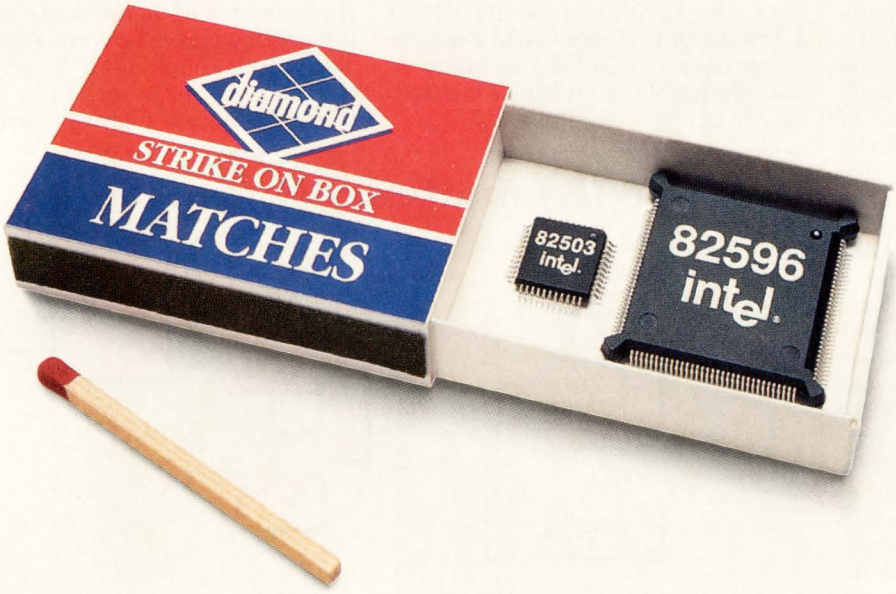

put Ethernet in any box.

a complete family of 82596 LAN coprocessors, each optimized to a specific Intel486™ CPU for maximum throughput. And our 82593 is the perfect LAN controller for Intel386™ SL notebooks.

Best of all, these true two-chip solutions give you the flexibility to simplify your design and deliver your product to market in the smallest of timeframes.

So look into today's hottest Ethernet chipsets. Call (800) 548-4725 and ask for Lit. Packet #YA23. And learn why we have the perfect match for your next box.

intel®

The Computer Inside.™

CIRCLE NO. 38

COMPUTERS & SUBSYSTEMS

Multimedia—here today, not tomorrow

Dave Wilson, Senior Editor

This month, AT&T Microelectronics (Allentown, PA) introduces the world's first chip set capable of handling all three of the major multimedia compression, decompression and transmission standards—MPEG, JPEG and Px64. Priced at less than $400, the chip set should enable true "all-singing, all-dancing" multimedia products to reach the desktop faster than anyone imagined. At least for now, the announcement puts AT&T ahead of multimedia chip set rivals—Intel (Santa Clara, CA), Texas Instruments (Dallas, TX), Cypress Semiconductor (San Jose, CA), and Motorola (Austin, TX), all of whom are working on MPEG/JPEG/Px64 multimedia chip designs, but none of whom have yet announced products.

A joint development between AT&T Microelectronics and the company's Bell Labs (Allentown, PA), the AVP1000 video codec (coder/decoder) chip set includes a decoder, a system controller and one of two encoders—one for low-cost applications and one for higher-performance systems. Working in conjunction with the AT&T DSP3210 digital signal processor, which provides MPEG/Px64 audio and JPEG still-image compression, the three-chip set provides a complete solution for teleconferencing and multimedia. In addition to supporting standards, the chip set provides a number of key system-level functions not covered by standards, including buffer management; synchronization of audio, video and user data; and acoustic echo cancellation.

MPEG, the hardest part

The hardest part of designing any multimedia chip set is building a cost-effective MPEG encoder. That's because the MPEG encoding algorithm requires processing power on the order of billions of operations per second, which until now has been difficult to architect cost-effectively on a single chip. To complicate the situation still further, although MPEG compression and decompression will be mandatory for designers of future real-time video systems, some applications may not need the power or require the expense of full MPEG encoding. To satisfy all camps, AT&T has built two versions of its encoder.

The less expensive 1300E performs full-motion H.261 compression for CIF (common intermediate format) and QCIF (quarter-resolution CIF) images at up to 30 frames/s. The encoded bit rate is user-definable over a range of 40 kbits/s to 4 Mbits/s, using either a constant bit rate or a constant quantization step size. To support digital editing and high-resolution still images, the 1300E can handle MPEG I-Frame-Only (Intra Coded Frames) compression. Spatial resolution is user-definable in multiples of 16 up to 720 pixels × 576 lines.

For those who need it, the fully functional 1400E encoder can perform both full-motion Px64 and MPEG compression. To enhance MPEG compression, the 1400E uses a hard-wired motion estimation processor that supports a full exhaustive search range of ±32 pixels with half-pixel accuracy. To achieve additional MPEG compression and improved picture quality, the device supports up to two successive interpolative frames.

Both encoders accept raster-scanned YCC (luminance and chrominance of signal) data via an input video bus or host bus interface and output compressed data via either the host or a serial bus interface. To absorb picture-dependent

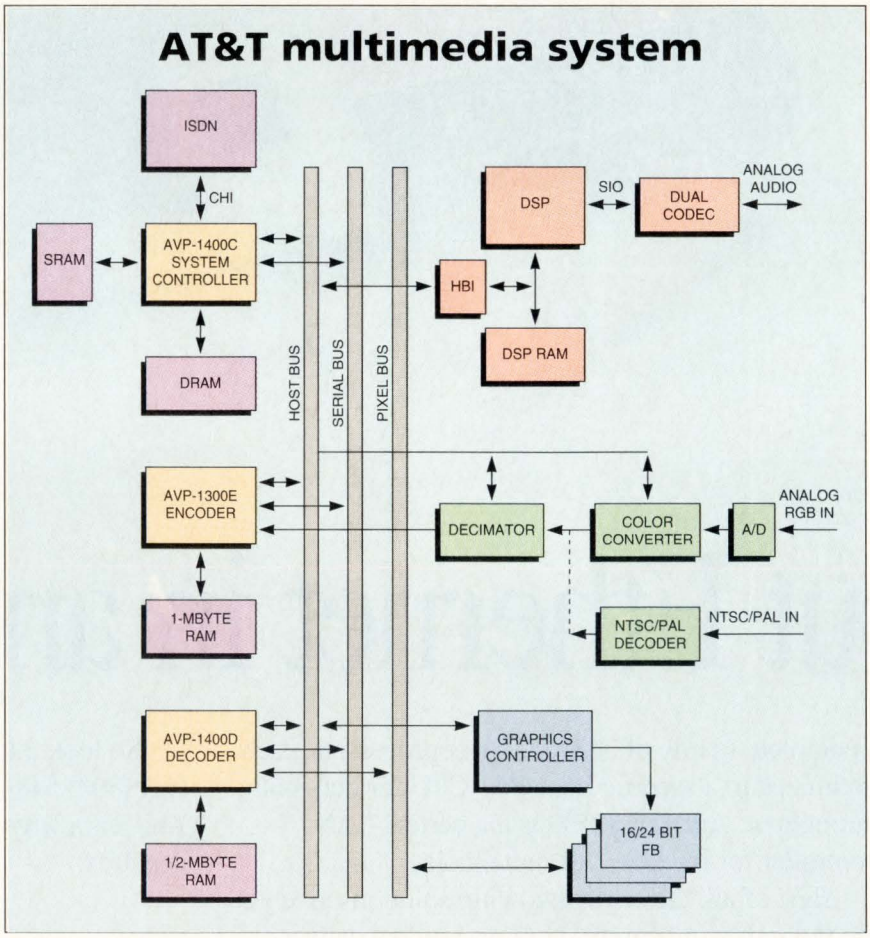

AT&T is the first vendor to offer an under-$400 multimedia chip set. Included in the offering are a system controller, decoder and two encoders, one of which can perform full-motion Px64 and MPEG compression.

NOTICE

WE'VE BEEN ASKED TO MAKE IT VERY CLEAR THAT
THERE'S A WHALE OF A DIFFERENCE BETWEEN QNX® AND UNIX®.

WE'RE GLAD TO OBLIGE.

EVEN THOUGH THE QNX REALTIME OPERATING SYSTEM IS *NOT* BASED ON UNIX SOURCE CODE, WE CAN UNDERSTAND HOW SOME DEVELOPERS MIGHT CONFUSE THE TWO. AFTER ALL, QNX FOLLOWS THE LATEST IEEE POSIX 1003.1 AND 1003.2 OPEN SYSTEMS STANDARDS, SO YOU GET THE SAME API AND UTILITY SET FOUND IN MANY UNIX SYSTEMS.

BUT LOOK BENEATH THE SURFACE AND YOU'LL SEE TWO FUNDAMENTALLY DISTINCT ARCHITECTURES. A MONOLITHIC OS ON THE ONE HAND, A MICROKERNEL OS ON THE OTHER. THEY'RE DIFFERENT SPECIES ALTOGETHER, AS DIFFERENT AS A WHALE AND A SCHOOL OF DOLPHINS.

ARCHITECTURE YOU CAN BUILD ON

QNX'S MODULAR, MICROKERNEL ARCHITECTURE GIVES YOU REMARKABLE FLEXIBILITY. OS SERVICES ARE PROVIDED BY A TEAM OF COOPERATING PROCESSES RATHER THAN BY A MASSIVE MONOLITHIC KERNEL. THE RESULT? YOU CAN STRIP QNX DOWN TO A ROM-ABLE EMBEDDED SYSTEM. OR BUILD IT UP TO A VAST NETWORK AND HARNESS THE POWER OF HUNDREDS OF CPUS.

OUR PRICING IS MODULAR TOO, SO YOU'RE NOT FORCED TO PAY FOR WHAT YOU DON'T NEED.

REALTIME PERFORMANCE YOU CAN COUNT ON

NEED THE SPEED OF A FAST, REALTIME EXECUTIVE? QNX CLOCKS IN AT 16 µSEC PER CONTEXT SWITCH ON A 33 MHZ 80486. BUT QNX DOESN'T STOP THERE. WITH ITS PRIORITY-DRIVEN, PREEMPTIVE SCHEDULING, YOU CAN BUILD *REAL* REALTIME SOLUTIONS.

DISTRIBUTED PROCESSING YOU CAN BET ON

QNX CAN TRANSFORM A BUNCH OF ISOLATED MACHINES INTO A SEAMLESS SUPERCOMPUTER,

ORCHESTRATING HUNDREDS OF CPUS WITH ITS NETWORK-WIDE IPC. YOU'RE IN COMPLETE CONTROL OF ALL RESOURCES AT ANY POINT, FROM THE PLANT FLOOR RIGHT UP TO YOUR DESKTOP.

SUPPORT YOU CAN DEPEND ON

DURING THE LAST TEN YEARS, WE'VE EARNED A REPUTATION FOR OUTSTANDING TECHNICAL SUPPORT. WE OFFER EVERYTHING FROM 24-HOUR ONLINE CONFERENCING TO ONSITE CONSULTING, SO YOU CAN EASILY REACH THE PEOPLE WHO MAKE UP THE QNX DEVELOPMENT TEAM. THEIR EXPERTISE CAN HELP YOU KEEP YOUR DEVELOPMENT PROJECTS RIGHT ON COURSE.

TO FIND OUT HOW YOUR APPLICATIONS CAN THRIVE IN THE QNX ENVIRONMENT, CALL 1-800-363-9001, (EXT. 103).

QUANTUM SOFTWARE SYSTEMS LTD.
175 TERRENCE MATTHEWS CRESCENT
KANATA, ONTARIO, CANADA
K2M 1W8
TEL: 613-591-0931
FAX: 613-591-3579

QNX®

REALTIME OPERATING SYSTEM

CIRCLE NO. 39

ARCHITECTURE MAKES THE DIFFERENCE

QNX IS A REGISTERED TRADEMARK OF QUANTUM SOFTWARE SYSTEMS LTD. UNIX IS A REGISTERED TRADEMARK OF UNIX SYSTEM LABORATORIES, INC. ©QUANTUM SOFTWARE SYSTEMS LTD. 1991

COMPUTERS & SUBSYSTEMS

fluctuations in the compressed data rate, the devices include on-chip FIFOs. When operated in the H.261 mode, they also use an adaptive quantizer/buffer control algorithm that lets you adjust picture quality, encoding delay and output frame rate. Both devices feature "glueless" interfaces with external DRAM and the system controller.

Both versions of the encoder are composed of 10 functional blocks. First, there's a host bus interface that synchronizes the host bus signals to the main input clock. Then, there's an uncompressed data FIFO used to hold the chrominance and luminance pixel data until it's written into the frame store DRAM. An on-chip memory controller arbitrates access to the frame store DRAM and provides a glue-free interface to DRAMs. The motion estimator compares current and previous frame macroblocks within a fixed search range, while it's the job of the quantization processor to determine the quantization step size and output frame rate. The signal processor itself comprises no fewer than six SIMD (Single Instruction Multiple Data) processors that work in parallel—performing functions such as discrete cosine transformations (DCTs), quantization and zigzag scan. The global controller sequences the operation of the other blocks to ensure that data dependency requirements between them are satisfied. A variable-length encoder operates on the information produced by the signal processor and generates a compressed bit stream that is compliant with the syntax requirements of the MPEG and H.261 standards. The compressed data FIFO section makes the external buffer DRAM operate like a 1-Mbit FIFO and absorbs picture-dependent fluctuations in the output data rate.

Decoding unlike encoding

Decoding MPEG signals, unlike encoding them, isn't as computationally demanding. The AT&T decoder, dubbed the 1400D, offers full-motion MPEG and Px64 video decoding, including an arbitrary number of bidirectional frames. It decodes compressed data streams at up to 4 Mbits/s, supports variable frame rates of up to 30 frames/s, and provides variable spatial resolutions for full-motion video up to CIF/SIF (standard intermediate format). It also performs MPEG still-frame decoding at variable resolutions up to 1024×1024. The decoder can accept compressed data via its parallel host bus or a serial bus. It outputs decoded raster-scanned 24/16/8 bit RGB or YCC pixels via either the host bus or a dedicated pixel bus. Key features include on-chip color conversion with programmable coefficients and dithering, a 4-kbit FIFO for storing compressed data and a glue-free DRAM and system controller interface.

Communications too

With any multimedia chip set, it's important to look for a communications solution as part of the multimedia mix. Next-generation systems won't simply require compression and decompression of live video, but will demand transmission of the multimedia data too.

Addressing this issue is AT&T's system controller. Known as the 1400C, it's a communications con-

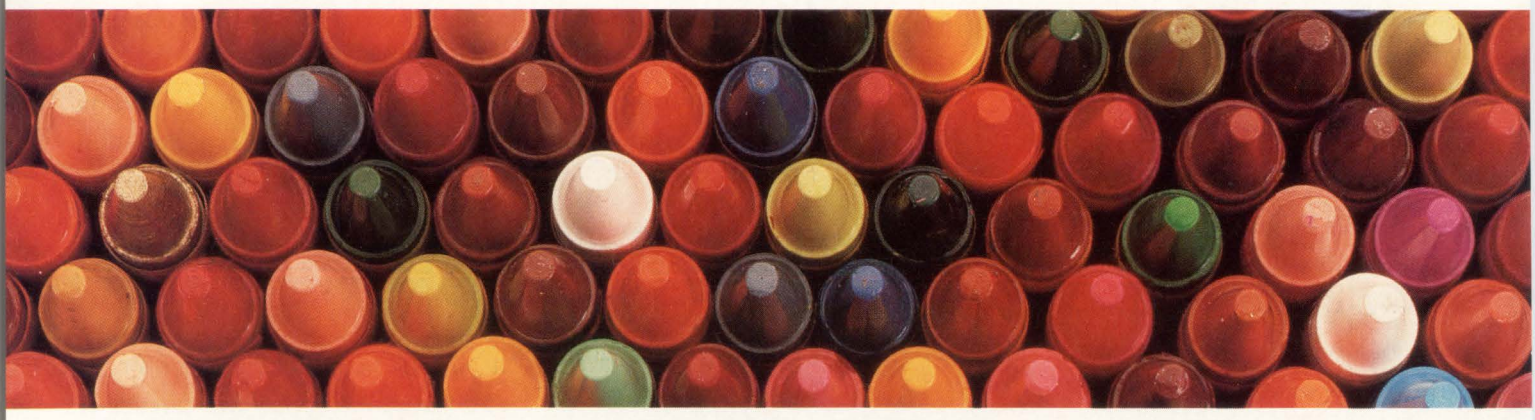

If your VMEbus or VAX needs 16.7 Million Colors and X Windows, call (510) 531-6500.

Peritek

Peritek Corporation
5550 Redwood Road, Oakland, California 94619
West Coast (510) 531-6500 fax (510) 530-8563
Mid West (303) 791-9161 fax (303) 791-9160

CIRCLE NO. 40

MULTI—The Source Level Debugger Fluent In Many Languages

MULTI gives you the ability to simultaneously debug in multiple languages: C++, C, FORTRAN, Pascal and Assembly code. You maintain seamless debugging even when your C++ program uses C and FORTRAN libraries. When the source language changes as you step into procedures, MULTI's expression evaluator and data structure displays change to reflect the current source language.

MULTI's X Windows™ customizable user interface (Motif™, OPEN LOOK™, etc.), lets you traverse your data and code in different views. Display source and assembly code intermixed. Click a mouse button to trace the hierarchy of C++ classes and member functions or disambiguate objects and functions. Change values interactively as you step through your code.

Your Passport For Native Or Cross Development

MULTI supports multiple targets such as the 680x0 (including 68040), as either a native debugger or as a cross debugger.

In cross mode, the same executable version of MULTI interfaces to your choice of debug execution devices (simulators, monitors, emulators) *without relinking*. Since MULTI is based on a standard UNIX ptrace interface with extensions, you can easily program your own interfaces to additional execution environments.

MULTI can be executed remotely from X-terminals across networks, or on any major UNIX workstation.

On most hosts, you can even debug multiple processes as your application forks.

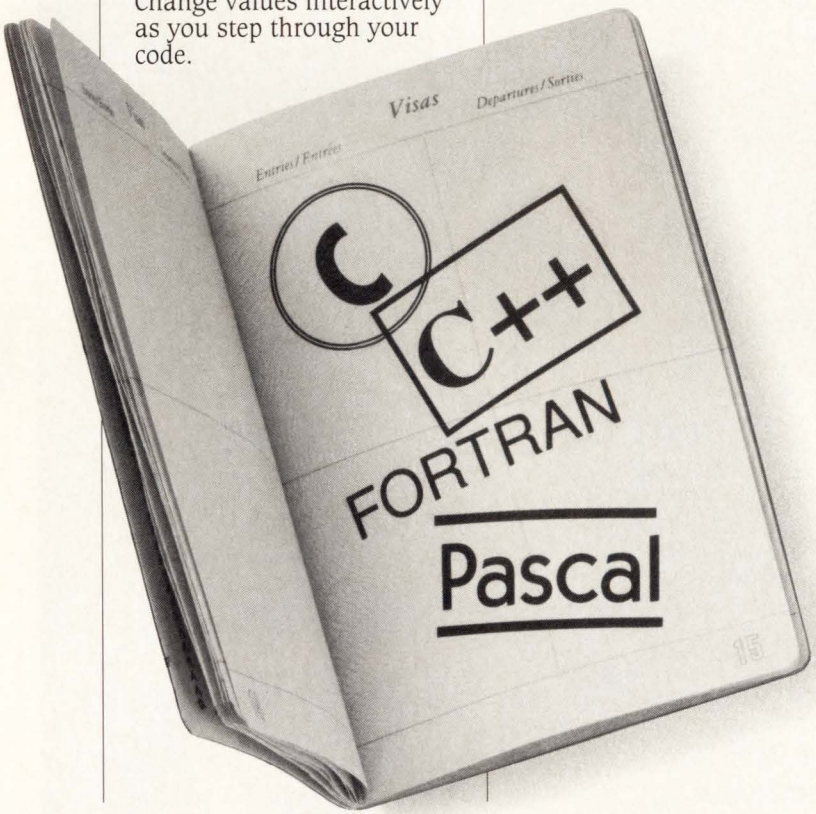

Let Oasys Be Your Guide

Oasys has years of experience with C and C++. We have the development tools you need today and the support, service and continuous technical advantage you need for a migration path to the future.

Oasys can guide you to:

- ☐ greater performance–quickly develop a competitive product that executes faster with smaller code size. Oasys delivers performance with the highly optimized Green Hills Compilers (C, C++, FORTRAN or Pascal). Our C++ Compiler takes you directly from source to executable with no translation required.
- ☐ productivity–debug native or cross applications on a workstation, a dumb terminal, across a network, or across the country using the same debugger commands and interface. MULTI helps you find and fix bugs fast.
- ☐ portability–Oasys delivers a complete software development tool kit including a common set of compilers, assembler/linker and debugger across major UNIX workstations, to take your application to any environment.

Call Us Today—
We Speak Your Language:

(617)862-2002
Or Fax: (617)863-2633

Trademarks are acknowledged to AT&T Bell Labs, Digital Equipment Corp., Green Hills Software Inc., MIT, Motorola Corp., Open Software Foundation, Sun Microsystems Inc., and XEL, Inc.

Call Now for a FREE MULTI Demo Disk!

(for Sun-4 SPARC only)

Oasys

One Cranberry Hill, Lexington, MA 02173

COMPUTERS & SUBSYSTEMS

troller based on the AT&T LAN Pacer architecture, designed to reduce audio and video system bus traffic and offload the host processor from a variety of system-level management functions. These include the multiplexing and demultiplexing of compressed MPEG audio, video and user data according to the MPEG system CD specification, including the MPEG reference clock, presentation time stamp and buffer and bitstream management. In addition, the device multiplexes, demultiplexes and frames Px64 audio, video and user data according to the H.221 specification, as well as performing forward error correction, channel management and setup. It buffers and synchronizes audio and video according to the MPEG and H.261 specifications, and can combine and synchronize multiple communications channels.

Configured from the host bus, the chip can operate in either MPEG or Px64 modes, receiving compressed audio and video information via either its host bus or serial interfaces. Full-duplex operation is supported for video teleconferencing applications. The device can send or receive H.221 framed data either over the host bus or via a 4-wire time-division-multiplexed link known as the Concentration Highway Interface (CHI).

The CHI provides a direct interface with external communications standards with channels such as Accunet Switched 56, Switched 64, Switched 56, T1, and ISDN. Later this year, look for AT&T to expand its communications support for the video codec chip set by introducing a silicon solution for FDDI II, which adds an isochronous transmission capability to FDDI's existing packet switching capability.

Most of AT&T's multimedia chip solution will be available in the third quarter. The exception is the 1400E encoder, which isn't expected until the fourth. So don't expect board or system products until later in the year. Nevertheless, many manufacturers of teleconferencing and multimedia equipment have already looked at the architecture on paper, and early indications are that the chip set will find many homes in such applications over the next few months. ■

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

AT&T Microelectronics
(800) 372-2447 ..... **Circle 296**
Cypress Semiconductor
(408) 943-2600 ..... **Circle 297**
Intel
(800) 943-2600 ..... **Circle 298**
Motorola
(512) 891-2000 ..... **Circle 299**
Texas Instruments
(800) 336-5236 ext. 700 ..... **Circle 300**

Mixing REAL TIME With REAL UNIX® Is Not Magic...

It Is Technology.

VenturCom's VENIX™ family of operating systems is the only realtime USL UNIX System V for Intel based platforms. Not a clone. Not just a realtime executive on top of UNIX.

Real UNIX provides designers with SVID and POSIX standards, access to NFS, X-Windows, Motif, streams, complete development tools and DOS tasks under UNIX. Plus, close tracking of future technologies.

Realtime adds performance and functionality to UNIX O/S with full kernel preemption, priority scheduling, contiguous file system, average interrupt latency of 50 μs, bounded context switches, memory locking, asynchronous I/O and much more.

VENIX and Embedded-VENIX products are used by Rockwell, Toshiba, General Electric, Honeywell and other leading edge companies. Ask us for our free Realtime UNIX/ VENIX Technical Overview.

VenturCom
Real Time/Real Unix
215 First Street, Cambridge, MA 02142 (617) 661-1230

CIRCLE NO. 42

BSO/Tasking Sets the Pace...

with the most comprehensive software development solutions for embedded applications!

Team-Up with BSO/Tasking

Jump out in front with BSO/Tasking's comprehensive software development solutions for embedded RISC, CISC, and DSP applications! With 150+ employees in 6 offices worldwide, BSO/Tasking's high quality tools, customer support, documentation, and training make application design and development a ride in the park!

Optimizing Compilers, Assemblers, and Linkers

BSO/Tasking's optimizing C Cross Compilers give you the performance edge by generating *fast*, accurate code for more than 40 microprocessors including complete support for the **Motorola 68040**, **68340**, and the **MIPS R3000** family of derivatives. All Cross Assembler/Linker/ Locators produce ROMable code and complete debugging information. Choose from 10 host development environments including PC/DOS, Unix workstations and VAX systems.

CrossView Source-level Debugger

CrossView's multiple windows display user-monitored data, registers, stack trace, C statement and instruction trace, and simulated I/O. Pop-up windows may also be activated to show current break-points, user defined functions, and two levels of help on all CrossView commands.

HP 64700 Emulator Support

Complete your development with the strong debugging capabilities of CrossView integrated with the Hewlett Packard 64700 In-Circuit Emulator (ICE), both available from BSO/Tasking.

100+ Combinations

Hosts

- PC/DOS
- Sun-3
- VAX/VMS
- HP9000
- Solbourne
- DECstation
- Sun SPARC
- VAX/Ultrix
- Apollo
- IBM RS/6000

Targets

- MIPS R3000 Family of derivatives
- 68000-040
- SAB80C166
- Intel 80x86
- 80C196 Family
- TMS320C10-C40
- NEC V20-V70
- 68302-340
- 68HC11, 16
- 68070/90C1xx
- Z80/64180
- Z8000
- 8051 Family

Call 1-800-458-8276

BSO/Tasking
333 Elm Street
Dedham, MA02026
P: (617) 894-7800
F: (617) 894-0551

BOSTON SYSTEMS OFFICE TASKING

HEWLETT PACKARD
Authorized Distributor

• Tasking the Netherlands 31 33 55 85 84 • BSO/Tasking UK 44 252 510014 • BSO/Tasking Italy 39 2 6698 2207 • BSO/Tasking Germany 49 7152 22090 • BSO/Tasking France 33 130 54 2222

CIRCLE NO. 43

Now you can and not get

Introducing new RISC System/6000 POWERstations

If you're interested in open systems but don't want to suffer the slings and arrows of outrageous prices, IBM is about to hit you where you live. The RISC System/6000™ POWERstation 220 gives you more wallop for your money, while delivering a hefty 25.9 SPECmarks.™ That's compared to the SUN IPC's™ 13.4 SPECmarks and the DEC5000's™ 17.8.

Model	Entry Grayscale Workstation**	Entry 8-bit Color Workstation†
IBM 220W	$7,185	$9,995
HP 705/710	$8,415	$14,065

Scientists see stars. CASE users can start with a grayscale workstation with a paging disk for just $7,185. If it's CAD clout you're after, you can get a workstation specially outfitted for mechanical design—with 2D color graphics and 400MB of fixed disk storage—for only $9,995. All models in the POWERstation 220 series come with two expansion

*In Canada, call 1 800 465-1234. **16MB, Paging Disk, Display, Operating System, Graphical User Interface. †16MB, 400MB Disk, Display, Operating System, Graphical User Interface.
IBM is a registered trademark and RISC System/6000 is a trademark of International Business Machines Corporation. SPECmark is a geometric mean of the ten SPECmark tests and is a trademark of Standard Performance Evaluation Corporation. All SPECmark figures listed are as published by their respective manufacturers. All prices listed are MSRP. Remarketer prices may vary. IPC is a trademark of Sun Microsystems, Inc. DEC5000 is a trademark of Digital Equipment Corporation. UNIX is a registered trademark of UNIX Systems Laboratories. HAGAR THE HORRIBLE Character(s) © 1992 King Features Syndicate, Inc. ©1992 IBM Corp.

get more clout, clobbered.

and POWERservers that pack more punch for less.

slots and upgradable components. And industry-standard memory upgrades and add-ons for both are affordable, so growing won't be a pain.

Striking a blow for business. The POWERserver 220 is great for commercial UNIX® solutions, too. You can configure it as a commercial server, to give your business the speed, muscle and openness of UNIX, for only $9,715. And the POWERserver 220 is as expandable as all our other models.

IBM is in your corner. Nobody else delivers the knockout support of IBM. An IBM customer engineer can install your machines, configure your network and integrate all your systems, whether they're made by IBM or not. And IBM Credit Corporation has flexible financing packages to meet your needs. Get hit with the details. Call your IBM marketing representative or Business Partner. For literature, call 1 800 IBM-6676, ext. 769*

And, for those who decide to shop for UNIX solutions elsewhere, a word of advice. Duck.

The RISC System/6000
For the Power Seeker.

IBM®

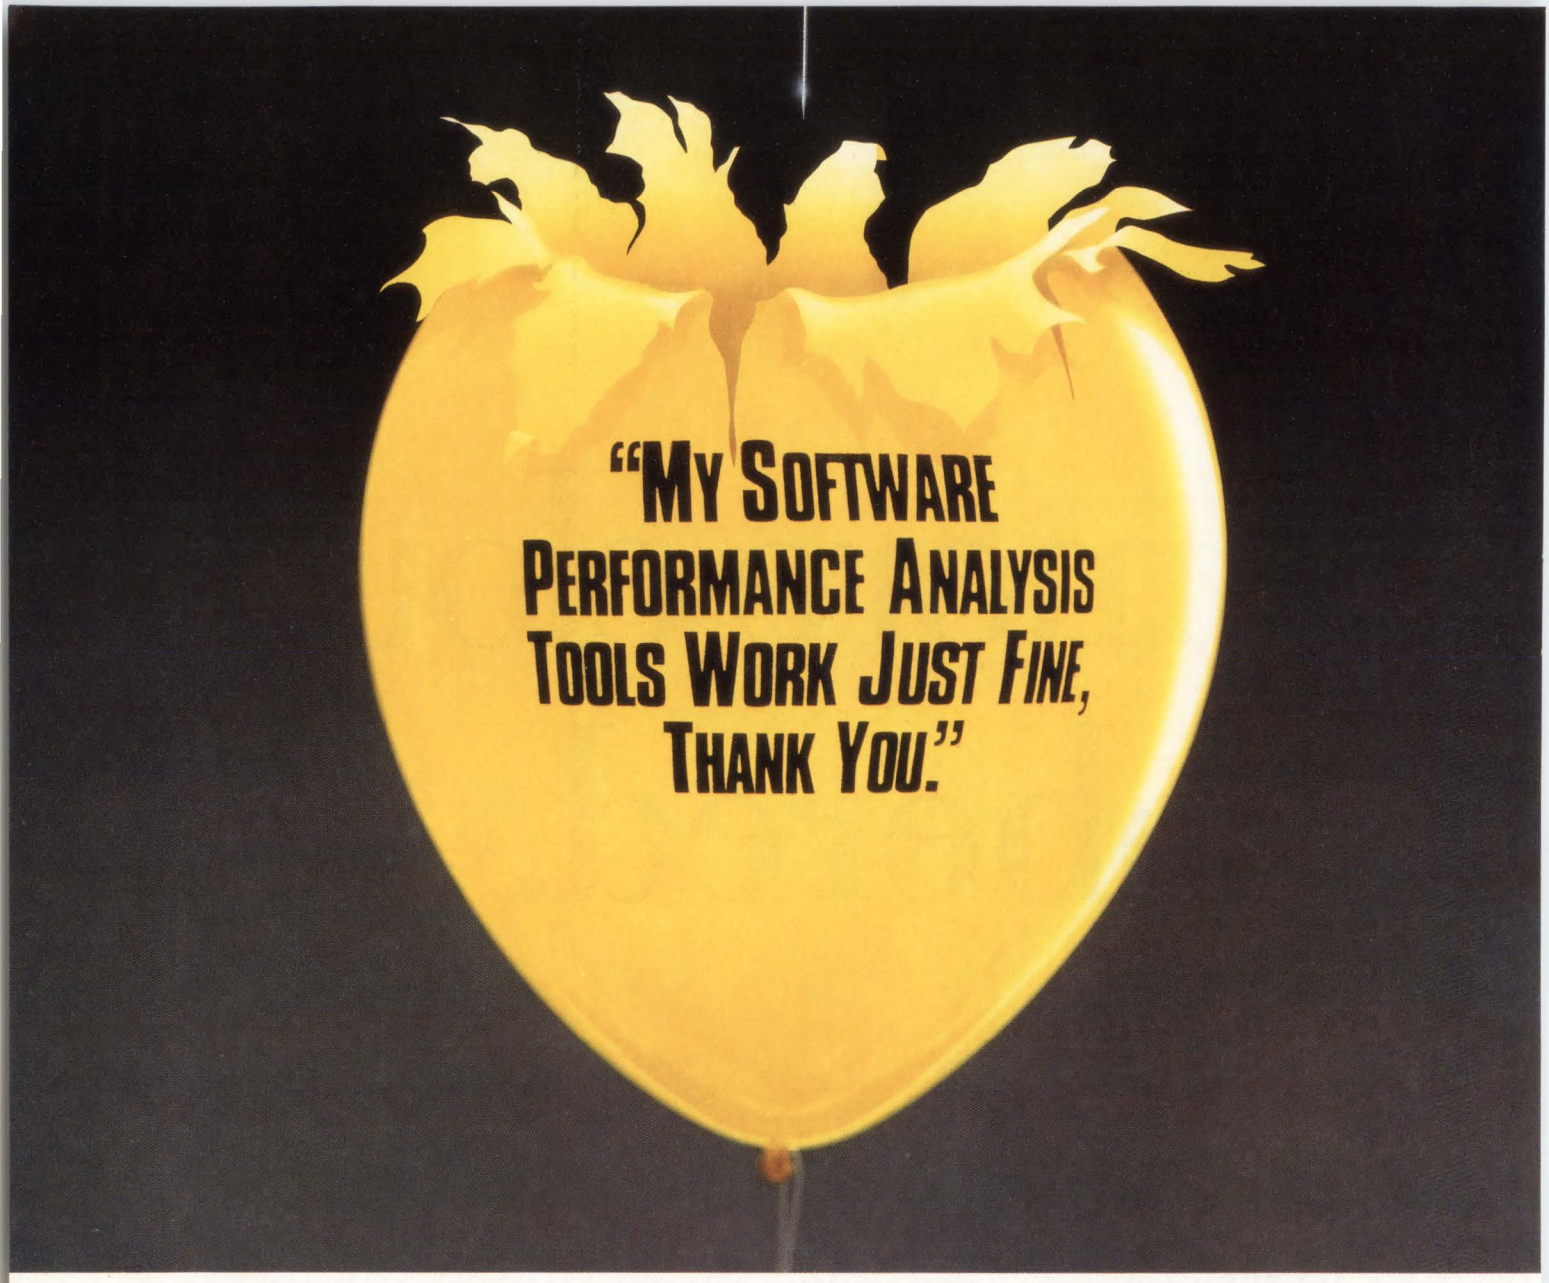

Truth is, without real-time performance analysis, you can only stretch things so far. Because unless you see every event, every bus cycle, and every routine, you may miss something important. At Tektronix, our PRISM 3000 logic analyzers can quickly detect performance problems that often creep into complex embedded systems.

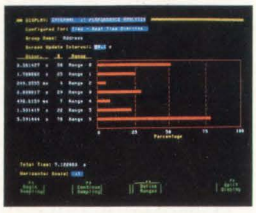

That's because they're real time. They don't miss a thing. And they're 100% accurate, so you can fine tune your software with absolute confidence. Now, if you're surprised to find all this in a multi-purpose logic analyzer, we understand. But once you see one in action, we think you'll appreciate our point of view. Call us for a demo today. **TALK TO TEK/1-800-426-2200 EXT. 70.**

Tektronix

Test and Measurement

5EW-188199 Copyright© 1991, Tektronix, Inc.

CIRCLE NO. 44

Is RISC or DSP best for your application?

RISC processors have already kicked CISC out of more than a few embedded sockets. Now, DSPs are trying to prove they can be just as useful in embedded applications.

Dave Wilson
Senior Editor

Spectrum Signal Processing's Ross Mitchell feels that DSPs have the edge over RISC processors in areas where the devices are used in numerically intensive parallel-processing applications. That's because some DSPs provide hooks on-chip for interprocessor communications; RISCs lack such hooks.

Traditionally, RISC processors have been built for Unix workstations. Their price tags alone have kept designers of embedded systems from giving them the time of day. Embedded design has consequently remained the realm of the CISC processor.

Today, things are different. Many vendors offer inexpensive RISC devices for the embedded marketplace. Lower cost makes them an attractive alternative to CISCs and their derivatives.

At the same time, the digital signal processor (DSP) has fallen in price and gained so much functionality that it may be considered a viable general-purpose embedded processor as well. "DSPs and RISCs are more alike than people think," says Michael Tighe, senior systems architect at Intermetrics Microsystems Software (Cambridge, MA). So, designers upgrading from CISC and needing to choose a high-performance embedded processor are now faced with a choice. Should you take a RISC or should you choose a DSP?

In many applications, such as mobile phones and Unix workstations, the design decision is clear-cut—cost, performance and time-to-market alone mandate the use of one technology or the other. A Unix system designer would no more build a workstation around a DSP chip than a mobile phone designer would choose a RISC processor. Indeed, it's not the extreme ends of the applications spectrum where the two technologies will meet face to face, but the gray area somewhere in between. In that nether land lies a large set of embedded applications equally suited to DSP or RISC solutions. These applications could be compute-intensive or, perhaps, aimed more at control tasks.

Philip Brownfield, manager of RISC applications design for Motorola

(Austin, TX), thinks that a RISC processor solution should be a viable alternative to DSP when a development team can take advantage of a standard compiler to generate application code. Opposed to this solution is what Brownfield calls the "coding gymnastics" of assembly language programming, which designers must use to optimize the performance of a system based on a DSP. "When people consider using RISC over a DSP, it's usually because of the ease of coding. A RISC lends itself to the maintenance and portability of the code," he says. Nevertheless, moving away from a DSP/assembly language programming environment to one based on RISC processors programmed in C isn't always so easy. Especially if 100 percent of an existing application is already written in assembly language—and not in C.

DSPs retain popularity

In fact, one reason that DSPs remain as popular as they do is that they actually lend themselves to assembly language programming. It's by programming in assembly language that you can squeeze every last drop of performance out of the silicon—and that's not as easy with a RISC chip.

In AMD's 29050 RISC chip, for example, programmers have no access to the multiply pipeline; it's impossible to examine partial results of the multiply operation. Faced with such a constraint, you have only two options: either stay with or return to a DSP solution, or change your programming methodology to gain the benefits of faster time-to-market using a RISC processor.

You might like to take advantage of the high-level language programming benefits of a RISC machine but it may not always be possible, simply because the performance of the DSP algorithm running on the RISC device may never match the achievable performance on the DSP chip. On the other hand, if the RISC machine can be shown to meet the performance requirements of the system, then you have some convincing reasons to use it, especially in place of earlier CISC/DSP combinations.

Although it may not be able to achieve real-time speed, a single-processor approach can meet the design specification and still reduce the number of processors that need to be programmed, lower system cost and reduce the design effort and time-to-market. Furthermore, such an approach may eliminate the need to employ a full-time DSP expert in the design group.

On the other hand, though, even RISC aficionados such as Bruce McKeever, an engineer at Fujitsu (San Jose, CA), aren't prepared to believe that RISC processors and C programming techniques will eliminate the need for DSPs altogether. Citing some poor results that he has seen, McKeever says, "You do pay the penalty for portability."

Even designers that try for the best of both worlds—programming DSP chips in C and making use of optimizing C compilers—note that there's some work involved in structuring C programs so that the optimizer creates the most streamlined code.

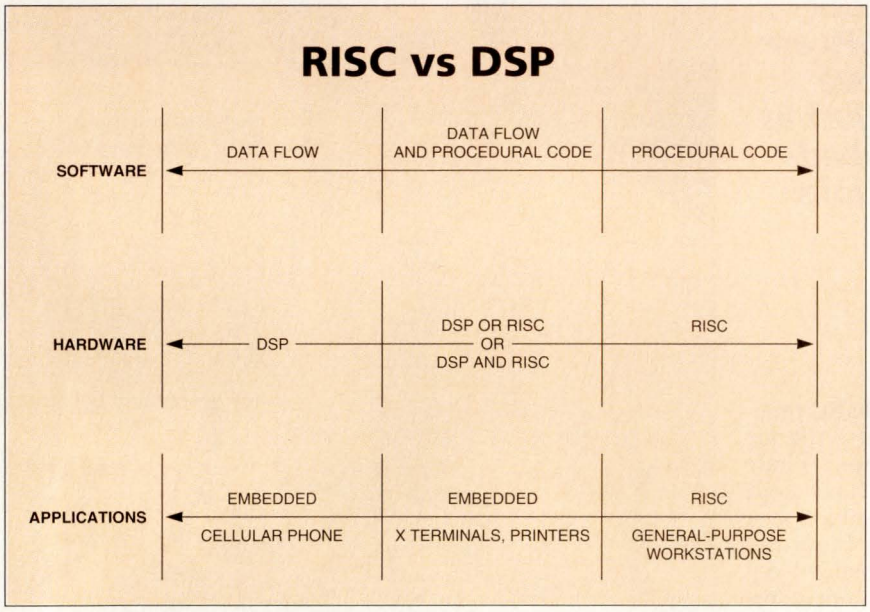

You can decide to use a RISC or a DSP chip based on the type of code running in the system. If the system software is mostly procedural, then RISC is the best choice. If the code is data-flow-oriented, then DSP is better. It's the area in the middle, where the two types of code are mixed, where you have the hardest time deciding which route to take. In some cases, where determinancy of DSP code and performance of procedural code are both very important, a RISC/DSP combination could give the best performance of all.

The need for tight code

There's no doubt that the size of the code generated on a RISC is larger than that for a CISC processor. And that's an important consideration in embedded systems, where keeping memory cost down is as important as keeping processor cost low. Intermetrics' Tighe performed some benchmarks and found that the size difference in embedded code depends on both the processor and the compiler used to optimize the code. "A 1-kbyte executable on a 68000 might be 2 or 3 kbytes on a 29000. That's without optimization. But once you apply a global optimizer to assign common subexpressions, you discover that you get as much as 50 percent code shrinkage on a RISC machine, whereas you only get 20 percent on a 68000. The results are similar for many DSP chips," Tighe says.

There are also new technologies being used in compilers, particularly in areas such as instruction coalescing, that can be used to produce even greater improvements.

One challenge that tool providers such as Intermetrics face is the need to bare as many hardware capabilities as possible to the programmer. And that, it seems, is more important for a DSP than for a RISC. Obviously, programmers writing in C don't care about such capabilities, but those designers that want to optimize their code to the architecture of the processor most certainly do.

To provide such capabilities means customizing the compiler for access to the chip. "On the 96002, we allowed access to the different memory models and we added keywords

Performance Integrated MIPS Module

40-MHz PIMM

The easiest & lowest cost path to 33 MIPS and 6.5 MFlops* with 5,000,000 Transistors inside a Single VLSI Package

*Double Precision LINPACK

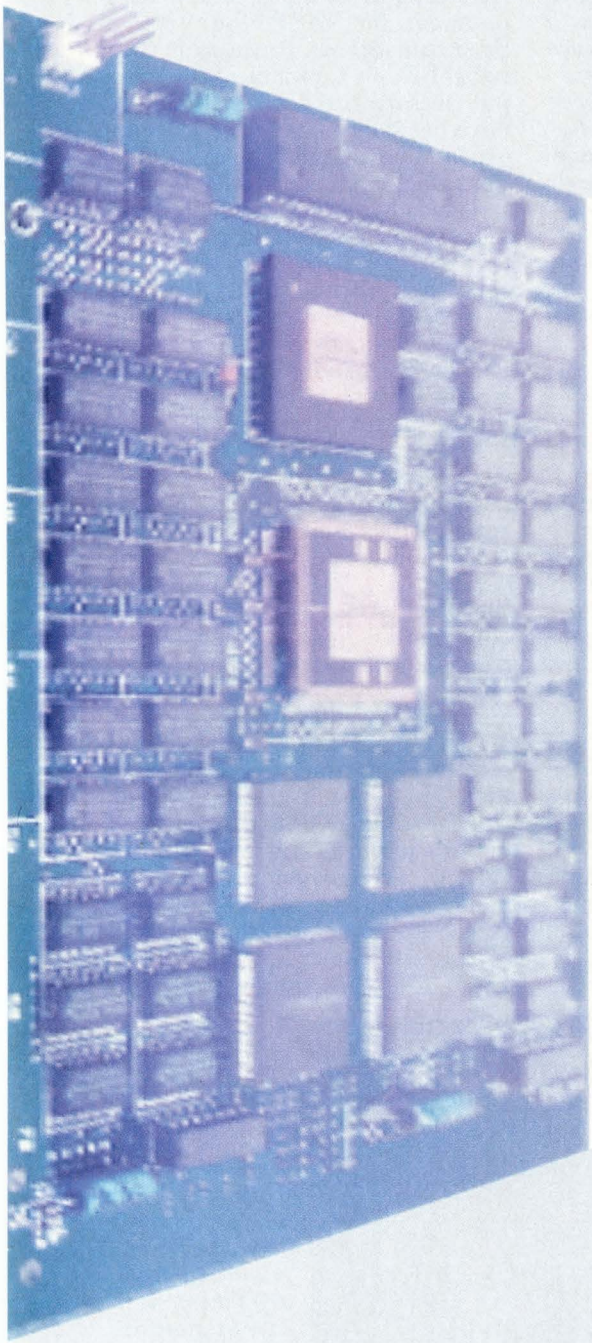

Actual Size
144-pin PGA
1.96x1.96"

Contains MIPS RISC CPU, Floating Point Accelerator, Write, Read & Parity Buffer, 32KB Instruction & 32KB Data Caches

Easy to Design-In

Designers no longer have to contend with the high-speed cache bus or other high-frequency signals because they are integrated in PIMM's VLSI package.

PIMM also provides a fully flexible and simple interface that helps designers to build high performance workstations, servers, and embedded systems.

Single Package with Minimum Board Space

PIMM integrates a CPU, FPA, 32KB I & 32KB D caches, 8-word-deep write buffer, and a 32-word programmable read buffer in a 2-inch square package — thereby reducing board space and cost ***dramatically***.

Full MIPS Software Support

Mature compilers and operating systems exercised on hundreds of thousands of production systems by dozens of manufacturers over three years.

For more information from the leading volume and speed supplier of MIPS RISC components, call

(408) 734-9000

PERFORMANCE
SEMICONDUCTOR CORPORATION
610 East Weddell Drive, Sunnyvale, California 94089

CIRCLE NO. 45

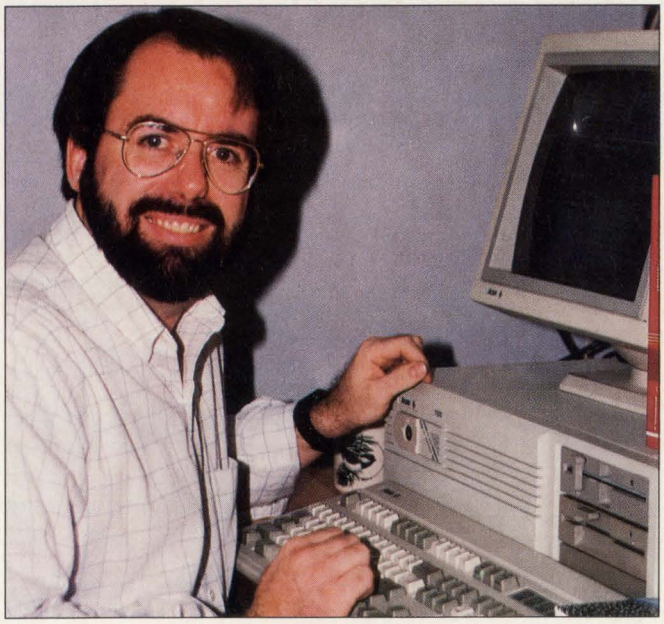

Michael Tighe of Intermetrics says that designers choosing between DSP and RISC processor solutions should pay careful attention to the quality of the compilers. In his experience, customizing a compiler for a DSP architecture is more important than doing the same for a RISC microprocessor because of each DSP's special capabilities.

to the language as extensions to let the programmer who knows what he's doing get every last ounce of performance out of his machine," says Tighe. To perform such customization on a RISC machine, Tighe admits, is less worthwhile. "On a RISC machine there's almost nothing that you can do—except make better use of the C 'register' keyword, because it's the registers where you spend a lot of time in a RISC machine."

The key differences

Clearly, the differences between the architectures of a RISC and a DSP processor are due to design tradeoffs made by the manufacturers in the early stages of the chips' development. But, aside from the high-level language versus assembly language debate, what is it about the architectures that make one better suited to a given class of applications than another?

One method that might have been used in the past to answer this question was to count the number of registers in the device—a DSP used to sport far fewer than a RISC. But today, DSPs such as the 96002 have so many registers that they're beginning to look like RISCs. Just counting registers may not be the best way to determine the difference between them anymore.

There are some other obvious differences. One important one is that, while a DSP can perform a single-cycle, high-precision multiply-accumulate operation in one cycle, that's still not true with RISC. On the other hand, comments Intermetrics' Tighe, "If you're less interested in the precision of a result in your algorithm, then a RISC can be a better choice than a DSP because of its higher clock rate." RISC machines offer the programmer a general-purpose CPU that isn't tuned for any specific application.

RISC processors may have the edge over DSPs in control-oriented applications, where the application code will be required to manage many different kinds of data streams. That's because the architecture of a DSP probably wouldn't be used to its fullest extent in such applications. On the other hand, Tighe feels that DSPs emerge as clear winners in applications such as cellular phones, where the machine needs to manipulate a stream of data with very high precision and little loss due to error.

DSP architectural features

According to Ross Mitchell, product manager for VASP and VME at Spectrum Signal Processing (Burnaby, British Columbia), the internal architectures of DSP devices have been optimized for high-speed execution of signal-processing functions, including dot products, correlation, convolution, and fast Fourier transforms (FFTs). Key architectural features include fast hardware multipliers, independent data and program storage areas, fast interrupt switches for efficient context switching, and multiple ports for the design of efficient parallel processing architectures.

DSP designers not only build devices with fast multiply-accumulators, but they also include multiple register files and datapaths to let you simultaneously access separate register files. This is to ensure that the processing rate in the DSP is only limited by the speed of data transfer to and from the processor and the memory subsystem. Such an arrangement, of course, will be most efficient while the DSP operates on a stream of data—like an FFT. Any break in the data stream will cause the DSP to pause, resulting in a performance loss.

RISC processors, on the other hand, are built around highly integrated and pipelined floating-point arithmetic logic units—multiply and add—operating at very low

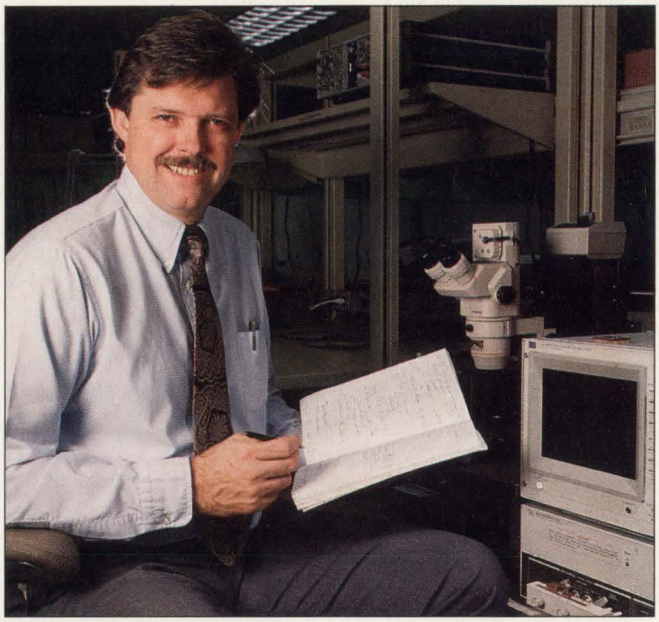

For certain applications, the RISC/DSP discussion is moot. "I'll be damned if someone would put a RISC in a cellular phone and have it be low-power and small. All of them use DSPs as cheap as they can get 'em," says Motorola's strategic applications manager, Garth Hillman.

RISC vs DSP in embedded applications

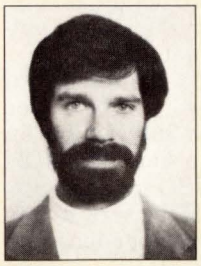

When deciding between a RISC and a DSP, you may find that cost, speed and software development time are more important than device architecture.

RISC processors were designed for efficient compilation of programs written in C or C++, and now DSP chips sport C compilers, too. C compilers for both, however, may produce code that's too inefficient to run many signal-processing applications in real-time.

It's true that RISC instructions are simple enough to run in a single clock cycle and serve as fundamental building blocks for a compiler. But several simple RISC assembly language instructions must often be written to perform an operation performed efficiently by one DSP assembly language instruction.

In many cases, it's simpler for an experienced programmer to write a signal-processing task in assembly language for a DSP than to write the same task in assembly language for a RISC processor. Nevertheless, a RISC processor may be the right choice if it can run compiled C or C++ code sufficiently fast for the applications involved. Algorithms may be developed with the DSP's C compiler and later optimized—or coded for speed in assembly language—when the algorithms appear to be stable.

Addressing modes

Some DSPs provide addressing modes that reduce computation time as well as the time it takes to code many signal-processing applications. Bit-reversed addressing is provided in modern DSPs for efficient computation of FFTs. Module addressing is useful in accessing tables, so that an address wraps around within the address range of a table instead of overflowing to memory above the table. Many DSPs also provide saturation arithmetic where the result of a computation is limited to the maximum positive or negative number rather than overflowing. Rounding toward 0 is another DSP feature useful in preserving signal fidelity. Few, if any, RISC processors provide such capabilities.

Both DSPs and many RISC processors suffer from numerical problems. To preserve a 16-bit signal-to-noise ratio, internal processing, particularly infinite impulse response filter structures, should be performed with high precision. Floating-point processors typically truncate the product's mantissa to 24 or 32 bits with every accumulation. The resulting nonband-limited error signal aliases into the useful band and may cause problems if there's sufficient gain in subsequent processing.

This and other subtle problems that can occur with floating-point arithmetic can be solved using double-precision arithmetic. With a minimum consumption of die area, double-precision arithmetic can be performed in four to six times the compute time of single-precision arithmetic. But most processors can't efficiently perform double-precision arithmetic. Motorola's DSP56001 has solved many of these problems by providing a sufficiently wide fixed-point accumulator (56 bits), so that the sum need not be truncated until all processing is complete.

The multimedia market

Texas Instruments and AT&T have recently targeted 32-bit floating-point DSPs at the multimedia market. Their DSPs provide more internal memory, a wider address space, DMA, and a better bus structure than the DSP56001. A floating-point multiplication may be performed every two clock cycles for $25 on a TMS320C31, which includes a synchronous serial port, two timers and over 8 kbytes of internal RAM. What RISC processor with an equivalent amount of RAM and a synchronous serial port can match this price/performance ratio?

DSPs such as Motorola's DSP96002 and TI's TMS320C40 provide the greater horsepower needed by multimedia applications, but at costs that would force many end products to be priced too high. Few RISC processors provide 24- or 32-bit-wide multiplication in one or two clock cycles. Intel's i860 is an exception in providing fast, wide, floating-point multiplication, but at a relatively high price. The i860 is also awkward to program because of its pipeline complexity. Overall performance of the i860 may be insufficient for real-time applications because you have to flush the pipeline and caches for context switches.

Although current DSPs can't execute one or more instructions per clock cycle, their Harvard architecture, multiple buses and multiport RAMs allow several accesses per instruction cycle. The TMS320C30, for example, can read an instruction from its cache while simultaneously reading two operands from on-chip RAM, and then can allow two more RAM accesses (e.g., for I/O) in the second half of one two-clock-cycle instruction. Data, parameters and instructions must be read from off-chip memory when there's insufficient memory within any processor. Since off-chip accesses often result in a speed penalty because of bus contention, it's desirable to provide as much on-chip RAM as price will permit.

Embedded controllers

Embedded controllers typically must provide many useful peripherals on-chip. Some DSPs provide an on-chip converter, usually with a limited signal-to-noise ratio. Most DSPs also provide high-speed synchronous serial ports for connection with 16- to 20-bit A-D and D-A converters, or they have external ports for standard data exchange (e.g., AES/EBU format for serial transmission of digital audio data). DMA units with many channels are useful for moving data between internal and external RAM and peripherals, while minimizing interference with internal signal processing.

DRAM controllers are beginning to be included in embedded RISC processors, and should soon appear in DSPs. Many designers, however, would prefer to use an external DRAM controller if the on-chip controller doesn't provide an efficient fast-page mode.

RISC processors offer the promise of efficient C and C++ compilers as well as high-level development tools. Many signal processing applications require real-time performance at a lower cost, however, than that of currently available embedded RISC processors. RISC processors with fast, wide multiplication, many kwords of internal RAM (or lockable caches with bus snooping) and more useful peripherals should eventually appear and drop in price. DSPs will also evolve, with single-cycle instructions and efficient compilation of C++.

John Snell, president, Timbre Engineering, San Geronimo, CA

clock periods—typically 20 ns. But it's just because of their speed that they can be programmed to provide throughputs comparable to that of their DSP counterparts. For example, as a single-processor accelerator to an embedded single-board computer, a RISC-based board provides an alternative to DSP boards, according to Mitchell. He likes to compare a 40-MHz i860 with a 40-MHz DSP96002. While the i860 provides 80 MFlops of raw computing power, the DSP96002 is specified at a peak of 60 MFlops. Both devices can perform a 1K complex FFT in just under 1 ms. Based on raw computing power, the i860 operates at 25 percent efficiency and the DSP96002 at 33 percent. But, Mitchell adds, the application software is much simpler in terms of size and complexity on the DSP96002 than it is on the i860, where efficient management of the internal pipeline is critical.

TMS320C3X comparison with embedded controllers

Same general-purpose features			With DSP capability		
General-purpose features	**C3X**	**RISC/ CISC**	**High-performance features**	**C3X**	**RISC/ CISC**
Unified address space	Yes	Yes	Single-cycle all instructions	Yes	No
Register-based CPU	Yes	Yes	Parallel MPY & ALU	Yes	No
Software stack	Yes	Yes	Delayed branch	Yes	Yes
Flexible addressing modes	Yes	Yes	Zero overhead loops	Yes	No
Cache	Yes	Yes	Circular addressing	Yes	No
Single-cycle bit test	Yes	Yes	Scan-based emulators	Yes	No
C language support	Yes	Yes			
O/S support	Yes	Yes			
C debugger/profiler	Yes	Yes			
W/S development support	Yes	Yes			

The people at Texas Instruments believe that their TMS320C3X family stacks up very well against RISC and CISC processors for embedded applications. In addition to the general-purpose functions (left), DSPs sport many additional features RISCs and CISCs don't (right).

A closer look

Designers of DSP chips are aware that their devices will spend most of the time running algorithms with very unique needs not found in general-purpose computing. For example, the devices perform multiply-accumulate operations in tight loops, one instance of this being a sine series expansion. This means you must build characteristics into your DSPs that will never be found in a general-purpose RISC machine. One common example of this is the 40-bit accumulator. Intermetrics' Tighe observes, "DSPs have super-wide registers of 40 bits or larger to maintain high precision. If you lost the low bits by truncating the result of a multiply-accumulate to 32 bits, you would get an accumulative error that, after several multiplies, turns out to be almost half the word." Features like 40-bit-wide registers can't be found in RISC chips, because they're not specifically designed for performing hundreds of multiply-accumulate functions. RISCs tend to perform a mixed bag of instructions—a multiply, an add, a subtract—all in random order.

A wide register file, of course, isn't the only case where a DSP is optimized toward executing a specific class of code. David Fair, marketing manager for sound processing at Analog Devices (Norwood, MA), notes that since loops typified by multiply and add operations are fundamental to many signal processing algorithms, it shouldn't be surprising to see on-chip support for looping operations on DSPs. When a DSP program can be expressed in loop form, the coding can be simplified and shortened.

When address looping is used on Analog Devices' ADSP-2100, for example, it's automatic and transparent. Without any extra checking cycles the processsor determines if a loop should terminate, either because it has run the specified number of cycles or because another termination condition was met. The processor then outputs the next instruction address. If loop iterations continue, the address of the first loop instruction is output. In one cycle, the last instruction of the loop is executed, and on the very next cycle the next instruction is executed, either within the loop or outside it when terminated.

By having on-chip support for looping operations, you're saved from issuing a new instruction for every iteration of a loop, as would be the case in a RISC implementation. That gains DSP users a two-processor-clock-cycle-per-loop advantage over a RISC design implementation.

According to Rick Rienhart, floating-point DSP program manager at Texas Instruments (Houston, TX), another special design feature found in a DSP but not in RISC is support for circular addressing. Circular addressing lets the processor treat a small region of memory as if it were an infinite memory buffer by looping through it. This is a useful way to deal with infinite streams of data, such as speech or motion control data coming from an A-D converter.

In operation, a direct-memory-mapped I/O device fills up a buffer. When the buffer is full, the DMA device returns to the beginning of the buffer and starts to fill it again. The processor can address that new data in a circular fashion without having to recalculate the starting addresses of the buffer. It's the responsibility of the programmer to guarantee that the DMA controller is following in lock step behind the processor as it fills the memory.

Supporting circular addressing in hardware removes the burden of programming it from the software developer, who no longer needs to manage the process. Not all DSPs have the feature, however. The difference between a DSP that has circular addressing and one that doesn't is easily measurable, though. In the DSP with circular addressing, an inner loop takes four instructions to execute. Without it, a loop will take six—a 50 percent overhead. That means that the six-instruction DSP will run 50 percent slower than the four-instruction device.

A RISC processor, on the other hand, has to handle circular addressing through a test at the end of the loop. This makes the RISC device no better than the six-instruction DSP without circular addressing.

Key commands accessible via simple pull down menus will give you the best view of DSP development.

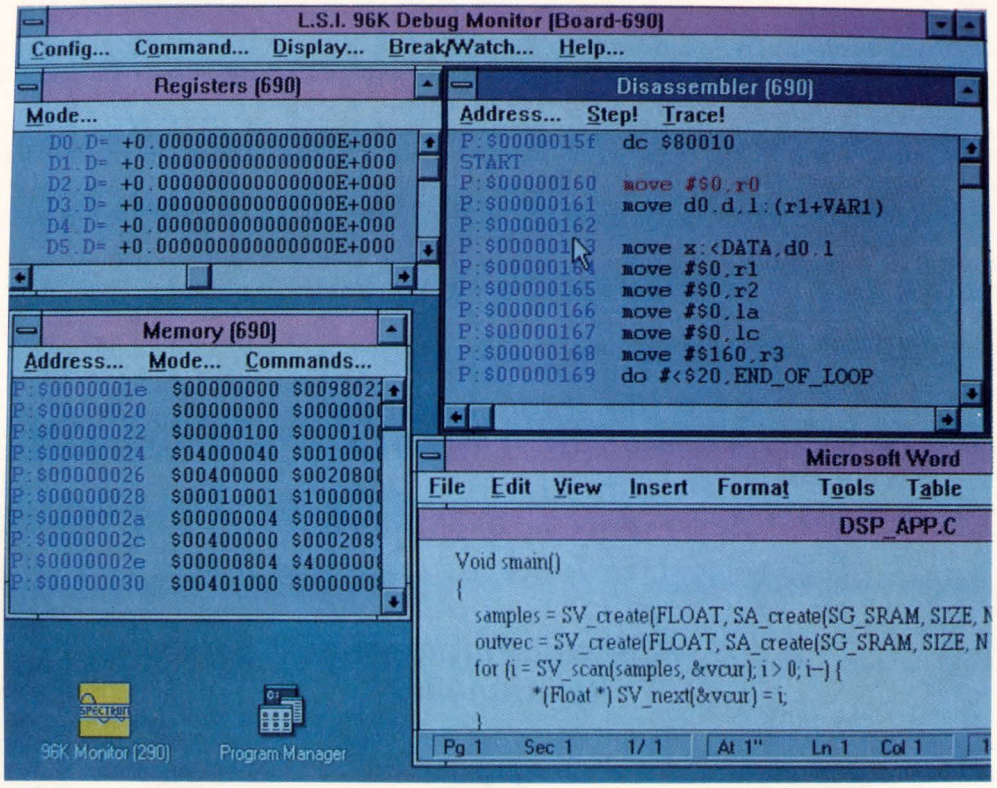

You can view registers, data memory and disassembled program memory simultaneously.

When using more than one DSP board, multiple instances of Views™ may be invoked.

Edit and view Source Code through programs such as Microsoft Word for Windows.

If you think DSP is hard to develop, maybe you should go Windows shopping.

When you shop around, you'll find only one DSP solution that includes a Microsoft Windows 3.0 debugger interface.

The one from Spectrum Signal Processing.

Spectrum is putting DSP to work making application development quicker and easier than you ever thought possible.

By taking DSP software development to the next level of productivity, Spectrum now gives you the simplest, most complete solution to your DSP applications.

Our View96™ and View5x™ development solutions include everything from Windows monitor debuggers, to a host of software tools including C compilers, assembler/linkers, libraries and operating systems.

Plus you'll find bus and I/O flexibility, on-board A to D and D to A conversion, our DSP-Link™ for peripherals and much more. All backed by the complete technical support you've come to expect from the leader in DSP technology.

So if you've been skeptical about putting DSP to work, call Spectrum now and find out why everyone is looking into our Windows: 1-800-663-8986 (Western U.S.), 1-800-323-1842 (Eastern U.S.) or 604-438-7266 (Canada).

Our software selection includes all the tools needed to get you up and running fast.

SPECTRUM

Putting DSP to work

© 1992. Spectrum, DSP-Link, View96 and View5x are trademarks of Spectrum Signal Processing, Inc. All other trademarks are trademarks of their respective holders.

CIRCLE NO. 46

Support for efficient hardware interrupt logic is another advantage Spectrum's Mitchell feels DSP devices have over RISC. Typically, the hardware implementation of a DSP provides several independent interrupt pins (with fast interrupt capability) and can be interrupted by the host communications port. Multiprocessor interrupt schemes including broadcast, narrowcast and semaphore can be constructed to minimize the overhead associated with interprocessor communication.

RISC devices, on the other hand, provide a single interrupt pin, and they're burdened with the tasks of servicing, prioritizing and handling interrupts, resulting in excessive context switch times—as high as 200 ms. Intermetrics' Tighe agrees. "Because of the large number of registers on some RISC machines," he says, "the speed of a context swap is slower than on a DSP."

Multiple internal DMA controllers are another salient feature of DSP devices that distinguish them from RISCs. Both the TMS320C40 and the DSP96002 have multiple on-chip DMA controllers to support overlapped data transfer and data processing, avoiding processor starvation during the data transfer. Furthermore, the DMA controller exercises the address generator units to support memory striding for accessing vectors or complex addressing schemes.

An edge in parallel processing

Spectrum's Mitchell feels that it's in the numerically intensive areas requiring multiple parallel processors that DSP devices have an edge over their RISC counterparts. The architectures of recent devices such as the TMS320C40 and DSP96002 encompass unique approaches to multiprocessing that RISC devices don't have. The TMS320C40, for example, is equipped with six high-speed communication ports for interprocessor communication in a range of topologies, including cubes, systolic arrays (matrix triangularization) and stars and rings. The DSP96002 has two fully independent external buses, with one dedicated to access to global resources and the other reserved for local memory processing, memory and I/O. Both devices also have serial ports for high-speed communication and control with a supervisory processor.

These processor architectures support efficient connection of multiple processing elements at both the board and system levels. Mitchell says, "The proliferation of boards targeted at the embedded systems market, based on multiple DSP devices, is testimony to the ease of use of multiple processor integration." A board with two DSP96002 processors operating at 40 MHz provides 120 MFlops of peak computing power and can effectively perform a 1K complex FFT in less than 0.5 ms. As a result, it successfully competes with a single-processor-based RISC board.

"We find that no off-the-shelf chip today addresses distributed computing requirements efficiently. The DSP features are not robust enough, and RISC chips lack certain functionality".

—Barry Isenstein, Mercury Computer

Barry Isenstein, product planning manager at Mercury Computer Systems (Lowell, MA), takes issue with Mitchell's belief that DSPs are better suited for parallel architectures than RISC. "We find that no off-the-shelf chip today addresses distributed computing requirements efficiently. The DSP features are not robust enough, and RISC chips lack certain functionality," he says. Isenstein feels that it makes more sense to provide added-value hardware to RISC to get access to its more general software aspects than it does to augment a DSP.

But even if one accepts Mitchell's argument that a DSP has an advantage in this area, that advantage may be short-lived. Embedded RISC vendors have also recognized the benefit of building their devices with hardware links specifically for parallelizable embedded applications. Jim Kearns, networking business development manager for Intel's i960 (Chandler, AZ), says he's seen customers interested in using his company's embedded RISC chip in such applications and he's actually working with such customers to define a next-generation architecture based on such a requirement.

DSP marries RISC

Given enough freedom, chip designers could build a RISC machine that performed a floating-point multiply in one cycle. So why don't they? The answer is that, faced with limited silicon real estate, they can't. They must instead choose which features to emphasize. RISC aficionados have chosen the data-driven, logic-oriented data flow architecture, whereas DSP designers have emphasized processing a continuous stream of data.

But will RISC or CISC ever meet DSP on a single piece of silicon? Motorola's Hillman observes that this has already happened. Indeed, numerous 16-bit microcontroller cores have already embraced limited DSP technology. On the 32-bit front, RISC processors are becoming more "DSP-like," and vice versa.

While RISC vendors have been putting more numeric capability into the chip, DSP builders have been adding more general-purpose capability into theirs. If a general-purpose RISC engine coupled with a powerful DSP coprocessor on the same piece of silicon is what you're looking for, however, you may have to be patient—unless, of course, you want to take predefined RISC and DSP cores from a vendor such as LSI Logic and do what many designers in the military have done, and build your own RISC/DSP solution. ■

For more information about the technologies, products or companies mentioned in this article, call or circle the appropriate

Analog Devices
(617) 329-4700 . . . . . **Circle 270**
Fujitsu Microelectronics
(408) 922-9000 . . . . . **Circle 271**
Intel
(800) 548-4725 . . . . . **Circle 272**
Intermetrics Microsystems Software
(617) 661-0072 . . . . . **Circle 273**
Mercury Computer Systems
(508) 458-3100 . . . . . **Circle 274**
Motorola
(512) 891-2000 . . . . . **Circle 275**
Spectrum Signal Processing
(604) 438-7266 . . . . . **Circle 276**
Texas Instruments
(800) 336-5236 ext. 700 . . . . . **Circle 277**
Timbre Engineering
(415) 488-4649 . . . . . **Circle 278**

RADISYS GIVES VXI THE INSIDE ADVANTAGE

EPConnect™ VXI Development Software
EPC-2 Controller
EPC-2AM PC Card Adapter Module
EPC-2e Controller (with Ethernet)

In just two C-size VXI slots, RadiSys Embedded PCs pack more power, versatility and flexibility:

- 16 or 20 MHz 80386 CPU
- 2, 4, 8 or 16 MByte DRAM
- 40, 100, or 200 MByte hard drive
- 3.5″ floppy drive
- 32-bit VXIbus interface, IEEE-488 (GPIB) interface and optional Ethernet interface
- 640 x 480 and 800 x 600 VGA graphics
- Complete VXIbus development and runtime software for the DOS/Windows, UNIX and OS/2 operating systems
- C, BASIC, PASCAL, FORTRAN and ADA programming language support*
- VXIbus Slot 0 and Resource Manager functions

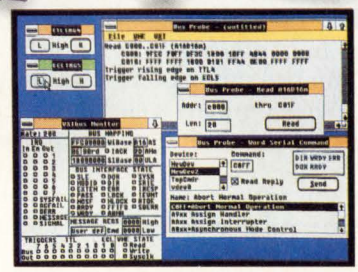

EPConnect software includes Configurator, Resource Manager, Message Delivery System, BusManager, BusMonitor, BusProbe, and Auxiliary Processor Interface.

Take the inside track. Plug in a RadiSys® EPC®-2. It's the preferred embedded controller for VXIbus. And no wonder. We've built nearly every Embedded PC installed in VXI systems worldwide. Since designing the first VXI embedded controller three years ago, we continue to develop and enhance them in partnership with major instrumentation manufacturers.

Inside more environments. Through the EPC's universal software interface, EPConnect, VXI systems are programmed with the most popular languages and run in more test development environments, including ATLAS, TekTMS, HP ITG, WaveTest, and LabWindows. And EPCs support DOS/Windows, UNIX and OS/2, with full access to their CASE tools, compilers and debuggers.

Inside vs. attached. The EPC-2 provides a direct 32-bit interface between the 386 and the VXIbus for optimum system throughput, leaving attached-PC alternatives out in the cold. EPCs lower overall system costs, too, eliminating the external interface baggage of attached PCs. And no attached PC can touch EPCs for compact size, fewer components, ruggedness and higher reliability.

The inside deal. A RadiSys EPC now costs under $6,000 including resource management software and VXI drivers. Call your VXI instrument supplier and ask for a demonstration of the RadiSys EPC-2 or EPC-2e. Or get "The Inside Advantage" direct from RadiSys at **800-950-0044.**

RadiSys® CORPORATION

THE INSIDE ADVANTAGE

RadiSys Corporation
19545 NW Von Neumann Dr.
Beaverton, OR 97006 USA
(800) 950-0044
(503) 690-1229
Fax (503) 690-1228

Copyright © 1991 RadiSys Corporation. All rights reserved. EPC and RadiSys are registered trademarks and EPConnect is a trademark of RadiSys Corporation. TekTMS is a trademark of Tektronix, Inc. HP ITG is a trademark of Hewlett-Packard Co. WaveTest is a registered trademark of Wavetek Corporation. LabWindows is a trademark of National Instruments Corp. Ethernet is a trademark of Xerox Corporation. 386 is a trademark of Intel Corp. OS/2 is a trademark of International Business Machines Corporation. UNIX is a trademark of AT&T in the USA and other countries.
*Languages supported vary with operating system.

CIRCLE NO. 47

70 Mbytes/sec...

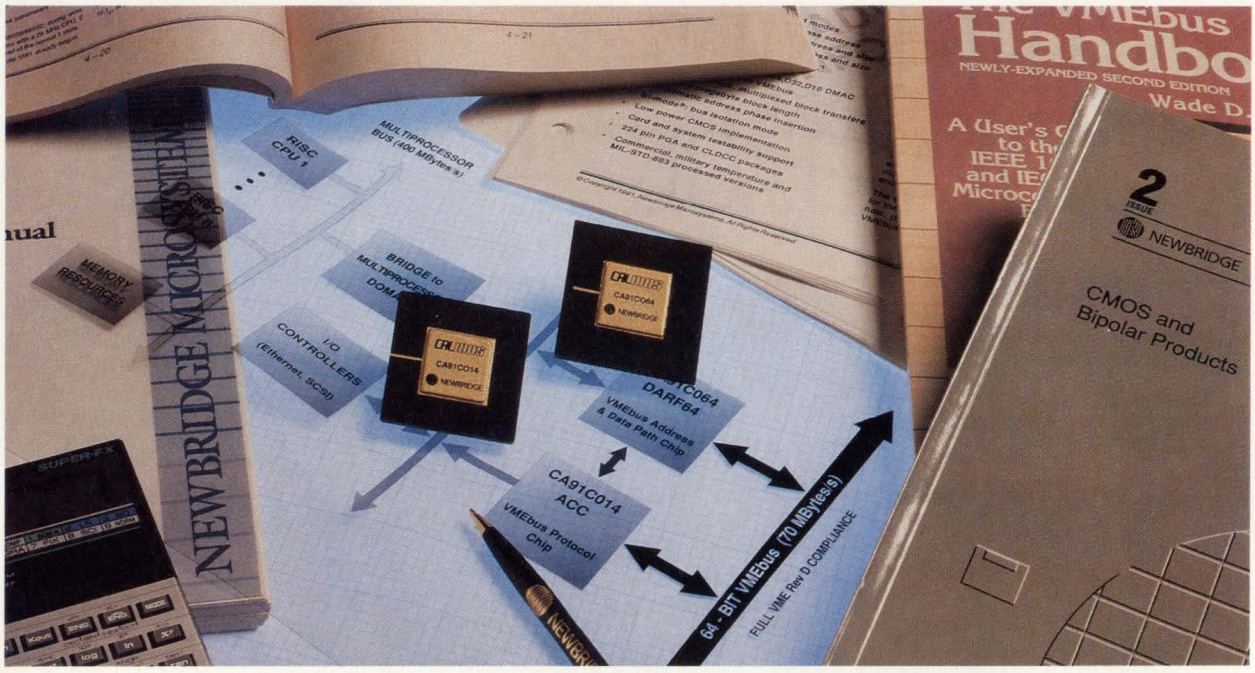

effortlessly.

"A second generation VMEbus interface chip set from Newbridge Microsystems (Kanata, Ontario) offers both a pin-for-pin compatible upgrade from the company's earlier chip set and full support for the proposed VME64 performance upgrade."

"Newbridge currently has the advantage, it would seem, in having incorporated the bells and whistles for distributed processing and multi-processing."

"It offers board makers and OEMs a shortcut to developing VME64-based boards, and a way to avoid the hassle of developing and testing a discrete interface."

"The Newbridge chip set is offered in a variety of packages, from economical plastic to full MIL-STD-883B screened versions. The CA91C064 DARF64 comes in a 224-pin grid, or in a PQFP with 1,000-lot prices starting at a little over $200."

Warren Andrews
Senior Editor, Computer Design

For high performance VMEbus interfaces give us a call at 1·800·267·7231.

NEWBRIDGE MICROSYSTEMS

Newbridge and logo and Calmos are trademarks of Newbridge Networks Corporation. © Copyright 1992, Newbridge Networks Corporation. All rights reserved.
The VMEbus Handbook was provided courtesy of VITA, Scottsdale AZ. Other brand and product names are trademarks or registered trademarks of their respective holders.
Quotations reprinted with permission from Computer Design (page 114-115) Jan. 1992, Copyright 1992.

CIRCLE NO. 48

Analyzers fine-tune high-performance buses

As standard buses increase in number and complexity, bus analyzers have progressed from simple logic analyzers to specialized, protocol-specific machines.

Warren Andrews
Senior Editor

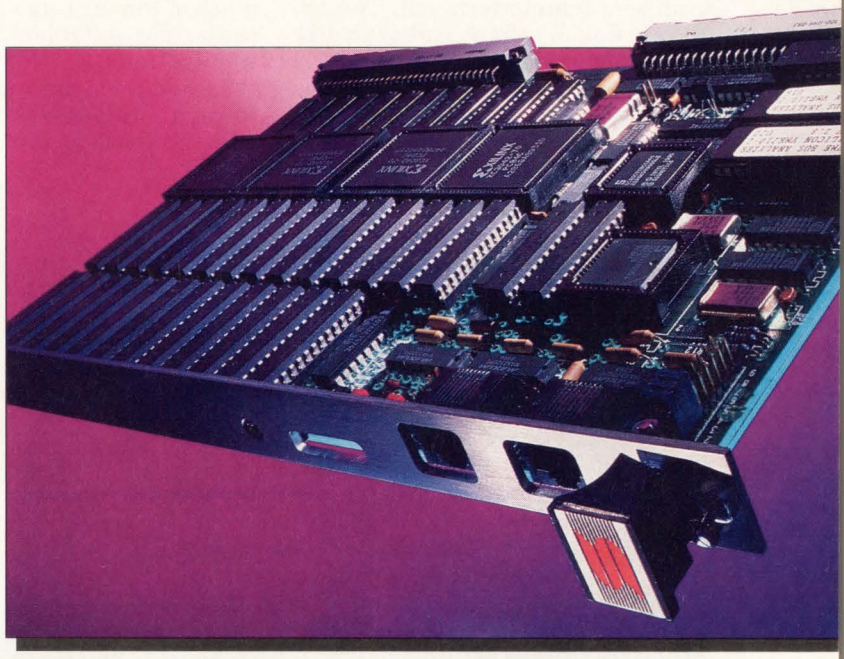

Silicon Control's VME210 VME analyzer provides most of the standard functions of a logic analyzer and works directly with the company's latest product, VMEWindow. This PC package provides a high-resolution interface with "point-and-click" mouse control.

A few years ago, the only tools available to the OEM system designer using a standard-bus architecture were conventional logic analyzers. Though these tools continue to improve and are effectively used at the board level, specialized, protocol-specific devices have almost become an essential ingredient in debugging and performance-tuning VME, Futurebus+ and Multibus II designs.

Such specialized analyzers usually take the form of single plug-in boards and are presently offered by a small handful of manufacturers. In addition to the single-board format, many board makers also offer mezzanine cards which provide added functionality. But despite the sophistication of current bus analyzers, there are some significant changes occurring.

Bus trends

Standard bus technology is on the move, looking to provide faster transfer rates, wider bus widths and higher performance. Also, time-to-market pressures are forcing OEM designers to minimize the time spent to debug systems and optimize performance. And as systems become increasingly complex, designers will want to simplify procedures—moving toward an intelligent front end, perhaps including a more complete graphical interface, consistent with many of the other design tools now available.

Analyzer makers have already started adapting to these trends, providing greater flexibility and utility—and even more changes can be expected in coming months. Increased hardware capability is being added, such as the capacity to analyze subbuses such as VSB, SCSI and proprietary P2 buses on VME and Multibus II. Other hardware features will include advanced timing analyzers, anomaly detectors and tracers. In addition, the newer analyzers are active participants on the bus. New hardware lets them 'tickle' the bus lines to look at how various parts of the system respond to critical bus-timing parameters and changes.

Bus-specific analysis tools are also getting friendlier. On one hand, they offer a clear path to testing and troubleshooting, while on the other, they provide a high level of flexibility. For example, hardware lets the analyzer

filter out unwanted transactions so you need to look only at the particular transactions causing a problem.

And at the design level, the tools let you focus on areas where performance may be in question. Major gains are being made in improving the front end and human interface. In some cases, this includes adding a personal computer as an intelligent host, rather than relying on the traditional dumb terminal. The addition of such a front end with disk storage capability lets you quickly set up patterns and exercises on the bus. It also provides a fast and easy way to look at results, especially with today's high-resolution screens. Still other products have the ability to store analyzed information for further review or postprocessing. In other cases, analyzers offer simple translations from bus-specific to more generic terms.

Finally, bus analyzers must address the latest modifications to standard buses, as well as the characteristics of emerging buses such as SBus and Futurebus+. VME changes that are beginning to emerge include those to be incorporated in Revision D—for example, VME64 and SSBLT—as well as other features expected to be grouped under "Recommended Practices (1014.2)." These will most probably include features such as hardware semaphores and multiplexed D32 and A40 on 3U VME.

On the Multibus side are many of the same kinds of changes that will have to be addressed by bus analyzers. These changes will include an increased transfer rate on the parallel systems backplane, which may include rates of 10, 16 and 20 MHz; a hot-swap capability; the ability to address more than the current 21-slot backplane; and the use of a higher-performance processor in place of the message-passing coprocessor for extremely high-traffic applications.

In addition, the rapid movement toward fault-tolerant computing and the need to provide hot-swap, live-insertion board capability will make new demands on bus analyzer makers. It's likely that in response to these demands, bus analyzers will be taking on new roles, performing constant system monitoring and diagnostics as well as analyzing performance for engineering development and debugging.

"And," says John Simpson, vice-president of VMetro (Houston, TX), "that's just the tip of the iceberg. What are now board-level analyzers will begin to take on a far larger role, analyzing not only a row of boards on a backplane but entire systems networked together."

John Simpson (left) of VMetro explains his company's Modular Analyzer System to a customer at this winter's Buscon West show. Simpson points out that the advantage of his company's product is that it lets you look at information in exactly the format it's specified in.

Bus sampler at heart

"At the heart of bus analyzers, whether VME, Multibus II, Futurebus+ or whatever," continues Simpson, "is some kind of a bus sampler including a bus interface, a qualifier or pattern-recognition circuitry and a trace memory. In addition, some kind of processor is needed to handle I/O functions to and from a terminal." Functionally, a bus analyzer is capable of looking at each transaction on a bus, determining what kind of transaction it is and deciding whether or not it conforms to the bus's specifications for that kind of transaction.

But providing timing analysis alone could easily be done by an advanced logic analyzer sampling each line on the bus. One of the principal things that differentiates bus-specific analyzers from conventional logic analyzers is their ability to perform a "state trace," where each transaction on the bus results in a single sample in the trace memory.

Protocol-sensitive sampling also provides other features which identify bus transaction types where there is no dedicated signal on the bus indicating that particular type of transaction. On VMEbus, for example, recognition of read-modify-write and block cycles depends on leaving the address strobe asserted between cycles. To capture this transaction, an analyzer must include internal circuitry to detect this condition and then create its own signal to store the information.

Still other information, such as the bus master level for each cycle in multiprocessor systems as well as an accurate identification of each cycle or transaction type, is often critical for system debugging. This type of information isn't available in conventional logic analyzers, and may not even be found on some bus analyzers.

It's what's up front

Being able to sit on a bus and capture and recognize bus transaction types is only the beginning of a bus analyzer's job. Getting that information to you is equally critical. Simpson points out that one of the advantages of his company's product is that it lets you look at information in exactly the format it's specified in. "A designer," he says, "can specify signals in a simple language he or she can understand, and the analyzer provides the necessary translation of the protocols. This lets you view the results the same way they were requested."

Seven Hot Reasons Why Your Next SBC Should Be a SPARCard

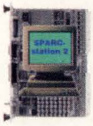

First

Our new SPARCard™ 2LC is the first and only VME SBC with SPARCstation 2™ performance. The 2LC has complete Sun compatibilty on-board to put your workstation software to work in demanding real-world target applications like robotics and voice recognition systems.

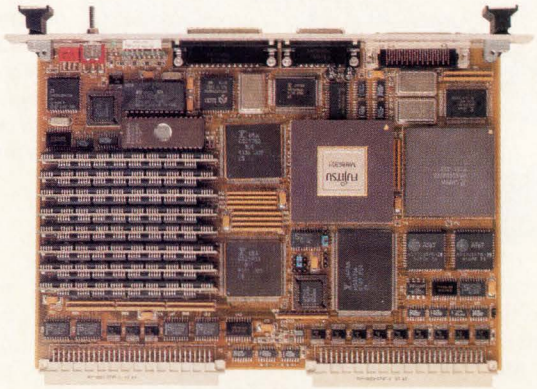

Real Time Choices

Choose between LynxOS, VxWorks and VADSworks to cover any real-time requirement. Wind River's VxWorks is the standard for embedded real-time applications. VADSworks speaks ADA. Or use LynxOS for a POSIX or UNIX compliant kernel in the target system. Any way you choose, Themis has the software for your real-time environment.

040 x 2 =28.5

With 28.5 Mips of integer performance, the 40 MHz SPARC MPU is more than twice as fast as the 68040. A larger 64KB cache and combined MPU/FPU put the SPARCcard 2LC in a class of its own.

64/32 64-Bit Architecture

The SPARCard is the only SPARC VME SBC built on the MBus, Sun's sophisticated 64-bit local bus. Used in the SPARCserver 600 MP Series, the MBus offers 64-bit paths to internal devices and memory for *fast* system throughput. Advanced ASICs designed specifically for the MBus guarantee performance and reliability.

The Future

The 2LC is just the beginning. Themis is producing an entire family of SPARC 2-based products for every embedded application, including those requiring SBus modules. For more information contact Themis Computer, the sole source for VME SPARCstation 2 performance.

SPARC 100% Pure SunOS

We're partnered with SunSoft, a division of Sun Microsystems and developer of Solaris, Sun's newest operating system. So we guarantee full SunOS compatibility including: SunOS 4.1.1, Open Look, Open Windows, and ONC/NFS. All Sun application software and development tools run on a SPARCard without modification.

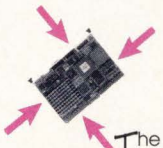

Fully Loaded

The SPARCard 2LC is a complete system on a card. Everything's on-board, including: 8-32 Mbytes DRAM, Ethernet, SCSI, VME interface, Open Boot PROM, Non-volatile SRAM, and two Serial Channels. Its MBus/VMEbus Controller has data FIFOs and hardware DMA for blazing VME performance.

THEMIS COMPUTER

THEMIS COMPUTER-Americas and Pacific Rim
6681 Owens Drive, Pleasanton, CA 94588
Phone (510) 734-0870, Fax (510) 734-0873

THEMIS COMPUTER-Rest of World
1, Rue Des Essarts, Z.I. De Mayencin,
38610 Gieres, France
Phone 33 1 76 59 60 61, Fax 33 1 76 63 00 30

SPARCard trademark pending approval by SPARC International.
All others are trademarks of their respective companies.

CIRCLE NO. 49

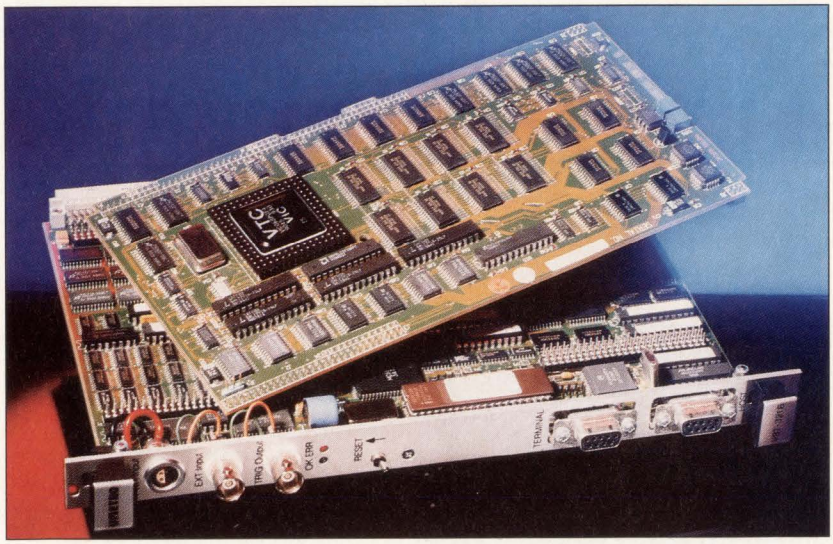

VMetro uses a piggyback approach for added flexibility. This add-on board provides both an expansion trace memory and a high-performance SCSI port. The larger trace memory increases memory depth, while the SCSI host adapter lets continuously sampled VMEbus data be dumped on to a disk.

"For example," Simpson continues, "an engineer can specify that he wants read cycles that occur at a specific bus level in a specific address range. The analyzer then goes off and configures all the control signals necessary to recognize that type of cycle and captures the data. When the data is presented, it's presented on the screen in the same context in which it was specified."

Without this capability, you're forced to actually do the analysis manually. "Without the analytical capability, such a product is simply a signal acquisition and display system, and it's up to the engineer to determine what's being displayed and whether or not it's useful," Simpson concludes.

Such a capability is becoming increasingly important, given such things as VME64 and SSBLT transfers. It should be possible, for example, to ask the analyzer to call up VME64 cycles or SSBLT cycles on VMEbus systems without knowing the bus-specific extensions.

The inevitable PC

While most analyzer products available today use a standard VT 100 terminal, the ubiquitous PC is starting to show up as an interface. PC interfaces have already emerged on at least two products—a Multibus II analyzer from ATM Computer GmbH (Konstanz, Germany) and a family of VME analyzers from Silicon Control (Deerfield, IL).

According to Laurie Burger, Silicon Control's director of sales and marketing, it's important to offer users more power on the front end. Users are getting increasingly creative in the way they're using bus analyzers, she says, and the type of information they want is even more sophisticated. "Particularly with the latest breed of analyzer that offers bus mastering capability, people are using them [analyzers] with embedded diagnostics. They're creating many of their own routines, and there have been a lot of requests by designers to enhance the front end," says Burger.

In response to this demand, Silicon Control has developed an IBM (and IBM-compatible) PC-based product to both simplify the front-end interface and provide increased flexibility. "VMEWindow is a stand-alone package that runs under DOS and is intended to offer users a windows-oriented approach to bus analyzers," Burger says. "In addition, it takes advantage of the high-resolution display available on the PC, but not available on the VT 100 terminal."

Burger goes on to point out that VMEWindow also lets you take advantage of disk storage capabilities for storing programs as well as data outputted from the analyzer. This feature provides a much friendlier interface, not unlike what some of the major logic analyzer companies such as H-P and Tektronix have done with their general-purpose products. The point-and-click approach lets you run through entire diagnostic routines more quickly and painlessly than having to manually enter things on a terminal keyboard.

The PC interface gives you even more flexibility when utilizing the included macro function. You can enter a series or sequence of diagnostic steps that includes anything that can be done on the analyzer. These steps can then be "played back" and the results stored on disk, printed out, displayed, and even manipulated in the processor to provide outputs like histograms. Steps can even include forcing the analyzer board into the 'master mode' and talking

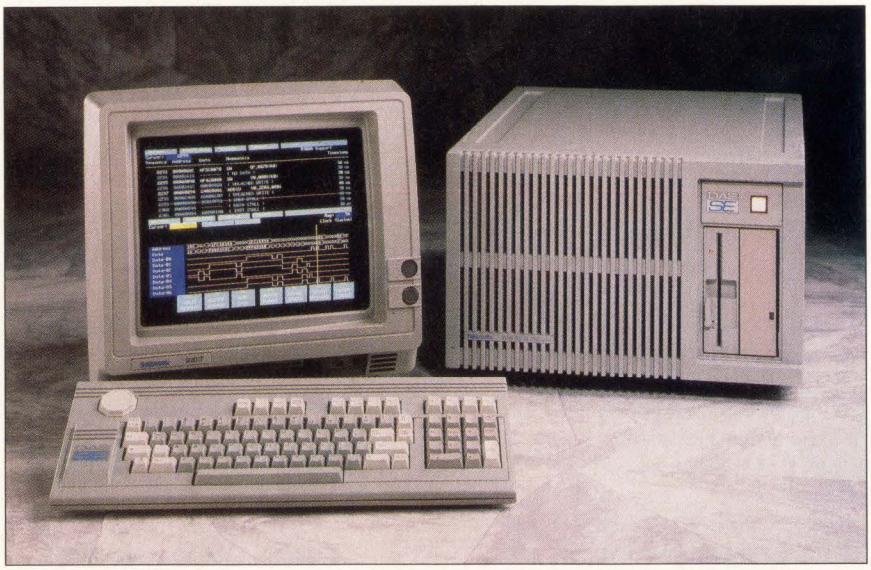

Inside this system resides Tektronix's DAS9200. The card, compatible with IEEE-896.2 profiles A, B and F, lets designers monitor high-speed timing or bus transactions on Futurebus+ systems.

32 MIPS 68040 & '030 • Starting At $2,695*

VME64 • VSB • Ethernet • SCSI • 256MB DRAM

Dual I/O Bus • 6 Serial Ports • 32 Parallel Lines • Centronics

Advanced Omnimodule • 1MB FPROM • 4 EPROM Sockets • 512KB SRAM

Single Slot • Watchdog • 8KB NVRAM • Snooping • Mailbox • Calendar Clock

Dynamic Bus Sizing • Designed for the Programmer • Software includes: VxWorks, UNIX, OS-9

Cross Code C, Freeform, VRTX, C EXECUTIVE, OMNIbug • 8 Status LEDS • 3-1-1-1 Burst Reads • 2-1-1-1 Burst Writes • 40MB/Sec '030 to '040 Transfers

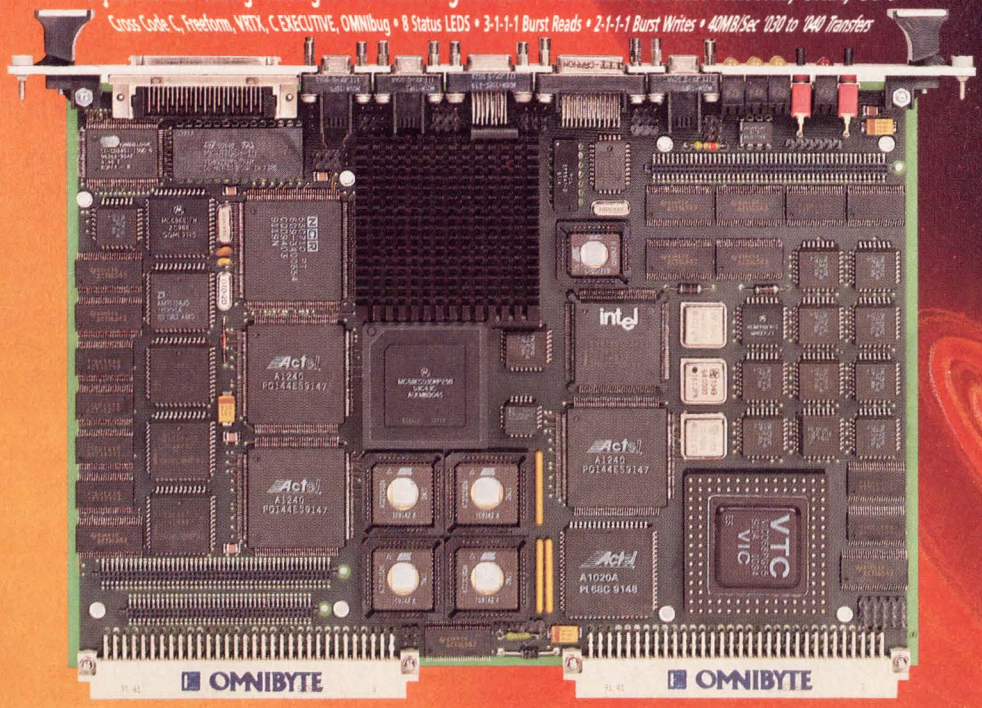

Taurus™: The New Standard for Performance, Value & Features

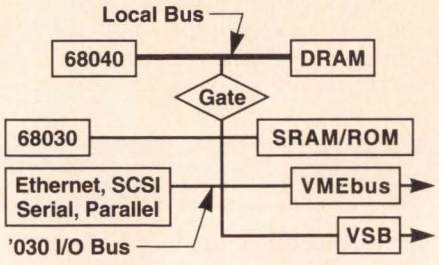

Taurus Dual Bus Architecture

The Taurus is a dual-processor and dual-bus single slot 6U VME board. Its dual-bus architecture allows the 68040 to execute code uninterrupted while the '030 processes all on-board I/O. This optimizes the 68040's performance. Writing code for the Taurus is very straightforward. The '030 can act as an I/O processor, communicating with the 68040 using control blocks. The '030 can also be used as a DMA controller with direct control of all on-board devices by the 68040. The '030 uses the 512KB of SRAM and the 128KB of code provided by Omnibyte in EPROM.

Performance	68040: 22 MIPS, '030: 10 MIPS, VME: 80MB/sec, VSB†: 80MB/sec
Intelligent I/O	Ethernet: i82596CA†, SCSI: NCR53C710†, 4 RS232D: CD2401
Standard I/O	2 RS232D: 68681 DUART, 32 Lines Parallel I/O, or 16 w/ Centronics Printer Port
Memory	4MB to 256MB† DRAM, 512KB SRAM, 8KB NVRAM, 1MB FPROM, 4MB EPROM
Other	VSB†, VME64†, Watchdog, Calendar Clock, Mailbox, (6) 16-bit Timers, Snooping, Advanced Omnimodule Socket
Software	VxWorks[1], OS-9[2], UNIX[3] VRTX[4], CrossCodeC, FreeForm[5] C EXECUTIVE[6], OMNIbug

† Denotes optional features.

Extensive use is made of intelligent, on-chip DMA devices for Ethernet, SCSI and serial I/O minimizing processor intervention. DRAM is contained on up to 2 stackable modules allowing upgradable options from 4-256MB.

Additional custom I/O is provided by several stackable Advanced Omnimodules. The Taurus can accept up to 1 of each type of module and still fit into a single slot.

To learn more contact Larry Snow:

800-638-5022

708-231-6880 in IL

© 1992 OMNIBYTE CORP.

OMNIBYTE™ Corporation, 245 W. Roosevelt Rd., West Chicago, IL 60185

Trademarks: 1-Wind River Systems, 2-Microwave Systems Corp., 3-AT&T, 4-Ready Systems, 5-Software Development Systems Inc., 6-JMI Software Consultants Inc., *Q1 price, some of the features above are optional & available at added cost i.e. Ethernet: +$125, SCSI: +$150

CIRCLE NO. 120

Analyzing Futurebus+

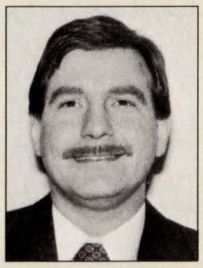

For Futurebus+ designers and system integrators, viewing bus transactions is critical to keeping the project on schedule. With the emergence of Futurebus+ as the industry-standard bus to succeed VME, new bus analyzer tools similar to those used on VME are starting to appear on the market. The differences between VMEbus analyzers and their Futurebus+ counterparts center mainly on the rigid electrical requirements adopted by the IEEE-896 Futurebus+ working group.

The electrical link

One key to analyzing a system that's failing is to view it in such a way that you don't add to or disguise the problem. The Futurebus+ standard requires a maximum stub length of 25 mm for each module in the backplane to guarantee incident wave switching for high performance. Connecting a logic analyzer probe directly to the backplane's BTL (backplane transceiver logic) signals violates this stub length requirement. Nimble-fingered debug engineers might manage to connect probes to the TTL side of the BTL bus transceivers on a module, but this will be difficult because of the 25- to 31-mil pitch of the transceivers.

Some VMEbus analyzer boards don't buffer the logic analyzer from the bus and violate not only stub length requirements but also capacitive loading. For Futurebus+, a reliable electrical link is needed from the backplane to the debug engineer's logic analyzer.

Bus analyzer boards

For VME there are two types of bus analyzer boards—single-board bus analyzers and bus analyzer boards that provide a connection from the backplane to your logic analyzer. A more general-purpose solution for analyzing a Futurebus+ system involves using a logic analyzer coupled with a Futurebus+ bus analyzer board. When not debugging the Futurebus+, the logic analyzer can be used for something else. (A single-board bus analyzer can't; it has a single, dedicated use.)

Bus analyzer boards—or preprocessors, as they're sometimes called—offer an easy and reliable connection to your backplane. The data from the backplane is latched at the appropriate time by latching BTL transceivers and is held until it's acquired by the logic analyzer. In this way it's *preprocessed* before it's acquired by the logic analyzer. To latch the data from the Futurebus+ correctly, protocol-sensitive clocking logic must be employed. Care must be taken when designing this logic, since the protocol is rich with exception cases. Signal integrity is also an issue, because a large number of signals will travel several inches from the transceivers before being acquired by the logic analyzer.

Not interfering with the system under test is a key feature of any good analysis tool. Futurebus+ requires all modules on the bus to handshake (assert and release) the bus strobe signals at certain times. A nonhandshaking preprocessor module must be designed with speed as a priority to ensure that no handshake will be missed. A bus analysis tool that handshakes the bus will affect the operation of the system under test and so make debug difficult. In addition, handshaking during the data phase on Futurebus+, even as a bystander, requires the participating modules to switch on glitch filters during this phase. This is prohibited in some Futurebus+ profiles, particularly Profile B, the I/O bus profile.

Bus disassembly

Once the data is acquired by the logic analyzer, it appears on the screen as lines and strings of 1s and 0s. A software disassembly package tailored to the preprocessor module and the logic analyzer can display bus activity in familiar bus syntax and abbreviations. For Futurebus+, all possible combinations of commands and transactions must be properly decoded on a phase-by-phase basis. This is complicated by the fact that certain signals mean different things in different phases. This requires that phase information be communicated between the preprocessor module and the logic analyzer on which the bus disassembly software runs.

A Futurebus+ preprocessor module coupled with a high-performance logic analyzer such as the new HP16550 will give Futurebus+ engineers both the timing and state analysis tools they need to solve the problems their systems are encountering.

Future+ Systems has codeveloped a Futurebus+ preprocessor module with Hewlett-Packard and is a member of HP's Preprocessor Connection.

Ed Aichenger, president, Future+ Systems, Westford, MA

to other boards in the system while looking for specific results. "This," says Burger, "is all without the necessity of writing any code or learning any fancy syntax."

"In addition," says Burger, "we've been racing to keep up with what's happening in the industry in terms of bus enhancements. Our just-released product, the VME310, incorporates support for a variety of features, such as VME64, not included in our earlier-version products. Unlike VMetro, which adds hardware to add new features, Silicon Control supports upgrades through software and firmware whenever possible. Since our boards incorporate a number of programmable gate arrays, such as the Xilinx devices, we're able to make a number of changes simply by changing a couple of PROMs to reprogram the programmable parts."

New products still needed

But that only works up to a point, she admits; then a new product has to be brought out to address major hardware improvements. "In our next-generation product," she continues, "we've added things like a separate trace buffer, a separate data collection mechanism for 200-MHz sampling and enhancement to the bus master capability to allow transmitting and receiving at the same time. This also lets the analyzer be a master on the P2 as well as on the standard VMEbus."

All of these features reside on a single board except for the P2 interface circuitry, which resides on a separate daughterboard. "This way," says Burger, "users that want

BUSCON '91 East Product of the Year.

Some people think we're power crazy, obsessed with quality and downright cheap.

And it's all true.

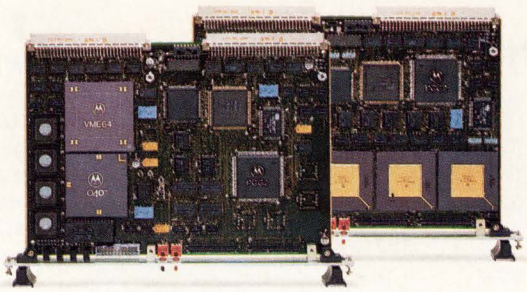

Okay, we admit it. We're hopeless perfectionists. Driven to produce the best boards at the best price.

Take our new MVME187 and MVME167 single board computers, for instance. Best in performance at 38 MIPS* (RISC) and 26 MIPS* (CISC), respectively. Best in value at $105 (RISC) and $154 (CISC) per MIPS.

Which comes out to just $3,995 per board. And which compares rather nicely with the $8,000 to $12,000 you could pay for the competition's boards. Based on the 88100 and 68040, Motorola's highest performance 32-bit microprocessors, the MVME187 and MVME167 employ VME D64 architecture. Boosting the VMEbus bandwidth to a full 40MB/s. And both boards come with four 32-bit timers, as well as SCSI and Ethernet connections.

Naturally, since they're made by Motorola, you can expect the best in quality. The Malcolm Baldrige Award is a powerful testament to that fact.

Malcolm Baldrige National Quality Award

So call 1-800-234-4VME for a free color brochure. And see why so many people are obsessed with Motorola products.

MOTOROLA
Computer Group

For a free color brochure, call 1-800-234-4VME.

*Motorola and the Ⓜ are registered trademarks of Motorola, Inc. © 1992 Motorola, Inc. All rights reserved. *MIPS rating based on Dhrystone 1.1 test results where 1757 Dhrystones/s is 1 MIPS (VAX™ 11/780). MIPS performance is based on the Diab 2.36E compiler.*

CIRCLE NO. 50

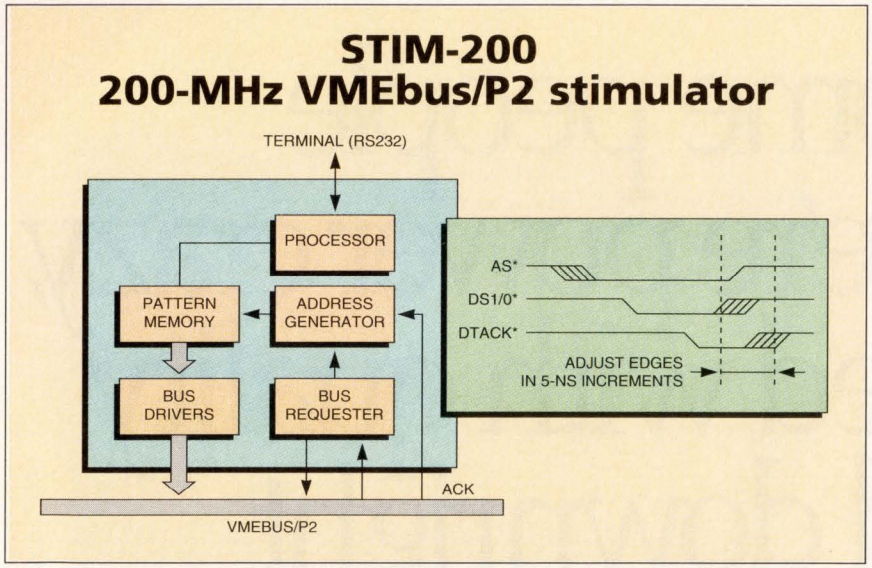

VMetro's latest add-on component for its VMEbus analyzer is its STIM-200, which is a VMEbus-specific pattern generator. Acting with the analyzer in a VMEbus master mode, it lets you shift around clock edges in 5-ns intervals to tightly control performance parameters.

to use VXI on the P2 or VSB, or their own proprietary bus simply have to plug in the right daughterboard."

For its part, VMetro offers a basic analyzer with a variety of plug-in modules to boost performance and add flexibility. Its Modular Analyzer System offers such features as an anomaly trigger designed to flag any direct violations of the VMEbus rules. Another module provides P2 capability, while still others offer a timing analyzer and memory expansion.

"One of our latest options," says VMetro's Simpson, "is support for a high-speed SCSI disk. This option lets engineers put together large audit files of system transactions for later review and more detailed analysis. There are a number of areas where this is of value, particularly in the isolation of hidden or obscure bugs, and in the analysis and monitoring of real-time systems."

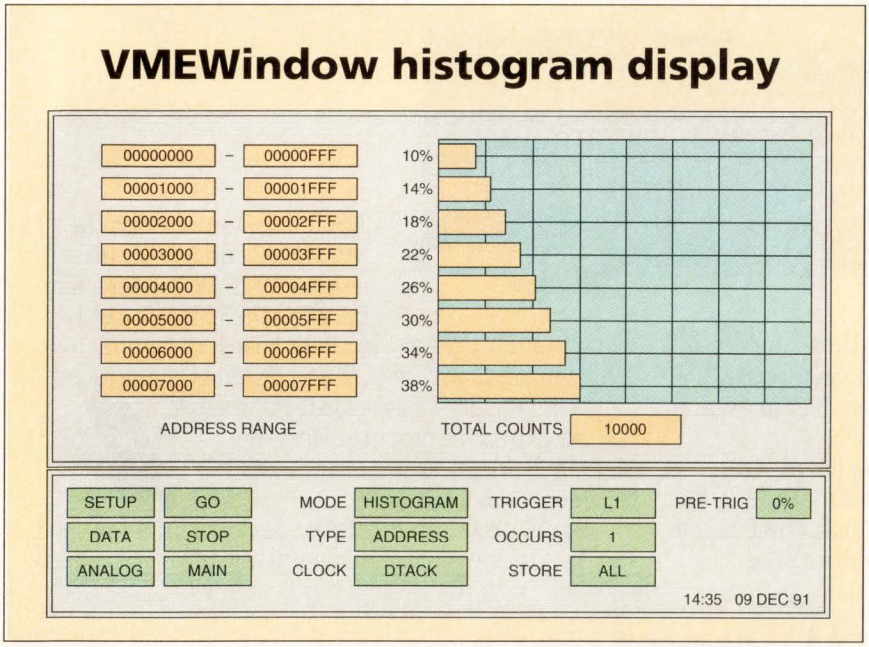

Silicon Control's VMEWindow uses an IBM-compatible PC instead of the traditional dumb terminal to let you look directly at a waveform generated by the bus. The "point-and-click" approach lets you quickly pull up the exact information required with a mouse rather than manually entering data on a keyboard. Using a PC as host interface also simplifies recording test sequences and results.

To the limit

One of the more significant add-on's VMetro offers is its VDRIVE-PB, which provides full capability to exercise the bus as if it were a master. Taking advantage of the VIC-068 VMEbus interface chip, it's able to precisely tailor timing to VMEbus timing specifications. "While Silicon Control may have been first to market with its product," says Simpson, "we believe our product offers true VME cycles. The Silicon Control product synthesizes cycles on the backplane that don't really look like true VME cycles. It tends to hold the bus after the signal is transmitted—a signal-oriented solution."

In comparison, VMetro's offering takes a cycle-oriented approach that works well for many tasks. "But," continues Simpson, "the first product didn't offer engineers the ability to test boards to the limit of the specification—and beyond. An add-on product, however—the STIM-200—provides the rest of the solution." STIM-200 is a bus-specific pattern generator which will let designers use standard templates for bus cycles, but lets them move the clock edges around in 5-ns increments to establish 'performance profiles' to determine specification compliance for specific boards.

The next generation

While the VMEbus community is gearing up and analyzer manufacturers are trying to stay with the pack, the Futurebus+ community is just starting to get off the ground. At least two Futurebus+-specific products are currently being offered, both as preprocessor front ends to more conventional logic analyzers. One comes from Tektronix (Beaverton, OR) and comprises a preprocessor extension board for the A, B and F profiles of Futurebus+. The back end ties into a Tektronix DAS9200 with the Centurion 92A96/D/XD acquisition modules.

The other product comes from start-up Future+ Systems (Westford, MA). Developed as part of Hewlett-Packard's "Preprocessor Connection" program, the board ties into an H-P logic analyzer. Both devices perform much the same task, providing the BTL transceiver logic, clocking logic and some decoding and latching logic to support the 896.1 protocol.

True portables are totally off the wall.

No outlets. No rechargers. No reliance on AC whatsoever. That's true portability. It's what the world is coming to. And it runs on easy-to-replace primary batteries. The next generation promises even smaller, lighter weight, more convenient portables. That depends on you, and you can depend on us.

Duracell is the primary source for primary power. We offer you a world of technical expertise and marketing experience in developing powerful solutions for computers, cellular phones and more. Let us help you select a primary battery system from our broad line that includes alkaline and high power lithium manganese dioxide batteries.

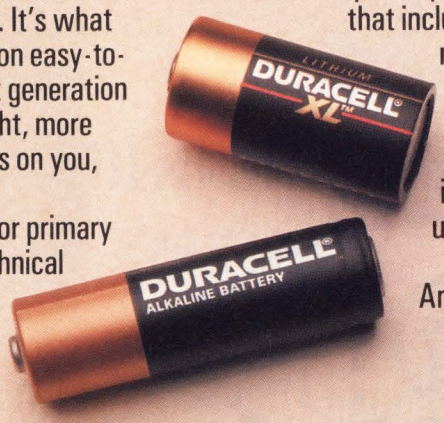

Call us for application-specific data, design-in assistance, or just more information. Our OEM hotline number is (800) 544-5454, Ext. 3281. Or fax us at (203) 791-3273.

True portability is the cutting edge. And it's in your power.

DURACELL®
PORTABILITY IS PRIMARY

CIRCLE NO. 51

"While the solution we offer now requires a separate logic analyzer," says Ed Aichenger, president of Future+ Systems, "it may provide a far more flexible solution as Futurebus+ systems progress through their performance iterations. For example, the BTL transceivers available at the time the product was first put together were National Semiconductor's second-generation Futurebus+ transceiver chips. The company introduced its third generation, the 'A' version, at Buscon in February. To step up to the next performance level, it's only necessary to change the preprocessor board. If the entire analyzer were part of a single-board solution, each iteration would call for an entire new analyzer. By separating the logic analyzer from the preprocessor, it's also possible to take advantage of the latest in analyzer technology—once again without starting from scratch."

The increasing complexity of analyzers permits them to look not only at straightforward bus cycles and transfer types but also to provide more detailed information on bus usage and efficiency. As more and more systems take multiprocessing approaches to increase processor performance, bus bandwidth and its utilization will become ever more critical.

The problem remains that we're not talking about one bus, though. The processor has its own internal bus, which interfaces with some kind of local or memory bus such as the Sparc version of M bus or the 88000 M bus, which may then interface with some other local, board-level bus, which in turn interfaces with the backplane system bus. And, to make things worse, that backplane bus is probably networked somehow with other similar—or different—buses into a larger system entity.

"Board-level analyzers will begin to take on a far larger role, analyzing entire systems networked together."

—John Simpson, VMetro

According to VMetro's Simpson, what's really needed is some type of bus analysis software and hardware that will let you look at, diagnose and analyze the system as a whole, from the top down or the bottom up. Simpson's dream, however, may be a while in becoming reality—each interface presents new obstacles as great or greater than those the backplane interface analyzer manufacturers have been wrestling with for the past several years. ■

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

ATM Computer GmbH (49)(07531)807-439	**Circle 290**
Future+ Systems (508) 392-9016	**Circle 291**
Silicon Control (708) 634-9313	**Circle 292**
Tektronix (800) 225-5434	**Circle 293**
VMetro (713) 266-6430	**Circle 294**

Is This The Language You're Using With Multibus II?

Not Anymore.

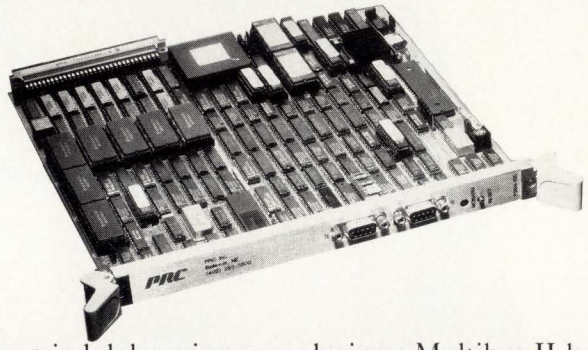

If you've tried debugging or analyzing a Multibus II-based system, you'll know why PRC created the PSBA-100 Parallel System Bus Analyzer. This easy-to-use productivity tool helps you make the most of Multibus II by showing you what's going on inside the computer – in a clear, simple format.

- ❑ New high-speed serial PC interface option
- ❑ Data capture and display based on Multibus II protocol
- ❑ Fully programmable filter logic with 16 trigger levels
- ❑ Supports all four Multibus II address spaces
- ❑ Stand-alone, terminal-based user interface
- ❑ Save captured data and capture set-ups

Wouldn't you like to deal with Multibus II in plain English? To receive further information and a free demo disk, call (402) 293-3900.

PRC Strength Through Understanding

PRC Inc. 1410 Wall Street, Bellevue, NE 68005
Multibus II is a trademark of Intel Corp.

CIRCLE NO. 52

It takes guts to design a great VMEbus system

GUTS

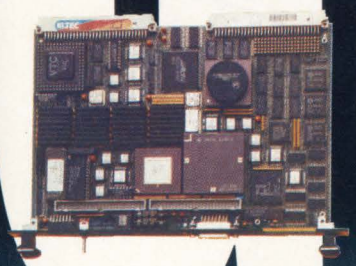

EUROCOM 7 CPU board; 25 MHz 68040

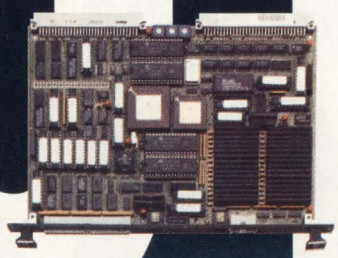

OPAC-200 Color Graphics Interface board

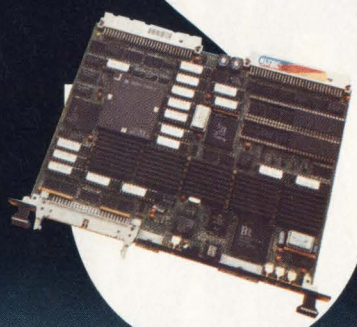

IC40 Image Processing Subsystem

and American Eltec is the ONLY SOURCE that offers a complete line of VMEbus boards: CPU, graphics, and image processing.

First, you create the idea for a system. Now you need to make it work.

You need to find all of the pieces. But why make it tough by using a lot of different sources? We have all of the pieces in your puzzle – just one source. And we are the only source with VMEbus CPU, graphics, I/O, and image processing boards, plus the necessary operating systems.

A single source makes your life easier, because you need to juggle fewer vendors, and you get better quantity discounts.

A single source gives you confidence in a successful application, because everything is compatible right from the start. You can take advantage of all special board features and capabilities. And you'll get better software and application support.

It is tough enough to design a complex and powerful VMEbus system. Buying compatible boards and operating software from one vendor makes your life easier, and gets your system functioning faster.

Call the one source for CPU, I/O, graphics, and image processing VMEbus boards: American Eltec.

We make your ideas work.

American ELTEC

USA American Eltec, Inc. • 4340 Stevens Creek Blvd., Ste. 204 • San Jose, CA 95129 • Tel. (408) 244-4700 • Fax (408) 244-5544
Germany Eltec, Elektronik GmbH • Galileo-Galilei-Straße 11 • D-6500 Mainz 42 • Tel. (0 61 31) 918-0 • Fax (0 61 31) 91 81 99
France Eltec International SARL • 1, Allée des Garays F-91120 Palaiseau • Tel. (1) 64 47 18 77 • Fax 64 47 09 33
U.K. Eltec International PLC • 45 Caldecotte Lake Drive • Milton Keynes, MK7 8LF • Tel. (09 08) 36 64 99 • Fax (09 08) 27 46 00

CIRCLE NO. 53

2ND ANNUAL ANALOG & MIXED-SIGNAL DESIGN CONFERENCE.

HYATT REGENCY
San Francisco Airport

OCTOBER 28-30, 1992

BURLINGAME
CALIFORNIA

As a design engineer, higher speeds and greater complexity of electronics systems are making your job more challenging. And with clock speeds climbing to 50MHz and beyond, analog characteristics are becoming major considerations for digital design.

According to Technology Research Group, 40% of all ASICs will be mixed analog and digital by 1994. With this kind of demand, today's designer must be informed and prepared. Finally, there's a source of objective and up-to-date information on mixed signal design–The *Analog and Mixed-Signal Design Conference*. This technical conference, sponsored by *Computer Design* Magazine, features over 50 lectures and workshops 100% dedicated to your needs. Nowhere else will you find so much information, technology and expertise under one roof.

HIGH QUALITY WORKSHOPS

Respected engineers will offer high-quality information that you can immediately use to complete your design projects more quickly, creatively, and elegantly. You'll learn about analog and mixed-signal simulators, different design techniques, analog implications of high-speed systems, and much more in workshops such as:

- Analog Effects in High-Speed Design
- Mixed Analog/Digital Design Modeling, Simulation, and Test
- Selecting the Optimum Analog Devices
- Digitizing Analog Functions
- Solving Convergence Problems in SPICE
- Configuring Mixed-Signal with ASICs

And much, much more.

PRACTICAL HANDS-ON INSTRUCTION

For years you've been subjected to press releases from vendors under the guise of objective discussions or exposed to academic papers that had little bearing on your design problems. The *Analog & Mixed-Signal Design Conference* is different. You'll find in-depth, practical and real-world information from industry experts who have been there. They'll give you knowledge you can bring to the workstation tomorrow!

WHO SHOULD ATTEND?

The *Analog & Mixed-Signal Design Conference* is dedicated to the unique needs of mixed-signal design. If you're a digital designer addressing the problems of high-speed designs, a digital designer involved with analog, or an engineering manager making decisions about CAE/CAD tools you must attend this conference.

QUESTION AUTHORITY

The *Analog & Mixed-Signal Design Conference* Vendor Products Exhibition features booths crammed with the latest in analog and mixed-signal design tools. As questions arise during the workshops, you can go directly to the vendor authorities on the exhibit floor for answers. They will be on hand to let you try out and compare products, give you a clear grasp of the practical tools available, and tell you where to look for new avenues of information exchange.

As an attendee at the *Analog & Mixed-Signal Design Conference* you'll learn how to improve your design skills, how to increase productivity, and where to turn in order to ensure design success. All in three days under one roof.

Join your colleagues this October in Burlingame. The *Analog & Mixed-Signal Design Conference* promises to be one of the most significant industry events of the year!

For more information call Tom Loftus at (415) 905-2354 or Fax (415) 905-2220.

Sponsored by Miller Freeman Inc.
P.O. Box 7843
San Francisco, CA 94120-7843

ANALOG & MIXED SIGNAL
DESIGN CONFERENCE

Tell me more about the *Analog & Mixed-Signal Design Conference.*
I'm interested in: ❑ Attending ❑ Exhibiting

NAME _____ TITLE _____

COMPANY _____

STREET _____

CITY _____ STATE _____ ZIP _____

PHONE _____ FAX _____

Mail or Fax to:
Analog & Mixed-Signal Design Conference
P.O. Box 7843, San Francisco, CA 94120-7843
Phone (415) 905-2354 / Fax (415) 905-2220

OCTOBER 28-30, 1992
Hyatt Regency
San Francisco Airport
Burlingame, California

CDAS1

FPGAs race for the gold in product development

Barbara Tuck
Senior Editor

Gate arrays and standard-cell ASICs are still the big winners when it comes to product development, but FPGAs are closing in with competitive benefits such as design flexibility, reduced cost and quicker time-to-market.

In these tough competitive times, system designers are racing to build cost-effective products and get them out the door before the next guy does. Just as designers are turning to hardware description languages and synthesis as design aids on the tool side, so are they turning in increasing numbers to FPGAs on the silicon side to help them meet tight schedules and to stay one step ahead of their end users and competitors. Wherever FPGAs can meet performance, density and functionality requirements, designers are looking to the programmable devices for at least a partial solution.

Designs weren't frozen

The need to respond to evolving specifications throughout the development cycle has made designers at Silicon Engines (Palo Alto, CA) turn to FPGAs. They used FPGAs to build the random logic portion of their new multichannel spectrum analyzer, which is key to NASA's Search for Extraterrestrial Intelligence (SETI) program. The spectrum analyzer, which looks at a 10-MHz bandwidth and analyzes down to a 1-Hz resolution, will be pointed at the heavens to search for intelligent radio signals. Richard Povey, senior digital designer at Silicon Engines, reports that the overall design requirement for the spectrum analyzer was evolving while the engineers were implementing the design.

To build the spectrum analyzer, Silicon Engines' designers had to do two full-custom chips—a digital signal processing engine and a very large synchronous memory module—and had to rely on 5-ns PALs from Texas Instruments (Dallas, TX) for clock multiplexers. But wherever they could use FPGAs to soak up functionality and to accommodate design changes, they did. "For much of the circuitry, we didn't have the option to prespecify," says Povey, "so we couldn't go to a gate array."

The designers started with a lot of large building blocks and then figured out how to interface those blocks. "Our approach was simply to make a design that fulfills the requirements—to use a particular piece for a specific function because we know it works," explains Povey. "With synthesis, we would have wound up asking what out there fits well with our synthesized design, and perhaps would not have found a good fit."

Povey says that the spectrum analyzer board he worked on has more than two dozen Altera (San Jose, CA) MAX 5192 and 5130 EPLDs, with probably 20 distinct designs among them, and five Xilinx (San Jose, CA) 3090 logic cell array (LCA) FPGAs, with three distinct designs. He chose Altera's MAX parts for address decoding over Xilinx components because there's a potential problem at power-up when Xilinx LCAs, or any other SRAM-based FPGAs, are used for this purpose. "The Altera devices are very logic-rich," says Povey. In contrast, the Xilinx LCAs are flip-flop-rich, making them very good for data multiplexing.

Changes over product lifetime

A three-month deadline and the likelihood that design changes would be required over the product's lifespan drove Scientific Atlanta (Norcross, GA) to use Xilinx LCAs for a piece of cable TV equipment that provides CD-quality digital music, permitting cable users to receive simulcast stereo sound with national TV channels.

Staff engineer David Burleson reports that three LCAs went into one system box and seven into another. "There are no PALs at all in the design, but we had to hand-route

some runs that are as high as 34 MHz," he says. In addition to Xilinx tools, the Scientific Atlanta designers used the Susie simulator from Aldec (Newberry, CA). "If we hadn't simulated everything, we'd still be debugging today," Burleson warns fellow users. Since the cable equipment is not a high-volume product, converting to an ASIC is not an issue for Scientific Atlanta.

Conversion an issue

For Bob Hackett, a Xilinx user and engineering manager in charge of VLSI designs at Integrated Network (Bridgewater, NJ), conversion to a full ASIC is very often an issue. "If time-to-market is the most important consideration," says Hackett, "I use Xilinx up front and then convert to a gate array if the volume's going to be more than a couple of thousand a year."

To Hackett, the cost consideration is more than the up-front NRE (nonrecurring engineering expense) compared to FPGA costs. "We have to consider the total cost of the product in the end. We can't afford to have a $20 part in a $100 product. Although FPGA vendors sometimes claim that FPGAs will save you money even when volumes hit the

tens of thousands, it's not true." Having been in the gate array business, Integrated Network has expertise at converting FPGAs to ASICs and can do it without difficulty, according to Hackett.

Hackett's designers choose a silicon implementation in the very beginning of the design cycle. Even when they know they'll be converting from Xilinx LCAs to a gate array or standard cell, they concentrate on getting the logic to work in the LCA, rather than beginning with the ASIC design and then proving it out in an FPGA. "If the design requirements involve more than 4,000 or 5,000 real gates," says Hackett, "we won't be able to get it into a Xilinx LCA and will have to go to an ASIC."

At what volume does it make sense to go with a standard-cell implementation? "It's not necessarily true that there's a volume number that can dictate your choice. Just get all the quotes in from the foundries," suggests Hackett, "and examine them for the cheapest path."

Though Hackett decides on a silicon implementation first, he goes through almost the entire design cycle before deciding on a silicon vendor. An HDL and synthesis fan, he uses the Boolean-equation-based Bool HDL and MIS II gate-level synthesis developed at the University of California at Berkeley.

Hackett's first synthesis step is pretty much generic. He reads data from several vendors' data books into the MIS II synthesis tool and synthesizes the design to see if the logic will operate correctly. Later, the design will be mapped into the technology that best meets the timing specifications of the relatively high-speed telecom devices Integrated Network builds.

Just make it work

Basically, Hackett's design philosophy boils down to: "Put your design in a black box and make it work. There's nothing magic; just make the logic design work and put it in a package."

Sandi Habinc is another proponent of the black-box approach. A design engineer at Saab Ericsson Space (Gothenburg, Sweden), he reports that a Saab design team decided they couldn't prototype with FPGAs after making a system-level black-box model of the Columbus project, an embedded processor board for the Freedom space station.

"We have better control over PLDs than FPGAs," explains Habinc. Speed was also a determining factor. "With a system clock of 40 MHz, it was obvious what to use."

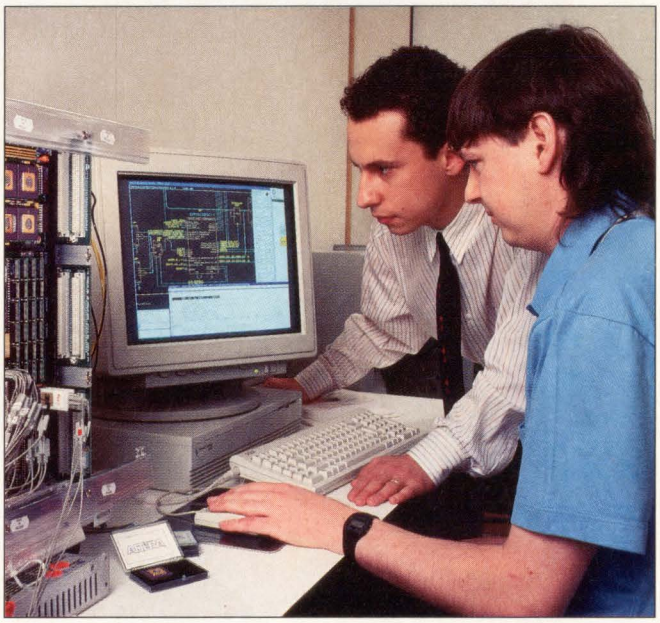

Silicon Engines' design engineers Richard Povey (right) *and Adrian Jeday* (left) *used Altera EPLDs and Max+Plus development tools to implement the random logic portion of their new multichannel spectrum analyzer, the heart of NASA's Search for Extraterrestrial Intelligence (SETI) program. The designers are running Altera's MAX+Plus software on a Sun workstation for hardware description language compilation, logic synthesis and design fitting.*

The Saab team used 40 Advanced Micro Devices (Sunnyvale, CA) 15-ns 22V10s to handle CPU cycles, I/Os and such. Since neither PLDs nor FPGAs are valid for space use, and since gate arrays are not classified by European space agencies, Saab is now in the process of converting the 40 22V10s into two cell-based ASICs—with 30 PLDs going into one ASIC and 10 into the other.

The Columbus project has been Saab's introduction to VHDL. Having previously used Abel from Data I/O (Redmond, WA) for PLD designs, Saab designers this time chose Hint, the VHDL option in the LOG/iC synthesis tool from Isdata (Monterey, CA), to describe and compile their PLDs. Hint can synthesize finite-state machines and combinational logic as well as structural descriptions. The designers then further synthesized their PLD-based design with LOG/iC, which does Boolean minimization and produces JEDEC files. For simulation, the designers had to move from a personal computer to a workstation environment, where they simulated one PLD at a time using QuickSim from Mentor Graphics (Wilsonville, OR).

To simulate the ASIC in VHDL, Saab designers are trying to move all of their VHDL code for theColumbus project from Hint to Model Technology's (Beaverton, OR) V-System VHDL simulator. They'll use it on PCs for design capture and some system-level simulation.

For most of the system-level simulation, the designers intend to move their VHDL code—which is "hopefully portable," remarks Habinc—to Mentor's VHDL-based System 1076. "When you go to an ASIC, you want to see how much of the code and the design work you can reuse," he says. "You don't want to lose code and work and money going from one tool to another. We're trying to figure out how portable Mentor is. We hope we're not doing the beta testing now."

The Saab designers, continuing up the synthesis ladder, are now planning to replace LOG/iC with either Synopsys (Mountain View, CA) synthesis or Mentor's AutoLogic VHDL Release 8.0. "Much of the problem is that you have to stick with what AutoLogic can do when you're going from a high-level description to where you can synthesize to code," says Habinc.

In the future, Saab designers will try to use FPGAs, and their engineers are currently involved in a project where they're comparing different FPGA technologies.

FPGAs meet space constraints

Billy Beckworth, president of systems-design firm Beckworth Enterprises (Mesquite, TX), has already made the switch from heavy reliance on 22V10s to the adoption of FPGAs. Beckworth discovered he

16 lines, No waiting.

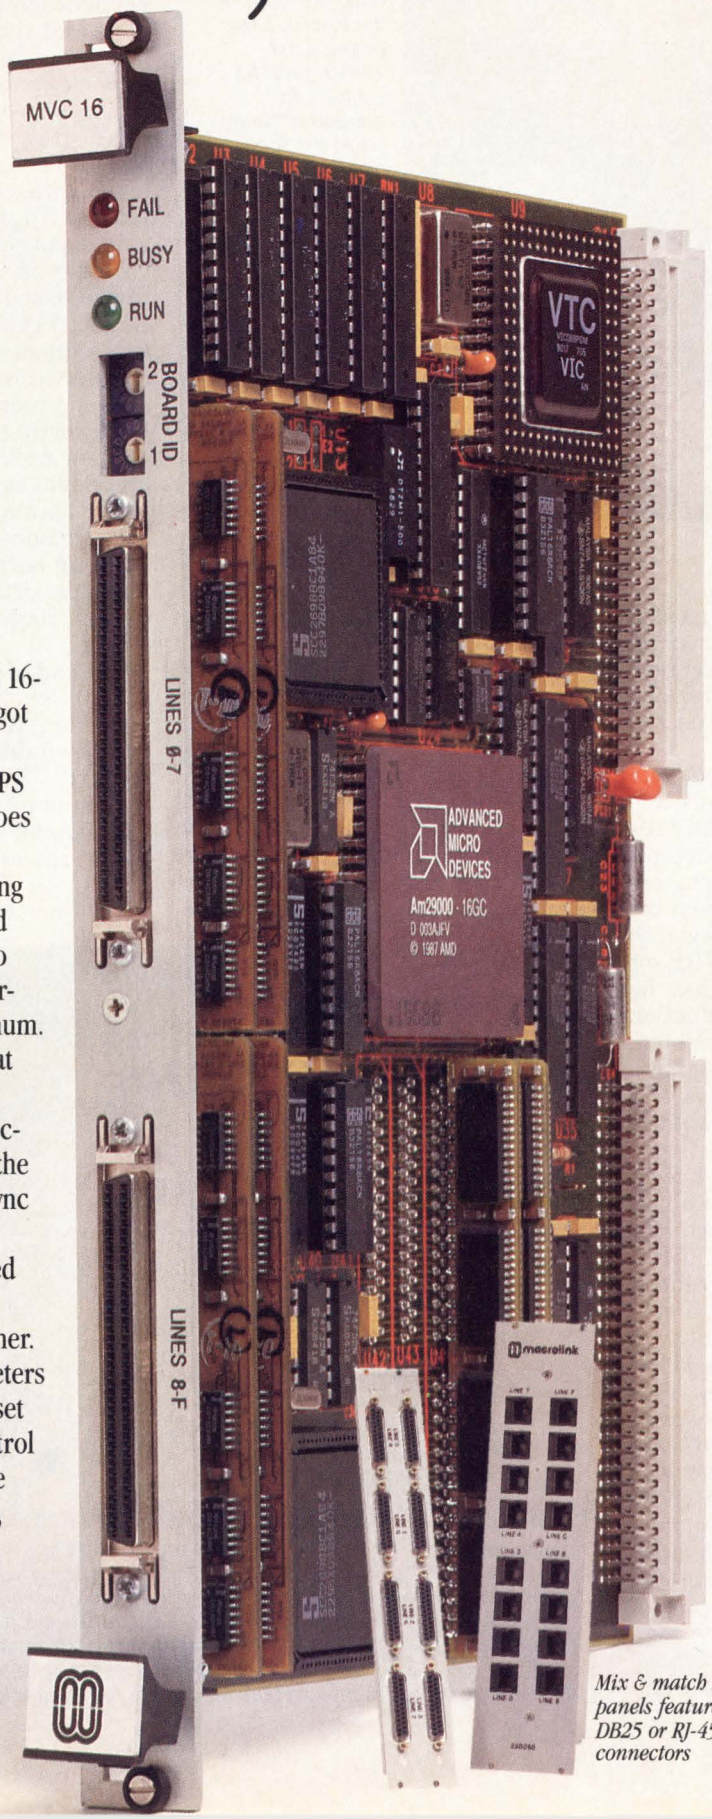

Mix & match I/O panels feature DB25 or RJ-45 connectors

Your application is first in line with the MVC 16-line Async Commux. It's got processing capacity to spare, thanks to a 16 MIPS RISC, so system power goes to your users—not I/O.

Character processing and buffering is managed on-board by our RISC, so driver calls and host overhead are kept to a minimum. All 16 lines can operate at 38.4 KBaud. That's over 61,000 characters per second throughput, double the rate of other VMEBus async controllers.

The MVC's advanced features benefit both the integrator and programmer. Port and VMEbus parameters are soft-configured and set line-by-line. Modem control is standard. Full software support is also included, along with diagnostics and a Streams driver.

Advanced memory architecture and high-speed buffering eliminate overruns and port domination. Memory is expandable from 128KB to 1MB, so the MVC handles today's requirements and future application needs.

With over 11 years experience producing advanced storage, communications and memory products, Macrolink delivers the powerful and flexible MVC in 8 and 16 line configurations. Call us today. We won't keep you waiting.

Macrolink Inc.,
1500 North Kellogg Drive,
Anaheim, California 92807.
Phone (714) 777-8800,
FAX (714) 777-8807.

macrolink® inc.

CIRCLE NO. 71

INSTANT DATA ACCESS (IDA)
DIAL (617) 494-8338 DOCUMENT NO. 1037

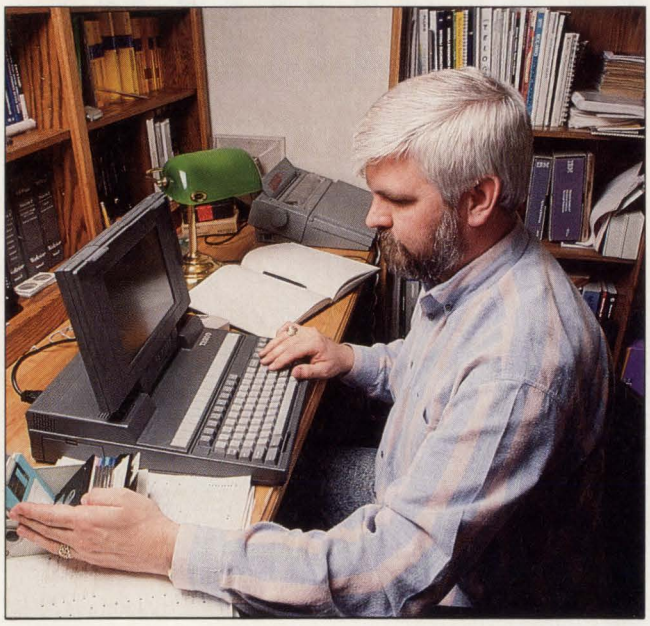

Billy Beckworth, Actel user and president of systems-design firm Beckworth Enterprises, used 22V10s until he couldn't fit enough of them on to the footprint for a single-slot SBus interface board he had to build. Beckworth used a pair of Actel's 2,000-gate A1020 FPGAs and later redesigned the board for extra functionality with Actel's new 8,000-gate A1280 FPGA.

had to go beyond 22V10s when faced with design requirements for a second-generation SBus interface board that would let a mainframe laser printer interface to a Sun Microsystems Sparcstation.

Even if he weren't limited to a 15-in.^{2} footprint, the cumulative propagation delay of the 22V10s would have made them unacceptable, he says. Although a custom gate array would have satisfied his real-estate concerns, Beckworth couldn't afford the up-front NRE for a masked device. On top of that, he didn't anticipate volume production since the board is a captive design. "You get locked in when you go to a full-blown gate array," he says. "I'd really have to have a hot design to do that."

After comparing the FPGA technologies of Xilinx and Actel (Sunnyvale, CA), Beckworth chose Actel for its security features and finer granularity. As the only vendor making SBus channel drivers, Beckworth needed to protect his design. "With an external SRAM to hold the device configuration program, as is the case with Xilinx, the risk of competitors mapping out the SRAM and determining the circuit design was too high," he comments. He chose Actel because the chip is completely programmed internally with antifuses, and it provides security elements that can be blown to protect the design by preventing the map from being read. He also felt that the finer granularity of the Actel architecture would allow a more flexible design approach.

Beckworth used schematic entry software from Viewlogic Systems (Marlboro, MA) and implemented his design in a pair of 2,000-gate A1020 FPGAs; he put the SBus interface and a FIFO memory interface into one A1020, and the logic to handle the printer channel interface into the other. Because he partitioned the board so the two FPGAs shared no common circuitry other than power and ground, Beckworth was able to drive the devices with separate asynchronous clocks, the first at 25 MHz and the second at 5 MHz.

A year later, Beckworth's customer requested a modified version of the board with a parallel input port plus two serial synchronous/asynchronous ports to be combined with the original channel controller logic, all on the same single-slot SBus form factor. He converted from the A1020-based design to one based on an Actel 8,000-gate A1280 FPGA. Although this time the requisite drivers and receivers for the various interfaces had to reside on an external interface board, Beckworth reports,"The arrival of the new A1280 with its 140 user I/Os and 8,000 equivalent gates was the key to making the design possible at all." He adds that the A1280's dual clocking networks made the device ideal, particularly since the original design required two independent clocks.

Instead of recompiling the A1020 design for the A1280, Beckworth took the A1020 building blocks and enhanced them for the A1280 architecture. "The conversion process was a snap," he says.

The Actel A1280 FPGA also came to the rescue of defense contractor Interstate Electronics (Anaheim, CA), which had barely gotten a 25-

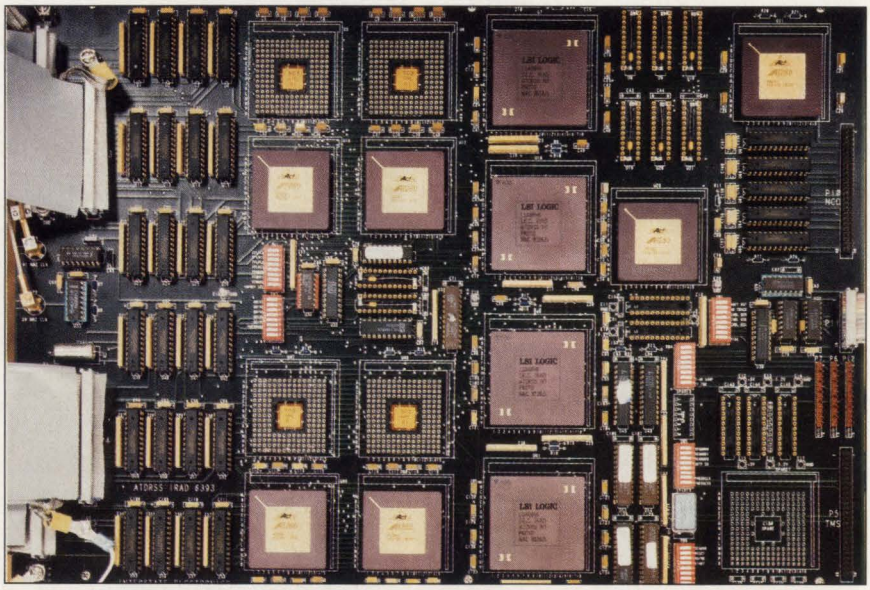

Defense contractor Interstate Electronics turned to Actel's 8,000-gate A1280 FPGA to implement quickly a 650-Mbit/s modem for NASA. Going with either an all-GaAs solution or all gate arrays would have been too expensive. Instead, Interstate used a parallel-processing solution with discrete GaAs for speed, a parallel data path of eight A1280s (four A1280 sockets open in photo) to store and distribute data, an LSI Logic gate array (used four times) to synchronize the sample data from the FPGAs, followed by another pair of A1280s for sample recombination.

Our GAL22V10-7 is Super Fast.

At 7.5ns, Lattice provides the world's fastest field-programmable 22V10. Based on our high-performance E²CMOS® technology, the **GAL22V10-7** combines 111 MHz F_{max} with extremely low power consumption. It even supports industry standard loads and pinouts. What's more, it's 100% tested to guarantee perfect programming yields, less board rework and fewer system failures.

For free samples and a GAL® Data Book, call **1-800-FASTGAL** and ask for information packet 107.

Lattice

Leader in E²CMOS PLDs.™

Lattice Semiconductor Corp.
5555 Northeast Moore Court
Hillsboro, Oregon 97124

Copyright © 1992, Lattice Semiconductor Corporation. E²CMOS and GAL are registered trademarks of Lattice Semiconductor Corporation.

CIRCLE NO. 56

Low Delays.

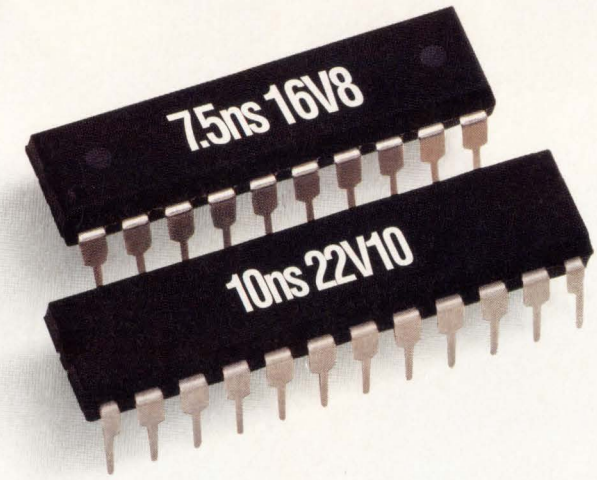

High-Speed 7.5ns CMOS PAL® Devices.

There's nothing we hate more than delays. That's why we developed high speed CMOS PAL devices that no one can beat—our CMOS 7.5ns 16V8H-7 and 10ns 22V10H-10 PAL devices.

In fact, nobody even comes close to our in-system performance, with the fastest set-up and clock-to-out times available. Both come in PLCC and DIP varieties. All on state-of-the-art submicron EE CMOS.

High-Volume, High-Speed Delivery.

Again, there's nothing we hate more than delays. You can get huge volumes of our new CMOS PAL devices now.

And they're on the shelf at your local dis-

No Delays.

tributor, too. So you can get the quantity and speed you need, whenever you need them.

What more can you expect from the company that sells more programmable logic than all of its competitors combined?

So pick up the phone and place your order today, or call **1-800-222-9323** for more information.

Because at AMD, we don't believe in long delays either.

Advanced Micro Devices

901 Thompson Place, P.O. Box 3453, Sunnyvale, CA 94088 © 1991 Advanced Micro Devices, Inc.
PAL is a registered trademark of Advanced Micro Devices. All brand or product names mentioned are trademarks or registered trademarks of their respective holders.

CIRCLE NO. 57

Mbit/s modem into production for NASA's second relay satellite subsystem ground terminal when it received word that NASA would be needing a 650-Mbit/s version of the device. The opportunity to win the NASA project motivated the Interstate designers to locate a solution to the speed problem at once.

If the designers opted for a parallel processing solution using discrete GaAs parts for speed at the input, they could then go with Actel FPGAs for a parallel data path. "Whatever logic lent itself to FPGAs, we put into FPGAs," says staff engineer Matt Thompson.

GaAs too expensive

An all-GaAs solution would have been too expensive, though. And, as for implementing the entire design in gate arrays, the designers had already spent considerable NRE dollars for a 60,000-gate custom gate array that housed high-speed multipliers plus surrounding logic. Though a single NRE charge actually covered four ASICs (the gate array would be used four times over in the design to synchronize the sample data received from the FPGAs), incurring a second NRE was rejected.

Also, it would have taken too many PALs or even higher-complexity PLDs, according to Interstate, to implement the array of 32 shift registers, each one 32 × 3 bits wide, needed to store sample data and distribute it to the bank of four custom gate arrays. By combining the outputs of the shift registers, Interstate could produce 16 channels, each containing 32 consecutive samples, and each running at 40 MHz.

The company used eight Actel A1280s for the array of shift registers and another pair at the back end of the system for sample recombination. Designers chose the Actel A1280 because of the speed requirements of the system and also because the Actel fuse map uses a nonvolatile technology. Designing with Xilinx SRAM-based LCAs would have required downloading all fuse map information each time the system was powered up. Interstate designers were afraid, according to Thompson, that with all the high-speed digital switching on their board, the slightest noise glitch might cause the Xilinx device to be configured incorrectly and malfunction.

Though they didn't have time to ramp up on VHDL or synthesis for either modem, Thompson says Interstate designers will use VHDL on their next project. He reports that synthesis might have alleviated the necessity to redo schematics to get the design information into a format that could be used to program the Actel FPGAs.

Racing competitors to market

The most significant benefit of prototyping with FPGAs is time-to-market, according to Dirk Hayden, senior project manager at Larse (Santa Clara, CA), a maker of data communication boards for high-speed digital networks. "It's not the risk of going through two sets of masks," says Hayden, "it's the time."

Acting as a beta site for Concurrent Logic (Sunnyvale, CA) FPGAs at present, Hayden's company looks for devices that are supported by Viewlogic on PCs, offer second sources and are flip-flop- and latch-intensive rather than being speed-intensive. Hayden eventually targets masked devices but initially uses FPGAs for any applications-oriented circuitry and PALs for high-speed microprocessor decoding.

Has migrating from FPGAs to gate arrays caused problems? "Yes," he says, "we've had problems, but of our own doing." Larse has gone with Viewlogic to alleviate some of those problems and relies on the conver-

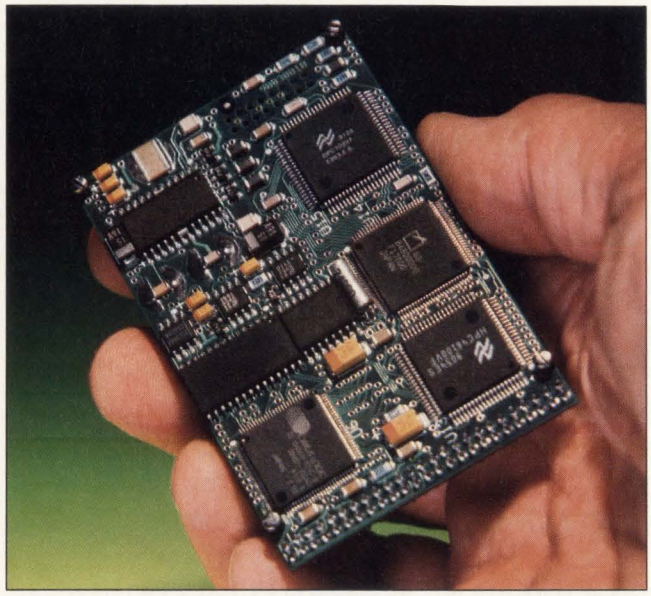

Stephen Cowen (foreground), VLSI development engineer at Integral Peripherals, and design engineer John Mount test an ASIC prototype of a 1.8-in., 20-Mbyte drive. Integral Peripherals developed its drive using half a dozen Xilinx 3042 FPGAs, and then migrated the design to an S-MOS gate array (middle chip on right side of pc board shown to the right).

National Semiconductor

CLOCK GENERATION AND SUPPORT

Minimize clock skew.

With a minimal loss of sanity.

Designing a system clock for a high-performance application can be the ultimate exercise in frustration.

Unless you begin by calling National. Our new family of advanced CGS™ (Clock Generation and Support) products includes bipolar, CMOS and ECL clock drivers tested and guaranteed for minimum skew. They meet the requirements of high-end microprocessors: duty cycle, skew, and edge rates.

MINIMUM GUARANTEED SKEW

And they're supported with the development tools you need to simplify your end-system design: SPICE models and a CGS design handbook that includes derating curves.

All of which means you can actually minimize skew while optimizing system speed and performance.

Without losing your wits.

FEATURES

- Typical pin-to-pin skew:
  - —ECL: 50-150ps
  - —CMOS: 100-200ps
  - —Bipolar: 200-400ps
- Extensive development tool support
- Future CGS products in development

Protect your sanity. Call National today.

For more information on our CGS2525, CGS2526 and 100115 clock drivers, call 1-800-NAT-SEMI, Ext. 177.

National Semiconductor

NORTH AMERICA: P.O. Box 7643, Mt. Prospect, IL 60056-7643 (Tel: 1 800 628 7364, ext. 177; Fax: 1 800 888 5113); **EUROPE**: Industriestraße 10, D-8080 Fürstenfeldbruck, Germany (Tel: 49 8141 103 0; Fax: 49 8141 103 515); **HONG KONG**: 15th Floor, Straight Block, Ocean Centre, 5 Canton Rd., Tsimshatsui, Hong Kong (Tel: 852 737 1654; Fax: 852 736 9921); **JAPAN**: Sanseido Building 5F, 4-15-3, Nishi-shinjuku, Shinjuku-ku, Tokyo, Japan 160 (Tel: 81 3 3299 7001; Fax: 81 3 3299 7000).

CGS is a trademark of National Semiconductor Corporation. © 1992 National Semiconductor Corporation.

sion services of American Microsystems Inc (AMI; Pocatello, ID).

Designers at Intégral Peripherals (Boulder, CO) also targeted a gate array for what they're claiming is the first 1.8-in. disk drive, but they chose to prototype and test the immature design in Xilinx FPGAs.

VLSI development engineer Stephen Cowen says, "We chose reprogrammable Xilinx parts because the design was so young and would be changing, and it would be interfacing to chips that would be changing. One-time-programmable Actel parts would be too expensive."

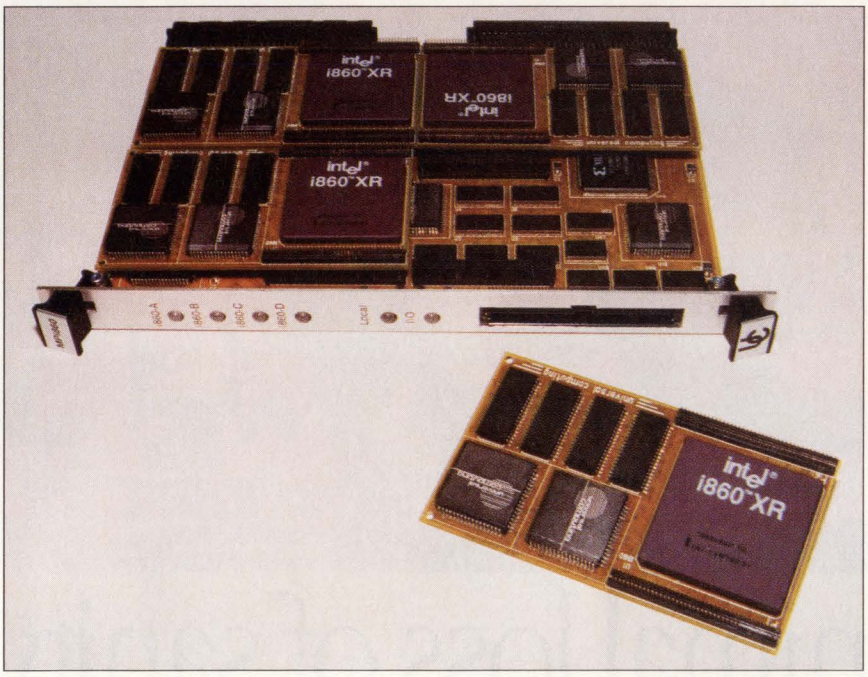

Designed for today's RISC- and CISC-based systems, a pair of QuickLogic QL8x12 FPGAs (shown here with logo of end-user Universal Computing [San Diego, CA]) are used to implement the memory and bus interface for a graphics and numerics accelerator based on Intel's 40-MHz i860XR. Last month, QuickLogic introduced its QL12x16 pASIC. With 2,000 usable gates, the new FPGA doubles the density of its predecessor.

Hierarchical approach

Since the designers knew up front they'd have to go to an ASIC to satisfy the footprint they were constrained to, they took a hierarchical approach and used Viewlogic tools to draw a huge schematic to represent the ASIC they were targeting. Within that schematic, they drew blocks using Xilinx library elements. Using Viewlogic tools, the designers broke up functionality into encode/decode for the read/write channel, servo sequencer and glue logic. Since there was lots of code involved, Cowen said it was very important to get hardware into the software developers' hands quickly, so top priority was assigned to whatever the programmers needed to be functional. Designers wound up fitting the disk-drive functionality into about half a dozen Xilinx 3042s. "Since cramming logic into the Xilinx parts means routing iterations, we kept the circuitry sparse," says Cowen.

The designers asked gate array foundry S-MOS Systems (San Jose, CA) to recreate the drawings with its own library elements. Because they had used a hierarchical drawing, Intégral's designers were able to enter the S-MOS library elements under the Xilinx library on Viewlogic instead of generating a new schematic.

Why did Intégral Peripherals choose a gate array over a cell-based ASIC? "Because we had an entirely new architecture, designed from scratch and very immature," says Cowen. "It will be a miracle if we don't have to return this design one more time, maybe in half a year. When the design matures and the architecture solidifies, we'll go to a standard-cell implementation."

Like system designers at Intégral Peripherals, Eastman Kodak (Rochester, NY) designers use Xilinx LCAs as early engineering prototypes when they're designing something they're unsure of for their high-end printers, reports senior engineer Greg Huber. Then the standard procedure is to migrate to a gate array. "In general, Xilinx parts are too expensive for production," he says. "Xilinx reprogrammability on-the-fly is a useful feature, though, when you have a number of different applications and interfaces. Then it's worth the extra dollars," he adds.

The Kodak designers wrote a Xilinx library for Synopsys but have been unsuccessful in getting the tool to synthesize their logic. "It's very wasteful," says Huber. "It seems to have a mindset of its own and doesn't use the options available."

The company's just installed the ArchSyn behavioral synthesizer, exclusively offered by Dazix/Intergraph (Huntsville, AL), and based on research and development at AT&T Bell Laboratories (Berkeley Heights, NJ). As soon as they're equipped with the necessary tools, the Kodak designers will write a library for Xilinx to run on ArchSyn.

In addition to using Xilinx parts for displaying and rendering functions, they use Altera's MAX EPLDs for smaller functional blocks and state machines.

Room for FPGAs and ASICs

A surprising number of projects wind up in both FPGA and ASIC implementations during a product's life cycle, says Mark Beal, vice-president of engineering at ASIC design house Nebula (Westboro, MA). "We always ask if we can use an FPGA as part of the design process. And if so, if it can be used long-term," says Beal. "Most end users intend to convert FPGA-based designs to an ASIC, but many never do."

According to Beal, the reasons his customers choose FPGAs include the possibility of changing a design downstream, flexibility for a design that may not be frozen and ease of prototyping to accommodate software developers.

"Sensitivity to the NRE costs for ASICs is also a factor," says Beal, "but to say there's no NRE with an FPGA is something of a fallacy. The cost differential between an FPGA and an NRE for an ASIC can actually be taken up in engineering."

The time element for a fully characterized FPGA that's been designed reliably is about the same as for an ASIC, according to Beal. "That's because it takes more time

Economics forces innovation in reusing existing designs

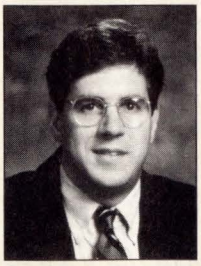

It could be argued that the saying, "Time waits for no man," was written to describe the dilemma of the electronic system manufacturer of the 1990s. Economic realities have hit this formerly explosive industry, forcing manufacturers to modify existing business models.

The most significant challenges for system manufacturers today are meeting time-to-market and cost constraints while product life cycles decrease and product complexity increases. Competition is fierce. Everyone on a new design's critical path is affected.

The introduction of FPGAs in the mid-1980s caused the first structural shift in the design cycle for system manufacturers. Turnaround time for ASIC prototypes can range from one to six months. By contrast, FPGA design implementations typically average one to six weeks. FPGAs offer additional advantages:

- No nonrecurring engineering expense (NREs)
- Total design cycle control, and
- Standard products.

By combining user programmability and logic capacity approaching that of "hard" ASICs, FPGA technology has become the strategic weapon for system manufacturers.

But that's only part of the story. The benefits of FPGAs inherently address the "back end" implementation phase of the design cycle. FPGAs offer tremendous assets to help reduce physical design steps. Further improvements to the back end of the design cycle will only offer incremental time-to-market leverage.

A second structural shift in the design cycle for system manufacturers comes from shrinks in the "front-end" design entry phase. The combination of powerful logic synthesis tools and FPGA technology is now becoming the strategic weapon of choice for electronic system manufacturers, both small and large.

Melding logic synthesis and FPGAs is being driven by these economic forces:

- A single FPGA can replace up to 50 PLDs,
- FPGA performance and cost can vary widely for a specific design, depending on the architecture, and
- Hard ASICs still offer significant cost reductions on a per-part basis for volume production.

While these economic factors are compelling, the redesign time required can eliminate the benefits. System manufacturers dread a missed market window. The ability to redesign quickly, automatically and without sacrificing performance is the imperative of the 1990s. Logic synthesis specifically designed for FPGAs makes possible these benefits.

No design entry revolution happens overnight. There are more than 100,000 discrete logic and PLD designers today. These designers don't want a new methodology. They want to solve today's design problems with "reuse." This means a logic synthesis tool must read JEDEC and Palasm—standard languages of PLD designers.

Today's tools must not only "synthesize," but optimize for the target device, using many of the same techniques you would utilize if redesigning the chip by hand.

At least 20,000 turnkey FPGA design packages from companies such as Altera and Xilinx are now in place. Designers using one or the other of these packages must be able to freely move between them. FPGA-specific synthesis is the enabling technology for this type of design flexibility—logic can be retargeted and reconfigured to best fit the device.

FPGA-specific synthesis will become a design mandate as FPGA designs reach 10,000-gate complexities. It was at this point in the ASIC design methodology where benefits of a hardware description language's synthesis entry methods were understood and accepted.

HDL entry of FPGAs requires sophisticated and technology-specific optimization and mapping. The solution is FPGA-specific logic synthesis and optimization tools. Logic synthesis software automatically translates high-level logic descriptions of digital circuits to the gate level. Logic optimization software minimizes a circuit to reduce silicon area, increases speed and makes trade-offs based on user-supplied constraints.

Exemplar's philosophy is to support standard interfaces and entry formats, supply a wide range of FPGA libraries (coupled with architecture-specific optimization algorithms) and focus on the details of reimplementing logic in target devices which have dissimilar library elements.

As FPGA design becomes as commonplace as ASIC design, tools will be developed to meet the FPGA designers' needs, while achieving business and economic objectives of electronic system manufacturers.

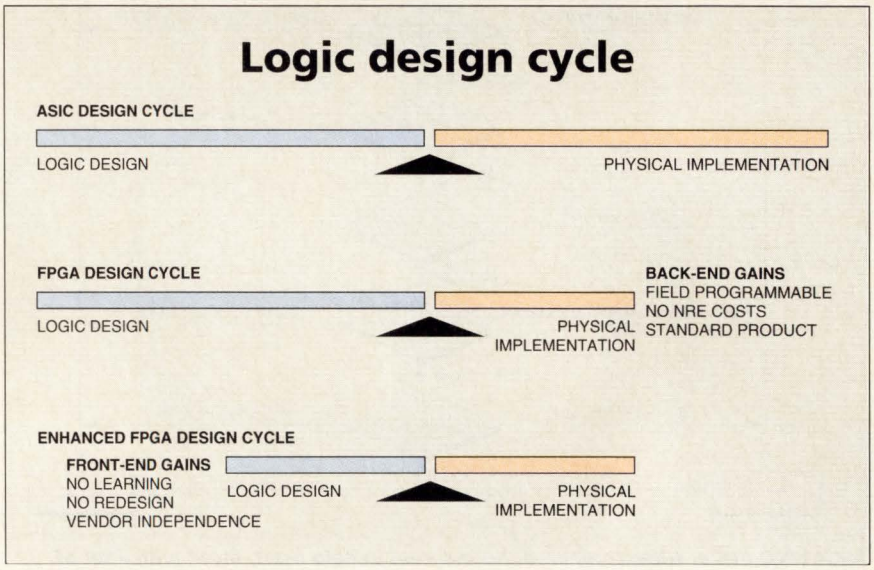

Ronald P. Ranauro, vice-president of marketing and sales, Exemplar Logic, Berkeley, CA

to meet timing constraints. FPGAs are so much slower, you have to design them more carefully."

What, then, is the point of FPGAs? "The real advantage to FPGAs," says Beal, "is the ability to change designs quickly. The equivalent of one NRE may be taken up by an FPGA, but three or four NREs for an ASIC would be prohibitive." When deciding between FPGAs and gate arrays for a silicon implementation, Beal says, "It's not so much an either/or situation but a both/and situation, where you might prototype with an FPGA and later migrate to an ASIC, or begin with an ASIC and validate your design in an FPGA. But for production you want an ASIC. The difference in price is enormous."

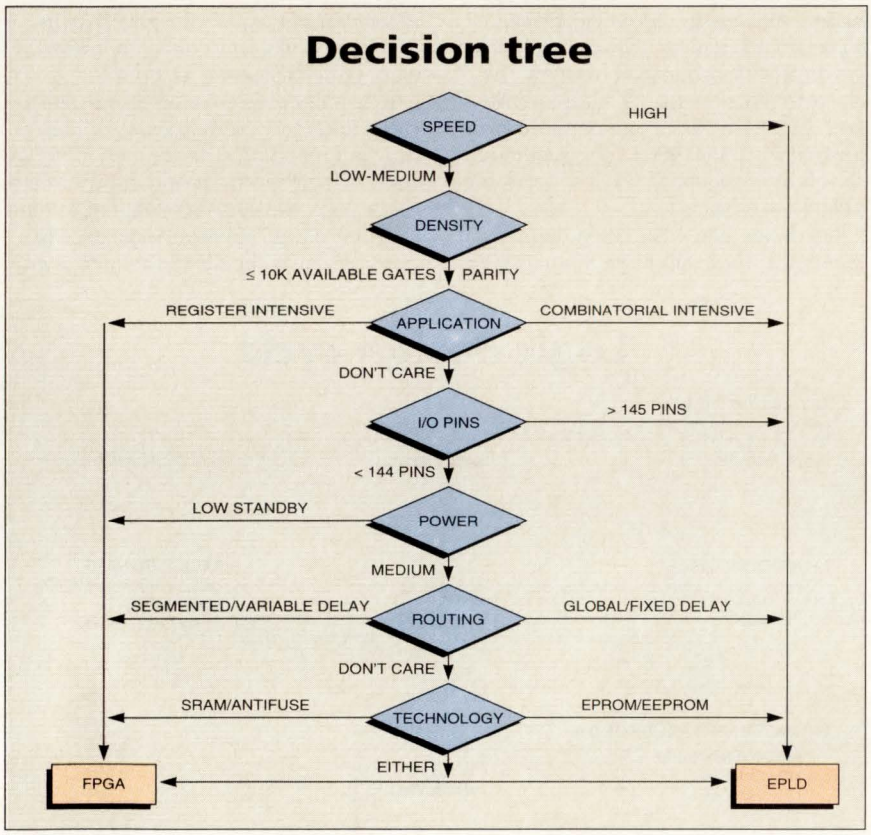

MAX EPLD maker Altera prepared this diagram to help users, faced with a set of design requirements, determine whether to use FPGAs or PLDs as a silicon implementation.

Nebula uses VHDL

At Nebula, designers use VHDL-based tools from Exemplar Logic (Berkeley, CA) to remap FPGA technologies to ASICs, or ASIC technologies to FPGAs. Though they can't use VHDL for existing FPGA designs that customers want remapped to an ASIC, designers do use VHDL for new designs. "If you specify the design in VHDL, the redesign effort is much smaller," reports Beal. "The I/O pads have to be dealt with differently, but most of the work has been done up front."

Nebula designers also use the Viewlogic toolset. Schematics are necessary for connecting blocks of synthesized logic as well as I/O pads, and also for control logic that's difficult to describe.

Though FPGA tools are catching up to FPGA silicon technology, San Diego-based consultant Mike Denie, who specializes in programmable logic design, says, "At this stage of technology, nobody can afford the slop that comes with synthesis for FPGAs."

Still, Denie wishes he had had some of the tools currently available a few years ago when he couldn't find any way around using 72 22V10s and 20V8s in addition to a trio of Xilinx 3042 FPGAs in a piece of ultrasound equipment. With production runs of 10 a month for the quarter-million-dollar equipment, end-user Philips Ultrasound couldn't justify the NRE for an ASIC. In addition to cost reasons, according to Denie, "It would have been impossible to know what the requirements would be. The design was a moving target because of competition." Denie says that designers ended up doing about a dozen iterations before finishing.

With today's tools, Denie says he might have mapped the dozens of PLDs into FPGAs. He might also have chopped out about 30 or 40 chips if he could have considered Xilinx and Actel FPGAs for what he calls 'dataflow garbage collection,' Altera's MAX EPLDs for control logic, Concurrent Logic's FPGAs for pipelining, and Atmel's (San Jose, CA) high-performance complex PLDs, for which he has a great deal of respect. "But in the real world, you don't have access to everyone's silicon," says Denie. "You go to the director of engineering and ask him for $10,000 for a certain complex PLD or FPGA vendor's development system, and you're stuck with that silicon."

Even with a vendor-independent tool like that from NeoCAD (Boulder, CO), Denie says that "It costs about $15,000 for the basic software and $5,000 for each new architecture. It's a hard sell to get designers used to Abel and Cupl to put forth a lot of money for software. It's less difficult for designers used to gate arrays."

Common design methodology

Designers at Cossor Electronics, Ltd (Harlow, Essex, England) are also frustrated by the need to use proprietary tools for FPGAs and complex PLDs, according to VLSI manager Mike Payne. "We're specifying our ASIC designs in VHDL, but we're still using proprietary tools for Xilinx, Altera and Atmel devices."

For ASIC design, Cossor has been one of the first to adopt LSI Logic's (Milpitas, CA) VHDL-based Silicon 1076. "Now we're looking to see if we can develop a common design methodology that will encompass FPGAs, PLDs and ASICs," reports Payne. "We'd like to start with VHDL and go to whatever silicon technology we like."

That sort of technology-independent design methodology is exactly what makes project engineer Chris Kulbida at Xerox (East Rochester, NY) a VHDL fan. A Xilinx user, Kulbida uses FPGAs wherever possible to prototype advanced technologies, but he says that VHDL lets him

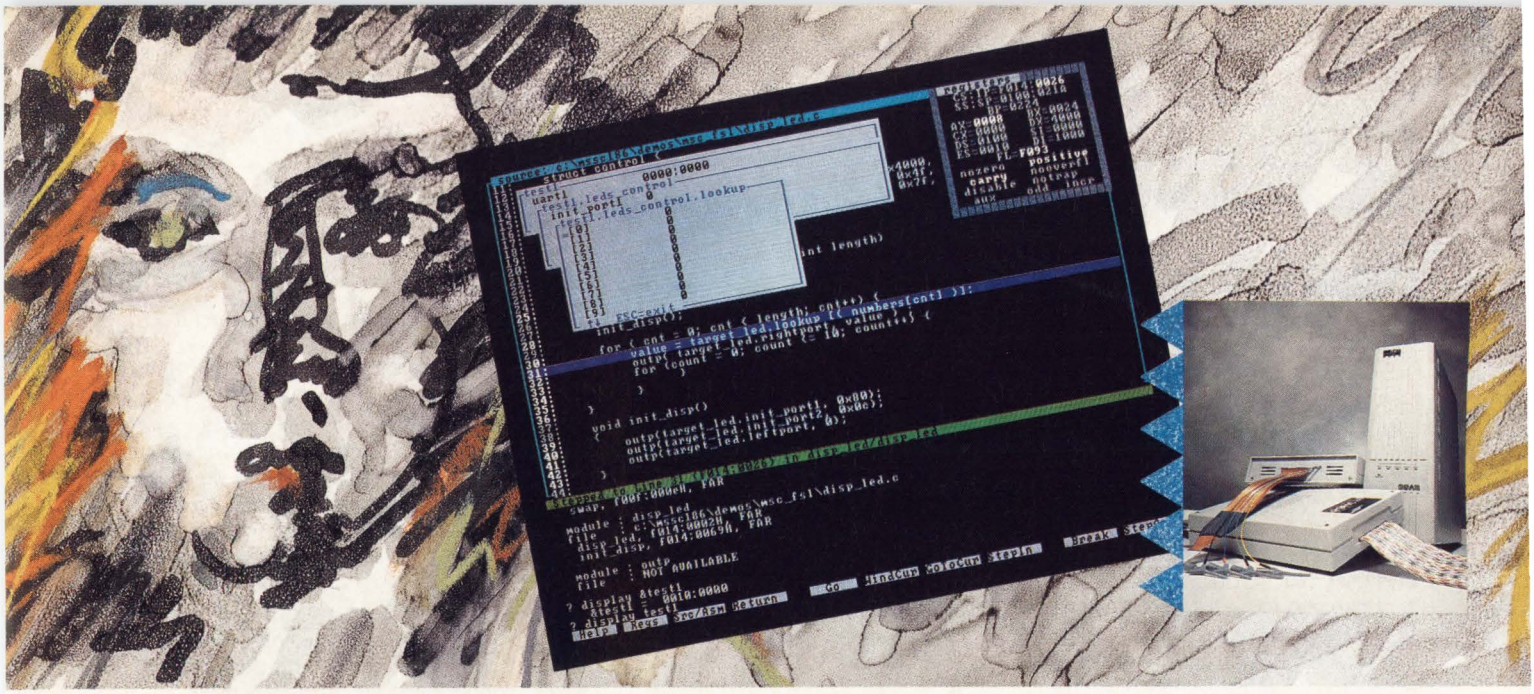

Power Debugging

. . . for Real-Time Targets

SOPHIA EMULATORS SUPPORT	
INTEL	
80C86/88	10 MHz
80C186/C188	20 MHz
80286	25 MHz
80386	25 MHz
80386SX	16 MHz
MOTOROLA	
68000/10/12	16 MHz
68020	20 MHz
NEC	
V20/V30	16 MHz
V40/V50	10 MHz
V25	16 MHz
V33	16 MHz
V53	16 MHz
86/V	16 MHz

. . . plus processors from AMD, Hitachi, Mitsubishi, Rockwell, Zilog.

Do you need to debug a real-time embedded system? Do you use C along with assembly language?

Our emulation systems can help you complete your project faster.

Real-Time Target Control

- ❏ Emulate without stealing target resources or adding wait states.
- ❏ Step through your actual source code, comments and all.
- ❏ Choose from many popular compilers.
- ❏ View code, registers, data, and commands all at once with our windowed interface.
- ❏ Capture precise information using our powerful trace capabilities, including sequential triggers, all in real time.
- ❏ Use full symbolic debugging capabilities, including high-level trace.

Real-Time Technical Support

- ❏ Experienced software and hardware engineers.
- ❏ Step-by-step troubleshooting assistance.
- ❏ 24-hour Bulletin Board for users. Upload problem software, questions. Download code fragments, new releases, examples.

Sophia Systems' emulators give you the *power* to debug your target in your environment. And we're continuing to expand our capabilities to match your processor and debugging needs.

Call and tell us about **your** target.

Sophia systems and Technology

1-800-824-9294

Sophia Systems and Technology, 777 California Ave., Palo Alto, CA 94304, TEL: (415) 493-6700 FAX: (415) 493-4648
Sophia Computer Systems, Ltd., Alpha House, London Road, Bracknell, Berkshire RG12 2TJ England, TEL: 0344-862404 FAX: 0344-861374
Sophia Systems Co., Ltd., Shinjuku NS Bldg. 8F, 2-4-1 Nishishinjuku, Shinjuku-ku, Tokyo 163, Japan, TEL: (03) 3348-7000 FAX: (03) 3348-2446

©1991 Sophia Systems and Technology

Choosing reprogrammability as a design requirement

With an ever-increasing number of players in the FPGA market, many factors come into play when choosing an FPGA vendor. Obvious to the designer are ease of use, proven technology, performance, and cost. Not so obvious is one factor that can eliminate the cost of all design iterations and provide some significant advantages in the system design—the ability to dynamically reprogram the FPGA device any number of times in the lab or in the field.

Exploiting reprogrammability

The ability to electrically alter FPGA device logic over and over during the design phase reduces component cost and design time. The ability to alter the device's logic during the manufacturing phase provides increased flexibility, reduced cycle times and lower board costs. Standard board designs of high-complexity logic can be created and customized for different customers or products by downloading the FPGA program during manufacture. This can be accomplished locally with the board, or the customized logic can be downloaded over phone lines. Configuring products to meet customer needs for specific features can provide competitive differentiation and reduce product costs.

Furthermore, the ability to design multiple FPGA configuration programs into the end system and allow for in-system reprogramming can facilitate innovative applications.

Reconfigurable logic can be used to implement different logic designs for a single set of FPGAs. System power-up diagnostics, for example, can be loaded, and, upon successful execution, the application logic could then be loaded into the same FPGA. This reduces board space and increases reliability. More important, it leads to reduced component costs by eliminating dedicated diagnostic logic devices that are only used "once" upon power-up.

The FPGA component is in effect being shared by multiple applications. In fact, many applications can share the same reprogrammable device as long as they can be executed nonsimultaneously. An example of this sharing is read and write logic for peripheral devices such as tape and disk drives.

Reprogrammable design issues

Why aren't more designers implementing reprogrammable FPGAs in these kinds of innovative applications? There are several reasons. Designers need to be certain in-system reprogramming time is fast enough to accommodate the system speed requirements. The different logic designs to be swapped in and out must fit into the same FPGA with the same pinout. The FPGA must also be cost-effective versus the alternative multichip fixed-function logic design, ASIC or discrete.

For those FPGAs offering partial reprogrammability, care must be taken to ensure that signal contention is not present when reloading only a portion of the FPGA. In addition, more designers need to be educated as to the competitive advantages possible through in-system reprogrammability.

Reconfigurable computing

The advent of high-speed, in-system logic configuration opens up entirely new applications. Now you can define your own computer architectures optimized for your application rather than rely on a standard general-purpose microprocessor with a fixed instruction set. In effect, you can design your own microprocessor.

Optimizing the processor architecture for specific applications isn't new. Custom processor designs are often developed and implemented in ASIC devices when system performance requirements can't be met with conventional microprocessors. This approach is time-consuming, expensive and rigid. High-speed image processing, which requires execution of many different software algorithms, is one example. Now, you can design your own optimized computer architecture and instruction set and have it downloaded in a high-performance reconfigurable FPGA for execution.

More important, many different computer architectures optimized for different algorithms or instructions can be developed and swapped in and out of the FPGA as needed by the application. It's like having an ASIC foundry with a turnaround time of milliseconds rather than weeks or months. Applications that don't fit conventional architectures can exploit reconfigurable FPGAs; such applications include signal processing, image processing, database acceleration, logic simulation, and many scientific computational problems. The reconfigurable FPGA is also ideal for large parallel-processing systems.

The advent of in-system electrically alterable programmable logic has freed the logic designer from the "hard" nature of traditional logic design, just as EEPROMs freed his or her counterpart in software design years ago. Like EEPROMs, reconfigurable FPGAs are being used in the most innovative applications now. But as the cost gap with conventional custom logic narrows, FPGAs will, like EEPROMs, be the mainstream design technology for logic design. New system architectures will emerge as reconfigurable FPGAs continue to narrow the performance and cost gaps with conventional ASIC technologies.

Charles A. Fox, director of component marketing, Xilinx, San Jose, CA

design in a more abstract manner. "VHDL gives me latitude in the choice of a target technology. I don't have to decide right up front how I'll implement a design. We have to demonstrate functionality very quickly," says Kulbida. "That keeps us from pursuing gate arrays and standard cells, and space constraints often prevent the use of PLDs."

Kulbida says that sometimes, however, he has to use PLDs for speed. To accommodate a system speed of 16 MHz in a recent design, he turned to Advanced Micro Devices' 12- and 15-ns Mach devices. "The design involved complex functionality going on within the clock cycle. With Xilinx, we had trouble getting the routing to give us good timing at 15 MHz," he reports. "It was almost an unknown what the timing would be with each routing."

FPGAs are ideal for small R&D outfits such as Teledyne Brown (Huntsville, AL), which does one-

68000 Developers Do It FreeForm!

FreeForm sets you free! FreeForm is the *first* remote debugger that gives you the freedom to incorporate *all* the resources and tools of your host machine *directly into your target debugging environment!*

FreeForm does it right! Now you can view source with your favorite editor, redirect debugging output into a file, pipe the output to another program, or incorporate your programs *right into FreeForm.* You never have to *exit or escape* from FreeForm in order to compile, link, and re-download your code.

FreeForm does it real time! Your application runs in real time because FreeForm is a nonintrusive, monitor-based debugger. And it's not necessary to turn off the code optimizer, because FreeForm debugs fully optimized, production quality C code.

There's only one thing better than FreeForm-- and that's FreeForm and CrossCode C together! CrossCode C is an advanced optimizing ANSI C compiler package used in over 1,000 applications. It's so good that three out of four clients are referrals from other satisfied users.

FreeForm gives you power! Together with your host (UNIX or MS-DOS) and your target, FreeForm and CrossCode C give you the industry's most powerful 68000-series development system!

FreeForm does it all! It lets you debug all ROMable 68000-series applications,whether you're designing around the 68000, the 68300 family, the 68040, or the new 68ECXXX.

Do it FreeForm Today! Call us at: 1-800-448-7733, EXT. 520 to receive your **FreeForm Demo Kit** for a 30 day free trial.
Do it FreeForm just once--and you'll never go back!

SOFTWARE DEVELOPMENT SYSTEMS, INC.
4248 Belle Aire Lane
Downers Grove, IL 60515 USA
Phone: 1-708-971-8170
FAX: 1-708-971-8513

The Most Responsive Software Company

FreeForm (TM), CrossCode (R) and UNIX (R) trademarks

CIRCLE NO. 58

of-a-kind prototypes to support optical computing contracts. Says design engineer Russ Dewitt, "PALs don't have the density we need, and ASICs are geared to production." For a decision-making-with-lights optical computing project he's working on, Dewitt is using Concurrent Logic FPGAs for their speed.

When FPGAs are not an alternative, and the choice is between gate arrays and standard cells, what makes end users go with a standard cell rather than a gate array? At Adaptive Networks (Cambridge, MA), designers went to standard cells because they needed a full-custom core element. Required was a shrink of the company's 3 × 5-in., surface-mount, multiapplication powerline carrier LAN module into a two-chip set consisting of a datalink layer (or communications processor) and an analog/digital physical layer chip.

After specifying requirements for the communications processor and studying alternatives in terms of functionality, cost and cost over time, the company decided to license a 65C02 microcontroller core and combine that with standard-cell library elements and a few custom blocks. "We work with a lot of different applications in the OEM business and need different interfaces," reports Adaptive Networks president Michael Propp. "It would have been a compromise with something else."

The communications processor has 12 kbytes of ROM and 768 bytes of RAM for resident OEM applications firmware. The first application of the two-chip set will be in meter readings, reports Propp, where the demand will be for tens of thousands of chips. "To modify the communications processor for later applications, we can take the microcontroller, cut out parts not used in that application and produce at low cost in high volume for a specific application."

Gate arrays for speed

What pretty much rules out the use of FPGAs at Convex Computer (Richardson, TX) is the performance-critical nature of the designs the company does, according to Harold W. Dozier, vice-president of advanced development. Most of the logic Convex puts into its machines is ECL and GaAs. Designers use a lot of high-speed PALs, as well as GaAs gate arrays from Vitesse (Camarillo, CA) and high-performance BiCMOS gate arrays from Texas Instruments.

Convex designers make the silicon implementation decision right up front, Dozier says. When doing system-level design, they're already making decisions concerning the physical design after considering the cost and performance targets of the machine they're building. "We ask what will go into gate arrays and what won't," he explains. "We ask how many parts there'll be, how much functionality will go into each, what the pinout limitations are, and finally what functionality will be left for glue logic." ■

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

Actel (408) 736-1030	Circle 215	**Mentor Graphics** (503) 685-7000	Circle 227
Advanced Micro Devices (408) 732-2400	Circle 216	**Model Technology** (512) 928-6000	Circle 228
Aldec (805) 499-6867	Circle 217	**NeoCAD** (303) 442-9121	Circle 229
Altera (408) 984-2800	Circle 218	**OrCAD Systems** (503) 690-9881	Circle 230
American Microsystems Inc. (208) 233-4690	Circle 219	**QuickLogic** (408) 987-2000	Circle 231
AT&T Bell Laboratories (800) 372-2447	Circle 220	**S-MOS Systems** (408) 922-0200	Circle 232
Concurrent Logic (408) 522-8700	Circle 221	**Synopsys** (415) 962-5000	Circle 233
Data I/O (206) 881-6444	Circle 222	**Texas Instruments** (800) 336-5236 ext.700	Circle 234
Dazix/Intergraph (205) 730-2000	Circle 223	**Viewlogic Systems** (508) 480-0881	Circle 235
Exemplar Logic (510) 849-0937	Circle 224	**Vitesse Semiconductor** (805) 388-3700	Circle 236
Isdata (408) 373-7359	Circle 225	**Xilinx** (408) 559-7778	Circle 237
LSI Logic (408) 433-8000	Circle 226		

Postscript

The programmable wonders we call FPGAs have become a very real and down-to-earth part of ASIC design methodology. With Xilinx having announced its 10,000-gate XC4010 logic cell array, the competition among FPGA silicon vendors is just heating up. At the same time, system designers are learning just how valuable FPGAs are to product development.

But don't be misled into believing that FPGAs are a panacea. Of the three FPGA boasts made most often—ease of design, no nonrecurring engineering expense (NRE) and time-to-market advantage—the last may be the only legitimate claim. If you've ever fully characterized an FPGA, you know it's not easy, and it takes a good deal of engineering time. And time costs money, whether an NRE label is attached or not.

As long as you consider the FPGA as a piece of the puzzle rather than as a total solution, you won't be disappointed. More often than not, you'll want to consider the FPGA as a design aid that offers the flexibility of going to an ASIC if volumes warrant, or of validating an ASIC design while waiting for silicon, or of reusing existing PAL and PLD designs.

If you have any leverage with your engineering director, say you need a design system based on a common design methodology (if you can find one) that encompasses PALs/PLDs, FPGAs, gate arrays, and standard cells. And make sure VHDL and synthesis are high on the list of specifications for that system. With today's design requirements, you can't have too many productivity tools!

Barbara Tuck

Why send your engineers through the VME maze?

MATRIX frees your engineers to do what they do best... solve application problems – not build VMEbus systems.

Reduce your development risk and speed your time-to-market with an integrated VMEbus system from MATRIX Corporation. Since 1985, we have specialized in bringing the power of the VME open-architecture to integrated systems. Now you can start with a pre-configured system and add modules to create a fully integrated system ready for immediate application code development.

Your engineers are experts in their field of application. They can't afford to spend the time or resources to install and bring-up operating systems, de-bug software drivers, configure boards, make cables, and resolve finger-pointing disputes between manufacturers. Let our investment in qualifying and integrating third-party hardware and software pay off for you. We offer:

DATA ACQUISITION SYSTEMS:
2 GHz sampling front-end.

DATA STORAGE SYSTEMS:
30 Mbytes/sec transfer rate.

DATA ANALYSIS SYSTEMS:
Multiple array processors.

MATRIX Corporation provides a fully documented system, and post-sale support includes training, installation, telephone, and on-site support.

A basic system contains a development host, a target CPU, and an operating system. Add modules from the MATRIX full-line VME catalog, or from over 20 pre-qualified vendors, or specify your own modules.

MATRIX Corporation will deliver the system you need, on time, ready to develop application-specific code. Call us at **1-800-848-2330** to discuss your specific system requirement or to receive our full-line VME catalog.

MATRIX CORPORATION

1203 New Hope Road, Raleigh, NC 27610
USA FAX (919) 231-8001

CIRCLE NO. 59

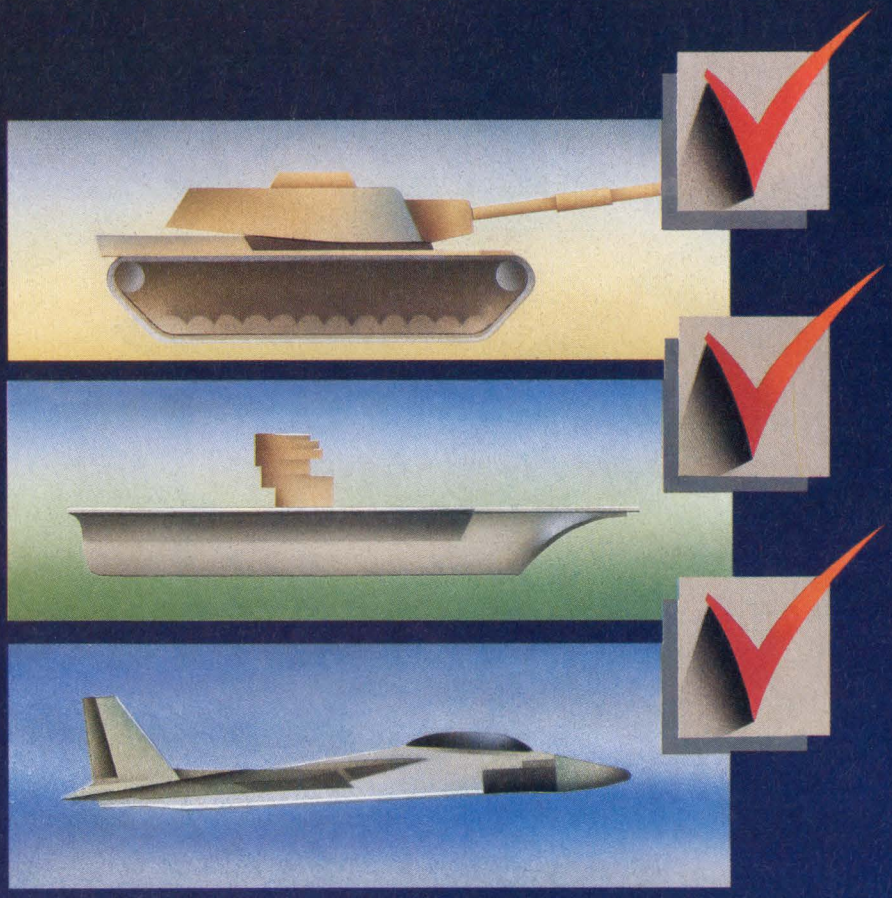

Mil-Spec NDI Computer Sub-Systems to Meet the Standards of the U.S. Army, Navy, Air Force, Marines ...and the U.S. Taxpayer.

The urgent need for reliability in Mil-Spec equipment and systems requires no further justification. If you're designing or integrating computer systems for *any* of the Armed Forces, reliability is paramount. And it must be achieved at reasonable cost levels.

At Radstone Technology, we have a proven strategy that enables you to control costs without compromising reliability, quality or performance. Our Mil-Spec computer boards and sub-systems meet the most stringent U.S. and international military requirements *and* the increasing mandate for NDI (Non-Developmental Item) products. Radstone NDI Mil-Spec VME provides proven economies in terms of hardware, software development and compatibility, plus reduced time to market and performance-proven reliability. That's why, now more than ever, the DoD, military and prime contractors are counting on Radstone for Mil-Spec NDI VME.

Radstone Mil-Spec VME. It's the right way to cut costs without cutting corners. Let us show you how *your* programs can take advantage of our technology, using an accepted standard and a winning strategy. We'll help you meet the requirements of the military *and* the taxpayer.

Radstone Technology Corporation
20 Craig Rd., Montvale, NJ 07645-1737
Call Toll-Free: (800) 368-2738
Eastern Region: (201) 391-2700
Central Region: (708) 304-0202
Western Region: (408) 727-4795

CIRCLE NO. 60

RADSTONE
TECHNOLOGY

SDI testbed relies on standard hardware to rein in costs

The SB-EV design and implementation team: (clockwise from bottom) Mark Sleeper, network analyst (sitting); Martin Durbin, SDI architecture analyst (sitting); Dr. Charles Eldridge, principal investigator; Steve Elliott, senior SDI architecture analyst; Mike Anderson, director of Addem project office; Craig Franklin, project engineer; Robert Harris, senior network analyst; Gordon Lee, project engineer; and Tim Wall, principal investigator.

The Strategic Defense Initiative (SDI) presents designers with a myriad of technical challenges. Perhaps the most significant is implementing battle management command, control and communications (BM/C^3) functions. SDI battle management includes such complex architectural issues as communications network management, human interaction, battle management coordination schemes, and system reliability. To study these issues, the U.S. Air Force Electronic Systems Division (ESD) requested proposals for a testbed called the Space-Based Experimental Version (SB-EV).

The ESD conceived the SB-EV as a "proof-of-concept" vehicle. Its purpose was to demonstrate SDI's distributed BM/C^3. Using the SB-EV, the Air Force can simulate the entire battle space, including SDI satellites, ground management stations and incoming missiles—and can do it in real time.

"The government wanted a software prototype that could be used as the basis for developing algorithms and to evaluate architectural alternatives," says Karl Warmbrod, chief engineer at Sparta (McLean, VA). "It also wanted something that could evolve; it didn't want something that would freeze the technology."

At the time, conventional wisdom pointed to monolithic solutions built around mainframes or super minicomputers. "We had already developed a software simulation for another project that contained all the elements required for the SB-EV," Warmbrod explains, "but we weren't happy with the results. Because of the large number of satellites involved, and the sheer volume of the simultaneous, independent, asynchro-

Bob Margolin, Contributing Editor

nous decisions that these satellites must make, the software was extremely complex. It was also unwieldy to maintain, extend and change. It's a simulation nightmare."

Sparta's engineers realized there was no way to accurately simulate SDI in software. The problem stemmed from the fact that no matter how fast the host computer was, it still had only one CPU. This meant that communication requests would have to be queued up to wait for the data bus. Because of this fact, communication between any two points could suffer from indeterminate delays that would not occur in the real system. "We looked at SDI itself and realized that it was really a loosely coupled computer network," says Warmbrod. "So we decided the best way to meet ESD's needs was to emulate the SB-EV in hardware."

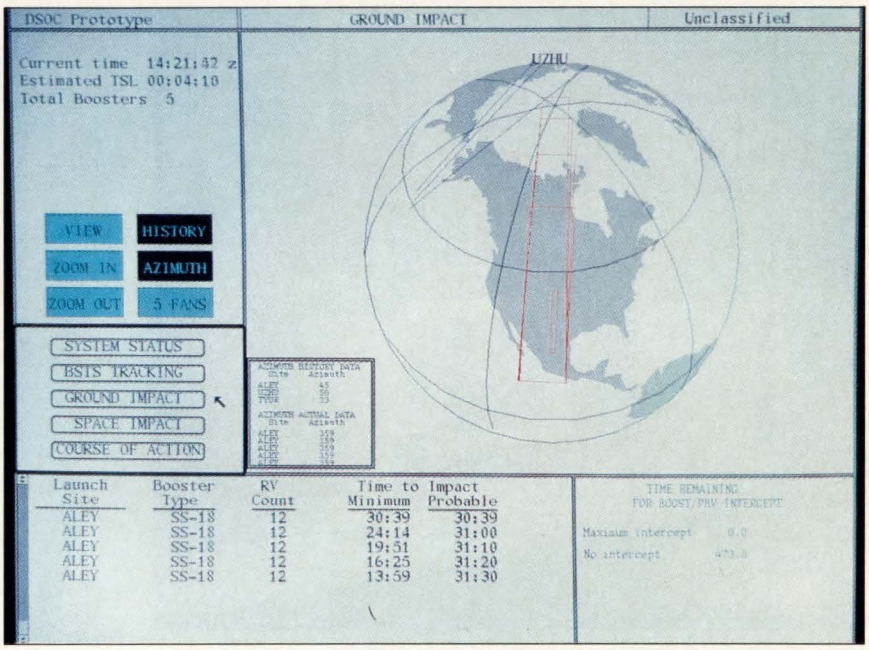

This screen shows a graphical interactive display of an experiment using the SB-EV testbed.

Hardware emulation

Emulating the SB-EV meant using a separate computer to represent each battle manager. According to Warmbrod, this approach gave the testbed great flexibility. "If we need to emulate a system with 20 battle managers, we'll use 20 computers; if we need 100 battle managers, we'll use 100 computers," he says. "We'll run the same operating system on each of them, but each will have its own unique database and all will run asynchronously, just as they will in the real system."

Because the SB-EV contract was awarded based on competitive bidding, cost was a major concern. "We wanted to develop a distributed-processor architecture that would be a faithful representation of the real system, but which used completely commercial, off-the-shelf solutions. We didn't want to custom design any hardware," explains Mike Anderson, Sparta's senior engineer.

One possible route to using off-the-shelf hardware was to tie a number of fast workstations together on a high-speed LAN. "The problem with this solution," Anderson explains, "is that the interworkstation communication would be too slow." That's because it takes slightly more than one nanosecond to propagate a signal through one foot of copper wire, regardless of how fast the communication mechanism is between the workstations. In a system constructed of 100 or more workstations, the distance between them could become substantial. A much better solution, says Anderson, is to use single-board computers (SBCs). A system comprised of 100 SBCs could fit in a single 19-in. equipment rack, with intercomputer communication handled through the VMEbus backplane.

"We chose the VMEbus environment because VME SBCs were readily available from several vendors. Also, at under $2,000 each, they were very affordable," Anderson says. "The initial design called for the SBCs to use the 68020 microprocessor. Although slow by today's standards, at the time the proposal was being developed the 68020 provided the biggest bang for the buck."

Cost, while a very important consideration, wasn't the only reason to go with off-the-shelf hardware. Using standard, commercial VME equipment greatly simplified the problem of repair and replacement in the field. There were more than 60 manufacturers producing VME SBCs, and many of the boards are essentially interchangeable. There are also well over 100 companies making VME components.

Choosing the operating system

Settling on a hardware architecture was only part of the solution; there was also the matter of an operating system to resolve. "We had to select the operating system based on the performance needs of the SB-EV, including Ada compatibility. But we also had to keep in mind the cost of developing the application programs," Anderson states. Sparta could have chosen a proprietary operating system, but that would have required the use of a half-million dollar VAX minicomputer for applications development. The other alternative was to use a standard, open operating system such as Unix.

However, for SDI to perform its intended functions, it requires a real-time, deterministic operating system—which Unix isn't. Even so, says Anderson, "We decided to go with a real-time operating system that looked enough like Unix that we could do our development under Unix on workstations, which at the time were going for about $15,000 apiece."

Sparta set about evaluating the available Unix-like real-time operating systems, but found few potential candidates. "Many of the Unix-like operating systems were not code-compatible with Unix or were too slow to process data representing the thousands of missiles SDI was originally intended to track and destroy," says Anderson.

Sparta finally settled on VxWorks and its Ada equivalent, VADS-Works. "The reason," explains Bob Harris, one of Sparta's principal engineers, "was that VxWorks gave us

DY 4 VME VETRONICS LEADERSHIP

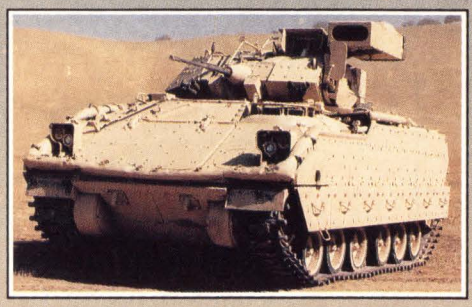

DY 4 is the industry leader in providing cost-effective VMEbus system solutions to vetronics system integrators. A full line of MIL-Spec, rugged, and commercial VMEbus modules are available today, including a comprehensive range of products compliant with vetronic architecture standards.

DY 4 is experienced in the full scale support of system integrator's flow down requirements with custom engineering, a military manufacturing facility, a 9858A quality program, foundation software, ILS and configuration management procedures. Contact DY 4 for your vetronics requirements.

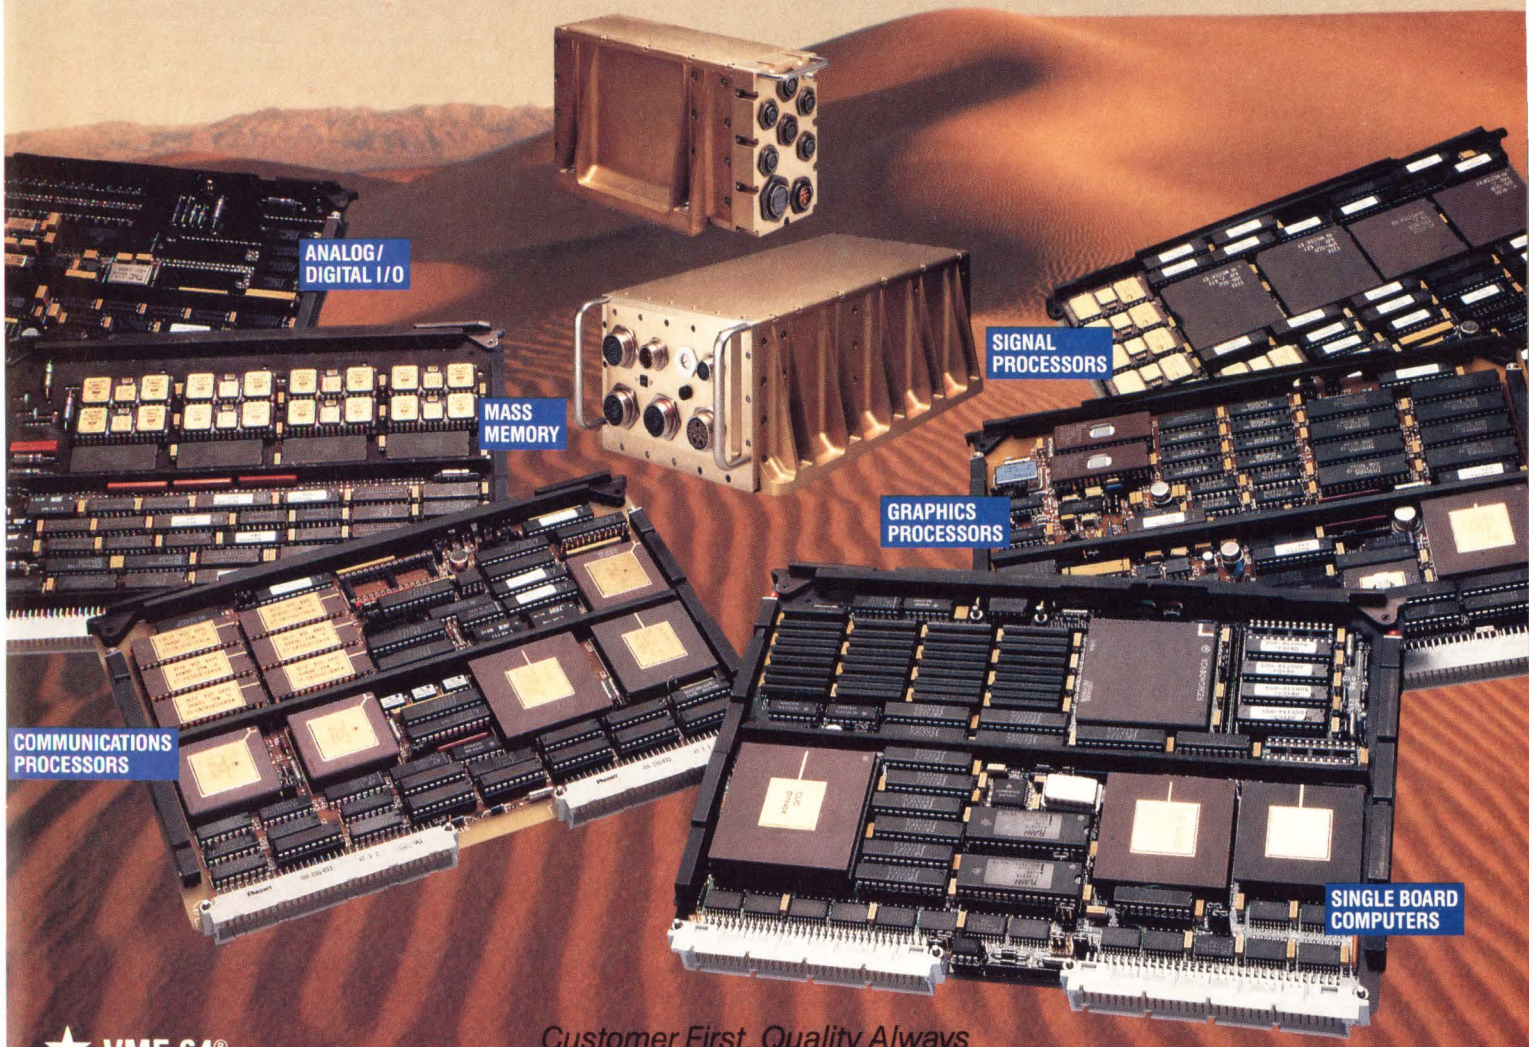

★ VME 64®

Customer First, Quality Always

DY 4 Systems Inc.

Nashua, New Hampshire
Tel: (603) 595-2400
Fax: (603) 595-4343

Campbell, California
Tel: (408) 377-9822
Fax: (408) 377-4725

Hammel, Denmark
Tel: +45-86-963624
Fax: +45-86-962575

Pennant Hills, Australia
Tel: +61-2-484-6314
Fax: +61-2-875-1665

Ottawa, Canada
Tel: (613) 596-9911
Fax: (613) 596-0574

®Registered Trademark of Performance Technologies Inc.

CIRCLE NO. 61

A COMPLETE MODEM!

CH1782 - 2400 bps/V.22 bis

- FCC Part 68 Approved
- 2" x 1-1/8" x 1/2"
- "AT" Commands
- 5V, LOW POWER
- Low Cost

$49.95 10K Qty.

Your Source for Modem Components:

Cermetek microelectronics

Tel: 408-752-5000
Fax: 408-752-5004

Cermetek Microelectronics, Inc.
1308 Borregas Ave. • Sunnyvale, CA 94089

CIRCLE NO. 87

RTXC™ & RTXC/MP™

Written in C, portable by design	Intel 80x86, 80x96, 8051 Motorola 68HC11/16, 68xxx TI TMS320C30/C40 Zilog Z80, Zx80 Analog Devices 2101 Inmos T2xx, T4xx, T8xx
Dynamic priorities Task Management Timer Management Memory Management Semaphores Message mailboxes FIFO queues Resource Management	From 8-bit microcontrollers to multiple 32-bit processors : same API !
Virtual single processor with RTXC/MP !	(distributed) Debugger System generation tool Workload monitor
Transputer networks C40 networks VME networks	Message based Transparent Distributed operation !
Free demo and evaluation kit New version 3.0	No royalties Source code included 3 months free upgrades and support

Intelligent Systems International
Lindenstraat 9, B-3210 Linden, Belgium
Tel.(+32)16.62 15 85. Fax.(+32)16.62 15 84

(List of distributors for USA and Japan on request)

RTXC trademark A.T. Barrett & Associates. RTXC/MP trademark Intelligent Systems International.

CIRCLE NO. 88

a real-time operating system base with a shell that looked like Unix. This let us develop applications that thought they were running on Unix, but which gave us the real-time operation the SB-EV needed." An added benefit was that the source code, in VADSWorks, was fairly compatible with Unix, so moving code back and forth between Unix workstations and the SB-EV was easy to do.

The choice of the operating system also affected the choice of SBCs. Although many VME SBC manufacturers support VxWorks, most rely on third-party vendors to provide it. "We wanted to use SBCs made by companies that provided their own software support," says Warmbrod. That desire pointed toward Heurikon (Madison, WI), which not only supports VxWorks but actually owns the source code.

Says Richard Peterson, a Heurikon applications engineer, "Because we own the VxWorks source code, we can correct bugs for our customers immediately. They don't have to wait for the next version of VxWorks, which may not be available for several months. Owning VxWorks source also benefits Wind River Systems. As soon as we uncover a bug, we alert Wind River and send them the fix, which they can incorporate into their next revision."

Support for VxWorks was very important, but Sparta also wanted to talk directly with the vendor's engineers so that technical problems could be quickly resolved. "Most vendors," says Anderson, "wanted us to deal with them through their sales force, or contact the OS vendor to resolve software and interoperability problems."

Heurikon actually did more than just talk to Sparta engineers. According to Anderson, Heurikon performed the initial integration. "They loaded the SBCs—68020-based HK68/V2Fs—into the VME chassis and initialized the system. Although we had to modify it to meet our specific needs, they gave us a working system."

Architectural problems

There were some problems associated with the distributed architecture Sparta was proposing. One was that, at the time, there was no validated Ada compiler for loosely-coupled architectures based on 68020-based SBCs. Development of a multi-SBC Ada compiler was under way, but its validation was still some time off. As Warmbrod recalls, "We didn't think it would be available in time, so we had to put some money into accelerating the validation process."

Another problem was that of rehosting the application code. In a monolithic simulation, the code is tailored to the specific host in which the simulation is being run. Regardless of how the simulated system configuration changes, the host remains the same. However, in the distributed architecture Sparta envisioned, the application code would have to be ripped apart and revised every time the system was reconfigured. "We had to define the interprocessor communications mechanisms for the underlying architecture, and that definition changes as the number of processors in the system changes," says Warmbrod. "To eliminate the need to revise the application code repeatedly, we had to separate the application code that tells the battle manager what to do from the host-dependent code that describes the system configuration."

Because Sparta is a relatively small company, it had to rely on subcontractors for some of the software development. One of those subcontractors, General Electric Space Systems Division (Philadelphia, PA), was assigned the task of developing a software architecture that would serve as the SB-EV's communications framework. "Within this framework," Harris explains, "each process or function is assigned a unique logical name. Thus, any process or function can communicate with any other in the system by addressing its unique name, regardless of which SBC is running the corresponding application."

Using subcontractors does shorten development time but, as Harris points out, "It's hard enough to integrate the efforts of two branches of the same company. When you're dealing with different companies, it's much more difficult because you also have to deal with separate chains of command. One of the reasons we were so successful with our subcontractors was that we gave each of them a lot of responsibility for handling their specific tasks."

SBC-to-SBC communications

The SB-EV's SBCs were to be housed in several rack-mounted VME chas-

sis, which meant that Sparta engineers had to address both intra- and interchassis communications. "Providing intrachassis communications was relatively straightforward," says Anderson. "Communications are conducted over the VME backplane bus, and VxWorks provides a backplane driver that uses the shared VME memory to emulate a TCP/IP local area network." Although the emulated LAN was relatively slow, at 250 kbytes/s it was more than fast enough for the task at hand.

Interchassis communication proved to be a little more challenging. According to Anderson, there were two paths Sparta could follow. One was to use backplane or bus repeaters. These repeaters were capable of transferring data at about 3 Mbytes/s, but they could only connect two chassis together. To this day there are no multichassis backplane repeaters. "The only other alternative available to us was to use a local area network. But to use a LAN for high-speed, deterministic interchassis communications, we needed to pull simulation artifacts out of the data stream."

Separating the simulation artifacts from the data stream also meant that in addition to the data channel, a control channel would be required. "Since it could be relatively slow," Anderson says, "we were able to use Ethernet for the control channel. We chose the 80 Mbyte/s Proteon ProNet 80 token-ring LAN for the data channel because, from the application's point of view, it appears to be another Ethernet running under TCP/IP." Thus, the entire suite of communications buses appears to the application as one large TCP/IP Ethernet, which really simplifies communication management.

Final steps

The SB-EV was conceived and developed during a period when the perceived threat was a barrage of long-range ballistic missiles launched from the Soviet Union. With the collapse of the Soviet empire, that threat has greatly diminished and, as a result, the role of SDI has changed. Today, the perceived threat is from a small number of missiles fired at shorter range. So, while the number of missiles that must be tracked and destroyed is much smaller, the flight time is also much shorter. "Although we would do things a little differently today," says Harris, "the original design is very well suited to today's reality."

Even so, some upgrading was required to increase throughput. Fortunately, the use of 68020-based SBCs proved to be a good choice. The SB-EV was borrowed back from the Air Force to perform a simulation for another project and was found to be slightly lacking in processing speed. The easiest way to increase the speed was to use faster microprocessors, so the decision was made to replace some of the 68020-based SBCs with 68030-based versions. According to Anderson, "the object code for the application software is virtually the same for the two microprocessors, so all that's required to upgrade the system with 68030-based SBCs is to swap the boards and recompile the code." ■

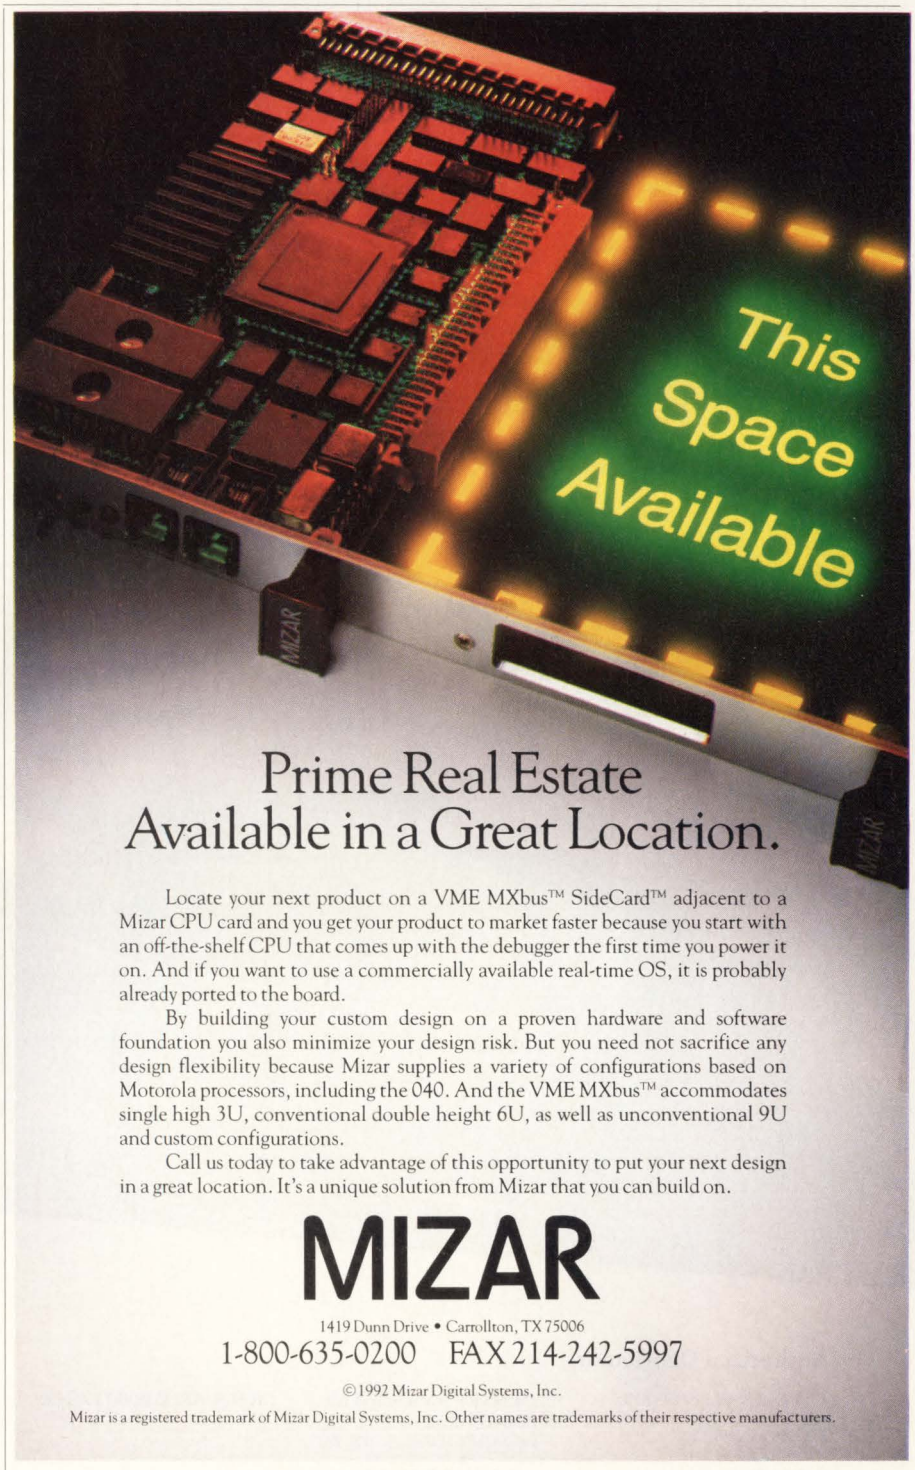

Prime Real Estate
Available in a Great Location.

Locate your next product on a VME MXbus™ SideCard™ adjacent to a Mizar CPU card and you get your product to market faster because you start with an off-the-shelf CPU that comes up with the debugger the first time you power it on. And if you want to use a commercially available real-time OS, it is probably already ported to the board.

By building your custom design on a proven hardware and software foundation you also minimize your design risk. But you need not sacrifice any design flexibility because Mizar supplies a variety of configurations based on Motorola processors, including the 040. And the VME MXbus™ accommodates single high 3U, conventional double height 6U, as well as unconventional 9U and custom configurations.

Call us today to take advantage of this opportunity to put your next design in a great location. It's a unique solution from Mizar that you can build on.

MIZAR

1419 Dunn Drive • Carrollton, TX 75006
1-800-635-0200 FAX 214-242-5997

©1992 Mizar Digital Systems, Inc.
Mizar is a registered trademark of Mizar Digital Systems, Inc. Other names are trademarks of their respective manufacturers.

CIRCLE NO. 89

Pushing the Edge of Capacity...

1.65 GB in a 5.25-inch Form Factor.

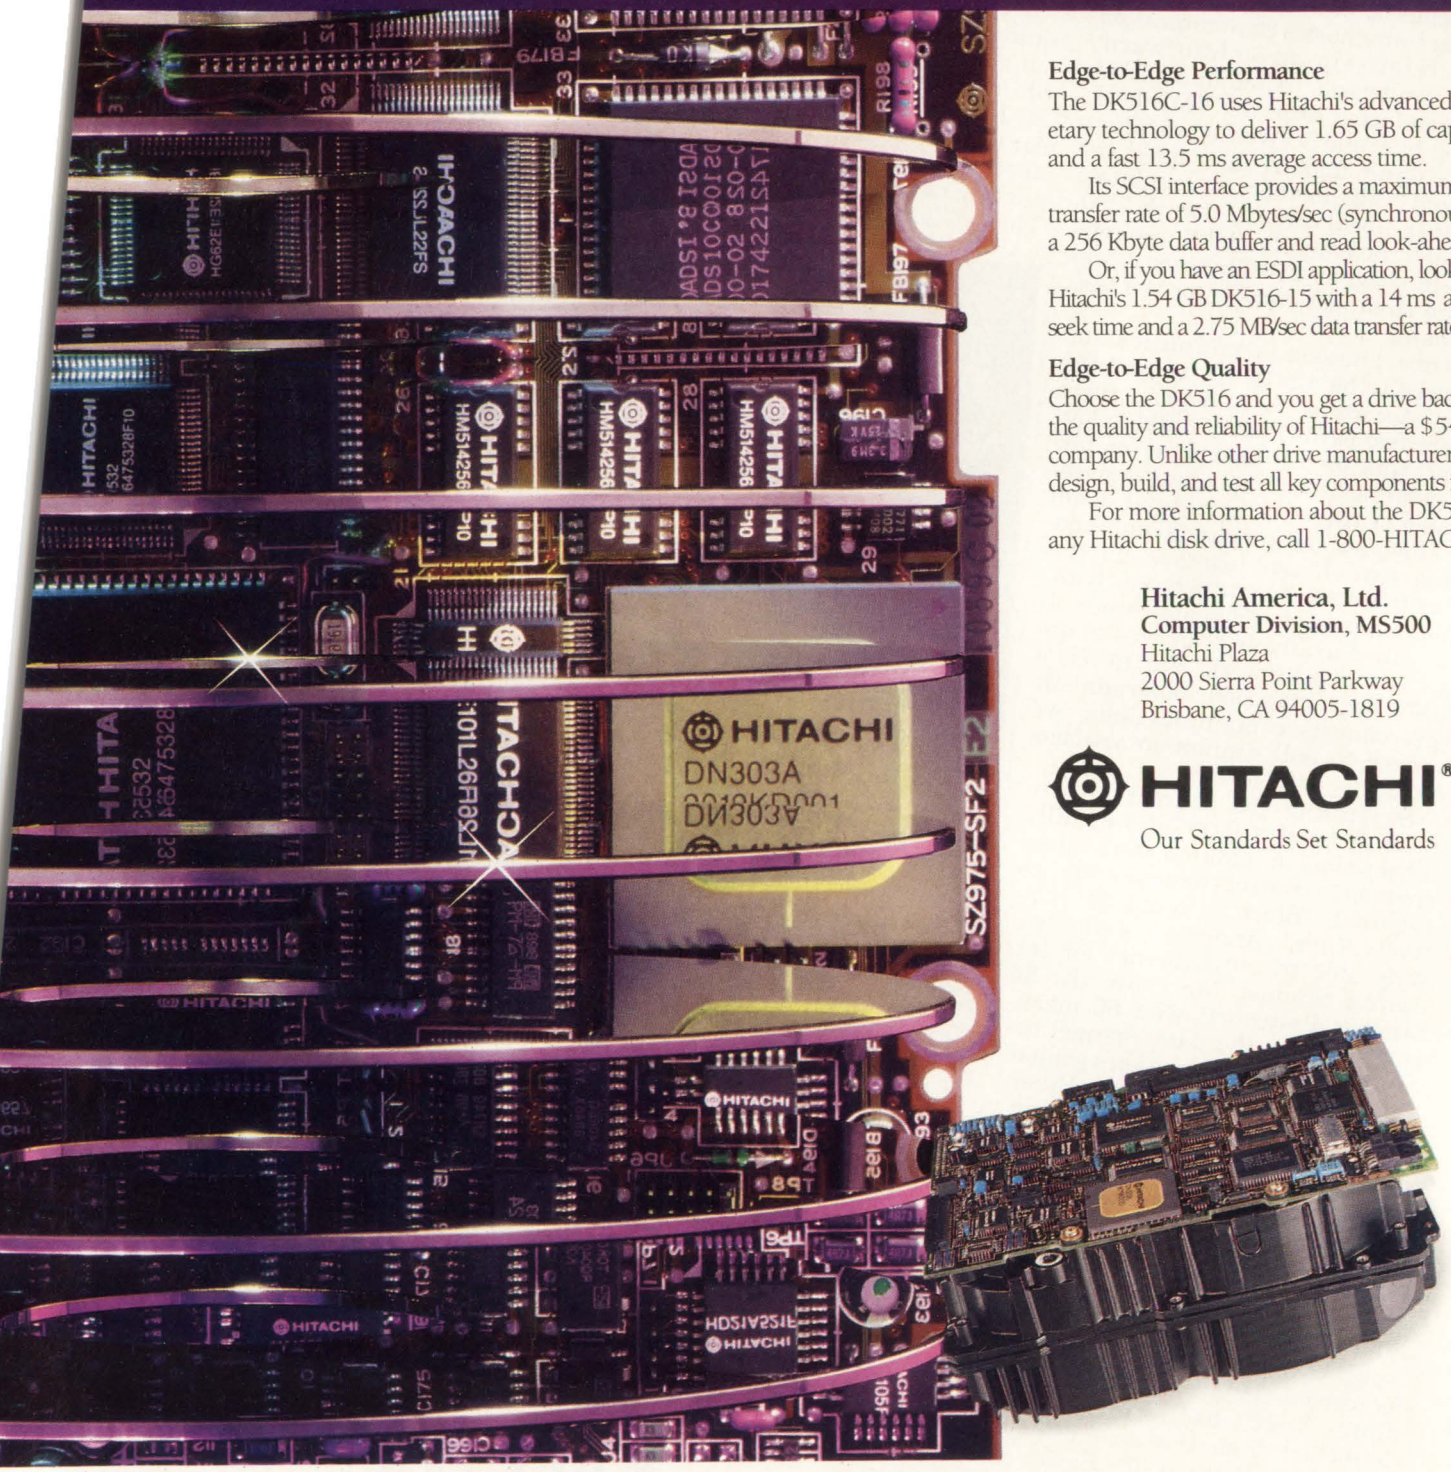

Edge-to-Edge Performance

The DK516C-16 uses Hitachi's advanced proprietary technology to deliver 1.65 GB of capacity and a fast 13.5 ms average access time.

Its SCSI interface provides a maximum data transfer rate of 5.0 Mbytes/sec (synchronous), with a 256 Kbyte data buffer and read look-ahead cache.

Or, if you have an ESDI application, look into Hitachi's 1.54 GB DK516-15 with a 14 ms average seek time and a 2.75 MB/sec data transfer rate.

Edge-to-Edge Quality

Choose the DK516 and you get a drive backed by the quality and reliability of Hitachi—a $54 billion company. Unlike other drive manufacturers, we design, build, and test all key components in-house.

For more information about the DK516, or any Hitachi disk drive, call 1-800-HITACHI.

Hitachi America, Ltd.
Computer Division, MS500
Hitachi Plaza
2000 Sierra Point Parkway
Brisbane, CA 94005-1819

HITACHI®
Our Standards Set Standards

Authorized Distributors:

CONSAN 612-949-0053 (IA, IL, IN, KS, KY, MI, MN, MO, ND, NE, OH, Pittsburgh, PA, SD, WI)

GENTRY ASSOCIATES 800-877-2225 (AL, D.C., FL, GA, LA, MD, MS, NC, SC, TN, VA)

R SQUARED 800-777-3478 (AZ, CA, CO, NM, OR, UT, WA, WY)

SIGNAL 800-228-8781 (CT, MA, ME, NH, RI, VT)

SPECIALIZED SYSTEMS TECHNOLOGY 800-688-8993 (AR, LA, OK, TX)

HITACHI (CANADIAN), LTD. 416-826-4100 (Canada)

CIRCLE NO. 62

INSTANT DATA ACCESS (IDA) DIAL (617) 494-8338 DOCUMENT NO. 1034

Fuzzy Logic is anything but fuzzy

Fuzzy logic is real, and its use is gaining popularity in the United States. Widely used in Japan, it can deliver results that are simpler to achieve, more precise and more reliable than conventional approaches to computing and control.

It often takes time for a good idea to catch on. First, it may be championed by a few "fanatics," while the mainstream considers it a cult phenomenon. Later, the press may pick up on it and popularize it out of all proportion, suggesting it as a cure for the world's ills. Finally, an idea will come to be seen in perspective for its strengths and weaknesses. If it's really a good idea, it may help in small part to solve the world's problems. Fuzzy logic is such an idea.

While it's still suffering from a negative image among English-speaking people because of its name, fuzzy logic has caught on in Japan and has been incorporated in some 2,000 products, ranging from rice cookers to subway systems. It's now poised to find its way into U.S.-designed products and systems.

Because the initial stages of a fuzzy-based design depend to some extent on intuition and the designer's "feel" for the nature of the problem, the development tools that have emerged recently are almost universally based on a graphical interface and the ability to state rules in simple textual form. These tools, many of which run on personal computers, may seem deceptively simple because they let you put together relationships of very rich complexity by merely drawing a few graphs and stating some rules. In doing so, they let you get your hands on real-world problems without

Fuzzy logic control is so precise that it could bring this fast bullet train to a smooth stop within inches of the stationmaster's toe.

By Tom Williams, Senior Editor

having to maintain a level of rigor and precision.

For example, a car can be considered properly parked when it's within 18 in. of the curb—or within 2 in. We need not specify 7.204 in., and we don't have to go to the trouble of computing and then homing in on such an exact result. Lotfi A. Zadeh, the inventor of fuzzy logic and currently professor emeritus of computer science at the University of California at Berkeley, calls this "exploiting the tolerance for imprecision." There are many problems in real-world control and fields such as finance and biology which lend themselves very well to such solutions.

On the other hand, fuzzy logic is anything but imprecise itself. The numerical results of its computations are exact and repeatable. The tolerance for imprecision comes from how we state the problem and establish the limits of what's acceptable in a solution. "Close to the curb" can be defined as anywhere between 2 and 18 in. because we say so, not because that's the best the system can do. We could as easily have told the fuzzy system that "close" was between 2.05 and 1.95 in., and the same system would have accepted this definition. The power of fuzzy logic comes from being able to state such a definition in natural language and by making simple changes to a graph, without having to recode in some computer language. As more designers realize this, we can expect increasing numbers to get very enthusiastic about using fuzzy logic—especially those who would rather concentrate on solving their problems than on mastering the intricacies of C or Fortran or some other language.

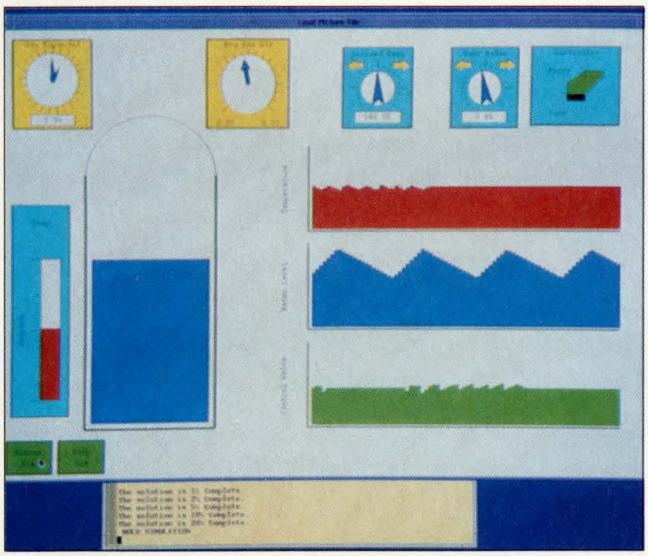

Fuzzy logic provides smooth control and can avoid the stress often caused by cycling components on and off. In this simulation of a water heater, the center chart (right) shows water level, which is affected by an operator opening and closing a drain. The top chart shows water temperature and the bottom the action of a gas valve to the heater. The first (left side) half of each plot illustrates conventional Boolean control, the second half fuzzy control.

Any new method of doing things, or even of looking at a problem, is going to require some changes in thinking. If you're talking about designing a control system using fuzzy logic, you'll definitely find differences, but there will also be similarities to what you're familiar with. "The first assumption," according to consultant Dr. David Brubaker, president of the Huntington Group (Menlo Park, CA), "is that we're designing a fully specified system." This is an assumption common to both conventional and fuzzy approaches, of course. Both types of control work best when you design to a specification of how something is to behave in the real world. It also helps in both cases to have a clear idea of what the inputs and outputs are to be. But, after the specification step, the methods of classical control and fuzzy design part ways.

For classical control systems, it's necessary to have a fairly complete mathematical model of how a system works—what equations govern the relationship between inputs and outputs. "You need to know your system very well," says Steve Houtchens, product manager for Integrated Systems (Santa Clara, CA), a software and development tools supplier for industrial control engineering applications. He says, "You have to know, for example, how your carburetor works. When I put in this amount of control—voltage—then how do my electrical and mechanical inputs relate to combustion? If I don't know it very well, then my control isn't going to work very well."

But, says Houtchens, there's no reason to go to that much trouble if you can use some other way of getting the answer that does not require you to know all the internal mathematics. "If I know intuitively that if I turn something up, I'm going to get this rate of change of output, I can develop a preliminary set of rules, assign membership sets to them and try that out. If it's not as close as I want it to be, I start adding more information."

So, if the carburetor in the example has been designed to take everything into account except air density, you can add some rules and membership functions that take air density into account and see how the system behaves. "It might take into account 28 different variables, and if I had to write out the equation, I'd have to find 28 different constants," says Houtchens. Using a fuzzy approach, you can add rules and membership functions to a design without having to radically revise a mathematical model.

Another characteristic of classical control strategy is the need to perform linearization on the model—that is, to find portions of the control curve that are approximately linear and characterize their operating behavior. Such a curve might have more than one region that needed to be controlled, in which case a control strategy would have to be defined for each linearized range. "That's why the classical method needs a very good model of what your plant is going to act like," Houtchens says.

It's a fundamental characteristic of fuzzy logic that it can handle nonlinearity. While you still want to have a control curve or surface that avoids abrupt "spikes," curves with "knees" that stray significantly from a linear function are no problem for the combinatorial power of fuzzy logic. That's because, no matter what the internal method of arriving at answers (which, by the way, mostly takes place on a digital computer), fuzzy logic involves mapping

points on one curve to points on another. That's the way the relationships of inputs to outputs of the system are defined, and that's how the system provides the results.

Design of a fuzzy system need not start with a well-defined mathematical model; it can start with a general idea of how the system should work. This idea may consist of defining input and output ranges, the shape of membership functions within those ranges and the beginnings of a rule base that defines behavior.

For example, you might specify that a steering angle for a vehicle is to be in the range of ±30°. That range would be normalized and mapped to the fuzzy system as a range of –1 to 0 to +1. Then, within that range, you have to define what constitutes a small negative, large positive, zero, or small positive angle. These are the individual membership functions that tell you that an angle of, say, –24° is 0.8 of the function negative large.

The Huntington Group's Brubaker points out two strategies for setting up membership functions. One, it's a good idea to group function sets closer together near where the action occurs. In something like a positioning system, this will result in finer control around the optimal point. Another strategy is to have functions overlap. "Overlap of input functions," says Brubaker, "guarantees that for any given input at least two and possibly more rules will fire. The combination of their actions will tend to smooth the executable output."

The next step is to construct a rule base made up of IF/THEN rules that will describe how the system is to work. The number of possible rules is a function of the number of inputs and their function sets. If a system has five inputs and each input has seven ranges ("very low," "medium-high," etc.) the input set will have 5^7 or 78,125 possible rules. Since the idea of fuzzy logic is to use a small number of rules, you want to strike a balance between having enough rules, some of which overlap, to adequately control the system, yet keeping the number of rules to a minimum. Having more than the absolute minimum number of rules—along with overlapping membership functions—helps ensure robustness. For example, a sensor could fail, removing a rule, and the system could probably still function acceptably.

Optimizing the program is the place where trial-and-error, simulation and plain old engineering are most important. Integrated Systems has developed a fuzzy logic module called RT/Fuzzy that fits into its SystemBuild product. SystemBuild is a large library of real-time control modules that can be combined hierarchically to construct complex control systems. RT/Fuzzy can, like other SystemBuild modules, use its animated simulation capabilities to verify designs. The RT/Fuzzy module in-

Simple language cloaks complexity

The power of fuzzy logic comes from its ability to combine immensely rich combinatorial operations with relatively simple and deliberately imprecise linguistic statements that can be expressed in plain English. Fuzzy logic expresses relationships by grouping large numbers of values into "fuzzy sets" described by adjectives such as "small," "medium slow," "warm," "hot," and "middle aged." Each such adjective covers a range of values, and the designer or programmer decides to what degree values falling within that range satisfy what's meant by the adjective. Thus, someone who is 6 ft 4 in. would satisfy the "tall" criterion 100 percent of the time, but someone only 5 ft 11 in. might satisfy the description of tall in only 50 percent of all instances. Each value's degree of membership in a set contributes to the final output value when the sets are combined and manipulated.

In the graph, therefore, heights of 4 and 6 ft show respective degrees of membership, as percentages, in sets named "short" and "tall." Having established these "membership functions," the designer may, on further thought, decide that the ranges and slopes of the sets don't represent what he or she really wanted to express with the terms "short" and "tall." All the designer has to do in this case is adjust the shape of the membership functions until the results meet requirements. He or she doesn't have to tediously rewrite an enormous set of rules.

The output of this set might overlap with the output of another set describing, say, weight—light, medium, heavy—to contribute to a combined output determining something such as "yards of material for a suit."

The combinatorial richness that's made possible by using fuzzy sets is shown in the graph. The two bottom axes are the ranges of two fuzzy sets. The surface in the Z direction shows all the points that result from combining the values of the two sets. Cursors set anywhere along the two axes will intersect at a point on the surface that is a resulting output value.

It's only possible to graphically represent three dimensions—the combination and output of two fuzzy sets at a time. But you can combine sets in many more dimensions, so that getting a clear picture of what's going on becomes impossible without a computer. The only way to intuitively appreciate what's in fact going on is through the linguistic expression of the relationships in terms of fuzzy adjectives.

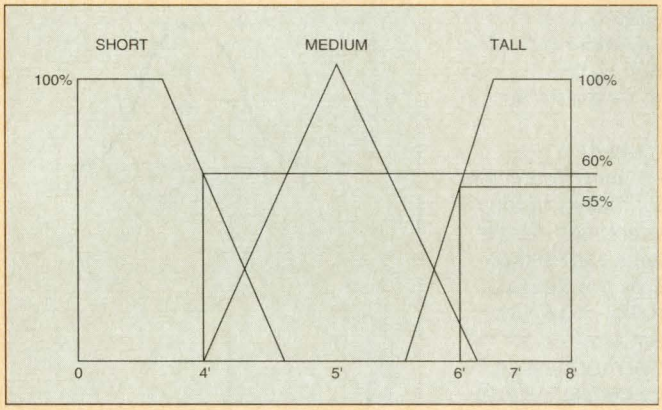

Shown here are fuzzy logic membership sets representing the heights "short," "medium" and "tall." A height of 4 ft has a 60 percent membership in the set "short," while a height of 6 ft has a 55 percent membership in the set "tall."

The seven noble truths of fuzzy logic

Truth one: There is nothing fuzzy about fuzzy logic.

The idea that fuzzy logic is fuzzy or intrinsically imprecise is one of the most commonly expressed fables in discussions of the subject. This widespread belief comes in two flavors: the first holds that fuzzy logic violates common sense and the well-proven laws of logic; the second, perhaps inspired by its name, holds that fuzzy systems produce answers that are somehow ad hoc, fuzzy or vague. The feeling persists that fuzzy logic systems somehow, through their handling of imprecise and approximate concepts, produce results that are approximations of the answers we would get if we had access to a model that worked on hard facts and crisp information.

Nothing could be farther from the truth. Fuzzy sets differ from classical or crisp sets in that they allow partial or gradual degrees of membership. We can see the difference easily by looking at the difference between a conventional (or "crisp") set and a fuzzy set. A fuzzy set can accommodate a bell curve showing a "normal" distribution of items falling within a membership category; a crisp set establishes firm boundaries which may lead to unrealistic yes/no distinctions.

In the example illustrated on the right, we must make a sharp distinction at the membership boundaries of the set using a crisp representation. Someone 34 years, 11 months and 28 days old isn't middle aged, by this definition. In the fuzzy representation, however, we see that as a person grows older he or she acquires a partial membership in the set of middle-aged people, with total membership at 40 years old. But there's nothing ambiguous about the fuzzy set itself. If we know a value from the domain—say, an age of 35 years old—then we can find its exact and unambiguous membership in the set—82 percent, in this case.

This precision at the set level lets us write fuzzy rules at a high level of abstraction. We can say, therefore, that: IF age is middle-aged, THEN weight is usually quite heavy. We mean by this that to the degree that the individual's age is considered middle-aged, his or her weight should be considered quite heavy. A weight-estimating function, following this very simple rule, might predict a weight from age through the fuzzy process shown in the illustration opposite.

Discomfort about fuzzy logic stems from the implicit assumption that a single "right" logical system exists, and if another system deviates from this it's in error. This "correct" logic is, of course, Aristotelian or Boolean logic. But, as a logic of continuous and partial memberships, fuzzy logic has a deep and impressive pedigree. Using the metaphor of the river, Greek philosopher Heraclitus aptly points out that a continuous reasoning system more correctly maps nature's logical ambiguities. From his dictum that all is in flux and nothing is stationary, he developed a rudimentary multivalued logic 200 years before Aristotle. Recently, Bart Kosko, one of the profoundest thinkers in fuzzy logic, has shown that Boolean logic is a special case of fuzzy logic.

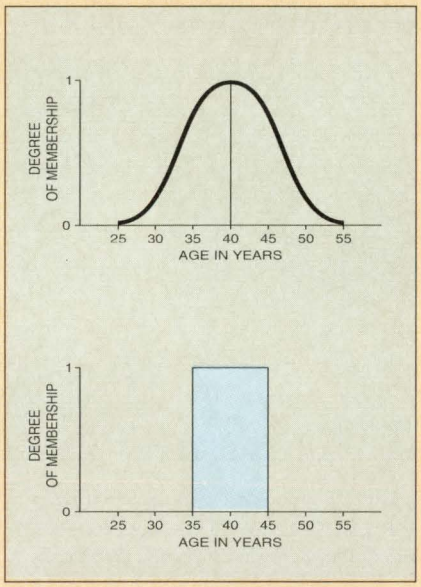

This bell curve graphically represents a "normal" statistical distribution of ages that might fit the membership class "middle-aged." Degrees of membership vary from 0 in the early 20s and late 50s to 1 at an age of 40. This fuzzy set is distinguished from a crisp representation of "middle-aged,", which classifies everyone between the ages of 35 and 45 as 100 percent in that class.

Truth two: Fuzzy logic is different from probability.

The difference between probability and fuzzy logic is clear when we consider the underlying concept that each attempts to model. Probability is concerned with undecidability in the outcome of clearly defined and randomly occurring events, while fuzzy logic is concerned with the ambiguity or undecidability inherent in the description of the event itself. Fuzziness is often expressed as ambiguity rather than imprecision or uncertainty, and remains a characteristic of perception as well as concept.

Truth three: Designing fuzzy sets is very easy.

Not only are fuzzy sets easy to conceptualize and represent, but they reflect, in a general "one-to-one" mapping, the way experts actually think about a problem. Experts can quickly sketch out the approximate shape of a fuzzy set. Later, after running the model or examining the process, the precise characteristics of the fuzzy vocabulary can be adjusted if necessary.

Truth four: Fuzzy systems are stable, easily tuned, and can be conventionally validated.

Creating fuzzy sets and building fuzzy systems can be done more quickly than working with conventional knowledge-based systems using "crisp" constructs. These fuzzy systems routinely show a reduction in rules of one or two orders of magnitude, since fuzzy logic simultaneously handles all the interlocking degrees of freedom. Fuzzy systems are very robust since the overlapping of the fuzzy regions, representing the continuous domain of each control and solution variable, contributes to a well-behaved and predictable system operation. Fuzzy systems are validated in the same manner as conventional systems. The tuning of fuzzy systems, however, is usually much simpler since there are fewer rules. Representation is visually centered around fuzzy sets, and operations act simultaneously on the output areas.

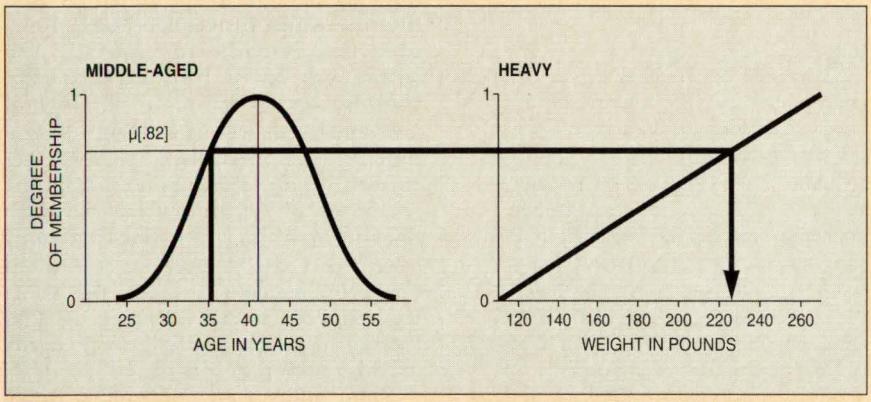

This example of graphical representations of the characteristics "middle-aged" and "heavy" shows how, by linking the membership sets, predictive value can be achieved. Here the average weight of a middle-aged person is predicted to be about 225 pounds with 82 percent probability.

Truth five: Fuzzy systems are different from but complementary to neural networks.

There's a close relationship between fuzzy logic and neural systems. A fuzzy system attempts to find a region that represents the space defined by the intersection, union or complement of the fuzzy control variables. This has analogies to both neural network classifiers and linear programming models. Yet fuzzy systems approach the problem differently, with a deeper and more robust epistemology. In a fuzzy system, the classification and bounding processes are much more open to the developer and user, with opportunities for explanations, rule and fuzzy set calibrations, performance measurements, and controls over the way the solution is ultimately derived.

Truth six: Fuzzy logic "ain't just process control anymore."

Historically, we've come to view fuzzy logic as a process control and signal analysis technique. But fuzzy logic is really a way of logically representing and analyzing information, independent of particular applications. The information management field in particular has ignored fuzzy logic, delaying its introduction into expert system and decision support technology. Recently, however, new types of knowledge-based construction tools have emerged. Such tools will make it easier for authorities who are not computer experts to intuitively represent and manipulateinformation.

Truth seven: Fuzzy logic is a representation and reasoning process—not the "magic bullet" for all of artificial intelligence's current problems.

Fuzzy logic is a tool for representing imprecise, ambiguous and vague information. Its power lies in its ability to perform meaningful and reasonable operations on concepts that are outside the definitions available in conventional Boolean logic. Fuzzy logic has been used in such applications as project management, product pricing models, health care provider fraud detection, sales forecasting, market share demographic analysis, criminal identification, capital budgeting, and company acquisition analysis. Although fuzzy logic is a powerful and versatile tool, it isn't a solution to all problems. Nevertheless, it opens the door for the modeling of problems that have generally been extremely difficult or even intractable.

Earl Cox, chief scientist, Knowledge Based Technologies, White Plains, NY

cludes a weighting scheme that can be applied to rules to adjust the amount of influence each has on the final results. Once the system looks like it's being controlled reasonably well, you can start playing with the weights to try to get the best answer. Says Houtchens, "Really, what's going to control things is how you tell which rules have the most impact on the final answer."

According to Lotfi Zadeh, there's still no theoretical or systematic method for eliminating unneeded rules or finding the optimum rule set, but this is an important area of investigation in fuzzy theory—an area for the art of engineering. You have to start with some knowledge and common sense, try out and simulate the design, add and modify rules and membership functions, and try again until you are assured that the system's operating satisfactorily.

There are, however, a number of methods beyond blind intuition that can help when designing fuzzy systems. One such method works very well when substituting fuzzy control for a system that previously required a skilled human operator. You take the knowledge of the skilled operator and construct rules based on how that operator used the human interface controls on the system you're now trying to automate.

Okumura Corporation, a Tokyo-based construction company, has applied fuzzy control to shield tunneling, in which a large, drum-shaped tunneling machine with a rotating cutting face is steered through the earth to dig subway, power and water tunnels. As the cutting face rotates, soil and rock are removed via a conveyor in the middle of the drum. As tunneling progresses, curved concrete sections are placed against the walls behind the machine to complete the tunnel. The machine is pushed forward by a set of 14 hydraulic "jockeys," or shafts, mounted radially around the rear of the machine. These jockeys push against the edge of the last completed concrete tunnel section.

Problems in controlling the speed and direction of the machine arise due to uneven hardness of the soil, which causes uneven pressure on the cutting face. A skilled operator is needed to monitor direction, pressure, cutter, speed and other factors, and to turn selected jockeys on and

NASA eyes range of fuzzy control ideas in space

NASA has recently reported encouraging research results in a number of areas using fuzzy logic according to an article in last October's *Aerospace America*". One of the most advanced projects is a controller for space shuttle proximity operations; that is, maneuvering around or keeping position with respect to another object in space. Work has been progressing on a fuzzy-based translational controller that deals with the parameters of azimuth and angle and their respective rates of change, and the range and rate of distance change with respect to another object.

NASA engineers developed natural language rules to run the controller and are testing it in a multivehicle simulation by substituting the fuzzy controller for the simulator's normal human inputs. The rule base was extracted from the experience of human operators and the efficiency of the controller was tuned based on flight profiles recorded from actual missions and simulations.

One of the main advantages in developing the fuzzy translational controller was that the engineers didn't need to construct a detailed mathematical model of the system in advance. Performance was honed through simulation and experience.

The results of simulations have been encouraging, especially in terms of fuel efficiency. In holding position with respect to another space craft, the fuzzy controller required significantly less acceleration—that is, smaller increments of position change—than did the human-controlled simulation. In overall maneuvers, the fuzzy controller has shown 20 to 70 percent better fuel efficiency than the currently used digital autopilot and the best simulation runs of human pilots.

NASA is also exploring other applications of fuzzy control in space. Among the projects being considered is the use of inexpensive cameras for constant tracking of objects around the space station. Fuzzy control can contribute to collision avoidance systems, robot arm control and traffic management. At the great distances found in interplanetary space, where it can take 20 minutes to send a signal and receive an answer, robotic systems will have to operate quasi-independently. A fuzzy controller on an unmanned Mars rover vehicle is expected to help the rover avoid obstacles and identify and collect soil samples based on imprecise sensor input and only partially known conditions.

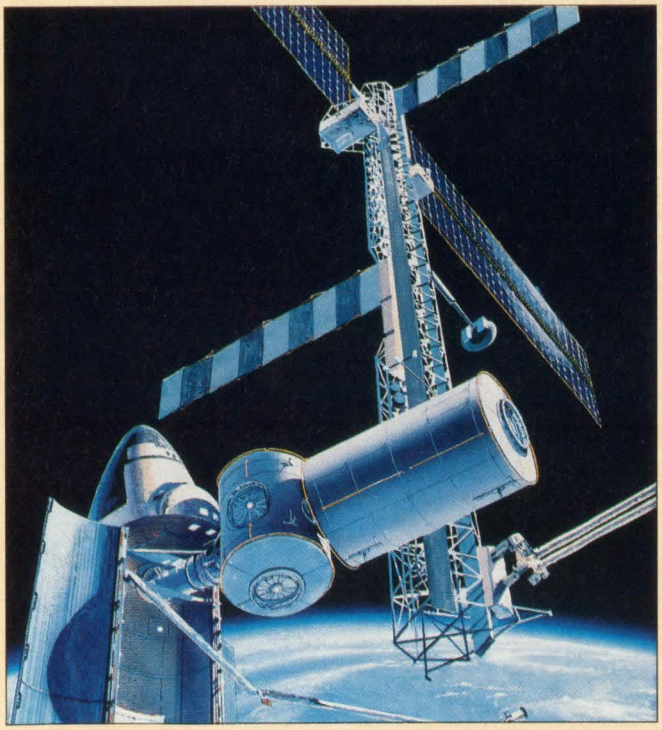

Fuzzy logic provides the ability to design systems that are much more efficient than those using conventional logic approaches. NASA found that, in simulations of space shuttle maneuvering and docking, a fuzzy logic controller performed much better than both a standard autopilot and experienced human pilots.

off to maintain the proper direction.

Using the knowledge acquired by skilled operators, Okumura was able to design a set of rules and membership functions that operated the machine in basically the same way as a human operator, thereby saving complete redesign of the control system to automate the machine. The results have been, according to Yasuhide Seno, assistant manager of Okumura's civil engineering division, "An improvement over even the results of a skilled operator, especially when negotiating curved tunnel sections." The tunneling machine can now be operated by a civil engineer.

Seno notes that the entire fuzzy logic program consists of about 10,000 lines of C code, which is mostly devoted to the graphical user interface. The fuzzy processing takes up about 1,000 lines of code.

Another method of characterizing behavior, especially when you're upgrading or modifying an existing system or converting an existing system to fuzzy control, is to gather operating data from that existing system. Integrated Systems has a tool called System Identifier that does exactly that. The data gathered by System Identifier can be fitted to curves and interpreted to define membership functions for use with fuzzy rules.

Yet another method is to interview a number of experts on a given process and use a mean of their responses—they need not exactly agree—to establish membership functions. For example, there might be some difference of opinion among experts as to what "cool" means in terms of jet engine exhaust, but an average of those opinions could establish a useful set of membership functions.

In fine tuning and verifying a fuzzy design, however, there's no substitute for simulation. Where classical control systems demand a lot of work establishing the model in advance of doing the design, fuzzy systems require a large amount of simulation and tuning to optimize their performance and verify their reliability. So, there's still no such thing as a free lunch.

Verifying a system's reliability has been held up as an obstacle to using fuzzy logic, especially in safety-related areas. This is becoming less of an im-

pediment, however, as engineers realize that the ability to mathematically prove the stability of classical control systems represents not proof but only a degree of confidence.

"What it means when they say proof of safety is really proof of stability, that the system will not oscillate wildly," says Fred Watkins, president of Hyperlogic (Escondido, CA), which manufactures and sells fuzzy logic and neural network development tools. "Nobody can prove something is safe."

"The only way the mathematical model fits reality is if all sorts of simplifying assumptions are made," continues Watkins. "Assuming that the only factors that work in a situation are this, this and this."

Of course, in the real world all kinds of other factors might come into play, so there's a philosophical leap that must be made to bridge the gap between an engineer's mathematics and reality. Claiming proof of reliability in an absolute sense—yes or no—is, therefore, impossible.

Advocates of fuzzy logic take a certain satisfaction in pointing out that the confidence level one gets from classical proof of stability is in itself a fuzzy value. They also argue that if a fuzzy system can produce the same control surface or matrix through exhaustive simulation that a corresponding classical control strategy can produce through equations, within the parameters established for the problem, both systems are equally stable.

Another factor in fuzzy reliability was described earlier as robustness. Overlapping membership functions and rules help assure that even if some input isn't presented, the system will operate. Robustness can be extended to the hardware level, notes Paul Basehore, vice-president of engineering for NeuraLogix (Sanford, FL), which offers fuzzy logic and software design tools and intergrated circuits. "With a conventional microprocessor, if you drop a bit the thing probably crashes. In a fuzzy controller, such as our highly parallel NLX230, you can completely eliminate a membership function or throw a rule away and it will still perform the functions to get to the end goal."

The basic activity of developing a fuzzy application consists of "drawing some graphs and typing in some English rules," as Hyperlogic's Fred Watkins puts it. Any fuzzy development tool worthy of the name is oriented toward helping you do just that. Of course, tools vary quite a bit in the range of features they provide, the kinds of membership functions and methods they support, the range of "concept-to-code" completeness they support, and, finally cost. For example, a tool might depend on third-party or user-supplied animation to aid simulation, or output its simulation results as static graphs. Another might allow definition of basic fuzzy input and output relationships, but not supply a debugger. All these considerations are less criticism than indications of cost vs. feature trade-offs.

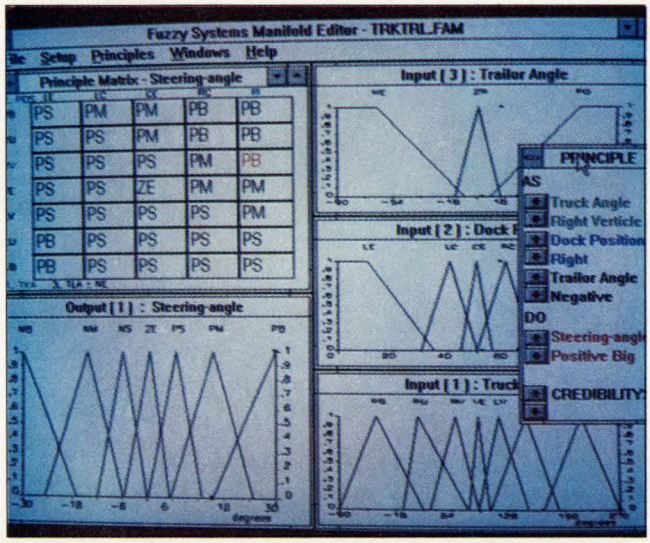

The Fuzzy Manifold Editor from Fuzzy Systems Engineering can represent a fuzzy associative memory (FAM). Here, a slice through the matrix is shown, with an output fuzzy set designated for every intersection of input sets. Also displayed are the membership function graphs for input and output values. Note how the membership functions are grouped more closely around the point of control.

Ease of use for a broad but well-defined range of practical control applications is the aim of the fuzzy microcontroller, development system and associated hardware components from NeuraLogix. NeuraLogix's Basehore advocates introducing a new technology such as fuzzy logic in a way that can add useful features to a product while maintaining a cost-effective, practical implementation.

The company's development system, therefore, doesn't try to cover "complete fuzzy logic theory." Rather, it contains features aimed at controller design, is PC-based, includes an AT bus board with an NLX230 FMC (fuzzy microcontroller) that lets you test your application right at the development workstation and cost $395.

Notes Basehore, "One of the seemingly difficult tasks of fuzzy logic design is in determining the optimum membership functions," or, to put it another way, in relating input sensor values to output control responses in a nonlinear way. "The shape of the membership function can be determined empirically or through adaptive methodologies," he continues, "or through algebraic means which could be quite complex. These shape determining methods can be time-consuming and require at least general knowledge of fuzzy logic. In most applications, however, the actual shape of a membership function is less important to the total solution than the degree of overlap of associated membership functions."

Basehore argues that, instead of concentrating on determining the shape of the optimum membership function, you might better spend your time selecting the more important center and width values. NeuraLogix has developed a number of devices that provide several fixed shapes that you can choose from and then select center and widths. The development software environment displays the functions and records the values selected. Since the fuzzy microcontrollers are true hardware devices, no further programming of these functions is needed.

Another low-cost tool, called the Fuzzy Systems Manifold Editor from Fuzzy Systems Engineering (Poway, CA), lets you define and check out

Making the interface transparent to sub-nanosecond rise times.

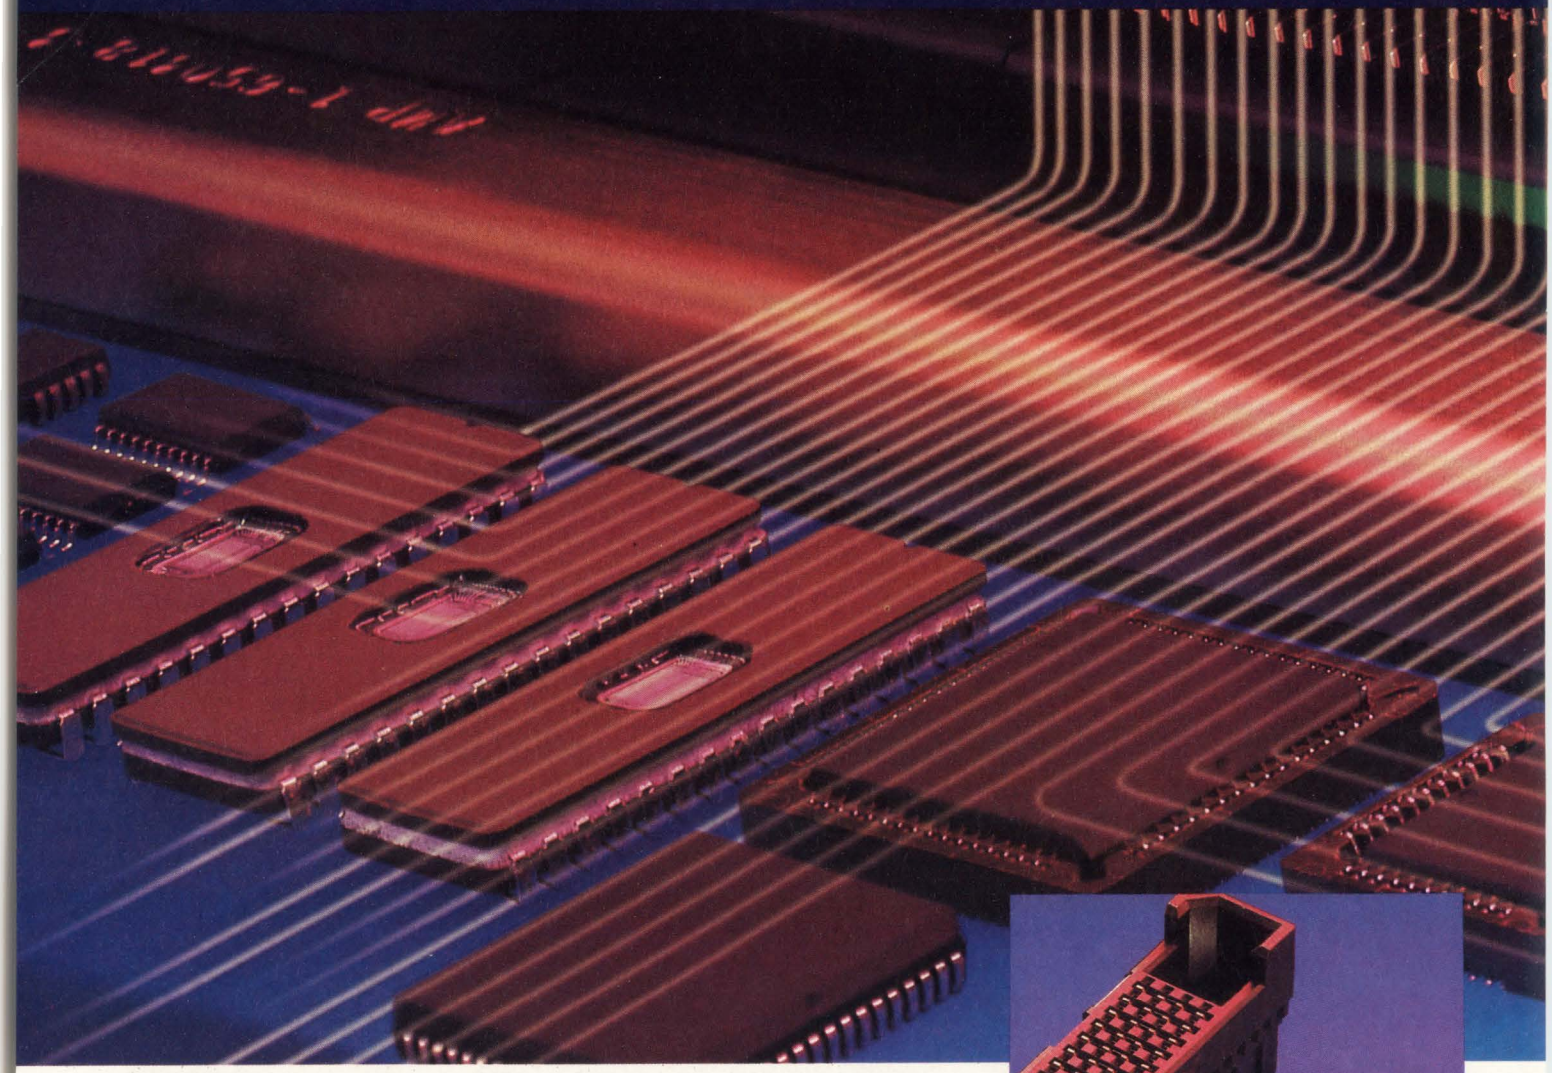

THIS IS AMP TODAY.

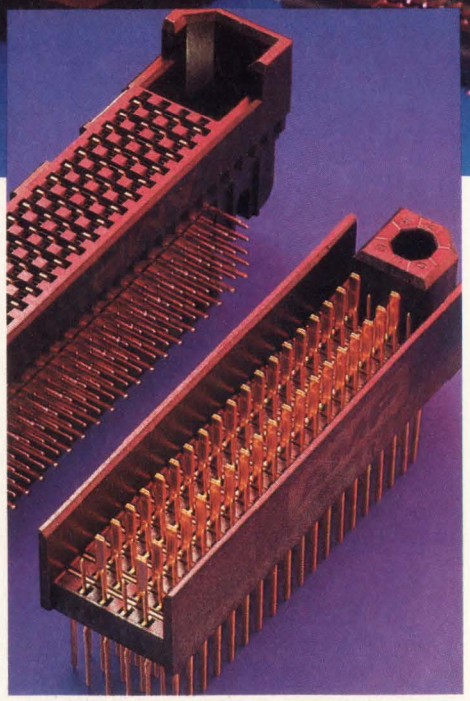

Stripline high-performance connectors.

AMP and ACTION PIN are trademarks of AMP Incorporated.

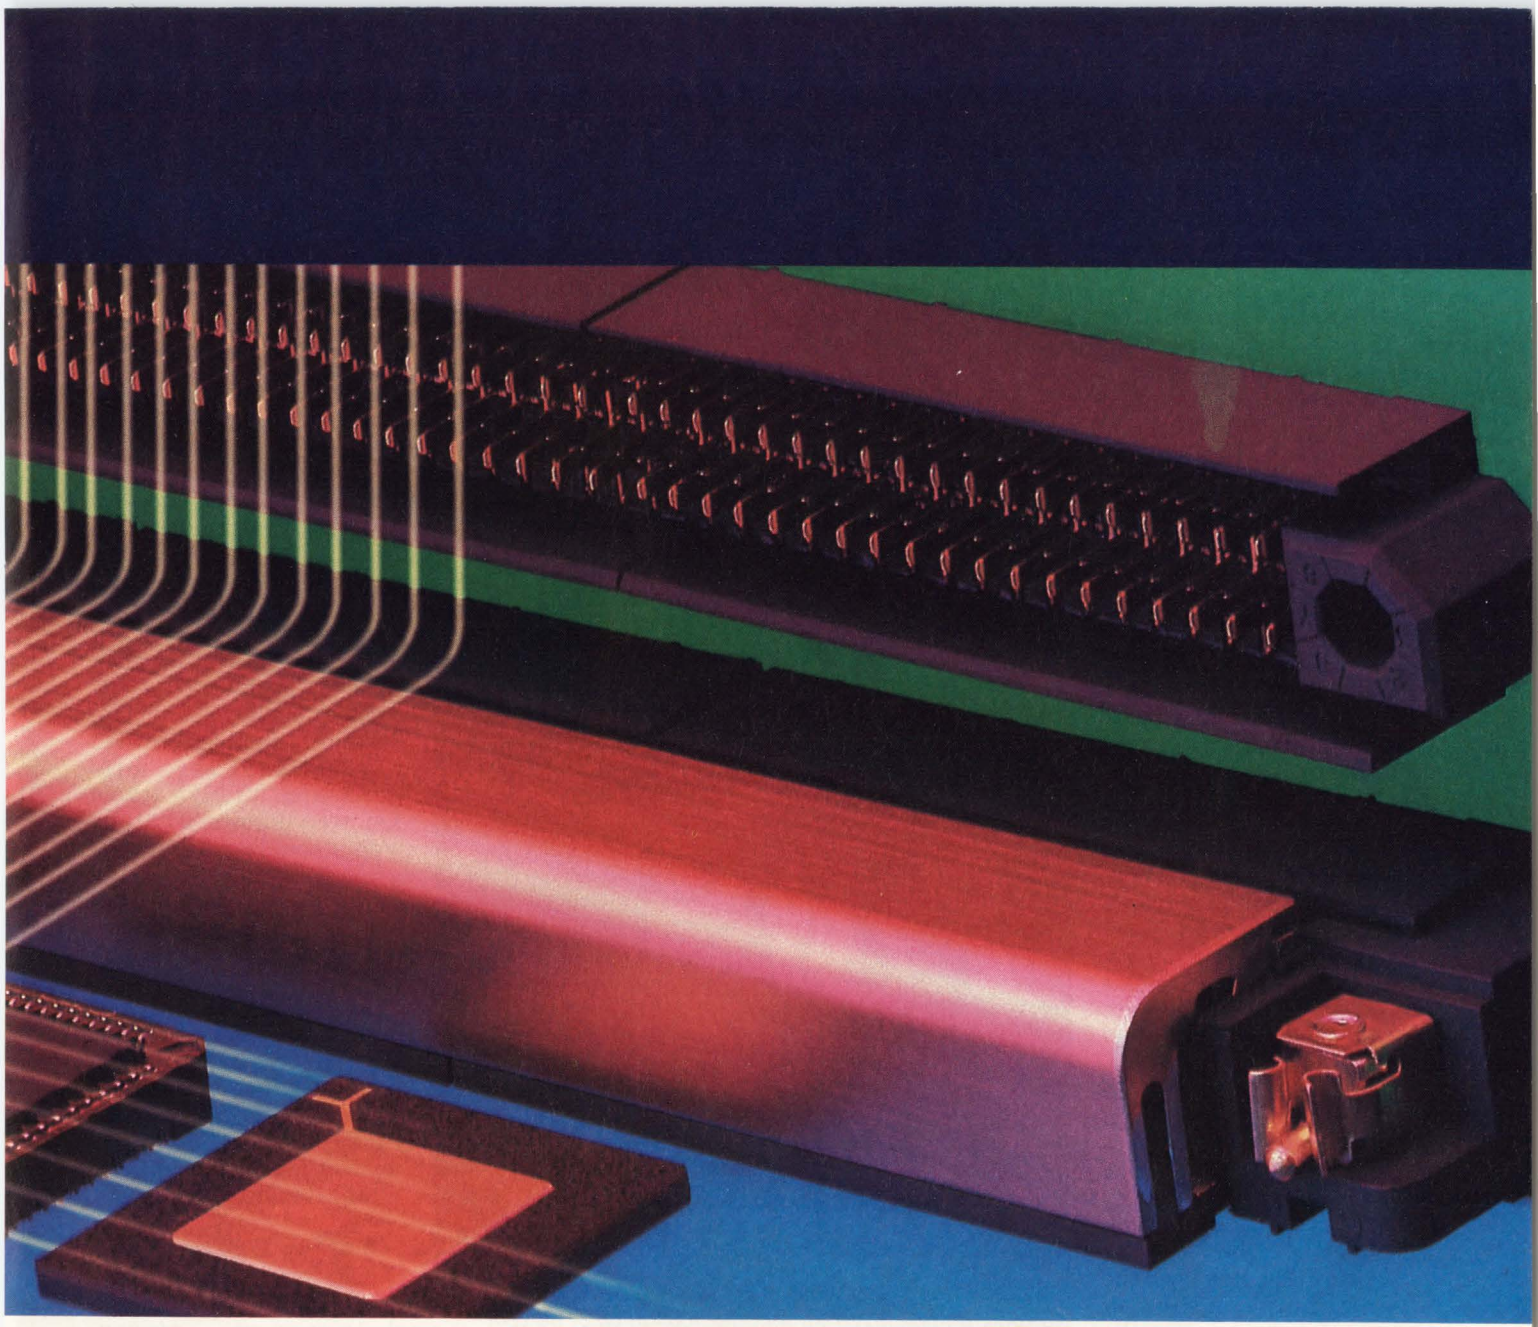

'Fast silicon' (rise time≤1 ns) requires strict impedance control. Conventional connectors give up half their pin count for this – a sacrifice you can do without.

Our modular, scalable Stripline 100 connector system can accommodate edge rates of 250 ps (500 ps at <3% crosstalk), and still give you 40 signal lines per inch – all four rows on a .1"x.1" grid. Reference planes isolate individual signal columns within the standard grid geometry, creating an interface completely transparent to high-speed logic.

Stripline 100 connectors deliver more than raw speed, too. Each reference plane can distribute three amps, and sequenced mating is available for ground, power, and two signal levels.

Manufacturing is easier as well. ACTION PIN compliant posts (for existing 0.040" pcb holes) simplify backplane assembly, and all materials are compatible with high-temp reflow processing.

In fact, sub-nanosecond logic just got easier all around, and there's an easy way to 'bring yourself up to speed' on this exciting technology: call our Product Information Center at 1-800-522-6752 (fax 717-986-7575). AMP Incorporated, Harrisburg, PA 17105-3608. In Canada call 416-475-6222. For design assistance in characterized backplane assemblies, contact AMP Packaging Systems, 512-244-5100.

AMP

CIRCLE NO. 64

the core of a fuzzy application in a Microsoft Windows environment. The concept of a manifold refers to the fuzzy associative memory (FAM), an n-dimensional combinatorial matrix that represents all the relationships between fuzzy input and output values. Any given display of this matrix is a two-dimensional "slice" through it, which lets you view two input dimensions—that is, the combinations of the membership functions making up two inputs, and the resulting matrix of outputs. You can successively select different slices to get an overview of the combinatorial richness of the design.

By clicking on the boxes in the output matrix, you can select which output membership functions will result from which combination of inputs. Therefore, for the hypothetical inputs of some system—which controls temperature and pressure with a heater, for example—you could select the box where the sets "temperature medium" and "pressure low" intersect and specify that heater output will fall in the "medium" membership function. The matrix represents the combinations from which IF/THEN rules can be written. To emphasize the continuous nature of fuzzy logic, however, the manifold editor expresses the relationships as AS/DO.

The $295 manifold editor lets you graphically edit membership functions and examine the system's behavior by simply setting input values and reading the resulting outputs. It doesn't support animation or true simulation, but lets you specify up to five input dimensions and two output dimensions with up to seven membership functions each. The matrix can be displayed, along with the membership function graphs. This lets you select output functions, change or edit the graphs that define the input functions, and then click to a manifold "walker." Here you can enter individual input values, click on a button and read the resulting output. The manifold editor can output a source-code include file in C, Fortran or Basic that can be built into other programs. The manifold editor doesn't support writing fuzzy rules in English however. That and the actual inference engine operations—interpreting the rules—are left to the application program.

Still another low-cost package, in versions ranging from $495 to $795, includes the ability to write English rules as well as necessary preprocessing and postprocessing operations. CubiCalc from Hyperlogic runs under MS-Windows. It lets you express rules as IF/THEN statements and also to use modifiers such as "very hot" and "somewhat heavy" with your variables. Graphical function set editing is supported, and there's also the ability to do some simulation, thanks to CubiCalc's preprocessing and postprocessing capabilities.

There's more to a fuzzy system than fuzzy logic operations alone, however. You've got to take in data from sensors or deal with other data in real-world formats—for example, angles, miles per hour or parsecs. These items have to be normalized and mapped into fuzzy sets. CubiCalc provides a preprocessing language to let you do this; there's also a postprocessing capability that lets you present data to a display, a larger application program or other output. This language also supports graphic simulation.

CubiCalc can output to on-screen scatter charts or strip charts so you can view the results of operations. Its preprocessing capability lets you set up data files or ranges of variables to use as inputs and then view the results in a graph on-screen. For a more precise examination of operations, you can log your output to a file and view it as a numeric display or you can output it to another program for further processing.

Hyperlogic president Fred Watkins sees the potential for such tools as CubiCalc as extending far beyond control systems, making possible natural language programming. "A high school kid who knows what he wants could use that program to build a nontrivial system," he says. "I think the linguistic aspect of things is really where it's at. We can expect machines in the future that will come knowing triangles, knowing what 'big' means. I can imagine a machine with a repertoire of fuzzy functions tucked away, knowing how to communicate with you."

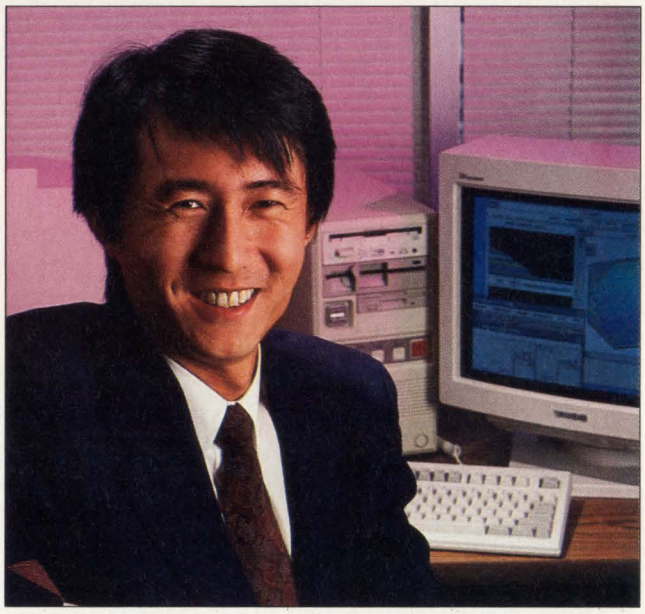

Aptronix president Wei Xu sees the application of his company's Fideé toolset going far beyond embedded control. "Computer programmers are not application engineers," he says. "We're trying to provide an engineering language for computer-aided algorithm design to give knowledge designers and experts the tools—such as CAD/CAE—that will let them sculpture their algorithms."

Two U.S. companies have produced ambitious CASE-like development environments for fuzzy logic that include design, simulation, debugging, compiler, and linking facilities. They are Aptronix (San Jose, CA) and Togai Infralogic (Irvine, CA). Aptronix has just introduced its Fuzzy Inference Development Environment, called Fidé, which Aptronix president and CEO Wei Xu describes as a tool for "computer-aided algorithm design. I'm trying to provide an engineering language instead of a language for software engineers. If I'm developing a car transmission control, I'm a transmission specialist. Do you expect me to be a C specialist? So who's going to do the algorithm programming, because all the algorithms must be embedded in code?"

The traditional alternatives are that either the transmission special-

ist must try to code the application, which will be less than wonderful because he or she is just learning C to help do the real job. Or, the specialist can let a software engineer do the code, after writing all the specifications while carrying on a dialog with the software engineer. "Software engineers are interested in how computers are programmed; they don't care about shifting gears," Xu says.

Fidé, priced at $1,495, consists of an editor for text and graphics and another for editing membership functions; static and dynamic simulators; a debugging tool called Tracer; compilers; and a graphical linking facility for linking fuzzy modules with other fuzzy and nonfuzzy code. Fidé offers different levels of complexity for implementing fuzzy designs. For example, in addition to the familiar triangle and trapezoid membership functions commonly used in control applications, lookup table and operator functions can be selected. (The operator function is an equation for a curved function.) Different kinds of operators can be used to logically combine membership functions such as "minimize" and "maximize"—the traditional AND and OR as well as 'probability sum and product' and 'bounded sum and difference,' which are variations of the AND and OR operators.

Fidé's static analyzer displays a three-dimensional response curve surface representing the combination of two input membership functions, along with side views of each function. You can move a crosshair cursor anywhere on the surface and see what the actual input values are to the membership functions that generate that point on the surface. If the surface shows a radical nonlinearity, for example, such as a spike that may indicate a bug, you can click directly into Tracer and see the rule that generated the point under investigation. From the rule, you can go on to examine the source code and look at the specific membership functions involved, all without losing your place in the analyzer.

Togai Infralogic, a pioneer in easy-to-use fuzzy development tools, has been offering tools since 1988, the most prominent of which is its TILShell object-based development environment, which runs on MS-DOS (Windows), Macintosh and Unix workstations. It starts out

Making a fuzzy decision definite

Fuzzy decisions are made using plain-language rules based on definitions of membership functions. Rules are stated in a simple IF/THEN format; values can be combined using AND or OR operations. By far the most common operator is the AND, which requires that the input value lie within all the sets stated by the rule.

For example, the degree to which a gas valve should be opened to maintain a vessel of water at a desired temperature might depend on the volume of water in the vessel and its current temperature. This could be stated as follows:

Rule 1: IF temperature is cold AND volume is large, THEN valve is wide

The exact value of "wide," called its crisp or "defuzzified" value, is determined by the degree of membership the actual or crisp values of "temperature" and "volume" have in their respective sets of "cold" and "large." Since an AND rule applies, the minimum value—the 75 percent membership which fits both input sets—would determine the degree of membership in the output set, "wide."

One of the most popular ways of extracting the final "crisp" output value from the resulting fuzzy set is the center-of-gravity or centroid method. The centroid method finds the line that divides the resultant shape—in this case, a trapezoid—into equal areas.

But suppose now that another rule considers the pressure inside the vessel in determining how wide to open the valve. It might be stated as:

Rule 2: IF pressure is high THEN valve is medium.

The output of this rule would fall into another membership function in the group associated with "valve."

The two rules have given different values for the valve setting that must be resolved or "defuzzified" to yield a crisp final output value. Here, as with the single output function, the centroid method can be used to find the center of gravity of the combined output shape, ensuring that both rules contribute to the final crisp result.

A system designed this way can be fine-tuned by adding rules and membership functions to cover other factors, and also by changing the shape of the membership functions within existing groups of functions without reworking a large set of rules and values.

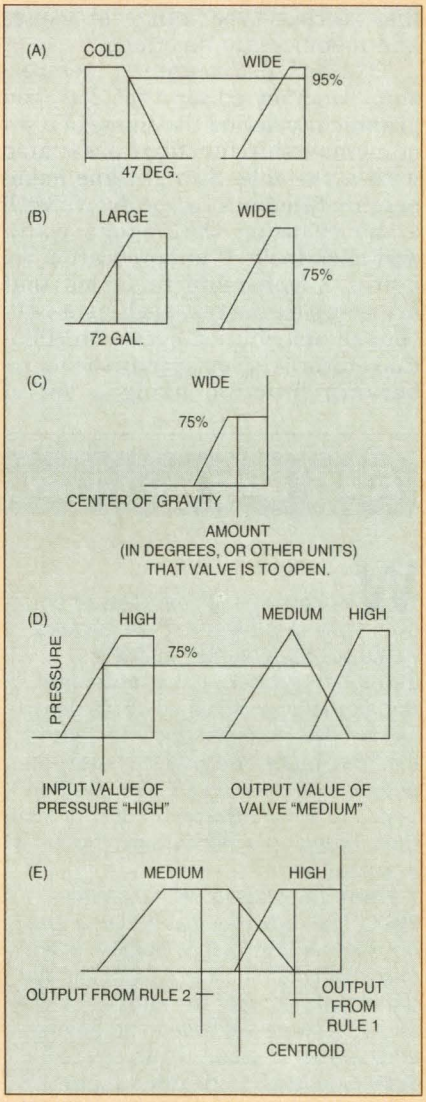

This diagram illustrates the graphical solution of a fuzzy control problem. Part (a) relates water temperature to the degree to which a gas valve is open. Part (b) relates container volume at 72 gallons, the membership set is 75 percent. The combined output, represented by part (c), takes the form of the volume membership set, since the AND rule calls for the minimum value that fits both temperature and volume sets. Adding a second, or pressure, rule to the temperature/volume rule gives the result shown in part (d). The output value of the rule, is shown on the right. In part (e) the graphical outputs of Rules 1 and 2 are combined. The vertical line between these two represents the combined solution to the problem.

with a facility that lets you graphically define the structure of a fuzzy system using input, output and processing objects that can then be assigned attributes. Attributes can define storage type, range of values and membership functions.

TILShell also contains a membership function editor that lets you graphically define the shape of a set of membership functions associated with a variable. Straight-line membership functions can be drawn with a mouse using the editor's point-and-click tools. If an application requires membership functions that are complex curves, such as a bell-shaped distribution curve, an equation editor lets you specify the membership function using a set of mathematical functions. TILShell produces files in Togai's Fuzzy Programming Language (FPL) and works with the Fuzzy C Development System (FCDS), which includes a C code generator which translates FPL files defined in TILShell into C. It also includes C editors and debuggers.

Recently, Togai has added TILGen to its repertory: this tool uses neural network technology that lets you automatically generate a basic fuzzy rule set. This is a useful tool for converting existing systems to fuzzy control. It can also be used to speed up rule generation for a system whose behavior is fairly well defined in terms of inputs and outputs and for which you want to extract a useful set of rules from the very large number that are theoretically possible. The rules generated characterize the control of the system based on the inputs and outputs.

To generate a rule base with TILGen, you input a description of the fuzzy object as an FPL format file, along with a file containing input and output vector data—real input and output values from an existing system or values expected from a system under design. The FPL file can be produced by defining inputs, outputs and their respective membership functions in the TILShell tool, but when used with TILGen the file won't contain any rules. TILGen applies the two files to a neural network which it sets

Fuzzy mug search helps cops catch crooks

Well known to all devotees of detective shows is the scene where a crime victim sits for long hours paging through mug books to try to identify some evildoer. Even when police departments have managed to computerize their databases of known criminals, the process of narrowing the identification search to a manageable number of mug shots based on a witness's description is a tedious one.

Knowledge Based Systems (White Plains, NY), however, has added a fuzzy front end to the image database of a major European police department that significantly reduces the number of look-throughs witnesses have to do before finding a set of pictures they can seriously work with to identify a subject.

In the past, if someone came in and said, "He was kind of tall and heavy-set and looked rather young," the police would have only a vague idea of what group of pictures to start showing the witness. Even if the system were computerized, someone would have to decide where the cutoff points were for "rather young" and "tall." If the person were described as 6 ft 1 in. but was really 5 ft 11 in., a system with a six-foot breakpoint might not catch the out-of-range number even if other factors in the description pointed to an overall match.

In addition to implementing a front end that uses fuzzy sets to describe characteristics such as "old," "thin" and "tall," Knowledge Based Systems built

in what it calls "perspective shifting" and "semantic plies." Perspective shifting changes the shape of the fuzzy set representing "tall" if the witness is, for instance, a 16-year-old girl or is Japanese. This quality lets the system search for "old" from a "young" perspective—or, for that matter, for any other characteristic as modified by the preconceptions of the searcher. Semantic plies perform a similar function, adjusting descriptions such as "tall for women" or "heavy for Samoans."

Successful searching using perspective shifting and semantic plies depends on the fact that the degree of reliability of a fuzzy concept is predicated on the witness's own characteristics; the system doesn't make crisp distinctions. A height of five feet, therefore, will have a greater degree of membership in the "tall" category for a 10-year-old than for an adult. Mug shots for a witness to look at are selected based on an overall degree of truth stemming from the combined described characteristics—in the instance described above this was set at 0.38.

According to Knowledge Based Systems, the old manual system required an average of 16 look-throughs to find a set of pictures that a witness could actually work with, while the fuzzy system reduced the number of look-throughs to two.

up on the host computer and outputs a new FPL file that also contains a rule base. You can then examine the generated rules and write additional overlapping rules for increased robustness if desired. A full system consisting of TILShell, FCDS and TILGen is priced at $5,000.

As the use of fuzzy logic becomes more widespread, especially in consumer products, optimizing cost and performance will become ever more important. Companies such as Togai, NeuraLogix and Omron (Kyoto, Japan) have been offering parallel fuzzy processors for several years. Because of their parallel architectures, such processors offer very high speed for fuzzy inferencing, as well as bringing robustness to the hardware level of a fuzzy design.

Fuzzy processors usually have to work as coprocessors with a general-purpose microprocessor, which is needed for preprocessing sensor inputs, handling keyboards and supporting other input/output. The race is on to lower the cost and increase the integration of such chips. NeuraLogix is aggressively pushing into the low-cost end of the market, and, recently, Omron and NEC (Tokyo, Japan) announced a joint effort to integrate functions of the Omron FP3000 fuzzy processor on an NEC microcontroller chip.

"There are many, many applications where software on a microcontroller will be plenty fast enough," points out Steve Marsh, director of strategic operations for Motorola's Advanced Microcontroller Division (Austin, TX). Fuzzy logic processing is always parallel in concept, if not in physical reality. That is, all the rules have to be evaluated before a result is reached. If the processing is done sequentially, then this "parallelism" just takes longer. Hyperlogic's Fred Watkins says, "No matter how big the system, if you have enough processors, you can render it [fuzzy logic] in three steps: evaluate memberships, combine the rules and defuzzify the output."

To meet the time budget of a control application, you've got to consider processor speed, number of rules and fuzzification/defuzzification methods. "If you can do it in software and it only takes a couple of hundred bytes, you're done," says Marsh. "If it doesn't work, it's probably because you've run out of time" Moving to a somewhat faster processor can often solve the problem.

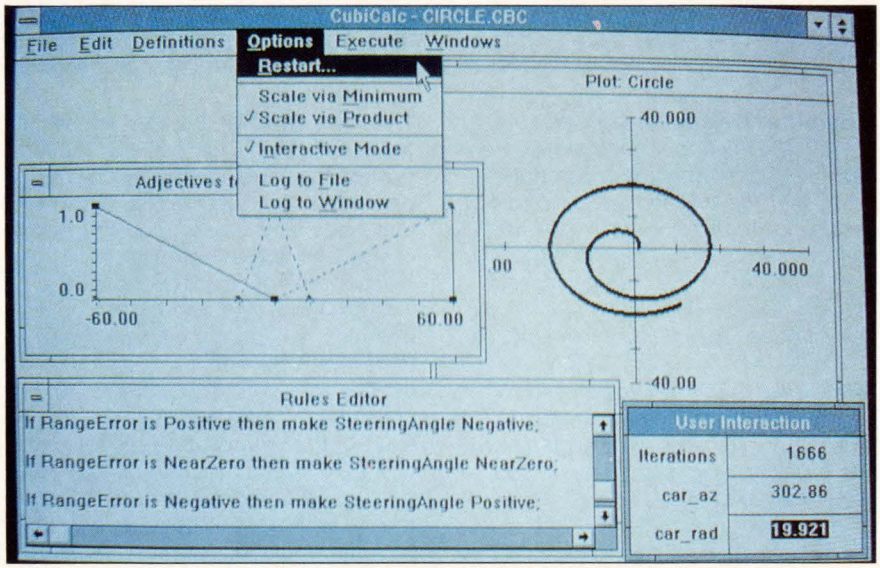

This screen shows graphical simulation from Hyperlogic's CubiCalc. The plot shows the path taken as a car is smoothly steered toward an objective. Fuzzy rules expressed in simple English appear in the rule editor. CubiCalc also provides a user interaction window where one could enter a new arbitrary starting point for the car, then observe its behavior on the graphical plot.

For example, fuzzification and defuzzification entail mapping an input value to a point on a curve and vice versa. Linear interpolation takes longer than finding the point in a lookup table, which takes more memory. So designers of high-volume products will want to look at speed-versus-cost issues. Using complex curves as membership functions is still popular in some noncontrol applications, but it's almost completely out of favor in control applications because of the complex formulas and processing overhead involved.

Motorola will be jointly marketing the Aptronix Fidé development environment, which directly supports Motorola's 8- and 16-bit microcontrollers. Rather than design specialized fuzzy processors, Marsh says Motorola will gradually add processors to its families that will implement certain customized instructions to speed up fuzzy operations such as fuzzy inputs, min/max operations for rule evaluation and defuzzification of outputs.

Motorola and Aptronix are also jointly working on an intermediate data representation and syntax standard for fuzzy systems that they plan to put in the public domain. This representation language will be similar to assembler code and will let tools output a text file that can be read by a code generator to produce processor-specific object files. Such code generators could be written for specialized fuzzy processors as well as conventional microcontrollers.

The possibilities for using fuzzy logic in practical expert systems derive from the advantages of incorporating skilled operator knowledge in a control system. Earl Cox, chief scientist for Knowledge Based Technologies (White Plains, NY), says that most expert systems today simply use von Neumann pattern matching. His company, a software applications development and consulting firm, took a traditional artificial intelligence insurance underwriting application that incorporated 181 rules and rewrote it using only 12, says Cox. "When you write fuzzy rules, you're writing at a much higher level of abstraction," he says. That means the fuzzy rules can be expressed without having to explicitly pick up every possible combination.

"In a conventional rule-based system," says Lotfi Zadeh, "a rule either fires if the conditions are met or it doesn't fire if the conditions are not met. Since there are many possible outputs, you must have a large number of rules to meet every contingency. Furthermore, such systems are not robust because, if a condition

The continuing evolution of fuzzy logic

Computer Design *interviewed Lotfi Zadeh, the inventor of fuzzy logic. He's now professor emeritus of computer science at the University of California at Berkeley and director of the Berkeley Initiative in Soft Computing. In 1965, Zadeh published the original paper describing the theory of fuzzy logic. Today, he continues to develop his theories.*

CD: *Today fuzzy logic appears to be most widely used in control applications, but still seems to be having trouble gaining acceptance. How do you view the situation?*

Zadeh: We have to realize that it's very natural for people, including myself, to be skeptical when they're presented with something that claims to provide a different way of looking at things. In 1965, my expectation was that most applications would be in the realm of "humanistic systems," such as linguistics, social sciences and biological sciences, where hard mathematics doesn't seem very effective. But then we began to see that fuzzy logic could be used in control. In control it's said that people want rigor and respectability. But then there are many realistic problems that cannot be rigorously defined.

Fuzzy algorithms for control policy will gain increasing, though perhaps grudging, acceptance because conventional nonfuzzy algorithms cannot in general cope with the complexity and ill-defined nature of large-scale systems. Control theory must become less preoccupied with mathematical rigor and precision and more concerned with the development of qualitative or approximate solutions to pressing real-world problems.

CD: *What do you tell people who express doubts about the reliability and stability of fuzzy systems?*

Zadeh: In the case of [classical] control systems, we do have a theory of stability. And, presumably, that theory can tell you that a certain kind of system will be stable. But that's much less significant from a practical point of view than one might think. Once you read the fine print, you find that what the theory can tell you is much more limited. It can tell you [things] if you linearize and if you do all sorts of things under certain assumptions. The trouble is it's very difficult to say whether those assumptions hold or not. So, you're left with something that isn't really comforting. You can't really sleep safely if someone using classical theory tells you that some control system is stable.

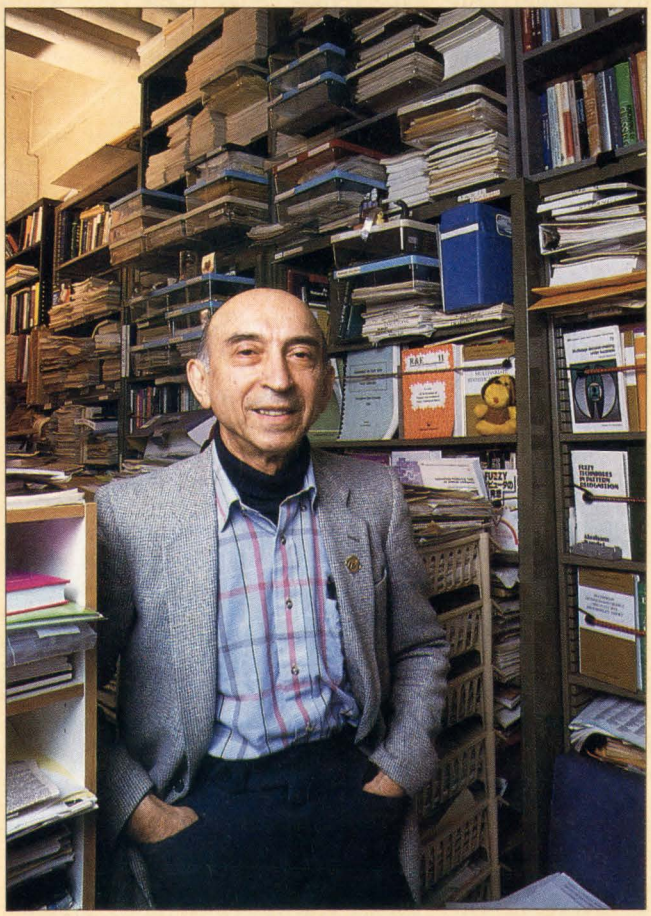

Fuzzy control is coarse control that exploits the tolerance for imprecision. So, if there's some imprecision and if the imprecision can be tolerated, you try to take advantage of it by making the system more robust and less susceptible to deviation. But still it's correct to say that at this point we don't have a theory for stability of fuzzy logic control that's nearly as well developed as for classical systems. Stability theory is really effective when it comes to linear systems, and fuzzy systems deal with nonlinearity.

In the case of fuzzy control, the systems are very complex. In many cases you cannot describe really what they do, so it's difficult to prove or disprove stability. People compensate for this with simulation. They perform many, many trial runs. In the case of the subway in the city of Sendai, Japan, I think there were some 300,000 simulations and 2,000 actual runs to prove the system, because you don't play with a subway system. So, I think the fact that the Sendai subway system has functioned perfectly since July 15, 1987, is a stronger testimony than theory.

CD: *Is the choice, then, between devoting a lot of time to establishing a mathematical model for classical control in advance, or, in fuzzy logic, designing the system and then proving and refining it in simulation?*

Zadeh: I think you put it well. The test of any theory is the ability to predict. So, if you cannot predict what will happen, you don't have much of a theory. Many so-called theories flunk this test, particularly in economics. In fuzzy systems, instead of performing some sort of analysis on paper or on computer that will predict how the system will behave, you simulate. So simulation is an alternative to prediction. It's not as desirable, but in the final analysis it may be more reliable.

There's always a possibility that your theoretical analysis didn't take certain things into consideration. Software is a good example. In the final analysis you have to run the program. Only actual use will tell you if there are bugs in the program or not.

Fuzzy logic will make cars smarter

It's estimated that by the end of the century, the electronic component in automobiles will double. Among the areas of interest to fuzzy-based system designers are an emission control system that would continuously monitor the exhaust gases and make adjustments to the carburetor and ignition system to keep levels of hydrocarbons and other toxic elements within specified limits. Interior climate control could adjust to the number of passengers, just as some fuzzy-controlled air conditioning systems do today. Fuzzy-controlled digital signal processors are being investigated for use in interior noise cancellation systems.

Today's digital antilock braking systems work by cycling the brakes rapidly on and off. A fuzzy controller would provide smooth antilock braking by adjusting hydraulic pressure in response to slight variations in wheel rotation.

One fuzzy item already under active development is a multiple-mode automatic transmission. By sensing rotation, torque, engine speed, and throttle position, a fuzzy controller could determine the proper gear shift points. By selecting a different set of rules and membership functions, the driver could set the transmission for smooth shifting, optimal economy shifting or sport shifting.

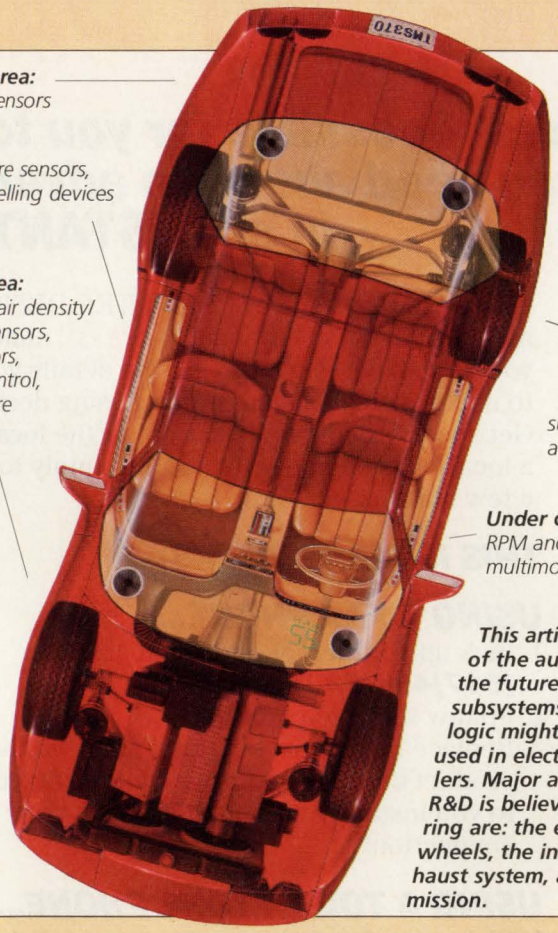

This artist's rendering of the automobile of the future shows those subsystems where fuzzy logic might be profitably used in electronic controllers. Major areas where R&D is believed to be occurring are: the engine, the wheels, the interior, the exhaust system, and the transmission.

is violated by even a small margin, the rule is not fired, even though that condition influences the situation. We know there are situations in the real world where the introduction of these artificial cutoffs results in all kinds of nonsensical situations. Something very drastic happens if you miss by even a small amount."

The other major advantage of using a fuzzy approach to expert systems, Cox notes, is that it greatly simplifies knowledge acquisition and representation. "Experts really like to talk in terms of fuzzy sets, whether they know it or not. It's just this high level of abstraction that lets an expert be an expert," Cox says. "Fuzzy logic reduces the cognitive dissonance between the machine representation of a problem and the way an expert thinks about it."

In the example of the insurance application, an actuarial table relating risk to age could be simply drawn as a U-shaped curve with high peaks at young and old. That's a fuzzy set because that's the way insurance people think about it. By writing a rule combining age with number of accidents and violations, Knowledge Based Technologies was able to reduce the number of rules for one part of the problem from 30 or 40 to one. "Because we didn't have to bracket all the cases," says Cox, "the fuzzy system did it for us by combining the fuzzy sets."

It turns out, then, that going from crisp rules to fuzzy rules buys you a lot. There are many situations where fuzzy IF/THEN rules provide you with a kind of language—that a system designer needs to express what he or she wants. A small number of rules really covers an enormous amount of combinatorial complexity. Says Zadeh, "Some people who are opposed to fuzzy logic don't realize that this is what it really is." ■

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

Aptronix (408) 428-1888	**Circle 201**
Fuzzy Systems Engineering (619) 748-7384	**Circle 202**
Huntington Group (415) 325-7554	**Circle 203**
Hyperlogic (619) 746-2765	**Circle 204**
Integrated Systems (408) 980-1500	**Circle 205**
Knowledge Based Technologies (914) 682-2139	**Circle 206**
Motorola (512) 891-3938	**Circle 207**
NeuraLogix (407) 322-5608	**Circle 208**
Togai Infralogic (714) 975-8522	**Circle 209**

Computer Design's

"INSTANT DATA ACCESS"

Makes it possible for you to get the INFORMATION you need about an advertised product NOW... in an INSTANT— by FAX!

If an ad in this issue of COMPUTER DESIGN includes an *"INSTANT DATA ACCESS"* number, you don't have to wait to get the details you need to make critical design or purchasing decisions. Get specs, pricing information, or the location of a local sales engineer or rep...by simply touching a few buttons.

THIS IS ALL IT TAKES—

USING A FAX PHONE...

1. Pick up your fax phone and dial ***(617)494-8338***.
2. Follow the voice prompts.
3. Replace the handset in the cradle and the product data you requested will be faxed back to you instantly without a break in your connection!

USING A TOUCH TONE PHONE...

1. Pick up your touch tone phone and dial ***(617)494-8338***.
2. Follow the voice prompts.
3. Hang up the phone and the information you requested will be faxed back to you immediately— in about 15 seconds— right to the fax number you specify.

WHAT COULD BE EASIER?

The next time you need product information right now, remember: Only COMPUTER DESIGN gives you *"INSTANT DATA ACCESS"* so you can get the details you need — INSTANTLY, or a moment later.

JUST CALL: (617)494-8338

INSTANT DATA ACCESS (IDA) DIAL (617) 494-8338 DOCUMENT NO. 1047

PARTICIPATING ADVERTISERS IN THE INSTANT DATA ACCESS PROGRAM IN APRIL:

ADVERTISER	"IDA" NO.
Atmel	1045
Brooktree	1026,1029
CMC	1025
Dynatem	1032
Hitachi	1034
Ironics	1046
KADAK	1036
Macrolink	1037
Mercury	1020
SBE	1040
Tech Source	1043
Teknor	1028

FROM THE MARCH ISSUE: (AVAILABLE THROUGH MAY 31)

ADVERTISER	"IDA" NO.
Brooktree	1026,1029
BSO/Tasking	1030
Canon USA	1021
CMC	1023
Dynatem	1032
Electronic Solutions	1033
Hitachi	1034
I-Bus	1035
KADAK	1036
Macrolink	1037
Mercury Computer	1020
Micron	1038
Newbridge	1039
SBE	1040
Spectrum Signal Processing	1041
Synergy	1042
Tech Source	1043
Teknor	1028
Ziatech	1044

Delivering Today's Electronics OEM

COMPUTER DESIGN• ONE TECHNOLOGY PARK DRIVE , WESTFORD, MA , 01886 • 508-692-0700 • FAX 508-692-7780

COMPUTER DESIGN Technology and Design Directions

A PennWell Publication

CIRCLE NO. 78

SOFTWARE & DEVELOPMENT TOOLS

C/C++ compilers bring faster code crunching to PC platforms

Jeffrey Child, Associate Editor

Developers at Honeywell used Borland C++ 3.0 to write software for the latest version of Honeywell's flight data test recorder. Because fast throughput was a priority in their application, they wanted a compiler that could produce fast, compact code. Borland's C++ produced code that was four times more compact than that of competing compilers.

The beauty of the C programming language is that it combines the ease of use of a high-level language with the ability to manipulate bits. For developers that write their code to adhere to the ANSI C standard, C also offers the advantage of portability. While porting code between platforms can be problematic, PC-based C compilers let you begin writing code in a comfortable, standard environment before the target hardware is even available.

As applications get larger and more complex, maintaining and organizing code has become a major issue. While maintaining programs written in C is much easier than maintaining miles of assembly code, even the original inventors of the C language admit that it was never intended for large programs. Enter C++: a superset of the C language that provides the basis for adopting an object-oriented approach to software design and enhances the original C language.

Until fairly recently, C++ was viewed by many as simply an "academic" language because of its lack of good tools. Now that tools for C++ are improving, many software development teams are ready to make the shift from C to C++. Several compiler vendors are currently offering full C++ compilers that generate object code directly from a C++ program as opposed to just translating it into C code and compiling it with a C compiler.

Thanks to the availability of powerful, low-cost PC software, more code is developed on a PC than on any other platform. And thanks to the increasing acceptance of Microsoft Windows, the PC will remain the most popular platform for software development for some time.

Today's selection of C and C++ compilers for the PC ranges from highly optimized compiler-only solutions to complete integrated development packages that contain debuggers, profilers and extensive class libraries for C++. Most of these compilers cost under $1,000, making it feasible to buy several and test them out at the same time.

Microsoft vs Borland

The most interesting C compiler news is the battle between the industry leaders, Microsoft (Redmond, WA) and Borland (Scotts Valley, CA). After a false start, Microsoft has finally come out with a C++ compiler, C/C++ 7.0. Despite the 7.0 designation, this is Microsoft's first C++ compiler release. (The 7.0 version number refers to the fact that the release is an upgrade from Microsoft C 6.0.)

Borland released version 3.0 of its C++ compiler last fall. As the industry waits for the 3.0 standard for C++ to be finalized, Microsoft and others have been content to implement the AT&T 2.1 specification of C++ in their compilers. In contrast, Borland chose to include support for templates, an element of the 3.0 standard already approved by the ANSI C++ committee. At least one other company, Metaware (Santa Cruz, CA), has an AT&T 3.0-compatible C++ compiler, to be released when the standard is finalized.

Besides a compile speed that's twice that of the previous version, Borland C++ 3.0 offers several other improvements. It supports the Direct Memory Programming Interface (DMPI), the standard for extended memory access which unlocks memory beyond the 640-kbyte barrier imposed by DOS. The new release also features a global optimization capability that lets you configure the compiler to optimize for fast execution speed or small code size. Capabilities specifically for C++ programming include class library support for both DOS and Windows.

Class hierarchies

To address the particular needs of object-oriented programming, Borland tailored its class libraries in hierarchical format. "Typically, when you're writing an application, you want it to be as general as possible," says Charles Dickerson, C++ product manager at Borland. "As soon as you start writing directly to the computer, you start limiting the portability of the application. Class hierarchies provide a layer of abstraction that ultimately lets you use the application on multiple platforms."

Object Windows is one example of a class hierarchy provided by the Borland C++ compiler. It provides objects important for Windows developers—such as dialogue and scroll-bar objects. These are pro-

Name	Version	Language	Compliancy	Modes (real, protected)	System requires (minimum)	System requires (recommended)	Price	Comments
Borland International 1800 Green Hills Rd, Scotts Valley, CA 95067-0001 (408) 438-8400								**Circle 301**
Borland C++	3.0	C/C++	ANSI, AT&T CFront 3.0	Protected	DOS 3.31+, 286, 24-Mbyte hard disk, 2-Mbyte RAM	Same as minimum with 3-Mbyte RAM (for Windows)	$495	Includes globally optimizing compiler, debugger, profiler, assembler, resource editor for Windows
Borland C++ with Applications Frameworks	3.0	C/C++	ANSI, AT&T CFront 3.0	Protected	DOS 3.31+, 286, 32-Mbyte hard disk, 2-Mbyte RAM	Same as minimum with 3-Mbyte RAM (for Windows)	$749	Same as above
Turbo C++ for Windows	3.0	C/C++	ANSI, AT&T CFront 3.0	Protected	DOS 3.31+, 286, 8-Mbyte hard disk, 2-Mbyte RAM	Same as minimum with 3-Mbyte RAM (for Windows)	$149.95	Same as above
Intel 5200 NE Elam Young Pkwy, Hillsboro, OR 07124-5961 (503) 696-4701								**Circle 302**
386/486 CodeBuilder Kit	1.1	C	ANSI	Protected	386, 7-Mbyte free disk space, 1-Mbyte extended memory	Same as minimum with 486 system	$695	Development kit for 32-bit DOS applications
JPI TopSpeed Consortium 117 Hunt Hill Rd, Rindge, NH 03461 (800) 448-4440								**Circle 303**
TopSpeed C Standard Edition	3.02	C	ANSI	Real mode DOS, protected mode DOS & OS/2	DOS 3.0+, hard disk, 512-kbytes RAM	—	$198	Fast ANSI validated, multitasking under DOS, segment-based overlaying of code & data
TopSpeed C++ Standard Edition	3.02	C++	AT&T 2.1	Same as above	Same as above	—	$198	Includes class library
Metaware 2161 Delaware Ave, Santa Cruz, CA 95060 (408) 429-6382								**Circle 304**
High C	2.4	C	ANSI	Protected	386/486, 1.5-Mbyte RAM	386/486, 3-Mbyte RAM	$795	32-bit, compatible with AT&T Unix porting tool
High C++	3.1	C++	ANSI, AT&T 3.0	Protected	386/486, 1.5-Mbyte RAM	386/486, 4-Mbyte RAM	$795	32-bit, available Q2 '92
Microsoft 1 Microsoft Way, Redmond, WA 98052-6399 (206) 882-8080								**Circle 305**
Microsoft C/C++	7.0	C/C++	ANSI, AT&T 2.1	Both	DOS 3.3+, 386, 15-Mbyte hard drive, 2-Mbyte RAM	DOS 5.0, 386, 20-Mbyte hard drive, 4-Mbyte RAM	$499 ($139 for upgrade)	Includes support for OLE and GDI, Windows 3.0/3.1 debug kernel and tools; uses p-code technology to reduce code size
Quick-C for Windows	1.0	C	ANSI	Both	286, 10-Mbyte hard drive, 1-Mbyte RAM, Microsoft Windows	386, 20-Mbyte hard drive, 2-Mbyte RAM, Microsoft Windows	$199	Integrated environment with Windows-hosted debugging; includes Windows development tools; Quick-Win libraries enable DOS applications to be recompiled as Windows applications
Microway Research Park, PO Box 79, Kingston MA 02364 (508) 746-7341								**Circle 306**
NDP C	3.2	C	ANSI, AT&T 2.1	Protected	386/387 or 486, 20-Mbyte hard disk, 2-Mbyte RAM	Same as minimum with 4-Mbyte RAM	$995 ($1,295 for 486)	Includes NDP tools and royalty-free DOS extender and a library that emulates Microsoft C DOS interface and graphics routines
NDP C++	4.1A	C++	ANSI, AT&T 2.1	Protected	Same as above	Same as above	$995 ($1,295 for 486)	Includes royalty-free DOS extender, features a full-function in-liner, classes, inheritance, function, and operating overloading
NDP C/C++ i860	4.1A	C/C++	ANSI, AT&T	Protected	Same as above	Same as above	$1,995	Free with Numbersmasher i860 board

Name	Version	Language	Compliancy	Modes (real, protected)	System requires (minimum)	System requires (recommended)	Price	Comments
Mix Software 1132 Commerce Dr, Richardson, TX 75081 (800) 333-0330								**Circle 307**
Power C	2.1	C	ANSI	Real	320-kbytes RAM	640-kbytes RAM	$19.95	Integrated make utility, 450 library functions, IEEE floating-point, supports x87, small/med/large models, near/far/huge keywords
Symantec 10201 Torre Ave, Cupertino, CA (408) 253-9600								**Circle 308**
Zortech C++ for Windows, DOS, & OS/2	3.0	C++	ANSI	Both	MS-DOS 2.0+ or OS/2 V1.x, hard disk, 640-kbytes RAM, PC-compatible	—	$699	Supports Windows, DOS, 32-bit and 16-bit protected mode DOS, and OS/2 V.1.x
Watcom 415 Phillip St, Waterloo, Ontario N2L-3X2 (519) 886-3700								**Circle 309**
Watcom C/386	9.0	C	ANSI, IBM SAA	Protected	PC-compatible, hard disk	386 or 486-based PC-compatible, 2-Mbyte RAM	$895	Supports development of 32-bit applications for OS/2 2.0, extended-DOS and Windows, royalty-free 32-bit DOS extender
Watcom C	9.0	C	ANSI, IBM SAA	Real	PC-compatible, hard disk	640-kbytes RAM, any x86 system	$495	Supports DOS, Windows, and OS/2, 286-DOS extender support; includes debugger, linker, profiler, and graphics library

Fuzzy Logic

You don't need to be a programmer to use CubiCalc. You can define and use fuzzy rules interactively to obtain immediate graphical and numeric results. And CubiCalc RTC provides tools for developers to bring fuzzy decision-making to their own programs – even real-time embedded systems.

CubiCalc is inexpensive and easy to use, so your investment is small but the potential payoff is large. (Fuzzy statements like that make a lot of sense, don't they?)

CubiCalc

The Third Wave in Intelligent Software

HyperLogic Corporation
1855 E. Valley Parkway, Suite 210
Escondido, CA 92027 USA
(619) 746-2765

CIRCLE NO. 66

REAL-TIME

SPARC

Integrated VME Workstation

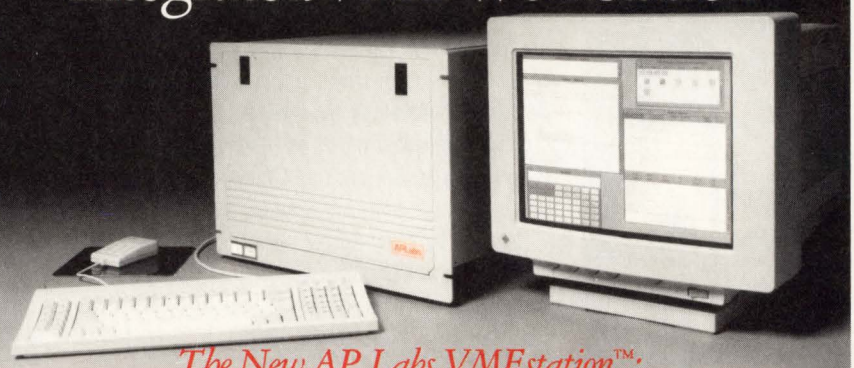

The New AP Labs VMEstation™:
Fully integrated systems for seamless real-time development...

- SPARCengine or 68000 series architecture
- SunOS and VxWorks with AP Labs' Asynchronous I/O
- RDBXworks™: C source-level SPARC debugger
- *Fully integrated* VME and SBus I/O real-time device driver support
- Data Acquisition; Communications; Signal Processing

AP Labs
6215 Ferris Square
San Diego, CA 92121
(619) 546-8626
Fax:(619) 546-0278

SPARC, SPARCengine, SunOS and SBus are trademarks of Sun Microsystems, Inc. VxWorks is a trademark of Wind River Systems. 68000 is a trademark of Motorola, Inc.

CIRCLE NO. 67

SOFTWARE & DEVELOPMENT TOOLS

vided at a very high level of abstraction. You can write to the Windows Application Programming Interface (API) when a lower level of control is required.

Borland eventually plans to release an OS/2 version of the compiler. OS/2 uses Presentation Manager (PM) as its graphical user interface. "Because we've provided this abstract layer, the developer can target an application for PM and the calls made into the class hierarchy don't change, because they're abstract enough not to be dependent on underlying API," says Dickerson.

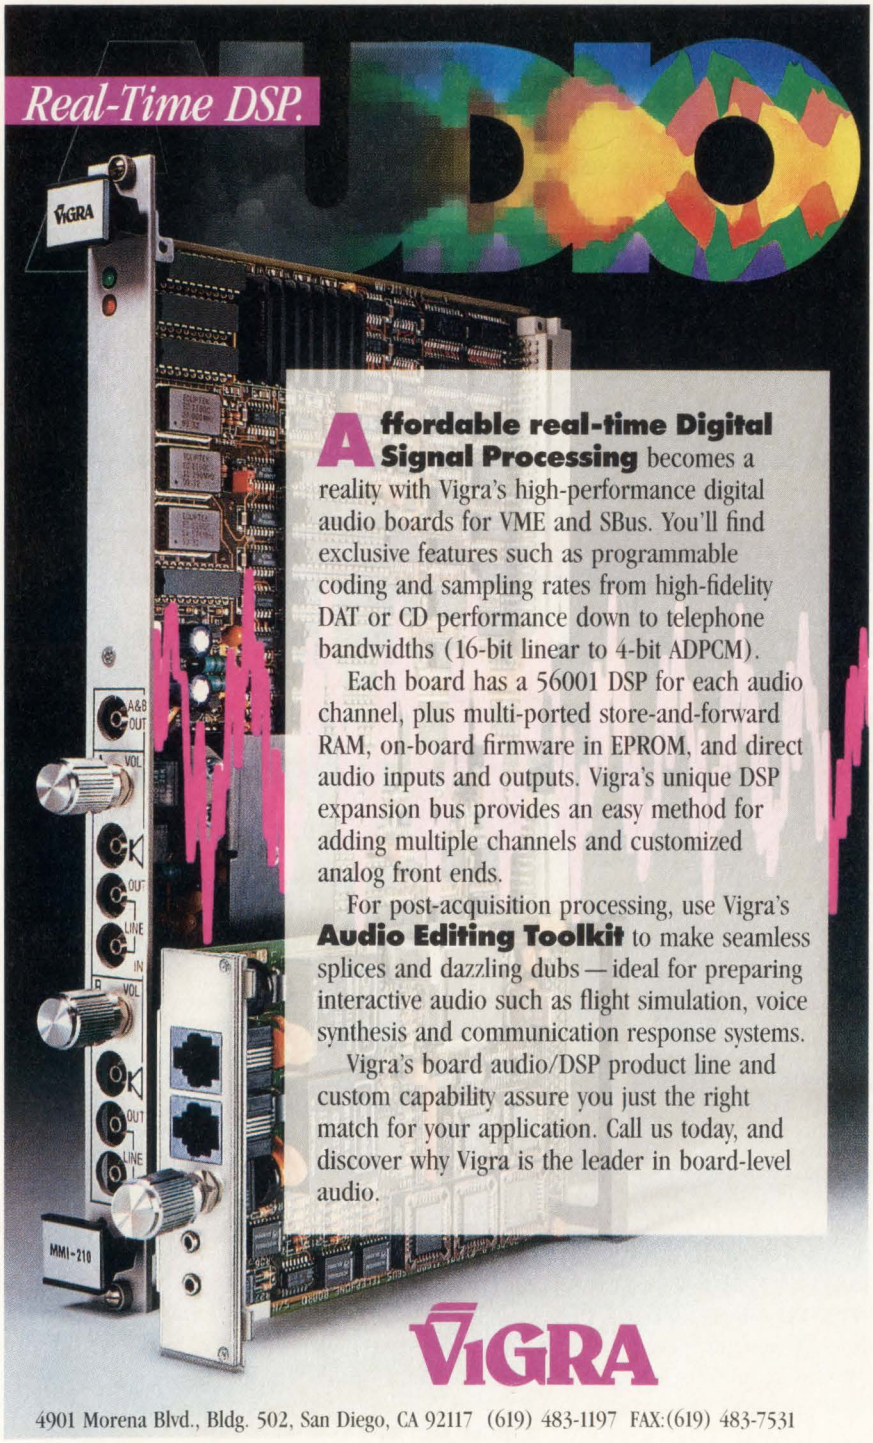

Affordable real-time Digital Signal Processing becomes a reality with Vigra's high-performance digital audio boards for VME and SBus. You'll find exclusive features such as programmable coding and sampling rates from high-fidelity DAT or CD performance down to telephone bandwidths (16-bit linear to 4-bit ADPCM).

Each board has a 56001 DSP for each audio channel, plus multi-ported store-and-forward RAM, on-board firmware in EPROM, and direct audio inputs and outputs. Vigra's unique DSP expansion bus provides an easy method for adding multiple channels and customized analog front ends.

For post-acquisition processing, use Vigra's **Audio Editing Toolkit** to make seamless splices and dazzling dubs — ideal for preparing interactive audio such as flight simulation, voice synthesis and communication response systems.

Vigra's board audio/DSP product line and custom capability assure you just the right match for your application. Call us today, and discover why Vigra is the leader in board-level audio.

VIGRA

4901 Morena Blvd., Bldg. 502, San Diego, CA 92117 (619) 483-1197 FAX:(619) 483-7531

CIRCLE NO. 68

Lee Pedlow, senior project engineer at Honeywell (Phoenix, AZ), used Borland C++ to develop test software for the latest version of Honeywell's flight data test recorder. The software runs on a PC which is used as a dedicated instrument. A mainframe computer models all the parameters of an aircraft in flight. This information is downloaded to the PC over a high-speed data channel. The test software on the PC sorts the information and divides it up to send to the various simulated aircraft buses. The information is then sent out to Honeywell's various systems—flight recorders, autopilots, indicators. Data also comes back through the PC from those systems and is then fed back into the simulation computer to update the aircraft model. All this I/O activity can be inspected visually and modified from the PC, without interrupting the modeling computer.

"We had a very limited amount of time in which to process all this information," says Pedlow. "Therefore, the degree of compactness of the code directly affects how much we're able to do in a fixed amount of time." The Honeywell team had been developing its system using another company's compiler before it made the switch to Borland C++ 3.0. "We were really impressed with the speed and efficiency of the Borland compiler," says Pedlow.

Tales of assembly

Because speed was a key design criterion, Pedlow tested various C/C++ compilers on the market by examining their assembly code output. "Most compilers provide the option to output an assembly listing," he says. "We used that to rate the efficiency of the compilers. We put one of our standard structures through each compiler, and then looked at the resulting assembly code to see just how tight it was. We found that the Microsoft compiler generated over 400 lines of assembly code. The code generated from the Borland compiler was only around 100 lines. That was significant—we reduced the size of the program by a factor of four just by switching compilers."

Typical of a growing number of system designers, Pedlow and his team are interested in migrating their code to C++ and using object-oriented programming techniques.

Currently, they only use a few features of C++. While they've used the C++ inheritance capabilities in some of the display routines and user I/O, the code wasn't utilized in the real core of the program. "Our particular application doesn't lend itself to the concept of inheritance," says Pedlow. "Our program has about 200,000 lines of code, 75 percent in C or C++, and 25 percent assembly code. C++ didn't help us in the assembly department at all. But, the inheritance concepts did help us move the assembly code into the C domain as callable routines. Basically, we used our assembly routines as the base of a library. We could then assign inheritance traits to the library."

Never too late?

A late entry in the C++ game, Microsoft offers its first C++ compiler in combination with release 7.0 of its C compiler. The Microsoft C/C++ 7.0 compiler features a sophisticated browser that's part of its Integrated Development Environment (IDE). Using one of 16 different types of queries, you can use the browser to help find classes, functions, macros, and members. The browser maintains a history of all the items found. You can look for a class, for example, find some of its members and then locate where that member was first defined—or you can go to some other place where that member is redefined or reused, and make modifications to it. This helps automate the task of organizing large, complex object-oriented programs.

Several new code optimization technologies are featured in Microsoft C/C++ 7.0, two of which are significant. The first is packed-code (p-code), a code-generation technology used by Microsoft in developing its other products. P-code is an intermediate object-code representation which compresses code size by 30 to 40 percent. You can configure the compiler to generate p-code selectively in all applications where code volume reduction can be made at the least cost in speed. User-interface code, for example, where the software is often waiting for user input, is a place where applying p-code makes sense. Although p-code runs slower than native code, its smaller size can mean less swapping of memory to disk, resulting in an overall gain in speed.

The second Microsoft optimization feature is in-lining. In-lining places an entire function's code "in line" in the program instead of jumping to another part of a program when a function is called. Microsoft's approach to in-lining is unique because it can handle functions containing control structures. This lets larger and more complex combinations of code be in-lined. You can set the compiler to automatically in-line all the code, or choose individually which functions to in-line. ■

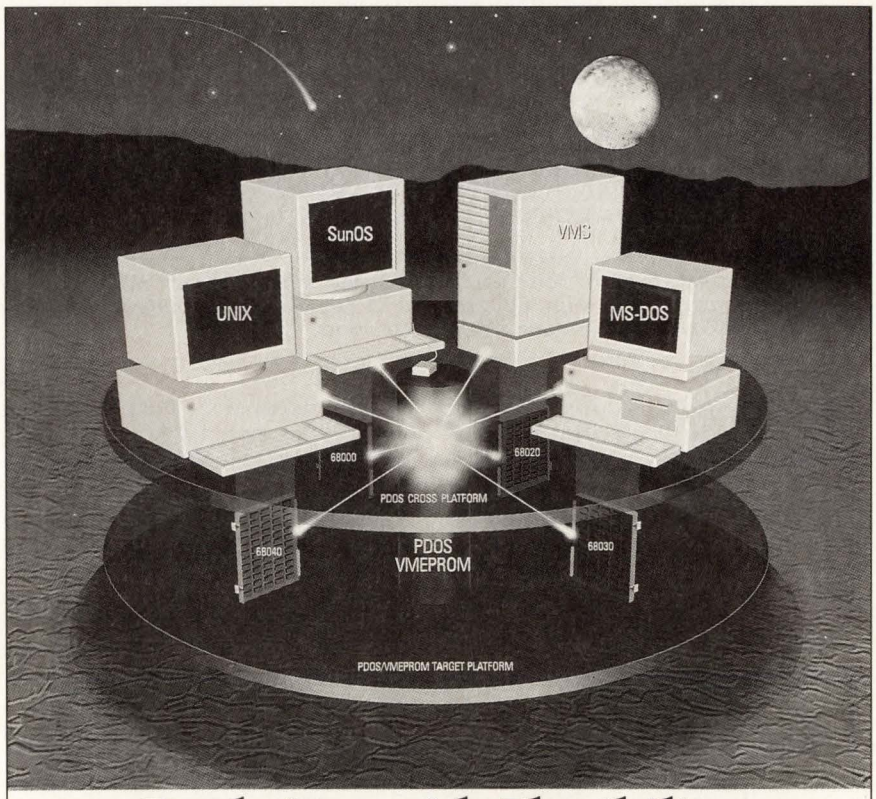

Real-time with Flexibility, Portability, and Maintainability.

PDOS Full Spectrum Development gives you open platform flexibility. You can write real-time/embedded software on your favorite cross development system—and/or the target system.

Generated code is source and object compatible across all platforms. The code is completely portable between cross environments, across host and target platforms, and between the same or different target boards. You simply move the code to the new platform. It's that easy.

With proven PDOS performance and Full Spectrum flexibility, you get your product to market on time and within budget. Plus, with VMEPROM options, you have a future for upgrades and easy maintenance.

Set your product worlds apart from the competition, write or call for informative literature to help you get started. Act now and we'll also send you a free benchmark report on the performance of several real-time operating systems.

CALL 1-800-YES-PDOS PDOS

PDOS Puts Power in Your Products – For The Present and The Future.

EYRING CORPORATION, SYSTEMS SOFTWARE DIVISION, 1450 WEST 820 NORTH, PROVO, UTAH 84601 801-375-2434 FAX 801-377-3850 EMAIL: pdos-info@Eyring.COM

GERMANY, AUSTRIA: Allmos Electronic GmbH, Martinsried, Germany, Tel: 089/857000-0 **UNITED KINGDOM:** Eyrisoft Ltd., Derby, England, Tel: 0332/384978

BENELUX: Interay BV, Friesland, Holland, Tel: 05116/4052 **JAPAN:** Hi-System Control Corp., Tokyo, Japan, Tel: (03) 3201-0511

CIRCLE NO. 69

COMPUTERS & SUBSYSTEMS

Two Sparc-based CPUs unveiled

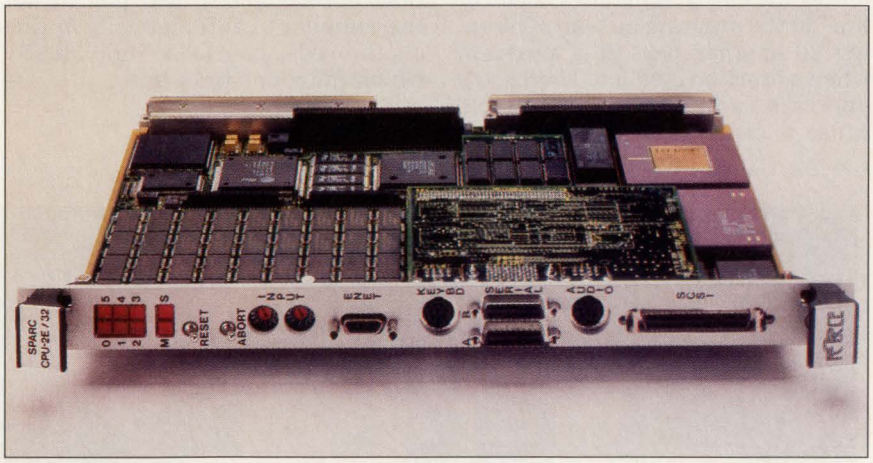

The newest members of the Force Computers family of Sparc-based CPU boards are a pair of full-featured, 28.5-Mips, 6U cards. The 2E (shown here) includes a full VME64 interface with SSBLT. Both include a pair of SBus slots.

Although Revision D of the IEEE-1014 specification, covering VMEbus, still has obstacles to overcome before becoming part of the official standard, companies have moved ahead to implement designs in silicon. One of these, Force Computers (Campbell, CA), is offering the first silicon implementation of SSBLT (source-synchronous block transfer) as part of a new Sparc-based single-board computer. Force's 28.5-Mips SBC comes complete with all the 'fixings' of a Sun Microsystems (Mountain View, CA) Sparcstation 2 workstation, including a pair of SBus slots.

Two products were announced by Force; both were 6U single-board Sparc-based CPUs, the only difference being that one includes the full VMEbus interface while the other is designed for stand-alone applications.

"By providing OEMs both alternatives," says Force vice-president of North American operations Fred Rehhausser, "it's possible to address both large VMEbus solutions as well as embedded applications with high CPU performance—but with more modest I/O requirements. The pair of SBus slots provides I/O solutions for those modest embedded applications. And the VMEbus board provides a low-cost I/O where a full-sized VMEbus board isn't required."

The basic boards have up to 32 Mbytes of memory, with a connector for an expansion card which can boost memory to 64 Mbytes. For local I/O, the board features a front-panel assembly with connectors for a pair of serial interface ports, a keyboard/mouse serial port and an audio port. The local I/O is implemented with a mezzanine board that provides both the logic and the connector. The board also includes Ethernet and SCSI II ports, which are located on the base board.

A floppy-disk interface, a loudspeaker port and a second SCSI II interface are available on the P2 connector. The two SBus expansion ports also provide the capability to plug any available expansion cards directly into the board.

Lots of EPROM

In addition to the main DRAM, the board includes a single 32-bit JEDEC EPROM socket that's used to hold a 1-Mbit boot EPROM. Though the EPROM containing the open boot (SBus) firmware isn't identical to that of the Sparcstation 2, Force claims it's 100 percent downward compatible with Sun's approach. The board also includes four flash EPROMs in TSOP (thin small-outline package) format. This additional one Mbyte of memory is mapped in place of the Sun Sparcstation 2's third SBus slot. The EPROMs are organized in 32-bit words and are accessible with non-burst cycles. They are erasable and writable using routines supplied in the boot ROM. Using 150-ns EPROMs, the access time is 4 clock cycles.

The VME-compatible board (CPU-2E) is identical to the non-VMEbus version (CPU-2S), except that it includes a complete VMEbus interface. This interface, including both VME64 and SSBLT, is part of a two-chip set to boost VME performance that's proprietary to Force.

SSBLT is a technique that permits the bus clock to be controlled by the data rate rather than an independent clock. Using SSBLT it's possible to boost the conventional VME transfer from a theoretical maximum of 40 Mbytes/s to just about 160 Mbytes/s.

The chip set not only provides the full VME64/SSBLT interface, though; it also lets the 2E board implement the latest SBus version. About a year ago, Sun introduced Revision B.0 of the SBus specification, which permitted multiplexing the data bus to allow 64-bit transfers—boosting SBus transfer rates to 160 Mbytes/s.

Both boards are designed to operate with Sun's SunOS, as well as other real-time operating systems designed to work with the Sparc architecture. The 2S with 16 Mbytes is priced at $7,495; the 2E with 16 Mbytes is $7,995. — *Warren Andrews*

CPU 2S/2E at a glance

- 28.5 Mips Sparc processor at 40 MHz with 64-kbyte cache
- 4.2 MFlops double-precision floating point unit at 40 MHz
- 16-, 32-, or 64-Mbyte DRAM
- 1-Mbyte Flash EPROM
- SunOS compatible
- Two SBus expansion slots, SCSI II port, two serial ports, keyboard and mouse port
- Floppy disk controller
- Ethernet
- VME64, SSBLT Interface (2E)

Force Computers
3165 Winchester Blvd
Campbell, CA 95008
(408) 370-6300
Circle 354

How Do You Make A Variety Of Computers Share Information In Real Time?

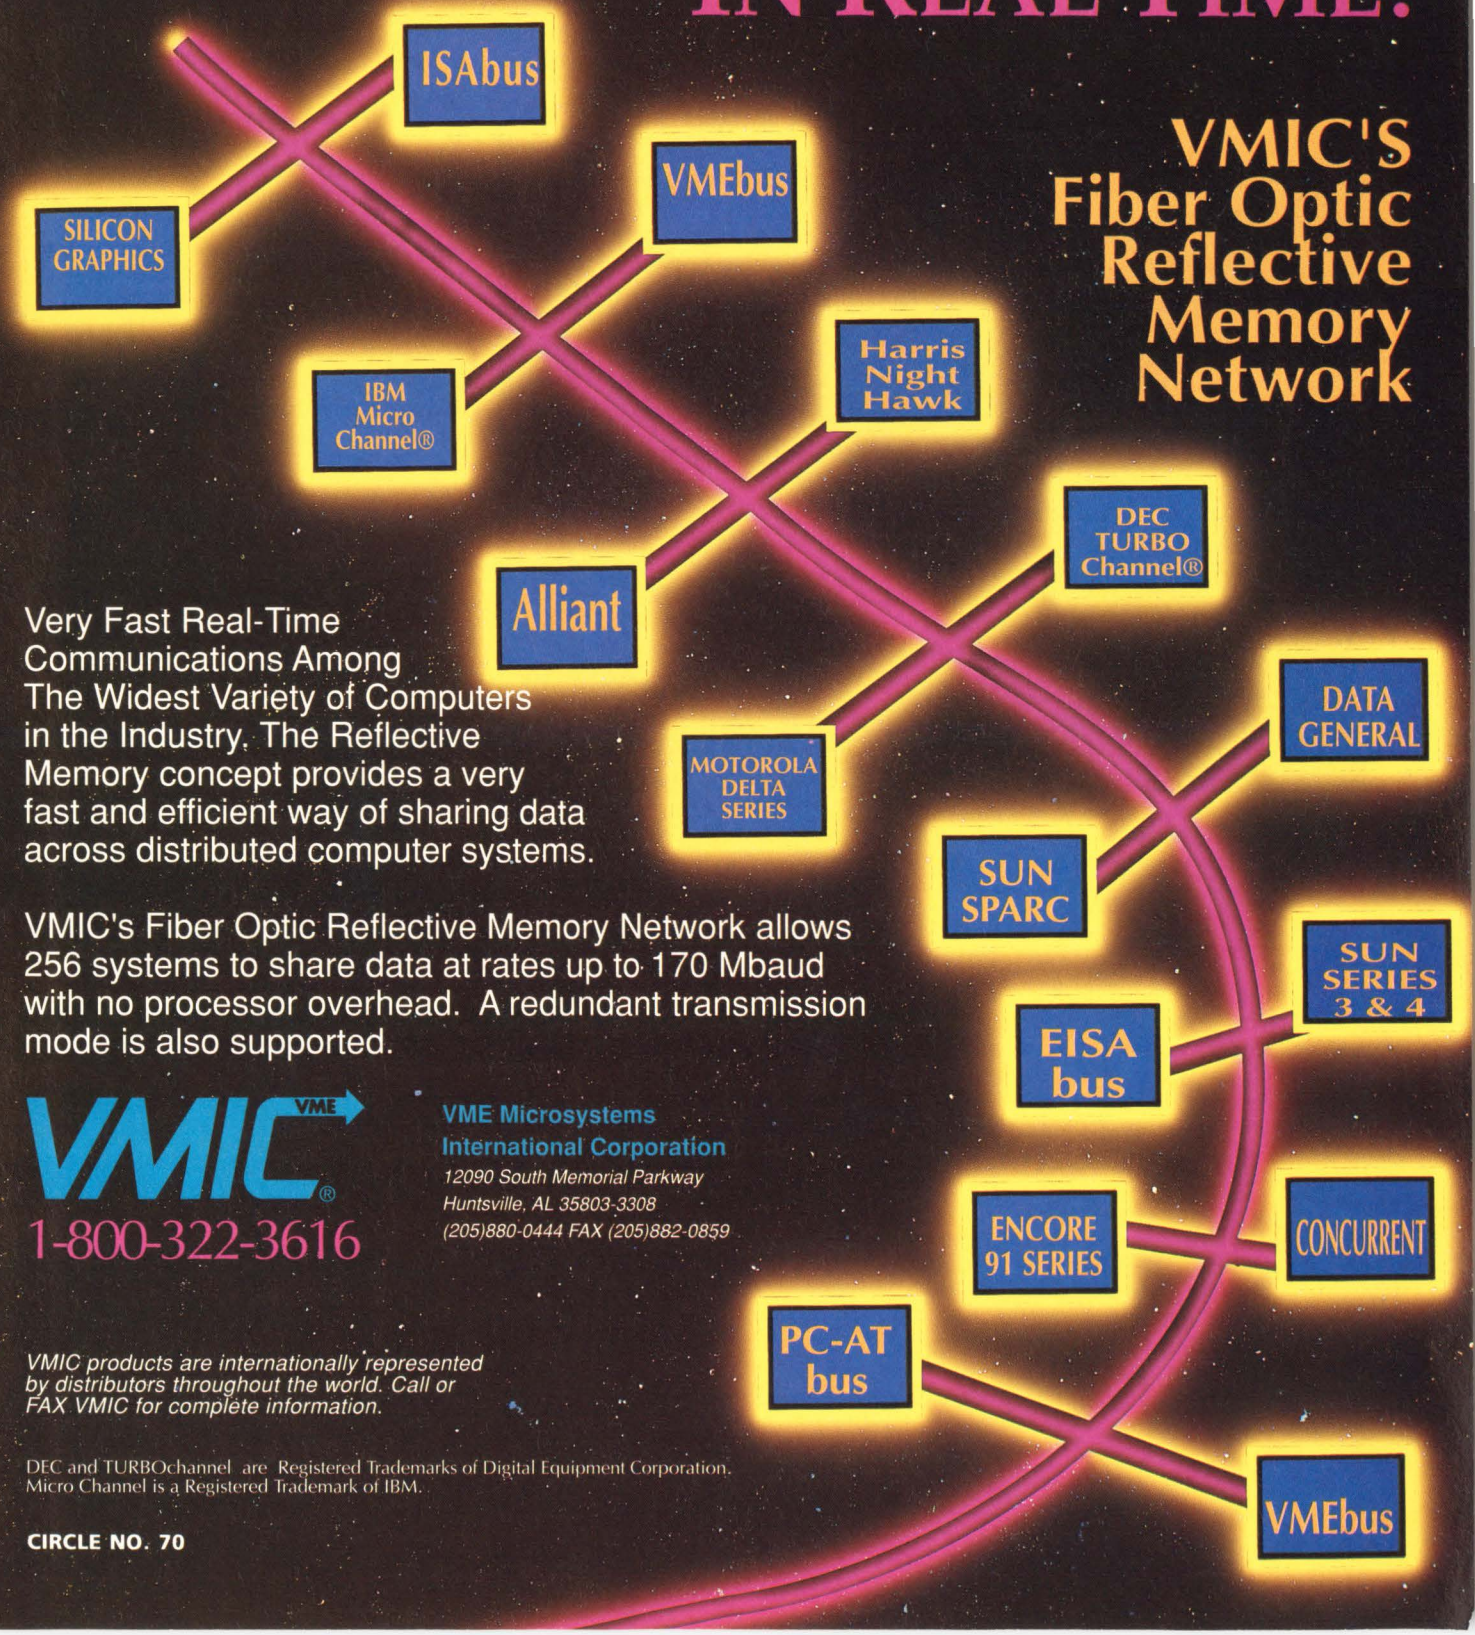

VMIC'S Fiber Optic Reflective Memory Network

Very Fast Real-Time Communications Among The Widest Variety of Computers in the Industry. The Reflective Memory concept provides a very fast and efficient way of sharing data across distributed computer systems.

VMIC's Fiber Optic Reflective Memory Network allows 256 systems to share data at rates up to 170 Mbaud with no processor overhead. A redundant transmission mode is also supported.

VMIC VME®

1-800-322-3616

VME Microsystems International Corporation
12090 South Memorial Parkway
Huntsville, AL 35803-3308
(205)880-0444 FAX (205)882-0859

VMIC products are internationally represented by distributors throughout the world. Call or FAX VMIC for complete information.

DEC and TURBOchannel are Registered Trademarks of Digital Equipment Corporation.
Micro Channel is a Registered Trademark of IBM.

CIRCLE NO. 70

SOFTWARE & DEVELOPMENT TOOLS

Graphic tool shows multiple views of real-time behavior

A CASE tool aimed at describing the actual behavior of real-time systems using graphic representation has been introduced by Verilog (Dallas, TX). Scenario is based on a formal language defined by the Consulting Committee on International Telegraph and Telephone (CCITT), an organization of the United Nations. The formal description lets you link the external behavior of a system with the internal graphical view of its architecture. Scenario features graphical capture of system event behavior, a semantic model manager and static analysis.

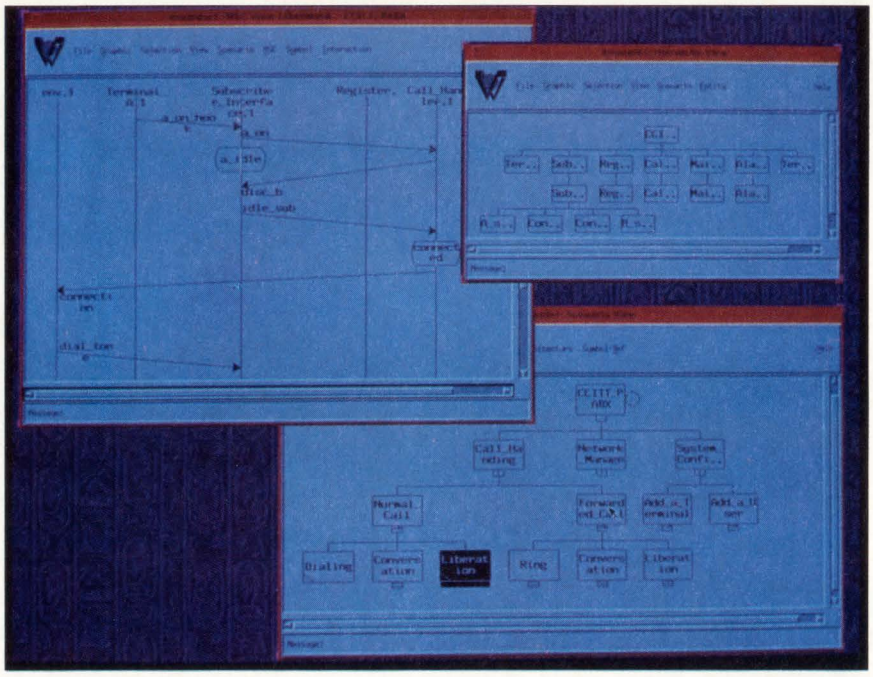

Several views of Verilog's Scenario show the message sequence chart (upper left), the service decomposition hierarchy (lower right) and the file structure (upper right) of a real-time system under analysis.

Three methods

In the graphical capture of system event behavior, Scenario describes the real-time event interaction between design components in the system in three ways. Message sequence charts show what happens within the system in response to a message, such as a message passed between concurrent tasks. A display of the system architecture shows the hierarchy of the actual components of the system, and a service decomposition display shows the hierarchy of services provided by the system. The display of the service decomposition, for example, shows a hierarchical chart with elements linked via "composition operators." These tell whether an element and its subelements are recursive, operating in parallel with other elements, are exceptions, or fall into some other category.

Consistency between views is maintained by Scenario's semantic model manager, which automatically updates the information presented in all views. You can confidently use a multiwindowed editing environment, therefore, and instantly see the effects that a change made in one view will have on other aspects of the system. In static analysis, Scenario provides a checking facility to identify errors associated with completeness and correctness of data. It does this by comparing the model content with the formal CCITT language description to verify type consistency, completeness, balancing of data, and correctness of entity decomposition.

Easy integration

Scenario is the first product release in Verilog's new generation of object-oriented technology for software development. The object-oriented architecture lets a tool such as Scenario be easily integrated into Verilog's Logiscope line of tools, which in turn can be integrated with other tools for such things as automated document generation, user-specified configurations and control integration frameworks such as Hewlett-Packard's SoftBench.

Scenario runs on DEC, Hewlett-Packard/Apollo and IBM Unix workstations, and it supports the Motif/X Windows interface, as well as DEC Windows. It also runs on IBM 43XX and 30XX systems, and on PCs under DOS/MS-Windows and OS/2.

Scenario is priced at $15,000 for single quantities. Multiple licensing arrangements in quantities of 15 or more can bring the price down to around $4,500. — *Tom Williams*

Scenario at a glance

- Object-oriented real-time analysis tool
- Multiwindowed views of system architecture, service hierarchy and real-time event behavior
- Automatic update of multiple views
- Automatic data consistency and correctness checking
- Easy integration into overall design environment

Verilog USA
3010 LBJ Fwy
Suite 900
Dallas, TX 75234
(214) 241-6595
Circle 351

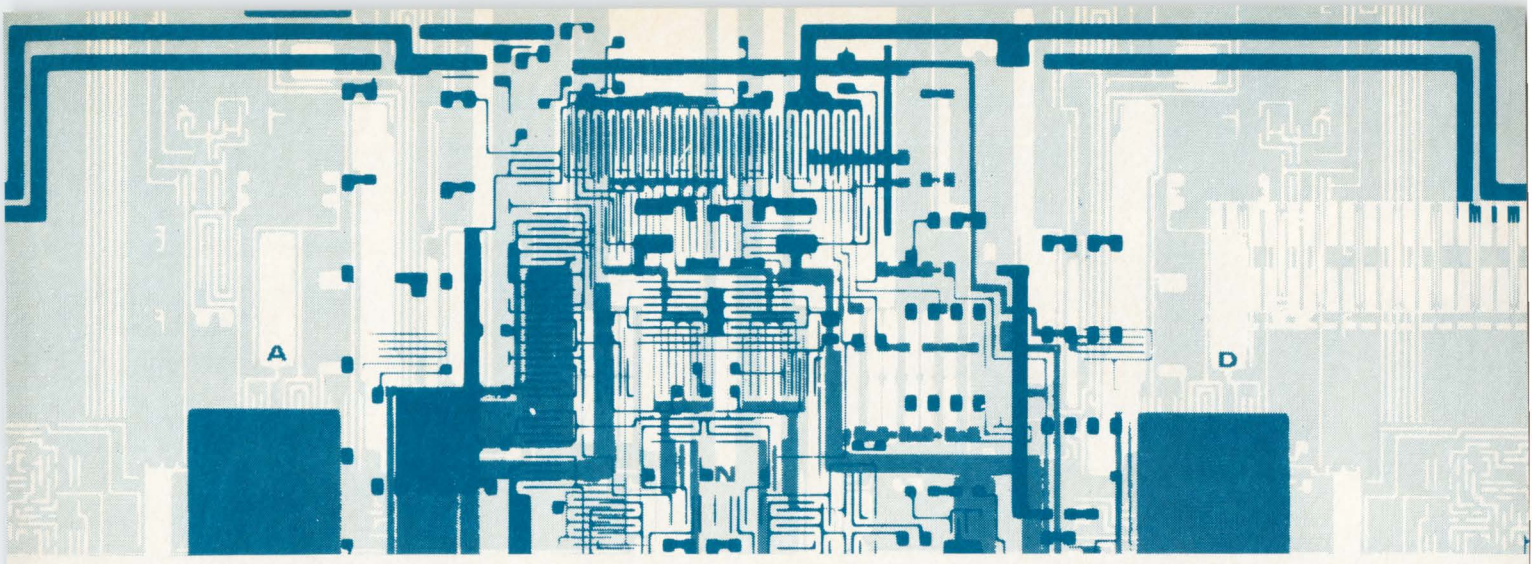

Fourteenth Annual

Custom Integrated Circuits Conference

Westin Copley Place & Boston Park Plaza Hotels
Boston, MA
May 3-6, 1992

Keynote Speaker: Dr. James E. Carnes, President & COO, David Sarnoff Research Center
High Resolution Video & ICs: Help or Hype?

Luncheon Speaker: Mr. Eric M. Howlett, President, LEEP Systems Inc.
Virtual Reality

Technical Sessions:

Over 170 papers presenting the latest developments in the areas of digital and analog circuit design, system analysis and simulation, CAD systems and methodology, fabrication technology, ASIC design automation, device and circuit modeling, data conversion and telecommunication circuits, testing, and circuit/system reliability.

Educational Sessions:

Two-hour primer courses, with extensive "course notes" in an interactive setting.

Managing the ASIC Design Cycle
Design for Reliability
High Speed Digital Circuit Design
Programmable Logic Design Choices
Packaging & Multi-chip Modules
Opto-electronic Circuits
Automating Analog Design & Layout
High Speed Analog Circuit Design
IC and System Simulation
Logic & High Level Synthesis
Testing and Design for Testability
Built-in Self-Test & JTAG

Evening Panel Discussions:

Intellectual Property: How Smart Is It
Patent laws: Are they an aid or barrier to innovation? How do little guys fare against big guys in patent wars? Europe and Japan's laws differ markedly from the US system. Their impact on global markets? Our line-up of CEOs and patent attorneys trade views.

Multi-Chip Modules: Dream or Reality?
MCMs are becoming pervasive. Or are they? Along with their obvious merits come the inescapable cost and testing issues. A panel of leaders for the IC and computer industry weigh the future of MCMs.

Eng. Skills for the Year 2000: What are they?
Think of it: multi-million transistor chips designed in just a few months. Who is the chip designer of the next decade? A programmer? System architect? A super-generalist, or a super specialist? Compare your crystal ball with our panel of experts.

Exhibits:

Dozens of companies show the latest developments in software & hardware in ASIC design and development.

For more information, including registration & housing forms, contact
Roberta Kaspar, Technical Program Coordinator:

Phone: 716-865-7164 FAX: 716-865-2639

CIRCLE NO. 82

29th ACM/IEEE DESIGN AUTOMATION CONFERENCE® AND EXHIBITION

ANAHEIM CONVENTION CENTER
June 8 - 12, 1992
Anaheim, California

acm SIGDA® IEEE COMPUTER SOCIETY DATC® IEEE CIRCUITS & SYSTEMS SOCIETY®

NEW DAC USER TRACK

ATTEND THE WORLD'S PREMIER CAD/CAM CONFERENCE!

DA Users and Developers, world-wide attend DAC to keep current in the latest Design Automation Tools and Algorithms.

- **Over 130 exhibitors of CAD hardware and software products.**
- **Over 60 exhibitor technical presentations Monday, June 8, 1992, many announcing new products.**
- **Monday Exhibits Only Passes. Free invitations available from participating companies.**
- **Twelve full-day tutorials, Friday, June 12, 1992.**
- **Over 140 papers, tutorials and panel presentations.**

ADVANCE REGISTRATION ENDS MAY 8, 1992!

HOTEL RESERVATION FORM
29th Design Automation Conference
June 8 - 12, 1992

Hotel reservations will be handled by the Anaheim Convention and Visitors Bureau, on a first-come, first-served basis. The following hotels are available:

HOTEL (prioritize your choices)	SINGLE (please circle one)	DOUBLE
___ Convention Center Inn	$ 55.00	$ 55.00
	(Queen Parlor Suites - $78.00)	
___ Anaheim International Inn	62.00	62.00
___ Castle Inn	64.00	64.00
___ Raffles	64.00	64.00
	(Parlor Suites - $79.00)	
___ Holiday Inn	78.00	78.00
___ Quality Inn	79.00	84.00
___ Jolly Roger Inn	79.00	85.00
___ Grand Hotel	82.00	88.00
___ Inn at the Park	85.00	95.00
___ Hyatt Regency Alicante	89.00	95.00
___ Pan Pacific	92.00	98.00
___ Disneyland Hotel	100.00	100.00
___ Anaheim Marriott	110.00	120.00
___ Anaheim Hilton	115.00	130.00

Above rates honored only if your reservation is made by May 8, 1992.

Complete the form below and send to:

DAC HOUSING BUREAU
P.O. Box 4270
Anaheim, CA 92803
(714) 999-8939 - FOR INFORMATION ONLY!

CAUTION: Hotel reservations sent to the DAC office will be discarded!

Name_____
(first) (last)

Company_____ Mail Stop_____

Mailing Address_____

City _____ State_____ Zip_____

Country_____ Telephone (____)_____

Arrival Date_____ Time_____ AM_____ PM_____

Departure Date_____ Time_____ AM_____ PM_____

Reservations will be confirmed by the DAC Housing Bureau and the hotel to which you are assigned. A one night deposit of $80.00 per room is required by check or credit card. Make check payable to the DAC HOUSING BUREAU.

☐ VISA ☐ MASTERCARD ☐ AMERICAN EXPRESS

Number_____ Expiration Date_____

Signature_____

CIRCLE NO. 77

ADVANCE REGISTRATION FORM
29th Design Automation Conference
June 8 - 12, 1992

To register, mail your payment with this form or a copy to:

29th Design Automation Conference
c/o DAC REGISTRATION DESK
7490 Clubhouse Rd. #102
Boulder, CO 80301 USA

For Information Only, Call (303) 530-4333. **NO TELEPHONE REGISTRATIONS**

IMPORTANT - Registration Instructions

1. Full payment in U.S. Dollars (drawn from a U.S. bank) MUST accompany registration. Any registration without payment WILL be discarded.
2. Advance registrations must be postmarked no later than **May 8, 1992.**
3. Register one person per form (copy this form as needed).
4. **Refund Policy**: Written requests for refunds must be received in the DAC office no later than **May 8, 1992.** Refunds are subject to a $15.00 processing fee.
5. Make checks payable to: **29th Design Automation Conference.**

PLEASE TYPE OR PRINT CLEARLY

CAUTION: This information will be used to print your ID badge!

Name_____
(first) (last)

Company_____

Street Address_____

Mail Stop_____

City_____ State_____ Zip_____

Country_____ Telephone (____)_____

MEMBER NO. ACM/IEEE_____

CONFERENCE REGISTRATION FEES

	Conference Only	Conference and Tutorial
Advance Registration (postmarked by May 8, 1992)		
Member IEEE/ACM	$140.00	$265.00
Non-Member	$180.00	$355.00
At Conference Registration (after May 8, 1992)		
Member IEEE/ACM	$180.00	$305.00
Non-Member	$220.00	$395.00

Students and One-day-only registrations will be accepted AT CONFERENCE only!

☐ VISA ☐ MASTERCARD

Number_____ Expiration Date_____

Signature_____

CD

COMPUTERS & SUBSYSTEMS

Ethernet multiplexer takes over network processing

Increased throughput for Ethernet terminal servers comes from a new Ethernet multiplexer from Systech (San Diego, CA). The Etherplex NTX-8023 multiplexer offloads the burden of interrupts and context switching from the host computer by relieving it of the burden of time-consuming interface tasks such as network packet interrupt and TCP/IP protocol processing. In addition, software such as the Telnet "daemon," which normally runs as a user task on the host, can execute on the Etherplex board.

The speed and processing power to handle multiple packets and make them look like simple multiplexed lines comes from the Etherplex's i960CA RISC processor, running at 25 MHz. The i960CA is the superscalar version of the i960, capable of executing multiple instructions per cycle. Tests run by Systech indicate that at up to 16 users, it's equal to Telnet in output processing. As the user base increases to 80, however, Etherplex maintains an output rate of nearly 4,000 characters per second (cps), while the host-based system drops off to around 1,000 cps.

The higher sustained data rate has less impact on CPU utilization, freeing more host power for processing applications. In tests of output efficiency with 80 users, Systech reports 26 percent CPU loading using Etherplex versus 95 percent under host-based Telnet. Under input conditions with the same 80 users, the testers report 17 percent loading versus 98 percent under Telnet.

Insulated for performance

Etherplex is fully insulated from the operating system as well as from terminal servers, acting like a "black box" as it optimizes terminal traffic. This means no changes are needed in Telnet-compatible terminal servers or host applications. Any application with a standard TTY interface can use Etherplex without changes, and any off-the-shelf terminal server that complies with TCP/IP or Telnet protocols can operate with the device. Systech maintains a list of terminal servers which it has tested and verified. Use of the proper transceiver lets Etherplex support any medium—twisted pair or ThinNet, as well as the standard Ethernet coax.

Central to the Etherplex board's performance is Systech's terminal control software (TCS), originally developed for its Unplug terminal I/O subsystem. TCS provides integration of standard networking protocols so that the host computer sees the subsystem as a single, multiline, asynchronous multiplexer. At the same time, networked terminals see the Etherplex subsystem simply as a host offering standard Telnet terminal devices.

The initial offering of Etherplex will be in its present VMEbus configuration, but Systech plans to introduce versions for EISA, Micro Channel and SBus in the second half of 1992. The price is $3,995 for single-lot quantities, with OEM discounts available. — *Tom Williams*

Etherplex NTX-8023 at a glance

- i960CA CPU at 25 MHz
- 2 Mbytes two-way bank-interleaved DRAM
- 16 kbytes dual-ported DRAM
- 128 kbytes EPROM
- 16- and 32-bit data transfers
- Supports both master and slave bus modes
- Includes Network Terminal Software: TCP/IP with Telnet, Terminal control software (TCS) utilities, TCS-compatible drivers for AT&T Unix 5.3 and 5.4 Streams

Systech
6465 Nancy Ridge Dr
San Diego, CA 92121
(619) 453-8970
Circle 355

68030 VME

3U efficiency.

Wouldn't it be great if you could get a compact, high performance 3U 68030 VME board with the features you need at a price you can afford? Now GreenSpring delivers the goods with the new workhorse SBC5.

NEW VME PRODUCT

We started with a speedy 32Mhz 68030 processor. Added a burst memory cache fill controller and 60ns DRAMs for 64MB/second throughput. Integrated our proprietary GLA™ bus interface silicon in an economical 100-pin quad pitch package. Then topped it off with 32-bit FLASH memory for in-circuit programming. Now *that's* efficiency.

Of course, we've included all the functionality you expect – EEPROM, RTC, serial ports and timers. A 32Mhz 68882 FPU is optional. Get OS-9 or CrossCode™ C and FreeForm™ for source-level debugging. Or use the board support package for your favorite kernel.

Best of all, you can get this unprecedented efficiency for the unexpected low price of $1695 list ($1295 OEM).

Call or fax GreenSpring today for a complete catalog of VMEbus products.

GREEN SPRING COMPUTERS

1204 O'Brien Drive, Menlo Park, CA 94025
(415)327-1200/FAX(415)327-3808

CIRCLE NO. 73

ASICs & ASIC DESIGN TOOLS

High-density PLD runs 80 MHz, programmable in-system

A 6,000-gate programmable logic device (PLD) from Lattice Semiconductor, claimed to have a maximum system-level performance of 80 MHz with a 15-ns propagation delay, is available in two versions: a programmable LSI device (pLSI 1032) and an in-system programmable device (ispLSI 1032). Lattice supports its 0.8-µm E^2 CMOS PLDs with a $995 pDS development system which runs on 80386- and 80486-based PCs with Microsoft Windows.

New benefits for FPGA designers

"Lattice's pLSI and ispLSI devices deliver predictable timing, E2 reprogrammability, performance, and quality to the experienced FPGA designer," says Lattice president and CEO Cyrus Y. Tsui. Dave Harper, senior design engineer at Convex Computer, reports that he's using multiple ispLSI 1032s running at 80 MHz in a high-speed, high-definition TV design. "They were the only devices to offer predictable 15-ns speed and the unique in-system programming capability," Harper says. Lattice claims that its ispLSI devices are the industry's only high-density PLDs that offer nonvolatile in-system programmability and reprogrammability.

The Lattice PLD has 192 registers, 64 universal I/O pins, eight dedicated input pins, four dedicated clock input pins, and a global routing pool that provides interconnectivity to all its elements. The basic unit of logic is the generic logic block, which has 18 inputs, four outputs, and the logic needed for most standard functions. The internal logic of the generic logic block is divided into an AND array, a product term sharing array, reconfigurable registers, and control functions. The four outputs of each of 32 logic blocks can be configured to be either combinatorial or registered.

Inputs to the generic logic block come from the global routing pool, and all logic-block outputs are brought back into the routing pool so that they can be connected to the inputs of any other logic block. Each of 64 I/O cells can be individually programmed to be a combinatorial input, registered input, latched input, output, or bidirectional I/O pin with three-state capability.

Eight of the 1032's generic logic blocks, plus an output routing pool and 16 I/O cells, make up a megablock. Each megablock shares two sets of resources. Two dedicated (nonregistered) inputs in each of the four megablocks are available only to generic logic blocks within that megablock. Common to all 16 of the I/O cells in each megablock is an output enable signal that can be generated using a product term in any of the eight logic blocks. Because of this shared logic, you will obtain the best utilization of the 1032's logic and most effectively eliminate routing bottlenecks if you group signals sharing a common function, such as buses or counters, within a megablock.

The 1032's clock distribution network generates five global clock signals—three for clocking all the generic logic blocks and two for clocking all the device's I/O cells. Four dedicated system clock pins can be directed to any generic logic block or any I/O cell that uses the clock distribution network.

Do-it-yourself design

To design the 1032, Lattice offers the pDS development system, which has a graphical user interface and features Boolean entry with Abel-like syntax. You can define your own macros or use Lattice's library of over 240 macros, which cover most TTL functions, from gate primitives to 16-bit counters. After the pDS software verifies the design and performs logic minimization, Lattice's proprietary place-and-route algorithm assigns pins and critical speed paths, ensuring 100 percent routability at 80 percent utilization. Harper of Convex says, "The pDS software has proven to be a very cost-effective development system. Its low price, easy-to-learn interface and quick design entry enabled us to implement the design quickly, without significant cost or timing delays."

To complete the design environment for in-system programmable 1032s, Lattice offers the isp Engineering Kit. You can program ispLSI devices directly from a PC, either stand-alone in a programming module or directly on a PC board. The kit includes a sample of the ispLSI 1032, a programming module, cabling, and a power adapter.

Production versions of the pLSI 1032 are available now at both 80 MHz and 50 MHz. In quantities of 1,000, the 84-pin pLSI 1032-80 and 1032-50 are $81 and $49 respectively. The ispLSI 1032 is available today in samples, with production quantities expected sometime this quarter. The Lattice pDS development system is $995, and the isp Engineering Kit is $395.

—Barbara Tuck

E^2 CMOS PLDs at a glance

- In-system programmability
- 80-MHz maximum operating frequency
- 15-ns propagation delay
- 192 registers
- High-speed global interconnects
- 100 percent routable at 80 percent utilization
- Boolean logic compiler

Lattice Semiconductor
5555 NE Moore Ct.
Hillsboro, OR 97124-6421
(503) 681-0118

Circle 353

Trademark Information

UNIX is a registered trademark of AT&T Bell Labs.

PAL is a registered trademark of Advanced Micro Devices, Inc.

SMART-POWER is a registered trademark of Nartron Corp.

CDA and BurstRAM are trademarks of Motorola, Inc.

SCOPE and ASSET are trademarks of Texas Instruments, Inc.

IRIS POWERVISION is a trade-mark of Silicon Graphics, Inc.

RealTimeX is a trademark of Concurrent Computer Corp.

Zone Bit Recording is a registered trademark of Seagate Technology, Inc.

ASICs & ASIC DESIGN TOOLS

Xilinx 10,000-gate FPGA supports 32-bit systems

With the 0.8-μm CMOS XC4010 FPGA from Xilinx, the company claims you can comfortably design a 32-bit, 33-MHz microprocessor peripheral control system based on FPGAs for the first time.

The third and densest member of the XC4000 family, the XC4010 has 10,000 usable gate-array-equivalent gates, according to its maker. Moreover, each XC4010 configurable logic block (CLB) can alternatively be configured as a 32 × 1-bit (or 16 × 2-bit) block of SRAM. Xilinx says that an XC4010 design using ⅓ of the blocks for memory and ⅔ for logic could exceed 20,000 usable gates. With the recently introduced X-Blox graphical high-level schematic entry system, you can use XC4000 system features such as fast carry logic and hard macros to build an XC4010 at the block-diagram level.

In addition to having many more routing resources than previous Xilinx FPGAs, the XC4010 has internal three-state buffers that make 40-bit on-chip buses at 45 MHz possible. Among other system functions are 32-bit arithmetic functions and 40-bit external address decoding in 11 ns.

The XC4010 contains 400 CLBs arranged in a 20 × 20 matrix. According to Xilinx, each CLB provides significantly more logic capacity than XC3000 devices, which have a five-input maximum. Any eight-input combinatorial function (and many nine-input functions) can be implemented in a single block of the XC4010, with each block being made up of four modules—two four-input combinatorial modules and two edge-triggered D-type flip-flops. Each module can be independently accessed by placement-and-routing software.

When an XC4010 logic block is used as a 32 × 1 SRAM, the access time is a single block delay, or 4.5 ns, for up to 32 × N-bit-wide on-chip memories. Up to 128 × N-bit memories have access times of about 15 ns, and even larger memories will be accessible in less than 25 ns. Since any number of the XC4010's 400 logic blocks can be used as memory elements instead of logic elements, the XC4010 has the capability to store up to 12,800 bits. Many functions, such as deep register files, small FIFOs, DMA counters, and multiple accumulators, can be implemented with the XC4010's programmable memory resource. Xilinx claims that a 256 × 9 FIFO will run in excess of 25 MHz on the XC4010.

Each internal logic block includes dedicated arithmetic logic for the fast generation and propagation of carry-and-borrow signals. With two counter bits in each internal block, the device can implement nonloadable counters with clock rates of 90 MHz, 16-bit loadable counters at 33 MHz, and 32-bit loadable counters at 28 MHz. That's twice the speed, with half the number of logic blocks, when compared with members of the Xilinx XC3000 family.

The XC4010 has 160 I/Os (7 ns). All outputs can be three-stated on a pin-by-pin basis. Since pairs of I/Os can be connected in parallel to sink up to 24 mA, the XC4010 can drive short buses on PCBs without the need for external drivers.

With the XC4010, Xilinx offers support for JTAG boundary scan. Available now in a 191-pin lead pin grid array, the XC4010 is $737 in quantities of 1,000. —*Barbara Tuck*

XC4010 at a glance

- Can provide 40-bit on-chip buses at 45 MHz
- 32-bit arithmetic functions
- JTAG boundary scan logic
- Library of hard macros
- Graphical high-level schematic entry system

Xilinx
2100 Logic Dr
San Jose, CA 95124
(408) 559-7778
Circle 352

68030 VME

FLASH flexibility.

Tired of hunting down blank EPROMs? GreenSpring's SBC5 with FLASH memory ends hunting season forever. Say "good-bye" to long delays. In-circuit programming means you don't even have to power down.

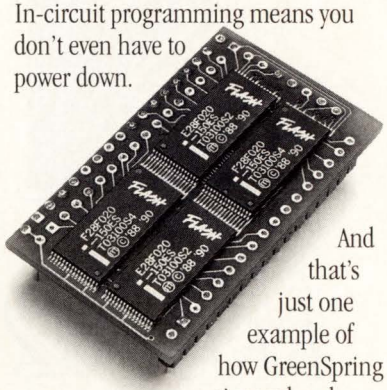

And that's just one example of how GreenSpring engineered real-world functionality into our 32Mhz 68030 VME board. Along with the features you expect, like 4MB dual access DRAM, serial ports, a four-level arbiter and mailbox interrupt.

And some you don't expect, but will learn to appreciate, such as our GLA™ custom bus interface silicon. Not to mention IndustryPack™ slots for almost unlimited I/O options.

NEW VME PRODUCT

Finally, we back you with a 30-day return privilege, a one-year warranty and the best technical support in the business. An unbeatable value at an unexpected low price.

Call or fax GreenSpring today to get a free catalog or a copy of the SBC5 User Manual. Or information on our other exciting VMEbus products. Try us. We're flexible.

GREEN SPRING COMPUTERS

1204 O'Brien Drive, Menlo Park, CA 94025
(415)327-1200/FAX(415)327-3808

CIRCLE NO. 74

Miss the Conferences? Don't miss the Proceedings!

SysComp

SysComp/92 West, the first conference devoted to OEM system/subsystem components, featured a Technical Program with more than 50 presentations on disk drives, power sources, display technology, bus architecture and embedded software.

The **PROCEEDINGS** are a must for systems engineers and engineering managers involved in OEM integration. Cost $150

RISC '92

RISC '92 was devoted *exclusively* to RISC microprocessors (in particular, the Sparc, 29000, 88000, R3000/4000 families and derivatives) and microcontrollers, RISC architectures and software (compilers, debuggers, etc.).

The **PROCEEDINGS** are now available. It contains a wealth of in-depth, applications-oriented information, from more than 60 sessions, for designers of next-generation hardware and software. Cost $150

ANALOG & MIXED-SIGNAL
DESIGN CONFERENCE

The Analog & Mixed Signal Design Conference presented tutorials on topics such as Spice, A-D/D-A converter technology, analog effects in fast digital circuits, as well as lectures on specific applications of mixed-signal technology, CAE/CAD tools for analog and mixed-signal design, ASICs and dealing with ASIC vendors.

The **PROCEEDINGS** contain more than 60 sessions of timely, immediately useful information. Cost $150

To order Proceedings of the **SysComp/92 West, RISC '92** or **Analog & Mixed Signal Design Conference** at $150 each, fill in and MAIL/FAX to *Computer Design.*

(#) _____ copies of the SysComp/92 West Proceedings x $150 =$_____
(#) _____ copies of the RISC '92 Proceedings x $150 =$_____
(#) _____ copies of the Analog & Mixed Signal Proceedings x $150 =$_____
Total =$_____

Check enclosed ❑ Charge to Mastercard ❑ Visa ❑

Card # _____ Exp. Date _____

Signature_____ Telephone _____

SHIP TO:
Name_____
(Please print clearly, this is your shipping label)
Company _____

Address _____

City/State/Zip_____

COMPUTER DESIGN Technology and Design Directions

FAX OR SEND TO:
COMPUTER DESIGN
C/O Patti Kenney
1 Technology Park Drive
Westford, MA U.S.A. 01886
FAX: (508) 692-7780
TEL: (508) 392-2124

CIRCLE NO. 55

CAE/CAD TOOLS

Scicards revamped for high-speed digital design

The latest release of the PCB design suite from Harris Scientific Calculations (Fishers, NY), Scicards Version 27, has been revamped to include design rules for high-speed digital and analog boards. The new version includes an enhanced gridless router which provides intelligent push/shove capabilities with online design rule checking (DRC). The improved DRC can perform physical and logical evaluation of split planes and metal areas. The area fill feature has also been re-engineered for faster response time. The Version 27 gridless router, Freestyle+, has been optimized for speed and can intelligently shuffle trace interconnects to find the best route.

For automatic testing of dense surface-mount and through-hole board technologies, automatic test point generation and output have also been added. Test points, defined according to user-specified features and sizes, are automatically inserted during editing and autorouting. In addition, test points are associated with the logical interconnect and may be output to automatic test equipment.

Communication enhanced

Version 27 also includes enhanced communication capabilities, with environments both within and outside the layout department. The system lets you specify and automatically transfer design rules and other vital design information along with board layout data.

Rules for controlling impedance, parallel trace crosstalk, tandem trace crosstalk, and shielded nets can be entered by the design engineer to control the entire design process. Harris' Freestyle+ autorouter may be directed to route shielded nets as well as nets by area and layer. By using a gridless, shape-based approach, the Freestyle+ autorouter produces trace layouts that are optimized from a global view—your viewpoint—as well as on an individual net basis.

To ensure consistency between design libraries, the communication link with other CAE tools has been improved to accept the logical library definitions generated by design engineers on the schematic. This feature lets you access physical library definitions to generate floor plans to pass on to the layout engineer.

Version 27 can also share databases generated by other layout systems. Translation software lets you port design information from CAE databases into the Scicards environment.

In addition to supporting other industry-standard platforms, Scicards Version 27 has been recently ported to the Hewlett-Packard 700 series workstations. The suite is available now, and is priced at $45,000. —*Mike Donlin*

Scicards Version 27 at a glance

- PCB design tools for high-speed digital and analog boards
- On-line design-rule checking evalutes physical and logical implications of split planes and metal areas
- Automatic test point generation, defined according to user-specified features and sizes, inserted during editing and autorouting
- Freestyle+ autorouter routes shielded nets as well as nets by area and by layer
- Communication links with other CAE tools accept the logical library definitions generated by design engineers on the schematic

Harris Scientific Calculations
7796 Victor-Mendon Rd
P.O. Box H
Fishers, NY 14453
(716) 924-9303
Circle 358

68030 VME

$1295 economy.

Oh sure, you say. Anyone can advertise low OEM pricing. But no one can beat GreenSpring's combo of rock-bottom prices, state-of-the-art performance, powerful features and legendary technical support. Not to mention software, IndustryPacks™, customization and reliability.

Our SBC5 68030 VME board doesn't load you up with "bells and whistles"– just enhanced functionality. Like 32Mhz. 32-bit FLASH EPROM. IP slots. Even GLA™ custom bus interface silicon. What it all adds up to is $1295 OEM/$1695 list – an unbeatable value in high-efficiency VMEbus computing.

Sure, you can settle for false economy… the bargain basement boards with the garage sale performance. Or take a flier on the high-end boards with the stratospheric price tags.

But if speed, cost, options, back-up, software, flexibility and stability mean anything to you, talk to our engineers about your application. Call or fax GreenSpring with your toughest VME questions. We'll send you a free SBC5 User Manual just for calling.

GREEN SPRING COMPUTERS

1204 O'Brien Drive, Menlo Park, CA 94025
(415)327-1200/FAX (415)327-3808

CIRCLE NO. 76

CAE/CAD TOOLS

CAE suite spans platforms, entry methods and technologies

Dazix, a subsidiary of Intergraph (Huntsville, AL), has released a suite of tools dubbed Titan—a software package that combines Cadnetix, Daisy and Intergraph technologies and ports to Sun Sparcstations, Intergraph Clipper workstations and PCs.

At the forefront of the suite is the ACEPlus Design Entry tool, a schematic entry package that targets ASICs, ICs and PCBs. ACEPlus is the front-end link to all other Titan design environments, including digital, analog and mixed-signal technologies for designs that contain any combination of devices from discrete components to MCMs (multichip modules) and PCBs. When a design is created in ACEPlus, a master design database is set up for all future tasks, including design specification, verification, test, layout, placement, and routing. You assign analysis data directly on the schematic for simulation and test and specify any layout design rules directly on the schematic before sending the design to production. After the design is laid out, changes are automatically back-annotated to the schematic to keep the logical and physical data in sync.

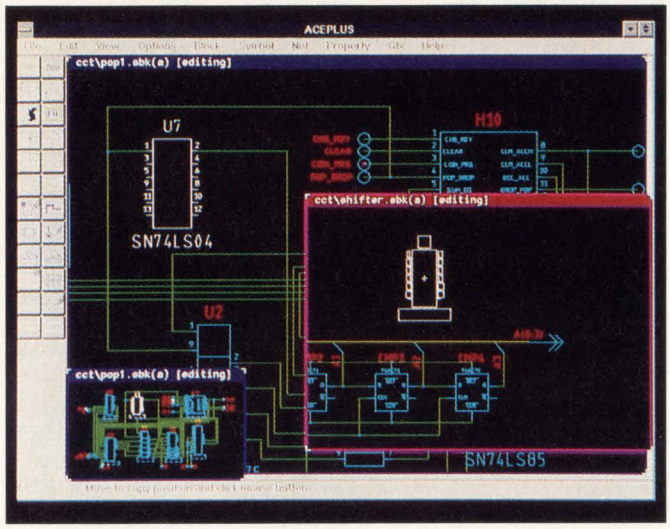

The ACEPlus design entry environment creates a master relational database to access information from various schematics. By using several displays simultaneously, as shown here, you can transfer design information across schematic views.

ACEPlus is tightly coupled to the Dazix Master Librarian, a centralized relational database where component information for all design tools resides. Graphic symbols can be found in the Master Librarian and placed on the schematic by using a combination of component attributes or by querying parametric specifications of parts.

The simulation portion of the package is the AdvanSim logic simulator, which supports VHDL and C language models, as well as software models from Logic Automation (Beaverton, OR) and hardware models. When used in conjunction with Dazix's A/D Bridgeway (ADB), AdvanSim supports mixed-mode simulation. ADB is an intelligent control processor that lets you simulate designs that are predominately digital but that incorporate some analog circuitry. It lets the simulator take a single-engine approach to mixed designs by simulating with its own analog and mixed-signal models or it can link AdvanSim to the Dazix analog simulator, Apex Plus.

Trio of simulation engines

The AdvanSim simulator incorporates three simulation engines—AdvanSim:DLS (Dazix Logic Simulator), AdvanSim:DTV (Dynamic Timing Verifier) and AdvanSim:CFS (Concurrent Fault Simulator). AdvanSim:DLS is a mixed-level logic simulator that verifies the correct logical operation of a circuit. The tool ensures that the test vectors developed and run against high-level descriptions can be reused to verify the synthesized gate-level descriptions. AdvanSim:DTV is a dynamic timing verifier that performs worst-case timing analysis. The tool features unlimited design size and operates from the same open graphical environment as the other AdvanSim simulators. The AdvanSim:CFS engine performs statistical fault simulations. In addition to the common features of fault simulators, AdvanSim:CFS provides a fault partitioning algorithm, which lets the simulator dynamically determine a workstation's resources in terms of swap space and memory and reduces or adds to the simulation fault sets accordingly.

The Titan suite also features the AdvanSim:DLAB waveform editor, a graphical user interface that compiles a simulation database from the schematic, which lets you simulate the design in any of the three modes—logical, timing or fault simulation.

All tools in the Titan suite are available now. ACEPlus for the Sun and Intergraph workstations is priced at $7,000; the PC version at $3,500. The AdvanSim tools are priced individually. AdvanSim:DLAB is $9,000, AdvanSim:DLS is $18,000, AdvanSim:DTV is $15,000, AdvanSim:CFS is $25,000, and A/D Bridgeway is $10,000. *—Mike Donlin*

Titan at a glance

- Comprehensive CAE tool suite for discrete ICs, ASICs, MCMs, and PCBs
- ACEPlus supports VHDL and C models
- The AdvanSim simulator features engines for circuit verification; dynamic timing verification and statistical fault simulations
- Includes the A/D Bridgeway, an intelligent control processor that links simulators for mixed-mode designs

Dazix, a subsidiary of Intergraph
Huntsville, AL 35894-0001
(205) 730-2000
Circle 357

INTEGRATED CIRCUITS

Chip set optimizes 486 caching

There are two ways to improve the performance of 486-based personal computer designs. One is to add a fast external secondary cache. The second is to provide the performance enhancements of a secondary cache by using logic instead of memory. Believing that the second route is a more cost-effective one, Headland Technology (Fremont, CA) has introduced the HTK340, a chip set for 80486 systems that optimizes the performance of internal cache memory in a PC environment.

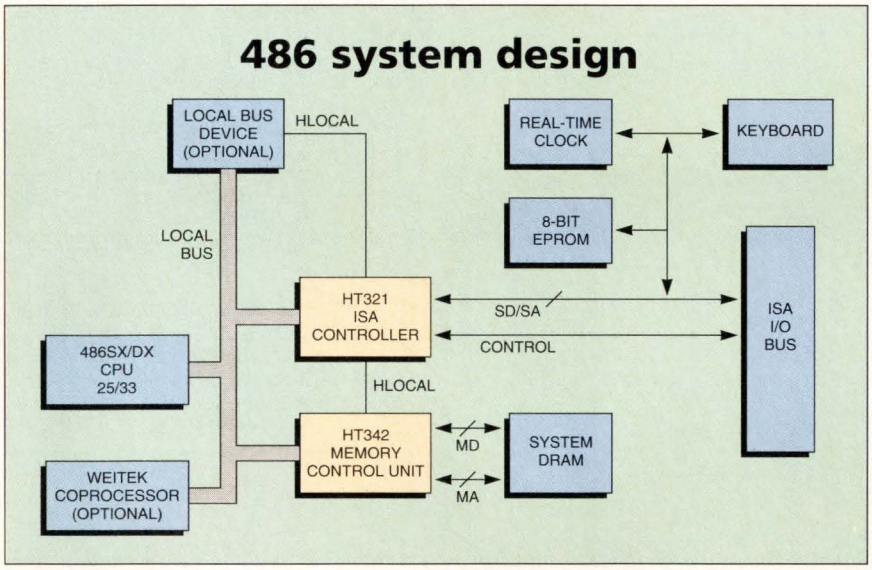

The HTK340 chip set is composed of two chips: the HT321 ISA controller and the HT342 memory control unit. Although benchmarks such as those used by personal computer magazines don't show it, Headland claims the two-chip set dramatically improves 486 performance on real-world applications.

The net effect is that the processor never has to wait to perform write operations, resulting in improved system performance. The HTK340 comprises two devices: an ISA bus controller chip (HT321) and the memory controller unit (HT342). They support system operation at frequencies through 33 MHz and provide local bus attachment of peripherals.

Eliminating the write bottleneck

The 486 architecture provides an efficient internal cache, so efficiently handling memory write operations determines system performance. Headland's write-optimized features eliminate the write bottleneck by integrating a four-level deep-write buffer with out-of-order operations and byte gathering. The result is a two- to four-fold improvement in DRAM write operation efficiency, providing the performance of a secondary cache through logic rather than costly additional SRAMs.

Byte gathering refers to the device's ability to take a stream of sequential addressed data items, either bytes or words (double-bytes), and assemble them into four-byte double words to be written into memory. Supporting features include out-of-order operations and the ability to have read or write hits on the buffer contents. This lets the processor perform a read without having to empty the write buffer. The net effect of these features is that the processor never has to wait for write operations.

Joe Nance, product manager at Headland, says that before the part was designed, his engineering team noticed that the 486 is very effective at performing read operations, because of the 8-kbyte on-chip cache. That means most of the activity on the 486 system bus is writes. And most of that write activity is in a data quantities smaller than 32 bits. "We rarely see 32-bit writes with current benchmarks—it's 8 or 16 bits," he says. With a 486 block move, a memory read is performed at 32 bits, but when it's written to the memory, that's done in bytes.

Boosts not seen in benchmarks

Don't expect popular benchmarks, however, to show the sort of performance boost Headland would like. In fact, the Power Meter V 1.7 benchmark shows no difference whatsoever between a 486SX system with 64-kbytes of cache and an uncached 486SX system built around the new Headland HTK340 device. Both systems rated a mere 8.9 Mips.

Although Nance admits that he was a little disappointed, he attributes the results to what he calls "the simple-minded nature" of PC benchmarks. "PC benchmarks do a lot of writes in tight loops. When they do that they break the pipeline within the processor; the processor itself inserts an idle state after the write so that you never have back-to-back writes to memory. This idle state masks the fact that our chip is never waiting for the processor," he says.

In real-world applications, Nance is quick to point out that his customers are very pleased with the results of the chip set design. Packaged in two 184-pin PQFPs, the Headland chip set has a U.S. domestic price of $45 in quantities of 1,000. One question you must ask yourself is this—which stands to come down faster in price, Headland's PC chip or 64K or more of static memory? For those betting on Headland, production quantities of the device are available now. —*Dave Wilson*

HTK340 at a glance

- First 486 cache enhancement hardware
- Optimizes 486 writes
- Byte-gathering on-chip
- 10 percent or greater performance increase

Headland Technology
46221 Landing Pkwy
Fremont, CA 94538
(510) 623-7857
Circle 356

A Complete Industrial-PC

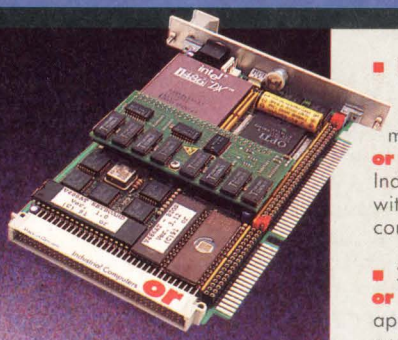

■ **Robust 19" Format**
Put an end to the infuriating problems of non-standard mechanics and poor quality. **or** now offers you a complete Industrial-PC in robust 19" format with reliable VMEbus DIN-connectors.

■ **286-AT, 386-SX and 486-SX/DX**
or has a solution for every application with the right price-performance ratio.

■ **Silicon-Disk, LCD, EL and Plasma**
Does your application demands a CRT, LCD, EL or Plasma-display, or do you need a silicon-disk? Just plug in a board – done!

■ **Low Power CMOS –25/+85° C**
Due to low power CMOS technology boards from **or** guarantee long life and high reliability. The typical temperature range of –70° to +130° C provides a dependable safety margin.

■ **More Quality for your Money**
High quality and newest technology "Made in Germany" must not cost more than other Industrial-PCs. Ask us for an offer, we will prove it.

or

Industrial Computers
Sieglindenstr. 19½, D-8900 Augsburg 1
Tel. +49 821 5034-0, Fax -119

In the US and Canada contact:
DYNATEM
15795 Rockfield Boulevard
Suite G, Irvine, CA 92718
Tel (714) 855 3235, Fax 770 3481

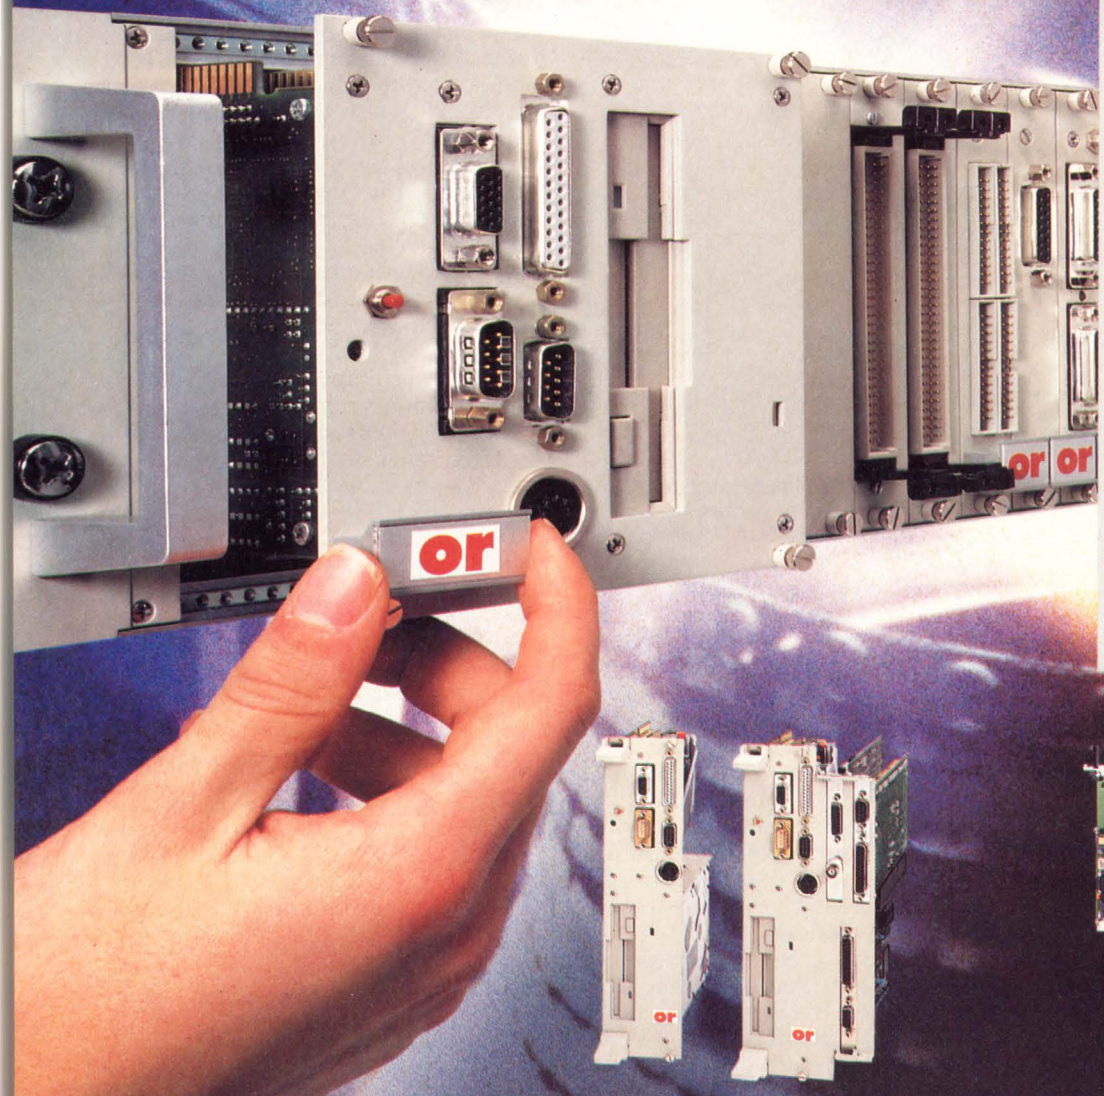

■ **Plug it in – finished!**
The 19" plug-in PC-Boxes are complete Industrial-PC's including floppy and harddisk. Just plug them in – finished. Three versions are available, one in 3U and two in 6U height. One of the 6U versions has 4 PCbus slots for your special PCbus boards.

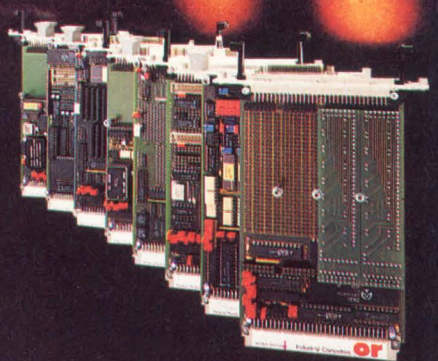

■ **Extendable on the PCbus and VMEbus**
With 440 VMEbus boards **or** has one of the largest industrial-I/O board ranges: digital I/O, analog I/O, counters, motion-controller and more. With Industrial-PCs from **or** both of the most widely used bus-systems are available for your applications: the PCbus and the VMEbus!

INSTANT DATA ACCESS (IDA) DIAL (617) 494-8338 DOCUMENT NO. 1032

CIRCLE NO. 72

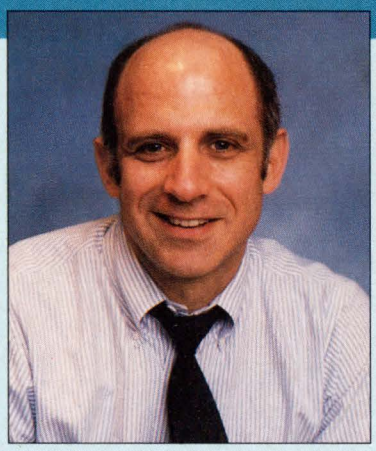

Where speed replaces precision

Consultant Andy Rappaport recently upset the analog design community with a newspaper column entitled "Whither analog." The argument advanced by "Whither analog" goes something like this: digital circuitry and design techniques will replace analog, and the proportion of the electronics world affected by analog design will be smaller and smaller. Frankly, that's no surprise to readers of *this* magazine, but Rappaport took this argument to a wrong and rather careless conclusion—he concluded that analog building blocks will be about as useful as 7400-series logic.

Prices for LM10s, OP-07s or LT-1028s may have dropped precipitously over the past five years, but implementing a circuit with one of these devices requires considerably more skill than implementing a logic circuit with a 7404. As I've suggested many times before, digital logic circuits do little more than register binary state changes (albeit at up to 100 million times a second), while analog circuits exercise *control* over changes in voltage and current.

The amount of control—or the degree of precision—is the major issue that affects analog design today. While analog circuitry takes up a smaller proportion of any system, one could argue that this proportion requires a higher degree of precision. Contrary to Rappaport's assertion, this isn't the playground for commodity op amps and 8-bit data converters.

Mixed-signal design has become of increasing concern to systems engineers, primarily because their work more and more requires manipulation of analog voltages and currents. In one Starch readership survey, 55 percent of *Computer Design* readers queried replied that they were involved in some sort of mixed-signal design. Examples of mixed-signal operations abound in the ordinarily digital world of computers—for example, disk drive read channels and modems transform a digital pulse stream into an analog waveform for recording or broadcast. Graphics and image processing, and even multimedia applications, are bounded by D-A converters which control the color and intensity of each pixel on a CRT screen. Designers of portable equipment, meanwhile, increasingly need to monitor the current consumed by system components and to power down unused elements in an effort to conserve battery life.

Whether these mixed-signal designs are classified as "big D, little A" or vice versa is hopelessly beside the point. They all need to manipulate voltages, currents and waveforms, and, typically, the digital designer needs to incorporate something more into the mixed-signal design suite than the data book parameters for a 741.

The "encapsulation of engineering knowedge"

Prior to the introduction of its Mixed-Signal Design System (MSDS) last year, a representative of Gould/AMI (Pocatello, ID) told me the company was looking for a way to "encapsulate" analog engineering knowledge. "There aren't enough analog designers," said Gould/AMI's CAE department manager Bryce Baker. "We're trying to build a system that captures what the analog expert knows."

The promotional literature for MSDS says it "enables systems developers with little or no experience in the design of analog ICs to incorporate analog functions precisely tailored to their application." The heart of MSDS is an analog model builder which works in conjunction with a configurable cell generator. At the time of its release, the cell library had 17 of the most common analog functions, with more in development. You simply need to specify the parameter requirements of your application, and the MSDS generates cells and simulation models. The Design Critiquer, an expert system within MSDS, will in fact check the generated cells for analog IC design errors.

Systems such as MSDS, I believe, are forerunners of the way analog and mixed-signal ASIC design will be performed in the future.

While there are toolsets to support precision analog design, there are currently no tools which will encapsulate this knowledge for use by the digital design community. Other than in magazine articles, columns and interviews, how does anyone capture the engineering knowledge of a Dave Fullagar of Maxim Integrated Products, or a Bob Pease of National Semiconductor, or a Jim Williams of Linear Technology, or a Gerry Graeme at Burr-Brown? So valuable is the

knowledge of these designers that Burr-Brown's management recently wondered aloud whether the real salable product of its company wasn't precision op amps and data converters at all but engineering man-hours.

The secret to encapsulating what these gurus know is modeling. We need to construct models for the simple but masterful techniques these designers use to effect precise control of a waveform. Modeling precise analog behavior will be more difficult than modeling digital behavior, because the control of voltages and currents forces you to get much closer to the components that produce them. Digital circuits can be described and modeled with 1s, 0s and clocks—it makes no difference whether you're describing transistor switches, clicking relays or flashing lights. The sequence and pattern of binary 1s and 0s are totally removed from the circuitry that implements them. With analog design, you're much closer to the resistors, capacitors and op amps. The genius of a Bob Pease, a Jim Williams or a Gerry Graeme is that they know instinctively what transistor type, what amp, what resistor, what capacitor values will yield a precise transfer function.

Changes in craftsmanship

Yet, there's a strong movement underway to put the equivalent of Bob Pease on a chip. "Analog designers have been 'craftsmen' for the longest time—but that is changing," says Ian Getreu, vice-president in charge of modeling at Analogy (Beaverton, OR), whose Saber behavioral-level simulator is used by mixed-signal ASIC vendors such as Gould/AMI and NCR's Microelectronics Division. "What's changing the demand for craftsmanship," he suggested recently, "is time-to-market pressures." Current-generation analog design tools will likely give up silicon real estate through the use of automated layout on grids, although analog designers will resist this. There will also be trade-offs in precision, but these sacrifices must be weighed against costs and the need to introduce a new product before the competition. "In general, the more precision you need," says Getreu, "the less you can use automation. But there is always a question as to how much precision is necessary."

One style of emulation craftsmanship looks to analog models and modeling activities, which describe designs in very abstract forms. Implementation alternatives are considered later. "In behavioral modeling," says Analogy's engineering manager for simulation, Ernst Christien, "the important issue is not the modeling of physical elements, as it is with Spice. The most important concept is the modeling of continuous time." If you accurately describe a transfer function, Christien suggests, you can find ways to implement it.

This is most apparent in the digital signal processing world, where toolsets from companies such as Comdisco Systems (Foster City, CA) let you define precise filter functions almost without regard to the silicon that implements them. The algorithm can then be downloaded into a Synopsys (Mountain View, CA) logic synthesis program and implemented with a gate array or FPGA. This design-before-implementing style is also used with Burr-Brown's Filter Pro software and the filter design set from Star Semiconductor (Warren, NJ).

MHDL arrives

The separation between an abstract algorithm or transfer function and its silicon implementation is growing for general-purpose analog as well as for DSP. DARPA has funded (through the U.S. Army Laboratory Command at Ft. Monmouth, NJ) the development of a millimeter and microwave IC (Mimic) hardware description language (MHDL). Many authorities believe that MHDL will be useful for describing precision analog as well as microwave circuits. Analogy's Christien, in fact, is a participant in the MHDL working group, and Analogy's Mast modeling language is a strong contender among the tools being considered by DARPA contractors for describing analog activity in the abstract.

With analog behavior abstractly modeled, it's possible and perhaps necessary, says Christien, to rethink how a solution's implemented. Solving analog problems with digital techniques—analog emulation—is the most obvious trendline, since it takes advantage of the most widely available design tools and manufacturing techniques.

Speed rather than precision

One example of analog emulation in which a precision data converter is rendered entirely with digital design techniques and processes—a lot of cheap CMOS gates—is sigma-delta conversion. Sigma-delta D-A converters produce 16- and 18-bit accuracy at a fraction of the cost of other converter technologies. The essence of the sigma-delta technique is clock speed; you raise the sampling frequency of an A-D converter so high—to 12 MHz—that the magnitude of the difference between one sample and the next becomes minuscule. In effect, you exchange resolution in amplitude for resolution in time. Instead of a number reflecting 16- or 18-bit precision—for example, a 96-dB dynamic range between the highest and lowest recorded amplitudes—each sample can be described with a 1-bit number, which only needs to describe the direction, up or down, the waveform is likely to take in the next sample.

With conventional data conversion processes, a sampled analog waveform is divided into a series of digital slices according to the amplitude of the signal. The bit resolution of the data converter is a measure of how fine the amplitude of the analog waveform must be sliced.

The focus on bit resolution assumes that the analog waveform is sampled at a rate (or frequency) which is just high enough to capture the highest analog frequency. The conventional sampling frequency is twice the highest frequency of the captured signal. For an audio bandwidth of 20 Hz to 20 kHz, the sampling frequency must be at least 40 kHz. Most

You probably need one.

The *1992 Arrow Systems Product Guide* features 12 product selection guides, a tutorial overview on each product group, and product warranty and repair information. For the fourth year in a row Arrow Electronics has chosen to make its catalog easily available to many of our *Computer Design* readers by mailing the guide with this April issue.

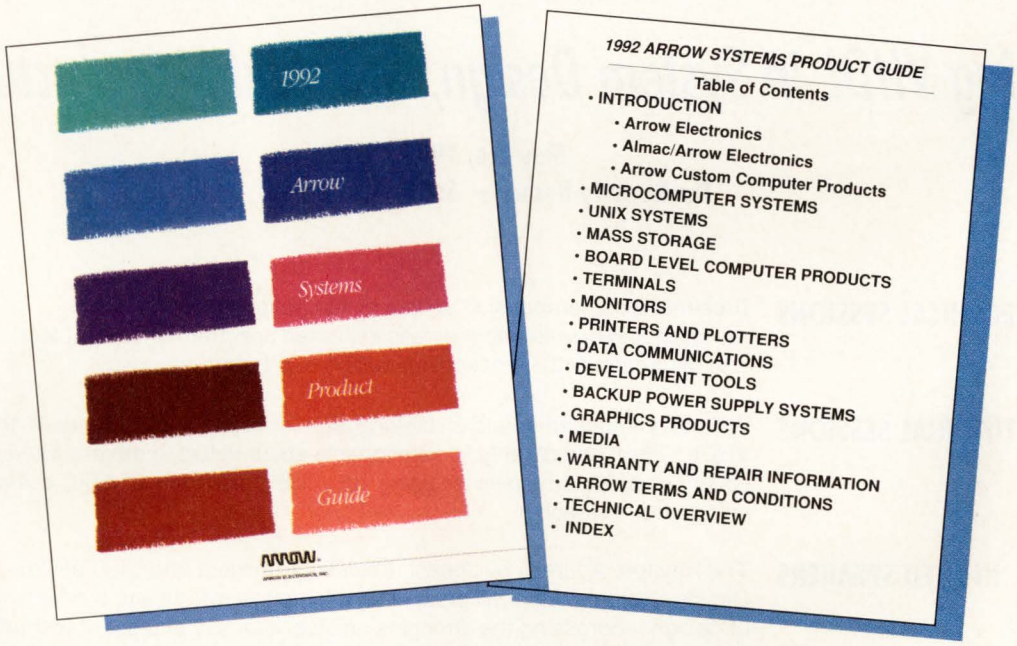

It's a very useful guide for anyone actively involved in the selection of computer products, subsystems, and/or boards. If you don't have your copy of the *1992 Arrow Systems Product Guide*, or someone has "borrowed" it, fill out and return the card below to receive your own copy of this important product resource.

TO: **Chris Johnson**
ARROW ELECTRONICS, INC.,
25 Hub Drive, Melville, NY 11747

[ ] Please rush me the
1992 Arrow Systems Product Guide.

Name_____

Company_____

Address_____ M/S___

City_____ State____ Zip_____

Phone _____

A91606
4/92

Dept I 04

Job Function:

[ ] Purchasing/Mgmt. [ ] Engineering/Mgmt
[ ] Executive Mgmt. [ ] Other

Please send me additional information on ARROW ADVANTAGE® programs:

[ ] Custom Computer Products
[ ] Programming Centers
[ ] Kitting Services
[ ] Arrow's In-plant Terminal Program

[ ] Please have an Arrow salesperson contact me.

CIRCLE NO. 65

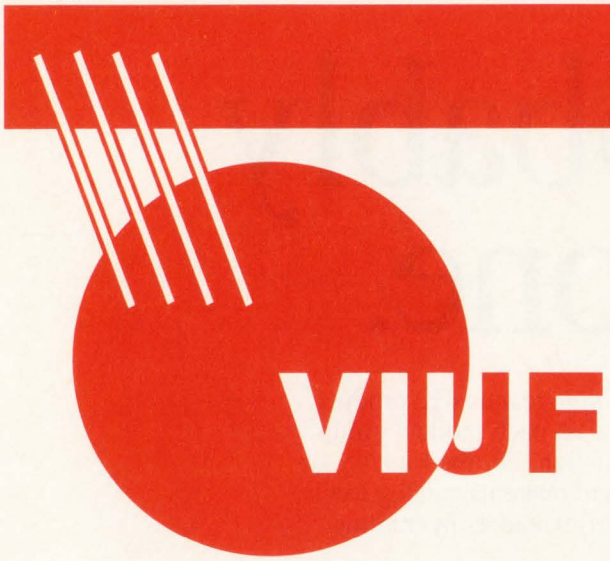

Spring 1992 Conference

Using VHDL in System Design, Test, and Manufacturing

May 3-6, 1992
The Registry Resort – Scottsdale, Arizona

TECHNICAL SESSIONS Presents a comprehensive update of international VHDL-related research and development across the broad spectrum of design, test, and manufacturing. Proceedings will be distributed.

TUTORIAL SESSIONS Will provide fundamentals of the language for those wanting to learn about VHDL. For those already knowledgeable about VHDL, there are tutorials covering advanced topics including 1992 IEEE revalidation, ASIC design, and guarded signals.

INVITED SPEAKERS The Keynote Address by Joseph Costello, President and CEO of Cadence Design Systems, Inc., will present an informative insight into the "linguistics" of design, addressing the strengths and weaknesses of present technology and the challenges that lie ahead. Representatives from the Japanese and European VHDL Spring 1992 conferences will give an overview of these international activities. Banquet speaker, Dr. Roger Lipsett, director of the initial language design effort, will offer an interesting historical perspective on the development of VHDL.

SUPPLIER SUITES Informal Supplier Demonstration & Hospitality Suites offer attendees the opportunity of gaining hands-on experience with the state of the art products from leading vendors.

REGISTER BY APRIL 17TH TO RECEIVE THE EARLY REGISTRATION DISCOUNT!

Sponsored By:

VHDL INTERNATIONAL

For Further Information Contact:

VHDL INTERNATIONAL USERS' FORUM
c/o Conference Management Services
407 Chester Street
Menlo Park, CA 94025
1-800-554-2550 (FAX) 415-324-3150

CIRCLE NO. 79

audio digitizers will sample at 44.1 kHz.

Suppose for a moment that the audio waveform was sampled at a rate much higher than 44.1 kHz. A sampling frequency of 11.2896 MHz, for example, represents a 256ö oversampling rate. The requirement for a 16-bit data converter effectively disappears, because the amplitude range of the signal within time slice would be practically flat. With strenuous oversampling, the only decision the data converter could make is whether the signal in any time slice is higher or lower than the signal in the previously-sampled time slice. This simple 'yes' or 'no' (1 or 0) could be handled with a 1-bit data converter called a "delta modulator." Sigma-delta builds on delta modulation by weighting (or summing) the results of previous conversions, and using this in the decision.

Using sigma-delta conversion techniques

As implemented by Philips/Signetics (Sunnyvale, CA), Crystal Semiconductor (Austin, TX) and others, the sigma-delta conversion technique builds highly accurate digital audio playback converters—essentially, 'one-bit' converters with 16-bit accuracy. It accomplishes this by transforming the 16-bit pulse-code modulation (PCM) serial data coming off a CD into the same kind of 1-bit pulse stream that would be created by a sigma-delta A-D converter. The oversampled 1-bit data stream is extrapolated from a binary data sample using DSP techniques.

Because the circuits of a device such as this are dependent on a system clock many times faster than the captured audio waveform, digital clock speed becomes the replacement for analog precision. But there are many mixed-signal applications apart from digital audio conversion in which heightened timing accuracy can replace precision voltage measurement. These include disk drive read channels and data modems in which a serial data stream is converted to a variable-frequency analog waveform.

If analog emulation is the wave of the future—where designers of precision analog circuits will inevitably be dragged "kicking and screaming"—the *encapsulation* of engineering knowledge may be less an issue than its *transformation*.

Stephan Ohr is a consultant with Indian Forest Research and editor of the monthly newsletter, Mixed Signals.

The Superior VMEbus Analyzer

Only VMETRO's VBT-321B and the Modular VMEbus Analyzer System offers a complete system solution to your development needs.

- **100 MHZ TIMING:** Waveforms and State listing shown on time correlated split screen.
- **Integrated VMEbus Anomaly Trigger (VBAT):** Provides on-screen explicit violation messages.
- **Separate analysis of P2-busses:** VSB, SCSI or user-defined P2-bus analyzed simultaneously with VMEbus.
- **256K deep Trace with dump to SCSI disk:** For archival or post-processing. Continuous SCSI download while sampling also possible.
- **VMEbus Master and Slave simultaneously with VMEbus analysis:** Implemented with VIC068 chip to provide real VME cycles. Powerful commands for test pattern generation.

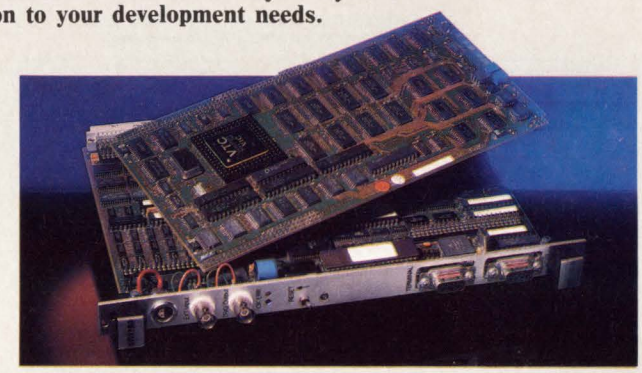

The VBT-321B Advanced VMEbus Analyzer with a proper piggyback module solves any kind of VMEbus development task. The piggyback functions are also available as standalone products, independent of the VBT-321B.

To develop the best products, you need the best tools!

VMETRO A/S
Professor Birkelandsvei 24, P.O. Box 213
Leirdal, 1011 Oslo 10, Norway
Phone: +47 2 322580 Fax: +47 2 322880

VMETRO
The Bus Analyzer Specialist

VMETRO, INC.
16010 Barker's Point Lane, Suite 575
Houston, Texas 77079
Phone: (713) 584-0728 Fax (713) 584-9034

CIRCLE NO. 80

Electro® International

May 12-14, 1992
Hynes Convention Center • Boston, MA

The largest electronics engineering expo in the Eastern United States

Recharged

The pulse of the hottest innovations in the electronics industry.

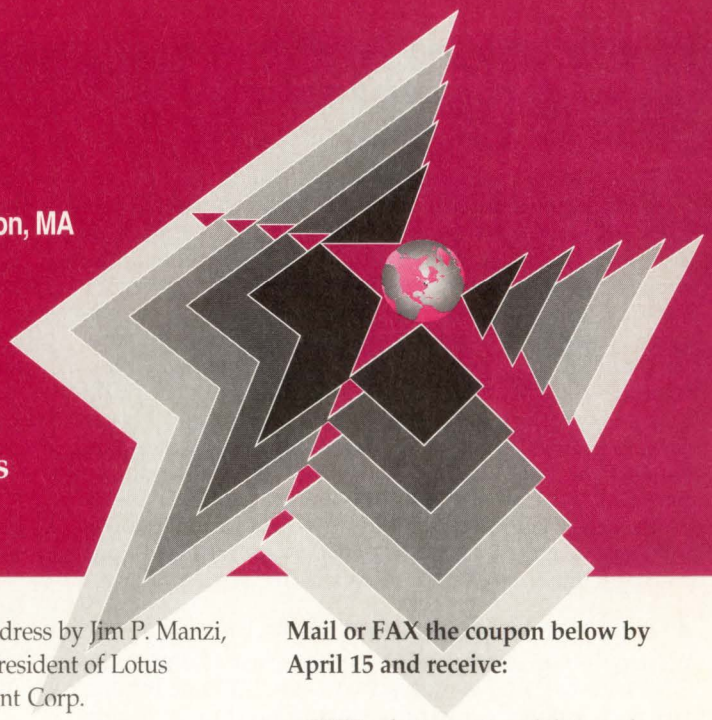

Electro returns to Boston with an expanded technology focus, featuring:

- CUTTING EDGE SOFTWARE for the engineering environment, including
- One-day Windows seminar
- CEO panel discussion of software and PCs
- All-day seminar on the Demeter method for object-oriented design
- Keynote address by Jim P. Manzi, CEO and President of Lotus Development Corp.

PLUS

- All-new state-of-the-art Semiconductor exhibit area on the show floor
- Special test and measurement section featuring up-to-the-minute developments in test and measurement equipment

Mail or FAX the coupon below by April 15 and receive:

- FREE admission to exhibits and technical sessions
- Full preview of Electro/92 technical program
- Chance to win FREE entry to a one-day Windows class taught by industry experts (a $345 value)

FORM MUST BE COMPLETE TO BE PROCESSED.
Please check ALL applicable selections

JOB FUNCTION

G1	Engineer
G2	Project Manager
G3	Senior Management
GB	Design
GC	Test
GD	Mechanical/Packaging
GE	Research & Development
GF	Manufacturing/Production
GG	Quality Control & Assurance
GH	Software
GI	Purchasing/Procurement
GJ	MIS
GK	Engineering Technical Support
GL	Consultant/Educator
GM	Student
GN	Sales/Marketing
GO	Rep/Distributor
GP	Other

BUSINESS CLASSIFICATION

HA Advanced Entertainment Electronics
HB Aerospace/Aircraft
HC Computers/Peripherals/Data Processing
HD Consumer Electronics
HE Contract Engineering/Manufacturing
HF Education
HG Electronic Components
HH Instruments & Controls
HI Medical Electronics
HJ Military/Government
HK Office or Business Machines
HL Research & Development
HM Tele-Data-Communications
HN Transportation
HO Utilities
HP Other

PRODUCT INTEREST

IA Semiconductors; ICs
IB Passive Components
IC Electro-mechanical Components
ID Mechanical/Packaging Components
IE Fiberoptic & Opto-electronic Devices
IF Microwave/RF comp., Instruments
IG Test, Measurement, Inspect., Instruments
IH Automatic Test Equipment
II Engineering Workstations/Peripherals
IJ EDA Tools (CAE/CAD/CAT)
IK CASE Tools
IL Local Area Networks
IM Computer Subsystems
IN Power Supplies, including Batteries
IO Production Equipment/Machinery
IP Production Supplies/Materials
IQ Contract Design/Contract Manufacturing
IR Other

Electro® International

COMPLIMENTARY REGISTRATION

9-90162

Please PRINT as you want shown on badge

NAME (FIRST OR INITIALS) LAST

POSITION AREA CODE NUMBER

COMPANY

MAILING ADDRESS DEPT. OR M/S

CITY STATE ZIP CODE

Important Note:
Please mail completed form to:
Electro Registration,
P.O. Box 92275 WPC
Los Angeles, CA 90009-2275
If received before April 15, your badge will be mailed to you. If received after April 15, please pick up your badge at Will Call. Or, bring this form to Electro for free admission

Help your company bring in the business at Electro.

Your company can also benefit from exhibiting its products to Electro's prime audience of highly-qualified decision makers.
For more information about attending or exhibiting in Electro/92,
call 1-800-877-2668 or FAX to 310-641-5117.

CIRCLE NO. 81

LETTERS, COVERS & OTHER DIVERSIONS

We had more than the usual amount of fun with the February cover except everyone around here kept saying "ribbit, ribbit, ribbit" all the time. Seriously, though (well, not really), the cover did say quite a lot and a couple of our readers identified some of the metaphors we dreamed up.

The simplest to decipher were probably the Siamese frogs and the pair of frogs with their arms together, representing tightly coupled and loosely coupled, respectively.

As Warren Andrews points out in the article, however, it's possible to have an architecture that's somewhere between loosely and tightly coupled—known in some circles as snugly coupled—and represented here by our two affectionate frogs. Tightly coupled systems are also characterized by shared memory, which is what the frogs looking at the photo album are doing.

The two frogs conversing over the telephone are engaged in distributed processing and the line-up of 680... frogs represents pipelining. The three frogs being propelled off the board by the 960 frog represent superscalar architectures, or simultaneous instruction launch. And would you expect the bees to represent anything but buzzwords?

I was happy to see another cover contest this month. I did decode the last one correctly but I didn't bother sending it in. So as not to make the same mistake again, here's my best shot for the February cover.

I see the concentric rings as ripples spreading through the computer industry caused by the over-inflated 386. The 960 has just launched parallel processing (the three side-by-side frogs) from the board-level arena. The Siamese frogs represent tightly coupled multiprocessing while the two frogs with arms interlocked are loosely coupled systems. The love birds (frogs) stand for compatibility. The blabbermouths on the phones, signal processing. The two frogs enjoying the family photo album must be image manipulation and the row of brownies are the march to multiprocessing. Why my mother-in-law is in the bottom left corner is anyone's guess.

The guy sitting down with the laptop is a software developer, paying close attention to the industry buzzwords (the bees). Last, but not least, are the three befuddled users in the corner, soaking it all in.

P.S. Better retire that 386, he looks old and tired.

Name withheld
to avoid a family squabble
Cambridge, OH

Working from the bottom up, the frogs are processors, rings are performance. The 386 is at the lowest performance level. The 960 is using "off-board processors." The two frogs with arms linked are loosely coupled. The two frogs joined at the back have shared backplanes and they are "leapfrogging" the single processor. On the next level, the two frogs reading the photo album are "shared memories." On the phone, they are "remotely networked" and "loosely coupled."

The man with the laptop is keeping his eyes on the latest buzzwords. The row of frogs labeled 680... are a rack of VME processor boards in a single backplane.

Yos Feit
Telerate Systems
Jersey City, NJ

Dave Wilson's "Wrestling with multimedia standards" and "DVI gets slicker software" in the January issue were both very well written and informative. Suffice it to say that your magazine is very useful in helping make decisions. Thank you.

M.J. Khisty, Principal Engineer
Geosystems
Omaha, NE

Good info in January. Thanks, too, for the 800 phone numbers, vendor locations and regular phone numbers.

Do Ethier, Manager
Ethier Enterprises
White Bear Lake, MN

Thanks for appreciating all the hard work that goes into confirming those numbers and locations.—Lisa Coleman, Assistant Editor

I really enjoyed your fine article "Wrestling with multimedia standards" in the January issue. You did a very thorough exploration of what's happening in the multimedia market today. The information was well presented and informative.

Nina Price
Digital Equipment Corp
Palo Alto, CA

Thanks for the SRAM article. What with processor speeds constantly going up, the RAM has become a critical path to a maximal design.

Robert Kropp, President
Catalina Research
Oro Valley, AZ

The article "Specialty SRAMs sprint in step with speedy CPUs" by Jeffrey Child was excellent, informative, great!

R. Venne, Manager
Venne Data Systems
Sudbury, MA

Glad you liked my SRAM article. Your feedback is important to us. What other SRAM topics would you be interested in reading about—Cache design trade-offs? SRAM cost trends? Ultra-fast SRAMs?—Jeffrey Child

THIS MONTH'S COVER:

There's a mixture of rebus writing and symbolism in this month's cover. There are several FPGA vendors represented in the illustration as well as some other references to programmable architectures, so look carefully. Since Barbara Tuck refers to all of these in her article starting on page 88, however, deciphering the cover should be a snap.

SYSTEM SHOWCASE

New for the 8051

Professional 'C' Developers Kit!

- Fastest "C" Execution, Most Compact Code.
- Mixed Memory Models.
- RTOS & Memory Banking Extensions.
- Accommodates All 8051 Derivatives.
- RTX-51 (Real-Time Multi-Tasking for 8051) included, and BL51 Banking Linker/ Locater (Manages up to 1 Meg of code) available!
- Source Level Debugger/Simulator with Emulator Interface and Monitor.

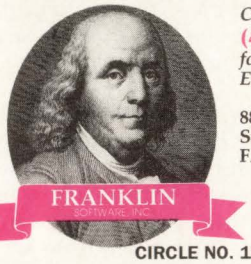

Call Today:
(408) 296 -8051
for your Free Evaluation Package!

888 Saratoga Ave #2
San Jose, CA 95129
FAX: (408) 296-8061

CIRCLE NO. 166

68HC16
68332

PC-based In-Circuit Emulator

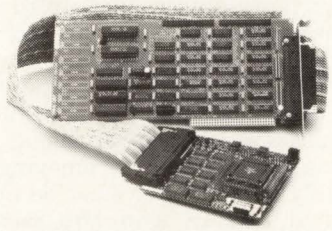

- Supports 68HC16Z1, 68331, 68332, 68340 with more members of the HC16 and 300 families to follow.
- User interface under Microsoft Windows.
- Memory contents available in real-time (shadow RAM).
- Up to 17 MHz real time emulation.
- Full control of TCNT at breaks and single step.
- High-level C support. In-line assembler and disassembler.
- 10-day free trials available to qualified customers.

nohau CORPORATION
51 E. Campbell Avenue
Campbell, CA 95008
FAX (408) 378-7869
(408) 866-1820

CIRCLE NO. 167

QUICK-TURN PC BOARDS

INSTANT QUOTES

1-800-234-1556
(After Hours (602) 377-7319)

Single/Double Sided/Multi-Layer
Gerber Input/Floppy or Modem

VISA MIL Approved U.L. Listed MasterCard

AVANTI CIRCUITS, INC.

17650 North 25th Avenue
Phoenix, Arizona 85023

(602) 863-7729 & 234-1556
Fax (602) 375-1909
BBS/Modem (602) 234-1737

CIRCLE NO. 168

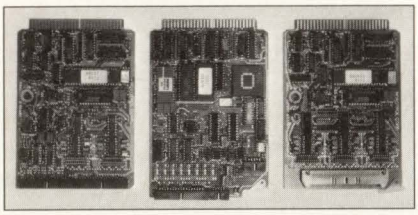

STD ANALOG I/O GET A COMPETITIVE EDGE!

Density—up to 16 A/D inputs and 8 D/A outputs on one card. **Speed**—throughput rates up to 59 kHz. **Intelligence**—many pre-programmed modes. Input filters, prog. gain amps, sample-holds, FIFO I/O buffers & many other features too numerous to mention here.

ROBOTROL CORP.
925 W. San Martin Ave, PO Box 990,
San Martin, CA 95046
(408) 683-2000

CIRCLE NO. 169

FULLY INTEGRATED, RACK MOUNT AND RUGGED SUN SPARC WORK STATION

STANDARD FEATURES INCLUDE:
- SPARK ENGINE 2 CPU WITH 16MB RAM
- 2107MB HARD DISK, 150MB TAPE DRIVE
- 644MB CDROM, REMOVABLE HARD DISK
- 1.44MB FLOPPY, 2 RS-232 AND S BUS PORTS
- SCSI-2 AND ETHERNET INTERFACE
- 16 INCH RACK MOUNT COLOR MONITOR
- KEYBOARD, MOUSE AND SunOS 4.1

FOR CUSTOM CONFIGURATIONS AND FURTHER DETAILS CONTACT:
IBI SYSTEMS INC.
6842 NW 20th AVE, FT. LAUDERDALE, FL 33309
305-978-9225 FAX: 305-978-9226

CIRCLE NO. 170

ORCAD OFFERS DESIGN POWER FOR DOS AND UNIX USERS

Our DOS-based *Schematic Design Tools* are known for being the best price/performance value in the EDA industry. With the introduction of OrCAD's new line of UNIX® based products for the Sun SPARCstation™, OrCAD defines a new standard for ease-of-use, power, and value. **Call (503) 690-9881 for information.**

OrCAD®

CIRCLE NO. 171

Instant Microcontroller

Instant C Programming

Don't use a microprocessor, use a **SmartBlock**™ microcontroller module to build your custom controller. Our $195 interactive **Dynamic C**™ development system makes programming easy. 3.5 x 2.5 inch module includes microprocessor, memory, time/date clock, eeprom, watchdog, serial ports and more. As low as $59 in quantity. The efficiency of a custom design without the headaches.

Z-World Engineering
1724 Picasso Ave., Davis, CA 95616 USA
Tel: **(916) 753-3722**
Regular Fax: (916) 753-5141
Automatic Fax: (916)-753-0618
(Call from **your fax**, request data sheet #14.)

CIRCLE NO. 172

BACKGROUND MODE EMULATOR

68332
68340
68HC16

and 68330, 68331, 68333

Complete hardware testing and C source level debugging for thousands less than an ICE. Call for a *free* trial unit.

(617) 828-5588

EST Embedded Support Tools Corp.

10 Elmwood St. Canton, MA 02021

CIRCLE NO. 173

OEM AIR PUMPS/ COMPRESSORS

- Only one moving part • Oilless
- Low power • Low vibration
- AC & DC models

Call us to discuss your application
TEL: 708-860-0500
FAX: 708-860-9473

MEDICAL/SURGICAL • LEAK TESTING
BUSINESS MACHINES
AERATION SOURCES
VENDING MACHINES
ENVIRONMENTAL
GRAPHIC ARTS

Send for FREE catalog

MEDO USA, INC.
808-A No. Central Ave.
Wood Dale, IL 60191

CIRCLE NO. 174

SYSTEM SHOWCASE

CMOS 186
Single Board Computer
MULTI-TASKING
Runs C or QuickBASIC™ Programs

Powerful 16-bit computer directly executes EPROM's containing any C or BASIC .EXE file. NO LOCATORS! Software includes multi-tasking, multi-drop comm, PID control, OPTOMUX.™

- *10, 12, 16 MHz 80C186*
- *CMOS design*
- *512K RAM*
- *384K EPROM*
- *STD BUS Expansion*
- *COM1 RS232/485*
- *COM2, LPT1*
- *RTC Avail*
- *80C187 Avail*
- *OEM discounts*

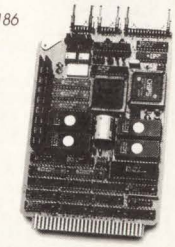

MICRO/SYS

1011 Grand Central Ave., Glendale, CA 91201
(818) 244-4600 FAX (818) 244-4246

CIRCLE NO. 175

T-10 UNIVERSAL PROGRAMMER

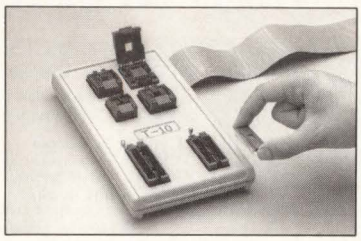

EPROM ■ PLD ■ Bipolar PROM

- Built-in LCC & SMT ZIF sockets to 44 pins
- Individual pin driving
- Point & shoot menus, with MOUSE support
- Device updates on disk & BBS
- Built-in margin testing
- Vector testing
- User definable test parameters
- Full screen editing in 21 formats
- Intelligent identifier

SUNRISE ELECTRONICS, INC.
524 S. Vermont Ave. ■ Glendora, CA 91740
(818) 914-1926

CIRCLE NO. 176

357 MHz 14-Pin Clock Oscillator

Connor-Winfield Corp. announces literature on a new line of high frequency ECL Logic clock oscillators. The ECLB Series oscillator covers the high frequency range of 8 MHz-360 MHz. AVAILABLE NOW IN A 14-PIN DIP. Literature for a double DIP version with frequencies to 500 MHz is also available. Frequencies to 90 MHz are now available with Voltage Control function (model EV53 series). Contact Barney Ill for details.

Additional specifications listed below:

Model	:	ECLB
Package	:	14-pin DIP
Frequency	:	8 MHz to 360 MHz
Supply	:	-5.2 or -4.5 Vdc

Connor-Winfield Corporation
1865 Selmarten Road
Aurora, IL 60505
PH: (708) 851-4722 FAX: (708) 851-5040

CIRCLE NO. 177

REMOVE HARDWARE LOCKS

Protect your investment!

Don't wait for your lock to get lost, stolen or fail and render your program useless.

Our **SAFEKEY'S** are guaranteed to unlock the following:

PADS 2000	MAXROUTE	CADKEY
PADS PCB	MASTERCAM	PCAD
SMARTCAM	MICROSTATION	& MORE

Limited Time Offer $399 plus S&H
(609) 390-2799
24 HR. FAX HOTLINE
(609) 390-3750
IMAGINE THAT
1 Evergreen Dr., Seaville, NJ 08230
TM-Trademarks property of respective owners.

CIRCLE NO. 178

THE FIRST PC BASED PBX BOARD

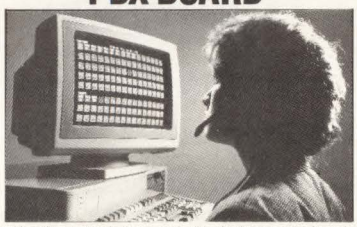

- Handles 4 trunk lines and 12 telephones per board, expandable to 48 telephones per PC.
- Can program any telephone set to individual specifications.
- Standard PBX capabilities at a fraction of the cost.
- Concurrent operation with other PC programs.
- Easy to use menu driven screen displays.
- Integrates with voice mail, fax or modem cards.

QUANTITY & DEALER/OEM DISCOUNTS AVAILABLE

(800) 755-PCBX
TEL: **(714) 668-1180**
FAX: (714) 668-0215
3730 S. Susan St., Ste. 250
Santa Ana, California 92704

PCBX SYSTEMS, INC.

CIRCLE NO. 179

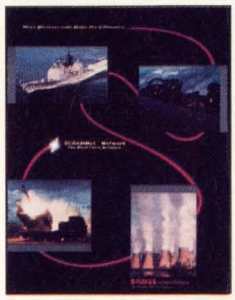

REAL-TIME NETWORK

The SCRAMNet™ Network combines the real-time speed of replicated shared memory with the flexibility of a fiber optic LAN to get microsecond response from multi-vendor computers. It offers distinct advantages in critical simulations. *Brochure highlights system's features.*

Systran Corp.
4126 Linden Avenue, Dayton, OH
45432-3068 USA.
Phone (513) 252-5601 or
1-800-252-5601.

CIRCLE NO. 180

PC Bus Data Acquisition
...over 50 exciting hardware & software products!

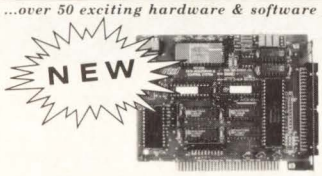

AD1200 *only* **$359**

- 12-bit 125 kHz A/D Conversion
- 16 Analog Input Channels
- 3 Timers and 16 DIO Lines
- Supports DMA and Pacer Clock
- AT Bus Version Available!

A/D, D/A, Digital I/O and Timer/Counters for PC XT and AT computers • 12-bit A/D rates to 200 kHz! • Eight boards for AMPRO form factor CPUs • 110 VAC relay, optoisolation, multiplexer and thermocouple boards • DSP and data acquisition software • Request your ***FREE 80 page catalog today!***

Real Time Devices, Inc.
State College, PA USA
Tel.:814/234-8087 • FAX:814/234-5218

CIRCLE NO. 181

Z80 STD USERS!

Cost Effective Upgrade
Clock Speeds to 10 MHz
1 Mbyte On-board Memory

Increase your system performance and reliability while reducing your costs by replacing three of the existing cards in your system with one **Superintegrated** Z80 Card from Zwick Systems.

A **Superintegrated** Card in your system protects your software investment, requiring only minor changes to your mature Z80 code. You can increase your processing performance by up to 300 percent in a matter of days!

Approximatly 35 % of each **Superintegrated** Card has been reserved for custom I/O functions including A/D, D/A, Industrial I/O, Parallel Ports, Serial Ports, Fax and Data Modems or almost any other form of I/O that you are currently using.

Call or Fax today for complete information on this exciting new line of **Superintegrated** Cards and upgrade your system the easy way!

ZWICK SYSTEMS INC.
Tel 613-726-1377, Fax 613-726-1902

CIRCLE NO. 182

EPROM PROGRAMMER

$750

- **Quick pulse** pgms. eight 1 Mbit EPROMs in 40 sec.
- Stand-alone or PC-driven
- RS-232, **parallel** in & out ports
- Binary, Intel hex, & Motorola S formats
- 100 user-definable macros
- Information, call **(916) 924-8037**
- 1 Megabit of DRAM
- Made in U.S.A.
- A9 Identifier
- 2 year warranty
- Single pgmr. $550

NEEDHAM'S ELECTRONICS
4539 Orange Grove Ave. • Sacramento, CA 95841
(M-F, 8-5 PST)
MasterCard VISA

CIRCLE NO. 183

SYSTEM SHOWCASE

FINALLY

One tool to satisfy all your firmware development needs

PROMICE is a universal system.
- Develops code for any microprocessor
- Complete, real time, source level debugging
- Host software for DOS, Unix, Mac, VMS
- Non intrusive on your target system
- Simply plugs into any ROM socket

PROMICE also supports Turbo Debugger, C_thru_ROM, FreeForm, GDB, and more.

PROMICE
...the affordable solution.

Grammar Engine Inc.

TEL: (614)471-1113
West: (415)750-0219
FAX: (614)475-6871

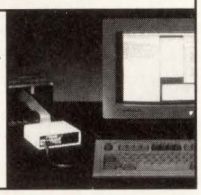

CIRCLE NO. 184

8051 68HC11 COP8

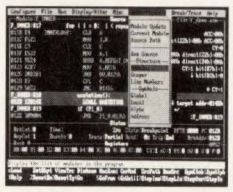

iceMASTER
Your Window To Emulation Productivity
- Easy to learn & use
- Windowed interface -- user configurable
- FAST! Download -- < 3 sec. typ. at 115KB
- Source Level debug

- **A 4K frame trace buffer** with advanced searching capabilities.
- **iceMASTER** connects easily to your PC, requires no disassembly, or expansion slots. Works on any PC (DOS or OS/2), MicroChannel or EISA. Even laptops!
- **iceMASTER is versatile:** iceMASTER-8051, iceMASTER-68HC11 and iceMASTER-COP8 support most family derivatives.
- **Rental and 10-day trials available.**
- 68HC11 A,D,E,F; 8XC528; 8XC552; 8XC515A and 8XC517A support.
- **Call today for free demo disk and ask about a free 8051 Macro Assembler!** **(800) 638-2423**

MetaLink Corporation

MetaLink Corporation P.O. Box 1329 Chandler, Az 85244-1329
Phone: (602) 926-0797 FAX: (602) 926-1198 TELEX: 4998050 MTLNK

HOT NEW PRODUCTS!

CIRCLE NO. 185

IND-286 SBC

AT Compatible DISKLESS SBC Includes DOS in ROM

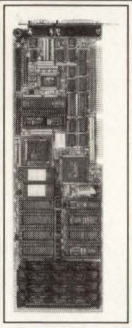

Complete 16MHz 80C286 Single Board Computer for embedded PC applications features a 4M-byte PROMDISK disk emulator with battery back-up and an MS-DOS 3.3 compatible disk operating system in ROM.

Features Include:
- ❑ 4M-byte DRAM
- ❑ XT Size Board
- ❑ Keyboard Port
- ❑ 80287 Socket
- ❑ 2 COM, 1 LPT
- ❑ WatchDog Timer
- ❑ IDE Disk Port
- ❑ 4M PROMDISK
- ❑ Floppy Port
- ❑ 100% PC/AT Compatible
- ❑ Optional Video Daughter Bd.

Other Products:
- IND-88 PC/XT Single Board Computers
- PROMDISK III & IV Disk Emulators
- EPROM/RAM Memory Board
- FLASHDISK Driver for Microsoft FFS
- FlexScan I & II Bar Code Decoders
- Custom PC Compatible Hardware & Software

MCSI micro computer specialists, inc.
2598-g fortune way vista, ca 92083
phone: 619/598-2177 fax: 619/598-2450

CIRCLE NO. 186

Reduced size 100 MHz Clock Oscillator

Connor-Winfield Corp. announces literature on a new line of High Speed 8-pin DIP oscillators. The A53 Series line of ACMOS logic clock oscillators covers the high frequency range of 50-120 MHz. Literature for a 14-pin version (AC53) with frequencies to 200 MHz is also available. The A53, 8-pin model will soon be available in a .48" square ceramic surface mount package. Contact Barney 111 for details. Additional specifications listed below:

Model : A53
Package : 8-pin DIP
Duty Cycle: 50/50 +/- 2%
R/F Times : 2-3 nS typical

Connor-Winfield Corporation
1865 Selmarten Road
Aurora, IL 60505
PH: (708) 851-4722 FAX: (708) 851-5040

CIRCLE NO. 187

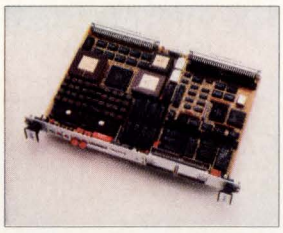

Intelligent VME I/O Controller w/Optional Mezzanine Modules

- 20/25MHz 68020 microprocessor
- 1 or 4Mbytes DRAM
- Mezzanine modules for analog, digital, I/O, memory, Ethernet, SCSI and more...

VPU-25 is a cost-effective design platform for adding I/O and networking capabilities to industrial/process control applications. Call SBE, Inc.:

1-800-347-2666

CIRCLE NO. 188

IT'S IN THE CARDS ...

More Leads, More Action, More Sales

Our readers are proven buyers of:

Computer Systems	93,256
System Boards	84,066
Integrated Circuits	77,922
Design/Development Equipment	85,312
Memory/Storage Equipment	72,237
Software	88,949
Terminals /Input-Output Equipment	81,894
Communications Equipment	70,342
Components	75,802
Test Equipment	65,035

Smart buyers depend on COMPUTER DESIGN DIRECT ACTION CARDS

Call Sue Shorrock at 800-225-0556 or 508-392-2185

COMPUTER DESIGN

CIRCLE NO. 188

EMPLOYMENT OPPORTUNITIES, OUTPLACEMENT AND POSITIONS WANTED

CAREER CONNECTIONS

Recruitment Advertising Manager
Tom Murphy, (213) 372-2744

Closing Dates for upcoming magazine issues

Issue	Closing
May	April 6
June	May 6

POSITION WANTED

DESIGN ENGINEER: 13 years experience with Fortune 50 companies in microprocessors, FPGA, RF, analog design, software development, CAD, high rel space design. Prefer position requiring some physical labor. Will consider sales, will relocate. Call Jerry at (206) 839-0747.

HELP WANTED

SENIOR SYSTEMS ANALYST Individuals must meet the qualifications listed below: Consult with clients in the preparation, design, and program based on client's requirements. Developments are on 4GL. DB2, CSP, CICS, COBOL-II on mainframe computers. This position is a full time opportunity that requires extensive domestic travel. Educational requirements: Minimum of a Bachelor's of Science or equivalent in Statistics/Mathematics. Experience requirements: Minimum of 4 years experience with the above mentioned languages, and client/consulting background required should apply. Salary $41,150.00. Must have proof of legal authority to work permanently in the U.S. Please forward your resume to: Illinois Department of Employment Security, 401 South State Street, 3-South, Chicago, IL 60605, Attn: Mr. James McBride Reference # V-IL 4761-T. No calls please. 2 copies of your resume required. An Employer paid ad.

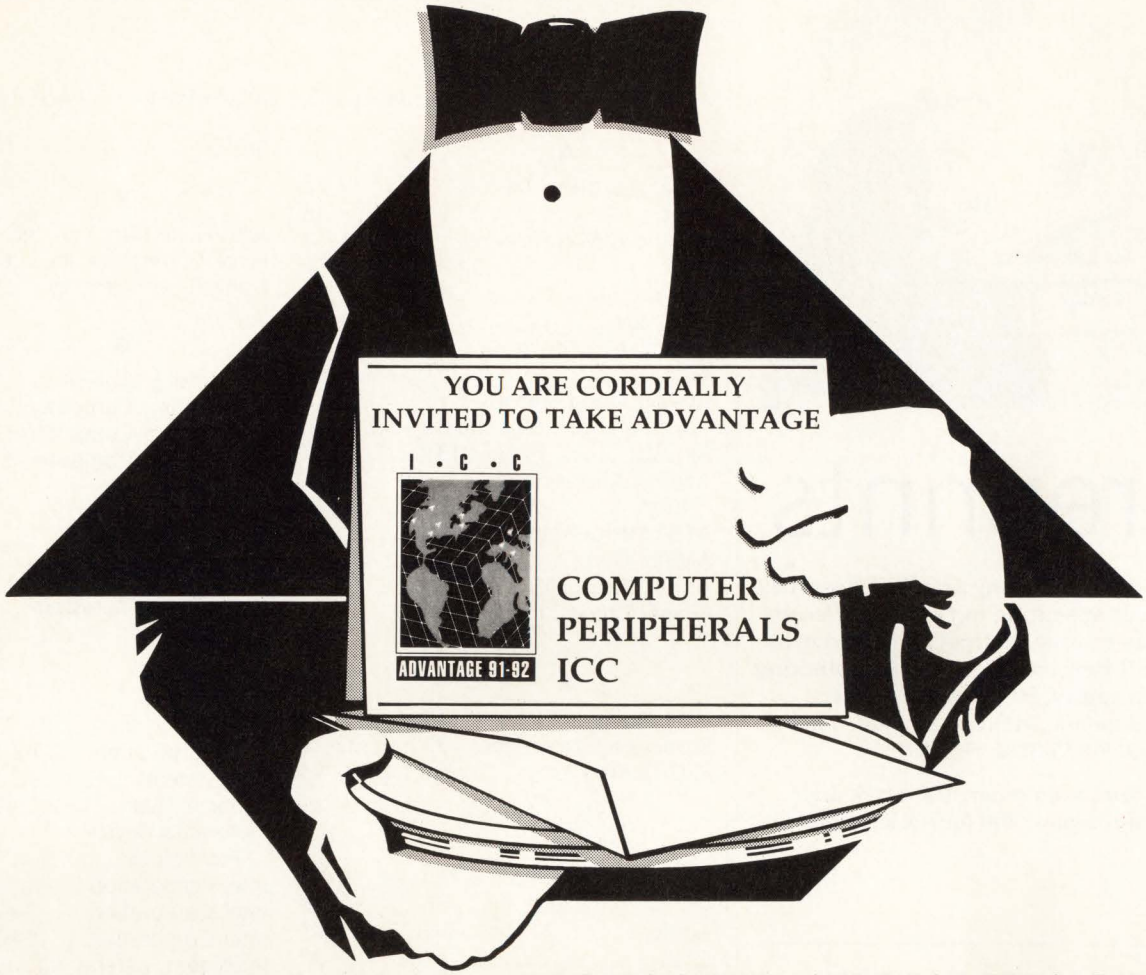

YOU'RE INVITED... TO TAKE ADVANTAGE

of the Leading Regional Peripherals Conference

featuring new technology presentations and product displays

Take your Advantage now.

Request your ***Complimentary Invitation*** and Computer Peripherals ICC Agenda

UNITED STATES

- ❑ Newton, MA - 9/5/91
- ❑ Tysons Corner, VA - 9/12/91
- ❑ Dallas, TX - 9/26/91
- ❑ Raleigh, NC - 10/3/91
- ❑ Portland, OR - 11/6/91
- ❑ Irvine, CA - 1/7/92
- ❑ Ft. Lauderdale, FL - 1/14/92
- ❑ San Jose, CA - 2/4/92
- ❑ Nashua, NH - 4/8/92
- ❑ Minneapolis, MN - 4/23/92
- ❑ Austin, TX - 5/14/92

EUROPE

- ❑ London, England - 9/12/91
- ❑ Copenhagen, Germany - 9/17/91
- ❑ Frankfurt, Germany - 9/24/91
- ❑ Barcelona, Spain - 10/2/91
- ❑ Amsterdam, Holland - 1/16/92
- ❑ München, Germany - 1/21/92
- ❑ Milano, Italy - 1/23/92
- ❑ Paris, France - 1/28/92

If you are involved in the purchase, design, or specification of peripheral products, you should be here!

Provide your conference selection by returning this section by Fax or mail.

Dataquest / ICC
3151 Airway Avenue, # C- 2
Costa Mesa, CA 92626

Place your business card here - Copy - Then fax this sheet back to us.

FAX: (714) 957-0903
TEL: (714) 957-0171

a division of Dataquest

a company of
The Dun & Bradstreet Corporation

Dataquest is a registered trademark of A. C. Nielson Company

CIRCLE NO. 75

reprints

Reprints of any article or advertisement appearing in *Computer Design* may be ordered from: June Bozarth, CSR Reprints, **PennWell Publishing Company**, P.O. Box 1260, Tulsa, Oklahoma, 74101, 1-800-331-4463 or 918-835-3161 Ext. 379.

Minimum order: 100 black and white copies, 500 four-color copies.

IT'S IN THE CARDS

More Leads, More Action, More Sales

Our readers are proven buyers of:

Computer Systems	93,256
Systems Boards	84,066
Integrated Circuits	77,922
Design/Development Equipment	85,312
Memory/Storage Equipment	72,237
Software	88,949
Terminal/Input-Output Equipment	81,894
Communications Equipment	70,342
Components	75,802
Test Equipment	65,035

Smart buyers depend on COMPUTER DESIGN DIRECT ACTION CARDS

Call Sue Shorrock: at 800-225-0556 or 508-392-2185

COMPUTER DESIGN

CIRCLE NO. 83

ADVERTISERS INDEX

ADVERTISER	PAGE NO.	CIRCLE NO.
A		
Advanced Micro Devices, Inc.	20-21	18
Advanced Micro Devices, Inc.	94-95	57
Alacron, Inc.	10	10
Alsys, Inc.	15	8
American Eltec, Inc.	85	53
AMP, Inc.	120-121	64
Analog/Mixed Signal Conference	86-87	54
AP Labs	131	67
Applied Microsystems Corp.	33	21
*AT&T Microelectronics	48-49	33
AMTEL Corp.	41	27
Arrow Electronics	149	65
Avanti Circuits, Inc.	154	168
B		
W.H. Brady	16	13
Brooktree Corporation	35	22-23
BSO/Tasking	61	43
C		
Cermetek Microelect.	110	87
Chrislin Industries, Inc.	36	24
Ciprico	CV2	1
CMC	S7	90
†Conference Proceedings	48-49	34
Conference Proceedings	142	55
Connor-Winfield Corp.	155	177
Connor-Winfield Corp.	156	187
CSPI	45	31
Custom Integrated Circuits Conference	137	82
Cypress Semiconductor	CV4	–
†Computer Design	112	63
Computer Design	128	78
D		
Design Automation Conference	138	77
Duracell	83	51
Dy-4 Systems, Inc.	109	61
Dynatem, Inc.	146	72
E		
EG&G Wakefield Engineering	42	28
Electro International	152	81
Electronic Solutions	1	2
Embedded Support Tools	154	173
Emulex Corporation	18	15
Eyring	133	69
F		
John Fluke Mfg. Co., Inc.	17	14
Force Computers, Inc.	6-7	5
Franklin Software, Inc.	154	166
G		
Grammar Engine, Inc.	156	184
Greenspring Computers	143	76
Greenspring Computers	139	73
Greenspring Computers	141	74
H		
Hewlett Packard	8-9	6
*Hitachi America, Ltd.	112	62
Huntsville Microsystems	4	4
Hyperlogic	131	66
I		
IBI Systems, Inc.	154	170
IBM Corporation	62-63	–
Ikos Systems	31	–
Imagine That	155	178
Integrated Device Technology	13	7
Intel Corporation	159	84
Intel Corporation	24-25	–
Intel Corporation	54-55	38
Intelligent Systems Int'l	110	88
Interactive Circuits & Systems	160	85
Intergon	52	36
Invitational Computer Conference	157	75
Ironics, Inc.	51	35
K		
KADAK	16	12
L		
Lattice Semiconductor Corp	93	56
M		
Macrolink, Inc.	91	71
Matrix Corporation	105	59
MCSI	156	186
MEDO USA	154	174
Mercury Computer Systems	11	11
Metalink Corporation	156	185
Micro/Sys	155	175
Microboards, Inc.	10	9
Micron Technology	37	25
Microware Systems, Inc.	43	29
MIZAR	111	89

ADVERTISERS INDEX

ADVERTISER	PAGE NO.	CIRCLE NO.
Motorola TSD Group	81	50
N		
National Semiconductor	97	–
Needham's Electronics	155	183
Newbridge Microsystems, Inc.	74	48
Nohau Corporation	154	167
O		
Oasys, Inc.	59	41
Omnibyte Corporation	79	120
OrCad Systems	154	171
P		
PCBX Systems, Inc.	155	179
Performance Semiconductor	67	45
Peritek	58	40
PRC Inc.	84	52
Q		
Quantum Software Systems, Ltd.	57	39
R		
RadiSys Corporation	23	19
RadiSys Corporation	73	47
Radstone Technology	106	60
Real Time Devices	155	181
Robotrol Corporation	154	169
S		
SBE, Inc.	156	188
SBE, Inc.	CV3	86
Software Development Systems	103	58
Sophia Systems	101	–
Spectrum Signal Processing	71	46
Sun Microsystems, Inc.	S3	91
Sunrise Electronics	155	176
Synergy Microsystems, Inc.	2	3
Systran Corporation	155	180
T		
Tadpole Technology	26	20
Tadpole Technology	47	32
Tech-Source, Inc.	S9	92
Teknor Microsystems,	53	37
Tektronix, Inc.	64	44
Themis Computer	77	49
Toronto MicroElectronics, Inc.	18	16
Toshiba America Elect. Components	38-39	26
V		
VenturCom	60	42
V.H.D.L. Consulting Group	150	79
Vigra, Inc.	132	68
VME Microsystems Int'l Corp.	135	70
VMETRO, Inc.	151	80
W		
Wind River Systems	19	17
X		
Xycom	44	30
Z		
Z-World Engineering	154	172
Zwick Systems, Inc.	155	182

Computer Design's INSTANT DATA ACCESS ADVERTISERS

To receive data sheets instantly from these advertisers via FAX, **dial (617) 494-8338** and when prompted, enter the 4 digit document number shown here.

ADVERTISER	Page NO.	IDA NO
ATMEL	41	1045
Brooktree	35	1026
Brooktree	35	1029
CMC	S7	1025
Computer Design	128	1047
Dynatem	138	1032
Hitachi	112	1034
Ironics	51	1046
KADAK	16	1036
Macrolink	137	1037
Mercury Computer	15	1020
SBE	CV3	1040
Tech Source	S9	1043
Teknor	53	1028

* Domestic only

† International only

The Advertisers Index is published as a service. The publisher does not assume liability for errors or omissions.

intel®

Access in Europe

Microelectronics Components Boards Systems

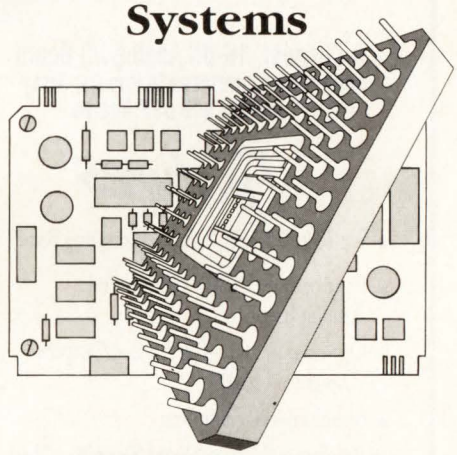

EUROPEAN SALES OFFICES

FINLAND

Intel Finland OY
Ruosilantie 2
00390 Helsinki
Tel.: (358) 0544 644
Fax: (358) 0544 030
Telex: 123 332

FRANCE

Intel Corporation S.A.R.L.
1, rue Edison - BP 303
78054 St.Quentin-en-Yvelines Cedex
Tel.: (33) (1) 3057 7000
Fax: (33) (1) 3064 6032
Telex: 699 016

UNITED KINGDOM

Intel Corporation (U.K.) Ltd.
Pipers Way
Swindon Wilts SN3 1RJ
Tel.: (44) (0793) 69 60 00
Fax: (44) (0793) 64 14 40
Telex: 444 447/8

NETHERLANDS

Intel Semiconductor B.V.
Postbus 84130
3009 CC Rotterdam
Tel.: (31) 10 407 11 11
Fax: (31) 10 455 46 88
Telex: 222 83

ISRAEL

Intel Semiconductor Ltd.
Atidim Industrial Park-
Neve Sharet
P.O. Box 43202
Tel-Aviv 61430
Tel.: (972) 03 49 80 80
Fax: (972) 03 49 18 70
Telex: 371215

ITALY

Intel Corporation Italia S.p.A.
Milanofiori Palazzo E
20090 Assago
Tel.: (39) (02) 89 20 09 50
Fax: (39) (02) 349 84 64
Telex: 341 286

SPAIN

Intel Iberia S.A.
Zubaran, 28
28010 Madrid
Tel.: (34) 1 308 25 52
Fax: (34) 1 410 75 70
Telex: 468 80

SWEDEN

Intel Sweden A.B.
Dalvägen 24
17136 Solna
Tel.: (46) 8 734 01 00
Fax: (46) 8 27 80 85
Telex: 122 61

GERMANY

Intel GmbH
Dornacher Strasse 1
8016 Feldkirchen
Tel.: (49) 089 90 992-0
Fax: (49) 089 90 439 48
Telex: 5-23177

Intel GmbH
Wandersmannstr. 64
6200 Wiesbaden-Erbenheim
Tel.: (49) 0611 97 70 3-0
Fax: (49) 0611 9 77 03 30

Intel GmbH
Zettachring 10A
7000 Stuttgart 80
Tel.: (49) 0711 72 87-280
Fax: (49) 0711 72 80 137
Telex: 7-254 826

DORN ADVERTISING

CIRCLE NO. 84

Optimum Data Acquisition Solution

ICS-140

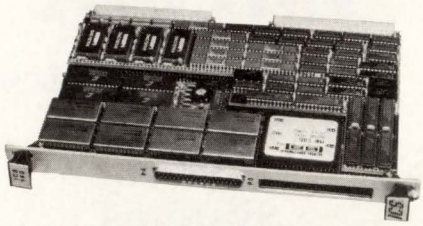

8-Channel, 16-Bit Analog I/O Board
With Direct Interfaces to Popular Array Processor and DSP Boards

Features

- Up to 100 KHz/Channel Sampling Rate
- Four or eight 16-Bit Differential Analog Inputs
- Optional Four 16-Bit Analog Outputs (Daughter Card)
- Simultaneous Sampling
- Programmable On-Board Sampling Clock
- Convenient VMEbus Access of ADC or DAC Data
- Direct Serial and Parallel Interfaces for Array Processor and DSP Boards

Applications

- Sonar Signal Processing
- Digital Audio
- Vibration Analysis
- Seismic Signal Processing
- Large BW Track And Hold Amplifiers Permit Undersampling of Signals (Band Limited) of Frequency up to 500 KHz
- Precision Applications Requiring Matched Channels, Negligible Sampling Aperture Jitter and Low Inter-channel Crosstalk

Ruggedized Version Also Available

CALL US TOLL FREE
1-800-267-9794
Tel: (613)-749-9241 Fax: (613)-749-9461

INTERACTIVE CIRCUITS AND SYSTEMS LTD.

CIRCLE NO. 85

COMPUTER DESIGN
Technology and Design Directions

Executive Office: One Technology Park Dr, PO Box 990, Westford, MA 01886
(508) 692-0700 Fax: (508) 692-7780 Telex: 883436
Publisher: David L. Allen (508) 392-2111
Associate Publisher/Editor-in-Chief: John C. Miklosz (508) 392-2114
Associate Publisher/National Sales Manager: Tim L. Tobeck (508) 392-2116
Administrative Assistant: Peg Alexander (508) 392-2112
Marketing Specialist: Betsy Anderson (508) 392-2209
Ad Traffic Manager: Kelly Rice (508) 392-2198
Recruitment Advertising: Tom Murphy (213) 372-2744
Postcards/System Showcase: Sue Shorrock (508) 392-2185
Circulation Director: Bob Dromgoole (918) 832-9213
Circulation Manager Paul Westervelt (918) 832-9287
List Rental: Deanna Rebro (918) 831-9551
Reprints: June Bozarth (918) 835-3161

SALES OFFICES

Associate Publisher
Tim L. Tobeck

National Accounts Managers

Eastern/Midwestern Region
Tim Pritchard
One Technology Park Dr
PO Box 990
Westford, MA 01886
Tel: (508) 392-2217
Fax: (508) 692-7780
Telex: 883436

Southwestern/ Northwestern Region
Eric Jeter
19627 I-45 N, Suite 220
Spring, TX 77388
Tel: (713) 353-0309
Fax: (713) 288-8350

Northern California
Diane Palermo
910 Campisi Way
Suite 1B
Campbell, CA, 95008
Tel: (408) 371-7551
Fax: (408) 371-6329

Strategic Accounts Representatives

Eastern U.S.
Sue Nawoichik
One Technology Park Dr
PO Box 990
Westford, MA 01886
Tel: (508) 392-2118
Fax: (508) 692-7780

Western U.S.
David Singer
One Technology Park Dr
PO Box 990
Westford, MA 01886
Tel: (508) 392-2109
Fax: (508) 692-7780

U.K./Scandinavia
David Round
Westmead House
123 Westmead Road
Sutton
Surrey SM1 4JH
Tel: (081) 770-1100
Fax: (081) 770-9779

France/Belgium/ Southern Switzerland/ Spain/Netherlands
Daniel R. Bernard
10, rue Michelet - B.P. 279
78502 Sartrouville Cedex, France
Tel: 01-39-14-67-80
Fax: 01-39-14-70-14

Germany/Austria/Northern Switzerland/Eastern Europe
Johann Bylek
Verlagsbuero Johann Bylek
Stockaeckerring 63
D-8011 Kirchheim/Muenchen
Federal Republic of Germany
Tel: 89-903-88-06
Fax: 89-904-35-26

Italy
Luigi Rancati
Rancati Advertising
Milano San Felice Torre 5
20090 Segrate, Italy
Tel: 02-70300088
Telex: 328601 RANCAD I
Fax: 02-7532354

Japan
Toshio Egusa
PubliNetwork
C407, 2-22-6
Tsukuda Chuo-ku
Tokyo 104, Japan
Tel: 03-3536-5404
Fax: 03-3536-5490

Southeast Asia Taiwan/Hong Kong/Korea/ China/Thailand/Philippines/ India/Pakistan/Singapore/ Malaysia/Indonesia/Australia/ New Zealand
Tom Gorman
CCI Asia-Pacific, Ltd
Suite 905, Guardian House
32 Oi Kwan Rd, Happy Valley
Hong Kong
Tel: 833-2181

COMPUTER DESIGN (ISSN-0010-4566) is published 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 Dr, PO Box 990, Westford, MA 01886. Subscription Prices: Free to design and development qualified engineers and engineering managers in the U.S. and Canada. Paid to all others. Qualified engineers and engineering managers outside the U.S. and Canada - air shipped - $85 one year. For non-qualified recipients in the U.S. - $88 one year; in all other countries - air shipped - $154 one year. Call (918) 832-9263 for subscription information. Microfilm copies of COMPUTER DESIGN may be purchased from University Microfilms, a Xerox Company, 300 N Zeeb Rd, Ann Arbor, MI 48106. POSTMASTER: Send change of address form to COMPUTER DESIGN, Circulation Department, PO Box 3466, Tulsa, OK 74101.

©1992 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; Joe T. Bessette, Senior Vice-President; John Maney, Vice-President/Finance; Steve Zimmerman, Vice-President/Corporate Services.

No Matter What the Application,

SBE Fits.

Matching your high-speed data communications requirements with a quality supplier has never been easier. Whether you're a manufacturer of mini/superminicomputers, workstations or high-performance data communications products, *only* SBE provides a perfect fit.

Only SBE offers a complete line of intelligent high-performance communications controllers for all major interface technologies: FDDI, Token Ring, Ethernet and High Speed Serial. Only SBE adds premium features, without a premium cost, for the best price/performance in the industry.

Add integrated hardware/software solutions; availability in VMEbus, Multibus and SBus; plus legendary development assistance and continuing product support.

Discover how SBE's intelligent high-performance controllers can meet your LAN and WAN interface requirements. Turn to SBE today.

For fast action, call: 1-800-347-COMM
Germany: 0130-810588
United Kingdom: 0800-378-234

SBE, Inc., 2400 Bisso Lane, Concord, CA 94520

CIRCLE NO. 86

INSTANT DATA ACCESS (IDA) DIAL (617) 494-8338 DOCUMENT NO. 1040

SBE
Communications & Real-time Solutions

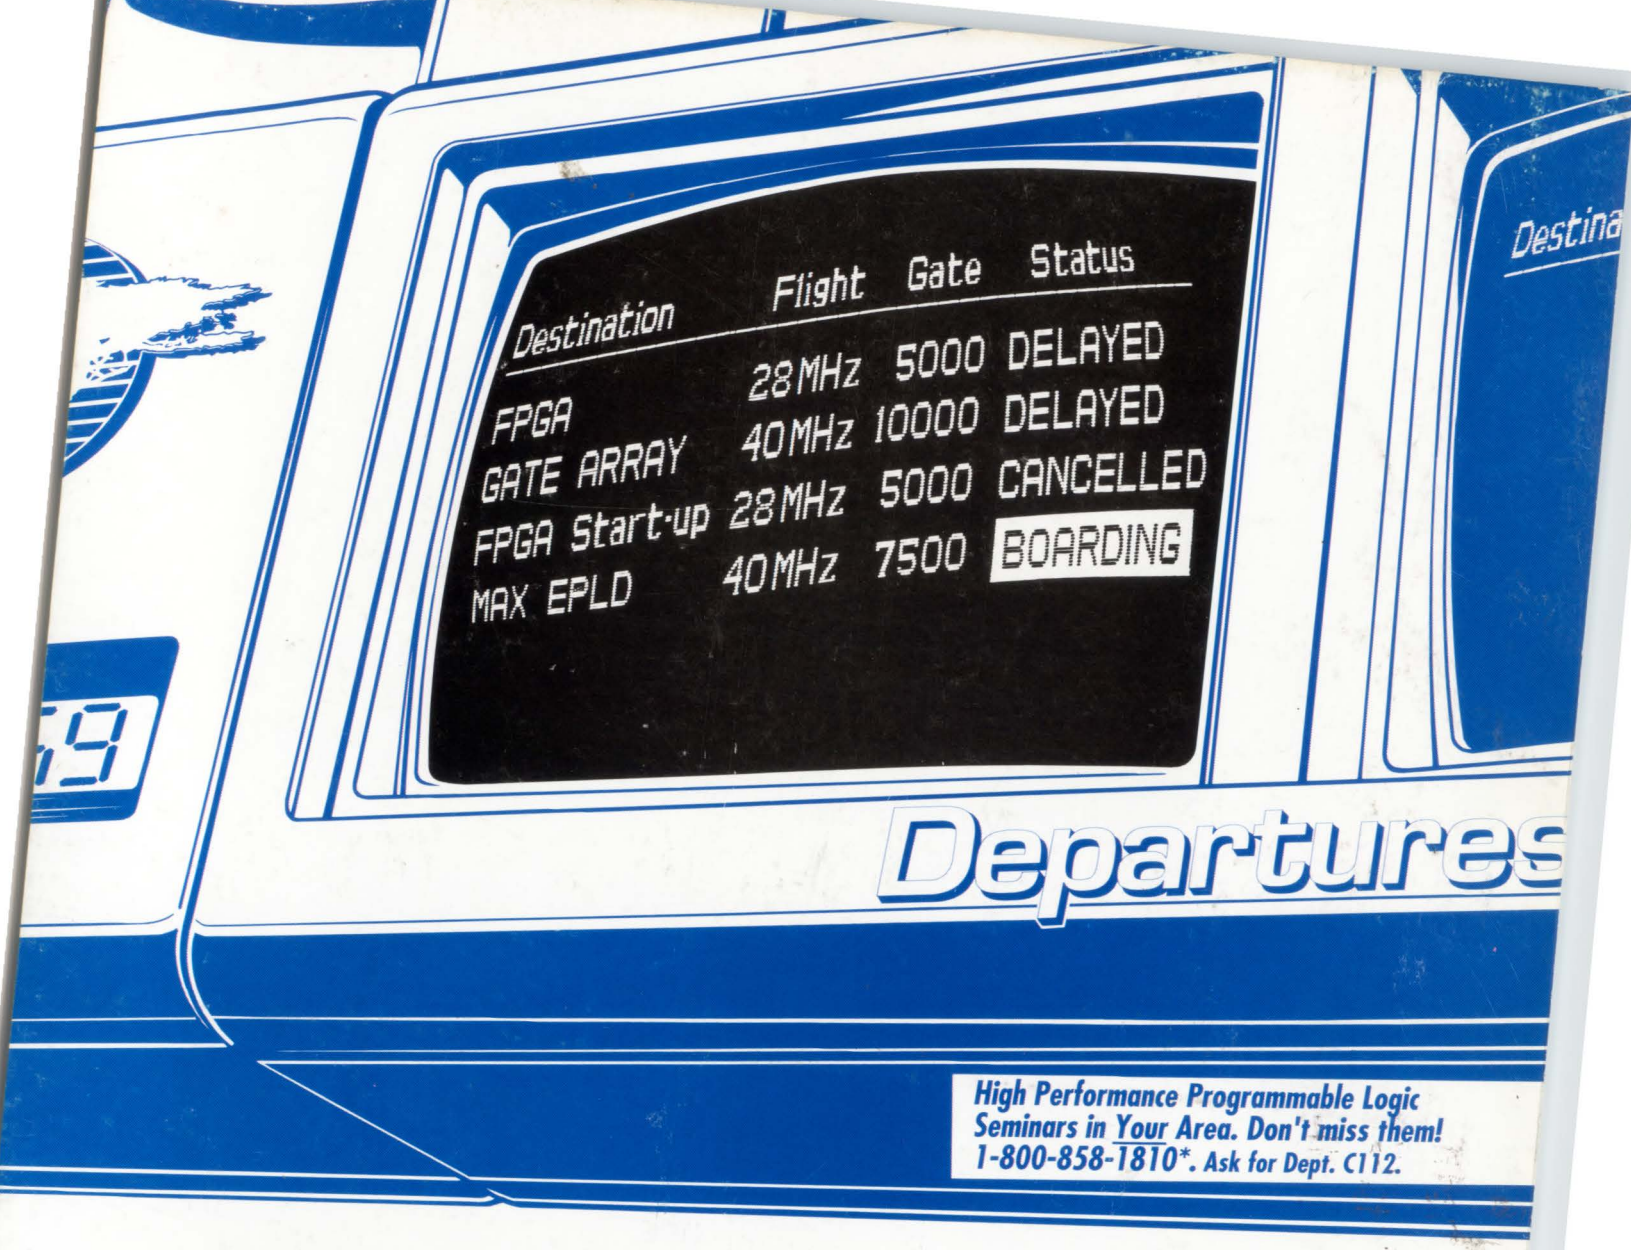

High Performance Programmable Logic Seminars in Your Area. Don't miss them! 1-800-858-1810*. Ask for Dept. C112.

NOW TAKING OFF–MAX EPLD.

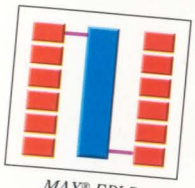

MAX® EPLD Architecture allows full utilization of 7,500 gates.

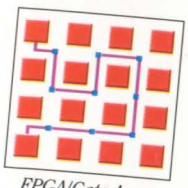

FPGA/Gate Array Architecture: prepare for lots of delays (speed depends on routing).

The 40 MHz PLD with Gate Array Capacity

If you want PLD design convenience with the density of FPGAs or gate arrays, Cypress has the best way to get you there without delay: our new MAX CY7C341 EPLD.

MAX 341's exclusive Logic Array Block (LAB) architecture always affords you the shortest interconnect path. Data glides between the 12 LABs with a singular, predictable interconnect delay. Count on full, 7,500 gate utilization, with no drop-off from MAX's 40 MHz performance. And MAX 341's nonvolatility and built-in security bit give you extra safety and stability.

Since MAX 341 is a field-programmable EPLD, you cruise above the NREs that can cloud gate array solutions. You save even more in design-in time with MAX's PC-based development kit, which enables you to have a programmed solution in hours, not months.

Catch the entire MAX family of 28-, 44-, 68- and 84-pin devices in both (erasable) windowed ceramic or plastic packages. Call the Cypress hotline today for your free MAX information, including brochure and Data Book. Our Field Applications Engineers are ready to show you how MAX 341 can make your designs fly.

FREE MAX INFO HOTLINE:
1-800-858-1810*
Ask for Dept C111.

CYPRESS SEMICONDUCTOR

*In Europe, fax your request to the above dept. at (32) 2-652-1504 or call (32) 2-652-0270. In Asia, fax to the above dept. at 1 (415) 961-4201. © 1992 Cypress Semiconductor, 3901 North First Street, San Jose, CA 95134. Phone 1 (408) 943-2600, TELEX: 821032 CYPRESS SNJ UD, TWX: 910-997-0753. MAX is a registered trademark of Altera Corporation.