A PENNWELL PUBLICATION FEBRUARY 1, 1991

COMPUTER DESIGN

Technology and Design Directions®

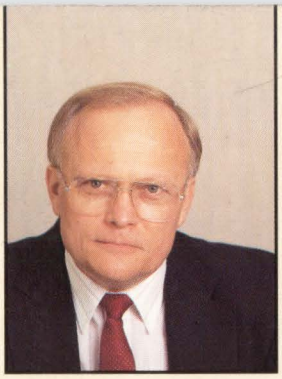

Pete Yeatman on board architecture

FOR ELECTRONIC ENGINEERS & ENGINEERING MANAGERS

Will performance win over sophistication in workstation buses?

FDDI chips struggle toward the desktop

Tools bring vital information to early stages of system design

Unparalleled, parallel performance.

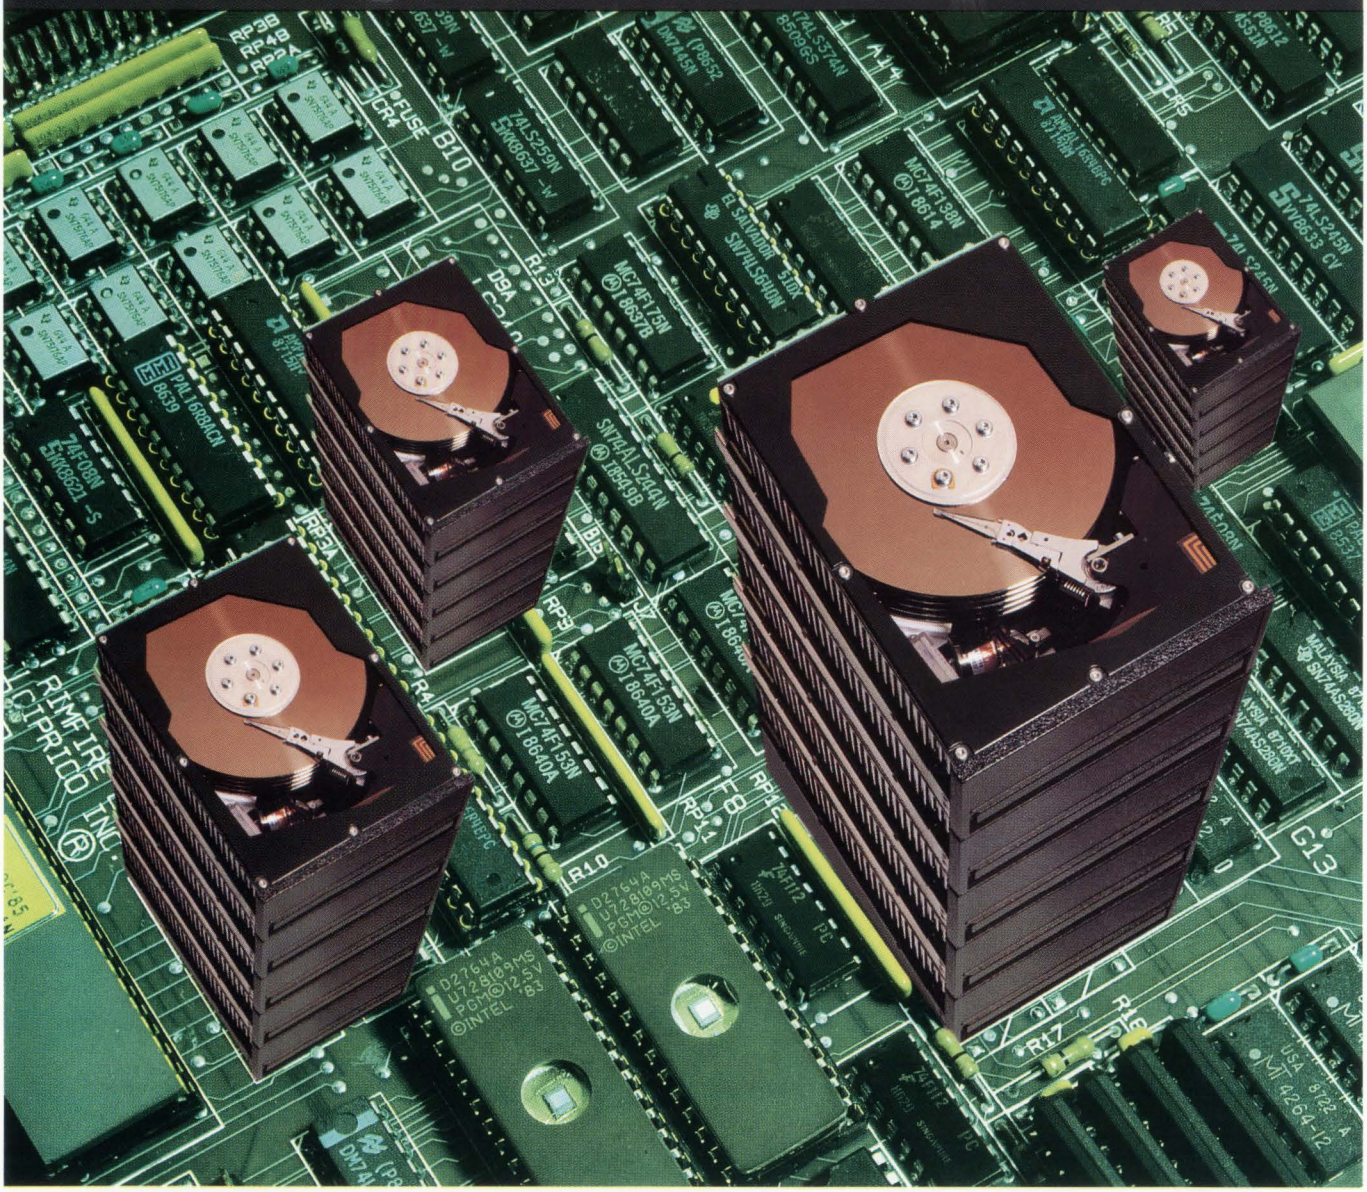

SCSI-2 is at your command with Rimfire® 6600 Series Parallel Disk Array Controllers.

Discover unmatched performance and data integrity for mass storage applications—the Rimfire 6600 series of Parallel Disk Array (PDA) Controllers.

The first array controller offering SCSI-2 as its host interface, as well as the first offered as a board level product, the Rimfire 6600 Series supports four data drives plus a single redundant drive. Because all disk data transfers occur in parallel, the array appears to the host as a single SCSI drive. *For you, that means excellent performance. Transfer rates and capacity four times that of an individual drive. And exceptional data availability.*

In addition, the flexibility of the PDA controller allows the OEM to select drives from many vendors. It's compatible with SCSI, operates with off-the-shelf host adapters and is priced significantly lower than other solutions of equivalent capabilities.

For detailed information on the Rimfire 6600 Series, or other SCSI-2 compatible products, call Ciprico, the industry leader in technology, technical support, customer service and respondability at 1-800-SCSI-NOW (1-800-727-4669). European customers call our United Kingdom office, (0703) 330 403.

CIPRICO
2955 Xenium Lane
Plymouth, MN 55441

CIPRICO LISTENS. AND RESPONDS.

Ciprico and Rimfire are registered trademarks of Ciprico, Inc. Disk drive courtesy of Seagate Technology.

CIRCLE NO. 1

New

10-Layer VME Backplane!

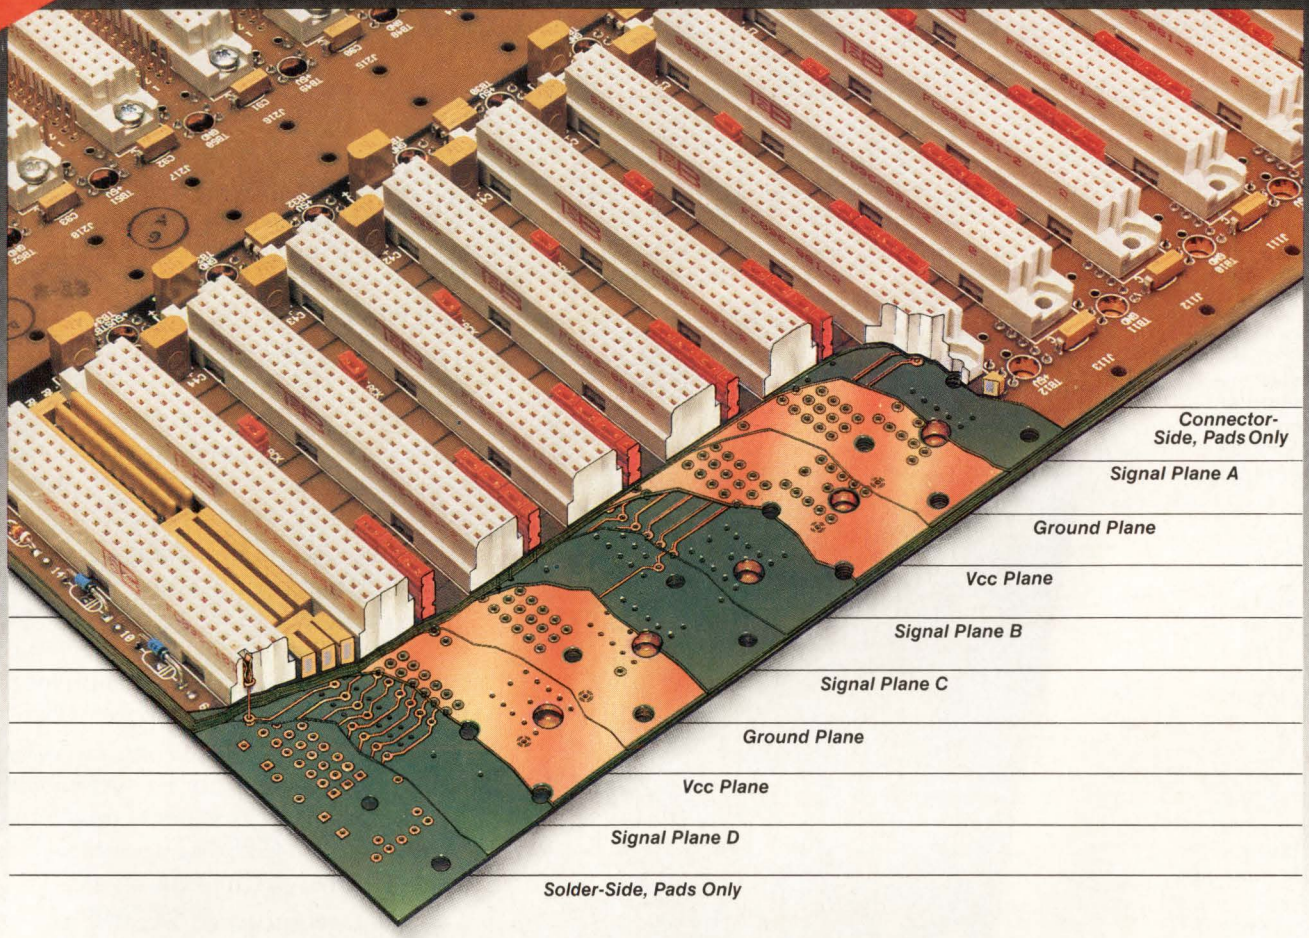

There's Logic To Every Layer!

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

At Last, a *Truly* Industrial 386 PC AT System for the OEM

Easy to mount Eurocard hardware with optional power supply. Can be rack mounted or bolted into any embedded system.

Three side mechanical bond into the card cage, yet easy to access for simplified maintenance.

Two boards make a complete 386 computer with your choice of EGA or VGA display.

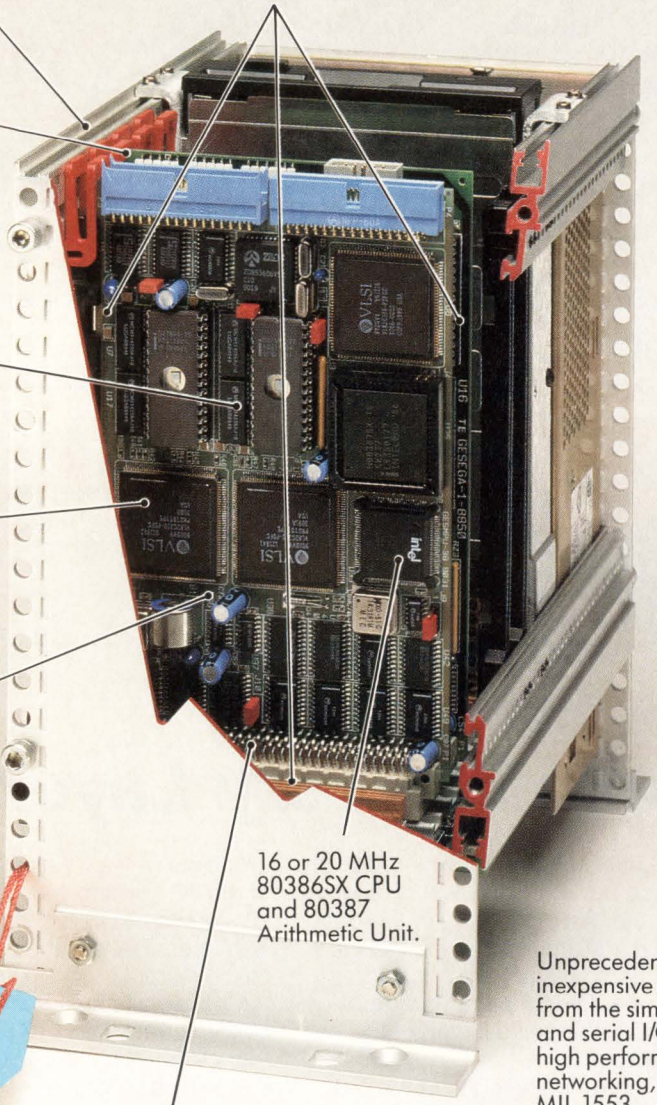

1, 2, 4 or 8 Mbytes of on board. Full support for LIM EMS and BIOS shadowing.

On board support for COM1, COM2, LPT, keyboard, clock, floppy disk, and IDE hard disk.

Compact single height Eurocard board format (100 x 160 mm) that doesn't flex or bend and withstands the toughest shock and vibration.

Small price to help you compete.

16 or 20 MHz 80386SX CPU and 80387 Arithmetic Unit.

Full 32-bit expansion bus for easy system expansion and customization with up to 20 slots.

Introducing the first 386 PC AT system originally designed to be embedded into your machinery or instrumentation.

With the GESPAC AT system, "ruggedized PC" no longer describes a desktop PC in a stronger black box. Instead it is an architecture made to last by design, with such features as small and robust board format, and ultra-reliable pin-in-socket connector.

What's more, the GESPAC System is built for performance with Intel's original 16 or 20 MHz 80386SX processor and 80387SX Arithmetic Unit. In fact, this system is the most compact implementation of the 386 in the market today.

The powerful G-64/96 bus interface lets you customize your AT system with any of the 200-plus system components available from GESPAC, to match virtually any system requirement.

Take advantage of GESPAC's 12 years of experience in serving the OEM marketplace. Call today to receive your free data sheet and GESPAC's catalog of board level microsystem products. Ask about GESPAC's exclusive Guaranty of OEM Long Term Delivery.

Unprecedented selection of inexpensive I/O functions, from the simplest parallel and serial I/O, to high performance networking, MIL 1553, motion control, and computer vision.

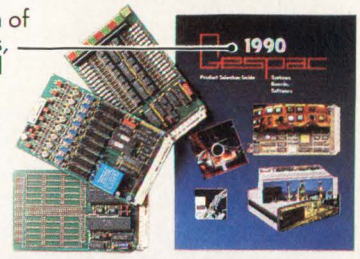

Call Toll Free 1-800-4-GESPAC or call (602) 962-5559.

Gespac

USA - CANADA
50 West Hoover Ave.
Mesa, Arizona
85210 USA
Tel. (602) 962-5559
Fax. (602) 962-5750

FRANCE
Z.I. les Playes
83500 La Seyne
France
Tel. 94 30 34 34
Fax. 94 87 35 52

INTERNATIONAL
18, Chemin des Aulx
CH-1228 Geneva
Switzerland
Tel. (022) 794 3400
Fax. (022) 794 6477

CIRCLE NO. 3

VOL. 30, NO. 3 FEBRUARY 1, 1991

COMPUTER DESIGN

Technology and Design Directions®

FOR ELECTRONIC ENGINEERS & ENGINEERING MANAGERS

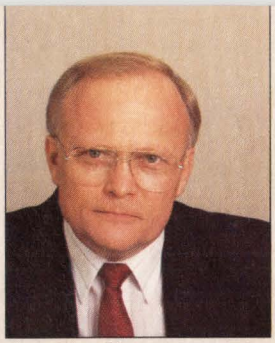

Pete Yeatman on board architecture
Page 23

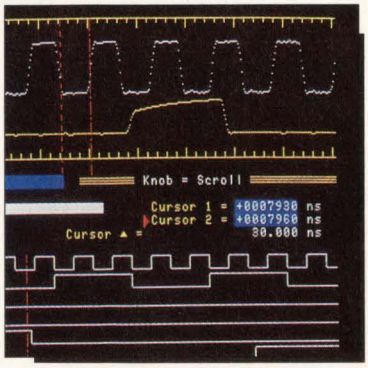

Page 54

TECHNOLOGY UPDATES

Integrated Circuits

MIPS rethinks RISC with superpipelining 28

Software

Japan's Tron initiative may be opportunity for U.S. manufacturers..... 35

Computers and Subsystems

National unveils Futurebus+ silicon 40

Semiconductor maker enters VMEbus board market..... 46

Design and Development Tools

Combined scope/logic analyzers detect hard-to-find analog problems.... 54

Windows 3.0 extends PC-based PLD design limits 56

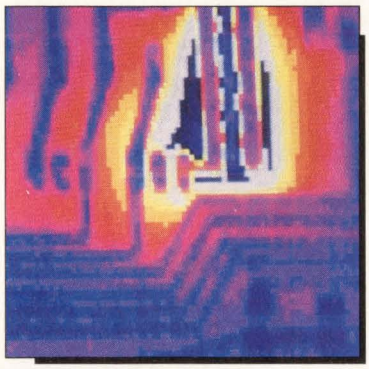

Page 71

TECHNOLOGY AND DESIGN FEATURES

FDDI chips struggle toward the desktop

The success of FDDI networking may stall at the high end unless workstations and PCs are included 61

Tools bring vital information to early stages of systems design

With shrinking design times and growing systems complexity, designers rely on tools that predict component behavior before prototyping 71

COVER STORY

Will performance win over sophistication in workstation buses?

A new class of standard bus has emerged to support the I/O requirements of compact desktop workstations. The continued growth of processor performance is driving these buses to higher speeds and minimal overhead at some cost in flexibility 78

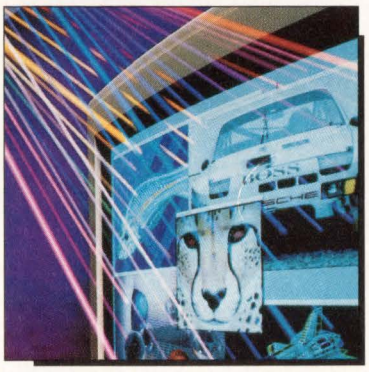

Page 93

NEW PRODUCT HIGHLIGHTS

PRODUCT FOCUS

High-speed DACs target waveform synthesis and video..... 93

Integrated Circuits

New chip offers nonvolatile analog memory 100

Multimedia chip compresses video and audio in synch..... 102

Xilinx backs enhanced FPGA family with improved development software..... 104

First BiCMOS PLDs are the fastest standard 22V10s yet..... 106

Software

Ada development environment is all Ada 107

Design and Development Tools

Analog layout editor links schematics to physical design 109

Design tool solves timing problems 110

DEPARTMENTS

News Briefs 8
Editorial 12
Calendar 16
Advertisers Index 119

FDDI.

From deskwork to network.

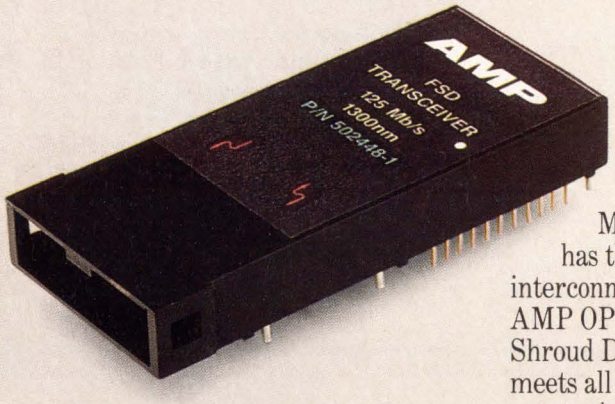

Good news for networks!

The X3T9.5 Task Group, under the procedures of ANSI Accredited Standards Committee X3, has reaffirmed approval of the Media Interface Connector (MIC) for the proposed FDDI (Fiber Distributed Data Interface) Physical Layer Medium Dependent (PMD) document.

More good news! AMP has the complete fiber optic interconnection system—the AMP OPTIMATE Fixed Shroud Duplex System—that meets all FDDI PMD requirements. And includes all the physical components you need to make your fiber optic network a reality.

Of special note: the transceiver is capable of operating at data rates up to 125 Mb/s. Available in standard or raised (+5v) ECL logic, it gives you a compact, board-mount data link in an industry-standard 22-pin package. Reliable duplex mating and electro-optic conversion are now easier than ever.

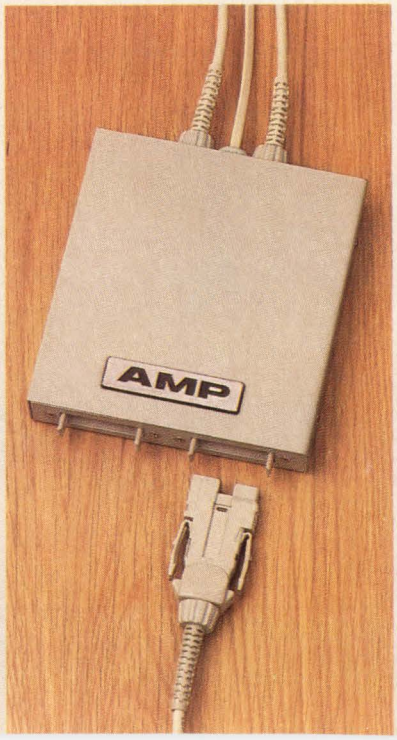

AMP and OPTIMATE are trademarks of AMP Incorporated.

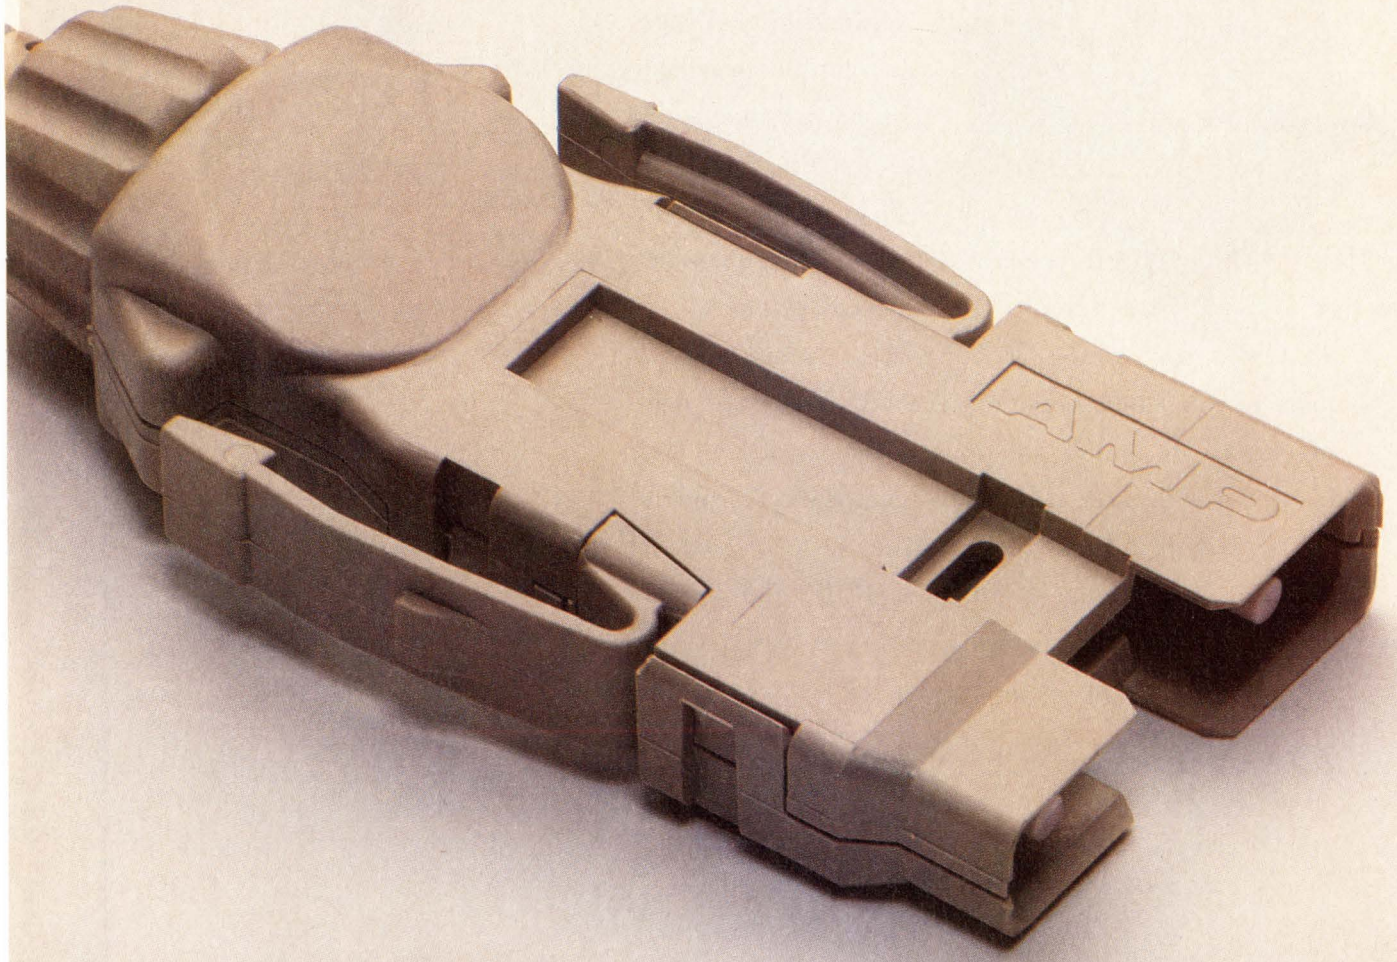

All system components, in fact, are easy to install and reconfigure. Our field termination kit makes short work of attaching duplex connectors to fiber cable. And because all interconnections use a floating interface, you get consistent, low-loss mating (0.6 dB typical) throughout.

You can also order complete, custom-built cable assemblies from us. Either way, you'll have the assured compatibility that comes from dealing with only one supplier for all your FDDI interconnection components. A supplier whose capability in fiber optic technology is everything you'd expect from the world's largest connector company.

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

CIRCLE NO. 4

THIS IS AMP TODAY.

AMP

Networking challenges can often be formidable, nearly insurmountable. But by joining forces and working together, we can succeed in ways that didn't seem possible. That's *"The Power Of Partnership."*

At CMC, The Power Of Partnership is our commitment to you – the system designer, integrator and manufacturer – to finding solutions to your most ambitious local and wide area networking enterprises. To providing you with a broad and flexible base of protocol software and intelligent Ethernet, FDDI and Primary Rate ISDN adapters, each linked to the expertise and support of a highly skilled technical team.

The Power Of Partnership™

For over a decade we have been providing high performance networking products to the world's leading computer companies. As a part of Rockwell International, our solutions-based approach to your most demanding networking problems means you can place your trust and confidence in our products and our people. We're here to help you succeed.

To scale new heights in connectivity call for *The Power Of Partnership.* Call 1-800-CMC-8023.

CMC *A Rockwell International Company*

125 Cremona Drive, Santa Barbara, CA 93117 • Phone: 1-805-562-3173, 1-800-CMC-8023
FAX: 1-805-968-6478 • In Europe: 44-81-577-2800

CIRCLE NO. 5

The Power Of Partnership and CMC are trademarks of CMC.

NEWS BRIEFS NEWS BRIEFS NEWS BRIEFS NEWS BRIEF

Force inks deal to manufacture Sun boards

Force Computer (Campbell, CA) has finalized a contract with Sun Microsystems (Mountain View, CA) to manufacture Sun's 1E VME board and other Sparc-based technology. Meanwhile, a second generation, the 2E, is apparently already under development. According to Peter Palm, board-level marketing engineer at Sun, the 2E version will boast more than twice the performance of the 1E, promising 28 Mips and 4.5 MFlops.

Look for whatever Sparc 2E technology is surfacing to appear on both VME and Futurebus+ almost simultaneously. According to Scott McNealy, Sun's president and CEO, "Sun anticipates that the combined marketing and sales impact of the two companies will make Sparc the leading RISC architecture on VMEbus by the end of 1991." The move certainly emphasizes the push Sun is putting behind its Sparc architecture.

But the Sun/Force team may have some problems making Sparc the dominant RISC architecture on VME. Motorola is heavily pushing its 88K board, and it's rumored a few other VME companies (besides Force, which already offers an 88K board) will soon have offerings. The major player in the RISC/VME market, however, might just end up to be MIPS Computer Systems (Sunnyvale, CA). At least two companies, Lockheed Systems (Nashua, NH) and Omnibyte (West Chicago, IL), already offer MIPS R-3000 VME boards. Lockheed initially offered only a mil version, but at the Buscon/90-East show it entered the commercial market. Lockheed, with a board full of big, expensive ASICs, will probably be looking to faster processors (their current offering is clocked at 25 MHz) for an immediate jump in performance, while Omnibyte hopes to get an early entry in the R-4000 market. It's guessed Lockheed will close in on the 50-Mips performance range with its next-generation product.

At this writing, it's expected that the Buscon/91-West show will see another entry in the MIPS Computer market: Performance Semiconductor (Sunnyvale, CA) will offer its first commercial board based on its own R-3000 chip set running at 25 MHz. Later in the year it expects to boost chip speed to 33 MHz and by year-end to 50 MHz. *—Warren Andrews*

Another HDTV team goes with all-digital simulcast system

After a year and a half of cooperative research, Zenith Electronics (Glenview, IL), AT&T Bell Laboratories (Murray Hill, NJ) and AT&T Microelectronics (Berkeley Heights, NJ) have expanded Zenith's spectrum-compatible high-definition television simulcast system. Unlike the original analog and digital approach, the new system incorporates new digital techniques for an all-digital simulcast HDTV system. Of the five contenders vying to be selected as the U.S. simulcast HDTV broadcast standard, the Zenith/AT&T team is the third to go digital. Japan has committed itself to a combined analog and digital approach.

Interference is the largest deterrent to an all-digital simulcast HDTV system—one that will broadcast HDTV on currently unusable channels in the VHF and UHF broadcast spectrum simultaneously with today's National Television System Committee television stations. Zenith, responsible for system definition and transmission technology, has eliminated interference from the standard television signal by placing a digital filter at the HDTV transmitter and a complementary filter in the HDTV receiver. The new technology offers noise-free HDTV signals not only for terrestrial broadcasting, but also for cable, satellite, fiber, and VCRs.

Because all-digital HDTV systems require significantly more video compression than analog counterparts, AT&T Bell Laboratories has developed a video compression algorithm, featuring motion compensation and adaptive quantization, that squeezes the HDTV data into a 6-MHz channel without loss of resolution. Among the digital signal processing chips being developed by AT&T Microelectronics for prototype Zenith HDTV receivers is one that performs about 4 billion operations/s—1,000 times the performance of a typical desktop computer. *—Barbara Tuck*

Toshiba buys into DRAM futures, Moto buys out

Illustrating both its commitment to the memory market and the staggering price of admission to this exclusive game, Toshiba Corp, represented in the United States by Toshiba America (Irvine, CA), has announced the start of a new DRAM plant site in Yokkaichi, Japan. The plant, according to Toshiba, will manufacture 4-Mbit and 16-Mbit DRAMs, beginning production in 1992.

The 135,000-ft^2 facility is expected to peak at 3.5 million chips per month. Toshiba estimates the cost of the plant will total 100 billion yen. The size of the plant and the level of investment suggest that Toshiba doesn't expect a worldwide softening of DRAM demand in the long term, nor does the company seem to feel that surplus CMOS manufacturing capacity will be a big issue by 1992.

In a perhaps unrelated move, Motorola Semiconductor (Austin, TX) has announced that it won't manufacture 4-Mbit DRAMs after all. The company says it will focus its attention on its existing 1-Mbit DRAM capability and on future SRAM products. Motorola will rely on its relationship with Toshiba to provide for the company's future 4-Mbit and 16-Mbit DRAM requirements. This represents a turnabout from previous Motorola strategy, in which the company sought to acquire DRAM technology from Toshiba in exchange for 68000-family CPU technology.
—Ron Wilson

Continued on page 10

Tough, rugged boards that handle shock, vibration, heat, cold and the budget squeeze

MATRIX VMEbus Rugged Series loves harsh environments. Designed to operate from -40° to +85°C, the Rugged Series surpassed these temperature specs and sustained continuous operation during a series of severe environmental tests.* At off-the-shelf prices, these boards are tough to beat.

Rugged 32-bit processors, memory boards, and a variety of specialty I/O boards make harsh environments manageable. Available in extended temperature and/or rugged versions, these products provide all the power and flexibility VMEbus can offer. And at a fraction of the cost of full Mil-spec products.

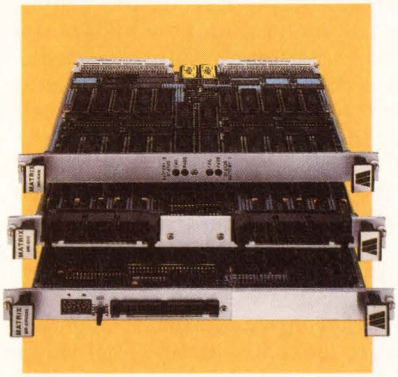

Call today for more information about our Rugged products. For less extreme applications, ask about our full line of VMEbus Industrial Quality systems, boards and enclosures.
Telephone: 1-919-231-8000.
FAX: 1-919-231-8001.

MATRIX CORPORATION

*Test results for the MATRIX Rugged Series show boards withstood 105 g's of shock for 6 ms and 10 g's of vibration at the first natural resonant frequency.

See us at Buscon West, Booth #809

CIRCLE NO. 6

1203 New Hope Road, Raleigh, NC 27610

Continued from page 8

U.S., Japan eye joint Object standards

The efforts of a large consortium of vendors to establish standards for object-oriented computing in the form of the Object Management Group (OMG) appear to be bearing commercial fruit. HyperDesk (Westboro, MA), a new company formed from a project started within Data General, has announced that it will develop and market a new class of software that will give users access to applications across computer networks via a single-user interface. HyperDesk is being financed by ASCII (Tokyo, Japan), one of Japan's largest vendors of systems and application software for personal computers, workstations and network servers.

HyperDesk's software will be based on distributed object management concepts, and products will be submitted in response to OMG's first Request for Proposal in its effort to establish a common framework for object-oriented software and standards. The framework envisions programs as freestanding components, or objects that can be combined and rearranged and can communicate with each other across networks, hardware environments and operating systems. *—Tom Williams*

Stacked-trench structures target 64-Mbit DRAMs

Researchers at Toshiba America Electronic Components (Irvine, CA) have developed a new memory cell structure that advances the level of integration beyond that necessary for the development of 64-Mbit DRAMs. According to Toshiba, the asymmetrical stacked-trench (AST) structure reduces the area of individual cells to less than 60 percent of that of cells in the simpler trench structure used for 4-Mbit and many of the experimental 16-Mbit DRAM chips.

Whereas Toshiba's prototype 16-Mbit DRAMs use 0.6-μm technology, company researchers have achieved a 0.4-μm lithography for the 1.53-μm^2 64-Mbit DRAM cell with Excimer Laser equipment. With trenches located asymmetrically in the AST structure rather than in parallel as in conventional trench structures, Toshiba achieves a 6-μm separation between each adjacent trench cell, as well as between the electron-flow path and the storage capacitor, thus maintaining the minimum distance required to avoid leakage between and within cells.

According to Frank Stefan Becker, press officer for Siemens (Munich, Germany), issues of process simplicity and height differences over chip topology are as significant as leakage when striving for the miniaturization level required for a 64-Mbit DRAM. Siemens, which has partnered with IBM (Armonk, NY) on 64-Mbit DRAM development, is experimenting with a 1.62-μm^2 stacked-trench capacitor cell characterized by extremely thin capacitor electrodes combined with a 5-nm dielectric. Specifications of the experimental 64-Mbit DRAM, which requires a 3.3-V power supply, include 0.4-μm lithography and a 40-ns access time. *—Barbara Tuck*

1991—The year of the cache?

1991 will be "the year of the cache," according to Ray Alderman, technical director of the VFEA International Trade Association (Scottsdale, AZ). He believes that cache technology will dominate the design front and will be an element of all new high-performance systems. And while cache is accepted as critical to many processor and systems today—and is one of the critical elements of the Futurebus+ specification—it's not universally considered a panacea for all high-performance systems.

Some tend to look at cache as being analogous to a turbocharger on an automobile engine. Not too many years ago, turbochargers were looked at as the ideal way to increase auto horsepower without the weight associated with a larger displacement. Turbocharged engines are still made, and in the right applications provide enormous benefits for both automobiles and aircraft. But the benefits are not universal.

So it will be with cache, some believe. "1991 will be the year of the cache just as 1986 was the year of the turbocharger," speculates one system maker. "Cache will be a significant component in memory architectures; however, it will have to be treated with extreme caution not to have the wrong effect on performance. Particularly with today's processors. With multilevel pipelines, cache could actually reduce, rather than increase, performance."

—Warren Andrews

U.S. semiconductor firms gain market share

For the first time in 12 years, U.S. semiconductor manufacturers gained ground against their Japanese counterparts. Falling DRAM prices, a new study says, is the foremost cause.

Figures from market research firm Dataquest (San Jose, CA) show that U.S. firms grew nearly 2 to 36.5 percent of the estimated $58.4 billion worldwide market for semiconductors. Intel and Motorola, the two largest domestic chip makers, drove the trend. Both companies produce microprocessors and special-purpose ICs and benefited from a 23 percent worldwide increase in these markets. According to Dataquest, Intel leaped from eighth to fifth place in the rankings, while Motorola held its spot as the fourth largest supplier.

While hurt by DRAM price drops, Japanese chip makers continue to be the world's biggest chip suppliers. NEC, Toshiba and Hitachi still hold the top three spots, although each saw an estimated 1 percent decline in revenue because of a heavy reliance on memory chips, according to Dataquest. Despite high demand for these chips, there are plenty in supply. This scenario lead to sharply falling prices during the past year. *—Jeffrey Child*

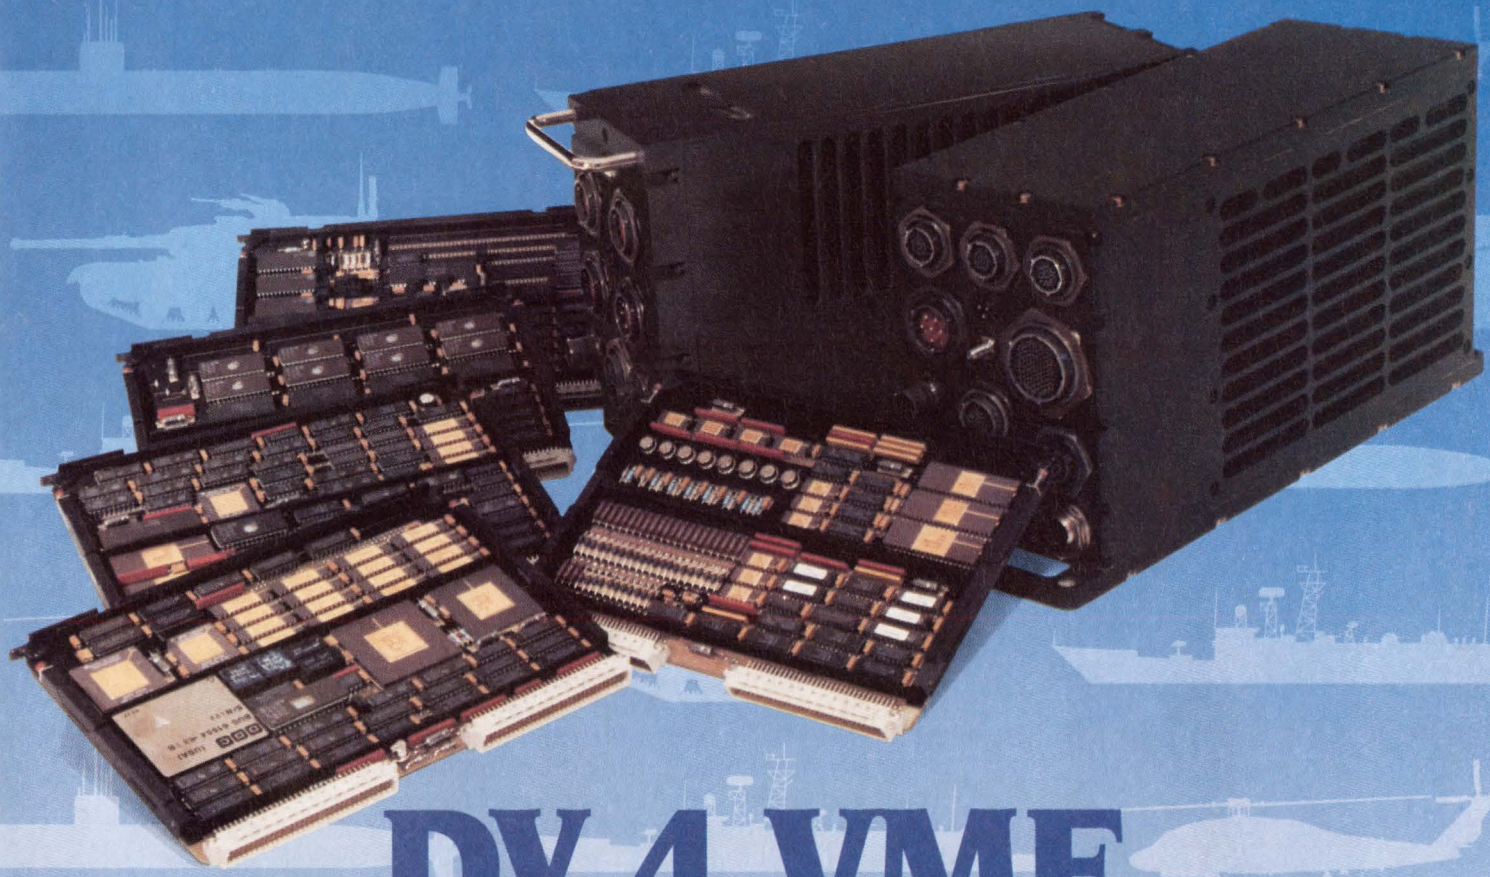

DY 4 VME

When you must win!

For your "must-win" defense and aerospace programs, put the world's leading supplier of MIL-SPEC VMEbus systems on your team. Ready-to-run VMEbus systems from DY 4 are selected overwhelmingly by system integrators for aerospace and defense programs world-wide.

DY 4 provides performance, reliability and cost-effectiveness through integration of a full range of open-system VMEbus products and services to military, ruggedized and commercial standards.

DY 4's system solutions incorporate non-developmental item (NDI) products from the broadest product line in the business - CPUs... memories... communications controllers... analog I/O... high-performance graphics engines... chassis... Ada* foundation software and built-in-test (BIT) diagnostics.

DY 4 provides a comprehensive quality program to MIL-Q-9858A and fully compliant configuration management to MIL-STD-483; design procedures conform to MIL-STD-1521 with manufacturing according to MIL-I-48604 (quality control) and soldering to MIL-STD-2000 in an ESD-controlled environment.

Customer First, Quality Always

DY 4 Systems Inc.

21 Fitzgerald Road, Nepean, Ontario K2H 9H4 Tel: (613)596-9911 Fax: (613)596-0574

Campbell, CA
Tel: (408) 377-9822
Fax: (408) 377-4725

Nashua, NH
Tel: (603)595-2400
Fax: (603)595-4343

Hanmel, Denmark
Tel: +45-86-963624
Fax: +45-86-962575

Pennant Hills, Australia
Tel: +61-2-484-6314
Fax: +61-2-875-1665

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

CIRCLE NO. 7

A confederation of fools

The United States isn't setting a standard; it's filling in an order form.

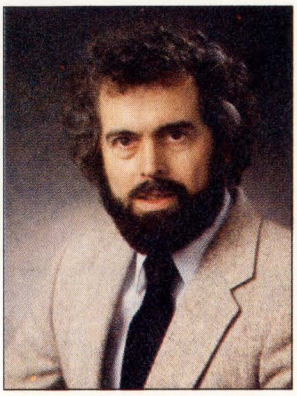

Tom Williams
Senior Editor

Perhaps a good way to view the differences between American and Japanese business strategies is to consider the games of chess and go. In chess, the object is to capture individual pieces and corner the king. In go, the object is to occupy strategic positions to control territory. The acquisition of MCA by Matsushita serves as an example of the devastating effects of a go mentality vs. a chess mentality, and bodes ill for the future.

The first reactions in the press were that the U.S. culture was so lucrative that the Japanese probably would not block the release of films like *Tora! Tora! Tora!* This is the judgment of a player who thinks he has just lost a rook when in fact he has lost half of his maneuvering area. The MCA acquisition, along with the earlier acquisition by Sony of CBS Records, is really about digital audio tape and high-definition TV—and even more ominous issues beyond those.

The situation leaves the impression that the U.S. government and business elite is a confederation of fools—fools who treat employees like the enemy; who measure success in three-month increments; who can't bear to invest in research without assurance of an immediate return.

In the past few years the recording industry was obsessed with defending its queen—an outmoded revenue and royalty structure—from the DAT threat. U.S. manufacturers, cowed by the threat of lawsuits, failed to design and develop DAT products. While Congress dithered, the administration recited platitudes about free enterprise. During this pointless melodrama, small white stones were calmly being placed on the board.

The result is that the recording industry's foolish resistance to DAT has been broken. DAT will emerge in the market, but only from Japanese companies. The next stone awaits FCC approval of a U.S. HDTV standard.

A significant research investment by the U.S. television industry has developed key technologies—to the ultimate economic benefit of those controlling the manufacture of television sets, video cameras and, now, vast libraries of program material. While the FCC will proudly proclaim a U.S. standard, the Japanese will view it in its true form—a manufacturing specification. The United States isn't setting a standard; it's filling in an order form.

The costs of fab for advanced submicron ICs are so high that only the volume of the consumer market can absorb the expense. Thus he who controls the consumer market will eventually control the semiconductor market. Another white stone is placed on the board. Advances in semiconductor technology will flow not from the military and aerospace down to the consumer but from consumer applications into military and space fields by reason of simple economics. And, should the Go Masters become so inclined, it's only a short step to domination of military and space technologies.

Editor's Note: Tom recently returned from a two-week editorial visit to Japan. His observations confirmed some of our loudly voiced positions.

"THEY LAUGHED WHEN WE PLUGGED A PC IN OUR VME SYSTEM...

"The pressure was on. Shorten our design cycle. Cut our software costs. Deliver the "gee-whiz" features customers wanted. But the old solutions weren't working. We needed a new approach.

We looked to the PC. Others looked at us like we were nuts. They sent memos. *'The PC is not a* real *computer.' ...'Not enough horsepower' ...'Just a pretty user interface' ...'It can't survive that environment.'*

But the PC is going places. Over 40 million are in use; 4 million in industrial environments. Another 1.4 million are expected on the plant floor this year. From the moment we plugged in a RadiSys Embedded PC, we understood why. **Software for PCs is abundant and inexpensive.** We had two dozen software houses fiercely competing to sell us high quality, man-machine interface software. Why should we reinvent the wheel? For multi-tasking operating systems, we could choose from Windows 3.0, OS/2, UNIX, VRTX—all with integrated VMEbus support from RadiSys.

There is a wealth of development tools available to shorten design time. No wonder the PC is the world's most popular software development host. And PC-based networking is light-years ahead. **PC versatility is unmatched.** We can give customers what they want, right now. More options. Proven features. Everybody knows how to use the PC, our own team and our customers. And it's a sure bet that future application programs, languages and OS standards will run on it. **PC horsepower is up there.** The new Embedded 386 and 486 PCs have more than enough power to handle our mix of user interface and control functions. **We got the best of both worlds.** Plugging a RadiSys Embedded PC directly into the VMEbus gave us the full performance, ruggedness and reliability of the VME form factor, plus all the PCs software advantages. **You know the rest of the story.** That Embedded PC has the *whole company* laughing. All the way to the bank."

Join the party! Call RadiSys at 800/950-0044 (fax requests: 503/690-1228) for a catalog of Embedded PC Products and brochure *"Open New Windows of Opportunity With Embedded PCs."*

RadiSys™
CORPORATION

The Inside Advantage

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

Copyright ©1990 RadiSys Corporation, Inc. All rights reserved. EPC is a registered trademark and RadiSys is a trademark of RadiSys Corporation. Windows 3.0 is a trademark of Microsoft Corporation. OS/2 is a trademark of International Business Machines Corporation. UNIX is a registered trademark of AT&T. VRTX is a registered trademark of Ready Systems, Inc. 386 and 486 are trademarks of Intel Corporation.

CIRCLE NO. 8

NEW

Introducing the Highest Standards in the Industry.

Gates New High Capacity VALUMAX™ Line.

At Gates, our standards have reached new heights. With our VALUMAX line of Ni-Cd cells—the best value cells on the market.

VALUMAX cells are available in all the sizes you need to make virtually any product perform like no other standard cell can. For example, the VALUMAX C cell gives you up to 25% more run time. And our 4/5 Cs, Cs and AA cells deliver as much as 33% more capacity than even our closest competitor.

In fact, with VALUMAX you get premium performance without paying the premium price. And that makes VALUMAX the best value—dollar for dollar—of any cell you can buy.

So isn't it time you raised your standards? Gates VALUMAX—we're setting the standard.

Gates Energy Products, Inc., P.O. Box 861, Gainesville, FL 32602, **1-800-67-POWER.**

Gates Energy Products

The power of great ideas.

CIRCLE NO. 9

NEW

You'd Have to Go to Great Lengths to Match the Capacity of Our New VALUMAX™ Cells.

No matter how you stack it, when it comes to capacity, other Ni-Cd batteries seem to come up short. That's because Gates' new VALUMAX cells pack more capacity than any other standard cells on the market.

At Gates, we've gone to great lengths to give you the longest possible run times in a wide range of cell sizes for virtually every application. For example, the VALUMAX C cell gives you up to 25% more run time than any other standard C cell. And the VALUMAX 4/5 Cs, Cs and AA cells deliver as much as 33% more run time than even our closest competitor.

In fact, with VALUMAX you get premium performance without paying the premium price. And that makes VALUMAX the best value—dollar for dollar—of any standard cell you can buy.

The way we see it, the others still have a long way to go. And, at Gates, we think you shouldn't have to go to great lengths when you want a standard cell that lasts a long time. Gates Energy Products, Inc., P.O. Box 861, Gainesville, FL 32602, **1-800-67-POWER.**

Gates Energy Products

The power of great ideas.

CIRCLE NO. 9

NEW & NOW

Mix & match

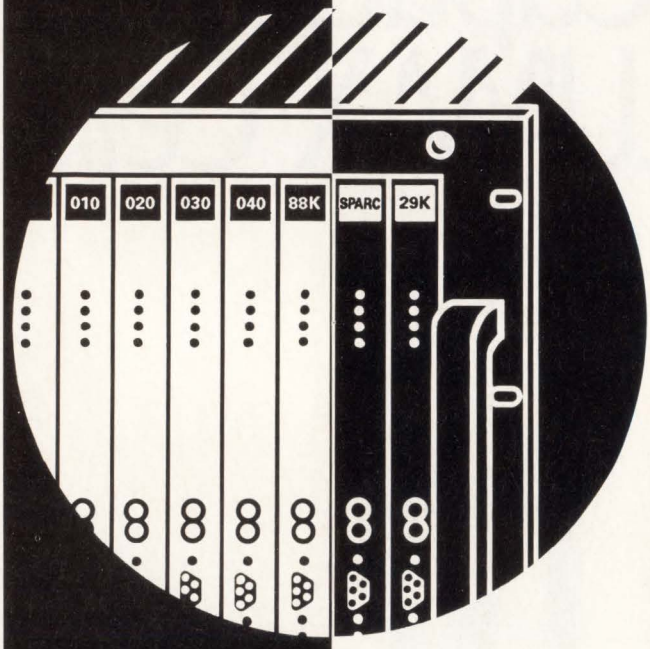

Ironics now enables you to mix and match processor families. The *ultimate* in performance and flexibility.

IRONICS

Call our Product Hotline today 800 334 4812

ADVANCING VME MULTI-PROCESSING

CIRCLE NO. 10

CALENDAR

CONFERENCES

February 5-7
AFCEA 2d Annual Military/Government Computing Conference and Exposition

Hyatt Regency Crystal City, Arlington, VA. The exposition will offer the latest in computer software and hardware products to link military and government needs with the latest industry technologies, applications and services. The conference will also feature tutorials, including "Open Systems Architecture" and "Information Engineering," and technical tracks, including "Technology Advances" and "Software Development/Maintenance." Information: AFCEA Programs Office, 4400 Fair Lakes Ct, Fairfax, VA 22033-3899, (703) 631-6125. **Circle 366**

February 24-28
Nepcon West '91

Anaheim Convention Center, Anaheim, CA. More than 1,000 exhibitors will display the latest in products and technology for the design, fabrication, assembly, packaging, inspection, and test of printed circuits and electronic assemblies at this 29th annual conference and exhibition. Special features include over 50 sessions on surface-mount technology, several sessions on applicable environmental issues and a session on planning for the 1992 European Economic Community. Information: Cahners Exposition Group, 1350 E Touhy Ave, Des Plaines, IL 60018, (708) 299-9311. **Circle 367**

March 13-20
Hannover Fair CeBIT '91

Hannover Fairgrounds, Hannover, Germany. The more than 4,000 exhibitors from 40 countries at CeBIT will attract computer and communications professionals from 100 countries. The show will include NetWorld Europe, an exhibition and conference presenting 200 networking and connectivity firms from around the world. NetWorld will feature demonstrations such as EurOSInet and Multinet, as well as conferences and seminars on developments in network computing, LANs and wide area networks. Information: Hannover Fairs USA, 103 Carnegie Center, Princeton, NJ 08540, (609) 987-1202. **Circle 369**

April 8-11
NEMDE '91

McCormick Place, Chicago, IL. The National Electronic Manufacturing and Design Exposition and Conference is designed to present a complete spectrum of products and equipment to the entire engineering and management team. The conference will include 14 technical sessions, 12 workshops and eight professional advancement courses. Sessions include "Design for Thermal Management," "Functional Test" and "Infrared Reflow Soldering Considerations." Information: Cahners Exposition Group, 1350 E Touhy Ave, PO Box 5060, Des Plaines, IL 60017-5060, (708) 299-9311. **Circle 370**

A World of Choice

When it comes to mass storage needs, your choice is clear: Ricoh optical disc drives. From WORM and rewritable to bare drive and entire library units or subsystems, Ricoh offers a world of choice. Each drive boasts quality you can trust to safeguard vital information. The kind of quality that comes from over 50 years of know-how in optics and image processing, in such products as cameras, copiers, facsimiles, telepress, laserprinters and scanners. Now you know why Ricoh is the world's leading supplier of 5-1/4″ optical disc drives. For more information, call 1 (800) 955-FILE.

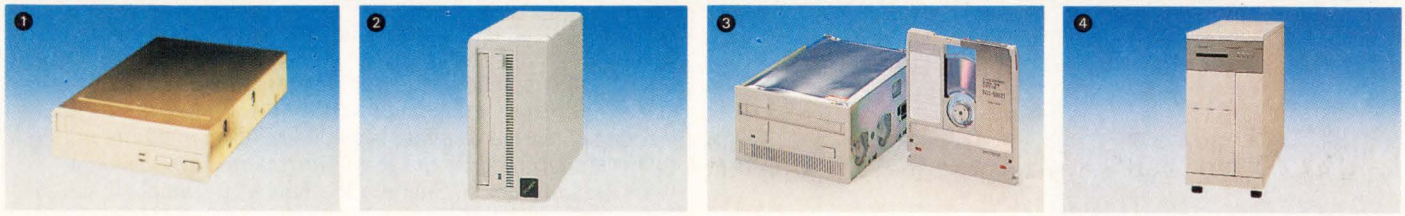

❶ RH5500/5-1/4″ REMOVABLE CARTRIDGE HARD DISK DRIVE
❷ RS-9100H/5-1/4″ HALF-HEIGHT WORM SUBSYSTEM
❸ RO-5030EII/REWRITABLE OPTICAL DISK DRIVE
❹ RJ-5330E/MAGNETO-OPTICAL DISK LIBRARY UNIT

BEST FOR QUALITY.
UNMATCHED FOR SERVICE.

RICOH

RICOH CORPORATION **(File Products Division)** 5150 El Camino Real, M/S D ~ 20, Los Altos, Ca 94022 U.S.A Phone: 415-962-0443 Fax: 415-962-0441
RICOH CORPORATION **(Eastern Region Office)** 59 Stiles Road, Suite 102, Salem, NH 03079 Phone: 603-893-0547 Fax: 603-893-4162

CIRCLE NO. 11

NEW & NOW

Multiple SPARCS

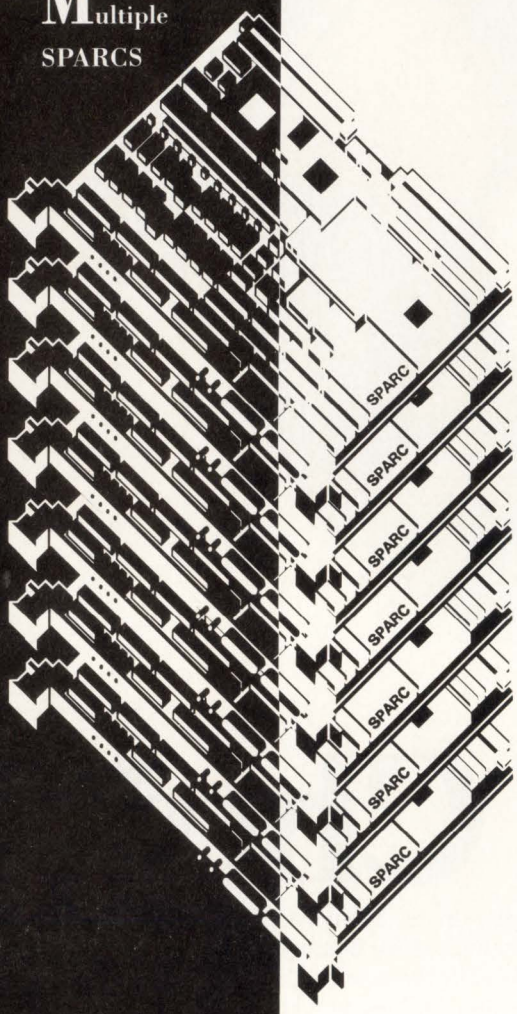

Ironics now offers the *latest* SPARC technology optimized for multiple processor VME systems.

IRONICS

Call our Product Hotline today 800 334 4812

ADVANCING REAL-TIME RISC SOLUTIONS

CIRCLE NO. 12

CALENDAR

CONFERENCES

April 9-12
SMTCON

Trump Plaza Hotel and Atlantic City Convention Center, Atlantic City, NJ. More than 250 suppliers will demonstrate products at the second annual Surface-Mount Technology Conference and Exposition. The exposition will offer a comprehensive range of components, equipment, materials, and services used in the design and manufacture of surface-mount and mixed technology products. The conference features advancement courses, technical forums, technical sessions, and workshops. Information: IC Management, 900 N Franklin St, Suite 700, Chicago, IL 60610, (312) 944-3434. **Circle 371**

April 30-May 2
Federal Computer Conference West

Disney Convention Complex, Anaheim, CA. This third annual conference will concentrate on advances in key computer and communications technologies. A Solutions Showcase, which combines classroom sessions with live demonstrations, will cover solutions to computing applications needs made possible by advances in microcomputers and LANs. The conference will also feature in-depth seminars on microcomputers and networking. Information: Federal Computer Conference West, 8455 Colesville Rd, Suite 700, Silver Springs, MD 20910, (800) 343-6944. **Circle 372**

SEMINARS

February 13-14
Concurrent Engineering Seminar

Radisson Plaza Hotel, San Jose, CA. This two-day event is targeted toward the management and engineering team, including those involved with quality assurance and design, test and manufacturing engineering. The one-day concurrent engineering seminar will focus on implementing the simultaneous design of products and the processes by which they're built, tested and supported. The one-day technical session will cover design for manufacturability, design for test and design for service. Seminar sessions in other locations will follow in March, April and May. Information: Logical Solutions Technology, 96 Shereen Pl, Suite 101, Campbell, CA 95008, (408) 374-3650. **Circle 373**

March 11-15
IEEE Software Engineering Seminars

Knickerbocker Chicago, Chicago, IL. The seven IEEE Standards Seminars will focus on software reviews and audits, quality assurance, testing, project management, verification and validation, requirements specifications, and configuration management. Each seminar is two or three days long, is led by a team of industry experts and is targeted to specific groups in the software arena. Information: IEEE Standards Seminars, 445 Hoes La, PO Box 1331, Piscataway, NJ 08855-1331, (908) 562-3805. **Circle 374**

ON JAN. 29, THOUSANDS OF VME DESIGNERS WILL TURN 040.

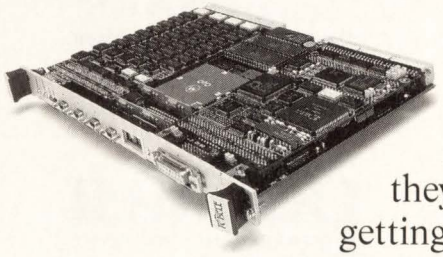

They're not getting older, they're just getting better.

Because that's the day FORCE unveils the first 040 single board computer. Available in quantity. For immediate delivery.

So get FORCE for better performance, speed, functionality and compatibility.

See us at BUSCON Booth #654, or call 1-800-BEST-VME, ext. 40. And give your competition some gray hairs.

FORCE COMPUTERS®

VME at its best.

FORCE Computers, Inc., 3165 Winchester Blvd.
Campbell, CA 95008-6557 (408) 370-6300 ext. 40
© 1991 FORCE Computers, Inc.

B U S C O N # 6 5 4

CIRCLE NO. 13

LAN Magnetics and Filters

Solutions to the High Cost of Real Estate

ETHERNET
10Base-T and
Token Ring 16 Mbps

- Miniaturization saves 40% on real estate
- Lowers production costs
- Designed to meet IEEE 802 performance standards
- Thru-hole and surface mount
- Discrete or integrated modules

Quality On The Leading Edge

FEE Fil-Mag
Filters • Magnetics • Delay lines

4787 Cardin Street
San Diego, CA 92111-1416
Phone: (619) 569-6577
Fax: (619) 569-6073

CIRCLE NO. 14

25 & 33MHz
486/386 Single Board Computers
Thousands Shipped

486

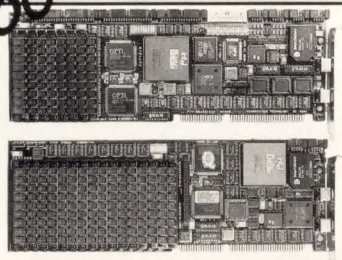

CAT1010
- Fully Integrated 80486 25/33MHz /AT CPU
- Up to 32Mb DRAM Memory
- 8Kb Internal CACHE
- 6,8,12MHz Bus Speed
- Phoenix BIOS
- Floppy Interface
- SCSI & IDE Interface
- 2 Serial/1 Parallel Port
- Up to 512Kb PROM
- On-board Battery
- 4.7 X 13.4"

CAT1000
- 25/33MHz 80486 /AT CPU
- Up to 32Mb DRAM Memory
- 8Kb Internal CACHE 128Kb sec. CACHE avail.
- 6,8,12MHz Bus Speed
- Phoenix BIOS
- On-board Battery
- 4.7 X 13.4"

386

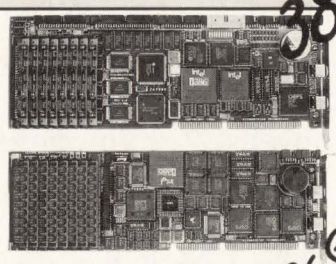

CAT990
- Fully Integrated 80386 25/33MHz /AT CPU
- Up to 32Mb DRAM Memory
- 64 or 128Kb CACHE
- 6 or 8MHz Bus Speed
- Phoenix BIOS
- Floppy Interface
- SCSI & IDE Interface
- 2 Serial/1 Parallel Port
- On-board Battery
- 4.7 X 13.4"

CAT980
- 25MHz 80386 /AT CPU
- Up to 8Mb DRAM Memory 20Mb w/daughter card
- 32Kb CACHE
- 6,8,10,12MHz Bus Speed
- Phoenix BIOS
- On-board Battery
- 4.2 X 13.4"

386SX Too!

'486', '386', '386SX', '286' & '8088'
Complete FCC/U.L. Recognized Systems Available

Diversified Technology
An Ergon Co.

1-800-443-2667
Outside USA Ph/201-891-8718 or Fax/201-891-9629

CIRCLE NO. 15

486 & 386 are trademarks of Intel Corp.
Product names are trademarks of their respective holders

COMPUTER DESIGN
Technology and Design Directions

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

SENIOR EDITORS
Warren Andrews, *Computers & Subsystems* (508) 283-2102
Barbara Tuck, *ASICs* (516) 744-2984
Michael G. Donlin, *Design and Development Tools* (508) 392-2123
Tom Williams, *Software* (408) 335-5514

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

EDITOR: *News Edition*
John H. Mayer (508) 392-2118

NEWS EDITOR: *News Edition*
Elizabeth A. Stein (508) 392-2116

MANAGING EDITOR: John Donovan (508) 392-2115

MANAGING EDITOR: *News Edition*
Barbara-Ann Scofidio (508) 392-2121

ASSISTANT EDITORS:
Sue Cavallaro, Anne Holm

PRESENTATION MANAGER: Richard D. Sarno

ASSISTANT PRESENTATION MANAGER:
Claire Coupal

EDITORIAL ASSISTANTS: Claire Ellis, Terrie Robertson

ART DIRECTOR: Jan Horner

ILLUSTRATORS: Fred Lomas, Manuel King

COVER ART: Bill Morrison

AD TRAFFIC MANAGER: Betsy Anderson

CIRCULATION DIRECTOR: Robert P. Dromgoole

CIRCULATION MANAGER: George Andrew

PUBLISHER
David L. Allen (508) 392-2111

ASSOCIATE PUBLISHER/NATIONAL SALES MANAGER:
Tim L. Tobeck (516) 883-8815

PennWell PUBLISHING COMPANY

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

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

BPA ABP

HITACHI'S DK515, 5¼", 780 MB WINCHESTER

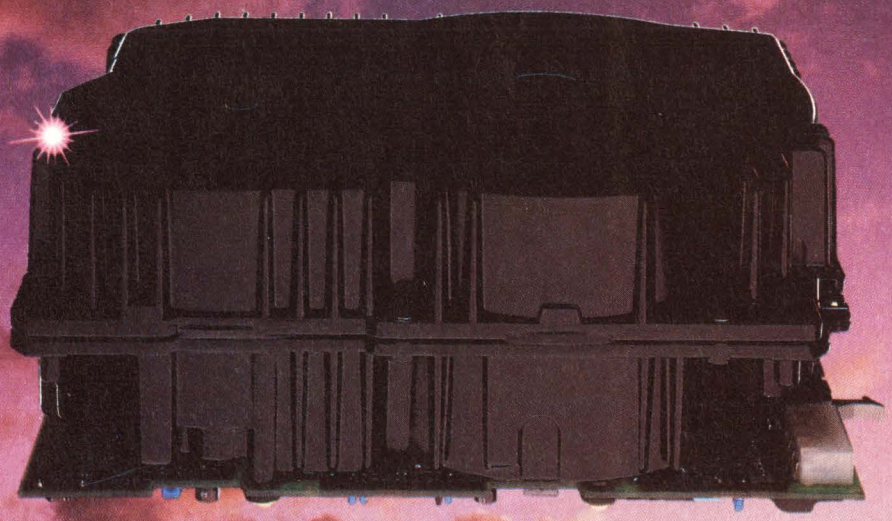

From any point of view, this 5¼" disk drive reflects Hitachi's superior technical expertise, high-performance, and quality.

The DK515 features a fast 2.46 MB/sec. data transfer rate, an average access time of 16 ms, and a choice of ESDI, SCSI, or ESMD interfaces.

Like all Hitachi drives, the DK515 reflects quality, because all critical components—including heads, media, and servo systems—are designed, engineered, and manufactured by Hitachi.

Then, to make sure that Hitachi's strict standards of excellence are maintained, each and every drive is 100% burned-in and tested. The final result: a clear reflection of Hitachi's commitment to providing the very best in reliable, high-quality, high-performance disk storage. All this from a $48 billion company.

A Reflection of Quality

Available now!

For more information about Hitachi disk drives, call your local Hitachi Distributor listed below, or Hitachi at **1-800-283-4080, Ext. 877.**

Hitachi America, Ltd.
Computer Division, MS500
Hitachi Plaza
2000 Sierra Point Parkway
Brisbane, CA 94005-1819

HITACHI®
Our Standards Set Standards

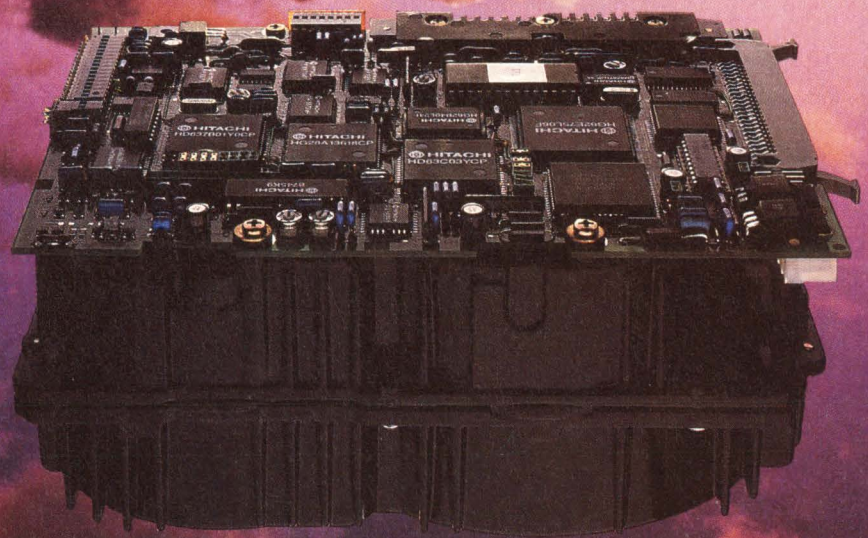

Authorized Distributors

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)

CIRCLE NO. 16

EPSON®
IC MEMORY CARDS

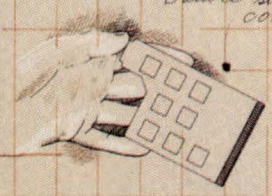

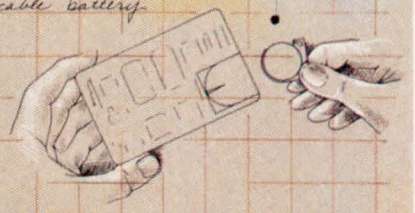

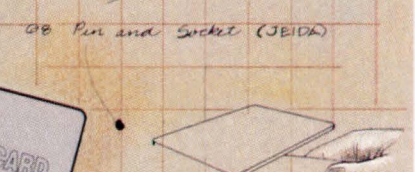

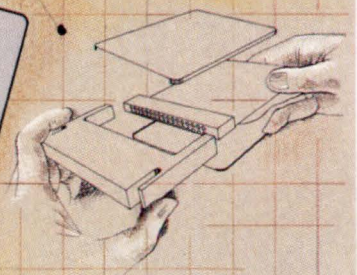

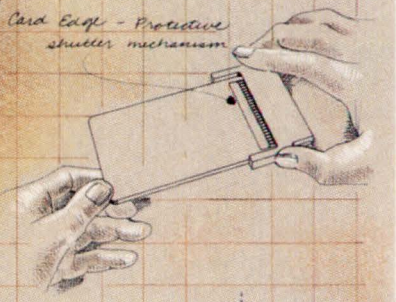

JEIDA/PCMCIA

SRAM	64kB	to	2M
OTP EPROM	64kB	to	2M
Flash EEPROM	64kB	to	1M
Mask ROM	256kB	to	8M
EEPROM	16kB	to	32kB

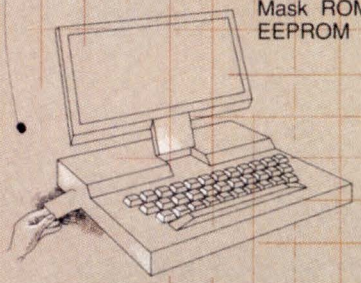

CARD EDGE

SRAM	8kB	to	1M
OTP EPROM	32kB	to	1M
Flash EEPROM	128kB	to	1M
Mask ROM	128kB	to	4M
EEPROM	8kB	to	32kB

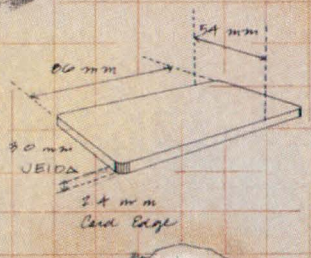

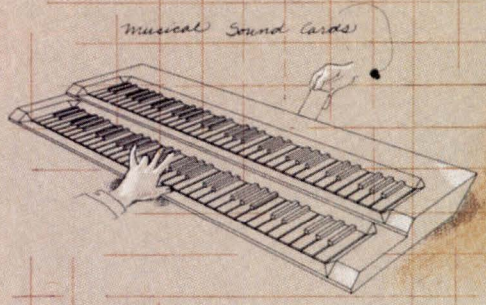

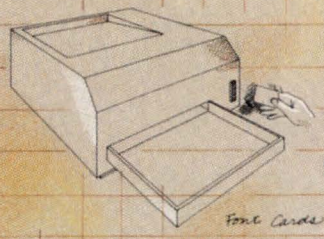

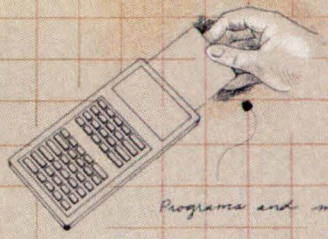

EPSON EPSON AMERICA, INC.
Component Sales Department Telephone: 213/787-6300

Call your sales rep today.

EPSON Sales Representatives: **AL-GA-TN-NC** Concord Components 205/772-8883 . **AZ-NM** Fred Board Assoc. 602/994-9388 . **CA-No.** Costar 408/446-9339 . **CA-So.** Bager Electronics 714/957-3367 . **CO-UT** Wn. Region Mktg. 303/428-8088 . **FL** Dyne-A-Mark 305/771-6501 . **IL-WI** LTD Technologies 708/773-2900 . **IN-KY** C C Electro 317/921-5000 . **KS-MO-IA** Microtronics 913/262-1444 . **MA-NH-CT** Rosen Assoc. 617/449-4700 . **MD-VA** Tech Sales Assoc. 301/461-7802 . **MN** Electro Mark 612/944-5850 . **NJ** JMR Sales 201/525-8000 . **NY** Elcom Sales 716/385-1400 . **METRO-NY** Niktronix 516/929-4671 . **OH-MI** J.D. Babb Assoc. 216/934-4454 . **OR-WA** Matrex 503/245-8080 . **PA** Omega Sales 215/244-4000 . **TX-OK** Component Tech. 214/783-8831

CIRCLE NO. 64

Pete Yeatman on: Bus-board architectures

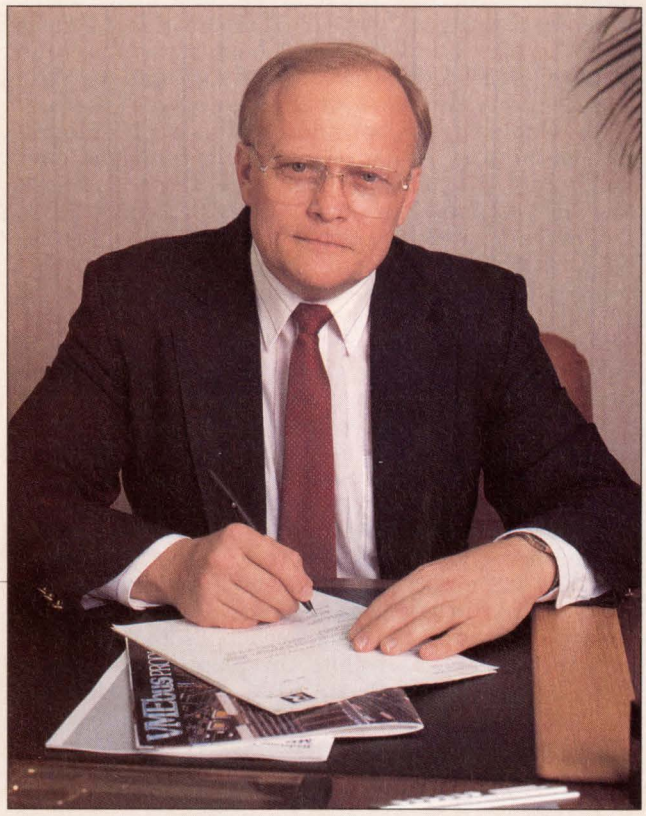

For the past 20 years I've heard all kinds of dire predictions about how the improvements in semiconductors—higher levels of integration, the combination of multiple functions and super high speed devices—were going to be the end of the standard board business. During that time, I've seen mainframe computers go from industry giants to near extinct dinosaurs; the once dominant minicomputer fade into obscurity to be replaced by desktop workstations; and the personal computer become the ubiquitous writing and graphic tool, replacing the typewriter, pencil and drafting tools.

What I haven't seen, though, is the demise of the board business. In fact, over that 20 years, I've seen the standard board business grow and mature from what some considered a stepchild of the computer business—a necessary evil—to a healthy self-sustaining business with a significant number of major players. And, according to current estimates, it's expected to grow at a double-digit compound annual rate over the next several years. No doubt the growth of the market for standard boards is a function of the growth of the larger market for computers. Outperforming the growth of the computer market, standard boards are capturing a growing share of that market's hardware spending by supplanting proprietary solutions.

The advances in standard boards haven't been free. They continue to be a direct consequence of the research invested in technological and design innovation by companies participating in this business. For the progress to continue, board and systems designers will have to be increasingly sharp to design products that take advantage of the latest silicon developments and serve ever-more complex applications. In addition, manufacturers will have to dig deeper and deeper into their R&D pockets to keep the magic happening. The challenge to the board manufacturing business, therefore, is to continue to provide the performance that systems makers would like to see with proprietary designs, on a standard bus platform. This will undoubtedly continue to kindle the best architectural genius in the industry.

The founding fathers of the standard buses—VMEbus, Multibus, STD, Futurebus, and the others—demonstrated remarkable foresight when they developed standard architectures that remain alive today. VME, for example, was developed in the late 1970s and first introduced in 1981—the reign of the 8-bit processor. And while there were many provisions in the specification to handle growing semiconductor technology, its authors could not anticipate the rapid advances in silicon speed and density and in design turnaround time.

Yet these standards have remained viable, partially because the authors anticipated certain developments, and because the standard was too important to discard just because new technology came along. Instead, board designers buckled down to the task of adapting the newer technology to older buses.

Today the effort continues, and continues to be successful. The techniques used now to put new technology on VME and Multibus will undoubtedly be useful with the next generation of boards. And so it will be with Futurebus+. Despite the Herculean effort of the working committee, it's almost impossible to foresee the future's technological developments. Thus, designers will be working around the "restrictive" areas of standard buses for a long time despite prophesies that semiconductors will obsolete standard buses.

Silicon spurs board growth

Sure, new generations of semiconductors—microprocessors, logic, memory, and application-specific parts—are having, and will continue to have, a pronounced effect on the board-level industry and the directions it will grow in. But unlike some industry analysts, such as Andrew Allison of RISC Management (Freedom, CA), who believe the board industry

is doomed by higher levels of IC integration, I believe just the opposite. The board business will see an unparalleled growth—because of the availability of new semiconductors—over the next several years.

Undoubtedly, the continuing advance of IC integration permits inclusion of ever-more functions on a single chip. Allison believes this progress is so pervasive that the functions now residing on a single board will be integrated on a single chip in the near future, obsoleting the need for a bus. I agree this level of compaction is, and will be, a reality. But all that means is that the functions of today's multiple-board system will be reduced to a single board. The fact is, some of the 68040-based boards available today provide complete system-level performance. Soon we'll have multiple system capability on one board. Next generation silicon won't hurt the board business; on the contrary, it will give us that much more power per square inch of board space and let systems designers achieve that much more performance and flexibility in a system.

Increasingly higher levels of IC integrations isn't the only force that will drive the board business to new heights in the coming years. Take a look, for example, at the computer powerhouses that are jumping on the open-bus-architecture bandwagon—DEC, Sun, Data General, Prime, IBM, Foxboro, GE, and the list goes on. For some of these companies, the buses will be used as I/O, not systems, buses. For others, the buses are the foundation of industrial control systems. Nonetheless, the participation of so many front-line companies—world-class computer companies—further legitimizes and gives impetus to the standard board industry.

The big computer companies don't want to waste their time or resources building boards to meet their unanticipated and often transitory board needs. In addition, to meet critical market windows, it's imperative to use the expertise of independent third-party vendors. Companies such as DEC, IBM, Sun, and others are willing to let the independent board manufacturers take those risks and enjoy that part of the business. It's so important that even when Sun and DEC developed their own new buses, SBus and Turbochannel, the companies made a major effort to enlist third-party board manufacturers and software companies. In fact, the board-level business grew partially by meeting those kinds of needs, as those of us from the DEC world well remember.

Standardization is key

What makes possible both the multiple-vendor sourcing of various computer functions (that is, boards) and the development of the open architecture board business is standardization. The social and economic climate is leaning increasingly away from proprietary solutions in favor of standards in all areas of the computer/electronics industry—hardware, software and even systems. With the fast-paced developments, nonstandard products are quickly made obsolete, and then must be totally discarded and replaced. But with systems based on standard products, upgrades can often be made within a fixed framework, causing only a small part of a machine, system or program to be discarded.

The standards issues are invading not only the commercial computer business, but also the industrial process and machine control arena as well as a broad array of stationary and mobile (aircraft-, vehicle- or ship-based) military systems. Benefits accrue from standardized training and maintenance procedures, and from fast and relatively painless upgrades as new systems components emerge.

In discussing the board-level business, the concept of an open architecture remains the key and is supported on several fronts. The development of open buses like Multibus I, VMEbus, PC bus, and others has provided designers with an array of possibilities in terms of architectures, mechanical configurations and performance capabilities. Open buses also provide designers with a broad and widening range of boards and suppliers from which to choose.

By the same token, however, such open architectures are, by necessity, restricted. First, the standardization process itself takes long enough to assure a certain degree of obsolescence. Second, standardization can restrict performance and product differentiation. Third, manufacturers need some protection of their designs—which is best provided by a proprietary, rather than standard, architecture—to be economically and competitively viable. Finally, standards tend to keep creative efforts from blossoming into new product ideas.

Despite these obstacles—or perhaps because of them—the standard board business is more vital than ever. It becomes a key to understanding today's implementation of the open-bus strategy, to realize that the disadvantages mentioned above are largely nonexistent thanks to increased levels of IC integration. Designers have the freedom to develop a proprietary architecture on-board for a board that will become part of a system. As mentioned, many single boards are complete systems that operate in consort with other "systems" across a standard bus.

Performance or price?

Performance is another issue. It's a particularly interesting issue because it's a driving force—but by whom? The media? Some percentage of the marketplace? Some fear within certain suppliers' organizations that they'll be left behind? Looking at real applications—what's really needed—the final measurement is usually based on price/performance evaluations, rather than performance alone.

And while there's a significant focus on performance, it will not necessarily require dramatically new and different bus architectures in the immediate future. More likely it will depend on clever board and systems designs on whatever bus standard is used, whether it be VME, Multibus II, Futurebus+, or something we haven't heard of yet.

It's easy to look at the bus specifications alone and blame it for substandard operation, saying "we need to move to a faster bus." Often, the problem is not the bus. Equally often, a clever architectural approach can bring the product up to speed on the older bus at substantially lower cost and with fewer problems than designing for an unfamiliar bus. As we look at a period of relatively tight budgets, there's going to be a tendency to try to make the old bus last another generation or two—or three or more. This not only avoids the need for scrapping existing hardware, but it also preserves software investments as well as other

assets such as parts inventory.

Arriving at the next level expected for systems performance requires an injection of additional computing horsepower. One of the favored approaches of the higher-horsepower processors today is multiprocessing. There are a number of variants on the approach, each with its advantages and disadvantages, and each with an edge in a different application. Looking forward, it's not clear that tightly coupled processors talking across a standard bus—whether simple VME or a cache-coherent Futurebus+—is either an efficient or practical approach to gain computer power in a standard bus system. Multiprocessing will undoubtedly be a major technique in systems based on standard board products, but it will probably comprise tightly coupled processors on a single board that will reside in systems where many such boards will be loosely coupled.

The key to making such systems work will be the individual board architecture that will include multiple internal buses of its own to prevent the main systems bus from being a bottleneck. In today's environment, these situations are handled by architectures designed specifically to accommodate the enormous throughput capability of the enhanced performance boards. To get out of these boards all of the performance that's been designed into them, processing bottlenecks have to be eliminated. Highly effective multiple processor/multiple local bus schemes have already been developed. In one, for example, a 68040 serves all the data processing functions while a separate independent on-board 68020 handles all I/O activities, leaving the 68040 free to flex its processing muscle totally unencumbered by I/O needs.

Boards like this, as well as others entering the market, illustrate how the board business can grow its capabilities without abandoning its existing architectures. And there's much to be said—for vendors and users alike—for upward migration within an architecture, or even with a single vendor's board family. Economies in software development alone make this kind of upward-performance mobility highly efficient, even if not as exciting as the emergence of a brand new bus architecture.

Software needs standards

Software remains one of the most critical, confounding and problematic issues in the standard bus/board industry. By rights the software, not the hardware, should be selected first for most applications that fall into the realm of standard boards. Unfortunately, though, this model is seldom exercised. Instead, the hardware is chosen first and the software selected on the basis of its appropriateness to the application—and, of course, its availability for a particular board. Inappropriate emphasis on hardware during initial system specification is a problem perpetuated by users.

Another set of problems begins with software vendors. The board business is characterized by a highly independent network of board software vendors. Each has its own bag of tricks, and, as effective as those tricks may be, there is virtually no capability for interoperability between different software vendors' products.

Not too many years ago, users pleaded with board makers to adhere to standards and specs in a way that would help assure plug-and-play interoperability between boards from different vendors. Today, that's largely been achieved, even if some tweaking and tuning is required. But on the software side, we don't even have standards yet. We should. The board business deserves that—needs that—to continue to serve its customers.

If the board-level business is to reach new levels of technical sophistication, the software wrinkles will have to be ironed out with new levels of software intercompatibility. We've already seen some effort in that direction within both the Posix group and the ad-hoc OBIOS (Open Basic Input/Output System) consortium. Progress, however, is painfully slow. Other efforts, such as Orkid (Open Real-time Kernel Interface Definition), failed because they didn't spark the minimum level of commonality. More needs to be done. Just as there are standard bus/board architectures, there should be the equivalent of standard software architectures so that the marketplace can avail itself of the benefits of multiple-vendor software, as now it does for hardware.

Design drivers

The computer industry, almost by definition, is in turmoil. And the standard board business, once a poor stepchild to the industry, is now a fully initiated member and well on its way to being a key—perhaps THE key—player in all new computer systems architectures. From workstations that have adopted their own set of open-standard buses—SBus, Turbochannel and Micro Channel—to the Futurebus+ platforms already on drawing boards for incredibly powerful compute servers, standard buses are on the move. In addition, vendors continue to integrate the latest IC technology onto board-level products providing a variety of on-board communications schemes to reduce potential bottlenecks. Even more clever techniques are on the horizon—extending the use of mezzanine boards with their own associated buses to increase transfer rates.

VMEbus, Multibus II and even the workstation buses are pitifully slow for the kind of data transfer speeds expected over the next several years. Networking standards such as Ethernet, for example, are under tremendous pressure to provide higher bandwidth communications between systems. The current work on Ethernet replacements such as FDDI (Fiber Distributed Data Interface) are already seen by some as falling short (at only 10 times Ethernet's performance). Higher-performance links, such as HIPPI (High-Performance Peripheral Interface), are on the horizon. The impact of networks operating at speeds over 100 Mbytes/s on systems bus transfer rates will be profound. The current generation of buses (even early versions of Futurebus+) won't be able to provide the required bandwidth. So, clever board architectures will have to provide the required speed, buffering the systems bus from the overwhelming demands of communications, memory transfers, and video data—just as they do now to satisfy other high-speed requirements.

Pete Yeatman is president of Radstone Technology (Montvale, NJ).

IN THE ERA OF MegaChip™ TECHNOLOGIES

A lot has been said about company is doing a lot about

testability, but only one it. Texas Instruments.

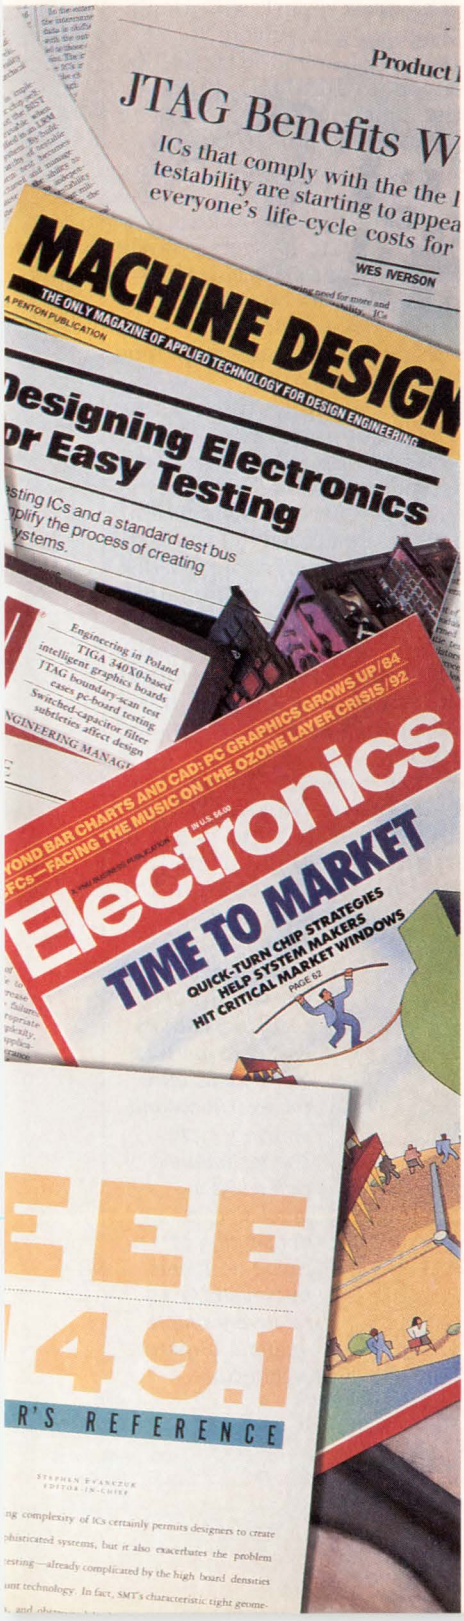

You've seen the headlines and read the stories. Design-for-test (DFT) is a challenge but one that's now easier to live with. The reason: Texas Instruments is the first to develop products for implementing the JTAG/IEEE 1149.1 testability standard quickly and effectively.

To market faster at lower cost

By implementing testability into your system from the outset, you can create one that uses high-performance circuits and is readily manufacturable, one that is lower in total cost and on the market faster. You can expect:

- Test integration — from silicon to system — that reduces debug and test time
- Reduced test software development time — generating test vectors is greatly simplified
- Reduced capital investment in test equipment
- Increased system fault coverage and reliability

SCOPE, our broad-based solution

To simplify and speed your design task, TI has developed its SCOPE™ (System Controllability/Observability Partitioning Environment) family. It is a coordinated, broad-choice set of commercial and military products compatible with the IEEE 1149.1 standard.

Included are bus-interface devices, standard cells, gate arrays, and digital signal processors, as well as our ASSET™ (Advanced Support System for Emulation and Test) diagnostics software.

On the way are diary memories, a series of IEEE 1149.1 stand-alone controllers, and microprocessors with boundary-scan and built-in self-test features.

We are in for the long haul

As a member of the Joint Test Action Group (JTAG), we contributed to the formulation of the IEEE 1149.1 standard and wholeheartedly support it. We are committed to growing our SCOPE family of products so that designing to the IEEE 1149.1 standard will be like second nature.

Your future competitiveness depends upon an engineering methodology where design teams bear the burden of testability, manufacturability, and reliability. The demands of concurrent engineering will be met in part by the extended capabilities accessed via the IEEE standard — from embedded system information that allows realtime availability of data throughout the design cycle to emulation and realtime system analyses capabilities built right into the silicon.

Get our floppy free, and learn more

Call 1-800-336-5236, ext. 3909, and we'll send you our unique floppy disk presentation. Just pop it into any MS-DOS™-compatible PC to find out more about DFT and TI's SCOPE testability family. What's more, the disk features a formula that allows you to calculate the cost-effectiveness of implementing testability in your system.

You will continue to read headlines about DFT. We intend to make many of them.

™ MegaChip, SCOPE, and ASSET are trademarks of Texas Instruments Incorporated.
MS-DOS is a trademark of Microsoft Corporation.
© 1990 TI 08-0205

TEXAS INSTRUMENTS

INTEGRATED CIRCUITS

MIPS rethinks RISC with superpipelining

Ron Wilson, Senior Editor

Phase one of the RISC revolution has run out of steam. The simple, single-CPU RISC architectures that started it all have wrung as much speed as they can out of existing CMOS processes, forcing vendors to look for other tricks to get more Mips. Some companies have turned to faster processes, like ECL or gallium-arsenide, to get higher clock rates. These efforts have for the most part been disappointing, both because the CPUs proved hard to fabricate and because the system designs at such elevated frequencies proved intractable.

Other vendors have attacked the CPU architecture itself. These vendors reasoned that, if a simple RISC CPU maxed out at one instruction per clock and if raising clock frequencies into the stratosphere was temporarily impractical, the best option was parallelism, or increasing the number of CPUs. The first steps in this direction were taken not by chip vendors, but by system designers, who created multi-CPU processor boards.

Intel followed close behind these initial efforts with its multiple-execution-unit 80960CA, and the term superscalar was born. The superscalar concept is simple: get several instructions out of the cache at once, and dispatch them simultaneously to several execution units. Use register scoreboarding to make sure that, even if the instructions finish out of sequence, data dependencies are respected.

In practice, superscalar proved difficult to implement. The dispatch and scoreboarding logic are so complex that no one has yet to do a completely general multiple-dispatch unit. Because the dispatch unit will only work efficiently with certain sequences of instructions, compiler developers have been left with a tremendous burden. In addition, the amount of silicon absorbed by the control logic and the several execution units leaves little room for an even-more-critical resource—on-chip cache.

Superpipeline approach

MIPS Computer (Sunnyvale, CA) has paved the way for the continuing evolution of RISC with a whole new set of ideas, embodied in the R4000 architecture. The company has developed a three-part manifesto against the general movement toward superscalar chips, focusing on multiple instruction issues, a true 64-bit architecture and integrated multiprocessing support.

On the first issue, MIPS has parted company with the superscalar advocates, introducing a concept it calls superpipelining. This idea is simple: by running the execution pipeline twice as fast as the instruction fetch logic, users can fetch two instructions at once and feed them both into the pipeline, half a cycle apart. The tricky part was implementing the idea in silicon.

To begin with, a 50-MHz R4000 would need a 100-MHz pipeline. Since the pipeline doesn't have to drive any external pins, such speeds were achievable with existing CMOS processes. "The circuitry necessary to let the pipeline take a new instruction every half cycle is fairly complicated. But at an architectural level, the design is much simpler than a superscalar approach," says MIPS vice-president of development Skip Stritter.

Additional complications came from data dependencies. What would happen, for instance, if the next instruction needed to use the results of the current instruction? The MIPS designers solved these issues not by scoreboarding, which would simply stall the pipeline until the key instruction finished, but by an elaborate set of bypass circuits. Thus, a following instruction can get at results while they are still in the pipe, before they're put back in the register file.

There are several advantages to the superpipeline approach. First, the dispatch logic is much simpler, and the single execution unit takes up less space than the

MIPS R4000 SUPERPIPELINE ARCHITECTURE

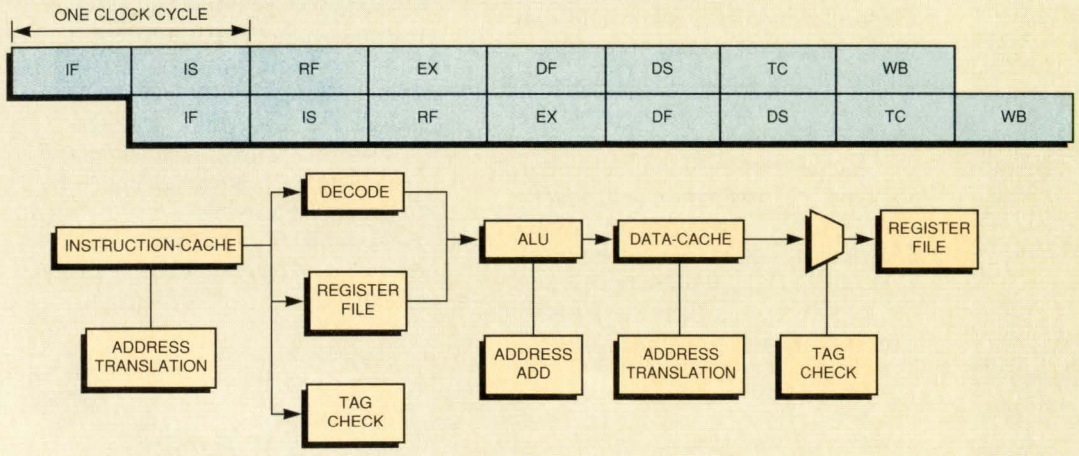

By running the execution pipe at twice the CPU cycle speed, MIPS is able to get two instructions into the execution unit on every cycle. The technique, called superpipelining, is demanding of circuit designers but results in a relatively simple multiple-dispatch architecture.

OS-9000. Bringing Economy of Design to Industrial 386 PC Developers

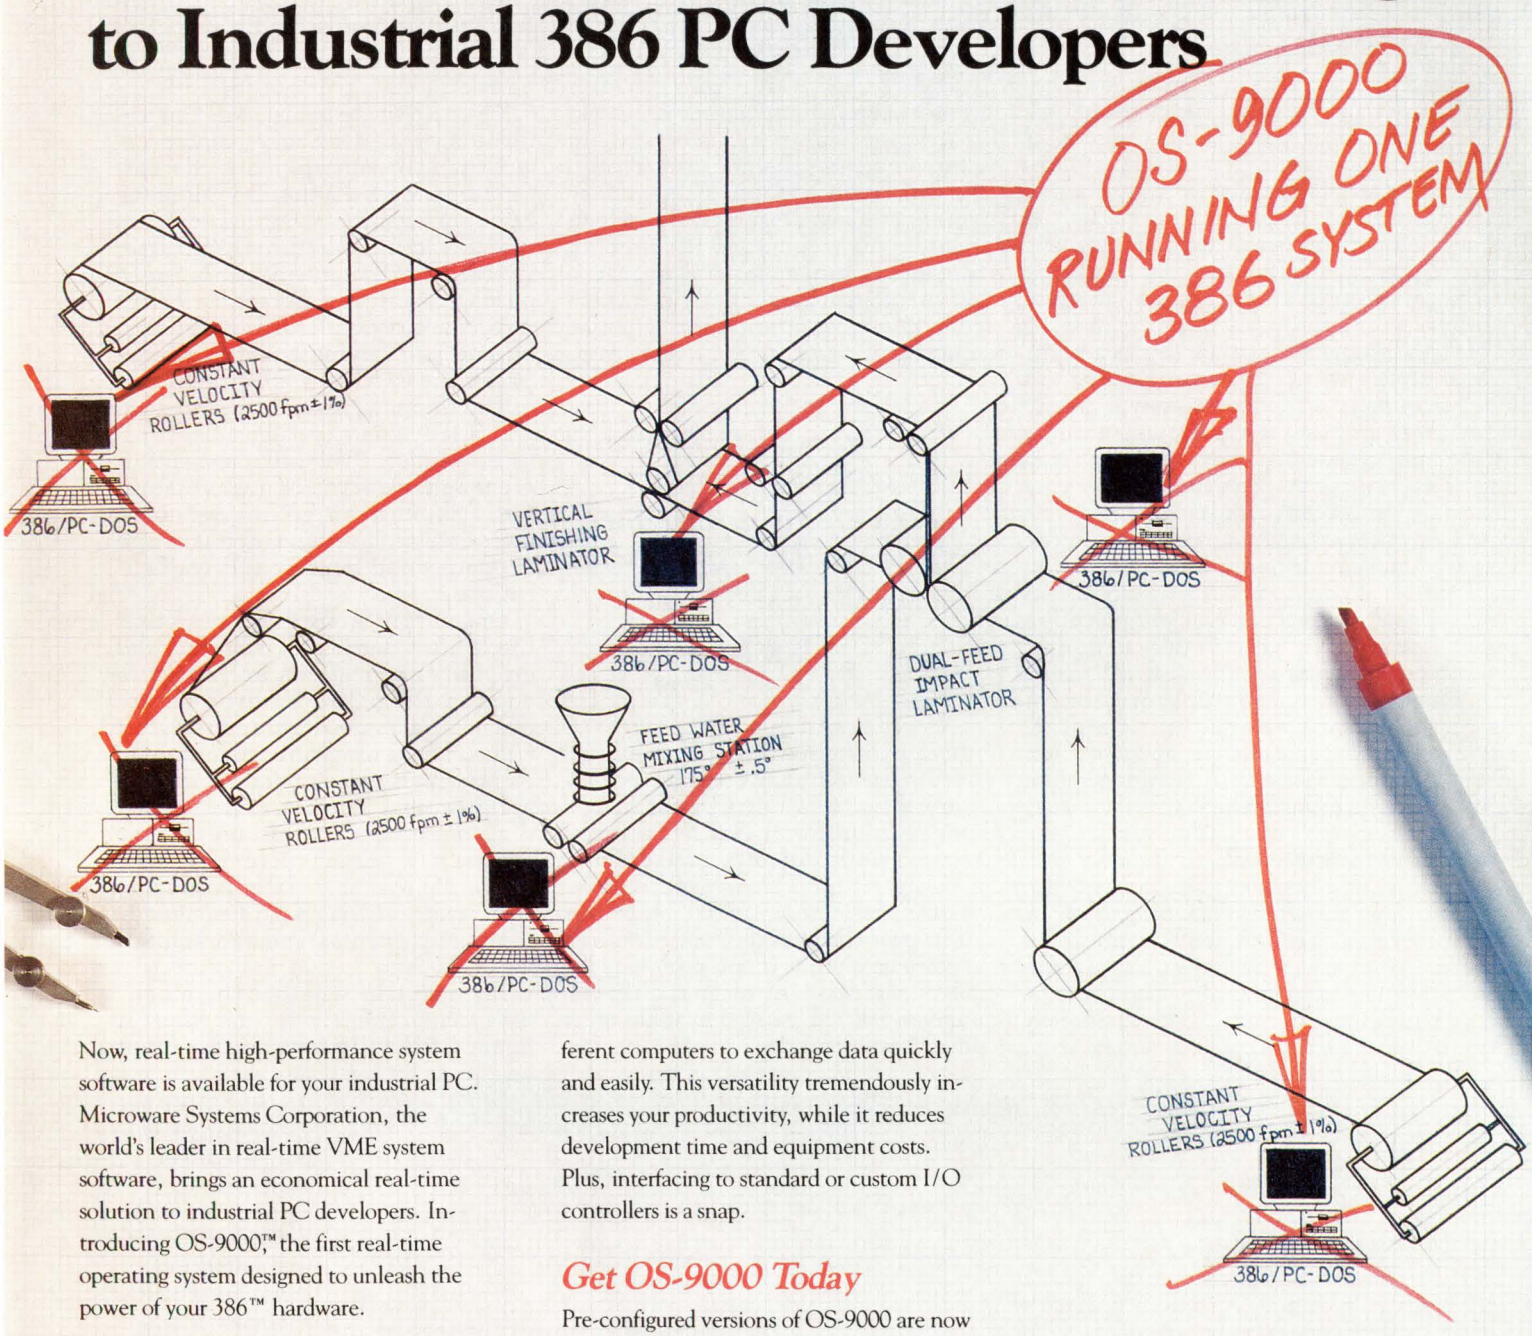

Now, real-time high-performance system software is available for your industrial PC. Microware Systems Corporation, the world's leader in real-time VME system software, brings an economical real-time solution to industrial PC developers. Introducing OS-9000,™ the first real-time operating system designed to unleash the power of your 386™ hardware.

Real-Time Response To Control Multiple Tasks

Why waste the power of your 386 by letting it control only one function at a time? Unlike PC-DOS, OS-9000 is a true multi-user and multi-tasking real-time operating system. Now you can utilize a single PC as a powerful time-sharing system able to control multiple processes simultaneously. And by using industry-standard protocols (like Ethernet and Arcnet), OS-9000 allows PCs to be interconnected and connected to different computers to exchange data quickly and easily. This versatility tremendously increases your productivity, while it reduces development time and equipment costs. Plus, interfacing to standard or custom I/O controllers is a snap.

Get OS-9000 Today

Pre-configured versions of OS-9000 are now available for your 80386™ PC-compatible hardware. These systems include integrated development tools such as a "shell" user interface, over 70 powerful utilities, a superlative C compiler, a C source-level debugger and a screen editor. Your full-featured, "plug-and-play" OS-9000/386 system is now available for just $995.

Call Microware® today to order your copy of OS-9000. Or order a FREE copy of the OS-9000 Catalog (your complete guide to the OS-9000 Real-Time Operating System).

Call Microware Today!

1-800-475-9000

In California, call (408) 980-0201

microware®

MICROWARE SYSTEMS CORPORATION

1900 N.W. 114th Street • Des Moines, Iowa 50322
Phone: (515) 224-1929 • Fax: (515) 224-1352

PC-DOS is a trademark of IBM. 80386 and 386 are trademarks of Intel Corporation. All other brand or product names are trademarks or registered trademarks of their respective holders.

MORE CHOICES • MORE OPTIONS • TOTAL SUPPORT

CIRCLE NO. 19

INTEGRATED CIRCUITS

multiple execution units of a superscalar machine, allowing more-complex circuitry to achieve the higher speeds. This makes more space available on the die for wide data paths and large caches, both of which are critical to the success of any multiple-dispatch architecture.

A second important advantage is the robustness of the approach. Superscalar architectures are notoriously sensitive to instruction order, and can drop well below one instruction per cycle on many code sequences. But the superpipeline architecture delivers a significant performance boost on existing R3000 binaries without reordering of the instructions, according to MIPS' claims. And users can get additional speed with the R4000 compiler, since the new compiler will know how to reorder multiple-cycle operations—integer multiplies and divides, for instance, and most floating-point codes—to avoid tripping the pipeline interlocks.

Less scalability?

On the minus side, superpipelining is theoretically less scalable than superscalar techniques. Running the pipeline twice as fast as the instruction cache is probably the limit for this generation of CMOS technology. So for now the approach will be stuck at around two instructions per clock.

A superscalar architecture, in contrast, could include an arbitrarily wide data bus and an arbitrary number of execution units, which would permit peak speeds of many instructions per clock. But practical considerations make this difference less significant than it might appear. In real silicon, the complexity of superscalar dispatch hardware increases so fast that more than two-instruction dispatch may be impractical. And as the number of simultaneous instructions goes up, the complexity of the compiler necessary to avoid perpetual stalls goes through the roof. As a final argument should these problems be solved, "We can always put multiple superpipelined execution units on a chip," says Stritter.

"At an architectural level, the design is much simpler than a superscalar approach."

—Skip Stritter, MIPS Computer

The second part of the MIPS manifesto may be even more controversial than the superpipeline concept: the company feels that it is time for high-performance microprocessors to move to full 64-bit architectures. By "full," MIPS means 64-bit data paths, 64-bit registers and 64-bit virtual addresses. This is a long step beyond just having a wide path between the CPU and the cache.

The company's reasoning is based on its previous experience as a systems vendor. "We have learned a great deal from the ECL R6000 system," claims Stritter, "not just about fast CPUs, but also about large memories and huge system configurations."

One of MIPS' observations is that the largest database, technical and graphics applications are about to outgrow the paltry 4-Gbyte flat-address space available on 32-bit CPUs. These programs, according to MIPS, must either undergo critical architectural changes in the way they organize data, or must have access to a larger virtual address range. Behind this claim is the fact that users of MIPS CPUs include companies such as Tandem and Silicon Graphics.

It's not clear whether the majority of customers who just want a faster CPU will be willing to accept

MIPS R4000 CHIP ARCHITECTURE

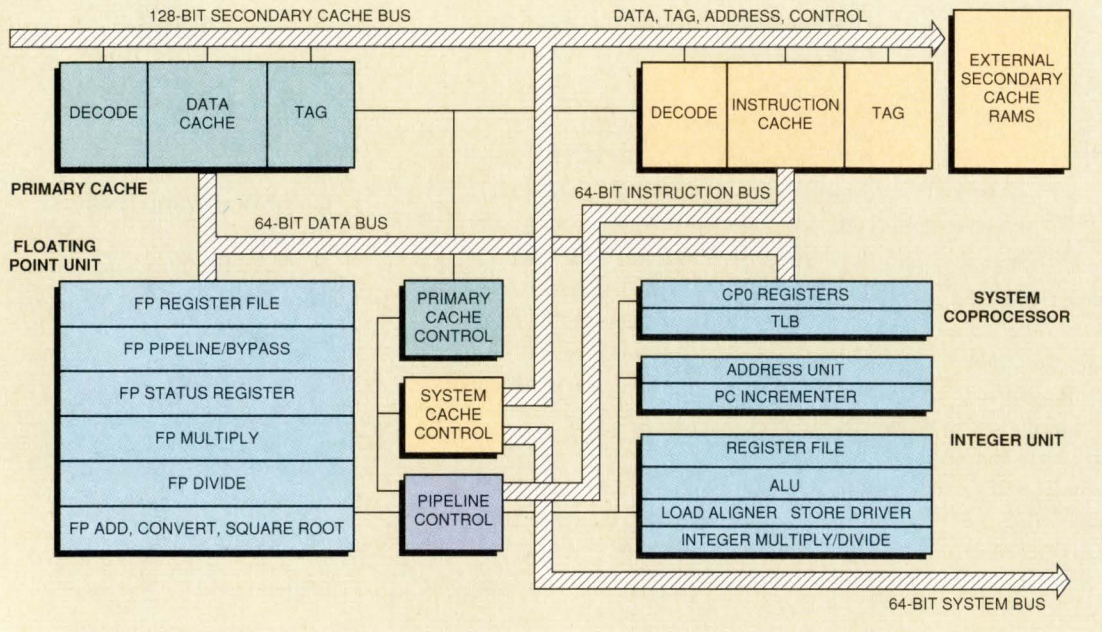

The first R4000 implementation will pack execution units, primary caches, secondary cache controller, and system bus controller onto a single die. In a departure from previous products the R4000 will have separate cache and system buses, which should ease many design headaches.

A 68040 for data, a 68020 for I/O... for real real-time performance on a single VME board.

Radstone's 68-41 Freeflow+ multiple microprocessor board with truly independent microprocessors for data and I/O gives you next generation VME performance...Now!

- 68040 with 16 Mbytes of dual-ported memory for maximum data throughput via concurrent, uninterrupted microprocessor operation up to 40 MHz
- 68020 with 4 Mbytes of dual-ported memory controlling extensive high performance on-board I/O facilities–all operating independently
- Multiple independent external buses–VME, VSB & APEX
- Multiple independent local buses–processor and I/O
- High performance DMAs
- Intelligent, high performance Ethernet and SCSI/SCSI-2
- ...and much, much more.

Radstone's Freeflow+ architecture takes VME to new performance levels. And now it's available with 040 processing punch. It's the very latest in Radstone's long line of leading edge commercial real-time VME board level products.

Extend your VME lead...and investment

For details on how to supercharge your VME system with Radstone's Freeflow+, and extend your current investment in VME hardware and software, call or write. Do it now, because your system is worth it!

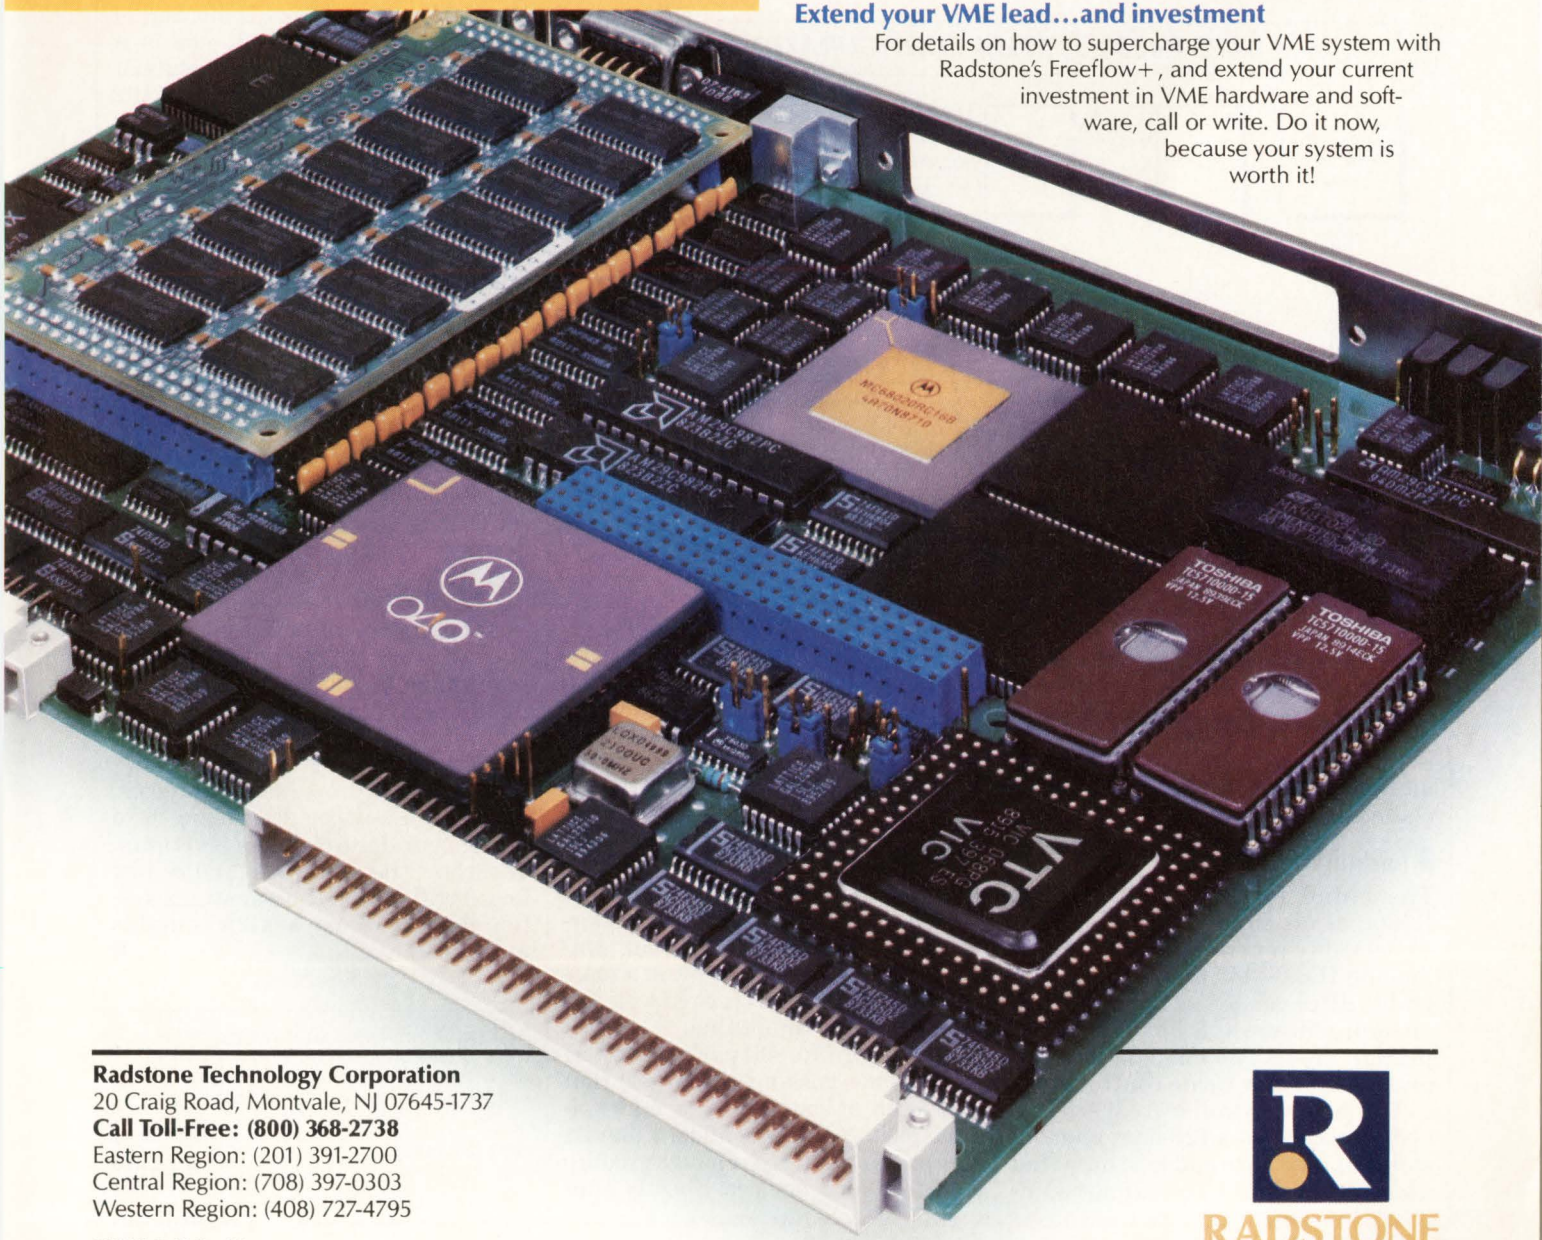

Radstone Technology Corporation
20 Craig Road, Montvale, NJ 07645-1737
Call Toll-Free: (800) 368-2738
Eastern Region: (201) 391-2700
Central Region: (708) 397-0303
Western Region: (408) 727-4795

CIRCLE NO. 20

RADSTONE
TECHNOLOGY

INTEGRATED CIRCUITS

the extra silicon burden of a full 64-bit design. The R4000 will run existing R3000 code without modifications, but it will have a complete additional set of instructions for 64-bit operation. Even if users don't need the 64-bit silicon, they have to buy the whole die.

Stritter argues that the actual overhead isn't that much. Once the large caches, bus control units, cache controllers, floating-point unit and associated hardware are packed onto the die, he says, it doesn't really make that much difference whether the integer ALU and register file are 32 bits or 64 bits. "If you look at a floor plan, you can see that the difference is a relatively small portion of the chip area—perhaps the equivalent of another 2 kbytes of cache," claims Stritter.

Packing the die

The initial implementation of the R4000 as disclosed by MIPS will be remarkably ambitious. The die will include the 64-bit integer unit and 64-bit FPU, an MMU for the 64-bit virtual addresses, 8 kbytes each of primary instruction and data cache, and a secondary cache controller.

The latter controller creates a secondary cache bus, a 128-data-bit monster that combines address, data, tag and control lines in a manner similar to that of the R3000 cache bus. The controller is highly configurable, allowing selection of access times, cache organization and size (up to 4 Mbytes is possible). The secondary cache can be either unified or split between instruction and data.

After listening to the problems system designers had in working with the cache bus on the R3000, MIPS decided to make a major change from the older architecture. The R4000 has a separate 64-bit system bus for connection to main memory, peripheral controllers and other R4000 CPUs, again with a highly configurable controller.

A host of features were included in the bus and cache controllers to solve the cache coherency problems that arise in multiprocessor configurations. Elaborate tag bits, snooping logic and primitives for virtually any sort of coherency protocol promise to make the R4000 compatible with almost anyone's notion of parallel architecture. In addition, control lines let the chip work in a master-slave relationship for fault-tolerant computing.

The initial R4000 package will be enormous—possibly the largest that has been attempted for a high-volume part. MIPS is confident that its silicon partners are ready to produce and test such devices. Shortly after it introduces the full-blown device, however, MIPS plans to roll out a reduced R4000 for more modest designs that will eliminate the secondary cache bus.

Next step: implementation

MIPS called its announcement this month an "architectural disclosure," not a chip release. The company had, as of the end of 1990, only taped out portions of the chip for testing purposes—no full R4000 die has been fabricated as of this writing. Many questions remain unanswered.

One is ac performance. MIPS is targeting a 50-MHz primary cache cycle time for the initial chip. Understandably, the power dissipation of the chip isn't stated yet, nor is it certain what the electrical behavior of a part that attempts to thrash 192 bus pins at that speed would be.

An added complication is that MIPS doesn't build its own silicon; the partners that supply its other CMOS CPUs will individually tape out and fabricate the parts. So a number of vendors will all be bringing the part up simultaneously on slightly different processes, probably with varying results. All of the devices, however, will be returned to MIPS for certification before they are declared to be R4000s.

When it does become available, the R4000 will represent a major rethinking of the evolution of RISC. The architecture represents a return to extreme circuit complexity on the die, but may actually make system implementation easier by reducing many of the critical timing requirements that characterized the R3000 bus.

By staking out new ground in 64-bit operation and hardware support for multiprocessing protocols, MIPS is asserting its vision as a major vendor of large systems. Whether the answers MIPS has found will match the questions system designers are asking remains to be seen. ■

MIPS R4000 SYSTEM CONFIGURATION

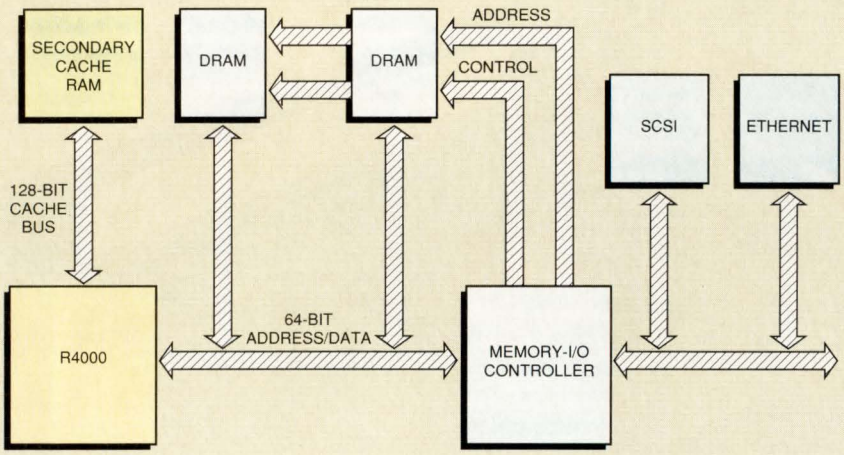

The full-blown configuration of the R4000 will attach a large secondary cache on one bus and main memory on a second bus. MIPS will introduce a simpler version of the chip for desktop applications that will eliminate the secondary cache bus.

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

MIPS Computer
(408) 720-1700 ..... **Circle 207**

If you're tired of pouring on SRAM to get 128Kbyte cache performance,

Take a load off.

Introducing an amazing new cache chip that gives crowded motherboards a new lease on life.

It's called 386™ Smart Cache, a family of chips that integrate a cache controller with up to 16Kbytes of SRAM onto a single, one-million-transistor chip. The result is a 16Kbyte chip that really cleans up, equalling the performance of 128Kbyte caches.

Intel's 386 and 386SX Smart Cache are one-chip solutions that give you extra space on your motherboards for other applications.

©1990 Intel Corporation. 386 is a trademark of Intel Corporation.

And when you stack them up against other caching solutions, you'll find their innovative architecture costs substantially less.

Give your motherboard a lift. Send in the attached card today for a free copy of our unique 386 Smart Cache video, performance report and documentation. Or if the card's been ripped off, call 1-800-548-4725, dept. HA47.

Because the reasons to look into 386 Smart Cache are really starting to pile up.

intel®

CIRCLE NO. 21

QuickFix™

The Embedded System Debugger

Sierra C™

The 68000 Family C Compiler

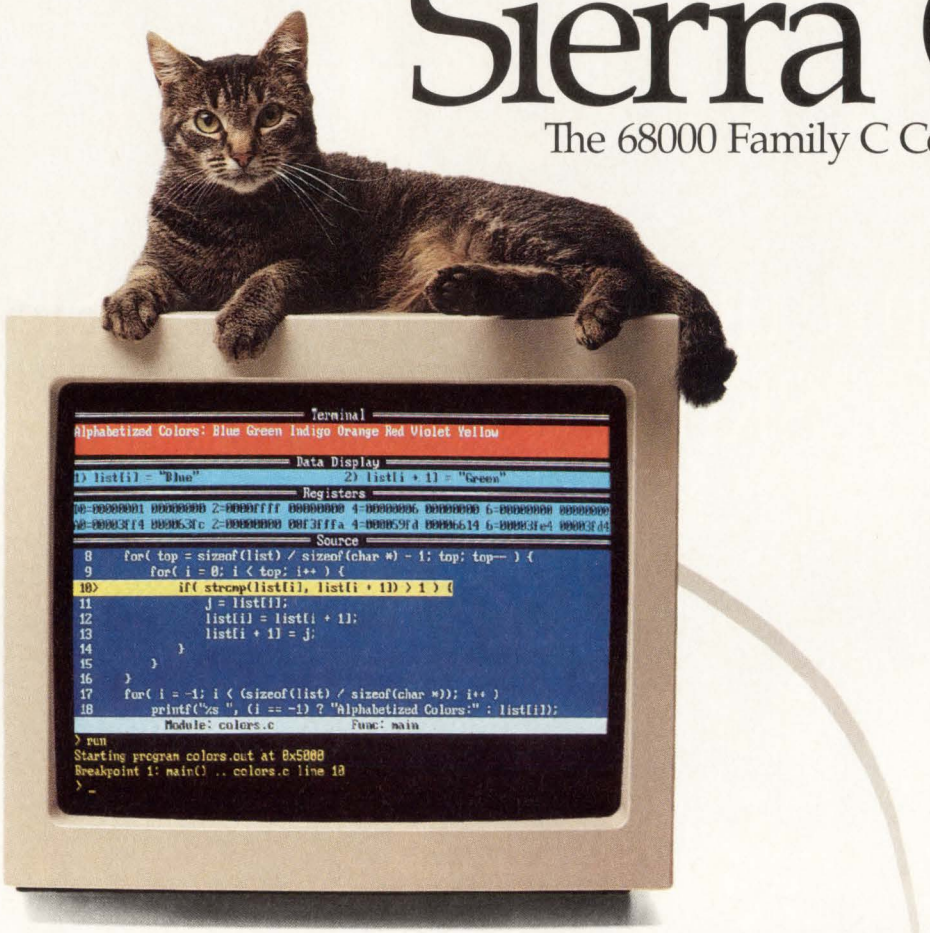

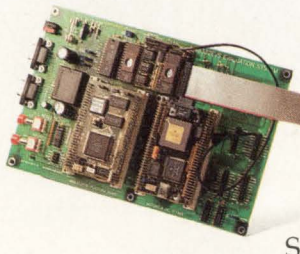

The Purrrfect Combination

Speed your products to market with the QuickFix Debugger and Sierra C Compiler, the perfect combination for all your 68000 family embedded applications.

The QuickFix Debugger is the fastest and most easy-to-use debugger available today. Besides the convenience of the multi-window display, commands feature a C-like syntax for extremely intuitive operation. Plus, a quick-connecting ROM communication cable and preconfigured drivers cut installation time from possibly weeks down to *minutes*!

The Sierra C Compiler generates the tightest code available today at blazing compilation rates. And it's not only fast — it's complete. Along with Sierra C's optimizing compiler, macro assembler and linker, you get a unique high speed parallel download system. Sierra C comforms to the ANSI C standard. Additionally, position independent, re-entrant and ROMable code can be generated to fit any memory configuration.

For more information on the tools that provide the speed, power and ease of use necessary to get better products out faster, call **800-776-4888**

SierraSystems

6728 Evergreen Avenue, Oakland, California 94611
Tel.(415) 339-8200 Fax (415) 339-3844

CIRCLE NO. 22

SOFTWARE

Japan's Tron initiative may be opportunity for U.S. manufacturers

Tom Williams, Senior Editor

Words often have amazing power to shape our perceptions and prejudices. Consider a vast Japanese undertaking called the TRON project. Since TRON stands for "the real-time operating system nucleus," most people are inclined to assume it's just another real-time executive. But TRON is much more than that, and it's beginning to take on real momentum in Japan.

TRON encompasses a real-time operating system that's the basis for higher-level operating systems for workstations, network servers and network management. It's also the basis for an instruction-set specification that has resulted in several microprocessor designs from Japanese manufacturers. Finally, TRON is the impetus behind a networking scheme to connect literally a billion or more processors in a highly functional distributed system (HFDS).

The brainchild of professor Ken Sakamura of the University of Tokyo, the TRON project is guided by a vision of a future computerized society in which computer-based "intelligent objects" are linked together via the HFDS. Such objects will include all kinds of microprocessor-based consumer products—TVs, VCRs and household appliances, for example. But intelligent objects will also include personal workstations, industrial robots and network databases. Nothing less than reorganizing the entire computer world—from microprocessor to operating system to user interface to network—is the goal of the TRON project.

Although that may seem a grandiose goal, it appears to be finding increasing resonance, at least inside Japan. Recently a group of 19 companies put up $17 million to build the TRON-concept Intelligent House, a prototype to test the TRON intelligent-object and networking concepts and a kind of laboratory to work out some of the mega-system issues inherent in networking together some 1,000 computers with their attendant sensors and actuators.

For instance, the house contains subsystems for security, lighting, entertainment, and temperature control. Ways have been worked out for different subsystems to negotiate with each other if their needs conflict, so that their functions can complement one another. So selecting a given mood for the lighting will make information available to the entertainment system, which will then select—or suggest—appropriate music. When the phone rings, the music is automatically turned down in the vicinity of the phone being picked up. When the climate control system decides to close the curtains, it won't do so if sensors tell it that someone is standing looking out the window.

More ambitious research projects are in the works: an intelligent office building designed to link up over a

TRON NETWORK ARCHITECTURE

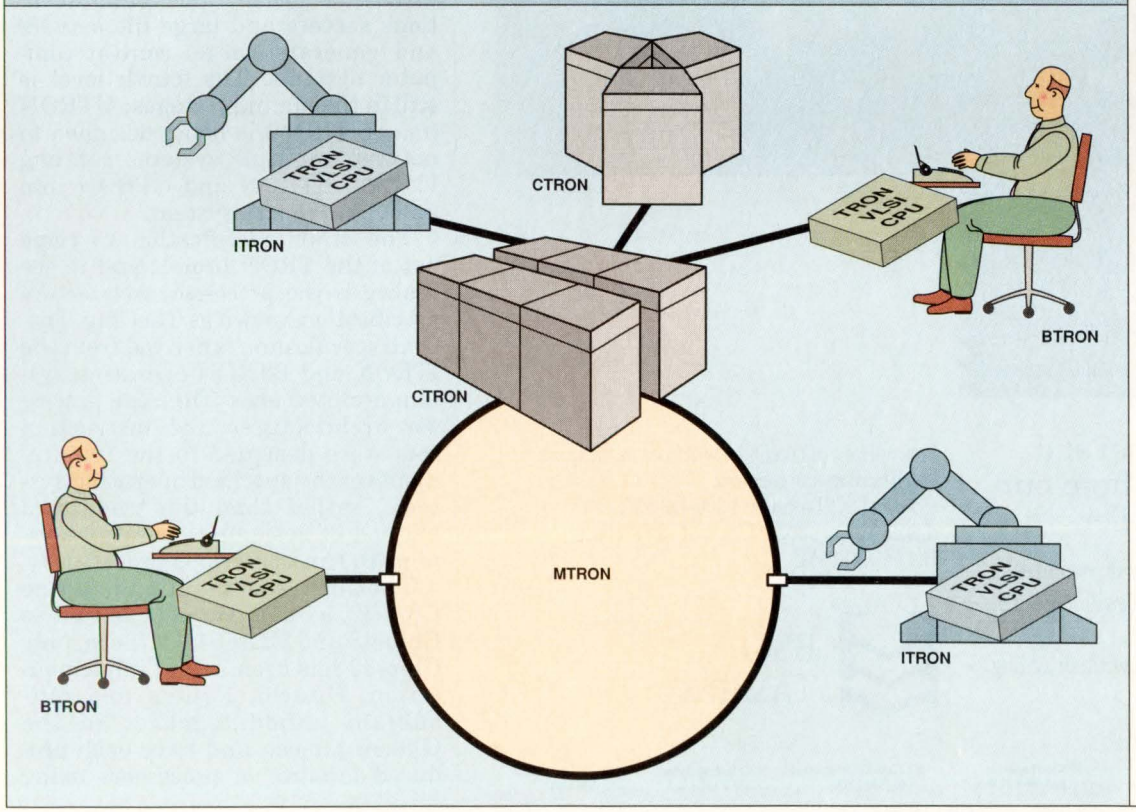

The universe of TRON ranges from the VLSI processor architecture to the ITRON and BTRON operating systems for embedded systems and personal computers. CTRON applies a standard architecture to communication, switching and information processing. MTRON will link systems based on the first three operating systems, allowing them to work in mutual cooperation. In addition, a media-independent data protocol and format support the wide-ranging network scheme envisioned by TRON.

SOFTWARE

million processors, and an intelligent city, to be built in Chiba Prefecture in the coming years. Such commitment to research reveals that a significant segment of Japanese industry is determined to make the TRON concept a reality.

TRON is first and foremost a hierarchy of specifications. The specifications set rules for interfaces but don't assume any specific hardware or software code. The hierarchy starts with the microprocessor instruction set, the operating system kernel (nucleus), operating system outer layers and applications. Since TRON specifies the interfaces between layers, companies have a great deal of leeway for different implementations of different layers. Thus functions are specified but not performance, leaving an open field for competition.

The many layers of TRON

TRON operating systems are specified at four basic levels. ITRON (industrial TRON) is a multitasking, real-time operating system for embedded systems. Versions have been specified for 8-, 16-, and 32-bit microprocessors, and implementations exist for the Intel iAPX86 series, Motorola 68000 and other microprocessors, in addition to the Gmicro (global microcomputer) microprocessors that support the TRON-specific instruction set.

The second level is intended for personal computers and workstations. BTRON (business TRON) has a file structure based on a network file system and is hierarchical, so it can be implemented at various levels for support of, say, multimedia or multilingual applications.

The third level, CTRON (central TRON), is specified for communications servers and large file servers and generally large, central computer systems. The fourth level is still in the planning stages; MTRON (macro TRON) is being designed to network multiple systems running ITRON, BTRON and CTRON into an overall macro system.

The other specification to come out of the TRON project and to see reality is the processor architecture specification known as TheChip. TheChip specification is derived from the ITRON and BTRON operating system architectures. Thus the processor architectures and instruction sets were designed to the requirements of the specified operating systems, rather than the traditional method of designing operating systems to run on various processors.

TheChip specifications define Chip-32, a 32-bit version, as well as Chip-48 and Chip-64. Of these, only Chip-32 has seen actual implementation. Hitachi, Fujitsu and Mitsubishi joined together in the Gmicro project, and have each produced families of processors using

Get more room and performance from our 8051 C compiler.

> Fit 30% more code in memory
> Compress your execution speed by a factor of 10
> Add more features without adding more memory

Find out how our compiler will find you the space you'll need in your memory.
Contact BSO/Tasking at **1-800-458-8276**.
(Inside MA. or outside U.S.A. 617-894-7800).
Fax 617-894-0551.

BSO TASKING

BSO/Tasking US 1-800-458-8276 | BSO/Tasking UK +44 252 10014 | BSO/Tasking Italia +39 2 6698 2207 | BSO/Tasking Deutschland +49 7152 22090 | BSO/Tasking France +33 1 30 54 222 | Tasking BV Netherlands +31 33 558584

CIRCLE NO. 23

the 32-bit TRON architecture, along with peripheral chips including floating-point units and memory-management units. The processors are no slouches at performance, either. The Fujitsu F32/300, for example, boasts 17 Mips at 25 MHz.

Although TRON hasn't been universally embraced by Japanese industry, it has gained some impressive support. The latest example is the Japanese Ministry of Education's decision to standardize (starting in 1992) on the BTRON operating system for all computers placed in Japanese junior and senior high schools. A key feature of TRON that may smooth its path is that it is open. All results of the project, in the

A significant segment of Japanese industry is determined to make the TRON concept a reality.

form of published specifications, are made available to the public domain. Specifications are also published in English, and anyone is free to develop and market products implementing the specifications.

U.S. firms slow to act

Several U.S. companies are members of the TRON Association, but actual U.S. activity in TRON has been minimal. According to Sakamura, a number of U.S. software companies are active—including Ready Systems, Wind River Systems and Green Hills. "Unfortunately, the systems makers and application makers aren't doing anything," he says. "Wind River and Ready Systems may have TRON capability, but their customers aren't using it."

What Ready Systems (Sunnyvale, CA) and Wind River (Alameda, CA) have done is to port their real-time kernels—VRTX32 and VxWorks, respectively—to the Gmicro architecture. This doesn't in itself yield any TRON capability; it merely allows VRTX32 and VxWorks applications to run on Gmicro processors. Both companies obviously feel that their real-time kernels are superior to what could be done using ITRON. Still, it's possible to implement ITRON calls as a shell over VRTX32 and get VRTX32 performance with ITRON compatibility, according to Ready Systems' senior software engineer Elie Grouchko. That would presumably hold true for other processors running VRTX32 and VxWorks as well—and be an easy way to implement high-performance ITRON.

Sakamura feels that the Japanese Ministry of Education decision rep-

PERFORMANCE + PORTABILITY ⇨ PRODUCTIVITY

You Only Have to Write it Once

Oasys software tools save you time and money. Port your applications quickly without recoding. Our Green Hills' compilers ensure source code compatibility across all major hosts, targets, and operating systems. Portability is built in with complete compliance to industry standards:

- ANSI C and K&R C
- ISO Level 1 Pascal
- Support for IEEE 695+ ICE's
- C++ - OOP plus optimizations!
- FORTRAN-77 (VAX VMS extensions)
- Interface to major RTOS environments

Optimizing Your Code is Key

Develop more efficient, high quality software that maximizes your hardware performance, without writing assembly code. Take advantage of 50+ advanced optimizations implemented in Green Hills' compilers. Let Oasys' high performance tools optimize your code to produce the fastest, tightest executable.

Oasys Integrated 680x0 & 88x00 Tool Kits

Green Hills Native & Cross Compilers (C, C++, Pascal, FORTRAN)
Oasys Cross Assembler/Linker System
Oasys Simulator
Oasys UDB™- Universal Remote Debugger

Other Tool Kits: 386™, i860™, TRON™, Microsoft Cross C™

Availability:

- Sun-3, Sun-4 SPARC™
- DEC (VMS/Ultrix)™
- IBM RT™ and RISC System/6000™
- HP/Apollo™
- More! Call for the complete list

Oasys

One Cranberry Hill • Lexington, MA 02173
FAX: (617) 863-2633 or (617) 862-2427

Tomorrow's Solutions Today

386 and i860 are trademarks of Intel Corp. UDB is a trademark of CaseTools Inc., Microsoft Cross C is a trademark of Microsoft Corp. All other trademarks are acknowledged to their respective manufacturers.

Worldwide Support: Belgium Retime 32-2-376-5146, **France** Real Time Software 33-1-6986-1958, **Germany** Stollman 49+40-89088-0, **Holland** Computing & Systems Consultant 31-40-434-957, **Israel** Ankor Computers Ltd. 972-3-5447356, **Italy** Instrumatic 39-2-353-8041, **Japan** MCM Japan Ltd 81-3-487-8477, **Spain** CIDISA 34-1-563-3649, **Scandinavia** Traco AB 46-893-0000, **Switzerland** ZuhlkeEngineering AG 41-1-730-7055, **UK** Real Time Products 44-21-236-8070

For Real Time Response,
Call (617) 862-2002
CIRCLE NO. 24

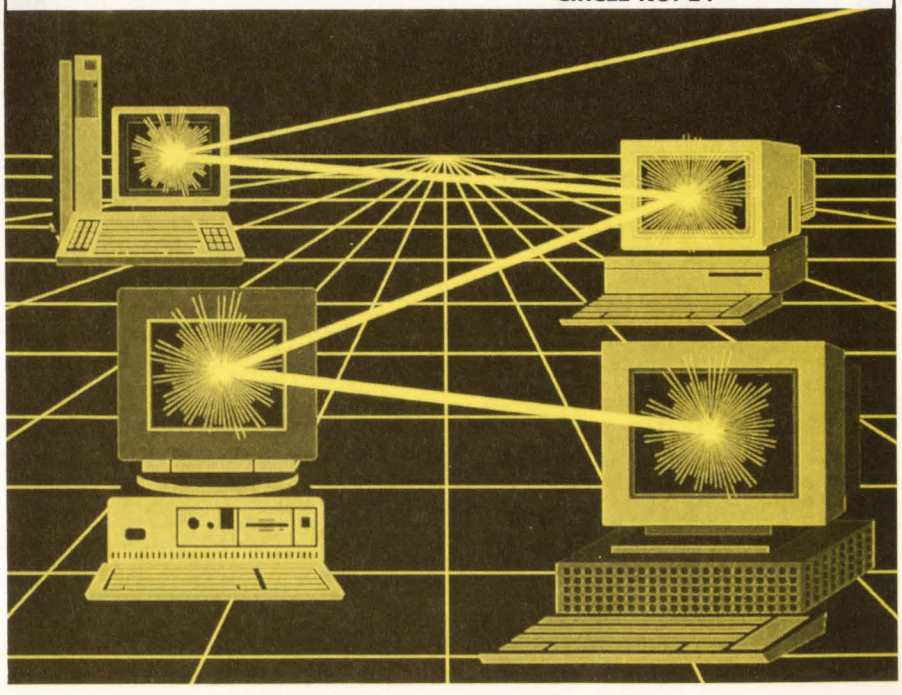

THE PERFECT STRATEGY FOR ANALOG & DIGITAL PCB DESIGN

- YOU want powerful, automated analog design features — ULTIboard's new Version 4.1 let's you use curved or any angle traces, Also, you specify automatic fill patterns for copper areas of any shape. Polygons are assigned to nets while traces of other nets are automatically ploughed through filled areas.
- YOU use power planes — With ULTIboard, power and ground planes with thermal relief pads are created automatically.
- YOU want lots of other automated features — Real-time DRC, trace shove, reroute-while-move, round tracing and board outlines, .DXF output, automatic component renumbering, and more are standard in all ULTIboard systems.
- YOU want an affordable upgrade path — ULTIboard's got a smooth upgrade path with a system in every price range, running under DOS or extended memory. Choose a system to design a small board of 50 equivalent IC's, up to our 32-bit 386 version with unlimited design capabilities.
- And don't forget schematic capture... — ULTIcap schematic capture featuring automatic guided wiring, automatic junctions, auto rewire after block moves, and on-line symbol creation are just some of ULTIcap's easy-to-use automated features. Or, use your own schematic package with ULTIboard's backannotation.
- YOU want to find out more — Call ULTImate Technology today for a **free demo disk**, and the name of the dealer nearest you.

*ULTIboard Professional 386 is a **true 32 bit** Design System; approx. 2x faster than 16 bit versions.*

Select your ULTIboard system from this table:	Model	Price	Nr. of 16 pins equiv. IC's	Memory required
	Entry Level	$ 995	35 - 50	640 K
	Advanced Level	$ 2.475	140 - 200	640 K + 1 Mb
	Professional 286	$ 5.950	unlimited	640 K + 2 Mb+
	Professional 386	$ 7.900	unlimited	640 K + 2 Mb+

USA/Canada Headquarters
ULTImate Technology Corp.
269 Mt. Hermon Road,
Suite 105
Scotts Valley, CA 95066
Tel. (408) 439-8944
Fax (408) 439-8966

UK/Ireland Sales-Office
ULTImate Technology Ltd.
1 Mars House Calleva Park
Aldermaston
Berkshire RG7 4QW
Tel. (0734) 812030
Fax (0734) 815323

Germany Sales-Office
ULTImate Technology
Vaihinger Straße 12
D-7032 Sindelfingen
Tel. (07031) 876028
Fax (07031) 85915

Europe Intl. Headquarters
ULTImate Technology bv
Energiestraat 36
1411 AT Naarden, Netherlands
Tel. (+31) 2159-44424
Fax (+31) 2159-43345
Telex. 73415

CIRCLE NO. 77

Trying to find the right ICs?

Turn to the Integrated Circuits section in New Product Highlights. You might find just what you need.

SOFTWARE

resents an opportunity for U.S. manufacturers to sell into the Japanese market. While some Japanese manufacturers have already implemented a BTRON operating system, there's nothing to stop U.S. manufacturers from doing so. "We have the specs in English, and U. S. software companies can easily get the specs and make an operating system," says Sakamura.

The U.S. Commerce Department, on the other hand, has denounced the Ministry of Education decision as a plot by the Japanese to lock up the market. "The U.S. government completely misunderstands our project," Sakamura says, indicating that the main reason was language. "Please remember, I think Unix and MS-DOS are good operating systems, but the Japanese language is complex and Japanese pupils and teachers wanted an operating system that could handle Japanese."

BTRON includes facilities to automatically translate English into Japanese, so English-speaking developers could do a great deal of their user interface implementations in English, and possibly have the final versions checked by native Japanese speakers. In addition, the translation facility is able to translate—legally, of course—English applications and automatically format the screen and menus to accommodate the resulting Japanese version.

The fact that BTRON is, as Sakamura says, "only a specification and not real code" and that it can be implemented at many different levels would appear to offer an opportunity to enterprising U.S. software companies. If the United States is so far ahead in software technology, it should be possible to create BTRON-based high-performance operating systems that would be strong competitors in the Japanese market, where, Sakamura assures us, they will be welcome. ■

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

Ready Systems
(408) 736-2600 ..... **Circle 237**
TRON Associates
(602) 249-3388 ..... **Circle 238**
Wind River Systems
(415) 428-2623 ..... **Circle 239**

16 lines, No waiting.

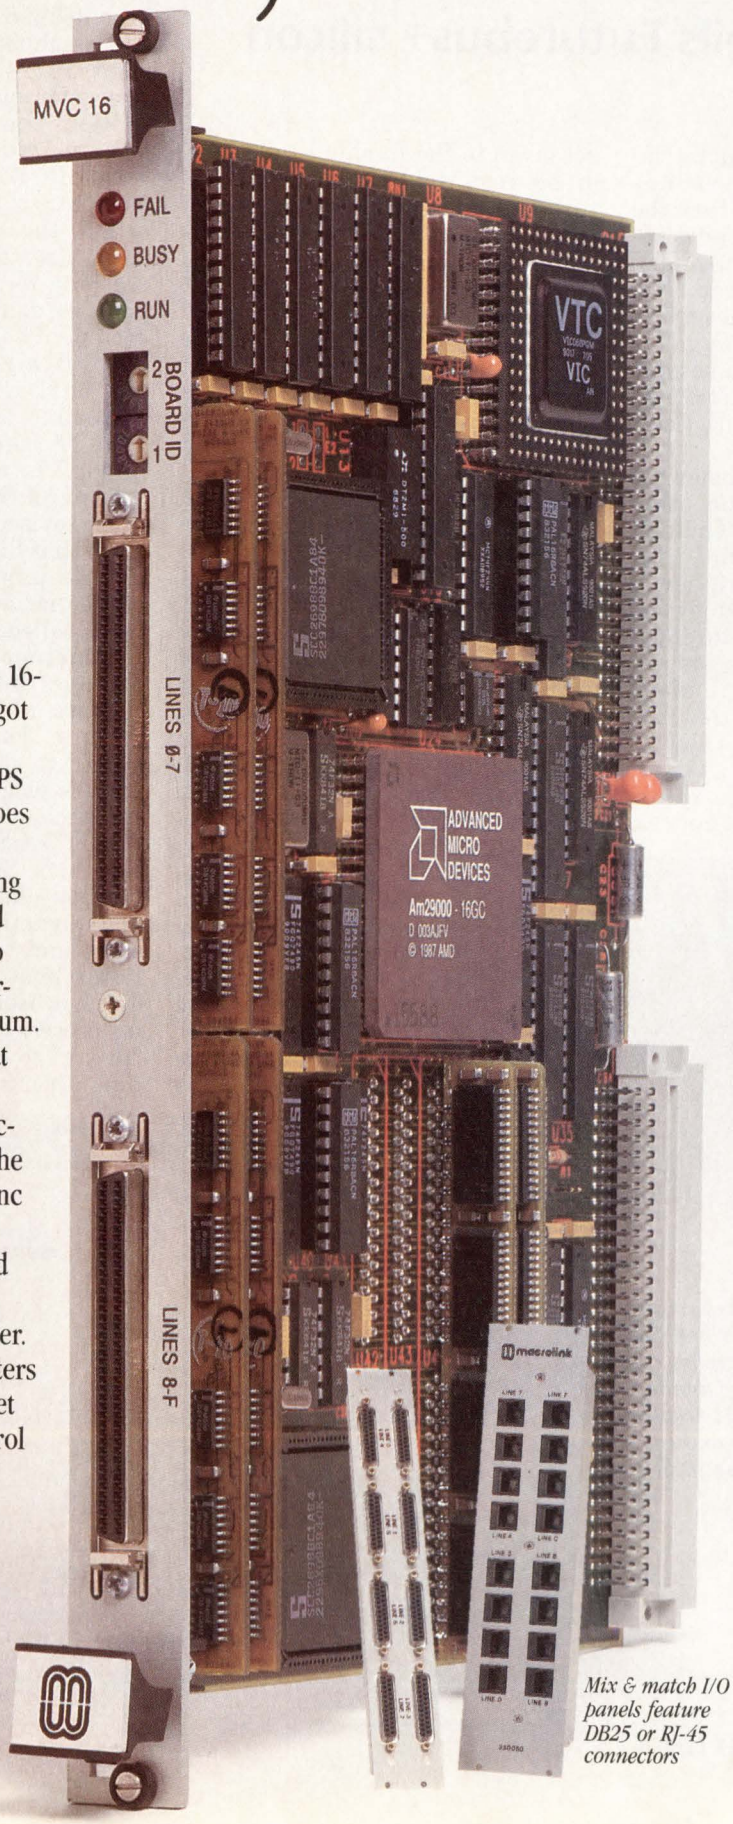

Mix & match I/O panels feature DB25 or RJ-45 connectors

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. 26

COMPUTERS AND SUBSYSTEMS

National unveils Futurebus+ silicon

Warren Andrews, Senior Editor

The recent announcement by National Semiconductor (Sunnyvale, CA) brings to three the number of major semiconductor vendors that have announced Futurebus+ interface silicon. National joins Philips-Signetics (Sunnyvale, CA) and Texas Instruments (Houston, TX), which has a second-source agreement with Signetics and a joint design agreement with Force Computers (Campbell, CA).

National's approach is somewhat different from that of the other firms, according to Brian Gillings, marketing director for the company's Interface and Peripherals Group. "First, we reserved formal announcement of our Futurebus+ products until the specification settled down enough so that major changes wouldn't be required," he says. "Second, we held back the announcement until our products were in volume production. And third," Gillings continues, "we kept our chip set down to the fundamental ingredients required to implement a Futurebus+ system, leaving many of the higher-level protocol functions to be implementation- and processor-specific."

The first five chips to come out of National's Futurebus+ program are a 9-bit data transceiver, a 9-bit latched data transceiver, a handshake transceiver, a 9-bit distributed arbitration transceiver and an arbitration controller. When combined with a protocol controller, these chips constitute a complete Futurebus+ solution.

BTL transceivers

Futurebus+ uses BTL (backplane transceiver logic) instead of the more-familiar TTL to speed transfers across the backplane. Initially developed by National, BTL defines a low-voltage-swing logic with an exceptionally fast rise time. Though various incarnations of BTL are being manufactured by National, TI and Signetics, there has been no standard defining what all the parameters are and how the logic should behave in a system. Now, however, parameters for the logic have been defined by an IEEE committee and established as an IEEE document, P1194.1. This document is in turn referenced by the P896.2 specification, the Futurebus+ physical-layer definition.

Not unlike the balance of the various parts of the Futurebus+ specification, the P1194.1 document has seen several iterations prior to reaching its present state. "At this point we don't see any change in the document before it reaches final approval," says Gillings. He expects the final IEEE blessing to come sometime around mid spring of 1991. "According to the present definition," he adds, "National has the only parts that completely conform to the revised specification." Gillings hints that Signetics may have to revise its process somewhat to have parts that conform to the fast rise times called for in the specification.

Sometime about mid last year, Signetics announced its Futurebus+ chip set, which includes BTL transceivers. Subsequently, Signetics agreed to share its Qubic BiCMOS process technology with TI in a design/technology agreement. This is the BiCMOS semiconductor process by which Signetics makes its Futurebus+ family of chips. Then TI and Signetics agreed to swap designs and technology on all future Futurebus+ products.

It has been suggested that Signetics' process may require only a small "tweak" to have parts up and running according to the revised speci-

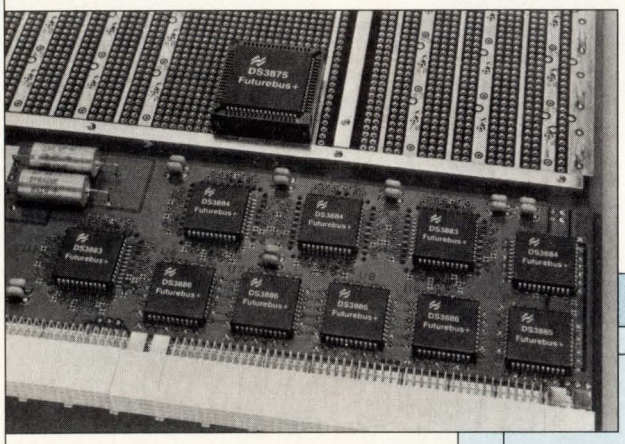

National Semiconductor's Futurebus+ chip set includes five chips that, when combined with a protocol controller, constitute a complete Futurebus+ interface. The set's implementation of backplane transceiver logic conforms to the IEEE P1194.1 specification.

NATIONAL'S FB+ CHIP SET

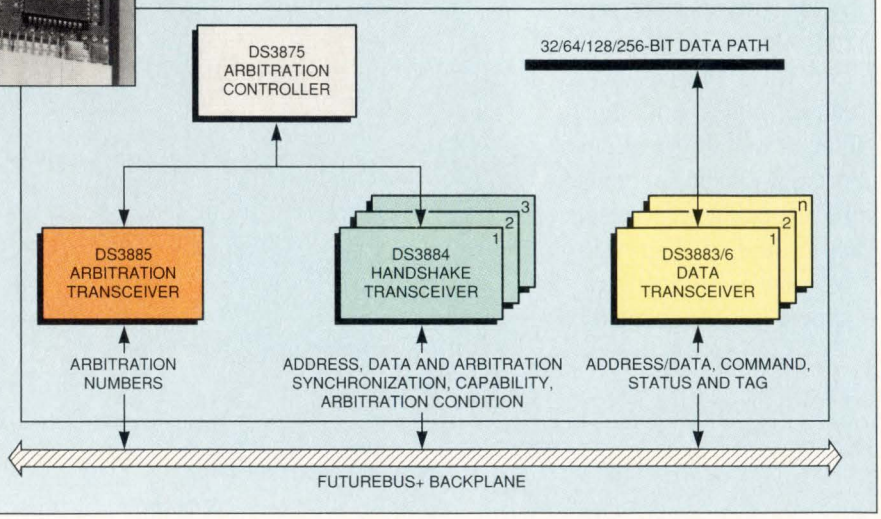

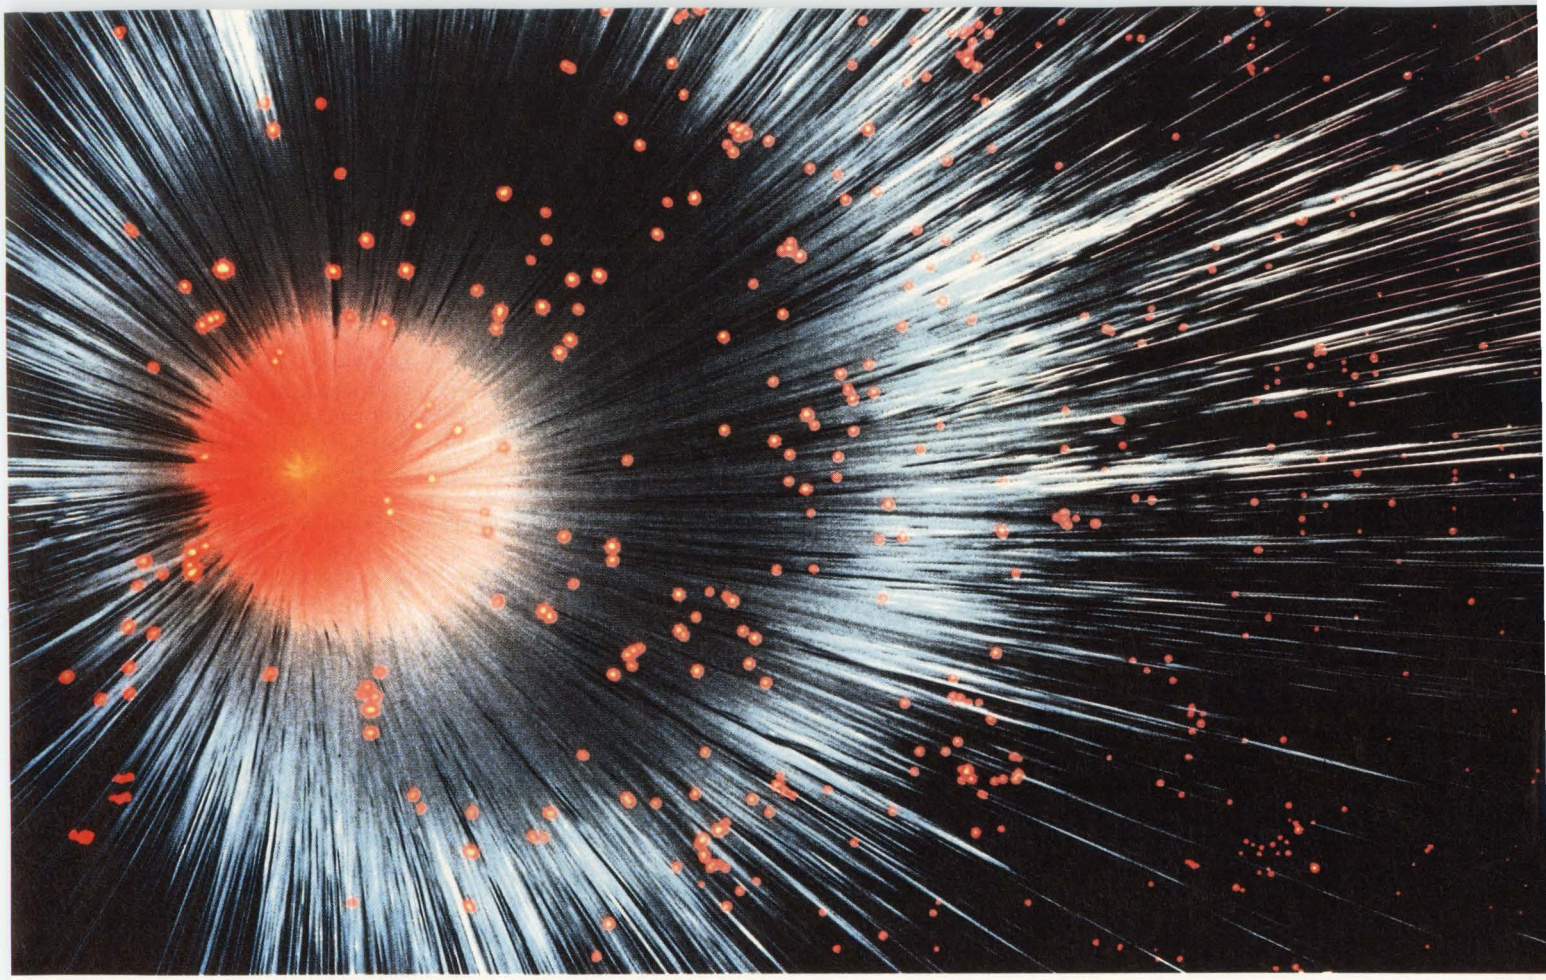

Speed + Performance . . .
The SuperCard™ Family of i860 Vector Processors

The second generation of i860-based SuperCards™ is here... vector processors with GigaFLOPS system capability, concurrent I/O via secondary buses, large math libraries, FORTRAN and C development software... and, best of all, the *best price* on the market. The SuperCard family is compatible with VMEbus, ATbus, EISAbus and TURBOchannel platforms for hosts like Sun, HP, Motorola, Compaq, and DEC... here's what you get:

- 1024 complex FFT in 0.8 msec.
- 80 MFLOPS to 1.5 GFLOPS
- 160 MB/sec I/O
- Large memories
- pSOS+ Multiprocessing
- CASE Tools

CSPI

To find out how SuperCard can work for you, call: 1-800-325-3110 or 617-272-6020, (Fax 508-663-0150) or write CSPI, 40 Linnell Circle, Billerica, MA 01821.

SuperCard is a trademark of CSP Inc., i860 is a trademark of Intel Corp.
TURBOchannel is a trademark of Digital Equipment Corp.
pSOS+ is a trademark of Software Components Group, Inc.

CIRCLE NO. 27

COMPUTERS AND SUBSYSTEMS

fication. However, given the internal turmoil at Signetics—and that of its parent, Philips—there's no indication of when that may happen, nor has TI indicated when it expects to have the Signetics process up and running. And there's some doubt as to whether TI will have the same problems Signetics had.

Three transceiver flavors

National is offering three flavors of its Futurebus+ transceivers, which are manufactured in two semiconductor technologies. The plain-vanilla variety, the 9-bit (one byte with parity) chip, is fabricated in straight bipolar. It provides a straightforward transceiver function to handle address, data, command, status and tag signals. The 44-pin chip is liberally sprinkled with additional ground and Vcc pins to reduce the effective inductance of bonding wires and leads, thus reducing noise from transients in the ground path.

In addition, the chip supports live insertion, has a controlled rise/fall time of 2 to 5 ns, has a built-in band-gap reference to provide accurate thresholds and has pinouts specifically designed for Futurebus+ applications. An option to the 9-bit data transceiver provides latched outputs. But the additional logic on-chip calls for the latched data transceiver—as well as the rest of the transceiver chips in the family—to be fabricated in National's BiCMOS process.

"National has the only parts that completely conform to the current P1194.1 definition."

—Brian Gillings, National Semiconductor

The handshake transceiver, designed for address, data and arbitration synchronization signals, is similar in function to the data transceiver but is a 6-bit device with three of the transceiver outputs providing a parallel wired-OR glitch filter output. A pair of pulse-selectable pins on the chip provide four settings for the glitch filters. Settings are established such that the worst-case pulse width for the wired-OR glitch is equal to the round-trip delay of the bus. This round-trip delay will depend on the length of the bus—the longer the bus, the longer the delay. The chip lets designers optimize the filters to the round-trip delay of a given backplane.

The arbitration transceiver is designed to support National's arbitra-

All you need to know about 5.25" Winchester disk drives, inside

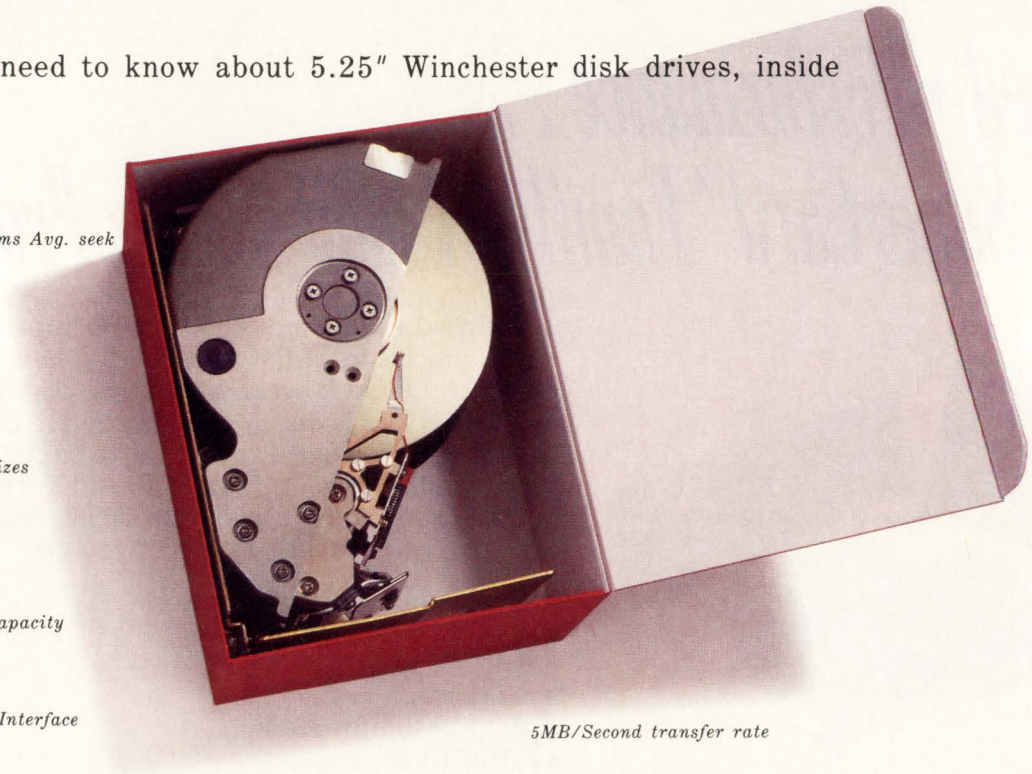

14ms Avg. seek

Variable sector sizes

1.6GB Capacity

SCSI-2 Interface

5MB/Second transfer rate

C&C Computers and Communications

It's nice to know that NEC disk drives have the most advanced technical features. And it's reassuring that they're consistently available, and with a DOA rate of less than 1%,and up to 100,000 hours MTBF rate that they're reliable.

tion controller and reduces additional on-board logic and possible delays by including the competition logic required in handling arbitration numbers. In addition to transmitting and receiving arbitration information signals between the bus and the controller, it implements a parity check to validate the information it receives, and generates status information to indicate the current bus conditions.

The arbitration controller—the only chip fabricated in straight CMOS—regulates access to the bus as specified in the acquisition, allocation and alignment protocol of the Futurebus+ standard. It's software-configurable to support unrestricted-mode, mixed-mode and one-pass operation. In addition, it allows for programmable arbitration timing delay, and, like the other chips in National's Futurebus+ family, it supports live insertion. Other features include built-in timers for all protocol-defined timeouts; dual-port, bidirectional data accessibility for arbitration number and messages; a hardwired register for emergency messages; and an on-board parity generator.

Protocol controller

Unlike its two major competitors in the Futurebus+ interface silicon business, National isn't offering a protocol controller, at least at this time. "We're recommending that designers implement the protocol control function with an ASIC, or that they use programmable logic to tailor it to the microprocessor, local bus and system architecture selected for any specific application," says Gillings.

Both Signetics and TI have announced protocol controllers, the most recent and most extensive of which has been the controller optimized for cache coherency—the joint effort of TI and Force Computers. This controller not only controls all of the allowable transaction types defined in Futurebus+ but defines a secondary local bus, called the H bus, which is basically a generic processor bus. In addition, the controller includes a cache buffer that can serve either as a buffer to an external cache or as a small secondary cache.

Though the inclusion of the complete protocol control functions on a single chip simplifies the design of a board to some extent, it also limits the maximum performance of that board. The local bus and cache architecture are the key elements in a system's performance and are expected to be one of the major differentiating factors between relatively similar products made by different manufacturers. This has largely been the case in VMEbus and Multibus environments.

and out.

© 1990 NEC Technologies, Inc.

But all you really need to know is that they're made by NEC, a 24-billion-dollar company, and the fourth largest manufacturer of disk drives in the world. For more information, call 1-800-NEC-INFO.

NEC

CIRCLE NO. 28

Buy The Left Side.
We'll Throw In The Right Side Free.

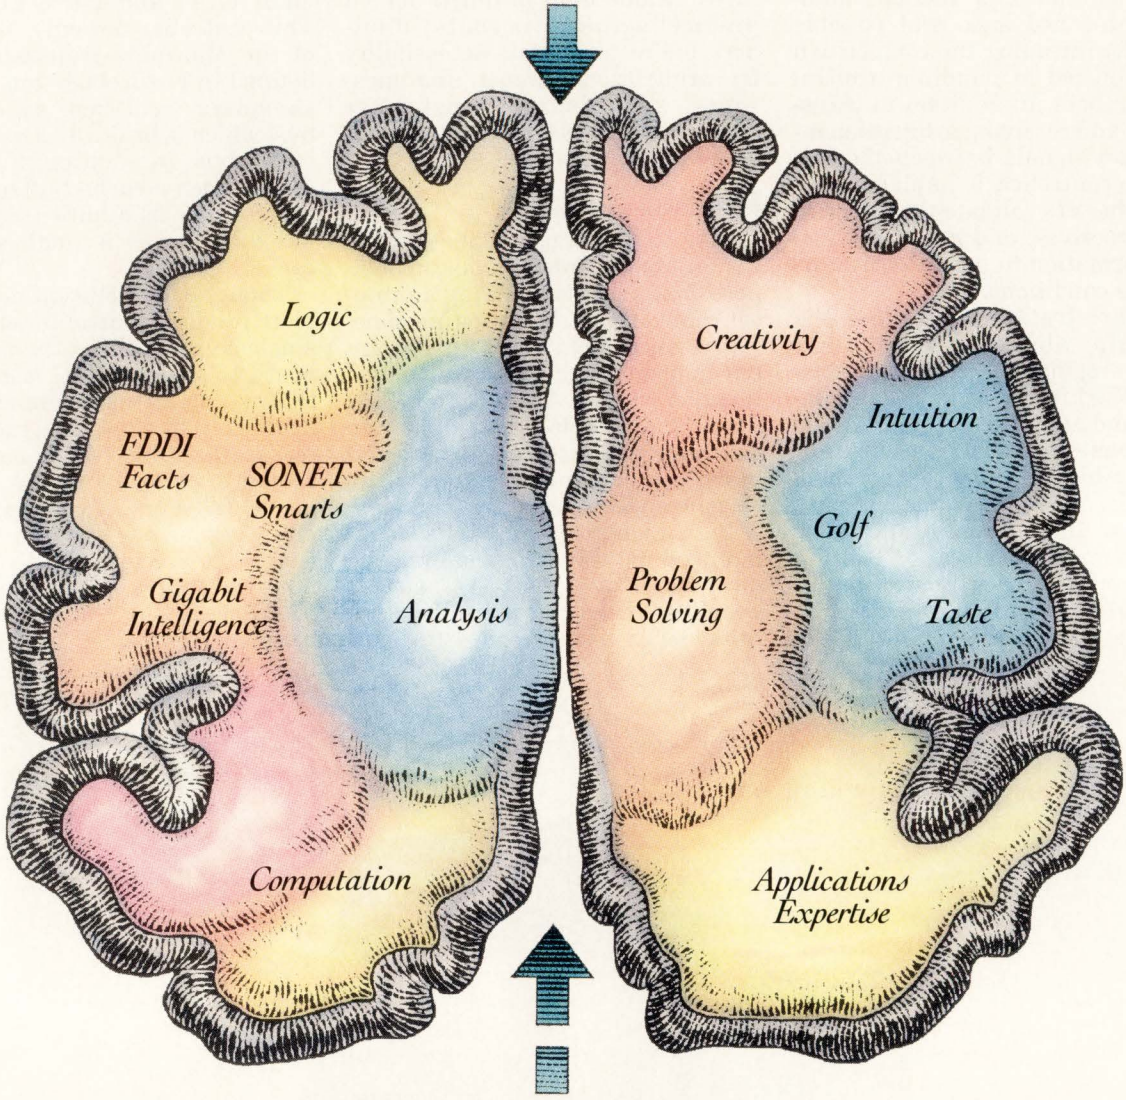

If you're trying to design a digital system, remember – there are two sides to the story: finding the best components, and knowing how to use them to your best advantage. Fortunately, BT&D has the brainpower for both.

With a global team of experienced sales and applications engineers, an advanced customer service lab, easy-to-use evaluation boards, product documentation and applications notes, BT&D offers services that help see you through even the most complex digital design challenges. As well as simplify component selection – no matter what your application or interface requirement.

From products for FDDI and other LAN applications to gigabit communications, our innovative components – manufactured in volume – are designed to meet your exacting needs. And, they're certified for some of the toughest quality specifications in the world. Including international quality standard, ISO 9001.

Plus, if one of our off-the-shelf components isn't precisely suited to your unique requirements, we can offer alternative solutions specifically for you.

So if you're racking your brain trying to keep up in today's ever-changing systems design environment, call BT&D at **1-800-545-4306.**

And put our brainpower on your side.

USA	Far East	Europe
Wilmington, DE	Yokohama, Japan	Ipswich, England
800-545-4306	(045) 593-4181	+44-473-742250

See us at OFC '91
San Diego, February 19-21

BT&D™
TECHNOLOGIES

CIRCLE NO. 29

COMPUTERS AND SUBSYSTEMS

"Futurebus+ is going in a number of different directions at once, and the particular implementation will have a lot to say about the design and system architecture," says Gillings. "I'm convinced that profile B—the I/O profile heavily backed by Digital Equipment Corp—will be the first to emerge in the commercial market after expected approval this spring." Others are in agreement, and some believe a formal announcement of profile B from DEC may come earlier.

National's chip set provides a real and immediate solution for Futurebus+ transceiver logic.

"Obviously," adds Gillings, "a full-featured protocol controller would be of no advantage in such I/O applications." Just how much silicon vendors will be pushing protocol controllers also depends on which Futurebus+ profile emerges after profile B. Gillings is betting that it will be the F (for fast) profile, which will once again preclude the desirability of using a standard protocol controller.

In the F profile, the total emphasis will be on performance. And while the profile will define a minimum transfer capability of 100 Mtransfers/s, the hot end of the system is expected to be the individual CPU-memory subsystem. To achieve the type of blazing performance expected—100 Mips per processor—it's likely that each processor, as well as each vendor, will have its own processor interface. And with each succeeding processor generation—perhaps as frequently as each year—that interface will probably change.

"Where the protocol controllers such as those developed by Force/TI and Signetics will fit," says Gillings, "is in the A profile." The A profile, still being defined, looks like a full-function profile of the bus, with more performance and more flexible functionality than the F profile. It's likely the A profile will be a follow-on for many applications, such as high-performance simulation engines currently using VME.

While National's chip set doesn't necessarily offer the level of integration of the chip sets proposed by either TI or Signetics, it does provide a real and immediate solution—particularly for Futurebus+ transceiver logic. Futurebus+ vendors will have to use the National transceivers to get product—even evaluation product—conforming to the expected final specification out the door in the first half of 1991. Even Force Computers, with its deal with TI, will be using National's parts on its proof-of-concept boards for its Navy contract. ■

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

Force Computers
(408) 370-6300 ..... **Circle 233**
National Semiconductor
(408) 721-5000 ..... **Circle 234**
Philips-Signetics
(408) 991-2000 ..... **Circle 235**
Texas Instruments
(800) 336-5236, ext. 700 ..... **Circle 236**

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.

Put Our Brainpower To Work For You.

Digital Communications Components from BT&D™

DLT1040/DLR1040 Transmitter/Receiver
- FDDI compliant
- Pin compatible with ODL 125®

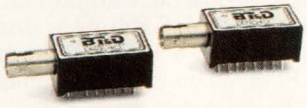

DLX2040 Transceiver
- FDDI compliant
- Pin compatible with ODL FDDI®

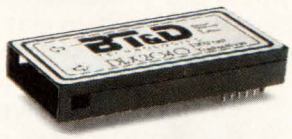

DLT1101/2-FC Single Mode Laser Transmitters
- DLT1101-FC, FDDI compliant
- DLT1102-FC, high power version

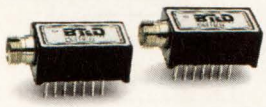

MFS2200 FDDI Bypass Switch
- Lifetime exceeds 10^7 cycles

XMT1300-1.2 Transmitter
RCV1201-1.2 Receiver
- 1.2 Gbit pair, ECL

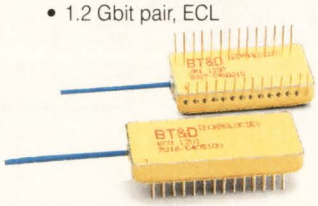

BT&D™
TECHNOLOGIES

ODL is a registered trademark of AT&T.

COMPUTERS AND SUBSYSTEMS

Semiconductor maker enters VMEbus board market

Warren Andrews, Senior Editor

Performance Semiconductor (Sunnyvale, CA), maker of high-speed semiconductor memory, logic and processors, has plunged into the board business, announcing its first commercial product—a MIPS CPU board. In addition, the company took the opportunity to announce an upgraded version of an earlier 1750A CPU evaluation card. "Our RISC-based single-board computer is a full-scale volume-manufactured board, not simply a prototyping tool for evaluation," says Nageen Sharma, marketing manager of RISC board products.

The company first stepped gently into the VMEbus business last year, when it offered its military customers a VMEbus CPU board incorporating its MIL-STD-1750A chip set. However, the test and evaluation vehicle for its chips was offered as a commercial version—as opposed to a full MIL-STD or industrial board—which would limit the product to evaluation only.

Since the 1750 specifications define a basically 16-bit processor with somewhat limited performance compared with that of more conventional commercial products, there is little commercial interest in the processor. The 1750 specifications were originally formulated by the Air Force as a "standard" military processor architecture. Since then, the specifications have been adopted, in one form or other, by other branches of the service. However, the military has recently offered contractors waivers from using the older processor in favor of far-higher-performance, 32-bit commercial processors.

In contrast to the military evaluation board, Sharma emphasizes, the new MIPS-based product is definitely targeted at the high-volume, commercial environment. "We're looking at any applications for a high-performance SBC," he says. "The board is general-purpose enough to appeal to OEMs, VARs and system integrators working on a variety of applications, from industrial controllers and imaging stations to simulation accelerators and robotic, military and scientific applications."

Will this attempt fly?

When other VMEbus semiconductor suppliers have competed with their own VME customers, the results, with a single exception, have been less than sterling. Of the three semiconductor makers that inked the initial VMEbus specification, for example, only Motorola has remained in the business. The other two ignominously bowed out subsequent to introduction. Mostek had to bail out after financial pressure forced it to

PERFORMANCE SEMICONDUCTOR'S RISC-BASED SBC

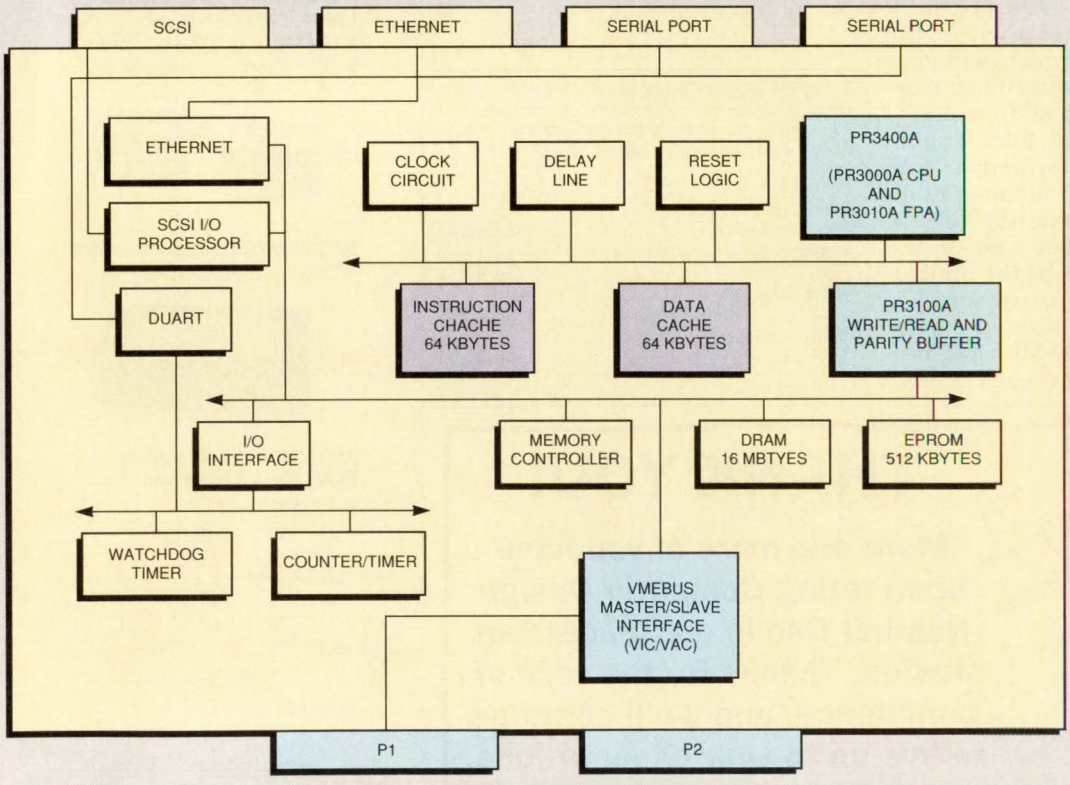

Performance Semiconductor's first commercial VME board features the MIPS R3000 RISC-based processor set and a high-performance architecture, including a 64-kbyte cache and 4 Mbytes of memory. The PaceRunner/3400 delivers up to 28 VAX MIPS performance in a single VME slot.

SOFTWARE CONFIGURABLE BOARD MAKES DATA ACQUISITION EASY

"With our 'Hands-Off' DT2831 board series for the PC AT, all set up and calibration is automatic via software."

–Fred Molinari, President

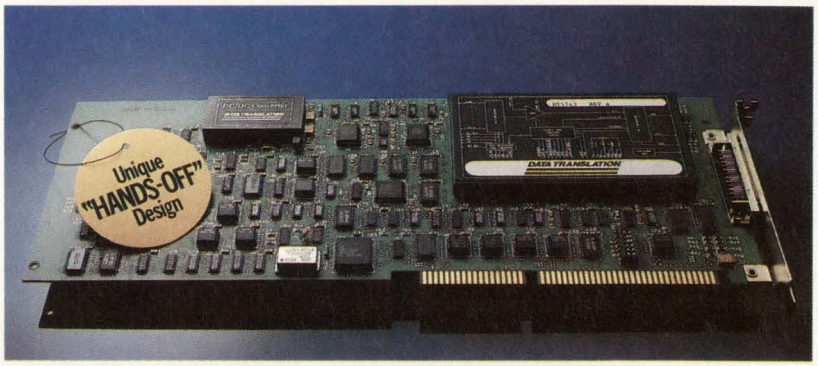

Total software control–No pots! No jumpers!
- Configuration of all board functions
- Self-calibration
- Gain and offset correction
- Independent selection of gains, sampling rates, DMA channels, and interrupts

Up to 250kHz gap-free data transfer

Simultaneous A/D, D/A

Up to 256 input channels with expansion panel

Two 16-bit counter/timers (AM9513)

Eight lines of digital I/O

Maximum noise immunity with shielded connector and analog modules

FREE software
- Drivers and subroutine library
- C, Pascal, BASIC and FORTRAN support

Application software available
- GLOBAL LAB™ *Data Acquisition*
- Leading software packages including LABTECH NOTEBOOK

Quantity pricing available; All prices U.S. list.

FAST 5 day delivery

Call for FREE Catalog

(508) 481-3700

In Canada, call (800) 268-0427

HIGH PERFORMANCE FEATURES

	DT2831	DT2835	DT2831-G	DT2836	DT2837
ANALOG INPUTS					
Channels	16	16	16	16	16
Resolution (bits)	12	12	12	16	16
Sampling Rate	50kHz	50kHz	250kHz	30kHz	30kHz
Gain	1,2,4,8	1,10,100,500	1,2,4,8	1,2,4,8	1,10,100,500
ANALOG OUTPUTS					
Channels	2	2	2	2	2
Resolution (bits)	12	12	12	16	16
Throughput	130kHz	130kHz	130kHz	100kHz	100kHz
PRICE	$1,750	$1,850	$2,450	$2,735	$2,795

DATA TRANSLATION®

World Headquarters: Data Translation, Inc., 100 Locke Drive, Marlboro, MA 01752-1192 USA, (508) 481-3700, Fax (508) 481-8620, Tlx 951646
United Kingdom Headquarters: Data Translation Ltd., The Mulberry Business Park, Wokingham, Berkshire RG11 2QJ, U.K., (734) 793838, Fax (734) 776670, Tlx 94011914
Germany Headquarters: Data Translation GmbH, Stuttgarter Strasse 66, 7120 Bietigheim-Bissengen, Germany 7142-54025, Fax 7142-64042
International Sales Offices: Australia (2) 699-8300; Belgium (2) 466-8199; Brazil 11 240-0598; Canada (416) 625-1907; China (1) 513-7766 x1222; Denmark 2 274511; Finland (0) 3511800; France (1) 69077802; Greece (1) 361-4300; Hong Kong (5) 448963; India (22) 23-1040; Israel 52-545685; Italy (2) 82470.1; Japan (3) 502-5550, (3) 5379-1971; Korea (2) 718-9521; Netherlands (70) 399-6360; New Zealand (9) 415-8362; Norway (2) 53 12 50; Poland (22) 580701; Portugal 1-7934834; Singapore 338-1300; South Africa (12) 803-7680; Spain (1) 555-8112; Sweden (8) 761 78 20; Switzerland (1) 723-1410; Taiwan (2) 3039836

GLOBAL LAB is a trademark and Data Translation is a registered trademark of Data Translation, Inc. All other trademarks and registered trademarks are the property of their respective holders.

Before you buy development

One glance at the advertising in this magazine and you know your embedded design involves mega choices. But too much of a good thing may lead to one thing: confusion.

We can help clear this up, for free. With the articles and technical notes you see above.

Because whether you're an old hand or just starting out, they'll provide timely insights and techniques. To help you get to market first.

Learn which development tools are right for your project. Do you need emulation or simulation? And just how can you debug less and enjoy it more.

These papers are based on Applied Microsystems' experience at setting up over 12,000 systems

© 1990 Applied Microsystems Corporation, P.O. Box 97002, Redmond, WA 98073-9702 USA. All rights reserved. AMC-9. Other names indicated by ® are registered trademarks of their respective holders.
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., Chiltern Court, High Street, Wendover,

tools, read the fine print.

worldwide for 8-, 16- and 32-bit designs. Our field engineers have seen enough to write a book.

But we've boiled some of it down to several thousand well-chosen words.

To order an article, or to receive a system demonstration, call 1-800-343-3659.

And get the bigger story.

Applied Microsystems Corporation

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.

CIRCLE NO. 65

COMPUTERS AND SUBSYSTEMS

sell its board division—and eventually the entire company. Philips-Signetics fared little better, wrapping up its board business a few years ago, and is now being faced by the massive layoffs implemented by its parent company.

But Performance's approach to the board business—offering a limited product family based on its leading-edge processors—may well help the company escape the fate of Mostek and Signetics. Performance is looking to capitalize on its semiconductor technology in an area where no other VME maker has yet jumped in—using Performance chip sets. Its premier product, the PaceRunner/3400, which was introduced at Buscon/91-West, is an SBC using the company's implementation of the MIPS R3000 RISC processor. Versions sporting 25- or 33-MHz chips will be offered.

Jumping on the VMEbus now is a lot different from what it was a few years ago, according to Joel Silverman, commercial product marketing manager for Radstone Technology (Montvale, NJ). Chips such as the VIC and VAC, developed by the VME Consortium, make it relatively easy for a vendor to design a VME board, says Silverman. Most of the work in putting together the interface, he says, is already done. In earlier-generation products, the VME interface not only took up a lot of relatively scarce board real estate but also required considerable engineering and design expertise. There was less time and space to devote to performance.

High performance, full function

"In designing the PaceRunner/3400, we took advantage of our latest processor, memory and logic chip technology so that we could provide a significant level of functionality. We believe it's the best-performing single-board VME machine," says Sharma. "Using the space savings realized from our chip sets, we can offer what the customer needs, with superior price/performance compared with that of discrete boards.

"The 6U, single-slot VME board uses our PR3000A/PR3010A integrated RISC CPU/FPA chip set, which is available at either 25 or 33 MHz," he says. "This results in performance of 21 and 28 VUPS, respectively. VUPS, or VAX units per second, is roughly equivalent to what's sometimes called VAX MIPS."

In addition to the MIPS processor, the board has 64 kbytes of data and instruction cache, an eight-word-deep write buffer, a programmable 32-word read buffer, 4 Mbytes of DRAM, a pair of RS-232 ports, on-board SCSI with a DMA controller, an Ethernet controller, 256 kbytes of EPROM, programmable counter/timers, a watchdog timer, and the VME interface for master/slave and interrupt operation. The board uses the VME Consortium's VIC and VAC chips for the VME interface.

Performance Semiconductor is looking to capitalize on its chip technology.

One of three

Of the three MIPS R3000-based CPU cards on the market, the PaceRunner/3400 is the only one with universal appeal and functionality, Sharma claims. "The other boards were designed for a particular application, and then adapted for the general-purpose market," he says. "Our board was designed to bring together the high performance of the MIPS processor and the most popular I/O options."

The other two contenders in the R3000 CPU market are the Lockheed Systems (Nashua, NH) board, available in both military and commercial versions, and the recently developed card from Omnibyte (Chicago, IL). Though Lockheed had the lead in offering the first MIPS board to the market, it was a militarized version—the company didn't release its commercial version until last fall. Omnibyte's board, claims company marketing manager Pete Czuchra, was the first commercial offering to hit the street.

However, Lockheed's approach uses only a single board, while both the Omnibyte and the Performance solutions are basically two-board solutions—though the Omnibyte system occupies two slots and the Performance approach only one.

"There are always design tradeoffs in any design," says Czuchra. An analysis of both two-board solutions reveals some significant differences in features—beyond the additional functions included in the Performance board. The Omnibyte board allows for more memory, 32 Mbytes instead of the 16 offered by Performance. In addition, it allows for 128 kbytes of high-speed cache instead of the 64 in the PaceRunner/3400.

One of the major reasons for vying for the two-slot solution, says Czuchra, was to take advantage of the additional power and ground pins available. With everything on the boards, it would have really been marginal using only a single slot. Also, a single-slot solution would have forced the use of double-sided surface-mount technology with increased cost and possible cooling problems.

"We've had no problems with the power and ground in a single slot," says Sharma. And, he adds, with components surface-mounted on both sides of the board, the company has had no problems with the board overheating in a standard VMEbus environment.

Part of the key to Performance's tight design is that this company, like Omnibyte, provides a piggyback design—only the double-sided surface-mount technology allows Performance to pack both boards into the width of a single slot. In Performance's approach, the processor chip set resides on its own card, while the other functions (VME interface, communications, and so forth) are on the standard VME board. Instead of one of the standard VME mezzanine bus standards, Performance elected to use its own approach, which comprises essentially the signal pinout of Performance's processor interface chip, the PR3100.

Upward migration

In addition to the more-compact design, Performance may have gained some other advantages by using the double-sided surface-mount construction technique. "At introduc-

THE EXTREME PERFORMANCE

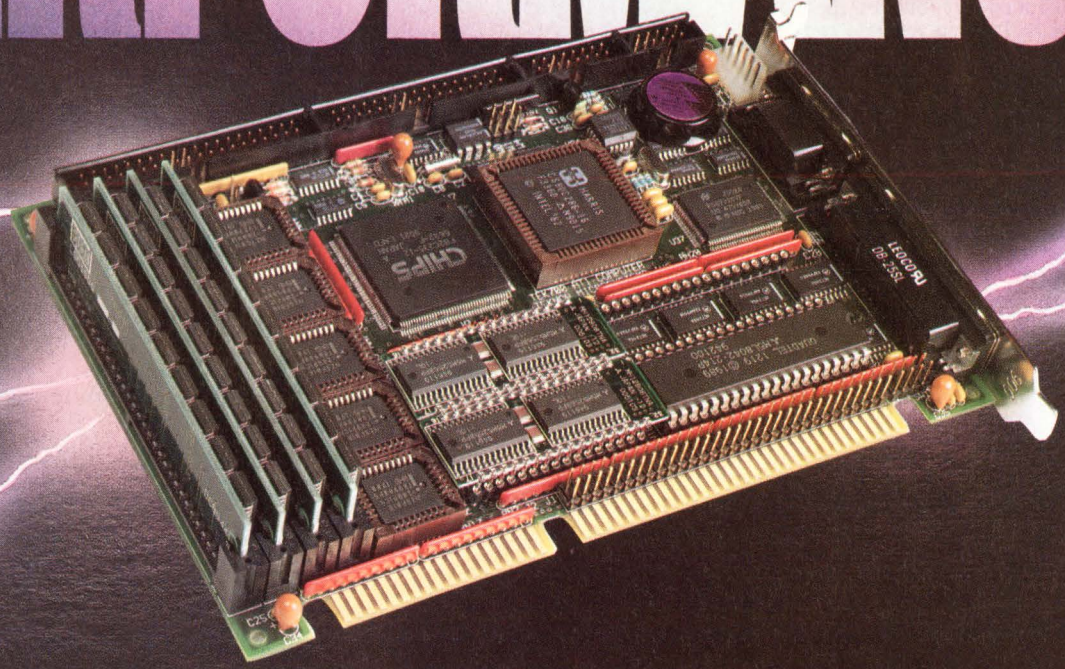

SO MUCH FROM SO LITTLE

High Level Of Integration. The TEK-AT1 features an 80C286 at up to 20 Mhz with system memory from 512 Kbytes up to 4 MBytes, two serial ports, one parallel port, a watchdog timer, a power failure detector, solid state disks with support for FLASH EPROMS, floppy and hard disk controllers.

Versatility. The TEK-AT1 can be used in a PC/AT passive backplane or as a stand alone computer for embedded applications.

Reliability. Teknor's products are built to operate in harsh environments. The TEK-AT1 can operate in extented temperature ranges and require very little power (sleep mode supported), typically less than 4 watts. The TEK-AT1 is backed by a two year warranty.

The solution: CALL 1 (514) 437-5682

TEKNOR MICROSYSTEMS INC.

CIRCLE NO. 31

The right connection

P.O. Box 455, Sainte-Thérèse (Québec) Canada J7E 4J8 • Fax: (514) 437-8053

COMPUTERS AND SUBSYSTEMS

tion, we'll offer 25- and 33-MHz versions of the boards, but by the second quarter of this year," says Sharma, "we expect to have a 40-MHz version ready. We already have silicon out and being sampled."

Omnibyte's current product uses a 25-MHz chip set, which the company will probably continue to offer until its next-generation product is ready. "We'll be putting some optional features on our board, such as SCSI, Ethernet and additional I/O ports, through the use of a mezzanine bus," says Czuchra, adding that it doesn't look as if the current board architecture will be able to accept a faster processor. "We'll probably wait until the next-generation processor, the R4000, to boost the clock speed," he says.

As processor speed increases, it becomes more important to keep the coupling of the processor section as tight as possible—and that means putting chips close together with very short printed circuit board traces. Using double-sided surface-mount technology means that interconnection distances can be kept to a bare minimum. "Keeping interconnection lengths to a minimum with this approach lets us offer faster processors on the same board," says Sharma.

Though Czuchra and Sharma declined to comment on features and performance anticipated from the emerging MIPS R4000 chip set, there's speculation that it will provide a higher level of integration and, undoubtedly, better performance. "The higher level of integration may eliminate the need for double-sided surface-mount for perhaps another generation," Czuchra says. ■

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

Lockheed Systems (603) 885-4321	Circle 201
Omnibyte (708) 231-6880	Circle 202
Performance Semiconductor (408) 734-8200	Circle 203
Radstone Technology (201) 391-2700	Circle 204

NEW FOR VME

The Next Generation SBC

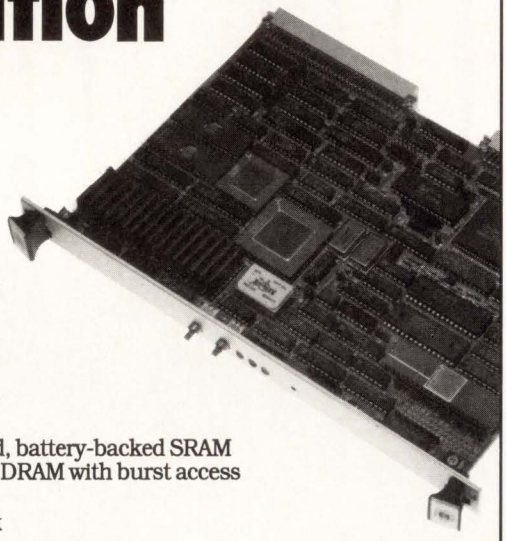

We listened carefully over the past 12 years as we've pioneered VME Single Board Computers. We now present the V36, a new generation of higher-performance, lower-cost SBCs designed for real-time applications.

The V36 Features:

- 68030 MPU at up to 40MHz
- SCSI Interface with a 32k Buffer
- 8 Channels of Serial I/O
- Ethernet™ with local dual-ported, battery-backed SRAM
- 2 or 8MB of two-way interleaved DRAM with burst access
- Up to 2MB of EPROM
- Battery-backed Real-Time Clock
- Optional 68882 Coprocessor
- Pin Compatibility with Motorola MVME147
- 68040 Version Available

Triple Local Buses for High-Performance

Through the use of buffers and component level DMA, the CPU, Ethernet and SCSI can operate simultaneously and independently for maximum performance and collision-free operation.

Expandable

Additional features can be added to the V36 via a local expansion bus, allowing system designers the flexibility to include proprietary features in an off-the-shelf computer, further reducing design costs.

Designed for Design-In

The V36 was designed for use in enclosed systems. All I/O on the board is routed through the P2 connector, which means no cable hassles during board swaps and upgrades. Additionally, the V36 makes extensive use of CMOS components, guaranteeing low overall power use (7W typ.) and the absence of "hot spots". The V36 can be used without a cooling fan in most applications.

Software Support

Software support is available for all major real-time operating systems, including VxWorks, OS-9, PDOS and PSOS+.

Call today to learn more about the V36 SBC, and our complete line of VME CPUs, I/O and memory boards.

GMS GENERAL MICRO SYSTEMS INC.

4740 Brooks St., Montclair, CA 91763 (714) 625-5475 FAX (714) 621-4400

CIRCLE NO. 32

Quit smoking.

American Heart Association

WE'RE FIGHTING FOR YOUR LIFE

33 MHz. 80960CA. VME.
35 nsec SRAM. DRAM. DMA.
SQUALL Module™. VxWorks.

Cyclone's Squall Hits 80960CA.

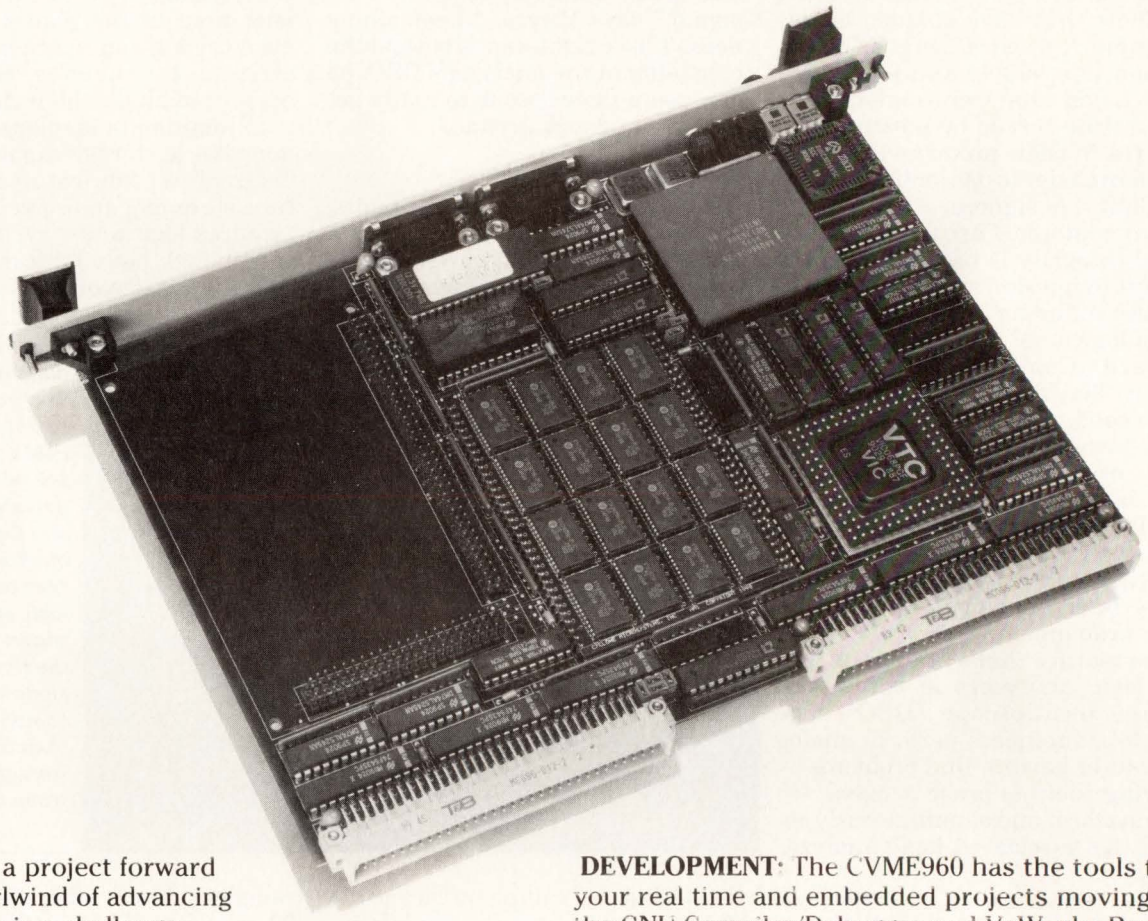

Propelling a project forward in the whirlwind of advancing technology is a challenge. Cyclone Microsystems can help accelerate your product development with the CVME960 Single Board Computer.

PERFORMANCE: Cyclone's CVME960 offers the 80960CA's SuperScalar performance of multiple instructions per clock cycle and an on-chip DMA controller. Cyclone couples the 80960CA's power to either 2 Mbytes of 35 nsec SRAM or up to 8 Mbytes of fast DRAM.

FLEXIBILITY: A SQUALL Module™ is an I/O module that utilizes the 80960CA's capabilities as an embedded processor and an I/O server. The interface to the SQUALL Module includes four 32 bit DMA channels as well as the 80960CA's address, data, and control signals. SQUALL Modules are available for High Speed Data Links, Communications, and Network Interfaces.

DEVELOPMENT: The CVME960 has the tools to get your real time and embedded projects moving. Like the GNU Compiler/Debugger and VxWorks Real Time Operating System with a symbolic source language debugger.

SUPPORT: Cyclone backs the CVME960 with comprehensive technical and software support. Engineering teams are ready to support your unique SQUALL Module requirements.

Call Cyclone Microsystems to find out how the performance and flexibility of the CVME960 and SQUALL Modules™ can help you survive the storm.

CYCLONE
MICROSYSTEMS
25 Science Park New Haven, CT 06511
(203) 786-5536

VxWorks is a product of Wind River Systems, Inc.

CIRCLE NO. 33

DESIGN AND DEVELOPMENT TOOLS

Combined scope/logic analyzers detect hard-to-find analog problems

Mike Donlin, Senior Editor

As microprocessor speeds grow and circuit board sizes shrink, digital engineers face the same problems that have plagued analog designers for years, such as transmission line effects and loading effects. Logic analyzer vendors have come to the rescue by putting more features in their products and making them easier to use for both analog and digital designers.

"Our customers are telling us that signal integrity is becoming too important to ignore," says Greg Peters, product manager for the 16500 series of logic analyzers from Hewlett-Packard (Colorado Springs, CO). "When they're dealing with these high-speed, closely packed circuits, every picosecond counts. But when many of today's digital designers were in college, they learned on a 1-MHz breadboard system, where they didn't have to worry about crosstalk and ringing, so logic analyzers have to help them address these problems."

One feature that's finding its way into logic analyzers is the digital storage oscilloscope (DSO). The DSO lets engineers perform analog analysis of hard-to-find problems by allowing them to probe a signal, see the waveform and simultaneously observe the associated logic analyzer views.

HP's 16500 series of logic analyzers offers an optional DSO module with a 400-Msample/s digitizing rate. This digitizing rate, with 4,000 samples per channel, lets users view events up to 10 µs before triggering, with better-than-1-ns time interval accuracy. Automatic display of waveform parameters on the scope lets users analyze a signal's behavior without having to count graticules. Because the scope can be triggered by the analyzer, or vice versa, error conditions that occur infrequently can be observed many times and accumulated over multiple acquisitions.

"If an engineer is looking at an address or data bus that has a trigger condition that's easily found with the state analyzer, he can use that to repetitively trigger the scope and probe while he's looking at the signal," says Peters. "If something doesn't look right and is beyond the capability of the analyzer's DSO, he can use a faster scope to really get in there and look around."

Storing the signals

When debugging complex circuits, it's often necessary to build up a library or history of waveforms, as well as of state and timing analyses. By storing and calling up the results of previous debugging sessions, engineers can assemble a portrait of a circuit's behavior, as well as trace problems back to changes made in the original design.

Tektronix (Beaverton, OR) offers a dual-channel, 8-bit digitizing scope module, which can be installed in its Series 1230 or Prism 3002 logic analyzers, that fills the need for storage capacity. Waveforms and logic analyzer information can be stored on the Prism's 720-kbyte floppy or 20-Mbyte hard drive. To accurately capture digital and analog information, the module hosts 32 kbytes of memory per channel.

"We decided to have that much memory per channel when we looked at a typical use for the scope/logic analyzer combination," says Roger Crooks, product marketing manager for midrange logic analyzers at Tektronix. "In debugging a circuit where you're looking for analog effects in a digital design, you need enough memory depth to store an entire event. What may be a short time in the analog world is a long time in the digital world. So even though you're tracing a digital circuit, you need a reasonable amount of analog information."

In addition to memory depth and nonvolatile storage capabilities, the Tektronix module features automatic time stamping and correlation of all acquired logic and waveform data.

Although both Tektronix and HP claim that their scopes' 400-Msample/s sampling rate is adequate for about 90 percent of debugging problems, they will undoubtedly be expanding this capability in the future. As circuit boards shrink and packaging alternatives such as multichip modules become popular, designers will need even more oscilloscope power to track down elusive analog effects. These combined scope/logic analyzers will undoubtedly help. ■

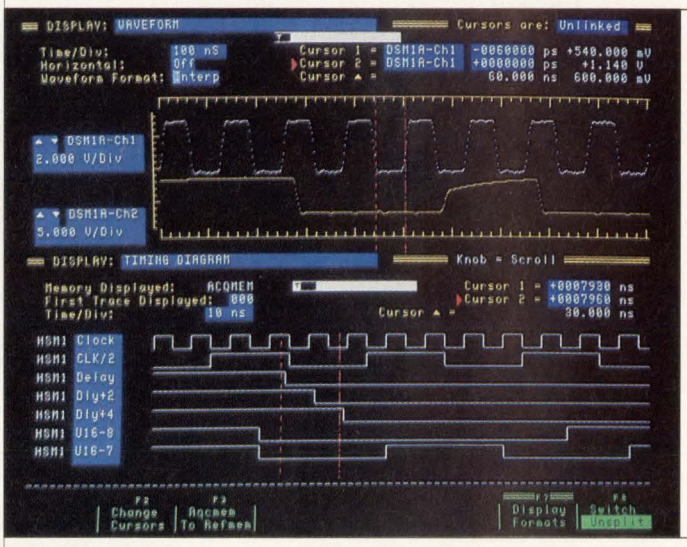

Simultaneous displays of information from different sources help engineers to see relationships between analog and digital information. Here, the Tektronix Prism logic analyzer shows data from the timing module along with two analog waveforms from the digital storage oscilloscope module.

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

Hewlett-Packard
(719) 590-1900 ..... Circle 205
Tektronix
(503) 629-1960 ..... Circle 206

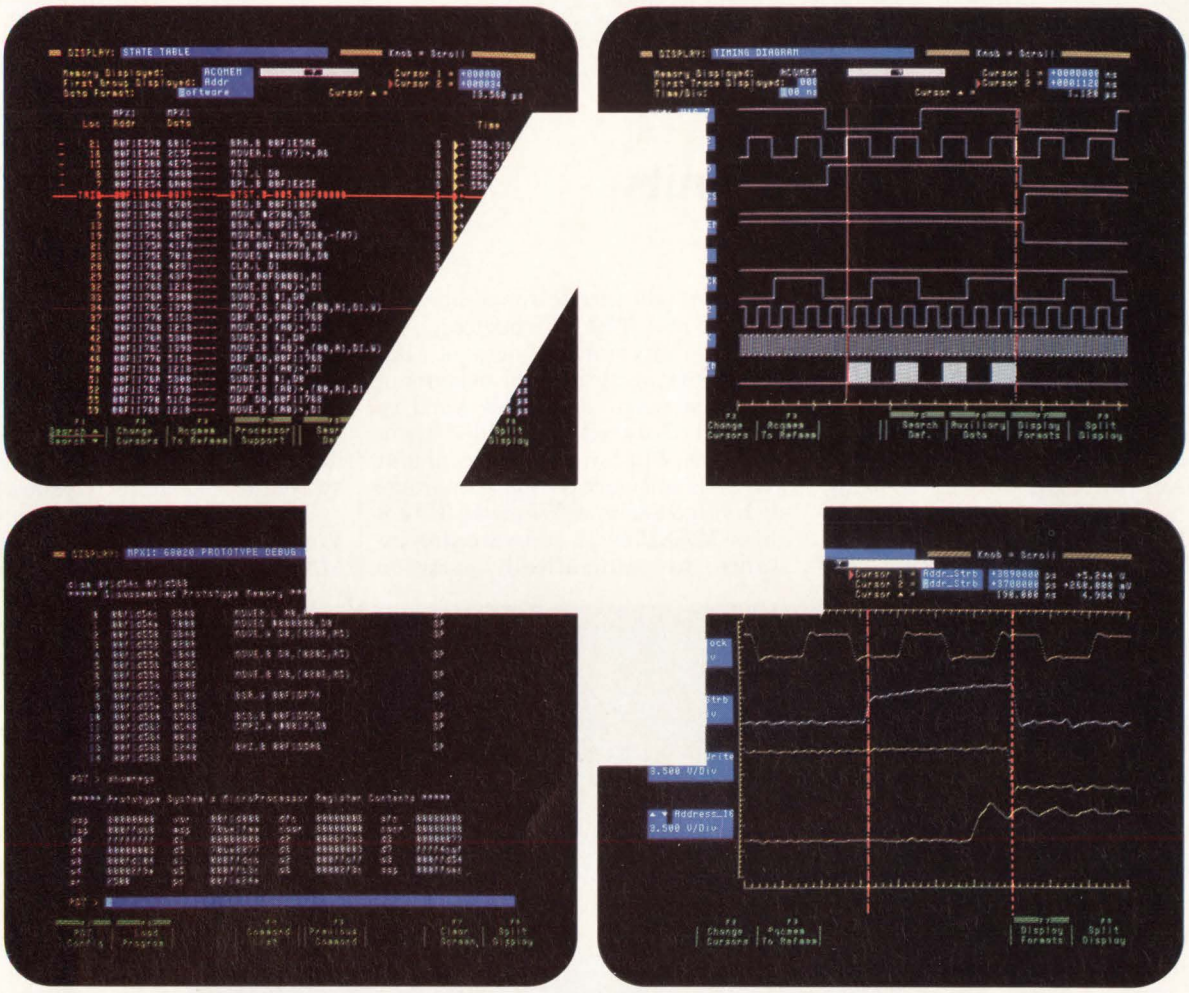

Prism. Four instruments in one logic analyzer. For faster, time-correlated results.

State analyzer. Timing analyzer. Emulator. DSO. Getting the complete picture is easier than you think.

The 4-in-1 Prism 3000 Series is unlike any other logic analyzer. From one keyboard and display, it can do the work of multiple instruments. It can time-correlate data acquired by one Prism module to data acquired by all other modules, via revealing split-screen displays.

Combine the triggering and channel resources of a logic analyzer with emulator-like ability to change registers, patch memory, and step through your code. View both signal timing and signal integrity at the same time with the integrated DSO.

For twice the power and convenience afforded by separate instruments, at half the cost, contact your Tek sales engineer. Or call 1-800-426-2200 for the complete multipurpose Prism story.

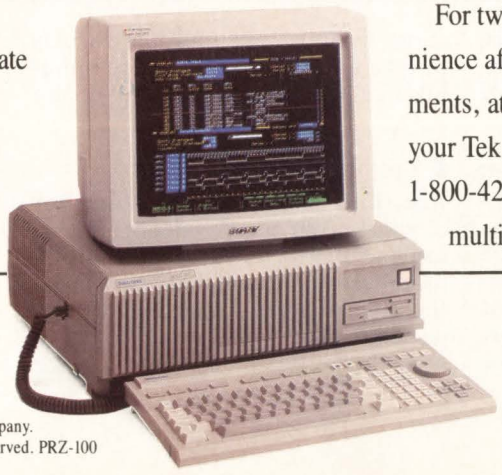

Tektronix
COMMITTED TO EXCELLENCE

CIRCLE NO. 34

Rental units available from your local rental company.
Copyright © 1990, Tektronix, Inc. All rights reserved. PRZ-100

DESIGN AND DEVELOPMENT TOOLS

Windows 3.0 extends PC-based PLD design limits

Barbara Tuck, Senior Editor

Programmable logic devices have traditionally been designed on personal computers, but a family of PLDs from Altera Semiconductor (San Jose, CA), with as many as 20,000 gates and 280 pins, came close to threatening that tradition. To help designers wring every ounce of productivity possible from PCs, Altera will ship in March the MAXPlus II programmable logic design toolset, the first CAE software to run under the Windows 3.0 graphical environment from Microsoft (Redmond, WA).

Altera silicon users will have a few months to explore the graphical user interface, multitasking, and virtual memory management of Windows 3.0 before samples of Altera's high-density MAX 7000 CMOS erasable PLD series become available at midyear.

Since Windows 3.0 uses virtual memory, Altera customers will be able to leapfrog the 640-kbyte DOS memory barrier without moving to workstations running Unix. And the MAXPlus II toolset gains enough support to handle design synthesis, automatic place and route, and simulation of even the densest MAX 7000 designs.

The support offered by the Windows 3.0 graphical environment will become increasingly critical to designers who continue to rely on PCs. The resources of Windows will be required for both device- and system-level complexity. Operating at speeds up to 70 MHz, the largest members of the MAX 7000 family, for instance, will integrate as many as 300 SSI and MSI TTL devices.

With individual devices of that complexity, system-level designs of 50,000 or more logic gates won't be unusual. Only with the added memory of the Windows 3.0 environment will CAE software be able to manage such huge designs. Windows 3.0 enables MAXPlus II software, for instance, to automatically partition large designs into multiple EPLDs and to support multiple-chip simulation for verification of the entire design.

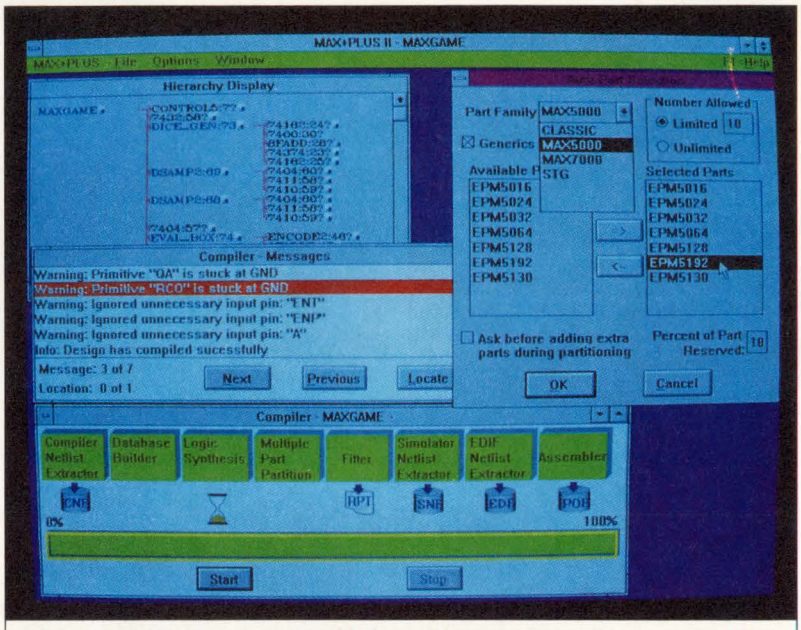

Altera's MAXPlus II programmable logic design toolset operates under Windows 3.0, and the combination provides a graphic user interface supporting multiple design windows, extended memory management, and multitasking. The MAXPlus II logic compiler takes advantage of Windows' memory management to automatically partition large designs into multiple erasable-programmable logic devices.

The next PC standard

According to Brad Fawcett, director of marketing for development systems at FPGA market leader Xilinx (San Jose, CA), Logic Cell Array customers are beginning to look for the support that Windows 3.0 would give them. When a PC is busy for perhaps a few hours placing and routing a dense FPGA, says Fawcett, customers would like the capability to do other tasks on the PC at the same time.

There's no argument that productivity would be increased if users were able to enter and compile new designs while waiting for a prior one to be laid out. Xilinx, having recently released the first members of its XC4000 family of Logic Cell Arrays, hasn't yet committed to modifying its XACT 4000 CAE development system to run under Windows 3.0.

PC-based design software vendor Viewlogic Systems (Marlborough, MA) is evaluating Windows 3.0, but hasn't yet decided whether to adopt the environment for its next major revision of software. According to a Viewlogic spokesperson, customers aren't screaming for Windows 3.0 yet. Although the EDA vendor has given its customers the functionality of a proprietary windows-based environment, Viewlogic is likely to back Windows 3.0 as it emerges as a standard, just as the company has done with EDIF (Electronic Data Interchange Format) and VHDL in the past.

Introduced by Microsoft in May 1990, Windows 3.0 is expected to dominate new PC-based applications development over the next few years. An indication of this environment's impact is that Microsoft shipped more copies of Windows 3.0 in the first six months of its existence than of the more sophisticated, more expensive and less portable OS/2 in the three years since it was released.

As a "shell" program that runs on top of DOS, Windows 3.0 can run on any current DOS platform. OS/2, on the other hand, is hardware-dependent and needs to be ported to the specific PC hardware if it differs at

Take the proven shortcut to VME serial communications.

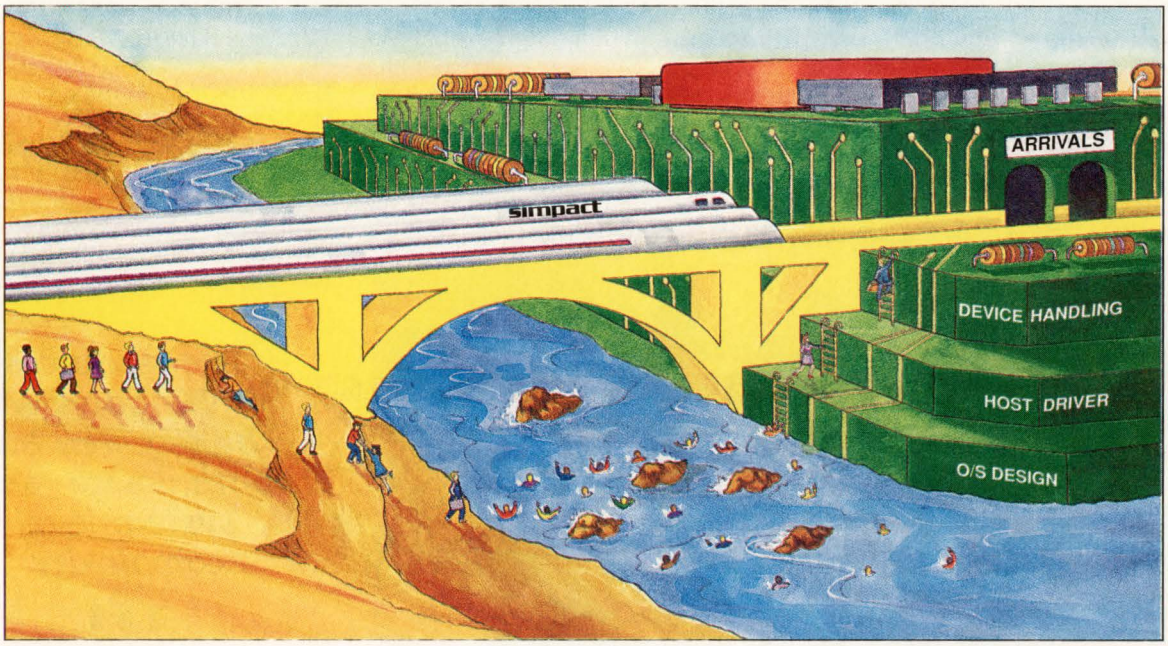

Simpact offers the UNIX® world's most complete Communications Software Toolkit for fast protocol porting.

With Simpact's Communications Software Toolkit, you can port your current protocols or develop new ones for VME systems faster, more economically, and with fewer engineering resources than ever before.

You're miles ahead—*before* you start.

Porting or developing a board-level protocol no longer means time-consuming design of an operating system and drivers. The Communications Software Toolkit gives you those resources *ready-to-use.*

Simpact's toolkit includes a sophisticated real-time operating system, UNIX host drivers, board-side host interface, and message handling routines. Templates provide source code for asynchronous, bit-synchronous and byte-synchronous protocols. Simply use the templates' source code, or modify it quickly and easily to adapt to your specific protocol needs.

The only source you'll ever need.

Simpact is the only single-source supplier of both a complete Communications Software Toolkit and intelligent processor boards for VME systems. Our multiport front-end boards with 16MHz 68020 processors offload communications from your host.

Support you can count on.

With our unsurpassed field support, you have a partner who can help you navigate rough water. Call today and we'll show you the only shortcut worth taking.

1-800-488-4188

Simpact's off-the-shelf protocols were developed using the toolkit.

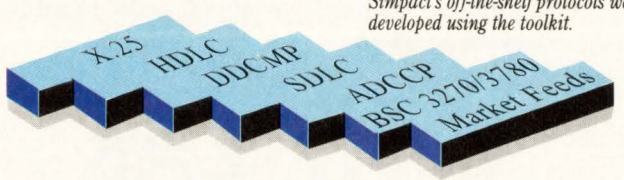

simpact

Simpact Associates, Inc., *9210 Sky Park Court, San Diego, CA 92123, 800-488-4188, FAX 619-565-4112*

Trade names referenced are the trademark or registered trademark of the respective manufacturer.

CIRCLE NO. 35

DESIGN AND DEVELOPMENT TOOLS

all from IBM's configuration.

For design entry, the Windows 3.0-based MAXPlus II software includes hierarchical schematic capture and the Altera hardware description language. It also features automatic error location, full timing analysis, and simulation. According to Altera, a 5,000-gate design can typically be compiled in 10 minutes or less with the new software. To help the logic designer organize, manage and execute multiple design tasks, the Windows 3.0 graphical user interface offers icon menu selections and configurable application windows. And because MAXPlus II supports EDIF, workstation-based designers can use their design entry and verification tools together with MAXPlus II for EPLD logic synthesis.

Users switching from Altera's MAXPlus to Windows 3.0-based MAXPlus II will be able to take advantage of graphics boards with much higher resolutions than the 640×480-pixel boards to which they have been confined with the older software. Windows 3.0 also supports over 150 printers and plotters.

The resources of Windows will be required for both device- and system-level complexity.

■ ■ ■

In the end, there's not much that can be said against a software standard that costs only about $150 and allows designers to squeeze functionality out of their PCs. Windows 3.0 will unquestionably not make design software less attractive. The question then is, can Windows 3.0 make a silicon vendor's software attractive enough to be the determining factor in the selection of a PLD or an FPGA vendor? Or are designers sold on a vendor's silicon to the point that they put up with the software for that silicon no matter how inferior? Or are all of these considerations academic since designers will take device-independent, top-down design seriously and will, en masse, abandon their PCs for workstations? ■

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

Altera Semiconductor
(408) 984-2800 ..... **Circle 217**
Microsoft
(206) 882-8080 ..... **Circle 218**
Viewlogic Systems
(508) 480-0881 ..... **Circle 219**
Xilinx
(408) 559-7778 ..... **Circle 220**

Small Wonders

Mizar packs up to 4 MB on its single-height 68020 and 68030 VME CPU boards.

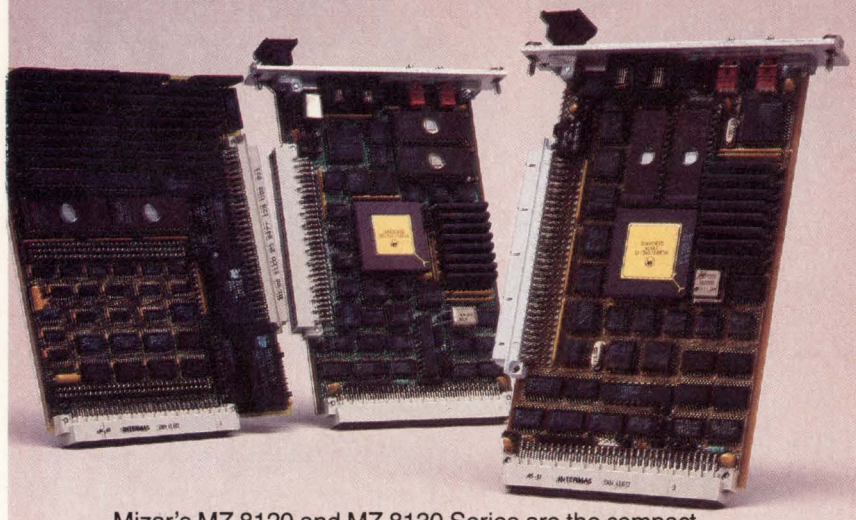

Mizar's MZ 8120 and MZ 8130 Series are the compact, flexible solutions for your application. With a 68020 or 68030 processor and one or four Mbytes of dual-access DRAM, they're perfect for applications where space is tight and performance can't be sacrificed. And with Mizar's MXbus™ expansion interface, you can tailor custom double-height boards based on a single-height CPU.

When you choose Mizar, you've got a VMEbus partner you can rely on. No one is more experienced than Mizar, because we were the first company to introduce single-height boards. Today, Mizar provides high-performance 3U CPU boards based on the 68010, 68020, and 68030, as well as memory, I/O, and peripheral controllers. Complete technical support and integration services are also available.

It's no small wonder that more and more engineers are turning to VME boards from Mizar for the toughest applications. Find out how you can benefit. Call today.
1-800-635-0200.

Mizar. The shortest distance between concept and reality.

MIZAR

1419 Dunn Drive / Carrollton, Texas 75006 / 214-446-2664 / FAX 214-242-5997

MXbus is a trademark and Mizar is a registered trademark of Mizar Digital Systems, Inc.

©1990 Mizar Digital Systems, Inc./ 1017c

CIRCLE NO. 36

"High Performance VME... Only From GEM"

Super High Performance Backplane

- Full VME Specification
- 10 Layer Stripline Construction
- Virtually Zero Cross Talk
- Controlled Impedance
- Noise Reduction
- 3-21 Slots Monolithic J1/J2

New Enclosure Design

- Aluminum Construction
- Easy to Access
- Direct Flow Cooling
- Easy to Modify
- 3-20 Slot Configurations

Call 1-800-628-4095

Graham Electronics Manufacturing Inc,
109 State Pl, Escondido, CA 92029,
(619) 747-9177 FAX (619) 747-9326

CIRCLE NO. 37

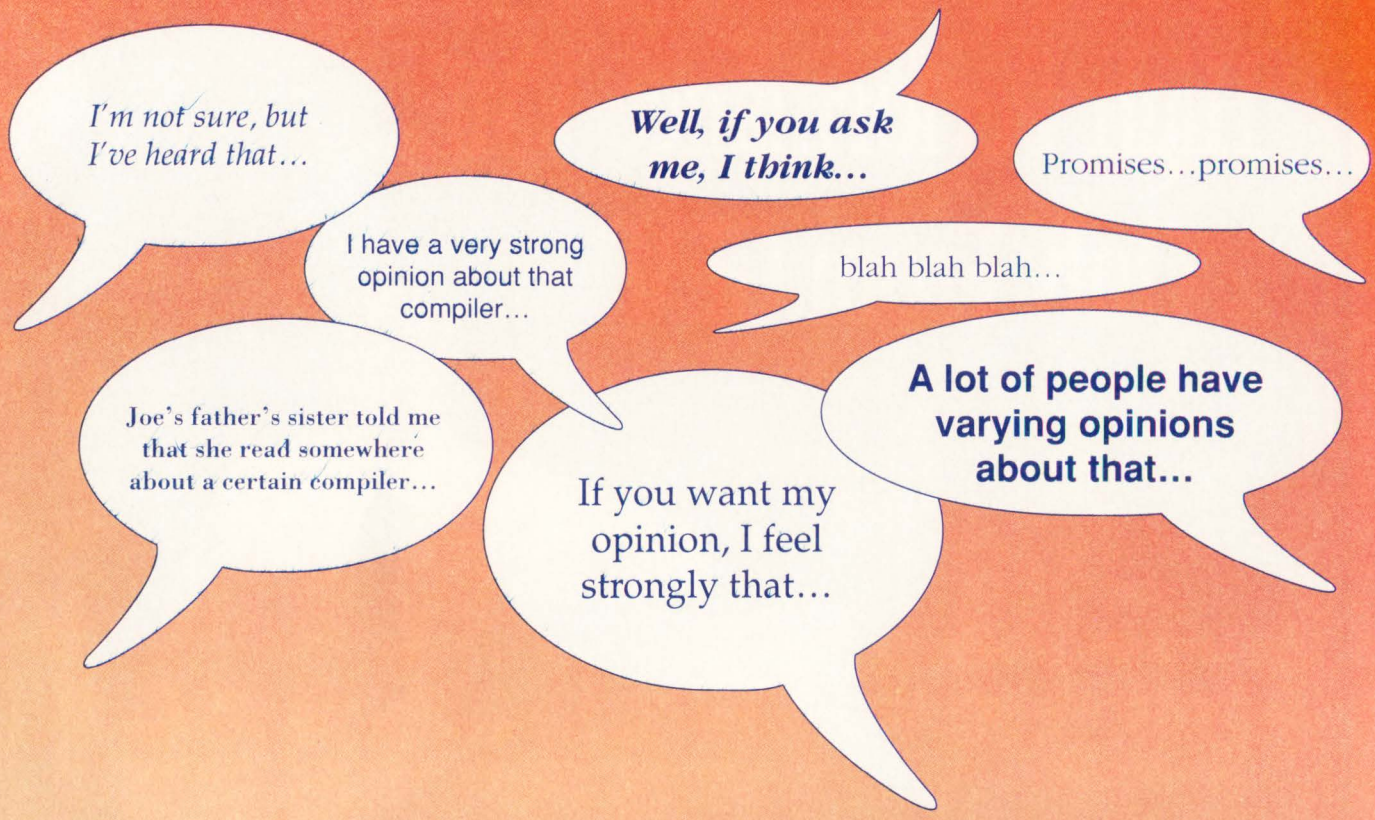

'NUF SAID.

Nowadays, it seems like everybody has an opinion on which 8051 C compiler performs best.

The truth is, Franklin Software's C51 8051 C Compiler Professional Edition generates the quickest, most responsive 8051 code in the world.

And that's not just our opinion or hearsay. It's fact. Just take a look at the benchmark results. Now you can see why professionals can't settle for anything less.

Franklin Software's C51 8051 C Compiler Professional Edition. When you need the best in the world demand Franklin 8051 software!

TO ORDER CALL
408-296-8051

FRANKLIN
SOFTWARE, INC.
888 Saratoga Ave. #2, San Jose, CA 95129

NEW PROFESSIONAL EDITION

BENCHMARK RESULTS

	FRANKLIN SOFTWARE C-51 v3.0 PROFESSIONAL	IAR/ Archimedes ICC8051 v4.01B[2]	BSO/Tasking CC51 v1.0.3
Sieve Benchmark[1]			
Memory model	*SMALL*	*TINY*	*SMALL*
Total execution time	*6.744 sec*	*8.388 sec*	*8.753 sec*
Module code size	*188 bytes*	*238 bytes*	*266 bytes*
Dynamic data size	*41 bytes*	*96 bytes*	*126 bytes*
Total code size	*1223 bytes*	*1626 bytes*	*1936 bytes*
Memory model	*LARGE*	*LARGE*	*LARGE*
Total execution time	*8.287 sec*	*17.295 sec*	*10.98 sec*
Module code size	*188 bytes*	*343 bytes*	*303 bytes*
Dynamic data size	*21 bytes*	*28 bytes*	*74 bytes*
Dynamic Xdata size	*8238 bytes*	*8259 bytes*	*10,351 bytes*
Total code size	*1292 bytes*	*2088 bytes*	*3467 bytes*
Dhrystone Benchmark[3]			
Dhrystones per second	*203*	*90*	*163*
Single Precision Whetstone Benchmark[4]			
Whetstones per second	*13*	*3*	*(Floating point not yet supported)*
Fibonacci Numbers Benchmark[5]			
Total execution time	*20.88 msec.*	*94.29 msec.*	*368.14 msec.*

1. Erathosthenes Sieve Program from BYTE (1/83), expanded with I/O and interrupt handling. (10 iterations)
2. Currently available only in Europe.
3. Dhrystone Benchmark Program C/1.1 from Reinhold P. Weicker, CACM Vol 27.
4. Single Precision Whetstone Benchmark Program (A001)
5. Fibonacci Numbers Benchmark Program, Embedded Systems Magazine, April 1989

FDDI chips struggle toward the desktop

The success of FDDI networking may stall at the high end unless workstations and PCs are included. The needed drop in cost awaits progress in standards and components.

Ron Wilson
Senior Editor

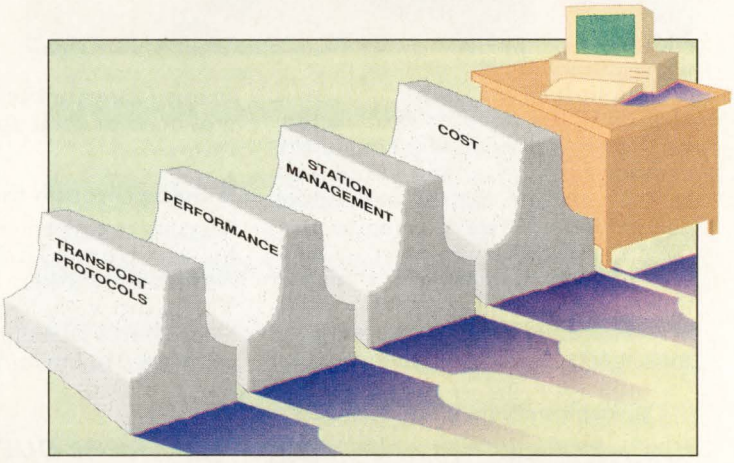

While FDDI (Fiber Distributed Data Interface) is solidly accepted as a backbone network, the technologies and standards supporting it face substantial obstacles before they achieve FDDI's goal of connecting every desk.

The 100-Mbit/s fiberoptic network is winning as a backbone link. But FDDI (Fiber Distributed Data Interface) is still a long way from its goal of a connection on every desk, and chip technology will have to help it get there.

As a corporate network backbone, a link between slower and less-expensive LANs, FDDI works. Many organizations are already pulling fiber through their walls, shopping for network bridges and setting up 100-Mbit/s links between work groups. Interoperability testing is pouring out a growing pool of adapters, concentrators and bridges that work together on the FDDI token ring. Even though there are still changes coming in some of the higher protocols, users can be confident that their adapters will play on the network and be upgradable as standards solidify.

But the big win for FDDI—the breakthrough that will pay back vendors' enormous investments in R&D—remains elusive. Relatively few workstations and virtually no personal computers are connected to the fiber networks. Until FDDI can move down from the corporate heights to the lowly desktop, sales won't take off, and the optical net will remain a specialty solution waiting helplessly till it's displaced by the next proprietary, superfast technology to come along.

Reaching the desktop is a vital step for FDDI, but most experts admit it's not ready yet. While there are many issues to be resolved—including shifting standards, a plethora of media choices, and serious concerns about the adequacy of workstation architectures—the biggest issue can be stated simply. Existing FDDI adapters can't achieve a cost/performance point low enough to ignite growth.

Seeking the right mix

The magic combination that will put FDDI on the desktop depends on both the right price and the right performance. Price is the most obvious problem. Most customers are resistant to paying more than 10 percent of the cost of

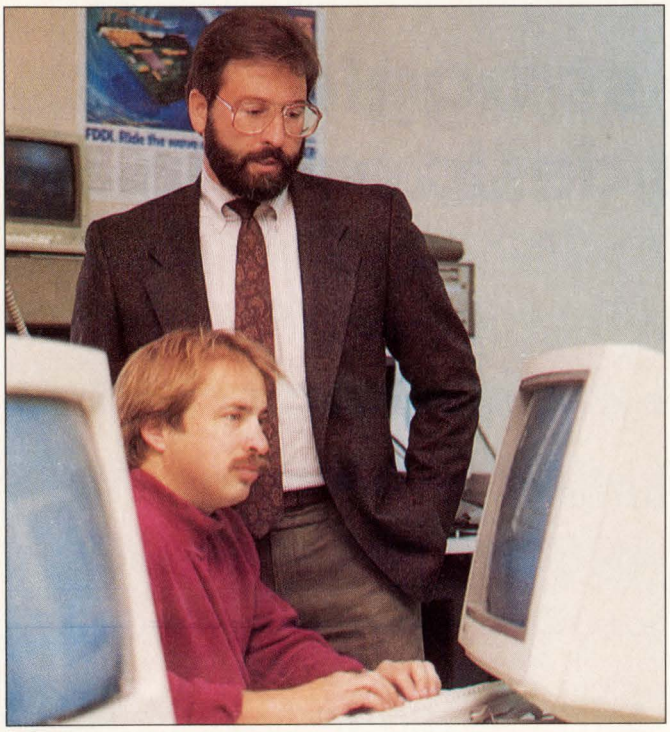

CMC director of engineering Jim Soriano (standing) points out that the high throughput of the Fiber Distributed Data Interface is required of the nodes as well as the media—and that slow nodes will likely be segregated on separate networks.

a workstation for connectivity. With today's FDDI adapter cards selling in the $6,000 to $8,000 range, this prejudice by itself limits the network to the very top end of workstation connections. "If we're going to meet our projected growth, we're going to have to get the cost of an adapter down to $1,000 by 1993," warns Karen Parker, strategic program manager at National Semiconductor (Sunnyvale, CA).

Performance can be a more subtle issue. After all, the speed of the FDDI ring is defined by the standard at 100 Mbits/s. But the overhead required to process the transport-layer protocol can slow things down considerably. If the adapter vendor saves cost by passing the protocol task off to the workstation CPU, the actual throughput can drop to Ethernet levels. And this severely compromises the attractiveness of the network. "If all it gets me is connectivity, not greater performance, I'll stay with Ethernet, thank you," says Ron Perloff, president of FDDI think tank XDI (San Diego, CA).

To get the mix right, silicon vendors, software vendors and adapter builders are all tinkering with a range of variables. It's in everyone's interest to sift out as much of the cost as possible from the FDDI silicon, but without losing the performance that makes the net attractive, and without sacrificing the flexibility the components will need to keep up with a shifting future.

Going through the budget

"If you look at the bill of materials for a current FDDI adapter," says National's Parker, "you see a 9U VME card completely full of stuff. There's an FDDI chip set. There are counters, timers and multiplexers to help the chip set do its job. There's a RISC CPU with plenty of support logic and a lot of memory, and there's a bus interface. Somehow we have to get that down to the space and power budget of a 3×5-in. SBus card."

"To meet our projected growth, we have to get the cost of an adapter down to $1,000 by 1993."

—Karen Parker, National Semiconductor

On examination, the problem neatly divides itself into three categories. "There are a few obvious high-cost parts," says Jim Soriano, director of engineering at adapter vendor CMC (Santa Barbara, CA). "First, there's the FDDI silicon itself—that needs to come down in price. Second, and related, is the cost of the processor and memory you need to support the chip set. Third, there are the optical components—which are very expensive. Getting the price down on them is a key to success."

Probably the least change will come in the FDDI controller chip set itself. Most existing adapters are built on the first-generation Supernet chips from Advanced Micro Devices (Sunnyvale, CA). Since that first effort, National has introduced a more-integrated set, AT&T (Berkeley Heights, NJ) has been sampling a physical-layer chip, and Digital Equipment Corp (Maynard, MA) has shared its own PMD (physical media dependent) chip set with both AMD and Motorola (Austin, TX). This latter effort has resulted in two designs, the just-announced Supernet-2 set from AMD and an as-yet-unannounced design from Motorola.

As a generalization, the chip sets are getting denser, more power-efficient and more attuned to the needs of the various FDDI protocol layers. But there is a limit. The speed of FDDI dictates a mix of CMOS for the upper layers and bipolar for the lowest layers—technologies that still don't mix on economy-minded chips. And there's so much diversity on both ends of the chip set—the physical media on one end and the CPU interface on the other—that further integration would risk cutting off potential market for the chips.

The most recent of the announced designs, Supernet-2, gives a good picture of the direction in which the silicon is moving. The new set uses DEC's two-chip bipolar PMD design. To these, it adds a physical-layer controller chip and a single-chip media access controller, both in CMOS. The latter chip combines functions from three Supernet-1 packages, and adds additional hardware to support the still-solidifying station management technology (SMT) protocol. AMD claims that the new set allows customers to implement an entire single-attach station on an AT bus half-card.

This small an implementation, of course, relies on the workstation's host processor for most of the SMT and all of the transport-layer processing. "We see people moving away from the idea of a dedicated

40 MB/SEC
DUAL PORT VME/VSB MEMORY

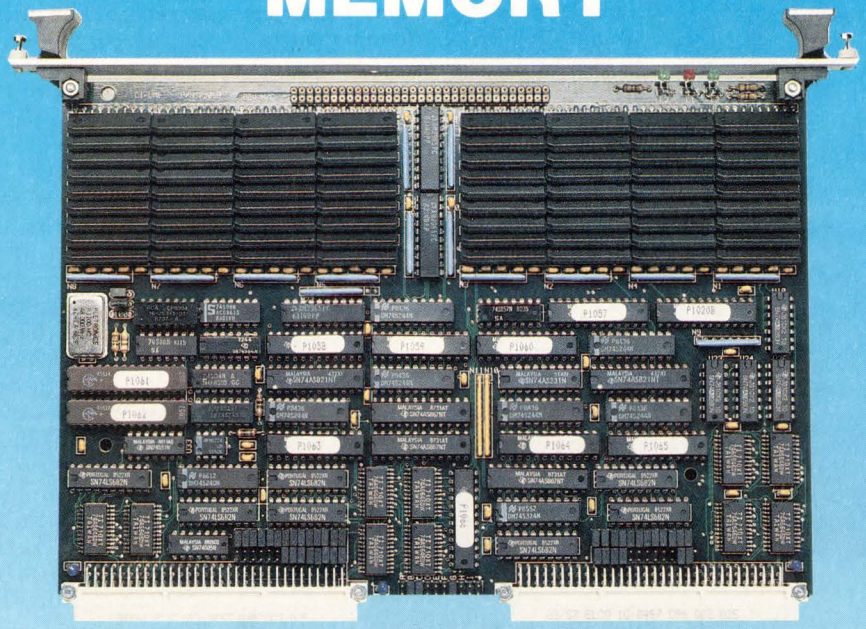

THE CI-VME40 FEATURES

ACCESS TIMES: 20/150NS
CYCLE TIMES: 95/195NS (BLT/SINGLE CYCLE)

- 4MB, 8MB, 16MB, 32MB, 64MB in one VMEbus/VSB slot
- Byte Parity Error Detection
- Addressable in 24 or 32 bit through 4 Gigabytes
- Memory start and end addresses selectable on 256KB Boundaries configured independently

ALSO AVAILABLE FOR THE VMEBUS ARE...

THE CI-VMEmory FEATURES

- 4MB, 8MB, 12MB, 16MB, 32MB, 64MB in one VMEbus slot
- VME Revision C compatibility
- Lower and upper memory addresses independently selectable in 64K byte increments
- Byte Parity Error Detection with Selectable Trap on Parity Error
- Control Status Register

THE CI-VSB-EDC FEATURES

- 4MB, 8MB, 12MB, 16MB, 32MB, 64MB in one VMEbus/VSB slot
- Dual VMEbus/VSB ported board
- Error Detection and Correction circuitry
- Addressable in 24 or 32 bit through 4 GIGABYTES
- Selectable on 256KB, boundaries
- VME and VSB configured independently

Chrislin Industries, Inc.

Call Toll Free: 800-468-0736 (pst.)

31332 VIA COLINAS #106
WESTLAKE VILLAGE, CA 91362
FAX NO. 818-991-3470
TEL: 818-991-2254

CIRCLE NO. 39

Advanced Micro Devices marketing managers Robert Sykes (left) and Basil Alwan show a complete single-attach station on an AT-bus card. As AMD moves into its second generation of Fiber Distributed Data Interface controller sets, the company is placing increasing emphasis on integration to reduce adapter cost.

node processor," claims AMD product marketing manager Basil Alwan. "It's been kind of a religious argument in the industry, but increasingly people are running everything on the host." AMD division marketing manager Robert Sykes adds, "It's moving that way because the power of workstation CPUs has gone well past what's needed to handle the protocol, and the appearance of such buses as EISA and Micro Channel gives machines enough bandwidth to handle the traffic."

Cogitating the protocol

If the use of a dedicated protocol CPU is a religious issue, then AMD customer CMC is clearly from a different sect. "Maybe the low end of the market, where all you're buying is connectivity, can get by with little or no processing power on the board," says CMC's Soriano. "But I'm afraid people who get into FDDI with that kind of solution will be disappointed—you'll still be spending an awful lot of money, but for a minimal improvement in throughput. And because of the way flow control works, one slow node can be a burden on all the other protocol stacks on the network. That leads to a lot of finger pointing, and to eventual segregation of the slow devices onto a separate net."

Several companies are trying to untangle the CPU cost/performance dilemma by driving down the cost of a powerful dedicated processor. One of the common themes in last fall's round of embedded RISC announcements was the opportunity to sell to the FDDI adapter market. Highly integrated Sparc, MIPS or 80960 CPUs could cut down on the amount of support logic on the adapter card, and could help minimize the number of fast signal runs that have to be laid out. But the fast CPUs by themselves can't help with the biggest cost of the protocol-processing hardware—the memory. In fact, the powerful chips could exacerbate the problem by substituting fast SRAM for cheap DRAM.

An alternative would be to cut down on the cost by reducing the processing load rather than by pumping up the CPU. In this vein, Soriano suggests that the solution may be more ASICs to help with the protocol processing. XDI's Perloff agrees. "The silicon vendors are coming under a lot of pressure to help out—there's a feeling that adapter vendors can't deliver the cost/performance they need unless the chips work harder." Perloff points out that in relatively high-priced VME adapters, the designer can add ASICs to do some speed-critical protocol functions, such as header splitting and check-sum computation. But in a really low-cost adapter, that luxury may be out of reach, and the FDDI chip set may have to shoulder the burden.

Silicon vendors resist that idea for practical reasons. They point out—accurately—that there's no agreement on transport protocols for FDDI now. In this situation it would be almost impossible to design hardware assists into the controller chips without making the whole controller protocol-specific—which is a financial impossibility.

"We're trying to look at the problem from a systems-design point of view," says Rhonda Dirvin, communications VLSI marketing manager for Motorola. "We know that we can incorporate into our controller chip set a lot of FDDI-specific features that will reduce the need for Mips in the CPU. But when we talked to prospective customers, we found people using ISO, TCP/IP (Transmission Control Protocol/Internet Protocol) and even some XTP (express transport protocol). So it was clear we couldn't lock any one protocol into the silicon."

If more-powerful CPUs are still too expensive, and if hardware as-

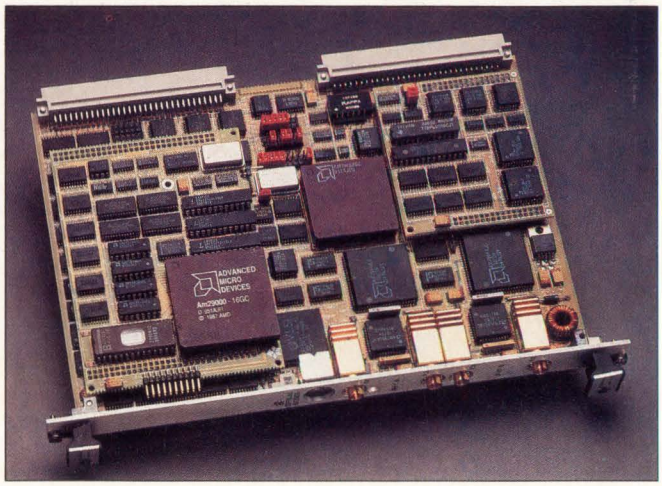

The 4211 Peregrine FDDI (Fiber Distributed Data Interface) adapter from Interphase typifies the latest high-performance boards for the 100-Mbit/s net. To achieve the throughput that FDDI promises, vendors must pack a controller chip set, a 32-bit RISC CPU and a lot of memory onto the board.

Debate over frame-based SMT functions continues

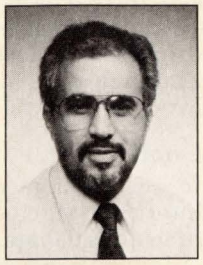

Now that standard hardware is available for the FDDI (Fiber Distributed Data Interface), designers are beginning to look seriously at integrating FDDI networks into their systems. Still, disagreement over FDDI's station management technology (SMT) functions remains a stumbling block. Until these differences can be resolved, concerns over interoperability and performance will hinder the deployment of FDDI networks.

So that a station may work cooperatively on a ring, SMT functions provide the station-level control necessary to manage the processes underway in the various FDDI layers. These layers include the PHY (physical layer protocol), the PMD (physical media dependent) and the MAC (media access control).

SMT services include connection-management functions such as station insertion and removal, station initialization, and configuration management, and fault isolation and recovery functions such as error detection and station isolation. SMT also provides configuration partitioning, a communications protocol for external authority, scheduling, station statistics collection, address administration, downline load, and upline dump.

Recently, the ANSI committee responsible for FDDI standardization submitted the FDDI standard for letter ballot. Preliminary results indicate solid agreement on FDDI's PHY, PMD and MAC layers. Agreement has also been reached on lower-level SMT functions such as ring management and connection management, which ensure that the ring comes up satisfactorily.

The primary hurdle in solidifying the standard is disagreement over SMT's higher-level frame-based functions, which are responsible for passing frames to and from other SMT nodes. At issue is how much of the SMT should be standardized, including the number of frame classes, the kind of information that should be included in frames, and the protocol that should be used to transmit frames.

One disputed issue is whether frame-based services should include a protocol that enables an SMT to ask its node questions such as whether frames have been transmitted or received, the error count, and how long the station has been active.

The issue is how much of the SMT should be standard.

Another disputed issue is what type of action the SMT should take over a failed transmission or other fault. For example, should it transmit a frame after a certain number of frames have been lost?

A third area of debate is over what type of information should be stored in each station's management information database, which stores the operational parameters needed to set up frames and package data. The SMT functions access this database prior to transmitting a frame.

Philosophical differences

Much of the dispute over which frame-based functions should be included in the SMT can be attributed to a basic philosophical difference over the role of SMT in an FDDI network. In one camp are those who believe that true interoperability is best achieved by standardizing as much of the protocol as possible. In the other camp are those who believe that SMT shouldn't address higher-level protocol issues. Most of these services, they argue, are already available in standard protocols such as the Open Systems Interconnection and the Transmission Control Protocol/Internet Protocol, and shouldn't be duplicated in the SMT.

Until SMT issues are resolved, interoperability will be a serious concern of systems designers. Because so much of the SMT is still optional, SMT implementations can vary considerably, making it difficult to integrate FDDI equipment from different vendors.

In addition to affecting interoperability, the lack of SMT standardization can affect performance. This occurs because networks with incompatible SMT protocols are more likely to be unstable, thereby increasing downtime and reducing throughput. If instability causes the ring to be down 50 percent of the time, for example, then the network's maximum throughput will be limited to 50 Mbits/s.

By the end of 1991, the SMT standard should be stabilized. Until it is, however, the safest approach for systems integrators will be to use the same SMT implementation for all of the equipment in their networks. This, in turn, means purchasing the entire FDDI solution from a single vendor who supports a wide range of equipment, from personal computers and workstations to minicomputers and mainframes.

Juan E. Figueroa, FDDI marketing manager, AT&T Microelectronics

sistance from the chip set can only go so far, that leaves the protocol software itself as a place to look for savings. "You might be able to save about a third of the computing work by putting check sums and the like in hardware," says Perloff. "And you can save maybe that much more by just organizing the system to avoid gratuitous data movement. That leaves you with a core of necessary code that you can really beat on to improve performance." In TCP/IP, for example, Perloff says there are about 30 lines of code that are the subject of constant tweaking because they form the inner loop of the transport layer.

More-radical approaches could reduce the load even further. Protocol Engines (Santa Barbara, CA), for instance, is designing a protocol specifically for implementation in custom silicon. This should give the best throughput per dollar of any transport-layer solution. But there are trade-offs here as well. An adapter hard-wired for a streamlined protocol such as XTP might be unable to communicate with whole branches of a corporate network.

Nailing down the SMT

At the moment, the interoperability question asserts nearly veto power over many efforts to streamline the

transport hardware. Not only is it necessary to accommodate lots of different transport protocols, but there are growing unresolved issues at the next level up—station management software.

The FDDI standards people are working to nail down a particular SMT definition. Until they can stabilize the SMT, silicon vendors are reluctant to move too fast in supporting the protocol. Most second-generation chip sets, including the National and new AMD offerings, support the connection management portion of SMT in hardware now. This is partly because it's well defined, and partly because it involves transactions too fast to handle on a separate CPU. But that's not enough for adapter vendors.

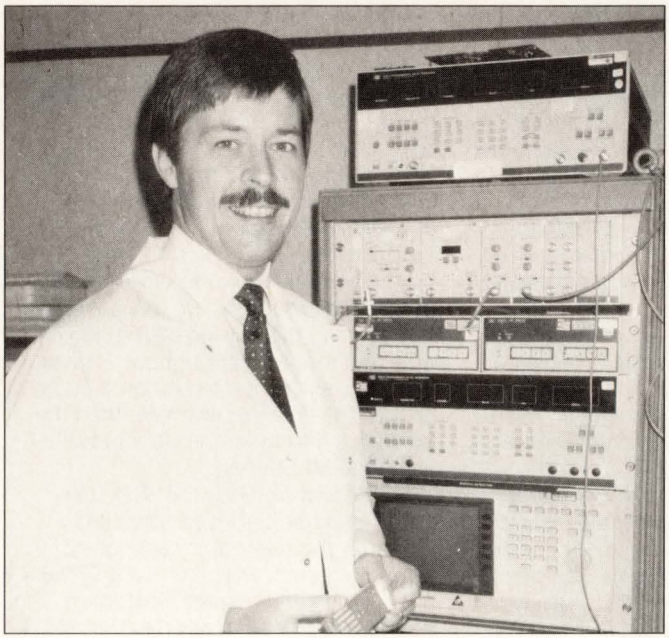

Hewlett-Packard marketing section manager Nate Walker examines a test setup for an FDDI (Fiber Distributed Data Interface) transducer. Walker says that the stringent requirements the FDDI standard places on the transducers stress both the manufacturing and the testing capabilities of vendors.

"We really need to get the whole SMT into silicon, for interoperability reasons," argues Jacob Hsu, product line manager for FDDI adapter products at Interphase (Dallas, TX). Meanwhile, Interphase, like other vendors, is moving more of the protocol load onto its resident processor. "We have a 16-MHz 29000 on the board, primarily for real-time housekeeping tasks," Hsu says. "But we're porting TCP/IP to this local processor. We found that customers with less-powerful hosts—68030s and the like—wanted us to off-load their CPUs as much as possible."

The need to provide a programmable environment for the changing SMT and to off-load protocol software dictates a big on-board CPU. Another consideration seconds the call. For the foreseeable future, FDDI rings will work in a heterogeneous environment, connected to all sorts of other networks.

"Connection to multiple types of LANs is important," says DEC marketing manager Karl Pieper. "This causes real problems with a low-level management facility like SMT. It can't cross bridges to other networks—in fact, you have to have separate management stations for two interconnected FDDI rings. So people need a tool that integrates network management over diverse nets—right now that means SNMP (simple network management protocol)." In fact DEC's FDDI products will be trilingual, supporting SMT, SNMP and DEC's own Remote Bridge Protocol. Like the other tools, SNMP runs at level three, above the chip sets. So it, too, requires CPU resources and demands a programmable local CPU.

From the cacophony of demands for computing resources, it's clear that FDDI can't move gracefully onto a modest workstation or PC without bringing its own CPU along. It's also clear that until the protocol and net management fields stabilize somewhat, there's little the chip set vendors can do to help—the job of cutting back on CPU and memory costs will be up to the RISC processor vendors. Having said this, it's time to move on to the third big area of chip cost on the adapter board, the optical components.

"At this point, the chips aren't the place to look for a big breakthrough in cost," says DEC's Pieper. "Shrinking the new chip sets further won't necessarily make huge savings in either cost or board space. The place to look for a breakthrough is in alternative media for the network."

Looking into the optics

In fact the search for cost savings in the network medium has started one of the most acrimonious debates in the industry. Adapter vendors charitably accuse their optical component suppliers of manufacturing naivete or, sometimes, of price gouging. Optical vendors respond that their hands are tied by ill-informed and artificially stringent performance requirements in the FDDI standard. The debate is raising everyone's interest in alternatives—less-expensive fiberoptic transceivers and twisted-pair wiring.

The problem lies in the original FDDI specifications, according to Nate Walker, marketing section manager at Hewlett-Packard's Optical Components Division (San Jose, CA). "FDDI's PMD specification pushed the limits of optical technology," he claims. "The 1,300-nm wavelength, the 2-km range at 125 Mbaud, the optical power budget, and the spectral and rise/fall requirements drive up the cost of the FDDI transceivers."

While Walker is optimistic that increasing volume will gradually lower the cost of the little nest of drivers, receivers, signal conditioners, LEDs and PIN diodes that make up the transceiver, he warns that big savings can come only from a compromise on some of the FDDI requirements. "There are 820-nm transceivers available for $30 to $40 in production quantities," he says, "and 665-nm plastic-fiber transceivers can cost less than $10 in volume. But these parts can't handle the range and speed demanded of FDDI.

"The plastic-fiber devices are used only in the 5-Mbaud range. But we have demonstrated FDDI eye patterns with glass-fiber 820-nm components—the only compromise would be in limiting the length of the connection." Yet, to Walker's chagrin, the FDDI standards group has turned its back on short-wavelength devices in favor of reduced specifications on 1,300-nm parts.

"Getting the price down on 1,300-nm components is just a

EMULATION MADE EASY!

In the time it would take to select an emulator from all of the different options, install a configuration that wasn't hostile to your environment, read the manual and study the tutorial, check the application notes and then learn the interface(s) *. . . you could have debugged your product with the MICE-III.*

MICE-III in-circuit emulators from Microtek. Easy-to-use 68000/20/30™, 80C186™ support. Source debug. Fully-qualified trace. High-speed download. A full megabyte of emulation memory. And on-line applications support. Just plug it in and go.

Microtek also has in-circuit emulators for 286™, 386™, 386SX/376™, i486™, Analog Devices ADSP-2101A, 2100, 2105 and other microprocessors.

So, if you'd rather be debugging than getting ready to debug, call Microtek. We've made emulation easy.

MICROTEK

The Leader In Development Systems Technology.™

MICROTEK INTERNATIONAL, INC. - Development Systems Division
3300 N.W. 211th Terrace, Hillsboro, OR 97124 • (503) 645-7333 • Fax (503) 629-8460

Australia: 61-3-7202399, Austria: 43-1-6271953, Belgium: 32-2-2523680, Denmark: 45-42-804200, France: 33-1-39611414, Germany: 49-89-8570000, Israel: 972-3-5447475
Italy: 39-11-503427, Hong Kong: 852-3-7443516-9, Japan: 81-3-7955171, Korea: 82-2-679-1161, Netherlands: 31-3403-91234, Norway: 47-2-289490, Singapore: 65-7468182
South Africa: 27-11-7863710, Sweden: 46-8-930000, Switzerland: 41-1-2010711, Taiwan: 886-2-7576939, Thailand: 66-2-2217040, U.K.: 44-276-685005

MICROTEK, MICE-III, MICE-V and The Leader... are trademarks of MICROTEK INTERNATIONAL. 286, 386, 386SX/376 and 80C186 are trademarks of INTEL CORP. 68000/20/30 are trademarks of MOTOROLA CORP.

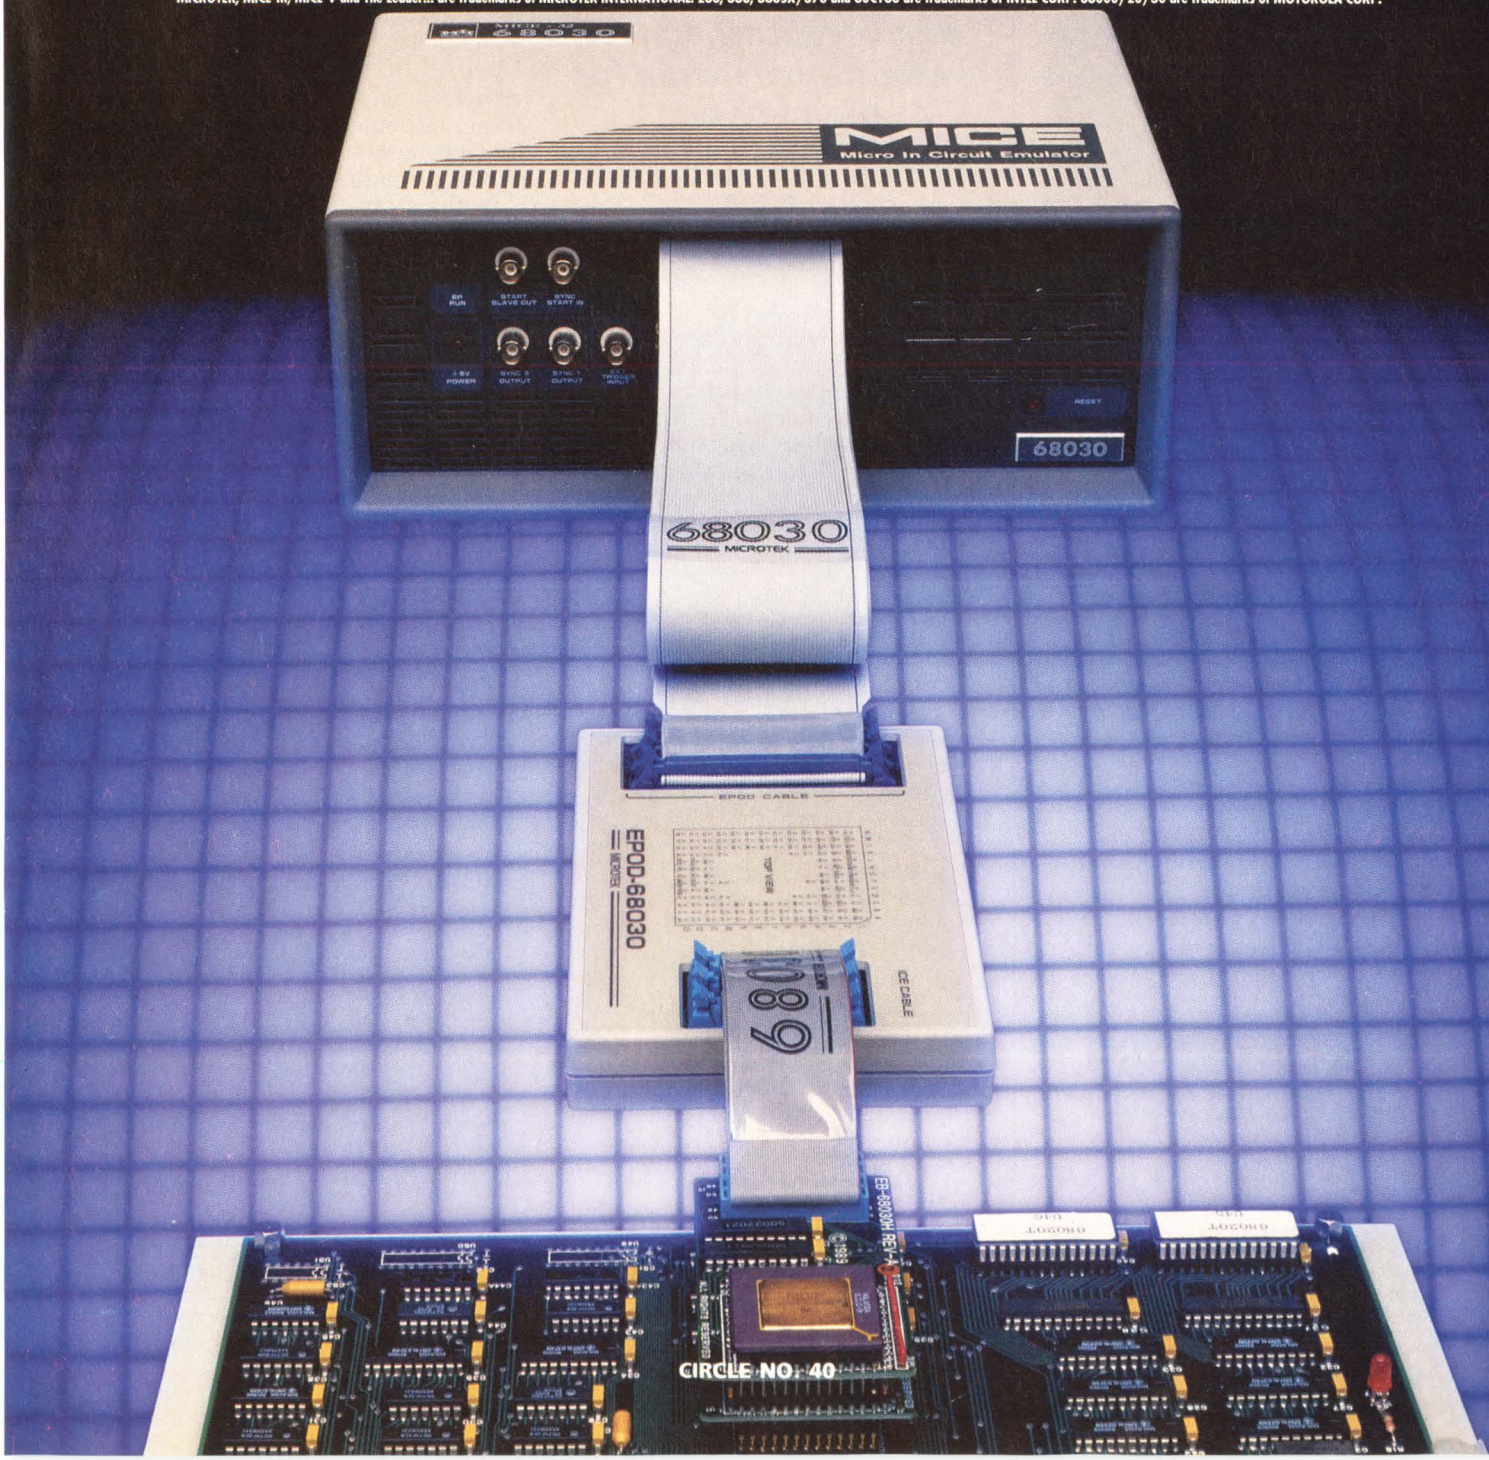

CIRCLE NO. 40

We'll Knock Your Socks Off and Pay You $100.

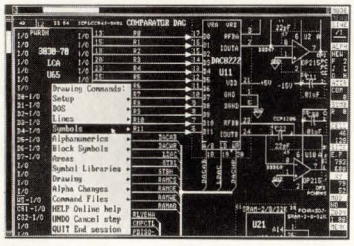

No matter what schematic editor you now use, Data I/O®'s new **FutureNet® 5 Schematic Designer** Software will give you more power and better results—in less time.

To prove it, we'll knock $100 off the already low price when you trade in your current schematic editor* for FutureNet-5.

NEW FutureNet-5, just $895.**
FutureNet-5 will knock your socks off with an enhanced design environment, dialog boxes, symbol browser, and other advanced features. An OrCAD translator is also available to translate OrCAD schematics and symbols into FutureNet-5's easy-to-use format.

We'll even **guarantee** your satisfaction within 30 days, or your money back.

1-800-247-5700

*Users of previous DASH™ or FutureNet versions from Data I/O may upgrade to FutureNet-5 for only $595.

**Price valid in U.S.A. only.

The Personal Silicon Experts

DATA I/O Corporation

CIRCLE NO. 41

matter of time. It doesn't make sense to switch to 850 nm for what will eventually be a $5 difference in cost," argues Schelto van Doorn, general manager of the fiberoptic components group at Siemens (Evansville, IN). "The problem isn't that the wavelength is wrong—1,300-nm components use more-stable indium phosphide technology and are much better able to handle the 125-Mbit/s data rate. The problem is that people started out thinking they wanted a Cadillac with a 2-km range and an 11-dB power budget. Now, they've decided they don't want to pay for it. If they'll settle for a 500-m range and a 7-dB power budget, they can have less-expensive components right now."

"The place to look for a cost breakthrough is in alternative media for the network."

—Karl Pieper, Digital Equipment Corp

■ ■ ■

But many users are losing interest in the debate over optical specifications. Instead, they're attending to a new concept, full-speed FDDI over a shielded copper twisted pair. The reason is simple. "Copper could be the answer to getting the price down," says AMD's Alwan. "You're talking $10 worth of passive components instead of $200 worth of electro-optics."

The feasibility of twisted pair has already been demonstrated, and the standards people are busily working on implementation details. But little issues—like whether to predistort the signal or postcompensate it to make up for the aberrations of the copper medium—could have a major impact on chip sets. "Since we've used the drivers from our Taxi parts in Supernet-2, our chip set can handle twisted pair now," Alwan claims. "We've already achieved acceptable error rates by using a postcompensation scheme in the lab. But to do the predistortion things that Cabletron (Rochester, NH) is proposing would probably require turning the chips again."

So the debate between fiber and shielded twisted pair could have major implications for controller IC vendors. Waiting in the wings is another possibility, a PMD specification for unshielded twisted pair. No one is even sure that this approach will work, but it's being investigated. This, too, would mean major changes in the chip set, since it would undoubtedly involve elaborate signal processing.

Taken together, the efforts to reduce the controller chip set, the protocol processor hardware and the transceiver components can probably take out half of the cost of today's FDDI adapters. Increasing volumes, as the streamlined adapters begin to appeal to more and more high-end workstation users, can probably accomplish another factor-of-two reduction. This should bring FDDI into the arena where it can win slots in bandwidth-critical nets—nets where graphics, large databases or distributed processing makes network speed crucial to throughput.

Yet it seems unlikely that FDDI can compete much below this level—Ethernet and token-ring solutions will be getting cheaper, too. So it's imperative that FDDI silicon vendors face the fact that they will always be living in a heterogeneous network, a world of rings and trees, of single-attach nodes and Ethernet loops, of multiple management protocols. In driving down costs and prying open bandwidth, the vendors must leave the flexibility and power for managing such diverse systems. ■

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
(800) 372-2447 . . . . . Circle 221
Advanced Micro Devices
(408) 732-2400 . . . . . Circle 222
CMC
(805) 968-4262 . . . . . Circle 223
Cabletron
(603) 332-9400 . . . . . Circle 224
Digital Equipment Corp
(508) 493-5111 . . . . . Circle 225
Hewlett-Packard
(408) 435-6010 . . . . . Circle 226
Interphase
(214) 919-9000 . . . . . Circle 227
Motorola
(512) 891-2000 . . . . . Circle 228
National Semiconductor
(408) 721-5000 . . . . . Circle 229
Protocol Engines
(805) 965-0825 . . . . . Circle 230
Siemens
(812) 422-2322 . . . . . Circle 231
XDI
(619) 487-9320 . . . . . Circle 232

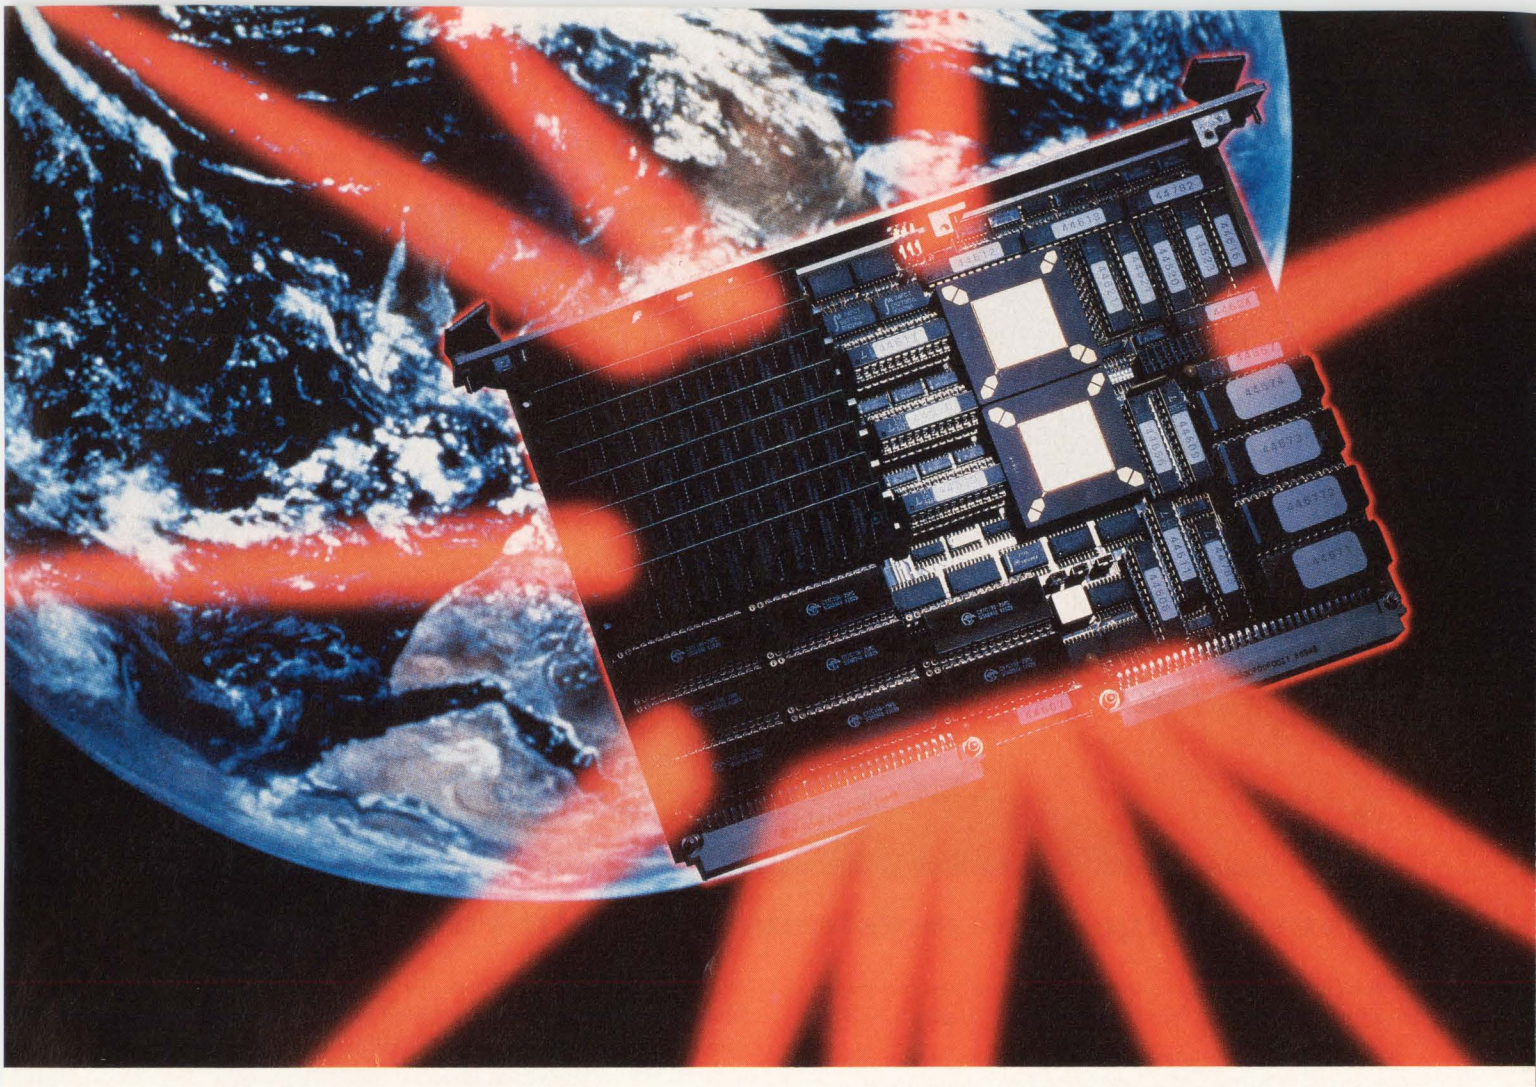

For Those Who Like It HOT

Announcing the MSP-6C30 66 MFLOP VME Array Processor

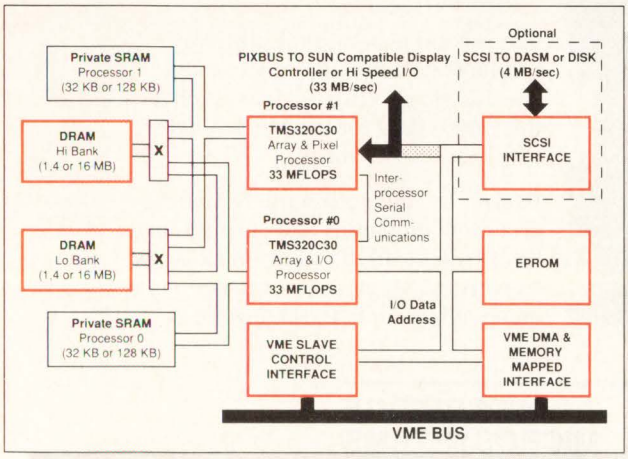

The MSP-6C30 is a 66 MFLOP, floating-point, VME array processor optimized for signal and image processing. High performance is assured through a unique, multi-banked, crossbar memory architecture, and high bandwidth is achieved with multiple, fast, industry standard I/O ports. Also available in a PC-AT version with optional VGA compatible display controller.

Features

- Under $100/MFLOP (single qty.)
- Dual TI TMS320C30 Processors
- 66 MFLOP, 32 MIPS Peak Rate
- Up to 32 Mbytes On-board Dynamic RAM
- High Speed Private SRAM
- Optional SUN Compatible Frame Buffer
- Optional SCSI I/O (disk, tape, A/D, D/A, Video, DR11-W)

Applications

- Medical
- Geophysics
- Defense
- Communications
- Digital Signal Processing

For Further Information Contact: Howard Cohen
Analogic Corporation/CDA Division
8 Centennial Drive, Peabody, MA 01960-7987
Tel. (508) 977-3030, (800) 237-1011, FAX (508) 977-9220

ANALOGIC®
CDA COMPUTER DESIGN AND APPLICATIONS

CIRCLE NO. 42

Whatever your destination on the open bus highway...

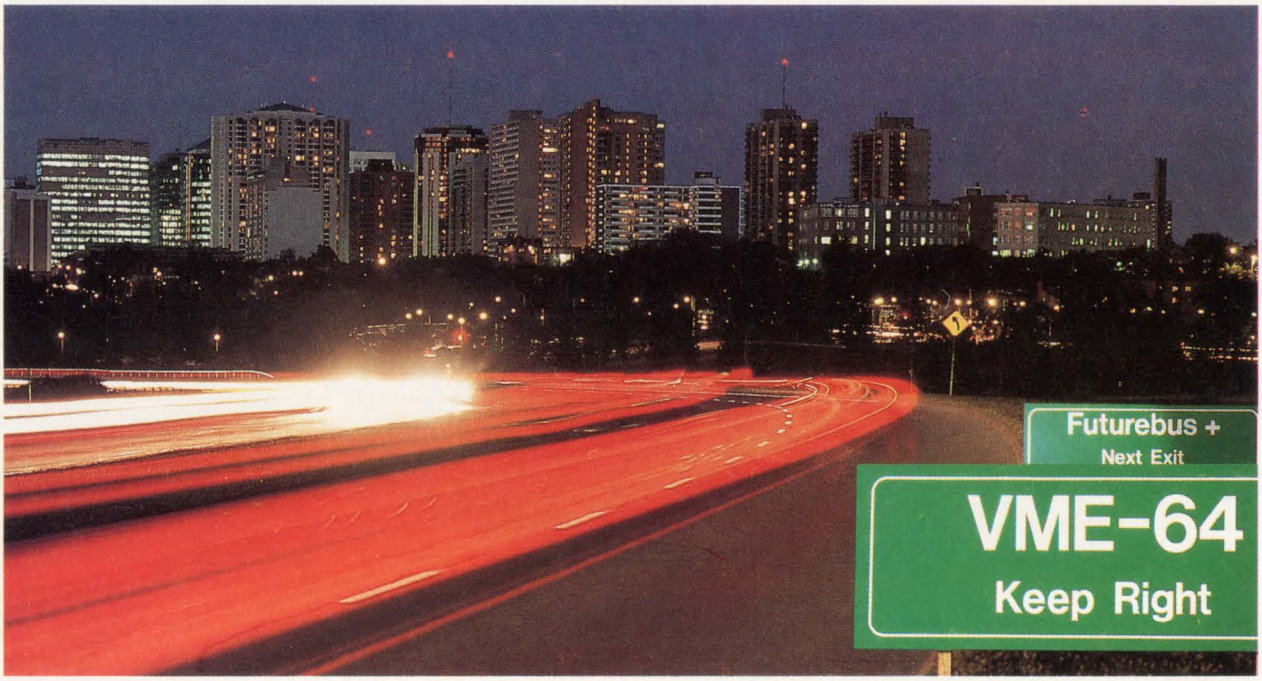

Newbridge Microsystems can help you get there.

Newbridge Microsystems offers the industry's most complete and flexible open bus silicon product line.

Leading the way with our 32-bit VME chipset that supplies a full high performance VMEbus interface, we provide a pin-for-pin compatible migration path to VME64. Our product direction extends beyond VME to devices and VHDL models for emerging Futurebus+ protocols.

The CA91C014 and CA91C015 VME chipset provides:

- IEEE 1014/VMEbus Rev C compliance
- Full VMEbus interface implementation
- High performance VMEbus decoupling with on-chip DMA functionality
- Bus Isolation mode for system diagnostics and fault isolation

For those applications demanding even higher performance, the CA91C064 VME64 device will offer users of our 32-bit chipset a pin-for-pin compatible upgrade to VME64.

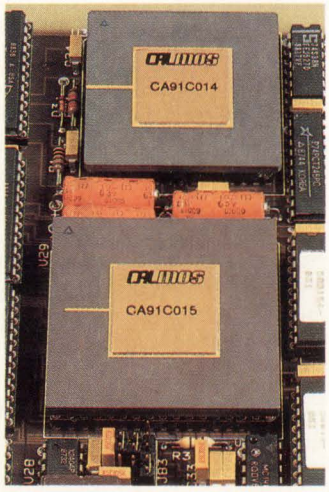

Newbridge Microsystems is also developing a highly flexible solution for entering the Futurebus+ market quickly. The CA91C896 and CA91C897 Futurebus+ chipset will feature:

- Generic local bus and decoupled Futurebus+ interfaces to support split transactions for multiprocessing and burst transfers for higher performance
- Support for scalable external performance enhancements for use across a spectrum of applications
- Portable VHDL models to minimize design effort — and your time to market

Our second generation Futurebus+ products will include processor specific single chip devices, such as the CA91C880 for the Motorola 88K RISC processor family.

Backed by the worldwide Newbridge group of companies and an international sales and distribution network, Newbridge Microsystems is laying out an open bus route to the future.

NEWBRIDGE

NEWBRIDGE MICROSYSTEMS

Newbridge Microsystems • A Division of Newbridge Networks Corporation.
603 March Road, Kanata, Ontario, Canada K2K 2M5
Tel: (613) 592-0714 • 1-800-267-7231 • Fax: (613) 592-1320

Newbridge and logo and Calmos are trademarks of Newbridge Neworks Corporation. © Copyright 1990, Newbridge Networks Corporation.
Other brand and product names are trademarks or registered trademarks of their respective holders.

CIRCLE NO. 43

Tools bring vital information to early stages of systems design

With shrinking design times and growing systems complexity, designers are relying on tools that predict component behavior before prototyping.

Mike Donlin
Senior Editor

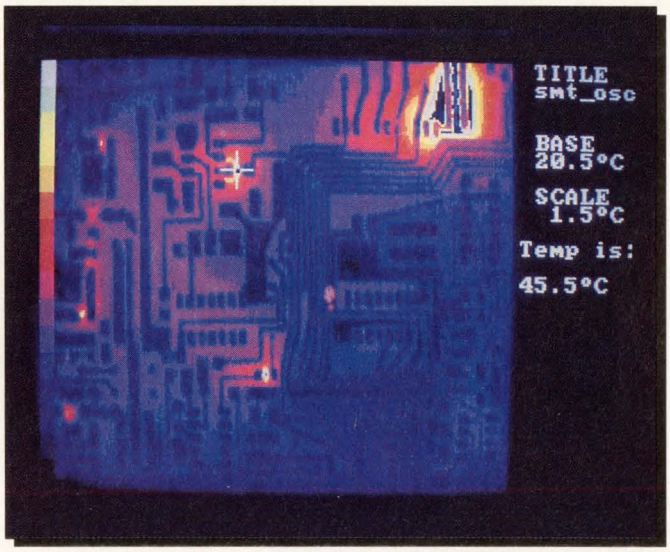

The Compix 6000 thermal imaging system lets design and quality assurance engineers measure the heat of components and printed circuit boards at prototype stage to detect potential heat problems before manufacturing. This thermal image of a circuit board shows a surface-mount device with possible heat problems.

Computer manufacturers are beginning to change the way they design systems. Problems that were once left until after the prototype stage of the design cycle are being moved forward to the early stages of design capture. And EDA vendors are fueling this transition with tools that let designers predict the electrical and mechanical behavior of ICs, circuit boards, cooling systems, and cabling layout before committing their designs to a physical prototype. In circuit board design, for instance, characteristics such as thermal effects, electromagnetic interference and airflow across components are being considered during layout and routing as opposed to the traditional methods of relying on numerous iterations of a prototype. Even manufacturing issues, which have often been addressed only in the last phases of a product design, are being considered during the initial phases of systems design.

"It's essential to address the manufacturability of a product as early as you can in the cycle," says Bryan Floyd, executive manager of strategic planning and marketing for the mechanical design engineering and manufacturing group at Intergraph (Huntsville, AL). "If you design a backplane that hits a brace in the cabinet, for instance, you've got a problem, but its not catastrophic. But if you make a prototype circuit board that can't be manufactured, it's disastrous. We've addressed that here by designing our manufacturing facilities so that they can produce the prototype. If it can't be manufactured, it won't exist as a prototype."

The driving forces behind this revamping of design capture are time-to-market and cost. Design cycles that stretch out beyond the predicted life of the system are unacceptable, and these are a direct result of numerous iterations of ICs, circuit boards and peripherals. Cost, of course, goes hand in hand with time-consuming rework. "Companies are getting better information about costs on the manufacturing floor," says Frank Binnendyk, product marketing manager for the simulation and test division at Mentor Graphics (Beaverton, OR). "When they start to analyze that data, they realize that the best solution is to move as much information as they can to the early stages of design."

Restructuring design capture is not simply a function of the available EDA tools; it also must encompass the underlying culture of a systems or IC house. Even if tools and frameworks are put into place, there must be

a commitment from those that use the tools and their managers to share critical information to ensure that at the end of a design cycle there are no nasty surprises. "One of the key indicators that there is a transition going on is in the restructuring of how different disciplines are managed," says Binnendyk. "Many companies are moving test engineers from under the arm of manufacturing up into the design segment. Having test engineers work side by side with design engineers is not just a technical decision, it's a managerial one. The end result is that the finished prototype will be testable, an event that wasn't always the case."

Because design and test data are often generated by incompatible tools, there has traditionally been a problem when it's time to generate test programs for prototype and manufacturing. To alleviate this incompatibility, EDA vendors are working with systems and IC houses to develop tools that will work toward a common goal—a product that can be tested during manufacture. "In the past, the design team would pass off their work to manufacturing, while the test group would try to develop a test program using their own tools," says Irv Christy, product manager for the manufacturing test division at Hewlett-Packard (Colorado Springs, CO). "If they saw something they didn't like, they'd go to the design group and try to get some information, but because the tools were incompatible, there would be a conflict. The communication wasn't very good."

Vendors teaming up

Mentor Graphics and Hewlett-Packard worked together to develop Quickfault, a deterministic fault simulator, and Quickgrade, a high-speed fault grading system. Quickfault provides an independent assessment of test patterns before they're used in the manufacturing process. This ensures that the test patterns are comprehensive in finding manufacturing defects. The tool automatically analyzes the results of the fault analysis and displays the fault location on the schematic. The results can be displayed at any level of the design hierarchy and for any selected portion of the design. Quickgrade uses statistical methods to predict total test coverage, provide information for developing test patterns and aid in the selection of test points.

This melding of design and test disciplines is proving valuable to companies like HP, who realize that design engineers aren't necessarily concerned with test engineers' needs in the early phases of a project. "In order to properly test a circuit board, a test engineer has to initialize it," says HP's Christy. "But while a design engineer can just assume a certain state for simulation and continue with the design verification, a test engineer can't. We've had situations where the test engineer saw that the traces necessary for initialization were missing from the design and was able to make sure that they were brought over to the edge connector. If he hadn't been a part of the design phase, the board would have been untestable."

Considering testability early in the design cycle is also becoming more critical for automatic test equipment (ATE) manufacturers. As circuit boards are packed with surface-mount, custom and standard package devices, the traditional bed-of-nails board tester has problems with irregular pin spacing and component density. Surface-mount parts, for instance, cannot be tested from underneath the board unless a via is routed for the test nail. Decisions have to be made during placement and routing to ensure that critical pins are tested. "Think of the nails on a tester as a scarce resource," says Peter Hansen, manager of board test applications at Teradyne (Boston, MA). "A board with 2,000 nets typically has about 5,000 device pins on it. The problem is that the person laying out the board doesn't necessarily know what the test engineer needs to ensure a test pin at critical points. Our figures indicate that if a designer makes careful decisions, the resulting board is 10 times more testable than if ATE is ignored."

To aid the design and test engineer in making these decisions, Teradyne has developed a software package that categorizes the device pins according to how critical they are to accurate testing. "These decisions aren't simple black-and-white choices," Hansen says. "There are some nets with boundary scan parts, for instance, and these don't need a test nail. There are decisions that aren't as easy, where you can leave off a nail, but you'll lose the ability to test resistors. These are decisions that have to be made among the design, test and manufacturing engineers. In the '80s we tested virtually every net on the board. In the '90s it's going to be impossible to do that, and the question is how to decide which nets have nails and which ones don't. It's

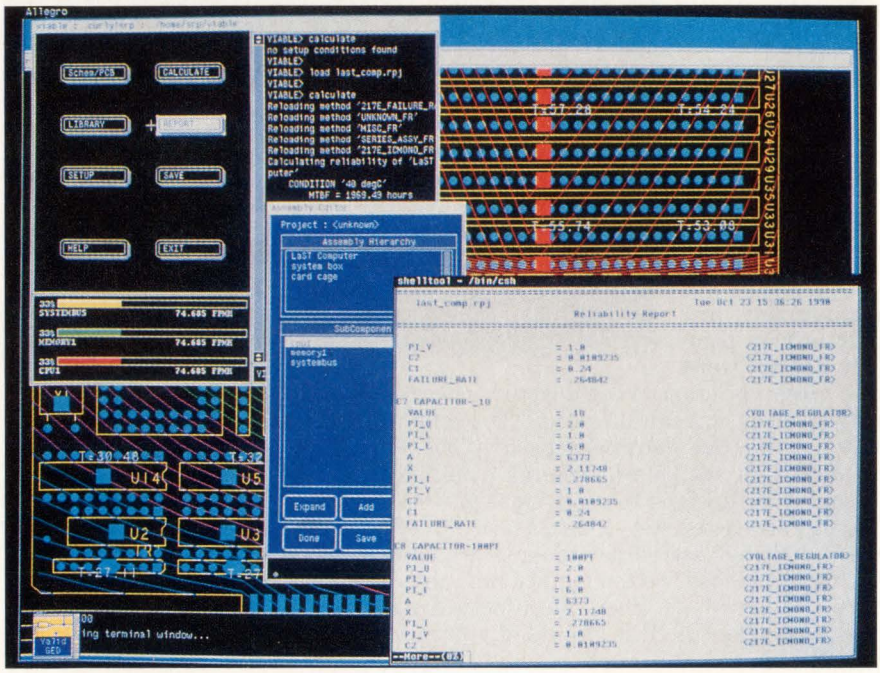

Providing component reliability information early in the design cycle lets engineers examine the consequences of design decisions. The Viable reliability analysis tool from Valid Logic Systems uses information from schematics or layout designs and MIL-HDBK-217D/E standards.

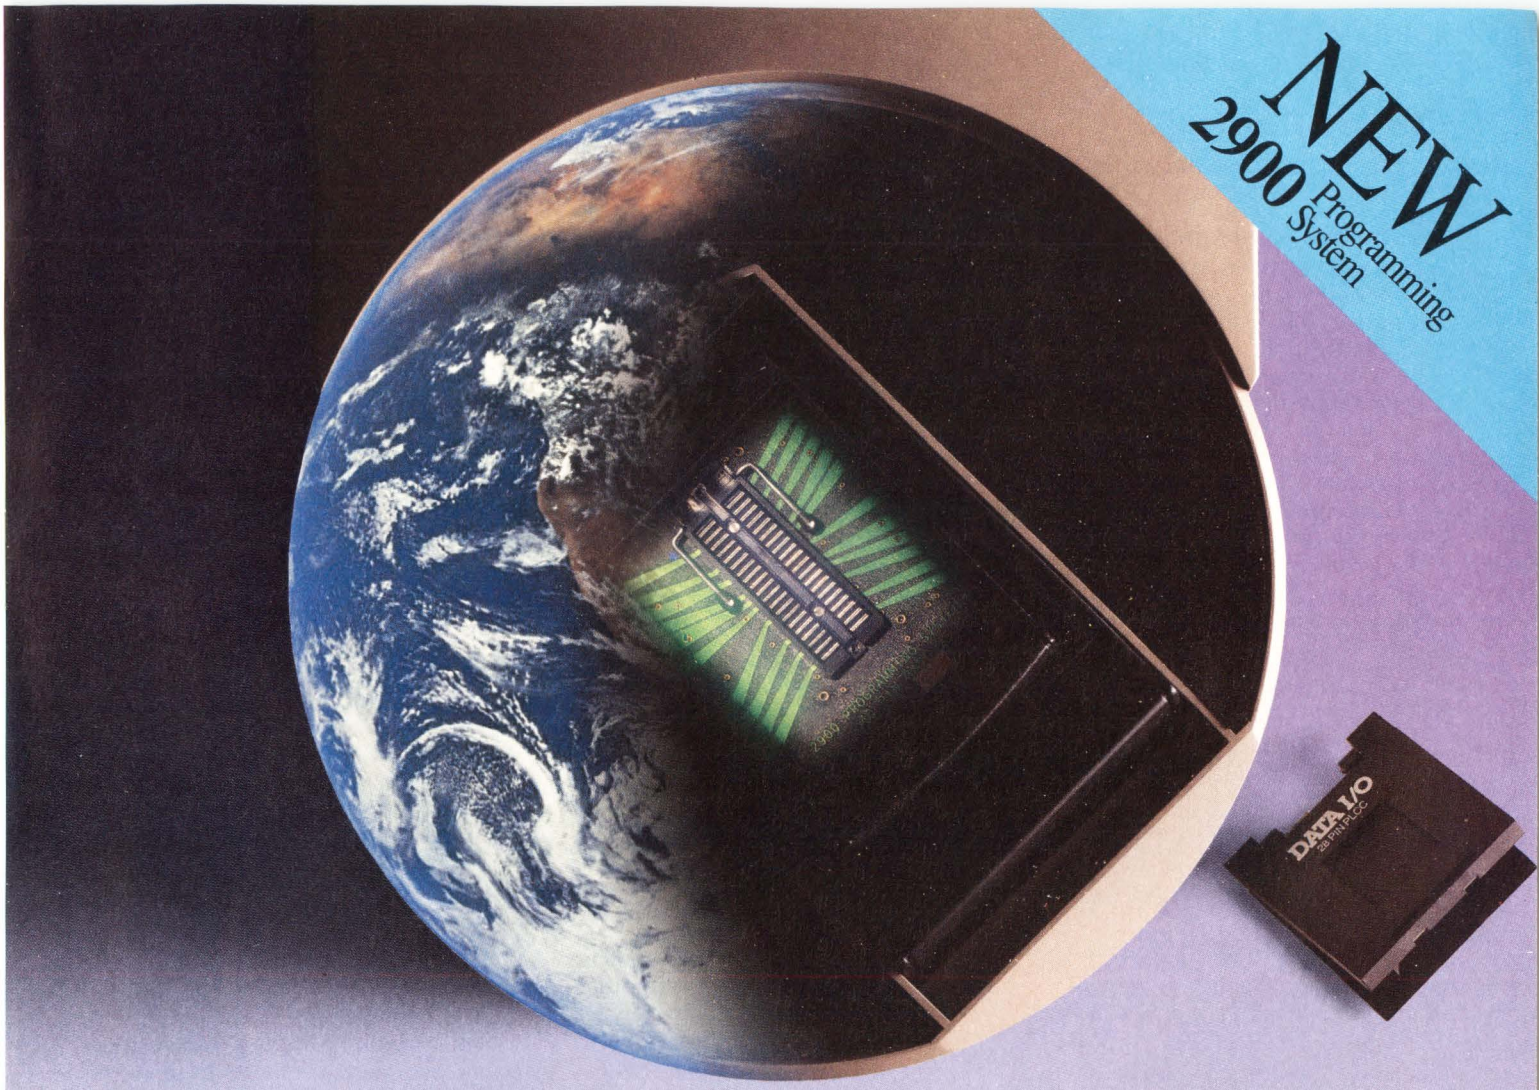

Out-of-this-world performance at a down-to-earth price.

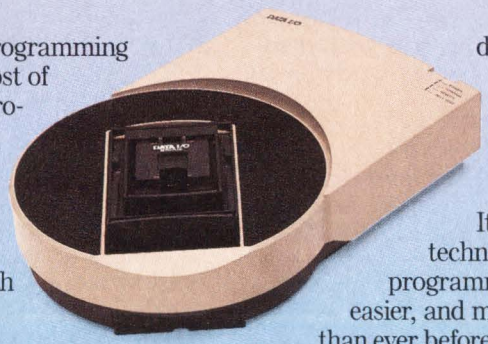

The new 2900 Programming System brings the cost of high-performance programming down to earth. Buy only the device libraries you need today and expand capability when you need it with simple software updates.

The 2900 supports virtually every programmable logic and memory device on the market—even surface-mount packages. Its innovative technology makes programming faster, easier, and more affordable than ever before. And with Data I/O®'s industry-standard design and testing software, you can create a complete PLD development solution.

For more than 15 years, Data I/O has set the standard in device programming. Call today to learn how the 2900 is setting a new standard for both price and performance.

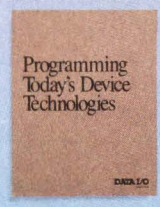

Call today for a FREE tutorial.

1-800-247-5700

Data I/O Corporation 10525 Willows Road N.E., P.O. Box 97046, Redmond, WA 98073-9746, U.S.A. (206) 881-6444/1-800-247-5700
Data I/O Canada 6725 Airport Road, Suite 302, Mississauga, Ontario L4V 1V2 (416) 678-0761
Data I/O Europe World Trade Center, Strawinskylaan 633, 1077 XX Amsterdam, The Netherlands +31 (0)20-6622866/Telex 16616 DATIO NL
Data I/O Instrumatic Electronic Systems Vertriebs GmbH Lochhamer Schlag 5A, 8032 Graefelfing, W. Germany, 089-85-85-80
Data I/O Japan Sumitomoseimei Higashishinbashi Bldg., 8F, 2-1-7, Higashi-Shinbashi, Minato-Ku, Tokyo 105, Japan
(03) 432-6991/Telex 2522685 DATAIO J

©1990 Data I/O Corporation

DATA I/O Corporation

CIRCLE NO. 44

too complex a problem without some software assistance."

Thermal considerations

In addition to testability concerns, systems vendors are examining the quality of their finished products by considering thermal effects of densely packed components. By moving thermal analysis to earlier stages of the design cycle, decisions can be made that affect reliability and manufacturability of circuit boards and peripherals. "Thermal analysis of circuit boards used to occur later in the design cycle, after the prototype was finished," says Bill Hostmann, product manager in the design creation group at Mentor Graphics. "Sending the prototype out for thermal analysis could take anywhere from two to three weeks. This, of course, would affect other elements of a system's design such as location of cooling fans, placement of boards in the card cage and cable routing. By using thermal models, these concerns can be addressed much earlier than before."

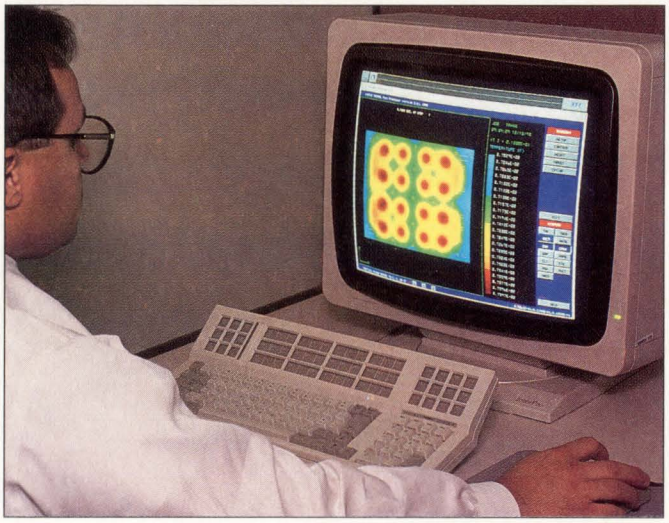

Thermal analysis tools, such as Hybrid Thermal from Intergraph, provide software models of potential hot spots on circuit boards, hybrids and ICs. By analyzing temperature models of components and connecting materials, the tool provides valuable data for component selection and layout.

Mentor's thermal analysis tool, Autotherm, is a finite element-based tool that creates thermal models from the Board Station circuit board design tool. Models include components, circuit boards, thermal rails, and heat sinks that can be used during layout to analyze thermal consequences of component placement.

Even with the sophistication of thermal models, some systems designers still require a post-prototype thermal analysis based on infrared images of the finished circuit board. Until recently, this meant purchasing a thermal imaging system that could cost well over $50,000, or waiting several weeks for outside analysis. Today, there are lower cost alternatives that produce thermal portraits that can be stored on PCs for analysis and archival purposes. Compix (Tualatin, OR) has developed a thermal analysis tool for under $20,000 that measures temperatures ranging from 17° to 150°C and records variations as low as 0.2°C.

Such thermal information not only lets systems designers make architectural decisions earlier in a cycle, it also helps control costs. "Certain component package materials, like ceramic, dissipate heat better, but are also more expensive," says Bill Johnson, president of Compix. "In the absence of real data, a cheaper package, such as plastic, might be chosen, but it could ultimately affect the reliability of the product. With real thermal data, decisions can be made about whether to cut costs or live with the more expensive option. It allows companies to design closer to real-world limitations."

Predicting reliability

Because the detrimental effects of component overheating eventually lead to system failure, thermal effects prediction and analysis leads naturally into reliability prediction. There are several reliability prediction packages on the market that use thermal characteristics, as well as other factors to calculate mean time between failure statistics for circuit boards, hybrids and multichip modules. "A designer needs as much information up front as possible," says Dick Albright, director of marketing for CAE products at Valid Logic Systems (San Jose, CA). "A lot of the environmental factors that were either considered late in the design cycle or ignored altogether are being brought in as early as possible. Thermal characteristics, noise analysis and parasitic effects not only cause problems with prototypes, they also show up in the field with disastrous consequences. Tools that examine these effects are indicative of a trend today to know as much as possible before committing to a total design."

Valid's thermal modeling tool, Thermostats, is integrated with its Allegro circuit board design system. The tool allows designers to run a suite of thermal, reliability and noise margin analyses after preliminary component placement has been determined. By answering questions about the board's environment, such as spacing between boards, the velocity and direction of airflow and the ambient air temperature, the user gets an automatically generated thermal model.

More recently, Valid has introduced its reliability analysis tool, Viable, which lets engineers predict the reliability of designs at the component, board and system level. Viable derives its predictions based on device operating voltage, current, power, temperature, duty cycle, and manufacturing processes. The tool then provides MTBF, as well as relative and absolute failure rates for components using the MIL-HDBK-217D/E standard.

Managing the team

Even though systems vendors are using tools to bring thermal, mechanical and testability concerns earlier in the design cycle, the problem remains of how to manage the individual design teams, so that everyone is working with current information. If the team designing the disk drive, for instance, makes a decision that affects the mechanical and power supply teams, it's essential that the change is made known as soon as possible. "One of the first questions that large systems houses ask us is not 'What can we automate next,'" says Isadore Katz, manager of business development, systems division at Cadence Design Systems (San Jose, CA), "but how can we improve the quality of the design process. That means identifying the bottlenecks that cause unnecessary iterations of any portion of the design. We have to look at communication between design groups as well

Graphical design environment builds on framework

The most pressing issues for systems designers are time-to-market and reliability. Concurrent engineering helps address those issues by coordinating the simultaneous work of all disciplines in systems design. All engineers responsible for a product can thrash out issues early in the design cycle, before significant expense has been committed.

Many companies, however, are still in a linear "over the wall" mode of product development. Concurrent engineering is a new paradigm for them. How do they implement it, control it, and measure its progress?

The answer is based on systems-design capture that starts at the earliest requirements stage and works throughout product development. This is accomplished with graphical, multidisciplinary design capture and new analysis tools embedded in a design automation framework.

This approach raises the level of analysis, so better specifications and designs are created earlier. It also enables design iterations prior to simulation, so designers can quickly evaluate evolving specifications and designs before initiating exhaustive verification efforts.

Start at systems level

Complex systems are typically specified and designed at the systems level with little or no automation. As a result, specifications are likely to change, negating existing design work and introducing the likelihood of errors. What's more, designers have few direct means for making sure that designs are actually implementing specifications until after much of the detailed design work is done.

The start of graphical design capture, therefore, is at the systems level. This level becomes the glue that pulls together all aspects of systems design: hardware, software, communications, project management, reliability, safety, and ergonomics. At this level, the systems designer polls users for a set of requirements, checks those requirements for consistency and feasibility, and then generates functional specifications from those requirements.

To accomplish this, systems designers need a method that accounts for all design information, and a language that all project members can understand. Today, commercially available tools can provide an environment and a method that lets multiple engineering disciplines develop a common system description in an on-line, shared database.

The automated systems-design tools work as follows: Designers enter existing requirements into a database to describe the system. They generate an initial graphical representation of the system, and add information in one of a variety of modes of expression. They submit this representation to analysis, including dynamic system simulation. Finally, they generate reports that serve as complete specifications for all subsequent design work.

Systems-level capture and analysis can provide a new level of automation to the design process. It enables all-important traceability, producing audit trails from each specification to the components that satisfy it. Systems designers can carefully control and evaluate changes. Equally important is executability to ensure that specifications are accurate and relatively stable. By initiating analysis very early in the design process, systems designers can lessen the likelihood of interruptive design changes later on.

Systems designers need a method that accounts for all design information, and a language that all project members can understand.

Systems-level capture streamlines the rest of the design process. Designers can minimize front-end errors by making sure that all requirements have associated functions and by exploring all avenues for failure. Overall, systems-level capture helps engineers work smarter, particularly in eliminating drudgery.

The verified specifications need to be broken down, however, into an architecture with hardware, software and mechanical components. As the architecture is created and detailed, designers must verify that it's matching the specifications. This verification is accomplished with traceability—from each component back to the specification that spawned it—and static constraint analysis using a new category of design tool.

Expressing design views

To fully detail the architecture, the graphical design environment needs to be able to express all views of a design—hardware, software and mechanical. It expands from the province of schematic capture for physical layout to the graphical representation of the entire system. Once all different views are established, the graphical design environment helps the design team decompose each view into lower levels of hierarchy until the views consist of pretested components from qualified design libraries.

Each view has a different set of parameters: hardware comprises devices and nets, whereas software has interfaces with associated functions and calls. Each view requires a specially tailored editor, with all editors operating simultaneously in windows in the environment. The editors must work within the same database system, so all views are consistent and design changes are immediately reflected in all relevant views.

To assist in evaluation of design alternatives, analysis tools check the parameters of the evolving system's design, taking a snapshot of its status. These constraint-analysis tools use all available design information to estimate performance in terms of static issues, such as weight, cost and size.

A design automation framework is the enabling technology for systems-level capture and analysis. It supports the design data, data management and high-level analyses that coordinate multidisciplinary systems design.

For all levels of analysis, the underlying elements are the component libraries. Here, component means any predefined and pretested building block, ranging from integrated circuits to software subroutines and even standard bus connectors. The libraries provide the raw data about existing designs and components to speed capture, constraint analysis and subsequent simulation.

Design technology still needs to evolve to support systems-level capture, but some of the pieces are already in place.

Bill Hostmann, MS in computer science and engineering, product manager for design creation, and **Gordon Hoffman,** MSEE, director of systems design, Mentor Graphics

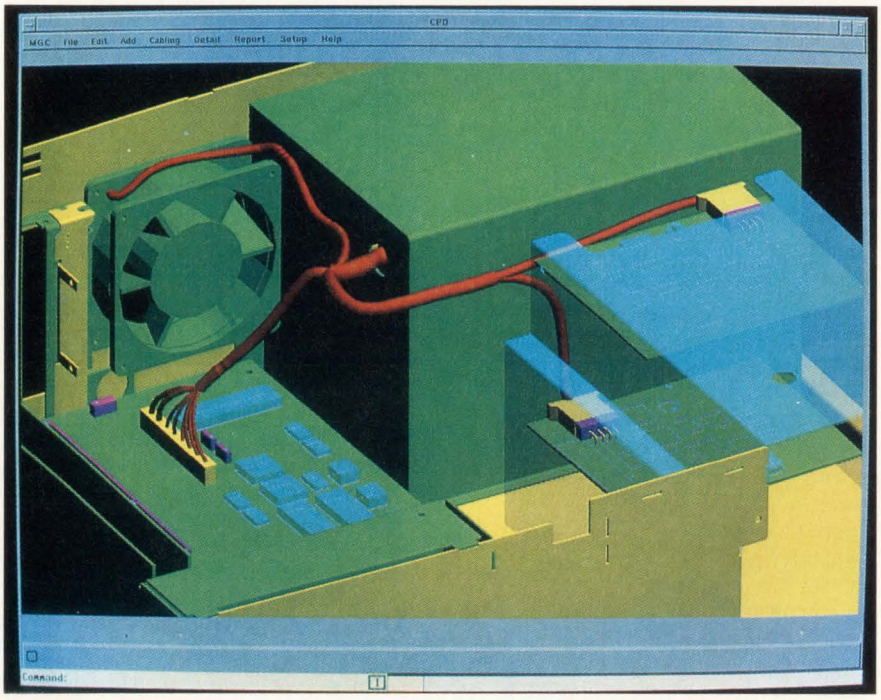

Mechanical and electrical considerations, such as cable routing, are finding their way into the early phases of system design. With tools such as Cabling Station from Mentor Graphics, designers can store cable-specific attributes such as wire gauge and color, while automating the generation of diagrams, bill of materials and parts lists. Design rule checks and calculations of wire lengths and diameters can be performed.

as tools to help them design their product."

Framework technology promises to solve some of these data management tasks in the future, but right now they seem to be limited to managing the suite of point tools that individual teams use to complete their portion of the design.

"We use Powerframe, our framework product, to manage the tools of design teams," says Fumio Taku, electronic design strategy manager for the engineering systems group at Digital Equipment Corp (Maynard, MA). "Then we use a worldwide network to communicate with all the various design groups that are working on a project. At each stage—from product definition, to introduction, to market retrieval—the groups communicate on a daily basis to review the progress of each phase."

In addition to tool management in select design groups, framework technology is being used today for data and process management. The possibility of one large framework to manage the entire design program is currently unfeasible. "The ideal situation would have a unified, giant database that could share all data about the product," says Martha Stone-Martin, frameworks marketing manager for DEC's systems group. "The problem is that it would be so unwieldy that it couldn't offer the performance that designers need. Sifting through gigabytes of data would take so long that no one would use it. Besides, it's simply not necessary that all design data is accessible by every member of the entire design team."

Looking beyond tools

The problem of design capture, then, is not only one of knowing as much as possible as early as possible, but efficiently managing the information flow between the various groups along a product's development path. Tools that give valuable thermal predictions are useless if the information doesn't get to the necessary teams. "We're being told by large systems houses and IC companies that the design teams are getting larger, the amount of data is enormous, and time-to-market deadlines are getting more critical," says Tony Zingale, vice-president of corporate marketing at Cadence. "They want to know how we're going to help them—not only as a tool vendor, but as a framework vendor—manage that huge design task. It's an intimidating problem, but it can't be avoided any longer."

As more and more disciplines, such as design, test and manufacturing get included early in design capture, there will be fewer unpleasant surprises when it's time to put the pieces together in the finished product. For this to happen, there has to be a commitment to inter-departmental communication as well as employing the latest design tools. "A lot of steps in the design process are still manual," says Benjamin Sprachman, director of system development at Prime Computer Systems (Natick, MA). "But I can see it changing. It's essential to keep control of the design, and that means an evolutionary change of the design process itself. We can't keep spending all this time carrying notebooks around to make sure that everyone is on the same track."

This restructuring of design teams, as well as the closer cooperation between EDA vendors and systems manufacturers, will undoubtedly result in higher-quality products with shorter time-to-market cycles. "The whole thought process about systems design hasn't changed yet," explains Cadence's Katz, "but designers are changing how they think of purchasing EDA. In the past, people bought a workstation or tool based solely on its power to perform a task. Today, people purchase an EDA suite of tools to get a design done. It's not unlike purchasing a phone system. Businesses don't just buy a phone to put on a desk—they buy a system, with features tailored to help their organization communicate. Likewise, an EDA system has to be a combination of communication and design tools to ensure that the finished product fulfills the promises of the initial design." ■

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

Cadence Design Systems (408) 943-1234	**Circle 208**
Compix (503) 639-8496	**Circle 209**
Digital Equipment Corp (508) 493-5111	**Circle 210**
Hewlett-Packard (800) 752-0900	**Circle 211**
Intergraph (205) 730-7504	**Circle 212**
Mentor Graphics (503) 626-7000	**Circle 213**
Prime Computers (508) 620-2800	**Circle 214**
Teradyne (617) 422-3309	**Circle 215**
Valid Logic Systems (408) 432-9400	**Circle 216**

100 WORLD-CLASS COMPANIES HAVE ALREADY TAKEN A SHORTCUT THROUGH THE X.25 MAZE.

TITN X.25 SOURCE CODE GETS YOUR SYSTEM TO MARKET FASTER.

You're building a product that connects to X.25 networks. Should you go it alone through the development maze? Or bypass it completely with proven, guaranteed source code that has all these features—licensed from TITN.

Fast. With TITN source code, you're buying time. Our streamlined software package makes porting to your platform easy. And our handy reference implementations and extensive documentation further speed development. Plus we've already been certified on 18 PSDN networks in 14 countries.

Field-proven. Our fully debugged source code is a known quantity. We've been licensing it since 1982, to over 100 of the world's leading computer companies.

Guaranteed. Only TITN guarantees—in writing—that your product will pass worldwide X.25 certification tests. And our family of complementary OSI software conforms to emerging European and U.S. GOSIP standards, to give you a competitive edge.

Accelerate product development. Contact TITN today and shoot through the maze.

TITN

U. S. HEADQUARTERS
7011 Koll Center Pkwy., Suite 200
Pleasanton,CA 94566-3101
Tel (415) 484-5674
Fax (415) 484-4078

EASTERN SALES OFFICE
1300 North 17th Street, Suite 1530
Arlington, VA 22209
Tel (703) 528-2662
Fax (703) 528-0388

ATA/EUROPE
1, rue Galvani
B.P. 110
91301 Massy Cedex, France
Tel 33 69 81 11 00
Fax 33 69 20 15 04

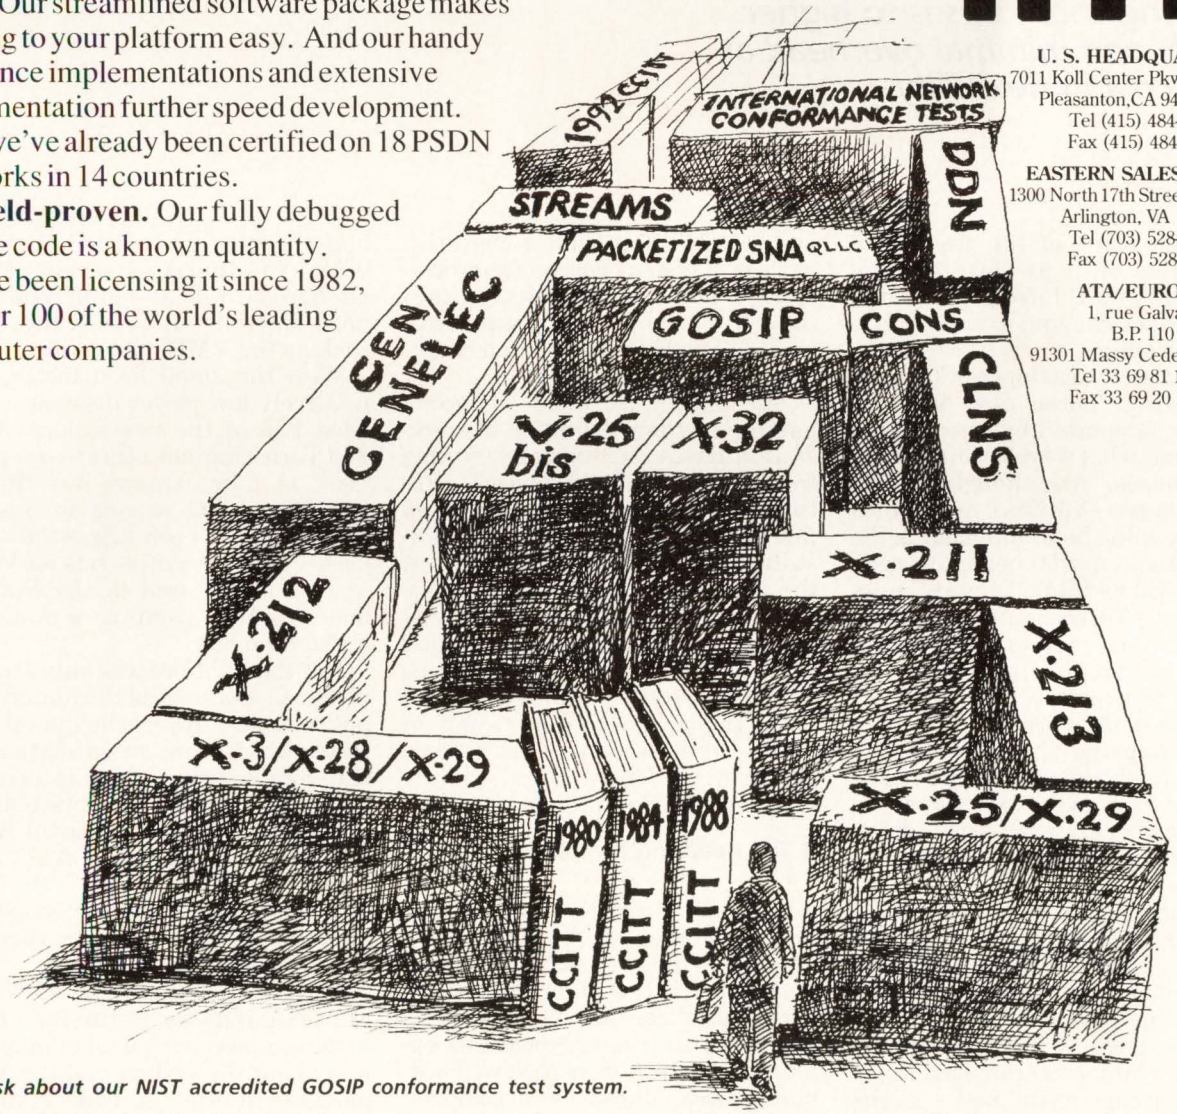

FLASH! Ask about our NIST accredited GOSIP conformance test system.

TWICE ® ***SOFTWARE MODULES*** /CMISE with ROSE / VTP / X.400 with Message Store and User Agent / FTAM/Application Services for FTAM, ACSE and Presentation protocols / OSI Transport and Session layers / CLNP / X.25/X.29 Network, Data Link and Physical layers / X.25 Options, such as X.3/X.28 and QLLC / UNIX Utilities and Streams interface / FAXWARE / Support for GOSIP, COS, NIST, MAP/TOP, CEPT AND CEN/CENELEC profiles

CIRCLE NO. 45

Will performance win over sophistication in workstation buses?

Warren Andrews
Senior Editor

A new class of standard bus has emerged to support the I/O requirements of compact desktop workstations. The continued growth of processor performance is driving these buses to higher speeds and minimal overhead at some cost in flexibility.

Though some of the traditional open buses such as VMEbus and Multibus I and II have been used in workstations ranging from Sun 3s to Silicon Graphics, a new breed of low-cost, compact, desktop workstation is emerging. These new workstations are supported by new bus architectures, what we refer to as "workstation buses." And though the traditional buses—VMEbus and Multibus—shouldn't be counted out in the workstation environment, the newcomers—Sun's SBus, Digital Equipment Corp's Turbochannel and IBM's Micro Channel—are taking a strong position with a growing following.

While the term workstation bus continues to defy precise definition, buses occupying the category are thought to share a few characteristics. But not everyone agrees on what these are. Paul Borrill, director of hardware engineering at Sun Microsystems (Mountain View, CA), says workstation buses are relatively small in form factor, so they can squeeze into desktop configurations; relaxed in power dissipation, minimizing power-supply requirements; extremely high in performance (50 Mbytes/s transfer rate) to take advantage of the power of the latest RISC and CISC processors; and scalable to provide a migration path to higher-performance processors. In addition, Borrill says, they have to be totally open standards, available to both system and board-level manufacturers.

Still others are more all-encompassing in their definition of workstation buses, including everything from personal computer buses to Futurebus+. And perhaps the bottom line defines any bus used in a workstation as a workstation bus. But the changing face of workstations, higher levels of semiconductor integration, increased performance, and the more-demanding applications expected of today's workstations are mandating a new generation of buses—and associated performance parameters. These agree closely with Borrill's definition.

The new and the old

The new breed of workstation buses—SBus, Turbochannel and Micro Channel—differ from the old variety—VMEbus and Multibus—in at least a few significant ways. First, all three of the new buses operate with a motherboard as opposed to being on a passive backplane. This allows for a compact, desktop form factor, and lets all the housekeeping as well as some universal I/O functions reside on the motherboard. It also eliminates the need for designation of a host slot, such as the VMEbus "slot 0."

With the small form factor, and relatively low power dissipation, at least two of the newcomers, SBus and Turbochannel, start to resemble more of a mezzanine bus than a full-fledged I/O expansion bus. In fact, SBus has been suggested—and used—as a mezzanine bus for VME, and even proposed to the Futurebus+ working group as a standard mezzanine bus.

But these differences only scratch the surface of some of the underlying philosophical approaches to bus architecture. "There are gradations in the complexity of bus interface," says Ray Alderman, technical director of the VFEA International Trade Association (Scottsdale, AZ). "One way to view it is in terms of the level of 'formality' of the bus interface. The more formal, the more complex and expensive the interface.

"Bus standards such as VME, Multibus and even Futurebus+ have to know and convey a lot of information about the system and the other modules in order to transfer information," Alderman continues. "As

such, these can be considered the formal buses, requiring such features as arbitration signals, complex interrupt schemes and even cache coherency."

Alderman considers these buses the "tuxedos" of the industry. "They're burdened with a lot of protocol, handshaking and rules of behavior," he says.

Less encumbered buses such as the new breed of workstation buses might well be considered the "sport jacket and slacks" set. "These buses are essentially extension buses rather than expansion buses," Alderman says. Designed primarily to handle very high speed I/O with as little extra baggage in terms of control signals and interface controllers as possible, the new generation workstation buses—if indeed they are buses—are basically easy and quick to design, yet provide real I/O horsepower. Workstation buses are still restricted by some level of formality and handshaking, but the rules are far less stringent than full system buses.

"The 'T-shirt and jeans' set of the bus business is the dedicated mezzanine bus such as iSBX for Multibus, or one of the many daughterboard approaches available for

VME," Alderman continues. These buses, he says, are even more elemental in their level of interface complexity than the workstation buses. In some cases, these buses represent little more than extensions of the processor's address and data lines. And, he says, these buses are almost bereft of protocol and call for only the most casual of greetings.

Key attributes

And while Alderman's definition is satisfactory for establishing some general levels of hierarchy, it only touches the surface of some critical elements of bus architecture. To look more closely, it's necessary to define the key attributes of a bus architecture and see how they apply to each bus. The relevant key attributes in workstation buses—or any bus architecture for that matter—include both technical and market considerations.

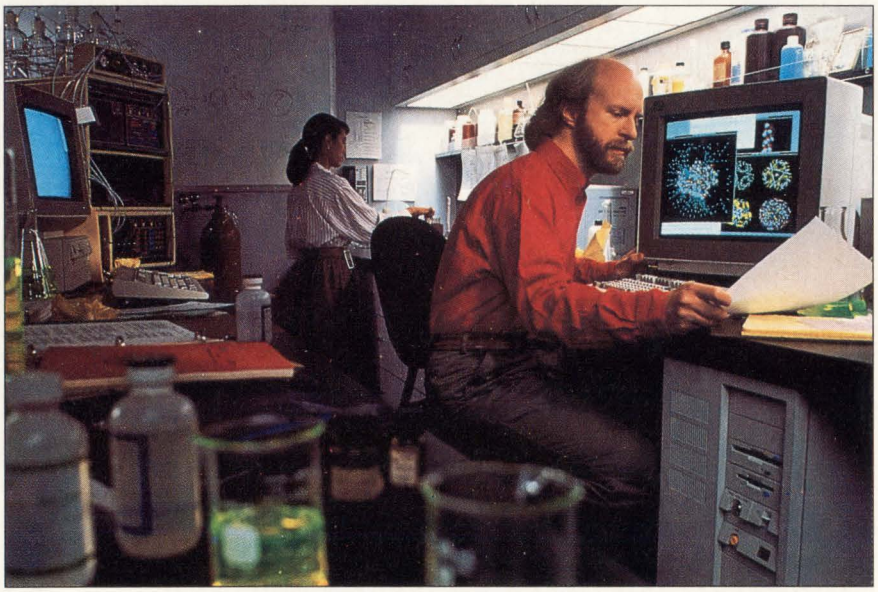

Compact desktop workstations such as IBM's RISC System 6000 (shown here with a server) are challenging systems designers by demanding high system performance without offering room for traditional bus configurations. A new class of systems bus is emerging to serve the special needs of workstations.

On the technical end, performance or bus transfer capability usually leads the list, followed by things like interface complexity, signal-pin and connector requirements, board size, power dissipation, and driver requirements.

On the business or market end are some of the same concerns, only with a different focus. It can be, for example, as important to know the cost as it is to know the speed and drive requirements for interface buffer chips. Similar concerns for support include the availability and cost of interface and controller (if required) ASICs, board size, layer count and complexity, corporate sponsor or industry organization support, and, of course, the all-important migration path.

In addition, business attributes include surface-mount requirements, process technology of onboard circuitry (with relation to available power), critical mass of vendors supplying requisite functions, and overall complexity (measured in design time) to hit critical market-entry windows. And, in addition to the migration path forward, systems must be able to provide backward capability with previous generations of equipment. In some cases, this requires software capability; in others, a complete bridge to older hardware.

Obviously with such a broad definition of attributes, no single bus will be ideal for every application. But it's likely that a quick review of some of the existing workstation bus characteristics will provide enough information to quickly narrow down the choices.

Analyzing by size

In looking at the three major workstation buses, it's difficult to determine which to start with. Chronologically, Micro Channel is the oldest, having its start in IBM's PS/2 personal computers. SBus, however, predates the present workstation incarnation of Micro Channel. To rank them on performance would be equally difficult, since all claim great performance specifications that either haven't yet been realized, or are "implementation dependent." Physical size, therefore, appears to be one of the more straightforward ways of cataloging the buses.

SBus is the smallest of the three with a card size, including connector area, of 3.3×5.76 in., resulting in a 19-in.^{2} area. Sun makes no bones about the bus being designed to operate with high-level integration components in CMOS to minimize size and power dissipation. Each slot in the SBus specification allows for a maximum current of 2 A at 5 V. In addition, the bus provides for 12 V at 30 mA if higher voltages are required for analog I/O functions or to accommodate communications protocols that require a higher voltage swing. This gives SBus a total power dissipation of 12.6 W.

SBus is a synchronous bus, regulated by the master bus clock at either 16.67 or 25 MHz, with all operations synchronous except for interrupts that are paced by the SBus controller. The bus specification allows for both master and slave adapter cards, capable of accepting up to eight masters. As initially released, the bus features a 32-bit data path with a 32-bit virtual address for masters and a 28-bit physical address for slaves.

With the introduction of the B.0 version of the SBus specification in January, the data path will be expanded to include full 64-bit transfers. To achieve the expanded bandwidth, the existing data pins will be multiplexed. The resultant 64-bit version will be fully upward and downward compatible with existing A.2 32-bit versions.

Using the master and slave configurations, SBus-based systems are capable of I/O communications to and from the host processor, direct virtual memory access (DVMA) to and from the host memory, and direct communication from one SBus master to another SBus master or slave. Translation of the virtual address to a physical address takes place on the SBus controller in the case of DVMA masters, and within a CPU in the case of SBus masters including their own CPU with a private means of performing virtual address translation. Typical SBus systems have only a single CPU master.

Turbochannel targets graphics, networking and large database storage

The Turbochannel I/O interconnect design originated from the same design philosophy as the RISC computer architectures. In the Turbochannel interconnect, as in RISC, the architecture is designed by analyzing where the problems are and solving only those problems that improve performance. The simplicity of the resulting design lets designers fix the problems both quickly and efficiently. Both architectures simplify designing, compared with other bus interfaces or computer architectures. Simplicity means very few signals. The whole design is optimized for input/output with all control signals being point-to-point.

The Turbochannel interconnect was designed to solve the I/O problems of graphics, imaging and animation, high-speed networking, high-speed data collection, and large database storage. Those applications benefit from a fast, low-latency I/O interconnect. In the special case of graphical computer design and usage, providing high bandwidth to memory for the processes to operate at fast speeds is the primary way to enhance performance. With a radial design architecture, the bus is not shared, so there's no loss in bandwidth to memory. Turbochannel brought the radial design to the desktop and enabled very fast graphics, animation and high-end imaging to operate on the desktop, rather than deskside, computer. The theoretical peak bandwidth of Turbochannel is 100 Mbytes/s. The achieved peak today is 78.1 Mbytes/s.

Simple is faster

It's easier to implement and debug an option because of the minimalist architecture of the Turbochannel design. With the market demanding more complicated options in the 1990s, a designer needs to spend the most time designing the option and the least time learning and designing the bus interface. That goal is supported by the simple design of Turbochannel. In other words, the less complicated the bus design, the more complicated the option can be and still respond to market demands rather than be limited by bus technology. The general simplicity of the Turbochannel design enables the use of more aggressive semiconductor process technologies to be used for the manufacturing of an option.

To solve high-throughput I/O interconnect problems, option vendors have available a range of implementation—from low-cost to high-performance. If an option implements DMA transfers, for example, a designer could make either the high-performance choice and use large data buffers for large DMA bursts, or choose less costly buffers for small DMA bursts. The choice is in the option designer's hands.

The Turbochannel architecture connects one system to perhaps several options in an asymmetric, radial design. This design trades off direct option-to-option data flow for the many benefits of simplicity and speed. The processor and memory are conceptually together. The Turbochannel option is at one end, and the system, defined as the processor and memory, is at the other. This configuration is less complicated than in a traditional bus structure and is more efficient.

The Turbochannel specifications are open to the industry through Digital Equipment Corp's TRI/ADD program.

Scalable design

The Turbochannel specification is flexible enough to take advantage of technological breakthroughs in workstations of the future, especially in the area of memory design. Quite a bit more board space is available than other comparable systems provide. Minimal space is taken up by the Turbochannel interface. If necessary, components can also be placed on both sides of the board. Turbochannel is a synchronous design that provides 32-bit wide data paths and 34-bits of addressing. Only 44 signal lines, together with 47 power and ground lines, result in high signal integrity. Each slot provides 26 W of power. Airflow is good across each option.

Turbochannel supports two types of transactions: programmed I/O—the processor writes to or reads from the I/O controller option; and DMA—the I/O controller option reads from or writes to memory.

Options that meet the specifications can function in a variety of system implementations, from low to high performance. The variables in system implementation include single-to-multiple I/O controller options. The address and data path can be bused or passed through a crossbar switch. For common noninterleaved memory, the clock rate can be 12.5 MHz. For interleaved memory, the clock rate can go up to 25 MHz. Memory can be single- or multiported.

Industry implementations

Lowest-cost system implementation might embed the Turbochannel interconnect within the system, and operate at 12.5 MHz, allowing the use of a single bank of page-mode DRAM to meet the timing requirements. Mid-level implementations might use multiplexed address/data paths between the I/O controllers and memory.

The higher-performance, and as of yet only theoretical, implementation would operate separate Turbochannels for each I/O controller, effectively granting each I/O controller the full bandwidth. The aggregate peak I/O bandwidth per second then would be 100 Mbytes times the number of I/O controllers. Such an implementation, along with Turbochannel's ability to support enormous physical address space, would let Turbochannel options developed today take advantage of tomorrow's technological advances.

The Turbochannel specifications are open to the industry through Digital Equipment Corp's TRI/ADD program. Use of the Turbochannel technology is governed by a "Technology Transfer Agreement" printed near the front of DEC's Turbochannel documentation. There's *no* fee to use the technology, and DEC encourages the use of Turbochannel interconnect by both system and option vendors.

Brenda H. Christensen, manager of TRI/ADD program, Digital Equipment Corp

New SBus specification helps developers prepare for the future

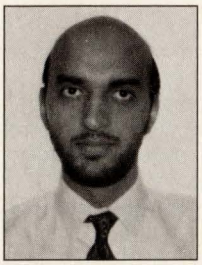

SBus, a high-performance, open bus architecture designed by Sun Microsystems, is emerging as a standard I/O interconnect for RISC desktop and server systems, optimized for the technologies of the 1990s. As desktop computers keep moving toward higher performance, SBus will be able to keep up with their I/O needs. Existing SBus solutions will become even less expensive, and bandwidth-intensive solutions will start appearing in greater numbers. Furthermore, Sun is committed to SBus as the I/O expansion interconnect for all current and future systems.

The SBus specification is available—free of charge—to any vendor who wants to design systems or add-in boards that use this low-cost, high-speed interconnect. Sun recently released a new version of this specification that will help system and add-in card developers dramatically improve the performance of SBus and SBus-based systems. A key feature of the new specification is the provision for a 64-bit transfer protocol.

Specification B.0 now lets 64 bits of data be transferred with each clock cycle, giving SBus a peak bandwidth of 160 to 200 Mbytes/s. In most systems, I/O bandwidth is the data bottleneck. Current 40-MHz CPUs, for example, can process data much faster than most I/O devices can supply it. The new specification makes SBus the only desktop expansion bus that can keep up with such fast-performing systems. It also eliminates the need for buffers in many applications. Low-cost SBus systems can now be designed for data-intensive applications such as graphics, digital signal processing, multimedia and high-bandwidth networking.

Moving to 64 bits

When platforms that implement the 64-bit architecture appear, SBus will be ready for them. The specification shows SBus systems designers how to create the new, wider data path without any increase in wires or transfer time. Data is loaded on the 32 existing data lines, 28 address lines and four control lines after the address has been latched.

In addition, the new specification is backward-compatible. Existing 32-bit SBus cards will continue to work in existing 32-bit systems as well as in future systems with 64-bit slaves; 64-bit masters will be able to communicate with 32-bit SBus slave devices.

Specification B.0 increases the maximum length for burst transfers from 64 to 128 bytes without any increase in latency. The specification also provides guidelines for building bridges from SBus to other buses, letting users with investments in older technologies migrate to SBus for new high-performance applications.

The specification discusses multiple SBus implementations on a single host system, so that several SBus devices can be accessed in parallel, resulting in an aggregate system I/O bandwidth higher than that of a single SBus. Such an implementation would be suitable for servers that must be able to handle the demands of several client workstations. The specification also covers atomic transactions and bus sizing in more detail than the previous version.

A major step toward making SBus cards off-the-shelf products is the inclusion in the B.0 specification of the SBus Profiles, which define the minimum criteria for certain functions—addressing, burst size, data parity, error handling, and time-outs—that every SBus expansion card must have to be considered a standard implementation of the SBus specification. SBus cards that adhere to the Profile specifications are completely compatible with any other platform or SBus card that meets the Profile specifications. (SBus developers can use Open Boot firmware to design systems or cards that exceed the minimum criteria spelled out in the Profiles.)

Support for SBus

Several chip vendors are working to provide silicon for the new specification. LSI Logic currently offers two SBus DMA interface chips that support the current A.2 specification; having such ASICs makes it much easier for developers to implement their solutions on SBus.

As of November 1990, approximately 63 vendors were offering more than 135 SBus products, including Ethernet cards, serial port extenders, and facsimile and modem boards. In keeping with Sun's strategy to make the SBus architecture and specification completely open, the SBus Developer's Group created a public SBus specification committee in September 1990 that lets third-party developers participate in developing the SBus specification. This committee provided substantial input into the B.0 specification.

Standard criteria such as the Profiles specifications will make it easy for developers to implement SBus cards in the high-performance systems of the future.

Kuljeet Kalkat, BSEE, SBus product manager, Sun Microsystems

Turbochannel, DEC's offering to the workstation environment, is probably the most recent entry into the workstation bus market and, perhaps, most closely resembles SBus. There are, however, some significant differences. Like SBus, Turbochannel is a fully synchronous 32-bit bus capable of high-speed I/O and DMA transactions. It's designed to be operated at synchronous speeds of 12.5 and 25 MHz. But, "we like to think of Turbochannel as an I/O channel rather than a traditional bus," says Brenda H. Christensen, manager of the program for third-party hardware development, TRI/ADD, at DEC (Palo Alto, CA).

"Traditional bus designs leave you dependent on the technology at the time the bus was designed. We tried to design a high-speed interconnect that would solve some I/O problems on the desktop and be radial in architecture and capable of taking advantage of new technologies as they emerge." Turbochannel, continues Christensen, doesn't solve all the problems of the standard bus world. "That's why we'll continue to have SCSI on the lower performance end and VMEbus adapters to handle higher-end problems," she says.

"And, bridges to other buses, including Futurebus+, which DEC is solidly behind, are already under consideration. We wanted to solve the problems of high-speed networking, large data storage and high-performance graphics on the desktop without having to go to larger systems. That's what Turbochannel does."

Heftier bus

In putting together the Turbochannel specification, DEC had perhaps some benefit of hindsight—Sun's SBus had already been on the market for almost a year. DEC elected a slightly larger form factor than the SBus card—DEC's measures 4.6×5.675 in., providing a 26-in.^{2} primary board area. "Despite the growing trend toward higher-level integration ASICs," says Christensen, "the greater board area leaves room for high-performance devices in addition to extra memory required by such applications as high-performance graphics subsystems."

In addition, Turbochannel allows a hefty 4.0 A of 5 V power for each slot and 12 V at 500 mA—in total, about 26 W—more than twice the available power per SBus slot. To handle the additional power dissipation, DEC also increased the airflow in its specification, calling for 150 linear ft/min—almost four times that of SBus.

The availability of the additional space and power, says Tom Furlong, DEC's RISC workstations manager, allows the development of modules with things such as high-speed static RAM that might otherwise be restricted on buses with less power and cooling. It also allows for the use of specialized high-speed processors such as graphics processors and advanced digital signal processing devices.

"And where even greater power is needed," says Christensen, "more than a single Turbochannel slot can be used. For example, using two slots in a double-wide configuration, a total of 8 A is available at 5 V and 1 A at +12 V. Similarly, triple-wide configurations allow for even greater power dissipation."

Yet despite the increased power available for Turbochannel options, the bus, like SBus, is intended primarily as a CMOS bus. The company suggests that a standard 1.5-µm CMOS is sufficient for on-board ASICs, and that relatively inexpensive 7.5-ns PALs are sufficient for interface and controller logic. And, though the bus drivers can be CMOS, the company recommends that the T versions with TTL threshold levels be used, rather than the rail-to-rail CMOS levels.

More of the same, but different

IBM's Micro Channel Architecture (MCA) shares many elements with the other approaches but also has some pointed differences. Similarities are in the relatively small form-factor and low-power CMOS implementations, and that they're basically motherboard bus architectures. But that's about where the likenesses end.

MCA is essentially a systems bus, rather than a simple extension bus. It transfers via asynchronous protocols, and it's the only one of the three "workstation buses" to exploit a full dual-bus architecture. "In developing the RISC System/6000 Power-station family," says Phil Hester, director of the Advanced Workstation Division, Engineering Center, at IBM (Austin, TX), "we wanted to focus the bulk of our creative energy in the most critical piece of the overall system—and that turned out not to be the I/O bus."

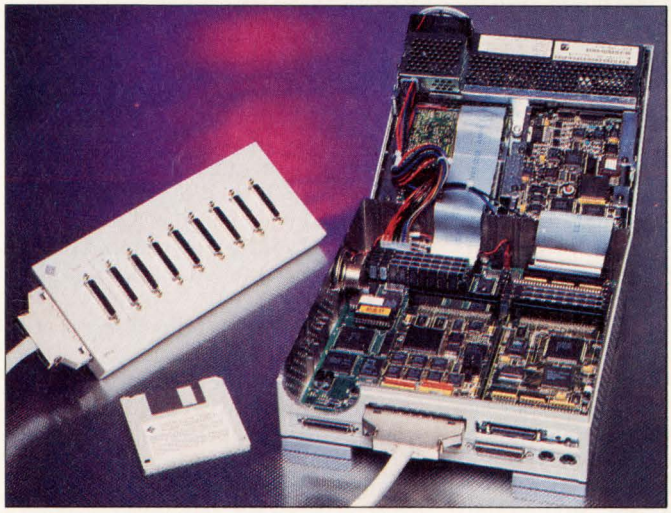

The compact configuration of desktop workstations puts option-board real estate at a premium. This SCSI adapter card (at the lower right) for Sun's SBus shares tight quarters with a serial-parallel controller card.

Hester says that in developing the workstation, IBM wanted to create an exceptional processing engine that would easily scale up with the technology available. The current generation, he continues, using the company's super-scaler CPU, boasts a Mips rating of 56, a floating-point performance of 23 Mflops and a 54.3 rating on the Specmark benchmark—one of the best performances turned out by any of the workstations in its category.

"But," adds Hester, "we see the technology moving forward at an accelerating rate. With forecasts that can be made even now, we envision CPU performance quadrupling within the next few years.

"In putting together the RISC System/6000, and examining the overall architecture, and looking at I/O bus alternatives," recalls Hester, "we were pretty sure of two things: we didn't want the I/O bus to get in the way of CPU performance either now or down the road; and, we didn't want to have to reinvent the wheel (another bus architecture) if we didn't have to." Hester describes looking at alternatives such as VMEbus, VME64, MCA and the other commercial alternatives.

MCA, he says, offers the clearest migration path of any of the existing approaches examined. (What he didn't say, but should be understood, is that MCA is basically an IBM proprietary bus and doesn't carry all the political baggage of an IEEE-endorsed architecture. This means the company can dictate changes or modifications of the bus without a costly and time-consuming round of approvals.)

Performance rules—or does it?

One of the keys to IBM's workstation strategy, Hester says, is to ensure that the system can be upgraded to significantly greater performance levels without forcing a major change—or expensive technological approaches—in developing new I/O modules. A second consideration, he says, is that the bus can be scaled to greater performance levels without major architectural revisions.

Interestingly, the same (or similar) claims are made by proponents

VITA's Ray Alderman was dubbed the unofficial "high priest" of Futurebus by the trade association. He groups buses by their complexity. Large buses, such as Futurebus and VME, have formal interfaces. "They're the 'tuxedos' of the field. Workstation buses are more casual; they're the 'sport jacket and slacks' of today's buses."

of both SBus and Turbochannel. According to Hester, both SBus and Turbochannel are involved directly, and synchronously tied to, the main processor bus. "This is acceptable as long as the processor technology doesn't outstrip the capabilities of standard bus driver and interface components," he says.

The result when the processor capability is significantly increased, Hester maintains, is that the I/O buses will either drag down the performance of the system, or they will have to be updated with relatively expensive—and perhaps scarce—components such as BiCMOS or perhaps even gallium-arsenide. Hester believes that processor technology will reach that critical level in only a couple of years.

IBM has already announced its migration path for MCA—one which will take it from the present 20 Mbytes/s to 160 Mbytes/s. Hester says the jump will take place in four stages, nominally labeled A, B, C, and D. The first jump from A to B, already implemented in the RISC System/6000 family, involves what the company calls its data-streaming mode. In essence, this high-speed burst-mode transaction lets a single address set up a sequential burst of data instead of a new address being sent for each data word. This effectively brings the data rate from 20 to 40 Mbytes/s.

64-bit question

The next jump moves the bus essentially into the 64-bit domain by multiplexing data over the address buses while in the data-streaming mode. Much like the approach taken in VME64, MCA's C version takes advantage of the idle address lines during data-streaming transactions and transmits data simultaneously with the information traveling across the data lines. This boosts the effective maximum data rate to 80 Mbytes/s. The final stage in the migration will double the maximum bus transfer clock speed, again doubling the theoretical transfer rate to 160 Mbytes/s.

In multiplexing the address and data lines, first VME and then Micro Channel doubled the data transfer rate by joining the 64-bit revolution (evolution). Sun Microsystems is next to make the jump, announcing this month its B version of the bus (see "New SBus specification helps designers prepare for the future," p 82). Sun also added some other enhancements to improve the flexibility and performance of the bus. According to Chris Wellens, manager of SBus technical support, the 64-bit extension is fully upward and downward compatible with existing hosts and modules.

Being able to handle 64-bit transfers in what is essentially a 32-bit system would, on the surface, provide little real advantage, except in some rather specific applications. What's largely suspected—and it's hardly secret—is a rash of 64-bit processors to hit the street.

In Sun's case, look for a 64-bit version of Sparc—not the Fujitsu H version, but a totally new version to come from a totally different vendor under contract with Sun. Rumors have it that the device, and probably a new "screaming" Sparcstation, can be expected later in the year. On IBM's side, there's little speculation that its proprietary RISC architecture won't jump to 64 bits in the immediate future.

Turbochannel doesn't have any painless 64-bit upgrade short of simply using a double-wide card to accommodate the 64 bits. But DEC doesn't seem concerned. "Right now we have proven performance near the 100 Mbytes/s rate," says Christensen, "which is far faster than anything Sun's SBus can do, even if it doubles its rate."

In its decision to multiplex the data and address lines, DEC vied for simplicity of the architecture over the possibility of expanding the bus to 64 bits. "Within the architecture of Turbochannel," says Furlong, "we have the capability of increasing bus transfer rate in the next generation from its current 100 Mbytes/s to as much as 400 or even 800 Mbytes/s, if needed."

Architectural vs. achieved

But there's always some debate over what the theoretical (architectural) transfer rate is versus what's actually achievable. And it's also sometimes interesting who put the numbers together, and how. Just as a

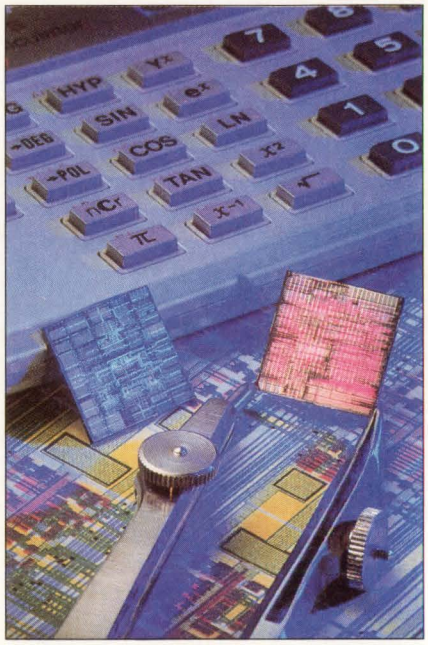

The super-scalar CPU behind the IBM RISC 6000 workstations currently leads the workstation pack with a 56-Mips, 23-MFlops rating. Many experts predict that the next wave of high-performance workstations will be based on full 64-bit processors and buses. Introductions of these systems are expected this year from more than one manufacturer.

Mips or an MFlops benchmark may not provide an accurate indication of a system's performance in any given application—even though these factors (instructions per second and floating-point performance) are fundamental to performance.

Similarly, looking at straightforward bus transfer rates is not always a key to what a bus can do. And, there's always the phrase "implementation dependent" mentioned earlier. This phrase seems to

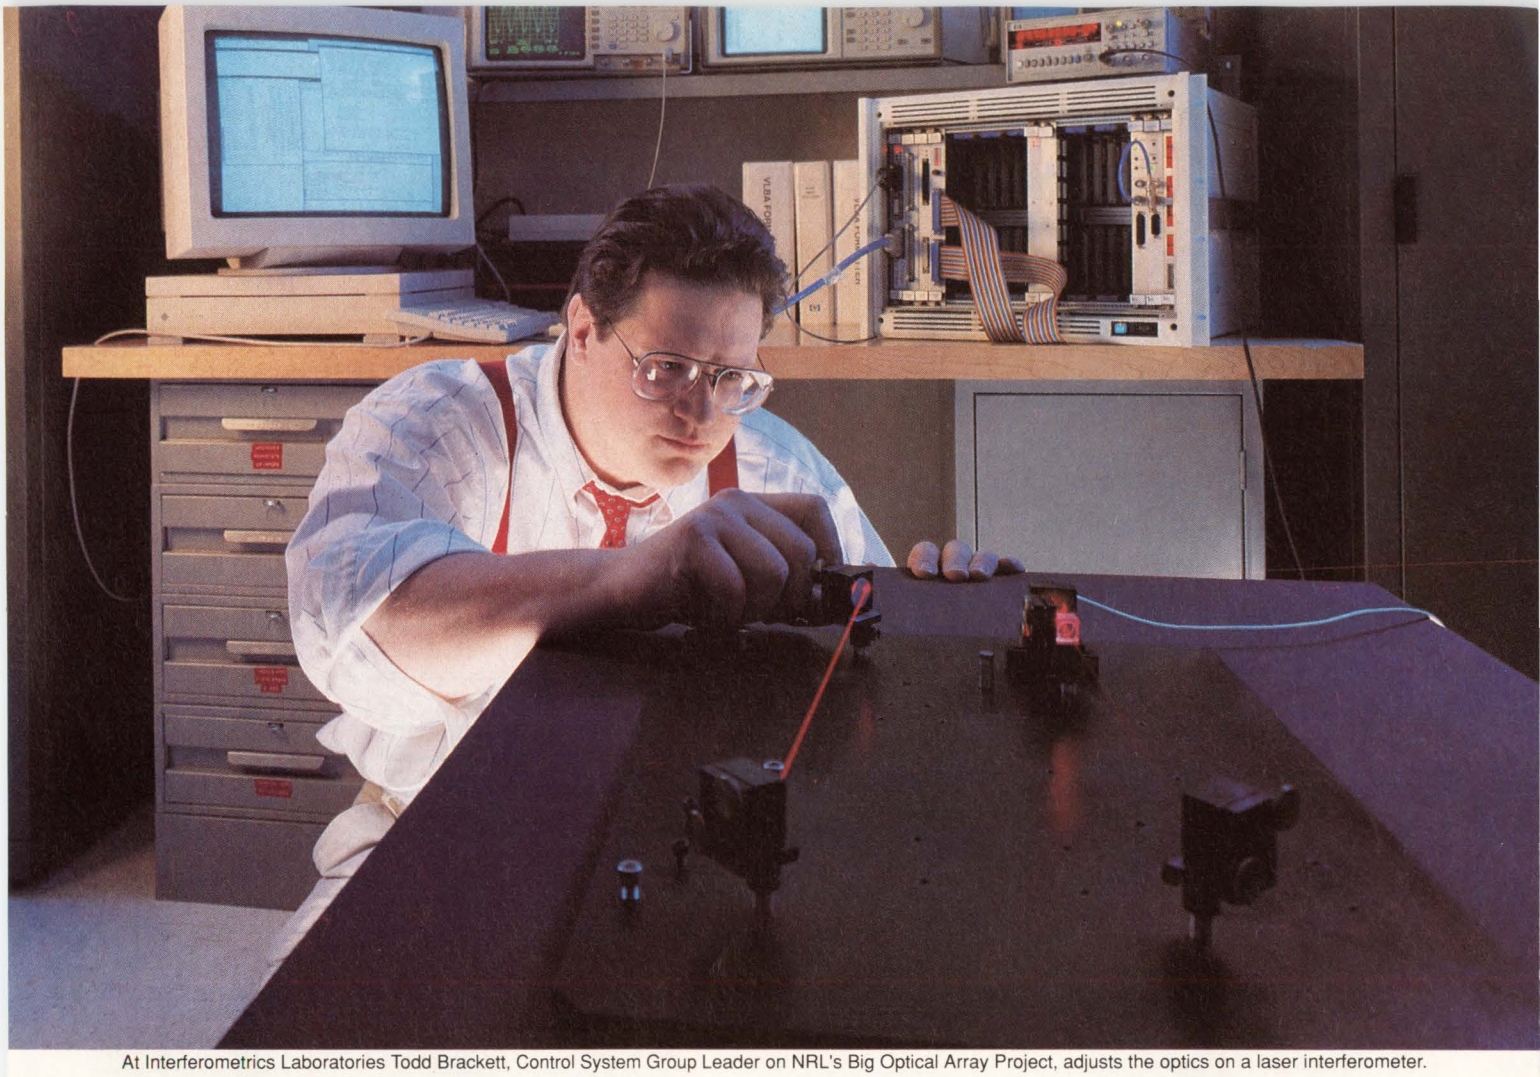

At Interferometrics Laboratories Todd Brackett, Control System Group Leader on NRL's Big Optical Array Project, adjusts the optics on a laser interferometer.

Opening a New Eye on the Cosmos

The Naval Research Laboratory needed a better way of seeing the Universe. They came to Interferometrics for help.

When Interferometrics needed computers to control NRL's revolutionary Imaging Optical Interferometer they came to Themis Computer.

The NRL Big Optical Array imaging interferometer requires a perfectly synchronized dance of mirrors to eliminate the delay and distortion of starlight passing through turbulence cells in the atmosphere. These mirrors require an intricate network of motors, sensors and actuators, linked through Ethernet, working together, to counter the earth's motion and distorting effects of the atmosphere. Themis Computer's broad, fully-compatible family of VMEbus CPU, I/O and Communication boards provides real-time control for this stellar project.

The result will be a new generation of optical telescope capable of imaging star systems to an angular resolution 1,000 times greater than conventional ground-based optical telescopes, and 200 times greater than the Hubble Space Telescope. Significantly, this will be accomplished at a cost of only $10 million, 1/200th of the Hubble's cost.

This remarkable project will result in more accurate mapping and measurement of the stars and improve our understanding of phenomena such as black holes, binary stars, and quasars.

We at Themis are proud of our involvement in this endeavor.

Whether your application is on the factory floor or out of this world, Themis has the hardware and software to put it all together.

Call (415) 734-0870, fax: (415) 734-0873, or in Europe, 33.1.69.86.15.25, fax: 33.1.64.46.45.50.

At Themis, we're making the open systems promise a reality.

THEMIS COMPUTER

Americas: 6681 Owens Drive, Pleasanton, CA 94588
Europe: 29, Av. de la Baltique, 91953 Les Ulis Cedex-France

CIRCLE NO. 46

Getting I/O buses to keep up with high-speed processors

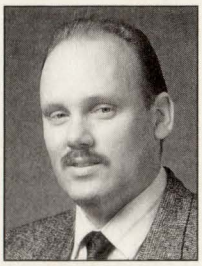

Today, very high speed CISC microprocessors, such as the 68040, can move data at almost 70 Mbytes/s. But this high-data bandwidth can be reduced tremendously by an inefficient I/O bus design, making a carefully engineered I/O bus interface an absolute must for high-performance systems.

The main goal of an I/O bus interface is to transfer data as quickly as possible between peripherals and the main processor's memory, while at the same time reducing the processor's participation in the bus transfer process. This lets the processor concentrate on other important high-level tasks without wasting bandwidth moving data over a relatively slow I/O bus.

There are four techniques for enhancing the performance of the I/O bus so that it can more effectively be utilized with very high speed processors. These techniques—byte packing, DMA, localized intelligence, and *optimized* block transfer—can all be incorporated at the same time to maximize performance.

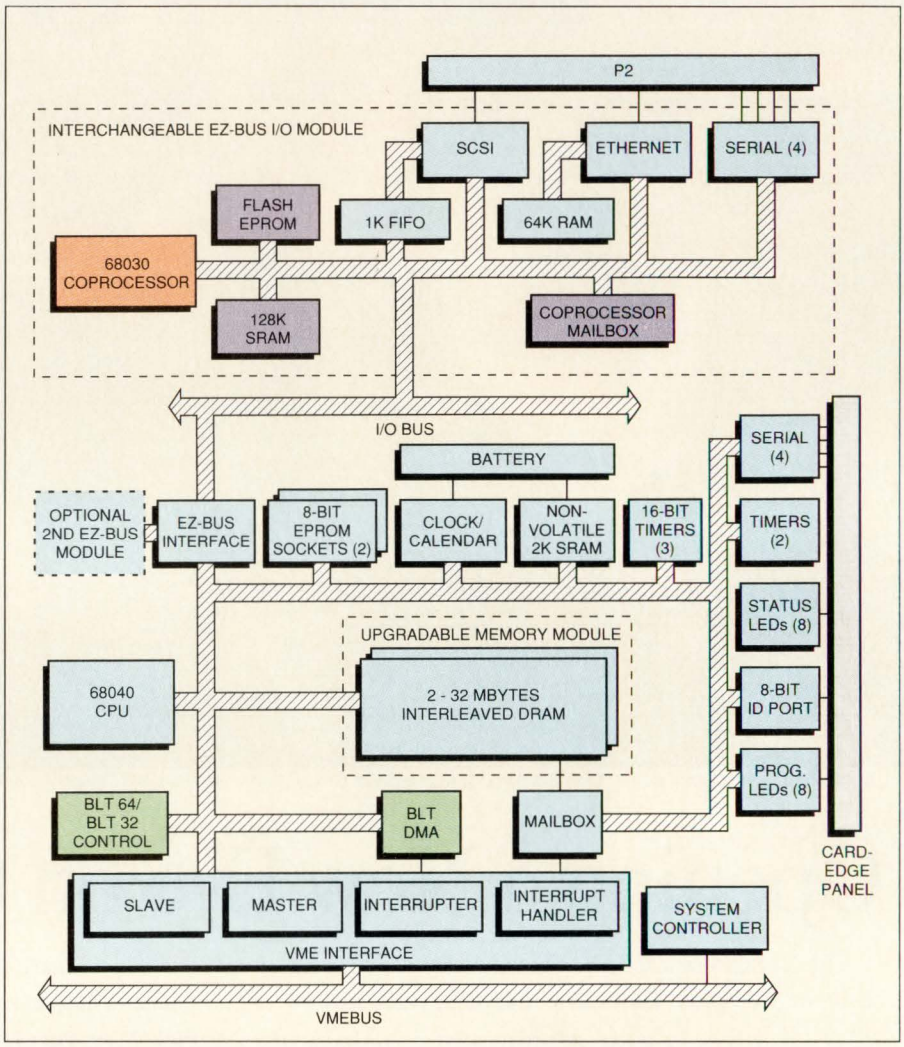

Reducing data transfers

Byte packing, a concept usually incorporated in all but the lowest performance systems, is relatively straightforward. The goal is to match the width of the data being transferred to the maximum width of the I/O bus. If an 8-bit-wide peripheral needs to transfer data across a 32-bit bus, byte packing circuitry would collect 4 bytes at a time from the peripheral and arrange them into a 32-bit long-word prior to having the main processor make the transfer across the bus. This reduces by a factor of four the amount of processor time wasted making four separate one-byte transfers.

Dedicating a DMA channel to I/O bus operations is another technique for minimizing the time the main processor is involved with I/O transfers. By letting the

Tom Powell, BSEE, vice-president of marketing, Synergy Microsystems

be a catch-all for "I think we could do it if we really wanted to—but we haven't quite done it yet."

One of the first limitations in transfer rates is a legitimate difference between theoretical and actual achievable performance. This has to do primarily with latencies in the processor, memory subsystem and associated controller hardware. Such latencies account for some discrepancies between theoretical and achievable transfer rates, but not all the differences that are often reported.

At its initial roll-out last year, DEC published some figures showing the superior performance of Turbochannel over some of the competing buses, showing both the theoretical and achievable performance. The figures stated that the achievable performance limits of Turbochannel's system with an architectural transfer rate of 100 Mbytes/s came out to 93 Mbytes/s. The figures indicated that with the same architectural rate, SBus achieved in the area of only 27 Mbytes/s, and IBM's MCA with an architectural rate of 40 Mbytes/s could achieve only 13. Actual transfer rates of a Turbochannel option, such as an FDDI (Fiber Distributed Data Interface) adapter, are in the 73 to 80 Mbytes/s range, says

DMA controller, which is optimized for moving data, take care of the transfer overhead, the processor is free to operate in parallel with the DMA by interleaving main memory accesses with I/O transfers.

The third technique for improved performance is to add localized intelligence to the boards on the I/O bus. Normally, an I/O board won't have its own on-board processor. This forces the main processor to handle high-level protocols and to communicate with the peripherals located on the I/O board. Not only does this use a great deal of compute time, but it also means the processor, rather than the DMA controller, must frequently be interrupted and forced to service the I/O bus to deliver those high-level commands to the I/O board. When the I/O board includes its own processor, memory and firmware, much of the overhead of running these protocols can be transferred to the I/O board itself. This frees the main processor to function in parallel with the I/O board and lets the on-board DMA controller handle the relatively low-level data-transfer functions.

The fourth method, optimized block transfer, combines several individual I/O data transfers into one high-speed burst transfer of several long-words. The bus is then held by the board making the block transfer until the entire transfer is finished. This eliminates the time-consuming bus arbitration overhead associated with making several individual data transfers. It also lets the I/O bus take advantage of burst-mode memory circuitry used by the latest microprocessors. Optimized block transfer techniques are crucial for reducing I/O bus impact on high-performance processors. With a properly designed block transfer interface, typical I/O bus performance can be increased by 300 percent.

Block transfers

Although byte packing, DMA and localized intelligence have a significant positive effect on processor performance, optimized block transfer techniques reduce processor overhead even more dramatically. A 68040 memory array, for example, typically supports a bandwidth of about 60 Mbytes/s. If an I/O bus board supporting optimized block transfer were transferring data across the bus at 5 Mbytes/s, it would use approximately 8 percent of the main memory bandwidth. But because the 68040 has large internal caches, it needs to access main memory only about 50 percent of the time. Consequently, the board supporting block transfer would only degrade processor performance by 4 percent. If an I/O bus board, however, were to continuously transfer 5 Mbytes/s of data without block transfer, the extra overhead involved would use 50 percent of the 68040 main memory bandwidth and by the same logic cause a significant 25 percent reduction in 68040 performance.

While many bus specifications define a block transfer protocol, few designers have fully optimized block transfer because of the significant cost and complexity it can add to on-board memory circuitry. Most of the latest microprocessors, like the 68040, now require burst-mode memory support. Burst-mode transfers are very similar to most block transfer protocols. For those memory designs supporting burst-mode accesses, supporting block transfers has become much more cost-effective.

Christensen.

What DEC failed to indicate in measuring its rate and reporting that of the other buses, claims Sun's Borrill, was that it called for 128-block DMA transfers in achieving the Turbochannel numbers, while the SBus number was derived assuming 4-block transfers. The existing SBus specification allows for 64-block transfers, says Borrill, which, if measured, would bring SBus into—and perhaps above—the Turbochannel performance range. The B version of the SBus specification allows for up to 128-block transfers. And, expect the next host—the latest Sparcstation—to be able to take advantage of the larger blocks.

In neither case is it indicated how the inclusion of additional options will affect performance and what the software and hardware overhead will do to overall system performance. It may be critical what specific applications the options address and how they could affect other than the straightforward DMA transfer rate.

Tuxedo junction

Within the limited category of defined "workstation buses," Micro Channel is the most formal: it has the most flexibility yet carries the most overhead. One way to look at the relative "formality" of a bus architecture is to look at the number of signal pins. DEC's Turbochannel is the most casual, sporting a minimum of 44 pins. As indicated, part of the parsimony resulted in multiplexing both address and data pins—somewhat limiting the expansion potential.

In comparison, SBus has 82 signal pins. Therefore, even eliminating the additional 32 signals for the address/data bus, SBus still has six more pins. And while this isn't a strain on the connector (both buses use 96-pin connectors, the remaining pins reserved for supply and ground potentials), the additional signals do require more on-board circuitry and software and hardware overhead.

IBM's Micro Channel, borrowed from the company's personal computer division, is the most formally attired of the workstation buses. It sports a whopping 136 signal pins to accomplish its job. In addition to 32 data and address lines, and the normal control signals, it includes an arbitration bus, transfer control bus and additional support signals. Designed as the system bus for IBM's PS/2 family of personal computers, MCA has full multiuser, multitasking capability and can accept multiple masters.

As a systems bus, MCA is totally independent of the processor/memory bus. In addition, the architecture permits any MCA master to get control of the bus and talk to any other adapter board—master or slave—or to the host processor. This flexibility with full arbitration and interrupt scheme provides the greatest degree of flexibility of any of the workstation buses—but at a price. MCA is the most complex of the workstation buses, calling for the most complex designs. In addition, its relatively small board (only 39 $in.^2$) is relatively large for a desktop, yet it's a tight fit for the complex interface required, especially so because the

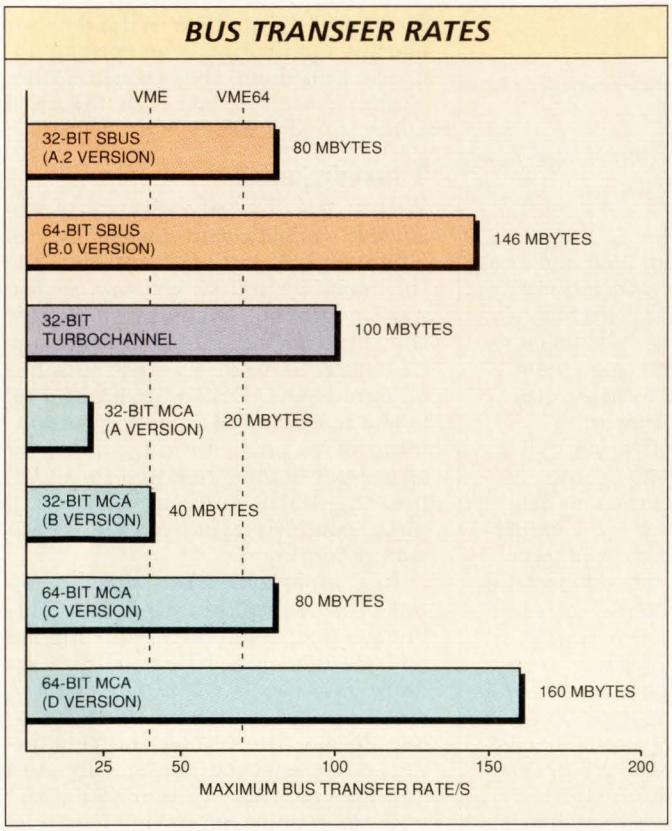

The theoretical maximum transfer rate sets a limit on bus performance that may be difficult to approach in applications. The mid-to-upper ranges of this group are achieved only with burst-mode transfers of 64-to 128-byte blocks. But these stripped-down buses are admittedly high-performers, as seen in comparison to the VME's transfer rates.

interface has to exist with equally complex option circuitry. The very low power available per slot adds considerably to the complexity of an option card and in many cases mandates the use of ASICs to curb power dissipation.

SBus is the middle of the three in terms of formality, providing some limited peer-to-peer communication and flexibility at the cost of some additional signal pins and circuitry. The bus is usually configured in a workstation environment such that the motherboard includes a number of SBus options internally with a few slots brought out for external option cards. At least two option cards already exist bridging SBus to VME, and a number of customers report using a VME subsystem to enhance a basic Sparcstation.

Turbochannel, as previously mentioned, requires the least in the way of interface circuitry and signal pins. Its relatively simple interface lets it interconnect with a host directly on the processor bus, through dual-ported memory, or in a crossbar configuration. This flexibility makes it easy to distribute processor power among a number of Turbochannel ports and can simplify integration into a host system.

Turbochannel is also the newest of the three workstation buses and has neither the critical mass nor support network of the others. "But," says Christensen, "we're rapidly developing a strong cadre of option makers including a number of vendors looking to provide Turbochannel systems. In addition, we will soon offer a new ASIC that will handle all the interface and control logic. This will not only simplify the option designer's task, but save board area and power."

"We tried to design a high-speed interconnect that would solve some I/O problems on the desktop and be capable of taking advantage of new technologies."

—Brenda H. Christensen, DEC

Micro Channel perhaps has the greatest support network, borrowed from its little PS/2 brother. IBM already boasts some 700 different adapter boards available from over 200 different vendors. And, though these boards were originally designed for the smaller form-factor personal computer, they are fully plug-compatible with the company's workstation. In addition, over a dozen different interface chips are already available for master and slave adapter boards as well as a variety of programmable interface chips from such vendors as PLX, Altera and others.

Despite IBM's apparent lead, Sun's SBus has all the earmarks of having the momentum on its side. SBus already counts over 150 option card vendors in its corner. While most are designing and producing SBus cards, some 25—and the number is growing—are currently designing both commercial and industrial SBus-based systems. These include a host of offshore companies developing Sparcstation clones.

Market decides

But ultimately, as with anything else, the market decides the winner(s). IBM has traditionally had problems in attacking the technical workstation market. Its earlier effort, the RT machine, survived a number of years before it succumbed to a lack of interest. In the almost five years of its existence, it managed to garner only about 20,000 sales, according to analysts. In only its first year of operation, it's estimated the RISC System/6000 will sell in excess of 40,000 units, with a majority in the high end of the price range. There's no indication its rate of sales will abate. Sun, as mentioned, has the momentum for the present and seems to be on the fast track. It will soon catch IBM with the number of option cards and vendors putting together workstations, but it will never approach the personal computer numbers. In terms of workstation sales, however, Sun will undoubtedly continue as the market leader, constantly boosting performance with each successive product generation.

DEC and Turbochannel are still the unknowns in the workstation business. Their Turbochannel-based workstations are still relatively new, having hit the streets only early last year. Turbochannel is slowly creeping up on Sun in terms of the number of board vendors it has attracted, but vendors are voicing caution before jumping on the new

bus, waiting to see which way the industry is going.

The particular bus architecture a company uses in its workstation may—or may not—determine the success of that company, or of the workstation. But, it can go a long way toward its acceptance by the industry. Sun, Silicon Graphics and others started out in the workstation business using a standard platform for an I/O bus (either Multibus I or VME). The use of the open bus encouraged third-party vendors to support the workstation platform and continue to develop additional modules. The fact that the bus was open, requiring no royalties, commitments or other allegiance, appealed to many of the third-party vendors, particularly those that were slowly strangled trying to sell into the DEC aftermarket.

Sun claims its SBus is as open as any of the IEEE-endorsed standards, requiring absolutely no royalties, registration or other paperwork. In fact, claims Wellens, Sun hopes to turn over control of the bus architecture to an independently formed group of third-party vendors as soon as possible.

DEC claims that Turbochannel, like SBus, is a totally open standard. It offers potential option board and system vendors complete specifications and details at only a nominal cost. The only restriction is that manufacturers must register their products with the company.

MCA is the only bus architecture that continues to be totally IBM proprietary and levies royalties on vendors using its specification. While the idea of royalties has frightened off some potential vendors, IBM's showing of vendors in its catalog supports its approach—design difficulties, design costs, royalties and all.

But the battle is far from over and the number of contenders certainly isn't limited. VMEbus, though not really a workstation bus according to our definition, has successfully been squeezed on a desktop, and its larger form factor lets a lot more circuitry be included. Similarly, Multibus II isn't out of the running. But, both of these approaches may be restricted to particular classes of applications because of relatively slow transfer rates and board size.

And Futurebus+ still looms on the horizon. Though the desktop profile was set back when its working subcommittee chairman was forced to step down, reports are that Hewlett-Packard, DEC and others are still actively pursuing the concept. This could change performance parameters and put a new set of players in the driver's seat. ■

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 Corp
(800) 678-6736 . . . . . **Circle 240**
IBM
(512) 823-0000 . . . . . **Circle 241**
Sun Microsystems
(415) 960-1300 . . . . . **Circle 242**
Synergy Microsystems
(619) 753-2191 . . . . . **Circle 243**
VFEA International Trade Association
(602) 951-8866 . . . . . **Circle 244**

The Best . . .

Starting at $85/MFLOP

The Best Price and Performance of Any off-the-shelf DSP/VMEbus in the world!

VALLEY ENTERPRISES, INC. offers you the BEST choice . . . The VE-32C-01V board starting at $85/MFLOP provides an excellent performance to cost ratio for VMEbus-based signal processing applications. This single board can perform up to 100 MFLOPS or 50 MIPS using four AT&T DSP32C digital signal processors. Designers can put up to 16 VE-32C-01V boards in a VMEbus chassis for a total system performance of 1.6 GFLOPS. A 10 MHz FIFO-buffered input data port and a 10 MHz non-buffered output data port allow the board to capture, process, and generate high-speed data streams.

CURRENTLY IN DEVELOPMENT FOR COMPLETION IN MARCH 1991: Companion A/D Module . . . *A 50 MSample/sec, 10-bit A/D converter on a single-slot 6U x 160mm VMEbus board. This board will interface directly to the high-speed FIFO port of the VE-32C-xxV family of DSP modules.*

". . . ***100 MFLOPS*** on a single board. Software support for target systems using VxWorks, OS/9, Motorola V/68 Unix, and SunOS. Call Valley about this unbeatable combination . . ."

". . . Software development tools and a DSP Library are available for the following host computers: SUN-3 and SUN-4, IBM-PC and compatibles, HP-9000 workstations, and Apple MacIntoshes. Target environments that are supported include VxWorks, OS/9, SunOS, and Motorola V/68 systems . . ."

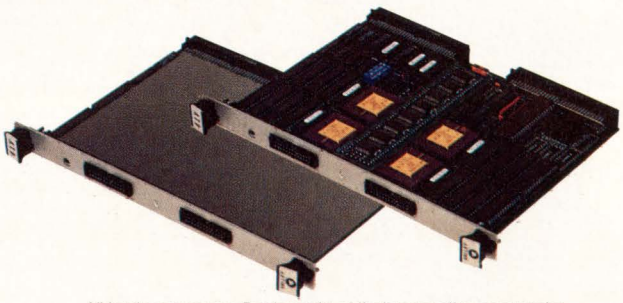

VALLEY ENTERPRISES, INC.
RD#4, ROUTE 309 • TAMAQUA, PA 18252 USA
Phone (717) 668-3737 • Fax: (717) 668-6360

All trade names are Trademarks of their respective companies.

CIRCLE NO. 47

40 FORESIGHT

TP41V DUAL VME/VSB FOR MULTIPROCESSOR APPLICATIONS

- MC68040 25-33MHz: 20MIPs, 3.5MFLOPs at 25MHz, Integral FPU, 2 x 4Kb Caches
- 2-16Mb DRAM onboard, Dual banked, Multiported to CPU, VME, VSB and LAN
- Segmented Memory Protection for VME/VSB access
- VME Rev C IEEE1014:
  - System Controller
  - Interrupter and handler
  - 16 vectored interrupts
  - DTB Master/Slave
  - BLT transfers • 4Kb Mailbox
- VSB Master/Slave
  - Interrupter and handler
  - 16 vectored interrupts
- Ethernet IEEE802.3 with 32bit DMA Controller
- 128Kb to 2Mb EPROM in 2 JEDEC sockets
- 2 x RS232 • 4 x CIO • 12 x Counter/Timers
- Software support includes:
  - pSOS+ᴹ • Velocity • VxWorks
  - TP-IX/68K V.3.2 implementation of UNIX

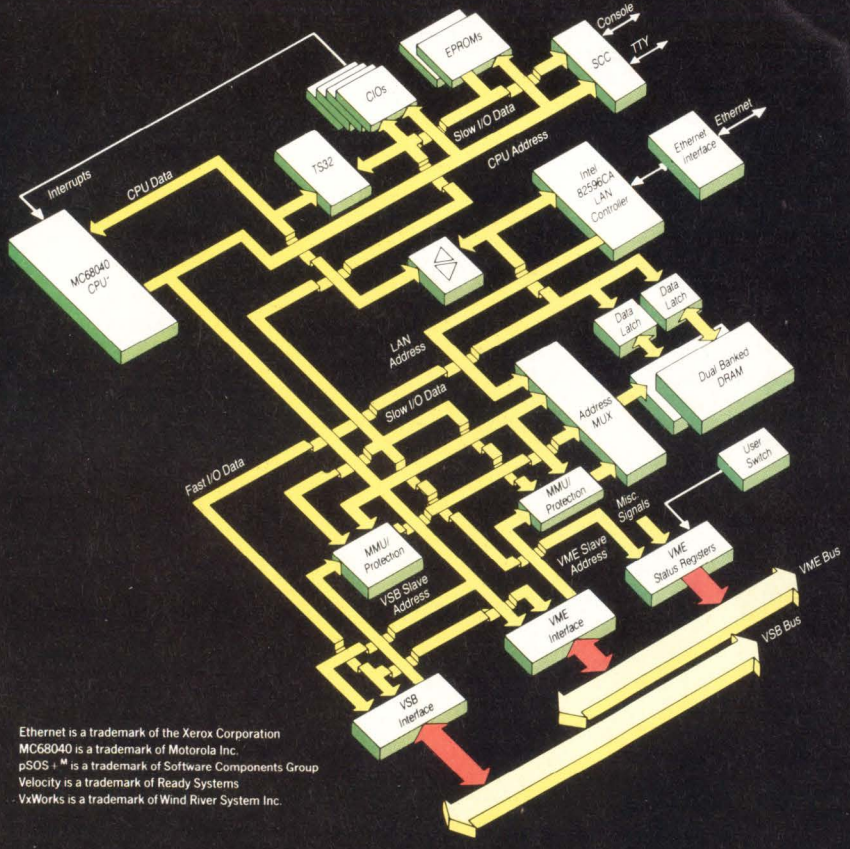

Ethernet is a trademark of the Xerox Corporation
MC68040 is a trademark of Motorola Inc.
pSOS + ᴹ is a trademark of Software Components Group
Velocity is a trademark of Ready Systems
VxWorks is a trademark of Wind River System Inc.

If you would like information on these, or any other Tadpole products, call us TOLL FREE on:

1 800 232 6656 for US enquiries

For European enquiries contact:

Tadpole Technology plc
Cambridge Science Park
Milton Road
Cambridge CB4 4WQ
ENGLAND

Tel: 0223 423030
Fax: 0223 420772

Tadpole Technology SA
44 Avenue de Valvins
77210
Avon-Fontainebleu
FRANCE

Tel: 331 60 72 50 60
Fax: 331 60 72 51 11

40 BY TADPOLE

TP40V VERY HIGH PERFORMANCE SINGLE BOARD VME COMPUTER

- MC68040 25-33MHz: 20MIPs, 3·5MFLOPs at 25MHz, Integral FPU, 2 x 4Kb Caches
- 2-32Mb Exchangeable DRAM Module, Dual banked, quad-ported to SCSI, VME, LAN and CPU
- VME Rev C IEEE1014:
  - System Controller
  - Interrupt handler
  - 4Kb Mailbox
  - DTB Master/Slave
  - BLT transfers
- SCSI interface on P2. 32bit DMA and NCR 53C90A achieving sustained data transfers up to 4.5Mb/sec
- Ethernet IEEE802.3 with 32bit DMA Controller
- RTC with 2Kb SRAM
- 8bit user I/O
- 2 x RS232 • 2 x CIO
- 6 x Counter/Timers
- Software support includes:
  - pSOS+™ • Velocity • VxWorks
  - TP-IX/68K V.3.2 implementation of UNIX

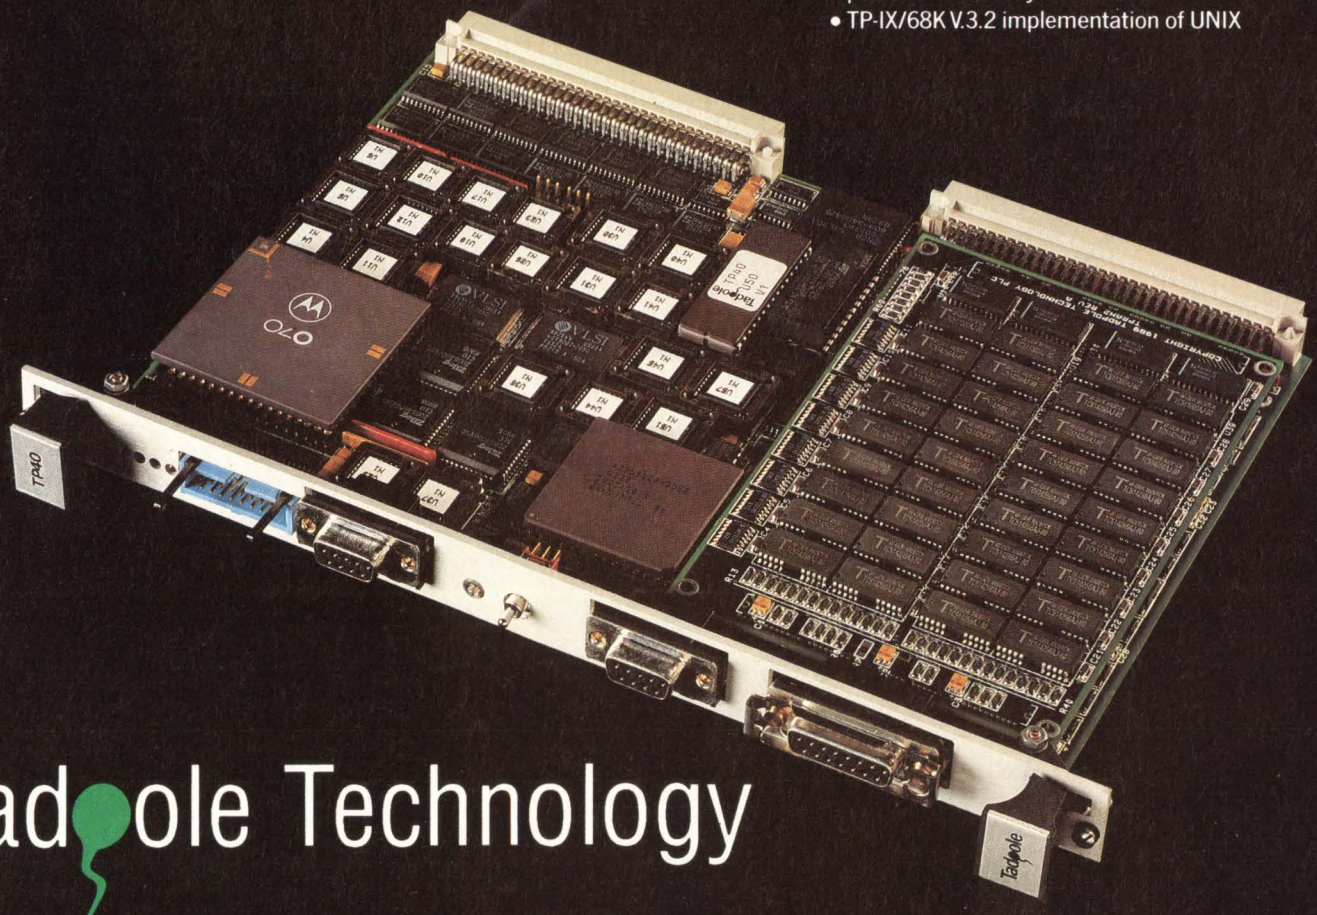

Tadpole Technology

CIRCLE NO. 48

WORLD CLASS ENGINEERING SOLUTIONS

That's what you'll find at Electro/International, the East Coast's only global electronics conference and exhibition for the entire engineering team with:

Over 700 exhibits of problem-solving products

38 professional/technical sessions and six tutorials organized in eight tracks: ■ Global Business ■ Computers, Software and Digital Systems ■ Medical ■ Communications ■ Manufaturing and Testing ■ IC Technology ■ Career ■ General Interest

Half-day short courses offering intensive learning opportunities ■ Keynote speech by Congressman Don Ritter ■ Purchasing conference ■ Blue ribbon panel exploring vital issues in technology, business and the profession ■ Newest developments in EDA tools in the Automated Design Center

Don't miss this outstanding educational opportunity with real bottom line returns for you and your company. Fax or write today for more information and plan now to attend.

April 16-18, 1991
Jacob Javits Convention Center
New York, N.Y.

Electro® International

Send me more information about Electro/International. I am interested in information about:
❑ exhibits ❑ technical/professional sessions ❑ short courses ❑ special events

Name _____ Address _____

Title _____

_____
City State Zip

Company _____

Fax: 213/641-5117 • Mail to: Electro/International Preview • 8110 Airport Blvd. • Los Angeles, CA 90045

CIRCLE NO. 49

INTEGRATED CIRCUITS

High-speed DACs target waveform synthesis and video

Jeffrey Child, Associate Editor

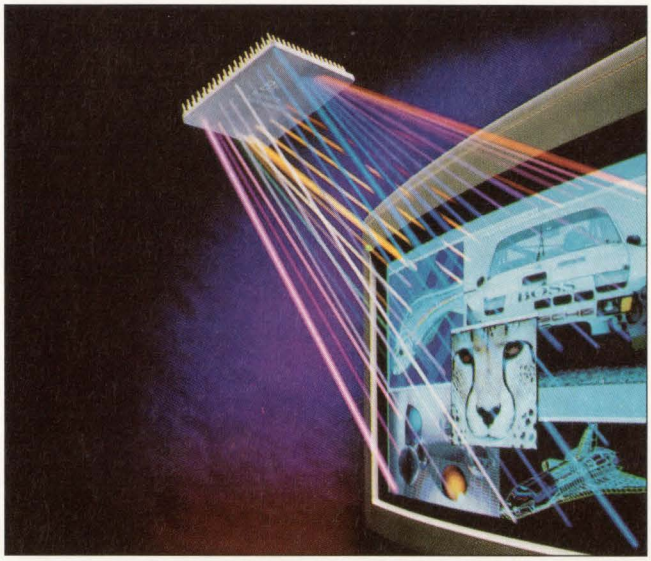

Supporting today's high-resolution displays, Brooktree's RAMDACs provide high-speed digital-to-analog conversion as well as sophisticated graphics and video functions. Application programs, stored in on-chip RAM, let the user custom-tailor images to fit requirements.

Applications that rely on high-speed digital-to-analog converters range from CAD to image processing to waveform synthesis. Across this wide range, these applications share one common aspect: they're becoming more complex. As faster, more-powerful processors and larger memories expand the practical limits of digital data manipulation, the critical link to the analog world provided by D-A converters is pressured for ever-higher levels of performance. Growing volumes of graphic and video information flow from computer systems to high-resolution displays. To fill the speed requirements needed to support these applications, today's video D-A converter manufacturers are pushing the limits of mixed digital and analog design, and asking the seemingly impossible of semiconductor processes.

Beyond raw speed and resolution, today's high-speed D-A converters also include critical support functions. Offering update speeds of 100 MHz and up, these devices are sporting more on-chip functions that support increasingly sophisticated video and graphics operations. To drive color displays, for example, video D-A chips contain three D-A converters, one for each of the primary colors. Design activity in sophisticated video applications is evident in the numbers: video D-A converters now make up the largest segment of the high-speed D-A converter market.

Innovations in process technologies, particularly in bipolar and gallium-arsenide, have enabled several D-A converter manufacturers to produce devices with blinding-fast settling times and update rates. While the need to support higher-resolution monitors has driven D-A converter speeds up, systems designers interested in direct digital synthesis (DDS) are looking to the latest technology to synthesize waveforms at frequencies higher than previously possible. In response, D-A converter manufacturers are building devices specifically for broad bandwidth waveform synthesis.

No noise is good noise

In D-A converters targeted for video, so-called RAMDACs, speed relates directly to resolution of the monitor driven by the D-A converter's signal. The faster the D-A conversion, the higher the resolution that can be supported. Any speed beyond that required for the desired resolution is unnecessary. Designers of video D-A converters agree, however, that achieving high speeds is not the biggest hurdle. The greatest challenge is deciding which digital functions to include on the chip—without producing noise to hinder the performance of the analog portion.

"If it's a stand-alone product with no digital circuitry nearby, you can usually work for and achieve almost any level of analog performance," says Dan Watson, applications engineering manager at TRW LSI Products (La Jolla, CA).

"In mixed-signal parts there's always a compromise to make," Watson says. "On a tiny piece of silicon you've got your precious D-A converters with a certain level of analog performance that can be degraded by digital circuitry a tenth of an inch away. Noise from that has the potential to couple over to the analog signal. You know you're going to have to pay a certain penalty in analog performance because of all the digital circuitry in the vicinity. And you have to play the technical games to trade that off."

A mixed-signal device, the TMC0458, is TRW's most recent D-A converter targeted for video. The device contains a 256×24 color palette and three 8-bit D-A converters. Structured with multiplexed inputs to the color palette, the converter supports pixel updates to 170 MHz without requiring any ECL circuitry. Dual-port RAM on the chip permits asynchronous access by the host computer.

Specialized functions

Because video D-A converters generate the display signal that the user actually sees, the converters lie in a critical path of system performance. Video D-A converter manufacturers must be sensitive to the user's evolving needs. The growth of windowing environments stimulates the need for specialized digital functions in video D-A converters.

With this in mind, Brooktree (San

Model	Resolution (bits)	Update rate (MHz)	Output (voltage, current)	Typical integral nonlinearity (± LSB)	Internal reference (Y/N)	Gain error percent (FSR)	Settling time (ns) (± 1/2 LSB)	Power dissipation (mW)	Price	Comments
Analog Devices 2 Technology Way, Norwood, MA 02062 (617) 461-3557 **Circle 301**										
AD9768	8	100	C	0.5	Y	—	5 µs	410	$18.75	monolithic, 20 mA output, ECL-compatible
AD9712	12	100	C	3	Y	4	30	700	$40	monolithic, very low glitch energy for direct digital synthesis
AD9720	10	300	C	0.5	Y	—	10	1,000	$40	—
Brooktree 9950 Barnes Canyon Rd, San Diego, CA 92121 (619) 452-7580 **Circle 302**										
Bt468	3×8	200	RS343	0.5	N	—	0.5 µs	2,000	$177	RAMDAC
Bt492	8	360	RS343	0.5	Y	—	0.5 µs	3,500	$293	RAMDAC
Bt107	8	400	RS343	0.5	N	—	0.5 µs	1,000	$100	video D-A converter
Bt109	3×8	250	RS343	0.5	Y	—	0.7 µs	2,000	$52	triple video D-A converter
Bt451	3×4	125	RS343	0.5	N	—	0.5 µs	1,000	$61	RAMDAC
Bt454	3×4	170	RS343	0.125	Y	—	0.7 µs	1,000	$81	RAMDAC
Bt455	3×4	170	RS343	0.125	Y	—	0.7 µs	1,000	$71	RAMDAC
Bt457	1×8	135	RS343	—	—	—	0.7 µs	—	$57-$72	RAMDAC
Bt458	3×8	165	RS343	0.5	N	—	0.7 µs	1,000	$105-$141	RAMDAC
Bt459	3×8	135	RS343	0.5	N	—	0.7 µs	1,100	$107-$120	RAMDAC
Bt460	3×8	110	RS343	0.5	N	—	0.9 µs	1,100	$139	RAMDAC
Bt461	8	170	RS343	0.5	N	—	0.5 µs	1,700	$160	RAMDAC
Bt462	8	170	RS343	0.5	N	—	0.5 µs	1,700	$143-$160	RAMDAC
Bt463	3×8	170	RS343	0.5	N	—	0.5 µs	2,000	$332	RAMDAC
Fujitsu Microelectronics, IC Division 3545 N First St, San Jose, CA 95134 (800) 642-7616 **Circle 303**										
MB40788	10	125	V	0.2	Y	—	3.0	450	$53.47	—
GEC Plessey Semiconductors 1500 Green Hills Rd, PO Box 660017, Scotts Valley, CA 95062 (408) 439-6077 **Circle 304**										
ZN454	3×4	100	V	0.25	Y	5	8	750	$3.25	video
ZN455	3×4	100	V	0.25	Y	5	8	750	$3.65	video
ZN515	8	100	V	0.25	Y	5	8	600	$19.82	video
SP98608	8	450	V	0.5	Y	1.0	2.2	650	$49.55	video
Signal Processing Technologies 1510 Quail Lake Loop, Colorado Springs, CO 80906 (719) 540-3900 **Circle 305**										
HDAC10180A	8	275	C	0.95	N	5	7	1,040	$27.10	—
HDAC10180B	8	165	C	0.95	N	5	7	1,040	$22.30	—
HDAC10181A	8	275	C	0.95	Y	15	7	1,040	$27	—
HDAC10181B	8	165	C	0.95	Y	15	7	1,040	$22.30	—
HDAC51400B	8	400	C	0.95	Y	5	4	1,040	$75.40	—
Sony Corp of America 10833 Valley View St, PO Box 6016, Cypress, CA 90630-0016 (714) 220-9100 **Circle 306**										
CXA1236Q	8	500	C	0.5	Y	7	1.5 µs	950	$95	—
CXA1156Q	8	360	C	0.5	Y	7	2.5 µs	1,200	$45	—

Model	Resolution (bits)	Update rate (MHz)	Output (voltage, current)	Typical integral nonlinearity (±LSB)	Internal reference (Y/N)	Gain error percent (FSR)	Settling time (ns) (±1/2 LSB)	Power dissipation (mW)	Price	Comments
Sony Corp of America 10833 Valley View St, PO Box 6016, Cypress, CA 90630-0016 (714) 220-9100										**Circle 306**
CX20201A-1	10	160	V	0.5	N	0.1	5.2 µs	390	$77	—
CX20202A-1	10	160	V	0.5	N	0.1	5.2 µs	420	$77	—
CX20201A-2	9	160	V	0.5	N	0.1	4.7 µs	390	$45	—
CX20202A-2	9	160	V	0.5	N	0.1	4.7 µs	420	$45	—
CX20201A-3	8	160	V	0.5	N	0.2	4.3 µs	390	$26	—
CX20202A-3	8	160	V	0.5	N	0.2	4.3 µs	420	$26	—
Tektronix ICO PO Box 500, M/S 59-420, Beaverton, OR 97077 (503) 627-2515										**Circle 307**
TKDA10C	6	100	C	0.005	N	—	4	470	$56	6-bit part with over 12-bit accuracy, data link escape 0.003 percent (typical)
TKDA30	12	100	V	0.25	N	1	—	1,000	$118	on-chip latches, analog multiplying input, 65-dBc spurious free dynamic range (typical)
TriQuint Semiconductor PO Box 4935, Beaverton, OR 97076 (503) 644-3535										**Circle 308**
TQ6113	8	600	V	1.0	N	—	2.0	2,500	$183	gallium-arsenide
TQ6114	8	1,000	V	1.0	N	—	2.0	3,500	$272	gallium-arsenide
TRW LSI Products PO Box 2472, La Jolla, CA 92038 (619) 457-1000										**Circle 309**
TDC1034	4	200	C	0.125	N	6	5	600	$10	synchronous, blank video out
TDC1334	3×4	200	C	0.125	Y	6	5	1,150	$15	synchronous, blank video out, RGB
TDC1018	8	200	C	0.5	N	5	8	676	$15	synchronous, blank video out
TDC1318	3×8	200	C	0.125	Y	7	8	1,820	$36	synchronous, blank video out, RGB
TMC0176-11	3×6	110	C	0.5	N	5	9	950	$6.39	256-word color palette
TMC0458-17	3×8	170	C	1	N	5	8	1,550	$97	256-word color palette, 4 overlay
TMC0458-12	3×8	125	C	1	N	5	8	1,550	$83	256-word color palette, 4 overlay
TMC0458-11	3×8	110	C	1	N	5	8	1,550	$79	256-word color palette, 4 overlay

Diego, CA) introduced the first RAMDAC that provides separate color maps for windows displaying different visual modes in a single frame. Developed jointly with Digital Equipment Corp (Maynard, MA), the 8-bit, 170-MHz Bt463 RAMDAC combines true-color and pseudocolor graphics with hardware windowing functions. The new RAMDAC features a flexible, window-oriented architecture that covers the full range of graphics, imaging and video display requirements for future platforms.

"The Bt463 has two key features," says Joe Santen, vice-president of corporate communications at Brooktree. "First, it has the ability to put up to 16 windows on a screen at one time. Second, you can put different types of visual information in each window. You can have, for example, a true-color image in one window, a gray-scale image in another, a solid-object model in another, and you could have live motion video in another."

The Bt463, a CMOS mixed-signal part, was built on a purely digital process, with no laser trimming and no linear CMOS, according to Santen. While Santen believes achieving 170-MHz speed in a complex mixed-signal CMOS part is a technical breakthrough, he insists the more significant success was in defining the part. "It's one of the few times when a component manufacturer or chip maker has actually had to get massive inputs from end-users in the design," he says.

Targeted for ultra high resolution monitors, the latest video D-A converter from Sony Corp of America (Cypress, CA) is the 8-bit, 500-MHz single-video model CXA1236Q. The device is built on Sony's ECL-3 pro-

INTEGRATED CIRCUITS

cess, the same technology on which the company bases many of its bipolar products. ECL-3 gives the chip both high speed and low-power dissipation. On-chip functions contained on the CXA1236Q include composite sync, blank, reference white, enhanced white/black, and bright (overlay) control inputs. With these functions the user can minimize memory costs, according to Sony, because some of the processing functions once required of memory can be performed by the video D-A converter. Glitch energy of the converter is less than 5 pV-s.

DDS renewed

While holding a smaller market niche than video D-A converters, devices targeted for DDS of waveforms are seeing renewed life. While DDS has been around for a long time, the speeds of the latest devices have enabled synthesis of waveforms at much higher frequencies. Besides speed, the key issue for D-A converters targeted for DDS applications is minimizing the noise of the signal. "The bottom line in DDS is how purely you can generate a sine wave," says Dave Buchanan, DDS product marketing engineer at Analog Devices (Greensboro, NC).

Actually, the characteristics of a good waveform D-A converter are desirable in a

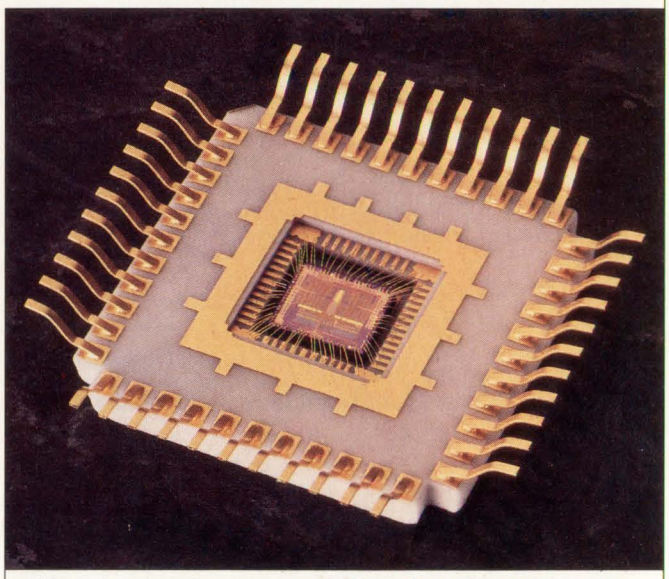

Although currently not a standard product, the 14-bit, 1-GHz digital-to-analog converter from TriQuint Semiconductor is optimized for direct digital synthesis (DDS) applications. The devices make possible DDS of high-bandwidth signals.

If you thought Multibus 1 couldn't handle complex tasks. . . *Think Again!*

From Robotics to CAD, Embedded Controllers to Host and Network Systems, MULTIBUS 1 by Zendex delivers Proven, Cost Effective Solutions that reach out and grab you!

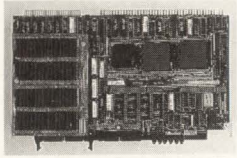

The ZX-386/20 pictured here features *20MHz* clock speed and on-board *SCSI*, supports up to 8MB of high-speed on-board *dual-port* RAM (or 16MB in PVAM), and has two 8/16-bit iSBX connectors. For Real-Time processing applications, Zendex's *TRU-32* mode turns MULTIBUS 1 into a *screaming 32-bit system* capable of addressing *256 MB* of RAM, yet is fully compatible with MULTIBUS 1 products.

For more information on our complete family of MULTIBUS 1 products, call or write Zendex today!

ZENDEX
MULTIBUS 1 SINCE 1979

6780 Sierra Court
Dublin, CA 94568
Phone: (415) 828-3000
FAX: (415) 828-1574

CIRCLE NO. 50

i860

SUPERCOMPUTER PERFORMANCE

The AL860 VME or AT Board
WILL GIVE YOU:

- An Intel i860 RISC microprocessor rated at 80 MFLOPS at 40 MHz (66 MFLOPS at 33 MHz)
- 2 to 64 Mbytes of 0 wait state local memory
- Optional SCSI and DT-Connect interfaces, and 120 Mbyte/sec interconnection bus
- DOS and UNIX interfaces, compilers, and scientific subroutine library software available

ALACRON Inc.

71 Spitbrook Road, Suite 204 • Nashua, NH 03060
Telephone: (603) 891-2750 • FAX: (603) 891-2745

Trademarks are property of their respective holders.

CIRCLE NO. 51

video system because they typically offer sufficiently high resolution—8 to 12 bits. And noise and distortion are usually low enough, since a video D-A converter needs to have good specifications in those areas to be useful.

"It would be pretty safe to say that waveform synthesis has come up to speed through using video D-A converters," says Buchanan. "For a long time there were very few D-A converters around that were suitable for waveform synthesis, and then as video D-A converters became available, people started doing DDS.

"Waveform synthesis at such high bandwidths wasn't possible before, because D-A technology wasn't fast enough," Buchanan continues. "Video D-A converters have come up to speed, and now it's the other way around: people are really interested in DDS now."

In response, companies are designing parts specifically for that application. The AD9712/13, Analog's 12-bit monolithic converter, offers a choice of output configurations. The version with ECL-compatible outputs can do updates at 100 MHz. A TTL-compatible version of the part updates at 80 MHz.

Designers of DDS systems need greater bandwidths so they can synthesize waveforms with components at higher frequencies. The fastest D-A converter available today is a 1-GHz, 8-bit part from TriQuint Semiconductor (Beaverton, OR). Designed on a gallium-arsenide process, the TriQuint device sampling rate allows simulating a 500-MHz bandwidth. TriQuint has also developed a 14-bit, 1-GHz D-A converter for proprietary use. A gallium-arsenide part, the 14-bit device could become a standard product in 1992, according to TriQuint.

Another speedy device is the latest D-A converter from GEC Plessey Semiconductors (Scotts Valley, CA). The SP98608 is an 8-bit device clocking speeds up to 450 MHz. One of this converter's design features is reduced glitch error. On-board capacitances and timing inconsistencies of incoming data, when presented to current switches, often result in glitch problems. By incorporating a latch/buffer section between the data inputs and the current switches, these effects are minimized as the latch enables time-synchronous data to be presented to the current switches. The output of the latches are strobed internally so that output is synchronized to achieve on-chip timing within 20 ps. In addition, the differential architecture, through to the current switches, reduces the glitch energy output to 20 pV-s. ■

RUN IT UP THE FLAGPOLE AND SEE WHO SALUTES.

FOR FREE.

Before Omation sells you any schematic design software, we offer you a free examination of the goods. Then, when you see how fast, easy and affordable PC-based design can be, you can take the plunge.

Using new SCHEMA III, you can create, verify, print and plot your designs, as well as communicate with PC or workstation board layout software. For $495, you get all the features you demand from design software with all the support and service you deserve, including our FREE 1-800 technical support hotline. Plus, a 30-day money back guarantee.

Call Omation for your FREE demo disk. And salute the top ranking PCB design software on the market today.

SCHEMA FROM OMATION

OMATION • 801 Presidential Drive • Richardson, Texas 75081 • 1-800-553-9119

CIRCLE NO. 52

ONE MILLION TRANSISTORS AND

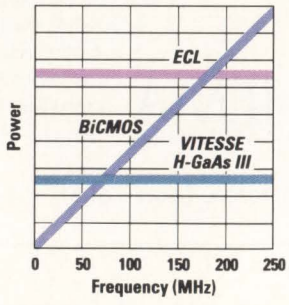

We're talking better than the best ECL performance at a fraction of the power. And for high frequency designs, Vitesse GaAs chips are lower power than BiCMOS.

Our prices won't sink your budget either– our GaAs chips can give you better performance for your dollar than BiCMOS.

THE BLINDING SPEED OF GaAs.

For more information on our ASIC and standard products with integration levels up to 350K gates, call Vitesse at (805) 388-7455. And leave the competition in your wake.

FX Arrays	Raw Gates	Usable Gates	Availability
FX100K	102,000	Up To 70,000	Now
FX200K	195,000	Up To 137,000	Now
FX350K	353,000	Up To 177,000	Summer 1991

VITESSE
The GaAs Company.™

CIRCLE NO. 67

INTEGRATED CIRCUITS

New chip offers nonvolatile analog memory

There are a lot of digital systems in which it would be handy to have a little bit of audio-frequency analog data stashed away. Audio I/O devices such as pagers and featurephones could have audio prompting. Hand-held computers could have speaking icons.

But the hardware necessary to achieve these conveniences has been forbidding. Analog storage has meant tape recording, involving serial-only access, lots of mechanical parts, lots of stuff to wear out or wrap around a capstan, and limiting environmental restrictions. Digital recording has meant preamplifiers, sample-and-hold amplifiers, analog-to-digital converters and, often, an increase in the design's memory budget. Either way, the growing complexity of the design quickly discourages managers from including analog storage facilities.

A new technology from start-up Information Storage Devices (ISD), however, offers an inexpensive, incredibly easy way to capture, randomly retrieve and replay audio-frequency analog data. The breakthrough is to use a more-or-less conventional EEPROM cell to store not a binary bit but an analog voltage.

Analog EEPROM cells

"It's a monumental technical challenge," says ISD cofounder, vice-president of engineering and CEO Dick Simko. "You're trying to store a roughly 2-V analog signal in a cell that has a 20-V write level, and then get the voltage back again when you read the cell." One measure of the number of tricks required to make it all work is that ISD expects to be granted about 10 patents on its first device.

In its first announced chip, ISD has built an array of analog storage cells configured to save and retrieve consecutive samples. Obvious applications are voice recording and playback, sourcing for test signals, or temporary storage of telemetry. The device can achieve a maximum sample rate of 8 kHz and can store up to 16 s of continuous data.

On the chip ISD has included all the components necessary to perform the store and replay functions: an analog preamp, automatic gain controlled input amp and output amp, 3.4-kHz anti-aliasing and smoothing filters, clock generators and sampling circuitry, and addressing logic. The result is a device so simple that you can wire it to a power supply, a microphone and a speaker and start recording. "We aimed for complete ease of use," claims ISD director of marketing Jim Oliphant. "You can literally learn to use the chip from a one-paragraph data sheet."

The address inputs to the chip select from 160 0.1-s segments of data. These data segments can be selected randomly or in sequence. During sequential operation, some clever timing and handshaking circuitry ensures that the chip can move from the end of one segment to the beginning of the next smoothly, preventing clicks, pops or dead spaces in a long message. To provide for storage capacities longer than 16 s, handshaking signals permit cascading of the chips, again without any interruption to the record and playback functions.

Some limitations

There are some limitations to the current devices and to the technology. For now, ISD is limiting the sampling rate to 8 kHz. This provides adequate bandwidth for voice, tone generation and similar applications, gives a useful recording time from a relatively small memory array, and doesn't stress the cell technology. Simko claims that in the long run, the technology should be able to achieve a 44-kHz sampling rate, giving enough bandwidth for high-fidelity music.

Noise in the reproduced signal is detectable, but it's not sufficient to interfere with voice applications. The designers estimate that the overall fidelity of reproduction in their current device is approximately equivalent to that of an 8-bit digital system.

Since the underlying storage element is an EEPROM cell, the questions of signal integrity and cell life are pertinent. "In principle, there's a degradation in the signal-to-noise ratio over time," admits Simko. "But in practice it's not a significant factor. We expect that the devices will retain useful signals for about 10 years."

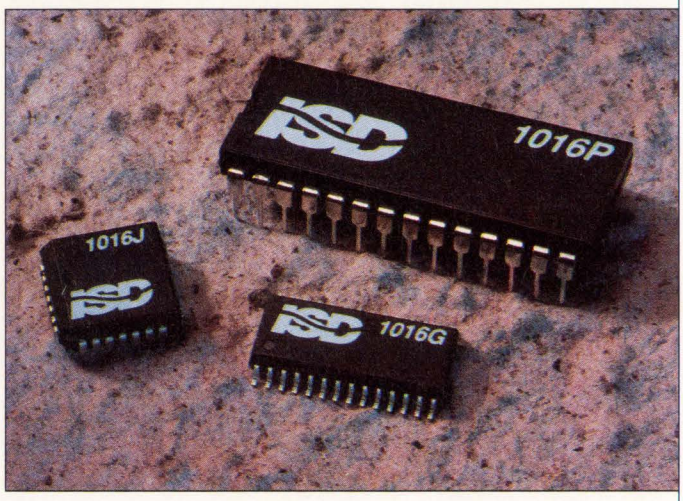

Like other EE devices, the ISD arrays have a finite cycle life. But, as in modern digital EEPROMs, it's only an issue for a few write-intensive applications. "We've seen 100,000 cycles without degradation of the cells on laboratory samples," Simko says. "But we'll probably spec the chips at around 10,000 cycles, just to be conservative." Such a specification might make the devices less than ideal in applications such as analog FIFO buffers, where continuous rewriting is necessary.

The ISD 1016 chip is available in samples now, and the company offers three packaging options. The device will be priced at about $20 each in 1,000-unit quantities.

—Ron Wilson

Information Storage Devices
2332 B Walsh Ave, Bldg G
Santa Clara, CA 95051
(408) 562-9550

Circle 351

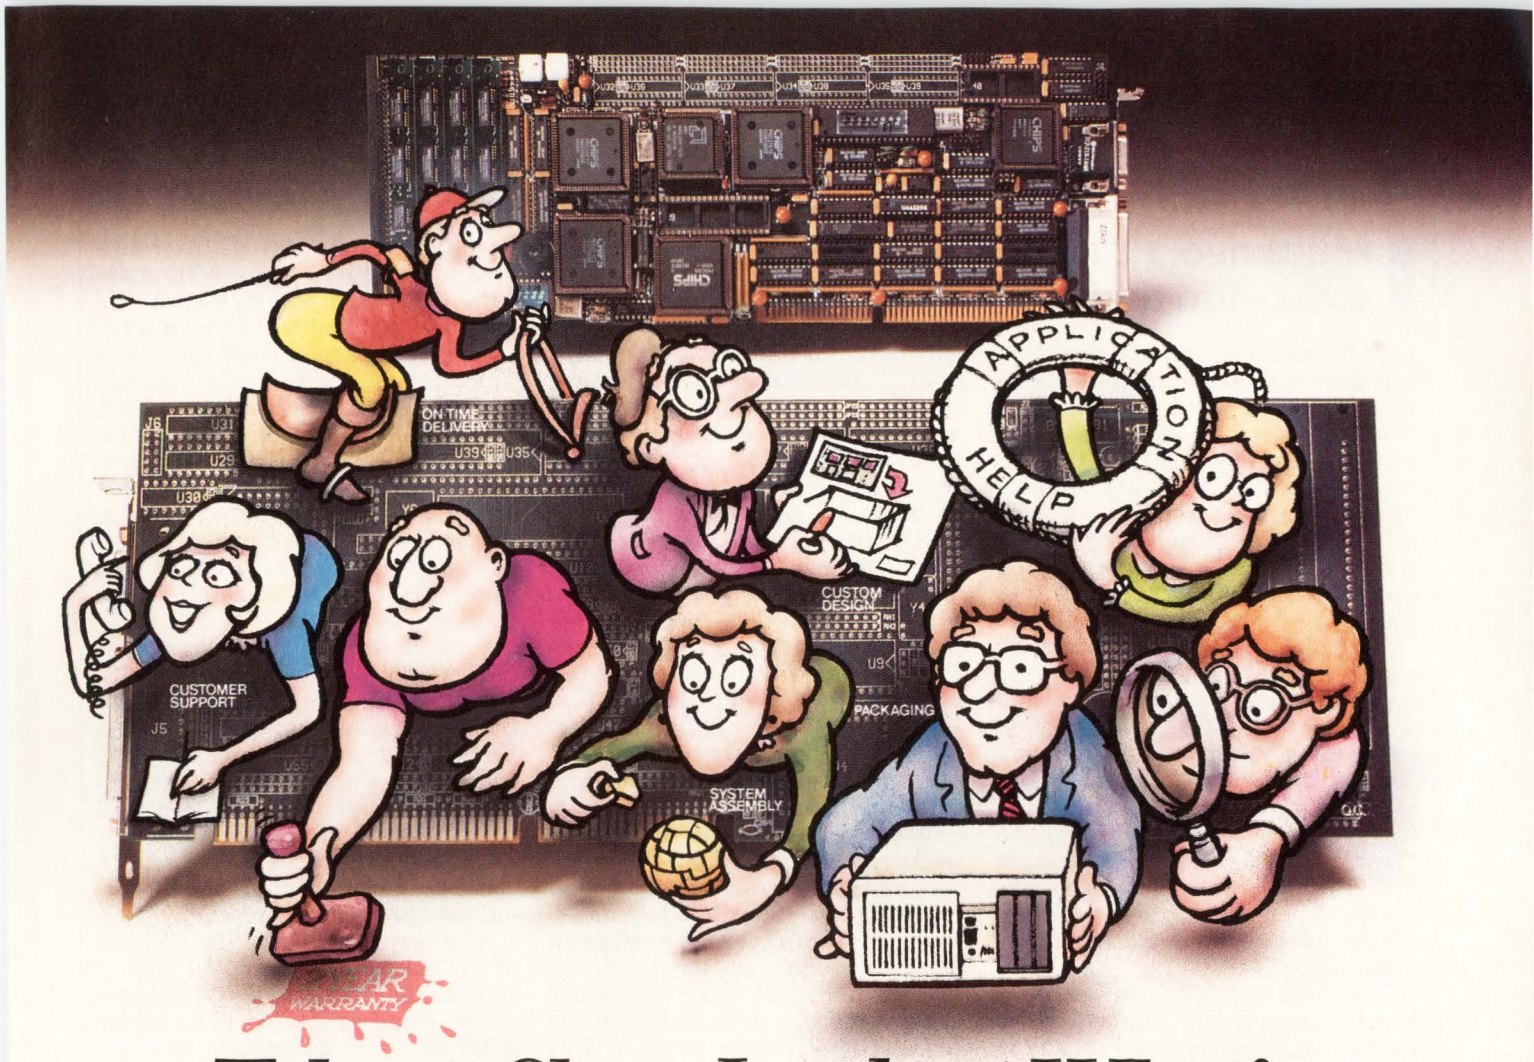

Take a Close Look at What's Behind Your PC CPU Board!

I-Bus has a lot more to offer the OEM and System Integrator than a family of CPUs with great price/performance. The I-Bus T.O.P.S.™ (Total OEM Program Support) Plan is designed specifically for OEMs and System Integrators:

- *Engineering and Manufacturing Services*
- *Program and Project Management*
- *Peripheral Sourcing*
- *OEM Developers Program*
- *Life Cycle Support Program*
- *Unprecedented 2 Year Warranty*

But that's only the beginning. We'll completely assemble and test your system (our boards and/or yours), modify or custom design an enclosure, and take it all the way to drop shipment to your customer.

We're 100% dedicated to the OEM and System Integrator, utilizing PC technologies in their application and products. Call today about the I-Bus OEM Developer's Program—a total support program that saves you time and cost in the design phase, and continuous support over the life of your system.

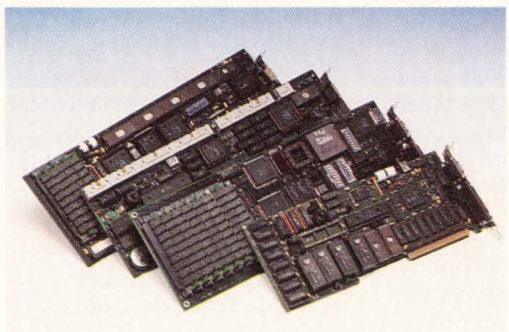

CPUs for the OEM and System Integrator

A286	8/12/16 MHz 80286 4 MB on-board DRAM Solid State Disk Option
J386SX	16 MHz 80386SX 4 MB on-board DRAM Floppy/IDE Hard Disk Controller
J386	16/20/25 MHz 80386 16 MB on-board DRAM
N88	10 MHz 8088 1 MB on-board DRAM Solid State Disk Option

I-BUS *PC Technologies*

Call toll free
800-382-4229
Ask for our new catalog!

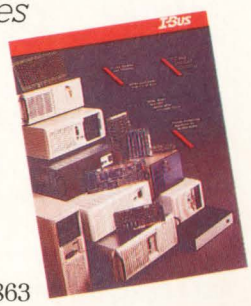

I-Bus, 9596 Chesapeake Drive, San Diego, CA 92123 (619) 569-0646 FAX: 619-268-7863

CIRCLE NO. 55

INTEGRATED CIRCUITS

Multimedia chip compresses video and audio in synch

A multimedia processor chip that can compress and decompress motion video signals in real time along with synchronized audio has been introduced by UVC. The UVC7710 integrated multimedia processor boasts compression ratios of 20:1 and 30:1 at 30 frame-per-second frame rates and even has a high compression mode with a ratio of 500:1.

The chip uses a patented compression technique that UVC claims is about 50 times less complex than the other conventional method of color video compression, discrete cosine transform (DCT). DCT relies on a tiled approach that works with the values of pixels surrounding the pixel of interest. DCT thus must first look at several scanned lines before it can operate on any data, and it uses complex mathematical algorithms.

The UVC approach, on the other hand, compresses a frame of video data by working on one scan line at a time. As each line is scanned, it's sampled, and the analog signal is digitized according to UVC's own algorithms. That digitized data is then compressed using data compression techniques. Digitization and compression are done "on the fly" because digitizing a frame by first storing it in a frame buffer and then processing it just doesn't work for motion video.

Video plus audio

In addition to video processing, the UVC7710 can accept and compress digitized audio using pulse code modulation with full 12-bit companding. The UVC7710 interfaces directly to the Texas Instruments TLC32045. In order to compress and decompress full video with audio, the UVC chip generates audio frame markers to synchronize stored audio and video. The marker code allows complete tracking of separately stored audio and video data. The marker code is automatically removed when the audio data is decompressed.

The simplicity of the compression techniques lets them be implemented completely in silicon and at a gate count far lower than other techniques. This in turn allows more system features—including bus timing and control, memory interface logic and a single clock with video timing for NTSC and PAL on the chip—to be placed on-chip. The UVC7710 can also be used to display video images on a 640×480-pixel VGA display with a scalable window. The chip can also capture and store compressed images in a 128-kbyte VRAM from which the images can later be transferred to disk storage.

The UVC7710 offers software-controlled selection of such parameters as size of a window on screen, frame rate and resolution to help system designers cope with bandwidth constraints. These parameters are also user-selectable at the receiving end. If during a teleconference, for example, one is watching a person speaking, one might select a high frame rate with somewhat lower resolution to have the feel of a living conversation partner. But when a chart is displayed, the user could switch to a lower frame rate and higher resolution to look at the details. The size of an on-screen display window can also be set by the receiver.

Multimedia 1 board

UVC will also be offering a Multimedia 1 audio/video processing board along with software that includes drivers for developing applications as well as some ready-to-run applications. The processor board is AT compatible and contains the UVC7710 and between 64 and 256 kbytes of VRAM. It's aimed at a wide range of multimedia, video imaging and teleconferencing applications.

The applications supplied with the board include a Multimedia PC application that lets users capture and edit real-time audio and video and convert between various video formats. A video conferencing application lets users initiate and receive calls using audio and video as well as select different privacy parameters (such as no video or video-only) during a conference. A third application is video mail, which works much like present electronic mail systems except that pictures are sent and received along with audio rather than just text messages.

The software also includes a set of authoring tools so that users can build custom applications. The tools include a set of drivers for audio, video, VGA, hard disk, scanner, VCR, and communications. Function libraries include facilities for LANs, modems, format conversion, editing, and database management.

In addition to the software that comes with the Multimedia 1 board, there's a stand-alone software package called Multimedia 1 Plus. This package lets users with an 80286-based or higher platform decompress and view/hear data that has been compressed by the UVC7710. Thus, some workstations in an operation that might not need to do full compression and capture but might only need to retrieve and view files could perform those functions without adding special hardware.

UVC says it's aiming its products at the mass market, and the pricing reflects this aim. The UVC7710 processor chip is expected to be under $100 in OEM quantities, and the Multimedia 1 board will be under $1,000. *—Tom Williams*

UVC
16800 Aston St
Irvine, CA 92714
(714) 261-5336

Circle 352

Trademark Information

UNIX is a registered trademark of AT&T Bell Laboratories.

PAL is a registered trademark of Advanced Micro Devices, Inc.

SMART-POWER is a registered trademark of Nartron Corp.

CDA and BurstRAM are trademarks of Motorola, Inc.

SCOPE and ASSET are trademarks of Texas Instruments, Inc.

IRIS POWERVISION is a trademark of Silicon Graphics, Inc.

RealTimeX is a trademark of Concurrent Computer Corp.

Zone Bit Recording is a registered trademark of Seagate Technology, Inc.

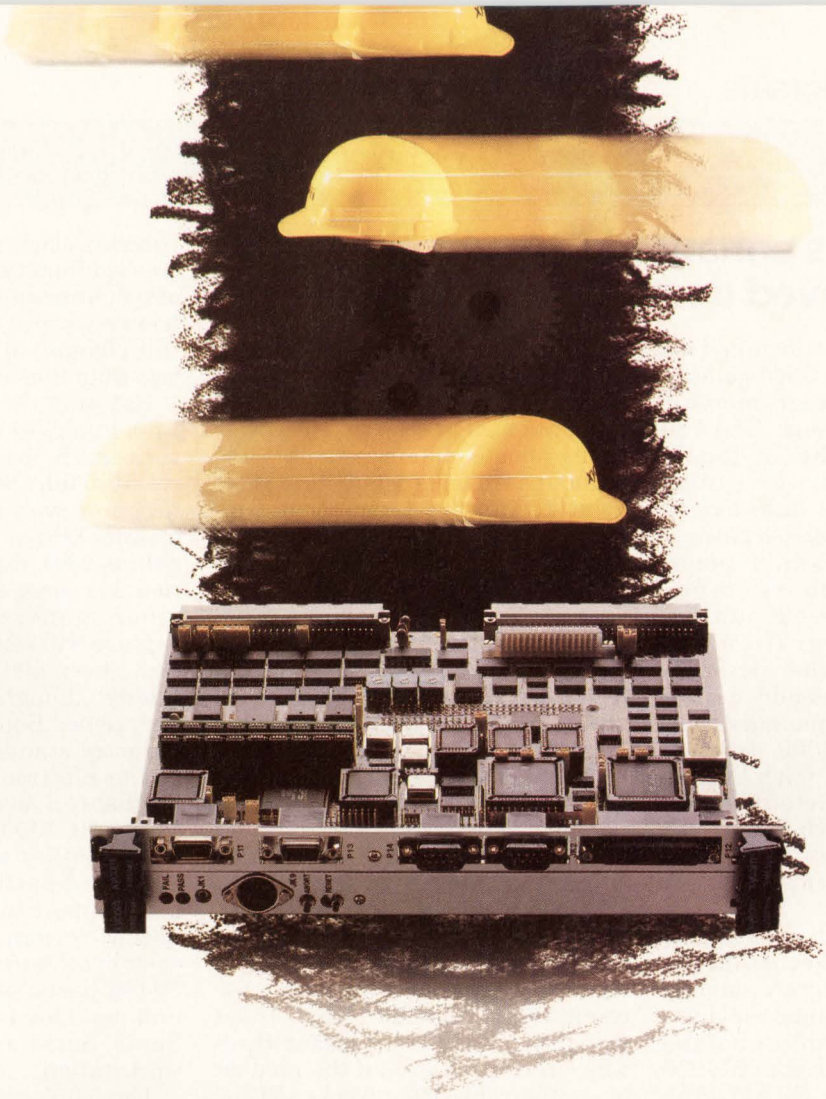

NOW AT A HARD PRICE TO BEAT

A HARD PC/AT TO KEEP UP WITH

The XVME-683—a versatile PC/AT processor in the industrial VMEbus format. It runs a true 25 MHz at a truly competitive price.

Now, product designers can take advantage of the high performance, multiprocessing capability of VMEbus and the software compatibility of the PC/AT architecture. Which means you'll be hard to keep up with when you're developing a new system.

The Xycom XVME-683 is a complete 80386 PC/AT processor on a two-card set. It features 25 MHz speed, 0 Wait State Cache Memory, and high resolution VGA color graphic capabilities. All in a VME format that's "hardhat hardened."

THE HARDHAT COMPUTERS

xycom

750 North Maple Road, Saline, Michigan 48176
(313) 429-4971 FAX: (313) 429-1010
In USA call: 1-800-289-9266 (1-800-ATXYCOM)
In Canada call: (416) 825-0281
In Europe call: +44 604 790767
In Asia call: 852 3-311 7855

CIRCLE NO. 68

Call us for a hard and fast quote on the Xycom XVME-683. We are shipping.

INTEGRATED CIRCUITS

Xilinx backs enhanced FPGA family with improved development software

The 5,000-gate XC4005 logic cell array is the first of the third-generation FPGAs (field-programmable gate arrays) from Xilinx. The full XC4000 submicron FPGA family will have 10 members, with system speeds up to 60 MHz and density from 2,000 to 20,000 usable gates.

The new FPGA family incorporates several on-chip system features for faster speed and increased logic integration. To accommodate the higher-performing devices, Xilinx has added hard macro support and a memory compiler to the XACT 4000 development system. And with XACT 4000, users can complete the place and route process at the push of a button with 95 percent utilization of all available logic gates, according to Xilinx.

Because logic cell arrays are based on an SRAM process, Xilinx has been able to integrate on-chip SRAM onto the XC4000 devices. Perhaps the most significant of the new system features of the XC4000 devices, the on-chip SRAM—the only such memory available on FPGAs—can be used for a range of applications, including FIFOs and register stacks. Users can configure chunks of memory from 2 to 28 kbits. With the Memgen memory compiler, engineers can create memories ranging from 1 to 32 bits wide and up to 256 words deep. The width and depth of the RAM can be specified with an entry language similar to a hardware description language. From a three-line design description, Memgen automatically creates a schematic symbol of the user's memory.

Hard macro library

Xilinx also provides a library of predefined, tested and fully characterized hard macros with the XC4000 FPGAs. These include fast counters, fast adders, RAMs, FIFOs, LIFOs, and register stacks. Each hard macro contains partitioning and routing information so that the XACT 4000 software can produce macro designs optimized for either density or performance. Since the development system already knows the relative placement and routing for a hard macro, the automated tools won't change the macro's functional performance.

To promote system testability and reduce board test costs, Xilinx offers Joint Test Action Group boundary-scan logic around the perimeter of the XC4000 architecture. The XC4000 family also dedicates resources for fast decoding of addresses. Up to 60 inputs plus complements can be decoded for on-chip logic in 10 ns. Moreover, decoded inputs can be fed back from off chip in a total of 15 ns.

Still another system feature that's been added to the XC4000 devices is the dedicated arithmetic logic for the fast generation and propagation of carry and borrow signals. The XC4000 fast-carry logic puts two counter bits into each logic block and operates at clock rates up to 50 MHz for 16 bits, regardless of whether the counters are loadable. For a 16-bit loadable counter, that means three times the speed in half the number of configurable logic blocks (CLBs) as in XC3000 devices. To help designers use the fast arithmetic logic, the XC4000 macro library has several hard macros that incorporate the carry logic.

No external buffers

By increasing the maximum output sink current of XC4000 devices to 12 mA, Xilinx has eliminated the need for external buffers, especially on bidirectional I/O lines. When two adjacent outputs are interconnected to increase the output current to 24 mA, the XC4000 devices can drive short buses on printed circuit boards or on plug-in printed circuit cards. For driving internal, bidirectional buses, XC4000 FPGAs have a pair of three-state buffers associated with each logic block in the array. These bus lines can also implement wide multiplexers or wired AND functions.

For connections between XC4000 blocks, Xilinx has increased the number of metal lines with programmable switching points and switching matrices that route resources. The number of globally distributed clock signals has been increased from two to eight. Designers of synchronous systems can now distribute several clocks as well as control signals all over the chip, with less than 2 ns of skew.

Not only do improved logic partitioning, placement and routing algorithms in the XACT 4000 software result in fully automated implementation of even the densest designs, claims Xilinx, but the new algorithms also make software execution five times faster than with previous generations. Schematic capture will remain the primary design entry method for the XC4000 family, though state-machine descriptions, Boolean equations and industry-standard interfaces such as the Electronic Data Interchange Format can also be used.

The XC4000 family will retain compatibility with the more than 100 third-party logic design tools that support the proprietary Xilinx netlist format. The development system software runs on 80386- and 80486-based personal computers and on Hewlett-Packard/Apollo, Sun-3, Sun-4 and DECstation 3100 workstations.

The 5,000-gate XC4005, with 112 user-definable I/O lines and a 14×14-gate CLB matrix, is available this quarter at $192.50 per sample. The 6,000-gate XC4006, with 128 user-definable I/O lines and a 16×16-gate CLB matrix, will be available next quarter. The 8,000-gate XC4008 and 10,000-gate XC4010 will be available by the third quarter.

Upgrades to the XACT 4000 software are available now at $1,750 for PCs and $3,300 for workstations. For new users, the XACT 4000, which includes XC2000, XC3000 and XC4000 support, is $7,990 for PCs with Viewlogic or Dash schematic editors. On workstations, it's $14,450 to $15,950, depending on the schematic editor and simulator interfaces. *—Barbara Tuck*

Xilinx
2100 Logic Dr
San Jose, CA 95124
(408) 559-7778

Circle 353

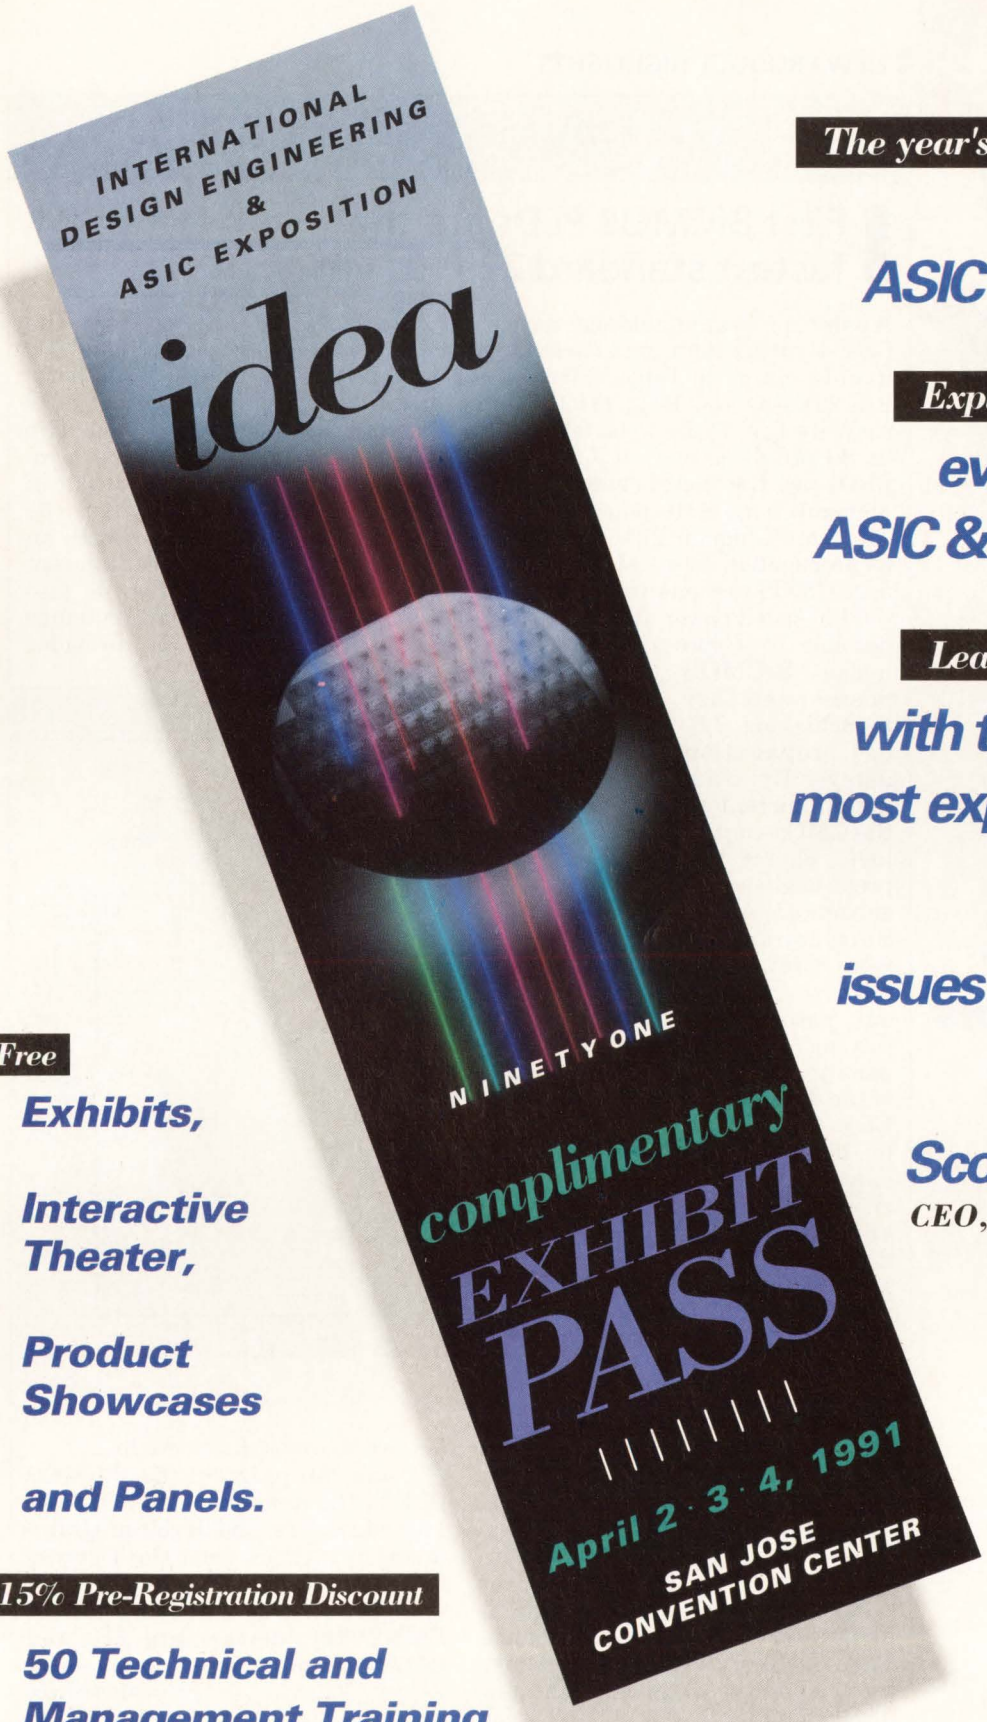

The year's most important event . . .

for ASIC & EDA users.

Explore and analyze . . .

every major ASIC & EDA company.

Learn and interact . . .

with the industry's most experienced users.

Today's . . .

issues and solutions.

Keynote . . .

Scott McNealy

CEO, Sun Microsystems

Free

Exhibits,

Interactive Theater,

Product Showcases

and Panels.

15% Pre-Registration Discount

50 Technical and Management Training Sessions and Seminars.

ASiC

Sponsored by ASiC Technology & News
480 San Antonio Road, Suite 245
Mountain View, CA 94040
1.415.949.2742

Register Now !

1·800·848·IDEA

1 · 8 0 0 · 8 4 8 · 4 3 3 2

CIRCLE NO. 69

NEW

High Performance Digital I/O

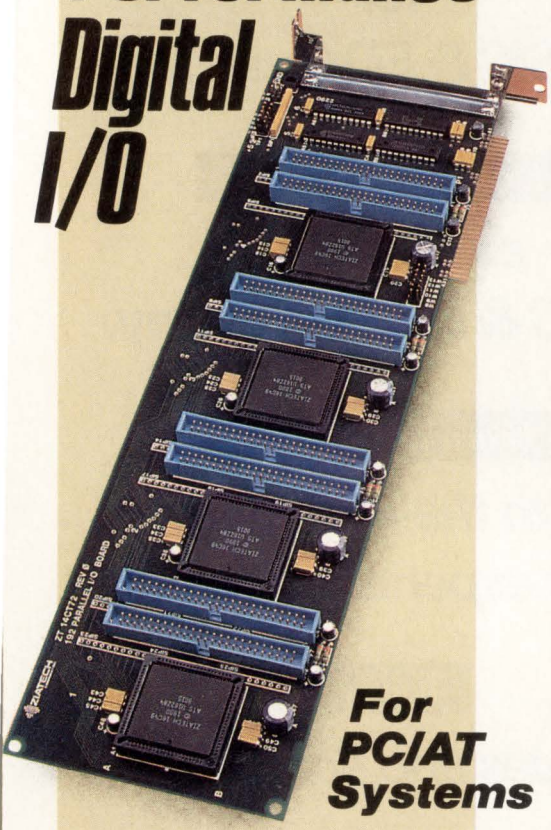

For PC/AT Systems

- **NEW** ASIC Provides Point By Point Input and Output Programmability
- High Current for OPTO 22 or Equivalent Isolation Modules
- Event Sense Interrupts— **8** (ZT 14CT73) **32** (ZT 14CT72)
- Device Drivers
- High Reliability —40°C to 85°C Temp Range
- 2 Year Warranty
- Sizes from 48 to 192 pts.

ZT 14CT73—48 PTS.
(half size card)
$240 List
$145 Qty. 100

ZT 14CT72—192 PTS.
(shown)
$540 List
$325 Qty. 100

ZIATECH CORPORATION

805-541-0488
FAX 805-541-5088

CIRCLE NO. 53

INTEGRATED CIRCUITS

First BiCMOS PLDs are the fastest standard 22V10s yet

A pair of programmable logic devices from Cypress Semiconductor has a double claim to fame. First, the PAL22V10C and PAL22VP10C are the very first PLDs to be fabricated in BiCMOS. Second, at 7.5 ns, the PLDs are the fastest standard implementations of the popular 22V10. The fastest bipolar 22V10s have 10-ns propagation delays, whereas standard CMOS versions are still at 15 ns.

The speed/power trade-offs made possible by Cypress' fuse-programmable BiCMOS process were key to achieving 7.5-ns propagation delays. By trading off current in the CMOS control logic, claims Cypress, engineers at subsidiary Aspen Semiconductor were able to maximize speed in the data path without paying a power penalty. The core of the data path is ECL, surrounded by bipolar level-conversion circuits. Since the supply current doesn't exceed 190 mA, the 22V10s can be housed in a standard 24-pin plastic package.

Another advantage of BiCMOS, claims Cypress, is that it lets users match output time constants to packages and loads, thus minimizing ground bounce. Users can program parts in low-inductance PLCCs, for instance, for faster edge rates than those of parts packaged in higher-inductance DIPs. Designers can also minimize ground bounce by opting for the new JEDEC standard 28-pin PLCC with isolated ground pins for outputs and an additional Vcc pin. With programmable edge rates and the new package pinout, ground bounce in the PLCC is held at 1.1 V. In the plastic DIP, ground bounce is 1.4 V.

The 22VP10 differs from the 22V10 by allowing bidirectional I/O on registered outputs. Both parts have 10 programmable output macrocells. On the PAL22V10C, two fuses (C_1 and C_0) can be programmed to configure the output in one of four ways. Each output can be registered or combinatorial with an active high or active low polarity. The feedback to the array is also from this output. An additional fuse (C_2) in the PAL22VP10C provides for two more feedback paths.

Users can program the BiCMOS PLDs with standard Data I/O, Stag, Logical Devices, and Sprint programmers, as well as with the Cypress QuickPro programmers. In 100-unit quantities and in PLCCs, PAL22V10C devices are $30 and PAL22VP10C devices are $39.45.

—Barbara Tuck

JEDEC 28-PIN PLCC

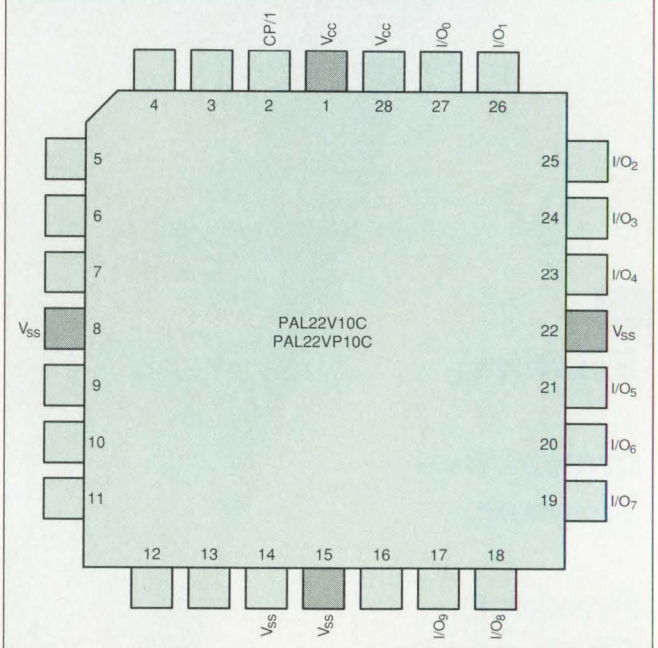

Cypress Semiconductor
3901 N First St
San Jose, CA 95134
(408) 943-2600

Circle 356

Ada development environment is all ADA

An integrated Ada development environment for real-time systems features a real-time executive that's also completely written in Ada. TeleSoft's Triad (Telegen2 real-time integrated Ada development) environment consists of three major component groups. It uses TeleSoft's Ada compiler technology targeted at the Motorola 68000 family of microprocessors; it contains a full suite of Ada productivity tools; and it has an Ada-based real-time execution environment including a new executive called TeleAda-Exec.

The entire package and all components are written entirely in Ada to ensure a seamless development environment, so there's no mixing of technologies or development paths. The tight integration of all components helps assure complete customer support from one organization, according to TeleSoft.

Three versions of Triad are available, and all use the Telegen2 optimizing compiler technology. There's a Sun-3 to embedded 68000 version, a Sun-4 (Sparc) to 68000 version, and a VAX to 68000 version. Both host and cross-compilation systems are included. The Sun-based versions feature window- and mouse-based user interfaces, SunOS and SunView Ada bindings and X Window bindings. The VAX/VMS version conforms to the digital command language interface standard and supports a VMS interface that lets Ada programs call any subprogram conforming to the VAX procedure calling standard.

Programming tools include a source-level debugger that uses the Hewlett-Packard/Microtec Research extensions to the IEEE-695 object file format, which is generated by the compiler; a global optimizer; a cross-referencer; a library manager; and a library toolset. In addition, there's an Ada source dependency lister, a compilation order tool, and a dynamic analyzer. With this analyzer, the developer can look at the percent of time spent in different modules in a program, and then decide where optimization would achieve the best return.

Triad supports Ethernet for both downloading and debugging activities during development and for implementing loosely coupled multiprocessor-based applications. The TeleAda-LAN facility is a set of reusable components, callable in Ada.

TRIAD REAL-TIME ADA DEVELOPMENT ENVIRONMENT

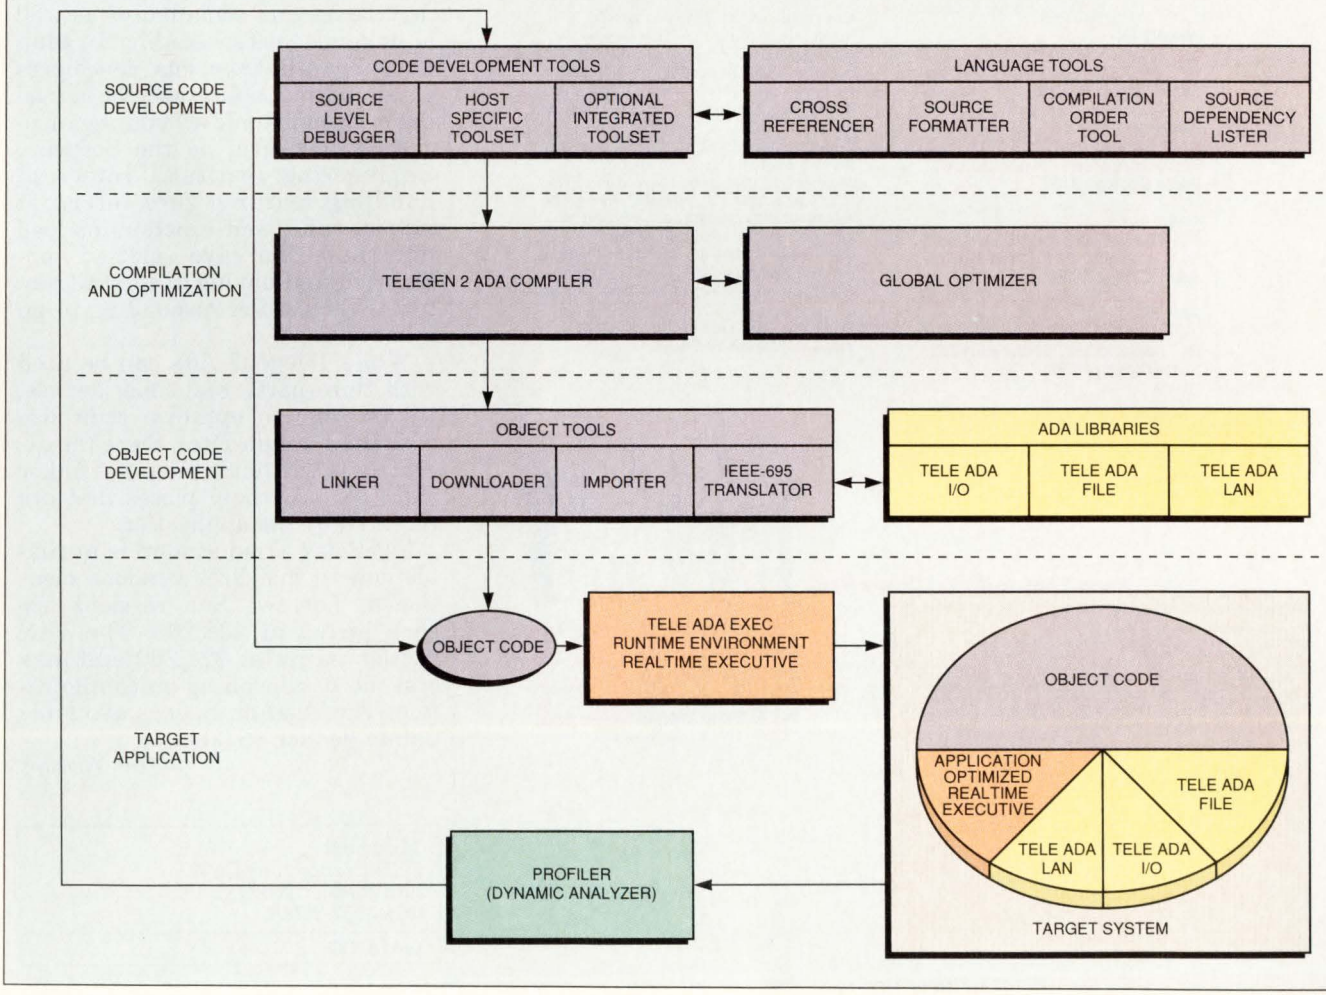

SOFTWARE

Its two most important features are task-to-task communication (TTC) and remote procedure calls (RPCs). TTC allows individual tasks within separately compiled Ada programs to communicate directly with each other or over the network. RPCs allow Ada subprograms to be called from programs running on remote computers on a network. In addition, TeleAda-LAN supports Sun-compatible network file system, file transfer protocol and simple mail transfer protocol interfaces.

Though Ethernet downloads programs to target systems many times faster than serial lines do, Triad has also implemented a method of minimizing large downloads by using phantom subsystems. A phantom subsystem is a separately linked subsystem of an application image that's referenced from other subsystems but not included in them. Thus stable code can be distinguished from modules under active development and only those modules need to be linked and downloaded when a change is made.

Another time-saver is the Recover feature of the Telegen2 downloader used to recover from test runs. Recover checks the downloaded image in the target system against what's already in the system, and only downloads code that has changed.

Real-time features

The TeleAda-Exec run-time system supports traditional real-time features such as events, interrupts, mail boxes and semaphores as well as dynamic task scheduling. In addition, TeleAda-Exec lets developers obtain hard deadline scheduling using rate monotonic scheduling algorithms developed by the Software Engineering Institute. Interrupt handling features zero interrupt lockout time, and function-mapped interrupts can save selected contexts without having to do a full context switch when responding to an interrupt.

While Telegen2 Ada can be used with third-party real-time kernels, it's possible to optimize code size with the TeleAda-Exec. Since the executive is written in Ada, the linker will link only those pieces that are required by the application.

TeleSoft's Triad system is available now in the three versions mentioned. The two Sun versions are each priced at $28,200. The VAX version starts at $28,200 and may cost more, depending on configuration. The TeleAda-Exec is available unbundled for $5,000.

—Tom Williams

TeleSoft
5959 Cornerstone Ct W
San Diego, CA 92121
(619) 457-2700

Circle 357

Concurrent Engineering is Here!

Don't Miss Any Part Of This Essential Time-To-Market Series

January 14:
DESIGN'S ROLE IN CONCURRENT ENGINEERING
Space Closing 12/31/90

February 11:
MANUFACTURING'S ROLE IN CONCURRENT ENGINEERING
Space Closing 1/28/91

March 18:
PURCHASING'S ROLE IN CONCURRENT ENGINEERING
Space Closing 3/4/91

April 15:
TESTING'S ROLE IN CONCURRENT ENGINEERING
Space Closing 4/1/91

It's the key to winning in the tough, competitive markets every company will face in the '90s. Concurrent engineering promises shorter development cycles, lower costs and higher quality products.

But getting it to work is no simple matter. It requires new thinking by management and a re-evaluation of everyone's role in product development. No one is unaffected — not Design, not Manufacturing, not Test, not Purchasing.

There are technical, organizational and personnel obstacles to concurrent engineering that must be overcome. This four-part series will describe some of the solutions that industry experts and OEMs at the forefront of product development have devised. It will become essential reading for every manager at every computer/electronics OEM.

Don't miss any part of this essential Time-To-Market Series.

Where Time-To-Market Comes First

PENNWELL PUBLISHING COMPANY • ONE TECHNOLOGY PARK DRIVE • WESTFORD, MA 01886 • 508-692-0700

For more information, contact any of us ...

PUBLISHER: David L. Allen, 508-392-2111 • ASSOCIATE PUBLISHER: Tim L. Tobeck, 516-883-8815 • SALES OFFICES: PORT WASHINGTON, NY, Eleanor Hobbs, 516-883-8815 • WESTFORD, MA, Tim Pritchard, 508-392-2217 • MATTESON, IL, George Mannion, 708-748-7565 • SPRING, TX, Eric Jeter, 713-353-0309 • SANTA ANA, CA, Tom Boris, Greg Cruse, Phil Cook, 714-756-0681 • PALO ALTO, CA, John Sly, Bill Cooper, 415-965-4334 • SURREY, ENGLAND, David Round, 81-686-7655 • PARIS, FRANCE, Daniel R. Bernard, 1-43-54-55-35 • MUNICH, GERMANY, Johann Bylek, 089-903-88-06 • MILAN, ITALY, Luigi Rancati, 02-7531445 • TOKYO, JAPAN, Sumio Oka, 03-502-0656 • SINGAPORE, REPUBLIC OF SINGAPORE, Anne Goh-Taylor, 734-9790 • SEOUL, KOREA, Jay G. Seo, 776-2096

COMPUTER DESIGN Time To Market News

A PennWell Publication

CIRCLE NO. 76

Analog layout editor links schematics to physical design

Increasing amounts of analog and mixed-signal components on today's advanced ICs pose some tough challenges for circuit designers. Even if the amount of silicon devoted to the analog portion of the chip is small, the time it takes to design that portion dominates most of the design task. Not only are the basic layout elements, such as resistors, capacitors and transistors, more complex in analog design, there are multiple design constraints, like thermal symmetry and complex parasitics, that must be addressed. This complexity makes it virtually impossible to make changes on the schematic and physical layout separately.

To help designers with these analog problems, Cadence Design Systems has unveiled the Analog Artist layout editor, an IC layout tool for advanced analog and mixed-signal ICs. The editor lets analog designers create analog cells and blocks through tight linkage between schematic design and physical layout.

Analog Artist is based on a 32-bit precision database that supports all-angle hierarchical design. The tool has unlimited library referencing, with an unrestricted editor selection list that works on multiple individually configured windows to allow editing of several designs concurrently. The editor's commands handle arcs, ellipses and Boolean operations, so that users can cut polygon shapes, merge polygons and convert paths to polygons.

To minimize the effort of design entry and editing, the Analog Artist editor lets designers specify parameterized cells (Pcells). With a transistor as a Pcell, for example, users can alter its length and width and fold it into multiple segments or other configurations by changing its parameters. For analog bipolar processes, Pcells can include arcs, circles, paths and texts. One of the key features of the tool lets designers take any drawn data, draw as many stretched parameters as needed, and compile them from a menu button. Several Pcell generators are supplied with the tool, including multigate transistors, bipolar NPN and PNP transistors, resistors, and capacitors.

Interactive layout

The tool includes a set of interactive layout functions for full-custom design of analog circuits. These functions permit layout at a higher level of abstraction (similar to schematics) than the polygon level. Instead of working with individual geometric shapes, designers work with device instances—transistors, resistors and capacitors—as well as with wires and contacts.

The editor automatically creates and maintains electrical connectivity information associated with the layout data. Imaginary thermal lines in the design can be created to enforce thermal symmetry of particular pairs of devices. These device-level editing functions can be used for interactive placement, routing, editing and advanced layout checking of a design. Because most of these functions work on both schematic and layout representations, there's no need to learn separate sets of commands.

Using placement functions, designers can interactively select a single device in the schematic and place a corresponding device in the layout. The editor can place all layout devices, using their relative placement in the corresponding schematic for initial placement.

With the tool, Pcells can be grouped in the layout and associated to a single device in the schematic. For example, a single resistor instance in the schematic might be drawn as several polygons in the layout, or a transistor might be drawn as multiple parallel transistors. Users can highlight mismatches between the schematic and layout, such as missing devices or Pcell parameter mismatches, as well as place a layout device with respect to a thermal line. Connectivity information is automatically transferred from the schematic to the layout.

Routing functions permit multilayer wires in the layout to be easily created. Contacts are created automatically when the designer changes layers, and connectivity is created and maintained automatically while a wire is routed. As a result, electrical violations are detected and reported immediately.

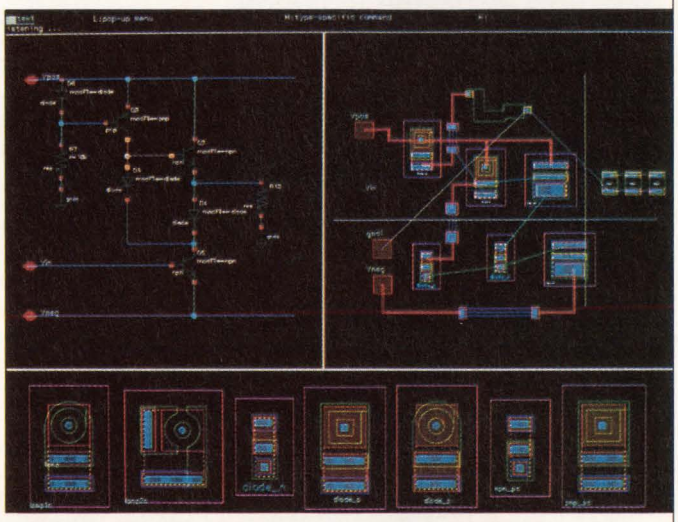

The Analog Artist layout editor is available as a stand-alone system or as an upgrade to existing Cadence layout editor installations. The tool supports most industry-standard Unix platforms and is available now, with single-quantity pricing starting at $25,000 per seat.

—Mike Donlin

Cadence Design Systems
555 River Oaks Pkwy
San Jose, CA 95134
(408) 943-1234

Circle 354

Coming March 1

Watch for the Special Report on fuzzy logic in embedded control by Tom Williams.

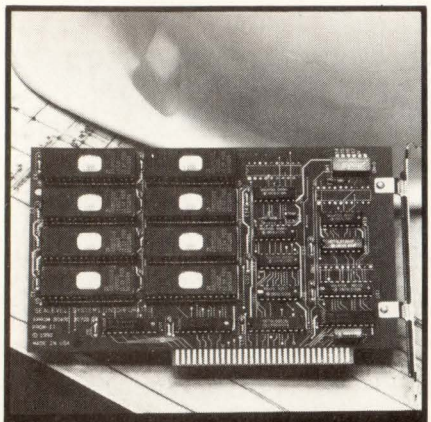

PROM-III

- PUT DOS AND APPLICATION IN EPROM
- ALLOWS DISKLESS OPERATION
- UP TO 1 MBYTE ROM-DRIVE WITH 16K FOOTPRINT
- PROMKIT SOFTWARE BY ANNA-BOOKS
- FLASH EPROM SUPPORTED
- BATTERY RAM MODULES SUPPORTED
- DELIVERY FROM STOCK

SEALEVEL COMMUNICATIONS & I/O

SEALEVEL SYSTEMS INC.
PO BOX 1808
EASLEY, SC 29641
[803] 855-1581

CIRCLE NO. 56

reprints

Reprints of any article or advertisement appearing in *Computer Design* may be ordered from: June Bosarth, 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.

DESIGN AND DEVELOPMENT TOOLS

Design tool solves timing problems

A software package called TimingDesigner from Chronology enables engineers to solve critical timing problems during creation and analysis of timing diagrams. Unlike a simulator, which can only be used to check circuit timing after a design is complete, the TimingDesigner is a front-end tool that can be used to specify and create a design, as well as to verify it. The tool lets users enter and modify all elements of a timing diagram, including waveforms, timing constraints, text annotation, and timing parameter tables.

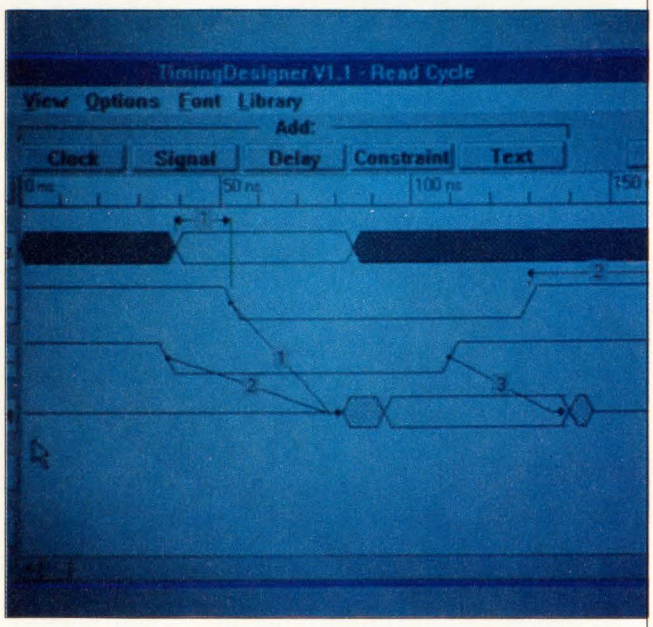

To draw timing diagrams, the operator uses a mouse to place edges on the diagram model and to select between logic signal states or valid/invalid bus states. Clock signals are automatically generated after the operator enters the values for duty cycle, duration, phase, and jitter.

Delays can be selected from a separate library window, with minimum and maximum delays for each part. Delay values can be entered directly or computed from formulas based on user-defined variables such as temperature and voltage.

Timing constraints are entered by the user's clicking the mouse on the two edges that must be constrained, then entering the values or formulas for minimum and maximum times. Text can be set in several sizes and placed anywhere on the diagram, or it can be attached to an edge so that it moves when the edge moves.

Relationships maintained

As a timing diagram is modified, TimingDesigner automatically maintains the relationships specified between waveform edges and delays. When a user moves the location of an edge that's dependent on other edges, all of the related edges in the diagram automatically move, so that the minimum or maximum delays specified are maintained.

The tool also automatically performs timing analysis calculations and displays them as the engineer draws the timing diagram. The software then displays the earliest and latest time that every edge can occur and computes timing margins for all timing constraints. Any timing constraints that have been violated are highlighted in red, and valid conditions are displayed in green.

TimingDesigner also features a database that can store and retrieve timing data associated with a design. This lets engineers save timing data in the same way that they now save connectivity data from schematic editors and layout data with circuit board editors.

Available now, the TimingDesigner runs under Microsoft Windows Version 3.0 on IBM PCs and compatibles with 640 kbytes of memory, an EGA monitor, a hard disk and a mouse. The software sells for $1,495. —*Mike Donlin*

Chronology
2849 152nd Ave NE
Redmond, WA 98052-5516
(206) 869-4227

Circle 355

Mitsubishi Fast SRAMs

When You Need Fast Cache.

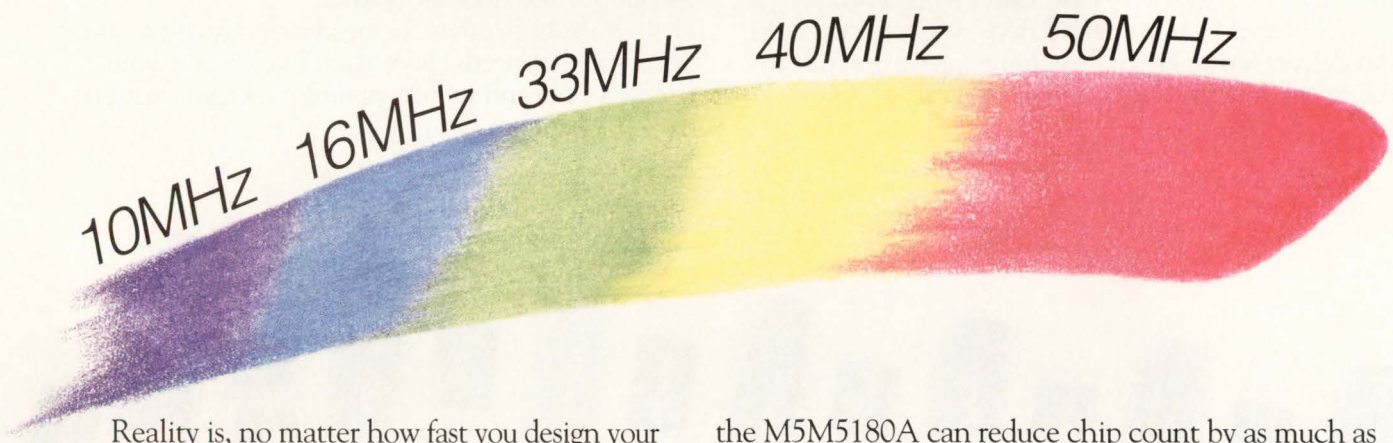

Reality is, no matter how fast you design your systems, customers always want them faster. To stay ahead, you need fast access to ultra-fast SRAM access times. And, Mitsubishi helps you lead the cache memory race with fast, 15ns Static RAMs in 28-pin SOJ, DIP and flat packages.

Mitsubishi's fast SRAMs let you cash in on your cache subsystem performance. Just match memory density and organization with any of the popular industry-standard microprocessors. You get the speed you need, and use fewer components. In fact, with its on-chip address latch, the M5M5180A can reduce chip count by as much as 50% in an 8K x 32 cache memory subsystem.

Mitsubishi's cutting-edge, 0.8 micron CMOS technology and mask/bonding wire option let you interchange/upgrade SRAMs without requalifying each device. So, your design and redesign processes are faster than ever.

Part No.	Organization	Access Time ns	Package (all 28-pin)
M5M5178A	8K x 8	15, 20, 25	SOJ, Flat Pack, DIP
M5M5179A	8K x 9	15, 20, 25	SOJ, Flat Pack, DIP
M5M5180A	8K x 8 (Latched)	20, 25	SOJ, Flat Pack, DIP

Set your sights on Mitsubishi's fast Static RAMs for cache subsystems and pick the speed that will get you to market quick. And, to get to market in strength, turn the page. You'll quickly discover our volume production capabilities.

MITSUBISHI
ELECTRONIC DEVICE GROUP

Mitsubishi Static RAMs.

When you need a lot.

Your design is almost finished. Or, maybe you're just starting. Perhaps you're ready to place production orders. In any event, if your design requires SRAMs you'll need a lot of them. Fast. Enter Mitsubishi.

Whether fast or slow SRAMs, you can count on us to deliver critically high-volume quantities, just in time to get your production lines moving.

From fast, 15ns 64K SRAMs and 35ns 1 megs, to standard 128K x 8 devices, we give you high performance in a steady flow, straight from dedicated Mitsubishi SRAM factory lines.

Mitsubishi offers the production quantities you need and the speeds you want. Plus, we give you through-hole and surface-mount packaging options.

Part No.	Organization	Access Time ns	Package
M5M5165	8K x 8	70, 100, 120	DIP, Flat Pack
M5M5255B/ M5M5256B	32K x 8	70, 85, 100, 120	DIP*, Shrink DIP, ** Flat Pack, TSOP
M5M51008	128K x 8		DIP, Flat Pack, TSOP***
M5M5178	8K x 8	15, 20, 25,	DIP, SOJ,
M5M5179	8K x 9	35, 45, 55	Flat Pack ****
M5M5180	8K x 8 (Latched)	20, 25	DIP, SOJ, Flat Pack
M5M5187	64K x 1	15, 20, 25, 35, 45, 55	DIP,SOJ
M5M5188	16K x 4		
M5M5189	16K x 4 (/OE)		
M5M5257	256K x 1		
M5M5258	64K x 4		
M5M51001	1M x 1 or 256K x 4	35, 45	
M5M51004	256K x 4 (/OE)		

* 600 mil ** 300 mil *** TSOP: Thin-small-outline-package; also available in reverse pin-out.
**** Flat package available in 25ns speeds and faster.
Some new products subject to availability.

When you need less.

When you need to pack more memory in less space and reduce power, Mitsubishi offers standard and soon, fast SRAMs in TSOP (thin-small-outline packaging). Static RAM packaging can't get any smaller than the TSOP, since it's virtually the size of the die.

Our TSOP SRAMs take up only half the surface area and are four times smaller in volume than equivalent pin-count, standard surface-mount packages. And, with a profile half that of standard surface-mount packages (1.2mm) and a low standby current (20μA to 50μA), TSOPs are perfect for portable, handheld, battery-backed applications.

So, if your design requires a lot of static RAM or just a little, call Mitsubishi for some action. We'll bring your vision to reality. Mitsubishi Electronics America, Inc., Electronic Device Group, 1050 E. Arques Ave., Sunnyvale, CA 94086. 1-800-624-8999 ext. 178C.

VISION SM
ACTION
REALITY

CIRCLE NO. 73

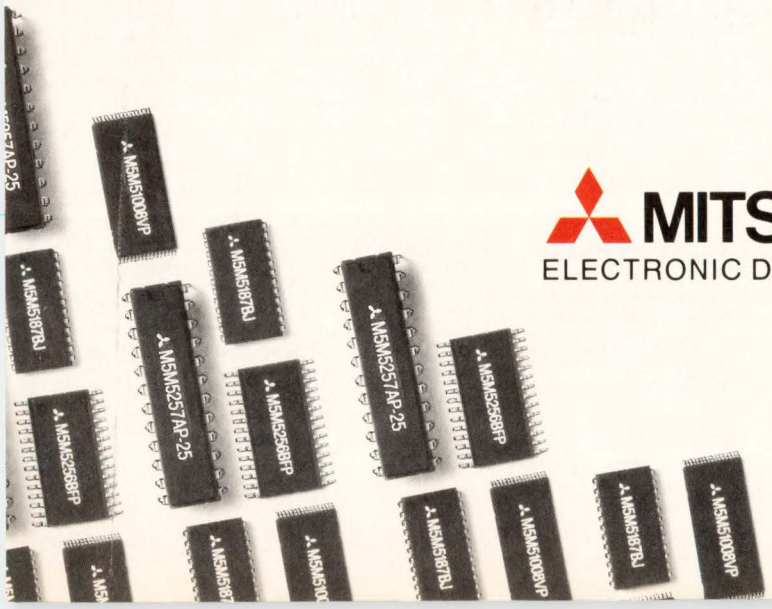

MITSUBISHI
ELECTRONIC DEVICE GROUP

EMPLOYMENT OPPORTUNITIES, OUTPLACEMENT AND POSITIONS WANTED

Recruitment Advertising Managers

Eastern U.S.
Sue Nawoichik (508) 392-2194

Western U.S.
Tom Murphy (213) 372-2744

CAREER CONNECTIONS

Closing Dates for upcoming magazine issues

Issue	Closing dates
March 1	February 8
April 1	March 8
May 1	April 5

POSITIONS WANTED

POSITIONS WANTED ADS

Free 1" ad to subscribers seeking full-time employment. Just include 50 words of copy and your subscription label. We'll run your ad in 2 consecutive issues.

Available to non-subscribers or consultants/companies at $125 per column inch. Mail your position wanted ad to:

COMPUTER DESIGN Technology and Design Directions®

Positions Wanted
One Technology Park Dr
PO Box 990
Westford, MA 01886

SALES/MARKETING MANAGEMENT. Strong technical background. Experience in Engineering, Manufacturing, Sales, Marketing. Semiconductors, Multichips, and Subsystems. Commercial and Industrial. 12 years experience. Desire progressive company and product line. West or Northeast location. Call (408) 377-3889.

SIX SIGMA/TQM QUALITY MANAGER. Experience in manufacturing, engineering, customer service with specialties in Training, Product Safety, Program management, Project management, Quality Engineering. Six Sigma certified instructor, experience in QFD facilitation. Extensive experience with International and national testing house and standards organizations such as UL, CSA, IEC etc... R. Minezzi 134 Center Rd. Shirley, MA 01464 (508) 425-4953.

PROGRAM MANAGER. Are you surrounded by Program Managers who are thorough, dedicated professionals with exceptional planning, analytical, and interpersonal skills? Who get things done *on time and within cost and meet your customer's expectations?* Congratulations if you are, and if you can use one more, your search has ended. Joe (516) 799-1229.

"HARDWARE/SOFTWARE DESIGN ENGINEER, 3 years experience with MSEE. Digital, Analog & Microprocessor based product design and embedded software. Proficient in Assembly, C and Basic. Schematic and PCB layout. Manoj Mehta, 83-35 116th Street, #2E, Richmond Hill NY 11418. (718) 805-3442."

SENIOR HARDWARE/SOFTWARE ENGINEER, BSEE +25 years experience in Peripheral controller, processor, midi computer, operating systems: specification design, implementation and integration. Recent experience with: C, VRTX, ECL, ASICs in high performance logic analyzer. Arizona preferred. Roger Newey, 2005 N 48th St. #4 Phoenix, Az 85008. (602) 244-2448.

DESIGN ENGINEER. 2 years design/prototyping experience of microprocessor-based systems (8051, 8085, z-80). Analog/Digital experience for chemistry research projects. PCB layout and manufacturing. Excellent troubleshooter and communicator. Self-motivated and needs a constant challenge. Call Troy at (316) 364-8831 ext. 8762 days.

ENTREPRENEUR. BSAE, MSEE, 20yrs Sales, Proj mgmnt, Test, Qual, RF, Telcomm, Consult. Bell Labs, Rockwell, Honeywell, clients. Seeks opportunity to feed six kids. prefers nice place to work and live. Also sings and writes poetry. Great new ideas. No risk call. Marty Grogan, P.O. Box 453, Forest Lake, MN 55025 612/464-7863.

Software programming experience in C, C++, ADABASE, MODULA 2, AI and EXPERT SYSTEMS, DBASEIV, and CLIPPER. Seeking independent programming projects. WERNER BUCHHOLZ, P.O. BOX 42, SARTELL, MINNESOTA 56377, 612-252-0767.

SALES/SALES MANAGEMENT - Extensive experience in successful selling and market penetration to high end graphics users. Market experience includes: Flight Simulation, Combat Mission Rehearsal, Image Processing/Analysis, Seismic Imaging, Animation, and Medical Imaging. Possess strong communication and presentation skills. Contact: Rick Tears, 4505 Fairway Ave., Dallas, TX 75219; (214) 559-2593.

SR. HARDWARE ENGINEER: 30 Years experience in Analog signal processing, Power supplies, Component Engineering, Worst-case analysis and Reliability derating, Radiation effects Consumer Audio Products & Loudspeaker Systems. Will relocate anywhere W. of Rockies. L. Garner, 5530 W. 190th St., #137, Torrance, CA 90503.

PARALLEL PROCESSING SPECIALIST. Let me help you deliver the promise. Senior Software Engineer with over 15 years experience in vector and parallel processing. Conversions and optimizations a specialty in Unix using Fortran and C. Vince Wayland, P.O. Box 3254, Boulder, CO 80307, (303) 499-9027.

SENIOR SOFTWARE ENGINEER 25 years experience with microprocessors, mainframes and science; 6 years software management. Extensive C, Fortran, Assembler, Object Pascal, PCDOS, Macintosh toolbox, MacApp, MVS. Know many more. Instrument development: embedded and computer interface. PhD Physics. Willing to relocate. David Tanner, 131 Altamont Ave., Tarrytown, NY 10591. (914) 332-0450.

SALES MANAGEMENT/TECHNICAL MARKETING. Real Time Single Board Computers, Data Acquisition, High technology sensors, signal conditioning and support magnetic & optical memory subsystems. Regional Sales Manager for over 20 years. Supported direct + Mfg rep salesmen. Engineering & MBA degrees. Industrial & Aerospace experience. NE region or will relocate. Contact Bernie Segal (607) 729-0118.

PROGRAMMER/ANALYST. Bachelor of Science, Computer Information Systems with minor in business. Achieved initiation of an internship program with an insurance company, designed a program to help solve agent problems and possess excellent written and verbal skills. Will relocate. Resume on request. Christopher Pasquali, 20 Fredonia Rd., Newton, NJ 07860. (201) 579-3647.

Seasoned executive from High Tech MFG, Operations and R & D. Start up, Turn around. QA skills as Director, Manager or Mfg. Mgr. Plant Mgr., Entrepreneur skills, P. E. licence. Contact Paul Savko, 5869 Rowanberry Drive, Elkridge, MD 21227. (301) 796-1049. Open Relocation, Salary 75k+ Range.

SENIOR QA/MECHANICAL ENGINEER seeks position which would benefit from over 19 years experience. Excellent experience in Configurations management, QA, Spec. writing and vendor and other departmental interfacing. Hardware, C/M procedures, ECP's, DEv. Waiver, As built lists and Retrofit requirements. Wish to stay in LA and surrounding area. Please call and leave message for resume...(818) 365-2200

SYSTEMS ENGINEER. 25 years experience in Signal Processing, Analog to DSP. RADAR, SONAR, SIGINT applications for parallel computer systems, vector processing. Hardware design integration VLSI chips to racks. Design analysis, proposal organization customer support. Permanent or temporary. Can travel. MSEE+ registered PE. For resume, references and salary contact (908) 572-4161 or FAX (908) 937-8190.

HARDWARE DESIGN ENGINEER. MS Computer Engineering, 10 years design experience in fields of computer controlled instrumentation, Robotics, Data Acquisition and motor control. Four Patents in CNC. Hands-on experience in digital (discrete/microprocessor) circuit design. Have knowledge of Assembly, C, dBaseIII+, MS-DOS. NY, NJ, CT areas. For resume, contact (718) 998-2155.

SALES ENGINEER or SALES MANAGER. Demonstrated success with high-tech components, equipment, systems, datacom/telecom, networks, services. Distribution and direct sales experience in 15 Southeastern states, calling on management, engineering, specifiers, purchasing. BS-degree. Strong "hands-on" technical training/experience. Peter Alexander, P.O. Box 957111, Duluth, GA 30136. or call (404) 945-6950 (Atlanta).

CONTRACT SOFTWARE/HARDWARE ENGINEER. BSEE Realtime systems specialist multitasking exp UNIX vxWorks OS/2 & DOS Desire contact S/W assignmnt CONTACT: Richard D. Dumas (214) 783-8455.

LET THE GOVERNMENT FINANCE your new or existing small business. Grants/loans to $500,000. Free recorded message: 707-448-0330. (PN4)

Industrial and manufacturing Engineer with fifteen years of experience in designing, manufacturing and marketing Audio equipment including experience with various CAD systems, computer analysis, test equipment, injection and vacuum molding. Truly understand the difficulties of trouble shooting and getting products on the market. Danny O. Quinn, 6322-A Seven Springs Blvd, Lake Worth, FL 33463.

Why the best CAE professionals call Compaq home.

At Compaq, our common-sense approach to business has helped make us a world leader in advanced personal computers. But its our commitment to people, our relentless dedication to building a quality work environment that's the heart of our success.

Mechanical CAE/CAD Engineers

You'll evaluate, develop, integrate and support the latest technological tools used to develop our market-leading computer products. You'll develop new tools and design methodologies, while providing direct user support, enhancing the tools in use and working closely with suppliers to adapt their products to meet user needs.

A BSME and at least two years' experience is required for the following positions. You should also have a good knowledge of CAE/CAD and be familiar with DOS and UNIX operating systems.

Engineer, Analytical Tools

Finite element; mechanisms, kinematics, and thermal analysis; structural correlation testing; training.

Analytical Engineer

Finite element; mechanism, kinematics, and thermal analysis; structural correlation testing; training.

CAD Tools Productivity Engineer

2-D and 3-D modeling; solids modeling; macro programming; expert systems applications; tool integration.

Engineer, New Tools

2-D and 3-D modeling; solids modeling; bench marking; specification writing; justification analysis.

Engineer, Manufacturing/Supplier Liaison

Information systems; computer-aided process planning; N/C programming; networking; 2-D and 3-D modeling; design for manufacturability and assembly.

Compaq offers competitive salaries, comprehensive benefits and an unequaled work environment. If you are interested in one of these unique opportunities, simply submit your resume, along with the positions for which you wish to be considered, to: *Compaq Computer Corporation, Dept. CAE110-JS, P.O. Box 692000, Houston, Texas 77269-2000.*

Compaq is an affirmative action employer, m/f/h/v.

COMPAQ

CIRCLE NO. 57

SYSTEM SHOWCASE

STD BUS PRODUCTS

8 & 16 Bit CPU's
Peripherals & Card Cages
Custom Design
Worldwide Service
OEM Discounts
Guaranteed Delivery

ZWICK SYSTEMS INC.

17 Fitzgerald Rd., Suite 104
Nepean, Ontario, Canada, K2H 9G1
Tel (613) 726-1377, Fax (613) 726-1902
Representatives Required!

CIRCLE NO. 176

Industrial 386sx

$3200.00

- 19" Rack Mountable
- Passive Backplane
- VGA Graphics
- 1.44Mb Floppy
- 2 Serial/1 Parallel
- Intel 80386sx-16MHz
- 1Mb RAM
- 14 Expansion Slots
- 44Mb HD
- 270w Power Supply
- One-Year Warranty

Custom Cases/Graphics Other Systems Available
Distributors Welcome
Call For Complete Information

Industrial Computer Solutions
200 Black Oak Ridge Rd., Wayne, NJ 07470
Phone (201) 696-1133 Fax (201) 696-0870

CIRCLE NO. 177

RGB CABLES

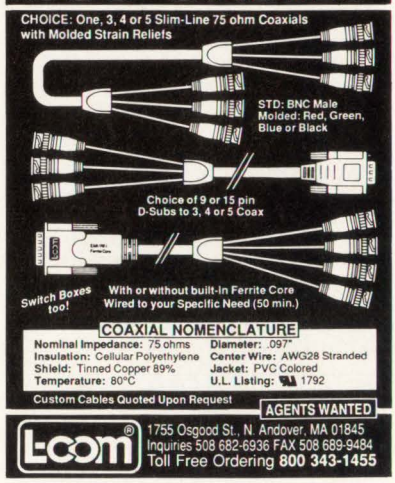

COAXIAL NOMENCLATURE

Nominal Impedance: 75 ohms
Insulation: Cellular Polyethylene
Shield: Tinned Copper 89%
Temperature: 80°C
Diameter: .097"
Center Wire: AWG28 Stranded
Jacket: PVC Colored
U.L. Listing: 1792

Custom Cables Quoted Upon Request

AGENTS WANTED

L-com®
1755 Osgood St., N. Andover, MA 01845
Inquiries 508 682-6936 FAX 508 689-9484
Toll Free Ordering **800 343-1455**

CIRCLE NO. 178

HIGH SPEED FIBER OPTIC LINK

for PC bus
for AT bus
for VME bus
for MULTIBUS I
for MULTIBUS II

New Procise Fiber Optic Transceivers feature:

- AT model for PCs and ATs
- SBX model for VME, Multibus I & II
- 11 Mbyte/sec full-duplex synchronous link
- Industry standard cable & ST connectors
- MS-DOS system drivers included
- $900 per transceiver (Qty 1-9)

For further information contact:

procise CORPORATION
PO Box 1011, Issaquah, WA 98027
(206) 392-0270

CIRCLE NO. 179

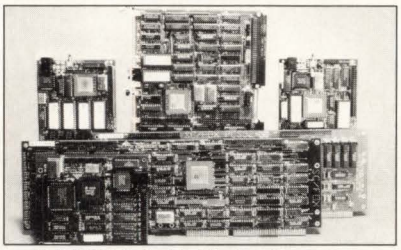

SC/FOX™ Embedded-System Computers

SC/FOX VME SBC (Single Board Computer) 18 MIPS, 60 MIPS burst, general-purpose, slot-1 Master/Slave System Controller. Up to 512K bytes 0-ws memory, 1 SCSI, 1 ptr, 2 serial ports. Uses 16-bit Harris RTX 2000.

SC/FOX PCS (Parallel Coprocessor System) 15 MIPS, 50 MIPS burst, general purpose PC/AT/386 plug-in board, 32K to 1M byte 0-ws static memory, multiple PCS operation, SCSI option. Uses Harris RTX 2000.

SC/FOX SBC (Single Board Computer) 18 MIPS, 60 MIPS burst, Stand-alone operation, Eurocard size, 1 ptr, 1 serial port, up to 512K bytes 0-ws memory, 2 50-pin user connectors, SCSI option. Uses Harris RTX 2000.

SC/FOX PCS32 (Parallel Coprocessor System32) 15 MIPS to 70 MIPS, general-purpose PC/AT/386 32-bit plug-in board with 64K to 1M byte 0-ws static memory. Uses 32-bit SC32 Forth microprocessor.

Ideal for embedded real-time control, data acquisition, or high-speed processing. Forth software and utilities included, C optional. OEM pricing.

SILICON COMPOSERS INC *(415) 322-8763*
208 California Avenue, Palo Alto, CA 94306

CIRCLE NO. 180

Precision Time Plug-ins

by TRUETIME

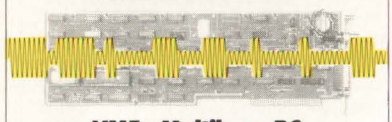

VME · Multibus · PC

TIME CODE
Time in various formats is furnished to the host computer over its particular bus. Time resolution to 1 microsecond is available.

SYNCHRONIZED CLOCKS
Available in a number of sizes and configurations, these boards provide 1 millisecond time synchronization referenced to UTC.

3243 Santa Rosa Avenue • Santa Rosa, CA 95407
(707) 528-1230 • FAX 707-527-6640 • TELEX 176687

KINEMETRICS | TRUETIME

CIRCLE NO. 181

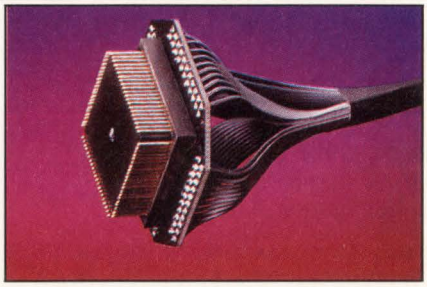

Easy Emulator Pods & Adapters

- Plug your PLCC and LCC packages into your PC board in minutes, with these easy-to-use adapters.
- Emulator/logic analyzer users: Adapt-a-Pod™ converts one package type to another (LCC, PLCC, PGA, and DIPs).
- Emulator pods and adapters are available in all standard pin counts, with ribbon or ribbon cable headers.
- Custom engineering services and do-it-yourself emulator pod converters. Free catalog.

Emulation Technology, Inc.
2344 Walsh Ave. Santa Clara, CA 95051
Phone: 408-982-0660 FAX: 408-982-0664

CIRCLE NO. 182

Passive Backplane CPUs

80386SX 16 MHz & 20 MHz, 512K through 16M DRAM

80386DX 25 MHz & 33 MHz, 1M through 32M DRAM, 128K cache optional

80486 25 MHz & 33 MHz, 1M through 32M DRAM, 256K cache optional

Call, write, or FAX for complete specifications and OEM terms.

Retail, OEM, and Private Label available

Designed, Manufactured, Sold and Serviced by:

PTech Inc.
Voice . . . (612) 345-4555
FAX . . . . (612) 345-5514
Modem . . (612) 345-4656

907 N. 6th St., Lake City, MN 55041

CIRCLE NO. 183

Program Your Chips

In Sets of 4 for $495.00

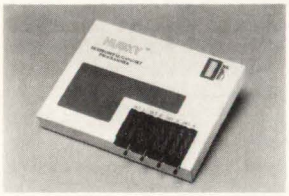

Special offer Now Includes:
Free UV eraser, CUPL starter Kit and a **$300.00 Rebate** with the PDT-1 Universal Programmer System\Kit.

LOGICAL DEVICES, INC. **1-800-331-7766**

CIRCLE NO. 184

SYSTEM SHOWCASE

C for the 8051 Compare:

Benchmark Results—Sample program: Eratosthenes Sieve Program from BYTE (1/83), expanded with I/O and interrupt handling.

	FRANKLIN SOFTWARE C-51 v2.1	MCC51 v1.2	Archimedes ICC51 v2.20A
Linkage time	*6 sec*	*9*	*29*
Execution time	*0.88 sec*	*9.00*	*11.45*
Total code size	*1726 bytes*	*3798*	*5318*
Compilation time	*17 sec*	*18*	*12*
Sieve module size	*541 bytes*	*1021*	*736*

Call now for your free DEMO disk.

FRANKLIN SOFTWARE, INC.

888 Saratoga Ave. #2 • San Jose, CA 95129

(408) 296-*8051* • FAX (408) 296-8061

CIRCLE NO. 185

IMAGE COMPRESSION

PARALLEL JPEG
UNPARALLELED PERFORMANCE

OPTIPAC - Flexible Compression Solution

- Color, Gray Scale and Binary Images
- Lossless and 10 Controlled Quality (CQ) Modes
- Varying CQ Compression Ratios in Excess of 100:1
- Lossless Ratios Up to 3:1 or More
- Extendible I/O Capability for Any I/O Device
- Fully Programmable Compression Engine
- A Single Software Interface

Configurations: • Up to Five (5) DSPs • Software Only Version

Standards: • JPEG for Gray Scale and Color • CCITT Group III & IV for Binary

Speeds: Compress 512 x 400 Full Color Images in Less Than One (1) Second

Advantages: • Increased Communication Rates • Increased Disk Storage Capacity

OPTIVISION, INC. 1-800-562-8934
1477 DREW AVE. DAVIS, CA 95616 FAX: (916) 756-1309

CIRCLE N. 186

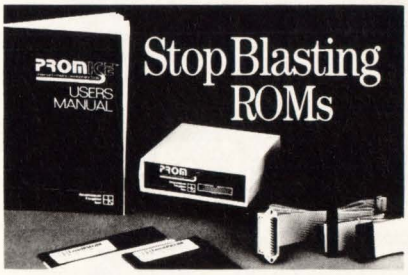

PROMICE emulates 8 bit ROMs from 2716-27080, or 16 bit ROMs 27C1024 or 27C2048. (Inquire about emulating other ROMs. Non-JEDEC ROMs require custom cable.) ■ Sophisticated LoadICE™ Host Software downloads, uploads and edits ROM contents, supports MS-DOS, UNIX, MAC & VMS. Software sources are included. ■ Bi-directional **Serial** link, autobaud to 57.6KB–loads 1 Mbit in 25 secs. ■ Bi-directional **Parallel** port (option)–loads 1 Mbit in 4 sec. ■ Emulate up to 2 ROMs per unit, daisy-chain up to 256 ROMs from one port! ■ **New! Analysis Interface™** (option) implements a ROM-based UART for sophisticated debugging.

Grammar Engine Inc
1161 Cherry Street
San Carlos
California 94070
415/595-2252

CIRCLE NO. 187

X.25 SDLC QLLC HDLC ADCCP PAD

- C source code
- ROM-able
- Full porting provided
- No OS required

GCOM, Inc.
1776 E. Washington
Urbana, IL 61801
(217) 337-4471

Specialists in Computer Communications
FAX 217-337-4470

CIRCLE NO. 188

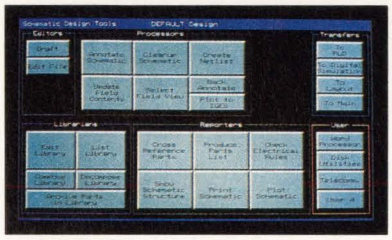

The Best Schematic Design on the PC

Call for a FREE demonstration disk of OrCAD's new framework for the DOS environment. See the ways that powerful-yet-affordable electronic design software has reached a new level of performance. For as little as $595, you can own OrCAD's Schematic Design Tools, the most popular CAE product in the world.

OrCAD
3175 N.W. Aloclek Drive
Hillsboro, Oregon 97124
Phone (503) 690-9881 Fax (503) 690-9891

CIRCLE NO. 189

QFP/PQFP Socket/Adapter System

Patents issued and/or pending.

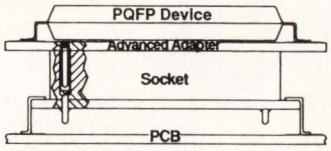

- Allows interface of PQFP to SMD footprint for testing.
- Allows interconnecting of PQFP device without soldering to PCB.
- Socket/Adapter system affords repeated use of PQFP socket.

ADVANCED INTERCONNNECTIONS®
5 Energy Way, P.O. 1019, W. Warwick, RI 02893
Tel (401)823-5200 FAX (401)823-8723

CIRCLE NO. 190

1991 UPCOMING ISSUES *Computer Design Magazine Edition*

MONTH	SPECIAL REPORT	TECHNOLOGY FOCUS	PRODUCT FOCUS
March 1*	Fuzzy logic in embedded control *Tom Williams*	Mixed CMOS, ECL & BiCMOS — *Barbara Tuck* LAN controllers ICs — *Ron Wilson*	VME CPU boards *Jeff Child*
April 1 Electro	PCB layout tools *Mike Donlin*	Communication with standard buses — *Warren Andrews* Designing ASICs for testability — *Barbara Tuck*	Static RAMs *Jeff Child*
May 1* CICC, Comdex	Superfast processors *Ron Wilson*	High-level design languages — *Mike Donlin* Object-oriented programming — *Tom Williams*	Emulators *Jeff Child*
June 1 DAC	Design synthesis *Barbara Tuck*	Mil-Spec standard buses — *Warren Andrews* Disk controller ICs — *Ron Wilson*	Op amps *Jeff Child*

**Starch Readership Study Issue*

PennWell Publishing Company • One Technology Park Drive • Westford, MA 01886 • 508-692-0700

COMPUTER DESIGN Technology and Design Directions

Break Through to the Future

Imagine a future with no walls. Product development teams working together to improve yields and shorten production cycles. No barriers. Leaving competition in the dust. Three words are making walls crumble—Design For Manufacturability. A process that's forcing electronic manufacturing disciplines to work concurrently. It's the only way companies can maintain leadership positions in the global marketplace.

Make the break to the future. Find working solutions to achieve DFM. Plan to participate in The National Electronic Manufacturing and Design Exposition & Conference. A first time event for engineering and management professionals involved in the manufacture of printed circuit boards, including design, component selection, assembly, production and final test. There is no other Midwest event that satisfies the need for a complete electronic manufacturing experience.

ACT NOW—Please call Michele Filippi, NEMDE '91 Show Manager at (708) 390-2458 for information on exhibiting and/or attending.

NEMDE '91 is part of National Manufacturing Week. Eight co-located Cahners Exposition Group events at one time with complementing technologies and industries.

NEMDE '91

NATIONAL ELECTRONIC MANUFACTURING AND DESIGN EXPOSITION & CONFERENCE

Exposition and Conference: April 8–11, 1991
McCormick Place, Chicago, IL

Straight out of Tomorrow

Cahners Exposition Group
Cahners Plaza • 1350 East Touhy Avenue • P.O. Box 5060 • Des Plaines, IL 60017-5060
Phone 708/299-9311 • Telex 82882 CEG/CHGO • FAX 708/635-1571

©Reed Publishing (U.S.A.) Inc., 1990

For Exhibitor Info Circle No. 58

For Attendee Info Circle No. 59

ADVERTISERS INDEX

ADVERTISER	PAGE NO.	CIRCLE NO.
A		
Advanced Interconnections	117	190
Alacron	96	51
AMP	4,5	4
Analogic	69	42
Applied Microsystems	48,49	65
B		
Boston Systems Office	36	23
BT&D Technologies	44	29
C		
Chrislin Industries	63	39
Ciprico	Cover 2	1
CMC	6,7	5
CSPI	41	27
Cyclone	53	33
Cypress	Cover 4	—
D		
Data I/O	68	41
Data I/O	73	44
Data Translation	47	30
Diversified Technology	20	15
DY-4 Systems	11	7
E		
Electronic Conventions	92	49
Electronic Solutions	2	2
Emulation Technology	116	182
Epson	22	64
F		
Fil-Mag	20	14
Force Computers	19	13
Franklin Software	117	185
Franklin Software	60	38
G		
Gates Energy Products	14,15	9
GCOM	117	188
General Micro Systems	52	32
Gespac	2	3
Graham	59	37
Grammar Engine	117	187
H		
*Hitachi Peripherals	21	17
I		
IBUS	101	55
Idea Show	103	69
Industrial Computer Solutions	116	177
Intel Microcomputer	33	21
Intel	119	60
Ironics	16	10
Ironics	18	12
K		
Kinemetrics	116	181
L		
L-Com	116	178
Logical Devices	116	184
M		
Macrolink	39	26
Matrix	9	6
Microtek	67	40
Microware	29	19
Mitsubishi	111-113	73
Mizar	58	36
N		
NEC	42,43	28
NEMDE 91 Exposition	118	58,59
Newbridge	70	43
O		
Oasys	37	24
Omation	97	52
Omnibyte	80A	119,120
Optivision	117	186
ORCAD	117	189
P		
PC Tech	116	183
Procise	116	179
R		
Radisys	13	8
Radstone	31	20
Ricoh	17	11
S		
SBE	Cover 3	62
Sealevel Systems	110	56
Sierra	34	22
Silicon Composers	116	180
Simpact	57	35
T		
Tadpole	90,91	48
Teknor	51	31
Tektronix	55	34
Texas Instruments	26,27	—
Themis	85	46
TITN	77	45
U		
Ultimate Technology	38	77
V		
Valley Enterprises	89	47
Vector	120	61
Vitesse Semiconductor	98,99	67
X		
Xycom	103	68
Z		
Zendex	96	50
Ziatech	106	53
Zwick Systems	116	176
RECRUITMENT ADVERTISERS		
Compaq	115	57

* Domestic only

† International only

The Advertisers Index is published as a service. The publisher does not assume liability for errors or omission.

intel®

Access in Europe

Microelectronics Components Boards Systems

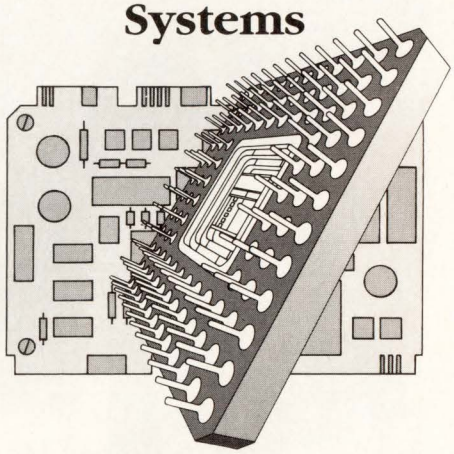

EUROPEAN SALES OFFICES

FINLAND
Intel
Ruosilantie 2
00390 Helsinki
Tel.: (358) 0 544 644
Fax: (358) 0 54 40 30
Telex: 1 23 332

FRANCE
Intel
1. rue Edison – BP 303
78054 St. Quentin-en-Yvelines Cedex
Tel.: (33) (1) 30 57 70 00
Fax: (33) (1) 30 64 60 32
Telex: 6 99 016

UNITED KINGDOM
Intel
Pipers Way
Swindon, Wilts SN3 1RJ
Tel.: (44) (0793) 69 60 00
Fax: (44) (0793) 64 14 40
Telex: 4 44 447/8

NETHERLANDS
Intel
Postbus 84130
3099 CC Rotterdam
Tel.: (31) 10.407.11.11
Fax: (31) 10.455.46.88
Telex: 2 2 283

ISRAEL
Intel
Atidim Industrial Park-Neve Sharet
P.O. Box 43202
Tel-Aviv 61430
Tel.: (972) 03-498080
Fax: (972) 03-491870
Telex: 3 71 215

ITALY
Intel
Milanofiori Palazzo E
20090 Assago Milano
Tel.: (39) (02) 89 20 09 50
Fax: (39) (02) 349 84 64
Telex: 3 41 286

SPAIN
Intel
Zurbaran, 28
28010 Madrid
Tel.: (34)-1-308.25.52
Fax: (34)-1-410.75.70
Telex: 46880

SWEDEN
Intel
Dalvägen 24
171 36 Solna
Tel.: (46) 8 734 01 00
Fax: (46) 8 27 80 85
Telex: 12261

SWITZERLAND
Intel
Zürichstrasse
8185 Winkel-Rueti bei Zürich
Tel.: (41) 01/860 62 62
Fax: (41) 01/860 02 01
Telex: 8 25 977

GERMANY
Intel
Dornacher Strasse 1
8016 Feldkirchen
Tel.: (49) 089/9 09 92-0
Fax: (49) 089/9 04 39 48
Telex: 5-23177

Intel
Abraham Lincoln Strasse 16-18
6200 Wiesbaden
Tel.: (49) 06 11/76 05-0
Fax: (49) 06 11/71 86 15
Telex: 4-186 183

Intel
Zettachring 10 A
7000 Stuttgart 80
Tel.: (49) 07 11/72 87-280
Fax: (49) 07 11/72 80 137
Telex: 7-254 826

DORN ADVERTISING

CIRCLE NO. 60

Vector And EISA Put You On Board First.

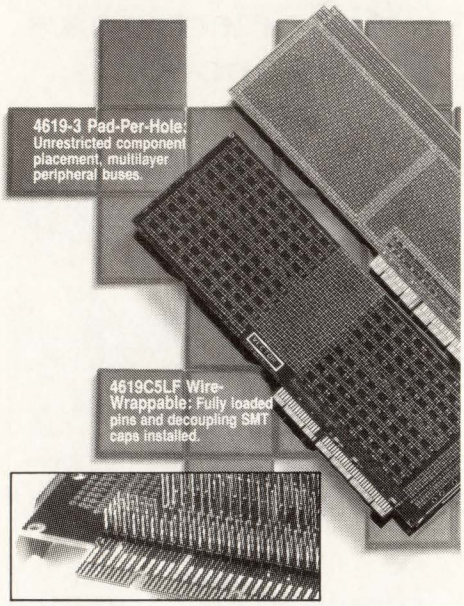

Vault into the surging EISA design market first with these new design tools; industry's first off-the-shelf EISA support line.

Whether you need a wire-wrappable, or pad-per-hole prototyping board with PGA and PLCC geometries, or an extender card to test your functioning EISA board outside its development system, you'll have the design edge.

3690-32
EISA Bus Extender: extends bus for easy access & testing.

Visit your local Vector distributor for these new, off-the-shelf EISA development tools. Or, call us direct for complete specs and pricing.

VECTOR
ELECTRONIC COMPANY
12460 Gladstone Ave., Sylmar, CA 91342
Inside CA (800) 426-4652
Outside CA (800) 423-5659
FAX (818) 365-5718
CIRCLE NO. 61

COMPUTER DESIGN
Technology and Design Directions

Executive Office: One Technology Park Dr, PO Box 990, Westford, MA 01886
(508) 692-0700 Fax: (508) 692-0525 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 (516) 883-8815
Marketing Communications Assistant/Research: Annette Petagna (508) 392-2209
Administrative Assistant: Peg Alexander (508) 392-2112
Ad Traffic Manager: Betsy Anderson (508) 392-2198
Recruitment Advertising: Sue Nawoichik (508) 392-2194, Eastern U.S.
Tom Murphy (213) 372-2744, Western U.S.
Postcards/System Showcase/Bulletin Board: Sue Nawoichik (508) 392-2194
Circulation Director: Bob Dromgoole (918) 832-9213
Circulation Manager: George Andrew (918) 831-9443
List Rental: Deanna Rebro (918) 831-9551

SALES OFFICES

Associate Publisher
Tim L. Tobeck

Mid Atlantic
Eleanor Hobbs
14 Vanderventer Ave, Suite 1
Port Washington, NY 11050
Tel: (516) 883-8815
Fax: (516) 883-6268

New England/New York State/ Eastern Canada
Tim Pritchard
One Technology Park Dr
PO Box 990
Westford, MA 01886
Tel: (508) 392-2217
Fax: (508) 692-0525
Telex: 883436

Upper Midwest
George Mannion
21517 Locust
Matteson, IL 60443
Tel: (708) 748-7565
Fax: (708)748-7695

South Central/South East
Eric Jeter
19627 I-45 N, Suite 220
Spring, TX 77388
Tel: (713) 353-0309
Fax: (713) 288-8350

Southern California/Southwest
Tom Boris, Greg Cruse, Phil Cook
2232 SE Bristol, Suite 109
Santa Ana Heights, CA 92707
Tel: (714) 756-0681
Fax: (714) 756-0621

Northern California/ Nevada/Western Canada
John Sly, Bill Cooper
1000 Elwell Ct, Suite 234
Palo Alto, CA 94303
Tel: (415) 965-4334
Fax: (415) 965-0255

Oregon
Tom Boris
2232 SE Bristol, Suite 109
Santa Ana Heights, CA 92707
Tel: (714) 756-0681
Fax: (714) 756-0621

Washington
Bill Cooper
1000 Elwell Ct, Suite 234
Palo Alto, CA 94303
Tel: (415) 965-4334
Fax: (415) 965-0255

U.K./Scandinavia
David Round
69 Imperial Way
Croydon, Surrey CRO 4RR,
England
Tel: 81-686-7655
Telex: 938420
Fax: 81-688-2134

France/Belgium/ S. Switzerland/Spain/ The Netherlands
Daniel R. Bernard
10, rue Michelet - B.P. 279
78502 Sartrouville Cedex, France
Tel: (1) 39-14-67-80
Telex: 214235F MIX
Attn: Missitex: PENNWELL
Fax: (1) 39-14-70-14

W. Germany/Austria/ N. Switzerland/E. 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
Sumio Oka
International Media Representatives
Room 100, 21 Bldg.
Z-2-22, Okusawa
Setagaya-ku, Tokyo 158
Tel: 03-502-0656
Telex: J22633 MEDIAREP
Fax: Tokyo 5706-7349

Southeast Asia
Jay G. Seo
Yong-Jin Park
Doobee International, Ltd
Center Building (Byulgwan)
1-11 Jeong-dong, Choong-ku
CPO Box 4557
Seoul, Korea
Tel: 776-2096
Telex: K27117 DOOBEES
Fax: 755-9860

COMPUTER DESIGN (ISSN-0010-4566) is published twice monthly by PennWell Publishing Company, 1421 S Sheridan, Tulsa, OK 74112. Second-class postage paid at Tulsa, OK 74112 and additional mailing offices. Editorial offices are located at One Technology Park Dr, PO Box 990, Westford, MA 01886. Annual subscription price: $88 in the U.S.A., $139 in Canada, $148 for Europe via air, and $192 other foreign via air. 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.

©1991 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.

Fasten Your Seat Belt... Token Ring Accelerates to 16 Mbps.

Two New SBE 16 Mbps Controllers Bring High-Speed Token Ring to VMEbus/Multibus Systems.

One advantage of Token Ring is that it provides an efficient, high-performance interconnect with IBM mainframes. In a multinodal LAN environment, Token Ring provides *four times the throughput of Ethernet.*

SBE delivers high-performance Token Ring with two new intelligent *16 Mbps* communications controllers that interface VMEbus/Multibus Systems with Token Ring LANs.

SBE's Token Ring Controllers include these features:

- Software-selectable interface for 4 or 16 Mbps.
- High-speed, on-board 32-bit 68020/68030 25MHz processors.
- 1 MB or 4 MB of DRAM.
- Support for IEEE 802.5 standards.

Turn to SBE and discover the difference these new 16 Mbps VMEbus/Multibus Controllers can make in your LAN application.

For fast action, call: 1-800-347-COMM.

CIRCLE NO. 62

SBE
Communications & Real-time Solutions

THE SPARC OF A NEW GENERATION.

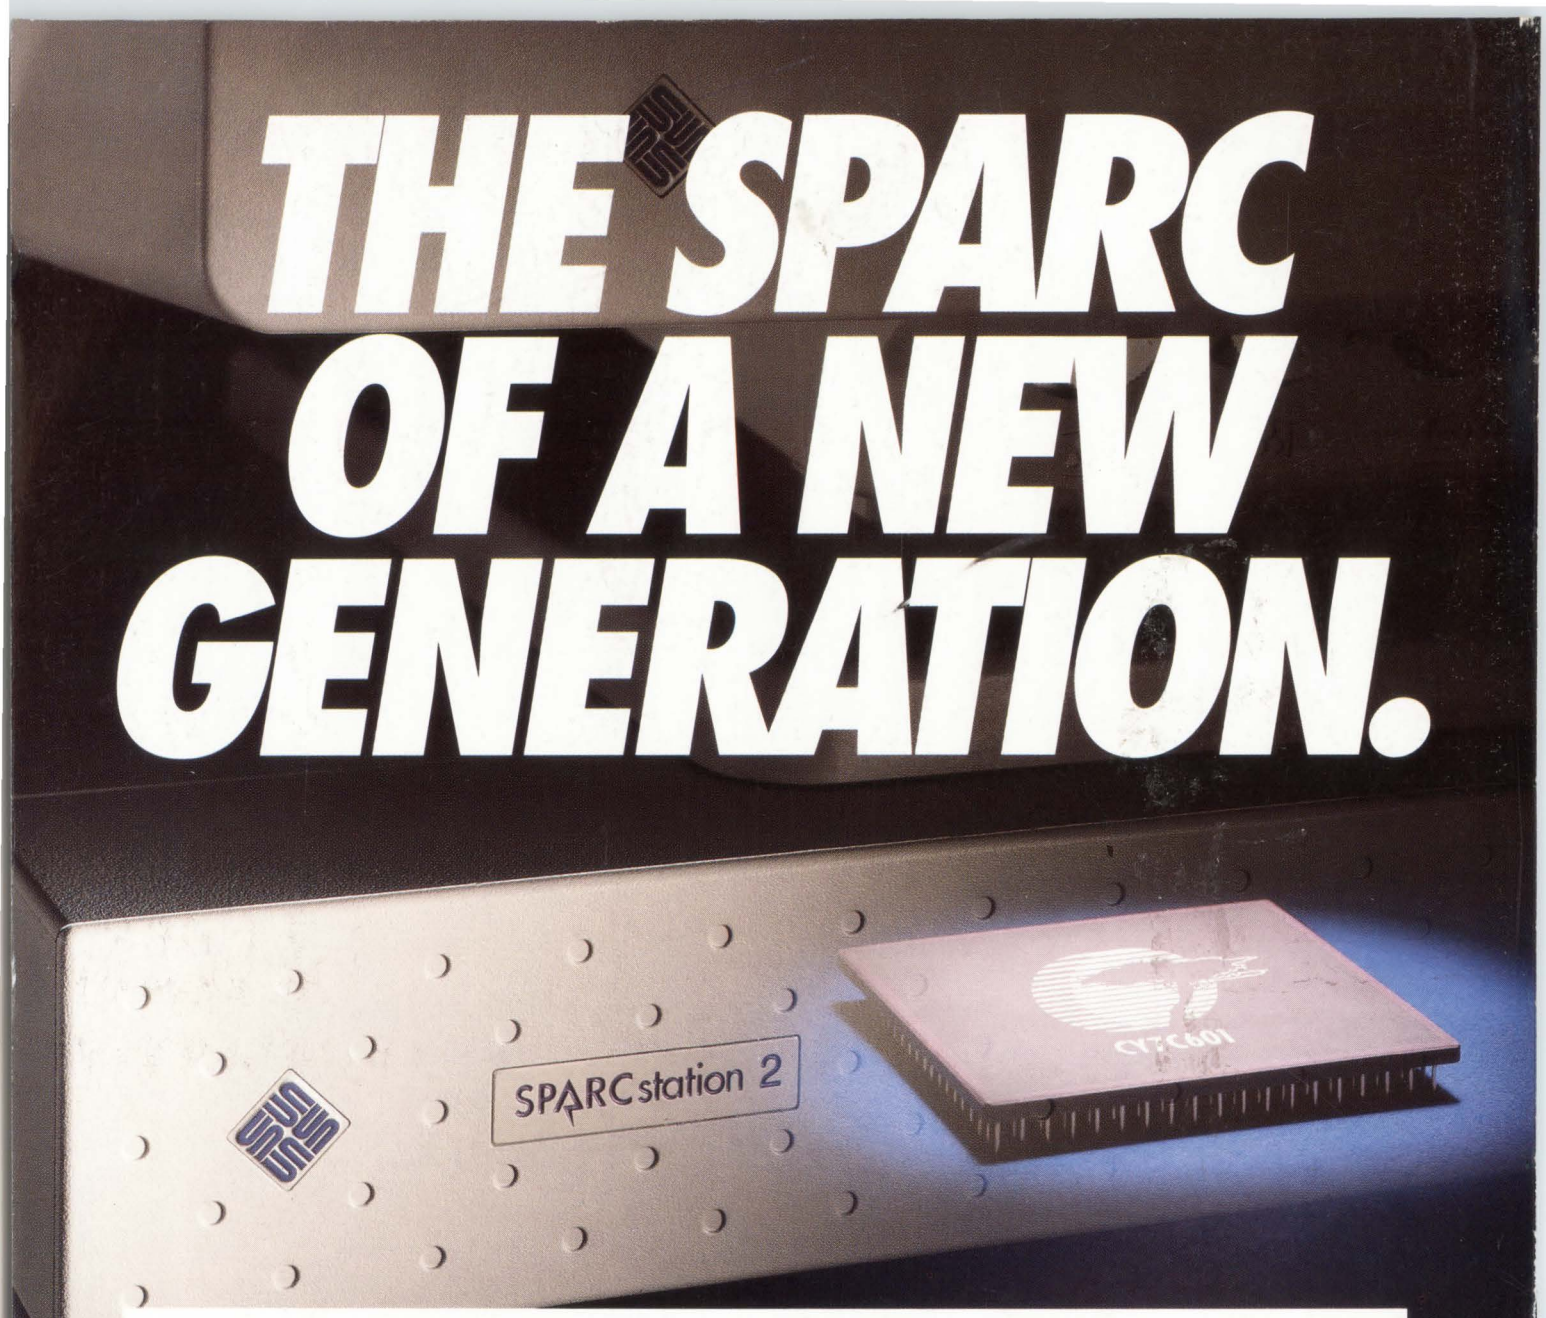

A new generation of performance.

The new SPARCstation™2 from Sun Microsystems delivers almost twice the CPU performance of previous generations – the highest performance in its class.

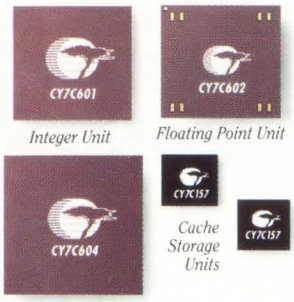

Integer Unit *Floating Point Unit*

Cache Storage Units

Cache Controller and Memory Management Unit

The complete SPARC chipset, designed by our ROSS Technology subsidiary, is now available for high-volume applications.

The reason for this blazing speed: The highest-performance CMOS SPARC® available. The CY7C601, 32-bit SPARC microprocessor and CY7C157 16K x 16 Cache Storage Units from Cypress Semiconductor. Cypress 0.8 micron CMOS gives the CY7C601 SPARC RISC microprocessor 40 MHz clock rates and a sustained throughput of 29 MIPS.

Next generation compatibility.

Because they're SPARC chips, they're completely compatible with over 2100 SPARCware™ products. They offer true scalability into a wide spectrum of applications. And a growth path for future high-performance products.

Get this level of performance, potential, and compatibility for your next-generation products.

The Cypress SPARC.

For the whole story on RISC, call the Cypress Hotline at 1-800-952-6300.* Ask for dept. C11M.

CYPRESS SEMICONDUCTOR

*1(800)387-7599 in Canada. (32)2 672-2220 in Europe. ©1990 Cypress Semiconductor, 3901 North First Street, San Jose, CA 95134. Phone: (408) 943-2600. Telex: 821032 CYPRESS SNJ UD. TWX: 910-997-0753. SPARC is a registered trademark of SPARC International, Inc. SPARCstation and SPARCware are trademarks of SPARC International, Inc., licenced exclusively to Sun Microsystems, Inc. Products bearing the SPARC trademark are based on an architecture developed by Sun Microsystems, Inc.