A PENNWELL PUBLICATION NOVEMBER 1, 1990

COMPUTER DESIGN

Technology and Design Directions®

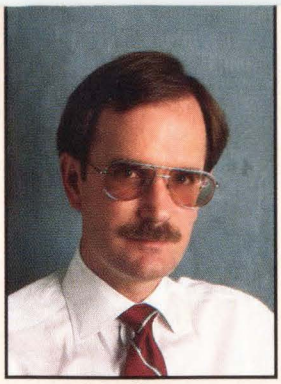

James Ready on real-time standards

FOR ELECTRONIC ENGINEERS & ENGINEERING MANAGERS

Are designers' expectations for synthesis realistic?

Cache controllers tread a rocky path to integration

U Interface silicon crystallizes ISDN standards

How to Design User Friendly Systems *in No Time*

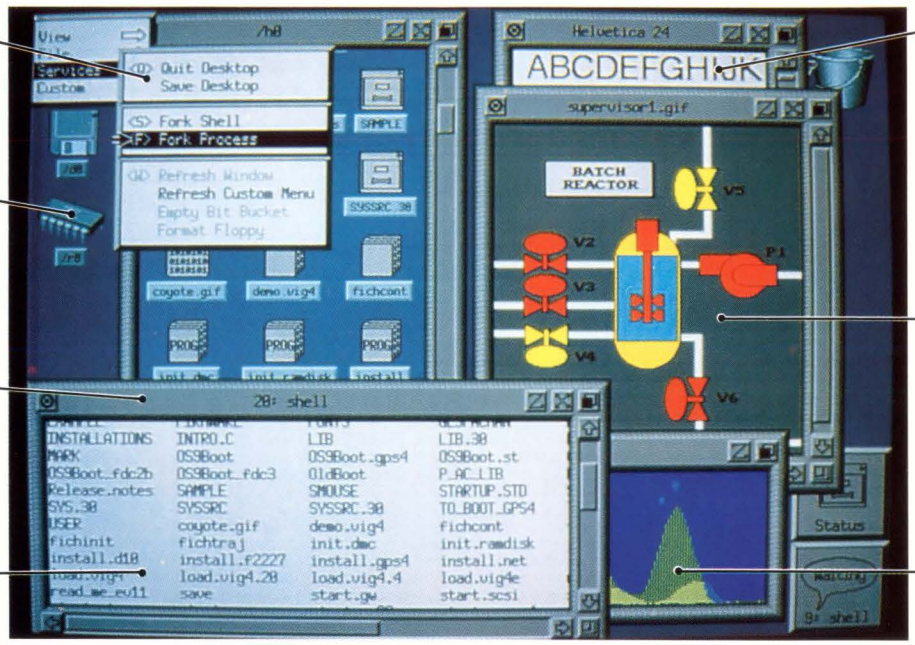

Introducing G-Windows, a complete Graphical User Interface for the OS-9 Real-Time Multitasking Operating system. G-Windows lets you build a visually attractive and powerful user interface using a set of simple and easy-to-use function calls to create windows, menus, and a variety of input/output gadgets. An application editor, G-View lets you draw screens with gadgets that link easily into your application program written in C. What's more, G-Windows is totally multitasking (all windows can be active and updated at the same time). It is extremely memory efficient and can be ROMed for embedded system applications.

G-64/G-96 or VME Bus Platform

G-Windows will operate on any G-64/G-96 or VME bus 680x0 system running OS-9. Just plug GESPAC's Graphics, mouse and keyboard interface card into your backplane, install the G-Windows software and your system is ready to shine with 256 colors and resolution of 640 by 480 and higher.

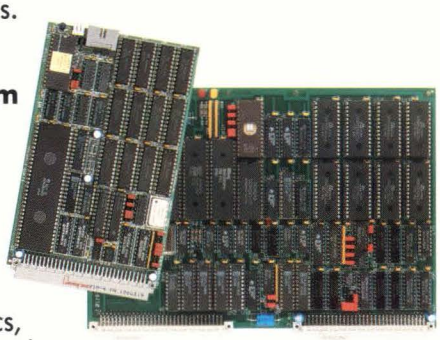

Seeing is believing

That's why we have produced a video that demonstrates the impressive features and capabilities of G-Windows. Call to request your free copy.

...And that's just the beginning

There are over 200 hardware and software microsystems components in GESPAC's catalog so that there is no limit as to what you can accomplish. Call GESPAC today to receive your free copy.

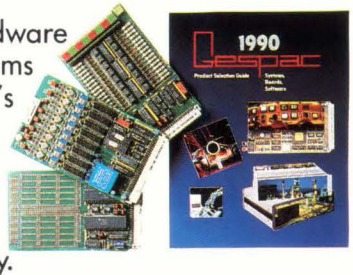

Toll Free 1-800-4-GESPAC or (602) 962-5559.

Gespac

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

EUROPE
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. 1

O M E G A

New

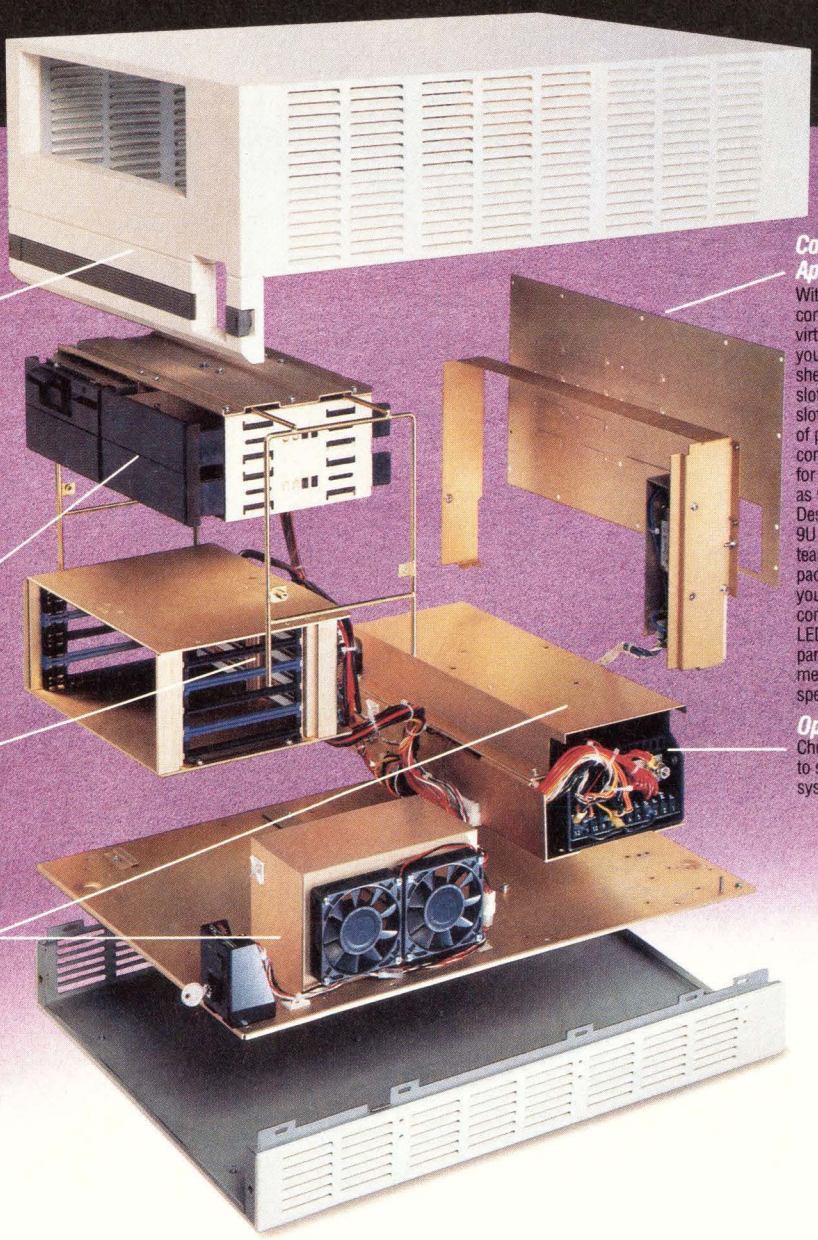

Easy Access = Maximum Accessibility
Quick access saves time in assembly, test, and maintenance. Loosen two captive fasteners and the front, top, and sides slide off for maximum access to all of the internal components. Each internal component comes out with two screws, or remove five hex nuts and the entire system lifts out of the bottom tray to provide a free-standing, fully functional development platform!

Peripherals Just Slide In
Up to eight half-height 5¼" peripherals simply slide into their mounting brackets.

The Quietest Backplane Available
Plug your VME system into the quietest backplane ever! Our new 10-LAYER VME backplanes handle the fastest VME signals with the absolute minimum distortion, crosstalk, and noise.

Engineered Cooling
Airflow is directed over the components that need cooling—the power supply and your boards. Separate chambers maximize cooling by channeling airflow and restricting recirculation.

Configured for Your Application
With over 2000 standard configurations you can virtually design the enclosure yourself—and still get it off the shelf. Choose from 3 to 14 slots for VME or 4 to 18 slots for Multibus. Loads of peripheral options and configurations are available for rack and tabletop models as well as floor-standing DeskMate™ verions that accept 9U × 400mm cards. Or let our team of engineers provide a packaging solution tailored to your specifications—from configuring the front panel LEDs to customizing the rear panel. Designing enclosures to meet your specific needs is our specialty.

Optimum Power Output
Choose from 175 to 1000 Watts to satisfy the demands of your system.

The Last Word in System Packaging

Now you can benefit from 12 years of design experience and innovation that has made Electronic Solutions the industry leader in system packaging. The Omega™ Enclosures will give your VME or Multibus system the attractive exterior your image demands with a rugged, well-designed interior to withstand the most demanding environments. And besides the Electronic Solutions' full 3-Year warranty, the Omega bears FCC approvals and is UL, CSA, and TUV listed to get your system to an international market in the shortest time possible.

Call Electronic Solutions for the latest information about the best choice for your system package. It's the last system enclosure you'll need to see.

ZERO **Electronic Solutions**
UNIT OF ZERO CORPORATION

6790 Flanders Drive, San Diego, CA 92121
(619) 452-9333 FAX: 619-452-9464
Call **TOLL FREE (800) 854-7086**
in Calif. **(800) 772-7086**

CIRCLE NO. 2

See Us At ELECTRONICA '90
Munich, Germany (Nov. 6 - 10)
Booth F31, Halle 16

SMC®. Above the rest.

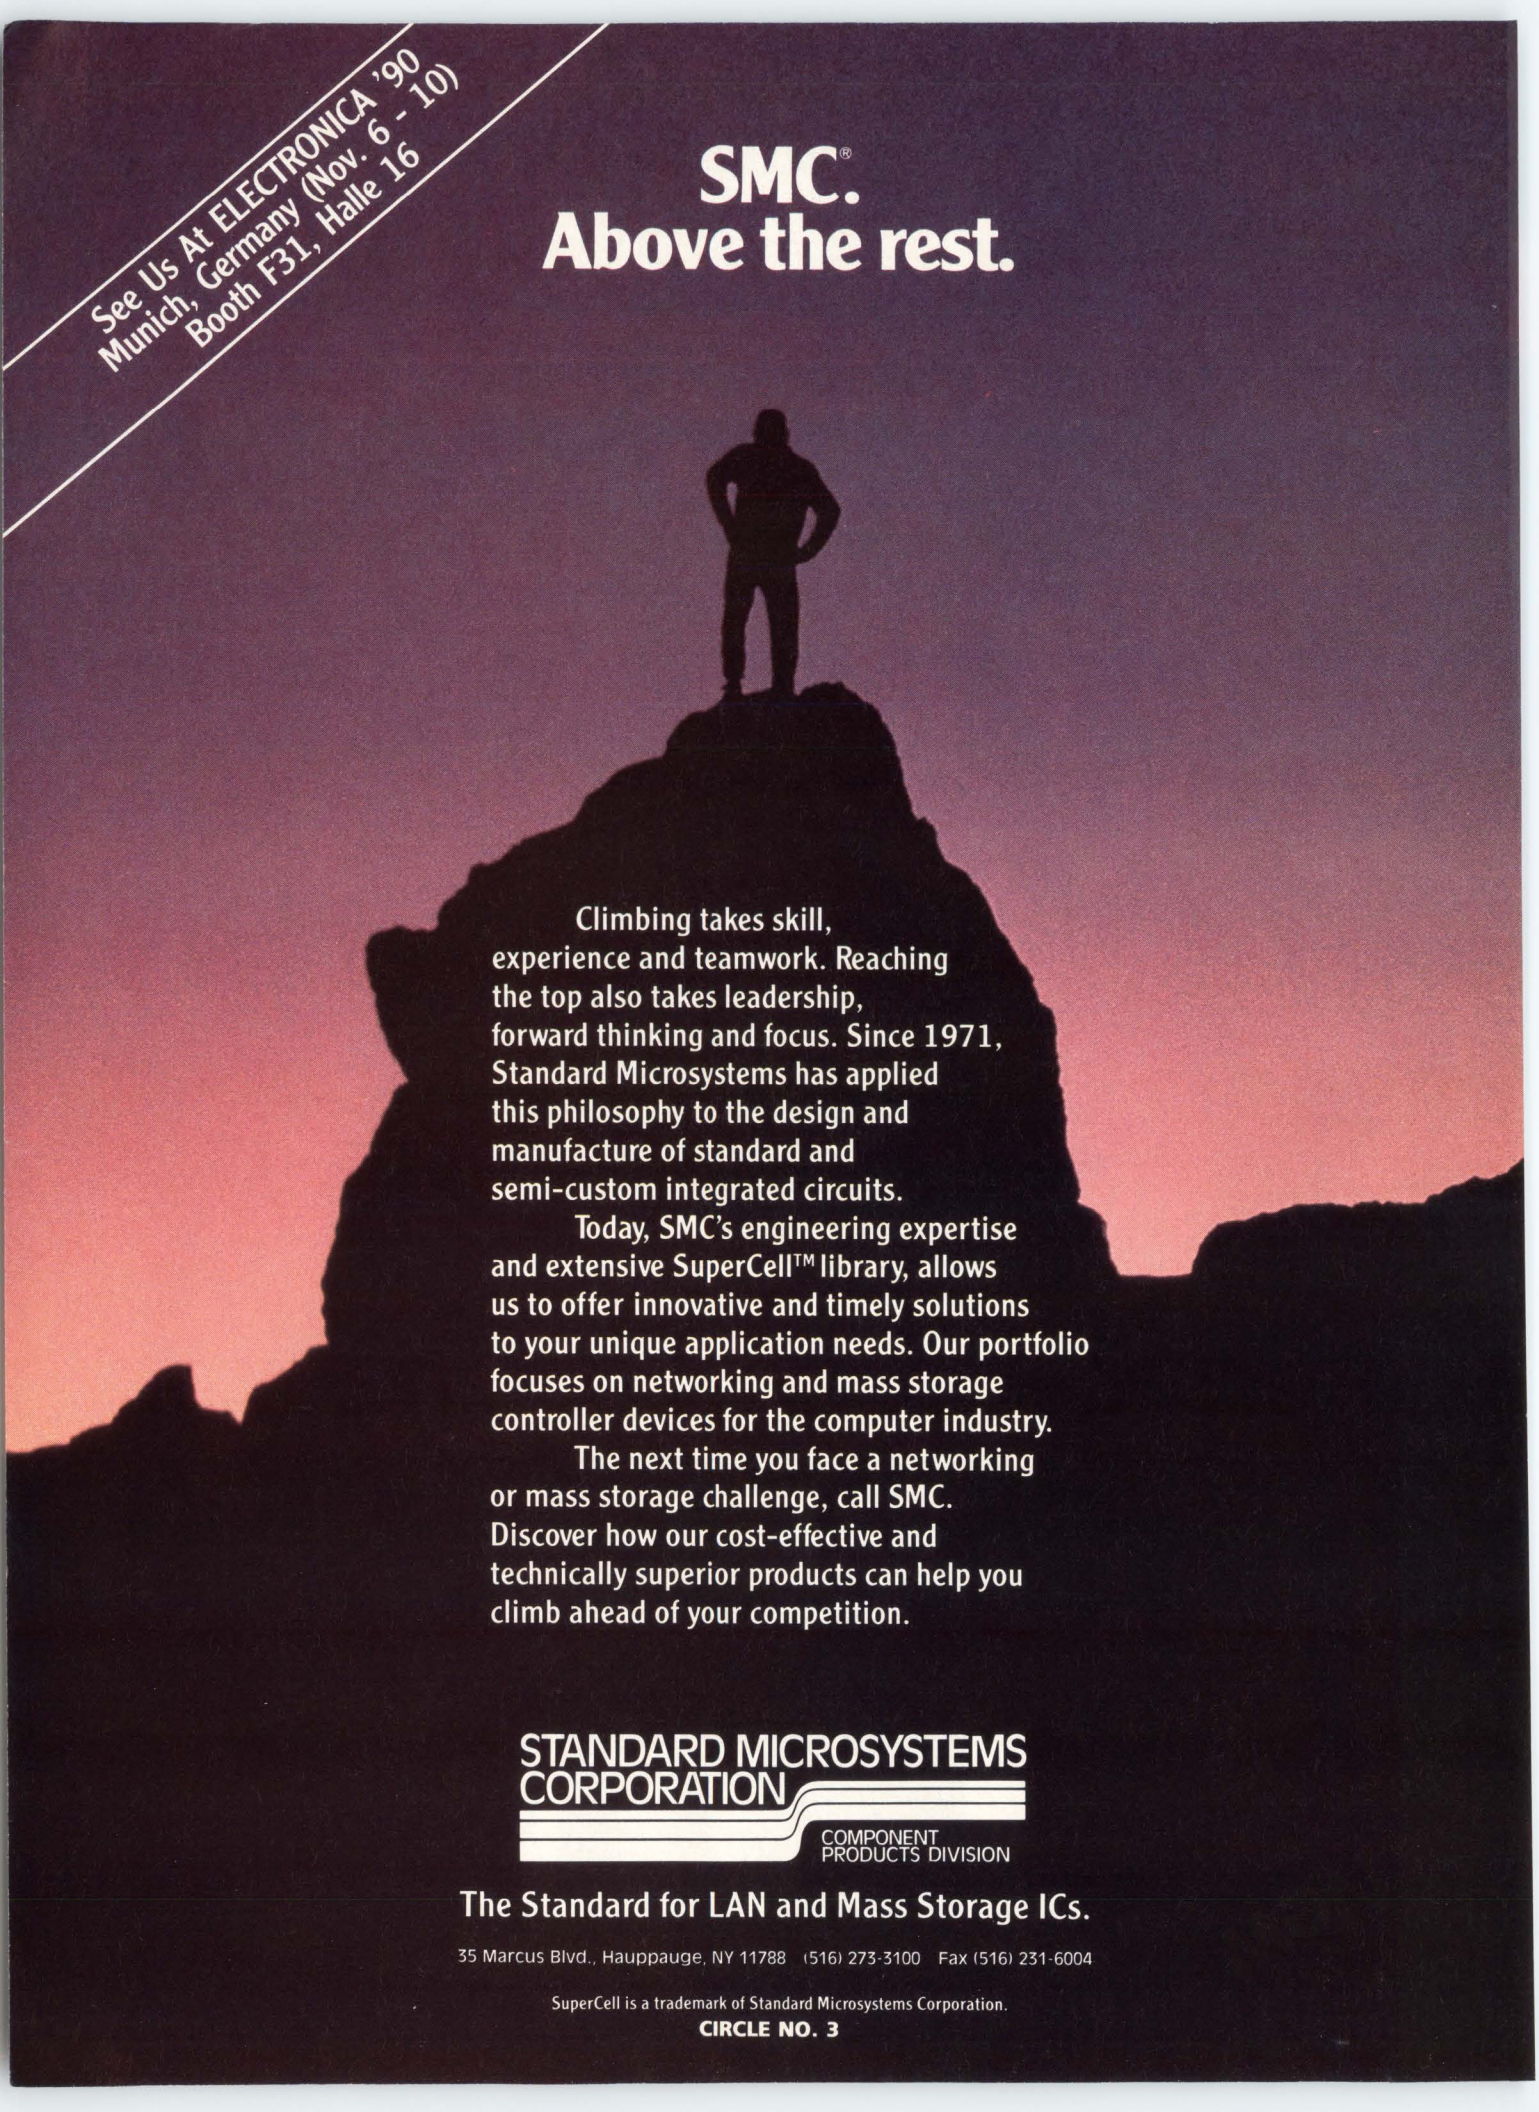

Climbing takes skill, experience and teamwork. Reaching the top also takes leadership, forward thinking and focus. Since 1971, Standard Microsystems has applied this philosophy to the design and manufacture of standard and semi-custom integrated circuits.

Today, SMC's engineering expertise and extensive SuperCell™ library, allows us to offer innovative and timely solutions to your unique application needs. Our portfolio focuses on networking and mass storage controller devices for the computer industry.

The next time you face a networking or mass storage challenge, call SMC. Discover how our cost-effective and technically superior products can help you climb ahead of your competition.

STANDARD MICROSYSTEMS CORPORATION
COMPONENT PRODUCTS DIVISION

The Standard for LAN and Mass Storage ICs.

35 Marcus Blvd., Hauppauge, NY 11788 (516) 273-3100 Fax (516) 231-6004

SuperCell is a trademark of Standard Microsystems Corporation.

CIRCLE NO. 3

VOL. 29, NO. 21 NOVEMBER 1, 1990

COMPUTER DESIGN

Technology and Design Directions®

FOR ELECTRONIC ENGINEERS & ENGINEERING MANAGERS

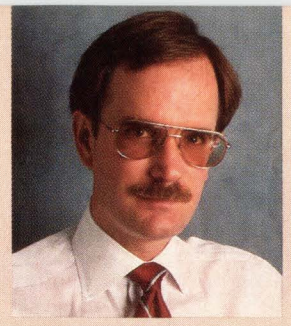

James Ready on real-time standards
Page 21

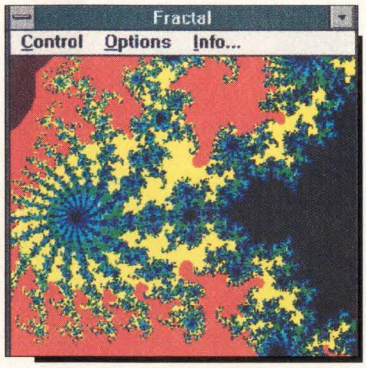

Page 74

TECHNOLOGY UPDATES

Page 99

TECHNOLOGY AND DESIGN FEATURES

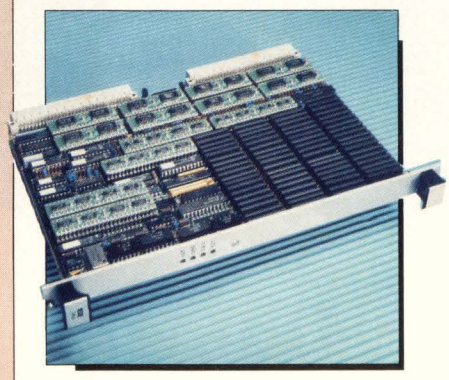

Page 131

NEW PRODUCT HIGHLIGHTS

DEPARTMENTS

If You Need Chips that Talk,

RealVoice™ Synthesizers offer Quality and Affordability

If you've talked about designing speech synthesizers into your system but think they sound like toys and are too complex and costly to work with, try Oki's RealVoice family of easy-to-use, *affordable* speech synthesizers, and listen to the difference.

Once you've heard the high-quality, realistic sound, you'll see why so many high-end consumer-interface products use Oki voice chips. Our proprietary ADPCM algorithm provides better than FM signal-to-noise ratio, making them ideal for sophisticated applications like PCs, home and car security systems, point-of-sale advertising, training tools, medical readouts, and more.

Plus these chips are easy to work with and more cost-effective than ever. On-chip features such as a low pass filter and 12-bit D/A reduce costly design time and IC count, yet increase system reliability. Choose from a range of devices: Synthesizers with 5- to 40-second message options, ROMless and OTP-based development chips, a pitch control device, and a solid-state recorder.

We also offer demo boards and our ASE88—a low-cost development system with IBM® PC AT™-based board and software for easy editing and waveform display.

Hearing is believing. Call **1-800-OKI-6388** for samples. See how easy and cost-effective it is to design RealVoice sound features into your system—it's quality that speaks for itself.

IBM is a registered trademark of International Business Machines. RealVoice is a trademark of Oki Semiconductor. IBM PC AT is a registered trademark of IBM.

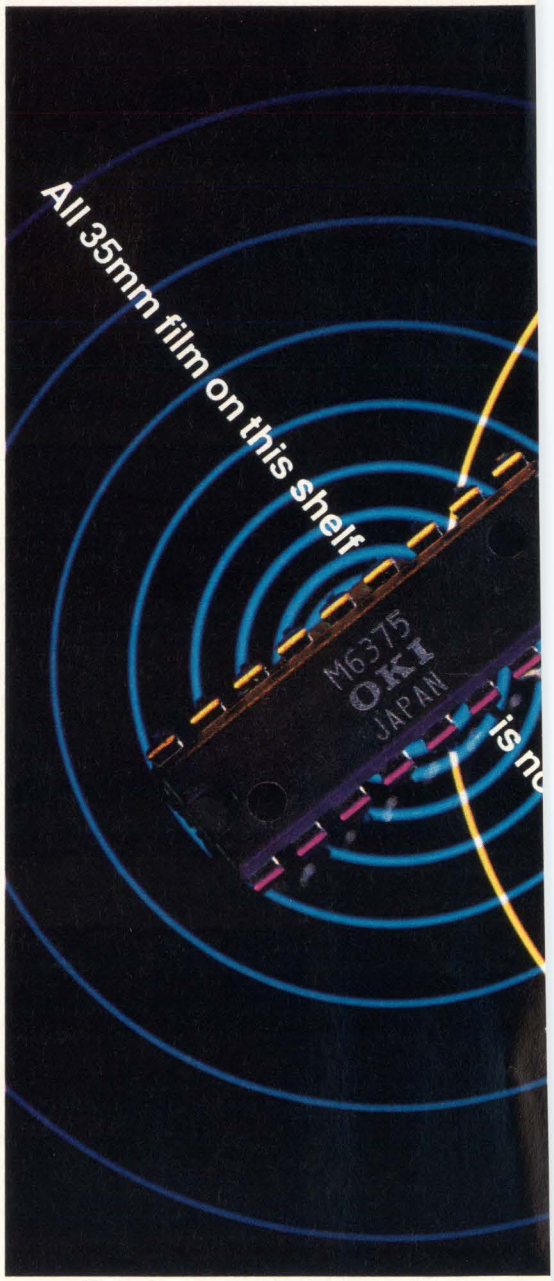

Transforming technology into customer solutions

Talk to Oki.

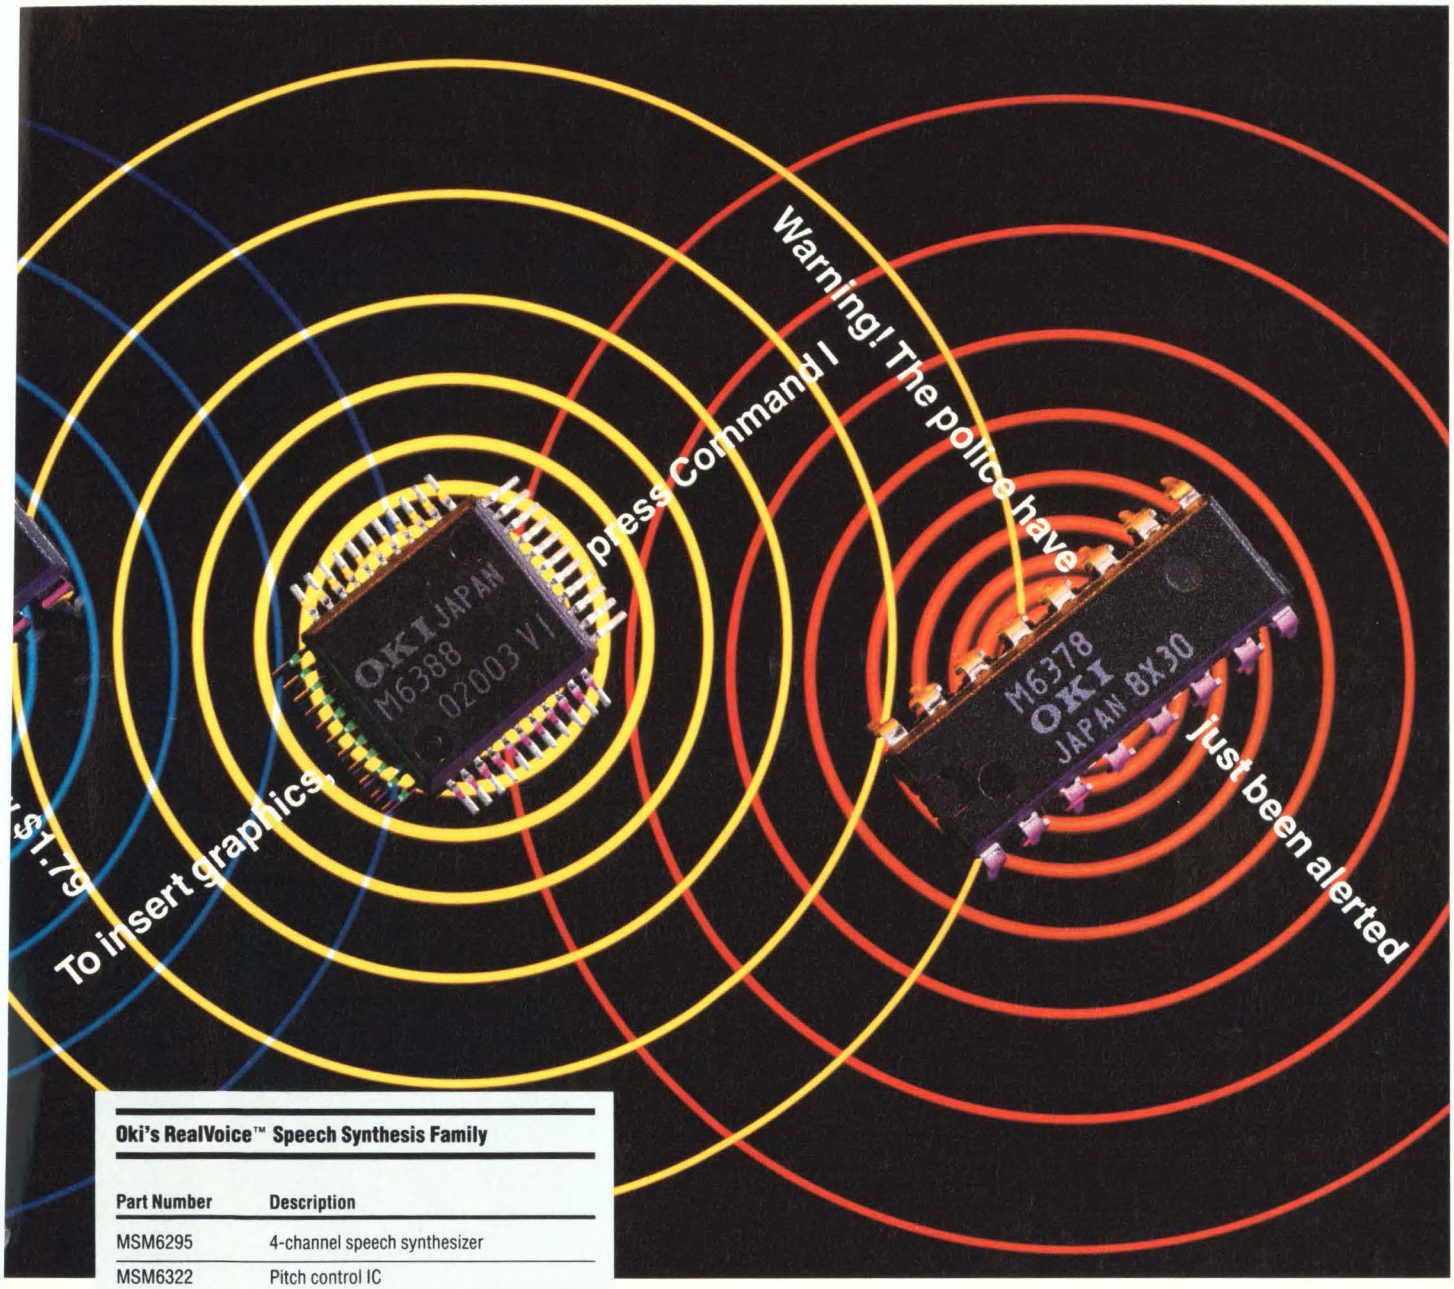

Oki's RealVoice™ Speech Synthesis Family

Part Number	Description
MSM6295	4-channel speech synthesizer
MSM6322	Pitch control IC
MSM6372	Speech synthesizer with 128K ROM, 5 secs
MSM6373	Speech synthesizer with 256K ROM, 10 secs
MSM6374	Speech synthesizer with 512K ROM, 20 secs
MSM6375	Speech synthesizer with 1M ROM, 40 secs
MSM6376	Evaluation chip for MSM6372/73/74/75
MSM6378	Speech synthesizer with 256K OTP ROM
MSM6388	Solid-state recorder/1M serial register I/F

OKI Semiconductor

CA Irvine 714/752-1843, Tarzana 818/774-9091, San Jose 408/244-9666 **FL** Boca Raton 407/394-6099
GA Norcross 404/448-7111 **IL** Rolling Meadows 708/870-1400 **MA** Stoneham 617/279-0293 **MI** Livonia 313/464-7200
NC Morrisville 919/469-2395 **NY** Poughkeepsie 914/473-8044 **PA** Horsham 215/674-9511 **TX** Richardson 214/690-6868

CIRCLE NO. 4

HOW TO TURN 040 WITHOUT LOSING A STEP.

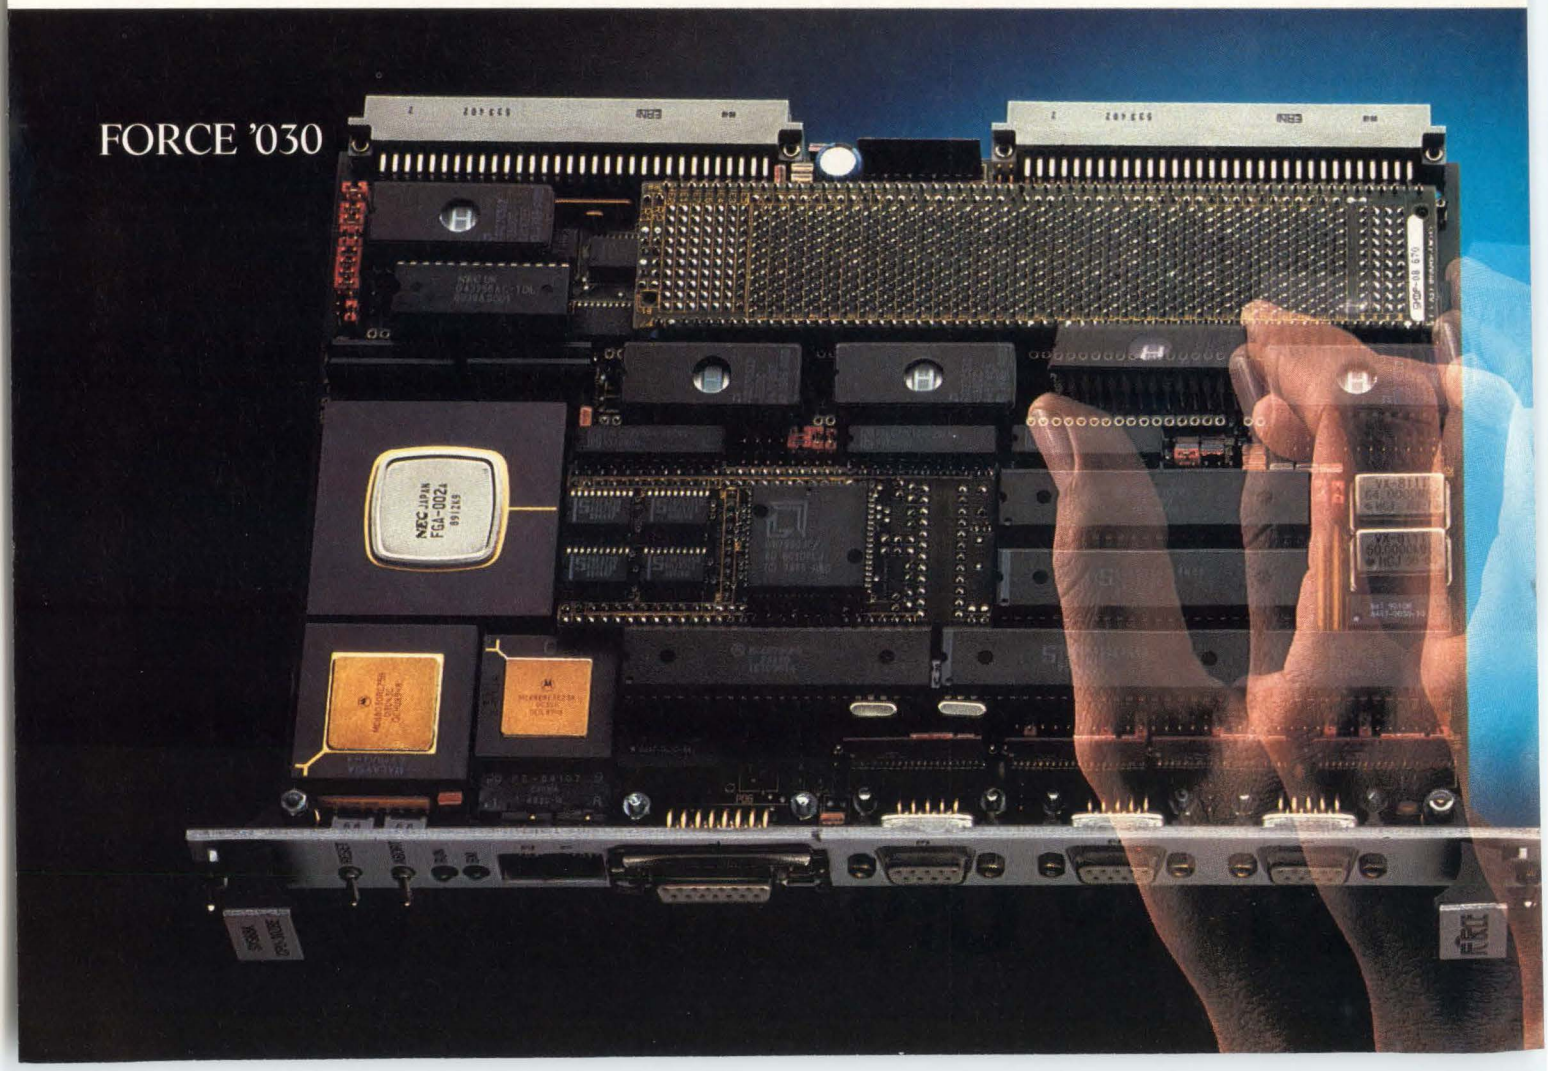

Turning 040 doesn't mean you have to give up the code you lived by when you were 030. Although that's what some manufacturers expect you to do.

But not FORCE. We guarantee that applications written for our 68030 VME boards will run on our 68040 boards. That's because we've built compatibility into our 030 and 040 address maps and on-board device drivers.

In fact, no one makes it easier to move your software from 030 to 040. The competition can't even come close. Just ask them.

Then ask us. We'll keep you from spending months writing new software drivers. So you can spend your time improving performance and functionality. Or getting to market months ahead of the competition.

What's more, you can start today on your 040 applications. Just develop them on a FORCE 030 board. When you're ready, we'll upgrade you to the highest performance 040 board you can buy.

So you can speed up your software without missing a step.

68030 CPU	*COMPATIBILITY*	*68040 CPU*
CPU-30	*DMA, SCSI, Floppy, Ethernet, Serial I/O*	*CPU-40+ EAGLE I/O Module*
CPU-31	*DMA, SRAM, VSB, Serial I/O*	*CPU-41+ EAGLE I/O Module*
CPU-33	*DMA, DRAM, Serial I/O*	*CPU-40*

Of course, we have all the tools you need to get started. Choose from the broadest range of real-time operating systems and kernels, including PDOS, OS-9, VxWorks, VRTX32 and pSOS+. We even give you VMEPROM, free of charge.

You can also take advantage of XRAY and the entire Microtec family of software tools. Including cross, native and embedded development environments.

Our performance advantage even extends to UNIX.® With the industry's top-rated Unisoft UNIX 5.4.

Finally, you get the industry's best-rated documentation, integration support, regional technical staff and a full one-year warranty.

Here's your next step: call 1-800-BEST-VME ext. 40 for details on our 030 to 040 upgrade offer. Or fax a request to (408) 374-1146 for an immediate response.

Because turning 040 doesn't have to slow you down.

FORCE COMPUTERS®

VME at its best.

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

CIRCLE NO. 5

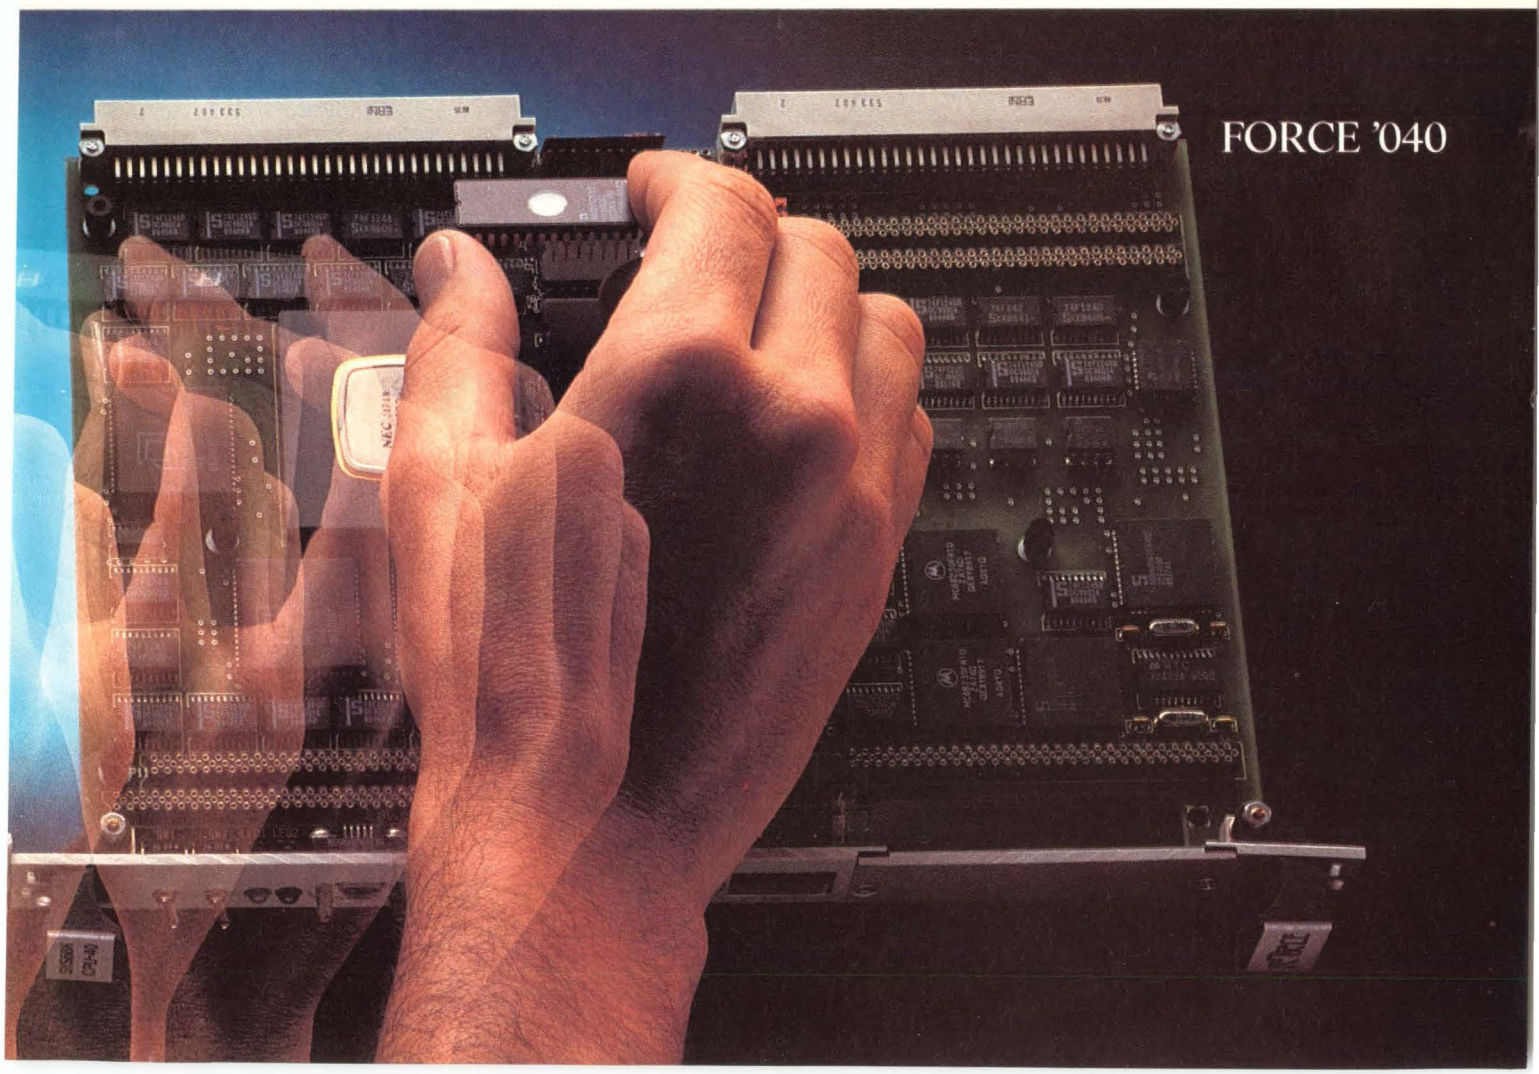

Arbitrator finds Intel breached AMD contract

As *Computer Design* went to press, Judge J. Barton Phelps, arbitrator of the long-standing dispute between Intel and Advanced Micro Devices over rights to the 386 microprocessor, had just handed down a decision. In the decision, he found that AMD failed to provide a number of products to be exchanged for Intel designs under the companies' technology exchange agreement. But more important, the judge found that Intel had, in the matter of the 8087 FPU, committed "a knowing, deliberate and unjustified breach of contract ordered or approved by the most senior officials at Intel." In addition, Phelps ruled that Intel had acted in bad faith on the agreement and had harassed AMD.

The judge declared that AMD would be entitled to recover damages from Intel in the Remedy Module of the arbitration process, which is scheduled to begin November 15. Significantly, the ruling specifically included the possibility that Intel might be ordered to turn over its treasured 386 design as compensation.

A statement from Intel, meanwhile, emphasized that the company did not believe it would have to surrender the chip. Monetary or technology awards to AMD could yet upset the balance of power in the microprocessor industry.

— Ron Wilson

Intel forming net of 960 alliances

In an aggressive attempt to establish a broad base of software and development support for its 960 family of RISC processors, Intel (Santa Clara, CA) is signing up with a crowd of third-party vendors. It announced an alliance with Wind River Systems (Alameda, CA), which will support the 960 with its VxWorks development environments, including the Wind object-oriented real-time kernel and a host of Unix-based development tools.

Cross-development leader Microtec Research (Santa Clara, CA) had been working with Intel early in the game to develop an optimizing C compiler and a version of its X-Ray debugger for the 960, both of which were announced in mid-October. A project to develop a 960 emulator is also under way with Applied Microsystems (Redmond, WA). Introduction of the emulator is expected by mid-1991.

And in what appears to be a coup attempt against Motorola, Intel reached an agreement to port the PSOS+ kernel and related software modules by Software Components Group (San Jose, CA) to the 960. Intel apparently hopes to lure developers used to the Open Real-Time Kernel Definition (Orkid) interface that's pushed by SCG and Motorola. Orkid is essentially the interface to PSOS+, which is used extensively on the 68000 family. By offering the same software interface to its 960 line, Intel seems to want to make it easy for defectors to switch from Motorola to underlying Intel silicon.

— Tom Williams

Mentor strengthens European foothold

European support for Mentor Graphics (Beaverton, OR) was reinforced recently when the Commission of the European Community chose Mentor's VLSI design software for VLSI Design Action, Esprit's Eurochip university program. Mentor, chosen from among a number of EDA vendors, will sell about 500 design systems at a substantial discount to Eurochip member universities throughout Europe. Those 500 seats will bring Mentor's installed university base in Europe to approximately 2,000 seats. Hewlett-Packard (Fort Collins, CO) has been chosen to supply the hardware platforms.

The significance of microelectronics to European high-technology industries prompted the EC Commission to launch the VLSI Design Action program under Esprit, a research program funded by the European Economic Community. The commission has allocated 13.5 million European currency units (about $17 million) to train an additional 3,000 university students in VLSI design each year. The software from Mentor will include schematic capture, simulation, IC layout and documentation.

Mentor is also involved in the European Development Center (EDC), an R&D organization based in Leuven, Belgium. Philips (Eindhoven, The Netherlands) and the Interuniversity MicroElectronic Center research organization (IMEC), also in Leuven, are equity partners with Mentor's Silicon Design Division in the EDC, which develops design tools and technology for European and worldwide markets. Its first project is a DSP design system based on research from Philips Research Labs and IMEC.

— Barbara Tuck

Unisys, Chorus target O/S development

The effort to develop operating systems to meet the demands of distributed multiprocessor systems is getting a shot in the arm from an agreement between Unisys (Blue Bell, PA) and Chorus Systemes, S.A. (Paris, France). The alliance with Chorus' Beaverton, OR, subsidiary aims at a three-year project to develop a high-performance system based on AT&T's Unix System V, Release 4 that will be modular and expandable over large networked systems. A main goal is to make the system compliant with standards such as Posix and X/Open's portability guide.

The Chorus system is based on a microkernel architecture in which a microkernel at each computing node provides a minimal set of operating system services. A set of system servers uses the microkernel to provide full operating system functionality to applications and users. The system's communications capability along with threads—tasks that can run anywhere in the system transparent to the user—make the operating system expandable among multiple-networked processors. Under the agreement, Unisys will have first rights to market systems using the new technology and Chorus will have exclusive rights to license the technology to third parties.

— Tom Williams

Continued on page 10

"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. 6

Continued from page 8

Photochemical process promises terabit storage

Researchers at the University of California (Irvine, CA) are experimenting with a new class of organic compound that could lead to high-capacity, fast-access optical storage systems. The technology calls for storage of data in 3-D polymer cubes, a technique that could store up to 1 trillion bits in a cubic centimeter. Current optical storage technology can store only about 100 million bits in a square centimeter.

The experiments revolve around spirobenzopyran (SP), a compound that changes states when stimulated by two independent beams of light. Embedded in a transparent polymer base, SP molecules absorb ultraviolet light until they are exposed to intersecting light beams, which alter them to absorb visible red-green light. A split beam of laser light is used to write bits to the compound. Reading is accomplished by sensors that pick up red fluorescent light emitted by the SP when it's exposed to ultraviolet beams. *— Mike Donlin*

Silvar-Lisco moving out of the shadow

Silvar-Lisco (Sunnyvale, CA), while having won the utmost respect for its technologically sound place-and-route software, has obviously become discontent with its difficulty in winning customers. A few recent moves demonstrate the company has recognized that technological brilliance must be combined with product promotion to meet with success.

Recently, in addition to naming general managers for central and northern Europe, Silvar-Lisco has chosen as director of U.S. sales Dick Bosenko, who recently held the position of director of Tangent products for Cadence Design Systems (San Jose, CA). As the industry-standard of place-and-route tools, Tangent software is the bull's-eye at which Silvar-Lisco must take aim.

Texas Instruments (Dallas, TX) recently chose Silvar-Lisco's Gards III place-and-route software for its triple-level-metal BiCMOS gate arrays, which is also likely to give Silvar-Lisco a boost. The TI order for Gards III was the result of a benchmark on a BiCMOS gate array with over 200,000 gates, according to Silvar-Lisco. The company claims that neither of two contending vendors completed the benchmark by the final deadline, while it delivered results six to eight weeks before the cutoff. Silvar-Lisco's new strategy will obviously be a winner—of both respect and customers. *— Barbara Tuck*

Pact targets Futurebus+ silicon

A joint agreement by Philips Components-Signetics (Sunnyvale, CA) and Texas Instruments (Dallas, TX) expands the scope and supply of silicon support for the proposed Futurebus+ standard. Under the agreement, the two companies will second-source a broad range of transceivers, controllers and datapath circuits that support the Futurebus+ standard, including high-performance packet mode and cache-coherent transactions.

Both companies will use their own processes to manufacture the products, which will find their way into applications ranging from workstations to supercomputers. Designated FB2000 by Signetics and TFB2000 by TI, the family of devices is designed to satisfy the bus bandwidth and performance requirements of high-speed multiprocessor systems. *— Mike Donlin*

Vitesse and Thomson join in GaAs agreement

Continuing the wave of restructuring that has marked the maturation of the gallium-arsenide IC market, continental specialty-IC power Thomson Composants Microondes (Paris, France) and U.S. GaAs pioneer Vitesse Semiconductor (Camarillo, CA) have announced a second-source agreement that could prove important to both the U.S. and European markets. In exchange for an equity investment in Vitesse, Thomson will hold exclusive European marketing rights to all Vitesse products, and will have the option to source any Vitesse GaAs ICs from a new facility in Grenoble, France.

Added to Thomson's existing position in microwave GaAs components, the new digital products will make an imposing array of high-speed solutions. Vitesse, meanwhile, gets cash—always a useful commodity in the GaAs business. In addition, the U.S. vendor gets a highly credible second source for at least some of its product line, and a powerful marketing ally inside the walls of fortress Europe. *— Ron Wilson*

Motorola plans to expand 88000 family

Since some early leaks in the press regarding its 88000 family of RISC processors (see News Briefs, June 1), Motorola (Tempe, AZ) has outlined some of its goals for the future, being careful to disclaim any similarity to a product announcement. According to Motorola 88000 product manager Jeff Nutt, the disclosure was made to show general future plans.

The first of the family to follow the 88100 and 88200 will be the 88110, a fully integrated unit combining CPU, MMU floating-point unit and cache in a single chip. In addition to combining the five major functions on a single chip, the part is expected to have 64-bit external and 80-bit-wide internal data paths, extended precision floating-point arithmetic, branch accelerations and bus snooping with separate tags. It will be available with low-cost package options.

The chip, which is expected to be fabricated in Motorola's 0.8-μm process, will count less than 1.5 million transistors, but will boast between three and five times the performance of the previous generation, according to Nutt. In addition, the company hopes to leverage its 88000 technology by offering a highly integrated version of the chip for embedded applications—not unlike some recent approaches with Sparc. Nutt offered no timetable for the release of any of the discussed technology, but sources said samples of the 88110 could be available as soon as spring of 1991. *— Warren Andrews*

Unparalleled, parallel performance.

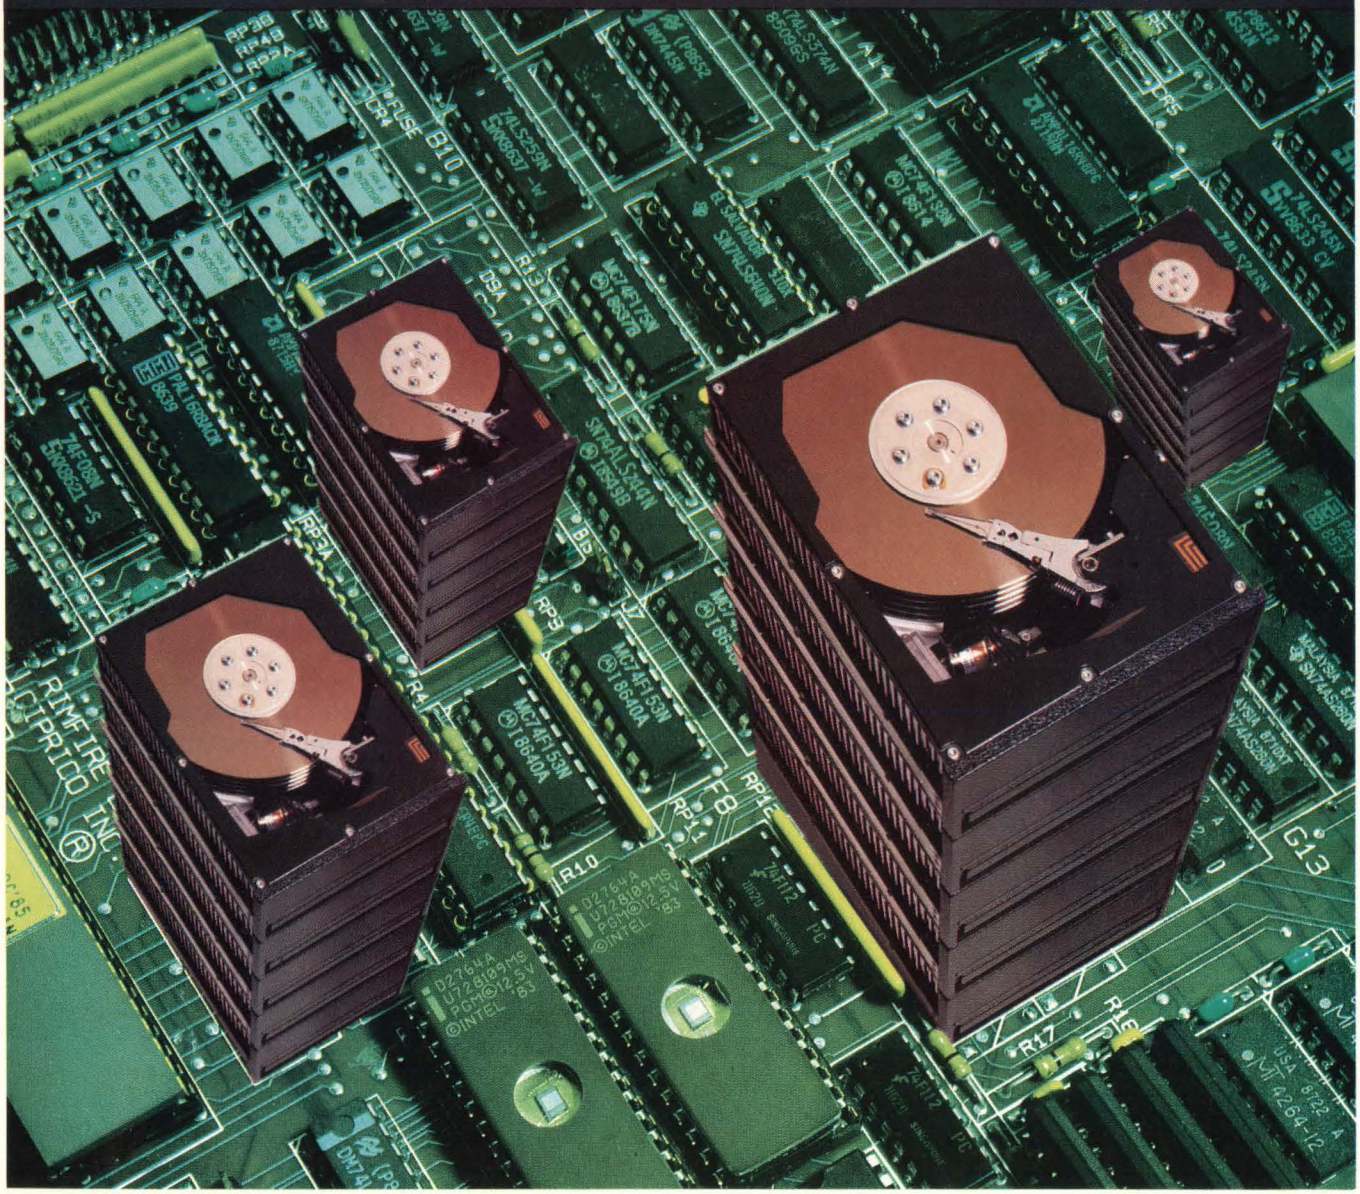

SCSI-2 is at your command with Rimfire® 6600 Series Parallel Disk Array Controllers.

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.

CIRCLE NO. 7

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

Protectionism is not a sin

The issue isn't simply protectionism but what we choose to protect and how.

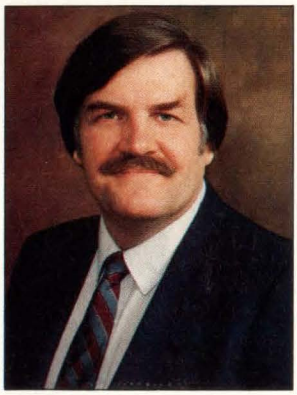

John C. Miklosz
Associate Publisher/
Editor-in-Chief

There's no question that protectionism is a sin for free-market, free-trade zealots, but it's no sin for the rest of us. Individuals practice protectionism in their personal lives, corporate executives practice protectionism to guard their companies and themselves, and governments practice protectionism in a host of ways to ensure the security and well-being of their citizens.

As individuals, we lock our doors at night and pay for police protection against robbery and rape. We set aside a portion of our incomes for accident insurance, health insurance, disability insurance and life insurance to protect ourselves and our families from the caprices of life. If we're wise, we also put something away to protect us from hard times in the event that a glitch in the free-market system eliminates our jobs. And if we're really wise, we keep our skills honed and up-to-date—maybe even develop alternative skills—to protect us from the obsolescence that free-market economies seem to bring.

Corporate executives—usually staunch advocates of free-market principles—see nothing wrong with trying to protect their companies from takeovers—mounted by equally determined free-market advocates. If stock manipulations fail to do the trick—God forbid they should try to build a more-competitive company—they're ready to protect themselves from hard landings with golden parachutes that will inflict financial punishment on the raiders.

Governments, too, protect their citizens in many ways from the vicissitudes of a free market. We have unemployment insurance for those who lose their jobs and a host of welfare programs for those who aren't employable. Even in this time of heightened deregulation, we'll see little diminution in the volumes of tax laws, banking laws, business laws, farm laws, zoning laws, licensing laws, conservation laws—the list goes on—that in one way or another provide protection for a segment of the population within the construct of a "free-market" economy.

Why then, must we consider free markets and free trade as gospel, and protectionism as sinful, when it comes to international commerce? Just as there are few individuals or nations who believe in unfettered free markets and free trade, there are equally few who don't believe in some form of protectionism. The issue isn't simply protectionism but what we choose to protect and how. Our dealings with trading partners shouldn't be governed by blind adherence to a free-market, nonprotectionist cult doctrine but by all that we really believe in and practice.

"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

CIRCLE NO. 8

Take Advantage Of Board-Level Partner

88open COMPATIBILITY CERTIFIED

Malcolm Baldrige National Quality Award

All That Our VME ship Has To Offer.

Once you've seen what Motorola brings to the table, we think you'll agree it's everything you need. Like the most complete line of VME products, services and engineering support available anywhere. Award-winning quality. Competitive pricing. All from the company that pioneered VME technology, and whose product line ranges from ICs to boards to full systems. And includes everything in between.

You decide exactly what you need from our more than 100 VME products at every level of price and performance. From CPU boards, like our new '040-based MVME165, to memory boards, to communications boards. And the industry's widest assortment of development tools, software resources, and technical support.

A partnership with Motorola not only helps you control costs, but even more importantly, speeds your time to market.

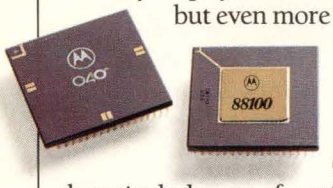

Our products include more functionality with a higher level of integration to accelerate your development efforts. And because of Six Sigma quality control you can be assured that our products will work right out of the box. It all adds up to the fact that getting you to market sooner is a promise only a company with the resources of Motorola can make good on.

Every Motorola product includes a built-in migration path, so your future product cycles are assured. Such as providing a way to upgrade from the 68020 to the 68030 to the 68040, or from a 68000 CISC board to an 88000 RISC board with a simple re-compile. Wherever you're headed, Motorola is going to be there.

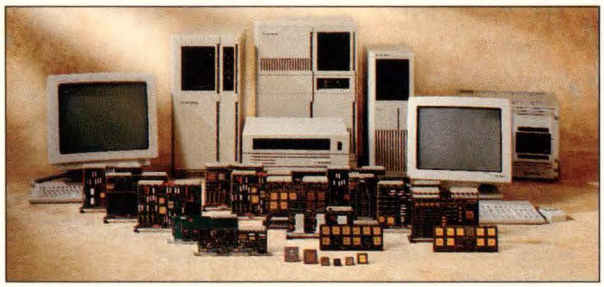

And we'll support you during the entire development process. Every board in our product line includes a full suite of board diagnostics available in both a run-time and a source package. This degree of flexibility also extends to our nationwide customer service programs, which run the gamut from total on-site maintenance to self-maintenance support packages.

For more information, call us today at **1-800-624-8999**, ext. 230.

Once you discover the advantages of a partnership with Motorola, you'll see that it's no ordinary board-level decision. It could ensure the future of your company.

At Motorola, Openness Is Standard Procedure

These days, openness has become something of a buzz word, and everybody seems to have a different idea on what is and isn't "open." To us, it's no big mystery. Openness means open architectures, open software, open networking. And open standards like UNIX,® as indicated by our role in founding 88open. It means we're open to helping solve problems with your point of view in mind, not just ours. And it's been that way ever since we helped introduce VME back in 1982.

That's why Motorola is committed to supporting official and de facto industry standards, interoperable computing between multiple vendors, and non-proprietary open system architectures. It's why we created VMEexec:™ to facilitate the interoperability of different real-time software modules within a common UNIX environment. And it's why we support virtually every networking protocol, including XNS, TCP/IP, DECnet,™ MAP/TOP/OSI, SNA, BSC, X.400, and X.25.

This philosophy of openness is the same reason we offer as many VME boards, products and services as we do. It's to our mutual benefit, and after all, isn't that what partnerships are for?

MOTOROLA

Computer Group

CIRCLE NO. 9

It Takes One To Know 1μm.

And, Mitsubishi is the only one who reduced 1μm gate width to a size that's small enough to provide a low, 5μW power dissipation (the lowest you can get), lower operating temperature and small die size. Plus, we're the only supplier to give you performance as fast as 370 picoseconds and four speed/power options: x1, x2, x3, and x4. So, you can optimize your circuit for *both* fast speed *and* low power. By taking advantage of Mitsubishi's smaller cell sizes, you achieve low operating temperature, more-cost effective silicon use

We're The Only One With 5μW And Four Speed/Power Options.

and smaller die size. And, all this means we can help you put your designs in smaller, lower-cost plastic packages and achieve higher reliability. Plus, our selectable output buffer drive and slew rate control help reduce noise and electromagnetic radiation. Mitsubishi offers 70,000 gates (50% utilization); patented* gate isolation structure for faster performance, and extensive packaging options, including plastic QFPs, very-fine-pitch QFPs (VQFPs), PLCCs and APGAs (adaptable pin-grid-arrays). And, of course, we give you 100% burn-in, 100% AC and DC testing, a library of more than 500 cells and design kits for the industry's leading-edge CAE design tools. It takes one solid, proven semiconductor supplier to know 1μm. We're one of the top ten worldwide semiconductor manufacturers, and the world's fifth largest MOS memory manufacturer. We've been producing gate arrays since 1976 and offer production capacities surpassed by none. So, we'll get you through 1μm design, prototype and production *fast* and *right the first time*. Also, keep in mind, a good 1μm will get you to a great 0.8μm! Call the one it takes to know 1μm today at **1-800-624-8999 #178.**

*U.S. Patent No. 4,562,453

MITSUBISHI GATE ARRAYS

CIRCLE NO. 10

CALENDAR

CONFERENCES

November 6-9
Software Development '90-Fall

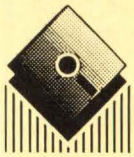

Omni Parker House, Boston, MA. The conference will offer an exhibition as well as over 60 lectures and workshops in object-oriented programming and design methodologies. A one-day seminar will focus on how to build, organize and manage effective project teams. Information: Software Development Seminars, 500 Howard St, San Francisco, CA 94105, (415) 995-2472. **Circle 361**

November 6-10
Electronica '90

Munich Trade Fair Centre, Munich, West Germany. The International Trade Fair for Components and Assemblies in Electronics will feature exhibits in three product groups: electronic components, electromechanical products, and quality-assurance and development equipment. Related events include a conference on microelectronic sensors and another on the use of semiconductor technology in switching and regulating applications. Information: Kallman Associates, 5 Maple Ct, Ridgewood, NJ 07450-4431, (201) 652-7070. **Circle 362**

November 12-16
Comdex/Fall '90

Las Vegas, NV, nine locations. The 12th Comdex conference and trade show will feature more than 1,800 exhibitors. It is targeted at computer distribution professionals, including volume resellers of small computers, peripherals, software and accessories. Ten tracks will encompass over 48 sessions on topics including portable computing, imaging systems, buyer/seller issues, communications, multimedia and industry perspectives. Information: The Interface Group, 300 First Ave, Needham, MA 02194, (617) 449-6600. **Circle 363**

November 13-15
Wescon

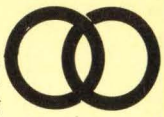

Anaheim Convention Center, Anaheim, CA. This year's conference will offer 36 technical sessions, 22 tutorials and panels discussing world engineering issues, professional development and technology initiatives. It will feature more than 1,000 exhibits, including a demonstration of the state of the art in PC-based design tools. Information: Electronic Conventions Management, 8110 Airport Blvd, Los Angeles, CA 90045, (800) 877-2668. **Circle 364**

January 14-17
ATE & Instrumentation Conference

ATE & Instrumentation Conference

Disneyland Hotel, Anaheim, CA. This conference for test professionals will offer almost 50 sessions in design, manufacturing, management, service test and systems integration. Session titles include "Testing in the 21st Century" and "Testing—the Competitive Edge." A four-day educational tutorial will follow the basics of design for test through to an actual production example of an integrated design and test board. Information: Miller Freeman Expositions, 1050 Commonwealth Ave, Boston, MA 02215-1135, (800) 223-7126. **Circle 365**

COMPUTER DESIGN Technology and Design Directions®

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

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

SENIOR EDITORS
Ron Wilson, *Integrated Circuits* (503) 284-3432
Warren Andrews, *Computers & Subsystems* (201) 825-1317
Barbara Tuck, *ASICs* (516) 744-2984
Michael G. Donlin, *Design & Development Tools* (508) 392-2123
Tom Williams, *Software* (408) 335-5514

SECTION EDITORS
Ron Wilson, *Technology Updates* (503) 284-3432
Jeffrey Child, *New Product Highlights* (508) 392-2126

ACTING MANAGING EDITOR: Kathleen A. Gow

ASST. MNG. EDITOR: Barbara-Ann Scofidio

ASSISTANT EDITORS:
Anne Holm, Michelle Keyworth

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: Steve Fuller

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. Tobek (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

ECS: The Major Player in the Ultimate Board Game

500,000 motherboards shipped in the past two years

In the big games there are lots of players, but only a few winners. And ECS has earned worldwide success as a PC motherboard manufacturer by being responsive, flexible, reliable.

The company's design skills and manufacturing technology provide fast turnaround, customization when required, and superior price/performance. Motherboards are built with SMT, and other advanced production and testing techniques.

In sum, ECS values long-term partnerships over short-term profits. And by offering a complete line of 286, 386SX, 386DX, and 486 motherboards, ECS has shipped half a million boards to its partners in the past two years. Partnerships which include major computer companies throughout the world.

So, count on ECS to provide what you need when you need it, plus prompt, dependable service.

Meet ECS at "Comdex", Booth #448 in the West Hall.

ECS ELITEGROUP

"We engineer success"

ELITEGROUP COMPUTER SYSTEMS U.S. Headquarters 365 Ravendale Dr.; Mountain View, CA 94043 Tel: (415) 969-1000 Fax: (415) 969-0343 **Subsidiaries—Los Angeles:** 12020 Mora Dr., Bldg. 8, Ste. 9; Santa Fe Springs, CA 90670 Tel: (213) 944-2881 Fax: (213) 944-4970 **New Jersey** 200 Centennial Ave., Ste. 205; Piscataway, NJ 08854 Tel: (201) 457-9500 Fax: (201) 457-8811

100% full-parallel flash ADC performance for 10% of the price.

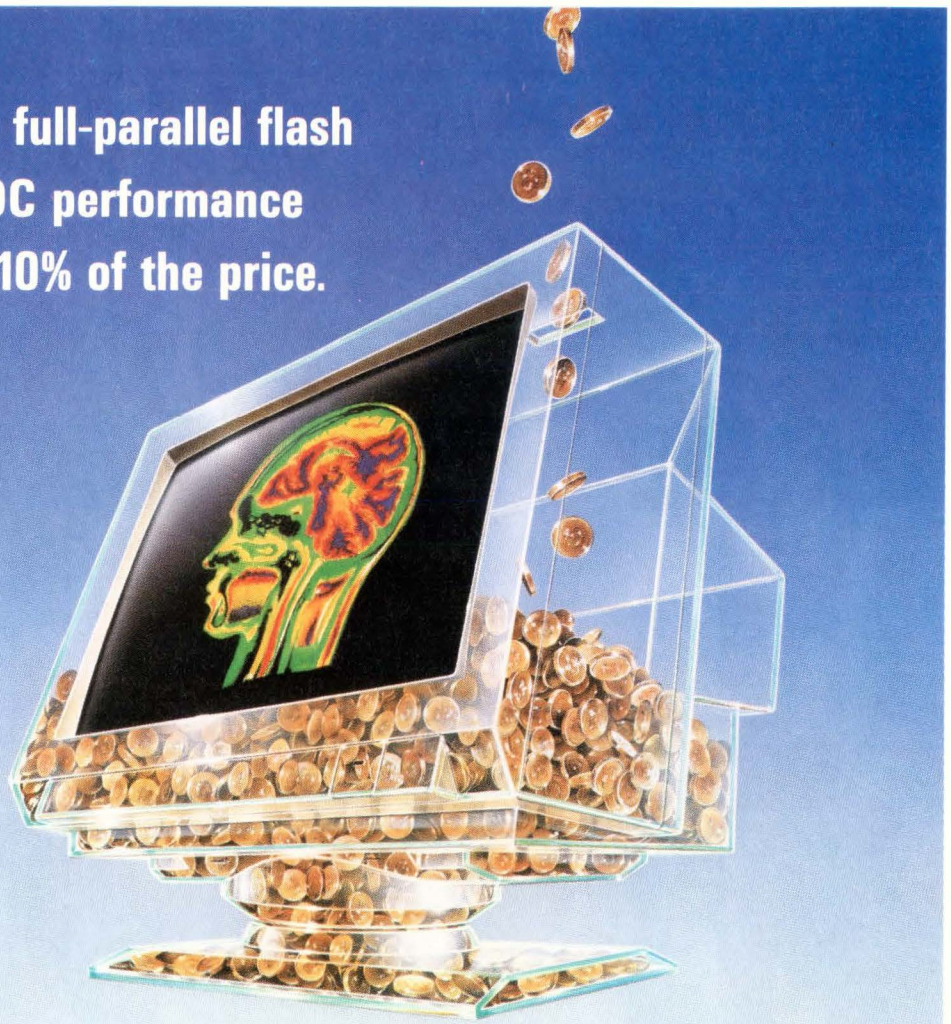

Save money and get an extra effective bit every time you convert video from A to D...with the new Philips 8-bit video ADCs. Giving full-parallel flash performance at up to 1/10 the cost of full-parallel flash converters, these ADCs provide professional performance at a consumer price.

How? Our new folding-and-interpolation technique slashes on-chip component count, thus reducing cost while enhancing performance. It also lowers power consumption up to 70%. For example, the TDA8713 consumes as little as 290mW and gives 7.5 effective bits with an input frequency of 4.43MHz (40MHz clock)...ideal for a wide range of professional and consumer video applications.

Type	MSPS	Diss.(mW)	Features
TDA 8703°	..40	290	Voltage Reg.
TDA 8713°	50	290	TTL I/O
TDA 8715°	50	325	ECL I/O
TDE 8715	50	325	Mil. Temp. Range
TDA 8708°	30	365	Clamp/AGC CVBS Signal
TDA 8709°	30	380	Clamp/Adj. Gain R,G,B,Y,U,V, C Signals

° in DIL and SO versions.

For more information, contact Philips Components, Marketing Communications Dept., Building BAF-1, P.O. Box 218, 5600 MD Eindhoven, The Netherlands. Telex: 35000phtc nl/nl je vmc. FAX:31-40-724825.

Philips Components

PHILIPS

CIRCLE NO. 15

PHILIPS

James Ready on: Real-time standards

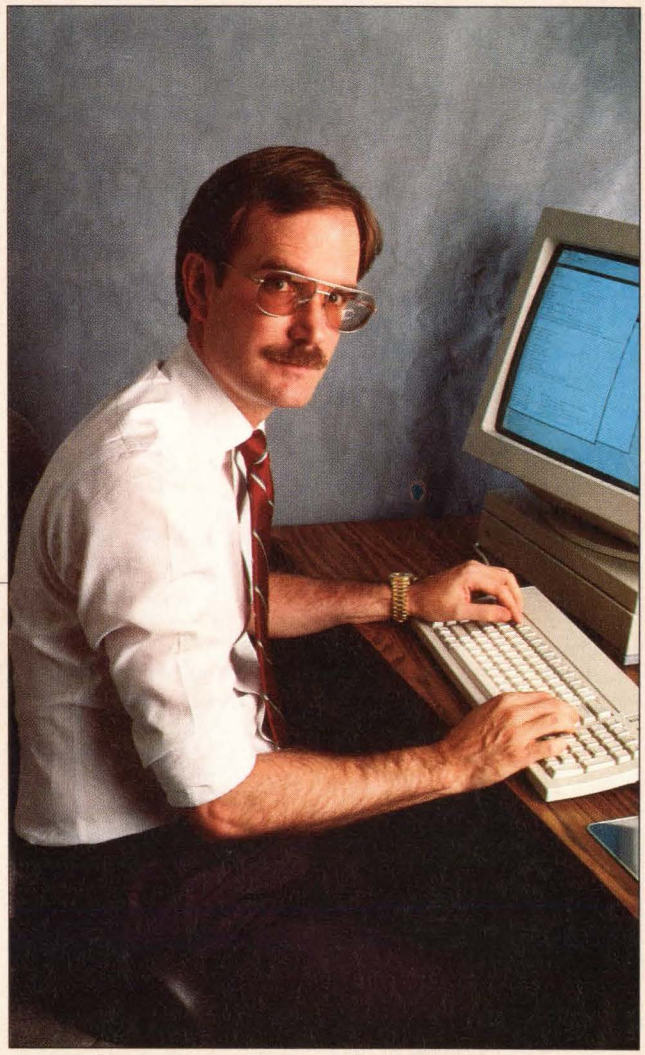

Some important work will unfold over the next several years to define and implement standards for real-time computing. Developers of real-time systems will benefit greatly if standards evolve for a reference model of layered interfaces, as well as for a standard set of application-oriented workloads to evaluate the suitability of a real-time system for a particular application.

The real-time system community lacks a reference model that can be used to frame discussions and structure the standards process. Reference models have proven quite uscful, notably in the area of communications. Here, I present a proposed real-time computing reference model that consists of seven layers.

The first layer represents the physical or mechanical aspects of the real-time computing system. Each of the widely used industry-standard buses, for example, such as VMEbus, Multibus II and STD Bus, has detailed physical specifications for board height, width, pin spacings, connector characteristics and so on. For more-rugged applications, these standards might also include specifications for temperature, vibration and humidity operating ranges.

Second is the hardware layer, which has two sublayers—a low-level signal layer and a high-level programmable layer. The low-level signal layer specifies the bus signal characteristics, including voltages, signal timings and bus arbitration policies.

This layer is largely invisible to software. The high-level programmable layer represents the hardware that's visible to the programmer. Such hardware generally includes the CPU, timers, message passing chips, serial I/O devices and any other programmable hardware.

The third layer represents basic services that raise the software interface of the hardware to a more abstract level. It has three major parts: a BIOS, a program loading (booting) interface and a low-level debugging interface. All services at this layer are independent of programming languages and operating systems, though adding bindings to make these services available from languages and various operating systems is usually a simple matter.

If a basic services layer is made available on a variety of hardware platforms, it will be simple for the upper layers of software to move across compliant systems.

The BIOS defines interfaces for abstract devices such as timers, counters, serial lines, network controllers, disks and SCSI controllers. This abstract interface exposes all of the functionality that's commonly expected for a given type of device. Programmers, for example, are presented with one interface to read, write and set baud rate and parity for all serial line devices.

The program-loading interface provides a standard means of loading software for execution from various types of devices including disks, tapes, serial lines, backplane devices, and shared memory and network devices. Standards are used wherever possible, both for the format of the software and the communications protocols used to

transfer the software. Booting from a network, for example, may involve using the TCP/IP (Transmission Control Protocol/Internet Protocol), the Trivial File Transfer Protocol (TFTP) and the standard S record format to load an executable image into memory.

The low-level debugging interface provides a device-independent, operating-system-independent means of controlling the execution of the system being programmed. It's properly viewed as a debugging kernel that provides read/write memory capability and start-of-execution capability. This interface is meant to be used by higher-level services and would most likely be hidden from the application programmer.

The kernel layer

The fourth layer is the kernel layer, which represents the interface for lightweight concurrency. Ready Systems' VRTX32 refers to this interface as tasks, while Mach, OS/2 and the Posix standard refer to it as threads. The Ada language tasking model is also at the same level of interface.

This layer is characterized by a program with multiple threads of execution, with all threads sharing the same address space and I/O environment (such as open files). The context of each thread is little more than the basic machine state, with just enough extra state to maintain the thread environment correctly.

Fifth is the process layer, which corresponds to the multiprogramming level of concurrency found in Unix, VMS and other time-sharing operating systems. Here the context is quite heavy and includes not only the basic machine state of registers but also the state of MMUs, floating-point units, I/O and other operating-system-dependent information on the that must be saved during context switching.

The sixth layer is the virtual machine layer, which builds upon the previous layers to present a model of computing to the application programmer that's closer to the application at hand. Building an application that uses a pipeline "virtual machine," for example, might be well matched to an image-processing application. An off-the-shelf pipeline virtual machine would make building the application much simpler, since the application programmer could concentrate on the unique properties of the application instead of on the implementation details.

Concentrating on the application

The seventh layer is where the application-specific programming gets done. The lower layers of the model are available to the application, which is important for cases in which the higher performance of the lower-level services is more important than increased functionality at a higher layer.

Several groups are making progress on a number of these layers. The IEEE Posix Working Group has focused on system interfaces for process control, directory and file management, and I/O. This standard is primarily concerned with the process layer. Newer Posix activities have included some proposed specifications for real-time extensions, although still at the process layer. The Posix.4a group is aiming at specifying the thread layer.

Another activity to develop standard interfaces for various parts of the real-time reference model is EPIS (Executive Processor Interface Specification), developed by the Embedded Systems Committee of the 88Open consortium. EPIS defines a low-level interface between hardware and operating systems or stand-alone programs. It fits at the basic services layer and represents a subset of a full BIOS interface. EPIS is concerned only with the hardware devices closest to the processor and doesn't have a general device-driver interface.

Another effort in this area is ITRON, an application interface to a real-time kernel, which would fit at the thread layer. ITRON was designed and promoted by Ken Sakamura of the University of Tokyo. The ITRON interface is now managed by the TRON Association, a worldwide group with members from industry and the academic community who are dedicated to supporting the TRON architecture.

Focus on basic services

While all the layers are important, it's critical to start with a firm foundation. The most beneficial layer to standardize today is the basic services layer. The reason is simple: if a basic services layer is made available on a wide variety of hardware platforms, it will be simple for the upper layers of software to be moved across compliant systems. The success of the PC BIOS in providing this kind of foundation is the best example of how such a standard can be used.

Companies building proprietary hardware are forced to build a basic services layer themselves, with no particular specification to follow.

The situation today is anarchic; lack of low-level standards forces each operating system developer to build its own set of services for each CPU board supported. Moreover, companies building proprietary hardware are forced to build a basic services layer themselves, with no particular specification to follow.

To help jump start the process of developing specifications for this layer, Ready Systems developed a "strawman" specification, which has been released to the industry. During development, it was important that Ready Systems avoid the temptation to specify features that might give itself a potential advantage. The specification, for example, might have called for an interface to an operating system primitive that looked similar to one found in VRTX32. If other operating system vendors perceived an unfair advantage to one player, they would reject the proposal outright. This happened to the ill-fated Open Real-Time Kernel Definition (Orkid), proposed as a potential standard by the VFEA International Trade Association. Most real-time operating system vendors quickly spotted a competitor's product in the Orkid specification and rejected it.

The OBIOS (Open Basic I/O System) specification

to standardize the basic services layer is in draft form and is being reviewed by a number of hardware and software manufacturers.

It's very difficult to judge whether a component of a real-time system, such as a specific-speed CPU or an operating system, can successfully support the performance demands of a given application. The selection process for a real-time operating system, for example, typically consists of a checklist comparing operating system functionality on a feature-by-feature basis, and perhaps a comparison of context-switching and message-passing times. As might be expected, there are significant problems with such simple comparisons.

Using workloads to simulate real-time systems

To see how an oversimplified comparison can lead to an incorrect conclusion, consider the following example: operating system A takes 25 μs to send a message; operating system B takes 75 μs to send the same message. If a choice is made based simply on this timing comparison, operating system A would seem to be the correct choice since it performs the operation three times faster. What wasn't documented, however, was that operating system A had side effects from other system calls that could result in worst-case delays of several milliseconds. This demonstrates that what's important is the behavior of the operating system under realistic conditions where all possible side effects and interactions are taken into account.

One way to make meaningful comparisons is to use standard benchmarks. For real-time systems, these benchmarks should take the form of loads on the system that simulate the behavior of real applications. This kind of benchmark is often called a workload because it tries to simulate the computing load and behavior of a real-time system.

From the point of view of standards development, the workloads can be divided into two broad categories: hard real time, where the results must always be delivered on time; and soft real time, where results can be delivered under less-stringent deadlines. One finds most hard real-time requirements in control applications where the system performance is measured in strict accordance to deadlines. Many soft real-time requirements are found in communications applications, where the measure of systems performance is more throughput-oriented.

By applying workloads, a system builder can select both hardware and software much more intelligently early in the design process. The alternative is to wait until the end of the development process to find out that either the hardware or the software didn't stand up to the applications performance demands.

Making progress toward general workloads

Various groups are making significant progress in the generation of useful workloads. At the seventh IEEE Workshop on Real-Time Operating Systems and Software in May 1990, Daniel Kiskis and Kang Shen of the University of Michigan gave a paper describing a synthetic workload generator that could produce workloads typical of real-time applications. Another presentation illustrated a workload typical of sonar signal-processing applications.

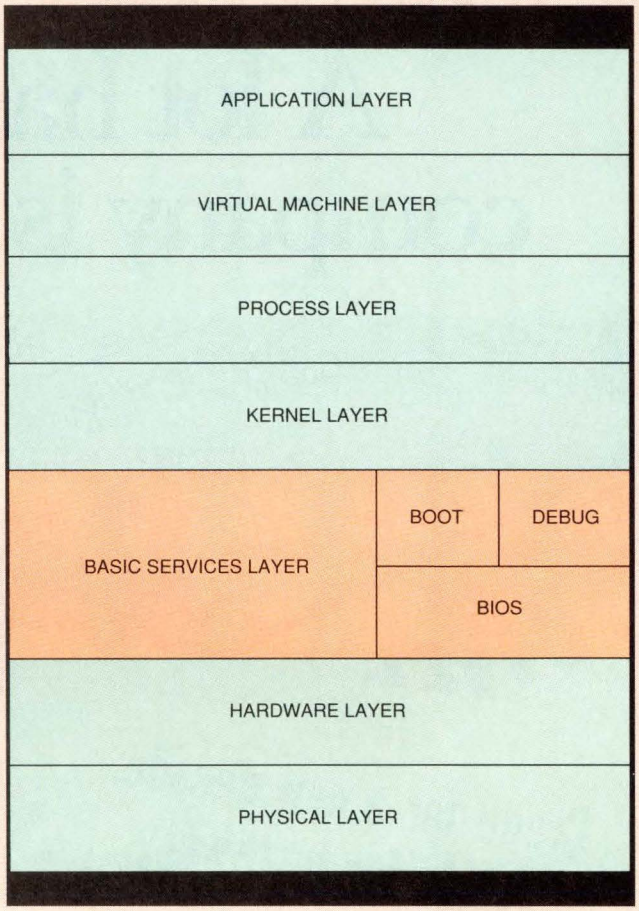

The real-time system community will benefit from a reference model such as this that can be used to frame discussions and structure the standards process.

More-generalized workloads have also been recently developed. The Rhealstone and Hartstone efforts, for example, are benchmarks designed to measure general properties of real-time systems. Recently, a Distributed Hartstone has been implemented in order to measure the behavior of real-time distributed systems.

There are a number of ways to develop standards for workloads. One possibility is that the IEEE Posix efforts for real time could include a standard set of workloads to validate that Posix compliance isn't just functional—for instance, a system meets the Posix real-time call interface, but it also behaves correctly according to the behavior that the Posix interface implies.

Various industry applications groups could also develop their own sets of workloads. The Department of Defense, for instance, could produce an avionics workload and a signal-processing workload. In all cases, it would be very useful for system developers to subject potential vendors' hardware, languages and operating systems to these workloads to see how these systems behave under realistic conditions.

James F. Ready is executive vice-president of Ready Systems (Sunnyvale, CA).

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.

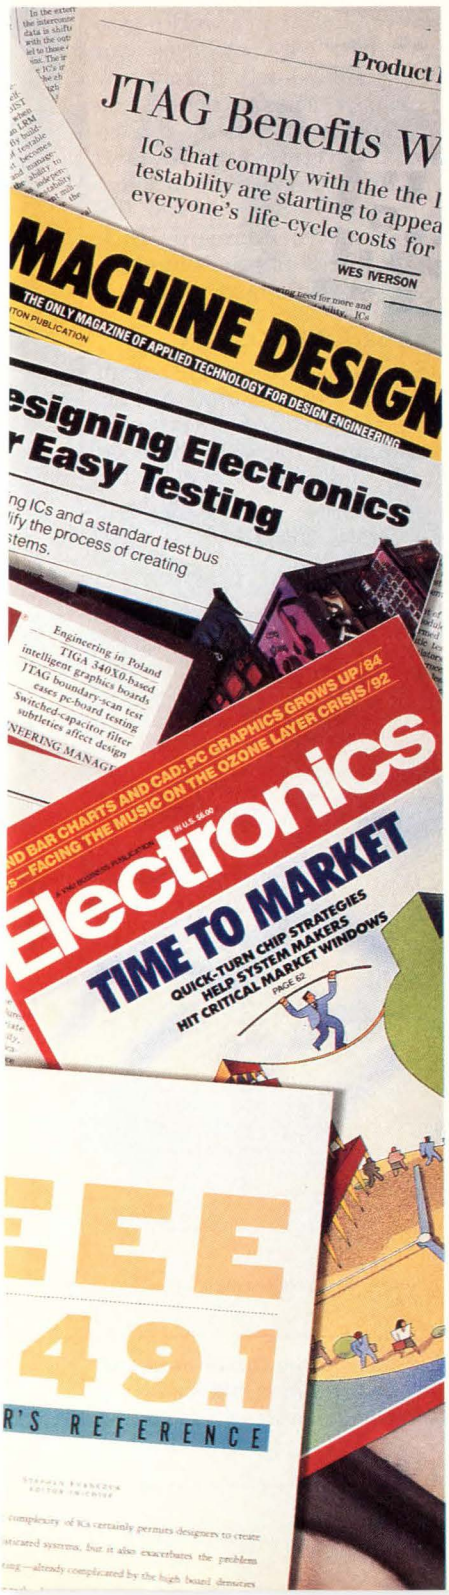

You've seen the headlines and read the stories. Design-for-test (DFT) is a challenge but one that's now easier to live with. The reason: Texas Instruments is the first to develop products for implementing the JTAG/IEEE 1149.1 testability standard quickly and effectively.

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.

TEXAS INSTRUMENTS

CIRCLE NO. 16

INTEGRATED CIRCUITS

European consumer products replete with mixed-signal, DSP technology

Barbara Tuck, Senior Editor

Unlike in the United States where a healthy computer market is driving the development of ASICs and application-specific standard products (ASSPs), IC technology in Europe is being driven by a flourishing consumer market that encompasses telecommunications and personal communications products. Though CMOS is still the workhorse technology in Europe, silicon vendors are increasingly looking to BiCMOS as an answer to the proliferation of mixed-signal designs. And as European countries unify behind standards, European semiconductor makers and U.S.-based companies with a presence in Europe are cooperating to implement products that conform to those standards.

Meanwhile, EDA tool vendors are addressing software problems that designers could meet with the mixed-signal technology and digital signal processing techniques employed in the slew of European communications products.

Such heavy reliance on technologies like DSP and high-performance analog/digital has created the need for a new generation of board-level and system-level tools, claims Graham Symonds, director of corporate marketing at Racal-Redac (Tewkesbury, England). Racal-Redac is developing tools designed to fully exploit technologies such as high-performance analog/digital and DSP and to make sure they can be properly integrated within systems, according to Symonds.

DSP gets top priority

The significance of DSP technology to European products has prompted the European Development Center (Leuven, Belgium) to develop as its first project a DSP design system that implements technology that had been funded by Esprit. (Esprit is a project involving the joint efforts of European countries.) Equity partners in the European Development Center (EDC) include, among others, the Silicon Design Division of Mentor Graphics (Beaverton, OR), Philips (Eindhoven, the Netherlands), and the Interuniversity MicroElectronic Center (Leuven, Belgium)—a research institute concerned with microelectronics process and design technology.

EDC general manager Herman Beke says that the workstation-based DSP design system, the DSP Station, encompasses the specification, simulation and optimization of the algorithm; emulation of the system; automatic synthesis of assembly code; and automatic creation of an IC through silicon compilation. The DSP Station, in mid-project now, is financed by Philips and executed by the EDC. It will be distributed in Europe and other markets throughout the world by Mentor Graphics and will be integrated within Mentor's Falcon Framework.

Meeting the DSP requirements of European applications is also a top priority at the digital business unit of SGS-Thomson Microelectronics' Semicustom Product Division (Agrate, Italy). Product marketing manager Christos Lagomichos reports that engineers are working to integrate a DSP core into a 0.7-µm BiCMOS standard-cell library for use in European cellular phones. So far, SGS-Thomson has worked with key customers to integrate the DSP core into 1.2-µm standard-cell designs. Telecom customers have driven the development of the submicron BiCMOS standard-cell family, slated for production sometime next year.

Vendors back BiCMOS

BiCMOS will be the mainstream technology for mixed-signal ASICs, according to Philippe Lambinet, business manager for SGS-Thomson analog cells and arrays. The analog

NICAM TV STEREO SOUND DECODING SYSTEM

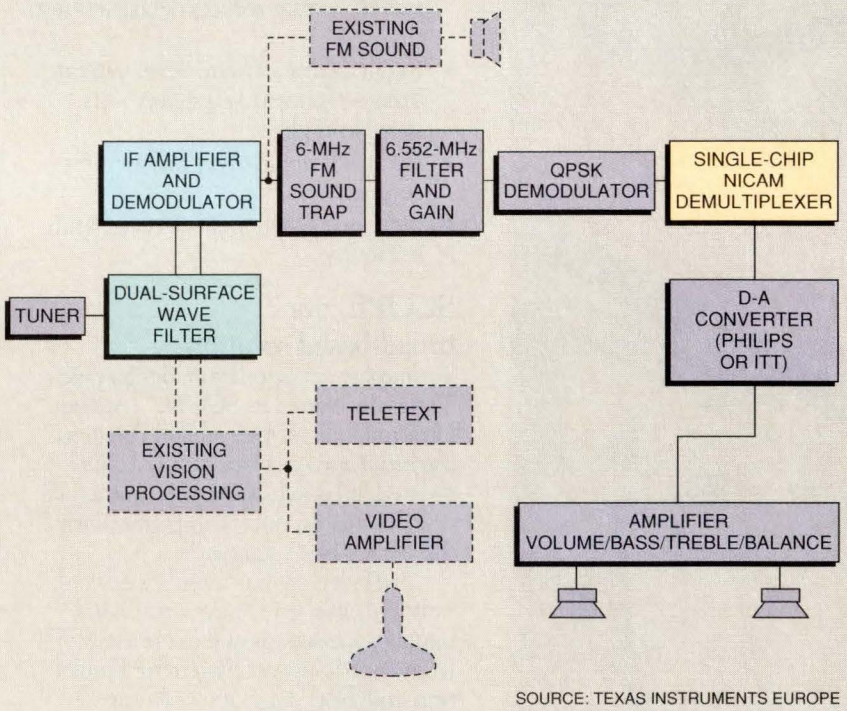

Among the ICs in a typical TV stereo decoding system that meets the European Nicam standard for digital stereo of CD quality are a Texas Instruments Nicam 728 demultiplexer, a Philips TDA2545A IF amplifier and demodulator, and a Plessey SW166 dual-surface wave filter.

We wrote the books on real-time.

When it comes to real-time system software, Microware Systems Corporation leads the way. Now, we offer two books that provide you with complete real-time solutions — the *OS-9* and *OS-9000 Catalogs*.

Microware offers the broadest range of real-time solutions in the industry. Data sheets and flyers just weren't enough to tell our story, so we wrote these books to tell you about our real-time operating systems. These catalogs are chock-full of more than 330 pages of useful information about our operating systems including networking, graphics, language compilers and productivity tools.

OS-9 is the world's leading real-time operating system for 680X0-based systems. OS-9 has been designed into thousands of embedded applications from industrial automation to consumer electronics.

OS-9000 is a portable real-time operating system written in C for advanced CISC and RISC processors. The ability to port OS-9000 to various hardware platforms "future-proofs" your real-time software investment through the 1990s.

Put the world's real-time leader to work for you today.

For Your Free Catalogs, Call

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

MORE CHOICES • MORE OPTIONS • TOTAL SUPPORT

CIRCLE NO. 17

INTEGRATED CIRCUITS

and mixed-signal business unit of SGS-Thomson's Semicustom Product Division (Grenoble, France) is offering customers a 2-μm, BiCMOS standard-cell library with 6-GHz NPN transistors for designing ASICs that can have an analog content of 10 to 90 percent.

The BiCMOS process, born as a bipolar process, is designed to accommodate high-performance analog circuitry, says Lambinet. For block-type design, the standard-cell library includes macrocell generators for PLAs, SRAMs, ROMs and—starting in first-quarter 1991—EEPROMs.

Though the Semicustom IC Unit of the Siemens Semiconductor Group (Munich, Germany) doesn't yet offer BiCMOS libraries, senior director of IC sales and marketing Hans-Peter Bette claims that experience in ECL gate arrays and state-of-the-art CMOS give Siemens the expertise required for a BiCMOS process. Bette says that Siemens will do a BiCMOS product announcement when market needs are more clearly defined. Siemens directs mixed analog/digital telecom designs, aimed at a specific application but multiple users, to a Siemens ASSP unit where IC and system design experts interact closely with customers.

The focus of engineers at the European ASSP unit of Analog Devices (Limerick, Ireland) is the integration of the analog acquisition and baseband conversion portions of a digital mobile radio that conforms to the European Standard of GSM (Groupe Speciale Mobile). By July 1991, Europeans are expected to have made the switch from bandwidth-limited analog cellular radios that aren't standardized from country to country to digital cellular radios conforming to the GSM's European standard.

Analog Devices' next step in the GSM digital mobile radio project—integrating the DSP portion with the combined acquisition and conversion circuitry—will require BiCMOS, says marketing manager Kevin Styles. Analog Devices has a fully integrated 2-μm BiCMOS process in production and is moving toward a higher-performance BiCMOS process. "Meanwhile, what we miss in performance, we make up for in functionality," he says. "Our strength is maximizing the performance of technology in the analog domain."

DESIGN SPECIFICATION FOR DSP STATION

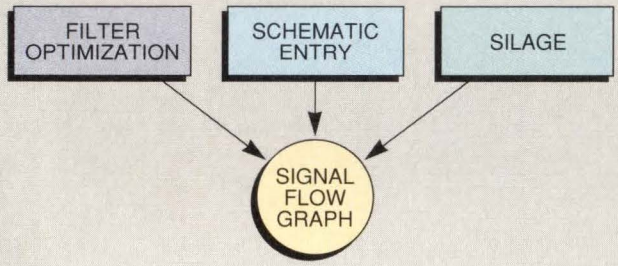

DSP Station, a workstation-based DSP design system, is based on a signal-flow graph representation of an algorithm. Users specify the algorithm graphically by using a schematic editor or textually by using a special algorithmic language, Silage.

Standardization efforts

Also working on the GSM digital mobile radio is Texas Instruments' European ASIC Product Group (Bedford, England). TI is addressing the basic voice-processing functions required for the pan-European cellular radio. Implemented in TI's 1-μm Epic CMOS process, a standard-cell GSM digital voice processor contains a DSP core optimized to execute the voice algorithms defined by the GSM specifications.

European ASIC product group manager Clive Hoggar reports that one-fifth of TI's business in Europe is consumer-oriented, with ASICs going into cellular phones and radios, as well as TV applications such as teletext and stereo sound. To meet the growing demand from customers for a European silicon source, TI is building a fabrication facility outside of Rome in Avezzano, Italy, and is also sourcing silicon in Freising, Ger-

Will FPGAs be hot in Europe?

Couple the fact that Europe, at any given time, is about 12 months behind the United States in ASIC density with the conservative spending habits of Europeans, and it would appear that the European market is ripe for FPGAs. According to Mike Inglis, manager for Texas Instruments' European ASIC Product Group (Bedford, England), Europeans are very excited about the FPGAs being made and shipped by TI through an agreement with Actel (Sunnyvale, CA).

"Europeans don't like to waste money," says Inglis. He reports that 70 to 80 percent of the ASICs designed in Europe actually go to market (in the United States, the percentage of ASIC designs that make it into products is quite small), and suggests that "FPGAs can address telecom applications where only a few thousand units are required a year and can also be used for breadboarding automotive applications before production."

FPGAs will open the ASIC market up to distributors, according to Inglis. He expects that new FPGA customers will come from among those not using ASICs right now.

Peter Bingham, FPGA marketing manager at GEC Plessey Semiconductors (Swindon, United Kingdom), agrees with TI's Inglis that FPGAs will broaden the ASIC market. Bingham sees interest in Plessey's Electrically Reconfigurable Arrays (ERAs) coming from customers in Spain designing ICs for the first time. He says that the FPGA market is "smaller than one would expect in Europe," and adds that the major market is in the United States.

Plessey, having realized that the fine-grained architecture of its ERAs is a disadvantage as far as usable gates go, is trying to exploit the ERAs' in-circuit reconfigurability. Though denser devices to be introduced next year will be similar in architecture, ERAs of the following generation will have a new architecture. Since Actel and Xilinx don't have a strong presence in Europe, Plessey expects its strongest European competition to come from TI.

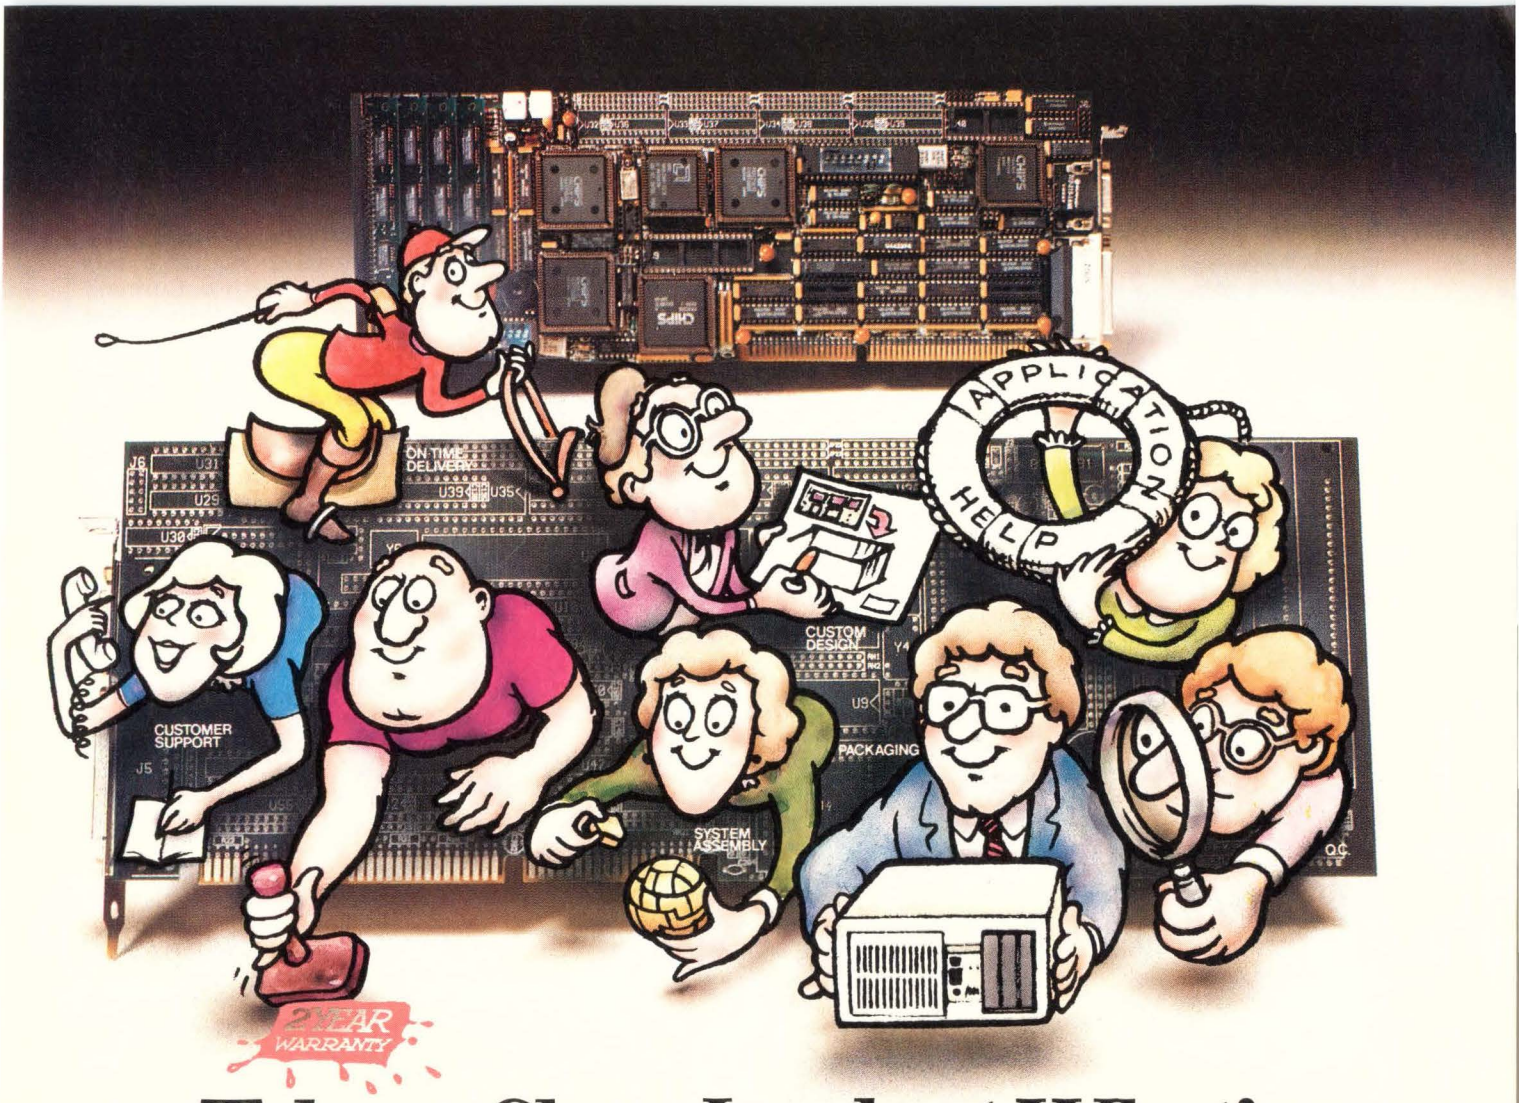

Take a Close Look at What's Behind Your PC CPU Board!

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.

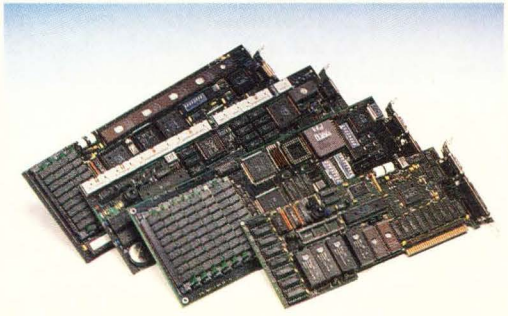

CPUs for the OEM and System Integrator

A286	8/12/16 MHz 80286 4 MB on-board DRAM Solid State Disk Option
SBC88	10 MHz 8088 1 MB on-board DRAM EGA/CGA/PGA/Hercules Floppy/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. 18

INTEGRATED CIRCUITS

many. "The 1992 scenario has driven investment into Europe," says Hoggar, "including equipment and manufacturing. Suddenly, Europe is the biggest market in the world—a single market with no trade barriers."

In fact, says Hoggar, TI is being pushed to add Eastern European languages to its 1-μm standard-cell Unitext decoder, which implements teletext processing. Teletext is a free European service that involves the use of broadcasting capabilities for displaying data such as news, flight information and stock prices on home TV screens. TI has designed its Unitext decoder to be upgradable to include Czechoslovakian, Polish, Turkish and Hungarian, should the need arise.

TI has also used its 1-μm standard-cell technology to develop a macrocell to form part of a voice coder for the new generation of digital cordless telephones that are emerging in Europe. That market, estimated by TI to be the fastest growing this decade, is being driven by the CT2 specification of the United Kingdom. The CT2 spec aims to create a Telepoint service whereby anyone with a CT2 handset can log on to a base station, located in a public place, with user-specific protocols and make a phone call via the public network.

European silicon vendors increasingly look to BiCMOS as an answer to the proliferation of mixed-signal designs.

To address the stereo TV market in Europe, TI has developed a standard-cell ASIC that descrambles and deinterleaves a Nicam 728 data stream. Nicam—or near instantaneous and companded multiplex—is the European standard for digital stereo of CD quality. The TI digital stereo demultiplexer provides output to industry-standard digital-to-analog converters. Also contribut-

The Multibus II product line has changed.

We've dramatically reduced your costs and added dozens of products. And that's great news when your designs call for more than a PC.

Today, our Multibus II system with a 386™ CPU, disk, tape, and 4MB of RAM is priced at just $9,995.

And the costs of our Multibus II single board computers are down as much as 47%.

Now you can build the Multibus II discipline right into your system for the price of VME or EISA. Multibus II's new MPI bus interface chip is perfect for building simple analog, digital, or serial interface boards. It costs just $40 in quantities over 100.

The Multibus II product line has expanded, too. In the last year alone, we've added over a dozen I/O and CPU board products. And you can choose from the more than 500 Multibus II products on the market, including 150 full-size Multibus II boards.

Need help migrating from VME? Our single-slot VME-to-Multibus II adapters will give you a jump on switching your custom VME boards to Multibus II.

To view the entire expanse of Multibus II products — from over 100 vendors — call Intel at (800) 548-4725, Dept. AA60. Ask for a free copy of the *1990 Multibus II Product Directory.*

So don't delay. Call now, and start a change for the better.

intel®

CIRCLE NO. 20

IN THE ERA OF MegaChip™ TECHNOLOGIES

APPLYING TI's BiCMOS

With more than 50 BiCMOS logic functions from Texas Instruments, you can beat tough bus-interface design challenges. Our free SamplePacs will show you how.

Specially designed for use in bus-interface applications, our growing BiCMOS logic family can make the difference in getting data on and off the bus faster. These advanced functions that combine the best of bipolar and CMOS can help you attain higher system performance levels.

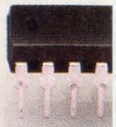

Lowering power, maximizing speed

For example, our BiCMOS devices can help you minimize power dissipation and maximize speed. Disabled currents are reduced by as much as 95% and active currents by as much as 50% compared to advanced bipolar equivalents.

DIFFERENCE

In fact, your system power savings can amount to more than 25%, and you should experience reduced switching noise as well.

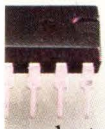

Yet you can maximize system speed. Switching speeds are comparable to advanced bipolar devices and provide the high drive current required for today's industry-standard buses (48/64 mA commercial, 24/48 mA military).

Gaining even greater performance

If you need even lower power and higher speeds, our submicron Advanced BiCMOS (ABT) family is the choice for you. Planned devices include 8-, 9-, and 10-bit buffers/drivers, transceivers, latches, registers, and registered and latched transceivers.

Our broad BiCMOS family also includes unique functions that can help you more quickly meet the design challenges involved with incident wave switching, driving MOS memories, and system testability.

Assuring incident wave switching

Wider word widths and additional cards on backplanes are requiring higher drive currents to assure incident wave switching.

Our BiCMOS family delivers. With our low-impedance line drivers, you get more "instantaneous" current even when impedances are as low as 25 ohms. You minimize transition "flat" spots that can degrade speed or cause oscillation at the receiving devices.

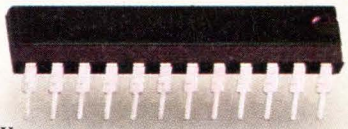

Managing MOS memory loads

MOS memory array interfaces create the high-capacitive loading environments that can result in overshoot and undershoot conditions. As a result, system reliability suffers. To handle this situation, our BiCMOS memory drivers incorporate a series damping resistor output structure that delivers advanced system performance when driving 256K, 1M, and 4M DRAMs.

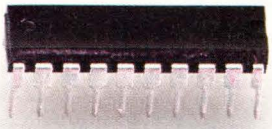

Building in testability with SCOPE

It is becoming more difficult to accurately test today's highly integrated boards and systems, but TI's BiCMOS family contains your solution: SCOPE™ (System Controllability and Observability Partitioning Environment) octals.

Used in place of standard octals, SCOPE devices allow specific circuitry within an assembled module, board, or system to be isolated for verification and debugging without manual probing. Currently, our BiCMOS family includes an octal buffer, transceiver, D-type latch, and D-type flip-flop.

TI's SCOPE products are the first to conform to the Joint Test Action Group (JTAG) specifications adopted by the IEEE 1149.1 Test Standards Committee.

Get your free SamplePac and sample our BiCMOS difference; call 1-800-336-5236, ext. 3008

You can take your choice of our BiCMOS SamplePacs containing a free BiCMOS device, our latest advanced logic brochure, plus appropriate product data. Just call the number given above, or use the return card to let us know which SamplePac you need to begin applying TI's BiCMOS difference.

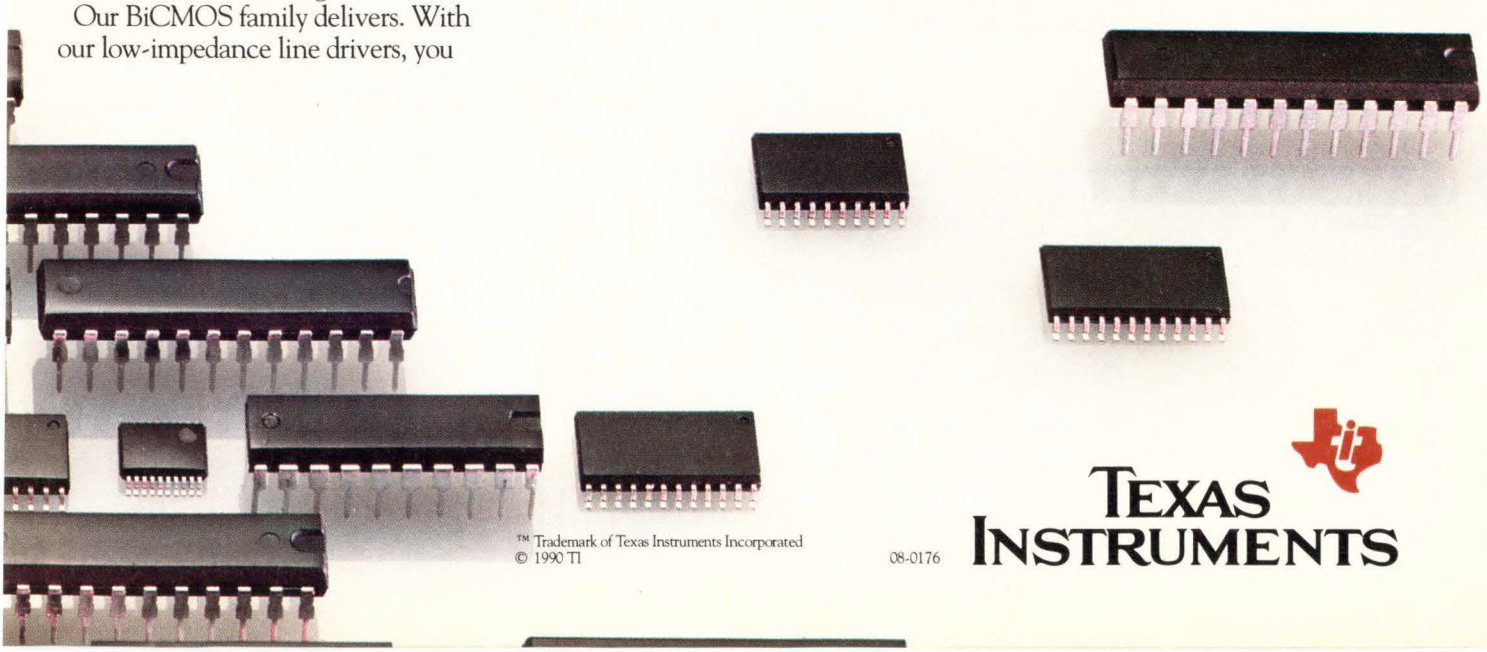

08-0176

TEXAS INSTRUMENTS

INTEGRATED CIRCUITS

ing to the development of the Nicam TV stereo sound decoding system are Philips with a dual-surface wave filter and GEC Plessey Semiconductors (Swindon, United Kingdom) with an IF amplifier and demodulator.

Cooperation can't hurt

The development of an interactive CD, specified by Philips with major components supplied by Motorola (Munich, Germany), represents another instance of mutual cooperation between European and U.S.-based silicon vendors. Motorola recently announced a CMOS 32-bit integrated processing unit with a pair of fast on-chip DMA channels that make it ideal for the high-speed data transfers required by the Philips CD-I. According to Gerd Westphal, ASIC operation manager for Motorola in Europe, "the CD-I is proof of excellent cooperative efforts and the sharing of silicon know-how."

Such cooperative efforts between U.S. and European vendors could give the Western world the edge it needs to compete in what has come to be a truly international semiconductor market. If the United States begins to look to Europe for technological partnerships as much as Europe looks to the United States, the United States might gain ground in the semiconductor race—or at least avoid losing any more than it already has. ■

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

Actel
(408) 739-1010 ..... Circle 262
Analog Devices
353-06129011 ..... Circle 263
GEC Plessey Semiconductors
44-0793518000 ..... Circle 264
Mentor Graphics/Europe
33-130671000 ..... Circle 265
Motorola
49-8992103118 ..... Circle 266
Philips
31-40791111 ..... Circle 267
Racal-Redac
44-0684294161 ..... Circle 268
SGS-Thomson/digital unit
39-6379133 ..... Circle 269
SGS-Thomson/analog unit
33-76585609 ..... Circle 270
Siemens Semiconductor Group
49-8941448080 ..... Circle 271
Texas Instruments
44-0234270111 ..... Circle 272
Xilinx
(408) 559-7778 ..... Circle 273

INTEGRATED CIRCUITS

Intel wades into palmtop fray with 386 SL

Ron Wilson, Senior Editor

Hard on the heels of the Advanced Micro Devices (Sunnyvale, CA) announcement of its 286 single-chip motherboard product, Intel (Santa Clara, CA) has joined the battle of the palmtop—the contest to minimize the chip count for an ISA bus personal computer. Intel's new product, the 386 SL Superset, has good news and bad news.

On the good side, the CPU is based on the 386 rather than the 286 architecture. On the bad side, the product is—as the name hints—a set: it takes two chips to complete the core logic.

The first Intel package comprises an extended but code-compatible SX CPU, cache controller with tags, DRAM controller and ISA bus controller. The second chip, the 82360SL, carries the remainder of traditional personal computer core logic, including timers, interrupt controller, DMA controller and real-time clock. The part adds parallel and serial interfaces and considerable additional hardware to support a new power-management scheme.

Based on a first look, the SL Superset won't automatically drive the competition out of even the palmtop PC market—the SL solution won't be much more compact than existing single-chip SX products from Opti (Santa Clara, CA), Headland Tech-

THE 386 SL SUPERSET

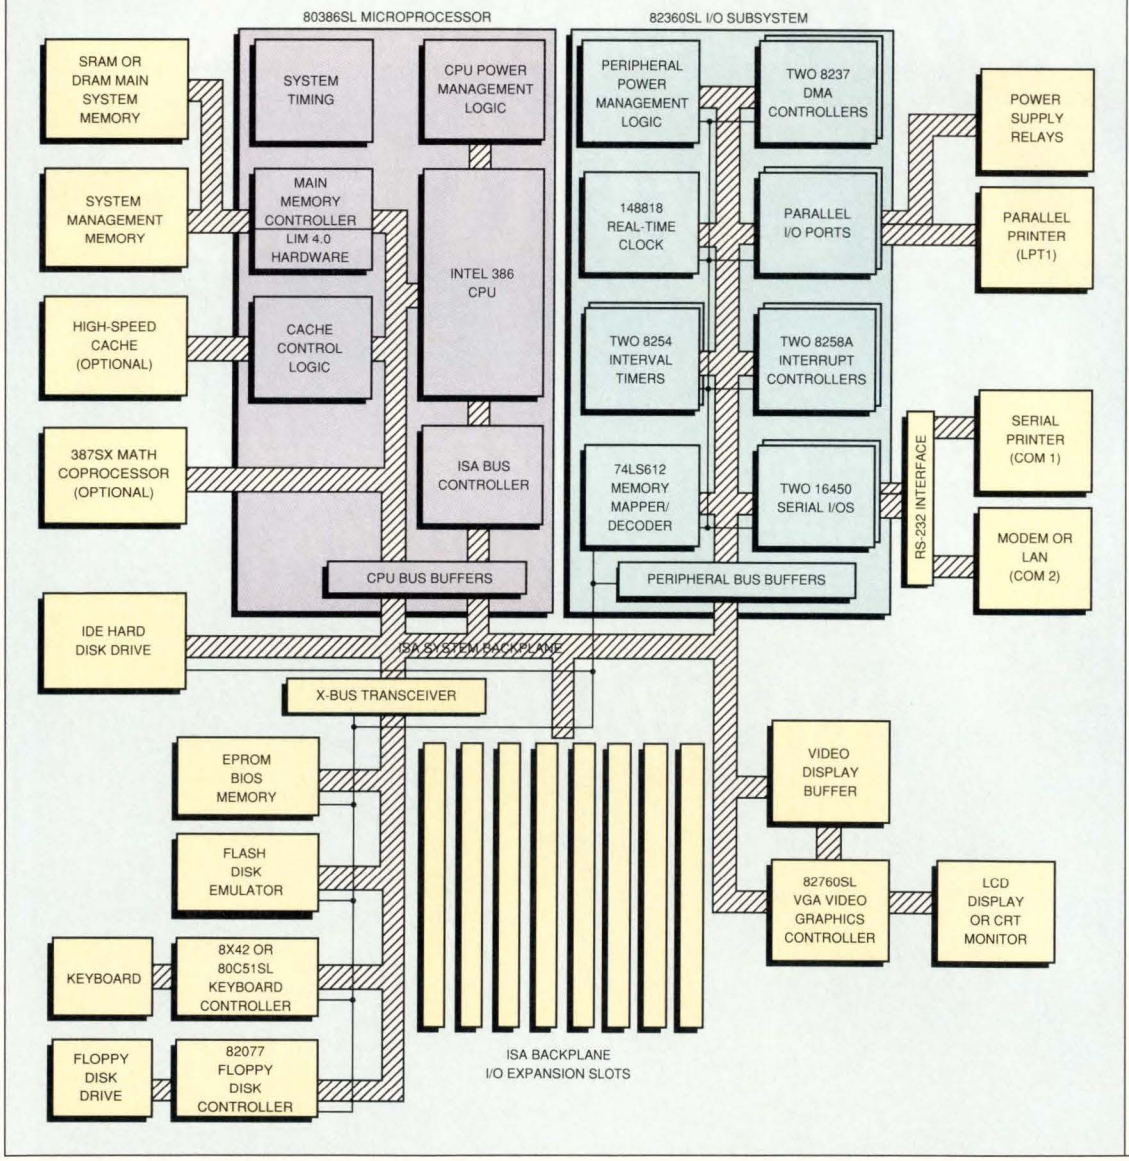

Intel's 386 SL brings a full 32-bit CPU to the space- and power-constrained world of notebook computing. The company shrank a whole motherboard chip set—including a cache controller—into two packages. Less obvious, but more important, the Intel architects have broken new ground in power management, including several new hardware features to help conserve battery life.

Unix Systems That Fit Your Image.

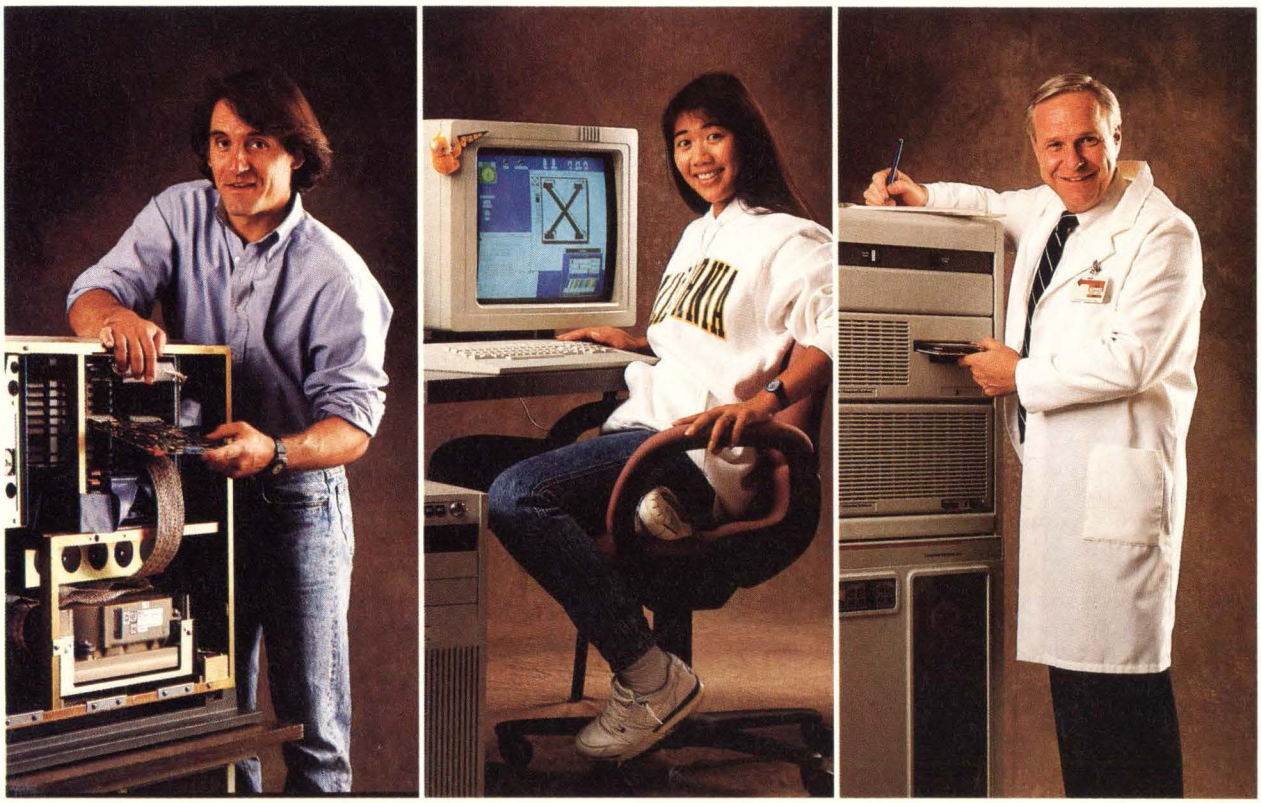

Mizar's Optimum Systems match your application style.

Do Unix® systems reflect their developers? Or are developers shaped by their machines? With the Optimum™ Family of Unix products from Mizar, you do the shaping, all according to your application needs and cost objectives.

Optimum Unix solutions range from pre-tested board and software kits to completely integrated systems. They're the kind of products you can easily tailor to meet unique application requirements. Best of all, the Optimum Family provides all the features and performance of the finest Unix workstations, but without their restrictive packaging and limited upgradability.

Based on the VMEbus, the 68030, and a high-speed memory architecture, Optimum systems are true performers. And our Berkeley 4.3 Unix with NFS™ and X.11 Windows™ support provides the ideal environment for both development and deployment of Unix applications. You can even develop on Sun workstations for later transfer to Optimum application systems with little or no change in source code.

But Optimum's capabilities don't stop at the basics. Designed specifically for high performance Unix applications, the Optimum family offers special features such as the Transparent Remote File System, TRFS,™ for file sharing and networking at twice the speed of TCP/IP over Ethernet.™ For multiprocessor applications, our unique clustering approach allows you to integrate multiple Unix processors in a single VME chassis.

Finally, Optimum products are backed by the best warranty, factory service, and technical support programs in the business. The Optimum solution doesn't just include hardware and software, but the personal attention you need to get your application running in record time.

If ordinary Unix products aren't your style, the Optimum Family probably is. Call today for more information on how we can serve your Unix application needs.

1-800-635-0200

MIZAR

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

INTEGRATED CIRCUITS

nology (Fremont, CA) or Chips and Technologies (San Jose, CA). But the SL will have a unique system-level solution to managing power consumption and a system-level attack on product size. These, rather than an integrated 386 processor, may be the telling features.

Compactness trade-off

Intuitively, a design that integrates much of the PC core logic onto the CPU chip ought to have the edge in compactness. But that isn't necessarily the case until you can do the whole job. Starting with the much smaller 286 core, AMD could get everything on one die. Starting with the bigger SX, Intel could not.

"There were a couple of reasons why we didn't go for a single-chip solution this time," explains Intel marketing manager Bruce Schechter. "To begin with, because of the type of design tools we used, and because the SX doesn't have any big areas of very regular patterns such as cache RAM, the SL CPU die is almost as big as the 486 die. So we couldn't economically put much more on the chip.

"Also, there was the matter of packaging. As it's configured, we fit the part into a new Intel-developed 227-lead land grid array (LGA). The package is like a pin grid array, but has metal pads instead of pins." The pin count would have gone up substantially if the functionality of Intel's two chips had been combined.

So Intel ended up with two packages: a 227-pad proprietary and a 196-pin plastic quad flat pack. The set doesn't require external ISA bus or DRAM drivers, but it will need external cache SRAMs—if the architect wishes to use cache—a new 8051SL keyboard controller and, of course, external DRAMs.

Existing SX solutions can just about match that chip count, and with standard packaging. "Our HT-21 chip needs only the CPU, three TTL packages, a real-time clock, keyboard controller and DRAM to complete the system," says Bert McComas, the director of single-chip products at Headland. "For notebook PCs, we have the highest-integration chip on the planet," he adds. Like the Chips and Technologies SCATsx, Headland's solution is designed to run with page-interleaved DRAM and will work with 20-MHz processors.

Neither of these single-chip logic products, though, contains an on-chip cache controller—that feature comes on the Opti 82C281, but at the cost of a total motherboard chip count around 20. "With the SRAMs and everything, cache nearly doubles the space taken up by the systems logic and substantially increases the cost," McComas says. "Customers who are after this level of integration usually aren't willing to spend that much for a small increment in performance."

But even with a cache controller, Intel hesitates to make performance an issue with the SL. Schechter argues that the 386 SX is an enabling technology for the very-small-computer market just because of its 32-bit architecture. "Packages such as Windows are very much an issue for notebook PCs," he maintains. "And handwriting technology is becoming important. These things require 386-level performance." Yet in discussing the design of the SL's unique, 16-bit line-size tag architecture, he explains, "This isn't a classical cache architecture. We've given priority to minimizing power consumption rather than to maximizing performance. So we make every attempt to minimize bus activity."

> ***"People should see this as an interim solution on the way to a real PC on a chip."***
>
> *—Bruce Schechter, Intel*

In fact, if there is a theme to the Intel announcement, it's not integration or performance, but power management. In designing the peripherals or the cache controller—even in modifying the system model of the CPU—the company has gone to great lengths to save every drop of battery energy.

The designers started by recognizing a software problem with existing power-management schemes. "The way systems are done today, DOS has the whole address space," Schechter says. "That means there's no space available for power-management code, so you have to do special drivers, and even then you run into conflicts—for instance, on a lot of laptops the suspend key won't work with protected-mode software such as Windows."

Intel's solution is a new interrupt—the System Management Interrupt (SMI)—and a complete set of hidden resources, including memory pages, I/O addresses and special registers, all reserved for power-management routines. When a situation arises that requires intervention of power-management code, the system can assert the SMI, giving control to entirely implementation-specific routines that are invisible to the normal DOS environment. Then the software can turn devices on or off, refresh memory or do whatever else is required.

In support of this concept, the 82360SL I/O chip brings some rather specialized hardware. There's the Ideaport—a set of six uncommitted lines with their own dedicated timers. The lines would ordinarily be used to power up peripheral devices, and the associated timers would announce, via SMI, when the peripheral had remained unused for some length of time.

In addition, I/O address-recognizer hardware stands watch for attempts to address peripherals that are currently powered down, again triggering an SMI. An executed-instruction FIFO in the CPU gives a record of activity, so power-management software can figure out just what the CPU did to trigger the recognizers.

Power-saving design

At a deeper level, Intel has done a completely static design for the 386 SL chip. Thus the processor can be shut down without having to save registers. Similarly, the I/O chip can be anesthetized until nothing in the system is running but the real-time clock and the DRAM refresh timer. Even the 80387 SX, notoriously resistant to power-saving measures, has been somewhat tamed. The numeric coprocessor lives on the cache bus in the SL system and receives its clock from the CPU chip. In this

Continued on page 42

WHEN MANY THINGS COME TOGETHER, BREAKTHROUGHS HAPPEN.

NOW AVAILABLE! A REAL BREAKTHROUGH IN DSP: 400MOPS ARRAY PROCESSING CHIP SETS.

The world's fastest DSP product family is now together as a complete system solution from Array Microsystems. We started by gathering world leaders in digital signal processing to build a new company dedicated to supplying the next generation of DSP. Our products are a response to the needs of E-Systems, Tektronix, United Technologies, and our other industry development partners. Together, we pioneered VAST™ technology, the first architecture to perform a complete DSP algorithm while operating on entire arrays of data in parallel—in real time, on a single silicon chip. Not only does array processing make 400 million operations per second possible, its high-level instructions make programming nearly transparent. For example, eight simple instructions execute a complex 64K-point FFT. Array Microsystems delivers complete solutions, not just DSP chips. The A66 family includes everything you need: chip sets, software development tools, complete array processor boards, and custom memory ICs and modules. The Digital Array Signal Processor (DASP) is the heart of the chip set, and executes 16 high-level functions, including FFT butterflies, windowing, complex multiplies, and general-purpose functions. The Programmable Array Controller (PAC) manages the entire system, including address generation for DASP and memory, and I/O up to 80 MHz. For even higher performance, you can cascade DASP/PAC chip sets (see table). You'll see A66 solutions in next-generation aerospace and defense systems, test equipment, medical instruments, and other breakthrough applications like HDTV. Join us by calling our DSP hotline. We'll help put the world's fastest DSP into your next application.

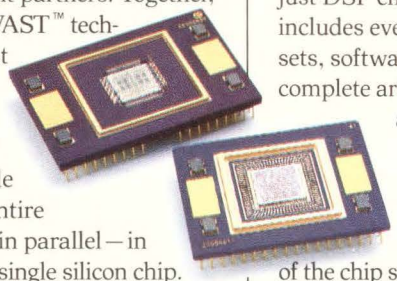

SOLUTION	1K complex FFT	64K complex FFT	2K tap FIR
Array A66			
1 processor	131 μSec	13.1 mSec	2.3 MHz
Fully cascaded	26μSec (5)	1.6mSec (8)	25MHz (13)

array™ MICROSYSTEMS

CIRCLE NO. 24

1420 Quail Lake Loop, Colorado Springs, Colorado 80906 • Call 719-540-7999, FAX 719-540-7950.

FLUKE®

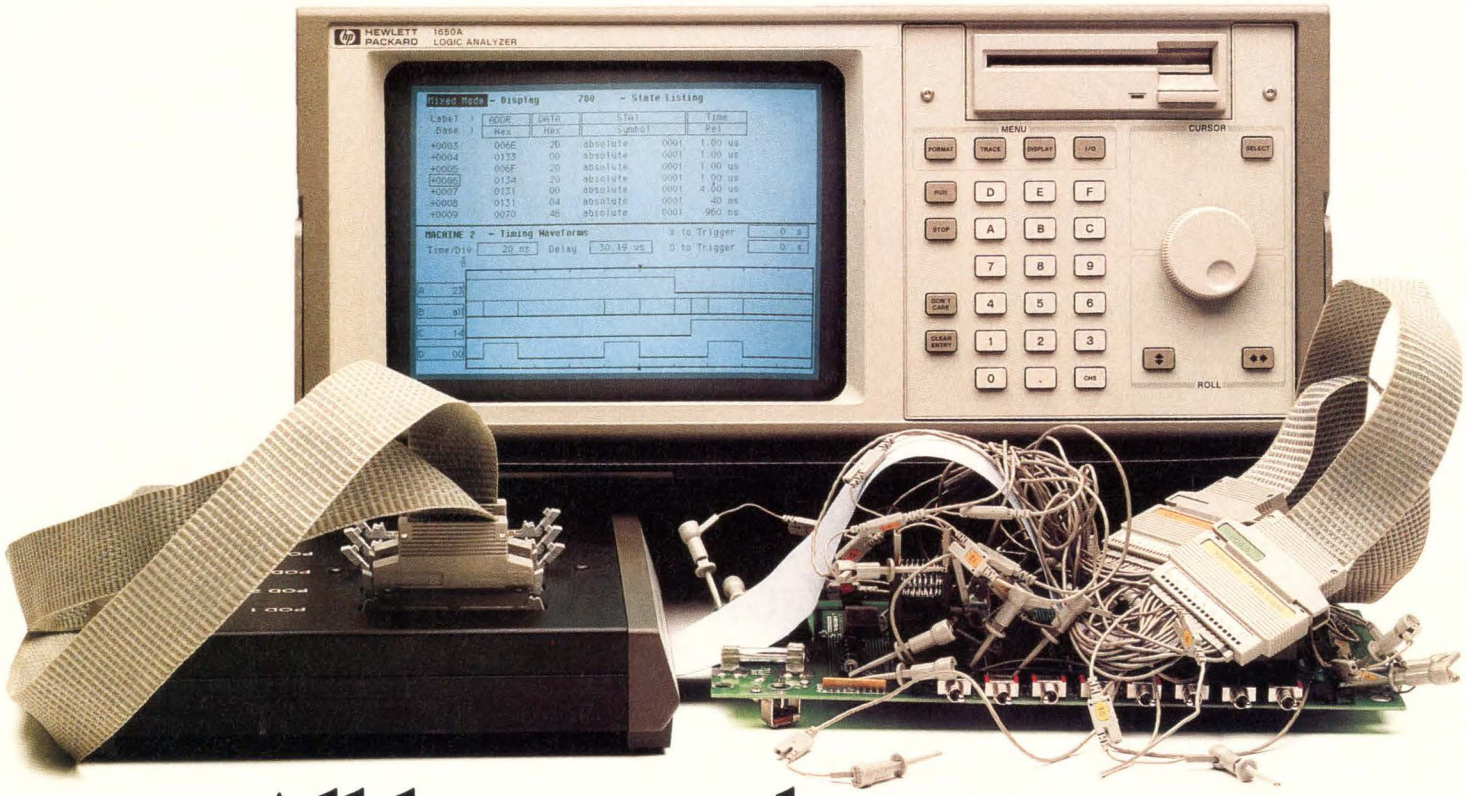

All logic analyzers give you integrated state and timing, sooner or later.

Why not now?

Introducing the new Philips PM 3580 family of logic analyzers from Fluke: the first instrument architecture to give you state and timing together on each channel—with a single probe.

Connect the probe to your board for state and you're automatically hooked up for timing. Or vice versa.

This means no more dual probing—a pain anytime and the source of loading problems—and no reconfiguration between state and timing. Which makes these analyzers simple to learn and use.

Plus, the pop-up menus and keyboard shortcuts guide you quickly through setup and data analysis. No matter if you use it every day or once a year.

What's more, capturing an elusive bug has never been easier with eight unrestricted trigger levels that let you select from state and

A basic comparison: Record all state and timing data on an 8-bit microprocessor with multiplexed bus, 8-bits for address, 3 control signals and a clock.	HP 1654B	Tek PRISM/MPM	Philips PM 3580/30
Probing:			
Channels used	48[1]	48[2]	20
One connection	No	No	Yes
Setups	Two	Two+	One
Interfaces to learn	Two	Two+	One
Integrated state & timing triggering	No, only one arming condition	No, only indirect 4-bit Teklink	Yes, 8 levels
State & timing data per pin	No	No	Yes
Price	$6700	$8600	$4250

Footnotes: 1) 8 channels lost to de-multiplexing
2) De-multiplexing requires double probing and only nine high-speed channels on basic unit

PHILIPS

PHILIPS

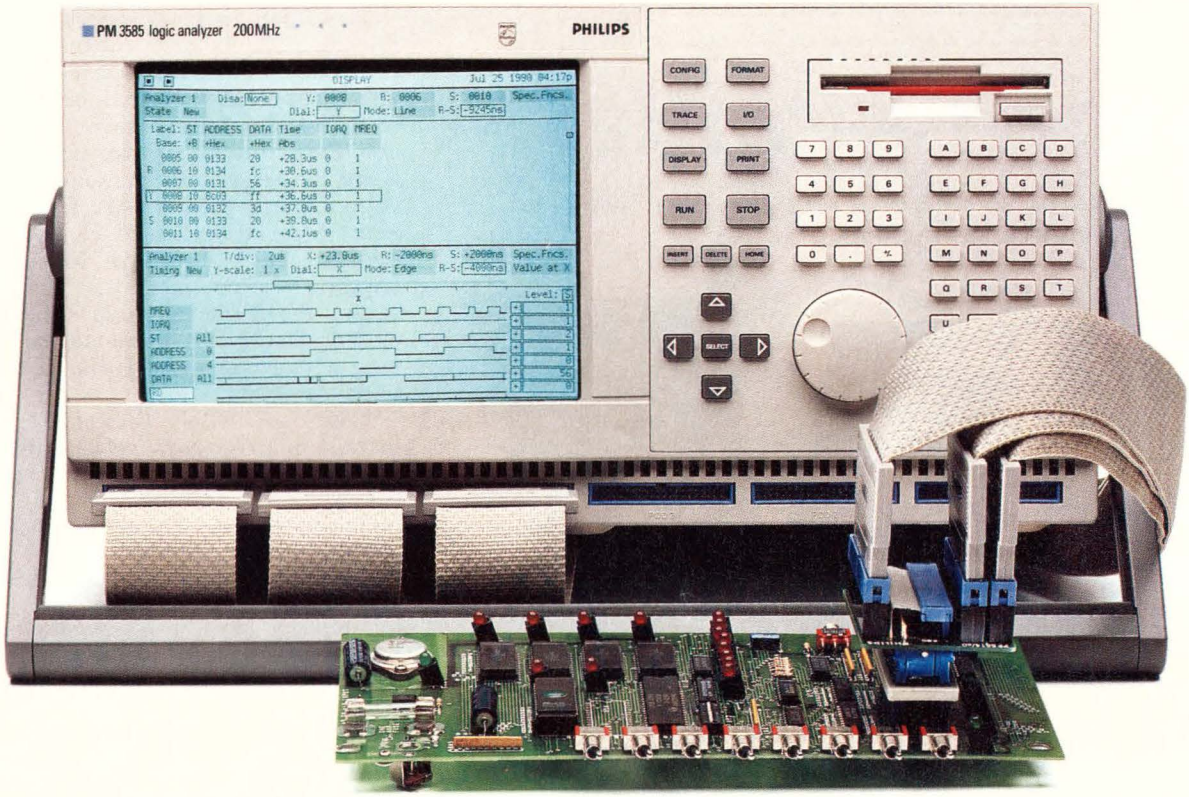

Why not sooner?

timing trigger conditions on each level.

But being simple doesn't mean simplistic. Basic performance of the PM 3580 family ranges from 32 to 96 channels, each with 50 MHz state and up to 200 MHz timing, plus 3 nanosecond glitch capture and 2K of memory per channel. For 8-, 16- and 32-bit processors.

And you get all this for nearly half the cost of comparable analyzers.

We'd like to send you a video. Or show you how to make state and timing measurements at your workplace—in 30 minutes or less.

We'll even bring the stopwatch so you can time us. And you can keep it to time the competition.

So call us today at **1-800-44-FLUKE.**
Ask for extension 720.
Because sooner is better.

FAST ANSWERS.

FLUKE®

CIRCLE NO. 25

INTEGRATED CIIRCUITS

Continued from page 38
configuration, the CPU can divide the FPU clock by 16 when no floating-point instructions are pending, and the power-management software can checkpoint the big chip and shut it off completely during a suspension.

Complementing its attention to the power-management issue, Intel has begun to address the other crucial need of notebook and palmtop system designers: compactness. The company has announced a series of products, drawn from a number of Intel divisions, to reduce the size of the rest of the notebook system. These include a power-managed 8051 keyboard controller, a single-chip VGA controller with LCD capability and a pair of modem chips.

The flash memory operation in Folsom, CA, is also contributing, providing both a flash-based BIOS and a series of flash memory cards in 1- to 4-Mbyte capacities, based on the Personal Computer Memory Card International Association/Japan Electronic Industry Development Association (PCMCIA/JEIDA) memory card standard. More intriguing, the company will use the same JEIDA format for a series of chip-on-board I/O functions, beginning with modem and 10BaseT connectivity cards.

Taken as a whole, then, the Intel introduction appears rather different than at first glance. Rather than betting on the technical appeal of a more-integrated CPU chip—not an unfamiliar approach for Intel—the company has structured a system solution for tiny computers around its 386 SX core. Individually, the parts may not be more compact—and certainly won't be cheaper—than competing single-chip SX solutions; but the package may offer new leverage in the control of both power consumption and space.

And there is another, more ominous message for competitors in the Intel system sell. "People should see this as an interim solution on the way to a real PC on a chip," advises Schechter. "It's already our assumption that a VGA controller is part of that future product, for instance."

How far will Intel take the integration process? Schechter hints: "We spent the time to develop the LGA package ourselves. For this product, the LGA is important, but in the next generation, it will be necessary." ■

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

Advanced Micro Devices
(408) 732-2400 ..... **Circle 205**
Chips and Technologies
(408) 434-0600 ..... **Circle 206**
Headland Technology
(415) 656-7800 ..... **Circle 207**
Intel
(800) 548-4725 ..... **Circle 208**

INTEGRATED CIRCUITS

For RISC, future may lie in integration

Ron Wilson, Senior Editor

Now that RISC processors are established in the workstation and departmental computer niches, vendors are turning their attention to new horizons—high-volume embedded applications, personal computers and ever-faster workstations. But it's becoming clear that existing RISC implementations—MIPS or Sparc—won't get into these new areas. For the embedded or mass-market PC arenas, workstation-oriented chips demand too high a system cost and are too difficult to work with. And the problems of chip interconnect are limiting the future of stand-alone CPUs at the high end. In all these applications, the future of RISC may lie in greater integration.

Three recent RISC introductions involving the MIPS R3000 architecture illustrate two paths to a new level of integration. Two, from Integrated Device Technology (Santa Clara, CA) and LSI Logic (Milpitas, CA), pick and choose among support devices to produce single-chip MIPS cores. The third, also from LSI, uses chip-on-board module technology to wring performance out of the R3000 while flushing complexity out of the design cycle.

"We aimed at the applications where the 68020 has ruled for years," says Bob Rowe, IDT marketing manager. "We intend to come in with a similar-sized part, at 15 Mips instead of 2 Mips, and with a similar price—as low as $30 in quantity."

Carefully selected chips

To achieve these goals, IDT picked carefully among the array of expensive chips that surround the R3000 CPU in a workstation. The designers eliminated parts that didn't contribute to the typical embedded application, such as the floating-point coprocessor interface, but left features that did get used—for example, the on-chip MMU, which is vital to X Windows terminal developers. Then the chip designers integrated essential parts onto the CPU die. "We talked to a lot of customers," Rowe says, "and came up with a general rule: don't constrain the chip with application-specific hardware, but use the silicon to build an execution engine."

The philosophy led IDT designers to integrate a 40-MHz R3000 core with instruction and data caches—4 and 2 kbytes, respectively, for the low-end 79R3051—on one die. They added a bus controller, thereby encapsulating the R3000's infamous critical paths within the silicon. "We were able to eliminate complexities such as the reset state machine and the clock sequencing logic," says technical marketing engineer Phil Bourekas. "And having the cache on the chip eliminates the need for external 8-ns SRAMs."

LSI Logic took a similar tack with its LR33000 CPU. The LSI chip also starts with a MIPS R3000 CPU and adds caches and a bus interface. But the 33000 uses a different mix of caches, provides a novel, embedded-system-oriented bus, and throws in a timer module as well.

"At 8 kbytes, ours is the biggest on-chip instruction cache of any chip out this year," says LSI marketing manager Rob Tobias. "Based on code samples from embedded designs, we determined that it was important to spend silicon area on a larger in-

IDT R305X ARCHITECTURE

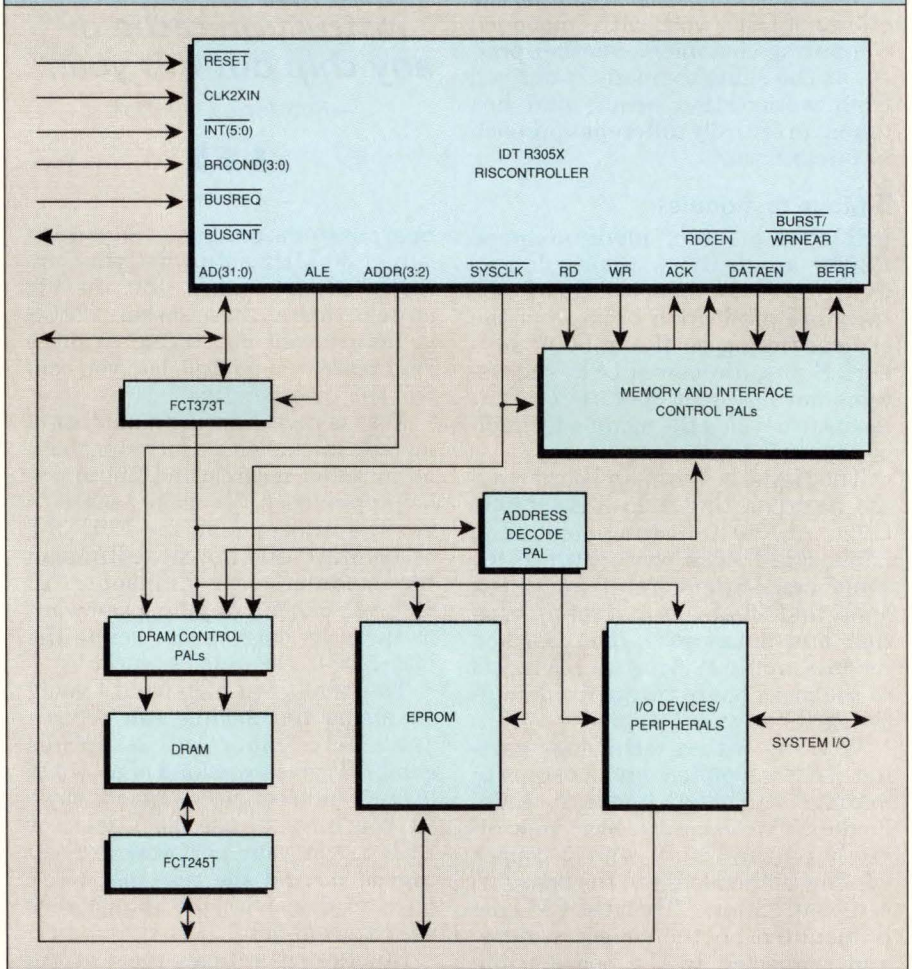

While embedded computing requires the speed of RISC, it can't always tolerate the complexity of the early workstation-oriented chips. IDT's answer is a high-integration single-chip CPU that meets the needs of control applications. The part squeezes small instruction and data caches onto the MIPS R3000 CPU chip, saving users from the legendary problems of the R3000's fast cache bus.

INTEGRATED CIRCUITS

struction cache, and that it was acceptable to cut the size of the data cache down to 1 kbyte."

To further reduce the problems designers must face in working with a 40-MHz parallel bus, the 33000 carries an on-chip DRAM controller that can directly drive modest-sized DRAM memories. The bus also has provision for multimaster DMA with bus snooping, including a quick DMA mode in which the CPU drives the DRAM lines, and the peripheral device simply provides data. Additional bus modes include more-traditional synchronous 32-bit operation and a direct 8-bit mode with gathering, primarily used with ROMs and peripherals. The chip provides chip selects and wait-state generation as well.

While the 33000 clearly shows the effects of LSI's work with embedded computing customers, another project at the same company is derived from workstation needs, and has taken an entirely different approach to integration.

Move to modules

LSI sees a lot of medium-speed R3000 applications taking longer than necessary. "People are slow getting into production because of the critical timing on the R3000," says Dirk Smits, director of LSI's subsystems development operation. "That led us to develop the module technology we call Ngine."

The Ngine is a chip-on-board module incorporating a 25-MHz R3000 CPU, an R3010 floating-point chip, LSI's 3220 read/write buffer IC, some bus buffers, reset and clock logic, and 32 kbytes each of instruction and data cache. The finished module is only 12.3 in.^{2} and connects to the main board through a demultiplexed bus on a 100-pin connector.

Working with a technology partner, LSI has pulled out the stops to manage lead length and heat on the module. "We mount a heat sink directly on the board, with a copper window embedded into the board itself," says Smits. "Then the CPU die is mounted on the copper window and connected to the board using conventional wire-bonding techniques. It's a process used extensively in consumer electronics, and by eliminating the lead frame it gives us better control over electrical, thermal and economic factors."

That control will become more important as LSI continues up the speed curve to 40 MHz. "At that speed, getting off-chip delays down to 1 ns is critical," Smits says. "We believe we can get there with this chip-on-board technology."

Cypress Semiconductor's Ross Subsidiary (Austin, TX) sees a similar use for module technology in the Sparc family. "With modules, you can control capacitance. And because you get to match the parts yourself, you can eliminate guard bands on speed," says Pete Simmons, Ross program manager for Sparc systems products. "So you can run at 50 MHz with what are nominally 48-MHz parts. But the big issue is that at these speeds, CMOS is just maxed out trying to drive long traces—a module lets you control the trace lengths."

"At 8 kbytes, ours is the biggest on-chip instruction cache of any chip out this year."

—Rob Tobias, LSI Logic

Ross is investigating a number of module technologies, including both alumina-on-ceramic and silicon wafer approaches. "It makes sense to put the integer unit, the FPU, the cache MMU and the cache RAMs in the module," says Simmons. "In fact, you can put anything you want in there as long as you define the interface to the outside world."

This ability to hide what's going on inside the module can become extremely important as speeds climb. "It gives you kind of a second layer of indirection," Simmons says. "If you have to use 3.3-V logic or ECL to get your performance, you can do it. And the customer won't have to deal with it—he just sees the external bus."

But the advantages come at the cost of technical problems. Primary among them is the need to test bare dice before building the module. This is necessary because, with the typical failure rates among untested dice, mounting a dozen of them on a module almost ensures that the final assembly will be bad. And the technologies in question aren't friendly to reworking or even probing. "We're really confident that we can test the dice over voltage and temperature," Simmons says. "But we won't actually know until we're in production next year."

Two-pronged approach

At the speeds Ross is addressing, modules are probably neither economical—the technology has yet to show any susceptibility to price reduction—nor labor-saving tools. But they are a necessity. "We're thinking that above 50 MHz, we'll sell only module products," Simmons says. "It's still possible for customers to do a PC board, but they'd have to do diode terminations and balanced-impedance lines—the whole thing's a black art at these speeds. It's just more viable to have a fast module with a familiar interface such as a 40-MHz Mbus coming out of it."

Simmons sees a bifurcation in the market, with different technologies pursuing high-performance and low-cost needs. It's a good bet that the technologies chasing the low-cost, embedded and PC-oriented applications will be high-integration single-chip cores, such as the IDT 79R305X family in the embedded market; or something like the Solbourne/Panasonic MN10501 single-chip Sparc core in Solbourne's new S4000 workstations. At the high end, we're likely to see increasingly elaborate core modules, using chip-on-board or chip-on-silicon technology in a protracted war against loading and heat. In any case, the original idea of a simple, low-integration RISC CPU die that could be quickly scaled to a new process seems to be fading quickly. ■

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

Cypress Semiconductor
(408) 943-2653 ..... **Circle 235**
Integrated Device Technology
(408) 727-6116 ..... **Circle 236**
LSI Logic
(408) 433-8000 ..... **Circle 237**
Ross Technology
(512) 448-8968 ..... **Circle 238**

And Now - Whitesmiths C Compilers and Source Level Debuggers for the 68HC11, 8051, Z80, 64180, 6809

C How They Run!

InterTools Products

Intermetrics' C cross development tools will help you clear the hurdles, so you can create a fully integrated environment that can deliver ***fast***, ***efficient code*** for any host-target configuration you choose.

The Intermetrics state-of-the-art ***cross compilers***, ***assemblers*** and ***utilities*** are built to work together as a team. Whether you're using a ***Sun***, ***Apollo***, ***DEC***, ***HP*** or ***IBM*** system, targeted to run on a ***680x0***, ***80x86***, ***V Series***, ***DSP96002***, or ***Am29000*** microprocessor, we are the single vendor that can help you get your development project across the finish line!

Our ***Source Level Debugger***, ***XDB***, works like a champion with our new low cost ***ROM Monitor***, and with our new ***hardware simulation board***, as well as with most in-circuit emulators.

Our technical support department will help you keep your development environment up and running by providing toll-free telephone assistance, free software updates and reduced rates for new releases.

Intermetrics' C cross development tools are the top tools in their class—call us to see how they ***run!***

1-800-356-3594

in Massachusetts or Canada (617) 661-0072

Intermetrics

Intermetrics, Inc. • 733 Concord Avenue • Cambridge, Massachusetts 02138-1002

CIRCLE NO. 28

NOW, LOGIC ANA AT THE SPEED

The NEW Tek Centurion:

400 MHz timing analysis.

100 MHz state analysis.

100 MHz complex triggering.

100 MHz synchronous clocking.

100 MHz timestamp.

RISC chips. CISC chips. High-speed busses. At the leading edge of clock speeds, channel counts and complex integration, one logic analyzer is equal to the test.

The Tek Centurion.

Centurion supplies an unprecedented 400 MHz timing. Plus 100 MHz synchronous acquisition on 96 channels at once—*all on one card.* You can expand to hundreds of channels in a single mainframe, with up to 32K memory depth per channel!

Add the most revealing disassembly in the industry, for any processor pictured here. Add the easiest, most compact probing. Get real powers of analysis, not just ordinary acquisition capabilities.

Call 1-800-426-2200 to learn more about this great leap forward in logic analysis. Or see your Tek sales engineer for a personal demo soon.

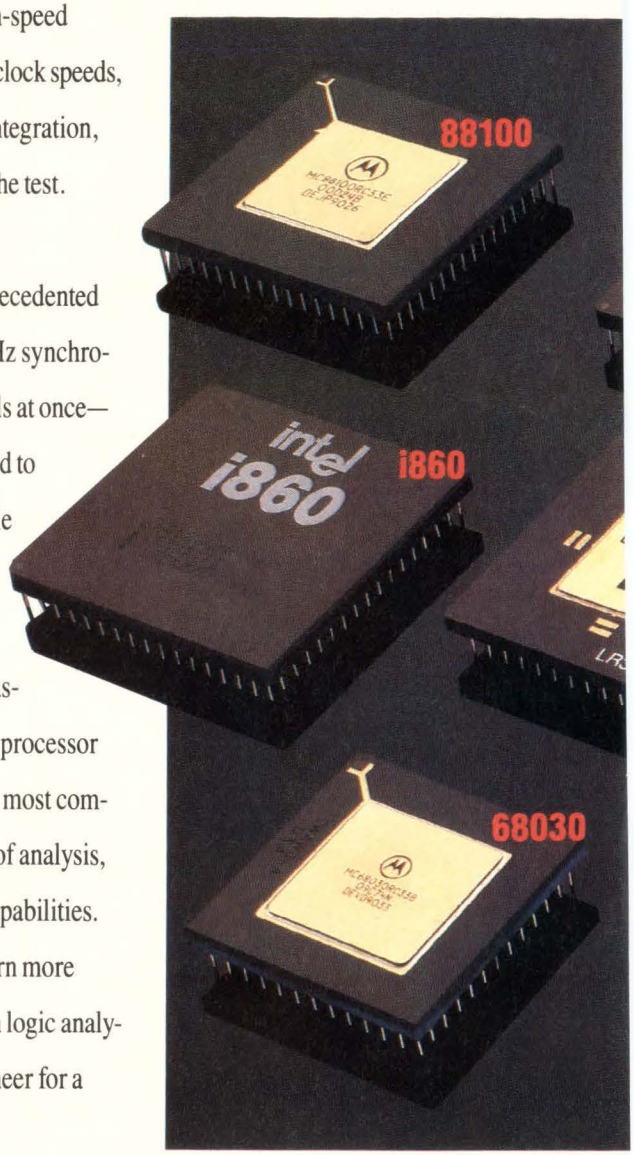

Rental units available from your local rental company.

LYSIS OF RISC.

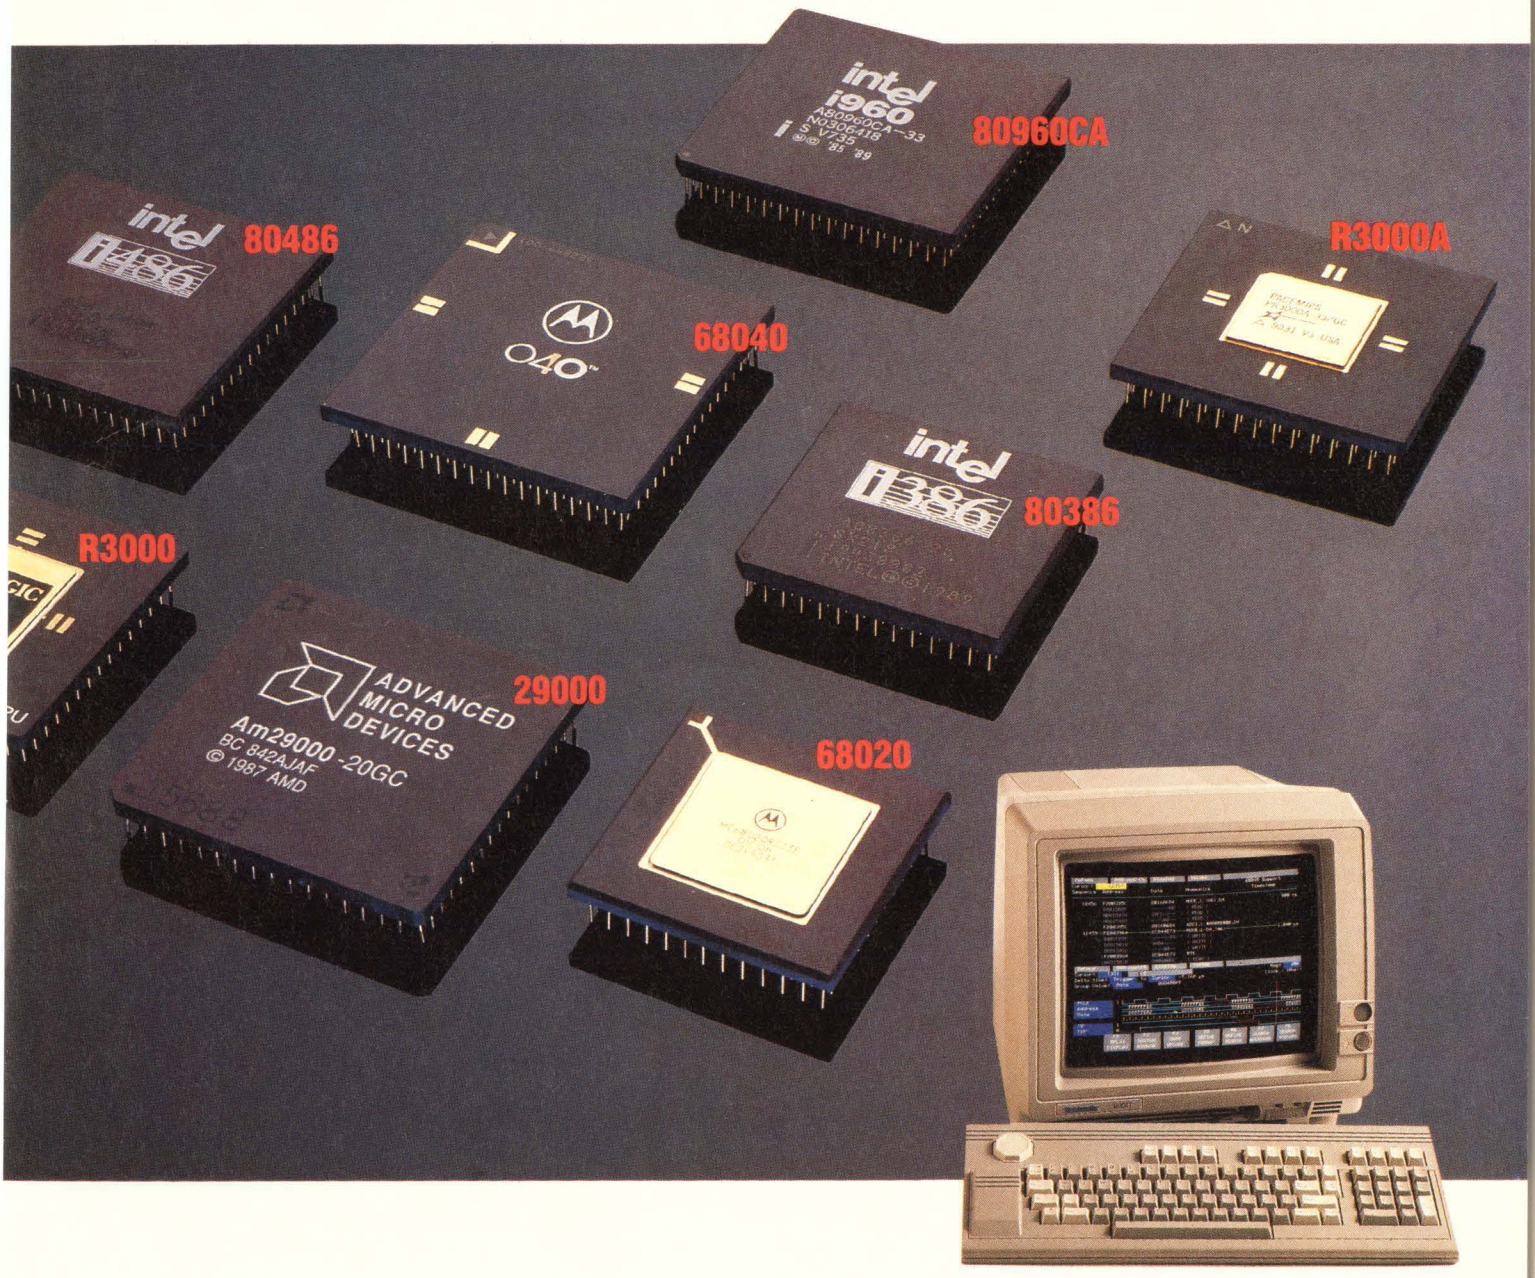

Tektronix®
COMMITTED TO EXCELLENCE

CIRCLE NO. 29

Frustrated with microprocessor system debugging?

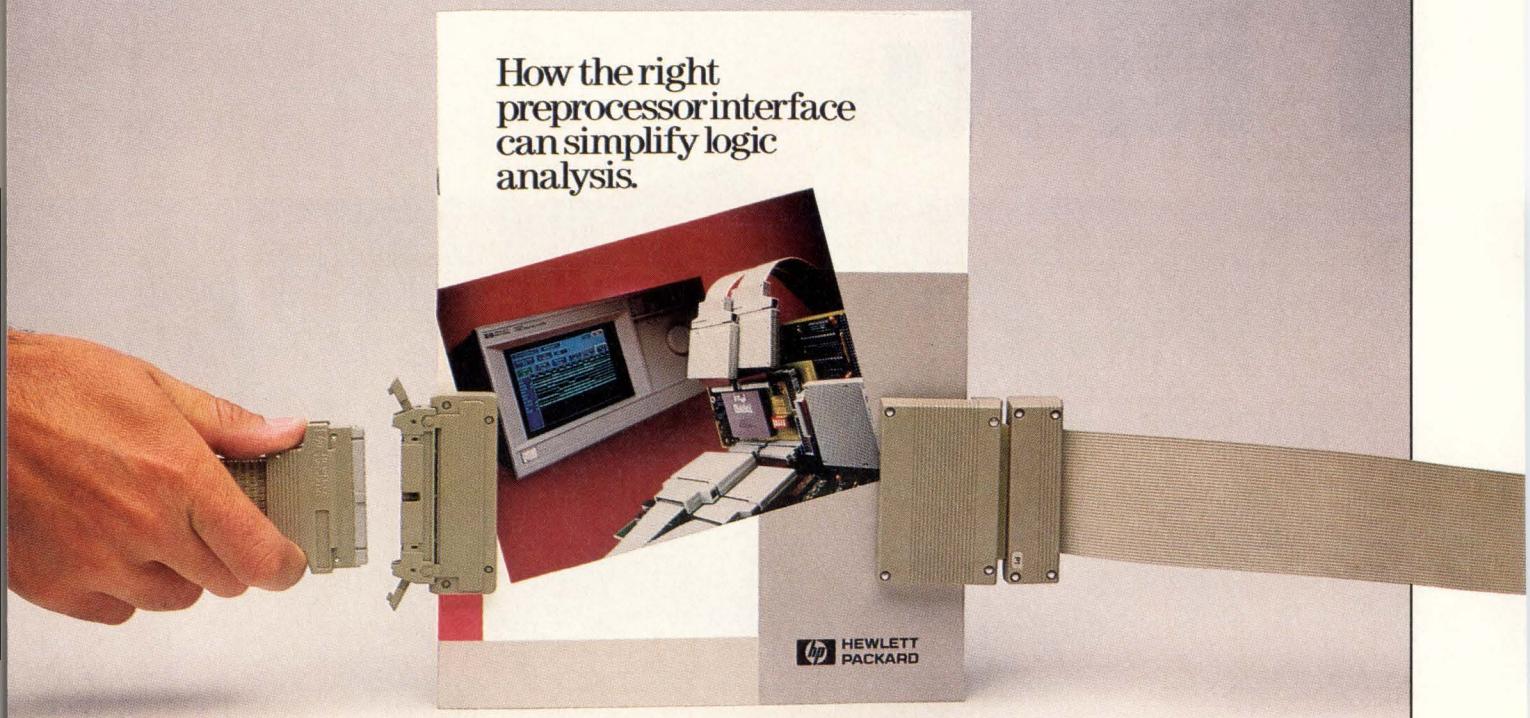

Here's free help in making the connections.

Plug into our new series of application notes today. And learn how to save time and aggravation when you're debugging microprocessor systems and busses.

Find out how an HP preprocessor literally makes it a snap to connect your microprocessor to an HP logic analyzer. And see how completing your HP logic analysis system with a preprocessor will speed and simplify measurements and data analysis. You'll also discover the industry's broadest selection of preprocessors – over 60 models support more than 70 microprocessors and busses including 100 MHz support for RISC and high-performance CISC processors.

So call **1-800-752-0900** today. Ask for **Ext. 1402** or mail the reply card and we'll send the application information you need from our FREE series of six industry-specific application notes.

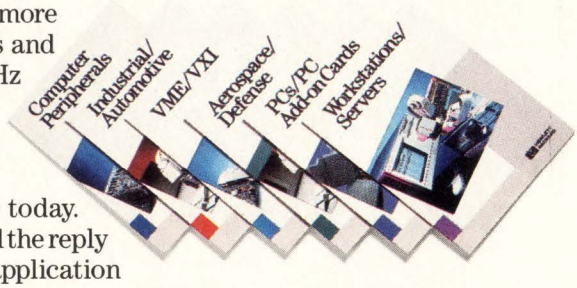

There is a better way.

HEWLETT PACKARD

INTEGRATED CIRCUITS

Image-compression chips advance on three fronts

Ron Wilson, Senior Editor

The ability to throw 100,000 transistors at a problem is rapidly advancing the cause of image compression. From fax machines to illustrations on computer screens to digital video cameras, silicon solutions are showing up to reduce the crushing bulk of image data. In some cases, the chips are out ahead of the standards that they're supposed to implement.

The simplest and most widely used of image-compression algorithms is employed in CCITT Group 3 and Group 4 fax machines. The algorithm simply looks for repeated white or black pixels—this is a bitonal system—along a scan line and translates the string of pixels into a bit and a pixel count. A two-dimensional version of the algorithm also compares consecutive scan lines to see if the previous line can be replicated. The algorithm is easy and lossless—the compressed and expanded image is bit-for-bit equal to the original—but compression works best on characters or line drawings.

Advanced Micro Devices (Sunnyvale, CA) has for some time marketed a video compression and expansion processor (VCEP) chip that handles the CCITT requirements. But in its continuing downsizing efforts, the company has spun off the product and its technology to Oak Technology (Sunnyvale, CA). "We see the VCEP as an excellent fit for Oak's direction," says Steven Gary, Oak director of marketing, who moved from AMD along with the chip he helped develop. "The device's throughput of about 50 Mbits/s, or about six typical business pages per second, opens a whole range of applications beyond just fax equipment. It will be used in intelligent printers, digital copiers and personal computers that routinely handle document images."

This range of applications gives Oak an idea about future elaborations of the existing part. "There's an opportunity to make the device faster and to integrate more circuitry onto it," Gary says. He points out, for example, that a current PC implementation for document compression is board-sized and requires about 20 chips. If that can shrink to a size where document compression is a motherboard option, users may become accustomed to expanding a document for viewing on a routine basis, and then recompressing it for disk storage.

Dealing with color images

While the CCITT fax solution is easy, it's specialized for two-tone images dominated by characters. Even halftone subjects can compress poorly, or even get bigger during the compression process. For the more general problem of compressing gray-scale or color images, the industry is turning to a different standard: that of the Joint Photographic Experts Group (JPEG).

The JPEG approach differs fundamentally from the fax approach in that JPEG uses the discrete cosine transform (DCT) for compression. This transform is fairly easy to compute and gives better compression ratios than the fax technique, but it isn't lossless—reconstructed images will degrade as the compression ratio increases. JPEG compression is being investigated for a wide variety of picture-related applications, from imaging workstations to still video cameras and even digital VCRs.

The most vocal of JPEG proponents has been start-up C-Cube Microsystems (San Jose, CA). This year the company introduced CMOS VLSI that implements JPEG compression at up to 30 frames/s. Now Next Computer (Redwood City, CA) has announced that its second try at a successful PC, the Nextdimension workstation, will include the C-Cube compression processor. While the win may not mean a lot of units, given Next's record with its first product, it could stimulate a lot of thinking about the issue of JPEG compression in imaging workstations and generate some applications that depend on the C-Cube hardware.

Meanwhile, another entrant has moved into the compression arena, not only with a JPEG solution but with a chip set for the more elaborate Motion Picture Experts Group

VIDEO COMPRESSION

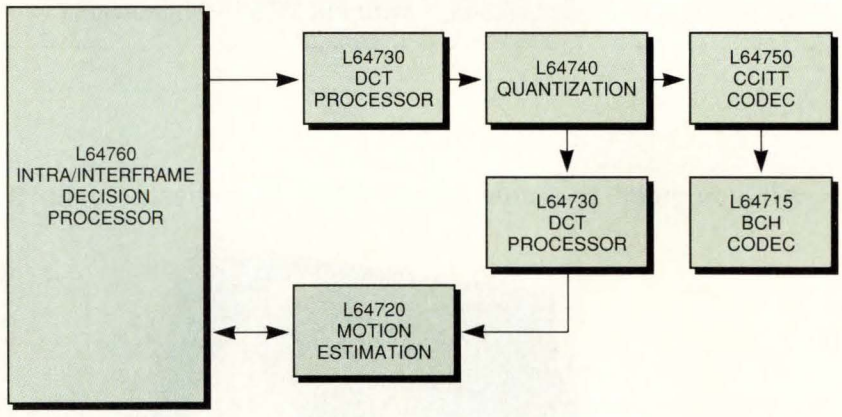

A new multichip set from LSI Logic offers the first attempt to condense the still-volatile MPEG graphics compression standard into silicon. The set performs—at video speed—the discrete cosine transforms, coding and interframe comparison functions necessary to achieve MPEG's impressive compression ratios. Initial applications will probably be in price-tolerant areas such as video conferencing.

INTEGRATED CIRCUITS

(MPEG) compression proposal. LSI Logic (Milpitas, CA) has announced that it's working on a series of very dense chips that will eventually implement both a two-chip JPEG engine and a 10-chip pipeline to deal with full-blown MPEG compression.

Better compression

The difference between the two compression techniques is, essentially, motion. While JPEG applies the DCT to a fixed image, MPEG assumes that it's working with a sequence of frames, each differing from the preceding one in only a few respects. So the MPEG algorithm first compresses a frame using JPEG-based DCT techniques, then re-expands the frame, compares it with the new frame coming in and determines how much of the difference is simply motion of the objects in the frame. This leads to a substantial reduction in the amount of data, since there's no need to regenerate areas of the picture that have only moved, rather than changed their contents.

Needless to say, all this regenerating and comparing takes additional hardware, but it can be worth it. "We can get about 20:1 compression out of spatial compression," says Peng Ang, LSI business and technical director, "and we get another 10:1 from temporal—that is, interframe—compression. Altogether, in applications such as video conferencing, we can expect data compression of 200:1 in a few years."

"Existing video conferencing equipment is already getting 100:1 compression, but with the sacrifice of some image quality," adds Simon Dolan, LSI product marketing manager. "So with the chip set, we should be able to do a $5,000 to $10,000 video conferencing station with full MPEG compression."

An MPEG encoder would normally start with the 64760 Intra/InterFrame Decision Processor. This chip looks at motion data from the 64720 Motion Estimation Processor and decides whether to use new data or references to previously encoded data to describe the new frame. From the 760, pixels flow to a 64730 DCT processor, which does the spatial compression, and then to a 64740 Quantization processor to prepare the transformed data for coding.

Next a 64750 CCITT 64750 codec (coder/decoder) produces the actual codes that will represent the transformed pixel blocks. This data moves on to a 64715 BCH endec (encoder/decoder), a chip specifically designed to deal with the error-correcting BCH code specified by CCITT H.261. Meanwhile, a copy of the quantized but uncoded data moves back, through a second DCT

chip for expansion into a frame buffer. From here the Motion Estimation chip will extract the data for comparison against the next frame. LSI estimates that the whole pipeline should run at a peak speed of 30 Mpixels/s.

Practical applications

Despite the performance and compactness of the solution, barriers still separate MPEG from wide application. First, there's the state of the standard. Revision zero has been available for review since August, but Ang doesn't expect a stable document until early 1992. In the meantime, chip vendors and users will be only betting on hardware compatibility with the final standard.

Second, in current technology the algorithm still translates into big, expensive chips. The MPEG pipeline from LSI will require 10 processing chips at a cost of almost $700 (1,000s), plus associated buffers, frame memories, clocks and glue.

For many applications now, customers may choose to give up the extra factor of 10 in compression and simply use JPEG hardware that's fast enough to keep up with video rates. LSI, for instance, provides an alternative to the C-Cube silicon. The company offers the 64730 DCT Processor with a JPEG companion, the LVC650 Image Compression Processor, as a 30-MHz, two-chip JPEG engine. "Thirty MHz is right for H.261 applications," says Ang. "In specialized algorithms, the chips can peak out at 40 MHz. And the set will cost about $100 to $150."

Even as the standards are coming together, competition is growing for image-compression implementations. And with big differences in cost and compression ratio, solutions overlap. CCITT Group 3 compression at 50 MHz is fast enough to compress two-tone video data inexpensively, but the algorithm will break down on halftone data. JPEG compression, with losses but with perhaps 20:1 compression, is cheap enough for some applications such as page proofing, but it's also fast enough to use on video. And full MPEG hardware, though expensive and not backed by a firm standard, can give much better compression. There should be a solution in there for just about everyone. ■

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

C-Cube Microsystems
(408) 944-6300 ..... **Circle 258**
LSI Logic
(408) 433-8000 ..... **Circle 259**
Next Computer
(415) 366-0900 ..... **Circle 260**
Oak Technology
(408) 737-0888 ..... **Circle 261**

INTEGRATED CIRCUITS

TTL, ECL face off over high-speed logic

Ron Wilson, Senior Editor

All of a sudden, we seem to need zero-nanosecond logic. As microprocessor clock frequencies go through the roof, designers have less and less propagation delay to budget for glue logic. Even an extra nanosecond in a logic stage can force the design into slightly faster—but dramatically more expensive—SRAMs. So a little spent on faster logic parts can save a great deal on cache RAMs.

But finding the faster logic to buy is a growing problem. At speeds in the 2- to 5-ns range, the inherent problems of CMOS logic are all but intractable. Even good old bipolar TTL—now reincarnated as blazingly fast BiCMOS—presents design issues that many teams are ill-equipped to handle. So some managers are taking another hard look at the long-time king of high-speed logic: ECL. And ECL vendors are rolling out a new generation of logic parts to reduce the entry barriers for migrating TTL users.

Nowhere have the problems of high-speed standard logic been more evident than in the CMOS families. Because of the inherently fast edge rates of CMOS drivers, the parts exhibited most of these problems—ground bounce, switching noise and transmission-line effects—before TTL parts did. But now it's evident that the CMOS devices just showed the problems first—they didn't have exclusive rights to trouble.

"Right now edge rates are the main concern," says Suneel Rajpal, director of marketing and sales at Quality Semiconductor (Santa Clara, CA). "At 33 to 40 MHz, people are still depending on CMOS logic with TTL I/O, but they have to be really careful. Not only is there the issue of ground bounce, but now there are transmission-line effects. When the rise or fall time of the driver gets shorter than the round-trip delay in the trace, you have to treat the trace as a transmission line, or you have to wait for it to stop ringing."

Coping with CMOS

This news hits designers like a one-two punch. After going through several years of package changes (corner ground to center ground, for instance), arcane layout rules and add-on load resistors to minimize ground bounce on CMOS parts, now we find there are more problems. In

fact, the very measures that helped most with ground bounce seem to worsen the transmission-line problems. "We have gone from 5-V logic swings down to 3.5 V to control ground bounce," Rajpal says. "Now I see that at the last JEDEC meeting, Texas Instruments (Dallas, TX) proposed 2.5-V logic. That would mean you would need to be very careful loading the outputs." But with lower output currents and noise margins, low-voltage gates might not be able to drive properly terminated transmission lines.

Quality Semiconductor aggressively attacks the CMOS problems with on-chip remedies. Its latest parts use reduced voltage swings, have on-chip load resistors and control edge rates on the drivers. All of these features make the parts considerably easier to use, but they can't take away the fact that 50-MHz board designs are hard work. "People want a 3-ns part with no bounce," Rajpal says, "and it just doesn't exist. You have to be more conscious of board layout, and at higher speeds, you have to do Spice simulation after layout. We're getting more requests from our customers for Spice models than ever before."

In their struggles with noise issues, customers are demanding not just Spice models but increasing vol-

QUALITY'S ON-CHIP SOLUTION

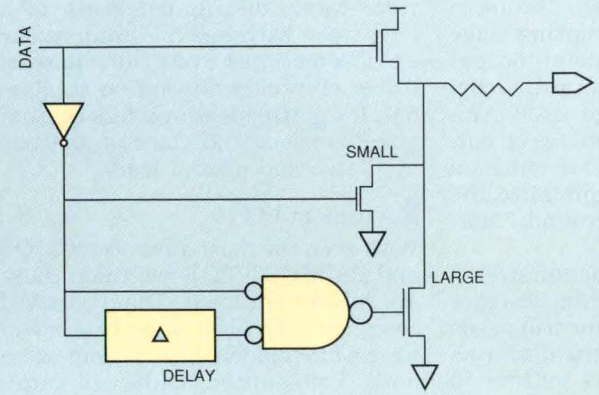

This output stage from Quality Semiconductor's FCT2000A family illustrates the use of two methods to contain switching noise problems. The combination of one small and one large output transistor reduces the edge rate, while the built-in 25-Ω series resistor manages current.

INTEGRATED CIRCUITS

umes of numerical data as well. "As parts get faster," says Michael Polacek, product marketing manager at Signetics (Sunnyvale, CA), "designers are finding that their boards don't work because of questions they didn't ask at the beginning. I mean things like, 'Is this part slower with four outputs switching than it is with one output switching?' Consequently, we're putting a lot of characterization data—stuff that we know about the part but don't actually guarantee in manufacturing test—onto the data sheets now."

Signetics is starting this policy on its new family of BiCMOS bus drivers, multiple-sourced via an agreement with TI. Starting with devices such as the ubiquitous '245 octal transceiver, the two companies are using BiCMOS to produce delays in the range of 4 to 6 ns with 64-mA drive capability and ground bounce typically under 1 V.

The new parts have gone to unprecedented lengths to maintain speed while reducing noise, according to Randy Morgan, applications engineering manager at TI. "You have to solve the problems on-chip now," Morgan says. "With the performance and space requirements designers are facing, you can't just ask them to add resistors to their board."

TI has not only taken the more traditional path of limiting edge rates through the bipolar output drivers, but it has significantly redesigned the driver stage itself. "We use circuitry that anticipates an output switch and shunts current from storage capacitors to minimize the current transients on ground," Morgan explains.

Even with these techniques, a vendor can't get the system designer off the hook. The designer still needs to understand the electrical characteristics of the load he's putting on the output. "No one family of parts addresses all the different driver needs in a system," says Morgan. "You may have a backplane that demands very high I^{OL}, or a bank of DRAMs that can't tolerate undershoot, or a Futurebus+ that's a whole different issue. In the worst case, every trace has to be considered separately. Do you have a short trace with evenly distributed loads, or a long trace with one big lump at the end? Sometimes even the choice of driver chip may depend on the layout. It's getting less and less reasonable to expect that you can just pop in a latch and have it work."

A look at ECL

With even the most advanced CMOS and BiCMOS TTL logic, then, there are still serious issues that the board designer must face. Some traces will have to be modeled as transmission lines. Unbalanced loads will cause

reflections, requiring settling time. The proper operation of the circuit may depend on parameters that have never before published in TTL data sheets, such as the temperature influence on thresholds and ground bounce. If all of this sounds like a nightmare to TTL designers, it sounds comfortably familiar to another group.

"ECL designers have always been aware of complex timing issues," says Joe Vithayathil, marketing manager at Synergy Semiconductor (Santa Clara, CA). "They've been working with specs like temperature dependencies, minimum/maximum ranges and multiple-output switching situations since the beginning. But the TTL guys are just discovering this stuff."

The nature of ECL I/O—switching current through nonsaturating transistors at relatively low voltages—makes the technology inherently more suitable to high-speed applications. And the ECL community's long experience with high-frequency board design has built up an infrastructure for handling the problems. "There's a lot more data on ECL data sheets," Vithayathil says. "And since ECL users started out as a pretty small community, there's a strong tradition of working closely with the vendor to characterize the devices you're using."

"ECL designers have always been aware of complex timing issues. But the TTL guys are just discovering this stuff."

—Joe Vithayathil, Synergy

Many designers admit that ECL offers a superior solution to the fastest design issues. "At these speeds, ECL can be much easier to handle at the board level," says Quality's Rajpal. "But the cost of ECL and the lack of standard parts have been a real problem for many people. I know a number of designers of large CPUs who did ECL vs. TTL evaluations and decided to do another TTL design based on cost/performance."

In fact, there haven't been, until recently, many of the major system building blocks available in ECL. ECL memories tend to be small and dear. Until last year, when Bipolar Integrated Technology (Beaverton, OR) announced ECL versions of the Sparc and MIPS CPUs, there were no ECL microprocessors. "But now,

INTEGRATED CIRCUITS

if you look above 50 MHz, there are no TTL CPUs—you have to go with ECL," Rajpal points out.

To meet what they see as a crystallizing mass market for ECL logic, some vendors are growing increasingly active. "If you're looking at the future of ECL, you have to ask yourself why Fujitsu, Hitachi and NEC are all in it, and why Sony is showing so much interest," suggests Vithayathil.

In fact, Synergy itself is starting to move. Using its 1.5-μm process that has produced 70-ps internal gate speeds and 3-ns RAMs, the vendor has decided to alternate-source Motorola's new, updated ECL family, Eclinps—a clever acronym for ECL in picoseconds.

"People were getting the latest microprocessors, RAMs and ECL gate arrays, and then hooking them up with antique glue logic," Vithayathil says. With maximum gate delays in the range of 500 ps and latch t_{pd}s around 1 ns, the new logic is certainly fast enough to work with leading-edge ECL LSI parts. In addition, this kind of speed can be a benefit in mainly TTL designs. "You can put ECL RAMs in a TTL system, and, even with two sets of level translators, it will be faster than a TTL SRAM," Vithayathil says.

The new family also addresses some of the practical issues that have kept designers away from TTL. The parts are voltage- and temperature-compensated and have 2-kV electrostatic discharge protection—ESD was a serious weakness on earlier ECL parts. And with the significantly lower power dissipation made possible by Synergy's process, the ECL devices will be available in plastic packages. That should make a big dent in ECL's historical cost problems.

As system speeds increase beyond 50 MHz, we may very well see the shift for which ECL designers have waited all these years. But ironically, people may change not because TTL won't do the job, but rather because the ECL system was easier to implement. With the increasing complexity of TTL design issues and the old impediments to large ECL systems slowly dissolving, it may be an ECL future after all. ■

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

Bipolar Integrated Technology
(503) 629-5490 ..... **Circle 253**
Quality Semiconductor
(408) 986-1700 ..... **Circle 254**
Signetics
(408) 991-2000 ..... **Circle 255**
Synergy Semiconductor
(408) 730-1313 ..... **Circle 256**
Texas Instruments
(800) 336-5236, ext 700 ..... **Circle 257**

COMPUTERS AND SUBSYSTEMS

First STD 32 boards target fast I/O applications

Warren Andrews, Senior Editor

Last month's Buscon/90-East saw the first STD 32 boards, backplanes, connectors—in fact, everything needed to assemble a complete system—come together under one roof in a formal debut. The focus, not unexpectedly, was I/O. STD has become widely known as an I/O bus, having perhaps the widest variety of I/O cards available—from conventional digital I/O to pneumatic switches mounted directly to a card.

But the development of STD 32 is bringing a new dimension to the basically 8-bit STD bus. While many in the STD community continue to claim there's no need for anything more than an 8-bit I/O channel, the latest introduction of STD 32 products, including the first ever 16-bit STD peripherals, belies that statement. And, of course, the power to drive the I/O—the CPU—stands in the forefront.

Ziatech (San Luis Obispo, CA), developer and most ardent supporter of the STD 32 specification, leads the parade of new STD 32 parts with a high-performance 286-like CPU card with 48 channels of high-speed digital I/O. At Buscon, Ziatech also showed off backplanes, connectors, cabinets and other components necessary for system implementation.

Versalogic (Eugene, OR) is offering a high-performance analog-to-digital converter board featuring full 16-bit capability across the backplane. And Technology 80 (Minneapolis, MN) has developed a sophisticated motion controller that takes advantage of some of the special EISA-like features STD 32 has over the more-conventional STD 80.

EISA compatibility

In addition, with the formal introduction at Buscon of the now-commercial connector, there's yet another surprise—the extension of the connector from 114 to 136 pins. "The connector has turned out to be a bigger advantage to STD 32 than we envisioned at first," says Ziatech vice-president Jim Eckford. The additional pins, he says, allow for a full EISA extension, increasing the memory addressing to the full 16 Gbytes called for in the specification. And the mounting configuration for both the old STD 80 and the new connector are identical, so in volume applications, STD 80 and STD 32 connectors can be used on the same backplane, Eckford points out.

But the EISA extension on the connector does more than simply increase the addressing range, Eckford says. By having the same pin count and configuration as the standard desktop EISA, STD 32 is fully compatible with any other EISA configuration. So any EISA board function can be adapted directly to STD with no further multiplexing or byte-lane swapping. This, believes Eckford, will make it easier and more attractive for conventional EISA makers to develop STD 32 boards. In addition, the connector remains compatible with all previous STD 80 boards.

At the heart of any system, of course, is the CPU function. The first full STD 32 CPU—though only 16 bits—to be offered is Ziatech's latest and is based on a 286-like processor, the NEC V-53. The V-53 incorporates the CPU function along with many of the peripheral functions, much the same way the 80188 included many of the functions of the 8088. The chip, though running at 16 MHz, performs like one much faster because the peripherals are internal, and signals don't suffer from delays getting on- and off-chip, says Eckford.

In addition, the CPU board includes 48 bits of digital I/O, taking advantage of Ziatech's first ASIC. It also contains three serial ports, 1 Mbyte of RAM, and an iSBX socket to plug in additional functions such as a graphics or communications subsystem.

Another Ziatech product adding to the STD 32 debut is a slot zero arbiter card that manages arbitration when more than a single bus master is used. The function, which would normally reside on the motherboard in a conventional EISA system, stands on a separate—though sparsely populated—card in the STD environment. "We could have included the function on the backplane," says Eckford, "but it would have added to the complexity and size of the backplane, and in many instances wouldn't be used." Similarly, the function could have been added to the CPU, but more than one CPU card used in a system would be redundant.

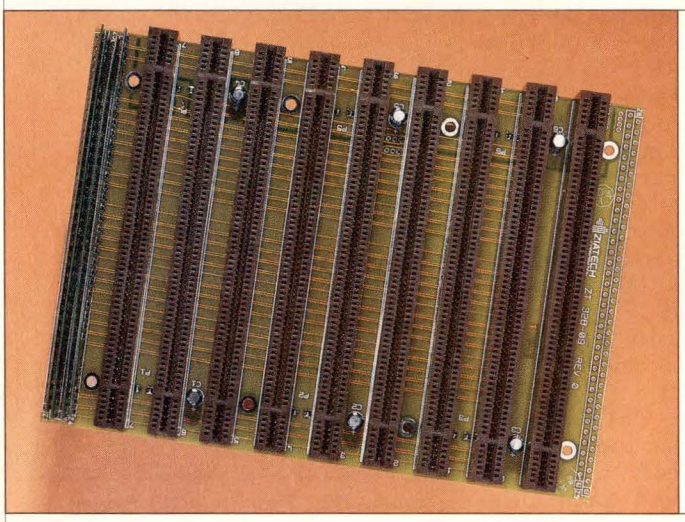

Ziatech's ZT 32B is the first STD 32 backplane on the market. It accommodates both older STD 80 and newer STD 32 boards, allowing 8-, 16- and 32-bit data transfers.

Bus mastering advantage

What the arbiter board offers—one of the key advantages of STD 32 over traditional STD—is the concept of bus mastering. This means that

COMPUTERS AND SUBSYSTEMS

more than just the main CPU can control the bus and the system. Though not critical in some traditional STD applications, bus mastering will become increasingly important as STD moves into more sophisticated applications, says Eckford. These applications can include complex robotics, in which each of a number of processors may be working on a single motion axis.

In addition, in many industrial-control applications, it's becoming increasingly important to have a DOS window into what's happening in a given environment. A programmable logic controller may be running ladder logic for machine control on one master, for example, while another provides a DOS environment for data entry and monitoring on another master, says Eckford. And though such multimaster approaches often call for complex programming, Ziatech is developing systems to minimize that effort.

"DOS MPX is a product we offer to support multimastering on STD 32. This software lets programmers make calls to masters at a high level, instead of requiring the programmer to deal directly with memory addressing at the assembly level," says Eckford. "Systems such as DOS MPX work with the slot zero arbiter to provide a smooth multimaster environment."

In addition to the CPU and arbiter cards, Ziatech also introduced at Buscon an Arcnet card, the first on STD to provide 16-bit communications on the backplane, and the fastest yet on STD. The company also showed a digital I/O card and a 9- and 15-slot backplane.

Fast 16-bit I/O

Versalogic will add lower-slot-count backplanes, along with its I/O boards, to the STD 32 market, although the products weren't announced at Buscon because they weren't formally released yet, says Gary Harris, Versalogic vice-president of engineering.

The backplanes aren't Versalogic's main STD 32 offering, says Harris. "Our first STD 32 board is a very high-performance analog card that needs the full 16-bit bandwidth afforded by the STD 32 specification. Actually, the board will operate in either an 8- or a 16-bit slot, but with a 3-μs conversion capability and high-speed DMA, it's more comfortable in an STD 32 slot," he says.

Although DMA is specified in both the 8- and 16-bit modes and can be selected on the board, the greater bandwidth is a major advantage, Harris says. In addition, the slot-specific interrupt feature of the STD 32—and EISA—greatly simplify wiring the board. "In the past, we used to have to run cables across the front of the card cage to get the signals where they belonged. Now we can bring other signals out the front where they belong," he says.

STD 80 VS. STD 32

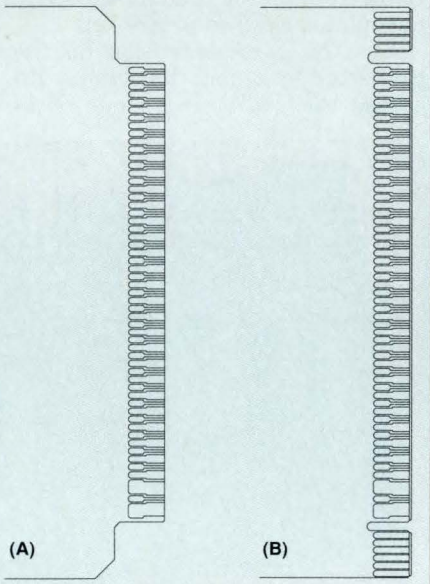

The new STD 32 connector (B) adds an extra set of wings to increase the original 114-pin count to 136. The additional pins represent the extra addressing range of EISA specification, making the STD 32 now fully electrically compatible with EISA. The connector is also upward- and downward-compatible with earlier STD 80 implementations (A).

Although the board is the first true STD 32 peripheral to use a 16-bit transfer across the backplane, says Harris, "It's really a simple I/O card. It's designed to be fully software-compatible with Analog Devices' STD 80 offering, the AD 1265; as well as with its PC card offering, the AD 820." He stresses, though, that Versalogic's analog I/O card is just a starting point for exploiting the advantages of STD 32. He envisions branching out into developing even faster converters using the full available 32 bits of backplane (once CPUs arrive). "Perhaps the next step out will be something in the digital signal processor area, which will provide the analytical horsepower for the converter products," he says.

Two-axis motion control

While Versalogic may hold the lead in data conversion, Technology 80 was the first to show a full 16-bit I/O card for two-axis motion control. Motion control systems in robotics, simulators and a variety of other applications are becoming increasingly sophisticated, according to Jim Burkett, vice-president of Technology 80. With only an 8-bit backplane and limited bandwidth, STD has been restricted from participating in many of the newer applications.

STD 32 will now be able to compete with some of the PC- or VME-bus-based systems, but in a smaller, rugged form factor and at a lower cost, says Burkett. "Our first product is only a 16-bit board. We haven't taken full advantage of the bus yet, but we wanted to get our feet wet," he says. "Moreover, until there are some 32-bit CPUs on the market, a 32-bit peripheral won't show much advantage."

Burkett admits that while a 32- or even 16-bit bus isn't called for in many traditional uses, applications are growing more complex at exponential rates. "We're already seeing applications that we thought would never require more than 8-bit performance moving to 16 and 32 bits," he says. Computer numeric control tools, for example, are now calling for complex number crunching and are passing a lot of information around in order to machine complex three-dimensional shapes.

"In addition," he continues, "even if a particular application doesn't currently call for 32-bit performance, the handwriting is on the wall that it will be called for in a short time. STD 32 gives STD users the confidence that they're not on a dead-end street. The message is clearly getting out that STD 32 is good for a few more performance gen-

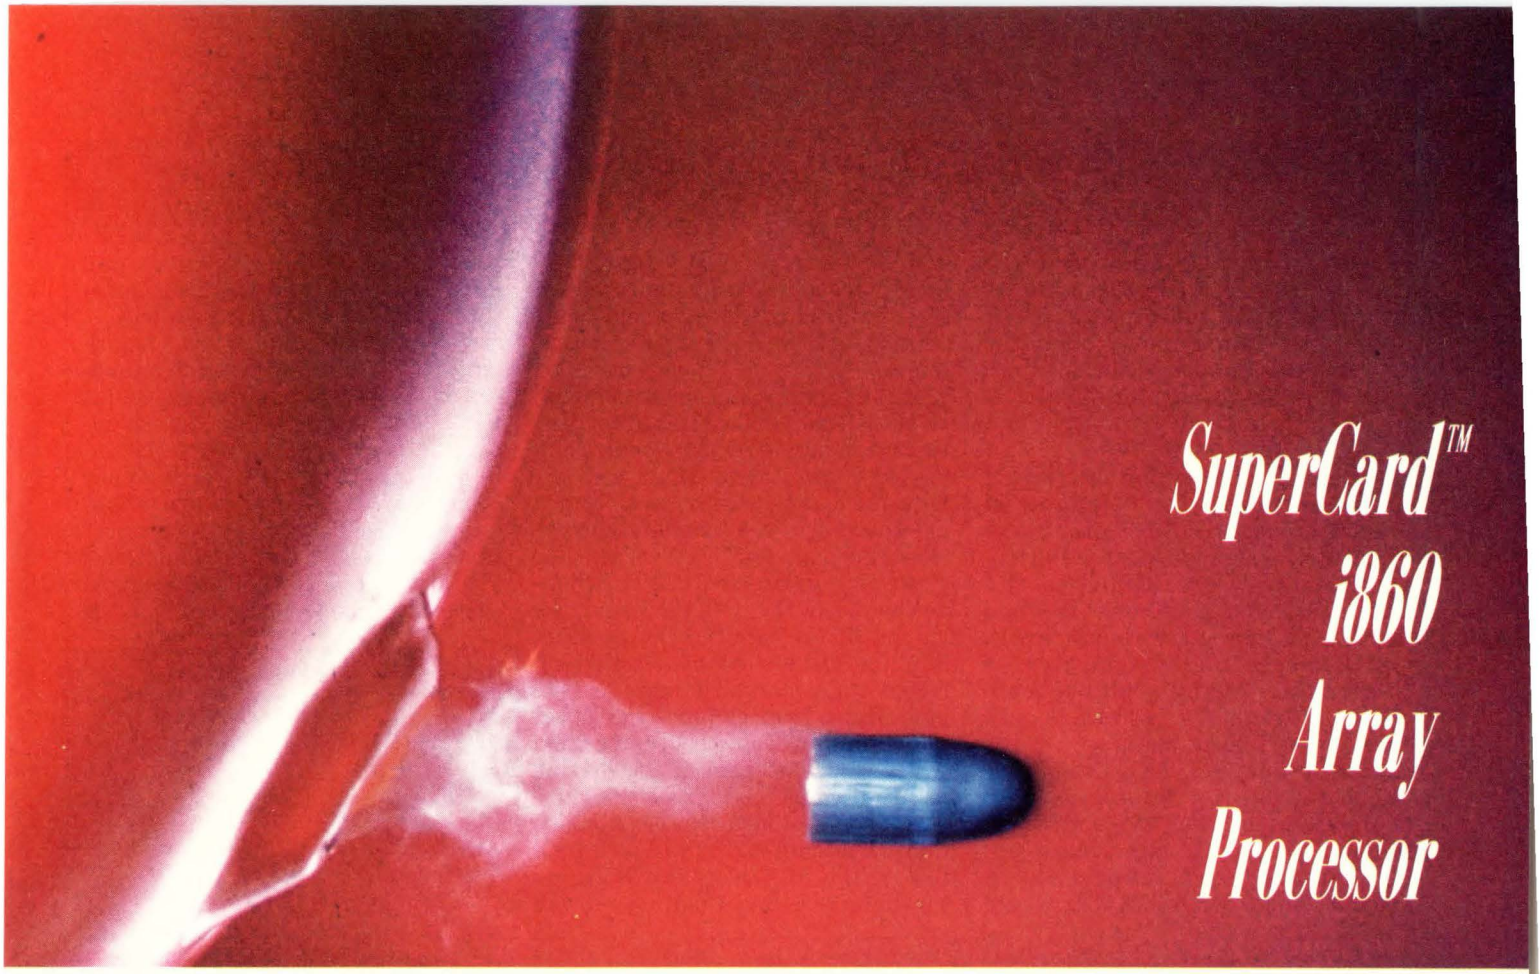

Faster than a speeding bullet . . . . . .

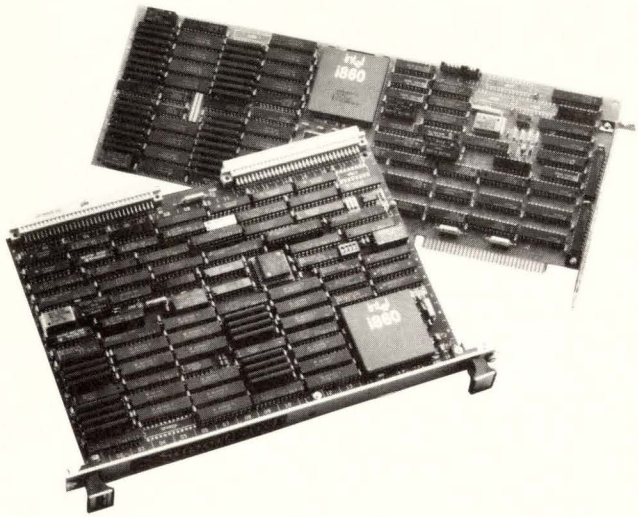

A millionth of a second . . . that's what it takes to stop a bullet in mid-flight. Yet in a millionth of a second, SuperCard can complete up to 640 floating-point operations.

SuperCard . . . an entire family of array processors based on Intel's revolutionary i860 processor. . . from CSPI, a pioneer in array and signal processing since 1968.

- **Configurable from 66-640 MFLOPS**
- **Direct I/O with industry standard VSB bus**
- **Up to 80 MFLOPS and 32 Mbytes in a single slot**
- **Support for most Unix and real-time kernel hosts**
- **Fully programmable with C and FORTRAN compilers**
- **Extensive signal and image processing libraries**
- **VAST-2 vectorizing FORTRAN pre-compiler**
- **Available on VMEbus, ATbus, EISAbus**

Call: 1-800-325-3110 or 617-272-6020, (Fax 508-663-0150) or write CSPI, 40 Linnell Circle, Billerica, MA 01821.

CIRCLE NO. 85

CSPI
THE ARRAY PROCESSORS

Top photo: ©1989 courtesy Palm Press. A .22 caliber bullet erupts from a toy-store balloon in a crater of gas and dust.
i860 is a trademark of Intel Corp. Unix is a trademark of AT&T.
VAST-2 is a trademark of Pacific-Sierra Research Corp.

COMPUTERS AND SUBSYSTEMS

erations." Burkett adds that the close tie with the EISA community may well turn out to be a major advantage—particularly with the full upward and downward compatibility the new connector offers.

The competitive picture

The recent introduction of STD 32 hardware is just the beginning of what promises to be an exciting generation of STD products—both hardware and software. The close alliance between conventional ISA and EISA, and the ongoing romance with the DOS operating system, promise to offer a continuing wealth of applications programs, software and hardware enhancements.

In addition, as EISA solutions are forced to vie in a competitive environment with other PC and workstation solutions, significant enhancements are expected. A second-generation EISA that would double the transfer rate from 33 to 66 Mbytes/s is already under discussion. Additional developments may be in the works for further enhancements.

EISA is in direct competition with IBM's Micro Channel Architecture. Both are aggressively looking for market positions—particularly within the workstation area, which is expected to see dramatic growth over the next four years. IBM has staked out its area with Micro Channel in both the PC PS/2 and the workstation (6000/AIX) product families. Others companies have selected other standard buses—VMEbus, SBus, Turbochannel and EISA.

The fallout from this competition will be at least part of the driving force pushing STD 32 into higher-performance, more-complex applications. Undoubtedly it will pick up a following of Unix applications through its EISA affiliation, as well as real-time DOS applications and other software that will filter down through EISA or up through ISA.

While last month marked a beginning for STD 32, it was only a modest one; still to come are the 32-bit CPUs and peripherals. Ziatech already has what it believes to be the first of the 32-bit CPU cards under design, an 960-based board—"Unless someone beats us to the punch," says Eckford. "I understand there are at least a few 80486 designs currently under way." ■

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

Technology 80
(612) 542-9545 ..... **Circle 232**
Versalogic
(503) 485-8575 ..... **Circle 233**
Ziatech
(805) 541-0488 ..... **Circle 234**

NEW...Warrior-III

The ultimate subroutine machine.

The ultimate. The best. A board level array processor whose i860, and i960 combination make it the most megaflop friendly array processor available. It is a full three to six times faster than its predecessor the Warrior-II. The two internal processors are supported by either fast SRAM or cached DRAM memory. Coupled with a 160 Mbyte per second internal bus you've got a real screamer. 80 Mflops, 40 MIPS peak performance.

Single slot 6U and 9U VME

Supercomputing on a single slot board. Memory from 1/2 to 64 Mbytes. The board includes a VSB bus interface and an additional expansion port. SKY also offers standard I/O daughter cards such as a 32-bit parallel interface. Or if you like, you can customize your own interface. And, multiple Warriors can be used without special cabling.

Real DMA

The i960 executive controls a powerful DMA engine capable of addressing host memory and external devices. There is no tie up of the i860 to accomplish this. And, since internal interrupts don't affect the i860, you get more usable megaflops. Real horsepower to solve problems like image processing, signal processing or matrix manipulation.

Software support

Add to this built in source code compatability with the Warrior-II library, advanced features like command chaining, user built command blocks, and tools to do custom routines. The result, the most competitive board level array processor on the market.

Want to find out more?

Circle the bingo, and we will send you more data or give us a call. We can have a SKY technical representative contact you immediately.

SKY Computers, Inc.
27 Industrial Ave.
Chelmsford, MA 01824
Tel. (508) 250-1920
FAX: (508) 250-0036

CIRCLE NO. 34

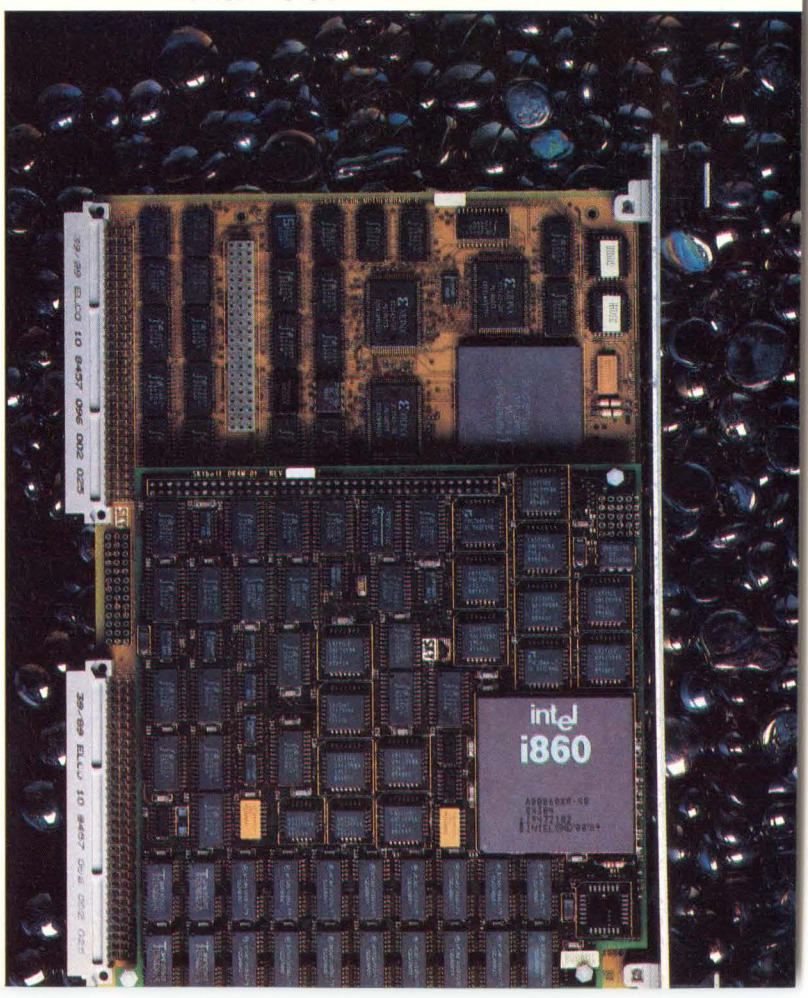

Signetics. Because com
isn't just a product

COMPUTING					
APPLICATION	**PRODUCT**	**APPLICATION**	**PRODUCT**	**APPLICATION**	**PRODUCT**
Workstations	● Advanced BiCMOS Logic (ABT) ● High-Speed ASICs ● Futurebus+ Chip Set ● High Speed PAL®-type Devices ● High Performance MCUs	Desk Top Video Personal Computers	● Video Data Converters ● Digital Color Decoders ● High Density ASICs/PLDs ● DRAM Controllers ● OTP EPROMs ● FLASH Memory	Peripheral Products	● 8-bit 80C51-based MCUs ● Zero Power PLDs ● Programmable Sequencers ● 3-State ECL Transceivers

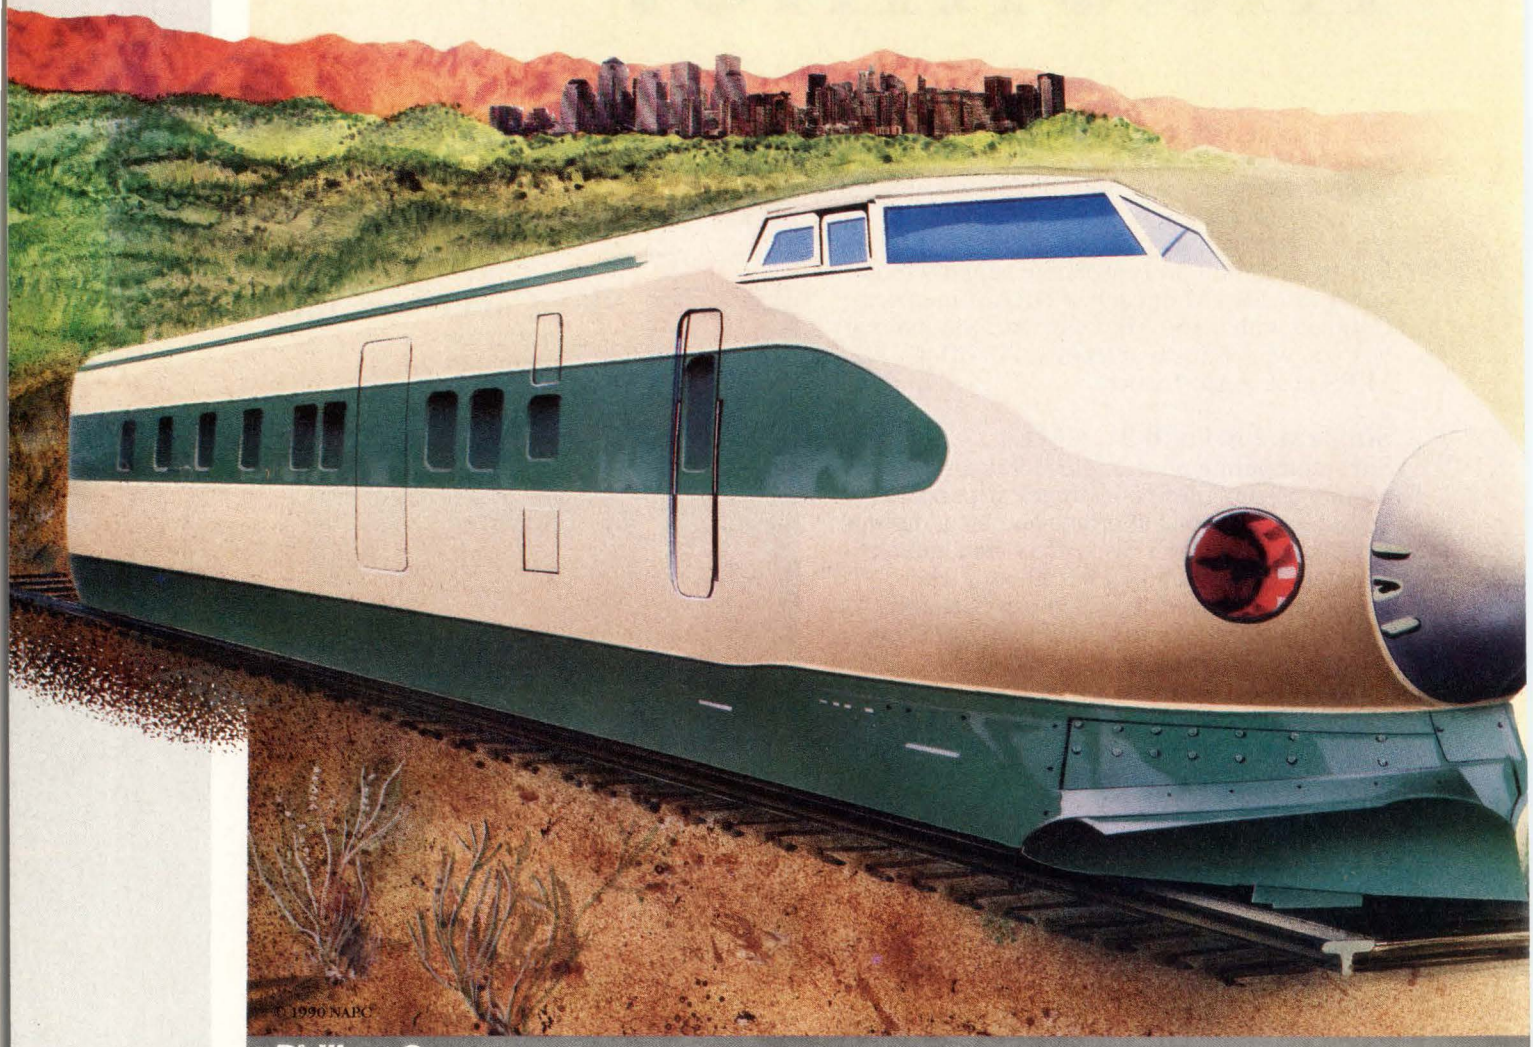

Philips Components

puting performance of your engine.

WITHOUT THE RIGHT SELECTION OF ICs, YOUR CPU COULD BE DEAD IN ITS TRACKS.

Today, everyone is using the same processors.

So, to separate your computer design from the rest, you need to get the full potential out of your processor. Full potential that's only possible through high-performance supporting subsystems.

To give you this performance we're offering a full range of ICs for major subsystem applications. Together they help you get the most from your processor, so your designs perform like never before.

For example, our bus interface logic devices are the industry's fastest. With our proven BiCMOS process —known as QUBiC—our ABT logic family is nearly twice as fast as the highest performing Bipolar ICs. This means you get the speed to keep pace with today's 16- and 32-bit systems, as well as tomorrow's emerging performance standards.

We also offer you a complete family of advanced PLDs. Including a full range of PAL®-type devices with speeds from 4ns to 7.5ns. As well as our innovative programmable logic arrays (PLA), programmable logic sequencers (PLS) and programmable macro logic (PML).

When you need microcontrollers, we offer the industry's most complete selection. Including devices from 8- to 32-bits and in OTP, EPROM, ROM and ROMless versions. Features include I²C serial bus, low voltage/low power, A/D, extended memory and more. All for EDP peripheral applications ranging from keyboards, disk drives and printers, to terminals and mouse devices.

And for desktop video applications, we've applied our expertise in digital video signal processing to offer you an 8-bit digital multistandard TV decoder subsystem, complete with data conversion and clock companion chips.

So when you need subsystem performance that lets your high-speed processors move at top speed, choose Philips Components–Signetics. For more information or our Computing Brochure, contact us today: 800-227-1817, ext. 713F.

PAL is a trademark of AMD/MMI

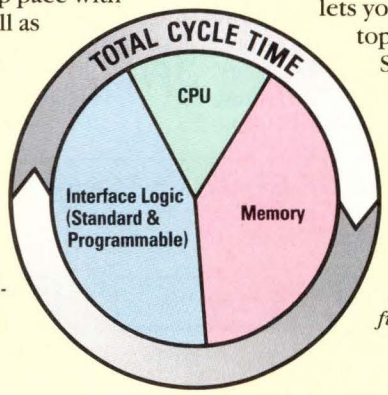

As processor speed increases, and total cycle time decreases, the percentage of time spent performing interface functions becomes more significant.

Signetics

EXTENDING THE DIMENSIONS OF PERFORMANCE

PHILIPS

CIRCLE NO. 35

PHILIPS

COMPUTERS AND SUBSYSTEMS

Multimedia graphics board eases human interface to VME systems

Warren Andrews, Senior Editor

High-resolution VMEbus graphics—not to mention those which include pointing device interfaces, audio functions and keyboard interfaces—have been relatively slow to appear in real-time control applications. As a result, human interfacing—entering, viewing and interpreting data—has been a complex task or has required additional hardware and often a significant software effort.

An experienced operator, for instance, has to make sense out of lines of data on a terminal and enter memorized control codes through a keyboard. The alternative is to provide a separate set of subsystems with their own software as windows into the real-time application. The set of subsystems is frequently made up of more than one card and can result in a programming headache.

"But," says Harry White, president of graphics board maker Vigra (San Diego, CA), "with emerging graphics systems such as X Windows and Microware Systems' Rave package, combined with new multifunction graphics boards, real-time VME systems can provide the same or better-quality graphics and human interface as systems requiring additional hardware and software.

"Through clever use of the VMEbus platform," he continues, "even high-resolution graphics, along with other functions, can be accommodated with little or no penalty to the VME system bus performance."

In the industrial arena, says White, the key to simplifying application development is to have some graphic interface to standard systems. But that isn't enough, he maintains. The graphics function must be accompanied by input ports for a mouse, trackball or touch screen, must provide a keyboard interface, and must offer some kind of tone-generation capability to alert operators of various conditions.

Vigra's latest contribution to this multimedia arena is the MMI-250, a VME color graphics controller that combines high-resolution bit-mapped graphics, keyboard and pointer inputs and a sound generator on a single 6U VME board. In addition, the board has a full VME subsystem bus (VSB) interface for applications that require high-resolution graphics without penalizing the system bus.

A perfect match

"Vigra's is the only VME board we've seen that, when coupled with a package like Rave, provides a standard I/O interface ideally adaptable to nontechnical users," says Drew

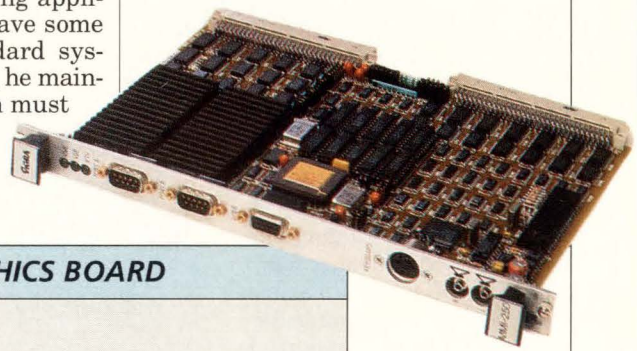

Vigra's MMI-250 is a multimedia graphics board whose functional subsystems tie together to provide a complete human interface for a VMEbus system without having to integrate another entire system.

THE VME HIGH-RESOLUTION COLOR GRAPHICS BOARD

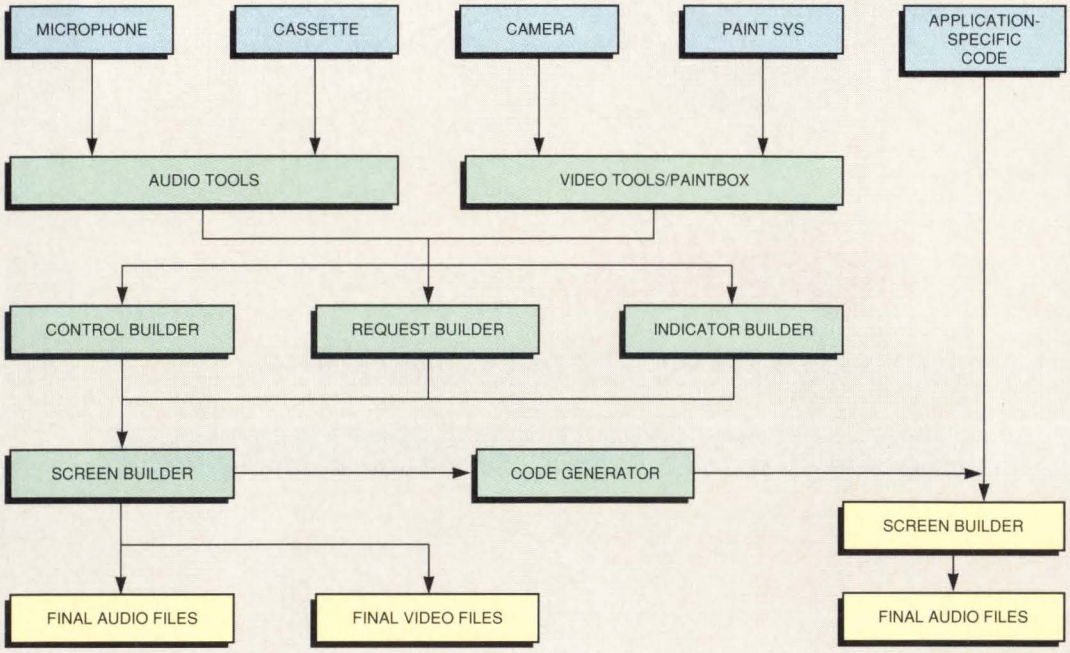

Crane, strategic marketing manager for Microware Systems (Des Moines, IA). Rave (real-time audio video environment) is a multimedia development tool and user interface that greatly simplifies the design of realistic man-machine interfaces for real-time process-control systems, says Crane. It lets designers combine high-quality audio and video, computer-generated graphics, and customizable menus in the same user interface.

cause the software just wasn't there before the introduction of Rave and X Windows—sort of a "chicken or the egg" situation. To perform functions equivalent to what can be done in a Rave or X Window environment, programmers used to have to spend countless hours programming individual bit-mapped images.

In most cases, it can take as long to write the interface portion of a program as it can the application part. And since the interface portion is often done last, it's frequently shortchanged—which explains why interfaces tend to be sloppy at best, and nearly unusable (except to the programmer who wrote them) at worst.

THE RAVE MULTIMEDIA TOOL

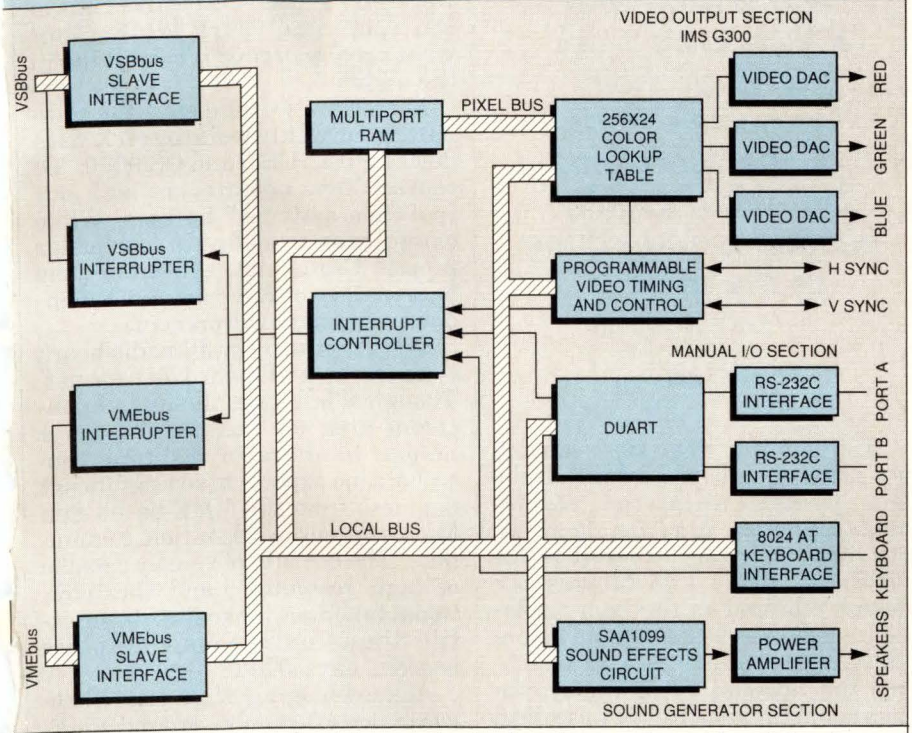

Microware Systems' Rave (real-time audio visual environment) provides all the elements required to develop and run a full multimedia interface. Designers can combine computer-generated graphics, audio and video and customizable menus.

The growth of Rave in the VME world, Crane continues, is expected to largely parallel Vigra's success. "Before Vigra's MMI-250 board (and its earlier lower-resolution graphics multimedia card), Rave was sitting there with no hardware to run effectively on VME," he says. But from early responses, he expects to see a major surge in interest.

While Microware Systems laments the shortage of multimedia hardware, Vigra's White maintains that multimedia boards haven't seen much market acceptance be-

To perform the functions Vigra's multimedia board handles in a single package, it's necessary to start with a VME system and then add a high-resolution graphics card. At least one other board for PC-compatible keyboard interface, mouse, trackball or touch-screen interface must be added, and perhaps a third custom board (or sophisticated daughtercard) with tone-generation and audio output capabilities. Special programming may be required to interface the additional boards.

While there are a variety of graphics boards for VMEbus on the market, there are few with any degree of multimedia capability. At least two off-shore VME companies have taken a stab at limited multifunction boards. And at least two major U.S. makers of VME graphics boards, Heurikon (Madison, WI) and SBE (Concord, CA), reportedly have work in progress. They're being pushed by some of their customers into developing mezzanine cards to provide multimedia capabilities for their existing graphics boards.

One of the driving forces behind multimedia products, says White, is the ever-present personal computer. Users see the type of interface that's readily available on their office PC and wonder why their industrial machine—costing orders of magnitude more than the PC—can't provide the same kind of interface. VME system vendors have already seen some of their applications move to PC environments—or have had to incorporate PC-on-VME subsystems for exactly that reason.

The basic multimedia VME card is composed of a video controller, I/O for input devices, and some kind of audio device. "The MMI-250 is the second-generation multimedia card we've developed," says White. The first, the MMI-100, had a full speech processor centered around a Motorola 56000 digital signal processor, but it had only limited graphics capability. The second generation provides far better graphics capability, but not speech capability.

Programmable resolution

To get the high-resolution video function in the MMI-250, Vigra uses an Inmos IMS G300 video controller, which includes a lookup table that provides 256 colors from a palette of 16.7 million. "We use the Inmos controller," says White, "because it greatly simplifies the board design. It drops chip count for the function from three to one, and it eliminates the need to route very high-frequency signals (100 MHz) on the board."

Resolution on the card is pro-

COMPUTERS AND SUBSYSTEMS

grammable, ranging from VGA-compatible 640×480 pixels up to 1,280×1,024 pixels, with each pixel programmable as 1, 2, 4 or 8 bits. In addition, the board has 2 or 4 Mbytes of video RAM that are ported to the video controller, the VME and the VSB. The video memory allows multiple frames to be stored, which lets software designers keep frequently used displays on board for ready access and permits easy screen changes in the background.

"Both Rave and X Windows work well at their respective ends of the spectrum, but what's really needed is something in the middle."

—Harry White, Vigra

The multiple framestore capability also helps to keep traffic off the bus. In applications that require high-resolution graphics, however, it's often necessary to resort to the VSB. At about 1.25 Mbytes per frame required in the high-resolution mode, the board can only store three frames at a time. And refreshing that memory using the system bus could significantly tie up critical real-time traffic.

The VSB interface can therefore handle all window data, says White, leaving the VMEbus free to take care of all application-related manipulation and accesses. The VSB interface, he adds, often has to be invoked when working in the X Windows environment.

In its short two-year history, Vigra has been one of the leaders in sound synthesis, having developed sound effects for almost every VME-based flight simulator product in production. Its multimedia board has stereo audio output and tone generators, letting users program a broad variety of realistic sound effects. The board provides 1 W/channel of audio output (into 8 Ω), and the sound is generated by six tone generators and two noise generators, as well as mixers and envelope controllers. It boasts an audio bandwidth of 30 to 7,800 Hz.

Rave to X Windows

X Windows provides a rich environment for designers, and many programmers prefer to work in it as well, White says. The Rave approach, however, provides a relatively painless and simple multimedia development and run-time environment. "Both approaches work well at their respective ends of the spectrum," says White, "but what's really needed is something in the middle."

Microware's Crane says the company is currently working on X Windows for both OS/9 and OS/9000. "Of course, Rave operates on both our real-time systems," he says. While careful not to define a particular project, Crane hints that there may be a merger of Rave and an X Windows program environment.

Vigra's existing multimedia board is made in a standard 6U format. Though White says the main application area for the board is as a human interface to real-time controllers, he expects to see significant business from the VME-based Sun Microsystems workstation community. "The board is obviously capable of high resolution, and the additional functions make it more versatile than conventional graphics boards," says White.

Acknowledging that the whole world doesn't revolve around VME, White says, "We plan to develop some SBus products in the immediate future. Sun seems to have its graphics capability pretty well in place, so I can't see doing a multimedia board similar to the MMI-250. Rather, we'll be looking at doing audio and speech boards." ■

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

Heurikon
(608) 271-8700 ..... Circle 201
Microware Systems
(515) 224-1929 ..... Circle 202
SBE
(415) 680-7722 ..... Circle 203
Vigra
(619) 483-1197 ..... Circle 204

ICS STACKS THE DECK.

Our VMEbus Multichannel Data Acquisition and Signal/Image Processing Cards are all Aces. Now, ICS-120, the fastest Multichannel Data Acquisition board can be the newest card in your winning hand.

ICS-120

A star is born! Fastest multi-channel Data Acquisition board with 1 MHz /Channel digitization rate and 2 Mbytes of data storage. Each of the 4 input channels has its own S/H and 12-bit ADC, a 128-tap FIR filter, 1-16 decimator, a band-shift multiplier and 256K output storage locations.

ICS-200

A solid member of the club! 384-Tap Video FIR Filter with 4 GOPs processing power for Radar, Communication and Instrumentation applications.

ICS-220

Our sweetheart card! Complete Image Acquistion, Processing and Display board with an RS-170 interface, input frame buffers, real-time 8 x 8 convolver, look-up tables, histogram circuit, pseudo-colour RGB out and MAXbus digital I/O ports.

ICS-105

The ace of spades! Versatile 256-Channel Digital FIR filter for Sonar Signal Processing and Multi-Channel Data Acquisition Applications.

ICS-100

A diamond in the rough! Widely used 4-channel Analog Input and Output board which features up to 400 KHz sampling rate, 12-bit digitization, 255-taps/channel FIR filter, and 4 analog outputs.

Complete Solutions

Find out more about our SCD/DAS-256. Complete 64-Channel (1MHz/Channel) Signal Conditioning and Data Acquisition System; SBF-256 Sonar Beamformer; MCF-256 Digital Filter; and more.

Call Toll-Free
1-800-267-9794
Tel (613) 749-9241 Fax (613) 749-9461

INTERACTIVE CIRCUITS AND SYSTEMS LTD.

CIRCLE NO. 36

DESIGN AND DEVELOPMENT TOOLS

Multichip modules test design tool limits

Mike Donlin, Senior Editor

Though the technology is still in its infancy, multichip modules (MCMs) are touted as an attractive design alternative when traditional methods of mounting ICs to a circuit board prove inadequate. But while the technique of interconnecting multiple bare dice to conductive and dielectric thin films can solve density and performance problems, it poses design dilemmas that can cripple existing CAD/CAE tools. The fine pitch of even the most rudimentary MCM architecture can stymie an autorouter, and the amount of data generated to simulate the performance of multiple ASICs can paralyze most simulators, while the thermal behavior of the tightly packed components hosts pitfalls for the unwary design engineer.

Depending on how the designer looks at it, an MCM's architecture resembles either a very small circuit board or a very large ASIC. Because of this, some engineers have tried to use either printed circuit board design tools, ASIC tools or a combination of the two to face the MCM's unique design challenges. Existing tools, unfortunately, have proven inadequate.

"If you're designing an MCM you're probably looking for high performance, lower power or a reduced footprint," says Milton Buschbom, product manager of high-performance ASICs at Texas Instruments (Dallas, TX). "If you're going for performance, you need to simulate down to the gate level of all the components. Four or five VLSI parts will choke the database of any simulator I've ever seen.

"If you're not pushing timing constraints or worried about internal clock skews, then the issue is the mechanical, electrical and thermal properties of the device. To really analyze these characteristics, you need a suite of interactive tools that use a common database. I haven't seen any that do that either."

TI's Defense Systems Electronics Group has produced one of the industry's most advanced MCMs to date. The module houses 45 devices consisting of EPROMs, SRAMs, digital signal processors and ASICs, all mounted with flip tape-automated bonding technology on a 2-×2-in. substrate. "There are individual tools that can assist in building a device like that with a lot of human intervention and experience," says Buschbom, "but I've been looking for an integrated tool for about four years now, and to my knowledge it doesn't exist."

Help on the way?

Valid Logic Systems (Chelmsford, MA) claims to have the answer to this challenge with the industry's first MCM design tool, Allegro-MCM, which is part of a design environment that includes design capture and logic simulation capabilities. "Up until now, many MCM designers have been making hybrids that are characterized by 5- or 10-mil traces separated by 5- or 10-mil spaces," says Shiv Tasker, director of product marketing in the printed circuit board division of Valid.

"CAD packages are therefore designed to work with these relatively large dimensions. So that the fine pitches of MCMs could be dealt with, we adapted and reworked our Allegro circuit board tool. We changed the database, for instance, from a 32-bit integer capability to a double-precision floating-point, 64-bit design to handle the finer pitches."

Because there are currently several methods for MCM production, including high-density laminated circuit boards, ceramic substrates and modules with deposited wiring on silicon, ceramic or metal substrates, Valid's tool supports different sets of design rules. This flexibility lets designers analyze and make decisions based on the trade-offs of each technology.

Once a design is complete, Allegro-MCM supports the chosen technique by providing multiple manufacturing outputs. Gerber output is available for circuit board manufacturing technology, while GDSII output is available for IC-based designs. Layout chores are handled by Valid's Insight autorouter, which can handle the variable-sized vias characteristic of most MCM designs.

There is debate, however, similar to that for printed circuit board

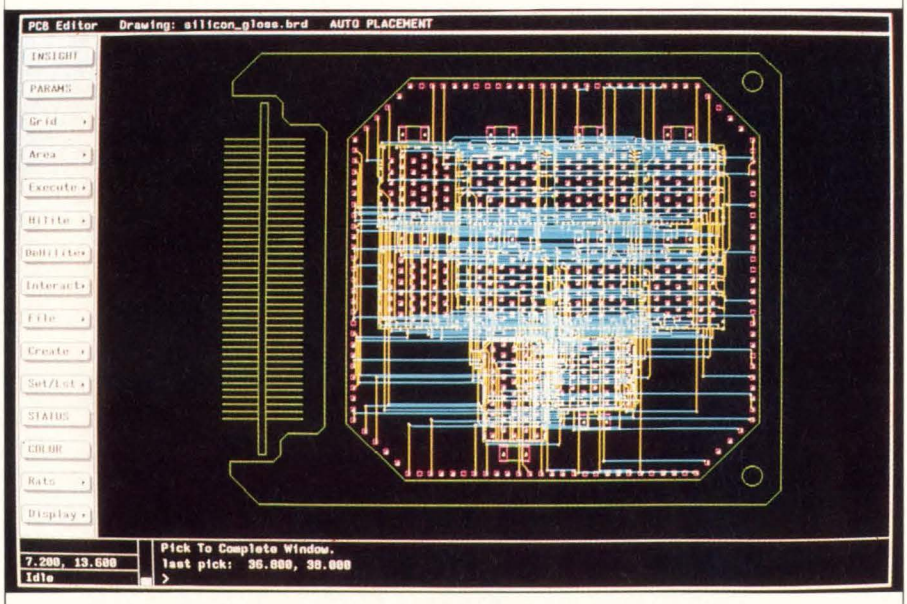

Valid Logic Systems' Allegro-MCM uses the Insight autorouter to simultaneously route up to 48 layers with different line widths. The tool utilizes blind, buried and staggered vias to complete the connections. Here, a two-signal-layer module with ultrafine lines and embedded power and ground planes has been routed to 100 percent completion.

Ethernet. Build with the power of FXP.™

CMC, world leader in VMEbus Ethernet products, has forged a second generation architecture worthy of our stature.

FXP, Full Throughput Architecture. The culmination of over five years' experience and tens of thousands of installed nodes. The power behind the highest performance Ethernet network processors available—the CMC-130 series.

▲ A 20MHz 68020 processor with its own separate program and data memory processes more than 2,000 packets per second.

▲ 256KByte Video RAM as global dual-ported data buffer memory to move packets at full Ethernet and VMEbus data rates.

▲ LANCE™ Ethernet interface reliably handles back-to-back packets to provide full 10Mbps throughput.

▲ Autonomous DMA engine bursts data in block mode at up to 38MBytes/sec providing maximum efficiency to VMEbus host processors with minimal interrupt loading.

▲ CMC's link level driver for the CMC-130 series supports UNIX® host-based protocols. CMC's FDDI products also employ this same driver, allowing developers to use common applications for both Ethernet and FDDI networks.

CMC, the first supplier of intelligent VMEbus Ethernet interfaces, remains committed to supplying fully compliant on-board TCP/IP and OSI protocol suites that off-load the host.

The CMC-130 series is backed by a complete two year hardware warranty, as well as a comprehensive development environment, superior training and unmatched support.

To connect with the power of FXP and join the Ethernet performance vanguard, call today.

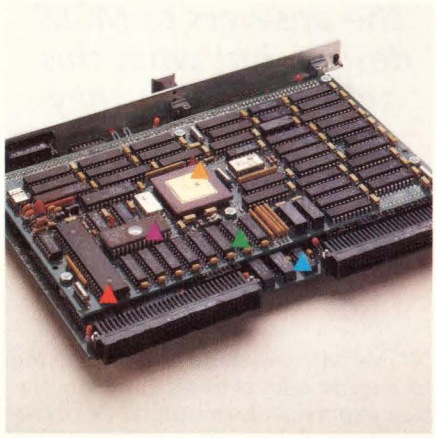

Featuring CMC's onboard TCP/IP.

1-800-CMC-8023

CIRCLE NO. 37

CMC

Designed well. Built well. Rockwell.
Rockwell CMC
125 Cremona Drive, Santa Barbara, CA 93117
PHONE: 805/968-4262 1-800-CMC-8023
FAX: 805/968-6478 TELEX: 240876

FXP

FULL THROUGHPUT ETHERNET

DESIGN AND DEVELOPMENT TOOLS

routers, over which is best for MCM design—gridded or gridless tools. "MCM track widths are going to be at the submicron level," says Hal Barbour, vice-president of marketing at Racal-Redac (Mahwah, NJ). "If you press most vendors who say they have routers that can handle MCM designs, you'll usually find that the routers have a grid, albeit a fine one, of about 1 mil. That means they're married to a channeled array of 1-mil widths. When you force that amount of data into a workstation, it gets tied up fast—there are just too many grid points on a large MCM substrate."

Though originally released as a printed circuit board design tool, Racal-Redac's Visula router can handle pitches down to 0.01 μm, a capability which positions it for handling MCM designs, but which far outstrips any available silicon technology. "Our system doesn't use a channeled approach or Lee's algorithm," Barbour says. "Instead, Visula uses a track-and-space approach in which each area is considered for routability based on rules governing how close an etch can be to any peripheral obstacle. This gives us not only fine pitch but the flexibility to handle the varying line spacings that exist on MCMs."

Vendors of gridded routers counter such claims by citing the speed advantage that gridded routers have traditionally held over gridless tools, as well as easier to use interactive capabilities. "I'll take issue with any vendor who claims that MCMs can only be routed with a gridless tool," says Valid's Tasker. "First of all we have techniques embedded in our tool, such as variable grids and off-grid routing, that offer the same flexibility as gridless tools. Second, since most editors work off a grid, a designer has a much harder time completing a design if the layout has been done by a gridless tool."

Though gridless routers seem to have the advantage when placing submicron pitch designs, gridded router vendors point out that their tools can handle today's MCM technology, which is still in the sub-mil pitch category. It's clear, however, that as MCM technology moves into submicron geometries, gridded router vendors will have to adapt their tools to keep up.

In addition to posing place-and-route problems, MCMs offer testability challenges that will force designers to produce working prototypes. "As we work on our top-down tool for MCMs, one of our biggest concerns is in the assembly and test area," says Jim Behrens, vice-president of packaging interconnect at Cadence Design Systems (San Jose, CA). "Remember, you're going to have ASICs that are difficult to test on a chip level, and there may be 50 of them on a 4-sq-in. package. Trying to test, rework and assemble

"A lot of people are claiming that they have the answers to MCM design, but what this tells me is that they don't understand the problem."

—Milton Buschbom, Texas Instruments

■■■

MCMs with any kind of yield is going to require a lot of design and simulation expertise. Essentially, we're facing some of the same problems today that we experienced about seven years ago with IC design when we needed an accurate simulation before prototype."

Feeding the simulator

The key to a successful prototype, of course, is in-process analysis tools that produce data that can be back annotated into a simulator before a silicon commitment is made. Because MCM technology is new, though, the transmission-line effects and thermal problems that designers encounter will still be unpredictable and thus remain a challenge to tool vendors. For now, most developers are relying on tools that have proven successful in determining these effects in circuit board and ASIC design.

Valid's MCM tool, for instance, works with the company's Signal Noise Analysis tool to compute and identify nets with excessive reflections, crosstalk, thermal drift and ohmic loss. According to Valid, these tools report overshoot and undershoot, backward and forward crosstalk, and provide a waveform display for either a rising or a falling edge. Etch delays extracted by the tools can then be fed back into Valid's Rapidsim simulator for re-simulation.

This need for accurate simulation data will undoubtedly narrow the gap that exists between the IC and circuit board designers. "Circuit board designers are getting into a level of sophistication that's been commonplace for IC designers for quite a while," says Richard Gordon, marketing director at Mentor Graphics' Silicon Design Division (Warren, NJ). "MCMs will get the traditional board designer working more closely with the IC designer. If you have a half-dozen chips on a single module with signals going to an outside bus with a comparatively simple structure, you'll have a big observability and testability problem unless you can tap into what's going on inside the silicon."

The nature of MCMs, then, will force design tool vendors into cross-pollinating the best of their ASIC and circuit board design tools. While much of the groundwork has been done in each of these disciplines, it's clear that a warmed-over version of existing tools simply won't do. "A lot of people are claiming that they have the answers to MCM design," cautions TI's Buschbom, "but what it tells me is that they don't understand the problem."

Whether tool vendors understand the problem or not, the real test of these emerging tools will be when designers have to rely on them to produce products that work and are cost-effective. ■

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 217**
Mentor Graphics
(201) 580-0102 ..... **Circle 218**
Racal-Redac
(201) 848-8000 ..... **Circle 219**
Texas Instruments
(800) 336-5236, ext. 700 ..... **Circle 220**
Valid Logic Systems
(508) 256-2300 ..... **Circle 221**

FDDI.
Build with the power of FXP.™

CMC, world leader in VMEbus Ethernet products, has forged a second generation architecture worthy of our stature.

FXP, Full Throughput Architecture. The culmination of over five years' experience and tens of thousands of installed nodes. The power behind the highest performance FDDI network processors available—the CMC-1000 series.

- A 25MHz Am29000™ RISC processor with its own separate program and data memory processes up to 30,000 packets per second.
- 512KByte Video RAM as global dual-ported data buffer memory to move packets at FDDI speeds and handle high VMEbus latency.
- Supernet™ FDDI chip set configured for dual- or single-attach operation coupled with high-speed 256KByte buffer memory provides full 100Mbps throughput.
- Autonomous DMA engine bursts data in block mode at up to 38MBytes/sec providing maximum efficiency to VMEbus host processors with minimal interrupt loading.
- Complete Station Management software fully complies with the latest ANSI X3T9.5 standards for maximum interoperability.
- CMC's link level driver for the CMC-1000 series supports UNIX® host-based protocols. CMC's Ethernet products also employ this same driver, allowing developers to use common applications for both Ethernet and FDDI networks.

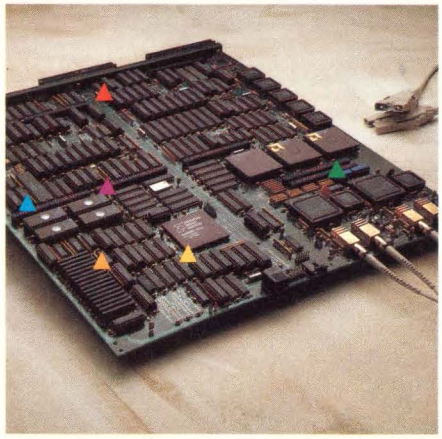

CMC, the first supplier of intelligent VMEbus Ethernet interfaces, remains committed to supplying fully compliant on-board TCP/IP and OSI protocol suites that off-load the host.

The CMC-1000 series is backed by a complete two year hardware warranty, as well as a comprehensive development environment, superior training and unmatched support.

To connect with the power of FXP and join the FDDI performance vanguard, call today.

1-800-CMC-8023

CIRCLE NO. 38

CMC

Designed well. Built well. Rockwell.
Rockwell CMC
125 Cremona Drive, Santa Barbara, CA 93117
PHONE: 805/968-4262 1-800-CMC-8023
FAX: 805/968-6478 TELEX: 240876

FXP

FULL THROUGHPUT FDDI

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

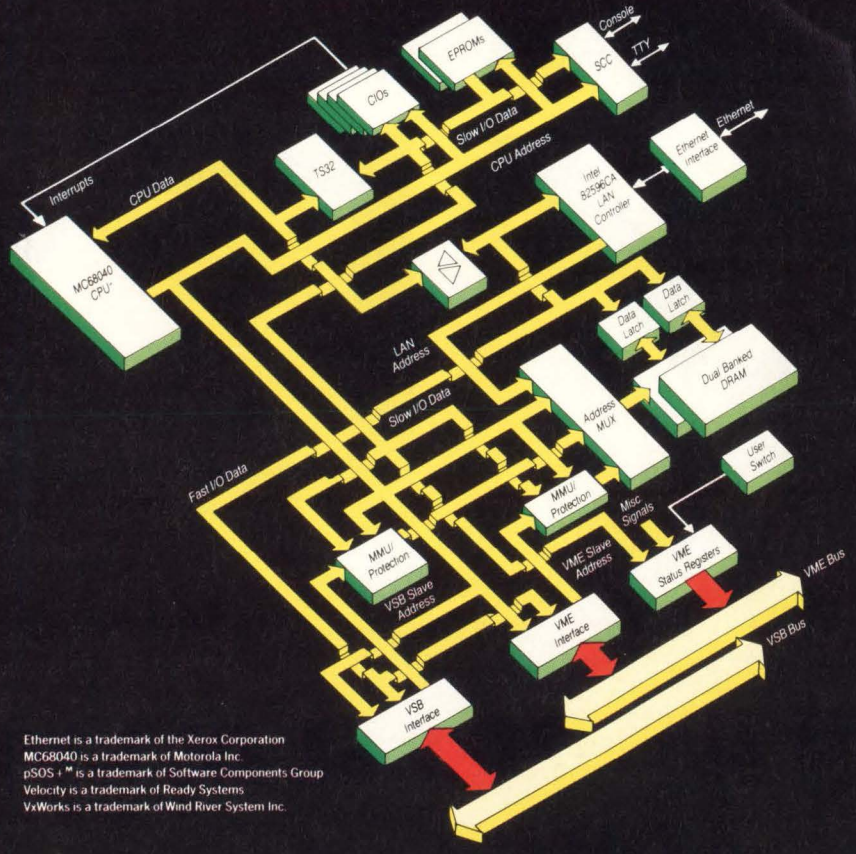

If you would like information on these, or any other Tadpole products, call us TOLL FREE on:

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+ᴹ • Velocity • VxWorks
  - TP-IX/68K V.3.2 implementation of UNIX

CIRCLE NO. 39

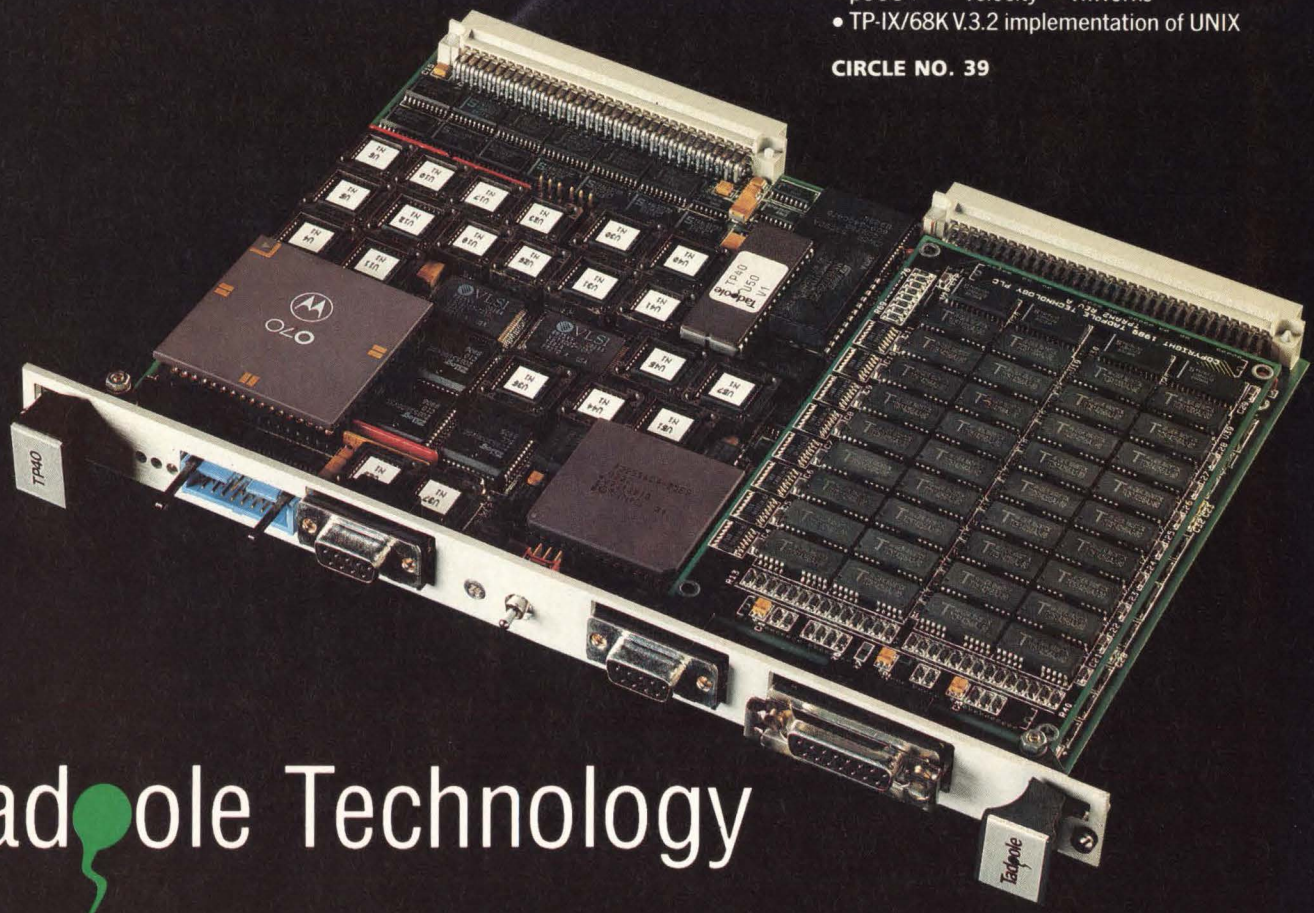

Tadpole Technology

SOFTWARE

Software architecture mates user applications with DSP functionality

Tom Williams, Senior Editor

The current generation of digital signal processing chips holds the promise of greatly increasing the functionality of all levels of computer systems. It will do this by means of software interfaces that will let application programmers easily integrate into their programs what were formerly very esoteric DSP operations. Traditional signal processing applications such as telecommunications, speech, image processing and instrumentation are poised to migrate to low-cost desktop systems. Other functions, such as multimedia information processing, are being developed by combining different traditional DSP tasks into integrated end-user applications.

New DSP chips such as Texas Instruments' TMS320C3x, Analog Devices' ADSP-210xx and Motorola's 96002 have reached a speed and complexity that make them the rivals of many general-purpose CISC and RISC processors. Not only do they perform at 12 to 17 Mips and 25 to 40 MFlops, they also sport many of the architectural features of general-purpose processors. These include multiple register files, large address spaces and integrated on-chip peripherals.

In order to work efficiently, these new chips must execute specialized DSP algorithms, manage memory, and perform system control and communications functions. Applications written directly to the hardware of subsystems based on these processors have in the past had to include such functions that are normally handled by the operating system on general-purpose processors. In short, DSP chips need a DSP-oriented operating system.

Spectron Microsystems (Santa Barbara, CA) has attacked the operating system part of the problem with its Spox signal-processing operating system, which provides for application portability among different DSP hardware. Currently, that would mean the TI320Cxx family and the Motorola and Analog Devices chips, for which support has been promised in the near future. But an additional bridge has been needed for user applications running on a host platform/operating system to easily access the signal-processing functions that run on the DSP hardware subsystem. To help forge this link, Spectron is providing an architectural concept called the Open Signal Processing Architecture (OSPA) to let applications running under operating systems that are resident on the host CPU, such as Microsoft Windows/DOS and X Windows/Unix, and that run separate from the DSP hardware, access DSP programs as if they were an operating system resource.

An operating system for DSP

The centerpiece of OSPA is the Spox operating system. Spox is a real-time multitasking operating system specifically designed to provide run-time support for compute- and I/O-intensive DSP functions. It consists

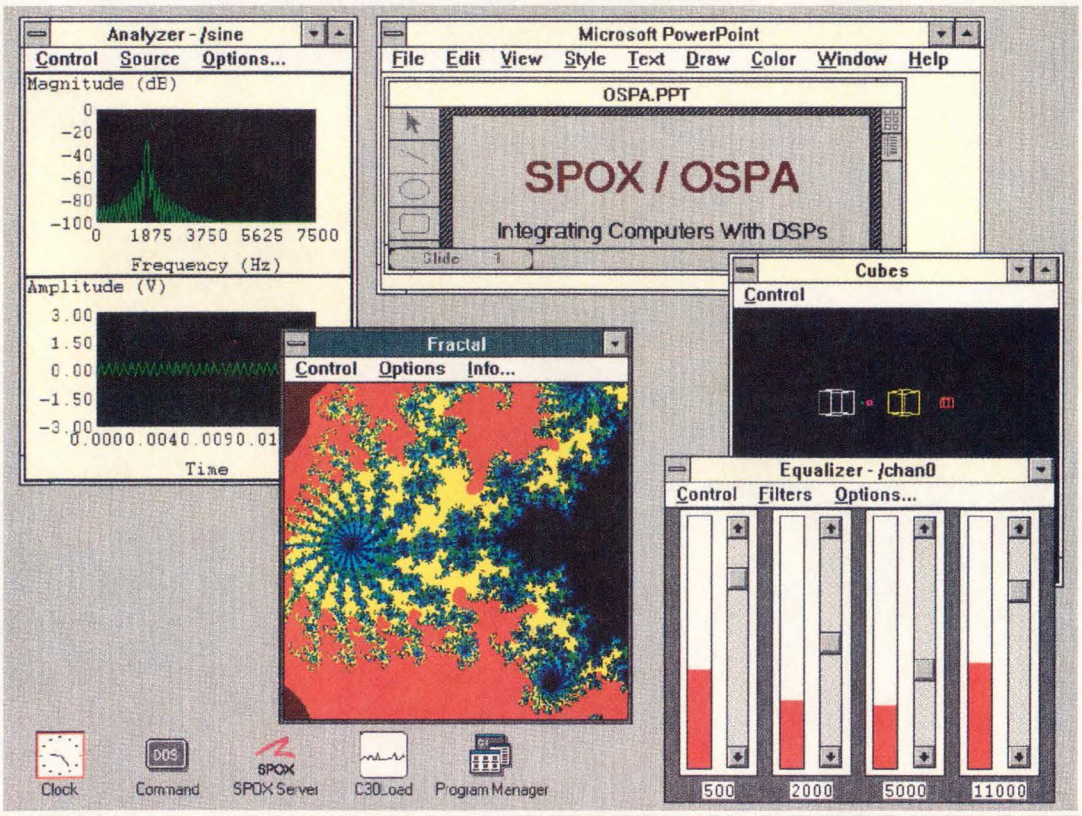

The Open Signal Processing Architecture (OSPA) is designed to let programmers build applications that use DSP technology by way of a familiar programming environment. Several applications appear within Microsoft Windows, including a spectrum analyzer, a fractal generator, graphics and an audio equalizer.

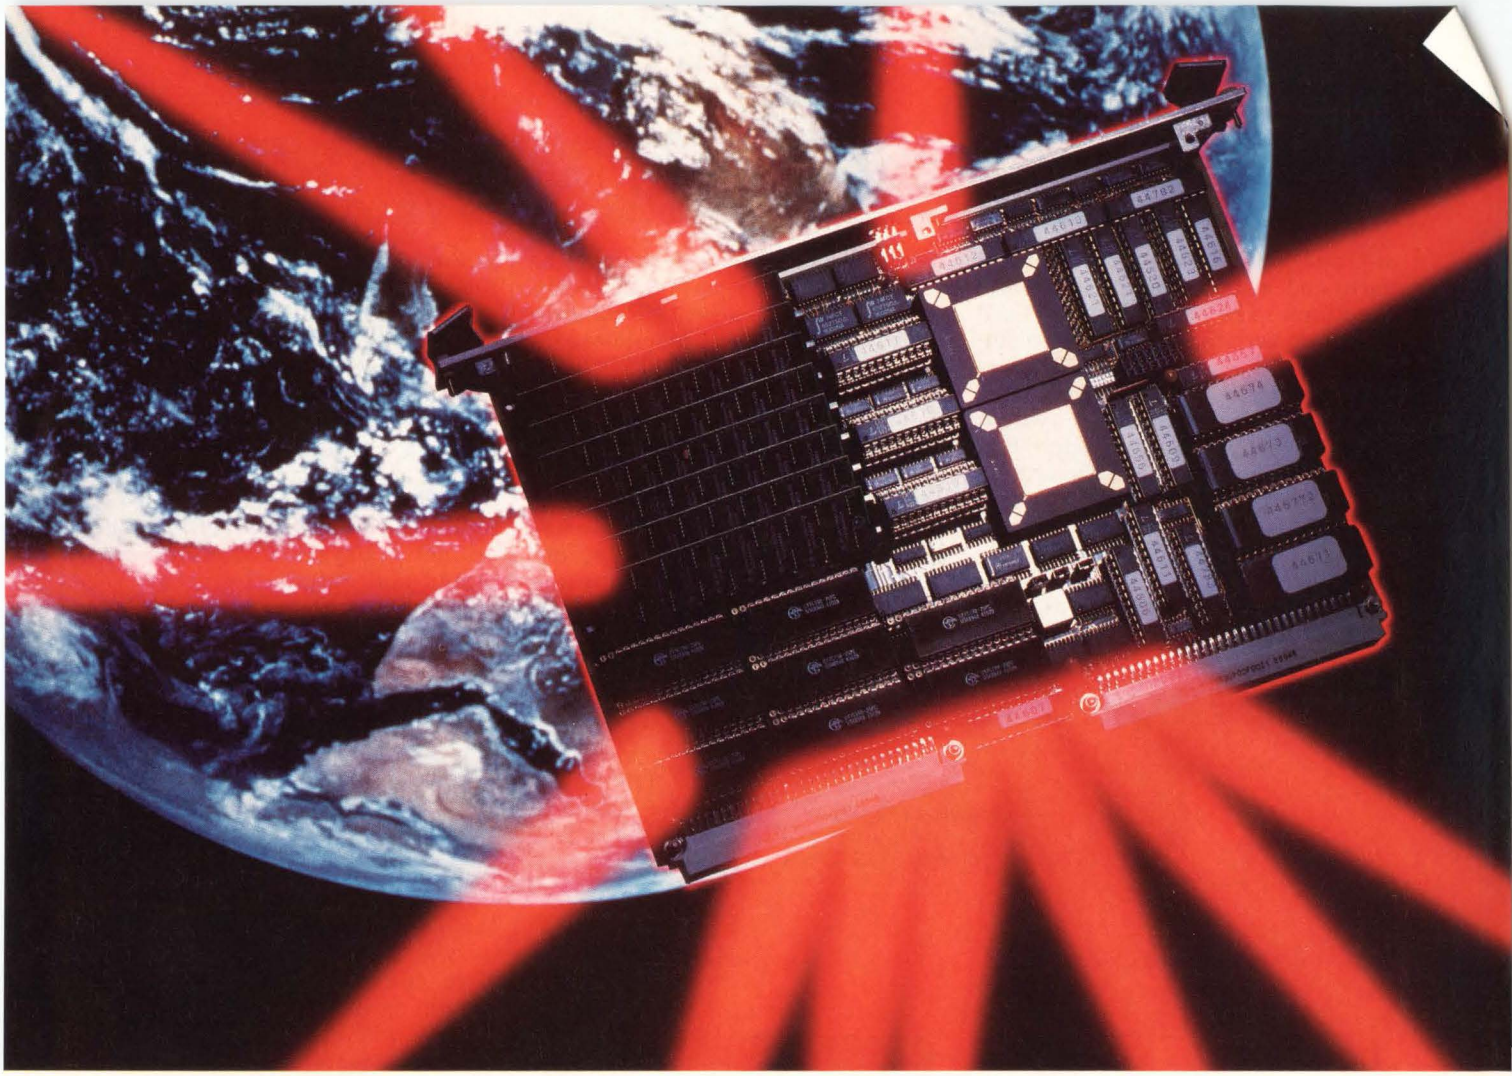

For Those Who Like It HOT

Announcing the MSP-6C30
66 MFLOP VME Array Processor

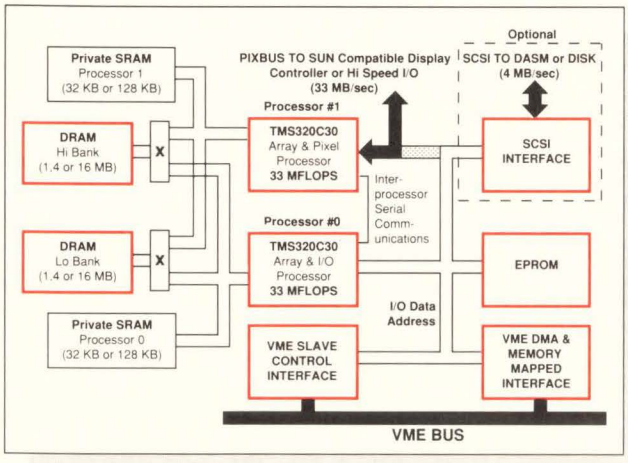

The MSP-6C30 is a 66 MFLOP, floating-point, VME array processor optimized for signal and image processing. High performance is assured through a unique, multi-banked, crossbar memory architecture, and high bandwidth is achieved with multiple, fast, industry standard I/O ports. Also available in a PC-AT 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. 40

SOFTWARE

of three main components that are packaged as libraries of C-callable functions. They are Spox-RTK, a real-time kernel; Spox-LINK, a communications component; and Spox-DSP, a set of general DSP functions.

These three operating system components are built around two programming interfaces: the system programming interface (SPI) and the application programming interface (API). At the SPI level, the system nucleus talks to hardware-specific device drivers via a device-independent I/O interface. The API provides application programmers with a standard set of high-level functions—both I/O functions and specialized DSP functions—that insulate them from the underlying hardware and let programs written to the Spox API be portable among Spox-based platforms.

The real-time kernel, Spox-RTK, includes many of the features of other real-time multitasking kernels, including timing services, synchronization of communications tasks and prioritized, event-driven scheduling. It's designed to service hardware interrupts as quickly as possible and thus keep overall latency to a minimum to meet DSP applications' need to process high-speed incoming data. On the TMS320C30, for example, interrupts are never disabled for more than 1.5 μs.

Spox-LINK lets the DSP subsystem communicate with the host operating system and interact with it through standard C I/O functions, letting DSP programs running under Spox on the DSP hardware subsystem get full access to the services of the host system. These include opening and closing named files, reading and writing byte streams, as well as formatting I/O operations. By using the device-independent I/O protocol, the I/O functions can talk to all kinds of device drivers. This is especially useful given the capabilities of today's DSP chips. "Some of our customers have written SCSI disk drivers for Spox because they needed a direct link from the DSP to disk," says Robert Frankel, Spectron vice-president of technology.

Spox-DSP addresses the traditional needs of advanced DSP programs, in terms of numerical processing algorithms and system-management functions such as managing multiple memory segments and device-independent I/O streams. "Memory allocation and device-I/O operations are equally important for meeting the time-critical requirements of advanced DSP systems," says Frankel. The fundamental data unit is the array, and the fundamental I/O entity is a stream, which inputs and outputs an array. Vectors, matrices and filters are used by the DSP math functions in Spox-DSP to manipulate data.

The library of vector and matrix math functions, along with management functions written in assembler, can be called by C programs to build custom DSP applications. Such applications would then be portable between DSP subsystems running Spox.

Another advantage is that DSP experts would be relieved of having to deal with the underlying hardware and assembly language programming of DSP operations. DSP experts will have access to ready-made tools to implement their ideas. The next step is to make the products of such expertise available to non-DSP experts who wish to take advantage of DSP functionality.

THE OPEN SIGNAL PROCESSING ARCHITECTURE

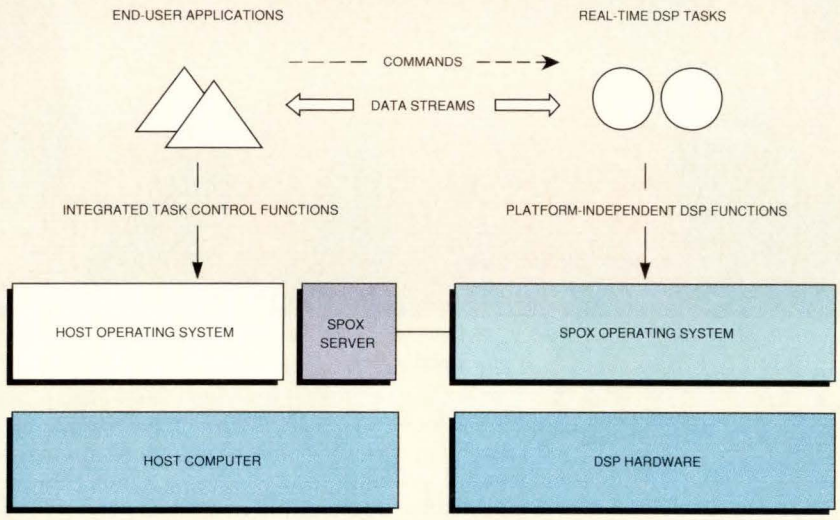

The Open Signal Processing Architecture (OSPA) includes the Spox operating system running on DSP hardware and an extension, the Spox-Server, resident with the host operating system on the host. End-user applications send commands to and send and receive data from the DSP subsystem via the Spox server, which appears as a host operating system resource. Thus, programmers wanting to incorporate DSP functionality in their applications can do so without having to know the details of DSP.

Bridge to broad-based uses

Engineers who want to process a certain incoming signal may know what type of filter they want to use and which parameters to apply, but not know how to implement such a filter on DSP hardware. The ability for a host-based program to send data to the DSP subsystem, request a service and then use the results later in the program is essential for the development of broad-based applications such as multimedia, audio, video and instrumentation. That type of application development depends on the availability of off-the-shelf DSP programs that are portable among hardware DSP subsystems (the problem addressed by Spox) and upon the ability of host programs to access those DSP functions via the host operating system, such as DOS or Unix.

To round out the OSPA—that is, to build the bridge between applications and DSP services—Spectron has developed the Spox-Server,

RESISTS EVERYTHING BUT A TOUCH

Introducing the DeeCO flat-panel SealTouch® Terminal. Now there's a VT220 compatible, NEMA 4- & 12-rated, touch-screen terminal that's also compatible with the realities of an industrial environment.

This incredibly tough little terminal encases an electroluminescent, flat-panel display in rugged, lightweight, completely sealed cast aluminum. No fans, no filters. Nothing gets in or out but information.

You get line graphics, downloadable fonts and easy pop-up touch menus. On-screen touch keyboards, and a keyboard port give you even more flexibility.

And because it's so small (10½" x 11½" x 3") and light, it mounts anywhere you need it – on a swing-arm, heavy machinery, pedestal, wall, desk, bench or cleanroom station.

If your environment eats up ordinary terminals, or if you'd like to put terminals where ordinary terminals won't go, just call (415) 471-4700.

Lucas Deeco

Formerly DEECO Digital Electronics Corporation

Lucas Deeco, A Division of Lucas Duralith Corporation, 31047 Genstar Road, Hayward, CA 94544-7831
(415) 471-4700

CIRCLE NO. 41

SOFTWARE

which is a host-resident extension to Spox. The server lets end-user applications communicate with real-time DSP tasks. It provides application developers with a high-level programming interface that insulates the application developer from many of the low-level details of the DSP subsystem. Spox-server differs from Spox-LINK in that the latter is a means for Spox to call host services on the DSP hardware. The server lets the host operating system access Spox services by sending commands and by sending and receiving data, making the host software independent of the DSP implementation. It merely calls a DSP service by name.

The first version of the Spox-Server runs with MS-Windows 3.0 and DOS. The next version will run with X Windows and Unix. Eventually, Spectron will supply a porting kit to let developers of embedded real-time systems communicate with Spox-based software via their real-time executives. "Controlling and communicating with real-time DSP tasks via the Spox-server should be conceptually the same as opening a disk file or reading from a terminal," says Frankel.

OSPA aims at building a bridge between "islands of expertise," as Frankel puts it. "Spox users are DSP experts and system programmers. Spox-Server users are knowledgeable about Windows, DOS, Unix, the host computer and end-user applications," he says. Spox-Server does impose a certain discipline on how applications are written, just as Windows calls for writing to a well-behaved convention.

The ability of applications to pass data via a mechanism such as MS-Windows' Dynamic Data Exchange (DDE) along with Spox real-time multitasking allows an integrated application using several DSP tasks to let a single DSP board take the place of several pieces of fixed logic. A single DSP processor board, for instance, with three serial ports and a TMS-320C31 could function as a fax, a modem and a telephone answering machine, all of which could operate simultaneously, taking up only about 25 to 30 percent of the DSP's bandwidth, according to Frankel. "And if it can even do those three things, it pays for itself," he adds.

For some, the real-time aspects of an OSPA system will be more important than the DSP part, says Spectron president David Wong. "DOS and Unix can't do real-time. Having a very intelligent DSP doing real-time takes the real-time responsibility away from the host operating system." In other words, one could build

Sierra C For The 68000 Family Creates The Four Second Coffee Break

Those endless compile-time coffee breaks are a thing of the past.

Compiles Faster

The Sierra C™ compiler gets your code into your target system 10 to 20 times faster than other high performance compiler packages. The dhrystone benchmark, on a 16 MHz PC AT® host, compiles, assembles, links and downloads to a 68020 target in less than 4 seconds.

Generates Better Code

Sierra C has resolved the paradox of fast compiles and optimized code. While compiling at blazing speeds, Sierra C generates faster and tighter code than any other optimizing compiler on the market today.

Complete Development Toolset

Sierra C is more than fast—it's complete. It includes all the tools you need for cross-development, built to perform with unsurpassed efficiency. Along with Sierra C's optimizing compiler, macro assembler, and linker/locater, you get a librarian, symbol table examination utilities, a high speed parallel download system and a multi-window source level debugger. Complementing the toolset is an extensive C runtime library supplied in both source and object form.

Designed for Embedded Systems

Sierra C conforms to the proposed ANSI C standard, supporting function prototypes and the important keywords *const* and *volatile*. Position independent, re-entrant, and ROMable code can be generated. Code can be optionally grouped into as many as 126 independently positionable sections. Assembly language can be inserted into your C program through the *asm* keyword extension. And the download utilities allow executable modules to be transferred directly into the target system or into a PROM programmer with optional even/odd/quad byte selection. The Sierra C toolset supports all members of the 68000 family including the new 68332. Sierra C is available now for PC AT and PS/2™ computers, DEC VAX® systems, Sun™ Workstations and other UNIX™ systems.

With Sierra C you get better products out faster. Call today for more information at **800•776•4888**

CIRCLE NO. 43

SierraSystems

6728 Evergreen Ave., Oakland, CA 94611 (415) 339-8200

SOFTWARE

a real-time control system with floating-point capability using a graphics-based user interface and standard host operating system linked directly to the real-time system.

The combination of DSP functionality with an almost free real-time capability has made Spox so attractive to DSP chip manufacturers such as Texas Instruments (Dallas, TX) that TI is bundling Spox "with every hardware/software DSP development system TI ships," according to Rick Rinehart, TI's floating-point DSP program manager. With the advent of OSPA, he says, "essentially you've got a platform-independent software system that really lets you begin using and proliferating this DSP technology."

Motorola (Austin, TX) has been jointly working with Spectron and DSP development tool maker Ariel (Highland Park, NJ) to implement a version of Spox for Motorola's new 96002. Spox itself will be available as a separate product for the Ariel MM96 board, which incorporates two 96002 chips. It will be able to work in conjunction with Motorola's development system.

Since Spox will be the operating system running on the final product, says Mike Collins, marketing manager for Motorola's DSP group, "You wouldn't want to run Spox on a development system, you would want to run it on a target system to develop applications." Thus the Motorola/Ariel arrangement won't bundle Spox as tightly as the TI package. But, according to Ariel director of marketing Les Listwa, "A C compiler will be bundled with Spox as soon as Spox is available for the 96002." The Spox-Server initially will be available from Spectron, with third parties still to decide how they'll offer or bundle it.

The Spox-Server will be an item primarily for the application developer to incorporate with the host computer's operating system. But the distinction between a DSP board-level subsystem and a host computer isn't expected to last long. Motorola's Collins says he has called on all leading PC and workstation vendors. "All are looking at or have active projects for adding DSP or multimedia capabilities (many on the motherboard)," he says. "And that is going to be the next generation of machines coming out." ■

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

Ariel
(201) 249-2900 ..... **Circle 213**
Motorola
(512) 891-2000 ..... **Circle 214**
Spectron Microsystems
(805) 967-0503 ..... **Circle 215**
Texas Instruments
(800) 336-5236, ext. 700 ..... **Circle 216**

SOFTWARE

Fast SCSI outpaces wide SCSI

Tom Williams, Senior Editor

SCSI-2, the long-awaited enhancements to the SCSI standard, will actually begin to be incorporated in real products as subsets of the specification are finalized, and as products that use the subsets are developed and brought into the market. SCSI-2 offers three basic improvements over SCSI-1: twice the speed, a wide data path supporting 16- and 32-bit transfers, and a standardized common command set that's applicable across a wide range of peripherals, including tape and optical drives. In addition, the SCSI-2 common command set supports queuing of commands and multithreaded operations—tasks that SCSI-1 could only do in limited form and in nonstandard ways.

The emergence of a standard really only takes on meaning, though, when products become available and designers have the means to integrate them into systems. "SCSI-2 is a conglomerate of features that you could implement," says Tom Martin, product marketing manager for the development systems operation of Adaptec (Milpitas, CA). Some of are relatively easy to implement; others are not so easy.

"Things like queuing and threading are logical extensions of SCSI-1. It's just a matter of when developers will decide to put these functions in their firmware and put enough memory on their peripherals. Caching falls in the same category. There's really nothing standing in the way," Martin says.

Wide SCSI slow to take off

However, Martin points out, "For wide and fast SCSI, you're talking real protocol changes, and you have to wait until chips are available to support such features. There are chips out for fast SCSI." Fast SCSI today means a 10-Mbyte/s synchronous transfer rate, but still over an 8-bit-wide data path—not a full-blown SCSI-2 implementation. Martin expects to see a slew of fast SCSI peripherals introduced later this month at the Comdex show in Las Vegas.

The mere existence of chips and peripherals, though, isn't enough for system designers; they also need the means to design with them, test and verify their designs and incorporate them into systems that can realize the performance potential that such chips and peripherals offer. To meet the needs of hardware and software developers, Adaptec has just announced its SDS-3F development system for fast SCSI.

SDS-3F is a superset of Adaptec's SDS-3 development environment for SCSI-1. Users of SDS-3 can upgrade to 3F for a fraction of the cost of a new development system, Martin says. In addition, all tests and software developed on SDS-3 will run on SDS-3F without recompiling or relinking, he adds.

Wide SCSI, the other major aspect of SCSI-2, is still being shaken out, according to Martin. "One of the problems with wide SCSI-2 is that it addressed the width issue by creating the A, B cable scheme," he says. One cable had 50 pins and the other had 68 pins to give the ability to transfer 8, 16, 24 or 32 bits at a time. The only problem was that people building drives with embedded SCSI controllers couldn't live with the amount of space required to contain the two cables. To accommodate their concerns, a single cable scheme has been devised that allows 16-bit transfers.

Compatible cabling

The single cable has a 68-pin connector but is still compatible with the 50-pin SCSI-1 connector. It's designed so that the middle 50 pins of the 68-pin connector are compatible with SCSI-1, and a SCSI-1 connector will mate with it. The extra 8 bits and a parity bit that supports 16-bit-wide SCSI-2 are placed on either side of SCSI-1's 50 pins so that SCSI-1 and SCSI-2 devices can be mixed and matched on the same cable. Address arbitration is compatible as well. The upper 8 bits of an address are of lower priority than the lower 8 bits, so a SCSI-1 device never has to be concerned with the upper 8 bits.

This situation makes it appear that 32-bit-wide SCSI-2 won't be widespread, even though it's listed in the official specification. "That's why fast SCSI came along first," says Martin. "It took a long time to figure out what scheme people would actually use. What we're seeing is that nobody is interested in the A, B split-cable scheme. But fast is very much alive."

In software, going to fast SCSI is

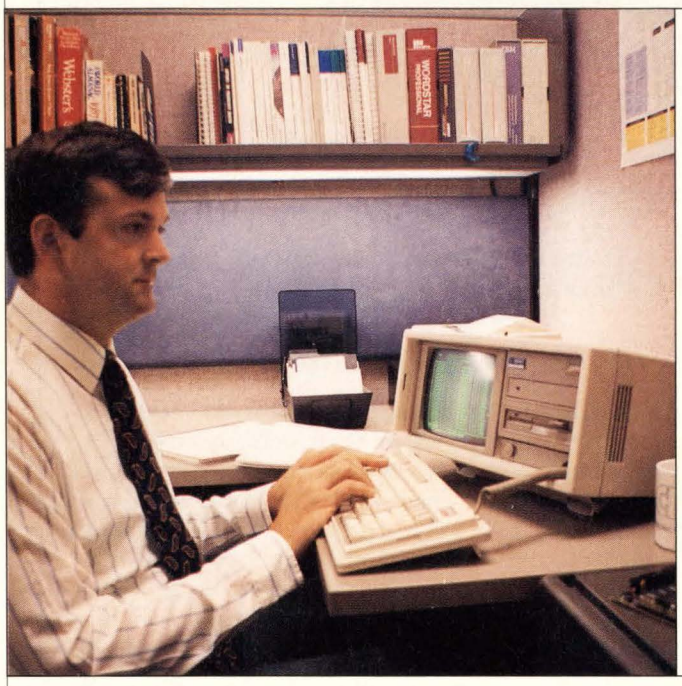

Tom Martin, product marketing manager at Adaptec, says that fast SCSI appeared first among the features of SCSI-2 because it can be implemented with existing cabling, unlike wide SCSI. "It took a long time to figure out what cabling scheme people were going to use," he says. While that issue is being sorted out, "Fast SCSI is very much alive."

SOFTWARE

a matter of changing one parameter—the synchronous period—from 200 to 100 ns. But that puts heavy demands on the development and test system. The higher speeds require test equipment that can not only resolve signals but can also verify transmission integrity.

"At 10 Mbytes/s, people are concerned about signal integrity and termination problems with single-ended transmission," Martin says. The suspicion is that those high speeds may require differential transmission, which, in turn, would require more chips and consume more power. Single-ended transmission is good for low-cost systems of up to 20 ft., while differential allows up to 50 m of cable, and the two cannot be intermixed. In order to allow developers to design and verify high-speed systems, the SDS-3F is capable of both single-ended and differential transmission up to the full 10 Mbytes/s.

SCSI's growth rate may increase with the advent of a standard common command set that addresses the needs of peripherals beyond just disk drives and speed increases. In terms of sheer volume, low-cost disk drives using the AT interface still far outnumber SCSI. But SCSI, being intelligent, has a need for more-complex test and development equipment. Despite that, Martin sees this intelligence—making SCSI a system bus for peripherals—the key to eventually making SCSI more widespread than the AT interface. The attraction of being able to hang printers, scanners and other storage devices on the same intelligent bus with disk drives is the feature that SCSI advocates think will push SCSI acceptance beyond the AT interface. That has, of course, been their hope for a good number of years. The question is, what level of functionality and cost will eventually make it come true? ■

"Queuing, threading and caching are logical extensions of SCSI-1. It's just a matter of when developers will put these functions in their firmware."

—Tom Martin, Adaptec

■ ■ ■

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

Adaptec
(408) 945-8600 ..... **Circle 222**

Optimize your 680x0 Application to its peak

...with Oasys Embedded/Cross Software Tools.

Reach the pinnacle of your development with Oasys Cross Tools

Oasys delivers full 680x0 Cross and Native development tool kits built around the industry standard Green Hills compilers.

The Green Hills 680x0 compiler's Dhrystone rating is 35% faster than Sun's native C compiler. Using the new procedural inlining optimization, the rating jumps to a staggering 80% improvement.

Advanced C, C++, FORTRAN and Pascal Language Support

- Full ANSI C XJ311 support with switch selectable K&R C support.
- A true *C++ Compiler* for cross development.
- FORTRAN 77 with VAX VMS and DoD MILSTD-1753 extensions.
- Pascal ISO Level 1 support.
- All compilers are inter-language callable.
- Native compilers are supported for 680x0, 88000, 80386 & i860 hosts.

Available Now!

Oasys tools are available on over 35 UNIX based workstations and systems including DEC (VMS/Ultrix), Sun, HP/Apollo, IBM (OS/2 and AIX), and DG. Call us for the complete list.

Other Cross Tool Kits include:

88000, 80386, i860 and Microsoft C.

Integrated 680x0 Cross Development Tool Kits

Green Hills 68000/10/20/30/40 Optimizing Cross Compilers (C, C++, Pascal, FORTRAN)

- Superior register allocation, inlining and loop unrolling techniques.
- Support for position independent code and data.
- 68881, 68882, 68332 support.
- ANSI C run-time source library included for cross development.

Oasys UDB Universal Debugger

- Source level remote debugger for embedded programming.
- Multi-targeted and multi-window (supports 680x0 and 88000 targets).
- Extensive macro language.
- Customize communications to your target.

Oasys 680x0 Assembler/Linker

- Macro Assembler, Linker, Librarian, Cross Reference and Symbol File Format Utilities.

Oasys 680x0 Simulator

- Invaluable for debugging software without target hardware.

Scale new heights with high performance tools from Oasys. Call us for detailed benchmark results.

CALL (617) 890-7889
FAX (617) 890-4644

Oasys

Worldwide Support: Australia—Forefront Computing Svcs 03-785-1122 **Benelux**—Computing & Systems Consultant 040-434-957 **France**—RTS 01-6986-1958 **Germany**—Stollman 040-890-880 **Israel**—Ankor Computers Ltd. 972-3-5483737 **Italy**—Instrumatic 02-353-8041 **Japan**—ERIC 03-668-8866, MCM 03-487-8477 **Norway**—OTE Komponent 02-30-6600 **Singapore**—Automated Systems Ltd 02-789-9566 **Spain**—Decisa 01-204-4500 **Sweden**—Traco AB 0893-0000 **Switzerland**—Zuhlke Engineering AG 01-730-7055 **UK**—Real Time Products 021-236-8070

230 Second Avenue, Waltham, MA 02154

Single Source, Total Solutions

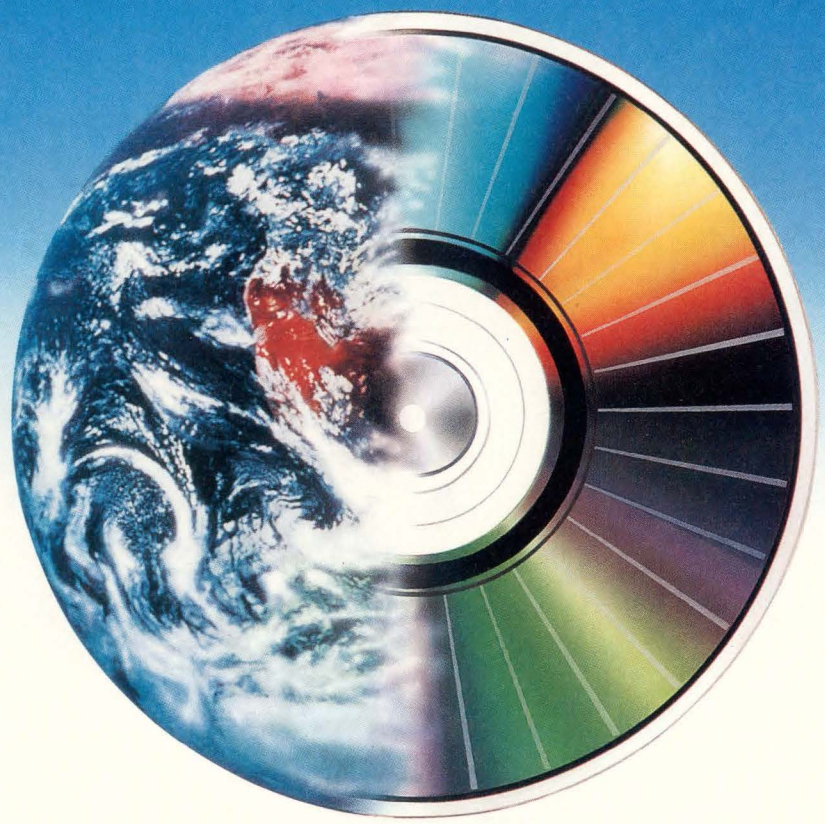

A World of Choice

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.

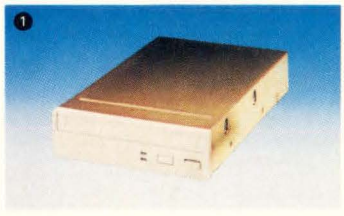

❶ RH5500/5-1/4″ REMOVABLE CARTRIDGE HARD DISK DRIVE
❷ RS-9100H/5-1/4″ HALF-HEIGHT WORM SUBSYSTEM
❸ RO-5030EII/REWRITABLE OPTICAL DISK DRIVE
❹ RJ-5330E/MAGNETO-OPTICAL DISK LIBRARY UNIT

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. 47

U Interface silicon crystallizes ISDN standards

ISDN still has hurdles to overcome before it's widely accepted. But the recent introduction of standard U Interface silicon and the settling of key standards issues are generating fresh interest in ISDN.

Ken Marrin
Contributing Editor

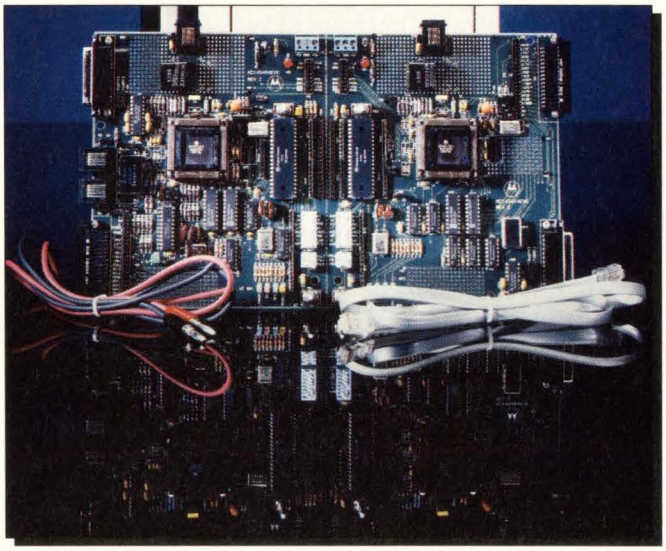

Providing two U Interfaces and an S Interface, Motorola's MC145494EVK development board can simulate the complete link between the customer premises and central office. Support for layers one through three of the OSI protocol is provided through a separate 68302-based protocol processing board, which may be used in conjunction with the development board.

Although the current realization of ISDN hasn't lived up to original market projections, over 2 million basic-rate ISDN lines are expected to be installed in North America by 1994. Fueling that estimate is the fact that most of the pieces of the ISDN technology puzzle have fallen into place in the past year.

One of the most important pieces of the ISDN (Integrated Services Digital Network) puzzle was the introduction of 2B1Q U Interface hardware, which provides the link between the central office and the customers' network termination (NT-1). To date, more than a half dozen vendors have announced U Interface devices, together with hardware evaluation boards and even embedded software support for the first three layers of the Open Systems Interconnection protocol.

Another key piece to fall into place has been the solidification of the 2B1Q U Interface standard. With the recent addition of three critical maintenance functions in September 1990, manufacturers and customers of NT-1 and terminal equipment should be less reluctant to make a serious commitment to ISDN.

To make ISDN's services available on a national level, manufacturers of central office switches are beginning to adopt the common-channel Signaling System 7 (SS7) protocol, a standard protocol that lets dissimilar ISDN central office switches communicate with each other. AT&T (Allentown, PA) and Northern Telecom (Nashville, TN) are already complying with SS7. Most other manufacturers are expected to comply by 1991.

To promote interoperability at all levels in the ISDN network, a number of organizations are beginning to provide conformance testing for a variety of ISDN equipment, including NT-1s, terminal equipment (such as phones and faxes) and even central office switches. Among these groups are Bellcore, ANSI, the National Institute for Standards and the Corporation for Open Systems. Many, including Siemens (Santa Clara, CA) and AT&T, would like to see Bellcore play a dominant role in administering tests and providing certification. Others, including Motorola (Austin, TX), would like to see this task distributed among a variety of organizations.

One of the final pieces of the ISDN puzzle to fall into place will be the establishment of tariffs to set standardized ISDN pricing and let the service be provided over lines regulated by regional Public Utility Companies. To

date, only Illinois Bell and Pacific Bell have filed tariffs, according to market research firm Dataquest (San Jose, CA). Bell Atlantic, Bell South and Nynex are expected to file by early 1991.

In search of applications

ISDN so far has been deployed primarily on a test basis as a voice/data LAN in campuses and large corporate environments. But it's unclear whether, in the long run, ISDN will be able to penetrate these applications on a large scale. Often, these environments are already serviced by higher-bandwidth LANs, together with gateways to digital private branch exchanges (PBXs) that not only handle voice and data but provide many of the same services offered by ISDN. Competition should stiffen with the deployment of FDDI II (Fiber Distributed Data Interface), a fiberoptic LAN that can handle voice and data at 100 Mbits/s.

While it's unlikely that ISDN will displace modern digital PBXs and high-speed LANs in large campus or business environments, it may fare well in applications where customers are upgrading from older analog PBXs. Even if the customer opts for a PBX and LAN over ISDN, ISDN may have a place as part of a hybrid network. LANs and PBXs, for example, might handle local data and voice traffic, while global networking might be achieved by providing a gateway between the office's LANs or PBX and either a basic-rate (U Interface) or primary-rate ISDN line.

"For the near term, one of the most lucrative markets for ISDN will be in replacing existing V.32 modems."

—Al Mouton, Motorola

For the near term, one of the most lucrative markets for ISDN will be in replacing existing V.32 modems, predicts Al Mouton, product marketing manager at Motorola. There are now over 60 million digital lines in North America, says Mouton, 8 million of which are dial-up and leased Centrex lines using V.32 modems. By upgrading to ISDN, customers not only eliminate the need for a dedicated modem line, but they increase their data rate from 9,600 bits/s to 64 kbits/s.

Ironically, until customer demand for ISDN services increases, says Ramesh Sirsi, product marketing manager at Siemens, the principle demand for ISDN hardware will be in non-ISDN applications such as digital pair gain, Digital Loop Carriers and rate adaption for T1 multiplexers. In such applications, local operating companies will use ISDN's U Interface to increase the capacity of existing analog lines, while at the same time partially defraying the cost of upgrading to ISDN.

In the long term, the majority of

ISDN overview

ISDN (Integrated Services Digital Network) is an integrated network that supports voice, data and video transmission over existing copper telephone lines. The CCITT recommendation defines two interface rates for ISDN: primary rate and basic rate. Both interfaces support simultaneous full-duplex voice and data using circuit-switched (voice) and packet-switched (data) connections on the same channel.

The primary rate interface (PRI) is typically used to interconnect high-bandwidth devices such as mainframe computers, private branch exchanges (PBXs) and groups of lower-bandwidth basic rate lines with the central office digital exchange. The North American PRI, as defined by ANSI I.431, is based on the DS1 (T1 trunk) transmission rate of 1.544 Mbit/s. It consists of twenty-three 64-kbit/s B channels for voice and data, and one 64-kbit/s D channel for signaling. Because a single D channel is used to handle all signaling, the other 23 channels are available for user data and voice transmission.

The European PRI, as defined by the Conference European Postal And Telecommunication (CEPT), specifies 32 total channels (30 B, one D and one control) with an aggregate bandwidth of 2.048 Mbits/s. The international CCITT specification incorporates both the ANSI 1.544-Mbit/s and CEPT 2.048-Mbit/s PRI rates.

The basic rate interface

Providing a composite bandwidth of 144 kbits/s, the basic rate interface (BRI) is typically used to carry data to and from small end-user systems—such as voice/data workstations—and terminal adaptors for non-ISDN devices. It consists of two 64-kbit/s information channels (B channels), used for voice and data, and one 16-kbit/s packet-switched data channel (D channel), used to carry either data or signaling information (such as call set-up and take-down). The B and D channels are full-duplex bit streams and are time-division multiplexed (two B and one D) into a common stream that contains both user and signaling information.

The BRI consists of both a four-wire S/T interface and a two-wire U interface. The S/T interface (S), an approved CCITT I.430 recommendation and ANSI T1.605/1989 standard, is wired inside the customer's home or office. Using the standard wall plug, this interface links customer equipment such as telephones, fax machines and computer terminals with the ISDN network. Up to eight devices may be connected to this four-wire interface, which forms a passive bus.

The two-wire U-interface connects the local telephone lines to the customer's home or office. U-Interface silicon is used in several places in the ISDN network: in the network termination that links the customer premises with the local telephone line; in the line termination that links the telephone line with either a primary rate multiplexer or the Central Office Digital Exchange; and in two-wire terminal equipment such as a two-wire ISDN telephone, in which the S Interface is bypassed.

New World Record VMEbus Data Dash!

VMIC's new NETWORKZ™ is a high speed, high performance, 16 chassis multidrop VME- to-VME parallel network that features data transfers @ 20 Mbytes/s*, reflective memory, 3 interrupts, and a price that won't overload your budget!

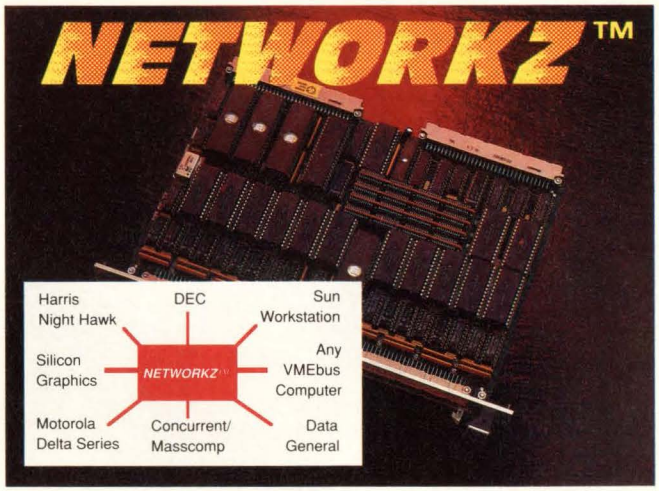

Now Available for Immediate Shipment and Backed by a 100% Satisfaction Guarantee. Call Our TOLL-FREE Hotline Today!

If your existing VMEbus network is too slow to keep up with your growing demands NETWORKZ™ is the fastest and most simple solution! In fact, newNETWORKZ™ is up to 50 times faster than comparable computer products. You don't have to be able to read minds to know the best choice for your VME-to-VME network needs. VMIC's new NETWORKZ™ will enable you to cross the finish line in world class, record breaking real time!

• 2 Year Warranty • 24 Hour Customer Service Hot Line • 20 Mbytes/s Data Transfers (*5 foot cable) • Multidrop Capability (up to 16 chassis) • No Processor Overhead or Involvement in Network Operation • Differential Line Drivers & Receivers provide ±7 V of Noise Immunity • A24:A32:D32:D16:D8 Memory Access • A16:A24:A32:D8 Status & Control Access • Single 6 U Board

New NETWORKZ™ transfers data by writing to on-board global RAM. Data written into the 1 Mbyte of reflective memory is broadcast to all nodes on the network without further involvement of the sending or receiving nodes. Data transfers from memory locations on sending nodes to corresponding memory locations on receiving nodes.

In addition to transferring data between nodes, NETWORKZ™ will allow any processor in any chassis to generate a VMEbus interrupt on any other chassis. Three interrupts are available. The user may define function, priority, and vector for each interrupt. Any processor can generate an interrupt on any other VMEbus on the network. In addition, any processor on the network can generate an interrupt on all VMEbuses on the network simultaneously. It's time you network in world class, record breaking real time - call VMIC toll-free and see how fast we are!

1-800-322-3616

VMIC®

VME MICROSYSTEMS INTERNATIONAL CORPORATION

12090 South Memorial Parkway
Huntsville, Alabama 35803-3308
(205)880-0444 FAX(205)882-0859

VMIC products are internationally represented by Distributors throughout the world. Call or FAX VMIC for complete information.

DEC, Delta Series and Night Hawk are registered trademarks of their respective companies.

CIRCLE NO. 48

U INTERFACE

ISDN applications will be in residences and small offices with only a few lines and no preexisting LAN, says Brian Jones, product marketing manager at Zilog (Campbell, CA). In these applications, ISDN will provide the fastest and least expensive means of networking devices, consolidating lines and increasing long-distance bandwidth.

Hurdles to ISDN deployment

While most of the technological pieces of the ISDN puzzle are intact, a number of hurdles will have to be overcome if ISDN is to experience the growth expected by many analysts. An important one is cost. While a network interface (such as Ethernet) for a personal computer costs only about $250, for example, an ISDN four-wire interface currently costs $1,500.

Upgrading to ISDN is also expensive for local telephone companies and switch manufacturers, as each upgrade of an analog line to an ISDN line (or the addition of a new line) requires a new line card. The chip costs for an ISDN line card are now $70, compared with $10 for an analog line card, says Pat Sullivan, communications marketing manager at SGS-Thompson (Phoenix, AZ). Unfortunately, without high-volume production, ISDN hardware costs will be slow to drop, points out Steve Bretoi, ISDN product marketing manager at Advanced Micro Devices (Sunnyvale, CA).

Another factor slowing ISDN's progress is a lack of consumer demand. While a growing number of customers are demanding ISDN equipment compatibility for fear of obsolescence, few are actually demanding ISDN services. Many vendor spokespersons, including Siemens' Sirsi, blame the lack of demand on a failure of local operating companies to properly promote the combined voice/data and centralized management capabilities of ISDN.

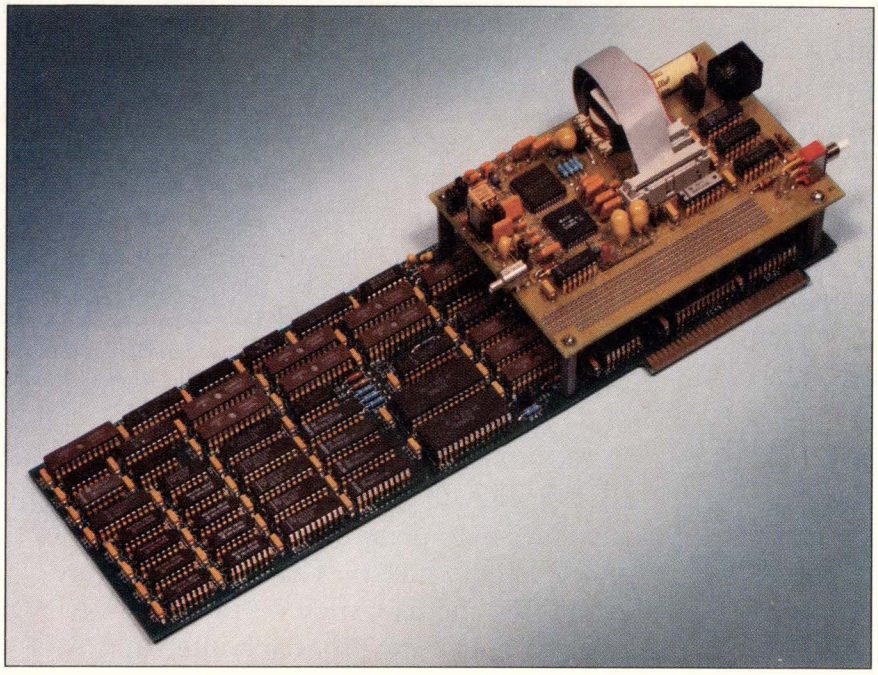

Based on AT&T's U Interface, the 2B1QCAK (Complete Adapter Kit) is a PC plug-in evaluation board that provides both a network termination and a line termination. By plugging a handset into both the NT and LT, designers can simulate a complete U Interface.

Mike Pauzer, product marketing manager at AT&T, agrees. Many of the ISDN services that operating companies are promoting—including conferencing, hold, transfer and drop—are already available under Centrex, he says.

ISDN standard a moving target

According to Siemens' Sirsi, "Other ISDN capabilities that are being promoted, such as ISDN's superior immunity to noise and crosstalk, are too technical and not easily recognized by the average consumer."

Another reason for the slow penetration of ISDN has been a lack of complete standardization for the U Interface. Fearful of quick obsolescence, manufacturers of chip, terminal and NT-1 equipment have been reluctant to make a serious commitment to the market.

Recently, substantial progress toward completing the standard was made when ANSI agreed on three critical maintenance functions. But as AT&T's Pauzer points out, the fact that the standard is still changing will continue to make both manufacturers and customers wary.

Some argue that a lack of complete standardization isn't critical, as future changes can be temporarily handled in software. But software is used only to interpret a new function, notes Pauzer. Typically, some type of hardware is needed to take action—for example, to generate a tone.

One way that Motorola is addressing the problem of tracking a moving standard is by designing its U Interface with a microcontroller interface. By using an inexpensive microcontroller such as the MC6805, says Motorola's Mouton, the company can respond to standards changes in software without having to redesign its U Interface devices.

U Interface silicon arrives

While flat customer demand has dampened ISDN enthusiasm, the recent introduction of standard U Interface silicon is generating fresh interest among equipment manufacturers, many of whom have been using proprietary interfaces and line codes until now. To date, more than six vendors have announced the availability of U Interface devices.

AT&T's two-chip device is currently in production, with a single-chip version that incorporates the latest ANSI revision, expected to debut in 1991. Siemens and AMD, who will second source Siemens' device, are in production with a single-chip device.

Other vendors who claim to be in production with single-chip devices are Motorola, which has developed its device in conjunction with Northern Telecom; and National Semiconductor (Santa Clara, CA), whose device is being second sourced by SGS-Thompson.

Mitel (Kanata, Ontario), working in conjunction with British Telecom, has announced samples of its single-chip device. NEC, Hitachi, Fujitsu and Oki Semiconductor are produc-

WE'D LIKE TO MAKE AN EXTREME CASE FOR BUYING OUR CPU CARDS. YOU COULD SAVE THOUSANDS

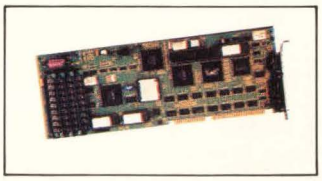

N286 high-speed 80286 CPU Board
The fastest AT equivalent processor of its kind
- Runs MS/DOS® PC/DOS® OS/2® XENIX,™ PC-MOS™ and other O/S and software.
- Choice of 10, 12.5, 16 or 20 MHz speeds.
- Surface mount technology and minimal parts count for high rel, high MTBF.
- Six layer board with discrete ground and VCC planes enhances signal integrity.
- Two serial and parallel ports, with built in LIM and EMS.
- Up to 4MByte RAM installable on-board.
- VLSI Chips and Technologies NEAT™ chip set for improved integration.

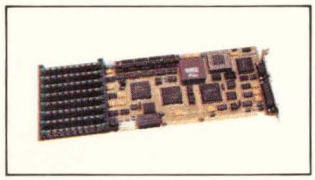

C386 80386 CPU Board with cache memory
Executes the most powerful programs with blazing speed.
- Runs MS/DOS® PC/DOS® OS/2® 386 UNIX,™ XENIX,™ PC-MOS™ and compatible O/S and software.
- Choice of 25 MHz or 33 MHz speeds.
- 80387/Weitek 3167 math co-processor socket.
- 64K on-board cache, expandable to 128K.
- 2 serial, 1 parallel port.
- Field upgradable "Flash" BIOS.
- 1 or 2MByte RAM upgradable to 16MBytes on-board.

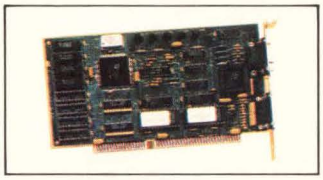

V2GA (1024x768) Video Controller Card.
With internal caching for superior speed.
- Up to 1024x768 pixels with 256 simultaneous colors from a palette of 262,144.
- 16-bit data path to the CPU.
- 512K and 1MByte DRAM memory configurations available.
- Up to 32-bit memory data bus.
- Analog and Digital outputs.
- Accommodates downward register compatibility with EGA/CGA/MDA and Hercules™ standards.
- Extended graphics mode drivers included for all major programs.

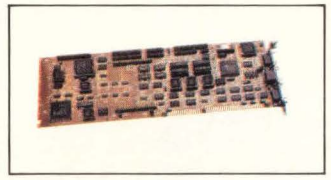

MACRO-4 Macrocard
Displaces 3 less-functional AT boards.
- Used with any TMS CPU card, it forms a complete computer, freeing slots for other function cards.
- User-programmable, menu-driven configuration with high-speed super VGA (1024x768) resolution.
- Up to 1MByte video memory.
- 2 RS232/serial and one parallel port.
- Floppy and hard disk controller support for multiple drives.

Extreme Value. No one makes as many rugged, mission critical CPU and option cards as we do, and no one can match us on quantity prices. So for quality cards at the best prices for quantities of 100 or more, call Texas Microsystems.

Extreme Support. No one offers a better level of support, including advice and consultation during system design. All our boards are guaranteed for twelve months. And no one has more application engineers to respond to your needs.

Extreme Performance. All our cards deliver performance equal to, and in many cases better than, cards used in leading desktop systems. And they deliver that performance under the most extreme conditions, even after hours enduring the industry's toughest test regime.

Extreme Choice. No one has a bigger choice of options for the industrial and mission critical market than Texas Microsystems. CPU cards from 8088 to 80486, at various speeds and on-board RAM configurations. Macrocards that can do the work of up to three standard AT cards. Solid State Winchester Disk cards that emulate hard drives, without the limitations. And a V²GA Video Controller card with internal caching for superior performance.

If you're looking for the best value in industrial CPU and option cards, and you're looking to buy in quantities in excess of 100, calling the 800 number below could save you thousands.

Mission Critical Micros is a trademark of Texas Microsystems Inc., all other trademarks mentioned are registered, trademarked or servicemarked by their respective manufacturers.
* Low volume resellers - ask about our special VAR programs.

Extreme Discounts for Large Orders

	C386/25 (CPU's priced with 4MB RAM)	N286/16	V²GA	MACRO 4
Qty. 1*	$3020	$1660	$450	$1295
Qty. 100 (per Unit)	$2085	$1145	$310	$895

1-800-627-8700

TEXAS MICRO SYSTEMS

CIRCLE NO. 49

ing devices for internal consumption, but have no immediate plans to make them available as standard products.

Selecting a U Interface device

One factor to consider when evaluating U Interface devices is operating voltage. Until 1992, the ANSI specification permits the use of a 2-V operating voltage. After that, chips must employ a 2.5-V operating voltage. At this time, though, it appears that most of the vendors who had originally announced 2-V devices have upgraded to 2.5 V.

Another feature to consider when selecting U Interface devices is power dissipation. The higher the power dissipation, the smaller the loop distance a device can support when operated from line power. Roughly 1,500 ft. of loop distance is lost for each additional 150 mW of power consumed by the U Interface, according to Siemens' Sirsi.

In Europe, low power is more important because the 4B3T standard specifies line powering for the U Interface. In North America, however, power dissipation isn't as critical because ISDN line powering isn't required; the U Interface derives its power locally at the customer premises.

Sirsi agrees with those who believe that line powering may yet be adopted for North America. But even if it's not, he argues, when the U Interface is used in non-ISDN applications such as pair gain and Digital Loop Carriers, line power is required and low-power dissipation is essential.

At 200 mW, Mitel's U Interface provides the lowest power consumption of any device, followed closely

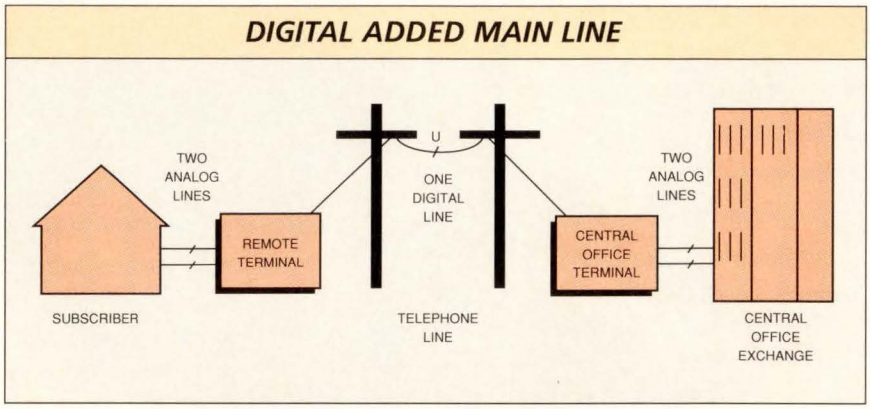

Until nationwide ISDN operability is achieved, the principal applications for U Interface devices may actually be in non-ISDN applications such as Digital Added Main Line (DAML). Also known as pair gain, this application lets the bandwidth of an existing analog line be doubled by incorporating a U Interface.

U Interface tackles non-ISDN applications

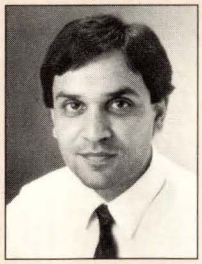

Until customer demand for ISDN emerges, the largest applications for ISDN U Interface hardware are likely to be in non-ISDN applications. Two of these applications are Digital Added Main Line (DAML) and Digital Loop.

DAML, also known as pair gain, is a technique that lets telephone companies increase the bandwidth of existing phone lines. With this capability, they can service new accounts—for instance, provide service to a new residence or add a second line—without having to run new lines from the central office.

Traditionally, pair gain has been achieved via frequency modulation, an analog technique that's subject to interference and noise. An alternative method is to use a U Interface, whose two 64-kbit/s channels can be used to double the bandwidth of existing copper wires. In this application, two analog lines from the central office are multiplexed onto a single U Interface line and transmitted using the 2B1Q format to a remote terminal. The remote terminal converts the digital signal back into two analog lines.

U Interface technology is ideal in pair-gain applications for a number of reasons. First, the U Interface's two 64-kbit/s channels support two subscribers on the same analog line. Second, the U Interface's independent 16-kbit/s channel lets a wide range of versatile signaling information be transmitted along with the data. And third, the U Interface's digital technology provides superior immunity to noise and crosstalk.

DAML lets telephone companies increase the bandwidth of existing lines.

Another non-ISDN application for U Interface technology is the upgrade of digital modems. In a typical 56-kbit/s modem, data is input to the modem from a PC or fax, for example, at 24 kbits/s.

Voice is input via a coder/decoder at 64 kbits/s and is compressed via ADPCM to 32 kbits/s. The two data streams are multiplexed in the modem's line interface and sent as a 56-kbit/s data stream to the central office via standard analog lines using a biphase protocol. There, the central office termination separates the analog and digital signals. By replacing the modem's line interface with a U Interface, the modem can now accept a 64-kbit/s data stream on the digital side.

On the analog side, the ADPCM compression circuitry can be eliminated, with the 64-kbit/s codec output run directly to the U Interface. The U Interface then transmits the two 64-kbit/s data streams over the existing analog line to the central office.

Because a 2B1Q, rather than a biphase protocol, is used over the telephone lines, transmission distances can be increased from 3 to 5.6 km.

Namraj Johal, product marketing specialist, communications ICs, Siemens

Solid SRAM Performance.

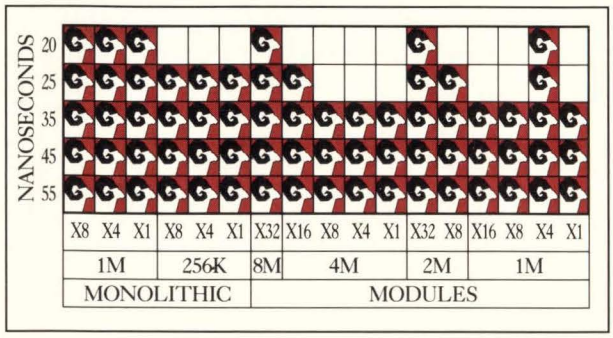

Experience and expertise that provide solid footing now, and a clear vision for the future. That's EDI.

Our commitment to Megabit Density SRAM Technology is continuously bringing high performance system designers uncompromising reliability and speeds as fast as any on the market today. And we can deliver them in a variety of package styles to accommodate your highest density requirements.

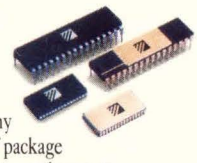

Innovation manufactured under the strictest quality control including MIL.-STD.-883 Para. 1.2.1 standards.

And for tomorrow? EDI is the company to watch. For a sneak preview of the new 1991 products call or write for our 1991 Data Book.

EDI
ELECTRONIC DESIGNS INC.
Committed to Memory.

Electronic Designs Incorporated: Corporate Headquarters: 42 South Street, Hopkinton, MA 01748 USA 508-435-2341, FAX: 508-4356302
Electronic Designs Europe Ltd.: Shelley House, The Avenue, Lightwater, Surrey GU18 5RF United Kingdom, 0276 72637, FAX: 0276 73748

CIRCLE NO. 50

Conformance testing drives ISDN compatibility

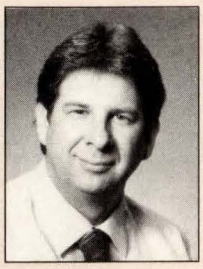

ANSI recently made great strides toward solidifying the U Interface standard when it agreed upon three critical maintenance functions. The first, Insertion Loss mode, requires that the network termination (NT-1) be able to generate a tone to the central office. The second, Quiet Mode, lets the NT-1 be temporarily disabled. The third, Metallic Loop Termination, provides for a signature at the NT-1 used by a loop tester at the central office to determine how the loop is terminated (ISDN or analog).

Unfortunately, though the ANSI U Interface standard is largely complete, true interoperability still may be a way off. One reason for this is that the standard has been a moving target, so equipment designed at different points in the standardization process may be incompatible. Another reason is that the U Interface standard doesn't provide an NT schematic. Instead, it includes a number of complex equations, performance targets and recommendations, each of which are open to some interpretation.

Conformance test necessary

A set of conformance tests that must be met to gain certification are needed to ensure interoperability at all levels in the ISDN network. The National Institute For Standards has been established in the government sector to develop and police such a set of tests. Eventually, all ISDN equipment used in government applications will have to be certified by this group first.

Though the U Interface is complete, true interoperability still may be a way off.

In the commercial sector, a set of tests is being developed by a group known as ISDN Conformance Testing Layer 1. Working under the ANSI umbrella, this group will establish conformance criteria for Open Systems Interconnection Layer 1, which specifies U Interface silicon functionality.

While the Layer 1 group is a start, what's really needed to drive interoperability in the commercial sector is conformance testing at all equipment levels, including terminal equipment, NT-1s, line terminations and even central office switches. To make conformance testing effective, a centralized group with some clout must be responsible for it.

In this regard, many vendors are looking toward Bellcore, an independent telecommunications R&D organization conceived at the time of AT&T's divestiture and funded by the seven regional Bell operating companies. Bellcore is already providing conformance testing services on a limited basis. But its role needs to be expanded.

In the meantime, the ANSI U Interface standard should be temporarily frozen until feedback from field trials is in and some of the bugs can be worked out. Once field data has been obtained, if additional refinements to the standard are indicated, they should be implemented with due regard to the economic ramifications.

Freezing the standard will let silicon and semiconductor vendors build and customers purchase equipment without fear of obsolescence. It will also promote a standardization process driven by both demonstrated need and theoretical considerations.

Mike Pauzer, product marketing manager, AT&T Microelectronics

by devices from Siemens and National/SGS-Thompson at 300 mW and Motorola at 550 mW. At 850 mW, AT&T's two-chip device consumes the most power. But a sub-400-mW single-chip device with a 50-mW power-down mode is scheduled for introduction from AT&T in 1991.

System-level considerations

Users should be wary when evaluating power claims, says AT&T's Pauzer. "Power consumption is heavily dependent on the characteristics of the loop that the device is driving," he says. Consequently, Pauzer recommends that vendors specify a loop configuration when quoting power dissipation. Motorola's Mouton agrees, pointing out that factors such as clock speed, loading and the number of support devices required for the U Interface all have a significant impact on system-level power consumption.

Though factors such as pulse voltage, cost and power dissipation are important considerations when evaluating U Interface silicon, designers must look beyond the vendor's specifications. One reason is that the ANSI specification doesn't establish strict performance levels of the U Interface. It specifies the performance and compatibility of the NT equipment that incorporates the U Interface silicon.

Because of distortion and attenuation introduced by components such as the U Interface transformer and lightning-protection circuitry, the NT's output signal will always be degraded relative to the U Interface device's output signal. As a result, the U Interface silicon must actually outperform the specification in many areas. The degree of performance depends on the quality of other components in the NT equipment.

Because the NT-1 design has such a significant impact on U Interface performance, many chip vendors have gone as far as providing application notes that define a complete NT-1. In this way, the vendor can guarantee compliance for at least one NT-1 design.

ISDN development support

Another factor that designers must consider when choosing a U Interface device is development support. Presently, most vendors are supporting their devices with evaluation boards and some type of bundled software. One of the advantages of Motorola's MC145494EVK development board is that it provides two U Interface devices as well as an S Interface. Because the board implements both a NT and LT, it may be used to simulate the complete link between the customer

WE APPROACH THE MEMORY BUSINESS FROM ONE POINT OF VIEW...

YOURS.

At Micron Technology, we offer a full line of leading-edge RAM components in speeds and packages for virtually any application.

But the memory business is more than just parts — it's people.

That's why we offer a total commitment to service and support, and a team of engineers and technical support personnel that are the most experienced memory professionals in the industry.

Because in the memory business there's only one point of view that counts. Yours.

Component Product Family	Memory Size (Total)	Org. (Bits)	Speed (ns)	Package						Special Features	Availability	Military Qualified
				DIP	ZIP	SOJ	PLCC	TSOP	PQFP			
SRAMs	1MEG	x1,x4,x8	25-45	X		X				x4 option: OE	Now	2H90
	256K	x1,x4,x8	20-45	X	X	X				x4 option: OE	Now	X
	64K	x1,x4,x8	12-45	X		X				x4 options: Separate I/O, OE	Now	X
	16K	x1,x4,x8	12-45	X		X				x4 options: Separate I/O, OE	Now	X
Cache Data SRAMs	288K	x9	14-34				X			486 Compatible, Self-timed write, Fast toe 7ns, 486 Burst and extended burst support	Samp: 1H91, Prod: 2H91	
	144K	x18	20-35				X		X	386/486 Compatible, Fast toe 8ns, Auto write completion, Parity bits	Now	
	128K	x16	20-35				X		X	386 Compatible, Fast toe 8ns	Now	
Synchronous SRAMs	288K	x18	15-25				X			Registered address, chip enables and write control; Data latch, Fast toe 6ns, Byte write capability	Samp: 2H90, Prod: 1H91	Samp: 1H91
	256K	x16	15-25				X					
SRAMs with Address Latch	288K	x18	15-35				X			Address, data and chip enable latches; Byte write capability, Fast toe 6ns, 3.0 Volt output buffer option	Samp: 2H90, Prod: 1H91	Samp: 1H91
	256K	x16	15-35				X					
	16K	x8	100	X		X				Intel 8051 and 8096 compatible	Now	X
	16K	x8	15-35	X		X				Compatible with high end micro controllers	Now	X
FIFOs	18K	2Kx9	15-35	X			X			Family options: 300 mil DIP package, Programmable flags	Samp: 1H91	2H91
	9K	1Kx9	15-35	X			X				Samp: 1H91	2H91
	4.5K	512x9	15-35	X			X				Samp: 1H91	2H91
DRAMs	4MEG	x1,x4,x8,x16	60-100	X	X	X		X		x4,x8 options: Write per bit x16 options: 2 WE/1CAS; 1WE/2CAS and 1WE/1CAS with write per bit	x1,x4 Samp: Now, Prod: 1H91; x8,x16 Samp: 1H91	Samp: 1H91
	1MEG	x1,x4,x16	70-120	X	X	X		X		x16 options: Byte write or write per bit	Now	X
	256K	x1,x4	100-120	X	X	X	X				Now	X
	64K	x1	100-150	X			X				Now	X
Quad CAS DRAMs	4MEG	x4	60-100			X				Separate CAS control for each DQ input/ output; Enhanced write per bit capabilities	Samp: 1991	
	1MEG	x4	70-100			X					Now	
Pseudo Static DRAM	1MEG	x8	80-120	X		X		X		Unmultiplexed addresses, Simple refresh control	Samp: 2H90, Prod: 1H91	
Dual Port DRAMs (VRAMs)	1MEG	x4,x8	80-120		X	X				CMOS, Fully static SAM, Serial input, Split read transfer	Now	Samp: 1H91
	256K	x4	100-120	X	X					CMOS, Fully static SAM, Serial input	Now	X
Triple Port DRAMs	1MEG	x4,x8	80-120			X	X			CMOS, Two fully static SAMs, Transfer mask, Split t.ansfers, Functional superset of 1MEG VRAM	Samp: Now, Prod: 2H90	

Module Product Family*	Word Size (Words)	Org. (Bits)	Speed (ns)	Package						Special Features	Availability	Military Qualified
				DIP	ZIP	SIP	SIMM					
DRAM Modules	2MEG,1MEG, 512K,256K	x36	70-120		X		X			Industry standard pin-out	256K, 512K: Now; 1MEG, 2MEG Samp: 2H90	
	4MEG,1MEG, 256K	x9	70-120			X	X			Industry standard pin-out	256K, 1MEG: Now; 4MEG Samp: 2H90	
	4MEG,1MEG, 256K	x8	70-120			X	X			Industry standard pin-out	256K, 1MEG: Now; 4MEG Samp: 2H90	
SRAM Modules	256K, 128K, 64K, 16K	x32	15-45		X					Industry standard pin-out with OE	16K, 64K: Now; 128K,256K: 2H90	1H91
	64K,32K	x16	30-45	X						Industry standard pin-out with OE	Now	1H91
	128K	x8	30-45	X						Compatible with 1MEG monolithic	Now	1H91

* Custom module and board-level product manufacturing services available.

MICRON TECHNOLOGY, INC.

2805 E. Columbia Rd., Boise, ID 83706 (208) 368-3900

CIRCLE NO. 51

U INTERFACE

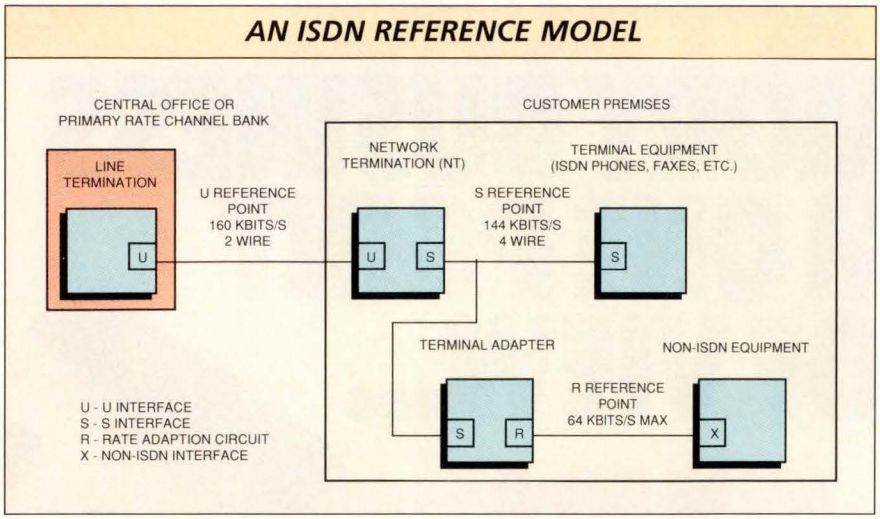

The U Interface links the customer's network termination with either a primary rate channel bank, which multiplexes several U Interface lines onto a single primary line, or directly to the central office. The NT converts between the two-wire U Interface and the four-wire S Interface, which connects ISDN equipment within the home or office. Terminal adapters are used to connect non-ISDN equipment with the S Interface.

premises and the central office.

Alternatively, the two U Interface sections may be separated to simulate either the LT or NT-1, or even to examine remote transmission characteristics. Support for layers one through three of the OSI protocol is provided through a separate 68302-based protocol-processing board, which may be used in conjunction with the development board.

Another vendor currently offering a development board with two U Interfaces and an S Interface is National Semiconductor. As with the Motorola board, the two U Interfaces let designers simulate a repeater or an entire subscriber loop. Alternatively, one U Interface and the S Interface may be used to prototype an NT-1. The board also includes a patch area that can be used for prototyping.

The availability of application software should prove crucial for the wide dissemination of ISDN technology.

Mitel's PC-compatible U Interface development board, known as the MB89010, not only includes the company's U Interface device, but emulates a full-featured digital telephone. The board may be connected directly to an ISDN network or be used in conjunction with another MB89010 to simulate a complete U Interface link between the Central Office and NT-1.

To support its U Interface device, Siemens offers a PC-compatible board known as the SIPB5000. Equipped to accept three daughterboards that interface with the board via Siemens' IOM2 communications bus, the SIPB5000 also provides three-layer OSI software support, including call control, link access protocol in the D channel (LAPD) and device drivers.

AT&T supports its device with a PC plug-in evaluation board known as the 2B1QCAK (Complete Adapter Kit). The board, together with an accompanying daughtercard, provides both an NT and an LT. By plugging a handset into both the NT and the LT, designers can simulate a complete U Interface. Using the accompanying PC-based software and working from a keyboard, designers can transfer data between the NT and LT at either the file or character level. AT&T also offers a loop simulator, which lets designers evaluate the U Interface over a variety of loop configurations.

While the availability of ISDN evaluation hardware is crucial for developing and debugging subscriber loops, the availability of application software should prove equally crucial for the wide dissemination of ISDN technology. Thus far, ANSI, working in conjunction with the ISO, has instituted a software architecture for ISDN based on the seven-layer OSI model.

Software support

Layer one is a hardware-specific layer whose device drivers let upper layers communicate with the B and D Channel control and data registers. These drivers are unique to particular ISDN hardware configurations. Layer two, the Data Link Layer, includes D Channel LAPD and B Channel LAPB support and is responsible for providing reliable, error-free transfer of user data to and from the target terminal.

Layer three, the Network Layer, is responsible for establishing, maintaining and terminating connections between network nodes. A variety of protocols, including X.25 and CCITT Q.931 (used by both AT&T and Northern Telecom) may be used to implement layer three. Layers four through seven are application-specific and are covered by a wide variety of standards.

While support for these layers is still limited, a number of vendors are beginning to port and upgrade software from other telecommunications applications for use in ISDN. But until hardware and interoperability issues are resolved, software vendors aren't likely to make a significant investment in ISDN tools and application programs. ■

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 223**
Advanced Micro Devices
(408) 732-2400 . . . . . **Circle 224**
Mitel
(613) 592-5630 . . . . . **Circle 225**
Motorola
(512) 891-2000 . . . . . **Circle 226**
National Semiconductor
(408) 712-5000 . . . . . **Circle 227**
Northern Telecom
(615) 734-4000 . . . . . **Circle 228**
Siemens
(408) 980-4518 . . . . . **Circle 229**
SGS-Thomson
(602) 867-6100 . . . . . **Circle 230**
Zilog
(408) 370-8000 . . . . . **Circle 231**

Is your emulator giving you the whole picture?

MICE-V-486. 33MHz Emulation. Real features. Real-time.

Without real-time emulation you never know how your product will perform until it has to fly. Traditional in-circuit emulators slow your target to collect, display or reprogram trace. Or even stop emulation (or your target) to load complex triggers. When your emulator can't show you what's actually happening you risk missing a bug that will sneak from your prototype to the finished product.

MICE-V-486 lets you see it all.

- ▼ Real-time emulation to 33MHz.
- ▼ Complex, sequential triggers, loaded without slowing the emulator or target.
- ▼ Access to the fully qualified trace buffer during full-speed emulation.
- ▼ High level language debug.
- ▼ Probe kits for 386, SX, 376 and 286 support.

Most in-circuit emulators require partially or completely functional hardware to operate correctly. MICE-V-486 has a unique *Isolation Mode™*, requiring only a working clock signal. Logic analyzer taps are conveniently located to give you access to critical timing information. *MICE-V-486 provides absolutely the fastest method for debugging non-functional 486-based hardware.*

Microtek also has real-time emulators and source-level debuggers for 68000, -020, -030 and 80C186.

So, stop wasting development time because your emulator isn't real-time. Call us, and get your product to market fast.

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

ASIA OFFICE – Taiwan - MICROTEK INTERNATIONAL 886-2-723-5577/Japan - CORE Digital 81-3-7955171
EUROPE OFFICE – Germany - ALLMOS 49-89-8570000/France - M.T.E. 33-1-39618228/U.K. - ARS Microsystems 44-276-685006

CIRCLE NO. 53

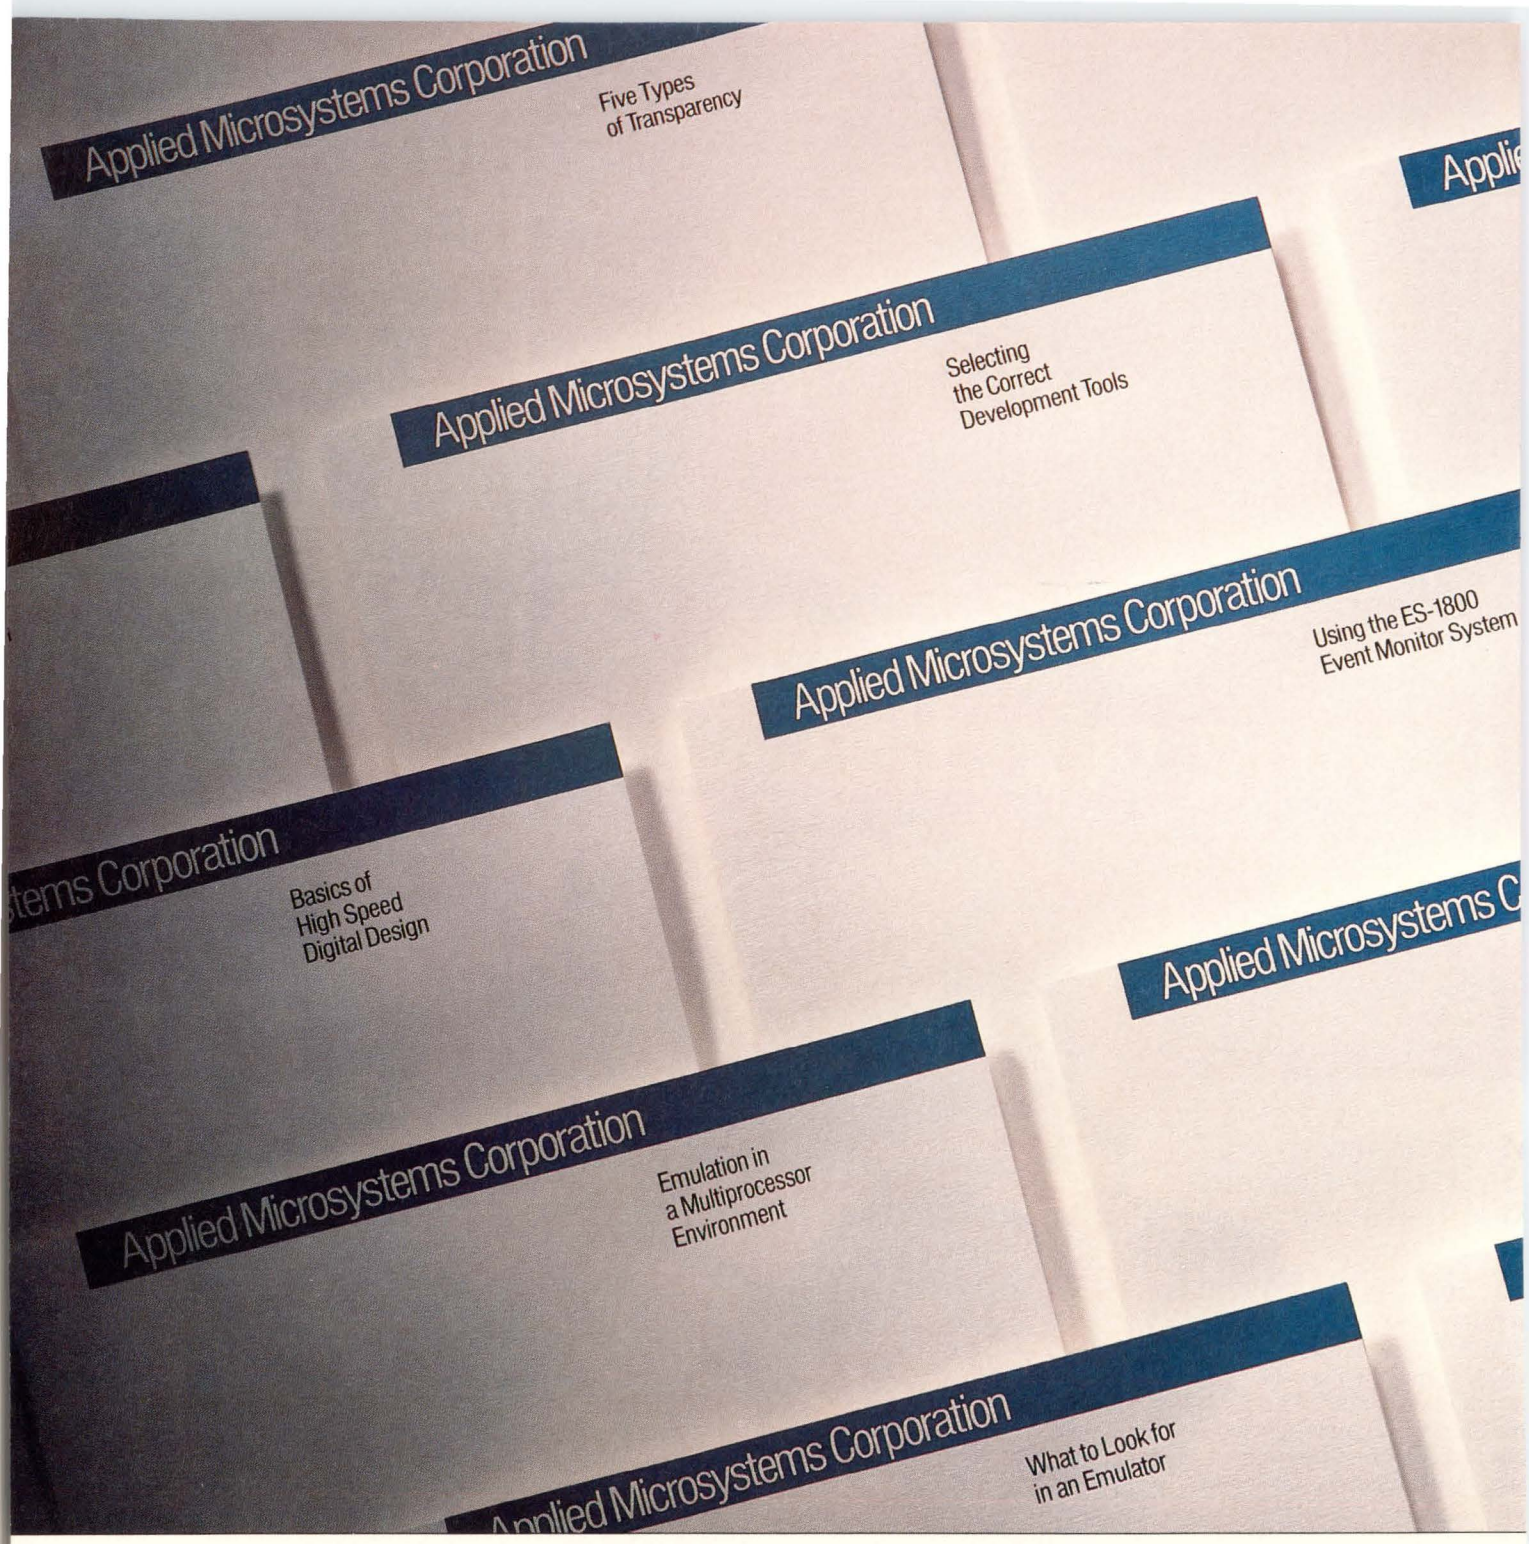

Before you buy development

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

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. 54

High Data Flow Requirements Clouding Your VMEbus Design?

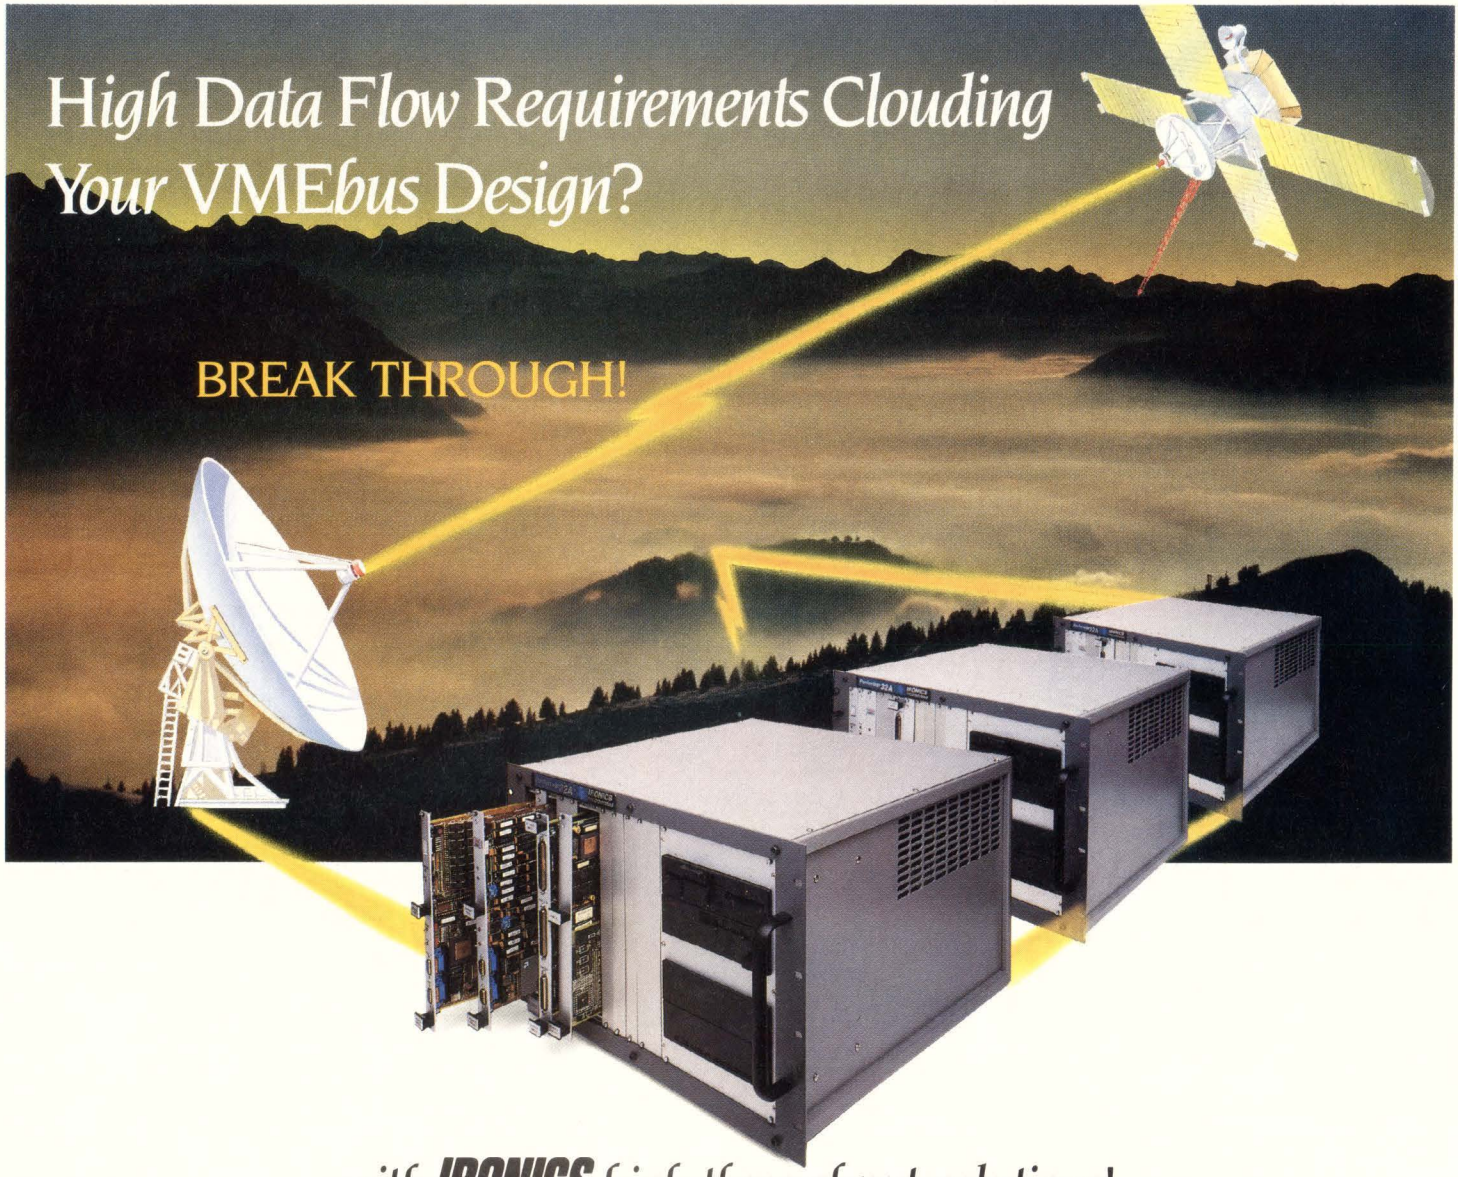

with IRONICS high-throughput solutions!

Selecting a High Performance CPU may be the easiest part of your VMEbus design. But how do you overcome the REAL BOTTLENECKS to SYSTEM performance... getting data in and out, across, and between chassis? — not to mention the headaches in the design process itself.

Ironics can help you burn through the high-data-flow haze... not only with hot CPUs, but with a complete line of BOTTLENECK BUSTING BOARDS.

HIGH-THROUGHPUT SOLUTIONS FOR REAL-TIME I/O BOTTLENECKS

HIGH SPEED DMA DAUGHTER BOARDS
- 92 MBytes/sec, sustained, real-time I/O data rates directly into processor memory.
- designed for the Ironics family of RISC processors

PARALLEL I/O DAUGHTER BOARD
- data rates up to 38 MBytes/sec
- designed for use with Ironics Full Speed Data Transporter ™

HIGH-THROUGHPUT SOLUTIONS FOR VMEBUS CONGESTION

FULL SPEED DATA TRANSPORTER
- block transfer rates up to 32 MBytes/sec across VMEbus
- optimized for Ironics RISC and CISC processor boards

VSB DAUGHTER BOARDS
- alternate path for high speed data transfers between CPU, Memory and I/O boards
- designed for the Ironics family of RISC and CISC processors

HIGH-THROUGHPUT SOLUTIONS FOR INTERCHASSIS DATA TRANSFERS

MULTICRATE PIPELINE ™
- 20-30 MBytes/sec, sustained, inter-chassis transfer rates
- multiple system link for VMEbus chassis and/or VME compatible SUN Workstations ™

HIGH SPEED DMA DAUGHTER BOARDS
- 92 MBytes/sec transfers between processor boards or VMEbus chassis
- designed for the Ironics family of RISC processors

HIGH-PRODUCTIVITY SOLUTIONS FOR REAL-TIME DEVELOPMENT

- CASEworks/RT ™
- Integration Toolkit ™
- High Level Debugger
- Application Simulator
- Real Time Kernels
  - VXWorks™
  - pSOS, pSOS+™
  - OS9™
  - VRTX™
- The most complete, highly integrated, and flexible Real Time tool set

For more information on Ironics' Hot Processor Boards, High Speed I/O Boards, High-Throughput Data Transfer products, Real-Time Target Systems, and High Productivity Development environments, contact:

Ironics Incorporated
798 Cascadilla Street
Ithaca, New York 14886

phone: ***607-277-4060***
fax: ***607-272-5787***
telex: ***705742***

IRONICS Incorporated

CIRCLE NO. 55

Cache controllers tread a rocky path toward integration

As cache memories multiply in PCs and workstations, cost and speed requirements demand integration while system designers plead for flexibility.

Ron Wilson
Senior Editor

The volume of personal computers from production lines like this Dell Computer facility is rewriting the book on cache design. Now, integration and cost management are the keys for a large part of the market, even at the expense of proprietary caching algorithms.

Not long ago only big, esoteric computers had caches. Today even personal computing microprocessors like the 386 and 486 run too fast for DRAM to keep up, so even PCs depend on cache architectures. This shift from technological oddity to mass-market commodity is forcing cache controllers into higher and higher levels of integration.

But at the current CMOS gate density, this integration has to come at the expense of flexibility. The more features you combine onto one die, the fewer options you leave to the system architect. This is a problem, points out Tom Hosey, senior product marketing engineer at SGS-Thomson (Phoenix, AZ). "You can't solve the cache optimization problem with just one solution unless you can get everyone to run the same software," he says. And software isn't the only issue. Differences in memory system design or addition of more CPUs to the system can radically alter cache requirements. So as vendors produce more highly integrated controllers, they are—intentionally or unintentionally—producing increasingly specialized parts.

To understand the significance of the chip vendors' decisions, it's important to look not only at the configuration of a particular chip but also at the more general question, what is a cache supposed to do, anyway?

In principle, a cache subsystem has to do only two things: minimize use of memory bandwidth and maintain data coherency. In practice, though, each of these substantial-sounding goals quickly dissolves into a morass of strategies, alternatives and trade-offs. How the goals are defined, how the cache designer sets about meeting them, and even how the results are measured vary from design to design.

Minimizing bandwidth

Fundamental to the notion of a cache is preserving memory bandwidth—that is, reading or writing DRAM as seldom as possible. To do this, the controller has several variables at its disposal: cache size, cache organization and write-cycle strategy.

Obviously, other variables being equal, the larger the cache the higher the hit rate. But since cache SRAM is much more expensive than DRAM, bit for bit, there's always an economic incentive to cut back on cache size. In addition, there are several common situations—including multitask-

ing—where a big but dumb cache can perform poorly. For these reasons cache organization is an important issue as well.

Big caches tend to be direct-mapped—that is, a particular line in the cache is capable of caching only memory locations with a particular pattern of low-order address bits. To increase the cache hit rate without making the cache gigantic, architects often build multiway set-associative caches. These are essentially several small direct-mapped caches sharing a single controller. On real programs, especially in multitasking environments, the improvement can be startling.

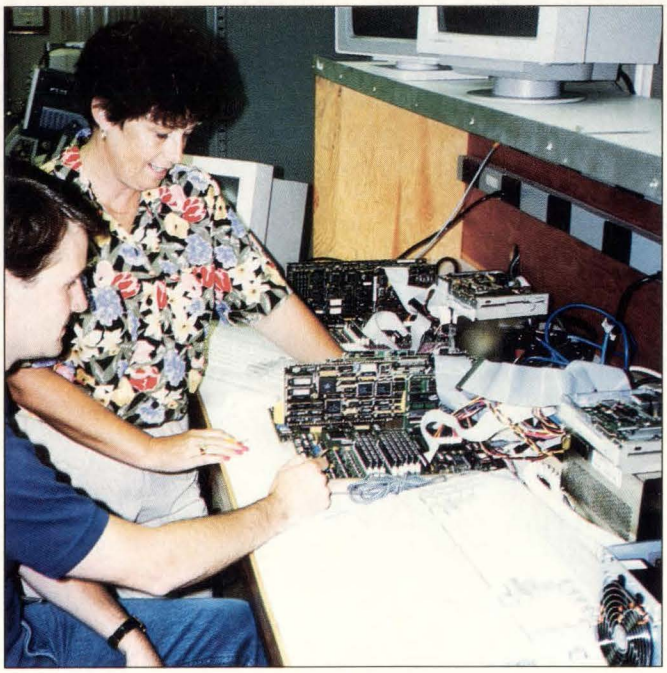

Intel cache marketing manager Sheila Cardno and product manager Mark Casey maintain that the PC market demands low cost and compact design, rather than the last grain of speed, from a cache. These considerations influenced the decisions Intel made in designing its cache.

"We've collected code traces as large as 30 million cycles, and run personal computing and SPEC benchmarks," says Mark Casey, product manager at Intel (Santa Clara, CA). "The results show that the 82395 cache chip, using four-way set associativity and 16 kbytes of data RAM, is as fast as, or faster than, any cache available on a PC today." Since there are PCs with 64-kbyte caches out there, the claim is a dramatic one.

Equally important to managing bus bandwidth is what the cache system does on a write cycle. The simplest caches just let the CPU write directly into main memory—so-called write-through. To save the CPU's having to wait for a slow DRAM write cycle, more-sophisticated cache controllers permit posted writes—the CPU writes into a FIFO, and the cache controller then writes into DRAM while the CPU continues execution.

Posted writing, given a deep enough FIFO, helps avoid stalling the CPU, but it doesn't do much to minimize bus usage—you're still doing just as many writes as before. The most aggressive cache controllers have a third option, copy-back. In this scheme, the CPU writes just into the cache. The data stays in the cache, only getting copied back into DRAM when the cache line is about to be replaced. Thus if a program uses memory locations for working storage, most of the read-write transactions to a particular location will be cache hits, and the number of write cycles to DRAM will be cut dramatically.

The coherency issue

As the cache strategy gets more aggressive, a related issue gets increasingly complex: data coherency. In even the simplest write-through cache, data coherency is a design consideration.

For instance, in a PC the DMA controller may write data into a section of memory that's currently being cached. The cache controller must somehow keep from using the old, cached copy of the data for further processing. "There are many levels of complexity you can apply to this problem," says Henry Wurzburg, strategic and technical marketing manager at VLSI Technology (San Jose, CA). "But most controllers take the simplest approach—they just declare regions used by DMA as noncachable."

Controllers that permit posted writes have a more difficult task. They must ensure that if someone attempts to read a location whose current data is in the write FIFO, the new data gets read. Many controllers do this by the simple but costly expedient of forcing all of the writes to be completed before the next memory read can take place. More-sophisticated caches will snoop the main memory bus, and take action only if one of the posted locations is actually addressed.

In a copy-back scheme there's no alternative—you have to snoop the bus. Cache controllers can accomplish this simply, by watching for cached addresses and forcing a write if one is detected. Or the problem can take on incredible complexity—in multiprocessing systems, several cache controllers may negotiate to determine who owns the current value of a particular variable, and whether the owner cache will supply data to another cache or to main memory.

The variety of options in size, organization, write strategy and coherency strategy lets system architects tune their performance to a particular application. But an arbitrarily complex controller with an arbitrarily large tag RAM won't fit on an arbitrarily small die—vendors have to make choices. The system designer's selection of a controller implementation thus also becomes a matter of trading off cost, integration and sophistication.

Levels of integration

The lowest level of integration commonly used today might be termed building block. Vendors such as Texas Instruments (Dallas, TX) and SGS-Thomson have combined SRAMs with comparators and some control logic to make integrated tag RAMs. And many vendors have added address latches, select logic and burst counters to conventional or synchronous SRAMs to make CPU-specific cache data RAMs.

The arguments in favor of the building-block approach center on flexibility. "People building high-performance systems don't want a commodity cache controller—they want to be able to differentiate," says SGS-Thomson's Hosey. "All the

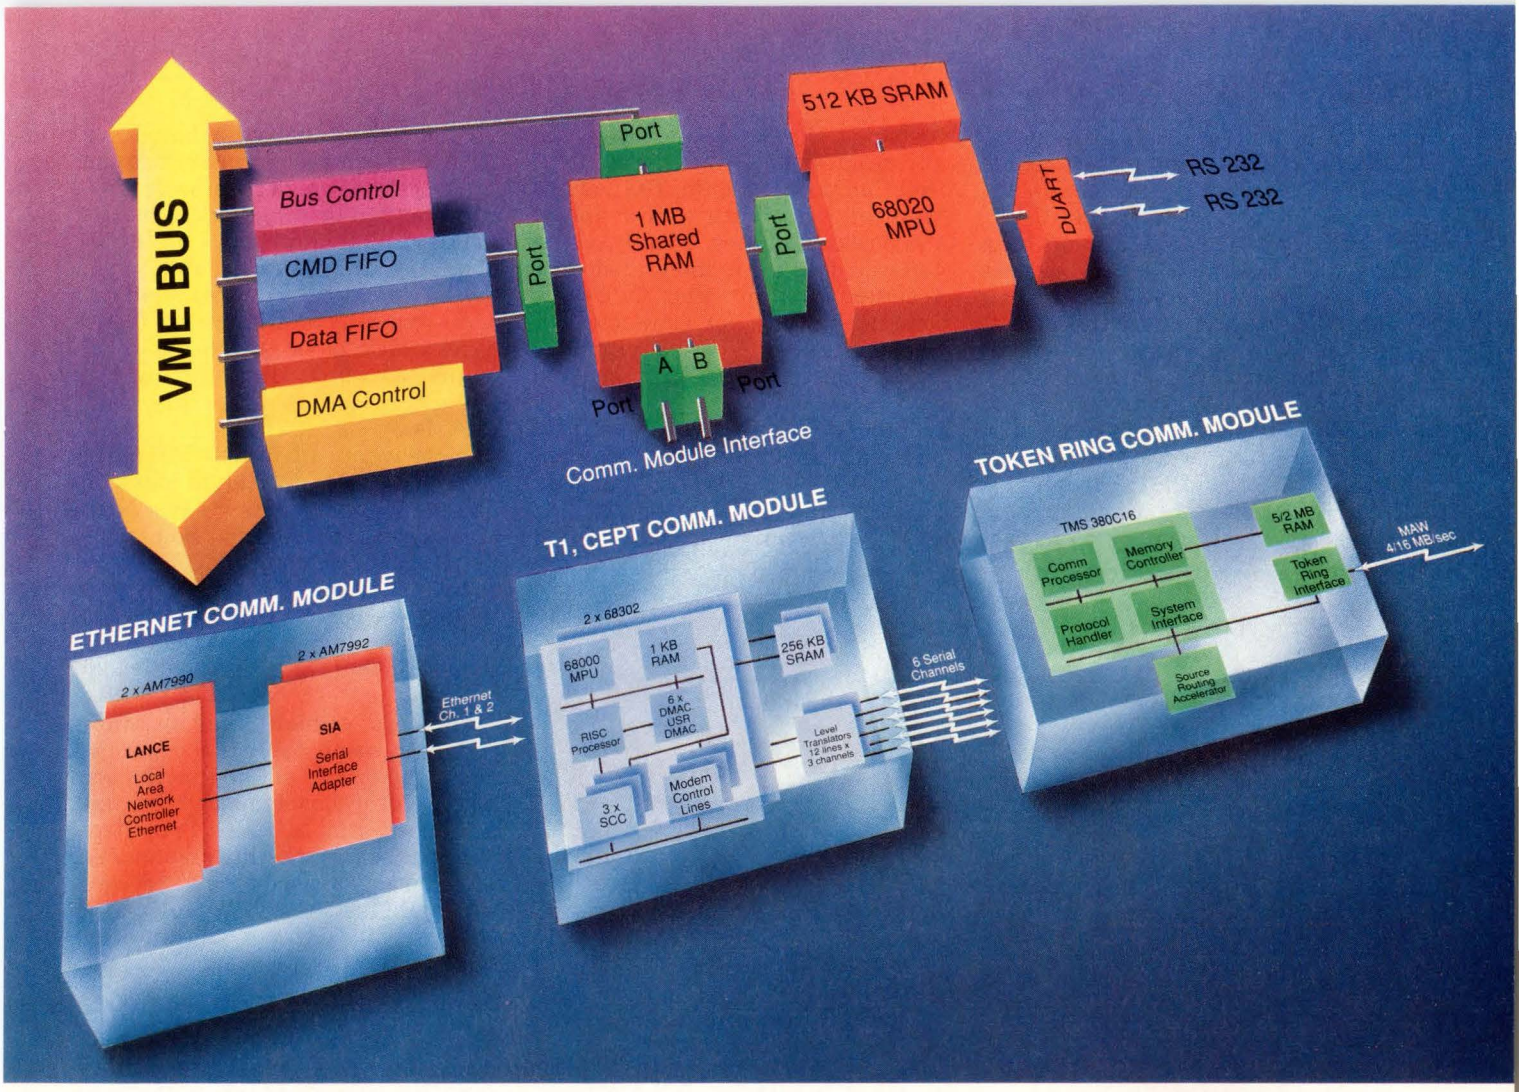

Great connections! Bridge, Router and LAN/WAN Gateway on one VMEbus board.

- Two Ethernet Ports (TSVME 551E).
- Six X.25 Ports, with rates up to T1 and CEPT 2.048 Mbits/sec (TSVME 551X).
- Token Ring (IEEE 802.5), with Source Router chip support (TSVME 551T).*
- Multiprocessor Triple-Bus™ architecture crunches protocol stacks up to 3 times faster.
- Unique DMA design provides VMEbus data block transfer rates in excess of 20MB/sec.

Your network systems design just got simpler. Smarter. More powerful.

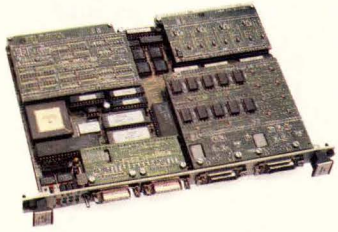

Themis' "open systems" strategy puts Bridge, Router and Gateway functions on one VMEbus board. The TSVME 551- improves network and system performance by providing a direct port-to-port path for Bridge and Router traffic. Now *only* Gateway traffic uses the VMEbus.

Triple-Bus™ means performance.

The TSVME 551- Triple-Bus and proprietary high-performance DMA engine make possible simultaneous data transfers between:

1. 68020 MPU and 512 Kbytes of local memory,
2. DMA data buffer and VMEbus,
3. Communications controller module and 1 Mbyte of shared memory.

Step up to the Themis TSVME 551-. It's multiprocessing, multi-port and supports multiple network technologies. All on one VMEbus platform.

Free Technology Report.

Ask for the TSVME 551- data sheet and a copy of our VMEbus Technology Report "Triple-Bus Architecture: The Advantage of Concurrent Traffic Flow."

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

*Availability to be announced.

CIRCLE NO. 56

high-end and mid-range workstation vendors I know of are using discrete solutions, as is nearly everyone working with the 486 CPU."

Using off-the-shelf tag and data RAMs with an ASIC controller does let system architects choose their own organization and write strategies, and it leaves room for whatever coherency algorithms are necessary. And building blocks make it much easier to choose the right cache size.

With controllers that include tag RAMs, the size of the tag memory is obviously fixed. But since there must be one tag entry for each line in the cache, the integrated controller fixes the number of lines in the cache. A bigger cache necessarily means longer lines. That can harm performance.

There are also risks in using the building-block approach, most obviously including longer design time and higher cost. But at least one design manager hasn't had these problems. Charlie Sauer, director of advanced systems development at Dell Computer (Austin, TX), relates his experience with a new line of 486-based machines. "Ideally, on the 486 you'd like an integrated part designed just for that CPU—it's coming, but it's not available yet. Lacking an integrated part, we found that a discrete solution had the lowest cost.

"I don't think there's a significant difference in design time between integrated and discrete implementations," Sauer continues. "It's a relatively short design either way."

A more serious challenge to the building-block approach, at least in the long run, is speed. Two critical loops in a cache system—one from the CPU address lines through the tag RAMs, comparators and logic and the other through the data RAMs—must both be short enough to fit into a CPU cycle. Given address skew allowances and data setup times, that means that both loops have to be somewhat faster than one clock cycle on modern processors—and that can be less than 20 ns.

The problem is that building-block designs put lots of pins in these two critical loops. And pin delays—typically on the order of 5 or 6 ns for CMOS devices—can eat up all of the allowable time, forcing the designer to use absurdly fast SRAMs.

"Somewhere above 40 MHz, it's going to be more difficult to continue using discrete parts," says Nick Mati, project manager at Solbourne Computer (Longmont, CO). "At 50 or 60 MHz it gets very squeaky getting by with 10-ns tags. And on top of the timing problems, at these speeds layout becomes extremely difficult."

Yet chip vendors remain confident. "To some extent, the performance limits on building blocks are a question of the technology the vendor has available," says Paul Reeber, strategic marketing manager at TI. "The way to cope with I/O delays is to build faster logic. We have people doing 50-MHz, no-wait-state designs today with our parts."

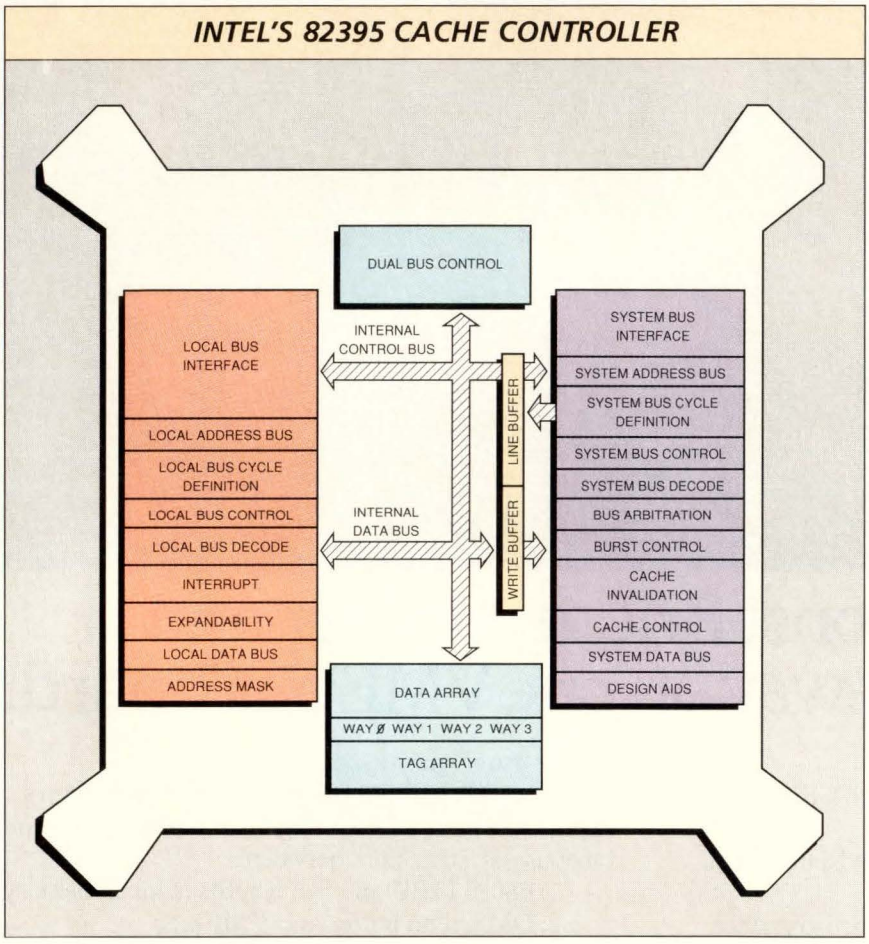

By combining control logic, tags and data RAM onto one die, Intel gets a lower chip count and saves system designers a lot of detail work. But whether the 395's four-way set-associative organization and write buffer will compensate for its silicon-imposed 16-kbyte capacity remains to be seen.

Pulling in the tags

Another way to eliminate some off-chip delays is to pull the tag RAMs and their comparators onto the controller IC. This also has the advantages of delivering the system designer from the perils of critical tag logic timing, and of reducing chip count.

It was probably the latter consideration that led Intel to introduce the pioneering 82385 cache controller for use with the 386 CPU. The part contains both the control logic and the tag RAMs for a 32-kbyte direct-mapped or two-way set-associative cache system, and it has become something of a standard in mid-range PCs.

The part also has its detractors among both competitors and prospective users. Criticism tends to focus on the alternatives Intel had to foreclose in order to achieve integration. For instance, the part is designed to work with 32 kbytes of data SRAM. It can support 64 or even 128 kbytes, but only with external glue and the penalty of increasing line size. At 128 kbytes of data RAM, the 385 must replace 16 bytes of data every time there's a cache miss. Also, the part is restrictive in its write strategy: it supports

NETWORKING SOLUTIONS

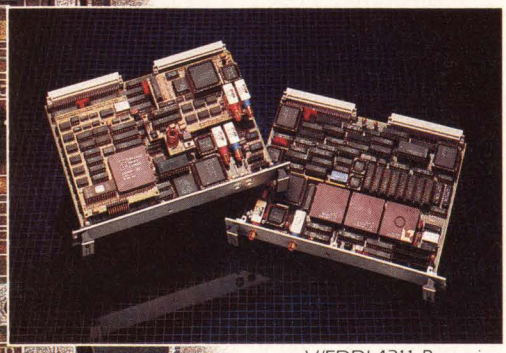

V/FDDI 4211 Peregrine
V/FDDI 3211 Falcon

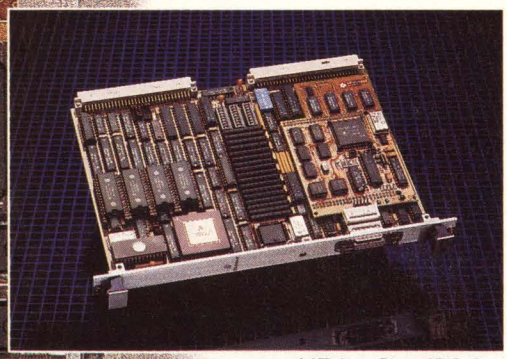

V/Token-Ring 4212 Owl

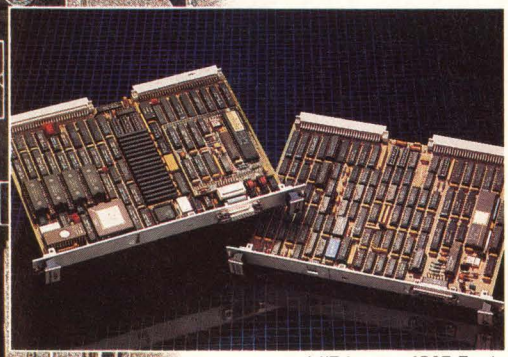

V/Ethernet 4207 Eagle
V/Ethernet 3207 Hawk

FDDI, TOKEN-RING AND ETHERNET COMMUNICATIONS

The need to network has never been greater. Diverse processing platforms, distributed architectures, client-server, departmental and workgroup environments all contribute to increased demands on the network. System and network designers need a proven source of technology solutions for the wide range of networking and communication application problems they face. Interphase delivers those solutions.

FDDI, TOKEN-RING AND ETHERNET SOLUTIONS

Interphase has long led the industry in high-performance VMEbus peripheral controllers, and that same leadership is now evident in networking node controllers. Interphase has FDDI, Token-Ring and Ethernet solutions for virtually any VMEbus system application challenge.

PROVEN FDDI SPEED AND INTELLIGENCE

Interphase's FDDI 100 Mb/s offerings are a logical choice for the industry. The V/FDDI 3211 Falcon received *UnixWorld* magazine's Product of the Year designation and was the industry's first 6U VMEbus FDDI solution. Interphase's newest FDDI product is the V/FDDI 4211 Peregrine, a RISC-based high-performance node controller capable of link level operation or on-board protocol processing. The Peregrine provides single or dual attach configurations, with SMT (Station Management Software) running on-board, all in one 6U VME slot.

TOKEN-RING RESULTS

The V/Token-Ring 4212 Owl is an ultrafast Token-Ring node controller based on the partitioned architecture of Interphase's proven Eagle class of controllers. The Owl facilitates connectivity of UNIX® systems, workstations, supercomputers or any other VMEbus system into an IBM® environment using IEEE 802.5 Token-Ring. This multiple processor design provides an elegant queued interface to the system supporting IEEE 802.2 LLC, and a flexible 4 or 16 Mbit interface to the Token-Ring network.

ETHERNET CHOICES

Interphase also offers two Ethernet design options. The V/Ethernet 4207 Eagle 32-bit protocol platform is the high-performance standard for the industry, and offers on-board TCP/IP support. The V/Ethernet 3207 Hawk is designed specifically for cost-sensitive VMEbus applications.

GET YOUR NET WORKING NOW

No matter what your networking need – FDDI, Token-Ring or Ethernet – Interphase is ready to provide *the* solution. For more information call today:

(214) 919-9000

INTERPHASE
CORPORATION

OPEN SYSTEMS CONTROLLERS™
Disk • Tape • Networking

13800 Senlac • Dallas, Texas 75234 • (214) 919-9000 • FAX: (214) 919-9200 • NASDAQ-NMS:INPH
Interphase International
Astral House, Granville Way • Bicester, Oxon 0X6 0JT • (01144) 869-321222 • FAX: (01144) 869-247720

CIRCLE NO. 57

Redrawing the borders in PC cache systems

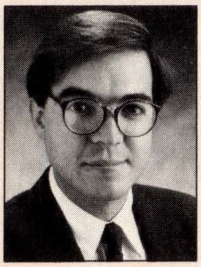

The advent of fast microprocessors has resulted in the widespread use of cache memory in personal computers. Two of the primary related trends are the appearance of processor-specific SRAMs and the resultant redrawing of borders in cache partitioning. A good example of this is the cache surrounding the i486.

Although more than 25 percent of the 1 million transistors of the i486 are used for cache, an imbalance exists between its small internal 8-kbyte cache and main memories that can easily exceed 4 Mbytes (greater than a 512:1 ratio). The results of this imbalance are a relatively high miss ratio and thrashing.

More than 80 percent of the i486 systems resolve these problems by using secondary caches as large as 256 kbytes—the most common size being 128 kbytes. This approach ensures a more balanced memory hierarchy. (Using the 128-kbyte and 4-Mbyte example, the ratios are S/P=16 and M/S=32.)

Zero-wait-state controllers

Applying the knowledge gained during the i386 generation, which predominantly used standard cache controllers, many designers are now developing their own discrete secondary cache controllers for the i486. These are typically implemented as a combination of ASICs, PALs, cache tags, fast CMOS logic and write buffers.

The most prevalent approach is direct mapping with write-through policy, because it's the easiest to implement and, at densities of 128 kbytes and larger, the difference in performance vs. set-associative controllers is small. Key benefits of these discrete controllers are shorter time-to-market and increased differentiation of the end product.

But this make-your-own strategy hasn't been easy to carry out; designers have had to struggle with the stringent timing of the i486 and the consequent difficulty in achieving zero-wait-state operation. This design challenge has been the main force behind the redrawing of borders in cache systems, whereby the overall controller and SRAM system is designed to achieve an optimal partitioning of functions, with many important functional blocks residing inside the SRAMs. Effective partitioning can make a surprisingly big difference in design complexity, performance, cost and time-to-market.

As we go beyond 50 MHz, the redrawing of borders will continue.

Many cache controller designers have eagerly embraced this approach and SRAM suppliers have obliged. The first example is the recently introduced IDT71589 32k×9-bit cache RAM from Integrated Device Technology (Santa Clara, CA).

Designed for i486 secondary caches, it includes a self-timed write, a burst counter and an i486 interface on-board. A 32-pin small-outline J-lead package offers the smallest possible footprint for a 128-kbyte secondary cache with parity. The various features of this device are the result of consulting with approximately 30 teams of designers that identified the main technical burdens they wanted to off-load to the SRAMs.

As a result, a 25-ns IDT71589 for a 33-MHz system offers benefits including zero-wait-state operation, a lower parts count, the replacement of 15-ns standard SRAMs, elimination of the narrow write pulse problem faced by designers and the ability to use lower-speed/lower-cost PALs and cache tags.

The definition process wasn't limited to OEMs; it also included most chip set manufacturers building standard cache controllers. This yielded excellent results, with several cache controllers being designed to take advantage of the IDT-71589, thus facilitating the design process faced by controller and board designers developing 33-MHz systems.

These standard controllers are often more sophisticated than the discrete ones, with features such as two- and four-way set associativity, write-back policy, Least Recently Used replacement, and so on. For board designers, the advantages of using such standard controllers are the simplicity of the design and the increased performance that's often achieved.

This kind of cooperation between designers of cache controllers (discrete or standard) and manufacturers of cache RAMs will proliferate as processor speeds increase. As we go beyond 50 MHz, the redrawing of borders will continue, with more functions being incorporated on the cache RAMs, especially when variables such as multiprocessing are factored in.

The challenges presented by the upcoming 50-MHz i486 are serious enough that some designers feel wait states will need to be added. This can certainly be avoided by use of the right cache RAM. Tight timing difficulties aren't unique to the i486 though—such difficulties are common to most RISC and CISC processors as they pass the 33-MHz barrier. Working closely with a supplier of processor-specific cache RAMs to carefully redraw the borders in the cache system could prove to be the only answer for cost-effective zero-wait-state solutions.

Manuel Alba, MSEE, corporate strategic marketing manager, Integrated Device Technology

only posted writes through an external write buffer. The chip offers bus snooping for coherency management, but with relatively simple algorithms.

Even the size and timing of Intel's part have been criticized by competitors. "The 385 SX cache controller comes in a 132-pin package, requires 35-ns SRAMs, and needs glue logic to reach 64 kbytes," says Mike Raghavan, president of start-up Nexel (Santa Clara, CA). "We're suggesting a design that fits in a 100-pin PQFP, uses 55-ns SRAMs, and can handle a 64-kbyte cache directly. In addition, a version of our part will form a plug-in upgrade path to the 386 SX for 286 motherboards."

In light of the competition surrounding the 82385, it's important to recognize that Intel's choices are a reflection of the company's view of the PC market, not inherent limitations in integrated controller chips. To see how far an integrated design can go, look at the 7C605 from Cypress Semiconductor (San Jose, CA): the device combines a cache controller, tag RAM and MMU for the Cy-

Available for Delivery
Now with VxWorks!

Top Real-Time Performance for VME

PERFORMANCE

Based on the powerful Intel i960CA, Heurikon's **HK80/V960E** Single Board Computer packs the best RISC or CISC price/performance on the market today. With 40,000 Dhrystones at 33MHz, the **HK80/V960E** is the *ideal* platform for critical real-time applications including embedded control, image processing and intelligent I/O.

The HK80/V960E sports 2 or 8 Mbytes of fast dual access, static column RAM. It also delivers a host of networking and I/O capabilities including on-card SCSI, Ethernet, Centronics compatible parallel I/O, and four RS-232C serial I/O ports. The VIC068 VME Interface Chip maximizes the VMEbus transfer rate and implements the full VMEbus interface with 4-level arbitration.

DEVELOPMENT TOOLS

Powerful real-time development tools are available for a variety of hosts with full support from Heurikon.

- **VxWorks** Real-Time Operating System allows the user to work in partnership with UNIX while running real-time targets. Code can be debugged using either a source-level debugger residing on the host systems or a symbolic debugger residing on the target. VxWorks, with driver support for SCSI and Mil Std 1553A, is available on Sun and Heurikon hosts with both backplane, and Ethernet target interface support.
- **GEN-960** for refining ROM or RAM code when setting up standalone embedded applications.
- **ASM-960** assembler allows macro-assembler users to fine tune application code as well as provide utilities for developing, debugging and maintaining application code.
- Intel and GNU compilers optimized for the 80960 family of embedded processors.

PRICE

Prices for the HK80/V960E series are very competitive with traditional CISC architectures...and it's available now! Call us now for more information on the HK80/V960E.

Heurikon Corporation
8000 Excelsior Drive
Madison, Wisconsin 53717

1•800•356•9602

Ext. 915
FAX: 608/831-4249
TEL: 608/831-0900

CIRCLE NO. 58

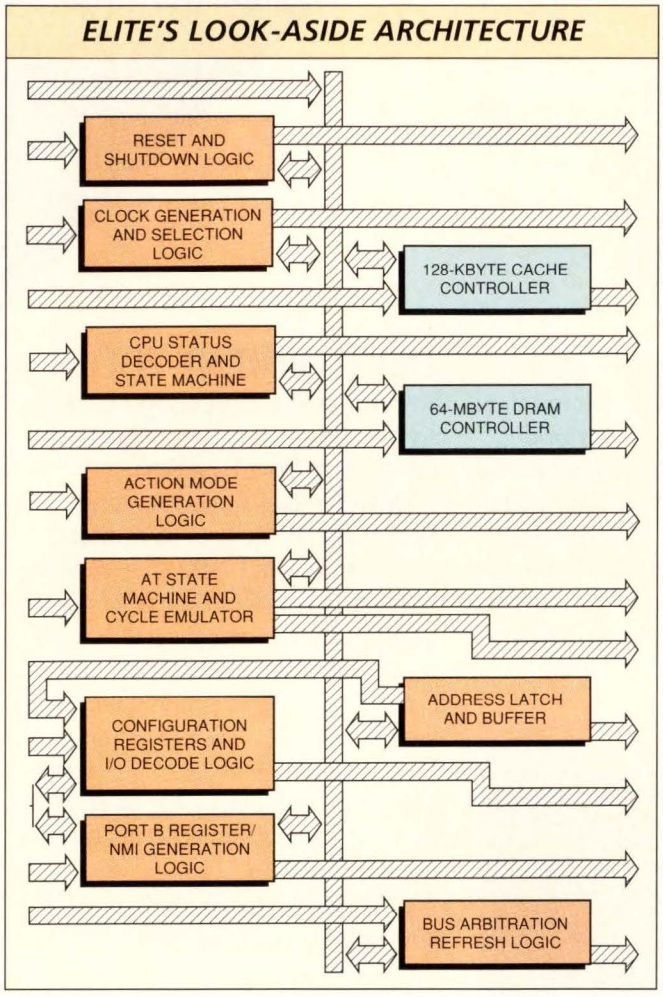

By combining a cache controller and a DRAM controller on a single die, Elite Microelectronics is able to improve the interaction of the two subsystems on its Eagle PC chip set. The look-aside architecture initiates a DRAM cycle while the cache controller checks the tag RAM, then aborts the DRAM timing if there's a cache hit.

press Sparc processor.

The 7C605 is intended to control a 64-kbyte direct-mapped virtual cache, as suggested by Sparc conventions. The chip provides either a write-through or a sophisticated copy-back strategy, and uses 32-byte read/write buffers to improve bus utilization.

Unique to the Sparc architecture is the 605's approach to multiprocessing coherency. The chip snoops the memory bus, and compares memory cycles not with the main tag RAM—which is virtually addressed, and therefore no help in recognizing physical addresses on the memory bus—but with a separate physical tag RAM maintained by the controller and the MMU circuitry. If the physical-address tag detects that a cached location has been addressed by some other CPU, the device can decide whether to source or invalidate the data in its data RAM.

This sort of sophistication is an impressive reminder that integrated controllers don't have to be simple. But it would be massive overkill on a PC. While Cypress has been stretching its integration muscles to achieve sophistication, other vendors have been working equally hard to achieve cost and space reductions.

Into the chip set

Chips and Technologies (San Jose, CA) has chosen to integrate the cache controller and tag RAMs into its high-end Peak chip sets, rather than into a stand-alone package. This gives, the company claims, significant cost benefits. "Our whole Peak DM chip set costs only a few dollars more than the 385 cache controller," says Nelson Chan, Chips marketing manager. He admits that the DM chip set, which uses a direct-mapped rather than a set-associative organization, doesn't include either tag or data RAMs, but he believes the system cost is still going to be substantially less than for a system using the 385.

"In the PC market, you have to look closely at the cost-performance ratio," he says. "You can do your own controller design, but how much more is it going to cost, and how much performance will you really gain? Often customers look at this question, and find they are talking about very expensive parts for a few percent gain in performance."

To date, Chips and Technologies has chosen to integrate the controller and tag RAMs (except in the DM product line) and leave the data RAMs to specialists. "We see data RAMs evolving CPU-specific features," Chan says. "Our role in that process right now is to keep talking to the SRAM vendors, and make sure our controllers take advantage of the logic in their parts. But as we get beyond 50 MHz, the off-chip delays may force us to reexamine that strategy."

While Chips has integrated primarily for cost reasons, other chip set vendors have had performance motives as well. S3 (Santa Clara, CA) has focused on PC families that start out with fast 386s and 486s and quickly move into multiprocessing. By the time the designers had incorporated a copy-back scheme and MESI multiprocessing coherency protocol, they found they were making choices that depended not only on the CPU but on how the memory controller worked. So the high-performance cache controller became chip-set specific. "At this level, if you don't design the cache controller in conjunction with the memory controller, you're almost sure to leave performance on the table," warns S3 marketing director Gary Baum.

Elite Microelectronics (San Jose, CA) sees similar advantages in having the cache controller and DRAM controller both available—in Elite's case, on the same chip. "By integrating the two, you can get the DRAM and cache controllers to work together. For instance, you can start them concurrently, and end up saving a cycle," says Peter Hsieh, Elite president and chief executive officer. In contrast to S3's multiprocessing approach, Elite's look-aside architecture is specialized for single-CPU systems.

At this level of integration, the needs of low-cost systems and high-end systems begin to diverge radically. Another vendor that sees this distinction is S-MOS Systems (San Jose, CA). "In a multiprocessing environment, the critical issue is bus bandwidth—you have to separate the CPU bus from the memory bus,"

What's the best kept secret in the VME industry?

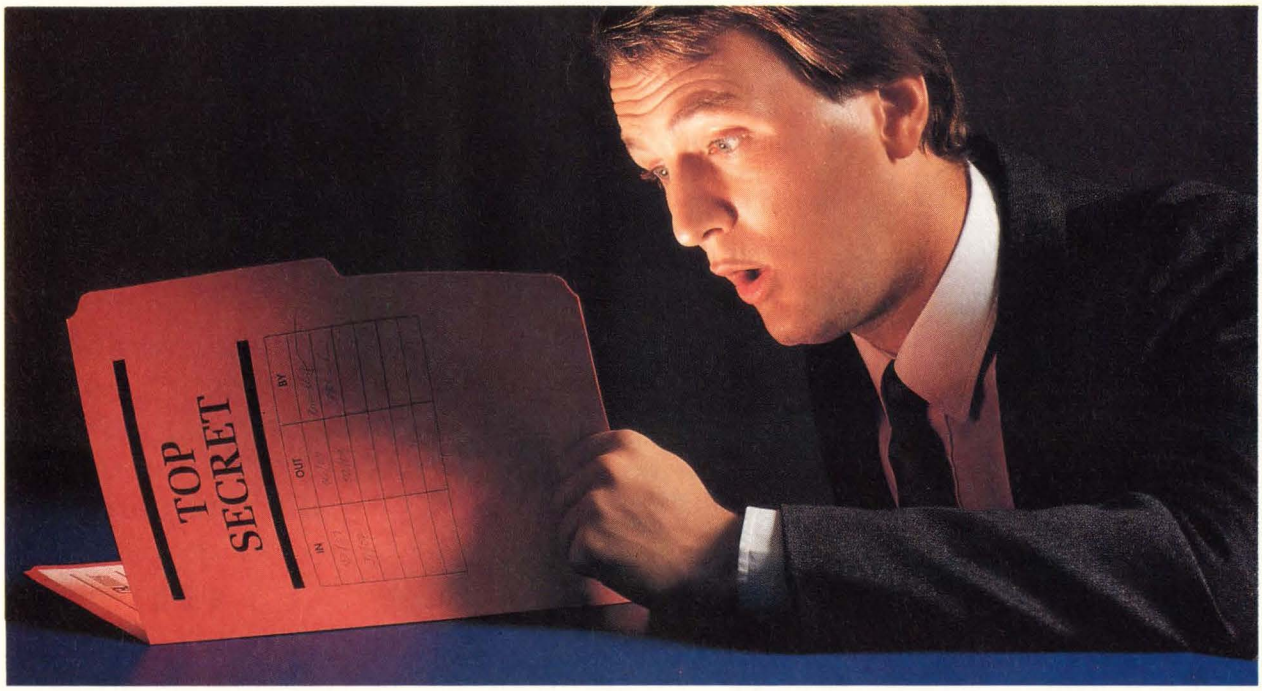

The Calmos VME chipset from Newbridge Microsystems.

The Calmos™ VME chipset has been in military and commercial production for over two years. With full compliance to industry standards, it provides high performance and flexibility for the future through a pin-for-pin compatible migration path to VME-64.

This unique VME chipset offers a versatile, complete solution. It can be used as a two chip set for a full VMEBus interface or as a single chip with external address and data bypass.

Complete compatibility with the IEEE 1014/VMEBus Rev C Interface ensures that all VMEBus signals are terminated, and all VMEBus Recommendations and Functional Modules, including master/slave interface, are implemented. The Calmos VME chipset also supports all cycle types and addressing modes.

The powerful Local Bus Interface provides compatibility with the industry standard 68020/030 protocol, and includes a 3 source Arbiter/Requester, Interrupt Controller (local and VMEBus), and Watchdog, Tick and Baud Rate Timers.

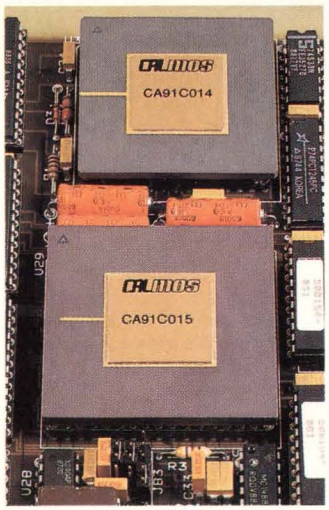

In addition, the Calmos VME chipset features BiMode Bus Isolation mode for system fault isolation, Auto ID and Auto SYSCON for jumper free design, plus 32 location monitors.

Part of the Newbridge Microsystems complete open bus product line extending all the way to Futurebus+, the Calmos VME chipset delivers the features you need today to build high-performance systems now and in the future!

NEWBRIDGE

NEWBRIDGE MICROSYSTEMS

Newbridge Microsystems • A Division of Newbridge Networks Corporation
603 March Road, Kanata, Ontario, Canada K2K1X3
Tel: 613-592-0714 • 1-800-267-7231 • Fax: 613-592-1320

CIRCLE NO. 59

CACHE CONTROLLERS

says director of marketing Joel Silverman. "But in single-processor systems, the right approach is to try to maintain continuity with Intel while offering a large enough data RAM." Continuity with Intel, in this case, means taking a hard look at Intel's 82395 controller with integrated data RAM.

Pulling in the data

"Our customers told us that they needed the performance they were getting from 385-based cache systems, but that in the next-generation system, cost would be of paramount importance," says Intel cache marketing manager Sheila Cardno. Developed in parallel with the 486 CPU's on-chip controller, the new part uses the four-way set associativity and buffered-write logic developed for the 486 to deliver strong performance from a 16-kbyte on-chip data RAM.

"The 395 represents a big challenge to the rest of the industry," admits Rajesh Vashist, senior product marketing engineer at Vitelic (San Jose, CA). "Intel is saying it wants the whole motherboard, including the SRAM business. And the company is going in the right direction; because of increasing clock speeds and power and cost considerations, this kind of integration is the way of the future."

But Vashist is confident: "You may see other combinations of controllers and cache RAMs on the market too, perhaps as SRAM specialists and logic specialists band together to produce an integrated product. Given the problems of high-speed SRAM design, like the need for high output drive, the need for self-timed synchronous parts, and so forth, you may get a better result starting out with SRAM expertise and adding logic."

Though the names of all the players aren't clear yet, the direction seems to be. Cache controller designs seem to be splitting into two branches: one to meet the cost needs of mass-market PCs, and the other to explore the challenges of multiprocessing and trans-50-MHz systems. Both branches will use high levels of integration, although the high-end designs may not integrate data RAM as quickly—the need for very large caches and synchronous, CPU-specific SRAMs may make such integration impractical.

"Somewhere above 40 MHz, it's going to be more difficult to continue using discrete parts."

—Nick Mati, Solbourne Computer

As to how far integration can go for the low-end systems, perhaps the answer lies in Intel's new 386 SL system-on-a-chip. Along with nearly all the logic of a conventional motherboard chip set, the 386 SX-derived chip includes a programmable, direct-mapped, two- or four-way set-associative cache controller and its tag RAMs. If Intel can sell the idea that small, well-organized caches give good performance, the next round of integration may put the cache data RAMs on-chip as well, giving a 386 SX CPU, core logic chip set and 395-like cache system on a single die. The product may lack the flexibility of today's building-block solutions, but its cost will be hard to beat. ■

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

Chips and Technologies		**S3**	
(408) 434-0600	Circle 239	(408) 986-8144	Circle 246
Cypress Semiconductor		**SGS-Thomson**	
(408) 943-2600	Circle 240	(602) 867-6100	Circle 247
Dell Computer		**S-MOS Systems**	
(512) 338-4400	Circle 241	(408) 922-0200	Circle 248
Elite Microelectronics		**Solbourne Computer**	
(408) 943-0500	Circle 242	(303) 772-3400	Circle 249
Integrated Device Technology		**Texas Instruments**	
(408) 727-6116	Circle 243	(800) 336-5236, ext. 700	Circle 250
Intel		**Vitelic Semiconductor**	
(800) 548-4725	Circle 244	(408) 433-6000	Circle 251
Nexel		**VLSI Technology**	
(408) 241-7585	Circle 245	(408) 434-3000	Circle 252

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. 61

Are designers' expectations for synthesis realistic?

Barbara Tuck
Senior Editor

Though still not perfected, synthesis tools are well on the way to changing design methodology and bringing technology independent design to life.

While the implication of behavioral synthesis is push-button silicon, synthesis is still very much in a growth stage. The dimension of optimization is changing as more and more synthesis is being done at a VHDL level. And because it permits users to describe a design at an abstract level and then target it to a specific technology, synthesis has been the catalyst for technology-independent, top-down design. Synthesis software is at the core of systems that are integrating programmable logic, ASIC and board design.

As it matures, synthesis technology is paving the way for a consistent design methodology that will let users explore technology alternatives, analyze trade-offs and then, from a single workstation seat, implement design descriptions in PLDs, FPGAs, gate arrays or even standard cells. Synthesis underlies the ability to prototype a masked gate array in an FPGA and to migrate from an FPGA to a gate array.

In five to 10 years, synthesis technology is expected to mature to the point where a design described at the most abstract level will be automatically partitioned into multiple chips, with the designer making no decisions beyond the architectural level. A portion of the design will perhaps be implemented in microcode, another portion in a microprocessor, another in a few ASICs and the remaining logic in a handful of PLDs and FPGAs.

Though synthesis is changing the entire methodology of design, an intimate knowledge of the technology is still confined to a relatively small group of people. To the rest of us, it's a mysterious, power-packed technology for which we have great expectations.

The question is, are those expectations realistic? Is synthesis, or can it be, a push-button technology? Dave Gregory, manager of optimization software at Synopsys (Mountain View, CA), thinks not.

"Synthesis tools let designers explore alternatives and experiment with decisions. Synthesis isn't a technology that will replace the designer," he says, "but it will leverage the designer. As time goes on, the amount of human intervention will lessen. Meanwhile, the tools let designers spend more time doing creative rather than mundane tasks.

"A synthesis tool has to do a task at least as well as a human," Gregory adds, "or the synthesis vendor loses credibility." As the leading synthesis vendor, Synopsys is extremely sensitive about raising false expectations for its tools, according to Gregory.

Since synthesis is somewhat esoteric, it's important to separate what it can do today from what it might do, or is promised to do, tomorrow. Among the questions that need to be answered are the following:

- At what level of abstraction are designs being synthesized?
- Is resource allocation a reality?
- How much designer interaction is required?
- What does a designer do about nonsynthesizable VHDL constructs?
- Are complex gate array designs automatically partitioned for optimization?
- Do silicon-specific tools have an edge on establishing a link to layout?
- What's the optimum design path for PLD and FPGA design?
- Are all the pieces in place for high-level FPGA synthesis?
- How feasible is it to go from an FPGA netlist to a gate array netlist?

Although all synthesis tools translate behavioral representations into structural representations, individual tools can differ sig-

nificantly from one another. When users are evaluating synthesis tools, it's important to question the level of abstraction for design description and synthesis and the range of synthesis tasks performed. Though abstraction levels actually form a continuum, they can be classified into discrete levels for the sake of discussion. The three levels most often referenced are behavioral, register-transfer level (RTL) and logic.

Levels of abstraction

A behavioral description is at the most abstract level. It addresses the functionality of a design without reference to its underlying structure, or how it might perform its behavior or the clock cycles required for execution. Users have to go beyond the constructs contained in a model or supported by a tool to judge the level of abstraction. The use of behavioral constructs such as computer-aided software engineering, IF-THEN, and FOR loops doesn't necessarily imply that the model in which they're contained is a behavioral model. Nor can the level of a VHDL synthesis tool be gauged by the constructs it supports. VHDL supports descriptions at most levels of abstraction, including structural, RTL

THE SYNTHESIS CONTINUUM

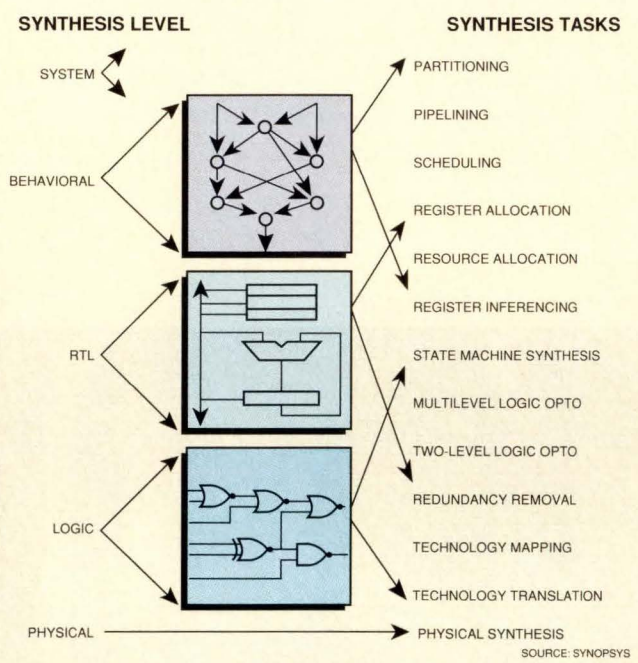

Though levels of abstraction for design description and synthesis actually form a continuum, it's often important to classify descriptions into discrete levels. The three primary levels of abstraction generally used for the sake of discussion are behavioral, register-transfer level (RTL) and logic. Each level of synthesis has specific associated tasks, some of which overlap, which can cause confusion.

and behavioral.

An RTL or dataflow description is at the next level of abstraction below a behavioral description. Most often characterized as a system definition in terms of registers, switches and operations, an RTL description does have a notion of an architecture and does incorporate a clocking scheme.

At the logic level of abstraction—the lowest, nonphysical level at which a design can be represented—descriptions retain the architecture of the RTL level and also show the Boolean architecture or the logical implementation of the function. Descriptions at this level are typically technology-dependent.

Some people in the industry see behavioral synthesis as a supercharged term. "Levels of abstraction can be used as marketing tools," says Aart de Geus, senior vice-president of engineering at Synopsys. "The synthesis world should be seen as a continuum, but people make artificial boundaries. Synopsys is now well anchored in RTL—ahead of us is the behavioral domain."

To be in the behavioral domain, says de Geus, optimizations need to be performed between clock cycles rather than within clock cycles as with RTL synthesis.

Responding to those who assume that synthesis done at a behavioral level has to be better than that done at an RTL level, de Geus says, "Credibility is our number one asset. Having capabilities at a high level without having capabilities at a low level will affect the quality of synthesis results."

Another innovator in synthesis technology, Silc Technologies, which was recently assimilated into Racal-Redac as the ASIC Design Division (Westford, MA), is adamant about categorizing the capabilities of its Silcsyn tool as behavioral synthesis. "We synthesize designs at a micro-architectural level, and as far as I know, we're the only ones that do," says vice-president of engineering Jeff Fox. A microarchitectural level is difficult to define, beyond that it seems to encompass some behavioral tasks and some RTL tasks. Silc's claim of synthesizing at the behavioral level is difficult to dispute since Silcsyn II with VHDL just began shipping last month. Users will now have the opportunity to evaluate the results of microarchitectural synthesis firsthand.

The first version of Silcsyn, based on the proprietary SDDL high-level language, was introduced before top-down design and high-level languages were popular. The shipping schedule for Silcsyn II slipped when Racal-Redac decided to do a major redraft of its VHDL. "The redrafted VHDL is somewhat less sophisticated and easier to understand," Fox maintains. "We've used the identical synthesis as with our own SDDL language and just built a VHDL front end. We have the same control as we do with SDDL; the only difference is the input language."

Through acquisition, Mentor Graphics (Beaverton, OR) obtained

HOW SYNOPSYS REDUCES RESOURCES

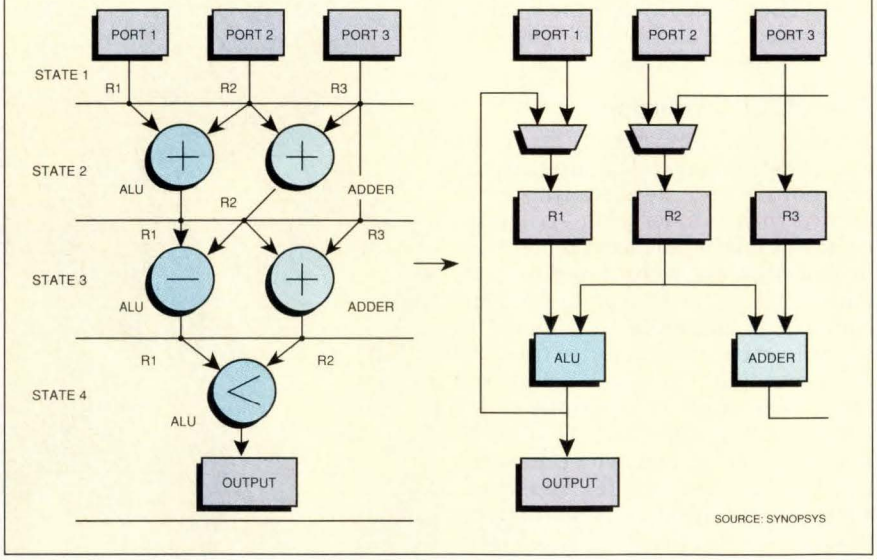

The use of resource allocation and sharing minimizes the resources required to two—an ALU and an adder— even though five arithmetic operations are performed. The ALU performs one of the "+" operations in state 2, and the "–" and "<" operations in states 3 and 4; the adder does the "+" operations in states 2 and 3.

68040 VME 33 MHz 0-Wait-State

Your Vision of High Performance at an Affordable Price is Now Real!

With the OB68K/VME40™ you no longer have to compromise on performance or price in your VME embedded control application. We start by giving you a very basic board which includes:

- 25-33MHz 68040.
- (8) 28-pin RAM sockets for up to 256KB of dual access 0-wait-state static RAM (32KB standard).
- (8) 32-pin sockets for up to 8MB of ROM.
- (2) asynch RS232C serial ports.
- (16) lines of parallel I/O.

You can configure it with just the right amount of RAM and ROM you need. And you do not have to sacrifice features. Our Omnimodule™ modular I/O connector allows you to implement a wide variety of serial, parallel, SCSI, GPIB, analog, digital and other I/O options - all fitting into one slot. Other features include:

- VTC's VIC068 VME interface chip with arbiter, interrupter, mailbox and more.
- Terminal monitor/debugger/diagnostic firmware program included.
- 2 year limited warranty.
- Worldwide availability.

All of this gives you a high performance board at a price you can afford with the features you need.

To learn more about our OB68K/VME40 contact our Marketing Manager, Pete Czuchra at 1-800-638-5022 or (708) 231-6880 in Illinois.

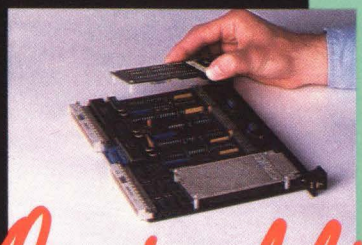

Omnimodules™

OMNIBYTE™

OMNIBYTE CORPORATION
245 W. Roosevelt Road
West Chicago, IL 60185-3790
Fax No. 708-231-7042
CALL TOLL FREE 1-800-638-5022
In IL 708-231-6880

CIRCLE NO. 120

Our VME and Multibus Product Lines Stretch for Over 124 Miles

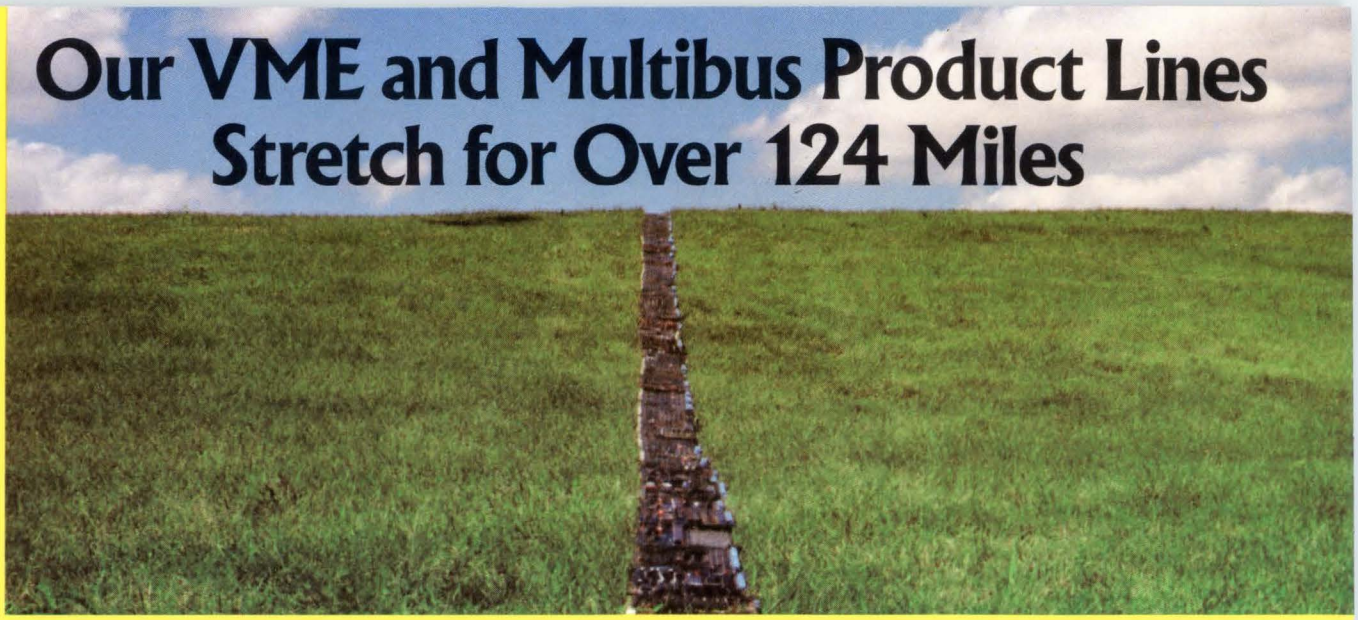

That's 854,738 uniquely configured boards to choose from and all from Omnibyte. You can choose from different processor types, RAM sizes, I/O options and other features to put together a board that gives you the features you need. With Omnibyte's quality, selection and 2 year limited warranty, you can count on finding exactly what your looking for.

Here are just a few of the boards we offer:

OB68K/VME20™ VME SINGLE BOARD COMPUTER

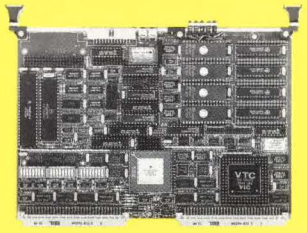

- 68020 16.66 - 33 MHz CPU
- (8) 28-pin RAM sockets for up to 265KB of dual-access zero-wait-state-static RAM
- (8) 32-pin sockets for up to 8MB of ROM, (4) sockets may be EEPROM
- (2) RS232C asynch serial ports
- (16) lines of parallel I/O
- (1) (OMNIMODULE socket for a wide variety of I/O (i.e. 2 serial ports, 20 parallel lines)
- VIC068 VME Interface Controller

OB68K/VSBC20™ VME SINGLE BOARD COMPUTER

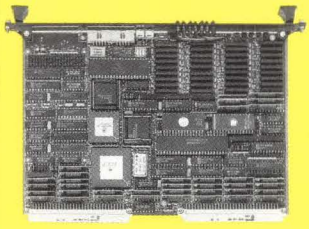

- 68020 16-33MHz, CPU
- 1-4 MB of dual-access, zero-wait-state DRAM with parity
- 68882 (optional)
- (2) 32-pin ROM sockets
- (2) RS232C serial ports
- (2) 8-bit parallel ports
- (1) OMNIMODULE socket for a wide variety of I/O (i.e. 2 serial ports, 20 parallel lines)
- 4 level bus arbiter (optional)

OB68K/VSBCI™ VME SINGLE BOARD COMPUTER

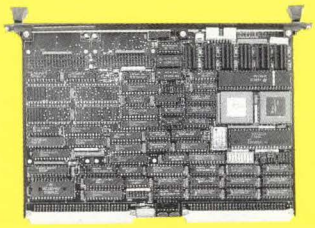

- 68000 12.5MHz 16/32 bit CPU
- 512KB of dual-access, zero-wait-state DRAM with parity
- (4) 28-pin ROM sockets
- (3) 16-bit counter/timers
- (2) Omnimodule™ I/O sockets for a wide variety of I/O (i.e. 4 serial ports, 40 parallel lines)
- DMA controller (optional)
- VME bus interrupt generator (optional)
- Optional 4 level bus arbiter
- Two year limited warranty

OB68K/VME1™ VME SINGLE BOARD COMPUTER

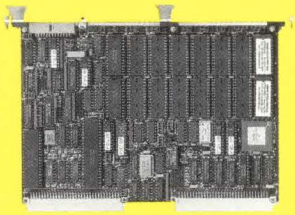

- 12.5 MHz 68000 CPU
- (8) pairs of 28-pin sockets for RAM or ROM
- (2) RS-232C serial ports
- (2) 8-bit parallel I/O ports
- System Controller

OB68K/VIO™ VME UNIVERSAL I/O BOARD

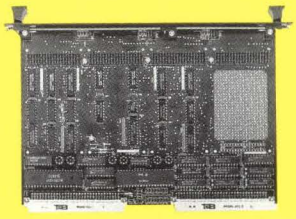

- (4) Omnimodule I/O sockets for a wide variety of I/O (i.e.8 serial ports, 80 parallel lines)
- One (1) interrupt per Omnimodule, two (2) optional

OB68K/MSBC30™ MULTIBUS I SINGLE BOARD COMPUTER

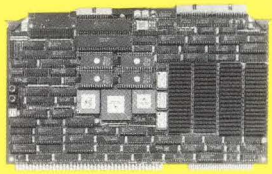

- 25-33 MHz 68030 CPU
- 4-32 MB dual access, zero-wait-state DRAM w/parity
- 68882 Math Co-Processor (optional)
- 2 channel DMA controller (optional)
- (2) RS232c sync/async serial ports
- (2) 8-bit parallel ports
- (1) OMNIMODULE™ socket
- (4) 32-pin ROM sockets

All our different configurations are built to give you the best in quality. And they are backed by our famous 2 year limited warranty. For more information call Pete Czuchra today. He'll help you pick the card you need.

OMNIBYTE™

A Look at Today...a Vision of Tomorrow

CIRCLE NO. 119

OMNIBYTE CORPORATION
245 W. Roosevelt Road
West Chicago, IL 60185-3790
Fax No. 708-231-7042
CALL TOLL FREE 1-800-638-5022
In IL 708-231-6880

Tying synthesis to layout: the missing link

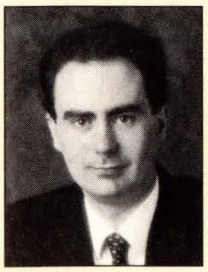

Logic synthesis tools have provided design engineers with major productivity enhancements during the last two years. Such tools take an existing netlist and optimize the area or performance of that netlist—and improvements of 30 percent aren't uncommon. With the added ability to synthesize from industry-standard HDLs, we are seeing the evolution of a design methodology that will allow the development of ASIC chips containing millions of gates in the latter half of the decade.

Today's synthesis tools are only the first generation of what will be a dynamic technology. The ability to accept higher levels of abstraction in a design description is frequently discussed, from register-transfer level to architectural to behavioral. For the designer, productivity and time-to-market will improve as the level of abstraction increases.

But advances in abstraction must be matched by advances in the physical implementation of a device. The goal of a synthesis system should be to produce a device that has the required functionality and performance with the least amount of silicon and power dissipation. And the system should achieve this goal by using the highest level of abstraction possible and the most advanced silicon technology.

The initial results from first-generation synthesis tools look good. They reduce the total gate count and critical path delays. Unfortunately, the advantage of these results can be lost when the post-route interconnect delays are plugged back into the simulation.

For example, a critical path that was predicted to be 10 ns before physical design can become 12 ns or worse when the extracted resistance-capacitance values are considered, making the final design of a critical circuit very time-consuming. The problem is that most synthesis tools exist in a CAE vacuum. Only performance and area are taken into account prior to placement and routing; physical implementation isn't a factor.

Today's synthesis tools use gate count to measure the area of a circuit, when in reality the actual area is made up of both cell macros and interconnect wires. This problem is compounded when the goal is performance optimization, because the circuit delay equals the macro base delay plus interconnect delay. Designers can use a predictive capacitance for nets, but this is empirical data based on a sample of previous designs. The interconnect delay becomes a more dominant delay as we move to submicron technologies.

In a typical scenario, the design engineer has verified that the latest design is functionally correct and will probably meet the timing specifications. The netlist is then handed over to the ASIC vendor for placement and routing.

A synthesized netlist can be difficult to place and route, though, causing the actual gate utilization to fall well short of the vendor's expected utilization. In the case of a gate array, this can cause the design to require a larger base array, or the critical path delay to exceed design specification. The productivity that had been gained in the design-capture stage can be lost in additional iterations of the design.

Productivity and time-to-market will improve as the level of abstraction increases.

There are two ways to get around this problem: to improve the synthesis tools so that they do an adequate job on the first generation of the system, or to tighten the link between front-end synthesis and back-end physical implementation to allow automation of the system's reiterative design.

Improving synthesis tools

One of the limitations of most synthesis tools is caused by the functional blocks with which they operate. Most systems can use only standard gates, either as gate array or standard-cell macros within a vendor's library. This isn't the most efficient way to implement RAM or complex blocks such as multipliers or datapaths.

One way that a synthesizer can improve on the implementation of a design (including synthesizing the system more accurately) is to access and use silicon compilers within a common software toolset. In this way, synthesizers can move from synthesizing gates to synthesizing ASICs.

A silicon compiler offers an efficient physical implementation of a single function. When using a silicon compiler, the designer knows exactly what structure is required, and can give the compiler the parameters for that structure.

VLSI Technology's ASIC Synthesizer is designed to use block and datapath compilers so that compiler output can be used in the synthesis of chips all the way down to layout. The parameters of these compilers are closely tied to the parameters of VLSI silicon, and these compilers produce layout from high-level blocks such as datapaths and memories, so the performance of the resultant system is much higher that what could be obtained by synthesizing the same system entirely from gate-level structures.

Tightening the link

The other method of enhancing the use of synthesis for critical designs is to generate the ASIC chip directly from the HDL description, to extract the timing information from the layout (for accuracy, this necessitates a resistance-capacitance extraction tool), and to back annotate the result of the extraction to the synthesis tools.

Consequently, updating the HDL description will ensure that the required system is generated. As process technology shrinks to the 1-μm level and below, the increased dependence of total circuit performance on parasitic elements delays requires a strong link between the tools performing the synthesis of the system and the tools creating the silicon layout from the synthesis netlist.

At VLSI Technology, ASIC synthesis involves the interaction of several highly integrated software tools encompassing compilation, verification (including timing verification), simulation, circuit extraction, floor planning and layout.

The attainment of a close synthesis-layout tie will be a combination of factors involving enhancement of synthesis tools and a closer tie to the layout. Only when these are accomplished will it be possible to develop software systems that can ensure that the extracted, synthesized and HDL-described circuit will demonstrate the performance required of submicron silicon circuits.

Michael O'Brien, BSEE, product line manager, Logic Design Automation, VLSI Technology

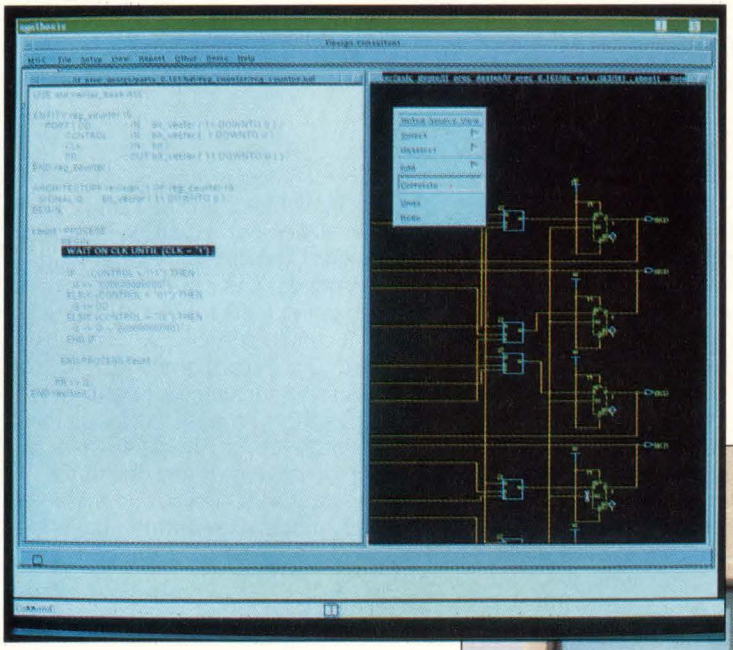

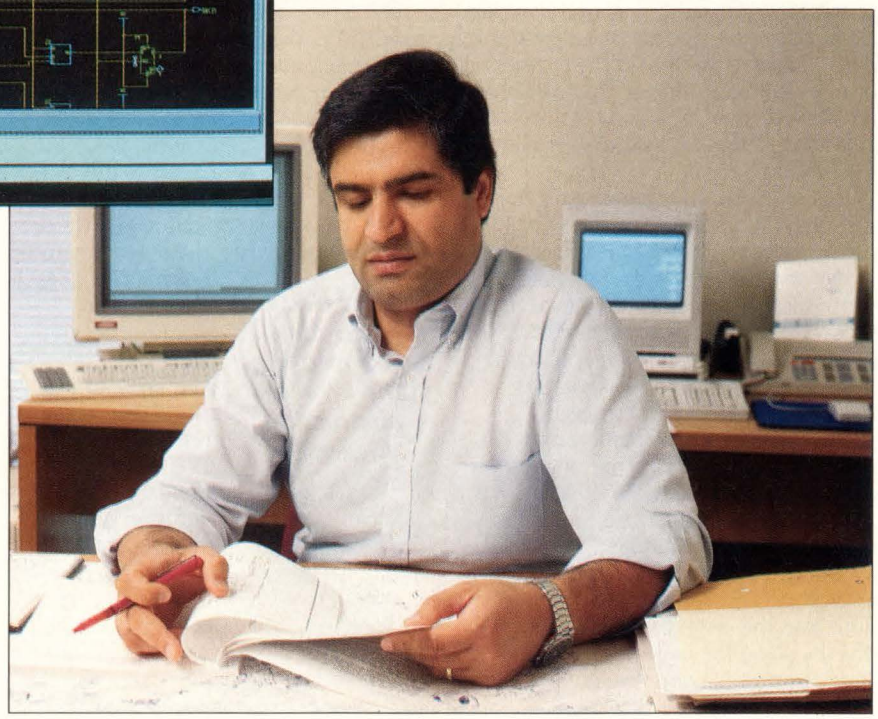

Rostam Joobbani, development engineer director at Mentor Graphics' newly created Design Synthesis Division, looks over a design he's synthesizing with Mentor's Design Consultant VHDL synthesis tool. With the Design Consultant, VHDL input can be graphically correlated with synthesized schematics (inset).

the VHDL synthesis technology of Trimeter and the ASIC synthesis technology of Silicon Compiler Systems. Mentor is integrating synthesis tools from both companies and a PLD synthesis tool from Minc into its top-down design environment Release 8.0. Mentor expects its integrated VHDL, ASIC, and PLD synthesis solution to be available the first quarter of next year.

Like Synopsys, Mentor makes no claims about being at the behavioral level. According to Rostam Joobbani, development engineer director at Mentor's newly created Design Synthesis Division, "We started at the gate level with logic synthesis, we're at the RTL level now, and we're moving toward the behavioral level. With VHDL, implementation details will disappear as we move to higher and higher levels until we reach a point where it won't really matter to the user in which technology his design is implemented. It will simply be an economic choice."

Minimizing resources

Each level of synthesis has specific associated tasks, some of which overlap levels of abstraction. Whether a tool can allocate and share resources and how it performs those tasks, which are associated with both RTL and behavioral levels, will be important for users to evaluate. Resource allocation can have a significant effect on the speed and area of the design being synthesized. When allocating resources, a tool selects components such as adders and ALUs to implement the operations of a behavioral description. A related task is resource assignment—or the selection of specific operations to be implemented on specific resources. When a synthesis tool shares resources, it attempts to rely on a single resource for more than one operation, as long as the operations are executed at different times.

Racal-Redac's Silcsyn was the first synthesis software to include resource allocation and sharing. Though Version 2.0 of the Synopsys synthesis tool incorporates resource allocation and sharing, Racal-Redac's Fox claims that Synopsys is doing it "in a limited sense—only in those situations where it's obvious from the behavioral description that resources can be shared, as in an IF-THEN statement. The Synopsys software doesn't permit users to do loops with noncomputable bounds, such as 'While enable is high, do such and such'."

Countering Fox's claim is Jeff Lewis, Synopsys marketing manager of HDL products. "I see it as two different issues," he says. "First, Racal-Redac's statement on noncomputable bounds involves the issue of how a designer writes VHDL. Second, the comment on 'while' statements involves the issue of a different coding style. Despite Racal-Redac's claim, Synopsys does share resources across different statements. But resource sharing can't be done indiscriminately—designer interaction is imperative."

The silicon-specific ASIC Synthesizer from VLSI Technology (San Jose, CA), of which a VHDL version is in beta site now, also synthesizes

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.

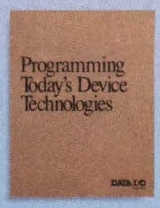

Call today for a FREE tutorial.

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

DATA I/O Corporation

CIRCLE NO. 62

A chip carrier socket that won't play "pop goes the circuit."

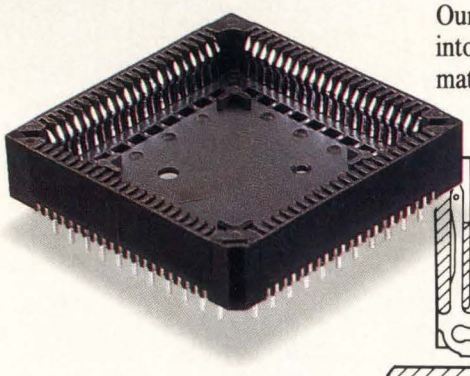

Our sockets are designed to get solidly into contact and stay in contact. No matter what the outside influences.

Pop-out is simply not a problem.

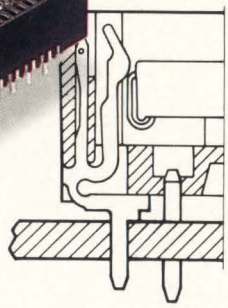

Controlled contact interface angle in AMP HPT sockets ensures positive chip carrier retention. Our exclusive removable housing allows direct inspection of solder joints, and fast repair/replacement of contacts.

The contacts are High Pressure Tin, an AMP proprietary design that creates very high normal forces — a minimum of 200 grams per contact — for maximum retention and reliable interconnection. Short-signal-path contacts float in the housing to accommodate thermal expansion.

Two basic styles of sockets are available: square or 32-position rectangular. Both come in solder tail or surface

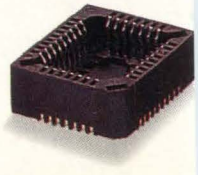

mount versions and feature all the important details. Tin-over-nickel plating is applied after the contacts are formed, to assure full plating. We've built in visual indicators for locating pin 1, and polarizing to aid correct insertion.

Orientation holes in the 94V-0 housing floor make registration to the pc board both fast and simple, ideal for hand or tube-loaded robotic insertion. And the high pin counts make very effective use of real estate.

Call the AMP Information Center at 1-800-522-6752 for literature on HPT PLCC Sockets. AMP Incorporated, Harrisburg, PA 17105-3608.

CIRCLE NO. 63

THIS IS AMP TODAY. **AMP**

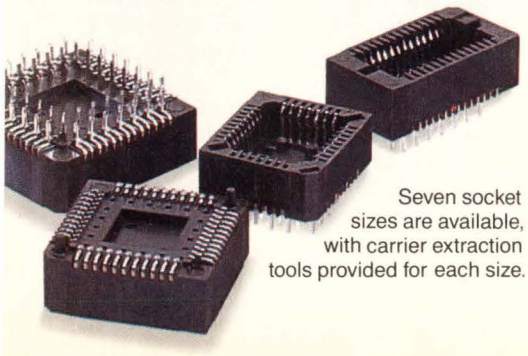

Seven socket sizes are available, with carrier extraction tools provided for each size.

at the RTL level and offers resource allocation. "Whereas most general-purpose tools struggle with the arithmetic involved in resource allocation, VLSI has the experience of the Datapath Compiler behind it," says Michael O'Brien, Logic Design Automation product line manager.

O'Brien says that VLSI has to go to a more abstract level of description before it can offer scheduling, a behavioral-level task related to resource allocation/sharing but at a higher level. Scheduling assigns operations in a behavioral description to a sequence of control states. "As we move to higher levels, we'll cease to define clocks," says O'Brien. "We won't care where or when something takes place. The software will move things around and optimize."

Blockbusting needed

Though technologists are working to make synthesis more automated and less restricted, O'Brien says, they're not there yet. Optimizing a complex gate array, for instance, requires a lot of human intervention. Designers have to feed complex designs to most tools in 3,000-gate chunks. Any more than that will take an unreasonably long time.

"When designers are writing a behavioral description of a design," O'Brien says, "they have to partition the description into modules through the use of syntax. The tools aren't in place to recognize 3,000-gate chunks." O'Brien attributes this restriction partially to the algorithms themselves, adding that a lot of algorithmic work is going on in the universities. "Also, the software tools croak," he says, "because of the size of the databases." Synthesis takes a lot of horsepower. Going to more-powerful hardware platforms will help, but it won't solve the problem. "To get to a million gates implies the need for good floor-planning tools," says O'Brien.

Another shortcoming of existing synthesis tools is the absence of a link to physical design. Though tools such as the VLSI ASIC Synthesizer and the LSI Logic (Milpitas, CA) Silicon 1076 VHDL synthesis tool, both closely tied to silicon, don't offer the freedom of technology-independent design and close the door to multiple sourcing, they do have an edge when it comes to establishing a link with layout.

"Right now, layout isn't integrated into the ASIC Synthesizer, but it will be in the future," O'Brien says. "Re-optimization after layout with different criteria isn't automatic," he admits. "But because the tool lets you maintain hierarchy, you can trace a problem that you've identified through simulation and timing verification to the block where it originates, modify the code and then reoptimize." O'Brien predicts that the reiteration of such a block will be automatic in the future.

Nonsynthesizable constructs

Bob Dahlberg, LSI Logic product marketing manager for system design, raises the issue that, for users synthesizing designs in VHDL, not all VHDL constructs are synthesizable. "At the board level, you can use full VHDL because you're not going to synthesize it. But if you're going to silicon through synthesis, you have to make sure that you haven't included constructs that are nonsynthesizable." LSI surrounded the Synopsys synthesis tool with proprietary modules to create Silicon 1076, currently in beta site and expected to ship the first quarter of next year.

No tool today can synthesize all of VHDL. There's no reason, in fact, to synthesize those parts of the language used only for simulation and logic modeling. But what happens if nonsynthesizable constructs were included in a design description?

According to Steve Wolfe, production manager in Mentor's Design and Analysis Division, "With our VHDL synthesis tool, the Design Consultant, users can set a switch in the compiler to flag nonsynthesizable constructs and syntax errors. Then they can go back into the VHDL description and make changes. This is all done early in the design cycle since the user compiles his VHDL as soon as he writes it." Mentor claims that the compiler switch ensures a synthesizable design. The Mentor synthesis solution may be the only one that features a common VHDL compiler for synthesis and simulation.

The Design Consultant solves another VHDL-related problem that is likely to have kept some designers from embracing VHDL. Up until now, once a VHDL design descrip-

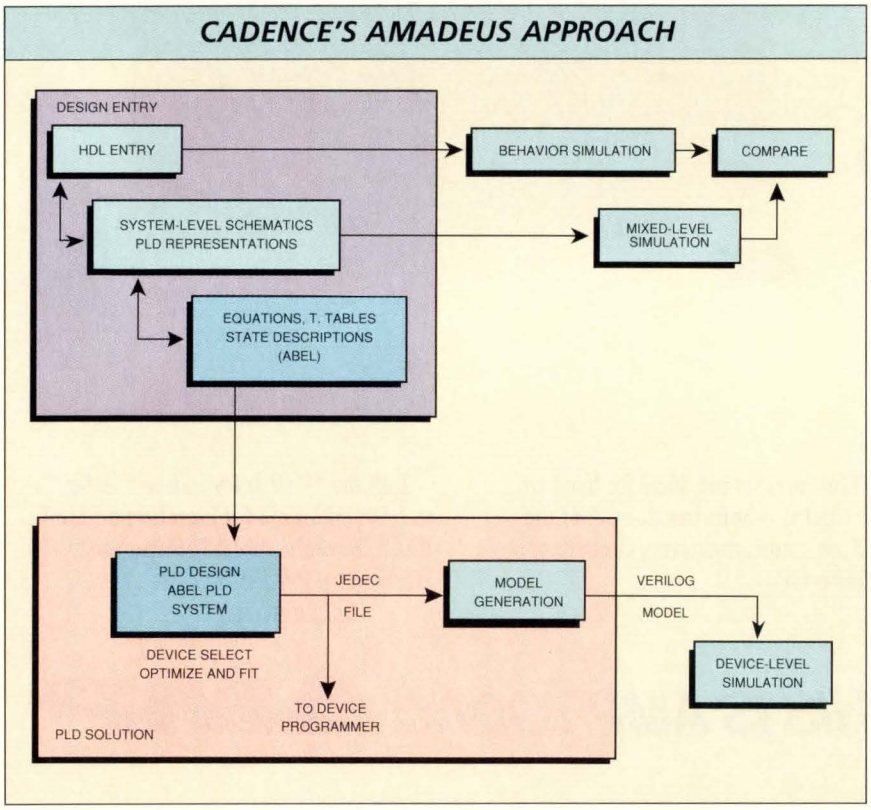

Cadence Design Systems' PLD strategy is to provide the best vendor-independent solution with extensive device support. Cadence has integrated Data I/O's Abel, the de facto industry-standard PLD language and synthesizer, into the Amadeus Systems Design Series. Highlights of the Amadeus solution include top-down/bottom-up PLD design, push-button Verilog modeling and a common user interface with board design software.

A more sensible approach to VHDL synthesis

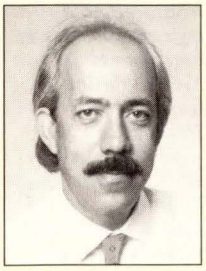

While most electronic designers agree that top-down design using VHDL and logic synthesis tools will provide significant productivity gains, many are concerned about how difficult it will be to adopt—and justifiably so. Design teams that have been using some of the early VHDL synthesis tools have encountered several problems.

These problems include the inability to correlate the synthesized schematics with the source VHDL description; the inability to determine up-front which VHDL constructs are synthesizable; the inability to target programmable logic device and field-programmable gate array (FPGA) technologies; simulation and synthesis libraries that are different; the inability to incorporate existing design data and leverage existing design expertise; and a relatively steep learning curve.

While these problems aren't insurmountable, they do require companies to consider a more sensible approach to VHDL synthesis—one that's integral to the design process, that's more tightly integrated into the EDA environment, and that's more flexible with respect to the designer's needs and expertise.

Technology-independent design

Technology-independent design has long been sought by electronic designers. The ability to defer technology decisions until later in the design process allows designers to respond more quickly to changing design requirements, and to make better trade-off decisions before they select an implementation.

Early VHDL synthesis tools were limited to semicustom technologies only, such as standard cell and gate array. But the 1990s call for a broader approach. What's needed are VHDL synthesis tools (such as Mentor Graphics' Design Consultant) that allow designers to target designs to multiple technologies—ASIC, custom IC, PLD and FPGA.

Although stand-alone synthesis tools can make engineers more productive in discrete areas, that productivity is lost if the tools don't work well with the rest of the EDA environment. What's needed is a VHDL synthesis solution that works in concert with other design and analysis tools, and is tightly integrated into the EDA environment. The level of integration necessary goes well beyond the netlist and user interface level. In fact, there are three levels of integration that are necessary.

At the first level of integration, a common user interface ensures operational consistency across all tools. A common database lets design data flow smoothly from one tool to another, and allows designers to re-use existing design investments. In addition, it enables capabilities such as graphical highlighting of VHDL code when errors occur and cross-correlation of VHDL source code with the synthesized schematic. More important, it supports true concurrent design.

What's needed are flexible VHDL synthesis tools that are tightly integrated into the EDA environment.

At the second level, all tools understand the data the same way, which ensures consistent results across all design creation and analysis tools. A VHDL compiler that's common to both simulation and synthesis assures that the simulator and synthesis tools don't disagree. More important, it tells the designer, before simulation, when nonsynthesizable VHDL constructs are used. This eliminates a typical problem related to existing VHDL synthesis tools: designers don't find out that their designs can't be synthesized until their functional simulations are complete.

At the third level, all tools view the same data—that is, they share a common library. Today's VHDL synthesis tools use ASIC libraries that are separate from the libraries used by the simulator and the other analysis tools, and sometimes the data between libraries don't match. Or the library itself is missing from either the simulator or the synthesis tool.

An effective VHDL synthesis environment is one that uses a single library scheme. This ensures data consistency and minimizes the work the ASIC vendor has to provide in order to support the various design tools. This, in turn, expedites the availability of new ASIC libraries as new semicustom technologies become available.

Mentor's Design Consultant VHDL synthesis tool, in conjunction with the new Concurrent Design Environment, supports all three levels of integration.

An incremental approach

One of the most universal concerns regarding high-level synthesis is the sharp learning curve for design languages such as VHDL. An evolutionary or incremental approach lets designers leverage the productivity benefits of high-level synthesis using familiar design techniques, such as graphics-based design and editing, state-machine descriptions and high-level macrofunctions. In other words, designers gain the benefit of top-down design methods before taking the full leap to VHDL synthesis.

Synthesis tools are available today that substantially increase productivity and allow designers to step up from gate-level design without the steep learning curve. Mentor Graphics' Autologic synthesis tool, for example, synthesizes a variety of high-level representations, such as state-machine descriptions and graphical, parameterizable macrofunctions.

These high-level design methods can be used in any combination to quickly and easily describe and synthesize high-level functions, short of using VHDL. Low-level descriptions—such as schematics, netlists, Boolean equations, and PLA tables—are also supported.

As the designer's experience and comfort level grows, the design environment should allow the designer to move smoothly to VHDL-based design. Mentor Graphics developed a completely top-down design environment that lets designers choose their level of involvement with VHDL and synthesis, depending on their level of expertise and comfort level.

VHDL-based design and synthesis promises to raise designers to new levels of productivity. But this promise could be undermined if companies adopt VHDL synthesis tools without considering the design process, the EDA environment, the designer's expertise, and the relationship of synthesis with other tools. A more sensible approach to top-down design allows designers to move to VHDL synthesis incrementally, without jeopardizing time-to-market.

Robert Mendes da Costa, BSEE/CS, director of marketing, Design Synthesis Division, Mentor Graphics

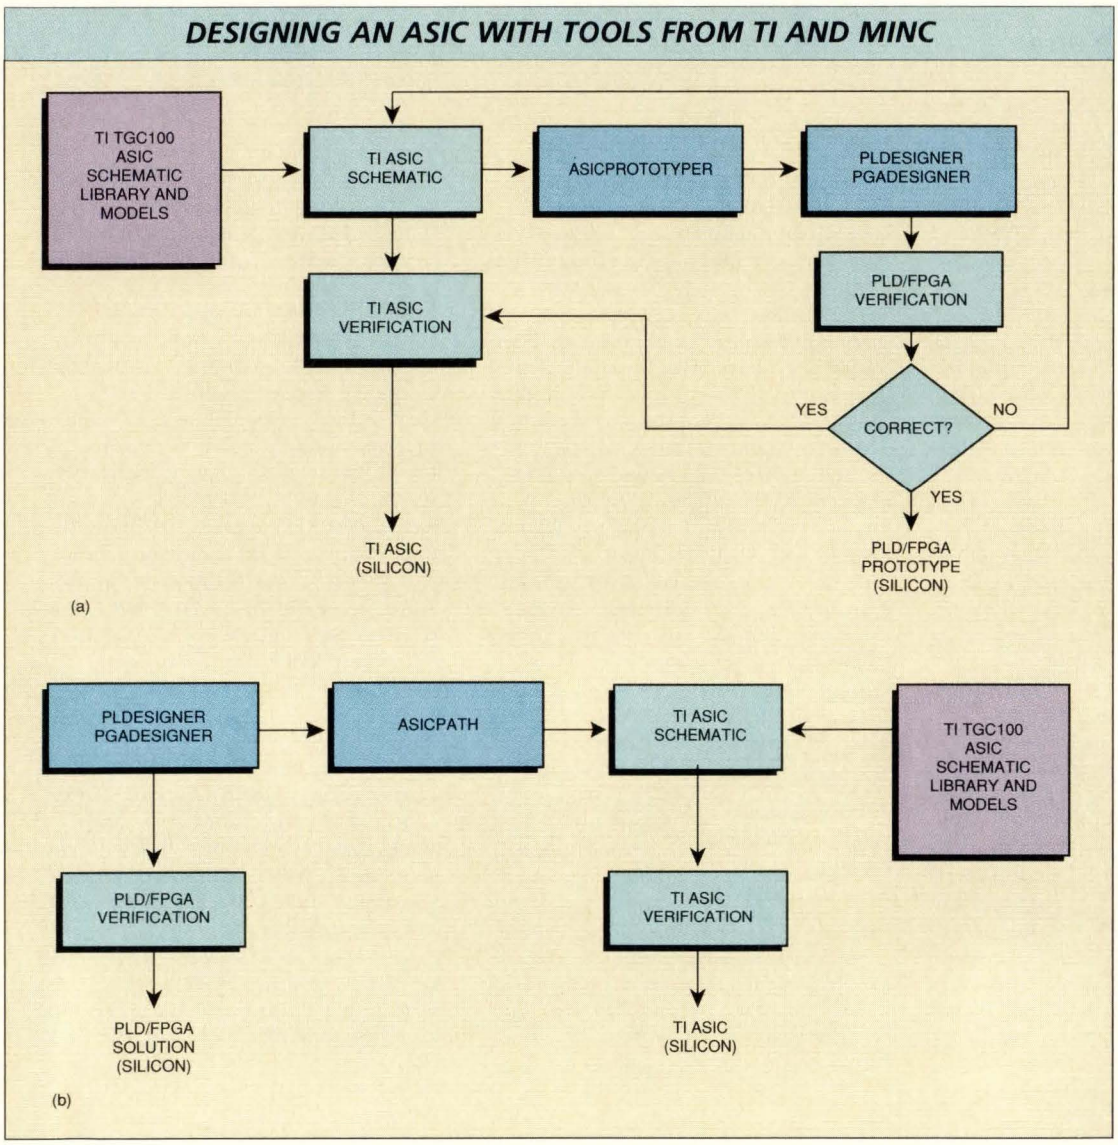

(a) Users who wish to design an ASIC first and then prototype it in programmable logic can use Texas Instruments' TGC100 design kit in the Mentor Graphics or Valid Logic Systems CAE environments. Staying within the same environment, they can then use the ASICPrototyper tool, jointly developed by Minc and TI, to generate a functionally equivalent PLD/FPGA design automatically from the ASIC design. The prototype device or devices can be implemented with Minc's PLDesigner or PGADesigner synthesis tools. (b) For users who have implemented a design in programmable logic and want to migrate the design to ASIC technology, Minc and TI offer ASICPath, which automatically converts a PLD/FPGA design from Minc's synthesis tools into the TI TGC100 ASIC environment.

tion was translated and synthesized to the gate level, it's been difficult to backtrack to the VHDL from the gates. A feature within the synthesis tool, called Cross Correlation, does double duty for the designer. Not only does it trace from gates back to the particular section of VHDL code that generated those gates, it points the designer from the VHDL code to the particular bunch of gates representing that code.

Wanted: FPGA support

The productivity advantages of VHDL are also motivating FPGA vendors, gate array vendors and users to seek support for FPGAs from Synopsys. Synopsys doesn't have synthesis libraries for FPGAs, though its Library Compiler can be used to generate an FPGA library.

Synopsys marketing director Robert Smith says that support for FPGAs will be customer-driven. "Though we're not in a position to serve PLD designers looking at FPGAs as the next step, we are in a position to serve gate array vendors looking at FPGAs as an alternative method," Smith claims. Because of the architectural peculiarities of FPGAs, supporting them with synthesis tools is more difficult than supporting gate arrays or standard cells. "In addition to partitioning, the tricky part is predicting the delay, especially with Xilinx Logic Cell Arrays," says Smith.

Andy Haines, marketing director at Actel (Sunnyvale, CA), reports that several customers designing with the company's antifuse-based FPGAs have interfaced Actel-specific design tools to Synopsys synthesis software, created libraries and synthesized with success. "The evidence is that we could use Synopsys [software] as is, since our FPGAs are so much like conventional gate arrays," says Haines. But the push is still on for FPGA vendors and synthesis tool vendors to establish interfaces from FPGA vendor-specific tools to high-level synthesis tools. And it's no wonder, he maintains. "Describing an FPGA design in a high-level language can compress design entry time from three days to three hours."

The Cadence Design Systems (San Jose, CA) strategy for FPGA design is to integrate vendor-specific FPGA tools. Vendors to be supported by Cadence include Xilinx, Actel, Plessey Semiconductors, Altera and Plus Logic. Cadence vice-president of marketing Anthony Zingale claims that as the novelty of FPGAs

Voice Processing From An Entirely New Perspective

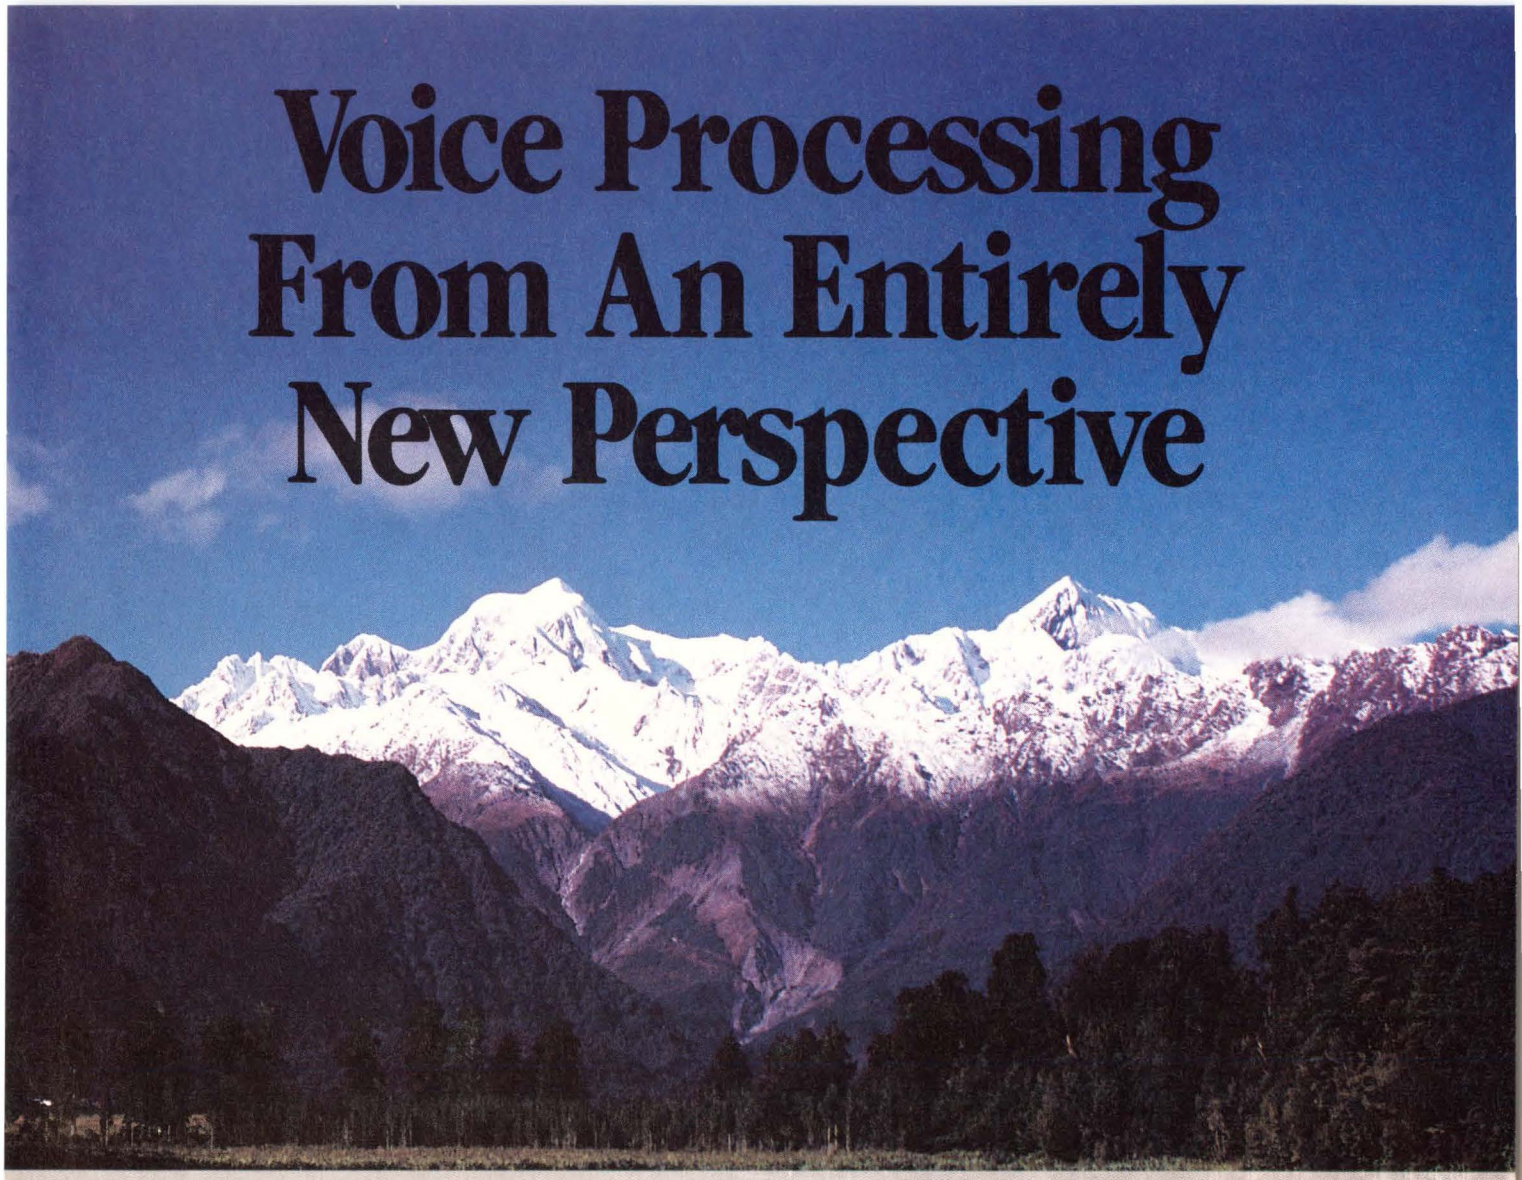

Discover the quality voice processing products, unparalleled distribution, superior service, and three year warranty of InnoVoice.

Innovative voice processing I/O boards from InnoVoice utilize our patented compression technology. The boards provide high quality voice/music capture and reproduction at low bit rates, giving you the building blocks to develop high performance systems at *competitive system cost.*

Innovative, Quality Products

OctaLink™

- 8 Channel Single Slot Voice I/O board;
- Supports industry-standard software;
- On board 8×12 analog crosspoint matrix switch;
- Applications include Interactive Voice Response, Voice Mail/Messaging and Automated Attendant.

StereoLink™

- 2 Channel Stereo Card;
- High quality voice/music for PC based multi-media presentations;
- Selectable storage rates provide sound quality from AM to near CD levels using only a fraction of conventional CD memory.

SpeechLink™

- 1 Channel Voice Card;
- Applications include Voice over E-mail, Digital Dictation and Annotation of Complex Graphics or Text.

Nationwide Service and Support

InnoVoice provides pre-sales support, applications and engineering assistance, and post-sales support. And all our voice processing boards are backed by a **full three year warranty,** the best in the industry! In addition to serving our customers via our 800 telephone service, all InnoVoice products are now available through Arrow Electronics, Inc.. In partnership with Arrow, we'll get you what you need, when you need it.

Our mission is to provide total customer satisfaction. We believe that by being responsive to your needs, we can provide better products, better service *and* better delivery than anyone else.

Call InnoVoice today. Let us show you what "A New Perspective in Voice Processing" can mean to your bottom line.

INNOVOICE™

Innovative Voice Solutions™

CIRCLE NO. 64

10700 Parkridge Boulevard
Reston, Virginia 22091
800-347-6007
A division of Extrema Systems International Corporation.

To help you design high dens the precise measurements o

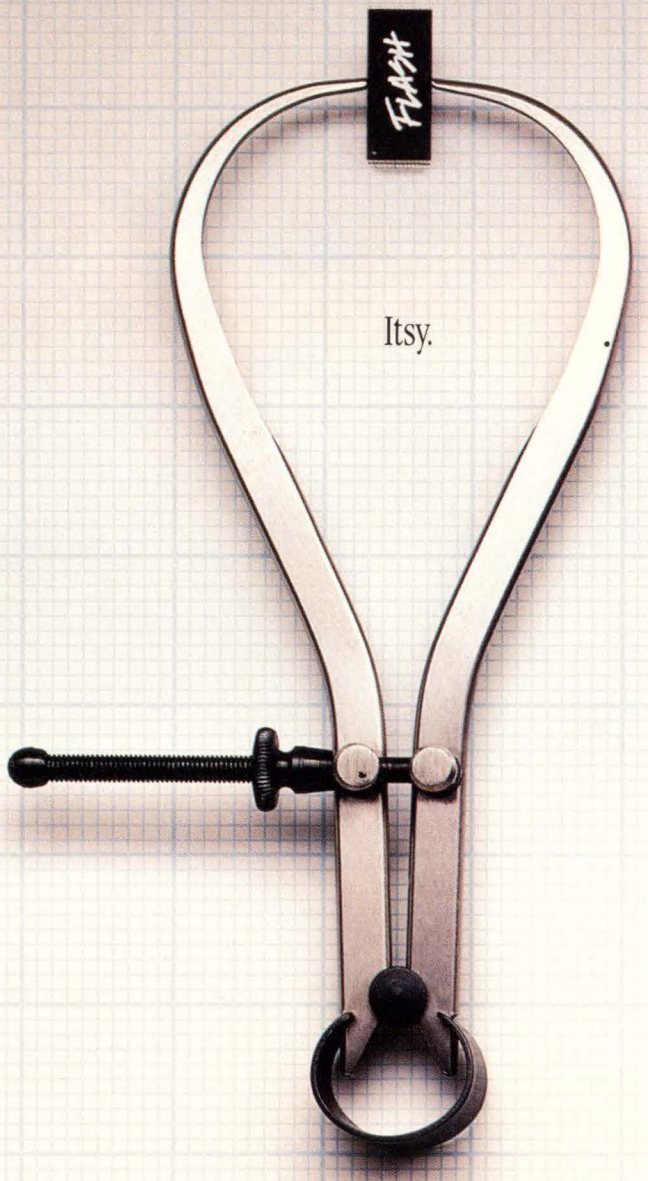

The biggest advancement in designing for portable and space-constrained products happens to be very, very small. Introducing our 1 and 2 Megabit Flash Memory in TSOP.

Now you can pack high density into a space one-third the size of a traditional 32-pin PLCC. And not just any memory, but a nonvolatile, rewritable, 100,000 typical cycle memory in a package that's only 20 x 8 x 1.2 mm in size. Or in non-technical terms, teeny x weeny.

And if that isn't small enough, our new TSOP Flash permits a compact serpentine layout. This side-by-side, end-to-end technique saves board space and layers. Which means you can design products with greater functionality into a smaller space – such as high-density IC cards. In fact, we used this layout to create our own Flash Memory cards.

ty into small spaces, here are)ur new TSOP Flash Memory:

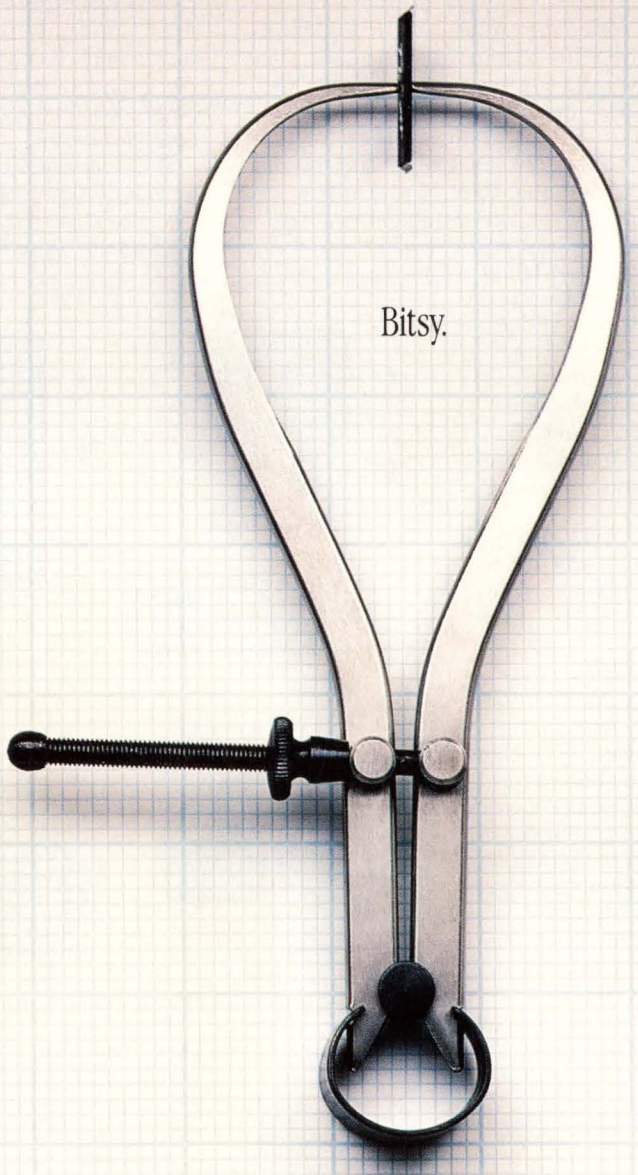

The benefits of designing with our flash go beyond its size. Since it combines RAM, ROM and mass-storage capabilities onto a solid-state, energy-efficient medium, it lets you eliminate slow hard disks and bulky battery packs from your design. In other words, you'll create products that are faster, lighter, more rugged and require less power.

But best of all, since it's manufactured on our proven ETOX™ II process, it's already shipping in high volume.

To find out more about this small wonder, call Intel at (800) 548-4725 and ask for Literature Packet #A6A31. You see, when it comes to designing innovative products, it's the only memory that really measures up.

intel®

CIRCLE NO. 65

FPGA devices require FPGA-specific synthesis tools

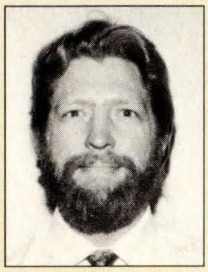

Logic synthesis for field-programmable gate arrays is different from logic synthesis for programmable logic devices or gate arrays because of fundamental architectural differences. Unlike the fixed sum-of-products structure of PLDs, the structure of FPGAs is flexible and multilevel. In this respect, FPGAs are similar to gate arrays, although the basic logic element of an FPGA is more complex than the basic logic gates commonly found in gate array libraries.

The architectural differences call for different requirements for synthesis. Blind application of PLD or gate array synthesis tools to an FPGA design results in a design that runs too slowly or is too large, or both.

Consider the architecture of the Xilinx (San Jose, CA) 2000 Logic Cell Arrays as an example that illustrates the requirements of FPGA synthesis. These devices have limited fan-in logic cells, called CLBs (configurable logic blocks), that can be any function of four inputs. A four-input XOR uses the same space and is as fast as a four-input AND gate. Multiple levels of logic cells must be used to implement functions of greater than four inputs. Logic design for Xilinx devices is therefore limited by fan-in—not by logic complexity as in PLDs.

Logic function options

One approach to Logic Cell Array (LCA) design is to decompose a function into its more-simple AND/OR equivalent representation, and then split the gates with large fan-in into multiple gates. Consider the following logic function as an example:

X=A(B+C)+(B×D)+(E×F×G×H×I)

Represented in AND or OR gates, the function would be as follows:

X=T1+T2+T3
T1=A×T4
T2=B×D
T3=E×F×G×H×I
T4=B+C

Next, reduce T3 to four inputs by creating another expression:

T3=E×F×G×T5
T5=H×I

Once the design is in AND/OR format, the Xilinx physical design software can be used to place the design into CLBs. This process is referred to as "partitioning." For this example, we can observe what options are possible and choose the best one. T5 or T3 may not be merged because of fan-in limitations. T1 may be combined with X, but it would be better to combine T1 with T4 and X with T2. This gives the following partitioning into four CLBs (see diagram, left side):

X=T1+(B×D)+T3
T1=A(B+C)
T3=E×F×G×T5
T5=H×I

The best partitioning for this example, however, isn't the best solution overall. The initial logic decomposition given to the partitioning algorithm strongly affects the final solution. A different decomposition yields a much better partitioning into three CLBs (see diagram, right side):

X=T1+(T2×E)
T1=A(B+C)+(B×D)
T2=F×G×H×I

It's clear that this is an optimal solution, since the three expressions share no common inputs.

The process of mapping arbitrary logic into simple AND/OR logic is trivial if you don't have to worry about optimizing for the underlying technology.

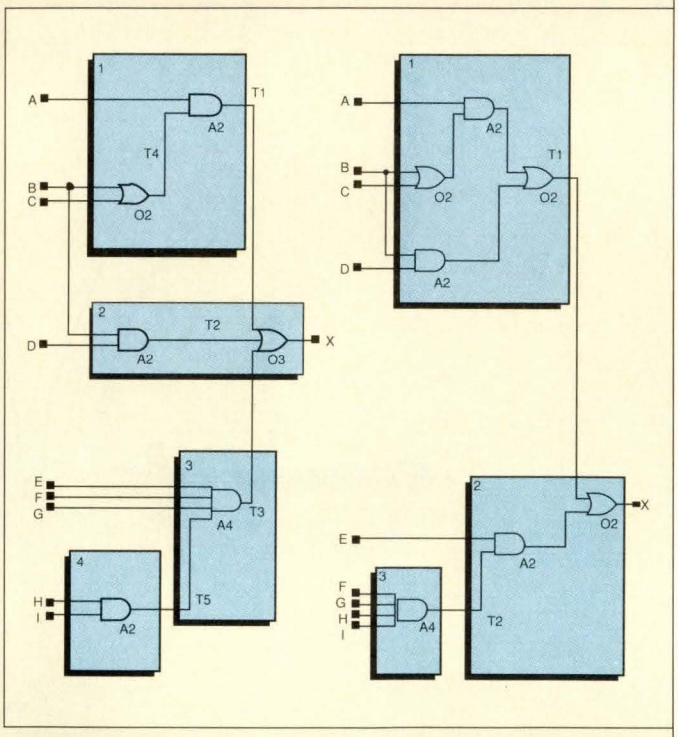

This technique can produce good results, however, when coupled with FPGA-specific optimization software. As the example shows, fan-in-limited decomposition is required for the Xilinx LCA architecture. In the case of the Actel (Sunnyvale, CA) ACT 1 architecture, where the basic logic building block is a multiplexer, a multiplexer-based decomposition is a major optimization technique.

FPGA-specific synthesis is the platform for future FPGA design technologies. As FPGAs move further into the mainstream, system designers will want performance information early in the design cycle, before implementation. FPGA-specific synthesis is the technology that will permit this flexibility.

In addition, FPGA designers will demand multichip partitioning and device selection aids. FPGA-specific synthesis techniques are crucial to these since they both require timing optimization with respect to the underlying device architecture.

Ewald Detjens, MSEE, president and founder, Exemplar Logic

THE NO COMPROMISE 22V10

Programmable logic always has been a give-and-take affair. If you wanted speed, the price was power—lots of it. And, if you tried to cut power, you lost the speed. It seems you could have either one or the other—but not both. **NOW YOU CAN HAVE THE BEST OF BOTH WORLDS.**

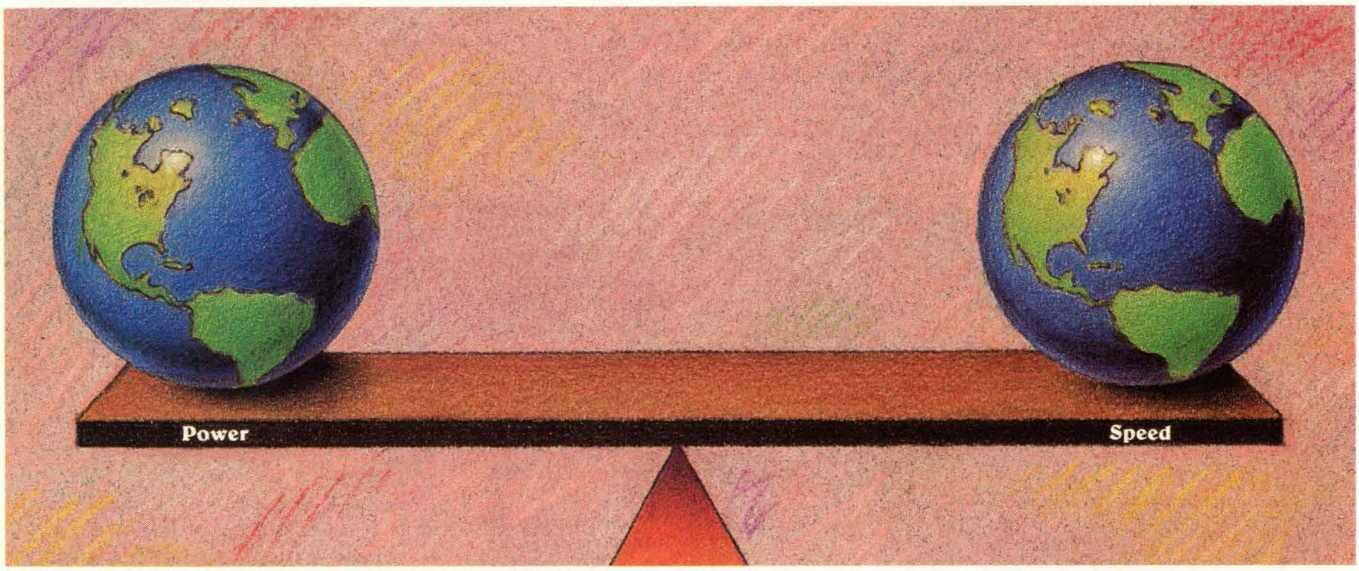

Announcing the AT22V10-15—the no compromise 22V10.

Talk about fast. A blazing 15 nanoseconds. That's fast enough for those advanced 32-bit systems you're designing today for tomorrow's machines.

And it's cool. When you plug in the AT22V10 you won't even think it's on. It typically draws a stingy 55 milliamps in standby and never asks for more than 90 milliamps.

So, if you're tired of having to compromise. Don't. Call Atmel, the home of the no compromise 22V10. If you're not sure yet that we make the best CMOS 22V10 in the whole world drop us a note on your company's letterhead, and we'll send you one. Or in the U.S. call us at **1-800-292-8635.**

ATMEL®

The people who make the difference.

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

Tel. (408) 441-0311
FAX (408) 436-4200

CIRCLE NO. 66

wears off, system designers will refuse to go outside top-down design environments to use proprietary tools. FPGA and PLD vendors may continue to offer proprietary tools, he says, but they will be integrated within design systems.

Cadence recently integrated the de facto industry-standard PLD synthesizer Abel from Data I/O (Redmond, WA) into its Amadeus board/system design environment as an option for PLD and FPGA designers. Though Abel-4 addresses device-independent design through its new hardware description language, Zingale doesn't think Abel-HDL will be enough. "Customers will drive the use of a single language," he predicts. "Cadence has customers who just want to work with Verilog, or just with VHDL."

About the same time Cadence began OEMing Abel, Data I/O initiated the Open Abel program. To accommodate FPGAs and complex PLDs such as the Mach family from Advanced Micro Devices (Sunnyvale, CA) and the Max family from Altera (San Jose, CA), Open Abel licenses chip vendors to write "fitters" to optimize their silicon. The Open Abel interface lets designers use the new Altera Abel2MAX software to enter their logic designs in the Abel language and then compile and simulate them with the Altera MAX+PLUS software.

Peculiarities of FPGAs

Commenting on the need for such optimization software for FPGAs, Ravi Ravikumar, product manager in the Cadence systems division, says, "Logic synthesis of FPGA-based designs can be different from that of gate arrays or other forms of masked ASICs. In the synthesis of an FPGA design, the optimization is both logical and physical. This requires intimate knowledge of FPGA device architectures and the available physical resources."

Clay Marr, vice-president of marketing and sales at Plus Logic (San Jose, CA), echoes Ravikumar's statement on the difficulty of synthesizing programmable devices such as FPGAs. "Synthesis tools as they are can't take advantage of the target architecture of FPGAs," he says. "A delicate balance is required to take a logic definition and do an optimized implementation—it's going to be different for every FPGA architecture out there. The idea is to optimize for performance and to minimize interconnect at the same time, to create logic elements as large as you can."

Software vendors probably aren't the best people in the world, says Marr, to optimize for an FPGA vendor's silicon. Through an improved optimizer/minimizer and automatic partitioning, the Plustran 2.0 version of the Plus Logic design system provides an average of 20 percent greater utilization of Plus FPGAs over the previous version, according to Marr. Plustran reads a circuit description, compiles the information into optimized logic patterns, and then maps those logic patterns onto the target device.

The first FPGA-specific, multiple-vendor VHDL logic synthesis and optimization toolset was introduced about six weeks ago by Exemplar Logic (Berkeley, CA), which focuses its R&D efforts on the development of FPGA synthesis solutions. Having been a consultant and supplier of FPGA synthesis tools to companies such as Xilinx (San Jose, CA) and Actel for the last few years, Exemplar will now OEM its FPGA Compiler so that it can be integrated with front-end tools from major EDA and FPGA vendors. Exemplar tools are currently integrated into design systems marketed by FPGA vendors Xilinx and Actel and EDA vendor Viewlogic Systems.

The FPGA Compiler accepts VHDL models, descriptions from the Palasm language and EDIF (Electronic Data Interchange Format) 200 netlists. The output is an FPGA-specific netlist, chip area and speed estimates and an EDIF 200 sche-

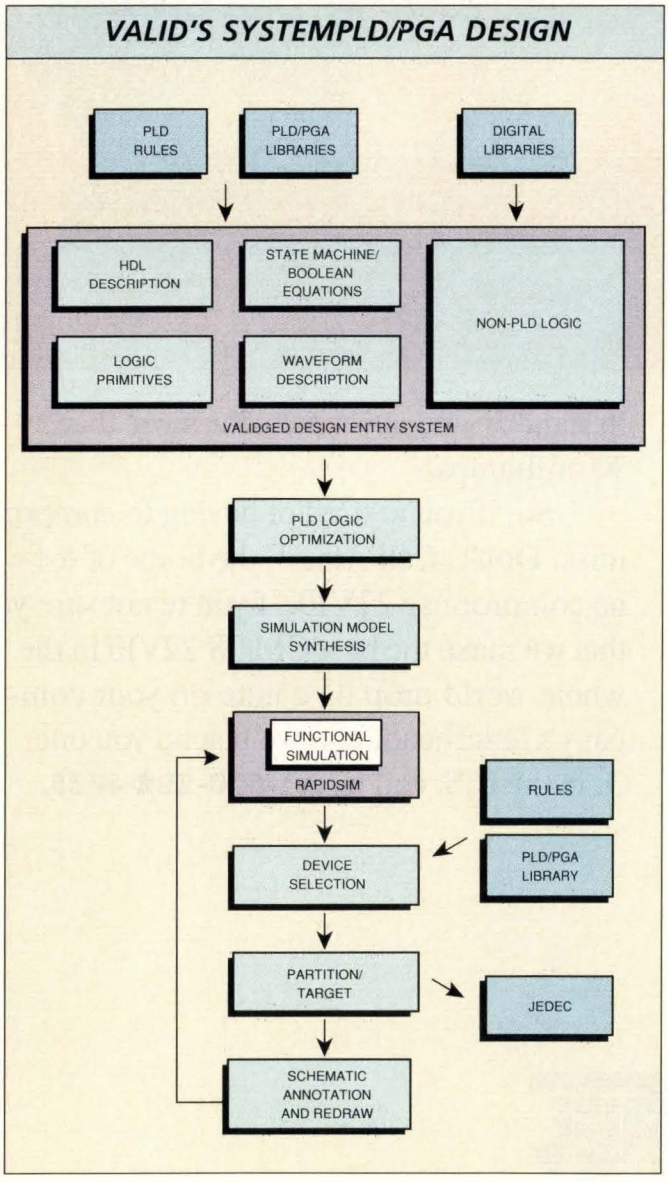

The Valid Logic Systems SystemPLD/PGA software fully integrates the PLD/FPGA design and synthesis technologies of Minc with the Valid Logic Workbench digital design environment. Users can mix PLD and FPGA logic throughout system-level schematics, perform complete system simulation with Rapidsim before device selection, and automatically combine or retarget PLDs and FPGAs without design description changes. A schematic redraw capability smoothes the integration with physical design, packaging and back annotation.

200 MOPS.

(Otherwise known as how to clean up in multimedia.)

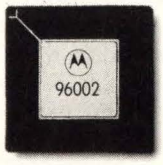

Finally there's a processor that's fast enough for the next generation of multimedia computers. Introducing the 96002 Media Engine™ from Motorola.

With a peak performance of 200 million operations per second (MOPS) at 40 MHz, the 96002 has no equal. It includes a high performance IEEE floating point engine. And more importantly, four independent function units and dual ports for extremely high parallelism. Which means it can handle image processing or graphics and stereo audio simultaneously.

And instantaneously.

To learn more, call 512-891-3230. With 200 MOPS you'll clean up faster than anybody.

MOTOROLA

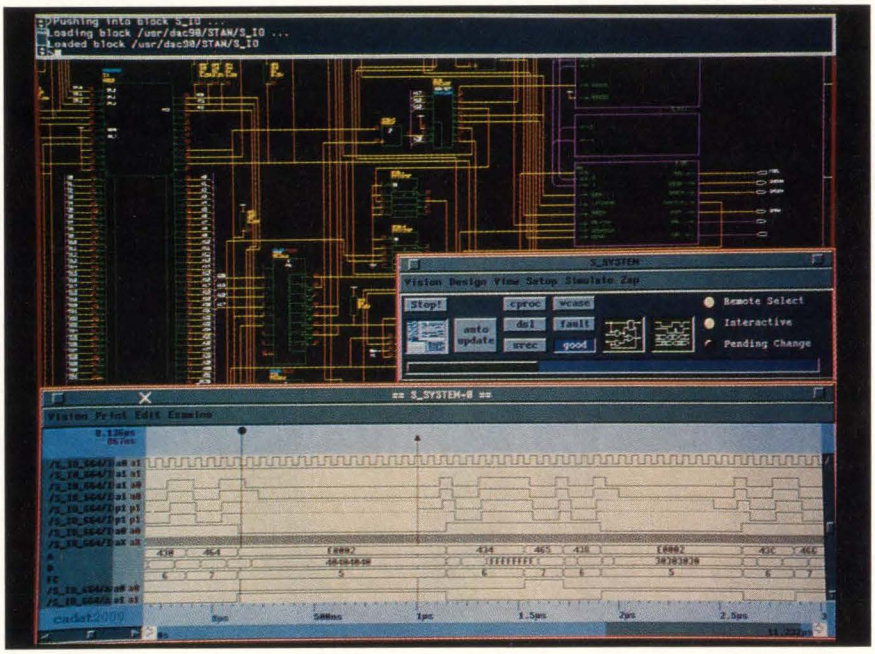

With Minc's synthesis tools integrated into Racal-Redac's design automation environment, PLD and FPGA designs can be described with several device-independent entry methods, including Racal-Redac's Visula schematic editor.

matic. The focus of the FPGA Compiler, according to Exemplar, is to provide an easy way to migrate PLD- or gate array-captured designs onto FPGAs. Users can evaluate the performance characteristics of the design in a particular FPGA technology before doing vendor-specific place and route.

VHDL compilation for PLDs, FPGAs and gate arrays is now available to users of the Isdata (Monterey, CA) LOG/IC synthesis software, which interfaces to a number of CAE design environments. After the company's Hint VHDL compiler defines designs, the LOG/IC software performs logic reduction and produces programming and test data for PLDs and netlists for FPGAs and gate arrays.

Customers and vendors alike are discovering that going from an FPGA netlist to a gate array netlist is fraught with difficulties. Steve Brightfield, CMOS marketing group manager at Plessey (Scotts Valley, CA), which takes its own Electrically Reconfigurable Arrays as well as competitive FPGAs into Plessey gate arrays, maintains, "There's a lack of visibility. You get functional equivalence but not observability equivalence. Some of the macrocells don't map well, and not all the test vectors are usable."

Tools that allow designers to go from ASIC to FPGA and from FPGA to ASIC will be available soon as the result of an agreement between Minc (Colorado Springs, CO) and Texas Instruments (Dallas, TX), which is now a maker of FPGAs as part of a relationship with Actel. (AT&T, which has an agreement with Xilinx, also began shipping FPGAs recently, which extends its product line to all but PLDs.) And customers are demanding tools similar to those from TI/Minc from other gate array vendors. Gary Hess, ASIC and custom products marketing director at Oki Semiconductor (Sunnyvale, CA), says, "Over the last six months, Oki has been feeling more pressure to provide a path from FPGAs to gate arrays."

> ***"Synthesis tools as they are can't take advantage of the target architecture of FPGAs."***
>
> *—Clay Marr, Plus Logic*

Hess is of the opinion that designers should select a gate array vendor up front, establish a gate array netlist and move from that to an FPGA. He claims that in a good percentage of situations where synthesis programs are used to map from an FPGA to a gate array, the parts don't work. "There's no way to get good timing data out of an FPGA," says Hess. "Viewlogic can produce a timing file for a Xilinx part, but it uses the spec for the part, assuming the design will be typical."

J. Scott Runner, manager of ASIC applications engineering at Fujitsu Microelectronics (San Jose, CA), shares Hess' concern about a sloppy migration from FPGA to gate array. "Just because a design works as a prototype doesn't mean that the designer doesn't have to verify timing and develop test vectors." Runner says the best method is to use a synthesis tool to map from the FPGA to a particular vendor's library. "Once you have synthesis libraries, as long as you can read the netlist format, you can go from FPGA to ASIC, or from ASIC to FPGA. From an ASIC to an FPGA, you might have the problem of partitioning into multiple FPGAs."

Runner says that working with tools that make VHDL available can have a significant impact. With VHDL, he says, designers can write a model with timing specified and then drive a synthesis tool with that model. But without VHDL, timing has to be specified in a constraint file and separate models created for synthesis and simulation.

VHDL impact

No one disputes the advantage of describing a design that may be retargeted to another technology in VHDL. "Whether a VHDL description is targeting an FPGA or a gate array, there's no difference in the input. It doesn't impact the design process at all," claims Jay Southard, principal engineer of synthesis at Viewlogic Systems (Marlborough, MA), which offers a VHDL synthesis tool for FPGAs. "A user can do a design once in VHDL, implement it as an FPGA, get to a production level, take the VHDL description and resynthesize it for a CMOS gate array," he says. On the other hand, if all you have is an XNF (Xilinx Netlist Format) file, it's an expensive proposition.

Though Valid Logic Systems (San Jose, CA) offers VHDL and has fully integrated the Minc PLD and FPGA synthesis tools into its Logic Workbench digital design environment, the company doesn't offer VHDL as an entry method for FPGAs. Valid says that a timetable for the availability of VHDL is up to Minc. Minc's

REAL-TIME.
REAL-FAST.
REAL-POWER.

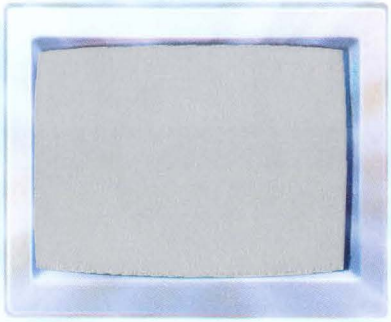

Scratch beneath the surface for a real-time speed and power breakthrough...

There's a new speed and power standard in Real-Time. It's VxWorks™ Version 5.0 featuring the **wind**™ real-time kernel. Preliminary, conservatively rated benchmarks for **wind** using an unmodified 68020 25MHz VME board, clock context switches at 17μs, semaphore give + takes at 8μs and interrupt latency at 8μs. Times for other processors can (and will) be faster.

With the blazing **wind** kernel, VxWorks 5.0 offers a commanding combination of performance, transparent connectivity and development cycle efficiency. For the full story on VxWorks 5.0, turn to the flip side of this bulletin.

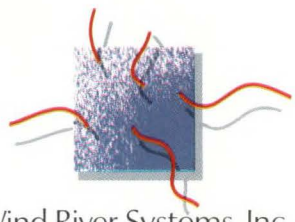

Wind River Systems, Inc.

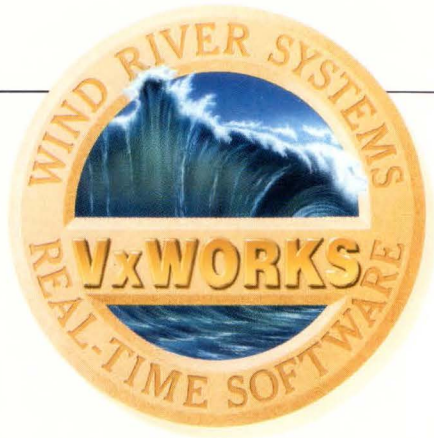

Real-Time.

VxWorks 5.0 is a powerful operating system built on a cross development partnership with UNIX. And leveraging on a broad selection of industry standards, VxWorks 5.0 supports a complete portfolio of CISC and RISC target processor architectures, host/development platforms and high-level languages including C and Ada.

Real-Fast.

The heart of VxWorks 5.0 is our new **wind** kernel. Incorporating object-oriented design concepts, **wind** offers unparalleled speed, functionality and determinism. New features include constant speed context switches regardless of the number of tasks waiting in queues, priority inheritance semaphores and time-outs.

Real-Power.

VxWorks provides transparent connectivity to access the power of distributed systems. Its integrated networking features include 4.3 BSD UNIX compatible sockets and support for protocols such as TCP, IP, UDP, *telnet, ftp, rlogin, rsh,* RPC and NFS. In addition, powerful symbolic and source level debuggers enable high-resolution access to all processes and tasks in an application.

VxWorks Real-Time Operating System

- ☐ High performance, multitasking **wind**™ kernel
- ☐ Integrated networking facilities
- ☐ UNIX partnership-based development environment
- ☐ Full-featured symbolic and source level debugging
- ☐ Easily configurable execution environments
- ☐ X Window System client-side support
- ☐ Device-independent I/O and file systems

Real-Commitment.

VxWorks has driven Wind River Systems to become the world's fastest growing real-time software company. We back up our users with a full range of services including training, technical support, maintenance, a users group, porting and consulting.

But the real-time bottom line is to provide developers with software tools as powerful as the advanced hardware technology they support. VxWorks 5.0 is our latest delivery on that promise. For more information on VxWorks 5.0 and **wind,** call 1-800-545-**wind.**

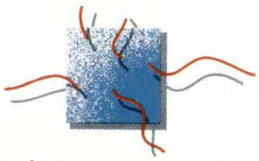

Wind River Systems, Inc.

1010 Atlantic Avenue,
Alameda, CA 94501
Phone: **1-800-545-wind**
FAX: 1-415-814-2010

Wind River Systems Global Service Network.

East Coast
Wind River Systems, Inc.
77 North Washington Street
5th Floor
Boston, MA 02114
Phone: 1-617-367-6567
FAX: 1-617-367-9295

Europe
Wind River Systems EC
Z.I. de la Bonde
15 Rue du Buisson aux Fraises
91300 Massy, France
Phone: 33 (1) 60 13 06 07
FAX: 33 (1) 60 13 04 39

Japan
Wind River Systems, K.K.
1-19-9 Toranomon
Minato-ku, Tokyo, 105
Japan
Phone: 813-502-5580
FAX: 813-502-5569

vice-president of marketing, Kevin Bush, says that it's Minc's intention to support VHDL, but that in the meantime, the company's proprietary high-level language is supplying a lot of the same capabilities. "If everyone supported VHDL as a synthesis language, you could bring the VHDL description to an ASIC vendor and be much closer to a workable solution," says Bush.

Before the end of the year, Racal-Redac will also be shipping software that integrates the Minc PLD and FPGA synthesis tools into its system that includes the CADAT digital simulation system and the Visula EDA design environment.

Though Mentor integrates Minc's PLDesigner into its system, the company doesn't offer PGADesigner for FPGAs. Instead, Mentor reports partnerships with major FPGA vendors that will result in coupling Mentor's VHDL synthesis to vendor-specific back ends for placement and routing. The Mentor generic library, called GenLib, will be the glue between the front and back ends.

At this point, it almost seems as if we have a hodgepodge of synthesis solutions—from Abel, Logical Devices' (Fort Lauderdale, FL) Cupl and AMD's Palasm, to Minc and vendor-specific tools, to high-level VHDL synthesis tools being integrated into workstation-based system-level design environments. Indeed, there are a lot of pieces to the puzzle, and it may be some time before we put it all together. A look at the missing pieces to a mature synthesis solution sort of takes the magic out of the technology. That's just as well, though, because when the puzzle is complete, synthesis may prove to be one of the most practical technologies ever. ■

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

Actel
(408) 739-1010 . . . . . Circle 274
Advanced Micro Devices
(408) 732-2400 . . . . . Circle 275
Altera
(408) 984-2800 . . . . . Circle 276
Cadence Design Systems
(408) 943-1234 . . . . . Circle 277
Data I/O
(206) 881-6444 . . . . . Circle 278
Exemplar Logic
(415) 849-0937 . . . . . Circle 279
Fujitsu Microelectronics
(408) 922-9000 . . . . . Circle 280
Isdata
(408) 373-7359 . . . . . Circle 281
Logical Devices
(305) 974-0975 . . . . . Circle 282
LSI Logic
(408) 954-4723 . . . . . Circle 283
Mentor Graphics
(503) 626-7000 . . . . . Circle 284
Minc
(719) 590-1155 . . . . . Circle 285
Oki Semiconductor
(408) 720-1900 . . . . . Circle 286
Plessey Semiconductors
(408) 438-2900 . . . . . Circle 287
Plus Logic
(408) 293-7587 . . . . . Circle 288
Racal-Redac/ASIC Design Div
(508) 692-4900 . . . . . Circle 289
Synopsys
(415) 962-5000 . . . . . Circle 290
Texas Instruments
(800) 336-5236, ext. 700 . . . . . Circle 291
Valid Logic Systems
(408) 432-9400 . . . . . Circle 292
Viewlogic Systems
(508) 480-0881 . . . . . Circle 293
VLSI Technology
(408) 434-7660 . . . . . Circle 294
Xilinx
(408) 559-7778 . . . . . Circle 295

128 MBYTE MEMORY, VME
OVER 30 MBYTES/SEC TRANSFER RATE

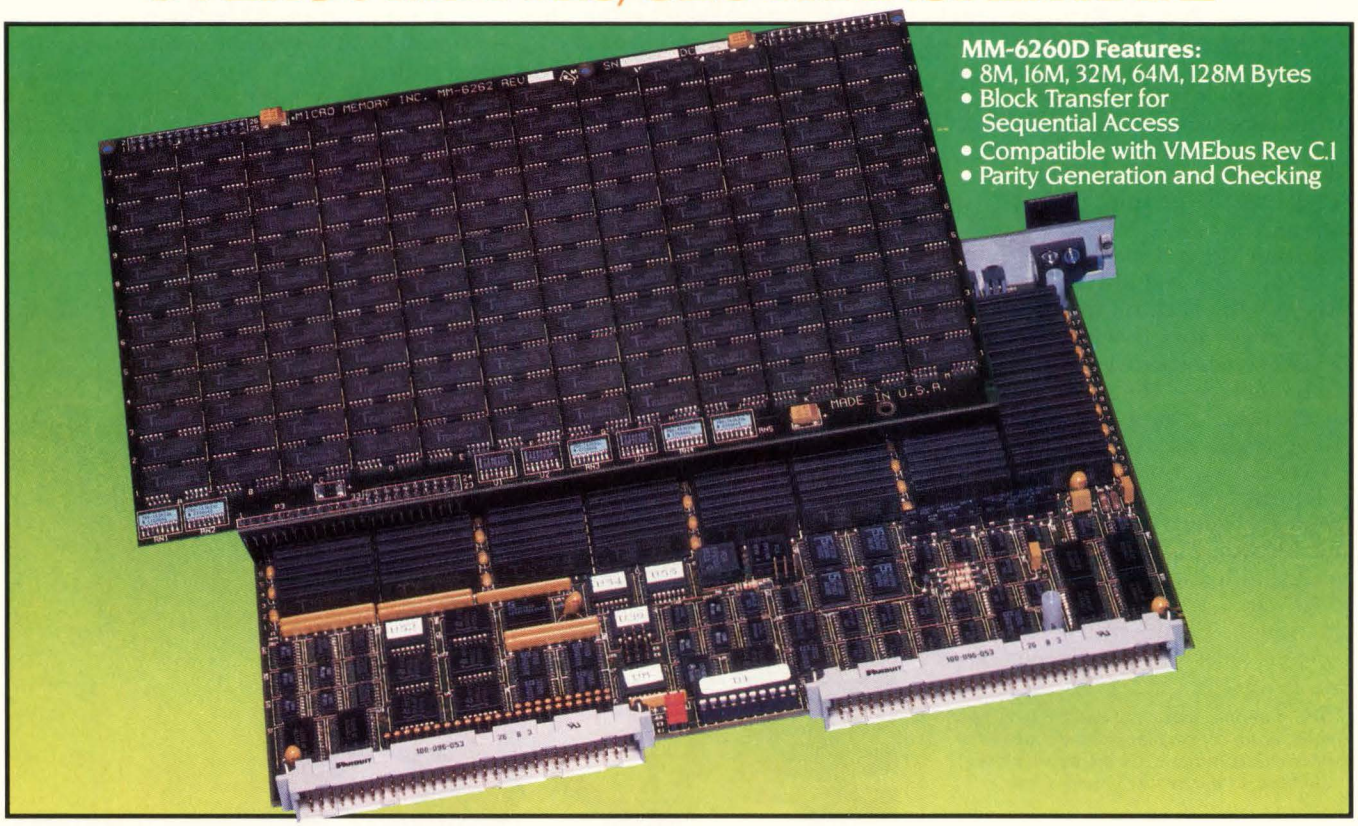

MM-6260D Features:
- 8M, 16M, 32M, 64M, 128M Bytes
- Block Transfer for Sequential Access
- Compatible with VMEbus Rev C.1
- Parity Generation and Checking

1M - 128M Byte Boards — For Every Application

VME DRAM Memory Boards (Lifetime Warranty)

Part Number	Capacity, Bytes	Cycle/Access, nsec	Operating Amps	Remarks (All DRAM Boards A32/D32)
MM-6340D	4M - 16M	240/220	6.5	ECC, VME/VSB, UAT, BLT, Page
MM-6326D	8M - 64M	125/75	5.5	Parity, Dual-Port, VME/VSB, UAT, BLT
MM-6320D	4M - 16M	275/200	4.8	Parity, Dual-Port, VME/VSB, UAT, BLT
MM-6316D	4M - 16M	275/220	4.8	Parity, Dual-Port, VME/VSB, UAT, BLT
MM-6260D	8M - 128M	120/90	3.8	Parity, UAT, BLT, Page, A32/A24, D32/D16/D8
MM-6240D	2M - 16M	250/130	4.5	Error Correction, 32-bit CACHE, UAT, BLT
MM-6230D	4M - 16M	240/175	2.9	Parity, Fast Write, UAT, BLT
MM-6220D	2M - 16M	240/175	3.8	Parity, 32-bit CACHE, CACHE Hits = 75 nsec
MM-6216D	4M - 16M	240/175	3.0	Parity, Fast Write, UAT, BLT
MM-6202D	1M - 4M	145/130	2.9	Parity, UAT, BLT, Page

VME CMOS Memory Boards: Non-Volatile, On-Board Batteries (One-Year Warranty)

Part Number	Capacity, Bytes	Cycle/Access, nsec	Data Retention Lithium/NiCad	Operating Amps	RAM/EPROM Mix for CMOS Boards
MM-6704C	512K - 4M	150/100	2 yrs./2 wks.	2.1	RTC, UAT, BLT, A32/D32
MM-6702C	512K - 8M	125/125	3 yrs./4 wks.	1.7	UAT, BLT, A32/D32
MM-6700D	64K - 1M	150/150	4 yrs./8 wks.	1.0	A32/D32
MM-6500C	32K - 512K	200/200	5 yrs./12 wks.	1.1	RTC, A24/D16

Micro Memory Inc. also supplies a full line of Multibus I and II, PC/AT, EISA and Sun Workstation Memory Boards.
Multibus is a trademark of Intel Corp., PC/AT is a trademark of IBM Corp., Sun is a trademark of Sun Microsystems Inc.

Micro Memory Inc

CIRCLE NO. 69

...First in Microcomputer Memories

VMEbus INTERNATIONAL TRADE ASSOCIATION VITA

9540 Vassar Avenue • Chatsworth, CA 91311 • (818) 998-0070 • FAX: (818) 998-4459

COMPUTERS AND SUBSYSTEMS

Dual-port designs help memory boards keep pace

Jeffrey Child, Associate Editor

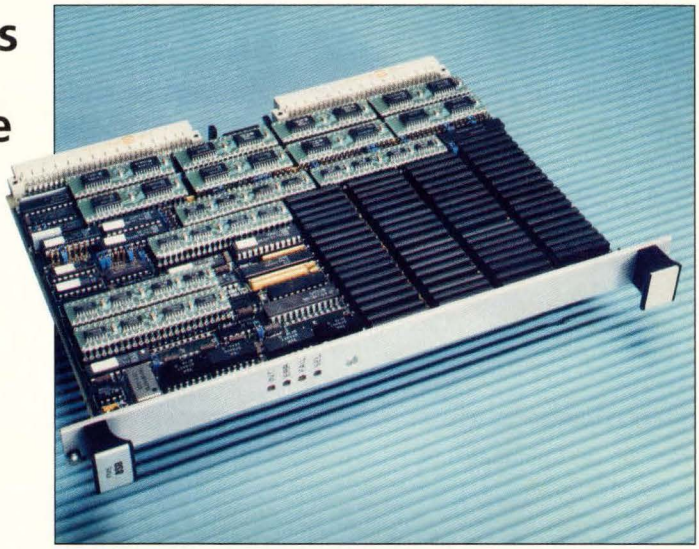

The 8SB, a dual-port DRAM board from Radstone Technology, provides up to 8 Mbytes of DRAM and has bidirectional data buffers that interface directly to the VMEbus and the VME subsystem bus. Each bus has its own set of latched buffers, so internal logic is free from the timing requirements of either bus.

Upgrading system performance is no longer a matter of simply tossing another memory board on the bus. The quicker access times of the newest disk drives and the faster clock speeds of the latest microprocessors have outdated such a simple solution.

Today, most VME users would rather squeeze DRAM chips onto the CPU card instead of using precious bus bandwidth to execute code. And with DRAM densities reaching 4 Mbits, CPU boards with from 4 to 16 Mbytes of on-board memory are commonplace. As a result, the role of today's VMEbus memory boards has moved from simple memory expansion to a global memory source to be shared between processor cards on the bus.

System designers have learned to overcome bus limitations by offering dual-port memory solutions, so several of the latest VME memory boards also use the VME subsystem bus (VSB). The VSB lets designers off-load some of the main bus activity and manipulate data independently. It also accelerates data transfer by eliminating the 256-byte block-move barriers of the VMEbus.

"We've seen a tendency in the industry to press the limits of VME and end up with a bottleneck—for many system designers, there's never adequate bandwidth on the bus," says Joel Silverman, marketing manager for commercial products at Radstone Technology (Montvale, NJ). "It's taken a while for the industry to accept VSB, or accept anything else, as a secondary bus. But in the past 12 to 18 months we've seen a lot more VSB implementations."

A dual-port DRAM board, the 8SB from Radstone Technology, provides up to 8 Mbytes of DRAM and a full implementation of the VSB specification. Enabling data transfer rates in excess of 20 Mbytes/s across either bus, the 8SB performs data read-ahead on block transfers and pipelined address and data cycles. The board has bidirectional data buffers that interface directly to the VMEbus and the VSB. Each bus has its own set of latched buffers, which frees internal logic from the timing requirements of either bus.

Control logic on the 8SB consists of an asynchronous part and a state machine part. Optimized for single random accesses, the asynchronous part offers short access times—230-ns read access times and 130-ns write access times. The first access of any cycle is handled by the asynchronous part, while subsequent access within that cycle is handled by the state machine logic as block transfers.

Another dual-port memory board, the MVME224A from Motorola (Tempe, AZ), uses a patented arbitration scheme in which the arbitration time for that board into the DRAM refresh circuitry varies. While it can't support block transfers, the board runs at 200-ns cycle time doing straight one-to-one cycles. Implementing block transfer on this board would require some ASIC development, according to Motorola.

Designed for use in VME-based systems to increase global memory, the MVME224A contains 4 to 32 Mbytes of DRAM for 8-, 16- or 32-bit data transfers. The MVME224A's arbitration scheme minimizes cycle time by burying the time required for refresh, VME arbitration and VSB arbitration within the cycle. This time reduction lets the board achieve 20-Mbyte/s transfer rates over the VMEbus and VSB.

Dual ports fall short?

Although more and more board designers are making use of the VSB to achieve greater data transfer rates, one company has found reasons for doubt. "We surveyed the marketplace," says Mike Strang, vice-president of Advanced Technology at SBE (Concord, CA), "and although we found some cards that were dual-port, none of them could really support the simultaneous transfer rates that we required on both the VSB and the VMEbus.

"Some of the manufacturers are claiming 30-plus Mbytes/s—but the reality is that when you plug one of those boards into a system in a real-world application, you find there are only certain circumstances under which you get 30-plus Mbytes/s." If a system has two masters—one on VME and one on VSB—accessing the memory at the same time, the data rate drops dramatically, according to Strang.

Model	Bus	Memory type	Capacity (bytes)	Access time read (ns)	Access time write (ns)	Cycle time read (ns)	Cycle time write (ns)	Price	Comments
Aitech Defense Systems 1250 Oakmead Pkwy, Ste 210, Sunnyvale, CA 94086 (408) 720-9400									**Circle 301**
1C203	VME	EPROM, EEPROM, flash, SRAM	4M to 8M	75 to 330	75 to 330	75 to 330	75 to 330	—	6U, 16-/32-bit data/address width
2C203	VME	EPROM, EEPROM, flash, SRAM	4M to 8M	85 to 330	85 to 330	85 to 330	85 to 330	—	same as 1C203 but with –40°C to +71°C operating temp
4C203	VME	EPROM, EEPROM, flash, SRAM	4M to 8M	85 to 330	85 to 330	85 to 330	85 to 330	—	same as 2C203 but with –55°C to +85°C ambient operating temp
5C203	VME	EPROM, EEPROM, flash, SRAM	4M to 8M	85 to 330	85 to 330	85 to 330	85 to 330	—	same as 4C203 but with full mil-spec components
Central Data 1602 Newton Dr, Champaign, IL 61821 (800) 482-0315									**Circle 302**
CD21/2511	Multibus I	DRAM	2M	210 max	210 max	245 max	245 max	$1,885	sync/async LBX interface, maps on 64-kbit boundary
CD21/2032	Multibus I	EPROM, SRAM	1M to 2M	160 + device speed	160 + device speed	—	—	$520	battery backup, maps on 64-kbit boundary
CD21/2035	Multibus I	EPROM, EEPROM, SRAM	128k to 512k	90 + device speed	90 + device speed	—	—	$620	battery backup option
Clearpoint Research 35 Parkwood Dr, Hopkinton, MA 01748 (800) 253-2778									**Circle 303**
VMERAM	VME	DRAM	2M to 16M	210	180	470	650	—	EDC, 64-bit cache
VSBRAM-EC1	VME/VSB	DRAM	2M to 64M	235	150	265	420	—	EDC, dual 64-bit caches
VMERAM-FP1	VME	DRAM	2M to 64M	138	89	181	181	—	fast parity memory
Compcontrol 15466 Los Gatos Blvd, Ste 109-365, Los Gatos, CA 95032 (408) 356-3817									**Circle 304**
CC109	VME/VSB	DRAM	16M to 64M	VME 190 VSB 240	VME 70 VSB 125	VME 240 VSB 275	VME 130 VSB 205	$3,750	—
CC81	VME	DRAM	4M to 8M	230	150	260	260	$1,095	—
CC86	VME	EPROM, EEPROM, ROM, SRAM	1M to 2M	100 to 350	100 to 350	130 to 380	130 to 380	$1,035	—
CC87	VME	DRAM	2M	290	200	320	320	$999	—
Datacube 4 Dearborn Rd, Peabody, MA 01960 (508) 535-6644									**Circle 305**
ROI-STORE	VME	PROM	2M	850	300	225	300	$3,750	part of MaxVideo family
MEGASTORE-8	VME	PROM	8M	850	300	225	300	$11,950	part of MaxVideo family
FRAMESTORE	VME	PROM	832k	600	600	—	—	$3,750	—
Duual Computer 26046 Eden Landing Rd, Ste 2, Hayward, CA 94545 (415) 785-8890									**Circle 306**
VMEM/4	VME/VSB	DRAM	4M	VME 230 VSB 220	VME 180 VSB 165	VME 275 VSB 260	VME 230 VSB 195	$2,150	parity checking, dual-port
VMEM/16	VME/VSB	DRAM	16M	VME 230 VSB 220	VME 180 VSB 165	VME 275 VSB 260	VME 230 VSB 195	$3,950	parity checking, dual-port

Model	Bus	Memory type	Capacity (bytes)	Access time read (ns)	Access time write (ns)	Cycle time read (ns)	Cycle time write (ns)	Price	Comments
DY-4 Systems 21 Fitzgerald Rd, Nepean, Ontario K2H 9J4 (613) 596-9911									**Circle 307**
DMV-536/ SVME-537	VME	EPROM	8M	200	200	—	—	—	mil-spec and conduction-cooled versions, EPROM programming
SVME-335	VME/VSB	DRAM	4M	300	300	—	—	—	EDC, mil-spec versions
SVME-530	VME/VSB	EPROM	8M	220	220	—	—	—	mil-spec versions, on-board EPROM programming
DMV-540/ SVME-541	VME	flash	8M	290	290	—	—	—	mil-spec and conduction-cooled, parallel I/O
Force Computers 3165 Winchester Blvd, Campbell, CA 95008-6557 (408) 370-6300									**Circle 308**
DRAM-8C	VME	DRAM	8M	230	80	320	320	$995	—
DRAM-8D	VME	DRAM	16M	230	80	320	320	$2,690	—
DRAM-8E	VME	DRAM	32M	230	80	320	320	$4,990	—
SRAM-6	VME	SRAM	2M	55	55	105	105	$4,990	battery backup or +5V standby
RR-2	VME	EPROM, SRAM, ROM	up to 8M	—	—	—	—	$995	dual-port VMEbus/VMXbus
General Micro Systems 4740 Brooks St, Montclair, CA 91763 (714) 625-5475									**Circle 309**
GMSV05	VME	DRAM	32M to 64M	—	—	—	—	—	dual-port, CPU expansion bus compatible, VME64, EDAC
GMS SRAM	VME	PROM, SRAM	256k to 1M	30	30	—	—	$995	dual-port, CPU expansion bus compatible
GMS DRAM	VME	DRAM	4M to 16M	100	100	160	160	$795	dual-port, CPU expansion bus compatible
GreenSpring Computers 1204 O'Brien Dr, Menlo Park, CA 94025 (415) 327-1200									**Circle 310**
DRAM3-1M	VME	DRAM	1M	280	190	320	280	$890	3U, with 6U optional
DRAM2-2M	VME	DRAM	2M	240	90	275	235	$1,295	6U
DRAM16	VME	DRAM	4M to 16M	280	90	310	240	$1,400 to $2,950	—
NVRAM4	VME	DRAM	4M	240	90	275	235	—	battery backup, 3U size
Matrix 1203 New Hope Rd, Raleigh, NC 27610 (919) 231-8000									**Circle 311**
MS-RPN Series	VME	PROM, RAM	384k to 1.5M	180	180	185	185	$245	automatic boundary generation
MS-RPC Series	VME	PROM, RAM	384k to 1.5M	180	180	185	185	$295	automatic boundary generation
MS-RPB Series	VME	PROM, RAM	384k to 1.5M	180	180	185	185	$275	automatic boundary generation, RAM battery backup
MS-DRM Series	VME	DRAM	up to 1M	170	170	285	285	$295	parity on/off, odd/even, unified global refresh
MD-RAM	VME	EEPROM, SRAM	up to 4M	275	275	515	515	$895	low-power CMOS, battery backup and status indicator
MX-RAM	VME	EEPROM, SRAM	up to 4M	275	275	515	515	$1,790	low-power CMOS, battery backup and status indicator
MR-RAM	VME	SRAM	up to 4M	275	275	515	515	$2,995	same as MD-RAM but with rugged design

Model	Bus	Memory type	Capacity (bytes)	Access time read (ns)	Access time write (ns)	Cycle time read (ns)	Cycle time write (ns)	Price	Comments
Micro Industries 691 Greencrest Dr, Westerville, OH 43081 (614) 895-0404									**Circle 312**
PG220X	VME	DRAM	0.25M to 8M	250	200	265	195	$650 to $950	protection mechanism permitting supervisor access only
PG221X	VME	DRAM	0.25M to 1M	—	—	300	300	$1,700 to $2,400	dual-port memory, ECC
PG222X	VME	DRAM	4M to 16M	—	—	200	100	$5,800 to $8,400	parity, dual-port with single-port select
PG2260	VME	EPROM, EEPROM, ROM, SRAM	1.25M	—	—	—	—	$950	universal memory board
PG2255	VME	SRAM	1M	—	—	—	—	$2,700	—
PG2206X	VME	DRAM	2M to 8M	230 to 290	150 to 200	260 to 320	260 to 320	$1,950 to $5,800	parity generation and checking
MBLC116	Multibus I	DRAM	16k	—	—	665	665	$695	48 I/O lines, battery backup, sync/async serial channel
MBLC300	Multibus I	DRAM	32k	1.0 μs	1.2 μs	—	—	$495	plug-in memory for BLC 86/123
MBLC304	Multibus I	DRAM	128k	—	—	—	—	$595	plug-in memory for BLC 86/30
MBLC416	Multibus I	PROM, ROM	16k	—	—	—	—	$495	on-board programming
MBLC0512A	Multibus I	DRAM	512k	275	110	400	400	$995	parity on-board refresh and control
MBLC0512B	Multibus I	DRAM, RAM	512k	350	400	500	500	$1,495	ECC, on-board refresh and control
MSBC012EX	Multibus I	RAM	512k	375	375	625	625	$595	on-board parity generator/checker
MSBC040EX	Multibus I	RAM	4M	375	375	625	625	$1,995	—
Mizar 1419 Dunn Dr, Carrollton, TX 75006 (214) 446-2664									**Circle 313**
MZ7210	VME	DRAM	up to 4M	170	55	225	225	—	—
MZ75708	VME	DRAM	up to 16M	—	—	—	—	—	—
EMX7230	VME	DRAM	up to 8M	120	120	240	240	—	—
EMX7240	VME	DRAM	up to 16M	—	—	—	—	—	—
MZ8205	VME	ROM, SRAM	64k to 256k	—	—	—	—	—	—
MZ8210	VME	DRAM	512k	—	—	—	—	—	—
MZ8215	VME	DRAM	up to 4M	170	55	225	225	—	—
MZ8315	VME	SRAM	up to 32k	—	—	—	—	—	—
Motorola Microcomputer Div 2900 S Diablo Way, Tempe, AZ 85282 (602) 438-3013									**Circle 314**
MVME216	VME	EPROM, EEPROM, SRAM	1M to 16M	170	190	230	230	—	supports block transfer on VME
MVME246-2	VME	DRAM	32M	170	60	200	200	—	—
MVME246-3	VME	DRAM	64M	170	60	200	200	—	—
MVME236-2	VME	DRAM	8M	170	60	200	200	—	—
MVME236-3	VME	DRAM	16M	170	60	200	200	—	—
MVME224a-1	VME/VSB	DRAM	4M	190	65	230	200	—	dual-port
MVME230-2	VME	DRAM	8M	220	65	360	360	—	ECC
MVME230-3	VME	DRAM	16M	220	65	360	360	—	same as above
MVME236-1	VME	DRAM	4M	170	60	200	200	—	—
MVME224a-2	VME/VSB	DRAM	8M	190	65	230	200	—	dual-port
MVME224a-3	VME/VSB	DRAM	16M	190	65	230	200	—	same as above

Model	Bus	Memory type	Capacity	Access time read (ns)	Access time rwrite (ns)	Cycle time read (ns)	Cycle time rwrite (ns)	Price	Comments
Motorola Microcomputer Div 2900 S Diablo Way, Tempe, AZ 85282 (602) 438-3013									**Circle 314**
MVME224a-4	VME/VSB	DRAM	32M	190	65	230	200	—	dual-port
MVME230-1	VME	DRAM	4M	220	65	360	360	—	ECC
MVME-214	VME/VSB	ROM, SRAM	1M	—	—	—	—	—	dual-port
MVME215-1	VME	CMOS, RAM	256k	225	225	230	230	—	Nicad battery supplied
MVME215-2	VME	CMOS, RAM	512k	225	225	230	230	—	same as above
MVME215-3	VME	CMOS, RAM	1M	225	225	230	230	—	same as above
Performance Technologies 435 W Commercial St, E Rochester, NY 14445 (716) 586-6727									**Circle 315**
PT-VME240	VME/VME64	DRAM	4M to 32M	43	43	108	83	$1,626	includes VME64 capability
PT-VME201	VME	DRAM	4M to 8M	290	140	340	260	$1,226	block-mode capability
PT-VME200	VME	DRAM	4M to 8M	290	140	340	260	$1,800	VME/VMX dual-port, block-mode capability
Radstone Technology 20 Craig Rd, Montvale, NJ 07645 (201) 391-2700									**Circle 316**
8SB	VME/VSB	DRAM	8M	100	100	200	210	—	—
8EP	VME	DRAM	8M	110	140	150	180	—	—
16EP	VME	DRAM	4M to 16M	220	120	300	320	—	—
SME-1A	VME/VSB	EPROM, EEPROM, SRAM	512k to 2M	280	280	320	320	—	mil-spec or commercial equivalent
DME-1	VME/VSB	DRAM	4M to 16M	125	130	225	230	—	same as above
SBE 2400 Bisso Ln, Concord, CA 94520 (415) 680-7722									**Circle 317**
VSRAM-2	VME/VSB	SRAM	2M	VME 71 VSB 63	VME 71 VSB 63	VME 81 VSB 72	VME 81 VSB 72	$4,750 (100)	jumper-selectable configuration, interleaved port access option

To avoid such drops, SBE uses 2 Mbytes of 25-ns SRAM on its VSRAM-2 memory board. Although more expensive than DRAMs, SRAMs don't require a precharge time or a row address strobe cycle. This allows sustained transfers even when both ports are accessing at full speed. Designed to meet the requirements of an application that would support interleave accesses to the data and the dual-port memory at high speeds, the VSRAM-2 is able to sustain transfers on both the VSB and the VMEbus in excess of 20 Mbytes/s without any interference between the two buses. The board can sustain these high data rates concurrently, resulting in an effective combined data transfer rate of 40 Mbytes/s.

While dual-port memory designs help in achieving full system performance, there are some advantages to using a local custom port—adding a third path for data transfers. As the VMEbus and VSB become saturated with data from external sources and interprocessor communication, the system CPUs barrage the arbiter with bus requests and data traffic. This can force the system designer into unwanted compromises in data transfer efficiency. Allowing direct access to the CPU, local buses can move data without the restrictions of system bus arbitration and clock speeds.

Taking this triple-port approach, the GMSV05 memory system from General Micro Systems (Montclair, CA) features a custom Special Application Module (SAM) bus. Operating through a DIN connector, the SAM-bus is basically a 68030 line that brings address, data, cache-fill and MMU lines from a GMSVX7 CPU up to the memory.

"The idea here is that an I/O board on the bus can be writing to one of these memory cards while a CPU is reading from another one," says Del Miller, director of marketing at General Micro Systems. "What we're trying to do is emulate a flow-through architecture in a batch-processing environment."

A GMSV05, for example, used in a two-board set with a GMSVX7 CPU provides 64 Mbytes of local memory to the CPU without going through the bus. This configuration also provides an additional VME and VSB

interface. Stackable connectors allow up to four GMSV05 modules addressable from either the VMEbus or the SAMbus. Supporting the use of burst-mode transfers over the SAMbus, the GMSV05 permits 4-1-1-1 bursts at CPU clock speeds up to 33 MHz.

Another rare feature of the GMSV05 is its implementation of VME64. The design allows 64-bit data transfers per the VME Rev. D specification. Data is transferred at twice the rate as block transfers over a 64-bit data path.

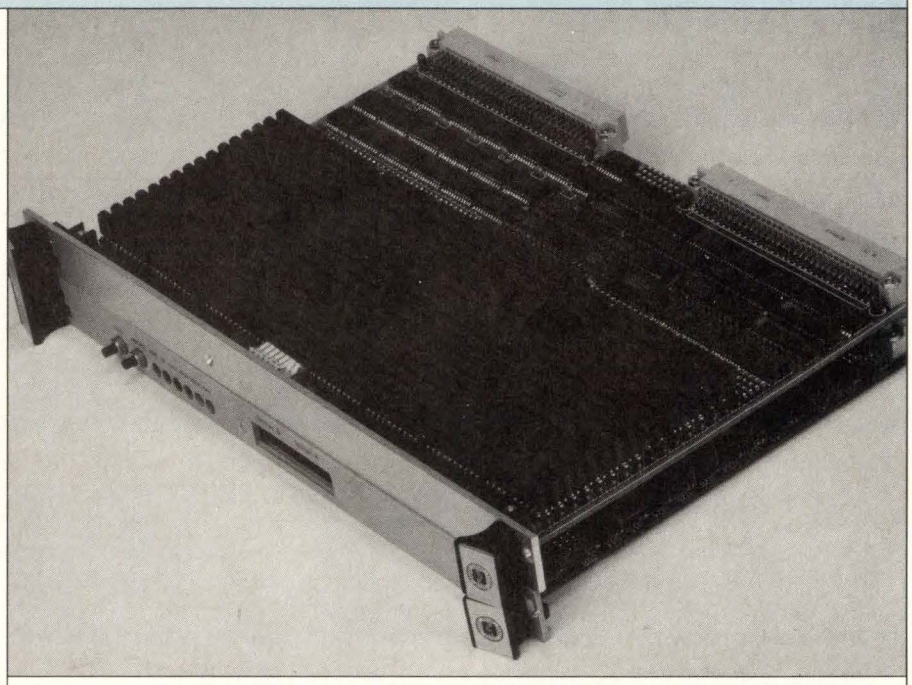

Accommodating 32 or 64 Mbytes of two-way interleaved, static column DRAM, the GMSV05 from General Micro Systems features triple-port design. This design allows data transfers along the VMEbus, the VME subsystem bus and a local Special Application Module bus. Used in a two-board set with a GMSVX7 CPU, the GMSV05 provides 64 Mbytes of local memory to the CPU without going through the bus.

Memory capacity crucial

While high data transfer rates are an important factor in memory board design, memory capacity can't be ignored—especially as more and more memory resides on the CPU boards themselves. Making decisions about which memory chips to use may be an afterthought in CPU or controller board designs, but it is a key factor for memory boards. To stay competitive, whether using DRAMs, SRAMs, EPROMs or a combination of memory types, memory board manufacturers need to understand the importance of keeping tabs on the densest, fastest memory devices available.

"The most difficult problem to deal with in the memory board business is keeping up with a highly volatile memory market," says Michael Curran, general manager at Micro Industries (Westerville, OH). "Since memory chips are a commodity, memory boards also tend to sell as a commodity. So the ability to sell memory boards is directly driven by the costs of those devices."

Accommodating up to 16 Mbytes of DRAM, Micro Industries' PG222X can run as either a dual-port or a single-port memory board. In dual-port operation, the board offers read cycle times from 210 ns to 270 ns and write cycle times from 110 ns to 179 ns. For faster access speeds, the board can run as a single-port memory device with 200-ns read cycles and 100-ns write cycles.

Recognizing the importance of using the fastest parts that are readily available, designers at Clearpoint Research (Hopkinton, MA) selected 70-ns DRAMs for their VMERAM-FP1 memory board. Careful not to exceed the maximum transfer rates inherent in the VMEbus, they also incorporated the latest 7.5-ns programmable memory components and tied in TTL logic to keep operation at 100 ns.

Available with up to 64 Mbytes of DRAM, the VMERAM-FP1 features a 181-ns cycle time for both reads and writes. The 6U single-slot card also offers parity error detection. It has a Motorola-compatible control status register and allows users to access parity error status information and control. In block mode the VMERAM-FP1 runs at a typical speed of 89 ns, which translates to a data rate of 35 Mbytes/s. Clearpoint will incorporate 60-ns DRAMs when they become available. ■

Mitsubishi Fast SRAMs

When You Need Fast Cache.

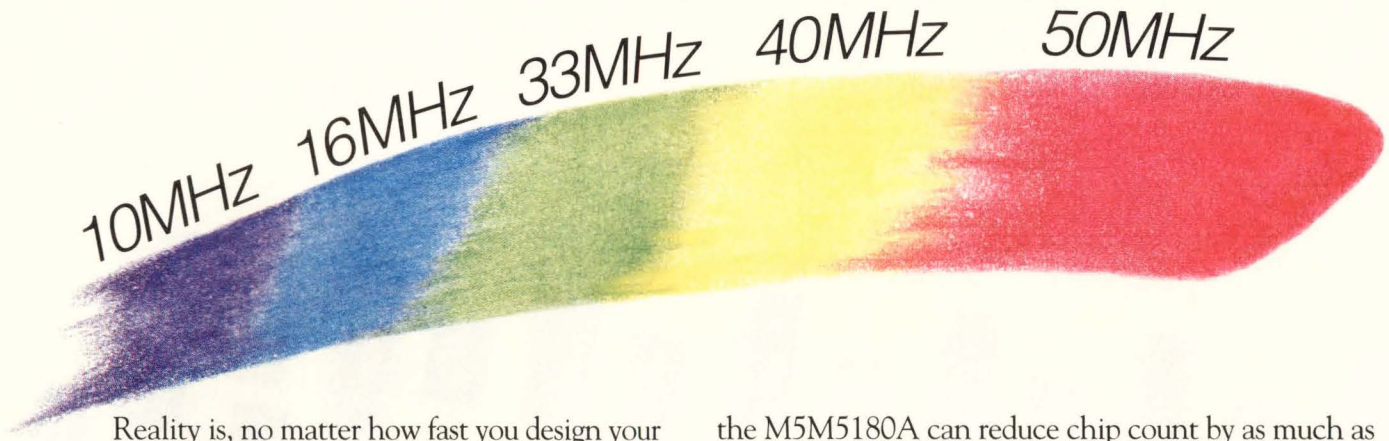

Reality is, no matter how fast you design your systems, customers always want them faster. To stay ahead, you need fast access to ultra-fast SRAM access times. And, Mitsubishi helps you lead the cache memory race with fast, 15ns Static RAMs in 28-pin SOJ, DIP and flat packages.

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.

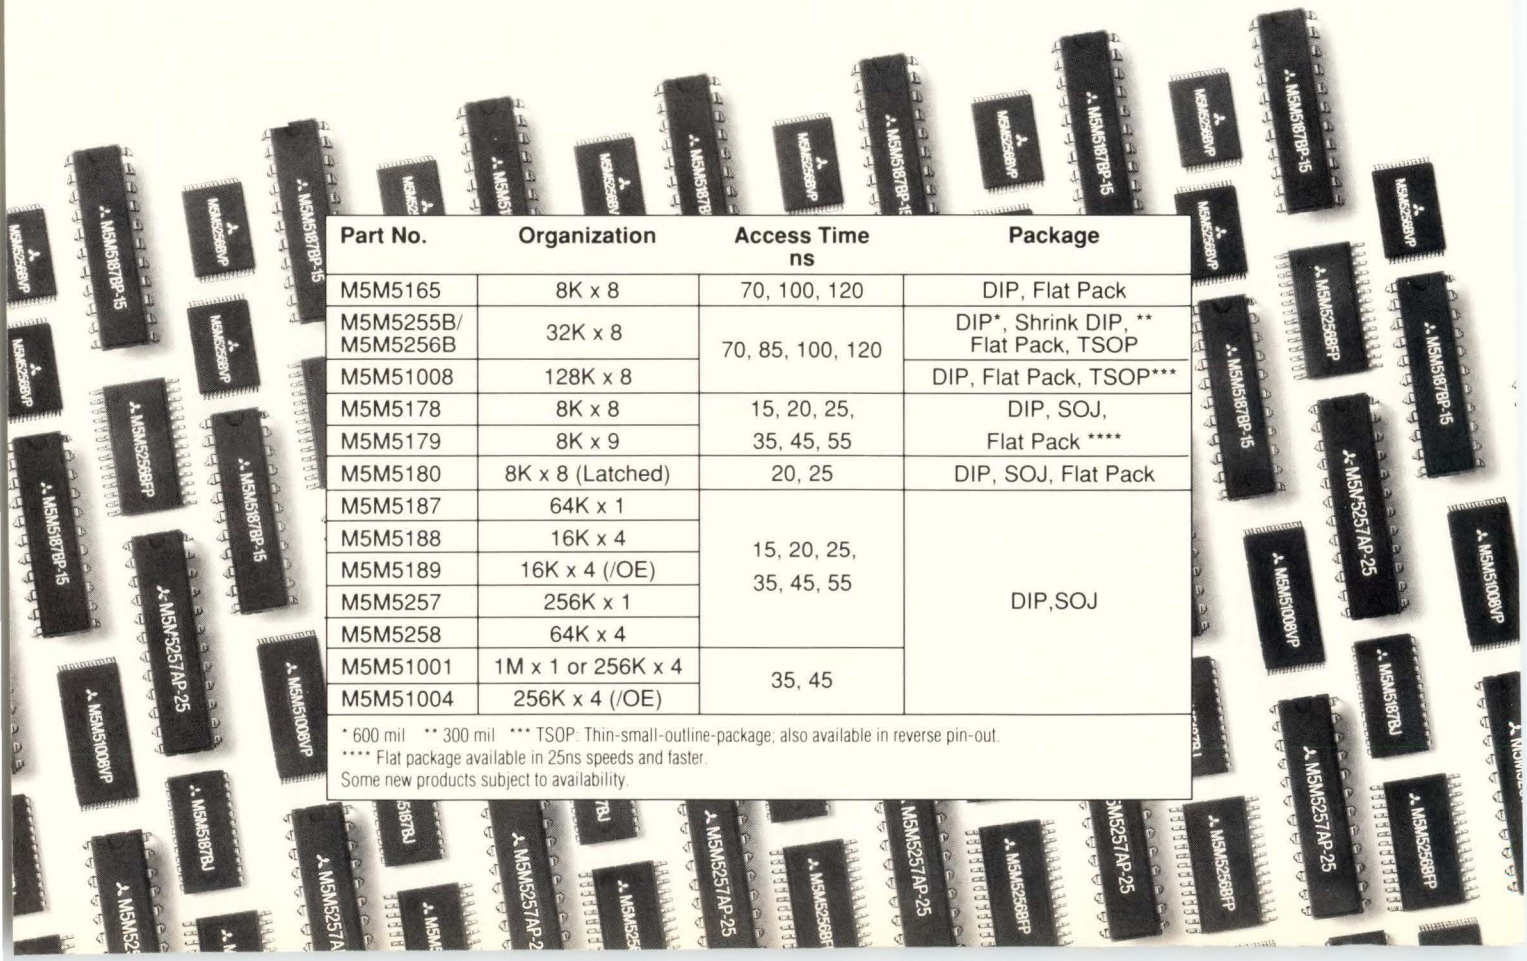

Part No.	Organization	Access Time ns	Package
M5M5165	8K x 8	70, 100, 120	DIP, Flat Pack
M5M5255B/ M5M5256B	32K x 8	70, 85, 100, 120	DIP*, Shrink DIP, ** Flat Pack, TSOP
M5M51008	128K x 8		DIP, Flat Pack, TSOP***
M5M5178	8K x 8	15, 20, 25,	DIP, SOJ,
M5M5179	8K x 9	35, 45, 55	Flat Pack ****
M5M5180	8K 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. 71

SOP

TSOP

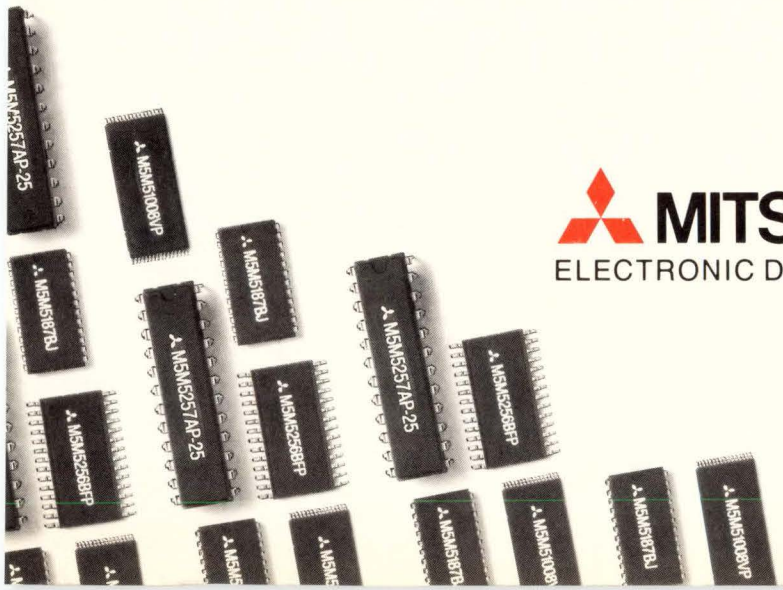

MITSUBISHI
ELECTRONIC DEVICE GROUP

INTEGRATED CIRCUITS

Motorola's 16-bit part offers surprises, resolves a few mysteries

After months of hints and pointed remarks about the lack of an upgrade path from 8- to 16-bit microcontrollers, Motorola has finally rolled out its new 16-bit part—the 68HC16. As expected, the device has a 16-bit ALU, a 16-bit data bus and software compatibility with the company's widely admired 68HC11 8-bit microcontroller.

There are both surprises in the new device and resolutions to a few mysteries. Among the surprises are some special-purpose additions to the register architecture. Among the resolved mysteries are just what Motorola meant by software compatibility, and which of the company's three silicon architectures actually ended up in the device.

Motorola settled on the Inter-Module Bus (IMB). Although it cost silicon area compared with the company's simpler HC11 layout, the IMB had important strategic advantages, according to vice-president and general manager R. Gary Daniels. "The IMB lets us use the sophisticated peripheral modules we've developed in the 68300 family, as well as simpler HC11 modules. That lets the HC16 serve as a bridge between the two product families."

A 16-bit CPU at the core

The most important module on the IMB is a new 16-bit execution unit, the CPU16. The core seems to be just an expanded version of the 8-bit HC11 core. The register layout is indeed similar, but in fact the CPU16 is a trimmed-down, sped-up version of the CPU32 core in the 68300 family, a fact that becomes evident from statistics such as a 600-ns multiply and several two-clock instructions.

There are two main 16-bit accumulators in place of the one pair of concatenated 8-bit ones in the HC11. The number of index registers has increased, and the index registers have all sprouted 4-bit bank registers to support the part's 1-Mbyte sort-of-linear address space.

More interesting to some will be the addition of rather specialized hardware: a pair of 20-bit multiply-accumulator (MAC) registers. The limited MAC capability is offered in recognition of the fact that many compute-intensive control loops are actually doing signal-processing algorithms, and would benefit from DSP hardware. Motorola is quick to point out, though, that the part isn't intended as a general-purpose DSP chip—the HC16 is outrun on most algorithms by TI's low-end 320C10.

Another important feature of the CPU16 is vastly improved interrupt latency. Motorola systems designer Mark Heene says, "The new CPU uses an interrupt controller based on the 68332 hardware. That improves latency from about 26 μs on the older HC11s to 3 μs on the HC16."

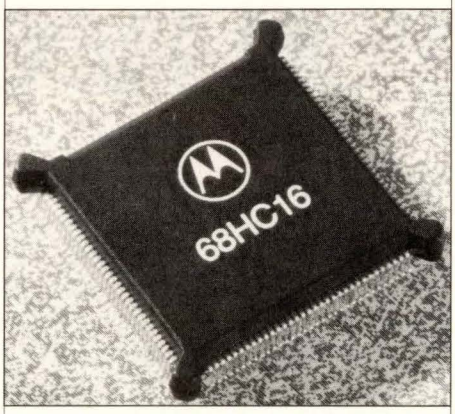

In the initial version of the new architecture, the MC68HC16Z1, the IMB will connect the CPU to a queued serial module, a general-purpose timer module and 1 kbyte of separately powered SRAM—all familiar components to 68300 enthusiasts. In addition, the part sports an eight-channel, 10-bit analog-to-digital converter.

As in other IMB-based designs, the silicon bus isn't brought directly off the chip. Instead, another new module, an HC16 system integration module, creates a demultiplexed external bus. With 20 address and 16 data lines, the bus runs at 16.67 MHz and provides the instruction stream for the ROMless Z1. The system integration module also provides the same sort of timer and chip-select features familiar from 68300 devices.

That familiarity makes up an important part of Motorola's migration strategy. The fact that the HC16 family will have access to both HC11-type and 68300-type peripheral modules means that much time-critical I/O code running on other Motorola microcontrollers won't have to be changed for the HC16.

Compatibility and familiarity

Motorola has attempted to salvage as much as possible of the rest of HC11 users' code, too. Nominally, the HC16 is assembly source-compatible with the HC11. That means that if you take a sequence of HC11 instructions and reassemble them, you will get a sequence of valid HC16 instructions, most of which will do what you expected.

Naturally, though, there will be exceptions. "The stack frame and the interrupt structure are different, for instance," explains Heene. "And there are a few cases where the assembler will replace an old HC11 instruction with a new more-flexible HC16 one. But the compatibility means that the source code will be familiar, and the old program will give you a good place to start."

Customers looking for a migration path to 32-bit will find an answer as well. With the C compiler Motorola will provide with the HC16, there will be a C-source migration path to the 68300 family, again taking advantage of common peripheral modules. So in principle, a design team could create a pool of C and assembler code that would be portable all the way from the HC11 to the 68332.

This path forms the most important asset of the new architecture. The HC16 may not be the fastest 16-bit MCU on many operations. It won't necessarily be the cheapest solution in town, either, with the ROMless Z1 costing $25 (1,000s). Nor will the part beat anyone to market, since it will not begin general sampling until the second quarter of 1991. But for an application that already has a great investment in HC11 code, or has already identified the 68300 family as its high-end solution, the HC16 is an important new alternative. —*Ron Wilson*

Motorola
6501 William Cannon Dr W
Austin, TX 78735-8598
(602) 962-2085

Circle 359

SOFTWARE

Tool models and simulates real-time O/S architecture

Ready Systems has introduced a software development tool for modeling the structure and behavior of real-time systems architecture. VRTX Designer uses graphical icons to model the software architecture of embedded systems built around the VRTX32 real-time executive. Developers can build an architecture model using icons to represent familiar real-time entities such as tasks, device drivers, mailboxes queues and semaphores. These graphic entities are connected via lines representing VRTX system calls such as SC_POST (place a message in a mailbox) or SC_RESUME (resume a suspended task).

The relatively small number of icons can be grouped together under a special icon to represent a subsystem. Thus the designer can step through a hierarchy from system overview to greater detail. The result is a representation on the screen of the concurrent programming entities and their intertask communications and synchronization, as well as VRTX scheduler mechanisms. This representation describes how the application code will interact with VRTX services. The actual code that users must write is represented by the various software building-block icons, which can be named to correspond to the task code to be written.

Supporting simulation

But VRTX Designer goes beyond a mere graphical representation. It also supports behavior modeling and simulation. Before being committed to actual code, the tasks' behavior can be modeled in a pseudocode that describes the control statements, loops and branches; the probability that each branch will execute; the average cycles for loops; the timing requirements and definition of the sequence; and the execution of system calls. The system does design-rule checking during creation of the graphic model so users can't inappropriately choose, for instance, to send a clock signal to a mailbox.

Users can opt to run a simulation of the entire system under design, but can also select any path desired and examine it in detail. A path is selected as a group of entities: tasks, devices or drivers connected by communication and synchronization components. Usually a path is selected as a group of entities that must achieve a response to a given stimulus in a given time. Paths don't need to be linear, and users can select several paths in a path set to view and compare their behavior in a single simulation run.

Since real-time systems typically interact with stimuli from the outside world, users can define stimuli that represent external hardware devices. External stimuli can be defined as periodic, such as a sampling rate from some sensor, or as distributed over some statistical time frame. In the latter case, a distribution curve might describe the probability of events occurring in either random, uniform or exponential distribution.

For simulation, VRTX Designer uses information about the specifics of a target hardware environment. Versions currently support the Motorola 680X0 and the Intel 80X86 families of microprocessors. VRTX Designer contains timing behavior information about the VRTX32 scheduler, which it uses along with the task timing information supplied by users in building their task timing models. Simulations can be run to test a design at different clock speeds. This is especially useful in determining that a slower, hence less expensive, CPU will do the job and reduce final product cost.

Users can set up a simulation by selecting the paths they wish to examine; the microprocessor they're using and its clock speed; and the number of cycles they want the simulation to run. Once the simulation is completed, the results are displayed as a time-line chart, with the individual paths shown one below the other. A dark horizontal bar shows when a given task has executed, and a CPU utilization summary is also displayed (for example, wait time, idle time, operating system overhead and application time). Users can zoom in on the output graph to improve the resolution and determine whether tasks are locked out, being starved or deadlocked—which often aren't discovered until the debug stage.

VRTX Designer is available now for Sun Microsystems Sun-3 and -4 platforms. Single-user prices start at $12,500, with a special price of $9,950 before Dec. 31. Discounts for volume or bundled purchases are also available.

—Tom Williams

VRTX-BASED APPLICATION ARCHITECTURE

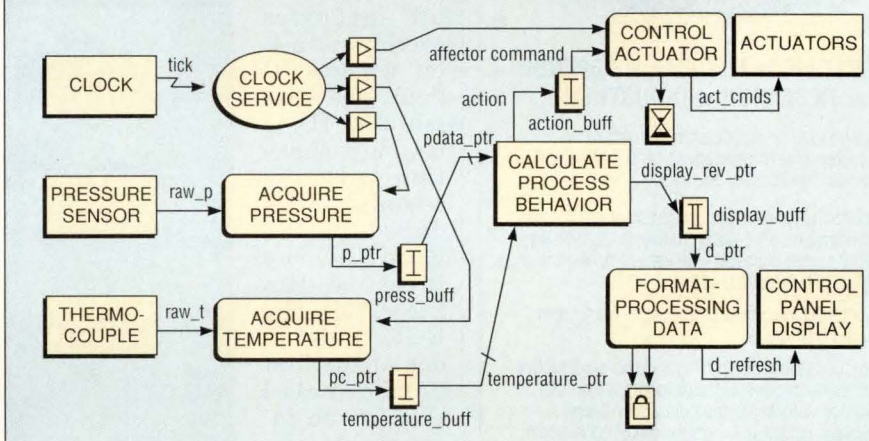

Ready Systems
470 Potrero Ave
Sunnyvale, CA 94086
(800) 228-1249

Circle 352

SOFTWARE

3-D EISA board aims at workstation graphics on 386/486 platforms

The Matrox EG3-1280 graphics processor board is designed to turn 80386- and 80486-based computers built on the EISA bus into 3-D workstations. The EG3-1280 combines Matrox's proprietary 2-D graphics engine with Texas Instruments' TMS320C30 digital signal processor to output 130,000 3-D vectors/s and 20,000 Gouraud-shaded polygons/s. By itself, the 2-D engine can output over 250,000 transformed vectors/s.

The EG3-1280 supports 1,280×1,024×24-bit-deep pixel display resolution, with a 16-bit depth buffer for 3-D. The total 32-bit depth of the display buffer includes the 24-bit pixel data, a 4-bit graphics overlay and a 4-bit window ID field. The 4-bit window ID is used in conjunction with the recently introduced Bt-463 RAMDAC by Brooktree (San Diego, CA). The EG3-1280 is one of the first commercial products to incorporate the Bt-463, which allows various applications to coexist in different windows on the screen. Users can run applications simultaneously with 8-bit pseudo-color, 24-bit true color and 12-bit dithered color along with 4-bit graphics overlays.

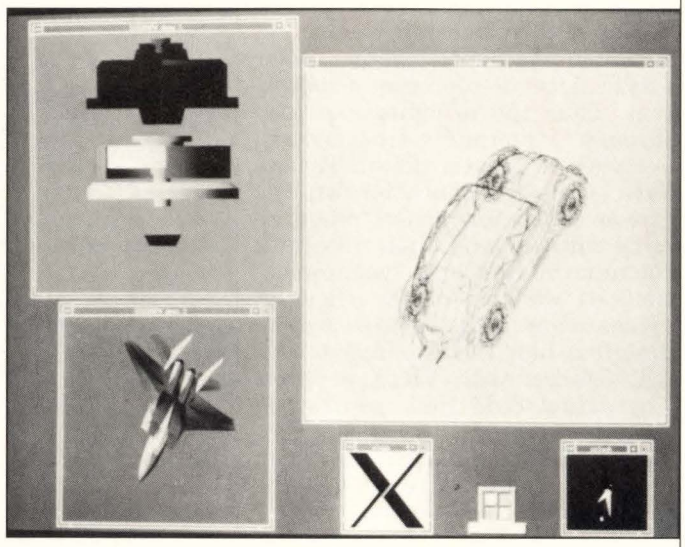

The EG3-1280 is designed to run under Unix System V, Release 3 and supports X Windows compatibility and an on-board, high-level 32-bit command interface called the Scalable X Command Interface (SCXI). SCXI lets users develop custom code by providing 100 2-D and 3-D drawing and control primitives. The code is downloaded to the board and accepts commands and parameters from applications.

In addition, the EG3-1280 includes a PEX server. PEX is the X Windows version of the Programmers Hierarchical Interface to Graphics Systems (PHIGS) 3-D software standard. Overhead is reduced by tightly coupling the PEX server with the on-board hardware resources. Other PEX applications running over a network can seamlessly run on an EISA system using the Matrox board because of the standard interface provided.

For OEMs with custom needs, the EG3-1280 can be adapted for stereo display, noninterlaced output or some nonstandard display resolution up to 1,600×1,280 pixels. The software interfaces supplied let OEMs use standard X Windows, PEX and a VGA pass-through capability. The SCXI interface allows the design of custom graphics software without the need to include numerous layers of interface code.

Matrox is targeting the EG3-1280 for applications in CAD/CAM/CAE at performance levels that haven't been available on 386/486 platforms. Other target markets include 3-D modeling and animation.

The EG3-1280 will be available this month at a single-unit price of about $8,000. *—Tom Williams*

Matrox
1055 St. Regis Blvd
Dorval, Quebec, Canada H9P 2T4
(514) 685-2853

Circle 351

The

Graphics Subsystem

Slash the cost of 3-D workstations

A Matrox EG3-1280 graphics adapter, for EISA-based systems, delivers workstation performance at a cost that's down-to-earth. Here's what you get:

- ★ *1280 x 1024 display resolution*
- ★ *X-Windows, PEX support*
- ★ *High-end 3-D workstation speed and power*

We've expanded the universe for 3-D graphics. Our EG3-1280 puts you up there with the stars.

Workstation Information

1-800-361-4903 *or* **(514) 685-2630**

matrox

THE LEADER IN VIDEO MICROTECHNOLOGY

SOFTWARE

Debugging system operates through ROM socket

The ROMport, a communication and debugging system from Embedded Support Tools, provides an alternative to in-circuit emulators while solving the speed and installation hurdles found in typical monitor debuggers. The system consists of the ROMport hardware device bundled with EST-Bug, the company's real-time monitor debugger. The device plugs directly into the 27xxx ROM socket of a 68000 or 8086-based system. Users can also install their own monitor debugger on the ROMport.

Although in-circuit emulators for the latest 16- and 32-bit microprocessors are more expensive than the ROMport, the less costly monitor debuggers also have drawbacks. They usually require their own serial port, are traditionally difficult to install and are generally restricted to slow serial communications.

ROMport tackles these problems by adding a serial port to the system through the ROM socket. When installed, ROMport provides developers with an added communications port for high-speed downloading and real-time debugging of their target software. Through a phone-style connector, the device provides a bidirectional serial connection to any host system or dumb terminal.

ROMport uses its own integrated serial driver with a simple auto-configuration utility to ease installation of the ROMport system. The driver program also permits designers to download software applications into RAM at speeds of up to 115.2 kbit/s. This can cut download time by a factor of six over standard serial communication, says EST.

Installed directly onto ROMport, EST-Bug provides a full-featured monitor debugging system in firmware. This system lets developers set conditional breakpoints on RAM or ROM code. Users can view a complete code and data trace, single-step the processor, display and alter memory or registers, set data breakpoints, disassemble memory and run menu-driven diagnostics.

The ROMPORT and EST-Bug system supports Motorola's 68000/010/020/030, 68302 and 68332, and Intel's 8086 and 80186. Debugging systems for other targets will be released soon, says EST. Available now, the ROMport and EST-Bug bundled package starts at $1,095. ROMport is priced at $495. —*Jeffrey Child*

Embedded Support Tools
10 Elmwood St
Canton, MA 02021
(617) 828-5588

Circle 358

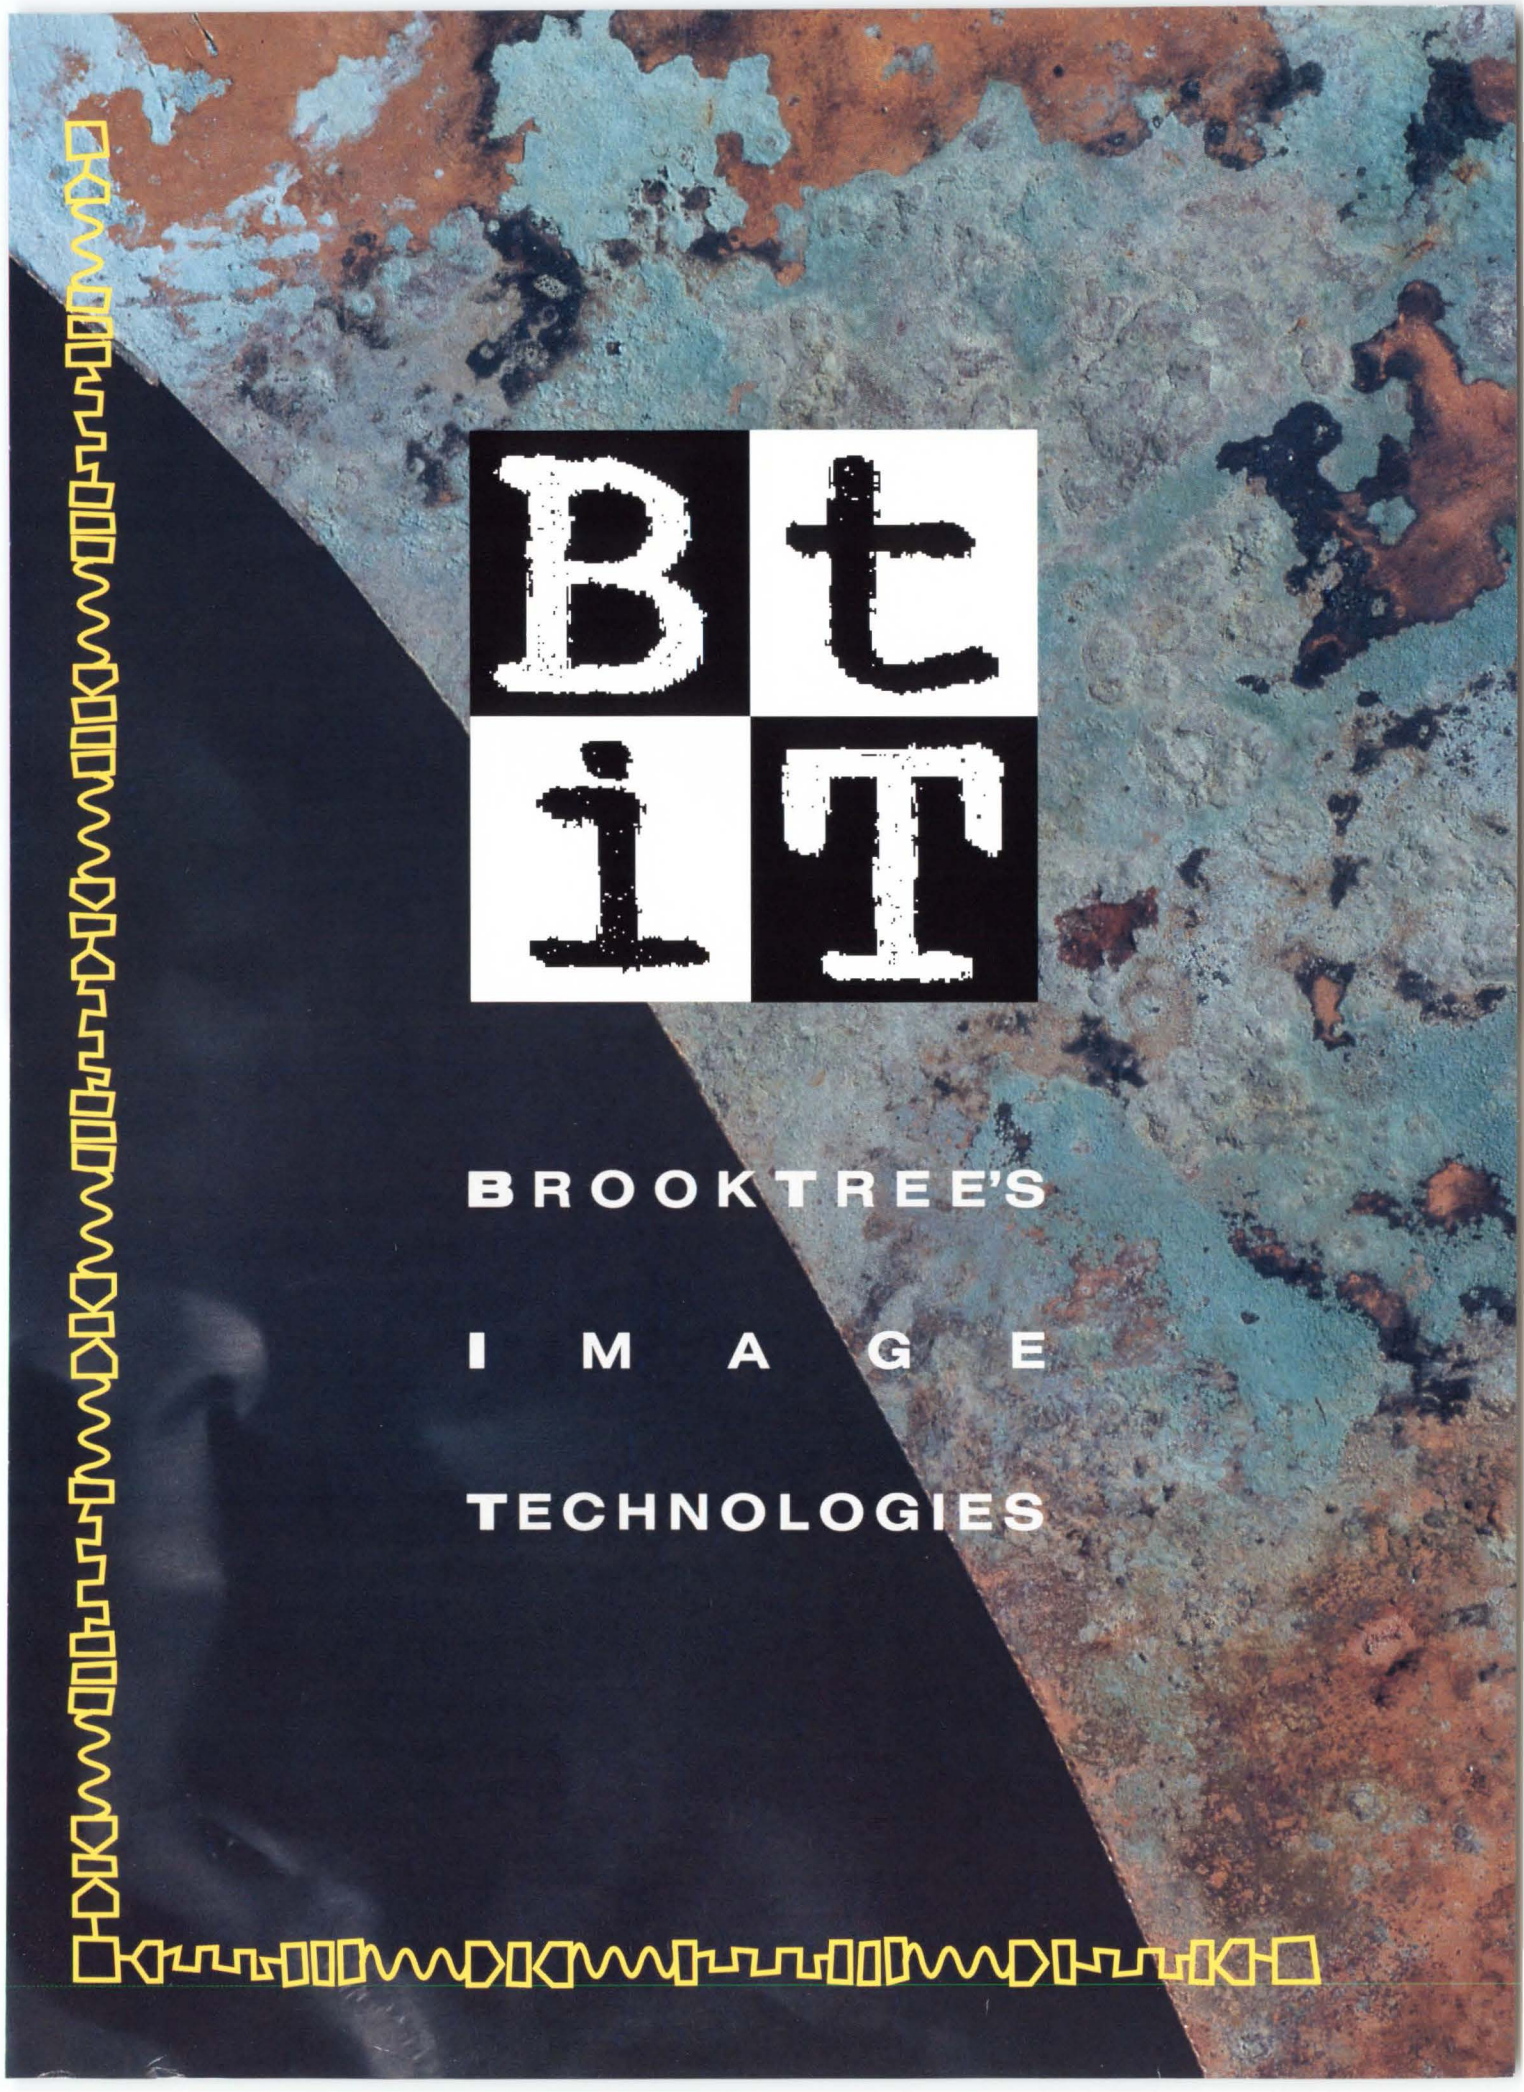
Bt
it
BROOKTREE'S
IMAGE
TECHNOLOGIES

Capture iT

Change iT

Capture an image, manipulate it, transmit it, display it—and do it all with Brooktree's Image Technologies. They'll give your system a visual edge in competitive markets.

In fact, Image Technologies can be the key differentiator in today's look-a-like world. At Brooktree, we're dedicated to creating the highly integrated devices designers need to set their systems apart with exciting imaging and graphics capabilities.

IMAGE ACQUISITION

Image Acquisition starts with our Bt251 Gray-Scale or Bt253 Color Image Digitizer chips. They're the easy, economical way to add image capture to your system, with flexible architectures and standard MPU interfaces.

And they link up to our new Bt261 30 MHz Line Lock Controller. It will change the way you bring video images into your system. It's flexible and fully programmable.

Program it to strip horizontal and vertical sync information from any incoming video signal it encounters. Program its sync noise gating to cope with noisy signals and to enable locking to horizontal sync.

In fact, you won't find a more flexible solution to the timing control section of your design. Or a better way to assure your Image Acquisition system is programmed for success.

IMAGE MANIPULATION

Once you capture an image, change it.

That's easy with our Bt281 real-time Color Space Converter/ Color Corrector chip. It's programmable and lets you convert from any color space to any other, including YIQ, YUV, RGB and Y, R-Y, B-Y, while capturing or displaying the image.

So now you can optimize the color space of your frame buffer for image processing independent of the video signal you're digitizing and the CRT's RGB needs. The Bt281 handles everything.

Since the Bt281 has programmable matrix coefficients and input look-up RAMs you can also use it for gamma correction, color correction or other image restoration techniques.

And if you think that's hot, you should see the Image Manipulation chips we'll be introducing this winter. Here's a hint: It will scale new heights.

IMAGE TRANSMISSION

How can you send your image from here to there? Digitally? In real time?

Enter two more pieces of Brooktree's Image Technologies story: The Bt291 and Bt294 VideoNet™ point-to-point Video Interfaces.

Send iT

Display iT

Simply put, the Bt291 and Bt294 let you ship and receive live color digital video using an 8-bit interface.

Which means you can replace about a square foot of board real estate with two highly integrated devices. And take the rest of the week off.

The two devices have, respectively, input or output look-up table RAMs to simplify the interface to the frame buffer and to add or remove gamma correction and scale signal levels.

So if you're working with CCIR601, SMPTE RP125, EBU 3246-E or other digital video standards, we've done our parts. You take it from here.

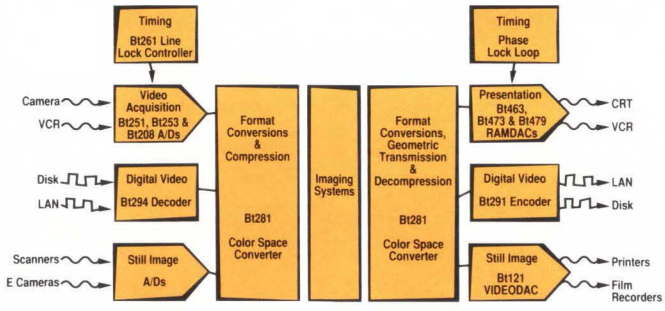

IMAGE PRESENTATION

When it comes to display technology, our true colors really shine. We invented RAMDACs. We understand the special needs of graphics systems designers. And we've never stopped innovating.

A perfect example is our Bt473, designed specifically for VGA true-color graphics. It has three 256x8 color look-up tables with 8-bit video D/A converters to support 24-bit true-color operation. And it can also support 8-bit pseudo-color, 8-bit true-color and 15-bit true-color operations. That makes it a perfect match for the Bt253 supporting the same formats.

Now our new TrueVu™ RAMDAC, the Bt463 is what's hot for designers of next-generation workstations eager to add windows capability, and delighted to do virtually everything with a single device. The Bt463 is the first monolithic true-color RAMDAC. That means it supports multiple display modes—both True Color and Pseudo Color—simultaneously. And with multiple windows, you get multiple colormaps, avoiding conflicts. Bt463 supports multiple plane depth, too, so a window can be 24, 16, 12 or 8 planes deep. And for a little frosting on the cake, it's flexible and easy to design in.

Dive into Image Technologies with Brooktree. We've got the growing family of highly integrated chips you need to handle graphics and imaging requirements in leading edge systems. Check these specs to see how easy we make it to capture, manipulate, transmit and display really neat visual stuff. Or call 1-800-VIDEO IC for complete product details.

Get With it

ACQUISITION

Bt208: 8-Bit Flash A/D Converter, 18 MSPS, External Zero and Clamp Control, On-Chip Voltage Reference, Overflow Output, No Video Amplifier Required, 28-Pin PLCC or 24-Pin DIP Package.

Bt251: 8-Bit Single Channel Image Digitizer, 18 MSPS, 4:1 Multiplexed Video Inputs, 256X8 Look-up Table RAM, MPU Adjustable Gain and Offset, Sync Detection, No Video Amplifier Required, 44-Pin PLCC Package.

Bt253: 8-Bit Triple Channel Image Digitizer, 18 MSPS, 2:1 Multiplexed Video Inputs, Output Format Logic, MPU Adjustable Gain and Offset, Sync Detection, No Video Amplifier Required, 84-Pin PLCC Package.

Bt261: HSYNC Line Lock Controller, 30 MHz Pixel Clock Generation, MPU Programmable Video Timing, Programmable Noise Gating, Generate HSYNC, Recovers VSYNC and FIELD, External VCO or High Speed Crystal Oscillator Clock Generation, 28-Pin PLCC Package.

MANIPULATION

Bt281: Color Space Converter, Three 256X8 Input Look-up Table, Programmable Matrix Coefficients, Optional Input Interpolation/Output Decimation, Standard MPU Interface, 36 MHz, 84-Pin Package.

TRANSMISSION

Bt291: RGB to CCIR 601/SMPTE RP125 Encoder, RGB Input Look-up Tables, RGB to YCrCb Conversion, Flexible Digital Filtering of YCrCb, 16-Bit YCrCb I/O Bus, Ancillary Input Port, Handles Video Timing Control, 100-Pin PLCC Package.

Bt294: YCrCb to CCIR 601/SMPTE RP125 Decoder, Handles Video Timing Recovery, Ancillary Output Port, Error Checking, 16-Bit YCrCb I/O Bus, YCrCb to RGB Output Look-up Tables, 100-Pin PLCC Package.

PRESENTATION

Bt463: TrueVu RAMDAC, 4:1, 2:1 MUX's, Switch on a Pixel Basis Between True Color and Pseudo Color of Multiple Plane Depths with Multiple Colormaps, Two 8 Plane Overlay Cursors, Variable Palette Size, Reconfigurable Pixel Port, Advanced Diagnostics including JTAG Port, 170, 135 and 110 MHz Operation, 169-Pin PGA.

Bt473: True-Color RAMDAC VGA Compatible, Compatible with Bt253 Output Formats—24-Bit, 15-Bit and 8-Bit True-Color, 6/8-Bit Pseudo-Color, Programmable Setup (0 or 7.5 IRE), Internal/External Voltage Reference, RS-343A/RS-170 Compatible Outputs, 80, 66, 50 and 35 MHz Operation, 68-Pin PLCC Package.

Brooktree Corporation
9950 Barnes Canyon Rd.
San Diego, CA 92121
1-800-VIDEO IC. FAX 1-619-452-7294

Brooktree®

DESIGN AND DEVELOPMENT TOOLS

IC tools target process-independent, top-down chip designs

Valid Logic Systems has introduced two tools, one for cell layout and one for floor planning, that implement process-independent and top-down IC design methods. Construct P.I. (Process-Independence) and Compose Architect let users architect and analyze a chip design at a high level and build a detailed physical representation of the device without concern for specific process design rules.

Construct P.I., the cell layout tool, features device-level compaction, graphical device generation and netlist-driven layout. The device-level compactor is at the heart of the tool's process-independent capabilities. It automatically spaces (shrinks or expands) the layout according to the design rules in the technology database. The compactor handles both orthogonal and encapsulated nonorthogonal geometries. It also maintains connectivity between wires and devices during compaction.

The single-technology database centralizes control over both physical layout and chip assembly tasks. It also eliminates scheduling impacts due to design rule changes. Design rules stored in the technology database can be modified when a technology advance or a foundry change necessitates a change in process. Designs can be updated from a 2-μm to a 1-μm CMOS process, for instance, with minimal manual layout work.

Generating C graphically

In addition to the device-level compactor, Construct P.I. implements process independence through the Graph-A-Cell graphical device generator. "This feature has drawn the most attention from users who previewed the product," says Dirk Wauters, Valid's director of IC marketing, "because it lets layout designers generate parametrized cell, or PCELL, programs rather than relying on a CAD programmer.

"In a nutshell," Wauters continues, "Graph-A-Cell is a graphical environment for generating C code. The designer starts by loosely sketching out a design that's independent of design rules. The main concerns during the sketch phase are putting the right material types in the design and making sure that the materials are in the proper relationship. That is, if the design calls for a metal overlap of a contact, the sketch should reflect that."

Graph-A-Cell incorporates unique scanning technology that automatically assigns layer-to-layer constraints. The designer can modify the automatically assigned constraints and assign user-defined constraints. The C code output from Graph-A-Cell can then be stored in a device library and called by the designer during layout editing to automatically generate PCELLs. These programs can also be manually modified to create complex PCELLs containing spirals and odd angles. Entire device libraries, as well as the designs containing the devices, can be automatically rolled over to new processes by changing the rules in the technology database and recompacting for design rule correctness.

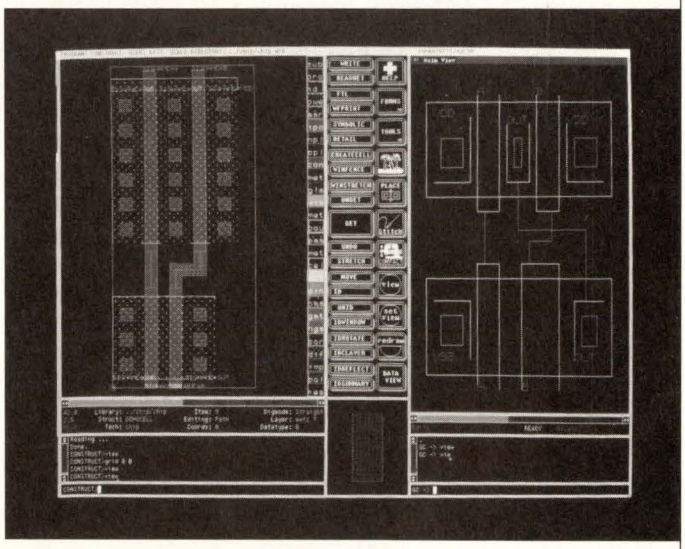

Coordinating IC design

Valid's other new tool, Compose Architect, is a netlist-driven chip floorplanning and analysis tool that lets engineers explore architectural trade-offs early in the design cycle. Based on top-down design methodology, the tool lets designers create chips using a combination of handcrafted devices, compiled modules, standard cells, synthesized layout and scaled existing layout.

At any level of design hierarchy, Compose Architect coordinates the assembly of IC design elements in varying stages of abstraction or completion. Each level of the design and each design element can be individually floor-planned and partitioned to let different design teams work simultaneously. Top-down floor planning is linked with bottom-up chip assembly and layout to minimize design iterations. Detailed placement, routing and compaction can be performed throughout the floor-planning process before all the underlying physical layout data is fully created. Compose Architect incorporates all of the design's physical layout data and automatically tracks the status of all levels and elements in the floor plan up to the end of the design cycle when the floor plan represents the finished chip design.

Both tools are available now on Sun Microsystems Sun-3 and -4 platforms with availability on systems from Digital Equipment Corp and IBM planned for early 1991. Construct P.I., with the device-level compactor, starts at $50,000. Current Construct users can upgrade to P.I. by purchasing the device-level compactor separately for $15,000. Compose Architect can be purchased as a stand-alone tool starting at $30,000 or bundled with the detailed routers of the Compose chip assembly tool for $65,000.

—Mike Donlin

Valid Logic Systems
2820 Orchard Pkwy
San Jose, CA 95134
(408) 432-9400

Circle 355

DESIGN AND DEVELOPMENT TOOLS

Analyzer puts 100-MHz RISC/CISC support on a single card

The Centurion logic analyzer from Tektronix features 100 channels (96 data, four clock), as well as 100-MHz sampling, clocking and triggering capabilities on a single card. Designed for the Tektronix DAS9200 logic analyzer mainframe, the Centurion can post asynchronous acquisitions of up to 400 MHz for 24 data channels, or for up to 384 channels with card expansions. Up to 1,536 data channels can be acquired at 100-MHz synchronous or asynchronous acquisition rates, with each channel backed by a minimum of 8 kbits of memory. A 32-kbit/channel version is also available.

The analyzer supports 32-bit RISC and CISC microprocessor analysis and can monitor multiple processors simultaneously. The system's 100-MHz performance starts with a front-end analog bandwidth of more than 200 MHz, giving the Centurion high-performance oscilloscope qualities for digital waveform fidelity.

Nichrome wire probe leads provide high bandwidth and transmission-line quality signal connections for ECL and TTL signals. These passive probes act like miniature podlets that can be clipped together in groups of eight to form a probe package that's about the size of a postage stamp. The probe's small size allows quick connection to wide buses without affecting timing performance.

Microprocessor analysis is also supported with real-time clocking from a 16-state, 100-MHz state machine. This machine is completely programmable, which allows it to emulate machine-state operation of complex buses and to place samples correctly for processor monitoring. Centurion's microprocessor support packages include complete preprogrammed setups for clocking, channel names, channel groups and symbol tables. A sample reference memory is also provided for quick familiarization with processor-specific capabilities.

The analyzer uses real-time state machine triggering, which provides full 100- MHz operation across all channels at all operating modes. Events can be recognized at the full 100-MHz analysis speed, with time left to make triggering for the next event. In addition, triggering can be preprogrammed from trigger libraries or can be user-defined with full symbolic triggering to simplify specification of trigger events.

Each acquired data point is time stamped from a 100-MHz clock. The 44-bit time stamp allows correlation of state, timing and microprocessor analysis displays with 10-ns resolution over a two-day range, without trading off memory depth or acquisition speed. Time stamping also allows the Centurion to be used with other DAS9200 card modules. Tight integration lets multiple cards work together and allows cross-triggering between modules.

The Centurion logic analyzer is available in two DAS9200 card configurations. The 92A96 has 8 kbits of memory per channel, and the 92A96D has 32 kbits per channel. Both cards are available now, with prices starting at $17,950.

—Mike Donlin

Tektronix
PO Box 12132
Portland, OR 97212
(800) 245-2036

Circle 354

CONCURRENT ENGINEERING: WESCON STYLE

November 13-15, 1990 • Anaheim Convention Center • Anaheim, Calif. USA

Gone are the days when you worked alone, developing your bit of a design. Wondering how to test this thing from the design department. Trying to figure out how to manufacture this whatsit that looked great on paper, but can't be produced.

Today the team works together. Simultaneously. In parallel. To shorten development time. And raise product quality. And increase profitability.

That's what concurrent engineering is all about.

And that's what WESCON is all about, too.

R&D. Design. Test. Manufacturing. Only WESCON covers them all. Advances in new technologies and applications. Real world solutions for real world problems. The latest products and best of proven products.

Semiconductors. Bus products. Special VME Future Bus+ Extended Architecture issues seminar and product demonstrations. RF components. Passive components. Instrumentation. Electronic design automation tools.

Production materials and services. WESCON has what you, and all of the members of your team, need.

Plan now to attend

Where World Markets Meet Wescon

For More Information call 800-877-2668.

Send me a Preview Program with full information about WESCON/90 (available August, 1990)

Name _____ Title _____

Company _____

Street Address _____

City _____ State _____ Zip Code _____

Mail to: Wescon Preview, 8110 Airport Blvd.
Los Angeles, CA 90045 or Fax to 213/641-5117

CIRCLE NO. 78

Just think ...

Just think of what the last 30 years have been like ... and what the next 10 will be ..... the next 10, when the transformation of electronics is virtually complete and the world has become digital.

Distributors and carburetors in automobiles have given way to microprocessor-controlled ignition and fuel-injection systems. By the year 2000, they'll be joined by automatic braking systems, collision avoidance systems, and navigation systems.

Radar, with its simple, sweeping CRT display, has given way to sophisticated digital signal processing systems that can find, identify, characterize and track aircraft in friendly or hostile skies.

The telephone and all its switches have given way to digital PBXs and FAXs, and by the year 2000 will have given way to the integrated digital services network (ISDN), fiberoptic transmission and digital videotelephones.

Music aficionados may deplore the passing of the LP but the Compact Digital Disc and DAT (digital audio tape) have brought a quality and functionality to home entertainment undreamed of in the traditional analog era. Couple this with digital signal processing and listeners will be hearing the real thing by the year 2000.

And when they're not listening to real music, they'll be watching digital television in the form of HDTV or some variation, and interacting with their TVs in a way that will make the distinction between TV and computers a blur.

And at work, in the factory, they'll be supervising computerized machines and robots that displaced older electromechanical controlled and operated tools: they'll be overseeing chemical processes where messages about temperatures and pressures and chemical mixtures are transmitted over digital networks, analyzed and then adjusted to yield optimum products. In hospitals, physicians and nurses will be capturing the tiny signals from spinning or splitting atoms and looking at complex, enhanced images of organs deep inside the human body and at the biological processes going on in those organs.

Just think ... what the last 30 years have been like ... and what the next 10 will be ... the next 10, when the transformation of electronics is virtually complete and the world has become digital.

THE

A little more than 30 years ago a transformation began in the electronics industry that would change forever the nature of electronic products and systems.

The transformation started with the invention of the integrated circuit and the introduction of the first simple, commercial digital ICs.

In the 30 years since **Computer Design** was chartered to cover the emerging computer industry, digital ICs have transformed computers from data processing monsters confined to special rooms into basic workhorses controlling aircraft, manufacturing, instrumentation, and more. These changes have been more awesome, more compressed in time, and more profound than

SYSTEMS ERA

The Last Thirty Years & The Next Decade

January 1991, Computer Design Celebrates Its 30th Year With A Major Analysis Of The Systems Era

those brought on by the development of the steam engine, the telegraph, the automobile or the vacuum tube which started the electronics age. What began 30 years ago was the digitization of electronics and, even more important, the digitization of the world we live in.

The next 10 years will see the completion of this digitization of our world.

Computer Design has the unique ability to provide this editorial perspective because it was created with the express purpose of covering the transformation of electronics into digital, computer-based products and systems. **Computer Design** has covered the transformation with 400 consecutive issues dedicated to pointing the way; to looking at the technology and design directions that were driving the transformation.

The editorial package for this 40-year look at the transformation of electronics —1960 to 2000—is being jointly developed by John Miklosz, **Computer Design's** Associate Publisher/Editor-in-Chief, Stephen Ohr, a well-respected longtime reporter of electronics and computer technology, and the entire **Computer Design** senior technical staff.

Don't miss this opportunity to celebrate with us. Show your company's commitment and the role it will play in the next 10 years: the decade in which the transformation will become virtually complete.

Vital Details

Issue Date: January 2, 1991
Space Closing: December 3, 1990
Material Closing: December 7, 1990
Bonus: Two personalized hard cover issues for each advertiser.
Extra Copies: Available by advance special order.

See your Computer Design District Manager for details and reservations.

Special offer for 1962/3 charter advertisers: If your firm is one of the farsighted firms that saw the advent of the systems influence and appeared in Computer Design's first issues., contact your Computer Design Salesperson on how you can have an extra ad in this January 1991 issue at 1962/3 rates.

Serving Today's Computer-Based Electronics OEM

COMPUTER DESIGN *Technology and Design Directions*

A PennWell Publication

CIRCLE NO. 79

COMPUTERS AND SUBSYSTEMS

VMEbus embedded PC workstation offers 20-Mips performance

The RadiSys EPC-5, an embedded PC workstation for the VMEbus, combines the rugged VMEbus form factor and the high-performance 80486 processor to create an embedded PC that can take advantage of PC-compatible application and development software. This system subassembly can be embedded into sophisticated computer-controlled equipment and machinery.

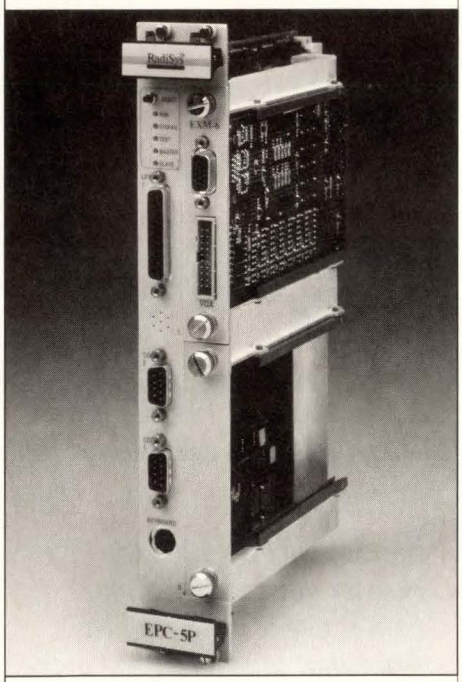

The EPC-5 occupies two 6U VME form factor slots. It consists of a main EPC-5 CPU module and two local EXMbus slots for optional expansion modules. In a standard configuration, one of these slots is normally used for the VGA display controller. To form a complete embedded PC, users can add a mass-storage module expansion containing a floppy and a hard drive with a 40- or 100-Mbyte capacity. At 6×9×3.2 in., the EPC-5 is the smallest 80486-based PC available, RadiSys claims.

A complete set of 80486 electronics and BIOS firmware is included, letting users run off-the-shelf PC application software without modification. In addition, the EPC-5 includes on-board byte-swapping hardware to convert data from little-endian to big-endian format and vice versa.

33 Mips at peak performance

At the heart of the EPC-5 CPU module is the 25- or 33-MHz 80486 microprocessor, with an on-chip 8-kbyte cache and numeric coprocessor. The cache memory permits near-zero-wait-state performance in most applications. The 80486 can execute its fastest instruction in one clock cycle and offers a peak performance rate of 33 Mips. With typical application code running, performance will average 20 Mips.

Memory on the board consists of 4, 8 or 16 Mbytes of DRAM and 64 kbytes of BIOS EPROM. In addition, the system's configuration parameters are contained in battery-backed CMOS RAM. The DRAM memory array has been optimized for 16-kbyte burst transfers between the 80486 CPU and its on-chip cache. The 80486's on-chip cache is automatically updated or invalidated during local and VMEbus writes to the on-board DRAM, which keeps the data coherent throughout. The entire EPROM BIOS contents are copied into DRAM when the system is turned on, which enhances the performance of BIOS calls. The BIOS supports operating system boot loading from local expansion devices on the local EXMbus as well as over the VMEbus.

A private EXMbus expands the capabilities of the EPC-5, allowing communication with disk units, PC add-in cards and other expansion modules. Driven from the CPU module's P2 connector, the EXMbus is electrically similar to the PC/AT bus definition but is modified to fit the EPC-5 mechanical arrangement. A sub-backplane routes the bus signals to other slots without affecting traffic along the VMEbus.

The RadiSys EXM-6 expansion module adds a VGA graphics controller for the EPC-5. Providing up to 800×600-pixel resolution with 16 colors, the EXM-6's VGA implementation includes 256 kbytes of video RAM. Standard VGA, Multisync or monochrome monitors are supported through a standard connector on the front panel.

Providing the full VMEbus interface, the EPC-5 has five RadiSys custom gate arrays. This interface allows master/slave capability and can generate and receive all VMEbus interrupts. The board's VMEbus interface supports single-slot control functions and a subset of the VXIbus. The VXIbus extensions include support for message passing and for dynamic configuration through software.

Using a hardware windowing technique, the 80486 can address the entire 32-bit VMEbus address range from within the DOS addressing mode. This hardware works by providing programmable registers, which add high-order bits to the 80486's basic memory address. The VMEbus address range is seen by the EPC-5 as a 64-kbyte logical page that can be accessed through a DOS page frame.

Besides the CPU, memory and bus interfaces, the EPC-5 includes the elements typical of a complete PC compatible. It provides two serial ports (COM1 and COM2), one parallel port (LPT1), a standard PC/AT speaker and a time-of-day clock with a battery.

Available this month, pricing for the EPC-5 starts at $7,495.

—Jeffrey Child

RadiSys
19545 NW Von Neumann Dr
Beaverton, OR 97006
(503) 690-1229

Circle 356

COMPUTERS AND SUBSYSTEMS

Interface board connects NTDS devices to VMEbus

The Hawke VME Navy Tactical Data Systems (NTDS) interface board from Sabtech Industries provides a high-speed intelligent 32-bit interface between the VMEbus and MIL-STD-1397 A and B I/O devices. This capability allows direct connection of military computer equipment in the open-architecture VMEbus for applications such as software development, simulation and training support. Though the VMEbus systems have long been used to support NTDS equipment, the Hawke is the first VME-to-NTDS interface that lets the full 32-bit bandwidth be used, Sabtech claims.

The Hawke is developed specifically for users, system developers and integrators of military equipment. It offers an alternative to costly and hard-to-get mil-spec hardware because its 1 Mbyte of on-board EPROM can be used to emulate costly peripheral devices. Each board can then represent a specific military peripheral, including such devices as reel-to-reel magnetic tape drives, teletypes and paper tape reader/punches.

Supporting VME Rev. C.1 (IEEE-P1014), the Hawke operates as A32/D32, A24/D16 and A16/D8 in master or slave modes. As a slot-1 master, the board can emulate multiple NTDS devices, handling up to four levels of arbitration. In addition, this single-slot 6U board has an RS-232C port, allowing monitoring and data collection of the I/O activity for an entire combat system.

Simultaneous access

The Hawke was designed to meet the requirements of a high-performance, versatile, intelligent NTDS interface. A 32-bit MC68020 processor coupled with 512 kbytes of fast video RAM delivers a sustained high-speed throughput over independent NTDS channels.

To meet the timing requirements typical of MIL-STD-1397 specifications, this 32-bit architecture permits simultaneous access to the memory, NTDS I/O and VMEbus. This differs from other DMA-based architectures in which the CPU has to relinquish bus control for every transfer. A parallel I/O architecture coupled with a powerful set of drivers lets the Hawke fit all NTDS type A/B applications.

Configurable as a MIL-STD-1397 type A or B interface, the board supports full-duplex NTDS transfers, with RAM buffers of up to 512 kbytes. This memory is organized as 64,000 32-bit words for each I/O channel. Buffers are dynamically allocated to suit user applications. Integration and diagnostics software is built into the monitor to expedite system installation and fault analysis. In addition, up to 1 Mbyte of customized firmware can be installed in the user EPROM. LEDs on the front panel display NTDS activity for all control signals. A battery-backed real-time clock and nonvolatile RAM support time-tagging and error-logging operations, which are useful for data collection and analysis.

Software drivers are supplied on system EPROM. A menu-driven monitor is provided to link the Hawke to a personal computer or terminal with any common terminal-emulation software package. During installation, the monitor facilitates downloading, execution and debugging of applications software through a remote terminal. The system software features comprehensive run-time data-processing and analysis capability, as well as a kernel for user-specific task structures. Unix drivers are available for use of the Hawke in a Sun Microsystems Sun-4 workstation environment; drivers for other operating systems, including OS-9 and PSOS, will been available soon. Under software control, the board can be configured in NTDS Computer, Peripheral or Intercomputer modes.

Available now, the Hawke VME-to-NTDS interface board is priced at $4,295.

—Jeffrey Child

Sabtech Industries
5411 E La Palma Ave
Anaheim, CA 92807
(714) 970-5311

Circle 357

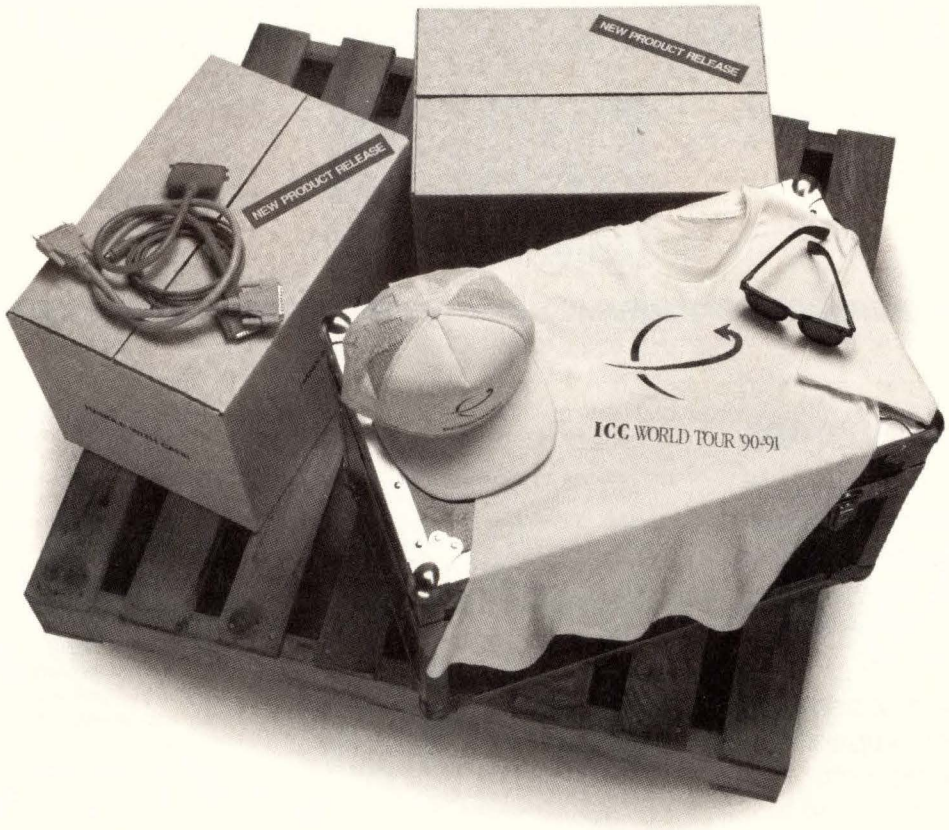

The Best Collection of Peripherals Ever Assembled Have Joined The ICC World Tour

Invitational Computer Conferences (ICC) collects the world's best suppliers for an event you won't want to miss. At a nearby location, you'll see the latest in technology *exclusively* for computer peripherals.

See what today's products offer through hands-on demonstrations ...learn what tomorrow's products will be. Attend technology seminars that give you expert answers to your peripherals questions. Take part in panel discussions and debates on current technology and business issues. Explore all your options so you can make your peripherals decisions.

Make plans now to attend. Watch for your invitation or call for additional facts and details on how you can receive an invitation.

For almost two decades, the ICCs have specialized in bringing the latest computer technologies to you, for your evaluation. With the research expertise from our parent company, Dataquest, we have the support of the world's leading computer analysts at each stop on the ICC tour.

Call for your OEM Peripherals invitation, ***Dataquest/ICC,***
U.S. TEL: (714)957-0171, FAX: (714)957-0903
EUROPE/UK TEL: (0895)835050, FAX: (0895)835260/1/2

U.S.			
Newton, MA	*06*	*Sep*	*90*
Los Angeles, CA	*12*	*Sep*	*90*
Atlanta, GA	*01*	*Oct*	*90*
Dallas, TX	*11*	*Oct*	*90*
Gaithersburg, MD	*23*	*Oct*	*90*
Portland, OR	*04*	*Dec*	*90*
Newport Beach, CA	*08*	*Jan*	*91*
Ft. Lauderdale, FL	*22*	*Jan*	*91*
San Jose, CA	*07*	*Feb*	*91*
Austin, TX	*12*	*Mar*	*91*
Nashua, NH	*01*	*Apr*	*91*
Minneapolis, MN	*21*	*May*	*91*

EUROPE			
Tel Aviv, Israel	*06*	*Sep*	*90*
London, England	*13*	*Sep*	*90*
Frankfurt, Germany	*18*	*Sep*	*90*
Stockholm, Sweden	*27*	*Sep*	*90*
Munich, Germany	*15*	*Jan*	*91*
Vienna, Austria	*22*	*Jan*	*91*
Milan, Italy	*24*	*Jan*	*91*
Paris, France	*29*	*Jan*	*91*

CIRCLE NO. 74

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
December 1	November 2
January 1	December 7

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

Positions Wanted
One Technology Park Dr
PO Box 990
Westford, MA 01886

HARDWARE AND SOFTWARE DEVELOPMENT SERVICES
Let ACT assist your design team.
• We specialize in developing microprocessor based products and embedded software. • We can help with part of your project or handle the entire development. • Facilities include Micro Development Systems, I²ICE, CAM/CAD and PCB layout system. • Analog and digital circuit design (video, FPGA, PAL, FLIR). • Bus development for Multibus, VME, PC, MCA, LAN & STD.
APPLIED COMPUTER TECHNIQUES, INC.
3914 Dekalb Dr. Orlando, FL 32809 (407) 851-2525

SALES MANAGEMENT/TECHNICAL MARKETING POSITION. Resourceful problem-solver with 15 years broad-based experience in planning, management, and coordination of projects. BA-BS-MBA. Contact: Robert A. Mazzola, 7251 Elk Circle, Apt #2, Huntington Beach, CA 92647; (714) 841-8978.

ENGINEERING COMPUTER SYSTEMS MANAGER. Implementation, support and management of CAD systems for electrical, mechanical and software development applications. Significant experience with DEC systems including: Microvax, workstations, and networks. Generate new technology proposals and perform start-up activity. Southeastern PA location. Resume on request. Charles Martocello, 634 Knoll Drive Lansdale, PA 19446. (215) 855-6189.

QUALIFIED FIELD SERVICE for Your Company in SW Missouri without high travel costs. 15+ years experience with mini-computer systems, wide variety of printers and other devices. If your customer knows how to use the equipment and you have in-house technical support, I can begin fixing your customer's problems today. Alan (417) 744-2214.

SENIOR SOFTWARE ENGINEER. 11+ years experience. Device drivers, communications, systems architecture in real-time process monitoring, point-of-sale and biomedical applications. CASE experience. Strong writing, requirements definition and project planning skills. Full time, will relocate. Request resume c/o E. Gorski, 12503 Brandywine Court, Austin, Texas 78727.

HIRING MGR'S GODSEND. BSEE + 10 yrs. IC & board-level design. CPU, EEPROM, Data Link, DMA. Discrete DSP, 68K Intfc, St Machines. Firmware. UNIX, VMS, TOPS20, DOS. FORTRAN, Assy, C. Test Stns, AT-Reconfiguration. Problem solver: Concepts to crying babies. Jim (415) 573-0737.

SALES ENGINEER or SALES MANAGER - demonstrated success with high-tech components, equipment, systems, networks, datacom/telecom. Experience with distributors and direct sales in 15 South-eastern states. BS-degree, technical training, heavy "hands-on" technical experience. Peter Alexander, Atlanta, Georgia (404) 945-6950.

HARDWARE/SOFTWARE ENGINEER. 20 years design engineering experience in Analog Instrumentation, Digital Controls, Data Acquisition, and Switching Power Supplies - Proficient in Worst Case Analysis, SPICE Modeling, Autocad, using IBM Compatible Computers - Some experience with Basic, Microsoft C, and Assembly Language Programming - Six Patents in Instrumentation, Active Filters, and Signal Processing - Work experience in Commercial, Consumer, and Military Markets - L.A. area only please. - George Dodson (818) 242-6925.

FIRMWARE ENGINEER. heavyweight assembly language programmer for 8051 family and similar microcontrollers. Experienced in high-volume consumer products and in complex communication and control systems. BSEE, nine years experience. Alfredo Ruiz, (407) 321-1860.

HELP WANTED

SYSTEM ENGINEER: Design 16 bit bus and micro channel video and input/output device, and programming with Assembler language in the BIOS level. Develop 32 bit and 64 bit microprocessor system PC board. Will use digital subsystem design techniques, subsystem design, data structures, and microprogramming techniques to perform. Requires M.S. degree in Electrical Engineering with major in Computer Engineering. Education to include completion of two courses in digital subsystems, one course in microprogramming, and one course in data structures. Hours: 8:30 a.m. - 5:00 p.m. 40 hours per week at $33,950 per year salary. Please send resume to: ILLINOIS DEPARTMENT OF EMPLOYMENT SECURITY, 401 South State Street - 3 South, Chicago, Illinois 60605, Attention: LEN BOKSA, Reference #1792-B, AN EMPLOYER PAID AD.

ELECTRICAL ENGINEERING TECHNOLOGY
The Department of Engineering Technology of the University of North Carolina at Charlotte seeks to fill a tenure track faculty position in the Department of Engineering Technology offering TAC/ABET accredited Bachelor of Electrical Engineering Technology program beginning August 16, 1991. Requires expertise in two or more of the following areas: microelectronics, digital logic, microprocessors, linear circuits, electronic devices and circuits, or control theory. Master's degree in Engineering, Engineering Technology or closely related field, with five or more years industrial electronics experience required. High priority will be given to applicants that additionally have P.E. registration, have prior upper level teaching experience in an accredited 4-year Engineering Technology program, and that have a record of recent scholarly activities including funded research. Rank and salary commensurate with qualifications. Applications will be accepted until position is filled; review of applications will begin January 1, 1991. Send applications, with detailed resume and names and addresses of at least five references to: Edwin R. Braun, Chairman, Department of Electrical Engineering Technology, University of North Carolina at Charlotte, Charlotte, North Carolina 28223. UNC-Charlotte is an equal opportunity, affirmative action employer.

SYSTEM SHOWCASE

Reach over 100,000 qualified engineers and engineering managers with System Showcase advertising. **Rates start at $765.**

For more information call Sue Nawoichik at (800) 225-0556 or (508) 392-2194.

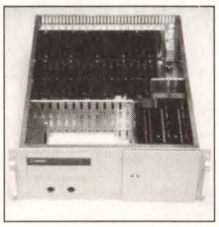

18 SLOT RACKMOUNT AT SYSTEM

The QPC-7300, an eighteen slot industrial rackmount 286/386 AT system, features a low profile (7" of vertical rack space) enclosure for a standard 19" rack. 300 Watt power supply, four half-height drive bays, and positive air flow cooling system are standard. Optional split backplane supports two systems. Fully featured single board computers support DOS and REALTIME operating systems.

Qualogy, Inc.
1751 McCarthy Blvd.
Milpitas, CA 95035
(408) 434-5200 (800) 722-4976
CIRCLE NO. 177

Embedded Control & Data Acquisition System

Easy to use SBS–2300H has everything to set up and run an embedded data acquisition and control system. Industrial control, operator interface, communications, and operating system all on a single card.

Built–in CAMBASIC II™, multitasking software supports all hardware features. No development system needed. Simply write your program and store it in nonvolatile memory that autoruns on power–up. Use your PC for writing and editing convenience.

- 10 Analog I/O, 12–Bit
- 48 Digital I/O Lines
- 128K RAM
- Keypad Port
- 2 Serial Ports
- EPROM/EEPROM Programmer
- Calendar Clock
- Drives Opto 22 Racks
- Network up to 32 Systems
- Low Power, –20° to 70°C

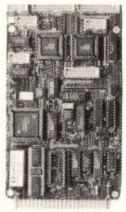

Call for 1990 Catalog
303–430–1500
OCTAGON SYSTEMS CORPORATION
6510 W. 91st Ave., Westminster, CO 80030

CIRCLE NO. 178

RUGGED AND HIGH PERFORMANCE COMPUTER SYSTEMS WITH FOLD DOWN KEYBOARD AND VGA MONITOR FOR RACK, BENCH OR PORTABLE APPLICATIONS

Standard features include:

* 12 slot passive back plane
* **80386** CPU card at **20/25/33** MHz, up to **8MB** of zero wait state RAM
* Sony Trinitron tube, high resolution VGA (640x480) monitor and card
* Room to mount three half height drives
* 2 serial, 1 parallel port, MS DOS/GW basic

Also available with **80486** or **80286** CPU cards in various configurations, for further details contact:

IBI SYSTEMS INC.
6842 NW 20 Ave., Ft. Lauderdale, FL 33309
(305) 978-9225 Fax: (305) 978-9226 Telex: 529482 IBI SYSTEMS

CIRCLE NO. 179

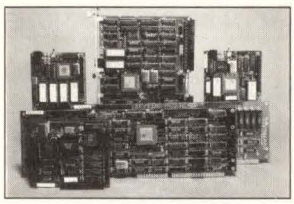

SC/FOX™ Embedded-System Computers

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, Standalone operation, Eurocard size, 1 ptr, 1 serial port, up to 512K bytes O-ws memory, 2 50-pin user connectors, SCSI option. Uses Hams 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

PC/AT™ PROCESSORS

FOR PASSIVE BACKPLANE SYSTEMS

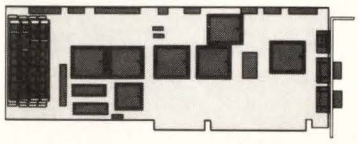

A COMPLETE FAMILY OF BOARDS FOR THE OEM AND INTEGRATOR:

- 286SBC - 12MHZ, 4MB, VGA, IDE, FDD, P, S
- 386SX - 16/20MHZ, 4MB, IDE, FDD, P, S
- 386DX - 20/25MHZ, UP TO 10MB

MANUFACTURING QUALITY PRODUCTS IN THE U.S. FOR 15 YEARS

See us at COMDEX/Fall '90 SANDS, BOOTH: N250

TRENTON TERMINALS INC.
2900 CHAMBLEE TUCKER ROAD
BUILDING 10
ATLANTA, GA 30341
(404) 457-0341 FAX (404) 457-0349

CIRCLE NO. 181

QUICK-TURN PC BOARDS

INSTANT QUOTES BY PHONE

1-800-234-1556

Single/Double Sided – 1 Day
Multi-Layer – 2 Days
MIL Approved
Gerber Input/Floppy or Modem

AVANTI CIRCUITS, INC.
17650 North 25th Avenue
Phoenix, Arizona 85023

(602) 863-7729 & 234-1556
Fax (602) 375-1909
BBS/Modem (602) 234-1737

CIRCLE NO. 182

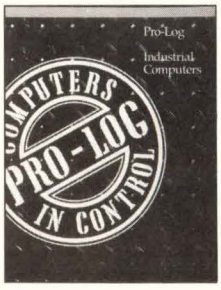

STD BUS BOARD LEVEL PRODUCTS

Pro-Log's line of STD Bus board level products are ideal for low power, high noise immunity and varying temperature range applications. Their reliable, high speed processing is ideal for use in manufacturing automation and industrial control. A variety of expansions and options enable the system to meet individual application specifications.

Pro-Log Corporation
(800) 538-9570
2555 Garden Road, Monterey, CA 93940

CIRCLE NO. 183

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. 184

SYSTEM SHOWCASE

COLOR, GRAY SCALE & BINARY IMAGE COMPRESSION

OPTIPAC is an extremely flexible image compression system providing both lossless and controlled quality modes. It is currently implemented as an add-in card for PC/AT compatibles and also as a software only version for DOS.

STANDARDS: JPEG (ANSI X3L2.8) for color & gray scale images and CCITT Group 3 and 4 binary images.

SPEED: the new OPTIPAC-ISA-Model B can compress/decompress full screen images in less than a second.

Optivision, Inc. 1-800-562-8934
Davis, CA Fax: (916) 756-1309

CIRCLE NO. 185

PC COMPATIBLE ENGINEERING

Annabooks gives you the **hardware, software, and firmware** information you need to design PC-compatible systems faster and better. And you have control of your design from the ground up - our firmware and software products include **source code!** Plus all the utilities you need. Do hardware design? The AT Bus Design book and the XT-AT Handbook replace a whole shelf of references. Start by getting these books:

- **AT BiosKit:** an AT Bios with source code in C you can modify. With setup & debug. 380 pages with disk, $199
- **XT BiosKit:** Includes a debug. 270 pages with disk, $99
- **Intel Wildcard Supplement** for XT BiosKit: Includes ASIC setup, turbo speeds, also useful with many other modern XT boards. 60 pages with disk, $49
- **AT Bus Design:** At last here are the complete timing specs to show you how to design ISA and 8/16 bit EISA, $69.95
- **PromKit:** Puts anything in Eprom or SRAM; DOS, your code, data, you name it! With source on disk, $179
- **SysKit:** Here's a debug/monitor you can use even with a brand X Bios in your desktop. Runs in ROM or TSR in RAM. Includes source, of course, $69
- **XT-AT Handbook:** The famous pocket-sized book jam-packed with hardware & software info, $9.95 ea. or 5 or more for $5 each.

Software tools:
You need MS C & MASM 5.1 for modifying the Kit products.

FREE Mention this ad when you order any publication and get a free XT-AT Handbook by Chesser & Foster!
Hurry, before we come to our senses and change our minds.

VISA 800-462-1042 MasterCard
619-271-9526

Annabooks
12145 Alta Carmel Ct., Suite 250
San Diego, CA 92128

FAX 619-592-0061
Money-back guarantee

CIRCLE NO. 186

Data I/O Programming Tools

50% OFF

280 Set Programmer **plus** PROMlink PC File Management Software just $995*.

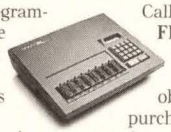

- Set/gang programming with the 8-socket 280
- Supports E/EEPROMs up to 512K
- Easy PC control with PROMlink™ software
- FREE one-year Data I/O® warranty

Call now for a FREE 15-day trial with no obligation to purchase,** and receive a FREE tutorial on programming today's device technologies.

*U.S. list price only.
**Offer expires 11/15/90.

1-800-247-5700

DATA I/O Corporation

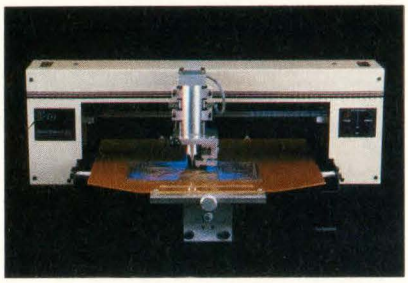

FROM CAD TO FINISHED PROTOTYPE PCB

- Without the wait
- Without chemicals
- Without outside expenses

Call Instant Board Circuits at (415) 883-1717 to order your free BoardMaker™ video and literature.

CIRCLE NO. 188

RAM CARD APPLICATIONS

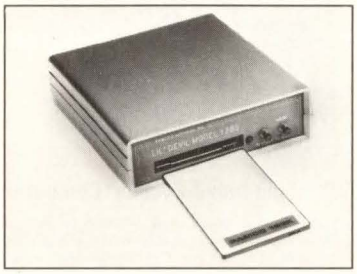

MS-DOS compatible, serial and parallel RAM CARDRIVE. Installable software drivers and format programs to make a RAM card "look like" an MS-DOS formatted drive. Compatible with XT, AT, PS/2, Zenith, Compaq and others.

PAMCO ELECTRONICS INC.
377 Carowinds Blvd.
Fort Mill, SC 29715
1-800-255-6265
(803) 548-6740

CIRCLE NO. 189

Precision Time Plug-ins

by TRUETIME

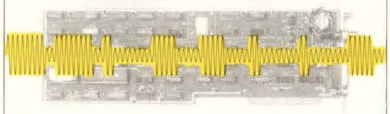

VME · Multibus · PC

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. 190

8086 8051 8096

ROYALTIES

Multi-Tasking EXECS

U S Software offers hi-performance software tools for embedded applications.

Get the full details by calling:
800-356-7097
503-641-8446
503-644-2413 (FAX)

U S SOFTWARE
United States Software Corporation
14215 NW Science Park Drive
Portland, Oregon 97229

68000 68HC11 Z80

CIRCLE NO. 191

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. 192

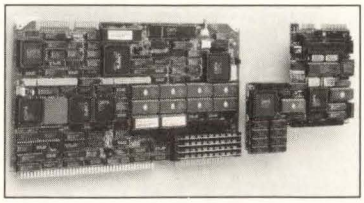

Multibus[tm] AT-COMPATIBLE SBC

Multibus I is now PC/AT[tm] Compatible with MAT 286 which includes all of the functions of a 10 MHz AT motherboard, plus 2 serial ports, a parallel port, two SBX interfaces, 10 28-pin EPROM sockets, and up to 4M of parity checked, dual ported DRAM. A daughter card, MATxSYSIO, adds EGA, floppy, and SCSI interfaces.

Single Board Solutions, Inc.
20045 Stevens Creek Blvd, Cupertino, CA 95014
(408) 253-0250

Multibus[tm] Intel Corp. PC/AT[tm] IBM

CIRCLE NO. 193

EMS 4.0
4 Meg
1/2 size card.
Simple. Reliable.
Very compatible.

$695.00

Designed, Manufactured, Sold and Serviced by:

Rtech Inc.
Voice . . . (612) 345-4555
FAX . . . . (612) 345-5514
Modem . . (612) 345-4656
907 N. 6th St., Lake City, MN 55041

CIRCLE NO. 194

Elegant, concise, fast & standardized
FLOATING POINT
libraries for embedded applications

Based on the IEEE 754 standard, FPAC (32 bit) and DPAC (64 bit) libraries are mature, well documented, and fully tested. The libraries are fully ROMable and include the following:

- Basic Operations
- Square Root
- Trigonometric
- ASCII Conversion
- Integer Conversion
- Logarithmic

U S Software supports most Intel, Motorola, Zilog and Hitachi micros, including 80X86, 80386, 680X0, 80960, 8051, 8096, 68HC11, Z80, 6809 and 6301.

For additional information, please contact:

U S SOFTWARE
United States Software Corporation
14215 NW Science Park Drive
Portland, Oregon 97229
800-356-7097
503-641-8446
503-644-2413 (FAX)

CIRCLE NO. 195

SCHEMATIC CAPTURE-FREE DEMO DISK

OrCAD/SDT III: the most popular electronic design tool in the world. Over 150 drivers for printers/plotters/graphic cards, 6200+ unique library parts, utilities for electrical rules checking, B-O-M generation and converting netlists to over 30 different formats are included for only $495.

Call (503) 690-9881

CIRCLE NO. 196

MIL-SPEC•COMMERCIAL•VME•DSP
A Classic Solution for Today's Challenges

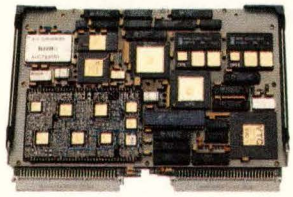

9000m Series Electronic Module

- MC68020 & DSP56001 Processors
- VMEbus Master/Slave
- Modular Analog Subsystem
- 64 16-Bit A/D Inputs - 50Khz Sampling
- 4 12-Bit D/A Outputs
- Full ADA & DOD-STD-2167A Support

Cadillac Gage TEXTRON
Control Systems Marketing • P.O. Box 1027
Warren, MI 48090 • 313/777-7100 • Fax: 313/776-9731

CIRCLE NO. 197

Program
Your Chips
In Sets of 4 for $495.00

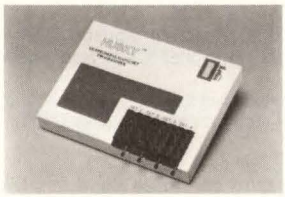

Special offer Now Includes:
Free UV eraser, CUPL starter Kit and a **$300.00 Rebate** with the PDT-1 Universal Programmer System\Kit.

LOGICAL DEVICES, INC. 1-800-331-7766

CIRCLE NO. 198

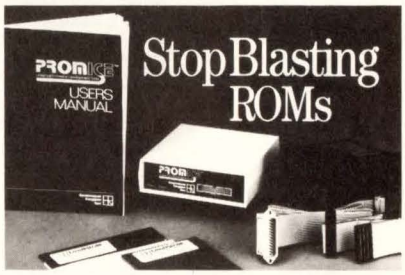

PROMICE emulates 8 bit ROMs from 2716-27080, or 16 bit ROMs 27C1024 or 27C2048. (Inquire about emulating other ROMs. Non-JEDEC ROMs require custom cable.) ■ Sophisticated LoadICE™ Host Software downloads, uploads and edits ROM contents, supports MS-DOS, UNIX, MAC & VMS. Software sources are included. ■ 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. 199

IT'S IN THE CARDS ...
More Leads, More Action, More Sales

Our readers are proven buyers of:

Computer Systems	93,360
Systems Boards	84,362
Integrated Circuits	78,181
Design/Development Equipment	85,552
Memory/Storage Equipment	72,236
Software	88,913
Terminal/Input-Output Equipment	82,726
Communications Equipment	70,389
Components	75,914
Test Equipment	65,400

Smart buyers depend on COMPUTER DESIGN DIRECT ACTION CARDS

Call Sue Nawoichik:
at 800-225-0556
or 508-392-2194

COMPUTER DESIGN

Don't Get Zapped!

High inrush current can **destroy** your sensitive VAX CPUs and peripherals **in less time than it takes to flip a switch.**

THE SOLUTION?
Power up with ***Z-LINE*** **TPC 115-10 MTD™ the smallest** power distribution and control system available.
POWER UP WITH — — —

Z-LINE

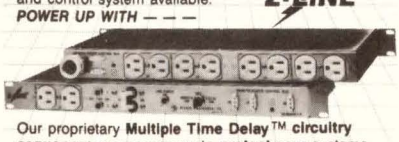

Our proprietary **Multiple Time Delay™ circuitry sequences** your power-up to **protect** your systems from the spikes and surges, EMI & RFI, that **destroy your hardware and erase your data.** And our remote on/off and emergency shutdown **gives the power control back to you.**

All Pulizzi Engineering MTD™ **controllers are compatible with DEC and UPS systems.**
PRICES FROM $453 TO $317

DON'T WAIT UNTIL IT HAPPENS, CALL TODAY!
PULIZZI ENGINEERING INC.
3260 S. Susan Street, Santa Ana, CA 92704-6865
(714) 540-4229 FAX (714) 641-9062

CIRCLE NO. 187

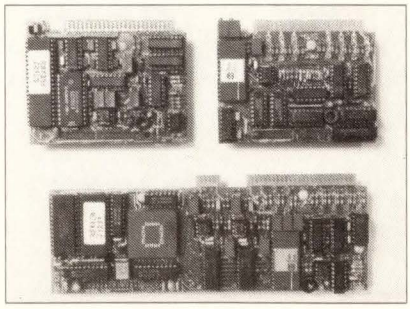

SBX ANALOG I/O
GET A COMPETITIVE EDGE!

Density - up to 16 A/D inputs & 8 D/A outputs on one single width card. **Speed** - throughput rates up to 59 kHz. **Intelligence** - many preprogrammed modes. Input filters, prog. gain amps, sample-holds, FIFO I/O buffers & many other features too numerous to mention here.

ROBOTROL CORP
925 W. San Martin Ave, PO Box 990
San Martin, CA 95046
(408) 683-2000

CIRCLE NO. 200

1991 Computer Design Magazine Edition UPCOMING ISSUES

Watch for these Special Features & Events in Computer Design Magazine

MONTH	SPECIAL REPORT	TECHNOLOGY FOCUS	PRODUCT FOCUS
January 1	30th Anniversary Issue — The Evolution Of Electronics Into Computers *COMPUTER DESIGN* editors		Logic analyzers *Jeff Child*
February 1 Buscon-W	Workstation buses *Warren Andrews*	FDDI ICs — *Ron Wilson* Design capture — *Mike Donlin*	High-speed D-A converters *Jeff Child*
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*
July 1	CASE for real-time programming *Tom Williams*	Mezzanine buses — *Warren Andrews* Device modeling — *Mike Donlin*	Multibus CPU boards *Jeff Child*
August 1* Siggraph	Mixed-signal ASICs *Barbara Tuck*	Display controller ICs — *Ron Wilson* Software-management tools — *Tom Williams*	DRAMs *Jeff Child*
September 2 Buscon-E ESC	Enhanced-performance standard buses *Warren Andrews*	CAD frameworks — *Mike Donlin* RISC in real-time — *Tom Williams*	Device programmers *Jeff Child*
October 1*	32-bit microcontrollers *Ron Wilson*	Fast PLDs — *Barbara Tuck* RISC-based CPU boards — *Warren Andrews*	Flash EPROMs *Jeff Child*
November 1 Wescon ITC	System simulation and verification *Mike Donlin*	High-density ASIC packaging — *Jeff Child* Multiprocessing in real-time — *Tom Williams*	STD CPU boards *Jeff Child*
December 2	Migrating PLDs to full ASICs *Barbara Tuck*	Accelerators to boost standard-bus performance *Warren Andrews* 8- and 16-bit microcontrollers — *Ron Wilson*	High-resolution A-D converters *Jeff Child*

**Starch Readership Study Issue*

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 Technology and Design Directions

A PennWell Publication

1991 Computer Design News Edition UPCOMING ISSUES

Watch for these Special Features & Events in Computer Design's News Edition

MONTH	SPECIAL REPORT	PRODUCT UPDATE UPDATE	TEST/ MANUFACTURING OPPORTUNITIES	MANAGEMENT
January 14* Buscon-W	Design's role in concurrent engineering†	Hard disk drives	Bus-based Testing	Massachusetts
February 11 Nepcon-W	Manufacturing's role in concurrent engineering†	Printed circuit boards	CIM systems & software	Southern California
March 18 SMT/Con-E	Purchasing's role in concurrent engineering†	Power supplies	Statistical process control	Northwest
April 15 Electro EDS	Testing's role in concurrent engineering†	SMT components	In-circuit testing	New York, New Jersey
May 13* Comdex	Choosing the members of your Time-To-Market Team	Floppy disk drives	Manufacturing inspection	Southeast
June 10 DAC Nepcon/E	Identifying the critical paths to product quality	DRAMs, SRAMs	ASIC prototype verification	Colorado
July 10 Siggraph	Integrating vendors into your Time-To-Market Team	Flat-panel displays	Functional testing	Texas
August 12	Designing for manufacturability	Interconnect systems inspection	Incoming	Arizona
September 9 Buscon-E ESC Nepcon/SE	Interfacing the Time-To-Market Team with customers	Hybrids	Using SMT components	Mid Atlantic
October 14* ITC SMT/Con-W	Designing for test	IC packaging & placement	Automatic insertion	Upper Midwest
November 18 Wescon	Using distributors to best advantage	PROMs, EPROMs, EEPROMs	Contract manufacturing	Northern California
December 16	Winning the Malcolm Baldrige Award	Tape drives	IC testing	Florida

**Starch Readership Studies*
†A four-part series

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

ADVERTISERS INDEX

ADVERTISER	PAGE NO.	CIRCLE NO.
Alacron	18	13
AMP	116,117	63
Analogic	75	40
Anna Books	156	186
Applied Microsystems	96,97	54
Array Microsystems	39	24
Atmel	125	66
Automatix	142	73
Avanti Circuits	155	182
Boston Systems Office	80	44
Brooktree	144A	–
Cadillac Gage Textron	157	197
Cadisys	152	87
Chrislin Industries	146	76
Ciprico	11	7
CMC	69	37
CMC	71	38
Conference Management Corporation	151	80
CSPI	5-9	85
Cypress	Cover 4	–
Data I/O	115	62
Data I/O	156	–
Diversified Technology	18	12
DY-4 Systems	109	61
EDI	91	50
Electronic Solutions	1	2
Elite Group	19	14
Force Computers	6,7	5
Franklin Software	156	192
GCOM	78	42
GCOM	155	184
Gespac	Cover 2	1
Graham	9	6
Grammar Engine	157	199
Green Springs Computers	144	75
Heurikon	105	58
Hewlett Packard	48	–
I-BUS	29	18
IBI Systems	155	179
Innovoice	121	64
Instant Board Circuits	156	188
Intel	31	20
Intel	122,123	65
Interactive Circuits & Systems	67	36
Intermetrics	45	28
Interphase	103	57
Invitational Computer Conference	153	74
Ironics	98	55
John Fluke	40,41	25
Kinemetrics	156	190
Lambda Electronics	64A	–
Logical Devices	157	198
Lucas Deeco	41	77
Matrox	143	72
Measurement Systems	30	19
Micro Memory	130	69
Microtek	95	53
Micron	93	51
Microware	27	17
Mitsubishi	16	10
Mitsubishi	137-139	71
Mizar	37	23
Mizar	42	26
Motorola Microcomputer	14,15	9
Motorola Semiconductor	127	–
NEC	50-55	30,31,32
Newbridge Microsystems	107	59
Oasys	83	46
Octagon Systems	155	178
OKI Semiconductor	4,5	4
Omation	17	11
Omnibyte	112A	119,120
Optivision	156	185
Orcad	157	196
Pacific Electro Data	82	45
Pamco	156	189
PC Technology	157	194
Philips Components	20	15
Pro-Log Corp.	155	183
Pulizzi Engineering	157	187
Qualogy	155	177
Radisys	13	8
Ricoh	84	47
Robotrol	157	200
SBE	Cover 3	84
Sealevel Systems	66	82
Sierra Systems	79	43
Signetics	62,63	35
Silicon Composers	155	180
Single Board Solutions	156	193
Sky Computers	61	34
Sorensen	80A	–
Standard Microsystems	2	3
Systems Manufacturing Tech	152	81
Tadpole Technology	72,73	39
Teknor Microsystems	60	33
Tektronix	46,47	29
Texas Instruments	24,25	16
Texas Instruments	32,33	–
Texas Microsystems	89	49
Themis	101	56
Trenton Terminals	155	181
Ultimate Technology	129	68
US Software	156	191
US Software	157	195
Vector Electronic	108	60
Viewlogic	35	22
VMIC	87	48
W.H. Brady	56	86
Wescon	147	78
Wind River	128 A	–
Zendex	78	27
Zoltech	34	21
RECRUITMENT ADVERTISERS:		
Applied Computer Techniques	154	–
Galaxy Technology	154	–
Univ. of NC, Charlotte	154	–

The Advertisers Index is published as a service. The publisher does not assume liability for errors or omission.

COMPUTER DESIGN Technology and Design Directions®

Executive Office: One Technology Park Drive, P.O. Box 990, Westford, MA 01886
(508) 692-0700 Fax: (508) 692-0525 Telex: 883436
Publisher: David L. Allen (508) 392-2111
Associate Publisher/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, P.O. Box 3466, Tulsa, OK 74101 (918) 831-9443
List Rental: Deanna Rebro (918) 831-9551

SALES OFFICES

Associate Publisher
Tim L. Tobeck

Mid Atlantic
Eleanor Hobbs
14 Vanderventer Avenue, 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 Drive
P.O. 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) 481-6969

South Central/South East
Eric Jeter
19627 I-45 North, 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 Court, Suite 234
Palo Alto, CA 94303
Tel: (415) 965-4334
Fax: (415) 965-0255

Oregon
Tom Boris
2232 SE Bristol, Suite 109
Santa Ana Heights, CA 92707
Tel: (714) 756-0681
Fax: (714) 756-0621

Washington
Bill Cooper
1000 Elwell Court, Suite 234
Palo Alto, CA 94303
Tel: (415) 965-4334
Fax: (415) 965-0255

U.K./Scandinavia
David Round
69 Imperial Way
Croydon,
Surrey CRO 4RR, England
Tel: 81-686-7655
Telex: 938420
Fax: 81 688 2134

France/Belgium/ S. Switzerland/Spain/ The Netherlands
Daniel R. Bernard
247, Rue Saint Jacques
75005 Paris France
Tel: (1) 43 54 55 35
Telex: 214235F MIX
Attn: Missitex: PENNWELL
Fax: (1) 46 33 63 54

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: 2 7531445
Telex: 328601 RANCAD I
Fax: 2 7532354

Japan
Sumio Oka
International Media Representatives, Ltd.
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
Anne Goh-Taylor
Seavex, Ltd.
400 Orchard Rd.
10-01, Orchard Towers
Singapore 0923
Republic of Singapore
Tel: 734-9790
Telex: 35539 SEAVEX RS
Fax: 732-5129

Jay G. Seo
Yong-Jin Park
Doobee International, Ltd.
Center Building (Byulgwan)
1-11 Jeong-dong, Choong-ku
CPO Box 4557
Seoul, Korea
Tel: 776-2096
Telex: K27117 DOOBEES
Fax: 755-9860

COMPUTER DESIGN (ISSN-0010-4566) is published twice monthly, by PennWell Publishing Company, 1421 S. Sheridan, Tulsa, OK 74112. Second class postage paid at Tulsa, OK 74112 and additional mailing offices. Editorial offices are located at One Technology Park Drive, P.O. Box 990, Westford, MA 01886. Annual subscription price: $80 in the U.S.A., $150 in Canada, $135 for Europe via air, and $175 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 North Zeeb Rd., Ann Arbor, MI 48106. POSTMASTER: Send change of address form to COMPUTER DESIGN, Circulation Department, PO Box 3466, Tulsa, OK 74101.

SBE...At the Core of WAN Interface Solutions

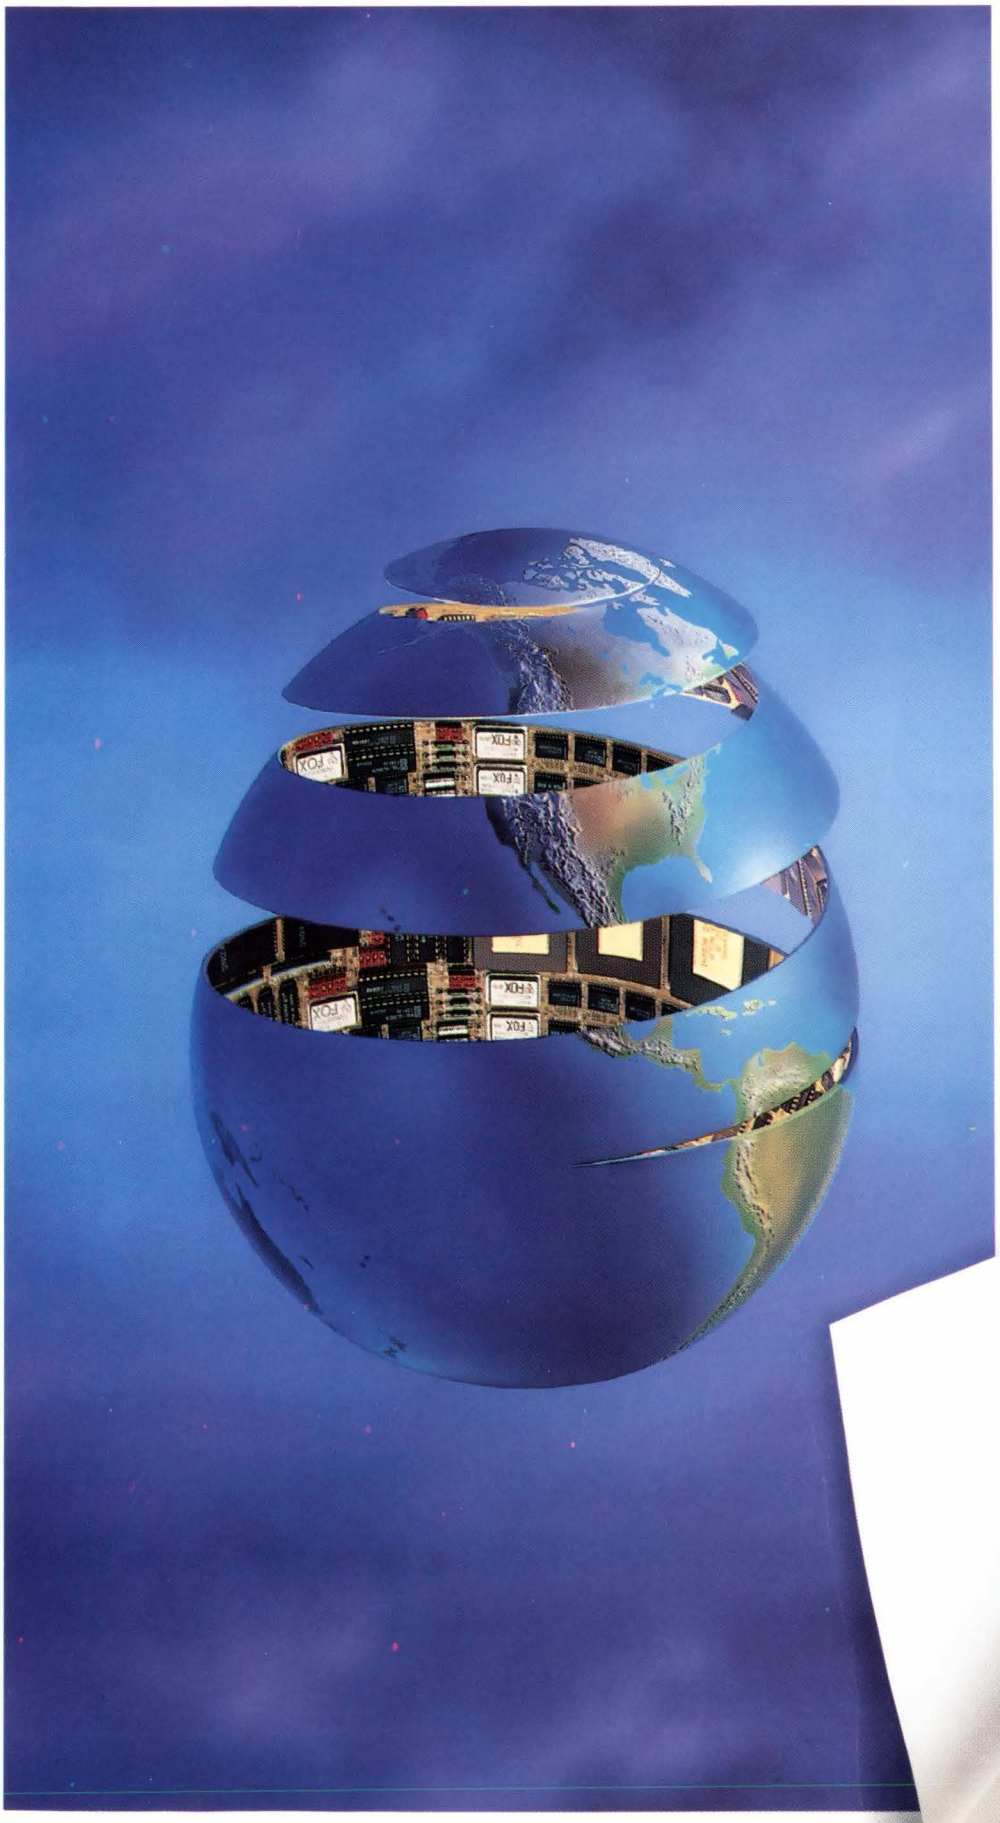

The SBE VCOM-4 Multiprotocol Communications Controller *. . . today's high-performance, cost effective WAN interface.* For price-performance in a *single* VMEbus communications controller, nothing equals the new SBE VCOM-4.

This exclusive SBE card features *four* full-duplex, independently programmable serial channels. Yet, it takes up only one VMEbus slot and provides twice the throughput of conventional boards. A complete implementation of X.25 is available ported to the VCOM-4, which speeds your product to market.

The result: an unmatched WAN interface for VME-based hosts, front-end processors, and data/voice networking systems, including: ■ Two channels at T1 speeds. ■ All four channels can operate with *sustained* throughput at speeds up to 768 Kbps, interfacing to fractional T1 services or 56/64 Kbps lines. ■ Each communications channel can be *independently* configured to support HDLC, SDLC, Bisync, Async.

Turn to SBE and the VCOM-4 for the core of your VMEbus WAN product design application. For fast action, contact SBE, Inc., 2400 Bisso Lane, Concord, CA 94520, or call **1-800-347-COMM.**

CIRCLE NO. 84

SBE

Communications & Real-time Solutions

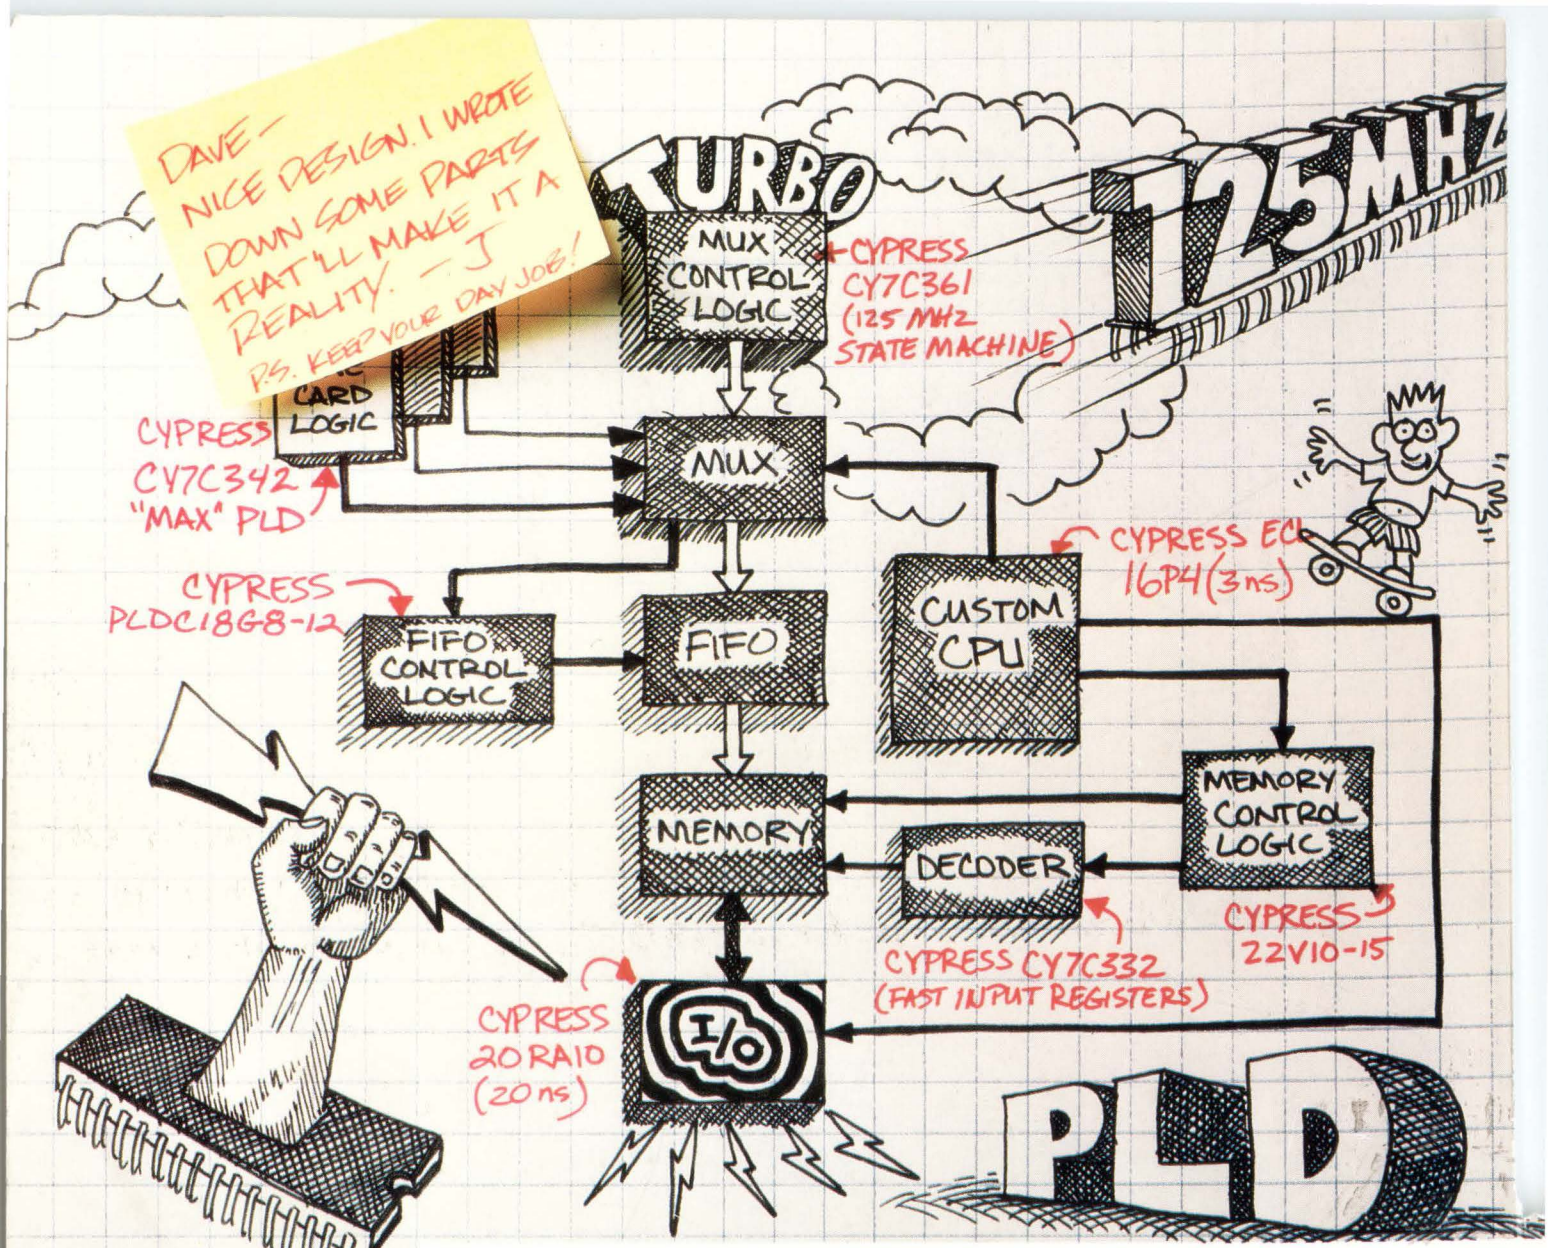

No matter where you're going, or how fast, we have the right PLD.

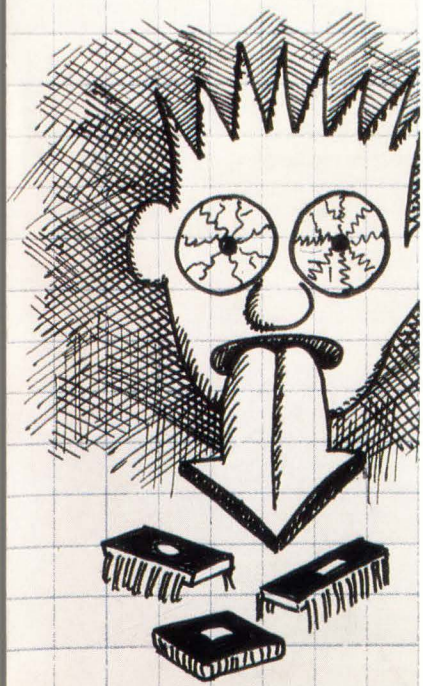

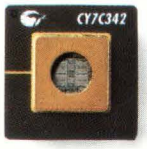

High Density Family: If you need density up to the level of small gate arrays, coupled with high performance and quick development times, our MAX™ family fills the bill. You get parts that can replace up to 50 TTL parts, or up to 15 PLDs, with performance to 50 MHz. Very flexible, very well supported.

Functionally Specialized Family: We've created new architectures tailored to key functions, to give you maximum performance. For example, for state machine functions, our CY7C361 employs an innovative 'split-plane' architecture to cut feedback loop delay and enable 125 MHz performance.

Standard Enhanced Family: If you like the 'classics' but want state-of-the-art performance, you'll find plenty of solutions in our Standard Enhanced Family. Consider our CMOS 18G8 Universal PAL at 12 ns. Or our CMOS 22V10 at 15 ns. Or our 20RA10 at 20 ns. Our ECL 16P4 (10E302) at 3 ns. To name a very fast few.

Call for your free Data Book. Hotline: 1-800-952-6300.* Ask for Dept. C11J.

CYPRESS SEMICONDUCTOR

*1-800-387-7599 In Canada. (32)2-672-2220 In Europe.