A PENNWELL PUBLICATION

COMPUTER DESIGN®

NOVEMBER 1, 1988

THE FIRST MAGAZINE OF SYSTEM DESIGN, DEVELOPMENT AND INTEGRATION

Robotic systems learn through experience

- Old and new contenders vie for position in 32-bit bus arena *Page 47*
- CMOS RAM brings industrial-strength data storage *Page 73*
- Compact dc-dc converters serve distributed power applications *Page 85*

You've just assembled the "perfect" EDA system. Now it's boxing you in.

Mentor Graphics delivers the only Electronic Design Automation system that grows in size and power at the same time.

Engineer or referee?

The time has come. You're about to take on a major design project, so you assemble a little of this and a little of that to construct the "perfect" EDA system. And quickly discover it's too little too late.

Why? Because unless all your tools are on speaking terms, you spend more time patching up arguments than you do designing. And the further you expand your system, the more boxed in you become.

Fortunately, there's a single, simple solution. Mentor Graphics' EDA tools.

The sum of the parts.

No other EDA vendors provide the rich inter-tool communications paths found within our IDEA Series™ EDA system. From schematic capture to output for manufacturing, your design data flows freely from one tool to the next.

Nor do they offer the broad array of tools you need to actually get your design out the door. Whether you're designing ICs, ASICs, PCBs, systems or embedded software, Mentor Graphics' tools cover the entire design cycle.

Nor can other vendors claim the experience we've accumulated. Currently, there are over 12,000 Mentor Graphics workstations producing everything from 32-bit VLSI microprocessors to 108 MIPS parallel computers. And our continuous dialogue with these customers gives us invaluable feedback in designing subsequent generations of EDA tools.

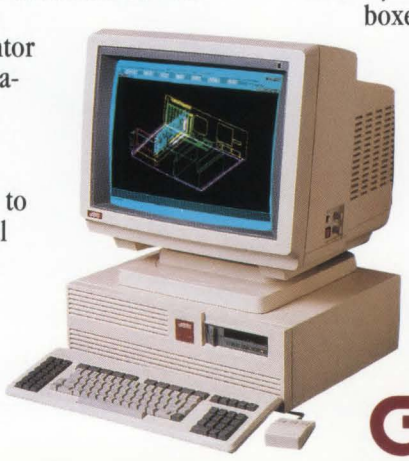

Mentor Graphics also provides a single source of service and support for your entire EDA tool set. No more finger pointing. No more delays. Our support organization takes care of all your hardware and software needs.

To be continued.

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

Before you open all those strange boxes, call us toll-free for an overview brochure and the number of your nearest sales office.

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

Mentor Graphics®

Your ideas. Our experience.

A New VME Enclosure with Benefits That Really Stack Up:

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

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

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

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

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

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

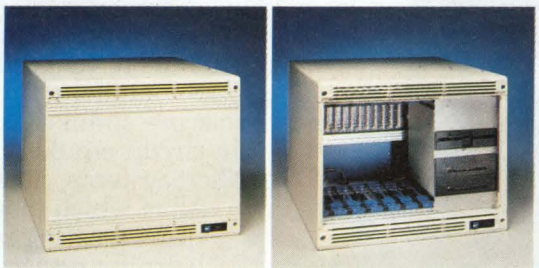

We'll FAX you the facts.

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

ZERO **Electronic Solutions**
UNIT OF ZERO CORPORATION

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

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

CIRCLE NO. 1

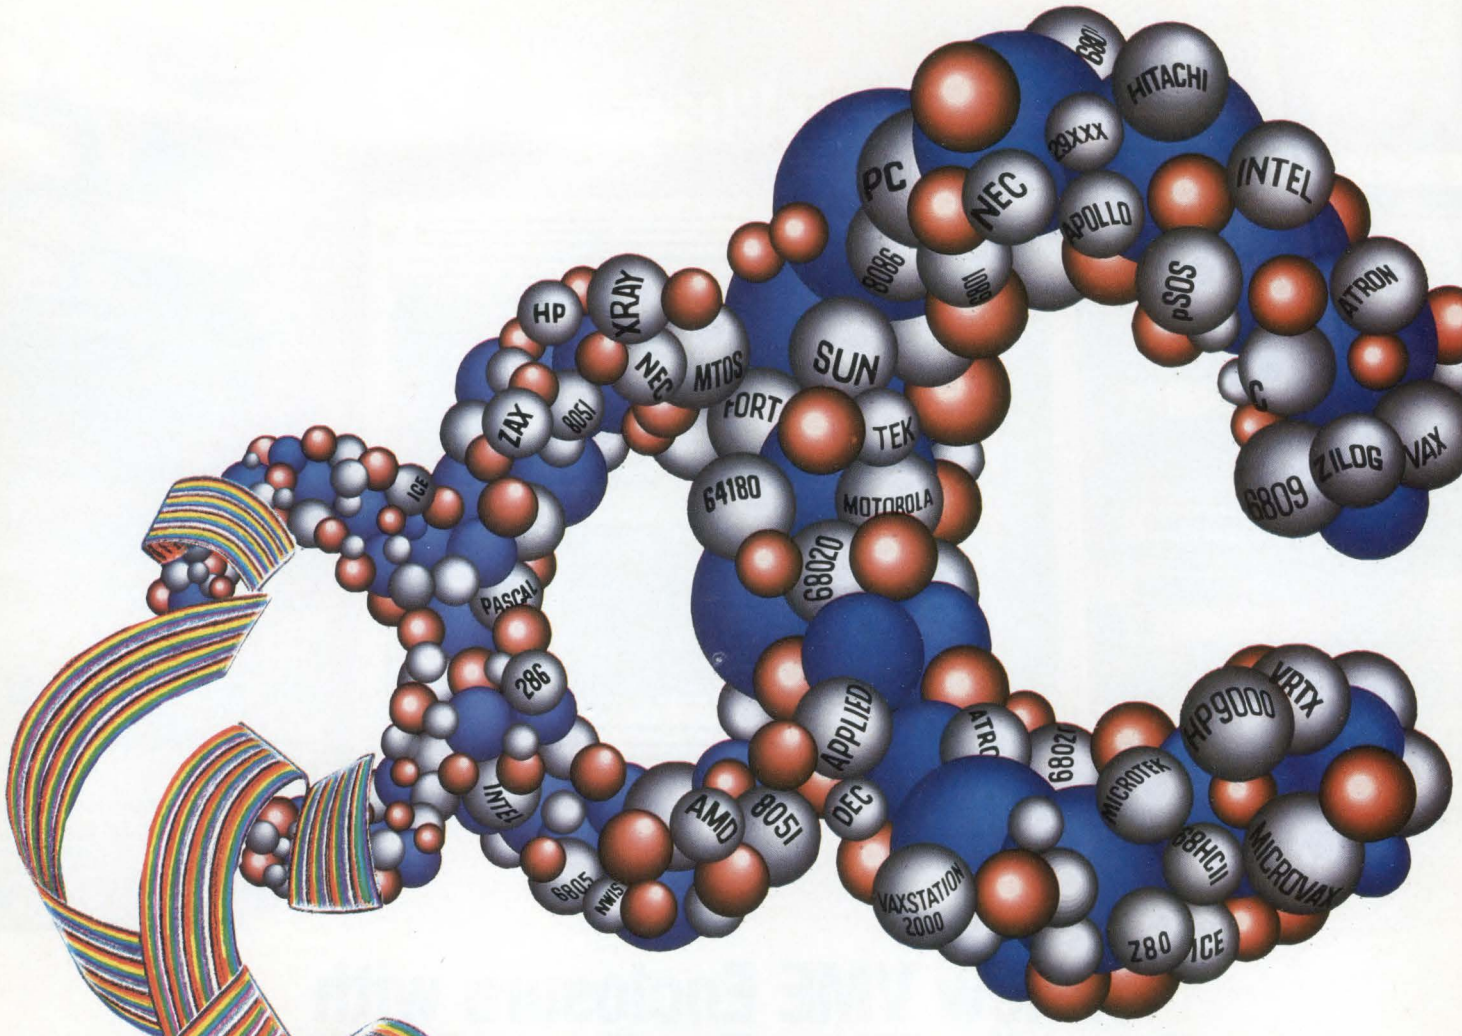

MICROTEC CROSS-COMPILERS DEBUGGERS ASSEMBLERS: THE DNA OF SUCCESS

Microtec Research has done a little genetic engineering on the cross-software development industry. Sure, our quality and customer support have become the industry cross-development standard. And what other company has the confidence to actually publish all fixes and ECOs for all products, for everyone to see? But our real legacy to you, our customer, is our commitment to providing *the same* cross-development software from more host processors to more target systems. From pSOS to VRTX to MS-DOS to UNIX to VMS. From assembler to C to Pascal to Fortran. From Motorola to Intel to Hitachi to AMD to Zilog. From IBM to DEC to HP to Apollo to Sun. From crossassemblers to cross compilers. From source debuggers running on hosts that are simulating target systems to the advanced in-circuit emulators from HP, TEK, Atron, Applied Microsystems, Sophia, ZAX and Taiwan's Microtek. All of which means you can spend your valuable time creating new silicon wizardry, rather than learning new cross-development user interfaces. Call Microtec at 800-551-5554 today (408-733-2919 in California). So you can do some genetic engineering on a competitor's physiology tomorrow.

MICROTEC® RESEARCH

3930 Freedom Circle • Santa Clara, CA 95054 • 800-551-5554 or 408-733-2919 in California

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

CIRCLE NO. 2

Vol. 27, No. 20

COMPUTER DESIGN

NOVEMBER 1, 1988

THE FIRST MAGAZINE OF SYSTEM DESIGN, DEVELOPMENT AND INTEGRATION

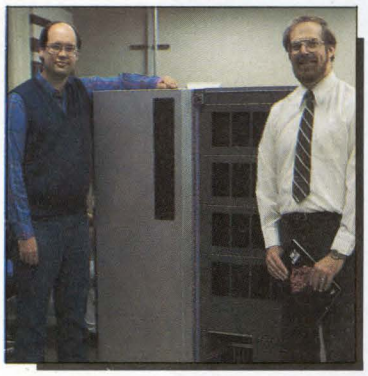

Page 23

■ TECHNOLOGY UPDATES

Computers and Subsystems

Parallel machine extends scalability to encompass I/O processing 23

IPI controller board talks to two disks at once 25

Integrated Circuits

Process advances bring new price, performance to GaAs 30

British microprocessors: a lesson in innovation 32

Acorn processor opens second front in battle for RISC acceptance 42

■ DEPARTMENTS/COLUMNS

News Briefs 11
Editorial 16
Technology Market Directions 112
Advertisers Index.....115
Last Bits.....116

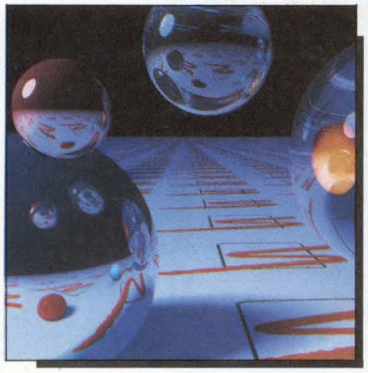

Page 47

■ TECHNOLOGY AND DESIGN FEATURES

Old and new contenders vie for position in 32-bit bus arena

The battle between the new 32-bit buses—particularly Micro Channel and NuBus—heats up, while VMEbus seems to have triumphed over Multibus II 47

COVER STORY

Robotic systems learn through experience

Real-time software, intelligent sensors, image-processing subsystems and neural networks let robots "think" and interact with their environments.....54

CMOS RAM brings industrial-strength data storage

Battery-backed memory cartridges are likely to become the storage medium of choice among designers of industrial systems 73

Electronic memories come in tough little packages

IC memory cards offer a fast, reliable, compact memory alternative that will open up new applications.....79

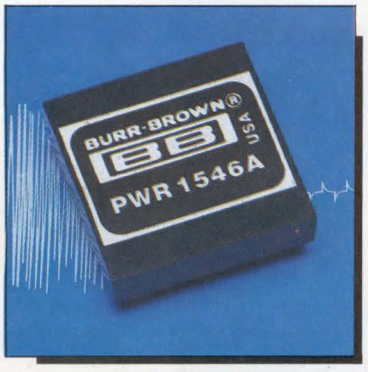

Page 85

■ NEW PRODUCT HIGHLIGHTS

PRODUCT FOCUS

Compact dc-dc converters serve distributed power applications.....85

Computers and Subsystems

DSP boards extend STD Bus speed to 10 Mips 101

Design and Development Tools

Software enhancements speed MC88000 development on Mac II 101

25-MHz microcontroller yields 64-Mips performance 102

PC-based system supports commercial in-circuit programming.....104

Memory Systems

First 2½-in. Winchester drive offers 20-Mbyte capacity.....105

Major System Components

Twelve-bit A-D converter pushes performance to 10 Msamples/s.....106

Testing and Manufacturing

Digital oscilloscopes feature 20-GHz sampling speed 107

Software

Compiler brings C language to custom microprogramming 108

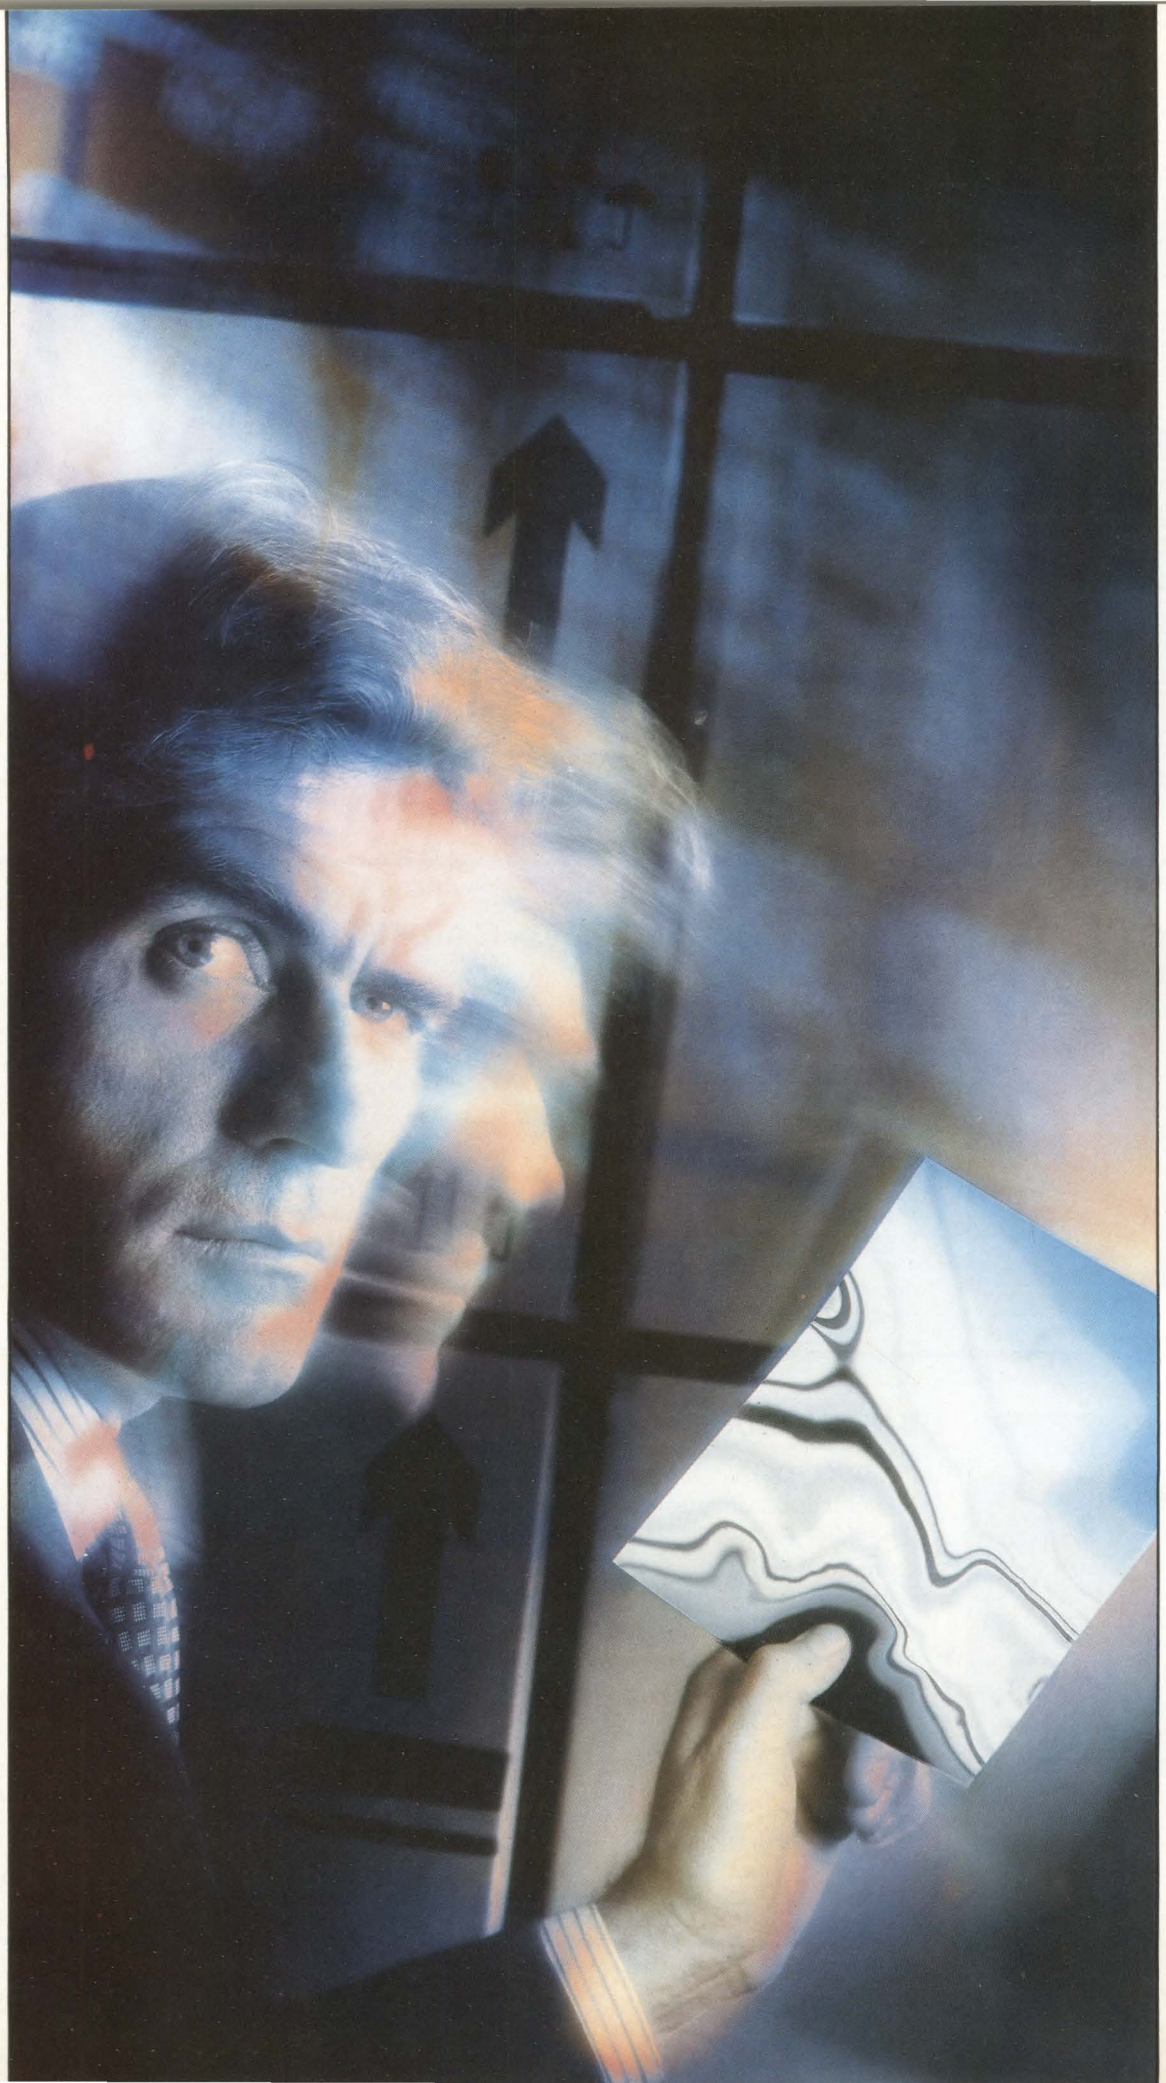

WE SPENT

10 MILLION DOLLARS ON A PIECE OF PAPER

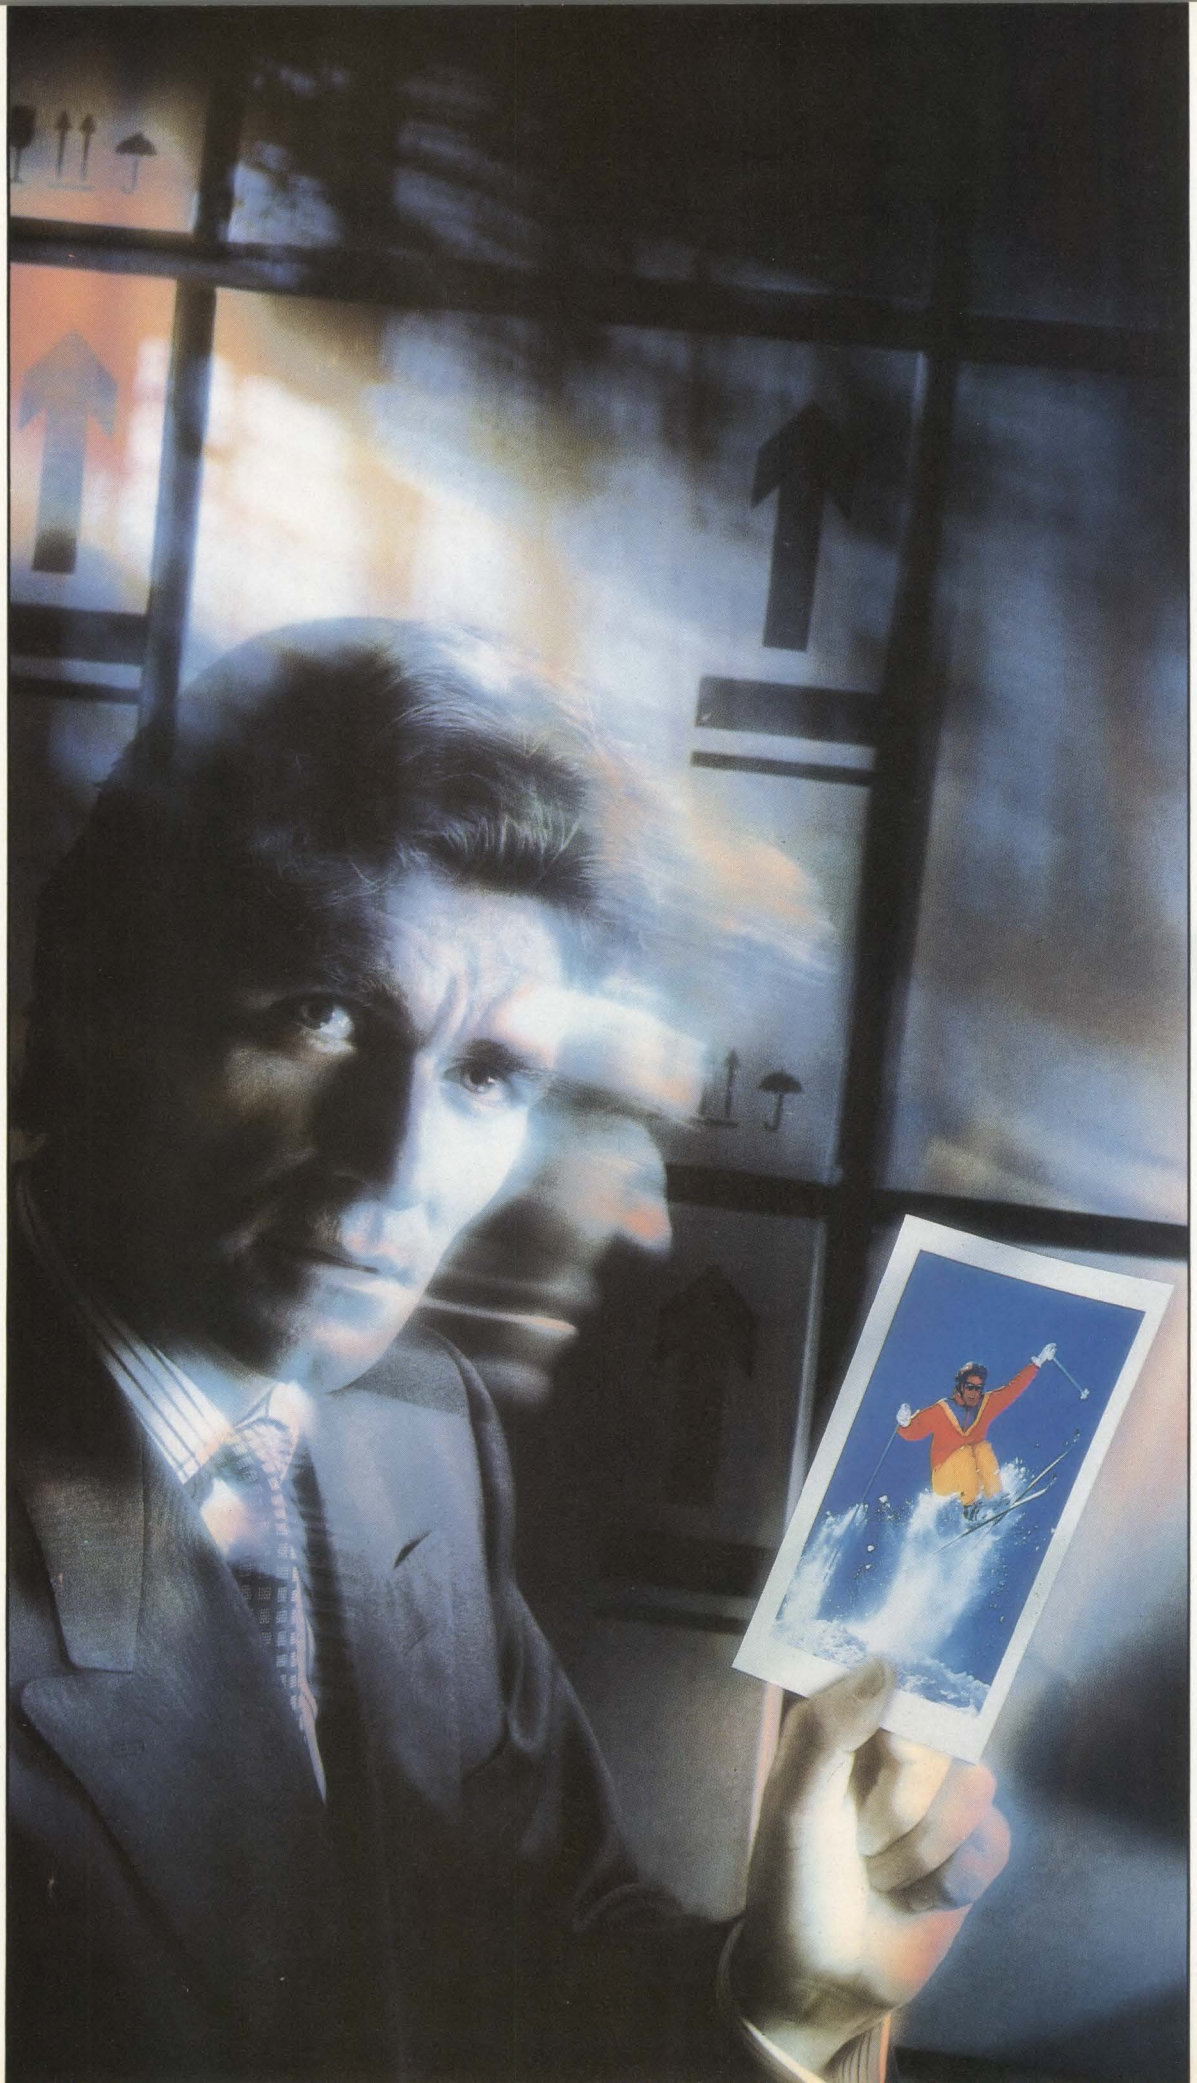

WE TOOK

150,000 HOURS TO DEVELOP A PICTURE

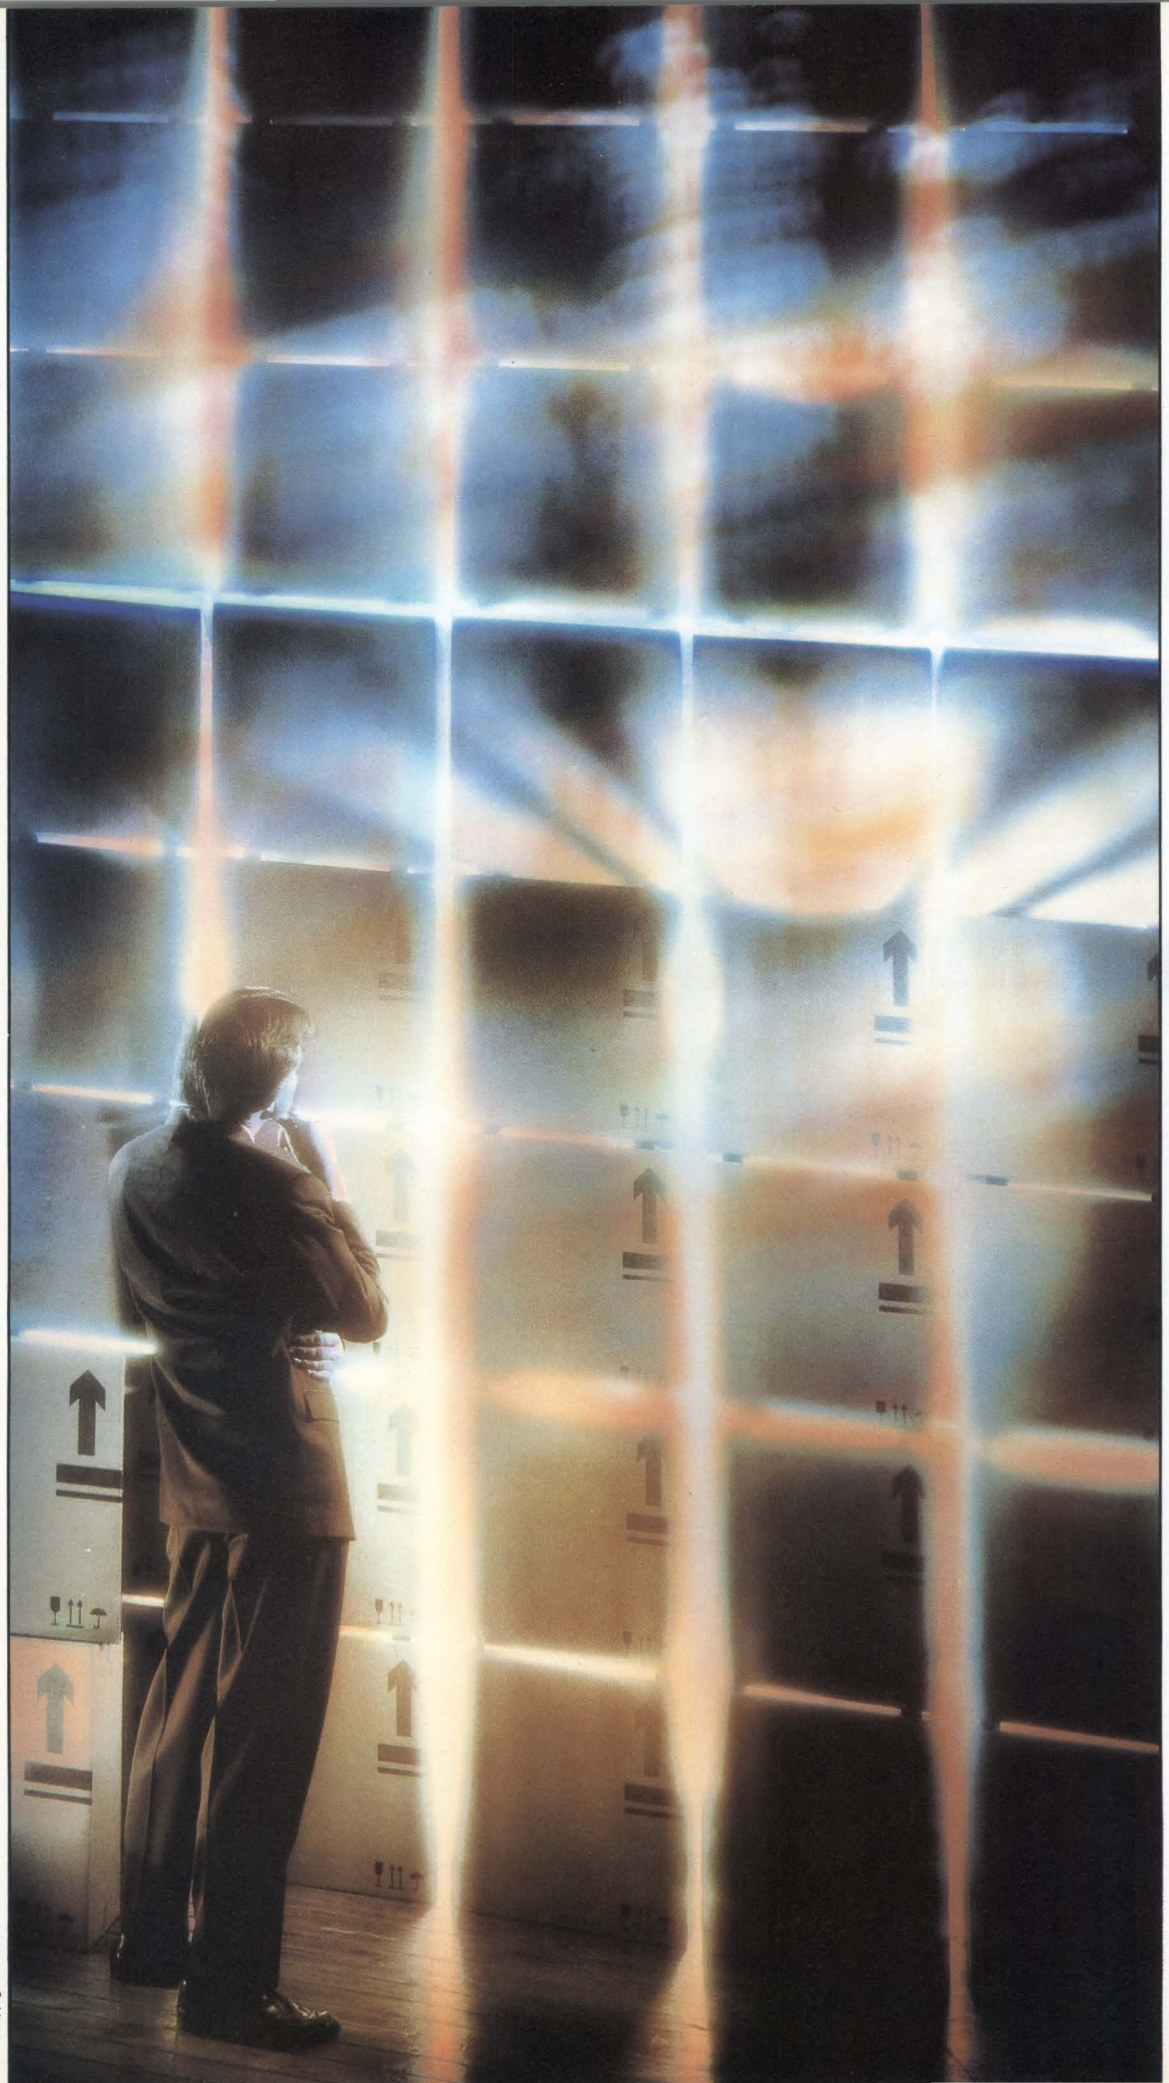

NOW WILL

©Copyright 1988 of ICI

YOU TAKE A SERIOUS LOOK AT THE FUTURE OF ELECTRONICS?

We predict the big money will be in hardware.

But only if it's of a kind suited to the most advanced electronic media.

That's where ICI Imagedata is currently leading the way, with two very important developments.

The first is *Digital* Paper, the data storage equivalent of traditional paper. It's low cost, indelible and permanent. Yet, unlike paper, Digital Paper offers the areal capacity and recording advantages of optical disks.

One 12" reel made from this unique material holds up to two Terabytes of data (that's 2,000,000,000 Kbytes), on-line if you like.

What's more, it's flexible. We can cut Digital Paper into disks of any size (one 5¼" disk provides 1.5 GB of data, fast). Chop it into tag form. Or use it as a tape for mass storage.

Our second development is no less impressive.

It involves a colored ribbon for use in a thermal printer, plus a new receiver paper which won't tear, scratch, curl or run.

The process is called dye diffusion thermal transfer (D2T2 for short). It produces brilliant color images from an electronic source (video camera, computer, fax etc.). The final picture looks as good as a photographic print (with a choice of millions of colors for every dot), even though it's produced on a simple desk top printer.

So what will these materials mean to you, the hardware manufacturer, software developer or system integrator?

A great deal.

Potential markets for both are enormous . . . everything from data networking and logging, on- and off-line storage and recording to publishing, computer graphics and business presentations.

ICI Imagedata is already inspiring new designs of disk drives, tape drives and thermal printers.

In fact, ICI spends $2.5 million per day on research and development.

But your contribution in electronics will further improve the performance standards we've set. That's why we'd like to discuss the future with progressive hardware and systems developers.

For a detailed brief on the technology and media of the future, call Rick Lamb (D2T2) or David Owen (Digital Paper) toll-free at 1-800-456-3669.

Or write to us at ICI Imagedata, Concord Pike, Wilmington, DE 19897 and we'll send in-depth information.

You'll learn in an hour what's taken us considerably longer to develop.

ICI **Imagedata**

CIRCLE NO. 5

68020 EMULATION NOW AT 25MHz.

68020-25MHz

68000-16.7MHz

80286-12.5MHz

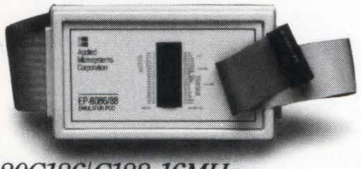

80C186/C188-16MHz

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

So now emulating faster moving targets is getting easier.

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

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

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

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

Applied Microsystems Corporation

CIRCLE NO. 6

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

AMC 247

Next is finally here, sporting on-board DSP

The long-awaited workstation from Next (Palo Alto, CA) has finally been unveiled by Steve Jobs, whose company has been developing the system under close security for the past three years. The entire system is on a single 1-ft² circuit board that carries 8 to 16 Mbytes of RAM plus 1 Mbyte of video RAM. The board fits in a box containing three empty NuBus slots. The CPU is a Motorola 68030 and a Motorola 68882 floating-point unit running at 25 MHz. A 10-Mips Motorola 56001 digital signal processor is also provided as standard equipment.

Key to the new architecture are two specially designed VLSI chips. One, the integrated channel processor (ICP), provides 12 separate DMA channels to manage all I/O in the system. It lets all devices on the system interact with memory without bothering the CPU. The other device is the optical storage processor (OSP). The OSP manages the interface between the system's standard 256-Mbyte read/write optical drive and the DMA channel assigned to the disk. It also manages all error-correction code for the disk, as well as managing two data buffers between disk and DMA channel that work simultaneously.

The Next workstation is the first commercial system to incorporate Display Postscript from Adobe Systems (San Jose, CA). A graphics user interface based on Display Postscript is coupled with a graphics-based, object-oriented application development environment.
—Tom Williams

Sparc partners hail chip for embedded control and real-time applications

Technology partnerships may prove to be the best strategy for Sun Microsystems (Mountain View, CA) in its bid to push the Sparc chip to the top of the RISC heap. Sun's latest target is the real-time/embedded control marketplace. In a tribute to the processor by Sun, Mizar (St. Paul, MN), Wind River Systems (Emeryville, CA) and Fujitsu Microelectronics (San Jose, CA) at the recent Buscon/88-East trade show, a spokesperson for Fujitsu predicted that 73 percent of all RISC designs will be in embedded applications by 1992.

Highlighting the event was the announcement of a real-time, Sparc-based application prototyping board from Mizar. The board is the first in a series of products expected to grow every four months. All the participants lauded the Sparc's fast interrupt response, context switching, data movement and processing capabilities. But the strongest argument for the processor may have come from Mizar president Joe Ramunni: "The preoccupation is with the architecture. What is important is that it works."
—John Mayer

One-Mbit flash EPROMs almost a reality

Although the in-circuit programmability of flash nonvolatile memories offers design engineers attractive cost advantages in testing and manufacturability, the relative newness of the technology has limited product acceptance. Look for the market to increase as vendors begin to unveil 1-Mbit parts with pin-outs compatible with the UVEPROMs they're aiming to replace.

At present, Seeq Technology (San Jose, CA) offers a 512-kbit part, and Toshiba (Irvine, CA) and Intel (Santa Clara, CA) have 256-kbit chips. But Seeq plans to introduce early next year both a 1-Mbit flash EPROM and a 1-Mbit flash EEPROM in the same 32-pin design used by 1-Mbit UVEPROMs. Toshiba and Intel have also indicated that they will jump to the 1-Mbit density in the same time frame.
—John Mayer

More fog over next-generation VMEbus

Reports from the latest meeting of the Next Generation Architecture (NGA) working group of the VME International Trade Association (VITA) did little to clarify what the follow-on bus architecture for the VMEbus will look like. Nor is there any indication yet whether the group will recommend migrating to some existing high-performance bus like Futurebus, extending the VMEbus to incorporate desired enhancements or developing an entirely new bus structure that subsumes the VMEbus. Neither is it yet known, as a group spokesperson puts it, how the next-generation bus will "maintain the integrity of VME," nor exactly what that phrase means.

The NGA group has, however, formalized its bus specification objectives, which include coexistence with VME products, a 256-bit data path, 2 to 3 Gbytes/s of bandwidth, cache coherency support, 64-bit addressing, scalability and fault tolerance. The specs are expected to firm up around mid-1989, with product availability to come sometime around 1991.*—David Lieberman*

Tektronix workstations run Sun-3 software without porting

Users of the Series 4300 workstations from Tektronix (Beaverton, OR) will discover that they have access to hundreds of applications written for Sun-3 workstations from Sun Microsystems (Mountain View, CA). Tektronix's UTek Unix-based operating system, like the Sun-3 OS, is based on the Motorola 68020 and the Berkeley version of Unix. Tektronix has built in a Sun-3 system call capability that lets users access Sun-3 binary series public domain programs, commercial software with Tektronix drivers, in-house applications, and graphics applications that are supported under X Window.

Programs that use proprietary extensions to the Sun-3 OS, such as NeWS or Sun Tool, or that try to use a mixture of Sun and Tektronix libraries in a software development environment won't run. However, a large number of commercial applications that run on the Sun-3 will run directly on the Tektronix 4300 systems. In addition, the compatibility gives third-party software developers access to a larger number of platforms without any porting effort.*—Tom Williams*

(continued on page 12)

Defining an ORKID is no bed of roses

At a recent meeting in Zurich, Switzerland, the VME International Trade Association (VITA) Software Subcommittee continued to wrestle with the development of the Open Real-time Kernel Interface Definition (ORKID). The committee focused its efforts on refining the draft document developed by Wind River Systems (Emeryville, CA), which will form the basis for a standard to bring software interoperability to real-time VME computing environments. With so many players with vested interests in the game, it looks like there may be no ORKID for some time to come, or, as one industry observer suggests, there may not be an agreement until the appearance of the Next Generation Architecture (NGA) bus, which is currently under development by another VITA committee.—*Michael Donlin*

TI introduces umbrella graphics standard for 340X0 boards

Vendors building graphics controller boards that use the 34010 or 34020 graphics processors from Texas Instruments (Dallas, TX) have had to choose from a variety of software interface protocols for resolutions above 640×480 pixels. Both board and applications vendors have then had to worry about supplying drivers that would let applications run on boards using different protocols. The TI Graphics Architecture (TIGA) specification is designed to offer a common programming interface for all graphics boards using the 340X0 processor family.

Software vendors will now be able to write extensions to their existing protocols; these extensions will let existing applications written to those specs run unaltered on TIGA-based boards. TIGA offloads certain functions from an 80286 or 80386 processor and distributes the graphics load between the host CPU and the TI graphics chip, according to Ahmed Nawaz, TI marketing manager. He reports that in a comparison between an IBM Model 70 with an 8514/A controller and a Compaq 386 with a TIGA-based board, the TIGA board did a redraw of an AutoCad drawing 16 times faster than the 8514/A.—*Tom Williams*

NuBus approaches form-factor hurdle

Time will soon tell how the burgeoning NuBus community deals with its first major internal controversy: the form factor question. Although NuBus already accommodates two factors (one, the "classic" Eurocard 9U form factor used by Texas Instruments; the other, the "desktop" form factor used by Apple Computer), a heated debate is brewing over the practicality of adding a third form factor: Eurocard 6U.

Many of the companies involved in industrial control are discontented with the rectangular shape of the desktop NuBus board (vs. the "squared-off" 6U), as well as with its limited room for extra connectors. "In our environment, it would be very advantageous to have a P2 connector so we could take transducer and control signals off the backplane instead of off the front panel," says Jim Eckford, vice-president of Ziatech (San Luis Obispo, CA). "We could also make good use of P2 to implement a VXI-bus for instrument control." Some believe, however, that a proliferation of form factors will only confuse the marketplace and will hinder general acceptance of the NuBus. —*David Lieberman*

DEC presses to own more of CIM market

Not satisfied with its already dominant base for computer-integrated manufacturing systems, Digital Equipment Corp (Maynard, MA) has announced a strategic business alliance with Allen-Bradley (Milwaukee, WI) that, in effect, combines a DEC MicroVAX computer and an Allen-Bradley programmable logic controller.

The Allen-Bradley agreement is the first in which DEC has licensed the integration of any VAX processor for a jointly developed product for commercial use. Allen-Bradley now has a programmable controller market share of over 40 percent. The single-chassis Pyramid Integrator combines DEC's VAX architecture, VMS software environment, and DECnet/OSI networking with Allen-Bradley's Pyramid architecture. To maintain compatibility with its Vista 2000 industrial computer products, Allen-Bradley will also manufacture a module that offers Unix functionality. —*Sydney Shapiro*

DEC looks outside for PC

Digital Equipment Corp (Maynard, MA) has announced that it will look to Tandy (Fort Worth, TX) for an IBM-compatible personal computer. The machine will reportedly be built by Tandy to DEC specifications. After years of attempting, and failing, to develop competitive desktop computing systems, DEC has now twice gone outside the company to obtain computing technology. The Tandy announcement closely follows DEC's decision to obtain microprocessors for desktop Unix workstations from MIPS Computer (Sunnyvale, CA).—*Ron Wilson*

CD-ROM on the move in optical publishing

Despite some feeling that the CD-ROM industry hasn't reached its expectations for optical publishing applications, a study from Infotech (Pittsfield, VT) indicates that reasonable projections of a market valued at $80 million in 1987 show an increase to $150 million in 1988. "We feel earlier expectations were based on scientific wild guesses that were made before there were any products available," says Richard Bowers, director of development for the Applied Information Technologies Research Center (Columbus, OH) and executive director of the newly formed Optical Publishing Association. "The Infotech prediction is conservative enough to be realistic. And any industry that can get to $150 million per year from a dead start two-and-a-half years ago isn't bad."—*Sydney Shapiro*

Object-Oriented Engineering

If you don't pick the right environment, you may come up dry.

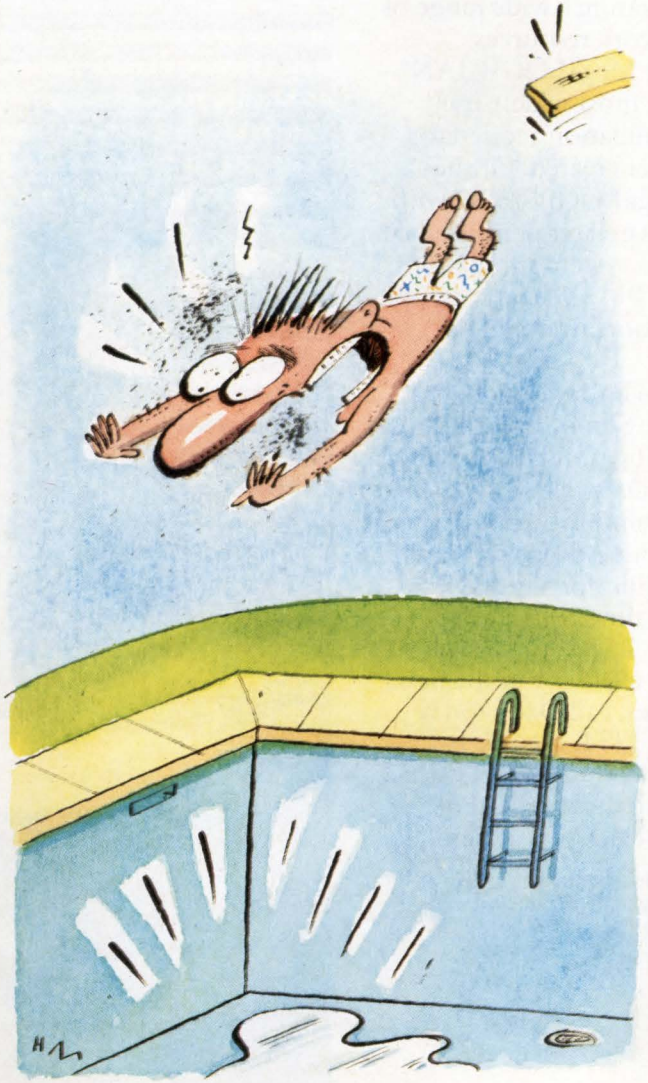

You've probably heard about object-oriented engineering. And the advantages it brings to software development.

You may even be ready to dive into it. We suggest you look before you leap.

Many companies today only offer one small bit of the support you need.

Not Stepstone. We give you more than just a compiler. We supply a whole object-oriented environment complete with powerful interpreter/debugger.

So you can shorten development schedules. Reduce maintenance costs. Build code which directly models your problem domain. And build truly re-usable software components.

Our products are based on the well known Objective-C® language and are designed for programmers who want the benefits of this new technology without giving up the advantages of C.

We even offer more than just a development environment.

We sell libraries of powerful, robust classes you can use directly in your applications. Like ICpak™ 201, a set of pretested Software-IC®s for building custom graphical user interfaces.

Our products operate in a wide variety of work-station environments. As well as on the PC-AT™ and compatibles.

In addition, our support is second to none. We design, build, market and service our own products. Hence, we will do all that's necessary to keep you completely satisfied.

So before you jump into object-oriented engineering, talk to the company that's got the fullest range of products. Stepstone.

We won't let you down.

Stepstone™
The Leader in Object-Oriented Technology

The Stepstone Corporation 75 Glen Road Sandy Hook, CT 06482 203 426 1875 1 800 BUY OBJECT (1 800 289 6253)

PC-AT is a trademark of International Business Machines Corporation. Objective-C and Software-IC are registered trademarks of The Stepstone Corporation. Stepstone and ICpak are trademarks of the Stepstone Corporation.

Daisy's standard now runs on an

Introducing the Advansys Series™ of high performance CAE workstations.

Now the electronic design environment you've dreamed of is here.

Because the most advanced design tools in the world now run on the world's most advanced standard platform.

The Sun 386*i*.™

We call it the Advansys Series. And it encompasses some of the most powerful design tools ever developed. Plus a variety of affordably priced workstations. Including the 20 MHz and 25 MHz Sun 386*i*, as well as Daisy's own LOGICIAN® 386 and the newly enhanced Personal LOGICIAN™ 386. All share a standard system level environment, featuring UNIX,™ advanced X Window System graphics, Sun's NFS™ distributed file system and standard TCP/IP communications.

Now you can get the workstation performance and flexibility you've demanded for your desktop. Including up to 5 MIPS of processing power, high resolution graphics display and an integrated UNIX/DOS environment.

Even better, all these advanced workstations run Daisy's field-proven Advansys software packages. Eight turnkey tool sets that meet all the demands of real world electronic design. Everything from design entry to digital and analog simulation, IC and PCB layout, fault simulation and test tools.

But with Advansys, your capabilities don't stop at your desktop. Because Daisy's unique network computing concept lets you create an affordable team design environment incorporating a wide range of powerful network resources.

Like Daisy's MegaLOGICIAN®—the most widely used simulation accelerator ever created. Or the brand new GigaLOGICIAN™ with up to 30 times greater performance.

For complex system simulations, there's Daisy's PMX,™ the most popular physical modeling system in use today.

You can also link with network servers like the Sun-4™ based XL Server for analog simulation or PCB routing. And Daisy lets you access all this power simply by opening a window on your Advansys desktop workstation. Eliminating file transfers and other time consuming bottlenecks.

To find out more about the new Advansys Series, call Daisy today at 1 (800) 556-1234, ext. 32. In California: 1 (800) 441-2345, ext. 32.

We're raising the standard of excellence for electronic design.

European Headquarters: Paris, France (1) 45 37 00 12.
Regional Offices: England (256) 464061;
West Germany (89) 92-69060; Italy (39) 637251.

daisy SYSTEMS CORP.

# of excellence excellent standard.

© 1988, Daisy Systems Corporation. Sun386i, NFS and Sun-4 are trademarks of Sun Microsystems, Inc. Ethernet is a trademark of Xerox Corporation. UNIX is a trademark of Bell Laboratories.

Bringing it all together

Together, Computer Design magazine and Computer Design News will provide the one-two punch you need to be successful.

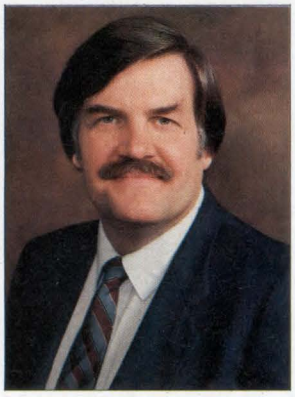

John C. Miklosz
Associate Publisher/
Editor-in-Chief

Since its inception 26 years ago, *Computer Design* magazine has devoted itself to technology, design issues and critical system components—in essence, to those factors that have the greatest impact on system performance and integrity. We took this approach because we appreciated the key role that senior design engineers and design engineering managers play in making crucial technical decisions. What's more, we believed that to make these decisions, our readers needed information that went beyond straightforward design details.

Design managers need information that gives them better insights into all the trade-offs they have to make. They need to know the advantages and disadvantages of a particular technology, system component or design approach. They need to know the directions in which key technologies are moving. And they need to know something about test, manufacturing and purchasing. But perhaps most important of all, they need to know about cost, cost/performance trade-offs, quality, source evaluation and selection, reliability, and everything that affects time-to-market.

In this need to know, design engineering managers aren't alone, however. Every other senior member of the product development team—test engineering managers, manufacturing engineering managers and purchasing managers—must also be concerned with cost, quality and time-to-market. Even though these concerns are universal, no publication serving the electronics/computer OEM community, either engineering- or business-oriented, has given them much play.

It's been a serious deficiency, and product development teams have been shortchanged by the neglect. Until now.

This month, the Advanced Technology Group of PennWell Publishing will introduce *Computer Design News*. The new publication will complement the technology and design coverage you find in *Computer Design* magazine. We'll bring together the key members of the system team—design managers, test managers, manufacturing engineering managers and purchasing managers—by focusing on the primary concerns they have in common: cost, quality and time-to-market.

Alternating every two weeks with *Computer Design*, *Computer Design News* will deliver the scoop on the latest developments at vendors and distributors; on the latest management approaches used to keep projects on track; on new products, with all of the details you need to get the information you need—fast; on the rules, regulations and procedures imposed by the military and the federal government; and on salaries, benefits and career opportunities.

We think that, together, *Computer Design* magazine and *Computer Design News* will provide the one-two punch you need to be successful. First, hitting hard with the "why-to" approach to technology and design that you'll find in every issue of *Computer Design*, and then following up with the essential "how-to": how to reduce costs, improve quality and shorten time-to-market.

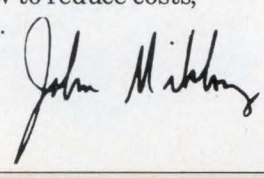

Disk-Caching SCSI for Multibus* II.

The CD22/4500 SCSI Controller from Central Data.

Central Data is committed to a leadership role in the Multibus II market. With special emphasis on SCSI support.

The CD22/4500 provides the ultimate Multibus II SCSI solution. It's fast, transferring data at the limits of the SCSI bus. And versatile, providing either direct SCSI commands or Intel compatible PCI commands.

On-board disk-caching firmware makes it even faster. Up to 2 megabytes of parity protected RAM and 80186 CPU speeds of up to 12.5MHz mean cache hits are frequent and fast.

For extra efficiency, tune your systems using the four provided disk-sorting algorithms, automatic read ahead and selectable write back or write through policies.

Design excellence, unparalleled account service, easy access to design engineers, even customizing for some applications. You'll find them in this and every Central Data product.

Call product manager Andre Felix today for a detailed product brochure on the CD22/4500.

Central Data

1602 Newton Drive, Champaign, IL 61821-1098

1-800-482-0315
(In Illinois 217-359-8010)
FAX 217-359-6904

*Multibus is a trademark of Intel Corporation.

CIRCLE NO. 8

We've Moved...

Computer Design's home office has moved to new quarters. Please change your records as follows:

PennWell Publishing Company
One Technology Park Drive
P.O. Box 990
Westford, MA 01886

Telephone: (508) 692-0700

Our address for subscription information, change of address forms, etc., stays the same: Computer Design Circulation Department, P.O. Box 3466, Tulsa, OK 74101, (918) 831-9401.

This issue's cover. . .

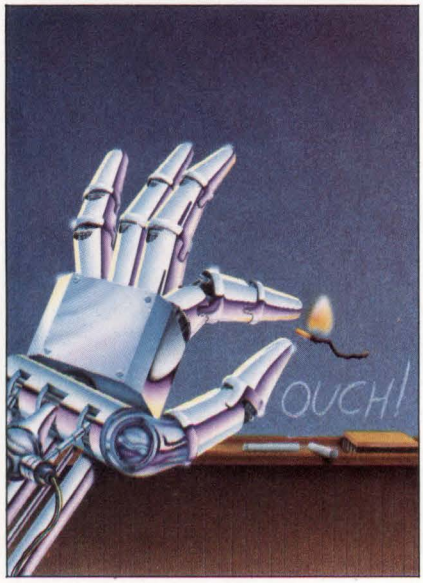

Cover design by Sergio Roffo

Robotic systems make great mechanics, so they're taking on chores in all kinds of manufacturing environments. While they're great mechanics, though, they're poor thinkers. In fact, they're downright dumb. They can only do what they're programmed to do, and if something goes wrong, they can only respond by stopping dead in their tracks or by continuing to do the same task, regardless of the consequences. New approaches that make use of exotic concepts, such as neural computers that mimic the human brain, could change all of that. Future robotic systems might be able to tell you what's wrong. They might even complain about it—which would really make them like humans.

COMPUTER DESIGN®

THE FIRST MAGAZINE OF SYSTEM DESIGN, DEVELOPMENT AND INTEGRATION

PUBLISHER
David L. Allen

ASSOCIATE PUBLISHER/EDITOR-IN-CHIEF:
John C. Miklosz

TECHNOLOGY EDITORS:
Ron Wilson, *Integrated Circuits*
Warren Andrews, *Integrated Circuits*
David Lieberman, *Computers & Subsystems*
Tom Williams, *Graphics & Imaging*
John H. Mayer, *Computers & Subsystems*
Sydney F. Shapiro, *Data Acquisition & Control*

CONTRIBUTING EDITORS:
S. Louis Martin, *Integrated Circuits*
Bill Harding, *Design & Development*
Howard Falk, *Software*
Art DeSena, *Testing & Manufacturing*

RESEARCH/SPECIAL PROJECTS MANAGER:
Sydney F. Shapiro

SECTION/SPECIAL PROJECT EDITORS:
Ron Wilson, *Technology Updates*
David Lieberman, *Contributed Articles*
Michael G. Donlin, *New Products*
Patti Kenney, *Directories*

MANAGING EDITOR: Susan R. Nykamp

ASSISTANT EDITOR: Kerstin L. Rosenquist

COPY EDITORS: Richard Dagley, Barbara-Ann Scofidio

EDITORIAL ASSISTANTS: Claire Coupal, Claire Ellis

WEST COAST EDITORIAL OFFICES
Tom Williams
75 Mount Herman Rd, Suite E-180
Scotts Valley, CA 95066, Tel: (408) 335-5514

Ron Wilson
1111 S.W. Gaines, No. 8
Portland, OR 97201, Tel: (503) 224-9396

ART DIRECTOR: Jan Horner
ILLUSTRATOR: Christopher Hipp

PRODUCTION DIRECTOR: Jan M. Lopez
PRODUCTION MANAGER: Mari Rodriguez
ADVERTISING COORDINATOR: Adrienne Adler
COMPOSITION MANAGER: Holly Kersey
COVER ART: Sergio Roffo

CIRCULATION DIRECTOR: Robert P. Dromgoole

PennWell PUBLISHING COMPANY

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

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

Fortunately, you don't have to invent something to perfect it.

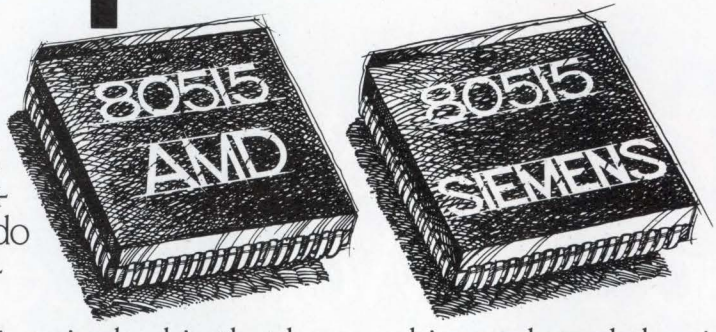

The 8051 was fine. Except that you had to add lots of extra chips to make it do everything you wanted it to do.

But now there's a single chip that has everything built in. It's the 80515 designed by Siemens.

With the 80515 you get the 8051 architecture you've depended on. Plus: An A/D converter, a watchdog timer, six 8-bit ports, 256 bytes RAM, 8K ROM and a pulse width modulation timer.

If you've ever tried to stay on a budget you can see how much money all these on-chip features will save.

If you've ever tried to fit four chips where there's only room for one, you'll also recognize a definite advantage. (With surface mount, you can save even more room.)

And with fewer chips on board, there's less chance of a chip failure.

You won't have to spend time learning a new instruction set with the 80515, either. It's fully compatible with the 8051.

One more very important thing you get with the 80515: Two suppliers. Siemens and AMD. Both offer all the support you need to get your project moving. From emulators to FAE's.

For more information on the 80515, give us a call at (800) 222-9323. Or drop us a line.

Advanced Micro Devices
and the
Siemens Semiconductor Group

Advanced Micro Devices, 901 Thompson Place, P.O. Box 3453, Sunnyvale, CA 94088.
Siemens Semiconductor Group, Integrated Circuits Division, 2191 Laurelwood Road, Santa Clara, CA 95054, (408) 980-4527.

©1988 Advanced Micro Devices, Inc.

CIRCLE NO. 9

There's a New Global L
It's Also The World's Leading So

With the recent acquisition of Borg-Warner Chemicals, GE Plastics can now supply worldwide markets with the broadest line of ABS resins in the industry. Combined with our total resources in high-performance engineering polymers, this expansion forms a global product offering that is unmatched in its scope and versatility —a product range totally supported by the technologies which made these ABS resins and engineering polymers the world leaders in sales.

GE Plastics

Now designers and manufacturers all over the world have one source for a wider choice of products and properties to fit specific application needs—products that provide a new dimension of design possibilities and manufacturing efficiencies. LEXAN® resin, NORYL® resin, VALOX® resin, XENOY® thermoplastic alloys, LOMOD® engineering elastomers, CYCOLAC® ABS, PREVEX® polymer,

® LEXAN, NORYL, VALOX, XENOY, LOMOD, CYCOLAC, PREVEX, ULTEM and GELOY are registered trademarks of General Electric Company or affiliates.

™ SUPEC is a trademark of General Electric Company.

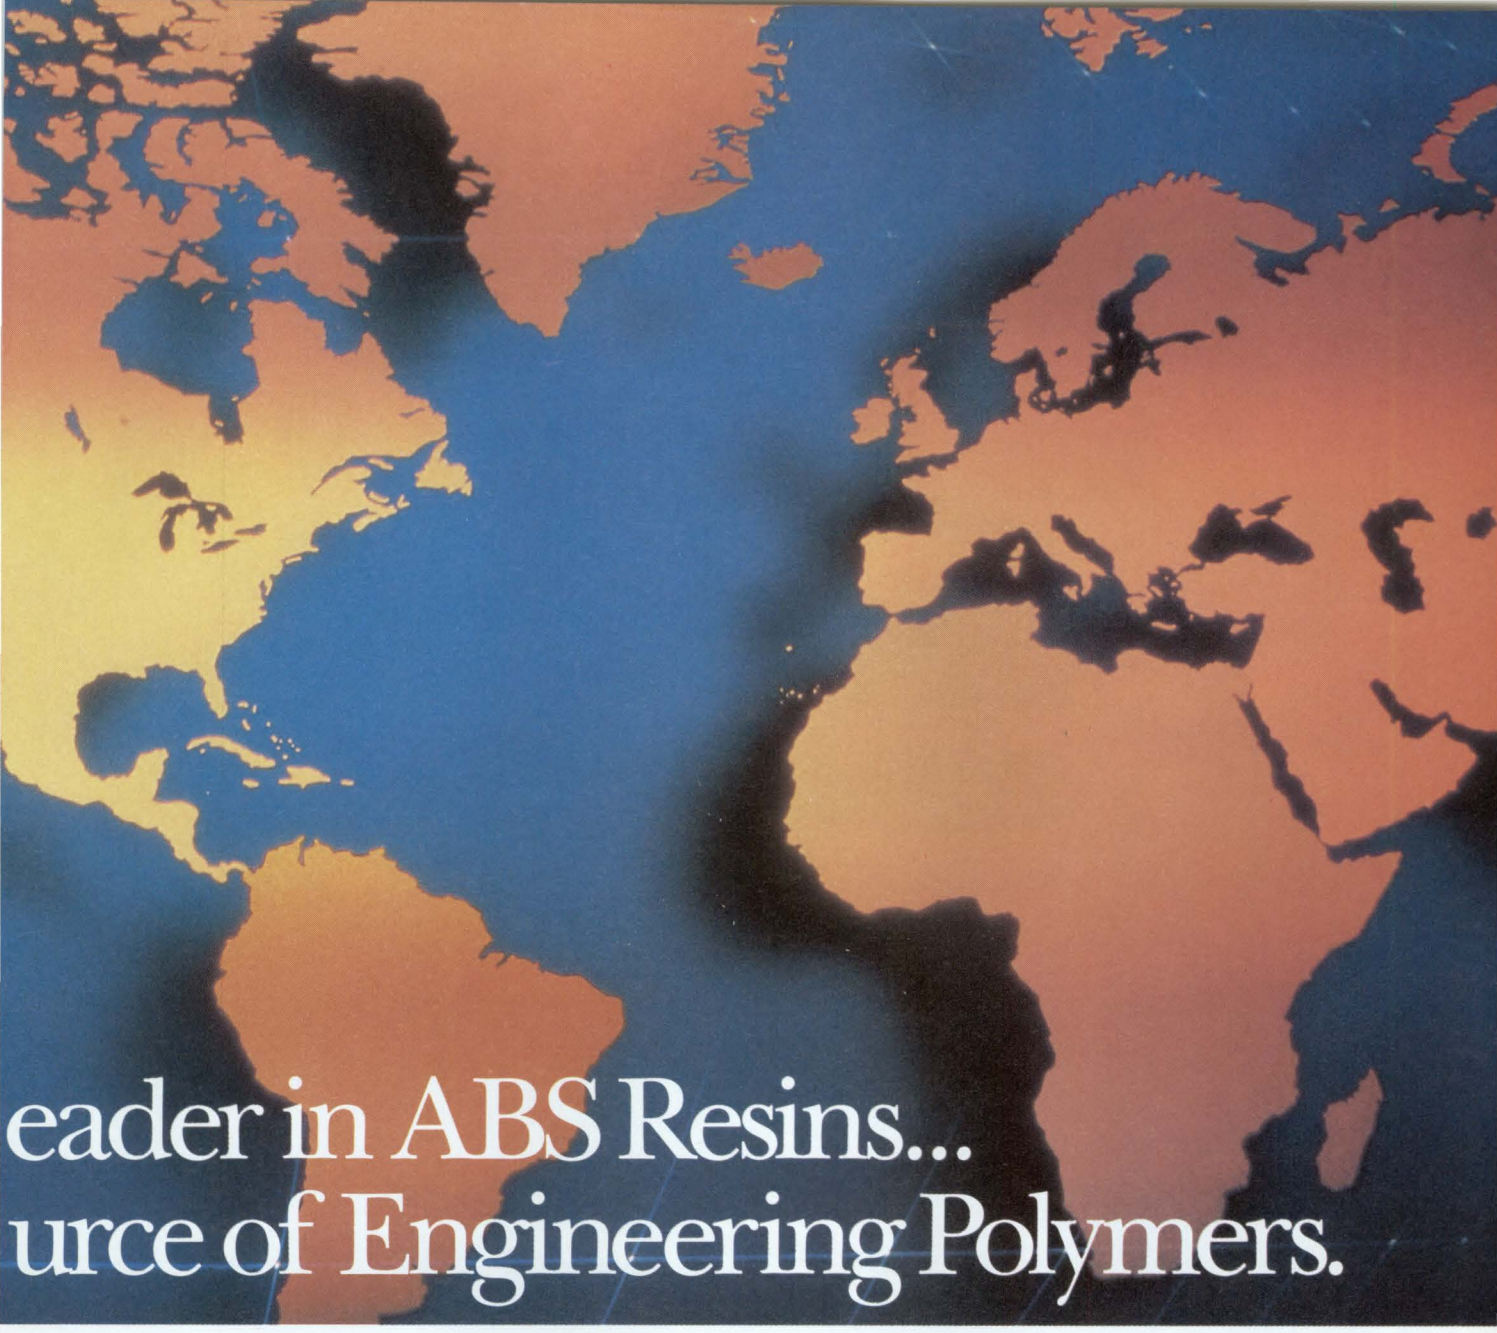

eader in ABS Resins...
urce of Engineering Polymers.

ULTEM®, SUPEC™ and GELOY® resins, and many more.

With a network of technical, R&D and manufacturing facilities and sales offices spanning the globe, GE Plastics customers get fast response, dependable supply and delivery, and hands-on support from the industry's largest technical staff and market development team.

GE Plastics: creating more options, more opportunities—all over the world.

For information on LEXAN, NORYL, VALOX, ULTEM, SUPEC and GELOY resins; XENOY thermoplastic alloys; and LOMOD engineering elastomers, call (800) 845-0600.

For information on CYCOLAC ABS and alloys and PREVEX polymer, call (800) 624-2611 (In Canada, call 416/372-6801).

CIRCLE NO. 10

GE Plastics

Unlimited Service.

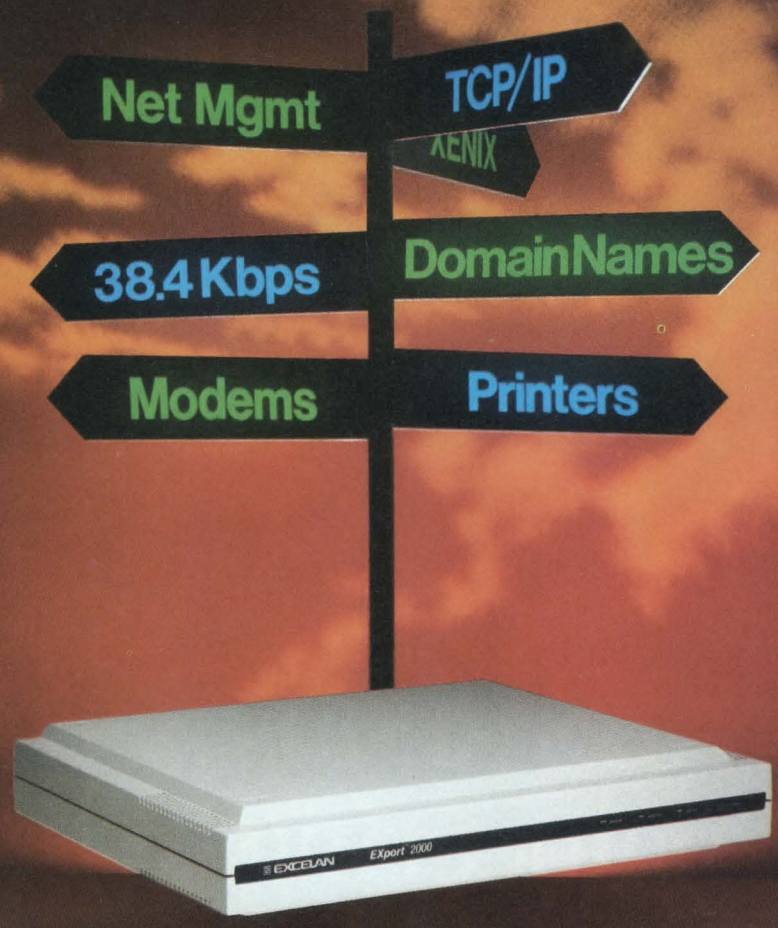

Now there's a server that does much more than connect terminals to your network. With Excelan's EXport™ 2000, you can communicate in a variety of directions. Modems and printers can be shared, as can incoming connections to hosts. With unparalleled flexibility, power, and convenience.

The EXport 2000 server connects all types of RS-232 asynchronous serial devices to Ethernet. And with Network Management Station™ (NMS) software that runs on VMS, UNIX or DOS, you can configure and control all communications servers on your network, no matter where they are located.

Moreover, Excelan's EXport 2000 communications server supports other industry-standard protocols such as TCP/IP, RS-423, and Domain Name Service.

Find out how the EXport 2000 extends Excelan's complete, high-speed LAN solutions for heterogeneous environments. Call today: **1-800-2-GET-LAN**.

CIRCLE NO. 11

EXCELAN

2180 Fortune Drive San Jose, California 95131 800-2-GET-LAN FAX 408-434-2310 Telex 176610
Excelan Europe, Ltd. Intec 2, Suite 17, Wade Road Basingstoke, Hampshire, RG24 0NE England, UK
Tel: (44)256-842296 FAX (44)256-842298

Excelan is a registered trademark; EXport and NMS are trademarks of Excelan, Inc. VMS is a trademark of Digital Equipment Corporation. UNIX is a trademark of AT&T Bell Labs. Copyright © Excelan, Inc., 1988. All rights reserved.

COMPUTERS AND SUBSYSTEMS

Parallel machine extends scalability to encompass I/O processing

David Lieberman, Senior Editor

The claim to fame of the hypercube approach to parallel computing is its scalability. That is, because the architecture doesn't use a single central resource for interprocessor communications, there's no essential strain in adding processing resources to a system and no inherent hindrance to achieving near linear performance improvements. The bane of all forms of parallel computing, however, has been difficulty of use, especially the onerous task of porting existing single-processor programs to run efficiently on a multiprocessor system. Last month, Intel Scientific Computers (Beaverton, OR) brought a new set of facilities to its iPSC/2 computer that addresses both issues.

The company has extended scalability and parallelism beyond its basic computing architecture to encompass I/O processing as well, and the new I/O facilities are accompanied by filing-system software that transparently manages I/O, hiding the complexities of parallel I/O processing. "This eliminates the last major barrier to the use of large-scale parallelism in a wider range of applications: getting I/O that's both powerful and easy to use," says Steve Cannon, the company's graphics and I/O product manager.

In the existing iPSC/2, each 80386-based compute node in a hypercube has eight full-duplex communications channels, each operating at 2.8 Mbytes/s. Seven of the channels are dedicated to the architecture's Direct-Connect internode communications network, while one channel on one node is used as the link to the Systems Resource Manager (SRM), an 80386-based front-end. The eighth channel of all but one compute node is, thus, idle, and the system has just one window on the outside world. In addition, external disk capacity in the current product is limited to that of the SRM-resident disk drive, 140 Mbytes in the current configuration.

With the new scheme, the eighth channels not used in the previous version of iPSC/2 can now each be connected to a SCSI peripheral subsystem consisting of as many as seven 5¼-in. Winchester disk drives per node. To make the connection, a SCSI module is placed on either a compute node or a new architectural element: the I/O node. The nature of the particular application at hand will determine where the SCSI module is most aptly located. "If a customer has a very I/O-intensive problem—such as a data base application whose whole point is doing I/O—he may want one or more disk drives feeding data directly into a compute node," Cannon says. Using an intervening I/O node, on the other hand, will be more appropriate for more traditional scientific computing problems, "where the system isn't necessarily I/O-bound but needs a lot of I/O, and it's desirable to distribute I/O tasks among as many separate processors as possible," he notes.

The I/O node of an iPSC/2 is basically the same as a compute node, except that it has only one Direct-Connect channel, which links it to a compute node. It's essentially irrelevant which compute node an I/O node is linked to since, no matter where the physical connection occurs, the I/O node becomes a part of the hypercube network. An iPSC/2 system can include up to one I/O node per compute node, providing a modularly upgradeable I/O system that can be tailored to the I/O bandwidth requirement of an application.

■ Managing the I/O channels

Intel certainly isn't the first computer vendor to replace a large disk drive with multiple small drives. This approach is merely an extension of the philosophy of scalability, which holds that because of economies of scale, configuring multiple microprocessors in parallel brings tremendous price/performance benefits to the computing arena; so too with multiple small drives and I/O channels. "Everyone's seen the benefit of that," says Cannon. "A big, fast disk drive system will probably run about $60/Mbyte, while smaller

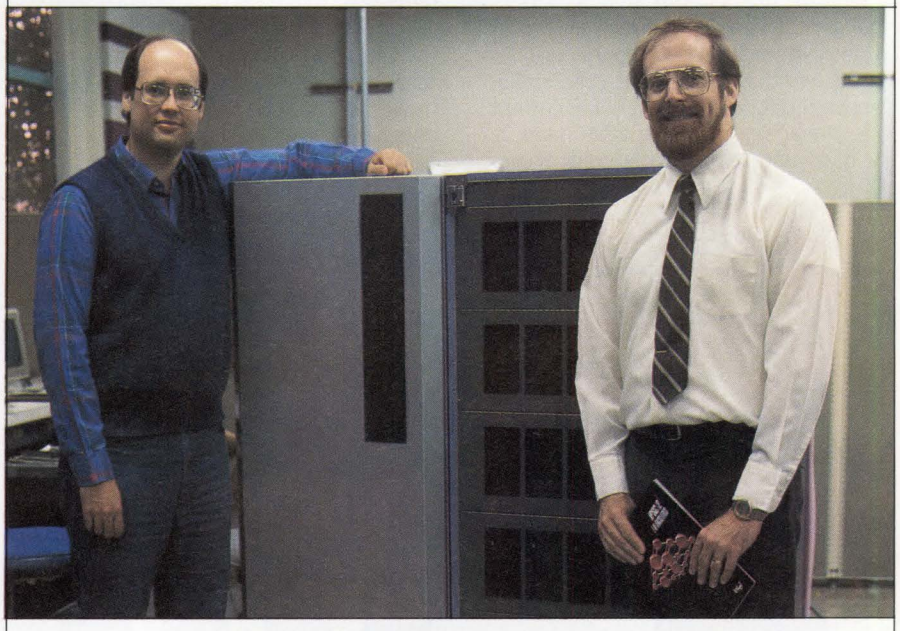

There are clear price/performance benefits of using multiple small drives and I/O channels rather than larger disk drive systems, according to Steve Cannon, Intel Scientific Computers' graphics and I/O product manager (right). But Paul Pierce, senior software engineer (left), adds, "Just putting a constellation of disks out there can create a real ease-of-use problem for the user." Intel addressed this problem by creating a distributed file system in which the specific details of the configuration are essentially transparent to an application program.

systems will cost more like $20 to $30/Mbyte, depending on how many drives are on each node, with cabinetry and everything else considered. But the trick is making the multiple drives easy to use."

Intel's scheme differs in its method of handling multiple I/O subsystems. "Just putting a constellation of disks out there can create a real ease-of-use problem for the user," notes Paul Pierce, senior software engineer at Intel Scientific and chief architect of the new Concurrent Filing System (CFS) software. "The common and most straightforward approach is to create a distributed file system by assigning a drive or set of drives to every processor, letting each processor run its own file system, and tying them all together at some higher level so one processor can get to the data on another processor's drive if it needs to, however inefficient that process may be. That's not much different from a network of workstations with multiple file servers."

■ The dedicated drive dilemma

There's a serious dilemma involved in the dedicated drive approach, though, in terms of deciding what data will reside on which drive. "Breaking up your data set into, say, 16 chunks for storage on 16 drives may be fine for disk striping with a single-processor system," claims Pierce, "but in a multiprocessing system, it's unlikely that the allocation of data for one job will be appropriate for the next. Since you'll never know until run time which node/drive or combination of nodes/drives a particular job will be assigned, either the right data isn't going to be immediately available, or you're going to have to reallocate the data every time you start a job."

"The way the data's divided in a distributed file system will be application-dependent and will vary with both the number of disks available and the number of nodes to be used," adds Cannon. "Creating this data jigsaw puzzle and then deciding where to put the pieces is a difficult task in itself, but what happens when things change? What if the job requires the use of additional nodes or additional disks, or it runs on different subcubes

Concurrent access vs. disk striping

The Concurrent Filing System (CFS) architecture of the iPSC/2 from Intel Scientific Computers (Beaverton, OR) is somewhat reminiscent of disk-striping I/O architecture in that both architectures segment files and store the pieces on multiple disks. But the similarity ends there. The intent of disk striping is to gang multiple disks to effectively provide a higher-speed I/O channel to a processor than a single disk drive could accommodate; the intent of concurrent file management, on the other hand, is to give multiple processors quick access to data by means of multiple parallel I/O channels.

Disk drives can be striped in a variety of ways. The most typical methods divide bits, words, blocks and so forth among multiple disks, which transfer their data in some coordinated fashion—either simultaneously or in a specific sequence. With its "deck-of-cards" data apportionment, the CFS also stores sequential blocks of data on different disks, but requests to those disks and the actual transfer of data are handled asynchronously. "This frees the system from having to queue and coordinate requests and from summoning data from a disk array in lockstep," says Steve Cannon, Intel Scientific's graphics and I/O product manager. "And you also get the power of parallel channels. When a node reads successive blocks, they're going to be flying in from different I/O nodes, and there will be multiple I/O channels running at the same time."

"Striping doesn't work well unless a single processor owns the whole disk array," says Paul Pierce, senior software engineer at Intel Scientific and chief architect of the CFS. If used with a multiprocessor system, striping would require a high-speed multiplexer to route data to its proper destination, whereas data distribution in the iPSC/2 is automatically handled by the Direct-Connect hypercube network.

Also, when used with multiple processors and/or multiple users, striping can become very inefficient. "Let's say processor A gets the drive's heads all lined up right and the data's coming off the disk as fast as it can go," says Cannon. "Then the job swaps out and processor B comes in and asks for a different file in a different location on the disk. Overall performance will quickly degrade unless you have a lot of striped disk arrays to choose from and processors can lock each other out. A concurrent file structure isn't as vulnerable to this problem, since there's little chance at any instant that there will be multiple requests to the same drive."

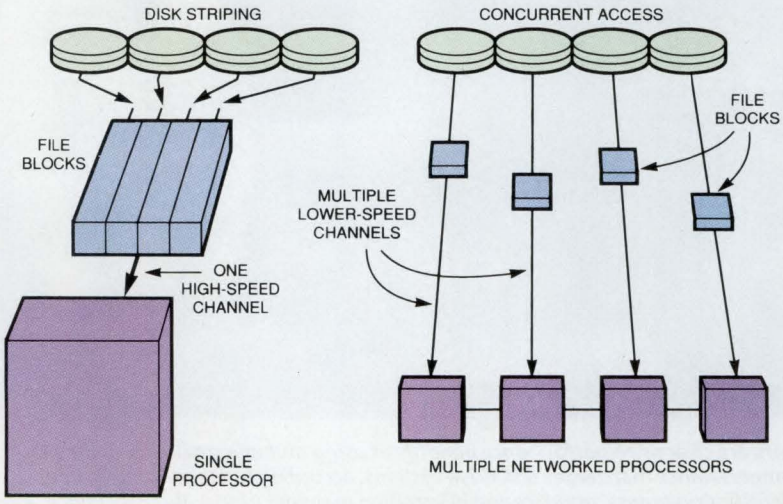

Although disk striping is effective for speeding data access for a single processor, it's less effective when multiple processors are involved. There's far less potential for bottlenecks to occur when there's concurrent access over multiple parallel channels.

of different sizes each day? The user is going to have to take the puzzle apart and then put it back together every time he runs a job. Obviously, this quickly becomes a very inconvenient way to manage things."

With Intel's CFS, in contrast, although specific disk drives are linked with a specific I/O node and that node is linked to a specific compute node, the particulars of the configuration are essentially transparent to an application program, which is aware of only one virtual resource. "Ideally," Pierce says, "just as the user is oblivious to which compute node is running his job, he should also be oblivious to where his data is stored. With the CFS, users don't deal with any of that unless they insist on it."

When a file is first read into the iPSC/2, the CFS breaks it up into blocks and deals these out in round-robin fashion, somewhat like a deck of cards, among all available I/O nodes and disks. "The benefit of this is twofold," says Pierce. "First, the work load is spread out among all the available resources. Second, because of Direct-Connect, no matter where a job is running on the hypercube, all the required data is, on the average, equally accessible to every node on the network without being run-time dependent."

For simplicity of use, the CFS programming model uses a Unix V System lookalike I/O system interface. The CFS also transparently manages cache coherency for the system, as well as other issues that would complicate the programming task.

Concurrent with the CFS introduction, Intel Scientific announced another new facility to enhance I/O: a choice of a Multibus II or VMEbus I/O bus for each node. "It's clear that the scientific computing community needs to be able to bring data in through the back door," Cannon says, "and that means the need for the disk to be tied in close to the hypercube as well as for standard external interfaces to take advantage of the wide variety of off-the-shelf peripheral devices that are available." □

MASS-STORAGE ARCHITECTURE

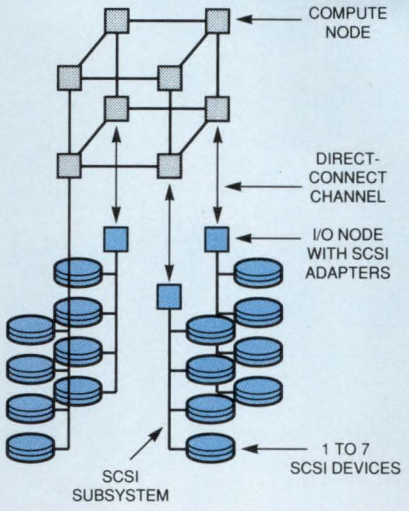

I/O capabilities for Intel Scientific Computers' iPSC/2 computer can now be scaled to the application by attaching SCSI subsystems to an I/O controller on either a compute node or I/O node. A system can accommodate as many subsystems as it has compute nodes.

IPI controller board talks to two disks at once

David Lieberman, Senior Editor

The Intelligent Peripherals Interface (IPI) is finally taking off as disk drive vendors push up their densities and data rates, and controller houses ready their boards to match the drives. This month, a new family of VMEbus IPI controllers from Xylogics (Burlington, MA) makes its debut with an architectural twist—the ability to communicate with two drives simultaneously. Until now, this ability could only be found in $20,000 to $60,000 formatters or head-of-string controllers used in mainframe environments.

What's the motivation for moving to IPI? I/O performance hasn't kept pace with the demands of high-performance CPUs. Also, as disk drive technology progresses to higher and higher bit densities, the resulting higher transfer rates off the disks are straining the limits of existing interface technology and causing reliability concerns. What's the value of dual IPI channels in the scheme of things? Double the potential data rate, of course, as well as the number of drives per controller.

The IPI edge

The device-level IPI-2 (the IPI standard also includes a host-level interface, IPI-3) is seen primarily as a replacement for the old SMD interface. Its major advantages over SMD include 16-bit parallel transfers (vs. one serial bit at a time); a maximum 10-Mbyte/s data-transfer rate (vs. 3 Mbytes/s) in its initial version and even higher rates on the horizon; a single-cable (vs. dual-cable) implementation; a daisy-chained data cable arrangement (vs. one data cable per drive); about a 164-ft maximum length (vs. 50 ft); and a degree of intelligence in managing status and configuration information.

"It's a much cleaner, more efficient way of running a string of drives," according to Chappell Cory, Xylogics senior vice-president for marketing operations.

The IPI itself already doubles the number of drives that a single board can reasonably control. Since the SMD interface requires one space-consuming data-cable connector for each attached drive, SMD controllers have been limited to the management of four Winchesters. The first IPI-2 controllers, in contrast—the Cougar series from Interphase (Dallas, TX) and Xylogics' SV series—each control eight drives via a single connector and cable. In the Xylogics architecture, however, it's also possible to add a daughter board containing a separate set of back-end resources and a second connector to manage two independent IPI subsystems and a total of 16 drives.

Although IPI controllers may initially be more expensive than SMD controllers, they offer inherently greater economies—from cutting down on expensive connectors and cabling to halving the number of controller boards needed to manage a particular number of drives. By managing twice the drives, they also halve the number of backplane slots

National Semiconductor

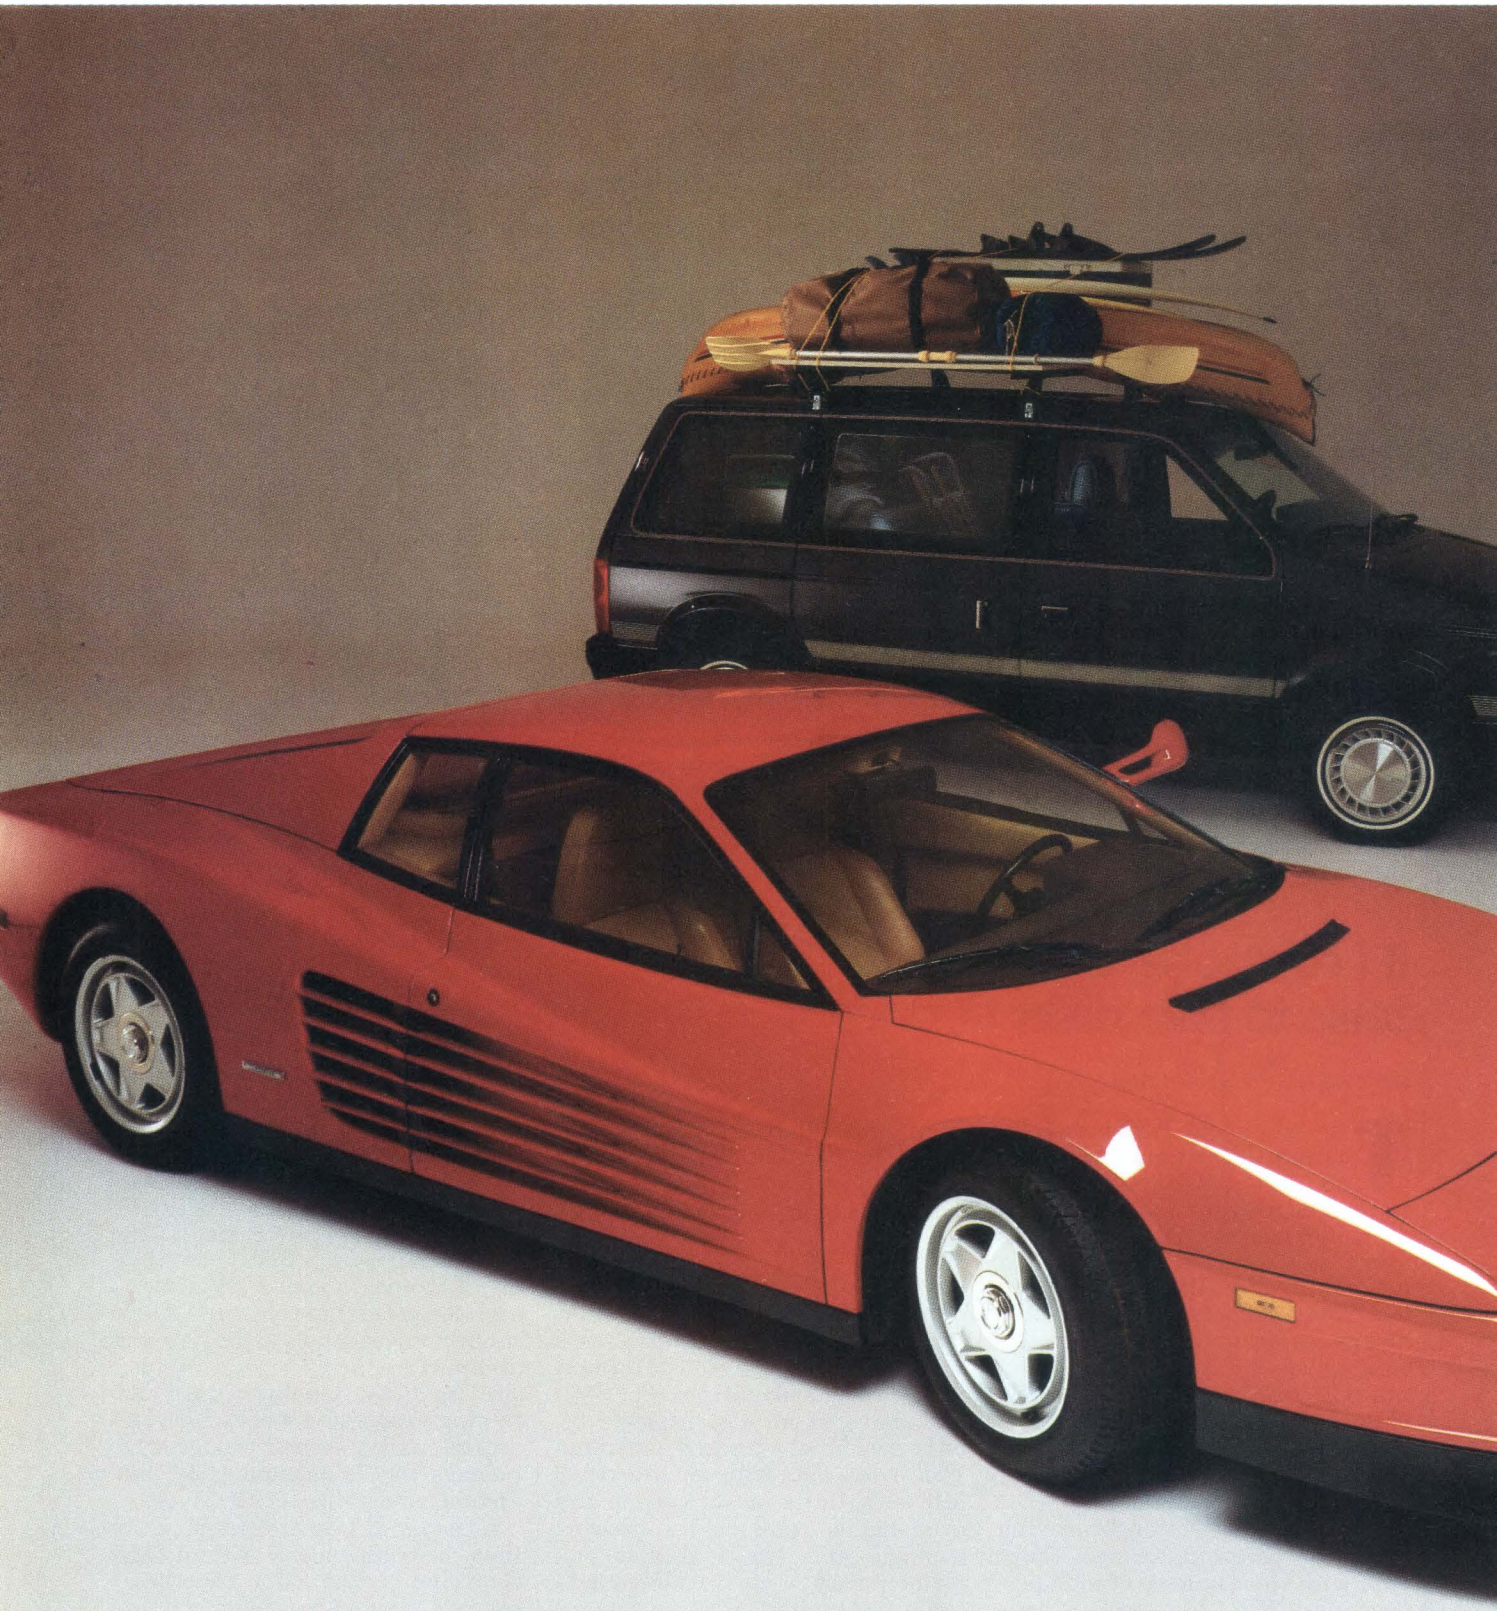

Now you don't have to choose between how fast you can go and how much you can carry.

OUR NEW 256-KBIT BiCMOS ECL SRAM GIVES YOU BOTH HIGH SPEED AND HIGH DENSITY

Finally, you can go all out for system performance without being forced to use low-density SRAMs. This new device has an access time of only 15 nanoseconds. Yet it has the storage capacity of four 64-kbit ECL chips in one-fourth the space.

Its modest power consumption makes air cooling not only possible but practical. It draws a maximum of 200 milliamps at 15 nanoseconds. Longer cycle times drop the current down to about 120 milliamps.

In worst-case testing, power dissipation stays below a single watt. In typical applications, it's only 600 to 700 milliwatts.

BiCMOS III IS THE KEY

What makes it all possible is National's proprietary BiCMOS III one-micron process, which combines the speed of pure bipolar with the high density, low power, and manufacturability of CMOS.

It gives you the best ratio of speed, power and cost per bit.

And we designed reliability into the process from the beginning. In fact, one-third of the development team had reliability as their sole responsibility.

The pay-off is a life test failure rate less than 50 FIT with more than 400,000 device hours. And a soft error rate less than 100 FIT. What's more, latch-up immunity is more than 200 mA. And ESD protection is greater than 2,000 volts.

IT GIVES YOU THE PERFORMANCE YOU WANT

Because of its stable memory cell, the 256-kbit SRAM tolerates skewed address signals without disturbed bits.

Customers—like Cray Research—also report low system noise and cleaner, quieter, faster signals. This high noise immunity is due to on-chip decoupling and minimal di/dt.

And because of the chip's 10-ns write pulse with zero setup and hold, it's easy to design-in in a 15-ns write cycle.

Our ECL SRAMs are input/output compatible with existing 10K and 100K ECL parts. Industry-standard pin-outs enable multiple sourcing.

AND YOU GET MORE MODELS TO CHOOSE FROM

In addition to our high-performance SRAMs, we're working on a family of specialty ECL RAMs and ECL RAMs in advanced packaging.

We also offer you the broadest array of memory product families of any domestic supplier—including TTL SRAMs, EPROMs, EEPROMs, and fuse-link PROMs.

This breadth of product line helps you minimize the number of suppliers you use. While National's large-volume manufacturing ensures availability.

CALL FOR OUR NEW BROCHURE

For more information on our new 256-kbit SRAM—or any other memory product—ask for a copy of our new memories brochure. Call us toll-free at (800) 252-4488, ext. 735.

National Semiconductor

CIRCLE NO. 12

© 1988 National Semiconductor Corp.

COMPUTERS AND SUBSYSTEMS

required to implement a mass-storage subsystem. Furthermore, they offer greater system efficiency in that, for example, only half the number of boards vie for the bus via arbitration and half the number contribute to the overall burden on system bus bandwidth. With the dual channel, one IPI controller board takes the place of four SMD boards.

Besides letting overlapped read/write operations occur at the same time, giving better overall subsystem responsiveness, a dual-channel capability also enhances system flexibility. Such a capability provides, for example, a relatively easy and economical path to disk striping (splitting a file among multiple drives to boost I/O transfer rates), as well as to disk shadowing (storing the same data on two drives) for building fault-tolerant systems.

A high-speed channel

The dual-channel capability of the SV series depends on three basic architectural characteristics: modularity, dynamically segmented caching and a high-throughput data path with no bottlenecks. The architecture, in turn, depends on an 80186 microprocessor and five newly developed application-specific ICs: a buffer management chip (BMC), a bidirectional first-in, first-out (FIFO), a DMA chip, a parallel disk back-end chip, and an error checking and correction (ECC) chip. One ASIC from Xylogics' earlier designs, the register read/write interrupt chip, implements the company's familiar software interface protocols and lets customers upgrade to IPI with only a minor driver modification.

"IPI-2 is a much cleaner, more efficient way of running a string of drives."

—Chappell Cory, Xylogics

■ ■ ■

The SV boards' internal data rate clocks in at up to 80 Mbytes/s, burst capability across the bus exceeds 35 Mbytes/s, and the company claims a sustained throughput of up to 35 Mbytes/s. The boards' internal bus is 36 bits wide, consisting of 4 bytes and one parity bit per byte, supplementing the parity that's provided on IPI (but lacking on VMEbus). A version with 16-bit data plus 2 bits of parity will also be available.

Multithreaded caching

The BMC chip contains four I/O channels for time-sharing its services among the 80186, the cache and the back-end resources (one set for the single-channel SV6800, two sets for the dual-channel SV7800). The chip is responsible for organizing and managing up to five independent cache buffer threads that are dedicated, as required, to different users, applications and/or disk locations. The 80186, which runs the boards' proprietary real-time kernel and a set of allocation algorithms, examines incoming operating system requests to determine the optimal buffer segmentation, which the BMC will then implement.

"Multithreaded caching is important because today's multitasking and multiuser environments have many independent I/O requests occurring simultaneously, which conventional caching designs don't support," says Cory. The adaptive feature of the SV boards' multithreaded caching, according to Cory, "greatly increases I/O subsystem performance by self-tuning to changing application environments. Intelligent caching is particularly important to Unix applications because they often involve many types of request profiles to the I/O system."

The SV series also features from 512 kbytes to 8 Mbytes of RAM, read-ahead algorithms, zero-latency reads and writes, scatter/gather capabilities, data concatenation, advanced Reed Solomon ECC, and dynamic command queuing and optimization. Pricing for the single-channel version starts at under $2,000 in large OEM volumes and $2,995 in single quantities; the dual-channel board is expected to cost about 40 percent more. Expect to see future versions of the architecture for different buses and different interfaces. □

XYLOGICS SV SERIES ARCHITECTURE

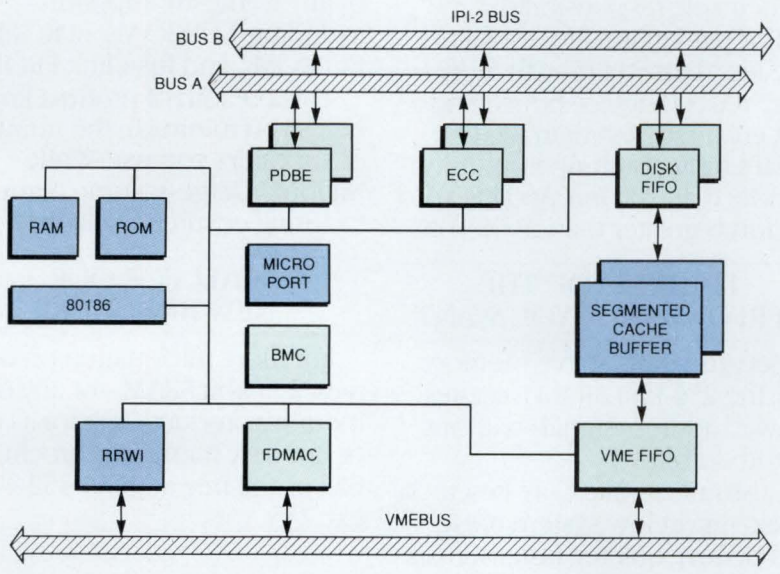

Xylogics' SV series architecture utilizes five new application-specific ICs—a bidirectional first-in, first-out (FIFO), a buffer management chip (BMC), a parallel disk back-end (PDBE), a fast DMA chip (FDMAC) and an error checking and correction (ECC) chip that offers four choices of correction methodology. The register read/write interrupt (RRWI) chip is a carry-over ASIC from previous designs. A set of transceivers between the microprocessor and data buses, called the Micro Port, lets the boards' 80186 communicate directly with the cache, ECC logic and so forth when necessary.

People are flocking to see the fastest machine in workstation printed circuit board design.

Don't be passed by. Find out why.

THE TOTAL CONCEPT COMPANY

ASI®

AUTOMATED SYSTEMS INC.

1505 Commerce Ave.
Brookfield, WI 53005
(414) 784-6400

CIRCLE NO. 13

Yes! I want to know more about ASI's Prance GT– the fastest software in automated PCB design.

☐ **Call me to set an appointment**

☐ **Rush information to:**

Name _____ Title _____

Company _____ Phone _____

Address _____

City _____ State _____ Zip _____

Mail this coupon today to: **ASI Marketing**
1505 Commerce Ave., Brookfield, WI 53005 • (414) 784-6400

CD

INTEGRATED CIRCUITS

Process advances bring new price, performance to GaAs

Ron Wilson, Senior Editor

New introductions from gallium-arsenide programmable array logic vendor Gazelle Microcircuits (Santa Clara, CA) have moved that company one step further ahead of its silicon-based competitors. And a new 4,200-gate array family from Triquint Semiconductor (Beaverton, OR) offers 1-GHz speeds with direct I/O compatibility to TTL, CMOS or ECL silicon circuitry.

The new super-speed products represent more than just a new alternative for microcomputer board designers: the parts are first indications of a new maturity in GaAs processing. As this once laboratory-bound technology continues to mature, it's changing the role of GaAs in digital and analog design.

The Gazelle announcement contained two messages: the introduction of a pair of 22V10-compatible 7.5-ns PALs and a major price reduction on the company's existing 10-ns 22V10. Like their predecessor, the new PALs are built from GaAs but are entirely TTL pin-compatible, including signal levels and impedances. The chips rely on Gazelle's patented circuit designs to provide strict TTL compatibility with the added benefits of controlled edge rates, which minimizes transient effects on the rest of the board, and extreme temperature stability compared to silicon devices. Like the original GaAs 22V10, the parts are factory programmed by Gazelle.

The first of the new PALs, the 22V10-7 is precisely what its name implies: a 22V10 with 7.5-ns maximum propagation delay and 110-MHz operating frequency. The second device, the 22VP10, represents a minor change to the 22V10 architecture. Instead of taking feedback from within the I/O macrocell, the 22VP10 takes its feedback from outside the output pin buffer. This lets the macrocell output pins, which are dedicated outputs on the conventional 22V10, serve as inputs or outputs. The change gives the 22VP10 slightly more flexibility than the conventional 22V10, particularly for connecting I/O pins directly to bidirectional buses. The 22VP10 also offers 7.5-ns propagation delay and 110-MHz operating frequency.

At 7.5 ns, the new parts offer designers an interesting alternative, according to Gazelle product marketing manager Robert Gunn. "Not only are these new PALs the fastest parts of their complexity on the market, they're the fastest available technology for implementing many types of microcomputer support logic. The parts are faster than the fastest versions of simpler silicon PALs such as the 16L8 and 16R8, and they're even faster than discrete logic for many functions. So the designer can implement the fastest available technology without giving up 22V10 flexibility," says Gunn.

■ GaAs processing matures

The new parts will certainly be exciting news to a number of design teams that are finding themselves in timing trouble because of higher microprocessor clock frequencies. But in the long term, the most significant part of the Gazelle announcement may not be the new parts, but the implication about the state of GaAs process technology.

"There have been significant advances in reducing the cost of GaAs processing in recent months," claims Gunn. "The process we're using now has 30 percent fewer overall steps than a typical silicon process and has only one-half the critical mask steps." In addition, the new parts are being fabricated on 4-in. wafers, a first for an industry that has been struggling with the overhead of 3-in. production lines, according to Gunn.

■ Geometries shrinking

Geometries are also shrinking and apparently have a long way yet to go. Gazelle's current parts are built on a two-metal, 1-micron enhancement-depletion process, and the company is moving to a 0.7-micron process. That may be just the beginning. "Because of the self-isolating nature of GaAs, you can pack circuits a lot closer," says Gunn. "And the accumulated experience of the monolithic microwave circuit people, who have been working at geometries as fine as 0.25 microns, means there's already a good deal of groundwork for smaller sizes of features."

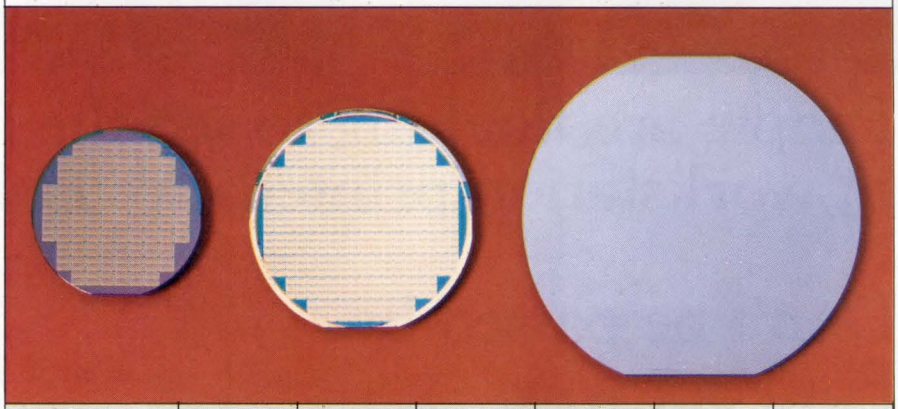

	Aug. '87	Jan. '88	May '88	Dec. '88	April '89	June '90
Wafer Diameter (inches)	3	3	4	4	4	6
Gross Die Per Wafer	1x	1.2x	2.5x	2.9x	3.6x	8.1x

Wafer processing cost has been a big factor in the price of gallium-arsenide ICs. Now the industry is successfully moving from 3-in. to 4-in. GaAs wafers, and Gazelle predicts that by June 1990, its GA22V10 will be available as 6-in. wafers.

Triquint Semiconductor, a GaAs application-specific IC house and foundry, has just announced a new GaAs gate array family with up to 4,200 equivalent gates and 1-GHz toggle frequencies, and is equally bullish on the progress of processing technology.

"The availability of quality crystals has been limiting our wafer sizes," explains Louis Pengue, Triquint product marketing manager. "Now, however, we're getting excellent 4-in. wafers, and we see the vendors on the learning curve toward producing 6-in. wafers."

Since GaAs costs are dominated by overhead rather than materials, the increase in wafer size will translate directly into a geometric reduction in die cost. And Pengue agrees with Gunn that geometries will continue to shrink as well. "We're able to move faster than the silicon processing people did simply because we can benefit from their experiences," explains Pengue.

Circuits boost performance

Pengue is also emphatic about the role of good circuit design in the advancement of GaAs. "Until recently, most of the expertise in GaAs houses has been in process engineering, not in circuit design," he says. "So a lot of the early products had primitive circuitry and used the brute speed of GaAs to compete with much more sophisticated silicon designs. Now that's changing."

The change requires three steps, according to Pengue. First, the process people have to stabilize their line so devices are consistently coming out with stable parameters. Then the device specialists have to do a good job of modeling the devices. "Our lead designer on our pin driver product spent the first months of the design just poring over transistor I-V plots, looking for discrepancies between the devices and the Spice models," notes Pengue. Only after detailed, reliable models are in place can top circuit designers come in and really take advantage of what GaAs has to offer. The goal is to use the circuit people's talents to boost product performance, not to overcome process variations.

So behind the scenes, GaAs is about to benefit from a number of simultaneous advances. Wafer sizes are increasing and geometries are shrinking, reducing costs substantially. And increasing experience with the process is translating into more sophisticated, faster circuit designs. In terms of cost and performance, easier-to-use, blindingly fast GaAs parts can stake out a big territory in digital design. □

10 MIPS and No RISC

Introducing the HK32/V532 VME microcomputer for **$4,995*.**

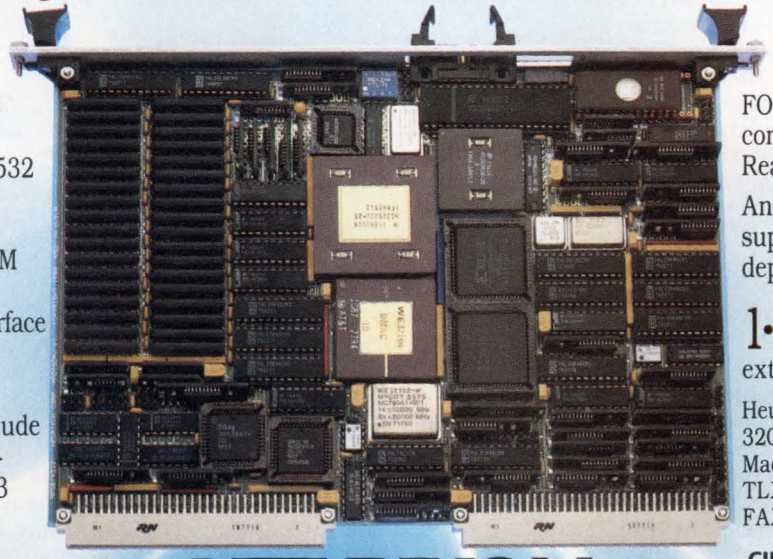

ON-CARD FEATURES INCLUDE:

- Up to 30 MHz NS32532 CPU
- NS32381 FPU (opt.)
- 4 or 16 Mbytes DRAM
- 4-channel DMA
- Sync/async SCSI interface
- 2 serial ports

Sun and Heurikon development tools include 680X0 to 32532 cross-assemblers, UNIX™ V.3 with optimized C, FORTRAN and Pascal compilers and VRTX®32 Real-Time Executive.

And all with the technical support you've come to depend on.

1·800·356·9602
ext. #569

Heurikon Corporation
3201 Latham Drive
Madison, WI 53713
TLX 469532
FAX 608-251-1076

CIRCLE NO. 14

HEURIKON CORP.

*Evaluation price includes HK32/V532 with 20 MHz CPU and 4 Mbytes DRAM.

INTEGRATED CIRCUITS

British microprocessors: a lesson in innovation

Ron Wilson, Senior Editor

Transputer and Viper: even the names seem a bit odd. On closer examination, these British 32-bit microprocessors may seem even odder. Each machine is based on a set of architectural decisions entirely foreign to the mainstream of the U.S. microprocessor market. But as these British-designed CPUs accumulate design wins throughout the world, it's becoming apparent that their foreignness reflects differences between U.S. and British design environments. In fact, the chips offer uniquely crafted solutions to some of microcomputing's nastiest problems—solutions that U.S. engineers have, perhaps for cultural reasons, tended to overlook.

To many, the T800 Transputer from Inmos (Colorado Springs, CO; Bristol, England) is a strange design. A quick glance at the chip's block diagram suggests that somebody has tried to build a 32-bit microcontroller with on-board RAM, four I/O ports and a floating-point unit. But in fact, a simple list of hardware features misses the point of the Transputer.

"Inmos set out to do something quite different from the traditional sequential VLSI processor," explains David May, one of the Transputer's architects at Inmos. "Our idea was to build a component for implementing concurrent systems." May claims that there had been discussion of concurrent processing in the British academic community since the early 1970s. "But the crunch always came when someone asked how we were going to program these machines."

A process/message architecture

A solution to the programming problem generated both the Transputer architecture and the Occam language in which the chips are usually programmed. The Inmos designers chose to regard a system as a set of relatively independent processes, connected to each other through synchronized messages. The Occam language was designed to make this process-and-message organization a natural programming style, and the Transputers were designed to be an efficient, and scalable, implementation vehicle for Occam.

The result was a toolset for exploiting the concurrency in systems, particularly in embedded systems. "Embedded computing is the place where parallel computing is easiest to apply," observes May. "The decomposition of the system is already there, in the application. The hard job is scheduling the tasks into a single machine." The Occam/Transputer combination lets the designer identify the natural independent tasks in a system, code them and then distribute them among however many Transputers happen to be in the implementation. By maintaining this independence between the Transputers and the tasks, the designer can increase the system performance simply by adding Transputer chips. Tasks that once shared a chip and communicated via memory can move to separate chips and communicate via the Transputer's I/O channels, and hence can run concurrently.

This concept caught on quickly in Great Britain, West Germany and Japan, but U.S. engineers have been much less receptive, according to May. He suggests that part of the reason may be a difference in academic environments: while the Europeans were concentrating on concur-

(continued on page 41)

THE INMOS T800 TRANSPUTER

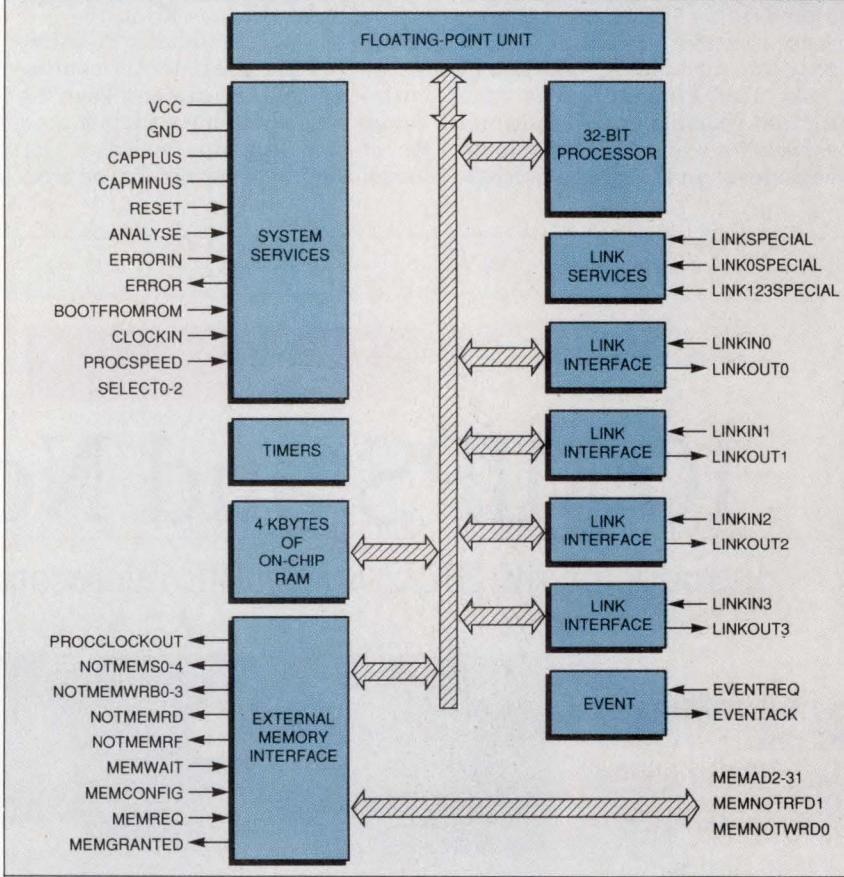

The Inmos T800 Transputer could be mistaken for a 32-bit microcontroller with floating-point hardware. In fact, the chip is intended to be an element in a multiprocessing system. The Link Interfaces form high-speed data paths between T800s.

WE'VE GIVEN THE WORLD THREE GREAT LEADERS

WREN™

THE 5¼" DRIVES THAT SET PERFORMANCE STANDARDS

Wren 5¼" disk drives give systems an extra performance edge. With an average seek time of a blazing 14.5 ms, Wren V is the fastest 5¼" drive we've made yet. Wren models are the fastest of their species because their patented, balanced, straight-arm actuator permits precision head positioning and super fast seek times.

ZBR INCREASES PERFORMANCE AND RELIABILITY

Wren models featuring Zone Bit Recording (ZBR), offer increased capacity of 30 percent or more compared with conventional recording techniques. In some cases, ZBR also means higher transfer rates that improve performance up to 30 percent. With ZBR's greater recording density, drives require fewer recording heads and less media for the same capacity. Fewer parts mean less potential for trouble. That's why Wren is a disk drive you can count on for years of trouble-free performance.

Model 94xxx	Capacity (Mbytes)	Avg. Seek (ms)	Interface	Transfer Rate (MHz)
Wren V	702	16	SCSI	12-16
Wren V	442	16	ESDI	10
Wren V	383	14.5	ESDI	10
Wren V H.H.	209	18	SCSI	10-15
Wren IV	376	17.5	SCSI	10-15
Wren III	182	16.5	ESDI	10
Wren III	182	16.5	SCSI	10
Wren III H.H.	106	18	ESDI	10
Wren III H.H.	106	18	SCSI	10
Wren II (RLL)	135	28	ST506-(RLL)	7.5
Wren II	96/86	28	ST506, ESDI	5
Wren II H.H. AT	80/74	28	AT	7.5
Wren II H.H. RLL	77	28	ST506-(RLL)	7.5
Wren II H.H.	51	28	ST506	5

All capacity figures are shown as unformatted.
H.H. = Half Height models

SWIFT™

200 MB + 15 ms = 3½" SWIFT DISK DRIVES PACKED WITH PERFORMANCE

The Swift 3½" disk drive sets the standard for small drive capacity and performance. It combines the two most important things people want in a storage device: capacity and speed. And packs it into the size everyone would like-3½".

CAPACITY PLUS PERFORMANCE

Swift features a blistering 15 ms average seek time. A 1-to-1 interleave feature allows read/writes of a full track in a single revolution. The result is a significant performance improvement due to the higher sustained data transfer rates.

Synchronous or asynchronous operation and a cache with read/look ahead capability further optimizes Swift performance.

DESIGNED FOR RELIABILITY

Low wattage, shock resistant design, rotary voice coil actuator, and a unique high torque, phase-locked, in-hub spindle motor are just a few of the advanced features that make Swift the industry's 3½" disk drive leader.

Model 9435X	Capacity (Mbytes)	Avg. Seek (ms)	Interface	Transfer Rate (MHz)
Swift	200	15	SCSI, ESDI, AT	10
Swift	150	15	SCSI, ESDI, AT	10
Swift	150	15	ST506/412 (RLL)*	7.5
Swift	100	15	ST506/412 (MFM)*	5

* Not available with 1-to-1 interleave

IN DATA STORAGE PRODUCTS.

IMPRIMIS

A Subsidiary of Control Data

"Recently Control Data made its Data Storage Products Group an independent subsidiary: Imprimis Technology Incorporated. As Imprimis, we will focus all our attention and energy on the data storage marketplace.

"Imprimis means 'in the first place.' That's fitting—because we are. And we intend to remain there. First in high-performance disk drives. First in quality. First in service and customer satisfaction.

"We were first to introduce a 3½" 200 MB disk drive, the Swift. Wren 5¼" disk drives were the first to feature sub-20 ms seek times, and they continue to be the fastest drives in the world. The Sabre 8" disk drive was the first of its size to offer more than 1,200 MBs of storage. But the biggest first of all is in commitment to customers.

"So, just because we changed the name doesn't mean we are changing our ways. We are continuing the commitment to leadership in the high-performance storage market, and we intend to remain the world's leading supplier."

Lawrence Perlman
President

WORLD-CLASS QUALITY WORLDWIDE CAPABILITY

Minnesota
Large Disk Division Development and Manufacturing

Canada
Sales and Distribution

Nebraska
Magnetic Media Development and Manufacturing

California
Small Disk Division
Development and Manufacturing

Oklahoma
Small Disk Division
Development and Manufacturing

Brazil
Elebra Informatica–
Licensed Disk Drive Manufacturer

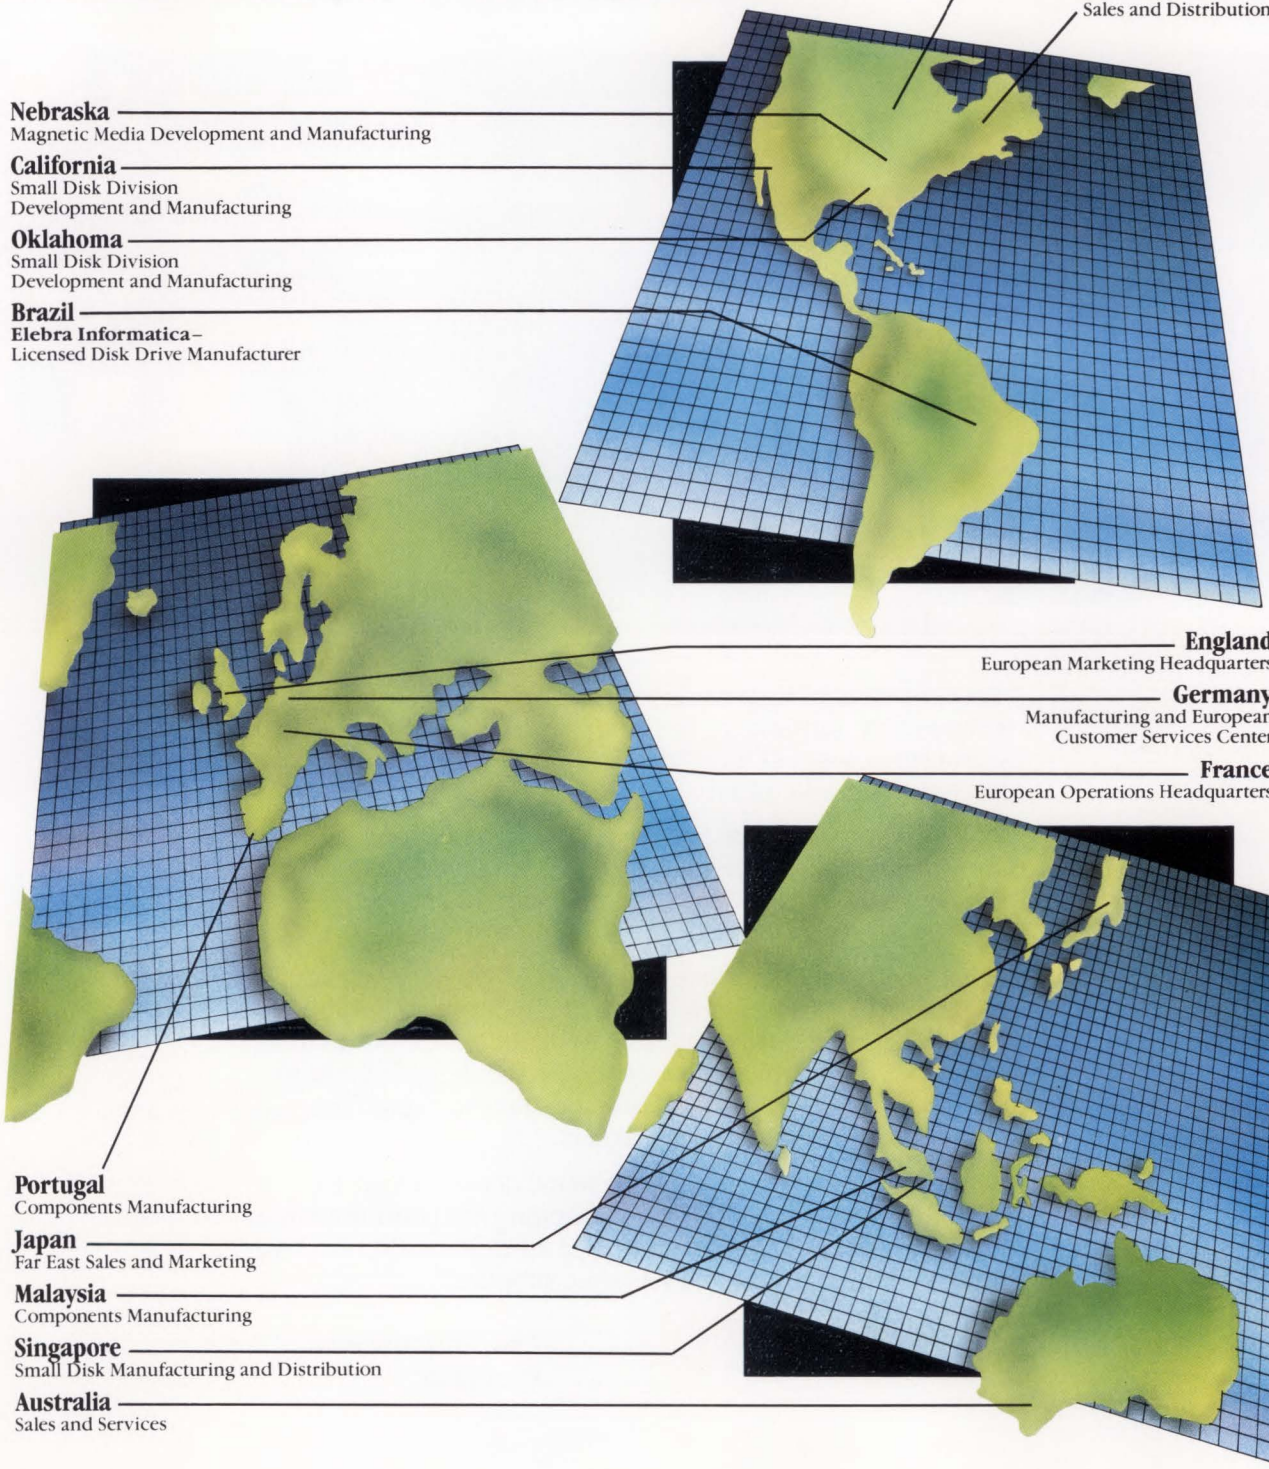

England
European Marketing Headquarters

Germany
Manufacturing and European Customer Services Center

France
European Operations Headquarters

Portugal
Components Manufacturing

Japan
Far East Sales and Marketing

Malaysia
Components Manufacturing

Singapore
Small Disk Manufacturing and Distribution

Australia
Sales and Services

SABRE™

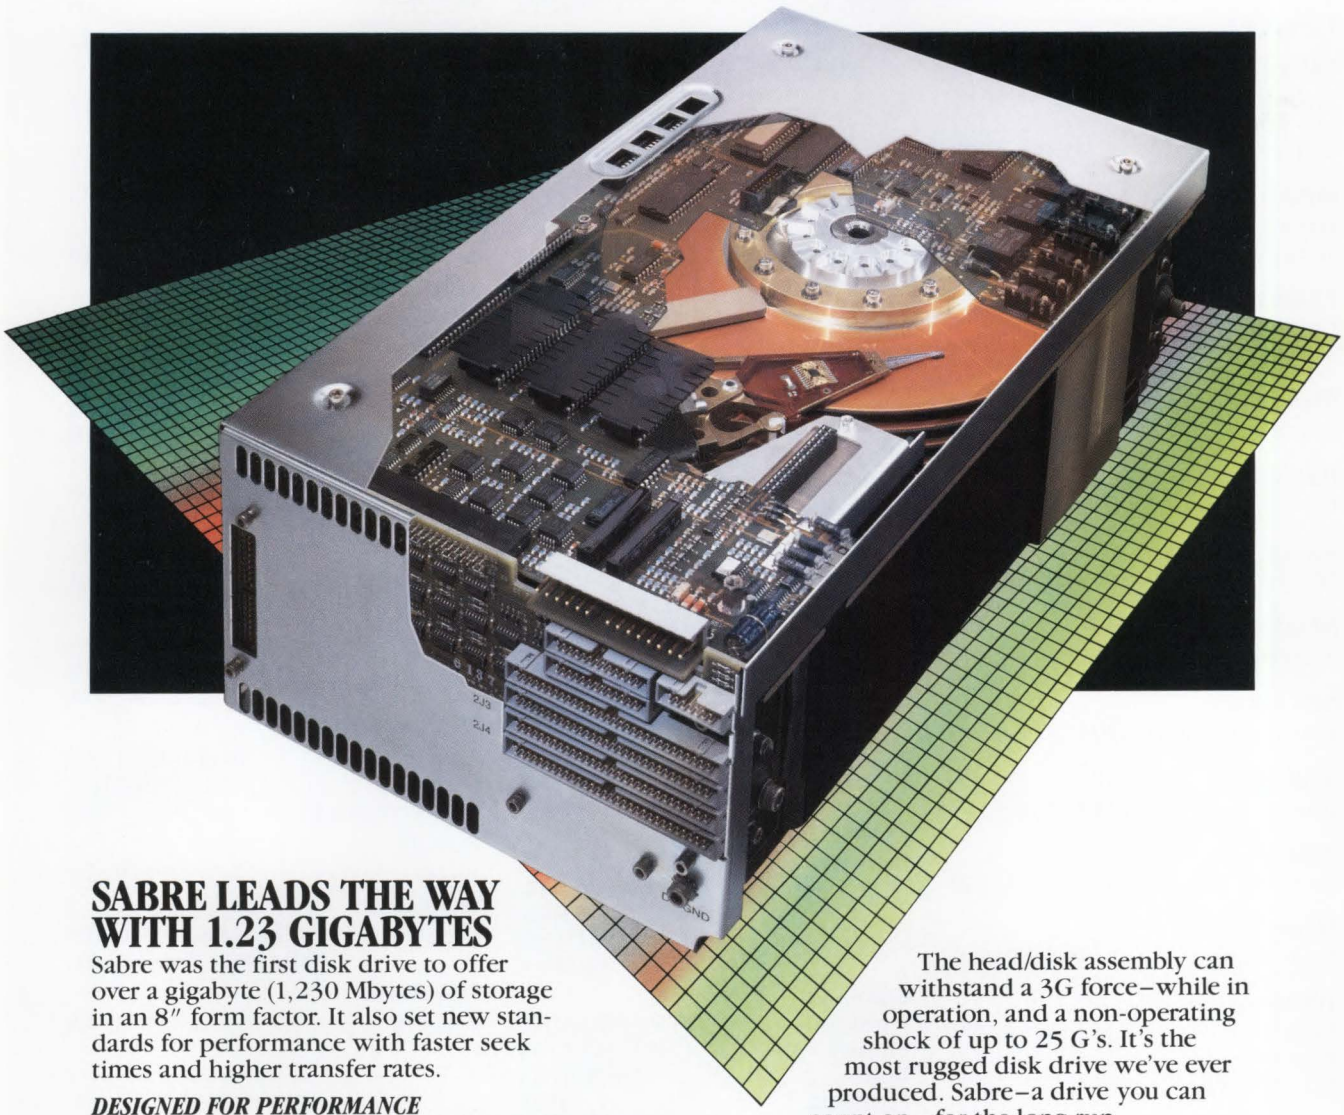

SABRE LEADS THE WAY WITH 1.23 GIGABYTES

Sabre was the first disk drive to offer over a gigabyte (1,230 Mbytes) of storage in an 8″ form factor. It also set new standards for performance with faster seek times and higher transfer rates.

DESIGNED FOR PERFORMANCE

Sabre is available in several capacities and with a variety of interface options. A 16 ms average seek time, coupled with transfer rates of 15 to 24 MHz, make Sabre drives truly high performance machines. Sabre's superior performance and capacity make it ideally suited for multi-user, multi-tasking systems.

BUILT FOR RELIABILITY

Sabre is designed to provide a 24-hour daily duty cycle ensuring the utmost in reliability.

The head/disk assembly can withstand a 3G force–while in operation, and a non-operating shock of up to 25 G's. It's the most rugged disk drive we've ever produced. Sabre–a drive you can count on– for the long run.

Model 9720x	Capacity (Mbytes)	Avg. Seek (ms)	Interface	Transfer Rate (MHz)
Sabre V	1,230	16	SCSI, SMD-O/E, IPI-2	24
Sabre IV	850	16	SCSI, SMD-O/E, IPI-2	19.7
Sabre III	736	16	SCSI, SMD-O/E	14.5
Sabre II	500	16	SCSI, SMD-O/E	19.7
Sabre I	386	16	SCSI, SMD-O/E	14.5

WORLDWIDE AVAILABILITY

OEM Sales and Support Services are in place around the world to help you.

U.S. AUTHORIZED DISTRIBUTOR
*ARROW ELECTRONICS INC.

For the location of the Arrow Sales Office nearest you, wherever you are...Arrow is as close as your telephone. In U.S. call toll-free

1-800-777-ARROW

U.S. OEM DIRECT SALES OFFICES

CALIFORNIA
Irvine (714) 851-5685
Los Angeles (213) 642-2222
San Jose (408) 280-2000

CONNECTICUT
Glastonbury (203) 633-6781
Greenwich (203) 622-2440

FLORIDA
Coral Springs (305) 753-2030
Tampa (813) 286-7672

GEORGIA
Atlanta (404) 399-2311

ILLINOIS
Chicago (312) 960-7545

MASSACHUSETTS
Waltham (617) 890-4600

MINNESOTA
Minnetonka (612) 931-8801

NEW JERSEY
Edison (N.Y. area) (201) 225-6100

NEW YORK
Albany (518) 785-3317

OREGON
Portland (503) 291-4358

PENNSYLVANIA
Valley Forge (215) 666-4540

TEXAS
Dallas (214) 385-5922
Houston (713) 778-5400
San Antonio (512) 736-2868

INTERNATIONAL OEM SALES OFFICES AND DISTRIBUTORS

AUSTRALIA
Melbourne (61) 3-268-9500
Sydney *IMAGINEERING

AUSTRIA
Vienna (43) 222-171733
*APH DATA
*W. MOOR DATA
Fischamend *PERCOM

BELGIUM
Merelbeke *BELCOMP

BRAZIL
Rio de Janeiro (55) 21-2663222
Sao Pablo *ELEBRA INFORMATICA

CANADA
Toronto (416) 826-8640
Montreal *ARROW CANADA
Scarborough *TENEX DATA

CHILE *SISTECO-CHILE

CHINA
Hong Kong 5-8683307
*TELTEC ELECTRONIC

COLOMBIA
Bogota *COLINVERSIONES

DENMARK
Copenhagen *DATA-VISUAL

ENGLAND
London (44) 1-848-1919
*RAPID SYSTEMS
*IDEAL HARDWARE
*SINTROM ELECTRONICS
*NORMAN MAGNETICS

FINLAND
Helsinki (358) 0-6926121
*INSELE

FRANCE
Paris (33) 1-6461-5000
*ALFATRONIC
*OMNILOGIC
*WINLINE
*ORBITEC

GERMANY
Frankfurt (49) 69-63050
Munich (49) 89-41790
Bonn *BONNDATA
Hamburg *IRTV
Tuebingen *HAMILTON & IMPEC
Filderstadt *TRIANGEL
Munich *TRIANGEL
*MCE

INDIA
New Delhi *CMC LTD.

ISRAEL
Tel Aviv (972) 3-543-4646
*MAG COMPUTERS
Herzlia *MINIX COMPUTERS
*I.E. MITTWOCH

ITALY
Milan (39) 2-2174221
*CELDIS
Modena *NAOS

INDONESIA *P.T. SKILL

JAPAN
Tokyo (81) 3-982-1051

KOREA
Seoul (82) 2-7564681
*TAESAN SYSTEMS INC.

MALAYSIA
Kuala Lumpur *MARKRITE SDN BHD

NETHERLANDS
Amsterdam (31) 70-119911
*MANUDAX

NORWAY
Oslo (47) 2-151400
*ICT DATAINFORM

SINGAPORE (65) 8614655
*IMAGINEERING MICRO

SPAIN
Madrid (34) 1-4560004
Barcelona *CEOI

SWEDEN
Stockholm (46) 8-7520020
*DATABOLIN
*ZETNER DATA

SWITZERLAND
Zurich (41) 1-2421434
Dietikon *INELTRO
Schwerzenbach *HSP

TAIWAN
Taipei (886) 2-7152222

THAILAND
Bangkok *THAI TECHNOLOGIES AND COMPUTER CO. LTD.

YUGOSLAVIA
Belgrade *PROGRES

* Authorized Distributors

Distributed by:

IMPRIMIS
A Subsidiary of Control Data

205,285A Litho in USA 88

INTEGRATED CIRCUITS

British microprocessors . . .
(continued from page 32)
rent systems, U.S. universities were focusing on single-processor RISC technology. Hence, while a generation of European and Japanese engineers were learning a rather abstract, computer-science-based approach to system architecture, U.S. engineers were studying the implementation of one particular architectural concept. "In America, parallel processing is emerging by tacking bits onto sequential machines. The approach works, and it doesn't violate the strictures of American design culture, but it's architecturally flawed," claims May.

This uncomfortable picture of the United States as a culture slow to adopt new concepts could sound a bit like sour grapes. After all, despite the fact that Inmos can boast about 1,000 design wins and serious evaluations worldwide, the Transputer has caught on very slowly in the United States. But confirmation of the image comes from a quite independent source with a quite independent notion of computing: Marconi Electronic Devices (Hauppauge, NY), a vendor of the Viper processor.

A formally proven CPU

The Viper CPU grew out of research work at the Royal Signal and Radar Establishment (RSRE) at Malvern, England. As David Scurr, marketing manager for space and defense products at Marconi, tells it, "In their work with military system designs and in the investigation of civil aviation disasters, the Malvern people became concerned about the industry's growing dependence on conventional microprocessors. They saw that, in the pursuit of higher performance, CPU designers were giving up the ability to test their designs." The result the RSRE feared was that designers could no longer predict with absolute certainty how a given CPU and a piece of software would behave in a given situation.

The RSRE felt that this risk of unpredictability wasn't acceptable in situations in which human life or the safety of the environment depended on correct functioning of a microcomputer system. The designers, therefore, turned to another field of academic research that's well developed in Europe yet little-known in the United States: formally proven logic.

The concept, if not the realization, of formal proof is straightforward. A complex circuit (or a piece of software) can be designed in such a way that its behavior can be described mathematically. The mathematical description can then describe the entire sequence of states through which the circuit can pass. And questions about the behavior of the circuit can be subjected to mathematical proof, rather than random testing. The correctness of a design can, thus, be determined with a level of assurance not possible with test vectors.

But the mathematics involved is rather primitive compared to the circuitry in conventional microprocessors. To design a 32-bit CPU that could be formally proven, the RSRE had to make some severe simplifying assumptions. Modules too complex for verification, such as multiply and divide hardware, had to be left out. So did the possibility of asynchronous events; consequently, the design doesn't support interrupts.

The RSRE design, christened the Viper, emerged as a very odd-looking 32-bit CPU indeed. The chip has a 32-bit, 16-function ALU and a small register set, reminiscent of early-1960s minicomputers or building-block signal processors. There's a 32-bit data bus, a 20-bit address bus and a lock-step instruction cycle so simple it can be described in an 11-bubble state transition diagram.

Viper has been turned over to private industry for implementation. So far, Marconi has released three versions: an initial 1-Mips CMOS part called Viper 1, a radiation-hardened SOS part, and a variant of Viper 1—the Viper 1A—intended for redundant-CPU applications. In the latter design, additional circuitry lets two Viper processors run in parallel, checking each other's work. Future versions envisioned by Marconi will use more sophisticated formal proofs in order to provide hardware multiplication and division, and it will use a protected mode in which unverified software may be run without compromising proven code.

The concern for reliability has to extend beyond just formal proof of the CPU circuitry, according to Marconi's Scurr. The entire hardware system incorporating the CPU must be verified, and then the software must also be formally proven. So far, these demands have meant that Viper designs can become a joint effort between the customer, Marconi, the RSRE and any of a number of British companies with expertise in the development and verification of proven hardware and software. As in the case with Ada's early years in the

THE VIPER PROCESSOR

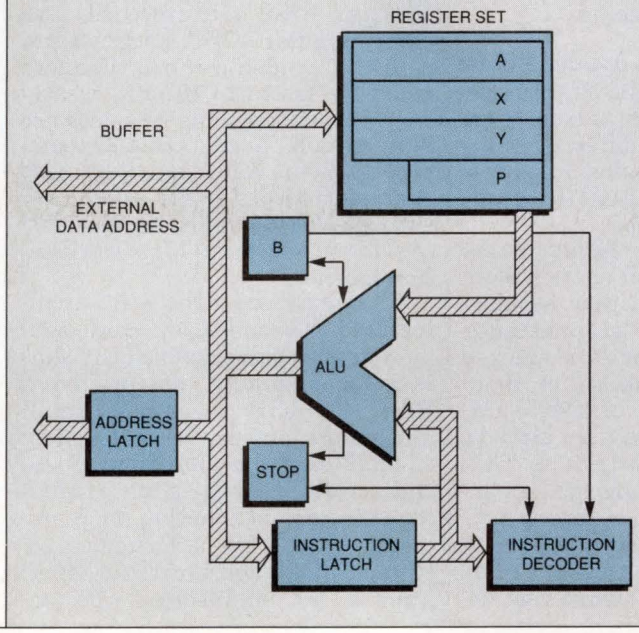

Evolving out of a fear that high-performance CPUs were not reliable, the Viper processor was developed around the concept of formally proven logic. Only logic that can be verified by formal mathematical proof is included in the CPU. The extreme simplicity of the Viper architecture reflects the chip's unique heritage.

INTEGRATED CIRCUITS

United States, the problem of using Viper is too complex for many individual design teams to go it alone. "There's more widespread interest now in high-integrity systems, particularly since some defense standards began to require formal methods," says Scurr. "What we see is a number of groups trying to form informal links to other groups."

While networks seem to be growing to address the formal verification problem in Great Britain, interest in the United States seems to be lagging far behind. There has been some work at a few U.S. universities on formally proven CPU designs, but apparently none has resulted in a prototype chip. There has been some interest in Viper and its supporting techniques from the U.S. military establishment, but so far the meetings have been organized by the RSRE, not by interested U.S. parties.

"Initially, we get a mixed response," says Scurr. "People who expect Viper to be another 32-bit RISC processor don't understand the lack of speed. People who have already thought about the problems of high-integrity systems tend to understand much more quickly." Scurr feels that, as the reliance on microcomputers increases in military systems, interest in Viper will grow. And such risks as lost life, failed missions or product liability suits can make integrity much more important than speed and can quickly bring designers to appreciate the slow, strange-looking Viper architecture.

■ Differences lead to opportunities

It seems clear that the Transputer and Viper have developed out of an environment quite different from that in the United States. In particular, the differences seem to be concentrated in four areas.

First, most British design teams avoid head-on competition with their U.S. counterparts. Either because they enter the market late or because they lack the huge domestic market and marketing resources of their U.S. competitors, the British usually won't duplicate a design already available in the United States. This bias has encouraged innovation, particularly in the cases of Inmos.

Second, the British seem to benefit from a much more successful relationship between academia and industry than do their U.S. counterparts. While U.S. semiconductor companies have tended to create architectures with little regard for current academic research, British architects seem to be keenly aware of the most recent European academic work in both architecture and software. This relationship may let important new methodologies move from theory to implementation much more smoothly in Great Britain.

Third, British designers seem freer to adopt new ideas. British companies aren't constrained by their own successes: Inmos had no prior generation of microprocessors with which it had to stay compatible. Also, the European market that the British designers serve does seem much more able to understand and adopt a new idea. The United States is developing a reputation as a hard place to sell a new concept.

Fourth, there are enormous differences in the capital environments between the United States and Great Britain. British observers and, increasingly, U.S. economists look upon the venture-capital domination of U.S. industry as a serious restraint on innovation. There's a tendency for only familiar ideas to get funding. This attitude of conservatism has even penetrated the planning departments of large semiconductor and computer companies, to the extent that one innovative product is quickly followed by a flood of imitators, all built on the same concept.

In the short run, these differences in climate have created a significant opportunity for U.S. designers. A team that recognizes the unique value of a device such as the T800 or Viper-1 may gain a considerable cost/performance advantage over U.S. competitors struggling to solve their problems with more conventional chips. But in the long run, one has to wonder whether the United States can afford to be in the position of a late adopter—a country that picks up innovations only after they've been applied by manufacturers in Europe and Japan. Conservatism, even as the result of success, can be a dangerous policy. □

Acorn processor opens second front in battle for RISC acceptance

Ron Wilson, Senior Editor

By now the arguments in favor of RISC architectures are familiar. RISC machines can be heavily pipelined, can therefore run very fast, can benefit greatly from optimizing compilers, and are quite simple. Most U.S. microcomputer vendors have exploited these characteristics to produce their own super-powerful 32-bit microprocessors for the 10- to 20-Mips workstation market.

But one aspect of the RISC architecture has been largely overlooked by vendors seeking double-digit Mips ratings: simplicity. Potentially, a RISC CPU chip can be very small and simple, and can be designed by a small team in a short time. These characteristics suggest the possibility of a very inexpensive microprocessor of moderate performance, positioned at the heart of a new generation of not engineering workstations but personal computers. In fact, such a microprocessor has been developed, almost unnoticed by U.S. engineers. It's the Acorn RISC Machine (ARM) from Acorn Computers (Cambridge, England).

ARM developed out of a dilemma, according to Roger Wilson, manager of software development at Acorn. "In 1983, when Acorn was busy selling 8-bit PCs, we had our first look at the IBM PC and the Apple Lisa, and saw that despite their 16-bit processors, they were no faster than the existing 8-bit Acorn machines.

"We studied the problem at some length and concluded that we could not build an economical 16-bit PC that was significantly faster than the 8-bit design. Nor could we get there with any of the existing 32-bit chips. So we set out to develop a 32-bit CPU that could serve as the basis of a family of PCs with $1,200 to $3,500

retail prices," says Wilson.

The company's resource constraints seemed overwhelming. The CPU design team consisted of four engineers, and they had only about 18 months to get a chip into production. With limited capital resources, the team had to rely on existing 6502-based computers and Apollo workstations for equipment.

Becoming a RISC machine

Early in the design process, managing director Hermann Hauser returned from a trip to the United States with two papers that would change the direction of the design: the IBM and Berkeley RISC reports. The team recognized that the RISC concept offered a combination of speed and simplicity that just might make the design possible within their constraints. So the Acorn 32-bit machine became a RISC machine.

There was still a lot of basic architectural work to do. The team had to work out some fundamentals—such as the implementation of delayed branching and virtual memory. And it had to create an instruction set that both served the needs of language developers and met the constraints that the miniscule design resources placed upon the silicon. In a three-way interaction, one designer would create an instruction set simulation on the 6502, a second would report on the implications of implementing the set in silicon, and the software people would evaluate it for their own purposes.

The team discovered another dilemma. The RISC CPU gobbled memory bandwidth: the 16-Mbyte/s speed achieved by putting four 8-bit Acorn buses together wasn't enough. The conventional solution of adding caches would drive the CPU board out of the target cost range. So the team developed a proprietary memory controller that allowed the use of dynamic RAMs in nibble mode, and also built circuitry into the CPU that detects whether the next memory cycle will be in sequence with the previous one.

Such a system-level approach to cost/performance paid off. The production chip, implemented in the 2-micron process at VLSI Technology (San Jose, CA) produces 5 kDhrystones and up to 20 kFlops from a single 25,000-transistor chip, using only DRAM memory. The CPU, with its accompanying memory controller, I/O controller and video controller chips, makes up the heart of Acorn's 32-bit Archimedes PC family. System prices start at around $1,500. The CPU itself is available in the United States, as a chip or a library element, from VLSI Technology.

ARM represents both the heart of a personal computer line for Acorn and a promising element in highly integrated ASIC 32-bit systems for U.S. designers. The chip may also represent the beginning of a new aspect of RISC application: low-cost systems where the simplicity and high Mips/transistor density of the architecture become more important than peak performance.

Future developments at Acorn will extend the performance envelope of ARM without substantial changes in cost. And in the United States some designers are evaluating the larger, more expensive, but faster Sparc design from Sun Microsystems (Mountain View, CA) as a possible basis for a PC. Either architecture should reach a price/performance point that would provide devastating competition to today's high-end PCs, if a sufficient software base can be drawn together.

The accumulation of an adequate software base can seem like an insurmountable barrier to a new architecture. But it isn't. Acorn is at last catching up with its customers' demands for application software. And in September, Wordperfect (Orem, UT) announced that its Wordperfect version 4.2, perhaps the premier word processor in the IBM PC world, was available on Sun workstations, running under Unix in a window-managed environment. Given the RISC architecture's untapped potential and the insatiable demand for desktop Mips in the PC world, the switch to a RISC PC platform could happen with unexpected speed. □

ACORN'S RISC MICROPROCESSOR

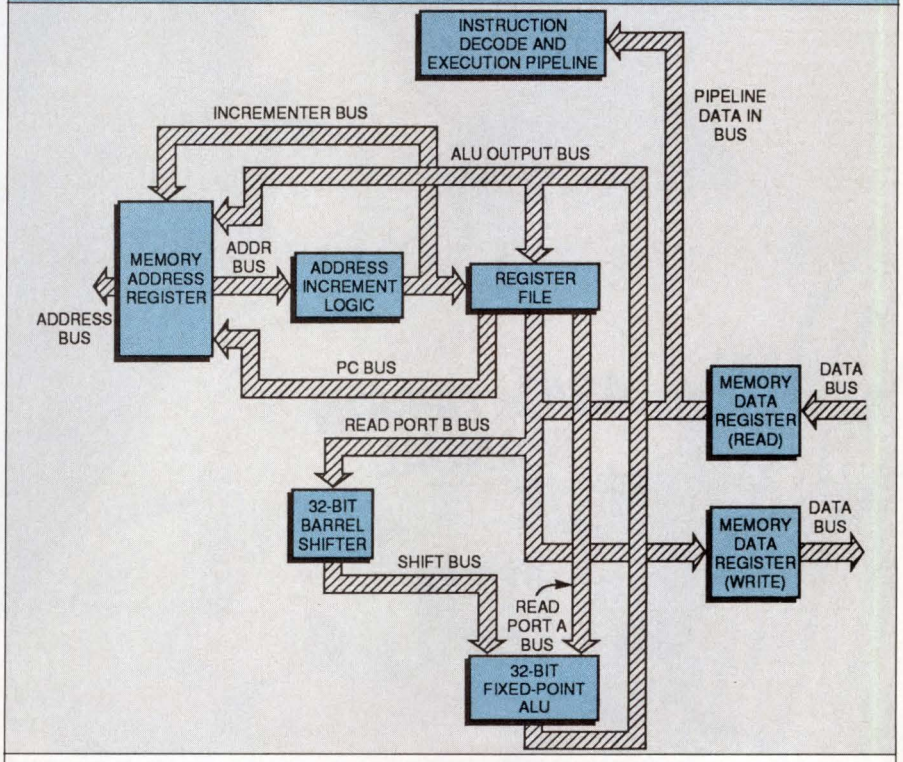

The simplicity of Acorn's RISC microprocessor permits its implementation in only about 25,000 transistors—about one-quarter the size of current Sparc integer units. The small die makes possible multi-Mips RISC computing at personal computer prices.

Our ASICs

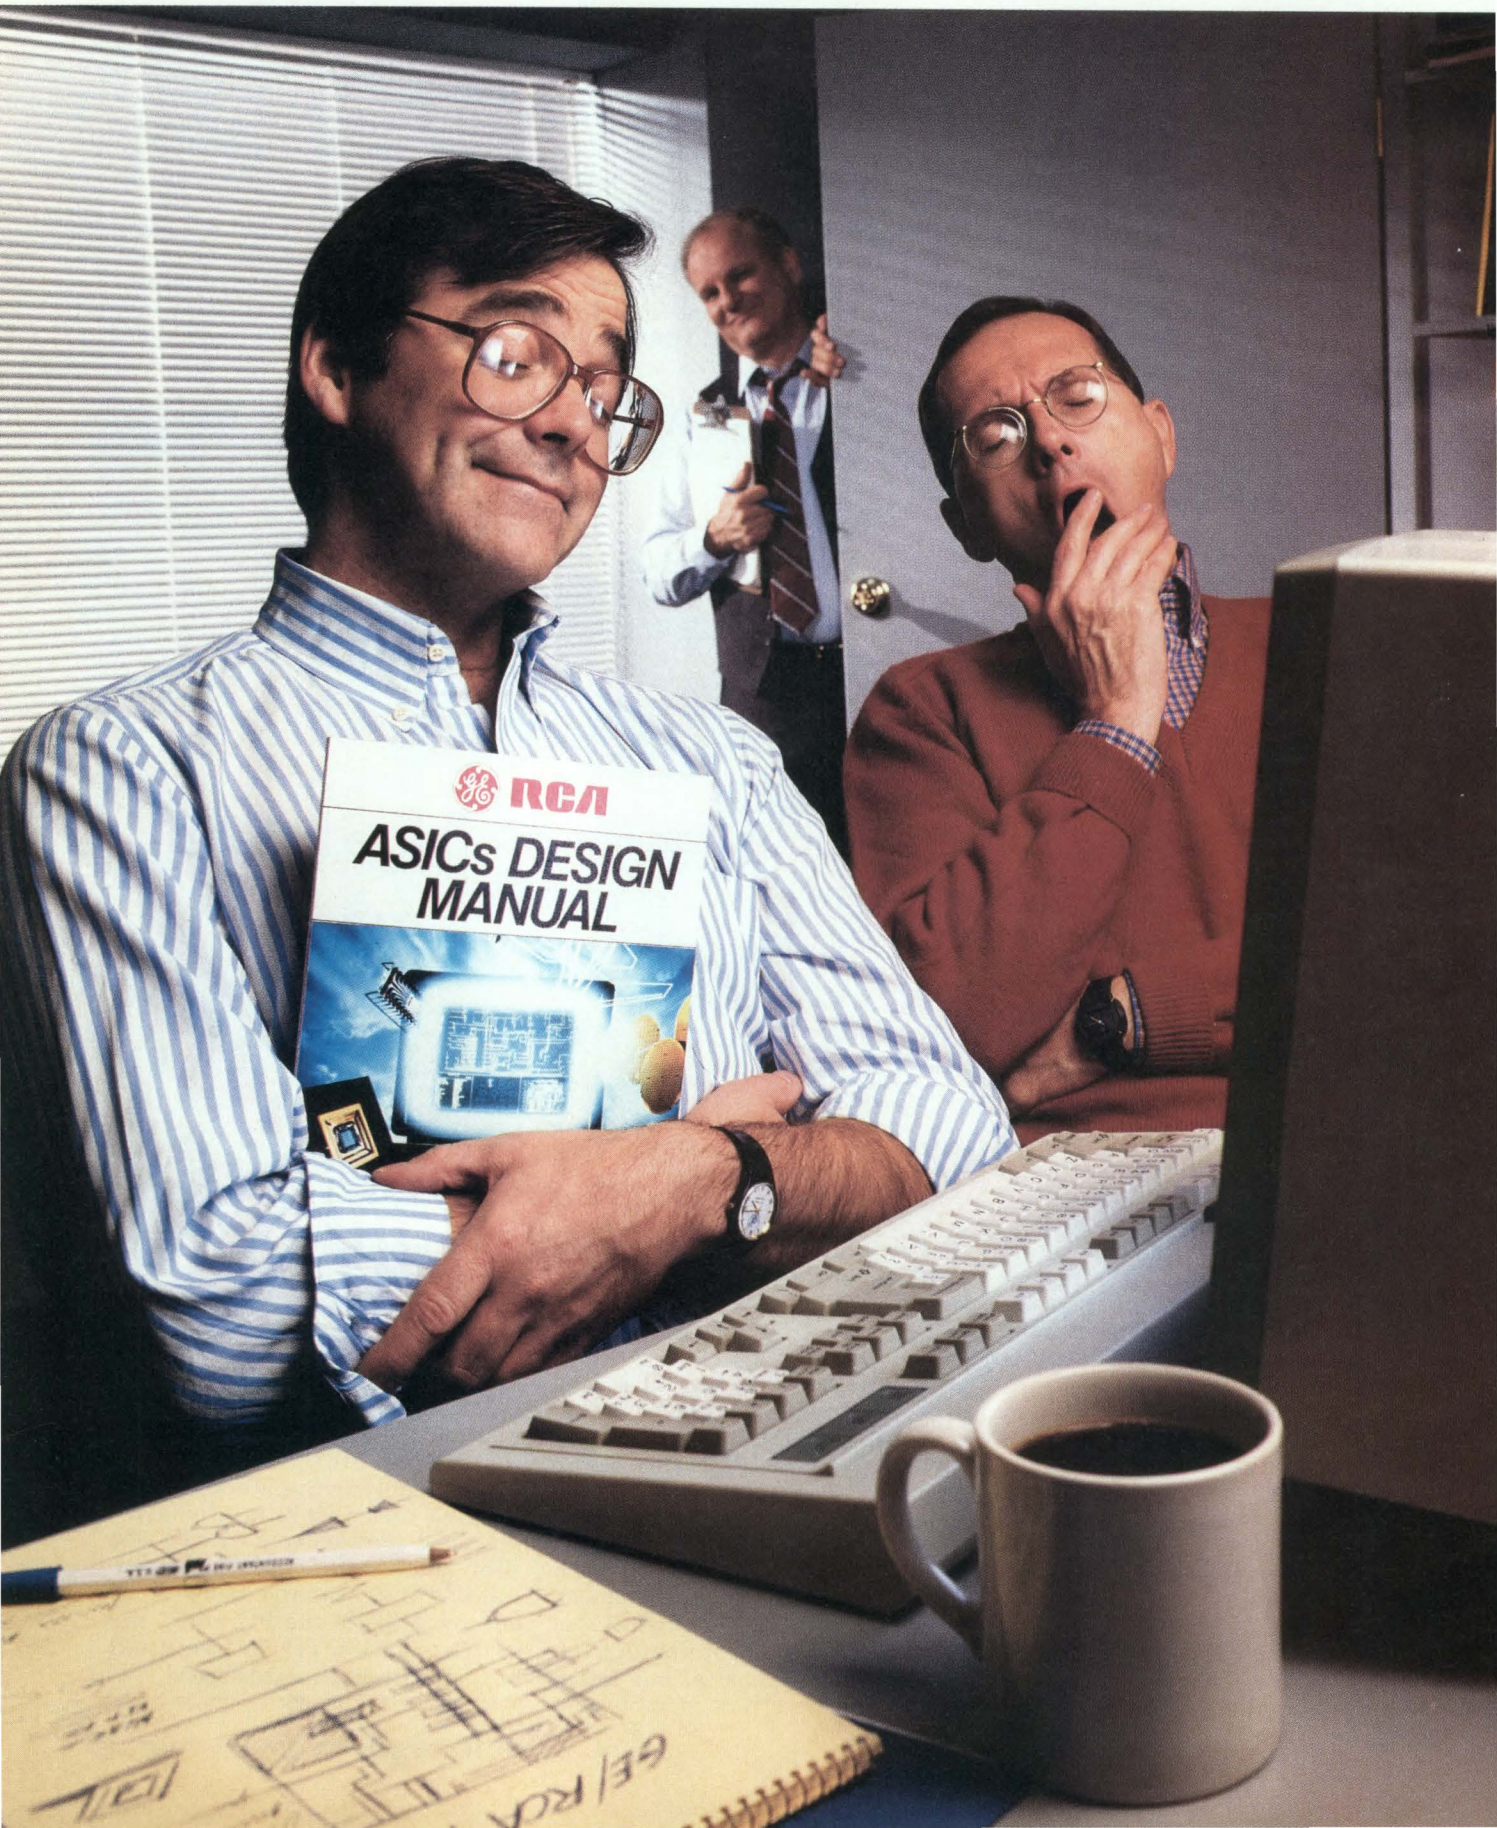

are boring.

They're easy to design. They're ready on time. And first-time success is virtually 100%.

You've heard all about the excitement of ASICs.

They improve performance, lower costs and make many new designs possible.

But, unfortunately, you've probably also heard about one big potential problem: while many ASICs pass the tests specified by the designer, they don't always work in the real world. And that causes excitement you can do without.

How to get first-time success.

It starts with our Design Simulation Software. It's been rated the best in the industry by the people who should know—designers who have used it. Within three days, you can be up to speed, working at any of the major workstations in the industry, creating and revising your ASIC with ease.

The standard cell advantage.

You'll really appreciate the power of our standard cells, which allow you to integrate a whole system, including macros, memories, logic and peripherals, onto a single chip.

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

You can choose from a wide range of Supercells, including the leading-edge RS20C51 core micro, RAMs, analog functions, bit-slice processors, HC/HCT logic, Advanced CMOS Logic, and high-voltage cells.

If they aren't enough, we can even generate Supercells to your specs.

And we're also in the forefront of silicon compiler technology. So we can offer you the ability to create designs that are heavily BUS-structured, with your ROMs, RAMs, PLAs and ALUs compiled right into the design.

We also bring you the resources of some very powerful partners, thanks to our alternate-source agreements with VLSI on standard cells; WSI on macrocells and EPROMs; and a joint-development agreement with Siemens and Toshiba on the Advancell® library of small-geometry cells.

Gate arrays, too.

If gate arrays are better for your design, you'll be able to choose from our full line up to 50,000 gates, with effective gate length as small as 1.2μ and sub 1 ns gate delays.

These gate arrays use "continuous gate" technology for up to 75% utilization. They are an alternate source to VLSI Technology arrays.

We also alternate source the LSI Logic 5000 series.

And we have a unique capability in high-rel ASICs, including SOS. Our outstanding production facilities here in the U.S. produce high-quality ASICs in high volume at very low costs.

It almost sounds exciting for something so boring, doesn't it?

For more information, call toll-free today 800-443-7364, ext. 25. Or contact your local GE Solid State sales office or distributor.

In Europe, call: Brussels, (02) 246-21-11; Paris, (1) 39-46-57-99; London, (276) 68-59-11; Milano, (2) 82-291; Munich, (089) 63813-0; Stockholm (08) 793-9500.

GE/RCA/INTERSIL
SEMICONDUCTORS

Network Troubleshooting?

How's it going?

Not well, from the look of it. It's not one of your better days. The system's down. All eyes are on you.

They're acting like it's your fault. It's not. But it's your job. You can go down with the system, or you can turn to the best Ethernet network analyzers on the market.

Before your day gets any worse, call Excelan and find out about the LANalyzer EX 5000 Series. We'll give you the details, and, for a limited time, a free poster version of this ad.

EXCELAN

2180 Fortune Drive San Jose, California 95131 800-2-GET-LAN FAX 408-434-2310 Telex 176610

Excelan Europe, Ltd. Intec 2, Suite 17, Wade Road Basingstoke, Hampshire, RG24 0NE England, UK
Tel: (44) 256-842296 FAX (44) 256-842298

Excelan and LANalyzer are registered trademarks of Excelan, Inc. Copyright © Excelan, Inc., 1988. All rights reserved.

CIRCLE NO. 16

Old and new contenders vie for position in 32-bit bus arena

The battle between the new 32-bit buses—particularly Micro Channel and NuBus—heats up, while VMEbus seems to have already triumphed over Multibus II.

David Lieberman
Senior Editor

Photo from SuperMac Technology; ray tracing by Jim Batson

Of the newer 32-bit buses, NuBus has attracted more than its share of graphics and imaging coprocessing boards, making it particularly suited for sophisticated image-processing applications. Some attribute this to the established graphics orientation of Apple's Macintosh computers.

You'd think that with the existence of VMEbus, Multibus II, NuBus, Futurebus, Fastbus, Micro Channel and others, designers would have a more-than-sufficient number of alternative 32-bit bus architectures to choose from. Yet, even as the 32-bit buses evolve and develop, there continue to be new 32-bit extensions to existing buses, as well as several new 32-bit buses—all looking to carve out a piece of the 32-bit processing pie. One begins to wonder how many 32-bit buses the market can support.

The latest 32-bit bus will be the Extended Industry Standard Architecture (EISA) bus—once the standard is completed, possibly by early next year. EISA was developed by a consortium consisting of AST Research, Compaq Computer, Epson America, Hewlett-Packard, NEC Information Systems, Ing. C. Olivetti, Tandy, Wyse and Zenith Data Systems, and is now supported by a host of other companies. It's hoped that EISA will rally those who have become disenchanted with IBM's radical departure from the PC, XT and AT buses. EISA's proponents claim to be the new standard bearers of the PC computing culture and hope to extend the technology without attempting to obsolete the AT.

Like EISA, another recent 32-bit effort incorporates AT compatibility. The C (for cache) bus developed by Corollary (Irvine, CA) is a very high speed (64 Mbytes/s) multiprocessing bus architecture equipped with a cache-coherency scheme. With a maximum 33-Mbyte/s data-transfer rate, EISA is better suited to single-CPU, though multiple bus master, designs. While the C bus is targeted at the low- to midrange multiuser system, EISA is aimed at the departmental computer serving as a shared resource for a network of PCs. Both EISA and C bus segment their architectures into memory and I/O buses; in both cases, the I/O bus is the AT.

Although these new buses will certainly stir up new bus battles, and plenty of VMEbus, Multibus II and Futurebus dramas are still unfolding, the most heated action in 32-bit buses today revolves around IBM's Micro Channel architecture. It's not so much that Micro Channel is hot—in fact, it isn't—but that there's a large office automation industry out there that's hot for a Micro Channel alternative.

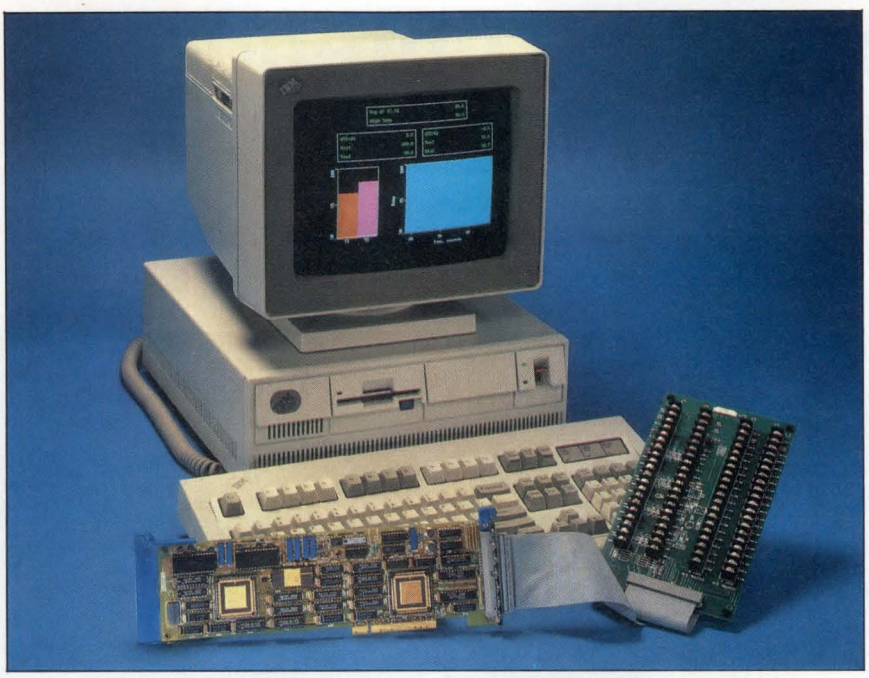

Micro Channel board makers face such difficulties as limited real estate, limited power allocation and rigorous mechanical demands. These difficulties have ironically driven some vendors to ambitious efforts, as in this Micro Channel board from Data Translation that simultaneously manages analog input, analog output and digital I/O operations.

Micro Channel's year-and-a-half-old brawl with "desktop" NuBus (as opposed to the older "classic" NuBus) continues in the desktop, industrial and laboratory arenas. Many companies have grown used to the very familiar benefits of building a system around the AT or XT, but they require higher performance than these industry standards will allow. NuBus supporters hope to attract that market, not by means of AT compatibility, but by combining 32-bit power with AT-like (though hopefully better) plug-and-play. These same supporters also have their eyes out for those who've become disenchanted with the complexity of VMEbus and Multibus II.

■ Clash of the buses

The primary vehicles for Micro Channel and NuBus are, respectively, certain models of the PS/2 family from IBM (Armonk, NY) and the Macintosh II from Apple Computer (Cupertino, CA). According to industry figures, from three to four times more Micro Channel machines have been shipped than Mac IIs. If the history of the AT is any indication, though, the bus that wins will be the one that attracts the greatest depth and breadth of hardware and software support, and support for the Micro Channel is noticeably lacking.

While IBM is apparently making things as difficult as possible for Micro Channel cloners, Apple has reportedly cooperated fully with industry efforts to develop and promote the NuBus, freely licensed by Texas Instruments, as a standardized open-architecture bus. "The trouble with Micro Channel is that no one is interested in getting into hassles with IBM," says Joe Ramunni, president and chief executive officer of Mizar (Carrollton, TX) and promotions director for the recently formed NuGroup, an industry association for NuBus manufacturers and users.

"NuBus delivers excellent cost/performance, and that's what it's all about, isn't it?"

—Bill Nowlin, National Instruments

■ ■ ■

The engineering community at large is somewhat enamored of NuBus and its much touted combination of power, simplicity, flexibility and ease of interfacing. On the industrial side, for example, NuBus has only one address space while Multibus II and VMEbus both have three, and only four control lines are required to define a NuBus transaction, while VMEbus requires 11 and Multibus II 10. "NuBus is a lot easier to use, but it still has high performance. And it's possible to build fairly inexpensive cards with very high bandwidth," says Bill Nowlin, vice-president of engineering for National Instruments (Austin, TX). "In other words, NuBus delivers excellent cost/performance, and that's what it's all about, isn't it?"

On the desktop side, NuBus is free of the more common difficulties encountered with Micro Channel, such as limited board real estate, stingy power-per-slot allocation and some rigorous mechanical demands. Some designers even take exception to the Micro Channel board layout. "It may not seem like a big thing," says Dick Pleau, product marketing manager at Data Translation (Marlboro, MA), "but the NuBus connector is in the middle of the board, while the Micro Channel connector is near the front panel, which means we have to have analog signals traverse quite a bit of the board. The NuBus board is simply laid out more logically; its definition was very carefully thought out."

Within a relatively short time, NuBus seems to have attracted more than its share of new graphics boards, frame grabbers, imaging boards, and array and DSP processors, while the Micro Channel has been almost entirely lacking in intelligent coprocessor boards. NuBus has become extremely popular for image-processing applications, although not for more cost-sensitive data-acquisition applications where the AT reigns supreme, according to John Molinari, Data Translation marketing director. "The great success of image processing on NuBus," he says, "is mainly because the Mac is a super platform for doing image processing. Its built-in graphics capability makes manipulating bit-mapped images much easier."

The graphics orientation of the Mac II, however, may or may not carry over into the world of NuBus computers at large, and the same is true of other Mac II features, some of which are perceived as shortcuts to accommodate the desktop form factor boards and the relatively low cost demanded of a desktop machine. One frequent criticism revolves around the Mac II's lack of a DMA facility, though NuBus has provision for sev-

eral types of DMA operations.

"DMA is crucial to high-speed data acquisition," says Molinari, "but not having DMA on the Apple box hasn't precluded us from doing very high speed applications." As Pleau explains, instead of taking up space with DMA hardware, Data Translation's NuBus boards include up to 1 Mbyte of dual-ported memory that's accessible on demand to the system CPU or other add-in board at reasonably high speeds. National Instruments, on the other hand, has implemented an eight-channel DMA server capability on some of its NuBus data-acquisition and instrument control boards.

Simplicity is key

Is there room in the market for NuBus? Many believe so. George White, Corollary president and a principal in the development of the bus, sees it as the only existing 32-bit structure that can comfortably bridge the industrial and desktop worlds. Mizar's Ramunni agrees.

"The unfortunate part of the other 32-bit buses is that they're all fairly complex and difficult to use," says Ramunni. "We're not talking about the need for another 32-bit bus in the industrial OEM markets that have traditionally innovated 32-bit buses; we're talking about an upsurge from the bottom, from the PC user and less sophisticated OEM who has gotten used to the form factor of the PC, likes the simplicity of plug-and-play and likes a simple operating system. NuBus is made for those people. And as soon as traditional OEMs realize that instead of having to hassle with industrial suppliers, they can just run down to Computerland to get their NuBus boards, the bus will really kick in. That will be its irresistible appeal."

As Berry Phillips, chief executive officer at Metrabyte (Taunton, MA) sees it, neither NuBus nor the Micro Channel enjoys a clear technical edge; rather, each represents a difference in orientation. "The basic NuBus spec is simple and elegant, yet very capable and versatile. But it tends to leave a lot more for the user to provide. The Micro Channel is more like the AT in that more is done for you by the system and you need to conform more to system requirements. Both accommodate 32-bit processors satisfactorily and provide a good 32-bit standard. NuBus does have an advantage in terms of high-speed transfers, and the Micro Channel has the edge in setup." That is, the Micro Channel uses hardwired interrupts and the NuBus doesn't.

National Instruments' Nowlin sees a trend in how his customer base is breaking into 32-bit computing: those who have been using the AT and those in "IBM only" shops move on to the Micro Channel, while those involved in start-up designs tend to evaluate all the available alternatives and wind up choosing NuBus.

The old battle fizzles out

While the new bus battles are just beginning, the old Multibus II vs. VMEbus battle has all but played out. In short, VMEbus won; it's the premier 32-bit bus today. Compatibility concerns have, for the most part, been put to bed as designers have standardized on the C.1 revision and board vendors and users alike have gained more experience with the bus. "At least now we can design to a target that isn't moving quite as fast," says George Schreck, product manager for Sky Computers (Lowell, MA).

That's not to say that VMEbus is stagnating. The level of new product introductions is extremely high, and there's a movement afoot to push performance closer to the limits of what the bus will allow. "We haven't seen many people using VMEbus to its full potential," says Schreck, "mainly because the 6U form factor doesn't provide enough real estate to support the extra functionality." But, better use of VMEbus will accelerate with the availability of interface application-specific ICs from a variety of vendors.

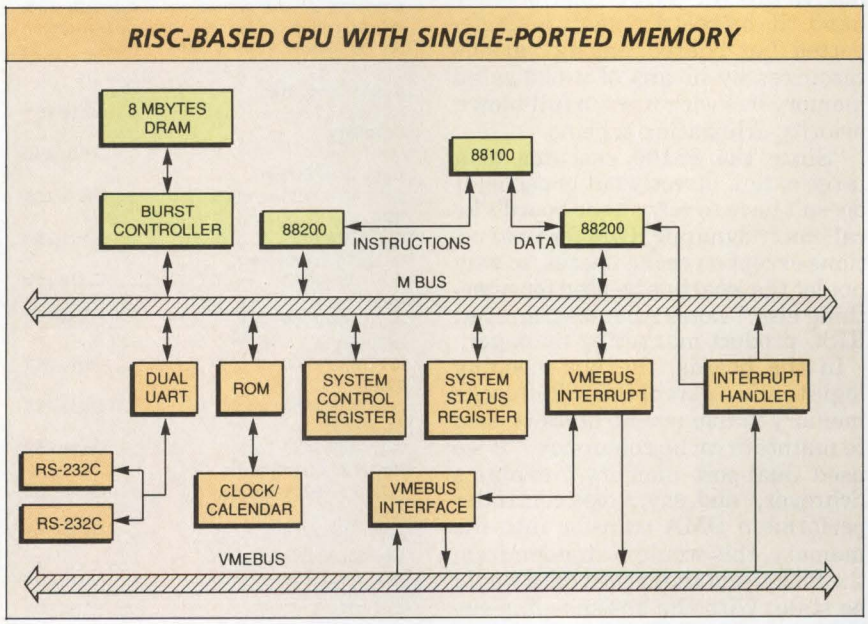

In a typical CISC CPU, on-board memory is dual-ported to allow access by the on-board processor as well as other masters on the system bus. Motorola's first RISC-based CPU boards, however, provide single-ported memory by means of the 32-bit M (for memory) bus. The M bus provides a route for sharing on-board memory without violating cache coherency of the 88200 CMMU (cache memory-management unit) chips.

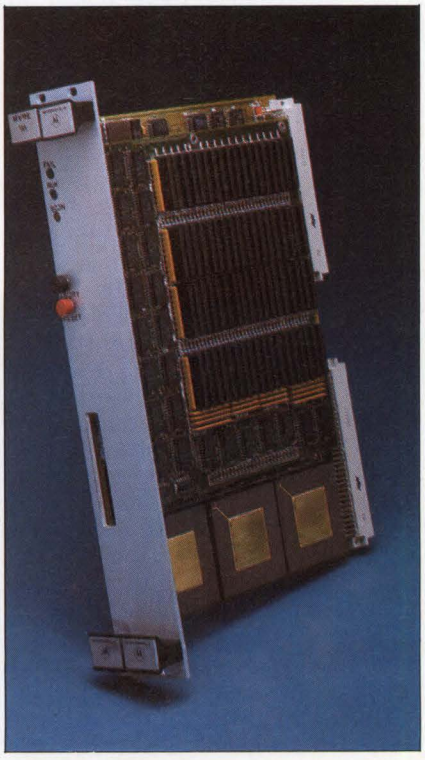

Besides freeing up valuable real estate, the bulk of the new chips also implement bus extensions such as message passing, which should have a significant impact on overall system performance. Vendors of VMEbus board products are also working to make better use of such bus functions as block-mode transfers, thus more efficiently using bus bandwidth and address modifier codes to support serious memory partitioning. "We're really just now learning how to make effective use of the tremendous amount of I/O bandwidth that VMEbus provides," says Bill Gage, president of SBE (Concord, CA).

There's also a significant movement afoot within the VME community to establish software standards for real-time systems. The rationale here—to simplify development and to speed time to market—is commendable, but it may be too late; there are too many vested interests among the players for the effort to be successful. John Rynearson, vice-president of engineering at Mizar and chairman of the VME International Trade Association (VITA) software subcommittee, is hopeful, however, that the clear market benefits of software standards will lead to a successful outcome for this effort.

The VMEbus community has also begun addressing itself to the bus requirements of future computing systems—particularly those based on RISC microprocessors. It's highly possible that, instead of developing a new bus from scratch, the community will eventually endorse a bus that's been long waiting in the wings: the Futurebus. Yet there's an equal chance that a compatible, extended VMEbus could be developed.

Though yet to gain a champion in the market or to be used to any significant degree, the Futurebus hasn't been stagnating either. This year, the IEEE Futurebus committee began to address itself to 64-, 128- and 256-bit extensions to the existing spec, as well as to a "noncompelled" handshake mode of operation that would let pipelined transfers occur between Futurebus boards.

To date, the acceptance of Multibus II has been disappointing, though proponents say that Multibus II activity is simply less visible than that in VMEbus. Next year will see a number of traditional minicomputer vendors move to Multibus II as their I/O bus standard, according to Steve Cooper, manager of the architecture proliferation group for Intel's OEM Modules Operation (Hillsboro, OR). It's also becoming clear that a good deal of activity in Multibus II boards and OEM systems is brewing, most of it in the multiprocessing arena.

■ Single-ported memory

Not all 32-bit bus activity this year took place on backplanes, however. Some very interesting work also was done on individual boards. In its first RISC effort, for example, Motorola's Microcomputer Division (Tempe, AZ) has implemented a new memory architecture on its entry-level 88000-based VME181 and VME181-1 CPU boards, employing the 32-bit M (for memory) bus.

"There's more of a show-me attitude among board companies today when a new bus becomes popular for whatever reason."

—Bob Birenbaum, Imaging Technology

■ ■ ■

In a typical CISC CPU, on-board memory is dual-ported to allow access by the on-board processor as well as by other masters on the system bus. In this RISC CPU, on the other hand, the on-board memory is single-ported but shared among various resources by means of a dedicated memory bus with its own full-blown priority arbitration scheme.

"Since the 88100 executes to a large extent directly out of cache, it doesn't have to rely on the board's local bus or dynamic RAM for instructions except on cache misses, so why not let the local bus be used for something else?" notes Andreas Schreyer, RISC product marketing manager.

In the boards, the bus-snooping logic that's a part of the 88200 cache memory-management units is used to maintain cache coherency. "If we used dual-port memory," explains Schreyer, "and, say, a disk controller performs a DMA transfer into the memory, this would be hidden from the 88200 and the cached data would be stale. With the bus-snooping circuitry continuously observing M bus traffic and detecting when a cached memory location is being modified, the appropriate cache entries will be invalidated."

Since the 181 boards don't implement cache coherency for the VMEbus as a whole, they'd be inappropriate for tightly coupled multiprocessing systems, Schreyer explains. But Motorola has another type of computing in mind: tightly coupled multiprocessing on a single board, perhaps with as many as four 88000 chip sets sharing one 6U-sized board.

■ The bottom line

With so many 32-bit alternatives popping up these days, a lot of companies are waiting and watching, trying to determine a clear market direction before they commit their resources to a new bus. Technical preferences notwithstanding, the decision to support a bus structure tends to be a hardheaded business decision about its market viability.

"The expense of adding support for a new bus has risen quite substantially," says Bob Birenbaum, director of marketing at Imaging Technology (Woburn, MA), "as has the risk to some extent. There's more of a show-me attitude among board companies today when a new bus becomes popular for whatever reason. Many of us would rather wait for demand to develop than proactively support any new bus that appears." □

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

Apple Computer (408) 973-2266	**Circle 251**
Corollary (714) 250-4040	**Circle 252**
Data Translation (508) 481-3700, ext 505	**Circle 253**
IBM (802) 769-1621	**Circle 254**
Imaging Technology (617) 938-8444	**Circle 255**
Intel (503) 696-7259	**Circle 256**
Metrabyte (617) 880-3000, ext 230	**Circle 257**
Mizar (214) 446-2664, ext 6412	**Circle 258**
Motorola (612) 438-3437	**Circle 259**
National Instruments (512) 250-9119, ext 72	**Circle 260**
SBE (415) 680-7722, ext 145	**Circle 261**
Sky Computers (508) 454-6200, ext 350	**Circle 262**

Please rate the value of this article to you by circling the appropriate number in the "Editorial Score Box" on the Inquiry Card.

High 500 *Average 501* *Low 502*

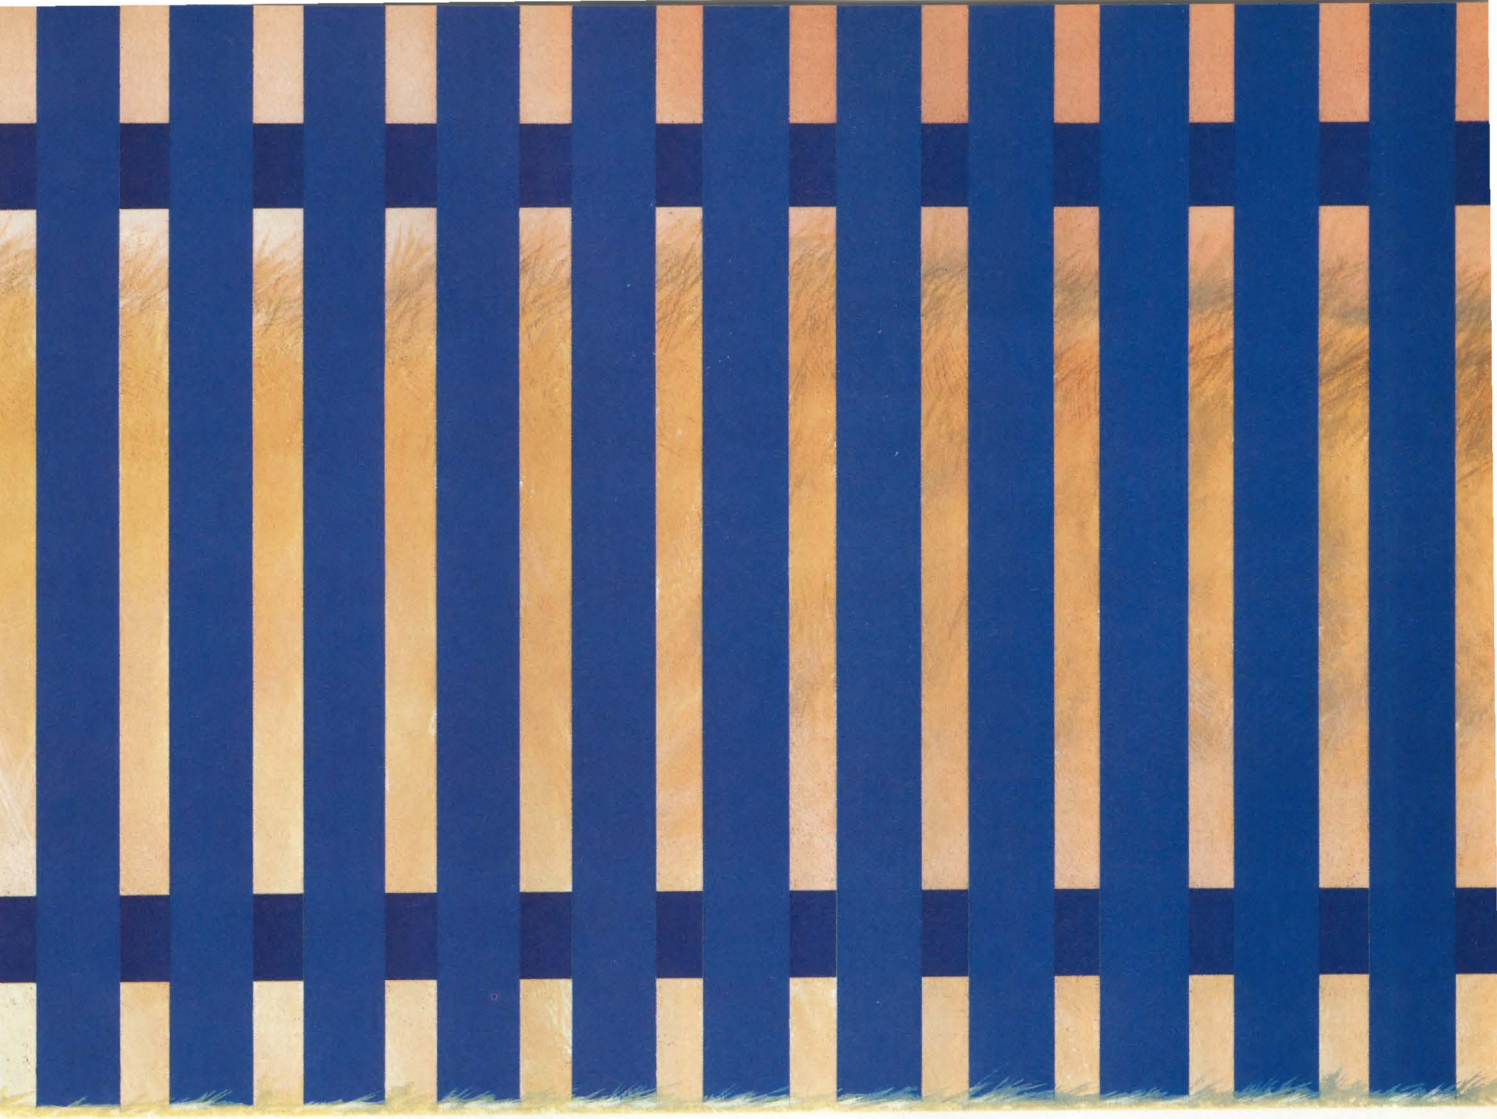

VGA

Others promise 100% register-level compatibility.

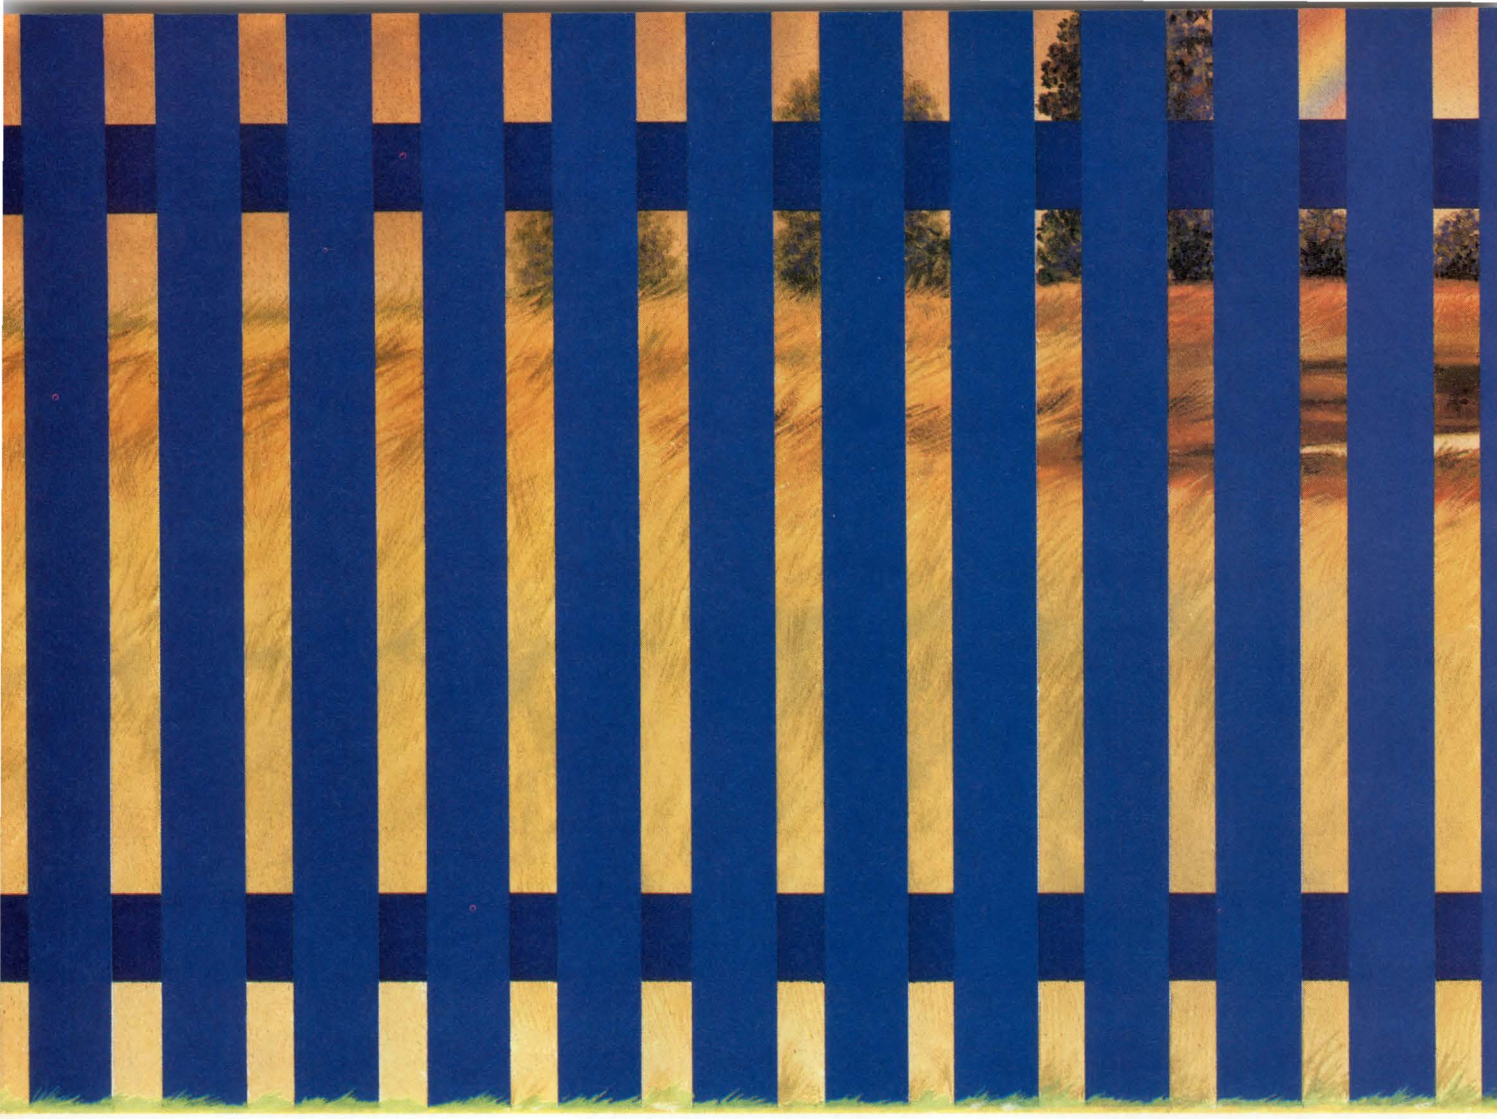

Only the CHIPS/450™ family gives you VGA™ compatibility at the gate level.

Everyone claims their VGA-compatible chips offer 100% compatibility with the industry standard. Which only means they've implemented VGA functions at the register level, right out of the IBM® reference manual.

But to achieve *real* VGA compatibility, it takes a much sharper focus. Down to the individual silicon gates used to build each and every register. Only one family of VGA CHIPSets™ does that: CHIPS/450.

That's why only the CHIPS/450 family can guarantee 100% VGA hardware compatibility, including implementation of hidden modes and other functions not thoroughly described in IBM's PS/2® Technical Reference. Features like a high-resolution mode, special zoom functions and readable CPU latches. Features that IBM and software developers might make use of in the future. Sending your design back to the drawing board if your "VGA-compatible" chip set isn't *gate-level* compatible.

Leave others at the starting gate.

Besides uncompromising compatibility, the CHIPS/450 family gives you gate-crashing graphics performance. Especially when tightly coupled with our industry leading system logic CHIPSets. Both our VGA and Super VGA controllers offer a proprietary 16-bit interface and *fast* VGA cycles—delivering *twice the performance* of the industry standard. Both include enhanced backward compatibility with all the popular video modes, and special protection

We give you the gate.

for sensitive VGA video control registers.

Get set for the future.

The CHIPS/450 family gets you ready to meet the performance challenges of today. But more important, CHIPS gives you a seamless upgrade path to tomorrow's standard: the OS/2™ Presentation Manager.™

Our Super VGA incorporates additional hardware enhancements that maximize Presentation Manager performance. Features like a hardware cursor to reduce windowing overhead. A paged memory interface that increases bandwidth. And since our VGA is pin compatible with our Super VGA, you can upgrade quickly to next-generation products. While protecting your original VGA investment.

So if you're checking out VGA chip vendors, get in touch with the only one that offers gate-level compatibility and high-level performance. Call us at **1-800-323-4477** for details on a hardware/software path that's not going to fence you in.

CHIPS

chip sets or CHIPSets.™ The difference *is* the difference.

Chips and Technologies, Inc. 3050 Zanker Road, Dept. M-3, San Jose, CA 95134. IBM, PS/2, OS/2 and VGA are registered trademarks or trademarks of International Business Machines Corporation. Presentation Manager is a trademark of Microsoft Corporation. CHIPSet and CHIPS/450 are trademarks of Chips and Technologies, Inc. © 1988, Chips and Technologies, Inc.

CIRCLE NO. 18

Robotic systems learn through experience

Sydney F. Shapiro, Senior Editor

Increasingly sophisticated real-time software, intelligent sensors, image-processing subsystems and neural networks will enable robots to "think" and interact with their environments.

Other than on the pages of science-fiction books or at the hands of television or movie writers, robots have never been portrayed as humanoid in appearance. Though they may have armlike appendages, real-life robots look exactly like what they are: machines. No matter how much glamour is built up around robotics, a robot is still nothing more than a sophisticated machine tool.

But unlike the robot's "dumb cousins"—ordinary machine tools—the robot has the potential for what can almost be considered a brain. But that brain can't really think, at least for now. It still can't operate on its own or order functions other than those it's told to order. Like its dumb cousins, a robot is entirely dependent upon input from some external source: a program that tells it when to start, what steps it should take while in operation and when to stop. But robots have the potential for artificial intelligence in the form of expert systems—or neural networks.

Depending upon the specific application, these sophisticated machine tools may function in a stand-alone mode or as part of a system made up of several other robots and their dumb cousins. Whichever configuration is chosen, commands must be relayed from a predetermined program to the servo motors or other devices that make the robot's manipulators perform programmed functions.

The basic requirement of robotics is software—both operating systems that translate commands to the robot and possibly some form of artificial intelligence that lets the robot make some controlled decisions of its own. Robotics also involves tactile and positioning sensors. Finally, a communications system is needed to carry the I/O signals between the robot controller and the servo motors on the manufacturing operation.

Software becoming standardized

Robot manufacturers, particularly those based in the United States, have found it nearly impossible to reach a break-even level of profitability. Part of their problem is that they've been producing robots that have been far more sophisticated than is really necessary for the relatively simple jobs that users have wanted them to perform.

But even more damaging has been manufacturers' tendency to base operation of their robots on proprietary software that usually would work when the robots were used in a stand-alone mode but wouldn't necessarily work when they were used as part of a system of robots. In addition, sometimes the program for a stand-alone robot would be on a separate magnetic tape that had to be inserted in the robot controller each time the robot had to perform a different task. This has deterred potential major customers, such as car manufacturers (where the majority of robots are used today), from purchasing such robots because it isn't cost-effective to run different software on different robots in their systems.

Those robot manufacturers that have managed to survive have reached a crucial point where they must consider the next generation of robots. Despite their large investments in proprietary software, many manufacturers realize that the only acceptable solution is to design their robots to work under standard operating systems and associated development tools. These so-called standard systems are based on established bus structures, such as STD Bus and AT bus. Although they still don't necessarily work with robots that have controllers that are built around proprietary buses, they're compatible with far larger numbers of machines than before.

Even more critical, however, is the

need for the software that lets robotic functions be performed in real time. A real-time computer system, whether it consists of a mainframe controller or an embedded microcomputer in a robot, must respond to events that occur in the process being controlled. It also must be able to prioritize any corrective actions and to reorder tasks within the process. Until recently, true real-time software for robotics was either nonexistent or so unique that it was impractical for routine manufacturing applications. But now, enough software tools are available for system designers to be able to choose the one that best fits the requirements of their systems.

Capabilities that must be part of real-time systems, unlike other types of computer systems, include guaranteed response to real-world events, priority-driven multitasking, reentrant programming languages and code, extremely fast response to interrupts, support for system programming and special I/O devices, and facilities for special real-time programming. Although the definition of "extremely fast" in relation to response time differs from system to system and from application to application, real-time system software typically responds to interrupts be-

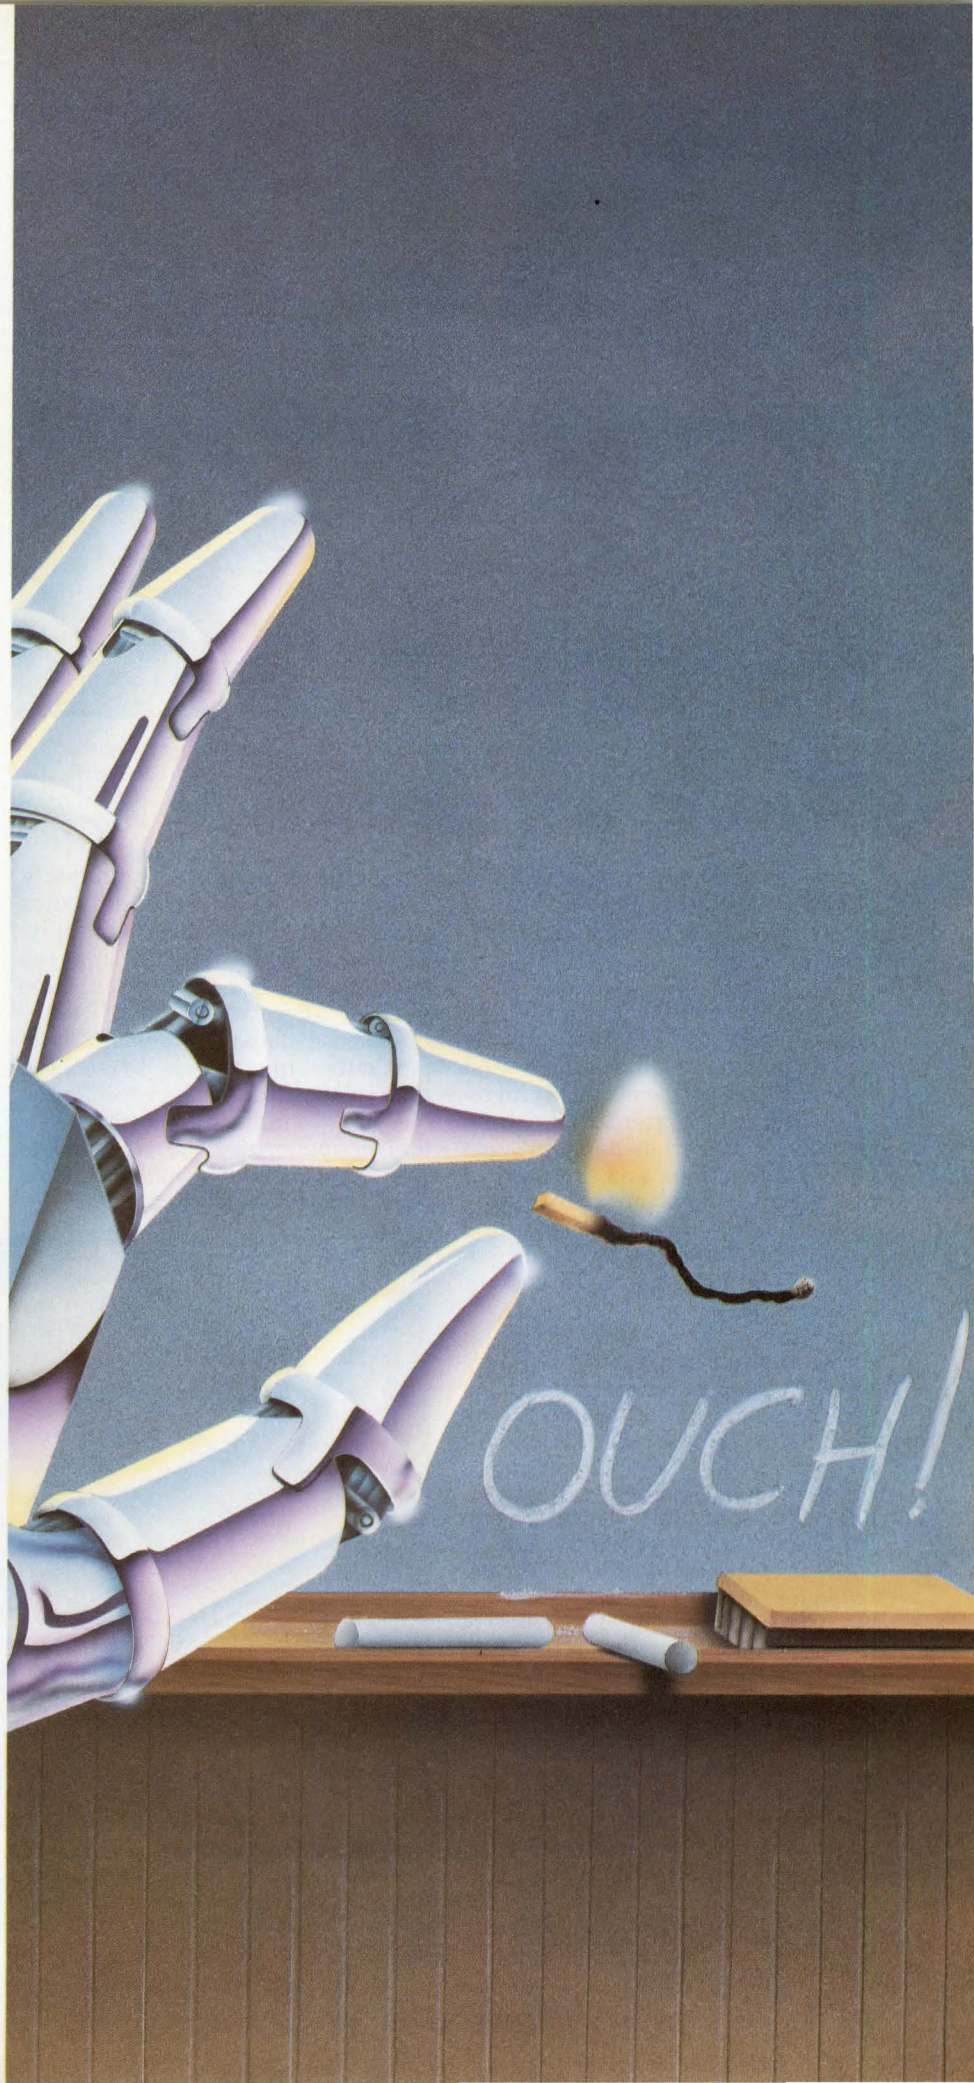

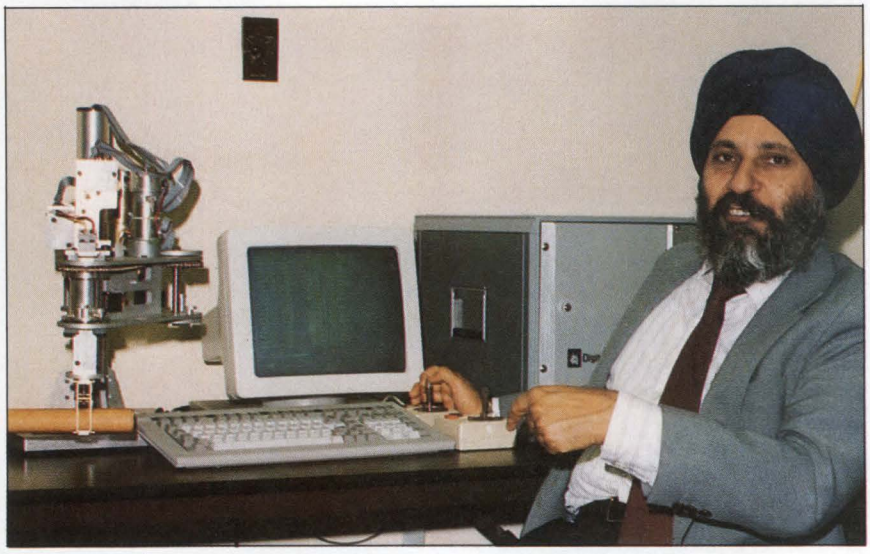

Lynx Real-Time Systems developed a real-time Unix-compatible operating system specifically for control applications that are relevant to robotics, according to Inder Singh, president. Based on a real-time, multitasking, multiuser kernel, LynxOS interfaces with a number of the networking schemes commonly found in factory environments.

tween 100 and 500 times faster than non-real-time system software, such as the standard time-sharing versions of Unix.

■ Working in real time

Real-time operating systems provide a variety of important functions and can be used in both embedded and reprogrammable systems. Real-time kernels, on the other hand, provide minimal functions and are used in embedded applications. Since they perform fewer functions, real-time kernels react even faster than real-time operating systems in certain applications.

Because of the dominant position its VAX and MicroVAX computers have attained in industrial and factory floor applications, Digital Equipment Corp (Marlboro, MA) is a familiar name to designers of systems for industrial applications. And that often influences their choice of real-time software—specifically DEC's VAXELN operating system—for developing robotic systems.

VAXELN, introduced in 1984, is intended for real-time applications from the lowest level of the factory floor, which includes the servo motors and embedded controllers in the robots, through the unit-control level, and up to the area-control level. Version 3.1, the latest release of the VAXELN real-time application development system and run-time kernel, which was introduced earlier this year, will run on any VAX processor with at least 1 Mbyte of memory. This nonpaging, completely memory-resident system eliminates the need for disk memory. One of its key features is its ability to treat the physical memory of any real-time VAX or MicroVAX as if it were a high-speed disk. This feature is of particular importance to robotic-system applications, because it enables use of a real-time operating system for applications in environments that are common to robots but too harsh for disk subsystems.

Intel's OEM Modular Systems Operation (Hillsboro, OR) offers both a real-time operating system and a real-time kernel. The iRMX II.3 operating system supports both the 80286 and 80386 microprocessors and the 80287 and 80387 math coprocessors. It has 16-Mbyte memory addressability and supports multiprocessing, as well as multiple tasks, jobs and users. It can function as a nucleus; a named, physical and stream file system; a basic or an extended I/O system; a bootstrap and application loader; and a system debugger.

A nucleus communications service provides a software interface between application code and a Multibus II message-passing coprocessor. This is intended to simplify message transmission between tasks on different boards and to provide a standard software interface to other Multibus II boards in the system. Programs and languages can be transported via a universal development interface (UDI) to or from iRMX II.3 and other operating systems that support the UDI standard.

A human interface, made up of a set of system calls, a set of commands and a command line interpreter, lets the operating system support multiple users. This interface lets the users develop applications, maintain files, run programs and communicate with the operating system. In addition, iRMX networking provides com-

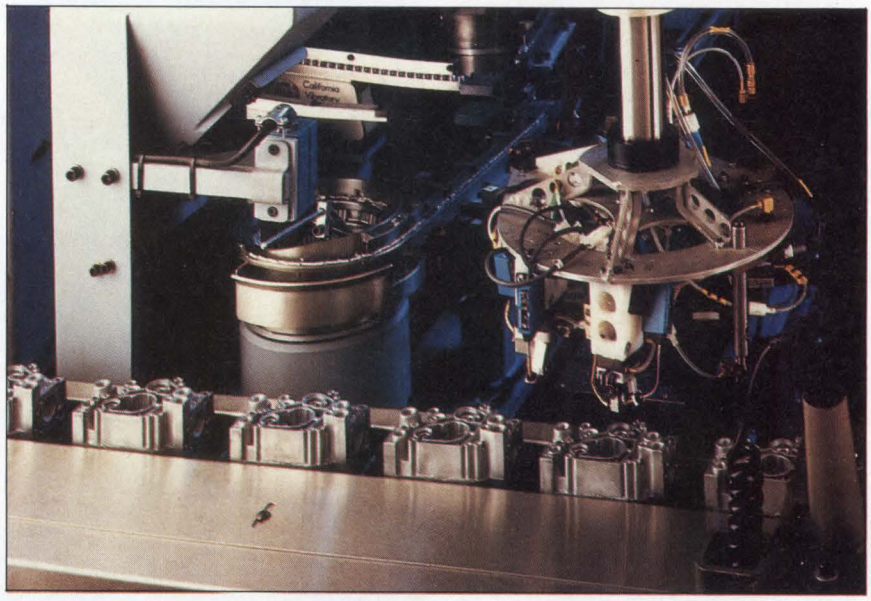

Micro Switch

A multitude of sensors are required for any robotic application. A robot arm at a Honeywell plant, for example, assembles gas valve housings with the help of fiberoptic photoelectric sensors (blue objects). A series of fiberoptic tips (not shown here) serve as the robot's eyes to determine proper positioning of the arm so that it can insert the components seen in the conveyor at the left.

Neural-network systems: the key to adaptive robotics

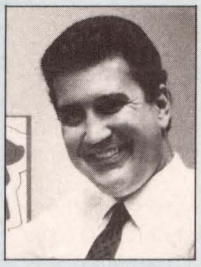

Over the past several years, a great deal of research and development time and money has been spent on adaptive robotic systems. Such robots are particularly valuable for performing tasks that involve a high degree of risk for humans, such as space exploration or the handling of toxic substances.

Until recently, however, the majority of work in the area of adaptive robotic systems has involved the use of rule-based logic. In such adaptive systems, the robot must know how to react or adjust to changes in its operating environment. A robotic manipulator on a spacecraft, for example, might need to recognize and handle a certain object or to learn a particular series of motions that it's never encountered before. Typically, the engineer defines a model that includes the logical and mathematical rules necessary to perform the task. The robot is manually guided through the initial series of maneuvers and remembers how to perform the task.

This approach works reasonably well for routine tasks. But in most cases, rule-based teaching approaches have fallen short of providing a robotic system that can truly learn to adapt to changes in the environment.

Several variances now must be considered in order to be able to provide a truly adaptive robotic system. First, the system must be able to orient a robotic arm in redundant degrees of freedom. There's probably no unique solution to this problem because there are more unknowns to be considered than equations to allow for them. Second, the system must compensate for changes that occur because of mechanical wear. Third, the system must integrate sensory inputs that may be noisy and highly variable depending upon the orientation and location of the sensors.

One of the most promising solutions to these and other complex issues in adaptive control is neural-network systems. Already in use for many complex pattern-recognition applications, they may be the key to developing true adaptive robotic systems.

The Nestor solution

A neural network is an information-processing system composed of a large number of interconnected processing elements (cells). Each element computes its output locally, based on the collective inputs from the cells to which it's connected and on the strengths of those connections. The Nestor Learning System (NLS), a neural-network computer model that emulates human learning and classification processes, facilitates the combining of multiple neural networks for effective pattern recognition.

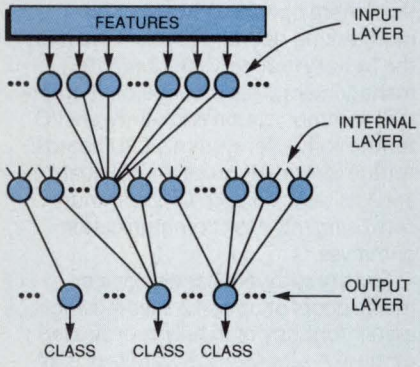

Actually, NLS is a hierarchy of parallel neural networks, each made up of three layers of cells: input, internal and output (as shown above). Cells of the input layer register values of features describing an input event. Cells in the internal layer of the network construct the mapping that ensures that the output cell for the correct classification fires in response to a given input pattern. The several neural networks in a given application may view different features of an input pattern and cooperate in its classification.

This NLS is particularly useful for pattern recognition and signal classification, domains where data sets are so large or complex that rule-based expert systems, programmed algorithms and public domain neural networks are impractical. It's also useful for data that's noisy, or highly variable over time.

The NLS preprocesses an input pattern to generate a set of features, the measurements of which characterize the pattern. Such measurements can be regarded as a collection of points in a feature space. A pattern of activity among the cells of the input layer corresponds to the location of a point in this feature space. All examples of a pattern class define a set of points that map out a region having some arbitrary shape.

NLS learning mechanisms

Learning involves two distinct mechanisms. From seeing training examples, the system dynamically allocates new cells to the internal layer of a network. These cells are allocated with an influence field surrounding them. The NLS generalizes by attempting to classify new examples that are close to stored internal layer cells (within the influence field, for example) as being in the same class as the stored cell. Each cell in the internal layer is connected to each cell in the input layer, but to only one cell in the output layer. This technique lets the system map out class regions even if they're disjoint and not linearly separable.

The second learning mechanism is the modification of the influence fields associated with cells in the internal layer. When a new example occurs that shows a conflict with an influence field of a different class, the conflicting influence field is modified to eliminate that conflict, thus mapping the boundary between the two class regions more precisely.

The allocation of cells and the modification of the internal cells' influence fields are controlled by training signals that move from the output layer back into the system. The system thus learns automatically by being exposed to examples of the various classes.

The NLS is a multilayer, feed-forward system with low connectivity within each layer and no relaxation procedure for determining an output response. This architecture lets the NLS operate in real time without special computers or custom hardware. Dynamic cell allocation lets the NLS be easily scaled up to any size problem or data base within the limits of the computer's size.

Michael G. Buffa, President and Chief Operating Officer, Nestor

Vision-guided robot system uses real-time machine vision

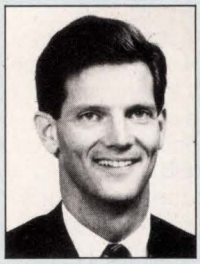

One of the major limitations preventing the widespread use of robotic systems in automated manufacturing has been the inability of conventional robots to correct for process variations or for part-preparation errors. To compensate for such variations or errors, systems must include some form of real-time sensor feedback.

The Automatix Vision Guided Seam Tracker utilizes real-time machine-vision processing to measure both the position and the dimensions of a weld joint. The joint measurements serve as feedback to control robot position and welding parameters, thus producing a quality weld.

The system uses a sensor consisting of a structured light source and a charge-coupled device television camera in an industrially ruggedized enclosure. The structured light source projects a thin stripe of laser light on the weld joint, highlighting its surfaces and key inflection points and creating a local topographic map of the weld joint. The laser-highlighted image of the weld joint is easier to analyze with machine vision than is an ordinary image of the weld joint. The camera and laser also can be easily calibrated to provide three-dimensional coordinates of any point on the laser stripe.

The robot and the welding process are controlled by a computer system that combines the company's robot control and vision-processing hardware in one cabinet. The integrated vision controller contains a master processor board that provides an operator interface and high-level control of the vision-processing and robot control systems; a vision-processing board that contains a frame grabber, 2 Mbytes of frame buffer memory and two microprocessors; a slave processor board for robot control; a robot analog I/O (servo) board; and a welding I/O board.

Software for the vision-guided seam-tracking system can be divided into four functional components: user interface, robot control, vision measurements and weld process control. Each of these functions is implemented as one or more separate software tasks that communicate with each other via message passing under the Ambre operating system. This multitasking operating system provides the basic system services of memory management, task management, intertask communication and hardware I/O for the computer system's CPU boards. Ambre is resident on each CPU board and lets tasks on each CPU communicate using intertask communication primitives.

Control software handles one or more robots of up to 12 axes moving asynchronously or in fully coordinated motion. A separate software task is associated with each robot in the system, letting multiple robots move asynchronously. If coordinated motion is required, the robot tasks pass messages among themselves to provide synchronization. A standard process/sensor interface lets each robot task control processes such as welding. These processes may need to be coordinated with the robot motion and to communicate with associated sensors that can be used to alter the speed or goal position of the robot. The sensor interface allows for modular addition of a variety of processes and sensors to the robot system.

The motion control tasks plan the trajectory for the robot or robots to be moved and send this information to a servo task, which resides on a separate processor board. The servo task continuously updates the position of each of the 12 axes under system control.

■ The image-processing pipeline

Seam-tracking software is functionally divided into a measurement task and a control task. The measurement task includes the steps of acquiring an image of the weld joint from the sensor, extracting a joint description from the image and then converting the image coordinates into a local 3-D model of the weld joint. The seam-tracker control task receives joint measurements from the measurement task and generates a control signal for the robot and the welder.

The measurement task is divided into seven software tasks distributed over three CPUs operating in a pipeline. The first CPU in the pipeline performs image acquisition and preprocessing, the second extracts the joint description, and the third constructs the 3-D joint model. The pipeline is divided into seven functional tasks to make it easier to adapt the system and add CPUs if a performance increase is required. Joint measurements are performed asynchronously relative to the other system-level tasks at a rate of 2 to 5 Hz in the current

Joseph Campbell, General Manager, Machine Vision and Controls Division, Automatix

munication between the real-time application and general-purpose systems, as well as other real-time systems.

The iRMK Version I.1 real-time kernel, a 32-bit real-time executive, was designed and optimized for protected-mode operation with the 80386 microprocessor. It also provides optional support for an 80387 math coprocessor and other peripheral devices. Although optional Multibus II message-passing support is provided, this kernel works with any bus. By default, the kernel and its application execute in a flat memory space of up to 4 Gbytes and in a single-privilege level.

Services provided by the kernel for real-time applications include task management with system calls to create, manage and schedule tasks in a multitasking environment; interrupt management by immediately switching control to user-written interrupt handlers when an interrupt occurs; and time management to provide single-shot alarms, repetitive alarms and a real-time clock.

■ Unix in real time

Unix originally was intended to be a general-purpose, time-sharing operating system with no provisions for the fast response to external events that's needed for real-time applications. Despite this, Unix has several advantages. It's available on a wide variety of computers; it's portable; it

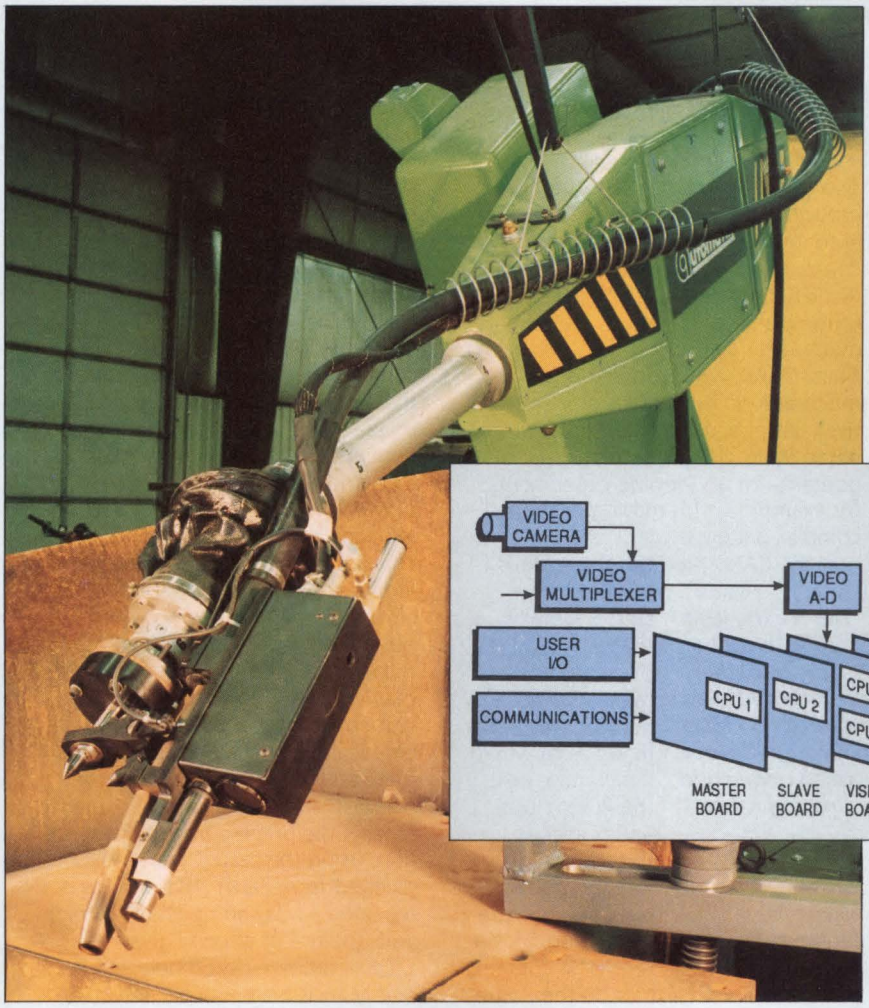

The Automatix integrated welding robot controller uses real-time pipelined vision processing to analyze images of the weld seam and to provide 3-D data that's used to control the robot's motion and the welding process.

hardware configuration.

The robot control signal contains an offset from the path that the robot was commanded to follow and an adjustment of robot speed. The weld control signal includes variables such as welding voltage and current. Robot offset is determined by finding the point on the commanded trajectory to which each visual measurement corresponds and then determining the vector distance from the commanded point to the visually observed point. Speed and welding parameters are determined from a user-defined table containing parameter setpoints for different ranges of weld-joint measurements.

Control tasks are executed synchronously with the robot motion tasks. For every new commanded position of the robot, offset, speed and weld parameter sets are generated by the seam-tracker control task. Because the seam-tracker system software is modular, the implementation of new features and capabilities is straightforward. This lets the system work in a wide variety of welding processes and weld-joint configurations.

provides a good development environment; and it provides a lot of conductivity support so that computers can be easily networked.

Several Unix-based operating systems have become available in recent years that either provided much greater performance than the original Unix or offered real-time kernels with libraries to implement a subset of the operating system's functionality. For control applications that are relevant to robotics, however, a real-time Unix-compatible operating system must contain additional features such as fast, predictable interrupt response, fast worst-case task switching, user-fixed task priorities with preemptive scheduling, and fast I/O.

Those features are provided in LynxOS, a fully Unix-compatible real-time operating system from Lynx Real-Time Systems (Campbell, CA), according to Inder Singh, the company's president. The first LynxOS systems delivered went into robotic applications in real-time situations, Singh says.

LynxOS is fully compatible with industry-standard versions of Unix, such as Berkeley 4.2, as well as with the IEEE Portable Operating System for Computer Environments and the AT&T System V Interface Definition. It's now available on the company's own 80286-based system and will be available in the future on 80386-based systems. It also can be used with 68010-based systems, as

Real-time software controls Mars Rover robot

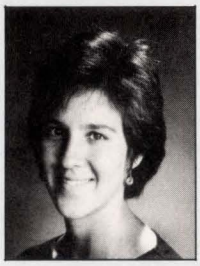

We on Earth have always been fascinated by our neighboring planet, Mars. But, despite space probes and flybys, we've had to depend on scientific conjecture as to what Mars is really like.

Before the year 2000, however, we'll be able to gather information directly on Mars through the use of an unmanned roving robot. The Robotics and Teleoperators Group at Jet Propulsion Laboratory (Pasadena, CA) is developing a stand-alone real-time robotic system with only one communications link with humans: a global map and path directions that it receives from Earth. Jet Propulsion Laboratory (JPL) has already tested navigational algorithms on an experimental version of the robot, known as the Rover, which has six wheels, each over 1 m in diameter; measures 2 m wide, 2 m tall and 4 m long; and has a mass of over 750 kg.

The Rover's interface to its immediate area will include camera eyes, a manipulator arm for collecting soil samples and performing experiments, position-sensing devices such as gyroscopes and accelerometers, and possibly a scanning laser range finder. This Rover will use on-board artificial intelligence or neural networks to make "smart" decisions in executing its tasks. It will generate a depth map, transform that map into an elevation map and then match it to a global elevation map sent from Earth. It will improve upon the global path in order to choose a safe route that avoids hazards only it can detect. And it will determine the movements it must make to travel the route, as well as monitor its own journey. Operation of the robotic arm while performing experiments and collecting samples will also be controlled by artificial intelligence.

Depending on planetary positions, it takes between six and 45 minutes for a radio signal to complete a round trip between Earth and Mars, making it impractical for an Earth-bound human to remotely drive the Mars Rover. It would be much more efficient if the Rover could drive itself at least part of the time. Although full autonomy of the Rover isn't yet possible, different degrees of partial autonomy are.

SAM vs. CARD

Two main methods of semiautonomous travel developed at JPL are semiautonomous mobility (SAM) and computer-aided remote driving (CARD). SAM offers a high level of autonomy and is planned as the primary method for controlling the Rover on Mars. CARD, which doesn't allow as much autonomy, enables more interactive control and may be used at times when more human instruction is desired—for performing experiments, for example, or for moving in short, complex patterns.

In the CARD method, the Rover is outfitted with stereo cameras that take pictures of whatever it "sees." When a decision has to be made as to where it will move next, the Rover stops and transmits three-dimensional images of the Mars terrain to Earth. A human operator views the images and designates a path for the Rover to travel. A computer, either on Earth or on the Rover, calculates appropriate turn angles and path segment distances for the Rover to take. The Rover executes movement commands from the computer while autonomously monitoring its own movement. When the Rover has completed this path, it stops, takes more pictures and then repeats the process. For each iteration, the Rover can cover roughly 20 m in 30 minutes, depending on the terrain.

Before the year 2000, we'll be able to gather information directly on Mars through the use of an unmanned roving robot.

■ ■ ■

The SAM method provides for a more autonomous Rover because the Rover is accompanied by a satellite that orbits the surface of Mars (as shown on opposite page). This satellite uses a high-resolution camera to take pictures of the Martian terrain from two different positions in its orbit. Those pictures are then sent to Earth, where they're used to form an elevation map of a large area surrounding the robot. This map can be generated with a resolution of 1 m. Next, a human operator draws an approximate path for the Rover to follow, in order to avoid large obstacles, hazardous areas and dead-end paths. This is drawn either manually on a stereo display of the pictures from the orbiter, or by using the map that was generated from the picture.

A SAM Rover also has stereo cameras, but it uses the view from these cameras to generate a depth map. From this depth map, the Rover generates an elevation map of its local area and finds the closest correlation between it and a portion of a global elevation map sent from Earth. Through various methods of determining its absolute position using sensing elements, the Rover can compare its position in relation to the path it must follow. Then the Rover creates a revised map with very high resolution in its immediate area and calculates a feasible local path based on the approximate global path sent from Earth. Finally, the Rover moves 10 m and repeats the process from its new position, using the map it previously received from Earth. This process is repeated all day and possibly all night with the use of strobe lights.

Less communication between Earth and Mars is needed with SAM than with CARD. Since a SAM Rover averages 10 km/day, it needs to receive only one map with a 10-km path each day. SAM is predicted to run at least four times faster than CARD, which requires continuous remote communication with Earth.

Multiprocessing on the Rover

JPL's current plan is to use different CPUs to run different types of tasks, such as vision, navigation and operation of the mechanical arm. To coordinate the various functions that the Rover must perform, a real-time operating system containing multiprocessing communication tools is essential. Powerful real-time control and data-acquisition capabilities are also very important. Furthermore, in order to

Leslie Kirby, BSCS, Chief Engineer, Wind River Systems

develop such an advanced application, equally sophisticated development tools that let engineers get their programs up and running quickly are also important. The real-time operating system selected for the Mars Rover is VxWorks from Wind River Systems (Emeryville, CA).

There are several possibilities for how the tasks can communicate with each other over the bus. One is to implement a shared memory scheme; another is to use the Transmission Control Protocol/Internet Protocol (TCP/IP). VxWorks runs TCP/IP, which allows easy communication between any two processes on any two processors in a network. The "network" in the Rover is the VME backplane.

TCP/IP allows a process to create a socket through which it can talk to another process. A socket is a virtual circuit offering error-free, bidirectional communications at typically up to 200,000 bytes/s. TCP/IP ensures data delivery between the processes so that a programmer needs to be concerned only with the information moving through the socket, not with the interprocess communication facility itself. TCP/IP isn't limited to the type of media on which it runs. VxWorks' TCP/IP runs over Ethernet and Pronet, in addition to the backplane, offering engineers a high degree of power and flexibility at development and test times, as well as during run time.

During project development, one or several host machines can be connected to one or several VME chassis. In such a network, programs can be developed on a Unix system and downloaded over an Ethernet via TCP/IP to a real-time target. Once the code exists on the target system, debugging can be done symbolically from the real-time side, or at a source-code level from the Unix side by communication through a socket between Unix and VxWorks.

Testing with sophisticated tools

Real-time test programs can be run on a Unix machine that uses TCP/IP to issue commands over Ethernet to VxWorks. An example of this is a menu-driven test program, modified by Gerald Roston, a member of the technical staff at JPL.

The program runs on a Sun workstation and communicates through a socket to the VME chassis containing a CPU running VxWorks and a vision board from Datacube (Peabody, MA). The programmer issues a command from a workstation, such as "grab frame" or "store frame." The command is parsed on the Sun workstation and sent across Ethernet to the real-time target. Vxworks receives the command and issues the appropriate control commands to the vision board. Testing moves quickly with such sophisticated tools and has the advantage that the application being tested can't crash the test system.

TCP/IP is very popular, and many high-level facilities have been built on top of it. Tools such as rlogin and telnet are used for remotely logging in to other systems on a network. Remote Procedure Call interactively calls routines to run on other machines. File Transfer Protocol is used to transfer files across a network, while the Network File System can be used to make files on one machine available to other machines on the network.

Another advantage of using TCP/IP to develop a multiprocessor real-time system is its transparency, or the fact that the engineer doesn't need to know where various tasks physically run. All that's required to open a socket is a network address and a port number for the task. This is important so that an engineer can easily restructure the system hardware without modifying the application software. Extra CPUs can be added just by plugging them into the backplane. When a task is moved over to the new CPU, the programmer only needs to specify a different network address in order to open a socket to a new location.

Network transparency is also very useful during development since it's usually more desirable and efficient to test application modes individually. When the modes appear to be bug-free, they can be merged together one by one until the final application is constructed.

It wasn't long ago that real-time development limited an engineer to using hex addresses, load maps and hex dumps. Object modules were converted to S-records and downloaded over serial lines. Real-time tools have finally evolved to a point that keeps pace with sophisticated technologies such as robotics and artificial intelligence. Now, real-time developers can use state-of-the-art tools to develop state-of-the-art applications.

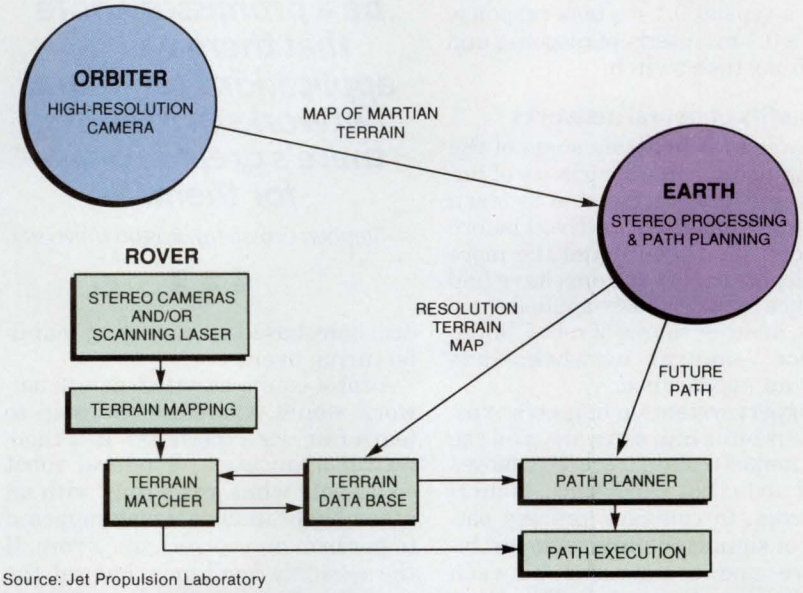

Source: Jet Propulsion Laboratory

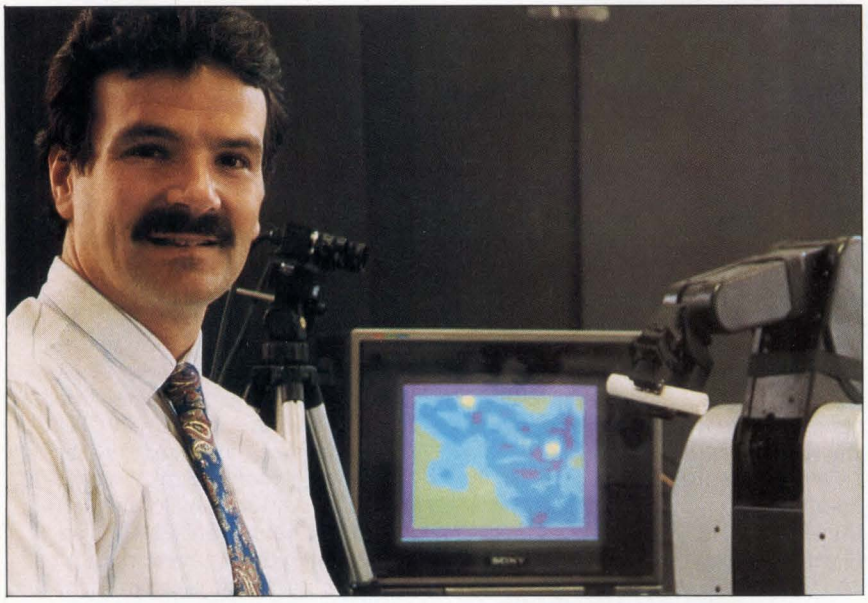

While many companies concentrate on theoretical studies of neural networks, Neurogen has developed a prototype neural-net robot with hand-eye coordination, according to Michael Kuperstein, president. The Neurogen robot learns coordinated behavior from its own experience and achieves adaptive visual-motor coordination of a multijoint arm without a teacher, says Kuperstein.

well as with various RISC processors, including SPARC. A version for the 68030 is in the works.

Based on a real-time, multitasking, multiuser kernel, LynxOS interfaces with a number of the networking schemes commonly found in factory environments, including Ethernet, Transmission Control Protocol/Internet Protocol, Manufacturing Automation Protocol and X.25. When used on a 1-Mips processor, it provides a typical 0.1-ms task response time, a 0.1-ms interrupt response and a 0.05-ms task switch.

■ Benefits of neural networks

If they're ever to attain some of the distinguishing characteristics of humans, robots must be able to learn from their experiences. Even before artificial intelligence and the more practicable expert systems have had a chance to be seriously applied to robotics, another source of robot "intelligence"—neural networks—has made an appearance.

In expert systems, a human's experience is built into software, and the robot makes certain corrective moves based upon that experience. Neural networks, in contrast, process patterns of signals to imitate the architecture and connections between brain cells, or neurons. The robot learns from what is happening rather than from what it has been told to expect. Much like a human brain, neural networks process information by recognizing patterns of signals. This differs distinctly from the action of digital computers, for example, which process information as individual symbols. Neural networks have the potential to recognize changes in environmental conditions and then react to them, or to make decisions based on changing manufacturing events.

"Although there can't be a promissory note that there will be applications for neural networks in robotics, there's great promise for them."

—Stephen Grossberg, Boston University

■ ■ ■

A robot equipped with a neural network would, therefore, be able to learn from its experience. In a theoretical application, a neural robot could note what was wrong with an assembly procedure it was supposed to perform and correct the errors. If the assembly conditions changed, the neural robot would adjust its assembly procedures to accommodate those changes.

"Although there can't be a promissory note that there will be applications for neural networks in robotics, there's great promise for them," says Stephen Grossberg, founder and director of the Center for Adaptive Systems at Boston University (Boston, MA) and one of the early investigators of cognitive and neural networks. "This is especially the case where you want self-adaptive and autonomous capabilities." Many new design ideas and new models of neural networks have been demonstrated and proven in simulation, according to Grossberg. Robots, however, require self-adapting sensing feedback systems and hand-eye coordination.

While other companies are still studying neural-net robots, Neurogen (Brookline, MA) claims that it has already developed a prototype neural robot with hand-eye coordination. Called the Neurogen robot, it has the ability to learn coordinated behavior from its own experience and achieves adaptive visual-motor coordination of a multijoint arm—without a teacher, according to Michael Kuperstein, Neurogen president.

Positioning accuracy of the Neurogen robot in an application where it reaches a cylinder arbitrarily positioned in space is now 3 to 4 percent of the robot arm's length, and its average orientation error is 4 degrees, which are far from the accuracies required for industrial applications. "This is enough to say our neural robot is working. And we're still improving the accuracy," says Kuperstein.

The Neurogen robot consists of a stereo-camera machine-vision system and an industrial robot arm. The robot's base, shoulder and elbow change positions while a gripper holds the object to be moved, which for now is a cylinder. In the learning process, the robot arm with a cylinder in its grip is moved by the robot's controller to various locations. The stereo-camera machine vision system senses where the object rests at each step and generates signals corresponding to the camera angles. During each of several thousands of trial runs, which involve the storing and studying of hundreds of such signals, the controller establishes the relationship between camera angles and arm-joint angles needed to move the arm to the target.

Another company conducting research into neural networks and robotics is Neural Systems (Vancouver,

INTEL'S 80960 ARCHITECTURE IS THE SUPERHIGHWAY FOR EMBEDDED CONTROL.

AND IN 1989 WE'RE RAISING THE SPEED LIMIT.

66* MIPS by 1989. Over 100 MIPS in the 1990's.

At those speeds, Intel's new 80960 architecture is going to be miles ahead of other embedded control solutions. By enabling the execution of multiple instructions per clock, the 80960 architecture goes beyond simple RISC. Thus raising both the speed limit and your expectations.

But you don't have to wait to get on this virtual superhighway for embedded control. Because Intel offers a family of processors that you can drive away with today. All based on the 80960 architecture.

Starting with the 80960KB. The only 32-bit embedded control processor with a complete Floating-Point coprocessor. All on a single chip.

Then there's the 80960KA. A single chip 32-bit embedded control processor with a direct upgrade path to 66 MIPS in 1989. Available in more economical versions, without Floating-Point, at 16-, 20- and 25 MHz.

And for military applications, Intel offers the 80960MC. The only 32-bit embedded control processor with hardware multi-processing capabilities on a single chip.

We also supply all the development support you need. Including high-level languages for a variety of industry-standard hosts, and the ICE™ 80960 development tool.

The result is the world's first fully supported 32-bit architecture designed exclusively for embedded control.

So why wait any longer? Our comprehensive 80960 brochure gives you the complete story. To receive it, just call (800) 548-4725 today and ask for Lit. Dept. #W475. With one phone call, your project can be on the road to 32-bit embedded control.

And ready to accelerate.

intel®

© 1988 Intel Corporation

*Sustainable MIPS on-chip.

CIRCLE NO. 19

Smart sensors play a crucial role in future robots

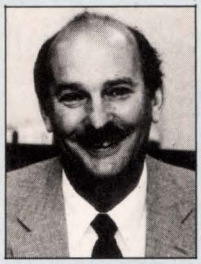

In theory, the ultimate robot would be a machine capable of emulating all human functions. In practice, however, most of today's robots are incapable of performing more than a basic few of those functions. For a robot to emulate even one of the five human senses, a combination of several functions may be required.

It's a relatively simple matter for a human being assembling a piece of machinery, for example, to distinguish one bolt from another. It's equally simple to pick that bolt up and move it to another location. A robot, in contrast, has to be equipped with several different types of sensors to accomplish the same simple feats. The robot must see the part, differentiate it from the others, grasp it without damaging it, identify the location where it's to be placed, place it there, and then release it. At the very least, a multitude of sensors—from vision-based sensors to limit switches to pressure sensors—are required in any robotic application.

Silicon-based sensors, such as photosensors or photoconductors, are used as the "eyes" of a robot and, when combined with sophisticated circuitry and image processing, are able to recognize patterns. Magnetic-responsive sensors are widely used in presence/absence and positioning types of applications, both to indicate position and as stops to limit motion. Pressure sensors let the grasping arms or fingers apply the right amount of tactile pressure—enough to hold firmly without crushing. Temperature sensors, silicon-based or not, aren't widely used in robotic applications today, but could be applied in any application where heat is generated and must be detected and controlled, such as in wave-soldering operations.

Designers are finding polyvinylidene fluoride, a form of plastic film, to be useful in robotic applications because of its relatively large piezoelectric coefficient. When force is exerted on the film, it generates its own voltage, which is proportional to the amount of force applied. Conformal coatings of this material can be wrapped around the gripping arm of the robot. After the voltage reaches a certain level, the grip can be mechanically relaxed.

Gas sensors are another form of sensor technology that can be used in robotics. A robot could be utilized, for example, in environments that are potentially hazardous to humans. The robot could contain a variety of environmental gas sensors to provide early warning against human intrusion into a contaminated environment. Available types of gas sensors include optical gas sensors, gas sensors that sense color changes, Nernst cells and electrochemical (electrolyte) cells.

A brand-new, promising sensor technology is biotechnology. Biologically based sensors use the characteristics of various enzymes and other organic substances and their reactions with other substances. Biosensors only detect the presence of specific materials. The enzyme glucose oxidase, for example, will only cause a chemical reaction to occur if glucose is present, and it isn't sensitive to any other blood sugars. In the future, robots may be able to use biosensors in simple labor-analysis applications.

The robots of tomorrow, with their ever-increasing levels of sophistication, may be even more dependent on sensors than the relatively simple robots of today.

■ ■ ■

■ Future robotic sensors

Smart sensors, or those that include logic, addressability, signal conditioning and multiplexing capability, will play an important role in future robots because they keep much of the decision making at the lowest possible level—the sensor level. As silicon IC technology becomes more sophisticated, it's becoming possible to merge all of the functions of a control system into smaller packages where the boundaries between these functions are no longer easily discernible or even important.

The advent of "smart power"—the most recent silicon IC technology advancement—lets logic circuitry and power-switching circuitry exist compatibly on the same piece of silicon. When combined with a silicon-based sensor, this can result in a sensing subsystem on a chip, an obvious advantage from a user's point of view because it's more cost-effective and provides higher reliability.

By incorporating smart power in a sensing subsystem, a number of other features can also easily be included on the same IC, such as over-voltage protection and surge protection, as well as certain types of diagnostic capabilities, such as sensor continuity. Since an actuation element is already included in the IC, another one isn't required, thus reducing the size. The increased functionality of the sensor subsystem lets the CPU or other central decision-making units concentrate more intently overall system parameters. The speed of the control system is also increased because of the inherently fast operation of solid-state devices.

Artificial-intelligence sensing networks, which combine sensors with sophisticated software, will most likely also be used in the robotic applications of the future. These networks will be able to sense intangibles such as taste or texture by combining the inputs from a variety of tangible parameters and then correlating the relationship between those variables that constitute an acceptable product. These artificial-intelligence sensing networks will be useful for testing food products, for example.

The robots of tomorrow, with their ever-increasing levels of sophistication, may be even more dependent on sensors than the relatively simple robots of today. They'll require smart sensors, sensors incorporating smart power for actuation and possibly even artificial-intelligence sensing networks in order to be most effective. The only question for the future is how quickly new types of sensors can be developed to meet the rapidly expanding requirements of the industry.

Paul Chapman, Director of Technology, Micro Switch, a division of Honeywell

It knows the Big Bang as well as the Big Buck Theory.

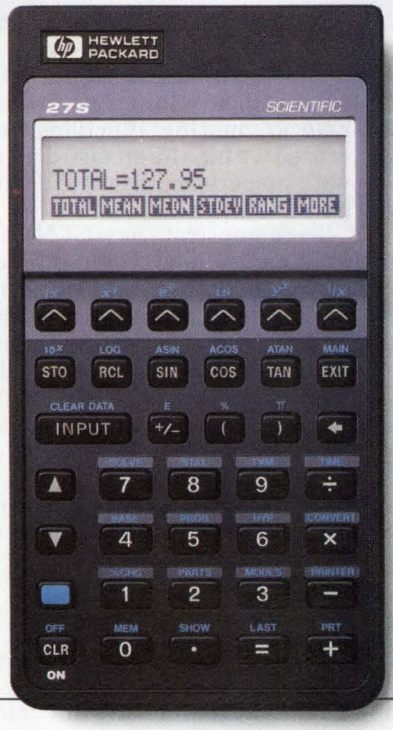

The HP-27S Scientific Calculator.

Has both scientific and business functions. And the unique HP Solve lets you enter your custom problems and solve for any variable without rearranging the equation. Algebraic Entry. Suggested retail price $110.

The HP-27S is the only scientific calculator available today with built-in business functions.

Featuring a wide variety of scientific functions that allow you to solve your technical problems, the HP-27S also offers built-in business functions like the time value of money and amortization. And it has a unique alarm feature that reminds you of your daily appointments. The HP-27S also has an infrared output that eliminates the need for a cable hookup to HP's optional printer.

Advanced, Yet Simple.

So you can concentrate on the problem not the manual, the HP-27S offers easy screen guides (alpha menus, prompts, labels and messages) that lead you step-by-step to your solution.

The HP-27S Scientific Calculator. Another member of HP's proud new family of quality calculators. For the name of the HP dealer nearest you, call 1-800-752-0900, Ext. 786W.

HEWLETT PACKARD

CIRCLE NO. 20

©1988 Hewlett-Packard Company PG 12810A

To test its proprietary neural network, known as the Graded Learning Network, Hecht-Nielsen Neurocomputers trained the system to balance a broom on end on a cart that's driven along a track. Using information about one point on the broom that it received from a video camera every 130 ms, the system learned to compensate for physical shocks, including those caused by a person tapping on the broom during the sequence.

Hecht-Nielsen Neurocomputers

BC, Canada). The company's strategy concerns inverse kinematics transformation procedures in which X, Y, Z coordinates are converted to robot-joint angles. Plans are in the works to integrate vision into the testing.

Neural Systems is now working with robotic companies to apply neural networks to real industrial robot arms, according to vice-president Dale White. None of these systems have yet reached production status, but White claims that repeatable positioning accuracies on the order of 2 thousandths have been attained.

Hecht-Nielsen Neurocomputers (San Diego, CA) has developed a proprietary neural network called the Graded Learning Network (GLN) that's relevant to robotic applications, although it hasn't yet been used for such. GLN is a real-time, closed-loop adaptive control system that learns how to perform a control operation by trial and error. The system learns on its own, without programming, initial data collection or extensive systems analysis.

The company's Anza-Plus single-board Neurocomputing Coprocessor System, which is available with either 2 or 10 Mbytes of on-board memory, transforms an IBM PC AT or a 386-compatible into a real-time neurocomputer. These boards are supplied with Release 2.0 of Hecht-Nielsen's Neurosoft software, which lets users treat neural networks as subroutines within C programs. All neurocomputing functions are executed within the on-board memory, with no performance overhead relegated to the host PC.

To demonstrate how GLN could be used with a neurocomputing coprocessor board in robot applications in harsh environments, the company has trained a system to balance a broom on end on a cart that's driven along a 30-cm track by a motor. No knowledge of falling brooms was built into the neural network.

In this training exercise, a video camera measures the position of a point near the top of the broom and passes the information to the network every 130 ms. The network also calculates a delayed speed and a delayed acceleration derived from the position data from the camera. These delays let the network function without receiving any information on positioning, velocity or acceleration from the stepper motor that moves the base of the broom.

During training, the neural-network system tries to balance the broom for a 15-s trial period. At the end of each trial period, the network receives a performance grade that it uses to improve its performance at the next trial. The network learns to bring the broom to the center of the track and balance it there with virtually no angular motion.

Robotics not a lost cause

In the past few years, several U.S. companies that had previously attempted to manufacture robotic systems and sell them to end users rather than to OEMs found that not only was this unprofitable, it was a fast road to Chapter 11. The U.S. manufacturers couldn't compete with foreign competition, particularly since most of them bought the robots for their systems from foreign companies. Several U.S. companies diversified their product lines by developing and offering peripheral systems to an OEM customer base.

At present, U.S. robotic technology has an advantage, particularly involving the use of neural networks for hand-eye coordination of robotic arms. But that advantage could be short-lived. Boston University's Grossberg reports, for example, that Japanese scientists were notably represented at recent seminars on neural networks. For now, the handful of U.S. and Canadian companies that are studying and developing neural networks, as well as real-time software, image-processing subsystems, intelligent sensors and embedded controllers for robotics, seem to be in advantageous positions. □

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

Automatix
(508) 667-7900, ext 2038 ..... **Circle 273**
Datacube
(508) 535-6644, ext 222 ..... **Circle 274**
Digital Equipment Corp
(508) 490-7737 ..... **Circle 275**
Hecht-Nielsen Neurocomputers
(619) 546-8877, ext 604 ..... **Circle 276**
Imaging Technology
(617) 938-8444, ext 309 ..... **Circle 277**
Intel
(503) 681-8080 ..... **Circle 278**
Jet Propulsion Laboratory
(818) 354-2755 ..... **Circle 279**
Lynx Real-Time Systems
(408) 370-2233 ..... **Circle 280**
Micro Switch
(815) 235-5731 ..... **Circle 281**
Nestor
(401) 331-9640, ext 44 ..... **Circle 282**
Neural Systems
(604) 263-3667 ..... **Circle 283**
Neurogen
(617) 232-8266 ..... **Circle 284**
Wind River Systems
(415) 428-2623 ..... **Circle 285**

Please rate the value of this article to you by circling the appropriate number in the "Editorial Score Box" on the Inquiry Card.

High 503 *Average 504* *Low 505*

GCR Tape Drive

+ "Reliability" = Maintenance Free

- High Reliability 9 CMOS-LSI
- ½" IBM Tape Format
- Auto-Loading
- Convenient LED Display
- Simple Mechanism
- Tri-density (6250, 1600, 3200)
- Tape Speed 100IPS, 50IPS (3200 BPI)
- Transfer Rate 20KB - 1MB/SEC
- Interface Pertec/SCSI
- Capacity 175MB/2400 feet
- Buffer Memory 512KB
- Rack Mounting/Desk Top

ANRITSU CORPORATION
10-27 Minamiazabu 5-chome, Minato-ku, Tokyo 106, JAPAN
Phone: (03) 446-1111/Telex: 0-242-2353 ANRITU J/Cable: ANRITDENKI TOKYO

ANRITSU AMERICA, INC.
15 Thornton Road, Oakland, NJ 07436, U.S.A.
Phone: (201) 337-1111/Telex: 642-141 ANRITSU OKLD

CIRCLE NO. 21

The meeting that should never have happened.

When a disk drive crashes, you're in trouble. To an irate customer, it doesn't matter that you purchased the drive from someone else. *Your* logo is on the system.

This means you need a disk drive supplier who is also a *partner*... a company who understands what it's like to compete in today's tough computer systems marketplaces. That company is Hitachi.

An intimate knowledge of the disk drive business.

When you decide to do business with Hitachi, you get more than just a high quality, incredibly reliable disk drive. You also get the support, resources, and expertise of a $40 billion dollar company who's one of the largest OEM computer systems manufacturers in the world. We use disk drives in our own systems, and risk losing customers if drives fail. Just like you.

This means we know what it takes to make the best. We design and build our own motors, heads, microprocessors, and custom LSI to our own stringent specs—all checked by the most stringent Quality Assurance program in the industry.

We have over 1,500 engineers who work on nothing but disk drives. They're supported by nearly two decades of experience manufacturing disk drives, and over 16,000 employees in R & D throughout the company.

A serious commitment to the disk drive business.

Hitachi continues to advance storage technology, through the ongoing development and manufacturing of new 5.25", 8", and 8.8"

Winchesters, magnetic tape sub-systems, and ultra-high density optical disk sub-systems. We plan on being in this business for a *long* time... making sure your customers call with orders, not complaints. Next time, specify Hitachi.

Fast Action: To get information or product literature immediately, CALL TOLL FREE 1-800-538-8157, Ext. 877. In California, 1-800-672-3470, Ext. 877. Ask for literature number APB-002.

Regional Sales Offices: Waltham, MA 617/890-0804, Dallas, TX 214/991-7983, Brea, CA 714/993-1610, San Bruno, CA 415/872-1902.

Hitachi offers a complete line of Winchester, magnetic tape, and optical storage products with speeds, capacities and interfaces for virtually any application.

Hitachi America, Ltd.
Computer Division
950 Elm Avenue, Suite 100
San Bruno, CA 94066

HITACHI®

CIRCLE NO. 22

A lot of companies claim complete VME capability.

Only Radstone can deliver it now.

Radstone Technology is the *only* company in the world that delivers proven VMEbus board and system level products for every application from commercial to full Military Specification...and everything in between. Plus more than 30 years of solid computer experience—more than *anyone* in the OEM board level computer market.

Commercial VME...highest performance, low cost, ready to plug and play.

Radstone's commercial VMEbus product line includes:

- 16 and 32 bit processor boards, including a complete family of 68030-based boards
- Memory boards covering all available capacities and technologies, CMOS versions and VME/VSB models
- The world's fastest and most popular SCSI boards; parallel, serial and analog I/O boards; and much more
- Development chassis with peripheral hardware and a wide range of advanced operating systems

Military VME...true Mil-Spec VME, not "militarized."

All Radstone Mil-Spec VMEbus boards fully comply with both VME and Military Specifications. These boards are built with MIL-STD-883C class B components and meet MIL-E-5400, MIL-E-4158 and MIL-E-16400. They feature low power CMOS components and conduction cooling via an on-board thermal management layer.

- 68020-based processor boards
- SCSI and 1553B interface boards
- MPCC, Ethernet and serial and intelligent communications boards
- Static and dynamic memory boards
- ATR boxes accommodating 15, 8 or 5 boards
- Complete software support, including Ada
- BITE Built-In Test Equipment

Ruggedized VME for severe environment, non-Mil-Spec, applications.

All Radstone Mil-Spec VME modules are also offered in electrically and mechanically compatible reduced environmental spec versions to give you low cost hardware for severe—but not full military—applications. Even lower cost versions are available for off-the-shelf development work.

Radstone Technology...Ultimate VME capability for you.

Radstone is the only company that produces all its boards in military qualified production facilities. And we back up our market-matched commercial, ruggedized and Mil-Spec VME products with technical support services second to none in the world.

A lot of companies claim complete VME capability. Only Radstone can deliver it now.

Radstone Technology Corporation
One Blue Hill Plaza, Pearl River, NY 10965-8541

Call Toll-Free: (800) 368-2738
Eastern Region: (914) 735-4661
Central Region: (817) 261-9988
Western Region: (408) 727-4795

RADSTONE TECHNOLOGY

CIRCLE NO. 23

CMOS RAM brings industrial-strength data storage

Rugged, fast and very cost effective, battery-backed memory cartridges are likely to become the storage medium of choice among designers of military, aerospace, industrial and mobile computer systems.

Gavin McLintock, M.Eng.

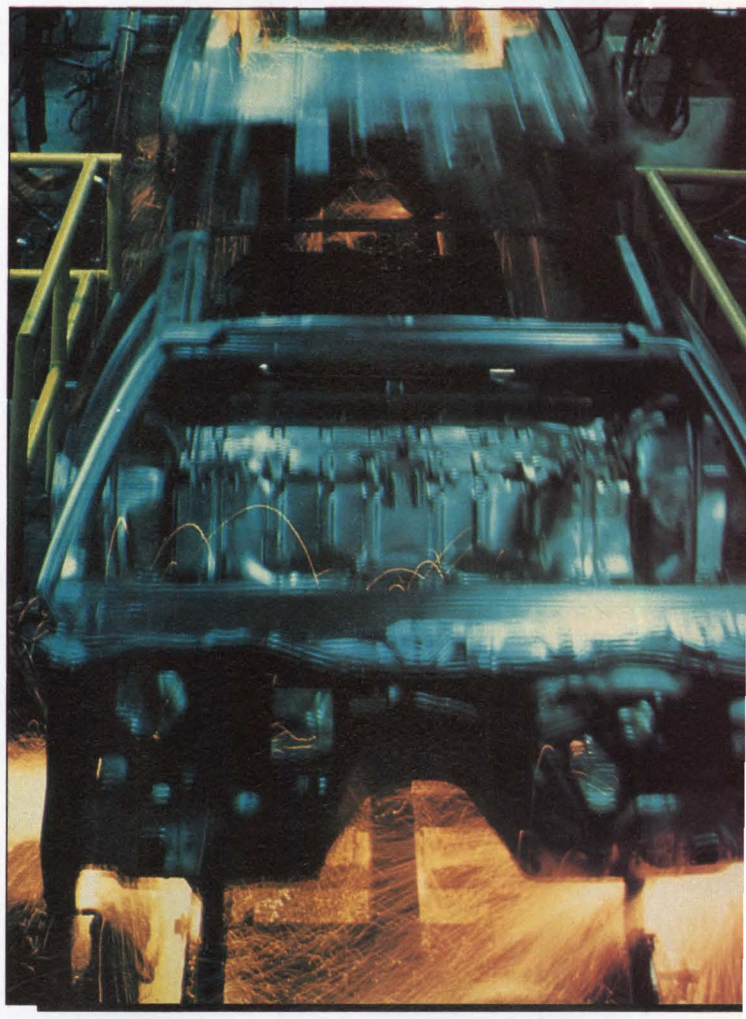

In factory-floor environments, ruggedized and reliable nonvolatile CMOS RAM products are used for loading and storing the graphics required for user-friendly interfaces, as well as for loading operating programs for industrial computers.

Those who develop computer applications for the military or industrial markets face so many more constraints than other designers that they must thrive on challenge, or on frustration. Those designing for harsh operational environments, for example, have been faced with the choice between equally unsatisfactory storage alternatives: the inexpensive but unreliable disk, or the costly but robust bubble memory. In the last year and a half, however, a third alternative, the CMOS RAM cartridge, has begun to provide the robust and inexpensive data storage that many computers require.

For any application, the selection of a storage peripheral revolves around four primary considerations: capacity, speed (access and transfer), cost and reliability in a given environment. The designer of office systems, for example, is concerned with only capacity, speed and cost, in that order of priority. The designer for the industrial environment, on the other hand, has different priorities: reliability is first, followed by cost and then speed.

The least important selection criterion for industrial storage is capacity. Factory-floor computers typically handle a limited number of functions such as controlling one machine or monitoring one process, and their storage requirements are small. In fact, very few plant-floor applications need more than 4 Mbytes of storage, and many are satisfied with less than 1 Mbyte of peripheral storage, according to a survey by Targa Electronics Systems.

In terms of reliability, magnetic storage has been a poor performer in industrial and many portable or mobile computer applications. Winchester disk, floppy disk and tape drives—with all their moving parts—don't oper-

Gavin McLintock is president of Targa Electronics Systems (Ottawa, Canada).

ate reliably in environments marked by extremes of vibration, humidity, mechanical shock, temperature variations, dust and chemicals. Bubble memory, on the other hand, is extremely reliable, but it's also very slow and expensive. Its entry-level cost is four times that of a Winchester disk, for example, and its per-bit cost is 80 times greater (see "A comparison of available media," p 75).

Most of the leading bubble memory producers have expended a great deal of effort attempting to develop faster, more moderately priced bubble memories. To date, though, no real breakthroughs have occurred, and bubble memory remains prohibitively expensive for most nonmilitary applications. Bubble memory's durability, nonvolatility and low data error rates, however, have created a niche market for rugged storage devices that can stand up to the harshest operating environments.

Although Memtech Technology (Santa Clara, CA), Targa, Magnesys (Santa Clara, CA) and Bubbl-Tec (Dublin, CA), to name a few, produce a variety of bubble memory boards and drives, some of these companies are avidly exploring CMOS alternatives, including CMOS cartridges that are interchangeable with bubble memory cartridges.

CMOS RAMs have been on the market for several years, characteristically delivering densely packed circuitry, low power consumption, and gate speeds that approach those of Schottky TTL devices. The development of peripheral storage subsystems using CMOS RAM, however, occurred only recently—partly as a result of technical advances in memory density, surface-mount technology and battery technology, as well as dramatic reductions in memory cost; and partly because hopes for improved bubble memory storage have dimmed.

Battery-backed CMOS RAM data cartridges are being shipped by Epson America (Torrance, CA), Mitsubishi Electronics (Torrance, CA), Dupont (Wilmington, DE), Targa and others (see "Two approaches to RAM cartridge design," p 75).

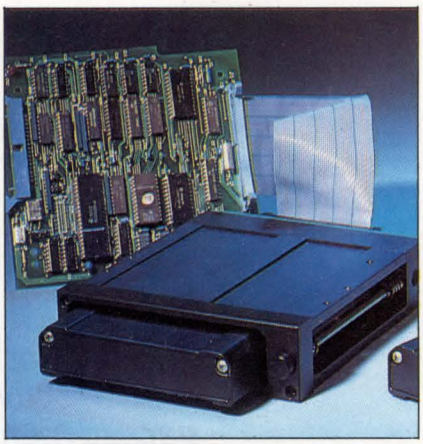

Ruggedized nonvolatile CMOS RAM data storage systems such as these from Targa Electronics Systems meet the three important criteria for use in harsh environments: reliability, cost and speed.

■ Rugged and reliable

Reliability, the primary selection criterion for industrial storage, is most commonly expressed in terms of mean time between failures (MTBF). Ruggedness is more difficult to express, but it generally refers to the environmental conditions under which MTBF figures are obtained. MTBF figures don't adequately indicate the performance of a storage system under environmental stresses such as shock or vibration, so the astute purchaser will always inquire how the figures were obtained.

The MTBF figures for Winchester disk drives—usually between 10,000 and 20,000 power-on hours—are most likely valid only under "typical" operating conditions: a 25 °C operating temperature and a 15 to 20 per-

Why bubble memory is slow

Many companies have invested great ingenuity and capital without finding a feasible solution to bubble memory's slow read rate. This technology stores data as cylindrical magnetic domains, known as bubbles, in a thin film of magnetic material. The presence of a bubble is interpreted as a binary 1; the absence of a bubble (a space) as a binary 0.

In a bubble memory, a rotating magnetic field generated by integral coils propels the bubbles through the magnetic film, circulating them past a pick-up point where data is read to the outside world. Note that the generation, destruction, movement and detection of bubble data takes place without mechanical motion. If power fails, the permanent magnets that shield the device from outside field sources are sufficient to maintain the bubbles, although all movement stops until power is restored.

The main factor making bubble memory a slow storage medium is the inherent necessity of serially manipulating the bubbles from a single generation point through a loop to the read point. Although the 256-bit loops used in most bubble memory designs work in parallel to speed the process, the bubble detector ultimately reads the data one bit at a time.

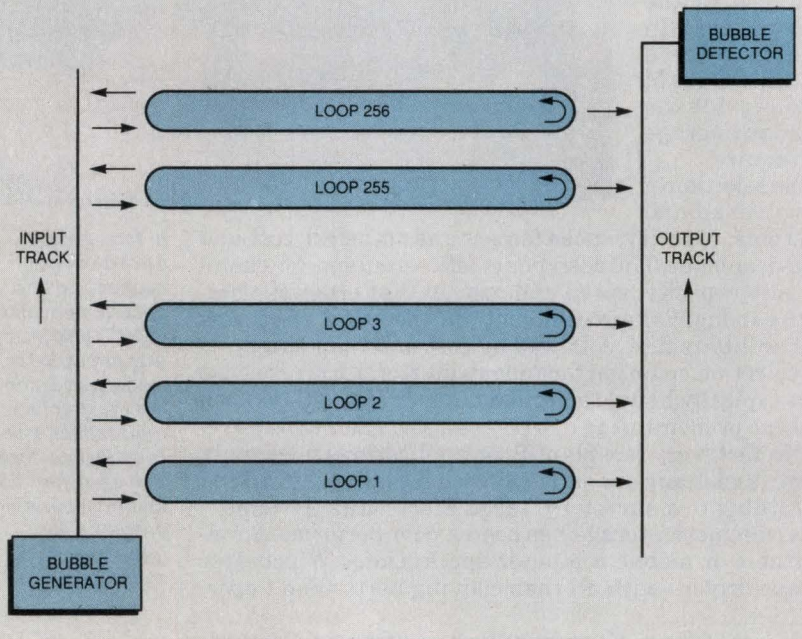

cent duty cycle. The MTBF figures for CMOS RAM cartridges and bubble memory cartridges—usually in the 100,000 power-on hours range—are obtained under harsher conditions: at maximum specified ambient operating temperature, for example, and a 100 percent duty cycle.

To comprehend the cardinal importance of reliability in the industrial environment, one has only to look at the estimated costs of plant downtime, which can run into several thousand dollars per minute. Yet every industrial process seems to produce a hazard for data storage devices developed for office use: tobacco dust in a cigarette factory can gum moving parts even in well-sealed units, airborne fluid particles at a chemical plant can destroy disks and corrode tapes, and so forth.

In a manufacturing environment, the cost of a storage peripheral—the second major selection criterion—is as much associated with the downtime occasioned by maintenance and repair as with capital cost. Since most computer service calls are due to disk drive problems, the cost advantage of RAM cartridges over magnetic storage derives largely from the elimination of downtime rather than from lower capital costs.

RAM cartridge proves fastest

The third priority for industrial storage is speed, and the RAM cartridge drive is typically faster than any of the alternatives. While the continuous transfer rate from a RAM cartridge is comparable to that of a typical Winchester drive, its typical access time is in the range of 350 ns—100,000 times faster than a Winchester. This advantage makes RAM cartridges a good candidate for military field test equipment, where the top three priorities, in order, are speed, reliability and cost.

The CMOS RAM cartridge, still in its early days, will doubtlessly expand its niche appeal as designers become familiar with its advantages. Point-of-sale terminals and portable computers, terminals and instruments, for example, would all benefit from this new type of storage peripheral—as a reliable alternative to a floppy disk drive, as a cost-effective alternative to a bubble memory, or as a practical storage medium for applications where none was practical before. □

Two approaches to RAM cartridge design

There are two different approaches to CMOS RAM cartridge design today. Epson America (Torrance, CA), Dupont (Wilmington, DE), ITT Cannon (Fountain Valley, CA) and Mitsubishi Electronics (Torrance, CA) have chosen to develop a cartridge about the size of a credit card with a maximum 1-Mbyte capacity. These cartridges are supplied with a simple connector, but the designer must develop his own hardware and software interfaces.

On a cost-per-unit basis, this approach offers the most inexpensive CMOS RAM cartridge, but savings are offset by the cost of developing the necessary interfaces. Depending on design requirements, these cartridges can be treated as peripheral devices or as part of the system memory, but in either case, appropriate packaging, electronics and software must be developed. Similar cartridges are also offered with EEPROM or ROM.

The other approach, which was taken by Targa Electronics Systems, has been to develop a cartridge that stores up to 4 Mbytes of data and is approximately the size of a video cassette. These larger cartridges also include battery backup to sustain memory contents in a nonvolatile state for more than five years, as well as a selection of electrical interfaces and software protocols for various bus structures. This type of cartridge is, of course, more expensive, but it reduces customer development time and design cost by providing pre-engineered interface and packaging options.

Most of these larger RAM cartridges can be used with the same versatility and portability as floppy disks; they can be write-protected, for example, and formatted in a manner similar to a disk. Some of the cartridges have interfaces that emulate a floppy disk drive interface, so they can be used with existing floppy controllers. Also, a number of the cartridge drives intended for internal mounting fit the common half-height disk drive form factor, further simplifying field and factory upgrades from floppy to RAM.

A COMPARISON OF AVAILABLE MEDIA

	Bubble Cartridge	CMOS RAM Cartridge	Floppy Disk	Hard Disk
Capacity	1 Mbyte	4 Mbytes	720 kbytes	40 Mbytes
Access Time	>10 ms	>100 ns	>100 ms	>10 ms
Transfer Rate	>10 kbits/s	>10 Mbits/s	>100 kbits/s	>10 Mbits/s
Entry Cost	$1,000	$250	$100	$250
System Cost (millicents/bit)	24	12	34	0.3
Operating Temperature	0 to 55°C	0 to 65°C	10 to 45°C	10 to 50°C
Operating Vibration	>10 g's	>10 g's	<1 g	<1 g

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

Bubbl-Tec
(415) 829-8705 ..... **Circle 266**
Dupont
(302) 774-3357 ..... **Circle 267**
Epson America
(213) 534-4500, ext 1227 ..... **Circle 268**
Magnesys
(408) 988-1881, ext 13 ..... **Circle 269**
Memtech Technology
(408) 970-8900, ext 312 ..... **Circle 270**
Mitsubishi Electronics
(408) 730-5900, ext 2201 ..... **Circle 271**
Targa Electronics Systems
(613) 731-9941 ..... **Circle 272**

Please rate the value of this article to you by circling the appropriate number in the "Editorial Score Box" on the Inquiry Card.

High 506 *Average 507* *Low 508*

ProDrive. Because

you can't afford to hurry up and wait.

HOW THE WORLD'S FASTEST 3½-INCH HARD DISK DRIVE CAN GIVE YOU THE COMPETITIVE EDGE.

It's a simple fact of systems life: your system is only as fast as it's slowest component.

So if your disk drive is slow, your system is slow.

And in today's competitive marketplace, every millisecond counts. No matter how fast your microprocessor, your customers just won't accept a system that bogs down in disk access. And neither should you.

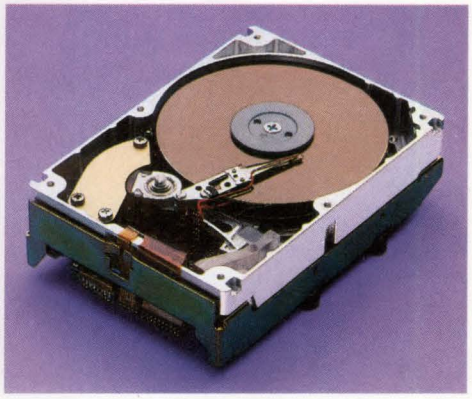

That's why you should know about the ProDrive Series™ of high-performance disk drives from Quantum.

ProDrive™ offers you the fastest seek time of any 3½-inch hard disk drive available today: 19 ms in the 42- and 84-megabyte models.

And with Quantum's exclusive DisCache® data buffer, you can cut that 19 ms to 12 ms or less, depending on your application.

But speed is only half the ProDrive advantage. Because speed is useless if your drive can't deliver it when you need it. That's why we built ProDrive with the highest MTBF on the market. 50,000 hours.

No one offers you higher performance with higher reliability.

	PRODRIVE 40/80	BEST OF THE COMPETITION*
Seek Time (ms)	19**	25
Buffer Size (bytes)	64K	16K
Buffer Management	DisCache®	None
Transfer Rate (MB/sec)		
AT-Bus	4	4
SCSI Sync	4	—
SCSI Async	2	2
MTBF (hours)	50,000	30,000

These are the best individual specifications of competitive drives. (This is not a real product.)

**Average seek time with DisCache is 12 ms in typical applications. (All specifications subject to change.)*

And no one offers it to you in a broader range of products. In 42- and 84-megabyte formatted capacities with SCSI and AT-Bus interfaces available now. And in capacities up to 168 megabytes with SCSI, AT-Bus, and ESDI interfaces available soon. And all at competitive prices.

But find out for yourself how the ProDrive Series can give you the competitive edge in your market.

Call Quantum today.

A few minutes on the phone with us will save you a few milliseconds on your system performance specs. And that could earn you a lot of dollars on your bottom line.

So we suggest you hurry up. Because your customers won't wait. And you can't afford to.

Quantum

Quantum Corporation, 1804 McCarthy Blvd., Milpitas, CA 95035 Western Region (408) 980-8555 Eastern Region (603) 893-2672
Europe (49) 69-666-6167 Authorized Distributors: Arrow (800) 777-ARROW Marshall (800) 522-0084

ProDrive and ProDrive Series are trademarks, and DisCache is a registered trademark of Quantum Corporation. © 1988 Quantum Corporation.

CIRCLE NO. 24

NOW SERVING 68030 EMULATION

MicroCASE supports the 68030 with 25 MHz emulation.

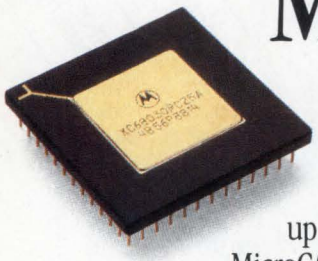

Step right up to the MicroCASE 68030 PROBE™ In-Circuit Emulator. The first, the fastest, the best 68030 support of its kind.

The 68030 PROBE incorporates the same unique features as our highly successful 68020 PROBE™ —the debugging tool chosen by hardware and software engineers in a wide variety of leading high technology equipment manufacturers world-wide.

Pre-fetch pipeline dequeueing

The 68030 PROBE provides in-circuit emulation— at speeds up to 25 MHz. PROBE also features pre-fetch pipeline dequeueing, so it's easy to figure out which instructions actually execute, and which bus cycles relate to those instructions.

And by merely changing the probe tip, the 68030 PROBE also supports the 68020.

True source-level debugging

The 68030 PROBE provides on-line debugging of high-level language software. Real source statements are displayed, eliminating lengthy translations from assembly language to your high-level language.

PROBE utilizes an IBM® PC AT as its instrumentation chassis, so you can get compiled code to its target via Ethernet, VAXNet, SUNNet, SCSI or RS-232. Whether you compile on a PC, a workstation or VAX,® MicroCASE supports more object code formats than any other vendor.

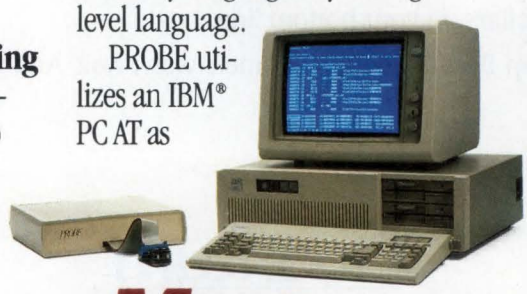

No need to wait any longer. Real-time emulation and dequeueing for the 68030 and 68020 are available now. From the number one supplier of hardware-assisted software debuggers for the 68000 series. The Atron Division of MicroCASE.

If you'd like more information, or a short product demonstration, call us today at **408-253-5933.** Or circle the number on the reader service card.

68020 PROBE and 68030 PROBE are trademarks of the Atron Division of MicroCASE. VAX is a registered trademark of Digital Equipment Corporation. IBM is a registered trademark of International Business Machines.

MicroCASE
ATRON DIVISION

Saratoga Office Center
12950 Saratoga Avenue
Saratoga, CA 95070
408-253-5933

CIRCLE NO. 25 FOR LITERATURE
CIRCLE NO. 26 FOR DEMONSTRATION

Electronic memories come in tough little packages

IC memory cards offer a fast, reliable, compact memory alternative that will open up new applications.

W. Judd Clark, BSME
Ben Sakai, MSEE

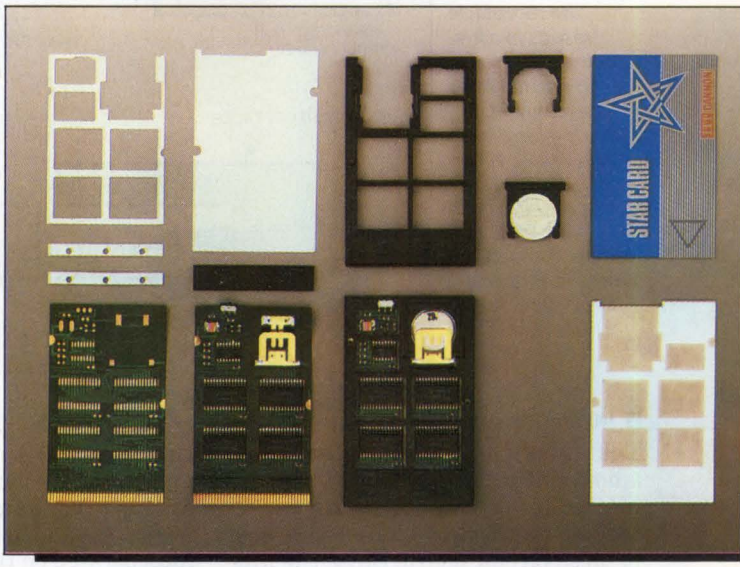

The ITT Cannon Star Card is an IC memory card that can be used as an internal memory device or as an external add-on memory. The card's assembly starts with the bare polyamide flat-flex circuit. Next, the components are installed using a surface-mount technique, and then the PVC frame is added.

Until recently, portable memory for data processing devices was limited to floppy disks, hard disk cartridges, cassette tape and bubble memory. These devices suffer from slow read-write speeds, and all except bubbles are sensitive to harsh environments, enduring the maintenance problems associated with electromechanical drive mechanisms. Advanced IC fabrication and packaging techniques, however, have made possible a totally electronic memory alternative that provides high reliability, reasonable capacity and high-speed access in an envelope the overall size of a credit card and about 3 mm thick. This superportable product, the IC memory card, has begun to carve an important niche in the portable memory market.

Because the technology of the IC memory card is entirely electronic, access times are in the 100- to 200-ns range. Faster access times can be attained by using faster read-write chips, but this alternative has so far had limited appeal because of the associated higher costs. Even at 200 ns, though, access time is far better than that of the most common portable medium, floppy disk, whose electromechanical characteristics keep access times above 120 ms.

IC memory card capacities range from as little as 8 kbytes to as much as 1 Mbyte, with larger capacities expected in 1989. As denser interconnect schemes for the cards become available, though, capacities beyond 1 Mbyte will appear. Once this occurs, further capacity improvements will depend on advances in chip technology and on fabrication and packaging developments. Also, expander cards will become available to drive and control groups of IC memory cards, allowing the configuration of large mass storage and retrieval systems to become feasible.

Available IC memory cards use static RAM, mask ROM, PROM, EPROM and, to a limited extent, EEPROM. Dynamic RAM isn't used in the cards because of its high current demands and the refresh circuitry that would have to be built in to maintain memory contents. The recent development of a new technology for memory cards—ferroelectric RAMs—will soon let cards be designed with the read-write speed of SRAM and the nonvolatility of ROM.

W. Judd Clark is director of new product development and Ben Sakai is product engineering manager for ITT Cannon/Components Division (Fountain Valley, CA).

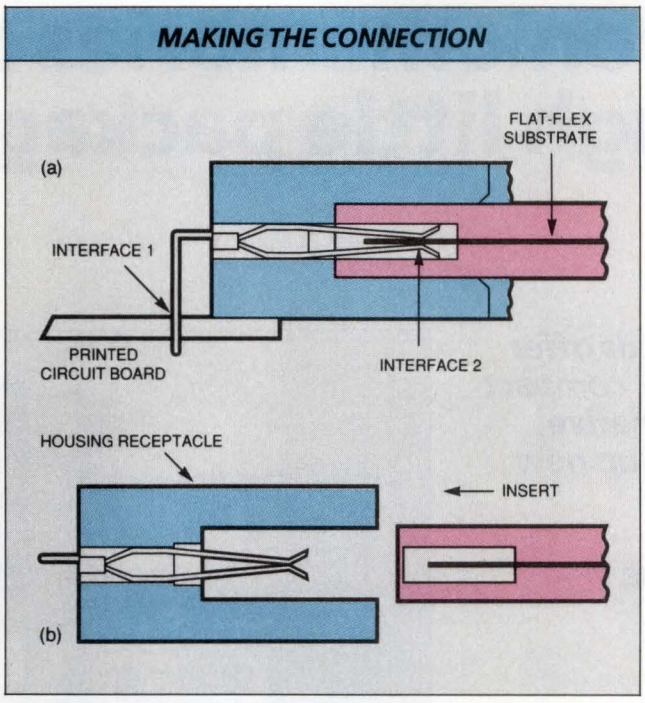

The IC memory card, mating to a receptacle connector on the host printed circuit board (a), should withstand 50,000 insertions and withdrawals without degradation. A double cantilever-type female connector in the receptacle results in a low insertion force, high normal force interconnect, enhancing long-term durability. When the card isn't inserted (b), both the male and female contacts are recessed for protection during handling.

Interconnect system vital

The IC memory card interfaces with a host computer by plugging into an interconnect receptacle that's attached to the I/O bus on the host printed circuit board. Since the ultimate success of the cards depends not only on reliable data retention but also on reliable mechanical performance over the life of the host, the interconnect system is a vital component. It should offer the user low insertion force combined with a high-contact normal force to ensure low resistivity and high durability. In addition, the physical design of the interconnect system must be such that both male and female contacts are well protected from electrostatic discharge and finger contact.

In the interest of reliability and low resistance, the number of contact interfaces between the card and the host printed circuit board should be held to a minimum. One way of doing this is to use a flat-flex circuit substrate in the card as the male interconnect to a double cantilever-type female in the receptacle connector. In this design, there are only two contact surface joints, compared to three required for a pin-and-socket type of design. The flat-flex circuit also provides durability against shock and vibration.

To ensure that the interface surface endures 50,000 insertions and withdrawals with no significant increase in resistivity, the interconnect system must be plated with 50 μin. of gold over 200 μin. of nickel. And to protect the memory chips within the card against damage by static electricity, grounded metal panels can be incorporated on the top and bottom of the device to carry off electrostatic voltages.

Additional key design features to look for in an IC memory card system include:

- A write-protect switch on the RAM, which can be inserted in the write-enable circuit. When the switch is open, the card can't accept input data, thereby protecting previously stored data.
- A "card in place" pinout, which indicates to the host that the card is properly inserted.
- A battery detect pinout, which the host computer can use to detect the condition of the battery.
- A replaceable battery, which is desirable for an IC memory card using volatile RAM. With mixed MOS technology in a SRAM card, battery life will be from eight to 30 months, depending on SRAM capacity. For a full CMOS card, which permits a battery life as long as ten years, a non-replaceable battery is adequate.

When a battery needs replacing, the card should be inserted in the host device with V_{cc} operating so that no memory loss takes place. A properly designed card, however, can use a large capacitor in the circuit that will give the user from two to three minutes to replace a battery while the memory contents are maintained by the slowly decaying capacitor current.

Other protection mechanisms are also important. Latchup prevention, for example, can ensure that the memory chips are excited with the proper sequences of interface signals to prevent overwriting data when the card is inserted into or removed from a host device that's powered up. Power-down protection circuitry, if not already designed into the host,

IC memory cards and smart cards: two different things

IC memory cards are sometimes confused with smart cards, but—despite their similar appearance—the two are entirely different types of devices. Smart cards, containing a microprocessor and a small amount of memory, are used primarily in the consumer marketplace as credit cards or for automatic teller machine access and control, medical record storage, security access and other such applications. The IC memory card, in contrast, is strictly a memory device serving the OEM market with essentially the same function as a floppy disk: portable or auxiliary data storage and retrieval.

The memory capacity of the two cards also differs, with the smart card storing about 64 kbits of data, and the IC memory card as much as 1 Mbyte. The smart card is just 0.76 mm thick, while the memory card comes in at 3.1 mm. The smart card's International Standards Organization interconnect system, which consists of eight surface contacts on the face of the card, is quite different from that of an IC memory card. The IC memory card interconnect hasn't been standardized and may consist of as many as 38 contacts.

Smart cards are also considerably less expensive than IC memory cards, since they're used by the billions in the retail marketplace. The smart card typically costs about $3 to $4, while an IC memory card ranges from $15 to $500, depending on its capacity, memory type and other characteristics.

should be incorporated in the IC memory card to prevent the loss of memory contents when the host is turned off.

Applications will multiply

A relatively new technology that has found use only in niche applications, the IC memory card is more costly than a floppy disk system with comparable capacity. As the capacity of IC memory cards increases and their cost declines, however, applications will multiply, though increased use of the cards will be motivated by their particular virtues, not by cost of the media.

Today, the cards are beneficial alternatives for data logging and mass storage for computers that operate in harsh environments. The cards can also be used as personality modules for customizing terminals, control modules and other peripheral equipment. As a personality module, the IC memory card can considerably reduce the cost and time required to configure and upgrade peripheral equipment. The existing card is simply removed and replaced with an updated one; there's no need to open the case of the host or call in field service personnel.

The IC memory card is also the ideal medium for printer customization and for storage of interchangeable character styles, pitches, and widths, point sizes, fonts, languages and so forth. Small, light and durable, the IC memory card can be installed and removed thousands of times without degradation. What's more, the cards require very little storage space and can be easily and safely handled by a user, so they're extremely well suited to use in the office environment.

The IC memory card is also well suited for machine control in hostile environments, where contaminants subject electromechanical storage media to constant maintenance and, sometimes, extended downtime. Because of the contaminant-resistant interconnect system of IC memory cards, hostile environmental applications will prove to be a major market for the cards.

Future applications for IC memory cards will include uses in graphics and desktop publishing workstations, medical records equipment, musical instruments, liquid and gas flow measurement devices, air pollution measuring equipment, digital facsimile and copying machines, automotive/aeronautical map display systems and even voting machines. IC memory cards will be used in virtually any application requiring portable data and program storage. □

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

ITT Cannon
(714) 964-8262 ..... **Circle 265**

Please rate the value of this article to you by circling the appropriate number in the "Editorial Score Box" on the Inquiry Card.

High 509 *Average 510* *Low 511*

AP LABS I/O SUBSYSTEM

IOS™

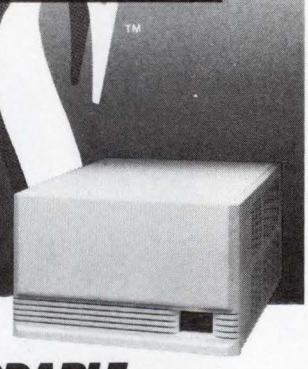

THE GATEWAY TO FAST, AFFORDABLE REAL TIME APPLICATIONS

The AP Labs IOS is a programmable I/O Subsystem enabling many general purpose and scientific host computers to perform data acquisition, storage and data routing tasks in real time environments. The IOS interfaces with Gould, DEC-VAX, Sun Microsystems, UNIX and NTDS Host computers, VMEbus computers and Aptec and FPS Processors.

The AP Labs IOS provides:
- LOW COST INTERFACE
- LARGE SELECTION OF PERIPHERALS—with drivers
- OFF LOADING OF THE HOST COMPUTER
- HIGH LEVEL SOFTWARE CONTROL
- TEST AND DEBUG PROGRAMS

AP Labs IOS supports Real Time Data Routing, Hardware In-The-Loop Simulation, Signal/Image Processing, Process Control, Data Acquisition, Data Switching and Real Time Graphics. Call today.

AP Labs
ADVANCED PROCESSING LABORATORIES, INC.
4411 Morena Blvd. Ste. 150, San Diego, CA 92117 (619) 272-8890

CIRCLE NO. 27

MULTI-MILLION DOLLAR VALUATION

MAJOR PUBLIC AUCTION

GE CALMA is now offering for auction, residual computer and office equipment, resulting from the 2nd quarter 1988 sale of its integrated circuit layout business.

Thursday, November 10th at 10am
Preview Wednesday, November 9th, from 9am to 4pm
501 Sycamore Drive, Milpitas, CA

CAD/CAM— Over 30 CAD workstations from Sun Microsystems and Apollo to include Sun models 3/50 & 3/60, Sun Stand-alone workstations models 160,260 and 3/60; Sun Servers model 3/180, never used 4/280 with 900 MB disk drive; large quantities of Apollo workstations model DN3000 color and monochrome; Apollo model 300 workstations with disk and floppy drives, large quantities of HP plotters 7580A,7585A,7585B, Calcomp plotter model 960 and much more!

Digital Equipment Corp. — DEC VAX 8650 with 32MB memory expansion to 64MB, cluster, S.I.A. and Unibus interface; 11/785, 11/780 with floating point accelerator, 16MB memory, Unibus adapter, 2 Massbuss adapters, DZII's with Ethernet I/F model DUENA and UDA 50 disc controller; 11/750 and others; TU77, TU78 master units; TK 50, RM05 & large quantity of RA81 disk drives and many others; CDC 300 & 600 MB drives, VT100 and VT101 terminals, DEC high speed/high quality printers model LP27, DEC LN03 laser printer, Imagen 5/480 high resolution laser printer with controller network interface, line printers, numerous DEC Microvax stations II/RC and II (some with graphic option), large quantities of DEC option/replacement modules EG memory, disc controller, communications, com boards, Versatec interface, complete VAX 785 floating point accelerator board set and much more!

Call for free brochures.
415-571-7400
Outside CA 1-800-445-DOVE

ROSS-DOVE COMPANY, INC
AUCTIONEERS SINCE 1917
MILLARD,ROSS & KIRK DOVE, Principals
CA License 616,1120,1265,1277

1118 Chess Drive
Foster City, CA 94404
TLX 510-100-6722 FAX 415-572-1502

Asset Conversion Consultants & Appraisers to Corporate America

CIRCLE NO. 28

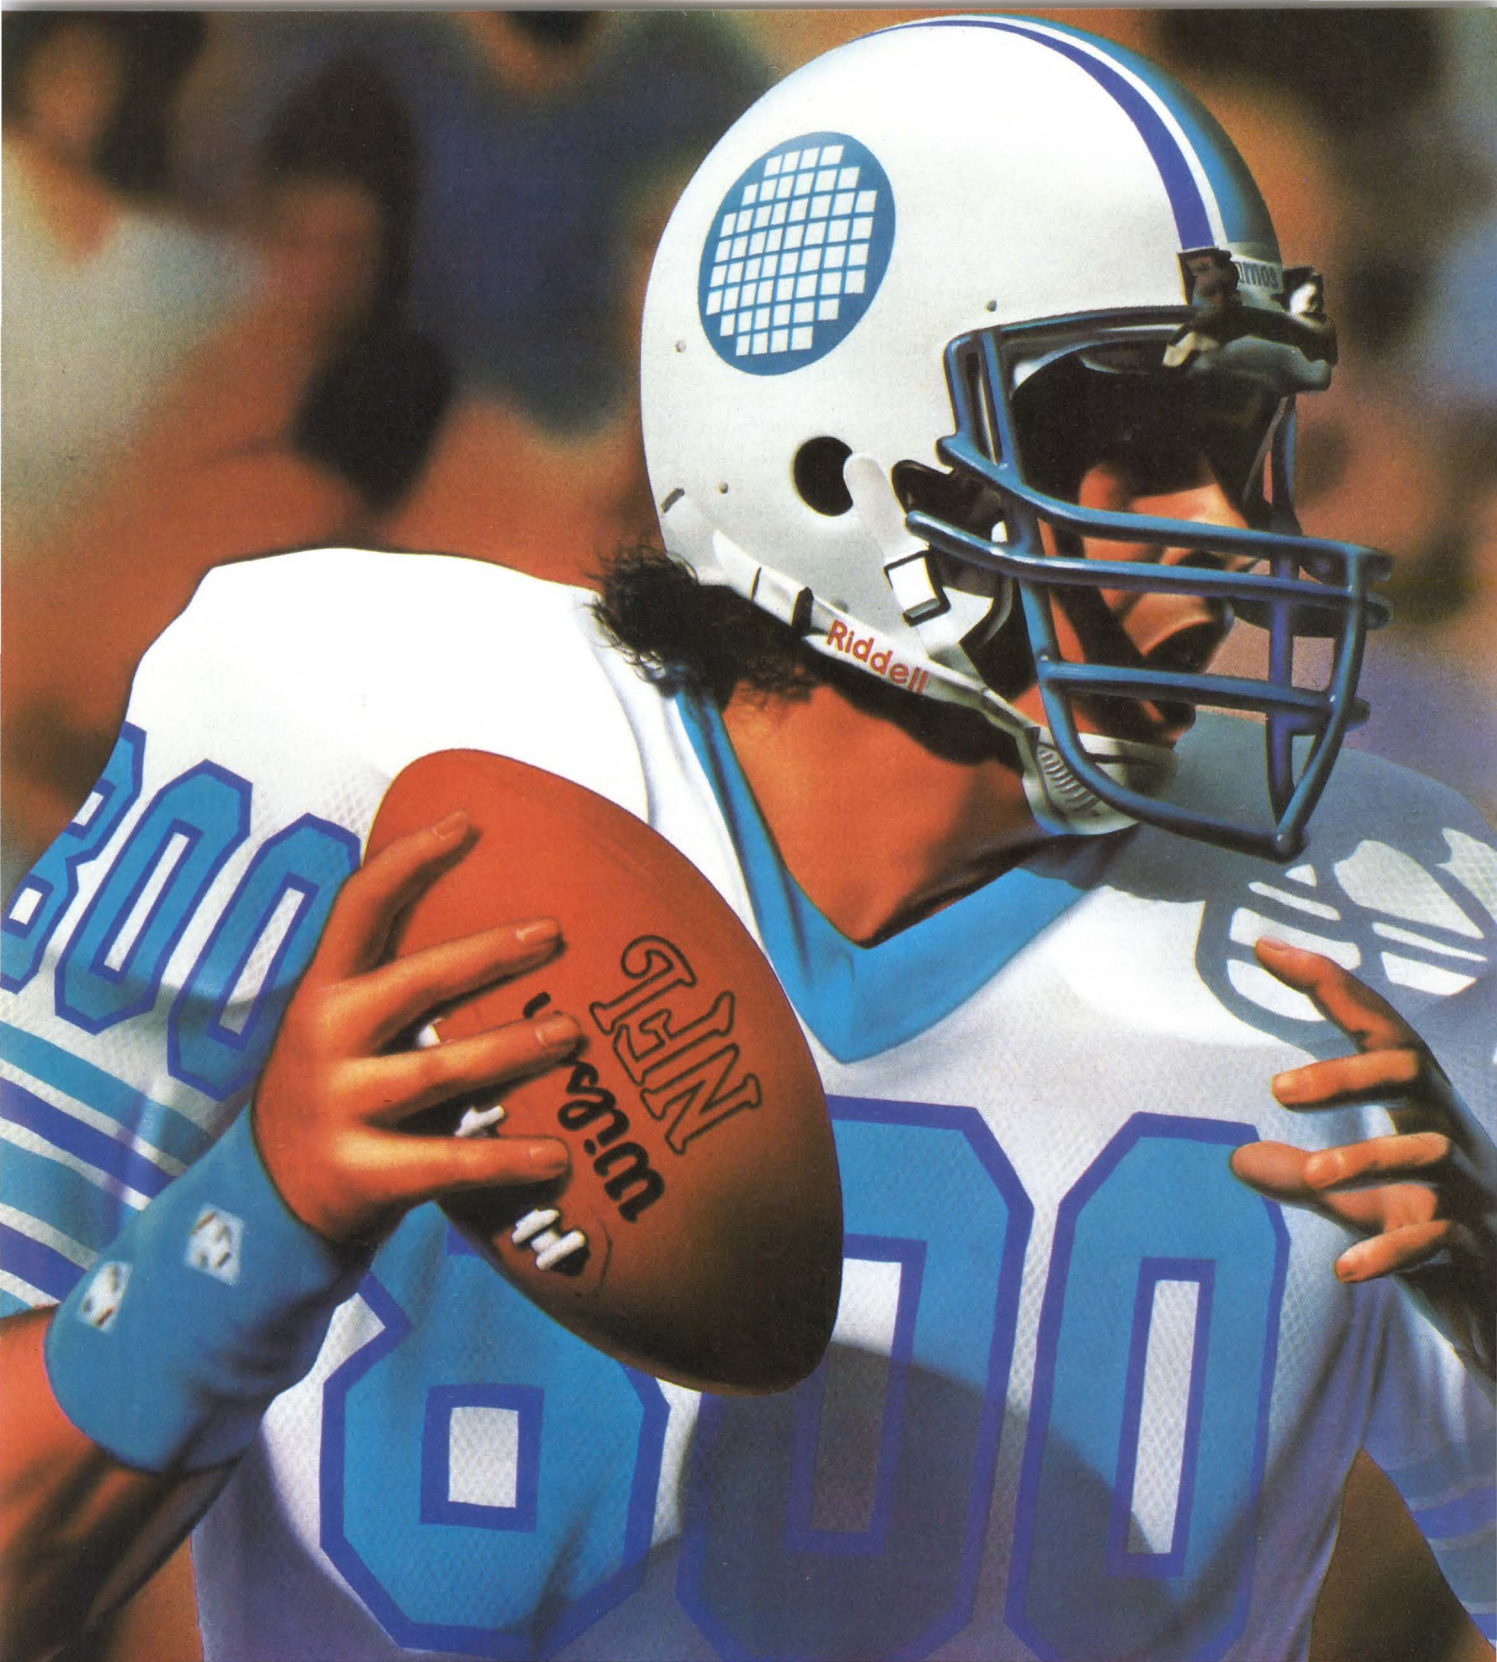

SINGLE PRECISION WHETSTONE LEAGUE

INMOS.
IMS T800 TRANSPUTER.
4.6 MEGAWHETSTONES.

INTEL.
386/387 16 MHz
1.8 MEGAWHETSTONES.

MOTOROLA.
68020/68881 20 MHz
1.5 MEGAWHETSTONES.

DEC.
VAX 11/780/FPA
1.1 MEGAWHETSTONES.

DOUBLE PRECISION WHETSTONE LEAGUE

ONE T800 TRANSPUTER GIVES
2.5 DOUBLE PRECISION
MEGAWHETSTONES...
SO WHEN IT COMES TO
PROCESSING POWER SEVEN
INMOS T800 CHIPS COULD
GIVE THE MIGHTY CRAY 1S,
RATED AT 16.1 MEGAWHETSTONES
A REAL RUN FOR ITS MONEY!

Motorola and Ⓜ are registered trademarks of Motorola, Inc. Intel is a registered trademark of Intel Corporation. VAX is a registered trademark of Digital Equipment Corporation. Cray is a registered trademark of Cray Research, Inc.

MORE MULTIPROCESSOR MUSCLE. MORE RAW PERFORMANCE.

When you're out in the trenches fighting it out with ordinary microprocessors, running out of muscle is all too easy. That's why you should look to the new T800 Transputer from INMOS.

The T800 is the fastest 32-bit, single chip, floating-point microprocessor available today. A quick glance at its statistics will show why nothing else is in its league...

32-bit enhanced RISC processor...64-bit on-chip IEEE floating-point processor...4K Bytes on-chip 50ns static RAM...Four 20 MBits/sec interprocessor communication links...Eight independent DMA engines. All on a single chip capable of sustained 1.5 MFLOPS...and 4.6M Whetstones!

And, if that's not enough raw power, the T800's links allow multiprocessor systems to be constructed quickly and easily – giving you 6 MFLOPS with four T800's...30 MFLOPS with 20...150 MFLOPS with 100...In fact, there's no limit to the number of Transputers you can use!

Programming Transputers couldn't be easier, with compilers for C, Fortran and Pascal, and the world's first concurrent programming language OCCAM.

Want to turbocharge your current system? No problem. Our exclusive Link Adaptor IC's allow Transputers to be connected to other microprocessors or peripherals.

Other team members include the pin compatible T414 Transputer, offering lower cost, 10 MIP performance and 0.75M Whetstones. Lined-up to provide all the I/O processing you need, the T212 16-bit Transputer is the ideal high performance controller and the M212 Disk Processor combines disk controller hardware and a Transputer on a single chip, supporting both Winchester and floppy disks. And the C004 Link Switch makes the design of software reconfigurable multiprocessor systems as easy as kicking an extra point.

Whatever field you're in – from real-time distributed systems to high-performance graphics, from fault-tolerant systems to robotics, Transputer technology can give you scalable performance at a cost you can afford.

Transputers are manufactured using an advanced 1.5 micron CMOS process which keeps the power consumption under one watt. So your system stays cool while under fire.

Transputers to MIL-STD 883C will be available in the first half of 1988.

If this all sounds like your kind of game, put the ball in play by contacting your local INMOS sales office today. And get ready to score.

DESCRIPTION			PERFORMANCE		AVAILABILITY		PACKAGE
Part No.	Word Length	Clock MHz	Integer Drystones	Floating Point Whetstones	Commercial	Military	
IMS T800-20	32-Bit	20	9500	4.6 Million	Now	Q2 88	84 PGA
IMS T414-20	32-Bit	20	9500	0.75 Million	Now	Q2 88	84 PGA
IMS T212-17	16-Bit	17	8000	–	Now	Q2 88	68 PGA
IMS T212-20	16-Bit	20	9500	–	Now	Q2 88	68 PGA
IMS M212-17	16-Bit	17	8000	–	Now	–	68 PGA
NETWORK SUPPORT PRODUCTS					AVAILABILITY		PACKAGE
Part No.	Description		Communication Speed		Commercial	Military	
IMS C004	Software configurable 32 way link switch		10 + 20 MBits/sec		Now	Q2 88	84 PGA
IMS C011	Link to system bus		10 + 20 MBits/sec		Now	–	24 Pin DIP
IMS C012	Link to system bus		10 + 20 MBits/Sec		Now	Q2 88	24 Pin DIP

THE TRANSPUTER TEAM

inmos®

INMOS, PO Box 16000, Colorado Springs, Colorado 80935. Tel. (719) 630-4000.

Orange County – 714-957-6018
Santa Clara – 408-727-7771
Denver – 303-252-4100
Minneapolis – 612-932-7121
Dallas – 214-490-9522
Boston – 617-366-4020
Baltimore – 301-995-6952
Atlanta – 404-242-7444

☐ Send me information on the Transputer Team. ☐ Send me the Transputer White Pages, a listing of third-party manufacturers' transputer-based products and services.
☐ Please have a Field Applications Engineer call. ☐ Please have a salesman call.

Name _____ Title _____

Company _____ Address _____

_____ Zip _____ Tel _____

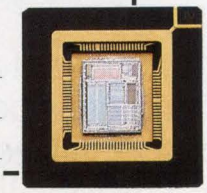

INMOS Transputer and IMS are trademarks of the INMOS Group of Companies.

CIRCLE NO. 29

Announcing another and another and another and another and another and another and another and another and another and another and another and another and another and another and another and another and another and another and another.

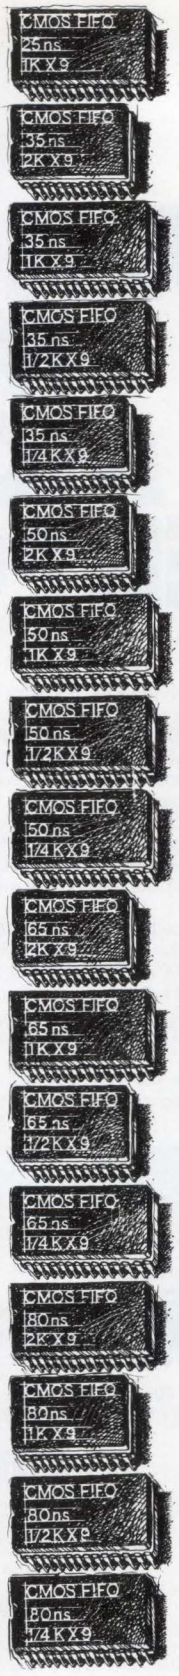

AMD is introducing 17 low power CMOS FIFOs.

They come in a broad range of speeds and densities. Even architectures for specialized applications.

And they're all available now.

Call us at (800) 222-9323. Ask for the Specialty Memory Data Book and find out about your favorite FIFO. And another and another and another.

Advanced Micro Devices

901 Thompson Place, P.O. Box 3453, Sunnyvale, CA 94088.

CIRCLE NO. 30

© 1988 Advanced Micro Devices, Inc.

MAJOR SYSTEM COMPONENTS

Compact dc-dc converters serve distributed power applications

John H. Mayer, Senior Associate Editor

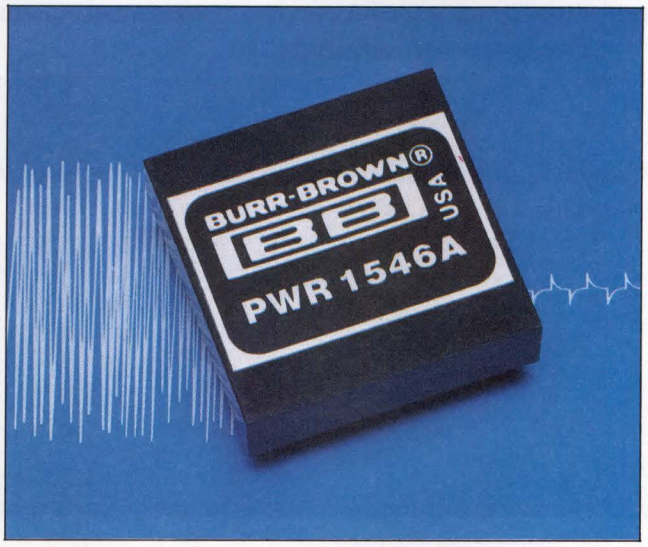

Designed to meet the sensitive noise requirements of high-accuracy data converters, gain amplifiers, and precision test equipment, Burr-Brown's PWR1546A 5-W dc-dc converter guarantees maximum output noise of 1 mV over a dc to 10-MHz bandwidth. The push-pull design keeps noise to a minimum by using multiple shields and extensive filtering.

By combining extensive use of surface-mount technology with higher switching frequencies, better thermal management and improved packaging techniques, vendors are developing a new generation of dc-dc converters in smaller, more efficient packages. These compact dc-dc converters will meet a market demand expected to grow annually in the double digits over the next few years.

What's driving this demand? Undoubtedly, the increased use of distributed power systems has played a major role. In telecommunications, industrial automation and computer systems, designers are deriving multiple benefits from building systems with a central supply that delivers power to a dc bus off which dc-dc converters power different parts of a system. Most important, distributing power in a system offers designers increased flexibility. They can now spread the cost of power in a modular system across the configuration.

"You can populate a system according to the usage of the customer," says Dave Roadruck, marketing manager for Burr-Brown (Tuscon, AZ). "With a distributed system, you can put in a very cheap central power supply and then put in precision supplies as you load boards. If the customer needs 30 percent of the system, he only spends a proportional amount of his cost on the power supply." Placing a dc-dc converter at the point of load rather than buffering 12 or 15 V across a system also offers better control and lower noise.

As usual, the primary goal in dc-dc converter design is to provide the smallest possible solution. Despite the advantages of distributed power, few dc-dc converters until recently could offer the size and efficiency needed to offset other power design alternatives. By integrating surface-mount technology and optimizing circuit design, the newest converters are offering power sources in packages half the size of their predecessors. Spurred by the growing availability of resistors, inductors and capacitors in surface-mount packages, low-power 1- and 2-W converters are moving to industry-standard, 24-pin, dual in-line packages. Even 5-W converters such as the AF series that was introduced last year by the Stevens-Arnold Division of Computer Products (South Boston, MA) are showing up in compact, 1-×2-in. footprints.

In fact, some small converters are available in single in-line packages. Once such product line is the NMA series from International Power Sources (Natick, MA). These small units provide 750 mW of power at an efficiency of up to 80 percent while occupying a scant 0.18 in.^{2} of board space. The parts family accepts inputs of 5, 12, 24 and 48 V and provides outputs of ±5, ±12 and ±15 V. "Everybody seems to be interested in the single in-line package," says Jack Swartz, vice-president of marketing. According to Swartz, his company will cut the size of the NMA series in half by next quarter. The converters will also be available in a surface-mount package.

Battery backup

One of the fastest growing applications for dc-dc converters is in battery backup systems for telecommunications. Fluctuating battery-powered systems demand dc-dc converters with a wide input voltage range. Products such as the 2100 series from Conversion Devices (Stoughton, MA) operate from 9 to 72 Vdc over ranges of 9 to 18 Vdc, 18 to 36 Vdc and 36 to 72 Vdc. By using a modified pulse-width topology and surface-mount parts, the designers squeezed the 20-W converters into a low-profile, 0.375-in.-high case. Power density is 13.3 W/in.^{3}, with operating efficiencies up to 83 percent. The units use the same pinout as the LP series from the Stevens-Arnold Division of Computer Products.

The 25-W LP-315 series from Power General (Canton, MA) offers the same wide input range in a similar low-profile package. Each of six models offers triple outputs of 5 Vdc at 4 A, ±12 Vdc at 0.25 A or ±15 Vdc at 0.25 A. Operation is specified from −25 to +85 °C, and mean time before failure is 300,000 hr.

Melcher (Natick, MA), a supplier of quasimilitary dc-dc converters, has brought that same wide input capability to a small 2-W unit. Intended to support 24-, 36-, 48- and 60-V batteries or compensate for long input line drops in telecommunications

Model	Power Rating (W)	Operating Efficiency (%)	No. of Outputs	Input Voltage (V)	Output Voltage (V)	Package Size (in.)	Price	Comments
Acopian PO Box 638, Easton, PA 18044 (800) 523-9478								**Circle 100**
Single output	—	65	S	—	1 to 185	—	$75 to $135	miniaturized, high isolation
Dual output	—	65	D	—	±12, ±15	—	$75 to $145	same as above
Burr-Brown PO Box 11400, Tucson, AZ 85734 (602) 746-1111								**Circle 101**
PWR 1726	1.5	—	D	7 to 18	15	1.6×1.1×0.5	—	medical isolation grade
PWR 1546A	5	60	D	5	±15	2×2×0.4	—	ultra low noise, regulated
PWR X00/X35	—	—	S, D	5 to 48	5 to 15, ±5 to ±15	—	—	1,000-V dc isolation
PWR 800/811	—	—	T	5 to 48	5, ±12, ±15	—	—	1,000-V dc isolation
PWR 70	3	—	D	15	±15	—	—	2,000-V dc isolation
PWR 74	3	—	D	5 to 22	±15	—	—	1,500-V dc isolation
PWR1017	3	—	D	15	±15	—	—	4 isolated channels
PWR 5038	2.5	—	T	5	5, ±15	—	—	unregulated
Calex 3355 Vincent Rd, Pleasant Hill, CA 94523 (415) 932-3911								**Circle 102**
5S12.450	5.4	60	S	4.75 to 5.25	12	2×2×0.4	$108	—
5D15.150	4.5	62	D	4.75 to 5.25	15	2×2×0.4	$105	—
24S5.5000XW	25	79	S	9 to 36	5	2.6×4.6×0.9	$160	—
Computer Products/Stevens-Arnold 7 Elkins St, S Boston, MA 02127 (617) 268-1170								**Circle 103**
PM600 series	0.5 to 1	50	S, D	5 to 12	5, 12, 15, ±12, ±15	24-pin DIP	—	miniature, regulated
EM600 series	1 to 1.5	60 to 80	S, D	5, 12	5, 12, 15, ±12, ±15	24-pin DIP	—	miniature, unregulated
CM600 series	1 to 2	50	S	5, 12	9	24-pin DIP	—	Cheapernet, Ethernet LANs
H series	1, 1.5	42 to 57	S	5, 12, 24, 48	5, 12, 15	1×2×0.4	—	low power, isolated
EA series	1.8	60	D	5	±12, ±15	1×2×0.4	—	low power
A/F series	4.5 to 6	58 to 69	D	5 to 48	5, 12, 15, ±12, ±15	2×2×0.4	—	same as PM900
LPS series	15	85	S, D	20 to 60	5 to 15, ±12, ±15	1.6×2×0.5	—	wide input, high density
SA-R400 series	15 to 48	74 to 90	S	8 to 40, 9 to 44	5.1 to 40	2.3×3.4×0.8	—	nonisolated
WRK series	60	80 to 84	S, D, T	9 to 72	5 to 15, ±12, ±15	3.5×5.5×0.9	—	2:1 input range
AF series	4.5	66	S, D	5,12	5, 12, 15, ±12, ±15	1×2×0.4	—	compact
WF series	5	75	S	7 to 32	5	2×2×0.4	—	regulated 5-V output
FW series	7.5	78 to 83	S, D	20 to 60, 36 to 72	5, 12	2×2×0.4	—	wide input
EM/PM9000 series	3 to 9	61 to 75	S, D	5, 12, 24, 28, 48	5, 12, 15, ±12, ±15	2×2×0.4	—	equivalent to A/F series
ES series	15	78 to 82	T	9 to 72	5/±12, 5/±15, 12	2.6×3×0.9	—	wide input triple
Conversion Devices 101 Tosca Dr, Stoughton, MA 02072 (617) 341-3266								**Circle 104**
100 series	1	50	S	5 to 48	5, 12, 15	1.3×0.8×0.4	$19	24-pin DIP, regulated
200 series	2	55	S, D	5 to 48	5, 12, ±12, ±15	same as above	$23	same as above
200 Hi series	2	55	D	5 to 48	±12, ±15	1.3×0.8×0.4	$26	high isolation
200 I series	2.5	56	S, D	5, 12	−9, 5	1.3×1.3×0.4	$32	miniature, for LANs
200 E series	2.5	65	S, D	5 to 48	5, ±12	2×1×0.4	$24	low profile, regulated
300 series	3	56	S, D	5 to 48	15, ±15	1.3×0.8×0.4	$29	24-pin DIP, regulated
300 Hi series	3	58	D	5 to 48	±12, ±15	1.3×0.8×0.4	$31	high isolation

key: D = dual; DIP = dual in-line package; S = single; SIP = single in-line package; T = triple

Model	Power Rating (W)	Operating Efficiency (%)	No. of Outputs	Input Voltage (V)	Output Voltage (V)	Package Size (in.)	Price	Comments
Conversion Devices 101 Tosca Dr, Stoughton, MA 02072 (617) 341-3266								**Circle 104**
300 E series	3	60	S, D	5 to 48	12, 15, ±15	2×1×0.4	$32	same as above
500 E series	5	65	S, D	5 to 48	5, 12, 15, ±12, ±15	2×2×0.4	$36	same as above
700 PW series	7.5	80	S, D	9 to 72	5, 12, 15, ±5, ±12, ±15	2×1×0.4	$47	wide input
600/900/1000 series	6 to 10	67	S, D	5 to 48	±12, ±15, 12, 15	2×2×0.4	$41 to $52	regulated
1500/2100 series	15 to 20	85	S, D, T	9 to 72	5, 12, 15, ±5, ±12, ±15, 5/±12, 5/±15	2×2×0.4	$48/$71	wide input, high density
2000/2500	20/25	80	S, D, T	9 to 72	5 to 15, ±5 to ±15, 5/±12, 5/±15	3×2.6×0.8	$63/$76.50	wide input, high density
Endicott Research Group 2601 Wayne St, PO Box 269, Endicott, NY 13760 (607) 754-9187								**Circle 105**
E1200	25	82 to 88	S	12, 15, 24, 48	5 to 250	2×2.5×1.2	—	—
E700R	3	65 to 75	S	5 to 28	5 to 250	1.4×1.5×1	—	regulated
E800U/HV	6	75 to 85	S	12 to 24	500 to 1,500	1.4×1.5×1	—	—
E900	7.5 to 12	75 to 85	S, D	5 to 24	5 to 500	2×2.5×1	—	—
E400	2.5	80	S	5 to 35	5 to 35	1×1.4×0.7	—	nonisolated, regulated
E500	3	75 to 85	S	5 to 28	5 to 500	1×1.4×0.7	—	unregulated
E700	3	75 to 85	S	5 to 28	5 to 500	1.1×1.2×0.9	—	unregulated
Integrated Circuits 10301 Willows Rd, Redmond, WA 98052 (206) 882-3100								**Circle 106**
DIP-DT	1	52	D	5 to 48	12, 15	1×0.5×0.4	$61	isolated, regulated
DIP	1.5	78	S, D	5, 12, 24	5, 12, 15	1×0.5×0.4	$51	wide range of inputs
MDP	1.5	78	S, D	5, 12, 28	5, 12, 15	0.9×0.8×0.3	$116	military screening avail.
DDR	3	52	S, D	5, 12, 24, 48	5, 12, 15	1.3×0.7×0.4	$65	isolated, 24-pin DIP
DCH	3	72	S, D	5, 12, 28	5, 12, 15	0.9×0,8×0.3	$79	military screening avail.
DCR	4	52	D	5 to 48	12, 15	1.8×0.9×0.4	$147	isolated to 500-V dc
MSR	4	75	S, D, T	28	5, 12, 15	1.1×1.1×0.4	$170	military screening avail.
DDC	5	70	S, D	5 to 48	5, 12, 15	1.3×0.7×0.4	$67	isolated, 24-pin DIP
DC	6.4	80	S, D	5 to 48	5 to 18	1.8×0.9×0.4	$69	isolated to 500-V dc
MFW	50	80	S, D, T	28	5, 12, 15	1.8×3.1×0.6	$695	military screening
MTW	30	86	S, D	28	5, 12, 15	2.7×1.4×0.5	$453	same as above
MHE	20	86	S, D	28, 12	5, 12, 15	2.1×1.1×0.5	$304	same as above
MHL	15	78	D	5	12, 15	1.9×1.4×0.5	$340	same as above
MTO	15	79	T	28, 12	12, 15	1.9×1.4×0.5	$409	same as above
International Power Devices 155 N Beacon St, Brighton, MA 02135 (617) 782-3331								**Circle 107**
G-series	5	72 to 78	S	9 to 72	5, 12, 15	2×2×0.4	—	battery-operated systems
H-series	7.5	75 to 85	S	7 to 32	5, 12, 15	2×2×0.4	—	nonisolated
L-series	45	78 to 87	S	9 to 40	5, 12, 15	2.6×3×0.8	—	nonisolated
C-series	0.5 to 1.5	32 to 78	S, D	5, 12	5, 12, 15, 9, ±12, ±15	1.3×0.8×0.4/ 1.3×1.3×0.5	—	PCB mountable
D-series	1.25 to 6	43 to 85	S, D	5, 12, 24, 28, 48	5, 12, 15, ±12, ±15	1×2×0.4	—	telecom
F-series	5 to 10	59 to 88	S, D	5, 12, 24, 28, 48	5, 12, 15, ±12, ±15	2×2×0.4	—	A-Ds, D-As, low profile
N-series	15	76 to 82	S, D, T	9 to 72	5, 12, 15, ±5, ±12, ±15	2.6×3×0.8	—	remote on/off
Q-series	25 to 30	80 to 84	S, D, T	9 to 72	5, 12, 15	2.6×4.6×0.8	—	4:1 input voltage range

Model	Power Rating (W)	Operating Efficiency (%)	No. of Outputs	Input Voltage (V)	Output Voltage (V)	Package Size (in.)	Price	Comments
International Power Sources 10 Cochituate St, Natick, MA 01760 (508) 651-1818								**Circle 108**
NMA series	0.75	80	D	5 to 48	±5, ±12, ±15	0.2×0.8×0.4	$11	miniature, SIP
BA series	15	70	S, D	5 to 48	5, 12, 15, ±12, ±15	0.4×1.3×0.8	$18	24-pin DIP
BC series	6	70	S, D	5 to 48	same as above	2×2×0.4	$29	industrial std pinout
DCA series	15	60	S, D, T	12, 24, 48, 110	5, 12, 15, 24, ±12, ±15	4.6×3×1	$45	low-profile chassis mountable
DCB series	25	75	S, D, T	same as above	same as above	5.3×3.3×1.2	$69	same as above
DCC series	50	75	S, D, T	same as above	same as above	6.3×3.8×1.3	$79	same as above
DCE series	100	80	S	12, 24, 48, 110	5 to 48, ±12, ±15	8×1.9×3.8	$145	—
Intronics 57 Chapel St, Newton, MA 02158 (617) 964-4000								**Circle 109**
K2439	100	80	T	20 to 60	5, ±15	6×4×1.1	$210	current mode, control 200 kHz
K2441	100	80	S	36 to 72	5	6×4×1.1	$175	same as above
K2444	100	80	S	36 to 72	24	6×4×1.1	$175	same as above
Kepco 131-38 Sanford Ave, Flushing, NY 11352 (212) 461-7000								**Circle 110**
ERD series	30, 60, 150	80	S	24, 48	5 to 48	3.7×5.1×1.7	$160	300-kHz FETs
Lambda Electronics 515 Broad Hollow Rd, Melville, NY 11747 (516) 694-4200								**Circle 111**
LW series	60	to 75	S, D	5 to 60	5 to 60	1.9×1.9×0.5 and up	$39 to $170	regulated
LJA-"D" series	45 to 90	60	S	42 to 58	5, 6, 12, 15, 24, 28	4.5×1.7×7.2	$274 to $658	narrow profile
LWS-389	250	75	S	42 to 60	5, 12, 15, 24, 28	1×4×6	$340	telecom, mobile equip.
Maxim Integrated Products 510 N Pastoria Ave, Sunnyvale (408) 737-7600								**Circle 112**
Max631/2/3	1	80	S	2 to 16.5	adjustable	8-pin DIP, SO	—	fixed/adjustable step-up
MAX635/6/7	1	85	S	2 to 16.5	–5, –12, –15, adjustable	8-pin DIP, SO	—	fixed/adjustable output
MAX641/2/3	10	80	S	2 to 16.5	adjustable	8-pin DIP, SO	—	fixed/adjustable step-up
Melcher 10 Cochituate St, Natick, MA 01760 (508) 653-9979								**Circle 113**
48ITS2-05-T	2	80	S	18 to 70	5	1×2×0.4	$48	wide input
PSR242-7	48	94	S	29 to 80	24	2×2.8×1	$55	–25 to 70 deg C
SMR56-8	30	81	S	10 to 35	5	6×3×1.3	$180	–40 to 85 deg C
BM1501	51	80	S	14 to 70	15	4.4×6.6×1.5	$222	2,500-V isolation
CM3020	42.3	82	T	28 to 140	5, ±12	4.4×6.6×1.5	$390	other outputs available
541WR1-2X1515-T	1	60	Q	4.5 to 5.5	±15	24-pin DIP	$50	–25 to 70 deg C
Nano Pulse Industries 440 Nibus St, Brea, CA 92621 (714) 671-7919								**Circle 114**
0.5 to 1 W	0.5 to 1	50	S, D	5 to 12	5 to 12	1.3×0.8×0.4	—	24-pin DIP, LANs
1.5 W	1.5	50	S, D	5 to 48	5 to 48	1.3×0.8×0.4	—	for LANs

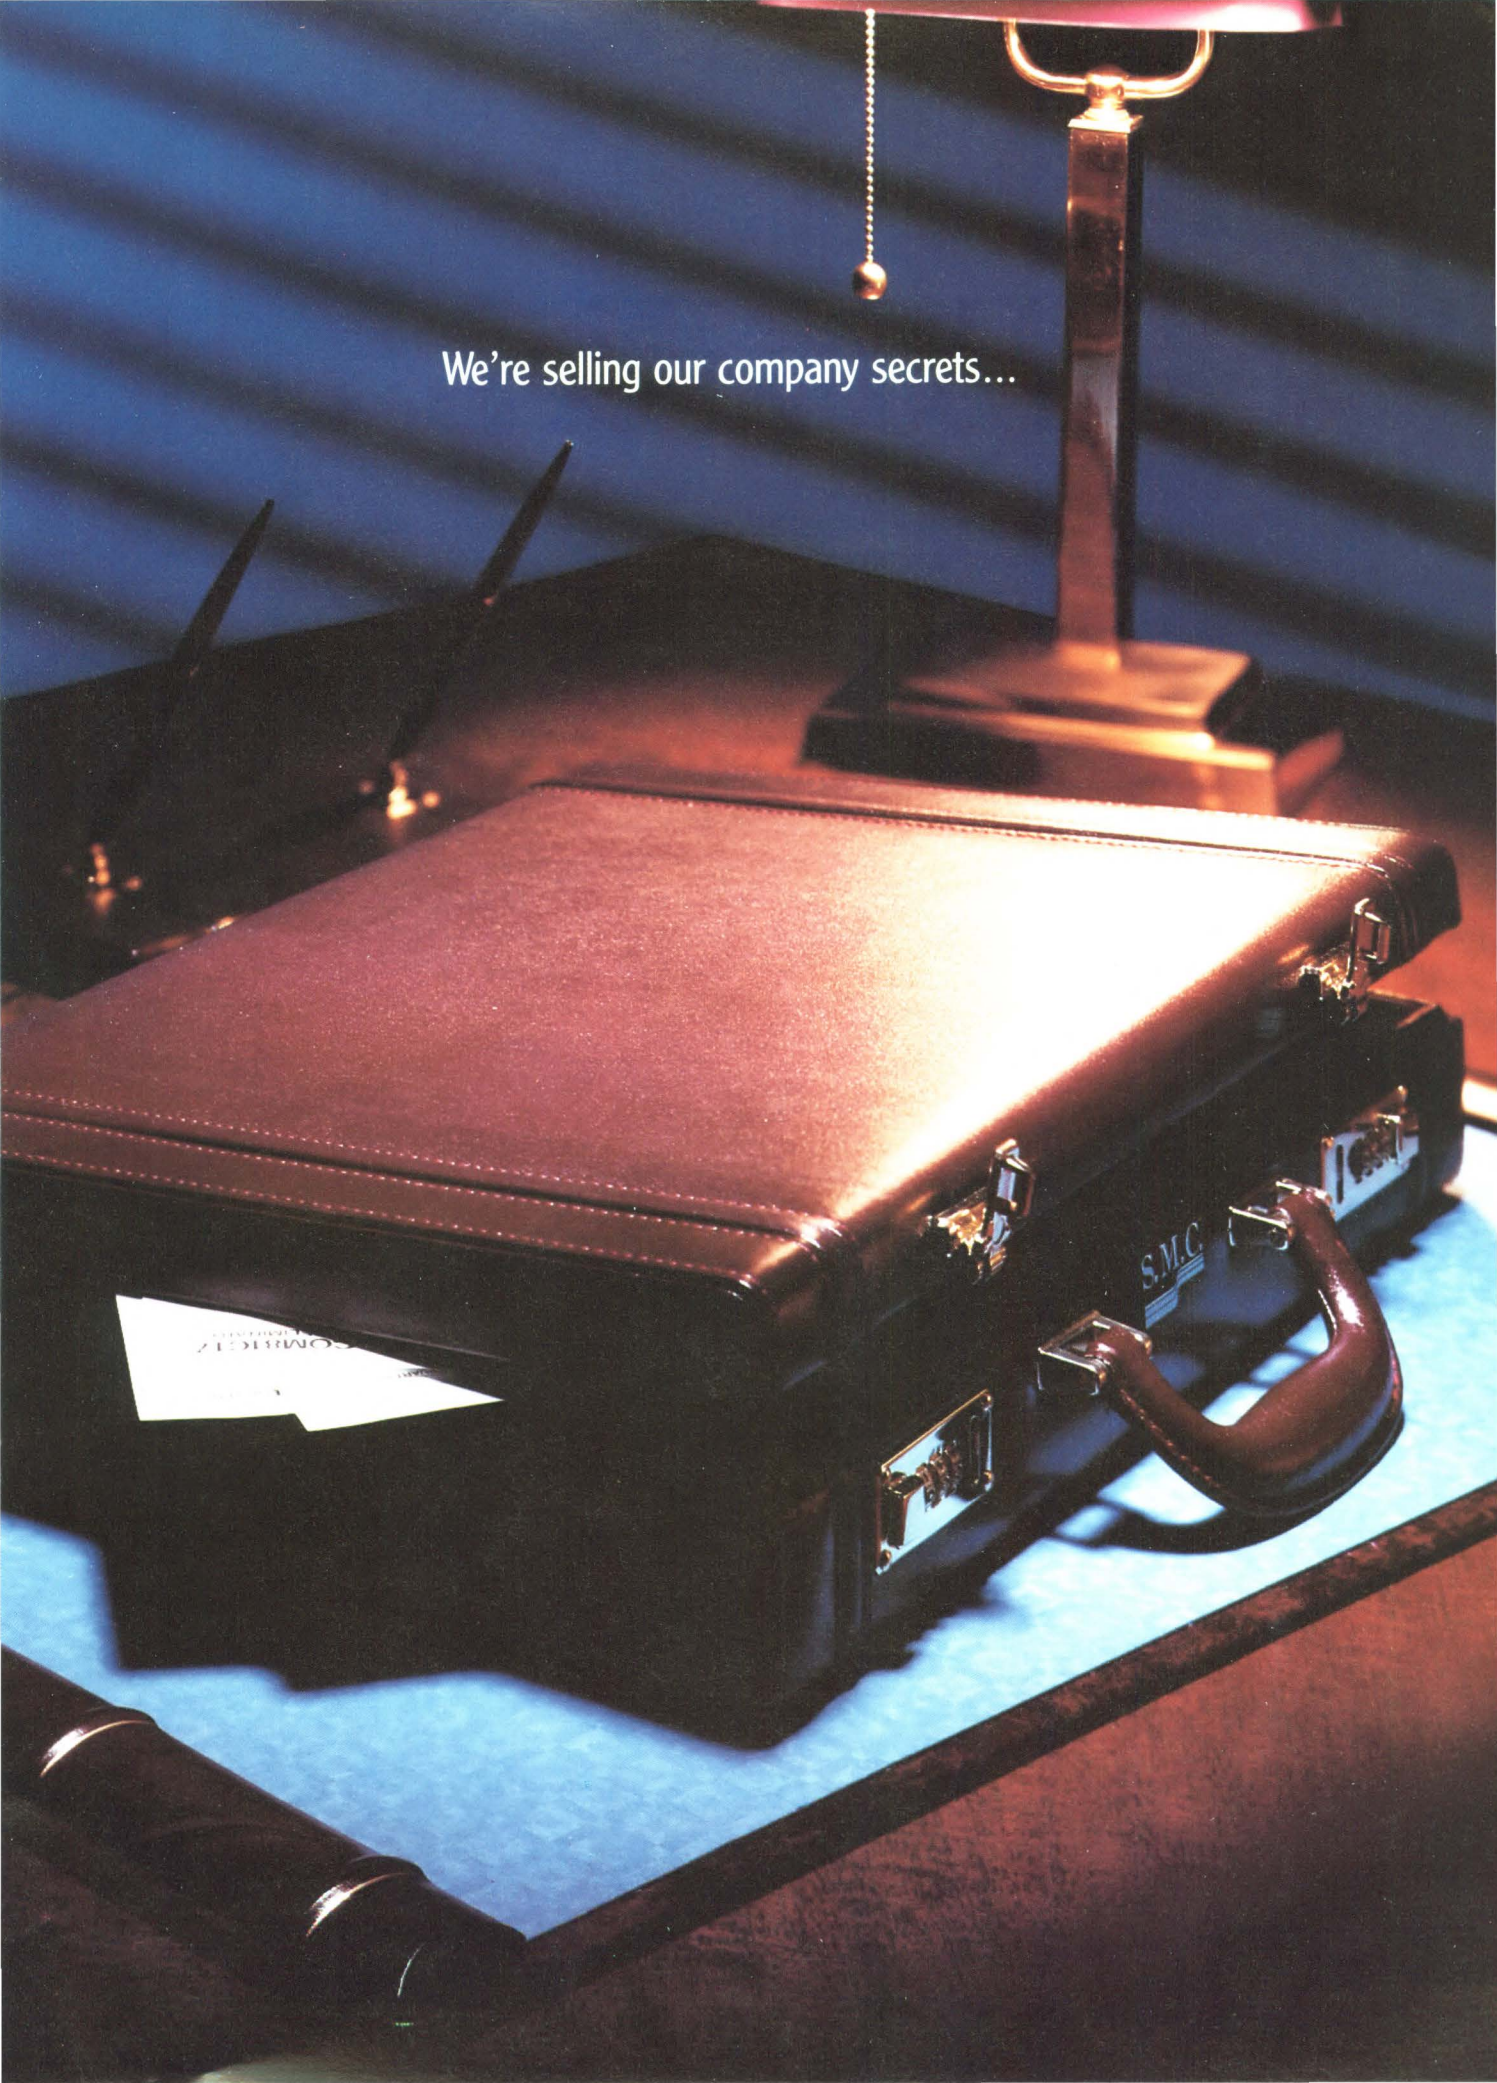
We're selling our company secrets...
S.M.C.

...AND SMC'S 17 YEARS O
CAN BE THE SECRET O

For nearly two decades SMC® has focused its technological expertise on the design and manufacture of application specific integrated circuits for the microperipheral industry. Our CMOS solutions for display, storage, data communications and local area networks begin with Standard Components such as:

Big Performance In A SCSI (Small Computer System Interface Controller)—SMC's MSD95C00.
This design features automatic arbitration for data transfer on an ANSI SCSI compatible computer bus. It also offers selection and reselection of devices, allowing simultaneous processor operation and transfer of entire commands from an initiator with minimal host processor intervention. This is due to the internal twelve byte buffer, which supports 5 MBytes/sec. synchronous and 3 MBytes/sec. asynchronous data transfer modes. It also features on-chip 48 mA high current drivers. It can be used with the MSD95C02 SunDAe™ companion chip in target applications.

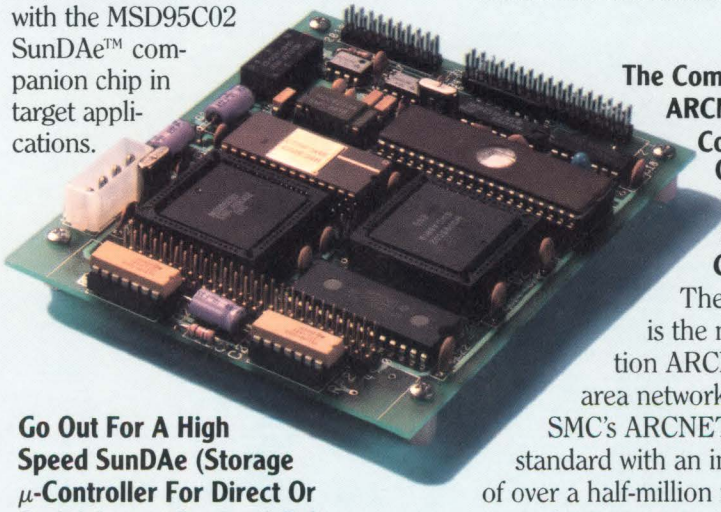

Go Out For A High Speed SunDAe (Storage μ-Controller For Direct Or Serial Access Devices) Drive—SMC's MSD95C02.
This is the next generation high speed micro-programmable mass storage controller ideal for embedded hard disk, QIC24 tape and optical drives as well as PS/2, ESDI and SMD disk controllers. It allows the user complete flexibility in design via the internal 32×32 writeable control store. A powerful on-chip 3 channel DMA controller arbitrates cache transfers without microprocessor intervention. This product is organized in a modular fashion, based on SuperCells™, for easy adaptation to special purpose applications.

Make IBM® Micro Channel™ Bus Interface Adapter Boards Spectacularly Simple—SMC's MCI94C18.
This CMOS device is fully compatible with the Micro Channel bus specification. It features support for 2 DMA channels with on-chip arbitration, shared local memory arbitration with the peripheral, a programmable wait state generator, as well as an on-chip timer. The MCI94C18 also provides a very flexible interrupt interface to the peripheral. Evaluation boards are available now.

The Complete ARCNET® Controller On A Single Chip—SMC's COM90C62.
The COM90C62 is the next generation ARCNET local area network controller. SMC's ARCNET, a defacto standard with an installed base of over a half-million nodes, has earned PC Magazine Editor's Choice Award for its high throughput, deterministic access and simple installation. Combining the functions of the COM90C26, COM90C32, and a crystal

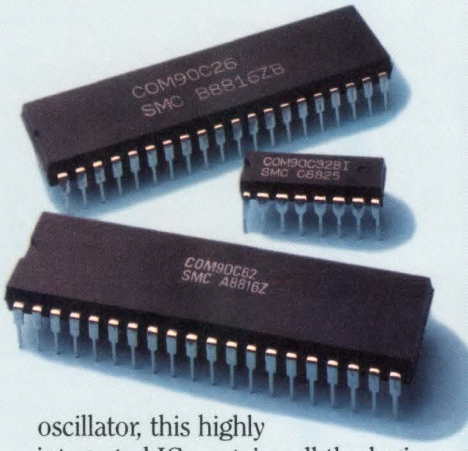

oscillator, this highly integrated IC contains all the logic necessary to control a token-passing ARCNET interface across standard coax, twisted pair or fiber optic cable. The COM90C62's high level of integration allows inclusion of a LAN interface directly on a PC mother board.

The Logical Choice For An ATLC (Advanced Terminal Logic Controller)—SMC's CRT92C07.
The CRT92C07 is a fully integrated CRT/attributes controller allowing implementation of VT220/320 terminals with less than 10 ICs. It features a 42 MHz video shift register, 80/132 column capability, a fully register programmable format and a writeable character font. The CRT92C07 can create multiple horizontal split screens capable of being smooth scrolled independently.

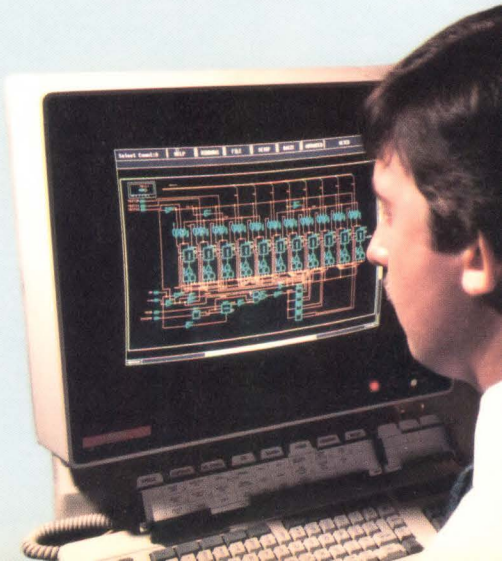

ASIC EXPERTISE YOUR SUCCESS.

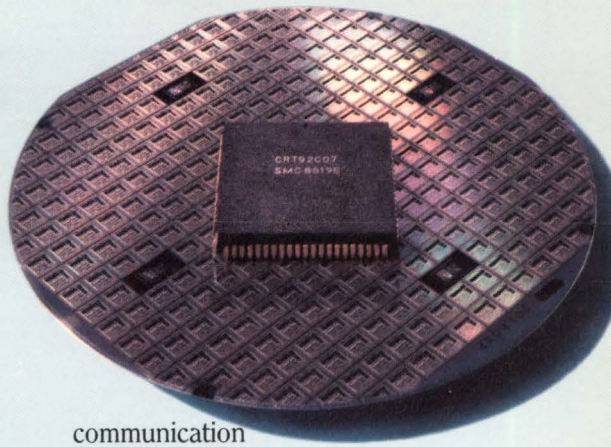

What's An SMC SuperCell?
Its foundation is the proven architecture of an existing SMC ASIC. Its agenda is the customization of capabilities to fulfill specifically defined application and performance objectives. Its methodology is a synergistic development process involving both user and SMC engineers. Its results are dramatically decreased developmental costs, reduced board space and a significantly accelerated production schedule. And we process our 2-micron SuperCell products with epitaxial silicon for greater latch-up resistance. It's SuperCell—another unique way SMC can meet your IC needs. Here are some examples:

Tune Into Expandable Micro Channel Capabilities—SMC's SuperCell MCI94C18.
Building on the MCI94C18 Bus Interface Adapter's capabilities (see above), a designer can integrate additional functions to enhance a specific component's performance. Yet, it is all achieved with a single, board-space-reducing, chip.

Two Ways To Customize Solutions In IBM Connectivity— SMC's SuperCell COM9064 And SuperCell COM52C50 (Twinax)
Our two popular IBM network controllers are now available as modular designs (SuperCell) for your specific adaptations. The COM9064 SuperCell is an MOS/LSI circuit used to facilitate high speed data transmission. It can implement an interface between IBM 3274/3276 compatible control units and 3278/3287/3289 compatible terminal units. Receiver and transmitter sections may be used independently of each other. With the COM52C50 Twinax SuperCell, communications interface to an IBM 5250 Twinaxial Bus can be combined with additional functions for connection to IBM's System 36 and 38 or to IBM's new Application System/400 Silverlake family.

Ace Design Problems With Our Asynchronous Communications Element (ACE)—SMC's SuperCell 82C50
The SuperCell ACE is ideal for integration within complex communication controller ASICs. Its soft-macro architecture can be modified for specific applications, with circuit features added or deleted as required.

Start With Six Storage μ-Controller Access Subfunctions In One IC And Build From There—SMC's MSD95C02 (SunDAe)
Create a proprietary device from what is already the next generation of intelligent programmable storage controllers for direct access and serial access devices. Composed of six peripheral controller subfunction SuperCells—Microsequencer, DMA Controller, ECC Generator/Checker and Processor, Disk and Host Interfaces—the MSD95C02 can be modified or interchanged with other SuperCell modules to increase functionality. This allows you to include proprietary circuitry for product differentiation and cost effectiveness.

STANDARD MICROSYSTEMS CORPORATION

35 Marcus Blvd., Hauppauge, NY 11788

Yes, I'd like to know more.

☐ Please send me your Data Sheet on: _____

_____

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

NAME _____

TITLE _____

COMPANY _____

ADDRESS _____

CITY _____ STATE _____ ZIP _____

SMC® is a registered trademark; SunDAe™ and SuperCell™ are trademarks of Standard Microsystems Corporation.
ARCNET™ is a registered trademark of Datapoint Corp.; IBM® is a registered trademark;
Micro Channel™ is a trademark of International Business Machines Corporation.

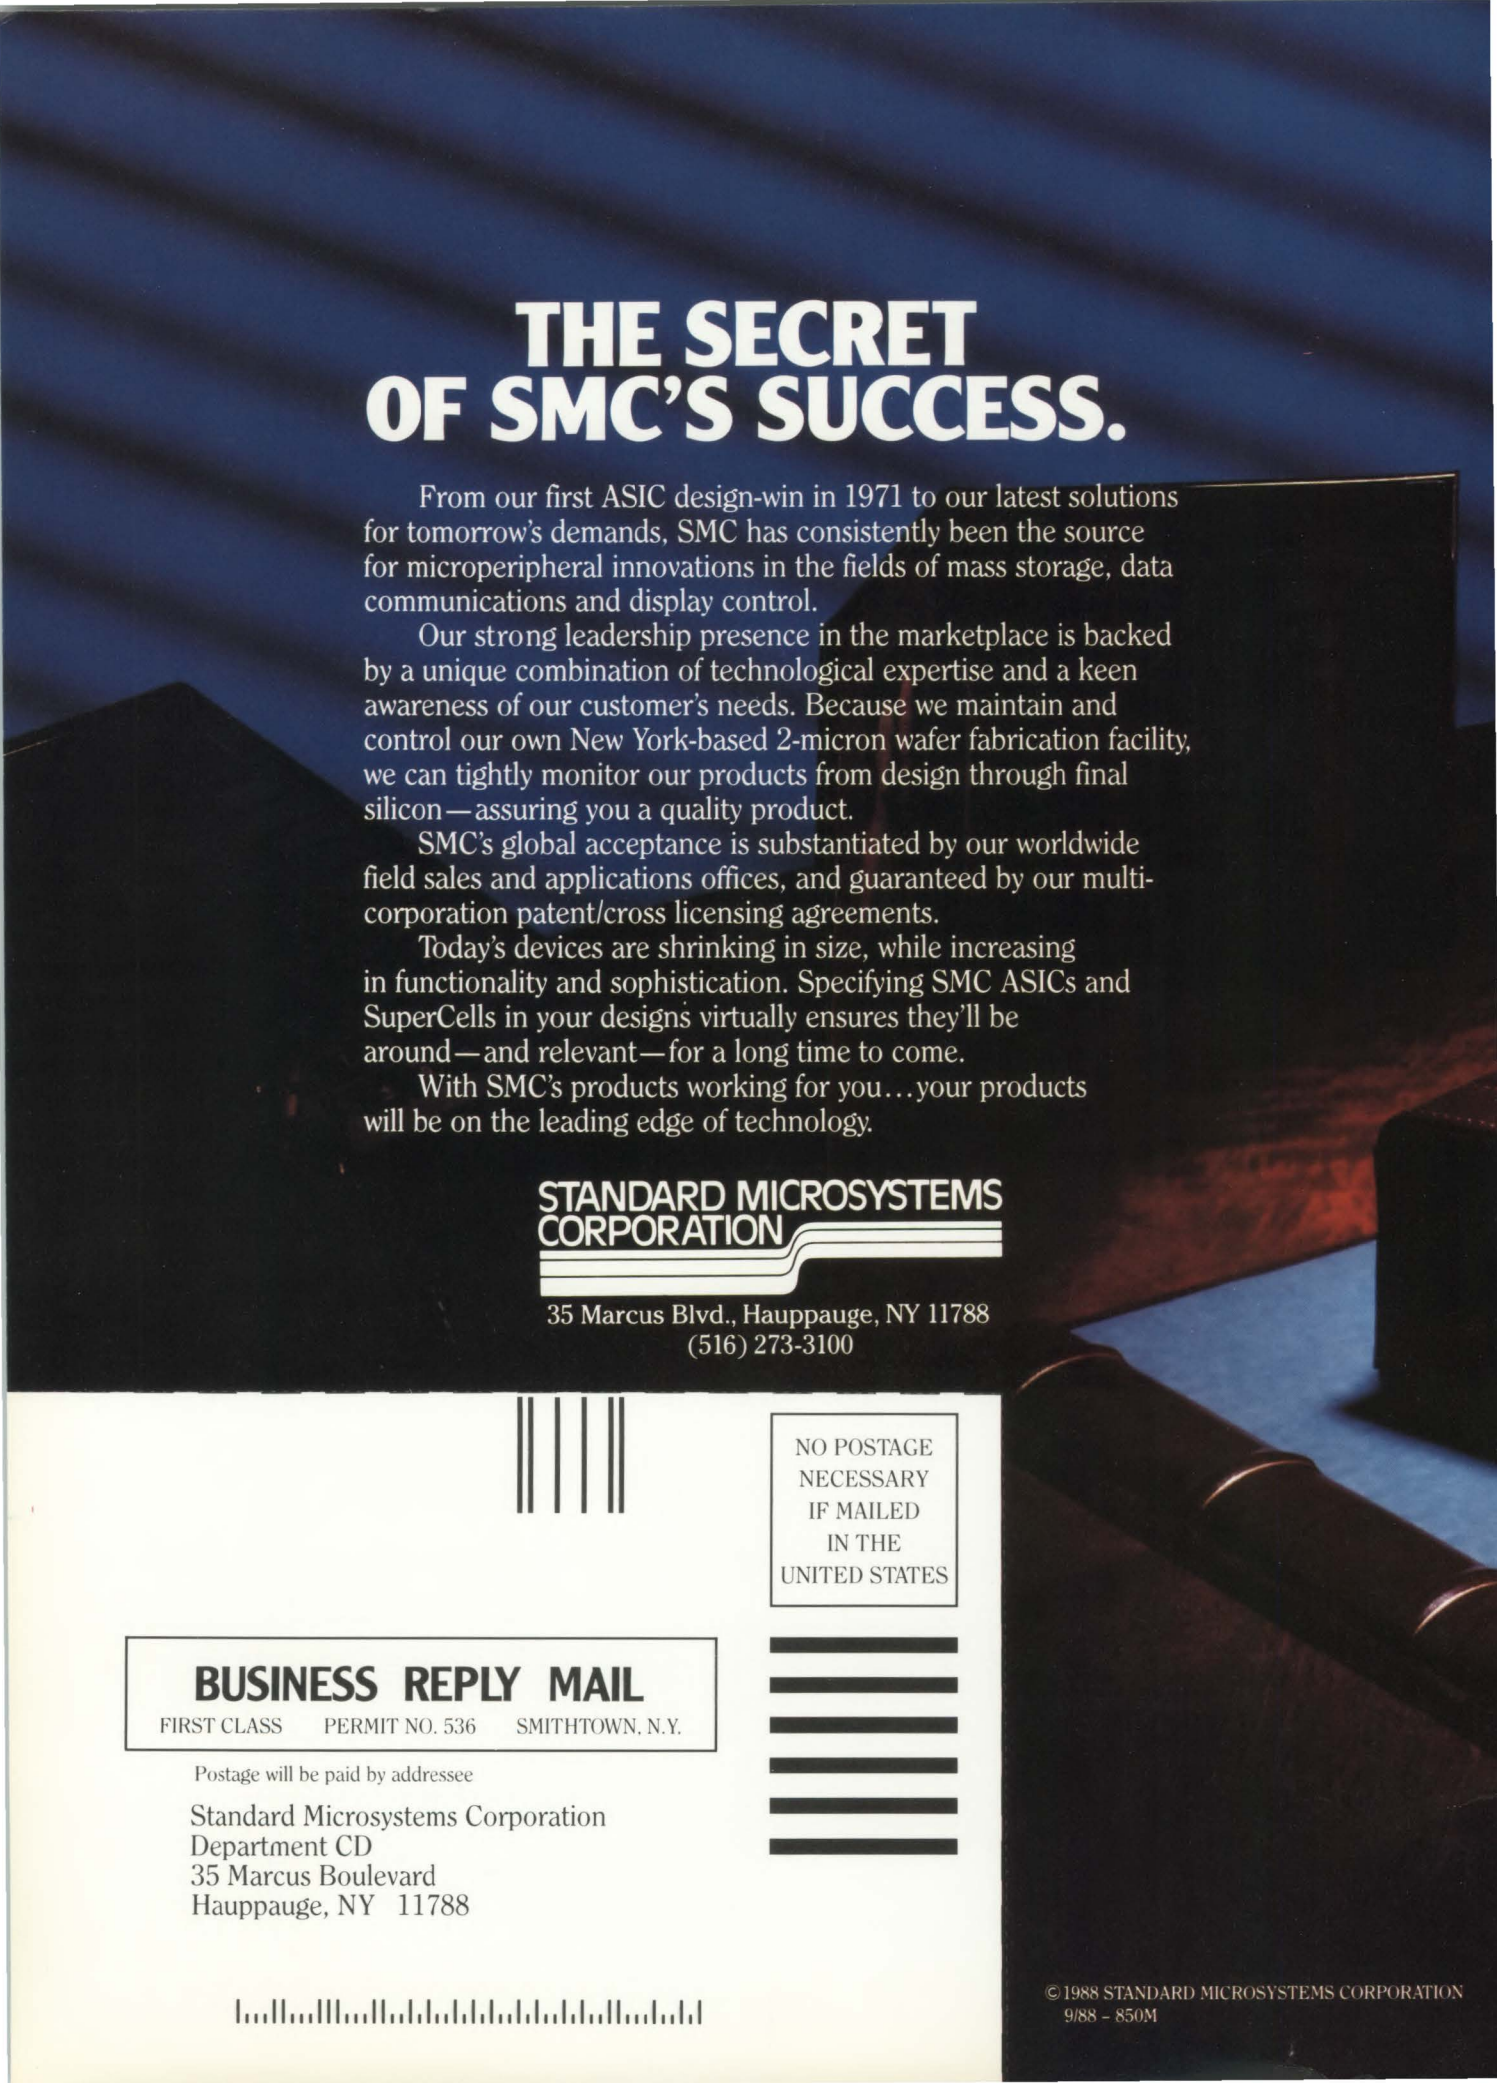

THE SECRET OF SMC'S SUCCESS.

From our first ASIC design-win in 1971 to our latest solutions for tomorrow's demands, SMC has consistently been the source for microperipheral innovations in the fields of mass storage, data communications and display control.

Our strong leadership presence in the marketplace is backed by a unique combination of technological expertise and a keen awareness of our customer's needs. Because we maintain and control our own New York-based 2-micron wafer fabrication facility, we can tightly monitor our products from design through final silicon—assuring you a quality product.

SMC's global acceptance is substantiated by our worldwide field sales and applications offices, and guaranteed by our multi-corporation patent/cross licensing agreements.

Today's devices are shrinking in size, while increasing in functionality and sophistication. Specifying SMC ASICs and SuperCells in your designs virtually ensures they'll be around—and relevant—for a long time to come.

With SMC's products working for you...your products will be on the leading edge of technology.

STANDARD MICROSYSTEMS CORPORATION

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

NO POSTAGE NECESSARY IF MAILED IN THE UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO. 536 SMITHTOWN, N.Y.

Postage will be paid by addressee

Standard Microsystems Corporation
Department CD
35 Marcus Boulevard
Hauppauge, NY 11788

© 1988 STANDARD MICROSYSTEMS CORPORATION
9/88 - 850M

Model	Power Rating (W)	Operating Efficiency (%)	No. of Outputs	Input Voltage (V)	Output Voltage (V)	Package Size (in.)	Price	Comments
Nano Pulse Industries 440 Nibus St, Brea, CA 92621 (714) 671-7919								**Circle 114**
3 W	3	50	S	5 to 48	5 to 48	2×2×0.4	—	for LANs
5 to 6 W	5 to 6	50	S	5 to 48	5 to 48	2×2×0.4	—	for LANs
10 to 12 W	10 to 12	50	S	5 to 48	5 to 48	3.5×2.5×0.9	—	for LANs
Power General 152 Will Dr, Canton, MA 02021 (617) 828-6216								**Circle 115**
400 series	1	55	S, D	5 to 12	5 to 15, ±12, ±15	0.8×1.3×0.4	—	input filter
520/530 series	1	—	S, D	5 to 12	5 to 12, ±12, ±15	0.8×1.3×0.3	—	isolated, SIP
LP-300 series	25	83	S	36 to 72	5 to 15	2.8×3×0.4	—	low profile
LP-315	25	75	T	9 to 72	5, ±12, ±15	3×3×0.4	$219	500-V dc isolation
730 series	25	82	T	9 to 72	5, –5, 12, ±12, ±15	2.6×4.6×0.8	$179	input filter
750 series	15	80	S	9 to 72	5, 6, 12, 15	2.6×3×0.8	$119	telecom, portable equip.
DC60 series	50 to 72	85	S	9 to 72	5 to 24	4×6×1.8	—	open PCB design
Powerline 10 Cochituate St, Natick, MA 01760 (508) 655-7987								**Circle 116**
5244	19	85	S	45 to 75	22 to 32	6.6×4.2×1.3	$106	—
5355	24	85	S	80 to 140	22 to 30	6.6×4.2×1.9	$128	—
C500	100	85	S	10 to 16	5	6.5×4.2×4.3	$336	—
Qualidyne Systems 3055 Del Sol Blvd, San Diego, CA 92154 (619) 575-1100								**Circle 117**
Series 50	125	80 to 84	S, T	24, 48, 155, 270	5 to 48, ±12, ±15	2.4×4.6×0.5	$140/$217	PCB mountable
QML	50	60 to 65	—	24, 48	2 to 30	2.3×5×var.	$1.25/W	custom-configured outputs
Reliability 16400 Park Row Rd, Houston, TX 77084 (713) 492-0550								**Circle 118**
S-Pac	1	50/48	S, D/S	5, 12	5, 10, 12, 15	1.3×0.6×0.3	$18.75 to $23	unregulated and regulated
V-Pac unreg.	1	50	S, D	5, 12	5, 10, 12, 15	1.3×0.6×0.4	$20.10 to $23.25	RS-232 drives, op amps, E/EPROM
V-Pac reg.	1	48	S, D	5, 12	5, 12, 15	1.3×0.6×0.4	$21.40/ $25.25	op amps, A-Ds, D-As, RAMs
2V-Pac reg.	2	50	D	5	±12	1.3×0.6×0.4	$25.25	op amps, A-Ds, D-As
2A module reg.	2	55	D	5	±12, ±15	2×1×0.4	$38.50	same as above
LAN-Pac	2	54 to 75	S	5, 12	–9	1.3×0.6×0.4	$17.60 to $22.90	LAN applications
LAN 2VA	2	70	D	5, 12	5, 10	1.3×0.6×0.4	$17.60	Intel LAN chip set
LAN 2E12 R10-5	2	39	D	10 to 15	5, 10	1.3×1.3×0.4	$24.15	same as above
2SPSU5.2	2.3	75	S	5	5.2	1.3×0.6×0.3	$20.85	ECL
VA5-5/2VA5-5	1/2	62/69	D	5	5, –5	1.3×0.6×0.4	$20.10	line drivers, receivers
5W	5	57	S, D	5, 12, 24	5, ±12, ±15	2×2×0.4	$58.25	battery backup
5A, D	5	50	S, D	5, 12, 24	5, ±12, ±15	2×2×0.4	$58.25	A-D, D-A, amps, data acquisition
3W	3	48	S, D	5, 12, 24	5, ±12, ±15	2×2×0.4	$50.50	battery backup, A-Ds, D-As
Telcor-Pac	5, 25	60, 80	S, D	20 to 72	5, 12, 15	2×2×0.4/ 3.4×2.7×0.8	$45/$95	two versions/telecom
GA12RN5.1-5.1	6	65	D	9 to 32	5, –5	2×2×0.4	$64.60	nonisolated

Model	Power Rating (W)	Operating Efficiency (%)	No. of Outputs	Input Voltage (V)	Output Voltage (V)	Package Size (in.)	Price	Comments
Rifa PO Box 853904, Richardson, TX 75085 (214) 480-8300								**Circle 119**
PI series	0.25 to 6	55 to 75	S, D	5, 12, 24, 48	5, 12, 15	1×0.5×0.4 to 2×2×0.4	$12 to $50	—
PKC series	15	85	S, D, T	18 to 36/ 36 to 64	5, 12, 15	3×2×0.4	$60 to $75	–45 to 85 deg C
PKA series	25 to 40	80 to 85	S, D, T	19 to 35/ 39 to 64	5, 12, 15	3×3×0.7	$70 to $90	–45 to 85 deg C
SGS Semiconductor 1000 E Bell Rd, Phoenix, AZ 85022 (602) 867-6100								**Circle 120**
GS-11209	2.4	75	S	11/16	–9	2×1×0.5	$18	I/O isolation, for LANs
GS-R400	140	75/90	S	9/46, 16/60	40	3.3×2.6×0.8	$20 to $30	soft start, reset output
GS-R400/2	140	80/85	S	9/40	5/12	2×2×0.5	$25 to $30	same as above
Teledyne Semiconductor 1300 Terra Bella Ave, Mountain View, CA 94039 (415) 968-9241								**Circle 121**
TSC962	0.5	50	S	3 to 18	1.5 to 15	8-pin DIP, 16-pin SO	$2.70	switched capacitor converter
TSC7662	0.3	50	S	3 to 18	1.5 to 15	same as above	$2	same as above
TSC7660	0.15	50	S	1.5 to 10	3 to 15	8-pin DIP, 8-pin SO	$1.25	same as above
Todd Products 50 Emjay Blvd, Brentwood, NY 11717 (516) 231-3366								**Circle 122**
DC series	250	75	S	40 to 60	5, 12, 15, 24, 28, 48	5×2.5×9	$296	covered or self-cooled versions
DCT-253-0515	250	75	T	40 to 60	5, 15, –15	5×2.5×9	$336	same as above
DCT-253-0512	250	75	T	40 to 60	5, 12, –12	5×2.5×9	$336	same as above
Toko America 1250 Feehanville Dr, Mount Prospect, IL 60056 (312) 297-0070								**Circle 123**
CPS	0.3	55	T	4.5 to 5.5	–22, –30, 4.5(ac)	1×0.6×0.5	—	power vacuum display, RS-232
DL	2	65	S	12	–5	1.2×0.6×0.3	—	miniature
MPS	4.9	65	D	4.8 to 5.3	+12, –5	2×2×0.2	—	comm. control
Valor Electronics 6275 Nancy Ridge Dr, San Diego, CA 92121 (619) 458-1471								**Circle 124**
24-pin reg.	2	50	S	5, 12	multiple	1.3×0.5×0.5	—	networks, varying line voltage
24-pin unreg.	2	65	S, D	5, 12, 15	multiple	1.3×0.6×0.5	—	networks, stable line voltage
24-pin unreg. w/shield	2	65	S, D	5, 12, 15	multiple	1.3×0.6×0.5	—	LANs, low radiated noise
24-pin reg. w/2-kV iso.	2	50	S	5, 12	multiple	1.3×0.6×0.5	—	network transceiver chip
24-pin unreg. w/2-kV iso.	2	65	S, D	5, 12, 15	multiple	1.3×0.6×0.5	—	network transceiver chip
24-pin reg. w/disable	2	50	S	5, 12	multiple	1.8×0.8×0.5	—	low power drain
8-pin SIL	2	75	S	5, 12	5, 12	1.3×0.6×0.3	—	Cheapernet, Arcnet
5-pin square	3.5	55	D	5, 15	multiple	1.3×1.3×0.5	—	transceiver applications
5-pin 15 W	15	75	D	5, 15	multiple	1.5×2×0.5	—	concentrator, file-server applications

Model	Power Rating (W)	Operating Efficiency (%)	No. of Outputs	Input Voltage (V)	Output Voltage (V)	Package Size (in.)	Price	Comments
Vicor 23 Frontage Rd, Andover, MA 01810 (617) 470-2900								**Circle 125**
VI-100	to 100	80 to 90	S	10 to 370	5 to 48	2.4×4.6×0.5	—	modules, high density
VI-200	50 to 200	80 to 90	S	10 to 400	5 to 48	2.4×4.6×0.5	—	same as above
Mega Module	50 and up	80 to 90	S	10 to 400	5 to 48	2.5×4.9×0.6	—	same as above
Master Module	50 and up	90	D, T	10 to 400	5 to 48	4.9×4.9×0.6	—	chassis mountable

High Speed Bridge Brings Transputer Parallel Processing Power to VMEbus.

BBK-V2 by paracom

10 MIPS RISC-based Transputer acts as a high performance VMEbus master

Unlimited expansion by plugging in additional 4-Transputer boards or by linking to external processor arrays

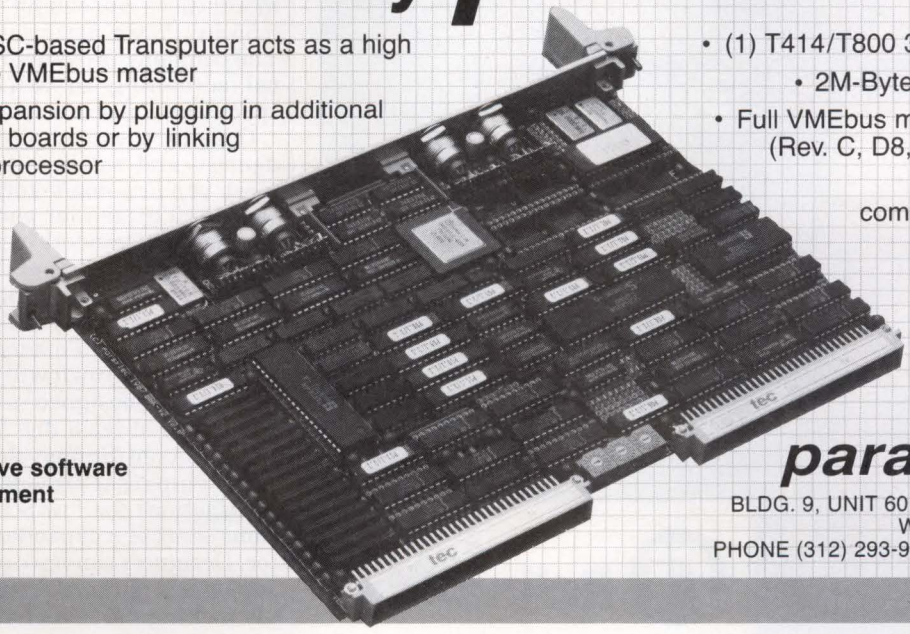

- (1) T414/T800 32-bit RISC processor
- 2M-Bytes of dual ported RAM
- Full VMEbus master/slave capability (Rev. C, D8, D16, D32, A24, A32)
- (4) 20M-Bit/sec. communications channels
- Supports ultrafast Transputer block move with D8/D16/D32

Extensive software environment

paracom inc

BLDG. 9, UNIT 60, 245 W. ROOSEVELT RD.
WEST CHICAGO, IL 60185
PHONE (312) 293-9500 FAX (312) 231-0345

CIRCLE NO. 32

MAJOR SYSTEM COMPONENTS

applications, the ITS series has an 18- to 70-Vdc input range in a 1-×2-×0.4-in. package. "We try to cover as many worst-case possibilities as we can so the user doesn't have to buy four different units with different input ranges," says James Stedman, applications engineer for Melcher. The single- or dual-output parts are built around the SI9100 regulator chip from Siliconix (Santa Clara, CA).

Burr-Brown is building converters for medical and high-resolution data-acquisition applications that require very high isolation and low barrier capacitance. One of the newest, the PWR1726, is a 1.5-W converter that's high-pot tested to 8,000 V, peak, for 60 s. The converter is a single-channel, dual-output 15-Vdc unregulated device with a 3,500-Vdc minimum isolation voltage and a 7-pF typical barrier current. Operating over a 7- to 18-Vdc input range, the unit offers an extremely high calculated mean time to failure—in excess of 100 yr at an ambient temperature of +25 °C with rated power.

■ High power, high density

At the high-power end, Vicor (Andover, MA) has pushed the density of its parts to an unprecedented 36 W/in.^{3} by using a proprietary zero-current-switching design. Operating at frequencies up to 2 MHz, the 35 models in the VI-100 series deliver 50 to 100 W in 2.4-×4.6-×0.46-in. cases. A second series, the VI-200, runs up to 200 W. Designers can evaluate the Vicor converters in off-line applications using a new evaluation board that houses a 750-W ac front end.

Densities are also growing in converters designed for military/aerospace applications. Accepting input ranges from 19 to 40 Vdc, the single-, double- and triple-output MFW models from Integrated Circuits (Redmond, WA) deliver up to 70 W in 3.5-, 5-, 12- and 15-V models. All the units are packed in hermetically sealed metal cases measuring 1.77×3.08×0.55 in. Power density is 23 W/in.^{3} The company provides environmental screening to MIL-STD-833C specifications as an option.

Despite these improvements in density and smaller module sizes, most vendors don't anticipate that the trend to smaller parts will slow down. Many vendors expect to manufacture smaller 1- and 2-W parts, some in single-wide packages. Application-specific parts will also become more prevalant. This month, for example, Burr-Brown will unveil a dc-dc converter designed specifically for ECL. And at least one IC manufacturer, Maxim (Sunnyvale, CA), offers ¼-in.^{2} chips that provide all converter functions for low-power applications short of a capacitor and inductor. Designers can have a customized supply in a minimal footprint by simply picking the inductor and capacitor their application requires. □

DID YOU KNOW...

. . . that Reliability is *the* alternate source for DC-DC converters supplied by Power General, Stevens-Arnold, Power Products, Burr-Brown and Intronics—among others?

We *do* have the expertise. Reliability introduced the miniature DC-DC converter 15 years ago. Today, our line includes over 150 one to thirty watt power sources. With worldwide sales and manufacturing operations, we can provide service anywhere.

NOW YOU KNOW. . .

. . . so call Reliability today at (713)492-0550. Most of our power sources are available from stock.

P.O. Box 218370 • Houston, Texas 77218
Phone (713)492-0550 • TLX 4620383 RELIUI

Reliability®

CIRCLE NO. 33

The Complete UNIX/Real-Time Connection.

OS-9 utilizes common UNIX features:

- Multi-user, Multi-tasking
- UNIX I/O model
- UNIX task model: (fork, set priority,...)
- Pipes and Signals
- Shell User Interface
- Hierarchical Disk File System

...and more to connect you to Real-Time:

- Fast task switching
- Data modules
- Easy system generation and reconfiguration
- Fast interrupt servicing
- Events and semaphores
- Resident OS-9 compilers, assemblers and debuggers

And no other operating system complements UNIX with real-time functionality better than OS-9. OS-9, like UNIX, is a complete operating system to provide extensive C and math libraries, independent file managers, inter-process communications, debuggers and resident compilers. Yet OS-9 is 100% ROM-able and executes in real-time to host thousands of imaging, process control, data acquisition, communications and robotics projects worldwide.

Best of all, every module of UniBridge and OS-9 is written entirely by Microware to provide you one source for technical support and service.

From UNIX and C, to OS-9 and VME, contact Microware today and let your next real-time application shine with UniBridge: The Complete UNIX/Real-Time Connection.

OS-9 UniBridge

Integrating UNIX and C applications into VMEbus systems has just been made easier with Microware's introduction of UniBridge. UniBridge is a complete development and communications package that allows you to connect your SUN 3, DEC VAX, HP 9000 or Motorola Delta workstation, to the world's premier real-time operating system: OS-9.

UniBridge provides a gateway so you can utilize UNIX in every phase of VMEbus system integration. From hosting C application development to monitoring large real-time networks, UniBridge connects UNIX to your OS-9 target system; whether it's ROM-based, disk-based, networked or stand alone.

UniBridge contains three high performance software modules that link UNIX and OS-9:

- OS-9/XCC Cross C Compiler allows you to produce compact, re-entrant and position independent C applications on your workstation for high speed execution on your OS-9 target.
- OS-9/ESP Ethernet Support Package provides industry standard BSD 4.2 sockets for TCP/IP communications with full FTP and Telnet support. UniBridge lets your workstation become a real-time server.
- OS-9/SRCDBG debugs resident VMEbus applications at the C source level, all interactively, across the Ethernet network from your UNIX terminal. And UniBridge makes all file access and transfer transparent from system to system while you debug.

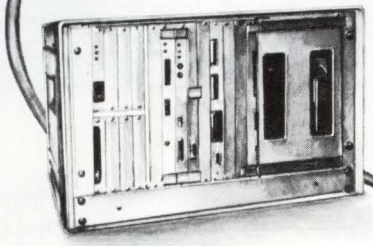

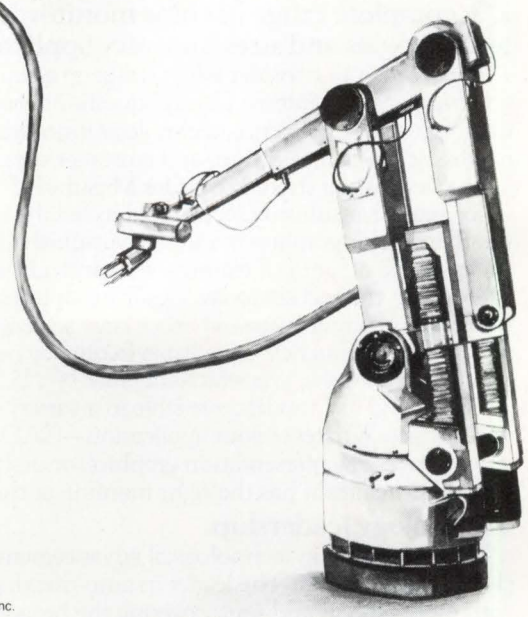

OS-9 and UniBridge are trademarks of Microware Systems Corporation.
VAX and VMS are trademarks of Digital Equipment Corporation. UNIX is a trademark of Bell Laboratories. Sun is a trademark of Sun Microsystems, Inc.

microware® OS-9

Microware Systems Corporation
1900 N.W. 114th Street
Des Moines, IA 50322
Phone: 515-224-1929

Microware Western Regional Office — Santa Clara, CA 408/980-0201;
Microware Japan 0474 (22) 1747

Distributors: West Germany 06221-862091; **France** (1) 43-33-96-38; **Italy** (6) 762221; **Switzerland** (056) 83-3377; **Sweden** (018) 138595; **Australia** (02) 919-4917

CIRCLE NO. 34

How To Size Up

A complete range of color monitor technologies and sizes for every application.

There's a lot to consider when sizing up a supplier of color graphic display monitors. Critical questions arise, such as manufacturing experience, technology innovation, proven reliability, product selection, and customer support.

That's why you should consider Mitsubishi.®

For years, Mitsubishi Electronics has led the industry in supplying color graphics monitors. Mitsubishi offers the widest range of monitor features, sizes and advanced technologies on the market today. Quality monitors to support your exact requirements—whether large screen, small screen, fixed-frequency or multiple-frequency performance.

A comprehensive line which includes 14″, 15″, 16″, 20″, 26″, 33″, and 37″ models, available in a variety of performance ranges. Whatever your application—CAD/CAM, image processing, presentation graphics, or desktop publishing—Mitsubishi has the right monitor, at the right cost.

Technology leadership.

When you size up technological advancements, Mitsubishi clearly leads the way. The leader in auto-tracking technology with more models and sizes, covering the broadest range of horizontal scan frequencies, than anyone else. The leader in dynamic beam focus (DBF) technology for sharper, clearer images to the edge of the screen. And the leader in microprocessor-enhanced, digital scan mode memory technology for optimum display size and clarity in *any* mode, text or graphic, or when switching between multiple modes.

OEM experience and commitment.

And when you size up Mitsubishi's continuing commitment to serving the OEM market, you'll find a company with the industry's broadest range of experience, service, applications assistance and resources to support you with high-quality monitors in volume.

To size up monitor technology, one name is all you have to know: Mitsubishi. Call or write Mitsubishi Electronics America, Inc., Computer Peripherals Division, 991 Knox Street, Torrance, CA 90502, (213) 217-5732.

The leader in auto-tracking technology.

Mitsubishi's auto-tracking monitors automatically track horizontal and vertical frequencies, eliminating manual frequency adjustments. From 14″ to 37″ display sizes, Mitsubishi offers you a total solution in auto-tracking convenience and versatility.

The leader in large screen technology.

Mitsubishi offers the two largest auto-tracking monitors in the industry today. Bright, vivid colors on our big 33″ or 37″ monitors result in greater impact and add a new dimension to the growing presentation graphics market.

© 1988 Mitsubishi Electronics America, Inc.
Mitsubishi is a registered trademark of Mitsubishi Electric Corp., Tokyo.

Screen images produced with permission from the following companies *(trademarked software package name follows company name)*: Autodesk, Inc. *(AutoShade)* and *(AutoCAD)*; MacroMind, Inc. *(VideoWorks II)*; Computervision Corporation *(Personal Designer)*; SuperMac Software *(PixelPaint)*; Software Products International *(Open Access II)*. Image printed with permission from Media Logic, Inc. *(Artisan)*.

Monitor Technology.

Actual unretouched screen images.

Screen Size	Mitsubishi Model	Horizontal Scan Frequency (kHz)
14″	XC1409C	15.7
14″	XC1410C/30C	22 or 15.75
14″	XC1429C	31.5
14″	AUM1381A*	15.7 ~ 36
14″	FA3415/25*	15.7 ~ 36
14″	HF1400/50	15.5–20
14″	HF2400/50	20–25
14″	HF3400/50	30–35
15″	FHF3500 (flat square)	30–35
16″	C8652	47–52
16″	FG6600	60–65
16″	HL/FL6605*†	30 ~ 64

Screen Size	Mitsubishi Model	Horizontal Scan Frequency (kHz)
20″	C3920/21/22	15-24
20″	C6920/21/22	28–35
20″	HA3905*	15.7 ~ 36
20″	HL6905*†	30 ~ 64
20″	HG6905□	40–67
20″	HJ6905*□	40 ~ 70
26″	C3510	15–18
26″	C6512	28–34
33″	XC3310*	15 ~ 35
37″	XC3710*	15 ~ 35

*Auto-tracking †Microprocessor-enhanced programmable display settings □Available with DBF

Sales Offices:
Carrollton, TX (214) 241-5300
Minnetonka, MN (612) 938-7779
Mt. Prospect, IL (312) 298-9223
Norcross, GA (404) 368-4845
Piscataway, NJ (201) 981-1001
Sunnyvale, CA (408) 730-5900
Torrance, CA (213) 217-5732
Woburn, MA (617) 938-1220

The leader in microprocessor-enhanced technology.

Mitsubishi was the first on the market with auto-tracking microprocessor-enhanced monitors. Digital scan mode memory features a microprocessor in the monitor which can remember up to 20 combinations of settings for horizontal width, phase, centering and pincushion correction, as well as vertical height and centering.

The leader in dynamic beam focusing (DBF).

With advanced DBF technology, Mitsubishi offers OEMs distortion-free, high-resolution displays. DBF technology changes the beam shape from elliptical to circular as it strikes the corners and edges of the CRT, resulting in the highest picture quality possible over the entire screen.

Performance Peripherals.

MITSUBISHI ELECTRONICS

See Us in COMDEX Booth #1128

CIRCLE NO. 35

ANNOUNCING: The EuroBus Conferences
The European Bus/board Conferences and Expositions

EuroBus/89 - Germany
8 - 10 May 1989
München Sheraton
Münich

EuroBus/89 - UK
4 - 6 September 1989
Novotel Hotel
London

WHY THE EuroBus CONFERENCES?

That rapidly-developing bus architectures are playing an increasingly important part in board-level microcomputer, sub-system, and systems development is a welcome fact of business life. Nowhere is this more true than in the European community, with emphasis added as Europe looks to the trade barrier changes due in 1992 and American companies turn to foreign trade for a competitive window.

And, board-level marketing is becoming increasingly competitive, as new buses enter the picture (e.g., Futurebus, VXI, PiBus, and NuBus) and others contemplate the possibilities of 64-bit architecture and higher (the "Superbus"?)

Perhaps even more so than in the U.S., Europe will be the "proving ground" for the various new and time-tested approaches. Certainly it is the best arena for applications comparisons.

WHAT ARE THE EuroBus CONFERENCES?

The EuroBus Conferences mean to be *the* information transfer vehicle bringing all of the architectures together under one roof, at one time, in a concentrated atmosphere of serious board-level customers, whose attention is not diverted by the sprawling mass of products found in horizontal trade shows.

The EuroBus Conferences are **technical meetings** with relevant exhibits. Meaningful sessions and seminars are designed to attract those who seek the information needed to make buying/specifying/design decisions. The educational environment will not only foster understanding among the user community, but also enhance the sharing and escalation of technology advances among academic and industry participants.

Direct contact with a full range of manufacturers of bus-oriented boards, systems and peripherals will precipitate on-the-spot responses from engineers and buyers. And, recognizing that not all European attendees can attend a given conferences in a specific location, the EuroBus/89 Conferences will be run in two sections - Germany and the UK - to maximize attendance potential, increase market exposure, and acknowledge the need for bus/board information from Israel to Spain to Scandinavia.

WHO SHOULD EXHIBIT?

Board and Systems Manufacturers * Manufacturers of Card Cages, Connectors, Packaging, Software * Manufacturers with Products in Bitbus, Exorbus, Fastbus, Futurebus, G-64 and 96, Multibus I and II, NuBus, PC Bus, QBus/Bibus, S-100, SCSI, STE/STD, Versabus, VME, VXI.

Produced by:
MultiDynamics, Inc.
13762 Newport Ave. #204
Tustin, CA 92680
714/669-1201
FAX 714/669-9105

(In London)
Roger Sherman
London House
243-253 Lower Mortlake Rd.
Richmond TW9 2LL England
01/940-4625
FAX 01/948-1442

WHAT ABOUT THE TECHNICAL PROGRAMME

Under the direction of noted industry expert, Dr. Paul Borrill (National Semiconductor), as Conference Coordinator, with able support from UK and German-based computer professionals, the seminar/session programme will be organized in four advanced-education tracks:

* **Board-Level Design:**	32-Bit Buses, Multibus I and II, VMEBus, VXI Instrumentation, Futurebus Applications, NuBus Report
* **System-Level Design:**	Intelligent Peripheral Buses, System Packaging, Crate to Crate Communications, High Performance Architectures, Development Tools for RISC, Stepwise Debugging, Bitbus.
* **Application Development:**	Real-Time Systems, Communications Technology, Real-Time Software, Designing for the Military, Real-Time Graphics, Industrial Networks
* **Electives:**	ASIC, Testability, Market Forecast, Transputer Technology, Ada, Surface Mount, Benchmarks.

And more as bus/board technology unfolds.

If your company would like to organize a session, present a paper or conduct a seminar, please contact Anne Weber at 714/669-1201 in the US (fax 714/669-9105) or Roger Sherman in the UK at 44-1-940-4625 (fax 01 948 1442). A Programme Call will be sent to you at once.

TO: Roger Sherman, London House
243-253 Lower Mortlake Rd.
Richmond TW9 2LL England

Please send me information on:

☐ Exhibiting ☐ Attending the EuroBus/89 Conferences

Name _____

Title _____

Company _____

Street Address _____

City/Province _____

Country _____ P.C. _____ Phone _____

COMPUTERS AND SUBSYSTEMS

DSP boards extend STD Bus speed to 10 Mips

Although the STD Bus standard has been accepted for widespread use in industrial electronics, relatively few dedicated DSP subsystems are available. It's precisely in the realm of industrial control and analysis, however, that the need for DSP functions is especially compelling. Severe number-crunching demands are placed on computer systems by applications such as analog signal processing, image recognition and enhancement, and closed-loop motion-control systems. The CS9320 Series of DSP boards provides such processing power to the STD Bus.

Based on the Texas Instruments 32020 family of DSP chips, the boards provide up to 5 Mips of processing speed with a bipolar processor and up to 10 Mips using a CMOS version. Memory options include program capacities of 8, 32 or 64 kwords, data memory of 8 or 32 kwords, and dual-port memory from 4 to 64 kwords in binary increments. The 28-pin JEDEC memory sockets can accommodate chips of various speeds and types, such as static RAM, battery-backed RAM, PROMs or EPROMs.

Jumpers on the board are used to configure the various sizes of program and data memory, as well as the amount of 8-bit dual-port memory accessible by an external processor. Additional groups of jumpers adapt the board to either fast or slow memory chips and let users choose either 16- or 20-bit shared memory addressing.

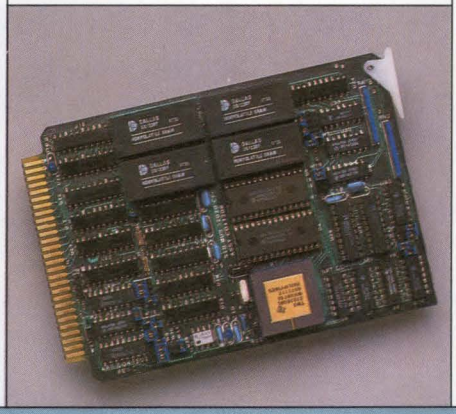

Two modes of operation are possible with the DSP: stand-alone or slave. In the stand-alone mode, the board doesn't generate bus requests or require bus acknowledgements to execute I/O operations. This mode accommodates full I/O interrupt handling. In a multiprocessor configuration using the device as a slave, communication takes place through the dual-port memory, letting only one processor assume the task of handling interrupts.

An on-board power-on-reset circuit ensures proper startup. Software support consists of complete documentation, a self-test program and source-code applications examples supplied on a floppy disk and in a PROM. These software aids not only let a user verify the board's performance, but they also provide a tutorial in DSP programming. Depending on the version (bipolar or CMOS) and memory options, prices for the CS9320 Series DSP boards range from $815 to $1,025.

Cheshire Computer
493 W Main St
Cheshire, CT 06410

Circle number 126

DESIGN AND DEVELOPMENT TOOLS

Software enhancements speed MC88000 development on Mac II

Several software enhancements to an Apple Macintosh-compatible development tool have increased its flexibity in designing RISC applications. The Tektronix TL88K-P Development Board lets users write applications for Motorola's 88100 RISC chip prior to that chip's market availability. The software improvements to the board include several libraries of prewritten, commonly used code, allowing the creation of programs that can take simultaneous advantage of both the Mac's MC68020 processor and the RISC chip.

The development board includes the 88100 processor chip set, three cache memory management units and 8 Mbytes of memory. The 20-MHz tool plugs into the NuBus on the Mac II and benchmarks at 17 Mips, or 34 kDhrystones and 6 MFlops. "The Apple Macintosh and TL88 combination yields system performance from 10 to 30 times faster than is possible on the Macintosh alone," claims Chip Schnarel, manager of the computer systems department at Tektronix.

The tool lets developers design programs in which the Macintosh and Motorola processors can share information and workload, with the more intensive portions of a program running on the RISC chip, enhancing system speed. These dual-processor applications can be programmed under the Macintosh's Programmer's Workshop, which is built out of standard Mac II toolbox routines.

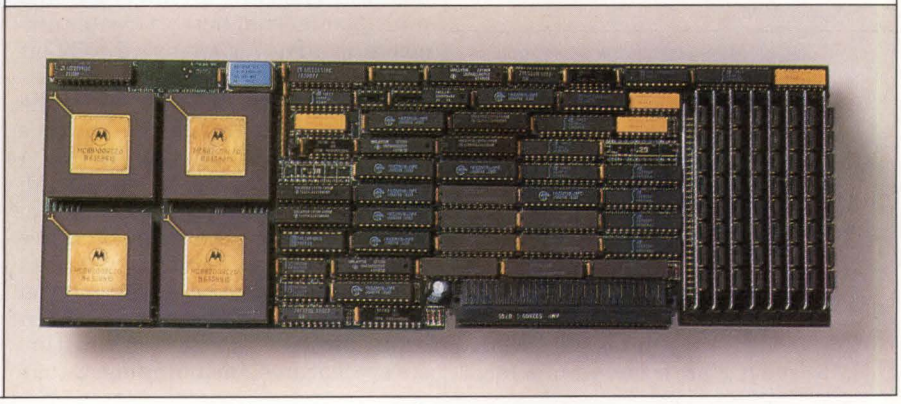

DESIGN AND DEVELOPMENT TOOLS

The new software tools include a set of four full libraries of code used for application development. The Client Library, running on the 8100 processor, provides resources to handle requirements such as interprocessor communication, memory management, and character and file I/O. A Host Library, which runs on the 68020, is used to connect the user interface to an application that takes advantage of both the 88100 and the 68020 processor. These two libraries function as a pair to facilitate communication between the two processors, and they're supported by two additional libraries. The Standard Input/Output (STDIO) Library contains Unix interface functions such as "read," "write" and "seek," which simplify moving applications from other computers to the development board. A standard Math Library is also available.

A variety of sample programs guide developers using the system through steps for designing their own RISC applications. These tools demonstrate interprocessor usage and can be used as basic templates for creating custom applications. The Host88 tool, for example, is a generic program that runs on the 68020 but calls on the 88100 to execute tasks. A second program, Client88, runs on the 88100 but calls on the 68020 to carry out operations. An "incorporate" menu command lets users knit the two programs together. A fully configured TL88K-P Development Board is priced at $14,995, with lower-priced configurations available from $10,995.

Tektronix Advance Technologies
PO Box 500
Beaverton, OR 97077
Circle number 127

Trademark Information

UNIX is a registered trademark of AT&T Bell Laboratories.

PAL and PALASM are registered trademarks of Monolithic Memories Inc.

GEOMETRY ENGINE is a trademark of Silicon Graphics Inc.

386, 386SX, 376, Intel386, iRMK, iRMX, and ICE are trademarks of Intel.

25-MHz microcontroller yields 6.4-Mips performance

Texas Instruments and Microchip Technology have jointly disclosed development of the 320C14, a 16-bit DSP microcontroller that combines DSP speed with the on-chip peripherals of a microcontroller. Operating at 25.6 MHz, the device features an instruction cycle time of 160 ns, allowing the execution of multiplication operations in a single instruction cycle while yielding a peak performance of 6.4 Mips.

The DSP engine of the 320C14 provides analog designers with a digital solution to servo-control applications through advanced control algorithms such as adaptive control, Kalman filtering and state controllers.

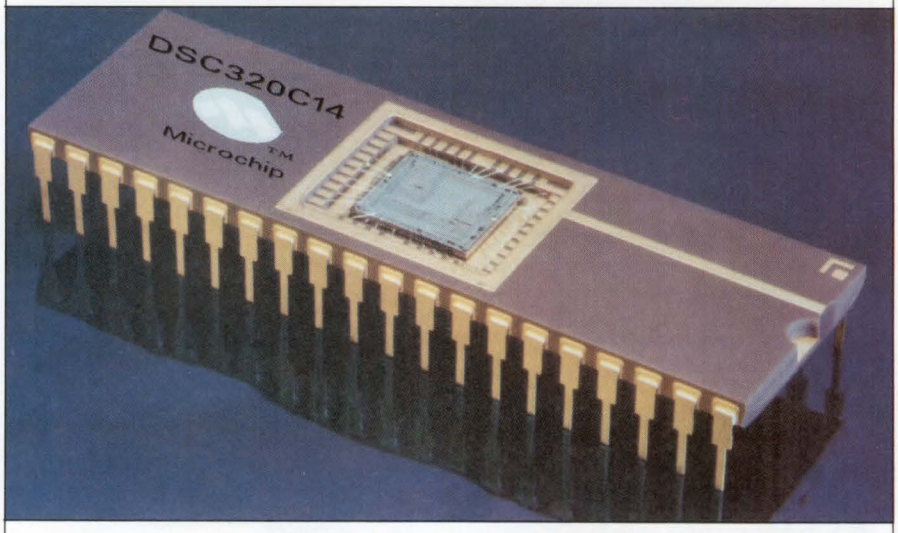

The high speed of the device is made possible through the implementation of a 16-×16-bit multiplier in hardware, as well as through the use of a Harvard architecture and multiple internal buses. Greater precision is achieved by using a 32-bit ALU and 32-bit registers.

In addition to the DSP CPU, the microcontroller provides 256 kwords (512 kbytes) of on-chip RAM, 4 kwords (8 kbytes) of on-chip ROM, and the ability to address 4 kwords of off-chip memory. An array of active on-chip peripheral and I/O functions work independently, effectively reducing the frequency of CPU interrupts. Four 16-bit incrementing timers are featured, including one watchdog timer used to prevent software hang-ups, two general-purpose timer/counters that can be configured to produce one 40-bit timer, and one baud-rate generator.

The on-chip event manager consists of two subsystems: a capture system that provides a direct optical encoder interface with 160-ns detection resolution, and a compare system that features six pulse-width-modulated output channels with 10-bit resolution at 20 kHz and 8-bit resolution at 80 kHz. The serial port can operate in synchronous mode (6.4 MHz), asynchronous mode (400 kHz) and codec mode for interfacing to analog systems.

Software support for the 320C14 includes an assembler, a linker and a simulator, while a full in-circuit emulator provides hardware support. An EPROM version facilitates prototyping. The 68-pin device will cost $75 for engineering evaluation units, with production quantities available in mid-1989 at less than $10.

Texas Instruments
PO Box 809066
Dallas, TX 75380
Circle number 128

Microchip Technology
2355 W Chandler Blvd
Chandler, AZ 85224
Circle number 129

Coming November 14th!

THE NEW NEWSPAPER FOR YOU AND THE REST OF YOUR TEAM MEMBERS. . .

. . .in test, manufacturing, quality control and procurement management.

NEWS EDITION

COMPUTER DESIGN®

A PennWell Publication | THE ONLY NEWSPAPER FOR THE SYSTEMS TEAM | November 14, 1988

"We've given design engineering managers their technology magazine; now, the systems team has its own newspaper."

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

Computer Design asked its readers which buses they're using in their current applications. Not suprisingly, VMEbus is way out in front, and it looks like it will stay there.

Page 2

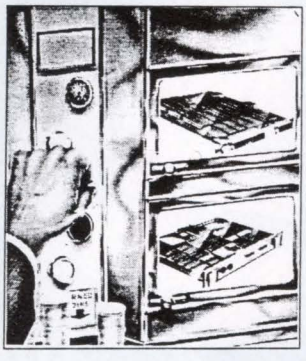

Computer Design hosts soire at Buscon East

The doors of the famous Horn and Hardart Automat (200 East 42nd Street) are being opened for exhibitors at Buscon East on Tuesday evening, from 6 to 10 p.m. Don't miss this chance to see a New York landmark and ...et some key members of ...n's editorial,

System teams get their own newspaper

Barbara-Ann Scofidio

LITTLETON, MA—PennWell Publishing Company's Advanced Technology Group will introduce next month the first newspaper to focus on what the company describes as "the systems product design and development team." *Computer Design News* will explore the issues and problems common to members of the design, test, manufacturing, purchasing and corporate sectors of computer and electronics companies. Its target readership are the individuals directly responsible for the cost and integrity of new products, and for expediting their movement into the market.

Computer Design and *Computer Design News* together will cover the full range of performance, quality, cost and time-to-market considerations that impact the entire product design and development team. The 26-year-old *Computer Design* magazine, targeted at senior systems engineers and design engineering managers, will be distributed monthly, beginning in January, 1989. According to a study com...ed by PennWell Publishing, an ...ority of readers be distributed monthly to *Computer Design's* present domestic circulation, bringing the combined 1989 frequency to 24 times per year. Beginning in January, *Computer Design News* will also be distributed to a qualified circulation of systems procurement team members at the major OEMs in the electronics/computer industry.

"The real bet-the-store decisions that are being made in companies today in terms of directions and vendors are increasingly becoming joint decisions incorporating inputs from quality-control, manufacturing, procurement management and engineering mana ment," says David L. Allen, publisher of *Computer Design* and its news edition. "We're going to wrap our arms around the systems community so that these people who are, by definition, encouraged and forced to work together, will have a common base of information on vendor viability, product availability, design direction and technical directions in the marketplace." *Computer Design News* will provide all of the pertinent information generated in the industry on major system components and how this affects a ...

Allen's predictions. Although the new newspaper was officially announced only two weeks ago, and no prototype is available to show potential advertisers, the charter issue has received commitment

"The real bet-the-store dec are being made in companie increasingly becoming join incorporating inputs from

DESIGN AND DEVELOPMENT TOOLS

PC-based system supports commercial in-circuit programming

The Boardsite 4100 and 4400 in-circuit programming systems connect to an IBM PC or compatible, and they let users reprogram PROMs, EPROMs and EEPROMS without removing the devices from the circuit board and without damaging surrounding devices. The 4100 can program up to eight boards; the 4400 can handle from eight to 32 boards simultaneously. Both are available in benchtop and portable models.

Until recently, most in-circuit programming systems were custom designed for military and aerospace applications that needed the reliability of soldered-in parts while maintaining design flexibility. The Boardsite systems are designed specifically for commercial use and can be configured and reconfigured to support a number of boards and ICs. The obvious advantage to in-circuit programming is reduced parts handling, which subsequently holds down labor costs and component damage. Further advantages include reduced labeling requirements for discrete components and the elimination of the need for sockets.

Both models connect via a high-speed expansion bus to a PC, which is responsible for data-flow control, disk storage and processing speed. The programmer will run as fast as the system to which it's attached, up to 16 MHz. A simple point-to-point, connector-to-connector interface attaches the circuit boards to the programmer.

Menu-driven board-profile development software is included to assist the user in creating profiles to display the address locations, bus width and data sources necessary to program the board. Prices on the model 4100 start at $9,500, while the 4400 costs $14,500.

Data I/O
10525 Willows Rd NE
Redmond, WA 98073

Circle number 130

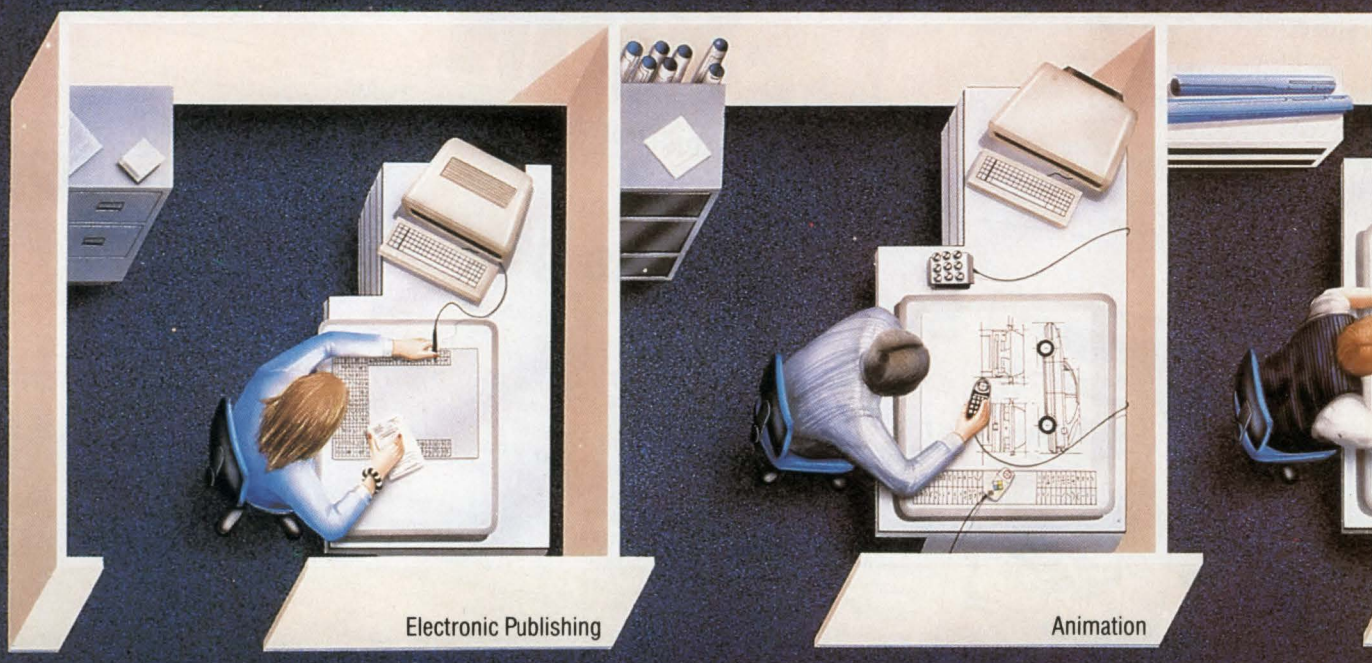

Microgrid® II. A solution

Why are more people specifying Microgrid II Series tablets for their applications? Simple. It's the best value.

The Microgrid II Series is competitively priced, yet it comes with several standard features (like a universal power supply and dual RS232 interface) that cost over $1,000 with other tablets.

Microgrid is compatible with the world's largest base of installed systems, so it can easily fit into most new or existing workstation requirements. You also

First 2½-in. Winchester drive offers 20-Mbyte capacity

Laptop computer designers are always on the lookout for ways to pack more performance into a small, rugged, lightweight package. A critical component of such design considerations has always been the disk drive, which is power hungry, sensitive to environmental stress and expensive. With the introduction of the industry's first 2½-in. 20-Mbyte Winchester drive, these restrictions seem to have been addressed. The Prairie 220 is a ruggedized disk drive that weighs a scant 9.6 oz, consumes only an average of 1.5 W of power and takes up only 12 in.^{3}, roughly one-half the space required by its nearest competitor.

The drive uses a ramp mechanism for loading and unloading the heads during power downs and periods of inactivity, a design scheme that provides several benefits. Since the heads don't land on the disk surface, the number of disk starts and stops has been increased to 250,000. In comparison, drives with landing zones on the disk surface are generally limited to 10,000 start/stops before significant media damage and head contamination occurs. Another benefit of the off-disk head placement is the drive's ability to withstand up to a 100-g force at 11 ms in a power-off mode. Head slap (the bouncing of head against disk surface) and stiction (adhesion of the head to the disk surface) is also eliminated by the ramp-load design.

Several strategies have been used to reduce power consumption. The ramp-loaded heads allow the use of a smaller spin motor, since no head/disk friction has to be overcome during power up. In addition, the integration of the drive electronics and the drive controller results in further power reductions compared to a drive using a nonintegrated controller.

The drive has three basic operating modes: active, power savings and standby. In the active mode, the device consumes 1.7 W in the ready state, 2.6 W when seeking and approximately 3 W when reading or writing. In the power-savings mode, the drive remains spinning with heads loaded and a major portion of the electronics powered down. This mode yields a power consumption of 1.4 W. The third mode, standby mode is a spun-down, deselected mode in which only the interface is alive. This mode has a power consumption of 0.1 W. All three operating modes are programmable depending on customer preference.

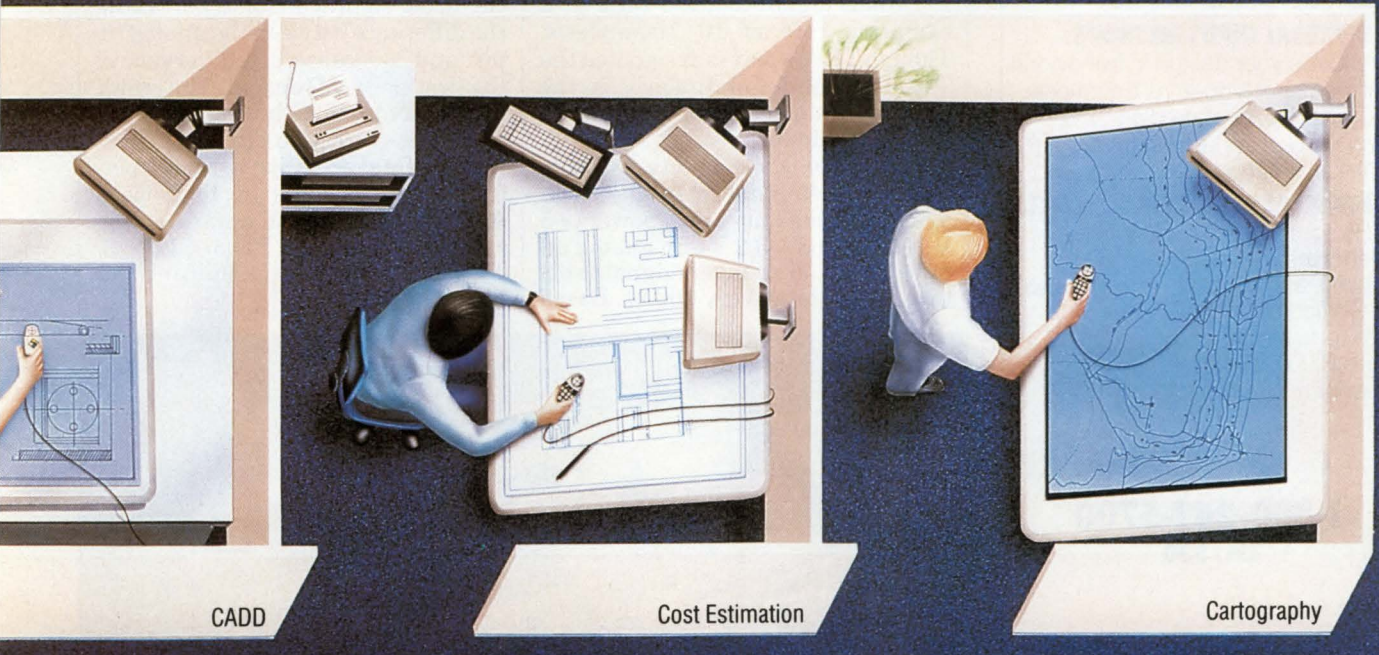

to fit any application.

get a choice of active work areas from 17" x 24" to 42" x 60"; high resolution; superior accuracy; a choice of manual or power lift stands, service plans and more.

Why pay more when you simply can't get more than Microgrid II. For more information and the name of your nearest distributor call our toll-free number.

CIRCLE NO. 39

Summagraphics® CORPORATION
777 State Street Ext., Fairfield, CT 06430
203-384-1344 800-221-9244
©1987 Summagraphics Corporation. All rights reserved.
See Us in COMDEX Booth #4224

ABEL: THE UNIVERSAL CHOICE FOR PLD DESIGN.

More design engineers choose ABEL™, Data I/O®'s high-level design software. Because this universal, industry-standard tool offers more PLD choices.

UNIVERSAL DEVICE SUPPORT. ABEL supports virtually every PAL®, PLA, FPLA, EPLD, FPGA, and GAL™ on the market, including both simple and complex architectures.

UNIVERSAL VENDOR COVERAGE. ABEL supports 100% of the semiconductor manufacturers.

UNIVERSAL INPUT METHODS. Describe your design in behavioral terms using high-level equations, truth tables, or state diagrams—whichever works best for the particular circuit.

UNIVERSAL PLATFORM SUPPORT. ABEL runs on multiple hardware platforms including Sun, PC, Apollo, and VAX.

To find out more, call today for our free PLD Development Kit.

1-800-247-5700
Ext. 330

DATA I/O
Corporation

CIRCLE NO. 42

■ **NEW PRODUCT HIGHLIGHTS**

MEMORY SYSTEMS

The drive electronics require no adjustments or calibrations, since an on-board microprocessor in the intelligent controller performs power-up diagnostics, compensates for the analog offsets of the drive electronics, and performs its own defect management during disk formatting and certification. Controller options include an intelligent IBM PC AT-compatible or a standard SCSI controller using the common command set.

A fully balanced rotary actuator and a servo system provide a 28-ms average seek time and use a bronze-filled nylon coil center in tandem with a proprietary zero-torque-flex circuit to minimize loading factors during operation. The actuator also assists in maintaining the heads in the detented position when they're not in use. The Prairie 220 is available in limited production quantities. The unit price is under $400 in large OEM quantities.

Prairie Tek
2120 Miller Dr
Longmont, CO 80501
Circle number 131

MAJOR SYSTEM COMPONENTS

Twelve-bit A-D converter pushes performance to 10 Msamples/s

The Comlinear-TRW alliance that was announced last May has produced an analog-to-digital converter that sets new standards in high-speed, high-resolution operation. Built with a synergistic merging of proprietary technologies, the converter provides 12-bit resolution and a +70-dB signal-to-noise ratio at a sampling rate of 10 Msamples/s. Though TRW's part is designated the THC1202 and Comlinear's is the CLC925, both are form, fit and function equivalents.

The device is a complete A-D converter in a 40-pin dual in-line package, requiring only apply power (+5 V, +15 V and −5.2 V), an analog input signal and a convert signal to obtain 12-bit data every 100 ns. Based on the subranging A-D architecture, the converter uses two high-speed, medium-resolution A-D converters and a fast, high-accuracy digital-to-analog converter to make two passes at converting the signal. The first pass digitizes the most significant bits, while the second pass digitizes the least significant bits by taking the difference between the analog input and a reconstructed version of the first-pass A-D output. The result is then sent to a 12-bit digital adder and error-correction circuitry.

In the final analysis, two basic performance measures of high-speed A-D converters stand out: sampling rate and accuracy. The CLC925/THC1202 performs admirably in both categories.

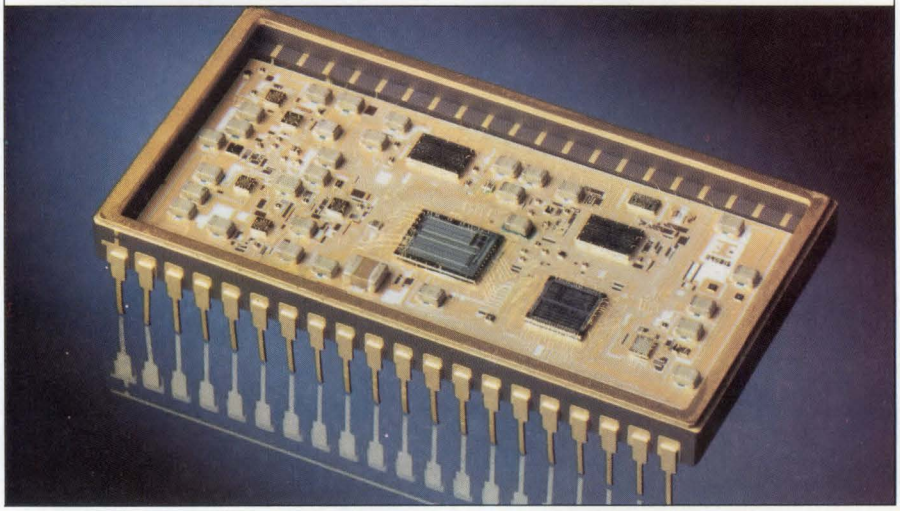

MAJOR SYSTEM COMPONENTS

Both the dc differential linearity error and the integral linearity error are ½ least significant bit. This level of accuracy is especially important in radar applications, since it's required in detecting and identifying low radar cross-section (stealth) aircraft. A similar need is present in medical imaging, where low distortion translates into better discrimination of tissue anomalies.

The converter also offers features that both improve performance and speed up the design process. A gain adjust pin, for example, lets the voltage gain of the circuitry be adjusted over a range of ±10 percent of the full-scale range, saving external circuitry and facilitating dynamic calibration of the user's system. Both Comlinear and TRW are offering evaluation boards for the converter. In sampling quantities of one to four, the A-D converter is priced at $1,120. In production quantities of 100 to 499, the price of the device is $875.

Comlinear
4800 Wheaton Dr
Ft Collins, CO 80525
Circle number 132

TRW LSI Products
PO Box 2472
La Jolla, CA 92038
Circle number 133

TESTING AND MANUFACTURING

Digital oscilloscopes feature 20-GHz sampling speed

Faced with ever-increasing processor speeds and higher component integration, manufacturers of today's fast digital and analog systems are using more sophisicated interconnect methods, such as balanced transmission lines, to provide precise timing, better noise immunity and more stable transmission characteristics. Testing and characterizing these systems, however, is proving to be a challenge. With the introduction of its 11800 Series of digital oscilloscopes, Tektronix is striving to address some of these design problems.

Boasting a 20-GHz sampling rate, a 17.5-ps or less rise time and an 8-bit vertical resolution, the scopes use differential time-domain reflectometry (TDR) for the evaluation and testing of high-speed computer components and ICs such as ECL, gallium arsenide, and fast bipolar application-specific ICs and VHSICs. The series includes the 11801 eight-channel mainframe, the 11802 four-channel mainframe, the SD-24 TDR/Dual-Channel Sampling Head, the SD-26 Dual-Channel Sampling Head plug-in unit and the SM11 Multi-Channel unit. A 9-GHz passive probe, the P6150, is also available.

Both mainframes can be configured with the sampling head plug-in unit to execute single-ended and differential TDR measurements. With these measurements, high-speed backplanes, balanced cables and circuit board runs can be evaluated automatically.

The eight-channel model can be configured with the multichannel unit to acquire and measure up to 136 channels for flexibility when designing high-speed computers and ICs. Both the four-channel and eight-channel mainframes provide a comprehensive measurement system with 16 continually updated automatic time and amplitude measurements. Mean and standard deviation statistics are provided for all measurements.

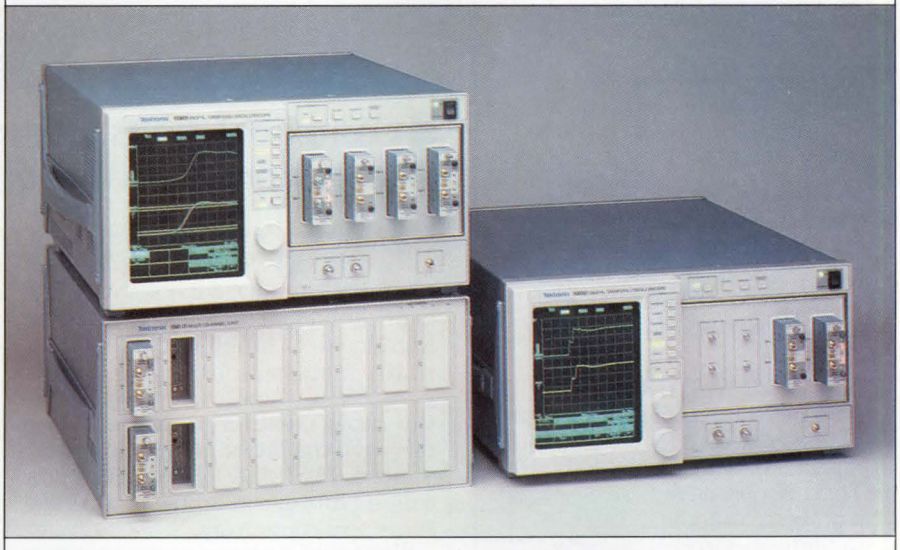

"With differential TDR testing, high-speed computer manufacturers and designers can detect impedance mismatches in backplanes and circuit board transmission lines that would otherwise lead to signal loss and would degrade system performance," says Bob Bousquet, Tektronix' product marketing manager for the 11800 series. "These mismatches can't be detected with single-ended TDR measurement techniques." Differential TDR requires that matched positive and negative pulses be sent, simultaneously, down a balanced line. Two of these pulses, each with adjustable amplitude and polarity, can be sent by the SD-24 dual-channel head, and either pulse can be inverted for differential TDR.

Up to 16 automatic measurements, including propagation delay and other pulse parameters, can be se-

FULLY APPROVED UNIVERSITY DEGREES Economical home study for Bachelor's, Master's, Ph.D., *fully approved* by California State Department of Education, not accredited. Prestigious faculty counsels for independent study and life experience credits (5000 enrolled students, 400 faculty). *Free information*—Richard Crews, M.D. (Harvard), President, Columbia Pacific University, Department 3E5N, 1415 Third Street, San Rafael, CA 94901. Toll free: (800) 227-0119; California: (800) 552-5522; or (415) 459-1650.

GATES: THE ADVANCED CHOICE FOR MULTIPLE PLD DESIGN.

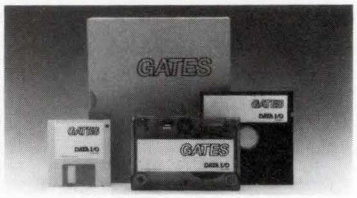

For complex PLD designs, choose GATES™ Logic Synthesizer, Data I/O®'s most advanced PLD design software.

ADVANCED DEVICE INDEPENDENT SPECIFICATION. The behavioral design specification in GATES gives you more flexibility by deferring device selection and implementation until the end of the design process.

ADVANCED INTERACTIVE ENVIRONMENT. GATES gives you immediate feedback during every step of the design process.

ADVANCED LOGIC SYNTHESIS. Use GATES factoring to combine common terms, and use smaller gates, so you can take advantage of the feedback and multi-level logic capabilities of the latest devices.

ADVANCED PARTITIONING AND DEVICE SELECTION. Design without concern for the size or number of devices needed. Then, simply select a group of outputs to partition; GATES does the rest.

To find out more, call today for our free PLD Development Kit.

1-800-247-5700
Ext. 331

DATA I/O
Corporation

CIRCLE NO. 43

TESTING AND MANUFACTURING

lected, and six of them can be displayed via touch-screen menus. Multiple windowing lets engineers simultaneously display a waveform and up to seven high-resolution portions for precise comparison. Using the Autoset feature, a user can set up screen settings and triggering to display a signal. Mean and standard deviations for all measurements are displayed to verify stability and to indicate measurement repeatability. These statistical results, together with all displayed measurement results, are continuously updated to reflect changing signal conditions and to provide an accurate picture.

The multichannel unit accepts up to 16 sampling heads for an additional 32 acquisition channels. With Option 1M, the 11801 can drive up to four of these units through a single general-purpose interface bus address, providing a total of 136 channels. One-half of all available channels can be displayed simultaneously. A hardware measurement mode lets timing measurements be made faster than in software, according to Tektronix. In this mode, it's not necessary to acquire and digitize entire signals. Instead, a Saved Trace Description feature lets unique sets of measurements be made on many signals without displaying their traces, improving throughput.

The scope's modular design lets customers adapt to future measurement and application requirements by adding sampling head plug-ins as needed. Higher bandwidth measurements can be performed, for instance, when higher bandwidth plug-ins are introduced. The prices for the series are 11801, $23,500; 11802, $22,000; SM11 (four-head), $10,000; SM11 with Option 1 (16-head), $15,000; SD-24, $5,000; and SD-26, $3,500. Delivery is two to four weeks.

Tektronix
PO Box 500
Beaverton, OR 97077
Circle number 134

SOFTWARE

Compiler brings C language to custom microprogramming

The Metastep Microprogram C Compiler lets the developer of custom microprogram-based systems write application and systems programs in high-level C language, as well as in the Metastep microprogram language. The compiler makes possible the rapid creation of microcode for use with the latest generation of advanced VLSI and DSP devices now coming into use in high-end workstations and computers.

"In the past, the microprogrammer was confronted with a difficult problem," says Darrell Wilburn, president and founder of Step Engineering, the compiler's developer. "Since the microprogram-based system was, by definition, a custom architecture, it was difficult to create a language that produced customized and optimized microcode for that architecture. Now we've added the ability to create instructions using high-level C to our Metastep standard."

The product consists of an advanced, optimized C language front-end compiler joined with an augmented version of the Metastep Microprogram Language. The C compiler front-end consumes the C source program and does all the processing of a standard compiler (analyzing the syntax, discovering errors, filling arguments), and then makes machine-dependent or -independent optimizations.

Machine-independent optimizations include recognition of expression reuse in the body of the code, strength reduction, common subexpression elimination and constants management. The Metastep back-end does most of the machine-dependent optimization. Prices range from $4,995 for a single-user MS-DOS configuration to $9,995 and up for Sun/Unix workstations and $19,995 and up for VAX/Unix installations.

Step Engineering
661 E Arques Ave
Sunnyvale, CA 94086
Circle number 148

we never stop asking "What if..."

TDR Yesterday

TDR Today

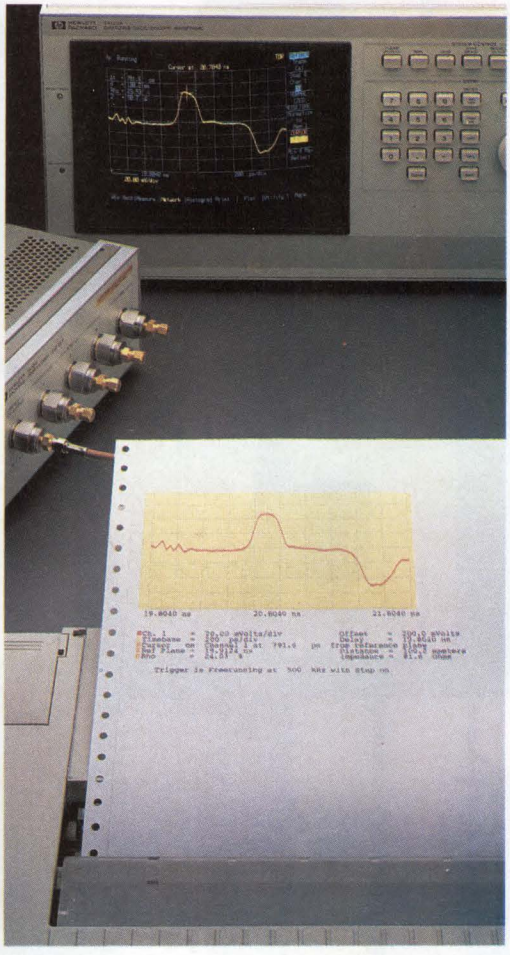

The HP 54120T: precise TDR measurements never before possible.

Throw away the calculators and slide rules. Toss out the CRT camera, the pencils and the scratch paper. Enter an unprecedented realm of TDR measurement that is powerful, fast, accurate, and easy to use. The HP 54120T Digitizing Oscilloscope and TDR with 20 GHz bandwidth and 10 ps resolution will make your life a lot easier.

Match bandwidth to system requirements.

The software we developed for the HP 54120T enables you to accurately match bandwidth to the speed of the device under test. Normalization* then corrects for reflections and imperfections caused by connectors and cables. You measure the device—not the connecting hardware—for error-free readings.

In addition, the HP 54120T features unmatched speed and convenience. HP Auto-Scale gets waveforms to the screen quickly, automatic measurements produce test results fast, English-like commands via HP-IB make programming easy, time and voltage markers customize measurement parameters, hardcopy documentation is available at the push of a button. And there are numerous other features we simply can't get into a single ad.

Call HP today! 1-800-752-0900, Ext. F786.

Application notes are yours for the asking. They include "Improving TDR Network Analysis," "TDR Fundamentals," and other technical data. You'll see for yourself what the HP 54120T can do for your design, characterization and test applications.

**Normalization is accomplished using the Stanford Bracewell Transform. The Bracewell Transform is under license from Stanford University.*

© 1988 Hewlett-Packard Co. EI15803/CD

HEWLETT PACKARD

THE BEST EVER H STILL INS

NEC is now delivering 1.4 GB of storage capacity in a 9″ disk drive.
All design engineers seem to want the same thing these days. Storage devices with larger and larger storage capacities. Is that too much to ask?

NEC doesn't think so. That's why the NEC 9″ disk drive now has up to 1.4 GB of storage capacity. That's more than any other 9″ disk drive on the market. But NEC has even more for you than raw capacity. Our 9″ drive now has a 50,000 hour

C&C Computers and Communications

IDEA WE AD IS TORAGE.

MTBF rate. And that makes it the most reliable disk drive in the industry.

Yet all the capacity and reliability in the world isn't going to help you get at your information more quickly. Which is why our 9″ drive has an average seek time of 15 ms with a data transfer rate of up to 3.07 MB/sec.

The NEC 9″ disk drive with 1.4 GB of storage capacity. It's the best news to ever come out of storage. For more information, call us at 1-800-343-4418.

NEC DISK DRIVES.
RELIABILITY
RUNS IN THE FAMILY.

NEC

NEC Information Systems, Dept. 1610, 1414 Massachusetts Ave., Boxborough, MA 01719.

A share of the market still down the road for GaAs ICs

Sydney F. Shapiro
Senior Editor, Business and Research

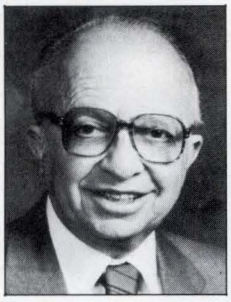

In the past year, gallium-arsenide technology and the applications for GaAs devices have matured only slightly. That doesn't necessarily mean that the future for GaAs in the IC industry is dim. But it does mean that GaAs still isn't a threat to the now-popular IC technologies.

The market for GaAs ICs has been continually troubled by high costs in processing and materials, as well as by the relative shortage of gallium. Excessive optimism regarding the "about-to-break-open" market for GaAs devices and the resultant build-up of production capacities that were left relatively unused have added to the problem.

Although predictions for the future of these ICs differ, the market for GaAs devices is beginning to develop. Integrated Circuit Engineering (Scottsdale, AZ), for example, finds that these ICs are finally "establishing a presence in the merchant IC market." That presence is expected to grow from a market of $100 million for a combination of device sales and development contracts in 1987 to $1.2 billion by 1992.

ICE bases its forecast on several factors, including advances in the technology and better production capacities. The development of new technologies also involves compatibility with current ones. For example, some GaAs IC producers are achieving GaAs transparency, by which their GaAs devices will be compatible with silicon ECL, TTL and CMOS gate arrays.

Improved performance is also key to the market success of GaAs ICs. ICE notes that lower power consumption may right now be more important than speed, although GaAs devices may be three to four times faster than silicon. But some major suppliers of GaAs gate arrays will offer products with up to 10 times less power consumption than silicon.

Still another important factor is that there are now enough competent digital GaAs IC suppliers with proven capabilities to produce the devices. Traditionally, new devices can't make a major impact on the market until second sources for the devices are available. The GaAs IC industry is no exception.

DOD controls the market

U.S. Department of Defense contracts, particularly those associated with the multimillion-dollar DARPA and MIMIC (Defense Advanced Research Projects Agency and Monolithic Millimeter and Microwave Integrated Circuit) programs, provide a market for over 80 percent of the world's merchant IC production. That should continue at least through 1992, when the commercial market will have matured enough to increase its percentage of the GaAs IC production. Estimated funding by the DOD for the MIMIC program alone amounts to about $225 million for material and technology development during the three-year period from mid-1988 to mid-1991, plus about $200 million for subsystem design and demonstration for the next three years to mid-1994. In addition, the DOD will spend another $100 million from mid-1988 to mid-1994 for technology support programs.

These and other DOD programs have financed much of the research and development that's been conducted by the captive GaAs manufacturers. These programs also helped finance some of the merchant GaAs vendors that split off from the captive manufacturers to market the devices to commercial companies. Now some of these captive manufacturers also offer their GaAs devices to the commercial market.

Many IC users still consider GaAs to be a material for the future.

The U.S. government has, in a sense, been financing GaAs device development for many years, although not necessarily for sophisticated ICs. A few companies have been supplying some forms of GaAs discrete devices to the DOD since the late 1960s. That financial backing kept the GaAs device industry alive and will continue to do so even though the prices for some GaAs ICs have fallen to low enough levels to make those devices viable commercial products.

Companies that develop GaAs devices invest tens of millions of dollars per year in the technology. Despite this, and despite some pessimism within the industry toward GaAs technology, the number of companies actively involved in GaAs device production has grown in the past five years.

In addition, the Japanese are now interested in GaAs technology. Unlike in the United States, however, where a large portion of the GaAs IC production has been controlled by small start-up companies, most of the Japanese GaAs activity is handled by large companies with substantial financial backing. Several of these major Japanese companies are reportedly each spending from $50 to $100 million annually to develop GaAs VLSI chips.

Diverse market for GaAs products

The analog GaAs device industry is mostly controlled by old-line companies, while most of the digital GaAs devices in the United States are developed and produced by start-ups. This tends, at least initially, to increase the levels of innovation because the start-ups aren't hindered by traditional methods of overcoming problems or by reticence toward initiating new ideas.

According to Henderson Ventures (Los Altos, CA), the market for monolithic GaAs ICs will increase at a 49 percent compound annual growth rate (CAGR) during the period from 1988 to 1996. Digital GaAs ICs lead analog GaAs ICs in the current market—$70 million for digital as opposed to $55 million for analog. Although the analog CAGR is predicted to be 49 percent by 1996, compared to 47 percent for the digital CAGR, the 1996 market for digital devices will be $1.487 billion, compared to a $1.365 billion market for analog devices. Analog GaAs ICs will be particularly strong in such applications as machine vision and robotics, as well as in the Strategic Defense Initiative and other space programs.

Optoelectronic GaAs ICs will experience even greater relative growth than other GaAs ICs during this eight-year period. A minor $4 million market in 1988 will grow at a 66 percent CAGR to $230 million by 1996, according to Henderson Ventures. This prediction is based on the integration of laser diodes with amplifying/modulating/signal-processing circuitry and the application of laser diodes in high-performance, fiberoptic LANs.

PERCENT OF TOTAL DIGITAL GAAS IC REVENUE

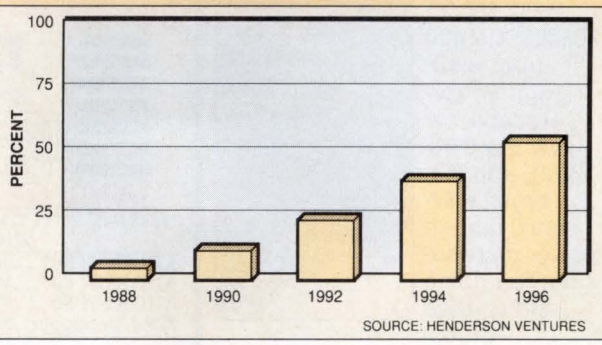

SOURCE: HENDERSON VENTURES

Gate arrays and other nonstandard logic devices are rapidly gaining importance in the total digital gallium-arsenide market. The market share for nonstandard digital logic monolithic GaAs merchant ICs is expected to increase from about a 3 percent share in 1988 to 53 percent in 1996.

Increase in GaAs gate arrays predicted

Gate arrays may well be the predominant GaAs products for 1988 and in the near future. Where there are too many production facilities for other GaAs devices, there may soon be a shortage of such facilities for gate arrays. In the worldwide market, bipolar ECL high-speed gate arrays are expected to increase at a 19 percent CAGR, from sales of $682 million in 1988 to $2.7 billion in 1996. For the same time period, sales of GaAs gate arrays will rise from a mere $2 million in 1988 to over $790 million in 1996, a CAGR of 111 percent.

With the increased use of GaAs gate arrays, a number of other problems will result, a major one involving the testing of these devices. Silicon-oriented automatic test equipment typically operates up to 100 MHz, which is only one-tenth of the actual GaAs gate array operating speed. Automatic test equipment capable of operating at such speeds is just beginning to appear. Another problem is packaging. The packages of 148 or more pins used in today's GaAs gate arrays are mostly custom products and, therefore, are expensive. To support lower cost GaAs gate arrays, lower cost packages will have to be made available along with the standard-priced packages.

Enhancement/depletion mode MESFET planar technology is expected to dominate GaAs gate arrays, at least for many years into the future. Although enhancement/depletion mode MESFETs were introduced in 1987, they aren't expected to make an actual impact on the market until about 1990, as is typical for many GaAs applications. A possible next-generation process is based on GaAs heterostructure IC technology.

The conflict between the current standard of silicon and the GaAs challenger is likely to continue for some time. GaAs is more radiation-resistant than silicon, which gives it an advantage for DOD projects. It's also a better optoelectronic electron emitter. The electron mobility of GaAs is about 5.7 times higher than that of silicon, and its electron velocity is twice as high. But the thermal conductivity property of silicon is more than three times better than that of GaAs material, which led to the introduction in 1986 of GaAs on silicon.

GaAs on silicon wasn't the ideal solution, however, because of the very slow and costly molecular beam epitaxy procedure used to deposit the layer of GaAs on the silicon. Efforts to solve this new problem led to still other difficulties caused by the differences in the thermal coefficient of expansion (TCE) between the two materials, which often results in warpage across the surface of the wafer. GaAs on sapphire, now being researched for some military applications, may offer an answer because the TCE mismatch is only one-third of that for GaAs and silicon.

A material for the future

Whatever advantages GaAs may have over silicon or over GaAs on silicon or on sapphire, or on some other material, many IC users still consider GaAs to be a material for the future. GaAs ICs will remain niche-type products for at least the next few years unless major breakthroughs occur in the technology. Their high cost, in part the result of low production yields, simply prohibits their use in applications other than military ones. And until manufacturers have adequate orders to warrant mass production, they aren't likely to be able to offer reliable supplies of cost-effective devices. □

For more information . . .

Market data referred to in this column was obtained from the following sources:

Status 1988, A Report on the Integrated Circuit Industry, $350. Integrated Circuit Engineering, 15022 N 75th St, Scottsdale, AZ 85260-2476.

Henderson Electronic Market Forecast, $750 annual subscription. Henderson Ventures, 101 First St, Suite 444, Los Altos, CA 94022.

Icecap Report, $395 annual subscription. Integrated Circuit Engineering, 15022 N 75th St, Scottsdale, AZ 85260-2476.

SYSTEM SHOWCASE

6-LAYER WIRE-WRAP SOCKET BOARDS 9U×400 mm; UHD; High Performance; for VME & SUN VME applications; PGA areas; unique power distribution; low noise and voltage drops.

Send for product data sheet

Hybricon

12 Willow Road, Ayer, MA 01432
TEL. EAST COAST: (508) 772-5422
WEST COAST: (602) 921-1824

CIRCLE NO. 221

RUN INTEL ON YOUR PC

DEVELOP...and...DEBUG

8 and 16 Bit INTEL CODE ON **YOUR PC** WITH **MICRUN**

MICRUN 80* .... 6 MHZ-8 Bit	$ 595
MICRUN 80* ....10 MHZ 8 Bit	995
MICRUN 86 .....16 Bit Development	450
8080 Debugger ..Intel Type Debugger	195

* Includes Hitachi 64180 co-processor

•Micrun 80/86 Package $ 895

Includes: Micrun 80(6mhz), 86 & 8080 Debugger

Call: 1-800-888-8086

Micro Interfaces Corp. 16359 NW 57th Ave., Miami, FL. 33014

Intel and Micrun are registered trademarks of their respective companies.

CIRCLE NO. 222

SOFTWARE VERSION CONTROL

The **POLYTRON** Version Control System (PVCS) provides precise & flexible configuration management for software development projects on Personal Computers, PC Lans and VAX minicomputers, including:

- Storage and retrieval of multiple revisions of source code.
- Maintenance of separate lines of development using branching.
- Merging of simultaneous changes.
- Modules can be reviewed by their own revision number, system version number, or specified date.
- Efficient disk storage. PVCS uses a very intelligent difference detection technique that minimizes the amount of disk space required to store a new version.

Personal PVCS — Offers most of the power and flexibility of Corporate PVCS. $149

Corporate PVCS — Offers features to maintain source code of very large and complex projects that may involve multiple programmers. $395

Network PVCS — Extends Corporate PVCS to use on networks. File locking and security levels can be tailored for each project. Call (503) 645-1150 for pricing on licenses for more than 5 stations. $995

VAX PVCS — Call for pricing

Requires DOS 2.0 or higher. Compatible with the IBM PC, XT, AT and other MS-DOS PC's. Works with any language.

TO ORDER: VISA/MC **1-800-547-4000.**
Dept. CPD Oregon and outside US call (503) 645-1150.
Send checks, P.O.'s to: **POLYTRON** Corporation, 1700 NW 167th Place, Dept. CPD, Beaverton, OR 97006.

CIRCLE NO. 223

IT'S EASY TO LOCATE ROM SOFTWARE SUPPORT

PC-LOCATE: Produce ROM-able code from your "exe" files. PC-LOCATE assigns physical addresses to the re-locatable image based on user inputs. PC-LOCATE supports the entire Intel processor family including the 8086, 8088, 80186, 80188 and 80286.

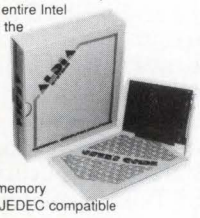

PC-PROMPAK: A PROM/ROM expansion board for IBM and IBM-compatible computers. PC-PROMPAK provides up to 384 Kbytes of expansion memory and supports most 28-pin JEDEC compatible devices.

PC-ROMDRIVE: Create a "Diskless PC" that can include MS-DOS and your application program. Program execution can be automatically invoked through the use of an "autoexec.bat" file.

ALDIA SYSTEMS INC

P.O. BOX 37634 PHOENIX, AZ. 85069
(602) 866-1786

CIRCLE NO. 224

HEWLETT-PACKARD TO IBM PC FILE COPY

IBM PC: PS/2, IBM AT, Vectra can read & write HP ASCII, Binary or Series 80 floppies. No additional hardware. Runs in the IBM PC, etc.

IBM PC FILE COPY $495

HP Series 80 PROGRAMS TRANSLATED TO IBM PC, etc.

Programs written on HP 85, 86, 87 translated to run on IBM PC, PS/2 & compatibles in Microsoft QuickBASIC 4.0. Includes file copy utility.

HP Series 80 to Microsoft BASIC 4.0 Translator . $495

CAD FILES A SPECIALTY

Our software treats CAD files like any other file. Experience with VersaCAD, Accugraph, Holguin & AutoCAD. We do custom file copy & program translation.

OSWEGO SOFTWARE, Inc.
507 North Adams Street, Oswego, Illinois, 60543, U.S.A.
Call us.
312/554-3567, FAX 312/554-3573, TELEX 858 757

CIRCLE NO. 225

CUSTOM DESIGN FOR CONTROL COMPUTERS

Why put this expensive specialist on your payroll when we already have the support you need? Vesta specializes in the design and production of control computers for OEM's.

- CREDIBILITY — We have **thousands** of installed systems
- SUPPORT — **Your** application engineer is as close as your phone
- ECONOMICAL — Your break even quantity may be as low as **50 units**
- FAST — Typically 3 to 4 months from ideas to fully functional **product**
- PROGRAMMING — From BIOS to application level
- PRODUCTION — If you don't want to, **we do**

Let us help you get your product into the market, ahead of your competition and at a reduced cost. Call us to discuss your requirements. Our prompt quote will make you happy you did.

VESTA TECHNOLOGY, INC.
7100 W. 44th Ave.•Suite 101•Wheatridge, CO 80033
(303)422-8088

CIRCLE NO. 226

IEEE-488 FOR PERSONAL COMPUTERS

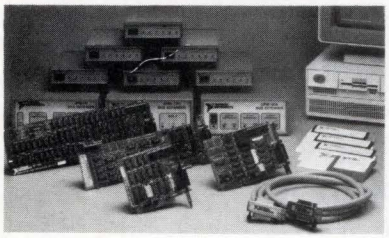

- IBM PC, XT, AT and PS/2
- DMA transfer rates up to 1M bytes/sec
- Industry's best and most extensive software
- Converters, Controllers, Extenders, Analyzer, and Buffer
- Free technical support and phone call
- 2-year warranty

NATIONAL INSTRUMENTS™
The Software is the Instrument™
12109 Technology Blvd.
Austin, Texas 78727-6024

Call for FREE Catalog
800-531-4742
512-250-9119

CIRCLE NO. 227

CCD LINE SCAN CAMERA

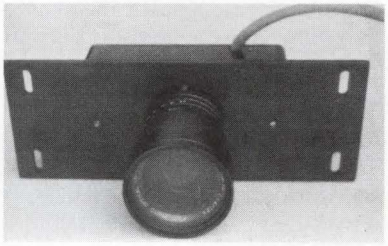

- 256, 512, or 1024 photo-elements
- Uses Standard "C" Mount Lenses
- Video rates to 2 MHz
- Rugged 8 x 3 inch Mounting Plate
- Low power—60 mA at +12V & –12V
- $800 for 512 element camera
- Lenses and Filters available
- VME Interface Board available

PiXoR Electronics
Box 2624 Roseville, CA 95746
916-791-2822

CIRCLE NO. 228

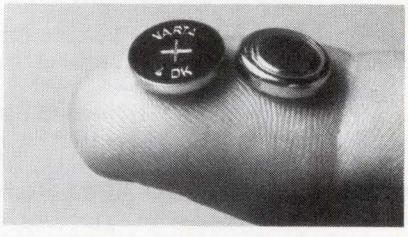

New for micro circuitry: the 4 DK, world's smallest rechargeable cell

Varta's new 4 DK—a tiny 4 mAh, 1.2 v NiCd cell—enhances miniaturization of a wide variety of applications, such as memory back-up, pulsing requirements, calculators and watches. Each cell is just 9.5x2.1 mm.

Built with Varta's unique mass-plate design, the 4 DK offers reduced charge requirements and superior charge retention. Only 400 micro amps are required for a full charge in 14-16 hours; or it can be trickle-charged continuously at 40 micro amps. The 4 DK retains over 40% charge after one year at room temperature. 4 DK cells can be assembled in battery packs yielding up to 12v. Contact Varta Batteries, Inc., 300 Executive Blvd., Elmsford, NY 10523, 914-592-2500, ext. 260.

VARTA

CIRCLE NO. 229

Circle 220 for rates and information, or call: Shirley Lessard or Barbara Kovalchek 800-225-0556. In Massachusetts 508-486-9501.

ADVERTISERS INDEX

ADVERTISER	PAGE NO.	CIRCLE NO.
Advanced Micro Devices, Inc.	84	30
	19	8
Advanced Processing Labs	81	27
Aldia Systems, Inc	114	224
Anritsu America, Inc	69	21
Applied Microsystems	10	6
Automated Systems, Inc	29	13
CalComp	CV3	40
Central Data Corp	17	8
Chips & Technologies, Inc	51-53	18
Control Data Corp	33-40	15
Daisy Systems Corp	14, 15	
Data I/O	106	42
	108	43
Electronic Solutions	1	1
Emulex Corp	CV4	41
Excelan, Inc	22	11
Excelan, Inc	46	16
GE Plastics	20, 21	10
GE Solid State	44, 45	
Heurikon Corporation	31	14
Hewlett-Packard	67	20
Hewlett-Packard	109	
Hitachi America, Ltd	70, 71	22
Hybricon Corp	114	221
ICI Image Data	4-9	5
Inmos Ltd	82, 83	29
Intel Corp	63-65	19
Mentor Graphics Corp	CV2	
†Mercanter International	100	36
Micro Case	78	25, 26
Micro Interfaces Corp	114	222
Microtec Research	2	2
Microware Systems Inc	97	34
Mitsubishi Electronics America	98, 99	35
*MultiDynamics, Inc	100	37
National Instruments	114	227
National Semiconductor Corp	26, 27	12
NEC Information Systems	110, 111	
Oswego Software	114	225
Paracom	95	32
†Philips	103	38
Pixor Electronics	114	228
Polytron	114	223
Quantum Corp	76, 77	24
Radstone Technology	72	23
Reliability Incorporated	96	33
Ross Dove, Co	81	28
Standard Microsystems	89-92	31
The Stepstone Corp	13	7
Summagraphics Corp	104, 105	39
Varta Batteries, Inc	114	229
Vesta Technology, Inc	114	226

*Domestic Issues Only †International Issues Only

The Advertisers' Index is published as a service. The publisher does not assume liability for errors or omission.

COMPUTER DESIGN

THE FIRST MAGAZINE OF SYSTEM DESIGN, DEVELOPMENT AND INTEGRATION

Executive Office: One Technology Park Drive, P.O. Box 990, Westford, MA 01886
(508) 692-0700 FAX: (508) 692-0525 Telex: 883436

Publisher, David L. Allen (508) 392-2111
Associate Publisher, John C. Miklosz (508) 392-2114
Marketing Services Manager, Mary M. Gregory (508) 392-2110
Sales Promotion and Research Manager, Judy Leger (508) 392-2113
Administrative Assistant, Tina LePage (508) 392-2112
Advertising Coordinator, Adrienne Adler (508) 392-2198
Postcards/System Showcase/Recruitment, Shirley Athorn (508) 392-2183
Barbara Kovalchek (508) 392-2184
Circulation Director, Bob Dromgoole (508) 392-2173
Assistant Circulation Manager, George Andrew (508) 392-2171
List Rental, Bruce McCarthy (508) 392-2176

SALES OFFICES

New England/Upstate New York/ Florida/E. Canada
Kevin Callahan
P.O. Box 605
North Scituate, MA 02060
Tel: (617) 545-6603
Fax: (617) 545-7821

New York/New Jersey/ Pennsylvania/Delaware/Maryland/ North-South Carolina/Georgia
Neil Versen
Park 80 West, Plaza Two
Saddle Brook, NJ 07662
Tel: (201) 845-0800
Fax: (201) 845-6275

Upper Midwest/Ontario, Canada
Robert D. Wentz
9501 West Devon, Suite 203
Rosemont, IL 60018
Tel: (312) 518-9191
Fax: (312) 698-2675

South Central
Eric Jeter
3050 Post Oak Blvd., Suite 200
Houston, TX 77056
Tel: (713) 621-9720
Fax: (713) 963-6285

Southern California/Southwest
Tom Boris, Greg Cruse
2082 SE Bristol, Suite 216
Santa Ana, CA 92707
Tel: (714) 756-0681
Fax: (714) 756-0621

Northern California/NV
Tom Boris, John Sly, Bill Cooper
1000 Elwell Court, Suite 234
Palo Alto, CA 94303
Tel: (415) 965-4334
Fax: (415) 965-0255

Oregon
Tom Boris
2082 SE Bristol, Suite 216
Santa Ana, CA 92707
Tel: (714) 756-0681
Fax: (714) 756-0621

Washington
John Sly
1000 Elwell Court, Suite 234
Palo Alto, CA 94303
Tel: (415) 965-4334
Fax: (415) 965-0255

U.K./Scandinavia
David Round
69 Imperial Way
Croydon
Surrey CRO 4RR, England
Tel: 01 686 7655 Telex: 938420
Fax: 01 688 2134

France/Belgium/S. Switzerland/ Spain/The Netherlands
Daniel R. Bernard
247, Rue Saint Jacques
75005 Paris France
Tel: (1) 43 54 55 35
Telex: 214235F MISIVZ
Attn: Missitex: PENNWELL
Fax (1) 4707 59 01

W. Germany/Austria/ N. Switzerland/Eastern Europe
Johann Bylek
Verlagsbuero Johann Bylek
Stockaeckerring 63
D-8011 Kirchheim/Muenchen
Federal Republic of Germany
Tel: 089 903 88 06
Telex: 529355 vbb d

Italy
Luigi Rancati
Rancati Advertising
Milano San Felice Torre 5
20090 Segrate Italy
Tel: 2 7531445
Telex: 328601 RANCAD I
Fax: 02 7532354

Japan
Sumio Oka
International Media Representatives, Ltd.
2-29 Toranomon 1-chome
Minato-ku, Tokyo 105 Japan
Tel: 03-502-0656
Telex: J22633 MEDIAREP
Fax: 03-591-2530

Southeast Asia
Anne Goh-Taylor
Seavex, Ltd.
400 Orchard Rd.
10-01, Orchard Towers
Singapore 0923
Republic of Singapore
Tel: 734-9790
Telex: 35539 SEAVEX RS
Fax: 732-5129

Jay G. Seo
Yong-Jin Park
Doobee International, Ltd.
Center Building (Byulgwan)
1-11 Jeong-dong, Choong-ku
CPO Box 4557
Seoul, Korea
Tel: 776-2096
Telex: K27117 DOOBEES
Fax: 755-9860

LAST BITS LAST BITS LAST BITS LAST BITS LAST B

IBM improves connectivity with Intel's Fastpath

IBM (Armonk, NY) will expand the application capabilities of its ES/9370 and 4381 mainframes by offering Intel's Fastpath 9770 Connectivity Control Unit as an integrated feature. Designed to fit into the system mainframe, the Fastpath unit provides an industry-standard open-bus architecture for all System/370-class processors, enabling them to connect to devices found in multivendor or networking environments.

Cooperative effort focuses on next-generation design-verification tools

Teradyne (Boston, MA) has announced the formation of an Electronic Design Automation Group to design and develop next-generation design and test tools. The group is composed of Aida (Santa Clara, CA), Case Technology (Mountain View, CA) and Teradyne's Design and Test Automation Group (Boston, MA). Products available from the group will include hardware and software tools for design verification of application-specific IC-based systems, CAE tools for circuit board layout, and design tools for medium-scale systems using commercial microprocessors.

Partnership targets ASIC test equipment

A strategic partnership between Gould (Cupertino, CA) and Cadic (Beaverton, OR) will closely align the sales, marketing and service organizations, as well as future product-development efforts, of the two companies. According to the terms of the agreement, Gould's Test and Measurement Division will distribute and service all of Cadic's digital VLSI test systems throughout the United States and Europe. Both companies will jointly participate in the development of future application-specific IC verification test equipment for designers of digital systems.

Development pact spawns VGA BIOS

Phoenix Technologies (Norwood, MA) will offer an IBM PC-compatible BIOS to support the new Video Graphics Array (VGA) chips from Cirrus Logic (Milpitas, CA). Designed for the PC AT and PS/2 markets, the VGA chip set features full compatibility with all popular graphics standards and monitors, as well as with VGA's special features. The BIOS software resides in the CPU and manages the interface between software, such as word processing, spreadsheet and desktop publishing programs, and the system graphics hardware.

Agreement links design and test efforts

An agreement between Genrad (Concord, MA) and Diagonal Systems (Zurich, Switzerland) will provide for the marketing of software that links the design and test phases of circuit board production. The first CAD/CAE software-integration products developed by Diagonal are programs that run on any Daisy workstation and link to Genrad's board tester families. These links let test engineers use design, layout and simulation data captured during the development cycle in test program preparation.

Harris and Kawasaki become VLSIC partners

A joint technology agreement will bring Harris Semiconductor (Melbourne, FL) and Kawasaki Steel (Tokyo, Japan) together for the development and marketing of CMOS very large scale ICs. Under the terms of the pact, Harris will supply design, marketing and production capability to Kawasaki, and Kawasaki will provide financial and marketing support to Harris during the initial three-year term. In the early phases of the agreement, a number of circuits will be jointly defined and developed by both companies based on regional requirements. The resulting products will be marketed exclusively in the Pacific Rim by Kawasaki, while Harris Semiconductor will target the European and American markets.

Marketing pact unites data-acquisition efforts

Burr-Brown (Tuscon, AZ) has entered into a strategic marketing agreement with Heath/Zenith's Computer Based Instrument Group (St. Joseph, MI). The alliance combines Burr-Brown's personal computer instrumentation products with Heath/Zenith's computer-based instrument products and will target the data-acquisition markets with a hardware/software package. The joint effort is intended to satisfy end users' needs for a complete integration of data-acquisition hardware, personal computer hardware and application software.

Licensing and marketing alliance brings 68000 to Edge Computer

Motorola's Microprocessor Products Group (Austin, TX) has finalized an agreement with Edge Computer (Scottsdale, AZ) that will bring the 68000 microprocessor to the Edge 2000 product family. These companies will cooperate in product development to ensure full compatibility between Edge systems and the 68000 line, and they will pursue joint marketing and patent cross-licensing projects. The 2000 Series uses a mainframe multiprocessor architecture to deliver a sustained performance of between 16 and 55 Mips, with an average instruction time of 1.2 clock cycles per instruction at the processor level.

COMPUTER DESIGN (ISSN-0010-4566) is published twice monthly, except in July and December, by the Advanced Technology Group of PennWell Publishing Company, One Technology Park Drive, P.O. Box 991, Westford, MA 01886. Second class postage paid at Tulsa, OK 74112. COMPUTER DESIGN is distributed worldwide to qualified subscribers free of charge. Annual subscription price for non-qualified subscribers: $70 in the U.S.A., $110 elsewhere. Call (918) 831-9401 for a subscription application. 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.

© 1988 COMPUTER DESIGN by PennWell Publishing Company. All rights reserved. No material may be reprinted without permission from the publisher. Officers of PennWell Publishing Company: Philip C. Lauinger, Jr., Chairman and Chief Executive; Joseph A. Wolking, President; John Ford, Senior Vice-President; Carl J. Lawrence, Senior Vice-President; John Maney, Vice-President/Finance; Leslie P. Cypret, Vice-President/Administration.

The Graphics Engine That Meets Government Standards As Well As Your Own.

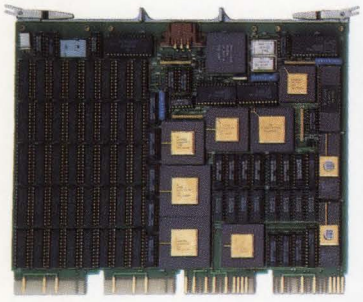

Now—your DEC MicroVAX or VME-based systems can meet a full range of critical government standards for high performance graphic workstations, thanks to three new versions of CalComp's hot Formula 1 graphics engine.

These single card Formula 1 solutions are engineered with the same "four on the floor" custom VLSI technology that runs up to four workstations from a single platform. With graphics superiority you can see at a glance.

Zoom, pan and update with turbo-charged speed and high-visibility 1280 x 1024 resolution. With on-board 32-bit floating point processor, 68020 command processor, display list and six custom VLSI gate arrays as standard equipment. Whatever your software interface is, CalComp has the solution ... graphic interfaces from CGI to GKS and language bindings—Ada, C and Fortran.

The powerful Q-bus Subsystem is designed both for the DEC MicroVAX and the new 3000 series. You get a full subsystem package including the Formula 1 card, enhanced for VMS and ULTRIX.

The VME Subsystem starts with an industry standard 6U160 VME card that supports X-Windows and UNIX. CalComp can supply you with VME modules or a total subsystem for integration into your system.

The Formula 1 meets your TEMPEST requirements, too. Choose a subsystem designed for DEC's TEMPEST MicroVAX II, or combine the graphics engine with CalComp's TEMPEST peripherals for your own unique solution. All are designed to meet NACSIM 5100A.

The Formula 1 will meet your standards for price and performance—without sticker shock. You can drive away with an entire Formula 1 Subsystem—card, high-resolution color monitor, keyboard and mouse—starting at $6995. Or choose the card alone, for as low as $3995. Attractive OEM discounts are available.

For more information, call: 1-800-CALCOMP. Or write: P.O. Box 3250, Anaheim, CA 92803.

We draw on your imagination.™

CalComp
A Lockheed Company

DEC and MICROVAX are registered trademarks of Digital Equipment Corporation. UNIX is a registered trademark of Bell Laboratories.
Formula 1 is a trademark of CalComp Inc.
© 1988 CALCOMP Inc.

CIRCLE NO. 40

DISK-PAC TO GO.

Removable Storage for Sun Microsystems.

Emulex has refreshing news for Sun users: We have removable storage. Our SN-ER2S Removable Disk Subsystem is fully compatible with Sun-2, Sun-3 or Sun-4 systems.

But that's not all. Onboard intelligence monitors critical drive functions, and senses movement in the Portable Data Modules (PDMs). This ensures optimum reliability. Should someone try to remove a module during operation, all activity is shut down. Previously entered data is fully protected, and heads are automatically locked down.

All of which gives you the capability of keeping important data locked up! You get all the security of removable storage, plus the convenience of portability.

You also get flexibility. PDMs are available with 180, 380 or 760 megabytes of unformatted storage. And since each subsystem is self-configuring, PDMs of differing capacities can be interchanged.

If you've been looking for a removable storage solution, we have a great one at hand. Pick one up on your way home!

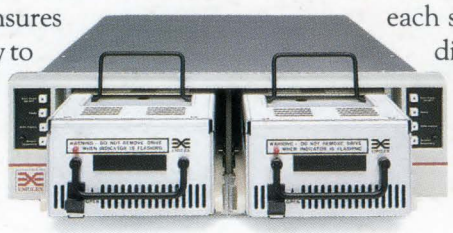

SN-ER2S Removable Disk Subsystem

Call (800) EMULEX-3, or
(714) 662-5600 in California
Emulex Corporation
3545 Harbor Blvd.
Costa Mesa, CA 92626

EMULEX

Regional Offices: Anaheim, CA (714) 385-1685 Atlanta, GA (404) 587-3610
Burlington, MA (617) 229-8880 Chicago, IL (312) 490-0050 Dublin, CA (415) 829-1170
Teaneck, NJ (201) 836-3717 Washington, DC (703) 264-0670

International Offices: Bracknell 344-484234 Munich 89-3608020 North Sydney 2-957-1669
Paris 134-65-9191 Tokyo 3-234-8951 Toronto (416) 673-1211

©1988 Emulex Corporation, Inc.

CIRCLE NO. 41