A PENNWELL PUBLICATION

COMPUTER DESIGN®

APRIL 1, 1989

THE FIRST MAGAZINE OF SYSTEM DESIGN, DEVELOPMENT AND INTEGRATION

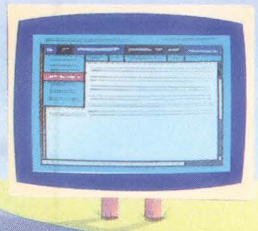

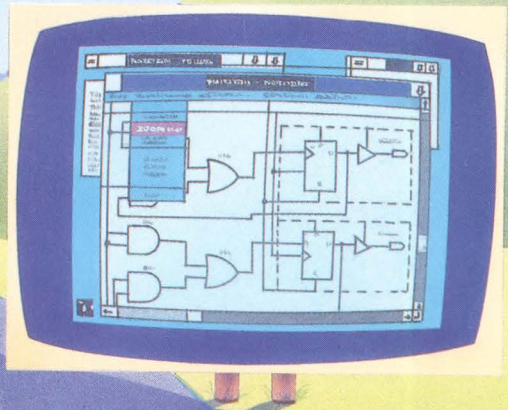

Software protocols smooth the path to a standard graphics interface

- 80860 CPU positions Intel to take on minisupers *Page 20*
- Emerging standards, hardware and software light the way to FDDI *Page 51*
- Static RAMs race to keep up with RISC *Page 59*

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

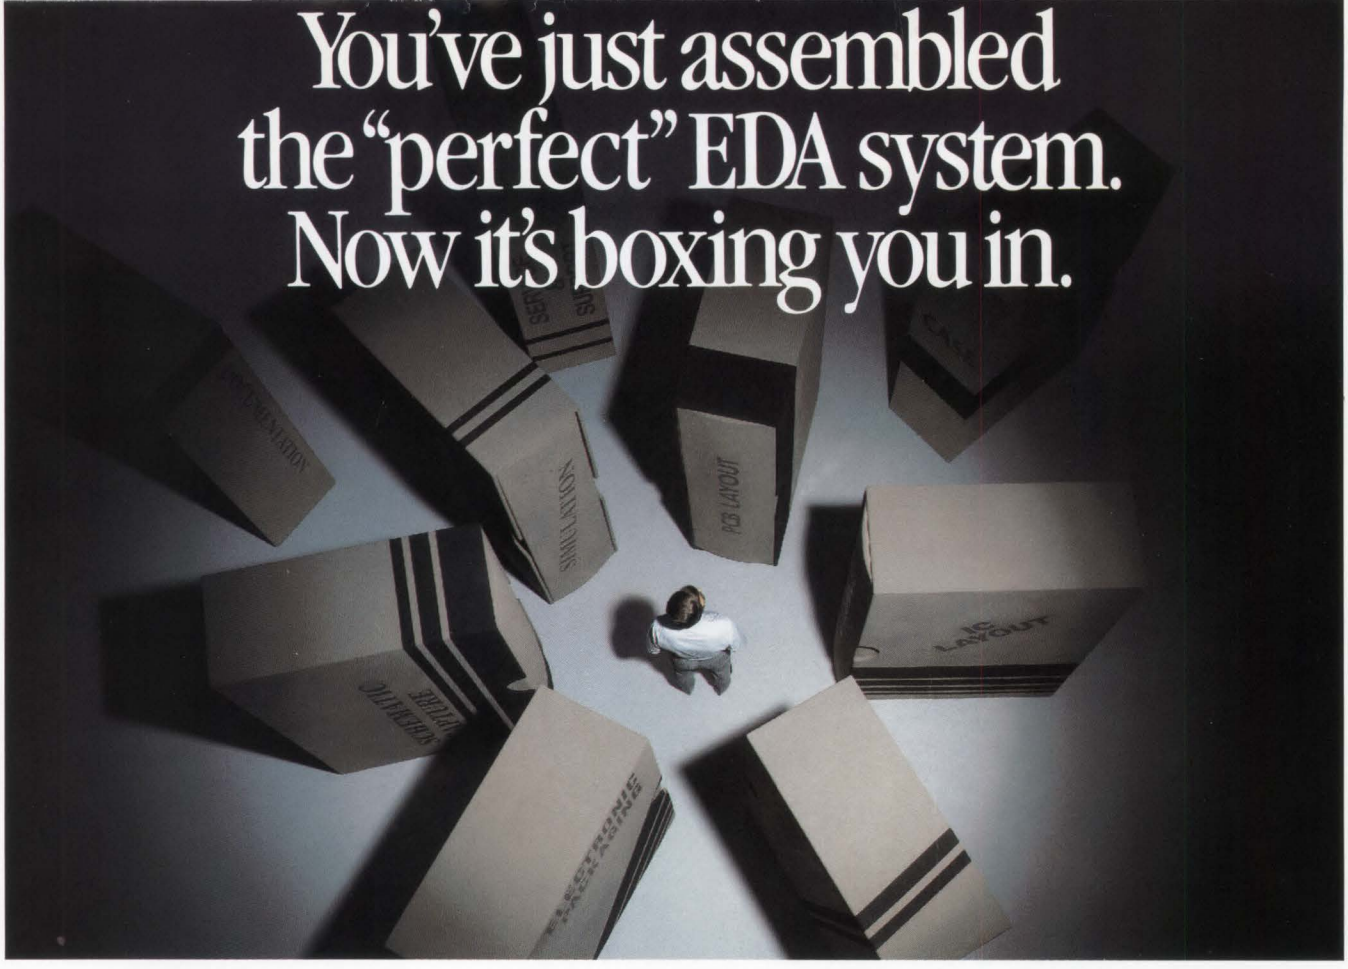

Mentor Graphics delivers the only Electronic Design Automation system that grows in size and power at the same time.

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.

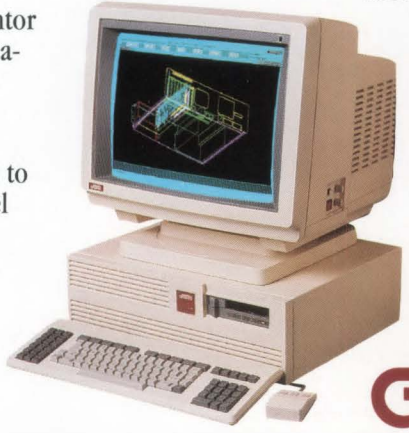

Mentor Graphics also provides a single source of service and support for your entire EDA tool set. No more finger pointing. No more delays. Our support organization takes care of all your hardware and software needs.

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.

Packaging Makes the Difference

If you want your system to look its very best, look to Electronic Solutions.

The first thing you'll find is that the beauty we add to your board level system is a lot more than skin deep. It will look better, sell better, and work better in an Electronic Solutions enclosure, for a lot of reasons:

Engineered Cooling for maximum performance from the new, high-density cards.

Built from Experience: We've been the enclosure experts since board-level systems began.

Safety and EMI/RFI Approvals by UL, CSA, TUV(IEC380) and the FCC.

IEC380 FCC

Off the Shelf or Built to Order: From simple needs like paint and labeling to complete reshaping for the most unusual requirements—we do it all. Most of the time we have a big head start.

See how packaging can make the big difference to your system. Give us a call today on our toll-free numbers.

Huge Selection of Complete Enclosures

Multibus I
- Tabletop 4,7,10,12,15 slots
- 3½" Rack 4 slots
- 7" Rack 7,10 slots
- 10½" Rack 12,15 slots
- DeskMate 7,10,12,15 slots

VME
- Tabletop 3,5,6,7,12, 20,32,40 slots
- 3½" Rack 3 slots
- 7" Rack 5,7 slots
- 10½" Rack 12 slots
- 14" Rack 12,20,32,40 slots
- DeskMate 5,7,12 slots

Multibus II
- Tabletop 6,7,12,20 slots
- 7" Rack 7 slots
- 14" Rack 12,20 slots

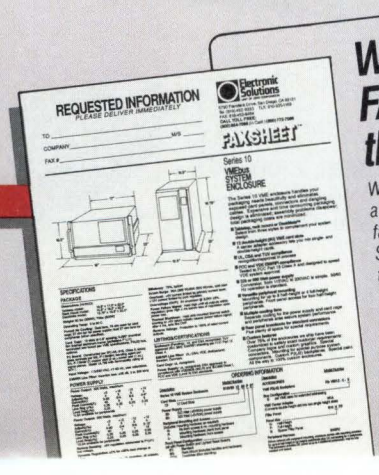

We'll FAX you the facts.

Want the latest data in a hurry? Nothing is faster than Electronic Solutions' new "FAX the FACTS" program. Call our "800" number and give us the information you need from us. We'll FAX it to you immediately.

ZERO **Electronic Solutions**
UNIT OF ZERO CORPORATION

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

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

CIRCLE NO. 2

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

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

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

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

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

RIMFIRE 3510
SCSI Adapter for
VMEbus

RIMFIRE 1500
SCSI Adapter for
Multibus® I

RIMFIRE 2500
SCSI Adapter for
Multibus® II

CIPRICO

2955 Xenium Lane
Plymouth, MN 55441

CIPRICO LISTENS. AND RESPONDS.

Multibus is a registered trademark of Intel Corporation.

CIRCLE NO. 3

VOL. 28, NO. 7 **COMPUTER DESIGN** APRIL 1, 1989

THE FIRST MAGAZINE OF SYSTEM DESIGN, DEVELOPMENT AND INTEGRATION

Page 44

TECHNOLOGY UPDATES

Integrated Circuits

80860 CPU positions Intel to take on minisupers 20

80860 blends vector processing into a RISC-like architecture..... 25

Computers and Subsystems

Languishing STD Bus receives a late-life kicker 28

SCSI-2 controller board builds parallel disk drive arrays 32

Major System Components

New twist in LCD designs makes them thin, light power-savers..... 37

Digital paper emerges as low-cost archival-storage option..... 41

Design and Development Tools

Niche simulators tackle special applications 44

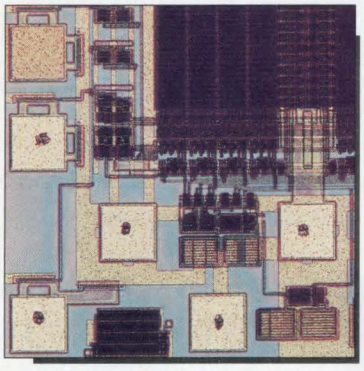

Page 59

TECHNOLOGY AND DESIGN FEATURES

Emerging standards, hardware and software light the way to FDDI

Within two years, FDDI will be established as the dominant high-end LAN for the 1990s..... 51

Static RAMs race to keep up with RISC

Cache memory is becoming an integral part of RISC-based computer systems. But these new architectures burden the memory system by calling for fresh data and instructions each cycle..... 59

COVER STORY

Software protocols smooth the path to a standard graphics interface

As designers steamroll toward a universal windows-based interface, PCs and high-powered graphics systems will converge on common ground..... 70

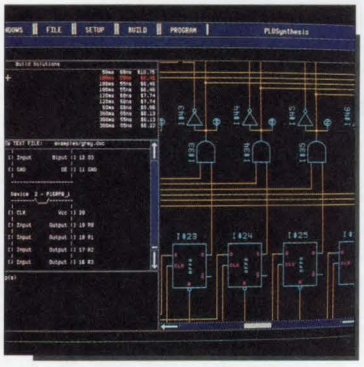

Page 100

NEW PRODUCT HIGHLIGHTS

Integrated Circuits

ICs bring monochrome and color image capture into single-chip solution 88

RISC processor posts burst speeds of 50 million operations/s..... 93

32-bit graphics microprocessor makes desktop peripherals 10 times faster 94

New two-chip set shrinks Micro Channel hard disk controllers 96

ICT unveils high-density programmable logic architecture 98

Design and Development Tools

Design tool selects PLD, partitions designs..... 100

Simulator accesses internal device registers 102

Hardware modeling system interfaces to any logic simulator..... 103

Computers and Subsystems

RISC coprocessor gives AT a boost.... 105

Intelligent NuBus I/O modules fit ruggedized Macintosh for factory use..... 106

VME SCSI adapter has four-bus architecture 107

DEPARTMENTS/COLUMNS

News Briefs 9
Editorial..... 14
Advertisers Index 112

National Semiconductor

Get the kit. Forget You'll move a whole

INTRODUCING THE ONE-CHIP, HIGH-PERFORMANCE SOLUTION TO FLOPPY-DISK CONTROL.

This new device is actually a complete floppy-disk controller kit in a single chip. It gives you everything you need—a precision analog data separator, pin-programmable write precompensation, and cable and bus interface buffers.

All you have to add is an address decoder, a crystal, and a few passive components.

Think of how quick and easy this kit will be to design with. The nitty-gritty work has already been done. And the DP8473 is fully compatible with the NEC μPD765A, the current industry standard. So all you have to do is drop it into your design. (And eliminate about 20 other chips.)

The DP8473 will control up to four floppy-disk drives in an IBM PC-XT,* PC-AT,* or any PS/2* model personal computer.

THE KIT.

National's one-chip solution to floppy-disk control.

AND THE CABOODLE.

Other solutions carry a lot of excess baggage.

the caboodle. lot quicker without it.

QUICKER DESIGNS, QUICKER PROFITS.

In addition to your tremendous savings in real estate, your savings in raw components, inventory, and assembly could add up to more than $10 a board. So you not only get to market sooner, you also make a quicker profit.

Even more important, the DP8473 gives you better performance than any other integrated floppy-disk controller solution on the market.

Unlike most competitors, its on-board data separator is analog rather than digital.

THE QUICKER WAY TO HIGH PERFORMANCE.

Because it has the continuous tracking ability of an analog phase-locked loop, the DP8473 is less sensitive to motor speed variation and bit shift. So it has a larger dynamic window margin.

In fact, its window margin is greater than any other integrated solution on the market—analog or digital. And, of course, the bigger your window margin, the better your data integrity. And the faster the data rates you can handle.

FLOPPY-DISK CONTROLLER PERFORMANCE

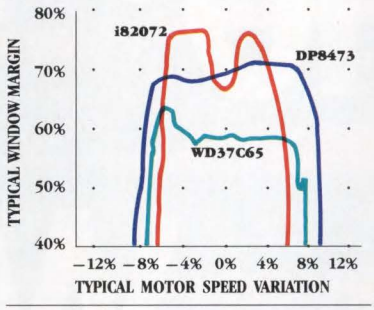

No other integrated data separator—analog or digital—can give you a 70 percent window margin at ±8 percent variation in motor speeds.

A QUICK PATH TO THE FUTURE.

The DP8473's analog data separator can handle transfer rates from 250k up to 1M bps. And disk drives with as many as 4,000 tracks.

You may not need that capacity right now, but the next generation of low-cost floppy-disk drives is already on the drawing table. And the DP8473 is ready to handle the data rates those drives will require.

It's also manufactured using low-power CMOS technology, so it's even ready for the next generation of battery-powered laptops. In fact, the DP8473 is the floppy-disk controller used in **Compaq's** new laptop computer, the SLT/286.

QUICK DELIVERY.

The DP8473 is in stock at your distributor now. So you can start taking advantage of its many benefits right away.

With the addition of this new floppy-disk controller, National's mass-storage chip set is now the most complete in the industry.

We can provide the ideal drive-control solution for every application—from PC to mainframe.

NATIONAL'S MASS STORAGE FAMILY

Floppy-Disk Data Controller
DP8473

Hard-Disk Data Controller
DP8466A

Data Synchronizers/Separators
DP8451 DP8455 DP8459
DP8461 DP8462 DP8465

Pulse Detectors
DP8464B DP8468B

2, 7 Encoder/Decoder
DP8463B

Async SCSI I/O
DP5380 DP8490

With data rates up to 25M bps. We also offer a complete line of hard-disk-drive support chips. For more information—quick—contact your National distributor. Or call us toll-free at 800-227-1817, ext. 17.

National Semiconductor

• ADVANCED • PERIPHERALS

©1988 National Semiconductor Corporation
IBM PC-AT, PC-AT and PS/2 are registered trademarks of International Business Machines, Inc. Compaq is a registered trademark of Compaq Computer Corporation.

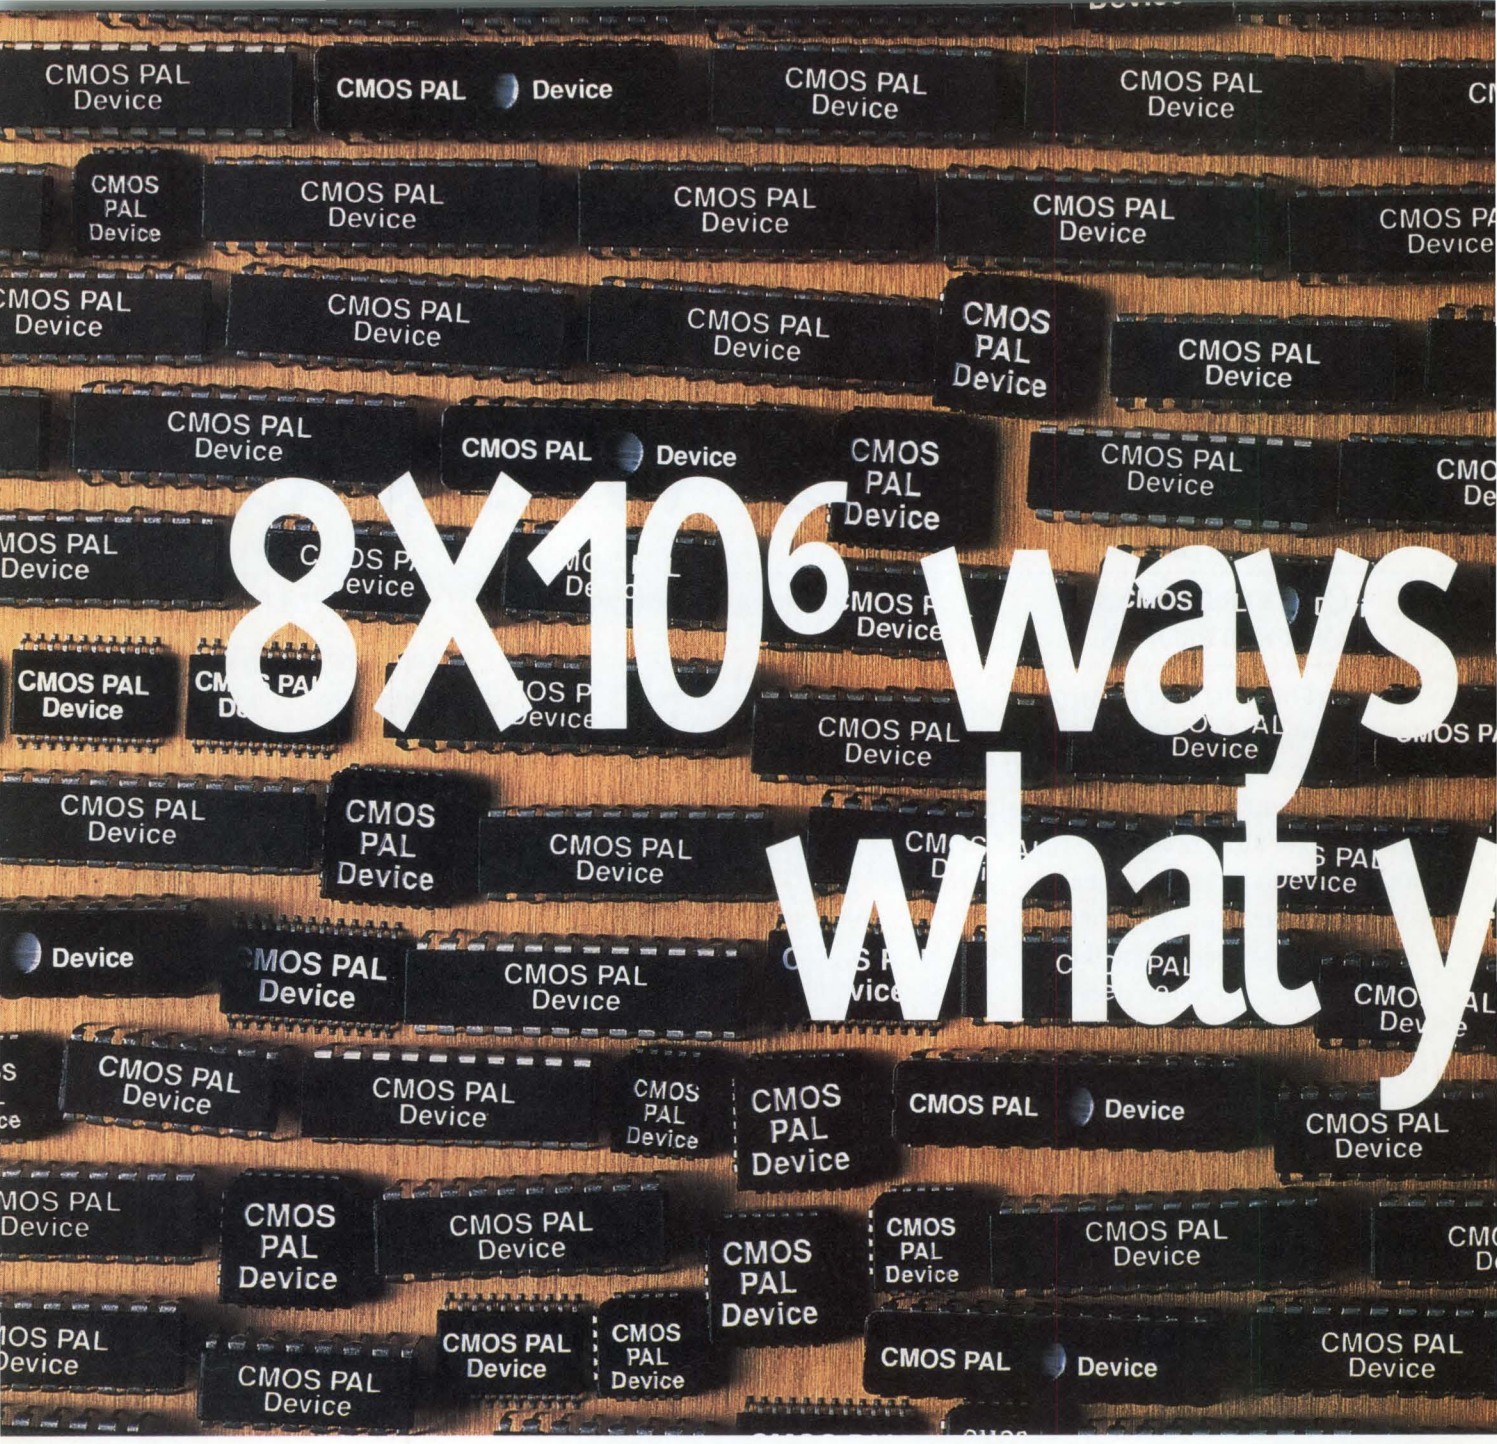

We have 8 CMOS PAL® chips. That totals up to 8 million different ways to program your logic. Quite enough of a choice to meet your most exacting low power logic needs. From universal PAL architectures to zero stand-by power PAL devices. You can have the system you want without compromise. And without costly delays.

Our all-purpose CMOS architectures give flexibility without draining precious power. The 24-pin industry standard PALC22V10 uses just 90mA—even when it's running at 25ns.

If you liked us in bipolar, you should see us in CMOS.

We have five CMOS ZPAL™ devices based on familiar bipolar architectures. You can use them in battery powered and low duty-cycle applications. And they use virtually no power when in stand-by. Two versions are electrically erasable. You can get the functions you're used to and you can change your design as fast as

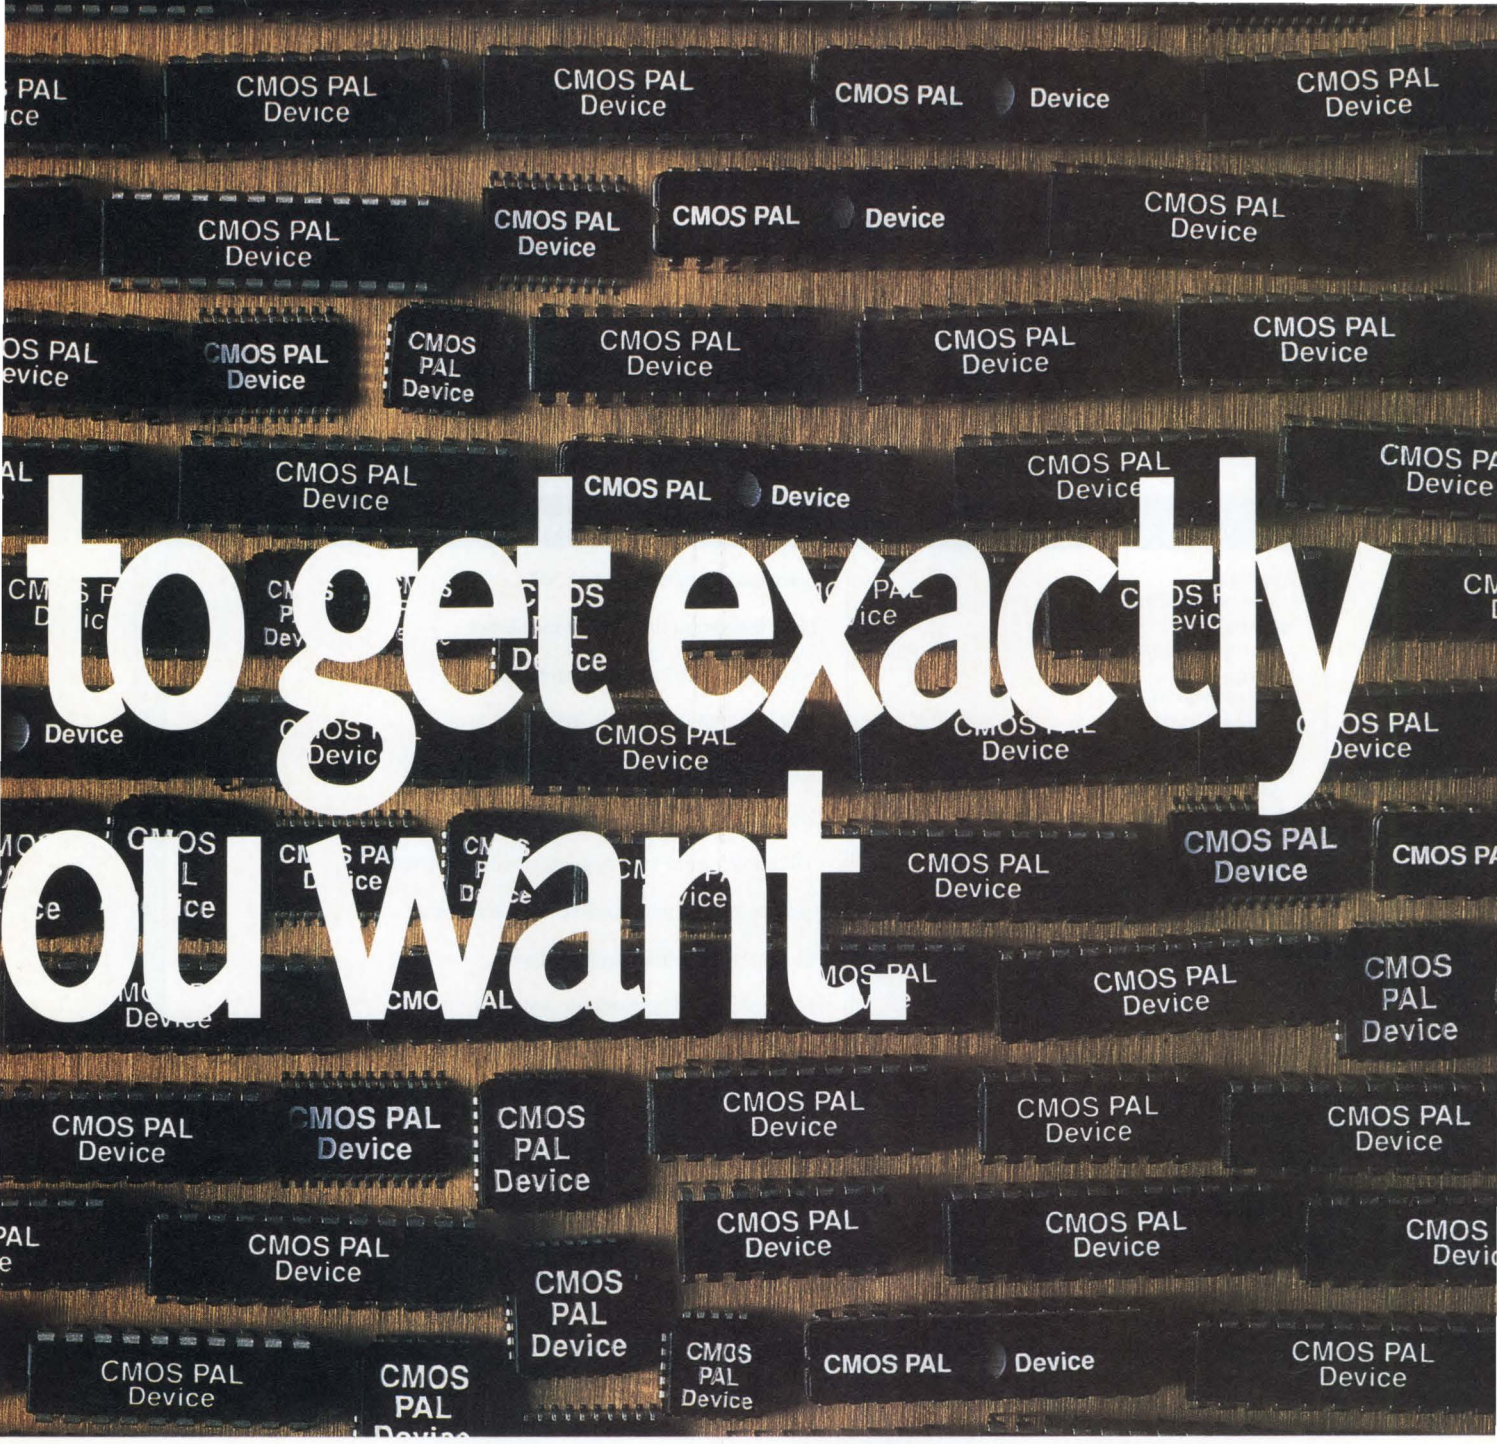

you change your mind.

What did you expect from the people who invented PAL devices?

The largest and most experienced Field Application Engineer force is as close as your phone all during your design-in process.

Our TestPro™ Centers give you access to superior testing and programming support. Just call your AMD sales office for the center nearest you.

And we've kept the classic PALASM® software up to date with our parts and your needs.

Most important, you can get your hands on our parts. They're all in high volume supplies.

Call (800) 222-9323. It's hard enough to get what you want in life without passing up 8,000,000 more chances.

Advanced Micro Devices

Monolithic Memories

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

PAL and PALASM are registered trademarks, TestPro and ZPAL are trademarks of Advanced Micro Devices, Inc. © 1989 Advanced Micro Devices, Inc.

CIRCLE NO. 5

VxWorks. More than Ready for Real-Time.

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

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

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

Only one real-time operating system satisfies these requirements.

VxWorks.

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

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

for VxWorks.

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

Become a Believer.

Wind River Systems, Inc.
1351 Ocean Avenue
Emeryville, CA 94608
CIRCLE NO. 6

NEWS BRIEFS NEWS BRIEFS NEWS BRIEFS NEWS BRIEF

Consortium to spark optical IC development

If a plan proposed by the David Sarnoff Research Center (Princeton, NJ) comes to pass, a U.S. fab will be producing optoelectronics ICs in volume by 1996. The proposed industry consortium, which requires about $155 million in funding, was presented to 30 chip and system vendors at the Commerce Department last month.

The proposal is designed to keep U.S. production in step with similar plans in Japan and Europe. If established, the consortium would coordinate an initial research and development phase aimed at integrating the optical transmitter and receiver into a silicon VLSI chip. A second, two-year phase would focus on prototype development, and a final phase would lead to manufacturing. —*John Mayer*

Arbitrator weakens AMD's hopes for 80386 rights

The battle between Advanced Micro Devices (Sunnyvale, CA) and Intel (Santa Clara, CA) over rights to Intel's 80386 microprocessor has moved a step closer to settlement with an initial ruling by arbitrator J. Barton Phelps.

In the dispute, AMD claims that terms of a 1982 technology exchange agreement give the company rights to the technology in Intel's 32-bit CPU, which AMD hoped to obtain in exchange for rights to the Quad Pixel Dataflow Manager (QPDM) chip. In the first phase of the disagreement under arbitration, Intel claimed that AMD had failed to meet its responsibilities under the agreement, thereby losing the right to 80386 data, and AMD claimed that Intel had failed to act in good faith. Both parties are claiming breach of contract.

In his initial ruling, Phelps found that AMD and Intel had in fact not reached an agreement to exchange the QPDM for the 80386 data. Further, Phelps found that Intel had failed to negotiate in good faith over the specifications of the QPDM. AMD, therefore, is entitled to some as-yet-undetermined damages, Phelps said. The judge's ruling included a statement that relief could be in the form of money, rather than manufacturing rights—apparently putting AMD further from its goal of obtaining rights to the 80386. A final settlement of the dispute isn't expected until the end of the year. —*Ron Wilson*

Joint effort yields seamless 1750A development environment

Four companies have teamed to produce an integrated environment for MIL-STD-1750A applications development. The 1750A 16-bit instruction set is mandated by the U.S. Air Force for airborne and spaceborne computer systems for avionics and control.

The catalyst for the joint effort is the Orbiter board from Sabtech Industries (Yorba Linda, CA). The Orbiter plugs into an AT-compatible platform and uses the two-chip 1750A implementation from LSI Logic (Milpitas, CA). Both board and chip set currently run at 20 MHz but are designed to go to 30 MHz. The Orbiter uses 2 Mbytes (1M 16-bit words) of dual-port RAM that can be mapped into a 64-kbyte window in PC address space. This lets the designer inject data to internal registers without the application under development being aware of it.

A complete set of real-time Ada development tools supplied by Ready Systems (Sunnyvale, CA) includes host-based CASE tools, an Ada compiler and debugger, a target-based real-time run-time environment and the RTAda/OS real-time operating system based on Ready Systems' VRTX real-time kernel. An 80386-based laptop computer from Grid Systems (Fremont, CA) includes a two-card snap-in tray to hold the Orbiter and a bus interface card. —*Tom Williams*

ISDN alliance to boost chip integration

With the integrated services digital network (ISDN) sitting on the brink of widespread implementation, Mitel (Kanata, Ontario) and International Microelectronics Products (San Jose, CA) have announced a joint technology exchange agreement that should not only broaden access to ISDN ICs, but also speed development of application-specific products.

Under the agreement, IMP will gain manufacturing rights and act as an alternate source for four Mitel chips. The devices include an S interface chip; a D-phone, or integrated digital telephone circuit; an H-phone, or integrated digital telephone circuit with high-level data-link control; and a crosspoint switch. Mitel benefits from access to IMP's ACL1.2-DCL1.2 cell libraries and its 1.2-micron CMOS process, which IMP claims is the industry's smallest mixed analog/digital process.

But the primary dividend coming from this alliance may lie in the development of future products. Once the basic products are entered as core cells in the library, both companies will be able to reconfigure the four Mitel chips. And both companies will codevelop macrocells to augment their design libraries for voice and data communications. "The real benefit will lie in our continuing cooperation in developing the macrocell library," says John Freeman, vice-president of marketing for Mitel's Semiconductor Division. —*John Mayer*

Analogy links Saber analog simulator with digital simulators

Analogy (Beaverton, OR) has reached agreements with Gateway Design Automation (Lowell, MA) and Genrad (Concord, MA) to link Analogy's Saber analog simulator with Gateway's Verilog and Genrad's Hilo digital simulators. Under the agreements, Saber will be linked with Verilog and Hilo in much the same way that it's currently linked with the Cadat digital simulator from HHB Systems (a subsidiary of Daisy Systems of Mountain View, CA).

These agreements help support Analogy's claim that Saber is the "emerging standard" among analog simulators. With links to Cadat, Verilog and Hilo, Saber now supports mixed-mode simulation with the digital simulators that control 90 percent of the stand-alone digital simulator market, according to

(continued on page 10)

(continued from page 9)

Doug Johnson, Analogy's vice-president of marketing.

Saber also provides analog simulation capability to a number of other companies, including Valid Logic Systems (through Valid's acquisition of Analog Design Tools, a Saber user), Racal-Redac, Computervision, NCR and Schlumberger.

—Bill Harding

GaAs op amp soars to 10 GHz

Engineers at Hughes Research Laboratories (Malibu, CA) have shattered previous op amp bandwidth specs with a gallium arsenide device that posts a 10-GHz unity gain bandwidth. Designed for use in high-speed data converters and line drivers, the amp shows a better than threefold performance improvement over existing devices.

Advances in MESFET processing let designers overcome typical GaAs problems of excessive backgating, low intrinsic transistor gain and light sensitivity. The new device uses a low-temperature buffer grown directly underneath the active channel to eliminate backgating and light sensitivity while improving intrinsic gain. Transistors on the chip are connected with air bridges to hold down overall circuit capacitance, and a single high-gain differential stage is used to maximize bandwidth. *—Mike Donlin*

Backplane bus daisy chains instruments

Although very similar to VMEbus, the MXIbus to be introduced this month by National Instruments (Austin, TX) differs in that it's multiplexed. MXIbus combines the GPIB form factor with the VME communication protocol—the reverse of VXIbus. But it extends VXIbus across multiple instruments and directly and transparently couples the VXIbus to industry-standard personal computers and other instruments that won't fit on a VXI module.

MXIbus has all of the advantages of an embedded VXIbus CPU, but at a much lower price, according to Ron Wolfe, strategic marketing manager at National Instruments. "MXIbus maps the hardware memory addresses of several systems together so that the individual systems think that they are one big system," he says. "And because the systems don't have to contend with communication protocol software, MXIbus obtains high performance at a relatively lower cost."

The first modules for this bus, not available before July, will consist of an MXI-VXI interface board and an AT-MXI interface board. Wolfe indicated that other bus interfaces can be expected in the future.

—Sydney F. Shapiro

On Technology and Nixdorf join Apollo's NCS bandwagon

Apollo Computer (Chelmsford, MA) has won two more converts for its Network Computing System (NCS). Nixdorf Computer (Waltham, MA) and On Technology (Cambridge, MA) are the latest in a large group of companies (including IBM, Digital Equipment Corp and Hewlett-Packard) to adopt what Apollo calls "the standard for creating open multivendor computer networks."

Apollo has offered NCS to the Open Software Foundation (OSF) as a base technology for the OSF open software environment. If chosen, the NCS influence would be widespread.

—Bill Harding

Neural net software lowers entry costs

Many of those who would like to explore the possibilities of neural networking have been put off by the high entry costs involved. Now, however, neural net pioneer Hecht-Nielsen Neurocomputers (San Diego, CA) has broadened the potential user base by introducing neurocomputing software packages that run on the omnipresent Sun family of workstations.

Although the Sun hardware is unlikely to run many neurocomputing applications optimally without the assist of a special-purpose coprocessor, the relatively low entry costs of a software-only solution should help justify a lot of new development activity and bring many programmers up to speed.

—David Lieberman

SBX PC comes to life

While it's not at all unusual when a board company builds a PC clone, a new entry from Zendex (Dublin, CA) breaks new ground by putting its PC on a little SBX module. The SBX, a daughterboard interface originally developed by Intel for Multibus I, has found long-term success among a number of buses, most particularly for prototyping purposes and easy customization of single-board computers.

The concept of the SBX PC has been around for several years. Today's technology makes it feasible, and today's market makes it worth implementing. The rampant activity in PC clone chips has given designers the needed functionality per square inch, while advances in application-specific ICs, packaging and board technology have further expanded their potential to overcome real-estate constraints.

—David Lieberman

Growth slows in office and computing equipment

Office and computing equipment sales in the United States grew by 7.6 percent—the lowest growth since 1972—according to figures released by market watcher In-Stat (Scottsdale, AZ). If it continues, the reduction in growth could have profound effects in the industry.

On the positive side, office and computing equipment vendors are major consumers of dynamic RAMs, and a softening in the demand for their equipment could ease DRAM availability problems for everyone. On the negative side, these same vendors are also consumers of a spectrum of other ICs, boards and modules that provide much of the industry's revenue.

A sustained slowing in the sales of office and computing equipment could put pressure on research and development, slow introduction of new ICs, and threaten the viability of some companies that have become overly dependent on a single market. *—Ron Wilson*

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

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

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

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

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

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

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

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

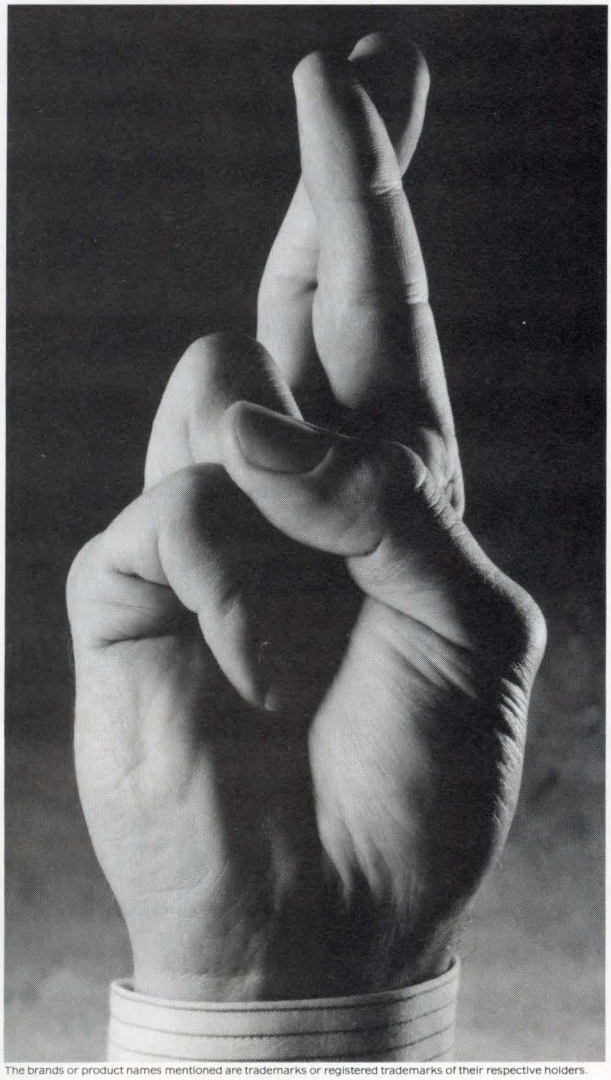

UP AND RUNNING.

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

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

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

PERFORMANCE NOT PROMISES.

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

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

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

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

◂EXCELAN▸

WE NETWORK NETWORKS.

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

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

We Supply Solutions Every Step of the Way

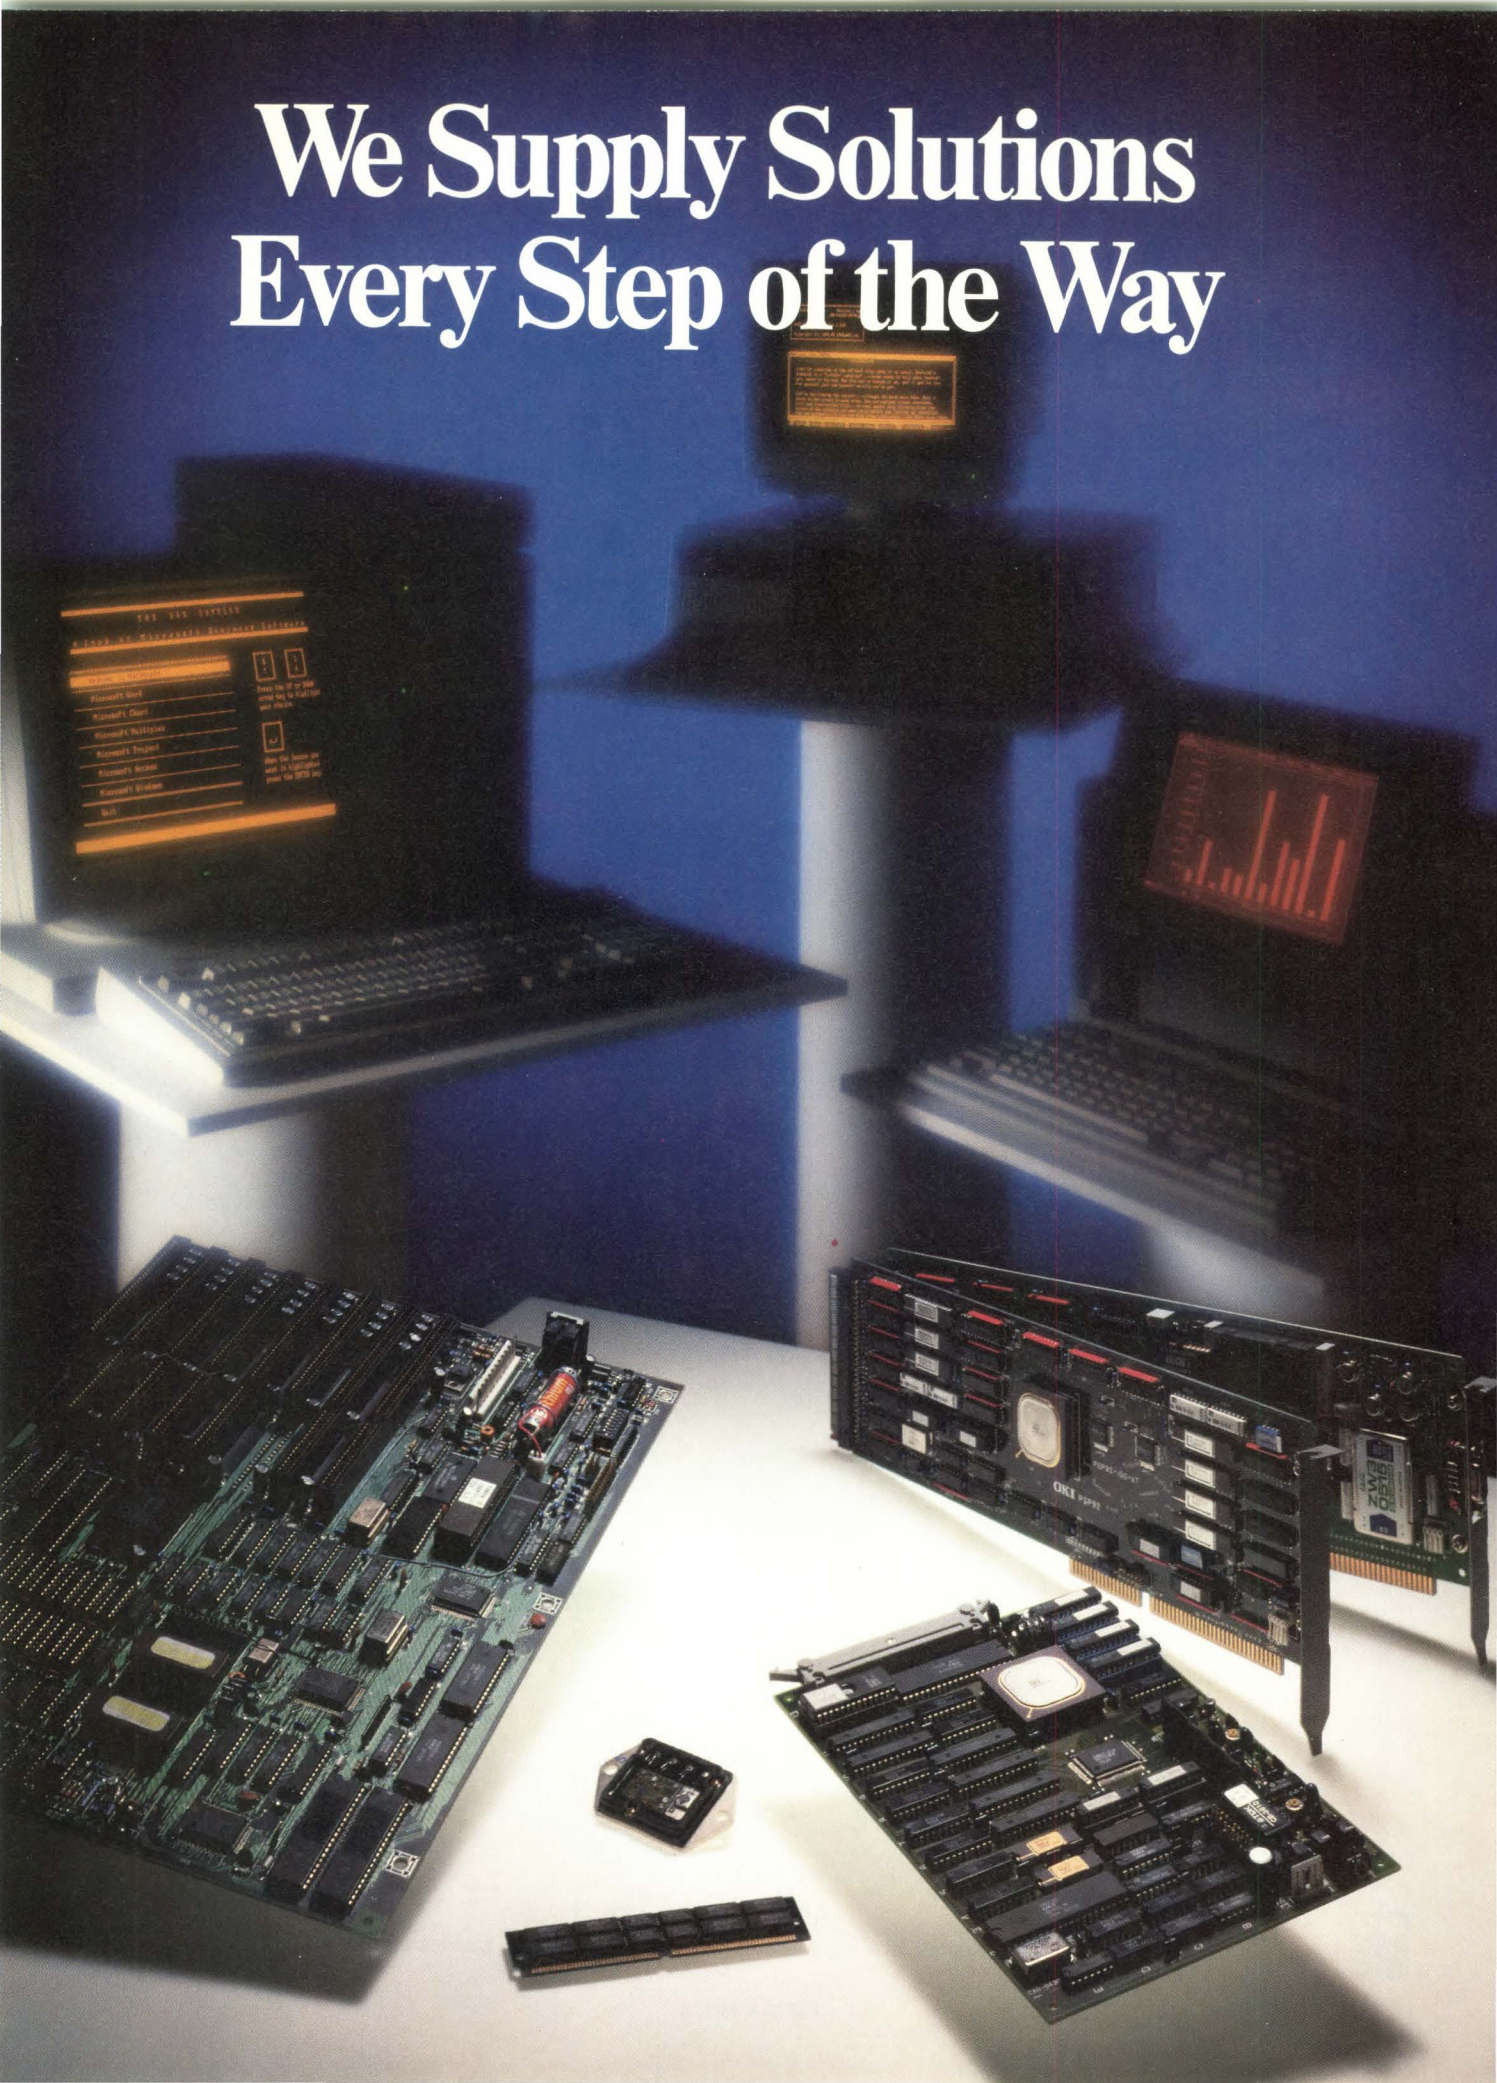

From single ICs to board-level development, Oki's systems approach gives you a competitive edge

Oki Semiconductor offers you a unique array of products, leading-edge technologies, and comprehensive support services. But what really sets us apart is our unique way of applying them. We call it *System Technologies for Customer Solutions.*

Oki systems thinking

Rather than just taking your order as other suppliers might do, we join with you as a partner and work to understand your product concept as a whole, before determining its component technologies. Then to meet your goals, we draw on a tremendous range of expertise and worldwide resources others simply cannot match.

Expertise at every level

As the only company that can deliver standard IC, ASIC, packaging, and value-added capabilities *from a single source,* we can uncover options and alternatives at many levels of system development and fabrication—opportunities others would not perceive. So we can assist you with everything from single ICs to value-added products, at any stage you require.

The result of this integrated, single-source approach is a solution that's precisely tailored to your needs. A solution that can help you achieve real competitive advantages—by speeding product development, getting you to market faster, lowering costs substantially, enhancing your product with new features, or making production more efficient.

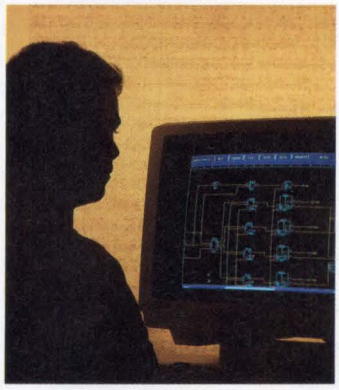

Partnerships that work

Unlike many companies, Oki is willing to share resources to co-develop products and technologies. We back this commitment with comprehensive design and support tools, as well as highly capable people who help you make the most of our resources. If you'd like to gain the kind of competitive edge Oki customers already enjoy, call us today or complete and mail the coupon below. And find out how Oki's unique blend of partnership and systems thinking can make a real difference for you.

Oki. Solutions from start to finish. CD0489

Please send detailed technical information on Oki products and capabilities:

☐ Memory
☐ DSP
☐ Microcontrollers & microprocessors
☐ ASICs
☐ Board-level solutions
☐ Packaging capabilities

( )
Business Phone

Name

Title

Company

Attach coupon to business card or letterhead and mail to: Corporate Customer Service, Oki Semiconductor, 785 N. Mary Avenue, Sunnyvale, CA 94086. Phone: (408) 720-1900.

OKI
SEMICONDUCTOR

CIRCLE NO. 8

Look, but don't touch?

An interpretation of "look and feel" should take into account the needs of both software users and software developers.

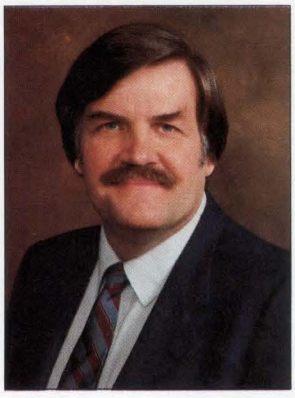

John C. Miklosz
Associate Publisher/
Editor-in-Chief

The dance card at the "look and feel" ballroom keeps getting crowded. Apple Computer is suing Microsoft and Hewlett-Packard; Lotus Development is suing Paperback Software International and Mosaic Software; Ashton-Tate is suing Fox Software; and in an ironic "twist," the developers of Visicalc are chasing after Lotus for the same reason Lotus is after Paperback Software and Mosaic. Each of the aggrieved parties is claiming copyright protection for what's come to be known as the "look and feel" of their respective software creations.

On one level, look and feel can apply only to the obvious user interface with a software package—that is, what the user sees on a display and how he or she provides input. Displaying information in a particular style of window, using a trash can as an icon and pointing to it with a mouse are examples, and they're what Apple Computer's suit is all about. On a deeper level, look and feel is also being used to describe clone software that has essentially the same user interface, logical constructs, language syntax, and algorithms as the original. Clones of this type—or perhaps we should avoid the euphemism and say copies—are at the heart of Lotus's and Ashton-Tate's suits.

The distinction is important because in the first case, a broad selection of products with identical, or similar, look and feel to the user interface directly benefits the end user. I underscored that point in an earlier editorial (see Nov. 15, 1988), when I described our experiences with a number of software packages developed to work with Ventura Publisher. (We still haven't found any with the same look and feel). It's the look and feel of the user interface—even for applications where the logical constructs, language syntax, algorithms, and so forth may be as different as those needed for a page-makeup package compared to those needed for a relational data base—that's at the heart of the Macintosh's success. That look and feel is still missing in the MS-DOS world, even though the seeds have been there in Microsoft's Windows (a threat that Apple is now addressing) and Digital Research's GEM (too little penetration to be a threat).

In the second case, users also benefit. Clones—or copies—that essentially duplicate every aspect of the originals are always cheaper than the originals. That's as true for software as it is for literature, art and designer clothes. But the cloner also benefits because he's saved the inspiration and perspiration (that is, time and money) that the creation of a truly original product requires. Let's all be honest and admit that cloning at this level isn't being done for the greater good of the user but, rather, for the good that comes to the cloners from greater profits.

With an interpretation of look and feel that takes into account both the needs of software users and those of software developers for copyright protection (that is, Apple should lose its suit, but Ashton-Tate should probably win), we'll ultimately end up with a much broader range of products that are all much easier to use.

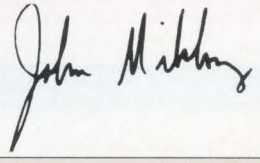

Check DY-4 for the widest choice in VME commercial, ruggedized and Mil-Spec NDI

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

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

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

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

DY-4 SYSTEMS INC.

CIRCLE NO. 9

DY-4 NDI

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

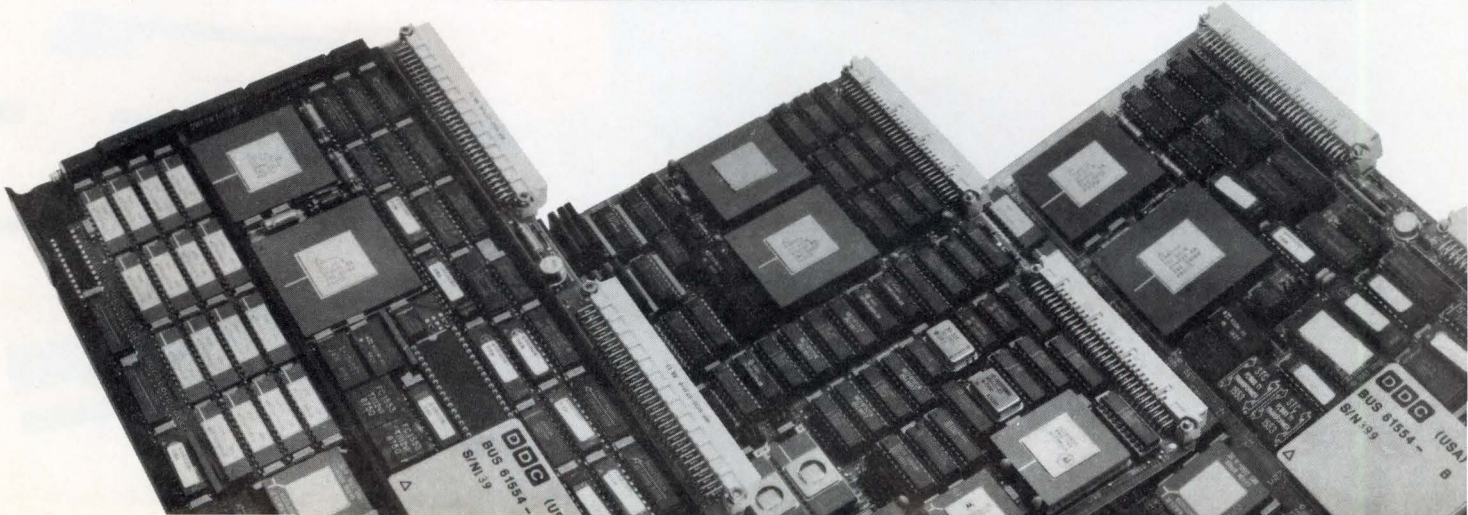

TEXAS INSTRUMENTS REPORTS ON HIGH-PERFORMANCE, IN THE ERA OF MEGACHIP™ TECHNOLOGIES

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

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

On your screen, see how our BiCMOS chips stack up

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

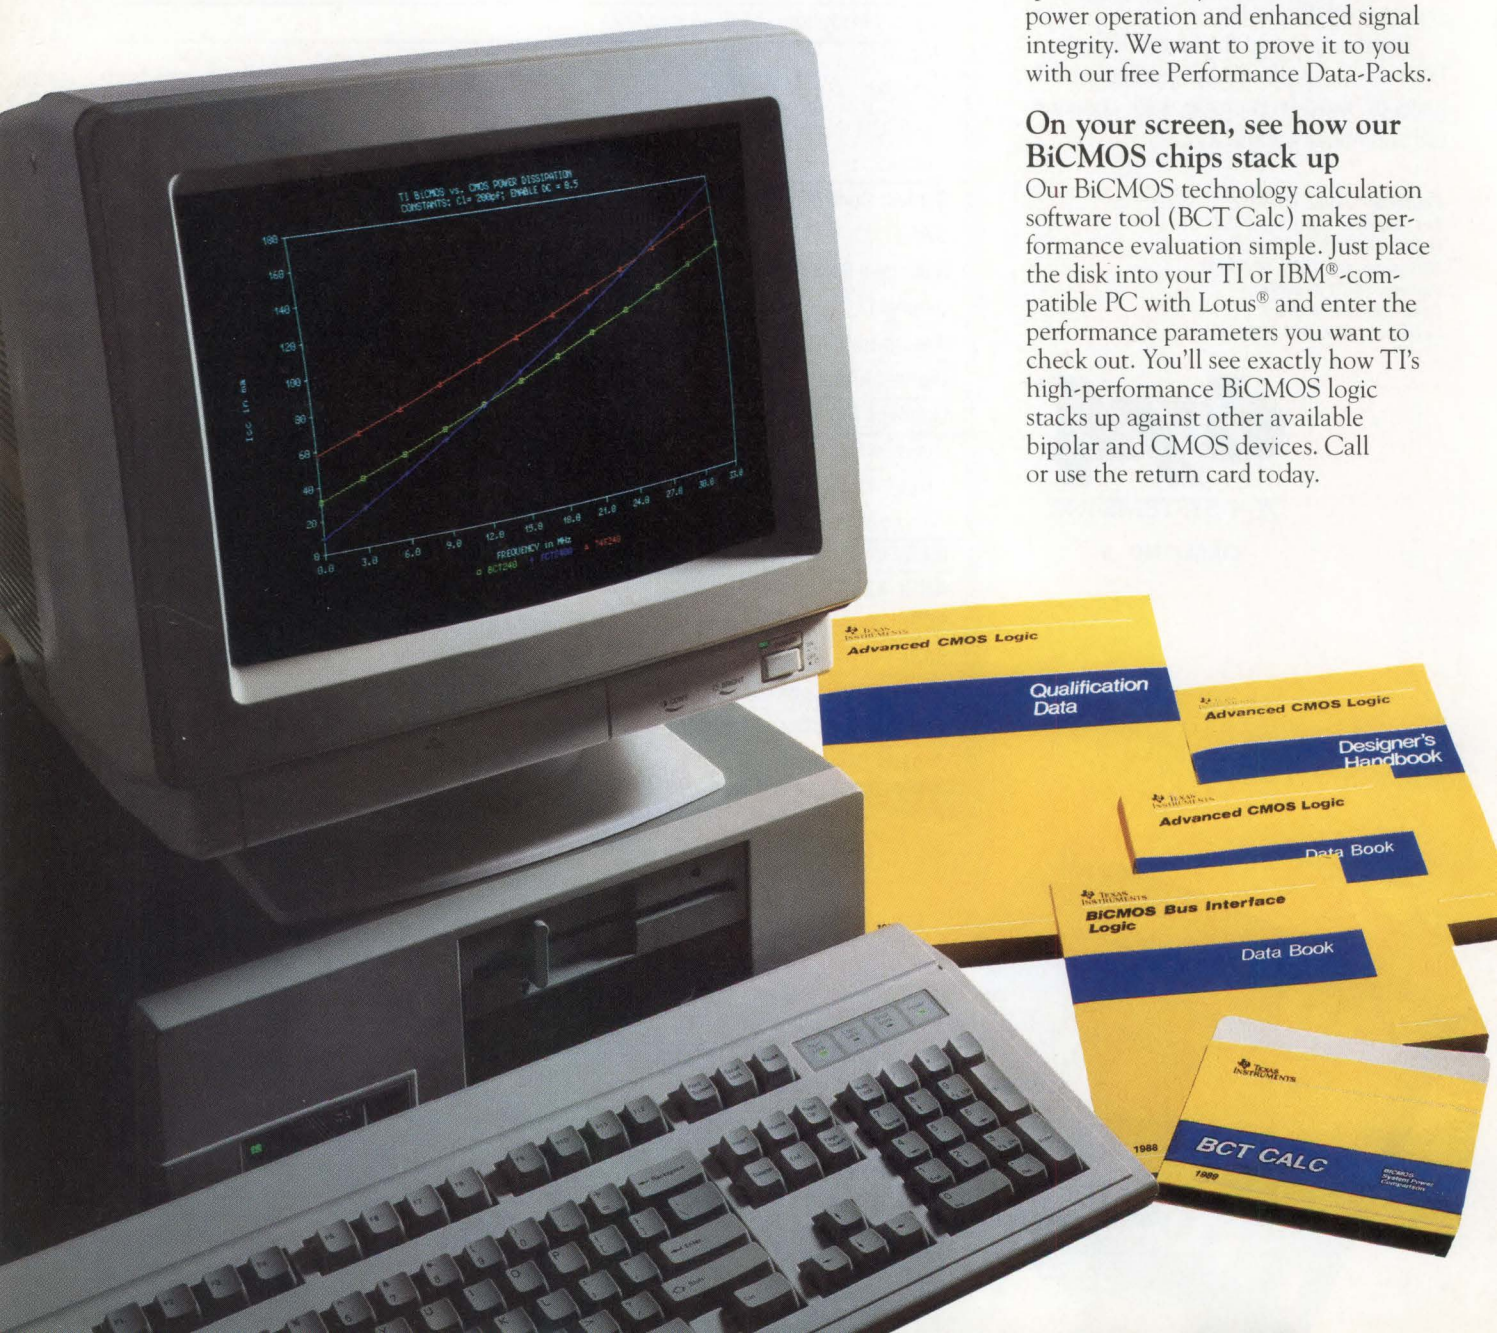

LOW-POWER LOGIC

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Our ACL family Performance Data-Pack includes:

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

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

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

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

TEXAS INSTRUMENTS

VME I/O THAT GROWS WITH YOU.

Start with just a few digital or analog I/O channels and grow to over a 1000. Burr-Brown's affordable, expandable MPV940 Series brings intelligent VMEbus I/O interfacing to process control, robotics, simulators, data acquisition and more. The MPV940 controller serves as a stand-alone computer, as a mono-master VME system or as a Master/Slave in a multi-processor configuration. RS232 or VMEbus control and software loading. Compatible with Burr-Brown's PCI-20000 termination panels. Call or write for MPV940 facts, plus new full-line catalog: Burr-Brown Corporation, 1141 W. Grant Road, MS#131, Tucson, AZ 85705. (602) 746-1111.

BURR-BROWN®

MPV940 10 MHz 68000 CPU Controller board can be fitted with two Burr-Brown analog or digital multichannel I/O modules and still occupy one VMEbus slot.

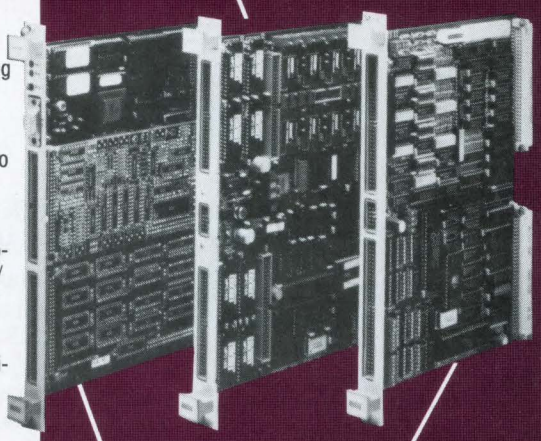

MPV941 Analog I/O expansion board provides 128 input channels & 8 output channels.

MPV942 Digital I/O expansion board provides 32 I/O lines plus 32 input lines with change of state monitor.

©1988 5800-8155

CIRCLE NO. 10

Eliminate the Jitters.

With VideoWindows, turn any IBM PC into a real-time, frame-grabbing, interactive digital video system and eliminate the interlaced video jitters. Affordably.

Combine any two video sources with VGA or EGA graphics on your 60Hz RGB monitor without jitter. Position multiple video windows — any size, anywhere. Create computer-controlled transitions, zooms, pans, and fades.

Multiple windows of digital video. Graphic overlays. No jitter.

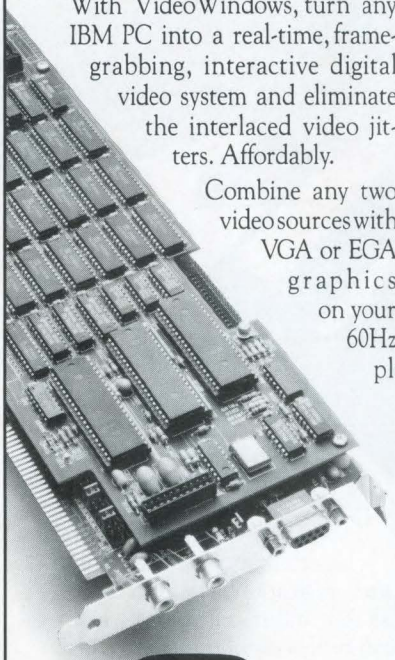

Digital video technology makes it flexible. New Media Graphics makes it affordable.

VideoWindows®

NEW MEDIA GRAPHICS
A MATRIX COMPANY

New Media Graphics
780 Boston Road
Billerica, MA 01821
(508) 663-0666

CIRCLE NO. 11

COMPUTER DESIGN®

THE FIRST MAGAZINE OF SYSTEM DESIGN, DEVELOPMENT AND INTEGRATION

PUBLISHER
David L. Allen (508) 392-2111

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

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

CONTRIBUTING EDITORS
Ernest Meyer, *Design & Development* (408) 378-6241
Howard Falk, *Software* (201) 783-4050

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

SECTION/SPECIAL PROJECT EDITORS
Ron Wilson, *Technology Updates* (503) 224-9396
Michael G. Donlin, *New Products* (508) 392-2123
Patti Kenney, *Directories* (508) 392-2124

MANAGING EDITOR: Susan R. Nykamp

ASSISTANT EDITOR: Kerstin R. Feldman

COPY EDITOR: Barbara-Ann Scofidio

EDITORIAL ASSISTANTS: Claire Coupal, Claire Ellis

ART DIRECTOR: Jan Horner
ILLUSTRATOR: Fred Lomas

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

CIRCULATION DIRECTOR: Robert P. Dromgoole

PennWell PUBLISHING COMPANY

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

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

BRINGING WORKSTATION 3-D GRAPHICS TO THE PERSONAL COMPUTER.

NTH 3D ENGINE.

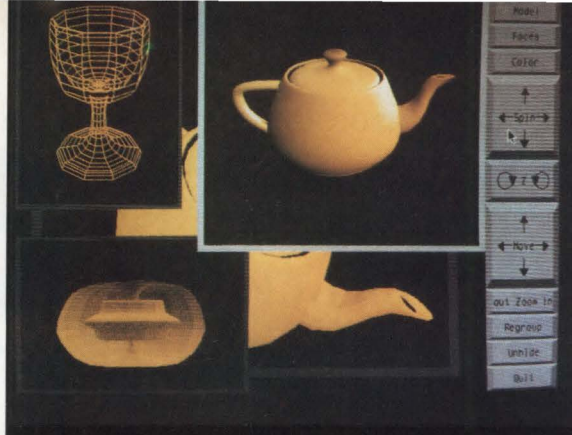

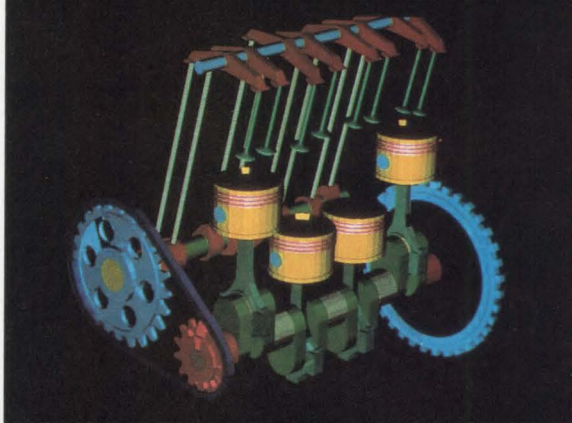

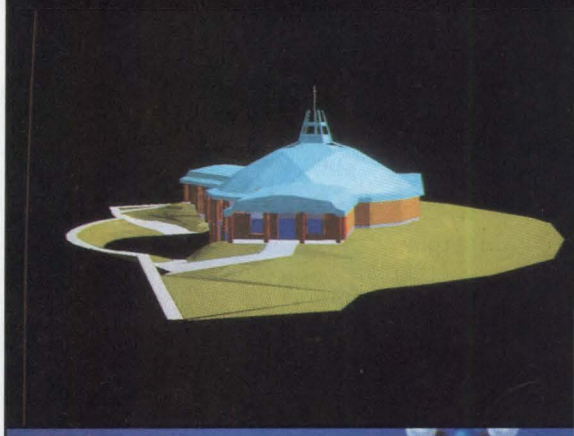

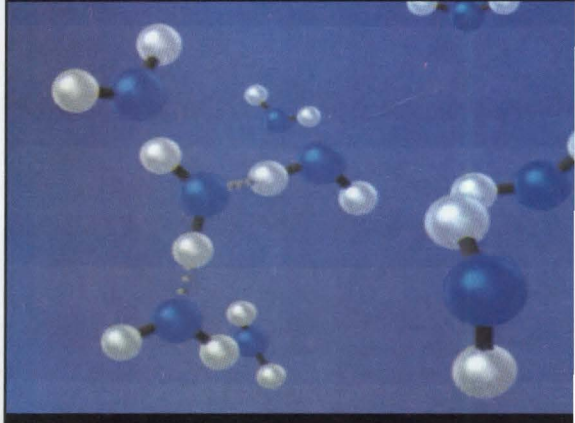

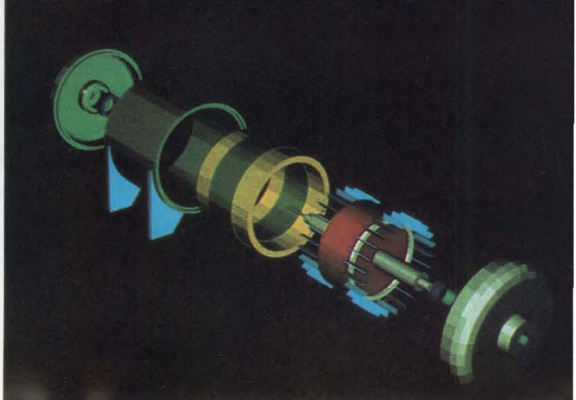

Interactive 3-D Performance on a PC

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

Realistic Display of Complex Objects and Data

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

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

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

Prices Start at $6995

For More Information, Call 800-624-7552

Nth 3D ENGINE™ DISPLAY CONTROLLER

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

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

CIRCLE NO. 13

INTEGRATED CIRCUITS

80860 CPU positions Intel to take on minisupercomputers

Ron Wilson, Senior Editor

The first thing that catches one's attention about the Intel 80860 CPU—formerly known only by rumor as the N10—is the sheer size of the numbers involved. Numbers such as 120 million operations/s, 90,000 Dhrystones or 21 Linpack MFlops sound as if they belong in the supercomputer world, not the microprocessor world. And numbers such as 1 million transistors on a single die sound as if they belong in science fiction. Yet the 80860 CPU is real, and in the hands of selected beta sites, according to Intel (Santa Clara, CA). The new processor, which can be used as either a Unix-based host CPU or a dedicated application engine, is likely to change one branch of computing forever.

That branch is personal supercomputing, the territory originally outlined by mini-Cray vendors, explored by startups like Ardent Computer (Sunnyvale, CA) and Stellar Computer (Newton, MA), and recently invaded by the high-end DN10000 from Apollo Computer (Chelmsford, MA). In this arena, enormous floating-point programs analyze mechanical structures, predict fluid flow or simulate airframes. Users interact with these huge programs through dynamic three-dimensional renderings on high-resolution displays.

The superworkstations that serve this market must provide a fast Unix platform, tremendous floating-point throughput and prodigious 3-D drawing capability. It's not unusual for such systems to offer 5 to 10 Linpack MFlops and to be able to draw 150,000 3-D polygons/s.

A chip for supercomputing

In the past, superworkstations generally contained at least three discrete processors: an integer CPU implemented with either a RISC microprocessor or a multichip proprietary design, a vector processor based on a specialized floating-point chip set, and a proprietary graphics board usually built from semicustom VLSI. Naturally, such a processor cluster was expensive, and the number of interconnections between chips, or even boards, in the design built unavoidable delays into execution cycles.

But the new Intel part offers to change this architecture entirely. Taking a cue from the multiple-execution design of the Apollo DN10000, the 80860 CPU contains three execution units, two register files, an 80386-compatible memory-management unit and a set of large caches on a single, million-transistor die. The result provides a cost and space savings, but more important, the ability for the entire multiunit processor to execute at 40 MHz or more over large sections of code, translating high peak Mips numbers into staggering benchmark results.

Where multiboard processors have been achieving 10 Linpack MFlops, the single-chip 80860 reportedly hits 21. Where separate proprietary graphics pipelines generate 150,000 polygons/s, a single 80860 CPU achieves 50,000 triangles/s, a somewhat comparable measure of transformed drawing speed. Since polygon generation seems to be an application that adapts well to parallelism, multi-80860 graphics accelerators could go well into the hundreds of thousands of polygons per second.

Not RISC, but RISC-like

Two sets of concepts are fundamental to this sort of execution speed. First, the 80860 relies on RISC concepts such as simple, one-cycle-per-

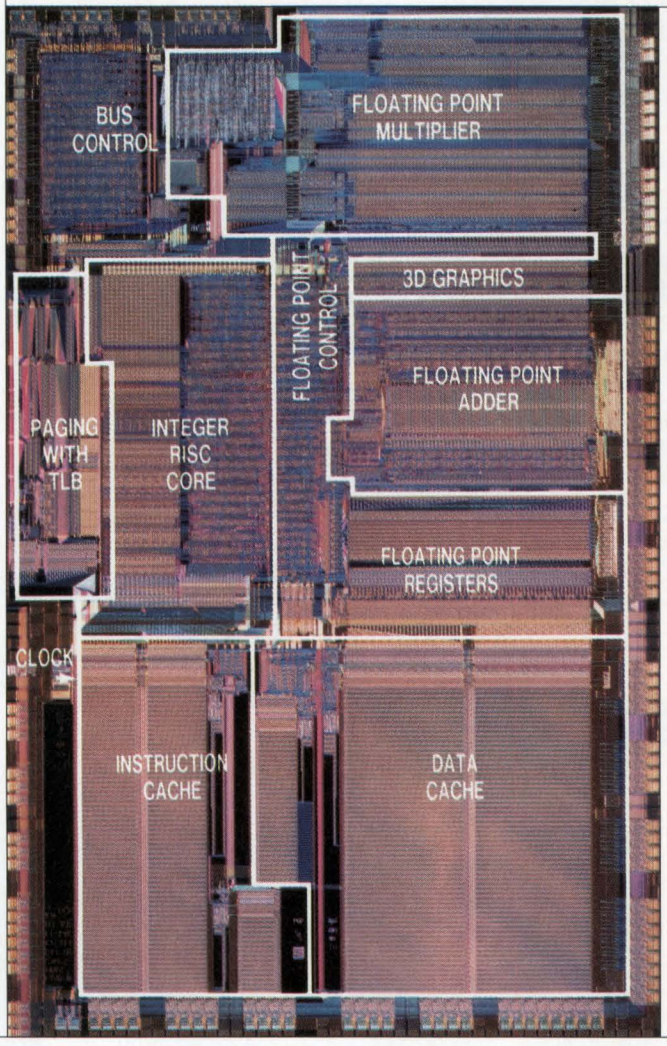

The Intel 80860 CPU packs 1 million transistors onto a 10- × 15-mm die. About a third of the die is used for the RISC-like integer unit, a third for the vector floating-point hardware, and a third for instruction and data caches.

You've always known about Hitachi's well-deserved reputation for high quality, high reliability disk drives. Now, there's more good news. It's the *super-high density DK515-78 780MB 5.25-inch drive,* the latest addition to Hitachi's comprehensive selection of Winchesters, magnetic tape subsystems and ultra-high density optical disk drive subsystems.

The performance you want.

Our new DK515-78 5.25-inch drive has it all. It features a full 780MB of storage capacity, lightning fast 16ms average access time, and a blazing 2.4MB/sec. data transfer rate. Plus, it offers a choice of SCSI or ESDI interfaces.

The reliability you demand.

When you specify Hitachi, you get more than high performance. You get unbeatable Hitachi reliability. Our DK515-78 is designed and manufactured to use our own motors, heads, microprocessors and LSIs. In-house design and production of all these major components and the most stringent quality assurance program in the industry, gives the DK515-78 a remarkable 40,000 hour MTBF.

The partnership you can trust.

Do business with Hitachi, and you won't get stuck with yesterday's news. You'll get a reliable business partner who can deliver the drives you need, when you need them. 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 America, Ltd.
Computer Division
950 Elm Avenue, Suite 100
San Bruno, CA 94066

HITACHI®

CIRCLE NO. 14

INTEGRATED CIRCUITS

instruction design to achieve high clock frequencies and high execution rates in the integer execution unit. "We don't view RISC as an architecture, but rather as an idea," claims Bill Rash, microprocessor marketing manager at Intel. "In the integer unit, we use one-clock instructions, delayed branches and a load-store architecture with a file of thirty-two 32-bit registers."

But Rash goes on to explain that the 80860 departs from pure RISC ideas in many ways. As examples, he cites the extensive floating-point hardware, which takes up about a third of the die area, and the dedicated graphics hardware. The graphics unit, which shares the floating-point data paths, performs pixel filling, Z-buffering and other hard-to-simulate operations in hardware. "The important thing is not that the instruction set is reduced, but that it supports real languages for real applications," Rash says.

Second, and equally important to the 80860's high processing speed, is the notion of multiple execution units. This concept, which has been gaining popularity since the introduction of the multichip Apollo DN10000 CPU last year, simply observes that, at a given clock frequency, the Mips can be doubled if two instructions are executed at once.

Of course, in practice it's not as easy as that. The CPU must be able to fetch two instructions at once and be sure that the instructions are independent: that one doesn't require data produced by the other. The 80860 normally fetches one 32-bit instruction at a time from its instruction cache, but it has an operating mode in which it, like the Apollo DN10000, fetches an integer instruction and a floating-point instruction in one 64-bit cycle. These two instructions then execute in parallel in their separate pipelined execution units. Once the instructions are fetched, though, keeping things running gets complicated.

Managing bandwidth

The ability to achieve sustained performance on complex benchmarks or actual application codes is more than a matter of execution-unit design. In a multiple-execution-unit CPU, the bandwidth necessary to keep instructions and data moving through the execution units can be a more serious problem than the instruction execution times. Consequently, the Intel design team lavished effort and silicon on the network of caches, register files and buses that surround the execution units.

The most obvious result of this attention is in the third of the die dedicated to caches. The 80860 uses a pair of 32-bit caches: an 8-kbyte data cache and a 4-kbyte instruction cache. "It was important at these speeds not to restrict the cache bandwidth by making an off-chip connection," claims Rash. "So we simply put in as much cache as we could on a manufacturable die." The importance of having the caches on-chip becomes obvious when one looks at the way these memories are connected to the execution units: the 64-bit path for the instruction cache and 128-bit path for the data cache would have put impossible strain on the pin count of any reasonable package.

In fact, the 80860 uses wide internal data paths heavily. While the RISC core processor works on 32-bit integers, the floating-point unit and its register file are 64-bit, as is the graphics execution unit. These devices are interconnected by a lattice of 64- and 128-bit paths. "One of our foremost challenges was sustaining the floating-point performance," says Rash. "We learned from the example of discrete floating-point units that offer high instantaneous speed but lack the I/O bandwidth to keep themselves fed."

The thicket of data paths on the 80860 can provide an instruction each to the integer and floating-point units, an operand to the integer unit, and two operands to the floating-point unit, while returning a result to the floating-point register file, on each cycle. That works out to an aggregate on-chip bandwidth of nearly 1 Gbyte/s.

The internal data paths can keep all the execution units fed as long as data and instructions are coming from cache. But in the event of a cache miss, the microprocessor is dependent on its external bus. Here,

Much of the benchmark performance of the 80860 is due not to the speed of the execution units, but to the bandwidth of the on-chip data paths. Internal connections let the integer and floating-point units run concurrently, taking an instruction and a set of operands on each clock cycle.

INTERNAL CONNECTIONS IN THE 80860

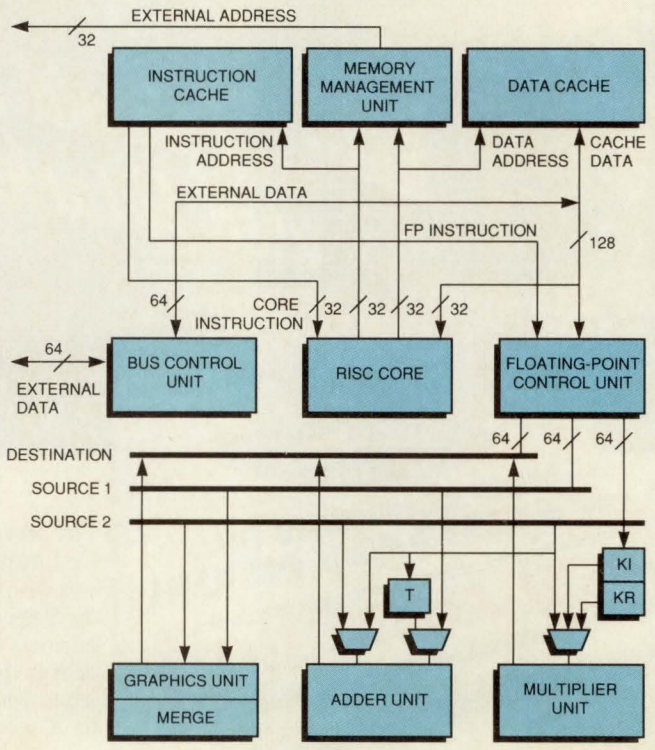

too, Intel has paid attention to speed. The external bus is 64 bits wide and provides a next-cycle-in-page signal for static-column dynamic RAMs. In addition, the bus has a pipelined operating mode that lets three read operations be pending at once. This gives the DRAM six clocks, or 150 ns at 40 MHz, to respond and keeps the bus controller running at full speed. But it does mean that for operations on large data sets, where operands will have to come from memory instead of cache, the 80860's throughput will drop substantially.

But is it manufacturable?

The greatest question hanging over the 80860 has to be manufacturability. At 1 million transistors on a 10- × 15-mm die, the part is probably the densest nonmemory standard product ever introduced. Yet the die isn't all that large, and the process is already well-characterized. "It's the same 1-micron process we use for the 80386," explains Rash. Certainly, Intel doesn't have a history of manufacturing failures.

As further proof of the company's confidence, Intel sources claim that 80860 beta units are already in the hands of selected accounts, and that much of the system software is already completed or in development. C and Fortran compilers and a Fortran vectorizer are said to be available, and a multiprocessing version of Unix System V release 4 is also underway.

Rash claims that the chip is already designed into some graphics accelerator applications where Unix availability wasn't a gating factor. This raises the possibility that in terms of actual application starts, as opposed to CPU-on-a-board products, the 80860 may be as far along as the Motorola 88000 or any of the newer Sparc chips. In terms of product availability, at least, the picture is good: 33-MHz samples are available, with production at that speed expected in the third quarter of 1989. Samples at 40 MHz should be available in the second half of 1989.

Taken in perspective, though, the 80860 isn't a direct challenge to most of the 32-bit RISC microprocessors on the market. Many of those chips are simply incapable of achieving the floating-point and graphics performance offered by the Intel design. Instead, the 80860 should thrive in the area of personal supercomputers, a market now dominated by proprietary multichip CPU designs with separate, multichip floating-point units and graphics processors. In its chosen area, the new Intel device should completely change the economics of workstation design. □

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

Apollo Computer
(508) 256-6600, ext. 7702 ..... **Circle 316**
Ardent Computer
(408) 732-0400 ..... **Circle 317**
Intel
(408) 765-4851 ..... **Circle 318**
Stellar Computer
(617) 964-1000, ext. 510 ..... **Circle 319**

Does your windowing package leave you hanging?

Frustrated with the limitations of your windowing package?

ÆWINDOS is a high-powered set of tools designed to interactively create windows at light speed with both DOS and UNIX systems.

Written entirely in portable "C", it includes the complete source code and even make files, so changes are quick and easy.

Don't wait until you're caught in a bind, call us today. Modular, portable, and efficient, ÆWINDOS won't leave you hanging.

æSOFT

1-800-634-5494 BBS (303) 449-5714

* UNIX is a trademark of AT&T.

CIRCLE NO. 15

PERFORMANCE
means 15ns wide word SCRAMs

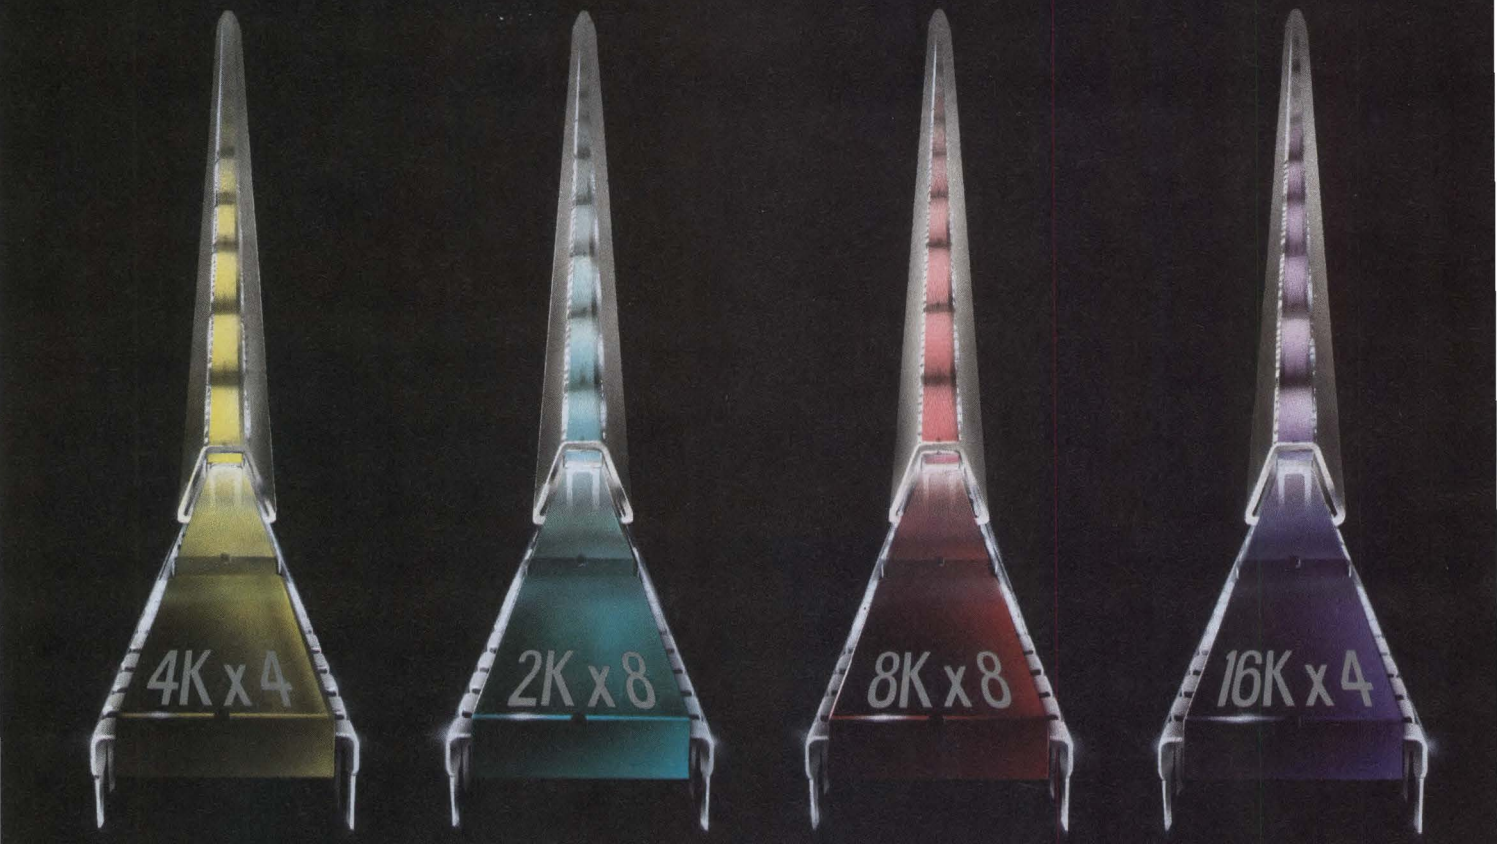

15ns

SCRAMs are Static CMOS Random Access Memories from Performance Semiconductor. At 15ns address access time these 64K's and 16K's are the world's fastest. SCRAMs are manufactured in Performance's six inch class 1 fabrication facility using PACE II 0.7 micron gate length technology which has set the standard for memory speed.

There is immediate availability of 15ns 64K and 16K bit SCRAMs compatible with JEDEC standard pinouts. Also available from stock are 17, 20 and 25ns speed versions.

15NS SCRAM PRODUCT GUIDE

PART	CONFIG.	SPEED	AVAIL.
P4C164	8K x 8	15ns	NOW
P4C188	16K x 4	15ns	NOW
P4C198	16K x 4	15ns	NOW
P4C198A	16K x 4	15ns	NOW
P4C1982	16K x 4	15ns	NOW
P4C1981	16K x 4	15ns	NOW
P4C168	4K x 4	15ns	NOW
P4C169	4K x 4	15ns	NOW
P4C170	4K x 4	15ns	NOW
P4C1682	4K x 4	15ns	NOW
P4C1681	4K x 4	15ns	NOW
P4C116	2K x 8	15ns	NOW

FAST, COOL, & AFFORDABLE

For further information or to order 15ns SCRAMs call or write:
Performance Semiconductor
610 E. Weddell Drive,
Sunnyvale, CA 94089
Telephone: 408 734·9000

PERFORMANCE
SEMICONDUCTOR CORPORATION

CIRCLE NO. 16

INTEGRATED CIRCUITS

80860 blends vector processing into a RISC-like architecture

Ron Wilson, Senior Editor

In the minisupercomputing world for which the new 80860 CPU from Intel (Santa Clara, CA) is intended, the heaviest computing tasks tend to be matrix operations. These processes involve multiplying an enormous vector or matrix by another enormous vector to compute the state of a structure, a flowing fluid, a three-dimensional image or a digitized waveform.

Taking advantage of the inherent order in these calculations, vendors have in the past designed specialized floating-point vector-processing hardware to boost the throughput of their systems. This hardware generally relies on a deeply pipelined floating-point unit, dedicated operand memories, and dedicated operand buses to achieve high computational rates.

By building a pipelined multiplier-accumulator configuration, designers can achieve several floating-point operations per clock cycle. And by keeping data in dedicated memories hard-wired to the pipeline inputs, the designers can ensure that the pipeline stays full.

But such a hard-wired approach to floating-point design seems incompatible with the architecture of a general-purpose computer, which lacks dedicated memory areas. The architects of the 80860 CPU have attempted to blend the data-flow notions of vector processing into the more general architecture of a RISC-like CPU, achieving an interesting compromise.

A flow-through model

Normally, a microprocessor has a scalar FPU: once an operation is started, the unit will accept no more operands until the first operation is completed. This simplifies both the hardware design and the programming model of the FPU, but at the expense of performance. In contrast, most building-block FPUs have a pipelined mode, which lets a new set of operands be loaded on every clock.

The 80860 FPU provides both modes, via two sets of floating-point instructions. Scalar instructions load a set of operands from the floating-point registers and store the result after one to four clocks. Pipelined operations load a set of operands and simultaneously store the result coming from the end of the three-stage pipeline, starting a new set of operations on each clock. (There are additional delays in double-precision mode, since the multiplier is effectively a two-stage, four-clock device in that mode.)

Since the floating-point multiplier and adder are independent units, they can be chained together to form a multiplier-accumulator pipeline. This capability lets the 80860 execute two floating-point operations, a multiply and an add or subtract, on each clock cycle. As long as the chip can keep its pipeline fed from the thirty-two 32-bit floating-point registers, it can perform multiply-accumulates at 80 MFlops.

Instead of the dedicated static RAM used in building-block vector processors, the 80860 relies on its floating-point register file and its data cache to provide operands to the pipeline. Extra-wide paths and a special instruction-fetch mode help with the monumental task of moving up to four operands per clock.

The path between the data cache and the floating-point register file is 128 bits wide, so two 64-bit operands can be loaded from cache in a single cycle. Three 64-bit paths connect the register file to the floating-point execution unit, so the pipeline can take in two operands and store one result in each cycle.

Further, the 80860 provides a dual-instruction mode, which fetches one integer-unit instruction and one floating-point instruction on each clock cycle. In this mode, the integer unit can execute a floating-point load or store while the FPU performs pipelined operations. One can thus think of the entire FPU as executing directly out of data cache, with input

A MULTIPLIER-ACCUMULATOR PIPELINE

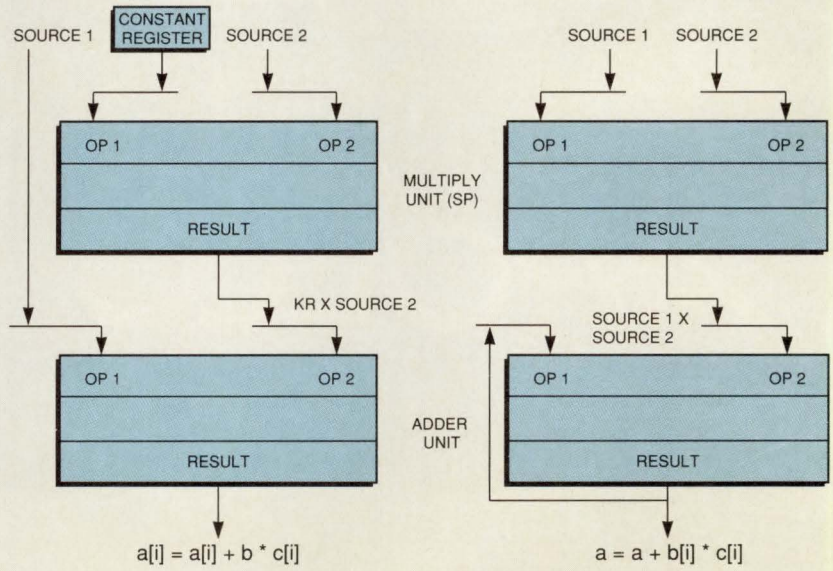

By chaining the floating-point adder to the multiplier, the Intel 80860 can effectively convert its floating-point unit into a pipelined multiplier-accumulator. In single-precision mode, the pipeline will produce a result on every cycle of the chip's 40-MHz clock. Consequently, some linear operations can be executed at the rate of two floating-point computations per clock.

INTEGRATED CIRCUITS

and output pipeline registers actually residing in the floating-point register file.

Precise exception handling

This level of concurrency is wonderful for processing speed, but such architectures are notoriously difficult when exception conditions occur. If an overflow occurs and the hardware can't report which of the operations in the pipeline caused the exception, it can take the software several milliseconds to unravel the mess.

The 80860 provides two features to help with this problem. "The chip conforms to all IEEE exception-handling requirements," claims Bill Rash, Intel microprocessor marketing manager. "Also, hardware marks the location of an exception in the pipeline, and the floating-point instructions are restartable. So the exception trap routine can quickly identify the problem, take corrective action and get the pipeline flowing again."

A more serious challenge to the usefulness of a deeply pipelined architecture is compiler support. If the pipeline is so complex that codes must be handcrafted, few users will realize the potential speed of the hardware. And the 80860 pipeline, with its multiple configurations, differing stage delays, and possibility of mixing single- and double-precision operations, threatens to be complex indeed.

But Intel is working on this issue as well. Green Hills Software (Glendale, CA), one of the most experienced optimizer houses in the industry, will supply an optimizing C compiler for the chip. Pacific Sierra Research (Los Angeles, CA), a firm with considerable experience on vector-supercomputers, will write a vectorizing preprocessor for Fortran code. These tools should help users exploit the pipelined hardware more effectively.

Naturally, there are limitations to the Intel approach. The limited size of the 80860's data cache, the limitation of one integer operation per clock, the relatively low speed of the external dynamic RAM connection, and the small floating-point register file will all impede the execution of large array operations. But as a compromise between the dedicated data-flow architectures of pure vector-processing machines and the unacceptably high overhead of scalar microcomputers, the 80860 represents a challenging step forward in silicon architecture. The chip may establish an entirely new level of MFlops per dollar on actual application codes. □

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

Green Hills Software
(818) 246-5555 ..... **Circle 313**
Intel
(408) 765-4851 ..... **Circle 314**
Pacific Sierra Research
(213) 820-2200 ..... **Circle 315**

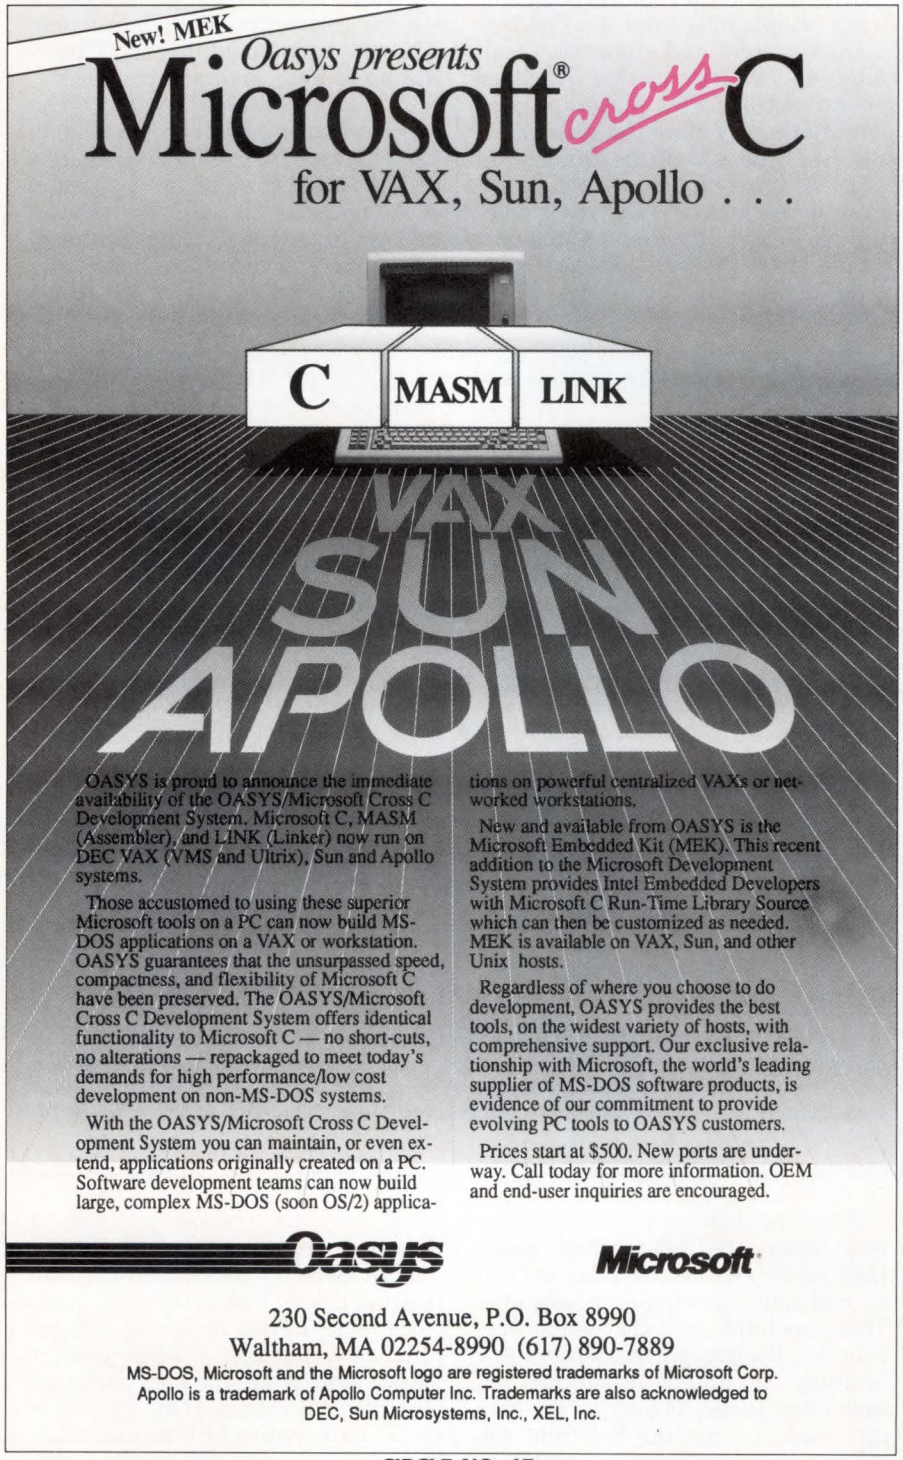

CIRCLE NO. 17

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

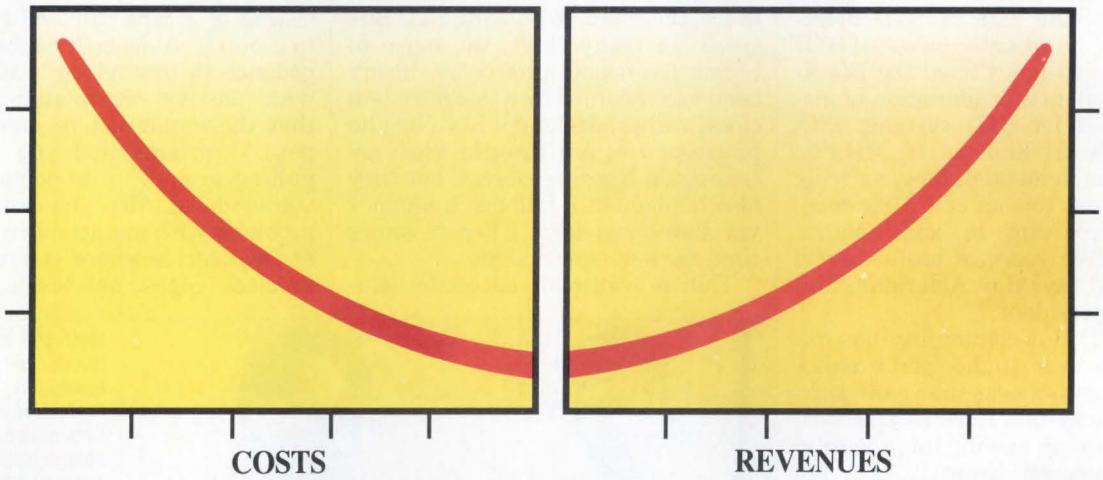

Lower costs and higher productivity add profits to your bottom line.

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

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

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

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

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

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

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

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

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

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

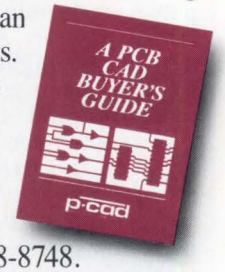

What we have to say will make you smile. P-CAD, 1290 Parkmoor Avenue, San Jose, CA 95126.

p-cad®
PERSONAL CAD SYSTEMS INC.

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

CIRCLE NO. 18

COMPUTERS AND SUBSYSTEMS

Languishing STD Bus receives a late-life kicker

David Lieberman, Senior Editor

It's no secret that the STD Bus is suffering from a decline in use as designers migrate to buses with wider data paths and bigger bandwidths. Matrix (Raleigh, NC), however, intends to halt, or at least decelerate, the rate of STD defections with its recently unveiled STD backplane design. Called the Blackplane, it offers a combination of elements that let STD systems with clock rates as high as 16 MHz be reliably implemented. Most existing STD systems top out at 6 MHz or so. "We're expecting to add another three to five years of useful life to STD Bus," says Ray Alderman, Matrix vice-president.

The STD Bus community has migrated to ever higher-performance microprocessors over the years, jacking up the system clock as appropriate, with many moving into the Intel family processor camp to take advantage of PC software compatibility. As many are quick to point out, however, the STD Bus specification, formulated back in the Z80 age, isn't geared for high-frequency operation.

"The STD specs aren't very tight," laments Alan Beverly, engineering manager at Ziatech (San Luis Obispo, CA). "It's got no particular capacitive loading spec, no distance limitations on how far drivers and receivers should be from the backplane, no termination recommendation and so on. As bus speeds grow higher and higher, we've been seeing a lot more problems with signal integrity, and those design rules that people have been somewhat lax about in the past are becoming extremely problematic."

Some place the blame for poorly performing systems not on any particular inadequacies in the bus spec, but on the generally poor quality of STD backplanes. "Not everybody has caught on to the fact that system speed can be severely limited by the design and construction of the backplane," says Kurt Priester, president of Computer Dynamics (Greer, SC). "Most don't pay a lot of attention to their backplanes and have to put up with severely limited performance because they don't want to put in the money for shielding and noise suppression that's necessary in very high speed systems."

There has been substantial resistance, therefore, to pushing STD Bus speed. In many cases, the move to higher-frequency processors hasn't been accompanied by a faster system clock; rather, designers have let the processor run full throttle when accessing on-board resources, but they also implement a fallback frequency via slow-down logic when it comes time for bus transactions.

This provides an adequate solution when enough of the crucial resources for an application can be fit on a single 6.5×4.5-in. STD Bus board—whether because of limited resource requirements or very dense board integration. But when it isn't possible to accommodate sufficient resources, off-board accesses undermine processor performance.

With the introduction of its Blackplane high-performance STD backplane, Matrix expects to expand STD's useful life by at least three to five years, according to Bill Harris (left), design engineer, and Ray Alderman, vice-president.

The 8-MHz frontier

About a year and a half ago, several STD companies began to see trouble coming in the form of a looming performance roadblock that would limit their ability to clock up to customer requirements. "In most instances, when people started to move to 8-MHz 8088s, there was no problem," says Alderman, "but beyond that you start to bump into two nasty phenomena—ground shift and metastability. To deal with these two engineering problems, you also have to face another giant monster: transmission line effects."

A large part of the problem with high-frequency STD, and much of the solution, lies with the backplane termination scheme. "An STD Bus backplane has passive terminators—that is, they're just snubbers, consisting of a capacitor and a resistor to ground," Alderman explains. "Impedance to ground on a backplane with passive terminators is such that there may not be enough of a path to ground and you wind up pulling ground up to some positive voltage level. Also, you can run into problems with metastability at higher frequencies where you're looking at clock edges, not levels. On the other hand, active terminators such as the Thevanin terminators used with VMEbus won't work well with STD Bus, because of the additional impedances of the gold-edge finger connectors."

With input from Ziatech, which resells Matrix backplanes and card cages, and several other STD companies, Matrix set out to solve all three problems. At the time, Ziatech was less concerned with higher-speed systems than with bigger systems. "Over a year ago, we started seeing some problems even at 8 MHz with the STD Bus, particularly with long card cages of 26 or so slots," says Beverly.

"The problems you see in a big

THE UNISITE PROGRAMMER: BECAUSE STATE-OF-THE-ART IS A STATE OF CHANGE.

PROGRAMMING TECHNOLOGY THAT SUPPORTS ADVANCED DESIGNS—TODAY AND TOMORROW. The UniSite™'s universal programming technology is the fastest and easiest way to keep up with new devices and packages. Its software-configured pin driver system provides a single site for programming any DIP device up to 48 pins, including PLDs, PROMs, IFLs, FPLAs, EPROMs, EEPROMs and microcontrollers.

The optional ChipSite™ module supports logic, memory and microcontrollers in the most popular surface-mount packages—PLCCs, LCCs and SOICs. Or add SetSite™ to gang program up to eight 40-pin EPROMs, or set program a single file into multiple EPROMs.

INSTANT ACCESS TO NEW DEVICES. The UniSite's device algorithms and system software are stored on 3½" micro diskettes, which are updated several times per year. To update the UniSite, simply load new master diskettes.

FAST, EASY PROGRAMMING. Menu-driven operation with step-by-step prompts makes programming simple. For added convenience, you can operate the UniSite from an ASCII terminal or from a PC using the provided terminal emulator. On-screen help messages are available throughout operation.

To speed parts selection, the UniSite provides a built-in list of devices by manufacturer name and number. You can also save your most frequently used programming parameters and jobs for instant recall.

DESIGN FREEDOM FOR TOMORROW. When leading-edge designers need to use the latest devices, they need the programming freedom only the UniSite provides. Call Data I/O® today, and ask about the UniSite. Because state-of-the-art never stops changing.

1-800-247-5700
Ext. 660

Data I/O Corporation 10525 Willows Road N.E., P.O. Box 97046, Redmond, WA 98073-9746, U.S.A. (206) 867-6899/Telex 15-2167
Data I/O Canada 6725 Airport Road, Suite 302, Mississauga, Ontario L4V 1V2 (416) 678-0761
Data I/O Europe World Trade Center, Strawinskylaan 633, 1077 XX Amsterdam, The Netherlands +31 (0)20-6622866/Telex 16616 DATIO NL
Data I/O Japan Sumitomoseimei Higashishinbashi Bldg., 8F, 2-1-7, Higashi-Shinbashi, Minato-Ku, Tokyo 105, Japan
(03) 432-6991/Telex 2522685 DATAIO J

DATA I/O Corporation

©1989 Data I/O Corporation

CIRCLE NO. 19

COMPUTERS AND SUBSYSTEMS

card cage because of the number of slots, capacitive loading and so forth are very similar to what you see on a smaller card cage at higher speeds," says Beverly. "You'll probably see the same problems in a 12-slot, 16-MHz system that you see in a 26-slot, 8-MHz system, though their severity may be more pronounced in the higher speed system."

How did the Matrix engineers attack the problems of ground shift, metastability and transmission-line effects? The Blackplane uses new layout techniques, a four-layer structure, heftier power and ground planes, precise materials control and a little black magic in the termination scheme.

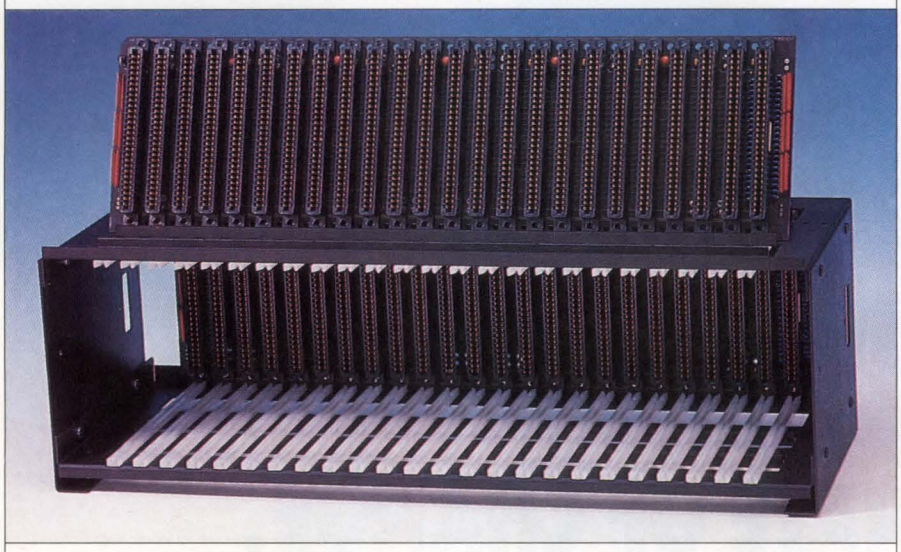

Matrix's Blackplane should help slow the decline of STD Bus by letting designers migrate to new high-speed processors. With the Blackplane, systems with clock rates of up to 16 MHz can be reliably implemented.

"To avoid metastability problems, you need good solid power distribution and balanced impedance," says Alderman. "First, you need to get the crosstalk down, which we do by maximizing the distance between any two active signal lines all along the backplane. That's pretty straightforward, as is the use of double-thick ground and power planes to minimize impedance. But our engineers took a fresh approach to balancing line impedance to eliminate transmission line effects. The average engineer references all the lines to ground, but that screws up impedance, so our guys referenced them to the closest plane."

Part of Matrix's solution to the problem of noise—induced by, among other things, high-frequency power supplies—is materials control. "We control the thickness of the copper, the width of the copper for balanced impedance on the lines, and the thickness of the printed circuit board material such that we actually make the backplane a capacitor, running about 140 pF/in.^{2}, which creates an extremely good high-frequency noise filter between the backplane layers," explains Bill Harris, design engineer at Matrix.

High-frequency operation not new

Computer Dynamics' backplanes have been capable of high-frequency operation for some time now, according to Priester. "We modified our backplane with proper shielding and termination frequency capacitors about five years ago," he says, "and we've successfully tuned it to 10- and 12-MHz Z80s and, more recently, 12- and 16-MHz 80186s and 80188s." In general, however, Priester considers running high-frequency signals over an STD backplane unwise.

"I don't believe the backplane in an industrial environment should be cluttered up with high-speed memory transfers," he says, "because that multiplies the probability of electrical glitches and noise transients hundreds of times and introduces probably a four to eight times greater chance of interference destroying your signals." Also, Priester insists, the major strength of the STD Bus lies in its wealth of I/O boards from multiple vendors, and these boards aren't equipped to communicate at the higher rates.

With I/O implemented in an interrupt-driven architecture, however, higher-frequency operation does provide gains, according to Alderman. "If instead of doing I/O on a polled basis, you make it all interrupt-driven," he says, "the processor's not looking at the I/O bandwidth but at the bandwidth of the interrupters. That way, though the I/O isn't operating any faster, the system looks a whole lot faster."

Beverly's judgment of the Matrix effort is unequivocal. "It solves all the ground shift and metastability problems we've seen and significantly improves the reliability of the bus," he says. "It represents a real contribution to the STD Bus, and we'll recommend it exclusively for any system over 16 slots."

Beverly admits, however, that a high-frequency backplane may not give STD an immediate performance lift. For those processors that slow down to talk to the bus, he says, "It'll only add a lot more system reliability. But it does provide a future performance upgrade path and if you do a CPU board today, you'll want to do it with a mapping ROM so that if you want to run full speed later on, it'll require only a PAL change."

"The STD is a mature bus, and it's hitting the wall because of performance limitations," concludes Alderman. "We're trying to open up a new realm of performance on the STD for manufacturers and their customers so they can migrate on to the 80286 and 80386SX level in the PC arena, run their equipment in the 10-, 12- and 16-MHz range, and increase throughput from the XT to the AT level." □

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

Computer Dynamics
(803) 877-8700 ..... **Circle 301**
Matrix
(919) 833-2000 ..... **Circle 302**
Ziatech
(805) 541-0488 ..... **Circle 303**

LSI LOGIC ACROSS THE BOARD.SM

When historians write the history of computing, the above five words may well deserve a chapter of their own.

With them, LSI Logic becomes the first company to provide an integrated design methodology for combining ASIC and RISC technologies.

Using our new approach, manufacturers can produce systems that are smaller. More powerful. More flexible. More reliable. And less expensive. Than ever before.

Computers that can be personalized to perform specific functions. And that can be changed as quickly as markets change.

LSI Logic Across The Board.

It's more than just a clever title. It's the beginning of a whole new chapter.

LSI LOGIC

©1989 LSI Logic Corporation

CIRCLE NO. 20

COMPUTERS AND SUBSYSTEMS

SCSI-2 controller board builds parallel disk drive arrays

David Lieberman, Senior Editor

A new disk-drive controller making its debut this month paves the way for system designers to configure very high capacity storage subsystems with very high data-transfer rates and a degree of fault resistance. The Rimfire 6600 SCSI-2 Parallel Drive Array controller from Ciprico (Plymouth, MN)—which, two months ago, announced the industry's first SCSI-2 host adapter—achieves this feat by simultaneously managing as many as five ESDI disk drives and combining their outputs into 20-Mbyte/s transfers over a "wide SCSI" bus.

In addition to the speed advantage of having multiple data channels operating simultaneously, the parallel drive approach offers a cost advantage for systems requiring substantial capacities. "Unlike many advances in disk technology that increase performance, a parallel drive array doesn't require technological advances in disk or controller hardware," says Bill Moren, Ciprico product manager.

Parallel array controllers don't, for instance, require more advanced, more expensive hardware components than do single-channel controllers; they just require more of them. And though parallel array controllers will cost more than controllers that can transfer data to and from only one disk drive at a time, they will cost less at the subsystem level on a cost-per-Mbyte basis.

The Ciprico 6600 manages five garden-variety ESDI drives, designating four for user data and one for parity information, which the board automatically generates. The board transparently distributes data files among the four drives on a byte-by-byte basis. "As far as system software is concerned, there's just one big disk drive," says Moren.

The board dedicates one 80186 microprocessor and associated memory to coordinate the activity among the different drives, manage a SCSI command queue, apply the familiar command-optimization techniques of the company's earlier controllers and handle an on-board 512-kbyte segmented cache buffer. An 80186, scratchpad RAM and a National 8466 serializer/deserializer chip are dedicated to each attached disk drive. A Western Digital SCSI controller chip manages arbitration and other housekeeping chores for the SCSI bus; discrete logic manages the extended 16- or 32-bit transfers defined by the SCSI-2 spec, as well as other SCSI-2 features that are foreign to the SCSI-1-vintage Western Digital chip.

With each write to the disk drive subsystem, the Ciprico board updates the parity information on the fifth drive. If one data drive fails, the board uses that information plus the data on the other three data drives to regenerate complete user data.

Despite a drive failure, the system can still write to and read from the subsystem as if no crash had occurred—though one-fourth of the data won't be immediately available. The 66000 stops short of supporting a "hot spare" capability for the subsystem, but provides several optional methods for reconstructing data. A good drive can later be installed in place of the failed drive, and the controller will place the appropriate data on the new drive.

Not quite fault tolerance

Although provision for a parity drive doesn't provide true fault tolerance, which requires redundancy in all system components, it does provide a good deal of fault resistance. "It's a cost-effective way to increase reliability, but the key feature is data availability," says Moren.

The three data-regeneration modes provided by the controller represent trade-offs between regeneration speed and subsystem response. With direct mode regeneration, the quickest method, the controller immediately dedicates itself to rebuilding data, blocking access to the subsystem until it's finished. With background mode regeneration, the controller continues to service user requests as a high-

(continued on page 36)

The evolution of disk drive arrays

Back in the early 1980s, when thin-film heads and disks began showing up on 5¼-in. Winchester disk drives and extraordinary capacities suddenly began looking feasible, designers started discussing the concept of ganging arrays of multiple small disk drives. By running arrays of small drives in parallel, synchronizing their spindles, and splitting data files up among them ("striping"), the argument went, it would be possible to build faster, more capacious and less expensive storage subsystems than was possible with larger, 8- and 14-in. disk drives.

Parallel-drive advocates also stressed the reliability advantages of disk drive arrays. Since a user's data would reside on multiple drives, they said, a drive failure wouldn't cause the complete and catastrophic loss of data that would occur if a subsystem with a single disk drive went down. Also, since the drive array concept allows for the use of a parity drive, the data on a drive that crashes can be easily reconstructed from the information stored on the parity drive.

The primary motivation behind drive arrays, however, is fast data-transfer rates. The past few years have seen a variety of schemes for enhancing the transfer rate of Winchester drives and subsystems, including multichannel parallel transfer drives from such companies as Fujitsu America (San Jose, CA), Imprimis (St. Paul, MN), Century Data Systems (Anaheim, CA) and Ibis Systems (Westlake, CA); parallel transfer drives and/or drive-array subsystems from Storage Concepts (Irvine, CA) and Pacstor (Los Gatos, CA); and, recently, a $20,000-plus drive array controller box from Maximum Strategy (San Jose, CA).

The just-announced $6,000 Rimfire 6600 SCSI-2 Parallel Drive Array controller from Ciprico (Plymouth, MN), however, appears to be the first offering that lets integrators cost-effectively roll their own subsystem and link it to their system over an industry-standard interface.

CROSS DEVELOPMENT SOFTWARE THAT'S MAKING NETWORK NEWS.

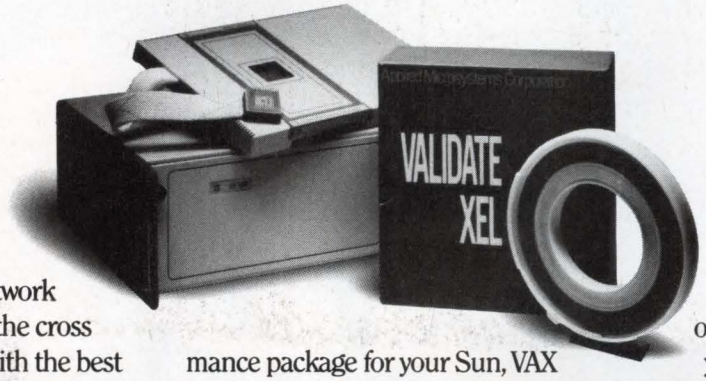

And the news is network pricing. Validate®/XEL, the cross development software with the best support in the business, has network licensing. So now you can have all the cross development power you need for all the engineers who need it. Without paying for individual licenses for each workstation.

With Validate/XEL your development team can stop playing musical chairs. And contribute more in less time for less money.

Validate/XEL is a high performance package for your Sun, VAX or PC. It includes a source level debugger, a compiler and an assembler that work with either a simulator or emulator. And whenever you need help, our large staff of qualified field and in-house applications engineers is at your service.

When you use Validate/XEL with one of our real time, transparent emulators, you'll be even more efficient and productive. You'll have the ability to debug your code inside the target system in real time. And with our high speed SCSI interface to your Sun or PC, you'll be able to download to your target system in record speed.

For more information or a demonstration on your network, call toll-free. In WA call (206) 882-2000. Or write Applied Microsystems Corporation, P.O. Box 97002, Redmond, Washington, USA 98073-9702.

Applied Microsystems Corporation

1-800-426-3925

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

CIRCLE NO. 21

Using the new HP 64000 AxCASE

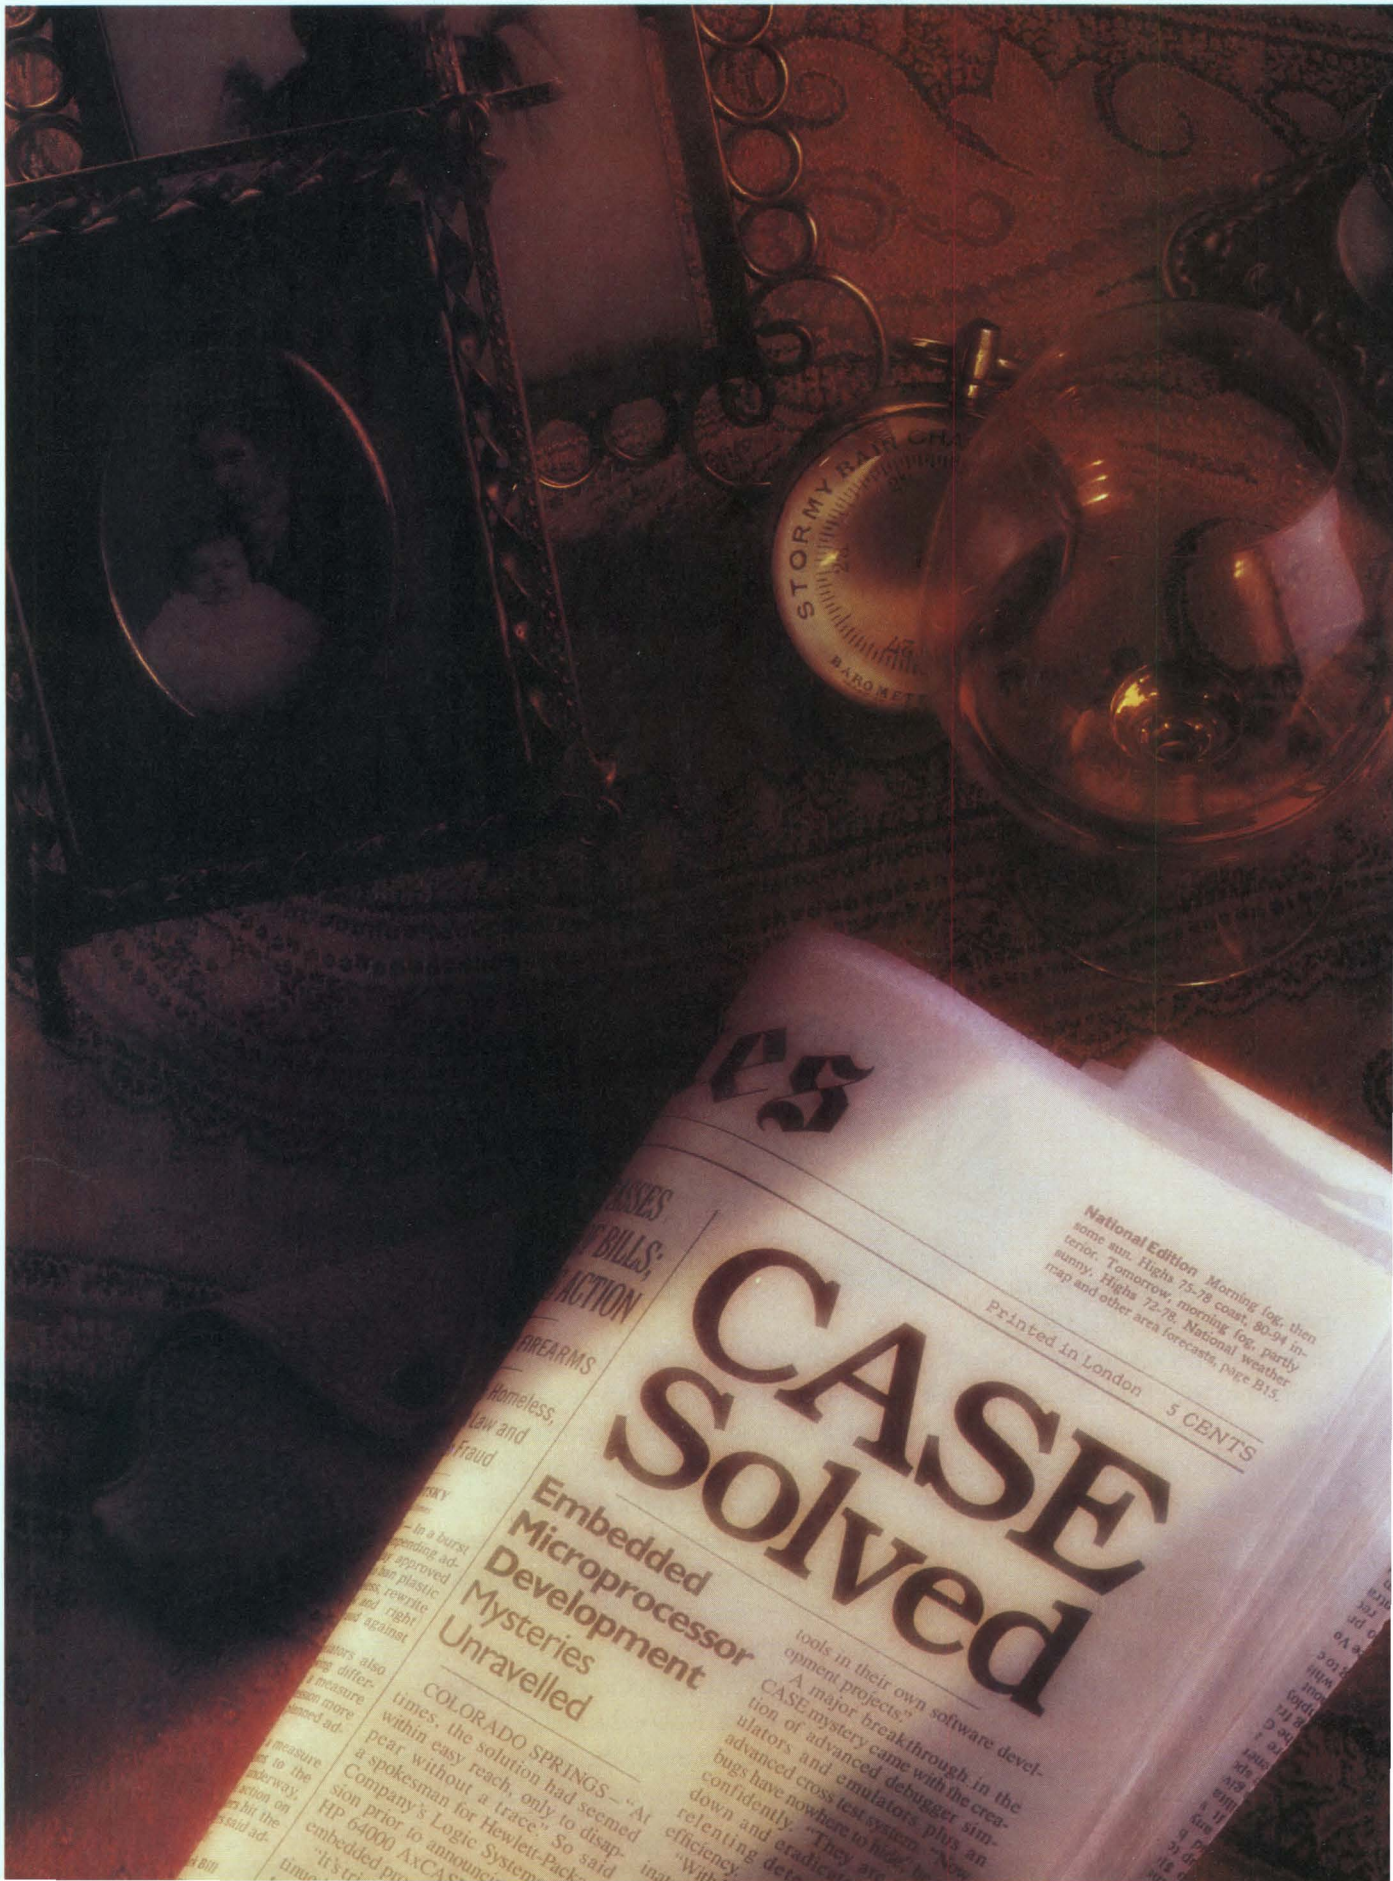

©1988 Hewlett-Packard Co. DS15908/CD

makes μP development elementary.

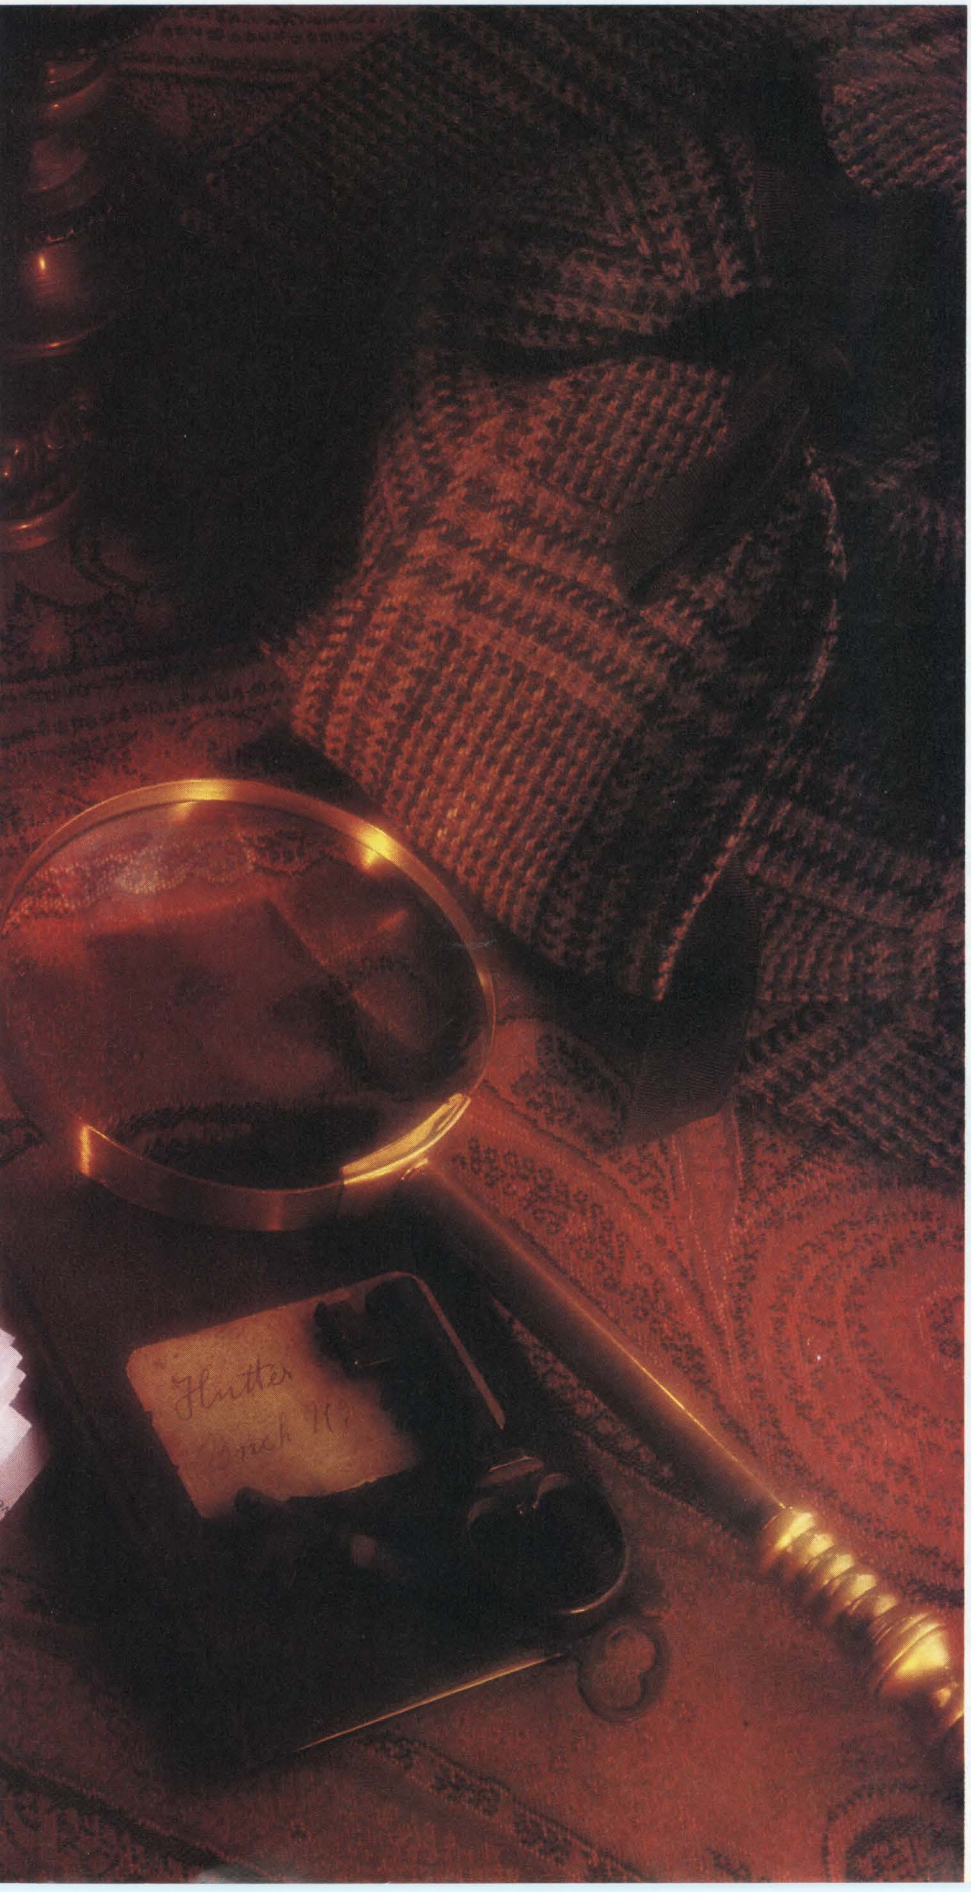

Introducing the first complete environment for embedded processor development. With the HP 64000 AxCASE environment, solutions to even the most complex problems come accurately and predictably. So managing embedded processor software development no longer has to be a mystery.

These powerful hardware and new software design tools put you in charge. Start with HP Teamwork's structured analysis and design tools for a solid definition and description of the software project. It'll help you uncover most defects right at the start. Instead of in final test.

Develop code with AxLS, our advanced, optimizing C cross compilers. Execute and debug on AxDB–an advanced-cross simulator debugger. Evaluate performance and locate system defects in real time with the latest emulator debugger. Finally, verify software testing strategy with the innovative AxTS branch analysis.

These solutions create an integrated, yet open system with standard interfaces that give you compatibility with other software tools and languages–even Ada. And it's backed by HP's worldwide support network. It's precisely what you need to be competitive in embedded processor development.

Call for our HP 64000 AxCASE Solutions Kit at **1-800-752-0900,** ext. 786U, and reach your own flawless conclusions. Better still, call to see it firsthand–after all, it's available now, so put it to work.

HEWLETT PACKARD

CIRCLE NO. 22

COMPUTERS AND SUBSYSTEMS

PRICE/PERFORMANCE COMPARISON

	Parallel Disk Array	IPI-2	SMD
Unformatted capacity	1.5 Gbytes	1.2 Gbytes	1.2 Gbytes
Burst transfer rate	10 Mbytes/s	6 Mbytes/s	3 Mbytes/s
Sustained transfer rate (80/75 percent format efficiency, IPI/SMD and PDA respectively)	7.5 Mbytes/s	4.8 Mbytes/s	2.4 Mbytes/s
Average seek time (typical)	16 ms	16 ms	16 ms
Average rotational latency (typical)	8.3 ms	8.3 ms	8.3 ms
Drive cost (number of drives)	$6,500 (5)	$7,000 (1)	$5,500 (1)
Controller cost	$4,000 (including host adapter cost)	$3,000	$1,500
Total cost	$10,500	$10,000	$7,000
Cost/Mbyte	$6.91	$8.09	$5.67

Source: Ciprico

A 10-Mbyte/s synchronous SCSI-2 subsystem (achievable by configuring either a 5-MHz, 16-bit interface or a 10-MHz, 8-bit interface) beats an IPI-2 subsystem in both performance and overall cost per Mbyte. A SCSI-2 subsystem also provides more than three times the performance of SMD at only about a 22 percent cost premium.

priority activity, regenerating the data as a background operation. With interleaved mode regeneration, the controller services requests and rebuilds data on a time-slice basis at a user-selected ratio. □

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

Century Data Systems
(714) 999-2680 ..... **Circle 304**
Ciprico
(612) 559-2034 ..... **Circle 305**
Fujitsu America
(800) 523-0034 ..... **Circle 306**
Ibis Systems
(818) 706-2505 ..... **Circle 307**
Imprimis
(612) 936-6270 ..... **Circle 308**
Maximum Strategy
(408) 729-1526 ..... **Circle 309**
Pacstor
(408) 395-8804 ..... **Circle 310**
Storage Concepts
(714) 852-8511 ..... **Circle 311**

ADVANCED COMMUNICATIONS PRODUCTS

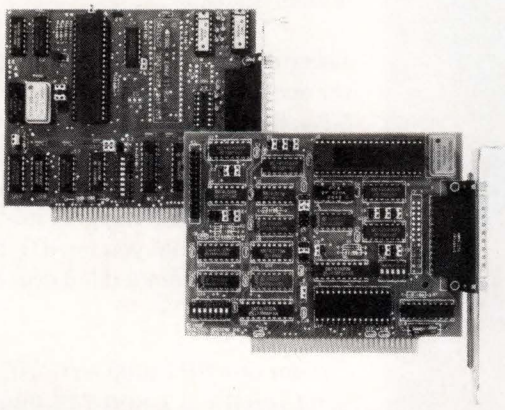

The Sealevel ACB-II and ACB-III Advanced Communications Boards provide SYNC/ASYNC high speed communications capabilities to PCs. Features include:

- SDLC/HDLC/ASYNC Protocols
- RS-232 and RS-422 Interfaces
- One or two channels with DMA
- Software Toolkit Provided

SEALEVEL COMMUNICATIONS & I/O

SEALEVEL SYSTEMS INC.
PO BOX 1808
EASLEY, SC 29641
(803) 855-1581

CIRCLE NO. 23

1+1 = CUPL™

PLD Logic Development System

The first device independent PLD compiler with the choice of schematic entry, state machine, or boolean equation design formats. With power tools such as simulation and minimizations. Supports most PLD manufacturers.

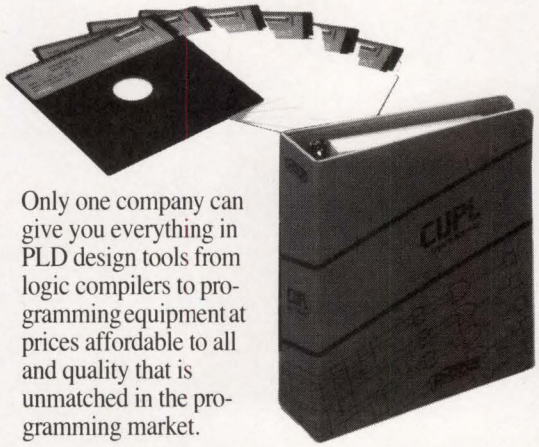

Only one company can give you everything in PLD design tools from logic compilers to programming equipment at prices affordable to all and quality that is unmatched in the programming market.

LOGICAL DEVICES, INC.
ELECTRO/89 #465 & #467

1-800-331-7766
305-974-0967
Telex: 383142
Fax: 305-974-8531

CIRCLE NO. 24

New twist in LCD designs makes them thin, light power-savers

David Lieberman, Senior Editor

The latest innovation in liquid crystal display (LCD) structure promises to drastically cut the cost, depth, weight and power drain of today's state-of-the-art displays. Developed by the Electron Tube Division of Hitachi America (Schaumburg, IL), the new structure essentially replaces one of the two LCD glass sandwiches of the popular double-twist LCD structure with an ultrathin polymer film that performs the same compensating function.

It's becoming more than a little difficult to keep the various LCD "twists" in perspective and, for both LCD makers and manufacturers of portable computers and instruments, to keep up with the pace of market change. The twisted-nematic LCD took the portable marketplace by storm in 1984. By 1986, though, the twisted-nematic LCD was being pushed out of the spotlight by the supertwisted display and its superior optical quality.

Mid-1988 then saw the double-twisted LCD take center stage as its black-on-white display capability and superior contrast attracted those with aesthetic objections to the blue or yellow tinge of the supertwists. Now, a mere six months or so later, Hitachi is aiming to shake the market up yet again.

The key to the advantages of the new structure lies in the elimination of the secondary LCD sandwich (polarizer, glass, liquid crystal material, glass, polarizer). Double-twists use the secondary sandwich to recapture certain components of the light spectrum that are lost in the primary LCD sandwich.

Despite the enhanced readability that the second, compensating LCD brings to the show, however, its contribution to light attenuation has made backlighting, hitherto an option, an essential component of a double-twist display module. In fact, the familiar thick-film electroluminescent backlight has been judged to provide insufficient output for optimal double-twist readability, and more-intense cold cathode tube (CCT) fluorescent backlighting has become the rule.

Less light attenuation

The polymer film used as the compensating layer in the new Hitachi structure attenuates light far less than a second LCD sandwich does. "Each of the polarizers in each of the two LCDs in a double-twisted LCD may be 40 percent transmissive," says David Ross, LCD product marketing manager at Hitachi. "Taken all together, that sucks up a lot of the brightness put out by a backlight. Our film layer, in contrast, may be more like 60 percent transmissive, and it's unnecessary to have a second set of polarizers in there."

By using a polymer film instead of an LCD sandwich, the Hitachi structure is considerably lighter and more shallow than a conventional double-twist—a real boon to portable design. Further, the structure's ability to make use of an EL backlight, or no backlight at all, also saves weight and depth—and power—in comparison to CCT-equipped double-twists.

As for optical qualities, the decreased depth of the new LCD structure provides a better viewing angle than that of typical double-twists, but the structure loses a bit in terms of contrast ratio. The benefit of using a thinner, lighter, less power-hungry EL backlight, however, may more

A NEW LCD STRUCTURE

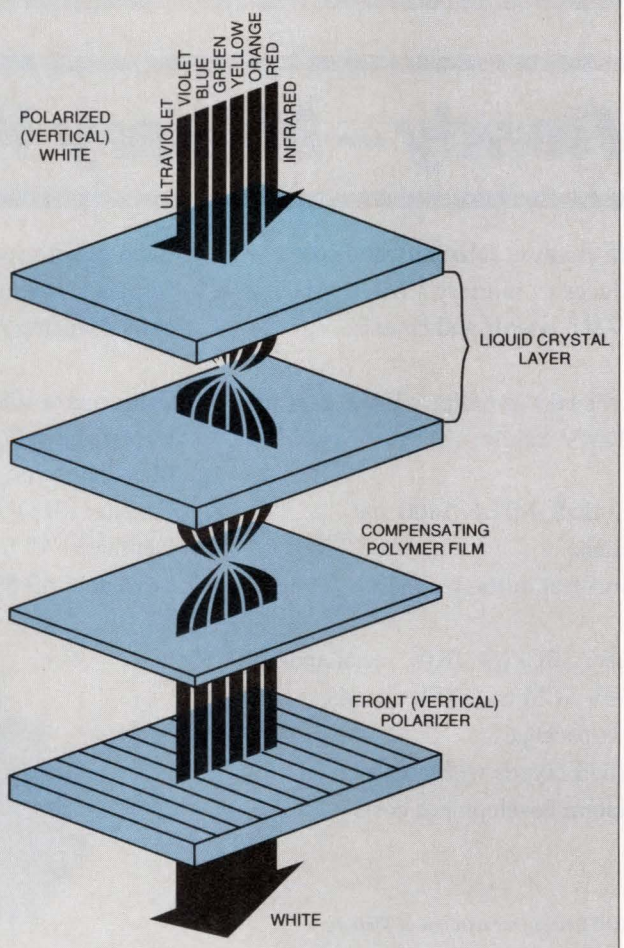

By replacing the secondary LCD sandwich—used as a compensating layer in conventional double-twist LCDs—with an ultrathin polymer film, the new Hitachi structure reduces depth, weight and light attenuation. Reduced attenuation decreases backlighting requirements so a planar electroluminescent lamp can be used, further reducing depth and weight.

MAJOR SYSTEM COMPONENTS

than compensate for a slight loss in contrast.

The new structure also brings cost savings to the display arena. Atypically, the product in which the structure makes its debut is already less expensive than the products it aims to replace. Hitachi's first offering, a 10-in.-diagonal display with a 640-× 480-pixel resolution, will go for $300 in OEM quantities, somewhat less than the $330 for a double-twist of comparable size and resolution, but somewhat more than the $275 for a comparable supertwist.

Ross expects the materials savings and simplified manufacturing processes of the new displays to result in a greater price advantage in days to come. No manufacturing changes are required to create the LCD sandwich for the new structure, since it's basically a garden-variety yellow-mode supertwist display, and the polymer layer is simply integrated into a display module with the same basic techniques used to integrate a polarizer. "Because less exacting tolerances are required when the film and glass are aligned during manufacturing," says Ross, "we anticipate higher throughputs and yields, as well as lower unit pricing." □

DOUBLE-TWIST LCDs

	Film Structure	Conventional
Pixel Resolution	640×480	640×480
Contrast Ratio	12:1	15:1
Viewing Angle	±60°	±45°
Weight	800 g	1,600 g
Response Time	250 ms	250 ms
Power Consumption with inverter circuit	3 W	7.5 W
Module Thickness	14 mm (max.)	28 mm (max.)
Module Size	270 mm×190 mm	295 mm×210 mm

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

Hitachi America, Electron Tube Div.
(312) 517-1144 ..... **Circle 300**

Mission Ready – Military VME from DY-4

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

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

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

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

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

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

dy4®
DY-4 SYSTEMS INC.

CIRCLE NO. 25

FLUKE AND PHILIPS – THE GLOBAL ALLIANCE IN TEST & MEASUREMENT

FLUKE®

PHILIPS

PHILIPS

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

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

UNRESTRAINED PERFORMANCE

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

SIMPLE OPERATION PLUS PC

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

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

UNRESTRAINED SUPPORT

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

TEST THE DIFFERENCE

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

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

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

PM 3655 • 100 MHz • LOGIC ANALYZER

FLUKE®

CIRCLE NO. 26

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

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

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

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

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

IBM is a registered trademark of International Business Machines Corporation.

MITSUBISHI ELECTRONICS

Applications-Driven Microcontrollers

CIRCLE NO. 27

MAJOR SYSTEM COMPONENTS

Digital paper emerges as low-cost archival-storage option

Tom Williams, Senior Editor

One gigabyte on a floppy diskette. One terabyte on a reel of tape. Those are the very realistic promises of digital paper, a flexible, write-once optical media developed by ICI Imagedata (Welwyn Garden City, UK) a subsidiary of the English chemical giant, ICI. As a medium, digital paper offers not only low cost-per-megabyte, high data density and flexibility, but other attractive properties as well. Chief among these is the ability to use less laser power to write data than is possible with other optical media, resulting in lower-cost storage systems and faster access times and data rates.

At least two companies have seized on the potential of digital paper and are developing storage products in cooperation with ICI. Bernoulli Optical Systems (Boulder, CO), a subsidiary of Iomega (Roy, UT), is developing a disk drive that draws on Bernoulli technology developed by Iomega and used in Iomega's high-density magnetic floppy drives. And Creo Products (Burnaby, BC) is developing an optical tape drive that's slated to store 1 Tbyte of data on one 12-in. reel of 35-mm-wide digital paper tape.

"Digital paper" is, of course, a misnomer. The material is based on a polyester film substrate, called Melinex, that's coated with a very thin sputtered layer of metal to make it reflective. But here the similarity with other optical media ends.

In rigid optical media, the laser burns a pit into the reflective metal and the difference in reflectivity between pits and nonburned areas is read as data. This process requires considerable energy from the laser since metal conducts heat well, and the amount of laser power required limits the speed at which the disk can be rotated, thereby limiting the possible data rate. Since the head-to-media distance isn't precise, an active focusing system is needed to keep a constant spot size. The addition of active focusing elements adds to the mass of the head assembly, limiting the actuator speed and, hence, access time. The characteristics of digital paper let changes in peripheral mechanisms overcome some of these limitations.

Polymer dye layer coats metal

ICI has developed a polymer dye coating that's applied to the mirrored surface of its medium. The dye is composed to absorb precisely at a wavelength used by infrared lasers. One version of the dye absorbs at 780 nm, and another at 830 nm. It's this polymer dye layer that's deformed by heat from the write beam.

In contrast to conventional write-once, read-mostly (WORM) media, material isn't blasted from a pit. Depressions are formed by a process called "pyroplasticity," says Michael Strelitz, ICI's marketing manager for data storage products. Pyroplastic deformation forms pits that don't strew debris out around them, but form small lips at their edges.

The thickness of the dye polymer layer is controlled such that it's a multiple of half-wavelengths of the laser light (for example, $n \times \lambda/2$, or in the case of the 780-nm material, $n \times 390$ nm). When the laser read beam passes through an unburned portion of the polymer dye layer, the part of the beam that's reflected by the dye layer is in phase with the part that reflects from the metallic layer. When the read beam hits a depression in the dye layer, it finds that the thickness of the layer is no longer a multiple of half-wavelengths. The components of light reflected from the dye surface and the metal surface are no longer in phase and tend to cancel, resulting in a difference that can be read as data.

As a good thermal conductor, metal tends to let heat flow away from the focus point, so melting pits in the polymer dye material requires much less laser energy than blasting pits into a layer of metal. The ability to keep heat concentrated at the focal point increases the time/power product that must be applied to that point to burn a hole. The difference in thermal characteristics between the two materials, combined with Iomega's Bernoulli disk technology that lets the heads fly much closer than other optical media permit, allows substantially higher data rates than with WORM technology that uses a rigid disk. The closer the heads are to the media, the more tightly the beam can be concentrated and the less energy is required.

Air flow used to stabilize media

The digital paper drive under development by Bernoulli Optical Systems uses the air flow created by the media spinning next to a solid plate,

HOW DIGITAL PAPER WORKS

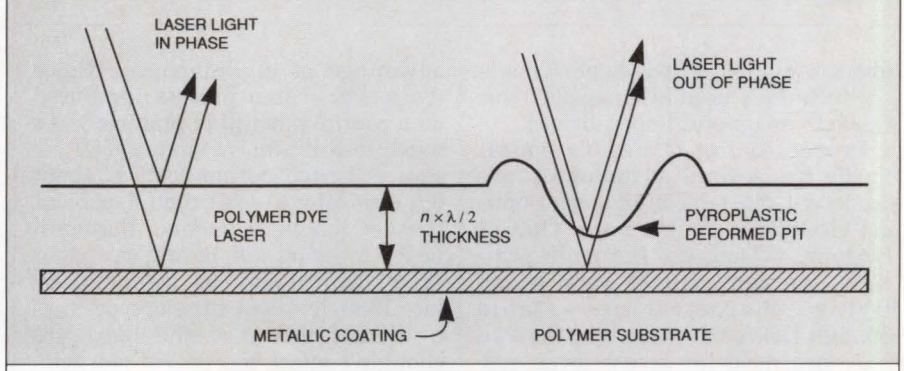

In digital paper, a layer of polymer dye deposited over a mirrored surface is tuned to absorb a specific wavelength of laser light. Components of light pass through a thickness that's a whole number of half-wavelengths. These components reflect in phase with light components reflecting from other surfaces. Light striking a deformed area is thrown out of phase with its other components and tends to cancel. The light can then be read as a data transition.

MAJOR SYSTEM COMPONENTS

called the "Bernoulli plate," to stabilize the media and fly it at an exact height from the head. As the air flows around the airfoil-shaped head sticking through the Bernoulli plate, it causes a small dimple in the medium around the head. The dimple stabilizes the shape of the flexible material and holds the surface about 1 micron off the head. This is the technique used in Iomega's high-capacity magnetic floppy drives.

Precise control of flexible optical media has several advantages, says Andrew Goldstein, manager of optical electronics for Bernoulli Optical Systems. "Using the Bernoulli principle, we can achieve very stable flying characteristics—stable enough that early signs are that we may not need an active focus system," he says. The exact, close head-to-disk interface will thus let Bernoulli cut down on the mass of the heads, allowing access times of less than 40 ms, Goldstein says.

The power efficiency, combined with the close flying height, lets a 10-mW laser write data at about 10 MHz, which in turn permits a rotational speed of 1,800 RPM. That translates to a data transfer rate of 1.5 Mbytes/s, compared with the 0.15- to 0.25-Mbyte/s transfer rates of conventional WORM drives.

Goldstein also mentions another advantage of digital paper. "Since it's a web-coated process (produced on a continuous roll of material), it's much less expensive than a CD," he says. ICI cites storage costs at about 0.5 cent/Mbyte—less than 1 percent that of magnetic media. Bernoulli hopes to bring out its digital paper drive, which may be sampling by late 1989, at about the same price as a WORM drive, so the hardware shouldn't affect the cost of storage.

BERNOULLI DISK TECHNOLOGY

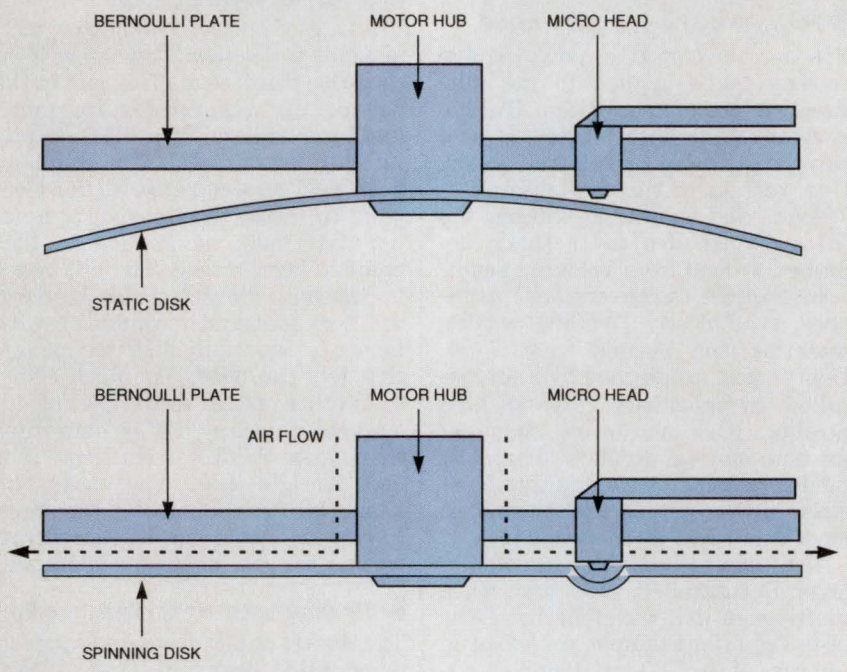

Bernoulli technology stabilizes flexible media by using the air flow around the specially shaped read/write head to form a small dimple with exact spacing from the head. The flowing air also makes the head/disk interface very resistant to particle contamination. Particles disturb the air flow just enough to pass though the gap without contacting the head or the media.

An optical tape drive

The other digital paper storage product under development, which will probably be the first commercially available, is the Model 1003 optical tape drive by Creo Products. The 1003 uses an 880-m-long, 35-mm-wide tape that fits on a 12½-in. reel. The drive's recording head uses a laser aimed into a movable focusing device that can write 32 bits across the tape. The data is written as four byte-wide words in 20-kbyte blocks, or blocks containing 80 kbytes. On the 880-m tape, the drive can access any byte in 28 s and boasts a 3-Mbyte/s data transfer rates.

ICI developed the 35-mm tape for Creo and produces the disk-format media for Bernoulli. In the case of the disk media, servo information is embossed in the form of grooves into the polymer dye layer. The grooves return an error signal to the heads, which is fed to the servo mechanism. At this point, the track spacing is 1.6 microns, the same as other optical media. But "the limits of capacity are defined basically by the lasers one can use," says ICI's Strelitz. The current spot size is about 1 micron, which is a function of the wavelength of the laser. Goldstein adds that the abilities of the servo system also can limit capacity.

Another important factor is the sharpness of the edges that define the pits. "Depending on the sharpness of that, you can use higher-order encoding schemes such as run-length-limited (RLL) coding," Strelitz says. RLL depends on the ability to burn pits of a precise length and on reading to exactly measure the length of pits in the time domain. It therefore requires precisely defined pit sizes—something much more difficult to achieve with ablated holes that eject material around their edges and thus don't have sharp transitions. Pits formed by pyroplasticity have small lips at their edges but don't eject debris that can blur the transition between written and nonwritten areas. □

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

Bernoulli Optical Systems
(303) 939-8611 ..... **Circle 321**
Creo Products
(604) 437-6879 ..... **Circle 322**
ICI Imagedata
(0707) 337284 ..... **Circle 323**
Iomega
(801) 778-1000 ..... **Circle 324**

G300 CVC • PROGRAMMABLE VIDEO TIMING GENERATOR • COLOR LOOK-UP TABLE

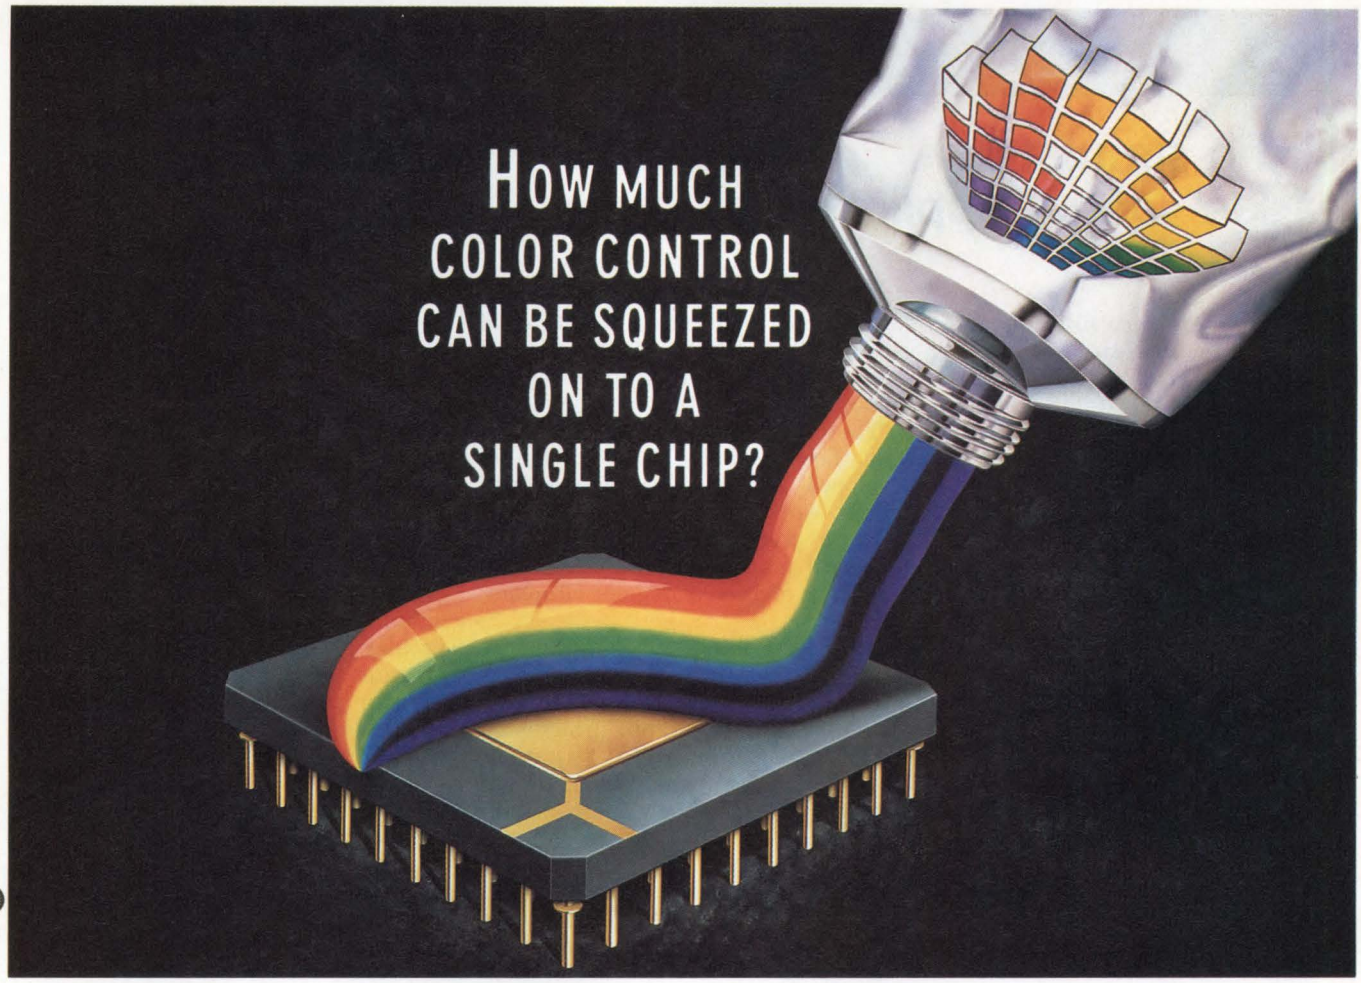

INMOS and IMS are trademarks of the INMOS Group of Companies.

FULL BITMAP MANAGEMENT • TRIPLE 8-BIT DAC • PHASE-LOCKED LOOP • G300 CVC

If the chip in question is the new IMSG300 Color Video Controller (CVC) from INMOS, the answer is quite simply, "all of it".

Through high integration, INMOS has combined total graphics control and data conversion onto a single 84 pin device. This includes a color look-up table, programmable video timing generator, full bitmap management, triple 8-bit video DACs and a phase-locked loop.

Because the IMSG300 is programmable, it is easily interfaced to any processor or monitor that may be used in your system. The IMSG300 promotes the technique of using software drawing algorithms as opposed to designing them in hardware. This dramatically reduces design-in time and allows for easy upgradability.

The IMSG300 also provides an optional 24-bit pixel mode for full color as well as pseudo color, an on-chip phase-locked loop for generating all high frequency clocks, and a 32-bit multiplexed pixel port for achieving pixel rates up to 110MHz.

Only the INMOS G300 CVC gives you all the options you need to tailor your systems exactly the way you want them – both now and in the future.

The INMOS G300 CVC single chip solution. Find out more today.

INMOS, PO Box 16000, Colorado Springs, Colorado 80935. Telephone: (719) 630 4000

inmos

Baltimore (301) 995 6952 • New York (914) 897 2422 • Denver (303) 368 0561 • Santa Clara (408) 727 7771
Atlanta (404) 242 7444 • Dallas (214) 490 9522 • Minneapolis (612) 932 7121 • Boston (617) 229 2550 • Orange County (714) 957 6018

CIRCLE NO. 28

DESIGN AND DEVELOPMENT TOOLS

Niche simulators tackle special applications

Bill Harding, Senior Editor

Simulators from full-range electronic design automation (EDA) vendors work well for the majority of electronic design applications. But some very specialized applications aren't adequately covered by the general-purpose simulators that most EDA vendors offer. Manufacturers are proposing various solutions to accommodate these new types of applications.

As technology becomes more complex, more applications that fall outside the general design cycle are likely to appear. "I see a lot of fragmentation in the simulator market today, and I see it increasing in the future," says L. Curtis Widdoes, president of Logic Modeling Systems (San Jose, CA). "Simulators are complex programs with large feature sets. But no one simulator can be the best in every feature. There are hundreds of features, and different customers want different features."

The first evidence of this fragmentation can be found in the simulators offered by the full-range EDA vendors themselves. Companies such as Daisy Systems (Mountain View, CA), Mentor Graphics (Beaverton, OR) and Valid Logic Systems (San Jose, CA) offer logic simulators to test the functional characteristics of digital designs, fault simulators to develop test vectors, analog simulators to test the functionality of analog circuits, and thermal-analysis simulators to determine whether a printed circuit board will run without burning up.

All of these simulators are technology-oriented, in that they address the hardware design cycle for ICs and printed circuit boards rather than addressing application algorithms or total systems (including software).

Niche simulators, usually from smaller companies, target specific applications or design cycle problems that larger companies don't adequately address. A niche simulator, or any niche tool, is one designed to solve a specific problem, and is not part of a full line of design tools that address the complete design cycle.

Once an application is addressed by a niche simulator, the larger companies may help integrate the niche simulator with their product lines. "The critical thing for Mentor is to let those niche market tools be developed and then make sure the tools work with ours in the full design process," says Geoff Bunza, general manager of the design and analysis division at Mentor Graphics.

Extensive algorithm development

One application that isn't adequately addressed by the larger EDA vendors is digital signal processing (DSP). "Designing DSP applications is different from most system design tasks," says Michael Walsh, senior vice-president of Comdisco Systems (San Francisco, CA). "One reason is that DSP engineers spend the first half of their design time perfecting an algorithm before they ever start designing system hardware or writing code." Once the algorithm is developed, it may be implemented in hardware, software or a combination of hardware and software.

The DSP algorithm development stage, therefore, isn't a hardware or software development problem, making it a difficult application for technology-oriented simulators. To address this problem, Comdisco Systems recently developed a suite of design tools, called the Signal Processing Worksystem (SPW), that target DSP applications exclusively. The SPW simulator is aimed at that first 50 percent of the DSP design cycle—the algorithm development phase rather than the hardware development phase.

Comdisco's SPW includes fully debugged DSP-oriented models of DSP algorithm elements, relieving DSP engineers of the need to build and debug models. But even if DSP engineers tried to develop their algorithms using a general-purpose simulator, they would likely run into problems. "The Comdisco simulator operates at a different level than Mentor's simulators," Mentor's Bunza says. "Some of our customers use our simulators for DSP applications. They use behavioral models and structural descriptions to model the detailed design, but they simulate at a lower level than the Comdisco simulator would allow them to simulate."

A back-end niche simulator

Niche simulators don't necessarily solve architectural or algorithmic problems at only the front end of the design cycle. They can also address problems that affect the overall design cycle. Cross-talk between traces on a printed circuit board, for example, wasn't a significant problem with system clocks that ran at under

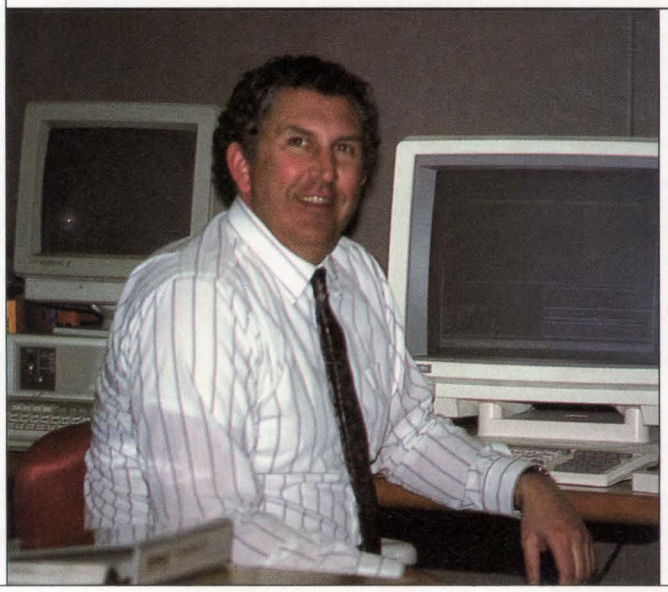

Digital signal processing (DSP) application design is unique because the first half of the design cycle involves developing a DSP algorithm, according to Michael Walsh, senior vice-president of Comdisco Systems. Comdisco developed a niche simulator that targets the DSP algorithm development phase.

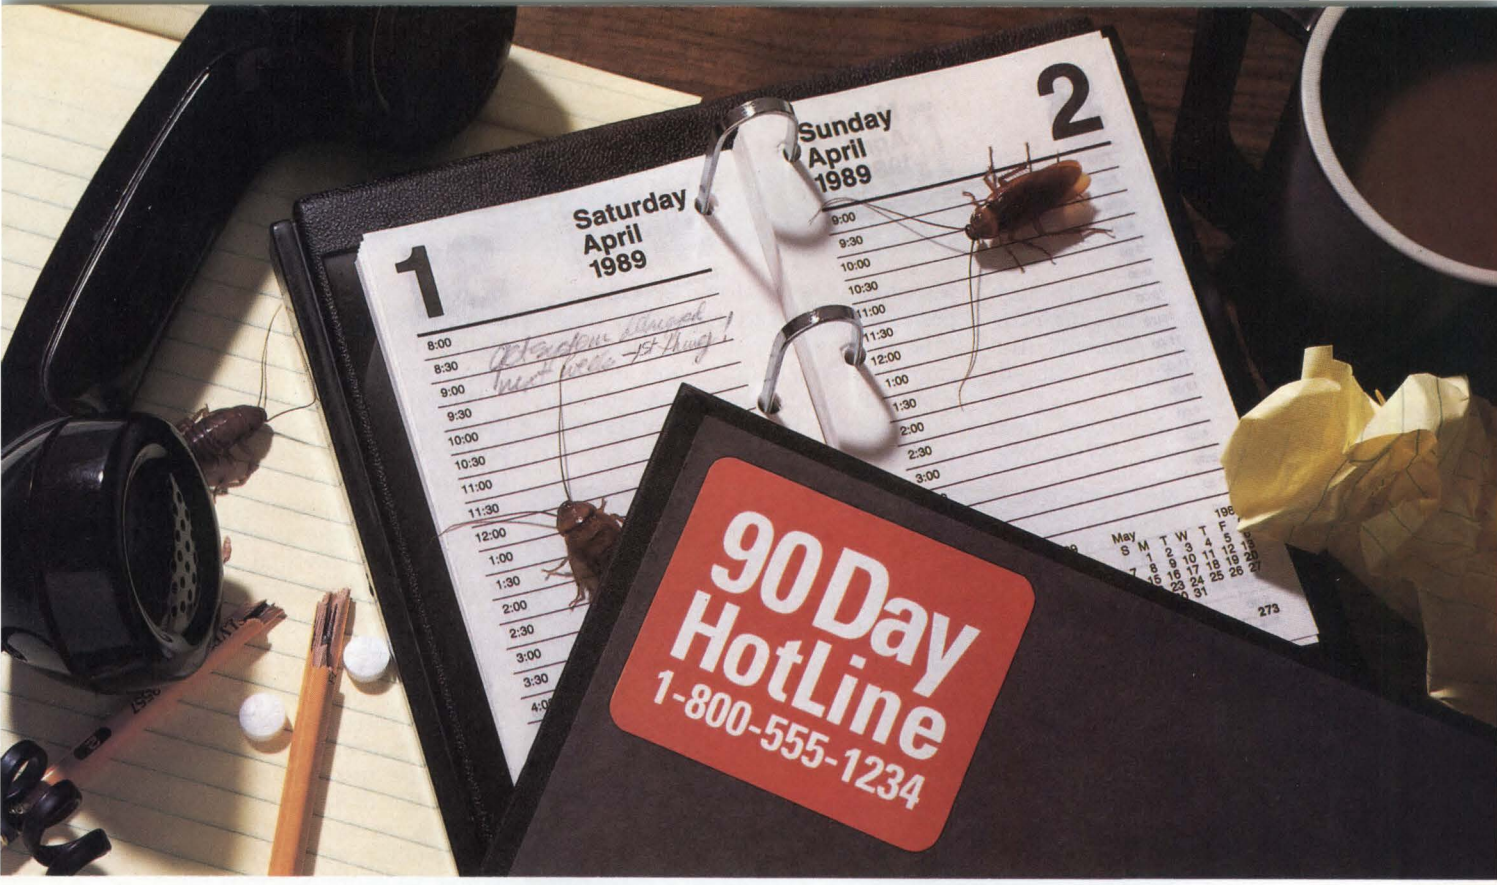

Isn't it funny how bugs appear 90 days after system installation?

It happens every time. You need a realtime operating system, so you evaluate the different vendors, read all the literature, compare the benchmarks, have dinner with the salesmen, listen to the hype. Nearing desperation, you choose one and go to work.

Things go well at first. You get excited. After getting the hang of the system you begin work on your project only to find you have big problems. Bug problems.

You call for help on the 90-day free consulting line only to discover that it's the 91st day. Of course, help is available for an additional fee. Excitement fades into disillusionment as project costs escalate.

Eyring has a solution to your problems – PDOS realtime software tools that are supported for a full year. And we mean real support – like hotline service, free upgrades, tech newsletters and on-line modem help. Training is included with development packages. And if you need us after a year? We'll still be there to assist you (for a very modest fee) as often as you like.

When the time comes for you to choose a Realtime Operating system; choose PDOS products that provide long term solutions from a team of experienced realtime systems engineers. Or get ready for a bug hunt.

1455 WEST 820 NORTH
PROVO, UTAH 84601
TEL: 801-375-2434
FAX: 801-374-8339
©1989 Eyring

EYRING
Creating Solutions™

Hotline help and service for a full year.

Call or write today for a free "Bugbusters" button and info packet on PDOS products.

PDOS Distributors in: Benelux
Interay BV
Lageweg 2A
9251 GM Bergum
The Netherlands
☎ 05116 / 4052

Germany, Switzerland, Austria
Systrix GmbH
Hindenburgring 31
D-7900 Ulm / Donau
West Germany
☎ 0731 / 37515

United Kingdom
Eyrisoft Ltd.
Etwall Street
Derby DE3 3DT
England
☎ 0332 / 384978

PDOS
Realtime Operating System®
PDOS is a registered trademark of Eyring

CIRCLE NO. 29

DESIGN AND DEVELOPMENT TOOLS

10 MHz. But today, a 10-MHz clock is considered slow, and high-frequency cross-talk is a problem. And there's a niche simulator that addresses that need.

The Greenfield simulator from Quantic Laboratories (Winnipeg, Manitoba, Canada) is a special-purpose analog tool that predicts crosstalk, ringing, time delays, parasitics, impedance mismatches and faults of a similar nature in high-speed digital systems. "In high-speed systems, printed circuit board traces can no longer be viewed as simple connections," says Al Wexler, Quantic president. "They must be considered as more complex transmission lines."

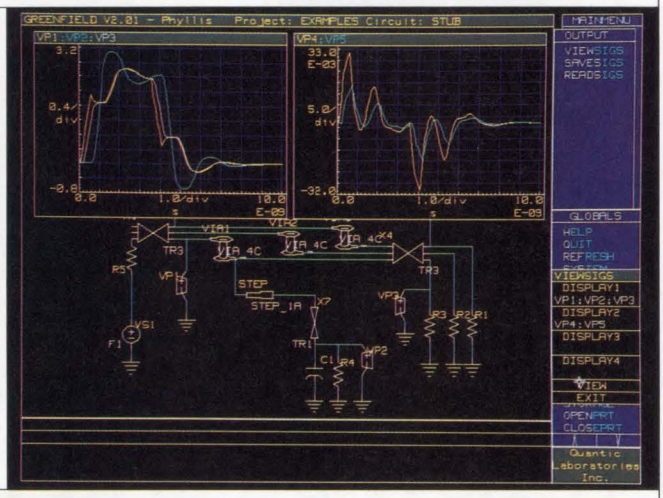

The Greenfield simulator from Quantic Laboratories analyzes transmission line characteristics of high-speed switching circuits to detect cross-talk, ringing, time delays, parasitics and impedance mismatches. Greenfield runs on Apollo, Sun 3, VAX, IBM, Prime, Hewlett-Packard and Cray computers.
Visual Sense

While engineers understand this phenomenon, it's very difficult to predict its effects without a tool such as Quantic's Greenfield, which consists of two programs. The first, Greenfield2, implements the boundary element method for solving printed circuit board transmission line problems. To analyze a board's transmission line characteristics, the user inputs the board's layout plus line width and spacing, the dielectric constant and thickness of the substratum, conductor cross-section shape and conductor placement within the dielectric. This information may be entered manually using Greenfield2's geometric modeler, or it may be contained in a file.

The second program in the Greenfield suite, Phyllis, is a physical load and line simulator that analyzes ringing and crosstalk in the time domain. Greenfield runs on Apollo, Sun 3, Hewlett-Packard 9000 and MicroVAX workstations, and on Prime, IBM 3090, Cray and VAX mainframe computers.

Useful within limits

Just how useful can a niche simulator be to an engineering manager with an applications problem? Opinion varies, but the general consensus is that niche simulators are useful—within limits.

"No engineers really think that they simulate their designs enough, but specialized simulators, targeted at specific applications, help cut down on the number of simulation cycles needed," says Hal Alles, vice-president and general manager of the IC CAD Division of Silicon Compiler Systems (Warren, NJ).

Mentor's Bunza agrees, but with a word of caution: "Such simulators are valuable and appropriate to the people who use them, but they aren't well integrated with the rest of the design cycle."

Bruce Bourbon, marketing vice-president at Cadence Design Systems (San Jose, CA), is more cautious. "The problem with most algorithmic simulators is that they don't link directly to physical implementation tools," says Bourbon. "It would be possible to bring simula-

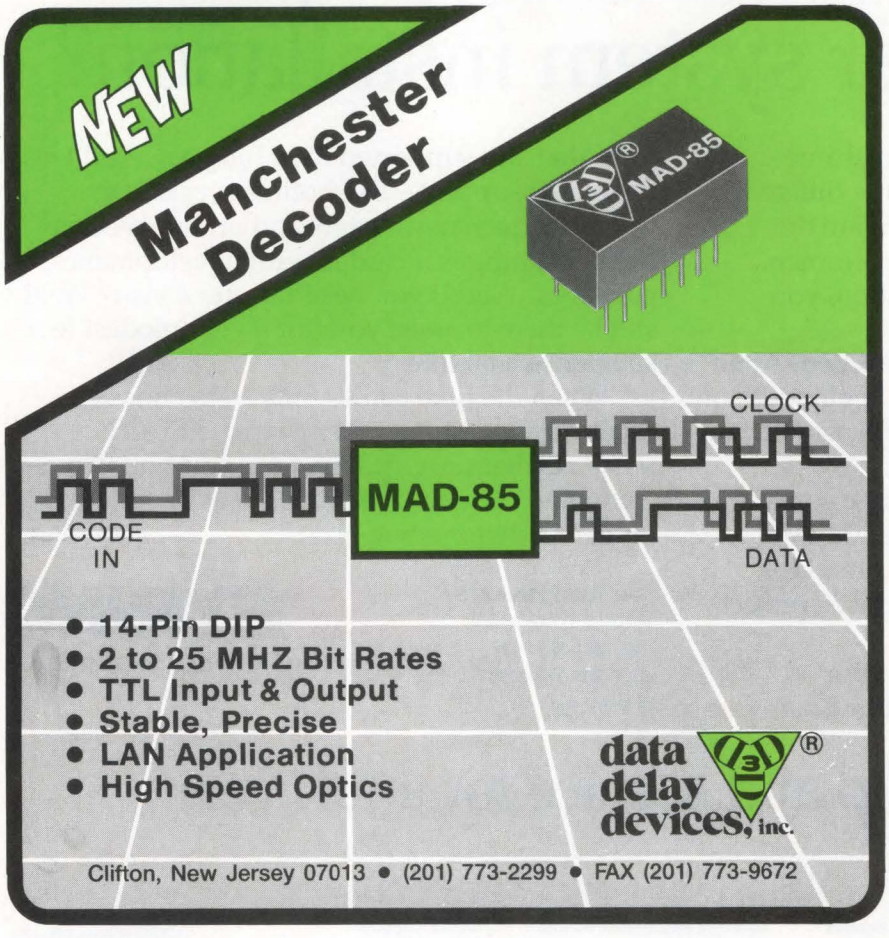

CIRCLE NO. 30

tors such as the Comdisco system under the Cadence framework, but why do that? If the link to the IC design tools isn't there, would the common interface be sufficient reason? I don't think so." The Cadence framework lets tools from various vendors work in the same environment as the Cadence CAE tools, and share a common user interface.

Tailoring a simulator

Not every application or gap in design cycle coverage demands a new simulator. Sometimes a general-purpose simulator can be tailored to meet specific user needs. If an application problem is technology-oriented, a custom model library supporting a general-purpose simulator may solve the problem. Daisy Systems solved a problem in just that manner.

Many of Daisy's customers design systems that must meet U.S. government military specifications. The component libraries that support commercial customers don't take into account the stringent requirements of U.S. Military General Specification for Microcircuits (MIL-M-38510). To let customers create designs that meet military needs, Daisy developed its MIL-SPEC 38510 Component Library. Containing some 600 models, Daisy's MIL-SPEC library works with the company's general-purpose EDA tools.

SCS's IC CAD Division recognized a similar need for customizing general-purpose simulators and developed a custom environment for designing radiation-hardened (rad-hard) ICs. SCS's approach is to provide cell development tools to run under the company's GDT (Generator Development Tools) application-specific IC library development system, and system design tools to run under the Genesil ASIC design system. Expert IC designers will use the GDT rad-hard extensions to develop rad-hard ASIC library cells, and system designers who need to design rad-hard ASICs will use the Genesil tools.

Custom models or simulators?

With the proper simulation models, general-purpose simulators can be very powerful tools. Custom model libraries, if they adequately address an applications problem, offer numerous advantages over niche simulators. A library can expand the scope of a simulator but can't change its characteristics; libraries can thus let a simulator address MIL-SPEC and rad-hard designs but would have difficulty tailoring it to simulate DSP algorithms. Libraries are automatically integrated into the supported simulator's family of EDA tools. Developing a custom library, while not a trivial task, is less demanding than creating a simulator from scratch; a custom library thus is in most cases less expensive than a special-purpose simulator.

On the other hand, a custom library can be a major undertaking for a limited market, so the user may be left with the job of creating custom models. Creating custom models can involve the same problems as creating a breadboard of a new design: if the simulation doesn't work, is the problem in the model or in the design? Also, limitations within the simulator may prevent full support of a given application, as is the case with DSP designs.

Niche simulators offer the distinct advantage of focusing directly on a given application or problem area. The user doesn't have to develop models for a given application since the models come as part of the simulator. If the simulator has been in use for a reasonable period of time, chances are that the models will be reasonably well debugged. In applications such as DSP algorithm development and transmission line analysis, model libraries added to a general-purpose simulator won't address the problem, and a niche simulator is needed. □

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

Cadence Design Systems
(408) 954-7673 ..... **Circle 330**
Comdisco Systems
(415) 421-1800, ext. 2728 ..... **Circle 331**
Daisy Systems
(415) 960-7040 ..... **Circle 332**
Logic Modeling Systems
(408) 922-0870, ext. 32 ..... **Circle 333**
Mentor Graphics
(503) 626-7000 ..... **Circle 334**
Quantic Laboratories
(800) 665-0235 ..... **Circle 335**
Silicon Compiler Systems, IC CAD Div.
(408) 371-2900, ext. 207 ..... **Circle 336**
Valid Logic
(800) 821-9441 ..... **Circle 337**

SPECTRUM

DSP

Development Tools and Integrated Systems for the IBM-PC.

Spectrum's DSP boards have a common expansion interface called **DSP~LINK**, letting you match analog or digital interface boards with the DSP chip of your choice!

DSP~LINK lets you select off-the-shelf interfaces, or quickly design a custom interface using the DSP~LINK Prototyping Module and Application Notes.

System & Processor Boards

Processor Boards have DSP processor, RAM, and DSP~LINK. System Boards have the same plus on-board analog I/O. Prices include monitors, debuggers, sample code, and 'C' libraries.

- TMS320C25 Processor Board $995
  TI's 'C25 with 16 Kwords RAM.
- TMS320C25 System Board $1,995
  50KHz 16-bit I/O, 16 Kwords RAM expandable to 128 Kwords.
- DSP56001 Processor Board $1,495
  Motorola's 56001 with 48 Kwords RAM expandable to 192 Kwords.
- DSP56001 System Board $3,495
  2 channel 153KHz 16-bit I/O, codec, 48 Kwords RAM up to 192 Kwords.
- ADSP-2100 System Board $2,595
  200KHz 12-bit I/O, 16 Kwords RAM.

DSP~LINK Peripherals

Off-the-shelf analog or digital interfaces for all System and Processor Boards.

- 4 Channel Analog Input / 2 Output $845
- 32 Channel Analog Input $845
- 16 Channel Analog Output $845
- Pro-audio (AES, Sony PCM, MIDI) $1,850
- Dual-Processor Comm. Module $95
- DSP~LINK Prototyping Module $95

DSP Development Software

'C' Compilers, Assemblers, and Simulators. Hypersignal Workstation, DADiSP, DISPRO, and Momentum Filter Design.

- LSI TMS320C25 'C' Compiler $995

CIRCLE NO. 31

Call for quantity discounts!

SPECTRUM SIGNAL PROCESSING INC

USA East: 800-323-1842 In Mass: (617) 890-3400
USA West: 800-663-8986 Canada: (604) 438-7266

"The VME world doesn't just need RISC performance.

It needs a total RISC computing environment."

Andreas Schreyer,
RISC Product Marketing Manager

The Motorola 88000 RISC architecture offers tremendous performance. But to serve a complex and well-established market like VME, it takes more. A total environment of hardware and software for both system development and applications. Software compatibility between RISC platforms from many suppliers. And a painless growth path from 68000-based systems to 88000 and beyond.

Creating VME hyper-performance.

To help people integrate the 88000 into their VME systems easily, we invented the HYPERmodule™—a miniaturized subsystem with one, two or four closely-coupled 88000s, cache memories, memory management and high-speed internal buses. Because all HYPERmodules have identical interfaces to the CPU board and identical programming models, you can plug in 17, 30 or 50 MIPS of RISC performance interchangeably. That's instantaneous scalability, with today's 20 MHz 88000. And when the 25 and 33 MHz versions arrive, you'll have up to 80 MIPS without rewriting a line of code. Scalability makes the future safe and sane for developers.

Standardizing 88000 software.

The real strength of our 88000 boards and systems comes from our compatibility in the software environment. 68000 applications can simply be recompiled to run on our 88000 systems. We have the optimizing compilers and software development tools in place. And we've strongly supported the industry standards for both multi-user and real-time 88000 software. Our UNIX™ System 5, Release 3, is compliant with BCS (Binary Compatibility Standard), as well as SVID and POSIX. Our VMEexec™ is RTEID compatible. This industry-wide compatibility has led to massive porting activity by software vendors.

Leading VME to 50 MIPS and beyond.

Today, Motorola has taken VME where it's never been before—to the frontier of 80 MIPS. We've provided a total RISC environment. And we've addressed issues like compatibility, scalability and long-term growth path. That's where Motorola's system experience, investment in VME software and overall attention to quality make a critical difference. That's what being the leader means.

Malcolm Baldrige National Quality Award

MOTOROLA *Microcomputer Division*

Approaching our technology from your point of view.

For reprints of this series, call 1-800-556-1234, Ext. 230; in California, 1-800-441-2345, Ext. 230. Or write: Motorola Microcomputer Division, 2900 South Diablo Way, Tempe, AZ 85282.

HYPERmodule and VMEexec are trademarks of Motorola, Inc. UNIX is a trademark of AT&T.

CIRCLE NO. 32

Pop Quiz.

Stop. This is a test. For the next 60 seconds, we will be conducting a quiz about Macintosh®II Videographics. Do not turn the page until you have looked at the visual clue and answered all the questions.

Which Macintosh II graphics card offers the widest range of capture and display resolutions–NTSC, PAL, Apple® Monitor, hi-res, interlaced, non-interlaced and other modes?

a) NuVista 2M b) NuVista 4M c) All of the above

Name the only videographics card which provides true-color, real-time capture *and* broadcast-quality display while occupying only a single slot in a Macintosh II.

a) NuVista 2M b) NuVista 4M c) All of the above

Which videographics card offers full QuickDraw™ compatibility at 1,2,4,8,16 or 32-bits per pixel?

a) NuVista 2M b) NuVista 4M c) All of the above

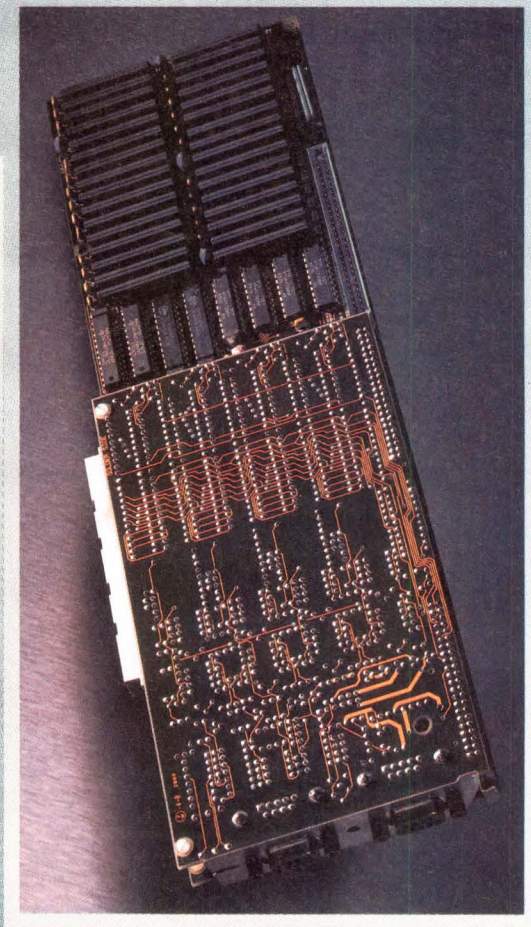

Visual clue for Videographics test.

If you chose (c) on all three questions, congratulations! You know that the NuVista series from Truevision is the answer to all your advanced videographics needs. The NuVista is available with either 2Megabytes or 4Megabytes of video memory, and creates professional video effects and computer graphics using any QuickDraw compatible software, now and in the future. No patches, no gimmicks, no hassles.

So whether your application is video production, digital prepress, presentation graphics or 3D renderings, you'll find the NuVista will pass *your* test with flying colors. Oh, and if you answered (a) or (b) to any question above, give yourself half credit. Then obtain even more NuVista information by requesting a copy of our educational brochure True color? True answers. or visiting your local Authorized Truevision Reseller. Either way, you can find all the answers with a NuVista. Call us at **800-858-TRUE.**

You may now return to your regular reading.

TRUEVISION® INC.

7351 Shadeland Station, Suite 100, Indianapolis, IN 46256
INTERNATIONAL: Canada 416/499-9400 France 33-13-952-6253 Italy 39-2-242-4551 Switzerland 41-1-825-0949 U.K. 44-1-991-0121
West Germany 49-89-612-0010 Other 617-229-6900

Apple and Macintosh are registered trademarks and Quick Draw is a trademark of Apple Computer. Truevision is a registered trademark and NuVista is a trademark of Truevision, Inc. © Truevision Inc. 1989

CIRCLE NO. 33

Emerging standards, hardware and software light the way to FDDI

Within two years, FDDI will be established as the dominant high-end LAN for the 1990s.

Ken Marrin
Contributing Editor

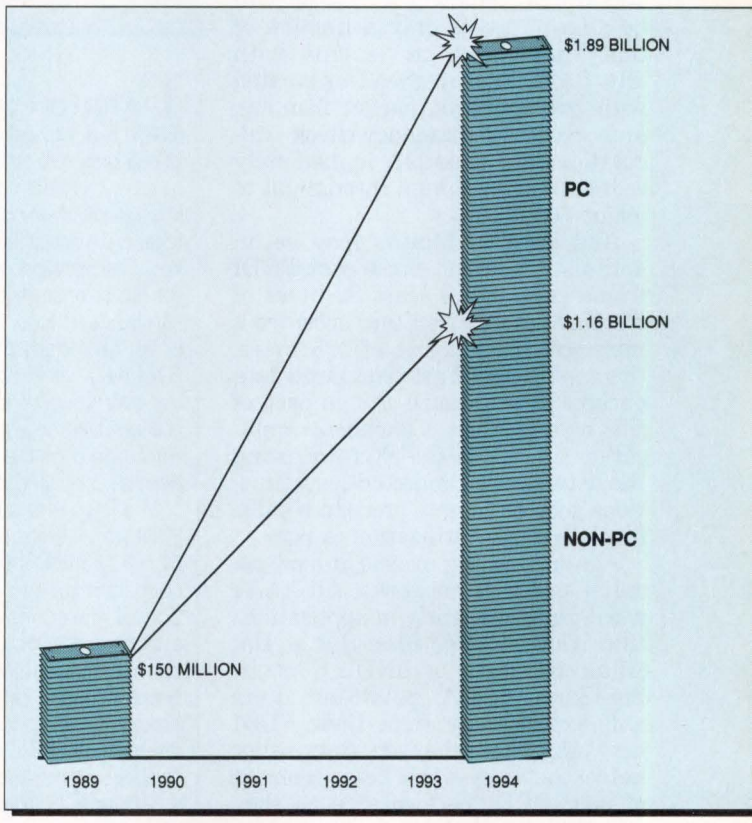

Advanced Micro Devices predicts that the market for FDDI stations will approach $2 billion by 1994.

Originally conceived in the late 1970s as a high-bandwidth point-to-point interface between mainframes and peripherals, the Fiber Distributed Data Interface (FDDI) is now poised to become the dominant high-end LAN of the 1990s. Within the next year, all of the pieces should fall into place. These "pieces" include ANSI ratification of the FDDI standard, including the Physical Layer (PHY), Data-Link Layer, and Station-Management (SMT) functions; hardware and software support; and cost-effective implementation of network and transport protocols that can support FDDI's full bandwidth.

Considerable progress has already been made in all of these areas. The minimum set of functions required to implement an FDDI node are established. The ANSI X3T9.5 standards committee has accepted the specification of the PHY and Data-Link layers, and acceptance of the SMT functions is expected within a few months.

Hardware support is also growing, with progress being made at the chip level. Cost continues to be a major hurdle, however, with the PMD, PHY and Media Access Control (MAC) silicon alone running about $2,000 per dual-attachment node. But higher volume, together with increased competition, are expected to drive the cost down by as much as 30 percent per year.

Another critical piece of the FDDI puzzle that's falling into place is the availability of standard network and transport protocol implementations that can support FDDI's 100-Mbit/s bandwidth. While current implementations of popular protocols such as TCP/IP (Transmission Control Protocol/Internet Protocol) are limited to 1 to 8 Mbits/s, vendors such as Communications Machinery Corp (Santa Barbara, CA) are preparing RISC-based communications controllers that can support TCP/IP throughputs of 60 to 80 Mbits/s. Entirely new protocols are also emerging. Designed for implementation in silicon, these protocols promise to be lower in cost, and to deliver throughputs as high as 1 Gbit/s.

FDDI

System-level support for FDDI is still in its infancy, but a variety of node adaptors, Ethernet-to-FDDI bridges, and development boards are already available. A much broader range of products will be announced later in the year, as standards crystallize and silicon implementations stabilize.

Performance increase

While FDDI theoretically supports a 100-Mbit/s bandwidth, users will be able to access only a fraction of that throughput, as is true with any LAN. The overhead associated with token delays, packet framing, and cyclical-redundancy-check calculations, for example, immediately reduces the maximum throughput to about 80 Mbits/s.

And even 80 Mbits/s may be an optimistic number. Since each FDDI frame contains at least 32 bytes of framing information and specifies a minimum packet size of 128 bytes, the application must send large data packets to minimize the impact of this overhead. In a backbone application, in which the PCs are using the network for command-level functions such as remote procedure calls, the bandwidth utilization is poor.

Again, framing overhead isn't peculiar to FDDI networks. All LANs work more efficiently in applications that transfer large files. But as Dal Allan, president of ENDL Consulting (Saratoga, CA), points out, users will expect more from their FDDI networks than they do from other networks. They won't be as tolerant of poor FDDI performance as they were with poor Ethernet performance.

Even if users can't take advantage of FDDI's full 100-Mbit/s bandwidth, FDDI still promises at least an order of magnitude increase in performance compared to LANs such as Ethernet. Recognizing the demand for this bandwidth, most major semiconductor vendors are rushing their FDDI silicon to market.

Several vendors, including PCO (Chatsworth, CA) and AT&T Bell Laboratories (Allentown, PA), provide the connectors and optical data links needed to implement FDDI's Physical Media Dependent (PMD) layer. Advanced Micro Devices (Sunnyvale, CA) is the only vendor, however, that offers a hardware implementation for FDDI's PHY and MAC levels. By the end of the year, AMD, National Semiconductor, Texas Instruments, Motorola, Intel and AT&T are all expected to make further announcements.

AMD's chip set, called Supernet, includes five chips: the Am7984 and Am7985 encode/decode (endec) chips, which implement the PHY layer; the Am79C83 Fiber Optic Ring Media Access Controller (FORMAC), which implements the Data-Link and MAC functions; and two buffer control chips, which implement a three-port memory controller.

The heart of the chip set is the FORMAC, which handles functions such as token management, system timer support, packet framing, and response to normal and system error conditions on the network. It also removes previously transmitted frames from the ring (all frames eventually return to the sender to

FDDI: an overview

The Fiber Distributed Data Interface (FDDI) is a 100-Mbit/s fiberoptic, dual token-ring network that can connect as many as 500 nodes with a maximum link-to-link distance of 2 km, and a total LAN circumference of about 100 km. The network uses its primary ring for data transmission; the secondary ring can be used to provide fault tolerance or for data transmission. Access to the FDDI network is controlled by a rotating token, which ensures deterministic, collison-free access in which the throughput of the network is independent of the number of stations.

The links that connect two adjacent FDDI stations are constructed of two 62.5/125 micron fibers with a duplex connector on each end. FDDI supports both single-connect stations, which attach to only one ring, and dual-connect stations, which connect to both rings and can be reconfigured around faults. A wiring concentrator interconnects the dual ring with single-connect workstations.

Through its dual-ring topology and connection-management functions, FDDI also defines a fault-recovery mechanism. If a fault (such as broken fiber cable or a malfunctioning station) occurs on the logical ring, that fault is isolated, causing the primary and secondary rings to collapse into one FDDI ring that maintains a logical path among users. Optical bypass switches let inactive nodes pass light directly from one neighbor to another without active power. If a fault occurs in a tree (subnetwork), recovery depends on the configuration. If the tree is redundant, the logical ring can heal itself; if not, the subtree is isolated.

The FDDI layers

The FDDI standard was developed in accordance with the Open Systems Interconnection (OSI) model, implementing the Physical and Data-Link layers of OSI. The four FDDI layers are Physical Media Dependent (PMD), Physical (PHY), Media Access Control (MAC), and Station Management (SMT).

The PMD Layer is the bottom half of the OSI's Physical Layer (Layer 1). It specifies duplex connectors, optical transceivers, and optional bypass switches. In the absence of power, these switches optically bypass the station.

FDDI's PHY Layer forms the upper half of the OSI's Physical Layer. In addition to providing 4B/5B encoding, it specifies a set of "line states" that perform a handshake between physical layers in adjacent stations. They are the basis for the connection-management protocols used in configuring FDDI's logical distributed topology.

The MAC Layer, which implements OSI's Data-Link Layer, provides peer-to-peer communications with the higher-level logical link control and SMT layers over the FDDI ring. In addition to token management, the MAC provides system timer support, packet framing, and response to normal and system error conditions on the network.

The basic data-link frame structure is similar to that specified for IEEE-802 frames, with the exception of starting and ending delimiters. The FDDI frame also adds a frame control field and frame status. The maximum frame length is 9,000 symbols (8 bits/symbol), with data packet sizes ranging from 128 bytes to 4,500 bytes.

SMT controls the activities of the MAC, PHY, and PMD. It includes functions such as connection management, fault detection, fault isolation and ring reconfiguration. The MAC and the PHY are published ANSI standards. The PMD should become a standard by spring. The SMT specification is about 80 percent complete, and should be finished by early summer.

FDDI is also being submitted as an ISO standard. The ANSI and ISO standards will be identical.

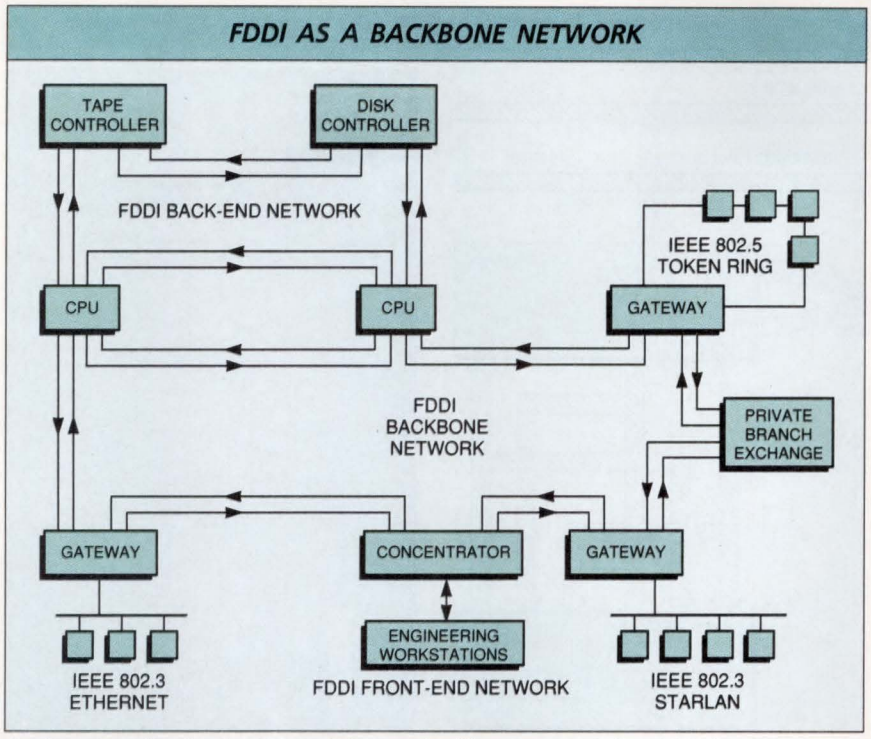

A primary application for FDDI will be as a backbone network, in which the FDDI network interconnects subnetworks, such as Ethernet, and other FDDI networks.

supply information such as transmission success). In addition, it handles the transmission and processing of beacon and claim frames, which are used to isolate ring faults and establish the token holding time for each node on the network.

The three-port memory controller supports up to 256 kbytes of buffer memory, organized as a FIFO. This buffer provides intermediate storage for data transmitted to and from the network, thereby preventing the loss and retransmission of frames and reducing network traffic.

The Supernet chip set is priced at $625 in quantities of 100.

AMD plans to introduce its next chip set before the end of the year, according to product manager Patrick Green. The new set will combine the MAC and buffer controllers in a single chip. In addition to supporting the current chip's 256-kbyte buffer memory, the new MAC will support a much larger buffer in system memory. It will also provide intelligent control that minimizes the movement of data from the intermediate buffer to other points in the system.

High-speed protocols the key

Silicon support may be the least of most designers' problems, however. To give users a painless migration path to FDDI, designers will have to preserve the existing base of application software available for popular networks such as Ethernet. This demands the porting of popular network and transport protocols such as TCP/IP.

The problem is that these compute-intensive protocols require expensive hardware support to sustain high throughput. Today's fastest communications controllers support TCP/IP throughputs of only between 1.5 Mbits/s (for an 80286-based controller) and 8 Mbits/s (for a 68030-based controller).

Until recently, many thought that TCP/IP's inherent overhead would make a cost-effective FDDI implementation impractical, but now it's not that far off, according to Russ Sherer, director of marketing at Communications Machinery Corp (CMC). By applying some of the queueing theory pioneered by Van Jacobson at the University of California at Berkeley, and using a RISC-based controller to execute TCP/IP, Sherer expects to achieve a TCP/IP throughput of 60 to 80 Mbits/s.

While CMC is extending the life of TCP/IP, Protocol Engines (Santa Barbara, CA) is working with 18 other vendors (including IBM, Boeing, Apollo Computer and Silicon Graphics) to build a new high-performance protocol. The protocol, known as XTP (Xpress Transfer Protocol), should be complete this month, with silicon support available by the first quarter of 1990, according to Larry Green, Protocol Engines' president. (Its initial price will be $600 for the four-chip set).

The protocol has already been accepted by the U.S. Navy as part of the Safenet (Survivable Adaptable Fiber Optic Embedded Network) standard, and will soon be submitted to ISO and ANSI committees for ratification.

XTP employs a number of techniques to outstrip the performance of protocols such as TCP/IP and TP-4. For example, TCP/IP calculates a frame's checksum and inserts it into the header prior to transmission.

The optimal FDDI chip set

Although most vendors are reluctant to provide details on what they're developing in terms of FDDI silicon, AT&T Bell Laboratories (Allentown, PA) has developed a description of what the ideal chip set would offer.

According to Juan Figueroa, member of AT&T's technical marketing and product planning group, the ideal FDDI solution would have the following attributes:

- A 32-bit external address and data path for higher speed communications.
- Reduced power to simplify the design of network concentrators, which will contain multiple nodes. This will require a CMOS implementation at the Physical (PHY) Layer level.
- A single-chip media-access solution, incorporating buffer management, glue logic, and a portion of the station-management functions, such as physical connection management and link error monitoring. Some station-management functions would also be incorporated at the PHY level.
- A single-chip PHY solution.
- A direct interface between the PHY and Physical Media Dependent layers, omitting the resistor-capacitor networks sometimes required to translate between the optical data link's ECL level and the encoder/decoder's TTL, CMOS, or quasi-ECL voltage levels.
- Lower cost (around $200).
- Support for larger buffer memories.

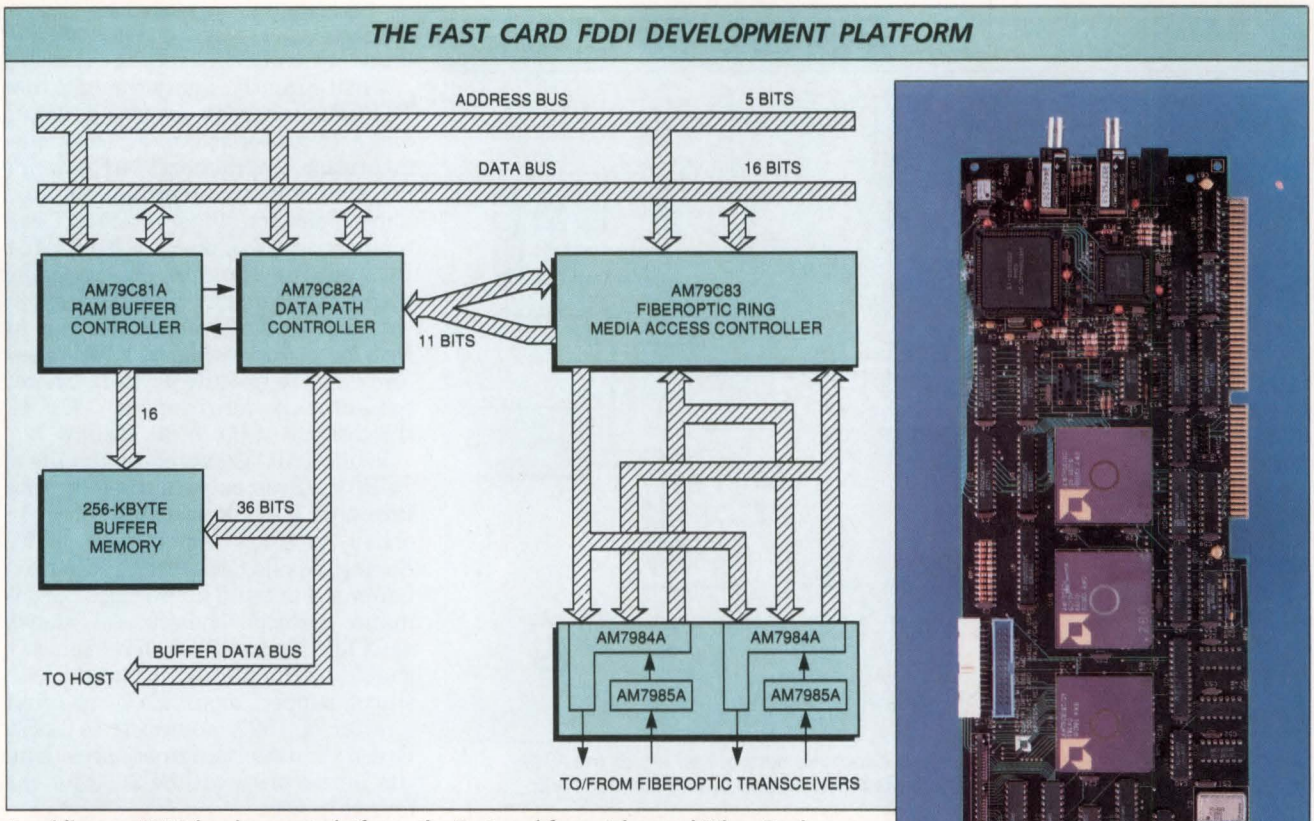

Providing an FDDI development platform, the Fast card from Advanced Micro Devices converts an IBM PC AT or compatible computer into a fully operational FDDI node that can be networked to other FDDI nodes.

XTP avoids this latency by initiating packet transmission without calculating the checksum. It calculates the checksum while the packet is being transmitted and appends it to the packet's trailer on the fly.

XTP is now running at about 8 Mbits/s on the Sun workstation. The new chip set, claims Green, will boost that throughput to 100 Mbits/s. Longer-term goals include a 1-Gbit/s implementation, which Green claims is achievable using current CMOS technology.

By implementing XTP in silicon, Protocol Engines may be able to provide a lower cost hardware solution than the RISC-based controllers currently being developed for TCP/IP. As with any new protocol, however, an initial lack of application software will be a major drawback. Its success will probably depend on how effective vendors are at extending TCP/IP. In any case, XTP seems to be a forerunner for FDDI's probable successor in the 1990s, a 1-Gbit/s LAN.

System-level support

While most vendors are awaiting the standardization of FDDI and the availability of low-cost silicon and protocol support, others are forging ahead. There are a number of FDDI products already available, from PC node adaptors to FDDI/Ethernet bridges. All are based on AMD's Supernet chip set.

To simplify the development of FDDI applications software for its chip set, AMD offers a plug-in evaluation board for the IBM PC AT called the Fast card. Priced at $5,000, this card converts an AT or compatible computer into a fully operational FDDI node that can be networked to other FDDI nodes.

AMD also offers FDDInet device drivers, which provide an interface to the link layer control, SMT and MAC layers. AMD licenses FDDInet at $6,000.

Simple-Net Systems (Brea, CA) also offers a single-attachment FDDI node adaptor card for the AT, known as LaserLAN. An optional daughter card supports a dual-attachment interface. Using an I/O-mapped PC interface, the board has 128 kbytes of packet memory, supports up to 1,500 m of cable between nodes, and up to 10 km for the entire ring.

To ensure ring integrity when the PC is turned on and off, LaserLAN includes a 12-V dc, 500-mA power transformer, which plugs directly into a wall outlet. To handle PC faults, LaserLAN controls an external relay, which bypasses the PC if it's disabled. The LaserLAN adapter card is priced at $4,995.

To support VMEbus military applications, Martin Marietta Aero & Naval Systems (Baltimore, MD) manufactures an FDDI Single Ring Interface (SRI) board that meets the U.S. Navy's Safenet II requirements. The card provides an extra level of redundant hardware that supplements the FDDI's inherent fault tolerance, according to Richard Steinberger, manager of network systems.

Providing a single-attachment FDDI interface, the SRI includes a custom DMA controller, 256 kbytes of 70-ns memory, and control for an optional optical bypass switch. A second board implements a dual-attachment interface.

FDDI-to-Ethernet bridge

In addition to developing simple node adapter cards, a number of

To Challenge Us.

Go on...put us to the test. Challenge our design engineers with your applications problems, and they'll design a quality Switchcraft component to satisfy your toughest product specifications.

Our product line is over 5,000 strong and includes success stories in all industries: we've developed EAC receptacles and high security interconnect devices for computer/datacommunications applications, locking switches for microwave ovens, even patch panels for a military satellite communications network. No matter what market, people look to Switchcraft for innovative design concepts to meet their unique demands.

The reason for Switchcraft's versatility is simple. Our team of design experts is supported by a fully automated manufacturing facility. Strict statistical process control standards support our high-volume production, and sophisticated CAD/CAM capabilities allow for flexibility in design.

We've thrown down the gauntlet. We challenge you to test Switchcraft components for their rugged dependability and unsurpassed quality. And we challenge you to push our engineers' imaginations to the limit.

Switchcraft

A **Raytheon** Company

5555 N. Elston Ave.
Chicago, IL 60630
(312) 792-2700

For instant response to your questions or for additional information, FAX us at (312) 792-2129, or call (312) 792-2700.

CIRCLE NO. 34

Moving toward FDDI II: integrating voice and data

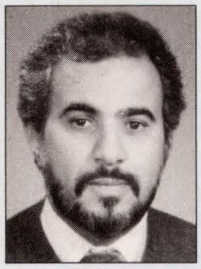

Even as FDDI is nearing ratification, the ANSI (X3T9.5) Committee is already considering two new extensions. One provides for a single-mode fiber Physical Media Dependent (PMD) layer, which will extend the maximum node-to-node distance from 2 km to 40 km—and even possibly greater than 60 km. The other, known as FDDI II, adds an isochronous data-transmission capability to the network, which will enable it to handle both voice and data.

One of the primary objectives of the FDDI II framers is to establish a standard that maintains compatibility with the existing FDDI standard. In fact, FDDI II will incorporate all of the standards already approved for FDDI, including the PMD, Physical (PHY), Media Access Control (MAC) and Station Management (SMT) functions. To support the circuit-switched transmission mode added to handle isochronous applications, vendors of existing FDDI silicon will have to make minor modifications.

Circuit switching for voice

FDDI was defined to handle the transmission of data in a packet-switched format. While FDDI is ideal for transmitting data, however, transmission delays associated with packetization and token rotation time make it less than ideal for voice transmission.

FDDI II adds a quality voice capability to FDDI by supplementing FDDI's packet-switching capability with synchronous circuit switching. Circuits are more suitable for voice transmission because they provide continuous bidirectional communications channels. Once a station that wants to transmit voice is allocated bandwidth and establishes a connection, the connection is maintained until the call is terminated.

By handling both voice and data, FDDI will be able to support integrated services applications and provide an efficient interface to digital public networks such as ISDN (integrated services digital network). It will also let organizations support local communications without relying on external facilities. External communications will be implemented via a private line link to the central office.

Basic and hybrid modes

FDDI II supports two modes of operation: basic and hybrid. In the basic mode, the ring supports only standard FDDI packet switching. In the hybrid mode, both packet-switched and isochronous data are transmitted within the same special frame structure, known as a cycle.

A cycle has a duration of 125 μs and nominally carries 3120 symbols and a T symbol preamble at 100 Mbits/s. This 125-μs interval enables FDDI II to support the full 8-kHz bandwidth needed for quality voice transmission and synchronization to external voice circuits (such as T1 trunk lines). The cycle supports a variable-rate, packet-switching service (using the FDDI token ring protocol) and a time-division multiplexed circuit-switching service.

A cycle's bandwidth is partitioned as a dedicated packet data channel, plus sixteen 6.144-Mbit/s wideband channels. The cycle master dynamically allocates these channels among the ring's stations for either packet data or isochronous use via a programming template that it transmits each cycle. The dedicated packet channel operates in conjunction with any wideband channels allocated to packet traffic. It uses the standard FDDI token-passing protocol and guarantees that a minimum packet channel bandwidth is 768 kbits/s (at 100 Mbits/s).

FDDI II defaults to the basic packet-switching mode. When a user application wants to transmit isochronous data such as voice, it makes a request to the SMT entity. In response to this request, the SMT entity takes the ring into hybrid mode. The exact process varies, but it generally involves allocating bandwidth to isochronous traffic, assigning monitor ranks, and ensuring that all stations on the ring are FDDI-compatible. Every station on the ring must be FDDI II-compatible for hybrid mode transmission to proceed. The actual transition to hybrid mode is initiated by a monitor station transmitting a cycle onto the ring.

One monitor station in an FDDI II ring acts as the cycle master, which is responsible for generating and maintaining the cycle structure and ring timing. It also inserts a latency adjustment buffer to adjust the ring size so that it's an integral multiple of 125 μs. The cycle master may be assigned by a system manager, or it may be selected through a bidding process that includes all ranked monitors on the ring. In the latter case, the monitor with the highest rank becomes the cycle master. A station's rank is a composite of an optionally assigned rank and its MAC address.

A station designated as a channel allocator (by default, the cycle master) allocates isochronous bandwidth to user stations. It maintains a record of which stations are using which isochronous channels and how much bandwidth they need. This record is updated as station requirements change.

Cycles are subdivided into 96 cyclic groups, each of which contains interleaved bytes from each of the 16 channels. When used for voice transmission, each channel may be subdivided into a number of transmission channels with varying bandwidths. To support ISDN, for example, each channel might be defined as an increment of 64 kbits/s. Any number of FDDI II stations may share the 16 channels, or any portion thereof. If all of the bandwidth is consumed by ongoing calls, a station trying to place a call receives a busy signal.

Move to standardization

The most recent meeting of the ANSI X3T9.5 Committee, held in February, included participants from major communications and computer vendors such as AT&T, Apple Computer, Prime Computer, IBM, Digital Equipment Corp, Sun Microsystems, Apollo Computer, and Hewlett-Packard. The meeting was marked by an unusually high level of cooperation.

The standard FDDI II document, known as Hybrid Ring Control, has already been submitted for letter ballot to the committee. The proposal should be submitted to ISO by September.

Because the standard hasn't yet been ratified, there is no silicon support for FDDI II. Once the standard has been accepted, the first wave of products will probably implement the PHY products.

Juan Figueroa, PhD, Team Member, Technical Marketing and Product Planning, AT&T Bell Laboratories

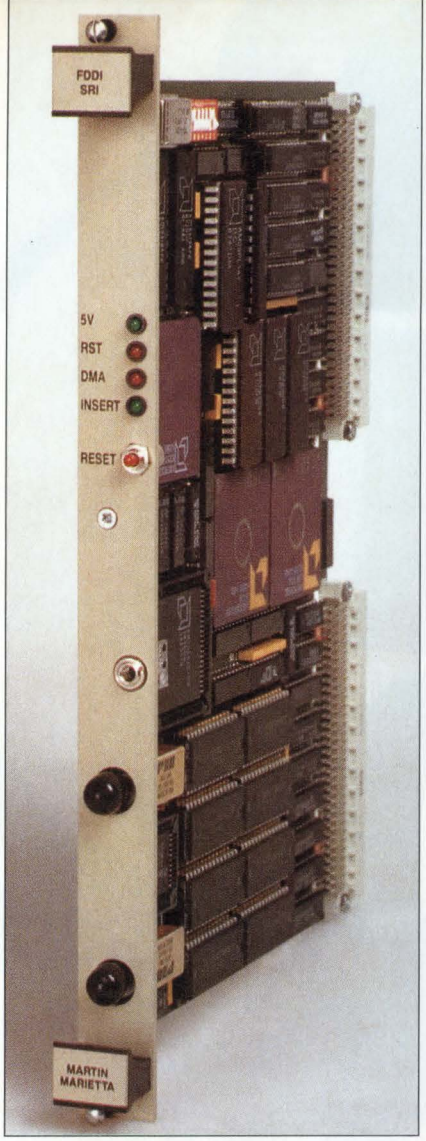

Targeted at VMEbus-based military applications, Martin Marietta's FDDI Single Ring Interface (SRI) board meets the U.S. Navy's Safenet II requirements and includes a custom DMA controller, 256 kbytes of 70-ns memory, and control for an optional optical bypass switch.

companies are working on bridge products that enable FDDI to be used with other LANs. Fibronics (Hyannis, MA), for example, offers an Ethernet-to-FDDI bridge, the FX8210. By combining a high-speed DMA controller with internal buffering, the FX8210 can transfer data to and from the Ethernet network at more than 10 Mbits/s. To reduce FDDI traffic, the FX8210 prevents the forwarding of local packets (packets not destined for remote Ethernets) onto the FDDI network.

Depending on how many Ethernet frames are forwarded to the network, the bridge can handle 5,000 to 10,000 frames/s (100 bytes each)—well within the practical rate supported by Ethernet, according to Greg Koss, corporate marketing manager. The bridge costs $25,000 per workstation. □

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

Advanced Micro Devices
(408) 749-5518 ..... **Circle 325**
AT&T Bell Laboratories
(215) 770-2922 ..... **Circle 326**
CMC
(805) 968-4262, ext.131 ..... **Circle 327**
ENDL Consulting
(408) 867-6630 ..... **Circle 328**
Fibronics International
(508) 778-0700, ext. 374 ..... **Circle 329**
Martin Marietta Naval Systems
(301) 682-0900 ..... **Circle 330**
PCO
(818) 700-1233 ..... **Circle 331**
Protocol Engines
(805) 965-0825 ..... **Circle 332**
Simple-Net Systems
(714) 529-8850 ..... **Circle 333**

Thoroughbred

Mizar's SPARC CPU keeps you ahead from start to finish.

You'll be favored to win with Mizar's MZ 7170, the VME processor with triple crown features. With the power of SPARC™ and the speed of zero wait-state SRAM, the MZ 7170 lets you run your application in record time. All for a price that keeps you on track.

The MZ 7170 was bred for demanding applications: a SPARC CPU at up to 25 MHz (15 MIPS), one MB of fast SRAM, up to four MB of EPROM, two RS-232 serial ports, mailbox interrupts, real-time clock, and optional T.I. 8847 FPU. And, the MZ 7170 is supported by a complete real-time operating system and UNIX®-based development tools.

Race into the homestretch with Mizar's thoroughbred, the MZ 7170. Call today. **1-800-635-0200.**

Mizar. The shortest distance between concept and reality.

MIZAR

1419 Dunn Drive • Carrollton, TX 75006 • (214) 446-2664

SPARC is a trademark of Sun Microsystems, Inc. UNIX is a registered trademark of AT&T. ©1989, Mizar, Inc.

CIRCLE NO. 56

SUPER FAST FIFO

25ns AVAILABLE NOW

Micron just broke another speed record with our new family of high performance FIFOs.

And that's just the beginning.

Micron FIFOs offer more than just incredible speeds, they're also packed with incredible features that solve your communication problems.

The Micron FIFO family offers industry standard parts, exciting new features, and a variety of configurations that run at speeds as fast as 25ns.

Expandable FIFOs are industry standard parts that can be combined to increase the depth and/or width of the buffer and provide standard flags.

Mailbox Register FIFOs feature an on-chip programmable register that bypasses specific words around the FIFO.

And **Variable Flag FIFOs** offer programmable full and empty flags.

We also employ high speed, low power CMOS designs using a true, dual-port 8-transistor memory cell. This produces a FIFO that operates at high speed with minimum power consumption and maximum data integrity.

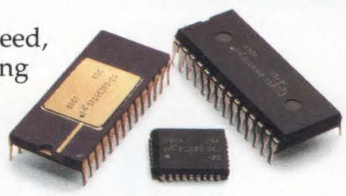

And like all Micron memory products, they're backed by a strong sales, customer service and technical support program designed to keep you on the leading edge.

So call Micron at 208-386-3900, and request more information about our Super Fast, Super Feature FIFOs.

Micron. A name worth remembering.

FIFO PRODUCT SELECTION GUIDE

Memory Configuration	Control Functions	Part Number	Access * Time (ns)	Package & Number of Pins			
				PDIP	CDIP	LCC	PLCC
512 × 9	E	MT52C9005	25	28	28	32	32
512 × 9	MB	MT52C9006	25	28	28	32	32
512 × 9	VF	MT52C9007	25	28	28	32	32
1K × 9	E	MT52C9010	25	28	28	32	32
2K × 9	E	MT52C9020	25	28	28	32	32
4K × 9	E	MT52C9040	25	28	28	32	32

MB .. Mailbox Register VF ... Variable Flags E .... Depth and Width Expandable
* Slower speeds also available

MICRON TECHNOLOGY, INC.

2805 E. Columbia Road, Boise, ID 83706 (208) 386-3900

CIRCLE NO. 35

Static RAMs race to keep up with RISC

Cache memory is becoming an integral part of RISC-based computer systems. But these new architectures burden the memory system by calling for fresh data and instructions each cycle.

Warren Andrews
Senior Editor

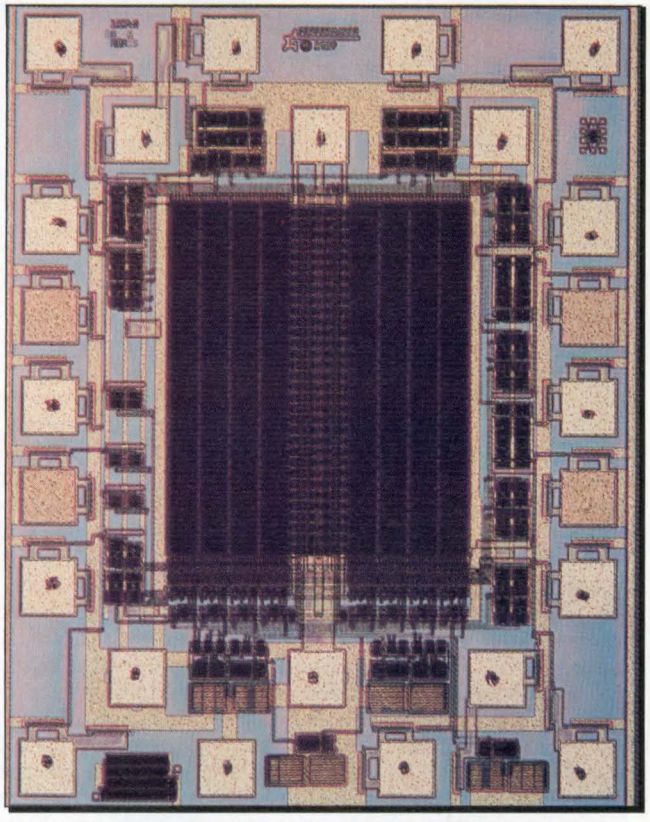

The demand for higher-speed SRAMs is pushing the limits of semiconductor technology. Performance Semiconductor leads the race in trimming feature sizes, bringing them down to 0.4 microns (l-effective). The company's 4-kbit all-CMOS SRAM boasts 6-ns access times.

High-speed static RAMs have traditionally been viewed as specialty parts, targeted at application niches where performance, not price, was the primary consideration. Not long ago, such niches included only specialized military applications, writable-control store for minicomputers and mainframes, and main storage for supercomputers. But the advent of very fast microprocessors with pipelined architectures, the widespread acceptance of RISC processors, and system designers' insatiable craving for higher throughput rates has greatly expanded the role of these "niche" memory parts.

Due to shrinking feature sizes and clever architectures, access times have been driven down by 50 to 70 ns and have passed the 25-ns barrier, with some SRAM vendors offering devices with speeds in the single-digit range. Parts are available in these super-speed ranges in all technologies—GaAs, bipolar ECL, BiCMOS and bulk silicon CMOS. But while access and cycle times of different technologies may appear the same on the spec sheets, each technology boasts subtleties that endear it to particular users with special applications. In addition, the differences in process technologies are reflected by significant differences in price.

The current demand for high-speed SRAMs was well anticipated by the industry, and the number of suppliers mushroomed from a handful in the early 1980s to well over 50 today. Although most of the leading-edge suppliers have been smaller start-up companies such as Integrated Device Technology (San Jose, CA), Saratoga Semiconductor (Cupertino, CA), Performance Semiconductor (Sunnyvale, CA), Cypress Semiconductor (San Jose, CA), Micron Technology (Boise, ID), Synergy (San Jose, CA), Vitesse Semiconductor (Camarillo, CA) and Gigabit Logic (Newbury Park, CA), many of the mainline semiconductor makers on both sides of the Pacific

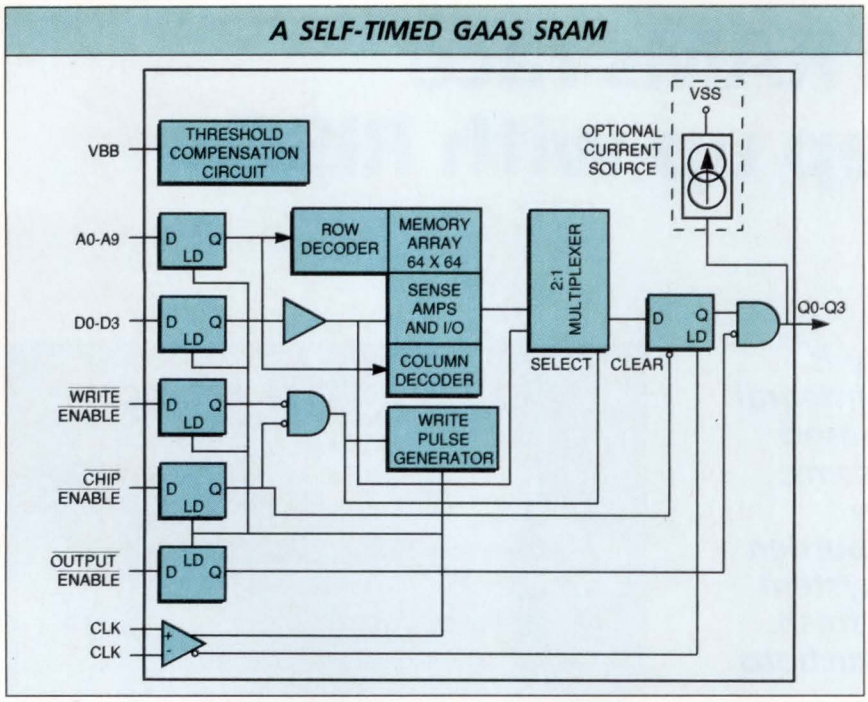

Gigabit Logic's 3.5-ns self-timed GaAs static RAM features balanced read and write cycles, latched I/O and provision to handle write-through cache operations.

have entered the market.

National Semiconductor (Sunnyvale, CA), Motorola (Austin, TX) and Advanced Micro Devices (Santa Clara, CA), for example, have expanded their high-performance SRAM offerings, while the traditional semiconductor vendors from Japan—Hitachi (San Jose, CA), Toshiba (Tustin, CA) and Fujitsu (Santa Clara, CA), for example—have been joined by other, less well-known silicon vendors such as Sharp Microelectronics (Mahwah, NJ) and Sony (Cypress, CA). And while much of the semiconductor industry is settling in for a slump, SRAM lines are clicking away at peak capacity and new lines are being built.

Demand driven

Although traditional requirements for minicomputers, mainframes and supercomputers continue to fuel the demand for high-performance SRAMs, the most recent catalyst for SRAM growth has been the emergence of RISC processors such as the Sun Microsystems Sparc, the MIPS 2000 and 3000 processors, the Motorola 88000 and the AMD 29000. For RISC processors to provide significant performance increases over CISC engines, they have to execute one instruction per clock cycle. But this rate starts to put heavy demands on the processor's memory system. A 33-MHz processor, for example, would have to gobble up one instruction every 30 ns. And a 50-MHz processor—some will probably start to emerge later this year—must have a new instruction waiting for it every 20 ns.

"But that doesn't necessarily mean that a 30-ns access-time SRAM will work with a 33-MHz processor," says Sam Orr, SRAM marketing manager for Cypress Semiconductor. "The propagation time in the logic to set up and latch the information takes between 7 and 9 ns, so a 25-ns part will just barely squeeze by for a 33-MHz processor." Similarly, processors operating at 40 MHz will require information to be available in 20 ns, calling for a sub-15-ns access device.

As manufacturers scramble to provide these lightning-fast access-time parts, they're also developing clever new architectures to reduce the amount of external logic used and, therefore, the external delays. Beyond that, they're being asked to provide other features, along with the SRAM, that will help reduce propagation delays through the copper traces of printed circuit boards and eliminate timing skews.

Leader of the pack

As expected, GaAs devices are at the leading edge of high-speed SRAM development, with some boasting access times as fast as 2.5 to 3.5 ns. The fastest GaAs SRAM is a 1k-x4-bit device from Vitesse Semiconductor (Camarillo, CA) that boasts a 2.5-ns access time. Vitesse's part is self-timed—it contains registers on inputs and outputs that eliminate timing problems, says Tom Dugan, Vitesse marketing manager. Since the data is latched in and out on the edge of a single clock pulse, there's no fear of data changing if a slight skew in timing occurs.

Another recent GaAs entry comes from Gigibit Logic in the form of a 1k-×4-bit GaAs SRAM. The part is a self-timed SRAM using an on-board write-pulse generator to provide an equal read/write cycle time, according to John Kemps, marketing director at the company. What's more, the read/write cycle time is the same as the read/write access time. In many high-speed static RAMs, the read time is often faster than the write time. In applications such as writable control storage, this difference presents only minor slow-ups, but in cache-memory applications using various write-through schemes, the total system performance can be degraded.

The chip also has an output-enable control and has a special provision to simplify write-through operation (a feature primarily used for high-performance cache memory applications). The chip can handle write-through operations because it makes a copy of the written data available at the output of the chip during the same cycle in which the data is written. Another departure from a conventional static RAM, says Kemps, is that the part is a self-timed, or clocked, part similar to Vitesse's. The clock is used to control the I/O latches. On the Gigabit part, the input is open and the output is latched on one side of the clock pulse, and the output is open and the input latched on the other side of the clock pulse.

While GaAs may hold the lead in terms of access time, that lead may be short-lived as other vendors using other technologies scurry to catch up. Fujitsu, for example, offers its industry-standard 474xx, which boasts a 5-ns read-access time and is one of the fastest ECL devices available. The write-access, however, climbs up to about 9 ns, however, limiting the part's performance in applications such as cache memory, where it's desirable to mix arbitrary reads and writes. In such applica-

I/O—as opposed to TTL I/O—gives the designer more flexibility and often simplifies design, Ford says. "We've seen many customers who have not traditionally been ECL customers looking closely at ECL I/O logic levels.

"What's happening," he says, "is that their systems have gotten to the point where the system clock is running in the 30- and 50-MHz range in RISC-based systems and perhaps twice that in CISC-based processors (resulting in a system clock rate of about half of the processor clock rate)." At these speeds (30 to 50 MHz), designers are finding out that TTL I/O isn't the way to go anymore, he says. Making a transition from 0 to 5 V or from 0.4 to 2.4 V limits the access time window.

"You can't switch 2 or 3 volts quickly enough to give useful access times at that rate," he says. At a 50-ns clock rate, there's a maximum of 20 ns available for memory access. The typical rise and fall time of TTL is somewhere around 10 ns, leaving only 10 ns to get the signals propagated, allow for timing skew, access the memory and propagate the signal back again. "That's just not enough time," he says.

Designers are looking for higher performance I/O, says Ford, and ECL is a good answer: it's been around for a long time, it's fully debugged and it's rule-based. Most PC-board layout programs comprehend only the logical and functional definition for TTL logic. But many of these same programs take into account these factors plus the ac characteristics of ECL logic. These programs automatically take care of wire lengths and terminations, actually making high-speed design easier than conventional design once the decision is made to go to ECL.

Interestingly, the need to go to fast I/O logic levels has taken on different meanings in some of the GaAs devices. Gigabit's Kemps comments that at least two of the company's customers—Prisma and Cray Research (Minneapolis, MN)—have elected to use their own logic levels rather than accept ECL or other standards.

Because of the increasing demand for ECL SRAMs, Aspen plans to introduce a variety of devices, including higher-density parts, in the near future. "The next jump in density will be to a 64-kbit device," he says. The denser part won't sport the same 3-ns access time, however. To achieve the greater density without going overboard in power dissipation, the part will include more CMOS functions and fewer bipolar devices. Ford expects the 64-kbit part to come in with an access time of around 8 ns.

For its BiCMOS devices, Aspen is using a 0.8-micron CMOS process, which Ford calls a very "manufacturable" process. This manufacturability, in fact, will let the prices for the BiCMOS parts be basically the same as current prices for slightly slower ECL parts. The next generation will take advantage of the 0.5-micron process under development at Cypress, he says. The BiCMOS process used for Aspen's high-speed SRAMs is completely scalable to finer geometries.

The advantages gained with the BiCMOS process translate into getting very high speeds with relatively mature process technology. But because three to five additional mask layers are required, the process is still relatively costly.

■ CMOS leadership

The alternative is to push the state of the art in all-CMOS, which is exactly what Performance Semiconductor has done. Performance is an all-CMOS house, says Sam Young, manager of memory and logic marketing for the company, and he believes that it's possible to do everything that has to be done (with the exception of ECL I/O) in CMOS. "It's by far the best technology from the points of view of density, manufacturability and cost," he adds.

The company stretched the limits of process technology to produce all-CMOS SRAMs, which it calls SCRAMS (for static CMOS RAMs), with access times dipping to 6 ns. To do this, Performance invoked its latest technology, called Pace III, which has 0.4-micron feature sizes and 0.75-micron line widths with a 1.75-micron metal pitch.

Pace III is one of the first commercial SRAM processes to use the recently approved 3.3-V JEDEC standard. The company's first parts to use the process are two 4-kbit SRAMs, each offering 6-ns access times. One is organized as 4k-×1-bit, the other 1k-×4-bit. Both parts use Performance's six-transistor-cell architecture, provide three-state outputs and offer balanced read and write cycles.

In addition, the chips are packaged with center-pin power and ground pins to significantly reduce noise and ground bounce. Both Texas Instruments (Dallas, TX) and Philips/Signetics (Sunnyvale, CA) recently used this technique in their high-speed CMOS logic devices. Performance's choice of this technique—along with the provision of extra power and ground pins—may have been prompted in part by the lower noise margins of the 3.3-V process. Performance, says Young, will also release a number of standard logic functions to mate with its high-performance SRAMs.

Though the 3.3-V standard has been JEDEC-approved, its widespread use probably won't be realized for a few years yet. This year, Performance plans to introduce a complete product line in its Pace III

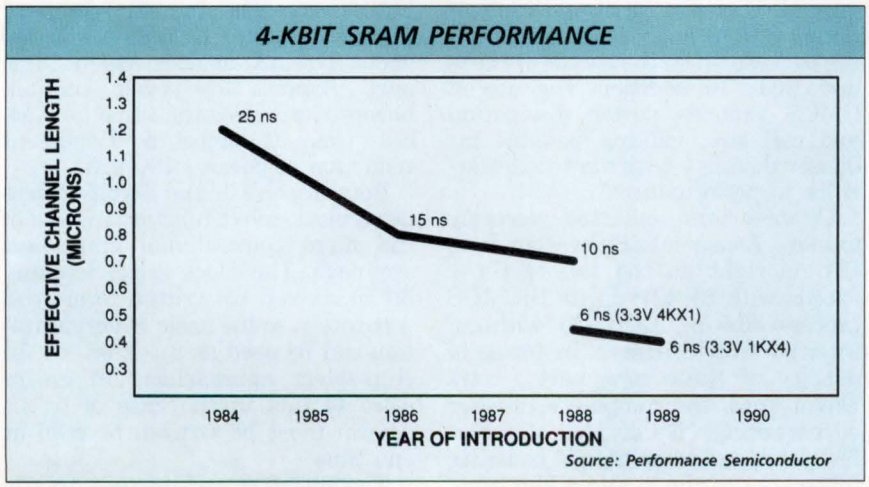

Scaling the feature sizes is critical to the performance of the latest-generation CMOS static RAMs. These two curves show the relative performance of a 4-kbit full CMOS SRAM with different effective channel lengths. Note the stair-step function as parts move to the 3.3-V JEDEC standard.

tions, the different read/write times slow operation down to the write-access time.

Fujitsu has also recently introduced a different ECL part with functions and pinout similar to the Gigabit GaAs part, which offered a 9-ns access time at its introduction. The company will soon announce a much faster version that will drop access times to the 3- to 4-ns range, however.

As SRAM access times continue to drop, self-timing and registered I/O become increasingly important. "The movement toward self-timed SRAMs with on-board latches is just starting to catch on with such devices as the Fujitsu part and some 16- and 64-kbit BiCMOS devices from Motorola and others," observes Kemps. "The more logic that can be incorporated on the SRAM, the fewer delays there will be from associated external logic in fast cache memory subsystems." And the need for super-fast cache memory isn't restricted to mainframe and supercomputer makers. "Many workstation makers are calling for memory that will cycle in the 3- to 4-ns range," adds Kemps.

Even as access times continue to fall in transparent memory products, some demanding applications call for more than even conventional GaAs logic speeds can provide. This is particularly true in some proprietary cache memory subsystems. Kemps cites Prisma Computers (Colorado Springs, CO), which is building a 250-Mips, Sparc-based machine using GaAs logic, as an example. Prisma's problem is to combine both memory and logic in a cache system that can be accessed in one or two cycles. "They were forced to abandon standard logic and use a custom GaAs chip to get the speed," says Kemps.

Kemps believes there may be an increasing trend toward the application of high-speed custom logic to surround SRAMs in cache-memory applications. "One popular approach to providing this logic is to use high-speed PLDs," he points out. But the delays associated with such approaches start to bring total cache memory cycle time to 6 or 10 ns—too slow, by at least a factor of two, and perhaps as much as four, for what some workstation and computer makers need.

Vitesse also believes the future for GaAs SRAMs will be in the semicustom arena. By using straight GaAs levels instead of converting to ECL or TTL levels, says Dugan, speeds can be reduced from the 2.5 ns area to 1.5 ns. It's unlikely, however, that such parts will be transparent SRAMs (industry-standard stand-alone RAMs), but rather will comprise parts of semicustom libraries. Vitesse is completing the development of a library of GaAs SRAM cells using a RAM compiler from VLSI Technology (San Jose, CA) through a technology exchange with that company.

Gigabit continues to push GaAs technology and wants to drop access times to the sub-3 ns range, according to Kemps. At the same time, he says, the company is trying to increase density, with activity already underway on a 16-kbit device. For such a device, Gigibit's 1-micron process will have to be shrunk to about 0.6 microns. The three levels of metal used in the process for the 4-kbit device will probably remain the same, he adds.

"The movement toward self-timed SRAMs with on-board latches is just starting to catch on."

—John Kemps, Gigabit Logic

■ ■ ■

■ CMOS and BiCMOS creep up

GaAs and ECL devices may be in the lead in terms of access time, but other approaches are rapidly gaining ground. Both BiCMOS and conventional CMOS approaches are starting to creep up on the speeds of GaAs and ECL. In addition, the use of CMOS reduces power dissipation and cell size, making possible far denser devices—with very little sacrifice in performance.

Cypress Semiconductor's recently founded Aspen subsidiary (San Jose, CA) is right on the tail of GaAs speeds with SRAMs in its BiCMOS process offering ECL I/O with an internal CMOS array. "In terms of density of these new parts," says David Ford, the company's director of marketing, "it's deja vu of only a few years ago when SRAM densities topped out at the 1- and 4-kbit area. We're the first company to bring out BiCMOS devices at such low-density levels." Others are looking at densities in the 256-kbit range, he says. The feeling has traditionally been that BiCMOS devices in the lower density area (1, 4 and 16 kbits) would be too slow—greater than 10-ns access time—providing little competition to ECL and GaAs devices.

But that's not the case, says Ford. One of Aspen's first parts is a soon-to-be-released 1-kbit device organized as a 256-×4-bit part that clocks in at a 3-ns access time, fully comparable to the GaAs devices. "One of the advantages of Aspen's approach," says Ford, "is that the device uses standard ECL levels and is produced with standard manufacturing technology in a high-volume process." Ford claims that the company could reduce the access time of its 1-kbit SRAM to 2.5 ns once the process gets rolling.

Ford sees many applications for Aspen's fast devices in high-performance mainframes as writable control storage, in fast parallel processors and array processors, and in many automatic test equipment machines. Most supercomputers are looking for ×1, rather than ×4, organizations and by and large, are using higher-density (in the 256-kbit range), albeit slower, devices. Today's high-performance RISC processors, for the most part, haven't reached a point where this kind of speed is required, he says.

Aspen is also introducing a 4-kbit device with the same 3-ns access time. Like Gigabit's device, the device provides balanced read and write cycles. In addition, the company offers low-power versions of the devices, which are pin-for-pin replacements for industry-standard ECL parts. The industry-standard Fujitsu 494, 5-ns ECL device draws about 275 mA, as does Aspen's 3-ns part. Aspen's low-power version, however, provides the same 5-ns access time at almost a 40 percent reduction of power—190 mA.

Both Aspen's 1- and 4-kbit devices use a block-select function instead of the more conventional chip-select approach. The block select lets any bit in a word be written to or read separately, so the basic ×4 organization can be used as a ×2 or ×1. In chip-select approaches, an entire word (4 bits in the case of a ×4 device) must be written or read at one time.

The BiCMOS architecture greatly reduces the power required to drive the memory array, leaving more power available to focus on the ECL I/O functions. Further, the ECL

Our new 256Kx1 and 64Kx4 Mighty Memory CMOS SRAMs are ready to take on your heavy-weight advanced architecture design projects.

Perfect for high performance applications such as Military 1750, DSP, ALU, and other 32 bit micro designs, our Mighty Memories are available now in DIP and LCC. Both devices feature speeds to 35ns in standard and low power versions, MIL-STD-883 Class B screened.

EDI is also listed as an approved source on 256Kx1 and 64Kx4 DESC-SMD drawings.

High density. High speed. Low power. The new Mighty Memory 256K Static RAMs from EDI, the high performance military Static RAM leader.

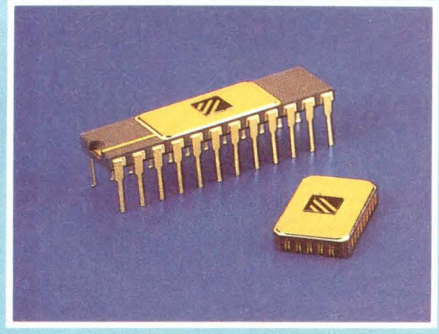

EDI's fast new 256K CMOS Static RAMs, available now in 256Kx1 and 64Kx4 with JEDEC standard pinouts, DIP and LCC.

CIRCLE NO. 36

EDI
The future...today.

ELECTRONIC DESIGNS INC. *42 South Street, Hopkinton, MA 01748, USA, (508) 435-2341, TFX: (508) 4356302, TELEX 948004*
ELECTRONIC DESIGNS EUROPE*, Shelley House, The Avenue, Lightwater, Surrey GU18 5RF, United Kingdom, Phone: 0276 72637 TELEX 858325*

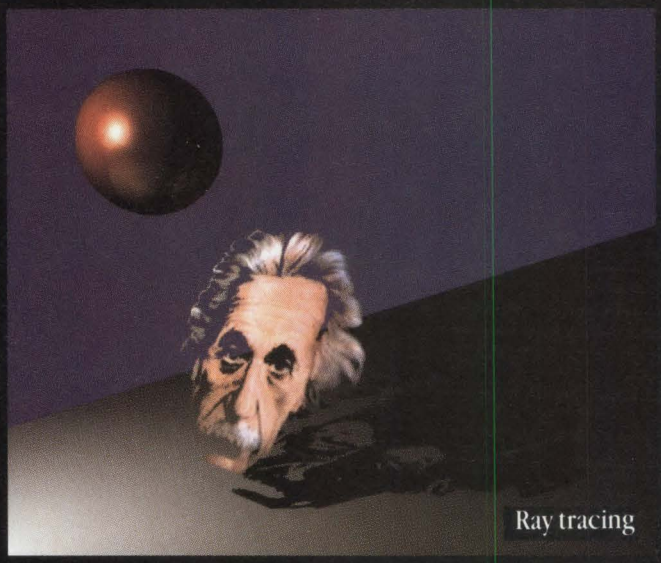

Start jumping to conclusions.

Pixel Machines' PXM 900 Series supercomputers deliver a quantum leap in image computing.

Pixel Machines' PXM 900s give you realtime, interactive performance for both 3D graphics and image processing—as much as 800 times faster than today's high-end machines.

Pixel Machines use a parallel array of up to 82 32-bit floating point processors—fully programmable to optimize the power of your image computing application as no other supercomputer can.

That means maximum functionality for photorealistic rendering and animation, visual simulation, volume display and analysis, and scientific visualization.

The Pixel Machine delivers. Up to 820 Megaflops. Image memory? Up to 4096x2048 at 32 bits per pixel. Plus an additional 16 Megabytes of general purpose memory. Software? A variety of tools tailored to your application.

Yet for all its supercomputer performance, Pixel Machines' PXM 900 Series is workstation-priced.

For more information, call Pixel Machines at **1 800 544-0097**. (In NJ, call **201 563-2294**.)

AT&T
The right choice.

CIRCLE NO. 37

© 1989 AT&T

process. In addition to the company's 4-kbit SCRAM, the line will include other SCRAMs with up to 256-kbit densities, popular logic devices and even a small gate array for users to incorporate random logic, says Young. As CMOS feature sizes continue to be scaled, the shift down to the 3.3-V standard is imminent. In the scaling process, both horizontal and vertical dimensions are scaled, resulting in oxide layers in the 150- to 129-A area.

Despite the company's lead in 3.3-V CMOS technology, however, Young stresses that Performance is still very much in the 5-V SCRAM business and intends to continue to improve speed and density in that family of products. The company offers some of the fastest 5-V CMOS SRAMs available, including 10-ns 64k-×1-bit, 15-ns 16k-×4-bit, 15-ns 4k-×4-bit, 15-ns 2k-×8-bit and 10-ns 4k-×1-bit SCRAMs.

Performance is also developing products targeted directly at the RISC processor market. The first of these is a standard 16k-×4-bit SCRAM but with a very fast output enable and tight write cycles. RISC machines, such as the 2000 and 3000 processors from MIPS Computer Systems (Sunnyvale, CA), want the output enable to be extremely fast—less than 50 percent of the address-access time—and the write pulse width and data set-up time on the rising edge of output enable to be less than 50 percent of the cycle time. While the MIPS architecture places tight demands on such specifications, these timing specifications are equally critical on other RISC machines such as most of the Sparc architectures in addition to CISC processors such as the 68030 and 80386, says Young.

Unlike what happens in the dynamic RAM market, lower density SRAMs never go away—they just get faster, says Young. Because of this, Performance plans to keep a large catalog of SRAMs of different densities, speeds and organizations.

Higher-density devices

While Performance holds the speed record for low-density CMOS SRAMs, clever architecture and the learning-curve experience of making 1-Mbit DRAMs have let Micron Technology, using an all-CMOS process, push the speed limit on higher-density devices. DRAMs traditionally lead the process technology, says Terry Walther, applications engineering manager for the company, and the experience gained through DRAM manufacturing let Micron come out of the chute with a 256-kbit full CMOS SRAM with a 25-ns access time—at the time, the fastest SRAM of that density.

Though Micron had to struggle with foreign competition in the DRAM market, the company managed to stay alive in the DRAM business and survived—with Texas Instruments—as one of only two U.S. DRAM makers. Early on, Micron distinguished itself in the DRAM business by having the smallest die size of any of the 64-kbit devices. It's been largely through clever designs resulting in small, manufacturable die (and a lot of gumption to stick in the business) that Micron survived the bloodbath that literally wiped out domestic DRAM manufacturing. Even at the 1-Mbit DRAM density, Micron still has one of the smallest die.

Micron has taken the same approach with its SRAMs and is producing the 256-kbit device on its 1.2-micron, 1-Mbit DRAM line. With a die shrink coming within the next four or five months, the company expects to reduce the access time to about 20 ns or less. Already under development, the next generation is the 1-Mbit SRAM, which will be produced on the company's 4-Mbit DRAM line using 0.8-micron feature sizes, says Walther.

"Though we're going head-to-head with some of the BiCMOS processes, which we admit have a 10 to 20 percent theoretical advantage in speed, we've been able to keep up with—and in some cases exceed—the speed of the BiCMOS," says Walther. He adds that Micron doesn't plan to go to a BiCMOS technology until it can be proven that CMOS can't keep up in performance. Further, he says, production costs are lower because of the high-volume production capability of the company's DRAM fab.

BiCMOS leapfrogs

Though Micron's 256-kbit SRAM was the industry's fastest when it was first introduced about a year ago, several new BiCMOS devices have since leapfrogged over it. One of these devices, from National Semiconductor, is a 256k-×1-bit BiCMOS SRAM with a 15-ns cycle time. The 1-micron, double-level metal BiCMOS III process is being ramped up in the company's Puyallup, WA plant, according to Robert Bernardi, BiCMOS memory marketing manager at National.

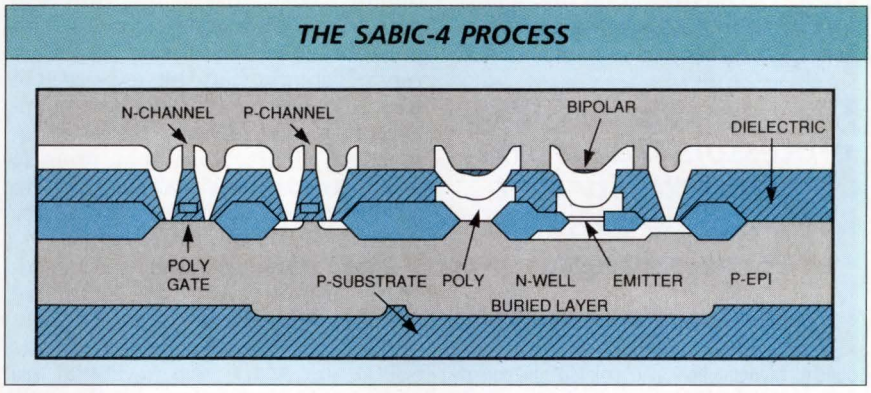

Saratoga Semiconductor's latest process, SABiC-4 (Self Aligned Bipolar CMOS), allows high-performance I/O with power-stingy CMOS arrays. Because speed and drive are managed by the bipolar transistors, the process is easy to manufacture.

While the first memory device to be run in the process is a 256-kbit part, Bernardi says the company will "backfill" its line with lower density devices. In addition, says Bernardi, a handful of specialty memory parts are on the drawing boards. These parts will include some cache memories that will come in with sub 10-ns access times and will probably be relatively small and wide—with 9-bit word widths. There are no present plans for wider word widths, such as the 16-bit width provided by IDT, he says.

National plans to reduce feature sizes in the near future from the present 1 micron to 0.8 micron. The target is to be able to reduce access time to below 12 ns, says Bernardi.

National, and others, will see some competition as Hitachi, Fujitsu and others ramp up their BiCMOS processes, however. Fujitsu is readying its 256-kbit BiCMOS device for volume production, and Hitachi, one

of the oldest players in the BiCMOS RAM business, is also expected to provide tough competition.

Still another contender in the BiCMOS arena is Saratoga Semiconductor, which just introduced its SABiC-4 (Self Aligned Bipolar CMOS) BiCMOS process. The new process represents an almost 48 percent shrink over the company's earlier SABiC-3 process, according to Steve Lau, technical marketing director. Lau says the drawn dimension of the CMOS transistors has been reduced from 1.5 to 1.2 microns, while the bipolar transistors' 0.15-micron base width has remained relatively the same. He notes that the bipolar transistors have a 7- to 8-GHz cut-off frequency.

With the move to the new process, Saratoga has been able to cut SRAM access times by 20 to 25 percent while reducing power dissipation and die area by 35 percent. With the die shrink, the company can produce 16-kbit SRAMs with 15-ns access times and 64-kbit devices with 20-ns access times.

The new breed of RISC processors may create an opportunity that dwarfs the existing SRAM markets.

■ ■ ■

While Saratoga will participate in the commodity SRAM market as a driver for its technology, the company has seen greater market success with its specialty SRAM products. These include a variety of cache RAM parts as well as FIFOs and other SRAMs aimed at capturing more of the RISC cache business. In addition, says Lau, Saratoga plans to expand its offerings of $\times 4$ and $\times 8$ organizations to include word widths of 16 and perhaps even 32 bits.

The list of SRAM makers is long and growing. Over the next few months, one manufacturer after the other will probably leapfrog over the previous generation of SRAMs in performance and density. Sony America's semiconductor group, for example, is about to introduce its family of MixMOS SRAMs, which includes 64- and 256-kbit devices. The 64-kbit part is expected to clock in with a 25-ns access time, while the 256-kbit device will have a more modest 35-ns access time. These times are likely to drop sharply as the company ramps up production. Sony has recently been a speed leader in ECL logic and will probably map that technology into a powerful BiCMOS process in the near future.

Similarly, Sharp Microelectronics will expand its fast SRAM line with products ranging from 16 to 265 kbits and some versions pushing access times down to 20 ns. Toshiba with a staggering share of the low-power SRAM market, continues to push speeds down and will undoubtedly be a major participant in the high-performance SRAM business.

■ No end in sight

While the SRAM market has traditionally been driven by military, mainframe and supercomputer applications, there's a strong undercurrent of thought that the new breed of RISC processors will create an opportunity that dwarfs the existing markets. Sparc chips that promise one instruction/cycle are already in production at 25 MHz, with faster versions just around the corner.

Mips machines with speeds already approaching 33 MHz are proliferating from a number of manufacturers. Motorola with its 88000 and AMD with its 29000 are right up there battling for the pole position. And Intel recently unveiled its long-awaited 80860 RISC-based processor at the International Solid State Circuits Conference. The chip is designed with its own sizable internal cache, eliminating the need for at least the first level of external fast SRAM cache.

Undoubtedly, more such parts will follow. But performance will depend on whether SRAM makers can keep up the pace. Part of the battle will always be raw speed—the lowest possible access time. To make up for shortcomings in SRAM access time, processor architectures are beginning to reduce the quantity—and in some cases the speed—required of caches. Other semicustom approaches are integrating logic with SRAM cells to dramatically reduce the overhead for cache-memory subsystems. Further savings in overhead may be accomplished through the use of modules or hybrid circuits where board-trace lengths can be minimized, thus reducing timing skews and giving parts a little more tolerance.

And while conventional TTL I/O served well for a couple of generations, it may no longer be up to the task of handling signals above 20 MHz, as Aspen's David Ford points out. ECL seems to be a logical choice, since BiCMOS processes are maturing and ECL transistors can be included without additional overhead. BiCMOS processes are more costly than straight CMOS, however. Performance Semiconductor demonstrated that straight CMOS can close in on ECL speeds and, with the 3.3-V supply, the TTL I/O approach may withstand another generation of performance.

To add to the confusion, multiprocessing techniques with a need for some kind of cache coherency are just starting to surface. And it's possible that with a strong push from Intel, an advanced Futurebus architecture may lead the way for many of these RISC-based products on a cache-coherent standard bus. Undoubtedly, this type of cache coherency will call for increasingly specialized SRAM products—with perhaps even faster read/write times to accommodate cache coherence cycles in between processor cycles. □

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

Advanced Micro Devices
(408) 235-8327 ..... **Circle 268**
Cypress Semiconductor
(408) 943-2600 ..... **Circle 269**
Fujitsu
(408) 562-1000 ..... **Circle 270**
Gigabit Logic
(800) GAA-SICS ..... **Circle 271**
Hitachi
(415) 244-7147 ..... **Circle 272**
Integrated Device Technology
(408) 749-1733 ..... **Circle 273**
Micron Technology
(208) 386-3900 ..... **Circle 274**
MIPS Computer Systems
(408) 720-1700 ..... **Circle 275**
Motorola
(512) 440-4505 ..... **Circle 276**
National Semiconductor
(408) 721-5724 ..... **Circle 277**
Performance Semicondutor
(408) 734-9000 ..... **Circle 278**
Prisma Computers
(719) 594-6018 ..... **Circle 279**
Philips/Signetics
(408)991-2373 ..... **Circle 280**
Saratoga Semiconductor
(408) 522-7523 ..... **Circle 281**
Sharp Microelectronics
(201) 529-8200 ..... **Circle 282**
Sony America Semiconductor Group
(714) 229-4197 ..... **Circle 283**
Synergy
(619) 753-2191 ..... **Circle 284**
Texas Instruments
(214) 997-6251 ..... **Circle 285**
Toshiba
(714) 832-6300 ..... **Circle 286**
Vitesse Semiconductor
(800) 848-3773 ..... **Circle 287**

VME-to-VME Links

2 YEAR WARRANTY

REPEATERS

SOFTWARE TRANSPARENT

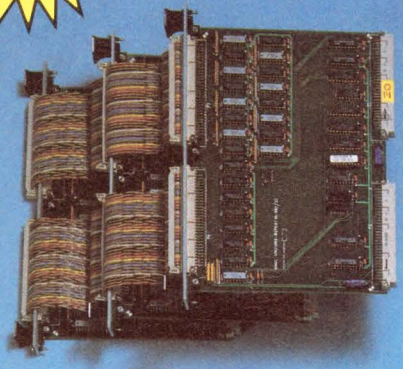

VMIVME-Repeat L-485/32
32 Bit Long Line Multi-Drop (A16,D32)
Connects up to 16 chassis
(up to 2000 feet)

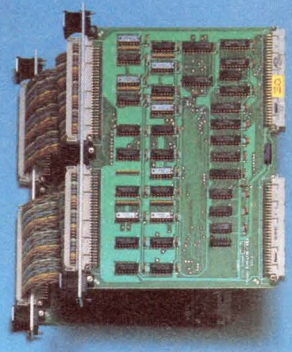

VMIVME-Repeat L-485
16 Bit Long Line (A24,D16)
(up to 2000 feet)

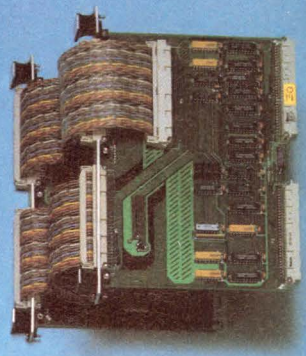

VMIVME-Repeat L
32 Bit High Performance (A32,D32)
(up to 50 feet)

DMA CONTROLLERS

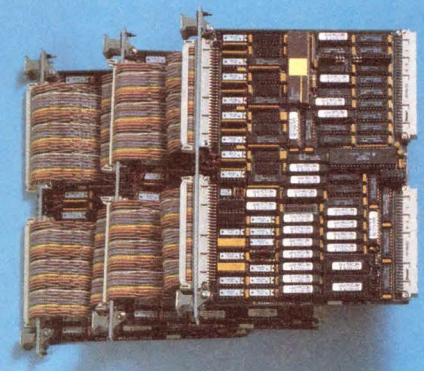

VMIVME-VMEnet II™
32 Bit Long Line Multi-Drop with
Arbitration and Broadcast (A32,D32)
Connects up to 31 chassis

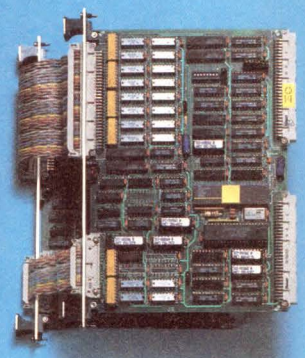

VMIVME-DMAL-485
32 Bit Long Line (A24,D32)
(up to 2000 feet)

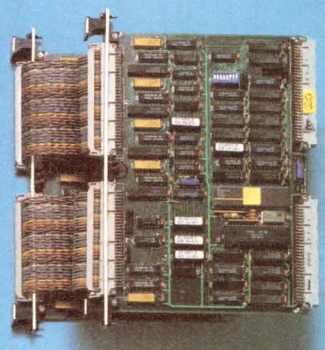

VMIVME—DMAL
32 Bit High Performance (A24,D32)
(up to 50 feet)

SYSTEM BUILDING BLOCKS

VMIC offers the widest variety of VMEbus high performance system building blocks in the industry. We pioneered the development of VMEbus Repeaters and DMA Link Controllers. VMIC is continuously developing high performance, state-of-the-art solutions to help you solve system expansion problems. When you need a VME-to-VME high performance solution, call VMIC, the VME CONNECTION!

VMEnet II™ is a trademark of VME Microsystems International Corporation.

OTHER SYSTEM BUILDING BLOCKS

- Host Computer Interfaces
- Digital I/O Boards
- Analog I/O Boards
- Synchro/Resolver I/O Boards
- Serial I/O Boards
- Intelligent I/O Controllers

Only VMIC offers you over 60 VMEbus board level products! You also get fast delivery, a two year warranty, and VMIC's 24 HOUR TECHNICAL SUPPORT TEAM for around the clock customer service!

FOR MORE INFORMATION CALL:

VMIC VME ™

VME Microsystems International Corp.
12090 South Memorial Parkway
Huntsville, Alabama 35803
(205) 880-0444
Toll Free 1-800-322-3616/Ext. 489
Fax (205) 882-0859

INTERNATIONAL DISTRIBUTOR: Diamond Point International, Inc. • 9 Enterprise Close, Medway City Estate, Rochester-upon-Medway, Kent (44) (0634) 722390 • FAX: (44) (0634) 722398

CANADIAN DISTRIBUTOR: Tracan Electronics Corp. • 1200 Aerowood Dr., Unit 4, Mississauga, Ontario L4W 2S7 • (416) 625-7752

CIRCLE NO. 38

One small step for PQFPs. One giant step for service

Introducing the right technology at the right time—our two-piece Micro-Pitch sockets for JEDEC plastic quad flat pack ICs.

Our patented low-profile (.400″) Micro-Pitch sockets are a remarkably simple, safe way to mount high-speed, high-performance PQFPs. And their simplicity—in engineering, and in the minimum use of material—brings them to you at a remarkably low price.

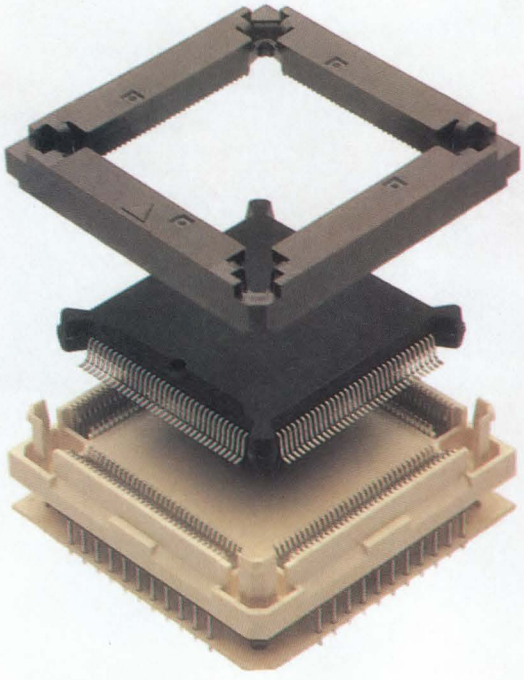

PQFP as easy as 1-2-3. .025″ centerline Micro-Pitch sockets ease handling, automate readily. High-speed contacts provide oxide cleaning wipe action, 200 grams normal force for positive mating. Footprint on .075″ x .100″ three-row grid.

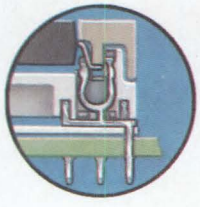

AMP is a trademark of AMP Incorporated.

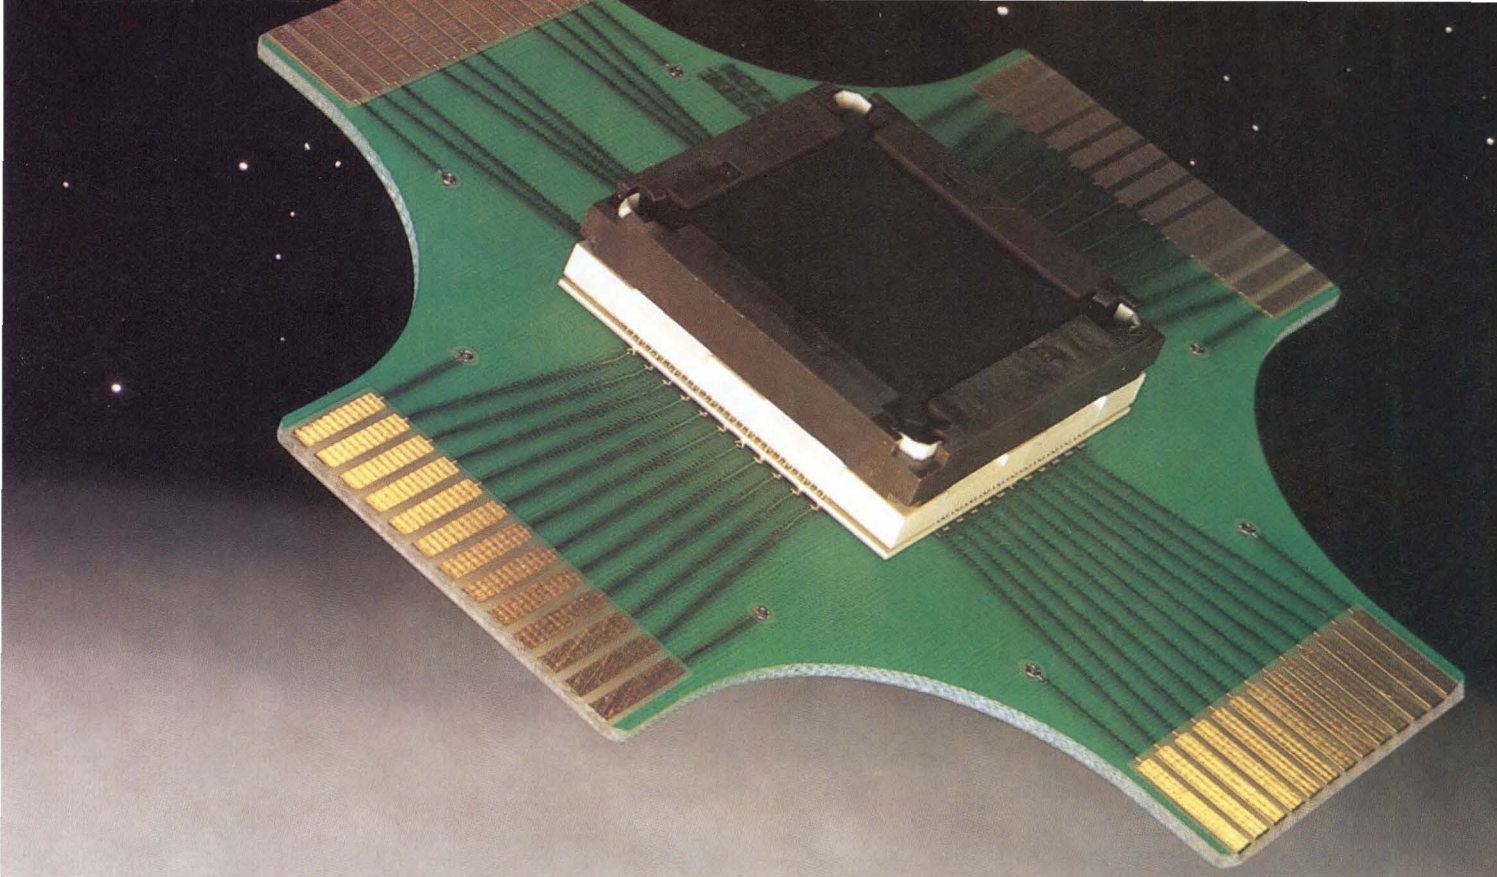

and manufacturing.

The cover aligns IC leads for easy insertion into the base. It also protects the leads during handling. In fact, your PQFP supplier can insert chips into covers at his site, and send completed units to you. Result: hands-off installation.

Our high-speed tin-plated contacts provide .020″ to .030″ of contact cleaning wipe during insertion, and a normal mating force of 200 grams on each contact. Very positive engagement. And just the kind of quality engineering you can count on from AMP.

Micro-Pitch sockets have been designed with automation in mind, as well. They're fully polarized cover-to-base, and base-to-board, and compatible with vacuum pickup, for fast robotic implementation.

With today's chip costs, and with the high quality and low installed cost AMP is known for, this is one option you'll want to pick up. Fast.

For technical literature and more information, call the AMP Product Information Center at 1-800-522-6752 and ask about Micro-Pitch Sockets. AMP Incorporated, Harrisburg, PA 17105-3608.

AMP Interconnecting ideas

CIRCLE NO. 39

Software protocols smooth the path to a standard graphics interface

Tom Williams
Senior Editor

As designers steamroll toward a universal windows-based interface, PCs and high-powered graphics systems will converge on common ground.

The computer industry is working to develop a common user interface that will reduce the user-perceived variations now readily apparent in different computers and operating systems. Whatever final form this common interface takes and whatever variations and grace notes individual vendors provide, it seems certain that users will interact with systems and their applications via multiple graphics windows on the screen. And users will invoke functions by means of graphical buttons, switches, menus and scroll bars.

Ultimately, this graphical user environment will be multitasking; users will access multiple tasks from their own machines or from a single-tasking computer or terminal over a network. Network technology and computational power that can provide these capabilities are already in place, but there's still a great disparity in the level of graphics performance between the high-end—mostly Unix-based—systems and the systems that are migrating up from what were once looked down upon as "mere PCs."

While disparity will always exist, there's a need to find a common ground, a level of graphics capability at which lower-end systems can look at some representation of the graphics output of the high-end machines (although maybe not in real time or in full-perspective three-dimensional form). High-end workstations and "personal supercomputers" have mostly proprietary ways of doing superfast, high-resolution 3-D interactive graphics for such tasks as real-time simulation and realistic modeling of physical objects. That's not likely to change because these methods are a major way in which manufacturers differentiate their products. Still, to meet the high-performance systems halfway, the lower performance machines—or high-end PCs—need better graphics capability than that offered by the standard IBM modes.

The need for graphics performance that surpasses that of the IBM standards has become especially evident as the new breed of systems based on 20- to 25-MHz versions of the Intel 80386 have arrived—machines in the class of Unix-based workstations. IBM-defined graphics controllers can't keep up with the computation power of these systems. The high-performance graphics subsystems developed for such machines are for the most part divorced from any standards and are aimed at specialized niche application areas.

The two highest IBM standards are the Video Graphics Array (VGA) with 640×480 pixels and 16 colors, and the 8514/A with $1,024 \times 768$ pixels and 256 colors. VGA was introduced on the mother board of IBM's PS/2 family, and 8514/A was offered as an add-on option for those doing CAD work or other applications needing higher resolution. While the resolution of 8514/A approaches that of higher-class systems—the majority are $1,280 \times 1,024$ pixels—resolution is only part of the issue. Performance is a much more pressing concern.

Performance is important because a window environment is interactive—most users may not be twirling a molecule around in 3-D space, but every user will be opening, closing, sizing and repositioning windows on the screen, scrolling text and calling graphics images into windows. These operations require the fast manipulation of blocks of pixels—bit-block transfers, or bitblts. But performance measured simply in terms of drawing or bitblt speeds unrelated to resolution can be misleading. A 1k- $\times$ 1k-pixel screen has about four times the pixels that

the 640-×480-pixel VGA resolution has, so a speed that seems high-powered for VGA will bog down at the higher resolution and be unacceptably slow.

CPU can't be overloaded

For very high-end systems, such as the Titan Graphics Supercomputer from Ardent Computer (Sunnyvale, CA) and the DN10000 from Apollo Computer (Chelmsford, MA), there's an advantage to having the main CPU do as much graphics computation as possible, leaving only the actual generation and writing of pixels into the frame buffer to be done by specialized hardware. This arrangement allows a tighter coupling of the graphical representation of an object to its underlying computational model and enhances interactivity between the user and the computational problem.

But to achieve graphics performance that can adequately reflect the computational capabilities of their CPUs, even fast 32-bit microprocessor systems can't afford to burden the CPU with large blocks of routine drawing operations. Even more important, such systems can't have system bus bandwidth tied up by the CPU spitting pixel data

across the bus to the graphics frame buffer. For resolutions beyond VGA to show acceptable performance, the CPU must communicate across the bus at a higher level of command than mere frame buffer addresses and pixel data to reduce bus overhead and free CPU cycles.

A discussion of "higher levels of command" means getting into the world of machine instructions, in which short commands invoke much longer sequences of activities from their respective devices. "Graphics engines are data amplifiers," says Steve Dines, strategic marketing manager at Advanced Micro Devices (Sunnyvale, CA). Given the command "draw circle x, y, r," for example, an intelligent graphics device will write pixels to the frame buffer, defining a circle centered at coordinates x and y with the radius, r, writing a sequence of pixels to the locations defining the circumference. The larger the single value r, the greater the amount of actual pixel data that's generated.

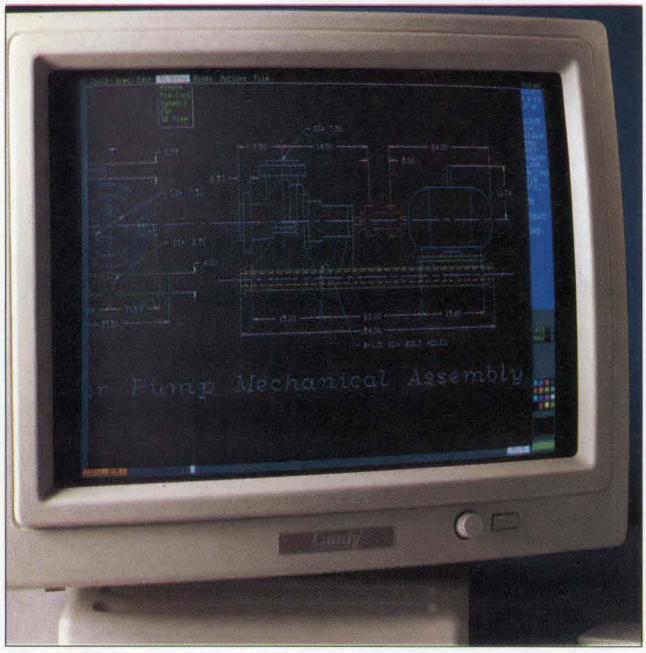

The Lundy Electronics 1612 graphics subsystem brings workstation-class graphics—1,600×1,200 pixels with 16 colors—to an AT controller card. Based on the Texas Instruments 34010 graphics processor, the controller supports the Direct Graphics Interface Standard (DGIS) from Graphic Software Systems and can run any application that has a DGIS driver.

And since there are many possible ways to actually implement a circle with hardware, the nature of the actual machine instructions invoked can vary when designers communicate with the graphics subsystem on even this fundamental level. Further, designers create a layer of software that bestows a degree of freedom for designs in silicon that isn't possible with dumb controllers such as the VGA. This software "membrane" shifts the major issue from one of hardware-register compatibility with some graphics device to one of being able to interface to a software protocol layer, which in turn may be adapted to a wide variety of graphics hardware engines.

While an applications device driver is written for a specific peripheral or controller board and must deal with all of the idiosyncrasies of its hardware interface, a protocol layer acts as a virtual device interface. The protocol layer masks the hardware details from the application and lets applications access, via a single driver, as many devices as the protocol supports. Among the various graphics protocols in use are the Direct Graphics Interface Standard (DGIS) from Graphic Software Systems (Beaverton, OR), the Renaissance Graphics Device Interface (RGDI) protocol from Renaissance GRX (Bellevue, WA), the Texas Instruments Graphic Architecture from Texas Instruments (Dallas, TX), and the 8514/A Adapter Interface (AI) protocol from IBM.

Sharing the computation burden

Interface protocols also perform an additional function. They distribute the computation burden between the graphics controller hardware and the main CPU, depending on the capabilities of the graphics hardware. Various graphics commands supported by a protocol (draw circle, draw rectangle, fill polygon and so forth) might all be offloaded to a very intelligent graphics device, for example. But when used with a controller of lesser capability, one or more of these same commands might be directed back to the main CPU, with a corresponding performance penalty. In this case, the CPU itself must translate the commands into a code sequence understood by the device. This is the case with the 8541/A, where the CPU must define a circle in terms of vectors because the AI doesn't support a circle draw command.

How to distribute the computation burden is a decision that's made when the OEM adapts a protocol to a selected controller. However the burden is distributed, though, an application program still uses the same driver for all devices supporting the protocol and notices only the difference in performance between them. This is a distinct advantage because it not only eases the task for the software application developer, but also lets the OEM offer a product line with a range of price/performance levels that still require only one driver from each application for a whole family of boards.

The obvious problem here is that while there are fewer device-level protocols than there are different controllers, there's still no industry-wide standard protocol. IBM's 8514/A hasn't generated wild enthusiasm among users of the technology, which has limitations—it doesn't support circles and arcs, for example.

There are, nonetheless, many who feel they have to support 8514/A—either as the main interface to their products or at least as an option. Among those are Stephen Andes, vice-president of Enertronics (St. Louis, MO). But Andes says that as a software developer, he doesn't "feel hemmed in at all. IBM sets the standard, and the standard is 8514/A." Still, Enertronics' Aurora 1024 product, which is based on the TI TMS34010 graphics processor, does offer the DGIS protocol as an option.

Another criticism of 8514/A is that it's more than an interface protocol. 8514/A refers to the adapter interface as well as to its underlying processor. But 8514 (without the /A) also specifies a monitor, which some see as a definite limitation. The monitor used to display 1,024×768 pixels is specified at 44.90-MHz interlaced scan, which places limits on the resolution that's practical with 8514/A. By contrast, Lundy Elec-

Achieving 8514/A compatibility: hardware and software trade-offs

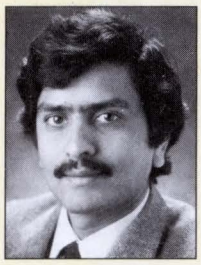

There are at least three software interface layers for the 8514/A: the 8514 Adapter Interface (AI) from IBM, and the Windows 2.1 and Microsoft Presentation Manager application interfaces from Microsoft. All three are independent of each other, but they provide similar functionality to an application program and can thus be considered parallel layers of software. There will undoubtedly be more such parallel layers in the future.

For manufacturers and developers, the major issues are requirements for an 8514/A-compatible product. IBM hasn't published a hardware interface specification to the 8514/A, and the company's position is that all software must use the 8514 AI to communicate with 8514/A hardware. This indicates that IBM is likely to have future products that won't be hardware-compatible with the current 8514/A. Such a product could have its own AI, offering compatibility with application programs at the AI level.

For those interested in making an 8514/A-compatible product today, there are three major possibilities. First, the manufacturer could develop a product using a different hardware implementation. Second, a manufacturer could develop a product that's 100 percent register- and gate-level compatible with the IBM 8514/A. Third, a manufacturer could develop a product that's compatible with the IBM 8514/A at the register level but also has significant new features. Each of these approaches offers advantages and disadvantages.

A product that uses a different hardware implementation wouldn't be IBM-compatible at the hardware level, but it could have an interface to make it compatible at the AI level. With this approach, the developer can optimize the system in terms of cost and performance, for example. In addition, he can choose the algorithms to do the low-level graphics functions.

For maximum performance and ease of implementation, he can move the hardware/software dividing line between the AI and the hardware. It's also possible to add new features in the hardware that the AI could exploit. And the product will run all software that uses the 8514 AI.

But there are several disadvantages that are associated with hardware noncompatibility. For one, the new product wouldn't be able to run IBM's version of the AI—although this is more a marketing handicap than a real problem. Because every 8514/A-compatible product is both a hardware product and a software product, developers must come up with their own AI that can be used with the hardware. But doing so may lead to customer doubts about product compatibility.

Ultimately, very few application programs will communicate directly with the 8514/A hardware.

■ ■ ■

In addition, such a product wouldn't be able to use the standard 8514/A display driver in Windows 2.1. But developers can ship a Windows 2.1 display driver for their products. The new product also wouldn't run other potential software that communicates with the 8514/A hardware directly. Theoretically, as new software arrives, the hardware developer could write a driver to make it run on the product.

■ Complete compatibility

One advantage of the second approach—developing a product that's 100 percent register- and gate-level compatible with the 8514/A—is that the product can use any software that runs on the IBM 8514/A, even IBM's AI. Unfortunately, however, the issue isn't actually that simple, since developers can't legally ship IBM's AI with their hardware and, therefore, must develop their own AI. But those developers will have an easier time convincing customers that the product is compatible if it has 100 percent register- and gate-level compatibility.

The disadvantage with this approach is that the developer must make sure the AI doesn't infringe on IBM's copyrights. What's more, hardware costs will increase if the developer tries to significantly improve performance while maintaining compatibility.

The third product-development scenario—developing a product that's compatible with the IBM 8514/A at the register level but has significant new features—offers the greatest promise. With this solution, the new product would be able to use all software that runs on the IBM 8514/A. Because of this compatibility, the developer will find it easier to convince customers that the product is compatible. In addition, the developer can add new hardware features, such as color expansion and reduction, that can be used by the AI.

Some possible shortcomings to this approach are that the hardware would be very complex, thus requiring longer design cycles. Maintaining compatibility at all levels is difficult.

■ Software interface important

It's becoming evident that an 8514/A-compatible product must be able to run all software that communicates with the 8514/A, including IBM's version of the AI. Ultimately, very few application programs will communicate directly with the 8514/A hardware. The AI is here to stay and will continue to be enhanced. Today's product developers won't be able to design a software interface, however, that takes care of all the needs of tomorrow's software.

In the future, systems using Micro Channel and Extended Industry Standard Architecture buses will feature multiple processors, which may not be architecturally compatible with each other. There will be coprocessor cards with Sparc, MIPS R2000 and R3000, 680X0 and even IBM RT processors, among others. These processors won't be able to run the standard AI, which has been written for the Intel 80X86 family of processors. There will thus be many versions of AIs and other interface layers such as Windows Application Program Interface (API), Presentation Manager API, and Computer Graphics Interface. There will also be different interface layers for Unix, X Window, Display Postscript and other operating environments. So although application programs won't communicate with the 8514/A hardware directly, there will be several system programs that will do so.

Arun Johary, MSEE, Technical Marketing Engineer, Chips and Technologies

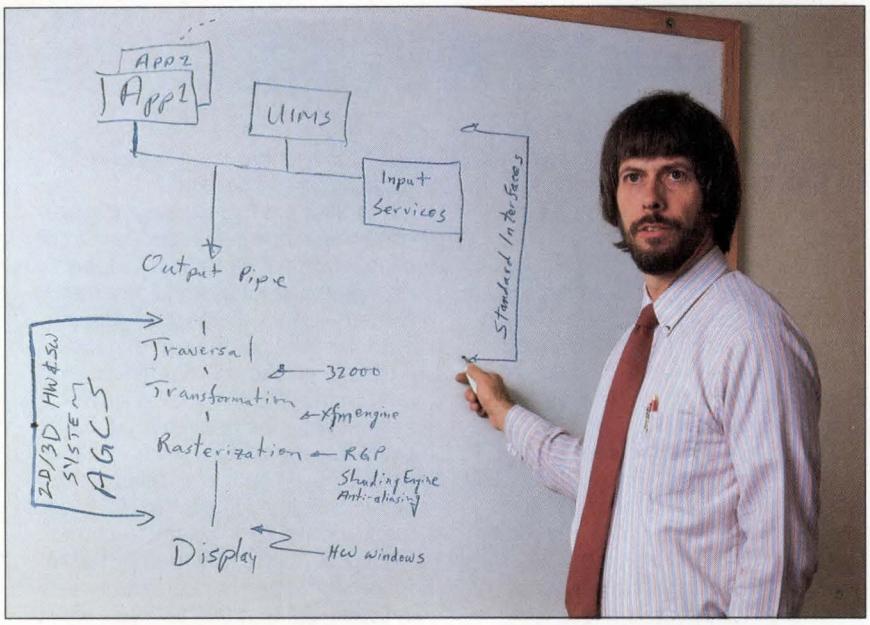

The challenge for manufacturers is to offer an interface that both supports current applications and prevents new silicon developments from outpacing its capabilities, says John Blair, National Semiconductor's graphics strategic marketing manager. National has chosen the Direct Graphics Interface Standard (DGIS) from Graphic Software Systems as the interface for its products.

tronics (Glen Head, NY) produces the 1612 graphics subsystem, which includes a TMS34010-based controller on an AT card and a noninterlaced 19-in. 160-MHz monitor. The 1612 achieves resolutions of up to 1,600×1,200 4-bit pixels (16 colors) and offers the DGIS interface along with the DGIS driver for Windows from Microsoft (Redmond, WA). "The killer [with 8514/A] is that it's interlaced," says Jeff Leibowitz, marketing manager at Lundy.

Keeping up with the silicon

It appears that IBM, while making a good deal of headway with 8514/A, hasn't solidified support because the standard isn't perceived as providing the path to even higher performance that may be required to work in the same environments as high-end workstations. IBM's advantage is that it presented the AI at the same time that it introduced the 8514/A silicon. Graphics processors that appear to have much better performance than the IBM silicon, notably TI's TMS34010 and the DP8500 raster graphics processor (RGP) from National Semiconductor (Santa Clara, CA), were introduced without a clearly defined interface protocol. Both TI and National now somewhat regret that decision because they realize that their processors would have been better accepted had the companies provided such an interface along with the silicon.

"The challenge is to have an interface that's rich enough not only to support applications and people's ideas of what they need to do now, but that's also able to shield new developments in silicon from outpacing what the interface can do," says John Blair, strategic marketing manager for graphics at National Semiconductor. National recently announced an agreement with Graphic Software Systems (GSS) to establish DGIS as the definitive interface for the DP8500 RGP and family of graphics support chips.

"We're relatively late in the game getting involved with DGIS," says Blair. "Originally, we saw ourselves based on Unix, high-end sort of things where DGIS wasn't a prevalent force." But National's sights haven't dropped to a lower performance target. Rather, the "low end" has come up to where it needs the graphics power offered by a controller such as the RGP.

Like similar interface protocols, DGIS offers a set of graphics operations, including line, circle, ellipse, polygon, splines and filled shapes, to name a few. DGIS operations are accessed by standard commands issued by the application driver—"output filled ellipse," for example. Each command is then implemented in the native instruction set of the target hardware, so that a line draw command invokes a routine of 8500 RGP code, or a circle command a series of 34010 operations.

Alternatively, some operations might be implemented in 80286 or 80386 code to compensate for graphics hardware with less capability or to balance the computation load between CPU and graphics controller. To implement these commands, the CPU must either translate them into a sequence the target controller can perform or have direct access to the bit map itself. Although these CPU-specific code modules are almost all implemented for the Intel instruction sets, they're also all written in C. It's thus entirely possible to adapt DGIS to other CPUs, primarily the 68000 family, says Jim Cochell, hardware business unit manager at GSS. No announcements have yet been made, however.

An interface such as DGIS provides even more than a shield to the graphics hardware. Since DGIS' standard commands invoke native code on the graphics controller side of the interface, one can alter the actual algorithms as long as their results are unchanged. "A hotshot board maker could try to optimize certain parts of the code, rebalance the load between host and board and not disturb the basic interface," says National's Blair.

But balancing is a dynamic process that can be quite complex, he warns. In a multitasking environment, for instance, the CPU is almost always busy, so it's probably better to take all the graphics processing off the host. Blair also notes that one of the reasons users haven't been getting optimum performance is the lack of adequate analysis tools such as profilers, which let the user see which routines are running more than others and how long certain processors are sitting around with idle cycles.

Still, changing applications or plugging a board into a different kind of system or network can change the demands placed on it. Since the wave of the future is toward multitasking, there will be an even greater demand to place more of the graphics workload on the controller and to increase intelligence on the graphics side.

Software compatibility

TI, whose 34010 graphics processor is used in more than 240 board designs—including designs on VMEbus and Multibus—decided to

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

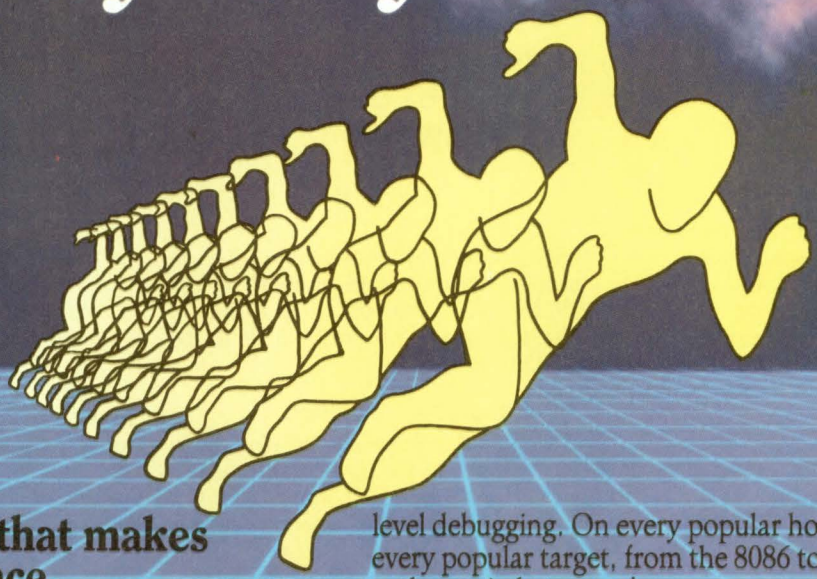

The 10% edge that makes a huge difference.

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

Work in the fast lane.

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

Get complete support.

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

Put us to the test!

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

The big difference in embedded systems development.

*And we have the benchmarks and data on file to prove it!
†• 68000/008/010/020/030; 68HC11/00/05/09; Z80/64180
• 8086/186/286; 8051; V20/25/30/33/40/50
• AM29000; V60/70; VAX; PDP-11; and many more.
‡In MA (617) 661-0072.

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

Intermetrics

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

CIRCLE NO. 40

support the DGIS interface a couple of years before National, but found that DGIS alone didn't render all 34010-based boards software-compatible. Users designing in the 34010 also need a common software communications protocol to mask off the exact hardware communications interface, according to Karl Guttag, TI graphics strategy manager.

"Some boards use the 34010 host interface, but some put in a shared dual-ported RAM, and there are many variations," says Gutlag. Even though applications might call a common set of commands—either DGIS or subroutines supplied by TI—their drivers thus had to be adapted to the hardware interface, or there had to be a different adaptation of DGIS for every board.

The Texas Instruments Graphics Architecture (TIGA) supplies a common software communications protocol that can be used with all 34010- and future 34020-based boards. The TIGA protocol consists of a small piece of memory-resident code that runs on the host and communicates with the software interface on the 34010 side of the bus. The protocol sends data packets back and forth, and the packets can be of variable size so they can be optimized for individual commands. The TIGA protocol also supports DMA transfers so that large display lists can be transferred quickly and the processor simply told to display them.

TIGA also supplies a set of graphics routines that are similar but not identical to those of DGIS and RGDI, a graphics protocol developed by Renaissance GRX specifically for the company's 34010-based boards. "The most important part of TIGA isn't the actual graphics routines," says Guttag. "It's the fact that it's a single communications protocol so that, regardless of how the hardware is implemented for communicating between the host processor and the 34010, there's one protocol setup for communications."

The existence of a common communications protocol gives designers several options: use TIGA commands for quick development time, use C calls to the TI graphics subroutine library, or implement their own optimized graphics routines. TIGA is thus automatically extensible since anything written in 34010 code will run through TIGA's communications protocol.

This sounds fine for applications written for the 34010, but as National's Blair points out, "The hardware developer needs a way to get all the applications out there up and running on his product quickly and efficiently." Software developers have a similar interest in wanting their applications to run on as many hardware platforms as possible. This brings a processor-independent interface such as DGIS back into the picture. Work is underway to embed the communications features of TIGA into DGIS so that programs with DGIS drivers will instantly be able to run on boards supporting the DGIS/TIGA protocol, according to Guttag. The same will be done for RGDI, Guttag says.

Designers must consider the trade-offs involved with using a software interface.

Getting down to silicon

Despite all the talk of device-interface protocols for intelligent processors, there's still a school of thought that maintains that inserting a software layer between the application and the hardware compromises performance. In strict terms this is true, of course, but designers need to consider trade-offs between how badly a software interface affects performance versus how well it lets them access high functionality and provide application portability. They also need to think of ways they can optimize between any performance penalties and the advantages bestowed by software interfaces—not a trivial problem.

The first caveat for board manufacturers is that by putting nonstandard hardware on the market, they have only limited choices. They can support the hardware themselves by writing and maintaining a vast array of application drivers, they can convince application developers that their board is so important that they'll do the work of writing the drivers, or they can support a software interface standard.

As long as higher performance graphics products remain a niche market, companies will have success designing specialized boards and supporting drivers for the most popular applications in their chosen markets. In addition, the advent of standard user-interface environments—such as Microsoft Windows and Presentation Manager, and the X Window System in Unix—will cushion the board maker from having to support too many individual drivers. A large number of applications written to a window environment will run through that window system, which can be adapted to a given manufacturer's hardware.

Still, it's safe to say that there will always be important applications, such as Lotus 1-2-3 and Wordperfect in the DOS world, for which the designer must choose to circumvent a window environment; many such applications already exist. In addition, the prices of high-resolution monitors as well as graphics controllers continue to drop, bringing the

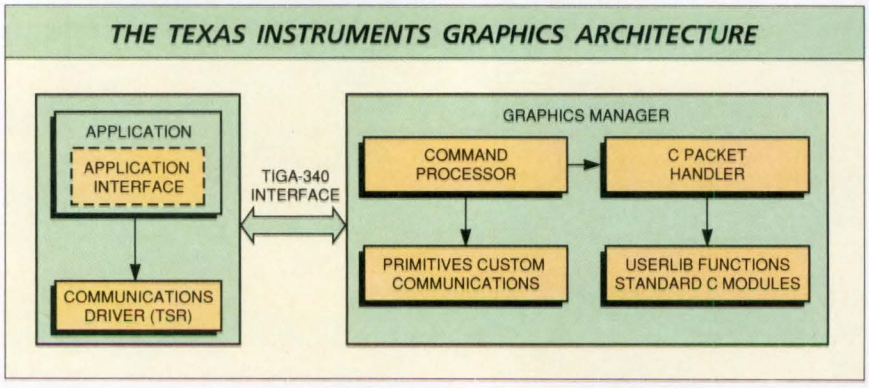

The Texas Instruments Graphics Architecture (TIGA) consists of three parts: an application interface, a communications driver and a graphics manager. The communications driver resides on the host and provides a common communications protocol with the graphics manager. The command processor handles the communications protocols on the 34010 side of the bus. The actual graphics functions of the TIGA standard are optional and extensible, since any application code written in 34010 instructions will run on the 34010-based graphics subsystem as well.

Moving beyond VGA for mainstream graphics

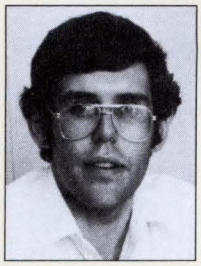

Today there are three main PC graphics alternatives to moving beyond the 640 × 480 × 4 bits-per-pixel resolution offered by IBM's Video Graphics Array. One choice—simply adding new resolution modes to the current IBM standard resolutions—doesn't involve adding a graphics processor. The other two alternatives—using IBM's 8514/A or using the Texas Instruments 340 family (used by Compaq, Hewlett-Packard, Wyse and others)—increase graphics speed with a processor.

While adding higher resolution modes to VGA sharpens the screen image, keeping up with the increase in display data requires more graphics processing. And although some of the extended VGA-type boards have improved the host's access to the display memory, they haven't dramatically changed the system's graphics processing power.

A graphics processor can greatly improve the responsiveness of graphics-intensive CAD and desktop publishing applications. Because the display memory, color palettes and monitors are common to systems with the same resolution, using a graphics processor increases system cost only slightly. The combination of higher performance without a significant effect on system cost is leading to more widespread use of graphics processors.

Fixed or programmable?

The graphics processor relieves the host from what can be massive data manipulations. A fundamental issue has been whether the graphics processor should have a fixed set of preprogrammed functions, as does the 8514/A, or be a fully programmable microprocessor, as are the 34010 and the 34020.

The attraction of a fixed-function processor is that it's initially easier to design and support. But a fixed-function processor can't support all the functions that an application requires. The host processor must therefore translate required operations into command sequences that the fixed-function processor can handle. This extra step can defeat much of the reason for having a graphics processor.

A fully programmable graphics microprocessor, on the other hand, can download and execute all graphics functions independently of the host, thus relieving the host of processing any graphics operations. In effect, the command set of a fully programmable processor can be extended to meet the needs of any graphics application. And, because a fully programmable processor can be used in many applications, it can cost less than a fixed-function one.

Software standards needed

The added capability of a graphics processor increases the need for a standard software interface. Most software developers don't want to become intimately familiar with a specific system architecture. A standard software interface lets developers work at a higher software language level and gives them access to a larger installed base.

A standard interface must not, however, add so much overhead that the advantages of a graphics processor are diminished. If the interfaces add too much overhead, applications will need to be directly ported to the specific hardware to avoid a performance penalty.

Another issue is the standard's level of processor independence. While standards that are completely processor-independent are more portable, they can lose the ability to take full advantage of a specific processor's capability. Ideally, therefore, a standard should give independence from the exact hardware implementation but shouldn't slow down graphics processing.

The 8514/A and 340 processors both have a software interface: AI (Adapter Interface) for the 8514/A and TIGA (Texas Instruments Graphics Architecture) for 340 processors. While both interfaces are potential software standards, they differ significantly because they reflect the types of processors for which they were designed. AI supports a fixed set of functions that isn't currently extensible, and AI's primary purpose is to maintain compatibility with future hardware that will supersede the 8514/A.

TIGA was designed for a programmable processor family to support such features as downloadable functions, dynamic linking and memory allocation. At the lowest level, TIGA is simply a communications protocol specifying how the PC's host processor transfers information to and from the 340 processor. Tasks are split so that the 80X86 CPU executes operations it does well, such as display list processing, and the 340X0 handles the drawing operations. The host thus can be relieved from processing graphics commands. TIGA also defines a standard set of graphics primitives to make it easy for developers that don't need more direct control. But at the same time, TIGA lets developers download their own custom extensions if they don't want to use the standard primitives.

Open architectures

When moving beyond VGA, developers need to make a choice between using a commercially available processor or cloning an existing system. IBM's 8514/A is a proprietary architecture that will have to be successfully cloned if it's to become a standard outside of IBM. As graphics processors are called on to bring more detail to the display at a faster rate, their architectures are getting as complex as host processors, which in turn makes them difficult to clone. This trend is indicated by the fact that there are more compatibility issues associated with VGA than there ever were for its predecessors, the Enhanced Graphics Adapter and Color Graphics Adapter standards.

For a graphics architecture to achieve mainstream acceptance, it needs an open architecture for wide usage. Using a commercially available processor avoids the compatibility issues that can come up when cloning proprietary architectures.

Now that TIGA provides 34010-based systems with a common high-performance software interface, the 340 architecture is in a strong position to become the next mainstream standard beyond VGA. TI plans to make the TIGA specification public in the near future, as well as support it with tools for both hardware and software developers.

Karl Guttag, MSEE, Graphics Strategy Manager, Texas Instruments

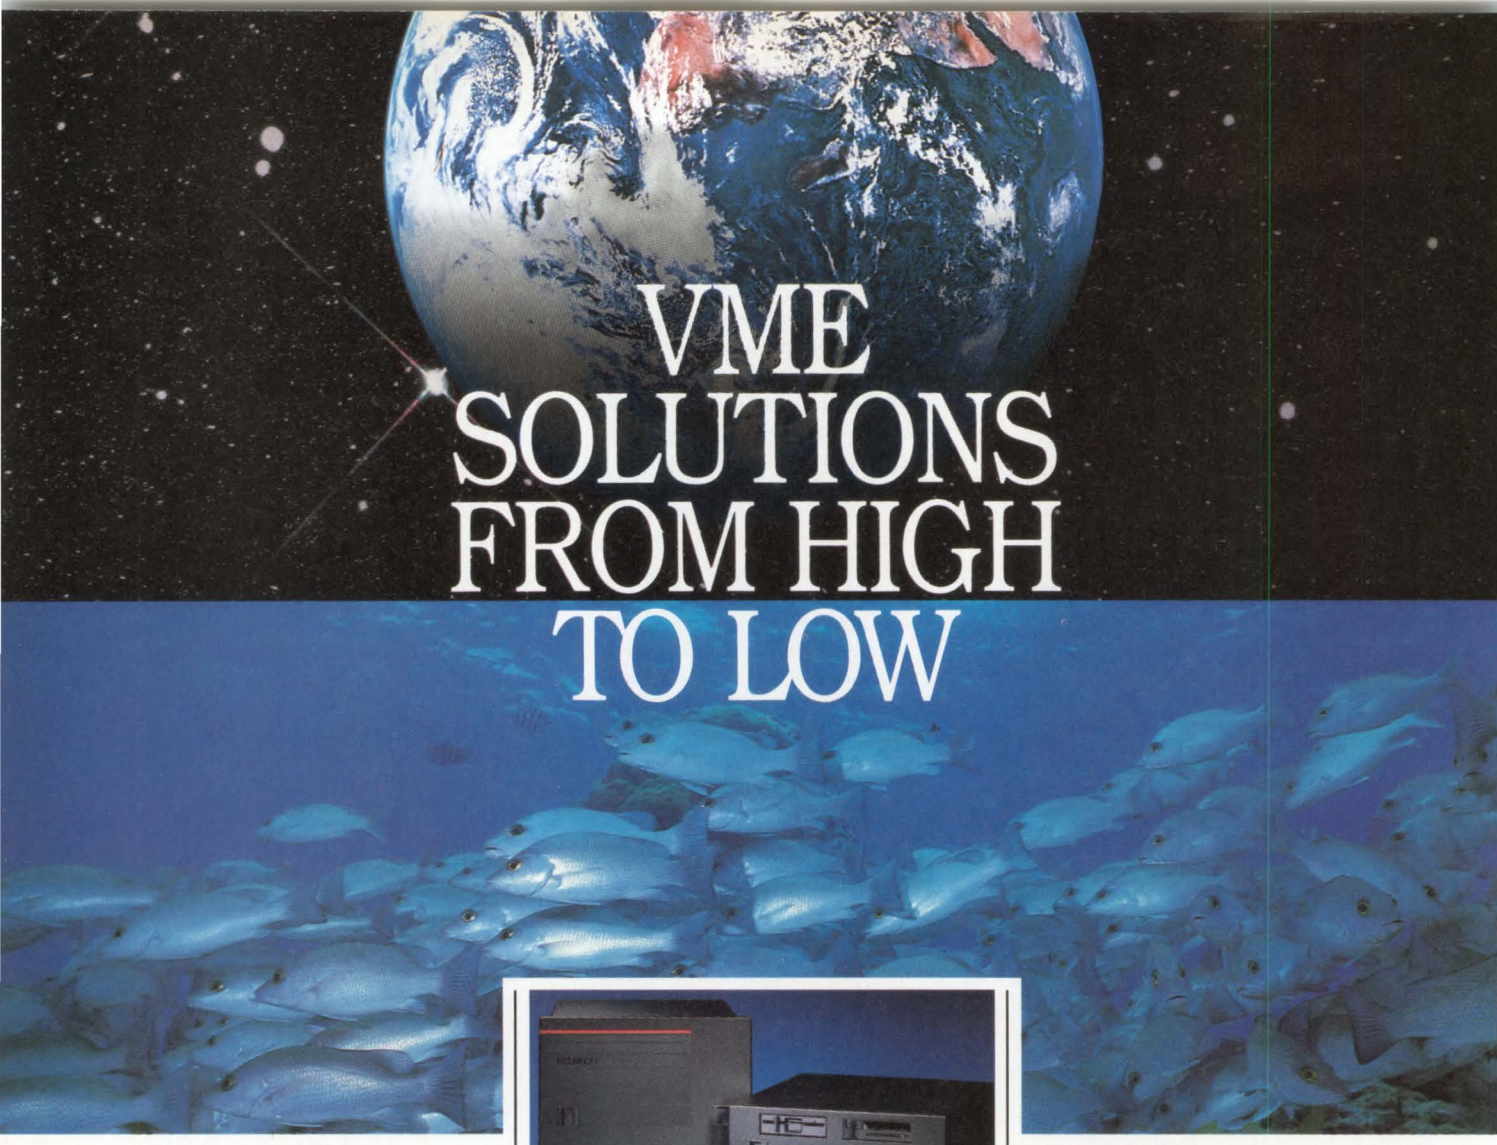

VME SOLUTIONS FROM HIGH TO LOW

The Power to Serve. Whatever your application, high or low-end, Heurikon has the VME power to drive it. From UNIX™ to real-time. From satellite communications to oceanographic research. We can solve the problem.

A New Dimension in Processors. Get into the fast lane with Heurikon's 16 or 32-bit microcomputers. Your choice of 68000-68030 or NS32532. Single boards or systems. Off the shelf or customized. With unmatched support from day one.

Potent Development Tools. Slash your development time. Choose from multiple networked or bus-based environments. Wind River VxWorks™ and Heurikon BusLink™. SUN, VAX and other hosts. Plus solid in-house support for UNIX V.3 and VRTX® or OS-9™ real-time operating systems ensures the integrity of your design.

From high to low, power your application with VME from Heurikon. Call us today.

1·800·356·9602
ext. #501

Heurikon Corporation
3201 Latham Dr., Madison, WI 53713
TLX: 469532 FAX: 608-251-1076

UNIX is a trademark of AT&T Bell Laboratories, Inc. VxWorks is a trademark of Wind River Systems, Inc. BusLink is a trademark of Heurikon Corporation. VRTX is a registered trademark of Ready Systems, Inc. OS-9 is a trademark of Microware Systems Corporation.
© 1988 Heurikon Corporation.

HEURIKON CORP.

CIRCLE NO. 41

day closer when displays that have 1,024-×768-pixel resolutions and beyond become a mass market with fierce competition.

Until now, companies in the higher resolution end haven't had to play to mass markets to be successful. Metheus (Beaverton, OR), for example, has chosen to distinguish itself in terms of performance by designing its own proprietary silicon for use on its Ultra Graphics accelerators, which support 1,024-×768-pixel displays. The company supplies drivers for selected applications. "Our strategy is to focus a lot of energy on the really important packages, and right now those are Autocad and Windows," says Gene Chao, Metheus president and CEO. The Metheus-supplied drivers will thus be written directly to the proprietary silicon to maximize performance.

Metheus has developed a communications interface protocol, called Magic, for those software developers whose packages Metheus doesn't support, but who want to write drivers to Metheus boards using the command set of the silicon. Protocols that provide standard graphics functions such as TIGA and DGIS "do have their place," Chao says, "but when it comes to performance, they're always going to fall."

Mass markets change the rules

But what were once niche markets are rapidly becoming mass markets where different sets of rules apply, among the foremost of which is wide-ranging compatibility. The question is, compatibility in what sense? Exact adherence to hardware, or functional software compatibility? Even in the mass-market arena, the game is shifting away from strict hardware compatibility toward compatibility with software functionality.

One company that has made a big success at IBM hardware compatibility in the VGA arena is Chips and Technologies (San Jose, CA). In moving into the emulation of the 8514/A specification, Chips is taking a cautious approach, vowing to maintain 100 percent gate-level compatibility with the 8514/A hardware, but realizing that the AI is probably the only standard that IBM will reliably support.

If IBM solidly commits to AI, Chips will be free to perhaps make innovations in the next generation of hardware that depart from strict gate compatibility, notes Sikander Naqvi, director of graphics at Chips and Technologies. "We believe in hardware compatibility, but we know this time that IBM is supporting AI. With that in mind, we can change direction and take advantage of a good software interface and optimize hardware," he says.

For its first generation, however, which has some major software written directly to the 8514/A silicon, Chips is committed to gate-level compatibility and claims it can achieve major speed improvements over IBM through hardware technology. Conventional wisdom has it that although Microsoft wrote Windows and Presentation Manager to the initial version of 8514/A hardware, IBM's next generation will feature a numeric coprocessor and an enriched AI. At that point, everybody—including Microsoft—is expected to write to the AI.

But Robert Butchko, president of Renaissance GRX, cautions against getting too caught up in the idea that protocol layers get in the way of speed. "A graphics engine, by its nature, can't possibly include all of the graphics library functions that are needed in a real high-performance graphics environment," he says. In other words, the controllers are programmable, and that means that *some* layer of software has to tell them what to do.

And the less programmable the graphics hardware is, the more likely that the CPU will have to intervene for some functions. Another alternative some companies use is to place a dedicated microprocessor on the board with the less intelligent graphics controller chip. But, as Butchko puts it, "The software interface is more than an interface; it actually does a lot of the things that can't be done in the silicon." It provides a curtain behind which the designer can develop optimized code and algorithms and still offer application developers a single, well-known programming interface.

But applications' drivers also have to play their part in improving performance by taking advantage of the power offered by standard software device interfaces. One application frequently used for performance comparisons is the Autocad drafting program from Autodesk (Sausalito, CA). Release 9.0 of Autocad has achieved significant speedups over earlier versions by requiring a math coprocessor on the host system to calculate vectors for line drawing. The only kind of command Autocad's drivers pass on to the graphics controller, however, are line-drawing commands. Thus, even if a board supports circles and arcs, all it gets from Autocad are circles and arcs

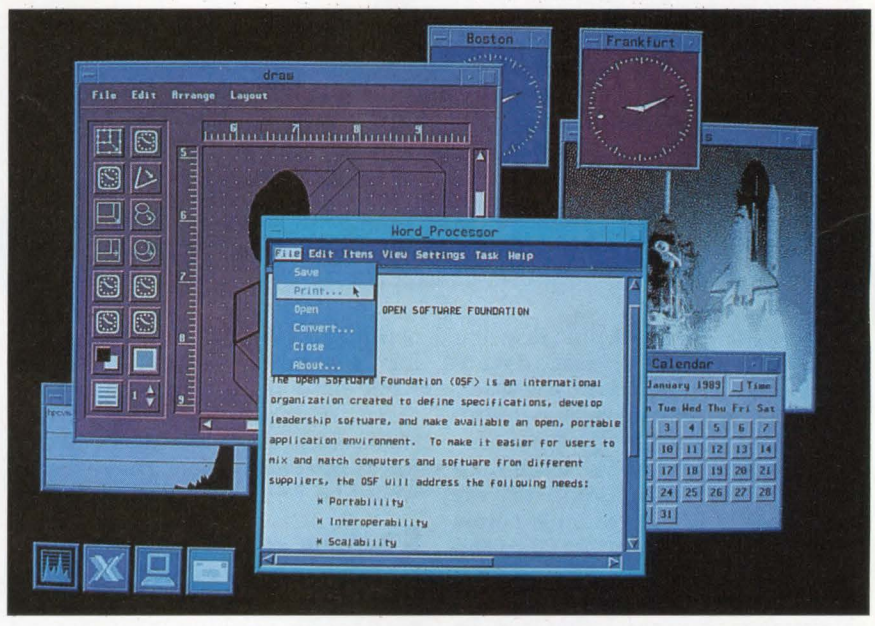

The OSF/Motif User Environment Component (UEC) was chosen by the Open Software Foundation for use with its Unix operating system. The UEC combines user-interface elements from Digital Equipment Corp, Hewlett Packard and Microsoft, maintaining the behavior of Microsoft's Windows and Presentation Manager with a three-dimensional appearance designed by HP. Because the UEC is based on the the X Window System Xlib graphics protocols, it can appear on different displays of varying resolution and performance as long as they contain an X Window server.

Interface brings Unix and DOS to high-performance PC graphics boards

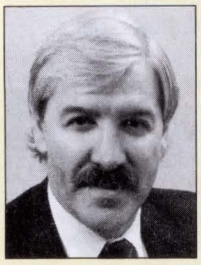

As Unix-based 80386 PCs continue to displace low-end dedicated workstations, there's a growing need for low-cost, high-performance PC graphics boards that are compatible with both Unix and DOS graphics applications. It's feasible to develop such boards by using graphics processors from Texas Instruments, National Semiconductor, Intel and Hitachi. The challenge, then, for many hardware designers is how to ensure software compatibility. High-end PC users need maximum software compatibility because they want to exploit the general-purpose applications advantage that their machines have over dedicated graphics workstations. They also need to use specific Unix-based software solutions.

To help designers in this task, Graphic Software Systems has developed a graphics protocol called the Direct Graphics Interface Standard (DGIS), a board-level interface for DOS- and Unix-based application software compatibility. More than 50 high-resolution graphics hardware manufacturers have implemented DGIS firmware.

DGIS firmware cuts the development cycle required to produce a software interface for graphics boards. As an off-the-shelf software interface, DGIS provides immediate support for application software packages in DOS and Unix. Porting existing DGIS firmware to a standard graphics coprocessor-based design requires about one engineer-month, including testing.

Without a standard graphics library such as DGIS, hardware manufacturers would need to dedicate as much as ten times more development resources. Testing would require additional engineer-months, and each application device driver would require even more. Maintenance and driver support would continue to draw on the manufacturer's resources. Furthermore, this effort would need to be repeated if the designer changed chips. The combination of these tasks would result in added development expense and increased design cycles, which would ultimately increase time to market.

DGIS provides application software compatibility for graphics boards based on coprocessors from TI, Intel, Hitachi and, soon, National Semiconductor. Graphics hardware designers are assured that the software interface will be preserved as they change from one chip to another, develop around more than one chip, or migrate to new-generation chips available.

How it works

The TMS34010 from TI (Dallas, TX) is among the graphics coprocessors supported by DGIS (see figure). Because the TMS34010 is a completely programmable device, the only requirement placed on the host is to pass the operations and data to the TMS34010 and, if necessary, to wait for information to be returned. All graphics algorithms and all graphics processing occur on the TMS34010. DGIS graphics functions (which in this implementation all run directly on the TMS34010) are implemented using these three methods:

- directly calling the TMS34010 graphics primitives;
- emulating functions that aren't directly supported by calling the constituent TMS34010 graphics primitives; and
- emulating other functions by touching the pixels directly in the bit map with the pixel instructions of the TMS34010.

When software requests a DGIS operation that matches the Graphics System Processor's graphics capability, the request is placed in the format of the TMS34010 instruction and executed. When the requested operation doesn't match the graphics capabilities of the TMS34010, the operation can usually be broken down into a series of TMS34010 operations. The DGIS polyline operation, for example, is turned into a series of TMS34010 line operations, and the DGIS text operation is turned into a series of mapped-to-full-depth bit-block transfer operations. Some DGIS operations are accomplished by software emulation, which uses advanced graphics algorithms and draws directly into the bit map.

Critical low-level DGIS graphics routines were written to take advantage of the TMS34010's instruction cache. Because DGIS ensures that the inner loops of the graphics operations are

Nigel Smith, PhD, Vice-President of Entry Systems, Graphic Software Systems

executed in terms of line-drawing commands that don't take advantage of the full hardware potential.

Reports are that developments in graphics hardware and software interfaces are prompting rethinking on how some of Autocad's drivers work, but the above example points out that even more performance can be achieved if applications take advantage of features on the graphics subsystems—a process that can be helped only by an awareness of standards. With even a large number of proprietary boards supporting popular packages, it's likely that the lowest-common-denominator approach would be widespread. The effort to optimize drivers for so many controllers is simply too much of a burden, while optimizing drivers for two or three standard interfaces, which themselves are fine-tuned to the hardware on which they reside, would be well worth the effort.

Bringing systems together

Unix-based workstations and minisupercomputers will increasingly interact with machines that run not only both DOS and various flavors of Unix, but the OS/2 operating system as well. Some of these machines—especially the 80386-based systems—will be able to boot and run all three operating systems, as well as communicate over networks with applications that are running under all three.

While common ground for user interfaces among Unix systems appears to be the X Window System, which was developed by the Massachusetts Institute of Technology (Cambridge, MA), Microsoft has developed Presentation Manager for the multitasking OS/2. Presentation Manager, which has a look and feel much like that of Microsoft Windows, uses a device-level protocol called the graphical program interface (GPI) to talk to the underlying hardware. GPI will therefore undoubtedly be the standard by which

fully contained in the cache, significant performance advantages are realized.

Matching graphics functions

Regardless of whether a board uses the TMS34010 or another DGIS-supported coprocessor, DGIS firmware provides compatibility to the same application software base. A key design goal of the DGIS interface is to accelerate a broad range of graphics environments. DGIS provides a comprehensive set of board-level graphics functions that closely match the graphics functions in the X Window System, Microsoft Windows, Gem, Computer Graphics Interface (CGI), Graphical Kernel System (GKS) and environments such as those used in AutoCAD, Harvard Graphics, Wordperfect 5.0, VersaCAD, Personal Designer and other popular applications.

Specific graphics functions were incorporated into DGIS to achieve the best possible performance for Microsoft Windows. DGIS functions such as Scan Left/Right, Flood Area, Output Pixel, and Draw Extended Text are critical to the performance of Windows running on a DGIS graphics device.

The X Window System includes a standard library of windowing functions and utilities used by a variety of window managers. DGIS functions such as Output Text, Copy Bitmap, Hardware Cursor Tracking and Output Polyline are key to providing workstation-like performance for X Window-based software on Unix- or DOS-based PCs.

X Window System compatibility to DGIS display boards is achieved through either of two X Window-based software packages from Graphics Software Systems (GSS): GSS*X/386 or PC-Xview. GSS has implemented the X Window System Version 11 on Unix-based 80386 PCs. The GSS*X/386 implementation includes a DGIS display server. Together with a DGIS Unix kernel driver, GSS*X/386 and a TMS34010-based DGIS board form a high-performance workstation environment for Unix-based PCs. PC-Xview turns a DOS-based PC into a low-cost X Window terminal.

DGIS also supports graphics toolkits such as the CGI-based Graphics Development Toolkit and the GSS*GKS Kernel System. The Graphics Development Toolkit includes a library of high-level bitmap and vector graphics functions, and it supports hundreds of different graphics devices and compilers such as C, Fortran, Pascal, Basic Compiler and Macro Assembler.

The GSS*GKS Kernel System is for programmers who want to write applications that require segmentation, window/viewport capabilities, and floating-point operations, and that conform to the International Standards Organization and American National Standards Institute GKS standard.

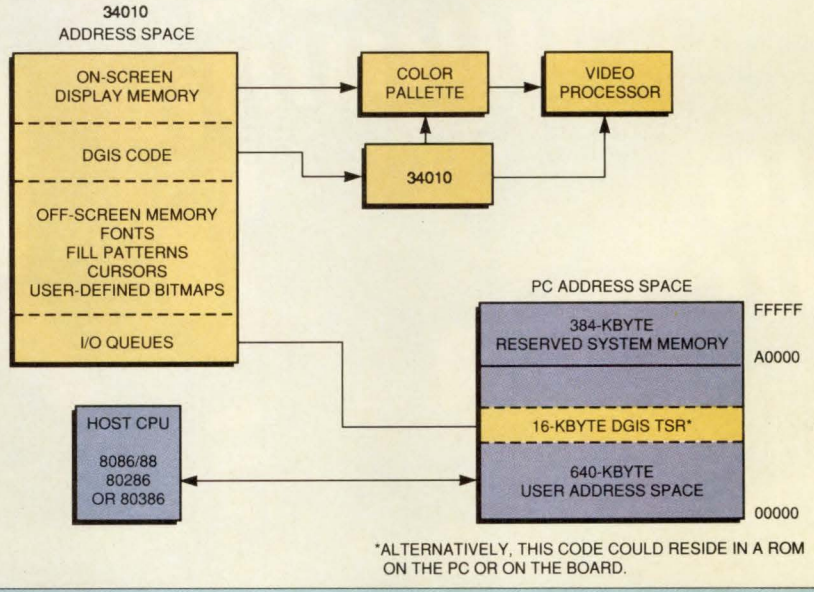

board vendors design products to run Presentation Manager at resolutions beyond VGA, since applications are supposed to communicate to Presentation Manager and not directly to GPI in the OS/2 scenario.

GPI is interesting in that it's organized as a dynamically linked library with a rich set of graphics primitives that aren't only extensible but also modular. A graphics display or peripheral controller thus could implement only those parts of GPI on, say, a plotter controller, that were needed. And certain portions of GPI lend themselves to implementation as hardware accelerators, which would increase their speed over software or ROM-based versions of the same protocol, according to Cochell of GSS. To date, however, GPI hasn't attracted much attention, but its potential for the future will grow with the acceptance of OS/2.

At any rate, a combination of interface protocols that use a single hardware controller appears to be a viable way of making a single board that can accept the graphic user interface conventions that are associated with different operating systems. A board that could run DOS graphics and the X Window System drivers through a DGIS interface, for example, could theoretically also include a GPI protocol in some combination of hardware and software that would take over when the system booted OS/2.

Looking in on the high end

The common ground that high-performance workstations and personal supercomputers are settling on is the X Window System environment. The recent announcement by the Open Software Foundation of its OSF/Motif User Environment Component (UEC), which is based on X Window, has accelerated the move in the Unix world toward X Window as a common platform for graphics shared among networked systems.

The UEC, called OSF/Motif, uses

Real Time Hit Man

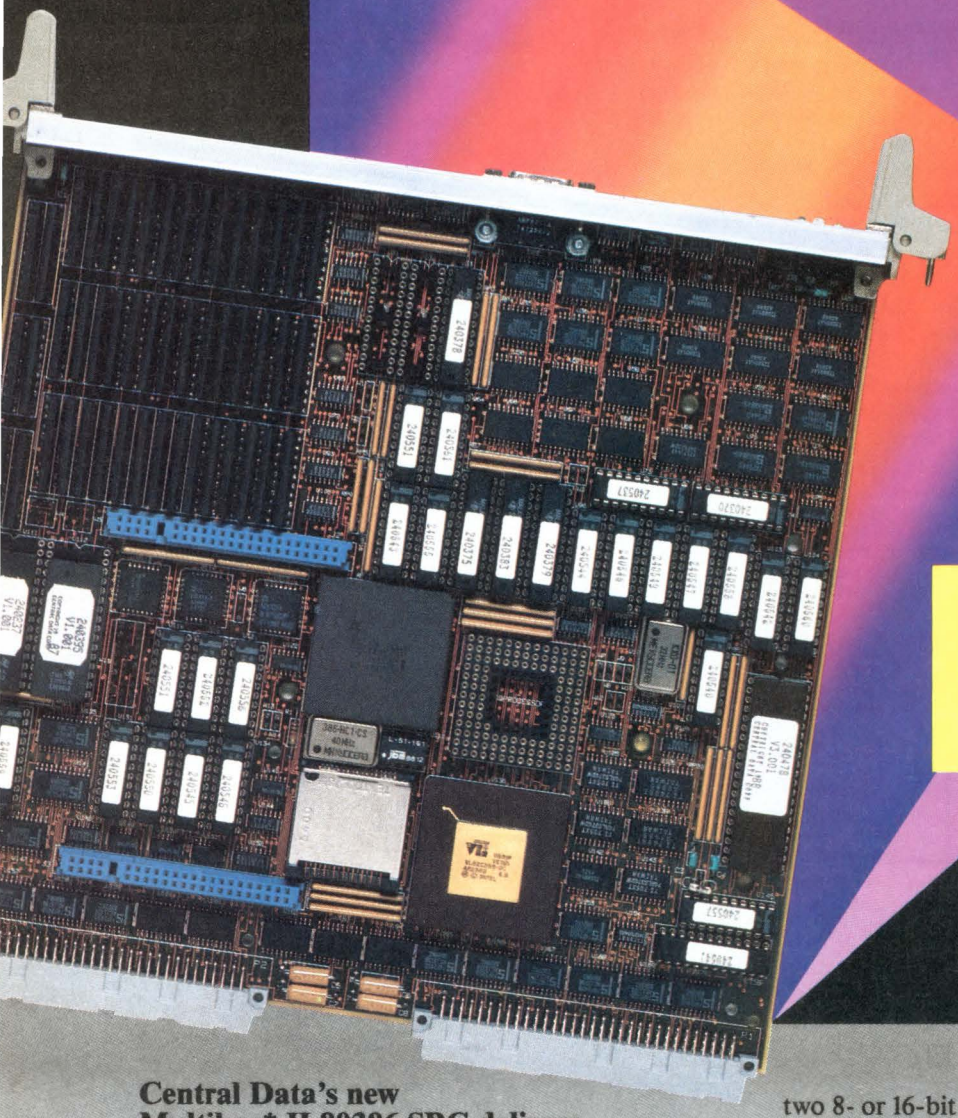

Multibus II 80386 SBC

Central Data's new Multibus* II 80386 SBC delivers deterministic performance in real time.

When you step up to 32-bit processing, you need a single board computer with memory performance to match.

That's why we built our new Intel-compatible CD22/1386 SBC with fast-page RAM—up to 16 megabytes** on board—in 4 interleaved banks. The result is a higher effective hit rate providing zero wait-state operation at 20MHz in most applications—significantly better than Intel's cache-based 386/120 in virtually every application.

Caching works great in some cases. But in real time applications, when cache hit rates can be low due to flushing and non-repetitive accesses, the memory architecture of the CD22/1386 outperforms caching—yielding consistently higher effective hit rates and reducing interrupt latencies.

And you get the full range of extra board options for which Central Data is famous. Such as two 8- or 16-bit double-wide SBX interfaces, a defined P2 interface allowing memory and/or I/O expansion for your custom applications, plus a socket that supports either an 80387 numeric data processor or a Weitek 3167 co-processor—separately or simultaneously.

Power. It makes the CD22/1386 the real time hit man your application deserves. And it's proof of Central Data's commitment to leadership in the Multibus II market. For more information and our Multibus II catalog, just call.

Central Data

1602 Newton Drive, Champaign, IL 61821-1098
Toll free 1-800-482-0315
FAX 217-359-6904
U.K. Office 0734-313877

*Multibus is a trademark of Intel Corporation.
**When 4-megabyte chips become available.

CIRCLE NO. 42

the behavior of Microsoft Presentation Manager combined with a graphic style and appearance jointly developed by Microsoft and Hewlett-Packard (Palo Alto, CA). It also combines the toolkit application program interface (API) from Digital Equipment Corp (Maynard, MA) with elements of the Hewlett-Packard API. The API is the interface through which applications communicate with the window system, not the protocol interface through which the window system talks to the hardware. That protocol, called Xlib, is part of the X Window System.

Since the X Window System is designed for networked systems, there's also a network communications protocol, called Xprotocol, that supports variable-sized data packets to transmit the Xlib graphics protocols over the network. The Xlib protocols, as distinguished from the communications protocols, are some 200 procedures that include drawing primitives as well as clipping and tiling operations and procedures to create and manipulate windows.

The Xlib protocols directly control a graphics controller called the X Server, which can be implemented as software driving a display on a workstation, as a graphics controller board plugged into a PC, or as a dedicated X Server terminal with no local computing power but with a network interface. The client/server model used in the X Window System lets client applications run on computer systems attached to the network. These systems also run the user environment, such as OSF/Motif. Only the actual display portion of the user environment needs to reside on the local node, or server, where the user is.

This partitioning lets a minimum of X Window-oriented hardware and software be lodged in the local server, making possible inexpensive graphics stations that can work with X Window applications. Network Computing Devices (Mountain View, CA), for example, has developed an inexpensive display station, the NCD16, which combines an Ethernet network interface and an X Window System server in a package with a 16-in. 1,024-×1,024-pixel monochrome display. With only the X Window server resident in the NCD16, the station can display applications appearing as any look and feel written to the Xlib protocols. Such display stations are intended for use on networks in situations where "an ASCII terminal won't do and a workstation is overkill," says Judy Estrin, vice-president at Network Computing Devices.

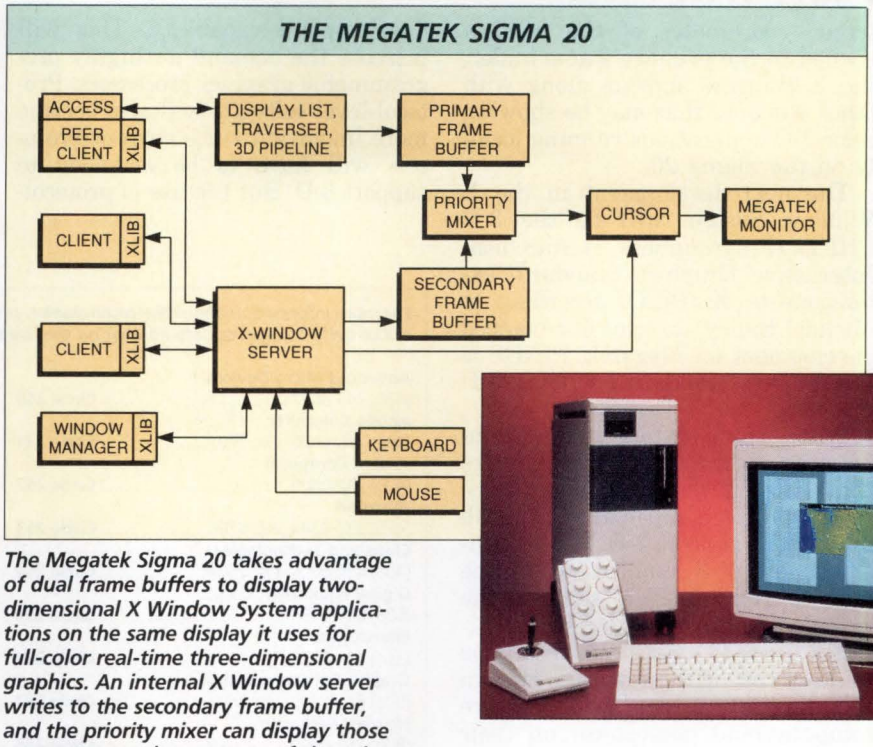

The Megatek Sigma 20 takes advantage of dual frame buffers to display two-dimensional X Window System applications on the same display it uses for full-color real-time three-dimensional graphics. An internal X Window server writes to the secondary frame buffer, and the priority mixer can display those contents over the contents of the primary frame buffer at the user's option. In such cases, the primary frame buffer supports a window in the same coordinates as the X Window application, so there's no conflict with other tasks in the primary buffer.

Graphics plug-in cards that run DOS, Windows and OS/2 applications on PC-type platforms can also be fitted with X Server firmware to run X applications when the system boots Unix—all with a change at the device interface protocol level. By upgrading the graphics subsystem and adding a network interface, a PC platform can thus be upgraded to a networked Unix workstation.

On the high end, systems designed to run real-time, 3-D shaded graphics have to break out a lower level of display sophistication that can be translated into Xlib protocols to work with X servers on a network. The ways of doing this are usually unique to the architecture and software philosophy of a specific system. Some, such as the Titan Graphics Supercomputer from Ardent, use a layered software model called the dynamic object-oriented rendering environment (DORE), which lets the user choose the level of rendering (wire frame, flat shading, Phong shading and so forth). A lower level on the rendering hierarchy can then be sent to an X Window System manager on the Titan and then to X Window servers on the network in the form of Xlib protocols carried by Xprotocol data packets.

Other systems, such as the Sigma 20 from Megatek (San Diego, CA), can include an X Window server along with their advanced 3-D pipeline. The Sigma 20 can communicate with clients on the network or other servers via its internal X Window System capability. Its local user can still use the 3-D real-time color graphics using the Sigma 20's own window-based user interface and can display remote applications in an X Window System at the same time.

This is possible because the Sigma 20 has two separate frame buffers whose contents can be mixed when the pixel data is actually being sent to the display. The priority mixer determines which pixels from which frame buffer are to be displayed. Via an access function, the 3-D portion can thus open a window in the main frame buffer to keep other processes from using that screen area. The X Window server then opens a corresponding window in the secondary frame buffer to display the X Window application. When the data is sent to the display, the pixels in the X Window area of the secondary frame buffer are given priority. Since this area corresponds to the

screen coordinates of the window opened in the primary frame buffer, the X Window appears along with other windows that may be showing color 3-D applications running locally on the Sigma 20.

The next development in the X Window System will be the 3-D PHIGS (Programmers' Hierarchical Interactive Graphics Standard) extensions to X (PEX), according to Michael Bailey, director of advanced development for Megatek. PHIGS is the de facto standard for doing 3-D modeling.

When those extensions are announced, we'll begin to see X servers that support the PEX extensions, which will be a superset of the Xlib protocols to support 3-D. While earlier versions of X applications will be able to run on PEX servers, the reverse won't be true.

For those PC platforms that hope to keep up with developments in 3-D, the only answer will be more computational horsepower on their graphics controllers—including local floating-point capability. This will increase the demand for highly programmable graphics processors. Protocol-level software will also become more important since current protocols will have to be extended to support 3-D. But the use of protocol-level interfaces will also play a vital role in preserving the existing software base as the power of silicon increases to accommodate the new demands placed on it. The common ground will constantly become higher ground. □

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

Advanced Micro Devices
(408) 749-3697 ..... **Circle 250**
Apollo Computer
(508) 256-6600, ext. 7702 ..... **Circle 251**
Ardent Computer
(408) 732-0400 ..... **Circle 252**
Autodesk
(415) 332-2344, ext. 4706 ..... **Circle 253**
Chips and Technologies
(408) 434-0600, ext. 2564 ..... **Circle 254**
Digital Equipment Corp
(508) 881-2934 ..... **Circle 255**
Enertronics
(314) 421-2771 ..... **Circle 256**
Graphic Software Systems
(503) 641-2200, ext. 169 ..... **Circle 257**
Hewlett Packard
(800) 752-0900 ..... **Circle 258**
Intel
(408) 765-8080 ..... **Circle 259**
Lundy Electronics
(516) 671-9000, ext. 374 ..... **Circle 260**
Megatek
(818) 357-9921, ext. 2165 ..... **Circle 261**
Metheus
(503) 690-1542 ..... **Circle 262**
Microsoft
(800) 426-9400 ..... **Circle 263**
National Semiconductor
(408) 721-2900 ..... **Circle 264**
Network Computing Devices
(415) 694-0650 ..... **Circle 265**
Texas Instruments
(713) 274-3297 ..... **Circle 266**

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

It Is Technology.

VenturCom's real time UNIX product family has the only real time kernel for standard hardware platforms which is AT&T's UNIX System V. Not a simple UNIX clone. Not just UNIX on top of a real time executive.

Real UNIX provides designers with SVID and POSIX standards, RFS and NFS, X-windows, Streams, complete development tools, multiple DOS under UNIX tasks. And early access to future technologies.

Real time adds performance and functionality to UNIX with preemptive and biased scheduling, contiguous file system, average interrupt latency of 50 μs, bounded context switches, memory locking, asynchronous I/O, and much more.

Ask us about VENIX™ for 80286/80386 PCs and Single Board Computers; RTX™ for other UNIX versions, such as Interactive System's 386/ix,™ and small, diskless, and ROMable UNIX kernels. Find out why Foxboro, GE, and many others are using VenturCom's real time UNIX products.

VenturCom
Real Time/Real Unix
215 First Street, Cambridge, MA 02142 (617) 661-1230
Nippon VenturCom, Inc., Tokyo 102, Japan 03-234-9381

UNIX is a registered trademark of AT&T
386/ix is a trademark of Interactive Systems Corp.
VENIX, RTX are trademarks of VenturCom, Inc.

CIRCLE NO. 43

Clip 'em and ship 'em

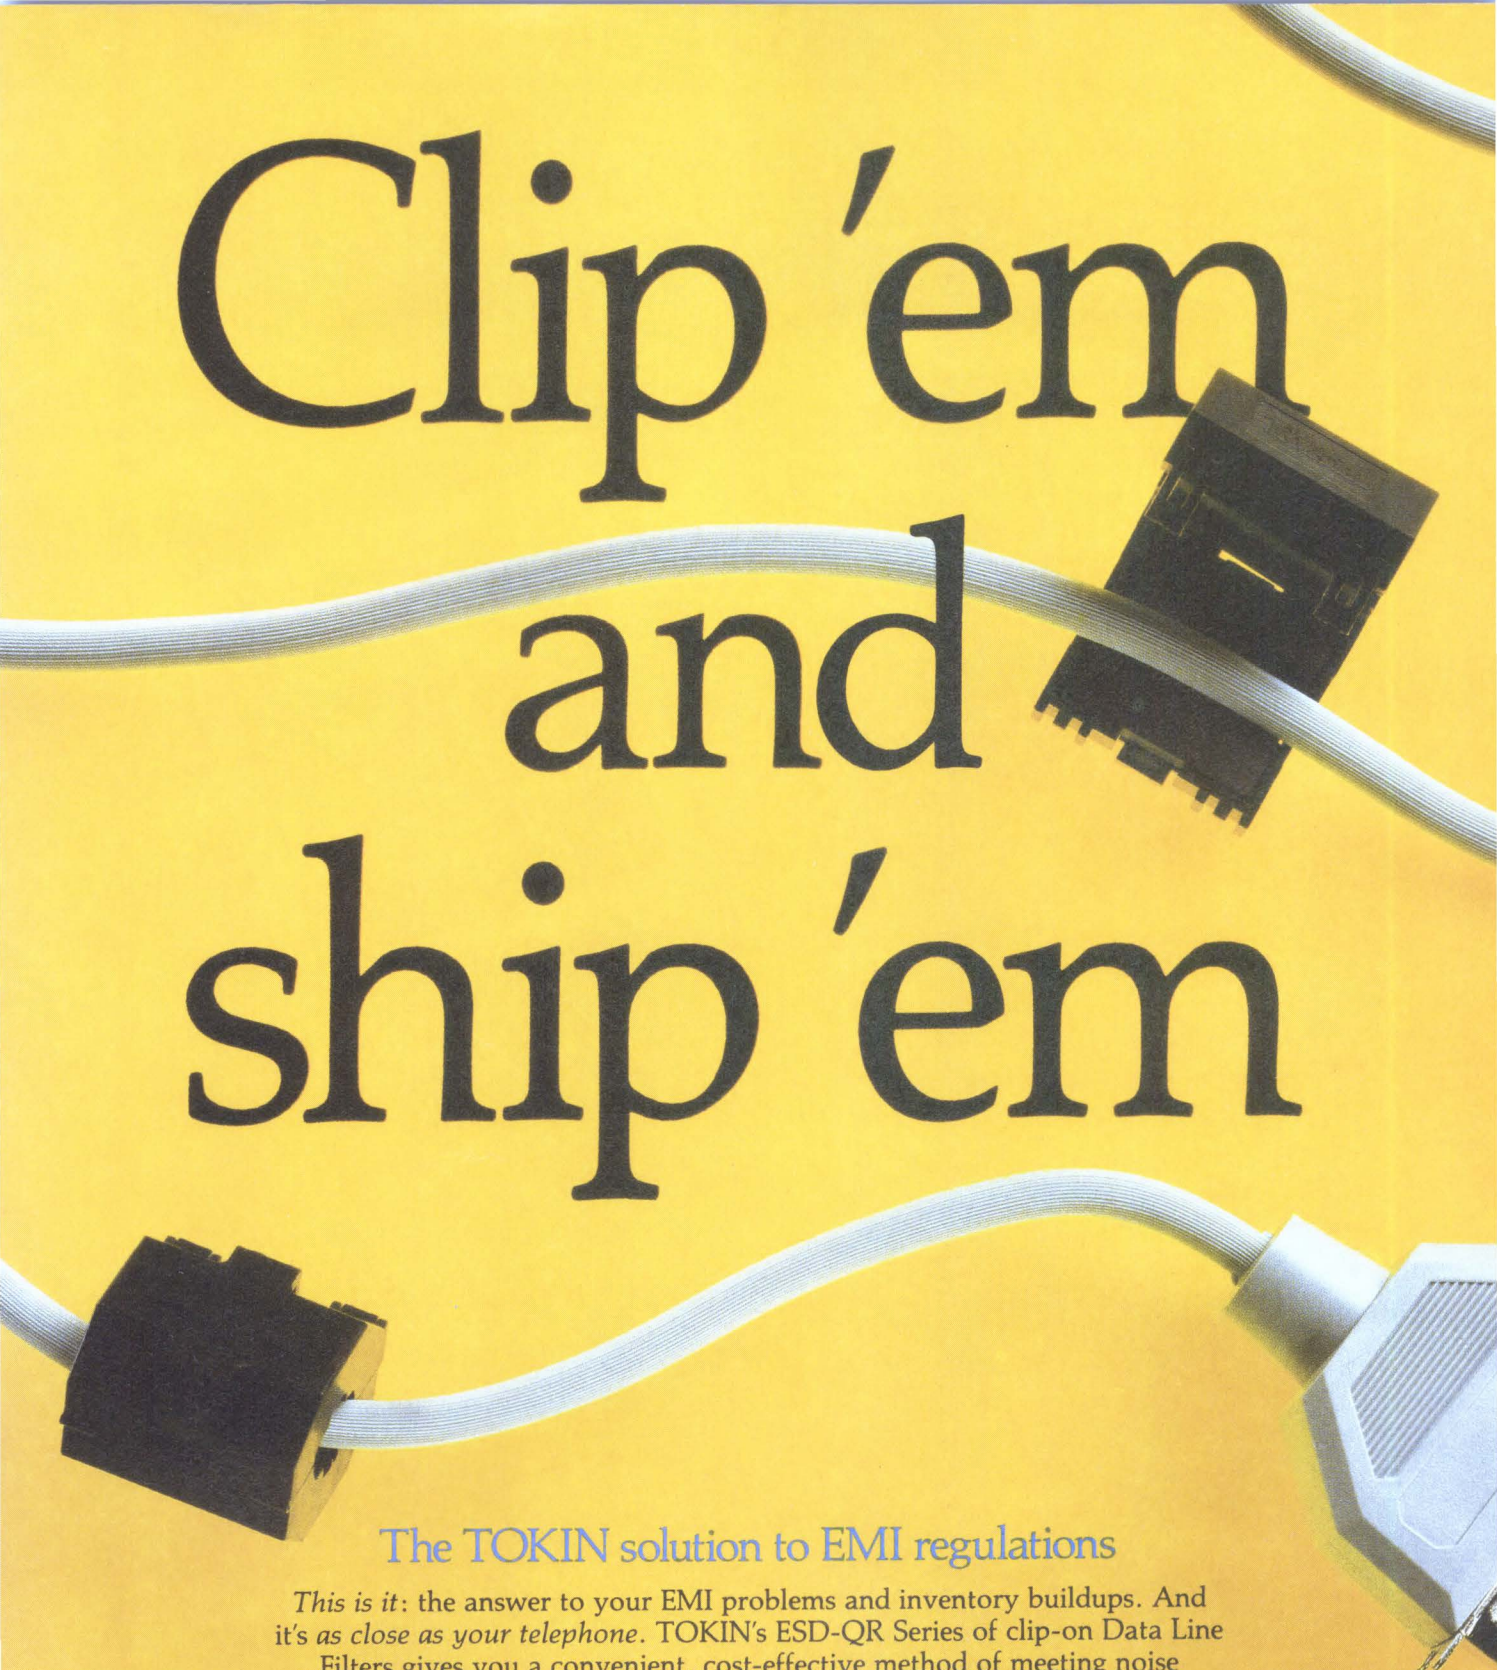

The TOKIN solution to EMI regulations

This is it: the answer to your EMI problems and inventory buildups. And it's *as close as your telephone*. TOKIN's ESD-QR Series of clip-on Data Line Filters gives you a convenient, cost-effective method of meeting noise emission and susceptibility regulations, for a wide range of electronic products. So give us a call. Overnight gets you the finest TOKIN filters. Clip them on, and watch your shipments move out.

TOKIN

Tokin America Inc.

155 Nicholson Lane, San Jose, California 95134, U.S.A. Phone: 408-432-8020 Fax: 408-434-0375
Chicago Branch 9935 Capitol Drive, Wheeling, Illinois 60090, U.S.A. Phone: 312-215-8802 Fax: 312-215-8804

Please see us at Electro'89 Electronic Show and Convention April 11—13 New York, NY #2071

CIRCLE NO. 44

Our 50,000 gate standard cell can reduce your production

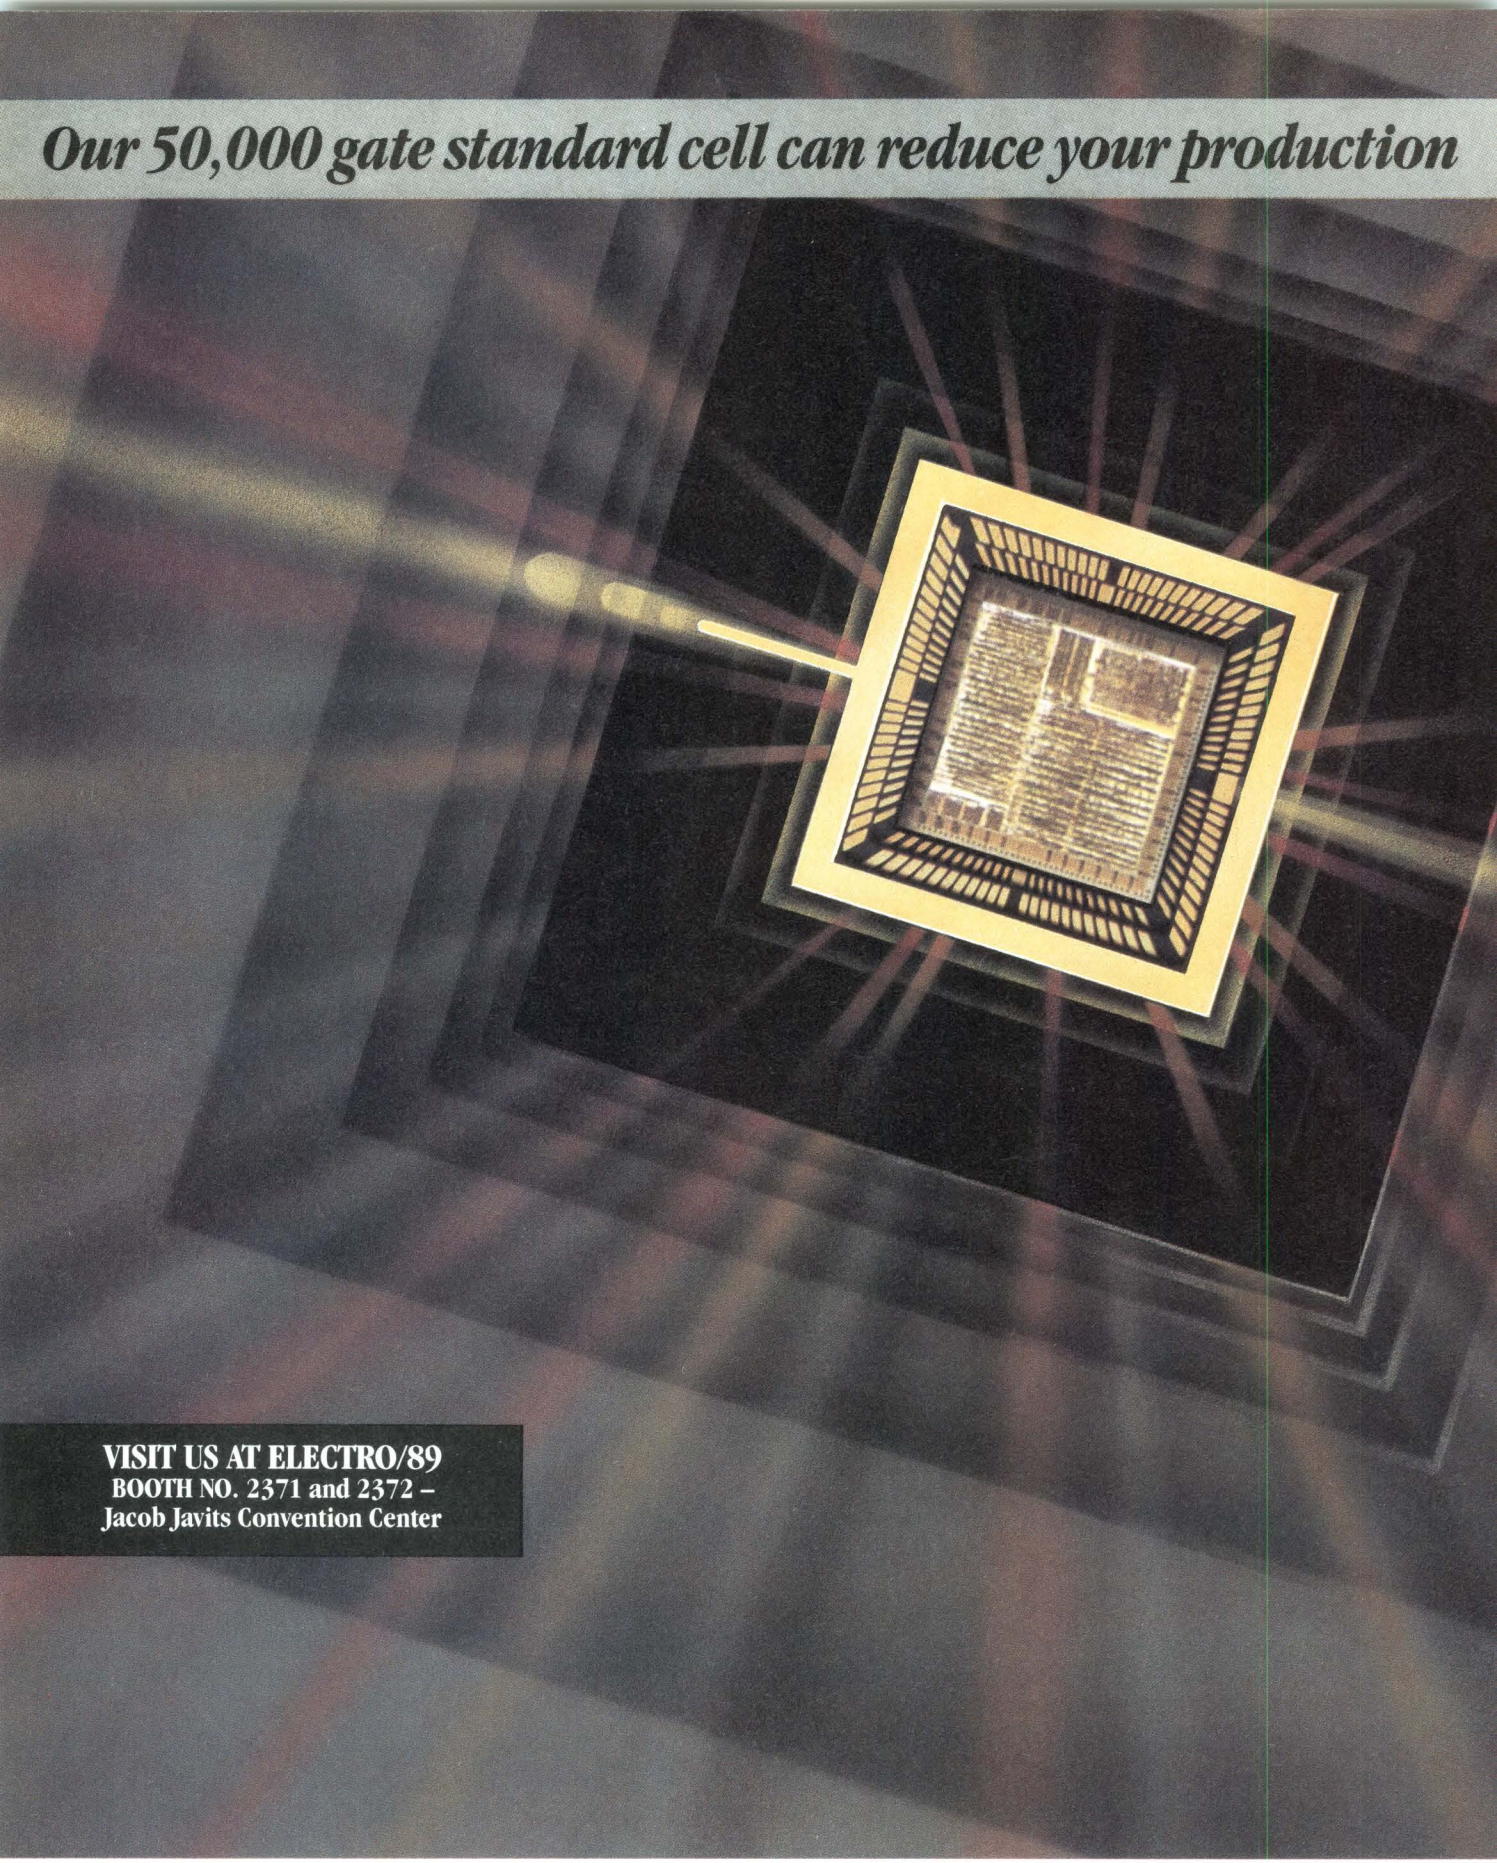

VISIT US AT ELECTRO/89
BOOTH NO. 2371 and 2372 –
Jacob Javits Convention Center

AREA SALES OFFICES: CENTRAL AREA, Toshiba America Electronic Components, Inc., (312) 945-1500; **EASTERN AREA,** Toshiba America Electronic Components, Inc., (617) 272-4352; **NORTHWESTERN AREA,** Toshiba America Electronic Components, Inc., (408) 737-9844; **SOUTHWESTERN REGION,** Toshiba America Electronic Components, Inc. (714) 259-0368; **SOUTH CENTRAL REGION,** Toshiba America Electronic Components, Inc., (214) 480-0470; **SOUTHEASTERN REGION,** Toshiba America Electronic Components, Inc., (404) 368-0203; **MAJOR ACCOUNT OFFICE, FISHKILL, NEW YORK,** Toshiba America Electronic Components, Inc., (914) 896-6500; **MAJOR ACCOUNT OFFICE, BOCA RATON, FLORIDA,** Toshiba America Electronic Components, Inc., (305) 394-3004. **REPRESENTATIVE OFFICES: ALABAMA,** Montgomery Marketing, Inc., (205) 830-0498; **ARIZONA,** Summit Sales, (602) 998-4850; **ARKANSAS,** MIL-Reps, (512) 346-6331; **CALIFORNIA** (Northern) Elrepco, Inc., (415) 962-0660; **CALIFORNIA** (L.A. & Orange County) Bager Electronics, Inc., (818) 712-0011, (714) 957-3367, (San Diego County) Bager Electronics, Inc., (619) 944-4188; **COLORADO,** Straube Associates Mountain States, Inc., (303) 426-0890; **CONNECTICUT,** Datcom, Inc., (203) 288-7005; **DISTRICT OF COLUMBIA,** D.G.R., Inc., (301) 581-1360; **FLORIDA,** Sales Engineering Concepts, (813) 823-6221, (305) 426-4601, (305) 682-4800; **GEORGIA,** Montgomery Marketing, Inc., (404) 447-6124; **IDAHO,** Components West, (509) 922-2412; **ILLINOIS,** Carlson Electronic Sales, (312) 956-8240, R.W. Kunz, (314) 966-4977; **IOWA,** Carlson Electronics, (319) 377-6341; **KANSAS,** D.L.E. Electronics, (316) 744-1229; **LOUSIANA,** MIL-Reps, (713) 444-2557; **MAINE,** Datcom, Inc., (617) 891-4600; **MARYLAND,** D.G.R., Inc., (310) 583-1360; **MASSACHUSETTS,** Datcom, Inc., (617) 891-4600;

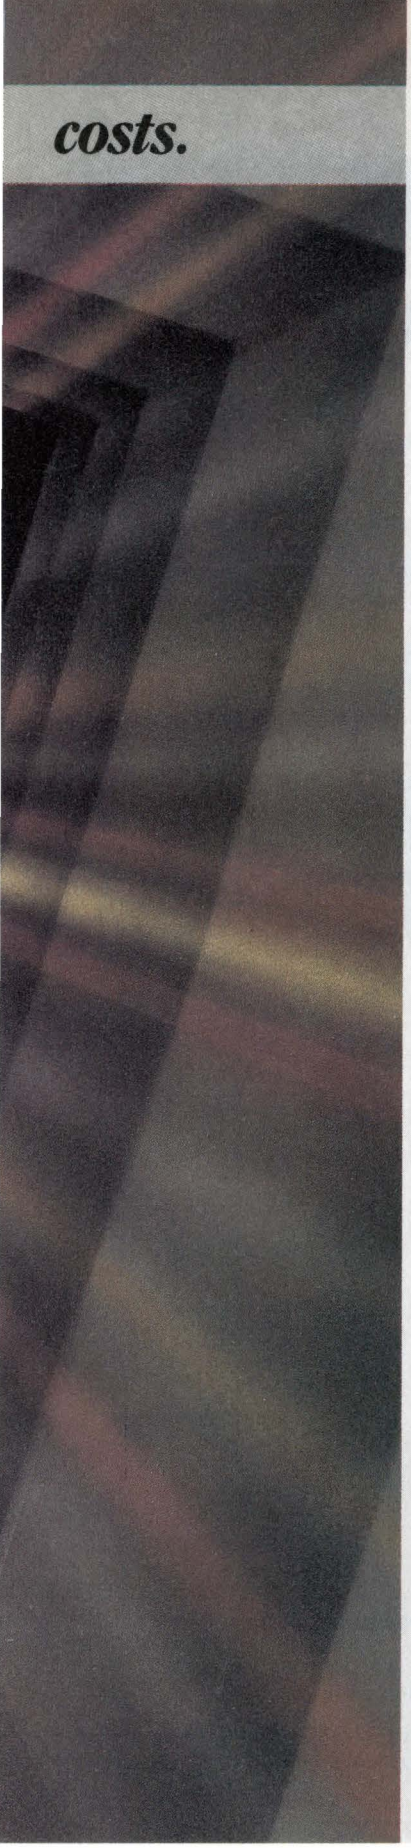

costs.

SERVICE IS OUR KEY COMPONENT

Our TC23SC Standard Cell is the most cost-effective solution for high-volume, memory-intensive ASIC applications. Because it uses less silicon. It's as simple as that. Less than an equivalent gate array. And our proven 1.5 micron process provides reliable high density.

Of course, the higher your production volume, the greater your savings. But the TC23SC can even cut costs on low volume applications, because it can simplify your ASIC design task immensely. You don't even have to need all 50,000 gates. the TC23SC Series includes a full range of sizes from 700 to 50,000 equivalent gates. It's simple to choose the optimum standard cell for your design.

Then there's the extensive library lineup. The TC23SC Series library, ADVANCELL™, is an ultra-high function LSI library. The full product lineup includes 166 primitive cell types, 108 74HC Series macrofunction types, and 67 I/O cell functional types. Macrocells include 2900 Series, CPU peripherals, multiplexers and more. Several RAM and ROM megacells are also available.

The TC23SC Series is fully supported by the VL-CAD system. The system is a CAD design tool developed by Toshiba with full support from logic design to autoplace and route.

BASIC CHARACTERISTICS		
PRODUCT LINEUP	TC22SC Macrocells & Macrofunctions	TC23SC
LIBRARY	74HC Series Compatible Macrofunctions RAM ROM Functional Macros (PLA**, ALU, MPY)	ADVANCELL™ * RAM ROM Functional Macros (PLA**, ALU**, MPY, Barrel shifter, adder, FIFO**) 2900 Series Macros*** CPU Peripherals** Analog Macros***
PROCESS TECHNOLOGY	2μm	1.5μm
	HC²MOS Si-gate double layer metal	
MAXIMUM TOGGLE FREQUENCY	100MHz	150MHz
GATE SPEED (INNER GATE)	1.5ns	1.0ns
SUPPLY VOLTAGE	5V	
GATE COMPLEXITY (MAX.)	10K gates	50K gates
AVAILABILITY	NOW	

*ADVANCELL is a trademark owned by Toshiba or licensed from Siemens or General Electric Co., U.S.A. in certain countries.
**Under development.
***Development complete, to be available shortly.

The 1.5 micron technology used in the TC23SC system ensures high integration and performance and it provides speeds of 1.0 ns typical gate delay.

To help you even further, we now have a total of five ASIC design centers around the United States for your convenience.

For complete information, contact your nearest Toshiba Regional Sales Office: Northwestern: (408) 737-9844, Southwestern: (714) 259-0368, Central: (312) 945-1500, South Central: (214) 480-0470, Eastern: (617) 272-4352, Southeastern: (404) 368-0203. Or call 1-800-888-0848 ext. 517 today.

In Touch with Tomorrow

TOSHIBA

TOSHIBA AMERICA ELECTRONIC COMPONENTS, INC.

CIRCLE NO. 45

©1989 Toshiba America, Inc.

MAS-88-005-1

MICHIGAN, Action Components Sales, (313) 349-3940; **MINNESOTA**, Electric Component Sales, (612) 933-2594; **MISSISSIPPI**, Montgomery Marketing, Inc., (205) 830-0498; **MISSOURI**, D.L.E. Electronics, (316) 744-1229, R.W. Kunz, (314) 966-4977; **MONTANA**, Components West, (206) 885-5880; **NEBRASKA**, D.L.E. Electronics, (316) 744-1229; **NEVADA**, Elrepco, Inc., (415) 962-0660; **NEW ENGLAND**, Datcom, Inc., (617) 891-4600; **NEW HAMPSHIRE**, Datcom, Inc., (617) 891-4600; **NEW JERSEY**, Nexus Technology, (201) 947-0151; **NEW MEXICO**, Summit Sales, (602) 998-4850; **NEW YORK**, Nexus Technology, (201) 947-0151; **NEW YORK** (Upstate), L.D. Allen, (315) 437-8387; **NORTH/SOUTH CAROLINA**, Montgomery Marketing, Inc., (919) 467-6319; **NORTH/SOUTH DAKOTA**, Electric Component Sales, (612) 933-2594, J.R. Thornberry, (216) 248-4995; **OKLAHOMA**, MIL-Reps, (214) 644-6731; **OREGON**, Components West, (503) 684-1671; **PENNSYLVANIA**, Nexus Technology, (215) 675-9600; **PENNSYLVANIA** (Western), J.R. Thornberry, (216) 248-4995; **RHODE ISLAND**, Datcom, Inc., (617) 891-4600; **TENNESSEE**, Montgomery Marketing, Inc., (205) 830-0498; **TEXAS**, MIL-Reps, (512) 346-6331, (713) 444-2557, (214) 644-6731; **UTAH**, Straube Associates Mountain States, Inc., (801) 263-2640; **VERMONT**, Datcom, Inc., (617) 891-4600; **VIRGINIA**, D.G.R., Inc., (301) 581-1360; **WASHINGTON**, Components West, (206) 885-5880, (509) 922-2412; **WEST VIRGINIA**, D.G.R., Inc., (301) 581-1360; **WISCONSIN**, Carlson Electronics, (414) 476-2790, Electric Component Sales, (612) 933-2594; **WYOMING**, Straube Associates Mountain States, Inc., (303) 426-0890; **CANADA, BRITISH COLUMBIA**, Components West, (206) 885-5880; **ONTARIO**, Electro Source, Inc., (416) 675-4490, (613) 592-3214; **QUEBEC**, Electro Source, Inc., (514) 630-7486.

INTEGRATED CIRCUITS

ICs bring monochrome and color image capture into single-chip solution

Two image digitizer ICs introduced by Brooktree are aimed at highly integrated, low-cost image capture systems for use with PCs and workstations. The announcement includes a single-channel monochrome digitizer and a three-channel part for capturing RGB color images. The parts also signal Brooktree's move to a line of integrated imaging components that will serve as building blocks for addressing problems in the capture, digitization and manipulation of both live and computer-generated images.

The Bt251 single-channel digitizer captures gray-scale video images at 15 million samples/s. The 8-bit analog-to-digital converter produces samples that can differentiate 256 levels of gray scale, with sampling rates that make it suitable for digitizing National Television Service Committee video signals.

In addition, the Bt251 contains a 256-byte programmable look-up table RAM that can be used to perform on-the-fly image-processing algorithms. A sampled pixel can be used to look up the address of a value that's then stored as the actual pixel in the captured image. This capability is useful in such applications as contrast enhancement, noise filtering and thresholding. In thresholding, for example, a samples pixel would have to equal a given value before the pixel could be stored in the frame buffer.

The Bt253 essentially contains the digitizing circuitry of three Bt251s, less the look-up table RAMs. The Bt253 digitizes the three channels of color analog signals via its three video input ports and can generate 8, 15 or 24 bits of color data per pixel.

Because RGB data doesn't lend itself to practical color image processing, the Bt253 is intended primarily as a means of bringing color video images onto a computer screen, according to Keith Jack, imaging product manager. Such an application—in a window, for instance—would be viewed by the user in conjunction with other operations. The Bt253 can, however, be used to digitize monochrome images by using only the green channel. Such captured monochrome images would then be available for normal image-processing operations.

Both the Bt251 and the Bt253 have on-chip sync detection and programmable gain and offset, eliminating the need to build these necessary features with discrete components. They also have a direct microprocessor interface for system access to programmable registers and on-chip RAM. Brooktree plans to introduce enhanced performance versions of both chips, which will let them capture European PAL and SECAM video signals.

Both the Bt251 and the Bt253 will be available in sample quantities this month with volume shipments beginning in late summer. The Bt-251 comes in a 44-pin plastic leaded chip carrier (PLCC) and is priced at $36 in quantities of 100. The Bt253 comes in an 84-pin PLCC and is priced at $48 in quantities of 100.

—*Tom Williams*

Brooktree
9960 Barnes Canyon Rd
San Diego, CA 92121

Circle number 100

THE BT251 SINGLE-CHANNEL DIGITIZER

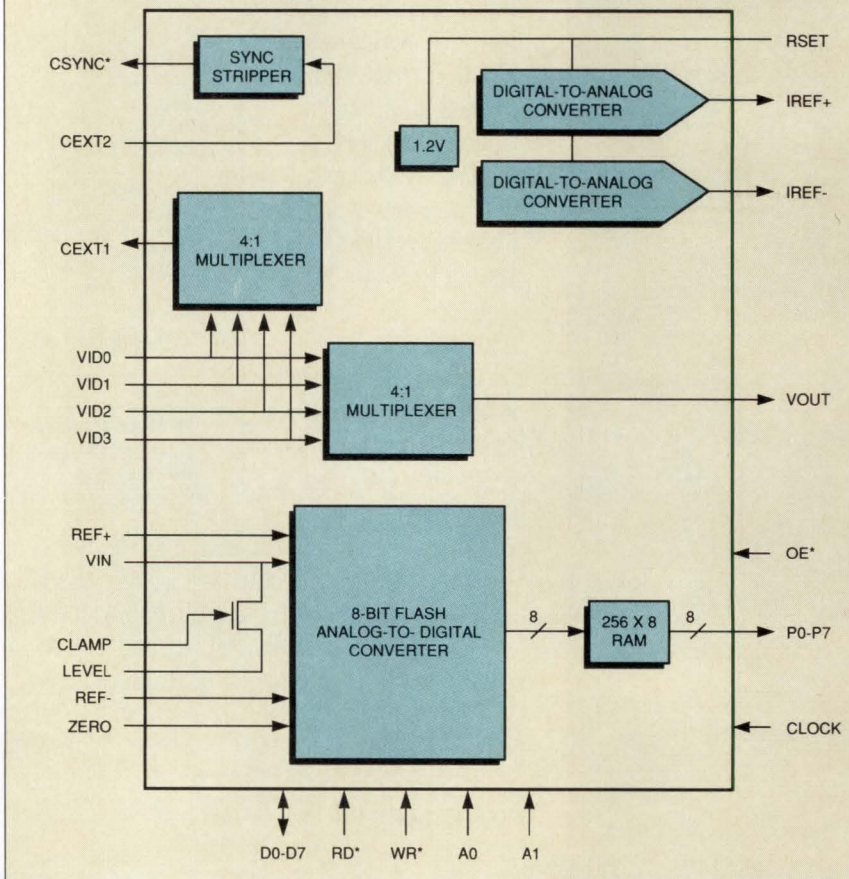

Trademark Information

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.

SMART-POWER is a registered trademark of Nartron Corp.

If you're in the market for a UNIX® workstation, we're the company you should be talking to.

$5490

DN3000 with 15" 1024 x 800 monochrome display

$6490

DN3000 with 19" 1280 x 1024 monochrome display

$9400

DN3000 with 15" 1024 x 800 color display

$8490

DN3500 MC68030-based workstation with 15" 1024 x 800 monochrome display

$9490

DN3500 MC68030-based workstation with 19" 1280 x 1024 monochrome display

$9990

DN3000 with 15" 1024 x 800 monochrome display and a 155MB ESDI disk

We can talk price.

We'd like to begin this conversation by making six points. Price points, that is.

Because within our family of UNIX® workstations are six machines whose prices are as low as their technology is high.

Take for example, our family of personal workstations. Machines that offer superior networking ability and even MC68030 performance for less than $10,000.

Prices that become even more impressive when you realize they include 4MB of memory, a floating point coprocessor, an IBM PC AT®-compatible bus, and a choice of Ethernet,® Apollo® token ring, or IBM token ring networks.

So when you're ready to talk price, ask us about our Series 3000™ and Series 3500.™ Or better yet, just let the numbers speak for themselves.

apollo

CIRCLE NO. 47

We can talk performance.

Price isn't the only attribute of our workstations that speaks for itself.

Recent tests pitting the Apollo Series 10000™ against workstations, minis, superminis and mainframes declared the 10000 an undeniable victor. Expertly running complex applications up to seven times faster than its challengers.

Such an outcome could hardly be achieved by conventional means. Instead, the Series 10000 required an entirely new architecture, making it the only office-sized UNIX® server and workstation to yield true supercomputer performance.

With up to four CPUs, the 10000's *PRISM*™ architecture delivers performance 100 times that of a VAX® 11/780. Yet costs as little as $6500/MFLOP, one-tenth that of the competition.

All of which suggests that if extraordinary performance is what you're after, the Series 10000 from Apollo will leave you speechless.

apollo

CIRCLE NO. 48

And we can talk to any computer you'd ever want to talk to.

The only story bigger than our performance capabilities is our open computing environment. As a founding member of the Open Software Foundation, Apollo is committed to providing a broad range of computing alternatives.

Today every Apollo workstation can communicate over OSI, DECnet,® Ethernet or IBM token ring networks and accept VME-, SCSI- or PC AT-compatible equipment. Our support for UNIX® (Berkeley 4.3 and System V.3 SVID compliant), MS-DOS™ and X Window System™ assures that your application software is truly portable. And our adherence to standards such as PHIGS and GKS streamlines graphics-intensive applications.

What's more, Apollo's de facto industry standard for distributed computing, the Network Computing System,™ allows users to share data and computing power in a single integrated environment. Resulting in a new way of computing that's both more elegant and more efficient than any other.

Give us a call soon. You'll find that our actions speak even louder than words.

apollo

For more information, call 1-800-323-1846 (Massachusetts, 1-800-847-1011) or write Apollo Computer, 330 Billerica Road, Chelmsford, MA 01824.

Apollo is a registered trademark and Personal Workstation, Series 3000, Series 3500, Series 10000, *PRISM*, and Network Computing System are trademarks of Apollo Computer Inc. UNIX is a registered trademark of AT&T. AT is a registered trademark of International Business Machines Corp. Ethernet is a registered trademark of Xerox Corporation. VAX and DECnet are registered trademarks of Digital Equipment Corporation. MS-DOS is a trademark of Microsoft Corporation. X Window System is a trademark of MIT.

CIRCLE NO. 49

INTEGRATED CIRCUITS

RISC processor posts burst speeds of 50 million operations/s

The SC32 Language Chip is a 32-bit CMOS high-level language processor implemented in 34,000 transistors and capable of 10 Mips with burst speeds of up to 50 million operations/s. The 10-MHz version of the device runs the popular Sieve and Fibonacci benchmarks in 2.3 and 7.0 s, respectively, as compared to an Intel 20-MHz 80386 processor with an 80387 math coprocessor that posts times of 23.18 and 42.48 s.

A unique hardware stack-caching mechanism coupled with stack registers lets the SC32 combine multiple instructions for optimized execution. Multiple operations such as arithmetic logic functions, stack and subroutine return operations are optimized into single machine instructions that can be executed in one 100-ns clock cycle. The chip directly accesses 16 Gbytes of data memory and 2 Gbytes of code memory.

The chip was developed by the Applied Physics Laboratory at John Hopkins University as a multitasking processor for controlling satellites. Operating under an exclusive license agreement, Silicon Composers will manufacture, distribute and sublicense the SC32 to OEMs, semiconductor and computer companies for use in embedded-control and general computing applications. The chip goes a step further than most RISC processors by supporting high-level language Forth directly in silicon instead of through a small or reduced instruction set.

Architecture speeds execution

The SC32 fetches an instruction while executing another by pipelining the two operations, resulting in an effective instruction execution time of one clock cycle. A data fetch or store anywhere within memory requires one additional clock cycle to execute.

Internal architecture of the device consists of 39 internal registers, three 32-bit buses and an ALU with support registers. There are 16 registers apiece for the parameter stack and the return stack, of which the top four registers are accessible from the instruction set. The parameter stack holds data during subroutine calls, while the return stack holds the subroutine return addresses. Two of the four user registers are used as pointers to main memory for the stack caching, whereas the two remaining registers are available for general use.

No stack overhead

Stack caching is performed in hardware with no software intervention required. Most applications remain within the overflow/underflow limits for long periods of time, so there's no stack overhead. When overflow/underflow occurs, the pointers are adjusted and the stack data is transferred to or from memory as required. Each stack cache overflow/underflow requires two clock cycles and is transparent to the software. Since most local stack action happens within a narrow depth range, overhead is typically less than 1 percent for both stack caches.

There are several advantages to this stack-caching mechanism. The first four stack values of both stacks are directly accessible by the SC32 instructions for operations such as double-precision (64-bit) math. Stack depths are limited only by the amount of memory installed in the system, making deep recursion possible with little subroutine overhead. As a result, most programs stay within a very narrow range of stack depth for a long time. The circular stack cache with its ten-element range indicates that very little stack caching actually occurs. For multitasking systems, a task context save takes from 40 to 80 clock cycles; a task context load takes only 35 clock cycles.

An SC32 development board, the SC/FOX32 PCS Parallel Coprocessor System, plugs into an IBM PC, PC XT, PC AT or 80386-class computer or compatible. The coprocessor system prices start at $1,995; the SC32-10 is priced at $295 each in lots of 100. —*Michael Donlin*

Silicon Composers
210 California Ave
Palo Alto, CA 94306

Circle number 112

Free: DB86 Debugger *and* ASM86 Assembler

Just buy an ICE™ or I²ICE™ emulator for the 8086/186/188 now. **(A $750 Savings.)**
Buy *two* of the above emulators and also receive two iC-86 C Compilers Free. **(An additional savings of $1,500.)**
To order, or for more information, call:
800-874-6835
Offer expires June 30, 1989.

© 1989 Intel Corp
ICE™ and I²ICE are trademarks of Intel Corp.

PLUG AND PLAY.

Intel 8086/186/188 languages, debuggers and emulators all work together. Right out of the box. And they work *better*. Our all new windowed source debugger, for example, is a snap to use. Our compilers generate efficient ROMable code. And, our I²ICE™ and ICE™ Emulators provide precise emulation to *16 MHz*. Next time, just Plug and Play. With Intel tools.

intel

Our Tools Work Better
Because They Work Together.

CIRCLE NO. 50

INTEGRATED CIRCUITS

32-bit graphics microprocessor makes desktop peripherals 10 times faster

The growing popularity of desktop graphics peripherals creates quite an opportunity for system vendors, but quite a headache for system architects. The latest generation of laser printers, image scanners and high-resolution facsimile machines presents a set of design problems almost guaranteed to send experienced microcomputer designers back to the Excedrin. Recently the pain has become so acute that silicon vendors are trying to figure out how to spell relief.

Some relief has arrived in the form of a new graphics processor, the 32GX32 from National Semiconductor. The chip is conceptually quite similar to National's earlier 32-CG16, which is a cost-reduced version of the 32000-family CISC architecture, but with added hardware for graphics, signal processing and pixel-moving operations. The 32-GX32, however, has a number of important architectural differences that reflect the more sophisticated target market.

"The clearest difference between the CG16 and the GX is performance," claims Giora Yaron, microprocessor group vice-president at National. "Using a range of page compositions in Postscript, the GX measures six to ten times faster than the 32CG16. That's about eight to ten times faster than the controller in the high-end Apple Laserwriter II NTX." Yaron's figures come from benchmark results that use Bauer PDL code on the GX.

The increased speed comes from a number of implementation changes in the new chip. Most obvious, the GX has a 32-bit external data bus, compared to the CG16's 16-bit bus. Like the CG16, the GX dispenses with the 32000 family's fine internal memory-management unit, so the CPU looks out on a completely flat 4-Gbyte address space.

Internally, the changes in the new chip are more dramatic. The GX execution unit is implemented as a four-stage pipeline, rather than the single stage used in the CG16. This permits operation at up to 30 MHz. This increased speed puts great pressure on memory bandwidth, which National has relieved by using on-chip caches. The GX has a 512-byte instruction cache and a 1,024-byte data cache.

The data cache is particularly im-

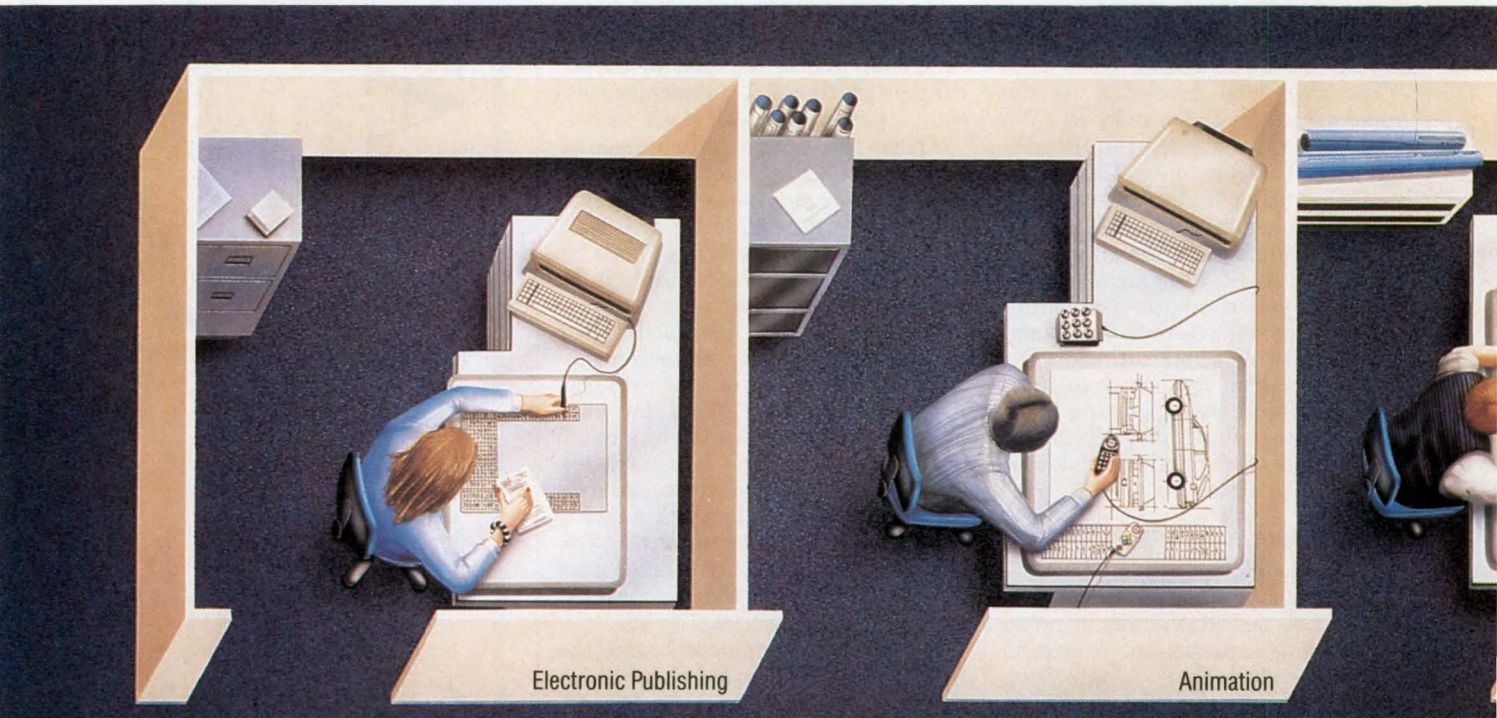

Microgrid® II. A solution

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

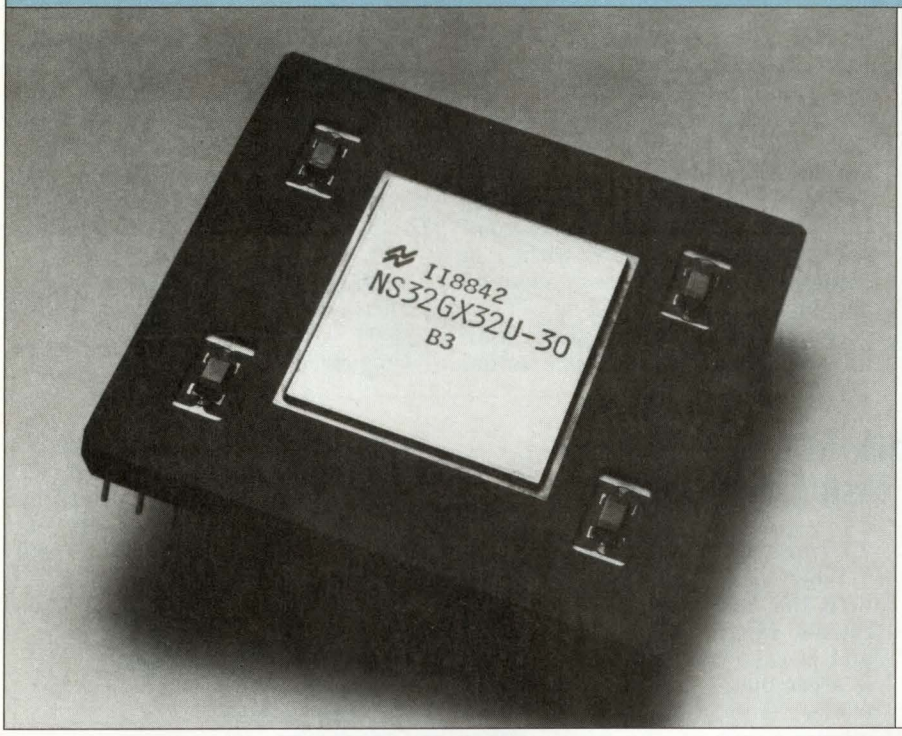

portant to character-generation operations, according to Yaron. At 1,024 bytes, it's large enough that a character can be converted from outline to bit map within the cache, greatly reducing the number of memory references in this critical operation. During data-moving operations, the cache can be disabled.

Managing cost

These enhancements greatly increase the speed of the GX, but they also increase the cost. Recognizing that even high-end, 60-page/min. printers are cost-sensitive, National has added a number of cost-control features as well. "The memory system accounts for the dominant cost in most printer designs," observes Yaron, "so we took steps to reduce both the size and the speed of the memory required for the GX."

Size reduction resulted from the

(continued on page 96)

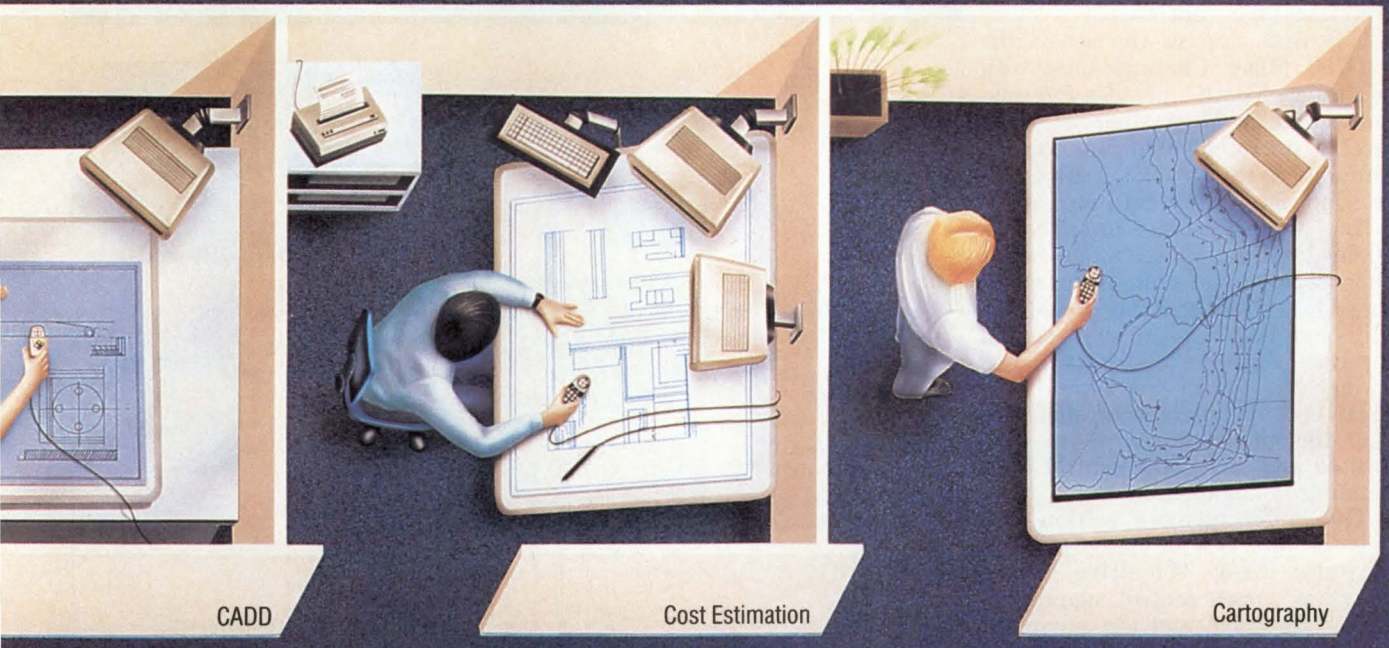

to fit any application.

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.

See us at NCGA Booth #4127

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. 51

Summagraphics® CORPORATION
60 Silvermine Road, Seymour, CT 06483
203-881-2000
©1987 Summagraphics Corporation. All rights reserved.

INTEGRATED CIRCUITS

GX's commitment to a complex instruction set. Because of the rich instruction set with graphics primitives, National claims a 30 to 40 percent advantage in the amount of code required for a printer, compared to 32-bit RISC processors.

But the company also worked to reduce the need for expensive high-speed memory chips. The size and even the replacement algorithms of the GX are tuned to printer-control applications, resulting in a very high hit rate. This reduces the GX's sensitivity to slow memory, and lets the chip run with one or even two wait states on external memory operations without significant loss of performance. "We pay a penalty in overall speed of about 8 percent per wait state," Yaron claims. "In comparison, the AMD 29000 without a cache suffers a speed loss of about 45 percent per wait state."

National has also tried to minimize the cost of the 32-bit external bus. The bus width is dynamically alterable to 8, 16 or 32 bits, so cheaper 8-bit memory devices and peripheral controllers can be used where performance isn't an issue.

The 32GX32 is available now and costs $99 for 20-MHz parts in quantities of 1,000. —*Ron Wilson*

National Semiconductor
2900 Semiconductor Dr
Santa Clara, CA 95052
Circle number 111

New two-chip set shrinks Micro Channel hard disk controllers

Chips and Technologies has made a living out of shrinking large chunks of personal computer motherboard into a small handful of ICs. Now the company has taken its pastime to a new ground: shrinking the mess of hardware lying between a hard disk drive and a PC motherboard. The result of these efforts is a new two-chip set that forms the heart of a six-chip hard disk controller.

The first chip in the set is the 82C780 Micro Channel hard disk controller. The IC includes five major sections: a Micro Channel interface, a microcontroller interface, a buffer memory interface, a drive control interface and a data formatter.

The Micro Channel interface includes Programmable Option Select support and programmable card ID. In addition, the 82C780 provides a full slave DMA controller. A second interface forms a multiplexed bus connection to an 8051-family or 68HC11-family microcontroller, letting the single-chipper move data around and access the control lines running to the disk drive.

The buffer memory section provides an interface for up to 64 kbytes of static RAM. The drive control interface passes control signals between the drive and the microcontroller via a 48-mA I/O port.

While the 82C780 handles interfacing and formatting functions, the actual serial data stream onto and off of the drive is handled by the 82C784 data separator, encoder and decoder chip. Operating from a 5-V supply, this chip combines digital and analog technology to provide modified frequency modulation or 2,7 run-length-limited capability.

When the two-chip set is combined with the necessary microcontroller, SRAM, Micro Channel bus buffers and passive components, the result is a complete embeddable disk controller. Through the extensive programmability of the parts, either ST506 or enhanced small device interface drives can be controlled.

The two-chip set is now sampling, and production quantities are expected in the third quarter of 1989. In volume, the set will cost $25. —*Ron Wilson*

Chips and Technologies
3050 Zanker Rd
San Jose, CA 95134
Circle number 110

THE 82C780 AND THE 82C784

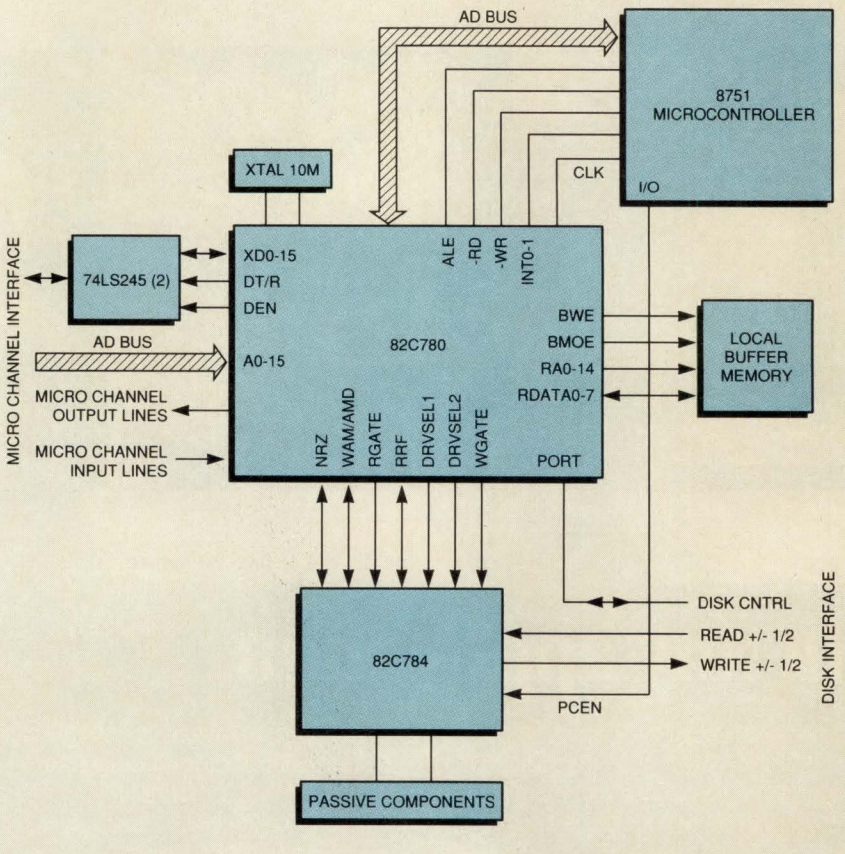

It's 2 in 1

New! Analog and digital I/O Signals on a single VME board.

Acromag gives you 32 analog inputs, 32 digital input/outputs, plus 2 analog outputs. . . all on one VME multifunction board. This also saves you card cage slot positions. It costs less, too.

Analog Features.

The analog input channels can be used differentially for 16 channels, or as single-ended inputs for 32 channels. Each individual channel has programmable gain for maximum performance.

Programmable ranges are ±5V, ±10V; 0 to ±10V, and gains X1, X2, X4, X8. This gives you input ranges from ±0.625 to ±10V. Conversion time is 32 μs.

The two analog outputs each have individually selectable voltage output ranges, ±2.5V, ±5.0V, ±10V; 0 to +5V, 0 to +10V. Settling time is 6 μs.

Digital Features.

Each of the 32 digital channels can be used as either an input or output. Each output channel sinks up to 100 mA, from up to a 30V DC source, and TTL input loads.

More Details?

For full details about our 2 in 1 VME boards, ask for your copy of technical bulletin 20-004.4.

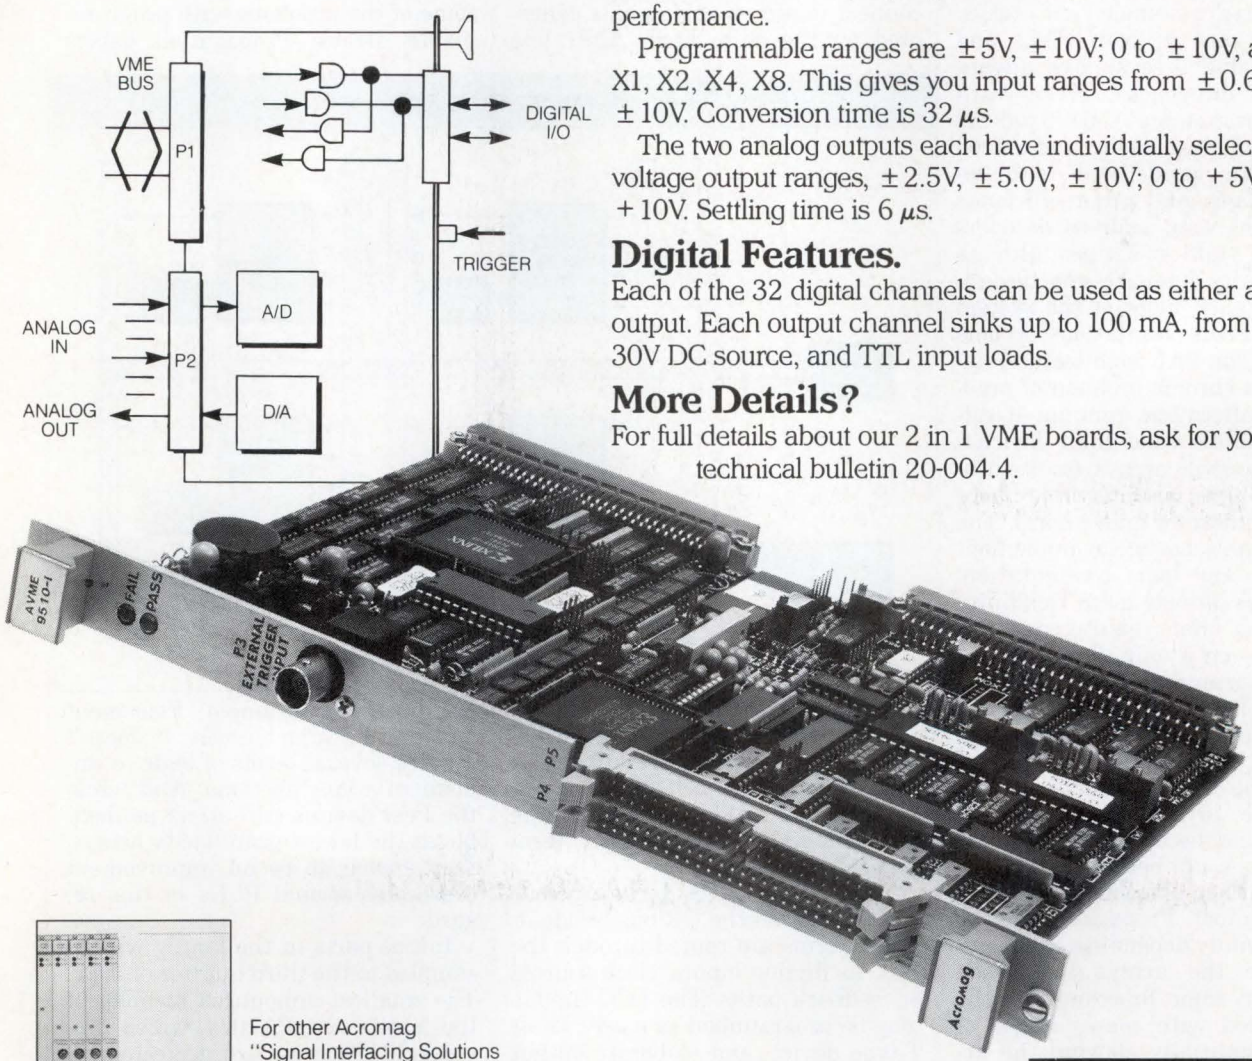

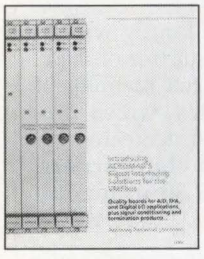

For other Acromag "Signal Interfacing Solutions for the VME bus" ask for bulletin 23-005.0.

OS-9* and PDOS** Drivers

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

CIRCLE NO. 52

Acromag®
LEADER IN SIGNAL CONDITIONING

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

INTEGRATED CIRCUITS

ICT unveils high-density programmable logic architecture

Another participant has joined the debate over architectures for high-density PLDs. The new entrant is International CMOS Technology, a firm well-regarded for its moderate-density standard PALs with electrically erasable programming technology. The Peel Array is ICT's attempt to synthesize a new architecture from conventional PAL layouts and programmable logic arrays.

As ICT sees it, both PALs and programmable arrays have advantages and drawbacks. PALs, with their programmable-AND, fixed-OR approach, offer an easily understood architecture, predefined input-to-output delays, and great efficiency for implementing address decoders and some state machines. But as gate densities increase and designers try to pack more of the system logic into PALs, limitations become apparent. The PAL architecture suffers from a chronic problem of product term allocation, making it difficult to handle wide-input devices.

Programmable arrays, on the other hand, have architectures more similar to those of gate arrays. The programmable parts can implement just about any logic configuration. The devices provide great flexibility for designs that require multiple subsystems on a chip, for example.

But programmable arrays have their problems, too. Like gate arrays, the devices must go through placement and routing steps. But because the logic elements in programmable arrays are somewhat complex, and because the interconnect schemes can become downright arcane, neither placement nor routing is foolproof. Propagation delays can vary wildly depending on layout. And since the arrays lack wide gates, many-input functions must be implemented with many stages of logic, substantially slowing the arrays down on some functions.

The Peel Arrays are PLAs—with programmable-ANDs and programmable-ORs—not PALs. But buried inside the arrays are one-flipflop cells called logic control cells (LCCs). The Peel Array chip looks almost like a conventional channeled gate array. But instead of routing channels, the Peel Array has buses carrying AND and OR terms.

Down each vertical channel on the chip runs an input bus, which carries all of the signals generated in the LCCs to its left. Across each vertical channel runs an AND bus, which crosses all of the input buses. Each line on each AND bus thus can connect to any of the signals generated on the chip. Each AND line forms a product term.

Down the vertical channels, beside the input buses, run OR buses. Since each line in each OR bus crosses all of the AND lines, each OR line can connect to any product term generated on the chip.

Each of the LCCs picks up four OR-terms from the OR bus beside it. These terms are routed through the LCC as flipflop inputs, clock sources or feedback paths. The LCC flipflop may be programmed as a J-K, D- or T-type device, and elaborate output multiplexers let each LCC simultaneously drive a line on the input bus and an output cell on the periphery of the chip.

A PLA with floating registers

In effect, the Peel Array is a very dense PLA with output macrocells that aren't dedicated to I/O pins. This lets the macrocells be used to drive outputs, act as buried registers, or simply serve as feedback paths. In fact, an LCC can perform both register and feedback functions simultaneously.

With this architecture, the Peel Array addresses many of the shortcomings of conventional PALs. Product term allocation isn't an issue, and arbitrarily wide inputs can be accommodated.

In addition, the parts address some of the problems with programmable arrays. Propagation delays are fixed, for example. Placement and routing aren't issues. It doesn't require several levels of logic to implement wide functions. And while the Peel devices still aren't as flexible as the lowest-granularity arrays, they appear to be an improvement over conventional PLDs in this regard.

Initial parts in the family will be sampled in the third quarter of 1989. The smallest announced member of the family, the PA7024, which provides 20 I/O pins, two dedicated inputs and 20 leadless chip carriers, will initially cost $22.50 in quantities of 100. —*Ron Wilson*

THE PEEL ARRAY

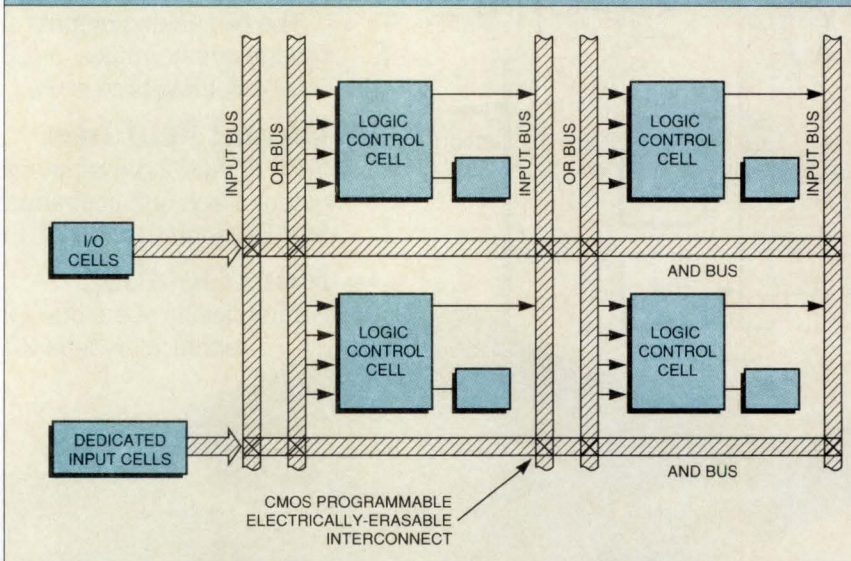

International CMOS Technology
2125 Lundy Ave
San Jose, CA 95131

Circle number 104

RESISTS EVERYTHING BUT A TOUCH

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

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

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

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

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

deeco®

SealTouch and DeeCo are registered trademarks of Digital Electronics Corporation.
VT220 is a trademark of Digital Equipment Corporation. © 1988.

Digital Electronics Corporation, 31047 Genstar Road, Hayward, CA 94544-7831 (415) 471-4700

CIRCLE NO. 53

DESIGN AND DEVELOPMENT TOOLS

Design tool selects PLD, partitions designs

PLDsynthesis, a new Mentor Graphics tool, reverses the usual programmable logic device design sequence. While most PLD design environments require the selection of a PLD family for implementation before creating the design itself, PLDsynthesis lets the designer create the design first, and then automatically selects the best combination of PLDs to meet the selected criteria. Without being constrained to a particular family, the designer can concentrate on functionality rather than PLD selection.

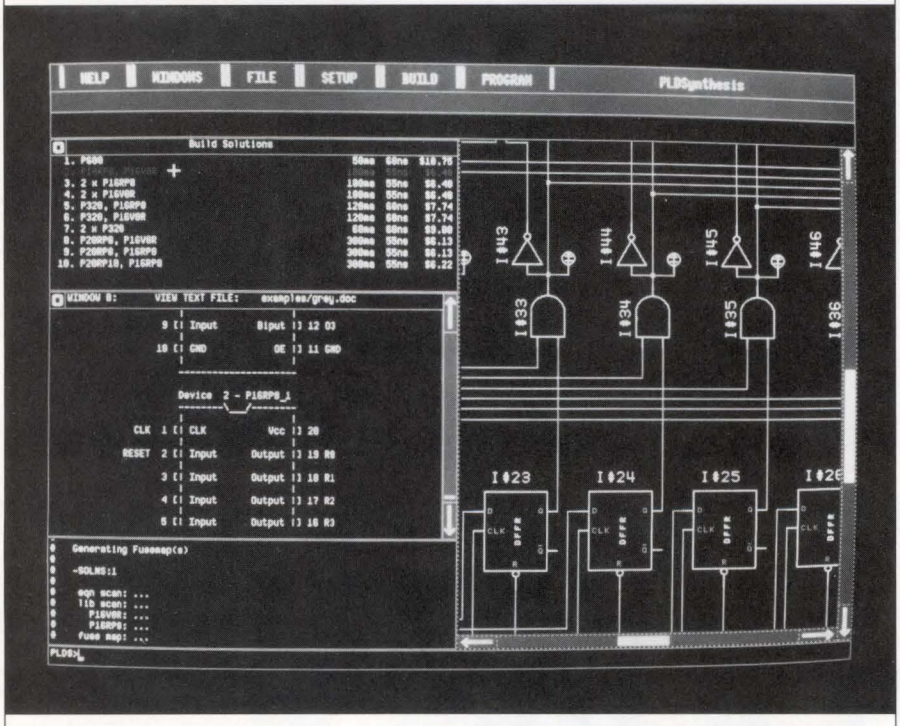

PLDsynthesis isn't a logic synthesizer; that is, it doesn't synthesize gate-level logic from some high-level design description. Instead, it synthesizes a PLD design for a specific PLD family from a generic design input (such as a schematic, a state machine description, a set of logic equations or a truth table) and from a designer's weighted PLD selection criteria.

Design entry first

PLDsynthesis separates the designer's task into two phases: PLD design and device selection. The designer creates the PLD design using Mentor's NETED schematic capture package or a high-level description language integrated with PLDsynthesis. Functional simulation proves operability of the design prior to selecting a specific PLD family. Functional simulation doesn't include timing analysis, since these characteristics can't be known until specific PLDs are selected.

Once the design is functionally verified, the designer accesses the PLDsynthesis device selection capability. Through a PLDsynthesis menu, the designer prioritizes and "weights" various PLD selection criteria, including cost, speed, power, temperature, package type, manufacturer and device count. PLD synthesis then selects, from a library of more than 3,000 devices, the ten devices or combinations of devices that most closely match the designer's selection criteria.

The designer can fine tune the selections by varying the criteria until an optimum selection is achieved. The designer may be willing to trade speed for lower power consumption, for example, but may be constrained as to the number and sizes of the devices that will fit on the board. By varying the weighted criteria, the designer will usually be able to obtain an optimum combination of devices.

Automatic partitioning

After device selection, PLDsynthesis matches the design to the selected PLDs and performs logic reduction to remove redundancies. It then automatically partitions the design among the PLDs, eliminating the need for time-consuming manual partitioning. In addition, PLDsynthesis generates complete documentation for the final PLD design, including reduced equations, simulation results, test vectors, a pinout diagram and a fuse map in the formats of the selected devices.

Mentor's full range of design and verification tools work with PLDsynthesis, including the Quicksim logic simulator, the Quickpath timing analyzer, and Quickgrade and Quickfault fault simulators. The company's Boardstation PCB design and layout system, and Packagestation packaging and thermal analysis systems also work with PLDsynthesis.

Designs that drive PLDsynthesis will also work with Mentor's AVIDesigner (ASIC Vendor-Independent Designer), letting a single logic design be targeted to a wide variety of PLD and application-specific IC solutions. A designer might choose to implement a design in PLDs to get a product to market quickly, and then in an ASIC to reduce component count, for example. The original design, used for the PLD implementation, may be used unchanged for the ASIC implementation.

PLDsynthesis is priced at $14,900, and is scheduled for June delivery.

—Bill Harding

Mentor Graphics
8500 SW Creekside Pl
Beaverton, OR 97005

Circle number 103

Coming May 1

Watch for
Dave Lieberman's
Technology Focus Report
on 16-bit buses.

"No, we were back up by noon. The support people at Summagraphics found our problem. Yeah . . . I'm glad we went with them, too."
Summagraphics. Tablet of choice for today's computer professional.
Summagraphics®
60 Silvermine Road, Seymour, CT 06483
(203) 881-5400.
See us at NCGA Booth #4127
CIRCLE NO. 54

DESIGN AND DEVELOPMENT TOOLS

Simulator accesses internal device registers

Logic Automation has added a new capability to its Smartmodel family of behavioral models for simulation of complex VLSI devices. Previously, Smartmodels let users view activity at the I/O pins of a device during simulation. Logic Automation's new Smartmodel Windows family lets users access registers inside the models as well. Because today's microprocessors can have 90 or more internal registers, and a peripheral controller another 40 or 50, it's very difficult to determine what's happening when simulating a system without being able to view the registers inside the complex device models. The enhanced Smartmodels provides a transparent view of the whole system during simulation, a capability that allows a level of debugging not previously possible.

In addition to viewing register contents and tracking their changes, Smartmodel Windows lets the user force values into registers, set breakpoints on register values and single-step through simulated operations, just as if the registers were modeled individually. Users can combine Smartmodel Windows register values with other system values to set breakpoints, according to the breakpoint algorithms of the supported simulator.

Many models supported

The initial Smartmodel Windows release supports the Motorola family of microprocessors, up to the 68020. In addition, it supports the 68230, the 68440, the 68442, the 68452, the 68851, the 68881 and the 68HC11. A full range of PLD models is included, supporting PLDs from Advanced Micro Devices, Altera and Texas Instruments, among others.

Smartmodel Windows depends on the user interface of the host simulator to view and manipulate internal registers. For this reason, most simulators must be modified to accept Smartmodel Windows. The initial release supports the Mentor Graphics Quicksim logic simulator.

Logic Automation doesn't sell individual models; instead, users buy a subscription service for the entire model library. The subscription for the Smartmodel library is $7,900 per workstation. The Smartmodel Windows option is an additional $4,000. The subscription entitles the user to all models introduced during the term of the subscription, and to other library updates. —*Bill Harding*

Logic Automation
19500 NW Gibbs Dr
Beaverton, OR 97006

Circle number 107

5.4 Gigabytes Unattended Backup

Digi-Data's GIGASTORE™ provides up to 5.4 Gigabytes of data storage on a single T-120 VHS video cartridge. That permits backup of your largest disk drive on off-hours without an operator.

Utilizing true read-after-write coupled with very powerful error correction, GIGASTORE gives you an unsurpassed error rate of 1 in 10^{23} bits. In addition, you get a high speed search capability not available in most 9-track drives.

GIGASTORE can be provided with an interface for IBM PC/XT/AT/PS-2™. It is available with DEC interface for VAX and MicroVAX. It is also available for backup of data on Novell LANs.

Call Digi-Data, an organization with a 25 year history of manufacturing quality tape drives, at (301) 498-0200.

™GIGASTORE is a trademark of Digi-Data Corporation. PC/XT/AT/PS-2 are trademarks of IBM Corporation.

DIGI-DATA CORPORATION
8580 Dorsey Run Road
Jessup, MD 20794-9990
(301) 498-0200
Telex 87-580

® ... First In Value

DEALER INQUIRIES INVITED

In Europe contact: Digi-Data Ltd. • Unit 4 • Kings Grove • Maidenhead, Berkshire England SL6 4DP • Telephone No. 0628 29555/6 • Telex 847720

CIRCLE NO. 55

Hardware modeling system interfaces to any logic simulator

The LM-1000 from Logic Modeling Systems is a universal hardware modeling system that readily interfaces to logic and fault simulators running on all popular engineering workstations. Designed to work as a resource that can be accessed by any workstation on the network, the LM-1000 works with simulators offered by Gateway Design Automation, Genrad, Valid Logic Systems, Vantage Analysis and Viewlogic.

An LM-1000 system consists of control and interface electronics, pattern memory, pin electronics for device sensing and driving, and the hardware models themselves. Core modeler software and a software shell for each installed model also reside in the system.

A simulator function interface (SFI) lets a simulator access the LM-1000. Along with a set of LM-1000 utilities, the SFI resides in a workstation with the simulator. While Logic Modeling quotes one month as the development time for a new SFI, a new interface can be developed in as little as one week.

Technical advances

According to Logic Modeling, LM-1000 performance and access time is typically ten times faster than that of other systems, with round-trip access of 7 ms or less. Round-trip access is defined as the time that it takes for a software simulator to request information from a hardware model in the LM-1000, for the LM-1000 to perform the necessary tasks to generate the information, and for that information to be presented to the simulator.

Pattern memory is 320 bits wide, corresponding to the maximum device pin-out in the current release. The pattern presentation rate is as high as 25 MHz over the full pattern memory width, with a memory depth of up to two million patterns.

The device to be modeled (the reference element) is mounted on one or more device adapters. Each LM-1000 device adapter supports a device of up to 80 pins. A single LM-1000 chassis supports up to 32 device adapters, arranged in eight slots of four device adapters per slot. A model must fit within one slot, so up to four device adapters can be ganged to construct models of devices having up to 320 pins. A single slot can support four 80-pin models; two 160-pin models; one 160-pin model and two 80-pin models; one 240-pin model and one 80-pin model; or one 320-pin model. Models can

(continued on page 104)

WHY DO YOU REGULARLY READ COMPUTER DESIGN?

"Because COMPUTER DESIGN gives the overall systems picture... it doesn't just emphasize one component."

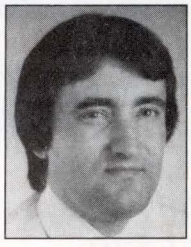

A. Karagiannis
Program Engineer
SUNDSTRAND CORP.

Mr. Karagiannis is a Program Engineer at SUNDSTRAND CORP., and has been a reader of COMPUTER DESIGN for 4 years. His engineering management responsibilities include system design, development and integration.

As a system design engineer, we'd like to know why you read and prefer COMPUTER DESIGN. Drop us a line at One Technology Park Drive, Westford, MA. 01886.

COMPUTER DESIGN®

THE FIRST MAGAZINE OF SYSTEM DESIGN, DEVELOPMENT AND INTEGRATION

DESIGN AND DEVELOPMENT TOOLS

use fewer pins; for example, an 80-pin device adapter can support a 64-pin device model.

ASIC verification

In addition to modeling VLSI components for simulation, the LM-1000 may be used to verify prototype ASICs. When an ASIC is designed using one of the simulators that the LM-1000 supports, the vectors that were used to verify the ASIC design may be used to verify that the prototype silicon matches the original design. The user simply creates a hardware model of the ASIC using the prototype silicon and runs the same simulation using the hardware model that was originally run using the ASIC design.

The original simulation outputs are compared with simulation outputs using the hardware model. If the results are the same, the prototype silicon matches the design.

Timing verification of the prototype ASIC is performed by actually measuring propagation delays in the hardware model. The user can view the timing results either as waveforms created by the simulator or as a device characterization file created by an LM-1000 utility.

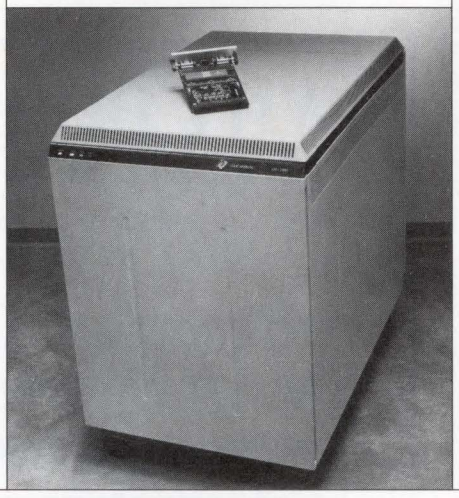

The LM-1000 supports workstations from Sun Microsystems, Apollo and Digital Equipment Corp, as well as 80386-based PCs. Operating system compatibility includes SunOS, Domain/OS, DOS and VMS. Compatible networking protocols include Ethernet and TCP/IP. Logic Modeling distributes the LM-1000 hardware modeling system and models through its OEM partners rather than directly to end-users. OEM partners include Gateway Design Automation, Genrad, Valid Logic Systems, Vantage Analysis and Viewlogic.

A base system, capable of modeling one 80-pin device, costs $50,000. Hardware models cost from $1,500 to $3,000. *—Bill Harding*

Logic Modeling Systems
2880 Zanker Rd
San Jose, CA 95134

Circle number 102

OBJECT MEASUREMENTS AT VIDEO FRAME RATES
APA512
APA512S

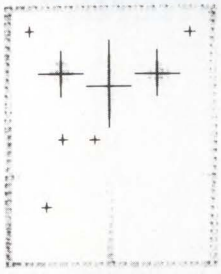

Multiple Objects

You can measure, identify, select or track up to **255 objects each 33 milliseconds!**

Real-time SRI Algorithms

16 hardware calculated parameters for each object provide **Connectivity, Moments, Perimeter, Area, Polarity** for each object each frame time. **34 Software Functions** perform SRI Algorithms in real time on multiple objects.

***New Development System:* APA512S**

VME Processor, Chassis; Image Processor, APA512 Development Software, Integration, Installation, Training, Call in Support.

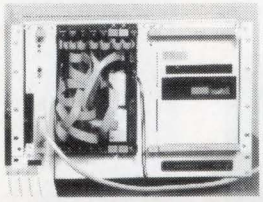

Call us for Details!
408 295-1661

Vision Security, Inc.
1762 Technology Drive #124
San Jose CA 95110

CIRCLE NO. 61

THANK YOU!

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

RISC coprocessor gives AT a boost

The AT-class MS-DOS machine has become the low-cost platform of choice for virtually every conceivable application area, and the continual appearance of leading-edge coprocessor boards for the platform will keep it vital for years to come. A new Yarc Systems board, dubbed the AT-Super, puts a powerful 25-MHz Advanced Micro Devices 29000-based engine on the AT bus, yielding a machine that's more powerful than a Sun 4 workstation, according to Trevor Marshall, company president.

"The computer marketplace has changed drastically in recent years," says Marshall. "Applications that required large minicomputers or mainframes can now be performed on personal computer workstations. Hardware that can cost millions of dollars to purchase and hundreds of dollars an hour to operate has been successfully challenged on the basis of delivered performance."

Scalable multiprocessing

Targeted specifically at number-crunching and high-performance graphics applications in the scientific and engineering arenas, the board has an architecture that's suited to building scalable multiprocessing machines with up to four of the boards in a PC AT or 80386-based machine. By means of its bus master mode, the AT-Super can gain direct access to PC memory cards, I/O cards and other system resources, including the memory on other AT-Super cards.

To keep the 29000 adequately supplied with data and instructions (and prevent the need to access the AT bus) the board packs in 512 kbytes of SRAM and 2 Mbytes of high-speed instruction memory. The board's memory architecture uses multiple interleaved banks and burst transfer techniques to outstrip the capabilities of cache memory architectures, according to the company.

Using a 50-MHz system clock, the AT-Super achieves 25-Mips peak performance, 17-Mips sustained. Using Metaware C, it executes over 25,000 Dhrystones/s. With an optional 29027 floating-point chip in place, the board is capable of 3 million Linpack operations/s, and boasts six-layer construction with special bypass capacitors to minimize radio frequency interference. An expansion connector is provided for links to external equipment, facilitating development work. The board requires a 5-V supply and draws 10 W maximum.

"The de facto acceptance of the PC workstation has caused a major shift in the thrust of VAR and OEM marketing programs," says Marshall. "As increased performance becomes available in the PC workstation, software packages that once required expensive large-scale computers will be marketable in volume to a larger base of users. The resultant lowering of software prices will expand the market for these workstations."

The AT-Super is priced at $4,595 in single-unit quantities. Samples are available now; volume quantities will be available second quarter. The Metaware C compiler, a DOS run-time interface and such tools as a macro assembler, a linker and a debugger are also available from the company. —*David Lieberman*

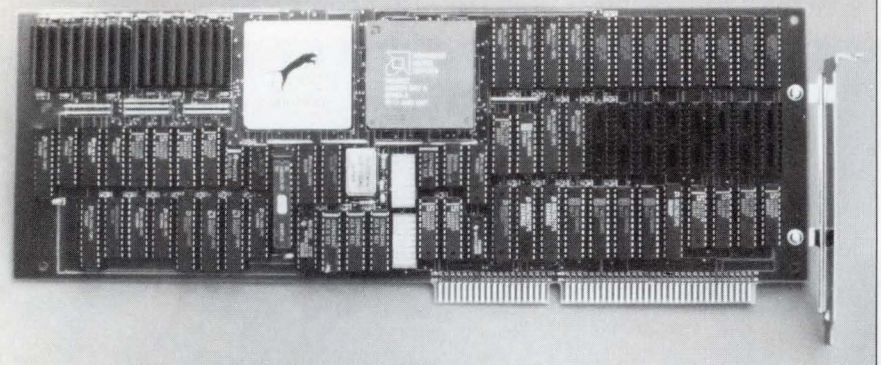

Yarc Systems
5655 Lindero Canyon, #721
Westlake Village, CA 91362

Circle number 109

Addendum

The telephone number for JRS Research Laboratories (Orange, CA) was missing from "VHDL opens the road to top-down design" (p. 58, February 1, 1989). The number is (714) 974-2201.

In "Local buffers minimize LAN data-rate demands" (p. 117, February 1, 1989), the term microseconds (μs) was inadvertently changed to milliseconds (ms) during the typesetting process.

100,000 + FAX Numbers — of — U.S. Businesses including Names & Addresses

- Quick reference
- Maximize FAX usefulness
- New advertising medium

Only **$49.95** Postage Paid

Order Now!

Entrepreneurs' Library
Box 17509, Fountain Hills, Az. 85269
☐ Please rush me a copy of the new Fax Directory. Enclosed is check or money order for $49.95.
Name _____
Company _____
Address _____
City _____
State _____ Zip _____

CIRCLE NO. 57

COMPUTERS AND SUBSYSTEMS

Intelligent NuBus I/O modules fit ruggedized Macintosh for factory use

A family of configurable, NuBus-based I/O modules lets ruggedized versions of the Apple Macintosh II be harnessed as a user interface for real-time industrial control applications. The Rackmac from Greenspring Computers is a repackaged Mac II in a 19-in. rack mount that can be fitted with a custom display controller that supports touch-screen operations. The Rackmac also has 40- or 80-Mbyte Winchester drives that can take 80-G shocks.

The new modules let the system be configured for a wide variety of I/O capabilities. Software and plug-in daughter cards called industry packs (IPs) reside on the boards, dubbed IP carriers. One IP carrier, called the Support Board 1270, is a low-cost platform that can accept two IPs, such as an analog-to-digital converter, or a 24-bit parallel I/O module. The other board, called the Spring Board 1260, is an intelligent I/O channel processor that sports its own 68020 microprocessor, an optional 68881 floating-point coprocessor and up to 4 Mbytes of on-board RAM.

There are about a dozen different IPs available. Each contains an ID PROM that identifies pack specifications on power-up for automatic system configuration. Since all IPs share the same logic protocol, are mechanically interchangeable and fit two to a carrier, users have access to a vast array of configurations.

Among the IPs available are Centronics, SCSI, RS-232 and IEEE-488 modules. An Opto-22 parallel I/O pack supports 24 I/O signals that can be individually designated as input or output and can be driven as TTL high current, CMOS, open collector or direct relay. The IP is plug-compatible with Opto-22 direct I/O isolation panels.

Other IPs include a 20-channel, 12-bit A-D converter with two channels of D-A, a hex D-A converter, a shaft encoder input module and a module supporting the Intel Bitbus. The A-D converter is factory calibrated in PROM and provides autocalibration in the driver software. Greenspring also provides a prototyping board with interface logic to let users develop their own IPs. The board connects with flat cables to the IP connectors on the 1260 or 1270 carrier board.

Real-time applications

For developing real-time applications using the intelligent Spring Board 1260, the developer can use the Macintosh Programmer's Workshop (MPW) development tools and a set of coprocessor-oriented tools provided by Greenspring. The MPW contains a shell for the user interface, an MC680X0 assembler, Pascal and C compilers, and link and make facilities.

The operator interface portion of the application is designed to run in Macintosh memory, and the real-time portions to run on the Spring Board's 68020, which can itself be a bus master. To create the Spring Board portion of the application, the developer uses a software tool provided by Greenspring to convert the code resources developed using the MPW facilities into coprocessor resources that can be loaded directly onto the Spring Board.

Spring Board resources can also be developed with other software tools that require knowledge of only C, Pascal or a 680X0 assembler, and not the Macintosh interface. These applications can later be linked with the user interface portion to run on the Rackmac. A completed application is thus a Macintosh application that automatically loads the Spring Board portion onto the coprocessor board and handles the communication between the real-time application and the user interface.

The Spring Board EPROM contains a debugger for both the 68020 and the 68881. The debugger can interface via the Macintosh interface, with a monitor application provided by Greenspring, or in a TTY mode with any terminal emulator. Development and debugging of real-time applications can thus be done outside the Macintosh environment.

The Rackmac with its new I/O subsystem boards and IPs is available now. Prices vary depending on the combinations selected. A support board with one IP costs about $500, while a fully loaded system with all options costs about $15,000.

—*Tom Williams*

Greenspring Computers
1204 O'Brien Dr
Menlo Park, CA 94025

Circle number 106

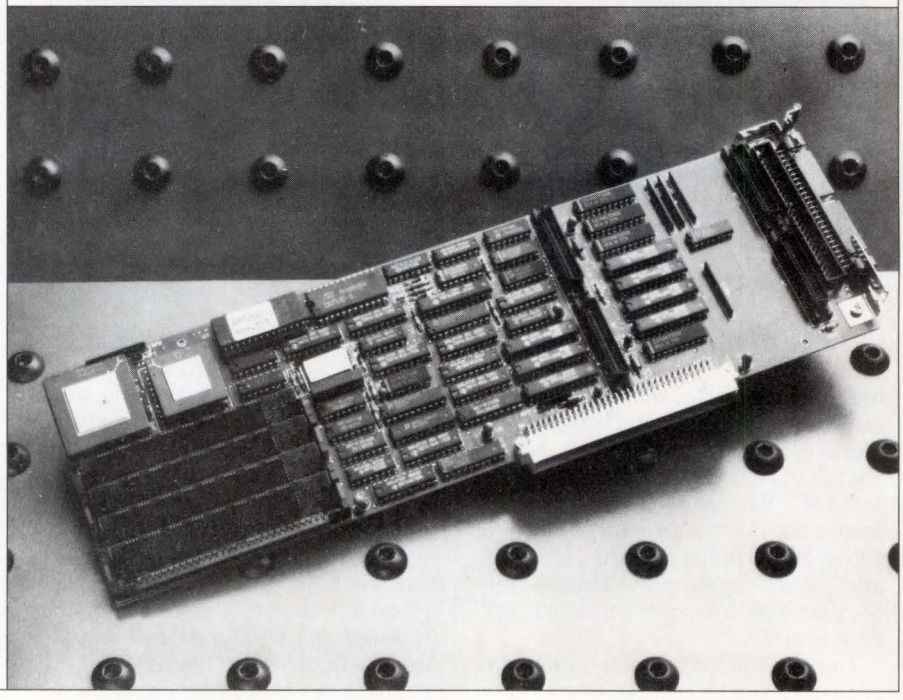

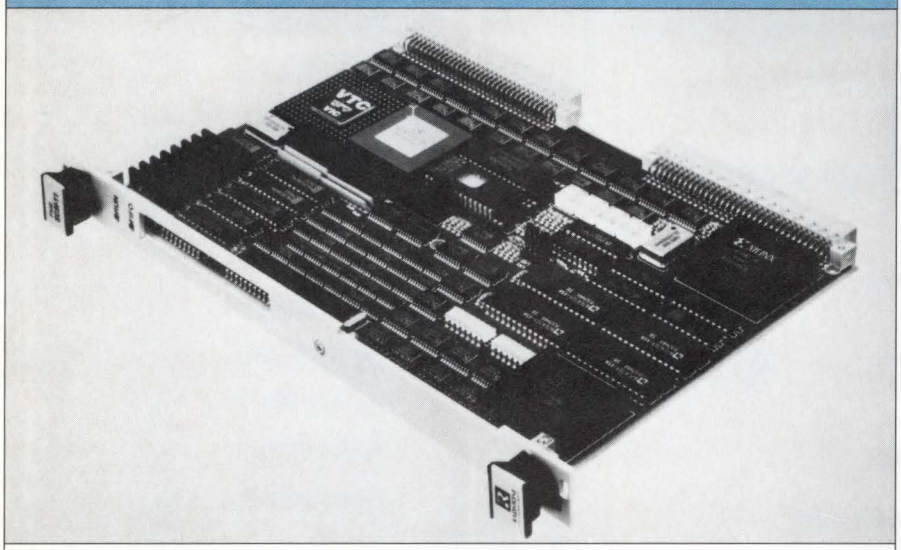

VME SCSI adapter has four-bus architecture

The use of multiple-bus architectures to prevent the logjams that can occur on single-bus systems is fairly common. The SCSI-11 host adapter from Radstone Technology, however, is unusual in the lengths to which it goes to provide multiple data pathways and avoid these bottlenecks.

The board aims to meet the needs of speed-critical multiprocessor designs and very high performance I/O systems by means of a four-bus architecture. At the same time, it strives to avoid what Pete Yeatman, vice-president of marketing at Radstone, calls the "moving I/O bottleneck syndrome": simply moving the I/O bottleneck from the bus onto the processor board.

The 68020-based SCSI-11 ensures an effective decoupling of the VMEbus system and SCSI subsystem with memory that includes 512 long words of very high speed FIFO buffer and a 1-Mbyte block of quad-ported DRAM that permits simultaneous data transfers across all four buses without interleaving.

The board includes a VSB (VME Subsystem Bus) interface, useful for offloading traffic from the VMEbus. The VSB interface on the board is a full master/slave implementation with its own high-speed DMA.

A proprietary 32-bit APEX (Advanced Processor Extension bus) connector on the board rounds things out with a local microprocessor path to a special-function daughterboard. Apex daughters are available for floppy disk control, a second SCSI bus and so forth.

The ability to use all four buses simultaneously provides both power and flexibility. "You can send data in one bus and out another while sending control and status down a third and possibly offloading removable data transfers via a floppy APEX interface," Yeatman says.

Both initiator and target

As far as the SCSI bus is concerned, the board functions as both an initiator and a target and can be used for synchronous or asychronous operations at up to 5 Mbytes/s and 3 Mbytes/s, respectively. The board can be used for single-ended or differential SCSI subsystems.

On the VMEbus side, the company claims the SCSI-11 to be the first VME board to put VTC's VIC068 VME chip to work pushing data across the VMEbus at burst rates up to 40 Mbytes/s. The board interfaces to a host CPU or CPUs via its quad-ported memory. The host constructs a high-level device-independent parameter block in VSB or VME global memory, then places a pointer to the command block in the SCSI-11's local memory. The host is freed from any involvement in subsystem operations until the operation is done.

The SCSI-11 is priced at $2,195.

—David Lieberman

Radstone Technology
20 Craig Rd
Montvale, NJ 07645

Circle number 105

Host Independent Multi-Channel Graphics

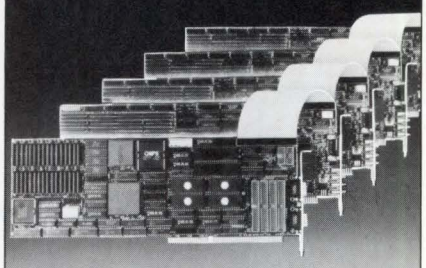

Now there is a unique graphics display system (GDS) that provides multiple display channels and display list processing for a single PC. The **OMNI 1500 GDS™** can also operate external to its PC host (via a passive backplane) and interface with a variety of platforms. It maintains software compatibility across a broad selection of host computers through its firmware-resident **OMNI*KERNEL™** system, which incorporates the functionality of the GKS level 2C+ standard.

- **Flexibility:** Supports as many as 8 graphics "heads", using four serial device ports, **DEC™** standard DMA and programmed I/O ports. Modular expansion with plug-in boards.
- **Resolution:** Up to 1280 by 1024, 60Hz non-interlaced.
- **Power:** Uses 80386 processor and 80387 coprocessor, firmware-based graphics kernel, up to 4 MBytes of display list memory and 512 KBytes of EPROM.
- **Host independent:** Through Omnicomp interface boards for **PS/2®, Multibus™, PC/XT®** and **PC/AT®.**

Display List Processing With The OMNI 1500 DLP™

DISPLAY LIST MICROPROCESSOR	80386 with 80387 coprocessor
LIST MEMORY	1, 2 or 4 MBytes
TOTAL EPROM MEMORY	512 KBytes
USER DEFINABLE MEMORY	256 KBytes
PORTS	Four RS232-C, plus DRV11-W(DMA)/DRV11(PI/O)
GRAPHICS KERNEL	OMNI*KERNEL in Firmware

Graphic Display Controls

OVERLAYS **OMNI 1600 GDC™:** **OMNI 1400 GDC™:**	 4 (Independent zoom in overlays) 2
MEMORY ORGANIZATION	2048 x 1024 x 4 or 8 planes
COLORS	16 or 256, simultaneous
PALETTE	4096 or 16.7 million colors
MANIPULATORS	Hardware cursor, pan and zoom

Omnicomp™ Graphics Corporation

1734 West Belt North • Houston, TX 77043
Phone: (713) 464-2990
Telex: 285801 OMNICO UR
Fax: (713) 827-7540

Copyright 1988, Omnicomp Graphics Corp.
Omnicomp reserves the right to change product specifications at any time without notice.
OMNI*KERNEL, OMNI 1500 GDS, OMNI 1400 GDC, OMNI 1600 GDC, and OMNI 1500 DLP are trademarks of Omnicomp Graphics Corporation. Other trademarks are as identified.

CIRCLE NO. 58

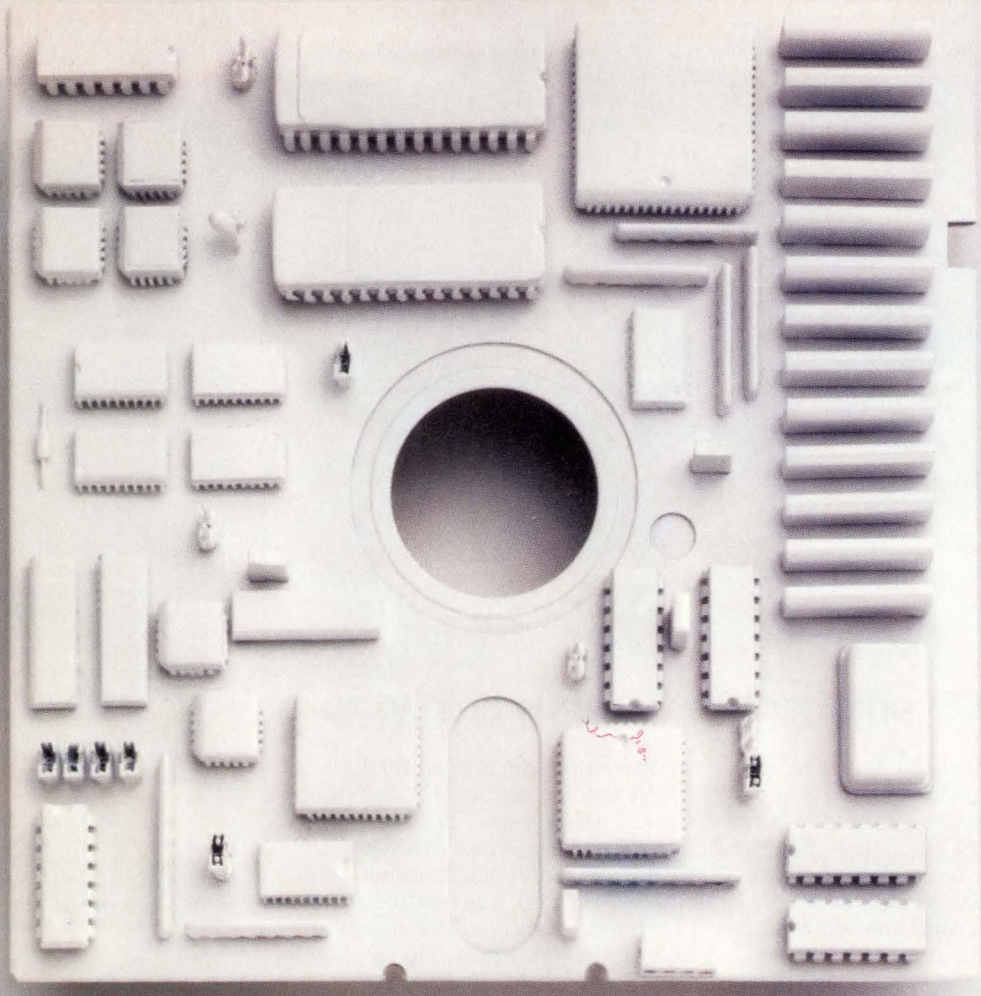

IF YOU'RE SPENDING TOO MUCH TIME ON SNA DESIGN, YOU'RE LEAVING OUT ONESTEP.™

Most likely, if you need SNA compatibility, you'll buy the software source code from one company and hardware from another. Then you'll spend several months desperately working out the integration.

Or worse, you'll design it from scratch.

Systech knows you've got better things to do with your time. So we've created a way to cut down that lengthy process. With OneStep.™

OneStep is a complete SNA/DCP board package. Which makes it the quickest and easiest off-the-shelf solution for IBM communication. In fact, with our help, the OneStep can be up-and-running in just five days.

To make this unique integration work, we teamed up with the most reputable and experienced company in IBM compatibility around—Systems Strategies, Inc. With their superb emulation software and Systech's modular design, OneStep gives you the wide-area links you need for VME and Multibus systems. What's more, the OneStep takes a load off your CPU. That way, you have more power and improved system performance overall.

Systems Strategies, Inc.'s extensive variety of popular protocols includes SNA 3270, RJE and LU6.2. BSC 3270 and RJE. Plus, X.25 and QLLC. All thoroughly tested and successfully ported to Systech's hardware.

Along with easy installation, the OneStep will keep your up-front costs down. And its modularity makes it easy to maintain and upgrade.

Most importantly, with the experience and expertise of Systech and Systems Strategies, Inc., you'll be in step with the latest software and hardware developments. And you'll always be OneStep ahead of the competition.

For a demonstration or more information, call Systech at (619) 453-8970. Or write: Systech Corp., 6465 Nancy Ridge Dr., San Diego, CA 92121.

SYSTECH

WE OPEN UP A LOT OF POSSIBILITIES.

OneStep is a trademark of Systech Corporation.

CIRCLE NO. 62

MAJOR SYSTEM COMPONENTS

Floppy drive pushes capacity to 16 Mbytes

With the introduction of the Model 3511, Panasonic Industrial has upped the ante in the competitive 3½-in. floppy drive market. The large-capacity drive can store 16 Mbytes of unformatted data and 11 Mbytes of formatted information on a metal disk. With 640 tracks on each side of the disk, the drive achieves a recording density of 35,000 bits/in.

To increase linear recording density, Panasonic developed the metal disk media along with two specially designed recording heads (a Sendust head and a Metal-In-Gap head), one of which will be chosen for mass production after final evaluation. The 3511 pushes track density to 542 tracks/in. by using a microstepping motor with an encoder as a head-positioning actuator. The motor increases high-vibration resistance, delivers high-speed access, improves head-positioning accuracy and lowers power dissipation. The drive uses a modified-frequency-modulation recording format and Reed-Solomon error-correcting code.

Measuring only $1.5 \times 4 \times 6$ in., the drive includes an embedded small computer system interface, a built-in exclusive controller and a self-format function to let the user initialize the physical format. In addition, the drive is read-only compatible with conventional 1- and 2-Mbyte 3½-in. drives.

Average access time, including settling time, is less than 60 ms. Track-to-track access is less than 10 ms, and average latency is 50 ms. The drive's precision track-positioning system falls within ± 2 μm. Power dissipation is +5 V at 140 mA and +12 V at 230 mA.

Samples of the 3511 will be available in second quarter 1989, with production scheduled for July 1989. In OEM quantities, the drive will be priced at $250. —*Michael Donlin*

Panasonic Industrial
1600 McCandless Dr
Milpitas, CA 95035

Circle number 113

Quality in. Quality out.

When you put your best into your products, it shows.

At Compaq, the value of our products is simply a reflection of the personal values inherent in each Compaq professional. Attributes like teamwork, mutual respect, leadership and a commitment to quality translate into some of the most advanced personal computers on the market.

Our continued success has created immediate opportunities for top CAE Engineering professionals at our Houston headquarters. Start with a degree and relevant experience. Add to that a strong set of personal values and the desire to do what's never been done before, and you could be part of our winning team.

Senior CAE Engineers

Bring new design and emulation methodologies on line. You'll have a hand in definition, evaluation and technical support of design methodology and CAE tools for system, board and ASIC designs. You'll also work with CAE/CAD vendors providing solutions to complex issues.

Start with a BS or MS in Electrical Engineering or Computer Science, along with five years' experience in CAE/CAD. Experience developing or supporting CAE tools and workstations is essential. Knowledge of UNIX and 'C' programming is a plus.

CAE Application Engineers

Integrate and test CAE tools for system, board and ASIC designs. You'll evaluate and integrate new CAE tools into new and existing CAE environments. Positions are available for CAE Engineers with experience in HDL, behavioral simulation, hardware modelers, PLD designs, logic synthesis, logic & fault simulation, timing analysis, analog/SPICE simulation, mixed analog/digital simulation, automatic test generation and design-for-test.

You'll need a BSEE or BSCS and two to four years' experience developing, supporting or using CAE tools. Knowledge of ASIC/board design, valid logic systems, gateways, Verilog-XL and UNIX/'C' environments is a plus.

CAE Application Support Engineers

Support CAE tools on Compaq workstations and super-mini/mainframe computers. You'll track, document and monitor CAE application issues, as well as perform regression testing on CAE tools before release. You'll work with design engineers, solving CAE problems and promoting more effective use of CAE tools. In addition, you'll provide user feedback and perspective for future revisions and selections of CAE tools.

Your BSEE or BSCS should be complemented by one to three years' CAE experience. Knowledge of UNIX and 'C' environments is a plus.

Make a quality career decision.

Compaq offers competitive salaries, comprehensive benefits and an unequaled work environment. If you're interested in one of these opportunities, please call us at **1-800-243-9003**. Or simply submit your resume and the position for which you wish to be considered to:

Compaq Computer Corporation,
Dept. CD315-MW, P.O. Box 692000,
Houston, Texas 77269-2000.

© 1989 Compaq Computer Corporation. All rights reserved. All products are registered trademarks or trademarks of their respective manufacturers. Compaq is an affirmative action employer, m/f/h/v.

COMPAQ

CIRCLE NO. 59

SYSTEM SHOWCASE

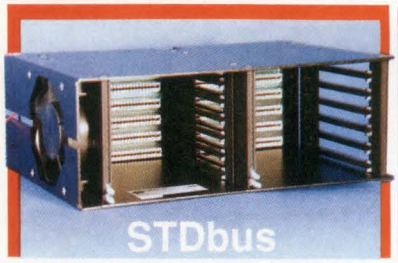

12-SLOT CARD CAGE/BACKPLANE ASSEMBLY

OFFERS MORE FUNCTION CAPABILITY IN LESS SPACE
INEXPENSIVE ALTERNATIVE TO MULTIBUS I
WITH SIMILAR FORM FACTOR
IDEAL FOR PC BUS EMULATION APPLICATIONS
CAN BE USED FOR 8 TO 16 BIT APPLICATIONS
BUILT-IN COOLING SYSTEM

Hybricon

12 WILLOW ROAD, AYER, MA 01432
FACTORY 508-772-5422
WESTERN OPERATIONS 602-921-1824

CIRCLE NO. 230

PAL™/PLD SOFTWARE
Sets The Standard

CUPL™ PLD compiler, the most powerful language for state machine logic design, now allows OrCAD™ schematic software as the front end design entry. CUPL supports all PLDs and carries the most extensive update program. Available on MS-DOS™, Apollo™, Sun™, VAX™ and most UNIX™ based platforms.

LOGICAL DEVICES, INC.
1201 N.W. 65th Place
Ft. Lauderdale, FL 33309
305-974-0967
1-800-331-7766

CIRCLE NO. 231

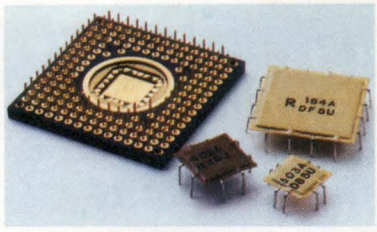

STOP NOISE IN PGA, LCC PACKAGES

MICRO/O 3000 capacitors decouple low-inductance, high-frequency noise for PGA, LCC packages on complex board layouts. They fit under PGA or LCC sockets, take no extra board space. Choose from many pinout configurations. Rogers Corp. 2400 South Roosevelt St., Tempe, AZ 85282. 602/967-0624.

CIRCLE NO. 232

REAL WORLD I/O
For PC/XT/ATs

DG24 • 24 line digital I/O; 10 MHz 8255. **$95**
AD500 • 8 channel 12-bit (plus sign) integrating A/D; prog. gains of 1, 10 & 100; 7 digital I/O lines. **$239**
AD100 • Single channel version of AD500; 10 digital I/O lines. Same programmable gains. 700 meg input Z. **$149**
AD1000 • 8 channel 12-bit A/D; 25 uS; sample & hold; 3 5 MHz timer/counters; 24 digital I/O lines. **$295**
ADA300 • 8 channel 8-bit 25 uS A/D; single D/A sample & hold; 24 digital I/O lines. **$239**
AD200 • 4 channel 12-bit 125 uS A/D; 3 5 MHz timer/counters; 24 digital I/O lines. **$239**
DA600 • Fast settling dual bipolar D/A. **$179**
PD200 • Prototype board w/ address decoder; manual **$99**

All boards include BASIC, Pascal, C, and Forth drivers. 30 day return; 1 year warranty. Call for "Real World Interfacing" application notes.

Real Time Devices, Inc.
P.O. Box 906 State College, PA 16804
(814) 234-8087

CIRCLE NO. 244

IEEE 488 Solutions

- Hardware & software interfaces for PC, AT, 386, PS/2, Macintosh, SUN, HP & DEC
- IEEE converters to SCSI, RS-232, RS-422, modem, Centronics, digital I/O & analog I/O
- IEEE extenders, expanders & buffers
- IEEE drivers for DOS, UNIX®, Lotus 1-2-3, Symphony & Quattro.

Call or send for your FREE Technical Guide
Demo disks and application notes available

IOtech

(216) 439-4091
Telex 6502820864 • Fax (216) 439-4093
IOtech, Inc. • 25971 Cannon Road
Cleveland, Ohio 44146

CIRCLE NO. 234

IEEE-488 for VMEbus & MULTIBUS

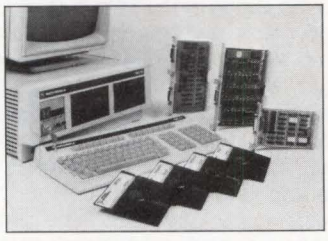

Hardware Options — High-speed DMA or low cost programmed I/O
Software Support — UNIX, PDOS, MTOS, VERSAdos, Aegis

NATIONAL INSTRUMENTS™
The Software is the Instrument™

Call for FREE Catalog
512/250-9119 or 800/531-4742

CIRCLE NO. 235

DATA I/O SET PROGRAMMER

NOW $1595.

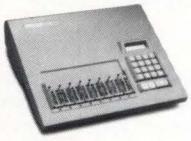

Save $400 on the eight-socket 280 Set Programmer, and get superior set/gang programming capabilities for all EPROMs and EEPROMs up to 512K. Plus:

- High-speed algorithms
- Stand-alone or computer remote control operation
- Full after-sales support from Data I/O®

1-800-247-5700
Ext. 815

DATA I/O Corporation

CIRCLE NO. 236

CDC HIGH PERFORMANCE DRIVES
LAN KITS for PC's, Work stations & DEC computers

106MB	5¼"	18ms	40,000MTBF	$995
182MB	5¼"	16.5ms	40,000MTBF	$1295
228MB	5¼"	18ms	40,000MTBF	$1595
383MB	5¼"	14.5ms	40,000MTBF	$2195
385MB	5¼"	10.7ms	40,000MTBF	*Call*
766MB	5¼"	14.5ms	40,000MTBF	$3295
1.2GB	5¼"	16.5ms	40,000MTBF	*Call*
1.154GB	8"	16ms	50,000MTBF	*Call*
1.123GB	8"	16ms	50,000MTBF	*Call*

- ❑ Complete LAN KITS, CD ROM NETS, WORM NETS, ERASABLE OPTICAL DRIVES. Custom configurations available.
- ❑ High performance drives from Maxtor, Priam and other reputed manufacturers available
- ❑ Let us configure and custom package storage to your specifications. Call us today for a fast quotation
- ❑ We accept International orders

EDS (714) 634-4441

CIRCLE NO. 237

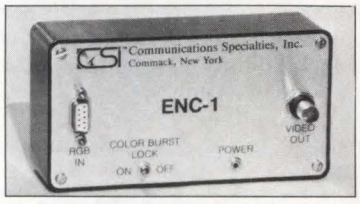

CONVERT RGB INTO COLOR COMPOSITE VIDEO

Now you can convert RGB computer signals into NTSC composite video with the ENC series of RGB Encoders. Three different models are available for compatibility with most analog and TTL RGB computers having 15 kHz scan rates. Output of ENCs can drive VCRs, projection TVs, and monitors. Interface cables and application assistance available. Only $395.00 each.

Communications Specialties, Inc.

Tel. # (516)273-0404, FAX # (516)273-1638

CIRCLE NO. 238

Circle 220 for rates and information, or call: Shirley Athorn or Barbara Kovalchek 800-225-0556. In Massachusetts 508-692-0700.

E/EPROM & MICRO PROGRAMMER

$895

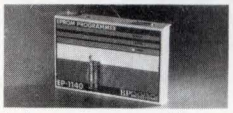

- EP-II4O includes: software, cable, user's manual, 2 free software update coupons, toll-free technical support, one-year warranty & a unconditional 30-day money back guarantee
- Programs 24-, 28-, 32- & 40-pin E/EPROMs
- Supports 874X & 875X series microcontrollers
- Connects to a standard parallel port
- Hardware to support future parts to 8 Megabit
- Software updates on floppy disk
- 32-pin model, EP-II32, available for $695

The Engineer's Programmer™

CALL TODAY 800-225-2102

BP MICROSYSTEMS
10681 Haddington, #190, Houston, TX 77043
713/461-9430 FAX 713/461-7413

CIRCLE NO. 237

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.

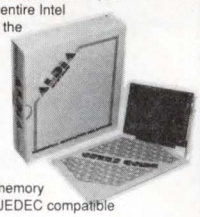

PC-PROMPAK: A PROM/ROM expansion board for IBM and IBM-compatible computers. PC-PROMPAK provides up to 384 Kbytes of expansion memory and supports most 28-pin JEDEC compatible devices.

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. 238

Don't Get Zapped!

High inrush current can **destroy** your sensitive VAX CPUs and peripherals **in less time than it takes to flip a switch.**

THE SOLUTION?
Power up with ***Z-LINE*** **TPC 115-10 MTD™ the smallest** power distribution and control system available.
POWER UP WITH — — —

Z-LINE

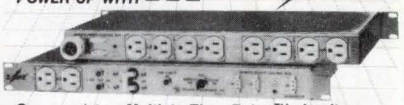

Our proprietary **Multiple Time Delay™ circuitry sequences** your power-up to **protect** your systems from the spikes and surges, EMI & RFI, that **destroy your hardware and erase your data.** And our remote on/off and emergency shutdown **gives the power control back to you.**

All Pulizzi Engineering MTD™ **controllers are compatible with DEC and UPS systems.**
PRICES FROM $436 TO $305

DON'T WAIT UNTIL IT HAPPENS, CALL TODAY!
PULIZZI ENGINEERING INC.
3260 S. Susan Street, Santa Ana, CA 92704-6865
(714) 540-4229 FAX (714) 641-9062

CIRCLE NO. 240

AutoKey,
400 Macro Keys

- Plugs inline with your PC Keyboard
- Uses no system memory.
- Requires no software.
- Has its own microcomputer.
- Eliminates complicated memory resident macro programs such as SuperKey and ProKey. Prices start $139.00

Mextel Corp. 159 Beeline Rd.
Bensenville, Illinois 60106
Phone (312) 595-4146.

M/C VISA AMEX **CIRCLE NO. 241**

Superkey and Prokey are trademarks of their respective manufacturers

9-Track Tape Subsystem for the IBM PC/XT/AT

Now you can exchange data files between your IBM PC and any mainframe or mini-computer using IBM compatible 1600 or 6250 BPI 9-Track tape. System can also be used for disk backup. Transfer rate is up to 4 magabytes per minute on PCs and compatibles. Subsystems include 7" or 10½" streaming tape drive, tape coupler card and DOS compatible software. For more information, call us today!

QUALSTAR®

9621 Irondale Ave., Chatsworth, CA 91311
Telephone: (818) 882-5822

CIRCLE NO. 242

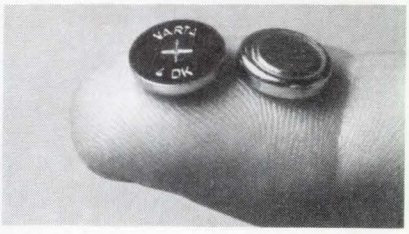

New for micro circuitry: the 4 DK, world's smallest rechargeable cell

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. 243

PROFESSIONAL OPPORTUNITIES

CIRCULATION 100,000 +

Target your recruitment message to system design engineers — with no waste circulation.

If you need engineers and engineering managers who are skilled in the design of computer or microprocessor systems, **COMPUTER DESIGN** *should be the number one choice for your recruitment advertising.*

Published Twice Monthly — Closing Dates Every Two Weeks

ADVERTISER'S INDEX

ADVERTISER	PAGE NO.	CIRCLE NO.
Acromag Inc.	97	52
Advanced Micro Devices Inc.	6,7	5
Aesoft	23	15
Aldia Systems Inc.	111	238
AMP	68,69	39
Apollo Computer Inc.	89-92	47,48,49
Applied Microsystems	33	21
	32A,32B	–
AT&T	64	37
BP Microsystems	111	237
Burr-Brown Corp	18	10
Central Data Corp	82	42
Ciprico	2	3
Communications Specialties	110	238
Compaq Computer Corp.	109	59
Computer Consulting Associates	110	237
Cypress Semiconductor Corp.	Cover-3	–
Data Delay Devices Inc.	46	30
Data I/O	29	19
	110	236
Digi-Data Corp.	102	55
Digital Electronics	99	53
DY-4 Systems Inc.	38	25
	15	9
EDI	63	36
Electronic Solutions	1	2
Emulex Corp	Cover 4	60
Entrepreneurs' Library	105	57
Excelan, Inc.	11	7
Eyring Research	45	29
Heurikon Corporation	78	41
Hewlett Packard	34,35	22
Hitachi America, Ltd.	21	14
Hybricon Corp	110	230
Inmos	43	28
Intel	93	50
Intermetrics	75	40
Iotech	110	234
John Fluke Manufacturing Co.	39	26
Logical Devices Inc.	36	24
	110	231
LSI Logic	31	20
Mentor Graphics Corp.	Cover 2	1
Mextel Corporation	111	241
Micron Technology Inc.	58	35
Mitsubishi Electronics America	40	27
Mizar Inc.	57	56
Motorola Microcomputer Division	48,49	32
National Instruments	110	235
National Semiconductor Corp.	4,5	4
New Media Graphics	18	11
NTH Graphics	19	13
Oasys	26	17
OKI Semiconductor	12,13	8
Ominicomp Graphics Corp.	107	58
Performance Semiconductor	24	16
Personal CAD Systems	27	18
Pulizzi Engineering Inc.	111	240
Qualstar Corp.	111	242
Real Time Devices	110	244
Rogers Corporation	110	232
Sea Level Systems Inc.	36	23
Spectrum Signal Processing	47	31
Summagraphics Corp.	94,95	51
	101	54
Switchcraft, Inc.	55	34
Systech	108	62
Texas Instruments	16,17	12
	16A,16B	–
Tokin	85	44
Toshiba America, Inc.	86,87	45
Truevision, Inc.	50	33
Varta Batteries	111	243
VenturCom	84	43
Vision Security	104	61
VME Microsystems Int'l Co.	67	38
Wind River	8	6

*Domestic Only Issues †International Only Issues

The Advertisers' Index is published as a service. The publisher does not assume liability for errors or omission.

COMPUTER DESIGN

THE FIRST MAGAZINE OF SYSTEM DESIGN, DEVELOPMENT AND INTEGRATION

Executive Office: One Technology Park Drive, P.O. Box 990, Westford, MA 01886
(508) 692-0700 FAX:(508) 692-0525 Telex: 883436

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

Associate Publisher, John C. Miklosz (508) 392-2114
Marketing Services Manager, Mary M. Gregory (508) 392-2110
Administrative Assistant, Tina LePage (508) 396-2112
Advertising Coordinator, Adrienne Adler (508) 392-2198

Postcards/System Showcase/Recruitment, Shirley Athorn (508) 392-2183
Circulation Director, Bob Dromgoole (508) 392-2173
List Rental, Bruce McCarthy (508) 392-2176
Circulation Manager, George Andrew P.O. Box 3466, Tulsa, OK 74101 (918) 831-9443

SALES OFFICES

New England/Upstate New York/ Florida/Eastern Ontario/Quebec
Mary M. Gregory
One Technology Park Drive
P.O. Box 990
Westford, MA 01886
Tel: (508) 392-2110
Fax: (508) 692-0525; Telex 883436

New York/New Jersey/ Pennsylvania/Delaware/ Maryland/North-South Carolina/ Georgia
Neil Versen
Park 80 West, Plaza Two
Saddle Brook, NJ 07662
Tel: (201) 845-0800
Fax: (201) 845-6275

Upper Midwest/Western Ontario/ Manitoba/Saskatchewan
Robert D. Wentz
9501 West Devon, Suite 203
Rosemont, IL 60018
Tel: (312) 518-9191
Fax: (312) 698-2675

South Central
Eric Jeter
19627 I-45 North
Suite 220
Spring, TX 77388
Tel: (713) 353-0309; (713) 288-8963

Southern California/Southwest
Tom Boris, Greg Cruse
2232 SE Bristol, Suite 109
Santa Ana Heights, CA 92707
Tel: (714) 756-0681
Fax: (714) 756-0621

Northern California/Nevada
Tom Boris, John Sly, Bill Cooper
1000 Elwell Court, Suite 234
Palo Alto, CA 94303
Tel: (415) 965-4334
Fax: (415) 965-0255

Oregon
Tom Boris
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) 47 07 59 01

W. Germany/Austria/ N. Switzerland/ Eastern Europe
Johann Bylek
Verlagsbuero Johann Bylek
Stockaeckerring 63
D-8011 Kirchheim/Muenchen
Federal Republic of Germany
Tel: 089 903 88 06
Telex: 529355 vbb d

Italy
Luigi Rancati
Rancati Advertising
Milano San Felice Torre 5
20090 Segrate Italy
Tel: 2 7531445
Telex: 328601 RANCAD I
Fax: 02 7532354

Japan
Sumio Oka
International Media Representatives, Ltd.
2-29 Toranomon 1-chome
Minato-ku, Tokyo 105 Japan
Tel: 03-502-0656
Telex: J22633 MEDIAREP
Fax: 03-591-2530

Southeast Asia
Anne Goh-Taylor
Seavex, Ltd.
400 Orchard Rd.
10-01, Orchard Towers
Singapore 0923
Republic of Singapore
Tel: 734-9790
Telex: 35539 SEAVEX RS
Fax: 732-5129

Jay G. Seo
Yong-Jin Park
Doobee International, Ltd.
Center Building (Byulgwan)
1-11 Jeong-dong, Choong-ku
CPO Box 4557
Seoul, Korea
Tel: 776-2096
Telex: K27117 DOOBEES
Fax: 755-9860

COMPUTER DESIGN (ISSN-0010-4566) is published twice monthly, by PennWell Publishing Company, 1421 S. Sheridan, Tulsa, OK 74112. Second class postage paid at Tulsa, OK 74112 and additional mailing offices. Editorial offices are located at One Technology Park Drive, P.O. Box 990, Westford, MA 01886. Annual subscription price: $70 in the U.S.A., $110 elsewhere. Call (918) 831-9401 for subscription information. Microfilm copies of COMPUTER DESIGN may be purchased from University Microfilms, a Xerox Compnay, 300 North Zeeb Rd., Ann Arbor, MI 48106. POSTMASTER: Send change of address form to COMPUTER DESIGN, Circulation Department, PO Box 3466, Tulsa, OK 74101.

© 1989 COMPUTER DESIGN by PennWell Publishing Company. All rights reserved. No material may be reprinted without permission from the publisher. Officers of PennWell Publishing Company: Philip C. Lauinger, Jr., Chairman and Chief Executive; Joseph A. Wolking, President; John Ford, Senior Vice-President; Carl J. Lawrence, Senior Vice-President; John Maney, Vice-President/Finance; Leslie P. Cypret, Vice-President/Administration.

BPA ABP

Make a quick killing in real estate.

Introducing our new high density, *highest performance* Multichip Modules.

Available now in a variety of SRAM configurations. Or choose *custom*. We'll work with you on defining mechanical and electrical requirements, as well as solve design, test and assembly issues.

Either way, we deliver fast, *tested* CMOS SRAM modules that can save you a fortune in real estate, design, and manufacturing, while speeding up system performance.

Our new, Cypress data book has one of the world's fastest multiple listings. Call today. It's free.

Multichip Hotline:
1-800-952-6300*
Ask for Dept. C117

CYPRESS SEMICONDUCTOR

70ns
CYM1461, 512K x 8 SRAM Module

25ns
CYM1821, 16K x 32 SRAM Module

25ns
CYM1611, 16K x 16 SRAM Module

WILL BUILD TO SUIT

25ns
CYM1621, 64K x 16 SRAM Module

25ns
CYM1610, 16K x 16 SRAM Module

*1-800-387-7599 In Canada. (32)2-672-2220 In Europe. © 1989 Cypress Semiconductor, 3901 North First Street, San Jose, CA 95134, Phone: (408) 943-2666, Telex 821032 CYPRESS SNJUD, TWX 910-997-0753.

EMULEX ALREADY HAS THE ONE THING VME SYSTEMS NEED FOR SMOOTH SAILING...

MORE

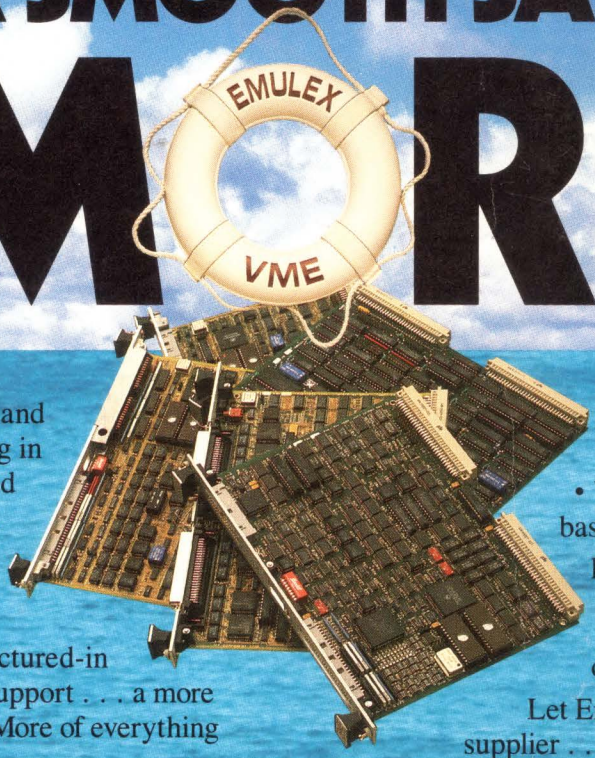

To assure system designers and integrators of smooth sailing in the technologically advanced VME marketplace, Emulex has more of what's needed most. More designed-in quality . . . more VLSI expertise . . . more manufactured-in reliability . . . more OEM support . . . a more dependable supply source. More of everything that's important.

And our "MORE" philosophy applies to every member of our new family of products for the VMEbus.

Introducing:

- VT01 – a 9-track Pertec tape formatter controlling up to eight 1/2-inch tape drives. Tape densities can range from 800 to 6250bpi, tape speeds to 200ips.
- VM21 – a disk controller that can handle four ESDI disk drives with the highest transfer rates available – up to 20Mbps.
- VM31 – a disk controller for four SMD-E disk drives with transfer rates up to 3.0MB/sec and no penalty on performance.
- VS02 – a communications controller designed around a proprietary I/O processor that can handle 16 asynchronous lines at 38.4Kbps simultaneously.
- VH01 – a SCSI host adapter based upon Emulex's own SCSI processor chip that supports synchronous transfer rates of 4.8MB/sec and asynchronous rates of 3.0MB/sec.

Let Emulex be more than your OEM supplier . . . think of us as your OEM partner. That means you can count on us for product quality and reliability, and for stability, capability and integrity. From initial design to final audit, Emulex is dedicated to meeting your needs . . . and MORE.

VMEbus PERFORMANCE CERTIFIED VME laboratories

Call toll-free (800) EMULEX-3, (714) 662-5600 in California.

EMULEX

I want to know MORE about VME products.

Name: _____

Company: _____

Address: _____

City: _____ State: _____

Zip: _____ Phone: _____

Emulex Corporation, 3545 Harbor Blvd., Costa Mesa, CA 92626

Regional Offices: Anaheim, CA (714) 385-1685 Atlanta, GA (404) 587-3610 Burlington, MA (617) 229-8880
Chicago, IL (312) 490-0050 Dublin, CA (415) 829-1170 Teaneck, NJ (201) 836-3717 Washington, DC (703) 264-0670
International Offices: Bracknell 344-484234 Munich 89-3608020 North Sydney 2-957-1669 Paris 134-65-9191 Tokyo 3-234-8951 Toronto (416) 673-1211

©1988 Emulex Corporation

CIRCLE NO. 60