A PENNWELL PUBLICATION DECEMBER 1992

COMPUTER DESIGN

Technology and Design Directions

$5 USA
$6 Canadian
®

FOR ELECTRONIC ENGINEERS & ENGINEERING MANAGERS

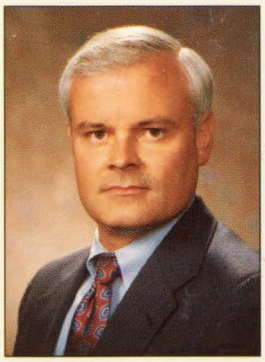

Data I/O's Tom Clark on: FPGA design

32-bit power and tools bring cheer to embedded system designers

FPGA vendors turn their attention to tools

GUIs move OSs toward object orientation

12
0 74470 80522 9

Ciprico Rimfire™ Adapters change the appearance of unique peripherals...

just about that easily.

Incorporate Rimfire Series SCSI Adapters into your system and unique non-SCSI devices will take on a new look. Your system will see them as one of the ten types of SCSI devices defined in the SCSI-2 standard.

Adapters integrate easily into host-to-peripheral and host-to-host operations, and support SCSI data transfers up to 20 MB/second. With our working-sample source code and documentation, you can put adapters to work readily in EISA, ISA, Micro Channel, and VMEbus environments.

Based on our 6-years experience in developing high performance SCSI/SCSI-2 adapters, Rimfire Adapters offer the most flexible "target" mode implementation available. And that's just one of the reasons there are so many satisfied Rimfire customers throughout the world.

If you want to know more of the reasons, contact us for details, plus a free ***Target Mode Technology Brief.***

CIPRICO

1-800-SCSI-NOW

2800 Campus Dr., Plymouth, MN 55441
In MN: (612) 551-4000

In Europe:
7 Clerewater Place, Lower Way
Newbury, Berkshire RG13 4RF
England (44) 635 8 73666

High performance solutions for mass storage management.

CIRCLE NO. 1

One Thing Separates a Great System Enclosure from a Good One.

It's the innovative application of cutting-edge technology to your specific requirements. It's what we do at Electronic Solutions.

If you're the engineer responsible for the system packaging, you have your hands full dealing with power distribution, cooling, and safety concerns such as UL, CSA and VDE approvals. That's where Electronic Solutions can make your job easier! You can rely on our experience and expertise to quickly resolve the design issues you're facing – at a price that even your bean counters will be pleased with.

Airflow testing. Our *in-house* airflow/temperature testing facility allows us to carefully examine the environment *inside* a system enclosure, testing the effects of wire harness placement, fan size and location, and the characteristics of the airflow path itself. We can even test your complete system to fine tune air direction, cool hot spots, and reduce audible fan noise and power consumption.

Safety Agency Approvals. Bypass the costly maze of paperwork and testing! We can provide enclosures that are already UL listed, CSA certified, and GS approved through TUV certification. That big step allows you to install recognized boards and peripherals and win agency approval of your finished system with no additional effort or cost! As a self-testing facility for the agencies, we're equipped to handle every aspect of certification to reduce your costs and shorten your time to market.

EMI/RFI Testing. We subject our enclosures to *in-house* testing per MIL-STD-285, providing an attenuation profile of each unit. When it comes time for you to pursue FCC or VDE approval of your system, we can perform these tests for you and resolve any problems immediately.

Whether it's off-the-shelf or something special that's off-the-drawing board, there's only one way to get the enclosure that's just right for your system. That's to call Electronic Solutions and put our experience to work for you.

VMEbus
VXI bus
Sun bus
Mixed buses
Proprietary buses
Multibus I
Multibus II
Futurebus+

Tabletop,
Rackmount,
DeskMate,
Bellcore and
Custom
Designs

ZERO **Electronic Solutions**
UNIT OF ZERO CORPORATION

6790 Flanders Drive, San Diego, CA 92121
(619) 452-9333 FAX: 619-452-9464
Call Toll Free: (800) 854-7086
In Calif: **(800) 772-7086**

CIRCLE NO. 2

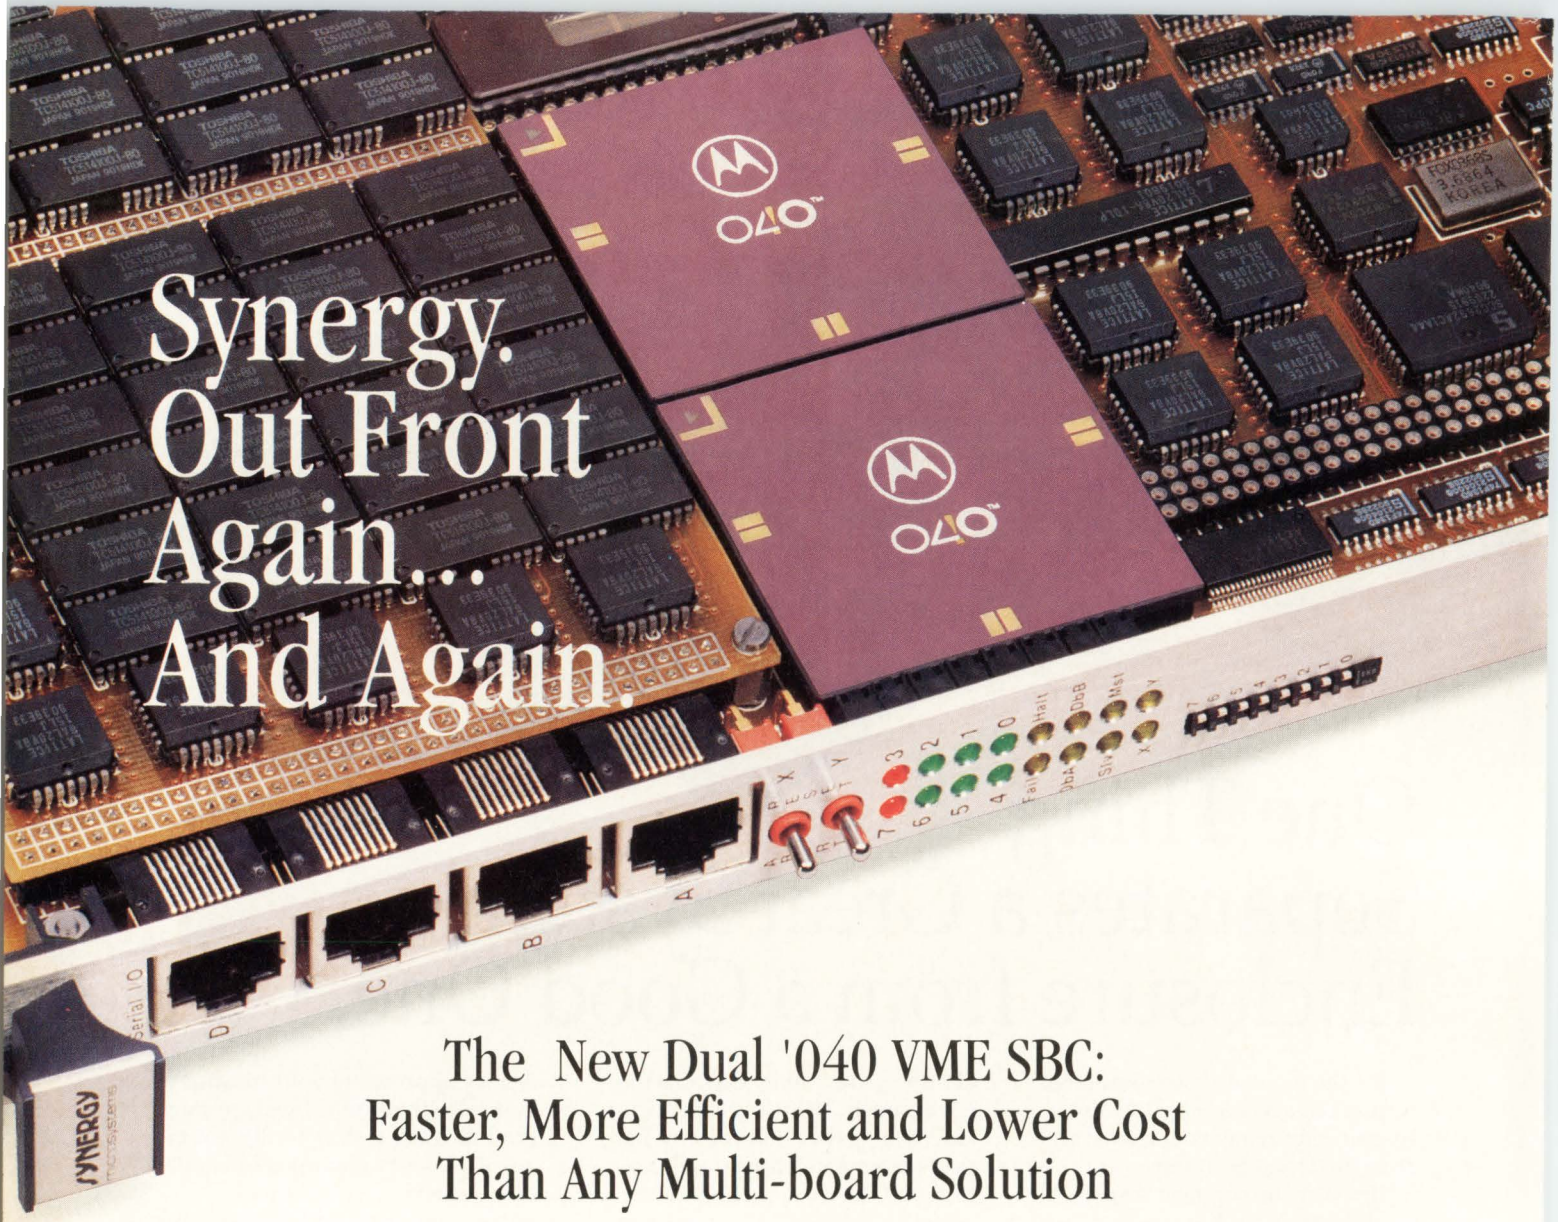

Synergy. Out Front Again... And Again.

The New Dual '040 VME SBC: Faster, More Efficient and Lower Cost Than Any Multi-board Solution

Get on-board multiprocessing and an astounding **40 MIPS** throughput, when you power your system with Synergy's new SV420 single-slot SBC. The SV420's dual-CPU design means you'll need fewer boards in your VME chassis, with no VMEbus bandwidth bottlenecks between the '040s, and system-wide cost saving of more than 40%.

And even if you don't need multiprocessing right now, the SV420 still puts you out front. Use the second '040 as a super-smart DMA controller. When combined with the SV420's **66 MByte/sec** VME64® circuitry, nothing communicates faster over the VMEbus.

Add even more on-board power by selecting from Synergy's big list of high-performance, intelligent plug-on I/O modules – such as our latest Super-VSB module offering 40+ MB/sec over the VSB bus. Or choose another module, from a T1/E1 controller to a super-fast graphics engine, that plugs onto any Synergy SBC.

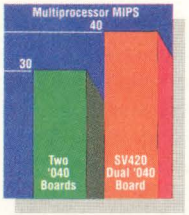

Better yet, just tell us what you need. We're the company you can talk to about your VME system design problems. You'll find that we listen *and deliver* (every Synergy I/O module on our list started as a customer request). We mean business when we say customer support is our most important mission.

So if you want to be out front in system performance, just be up front with Synergy. Call us today.

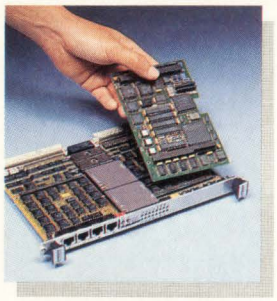

Over a dozen smart, powerful I/O modules fit Synergy's '020, '030, '040 and Dual '040 SBC's.

SYNERGY microsystems

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

CIRCLE NO. 3

A PENNWELL PUBLICATION VOL 31, NO 12 / DECEMBER 1992

COMPUTER DESIGN

Technology and Design Directions

FOR ELECTRONIC ENGINEERS & ENGINEERING MANAGERS

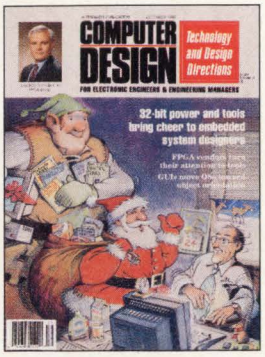

Designers of embedded systems can enjoy increased compute power, higher integration and high-level programming languages if they wish for, and can use, 32-bit microcontrollers in their next-generation designs..... **91**

COMING NEXT MONTH

Verilog and VHDL

DSP development tools

Disk drive controllers and ICs

C cross compilers

Networking

NEWS BRIEFS

TECHNOLOGY DIRECTIONS

Integrated Circuits

Software & Development Tools

Computers & Subsystems

CAE/CAD Tools

ASICs & ASIC Design Tools

NEW PRODUCT DEVELOPMENTS

CAE/CAD Design Tools

Software & Development Tools

Computers & Subsystems

Page 62

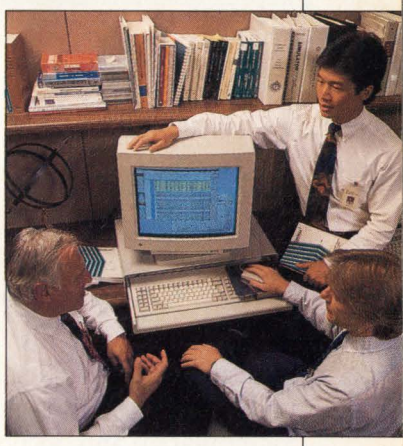

Page 70

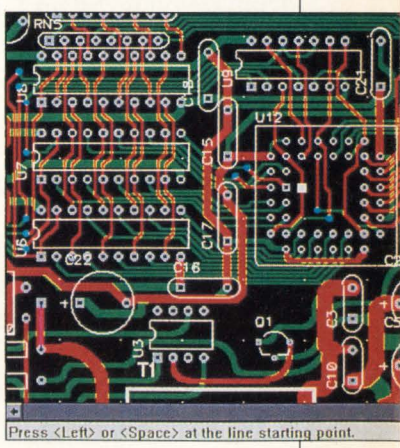

Page 123

CONTINUED ON PAGE 5

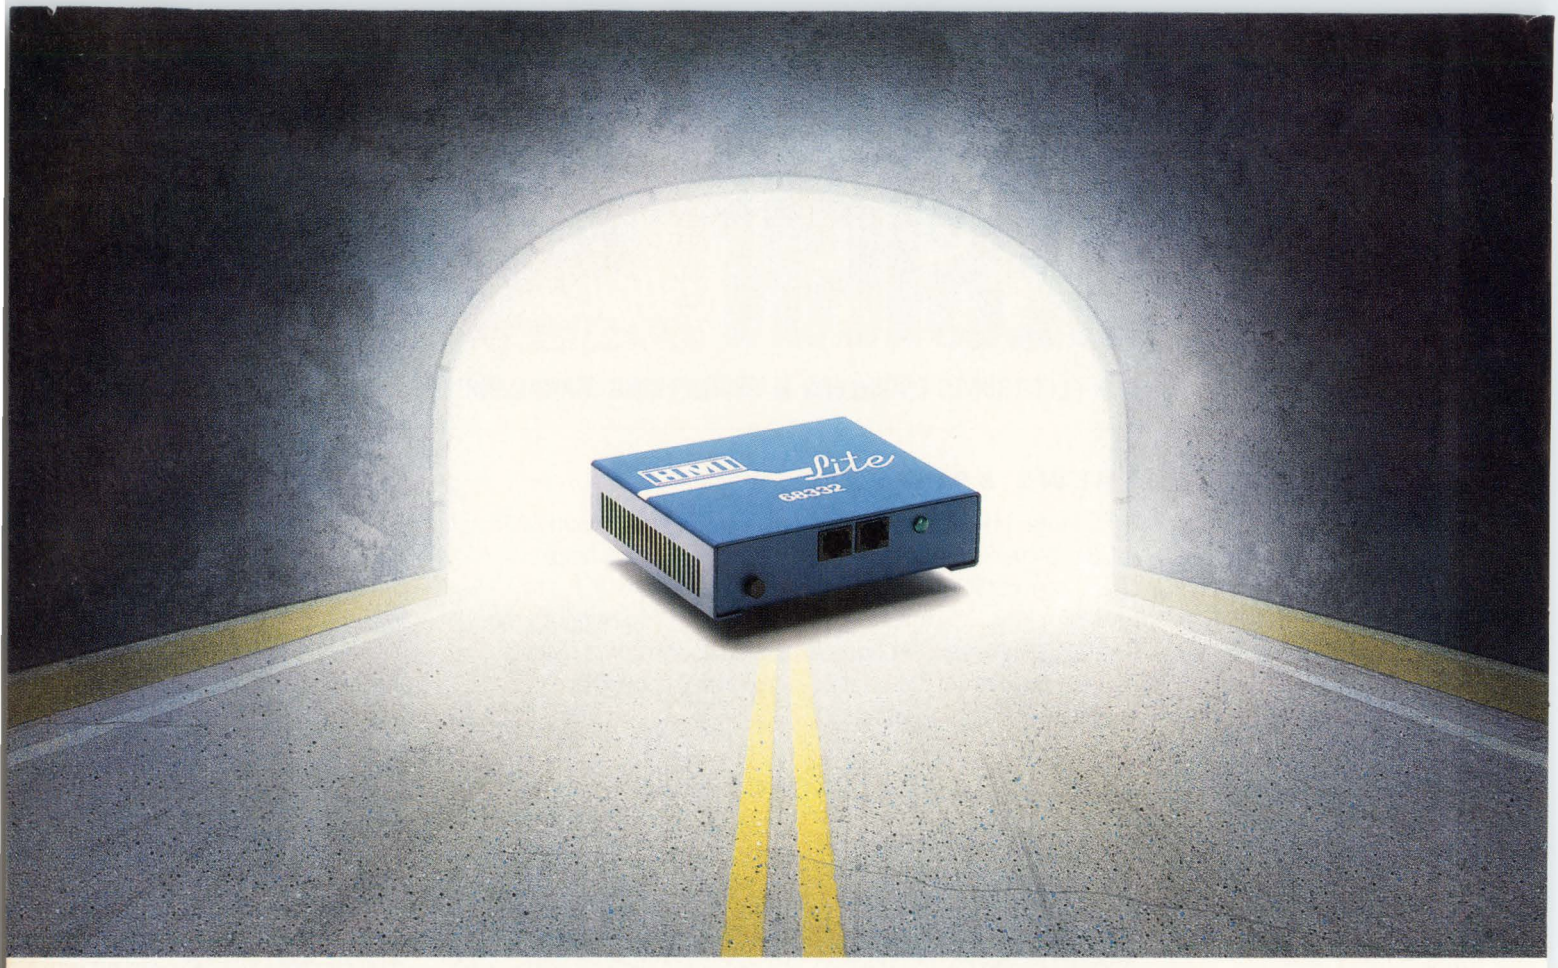

Seeking an affordable, high-performance development system? Our Lite's at the end of the tunnel.

It's traditionally presented a pretty dark scenario, the quest for emulation and debugging tools that are integrated and proven. That are powerful and, perhaps most elusive, highly affordable.

No more. At under $5,000, HMI's new Lite development system, integrating real-time emulation with our renowned SourceGate high-level-language debugger, offers a brilliant end to the search. And lights up the industry with a new price/performance standard.

Field proven by the thousands, SourceGate supports all major industry compilers: C, PASCAL and Ada. And the task is streamlined by menus, customizable windows, single stepping, and variable watch windows all controlled at the source code level.

No dead-ends in emulators, either; the Lite series features multiple hardware breakpoints, full trace capabilities and dual ported RAM for real-time monitor of critical variables. And, you'll ensure a smooth migration path to our premier Series 200 systems which feature the same SourceGate user interface.

Better write or phone for details on our new Lite. Ready for use with IBM PC family and UNIX based computers, it really is just what you've been looking for.

HMI

HUNTSVILLE MICROSYSTEMS, INC.
3322 South Memorial Parkway
Huntsville, AL 35801 USA.
Tel: (205) 881-6005
Fax: (205) 882-6701

IBM is reg. T.M., International Business Machines, Inc.
UNIX is reg. T.M., Bell Laboratories, Inc.

Supporting 683xx and 68HC16 families. Call for specific microprocessors.

CIRCLE NO. 4

COMPUTER DESIGN Technology and Design Directions

CONTINUED FROM PAGE 3

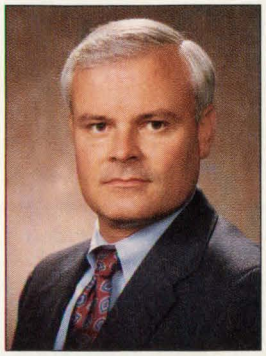

Tom Clark on: FPGA design

Logic designers using field programmable logic devices have two basic choices when it comes to design methodology — structural and behavioral. **31**

Look for

OEM INTEGRATION

next month, featuring secondary and mezzanine buses, backplanes and enclosures.

TECHNOLOGY & DESIGN REPORTS

FPGA vendors turn their attention to tools

In the face of stiff competition, FPGA vendors are enhancing proprietary tools, enlisting third-party support, adding text-based entry methods, and backing standards.
— *Barbara Tuck* **75**

GUIs move OSs toward object orientation

Computer graphics, originally meant to display data, is now the way users interact with systems. A new generation of operating systems with fully integrated GUIs is meeting the demands of those who want to work with their systems in a real-world environment of objects and actions.
— *Tom Williams* **85**

COVER STORY

32-bit power and tools bring cheer to embedded system designers

Because of their increased compute power, higher integration, extensive tool sets, and a desire for the friendliness provided by high-level languages, more designers are putting 32-bit microcontrollers on their wish lists for next-generation products.
— *Don Tuite* **91**

DESIGN STRATEGIES

Truck simulator integrates off-the-shelf subsystems

A defense contractor saw an opportunity to create its first commercial product, the TTI50 Truck Driving Simulator.
— *Jeffrey Child* **105**

PRODUCT FOCUS

STD Bus CPUs focus on solutions

Unlike VME, which has Motorola supporting it, or Multibus II, which is backed by Intel, STD Bus may suffer from not having a major semiconductor manufacturer behind it. But in reality, STD has a strong position as a low-cost workhorse bus for embedded control. — *Jeffrey Child* **115**

COLUMN

MIXED-SIGNAL DESIGN — *Stephan Ohr*

"Design for X:" A new push for manufacturability, testability and reliability **132**

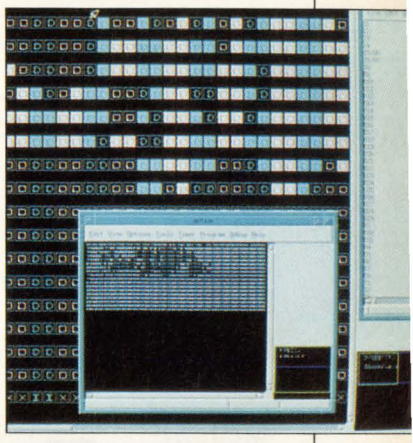

Page 75

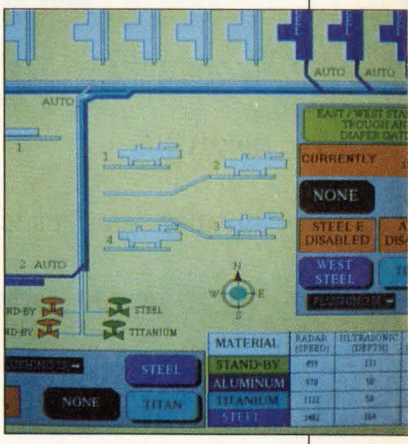

Page 85

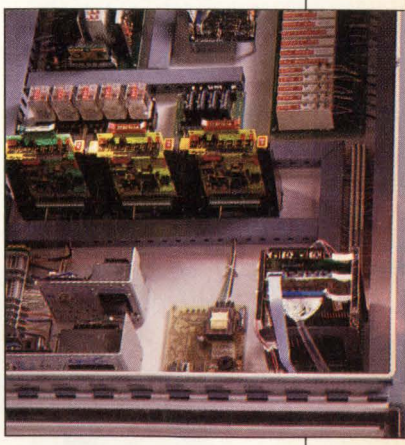

Page 115

No logic analyzer is, for that matter.

Because at speeds above 25 MHz, even the best designer needs a scope to handle nightmares like race conditions, ground bounce, and crosstalk. And not just any scope, but one tailored specifically for high-speed digital design.

At Tektronix, we understand this need all too well. As proof, we've not only designed and built a great logic analyzer — our powerful new GPX —

EVEN THOUGH OUR NEW LOGIC HIGH SPEED DIGITAL DESIGN,

but also a perfect companion. The TDS 640 digitizing oscilloscope.

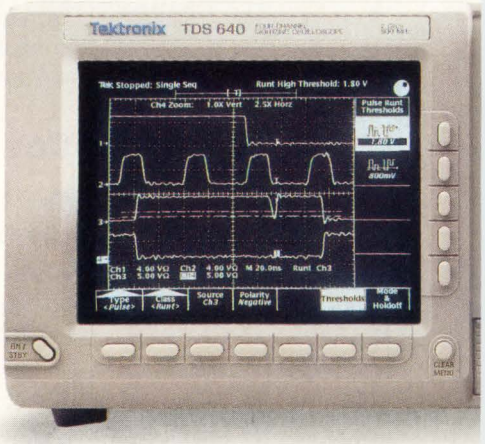

To get you started, the GPX provides more channels than cable TV. Up to 160, for instance, of 80 MHz state analysis, and 32 channels of 1 GHz timing or 160 channels of 200 MHz transitional timing.

In short, enough to handle the world's fastest microprocessors. And with that kind of performance, you can easily track your system at clock rates well beyond 50 MHz, which allows you to locate complex coding errors quickly and accurately.

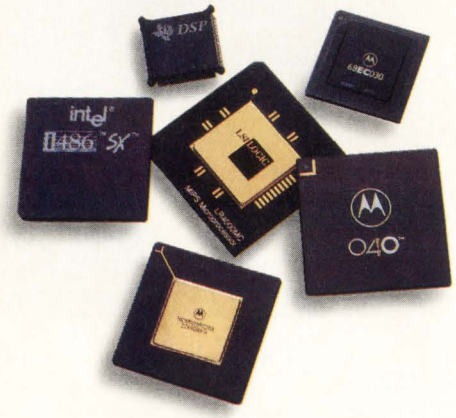

Enter the TDS 640.

With a 500 MHz bandwidth and 2 GS/s real-time sampling on four channels, the TDS displays logic and timing errors with absolute accuracy. And because it was created with the digital designer in mind, the TDS lets you trigger directly on common digital circuit events like glitches, runt pulses, and excessive clock jitter or skew. Put all that together with the power of the GPX and you have a remarkably effective solution.

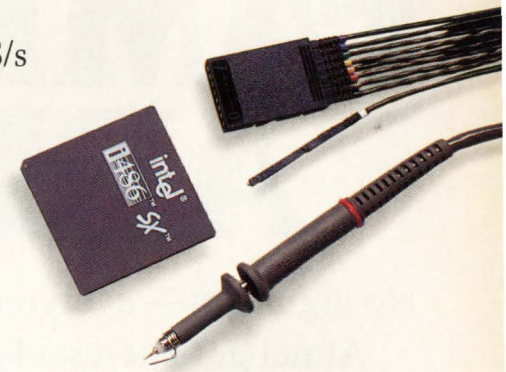

ANALYZER IS GREAT FOR IT'S NOT NEARLY ENOUGH.

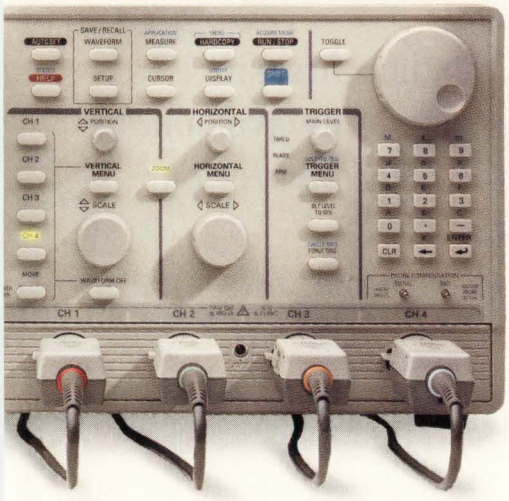

Cost effective, too. In fact, the GPX and TDS together sell for less than competing scope/logic analyzer combinations. And if you buy the pair between now and February 28, 1993, we'll take an additional 10% off the retail price.

Enough, already.

For more information on the very best high-speed digital design tools available, call Tektronix today at 800-426-2200.

Tektronix
Test and Measurement

CIRCLE NO. 5

TODAY'S FAMILY VALUES ARE 3 VOLTS.

No argument—everyone loves a good 3-volt system.

Atmel got interested in 3-volt components when we saw the pains our customers were having with their growing battery needs. Too many batteries hurt portability and pocketability. It was time for new devices with new values.

They're here. We make the world's largest collection of low-voltage logic and memory devices. We'll cut your battery count or give added life to the ones you keep.

We've got the world's only 3-volt Flash and 3-volt 22V10. A bunch of 3-volt EPROMs, 15 parallel and serial interface EEPROMs (most of which operate to 1.8 volts) and two full 3-volt gate array families. That's seven families and more than 50 circuits.

You know us already. Maybe because of our industry-leading first single-voltage Flash, or our serial interface EEPROM's (the largest family available). Or maybe because of the blazing speed of our non-volatile memories, or our families of user- and factory-programmable logic devices.

Now know our low-voltage families. Check out the book. Write, fax or call us and we'll send you our low-voltage catalog.

And, when the talk turns to family values remember ours are the same as yours—3 volts.

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

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


3
VOLTS
3
VOLTS
GATE ARRAYS
EPROM
EEPROM
EPLD
FLASH
SERIALS
SRAM
CUSACK
3
VOLTS
3
VOLTS

MIPS forges ahead

Last month, MIPS Technologies (Mountain View, CA) announced the R4400, the latest 64-bit MIPS microprocessor. Designers experienced with the R4000 won't have to make any changes in hardware or software when using the new R4400. The only differences are cache size, clock rate and a new write buffer. This buffer takes the output from a graphics loop and runs it in parallel with the next loop, resulting in a substantial increase in graphics performance.

Although the cache size on the R4400 is double that of its predecessor, a 20 percent shrink in process technology has returned that real estate, resulting in a die size no larger than the R4000. The clock rate has been raised to 75 MHz, with 67 MHz and 50 MHz for backward compatibility. Transistor count went up by a million on the R4400, with the increase fundamentally in cache. On-chip primary cache has been doubled from an 8k/8k instruction/data cache to a 16k/16k cache on the R4400. The secondary is the same. The processor is available in 3.3- and 5-V versions.

How does this new chip fit into the emerging RISC PC world? Even when Intel's P5 finally hits the streets, the MIPS crowd already has the more powerful R4000—and the R4400 is an upgrade to that. When compared with the 486, the R4400 offers more performance in the same price range.

—Jeffrey Child

Chorus spreads microkernel-based UNIX

Chorus Systemes (Paris, France) appears to be the most agreeable operating system company around. Last year Chorus concluded arrangements with ISI Software Components Group (Santa Clara, CA) to link its microkernel-based UNIX-compatible distributed operating system with ISI's pSOS+ realtime kernel. Recently, Chorus moved to support SCO's (Santa Cruz, CA) PC-based UNIX.

Now Chorus has entered into an agreement with UNIX System Labs (Summit, NJ), the guardian of UNIX System V Release 4 (SVR4). Its aim is to let Chorus microkernel technology evolve in step with SVR4, so that large system and realtime system vendors have an SVR4-compatible microkernel migration path for future development. Topping this off is an agreement with Tandem Computers (Cupertino, CA) to develop microkernel-based SVR4 fault-tolerant operating system technology. According to Chorus, such technology will be scalable from embedded realtime systems to large mainframe computers.

—Tom Williams

Benchmarks proposed for fuzzy logic

A suite of benchmark programs developed by Togai Infralogic (Irvine, CA) has been proposed as a means of measuring the performance of processors executing fuzzy logic inference code. The benchmarks are three fuzzy rule bases at different levels of complexity: simple (with seven rules, each having two input variables and one output variable); medium (14 rules of three inputs and two outputs); and complex (25 rules of seven inputs and three outputs).

Togai has released results of tests run on four processors, the Motorola 68HC11, Hitachi H8/300 and -500 and Intel 8051 (which showed the highest performance). There will no doubt be many questions from vendors as to how the code was produced, what inference methods were used and whether code was optimized. When this information is available, vendors may be able to use these benchmarks as a starting point for a common suite that will help designers pick price-performance points for fuzzy-based designs.

—Tom Williams

ViaLink helps QuickLogic cut FPGA prices

QuickLogic's (Santa Clara, CA) ViaLink process technology has been moved from codeveloper VLSI Technology's (San Jose, CA) pilot line to its high-volume production facility in San Antonio, TX, producing improved yield and lower wafer cost that's enabled the FPGA vendor to reduce pASIC 1 prices by up to 33 percent. "As we move to high-volume production of our products, we are pleased to be able to pass the cost savings on to our customers," said David A. Laws, QuickLogic's president.

VLSI Technology recently announced that it will use the ViaLink element as the basis for what it calls programmable Functional System Blocks (pFSBs), cells that provide field-programmable capabilities embedded in ASIC devices. These pFSBs will be used to develop embedded memory elements (vROMs or ViaLink ROMs), embedded logic elements and custom ViaLink products. ASICs built with pFSBs will offer visual and electrical security.

Don Ciffone, vice-president and general manager of VLSI's product divisions, says, "With the ability to include field-programmable structures directly on ASIC and ASSP chips, customers now have unprecedented flexibility in the design of secure, high-performance systems that can be highly differentiated from those of their competitors."

—Barbara Tuck

Vitesse prices GaAs ASICs to beat BiCMOS

With 1,500, 7,000 and 13,000 gates, the new 0.6-µm VIPER GaAs gate arrays from Vitesse Semiconductor (Camarillo, CA) are sized for the majority of designs today and offer two to three times the performance of competing BiCMOS arrays at comparable cost, claims the company. "With the introduction of our VIPER arrays, cost is no longer the issue when comparing GaAs with competing BiCMOS technologies," says Lou Tomasetta, Vitesse president and CEO. "Our H-GaAs technology is more mature than most BiCMOS processes currently offered. The real issues today are reducing design time through simplified system architectures and increased performance margin."

Housed in plastic and aimed at system designs running at 50 MHz and above, the VIPER gate arrays deliver shorter gate delays, lower power requirements and better design margins than BiCMOS. According to Vitesse, a two-input NOR gate has a typical unloaded delay of 60 ps, while dissipating only 0.18 mW. "For applications above 75 MHz, there

Continued on page 12

Simultaneous Access

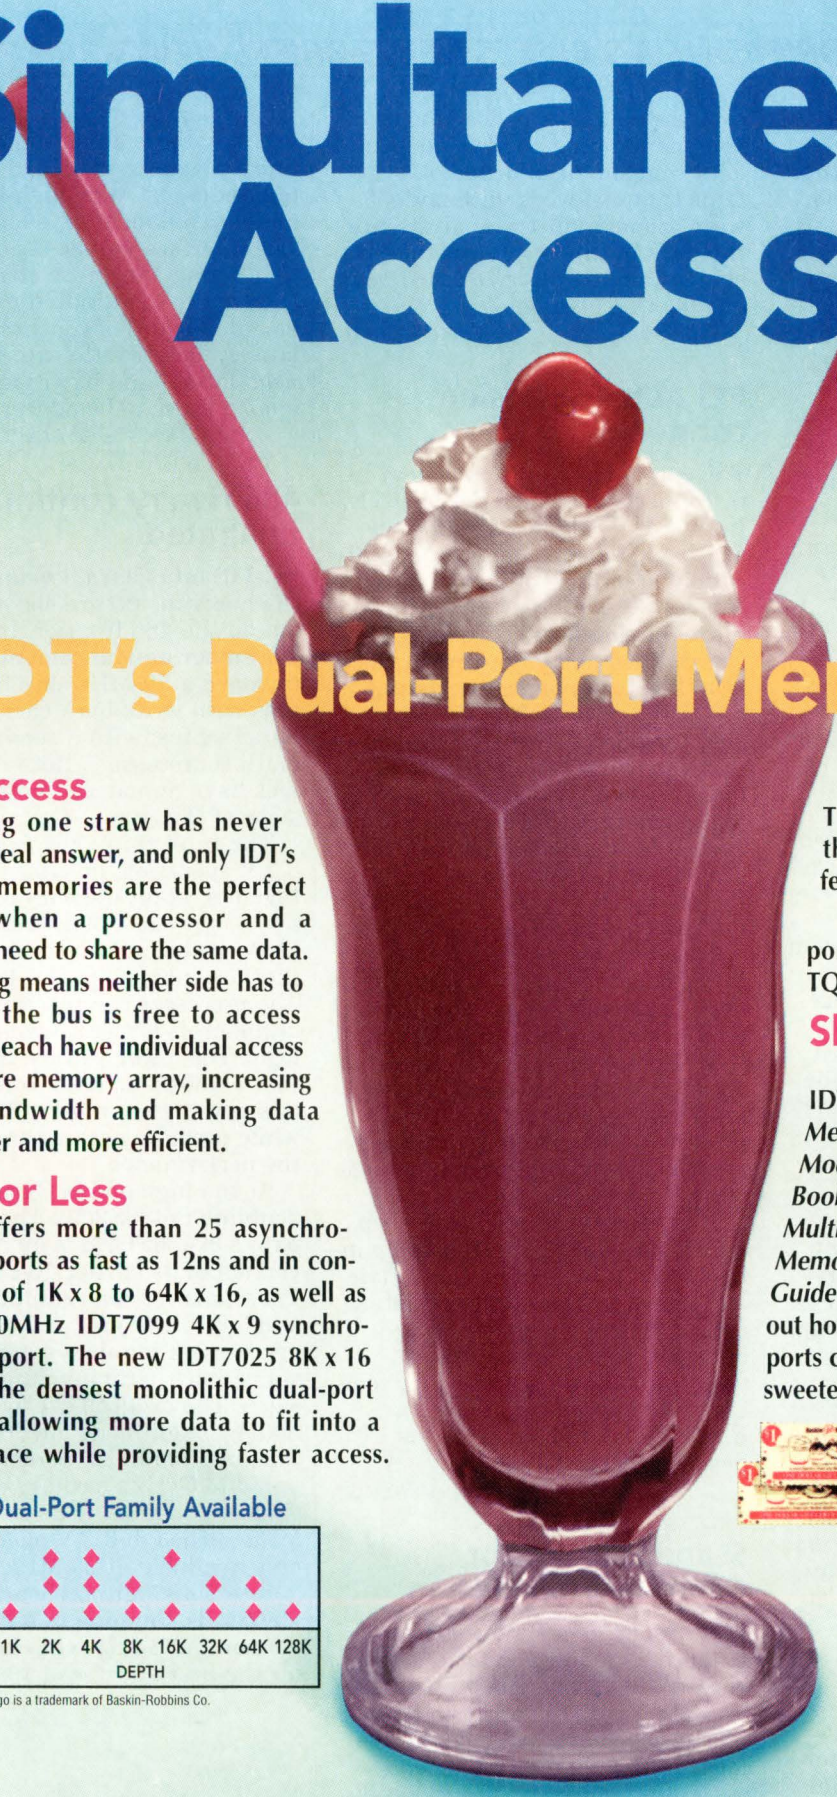

IDT's Dual-Port Memories

Fast Access

Sharing one straw has never been the ideal answer, and only IDT's dual-port memories are the perfect solution when a processor and a peripheral need to share the same data. True sharing means neither side has to wait until the bus is free to access data—they each have individual access to the entire memory array, increasing system bandwidth and making data access faster and more efficient.

More for Less

IDT offers more than 25 asynchronous dual-ports as fast as 12ns and in configurations of 1K x 8 to 64K x 16, as well as the new 50MHz IDT7099 4K x 9 synchronous dual-port. The new IDT7025 8K x 16 device is the densest monolithic dual-port available, allowing more data to fit into a smaller space while providing faster access.

Broadest Dual-Port Family Available

This translates into increased throughput for wider buses using fewer parts!

IDT also offers innovative dual-port packaging solutions, including TQFPs and modules.

Share a Sweet Deal

Call today and get a copy of IDT's new *Specialized Memories and Modules Data Book* and *Multi-Port Memory Design Guide* to find out how dual-ports can simplify your design! We'll sweeten the deal by sending you a pair of *free* Baskin 31 Robbins™ *coupons so you can share your next ice cream treat.*

(Limit 2 per customer.)

The Baskin-Robbins logo is a trademark of Baskin-Robbins Co.

(800) 345-7015 • FAX: 408-492-8674

ASK FOR KIT CODE 7061

CIRCLE NO. 7

Integrated Device Technology, Inc.

Continued from page 10

are no other ASIC technologies that can compete with the price-performance of VIPER," boasts Bob Nunn, vice-president and general manager of ASIC products at Vitesse.

—Barbara Tuck

A kinder, gentler EDIF?

The Electronic Design Interchange Format (EDIF) committee recently unveiled version 2.9.0 of the beleaguered EDIF standard, in hopes that the updated release will rectify the inadequacies of previous versions. For the past four years, the former release, EDIF 2.0.0, has been the primary means of exchanging data among CAD, CAE and test tools, but ambiguities in the standard's syntax have resulted in incompatibility among so-called "EDIF-standard" tools.

"Release 2.9.0 is the result of years of engineering staff hours and millions of dollars," says Rich Goldman, engineering manager of the semiconductor vendor program at Synopsys (Mountain View, CA) and chair of the EDIF technical committee. "We've gone to great lengths to clarify the EDIF syntax to make life easier for the EDIF reader. This is where the majority of the problems existed in version 2.0.0."

—Mike Donlin

Modeling information comes online

Electronic transfer of component modeling information may yet become a reality, thanks to a distribution agreement between start-up ViewPoint Information Systems (Waltham, MA) and Mentor Graphics (Wilsonville, OR). The new company's offering is different from other component information products because its data is machine-readable. Users can extract symbols and attribute information for schematic, simulation and layout programs. ViewPoint has also signed deals with Hitachi (Brisbane, CA), Intel (Hillsboro, OR) and National Semiconductor (Santa Clara, CA); these agreements should give the company component data faster than other vendors who rely on databook information.

The alliances will come as welcome news to EDA users, who've been demanding vendor-independent component information systems that can work with EDA software. Some industry analysts cite this lack of timely component data as a stumbling block to concurrent engineering strategies.

—Mike Donlin

PCI gets expansion connector

The PCI (Peripheral Component Interconnect) definition developed by Intel (Santa Clara, CA) will soon get an expansion connector specification, according to the PCI Special Interest Group Steering Committee. Introduced in June, PCI was defined as a high-performance local bus to supplement existing bus architectures. The definition only provided for an electrical specification, with the belief that PC makers would be soldering such high-performance peripheral chips as graphics and communications directly on a motherboard using the PCI specification for electrical interconnection.

The proposed connector uses a Micro Channel-style edge-card connection providing 32- and 64-bit interconnection. The configuration was selected to keep cost down while providing a high-reliability, high-density connection. The approach will let PCI boards be used in EISA, ISA and Micro Channel systems.

PCI's acceptance will be further enhanced by a PCI interface chip introduced by Intel (Folsom, CA) late last month. Compatible with standard I/O buses such as ISA, EISA and MCA, as well as the new PCMCIA standard, PCI will offer commercial and industrial users a new high-speed conduit to PC-based processors.

—Warren Andrews

Standards set for memory interface

While some vendors pursue PCI and others follow PCMCIA, a group within the IEEE Computer Society is dealing with advances in technology that have made it possible for traditional storage elements such as disk drives to be reduced in size so they can be directly soldered to PC boards.

The Computer Society's P1285 is a standards activity meant to define a new IEEE standard interface to handle just such high-latency, non-volatile memory elements. The interface will be used with either a single memory element or with many coordinated memory elements. Issues of concurrency, latency, bandwidth, extensibility, negotiation, and partitioning are among those to be addressed.

—Warren Andrews

486 rivalry continues unabated

Intel (Santa Clara, CA) and Cyrix (Richardson, TX) are shooting it out on the 486 frontier. In a November announcement, Cyrix unveiled a 50-MHz chip for desktops with write-back caching and burst writes (with a separate math coprocessor). Both companies have announced notebook versions of the chip.

The Intel chip's key features are 3.3-V operation, on-chip integration of a 32-bit memory controller, an ISA bus controller, and the 8-kbyte cache and math coprocessor that are integral to the 486DX. The processor can interface to 5-V peripherals without translation logic. The lower voltage means that a 486SL actually consumes less power than a 5-V 386SL, while providing more than twice the performance.

At the high end is Cyrix's clock-doubling cx486S2/50. The company's cx486SLC/e is an enhanced version of its earlier 16-bit chip. There are versions running at 5 V and 3.3 V, although the low-voltage version requires external translators in dual-voltage designs. The cx486SLC/e system management mode lowers power consumption by 25 percent, compared to Cyrix's earlier version. Chips are sampling at 25 and 33 MHz, although 3.3-V operation is only available at 25 MHz.

Intel's 486 for notebooks, the i486SL, is shipping in a 25-MHz version. A 33-MHz chip is slated for the first quarter of 1993.

—Don Tuite

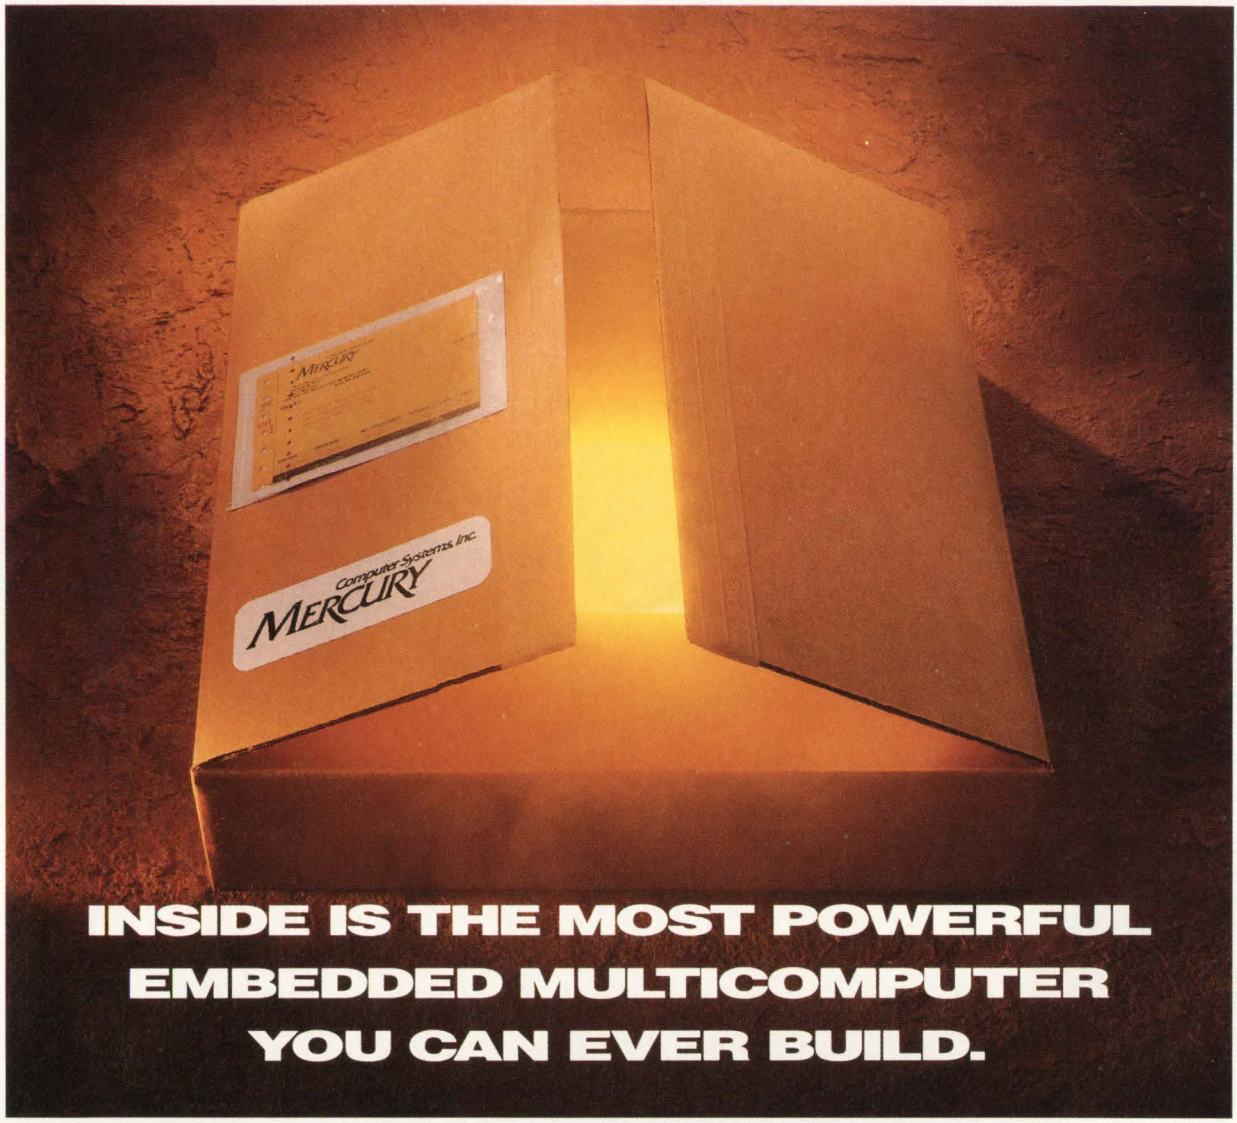

INSIDE IS THE MOST POWERFUL EMBEDDED MULTICOMPUTER YOU CAN EVER BUILD.

Open up to Mercury's embedded realtime multicomputer solutions. Because of a balanced architecture, our multicomputers deliver scalable performance into the tens of GFLOPS, with GBytes of I/O to match. With interboard and interchassis communications, you can build a solution to satisfy even the most demanding applications within your constraints of cost, space, and power consumption.

Fiber Optic I/O
VME64
Interconnect Crossbar
External I/O

Mercury's commercial and ruggedized multicomputers are supported by our standards-based realtime operating system software, including advanced multicomputing development tools that enable you to harness the power of hundreds of processors.

For more information on our commercial off-the-shelf (COTS) products and to receive our White Paper: **Balancing I/O and Computation, call Mercury toll free today at 1-800-229-2006 or (508) 458-3100.**

Computer Systems, Inc.
MERCURY
The Ultimate Performance Machine.
600 Suffolk Street, Lowell MA 01854-3608

CIRCLE NO. 8

COME TO MA

Media~Link is a registered trademark of Spectrum Signal Processing, Inc.

Together as they should be. On Spectrum's new DSP/PC.

The first to integrate a 33 MFLOPS DSP and a fully functional 386 PC on a single PCB.

The DSP/PC lowers your board count. Increases reliability. Reduces cost. And it's available right now, today.

MA, BOYS.

Still better, its Media~Link® architecture bypasses the ISA bus to accelerate interprocessor communications—and speed up your applications.

Plus, the DSP/PC board comes with integrated software development tools, including libraries, high-level language support, and source-level debugging.

So discover an entirely new meaning for motherboard. Call Spectrum Signal Processing today for your DSP/PC Applications Booklet, or for the name of your local distributor.

In the U.S., call 1-800-663-8986; in Canada and worldwide call 604-421-5422. Fax: 604-421-1764.

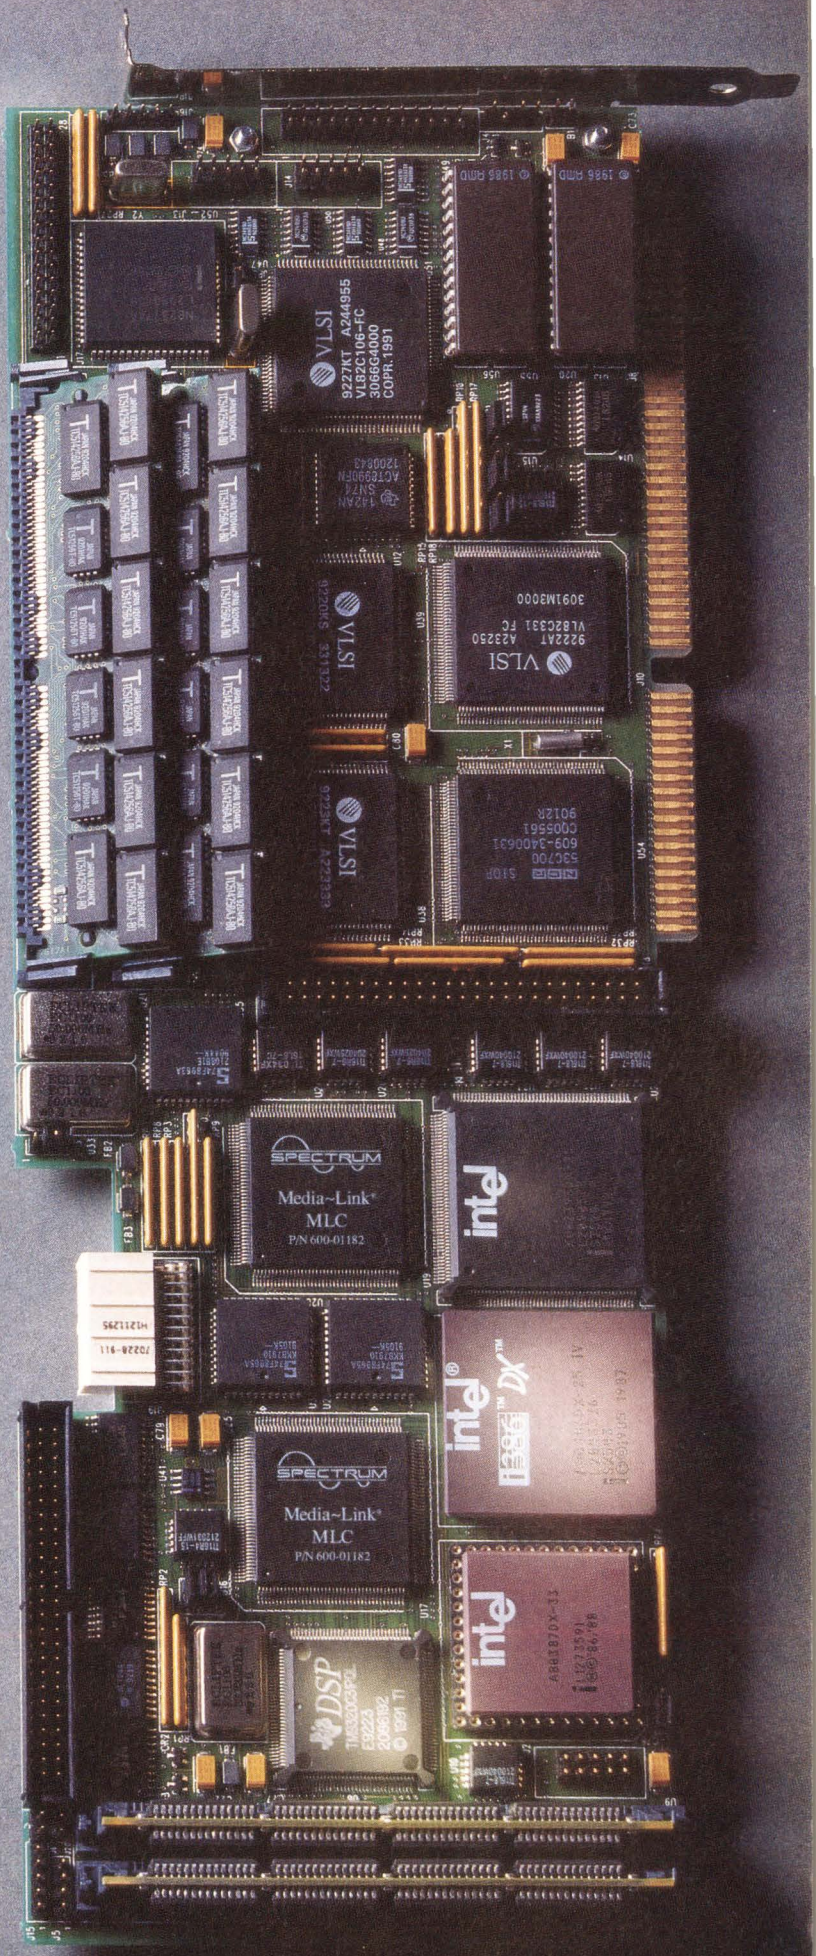

SPECTRUM

Putting DSP to work

CIRCLE NO. 9

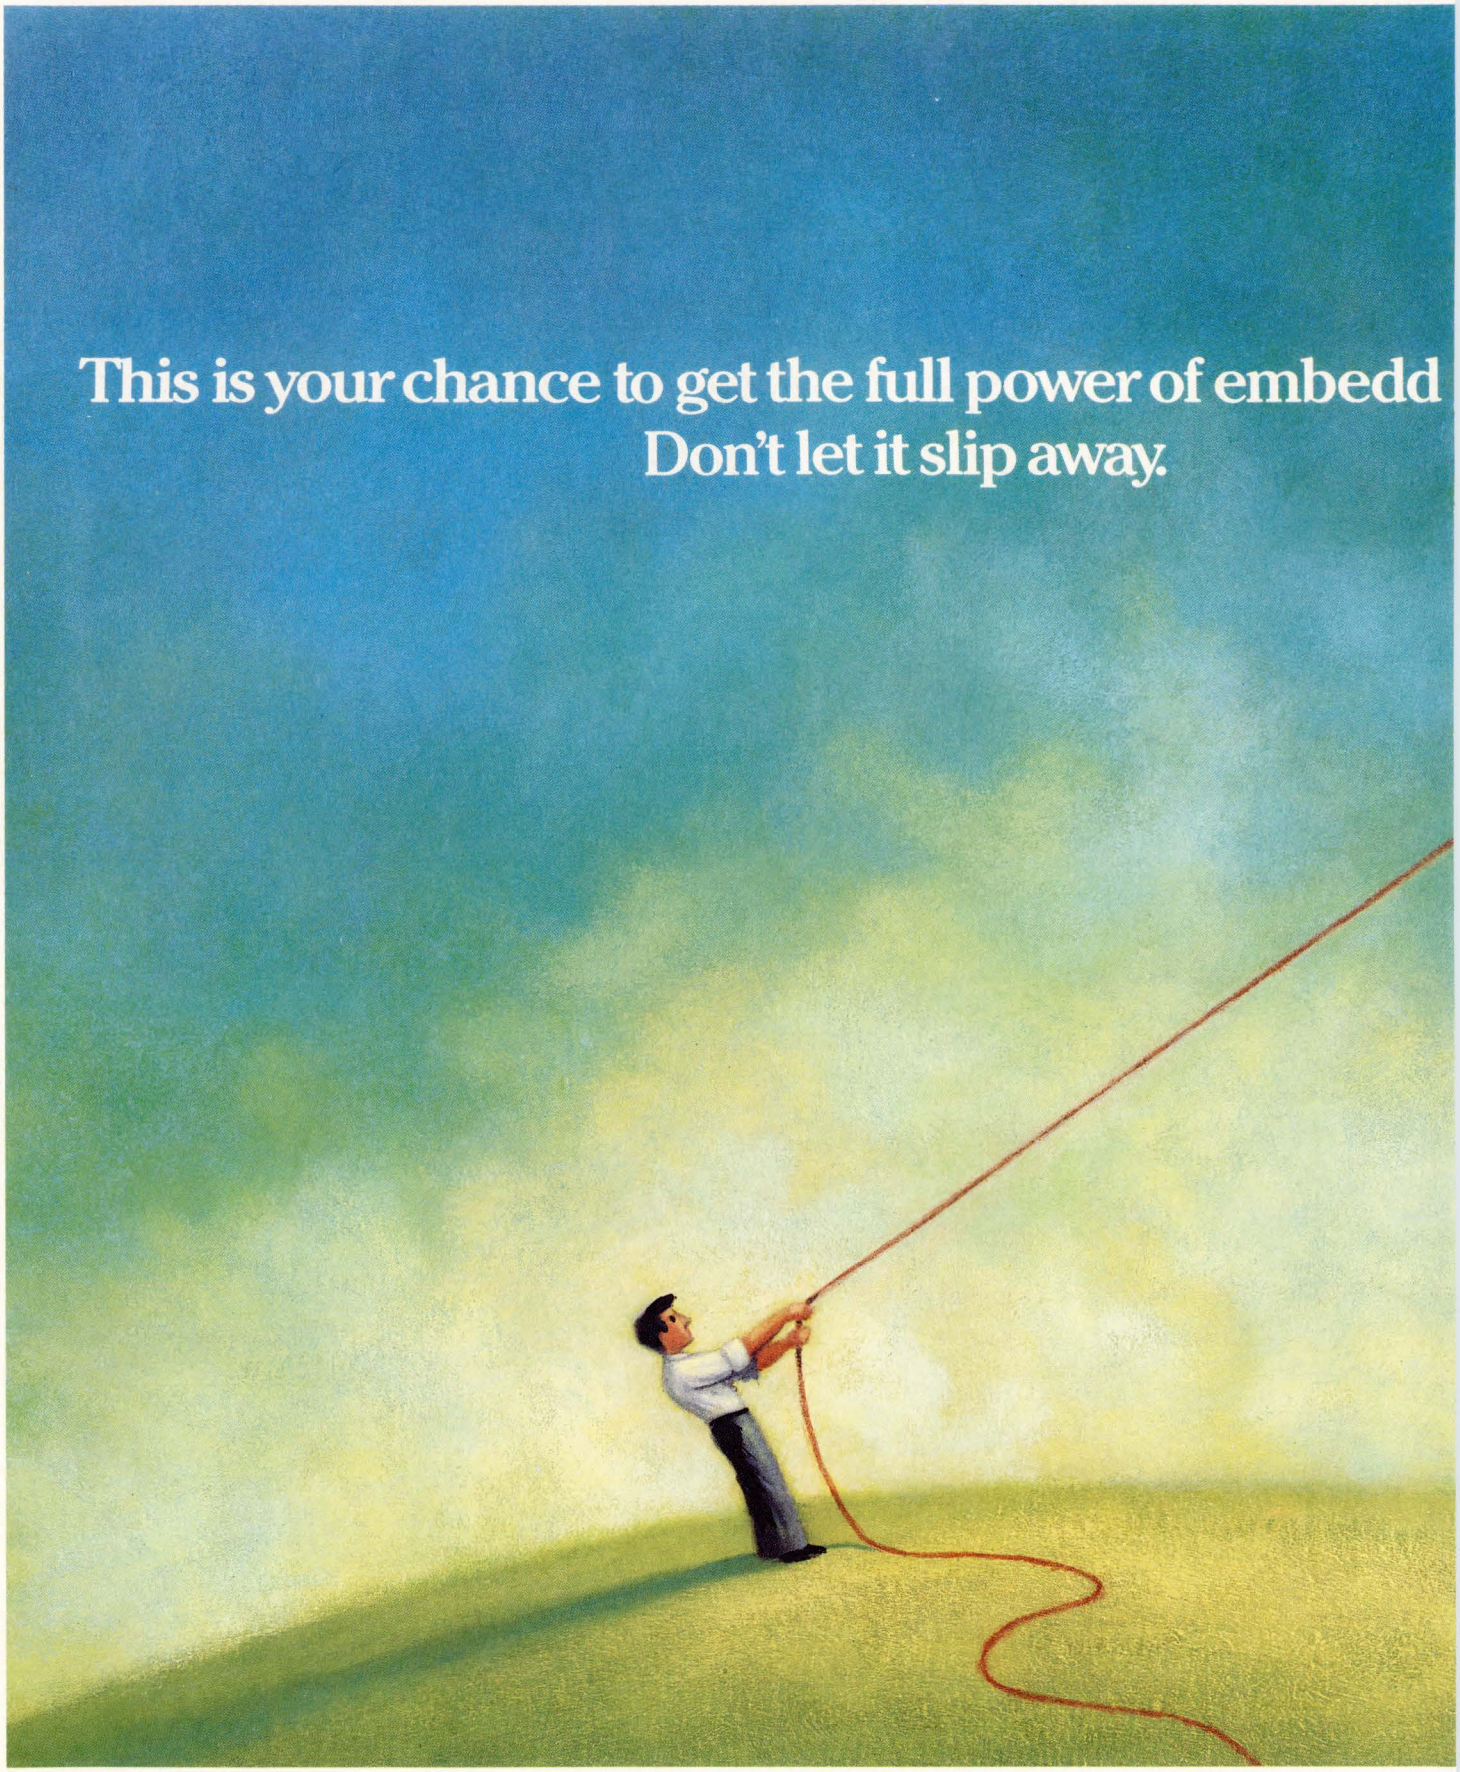

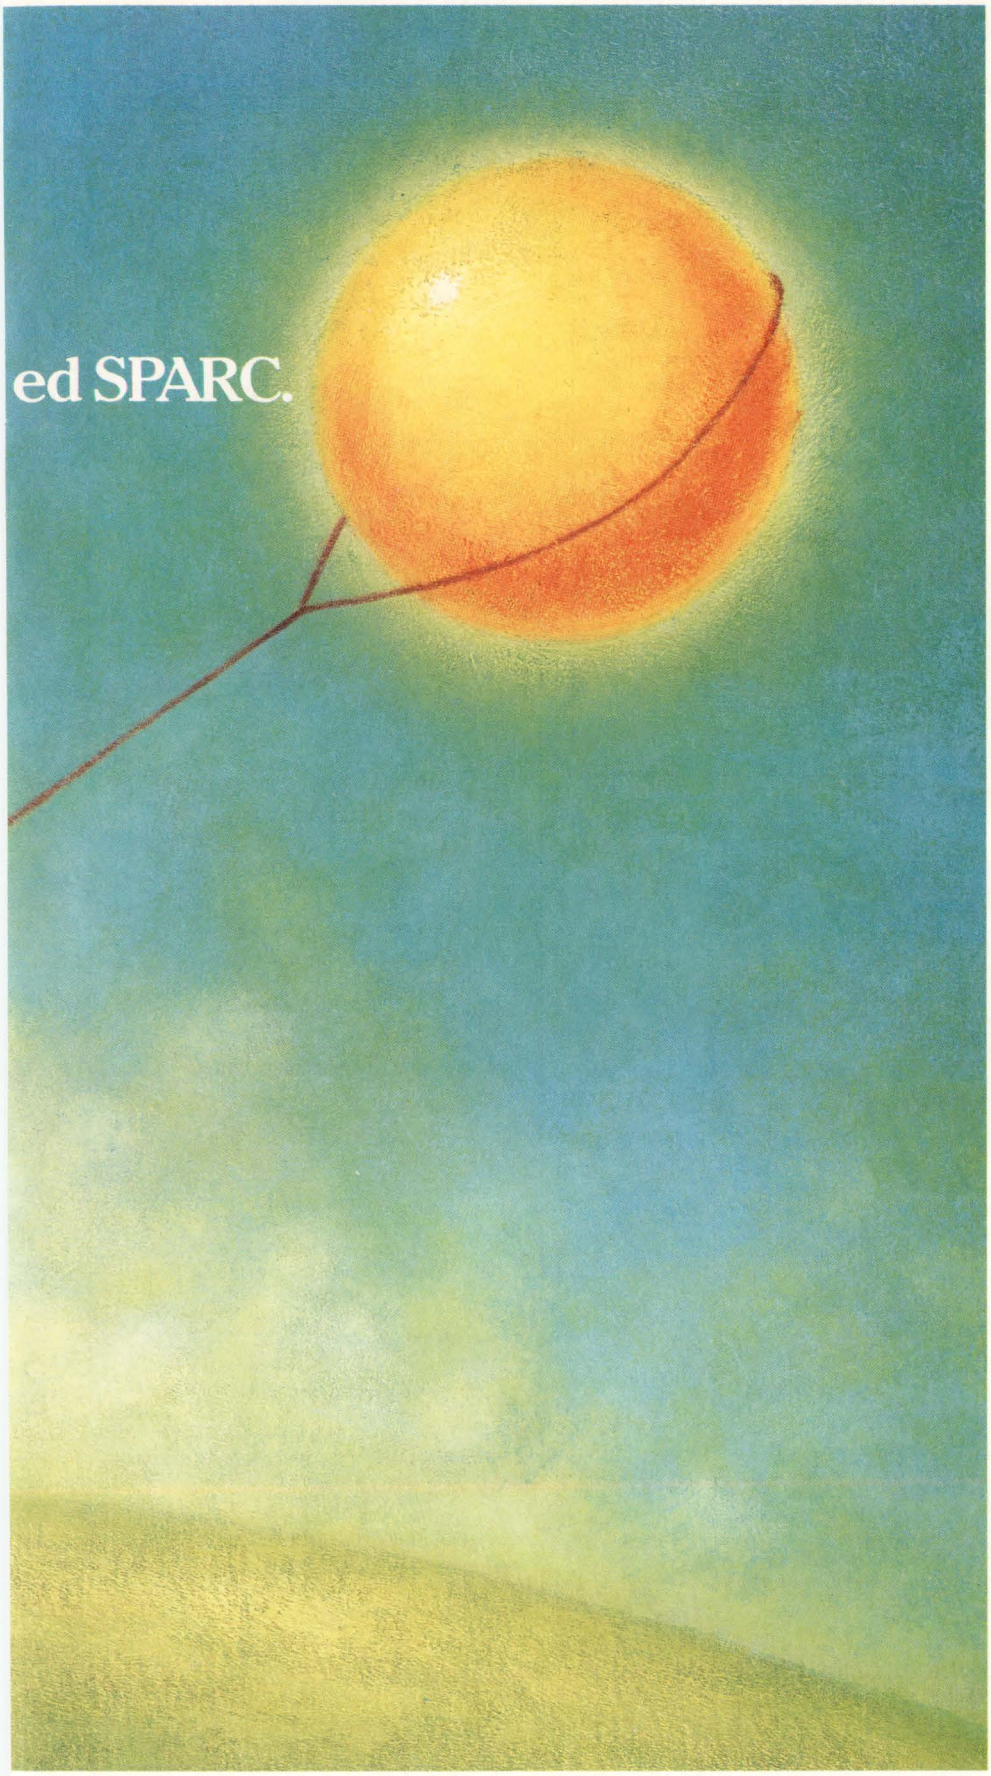

This is one of those times when you just have to seize the day. Because opportunities like this are rare indeed.

You see, FORCE is the only company licensed by Sun to put SPARC™-station 2 technology on VME.

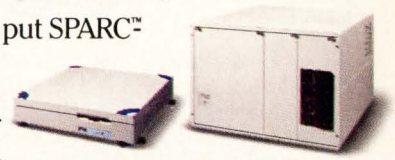

Choose from a range of systems, for a highly integrated solution.

So we're the only ones who can give you validated hardware and software compatibility. Allowing you to run SunOS™ with any of your SPARCstation 2 applications and peripherals. Without a hitch.

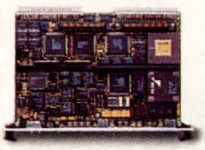

Our CPU-2CE features a 40MHz SPARC RISC microprocessor, delivering 28.5 MIPS.

We also have a whole new family of SPARC 2 products. With everything from our CPU-2CE (6U) board to the *teraforce* 20-slot system. Giving you the perfect combination of real-time and UNIX.

And our SPARC products provide the broadest software offering of any RISC architecture.

So if you want a better grasp of embedded SPARC, call for a free brochure. 800-237-8863, ext. 5. Or in Europe, at 49.89.608-14-0.

Because there's no reason to let all this power slip through your fingers.

The SPARC Family

FORCE COMPUTERS®

CIRCLE NO. 10

COMPUTER DESIGN *Technology and Design Directions*

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

SENIOR EDITORS
Warren Andrews, *Computers & Subsystems*
(508) 283-2102
Barbara Tuck, *ASICs & ASIC Design Tools*
(516) 473-1661
Michael G. Donlin, *CAE/CAD Tools*
(508) 392-2123
Tom Williams, *Software & Development Tools*
(408) 335-5514
Don Tuite, *Integrated Circuits*
(415) 365-5656
Jeffrey Child, *New Product Developments*
(508) 392-2126

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

OEM INTEGRATION
Senior Editors
John Mayer, *Technology & Product News*
(617) 484-5804
Frank Caruthers, *Special Reports*
(215) 993-9595

MANAGING EDITOR:
Arlyn S. Powell, Jr.

CHIEF COPY EDITOR:
Annette M. Staron-Wilson

COPY EDITOR:
Kym Wilson Gilhooly

PRESENTATION MANAGER:
Richard Sarno

EDITORIAL ASSISTANT:
Claire Ellis

ART DIRECTOR: Jan Horner

ILLUSTRATOR: Meg Benner

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

CIRCULATION DIRECTOR:

Robert P. Dromgoole, (918) 832-9213

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

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

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

PennWell PUBLISHING COMPANY

Advanced Technology Group
Dr. Morris R. Levitt, Senior Vice-President
Leslie P. Cypret, Vice-President, Administration
One Technology Park Drive
P.O. Box 990, Westford, MA 01886
Tel: (508) 692-0700 ❑ Fax: (508) 692-0525

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

ATTENTION RETAILERS: To carry *Computer Design* in your store, contact International Periodical Distributors at 1-800-999-1170, ext. 447 (dealer calls only).

BPA ABP

LEGENDARY RELIABILITY

Hitachi's 5.25" 1.65-Gbyte Hard Disk Drive.

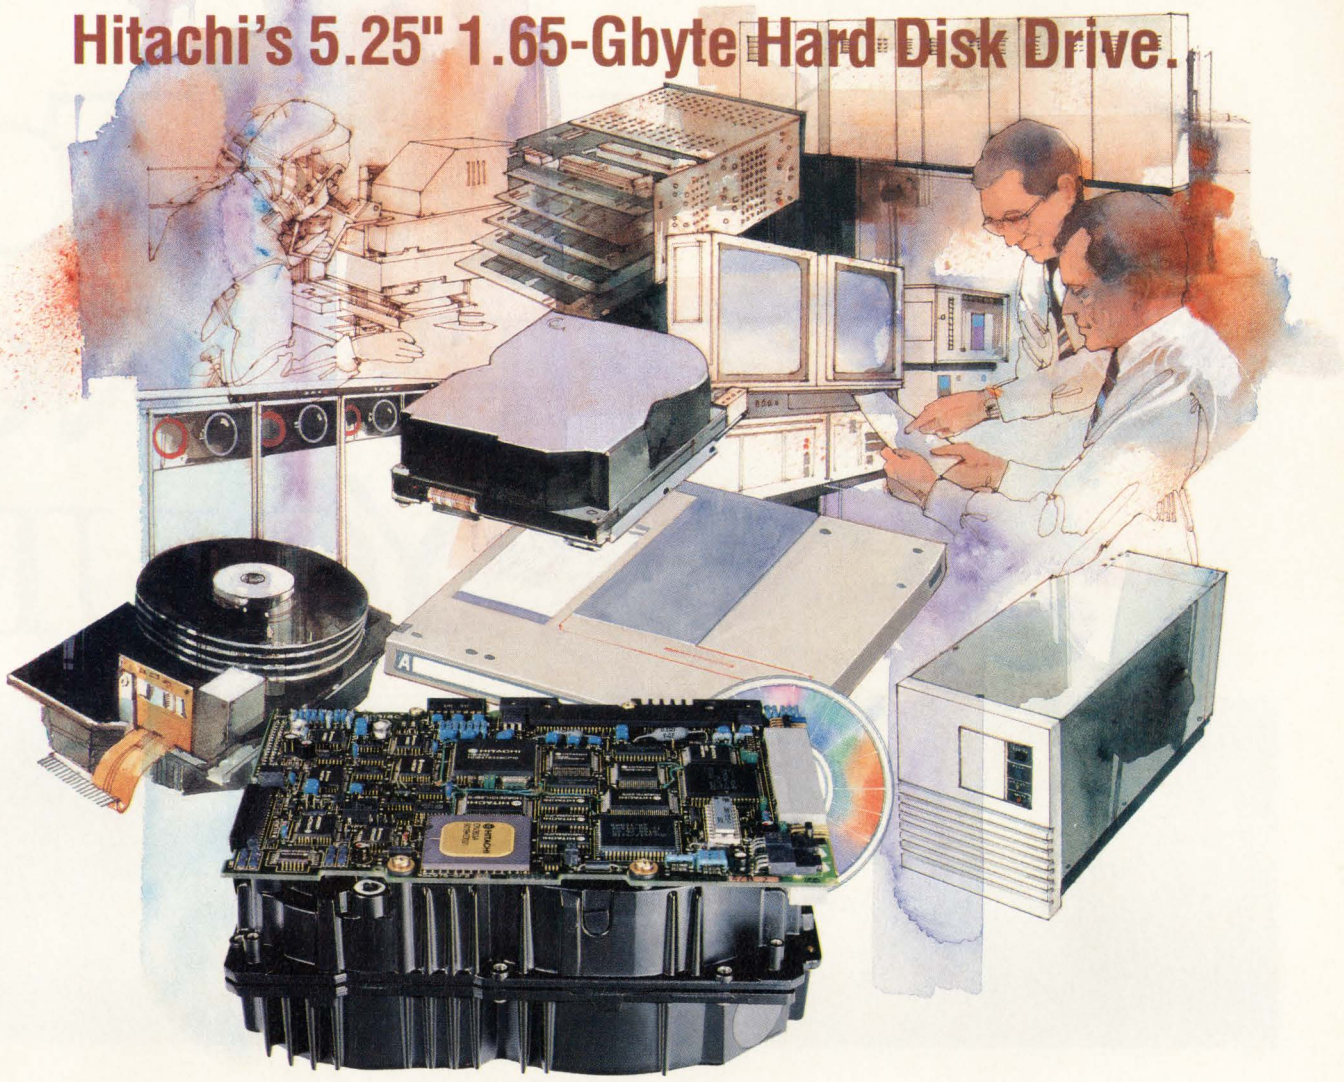

The DK516C-16 Winchester is another legend in a distinguished tradition of Hitachi mass storage products. Hitachi believes that product development starts with the pursuit of maximum reliability. That's why all of the DK516's key components are designed, built, and tested in-house by Hitachi.

Legendary Performance

Hitachi backs up this reliability with equally-impressive performance. The DK516C-16's SCSI interface provides a maximum data transfer rate of 5.0 Mbyte/sec (synchronous), with a 256-Kbyte data buffer and read look-ahead cache. Average seek time is a quick 13.5 ms.

For ESDI applications, choose the DK516-15. This 1.54-Gbyte drive provides a 14 ms average seek time and a 2.75 Mbyte/sec data transfer rate.

From a Legendary $62 Billion Company

Both DK516 drives are brought to you by Hitachi, a company renowned for mass storage reliability and innovation. According to a recent independent end-user site survey by Reliability Ratings, Hitachi drives had the highest Field MTBF and the lowest percent failure rate for any OEM drive rated. Fully 100% of the users surveyed would purchase the Hitachi drives again.

And to back our reliability, we offer one of the longest warranties in the industry. The legend continues. For more information on the DK516, or our new 3.7-Gbyte 5.25" or 1.4-Gbyte 3.5" drives, call 1-800-HITACHI.

Hitachi America, Ltd.
Computer Division
Peripheral Sales & Marketing, MS:500
2000 Sierra Point Parkway
Brisbane, CA 94005-1835

HITACHI®

Our Standards Set Standards

Authorized Distributors

CONSAN STORAGE SOLUTIONS
1-800-229-DISC IA, IL, IN, KS, KY, MI, MN, MO, ND, NE, OH, PA (412), SD, WI

GENTRY ASSOCIATES, INC.
1-800-688-2223 AL, DC, FL, GA, LA, MD, MS, NC, SC, TN, VA, WV

R SQUARED DISTRIBUTING
1-800-777-3478 AZ, CA, CO, ID, MT, NM, NV, OR, UT, WA, WY

SPECIALIZED SYSTEMS TECHNOLOGY
1-800-688-8993 AR, LA, NM, OK, TX

EMJ AMERICA INC.
919-460-8861 DC, DE, MD, NC, NJ, SC, TN, VA, WV

LAW-CYPRESS DISTRIBUTING CO.
1-800-310-6220 AK, CA, HI, OR, WA

SIGNAL COMPUTER PRODUCTS
508-263-6125 CT, MA, ME, NH, NJ, NY, PA, RI, VT

HITACHI CANADIAN
416-826-4100 CANADA

Survey was conducted independently by Reliability Ratings, Needham, MA 02191. The data is from a publicly available report. Reliability Ratings is not affiliated with Hitachi, Ltd., or its subsidiaries, and does not endorse its products.

Call Hitachi for a free copy of the survey.

CIRCLE NO. 13

How To Everything Your

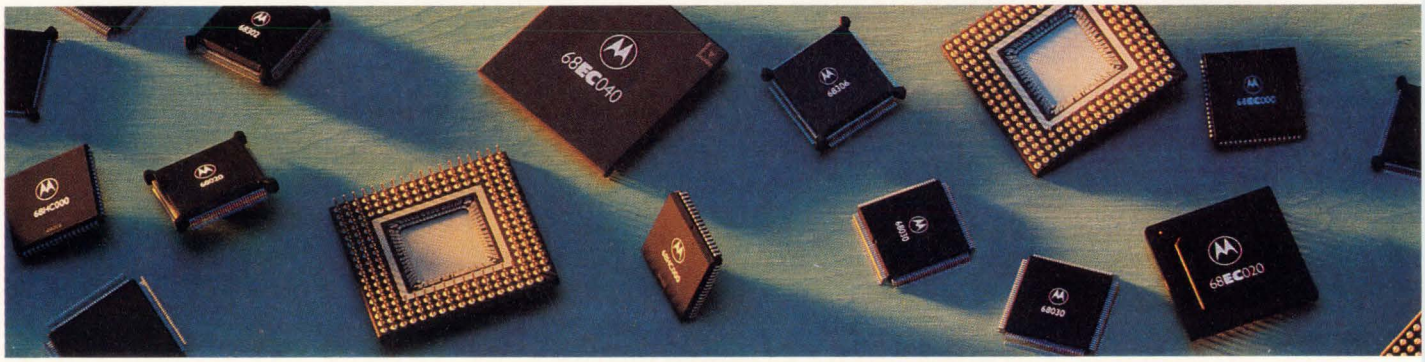

Do what most everyone else does.

Control your next embedded system with a microprocessor from Motorola's 32-bit 68000 family.

Whether your system calls for high performance, low cost, functional integration or a combination of all three, we've got chips that'll meet your needs.

If cost is your first concern, take a look at our industry standard 68EC000 microprocessor. It delivers 32-bit performance for $2 and change.

Which makes it perfect for everything from low-end printers to consumer electronics.

Or if you're in the "speed is everything" camp, consider our 68EC040. A streamlined screamer delivering 29 MIPS of sustained performance, it'll run the pants off RISC controllers in high-end applications. At a lower overall system cost.

Or maybe you need

CONTROL INCLUDING COSTS.

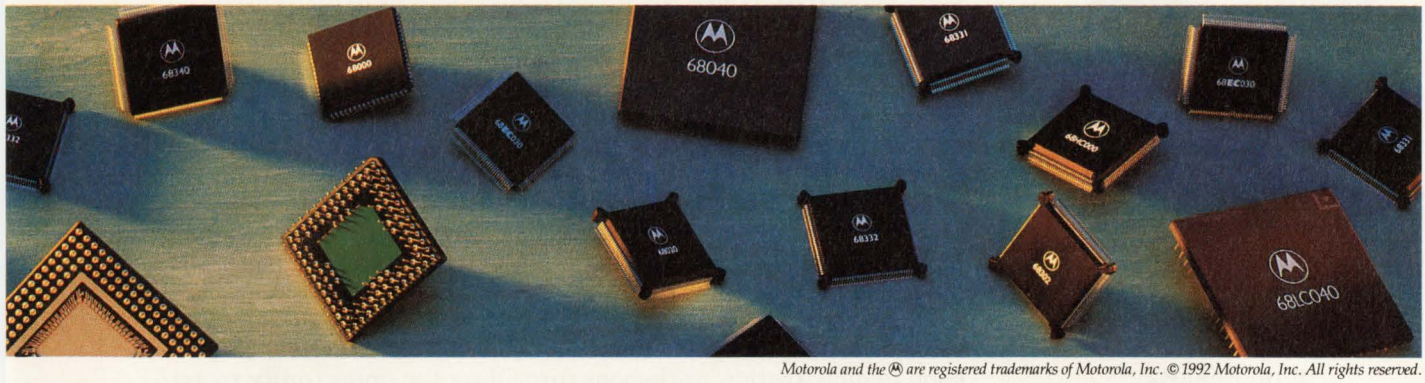

special functionality. Like built-in multiprotocol communications, direct memory access, or sophisticated timers. Our 68300 line of integrated processors comes fine-tuned for a wide variety of applications, from telecommunications to hand-held computers.

And because they're integrated, they take up very little space.

A quick glance across the page will show you there are many more 32-bit solutions where those came from.

Namely, from Motorola. The company that controls more 32-bit embedded systems than the rest of the world combined.

For a free copy of our 68000 Family Brochure, call 1-800-845-MOTO.

And get everything under control.

MOTOROLA

Computer Design takes over A & M-S Design Conference

"A technical conference in these times has to offer enough...to make your sacrifice in time and money worthwhile. We think we've come up with an approach that does that."

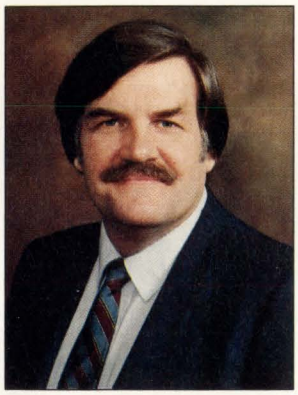

John Miklosz
Associate Publisher/
Editor-in-Chief

Well, the second annual Analog & Mixed-Signal Design Conference is behind us and, with so many of us away from the Westford office on the day of our Halloween party, we again failed to walk away with any prizes. This was the second year in a row that we came up empty-handed and, determined not to let that happen again, we entered into an agreement with Miller Freeman, the original sponsors of the A & M-S Design Conference, to take over full responsibility for future A & M-S conferences. Our first move was to change the scheduled date for next year's conference from the last week in October to the last week in January, 1994!

But we're doing more than just changing the date—we're revamping the entire format of the conference. As A & M-S was originally conceived, it was the conventional conference/exhibition, with technical presentations and sessions running during the same hours that the exhibit floor was open. We think we've come up with a better way to stage conferences as tightly focused as A & M-S. We're testing the approach with RISC '93 in March (see pages 38 and 39) and Fuzzy Logic '93 in July. The basic idea is to eliminate the traditional exhibits and place the conference attendees, as well as the presenters, in a close-up, face-to-face, "total immersion" environment that will expose them to the relevant technologies, tools and applications throughout the entire three days of the conference, morning till night.

With this new format, there'll be half-day tutorials, one-hour lectures, multipaper application-oriented sessions, two-hour afternoon demonstration workshops (providing the opportunity to get some hands-on experience with various design and development tools), and evening rap sessions (with beer and snacks provided). What's more, we're throwing lunch into the package on each day of the conference, with each lunch session featuring a distinguished speaker—a design guru, if you will.

These are tough times for everyone—for designers and design managers at OEMs and system houses, and for vendors of ICs, ASICs and design and development tools. Money is tight and, probably more important, time is tight, with many designers, design managers, product managers, and marketers doing one-and-a-half jobs. A technical conference in these times has to offer enough—and enough value—to make your sacrifice in time and money worthwhile. We think we've come up with an approach that does that. What do you think?

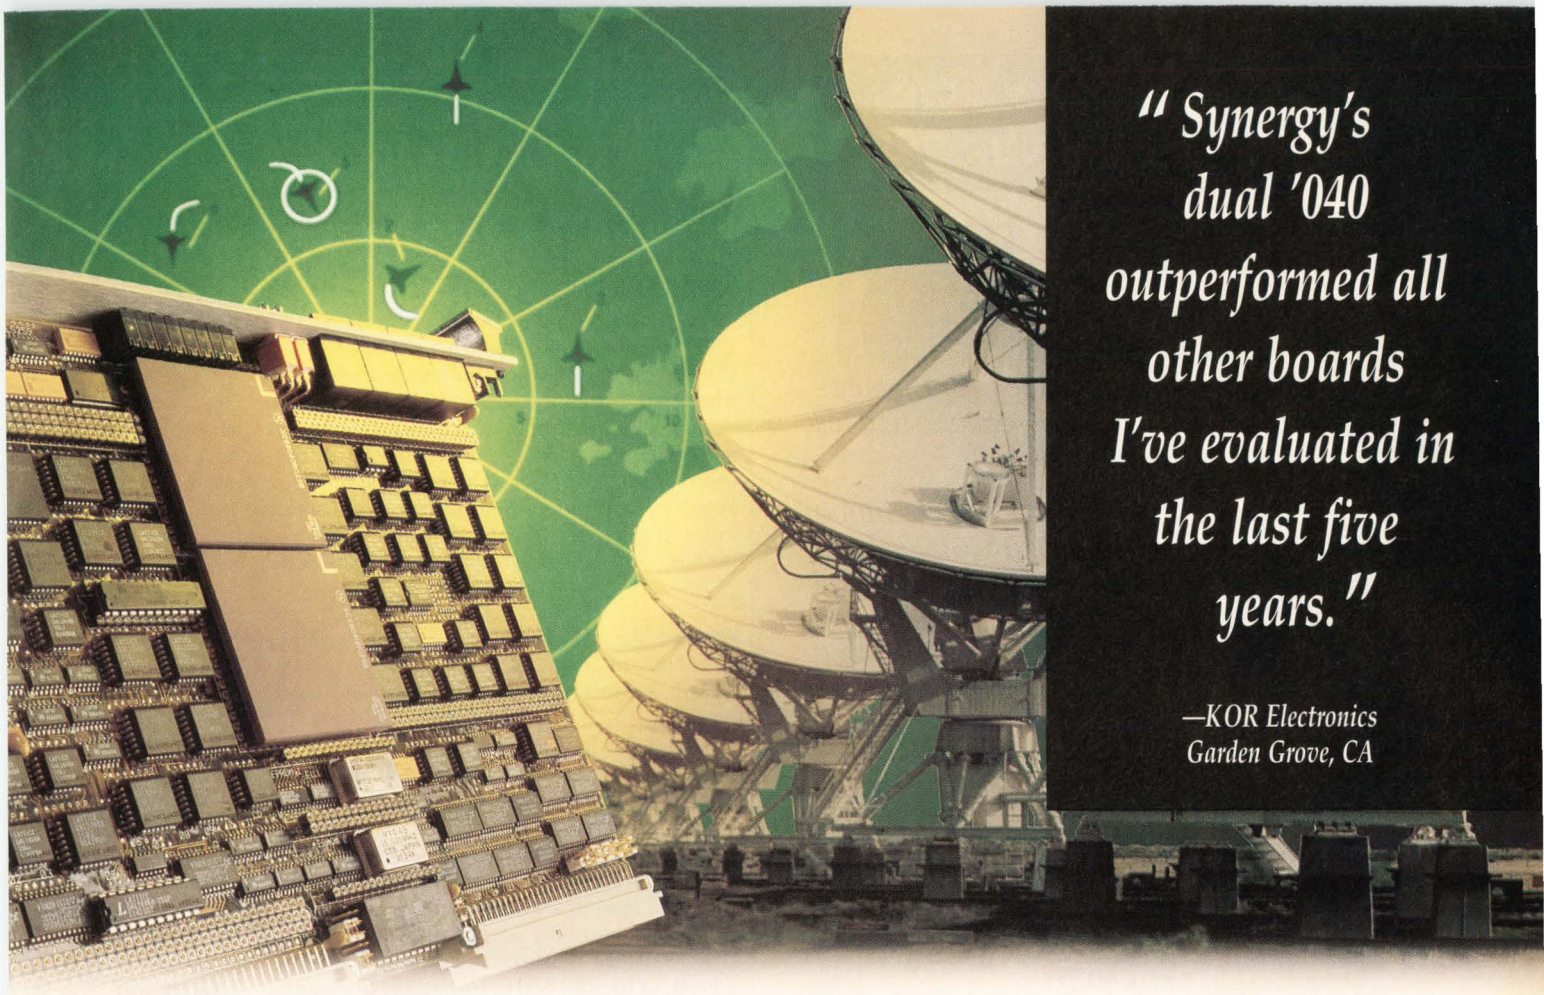

"Synergy's dual '040 outperformed all other boards I've evaluated in the last five years."

—KOR Electronics Garden Grove, CA

You'll find Synergy products built into the world's most demanding applications.

The Goal:

Design the world's only multi-frequency Radar Target Generator System able to simulate hostile threats on the military's diverse radar systems.

The Problem:

To get the needed horsepower, sixteen 68040 CPU boards were required. However, with this many boards, VMEbus bandwidth limits would severely degrade system performance making the project unfeasible.

The Solution:

Synergy's V420 dual 68040 SBC. After evaluating several products, KOR Electronics selected eight V420s which could deliver 320 MIPS without VMEbus bandwidth degradation.

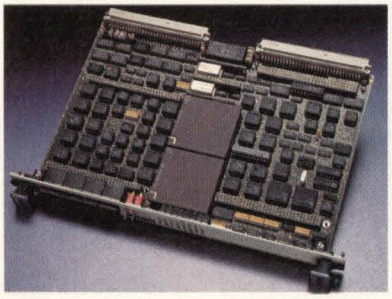

"Not only did we meet our project performance goals, we were able to reduce our costs by 40%."

Unexpected Benefits:

Synergy's dual '040 solution cut KOR's hardware requirements by 50% while vastly increasing system reliability. These unexpected benefits reduced system costs by 40%.

In Synergy, KOR also found a design partner with strong integration expertise and dependable customer support.

"I recommend you call Synergy today."

Next time you need high performance SBCs, do as KOR Electronics did. Call Synergy Microsystems. You'll be as satisfied as they are.

SYNERGY microsystems

High performance SBCs for demanding applications.

SYNERGY MICROSYSTEMS, Inc. 179 Calle Magdalena, Encinitas, CA 92024, 619-753-2191, Fax 619-753-0903

CIRCLE NO. 16

CALENDAR

CONFERENCES

December 13 - 16
1992 IEEE IEDM

Hilton Hotel, San Francisco, CA. The 1992 IEEE International Electron Devices Meeting brings together engineering professionals from industry, government and academia. The meeting features 36 sessions on such topics as solid-state technology, integrated circuits and quantum electronics. Also offered are several short courses, plenary sessions and panel discussions. Contact: Melissa Widerkehr, IEDM, Ste 610, 1545 18th St NW, Washington, DC 20036, (202) 986-1137, Fax (202) 986-1139. **Circle 366**

January 3 - 6, 1993
VLSI Design '93

Taj Intercontinental Hotel, Bombay, India. The Sixth International Conference on VLSI Design, with the theme *Chip, Board and Systems Design in the '90s*, brings researchers and designers to the west coast of India. The four-day program consists of paper sessions, posters, tutorials, and industrial CAD exhibits, covering such topics as CAE/CAD systems, logic synthesis, design for testability, circuit simulation, analog devices, and economic issues. Contact: Rochit Rajsuman, Dept. of Computer Engineering & Science, Case Western Reserve University, Cleveland, OH 44106, (216) 368-5510, Fax (216) 368-2801. **Circle 367**

January 6 - 8
WEST '93

San Diego Convention Center, San Diego, CA. The AFCEA and U.S. Naval Institute Western Conference & Exposition features a technical program directed to military, government and industry professionals in the fields of military weapon systems, computers, communications, aerospace, and electronics. The conference focuses on military- and space-related issues, joint requirements and naval and drug enforcement applications in imaging. Also offered are technical panels, development courses, career transition seminars, and more than 160 exhibits. Contact: Ginny Bracken, J. Spargo & Associates, 4400 Fair Lakes Ct, Fairfax, VA 22033, (800) 336-4583, Fax (703) 818-9177. **Circle 368**

February 22 - 25
EDAC-EUROASIC

EDAC-93
EUROASIC-93

CNIT Conference & Exhibition Centre, Paris, France.
EDAC, the European Conference on Design Automation, and EUROASIC, the premier European event in ASIC design, will be held jointly this year to provide a forum for a common discussion on ASIC design and design automation. Sponsored by the EDAC Association, the event is expected to draw more than 80 exhibitors, and covers such topics as design techniques and methodologies, high-level design tools, simulation, and testability. Contact: CEP Consultant Ltd, 26-28 Albany St, Edinburgh, EH1 3QH, UK, +44 31 557 2478, Fax +44 31 557 5749. **Circle 369**

With AMD's New PCnet™-ISA Solution On Your Mother—

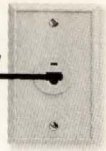

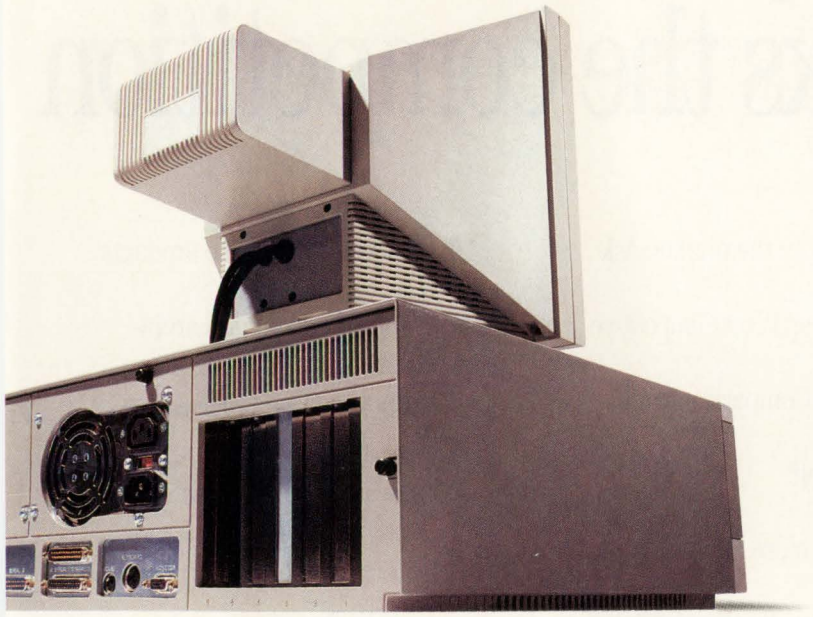

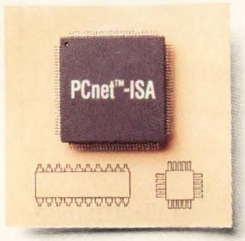

AMD's complete PCnet-ISA solution covers less than five square inches.

Look in almost every office and you'll find a twisted pair outlet right on the wall. But look for an efficient IC solution to make your PC design Ethernet-™ ready and it could drive you right *up* the wall.

Unless you look to AMD. Our new PCnet-ISA—a true one-chip Ethernet controller—gives you the integrated features you need to make your next PC Ethernet-ready. So you get a complete 10BASE-T design in less than five square inches of real estate.

PCnet-ISA is going to save you more than just space. Since your total solution will run you less than $25 in volume, PCnet-ISA will shine on the bottom line.

You'll also save design time, because no additional memory is needed. And you'll deliver higher performance, because data transfers directly to the host memory, instead of through a local buffer.

The software's ready to go too. Driver support is already available for Novell NetWare,® Microsoft LAN Manager, Banyan Vines,® Artisoft LANtastic,® SCO UNIX® and others.

AMD has over 15 years of experience in networking. PCnet-ISA is yet another in a long line of networking IC solutions from AMD for both Ethernet and FDDI LANs. Our cooperative partners include such industry leaders as DEC, HP, and SynOptics. So when you're working with any of AMD's networking solutions, you'll get the engineering support you need just by picking up the phone.

And when you're ready to go to work with PCnet-ISA, call **1-800-222-9323** and ask for Literature Pack 16M. Then plug into a networking leader called AMD.

Advanced Micro Devices

Introducing the With twice the performance and it sinks the competition

The ADSP-21020: with an enhanced Harvard architecture, you benefit from the highest performance and flexibility available in a floating-point DSP today.

Compared to the ADSP-21020, everything else is substandard. It's got the performance DSP designers want for ultra-demanding applications, like image processing, graphics, military, voice communications, and speech recognition. At 33.3 MIPS/100 MFLOPS and a 1024-point complex FFT benchmark in 0.58 ms, the ADSP-21020 is at least twice as fast as anything else on the market. Ask for our "Floating Point Competitive Benchmarks Comparison" and see for yourself.

For even more affordable performance, consider the ADSP-21010. For just $49.90 (in 100s) you get a floating-point DSP processor which is code-compatible with our ADSP-21020. Both are part of our new family of floating-point DSPs offering high performance and a path to higher on-chip integration. Just as with the ADSP-2100 family of fixed-point DSPs, these products use an enhanced Harvard architecture for optimum performance and flexibility.

ADSP-21020 Performance Benchmarks

- 1024-point complex FFT (radix-4 with digit reverse): 0.58 ms
- 1024-point real FFT (radix-2): 0.34 ms
- Divide: 180 ns
- $1/\sqrt{x}$: 270 ns
- Cubic Bezier evaluation: 300 ns
- Cubic B-Spline evaluation: 450 ns

We help you get started fast, too. With development tools that are easy to use and accelerate your code-writing and time-to-market.

Authorized North American Distributors: *Alliance Electronics* 505-292-3360 · *Allied Electronics* 800-433-5700 · *Bell Industries* 310-826-2355 · *Future Electronics* 514-694-7710 (Canada) 508-779-3000 (USA) · *Hall-Mark Electronics* 214-343-5000 · *Newark Electronics* 800-367-3573 · *Pioneer-Standard Electronics* 800-874-6633 · *Pioneer Technologies Group* 800-227-1693 · *WYLE Laboratories* 800-866-9953 · *Zentronics* 416-507-2600 (Canada)

ADSP-21020. a complete tool kit for $995.00, in floating-point DSP.

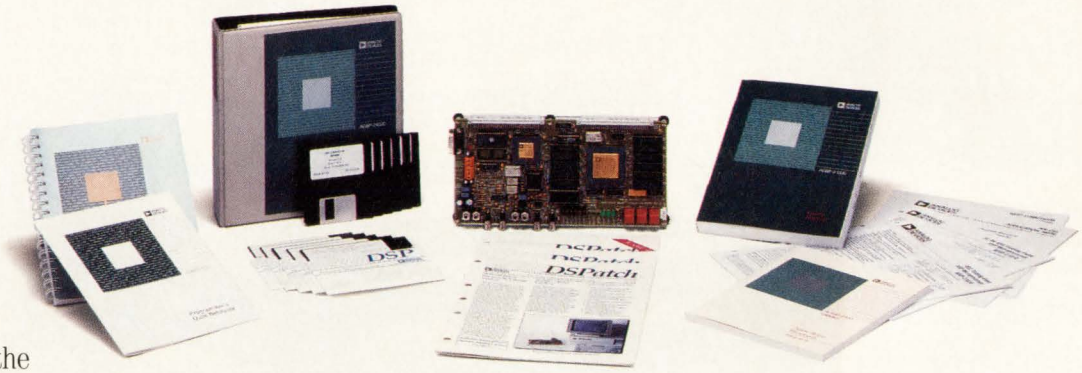

Analog Devices' EZ-KIT: the most extensive (and least expensive) DSP development tool kit on the market today, with software development tools, a powerful simulator, the EZ-LAB board for prototyping, and (in the EZ-KIT-PLUS), a C Compiler and source-level debugger.

The EZ-KIT enables you to evaluate a prototype design quickly, with modules to program and simulate both the ADSP-21010 and ADSP-21020; link in optimized library routines for key DSP algorithms; and run the application on the EZ-LAB Evaluation Board that's included. At $995, it's a very cost-effective tools package. For all of the above, plus C Compiler, runtime library, and source-level debugger, our EZ-KIT-PLUS is available for just $1,995. You can do real-time in-circuit emulation with EZ-ICE, and soon you'll be able to run Ada code or the SPOX Operating System on the 21020/21010 as well.

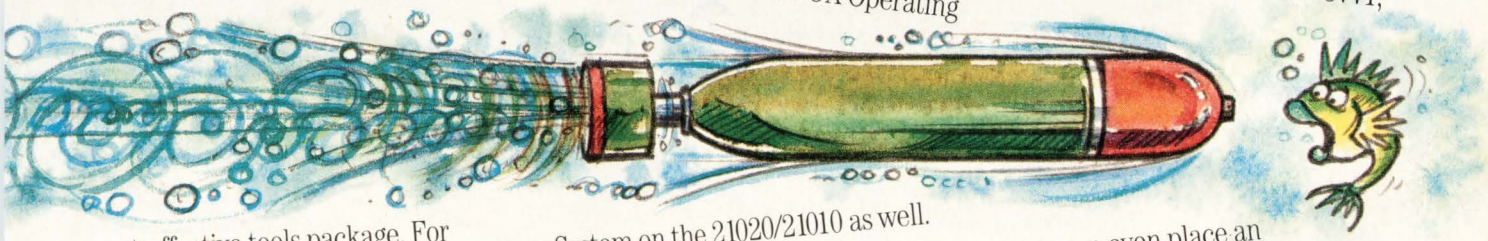

Now is the time to get on board with floating-point DSP. Because these great prices on our EZ-KIT and EZ-KIT-PLUS are only valid through March 31, 1993. To order your EZ-KIT or receive a product information packet, call us at (617) 461-3771; you can even place an order with your credit card.

ANALOG DEVICES

Analog Devices, Inc., One Technology Way, P.O. Box 9106, Norwood, MA 02062-9106. Distribution, offices and application support available worldwide.
EZ-LAB and EZ-ICE are registered trademarks of Analog Devices, Inc. SPOX is a registered trademark of Spectron Microsystems Corporation.

CIRCLE NO. 20

VMEbus • VXIbus • Multibus II • PC XT/AT • EISA • Custom Designs

HIGH-PERFORMANCE BACKPLANES

Optimum system performance lies with the backplane — the heart of your system and the key to error-free interboard data transfer.

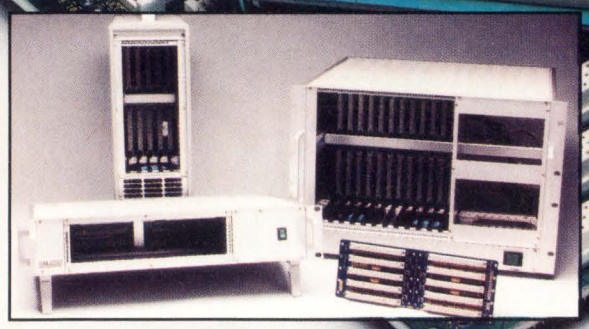

BUSTRONIC's high-performance backplanes are designed to meet and exceed today's high-speed computing requirements. With incremented CPU and logic speed, there is a greater concern for low noise and improved reliability. BUSTRONIC has developed the best possible products available for today and the foreseeable future. Each signal line is matched and balanced for the optimum in characteristic impedance. All backplanes offer versatile power distribution and a connection scheme with distributed decoupling capacitors for both high and low frequency. Discuss your custom or MIL-spec requirements with our engineers.

bustronic

Please contact us for more information.

Bustronic Corp. • 44350 Grimmer Blvd. • Fremont, CA 94538
TEL (510) 490-7388 • FAX (510) 490-1853

CIRCLE NO. 21

Tom Clark on: FPGA design

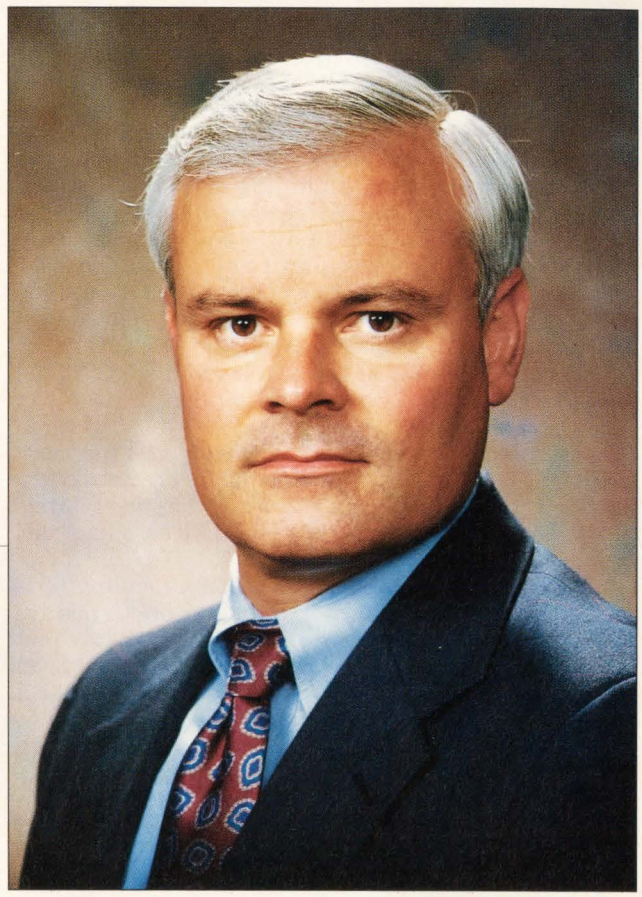

Logic designers using field programmable gate arrays (FPGAs) and complex programmable logic devices (CPLDs) have two basic choices when it comes to design methodology—structural and behavioral. Until recently, the preferred choice among most designers was structural design—namely, schematic capture. But a strong movement is underway toward a shift in design methodology.

Today, many systems designers are moving away from schematic design entry and are taking advantage of the higher level of abstraction offered by behavioral design methodologies. This doesn't mean that engineers are phasing out their schematic capture products. Rather, it signifies a shift toward a new design environment that utilizes hardware description languages (HDLs) as its core, and alters the role of schematic capture to provide less design and more documentation.

New architectures push change

Pushing this shift in methodology into the spotlight is the immense popularity of FPGAs and CPLDs. These new device architectures combine the best features of ASICs with the best characteristics of PLDs. Initially, schematics were the entry vehicle of choice, primarily because many of the early adopters of FPGAs were ASIC designers. Schematic entry is also inherently a lower-level (and so more detailed) representation of a design than a behavioral description, and offers better control over the silicon.

In addition, FPGA vendors have realized that design at the schematic level effectively locks a designer into one supplier's technology. Anyone who's changed a gate-array vendor in midstream immediately knows why—to retarget a schematic-based design to another vendor, you have to redo much of your work. With FPGAs, the situation is even more severe. Whatever one-to-one remapping you might get away with in gate-array retargeting is impossible, or at least hopelessly inefficient, for moving between FPGAs. This is because each FPGA vendor's architecture and granularity are dramatically different.

Thomas R. Clark, president and CEO, Data I/O Corporation, Redmond, WA

The move to behavioral methods

Several trends are weaning FPGA and CPLD users from schematic capture, and shifting them to behavioral methods. The first is the increase in the number of device vendors; the second is the increase in device density; and the third is the advent of new behavioral methodologies capable of handling the complexity offered by FPGAs.

When FPGAs were first made available, there was only a single vendor of choice—Xilinx. Later, Actel, AT&T and Texas Instruments entered the market. With these limited choices, designers weren't averse to a vendor-specific toolset. Today, however, a growing number of silicon vendors have FPGA offerings. Add to this the growing number of complex PLD solutions, many of which offer comparable density and higher performance than their FPGA counterparts, and it's obvious you have many choices.

In June's *Computer Design* Technology Viewpoint, Cyrus Tsui provided an excellent profile of our industry, and we agree with him that there's little chance of an industry architectural standard in the near future. As a result—and to combat the hype and false promises accompanying some device introductions—many designers have adopted one or two "pet" architectures. This benefits the few vendors who were first to market with FPGAs, and is a serious obstacle to those who've come later.

The problem with this defensive strategy is that all

FPGAs aren't created equal. In fact, there are significant differences in how well particular applications map to various architectures. Some FPGAs are significantly better at certain kinds of circuits than at others, and even in a particular circuit class (datapaths or state machines, for example), there are typically only one or two architectures that best match a given, specific circuit. Narrowing down the number of candidate architectures may be a perfectly rational way to deal with the complexity and economic rigor of a market featuring so many choices, but it's not the best way to match silicon to a design and ensure high utilization and top performance.

This abundance of choice is a leading factor in the shift to behavioral entry. The higher the level of abstraction designers work at, the less they're tied to any particular piece or class of silicon, and the more readily the design can be retargeted. Behavioral entry lets you migrate designs between different architectures, letting you choose the one that best fits your application.

Perhaps the most important factor leading to a methodology shift is the fact that FPGAs and CPLDs are already meeting the density requirements of the majority of today's gate-array designs. At this level (around 10,000 gates), design by schematic is still practical, but efficiency is questionable. The higher level of abstraction offered by behavioral design methodologies becomes a major factor in overall design efficiency. The use of HDLs for design entry lets you describe a design in a device-independent and higher-level fashion, with synthesis routines completing the tedious and time-consuming work of matching appropriate design elements into correct architectural features.

The third force in the shift from schematic to behavioral entry is the development of logic synthesis for FPGAs. FPGAs present a more difficult technological challenge for synthesis than the less constrained ASIC. The architectural variations among different vendors' FPGAs are vast in comparison to those evident in gate arrays.

Device fitters to the fore

It's only in the last two years that synthesis has become sufficiently practical to work its way into the mainstream, but it now offers a viable solution for FPGAs. At the forefront of this synthesis movement are software algorithms called device fitters. Device fitters, in the same vein as the FPGAs they support, have borrowed techniques and algorithms from traditional ASIC design software, and have melded them with generic PLD optimization routines and derivatives.

A device fitter is a program that synthesizes a generic logic description into an implementation that is optimal for a particular architecture. It works in a fully automatic mode, but also lets you manually specify placement, routing criticalities, buffering, and other characteristics. The fitter, rather than the designer, takes on the burden of knowing the low-level details of the target silicon intimately enough to efficiently implement an application into a circuit.

Today's device fitter synthesis capabilities dispel many previously held beliefs that prevented FPGA behavioral methodology from being widely adopted. Many assume, for example, that the best way to run synthesis is on a complete design. This is not true—synthesis performs far better if run on a design's submodules. Many also believe that synthesis will always improve a circuit's size or speed. Again, this isn't true. There are classes of circuits upon which synthesis nearly always fails, and upon which it should never be run. Many think that schematics are as good a candidate for synthesis as behavioral designs. This also isn't true—schematics contain valuable knowledge about how a circuit should best be structured, knowledge that's often beyond a synthesis algorithm's ability to divine.

The ideal design solution includes the ability to use behavioral entry as well as schematics, applying each to the portions of the design for which it's best suited. Making these points clear to the design community—educating users through design examples and training—is a prerequisite for achieving full acceptance of behavioral entry in the mainstream.

Applying benchmarks

Data I/O, as well as other vendors, have been involved in the past year in the PREP benchmarking effort. This Programmable Electronic Performance consortium is a group of PLD manufacturers and tool vendors whose objective is to develop a suite of benchmarks to accurately measure the functional capacity and speed performance of either PLDs or FPGAs. PREP was not formed to standardize such items as physical interfaces, pinouts, architectures, or deal with spec sheet issues, such as electrical characteristics. The PREP solution is an important and useful first step toward helping designers understand how the FPGA architecture announcements they see in the press every month translate into gains in their ability to design and deliver the products their markets demand.

Ultimately, we believe benchmarking will move to the desktop, where you can decide for yourself which devices give the best gate utilization, the best in-system speed, the best economic choice. That's why we've spent considerable energy in the past two years developing and constantly improving and adding to our list of device fitters. Our goal is to see designers armed with a universal front end that allows device-independent design, along with a rich back end, powered by device fitters, that can intelligently retarget designs between the full spectrum of architectures—and that lets you benchmark your circuits in any candidate device.

Movement in the CAE market, however, is slower than we've hoped for. Behavioral entry is important and its momentum is building, but it will not totally replace schematics. Automatic device fitting has appeared, but designers will always want additional control. As we've learned elsewhere, the best revolutions are those that augment and build upon the past.

FASTEST IN.

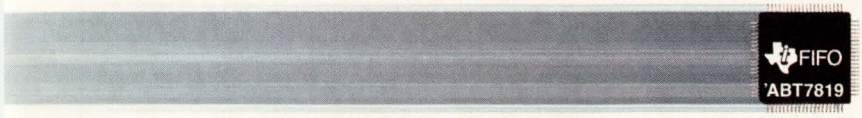

FASTEST OUT.

THE NEW 80-MHz 'ABT7819 FIFO FROM TEXAS INSTRUMENTS.

Introducing the industry's first 80-MHz FIFO, the SN74ABT7819. It's the newest addition to a full line of high-performance, clocked, 18-bit FIFOs from Texas Instruments. And it can run circles around the FIFOs you're using now.

Blazing speed with excellent metastable characteristics

Based on our power-saving Advanced BiCMOS technology, the 'ABT7819 has a maximum access time of 9 ns (measured at C_L = 50 pF). It's the first to offer such performance. And its unique multistage synchronization circuitry greatly improves the metastable characteristics of status flags. That makes it an excellent choice for today's high-speed multiprocessing systems—or any system where speed and reliability are paramount.

BIDIRECTIONAL FIFOs

Device	F max (MHz)	Max access time (ns)	Data setup time (ns)	Data hold time (ns)
'ABT7819	80	9*	3	0
'72615†	40	15	6	1
'5420‡	40	16	12	0

* C_L = 50 pF vs. standard 30 pF test load.
Specs based on manufacturers' data as of April 1992,† July 1992‡

Highly integrated in a choice of standard packages

The bidirectional 'ABT7819 is more than just two 512-word-deep by 18-bit-wide FIFOs sitting side by side. It's bidirectional control logic, flag logic, latches and synch circuitry integrated into a single package to simplify your designs.

The 'ABT7819 comes in compact 80-pin SQFP (actual size 14 mm x 14 mm tip to tip) and standard QFP packages to help save valuable board space. Plus, it's available in −12-, −15-, −20- and −30-ns cycle times for easy performance upgrades.

What are you waiting for—call today

For more information on the SN74ABT7819, call **1-800-477-8924, ext. 3011**, or contact your local TI sales office. We'll have you up to speed in no time.

TEXAS INSTRUMENTS

00-7536

INTEGRATED CIRCUITS

New clock chips are analgesic for run-length headaches

Don Tuite, Senior Editor

Synchronizing the edges of a distributed clock so that every chip in a system is clocked at the same instant has been made easier with new chips from Vitesse (Camarillo, CA), TriQuint (Santa Clara, CA) and Cypress Semiconductor (San Jose, CA). All the chips give you flexibility in routing clock lines on a circuit board. In some cases, prototype timings can be trimmed ad hoc using the programming inputs to the chips.

Some of these clock distribution chips also offer frequency multiplication and division. Division lets you reduce the frequency of a clock that is being sent off the board, reducing potential EMI problems; multiplication lets you increase it again. Because these chips use phase-locked loops (PLLs) that feed selected inputs back into the loop, they can be operated as zero-propagation-delay clock distribution chips. Propagation delays are 1 ns or less.

Cypress breaks rules

Despite the similarities between the purposes of these chips, there are differences among them—most notably, between the gallium-arsenide (GaAs) devices from Vitesse and TriQuint and the silicon device from Cypress. Not only did Cypress violate the "rule" that says you can't run a silicon PLL that fast, but the company's engineers designed trilevel inputs that simplify programming while providing a larger array of timing options than bilevel inputs.

It's easier to understand the differences between the chips by thinking of them in terms of the voltage-controlled oscillator (VCO). In all these devices, the VCO runs at a multiple, N, of the clock frequency. The higher the VCO frequency, the better, in terms of the precision with which you can control the outputs. In the GaAs parts, the oscillator runs at a speed as high as 840 MHz. It would be difficult for a silicon part to run that fast, but Cypress uses a trick to produce a precision equivalent to running a straight oscillator at 1,300 MHz.

You control the outputs of all these chips in two ways: by programming skew-control input pins, or by selecting which output is fed back to the input of the PLL. In these ways you are selecting values for N, and for edge placement in various channels.

Flexible output clocks

TriQuint's 80-MHz GA1000 has six outputs that can be programmed to run at one or two times the input frequency. This means the output frequency can be as high as 160 MHz. There are two skew-control inputs. With all the outputs in phase, the maximum skew between any two of them is guaranteed to be 500 ps. Typical values are 250 ps.

In programming the chip, you select a value for N between 4 and 22. You also determine where the rising and falling edges of the output fall with respect to the rising and falling edges of the internal VCO clock. Because the VCO frequency must be between 320 and 440 MHz, your system clock frequency determines what values of N are possible, and so what degree of timing precision you can achieve. For example, at 66 MHz, N must be 5 or 6—that is, 330 and 396 MHz are the only values of N multiplied by 66 MHz that fall in the allowable range. Depending on which value of N you choose, you can delay edges in multiples of 2.53 or 3.03 ns. At an input frequency of 55 MHz and a value of 8 for N, you get the smallest increment of edge placement, at 2.27 ns.

TriQuint's 50-MHz GA1110E has six outputs that you can shift in increments of 2.5 ns. The guaranteed maximum input skew is 500 ps, with the typical being 250 ps. All six

The basics of skew control

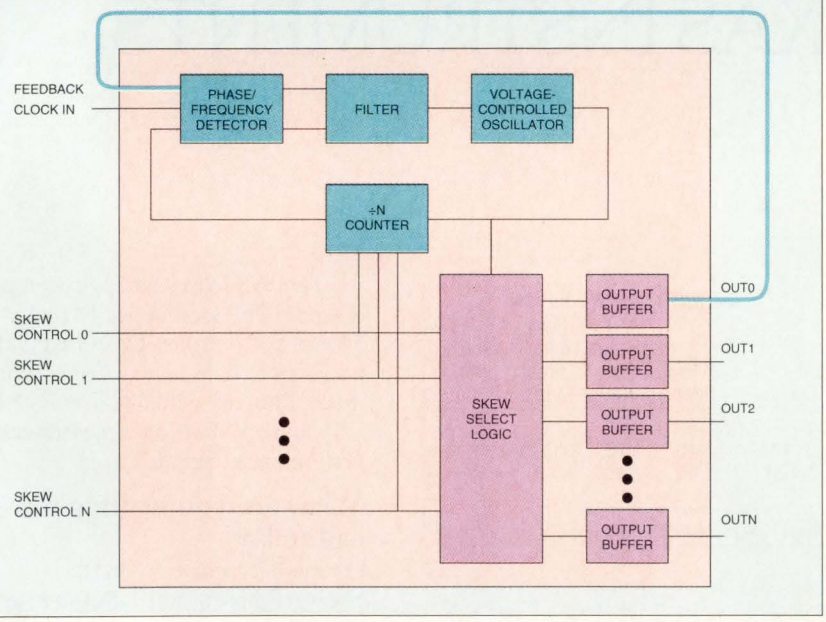

Clock distribution chips that let you control the relative skew of their various outputs consist of two parts. The first part (green) is a phase-locked loop that samples one of the outputs via an external link. This locks the output to the input, providing essentially zero propagation delay. It also supplies skew-shifting reference signals, either from running the PLL's voltage-controlled oscillator at a multiple of the input frequency or from tapping selected stages of a ring oscillator running at the clock frequency. The second part (purple) is a skew-control block that uses the signals from the first part to generate outputs that are multiples or submultiples of the source clock, or that are shifted in time by controlled phase differentials.

DOS, Windows *and* real time. At the same time.

The most widely-used application software runs on DOS and Windows. But they can't provide real-time capabilities for demanding industrial applications.

Microware's OS-9000 Real-Time Operating System with Virtual PC (**VPC**) lets you run DOS and Windows as tasks under OS-9000. Data and I/O devices can be shared between the two environments. DOS or Windows applications can be used to monitor and control real-time processes running under OS-9000. And **VPC** is included **FREE** with the OS-9000 Real-Time Operating System.

OS-9000 provides complete real-time development and run-time environments for hard real-time applications, including Microware's Ultra C advanced, high-technology ANSI C compiler. OS-9000 also offers extensive options for networking and graphics.

Do you want DOS, Windows and real time at the same time? Call our toll-free number to order your complete "plug-and-play" OS-9000 386/486 package for only $995. Or ask for your free copy of *Questions & Answers for Serious 386/486 PC Users.*

Call Microware Today!

1-800-475-9000

microware®

MICROWARE SYSTEMS CORPORATION

1900 N.W. 114th Street • Des Moines, Iowa 50325-7077
England/Benelux/Scandinavia: (44) 703 601990
Germany: (49) 6221-862091 • Switzerland: (41) 56-83-3377
France/Spain/Italy: (33) 42.58.63.00

CIRCLE NO. 24

INTEGRATED CIRCUITS

Config. no.	Pins		Output fed back	Output phase shift					
	S1	*S0*		*Q0*	*Q1*	*Q2*	*Q3*	*Q4*	*Q5*
1	0	0	Q0 . . . Q4	0	0	0	0	0	i
2	0	0	Q5	i	i	i	i	i	0
3	0	1	Q0, Q1, Q4	0	0	t	-t	0	i
4	0	1	Q2	-t	-t	0	-2t	-t	i-t
5	0	1	Q3	t	t	2t	0	t	i+t
6	0	1	Q5	i	i	i+t	i-t	i	0
7	1	0	Q0, Q2, Q3	0	t	0	0	-t	-2t
8	1	0	Q1	-t	0	-t	-t	-2t	-3t
9	1	0	Q4	t	2t	t	t	0	-t
10	1	0	Q5	2t	3t	2t	2t	t	0
11	1	1	Q0	0	t	t	-t	-t	-2t
12	1	1	Q1, Q2	-t	0	0	-2t	-2t	-3t
13	1	1	Q3, Q4	t	2t	2t	0	0	-t
14	1	1	Q5	2t	3t	3t	t	t	0

The combination of two bilevel inputs and one of six outputs fed back to the phase-locked loop produces an array of positive and negative skews and inversions.

outputs of the GA1110E operate at the clock fundamental, but you can use the part with TriQuint's GA1086 or GA1086E 10-output clock buffers, which have nine outputs at the clock frequency and one at half the clock frequency. Both the 1086 chips have very tight output-to-output skew variation—250 ps maximum, and 125 ps typical. The GA1086 specifies a propagation delay of 500 ps and the E-version specifies twice that.

Vitesse's 70-MHz VSL4485 and VSL4586 each have eight outputs. On two of the 4485's pins—or on six of the 4586's—you can multiply the input clock frequency by 2 or 4, or you can leave it at the fundamental. As with the TriQuint chip, combining two skew-control inputs with a choice of which output is fed back to the PLL gives a range of phase shifts at the various outputs. The actual value of the phase shift, relative to the output clock frequency, is determined by three other inputs that Vitesse calls divide inputs.

Like the TriQuint clock chips, the Vitesse chips guarantee 500 ps maximum skew between outputs. Typical values aren't currently specified. Because the Vitesse VCO operates at frequencies of up to 840 MHz, these chips provide more precise skew control than the TriQuint parts. The smallest possible increment is 1.25 ns.

Lots of control

Superficially, Cypress's 80-MHz CY7B991 (TTL I/O) and CY7B992 (CMOS I/O) resemble the TriQuint and Vitesse parts. There are several interesting differences, however. For one thing, the eight outputs of the two chips are arranged in pairs, and Cypress guarantees skew between members of the same pair to be 250 ps, while claiming that typical values are half that. For another, each pair of outputs has its own pair of skew-control pins, which may be pulled high or low or left unconnected. (An internal voltage divider prevents static buildup.)

There's an advantage to having a larger number of control pins. For each output pair, if the input pins are open, the output is in phase with the clock. Using other input values, you can trim the first two output pairs in increments of one, two, three, or four timing units, and you can trim the second two output pairs in increments of two, four and six timing units. The third pair can also be set for one-half or one-quarter of the applied clock frequency, and the fourth pair can be set to one-half the applied frequency or to the inverse of the applied clock.

Because of this simple control method, you can design a board using 991/2 chips to distribute clocks without paying a great deal of attention to trace lengths. Then you can use DIP switches on the skew-control inputs of the chips to trim edge arrival times on your prototype board while you observe edge placement on a scope. On production boards, you simply replace the switches with jumpers.

Lacking the raw speed capabilities of GaAs, Cypress implemented the VCO in its PLL as a ring oscillator. This device uses a cascade of gain stages to achieve the 180° of feedback an oscillator needs in order to work. Each stage introduces a constant increment of phase delay, and the signal tapped off any stage is delayed or advanced relative to a reference by an incremental amount. Instead of selecting a value, N, by which the VCO frequency will be divided, you select a number of delay stages, after which you will obtain your tap. In the Cypress parts the values are 16, 26 or 44. The precision, or degree of skew control, is determined by the clock period divided by the number of stages. The smallest increment of stage delay is 700 ps.

TriQuint and Vitesse unequivocally specify their output rise times: 1.5 ns (max) for 0.8 to 2 V. (For the 10-output clock distribution chips, it's 1.4 ns.) According to the manufacturers, these values are essential for Intel's Pentium P5 or 586 processor. Cypress's spec sheet cites a rather lackluster 3 ns for the TTL I/O part and 5 ns for the CMOS I/O part. According to the company, however, the parts actually slew at about 1 ns/V. ■

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

Cypress Semiconductor
(408) 943-2902 . . . . . **Circle 201**
TriQuint Semiconductor
(408) 982-0900 . . . . . **Circle 202**
Vitesse Semiconductor
(805) 388-7455 . . . . . **Circle 203**

Now You Can See True Colors Without Getting Soaked.

New full-featured RAMDACs from Brooktree deliver 24-bit true color for cost-conscious PC designers. Introducing four new RAMDACs that span the spectrum of PC applications, from 640x480 VGA systems to 1280x1024 workstation-quality graphics. They've got the right features, the right prices and are available right now.

True Color in a VGA Environment

Unplug Sierra. Plug in our totally compatible Bt481 or Bt482. You'll get 24-bit performance at 16-bit prices.

These new RAMDACs support 15-bit TARGA, 16-bit 5:6:5 and 24-bit true-color formats. They even allow you to switch between VGA and true color on a pixel-by-pixel basis.

Choose the Bt481 if you prefer an external hardware cursor. Or pick the Bt482 for its on-board 32x32x2 cursor — ideal for faster windowing environments.

Workstation Graphics at PC Prices

Introducing the Bt484 and Bt485, our newest mouth-watering RAMDACs. They combine true color with higher resolutions for higher performance graphics subsystems. And they're economical, too.

Bt485 operates at up to 135 MHz to drive pseudo color to 1280x1024 resolutions and gamma correct true color to 1024x768. It has a 64x64x2 cursor and all the serialization and timing to directly interface to a VRAM frame buffer.

Bt484 provides maximum flexibility with its programmable pixel port to provide 256 to 16.8 million colors, on-board 32x32x2 cursor and supports both interlaced and non-interlaced monitors.

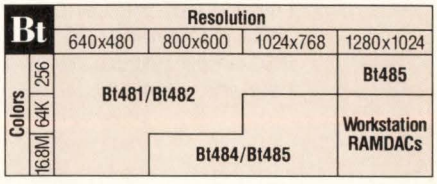

Call Brooktree at **1-800-VIDEO-IC** for technical details and pricing today.

Brooktree Corp., 9950 Barnes Canyon Road, San Diego, CA 92121, (619) 452-7580, FAX (619) 597-0673.

Brooktree®

RISC '93

MARCH 16-18, 1993
HYATT REGENCY SAN FRANCISCO AIRPORT
BURLINGAME, CALIFORNIA

WHAT IS RISC '93: RISC '93 is the premier annual gathering of hardware designers and software developers involved with RISC system architectures; RISC processors, ASIC cores and peripheral ICs; RISC programming and debugging; and the design of RISC-based systems for general-purpose computing, embedded and realtime applications.

THE ATTENDEES: RISC '93 is a down-to-earth conference aimed at both hardware and software engineers and engineering managers responsible for the design and development of products based on RISC processors. Specifically, RISC '93 is targeted at individuals with a solid background in electronic/computer system design and development (both hardware and software), as well as applications programming, who are looking for the most up-to-date information on the design of RISC-based computers, the design of embedded and realtime RISC systems, programming RISC-based systems, designing ASICs with RISC cores, and specifying and integrating RISC development tools into the overall product development environment.

THE FORMAT: TOTAL IMMERSION

RISC '93 has been designed as an intensive 3-day program that will place all participants in a total immersion environment where they'll have the opportunity to attend:

- Half-day tutorials that will help clarify the intricacies of the major RISC processor architectures.
- A variety of one-hour lectures that focus on specific aspects of RISC architectures, RISC-based hardware design, programming techniques, compilers, and development tools.
- Multipaper, application-focused sessions consisting of 20-minute presentations by system designers and software developers who've already implemented RISC designs in applications ranging from workstations, graphics, image processing, communications, signal processing and embedded control.
- A luncheon on each day of the conference that will be highlighted by an address from a recognized authority on RISC architecture or RISC-based designs.
- Demonstration workshops where attendees will be able to get hands-on exposure and detailed information about the latest RISC processors, compilers and development tools offered by leading hardware and software vendors.
- Free-wheeling "rap" sessions, led by RISC experts, on the first two evenings of the conference. Conference participants will have the chance to explore and exchange ideas about any aspect of RISC and RISC-based system design in an informal atmosphere of shirt sleeves, beer and pizza.

MORNING

THEORY TOOLS TECHNIQUES APPLICATIONS

LUNCHEON/SPEAKERS

AFTERNOON

THEORY TOOLS TECHNIQUES APPLICATIONS

DEMONSTRATION WORKSHOPS

DINNER

RAP SESSIONS LED BY RISC EXPERTS

RISC '93 Conference is sponsored by COMPUTER DESIGN ▪ One Technology Park Drive, Westford, MA 01886
TEL: 508-392-2124 ▪ 800-223-4259 ▪ FAX: 508-692-7780 ▪ Contact: Patti Kenney, **RISC '93** Conference Coordinator

COMPUTER DESIGN Technology and Design Directions

1993 INTERNATIONAL TECHNICAL CONFERENCE SERIES

PARTICIPATION: Suppliers of RISC processors, ASICs, peripheral ICs and memory, compilers, and development tools, as well as system designers and programmers working on RISC-based designs and applications, are invited to participate in RISC '93 by presenting tutorials, lectures, 20-minute application-focused papers or leading one of several evening "rap" sessions. A variety of topics that would be of interest to attendees is suggested in the Call for Papers for **RISC '93**. All presenters receive free admission of all tutorials, lectures, applications sessions, luncheons, rap sessions and a copy of the proceedings.

Suppliers of RISC hardware/software or development tools may also participate in **RISC '93** by sponsoring a demonstration workshop. These workshops will be two hours in length and will be the last formal sessions of the day. They're intended to provide attendees with an intimate, hands-on exposure to a supplier's product or service. The cost of sponsoring these sessions is $1500 for one two-hour workshop, $2700 for two workshops on successive days and $3500 for a workshop on each of the three days of the conference. One complimentary admission to **RISC '93** for each workshop sponsored is included in the fee.

RISC '93 is open to all system designers and software developers and the tuition fee covers admission to any tutorial, lecture, application session, demonstration workshop, evening rap session, and the three conference luncheons.

TUITION:

$495. Early Bird— Register by Dec.3, 1992

$595. If registered between Dec 4, 1992 - Jan 21, 1993

$695. After Jan 22, 1993

(Tuition includes lunch on all three days as well as snacks and refreshments during the rap sessions.)

Company group discounts available.

HOTEL ROOMS:

The single room rate at the Hyatt Regency is $108.
Call the hotel directly at 415- 347-1234
Be sure to tell them that you are attending **RISC '93.**

RISC '93

MAIL OR FAX TO:

Patti Kenney
RISC '93 Conference Coordinator
COMPUTER DESIGN
One Technology Park Drive
P.O. Box 990
Westford, MA 01886

FAX: 508-692-7780
TEL: 508-392-2124
800-223-4259

YES, I'M INTERESTED ! PLEASE SEND MORE INFORMATION

☐ **I AM A VENDOR I'd like to present a:**

☐ Tutorial ☐ One-hour lecture ☐ Application session

☐ I would like to lead a rap session

☐ My company would like to discuss sponsoring a workshop —

Please contact _____

☐ **I AM A POTENTIAL ATTENDEE:**

☐ I'm interested in attending, send me more details

☐ I'd also like to lead a rap session

☐ I am particularly interested in sessions on _____

NAME _____ TITLE _____

COMPANY _____

ADDRESS _____ M/S _____

CITY _____ STATE _____ ZIP _____

PHONE _____ FAX _____

CIRCLE NO. 25

INTEGRATED CIRCUITS

IEDM gets relevant

Stephan Ohr, Contributing Editor

This year, the IEEE conference organizers of the annual International Electron Devices Meeting (IEDM), held this month in San Francisco, have added a new theme. They're asking authors to respond to current economic conditions, and to assess what the costs of volume manufacturing might look like for the dramatic devices and processes they discuss. Coming from well-endowed and protected research laboratories, however, many of the best papers may still have a decidedly pie-in-the-sky feel to them.

One of the most attention-getting papers, for example, is a Matsushita presentation exploring how laser lithography might be used to produce 256-Mbit DRAMs, devices that even the authors acknowledge won't appear until 1997 or 1998. (At previous IEDMs, we should note, the electronics industry heard the first discussions and saw the first chip photos of 4-Mbit, and, later 16-Mbit DRAM architectures.)

Two current papers from NEC's Microelectronics Research Laboratories (Kanagawa, Japan), in fact, describe a capacitor structure that the authors, Hamada and Watanabe, suggest may be useful for 256-Mbit DRAMs. "This chip will be so advanced," says the promotional material for the conference, "it will be able to store 16 photograph-quality images, or an entire encyclopedia of text."

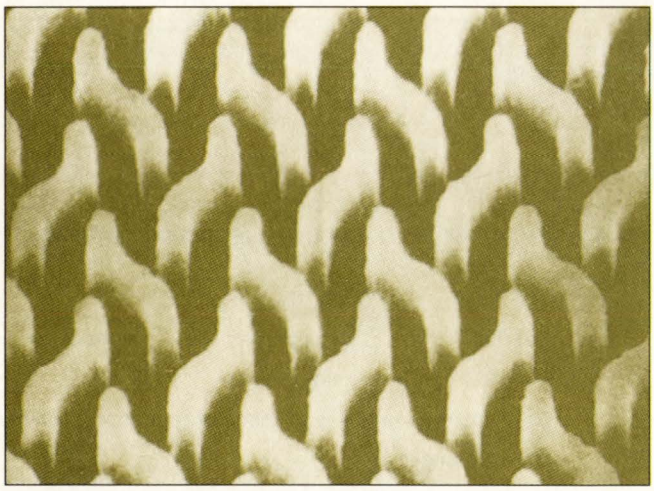

Researchers at Matsushita Industrial Electric Company have used excimer laser lithography to produce 0.25-μm isolation patterns, the geometry required for 256-Mbit DRAMs. As is the case with many IEDM presentations, it remains unclear whether the Matsushita researchers have hit upon a practical manufacturing method or a laboratory curiosity.

The chip will have a feature size of 0.25 μm, say researchers Endo, Hashimoto and Yamashita, all of Matsushita Industrial Electric Company (Osaka, Japan), in their abstract. Current semiconductor manufacturing relies on photolithography, but as the feature size of the devices decreases, the dimensions of the photomask—and even the wavelength of the light used to expose the photoresist—must also decrease.

The use of electron beams or X-rays, the authors believe, may be prohibitively expensive. An IBM/Toshiba/Siemens consortium, for example, has already committed to X-ray lithography to manufacture 64-Mbit DRAMs with 0.35-μm features, with the companies' investment expected to exceed $600 million by 1995. The investment in 0.25-μm, 256-Mbit DRAM manufacturing, is expected to exceed $1 billion by 1999. The Matsushita researchers feel that excimer laser operating in the frequency range of visible light are a much cheaper approach.

The excimer laser tested by Matsushita relies on krypton-fluorine and argon-fluorine, as well as a high-resolution stepping method, chemically amplified positive photoresists and a carefully controlled depth of focus for the laser. While the researchers have successfully produced 0.25-μm and even 0.20-μm dimensions using this process, it's unclear how difficult it will be to adapt the process to high-volume memory manufacturing. For example, should the laser light be projected through a photomask that mass produces the trenching in the photoresist layer? Or should the photoresist layer be scored directly by the swath of the laser beam? If the latter is the case, then it will take many minutes to score each chip. The process may be cheap, but it will be very time-consuming, and the excimer laser process may turn out to be as practical for high-volume manufacturing as E-beam writing proved to be.

Frank discussion needed

Of the 143 IEDM papers from commercial companies, IBM had the most accepted (15), followed by AT&T, Motorola and Japan's NEC (10 each). Universities contributed 72 papers, while eight came from government agencies. But it remains to be seen whether these papers offer practical manufacturing tips. "IEDM is traditionally seen as a showcase for new technologies leading to products that will hit the market three to five years down the road," says Hans Stork, manager of exploratory technology at IBM's T. J. Watson Research Center (Yorktown Heights, NY), "but IEDM 1992 will also feature some frank discussion about the direction of the industry and the research work driving it forward."

An evening panel discussion led by Lew Terman of IBM, for example, will examine the impact of slow growth and increased costs on technology research. Does it pay to be a leading-edge manufacturer? Or does it make more sense to focus on application-specific products with immediate sales potential?

This is a serious issue for the computer giant. William Bowles, director of IBM's OEM products marketing group and keynote speaker at *Computer Design*'s SysComp forum last February, used a Silicon Valley analog conference last month to announce the company's entrance into the market for mixed-signal ASICs.

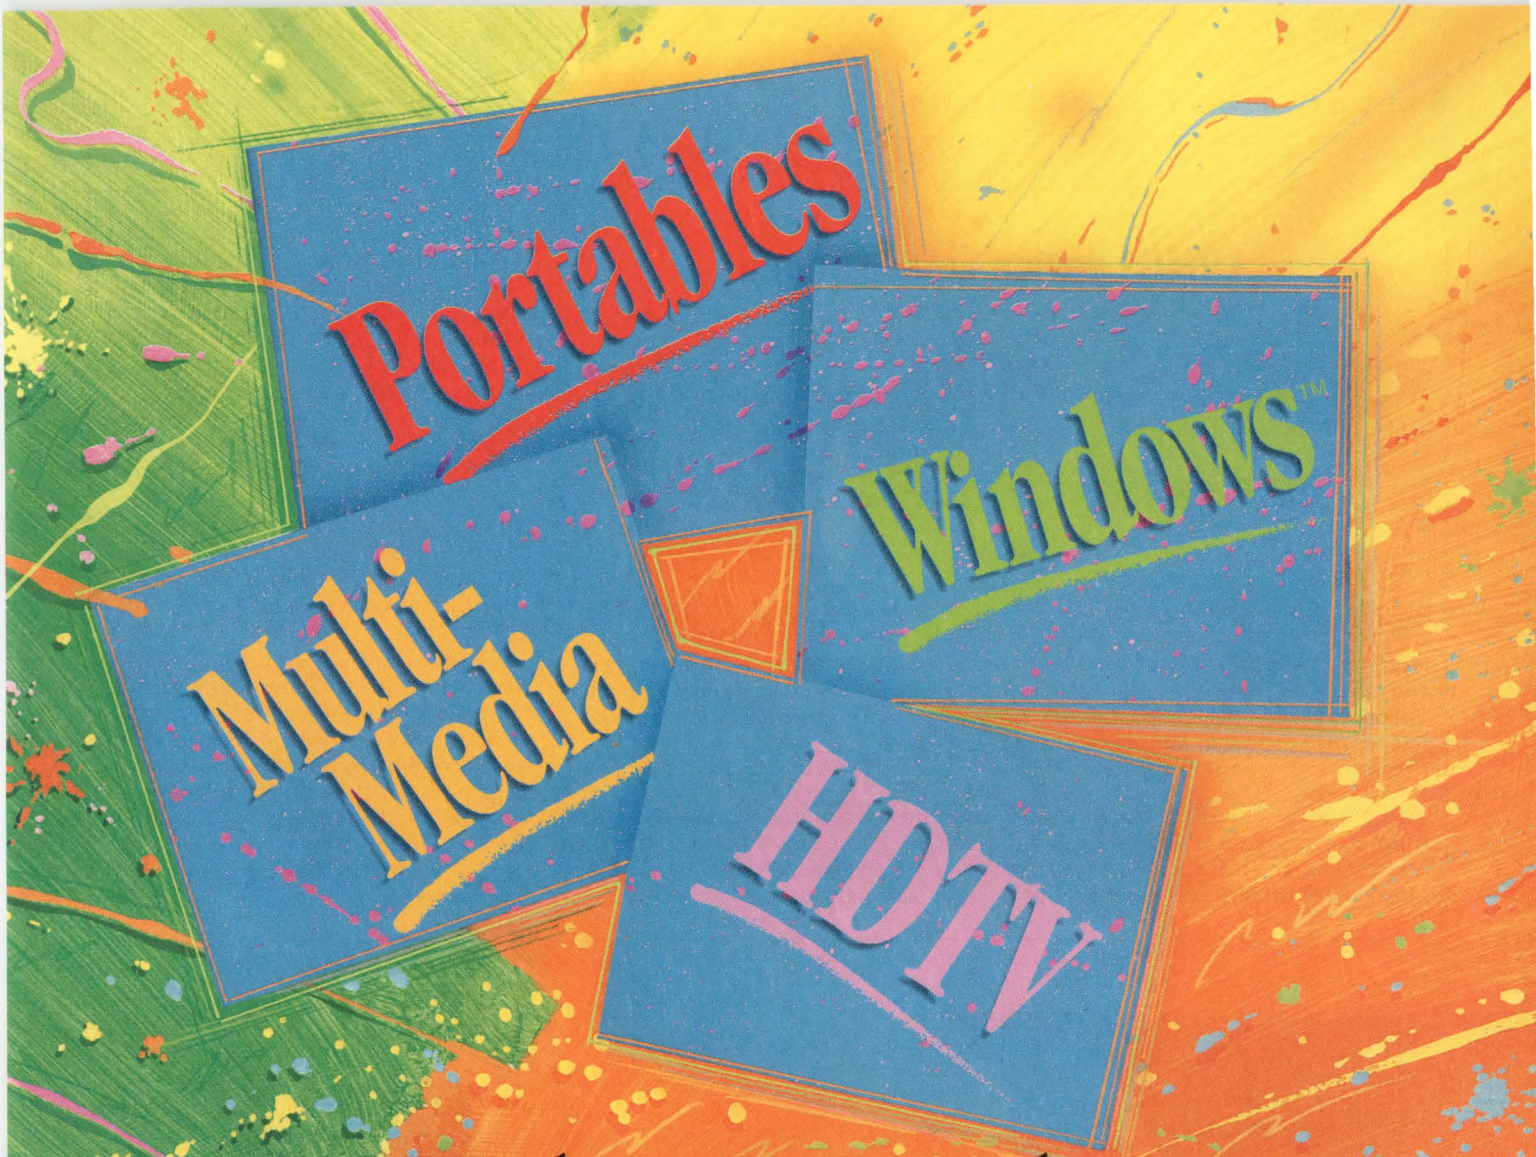

Graphic memory that lets your imagination run wild.

Memory Configuration	Features	Part Number	Access Time (ns)	Packaging Options	Availability
Triple Port DRAMs					
256K x 4	Fast Page Mode/Block Write	MT43C4257/8	80,100	SOJ	Now
128K x 8	Fast Page Mode/Block Write	MT43C8128/9	80,100	PLCC	Now
Dual Port DRAMs (VRAMs)					
256K x 4	Fast Page Mode/Block Write	MT42C4256*	70,80,100	ZIP, SOJ	Now
128K x 8	Fast Page Mode/Block Write	MT42C8128*	70,80,100	SOJ	Now
256K x 8	Extended Data Out/Block Write/Programmable Split	MT42C8256	70,80	SOJ, TSOP	Now
256K x 8	Fast Page Mode/Block Write	MT42C8255	70,80	SOJ, TSOP	Now
256K x 8	Fast Page Mode/Block Write/ Dual Write Enable	MT42C8254	70,80	SOJ, TSOP	Now
Specialty DRAMs					
256K x 16	Fast Page Mode	MT4C16256/7/8/9*	70,80	ZIP, SOJ, TSOP	4Q92

*Low power versions also available.

Windows™ is a trademark of Microsoft Corporation.

MICRON SEMICONDUCTOR, INC.

Technology that works for you.

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

CIRCLE NO. 26

INTEGRATED CIRCUITS

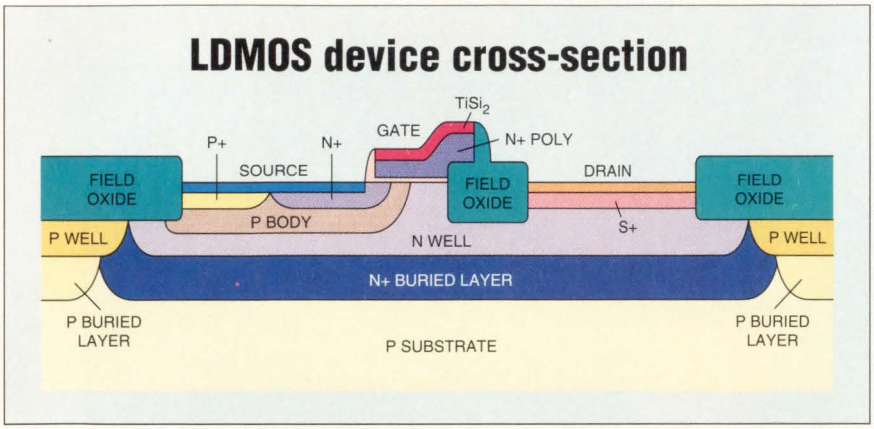

Motorola, in one of the most practical papers at IEDM, presents a way to integrate lateral double-diffused MOS on a BiCMOS substrate with 0.5-μm devices and maximum frequency (f_ts) on the order of 26 GHz. These smart power structures will contribute toward the portable RF transmitters of the near future.

The same week, the company told financial analysts it will close semiconductor manufacturing plants and lay off 40,000 workers. In light of these announcements, Lew Terman's panel may be more dramatic than a polite conversation among wealthy gentlemen.

"The concern about profitable manufacturing is also evident in a number of papers that describe cost-effective ways to produce new circuits," insists IBM's Stork, who heads the publicity effort for IEDM 1992. He cites as evidence an invited paper by Dr. Yoshio Nishi of Hewlett-Packard (Palo Alto, CA), a man who earlier led Toshiba's 1-Mbit DRAM development effort. The paper questions the conventional view that memory circuits are the best proving ground for advanced fabrication techniques for microprocessors and other digital circuits.

If anything, Nishi's paper, "ULSI technology toward the next century: driven by DRAMs or MPCs?," forecasts a split between the processes used to manufacture ultra-large-scale integrated (ULSI) circuits. One process will be used for memory, the other for RISC microprocessors. These chips require two fundamentally different technology drivers, says Dr.

Nishi. Where memories require sophistication in front-end fabrication, RISC microprocessors require expertise in back-end packaging techniques. It may be too costly to attempt both, Nishi warns.

Thought-provoking sessions

The plenary session on Monday, December 14th, of which Dr. Nishi's presentation is a part, may be the most thought-provoking event for the general engineering community. "semiconductor devices save the earth," by Yukinori Kuwano of Sanyo Electric (Tokyo, Japan), explores the possibilities of increasing energy and reducing pollution through the proliferation of semiconductor-based solar cells. In the same session, Alan Heeger of the Institute for Polymers and Organic Solids of the University of California (Berkeley, CA) will give a presentation demonstrating that "conducting polymers" (or electrically-conductive plastics) isn't an oxymoron, and suggesting applications for such materials.

A Tuesday, December 15th luncheon session will feature a talk on virtual reality by Jaron Lanier, chief scientist and founder of VPL Research (Foster City, CA). Also, check out a Tuesday evening panel session entitled "the electronics industry and the new world order: technology, Politics and Industrial Competition." Organized by the Berkeley Roundtable on the International Economy (BRIE), the panel will discuss the impact of international politics (for example, military conversion) and economic competition on the semiconductor industry.

In spite of its focus on relevance, IEDM will still feature its share of gee-whiz ideas that may indeed represent breakthroughs, if somebody can just figure out what to do with them. One of these gee-whizzes is a paper on transistors that emulate neurons in the human brain. Presented by K. Kotani of Tohoku University (Tokyo, Japan), "Neuron-MOS binary-logic circuits" describes ICs that embody real-world "fuzziness" by representing transistor inputs as weighted sums of multiple inputs. The authors claim that these fuzzy-logic ICs reduce the number of transistors required to implement machine vision, pattern learning, fault-tolerance, and other ambiguous decision-making systems.

There are several sessions, however, that are bound to have more immediate practical impact. The two one-day short courses, always popular with IEDM attendees according to the IEEE, will be held on Sunday, December 13th. One course, "Interconnect for the '90s," focuses on interconnection issues for on-chip, off-chip, module, and PCB systems. Interconnections are seen as the weak link when it comes to increasing performance. Presented by

INTEGRATED CIRCUITS

Brian Bakoglu of IBM, Bob Havemann of Texas Instruments (Dallas, TX) and Arjun Saxena of Rensselaer Polytechnic Institute (Troy, NY), the course examines the impacts of material, thermal effects and transmission-line behavior on packaging choices. The other short course, "Reliability: silicon to system considerations," promises a comprehensive look at reliability constraints on manufacturers of silicon devices. Given by Japanese and American industry experts, the course includes four presentations and will focus on the wear-out mechanisms of MOSFETs, including corrosion, electromigration and stress voiding effects.

"Transistor-level reliability" will be discussed by Akira Toriumi of Toshiba (Tokyo, Japan); "Reliability in multi-level interconnects" will be presented by Ronald Schutz of Sematech (Austin, TX), the U.S. chip-making consortium; "Reliability in packaging" will be covered by Paul Totta of IBM; and Daniel P. Siewiorek of Carnegie Mellon University (Pittsburg, PA) will be the instructor for the section on "fault-tolerant computing."

BiCMOS technology

Because it's so fashionable among designers and manufacturers of mixed-signal ASICs—even those involved in high-speed digital design—we should look at some of the papers in the IEDM session on BiCMOS technology. Once again, IEDM exposes us to the dramatic processes and capabilities we can only dream about using—combination CMOS and ECL devices with gate delays of 20 or 30 ps. D. Harame, E. Crabbe and other researchers at IBM (Yorktown Heights, NY) describe a silicon-germanium (SiGe) epitaxial base that produces 18.9-ps gate delays at 7.7 mW from the same substrate occupied by 0.25-μm-geometry CMOS devices. Their paper is called "A high-performance epitaxial SiGe-base ECL BiCMOS technology."

In the same session, T. Lui, G. Chin and other researchers at AT&T Bell Laboratories (Holmdel, Princeton and Murray Hill, NJ) show a 0.5-μm BiCMOS circuit with 31-ps ECL gate delays and 58-ps CMOS gate delays. Their paper is called "A Half-Micron Super Self-Aligned BiCMOS Technology for High-Speed Applications."

However, the most practical paper in the Monday session on BiCMOS (as well as one of the most dramatic) comes out of experience in smart power and RF circuits. Authors P. Tsui, P. Gilbert and S. Sun of Motorola (Austin, TX) describe a way to integrate 60-V lateral double-diffused MOS (LDMOS) on a BiCMOS substrate with 0.5-μm devices and f_ts on the order of 26 GHz. With luck, these structures will contribute to the portable RF transmitters of the near future. ■

Just What Your Customers Need, Another Outlet For Their Creativity.

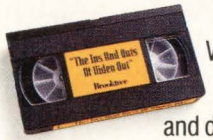

What's in? Video Out. Outputting video to a VCR and displaying video on a composite monitor are the newest capabilities every computer will need to compete in the Multimedia Age.

Now you're just a single chip away from adding Video Out to your very next computer design. Introducing Bt858, a monolithic digital device that packs in a board full of analog circuitry and puts out studio quality composite video.

Bt858 is a tweakless all-digital chip that bridges the video gap between RGB computers and composite or S-VHS outputs in the NTSC/PAL formats. It accepts multiformat digital inputs from 24, 16 or 15-bit RGB, 24 and 16-bit YCrCb and 8-bit VGA.

And because it has a programmable clock rate it adjusts for the 1:1 square pixels in computers and 4:3 rectangular pixels on TV without distortion.

Bt858 gives your system an image quality advantage, too. Studio quality output is a step above tape decks and TV monitors so images always look "first generation."

You've read the book. Now see the picture. Call **1-800-VIDEO IC** and we'll send you "The Ins and Outs of Video Out," a revealing presentation of Bt858's capabilities.

That's all folks.

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

Brooktree®

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

INTEGRATED CIRCUITS

Analog Devices courts designers with open architecture DSP

Stephan Ohr, Contributing Editor

The control and manipulation of analog signals is certainly among the most fruitful uses of DSP technology. But because many DSP components were initially architected as microcontrollers and required extensive programming or coding, traditional analog designers were among the slowest to grasp the potential of this burgeoning technology. Apart from programming difficulties, there is also the issue of price. Many frequency filtering and timing-window control functions can be performed much more cheaply with a fistful of inexpensive analog components.

Manufacturers of DSP components and development tools have tackled these problems head on. They're improving the performance and pushing down the cost of their components, while perfecting easier-to-use toolsets. And as a result they're witnessing an exponential growth in the number of applications that take advantage of DSP solutions.

But the conversion process is far from complete. A recent *Computer Design*/Indian Forest Research survey of analog and mixed-signal system designers found that far fewer than 50 percent were considering DSP solutions. Many analog designers seem resistant to using DSP, and chip-set vendors are taking new steps to convince this group that DSP isn't just some sort of fancy microprocessor.

Initiative proposed

Against the possibility that real-world system designers are finding it difficult to choose among competing DSP architectures and vendors, Analog Devices (ADI—Norwood, MA), an acknowledged leader in analog signal-conditioning and data-conversion technology, has proposed what it calls a Signal Computing Initiative. Like the CAD Framework Initiative talked about by EDA tool vendors, the Signal Computing Initiative calls for an open architecture that lets designers mix and match DSPs, signal-port chips, drivers, and operating systems. Ideally, you can pick and run off-the-shelf components and algorithms for speech recognition, audio and video image compression, and modem line conditioning.

"Front-end ports geared for audio or telephone can be paired with DSPs from Analog Devices or Motorola," says David D. French, Analog's vice-president and general manager. "Nothing is proprietary."

Anticipating critics who may suggest that the Signal Computing Initiative is a response to the dominance of competitors' components—specifically those of Texas Instruments (Houston, TX) and Motorola (Austin, TX)—Analog Devices announced an impressive series of design wins at DSPx, a new conference and trade show dedicated to digital signal processing, in October. Computer maker Olivetti (Ivrea, Italy and Cupertino, CA), for example, uses the Analog Devices ADSP-2111 processor and voice codecs to implement digital voice recording and playback on its Quaderno portable PC. Siemens (Munich, Germany) uses the ADSP-2111 as a speech recognition device for neural networks. Lernout & Hauspie (Ieper, Belgium) and VTech Systems (Hong Kong) have gotten together on a talking multilanguage dictionary using the Analog Devices ADSP-2105, and Digianswer of Denmark will OEM digital answering machines and boards using ADI components. All these manufacturers have endorsed the Signal Computing Initiative.

Still a custom business

Yet, the true impact of a Signal Computing Initiative may not be felt for some time. It isn't just that Analog Devices has a smaller number of design wins than its competitors, which makes it a less powerful lobbyist for open systems. The acceptance of the initiative may be limited more by the fact that the largest part of DSP design has been a customization effort requiring close coupling between silicon vendor and customer. It's only minimally a standard-parts business supported by off-the-shelf components and software.

According to Kun Lin, DSP marketing manager for Texas Instruments, the company with by far the

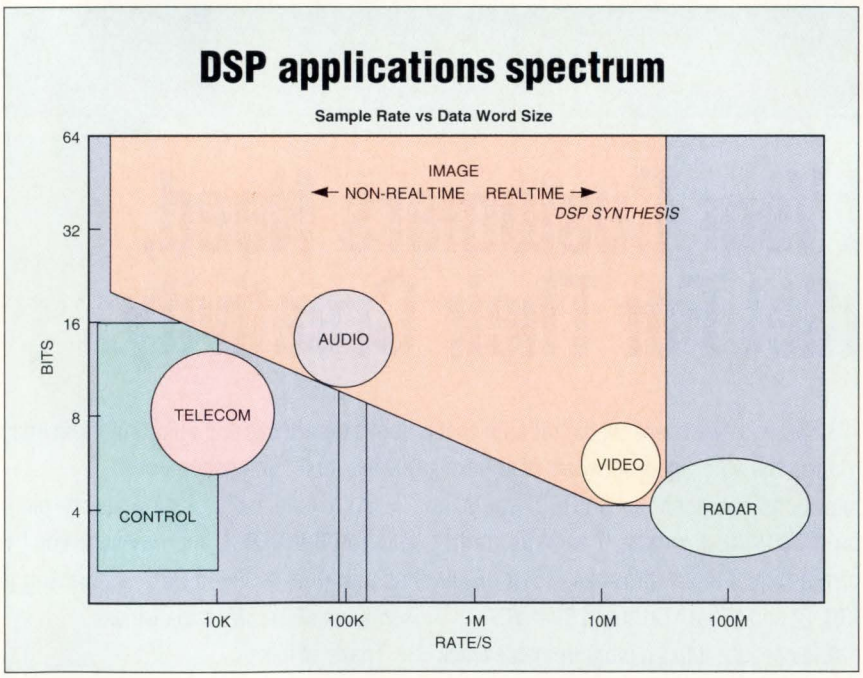

The increasing bandwidth and precision of DSP cores let them be utilized over a wide range of low- to high-end applications, says TI's Kun Lin. The majority of these designs, however, are custom projects rather than off-the-shelf drop-in replacements for analog components.

WHICH OPERATING SYSTEM DO YOU RELY ON EVERY DAY?

Windows

pSOSystem

Solaris

OS/2

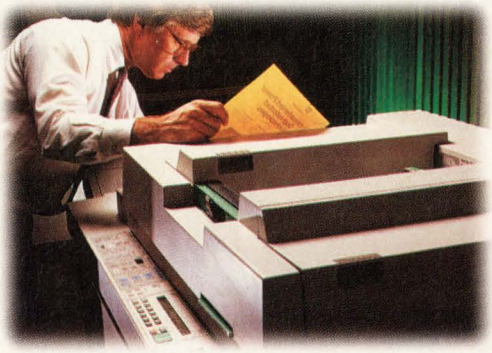

pSOSystem.

Did you know that pSOSystem is part of your everyday life? When you make a phone call; when you use a copy machine; when you fill your gas tank. In fact, pSOSystem is at the heart of over 2,000 embedded products — from air traffic control to patient-monitor systems.

From a reliable, compact kernel to the most sophisticated, UNIX-compatible networking and graphics capabilities, pSOSystem has the versatility to satisfy the widest range of requirements with the latest technology. C++; NFS client/server; RPCs; X Windows; transparent multiprocessing.

For a real-time operating system you can rely on, call 1-800-543-pSOS. You can also contact us by fax (408-980-0400) or e-mail (scg_sales@isi.com).

integrated systems

CIRCLE NO. 33

INTEGRATED CIRCUITS

largest proportion of DSP design wins, the cDSP product line has been the most aggressive in driving down the costs of DSP solutions. (The "c" in cDSP stands for "customizable.") "In volume," says Lin, "cDSP can be lower in cost than a boardful of separate analog and DSP components."

The cDSP line offers an ASIC methodology in which signal-conditioning components and A-D and D-A converters can be combined with C1X or C2X processor cores, versions of the TMS320 architecture, to provide parts with analog inputs and outputs. The increasing bandwidth and precision of the DSP cores, moreover, let them be used over a wide range of low- to high-end applications.

But these are hardly drop-in replacements for analog components. The specialized requirements of disk-drive manufacturers, for example, militate against the use of standard off-the-shelf components. Yet drive makers are gravitating toward DSP solutions to the problems of head positioning and spin control.

The DSP core provides specialized acceleration and deceleration algorithms that ensure greater tracking ability with tighter track densities and shorter seek times. The analog inputs make it easy to read embedded servo bursts, while the analog outputs drive power transistors and motor coils. The cost of an integrated DSP solution won't be competitive with the op amps and filters currently used for this application, but the increased precision, reliability and lack of drift mean that the overall cost of ownership is very reasonable. For this reason, DSP ASICs are finding their way into a wide variety of custom applications, including industrial motor controls and digital answering machines.

Zilog prefers DSP ASICs

Zilog (Campbell, CA) also supports a DSP ASIC methodology. Its Z89120 modem controller, for example, consists of an 8-bit microcontroller, a 16-bit DSP core and data converters geared toward pulse-width modulation. The part handles 9,600-bps modem, fax and voice interactions between a telephone line and a computer host.

Bryant Wilder, Motorola's DSP operations manager, sees little need to court analog designs with drop-in solutions to applications currently served by analog components. "There are many applications which will benefit from [the]...programmability...precision and reliability of a DSP solution," he says, "but it's always better to start with a clean sheet of paper." Although there are DSP chips that sell for less than $3, he adds that "You don't get a lot for $3." DSP becomes cost-effective in a design that was meant from its inception to offer high precision, reliability and programmability.

The improvements in price-performance of DSP components, and especially improved bandwidth, pave

NEW & NOW

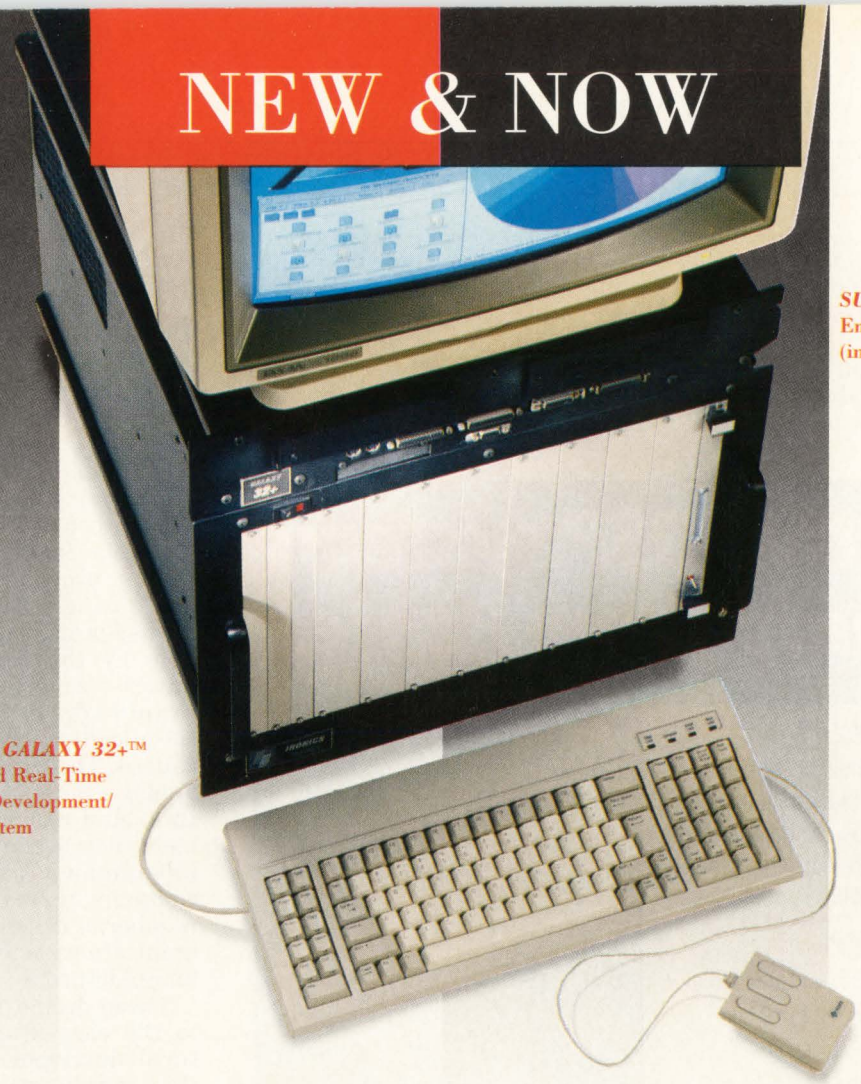

SUN SPARCengine™2
Embedded SunOS Host
(in Swing-Away housing)

IRONICS GALAXY 32+™
SunOS and Real-Time
VMEbus Development/
Target System

Real Sun. Real-Time. Real Tough!

The *GALAXY 32+* combines the heart of a Sun Workstation™ with the power of a VMEbus Real-Time Multiprocessing System. You get all the speed of a system where host and target share the same backplane...with none of the headaches.

A Complete Rugged Workstation — SunOS, color frame buffer and monitor, keyboard, and mouse PLUS:

40 MHz SPARCengine 2 — Sun's own powerplant within the *GALAXY 32+* ensures 100% SunOS compatibility.

Fast SBus to VMEbus Bridge — links and safely buffers SunOS from the real-time target debug process.

Real-Time VMEbus Target Options — a full line of compute, imaging, and I/O engines.

VxWorks™ or pSOS+™ Options — the best SunOS compatible real-time software environments.

20-Slot ICEBOX™ — the ultimate VME enclosure optimized for accessibility, power, and cooling.

Call our VME Product Hotline Today. **1-800-334-4812**

IRONICS Incorporated

IRONICS Incorporated
798 Cascadilla Street
Ithaca, New York 14850 USA

607-277-4060
Fax 607-272-5787

sun microsystems Value Added Reseller

CIRCLE NO. 30

INTEGRATED CIRCUITS

the way for applications that didn't previously exist. In addition to successes in digital cellular telephone applications in Europe (and, increasingly, in the U.S.), Wilder points to the use of 56000 products in Sony's newly introduced 5-¼-in. erasable optical disk products. The 56000 in these peripherals controls the focus and tracking of the optical laser. Because of their potential for drift with temperature, analog components could never provide the degree of control possible with a 24-bit DSP. "Analog runs out of gas," says Wilder, "but, with new DSP technologies, we have the opportunity to convert a $500 or a $1,000 system into a $50 or $100 board."

Success at Star

As it turns out, the success of Star Semiconductor (Warren, NJ) in winning over analog designers to DSP solutions is related more to its ability to customize parts for an application than it is to any low-cost drop-in replacement. While similar to the Motorola 56000 in capabilities (in that it provides 24-bit fixed-point processing), the architecture of Star's SprocChip is totally customizable by the user. Unlike microprocessors, which need to be programmed, SprocChips can be specifically crafted for analog applications. These chips are like PLDs for analog—the program is the architecture. And Star's PC-based development tools are intended to make the programming job easy. This start-up company now has an impressive number of design wins among communications and audio equipment manufacturers.

These design wins are due more to the ease-of-use of the architectural development tools, however, than to price or other factors. While the company is working toward a $25 part, current versions of the SprocChip are in the $100 range, and they're targeted toward the same applications that TI and Motorola say they can do for $10.

While DSP offers innovative solutions to the problems of interfacing electronics to the real world, the designer's preference for custom DSP may seem to negate the multivendor drop-in solutions proposed by Analog Devices. And while "cost-effective" is a part of the company's vocabulary, if analog design is to be synonymous with "cheap," DSP vendors still have a way to go. ■

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

Analog Devices	(617) 461-3881	**Circle 204**
Motorola	(512) 891-2030	**Circle 205**
Star Semiconductor	(703) 689-4400	**Circle 206**
Texas Instruments	(713) 274-2320	**Circle 207**
Zilog	(408) 370-8000	**Circle 208**

A 4X Improvement In High Density PLDs.

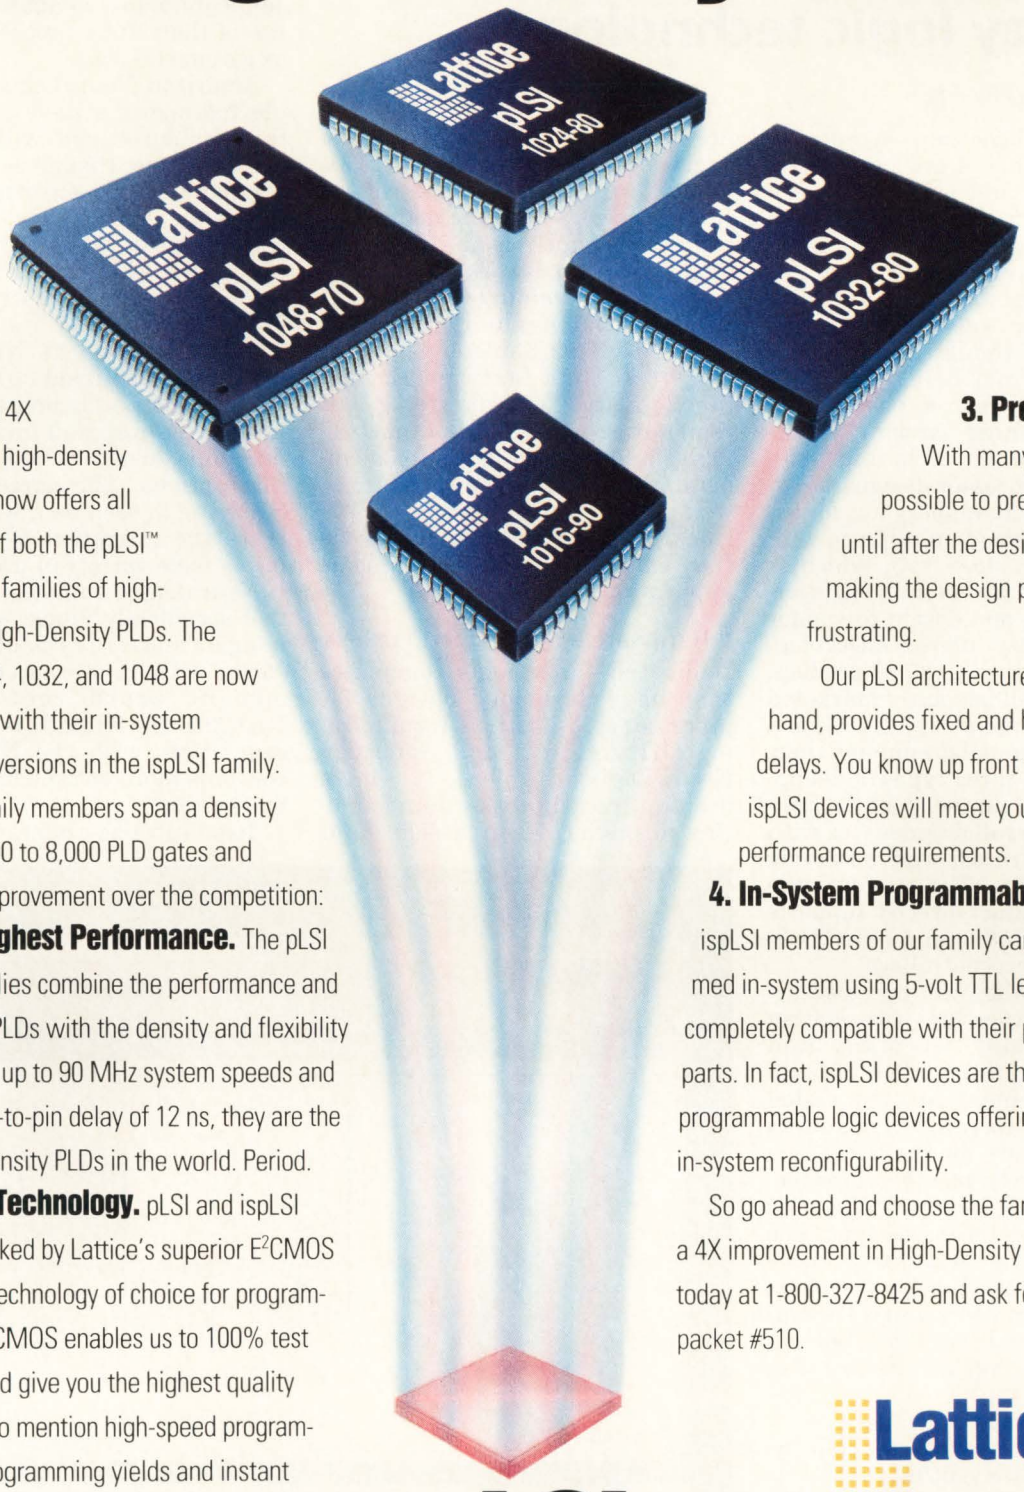

To achieve a 4X improvement in high-density design, Lattice now offers all four members of both the pLSI™ and the ispLSI™ families of high-performance, High-Density PLDs. The pLSI 1016, 1024, 1032, and 1048 are now available along with their in-system programmable versions in the ispLSI family.

The four family members span a density range from 2,000 to 8,000 PLD gates and provide a 4X improvement over the competition:

1. World's Highest Performance. The pLSI and ispLSI families combine the performance and ease of use of PLDs with the density and flexibility of FPGAs. With up to 90 MHz system speeds and a maximum pin-to-pin delay of 12 ns, they are the fastest High-Density PLDs in the world. Period.

2. E²CMOS® Technology. pLSI and ispLSI devices are backed by Lattice's superior E²CMOS process—the technology of choice for programmable logic. E²CMOS enables us to 100% test every device and give you the highest quality available. Not to mention high-speed programming, 100% programming yields and instant reprogrammability.

3. Predictable Delay. With many FPGAs it's impossible to predict performance until after the design is complete, making the design process long and frustrating.

Our pLSI architecture, on the other hand, provides fixed and highly predictable delays. You know up front that pLSI and ispLSI devices will meet your system performance requirements.

4. In-System Programmability (isp). The ispLSI members of our family can all be programmed in-system using 5-volt TTL levels, and are completely compatible with their pLSI counterparts. In fact, ispLSI devices are the world's only programmable logic devices offering non-volatile, in-system reconfigurability.

So go ahead and choose the family that offers a 4X improvement in High-Density PLDs. Call us today at 1-800-327-8425 and ask for information packet #510.

Lattice

Leader in E²CMOS PLDs.

pLSI and ispLSI.

CIRCLE NO. 29

SOFTWARE & DEVELOPMENT TOOLS

Alliances to speed acceptance of fuzzy logic technology

Tom Williams, Senior Editor

As fuzzy logic gains acceptance as a viable technology for embedded control, alliances are starting to form between large merchant semiconductor companies and smaller companies that were expressly founded to develop and market fuzzy logic software and hardware. Most of the latter companies are moving from being primarily consulting and custom engineering firms to becoming vendors of software development tools and dedicated fuzzy processor designs.

Motorola (Schaumberg, IL), for example, has formed an alliance with Aptronix (San Jose, CA), which is making the transition from engineering consulting and design to vendor of its new Fuzzy Inference Development Environment (FIDE) product. FIDE, which will be jointly marketed by the two companies, includes editors for fuzzy membership functions and rule sets, three types of debugging tools, a composer tool for linking modules created under FIDE with other C programs, and assembly code generators. The current version of the product directly supports Motorola's MC6805 and 68HC11 8- and 16-bit microcontrollers.

Intel (Chandler, AZ) has also arranged a partnership with Inform GmbH (Aachen, Germany and Evanston, IL). Inform offers a fuzzy logic development system called fuzzyTECH. FuzzyTECH supports graphical editors for the fuzzy rule base, as well as for membership functions, and it provides an interface for graphically simulating designs. The development tool generates C code and optimized assembler code for selected microprocessors. In collaboration with Intel, Inform is supplying a version of fuzzyTECH, Release 3.0, that directly supports Intel's 8xC196 line of 16-bit microcontrollers.

One of the unique features of fuzzyTECH is its online editing capability. With a serial line, you can modify the membership functions or rules of a running system through the graphic editor, and then you can directly observe the results in system behavior. Further, the Inform product supports the integration into designs of neural net technology. This feature may come in handy, given Intel's recent activity in neural net processors—such as its 80170NX analog neural network chip and PC-based neural training software.

For initial evaluation, Inform supplies a $199 Explorer version of fuzzyTECH that features limited functionality. You can create systems with 2-input and 1-output variables and up to 5 labels (or membership functions) per variable, as well as one rule block with 125 rules. The Explorer outputs C code that can be integrated into some applications.

Hardware/software partnerships

Dedicated fuzzy logic companies that have developed designs for processor chips, such as Togai Infralogic (Irvine, CA) and American NeuraLogix (Stanford, FL), have been contracting with silicon foundries to produce their designs. But both these companies have also signed licensing agreements with semiconductor companies for the use of their fuzzy processor designs as core technology.

American NeuraLogix, for example, has signed a 10-year technology transfer agreement with Samsung (Seoul, Korea) that gives the Korean company manufacturing and use rights to NeuraLogix's core chip technology. Central to that technology is the NLX230, a high-speed, low-cost 8-bit fuzzy microcontroller. The NLX230 and its newly enhanced version, the NLX231, are both capable of about 30 million rule evaluations per second and contain hardware-defined membership functions, fuzzifiers and defuzzification options. The agreement with Samsung is seen as a way to accelerate the migration of fuzzy logic into a wide variety of products.

NeuraLogix supplies a software tool specifically aimed at programming its processor products. The tool comes with an AT-compatible development board containing an NLX230. The NLX231 is downward-compatible with the NLX230 but contains more options and can hold more rules.

Togai Infralogic offers a full line of

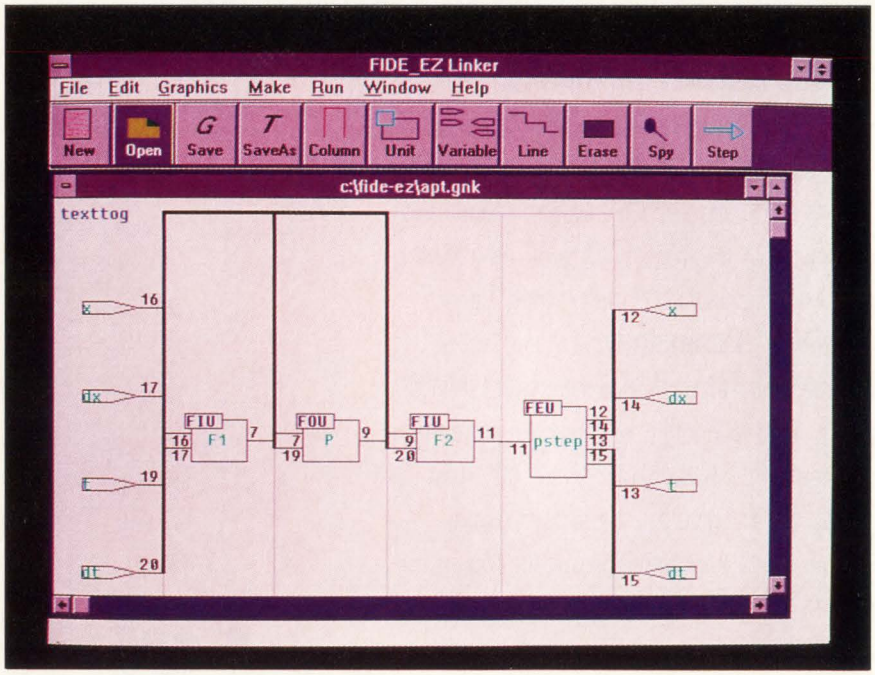

Aptronix's FIDE includes a facility called Composer, which can be used to graphically link fuzzy inference units (FIUs) to C code modules from existing or conventional applications. This process generates applications containing a combination of conventional and fuzzy logic code.

Learn the Only Embedded Debugger You Will Ever Need Without Turning the Page

TARGETS:
68xxx, 386/486, R2000/R3000/R4000, Sparc, 88000/88110, Gmicro, V810, more...

C • C++ • Fortran • Pascal

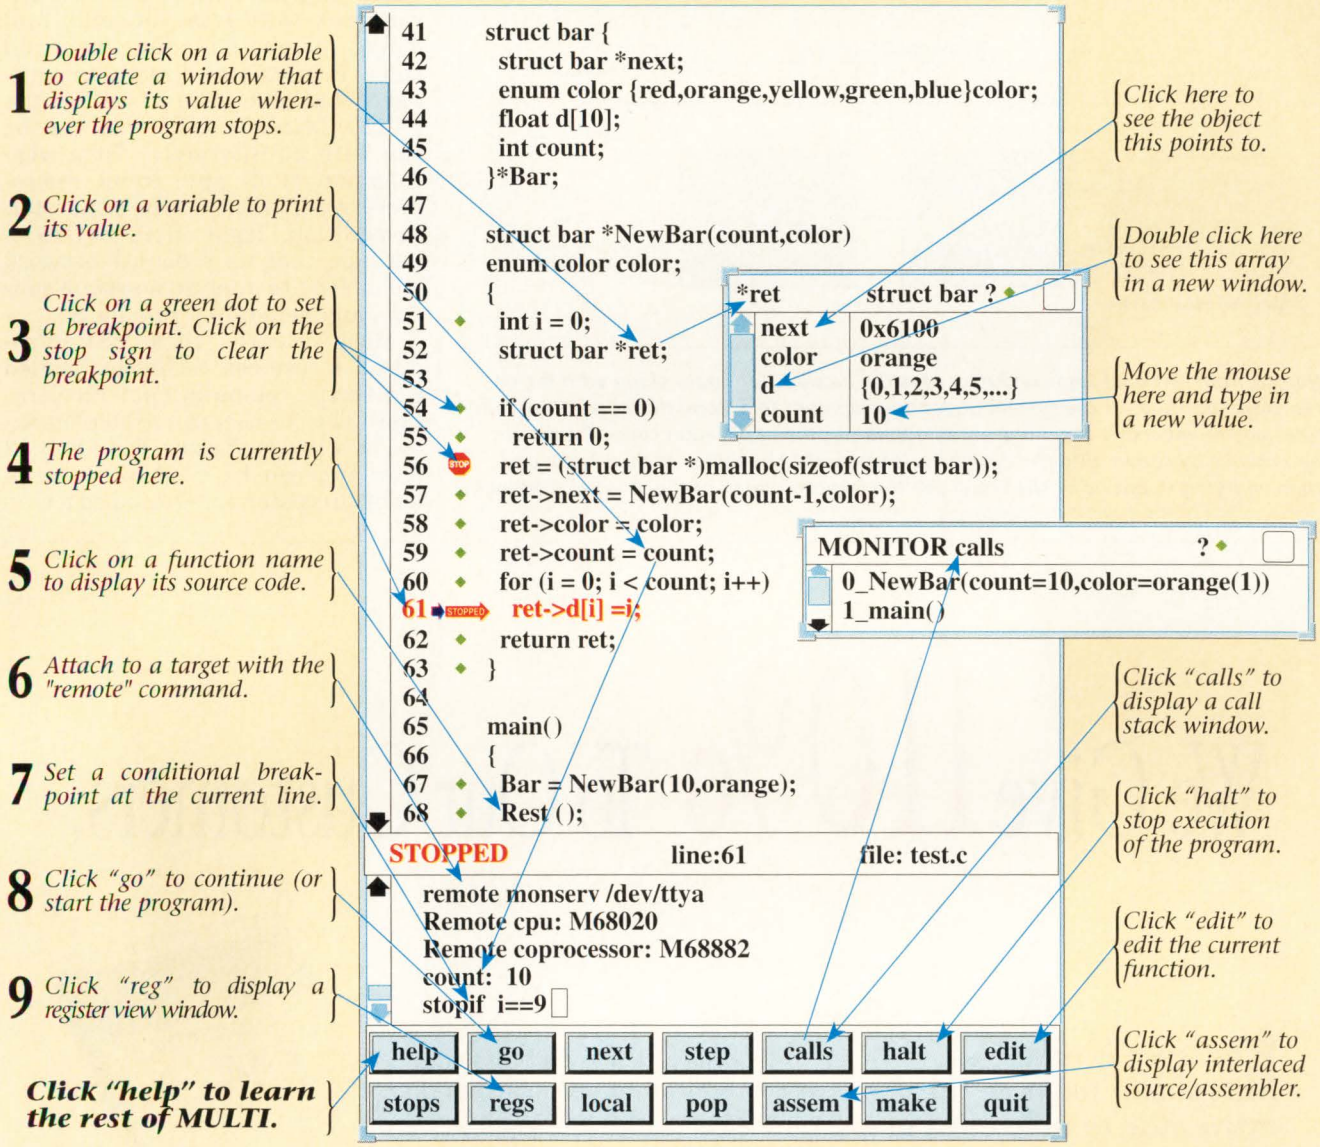

It's worth learning MULTI to fix one bug.

Hosts: Microsoft Windows, Sparc, Sun-3, DECstation, Unix/386, Motorola Delta, more...

Green Hills
• SOFTWARE, INC. •
510 Castillo, Santa Barbara, CA 93101
Tel: (805) 965-6044 • FAX: (805) 965-6343

Oasys
1 Cranberry Hill, Lexington, MA 02173
Tel: (617) 862-2002 • FAX: (617) 863-2633

CIRCLE NO. 34

SOFTWARE & DEVELOPMENT TOOLS

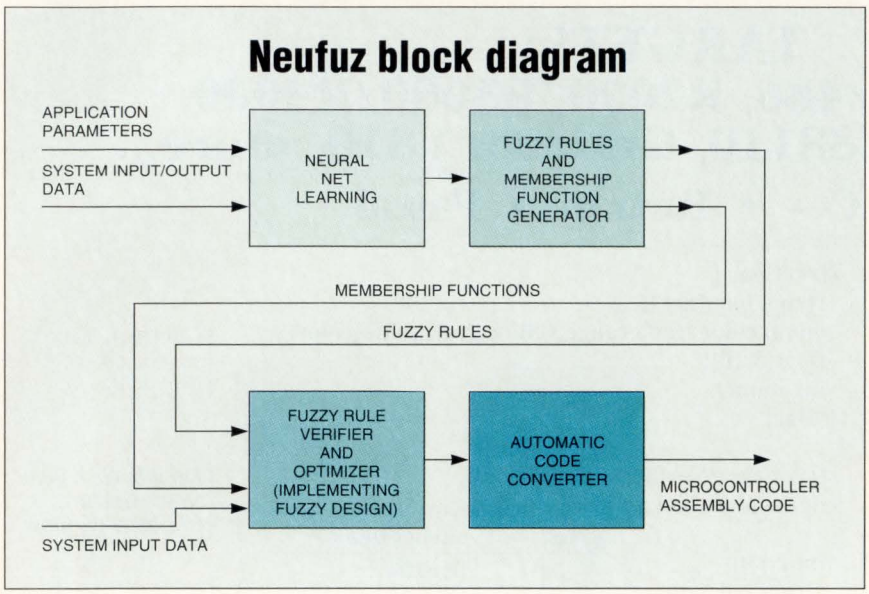

Neufuz from National Semiconductor uses the input/output data along with the parameters specified for the system under development to automatically generate fuzzy rules and membership functions—and, ultimately, the application code. The developer works by converging the design on the expected behavior of the system and then verifying it, but he or she won't see the intricate detail of the internal workings.

fuzzy logic software development tools and hardware. Its FC110 fuzzy processor is available as a chip or integrated on board-level products for AT, VME and Multibus. Togai also offers a line of software development tools, of which the centerpiece is TILShell. TILShell is a Windows-based graphical development environment that lets you edit and debug membership functions and rules. It interfaces directly to packages that generate code for the FC110, MicroFPL code that can be run with an interpreter for a wide selection of 8- and 16-bit microprocessors, and a C code generator. In addition, Togai offers TILGen, a package that uses neural network technology to analyze a system's inputs and generate a rule base.

Most recently, Togai has introduced a core cell technology called FCA (fuzzy computational acceleration). The FCA core can be implemented in a range of sizes, from 8 to 32 bits, and can be used as a standalone processor, integrated on a chip

DISTRIBUTED PROCESSING WITH THE QNX® FLEET™ NETWORK.

BECAUSE COMPUTERS WERE MEANT TO RUN TOGETHER.

Maybe it's not natural to expect a network of microcomputers to perform like a supercomputer.

With the QNX operating system, it's not natural to expect anything less.

POSIX AND MORE

Thousands of VARs and OEMs choose QNX for mission-critical applications – from POS to manufacturing to medical instrumentation.
And with good reasons.

Like POSIX compliance.
And *real* realtime performance.
And true microkernel architecture.
And message-passing IPC.
Not to mention our technical support and customer services.

Now add "FLEET" to the list.

FAULT-TOLERANT NETWORKING

With most networks,
a hardware failure spells disaster.
But not with QNX.
If a card or cable fails on a dual-net FLEET setup, QNX will automatically re-route data through the other network before you – or your application – can even blink.

LOAD-BALANCING ON THE FLY

For greater throughput, the FLEET network puts all available network hardware to work *at the same time.* And it will dynamically distribute the load by choosing the best route for the job.

EFFICIENT PERFORMANCE

FLEET uses network hardware to full advantage for maximum throughput. Whether you're running Ethernet for speed (application-level throughput at just under 1 Mbyte per second) or Arcnet for deterministic transactions – or both network cards in the same machine – you can count on FLEET for optimum efficiency.

QNX®

REALTIME OPERATING SYSTEM

EXTENSIBLE ARCHITECTURE

You can support new networks simply by adding new drivers. And you can start and stop drivers *dynamically,* without even rebooting.

TRANSPARENT DISTRIBUTED PROCESSING

In QNX there's no difference between local execution and network-remote execution. Which means you don't need to modify your applications in order to distribute them across the FLEET network.

To sum up: networking with QNX is fault-tolerant, load-balancing, efficient, extensible, and transparent. But it's a lot easier to just say "FLEET."

Go with the QNX FLEET network. Nothing runs like it.

To find out how your applications can thrive in the QNX environment, call **1-800-363-9001** (ext. 103).

ARCHITECTURE MAKES THE DIFFERENCE.™

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

SOFTWARE & DEVELOPMENT TOOLS

with a conventional processor core, or put on a chip with custom logic. Different mixes of rule-base and scratch-pad memories can be incorporated as well.

Togai has entered into two agreements, one with VLSI Technology (San Jose, CA) and another with Hitachi America (Brisbane, CA). The agreement with VLSI has resulted in the first implementation of a functional system block (FSB) for fuzzy logic applications. It's a 12-bit implementation of the FCA technology that's been dubbed the VY86C500, and it's capable of some 850,000 rule evaluations per second at 20 MHz. The fuzzy FSB can be combined with other microprocessor FSBs to form complete conventional/fuzzy processor units for embedded systems.

Togai's agreement with Hitachi America covers the use of fuzzy logic on conventional microcontrollers, primarily Hitachi's H8. It also includes support for training, documentation and hardware evaluation products, as well as for adapting software development tools to use with Hitachi processors.

Going it alone

The one U.S. semiconductor vendor that seems to be going it alone is National Semiconductor (Santa Clara, CA). National has set out on an ambitious project to develop fuzzy logic products simultaneously with neural network technology. It will soon be introducing a development tool called Neufuz. As this software technology matures, National plans to migrate it into silicon products.

Neufuz uses neural net learning to non-heuristically generate fuzzy rules and membership functions, at the same time using only the specifications of the system—that is, which outputs are expected from what inputs. The neural net learns by converging the inputs and outputs with the application parameters, after which its output is used to generate fuzzy rules and membership functions. These are run through a proprietary rule verifier and optimizer and eventually generate assembly code. ■

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

American NeuraLogix (407) 322-5608	Circle 209
Aptronix (408) 428-1888	Circle 210
Hitachi America (415) 589-8300	Circle 211
Inform GmbH (708) 866-1838	Circle 212
Intel (602) 554-2374	Circle 213
Motorola Microprocessor Division (708) 576-7000	Circle 214
National Semiconductor (408) 721-5000	Circle 215
Togai Infralogic (714) 975-8522	Circle 216
VLSI Technology (213) 931-0009	Circle 217

So you think you can design a better embedded computer than we can?

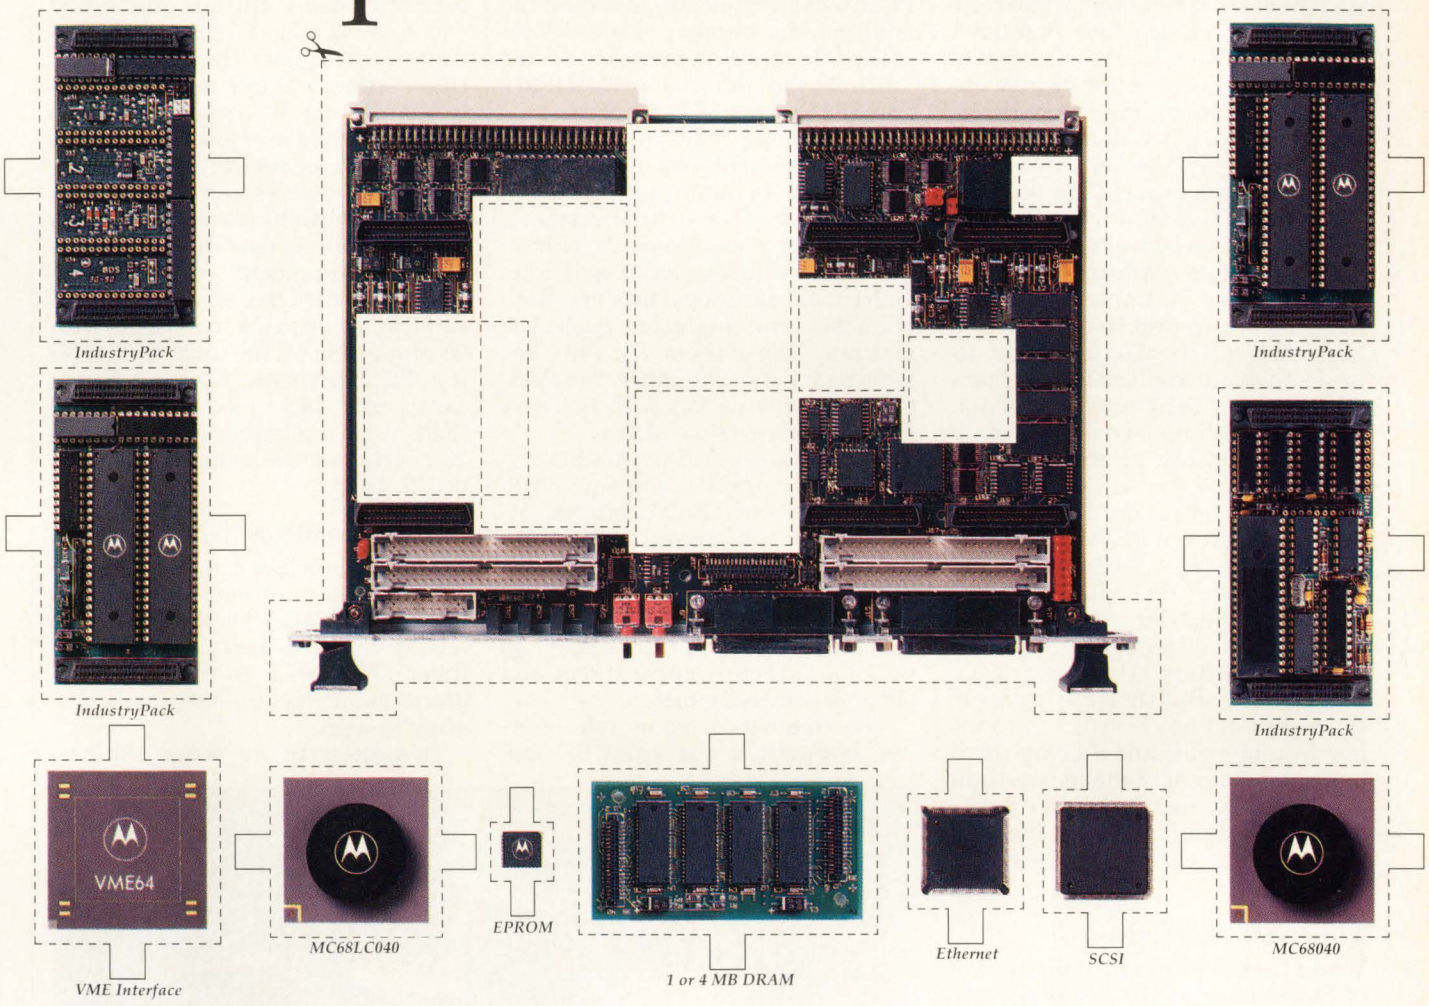

Of course you can. And we'll even help you. With the new MVME162 Embedded Computer.

The MVME162 is modular. Adaptable. So that you can select the exact functionality you want, without paying for features you don't need.

Now, would you like your MVME162 with a 68040 CPU (with floating point) or a 68LC040 CPU (no floating point)? VME bus or no bus? One, two, three, or four IndustryPack™ modules? Synchronous or Asynchronous serial ports? Ethernet or SCSI? 1MB DRAM or 4MB? 512K or 2MB of SRAM? And what other functions would you like with that?

Decisions, decisions. And they're all yours.

About the only thing that's not optional is Motorola's unmatched single board computer expertise, breakthrough engineering, outstanding performance and unwavering commitment to quality.

That's standard with every board.

With all that the MVME162 offers you, the easiest decision to make is to call the number below for a free information pack.

MOTOROLA
Computer Group

1-800-234-4VME

CIRCLE NO. 38

COMPUTERS & SUBSYSTEMS

HP debuts VME, realtime solutions

Warren Andrews, Senior Editor

Hewlett Packard (Colorado Springs, CO)has just joined the growing cadre of workstation makers bidding for the realtime industrial market. It follows closely on the heels of Digital Equipment Corporation, which announced its entry into the realtime market with a VMEbus-based approach and a POSIX 1003.4-compatible realtime operating system (RTOS). HP's announcement last month introduced a whole family of new products coming from its Measurement & Control Systems Division, a new unit that focuses on the needs of the factory floor, as well as control applications in manufacturing, aerospace, telecom, and commercial electronics. In addition to industrial workstations and boards, HP also debuted the HP-RT RTOS for its hardware.

At the center of HP's thrust is the company's PA-RISC 7100 processor. (PA stands for Precision Architecture.) The 50-MHz version of the chip set lets systems operate at 61 Mips and betters 60 SPECmarks in performance. But that's only part of the story—HP has boarded the VMEbus in a big way, announcing three major VMEbus board makers as complementary hardware vendors (CHVs) and signing on with the VMEbus trade association, VITA, as a senior member.

The VME alternative

While there's been much discussion of Futurebus+ as the latest-generation backplane bus for a variety of applications, including industrial automation, communications and aerospace, it's interesting that both HP and Digital have heavily supported VMEbus for industrial applications. Aside from the large number of vendors offering I/O for VME, recent activity within the trade group pushing the bus to higher levels of performance should keep it viable for a number of years.

First, the addition of VME64 doubled the effective bandwidth of the bus, and the possibility of using a source-synchronous transfer protocol holds promise for yet another twofold increase in performance. Now, it looks as if a viable live-insertion technology is available, and prototypes will soon be up and running. Future enhancements, including the addition of a number of serial lines to the P2 connector, promise to bring VME up to the performance level that's now the domain of Futurebus+.

While HP is firmly backing VME as its industrial platform, it still hasn't abandoned its other approaches. In announcing its realtime strategy, the company mentioned a pair of box-level workstations, as well as its 742rt VME board-level product. Both of its box-level products support the EISA bus. One of them, the 745i, has four EISA slots. The other, the 747i, offers a combination with two EISA slots and six VMEbus slots.

The two workstations are binary-compatible with the company's S700 and most current S800 systems, letting them run many applications currently supported on HP's UNIX platform, HP-UX. These applications include a variety of specialized manufacturing automation tasks, as well as database applications and office productivity tools.

The company's major VME offering, however, is the 742rt 6U VME board. Designed as a single-board system, the 742rt takes up two VMEbus slots and boasts one of the highest performance ratings of any VME CPU, topping off at 61 Mips. It differs from most competitive products in that it offers ECC memory (8 Mbytes standard, with 16, 32 and 64 Mbytes optional) in place of parity-protected memory. According to HP, this type of memory detects and corrects single and multibit errors, providing a more reliable system than a parity-protected approach.

The VME card has all I/O coming off its front, as opposed to systems with I/O connection off the card itself or off the P2 connector. Like most CPU cards, HP's 742rt has a variety of built-in I/O, including a pair of RS-232 ports, a parallel port, Ethernet, and SCSI-2.

HP-RT completes the picture

The 742rt includes a run-time license for HP's realtime operating system, HP-RT. Unlike Digital, which selected a third-party RTOS, Wind River's VxWorks, to modify for its processor, HP has built its own operating system.

Designed to be POSIX 1003.1-,

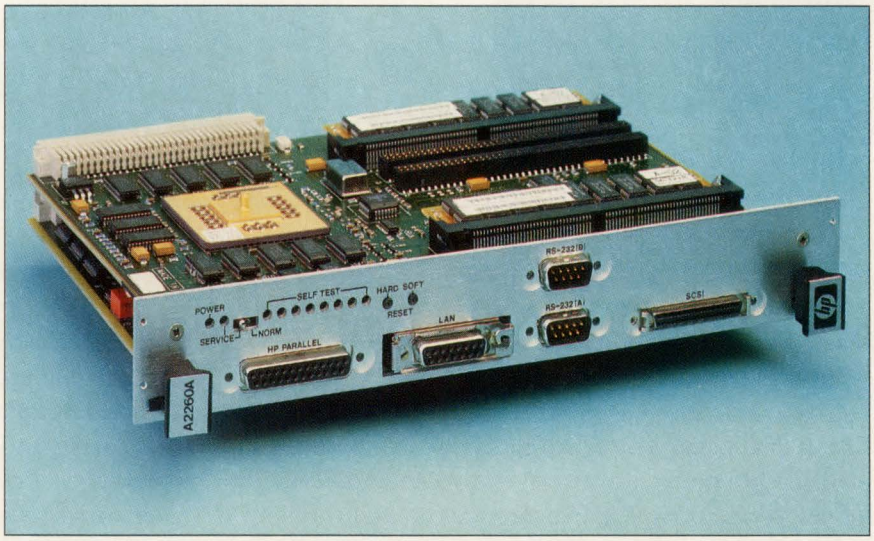

Featuring all-front-panel I/O, Hewlett Packard's HP9000 Series 700rt, model 742rt, is a 6U VMEbus card that occupies two VME slots. Its 50-MHz HP PA-RISC chip set gives it more than 60 Mips, one of the most powerful VME CPUs to date. It includes standard CPU features such as serial and parallel port, Ethernet and SCSI-2 ports, and comes with a run-time copy of HP's latest realtime operating system designed for the PA architecture, HP-RT.

INTRODUCING THE TP810V
TADPOLE'S HIGH-PERFORMANCE COMPUTING ENGINE FOR REAL-TIME APPLICATIONS

200 MIPS and VME-64

Real-Time Performance. No Compromises.

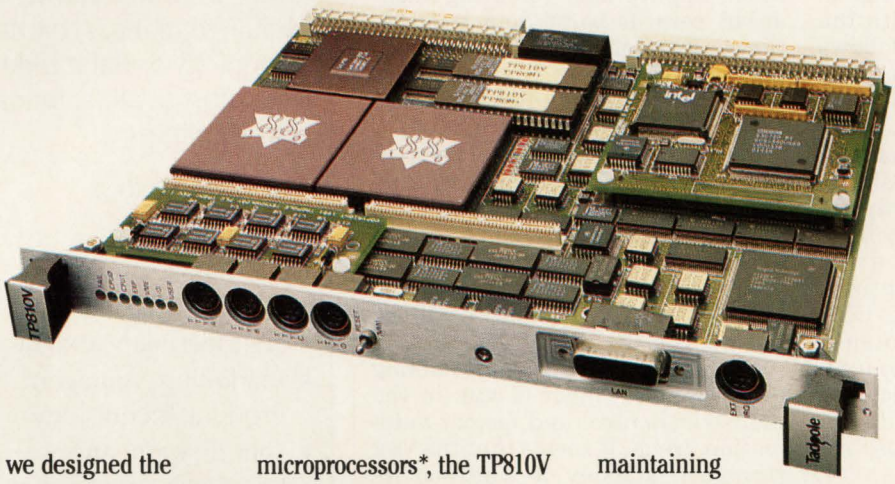

When we designed the TP810V, we approached it with total system performance in mind. With state-of-the-art microprocessors, high-speed memory, the latest generation I/O components, VME-64 and optimized real-time software, nothing gets in the way of information throughput on the TP810V.

Issuing up to 2 instructions on every clock, Motorola's 88110 Symmetric Superscalar™ RISC microprocessors provide the TP810V's horsepower. With dual 50 MHz 88110 microprocessors*, the TP810V achieves overall system performance of up to 200 MIPS/ 200 MFLOPS — that's an impressive $50/MIPS.

Memory options include 1 MB of high-speed synchronous SRAM with a sustained bandwidth of 228 MB/s and an expansion board with up to 128 MB DRAM. The TP810V accesses the additional DRAM at sustained on-board memory speeds of 133 MB/s.

And because cache coherency is essential for maintaining performance in multiprocessor systems, we've built in bus snooping support — even for the VME-64 interface.

The TP810V uses the latest in controller technology to eliminate common I/O bottlenecks: you can choose from two standard I/O modules with a 32-bit Ethernet controller and either one or two Fast and Wide SCSI-2 controllers. If neither of these options fit your system requirement, we'll design a custom module to your specifications.

To complete the system, the TP810V supports Integrated Systems' pSOSystem™, the modular real-time operating system with transparent multiprocessing capability. And with high-level language compilers readily available for the 88110 microprocessor, you can shorten your product development cycle without sacrificing performance.

If you're looking for a real-time computing engine with superb overall system performance at a great price, take a look at the TP810V. We think you'll like what you see.

Call Today For A Free Brochure

800-232-6656

FAX 512-219-2222

*Single processor TP810V models are also available.

Computing Without Compromise

T A D P O L E

CIRCLE NO. 39

1003.4- and 1003.4a-compatible, HP-RT is designed from the ground up to provide hard realtime capability tuned to the HP PA-RISC platform. A native POSIX application programming interface (API) is implemented for system calls, realtime extensions and process threads. The OS also incorporates some of the best UNIX features, including protected address spaces, multiprocessing and graphical user interfaces, into HP-RT.

In addition, HP-RT can be scaled to balance memory and performance requirements. With a small kernel, overhead is kept to a minimum, and optional services can be invoked as required. The OS complies with the POSIX 1003.1 standard, and it follows the POSIX 1003.4 draft 9/10 for realtime extensions, as well as the POSIX 1003.4a draft 3/4 for process-level threads. POSIX compliance has remained an elusive thing for realtime extensions because of the rapidly changing drafts. Some vendors have elected to use a particular draft, such as draft 4, rather than continue to shoot at a moving target. (The current version, which is undergoing the approval-or-rejection cycle, is draft 10.) The HP-RT includes many SVID/BSD commands and supports C, ANSI C, C++, and PA-RISC assembly.

Third parties party on

While both HP's box-level workstation products have at least some EISA capability, the company's support of VMEbus is emphasized with an industry-leading CPU, membership in industry trade groups and its third-party agreements. The latter include arrangements with the CMC Network Products Division of Rockwell International (Santa Barbara, CA), SBE (Concord, CA) and VMIC (VME Microsystems International—Huntsville, AL).

According to Hewlett Packard, CMC has been selected as a CHV to provide Ethernet and FDDI VMEbus local-area network connectivity for its realtime industrial workstations. CMC plans to work closely with HP to supply the necessary drivers for both HP's realtime and UNIX operating systems.

SBE will be offering HP's OEMs and integrators its eight-port VCOM-34 X.25 controller, with eight high-speed serial ports in a single VMEbus slot. Its eight full-duplex, independently programmable serial channels can transmit and receive asynchronous, X.25-compatible HDLC and bisynchronous protocol data at E1 (2.048 Mbps) rates. SBE's VCOM-34 card includes an on-board 68030 processor to ease the burden of the HP PA-RISC processor.

The company will also provide its VPU-25 intelligent I/O controller to handle a variety of I/O, memory and custom options. The board is meant for OEMs requiring a high-performance interface between VME systems such as a mini/super minicomputer and wide-area networks or other applications requiring high-speed point-to-point data transfer. Serial interface modules let each port be separately configured for EIA-232-C, EIA-422, EIA-449, EIA-530, X.21/V.11, or V.35 standards.

VMIC will offer a variety of I/O products, including host adapters, VME-to-VME links, repeaters, reflective memory, digital and analog I/O and synchro/resolver I/O. Transition panels, power supplies, chassis, and other supporting products will be offered as well.

With both HP and Digital making firm stands in favor of VME in the industrial-control and factory-automation arena, it seems that the VME standard is likely to continue its dominance in this area. With the inclusion of HP's PA-RISC, every major RISC architecture now has strong VME support—with the exception of IBM's RS/6000. SPARC is represented by, among others, Force, Themis and Ironics. ■

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

CMC
(805) 968-4262 ..... Circle 218
Hewlett Packard (U.S.)
(800) 637-7740 ..... Circle 219
Hewlett Packard (Canada)
(800) 387-3867 ..... Circle 220
SBE
(800) 347-2666 ..... Circle 221
VMIC
(800) 322-3616 ..... Circle 222

Live Insertion Now for Any VME Board!

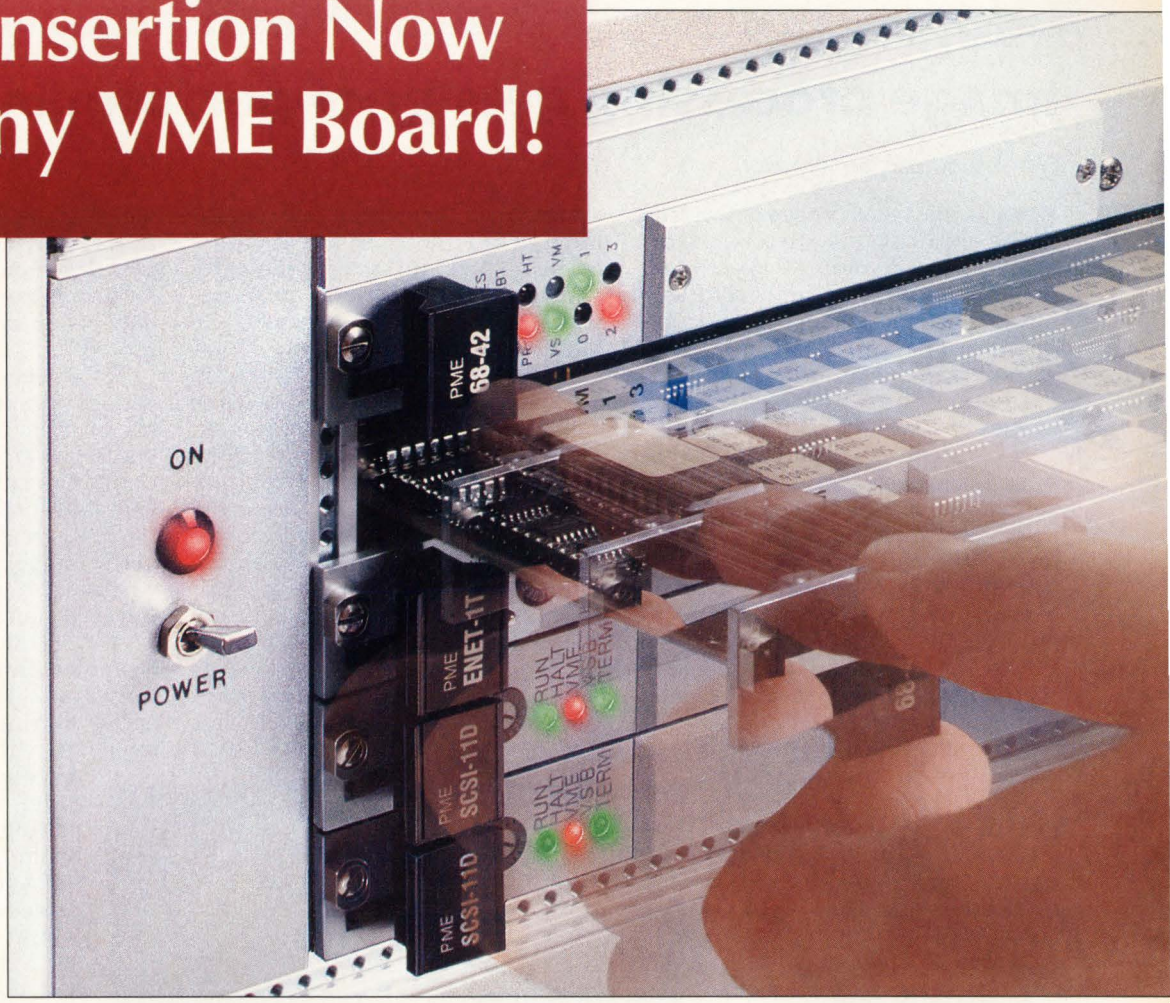

The Radstone HS-1 provides an immediate solution to all your VME Hot Swap needs.

Count on Radstone to deliver real world solutions to all your VME system needs. Need to insert or remove a board without powering down the system? The HS-1 is ready today to add this live insertion capability to your existing VME boards. No smoke, mirrors or complications. Just a clean, clever solution.

- Works with any standard VME board and backplane
- Lets you use off-the-shelf hardware to configure your system
- Provides controlled ramp up/down of power rails
- Assures complete isolation of the VMEbus from the board during power up/down sequences
- Automatically links daisy chain signals when a board is absent
- Sub-nanosecond delay times for all pertinent VMEbus signals

Radstone's done it again... this time with live insertion. Let us do it for you. Call or write for the details.

Radstone Technology Corporation
20 Craig Road, Montvale, NJ 07645-1737
Call Toll-Free: (800) 368-2738
Eastern Region: (201) 391-2700
Central Region: (708) 304-0202
Western Region: (408) 727-4795

CIRCLE NO. 40

BUSCON/93-WEST
See us in booth #826

RADSTONE TECHNOLOGY

COMPUTERS & SUBSYSTEMS

Image processing gets price-performance boost

Warren Andrews, Senior Editor

Because most image processing applications are too narrow for sales volume to reach even a minimum payback level, solutions have remained expensive and proprietary, or they've been composed of large and bulky collections of standard and custom boards assembled to satisfy a specific application.

Now, next-generation hardware from Datacube (Danvers, MA) takes advantage of finely tuned ASIC technology to cut size and cost while significantly boosting performance. At the same time, the modular architecture of the hardware, based on a technology called VSIM (Virtual Surface Image Memory), lets it accommodate a broad variety of applications while using relatively standard, low-cost components. VSIM is similar in approach to the company's MaxVideo 20 board, but the new module has a far larger memory capacity, as well as on-module processing circuitry.

At the heart of each module is a 225,000-gate ASIC Datacube calls the D52. Not only does it manage image pipelines and perform many image processing tasks, but it lets the system use relatively inexpensive standard DRAMs in place of more costly multiported video RAMs. Each module can support multiple 40-MHz image pipelines and virtual memory, and each includes integral ALU, crosspoint, look-up table, and statistics.

Development and target

The modules will be the heart of a new MaxVideo VMEbus board that Datacube will introduce early next year. A follow-on to the MaxVideo 20, it will be known as the MaxVideo 200. Datacube has already released the product's development platform, the MaxTD, which includes the new board, and will be available this month.

The development platform comprises a 5-slot VMEbus backplane; a Motorola MVME167 68040-based CPU; a MaxVideo 200 with a full complement of modules and 24 Mbytes of on-board virtual image memory; and Lynx's realtime operating system, LynxOS, with X-Windows and Motif; and a SCSI hard disk, all installed in a compact enclosure. In addition, MaxTD provides a tape backup unit, keyboard and mouse. The MVME167 supplies many of the housekeeping functions and provides industry-standard serial and parallel ports, as well as support for Ethernet, SCSI and video I/O. MaxTD is designed to be a superset of whatever target system a particular application calls for.

"We selected LynxOS," says Datacube marketing services manager David Wright, "because it has the look and feel of SunOS, with which many developers are familiar. It also provides modular services so that developers can select as little as a basic kernel, or the entire OS, or only as many services as are required."

The object of the development/target environment is to let you assemble an application using the hardware platform along with the company's software tools, which include an interactive graphical user interface and a library of accelerated image-processing functions called SILL (Standard Imaging Layered Library).

DIGIT (Datacube's Interactive Graphical Imaging Tool) is one of the software tools, comprising an interactive X-Windows- and Motif-based application that helps you quickly develop image-processing code. It eliminates lengthy editing and compiling cycles by letting you simply point and click on desired image-processing functions.

And when combined with SILL, DIGIT helps you develop and check high-performance C-callable functions and complete applications in an interactive environment without having to recompile at each stage. SILL and DIGIT are layered on top of, and are compatible with, Datacube's ImageFlow language for pipelined processors.

The object of MaxTD is to supply a full, rich development environment from which components can be added to an economical, high-performance target system. Target systems can be ROM-, disk- or network-based and can range from simple systems with few operating services and peripheral devices to complex hardware systems with all OS features.

The VSIM and the D52 at its heart are the brainchild of Datacube's principal design engineer, Shep Siegel, who says "It's the first 40-MHz image processor on a single chip." The ASIC not only speeds data on and off the new module, but also provides a large virtual memory and integral on-chip processing functions, permitting the use of relatively inexpensive DRAM.

ASIC is key

The third generation of Datacube's MaxVideo technology, the D52 and its associated VSIM technology ad-

OrCAD: the next chapter

As designs become more complex and time-to-market decreases, new tools are needed to keep up with the technology.

OrCAD's new line of products offer the same user interface that has made them the world's most popular EDA tools, plus unprecedented speed and capacity.

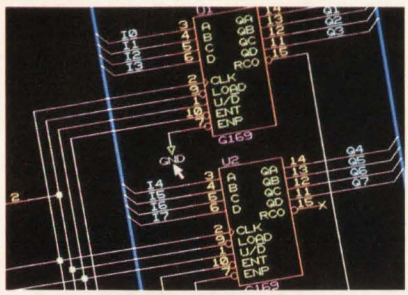

Schematic Design Tools 386+

Easy enough to get a single "A" size design done quickly, powerful enough to produce complex, 200 layer hierarchical designs. SDT 386+, designed to take advantage of 32 bit data structures and addressing on 386/486 computers, includes over 20,000 unique library parts in 60+ libraries, 30+ netlist formats, and much more.

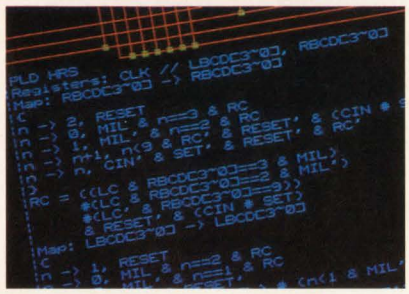

Programmable Logic Design Tools 386+

The best tool on the market for writing logic for programmable chips (PLAs, PLDs, GALs, etc.). The compiler accepts seven forms of input including Boolean equations, truth tables, state machine procedures, logic synthesis; even schematics from Schematic Design Tools 386+. Output can be accepted into OrCAD's digital simulator for accurate simulation of designs with multiple programmed PLDs.

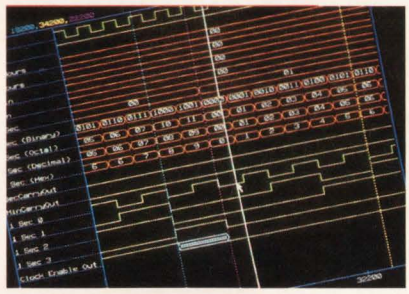

Digital Simulation Tools 386+

This timing based, 12 state, event driven simulator with support utilities allows easy simulation of digital designs of up to 200,000 gates. The graphical interface will be familiar to any designer who has worked with a logic probe. Thanks to OrCAD's intuitive ESP Framework, Digital Simulation tools 386+ works seamlessly with other OrCAD products.

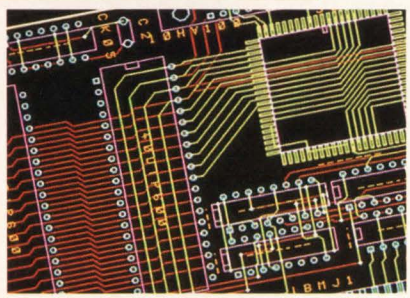

OrCAD PCB II and MASSTECK

OrCAD has joined forces with MASSTECK to bring you a powerful solution at an affordable price. Printed circuit board placement and routing can now be performed on a 386/486 PC with outstanding results.

Call or write today for a FREE demo disk and other information

OrCAD

The Better Solution

(503) 690-9881 • Fax (503) 690-9891

Europe 33-1-45 75 50 00 • Fax 33-1-45 77 82 89

CIRCLE NO. 41

COMPUTERS & SUBSYSTEMS

vance the company's resolution-independent, region-of-interest (ROI) concept, providing faster pipelines, a virtual memory and the capability to perform a variety of processing steps on each stage of the pipeline.

"The D52 is fabricated on a 225,000-gate sea-of-gates array in 0.8-µm CMOS capable of handling input, processing and output at 40 MHz," says the company's principal design engineer, Shep Siegel. "It's the first 40-MHz image processor on a single chip." With the capability to handle multiple frame memories, and with its own ALU, LUT and other features, it can off-load many functions from other parts of the system to permit a more compact implementation of complex systems.

The flexible architecture lets multiple VSIMs be wired in parallel, resulting in even greater pipeline bandwidths. Because much of the processing can now take place in the 40-MHz pipeline of the D52 chip, it's possible to provide functions that are difficult, if not impossible, to implement at the board level, because of buffer delays and interconnect capacitance.

By implementing virtual memory, the system can support image arrays that are larger than the amount of physical memory available, and so it can easily handle complex applications. The D52 can address a virtual memory space of up to 96 Mbytes. In addition, VSIM modules can handle 40-Mbyte/s block transfers into and out of the module simultaneously.

Diverse applications

Datacube has found wide-ranging applications for its technology, from image-enhancing systems for remote inspection of objects in hostile environments to controlling robotic equipment in food-processing operations. The company's systems are even used in high-level preprocessing applications for images that will later be worked on with high-powered DSP or even Cray-type supercomputers.

"But the bottom line," says Siegel, "is that for image-processing technology to become more prevalent in more applications, costs are going to have to drop significantly from their present levels and systems are going to have to be more compact." This has to happen at all levels of the system, from the hardware and software development right down to the rudiments of power supplies and packaging. ■

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

Datacube
(508) 774-9500 ..... Circle 223
Motorola
(602) 438-3576 ..... Circle 224

IN REAL-TIME DEVELOPMENT, ONLY VxWORKS GIVES YOU THE FREEDOM TO CHOOSE.

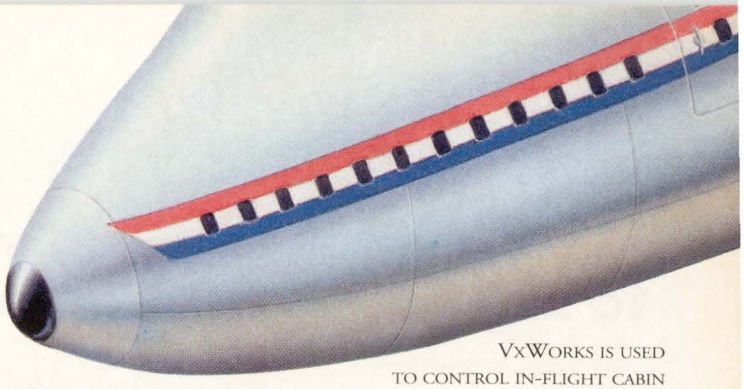

VxWorks is used to control in-flight cabin management systems.

Whether you're looking for speed, flexibility, or compatibility, only VxWorks® provides this much freedom:

1. MULTI-PLATFORM. VxWorks gives you the freedom to choose from over 200 commercially available platforms, including all of today's most popular architectures.

TARGET CPUs
68K
SPARC
MIPS
i960
683XX

2. SCALABILITY. It gives you the freedom to develop applications of any size - small enough to control a fax machine, or large and complex enough to handle electrical power distribution for a major metropolitan area.

VxWorks controls a new line of world famous laser printers.

3. PROVEN PERFORMANCE. VxWorks' high-performance kernel is fast, powerful, and easy-to-use. So you're free to concentrate on building your application and getting it to market quickly.

4. PORTABILITY. With VxWorks you have the freedom that comes from using one of the world's most widely ported real-time operating systems.

5. NETWORKING. VxWorks offers industry standard networking for real-time applications. To this day, it is the only real-time OS that gives you the freedom inherent in fully integrated networking.

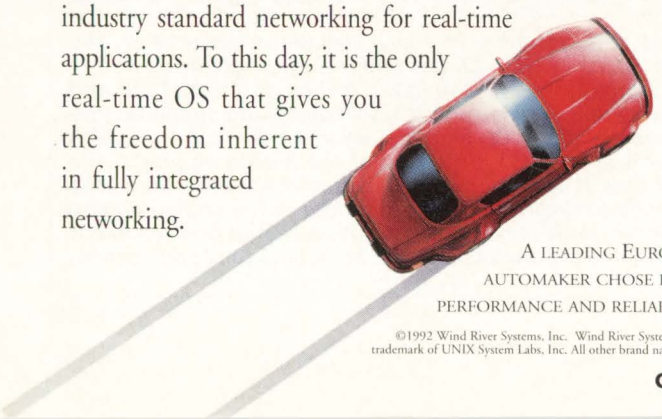

A leading European automaker chose it for performance and reliability.

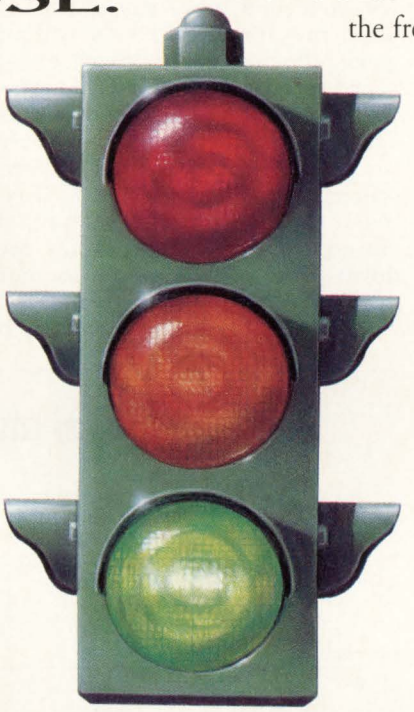

VxWorks controls the traffic signal lights at 12,000 of New York's busiest intersections.

6. TOOLS. VxWorks gives you the freedom to choose from a suite of powerful development tools. And the freedom to use only what you need.

7. SUPPORT. All VxWorks customers have the freedom to call upon our worldwide team of specialists, who are there to help with everything from training to documentation to technical support.

8. ACCEPTANCE. It gives you the freedom and peace of mind in knowing that VxWorks has been chosen by major OEMs, such as Intel® and Motorola®, as one of their own real-time operating systems.

9. STANDARDS. VxWorks gives you the freedom to work with all the leading open system standards.

10. OUR MISSION. We are pledged to provide developers with software and service solutions by giving them the freedom to create real-time applications quickly, economically, and with minimal risk.

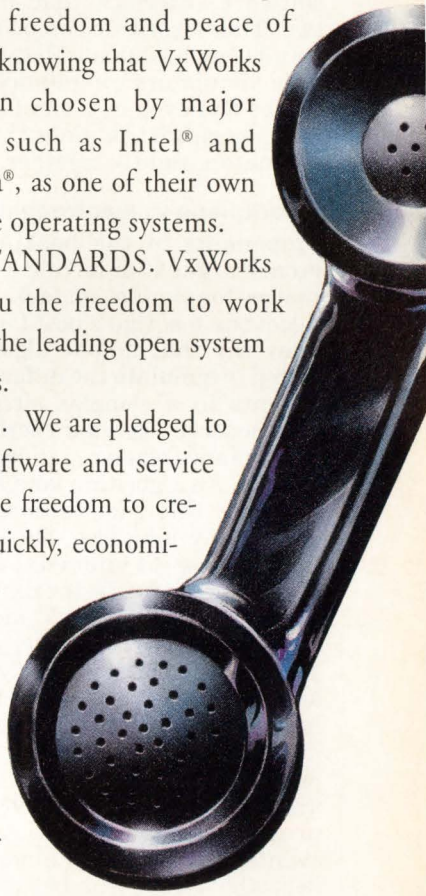

VxWorks is found in telecommunications at every level from PBX switches to handsets.

WIND RIVER SYSTEMS™

TO GAIN YOUR FREEDOM CALL 1-800-677-1586

CAE/CAD TOOLS

No agreement on best way to link digital and analog simulators

Mike Donlin, Senior Editor

The topic of mixed-signal simulation often causes heated debates among EDA vendors. Simply stated, the arguments revolve around which linking algorithms are best—lockstep, leapfrog or backplane. In all of these environments, digital and analog simulators perform separate tasks, each simulating the portion of the circuit that it's assigned. Usually, the faster digital simulator has to wait for the slower analog engine to perform its duties.

Whenever one simulator needs to pass information to the other, the communication is done through a software interface that links the two. This algorithm can be part of a simulation backplane that lets multiple simulators communicate, or it can be achieved by a simulator-specific algorithm that links the chosen simulators tightly together.

Backplane vs leapfrogging

Proponents of the backplane approach argue that having a range of choices for simulation tasks lets you select the tool that's best for the job. They say that no one algorithm is suited to simulate the different components in a complex circuit, and that accuracy is often sacrificed for performance when a simulator is too generic. An algorithm for a switched capacitor filter, for example, won't handle anything else, but it's accurate because it's tailored to simulate one device and is faster than Spice.

EDA vendors such as Cadence Design Systems (San Jose, CA) and Mentor Graphics (Wilsonville, OR) use the backplane method, which integrates digital, analog and mixed-signal algorithms into a homogeneous environment. But there are problems with such an approach—namely, in timing the events of multiple simulators to reflect the real-world behavior of a circuit as its digital and analog components interact.

"Many people who use the backplane method set up a time-based intercommunication," says David Smith, vice-president of engineering at Analogy (Beaverton, OR). "This means that each simulator works within a time slice—say, of 1 ns—to perform a function. That's well and good, but with that approach, events aren't necessarily simulated when they actually happen in the circuit. They happen plus or minus an error term, which makes the simulation efficient but lowers accuracy. If you make the time slice very small, accuracy goes up, but efficiency goes down because there's a lot more event traffic on the backplane."

Analogy solves this problem by avoiding the backplane approach altogether and using its Calaveras algorithm to tie its analog simulator, Saber, to the Cadat digital simulator from Racal-Redac (Mahwah, NJ). The tools are linked together in a master/slave configuration, with the user deciding which simulator will assume the dominant role. The simulators run concurrently and are synchronized via the Calaveras algorithm, which lets the analog simulator go beyond a digital event and run at full speed. If a digital event

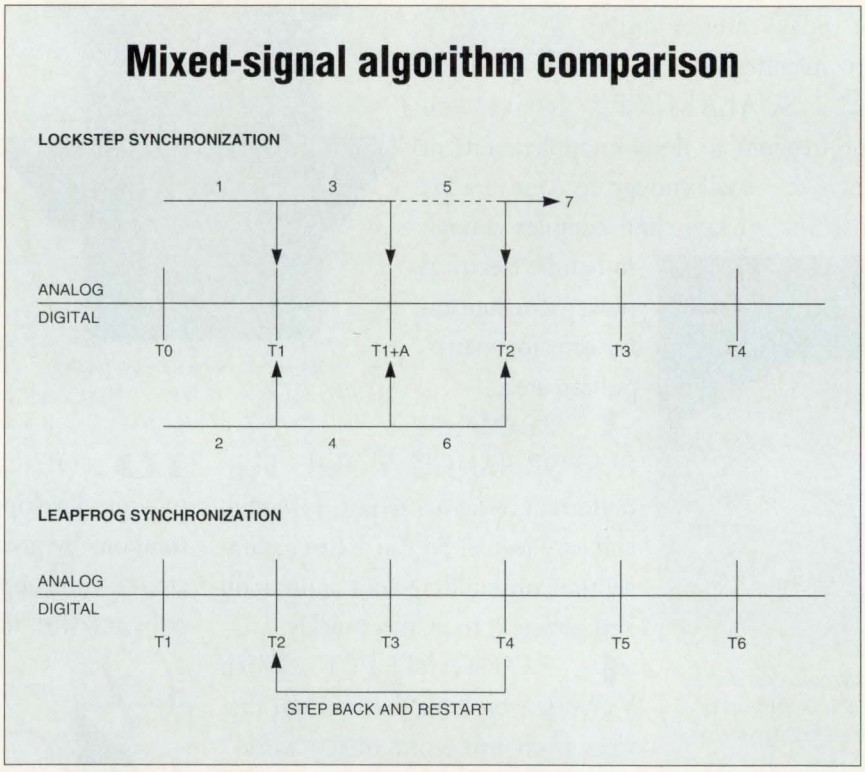

GenRad's Shado simulator uses the lockstep algorithm for synchronization of analog and digital simulation engines. With this approach (top), the digital engine is the master, the analog the slave. At time step t0, the analog engine is asked to simulate up to the time of the next event in the digital queue, t1. After, the digital engine simulates the event at t1, resulting in a new scheduled event at t2. The analog engine now tries to simulate up to t2, but finds an event occurring at t1 + a, short of t2, which will affect the digital simulator. It then stops at t1 + a and passes control back to the digital engine, which reschedules and simulates the event at t1 + a. GenRad compares this to the leapfrog algorithm (bottom), in which the faster digital simulator is allowed to run ahead of the analog engine. If the digital simulator arrives at time t4 with an event to pass to the analog simulator, it waits until the analog engine reaches t4 to transmit data. If the analog engine is at event t2 when the digital event t4 occurrs, the analog engine finds an event relevant to the digital engine, and the digital simulator steps back to t2 and restarts.

Move Up to the Power of XRAY MasterWorks

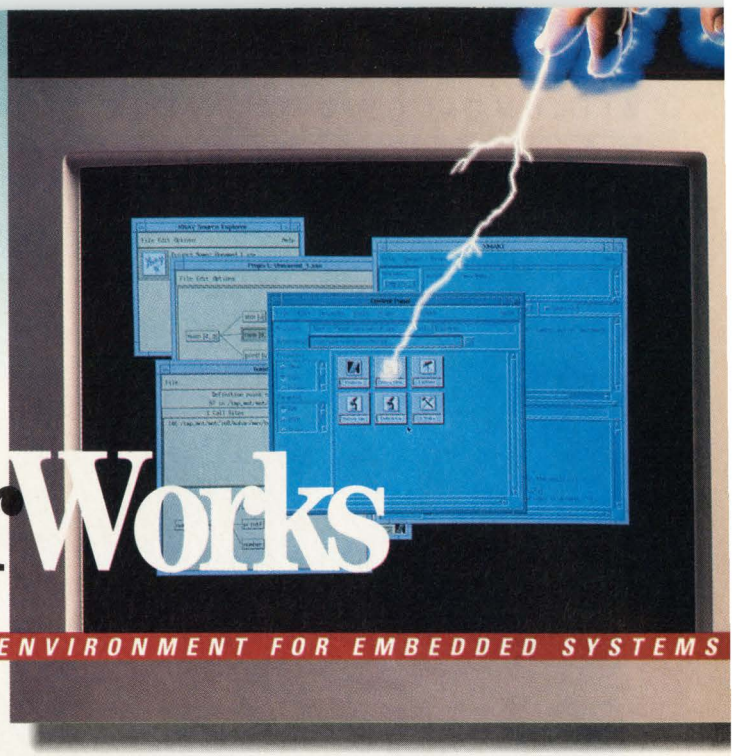

THE C & C++ SOFTWARE DEVELOPMENT ENVIRONMENT FOR EMBEDDED SYSTEMS

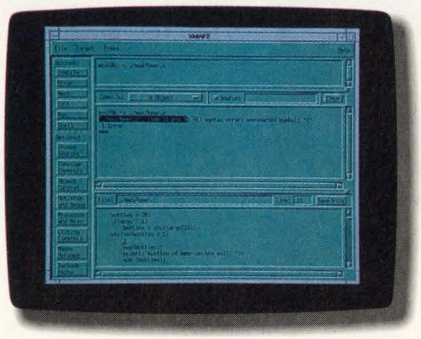

XRAY Make

Automate program build, browse errors graphically, and access build options from notebooks

You've spent years mastering your profession . . . moving up to where you are today. And now XRAY MasterWorks makes it easy to move up to the kind of power you can only get from unified, seamless development tools.

XRAY Debugger

See multiple code and data views, and debug optimized code with a variety of execution environments

XRAY Source Explorer

Graphically browse program call structure and rapidly navigate source code detail

INSTANT ACCESS

Now you can have instant access to any tool, any time . . . all through an industry-standard Graphical User Interface. With **XRAY MasterWorks,**™ the right way to use each tool is instantly obvious. No more guesswork . . . no more fumbling with manuals in your lap . . . all you need to know is right there on the screen.

PUSH-BUTTON CONSTRUCTION

You have enough to think about without trying to remember countless compiler switches. **XRAY MasterWorks**™ lets you access Make from your environment – just by pushing a button – rather than typing a command line. Now you can automatically rebuild your program – and reload it into the debugger – with a single command.

QUICK TO MARKET

XRAY MasterWorks™ lets you see previously hidden aspects of your program structure. For the first time, you'll see a graphical representation of the state of your program as you're debugging. Product quality will improve – and your products will get to market quicker – because you'll find problems faster.

To Move Up to XRAY MasterWorks™ call 800 950-5554.

Microtec Research

Setting the Standard

CIRCLE NO. 73

CAE/CAD TOOLS

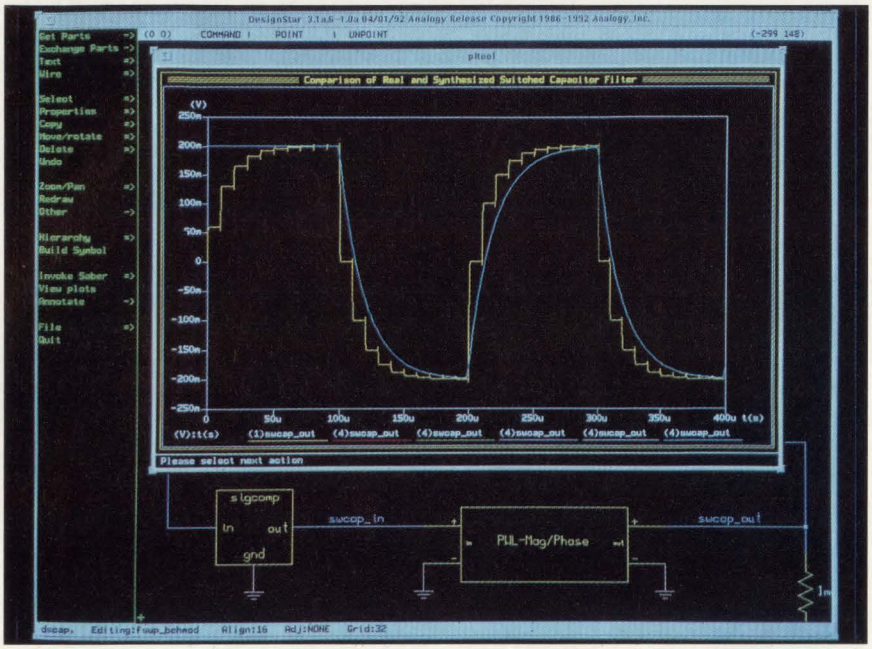

Analogy's Saber analog simulator can use synthesized behavioral models to shorten simulation times. The behavioral model of a switched capacitor filter (blue line) produced from Saber's model synthesis runs over 250× faster than the equivalent primitive circuit representation. The model closely correlates with the actual device output (yellow line) over the entire operating range.

occurs that will influence the analog circuitry, the analog simulator rolls back and resimulates.

"This approach is different than some mixed-signal simulators that backtrack," Analogy's Smith explains. "Those simulators throw out the last analog values computed and recompute with the digital event as the next time step. This wastes valuable CPU time because it throws out previously computed solutions. The roll-back algorithm, on the other hand, doesn't throw away the information it's already calculated. It knows, by the mathematics of calculating the solution at the end point, all the solutions along the curve from the previous time point. Instead of throwing out the current solution and starting over from the previous event, then, it simply rolls back to the time where the event occurred and uses that value without recomputing the whole event."

Critics of this approach say that any backtracking wastes CPU time and is inefficient. "The leapfrog approach is memory-intensive and expensive in CPU time," says John Palmer, mixed-signal product manager at GenRad (Concord, MA). "With that approach, the faster digital simulator arrives at an event and waits for the analog simulator. But if the analog engine determines that an event relevant to the digital engine needs to be transmitted across the A-D interface, then the digital engine must step back and restart. This means that the fastest engine must store large amounts of event history, which uses a lot of memory, and must constantly load and reload data, which is expensive in CPU time. In general, leapfrog algorithms only offer a benefit if the analog and digital engines run on separate CPUs, and even then, they're very circuit-dependent."

GenRad introduces lockstep

GenRad's recently released Shado mixed-signal simulator uses a lockstep algorithm to tie its digital simulator, HISIM, with an analog simulator, Eldo, from AnaCAD, a German company with offices in Fremont, CA. At any given time step, the analog engine is asked to simulate to the time of the next event in the digital queue. After doing so, the digital engine is free to simulate to the next event. If the analog simulator arrives at a solution that will affect the digital simulator, it passes control to the digital engine, which simulates up to the event at which the analog engine stopped, using the new information computed by the analog simulator. GenRad claims that this eliminates a lot of the inefficiency associated with leapfrog algorithms.

"I'd have to take issue on the efficiency question," argues Analogy's Smith. "Because with the lockstep approach, the overall simulation speed depends on the slower of the two engines—namely, the analog. It's true that there's a cost for rolling back, but if you do a statistical analysis, you'll find that, for many circuits, that doesn't happen very often. When you do have to roll back, our algorithm lets you save unchanged data, so the recomputation is reduced. I guess if your circuit behaves in such a way that every digital event affects an analog event, then lockstep would make sense."

Proponents of the backplane approach say that, while simulation speed and accuracy are important, the real issue is flexibility and ease of use. "We think it's important not to lock your design to any one or two simulators," says James Spoto, vice-president of R&D at Cadence. "By keeping a simulator environment open, you can pick different levels of simulation and tie them together hierarchically, depending on the level of detail you need. That means you could use every level of simulation, from behavioral- to circuit-level, and tie them together through one user interface."

Regardless of the claims and counterclaims, however, most vendors agree that the real key to choosing a simulator is knowing what you're going to simulate. Each of these methods has drawbacks in speed, accuracy or ease of use—and each has strengths in the same areas. The real winner will be the one best suited to your circuit. ■

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

Analogy
(503) 626-9700 ..... **Circle 225**
Cadence Design Systems
(408) 943-1234 ..... **Circle 226**
GenRad
(408) 432-1000 ..... **Circle 227**
Mentor Graphics
(800) 547-3000 ..... **Circle 228**
Racal-Redac
(800) 526-0680 ..... **Circle 229**

COMPASS Is All You Need.

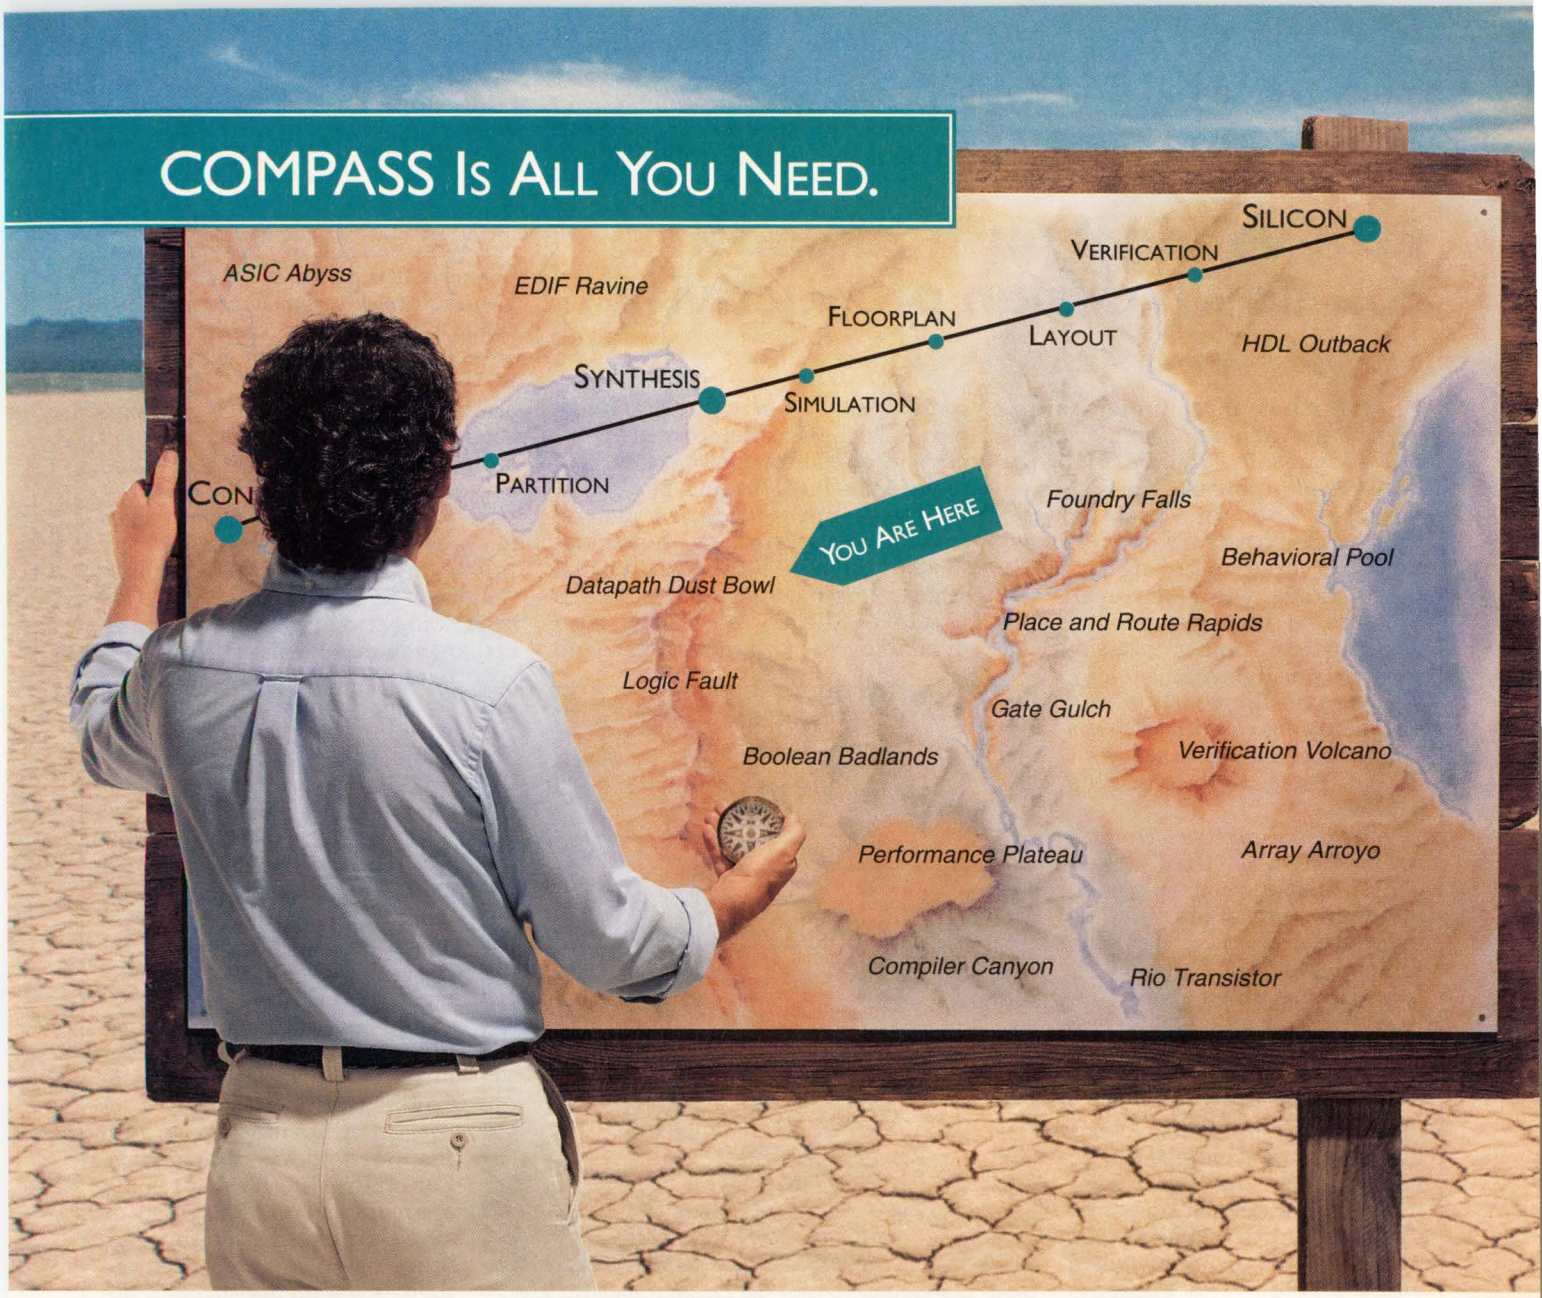

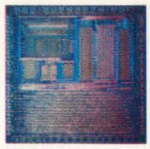

The path to top-down design can be anything but straight and narrow. And once you get off track, you can get into all sorts of trouble. The ASIC wilderness is littered with good point tools that somehow fail to provide an integrated solution.

COMPASS® Design Automation knows all about the rules of the road. That's because we've been down it more than any other EDA company. We've developed the most integrated top-down ASIC solution in the industry, from first concept to first silicon. Supporting even the largest and most complex designs that require a million gates and more.

Simply put, we save you time. Starting with an easy-to-use integrated graphical HDL and synthesis environment through comprehensive cell-based and gate array libraries. And all the automated tools you need to move on down the design road.

So chart your course today. For more information, please circle the reader number, call **(800) 433-4880, ext. 7097** or fax your request to (408) 434-7977.

COMPASS
Design Automation

1865 Lundy Avenue
San Jose, CA 95131
Tel: 800-433-4880
Fax: 408-434-7977

Caution: ASIC design may be hazardous to your health. Design delays due to poorly integrated tools may have a detrimental effect on your career. COMPASS assumes no responsibility for those who fail to follow our path to Silicon Success.

CIRCLE NO. 44

ASICs & ASIC DESIGN TOOLS

EDA vendors push to boost top-down design productivity

Barbara Tuck, Senior Editor

With real technological innovation apparently stalemated, suppliers of design automation tools are zeroing in on business opportunities that fall under the category of services. Vendors have discovered that it takes more than sophisticated toolsets to increase productivity with an HDL- and synthesis-based top-down design methodology. It also takes libraries that can be quickly characterized to the latest processes, a strategy for design reuse, and consulting services.

When Compass Design Automation (San Jose, CA) recently announced the commercialization of its physical layout libraries of low-level components—the first such offering from a major EDA vendor that provides foundry flexibility and support for multiple toolsets—Dan Skilken, worldwide product marketing director said, "Commercial libraries are a bit of a change for the EDA industry. We see it as a more complete solution, one that will let the industry focus more on productivity and value added." The Compass Liberty Series of gate-array and standard-cell libraries and compilers for CMOS ASICs and ASSPs for example is integrated with the Compass ASIC Navigator top-down design system.

Trend toward library sales?

What is the company's rationale for going into the library business? "IC designers and suppliers need more from EDA companies than good design tools to successfully respond to today's market pressures," says president Dieter Mezger. "To keep pace with rapid advances in production capability and design complexity, they require a robust set of library products that provides a solid foundation on which to build new and more advanced designs. This trend toward library commercialization in the '90s closely parallels the acceptance and proliferation of commercial CAE and CAD tools in the late '70s and '80s."

The Compass libraries include schematic descriptions, functional models, physical layouts, footprints, simulation models, and icons for low-level library elements. Also available are RAM, multiplier and datapath compilers. After a foundry is selected, Compass uses automated tools and techniques to generate timing models and layout for a specific set of design rules. On average, according to Compass, the customization of a subset of the full Liberty Series takes six to ten weeks. As part of its customization services, simulation models are generated for most popular simulators. Both VHDL and Verilog are supported.

Sierra Semiconductor uses the Compass physical layout libraries and compilers in the design of its DSP-based communications chip sets, which integrate voice, facsimile and data functions. Andy Varadi, vice-president of R&D at Sierra (seated at left), discusses the compiled datapath layout (on screen) for the company's new V.32 modem chip with design team members Neil Becker and James Tan. "Using Compass digital libraries enables us to spend more time on what we're unique at—analog and mixed-signal design," Varadi says.

Will semiconductor vendors embrace the concept of physical layout libraries being sold by EDA vendors? For fabless semiconductor vendors such as Sierra Semiconductor (San Jose, CA), vice-president of R&D Andy Varadi says that the Compass library service is useful. "Using Compass digital libraries lets us spend more time on what we're unique at—analog and mixed-signal design. We can differentiate ourselves in a new dimension," he says.

Although Sierra's been using Compass libraries extensively, it's been doing so on a contractual basis rather than purchasing the libraries outright. Now that Sierra is buying Compass general-purpose libraries that can be targeted to any silicon vendor's process, Varadi says that "We have a greater degree of freedom. We didn't have foundry flexibility before. Now we can change process or vendor, and we can move products from one process to another by recharacterizing and resimulating."

For semiconductor companies with in-house fabs but not much ASIC technology, the Compass libraries may provide an opportunity to get into the ASIC and ASSP businesses. For others, the libraries could augment current offerings or save characterization time. Library users can also leverage their investments through design reuse.

Design reuse a goal

The methodology, tools and relationships enabling smart reuse of designs is the focus of a new business segment called DesignWare at Synopsys (Mountain View, CA). Intelligent design reuse and the consequent leveraging of industrial intellectual property will yield the productivity required to remain competitive, according to Synopsys, which claims that previous strategies for design reuse haven't taken advantage of synthesis as an enabling technology and VHDL as the worldwide language standard. "VHDL gives the practical promise of being the standard for capturing and describing everything," says Aart de Geus, Synopsys' senior vice-presi-

A Complete Industrial-PC

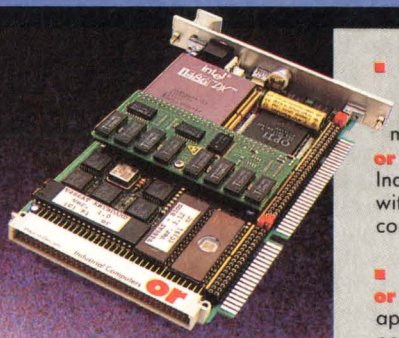

■ **Robust 19" Format**
Put an end to the infuriating problems of non-standard mechanics and poor quality. **or** now offers you a complete Industrial-PC in robust 19" format with reliable VMEbus DIN-connectors.

■ **286-AT, 386-SX and 486-SX/DX**
or has a solution for every application with the right price-performance ratio.

■ **Silicon-Disk, LCD, EL and Plasma**
Does your application demands a CRT, LCD, EL or Plasma-display, or do you need a silicon-disk? Just plug in a board – done!

■ **Low Power CMOS –25/+85° C**
Due to low power CMOS technology boards from **or** guarantee long life and high reliability. The typical temperature range of –70° to +130° C provides a dependable safety margin.

■ **More Quality for your Money**
High quality and newest technology "Made in Germany" must not cost more than other Industrial-PCs. Ask us for an offer, we will prove it.

or

Industrial Computers
Sieglindenstr. 19½, D-8900 Augsburg 1
Tel. +49 821 5034-0, Fax -119

In the US and Canada contact:
DYNATEM
15795 Rockfield Boulevard
Suite G, Irvine, CA 92718
Tel (714) 855 3235, Fax 770 3481

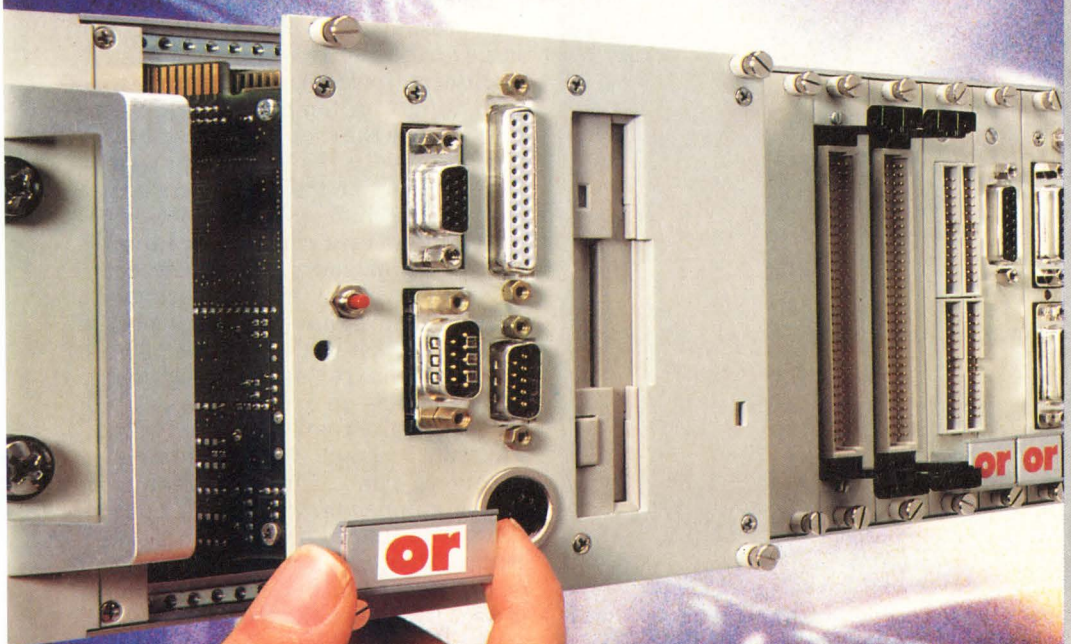

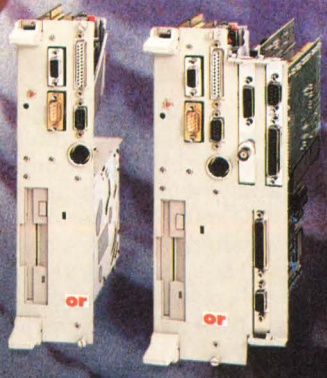

■ **Plug it in – finished!**
The 19" plug-in PC-Boxes are complete Industrial-PC's including floppy and harddisk. Just plug them in – finished. Three versions are available, one in 3U and two in 6U height. One of the 6U versions has 4 PCbus slots for your special PCbus boards.

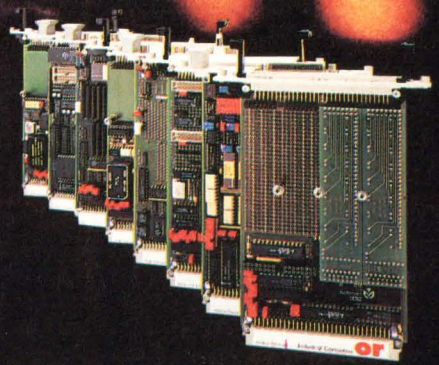

■ **Extendable on the PCbus and VMEbus**
With 440 VMEbus boards **or** has one of the largest industrial-I/O board ranges: digital I/O, analog I/O, counters, motion-controller and more. With Industrial-PCs from **or** both of the most widely used bus-systems are available for your applications: the PCbus and the VMEbus!

ASICs & ASIC DESIGN TOOLS

dent and chairman of the board.

To be available by year's end, the first two design-reuse products from the Synopsys DesignWare division are Synthetic Designs, consisting of off-the-shelf synthesizable component libraries, and software called DesignWare Developer for designers wanting to capture their own synthesizable and reusable modules. Both are integrated with the company's Version 3.0 high-level design tools—VHDL System Simulator, Design Compiler and Test Compiler.

Synthetic Designs are technology-independent, parameterized libraries of commonly used functions such as ALUs and multipliers that have been preverified with the Synopsys VHDL simulator. These functions can be instantiated in a VHDL or Verilog description or inferred from HDL operators. The user configures the Synthetic Design library component for a particular system design by supplying specific parameters and compiling to a target technology. The Synopsys synthesizer then optimizes for the context of the surrounding system, choosing appropriate drives and loads.

If you want to build your own inventory of reusable modules tailored to your particular design requirements, you can use the new DesignWare Developer to create, manage and protect your proprietary design data. The software gives you a way to key and encrypt designs that can then be distributed to users while intellectual property is safeguarded.

Harvey Jones, Synopsys president says, "We expect third parties and end customers to use DesignWare Developer to share their intellectual property with their leading customers and with the general marketplace. As our customers implement increasingly complex systems in silicon, the ability to reuse design data will provide the reduction in design time required to remain competitive."

To help you decide on the appropriate system architecture to target, Synopsys has entered into a partnership with Compiled Designs (Munich, Germany) to provide through the DesignWare program a wide range of system-level models. Also, Synopsys, Texas Instruments (Dallas, TX) and Comdisco Systems (Foster City, CA) will be focusing on optimizing design productivity using core DSP architectures.

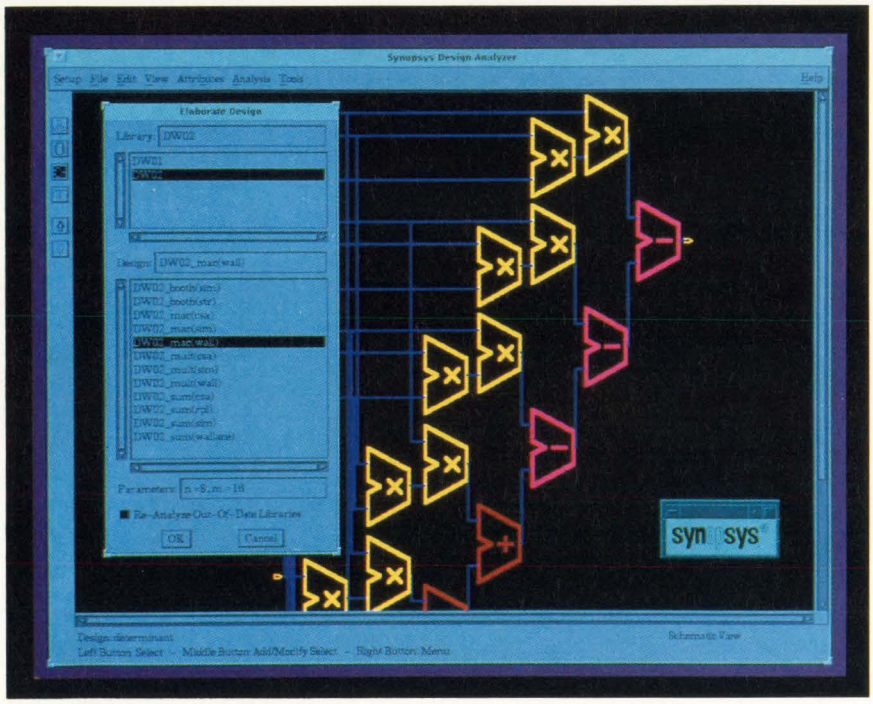

The Synopsys DesignWare synthesis-based design-reuse strategy permits the reuse of existing designs through automatic selection based on design constraints or by manual selection of a component through a menu-driven interface. In this photo, the menu for selecting a Synthetic Designs library family is shown (top left), and the desired component within that library family may be selected (bottom left). The existing design is shown with multipliers (yellow), subtracters (pink) and adders (red).

Hands-on training

To ensure productivity for those adopting VHDL-driven top-down design, Mentor Graphics (Wilsonville, OR) has opened several design centers in the United States, Japan and Europe and has teamed with leading ASIC and FPGA vendors, as well as workstation vendor Sun Microsystems, in a worldwide training program called SmartStart. Mentor's vice-president of corporate marketing, David Chen, says, "In the past, engineers have been reluctant to transition to a top-down methodology because the industry didn't offer a solution with the level of integration, support or training necessary to make the change. The SmartStart program offers a complete solution consisting of software, hardware, silicon fabrication, and support services to ensure that customers meet time-to-market goals."

During SmartStart training, you get the hands-on experience of taking a design from VHDL to layout. Fujitsu, LSI Logic, Mitsubishi, VLSI Technology, and Xilinx have teamed with Mentor to provide a fully qualified and endorsed design flow within Mentor's Version 8.0 design environment. Each of these vendors is offering design kits based on Mentor's Advanced Modeling Process (AMP) ASIC modeling technology to support the company's recently announced VHDL-based, fully integrated, top-down toolset called Design Solver. Mentor's System-1076 VHDL simulator will fully comply with the IEEE 1076 specification by year's end. ■

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

Compass Design Automation
(408) 433-4880 ..... Circle 230
Mentor Graphics
(800) 547-3000 ..... Circle 231
Synopsys
(415) 962-5000 ..... Circle 232

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

MATRIX Rugged VMEbus boards provide commercial-off-the-shelf (COTS) solutions at a fraction of the cost of full MIL-STD products. Test results verified operation from -40° to +85°C, 50 g's of shock at 11 ms, and continuous operation during 5 g's of vibration. Rugged 32-bit processor boards, memory boards, and signal I/O boards give you the flexibility and price advantage to win.

Selected for the Navy's High Speed Fleet Broadcast program, the Rugged MR-CPU330 board supports the Ada environment. You can choose from several host platforms, cross-compile to the target CPU, and begin immediate development using standard industrial products, which are fully compatible with our Rugged boards.

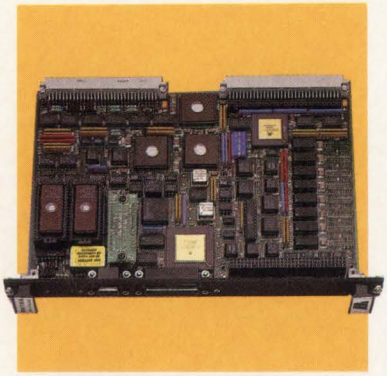

MATRIX CORPORATION

Call today to find out how MATRIX Rugged boards, shock and vibration isolated enclosures, and extended temperature products can control your harsh environment.
Phone: **(800) 848-2330**
FAX: (919) 231-8001

CIRCLE NO. 47

FPGA vendors turn their attention to tools

In the face of stiff competition, FPGA vendors are enhancing proprietary tools, enlisting third-party support, adding text-based entry methods, and backing standards.

Barbara Tuck, Senior Editor

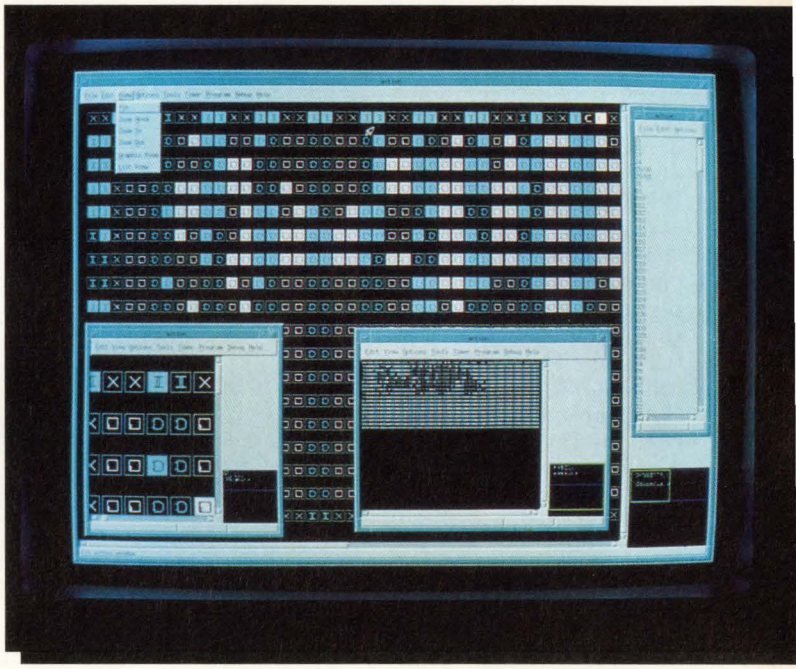

Actel's Action Logic System 2.2 uses a Motif-compatible X-Windows interface and features standard X-Windows menus, a list box (upper righthand text box) that shows all "instance" names in a flattened versus a hierarchical mode, a context window (lower righthand window) that shows the current zoom area on the full chip layout, and a main viewing area that shows three windows containing different zoom perspectives of the same chip layout.

With FPGAs such hot items, silicon vendors have had a captive audience. In concentrating on silicon, however, FPGA vendors haven't always delivered toolsets that were easy to use or boasted good results. But as more and more vendors enter the marketplace, a competitor with tools that fail to route a part or whose entry methods don't include VHDL is at an ever greater disadvantage. As a result, FPGA vendors are finally focusing on creating solutions for users.

To satisfy the demands of electrical designers such as Fred Rakvica at Kodak (Rochester, NY) for interactive place-and-route software, Actel (Sunnyvale, CA) just announced its Action Logic System (ALS) Release 2.2. With interactive placement and incremental place-and-route, ALS 2.2 will be available next quarter for X-Windows-based workstations. Rakvica, who'll be glad to finally have some control over placement, is using nine Actel 1280 FPGAs in a 30,000-gate project, an enhancement for Kodak's photo CD program.

Interactive placement

Although Actel has always stressed automatic place-and-route as the strength of its design system, director of marketing Andy Haines acknowledges that the extra knowledge a seasoned FPGA designer such as Rakvica has can be important in pushing performance. "With ALS 2.2," he explains, "users will be able to select the exact degree of control they'd like to have over their FPGA designs."

The interactive placement feature lets you manually fine tune the results of automatic place-and-route by viewing, moving and editing the location of Actel logic modules; to do so you use ALS 2.2's graphical interface with icon commands. You can also rely on the automatic features of the software to optimize a design incrementally, without disrupting the placement or disturbing timing optimization. On top of all this, ALS 2.2 offers improved macro modeling for up to 25× faster simulation speeds.

Newcomer Concurrent Logic (Sunnyvale, CA) wooed away Actel user Tom Minnis, senior project engineer at Larse Corp (Santa Clara, CA), before Actel introduced its improved toolset. (Larse manufactures communications products.) Though Minnis' FPGA choice was ultimately based on silicon, (according to him, Concurrent's CLi6000 Series is the only FPGA with the freedom

to distribute multiple clocks without introducing skew), design tools, may also have influenced the change in vendor. "Concurrent has a beautiful manual place-and-route program with a very good interactive editor," says Minnis. "Up until now, Actel's software has been push-button."

About two months ago Concurrent Logic also made an incremental design change, adding a feature to its PC-based CDS2100 Development System, which combines design entry from Viewlogic (Marlborough, MA) with proprietary back-end tools. With the new feature, you can change a schematic that's been completely or partially laid out, and implement the change in the layout without disturbing the existing placements and routes. By year's end, Concurrent will offer a Sun workstation interface, and by early next year CLi6000 designers will have the option of using Verilog and Synopsys (Mountain View, CA) synthesis.

Not wanting to be locked into a single FPGA vendor, especially one as new to the field as Concurrent Logic, Larse is also buying silicon and tools from Xilinx (San Jose, CA). But Minnis is no booster of Xilinx place-and-route tools. "It's really tricky to use Xilinx's interactive editor to move stuff around," he believes.

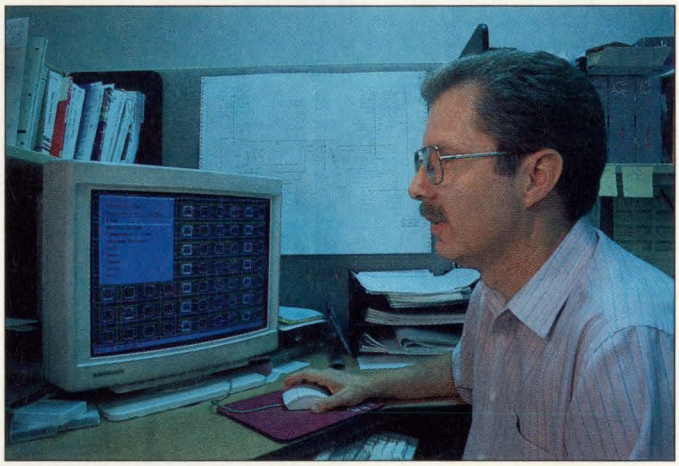

A Concurrent Logic FPGA user, senior project engineer Tom Minnis of Larse Corp uses the Concurrent interactive editor, Interact, to place and interconnect functional cells in a CLi6000 FPGA design destined for a T1 transmission product. "Concurrent's software shows you all the possibilities," reports Minnis.

Specifying timing up front

Xilinx began shipping a tool two months ago it says will significantly reduce the need to manually partition and route critical portions of logic cell array (LCA) designs, eliminating three to five design iterations in the process. As part of Xilinx's latest version of its XACT 4000 development system, the new XACT-Performance tool lets you enter your register transfer requirements (clock-to-setup), I/O transfer requirements (pad-to-setup and clock-to-output), and combinatorial logic requirements (pad-to-pad) in your schematics. The tool will inform you early if performance requirements are unrealistic for the design.

An early user of XACT-Performance, consultant Rocky Awalt, president of Highgate Design (Saratoga, CA), relied on the new Xilinx software, which he refers to as deadline timing, to show engineers at client Boeing Aircraft (Renton, WA) why one of their 4005 LCAs wasn't working properly. Boeing designers are using dozens of Xilinx FPGAs on a 777 aircraft. "Before XACT-Performance was available," says Awalt, "Xilinx users complained about having to do a detailed analysis after compiling a design to determine whether or not the desired result had been achieved. The new software is quite powerful and will meet your timing specifications."

About FPGA tools in general, and Xilinx tools in particular, Awalt says the routing delays and learning curve involved make it "nearly an oxymoron to speak about ease of use and performance at the same time. But there's no such thing as a slow Xilinx LCA. There are just slow engineers."

Senior design engineer Gene Jones at Universal Computing (San Diego, CA), a maker of bus boards and design consultants doing customer-driven designs, would probably disagree with Awalt. Jones designed a Xilinx 3090 LCA into a multiprocessor board for the VMEbus and had to take it out and replace it with a QuickLogic (Santa Clara, CA) pASIC because the Xilinx part wouldn't run at 20 MHz. "After you route a Xilinx part," explains Jones, "you can have internal delays of hundreds of nanoseconds. I don't have a person who can be a master of Xilinx tools. With QuickLogic, I don't need an expert."

Jones says he gets both performance and ease of use with QuickLogic. "And with QuickLogic," he adds, "there's the predictability factor. I can understand and predict what a part's going to do, whereas with Xilinx, I don't know until after place-and-route." Just about the time Jones was giving up on the Xilinx 3090, Xilinx was introducing its 3100 family, which is pin- and software-compatible with 3000 parts and yet is up to twice as fast.

QuickLogic has sought to expand its customer base by enlisting third-party support. For the broad installed base of ABEL users, now numbering over 30,000 PLD and FPGA designers worldwide, QuickLogic has partnered with Data I/O (Redmond, WA), which developed a device fitter for the pASIC 1 family. For designers wanting to stick with device-independent tools, QuickLogic has shared technology with third-party NeoCAD (Boulder, CO), so that NeoCAD's FPGA Foundry can support the pASIC 1 family. FPGA Foundry also supports Xilinx devices.

A heavy user of FPGAs, Derek Rowe, CEO and chief engineer at Defence Products (Lower Kingswoo, Surrey, U.K.), prefers to use FPGA Foundry rather than Xilinx tools to place and route the LCAs he designs into aircraft navigation systems. "NeoCAD is an order of magnitude faster, and it's right every time," says Rowe, who's now awaiting the second release of FPGA Foundry, which will include as an option a timing-driven tool called Timing Wizard. "This new tool will remove the need for manual iteration to achieve a maximum delay for a design," says Rowe, explaining that up until now designers have had to be conservative on the chip's behalf because of the unpredictability of timing delays.

Partitioning options open up

NeoCAD announced a few weeks ago that the second release of its FPGA Foundry would also include an optional timing-driven tool called Prism, which provides automatic post-mapped partitioning of logic functions into multiple FPGAs. Un-

Selecting FPGA design tools

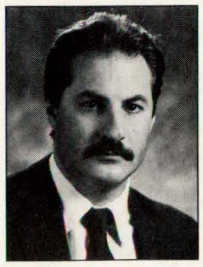

Because of the advantages offered by programmable devices, many engineers are finding themselves designing their first FPGAs. As experienced users will confirm, having the right tools is just as critical to this process as selecting the right device.

The right tools offer a minimum learning curve, shorter design cycle, maximum chip utilization and performance, and support for the device that's best suited to the design. The following guidelines will help FPGA designers select the proper tools for their needs.

What makes a good toolset?

1) A good toolset should support the existing CAE design environment. You should be able to use your existing capture and simulation tools.

2) It should be complete. A minimum toolset should include:

- Entry from all popular design methodologies (Palasm/ABEL, schematic capture, VHDL), as well as translators from industry standards such as EDIF and LPM.
- Mappers or fitters to convert your original design elements into the logic elements available in the selected FPGA.
- Device-specific optimization to provide efficient utilization of the logic within the FPGA without requiring you to manually perform device-specific optimization during design capture.
- Automatic placement and routing.
- A graphical editor with online design rule checking that can be used for preplacing and routing critical signals, or debugging the design after the automatic tools have finished.
- Timing analysis, with the ability to compare the completed design against user-specified requirements and report back on potential problems.
- Automatic back-annotation of timing delays to the simulator of choice.

3) A toolset must deliver shorter time-to-market. To make this a reality, the best FPGA design tools:

- Keep the FPGA design cycle to the fewest iterations possible. This can be achieved using sophisticated algorithms which converge on the best solution, and a rules-driven approach where requirements are set up front to reduce the amount of cleanup required in the end.
- Provide the shortest time per iteration. This can be done by using fast-executing algorithms, as well as the support of an incremental design capability where small changes don't require a total relayout of the design.

4) The ability to optimize the performance and utilization of existing devices is essential. It's vital that the toolset provides efficient optimization routines to best fit a design into the specific architecture. Utilization and performance results can vary considerably from fitter to fitter. Fortunately, the results are quantifiable, which lets you make comparisons easily through benchmarks.

5) The toolset should provide complete control over itself and results. Automated solutions will only give satisfactory results if you can truly direct the tools. You should be able to:

- Specify preferences up front, with the tools adhering to these rules. Preferences include physical constraints such as pin-outs and floorplanning, and timing requirements such as clock frequencies, skew and path delays.
- Prioritize trade-offs. Tool developers constantly have to make decisions such as whether they need algorithm speed vs sophistication, with more complex algorithms requiring more time to execute.

6) It should have timing-driven capabilities. By specifying exact timing requirements up front—frequency and path delays, and not just routing priority or critical versus non-critical—timing problems can be eliminated before they even occur. As a result, timing-driven tools shorten the design cycle while providing significantly faster clock speeds. If for some reason the tools can't meet the requirements automatically, they can pinpoint exactly where the timing problem occurs, as opposed to the search-and-find tactics otherwise required.

7) Make designing FPGAs simple. To do this the tools must:

- Let you think in terms of your original design methodology. An ASIC designer using an FPGA for prototyping, for example, shouldn't be required to design as if the FPGA and not the ASIC was the primary focus.
- Not require you to become an expert on the chip architecture to take full advantage of the device. This can be done using a combination of powerful automatic tools to handle all but the most difficult or unusual design problems, while letting you set requirements by specifying the desired result instead of detailing a specific implementation methodology (such as stating that the maximum path delay between registers isn't to exceed 20 ns, as opposed to "place the design in these logic blocks, route the nets in this order, and use these specific routing resources").

8) Device independence is becoming increasingly important as FPGA vendors continue to enter the market and the tool must support this. Since each FPGA architecture has unique advantages and disadvantages, this trend is good for the user. You shouldn't have to buy and learn multiple sets of tools, however, to take advantage of newer devices which better fit your design requirements. In addition, tools should support the ability to retarget existing designs.

9) Support of a technology-transparent design methodology is also important. This involves the ability to perform design capture and functional verification independent of the final implementation technology (whether it be PCB, FPGA or gate array). The tools should:

- Provide the ability to capture a design using generic libraries and attribute files, freeing designers from being locked into a specific architecture.
- Handle device-specific constraints and rules through files separate from the schematic itself. If the information exists in the schematic, then the schematic itself will have to change to target a different device.
- Support the ability to easily transition from one FPGA architecture to another, as well as between FPGAs and ASICs.

10) Finally, the toolset must run on popular platforms, such as PCs and engineering workstations, and support standards such as EDIF and LPM.

Bob Anastasi, senior vice-president, NeoCAD, Boulder, CO

Before you link up with a pro

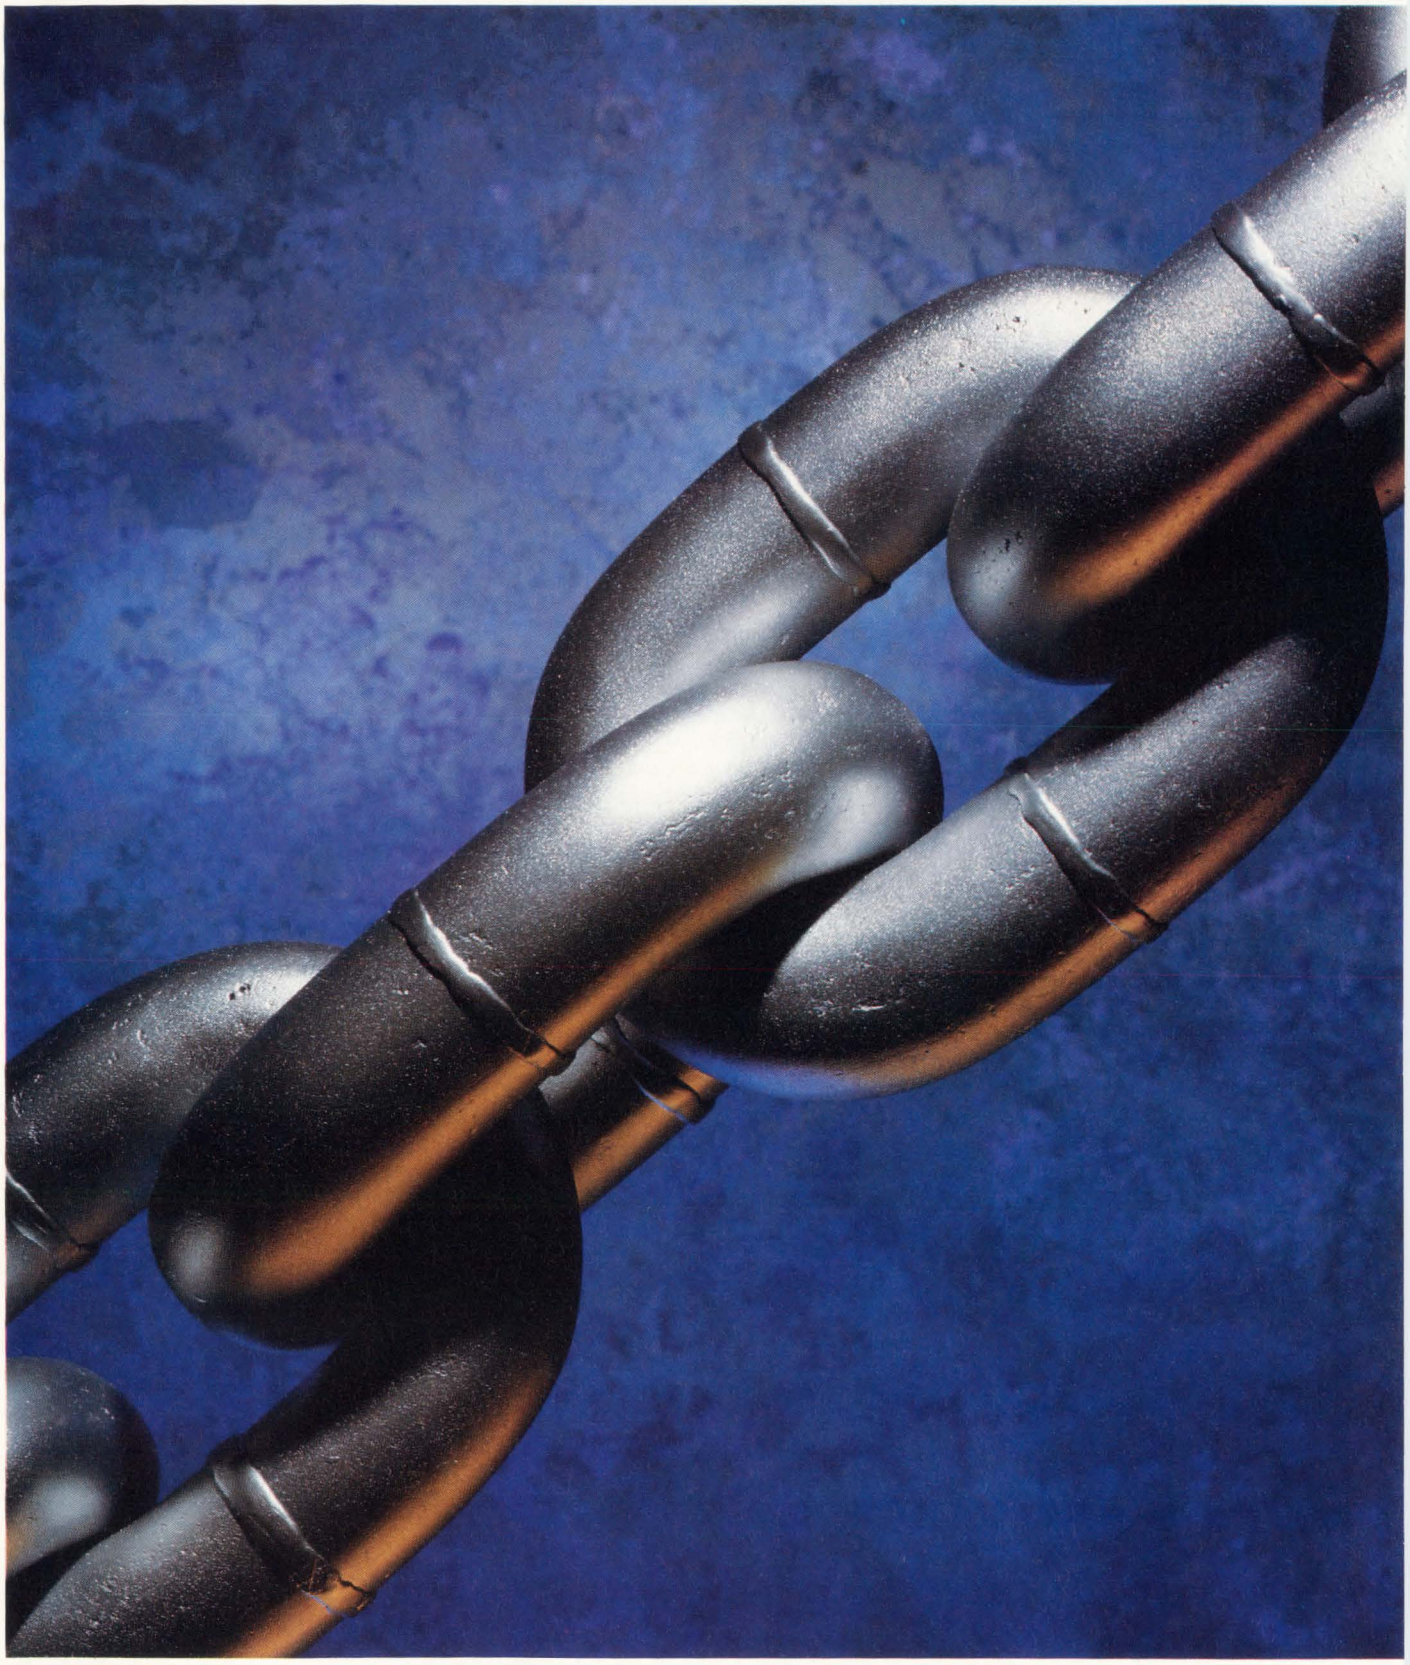

grammable logic company, check our strengths.

At Altera we continually ask, "What more can we do to become your best supplier for programmable logic?"

The answer comes back over and over, "Keep building high quality products, and deliver them on time at competitive prices." That's how the strongest partnerships are built. It's the only way they'll last. That's why so many companies are linking up with Altera to meet their programmable logic needs.

Check our lower prices. We're committed to providing the most attractive pricing for programmable logic devices. Our bottom line is meeting your bottom line — making sure you get your money's worth from every Altera device you choose. One example you'll find the next time you call is the new low prices on our entire MAX 5000 family.

Check our broad product line. Altera offers a wide variety of programmable logic products ranging from 300 to 24,000 usable gates. We started with our industry-standard Classic family of EPLDs almost 10 years ago. Then, we followed with the MAX 5000 family, and last year we introduced MAX 7000, the fastest high density programmable logic you can buy.

But, we didn't stop there. Most recently, we introduced the perfect marriage of EPLDs and FPGAs in FLEX, our newest family of programmable logic.

Check our leading-edge technology and software. Altera gives you the flexibility of multiple architectures and technologies, so you're never limited when designing with our products. Plus, every Altera architecture is compatible with our easy-to-use MAX+PLUS II design system and major EDA tools, allowing you to make the most of the investment made in your existing design environment.

At Altera, our idea of building a solid partnership is to provide you with everything you need to make your job easier and do it faster. In short, to make you more successful. Link up with the most advanced programmable logic technology at the best price.

Connect with Altera today. Call (800) 9-ALTERA (800-925-8372).

CIRCLE NO. 48

like the traditional approach of dividing the design at the schematic capture level, Prism partitions after the technology-mapping process, taking into account the timing requirements and technology-specific characteristics of the targeted FPGA architecture.

Quickturn Systems (Mountain View, CA) has also announced that it will begin shipping an FPGA partitioning tool this month based on technology developed for its hardware emulation systems. Quickturn's Automatic Design Partitioner divides large designs into multiple partitions based on user-specified gate and pin-count parameters. Quickturn's European counterpart, ASIC emulation company InCA (Ascot, Berkshire, U.K. and Campbell, CA), has been leveraging partitioning technology developed for its Virtual ASIC emulation system by selling a stand-alone partitioning tool since last summer. InCA's Concept Silicon compiles and partitions a complex digital design into hardware containing multiple FPGAs from different manufacturers.

FPGA vendors have also had to recognize that users are beginning to couple a hardware description language (HDL) synthesis design methodology with FPGA technology. Specifying functionality, area and performance goals at the register transfer level, they're looking to synthesis tools to generate the optimal implementation. But technology translation and optimization for FPGA architectures, which vary significantly from vendor to vendor, present tremendous challenges to synthesis tools, particularly tools such as those from Synopsys, which are geared to synthesizing the fine-grained architectures of gate arrays.

Synthesis challenged by timing

Commenting on these challenges, Jerry Rau, marketing manager at Synopsys, says, "One of them is critical-path timing. To give predictable results, placement and routing tools must know which nets are most critical, and they also need to know the delay slack available on each net. Logic synthesis tools can easily identify the delay slack, but to do so accurately requires predictable pre-layout wire load estimates. This means that a vicious circle exists today; both layout and synthesis tools are waiting for timing data that the other can provide only when the process is complete."

Rau sees two viable solutions to this problem emerging. "One solution," he says, "calls for FPGA vendors to use hard macros—large functional blocks, such as adders and multipliers, which are pre-characterized for speed and area. A logic synthesis tool such as Synopsys' Design Compiler can automatically evaluate different implementations of a particular function and choose the one that best meets overall design goals. A second solution is constraint-driven layout, which requires cooperation between FPGA and EDA vendors. Logic synthesis can provide rich detail about the criticality of every delay path. This information, when passed to place-and-route tools, can be of great help in fulfilling the designer's intent."

Some FPGAs have a finer granularity than others, making them better candidates for Synopsys synthesis. The Crosspoint Solutions (Santa Clara, CA) CP20K FPGAs, for example, are ideal for Synopsys synthesis because they have a fine-grained, gate-array-like architecture. Vacit Arat, director of marketing for Crosspoint, says, "Our very first customer fine-tuned his design for performance using Synopsys tools. It was 100-percent automatically placed and routed on our 4,200 FPGA, the CP20420, with no manual intervention, and it ran at 50 MHz." Crosspoint also offers a Design Kit for Mentor Graphics' (Wilsonville, OR) AutoLogic 8.0 logic synthesis tool.

Unlike the Crosspoint FPGA, the Xilinx LCA's architecture presented a real challenge for Synopsys synthesis until finally Xilinx and Synopsys together developed XSI, or the Xilinx/Synopsys Interface. The XSI library correlates to the multiple-input, look-up-table (LUT) architecture upon which Xilinx LCAs are based. With more than 65,000 functions needing representation, mapping logic functions as LUTs is more effective than mapping them as primitive gates. "Xilinx XSI provides designers access to the familiar VHDL and Verilog/HDL languages," says Jacob Jacobsson, the company's vice-president of development system products, "while optimizing the design for the Xilinx FPGA architecture. It may have taken longer to develop the library for an optimized solution, but for the user, it's worth the wait."

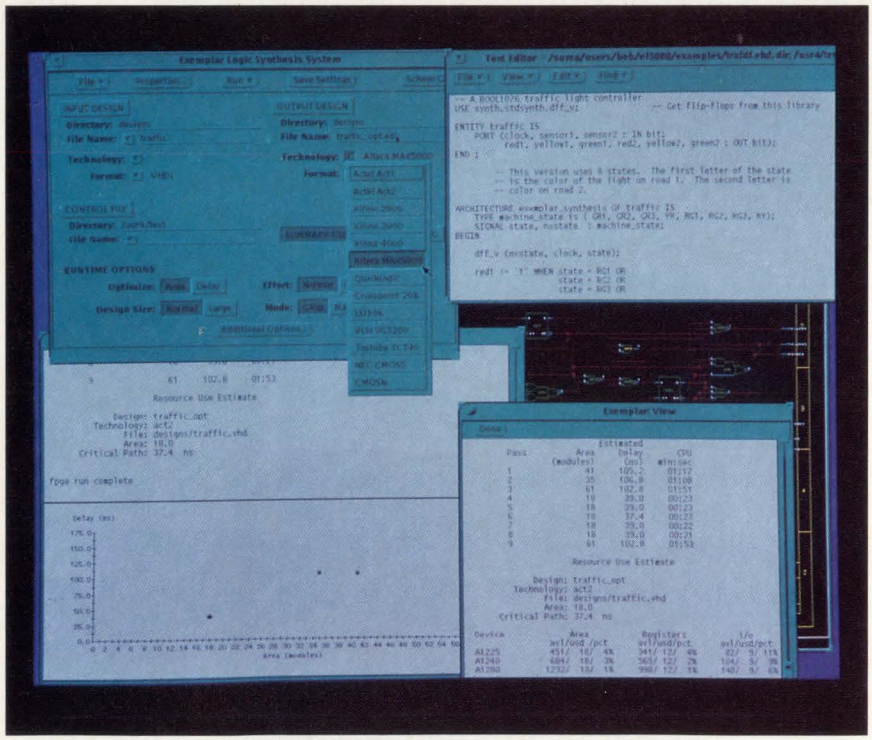

Exemplar Logic's CORE Solution offers vendor-independent design synthesis and technology-specific optimization for FPGAs. This screen shot shows (clockwise from the upper left): CORE's user interface, VHDL, background schematics, CORE's Design Summary Report, and CORE's run-time window. Exemplar offers a growing range of design kits, all leveraging industry or de facto standards and covering Boolean, PLD, Verilog, and VHDL entry.

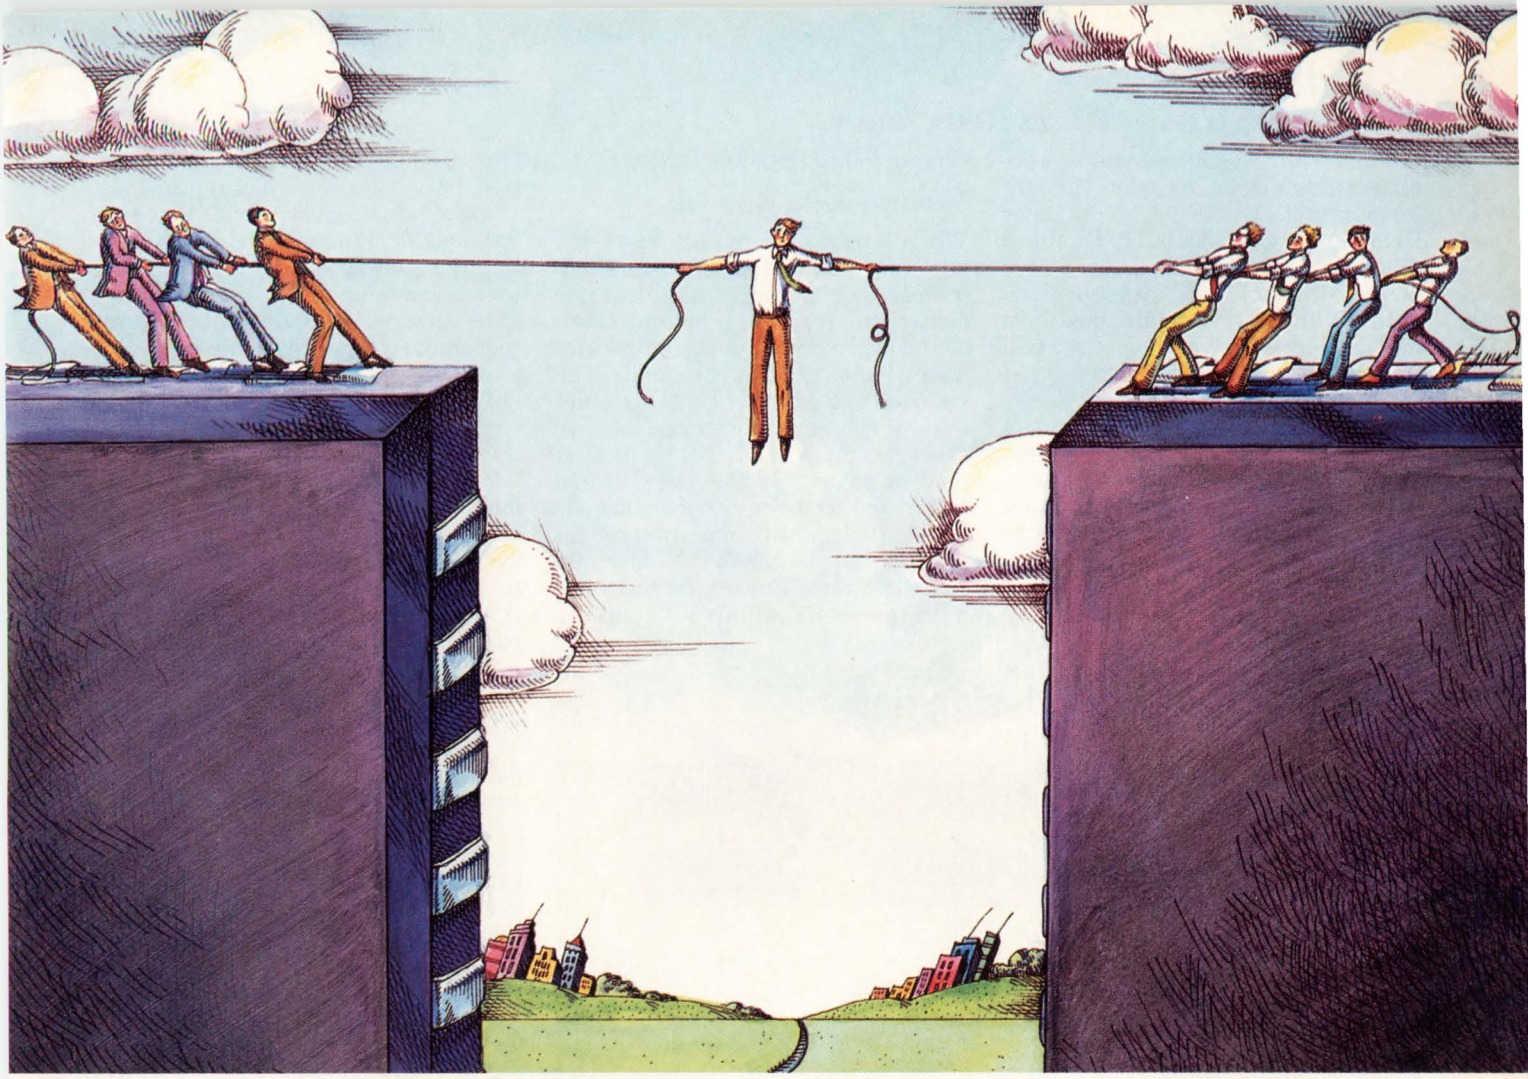

Have your bipolar PLD suppliers left you hanging?

Your product design is simple and elegant and relies on bipolar PLDs. But suddenly your supplier doesn't make them anymore. What do you do? On one hand, it's not cost-efficient to redesign for more expensive CMOS devices. On the other hand, most suppliers seem to be deserting the market. Turn to Texas Instruments. We're one of the largest suppliers of bipolar PLDs in the world. And we're dedicated to serving the market for years to come.

TI has the bipolar PLDs you need

From our TIBPAL16L8 and '20L8 families (the industry's first 5-ns bipolar PLDs) to our latest 7.5-ns '22V10, our PLDs feature familiar universal architectures and speeds from 5 ns to 25 ns. No matter which manufacturer's PLDs you're using now, chances are we've got a perfect match. What's more, each device is available in a variety of industry-standard packages, so there's no need to redesign. Just plug them in and go.

TI helps you get to market faster

As your supplier, our most important job is helping you get your product to market as quickly as possible. That's why we offer a variety of flexible ship-to-stock and just-in-time delivery programs to suit your needs. Plus, our worldwide manufacturing base and complete technical support network are at your service ... anytime ... anywhere.

Our PLDs are designed for quick and dependable programming with your present tools. Or if you wish, we can program your PLDs for you—we take care of the programming risks, while you avoid the time, money and frustrations of running an in-house operation.

Send for your free information pack today

Use the business reply card or call **1-800-477-8924, ext. 3717**, and we'll send you a pack filled with information about our complete range of bipolar PLDs. Or call your nearest Texas Instruments sales office and we'll put you back on solid ground.

00-7517R1

TEXAS INSTRUMENTS

The wait for a Xilinx/Synopsys solution may not be over yet, though, for senior electronic design engineer Brian Box at Lockheed Sanders (Nashua, NH), where engineers have done upward of 400 distinct designs with Xilinx FPGAs. While Box says his hopes are very high that XSI will be usable for mainstream projects at his company, at present its use is limited to research. For engineers to use XSI on the Xilinx parts they design into electronic warfare systems, area and timing optimization have to be enhanced. According to Box, "Synopsys has to have the control to feed timing constraints to the Xilinx place-and-route tools. The biggest achievement of XSI so far is that it gives you the ability to instantiate hard macros and so take advantage of the fast carry logic inside the CLBs [configurable logic blocks] of Xilinx 4000 FPGAs."

Box recently began using VHDL. If and when XSI is ready for mainstream use at Lockheed Sanders, he expects it will knock two weeks off a three-week FPGA design cycle. (This three-week period doesn't include VHDL specification time.) "Today," Box adds, "I have a person sitting in between the VHDL code and the completed chip. We lose two weeks actually constructing the chip, instead of letting a synthesis tool do it. XSI will give us a link between VHDL simulation and chip implementation."

FPGA-specific synthesis

When hardware engineer Van Oler, working at SpaceLabs Medical (Redmond, WA), made a simple logic change to an LCA in a patient monitor he was working on, he found that Xilinx tools wouldn't reroute it, so he resorted to Exemplar Logic's (Berkeley, CA) FPGA-specific synthesis. "We sent Exemplar the Xilinx netlist file," he says, "and they synthesized it and reduced the number of macrocells, freeing up some routing resources so the Xilinx tools could route the part." SpaceLabs has since purchased Exemplar's Complete Optimization/Retargeting Environment (CORE) Solution, which runs on UNIX- and MS-DOS-based operating systems.

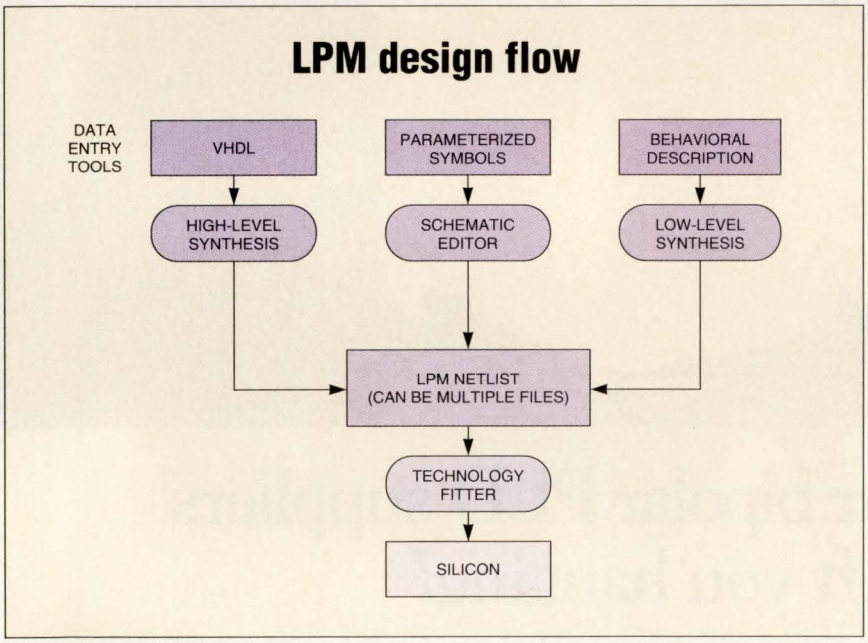

This diagram illustrates the basic tool structure that will implement the Library of Parameterized Modules (LPM) standard. The master netlist file is an EDIF netlist that uses only the components defined in the LPM specification. The technology fitter maps a logical netlist onto a physical implementation, including placement and routing.

After trying to use an interface to Synopsys from Altera (San Jose, CA), Brent Meyer, a senior member of the technical staff at Sandia Labs (Albuquerque, NM), says that, if handed a blank check by management, he too would purchase Exemplar's CORE Solution synthesis, which now supports MAX 5000 and 7000 devices. "The Altera library that runs on Synopsys has no timing information," complains Meyer, who's using a single MAX 5000 part as a memory interface in a system for arms control verification. "I had no control with VHDL. I couldn't do gate-level simulation, only behavioral simulation."

Meyer was concerned because the finished design had to match the original source code he had brought to his customer. For timing problems that showed up after Altera reoptimized the design, Meyer could either go back to the VHDL code, rewrite constraints and resynthesize with Synopsys, or he could rewrite the VHDL code at a lower level.

How long was Meyer's design cycle for the single Altera MAX part? A month, he says, from the day he had his VHDL code ready for his customer to the day he had his timing problems resolved. In the meantime, he sent the project as a test file to Exemplar. "I received a report back within three days. Although I didn't get the entire design back, it looked as if the timing would be okay. On top of that, Exemplar had fitted the design into a smaller part," reports Meyer.

The Exemplar approach to synthesizing Altera parts involves development of a compiler module optimized for the Altera architecture, as well as special libraries for the MAX 5000 and 7000 parts, with timing information built in. "Exemplar synthesis produces something we don't really have to resynthesize," says Craig Lytle, applications manager at Altera, "maybe just tweak a little. It produces code that maps almost directly to our logic." Exemplar outputs an AHDL file, which is then read into Altera's MAX+Plus II toolset. Altera reports that it's developing the capability to accept directly both VHDL and Verilog into its compiler, and it expects that early next year MAX+Plus II will be able to produce VHDL-formatted output files and, by the end of the year, Verilog-formatted output files, with or without timing annotation. At present, Altera users are limited to an EDIF-formatted file.

Standards on the way

As part of the effort to create a solution rather than a nightmare for users, programmable logic silicon and software vendors have been trying to ease the PLD and FPGA design tool dilemma by establishing standards through consortiums. The founding and supporting members of one such consortium have established a technical standard for logic design called the Library of Parameterized Modules (LPM), a generic, technology-independent set of logic primitives to be embedded in EDIF so

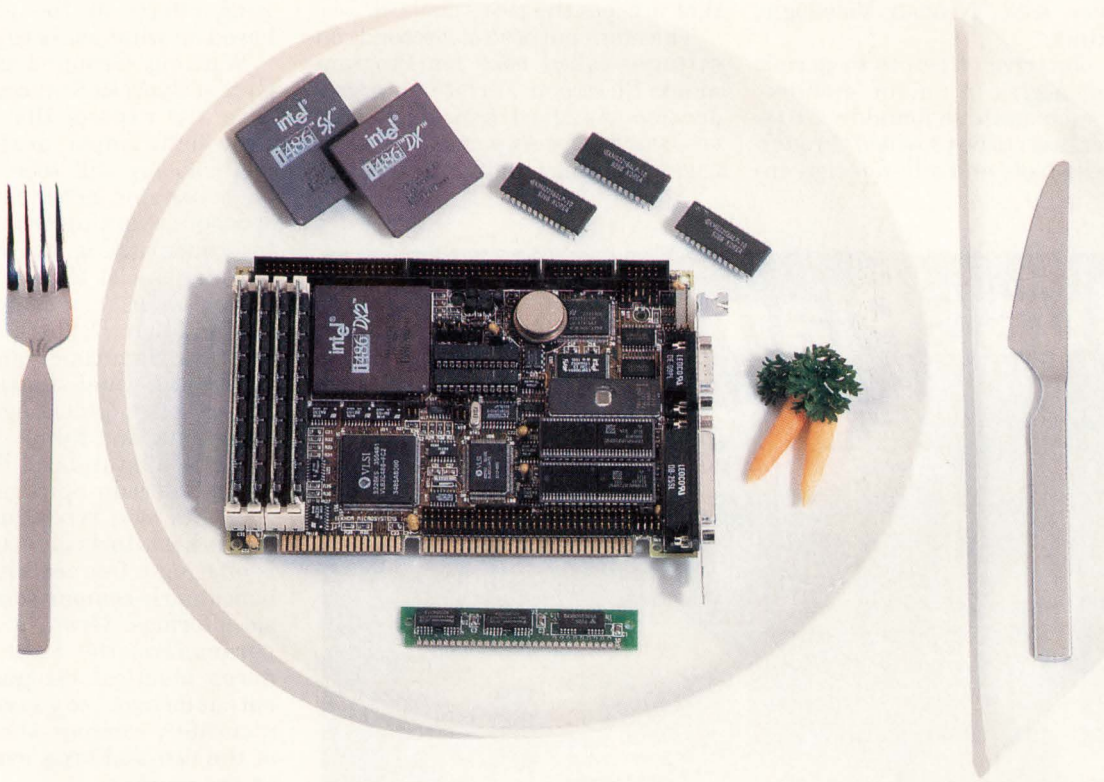

Food For The Power Hungry.

Power seekers demanded it. We delivered. Presenting the 486 TEK-AT4 single board computer.

The first half-size 486 PC/AT to offer a cornucopia of essential features and Teknor's unique approach to modularity: **the interchangeable micro-processor.**

Here's how it works: start with a low cost 486SX configuration. Then, when you're ready to upgrade, sink your teeth into the more powerful DX computing environment or the ultimate 66Mhz 486DX2. Simply by changing the microprocessor.

And since they are all 100% compatible, the cost of upgrading won't cause you heartburn.

So call us today at **1-800-387-4222** to find out how you can sample the TEK-AT4. We're certain it'll satisfy even your most power-hungry applications.

TEKNOR MICROSYSTEMS INC.

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

CIRCLE NO. 50

that a netlist can be created in a standard way. Active participants in the LPM effort include Actel, AT&T, Data I/O, Exemplar Logic, Mentor Graphics, MINC, NeoCAD, Viewlogic, and Xilinx.

The objective of LPM is to permit efficient access to unique architectures, such as those found in FPGAs. This access is to be provided through synthesis tools and other design entry systems. The LPM standard lets any logic synthesis program that generates LPM map a design efficiently onto any FPGA having a fitter that accepts the LPM standard.

The main purpose of a second consortium, called PREP for Programmable Electronic Performance, is to develop standard benchmarks for FPGAs and PLDs. As part of this effort, PREP's Working Group 2 is concerned with establishing a design entry standard for consistency in benchmarking. PREP president Stan Baker says, "PREP is starting its standardizing efforts at the ground level, based on what users say they need."

Working Group 2 chairman Al Graf of Cypress Semiconductor (San Jose, CA) expects that a standard VHDL behavioral design capture methodology will soon be proposed for FPGAs and PLDs. PREP may also sponsor a working group committed to standardizing synthesis benchmarking tools. Companies participating in PREP include Actel, Advanced Micro Devices, Altera, AT&T, Cypress Semiconductor, Data I/O, Gould AMI, Intel, Lattice Semiconductor, MINC, QuickLogic, Texas Instruments, and Xilinx.

Until a standard is available, MINC (Colorado Springs, CO), which offers its own PLDesigner-XL family of synthesis tools for PLDs and FPGAs, is offering a free toolkit to facilitate benchmark comparisons of synthesis offerings. Given a set of input parameters, the MINC utility produces identical designs across all output formats, so you can fairly and accurately compare the capabilities of the PLD and FPGA synthesis tools on the market. ■

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

Actel (408) 739-1010	Circle 233
Altera (408) 894-7000	Circle 234
Concurrent Logic (408) 522-8718	Circle 235
Crosspoint Solutions (408) 988-1584	Circle 236
Data I/O (800) 332-8246	Circle 237
Exemplar Logic (510) 849-0937	Circle 238
InCA (408) 986-8800	Circle 239
Mentor Graphics (503) 685-7000	Circle 240
MINC (719) 590-1155	Circle 241
NeoCAD (303) 442-9121	Circle 242
PREP (408) 356-2169	Circle 243
QuickLogic (408) 987-2007	Circle 244
Quickturn Systems (415) 967-3300	Circle 245
Synopsys (415) 962-5000	Circle 246
Viewlogic (508) 480-0881	Circle 247
Xilinx (800) 255-7778	Circle 248

GUIs move OSs toward object orientation

Computer graphics, originally meant to display data, is now the way users interact with systems. A new generation of operating systems with fully integrated GUIs is meeting the demands of those who want to work with their systems in a real-world environment of objects and actions.

Tom Williams, Senior Editor

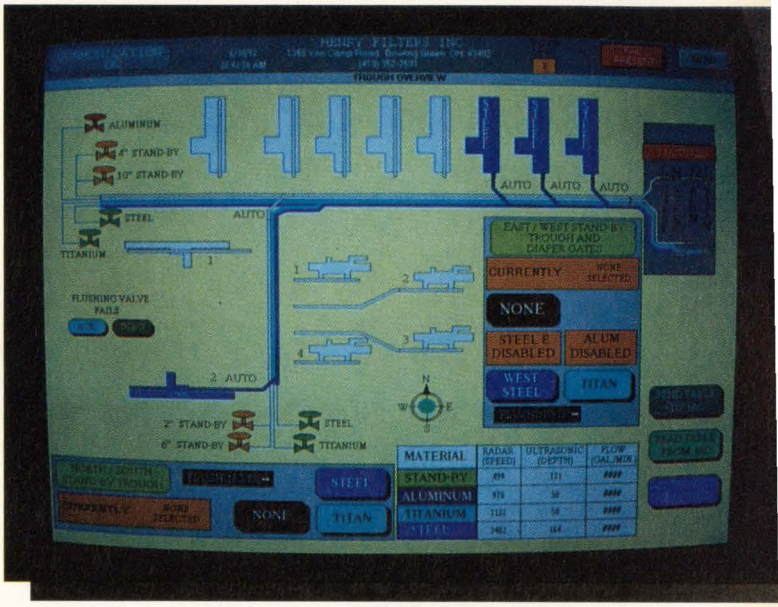

Graphical user interfaces now provide environments where applications controlling objects can represent those objects. GUIs also provide client/server models and interprocess communications that let networks of nodes communicate easily with screens to dynamically update displays. This plant display using Wonderware's In-Touch lets you monitor and interact with an animated model of the plant.

If it's true that the eyes are the windows to the soul, then graphical user interfaces (GUIs) are the windows to the operating system. We're in transition from one generation of operating systems to the next, and the latter will have fully integrated GUIs. Not only will these GUIs let you work more easily with the system, but they'll also define the functionality of the system and enforce programming discipline through application programming interfaces.

Today's GUIs, in fact, are having an even more profound effect on the community of users and software developers. Representing applications and utilities as little telephones and trash cans promotes the idea that we can manipulate our systems in terms of these objects. If you can drag and drop a file into a trash can, why can't you hook an incoming data stream to a meter or strip chart? In some cases you can do this simply—and in more cases if you're willing to write some code to run behind the graphic objects.

These expectations have let GUIs push us steadily toward a more object-oriented environment, both for application development and user interaction with the system. In realtime and embedded systems that require a user interface, operating systems (OSs) such as Windows NT by Microsoft (Redmond, WA), OS-2 from IBM (White Plains, NY), Nextstep from Next NeXT? Computer (Redwood City, CA), and even the iRMX realtime operating system for Windows from Intel (Santa Clara, CA) are including GUIs or direct support for GUIs as an integral part of the product.

This is in contrast to other major operating systems, such as DOS and UNIX, which are about 11 and 15 years old respectively. The main GUIs associated with these systems, MS-Windows 3.1 and X-Windows, were essentially bolted on to the underlying OSs. They have greatly increased the functionality, ease of use and ability to represent and interact graphically with data of the underlying systems, but the add-on nature of the present GUIs, especially in the UNIX world, has created certain limitations. With UNIX, for example, there are many flavors of the OS and several implementations of the basic X-Window technology floating around, so that it's difficult to

realize many of the benefits that a truly integrated GUI can bestow, such as application portability. In the case of MS-Windows, the limitations are less severe because Microsoft has control of both the GUI and the underlying operating system.

Integral GUIs add benefits

A big benefit of an operating system with an integral GUI is that it presents a solid application programming interface (API) that doesn't let programmers circumnavigate the OS and directly manipulate the hardware. An inviolate API ensures portability of applications by enforcing programming discipline, especially if the OS is itself designed to be portable to different processor environments (as is the case with Windows NT). Such a situation is bound to evoke howls of protest from realtime developers, however, because they insist that they need to get directly to the hardware.

There are two basic ways to meet the needs of designers who really have to tickle registers or I/O devices directly, and yet who want to take advantage of the user interface offered by a large operating system such as Windows NT. One way is to write hardware drivers that talk to the OS via its API protocols; another is to encapsulate a realtime kernel such as iRMX, pSOS+ or VRTX on a board with a dedicated processor. The latter lets the kernels respond to interrupts as fast as possible, since a user interface isn't realtime anyway, they can still communicate with the GUI on the larger operating system via a communication protocol supported by the OS's API. Or, if a large OS represents overkill, there's the alternative of iRMX for Windows, which provides a compact realtime kernel as well as access to an industry-standard GUI.

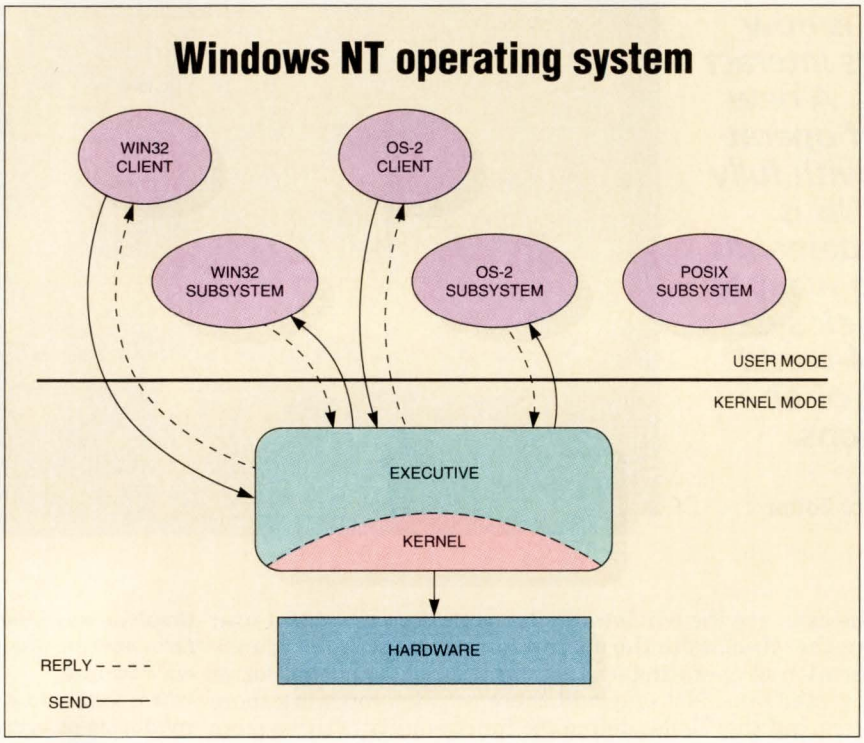

The Windows NT operating system will provide a client/server architecture to make it modular over networks. It will also provide compatibility modules, so that applications written to programming interfaces other than its own will run in the NT environment.

A new generation of OSs

"There's still a strong contingent that believes that UNIX is the wave of the future," says Dennis Morin, president of Wonderware (Irvine, CA). "Imagine something that's been around 15 years still being the wave of the future." The candidates for the "wave of the future" epithet appear to be Windows NT, OS-2 and, perhaps surprisingly, NextStep. All have well integrated GUIs, as well as different capabilities when it comes to realtime and embedded systems. And all have liabilities. OS-2 2.0, for example, is shipping, but sales have yet to reach a million copies. Windows NT was announced for the first part of 1993, but is now delayed until summer. NextStep was once thought to be specialized and restricted to the Next machine, but a 486 version will put it on mainstream hardware where its commitment to object orientation and visual application development could very well make it a strong competitor.

The most likely winner in this race is Windows NT. People seem to be waiting for it rather than opting for the already available OS-2. One thing that will make Windows NT suitable, at least for larger realtime applications, is that it is preemptable not only at the task level but at the kernel level as well. Even if the OS is manipulating some kernel data structure, then, it can respond to an external interrupt in a deterministic manner.

Microsoft hasn't partitioned Windows NT to have a self-sufficient microkernel. "Most folks working in process control are doing it with the assumption that they're going to take the whole thing and embed it in," says Microsoft's NT group product manager, David Thacher. "There's a class of applications where a full-up version of the OS is fine, but you wouldn't use it to control a refrigerator." A full-sized OS with realtime capability is useful, of course, where you need a user interface, a powerful development environment and a wide choice of development tools.

Another factor working in favor of Windows NT is that its user interface is already familiar to millions because it's virtually identical to that of Windows 3.1. The actual API is a 32-bit superset of the present Windows 3.1 API, which means that well-behaved Windows applications will run under NT without modification—as 16-bit operations. "Well-behaved" in this context means the applications are written faithfully to the API and don't directly manipulate the bus or other hardware. The government-mandated security features of NT will forbid such manipulation in any case.

Both OS-2 and Windows NT will aim to run applications written to other APIs. In addition to the Win 32 interface, for example, NT will include modules to let it run programs written for DOS, OS-2 and POSIX, the standard interface for UNIX programming demanded by the government to clear up some of the confusion in the UNIX world. OS-2 includes

a module to run Windows applications, and will also be adding a POSIX interface at some point.

The third potential player in this contest is NextStep, which makes no pretense of being able to run applications written to other APIs, since its own user interface is based on Display PostScript. Originally conceived for Next's proprietary 68030 machines, NextStep was written in Objective C and is being ported to the Intel i486 architecture. "Moving to the 486 is merely a recompile," says Avidis Tavanian, Next's director of systems software. That goes for the applications and object-oriented class libraries as well. The object-oriented nature of the NextStep environment and its Interface Builder graphical programming tool are examples of how GUIs are moving the world closer to object orientation.

Portability among different hardware platforms is something built into Windows NT and NextStep. At the moment, however, OS-2 is restricted to Intel X86 series processors, since portions of it were written in assembly language. IBM has stated that the next version of OS-2, Version 3.0, will be based on the Mach kernel, an improved version of the UNIX kernel developed at Carnegie-Mellon University (Pittsburgh, PA). The Mach kernel is the same one that's at the base of the NextStep operating system; presumably it will make OS-2 portable to other platforms.

There's one caveat about the Mach kernel when used in realtime systems—while it does allow task preemption at the task or thread level, the kernel itself isn't preemptable. OS-2 is similar in this regard in its present version, and will presumably be the same in its future 3.0 version. This limits its applicability for some realtime uses that require both very fast response and strict deterministic behavior. Still, a large number of process control and manufacturing applications don't require this extremely tight scheduling. "If you understand your application and it's well-behaved," says Tavanian, "[NextStep] provides all the capabilities to set up fixed priority scheduling on processes or threads. Or you can choose to time share." Next itself uses its own machines running NextStep to control the robots in its automated assembly operation.

Windows in embedded systems

Window environments, both MS-Windows and X-Windows, have been designed to support client/server architectures. Among the most powerful capabilities to come along with the Microsoft Windows and the OS-2 GUIs is the dynamic data exchange (DDE) interprocess communication protocol, which makes it easy to share data among applications. An enhancement to DDE, object linking and embedding (OLE), lets data created by one application be embedded and even edited in another. OLE is aimed mainly at creating compound documents and appears to the user much as a DDE does. Another capability provided by both MS-Windows and OS-2 is the ability to write dynamic link libraries (DLLs). DLLs are executable code modules that can be loaded on demand and linked at run time, then unloaded when they're no longer needed.

DDE was originally conceived as a way for applications running on the same machine to share data, but it's also served as an easy means of transferring data from a realtime node into the Windows-based user interface application. Conversely, an operator can interact with graphics objects on a Windows screen—a slider control, for example—and send values back to a realtime module via a DDE link; these values can serve as input commands to the realtime process. Process control designers have made creative use of off-the-shelf Windows applications, such as the Excel spreadsheet. Data from an embedded node, for example, is sent to a spreadsheet cell, and that cell is then linked to an Excel bar graph that changes as the cell data is updated.

A natural follow-on step to using off-the-shelf business applications is to supply graphical tools that let you actually model and control your system in terms of pumps, valves, sensors, actuators, and other physical objects in a plant, factory or other system. Wonderware supplies an MS-Windows-based package called InTouch that lets you design screens with dynamic graphic elements and connect them to control systems via DDE drivers.

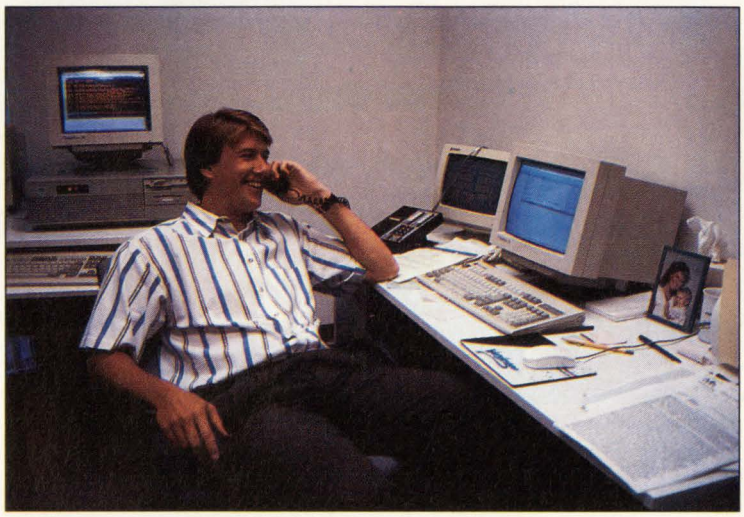

"The industrial automation business," says Wonderware's Phil Huber, "was once the bastion of proprietary equipment. Everything was proprietary. Now it's turning into a business built on commodities, and the differentiation is software."

Windows, of course, isn't itself an object-oriented environment, but Wonderware supplies object-oriented graphical tools that let you build dynamic objects. Once the object has been drawn, you can assign it dynamic characteristics, including animation. Such an object, when saved, can be recalled, modified and saved as a separate object, or built into some larger object. In this way, the user interface to a plant can easily be modified as the physical plant is reconfigured.

Normally, DDE links applications running on the same machine, but Wonderware has developed a network version called NetDDE that permits peer-to-peer communications over a network. Each node on the network must have a NetDDE driver. Wonderware has created drivers that support not only different OSs such as OS-2, UNIX and VMS, but also a range of programmable logic controllers commonly used in industrial process control, such as those made by Allen-Bradley, General Electric, Reliance, and others.

And Wonderware has built intelligence into its DDE drivers. The drivers have to poll various points in a device. Because most drivers send data out over a serial line, polling all the data points every time one of them changed could flood the network with data in a large system. If a system has 50 different screens, for example, each screen might call up a different subset of the data from a given PLC, or might want to poll some points more frequently than others. "All we want are the data points being displayed at the time, data points that might affect an historical database," says Phil Huber, vice-president of engineering for Wonderware. "So every time you change a screen, you've got to redefine the polling list."

Wonderware has built algorithms into the DDE servers that can dynamically allocate the polling list to determine how often and in what order points are polled. The server monitors changes and only notifies the user interface if there is one. This report-by-exception dynamic polling makes feasible systems with many nodes and a large number of screens. "This gives us the ability to develop an application in 10 to 15 percent of the time it used to take," says Huber. InTouch doesn't care whether you're dealing with a DDE server to a device that's local or across the net. "The whole purpose of NetDDE is to completely hide what's on the other side," Huber adds.

Building GUIs and applications

The one thing anyone who's written applications for a windowed environment will agree upon is that a GUI involves a lot of complex code. Writing C code from scratch to specify the pixel dimensions of a window or to design a dialog box gets old quickly and cuts into a programmer's productivity. X-Window implementations such as Motif and Open Look have established toolkits to help programmers with these repetitive tasks, but these toolkits by themselves have turned out to be too limited in functionality to let the programmer get on with the job.

The next stage of development for GUIs has been to automate the process of building the user interface to an application in terms of the look and feel of the operating system GUI. This level of programming aid now comes in two flavors. The first consists of interface construction packages that run as programs under the OS. With these programs you can create the user interface first and then link it to the application code. The second type of aid is composed of a growing number of programming language and development products including enhancements to help automate building the GUI as part of the application code.

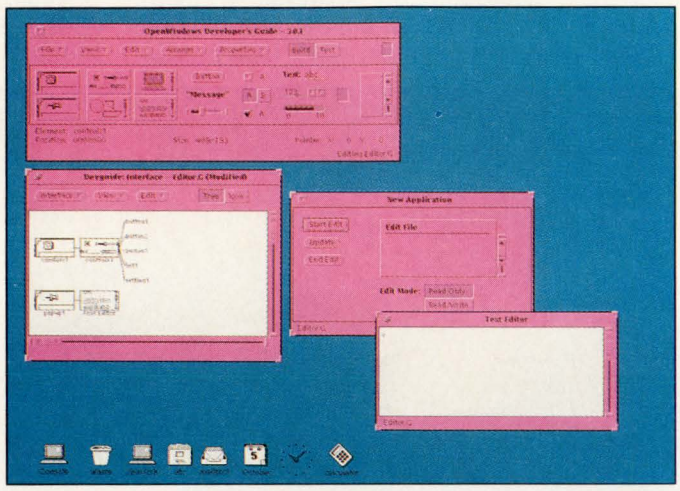

Devguide from SunSoft runs as a program under Sun's Solaris version of UNIX. It lets you create the menu and dialog items needed for a user interface to the application. To do so, it generates an intermediate code that can be used to generate source code, which can in turn be linked to the application and compiled.

An example of the first flavor of aid is Devguide by SunSoft (Mountain View, CA). Devguide's latest version runs under Sun's UNIX-flavored Solaris, and uses the principles of Sun's OpenWindows environment to permit design and testing of user interfaces for Open Look applications that are compatible in look and feel (as well as interoperability) with the Solaris DeskSet environment. The complexity of this situation illustrates why a new generation of OSs with integrated GUIs is emerging.

To develop a user interface under Devguide, you select and combine graphic elements such as windows, dialog boxes and buttons. You can give buttons names and assign help text to them. Devguide includes a connection manager to let you specify source and target objects. A source button, for example, can be associated with a target pop-up box, so that when the button is clicked the box appears. Each item also has a property sheet associated with it to specify name, footer, label, and other characteristics.

One advantage of designing the GUI in an interactive manner, such as that provided by Devguide, is that you're setting up the specification of the program either by following some prewritten specification or by intuitively defining how you want the application to work before you write any code. Devguide produces an intermediate code that can be used with code generators to produce C, C++ or Ada source code. Interface objects produce code stubs that connect to corresponding functions in the application. The GUI source can then be linked and compiled with the application code.

Bundled GUI tools

In operating system environments that have a single, standard GUI—Windows, Windows NT, OS-2 and NextStep, for example—vendors of programming languages can afford to include GUI development tools with their products. The Turbo C++ compiler and programming environment from Borland International (Scotts Valley, CA), for example, includes tools to create Windows user interfaces for applications on three levels. One can, of course, always write C code to call the Windows API directly, but Borland includes a facility called ObjectWindows that uses a library of C++ objects to encapsulate GUI functions. So, instead of writing 50 lines of C code to set up data structures, register the application with windows and perform other functions, you write a few lines of C++ with some parameters that call the ObjectWindows library.

The third method is to create windows, boxes, menus, and other objects interactively on the screen using a tool called Resource Workshop. The objects then generate C++ code as part of the application.

"What Windows needs to be object-oriented," says Borland's C++ product manager, Charles Dickerson, "is to encapsulate the API so that you have functional objects to plug

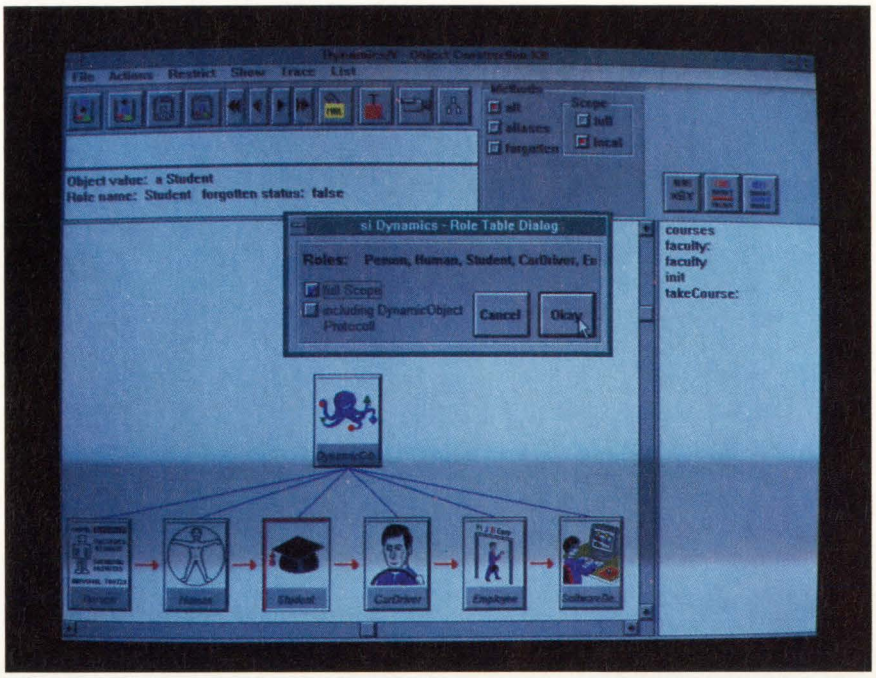

A dynamic object created with the Dynamics/V Object Construction Kit can have a number of roles assigned to it by embedding other objects. Then, depending on the type of message sent to it, it will respond in the proper role. Objects embedded in a dynamic object can be changed, added to or replaced as the situation demands—and without disturbing the other areas of functionality of the dynamic object.

into the application instead of doing all this C interaction." Turbo C++ ObjectWindows does this with a mechanism called dynamically dispatched virtual tables (DDVTs), which replace the nested switch statements required by the API for objects that can inherit response functions, such as "Open a box." ObjectWindows, then, lets the application and its user interface be object-oriented, even if the OS and its GUI aren't.

NextStep the next step?

Taking a step closer to a GUI-based, object-oriented environment that includes object-oriented program development is the NextStep Interface Builder. With it, objects supplied by Next or third-party software developers are represented as graphic icons and can be interactively connected on the screen. Interface Builder is run-time bound, so that as soon as two objects—a slider bar and something it controls, for example—are connected together and their relative scales defined, the slider can directly control the object. In this way you can test the functioning of the program as it's developed.

With NextStep, objects are created in Objective C and can be represented as icons, as can user interface objects. The difference is that the former won't appear on the screen when the application is run. This has given rise to a group of software vendors whose products are class libraries rather than packaged applications. You can put together the major parts of a generic application by simply hooking together these objects. Then you can add value or address your special needs by either modifying one or more of the objects in the library or writing new objects of your own.

A two-way street

Object-oriented languages are taking cues from GUIs as well as the other way around. Smalltalk, for example, is a language that was designed from the ground up to be object-oriented. Objects are retrieved from class libraries as pieces of source code and built into applications. Recent versions of Smalltalk for Windows and OS-2, developed by Digitalk (San Diego, CA), have included objects that let you graphically construct windows, dialog boxes and other user interface elements for the application under development. The GUI elements exist in a class library, just as other Smalltalk objects do. Recently, Digitalk introduced a product called PARTS (for Parts Assembly and Reuse Tool Set) that lets Smalltalk objects be represented graphically and assembled using an environment called the PARTS Workbench. Currently available under OS-2, PARTS will also be ported to Windows and Windows NT.

A third-party vendor, SI Dataservice (Munich, Germany), has developed a similar graphical development environment. Called Dynamics/V, it runs on top of Smalltalk/V and other flavors of Smalltalk. Dynamics/V uses an object construction kit to browse class libraries, connect graphical icons that represent classes, form objects, and ultimately create applications. The construction kit lets you create a dynamic object, or a "software IC," with as many "roles," or modes of behavior, as can be assigned to that class.

The implication of this capability for general-purpose interrupt drive systems is that response can be built easily into automation systems, because the interfaces between objects are already defined. As the requirements of the application change, roles can be added or subtracted—in terms of existing objects, objects constructed or modified with multiple objects, or custom-created objects that then add to the existing repertory. In this context, there could be interrupt service objects that let an application be set up with as many roles as needed and without having to be altered in other ways. ■

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

Borland International	(800) 331-0877	**Circle 249**
Digitalk	(213) 645-1082	**Circle 250**
Intel	(800) 547-8806	**Circle 251**
Microsoft	(800) 426-9400	**Circle 252**
NeXT Computer	(800) 848-6398	**Circle 253**
SI dataservice	(619) 322-2761	**Circle 254**
Sun Soft	(415) 336-0678	**Circle 255**
Wonderware	(714) 727-3200	**Circle 256**

DOS Engine.™

A crushing advantage for the DOS Engineer.

Imagine a single chip that wipes out the need for the individual components of a PC-XT motherboard. That's DOS Engine. The familiar DOS environment, including LIM 4.0 support, keeps software development running on time.

More powerful than a locomotive.

DOS Engine is a ready-to-use standard part. But, you can also add up to 20,000 gates of your own on-board logic. Or even use DOS Engine as the core of your cell-based application (CB-C7), choosing peripherals from our huge ASIC library.

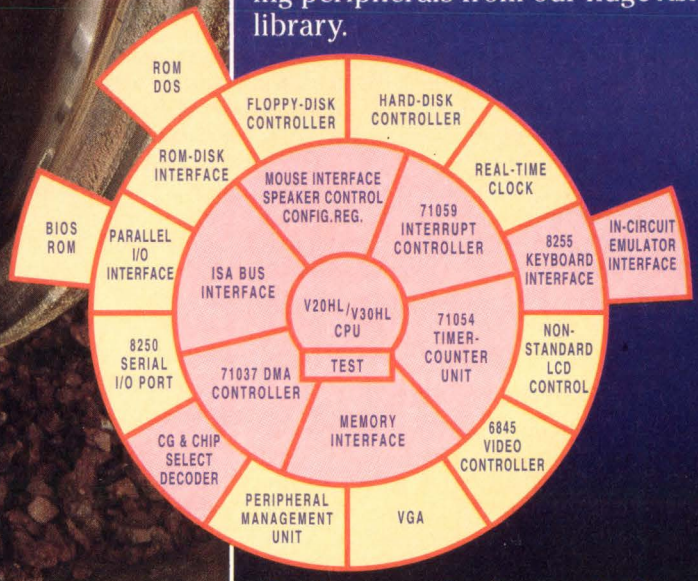

Order a whole trainload.

Need PC functionality in a standard part? Only DOS Engine is powered by NEC's astonishing cell-based technology. When high volumes are your destination, your safest route is NEC. You know we'll never run out of steam.

Hop on board.
Ask for Info Pack 160.
Call 1-800-632-3531.
FAX 1-800-729-9288.

NEC

401 Ellis Street P.O. Box 7241
Mountain View, CA 94039-7241
1-800-632-3531

CIRCLE NO. 52

32-bit power and tools bring cheer to embedded system designers

Because of their increased compute power, higher integration, extensive tool sets, and a desire for the friendliness provided by high-level languages, more designers are putting 32-bit microcontrollers on their wish lists for next-generation products.

Thirty-two-bit processors and microcontrollers have moved beyond PostScript printers into industrial controllers, multimedia devices and telecommunications bridges and routers. A number of new microcontrollers, support ASICs and boards are making it easier than ever to bring a 32-bit product to market at an economical price. At the same time, an equally impressive array of new products from software development companies and emulator suppliers is simplifying development, hardware integration and debugging.

This doesn't signal the end of the 8051-based household appliance controller, it's just that these newer applications require the processing and data- handling speed, as well as the address space, of 32-bit engines. In some cases, too, it isn't raw power that drives development

Don Tuite, Senior Editor

teams to 32 bits, but the need for the programmer-friendliness and documentability that high-level languages provide. The price paid for high-level programming, however, is larger object modules.

Some observers find that the shift from a hardware-driven approach to embedded systems design to a software-driven approach signals a watershed opportunity for both tool and chip makers. "I've seen a lot of 32-bit applications that would run just fine on 16- or even 8-bit processors if the code were written in assembler," says Charles Davis, president of Huntsville Microsystems (Huntsville, AL), half-jokingly. Davis also notes another shift in the development scenario. At the same time that company or contract requirements impose C++ or Ada programming restrictions, the complex nature of the embedded task creates a need for developers who are more familiar with the end application than with silicon architectures. Vendors of chips and development tools have noted this trend and are standing by with "solutions," as the marketers say.

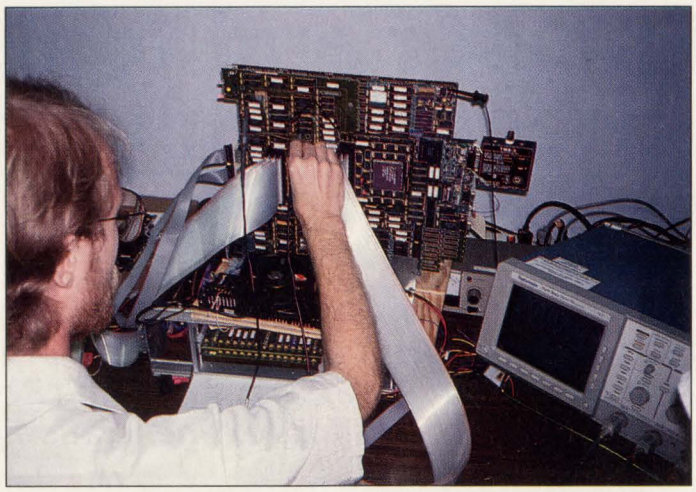

Using his company's present-generation emulator to design the next generation, Embedded Performance's Ted Conard tracks overshoot on a clock signal. A senior hardware development engineer, Conard agrees that the type of person designing embedded systems is changing. "We see a different kind of customer using [AMD's] 29200 than we did using the older chips," he says.

Not just RISC

The technical press's fascination with RISC sometimes obscures the 32-bit design wins that Motorola's 68000 and its 68000-based 683XX microcontroller family have piled up. Motorola's (Austin, TX) continuing success is at least partly due to the length of time the 68000 and its offspring have been around. As Kenneth Greenberg, director of technical marketing at Microtec Research (Santa Clara, CA) says, "The overwhelming majority of embedded system designers said they used the same processor used in their last project, or at least a member of the same processor family." And as Ed Rathje, vice-president of JMI Software Consultants (Springhouse, PA) notes, "Most of the engineers now doing embedded designs graduated from college knowing 68000 assembler."

The two most general-purpose members of the 68300 family are the 68331 and 68332. The 68331 strips out the 68332's sophisticated time-processing unit and its 2 kbits of SRAM. Both are aimed at applications such as automotive controllers.

In contrast, the 68340 has a different timer module than the 330 parts, plus a dual-channel DMA controller capable of 33-Mbyte/s speeds. It targets multimedia applications. There are now 5-V and 3.3-V versions of the 68340. At the top of the 68300 performance scale, the 68302 has a DMA controller and a separate three-channel communications processor. It's aimed at telecommunications.

Motorola offers embedded control versions of its 68000/020/030/040 microprocessors as well. In this case, embedded control means that the part is stripped of MMU, which results in a price about half that of a full-featured chip.

Adding field-programmable elements to 32-bit microcontrollers was pioneered by Motorola in 1991. There are currently two objects of programmability: to provide for design changes late in the cycle, or to supply security features, such as the ability to program serial numbers at the time the chip is assembled onto a board or to install data encryption keys.

With programmability in mind, Motorola's 68300 microcontroller family includes the 68F333, with a flash EEPROM on-chip. The 68F333 continues to be the only 32-bit microcontroller with on-board reprogrammable memory. This October, however, VLSI Technology (San Jose, CA) introduced programmable functional system blocks (pFSBs), ASIC PROM and PLD elements that can be embedded along with ARM (Advanced RISC Microprocessor) 32-bit processor cores. These are one-time programmable elements, based on the ViaLink antifuse technology jointly developed by VLSI Technology and QuickLogic (Santa Clara, CA). The 68F333 incorporates 64 kbytes of flash. The largest of VLSI Technology's pFSB ROMs is

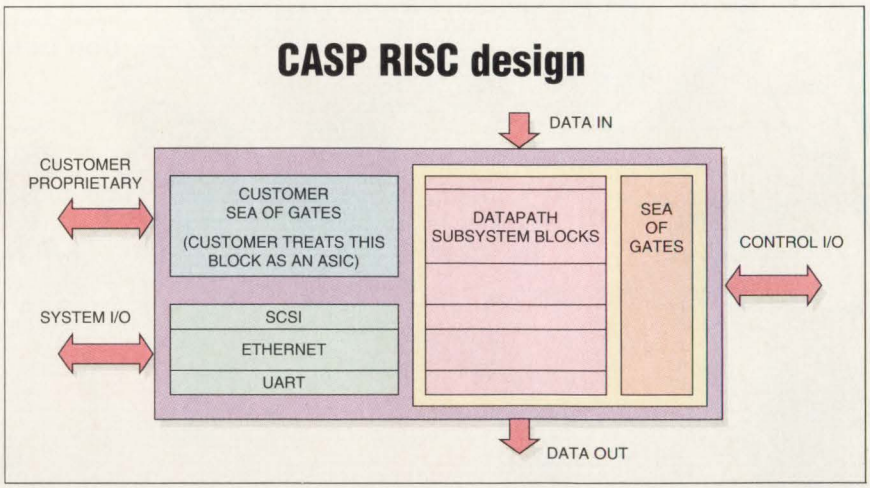

Not all 32-bit embedded design is for the hardware-challenged. The alternative to letting your silicon vendor do all the fun stuff is epitomized by S-MOS's CASP. If you have an unusual application that requires finesse, you can fine-tune the RISC architecture of the CASP ASIC's datapath. You can, for example, select exactly which instructions will make up your instruction set, and you can even provide for the launch of up to four instructions per clock cycle.

Selecting a 32-bit microprocessor for embedded applications

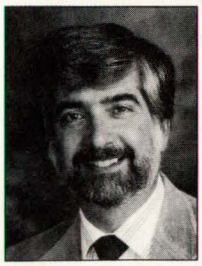

So you've decided that your next project will be based on a 32-bit microprocessor, but which one will you use? There are a number of factors to evaluate, and the most significant may be the processor used in your last application. The overwhelming majority of embedded system designers I've interviewed recently said they reused the same processor used in their last project, or at least a member of the same processor family.

But what if you're currently using an 8- or 16-bit processor and you have the freedom to select any architecture that meets the needs of your application? You'll have to analyze those needs carefully, and ask yourself some questions.

Price, performance important

Historically, price and performance have been primary criteria in processor selection. While price may still be an overwhelming reason to select a processor, performance may not. If you're moving up from an 8- or 16-bit application, you're likely to find that any of the 32-bit processors available will give you more than enough power to meet the needs of your design. Despite the ongoing arguments about CISC vs RISC, either approach is likely to satisfy the performance needs of most applications. RISCs tend to run somewhat faster, since their large register sets minimize memory accesses. Their simple instruction sets, though, mean they require more instructions to do the same work as CISC processors. Plan on having more memory available for storing your code if you base your design on a RISC chip.

A more interesting point to consider may be the level of integration available to you. The simplest microprocessors consist of only the CPU itself. Many microprocessor families include some members called high-integration parts, which add various kinds of peripherals to the basic CPU design. Some peripherals are general-purpose, while others may be tailored to a specific kind of application. Higher integration means a lower chip count and less real estate needed on your board. For embedded applications, where physical size is often an important limitation, this may be critical. It also may help reduce the cost of your product. At the very least, fewer parts will simplify your design.

At the most basic level, higher integration may simply be the presence of an on-chip floating-point coprocessor. If your application makes extensive use of floating-point computation, your product may not be competitive without a coprocessor. But the FPU is just the most obvious in a long list of available devices. Do you need counter-timer circuits? Most embedded applications do, at least for realtime clock generation or watchdog timers. Some microprocessors will have these available on-chip, saving you the trouble of adding them externally.

Watching for gremlins

How critical is the ability to catch memory access violations before they destroy valuable data—or perhaps the system itself? For some mission-critical applications, you may wish to consider processors that have an on-chip MMU. You may not need it for managing virtual memory, but the MMU can provide basic protection services that detect attempts to write into memory that you really intended to be read-only. In a small system, you might wish to select a processor with the capability of providing chip-select signals for memory or peripheral devices. You may be able to eliminate address decoding logic on your board if the microprocessor can do it.

While applications that run on 32-bit processors are almost always written in high-level languages such as C, most embedded applications contain at least some assembly language. While any microprocessor can be programmed in assembly language, some are easier to deal with than others. If your software engineers know Intel 8086 assembly language really well, then assembly language for the 386 or 486 is a relatively small step to take. Motorola 68000 assembler is widely known and easy to learn, and is similar to the assembly language used by several 8-bit CPUs.

In contrast, the three-operand style of assembly language used by most RISC processors takes some getting used to. For most RISCs, specifying a memory address and fetching its contents takes three separate instructions. You should plan to study the assembly language produced by your compiler to learn how to perform such basic operations.

Another basis for selecting a processor is the availability of support tools. Do you need in-circuit emulation? Many vendors provide emulators for CISC chips such as the Motorola 68000 family. You can select the one that best meets your needs. For a RISC processor, your choices will be more limited—or, there may not be an emulator available at all. There are certainly alternatives for debugging if no emulator is available, such as in-circuit monitors. These can even be left in your product for debugging in the field. However, emulators are particularly good at solving hardware/software integration problems. It's generally useful to have one around for debugging problems when your hardware isn't fully working yet.

Using what's available

Are you familiar with one of the commercially available realtime kernels? If so, you may wish to consider a processor for which that kernel is available. This will save you the trouble of learning a new operating system. Even if your application isn't going to use a realtime kernel, your next project might. Remember, you're likely to use the same processor you choose today in tomorrow's project, unless you want to start over with selecting hardware and software development tools.

You may wish to take advantage of the power and larger address space provided by 32-bit processors and implement your application in an object-oriented language such as C++. You can expect to get future products to market more quickly if you take advantage of the reusable code provided by C++, but it isn't available for all 32-bit processors yet.

Finally, you may wish to consider that software development is moving from the PC environment to UNIX workstations in many companies. This has improved productivity, but has caused some problems with compatibility. Almost any PC can run a PC application, but all UNIX workstations are different. Make sure the tools you need are available in the development environment you have chosen.

Kenneth F. Greenberg, BA in computer science, Microtec Research, Santa Clara, CA

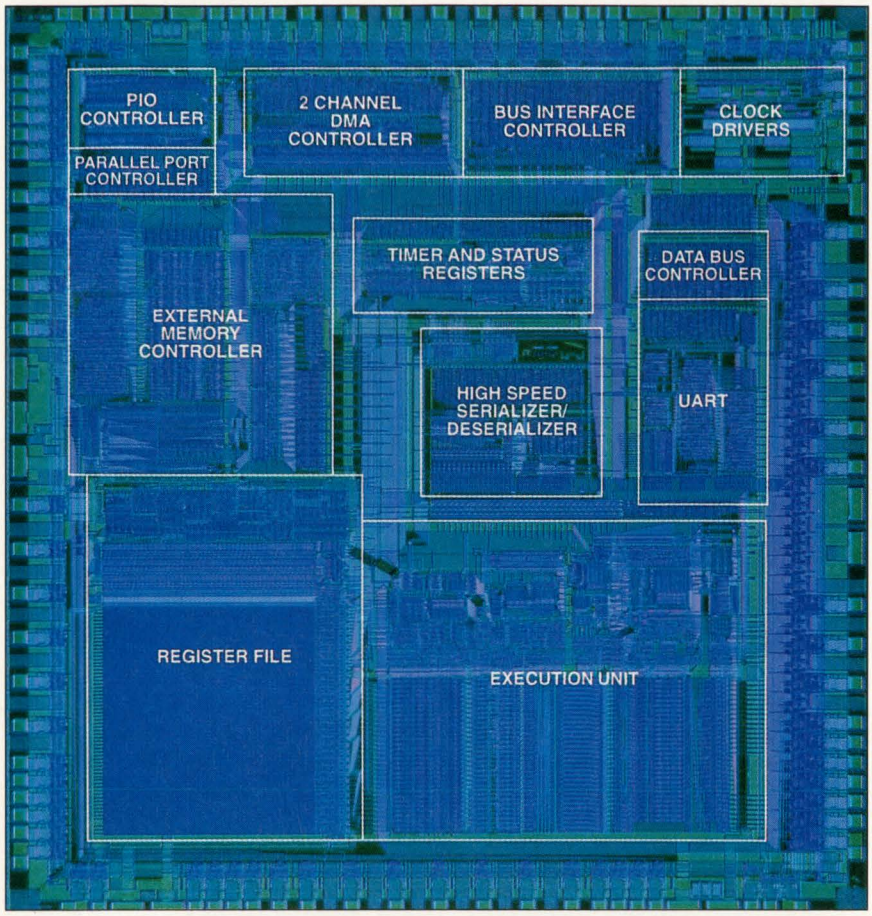

AMD's AM29205 RISC microcontroller is typical of the latest generation of highly integrated chips. The 205 is also an unabashed effort to capture design wins from 16-bit competitors. While it retains the 32-bit 29000 core, making it compatible with other members of the family, the datapath width has been trimmed to 16 bits. Despite having to use two clock cycles to load an instruction, the 205 offers better price-performance than 16-bit CISC competitors, according to AMD.

512 × 32 bits. Details on the company's PLD elements haven't yet been announced.

A new generation for the 29000

Another mature 32-bit product with a large number of design wins and an assortment of new microcontroller chips is the 29000 from Advanced Micro Devices (AMD—Austin, TX). The company's first 29000-architecture microcontroller was the 16-MHz 29200, with a full 32-bit core plus timer and memory and interrupt controllers, but no on-chip cache. In September, AMD introduced an economy version of the 29200, the 8-Mips 29205, which couples the family's 32-bit core with a 16-bit bus and an inexpensive 100-pin PQFP package. Intended to lure designers away from 16-bit CISC chips such as the 80186, the 29205 has greater performance than the 186 at a competitive price.

Even at its introduction in 1988, Intel's (Chandler, AZ) i960 offered an integrated half-kilobyte instruction cache along with the processor core. At that time, the civilian members of the family consisted of the i960KA and KB. The KB has a floating-point unit; the KA doesn't.

The lower-cost members of the family, the i960SA and SB, trim the external data bus to 16 bits, but retain the instruction cache. There's also an integrated interrupt controller that can handle up to four direct interrupts. The SB has an FPU.

The newest members of the family, the i960CA and CF, are superscalar, launching two instructions per clock cycle, at 16 and 33 MHz respectively. Relative to the earlier chips, the i960CA upgrades the cache from direct-mapped to two-way set-associative and increases its size to 1 kbyte. The i960CF has a 4-kbyte, two-way set-associative instruction cache and a 1-kbyte, direct-mapped data cache. Both of the superscalar chips integrate 1 kbyte of data RAM, four DMA channels and an interrupt controller that can handle up to 248 external interrupts.

R3000-based controllers

AMD's 29000 seized an early advantage in the cost-sensitive embedded arena thanks to its ability to interface directly to DRAM. This advantage is now eroding as members of other silicon vendors' families, such as the R3041 from Integrated Device Technology (IDT—Santa Clara, CA), offer similar capabilities.

IDT's MIPS R3000 controller family comprises the high-end R3081, with FPU and instruction and data cache (16 and 4 kbytes, respectively); the midrange R3051 and R3052, without the FPU and with smaller caches (4 or 8 kbytes of instruction cache, 2 kbytes of data cache); and, just introduced, the low-end R3041, with 2 kbytes of instruction cache, half a kbyte of data cache and some interesting special features.

The special features include programmable bus widths. The idea is to simplify the task of accessing 16-bit printer font caches and 8-bit boot ROMs, without sacrificing the advantages of a 32-bit bus for addressing memory and loading instructions in a single cycle. Four-level read and write buffers speed up memory operations.

The chip holds addresses longer—half a cycle past address latch enable—than conventional RISC chips, simplifying the interface to ASICs and FPGAs. The R3041 also provides fixed-map address translation instead of a translation look-aside buffer (TLB). The silicon real estate that was recovered from the TLB is used for an event timer.

IDT's 3081 illustrates one of the dilemmas inherent in creating chip families. As the highest-power, most fully featured member of the family, it seems logical that it be the one to incorporate the FPU. The reality of the market, however, is that the high end is dominated by telecommunications applications such as routers and bridges, which don't require floating-point computations.

"We're finding that telecommunications customers are ignoring the FPU and concentrating on the raw power of the chip," says IDT's director of RISC marketing, Bob Rowe.

For the SPARC RISC camp, the microcontroller candidate of choice is

continued on page 99

HOW MANY AMPS DOES IT TAKE TO POWER YOUR PRODUCT DEVELOPMENT?

ONLY ONE.

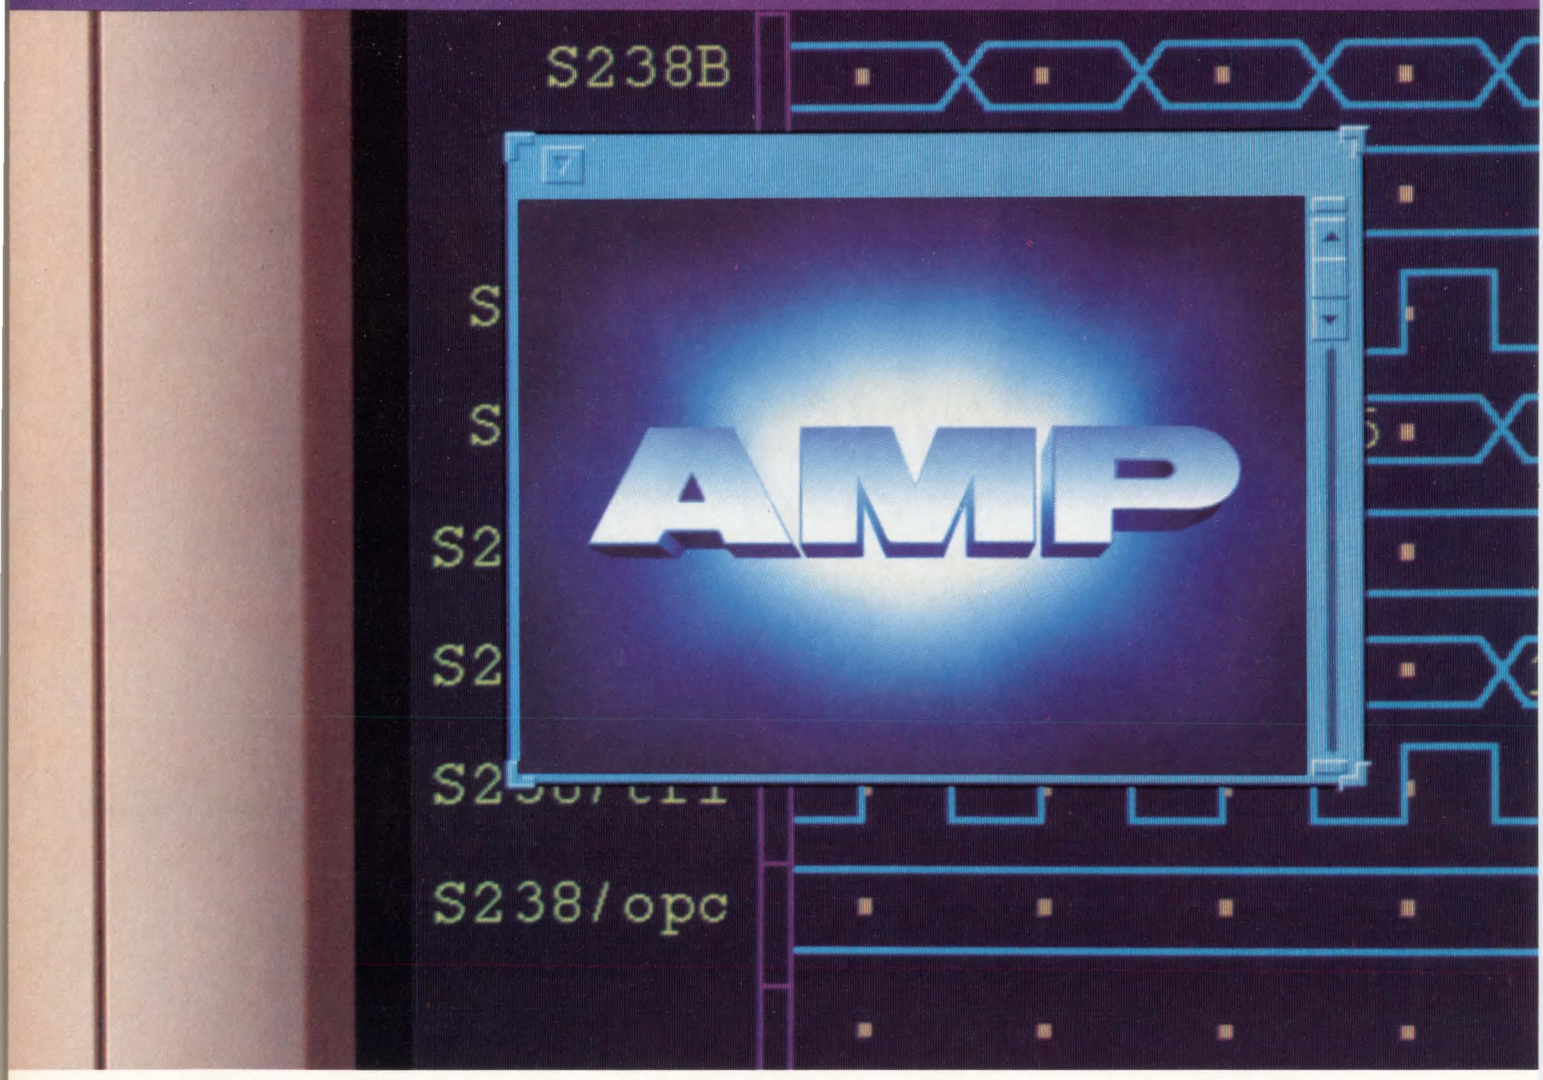

FROM IDEA TO MARKET.

The resources you need for that journey, especially with today's tight development schedules, can tax even the largest companies. AMP can help make your trip to market easier, and faster.

Our capabilities today go far beyond connectors . They include interconnect design, analysis, simulation and validation, systems development, and packaging design. And unequalled production experience in everything from 'metal-and-plastic' to heavy machinery, precision optics to SMT boards to complete plug-and-play systems.

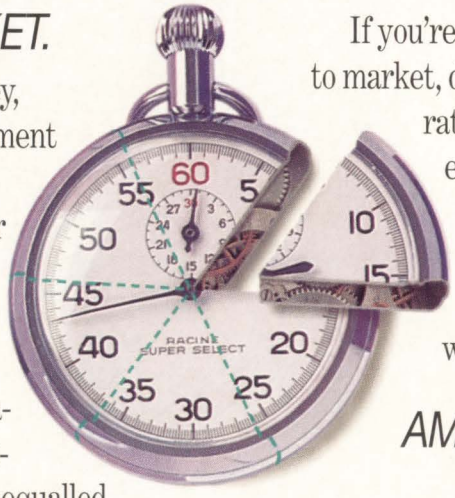

If you're concerned with reducing time to market, dealing with higher clock rates, coping with shrinking envelopes, and generally keeping pace with fast-changing technologies and design methodologies, look into what AMP has to offer, today.

AMP DESIGN SERVICES: EDA TO THE MAX.

One of the toughest jobs you face is analyzing and proving your initial design, and performing board placement and routing. We'll take you from

schematic or logic diagram all the way to a valid software first article.

Our comprehensive EDA toolset includes current, state-of-the-art software programs, and our own interconnect and device models–developed specifically for critical aspects of current design such as net simulation, timing verification, analog transient noise effects, system impedance, and transmission line layout rules.

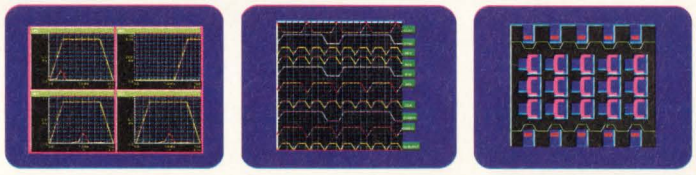

Characterized device and connector models are used in waveform analysis, timing verification, and board design.

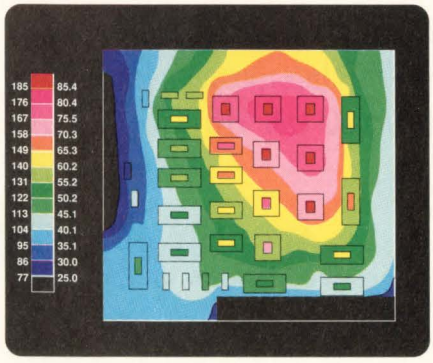

Placement optimization uses thermal analysis, along with electrical, power, and manufacturing considerations.

(Courtesy of BETAsoft, Dynamic Soft Analysis Inc., Pittsburgh, PA.)

Talk to AMP early in your design cycle, and we'll put you on board faster–with flexible service offerings that range from help with specific problems to complete project management. Tap us for the years of experience we've put into engineering under critical electrical, mechanical, power, thermal,

and manufacturing constraints. And shave a lot of time-to-market and development cost off your project.

AMP FABRICATION SERVICES: MAKING IT REAL.

Characterized models speed simulation; 2D and 3D CAD files simplify board layout and package design. (Models available on CD-ROM.)

Major elements of your job that add up to major time include printed circuit boards, backplanes, card cage design and fabrication, and systems packaging. Our help here can put you ahead of the crowd and let you spend that time on your real business.

Packaging is a crucial element of systems performance, and our modelling, simulation, power and thermal analyses, and noise budget balancing are crucial for packaging success. When you're ready, our manufacturing experience – from clean room to global shop floor – can pay off at any level you choose: board, subsystem, or finished product.

Look to AMP for the help you need to meet tight production schedules and narrow market windows, and the resources you want to move into next-generation technologies. Our knowledge base in fast boards, high-speed cable assemblies, surface mount technology, and related subjects can simplify the transition to faster logic and smaller form factors.

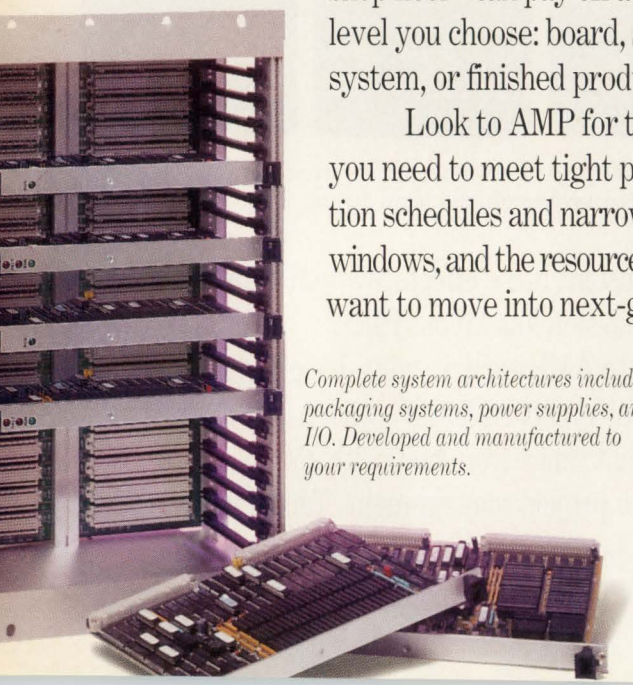

Complete system architectures include packaging systems, power supplies, and I/O. Developed and manufactured to your requirements.

Custom high-performance cable assemblies represent value-added capabilities in advanced design technologies.

THE BIG PICTURE: FROM CONCEPT TO PRODUCTION.

Your program, and your problems, are unique. Your program is yours alone, but your problems don't have to be. We can help in any of the areas we've listed here.

Or let us put it all together for you, and move you from start to finish. Our specialists can bring our services to you anywhere in the world, and coordinate the AMP resources you need to bring your idea into production. Cost-effectively. And on time.

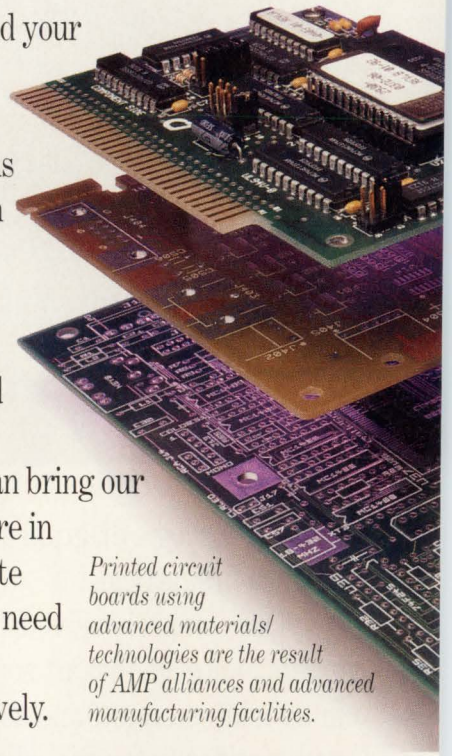

Printed circuit boards using advanced materials/ technologies are the result of AMP alliances and advanced manufacturing facilities.

For more information, call our Product Information Center at 1-800-522-6752 (fax 717-986-7575). AMP Incorporated, Harrisburg, PA 17105-3608. In Canada call 416-475-6222.

THIS IS AMP TODAY.

AMP

continued from page 94

Fujitsu's (San Jose, CA) SPARClite. Two years ago, Cypress Semiconductor (San Jose, CA) introduced its CY7C611 SPARC processor for embedded applications. According to Joe Nichols, marketing director for Cypress' Ross Technology subsidiary, this part "has a number of respectable design wins." Cypress, however, hasn't yet entered the microcontroller arena.

At $179 in 10,000-unit quantities, microSPARC from Texas Instruments (Dallas, TX) targets low-end workstations. The chip, however, has integer and FPUs, 4 kbytes of instruction cache and 2 kbytes of data cache, MMU and DRAM, SBus, and I/O controllers on-chip, so it isn't hard to envision microSPARC winding up in high-end embedded applications.

Below the $100 price barrier, Fujitsu's SPARClite microcontroller family offers the greatest number of choices. SPARClite represents the other horn of the controller family dilemma. In its processor core, it provides an integer unit only—no floating point, even though PostScript manipulations involve a great deal of floating-point manipulation. Fujitsu product marketing engineer Peter von Clemm says that PostScript printers will pay a modest performance penalty for handling floating point in software, but that this is outweighed by the price advantages of the family.

Three new members of Fujitsu's SPARClite family have joined the MB86930 processor, which debuted in 1990. The 20-/40-MHz 931 retains the first-generation chip's 2-kbyte instruction and data caches, bus interface with DRAM control, and interface to an emulator bus, while adding four counter/timers, interrupt controllers, and two serial channels. (The 930 used a companion chip, the MB86940, for these functions.) Gone is the 930's translation look-aside buffer.

In contrast, the 932 is more like a 930 with bigger instruction and data caches (of 8 and 2 kbytes, respectively). It retains the TLB of the 930 but does without the timers, interrupt control and USARTs of the 931.

Priced under $25 in quantity, the 20-MHz MB86933 is the price leader of the SPARClite family. To achieve its low price, it strips out cache, counter, timers, USARTs, and interrupt control. Also gone is the emulator interface, under the assumption that you would develop your application using a 930 or 931.

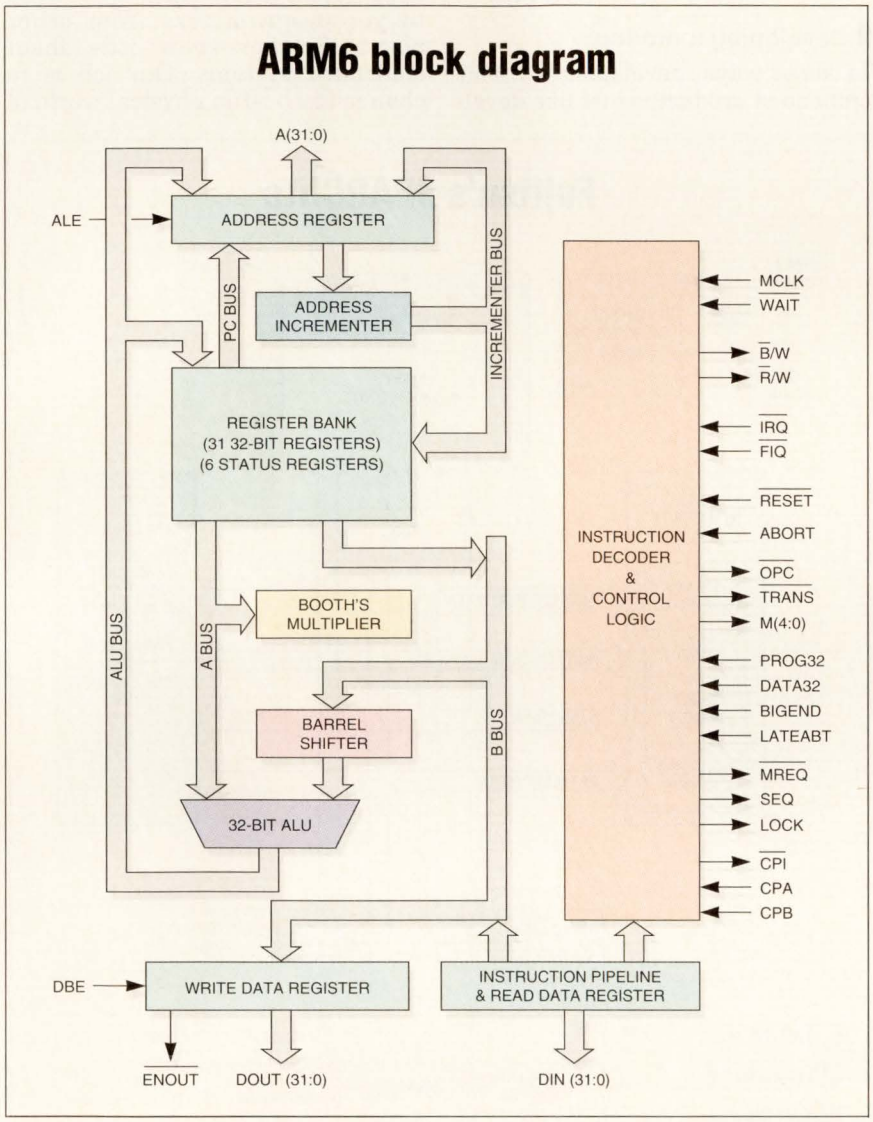

The architecture of the advanced RISC microprocessor (ARM), manufactured by VLSI Technology, is less familiar than other 32-bit options. Shown is the ARM6 CPU core that can be integrated as an ASIC cell. It has a full 32-bit address bus, rather than the 26-bit bus of its predecessor, with which it maintains backwards compatibility. ARM's strengths are small die size, low power consumption and the ability to handle either big-endian or little-endian data. The instruction set is large for a RISC processor; to handle it, VLSI Technology offers a C compiler, symbolic debugger and an emulator. Of the 31 general-purpose registers in the block diagram, 16 (including the program counter) are available in user mode. The remainder are mapped across ARM's interrupt, fast interrupt and supervisor operational modes.

Third-party choices

The level of integration that microcontrollers add to bare processors is just one advantage these new parts introduce. Third-party vendors can further reduce time-to-market by providing application-specific hardware and software using various silicon vendors' microcontrollers, but at a price. The ImageCard 8500 family from Adaptec (Milpitas, CA), for example, combines the company's printer ASICs with various AMD 29000-family processors and microcontrollers. VLSI Technology's LPIC is a controller chip you can use on your own board.

Alternatively, you can use DP-Tek's (Wichita, KS) TrueRes ASICs to create printers with software-selectable multiple levels of resolution. Phoenix Technologies (Norwood, MA) offers software PostScript and PCL-5 emulators that are 100-percent compatible with the Apple LaserWriter IIg, the Tektronix Phaser II PXi and the HP LaserJet III. Power-

Page, from Pipeline Associates (Morris Plains, NJ), also supports PostScript.

If you're not developing a printer, TeleSoft International's (Collierville, TN) FRAME Relay software works with the same processors to create WAN products. XLNT Designs (San Diego, CA) offers to do the same for FDDI networks.

There are also third parties lined up to give you a hand with manufacturing. Circuit Components (Tempe, AZ), for example, has under-chip decoupling capacitors personalized for popular processor and microcontroller packages. For prototyping, McKenzie Technology (Fremont, CA) provides an adapter that changes the 29205's PQFP pin-out arrangement to a pin grid array.

Getting your hands dirty

Many silicon manufacturers are going out of their way to insulate customers who are from the software side from exposure to the raw architectures of their products, but ASIC vendors with 32-bit cores are taking the opposite approach.

The 16-MHz, 7-Mips ARM is a product for mass-market applications created by Acorn Computers (Cambridge, England), Apple Computer (Cupertino, CA) and VLSI Technology. To appeal to both hardware and software types, VLSI Technology lets you either create your own custom ASICs with ARM processor cores or use ARM microcontroller products. There are a number of ARM ASIC cores, for example, but the ARM250 microcontroller, announced in September, contains a 32-bit RISC processor, memory, video and I/O controllers, and a PC bus interface. It interfaces directly to DRAM and ROM. Video and sound (SVGA and stereo) capabilities include on-chip digital-to-analog converters.

The ASIC vendor who goes farthest in encouraging customers to play with chip architecture is S-MOS (San Jose, CA). You can actually fine-tune the RISC processor core in the company's CASP (Configurable Application-Specific Product) ASICs. In fact, you can even go so far as to add instructions to the RISC instruction set if it helps your application run better, according to S-MOS manager of strategic marketing, Dr. Richard Ahrons.

You don't need an in-depth knowledge of silicon processing to use CASP, either. Even for the small development team, in fact, Ahrons says there's no need to be daunted by the prospect of creating a "tuned architecture" processor core or other high-level functional block. S-MOS is a Seiko Epson affiliate, and the company is prepared to make its considerable design resources available to customers.

To keep development time under control, manufacturing CASP products follows the semicustom ASIC model. S-MOS starts fabbing wafers as soon as you have defined the architecture, but the company holds them short of metallization until you've designed the control logic.

Developing a product

In some ways, developing a 32-bit embedded product is just like developing any other embedded product. At least the steps are the same. However, says Applied Microsystems' (Redmond, WA) vice-president of new business development, Dick Jensen, "Managing the development effort is tougher. The team is bigger, the code's bigger, and you need more discipline."

Jensen agrees that programmers for 32-bit systems are different from their assembly-writing counterparts, who work with narrower buses and longer instruction sets. He says, "The complex nature of the applications means that we're forced to get programmers from other places. So they know little about embedded systems. Our job is to change the hostile physical world of

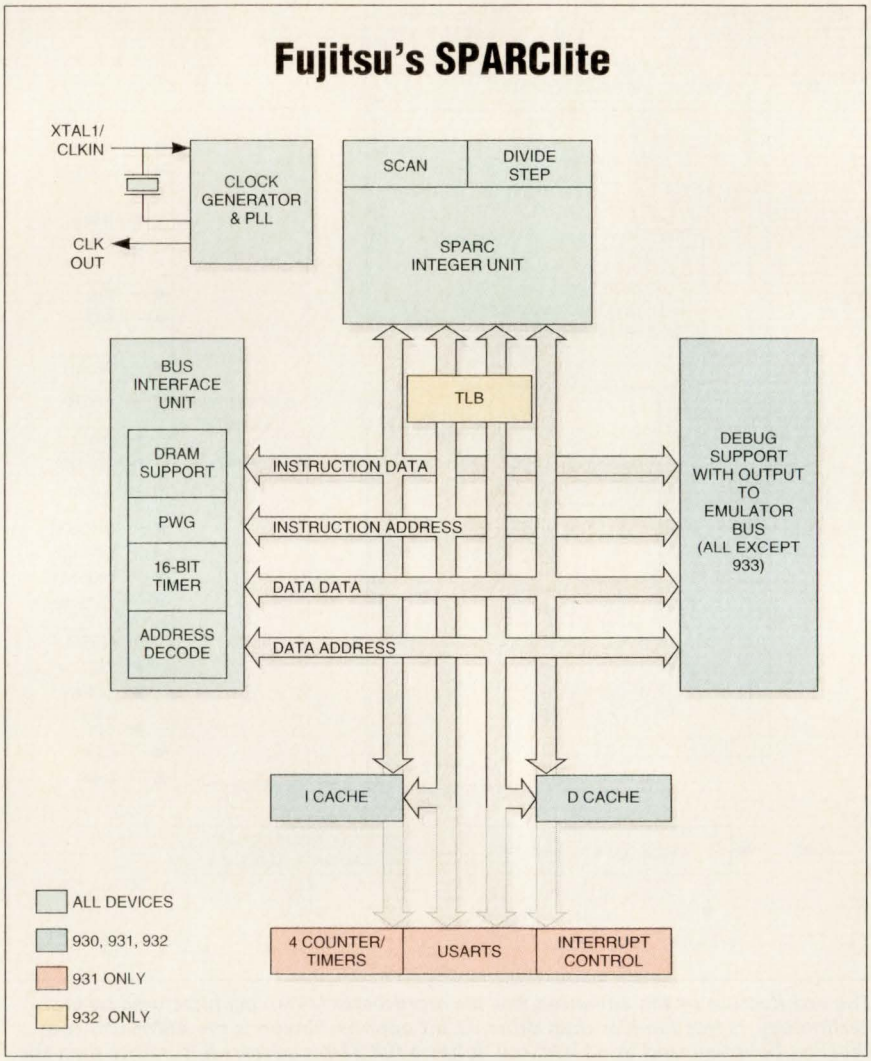

Using chips with a high degree of integration makes designing your application easier—once you've decided what you want integrated. For instance, every member of Fujitsu's SPARClite family integrates the integer unit and bus control. Most integrate the emulator bus interface. After that, you get to choose: cache (the 930 has 2 kbytes, the 932, 8 kbytes) or no on-chip cache, data cache size (2 kbytes for both the 930 and 932) if you choose on-chip cache. Serial interface, timer, interrupt control, and TLB are further options.

Choosing development tools: a microprocessor vendor's viewpoint

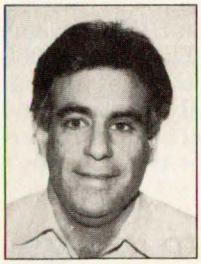

Developers of embedded systems who use Motorola's 68000 and 68300 families have relied on a wide range of development software and hardware. Helping them bring their products to market has shaped my own ideas about tools and the reasons that people choose them.

Often, companies pick tools simply because they fit into their existing installed base of development tools. Unless they are really dissatisfied with what they've been using, they find it desirable to upgrade what they already have. That only makes sense. The cost of starting from scratch is high, and a learning curve that takes months puts a company at a disadvantage in getting a new product to market.

Even if the tools represent a new generation, the deciding factor can be as simple as whether a particular tool runs on the platforms the development team already has (or has decided it can afford).

Checking out competing RTOSs

Embedded systems that require a predictable response time to system events or that involve a risk of data loss, damage to equipment or injury to personnel obviously must run under a realtime operating system. A sensible starting point for evaluating an RTOS is to ask how well it's regarded in the embedded systems community. How long has it been around? What is the general opinion regarding the quality of the kernel code?

Something worth considering at the same time is the reputation of the company offering the tool with respect to technical support. Applications development isn't a suit-and-tie, 8-to-5 business, and engineers and programmers often do their most productive work at odd hours of the night. That's why a 24-hour hotline can be important.

In matching an RTOS to the application, the trick is not to reinvent the wheel by writing your own code for things such as task swapping and I/O. If the company supplying the kernel can also supply the right I/O drivers for a particular application, that substantially reduces time-to-market for the application developer. Similarly, if the application requires higher-level services, such as a file system or a networking layer, it can save even more time if those are available with the RTOS. Also, a high-level debugger for the RTOS kernel is just as important as a symbolic debugger is for the compiler. The issue again is cutting the fat out of the development cycle.

Obviously, an important key to an RTOS is predictable response time. Vendors of RTOSs provide tables or formulas you can use to determine response time for specific tasks, based on what other tasks are currently running. In some cases, the formulas may include variables; this means that, for certain kernel services, response times wouldn't really be deterministic. This may or may not be critical.

On a more mundane note, the OS licensing structure can also influence the purchasing decision. Is there a fixed price for a single copy, with a sliding scale for multiple copies? Are site licenses available?

At a level above the OS, the choice of a compiler depends mainly on contract requirements or on the philosophy of the group developing the application. Ada may be a requirement on a government contract. C++ may suit a company committed to reusable code and the advantages of object-oriented programming. On the other hand, because of the large number of programmers who are familiar with it, C continues to be the language of choice for most new systems. Regardless of the HLL (high-level language) symbolic debug is necessary—that is, debug tools that understand the compiler is essential.

The emulation trap

To rely or not to rely on in-circuit emulation boils down to a philosophical issue. These days, we're seeing a trend in which our most sophisticated customers use far more up-front simulation, and depend far less on target debugging using emulators. It's a sign that more attention should be paid earlier in the cycle as to how the design is going to work.

From a time and a dollar standpoint, getting it right in simulation before you commit to hardware is easier if there are ASICs as well as a microprocessor on the board. It may be every bit as expensive in terms of dollars and time to revise a board as it is to revise silicon, but not being able to back-wire around a problem on an ASIC the way you can on a board makes the situation more absolute.

Simulation notwithstanding, developers can't completely walk away from emulation, and this presents something of a challenge to microprocessor manufacturers. A highly integrated microcomputer, or even just a powerful processor with on-chip cache, can operate for some time by accessing internal resources, without giving any indication to the outside world as to what's happening internally.

To deal with this, we provide hooks for debugging. For example, we have a mode in which we essentially put the processor to sleep and let the emulator mirror the processor externally. Our newer 68300 family of processors has a special debug-mode port, an alternative to using an emulator probe when the device is in a surface-mount package. Depending on the processor you select, the lack of silicon-vendor-supplied features like these may limit your choice of emulators. In-circuit emulation forms a small market, populated with many smaller manufacturers with limited resources, and not every vendor can support every processor.

Future scenarios

At this time, the 32-bit embedded systems marketplace is demanding more integrated solutions. Currently, manufacturers such as Motorola are answering the demand with specialized chips which target specific application areas—for example, data communications in the case of the 68302. However, the day isn't far off when customers will be able to specify exactly what they want on their chips. When that day comes, some of the considerations outlined above will be even more important. Making sure that operating systems and compilers have the necessary extensions that make it simple to exploit unique chip resources will add to the complexity of choosing development software, and the need to get it right the first time will shift the development model farther away from emulation and closer to simulation.

Art Parmet, senior factory representative, high-performance processor group, Motorola SPS, Woburn, MA

the chip into the virtual protected world that a software engineer can be productive in."

A choice of operating systems

One decision 32-bit developers face that their narrow-bus counterparts generally have an easier time with is the choice of an operating system (OS). As is the case with hardware, there are many choices here, too.

Industry analyst Andrew Allison says, "One of the key attributes of a realtime executive is that it be ROM-able–that is, that it execute out of read-only memory, because the typical embedded application doesn't have a mass-storage system. This requirement obviously places a premium on compactness of code, which has the serendipitous benefit of speeding execution." Allison also notes that the realtime operating system can mitigate the impact of processor context switching and interrupt latency on deterministic behavior of the system.

Among the choices there are some old and some new items. One of the more familiar is C EXECUTIVE from JMI Software Consultants, whose applications also run under UNIX for development. C EXECUTIVE lets you hand-optimize code using assembler for context switching, task scheduling, interrupt handling, and block data moves. It has a fully preemptive, prioritized task scheduler, standard and device-driven I/O, and message queues for data transfer. For RISC processors, it can fit into less than 64 kbytes.

Accelerated Technology's (Mobile, AL) Nucleus RTX provides a realtime multitasking executive for AMD's microcontrollers. Other new developments include Ready Systems' (Sunnyvale, CA) Spectra cross-development environment, which currently is only available for the 68000 family. Spectra facilitates debugging without a hardware target by providing a "virtual" software target. Spectra updates the company's VRTX32 kernel with VRTXsa, providing a 30-percent improvement in time-critical calls. The Spectra toolset is open to third-party tool developers.

"Our job is to change the hostile physical world of the chip into the virtual protected world that a software engineer can be productive in."

—Dick Jensen, Applied Microsystems

At about the same time that Ready announced Spectra, Wind River Systems (Alameda, CA) introduced MicroWorks, its development environment for the 68000 and 80960 families, with SPARC and R3000 support coming. The target can be connected to the processor via an in-circuit emulator (ICE) cable, or more economically by an RS-232 serial port.

Integrated Systems' (Santa Clara, CA) pSOSystem is an update that integrates the company's pSOS realtime kernel with C, C++ and Ada compilers. User options include complete UNIX networking, file management and graphical interface software. The company's XRAY+ debugger now handles C++ language features.

When embedded programmers coded in assembly language, debugging was simpler. Optimizing compilers have changed all that. Compiler optimizations include moving invariant expressions outside of loops; loop rotation, in which the test is placed at the bottom of the loop (saving one branch instruction per loop iteration); and replacing multiplications with much faster adds and shifts. Expressions that can be computed are turned into constants,

Software trick receives patent

Illustrating the tight coupling between silicon and tools that embedded systems foster, the U.S. Patent Office has issued a patent to JMI Software Consultants' Susan Wainer for a "trick" used in the company's i960 version of its C EXECUTIVE RTOS.

The idea covered by the patent saves 13 instructions and 33 memory accesses every time there's a process interrupt with task preemption. It even saves one instruction and 20 memory accesses on interrupts without preemption.

Copying takes time

Wainer's trick involves "lying" to the processor through the processor state flag in the i960's process-controls register. The two states the flag can have are *executing* and *interrupted*, and normally, the names are apt descriptions of the flag's state. Each user task has its own procedure stack. The first interrupt that occurs creates an interrupt record and frame at the top of the interrupt stack. Succeeding or nested interrupts create additional records and frames on the interrupt stack.

This creates a problem when the kernel, processing an interrupt, determines that there must be a context switch to a higher priority task. The interrupt stack will have to be overwritten, so the current interrupt record and frame must be copied into a task control block, using up cycles and time.

JMI's approach is faster. When a process is executing, C EXECUTIVE lies and sets the processor state flag to the interrupted state. Then, when an interrupt occurs, the processor creates the interrupt record and frame on the current stack—the task procedure stack.

If the interrupt occurs while a user task is executing, the kernel lies again and sets the processor state flag to the executing state. Subsequently, if another interrupt occurs during the processing of the first interrupt, the processor switches to the interrupt stack and the kernel leaves the processor state flag set to *interrupted*. All succeeding interrupts create frames on the interrupt stack.

As a consequence of this use of the state flag, when there's a preemption to a higher-priority task, the kernel doesn't have to copy data from the interrupt stack, because the user task information is already on the interrupted task's procedure stack. The savings, in instructions and memory accesses, is substantial.

There is also a further benefit in saving global registers in local registers during interrupt processing. Normally, the global registers would be saved on the stack, using loads and stores. With JMI's innovation, however, global registers can be saved in local task registers, where they will be safely preserved in the event of a context switch.

and the most frequently used variables are kept in internal registers.

Obviously, closer links between the OS and the debugging software help you understand what your debug traces are showing you. In one example of closer linking, Applied Microsystems and Wind River recently got together to let information about data structures, flags and semaphores pass between the kernel and debugger.

One problem in using source-level debuggers and ICEs in 32-bit designs is the inability to examine the contents of processor registers directly in realtime. You have to either single-step or look at the entire realtime trace buffer and calculate changes in register values manually. Applied Microsystems has attacked this problem with an inference engine in its XICE source debugger for Motorola's 68330, 68340 and 68F333 processors. The so-called "intelligent trace disassembler" builds a model of the processor state and modifies it as instructions execute. The result is an accurate history of hardware register values and corresponding instructions at any place in the trace.

Target integration

Every time processors get wider buses or faster clocks, the engineers who design in-circuit emulators get a new set of challenges. So far, however, they haven't been beaten.

The Applied Microsystems EL3200 supports Intel's i960 and Motorola's 68000 families. STEP Engineering's (Sunnyvale, CA) Excell and Eclipse emulators support, respectively, Fujitsu's SPARClite and AMD's 29200 and 29205 with full ICE. In the case of SPARClite, the availability of an emulation interface facilitates the implementation of an in-circuit emulator.

Also worth noting is that, even though it runs at 40 MHz, STEP's Excell doesn't buffer the pins of the chip, letting the probe run with the same timing as the actual chip in the target system.

More economical than full ICE, STEP provides a JTAG 29K emulator for the 29205 which takes advantage of the boundary-scan capabilities of the new AMD chips. Applied Microsystems' low-end CodeTap emulator supports i960 and 68330 families.

Other ICE suppliers include Embedded Performance (Santa Clara, CA), which provides emulators for SPARC, MIPS and 29000 chips. Its most recent announcement is the SYS29K-PUMA for the 205. ■

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 RISC Machines
(408) 399-5195 ..... Circle 257
Accelerated Technology
(205) 661-5770 ..... Circle 258
Adaptec
(408) 945-8600 ..... Circle 259
Advanced Micro Devices
(800) 292-9263 ..... Circle 260
Applied Microsystems
(206) 882-2000 ..... Circle 261
Circuit Components
(602) 967-0624 ..... Circle 262
Cypress Semiconductor
(512) 892-7802 ..... Circle 263
DP-Tek
(316) 687-3000 ..... Circle 264
Embedded Performance
(408) 980-8833 ..... Circle 265
Fujitsu Microelectronics
(408) 456-1000 ..... Circle 266
Integrated Device Technology (IDT)
(408) 492-8631 ..... Circle 267
Integrated Systems
(408) 980-1500 ..... Circle 268
Intel Corporation
(800) 548-4725 ..... Circle 269
JMI Software Consultants
(215) 628-0840 ..... Circle 270
McKenzie Technology
(510) 651-2700 ..... Circle 271
Microtec Research
(408) 980-1300 ..... Circle 272
Motorola, Microprocessor and Memory Technologies Group
(512) 891-2037 ..... Circle 273
Phoenix Technologies
(617) 551-4000 ..... Circle 274
Ready Systems
(408) 736-2600 ..... Circle 275
S-MOS
(408) 922-0578 ..... Circle 276
STEP Engineering
(408) 733-7837 ..... Circle 277
TeleSoft International
(901) 854-5267 ..... Circle 278
Texas Instruments Semiconductor Group
(214) 995-6611, ext. 3990 ..... Circle 279
VLSI Technology
(408) 434-7520 ..... Circle 280
Wind River Systems
(510) 748-4100 ..... Circle 281
XLNT Designs
(619) 487-9320 ..... Circle 282

Postscript:

Looking ahead, 64-bit embedded systems are on the horizon. IDT will probably be first to market with an equivalent of the R3051 for the R4000. The internal code name is Orion. Version 9 of the SPARC standard, which defines a 64-bit architecture, has been announced, although no SPARC licensee has yet said it's working on silicon.

For very powerful embedded systems—for example, in medical imaging—TI's SuperSPARC and Cypress' hyperSPARC are possibilities. By integrating superscalar IU, FPU and cache on pretested modules and providing the relatively easy-to-design-to M bus as a system interface, these chips supply an upscale version of less potent monolithic microcomputers. They also permit multiprocessing—with up to four processors, at least—for applications that can take advantage of it.

In general, 32-bit and wider embedded systems will maintain their differences from more traditional 4- and 8-bit embedded applications. Code will be written in high-level languages by engineers who are less familiar with fine-grain architectural details than they are with the applications for which they're coding. At the same time, chips will be more highly integrated to help insulate developers from hardware design decisions. Design teams will be bigger, and the design effort will be a more complex management task. And finally, for most companies, time-to-market and pricing constraints will continue to shrink.

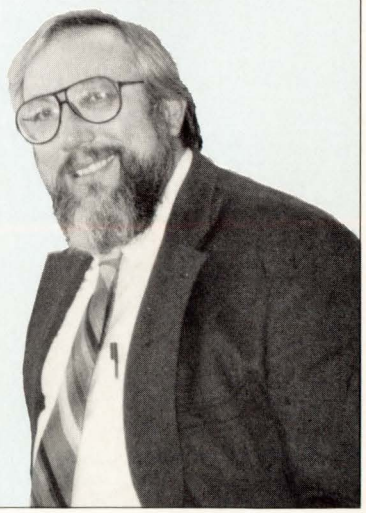

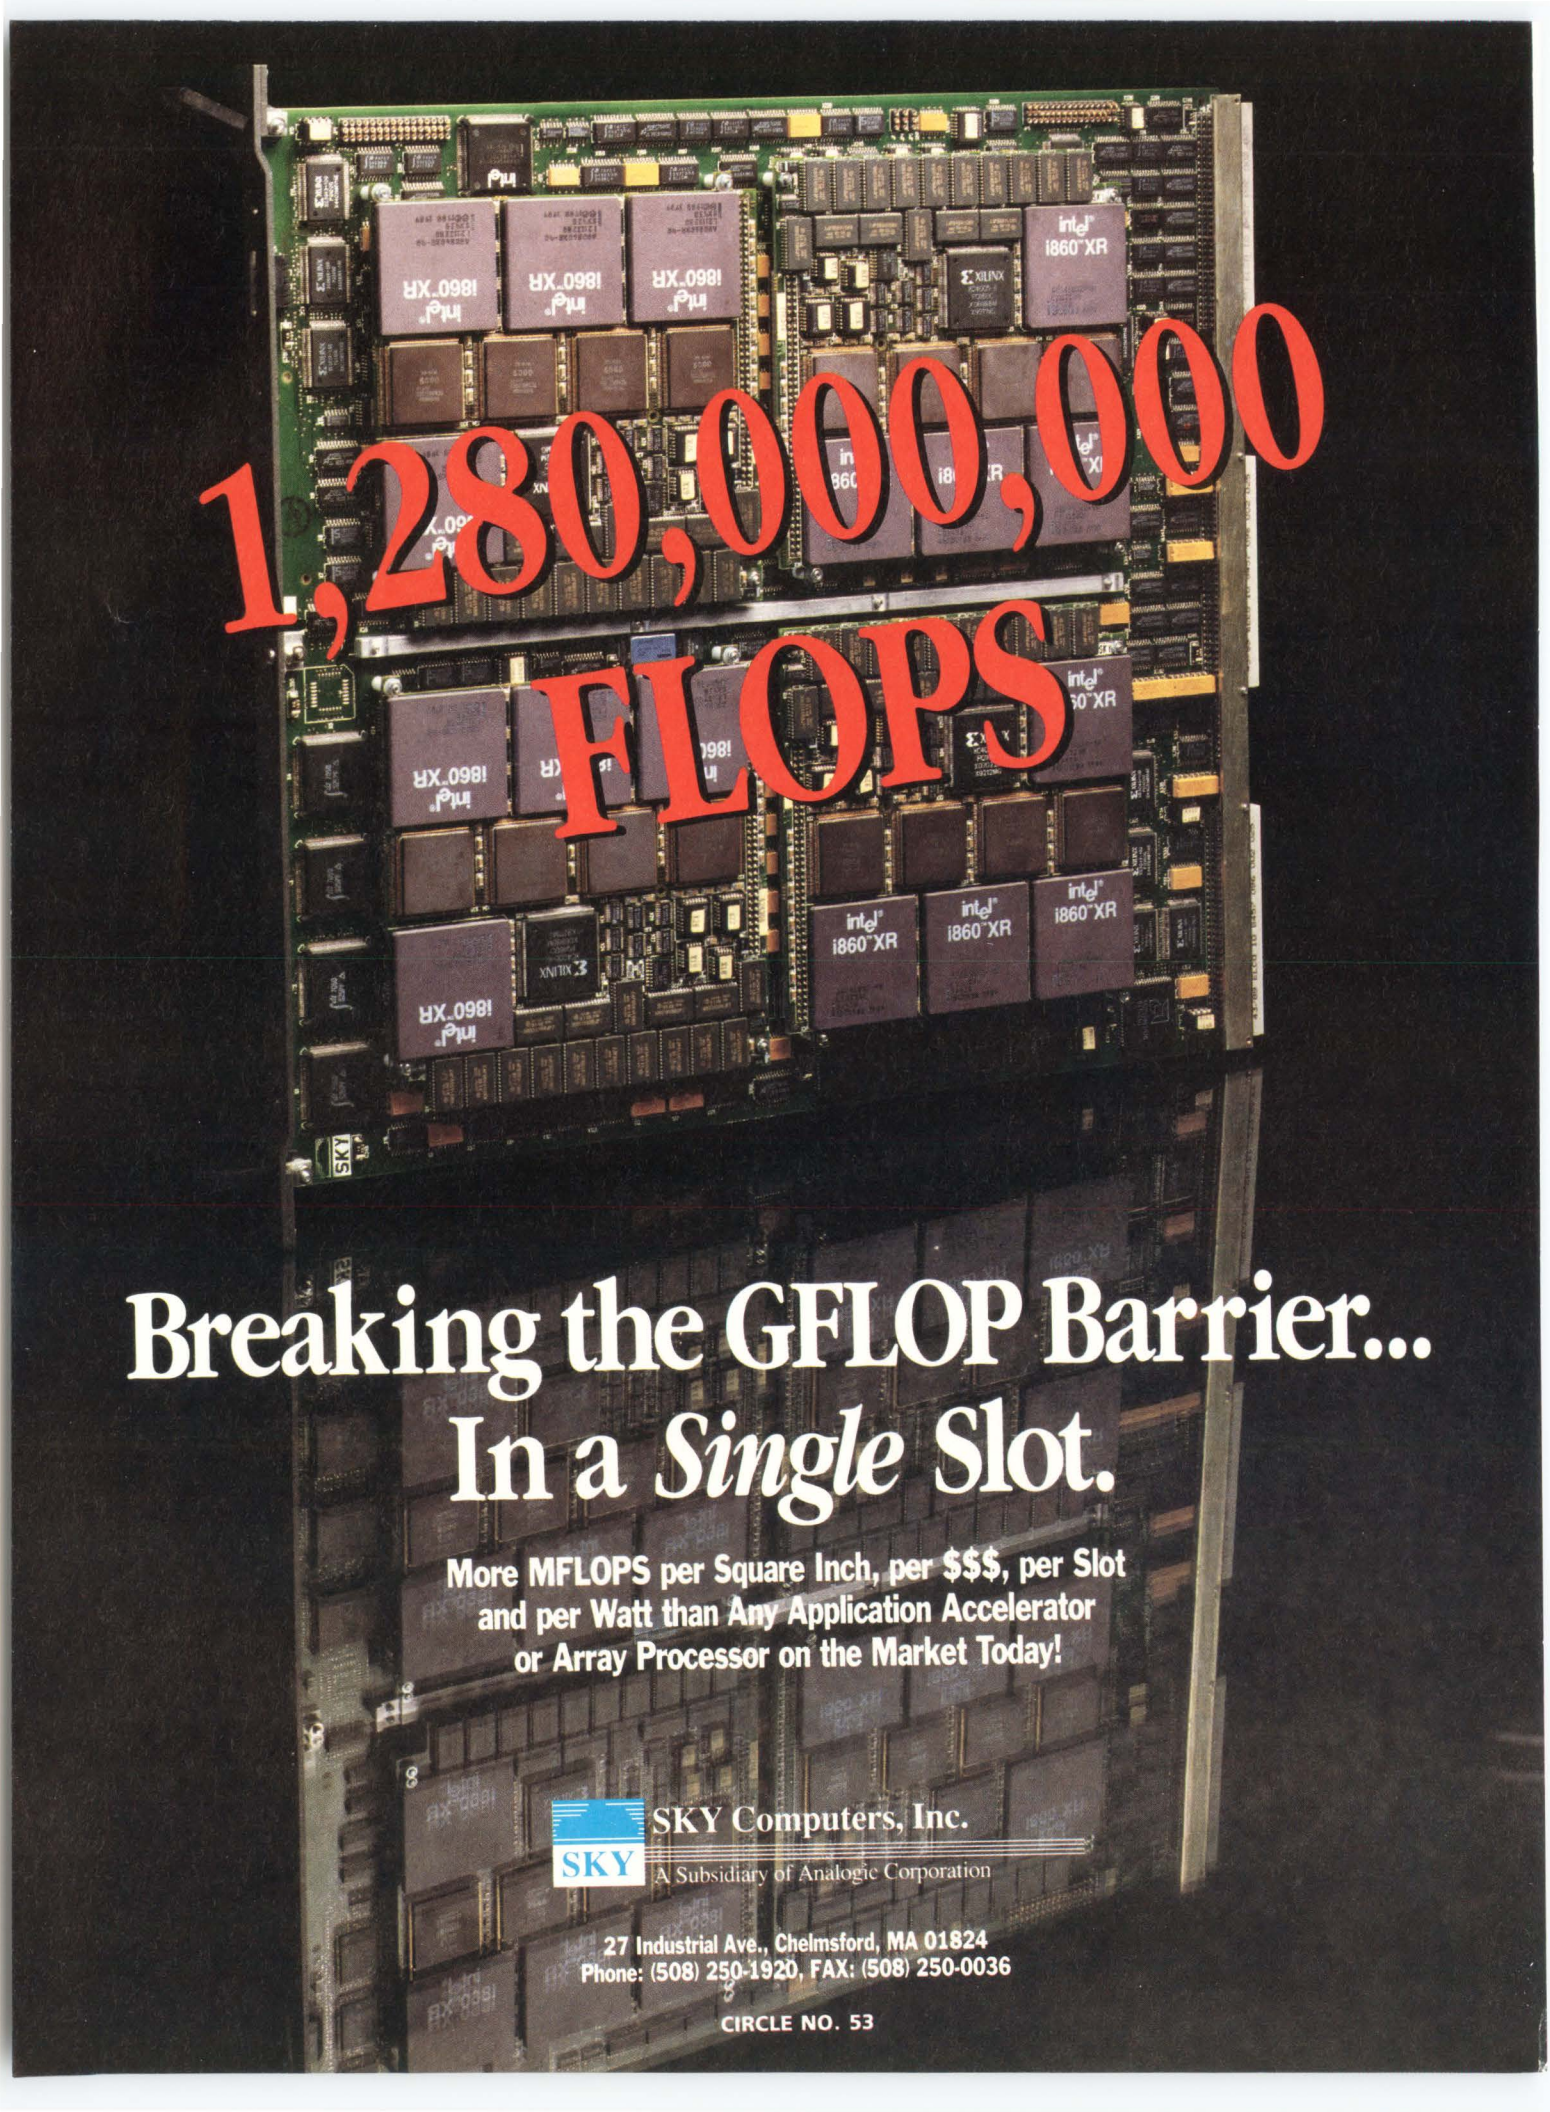
1,280,000,000
FLOPS
Breaking the GFLOP Barrier...
In a *Single* Slot.
More MFLOPS per Square Inch, per $$$, per Slot
and per Watt than Any Application Accelerator
or Array Processor on the Market Today!
SKY Computers, Inc.
A Subsidiary of Analogic Corporation
27 Industrial Ave., Chelmsford, MA 01824
Phone: (508) 250-1920, FAX: (508) 250-0036
CIRCLE NO. 53

Truck simulator integrates off-the-shelf subsystems

The TT150 Truck Driving Simulator design team from left to right (sitting): Rich DeFouw, software program manager; Perry Paus, engineer; Darrell Davey, senior engineer; from left to right (standing): John Eisenhardt, chief scientist; Karen LaFond, engineer; John Dibbs, senior engineer; Al Berrie, director of engineering.

Four years ago, FAAC Incorporated (Ann Arbor, MI), a software company in the defense industry, saw an opportunity to create its first commercial product. A phone call from a former employee sparked the idea of designing a truck driving simulator—one realistic enough for trucking companies to use as an alternative to on-the-road training. After keeping the idea as a back-burner market research project for a year, FAAC decided to pursue a design concept that would eventually become its TT150 Truck Driving Simulator.

While FAAC has a reputation for success in creating training simulators for the military, the company was cautious about jumping into the commercial market. "We didn't want to be a defense company telling the commercial world, 'This is what you need,' " says Albert Berrie, director of engineering at FAAC. "So we created an industry advisory council."

Advisory council provided expertise

FAAC brought in people from both the trucking industry and academia who were involved in driver safety and trucking issues. The council included members from large trucking fleets, such as Federal Express, North American Van Lines and Roadway Express, as well as the University of Michigan's Transportation Research Institute. "We met with these people and bounced our ideas off them," says Berrie. "They told us what they thought the simulator product should do, and where it fitted into the whole training scenario."

Jeffrey Child, Senior Editor

16 lines, No waiting.

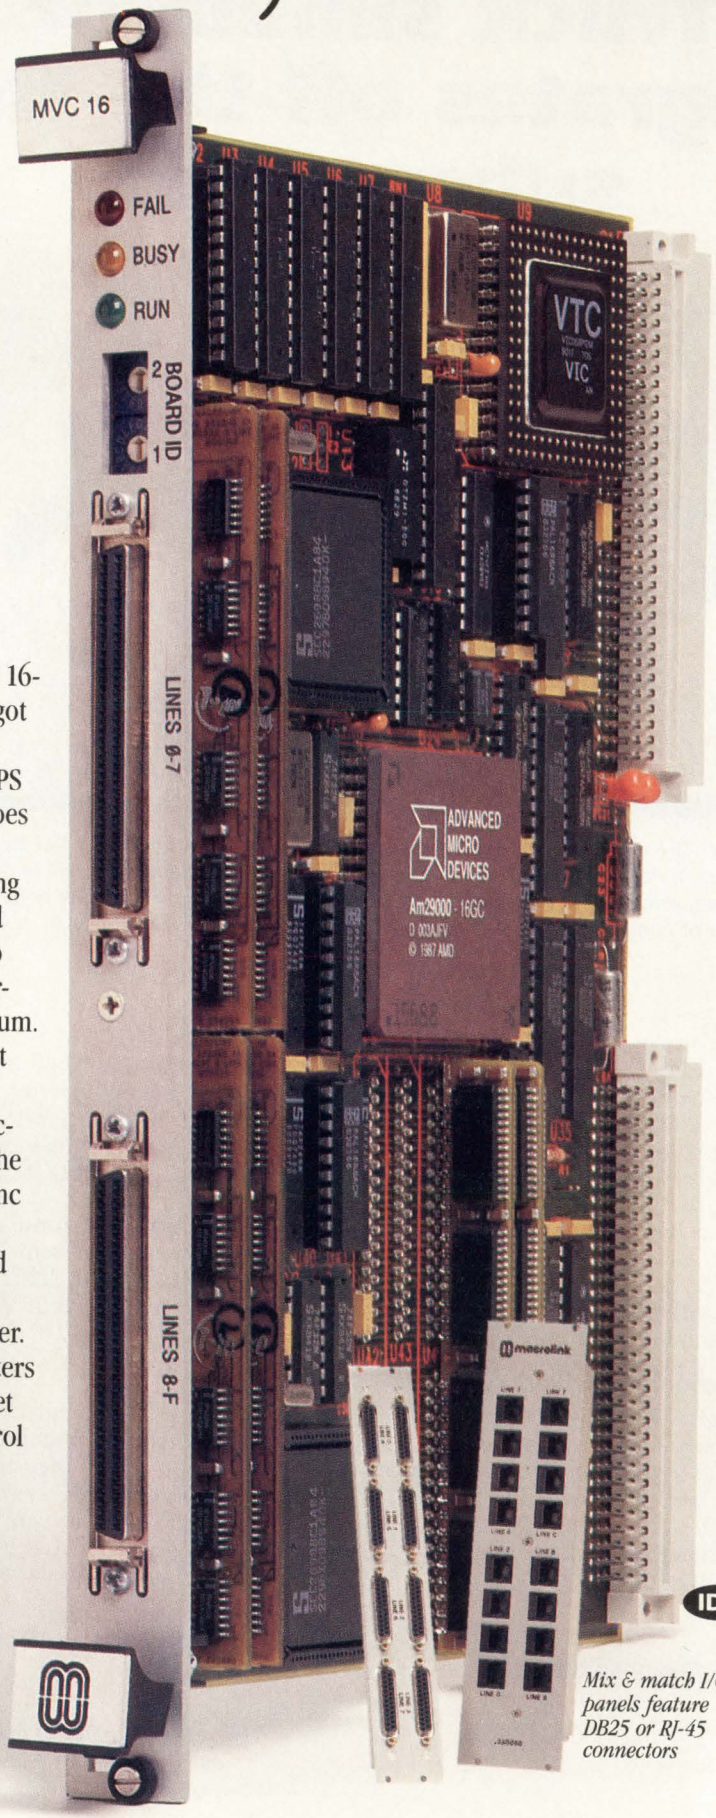

Mix & match I/O panels feature DB25 or RJ-45 connectors

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

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

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

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

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

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

macrolink® inc.

CIRCLE NO. 54

IDA

INSTANT DATA ACCESS (IDA)

DIAL (617) 494-8338 DOCUMENT NO. 1037

TT150 Truck Driving Simulator: how it works

The TT150 uses computer-generated imagery to provide a fully interactive truck driving simulation. Fundamentally, the system consists of three computers: a VME dual-board Motorola Delta system, a Ball 994 graphics system and an Amiga 2000 HD desktop system.

The Ball system generates the images on three projectors. It runs rendering software and processes display lists. A 68040 board in the Delta system performs all the paging of terrain database information, loads polygons and provides eyepoint information and vehicle position information. The 68040 receives realtime data concerning the position and orientation of the vehicle from a RISC 88000 board in the Delta. This data is passed to the Ball, along with a list of the current polygons within the view volume.

Residing in the Motorola Delta system, the RISC 88000-based MVME181 board performs several duties. It contains a sophisticated simulation model of a tractor trailer rig, including a full 14° of freedom. The 88000 handles height-above-terrain (HAT) calculations, keeping track of six contact points on the road, which approximates 18 wheels.

Finally, all the parameters that drive the sound system are defined through the 88000. When the brake is applied, for example, the sound system generates an air-release sound.

In the Ball graphics system, software performs rendering processing and forms the images out of the polygons in the immediate area, drawing pictures based on what the Motorola boards tell it. The 68040 board tells the Ball where the truck is located and passes it a list of polygons. The 68040 also controls two VMIC I/O cards, one digital and one analog, that interface with hardware in the cab such as the steering wheel, gauges, clutch, and gearshift levers. The truck model responds to changes in throttle and passes new orientation data back to the 68040's UNIX system.

The Motorola Delta system accumulates scoring information—current speed versus the speed limit, for example. At the end of the run, this information is processed and passed back to the Amiga, where the data is graphically represented.

Over time, FAAC's ideas and the advisory council's feedback developed into a design strategy. Because the company was entering the commercial market, the driving force behind its strategy was cost. "Everything we did we did with an eye towards minimizing the cost," says Berrie. "We tried to target the cost at roughly what it would cost a fleet to buy one full truck rig and trailer—about $150,000."

Partnership aided process

With itself as the managing partner, FAAC formed a partnership called Professional Truckdriving Simulators. The other partner, Perceptronics (Woodland Hills, CA), did the mechanical design of the simulator, building the cab, the instrument panel and the enclosure. In addition to overall program management, the FAAC team was responsible for the computer design, developing the software and evaluating, selecting and integrating the TT150's computer hardware.

Users of the TT150 sit in a mock-up of an actual truck cab. It contains a seat, steering wheel, clutch, gearshift lever, and all the gauges found in a typical truck. Three large graphic displays provide a 180° panoramic view of the road, along with inset images representing two rearview mirrors. Large bass speakers in the TT150's cab create the sound and vibration of a truck's diesel engine. The TT150 also has an Amiga 2000 HD computer as an instructor station from which driving conditions can be varied and the driver's performance monitored.

The most fundamental requirement stipulated by the advisory council was that the simulator behave as much like a real driving simulation as possible. "You can't create emergency risk-of-death situations at a training school or in a vehicle," says Berrie. "The simulator gives the instructor the ability to create failures in situations that you

Announcing CHANNEL EXPO- The Micro Channel® Product Showcase and Technical Conference running concurrently with BUSCON and SOFTCON

SOFTCON WEST

FEBRUARY 9-11, 1993
SAN JOSÉ CONVENTION CENTER
SAN JOSÉ, CA

The National Software Conference For System Engineers

and

BUSCON/93 WEST℠

FEBRUARY 9-11, 1993
SAN JOSÉ CONVENTION CENTER
SAN JOSÉ, CA

Because There's a Soft Side To Every Hard Sell

Today's system solutions call for an integrated approach to hardware and software. That's why we've put BUSCON and SOFTCON together.

BUSCON, the leading trade show and conference for bus boards and other elements of systems architecture, now runs concurrently with SOFTCON, the national software conference for system engineers. The combined focus of BUSCON and SOFTCON bridges the gap for system integrators, engineers, and their managers seeking total solutions.

Coordinated by *Electronic Design*, BUSCON and SOFTCON offer the hardware and software solutions you need now. Our flexible format allows you to mix and match sessions from either program. And both conferences offer entry to the combined BUSCON/SOFTCON exhibit hall.

If you're looking for system solutions you can use now, then plan to attend BUSCON '93 West and SOFTCON '93 West, February 9-11, 1992 at the San Jose Convention Center, San Jose, CA. For more information, just return the coupon or call (800) 243-3238 or fax (203) 857-4075.

SOFTCON WEST

BUSCON/93 WEST℠

Yes, I want more information about integrated system solutions at BUSCON and SOFTCON.

Name _____

Title _____

Company _____

Address _____

City _____ State ___ Zip _____

Phone _____ Fax _____

Return to: CMC, 200 Connecticut Avenue, Norwalk, CT 06856-4990
(800) 243-3238, fax (203) 857-4075.

BS5

A CMC EVENT

CIRCLE NO. 56

At the center of the TT150 Truck Driving Simulator is a mock-up of an actual truck cab—complete with seat, steering wheel, clutch, gear-shift lever, and all the gauges found in a typical truck. Three large graphic displays provide a 180° panoramic view of the road, along with inset images representing two rear-view mirrors. Large bass speakers inside the cab create the sound and vibration of a truck's diesel engine.

can't recreate on the road, and to evaluate how the driver reacts to those situations." The key to providing a realistic, fully interactive driving simulation was computer-generated imagery.

Getting good graphics

The biggest technical challenge for the FAAC engineers was finding and integrating a graphics subsystem that met their requirements. The TT150 required three forward channels for the three displays, with two channels for the rear-view mirrors. A realistic simulation required a system that could process 5,000 polygons in a scene at a minimum update rate of 15 Hz. At certain critical periods, the update speed had to reach 30 Hz.

Selecting a graphics system that met these requirements at an acceptable cost proved difficult. Several such systems were evaluated by the FAAC engineers, including multiboard products hosted in Sun systems and Silicon Graphics systems. "The ones that were mounted in workstations, to a great extent, didn't have true realtime capability," says Berrie, "so we started zeroing in on companies that concentrated on graphics and realtime simulation, as opposed to workstation-based products."

Eventually the team chose the Ball 994 graphics system made by Ball SED (San Diego, CA), an aerospace company. One reason FAAC liked the system was because of its scalability. "We were looking for a modular type of system that we could scale up or down based on the horsepower we thought was necessary," says Berrie. For FAAC's purposes, the system configuration consisted of three frame-buffer cards, one for each of the channels. Other boards included a pixel processor, span processor and, depending on the polygon load, from one to three display-list processors (DLPs).

The DLPs take a list of polygons and sort them. From the polygon data, the span processor computes which pieces of each polygon belong on each span line. The pixel processor receives span-line information and breaks it down into specific pixel data, such as color, luminance and chrominance. Pixel information which represents the image displayed on the screen is then stored in the frame buffers. Multiple sets of these boards may be used to achieve more horsepower.

Choosing the Ball system forced some trade-offs on the software side of the project. In the final system, the FAAC team had to create software that hadn't been planned. Many of the graphics systems FAAC considered had the capability to perform key simulation functions. The detection of a collision between the truck and another object in the database, for example, was handled directly by many graphics systems. Likewise, the height-above-terrain (HAT) function was taken care of by

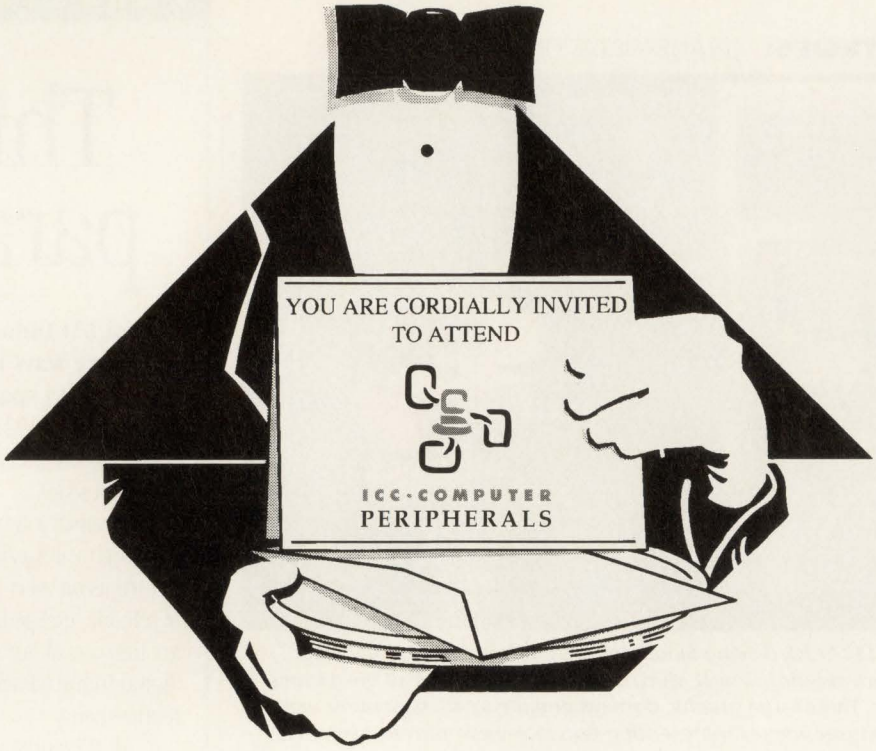

YOU ARE CORDIALLY INVITED TO THE. . .

Computer Peripherals Invitational Computer Conference

featuring new technology presentations and product displays

Request your ***Complimentary Invitation*** and Agenda

OEMs, Systems Integrators, VARs — If you are involved in the purchasing, design, or specification of peripheral products... you should be here!

Copy your business card here
FAX to: (714) 476-9969

icc
a division of
Dataquest®
a company of
The Dun & Bradstreet Corporation

Dataquest is a registered trademark of A.C. Nielsen Company

Indicate your conference selection below and return this section by FAX or mail.

❑	Newton, MA	Sep 09 '92
❑	Dallas, TX	Sep 22 '92
❑	Portland, OR	Oct 22 '92
❑	Irvine, CA	Jan 07 '93
❑	San Jose,CA	Feb 09 '93
❑	Minneapolis, MN	Mar 04 '93
❑	Austin, TX	Mar 25 '93
❑	Denver, CO	Apr 08 '93
❑	Nashua, NH	Apr 20 '93

3990 Westerly Place, Suite 100, Newport Beach, CA 92660 Tel: (714) 476-9117 Fax: (714) 476-9969

CIRCLE NO. 58

some graphics systems transparently. (The HAT function determines the location of the physical contact point between the tire and the polygon it's driving on.)

The Ball system, unlike the other graphics systems considered, is basically only a rendering box. The FAAC team, as a result, had to create additional software to define HAT and collision databases, and then had to store those databases on other hardware. The group also had to develop code to perform sorts and searches, as well as algorithms to provide the output.

VME for flexibility

Early in the design cycle, the FAAC team had to choose a main computer engine upon which to run the simulation. It decided on VME boards with Motorola microprocessors. "We didn't care for the memory addressing of the Intel chips," remarks senior engineer John Dibbs. "We preferred the linear address space that you get from the Motorola architecture. As for the choice of VME, that fitted our strategy of going for off-the-shelf hardware. We thought that VME would give us a wide range of third-party vendors to provide what we needed. Also, the whole development cycle supported the approach of keeping options open."

The designers settled on an architecture consisting of a pair of single-board computers, a 68040-based MVME167 board and a RISC 88000-based MVME181 board. The 68040 board, running UNIX, acted as the development environment and controlled peripheral I/O in the TT150 system. The RISC board, running the pSOS realtime operating system, provided the power needed to drive all the autonomous vehicle models, and to run the truck model itself. By using SBCs and the VME architecture, FAAC gave itself a path to easily upgrade to higher-performance boards as they become available.

RISC muscle critical

Among the reasons for including a RISC processor in the TT150 was the need to keep high-frequency models numerically stable. Models of the trailer hitch and suspension systems are both examples of high-frequency models. In a dynamic simulation, loops can occur in the simulation in the space of microseconds. During that interval, certain assumptions are made about which variables remain constant. Some weights, for example, are defined and integrated. In a dynamic system that is oscillating, problems can occur if the integration interval is too large. In such cases, the integration can result in a value that's unrealistic. This can create a numerical instability, causing the system to oscillate out of control.

"The best way to resolve this in a given model," says Berrie, "is to shrink your integration interval down to a small enough piece—small enough so that you can't extrapolate far into the future with bad information. We were concerned that we would have to run a vehicle model at a very high-frequency rate—30 Hz or higher—rather than the 15 Hz of the display system. We wanted to make sure we had sufficient horsepower and capability in a RISC realtime computer to handle that. That steered us toward the RISC 88000 processor and pSOS."

For its part, the 68040-based board performed as the overall system executive for the TT150. The 68040 handles the realtime paging function that pulls terrain database information off the disk and delivers it to the Ball system. Through a shared memory arrangement, the 88000 and the 68040 pass information back and forth.

Lessons learned

The TT150 Truck Driving Simulator design provided the FAAC design team with some insights into the importance of cost for commercial designs. "We were attempting to avoid new development, new products, new concepts," says Berrie. "We wanted to take off-the-shelf, currently available pieces of hardware and integrate them into a new application. Our market research told us that if we built a system that cost $300,000, the market just wasn't going to support it. We worked very hard to keep the price down, and we identified the components and computers that would help us achieve our target cost."

Because of the graphics system, the final cost of the TT150 was $230,000. "At first, the graphics system was the only component not available off-the-shelf," adds Berrie. "We had to settle on a price-performance point that we needed for our product. In the end, that made the TT150 commercially acceptable." ■

CALL FOR PAPERS:

APPLICATIONS MAKE THE DIFFERENCE AT ECC

THE ATTENDEES

ECC differs from other industry events because it focuses on providing solutions for a variety of embedded computer applications rather than focusing on technology alone or products alone. Developing these applications and finding the best solutions brings into play single-chip microprocessor and microcontroller implementations; custom and semicustom board designs; standard bus-based single-board computers and peripheral boards; operating systems, real-time kernels and compilers; development systems and debugging tools; embedded PCs, embedded workstations and even embedded microcomputers.

ECC attendees are intimately involved in the design, development and integration of a broad range of products and systems based on embedded computers.

ECC attendees are working in all major industries, where they are designing, developing and building a full spectrum of products and systems.

- In the computer industry, developing workstations and larger computers and making decisions about proprietary buses, open architecture buses or bus-less approaches.
- In process control and automation, developing systems for motor control, process measurement and control, machine vision, robotics and manufacturing automation, traffic control, etc.
- In communications, developing systems for use in cellular communications, PBXs, multiplexers, local area networks and wide area networks, etc.
- In the military/aerospace and avionics industries, building equipment and systems for communications, command and control (C3), weapons guidance and control, simulation, airborne and ground-based flight-control systems, etc.
- In the test, measurement and instrumentation industry, where equipment and systems are being built for product testing, maintenance and service applications, diagnostics (including medical), resource exploration, etc.
- Research and development, where scientists and engineers are designing equipment and systems for data acquisition, analysis and simulation that range from benchtop systems to space stations.

THE PROGRAM

ECC's technical program has been designed to provide attendees with the practical information they need to incorporate embedded computers in an end product or subsystem. These embedded computers can take the form of dedicated microcontrollers; sophisticated 32-bit CISC or RISC processors; off-the-shelf or customized SBCs; standalone SBCs; standard bus-based subsystems; or OEM workstation, desktop or industrial computer platforms.

While the Technical Program Committee will entertain proposals for presentations covering a broad range of embedded computer approaches, special consideration will be given to proposals dealing with the following major areas of concern:

SYSTEM ARCHITECTURES

Presentations in this category will deal with interprocessor and memory architectures, custom and semicustom system implementations, and loosely and tightly coupled software and hardware models. This section will include a discussion of chip sets for "standard-architecture" machines such as the X86-, SPARC- or MIPs-based workstations and the features that make them suited, or unsuited, to embedded applications. Some other areas of interest are:

- PCs (80X86, P5) in embedded applications
- RISC architectures (Alpha, HP-PA, MIPS, SPARC, 88K) in embedded applications
- Integrating processor, memory and I/O on stand alone SBCs
- Memory architectures with or without cache
- Symmetrical and asymmetrical multiprocessing
- Live-insertion, fault-tolerant, high-availability computing

INTERFACES AND STANDARDS

This track will deal with local on-board interfaces, addressing issues such as architectural considerations, transceiver considerations (cost, power, space, time-to-market) and implementations, current and emerging standards such as PCI and PCMCIA as well as more conventional mezzanine-I/O buses such as IndustryPacks, SBus, MX bus and others. Particular areas of concern include:

- Standard mezzanine/daughter boards
- SBus, TURBOchannel and other workstation I/O buses
- 80X86 (P5) peripheral interfaces
- Networking interfaces such as Ethernet, ATM and SONET
- Peripheral interfaces such as SCSI, HiPPI, and FiberChannel
- MCM standards and interprocessor module buses

ECC
EMBEDDED COMPUTER CONFERENCE & EXPOSITION

Embedded Computer Conference

April 14-16 1993 • Santa Clara Convention Center • Santa Clara, California

A FOCUS ON SOLUTIONS

INTERCONNECT ARCHITECTURES

The approaches to interconnecting board-level subsystems are undergoing a revolution, including the development of new standards such as PC/104, SCI (Scalable Computer Interface) and Futurebus+, as well as major changes and enhancements to such well-established standards such as VME and Multibus II.

In addition, there are many other approaches vying for the embedded market such as ESP (Extra Small Package), various STD approaches, G64 and others. Rapidly changing semiconductor technology is forcing changes in bus and interface technology which are reflected in:

- VMEbus, Multibus, Futurebus+, STD/STD32
- Custom and semicustom implementations
- High-performance backplanes
- Bridges and intercrate communications

SOFTWARE AND DEVELOPMENT TOOLS

Because embedded computer applications involve real-time processing, the major software focus of ECC will be real-time issues. Special emphasis will be given to multiprocessing in real-time using both traditional and Windows-based operating systems. Also of interest to attendees are:

- Real-time OS and kernels
- Real-time DOS
- Windows for embedded and real-time applications
- Multiprocessing with DOS and other real-time OSs
- High-performance optimized compilation
- POSIX and POSIX compatibility
- Communications protocols/standards
- Ada in military and nonmilitary applications

Proposals will be considered for presentations on any of the above topics, as well as on any other topics related to the design, programming and application of embedded computer products, subsystems or systems.

THE FORMAT

A combination of one-hour presentations, application-focused multi-paper sessions and tutorials will address a broad range of topics of importance to engineers and engineering managers designing both the hardware and software for embedded computers and subsystems.

- Tutorials are extended presentations intended to provide attendees with an in-depth understanding of core technologies used in embedded computers.
- Individual presentations are focused on various aspects of applying new or emerging technologies and products to solving specific problems in designing products or subsystems using embedded computers. Special consideration will be given to presentations that emphasize make-or-buy trade-offs in specific applications.
- Multipaper sessions consist of several shorter papers that focus on the implementation of embedded computing in specific application areas, including:
  - Medical instrumentation
  - Vehicular traffic control
  - Signal processing/data acquisition/DSP
  - Automated vehicles
  - Machine control
  - Graphics
  - Low-power and portable applications
  - Imaging
  - Visual inspection
  - Virtual reality
  - Military C3I
  - Communications
  - Peripheral interface and control
  - Process control
  - Laboratory automation
  - Multimedia
  - Networking

HOW YOU CAN PARTICIPATE

You may participate in the ECC Technical Program by submitting a proposal for either a one-hour lecture-type presentation, a 20-minute application-focused paper, or a longer tutorial. Please submit your proposal no later than November 20, 1992.

The proposal should be no longer than one page and consist of a short abstract that summarizes the content and goals of the presentation, and a brief outline of the major topics covered by the presentation. Presenters must be technically qualified and able to answer questions from attendees. A short biography of the presenter, detailing his or her technical background and accomplishments must accompany the proposal.

Acceptance of proposed presentations will be made by December 2, 1992. A complete copy of the presentation, including all visuals and graphics, for publication in the Conference Proceedings must be provided by March 1, 1993. Presentations given at ECC will be published in the Proceedings and copyright shall be assigned to Computer Design/PennWell Publishing Company.

COMPUTER DESIGN Technology and Design Directions

1993 INTERNATIONAL TECHNICAL CONFERENCE SERIES

COMPUTER DESIGN • One Technology Park Drive • P.O. Box 990 • Westford, MA 01886
Tel: 508-392-2124 • 800-223-4259 • Fax: 508-692-7780

Imagine Getting VME Performance . . .
. . . at Half the Size and Cost.

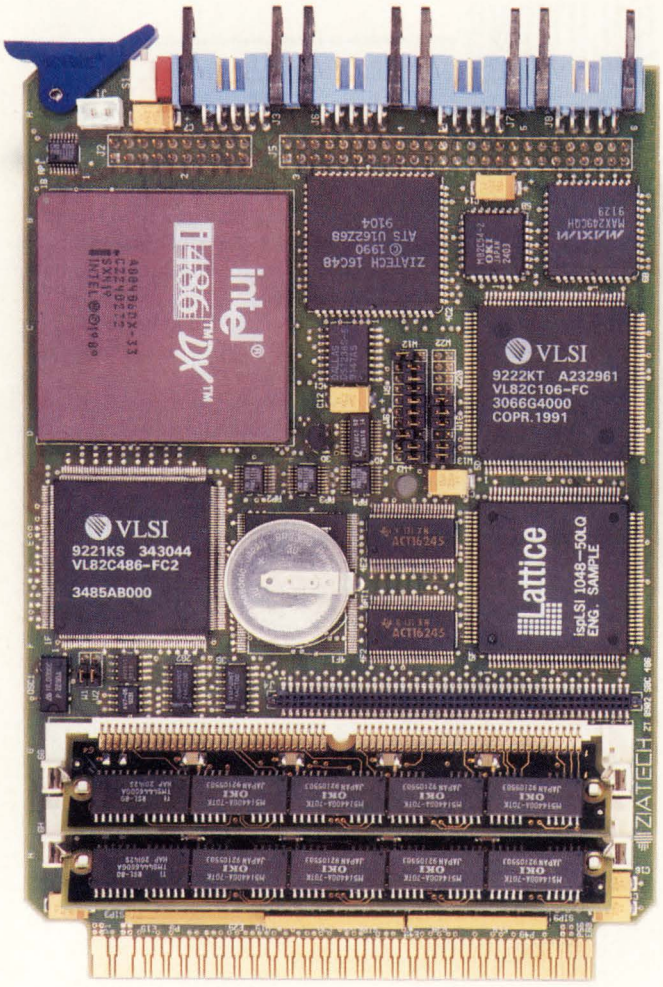

ZT 8902 Single Board 486 Computer

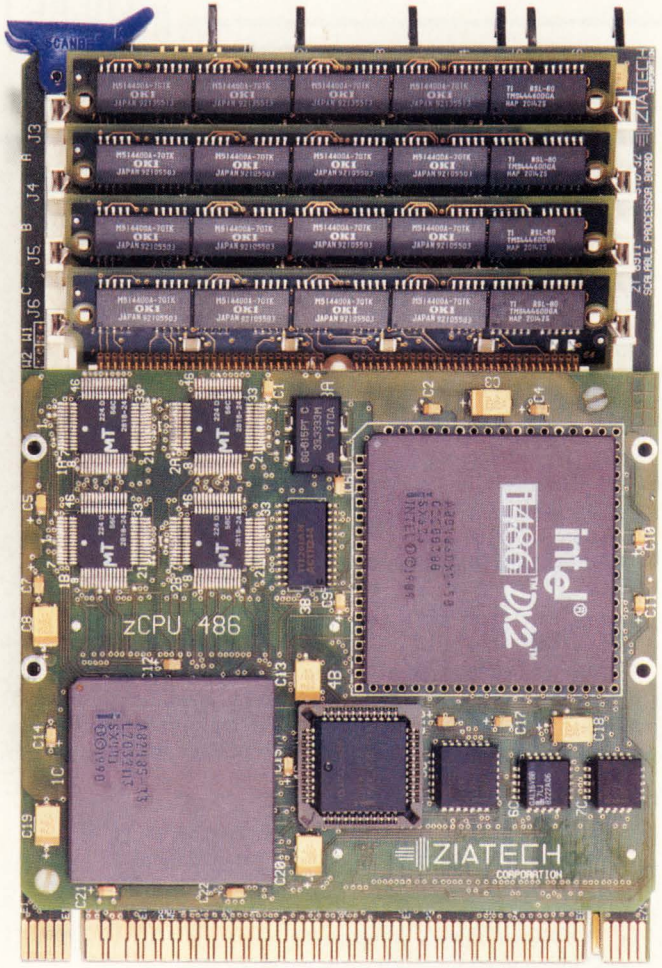

ZT 8911 Scalable Processor Board

Imagine No More.

Intel 486 Performance

Two new 486 industrial computers from Ziatech offer performance choices ranging from the 25 MHz 486SX to the 66 MHz 486DX2 in a very embeddable format.

PC Software, Multiprocessing Control

These high speed processors will run any PC operating system you need - MS-DOS, OS/2, QNX, UNIX, and more. And the STD 32 format provides something a PC or VME solution can't give you – multiprocessing capability that lets up to seven DOS-based processors share peripherals in a single system.

32-bit, Local Bus Video

Optional SuperVGA local bus video on the ZT 8902 SBC runs high resolution graphics at 486 CPU speed without using an extra I/O slot or any backplane bandwidth.

Small, Industrial Format

Ziatech's 486 computers provide all of the above and more in the compact 4.5- by 6.5-inch STD 32 format, the embedded computer that doesn't waste your money on board space and features you don't need.

Free Data Book

Call or Fax today for Ziatech's brand new STD 32 technical data book, which includes full specifications on our new 486 offerings and the rest of our industrial computer product line.

Phone: 805·541·0488
FAX: 805·541·5088

STD 32®
The Embedded ComputerSM

ZIATECH CORPORATION

CIRCLE NO. 61

COMPUTERS AND SUBSYSTEMS

STD Bus CPUs focus on solutions

Jeffrey Child, Senior Editor

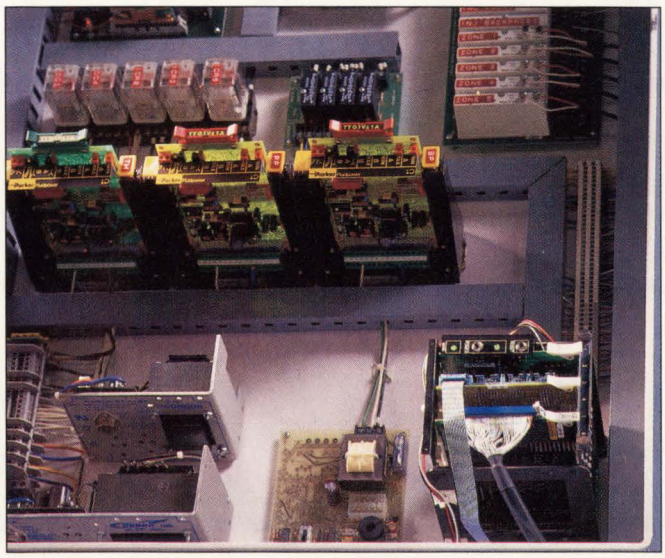

STD 32 CPU boards from Ziatech provided a solution for Hettinga Equipment when it needed a multistation injection molding machine. The system uses a PC as a master to control Ziatech's 8901 boards as slaves. These boards (lower right) perform the realtime control and data acquisition functions of the system.

Unlike VME, which has Motorola supporting it, or Multi-bus II, which is backed by Intel, STD Bus may suffer from not having a major semiconductor manufacturer behind it. But in reality, STD has a strong position as the low-cost workhorse bus for embedded control. STD Bus cards offer rugged hardware in a 4.5 × 6.5-in. form factor, as well as access to many off-the-shelf software development tools.

To keep the cost of their products low, STD CPU board makers use the inexpensive microprocessors and chip sets targeted for the PC market, where extreme price pressures and high volumes have had a favorable impact on the price and availability of components. Even more important, the exploding notebook PC market is driving the need for ever tighter integration of electronics, a trend that fits nicely with the requirements of a small-form-factor bus such as STD.

The newest STD CPU boards reflect aggressive efforts to leverage these hardware developments in the PC world. The current selection of STD Bus CPUs ranges from highly integrated, single-board computers to high-performance engines for handling the central computing tasks of a multiboard system. Among these are six new 486-based boards.

Trend toward solutions

If there's a trend in STD boards these days, it's toward a focus on solutions for the customer. This trend covers several areas, including support for multiprocessing, special I/O requirements, add-on modules, and future performance upgrade paths.

While many STD CPU board makers are pushing for increased performance, there's disagreement among these vendors over what future STD Bus performance should look like. Ziatech (San Obispo, CA) and a handful of other vendors are focused on STD 32, the 32-bit extension of the STD Bus standard. They argue that the bus should continue to act as a transfer medium, and so 32-bit access to memory or disk interfaces should occur over the bus.

Most of the STD Bus board makers take a different view, however. Their approach, and the approach supported by the STD Manufacturers Group, is to put all the computing power on the CPU card. Technology advances have let designers integrate a complete SBC, with ample memory and peripheral functions, onto a single STD Bus card. Such an approach relegates STD to acting essentially as an I/O bus channel. But, at speeds of 5 or 8 MHz, and using 8- or 16-bit specifications, STD is more than fast enough as an I/O bus.

Despite contrasting opinions over the need for a 32-bit STD Bus, some STD users take comfort in the fact that STD 32 provides an avenue for future performance upgrades and 32-bit multitasking—whether they need these capabilities or not.

Designers at Hettinga Equipment (Des Moines, IA), an OEM of injection molding machinery, used STD 32 CPU boards from Ziatech to control its latest multistation machine. While most molding machines can mount only one tool at a time, Hettinga's is unique in that it has up to six molding stations, all serviced by the same injection unit. This lets up to six different parts be molded on a single machine.

Hettinga's system uses an IBM PC as a master computer. This PC is connected to one or more Ziatech 8901 STD 32 CPU boards. These boards in turn run the realtime control and data acquisition functions of the system. Specifically, the 8901 works with STD 32 A-D interfaces and an 8-bit PAMUX card to provide standard I/O channels to the control process. A flat-panel display may be added using a Ziatech video interface.

Although Hettinga's machine design doesn't take advantage of the 32-bit features of the STD 32 8901 card, Richard Osborne, software project leader at Hettinga, plans to use the multitasking capabilities of STD 32 in the future. "We have some applications in preliminary design in which we have a multitasking type of control system," says Osborne. "We want those processors to share resources, such as the video card and a serial port. On the 8-bit STD Bus you could have two 386 cards, each with its own video, but you'd need to have two monitors. To me, that's one case where STD 32 makes a difference."

Model	CPU(s)	CPU clock speed (MHz)	Math coprocessor	RAM (bytes)	ROM (bytes)	DMA channels (no. and witdth)	Operating system support	I/O ports	Price	Comments
Advanced Micro Systems 2 Townsend West, Nashua, NH 03063 (603) 882-1447										**Circle 301**
BAS-52	8031/ 32/52	11.05	no	8-16k	8-16k	—	Intel BASIC	48 general purpose I/O	$329-$450	RS-232, printer port
BIB-52	8031/ 32/52	14.4	no	32k-128k	256k	—	modified Intel BASIC	48-144	$320-$495	RS-232; RS-485; optional 12-bit A-D
Antona 1643 1/2 Westwood Blvd, W Los Angeles, CA 90024 (310) 473-8995										**Circle 302**
ANC-7850	8032	11	no	32k	32k	1	no	2 8-bit parallel; 1 RS-232C serial; 1 serial printer	$195	Stand-alone or STD Bus
ANC-7852	8052 BASIC	11	yes	32k	32k	1	BASIC	2 8-bit parallel; 1 RS-232C serial; 1 serial printer	$285	Stand-alone or STD Bus; on-board PROM programmer; real-time clock
C-Matic Systems Ltd 2 Millbrook Business Park, Crowborough, E Sussex TN6 3J2 44-892-665688										**Circle 303**
STD 1804	Z80	4, 6	no	64k	64k	—	—	PIO/SIO	$290	
STD 1878	68030	12	no	16-32M	64k	—	—	—	$1,650	
Computer Dynamics 107 S Main St, Greer, SC 29650 (803) 877-8700										**Circle 304**
CPU-XT	8088/V20	5, 8	optional	640k	256k	2	DOS	2 RS-232; 1 parallel	$995	100 PC/XT compatible; directly drives flat-panel displays or CRTs
CPU-AT	80C286	12, 16, 20, 25	optional	4M	256k	2	DOS	2 RS-232; 1 parallel printer	$1,995	100 PC/AT compatible; directly drives flat-panel displays or CRTs
Cubit Division, Proteus Industries 340 Pioneer Way, Mountain View, CA 94041 (415) 962-8237										**Circle 305**
8650	80C186	16	no	256k	256k	1 8-bit	—	4 8-bit parallel, 2 RS-232	$595	C-Engine S/W, watchdog timer, 16-bit STD Bus
8660R	80C186	12.5	no	64kM	256k	1 8-bit	—	4 8-bit parallel, 2 RS-232	$775	12-bit A-D; –40° - +85°C; NAVMAT P-9492 vibration spec; C-Engine S/W; battery-backed RAM/clock
Datricon 31069 Genstar Rd, Hayward, CA 94544 (510) 471-9717										**Circle 306**
ACS-09	6809	1, 2	no	40k	40k	1 8-bit	—	2 RS-232, 2 RS-485	$215	
ACS-685BC	68008	8	no	1M	1M	1 8-bit	—	2 RS-232, 2 RS-485	$545	
DuraSys Box 814, Dover, NH 03820 (603) 742-7363										**Circle 307**
ESC 10809	6809	1, 2	no	48k	48k	—	ROM monitor	2 parallel, 1 serial	$295	4 timers; power-fail detect; interrupt generation; CMOS RAM backup
ESC 10812	6502	1, 2	no	CMOS to 48k	EPROM to 48k	—	ROM monitor	2 parallel, 1 serial	$295	Same as above
HiTech Equipment 9400 Activity Rd, Ste 1, San Diego, CA 92126 (619) 566-1892										**Circle 308**
STD-65F11 Forth card	R65F11	1, 2	no	64k	16k	—	—	3 RS-232, 2 parallel	$299-$349	

Model	CPU(s)	CPU clock speed (MHz)	Math coprocessor	RAM (bytes)	ROM (bytes)	DMA channels (no. and witdth)	Operating system support	I/O ports	Price	Comments
JF Microsystems 3641 Frontier Rd, Pasco, WA 99301-9619 (800) 532-2737										**Circle 309**
4188	8088	5	no	16k	32k	—	—	4 parallel	$400	I/O processor
8759	8088	5, 8	OPT 8087	—	32k	1 8-bit	—	—	$500	STD Bus master
Magnon Engineering 8739 Lion St, Cucumonga, CA 91730 (714) 466-0199										**Circle 310**
13000	68B09	8	no	—	—	—	Step motor controller 1-11 axis	1 IEEE, 2 RS-232	$400	For step motor applications
Matrix 1203 New Hope Rd, Raleigh, NC 27610 (800) 848-2330										**Circle 311**
MF9B	6809	2	no	64k	32k	—	OS-9	2 serial	$305	3 programmable counter/timers
SP9B	6809	2	no	80k	32k	—	OS-9	1 serial	$255	configuration register & software controlled memory mapping
Micro-Aide 685 Arrow Grand Cir, Covina, CA 91722 (818) 915-5502										**Circle 312**
80-0046	V40	8/10	no	1M	512k	1 8-bit	DOS	1 printer, 2 serial	$375	Battery-backed calander/clock and CMOS RAM; 5 32-pin sockets for memory; power-fail detect
80-0386	386SX	16, 20	no	8M	512k	4 8/16-bit	DOS	1 printer, 2 serial	$795	battery-backed calander/clock and CMOS RAM; iSBX expansion
Micro-Link Products Div. 401 Pennsylvania Pkwy, Ste 205, Indianapolis, IN 46280-1385 (800) 428-6155										**Circle 313**
STD224	68020	12, 20	6888	1-16M	1M	local	OS-9, PDOS	32-bit local	—	
STD247	Z80	3.6	no	64k	32k	—	—	1 RS-232	—	
Micro/sys 3447 Ocean View Blvd, Glendale, CA 91208 (818) 244-4600										**Circle 314**
SB8386	80386SX	16, 20, 25	80387SX	8M	1.8M	4 8-bit, 4 16-bit	RUN.EXE, DOS, Windows, OS/2	2 serial, 1 printer, iSBX	$845	Complete 386AT on card; VGA, IDE drive and STD80/MPX options
SB8486	CX486SLC	25	80387SX	8M	1.8M	4 8-bit, 4 16-bit	RUN.EXE, DOS, Windows, OS/2	2 serial, 1 printer, iSBX	$995	Complete 486AT w/1-kbyte cache on card; VGA, IDE drive, and options
Microcomputer Systems 1814 Ryder Dr, Baton Rouge, LA 70808 (504) 769-2154										**Circle 315**
MSI-C988	V20	5.0	no	64k	64k	—	—	2 serial, 3 digital, 1 analog	$395	Real-time clock; 3 16-bit timers; interrupt controller; battery-backup RAM; 10-bit analog input; 32 digital I/Os
MSI-CZ82	Z80A	4	no	8k	8k	—	—	2 serial	$285	Counter-timer clock
Miille Applied Research 1730 S Richey, Pasadena, TX 77502 (800) 729-0818										**Circle 316**
127-031-0	6809	1, 2	no	8-32k	8-32k	—	—	2 serial, 16 digital I/Os	$410	3 16-bit counter timers; watchdog timer; power-fail detect
127-031-1	6809	1, 2	no	8-32k	8-32k	—	—	2 serial, 16 digital I/Os, 16 analog IN	$595	

Model	CPU(s)	CPU clock speed (MHz)	Math coprocessor	RAM (bytes)	ROM (bytes)	DMA channels (no. and witdth)	Operating system support	I/O ports	Price	Comments
Mitchell Electronics P.O. Box 2626, 180B Mill St, Athens, OH 45701 (614) 594-8532										**Circle 317**
M/E 200	Z80	2.5, 4, 6	no	16/32k	48k EPROM	—	Forth	2 serial, 1 IEEE 488	$350	IEEE 488 talker/listener controller interrupt support for GPIB; portable battery-backed RAM; cal/clock option
CPU816	65C816	4, 8	no	8-128k	48k	—	Forth	2 serial, 2 8-bit parallel, 12 bit A-D	$450	8 channel A-D; flex memory mapping, 6502 opcode compatible, batt- backed RAM, cal/clock option
Mizar 2410 Luna Rd, Carrollton, TX 75006 (214) 277-4600										**Circle 318**
MZ77851	Z80	2.5, 4	no	1-2k	8k	—	—	3 8-bit OUT, 2 8-bit IN	$375-$475	
MZ77855	Z80	2.5, 4	no	256k	4k	—	—	—	$415-$435	4 counter/timers
Octagon Systems 6510 W 91st Ave, Westminster, CO 80030 (303) 430-1500										**Circle 319**
7100	64180	6	none	32k	8-32k	—	STD BASIC II	8 analog IN, 1 analog OUT, 38 digital I/Os	$445	
9600	V25	8	no	32k	32-128k	—	STD BASIC II	24 digital I/Os	$595	
Pro-Log 2555 Garden Rd, Monterey, CA 93940 (800) 538-9570										**Circle 320**
7874	486SX/DX	25, 33	in CPU (DX)	16M	512k	2aw	DOS, Windows, OS/2, QNX	2 serial, 1 parallel	$2,495	Multiprocessor support; CPUs per system; disk interface; VGA
7873	386SX	20, 25	387	8M	512k	2	DOS, Windows, OS/2, QNX	2 serial, 1 parallel	$1,895	Same as above
Quasitronics 211 Vandale Dr, Houston, PA 15342 (412) 745-2663										**Circle 321**
CPU-Z80	Z80	2.5	no	8k	8k	—	CP/M	—	$196	Single +5-V power required; DMA control
R.L.C. Enterprises 4800 Templeton Rd, Atascadero, CA 93422 (805) 466-9717										**Circle 322**
TSBC-C186	80C186	10, 12.5, 16	80C187	256k	256k	2	—	2 RS-232	$675 (10MHz, no co-processor)	16-bit stand-alone or STD operation; 6 16-bit counter timers; flash memory
SBIO-186	80186/ 80C186	8, 10, 12.5, 16	no	256k	128k	2	—	2 RS-232/ RS-422, 32 parallel I/Os	$555 (8 MHz)	Real-time clock; Opto-22 interface; RS-422 drivers on RS-232 ports
Robotrol 925 W San Martin Ave, San Martin, CA 95046 (408) 683-2000										**Circle 323**
RSD7832	UPI42 (8051)	12	no	32k	64k	—	—	1 8-bit parallel, 1 serial	$595	Intelligent analog I/O; 16 analog IN; 2 analog OUT
RSD7808	Z80	4	no	64k	64k	—	—	20 parallel, 2 serial	$245	Master/slave; 16-input vectored interrupt

Model	CPU(s)	CPU clock speed (MHz)	Math coprocessor	RAM (bytes)	ROM (bytes)	DMA channels (no. and witdth)	Operating system support	I/O ports	Price	Comments
Systek 415 N Quay St, Ste 6, Kennewick, WA 99336 (509) 735-1200										**Circle 324**
8825	V25	5, 8, 10	no	512k	512k	2 8-bit	MS-DOS	8 analog IN, 2 analog OUT, 22 digital, 2 serial RS-232, iSBX	$395	Master/slave for multiprocessing; watchdog timer; clock/calendar
8850	V50	8, 10, 12	8087	512k	256k	3 16-bit	MS-DOS	2 RS-232, 1 RS-485, 1 Centronics	$575	Multi-master; watchdog timer; clock/calendar
Versalogic 3888 Stewart Rd, Eugene, OR 97402 (800) 824-3163										**Circle 325**
VL-186-1	80C186	16	no	1M	1M	2 16-bit	DOS	12 parallel, 1 RS-232, 1 RS232/422	$645	STD32; DOS-compatible; hard-ware; multiasking; RT embedded DOS (trade-mark)
VL-188	80C188	5, 8, 10	no	512k	512k	2 16-bit	—	16 lines parallel, 1 RS-232/422	$325	Opto-22 compatible parallel port
WinSystems 715 Stadium Dr, Arlington, TX 76011 (817) 274-7553										**Circle 326**
MCM-SX386	80386SX	16, 25, 33	80387SX	4M	512k	4 16-bit, 4 8-bit	DOS, QNX, OS/2	2 RS-232/ 422/485, 1 parallel	$850	PC/104 expansion; floppy cntrl; keyboard cntrl; IDE interface; optional –40° to +85° C operation
MCM-486DX	80486DX	33, 50, 66	in CPU	8M	128k	4 16-bit, 4 8-bit	DOS, QNX, OS/2	2 RS-232, printer	$1,995	iSBX expansion keyboard controller
MCM-486SLC	80486SLC	25, 33	80387SX	4M	512k	4 16-bit, 4 8-bit	DOS, QNX, OS/2	2 RS-232/422/ 485, 1 printer	$995	PC/104 expansion; floppy cntrl; keyboard cntrl; IDE interface
XYZ Electronics 4700 N 600 W, McCordsville, IN 46055 (317) 335-2128										**Circle 326**
CPU68k16	68000	16	no	3M	3M	—	OS-9/68000	2 RS-232/RS422, 1 SCSI-compatible, 1 parallel	$495 (no RAM/ ROM)	Includes real-time clock & battery. All RAM can be battery backed.
SB-68k16	68000	16	no	750k	1M	—	OS-9/68000 in ROM	2 RS-232/422, 1 SCSI-compatible, 1 parallel	$995	Same as CPU68k16 plus the ROMed operating system; non-volatile RAMDISk and ROMDISk; C compiler; screen editor
Ziatech 3433 Roberto Ct, San Luis Obispo, CA 93401 (805) 541-0488										**Circle 327**
ZT8902	486SX/ DX/DX2	25, 33, 50, 66	in CPU	4-8M	1-2M	2	MS-DOS, OS/2, UNIX, QNX, Windows	2 serial, 3 parallel, 1 Centronics	$1,850	STD32; optional local bus video module
ZT8911	486DX/ DX2	33, 66	in CPU	4-32M	2M	6	MS-DOS, OS/2, UNIX, QNX, Windows	2 serial, 3 parallel, 1 Centronics	$3,350	Norton index of 103.5; replaceable CPU module; full 32-bit backplane transfers; STD32 multiprocessing
Zwick Systems 17 Fitzgerald Rd, Ste 104, Nepean, Ontario, k2H 9G1, 613-726-1377										**Circle 328**
ZSTD-100	Z84C13 (Z80)	4, 5, 8, 10	no	512k	512k	—	—	2 RS-232C	$529	Dual MMU; counter-timer clock; watchdog timer; wait-state generator; debug monitor
ZSTD-110	Z84C13 (Z80)	4, 5, 8, 10	no	512k	512k	—	—	16 parallel I/Os, 2 RS-232C	$589	Same as above

COMPUTERS AND SUBSYSTEMS

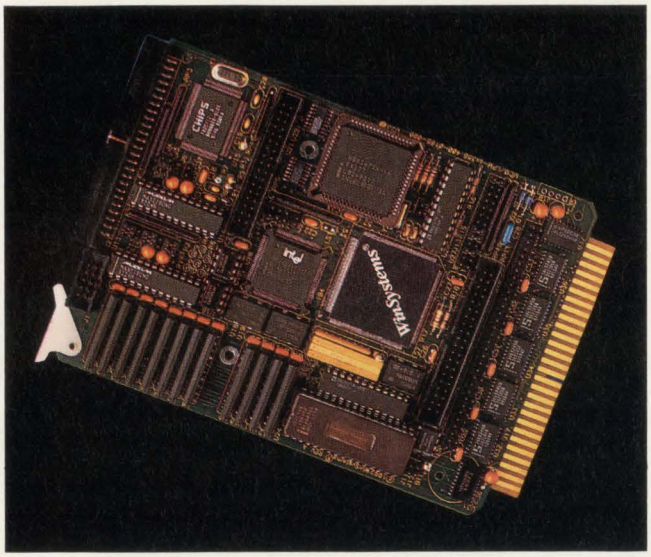

The MCM-SX386 is the first STD CPU board to use PC/104 as a mezzanine bus. Offering a choice of expansion solutions, the board provides a common computer core to which you can add off-the-shelf or custom-designed PC/104 expansion modules. Based on a 33-MHz 386SX microprocessor, the MCM-SX386 has up to 4 Mbytes of on-board RAM.

Ziatech's latest single-board STD 32 computer is the ZT 8902, a 486-based board with high-speed local bus video. The 8902 is available with a choice of Intel 486 processor options. These range from the 25-MHz 486SX to the 66-MHz 486DX2. Local bus 32-bit Super VGA capability is offered on a plug-in module. According to spokespeople for Ziatech, this video module lets the ZT 8902 run high-resolution graphics faster than most high-end PCs. The module operates at 25 or 33 MHz, increasing video performance up to 70 percent over 16-bit boards that operate at the ISA bus speed of 8 MHz.

Together, the ZT 8902 and the local bus video module occupy a single slot in an STD 32 card cage. By offloading the video function, the module helps prevent bottlenecks on the STD 32 bus. The board also features 1 Mbyte of flash memory, up to 8 Mbytes of DRAM, two serial ports, a printer port, 24 channels of industrial digital I/O, and extra counter/timers.

Beat 'em or join 'em

As STD Bus CPU board vendors move to a solutions-oriented strategy, they're finding that some embedded control solutions compete for the same types of applications. In response, some vendors have quite

sensibly offered avenues of compatibility from their STD Bus boards to competing embedded computer technologies.

Exemplifying this trend, Pro-Log (Monterey, CA) has, over the last year, formed a number of strategic relationships with makers of programmable logic controllers (PLCs). These partners include such firms as Allen-Bradley, GE and Modicon. The agreements let Pro-Log build interface cards to the various PLC companies' architectures. In the case of Allen-Bradley, for example, Pro-Log's STD boards can have direct access to Allen-Bradley's communications channel schemes: Data-Highway and Data-Highway Plus. In addition, Pro-Log cards can access remote I/O schemes such as Allen-Bradley's 1771 I/O subsystem.

"We're not going to replace PLCs in applications for pure I/O," says Paul Virgo, director of marketing at Pro-Log. "There are applications where system designers want to tie their PLC I/O control into systems the PLC doesn't support, such as disk management or alternative networking choices."

A system may require high-resolution graphics, for example, as well as some kind of realtime control. A PLC ladder processor can do some level of operator interface by sending status over the serial link to a terminal, but there's a time lag between the realtime I/O control action and when it appears on the screen. "Offering the ability to connect to their [PLC's] I/O scheme and those PLCs is a way of making our multiprocessing capabilities as useful as possible in the controls market," says Virgo. "In multiprocessing you're not overlaying one environment with another; you're keeping the DOS totally separate from the realtime control environment. As a result, there's no risk of crashing your program because something unexpected happened in the realtime control in the middle of doing a disk access."

Multiprocessing from Pro-Log

Pro-Log's Multimaster multiprocessing scheme lets up to seven 286, 386 or 486 CPUs share both data and I/O over the backplane. The company's most powerful board with Multimaster support is the 7874, a 486-based STD Bus computer. Available with a 25- or 33-MHz 486, the 7874 offers full 32-bit data transfers to the 4 or 16 Mbytes of on-board RAM. The CPU board may be used with any STD I/O card that conforms to the STD-80 Series specification for adding on application-specific I/O functions.

Microsoft's MS-DOS 5.0 operating system resides on a flash memory solid-state disk drive, permitting the system to operate in shock, vibration and temperature conditions that are unsuited to rotating disk drives. A true SBC computer, the 7874 includes

COMPUTERS AND SUBSYSTEMS

standard AT peripheral devices, including keyboard interface, two RS-232 serial ports, a parallel port, interrupt controllers, counter/timer, DMA controllers, floppy and IDE hard-disk controllers, and a battery-backed, realtime clock.

At last fall's Buscon, PC/104 standard modules generated some excitement. A spokesman for Motorola, for example, remarked that PC/104 represents a potential threat to VME over the long term. Clearly, the 3.85 × 3.6-in. PC/104 form factor offers direct competition to STD. Reflecting yet another example of the trend toward solutions, one STD maker, WinSystems (Arlington, TX), has decided to make use of PC/104 as a mezzanine bus. "There will be some competition between PC/104 and STD," remarks Bob Burckle, vice-president of WinSystems.

Targeted for low-cost embedded system applications, the WinSystems MCM-SX386 offers both stand-alone operation and expansion options with either the STD Bus or PC/104 bus. The MCM-SX386 provides a common computer core from which customers can add off-the-shelf or user-designed IEEE 961 STD Bus or PC/104 expansion modules.

The two expansion methods let you conform to open-standard architectures without having to pay for expensive proprietary solutions. At the heart of the board is a 33-MHz 386SX. Up to 4 Mbytes of parity-checked DRAM can be installed on the board. An on-board EPROM socket supports the BIOS and extensions, plus a bootable ROM disk of up to 440 kbytes. All the basic AT peripheral functions are included on-board. ■

CAE/CAD DESIGN TOOLS

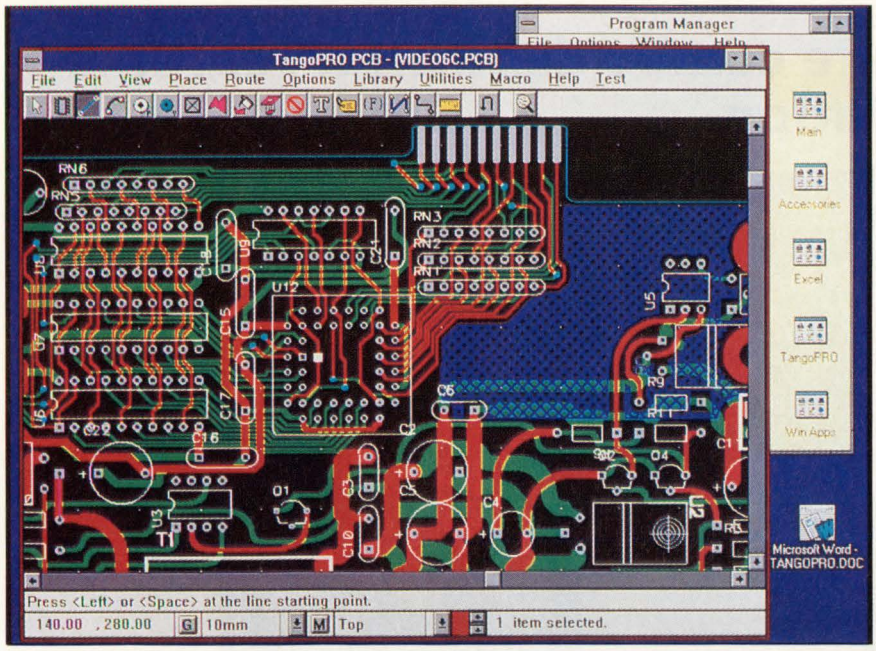

The 32-bit architecture of the TangoPRO PCB provides features such as copper pour (blue area in upper right) and routing of curved traces. The submicron database supports routing of fine-pitch traces and lets you control component rotation down to 0.1°.

Windows-based PCB tool suite boasts workstation features

As personal computers and the software that drives them become more powerful, PC-based EDA tools are closing the gap that separates them from their sophisticated workstation-based counterparts. The latest example of this trend, TangoPRO PCB from Accel Technologies, is a suite of printed circuit board (PCB) design tools that uses a 32-bit architecture running under Microsoft Windows 3.1 to provide workstation-caliber features on the PC.

The TangoPRO suite consists of a PCB layout tool, an optional high-speed autorouter and a library manager, which provides library integration of schematic and PCB component data. The system complies with the Windows memory management and user interface standards to ensure that the tool looks and feels like other popular Windows applications. The submicron database resolution permits full imperial and metric support down to 0.0001-in. or 0.01-mm grids and lets you control the rotation of components down to 0.1°.

TangoPRO's object-oriented database and its support of up to 16 Mbytes of extended memory effectively eliminate component limits on your designs. The tool also provides dialog boxes and menus to assist you in implementing difficult design concepts such as pad stacks and blind/buried vias. Prompt and status lines give online information, while the graphical icon-based toolbar provides access to often-used commands. Context-sensitive help gives you nearly the entire contents of the tutorial and reference manuals online at any time.

Attributes defined

The tool suite is flexible enough to let you define every attribute on the board, including grids, tracks, pads, polygon fills, text, and layers. Up to 99 layers are supported, 88 of which may be defined by the user without limit to the number of copper planes. Polygons may be placed on signal layers and poured with copper in solid or hatched fills. Traces may then be plowed out of the fills. In addition to 45° and 90° orthogonal traces, curved traces may also be manually routed on the board for high-speed digital and analog designs.

Available as an option to TangoPRO PCB is TangoPRO Route, a 32-bit autorouter with a rip-up and reconstruct algorithm letting it iterate to 100-percent completion of complex PCBs. Completely integrated within TangoPRO PCB, the router's options are set from menus within the PCB tool, where they may be executed, paused or halted. The autorouter supports the full capability and capacity of TangoPRO PCB, including display options, zoom and pan control, blind and buried vias, and number of components, nets and pads.

TangoPRO requires a system configuration with MS-DOS 5.0 or higher, Windows 3.1 or higher, a 386 or 486 IBM PC or compatible, 8 Mbytes of RAM (16 Mbytes with the autorouter) and 10 Mbytes of hard disk space. A mouse is also recommended. TangoPRO PCB is priced at $5,950, and the autorouter at $10,950. Both will begin shipping in January of 1993. —*Mike Donlin*

TangoPRO PCB at a glance

- PCB layout tool suite with library manager and optional autorouter capabilities, all based on 32-bit architecture and Windows 3.1
- Submicron database provides resolution down to 0.0001-in. or 0.01-mm grids and allows user-definable component rotation down to 0.1°
- Up to 99 layers 0.1° supported, 88 of which may be user-defined without limit to the number of copper planes
- TangoPRO Route is a 32-bit autorouter that supports uniform and non-uniform grids and simultaneous routing on up to 30 layers

Accel Technologies
6825 Flanders Dr
San Diego, CA 92121-2986
(619) 554-1000
Circle 287

SOFTWARE & DEVELOPMENT TOOLS

Hardware/software combo debuts for realtime design and simulation

An integrated development environment for realtime control systems is the result of a cooperative effort between The Math Works and dSPACE GmbH. The Math Works is supplying its Simulink system-simulation and code-generation software, which will be integrated with dSPACE's digital signal processor (DSP) hardware and software. The combination of simulation software with fast DSP hardware provides a platform for dynamic realtime testing and analysis in an environment that can automatically generate code when the design phase is complete.

Simulink is based on Matlab, The Math Works' interactive environment for numeric computation.

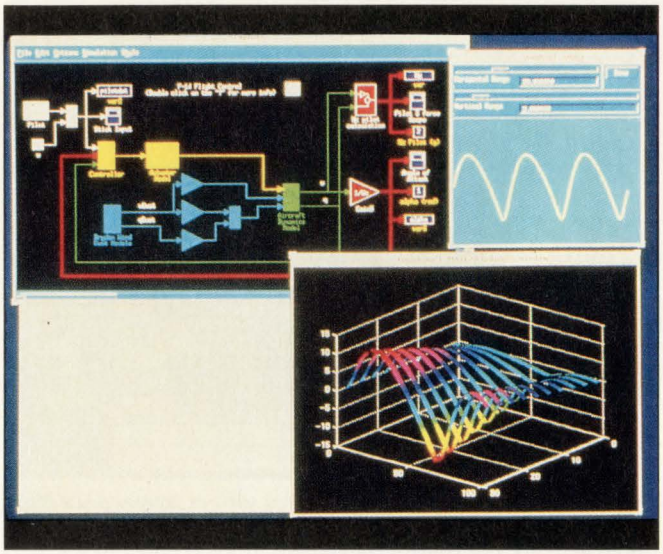

Matlab/Simulink from The Math Works can, running on a host computer, generate C code that compiles and runs on DSP hardware from dSPACE. This lets you simulate a realtime system on your host and observe its actual behavior on the target board by feeding back data to the simulation model.

Matlab contains algorithms and functions for digital signal processing, control system design, neural networks, dynamic system simulation, and optimization, among other functions. Simulink offers a graphical user interface (GUI) that lets you build block diagrams and hierarchical models. Connected blocks can be defined as a superblock, and can then be combined with other blocks at the next higher level.

In addition to functional blocks, such as integrators and filters, Simulink provides sources and sinks—blocks that can generate signals to flow through the model and blocks that receive signals from the model for use in analysis of the simulation. Simulink uses industry-standard GUIs, including the OSF/Motif flavor of X-Windows, Macintosh and Microsoft Windows.

As originally designed, Simulink stops at the simulation stage. But now The Math Works supplies a C code generator that can produce application code for a target system once the design has advanced sufficiently. Using the Simulink C code generator, you can output C source code directly from Simulink block diagrams. When the underlying Matlab function and S-function blocks are contained within a block diagram model, you can fill in the corresponding functions with your own code. The C source can then be compiled for the target hardware environment, and the data compiled from running on the target can be fed back to the Simulink model to observe the behavior of the real system.

Partnership enhances product

The partnership with dSPACE gives Simulink users access to a very high-speed target environment based on Texas Instruments' DSPs. dSPACE hardware environments range from plug-in AT bus boards to full-sized systems to specialized target processor systems such as Autobox. Autobox is used for testing designs in automobiles under real-world conditions, such as on the test track. Changes made to parameters in the Simulink model are automatically fed to the compiled code in the hardware, so that the behavior of the system can be observed in real-time.

The Math Works and dSPACE have collaborated to provide interfaces between the dSPACE hardware and Simulink. In addition to processor boards, dSPACE provides DSP-based I/O hardware, with specialized I/O software for real-world devices such as optical encoders and other sensors and actuators.

You can put together what amounts to a personal simulation system for about $30,000. Matlab and Simulink, along with the C code generator, are available from The Math Works for about $22,000. Ap-

Continued on page 126

SOFTWARE & DEVELOPMENT TOOLS

Continued from page 125

propriate DSP-based simulation hardware, along with the interfaces to The Math Works environment, are available from about $8,000, depending on configuration and options, from the U.S. distributors of dSPACE products.

— Tom Williams

Simulink/dSPACE at a glance

- Supports digital signal processing, neural networks, dynamic simulation, etc.
- GUI for block and hierarchical block diagrams
- Runs under OSF/Motif, Macintosh and MS-Windows
- Supports C code generation from block diagrams
- Supports dSPACE high-speed DSP processor and I/O boards
- Interfaces allow interactive debugging and code modification in realtime

The Math Works
24 Prime Park Way
Natick, MA 01760
(508) 653-1415
Circle 288

dSPACE GmbH
An der Schoenen Aussicht 2
4790 Paderborn, Germany
011-49-05251-1638-0
Circle 289

Optimizing C compiler tightly coupled to realtime OS

An ANSI C compiler with both global and interprocedural optimizations for realtime applications is also the first C compiler to be produced by a realtime operating system vendor. The Ultra C compiler from Microware is tightly coupled with the company's OS-9 and OS-9000 realtime operating systems, incorporating as it does a library of realtime system calls for those OSs. Ultra C supports language and target processor independence with an intermediate code (I-Code) architecture through most of the compilation process.

Currently, Ultra C supports all common versions of ANSI C, and will soon support C++, with other languages to follow in the future. The compiler also supports the full Intel 80X86 and Motorola 68XXX lines of processors and it will soon support RISC and 64-bit processors. Between the compiler's language-specific front end and its processor-specific code generator are linking and optimizing operations that are performed on the I-Code. Then there is a fourth assembly level "clean-up" stage of operation prior to final linking and object code generation.

The I-Code optimizer and linker is called Ilink and it produces an intermediate code representation of a program rather than an object code representation. The next stage is called Iopt, which is the intermediate code optimizer. Since the Ilink provides an I-Code representation of an entire program, the optimizer can oversee the whole program, optimizing it on several possible levels. It can perform local optimization within straight-line code, or it can perform global optimization across straight-line code, but within a function. Thanks to the I-Code representation, Ilink can also reach across all functions and data to optimize the program as a whole.

Source files can be linked in two ways. They can be compiled into relocatable object files and then linked, or I-Code files can be linked with other I-Code files and then optimized. The former method lets you optimize within source files, but not across them, while the latter lets you optimize across the whole program. The latter option also increases the compile time. A new feature with Ultra C is that C library modules and operating kernel system call modules are supplied in I-Code as well as in object code. This gives you the option of linking these library modules at the I-Code level and performing optimization across not only the application code, but the library code as well.

The back end of the compiler translates the I-Code into target-specific assembly language. It then lays out data areas, assigns registers and generates code to access global data. There's an assembly-level stage of optimization that cleans up some of the code by doing such things as merging sections of duplicated code. Then the tool performs the final step of generating machine language.

Ultra C complies with ANSI X3J11 1989 and ISO/IEC9899:1990 specifications. Libraries include system calls for OS-9 and OS-9000 kernels. The language processor can accept a variety of different source programs, including Microware's previous compilers, strict ANSI code and ANSI-extendible code.

Ultra C is priced at $1,250 and is shipping now. *— Tom Williams*

Ultra C at a glance

- Language-independent front end
- Processor-independent back end
- Implements optimization in intermediate code (I-Code)
- Optimization at the local, global and interprocedural levels
- C libraries and OS-9 and OS-9000 system call libraries supplied in I-Code
- Can link code as I-Code or as object code modules

Microware
1900 N.W. 114th St
Des Moines, IA 50325
(800) 475-9000
Circle 292

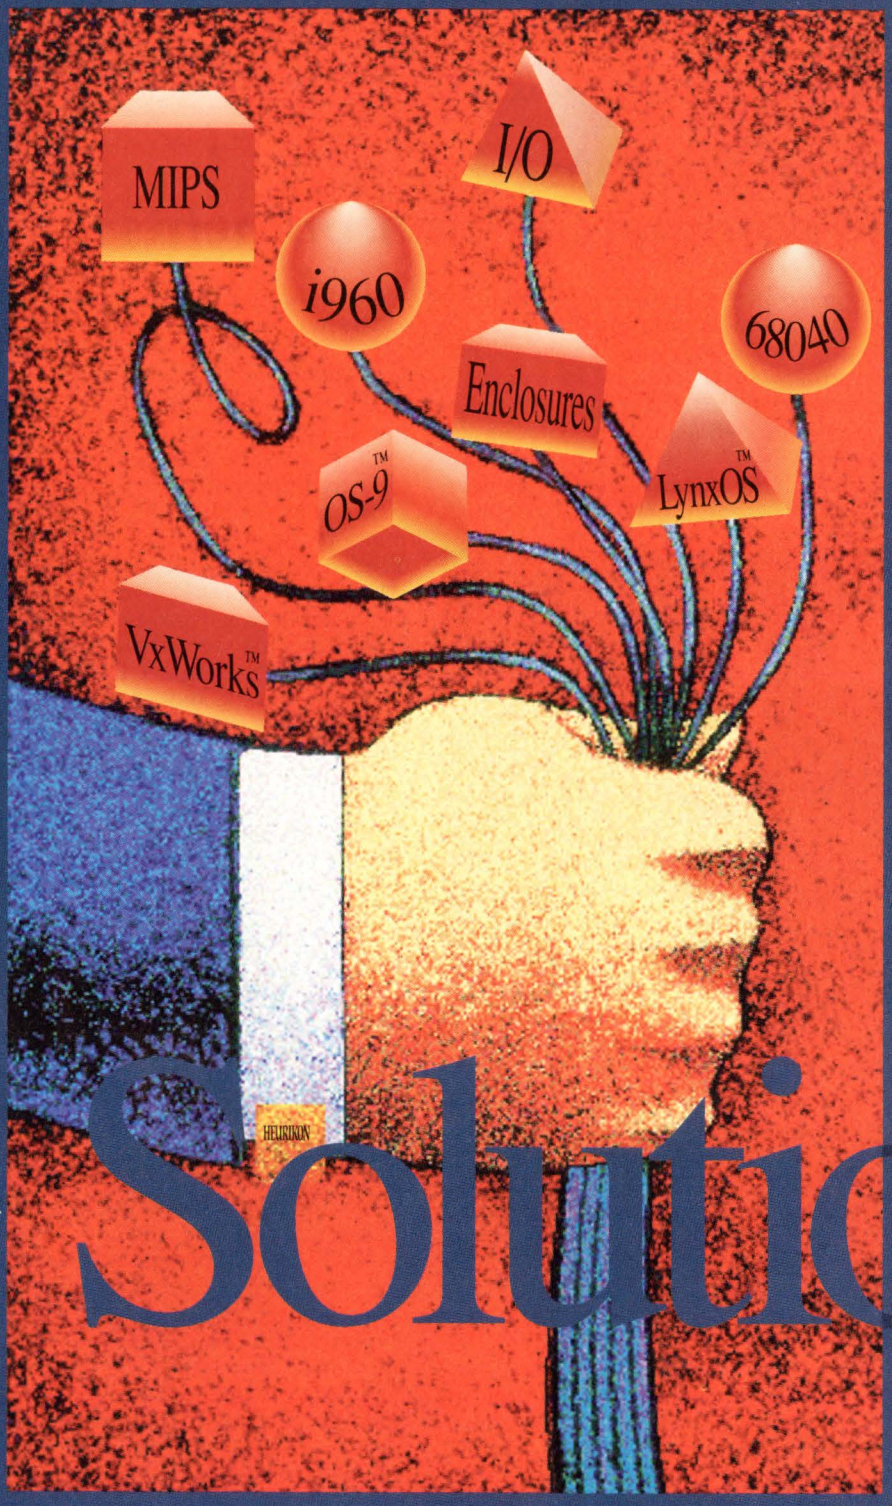

With Heurikon Integrated VME Systems...

Get A Grip On Real-time Solutions

Your success depends on having a firm grasp on today's performance choices. Technology, however, is only the beginning. Effectively integrating those choices demands performance from hardware and software alike. You should have the flexibility to select from standard or custom configurations.

Heurikon provides you with total solutions. Whether you use VME, Multibus or a proprietary platform we deliver you the real-time performance for your application. With Heurikon you have more than a grasp on the problem, you've got a handle on the solution!

HEURIKON CORP.®
A COMPUTER PRODUCTS COMPANY

Heurikon Corporation
8310 Excelsior Drive
Madison, WI 53717-1938
800 356 9602

Heurikon Europe
+44 (494) 88 32 32

CIRCLE NO. 65

COMPUTERS & SUBSYSTEMS

ACCESS.bus hardware released for industrial and commercial environments

Introduced over a year ago, ACCESS.bus has been relatively slow in gathering momentum. The basic idea was to bring some order to the interconnection of a broad range of accessory devices, such as keyboards, locators and bar-code readers.

Digital Equipment Corp and a handful of other OEMs have been implementing ACCESS.bus, but only recently has it started to catch on. Earlier this year, for example, an ACCESS.bus industry group was formed by 22 manufacturers. In addition, a company has been founded to exploit the benefits of ACCESS.bus. Computer Access Technology Corp (CATC—Sunnyvale, CA) has begun introducing products that let up to 125 peripheral devices connect to a single communications port in a PC/AT system.

ACCESS.bus is designed to handle relatively slow computer input from accessory devices such as keyboards, mice, bar-code readers, magnetic-card readers, modems, and some signal transducers for real-time control applications. Initially it was intended to handle up to 14 different devices on a single serial cable, which could be as long as 8 meters. The early specification, however, left some leeway for implementation options, and CATC's approach handles up to 125 devices, while extending the 25 ft to 250 ft with an ACCESS.bus adapter.

Computer Access Technology's ACCESS.bus standard AT board provides for the connection of up to 125 peripheral devices to a single communications port in a PC/AT computer. An ACCESS.bus development support kit is also available. The controller board is co-marketed with Philips.

The board uses the basic Signetics I[2]C configuration and interface, with an integral Signetics 8XC654 microcontroller; it can handle data rates of up to 80 kbits/s (100 kbits/s minus overhead). In addition, the ACCESS.bus cable carries a +12-V supply voltage for powering each device. The cable carries up to 500 mA.

CATC's board is supported by an extensive software package, including on-board microcode to control physical ACCESS.bus devices and an ACCESS.bus manager program that runs as a TSR under the PC/AT DOS operating system, routing control and applications messages between the physical devices and their software drivers. A Windows 3.1 version of the manager program is optional.

Commercial and industrial

The primary offering of CATC is a standard 16-bit AT/ISA half-board (4.2 × 6.5-in.) design that fits in a single slot and provides two ACCESS.bus connectors. An 8k × 8-bit SRAM buffer memory is included on the board. The unit controls a standard ACCESS.bus network, which lets you add peripherals in several ways.

CATC has also recently added to its product line a compact PC/104 controller to handle industrial applications. The unit has a smaller form factor (3.6 × 3.8-in.) but is functionally the same as the AT/ISA version. Says CATC president Dan Wilnai, "The PC/104 ACCESS.bus module exploits two new open standards ideally suited to embedded control applications. The ultra-compact, stackable module architecture of the PC/104 standard makes it easy to design the full capability of a PC into all kinds of equipment and instrumentation, while the ACCESS.bus serial bus communication standard based on the I[2]C physical layer protocol lets the designer connect multiple sensors or actuator devices to a single I/O port."

At the electrical level, ACCESS.bus functions as Philips initially defined its I[2]C setup. The host and devices are connected to both the data and clock lines in a "wired-AND" logical configuration. The wired-AND is implemented by connecting the data and clock output stages of each bus node to the lines through open-collector or open-drain transistors. These devices are included on existing I[2]C components. The wired-AND configuration lets any of the bus nodes force either line low. When there's no output from any bus node, the lines are held high with pull-up current sources in the host. All devices sense the level on both the clock and data lines.

Because of its relatively straightforward implementation and arbitration based on the same wired-AND configuration, the bus is basically immune to disruptions from the live insertion of additional peripherals.

While this is of some value in a commercial or desktop application, it's very important in industrial applications, where it could greatly simplify servicing, monitoring, updating, and downloading information. To avoid rebooting a system in such situations could result in saving appreciable downtime.

CATC offers a development support package that contains a controller board, an ACCESS.bus mouse, expansion box, cables, and an 87C751 microcontroller. The kit also includes a comprehensive software package and user's manual. The software package consists of on-board microcode and the ACCESS.bus manager, as well as a sophisticated ACCESS.bus monitor and control program. Source code for generic device driver interfaces to the PC and for ACCESS.bus device software modules is also included.

CATC claims that the low price of the controller boards—$86 for the AT version in lots of 1,000—will go a long way to assure the rapid acceleration of ACCESS.bus technology. Available now, the development support package sells for $1,500 and the PC/104 modules for $245.

— Warren Andrews

ACCESS.bus at a glance

- ACCESS.bus controller board
- Up to 125 peripherals
- Operates at up to 25 ft (250 ft with extender)
- Connects multiple interactive I/O devices
- Uses low-cost I[2]C microcontroller technology
- Compatible with open industry standard
- Connects keyboards, mice, trackballs, digitizers, scanners
- Available as AT/ISA or PC/104
- Development kit available

Computer Access Technology
949 Hillsboro Ave
Sunnyvale, CA 94087
(408) 732-8910
Circle 290

COMPUTERS & SUBSYSTEMS

Intel beefs up Multibus line using 486 processors

While the board industry speculates on Futurebus+ and enhancements to Multibus and VME that haven't yet achieved commercial viability, Intel is moving ahead with new families of Multibus I and II products. The products—designed to provide existing Multibus users with higher performance levels and new users with the flexibility of Multibus's fast processing capability—take advantage of Intel's latest family of 80486 processors.

The product introductions comprise three Multibus II single-board computers (SBCs) and a high-performance 486-based Multibus I card. All three SBCs are PC-compatible boards; the most complete is the Embedded Workstation. In addition, the company has introduced a new customer service capability to help designers and users of Multibus and iRMX products.

The three new SBCs are the iSBC 486DX66, the iSBC 486/166SE and the iSBC 486/150. The 486DX66 and 486/166SE use Intel's 66-MHz 486DX2 processor, while the 486/150 uses the 50-MHz version. Intel's DX2 technology lets the microprocessor operate at double the system clock speed. The 66-MHz DX2 processor, for example, works with 33-MHz-based systems, while the 50-MHz version works with 25-MHz systems. Intel's 66-MHz 486DX2 is rated at 54 Mips, while the 50-MHz device operates at 41 Mips.

"DX2 technology gives new customers the highest-performance Multibus II single-board computers on the market," says Dick Binns, Multibus marketing manager at Intel. "For our existing customers, the new Embedded Workstation, the iSBC 486/166SE and the iSBC 486/150 are plug-and-play upgrades to existing products. They are so compatible that customers merely need to check their timing loops and go. No other changes to the system hardware or software are required."

The Embedded Workstation is the most highly integrated SBC of the new offerings. It includes high-end PC graphics, integrated I/O and standard PC-compatible software that uses the high-bandwidth multiprocessing power of the Multibus II architecture. It can carry up to 32 Mbytes of on-board SIMM DRAM, IDE or SCSI peripheral interface and optional Super VGA graphics or PC Ethernet networking. On the software side, the Embedded Workstation supports DOS, iRMX for Windows and UNIX System V Release 4.0, Version 3.

The top performer of Intel's traditional Multibus II SBC releases is the 166SE, a superset of the iSBC

COMPUTERS & SUBSYSTEMS

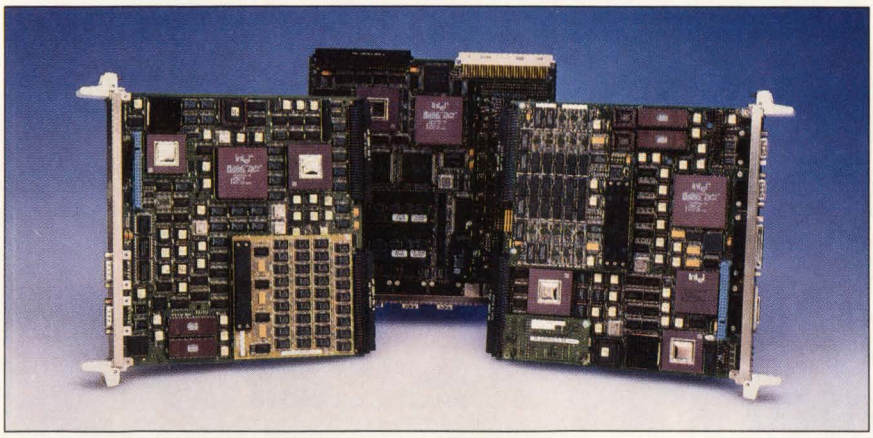

Intel introduced three new Multibus II single-board computers at this fall's Buscon East, becoming the first company to offer a Multibus II SBC based on DX2 microprocessor technology. The iSBC 486/150 (left) utilizes the 50-MHz Intel 486DX2, rated at 41 Mips. The iSBC 486DX66 Embedded Workstation (center) is PC-compatible and uses the 66-MHz Intel 486DX2 microprocessor. The iSBC 486/166SE (right) is a highly integrated CPU board that also uses the 66-MHz processor.

486/133SE. The latest version includes SCSI and Ethernet connectors, can handle anywhere from 8 to 64 Mbytes of byte-parity-protected, fast-page DRAM, and includes sites for EPROM and flash memory support.

The Multibus I board, the iSBC 486/12SDX2, is the highest-performing Multibus I SBC from Intel—or from anyone, for that matter. The board takes advantage of the company's DX2 technology and features a 66-MHz 486 that brings the board up to the 54-Mips performance level.

Despite the fact that Multibus I has passed its peak, continual board-level upgrades have kept the bus viable in many applications. "This new board," says Binns, "demonstrates Intel's continued commitment to the Multibus I product line. We have shown over the years that, as Intel provides new microprocessor technology, we will make that available in our Multibus I designs. Because the X86/12 single-board computer architecture is arguably the most popular and prevalent in the total single-board computer installed base, we pay particular attention to our existing customers."

Intel is not only bringing a new family of boards to the Multibus table, but it's also adding a new service organization to its existing sales operation. The new organization for Multibus I, Multibus II, iRMX, and iRMX for Windows products provides a direct line for customers in need of technical or business information.

All four iSBC boards are available now. For quantity one, the 486DX66 is priced at $5,250 (no memory), the 486/166SE at $7,515 (8 Mbytes memory), the 486/150 at $5,737 (8 Mbytes memory), and the 486/12SDX2 at $7,120 (8 Mbytes memory) and OEM pricing applies.

— Warren Andrews

Multibus SBCs at a glance

- 66-MHz 80486 processor-based
- Uses Intel's DX2 clock-doubling technology
- High-end PC graphics
- Integrated I/O plus IDE and SCSI peripheral interfaces
- PC-compatible software
- Optional Super VGA

Intel
5200 N E Elam Young Pkwy
Hillsboro, OR 97124
(800) 438-4769
Circle 291

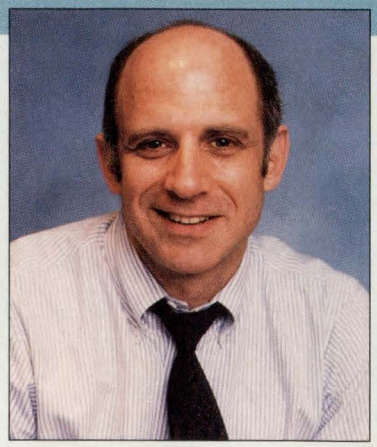

"Design for X:" A new push for manufacturability, testability and reliability

I learned a little about reliability analysis working for the Department of Defense 25 years ago. Let's say you're working on a massively destructive weapons system. You can't fully test the thing. You hope and pray, in fact, that it will never be used. But you must be absolutely certain—certainty, according to the Department of Defense is 99.968 percent—it'll work if you need it.

Reliability is but one of the product-support variables CAE tool vendors are considering under the banner of "Design for X." Joseph Costello, president of Cadence Design Systems (San Jose, CA) and co-chairman of EDAC, the CAE tool manufacturers' trade association, called attention to this during a press luncheon at the Design Automation Conference in July. The "X" stands for all the back-end processes designers and design tool vendors are just starting to consider—assembly, manufacturing, reliability, repairability, serviceability, and test.

Focusing on the nuts-and-bolts

Most of the attention of CAE software developers (and of trade magazine editors) has been concentrated on the glamorous, performance-driven design side of the engineering world. Less attention has been placed on bread-and-butter issues, such as the cost of manufacturing the supersystem and the need to shave pennies—if not dollars—from the total. We can design great electronic systems, but can we test them? Can we manufacture them cheaply and in volume? Can we service them when they break? Do we know how long they'll last? Attention to the stringent, less glamorous requirements of Design for X is where we'll need to look for answers to such questions.

Racal-Redac (Mahwah, NJ and Tewksbury, U.K.) was among the first companies to emphasize design for manufacturability as an approach to PCB layout. Mentor Graphics (Wilsonville, OR) is a more recent convert; its Manufacturing Advisor was introduced late last year. Cadence, with solid strength and experience in ASIC design, is just beginning to enhance its system- and board-level tool offerings with improved links to manufacturing.

Design for manufacturability acknowledges the gulf between those electrical engineers who design electronic circuits and those who must implement them on a PCB. If the truth be known, this clash of cultures isn't as great between ASIC designers and manufacturers as it is between system designers and manufacturers. The ASIC designers use design rules, cell libraries—an entire toolset, in fact—that's been tweaked for the fabrication of semiconductors. On most frameworks you can't invoke a simulator (analog or digital) without specifying which fabrication process you're using. Even where the IC is designed entirely at a customer's plant, it's done with the ASIC vendor's tools, model libraries, layout contingencies—in other words, with an intimate knowledge of what the IC manufacturer is capable of producing.

The laws of physics apply

The gap between design and manufacturing, however, is much wider for board-level products. There are differences between how designers want their board-level system to perform, and what you can practically stuff and solder onto a piece of copper-coated phenolic. Manufacturing and assembly engineers worry not just about the thickness or thinness of your traces and the orientation of pads and footprints, but also about the accuracy of hole-drilling equipment, component inserters and pick-and-place machinery. They think about how often this equipment will bend a component lead that fails to meet a throughhole, the relationship between solder defect rates and lead-frame spacing and the percentages attached to manufacturing-induced shorts and opens.

While the people who design high-performance digital systems typically worry about the number of nanoseconds between the rising edge of a clock pulse and the appearance of valid data on an output line, the people who actually build these systems worry about the number of minutes the board sits in a

THROW AWAY YOUR EXPENSIVE 1488s, 1489s, 75175s...

Eliminate 3 ICs with One RS-232 IC at Less Cost

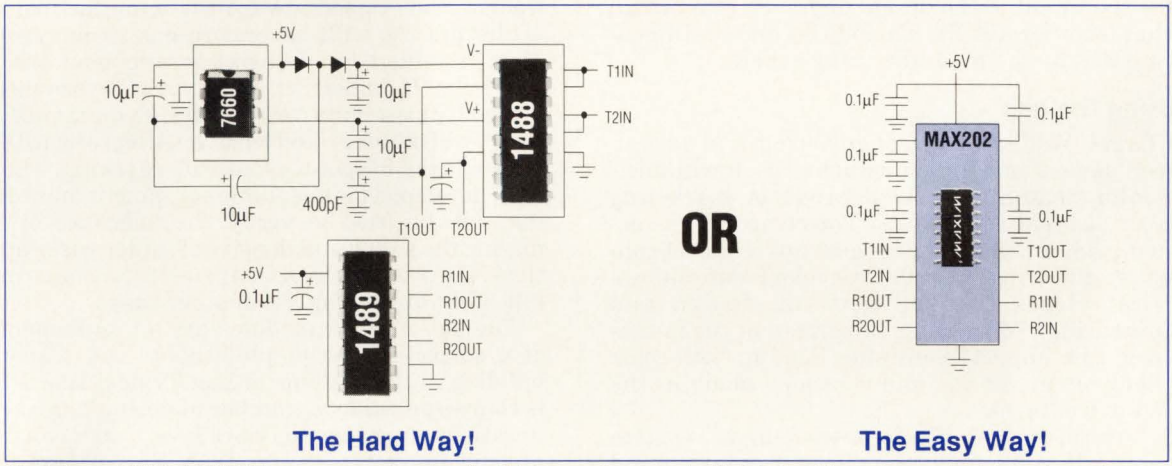

Maxim's new MAX200-MAX211 use space and cost-saving 0.1µF capacitors operate off of single +5V supply, operate at data rates up to 120kbits/sec, are available in the smallest packages possible, and are slew rate limited so no filter capacitors are required.

Call Maxim for A Budgetary Quote Today! TOLL-FREE 800-998-8800

Product	RS-232 Drivers/Receivers	External Caps (0.1µF)	Description
MAX200	5/0	4	5 drivers with shutdown
MAX201	2/2	2	Standard +5V/+12V or battery supplies
MAX202	**2/2**	**4**	**Industry standard MAX232 pinout**
MAX203	2/2	none	No external caps
MAX204	4/0	4	4 drivers; replaces 1488 in +5V system
MAX205	5/5	none	No external caps
MAX206	4/3	4	Ideal for modems
MAX207	5/3	4	Ideal for modems
MAX208	4/4	4	Ideal for laser printers
MAX209	3/5	2	Standard +5V/+12V or battery supplies
MAX211	4/5	4	Complete +5V AT port; SSOP package

FREE Interface Design Guide – *Sent Within 24 Hrs!*

Includes: Data Sheets and Cards for Free Samples

CALL TOLL FREE 1-800-998-8800

For a Design Guide or Free Sample

Master Card® and Visa® are accepted for Evaluation Kits or small quantity orders.

MAXIM

Maxim Integrated Products, 120 San Gabriel, Sunnyvale, CA 94086, (408) 737-7600, FAX (408) 737-7194.

Distributed by Arrow, Bell/Graham, Elmo, Hall-Mark, Nu Horizons, and Pioneer. Authorized Maxim Representatives: **AL**, M2i Montgomery Marketing, Inc.; **AZ**, Techni Source Inc.; **CA**, Mesa, Pro Associates, Inc., Centaur Corporation; **CO**, Component Sales; **CT**, NRG Limited; **DE**, TAI Corporation; **FL**, Sales Engineering Concepts; **GA**, M2i Montgomery Marketing, Inc.; **ID**, E.S. Chase; **IL**, Heartland Technical Marketing Inc.; **IN**, Technology Marketing Group; **IA**, JR Sales Engineering, Inc.; **KS**, Delltron; **LA**, BP Sales; **MD**, Micro-Comp, Inc.; **MA**, Comp Rep Associates; **MN**, Mel Foster Technical Sales, Inc.; **MS**, M2i Montgomery Marketing, Inc.; **MO**, Delltron; **MT**, E.S. Chase; **NE**, Delltron; **NV** (Reno, Tahoe area only) Pro Associates, Inc.; **NH**, Comp Rep Associates; **NJ**, Emtec Sales, Inc., TAI Corporation; **NM**, Techni Source Inc.; **NY**, Parallax, Reagan/Compar; **NC**, M2i Montgomery Marketing, Inc.; **OH**, Lyons Corporation; **OK**, BP Sales; **OR**, E.S. Chase; **PA** (Pittsburgh area) Lyons Corporation, (Philadelphia area) TAI Corporation; **SC**, M2i Montgomery Marketing, Inc.; **TN**, M2i Montgomery Marketing, Inc.; **TX**, BP Sales; **UT**, Utah Component Sales, Inc.; **VA**, Micro-Comp, Inc.; **WA**, E.S. Chase.

Distributed in Canada by Arrow and Future. Authorized Maxim Representative in Canada: Tech Trek.

CIRCLE NO. 70

chemical etching bath. Small details, easy to overlook, can make an enormous difference to the manufacturing engineer. Teardrop-shaped pads, for example, prevent the copper etchants from creating an electrical open by eating away at the sharp corners at junctions of pads and traces.

The goal of design for manufacturability is to bring these cultures closer together: To provide a layout tool to the manufacturing engineer that somehow communicates the intent of the design engineer, or a design tool that incorporates the capabilities and contingencies imposed by the manufacturing expert.

Closing the gap

"The Great Wall between engineering and layout, between layout and manufacturing, is inevitable," says John Seaton, advanced products marketing manager at Mentor Graphics. The company has positioned itself among the strongest advocates of concurrent engineering, a product development philosophy that closes the gap between design and manufacturing. "More than 70 percent of the manufacturing cost impact," continues Seaton, "will come from component choices and decisions made in the design environment."

But Seaton has no illusions that digital system designers will suddenly become expert in layout and manufacturing. "You don't want them to worry about that," he says. The better solution is to rely on a tool such as Mentor's Manufacturing Advisor/PCB, an addition to the company's BoardStation product that will provide direct, online feedback on the manufacturing impact of particular board design choices.

Racal-Redac, I believe, was among the first companies to offer PCB layout tools truly intended for manufacturing. Its design rule checks incorporate long lists of manufacturing contingencies. If you lay down the footprint for an SO-packaged device, for example, the system verifies not only whether the pad sizes and lead pitch are correct for the device you're placing, but also whether the pads are oriented properly for the soldering technique you're using, whether the placement of throughholes matches the resolution of your hole-drilling equipment, and even whether you've left enough space on a tightly packed board for the fingers of your component inserters and other pick-and-place machinery. Racal-Redac's post-processors provide not just the standard Gerber photoplots, but also numerical control (NC) tapes for drilling and milling machinery.

Key to the utility of Racal-Redac's Visula toolset is a manufacturing database that lets you assign information to each pin or connection point. Based on the Informix relational database, Visula supports up to 256 layers of information for each point, according to Dave DeMaria, Racal-Redac's director of product development in Tewksbury. Electrical information, voltages and currents, materials, pad dimensions, machine tool access points, photo mask tolerances, preferred vendors, cost data, NC file formats—all types of practical information can be accessed with a query from any point on the PCB layout. It's this type of database, says DeMaria, that's made Visula a preferred layout tool for complex new board designs with computer systems from companies such as Apple Computer and Data General.

While implementing rules-driven design, Cadence's Allegro board layout tools have traditionally been tilted toward the design part of the engineering process. You use them to complete a layout to get a better look at the layout-dependent parasitics. But Cadence has tipped toward the manufacturing part of the process with several recent announcements. Most prominent is a licensing arrangement with AT&T Bell Labs (Holmdel, NJ). Two AT&T software modules, design for assembly (DFA) and design for manufacture (DFM), will be integrated with the Allegro board design set and used for post-layout rule checking. The modules, developed for telephone equipment manufacturing, will be used to verify the tolerance of solder masks, the shape and density of solder paste applications, the orientation of components for soldering and assembly, and a long list of other tasks.

The DFA and DFM modules are an implementation of a correct-by-design philosophy that Cadence is building in the systems arena. Typical board design is characterized by a number of design/analyze/fix-it iterations, according to Debra Ives, Cadence's Allegro product marketing manager in Chelmsford. "If the system constraints are properly specified to the toolset," she says, "the layout tools will always make these visible to the user." Manufacturability analysis, in fact, becomes just one of many checks the tools will perform to assess the viability of a design. Cadence's board-level reliability assessment tool, Viable, introduced in December of 1990, for example, will assess the reliability of a system-level product, using thermal stresses and cost data in its calculations.

Reliability modeling technology

For the semiconductor industry, the ability to model reliability would mean quicker time-to-market. The normal process of burn-in and reliability testing can delay the introduction of a new product until many months after the piece has been designed and a test program developed for it. Reliability modeling lets a manufacturer put together statistically significant data about a product even before it goes into full production.

The implications of this aren't just for weapons and semiconductors. An engineer at Ford Motor Company told me he was doing extensive work on modeling the in-rush current of headlamps. If you think about it, the filament of an electric lamp will burn itself out just at the point you flip on the switch. The goal of reliability analysis, like many other projects under the Design for X banner, is a less expensive product with a longer product life. "You want to start this on a cold day in Minnesota or on a hot day in Texas," says Ford modeling engineer Gary Zack, "and know that it's going to work."

Stephan Ohr is president of Indian Forest Research and publisher of the monthly newsletter, Mixed Signals.

Announcing Two New Conferences

ISO 9000

MAKING IT WORK FOR YOU

THE BENEFITS OF ISO 9000 CERTIFICATION

- *Increases your competitive advantage*
- *Ensures control, consistency and assurance of high standards*
- *Improves product quality and productivity*
- *Reduces costs*
- *A listing in the registry of suppliers*
- *Use of recognized certification mark in marketing and promotion*
- *Access to the European Economic Community*

QUALITY- A Matter of Survival

There was a time when higher performance and a steady stream of new product introductions were all that it took for an electronics or computer company to win the hearts of both business and consumer customers. Those days are long gone and everyone is demanding more. The key to success today lies as much in providing quality products and quality customer service as it does in putting the most sophisticated, leading-edge technology to work.

A SPECIAL CONFERENCE SERIES PRESENTED BY PENNWELL'S ADVANCED TECHNOLOGY GROUP

Computer Design • Laser Focus World • Industrial Laser Review • Solid State Technology • Lightwave • Military & Aerospace Electronics

SANTA CLARA, CA • BILTMORE HOTEL - MARCH 29, 30, 31,1993
BOSTON, MA • MARRIOTT LONG WHARF - APRIL 12,13,14, 1993

ISO 9000
Making It Work For You

Unfortunately, early attempts at improving quality didn't always prove successful. First there were quality circles and quality councils. Then came concurrent engineering, cross-functional workteams, worker empowerment and, now, total quality management (or TQM). For many companies, the attempts to improve quality failed — sometimes miserably — and all of these high-sounding terms — and the ideals they represented — become nothing but buzzwords.

Perhaps more than anything else, companies failed at improving the quality of their products or services because there was no consistent, generally accepted framework within which they could develop their quality programs. Even more, there was no generally accepted performance standard by which they could measure the effectiveness of their quality management systems. In short, companies didn't know what to do, and didn't know what to do about it.

The International Organization for Standardization (ISO) and ISO 9000 are changing all that!

Developed over a period of several years by representatives from the 91 countries that are members of the ISO, ISO 9000 has become an internationally recognized family of specifications for quality assurance management systems.

The ISO 9000 standard (really a set of five standards, numbered 9000 through 9004) isn't a set of product specifications and it's not specific to any one industry. Rather, ISO 9000 defines a process — one that ultimately leads to certification — to ensure that a company's procedures for quality management and quality assurance for procurement, material management, design control, production process control, customer interface, and servicing are under control and are producing consistent results. It also defines guidelines for the activities needed in each of these areas to identify, analyze and correct problems.

Some of the benefits of adhering to the applicable ISO 9000 standard are obvious. First, quality and productivity are improved. Second, consistency in following other established standards is ensured. Third, costs are reduced by doing things right the first time and every time. The bottom-line result is that your company's ability to compete in the world market is improved. In fact, it may soon be impossible to compete globally without ISO certification.

There's an irreversible trend in political and trade policy toward quality system registration and more than 50 countries have already adopted ISO 9000 as a national standard. In the U.K. alone, more than 20,000 companies are now ISO certified, and certification will soon be a requirement for access to the European Economic Community (ECC). While only a few hundred companies in the U.S. are currently certified, and only a handful of these are electronics or computer

GLOBAL IMPACT OF ISO 9000

- Worldwide political and trade policy trend toward quality registration
- More than 50 nations have adopted the standards
- 350 companies in the U.S. are now registered
- 550,000 U.S. companies within the next 5 years
- Within 5 years, quality system registration will be fundamental to doing business on a global basis

"ISO 9000 is documented common sense. There is no black magic, no voodoo, no great complicated indoctrination program. Just plain, old-fashioned common sense - Doing Right Things, Right!"

Jim Carras,
The Carman Group

ISO 9000 — DEFINED FUNCTIONS AND ACTIVITIES

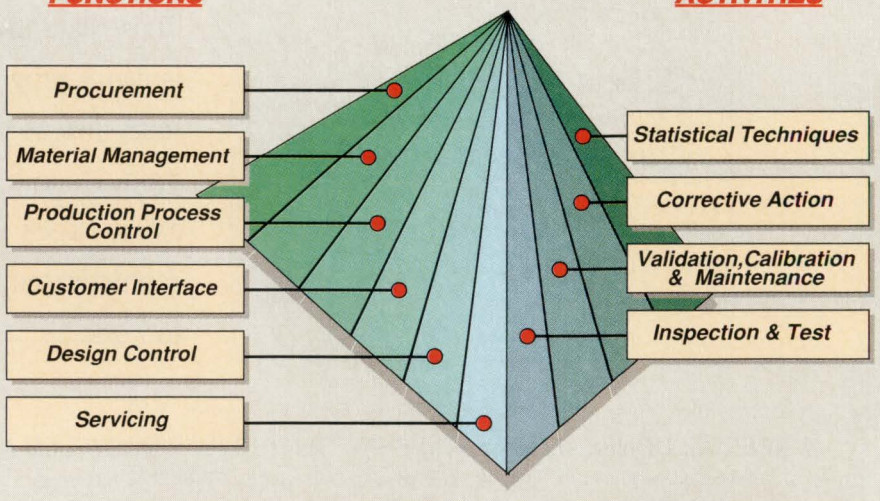

ISO 9000
Making It Work For You

companies, it's estimated that more than a half-million will be certified in the next five years.

Make no mistake about it, By the end of the decade, and most likely before, ISO 9000 certification will be a requirement for doing business on a global basis.

As you might expect, obtaining ISO 9000 certification isn't a cake walk. The procedures are rigorous and the auditing teams sent in by the organizations authorized to qualify a company for certification are thorough and soul-searching. Contacting the ISO for its list of accredited certification organizations, signing a contract with one of them, signing a check, and then exposing yourself to scrutiny isn't the best way to go about certification.

Grayhill Inc. which manufactures switches, solid-state relays and distributed control components, has just been registered to ISO 9001. Because of the registration, Hewlett-Packard, which audits Grayhill annually using its own three-person team, has said its audit would be unnecessary. According to Brian May, marketing manager for Grayhill, his company sought registration for several reasons. One was to gain an edge over competitors who weren't registered.

"Then there's the issue of making Grayhill a better company and being able to reduce redundancies in our system and reduce our scrap rate. And then there's the financial payback from doing all of this," says May.

A better way is to get educated about ISO 9000 first. Start by learning the details of the 20 specific ISO 9000 quality management system requirements, and the most important aspects of each of these requirements. Learn what it takes to define a quality management system, how to prepare a Quality Manual and how to select, train and qualify your own in-house auditors. Then get organized for implementation and implement an ISO 9000 Quality Manual. Then contact an accredited certification organization.

But how do you get educated about ISO 9000 and what it takes to get certification? Fortunately, there are consulting and training organizations, such as The Carman Group in Richardson, TX, that can help you through these early stages. They play an important role in the overall certification process, because the ISO expressly forbids organizations that perform ISO 9000 assessments and audits from advising companies on how to set up their quality systems or to write their quality documents. Companies seeking ISO 9000 certification that don't pursue pre-audit or preparation services suffer a 70 percent failure rate the first time through their audit. That's an expensive way to learn what ISO is all about.

Before long, trying to survive in a marketplace that demands quality as much as high performance products — and want some proof that you can provide the quality you claim — will be impossible. ISO 9000 certification is a critical part of that proof and the time to take the first step toward certification is now. That first step is learning what it takes to become certified and The Carman Group can help with it's ISO 9000 seminars created specifically for high technology companies, and in particular, for those companies in the electronics, computer and electro optic industries.

THE ISO 9000 STANDARDS:

- **ISO 9000:** Quality Management and Quality Assurance Standards — Guidelines for Selection and Use
- **ISO 9001:** Quality Systems — Model for Quality Assurance in Design and Development, Production, Installation, and Servicing
- **ISO 9002:** Quality Systems — Model for Quality Assurance in Production and Installation
- **ISO 9003:** Quality Systems — Model for Quality Assurance in Final Inspection and Test
- **ISO 9004:** Quality Management and Quality Systems, Elements and Guidelines

YOU'RE IN A SELECT GROUP IF YOU'RE ISO 9000 CERTIFIED

- AT&T
- Bell Northern Research
- Compaq Computer
- Data General
- Digital Equipment Corp.
- Dupont
- GE Fanuc
- Grayhill
- Hewlett-Packard
- Lasertron
- Northern Telecom
- Prime Computer
- Signetics
- Square D
- Sun Microsystems
- Texas Instruments
- Unisys
- Xerox

ISO 9000
Making It Work For You

The Carman Group Company, in association with PennWell Publishing, has created a specially tailored version of its successful three-day ISO 9000 seminars that's intended to meet the needs of corporate management (Presidents, General Managers, Sales Managers, Quality Managers, etc.) and design, test and manufacturing management (Directors of Engineering, VPs of Engineering, Product Managers, Test Engineering Managers, Senior Design and Test Engineers, Manufacturing Managers and Engineers, etc.) at computer/electronic/electro optic OEMs and OEM suppliers.

DAY 1:

Recommended for corporate managers who need to learn about ISO 9000 and its impact on their organizations.

- Benefits of ISO certification
- Six functions covered and the activities within each function
- The steps to certification and maintenance
- Typical timeframes and fees
- Application of ISO standards to your company requirements
- The Corporate Needs Assessment
- Strategic and operating plans
- Comparison of ISO to Malcolm Baldrige Award and MIL-Q-9858A

DAYS 2 & 3:

Recommended for those technical managers and engineering personnel who will actually be writing the documentation and procedures for ISO certification.

- Definition of a Quality Management System
- Global objectives of procedures
- Relationship of each ISO 9001 paragraph to key points, manuals and procedures
- Write a procedure to select, train, qualify, and certify your own in-house auditors
- Plan an audit/compliance schedule
- Provides a workshop environment to gain insight from other company representatives

REGISTER TODAY!

A SPECIAL CONFERENCE SERIES PRESENTED BY PENNWELL'S ADVANCED TECHNOLOGY GROUP

Computer Design • Laser Focus World • Industrial Laser Review • Solid State Technology • Lightwave • Military & Aerospace Electronics

SANTA CLARA, CA • BILTMORE HOTEL - MARCH 29, 30, 31,1993
BOSTON, MA • MARRIOTT LONG WHARF - APRIL 12,13,14, 1993

MAIL application with your check payable to:
Seminar Coordinator-The Carman Group, Inc.
P.O. Box 835088, Richardson, TX 75083-5088

or REGISTER by PHONE: 800-942-6880
or FAX this application to: 214-669-9478

Seminar Location: _____ Seminar Date: _____

Company Name: _____

Address: _____ City _____ State _____ Zip _____

Business Phone: (____) _____ Extension _____ Fax: (____) _____

Name of Attendee

1. _____	☐ **Day 1 $295**	☐ **Days 2&3 $795**	☐ **Days 1,2&3 $995**
2. _____	☐ **Day 1 $295**	☐ **Days 2&3 $795**	☐ **Days 1,2&3 $995**
3. _____	☐ **Day 1 $295**	☐ **Days 2&3 $795**	☐ **Days 1,2&3 $995**
	Total $_____	Total $_____	Total $_____

IF YOU REQUIRE ADDITIONAL INFORMATION, CALL (214) 669-9464

Grand Total $_____
(Deduct 10% if registering three or more from the same company)

If seminar is cancelled for any reason, the liability of the Carman Group is limited to a refund of the course fee paid. Apply a 10% discount if registering three or more from the same organization. Registration is not guaranteed until check or Purchase Order has been received by the Carman Group. Full refund for cancellations up to 2 weeks prior to the beginning of the seminar. Cancellations between 2 weeks and 3 working days prior to the start of the course are subject to a $150 cancellation charge. Cancellations less than 3 days prior to the start of the course are subject to the full course fee.

CIRCLE NO. 71

SYSTEM SHOWCASE

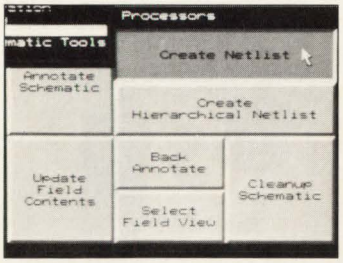

IT CAN BE THIS EASY!

OrCAD's innovative ESP Framework can make creating a netlist from your schematic as easy as pressing an on-screen "button."

Call **(503) 690-9881** for a ***FREE*** demonstration disk and information about the entire **OrCAD** product line: schematic design, simulation, pc board layout, programmable logic IC design and more.

CIRCLE NO. 166

In-Circuit Emulators For

8051 683xx 68HC11 68HC16

SEE EEM 92/93 Pages D 1286-1291

- Rental and 10-day trials available.
- Full Source-level Debugger *w/complete* C-variable support.
- Supports virtually all members of the 8051, 68HC11, 68HC16 and 68300 families.

CALL OR WRITE FOR FREE DEMO DISK!

nohau CORPORATION 51 E. Campbell Avenue Campbell, CA 95008 FAX (408) 378-7869 **(408) 866-1820**

Call 408-378-2912 Nohau's 24-hour information center to receive info via your FAX

CIRCLE NO. 167

8051 IN-CIRCUIT EMULATOR

$995.00

- 8031/32/51/52 support
- 64K Memory/64K breakpoints
- Source Level C and Assy debugging
- Trace Buffer ■ Programmers
- Money Back Guarantee

VAIL SILICON TOOLS
305 570 5580

CIRCLE NO. 168

REMOVE HARDWARE LOCKS

Protect your investment!

Don't wait for your lock to get lost, stolen or fail and render your program useless.

Our **SAFEKEY'S** are guaranteed to unlock the following:

PADS 2000	MICROCADAM	TOPAS
P-SPICE	MICROSTATION	PCAD
CADKEY	3D STUDIO	& MORE

From $399 plus S&H

(609) 390-2799

24 HR. FAX HOTLINE (609) 390-3750

IMAGINE THAT

1 Evergreen Dr., Seaville, NJ 08230

TM-Trademarks property of respective owners.

CIRCLE NO. 169

The Ultimate SBus Extender

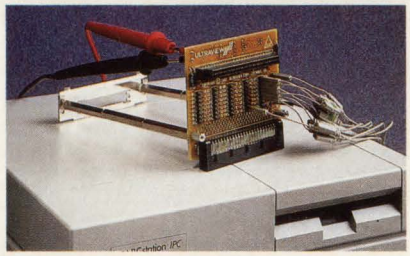

- All 82 SBus test points in Logical order for *fast* logic analyzer hookup.
- Voltmeter output jacks indicate current consumed by board-under-test, and LEDs warn if I >2A or V< 4.75.

ULTRAVIEW

475 Yampa Way, Fremont CA 94539
Phone 510-657-9501 Fax 510-657-0927

CIRCLE NO. 170

QUICK-TURN PC BOARDS

INSTANT QUOTES

1-800-234-1556

(After Hours (602) 377-7319)

Single/Double Sided/Multi-Layer

Gerber Input/Floppy or Modem

MIL Approved U.L. Listed

AVANTI CIRCUITS, INC.

17650 North 25th Avenue
Phoenix, Arizona 85023

(602) 863-7729 & 234-1556
Fax (602) 375-1909
BBS/Modem (602) 234-1737

CIRCLE NO. 171

386/486

Free Catalog!

Bolt-in Computers™

for embedded applications

High performance CPUs for aggressive applications. From full blown VGA flat panel, 8MB RAM, 510 MB IDE systems. To number crunching high throughput ROM-based systems.

- *16-25 MHz 386SX 486SLC*
- *8MB RAM, 1.8 MB ROM*
- *Solid state or IDE disk for DOS 5.0, Windows 3.1*
- *Multitasking, multidrop comm, PID control, multi-CPU software*
- *Optional STD BUS expansion*

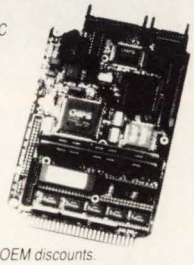

From $995, quantity 1, great OEM discounts.

MICRO/SYS

3447 Ocean View Blvd., Glendale, CA 91208
(818) 244-4600 FAX (818) 244-4246

CIRCLE NO. 172

Half size ECL Clock Oscillator

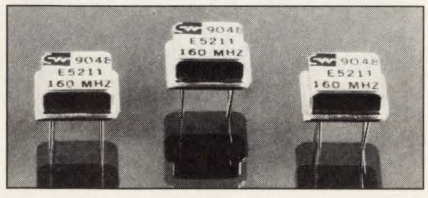

Connor-Winfield Corp. introduces its new 8-pin DIP ECL clock oscillators. The E500 Series is available in frequencies from 24 MHz - 180 MHz. Options available include industrial temp range of -40C to +85C, complimentary output, enable/disable function, and surface mount header. Frequency stabililites are available to +/-25 ppm over temperature.

These models are available with supply voltage of either -5.2 V dc (10K equivalent) -4.5 Vdc (100k equivalent) or +5Vdc. Prototype quantity pricing at 120 MHz is $43.90 each. Delivery is stock to 7 weeks A.R.O. Contact Barney Ill at: **Connor-Winfield Corporation**
1865 Selmarten Road, Aurora, IL 60505
PH: (708) 851-4722 FAX: (708) 851-5040

CIRCLE NO. 173

NEW! **Little PLC™** $195

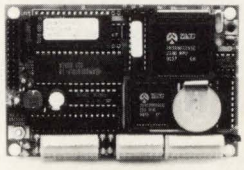

Program It In C

Our new **Little PLC**™ measures only 4.33 x 2.85 inches and can mount on standard DIN rail. This miniature controller costs only $195, including 8 optically isolated inputs and 8 relay driver outputs. Low cost expansion cards allow you to add more inputs and outputs: digital and analog. It has dual RS-485 serial I/O, battery backed memory and time/date clock, programmable timers and a watchdog. Our easy to use and affordable **Dynamic C**™ integrated development system also costs $195. You can write simple programs in an hour, or you can develop major applications with 20,000 lines of C language.

Z-World Engineering

1724 Picasso Ave., Davis, CA 95616
(916) 757-3737 Fax: (916) 753-5141
24 hr. Automatic Fax: (916) 753-0618
(Call from **your fax**, request catalog #18)

CIRCLE NO. 174

SYSTEM SHOWCASE

EPROM PROGRAMMER

- **Quick pulse** pgms. eight 1 Mbit EPROMs in 40 sec.
- Stand-alone or PC-driven
- RS-232, **parallel** in & out ports
- Binary, Intel hex, & Motorola S formats
- 100 user-definable macros
- Information, call **(916) 924-8037**
- 1 Megabit of DRAM
- Made in U.S.A.
- A9 Identifier
- 2 year warranty
- Single pgmr. $550

NEEDHAM'S ELECTRONICS
4539 Orange Grove Ave. • Sacramento, CA 95841
(M-F, 8-5 PST)

CIRCLE NO. 175

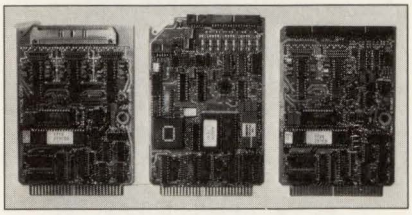

STD ANALOG I/O GET A COMPETITIVE EDGE!

Density—up to 16 A/D inputs and 8 D/A outputs on one card. **Speed**—throughput rates up to 59 kHz. **Intelligence**—many pre-programmed modes. Input filters, prog. gain amps, sample-holds, FIFO I/O buffers & many other features too numerous to mention here.

ROBOTROL CORP.
925 W. San Martin Ave, PO Box 990, San Martin, CA 95046
(408) 683-2000

CIRCLE NO. 176

High Speed Interconnection

New flexible circuit cable features integrated shielding, grounding and controlled impedance yet can be bent or folded to interconnect multiple planes.

- Designed for high-speed applications
- Available in Stripline or Microstrip
- Compatible with all available connectors, including ZIF
- Meets FCC emission standards

Sheldahl
Northfield, MN 55057
507/663-8000 ext. 246

CIRCLE NO. 177

WHY?!

BURN EPROMS?
when
Quality, FAST-LOADING* EPROM EMULATORS
start at ONLY

$129

(27256 - LARGER MODELS AVAILABLE)

BEAT YOUR NEXT DEADLINE -
Call or FAX for Engineering Specs
Voice or FAX: (214) 272-9392

Technical *Solutions*
PO BOX 462101 Garland, TX 75046-2101

*DOWNLOAD 27512 In UNDER 3 SECONDS!

CIRCLE NO. 178

357 MHz 14-Pin Clock Oscillator

Connor-Winfield Corp. announces literature on a new line of high frequency ECL Logic clock oscillators. The ECLB Series oscillator covers the high frequency range of 8 MHz-360 MHz. AVAILABLE NOW IN A 14-PIN DIP. Literature for a double DIP version with frequencies to 500 MHz is also available. Frequencies to 90 MHz are now available with Voltage Control function (model EV53 series). Contact Barney Ill for details.

Additional specifications listed below:

Model	:	ECLB
Package	:	14-pin DIP
Frequency	:	8 MHz to 360 MHz
Supply	:	-5.2 or -4.5 Vdc

Connor-Winfield Corporation
1865 Selmarten Road
Aurora, IL 60505
PH: (708) 851-4722 FAX: (708) 851-5040

CIRCLE NO. 179

FINALLY

New Lower Prices

One tool to satisfy all your firmware development needs

PROMICE is a universal system.
- Develops code for any microprocessor
- Complete, real time, source level debugging
- Host software for DOS, Unix, Mac, VMS
- Non intrusive on your target system
- Simply plugs into any ROM socket

PROMICE also supports Turbo Debugger, C_thru_ROM, FreeForm, GDB, and more.

PROMICE
...the affordable solution.

Grammar Engine Inc.
TEL: (614) 899-7878
West: (415)750-0219
FAX: (614) 899-7888

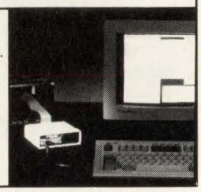

CIRCLE NO. 180

CapFast Schematic Capture and Interface Tools

Price Starts at $595.00

Available on both DOS, Windows and UNIX (SPARC, DEC; HP, IBM)

CapFast is an advanced hierarchical schematic design and interface tool for PCB, PLD/FPGA, as well as ASIC designs.

The optional EDIF 200 tools allow the designer to translate CapFast schematics to and from other EDA systems, such as Mentor Graphics and Cadence. For further information, please call **(503) 645-0313.**

Phase Three Logic

CIRCLE NO. 181

Multibus II Parallel System Bus Analyzer

PSBA-100 ***A productivity enhancement tool for the system integrator, field service engineer, software, and hardware engineer.***

- Stand-alone, single board computer
- 6U form factor, installs in a single slot
- Built-in, terminal-based user interface
- Supports all four address spaces
- Data capture based on Multibus II protocol
- Fully programmable filter logic uses templates based on Multibus II protocol
- Filter logic supports 16 trigger levels

Multibus II is a trademark of Intel Corp.

PRC
PRC Inc.
1410 Wall Street
Bellevue, NE 68005

For more information and free Demo Disk, Call .... **(402) 293-3900**

New Features!

CIRCLE NO. 182

BLAZING GRAPHICS

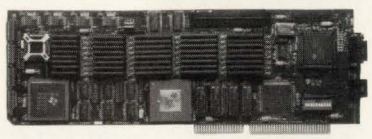

- Windows 3.1 compatible
- 640x480 to 1600x1200
- 15, 16 & 24-Bit True Color
- 3D Modeling, Fractal graphics
- User Programmable
- RGB & NTSC/PAL outputs
- 3 Year Warranty
- TMS34020, 32/40 MHz
- Floating Point Processor
- 512 KB Local SRAM for FPP
- 4 MB VRAM, 4 MB DRAM
- 512 KB EPROM site
- User Installed Upgrades
- Free TIGA Driver

ZWICK SYSTEMS INC.
17 Fitzgerald Road, Suite 104, Nepean, Ontario, Canada, K2H 9G1 (Tel) 613-726-1377 (Fax) 613-726-1902

CIRCLE NO. 183

EMPLOYMENT OPPORTUNITIES, OUTPLACEMENT AND POSITIONS WANTED

CAREER CONNECTIONS

Recruitment Advertising
Sue Shorrock, (508) 392-2185

Closing Dates for upcoming magazine issues

Issue	Closing
JAN	DEC 4
FEB	JAN 5

POSITION WANTED ADS

Free 1" ad to subscribers seeking full-time employment. Just include 50 words of copy and your subscription label. We'll run your ad in 2 consecutive issues. Available to non-subscribers or consultants/companies at $125 per column inch. Mail your position wanted ad to: COMPUTER DESIGN Positions Wanted, One Technology Park Drive, PO Box 990, Westford, MA 01886

COMPUTER DESIGN Technology and Design Directions

Fault tolerant computing, digital, analog , software hands-on design and analysis. Reliability, cost and performance modeling. Failure Modes and Effects, worst-case analyses. Multi-tasking software and hard real-time systems design. SOA system methodologies. Don Uhrich, DEPENDABLE SYSTEMS, 7088 Olde Stage Road, Boulder, CO 80302, (303) 440-4696.

SENIOR ENGINEERING MANAGER/ENGINEERING DIRECTOR: Over 20 years of successful and professional experience in major industrial electronic companies. Expertise in planning conceptual electronic design, with emphasis on managing, directing and developing people for project and product success. BSEE, MSEE and MBA. Resume on request. Bruce Carlton, 2045 E. Knox Rd, Tempe, AZ 85284, (602) 831-9699.

PROJECT MANAGER: Highly versatile professional with over 15 years of broad-based experience in Project Management, Economic Analysis and Engineering gained with blue chip firms (GTE, PacBell, FHP, etc.). Technically oriented, MBA, BS, and BA degrees with strong functional skills embracing all areas of research, applications engineering and business analysis. Call Bob Mazzola, (714) 841-8978.

SENIOR LEVEL MULTIMEDIA SYSTEMS MARKET ANALYST, STRATEGIC PLANNER AND PRODUCT DEVELOPER. Member of ANSI and ISO Experts Groups on: digital Image, Motion Video, Audio compression, and ISO Electronic Still Pictures Imaging. Experience developing MPEG Video chip, DSP-based MPEG-Audio software solution, Laser Disc-based Mulitmedia Applications. Call Alfred Riccomi (214) 644-8875.

GREATER BOSTON, MA AREA: Interested in adding an experienced COMPONENT/PRODUCT ASSURANCE or MANUFACTURING ENGINEER to your staff? Look no further! Over ten years experience: BS Electronic Engineering Technology; Supervised staff of 8; Military/commercial manufacturing; Call George J. Alexandropoulos (508) 692-7301. I'll FAX resume.

TECHNICAL WRITER. Eight years publications experience. Documentation and promotions including copy and scriptwriting. Knowledge of popular PC packages. Clients include CAD, robotics, software developers, design. Offer free diagnosis of your publications needs. Samples. Call Jay at (617) 523-4217.

COMPUTER DESIGN Technology and Design Directions 1993 EDITORIAL CALENDAR

A PennWell Publication

ISSUE	SPECIAL REPORT	TECHNOLOGY FOCUS	PRODUCT FOCUS	DESIGN STRATEGIES	OEM INTEGRATION
JANUARY	Verilog and VHDL	•DSP development tools •Disk drive controllers and associated ICs	C cross compilers	Networking	•Secondary & mezzanine buses •Backplanes & enclosures
FEBRUARY* *Buscon West, EDAC/Euro ASIC*	Bridging the buses	•Software tool integration •ICs for desktop video	Flash EEPROMs	Process control	
MARCH *RISC '93, PCB Design*	Memory architectures	•Multithreaded operating systems •PC-Based CAE/CAD Tools	SCSI host adapter boards	Data acquisition	•Input devices •Printers & output devices
APRIL*	*EMBEDDED COMPUTER CONFERENCE (ECC) SHOWGUIDE***				
Embedded Computer Conference	PC/AT architectures in embedded applications	•Developments in 3-V ICs •Evaluating simulation strategies	Device programmers	Peripherals	
MAY	*SPECIAL REPORT ON FUTURE COMPUTING: Virtual Reality*				
CICC	New applications for DSP	•Benchmarking programmable devices •Bus standards	Video D–A converters	Portable computers	•Networking interfaces, standards & components
JUNE* *DAC*	High–level synthesis and architectural design	•Small form factor VME •Data compression standards and ICs	Ultra–fast SRAMs	Graphics	
JULY *Fuzzy Logic '93*	Advances in IC packaging	•Interfaces for DSP •Fuzzy/neural update	Real–time kernels and operating systems	Imaging	•Display devices & monitors
AUGUST*	Trade–offs in programmable devices architectures	•RISC in real time •Integrating CAE and CAD databases	Emulators	Robotics	
SEPTEMBER *EuroDAC, Buscon East, Embedded Systems,* Wescon	Software testing and quality	•Futurebus+ •Integrating testability into the ASIC design process	Low–power DRAMS	Instrumentation	•Power sources •Interconnects
OCTOBER*	*ANALOG & MIXED-SIGNAL DESIGN CONFERENCE SHOWGUIDE***				
Analog & Mixed-Signal Design Conference	Designing mixed digital/RF systems	•Hardware/software trade–offs in multiprocessing •Network interfaces and interface ICs	Logic analyzers	Simulation	
NOVEMBER	*SPECIAL REPORT ON FUTURE COMPUTING: The merging of computers and communications*				
Comdex	Designing for testability and manufacturability	•PLD design tools •Mezzanine buses	High–resolution A–D converters	Communications	•Mass storage (disk, tape, CD-ROM, memory cards)
DECEMBER*	Designing the next generation of portable computers	•Full–system simulation •FAX/modem ICs	VME CPU Boards	Supercomputers	

**Starch Readership Research Issue* ***Contact Computer Design for circulation, rates and closing dates*

PENNWELL PUBLISHING COMPANY • ONE TECHNOLOGY PARK DRIVE • WESTFORD, MA 01886 • 508-692-0700 • Fax 508-692-7780

For more information, contact any of us: PUBLISHER: David L. Allen, 508-392-2111 • ASSOCIATE PUBLISHER Tim L. Tobeck, 508-392-2116 • NATIONAL ACCOUNTS MANAGERS: • WESTFORD, MA, Tim Pritchard, 508-392-2217 • SPRING, TX, Eric Jeter, 713-353-0309 Fax 713-288-8350 • CAMPBELL, CA, Diane Palermo, 408-371-7551 Fax 408-371-6239 • GLEN ELLYN, IL, David Ginter, 708-469-2388 Fax 708-469-2292• STRATEGIC ACCOUNTS REPRESENTATIVES: • WESTFORD, MA, Sue Nawoichik, Sales Manager, 508-392-2118 800-223-4259 • David Singer, 508-392-2109 800-223-4259 • INTERNATIONAL SALES OFFICES: • SURREY, ENGLAND, David Round, 81-770-1100 Fax 81-770-9779 • PARIS, FRANCE, Daniel R.Bernard, 1-39-14-6780 Fax 01-39-14-7014 • MUNICH, GERMANY, Johann Bylek, 089-903-8806 Fax 89-904-3526 • MILAN, ITALY, Luigi Rancati, 2-70-30-0088 Fax 2-70-30-0074 • TOKYO, JAPAN, Toshio Equsa, 81-33-536-5404 Fax 03-536-5490 • ASIA, AUSTRALIA, NEW ZEALAND, Tom Gorman, 852-833-2181 Fax 834-5620

COMPUTER DESIGN Technology and Design Directions

A PennWell Publication

ADVERTISERS INDEX

Computer Design's INSTANT DATA ACCESS ADVERTISERS

To receive data sheets instantly from these advertisers via FAX, **dial (617) 494-8338** and when prompted, enter the 4 digit document number shown here.

ADVERTISER	Page NO.	IDA NO
Brooktree	45	1026
Dynatem	71	1032
Macrolink	106	1037
NEC Technologies	125	1055

ADVERTISER	PAGE NO.	CIRCLE NO.	**IDA NO.
A			
Advanced Micro Devices	25-27	19	
Altera	78-79	48	
AMP, Inc.	95-98	—	
Analog Devices, Inc.	28-29	20	
Ariel Corporation	44	32	
ATMEL Corporation	8-9	6	
Avanti Circuits, Inc.	139	171	
B			
Bancomm	24	18	
Brooktree Corporation	37	—	
Brooktree Corporation	45	—	1026
BUSCON '93 West	108	56	
C			
Chrislin Industries	144		
Ciprico, Inc.	CV2	1	
Compass Design Automation	69	44	
Compcontrol, Inc.	141	184	
Computer Design	142	78	
Computer Design	73	46	
† Computer Design	19	14	
* Computer Design	143	80	
Connor-Winfield Corp.	140	179	
Connor-Winfield Corp.	139	173	
CSPI	64	42	
Cypress Semiconductor	CV4	—	
D			
Diversified Technology	42-43	27	
Dynatem, Inc.	71	45	1032
E			
Electronic Solutions	1	2	
ELMA Electronics	30	21	
ELMA Electronics	84	51	
Embedded Computer Conference	112-113	60	
Eyring Corporation	122	63	
F			
Force Computers	16-17	10	
G			
Grammar Engine	140	180	
Greenspring Computers	107	55	
Greenspring Computers	109	57	
Greenspring Computers	111	59	
H			
Heurikon Corporation	127	65	
Heurikon Corporation	129	67	
Heurikon Corporation	131	69	
* Hitachi America, Ltd.	19	13	
HM Systems	120-121	62	
Huntsville Microsystems	4	4	
I			
Imagine That	139	169	
Integrated Device Technology	11	7	
Integrated Systems, Inc.	47	33	
† Intel Corporation	143	79	
Invitational Computer Conference	110	58	
Ironics, Inc.	49	30	
ISO 9000	135-138	71	
K			
KADAK	18	11	
L			
Lattice Semiconductor	51	29	
Logical Devices, Inc.	141	186	
M			
Macrolink, Inc.	106	54	1037
Matrix Corporation	74	47	
Maxim Integrated Products	133	70	
Mercury Computer Systems	13	8	
Micro Alliance	141	185	
Micro-Aide, Inc.	122	81	
Micro-Sys	139	172	
Micron Semiconductor	41	26	
Microtec Research, Inc.	67	73	
Microware Sysetms Inc.	35	24	
Mitsubishi Electronics/EDG	54	35	
MIZAR	130	68	
Motorola Semicond Products	20-21	15	
Motorola TSD Group	57	38	
N			
NEC Electronics	90	52	
NEC Technologies	124-125	64	1055
Needham's Electronics	140	175	
Nohau Corporation	139	167	
O			
Oasys, Inc.	53	34	
OrCAD Systems	139	166	
OrCAD Systems	63	41	
P			
Pentek	50	28	
Phase Three Logic	140	181	
PRC, Inc.	140	182	
Q			
Quantum Software Systems, Ltd.	55	36	
R			
RadiSys Corporation	CV3	82	
Radstone Technology Corp.	60-61	40	
RISC '93 Overview	39	25	
Robotrol Corporation	140	176	
RTMX, Inc.	24	17	
S			
Sheldahl	140	177	
Sky Computers, Inc.	104	53	
Spectrum Signal Processing	14-15	9	
Synergy Microsystems, Inc.	2	3	
Synergy Microsystems, Inc.	23	16	
T			
Tadpole Technology	59	39	
Technical Solutions	140	178	
Teknor Microsystems, Inc.	83	50	
Tektronix, Inc.	6-7	5	
† Texas Instruments	81	74	
† Texas Instruments	33	22-23	
* Texas Instruments	33	—	
* Texas Instruments	81	—	
Toronto MicroElectronics	18	12	
U			
Ultraview	139	170	
V			
VenturCom	56	37	
VMetro, Inc.	48	31	
Vail Silicon Tools, Inc.	139	168	
W			
Wind River Systems	65	43	
Z			
Z-World Engineering	139	174	
Ziatech Corporation	114	61	
Zwick Systems, Inc.	140	183	

* Domestic only

† International only

** These advertisers have made arrangements to provide you with instant product data via Computer Design's Instant Data Access (IDA) System.

The Advertisers Index is published as a service. The publisher does not assume liability for errors or omissions.

COMPUTER DESIGN
Technology and Design Directions

Executive Office: One Technology Park Dr, PO Box 990, Westford, MA 01886
(508) 692-0700 Fax: (508) 692-7780 Telex: 883436
Publisher: David L. Allen (508) 392-2111
Associate Publisher/Editor-in-Chief: John C. Miklosz (508) 392-2114
Associate Publisher/National Sales Manager: Tim L. Tobeck (508) 392-2116
Administrative Assistant: Peg Alexander (508) 392-2112
Marketing Communications Manager: Betsy Anderson (508) 392-2209
Ad Traffic Manager: Kelly Rice (508) 392-2198
Recruitment Advertising, Postcards/System Showcase: Sue Shorrock (508) 392-2185
Circulation Director: Paul Westervelt (918) 832-9287
List Rental: Bob Dromgoole (918) 832-9213
Reprints: June Bozarth (918) 835-3161

SALES OFFICES

Associate Publisher
Tim L. Tobeck

National Accounts Managers

Eastern Region
Tim Pritchard
One Technology Park Dr
PO Box 990
Westford, MA 01886
Tel: (508) 392-2217
Fax: (508) 692-7780
Telex: 883436

Southwestern/ Northwestern Region
Eric Jeter
19627 I-45 N, Suite 110
Spring, TX 77388-6030
Tel: (713) 353-0309
Fax: (713) 288-8350

Northern California
Diane Palermo
910 Campisi Way
Suite 1B
Campbell, CA 95008
Tel: (408) 371-7551
Fax: (408) 371-6329

Midwestern Region
David Ginter
800 Roosevelt Rd
Suite E-120
Glen Ellyn, IL 60137
Tel: (708) 469-2388
Fax: (708) 469-2292

Strategic Accounts Representatives

Eastern U.S.
Sue Nawoichik, Sales Manager
One Technology Park Dr
PO Box 990
Westford, MA 01886
Tel: (508) 392-2118
Fax: (508) 692-7780

Western U.S.
David Singer
One Technology Park Dr
PO Box 990
Westford, MA 01886
Tel: (508) 392-2109
Fax: (508) 692-7780

U.K./Scandinavia
David Round
Westmead House
123 Westmead Road
Sutton
Surrey SM1 4JH
Tel: (081) 770-1100
Fax: (081) 770-9779

France/Belgium/ Southern Switzerland/ Spain/Netherlands
Daniel R. Bernard
10, rue Michelet - B.P. 279
78502 Sartrouville Cedex, France
Tel: 01-39-14-67-80
Fax: 01-39-14-70-14

Germany/Austria/Northern Switzerland/Eastern Europe
Johann Bylek
Verlagsbuero Johann Bylek
Stockaeckerring 63
D-8011 Kirchheim/Muenchen
Federal Republic of Germany
Tel: 89-903-88-06
Fax: 89-904-35-26

Italy
Luigi Rancati
Rancati Advertising
Milano San Felice Torre 5
20090 Segrate, Italy
Tel: 02-70300088
Telex: 328601 RANCAD I
Fax: 39-270300074

Japan
Toshio Egusa
PubliNetwork
C407, 2-22-6
Tsukuda Chuo-ku
Tokyo 104, Japan
Tel: 03-3536-5404
Fax: 03-3536-5490

Southeast Asia Taiwan/Hong Kong/Korea/ China/Thailand/Philippines/ India/Pakistan/Singapore/ Malaysia/Indonesia/Australia/ New Zealand
Tom Gorman
CCI Asia-Pacific, Ltd
Suite 905, Guardian House
32 Oi Kwan Rd, Happy Valley
Hong Kong
Tel: 833-2181
Fax: 834-5620

COMPUTER DESIGN (ISSN-0010-4566) is published monthly by PennWell Publishing Company, 1421 S Sheridan, Tulsa, OK 74112. Second-class postage paid at Tulsa, OK 74112 and additional mailing offices. Editorial offices are located at One Technology Park Dr, PO Box 990, Westford, MA 01886. Subscription Prices: Free to design and development qualified engineers and engineering managers in the U.S. and Canada. Paid to all others. Qualified engineers and engineering managers outside the U.S. and Canada - air shipped - $85 one year. For non-qualified recipients in the U.S. - $88 one year; in all other countries - air shipped - $154 one year. Call (918) 832-9263 for subscription information. Microfilm copies of COMPUTER DESIGN may be purchased from University Microfilms, a Xerox Company, 300 N Zeeb Rd, Ann Arbor, MI 48106. POSTMASTER: Send change of address form to COMPUTER DESIGN, Circulation Department, PO Box 3466, Tulsa, OK 74101.

Officers of PennWell Publishing Company: Frank T. Lauinger, Chairman and Chief Executive; Joseph A. Wolking, President and Chief Executive Officer; John Ford, Senior Vice-President; Carl J. Lawrence, Senior Vice-President; V. John Maney, Vice-President/Finance; W. Steve Zimmerman, Vice-President/Corporate Services.

RADISYS GIVES VME THE INSIDE ADVANTAGE

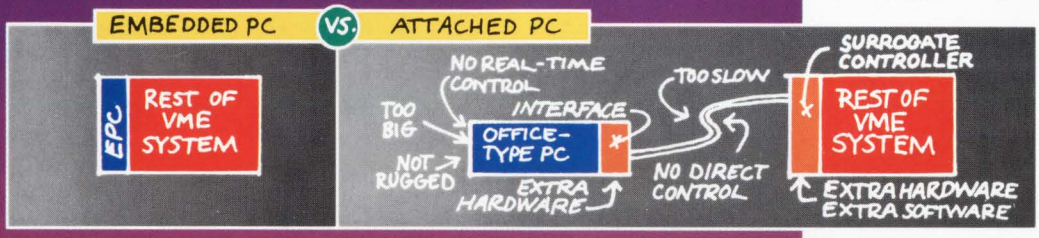

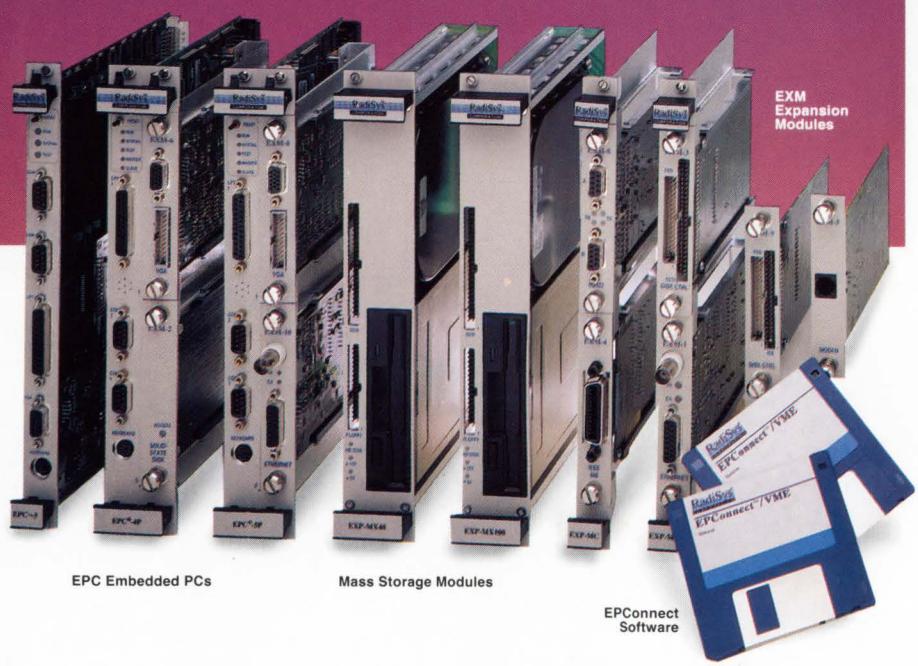

EPC Embedded PCs | Mass Storage Modules | EPConnect Software

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

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

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

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

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

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

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

intel inside™

RadiSys® CORPORATION

THE INSIDE ADVANTAGE

RadiSys Corporation
15025 SW Koll Parkway
Beaverton, OR 97006 USA
(800) 950-0044
(503) 646-1800
Fax (503) 646-1850

CIRCLE NO. 82

TO BREAK SPEED RECORDS, TURN TO THE SRAM LEADER.

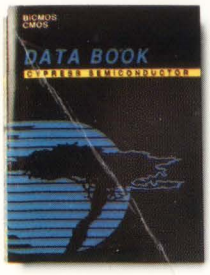

To break performance records, you need extraordinarily fast SRAM solutions. According to In-Stat, Inc., no U.S. company sells more SRAM parts worldwide than Cypress. Period.

If you need high-performance SRAMs, this broad range of fast memories is at your service. CMOS or BiCMOS, TTL or ECL, from 64-bit to multi-megabit, multichip modules, from standard SRAM pinouts to specialty dual-ports, FIFOs, or Cache SRAMs, we've got your high-speed SRAM solution. Our Data Book will show you. It's a hit. And it's free.

FREE 1992 DATA BOOK HOTLINE: 1-800-852-1810*
Ask for dept C115.

CYPRESS SEMICONDUCTOR

*In Europe, fax your request to the above dept. at 1 (415) 961-4201 or call (32) 2-652-0270. In Asia, fax to the above dept. at 1 (415) 961-4201. © 1992 Cypress Semiconductor, 3901 North First Street, San Jose, CA 95134. Phone 1 (408) 943-2600, Telex: 821032 CYPRESS SNJ UD, TWX: 910-997-0753.
Statistic from In-Stat, Incorporated's 1991 Worldwide Merchant Market Dollar Shipments data.