A PennWell Publication

COMPUTER DESIGN

THE MAGAZINE OF COMPUTER BASED SYSTEMS

SEPTEMBER 1984

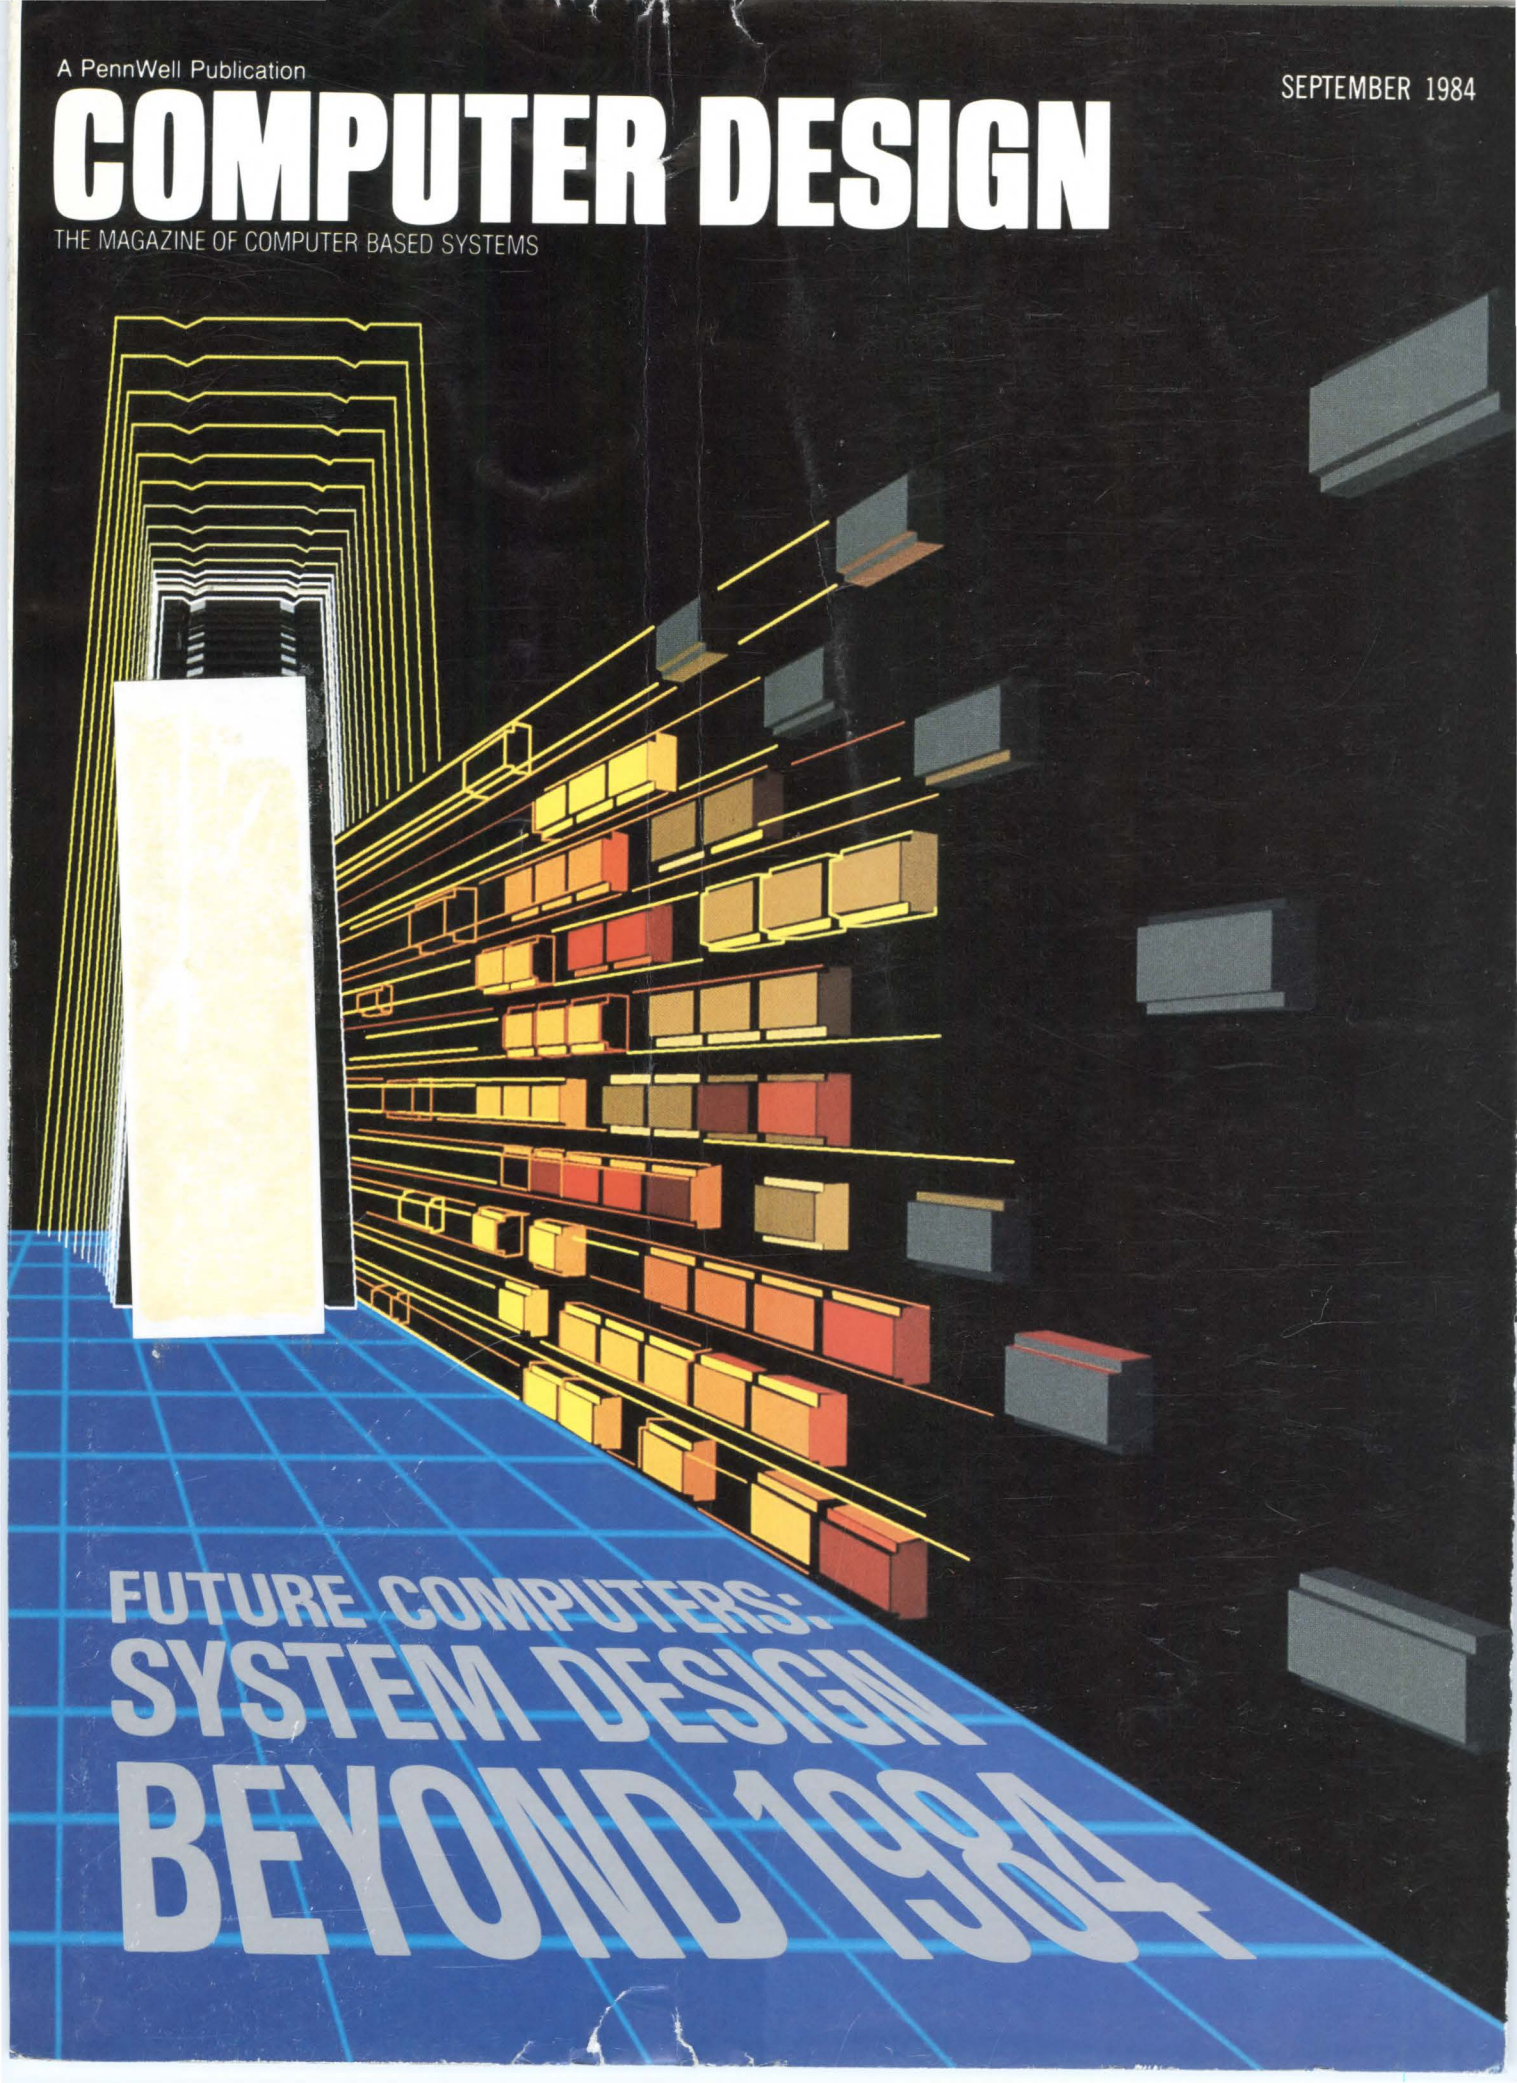

FUTURE COMPUTERS: SYSTEM DESIGN BEYOND 1984

AS BRIGHT AS YOU ARE?

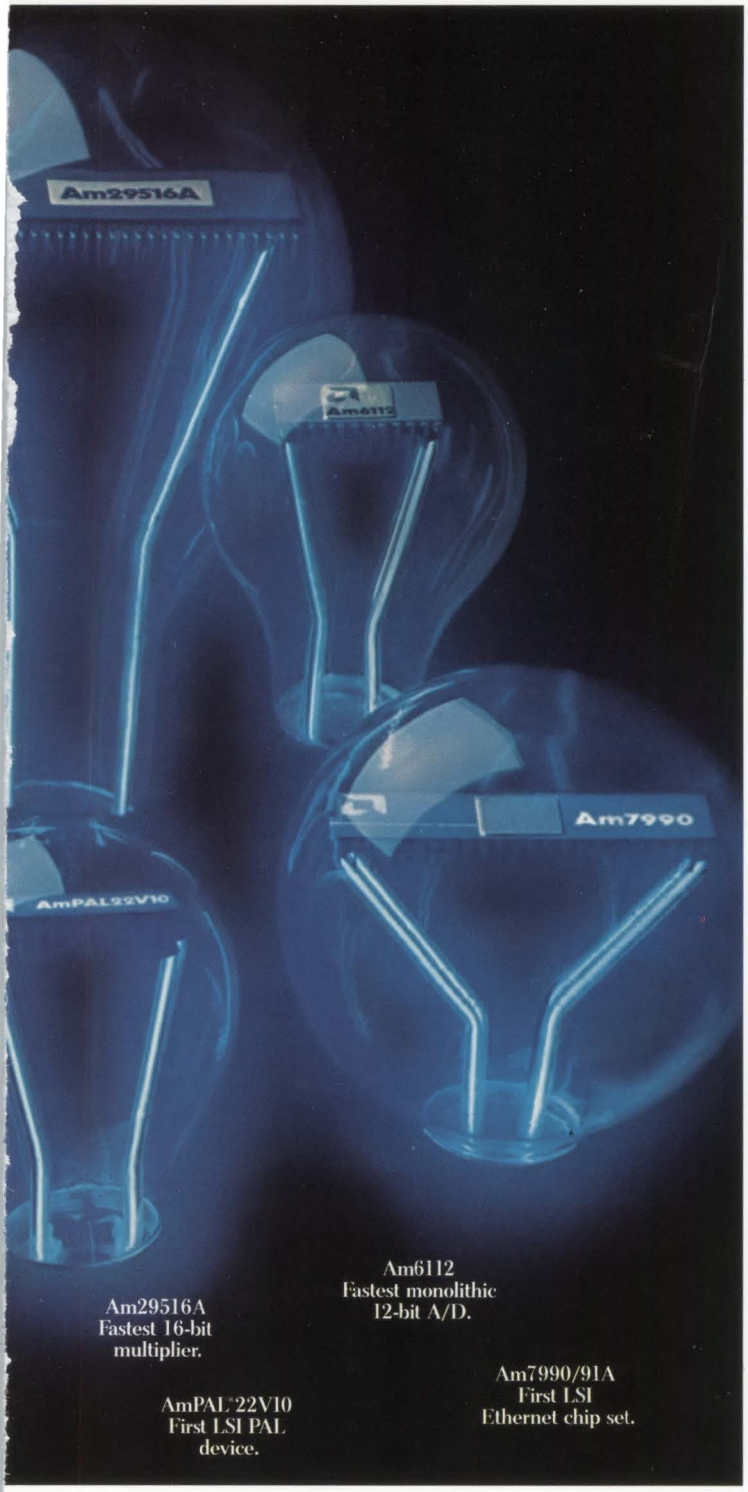

You're doing everything you can to stay ahead of the competition. Shouldn't your IC company do the same?

We think so.

In 1983 we spent a record-breaking 18.7% of sales on research and development. That's more than any other major IC company.

And that's why 40% of our total sales come from products that were invented here.

Our Am8052 CRT controller lets you squeeze every last drop of performance out of a video tube.

And our Am7910 is the only modem you can program for any major telephone system anywhere in the world.

We make the first perfectly matched VLSI Ethernet chip set and the only complete kit solution to super high speed digital signal processing.

We make the world's first 512K EPROM and the world's fastest microprocessor.

We're even an innovator when it comes to quality.

While other guarantees run on and on, ours is short and sweet.

The International Standard of Quality guarantees a 0.1% AQL on all electrical parameters, AC and DC, over the entire operating range.

INT·STD·1000

So, if you've got a product you want to really shine, call us. We can put you as far ahead of the competition as we are.

Advanced Micro Devices

901 Thompson Place, P.O. Box 3453, Sunnyvale, CA 94088
(408) 749-5000, outside California, call toll free (800) 538-8450, ext. 5000.

CIRCLE 1

IS YOUR CHIP MAKER

Am8052
Most advanced
CRT Controller.

Am7910
First multi-spec
single chip modem.

Am29540
First bipolar
LSI FFT processor.

Am7990/91A
First LSI
Ethernet chip set.

Am29818
First IC
with diagnostics.

Am7901
First programmable
CODEC/Filter.

Am29116
Fastest 16-bit
microprocessor.

Am27S43A
Fastest 32K
bipolar PROM.

Am9150
First clearable
static RAM.

Am27512
First
512K EPROM.

GOOD IDEAS ARE HARD TO COME BY.

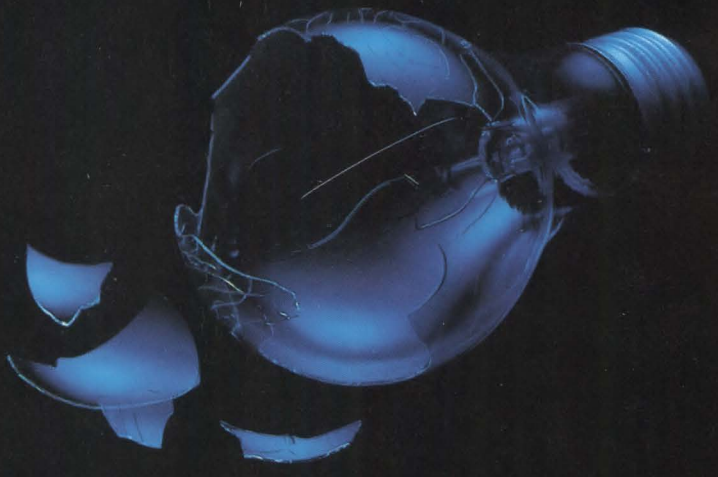

Introducing Model 6500 —

'the handful'

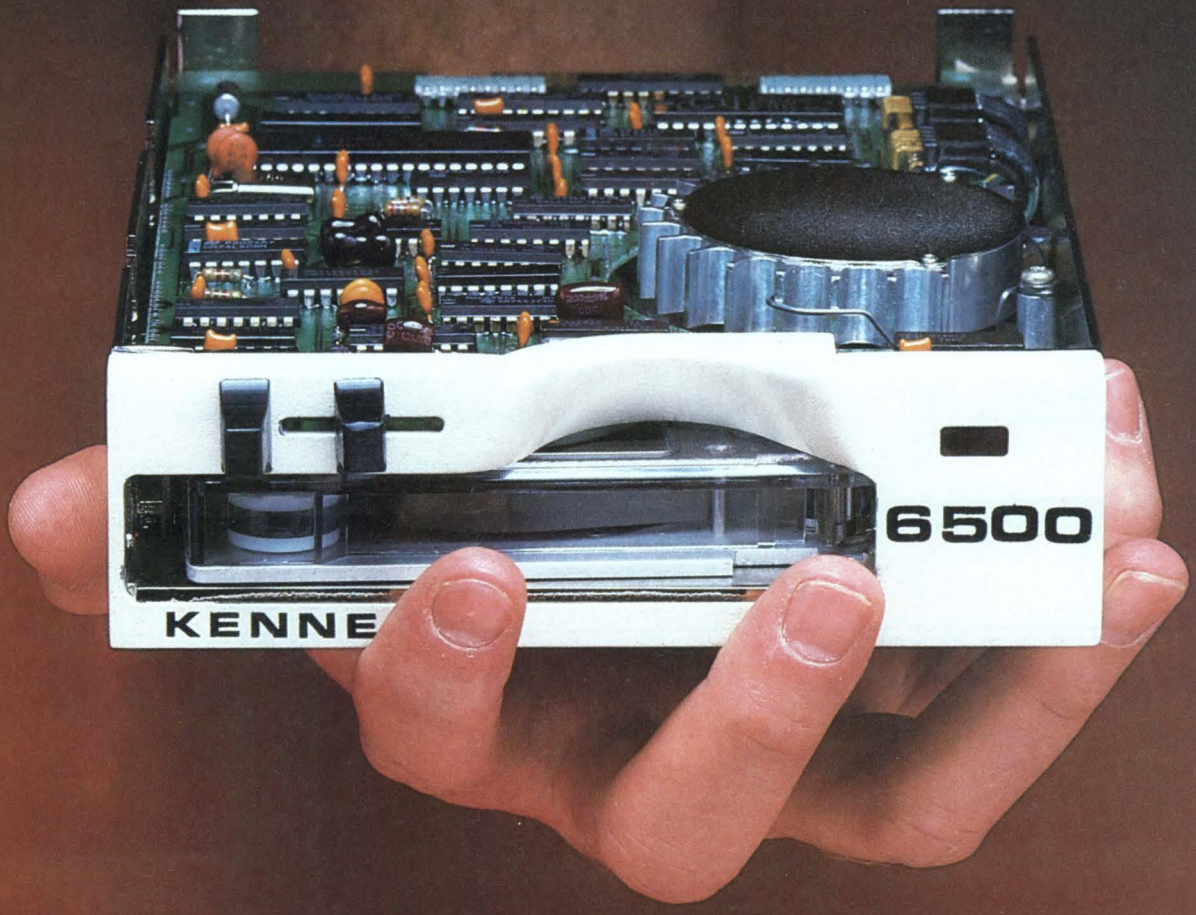

Kennedy's Model 6500 5¼" Cartridge Recorder — it fits in your hand, but it's more than a handful of new and unique features, such as:

- 16 KBYTE Buffer — Model 6500 can transfer data at a burst rate of up to 3 MBYTE/sec., and maintain streaming longer in systems with little buffering.
- 150 MS Ramp Time — Data access time and reposition time is cut in half; and it can better operate in a Pseudo s/s mode.
- Direct Drive Motor — Assures much higher reliability and lower I.S.V.

Finally, and best of all, Model 6500 features a simple mechanical design — loading is simple, moving the tape to the head instead of the head to the tape.

Model 6500 is one big handful of what you've been looking for in a 5¼" cartridge recorder. Write or call us.

KENNEDY
An Allegheny International Company
1600 Shamrock Ave., Monrovia, CA 91016
(818) 357-8831 • ITT TELEX 472-0116 KENNEDY

KENNEDY • QUALITY • COUNT ON IT

CIRCLE 2

COMPUTER DESIGN®

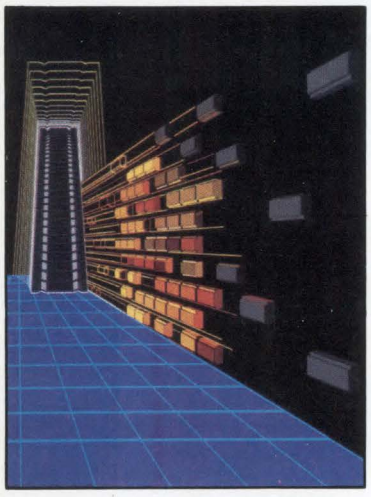

This month's cover was designed at Coddbarrett Associates, Inc, by Mary Codd and Steve Branch. It was illustrated by Mary Codd, using a DICOMED D-38 design station and D-48 high resolution film recorder.

SPECIAL ISSUE: FUTURE COMPUTERS

In the fast-paced and competitive computer industry, ongoing events are already startling enough to induce future shock. Yet, hard as it may be, today's engineers and engineering managers must forecast technology as far ahead as possible to capitalize on new developments ahead of worldwide competitors. Therefore, in this issue, the editors of *Computer Design* peer into the research labs and consult an army of experts to bring you the clearest possible picture of trends that will shape system design beyond 1984. However, we don't believe it is feasible, in a single issue, to present a complete picture of a technology whose potential is almost limitless. So, this is just the opening salvo of a continuing barrage of forward-looking information you will need to win the battle for technological survival.

SYSTEM TECHNOLOGY

AUTOFACT 6, ISA/84, IECON '84

COMPUTER DESIGN©1984 (ISSN-0010-4566) is published monthly, with a thirteenth and fourteenth issue respectively in June and October by PennWell Publishing Company, Advanced Technology Group, 119 Russell Street, Littleton, MA 01460. Second-class postage paid at Littleton, MA 01460 and additional mailing offices. COMPUTER DESIGN is distributed without charge to U.S. and W. Europe-based engineers and engineering managers responsible for computer-based equipment and systems design. Subscription rate for others is $50 in U.S.A. and $75 elsewhere. Single copy price is $5.00 in U.S.A. and $7.50 elsewhere. Microfilm copies of COMPUTER DESIGN are available and may be purchased from University Microfilms, a Xerox Company, 300 North Zeeb Road, Ann Arbor, Michigan 48106. POSTMASTER: CHANGE OF ADDRESS-FORM 3579 to be sent to COMPUTER DESIGN, Circulation Department, P.O. Box 593, Littleton, MA 01460 (USPS 127-340).

Officers of PennWell Publishing Company, 1421 S. Sheridan, Tulsa, OK 74101: P. C. Lauinger, Chairman; Philip C. Lauinger, Jr., President; Joseph A. Wolking, Senior Vice President; H. Mason Fackert, Group Vice President; Carl J. Lawrence, Group Vice President; V. John Maney, Vice President/Finance; L. John Ford, Vice President.

ABP BPA

Page 137

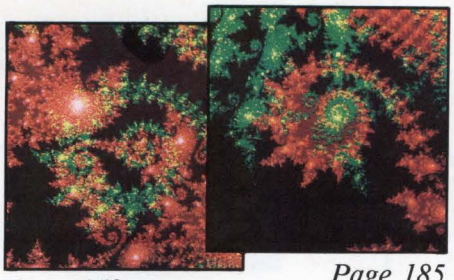

Page 169

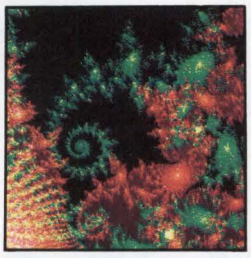

Page 185

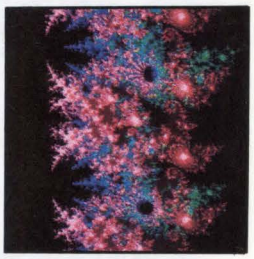

Page 197

Page 231

The photos on the lead page of each contributed article represent an iteration of cubic polynomials in the complex plane. They were contributed by mathematician Dr John Hubbard and associate Homer Smith of Cornell University, Ithaca, NY.

SYSTEM COMPONENTS

DEPARTMENTS

METACOMP LAUNCHES

COMM-2000™

A Major Breakthrough in Multibus Data Communications

BREAKING THE COMMUNICATIONS BARRIER

The COMM-2000 family is a powerful assortment of software and firmware products that perfectly integrate with METACOMP's MPA-2000™ CPU/Intelligent I/O Controller to provide the widest possible range of data communications protocol and networking support available to Multibus system integrators.

COMM-2000 plus MPA-2000 equals the complete hardware, software, and firmware solution to all your Multibus data communications problems.

Comm-2000 offers packet-switching, terminal I/O concentrator, network gateway emulation, and DLC protocol support.

Available capabilities include:

- X.25 with full CCITT 1980 recommended features
- IBM 3270 SNA protocols
- SDLC, HDLC (LAPB), BISYNC, and other DLCs
- LAN (XNS, TCP/IP, and others)

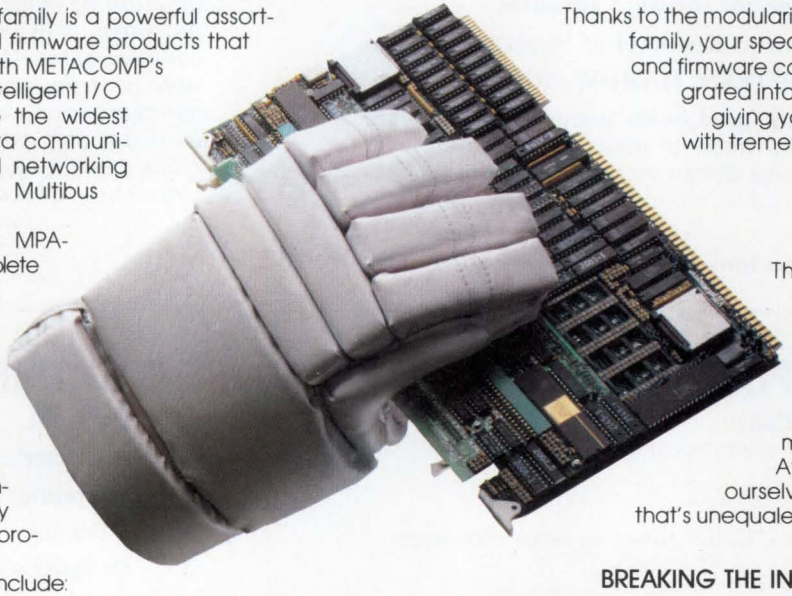

BREAKING THE TIME BARRIER

With the COMM-2000 family at your disposal, you can design Multibus systems that provide standard data communications facilities in the shortest possible time and with the least possible hassle. Because COMM-2000 combinations are endless . . . easy, and readily available.

BREAKING THE SPACE BARRIER

Thanks to the modularity of the COMM-2000 family, your special hardware, software, and firmware combination can be integrated into a single board solution, giving you the most capabilities with tremendous savings in system capacity.

BREAKING THE SERVICE BARRIER

The ultimate ingredient of COMM-2000 is commitment. You can absolutely rely on METACOMP for service and customer support guaranteed to meet your requirements. At METACOMP, we pride ourselves on customer service that's unequaled by any other Multibus manufacturer.

BREAKING THE INFORMATION BARRIER

Here's an easy breakthrough you can make today. Write or call for complete details on COMM-2000. We will respond with a package of information that will promote your data communications breakthrough tomorrow. METACOMP, INC., 9466 Black Mountain Road, San Diego, California 92126, 619/578-9840, TWX 910-335-1736.

METACOMP™

THE MULTIBUS BREAKTHROUGH PEOPLE

CIRCLE 3

UP FRONT

Personal computer will keep its electronic mouth shut

IBM has started taking orders for yet another version of its personal computer. This one is geared to the needs of military system and network designers. However, it is not being produced at the IBM Boca Raton, Fla operation that turns out PCs and XTs by the carload for the commercial marketplace. Instead, it is being made at IBM's Federal Systems Division and conforms to the U.S. Department of Defense Tempest specification. Tempest defines, among other things, radio frequency interference security requirements. For example, if a scanning receiver-computer eavesdrops outside a building containing Tempest-specified XTs, it would not be able to pick up enough XT-generated electromagnetic radiation to be able to "read" (after signal processing) what the XTs are doing.—*H.H.*

Intel and AMD move to ensure availability of 80286

One potential stumbling block for the new PC AT—and the expected AT compatibles—may be removed. For the IBM PC, the lack of available 8088 microprocessor chips has been a problem. The microprocessor for the PC AT is the 80286, developed by Intel Corp (Santa Clara, Calif). However, any possible supply problem can be fixed this time around by an agreement between Intel and Advanced Micro Devices (Sunnyvale, Calif). In this arrangement, Intel will transfer its 80286 manufacturing package to AMD as a second source. AMD expects to be in pilot production by first quarter, 1985. As might be expected, volume production is scheduled for the following quarter.—*J.V.*

Personal computers to function as digital signal processors

Now, the IBM PC can serve as a digital signal processor. Ariel Corp (New York, NY) has introduced a plug-in, dedicated fast Fourier transform (FFT) processor card for the PC. (A similar card is available for Hewlett-Packard series 200 computers.) Space and power requirements for the host preclude its use as a general-purpose array processor. However, the boards provide a 16-bit integer, 1024-point FFT in 9.2 ms. Fewer points take a shorter time, and the inverse FFT algorithm is also available. Power spectral density and Hamming window algorithms are standard. Others, such as Kaiser, Bartlett, Hanning, and custom algorithms, are available. FFT board usage requires no FFT knowledge; a single-program line in one of the PC languages is sufficient to call up the program by assembly language subroutines. To the user, it looks like a software-driven FFT is being performed—only much faster. Since algorithms are held in PROM-based microcode, no algorithm knowledge is required. Disk-based, machine-language, user-modifiable driver routines provided with the boards guarantee a host/peripheral interface that is transparent to the user.—*H.H.*

UP FRONT

Board-level computer I/O spec set to debut

Designers are used to adding I/O to their board-level systems through backplane buses. But, if the IEEE P959 I/O Expansion Bus Standard (based on the so-called P959 bus) is approved, it will be possible to add small I/O increments directly to a single-board computer in a well-defined manner. This standard allows small expansion boards to plug directly into the main computer board through a pin-and-socket connector. The new I/O expansion technique features low cost, minimal space requirement, and compatibility with existing I/O standards. The standard defines the logical, electrical, and mechanical aspects of the connection. For example, it requires that all P959 bus signals be derived from the address and control buses of the single-board computer. The P959 standard is now available for informal public comment before official submission to the IEEE Standards Board. Those interested in the complete specification should contact P959 Committee Chairman Leon Adams at Texas Instruments in Houston, Tex.—*H.H.*

Products support seven-layer OSI model

Local area networks have been on the scene for some time. But now, they are taking the form of more mature, second-generation products. Witness the latest introductions of three-year old LAN firm, Interlan, (Westford, Mass). Interlan's NP series of intelligent Ethernet processors can support all seven layers of the ISO/OSI communication protocol model. Further, they can off-load protocol processing activity that is burdensome to the host system. Versions aimed at Unibus, Q-bus, and Multibus systems now run the Xerox Network Systems (XNS) Internet Transport Protocols (ITP) package. Designers can still configure custom transport and application-level protocols. Interlan has also announced a plug-in Ethernet controller package that links PC users to distributed applications on an Ethernet.—*J.V.*

Wafer scale integration—another approach

There has been considerable interest lately in high density circuit packaging techniques. Much of this interest has been spawned by the publicity given Trilogy Systems' ill-fated effort to leapfrog into wafer scale integration technology. Now comes Mosaic Systems, Inc (Troy, Mich) with a hybrid approach called WHIPs—for wafer hybrid interconnection packages. Says Dr Robert Johnson, company president, "Currently, Mosaic's proprietary interconnection technology is designed to provide the performance benefits of wafer level integration while utilizing off-the-shelf logic and memory chips. Further, it is fully applicable to integrated, monolithic wafers." Johnson goes on to say that the company's cooling and packaging approach solves problems such as head dissipation, encountered with other high density packaging techniques. Mosaic will begin shipments of initial WHIP designs late this year. Developments downstream will include application of Mosaic's interconnection technology to monolithic wafers.—*J.H.*

Gould...Innovation and Quality in Superminicomputers

We've drawn the line on computer price/performance.

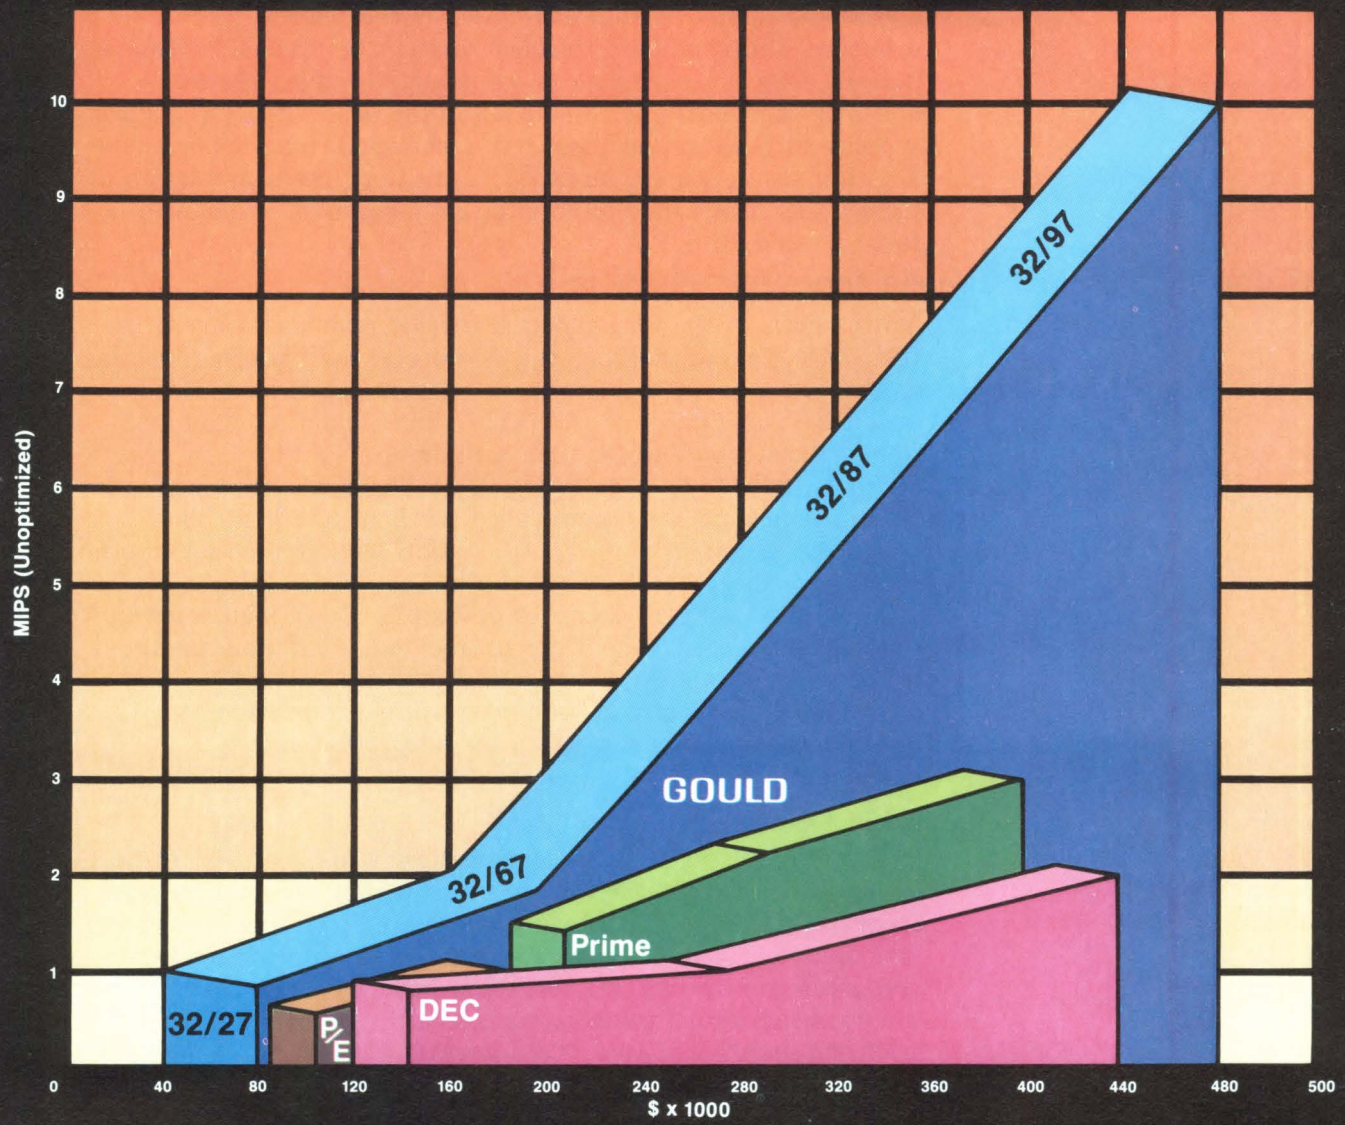

Gould has set new supermini-computer performance standards with its CONCEPT/32™ family of 32-bit machines. The cost-effective, wide-ranging capabilities of Gould minicomputers make Gould Computer Systems the dominant source for the compute power you need, at a price you can afford.

The competition just doesn't tow the line in either price or performance. Whatever the requirement. The Gould CONCEPT/32 family offers the widest range of superior performance while keeping the price in line. Our low-end CONCEPT 32/27 incorporates high density packaging for lower cost. The mid-range 32/67 combines a minimal footprint and cost with superior computational power. For heavy duty scientific and engineering applications, the Gould CONCEPT 32/8780 offers mainframe performance at a fraction of the cost. And if you're worried about where your application falls on the line, don't be. Upward compatibility and software transportability allow you to move up our line as far as you need to go.

Gould has drawn a new price/performance line. One that shows it takes more than 32-bits to make a supermini. A line the competition can't cross. Call or write for more information.

Gould Inc., Computer Systems Division
High Performance Systems Operation
6901 West Sunrise Boulevard
Ft. Lauderdale, Florida 33313
1-800-327-9716

All chart data from published competitive information.
™ CONCEPT/32 is a trademark of Gould Inc.

GOULD
Electronics

CIRCLE 4

UP FRONT

What does it take to make a standard?

A joint agreement between Cipher Data Products (San Diego, Calif) and IBM (White Plains, NY) signed in July may give Cipher the right to manufacture and market several tape products. Terms of the agreement will not be disclosed. However, industry watchers speculate on its potential impact if IBM's half-inch tape cartridge is included in the deal. Cipher is a leading supplier of both quarter- and half-inch tape drives. While "IBM does not set standards, industry groups do," the combination of Cipher with IBM would almost certainly be potent enough to force a standard.—*P.K.*

Motorola provides support tools for the 68020

On the heels of the long-expected official announcement of its full 32-bit 68020 microprocessor, Motorola, Inc (Austin, Tex) has announced a line of hardware and software system support products. These hardware tools are built around a 68020-based single-board computer that can include up to 1 Mbyte of RAM, and a 1- to 4-Mbyte RAM board. Both board-level products can fit into an evaluation system called the Benchmark 20. This system has a total of four Versabus slots to accommodate more memory or disk controllers. It can be used with Motorola Exormacs or CME/10 systems to develop or down-load software. In addition, there will soon be an in-circuit emulation module and cross development software to run in a DEC VAX environment. This cross C compiler will generate 68020 code to down-load to the Benchmark 20 for integration and debugging.—*T.W.*

Unix System V kernel speeds computer system I/O

Software developers writing Unix System V applications for host computers can speed their applications if they off-load the host from I/O chores so that it can spend more of its time calculating. One approach to the off-loading chore is the use of intelligent communication peripheral cards that have a Unix System V-compatible kernel ported onto the card's controlling microprocessor. For Intel microprocessor-based peripheral cards, Network Consulting, Inc (Burnaby, British Columbia, Canada) has come up with the first System V-compatible Unix kernel. The kernel allows IBM XTs to serve as process control or database machines with minimal I/O restraint while they handle Unix, Coherent, or any other Unix-compatible operating system. For example, tests have been conducted with the Persyst intelligent communication card (with an onboard NCI kernel) for serial line driving. These tests show that four terminal-to-XT users can keep going at their maximum 19.2-kbit/s throughput rate. This rate represents only about 4 percent of the host time usually needed. Software "hooks" (pointers) now under development will allow parts of host application programs to run off the kernel and speed matters up even more.—*H.H.*

Crack the Nut.

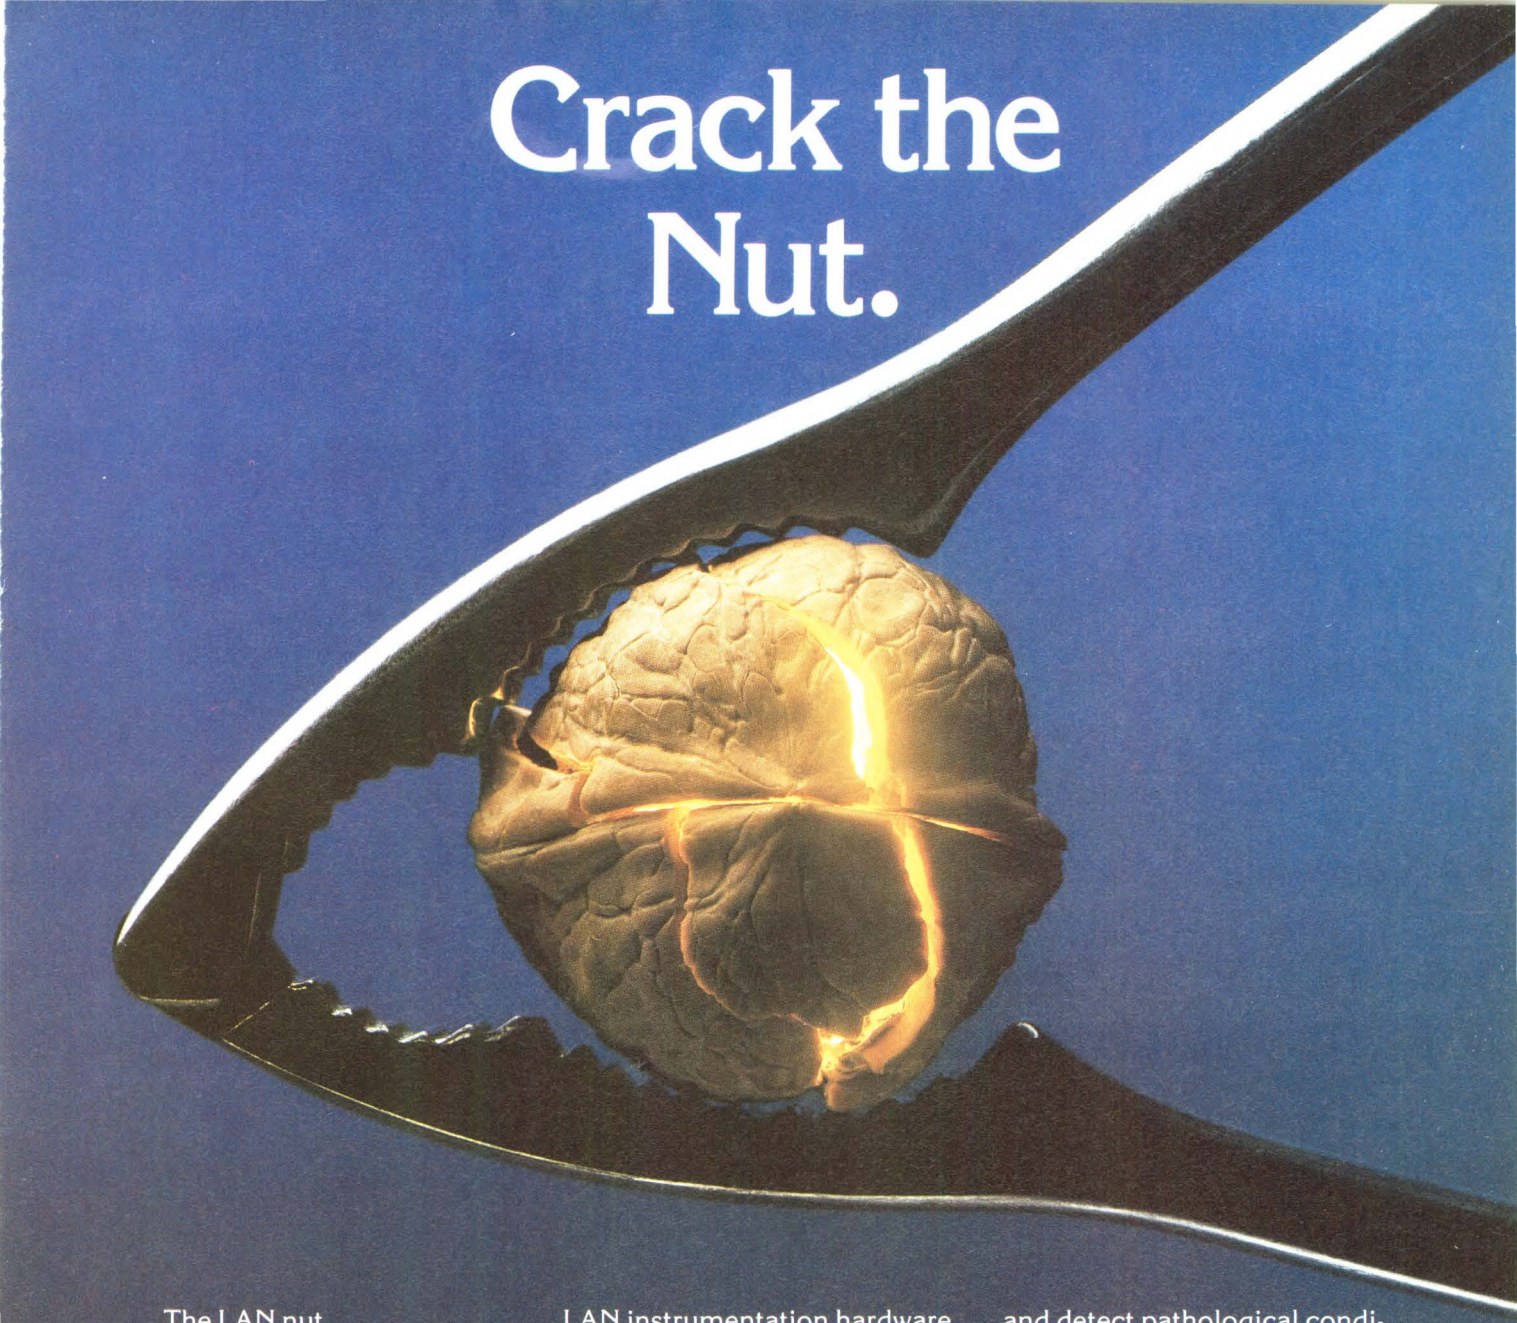

The LAN nut.

When you're developing a LAN-based computer system, you're faced with a tough nut to crack. You need to know what's going on in your Local Area Network system. Many computers are carrying on many conversations. Simultaneously. And at speeds 10,000 times faster than traditional data communications. How do you test it? How do you debug it? How does the system really perform?

Excelan can help you crack that nut with the Nutcracker™, the world's only comprehensive analyzer/simulator for Ethernet systems. Packaged as an integrated workstation, the Nutcracker provides advanced LAN instrumentation hardware capable of making real-time diagnostic decisions at 10 million bits per second. Also included is an 8086-based CPU, about 1 MB of RAM, a 20 MB disk, a 600 KB floppy, keyboard, 12″ CRT and 100 cps external printer, and complete menu-driven software that brings the power of the Nutcracker to you.

With the Nutcracker connected to the Ethernet cable, you're in control. You can create and detect pathological conditions for shake-down testing. Powerful filtering facilities allow you to extract precise packet substreams for triggering and tracing functions to find those elusive bugs. You can generate traffic for simulation of various load conditions. And high resolution time-stamping means you can track system performance against this regulated traffic load.

Increase your productivity, speed development time, and get your network system up and running — with confidence in it's performance and integrity.

Join computer industry leaders. Crack the LAN nut with the Excelan Nutcracker.

EXCELAN

"Excellence in Local Network Technology"

2180 Fortune Drive
San Jose, California 95131
(408) 945-9526 TELEX 176610

CIRCLE 5

WHEN IT COMES TO WINCHESTERS, WE RUN CIRCLES AROUND THE COMPETITION.

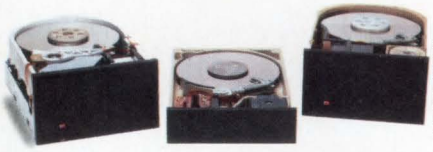

Tandon's made so many improvements in Winchesters that the rest of the drive industry is dizzy just trying to keep up.

We build a wide range of drives. Open-loop and high performance closed-loop models. In both full and half-height versions ranging from 12.8 to 36.2MB unformatted capacities.

All are available in high quantities now, because Tandon's been ramping up capacities at a pace like never before. And we've kept up the pace of improving Winchester performance as well, so the most advanced technology is ready when you need it.

As far as our current line is concerned, the low-cost TM500 line has become the most reliable full-height 12.8 and 19.1MB drives you can get. The new TM252 half-height 12.8MB surpasses the industry's most critical specifications. And our TM703 has been increased to 36.2MB to make it the ideal high capacity drive for desktop business systems.

Another part of our success revolves around our highly automated plated media plant in Northern California. Its buttoned-down efficiency and technological sophistication guarantee higher quality and higher volume production than ever.

What's more, our proprietary plated media is used in all our Winchesters, no matter what the capacity. It's more durable, more reliable. And it has six times more storage capacity than oxide media. At less than half the cost.

In fact, it's so good, other drive manufacturers are buying their plated media from us.

So all around, Tandon Winchesters have more capacity in less space at less cost.

And that's a pretty good circle to run around in.

Tandon

THE DRIVING FORCE BEHIND THE SMALL COMPUTER INDUSTRY.

Tandon Corporation, 20320 Prairie, Chatsworth, CA 91311. (818) 993-6644, TWX: 910-494-1721, Telex: 194794. Regional Sales Offices: Boston (603) 888-8612 • New York (201) 851-2322 • Atlanta (404) 934-0620 • Chicago (312) 530-7401 • Dallas (214) 423-6260 • Irvine (714) 669-9622 • Santa Clara (408) 727-4545 • Kelsterback/Frankfurt, West Germany 6107-2091, Telex: 411547 • Reading/London, England (0734) 664-676, Telex: 848411. Distributors: Hall-Mark, Kierulff, Schweber.

CIRCLE 6

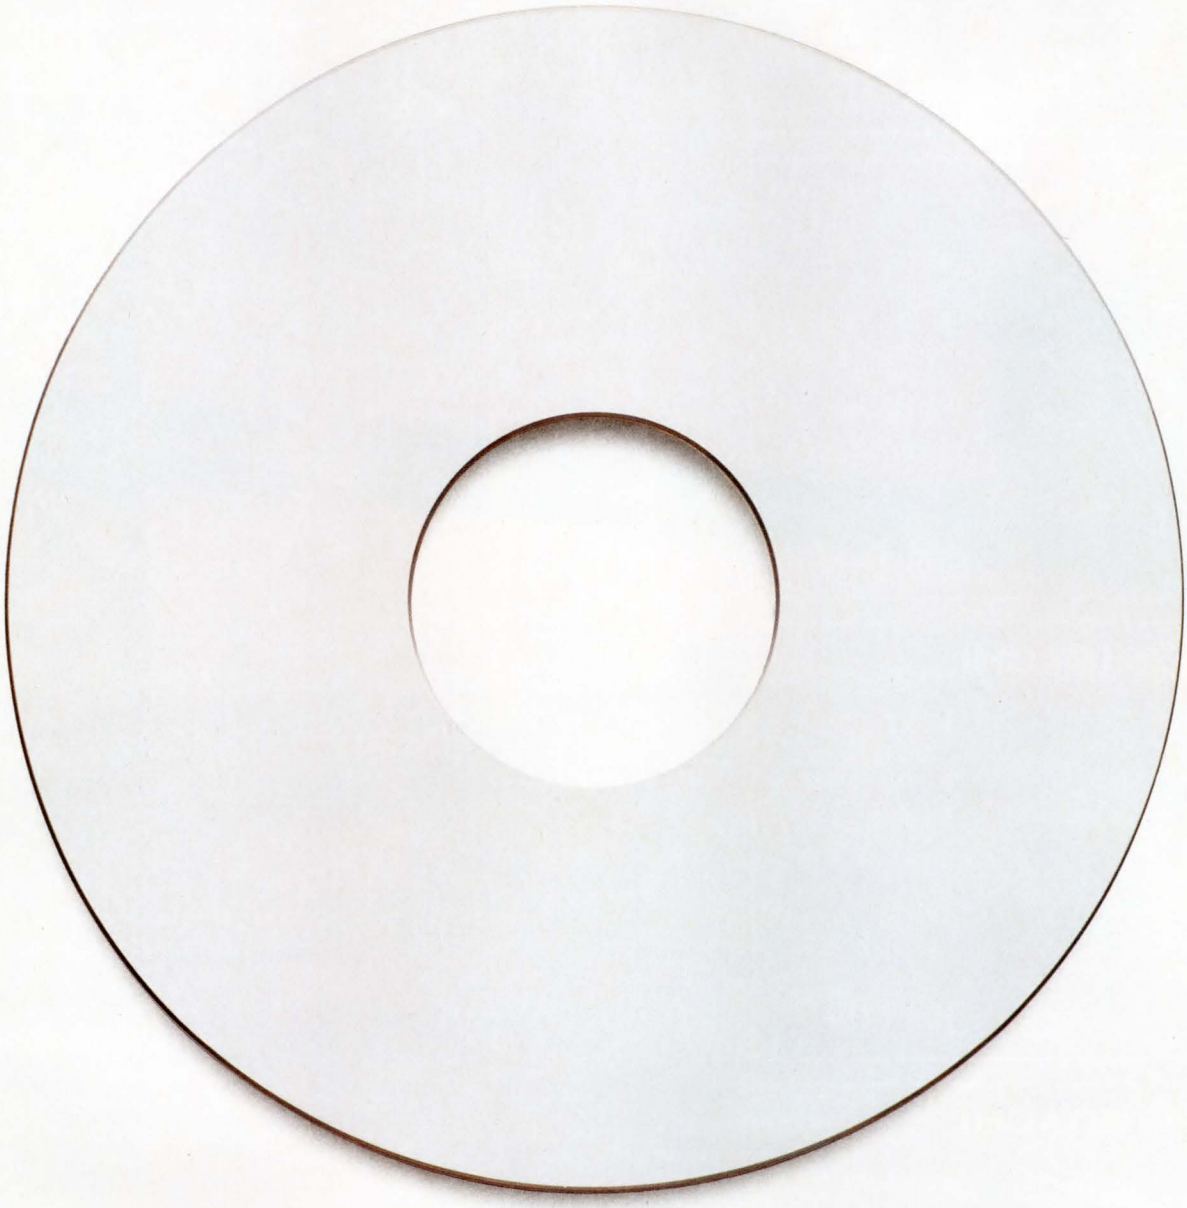

HITACHI

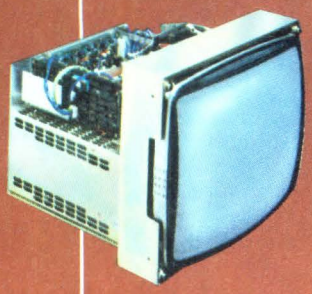

- 1280 × 1024 Resolution
- 60Hz, Non-interlace
- 100 MHz Bandwidth
- Compact, modular design

HM 4619

HITACHI

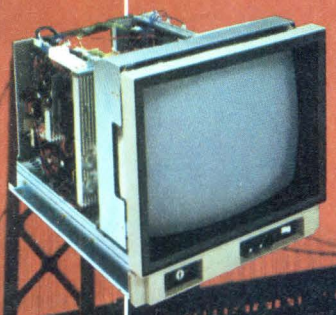

- 1280 × 1024 Resolution
- 30Hz, interlaced
- 45MHz Bandwidth
- 0.1/0.3 mm convergence

HM 3619A

HITACHI

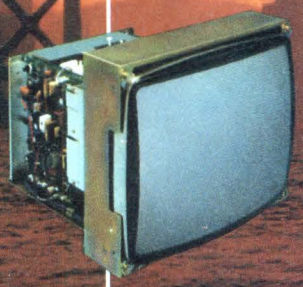

- 1024 × 1024 Resolution
- 40 MHz Bandwidth
- Compact

NEW

HM 3719

HITACHI

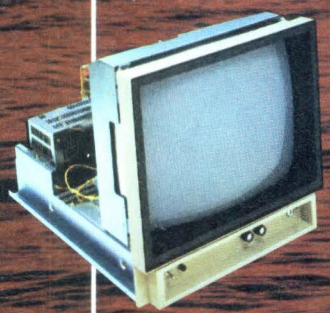

- 640 × 512 Resolution
- 25 MHz Bandwidth

HM 2719

HITACHI Bridges The Gap In OEM RGB Color Monitors

The world of computer graphics is moving at incredible speed. Software advances dictate hardware advances; new applications demand new engineering specifications. "Limited standards" don't always fit the need.

Hitachi is the totally integrated OEM source. Our complete range of RGB Color Monitors provide a very wide selection of power, resolution, size and price. More important, we can and do modify specifications to meet highly specialized individual OEM requirements.

Best of all, our "standard" monitors and our "not so standard" monitors all reflect clearly the world renowned Hitachi technology and our dedication to design excellence.

HITACHI.....the image speaks for itself.

HITACHI

The Image Speaks For Itself.

Hitachi America, Ltd.
59 Route 17S, Allendale, NJ 07401
(201) 825-8000

Hitachi America, Ltd.
3540 Arden Road, Hayward, CA 94545
(415) 783-8400

Hitachi Europe, Ltd., Lee House (10th Floor)
London Wall, London EC245AS
(01) 606-7831

Hitachi Denshi Canada, Ltd. 65 Melford Drive
Scarborough Ontario, Canada M1B 2G6
(416) 299-5900

CIRCLE 7

A NEW CALL TO ARMS

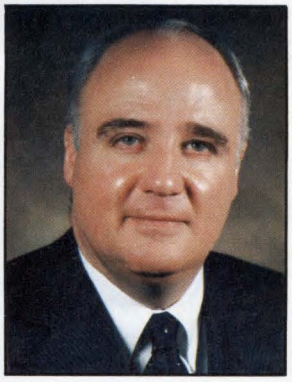

Today's computer design engineering managers are the generals in the most important battle of our time. The outcome of the struggle between Japan, Europe, and the United States for supremacy in information technology will determine the balance of world economic power as well as the security of the free world. This is why we have chosen to look even further ahead than usual in this issue of *Computer Design*.

Computers, communication equipment, and services are forecasted to account for 40 percent of world industrial value added by the year 2000. Success here will translate into widespread employment and rapid emergence of growth technologies. It will also mean significant gains in productivity and life-supporting technologies. Computers will not be merely improved—they will be transformed. They will begin to reason and to apply logic. They will be voice interactive and symbol manipulating. Their roles will change and these new roles will have a dramatic effect on our way of life.

The computers of the not very distant future will rely primarily on continuing technological progress in VLSI circuitry, high density storage, advanced communication and processing architectures, and infinitely more powerful software. Advances in design automation will help provide the needed push to shorten design and development cycles. Yet, in the final analysis, the battle will be won by the skill and dedication of computer design engineers. The responsibility for the technological breakthroughs required to win the contest for national security and economic prosperity is clearly theirs.

Currently, there is a dearth of computer design engineers in the United States. Talent turnover is high within the lopsided seller's market of the boom-or-bust computer industry. Vital military research siphons considerable numbers of engineers from the private sector. Business, in turn, attracts significant talent from the academic research and educational arenas.

We must remedy this, perhaps in a manner similar to that of U.S. education after Russia's successful Sputnik launch. Today, only 1 percent of U.S. college graduates are engineers, as compared to about 4 percent of Japanese college graduates. The proportion of engineers in Japanese industry is rising while it is declining in the United States. Government and education must be called to action. They must encourage students to enter technical fields. Enhancing high school science and mathematics programs, and helping colleges and universities to buy costly, state-of-the-art equipment would be a start.

Commercial research and development needs stimulation and support also. Rapidly changing technology is shortening product life cycles, often to the point where recovery of research and development costs is impossible. Bold measures, such as investment credits and liberalized depreciation allowances for research and development spending, would fuel progress.

President Reagan, in his letter to *Computer Design* readers that appeared in our 20th Anniversary issue two years ago, said that America looks to this industry "...to continue to create new opportunities for our citizens, to make our lives more productive, and to give us the tools we need for our national security."

Our generals in the information technology battle are facing stiff competition on a global front. To win this battle, the growth of our talent pool of design engineers must be assured, and a constructive, encouraging environment must be created for them.

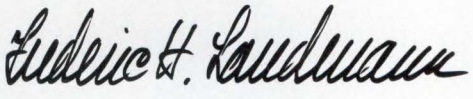

Frederic H. Landmann
Publisher

FROM A PERFORMANCE STANDPOINT, YOU CAN'T BEAT THE PRICE.

Cipher gives you something you won't find with any other GCR drive: high performance for low cost.

You can pack up to 180 megabytes of storage on a single reel. With average transfer rates as fast as 790 KBS, and burst rates as high as 904 KBS.

No other drive in this price category even comes close. In fact, the kind of performance you get with Cipher is usually reserved for drives costing twice as much.

What's the cache?

Cache is the heart of Cipher's GCR drive. This high-speed, solid-state memory is the perfect replacement for more costly—and less reliable—vacuum column and compliance-arm GCR mechanics.

Not only do you get higher performance from a smaller-sized package, but cache memory makes four-track error correction possible for the first time in a GCR product. This is accomplished by the GCR CacheTape's ability to make corrections on the fly, as well as in both forward and reverse directions independent of host interaction. Other GCR drives only give you two-track error correction.

FROM A PRICE STANDPOINT, YOU CAN'T BEAT THE PERFORMANCE.

The Cipher Advantage	
Feature	***Benefit***
Cipher Microstreamer compatible CacheTape™ interface	*Ease of integration start/stop performance*
395-790 KBS average transfer rate (450-904 KBS burst)	*High data throughput*
Error-free interface with 4-track error correction	*Improved read/write data reliability*
Compact 14" height	*Saves rack space*
Autoload/autothread and 8-character operator display	*Ease of operation*
True 0.3" inter-record gaps	*Higher formatted capacity per reel than long gap machines*

Total GCR compatibility.

Cipher's GCR CacheTape™ is a plug-compatible replacement for existing vacuum column and compliance-arm tape drives. It's completely software transparent. And it works with industry-standard tape adapters. The GCR CacheTape interface is compatible with Cipher's other ½" tape drives.

Easy to use.

Not only is GCR CacheTape easy to integrate, but it's also a snap to use. Cipher's patented automatic front loading process automatically loads and threads tape on any standard-size reel.

An easy-to-understand front panel displays literal English language messages—not coded references. This keeps things simple—even for inexperienced users.

If you're looking for unbeatable GCR performance at an unbeatable price, call 800-982-8808. Or write Cipher Data Products, P.O. Box 85509, San Diego, CA 92138.

cipher® data products, inc.

WE PUT THE BEST IDEAS INTO STORAGE.

GCR CacheTape is a trademark of Cipher Data Products, Inc.

See us at COMDEX, Booth #4127 in Las Vegas

CIRCLE 8

COMPUTER DESIGN

The PennWell Building,
Littleton, MA 01460, Tel: (617) 486-9501
Editorial/Executive Offices

Editor in Chief, Michael S. Elphick

Issue Editor, Nicolas Mokhoff
Executive Editor, John Miklosz
Managing Editor, Sydney F. Shapiro
Senior Editor, John Bond
Senior Editor, Peg Killmon
Special Features Editor, James W. Hughes
Senior Associate Editors,
Malinda E. Banash, Deb Highberger
Associate Editor, Jack Vaughan
Assistant Managing Editor,
Leslie Ann Wheeler
Copy Editors,
Helen McElwee, Leah A. Rappaport, Nancy E. Purcell, Jane E. Shattuck
Editorial Assistants, Julia E. Cote, Lisa M. Stephens
New York Field Office:
230 Park Ave, Suite 907
New York, NY 10169, Tel: (212) 986-4310
Senior Editor, Nicolas Mokhoff
Special Features Editor, Harvey J. Hindin
Western Field Office:
540 Weddell Dr, Suite 8
Sunnyvale, CA 94089, Tel: (408) 745-0715
West Coast Managing Editor,
Tom Williams
Senior Editor, Bill W. Furlow
Field Editor, Richard Goering
Editorial Assistant, Robin Mock

Publisher, Frederic H. Landmann

Marketing Director, Robert A. Billhimer
Circulation Director,
Robert P. Dromgoole
Promotion Director, Steve Fedor
Marketing Services/PR Manager,
Linda G. Clark

Production Director, Linda M. Wright
Production Manager, Philip Korn
Art Director, Lou Ann Morin
Technical Art, Designline
Ad Traffic Coordinator, Debra L. Stone
Printing Services, Padraic Wagoner

PennWell
Advanced Technology Group
119 Russell St, Littleton, MA 01460
Tel: (617) 486-9501

H. Mason Fackert, *Group Vice President*
Saul B. Dinman, *Editorial Director*
John M. Abernathy, *MIS/DP Director*
Patricia M. Armstrong,
Administrative Services Manager

TEK MICROPROCESSOR DESIGN SOFTWARE

FLEXIBLE SOLUTIONS FOR A FUTURE OF CHANGE

Tek software sheds new light on your VAX.

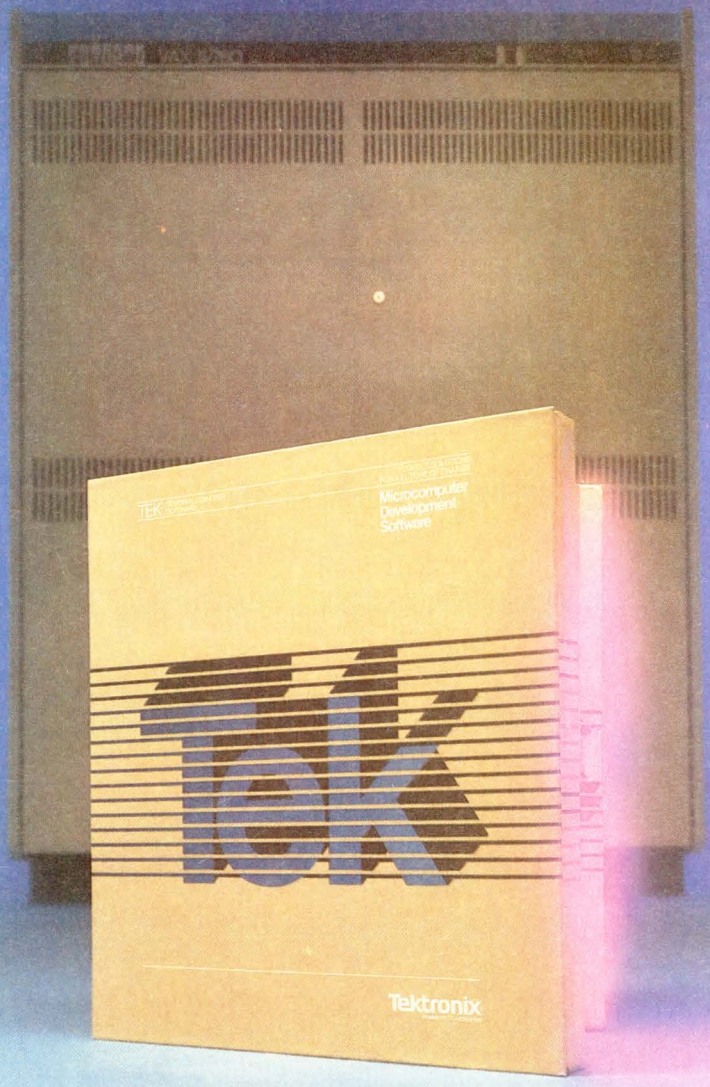

Turn your VAX™ computer into a powerful microprocessor development system.

Tektronix software. The same powerful tools that set the standard for high-level programming on Tek's 8500 series of microprocessor development systems are now available for use on your VAX. Get the sophistication of Pascal and C Language Development Systems (LANDS). Plus real-time emulation and debug when you integrate Tek's 8540 emulation unit to your system. Add 4105 Color Graphics terminals to access Colorkey+, Tek's single-key interactive user interface. All fully integrated with VAX-specific communications software.

Support from the first line of source code to the last line of debug.

Call your Tektronix sales engineer. With your VAX, we'll help you create a system that suits your engineering environment, and show you how Tek and your VAX have met the challenge of microprocessor software design.

Tektronix Microprocessor Development Systems.

Look at us now!

Call 1-800-547-1512

™ VAX is a registered trademark of Digital Equipment Corporation.

U.S.A., Asia, Australia, Central & South America, Japan: Tektronix, Inc., P.O. Box 1700, Beaverton, OR 97075. For additional literature, or the address and phone number of the Tektronix Sales Office nearest you, contact: Phone: (800) 547-1512. Oregon only: (800) 452-1877. TWX: (910) 467-8708. TLX: 151754. Cable: TEKWSGT

Europe, Africa, Middle East: Tektronix Europe B.V. European Headquarters, Postbox 827, 1180 AV Amstelveen, The Netherlands. Phone: (20) 471146. Telex: 18312-18328

Canada: Tektronix Canada Inc., P.O. Box 6500, Barrie, Ontario L4M 4V3. Phone: (705) 737-2700

Tektronix®
COMMITTED TO EXCELLENCE

Circle 13 for Literature
Circle 14 for Sales Contact

IT WILL GROW ON YOU.

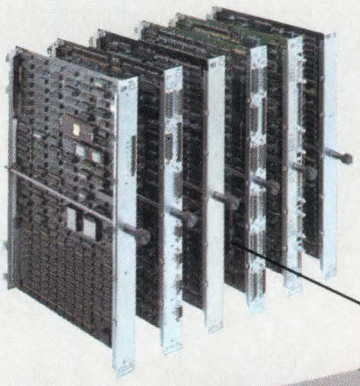

CONFIGURABLE ACCORDING TO USER NEEDS. As many as six processors can be installed in each enclosure. CPU-intensive jobs utilize multiple Applications Processors. Systems with heavy disk usage can distribute this load among several File Processors. The number of user ports can be increased by adding more I/O Processors.

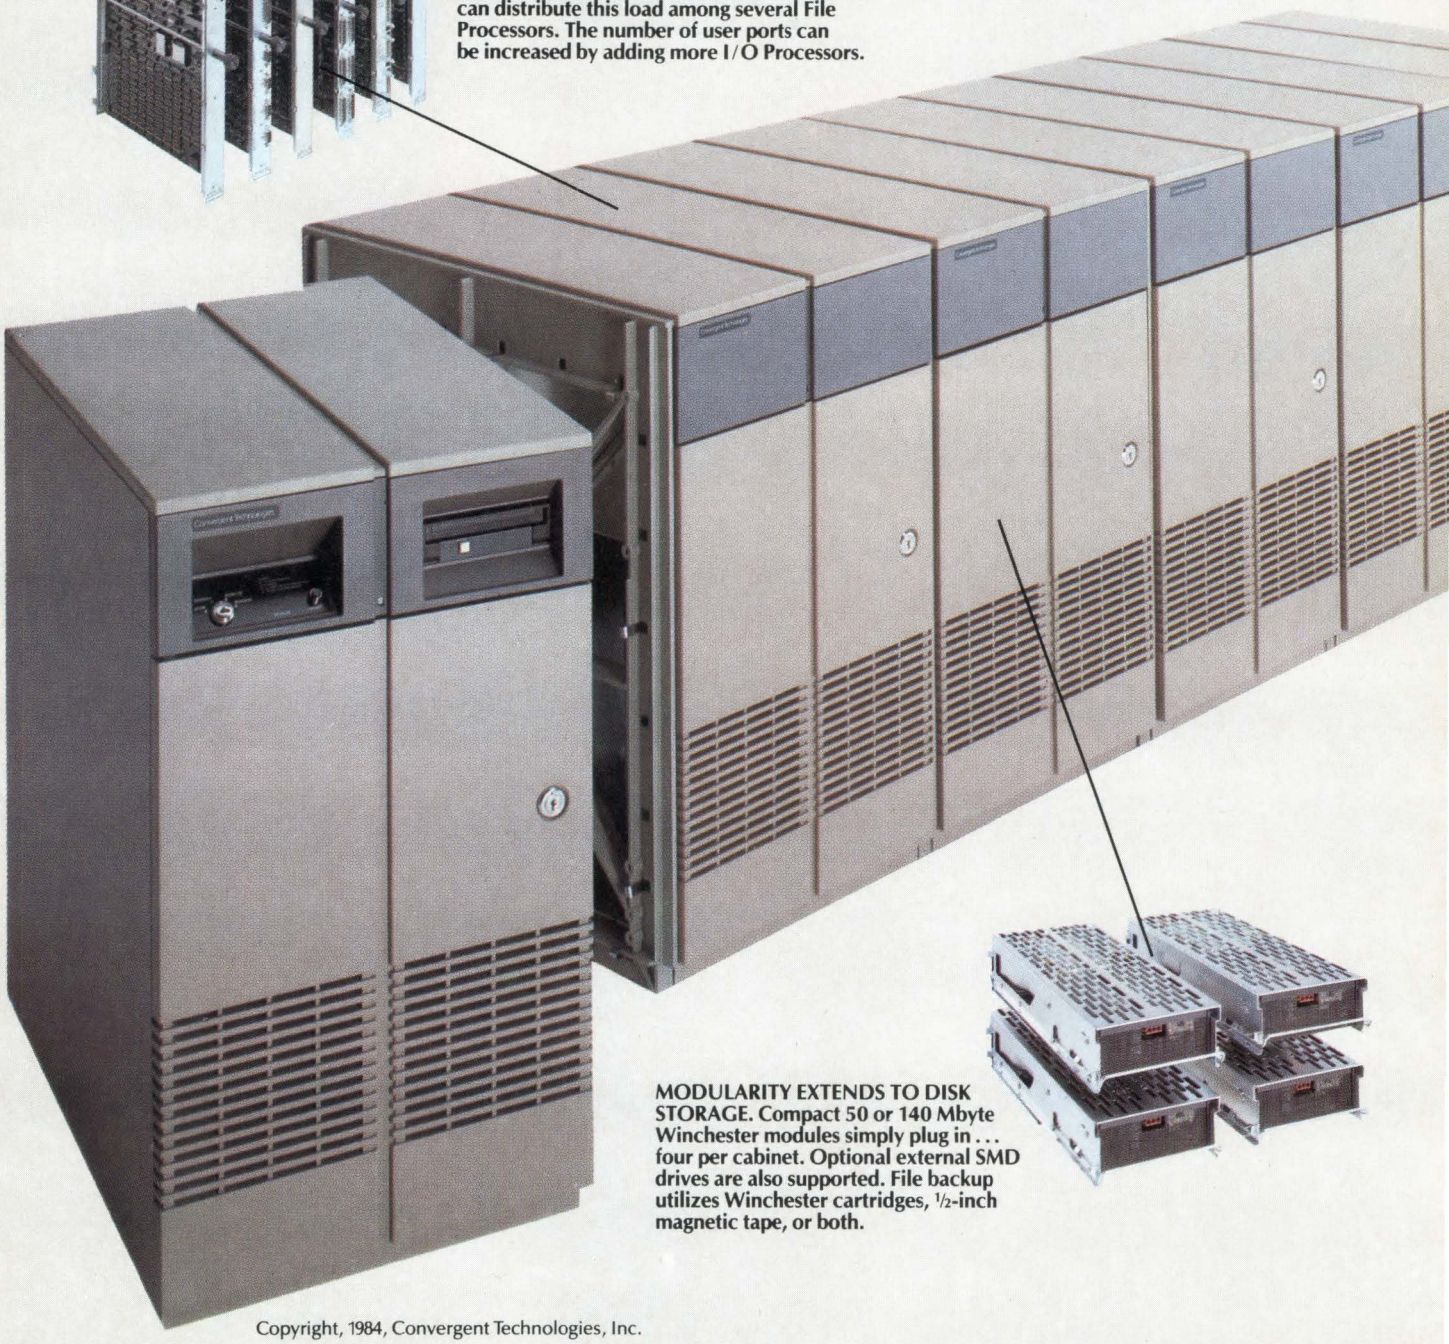

MODULARITY EXTENDS TO DISK STORAGE. Compact 50 or 140 Mbyte Winchester modules simply plug in . . . four per cabinet. Optional external SMD drives are also supported. File backup utilizes Winchester cartridges, ½-inch magnetic tape, or both.

MegaFrame™. Now OEMs can offer a high-performance UNIX™-based system that can't run out of performance.

OEMS can now deal cost-effectively with the problems encountered when user applications produce computing demands that outstrip the capabilities of conventional systems.

Convergent Technologies' MegaFrame is a revolutionary new UNIX-based super-minicomputer—so innovative in its architecture that it represents the ultimate in multiuser systems design. It grows exponentially from a system offering minicomputer-level performance to an enormously powerful engine serving as many as 128 users with 36 parallel processors, 24 megabytes of RAM and gigabytes of disk storage.

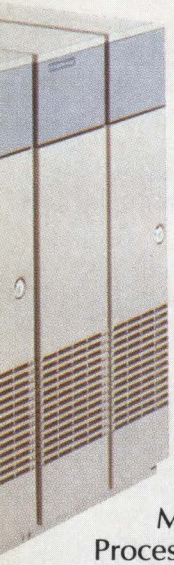

No other system can match the MegaFrame's potential for field expansion. It enables manufacturers and systems builders to keep pace with today's requirements for more and more computing services . . . but *not* at the cost of discarding hardware or performing expensive CPU upgrades.

MegaFrame's architectural breakthrough. Dependence on traditional single-CPU shared-logic architecture is the root of systems bottlenecks.

Convergent's response: a novel system utilizing *multiple* specialized processors to distribute workloads for optimum performance—even if user needs are unpredictable or subject to rapid change.

MegaFrame's virtual memory Applications Processors each have a 32-bit CPU, up to 4 Mbytes of RAM and run a demand-paged version of UNIX System V. Up to 16 of them can operate in parallel.

The File Processors effectively function as back-end machines providing DBMS, ISAM and other disk-related services. Up to six File Processors each with four disks can operate in parallel.

Terminal and Cluster Processors can also be added—the latter serving front-end communications needs. They off-load communications from the other processors by running protocols such as SNA and X25 networks.

MegaFrame's daisy-chained cabinets offer total expansion potential of up to 36 slots. OEMs configure the system needed for specific applications simply by adding the correct number/combination of processors.

Flexibility in applications development. Inclusion of one or more Applications Processors allows running UNIX System V. All standard UNIX tools are provided, along with COBOL, FORTRAN-77, BASIC interpreter and compiler, plus Pascal.

The "least-cost solution" to serving a wide range of UNIX-systems needs, MegaFrame has won acceptance from OEMs in the U.S. and abroad. The uniqueness of its modular design, its versatility in providing upgrade-path options and its price/performance advantages give it market-share potential of outstanding dimensions.

The system that will grow on you starts at a very attractive price: about $20,000 for a system that effectively supports 16 users. Send now for a comprehensive Information Package including reprints of magazine articles. It explains how MegaFrame's growth potential can impact favorably on your plans for growth in the UNIX market.

Convergent Technologies, Data Systems Division, 3055 Patrick Henry Drive, Santa Clara, CA 95050. Phone: 408/980-0850. Telex: 176-825.

MiniFrame™: the entry-level multiuser UNIX system.

Starting at under $5,000 for a single-user system, Convergent's MiniFrame offers outstanding capabilities for small to medium sized organizations running large UNIX-based applications. Utilizing an MC68010 microprocessor operating at 10Mhz, with no wait states, it provides impressive CPU speed—comparable to VAX™-11/750 running the AIM™ Benchmark. MiniFrame features virtual memory management, with demand-paged implementation of UNIX System V. It runs as many as eight terminals, with up to 50 Mbytes of integral mass storage. MiniFrame and MegaFrame are object-code compatible, allowing OEMs to offer a *complete family* of systems unrivaled in price/performance characteristics.

Convergent Technologies

Where great ideas come together

MiniFrame and MegaFrame are trademarks of Convergent Technologies, Inc. UNIX is a trademark of Bell Telephone Laboratories, Inc., VAX is a trademark of Digital Equipment Corp.

CIRCLE 15

In 64K DRAMs, INMOS high performance and flexible ...including 8K x 8.

Look to INMOS for a wide selection of 64K DRAMs. We offer access times from 100ns. By-1, by-4, and by-8 organizations. Plus features that boost performance and simplify design. This is the family that lets you avoid compromises by giving you the performance and organization you want...at costs you can afford.

8K x 8 — The cost and space saver. This newest member of the INMOS 64K DRAM family, available in a plastic package, offers access times of 120 and 150ns. The by-8 organization is a natural for microprocessor designs and other applications that require word width rather than memory depth...in process controllers, intelligent terminals, and buffer memories for example. And its pin 1 refresh makes it a very attractive alternative to 8K x 8 static RAMs. Because it combines low power and low cost with minimal support circuitry.

16K x 4 — The performance chip. With 100, 120, and 150ns access times, this organization makes a lot of sense in high-performance systems, such as high-resolution graphics, where high data rates are required. Packaged in plastic DIPs, the IMS2620 also gives you the right combination of cost, organization, and speed for microprocessor systems, terminals, and even arcade and home games. If you're upgrading from 16K x 1 chips, it provides a factor of four reduction in chip count. And its "$\overline{CAS}$ before $\overline{RAS}$" refresh assist function minimizes required support circuitry for cost and space savings.

64K x 1 — The big-system choice. This DRAM also gives you a choice of 100, 120, and 150ns access times. What's more, it includes "Nibble Mode," which allows effective cycle times below 85ns. Available in a variety of packages, including plastic DIPs, ceramic DIPs, and chip carriers, it's ideal for systems requiring lots of memory depth.

Check the chart. You'll find the industry's broadest family of 64K DRAMs...and the right device for your requirements. Then call an INMOS distributor for all the details.

Organization	Access Time (ns)	Cycle Time (ns)	Part No.
8K x 8	120	190	IMS2630-12
	150	240	IMS2630-15
16K x 4	100	160	IMS2620-10
	120	190	IMS2620-12
	150	240	IMS2620-15
64K x 1	100	160	IMS2600-10
	120	190	IMS2600-12
	150	230	IMS2600-15

INMOS Distributors: Anthem Electronics, Arrow Electronics, Falcon Electronics, Future Electronics, Lionex Corp.

inmos

P.O. Box 16000 • Colorado Springs, CO 80935 (303) 630-4000
TWX 910/920-4904 • Burlington, MA (617) 273-5150 • San Jose, CA (408) 298-1786 • Torrance, CA (213) 530-7764 • Minneapolis, MN (612) 831-5626 • Baltimore, MD (301) 995-0813 • Dallas, TX (214) 669-9001 • Atlanta, GA (404) 475-0709 • Whitefriars • Lewins Mead • Bristol BS1 2NP • England Phone Bristol 0272 290 861 • TLX: 444723

CIRCLE 16

means organizations

We've taken out operating

We've already done it for 200 American giants.

Companies who could well afford to develop their own real-time software. But instead called us.

Hunter & Ready.

Another software company, you ask?

No. But that brings us to the reason they called.

Simply put, Hunter & Ready components take the hard part out of real-time software.

Our software components handle the trickier parts of real-time. Functions like interrupts, I/O and task management. Which have nothing to do with making your application unique, but nonetheless take 60% of your software development time. Not to mention years of debugging and testing.

And if at this point you're wondering why components, the reason is simple.

Due to their standard nature, Hunter & Ready components provide the reliability strived for—but seldom attained—in real-time software. The kind that comes from 10 years of continual testing.

The same goes for any of our software components. The notable VRTX real-time kernel. The multi-stream IOX component for character and disk I/O. The high-speed FMX component for file I/O. The TRACER component for real-time debugging.

Every one tried and proven. But that's not to leave you out of the picture.

Because it's up to you what you make of any one Hunter & Ready component. Whether it's a robot, PBX system or medical instrument. It merely acts as a real-time extension for any of today's popular microprocessors.

If all this sounds too good to be

the hard part of systems.

true, we thought that it might. That's why we've prepared a packet with more details and some actual case histories.

We'd be glad to send it to you.

Just write Hunter & Ready, 445 Sherman Ave., Palo Alto, CA 94306. Or call (415) 326-2950.

If you can do that, we can do the hard part.

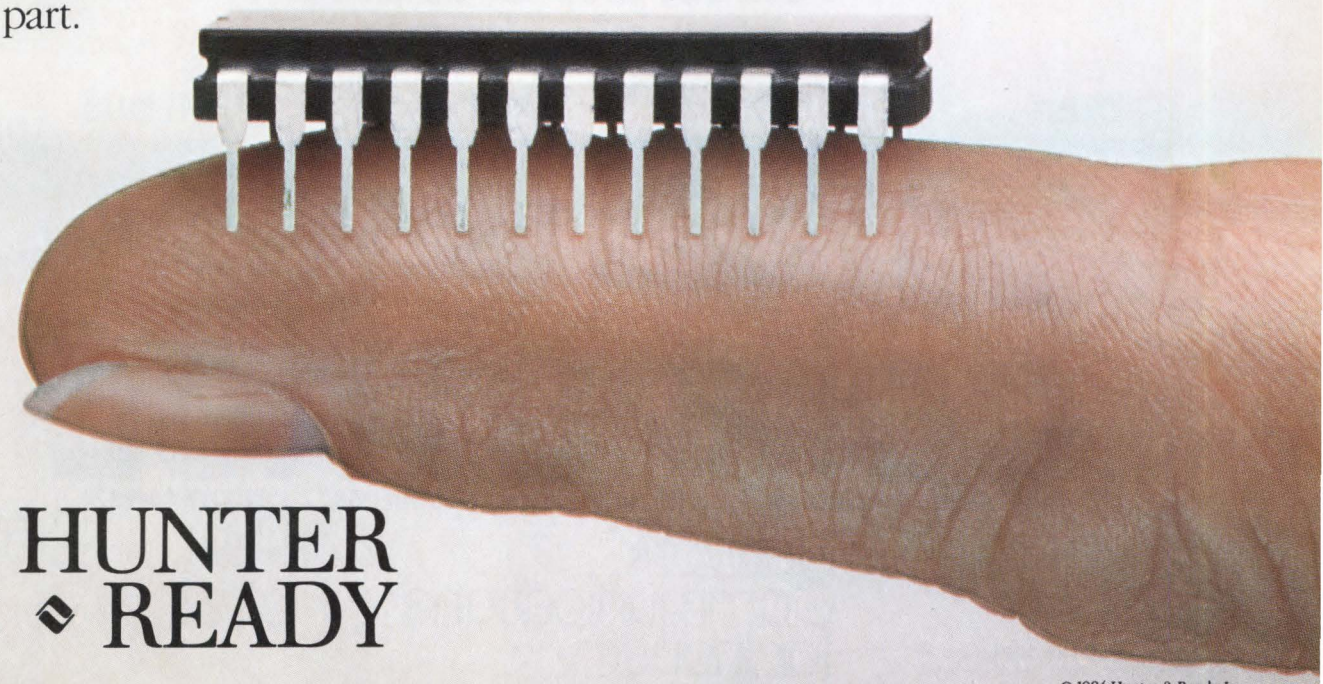

HUNTER ◆ READY

CIRCLE 17

KIDDE
AUTOMATED SYSTEMS
COMPUTROL DIVISION

Why build modems for 802.4-compatible networks? When you can buy 'em. In LSI design.

For $85.00

- Data rates to 2 Mbps without readjustment.
- Lower cost and power requirements with LSI technology.
- Multidrop (to 100 nodes) reduces cable cost.
- Simple T-connection to cable. No directional couplers.
- Increased throughput. No preamble.
- Low noise and low error rate in noisy environments.
- Cable shield can be grounded at any point.
- Transmits up to 50,000 ft. without amplifiers or repeaters.

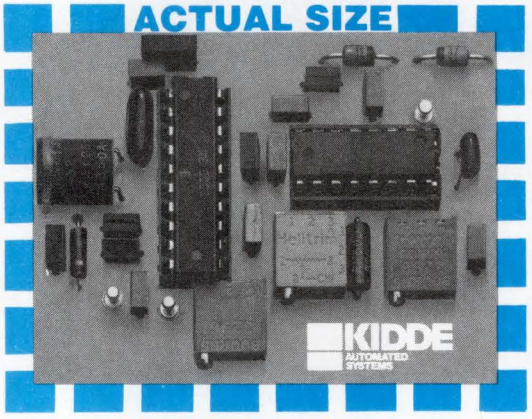

Modem M-I represents a major advance in high speed industrial communications. The unique LSI chip set is based upon a concept originated by Computrol. The result is a modem offering data rates of up to 2 Mbps. In remarkably compact 2½ x 2-inch design. At costs as low as $85**. Call Computrol. Detail your application to Carl Rohr or Don Babbitt. Let them tell you how Modem M-I can do the job more effectively. Call today.

Rate (Mbits/s)	Carrier freq. (Mhz)	RG-59 (ft.)	RG-11 (ft.)	JT4412J (ft.)	JT4750J (ft.)	T41000J (ft.)
1.0	3	5K	12K	20K	32K	50K
2.0	5	4.2K	9.5K	15K	25K	37.5K

**Quantities of 5000.

Modem M-1

FSK/LSI MODEL 30-0194
with Megalink interface to

Q-bus* . . .

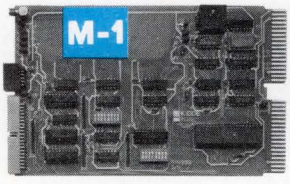

MULTIbus* . . .

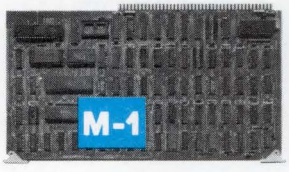

UNIbus* . . .

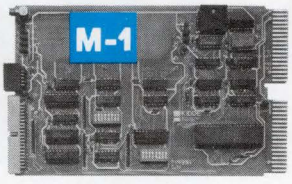

VMEbus . . .

FROM

COMPUTROL

15 Ethan Allen Highway / Ridgefield, Connecticut 06877 / (203) 544-9371

Network Communications for Industry

*Q-bus and UNIbus are trademarks of Digital Equipment Corporation. MULTIbus is a trademark of Intel Corp.

CIRCLE 18

Local area network chips stress reduced connection costs

The key to the success of local area networks continues to be connection economics. This theory of economics depends directly upon how far semiconductor chip manufacturers can integrate the three interface chips—controller, encoder/decoder, and transceiver—and thus, how much the cost per computer connection to the network can be lowered.

Baseband and broadband LANs

Both baseband carrier sense multiple access with collision detection (CSMA/CD) and broadband token-passing LANs are staking out their own territories: the former in networking computers and peripherals within a building, and the latter in tying together several buildings. Thanks to the efforts of more than six chip suppliers, LANs using baseband CSMA/CD continue to hold a leading position in the industry. Even the initial negative reaction of semiconductor chip manufacturers to the best example of a CSMA/CD baseband network, Ethernet, has not undermined its achievements. In fact, Ethernet is coming into its own.

At the same time, token-passing networks, despite their dearth, continue to be widely used where large distances must be spanned, where broad bandwidths are required for video and teleconferencing, as well as where harsh factory environments are encountered. Many broadband users are adapting their networks with board-level interface solutions, awaiting the day when monolithic chips are readily available so they can reduce their interface costs.

The list of CSMA/CD bus network interface chip suppliers is fairly long, with many supplying some of the chips on a limited production basis and others sampling them. These chips conform to both the Ethernet and IEEE 802.3 standards adopted by the IEEE 802 Local Area Network Standards Committee. (The 802.3 specification is a slightly modified version of the Ethernet specification, but in essence it is the same.)

Except for National Semiconductor Corp (2900 Semiconductor Dr, Santa Clara, CA 95051), which expects to have samples of its DP8392 bipolar transceiver by the fourth quarter, no company has supplied a monolithic transceiver chip for CSMA/CD bus networks. Nevertheless, Advanced Micro Devices, Inc (Sunnyvale, Calif) plans to sample its bipolar Am7995 transceiver chip late this year, and Intel Corp (Santa Clara, Calif) plans to introduce a monolithic transceiver chip for CSMA/CD bus networks either late this year or early next year.

It should be noted that the National transceiver chip is designed for the Cheapernet, a 10-Mbit/s, low cost network version of Ethernet, with relaxed performance specifications. The Cheapernet chip is functionally equivalent to an Ethernet, however, and can be used in such networks. The chip is simply optimized in layout and design for Cheapernet.

One reason for the difficulty in producing a CSMA/CD transceiver chip has been the high voltage of at least 500 V rms required by the IEEE 802 Committee. That is the minimum potential the chip must withstand in case of high voltage transients and lightning surges on the transmission cable, and in case surges occur when the cable accidentally touches an ac power line, or vice versa. In the National transceiver chip, a dual inline isolation transformer is used with the chip for this isolation.

Contention delays under debate

The main objection to networks such as Ethernet with CSMA/CD accessing has been (and continues to be) the nondeterministic aspect of their accessing. This leads to unacceptable message delays. For applications such as realtime process control in a factory, knowing the precise maximum delay that can occur when a message is sent from one network node to another is an absolute necessity to prevent fouling up the entire process.

Yet, there is wide debate about how heavy a load a network can support before its CSMA/CD accessing method begins to suffer delays. Everyone agrees that such networks perform well and experience no delay problems when lightly loaded. The differences of opinion center on at what point during loading do delays begin; some contend they occur at as little as 60 percent and others believe they happen at as high as 90 percent or higher. Some feel any delay at all is unacceptable. As Ethernet's resurgence demonstrates, particularly in office automation applications CSMA/CD delays are obviously not much of a problem.

Much research is now underway to minimize those delays through improved contention accessing schemes and CSMA/CD modifications. For example, delays in CSMA/CD accessing have been shown to be less when message packets of longer lengths are used. And, some researchers are combining CSMA/CD accessing (used when the network's traffic is low) with deterministic accessing such as token-passing and polling (used when collisions rise above a certain level) to control delays better. In fact, one commercially available network, UniLan, from Applitek Corp (Wakefield, Mass) claims to offer automatically the best of all three major accessing schemes—namely CSMA/CD, token-passing, and time-division multiplexing—without any of the disadvantages.

The key is UniLan's accessing protocol. It operates in the CSMA/CD

(continued on page 26)

LAN chips

(continued from page 25)

Status of Semiconductor Chips for LANs

Standard and/or LAN	Company	Controller (process, package, and availability)	Encoder/decoder (process, package, and availability)	Transceiver (process, package, and availability)	Comments
IEEE 802.3, Ethernet (CSMA/CD bus)	Advanced Micro Devices	Am7990 NMOS, 48-pin DIP, in production	Am7991 bipolar, 24-pin DIP, in production	Am7995 bipolar, in samples late 1984	Joint AMD, Mostek, and DEC development. LANCE chip set. Controller has DMA channel and no FIFO registers.
IEEE 802.3, Ethernet (CSMA/CD bus)	Fujitsu Microelectronics	MB8795A CMOS, 64-pin DIP, in samples	MB502 bipolar, 24-pin DIP, in samples	In development	Joint Fujitsu and Ungermann-Bass development. Controller has separate byte-wide data ports, and no FIFO registers. Dual port allows simultaneous packet transmission and reception.
IEEE 802.3, Ethernet (CSMA/CD bus)	Intel	i82586 NMOS, 48-pin DIP, in production	i82501 bipolar, 20-pin DIP, in production	In development, samples in late 1984 or early 1985	Joint Intel, Xerox, and DEC development. Controller has 4 DMA channels.
IEEE 802.3 Ethernet (CSMA/CD bus)	Mostek	MK68590 NMOS, 48-pin DIP, in samples	MK3891 bipolar, 24-pin DIP, in samples	—	Joint Mostek, AMD, and DEC development. LANCE chip set.
IEEE 802.3, Ethernet (CSMA/CD bus)	National Semiconductor	DP8390 CMOS, 48-pin DIP, in samples	DP8391 bipolar, 18-pin DIP, in samples	DP8392 bipolar, 14-pin DIP, in samples	Also supports IEEE 802 Cheapernet specifications. Controller has two DMA channels and buffer memory management. Encoder/decoder uses digital phase locked loop for 2-ns resolution.
IEEE 802.3, Ethernet (CSMA/CD bus)	Philips B.V.	In development	—	—	—

mode when the network is lightly loaded, and automatically adjusts the network's accessing to a more deterministic accessing form as traffic on the network increases.

CSMA/CD networks are growing in popularity, as can be seen from the increasing number of firms taking the available CSMA/CD chips and improving them with board- and box-level interface solutions. They include Interlan, Inc (Westford, Mass), 3Com Corp (Mountain View, Calif), Excelan, Inc (San Jose, Calif), Bridge Communications, Inc (Mountain View, Calif) and Communications Machinery Corp (Santa Clara, Calif).

CSMA/CD is not limited to baseband network. Sytek, Inc (Mountain View, Calif) feels that CSMA/CD is a perfectly viable approach for its LocalNet 20 LAN, a broadband 120-channel system, with each channel offering a 300-kHz bandwidth. The firm is also planning to whittle down the data rate of its earlier and higher speed LocalNet 40 broadband system to 2 Mbits/s. It will further integrate the CSMA/CD interface, compressing its current three PC boards to one board, forming a personal computer broadband network interface. Sytek plans to use LSI chips to make the reduction.

Voice networks go another route

Not everyone is convinced that CSMA/CD accessing is the best alternative, particularly some private branch exchange (PBX) manufacturers such as Ztel, Inc (Wilmington, Mass). The firm feels that for data transmission, and particularly for voice transmission, some form of deterministic accessing is absolutely necessary, and thus has chosen a token-passing approach for access to its private network exchange (PNX) LAN. Ztel feels that its ring network is the way to go, especially since IBM's yet-to-be-announced token-passing ring network will be the standard LAN of the future. Although IBM has said its network will not enter the market for at least two to three years, it has announced the cabling method to be used for its token-passing ring network. Thus, IBM has provided a sign that it intends to produce the network.

IBM's delay in announcing its network may be due to the difficulty Texas Instruments, Inc (Dallas, Tex) appears to be having in producing the first token-passing chip that conforms to the IEEE 802.5 proposed standard for token-passing ring networks. Although IBM and TI have an arrangement in which the latter will

Standard and/or LAN	Company	Controller (process, package, and availability)	Encoder/decoder (process, package, and availability)	Transceiver (process, package, and availability)	Comments
IEEE 802.3, Ethernet (CSMA/CD bus)	Rockwell International	R68802 NMOS, 40-pin DIP, samples in late 1984	—	—	Controller has DMA channel and no FIFO registers.
IEEE 802.3, Ethernet (CSMA/CD bus)	Seeq Technology	8001/8003 NMOS, 40-pin DIP, in production	8002 CMOS, 20-pin DIP, in production	In development	Joint Seeq, Silicon Compilers and 3Com development. Controller has no FIFO registers.
IEEE 802.4 (token-passing bus)	Motorola Semiconductor Products	In development HCMOS, 84-pin chip carrier, samples in late 1985	In development bipolar, samples in late 1985	—	Controller has 4 DMA channels. Broadband and baseband networks will be supported.
IEEE 802.4 (token-passing bus)	Signetics	In development	In development	NE5080/NE5081 bipolar, in production	—
ARCnet	Standard Microsystems	COM9026 NMOS, 40-pin DIP, in production	COM9032 NMOS, 16-pin DIP, in production	Available in hybrid form	Joint Standard Microsystems and Datapoint development for the latter's ARCnet.
Token-passing bus	Western Digital	WD2840 NMOS, 40-pin DIP, in production	—	—	Controller has two DMA channels, three microcontrollers, a transmitter/receiver FIFO buffer and self-diagnostics. General purpose token-passing chip.
IEEE 802.5 (token-passing ring)	Texas Instruments	In development	In development	In development	Joint TI and IBM development.

produce the chip, TI has yet to comment on when the first controller chip will be available. Delays may have resulted because the IEEE 802.5 token-passing ring working group has debated, but has yet to resolve, whether the controller chip's LSB or MSB should come first. The 802.5 proposed standard's MSB first approach is at variance with the LSB first approach for the already approved 802.3 CSMA/CD bus and 802.4 token-passing bus standards. According to some designers, however, a minimum of hardware and software can be used to bridge the differences between MSB and LSB approaches for the different networks. That view is not shared by everyone, however.

Despite IBM and TI's delays, interest in the token-passing accessing technique is great, particularly for factory network applications. An indication of this is the push General Motors Corp (Detroit, Mich) is putting on manufacturers of semiconductors, computer networks, and programmable controllers to develop its proposed Manufacturing Automation Protocol (MAP). It will be a token-passing bus network that will allow dissimilar factory automation elements to talk to one another, and is compatible with the IEEE 802.4 standard. Concord Data Systems, Inc (Waltham, Mass) has already produced the first token-passing bus broadband network to meet the IEEE 802.4 standard. Token/Net was used as part of a demonstration at this year's National Computer Conference, sponsored by the National Bureau of Standards. (GM also participated in the conference.)

In support of GM's MAP, Motorola Semiconductor, Inc (Phoenix, Ariz) is planning to manufacture the first single-chip controller to meet the IEEE 802.4 standard. Motorola is already in the advanced stage of designing the chip to be used in GM's MAP. While awaiting the chip's availability in late 1985, Motorola will offer a three-board, token-passing bus interface approach using gate arrays. The firm is actively developing an encoder/decoder chip as well, with samples expected late next year.

As the Table, "Status of Semiconductor Chips for LANs," indicates, token-passing chips are presently available from both Standard Microsystems Corp (Hauppauge, NY) and Western Digital Corp (2445 McCabe Way, Irvine, CA 92714). Neither of the firms' chips meet any IEEE 802 LAN standards, however. The Standard Microsystems COM9026 controller

(continued on page 28)

LAN chips

(continued from page 27)

and the COM9032 encoder/decoder chips implement ARCnet—the token-passing bus network protocol of Datapoint Corp (San Antonio, Tex). The network employs a modified token-passing and self-polling accessing technique. The two chips have been available for quite some time, and many users have gained experience in implementing them for networking applications. For the transceiver function, a hybrid circuit is available from Zenith Corp (Chicago, Ill). The Western Digital WD2840 controller IC is also a token-passing chip, but does not conform to any particular networking protocol, nor to any network standards specification.

It should be noted that the design and fabrication of a token-passing chip is much more complex than that of a CSMA/CD device. Thus, it is not surprising to see the availability of token-passing chips lagging behind that of CSMA/CD devices. Designers of token-passing chips have to be concerned with more functions than those encountered with CSMA/CD chips. For example, the controller chip must not only determine how a particular token is put on and passed around a ring network, but also how to detect a duplicate token, how to detect the loss of a token, and how to duplicate that lost token.

Semiconductor manufacturers of CSMA/CD chips have also had a head start on their token-passing counterparts, gaining valuable design and production experience in the process. As manufacturers of token-passing chips such as TI fine tune their design and production processes, it seems certain that they will develop relatively inexpensive token-passing chips that conform to the various standards. At that point, the competition between token-passing and CSMA/CD network interface chips should become keener.

Although in general CSMA/CD networks are finding greater use in office automation applications and token-passing devices in factory automation applications, there are some notable exceptions. For example, Ethernet is trying to penetrate the factory automation scene, and in fact is used in many automatic test equipment (ATE) data highways in the plant. Anticipating the network's wider use in the factory, companies such as Bridge Communications have begun to offer diskless communication servers for Ethernet. For the last few years, General Electric's Intersil subsidiary (Cupertino, Calif) has been offering a broadband LAN with CSMA/CD accessing. GEnet is also aimed at factory automation applications. And as mentioned previously, PBX vendors such as Ztel are using a token-passing ring approach for a network that is essentially aimed at office automation applications.

IBM, which has long championed the token-passing ring network for office automation—an approach for which no product yet exists—is also committed to the token-passing bus network approach. The company is one of more than a dozen manufacturers backing the GM MAP effort demonstrated at the recent NCC.

The choice between baseband and broadband, CSMA/CD and token-passing, and bus and ring topologies (or even star topologies) eventually may not be the determining critieria, however. For example, while Ethernet attempts to move out of the single building office automation arena, broadband networks—both token-passing and CSMA/CD, in ring, bus and star topologies—are going the other way. These networks not only link up several buildings within a campus environment, but also penetrate each building as subnets attached to the larger backbones that span the buildings. What this all means is that potential users will have a wealth of technology choices, and will be able to choose whatever approach will provide the most cost-effective solution.

—Richard Parker,
Contributing Editor

SYSTEM TECHNOLOGY

(continued on page 32)

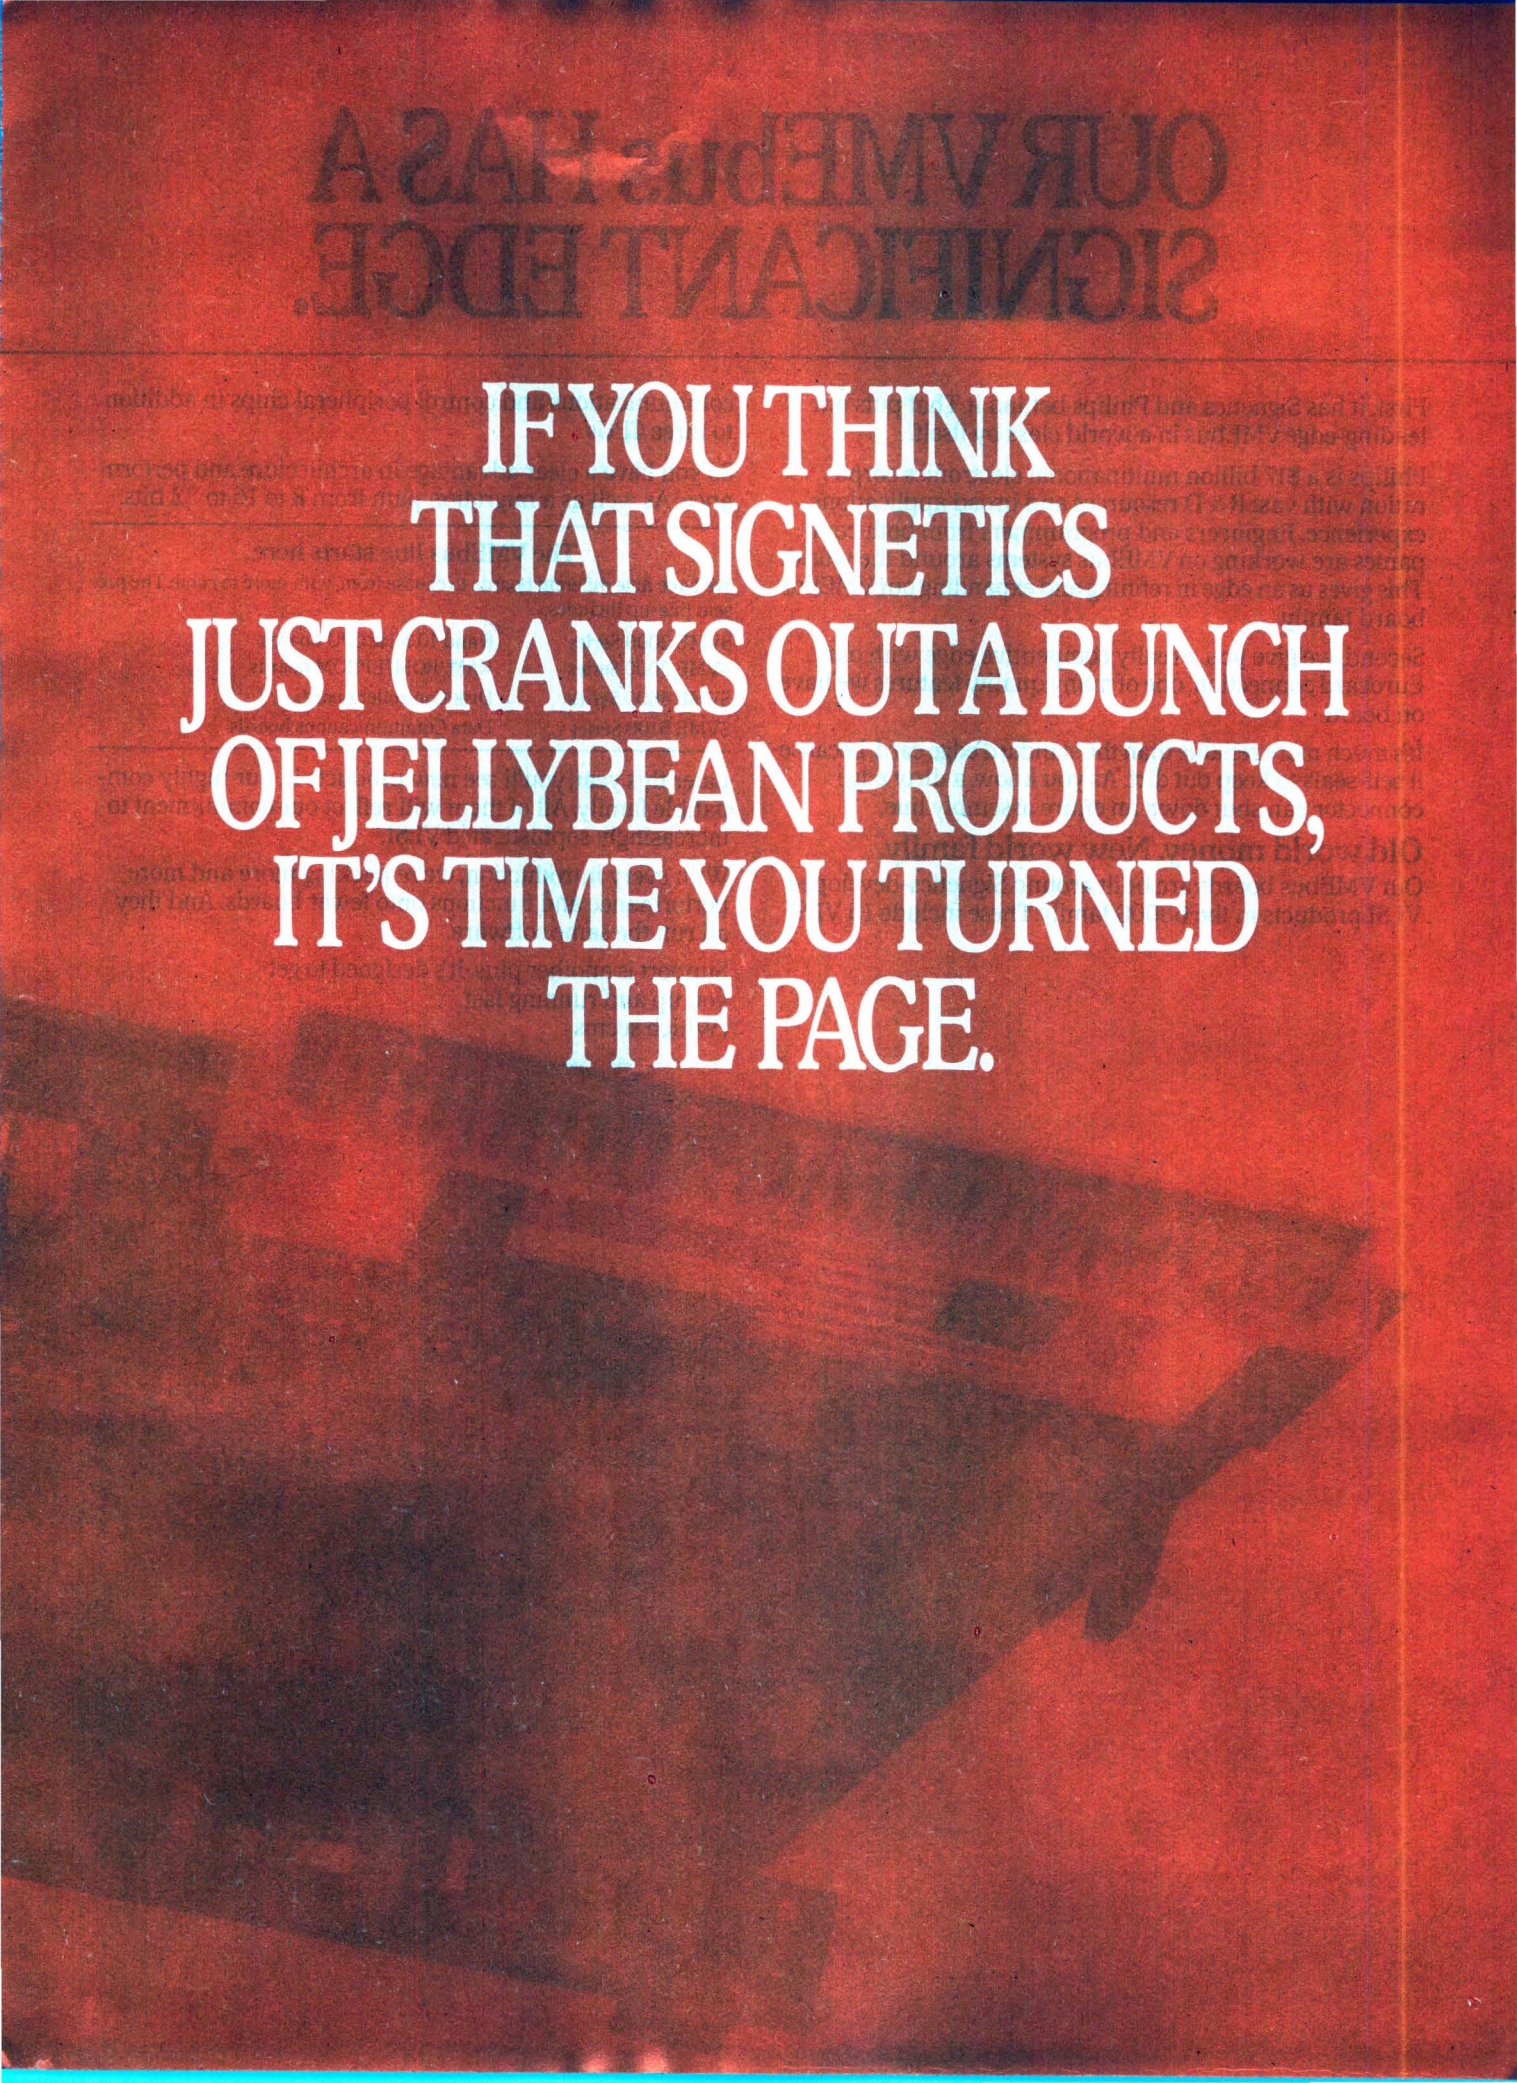
IF YOU THINK
THAT SIGNETICS
JUST CRANKS OUT A BUNCH
OF JELLYBEAN PRODUCTS,
IT'S TIME YOU TURNED
THE PAGE.

OUR VMEbus HAS A SIGNIFICANT EDGE.

First, it has Signetics and Philips behind it. That puts our leading-edge VMEbus in a world class by itself.

Philips is a $17 billion multinational electronics corporation with vast R & D resources and broad applications experience. Engineers and programmers from both companies are working on VMEbus systems around the world. This gives us an edge in refining and expanding our VMEbus board family.

Second, we give you a really competitive edge with our Eurocard connection, one of many quality features we have on board.

It's much more reliable than the standard edge card, because it self-seals to keep out dirt. As you know, a dirty edge connector can shut down an entire assembly line.

Old world money. New world family.

Our VMEbus boards are built around Signetics-developed VLSI products in the 68000 family. These include 15 VLSI communications and control peripheral chips in addition to three CPUs.

So you have a clear advantage in architecture and performance. As well as a migration path from 8 to 16 to 32 bits.

The VMEbus line starts here.

We have nine different boards to choose from, with more to come. The present line-up includes:

SVME 2000 Series	8 and 10 MHz CPU boards
SVME 3000 Series	RAM/ROM/EPROM boards
SVME 4000 Series	Disk Controller boards
SVME 5100 Series	Data Communications boards

Later this year, you'll see new products in our highly compatible family. All of them will reflect our commitment to increasingly sophisticated VLSI.

With every introduction, we're packing more and more performance and functions onto fewer boards. And they all run the same software.

Support is another plus. It's designed to get you up and running fast. (Which seems

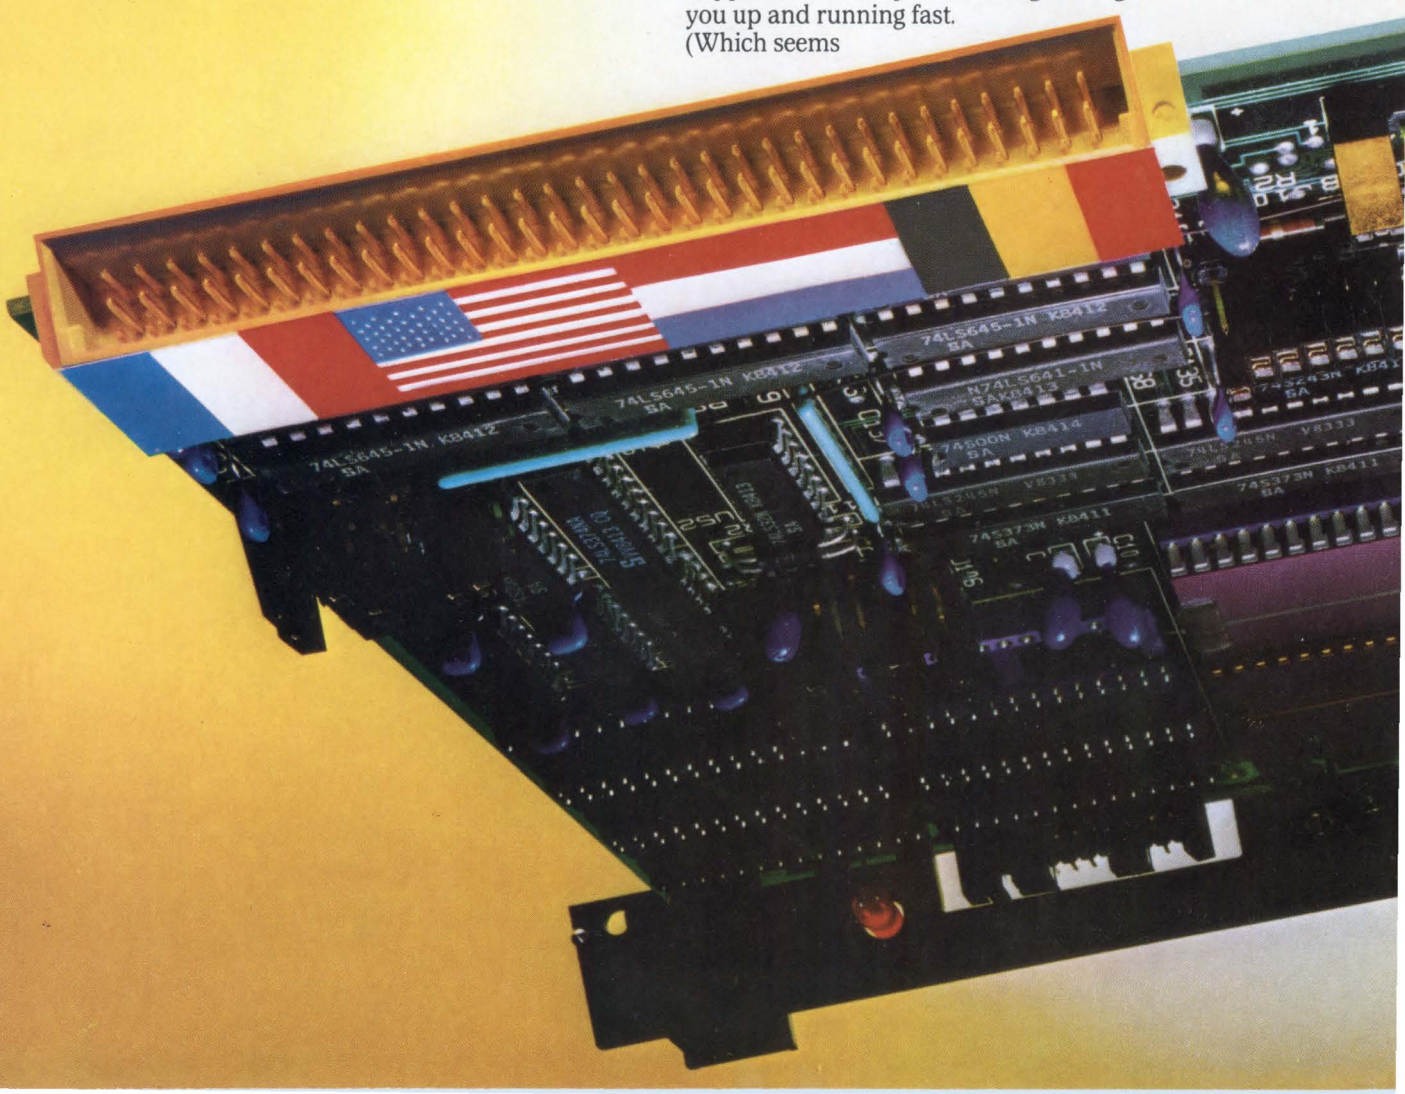

appropriate for a bus system that runs two to five times faster than the competition.)

We offer the pSOS-68K Real-Time Executive, together with a pROBE-68K debugger. Cross compilers and macro assembler. And our User Work Station, which interacts with existing computers for hardware/software integration and software debugging. If you need more help, our Field Application Engineers are on call.

An Evaluation Kit that tells all.

This is the quick, easy, low-cost way to assess the benefits of the VMEbus and our compatible VMEbus products. Our Evaluation Kit comes fully assembled and ready to go. You plug it in, hook up a terminal and power it up.

It includes all the hardware, software and support you need for a thorough evaluation. Afterwards, you can expand it into an economical resident development system.

With a full line of VMEbus boards at your local Signetics distributors, you can get a jump on the rest of the world right now. Contact them for details, including information on the Evaluation Kit. Or call us toll free for literature and the phone number of your nearest Signetics representative.

You'll soon see why our VMEbus gives you such a clear edge.

800-227-1817, Ext. 902F

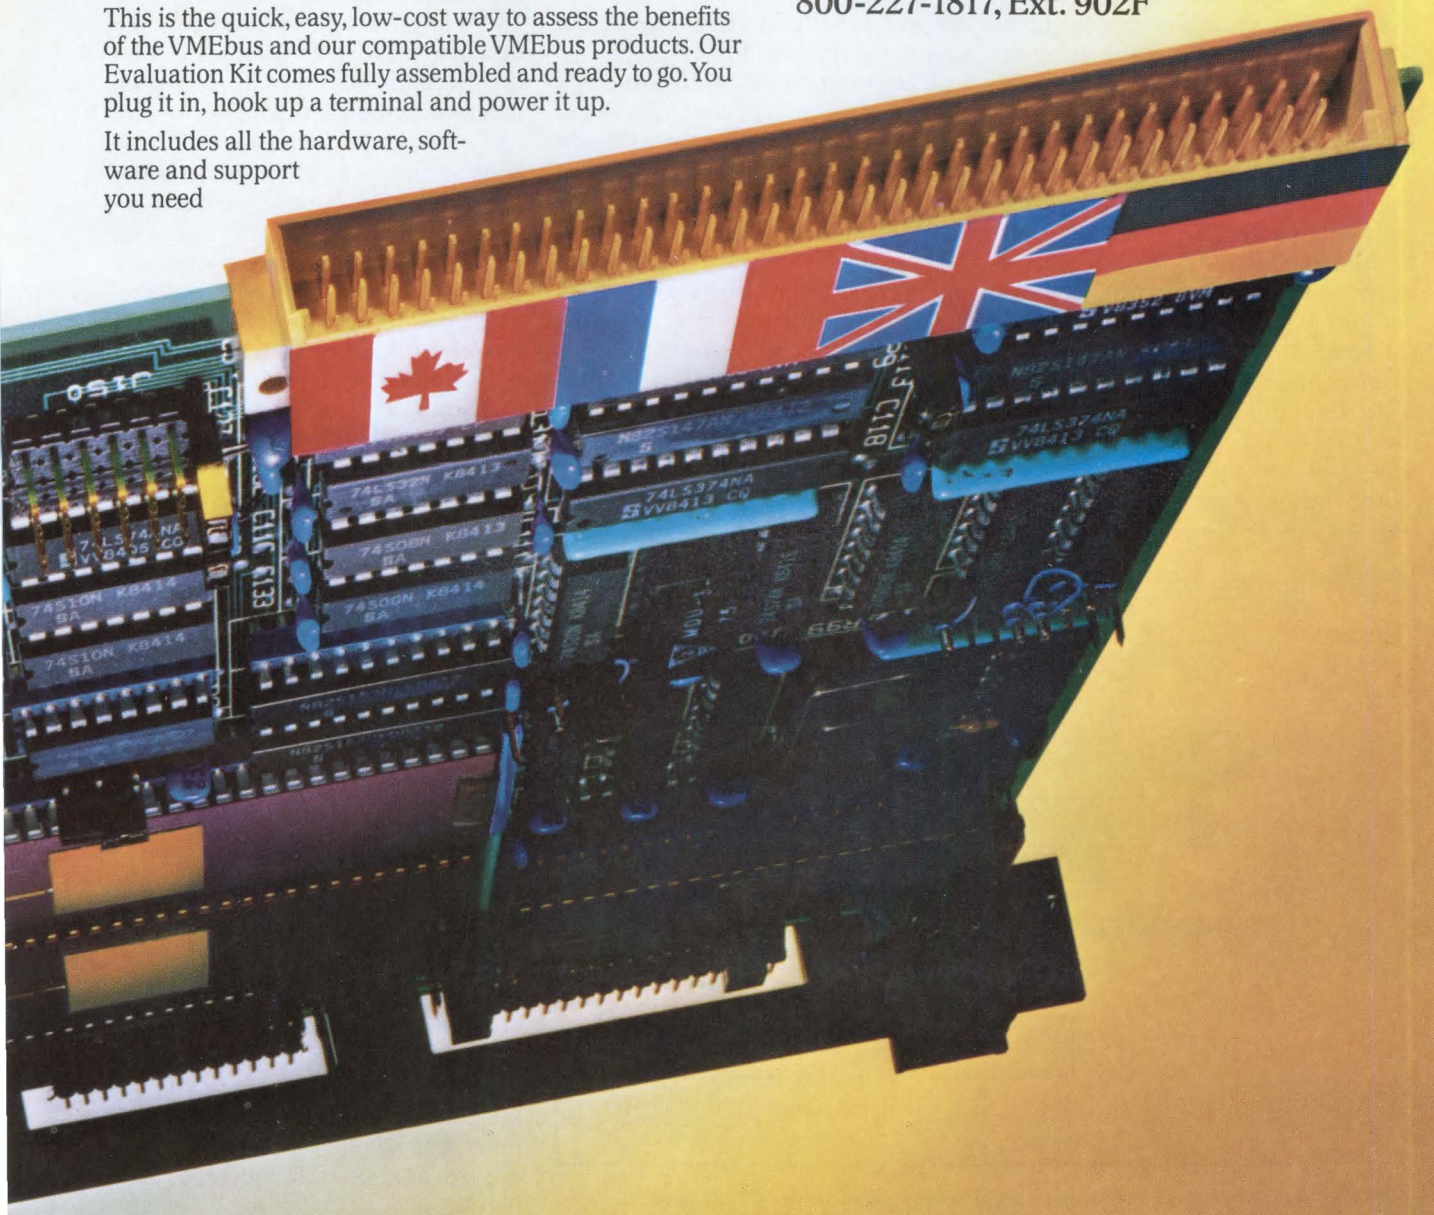

VLSI from Signetics

a subsidiary of U.S. Philips Corporation

VLSI memory access controllers support 32-bit processors

Computer designers know that the theoretical specifications of a naked microprocessor do not directly translate into those of a practical computer system. For example, for memory management-based systems, there is always the delay inherent in the processor's memory management (wait states), as well as the delay in the computer system's memory and processor interface. Furthermore, microprocessor performance touted at 10-, 12-, or even 16-MHz rates for engines like the Motorola 16/32-bit 68010 and the 32/32-bit 68020, has not been extendable to a computer's memory, memory protection, and I/O without expensive PC boards full of SSI and MSI parts.

But now, Signetics Corp of Sunnyvale, Calif has taken the heat off the designer by introducing two chips—the 68910 and the 68920. These perform memory management (with protection) and much more (like cache memory control to avoid wait states, and virtual I/O control) for the 68010 or the 68020 processors, respectively. As a bonus, the chips are designed to accommodate multiprocessors on the same memory access bus without quickly depleting scarce bus bandwidth.

The 84-pin, 90,000-transistor, 2.5-micron HCMOS VLSI part (the 68910) and the similar architecture 120-pin part (the 68920) allow the theoretical performance of the processors to be closely achieved in the virtual memory microcomputer systems that are popular today. Such systems can actually compete with VAX-like minicomputers in throughput (see Table, "Performance Comparisons").

The chips, known as memory access controllers (MACs), will be sampled in 1985. The second and third quarters of that year will see the 68910, which handles the 68010's 24-bit (16-Mbyte) addressing range by mapping it to a 24-bit physical address space, and the 68920, which accommodates the 32-bit (4-Gbyte) addressing range 68020 with mapping to a 28-bit physical address space. A chip designated the 68905 will appear first (in December). Geared to virtual memory-based microcompter system software design, it has most of the 68910's features and will aid designers preparing MAC-based systems.

If all Signetics did was come up with a single-chip memory manager for the 68000 series, it would be interesting but hardly a great architectural feat given that the memory management unit (MMU) chips already exist. Unfortunately, among their other problems, these MMUs add wait states and are not designed for a multiprocessor environment. Like an MMU, a MAC divides the microprocessor address space into pages, segments, or paged segments. But, MACs also accommodate a cache memory based microcomputer design and take care of virtual I/O.

For the former, the MAC contains all the logic and identification memory to control a user-configurable cache (both width and depth) from 1 to 32 Kbytes in size. For the latter, the MAC creates a virtual address bus between the microprocessor and its I/O controllers. The bus's "presence" results in increased system throughput.

The MACs are coprocessors, and the company could have opted for an MMU and cache onboard the microprocessor as, for example, Zilog has done. But, with the coprocessor, computer designers can specify the paging or segmentation technique, the cache hit rate, and more. It also allows multiple MMUs when designers feel it is necessary to have them.

(continued on page 34)

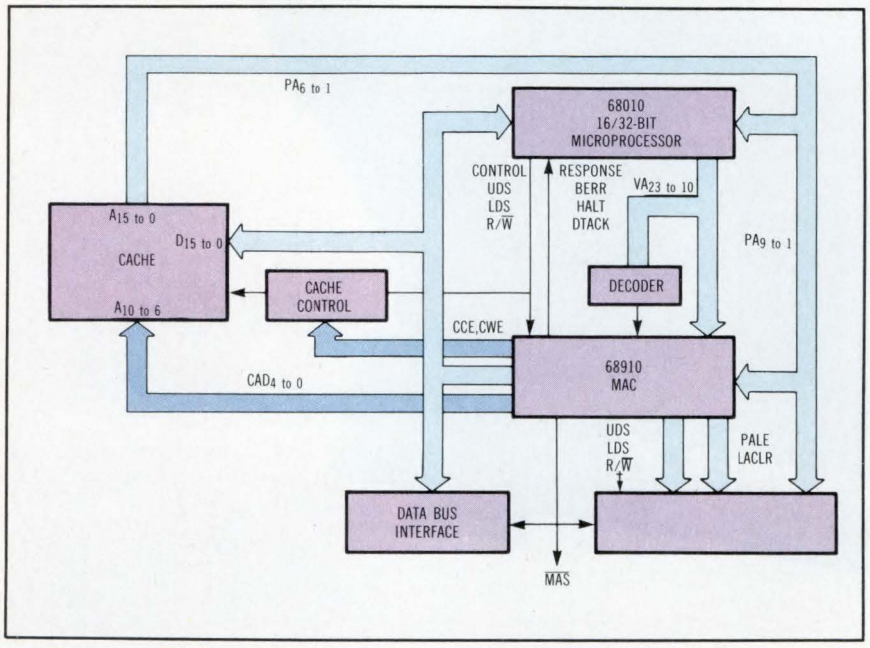

The 68910 memory access controller (MAC) works with the 16/32-bit 68010 microprocessor to establish a virtual memory management-based microcomputer system. The MAC handles cache memory, memory management and protection, and more for a single- or multiprocessor-based system.

Introducing a family of plotters... all fast, and all friendly.

Hewlett-Packard's plotters bring you ease of use...at a very affordable price. The HP 7580B, for A-D size plotting, is $13,900. The HP 7585B, for A-E size plotting, is $16,900.* And the new HP 7586B, for roll-feed and single-sheet plotting, is $21,900.**

*Domestic U.S. prices only

Hewlett-Packard drafting plotters... so easy to use, they almost run themselves

The plotters you don't have to babysit.

Hewlett-Packard's high performance drafting plotters are designed to make your professional life a little easier. So you can concentrate on doing your job, not figuring out how to run your plotter.

With the HP family of plotters, plotting has never been easier:

■ **Operating simplicity.** Just four buttons on the front panel run the entire plotter. And HP's joystick control moves the pen quickly and effortlessly.

■ **Quick and easy paper loading.** Our no-fuss, no tape, loading methods make single-sheet paper loading as easy as rolling paper into a typewriter. And HP's streamlined, non-sprocketed roll media lets you load rolls in less than 60 seconds.

■ **Compact and portable.** All HP drafting plotters can be moved easily from one area to another, letting you share one plotter among several users.

HP features let you forget the details.

And our automatic features further simplify plotter operation:

■ **Automatic paper size sensing** sets the correct margins for your paper automatically, so you'll never have to worry about "plotting off the paper."

■ **Automatic pen capping** prevents your pens from drying out and skipping, because HP plotters *never* forget to cap your pens.

■ **Automatic pen settings** always set the correct pen speed and force for the types of pens you're using–so you don't have to worry about these details.

HP designed-in quality and reliability.

And Hewlett-Packard's designed-in quality and reliability means plotting performance you can rely on, job after job. So your plotter will always be ready when you are.

Hewlett-Packard. Your best choice.

Hewlett-Packard drafting plotters bring you the plotting ease, performance and reliability that have made us leaders in the plotting industry. So when you make the decision to go with Hewlett-Packard, you know you've made the best choice.

If you'd like more information about our family of friendly drafting plotters, write to: Hewlett-Packard, Marketing Communications, 16399 W. Bernardo Drive, San Diego, CA 92127. Or call Craig Schmidt at (619) 487-4100.

HEWLETT PACKARD

1101402

CIRCLE 20

See us at Booth #568 at the Autofact Show

VLSI memory access controllers

(continued from page 32)

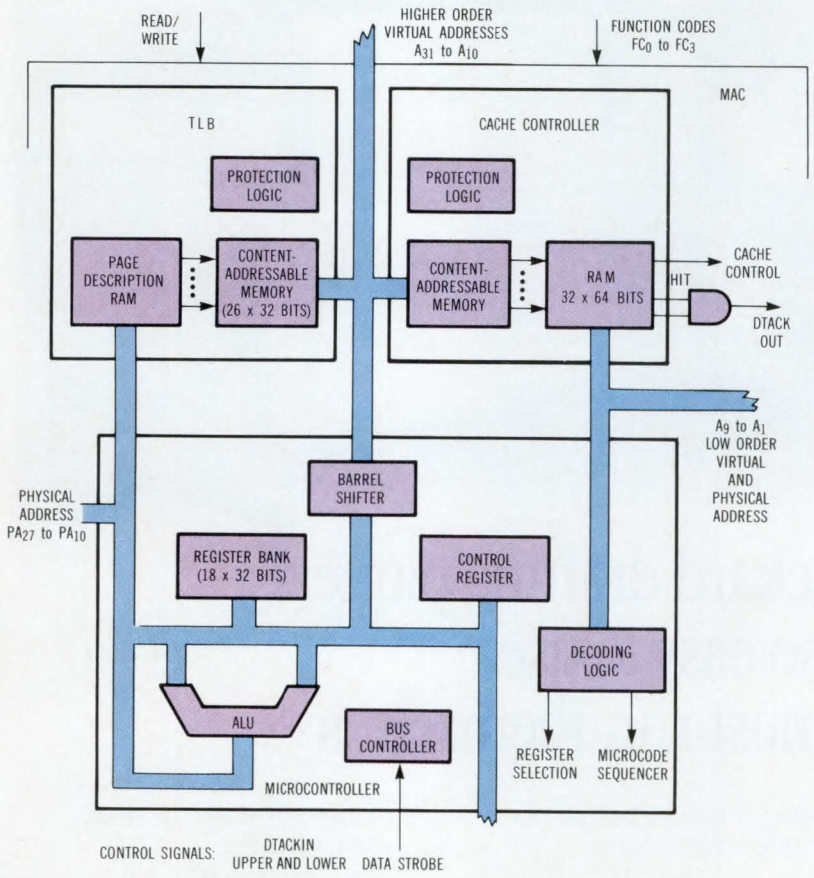

Both the 84-pin, 90,000-transistor 68910 and the similar architecture 68920 allow virtual memory-based microcomputers to compete with VAX-like minicomputers. These coprocessors have three major subsystems—a cache controller, a translation lookaside buffer (TLB), and a microcontroller.

There is yet another reason for adopting the coprocessor concept. To add the functionality that the MAC supplies onto the 68010 would require a chip with about 150,000 transistors and sub-2-micron process technology. While a device of this sort is possible, it would lock computer designers into the Motorola/Signetics approach to virtual memory microcomputer architecture. This is undesirable given the need for design flexibility in today's systems.

The reason that minicomputers and mainframes perform the way they do is that they can accommodate processor/memory speed mismatches, translate virtual to physical addresses, control a hierarchy of memories, buffer bus accesses, and control virtual I/O. The MACs perform these functions for the 68010 and 68020. They allow user-defined memory protection schemes such as no protection, normal protection, the rings that set up protection hierarchies, and, for the first time in silicon, a so-called guarded-region design.

Three-part architecture

To do their many jobs, the MACs contain three major subsystems—a cache controller, a translation lookaside buffer (TLB), and a microcontroller. The cache controller provides the control logic for an external static RAM cache; the TLB translates virtual to physical addresses, and, as in most coprocessors, the microcontroller runs the whole show. The controller and the TLB work together since most read operations are cache-supplied, many write operations are overlapped with subsequent reads, and the MAC's MMU and cache operations are often in parallel.

When a virtual address is given to the MAC (by its microprocessor), the virtual address accesses a content addressable memory in the cache controller. If the required program or data is in the cache (and the user-defined access permissions are not violated), the coprocessor MAC furnishes the requested information to the processor. While this is taking place, the TLB is also accessed. But, it is not activated unless the desired information is not in the cache. If it is not, then the associated physical address is output and the system memory is accessed. The acquired data is given to the processor and stored in the MAC's cache as it is likely to be needed again soon.

If the needed virtual to physical address mapping information is not in the TLB, the MAC's microcontroller takes the microcomputer system bus from the 68010 or 68020. It then fetches the necessary information

(continued on page 36)

Pick a number from 1 to 5

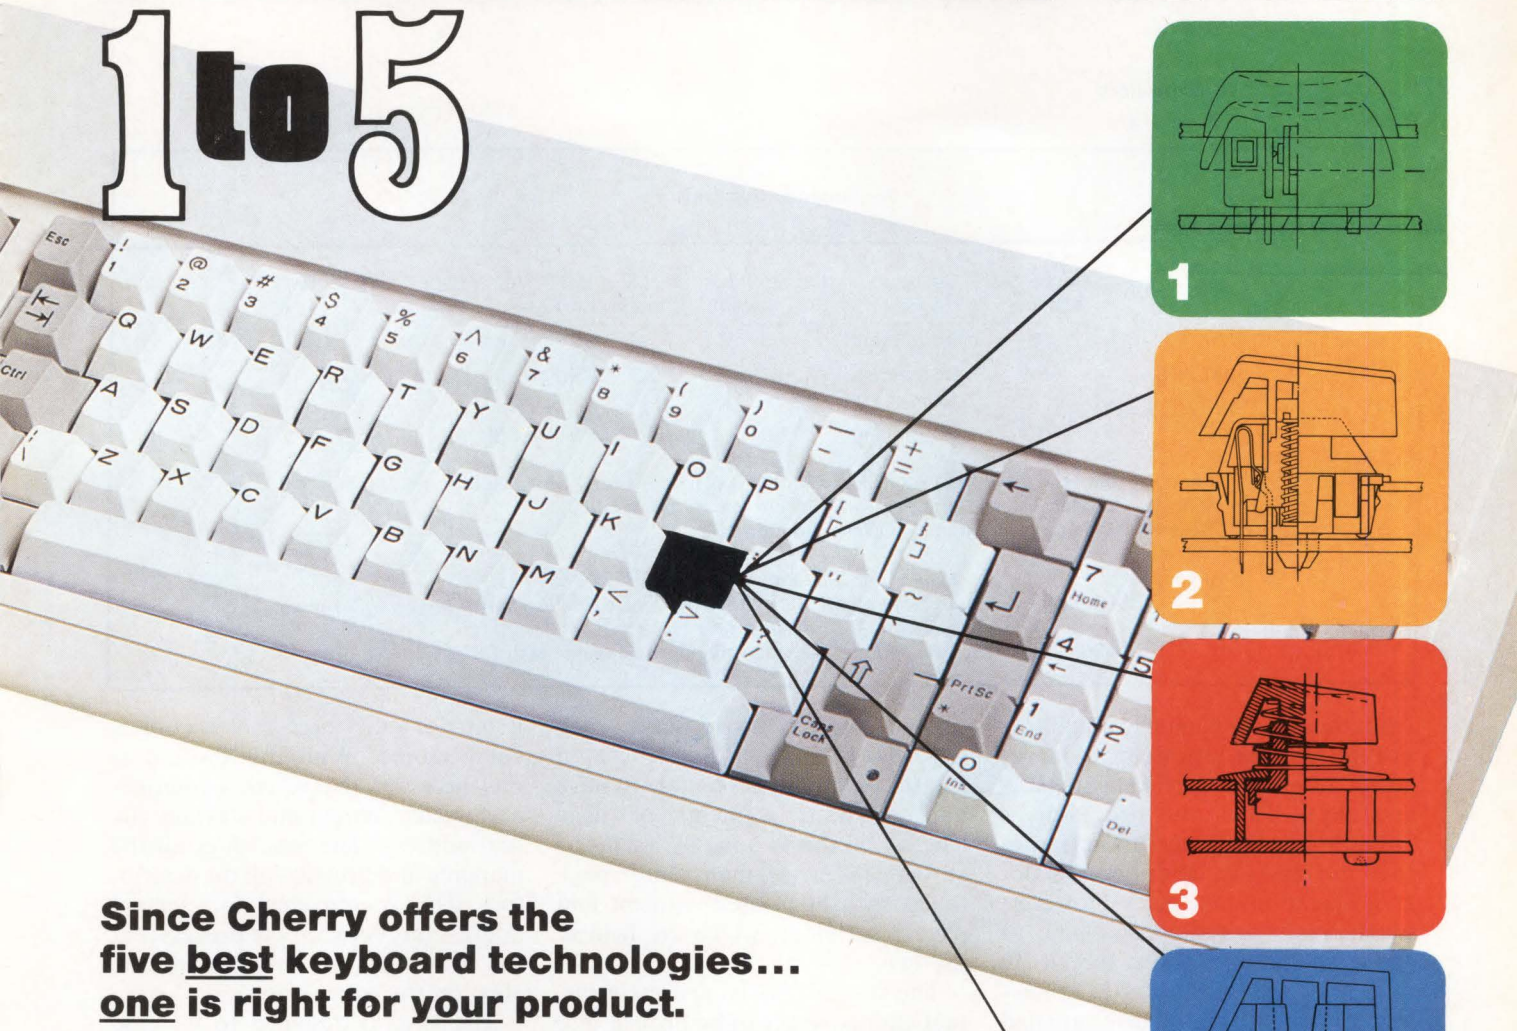

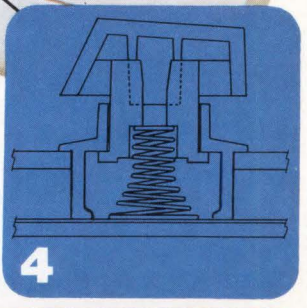

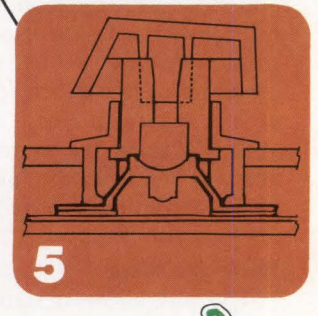

Since Cherry offers the five best keyboard technologies... one is right for your product.

All are low profile. **All** are DIN compatible. **All** are current state-of-the-art. **All** the most cost-effective technologies. **All** in standard or custom designs. **All** backed by more than a decade of Cherry keyboard know-how. **All** available now. From Cherry. Your Keyboard Headquarters.

1 **Mini-Travel Hard Contact**
Series M8. Mini, .100″ travel versions of our gold crossbar contact configuration. Same low profile. Even lower cost. Super low rest height: .426″.

2 **Full Travel Hard Contact**
Series MX. Gold crossbar contact configuration relied on in millions of applications. Worldwide. Travel: .160″. Rest height: .703″.

3 **Full Travel Capacitive**
Series MF. Pad capacitive in a uniquely simple design requiring only five parts and a snap-in pad. Travel: .150″. Rest height: .703″.

4 **Full Travel Linear Feel Sealed Contact**
Conical steel spring action plus the quality and reliability of sealed silver contacts for long life, low cost. Top mount ergonomic sculptured keyshape. Travel: .150″. Rest height: .665″.

5 **Full Travel Tactile Feel Sealed Contact**
Elastomeric action combined with sealed silver contacts for long life, low cost. Top mount ergonomic sculptured keyshape. Travel: .150″. Rest height: .703″.

NOW! WIRELESS KEYBOARDS, TOO!
To get facts fast, call Eric Olson at
1-312-578-3522.

CHERRY®
KEYBOARD HEADQUARTERS™
CHERRY ELECTRICAL PRODUCTS CORP.
3614 Sunset Avenue, Waukegan, IL 60087 • 1-312-578-3500

CIRCLE 21

VLSI memory access controllers
(continued from page 34)

Performance Comparisons

	Cache Size (Kbytes)	Cache Hit Rate	Cache Block Size (bytes)	# TLB Entries	TLB Hit Rate	Page Size (Kbytes)	MIPS
VAX-11/780	8	95 percent	8	128	99 percent	0.5	1.5
IBM 3084	64	95 percent	128	128	99 percent	4	8*
68020 plus MAC	32	90 percent	8	32	99 percent	4	1.7**

*For one processor; four processors = 27 MIPS. This means that 4 processors should equal 4 x 8 MIPS or 32 MIPS, but in practice, 5 MIPS are lost to contention.
**Based on ratioed computation, without floating point.

Sources: VAX 11/780— *Mini-Micro Systems,* Oct 1982, p 202
IBM 3084— *IEEE Spectrum,* Jan 1983, p 30; also *IBM Journal of Research & Development,* 25th Anniversary Issue
68020 plus MAC Calculation based on ratioed comparison with 68010

from the system memory and puts it in the TLB. The processor to coprocessor access is then tried again (successfully). When the processor is attempting a write operation (rather than the read just described), analogous operations take place. However, overlapping writes are possible to speed system operation. About 50 percent of the time, accesses following writes will be cache-satisfied reads. Thus, the MAC not only speeds reads, but allows many writes to proceed without slowing the processor.

While most cache systems use a serially implemented, set-associative organization (each RAM-implemented set is indexed by part of the applied address to yield a tag and data word), the 68910 and 68920 use a parallel-implemented, sector-associative cache because it is cheaper for their architecture. In this design, the MAC's address tag matching action and the presence bit access action proceed simultaneously.

Divide the space

The MAC allows designers to divide the computer's address space in three different ways. First, there are segments—contiguous in memory and multiples of a page size. The maximum segment length is between 8 Kbytes and 16 Mbytes for the 68910, and between 128 Kbytes and 4 Gbytes for the 68920. With the second division scheme, fixed-size pages, each page can be 1, 2, 4, or 8 Kbytes in length. Finally, it is possible to have noncontiguous segments of page length. All this is done by means of MAC-resident segment and page tables with associated segment and page descriptors, which are held in the TLB.

The MAC allows its processor virtual address space to be broken into up to four regions. These regions can be set up, for example, to store and to share code and data for an operating system kernel, a system library, and process-specific code and data. The computer system designer can use this feature to keep program sizes down. For example, in a Unix-based system, the addition of subroutines from the system library can greatly increase the size of a normally small program. But with the MAC, the system library region stores the subroutine code with a link to every program that calls it. The result is more efficient system operation since each subroutine entry point is a fixed member of the virtual space of each calling program.

The MAC implements its virtual I/O using virtual addresses because this is an efficient process in a paged system. Since pages need not be in consecutive locations in main memory, the computation of the physical address of each main memory page and the issuing of a disk access instruction is cumbersome. With the MAC, it is only necessary to give a disk conroller a segment's length and starting virtual address. The MAC does all the mapping and provides all the descriptors. This is equivalent to a virtual address bus between the process and I/O controllers on one side and the MAC on the other side.

The MAC is designed to increase the cache hit rate in a single microprocessor system. It does this by using a mainframe technique—multiple memory bus transfers. The company has determined just how efficient the two designs it has for accomplishing bus transfers can be—throughputs close to ideal are possible. Similarly, in a multiprocessing environment (eg, four 68010s accessing the same memory), a cache-based system can be designed with no significant degradation compared to a one-processor system.

—Harvey J. Hindin,
Special Features Editor

SYSTEM TECHNOLOGY
(continued on page 42)

Talk to the editor
Have you written to the editor lately? We're waiting to hear from you.

Dolch.

Compact-Logic-Tester

COLT 300

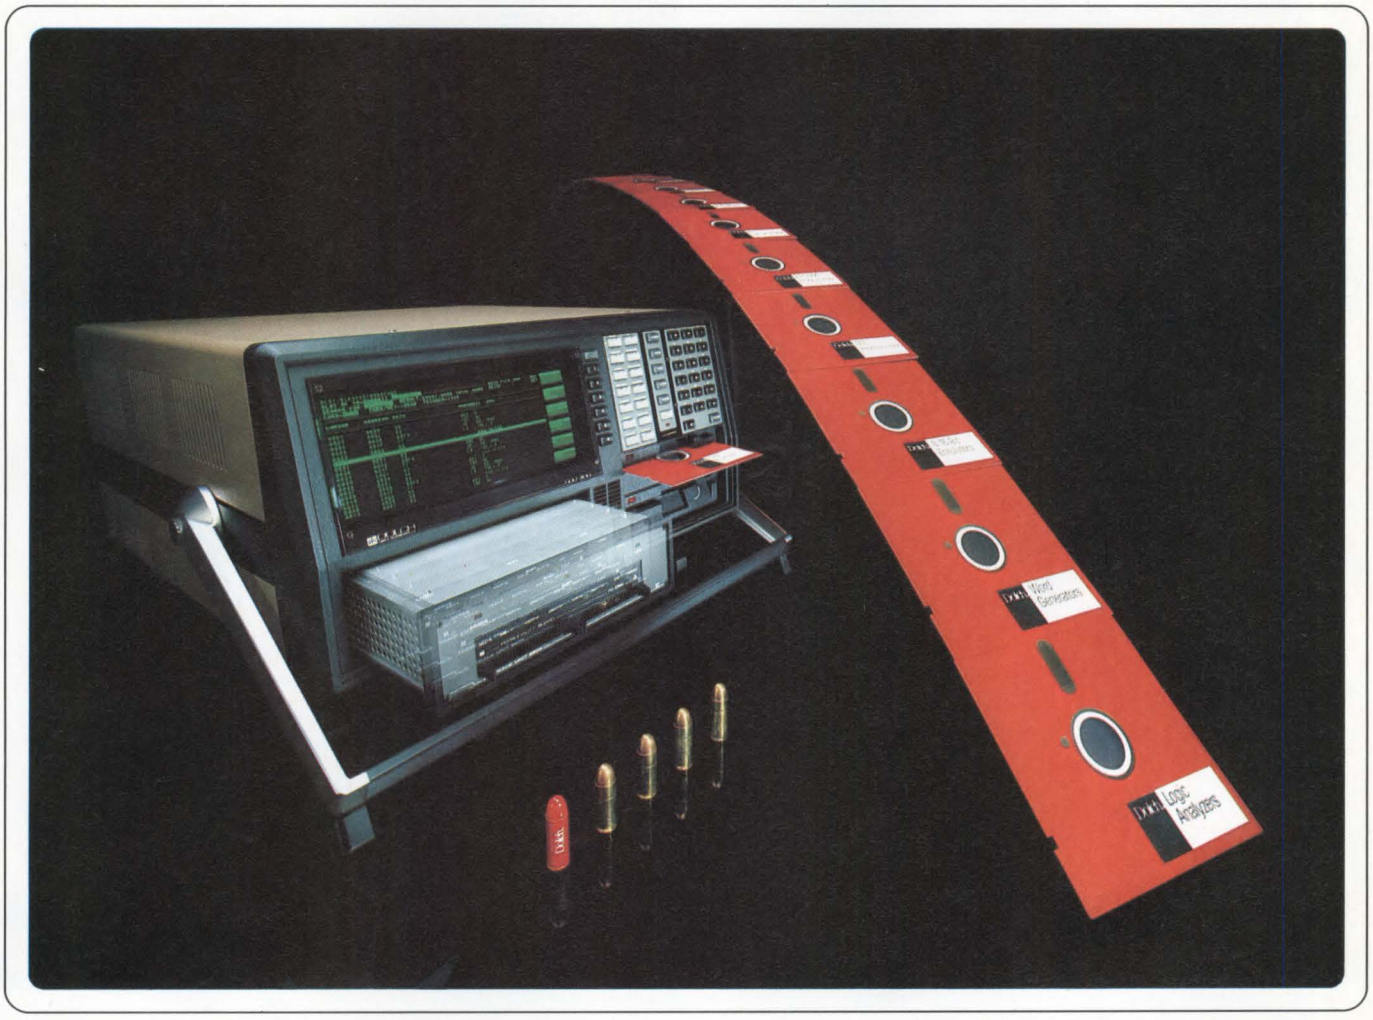

More Bang For Your Bug

Plug-Ins give you the ultimate in flexibility
Dolch's COLT/ATLAS concept is the most logical step into the future of logic analysis and digital test instrumentation. The integration of a general purpose computer with an instrumentation Plug-In creates the ultimate virtual instrument, providing flexibility for today's debugging needs and expandability for tomorrow's logic testing.

Plug-Ins expand your logic analyzer power
As a pioneer in the field of logic analyzers Dolch provides you with the most complete range of logic analyzer plug-in modules for state (48 channels) or timing (300MHz). μP-dedicated analyzers provide you with the ultimate logic analyzer power, carefully tailored to fit the needs for all popular 8 or 16 bitters.

Load up your COLT today and shoot down system bugs
with Plug-Ins now available for:

- Logic Analysis
- In-Circuit Emulation
- Word Generation
- Serial Data Analysis
- In-Circuit Testing
- PAL/IFL Development
- EPROM Programming

For more Information: CIRCLE 22

Write or call for a demonstration today
2029 O'Toole Avenue, San Jose, CA 95131
(800)538-7506 (408)945-1881 (in California)

dli DOLCH LOGIC INSTRUMENTS

Kontron has integrated the micro design tools, so your engineers don't have to.

With the time you save, you'll

Engineers can work together more efficiently and design better products when their tools work together —which is one reason why Kontron's integrated Tool Set for microprocessor-based product development makes so much sense.

Full integration and networking increase productivity

It has a single, integrated operating environment and database, so everyone operates from the same information, whether they're designing hardware, programming, or producing documentation.

It's fully networkable, so if you add new team members you can easily expand the system. All the specs, memos, and code in other people's machines are immediately available, so the new people can become immediately productive.

Comprehensive CAD software

Comprehensive CAD software not only eliminates hundreds of man hours per project (by eliminating tedious, repetitive drafting tasks); it also brings gate array designs within the range of projects where only a few hundred units will be built.

Faster information flow

In addition to all these engineering features, the Kontron Tool Set offers convenience features like word processing and electronic mail that simplify information flow. It's modular in design, and all the tools can be used independently. You can buy just those capabilities you need and expand as required, without being locked into a particular approach.

Complete product line

The Kontron Tool Set includes UNIX®- and CP/M®-based microprocessor development workstations; interfaces to IBM PCs and VAX computers via Kontron's TEAMWORK BUS™; logic analyzers with up to 64 channels and 100 MHz capture capability; real time in-circuit emulators for 8-

TYPICAL DESIGN CYCLE

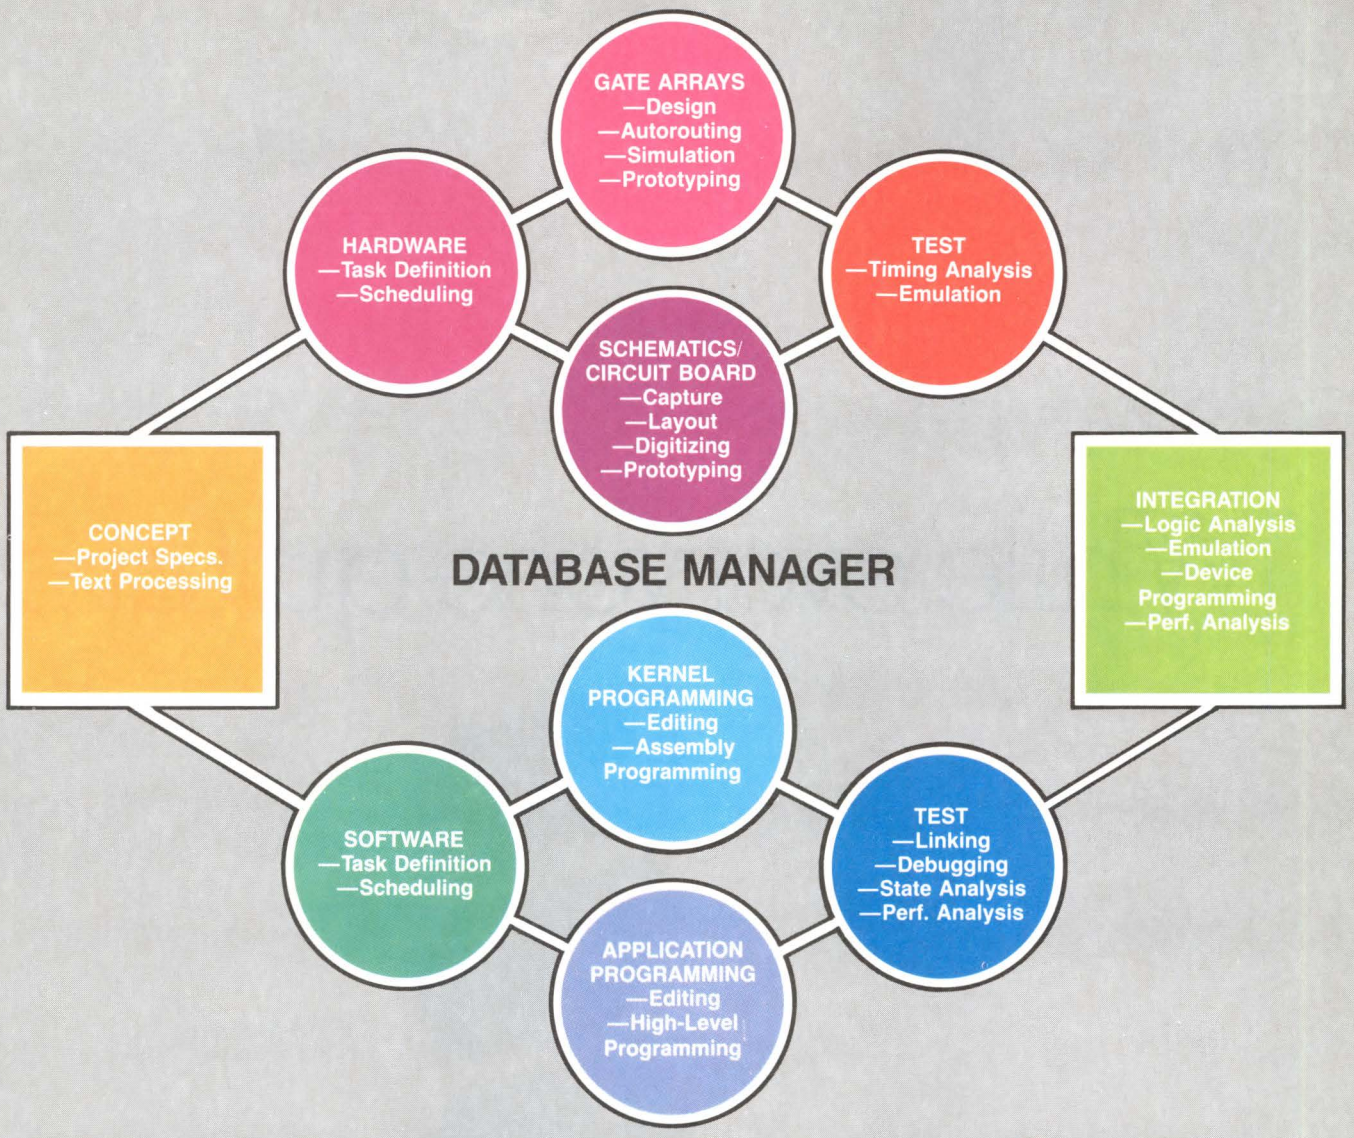

get better products—faster.

and 16-bit microprocessors; PROM programmers for over 500 devices; and MICAD software for computer-aided design of both PC boards and gate arrays.

Write for a 32-page booklet with complete technical information:

KONTRON ELECTRONICS INC.
630 Price Avenue
Redwood City, CA 94063
800/227-8834
(415/361-1012 in California)

D-8057 Eching/Munich,
West Germany, Breslauer Str. 2,
089/3 1901-0.

®UNIX is a registered trademark of AT&T.
®CP/M is a registered trademark of Digital Research, Inc.

Circle 23 for Literature
Circle 182 for Demonstration

ADVANCED ELECTRONIC INSTRUMENTATION

KONTRON ELECTRONICS

See our booths 526 & 528 at Midcon

While everyone quit at 5,

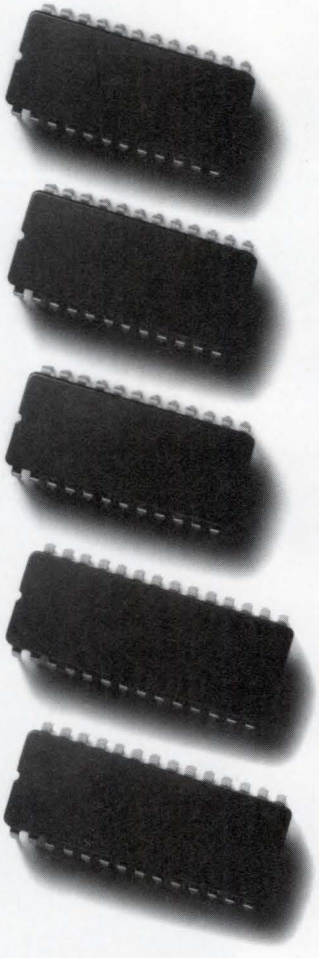

2K x 8 latched

2K x 8 latched timer,
2816A compatible

2K x 8 latched,
full military temp.

2K x 8 latched timer,
2817A compatible

8K x 8 latched timer

It's not that we're gluttons for punishment.

It's just that at Seeq we take the E²ROM business seriously. So seriously that our engineers insist our parts fit your designs, not vice versa.

Maybe they got a little carried away.

At last count, they'd developed 12 different E²ROMs. Outperforming their counterparts at other companies two to one.

And they're responsible for more E² improvements than any other design team.

Most of these innovations quickly became industry standards. Like 16K densities. Single 5-volt power. And leadless-chip-carrier packaging.

The rest are still unduplicated. Like our high-density, high-performance 64Ks. Exclusive DiTrace.™ Our industry-leading 1ms byte-write. And the world's first E² with million cycle endurance.

Plus there are plenty of ideas still on the drawing board. Because when it comes to designing new E²s, our engineers just hate to quit.

Especially while they're ahead.

For information on the industry's most innovative E²ROMs, write or call Seeq Technology, 1849 Fortune Drive, San Jose, California 95131. Telephone (408) 942-1990. Or circle number 11 on the reader service card.

we worked until 12.

2K x 8 latched

2K x 8 latched, 1ms write

2K x 8 latched timer, 2ms write

2K x 8 latched timer, 1M cycle endurance

2K x 8 latched, full military temp.

2K x 8 latched, 2816A compatible

8K x 8 latched

8K x 8 latched, 1ms write

8K x 8 latched, full military temp.

2K x 8 latched timer, 2817A compatible

2K x 8 latched timer, 2817A w/1M cycle endurance

8K x 8 latched, 1ms write, full military temp., 1M cycle endurance

Flat-panel technologies vie to displace CRT in terminals

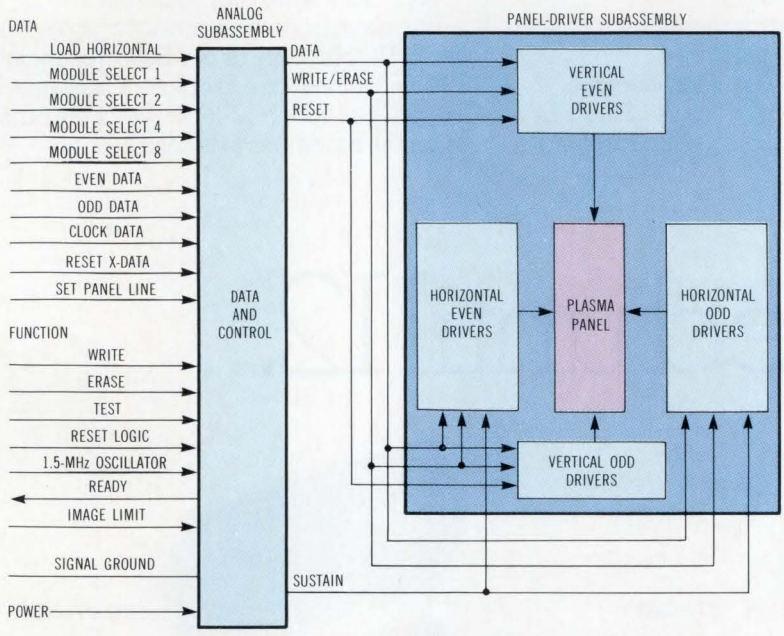

The 10,000-char capacity of IBM's ac plasma panel display implements a subassembly to accommodate all computer interfaces. Ten data-address lines are complemented with five function lines that allow a maximum of 960 pixels in the X direction and 16 pixels in the Y direction at any one time. Both vertical and horizontal lines are broken into odd and even drives to limit the density of interconnections.

Cathode ray tubes have been the predominant display medium for computer terminals since the latter became essential extensions of the computer mainframe. With distributed processing becoming the norm as intelligent workstations embed more local processing power, greater emphasis is being placed on the quality and size of the display. Workstation designers must choose a display that will transmit clear, detailed text and graphics information to the viewer at a high resolution. And, as computers become smaller, it becomes imperative to find an appropriate display alternative to the ubiquitous, bulky CRT.

Research in flat-panel alternatives to the CRT for general-purpose computer workstations has become concentrated in such technologies as gas plasma, electroluminescence, liquid crystals, and flat CRTs. These technologies must, of course, combine the best characteristics of the CRT as a general-purpose display, while sporting a smaller dimension and a CRT competitive price. In addition, the alternatives must supply the desirable characteristics of the CRT—high brightness, good contrast, a stable image, a dense character size, sharp character edges, no glare or extraneous reflections, a fast write/erase feature, and, of course, low cost (see Table, "Worldwide Demand for Large Area Flat-panel Displays").

Currently, the general-purpose workstation display is characterized by a minimum 2000-char content. This is usually divided into 25 lines of 80 chars/line, and spans the entire computer equipment spectrum from mainframe display terminals to personal computer displays.

The flat-panel technology that comes the closest to matching the needed character count is the plasma panel. As opposed to dc plasma displays, ac plasma displays allow information written on the display to remain until erased and rewritten. The average brightness is not governed by duty cycle considerations as with dc panels, and no perceptible flicker exists because the image is sustained at about two orders of magnitude higher than the flicker rate. A recently much-touted ac plasma display is the IBM (Neighborhood Rd,

(continued on page 44)

Worldwide Demand for Large Area Flat-panel Displays ($ millions)

By Technology	1982	1987	1992
LCD	2.3	1282	3060
EL	0	280	715
AC Plasma	2.5	174	530
Flat CRT	1	54	200
Total	5.8	1790	4505

Source: Arthur D. Little, Inc.

68000
16000
8086/88
Z8000

SOFTWARE DEVELOPMENT TOOLS

"One-Stop Shopping"

OASYS provides a **"One-Stop Shopping"** service for software developers and managers in need of proven, cost effective, **cross-** and **native-** development tools.

OASYS can save you time, energy and money! We understand what it means to be a developer. Over the past 3 years, we've built over 1MB of working code.

We not only develop our own tools, but also specialize in evaluating, selecting and distributing the best complementary tools from other suppliers.

Our tools are currently in use in over **1,000** installations worldwide on **micro-, mini-,** and **mainframe** computers for a variety of **8-, 16-** and **32-** bit UNIX (and non-UNIX) systems.

Most likely, we have what you're looking for (even if it doesn't appear in the tables shown). But, if we don't, we'll be glad to tell you who does.

So, call or write today for more information and start shopping the smart way, the fast way, the economical way.

"The One-Stop Shopping Way."

OASYS

60 ABERDEEN AVENUE
CAMBRIDGE, MA 02138
(617) 491-4180

CROSS TOOLS

PRODUCTS (1)	HOST (2)	TARGET (3)
C COMPILERS	VAX, PRIME	68000 16000 8086/88
PASCAL COMPILERS	VAX PDP-11, LSI-11 PRIME	68000 16000 8086/88
FORTRAN COMPILERS	VAX PDP-11, LSI-11	68000 16000 8086/88
(4) ASSEMBLERS	VAX, PDP-11, LSI-11, PRIME, IBM/PC, IBM 370	68000, 16000, 8086/88, Z8000, 680X, 808X, Z80
SIMULATORS	VAX, PDP-11 LSI-11, PRIME, IBM/PC, IBM 370	68000, 8086/88 808X, Z80

(1) WE DISTRIBUTE PRODUCTS FOR: GREEN HILLS SOFTWARE, VIRTUAL SYSTEMS, COMPLETE SOFTWARE, PACER SOFTWARE; SOFTWARE MANUFACTURERS
(2) HOST OPERATING SYSTEMS INCLUDE: VMS, RSX, RT-11, PRIMOS, UNIX V7, III, V, BSD 4.1, 4.2, UNOS, IDRIS, XENIX, MS/DOS, VM/CMS, CPM 68K
(3) OTHER TARGETS ARE: M6801-6803, 6805, 6809, 8080, 85, 28, 35, 48, 51; Z-80
(4) ALL ASSEMBLERS INCLUDE LINKER, LIBRARIAN AND CROSS-REFERENCE FACILITY
(5) AVAILABLE ON: CALLAN, OMNIBYTE, CHARLES RIVER DATA, PLEXUS, SAGE, FORTUNE, WICAT . . . to name a few.

C/UNIX NATIVE TOOLS

- **NATIVE ASSEMBLERS FOR 68000s** (4, 5)
- **SYMBOLIC C SOURCE CODE DEBUGGER**
- **C-TIME PERFORMANCE UTILITY**
- **UP/DOWN LINE LOAD UTILITIES**
- **COMMUNICATION UTILITIES**
- **C-BASED FLOATING POINT MATH PACKAGE**
- **BASIC-TO-C TRANSLATOR** . . . AND MORE

Flat-panel technologies

(continued from page 42)

Kingston, NY 12401) 581 Plasma Display. The developers report a 1000-char capacity. Its 960 vertical lines and 768 horizontal lines at a 71.4-line/in. resolution yield a pixel count of 74,000 that can be drawn on an approximately 1 ft^2 viewing area of the panel. It takes 200 ms to fill a screen. To achieve this large pixel count, IBM researchers minimized the panel process defect densities by using projection printing and low temperature glass materials. They also found an effective way to repair open lines.

On the 581, the vertical lines are brought out half on the top and half on the bottom of the panel as even and odd lines. The horizontal lines are handled in a similar fashion on the right and left sides of the panel. This is done to reduce the density of the interconnections to the panel and affects the architecture of the interface. The 581 interface consists of 10 data-addressing lines, five functions, and two lines with signals sent to the using system.

IBM researchers have operated the 581 display with the IBM PC to test the display's feasibility as a substitute for the CRT monitor. A 720-x 350-pixel element section of the 960 x 768 plasma panel was driven at a 50-frame/s refresh rate by the IBM PC's CRT adapter with three intensity levels (normal, dim, and off). The panel's video interface responded to the vertical and horizontal syncs, video, intensity, and clock signals originating from the PC's monochrome adapter. The standard IBM PC CRT monitor and the plasma display monitor shared the same interface signals. The CRT monitor was used throughout the project to generate and verify test patterns while the plasma was operational.

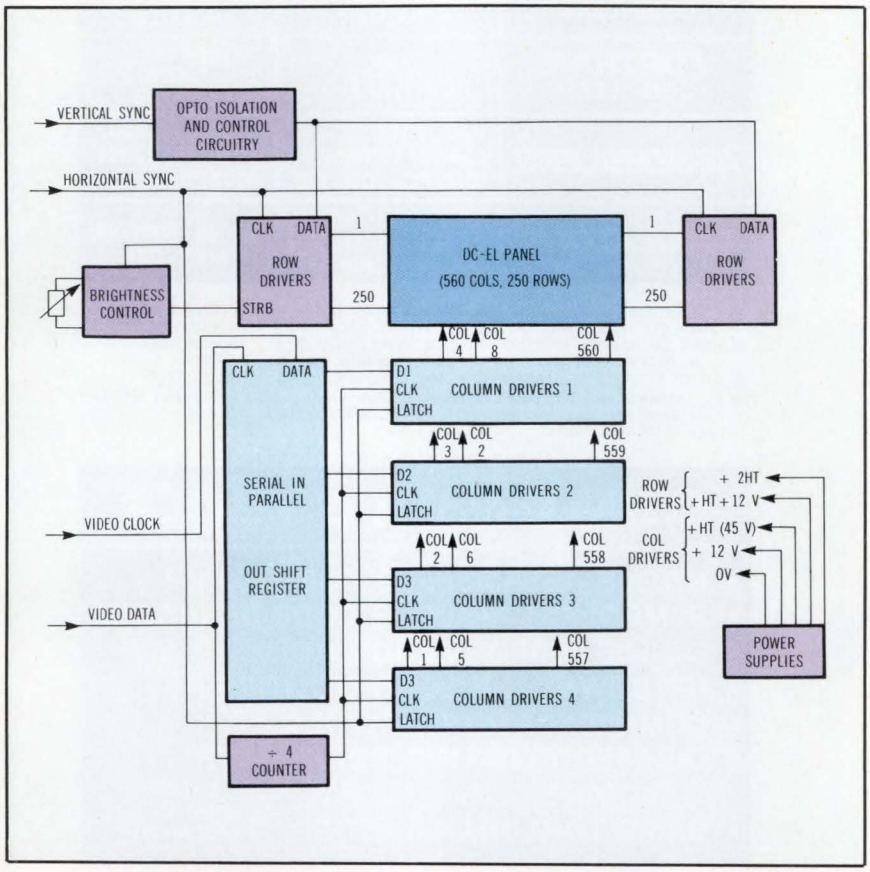

In this 2000-char dc electroluminescent panel, the video interface compensates for the fast 14-MHz data rate by splitting the data into four parallel streams via a shift register. Four 140-bit column shift registers are then operated on by one of the four data streams. The row scan is done by shifting 1 bit through a 250-bit shift register.

Replacing the CRT

IBM's goal was to prove that an ac plasma display would be a more flexible device if it could be made to operate from the video interface as well as from an XY address interface. The results of the experiment were that the ac plasma panel can be used to replace the CRT as a display component in a computer terminal or monitor, and that in the video interface mode, it operated better than it does in XY address mode. IBM does not expect the plasma display to dislodge the CRT in all computer applications since the production costs of ac plasma panels are not competitive with those of CRTs. But, for generating flicker-free, high quality images on a flat panel, the ac plasma panel can be a practical alternative.

Another flat-panel alternative to the CRT is also making large strides to be a CRT replacement. Plasma Graphics Corp (PO Box 1403, Warren, NJ 07060) engineers combined features of dc and ac plasma technologies into their Self-Scan memory plasma display. This model 120 panel combines the inherent memory features of ac plasma and the efficient addressing of dc plasma to yield a 12,000-addressable pixel display arranged in 480 cols x 250 rows. The hybrid plasma display effectively increases the luminance over conventional dc or ac plasma displays.

Addressed by a 5- x 7-char matrix, the panel displays 25 lines of 80 chars with all the necessary electronics to address and time the display housed in an integral drive board. A number of computer terminal manufacturers are evaluating this product.

Meanwhile, electroluminescent flat-panel displays have matured to the point of product applications. This increased activity has been greatly assisted by the development of custom LSI drivers using double-diffused MOS technology to meet the

(continued on page 46)

GOOD & PLENTY

The Fujitsu MB81416 16K X 4 DRAM. Now available in quantity.

Nobody puts tastier technology into a DRAM than Fujitsu.

Like the fastest access time available – 100ns. A TCAC/TRAC ratio of 2:1 for easier timing. Low power consumption, too.

And in a variety of packages. Like plastic DIP. JEDEC-approved compatible with the TMS 4416 and INMOS 2620.

Mmmm mmmm, good.

Best of all, we've got plenty in stock. So call the nearest FMI sales office listed below. For immediate gratification.

	100ns	120ns	150ns	Page Mode	TRAC/ TCAC	CAS before RAS refresh
FUJITSU	yes	yes	yes	yes	2:1	yes
TI	no	no	yes	yes	1.66:1	no
INMOS	yes	yes	yes	no	1.66:1	yes

FUJITSU MICROELECTRONICS, INC.

Technology that works.

3320 Scott Boulevard, Santa Clara, CA 95051 · 408/727-1700

FMI Sales Offices. Atlanta 404/449-8539. Austin 512/343-0320. Boston 617/964-7080. Chicago 312/885-1500. Dallas 214/669-1616. Houston 713/784-7111. Minneapolis 612/454-0323. New York 516/361-6565. No. Calif. 408/866-5600. S. Calif. 714/720-9688.

CIRCLE 25

Flat-panel technologies

(continued from page 44)

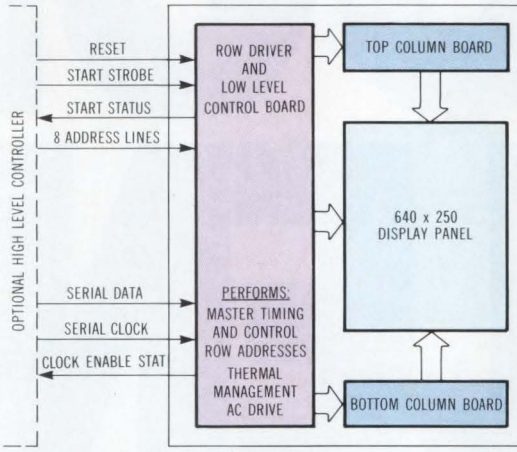

The LCD flat panel has the 640 odd and even bits clocked serially from two column driver boards. At 1 MHz, the loading is accomplished in 640 μs. After loading the shift registers with data, supplying an 8-bit row address, and enabling the controller to start the cycle; a writing speed of 2000 chars/s is achieved.

high voltage requirements. Electroluminescent displays fall into two categories depending on the materials and fabrication technology implemented—powder EL, and thin-film EL. Researchers at Phosphor Powder Co, Ltd, (Poole, England) have developed a powder direct current EL based on powder phosphor. It displays a matrix of 560 cols x 250 row electrodes and accommodates 2000 chars, the bare minimum to be a serious contender for CRT replacement in computer terminals. Compatible with video signals, the panel displays 25 rows x 80 chars at a brightness of 15 to 20 ft-L at 100 to 120 V, and a 60-Hz frame rate.

The basic drive philosophy that the English researchers used was to sequentially scan the smaller number of electrodes while feeding the data into the orthogonal electrode set. Both types of drives can be achieved using serial shift registers. The 60-MHz frame rate has a display driven with 67-μs pulse and a duty cycle of 0.4 percent. Thus, the 250 bits are loaded into each column shift register in 67 μs, giving a data rate of 3.5 MHz.

Typically, the display requires 100 V to drive it; this is split between rows and columns. Therefore, any drivers for this display must work at 3.5 MHz, be capable of greater than 50 V outputs, and both sink and source current transients. Of the drivers currently available, Phosphor Products' engineers consider the Texas Instruments SN75553 closest to meeting these requirements. Thus, they have chosen it to drive the display. The company's engineers have solved one of the major problems in designing a high resolution dot-matrix display—proper connection between the drivers and the display substrate. They devised flexible circuits that have gold-plated tracks at both ends. These match the pitch of the indium tin oxide electrode pads on the substrate and corresponding connection fingers on the drive PC board assembly. These are held in place using U-shaped sprung plastic strips and a compression rubber filler. This provides the appropriate pitch contact to achieve high resolution.

A larger display

While liquid crystal flat panels have usually been restricted to small applications, CrystalVision, Inc (Sunnyvale, Calif) has recently developed an LCD that has 256 rows x 640 cols, and can typically be used in a 25-line x 80-char format. Although the current writing speed is roughly 2000 chars/s, special cases where the information is repetitive (eg, displaying the grid for a graph or erasing the screen), can approach a writing speed of 1000 chars/s.

Basically, the display consists of two parallel glass plates separated by approximately 165 μm, which

(continued on page 48)

PLESSEY JUST GAVE VMEbus WHAT IT NEEDS MOST. PLESSEY.

Introducing a full line of board level hardware, software and technical support for all your VMEbus systems.
Take a bus as powerful, versatile, reliable and future compatible as VME. Add a company as powerful, versatile, reliable and future compatible as Plessey and you have everything it takes to board the VMEbus with complete confidence.

Introducing Plessey VME. Plessey has a full range of VME board level hardware, software and technical support services. We'll give you the languages, operating systems, the drivers, complete technical assistance and, of course, that famed Plessey guarantee with worldwide backup. All to support our superior VMEbus products including . . .

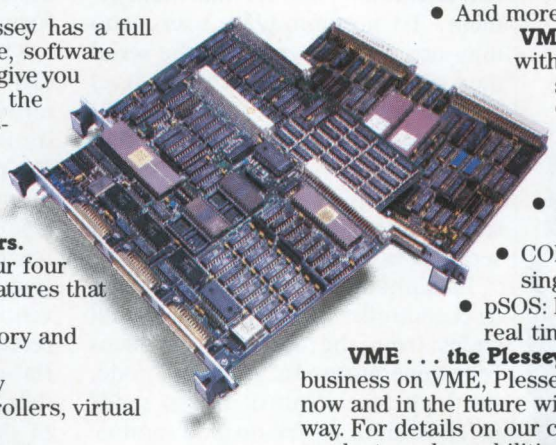

VMEbus Single Board Computers.
Build your systems around any of our four new single board computers with features that include up to:

- 10 MHz 68000, 512 Kbytes of memory and 128 Kbytes EPROM
- 512 Kbytes of dual ported memory
- Memory management, DMA controllers, virtual memory processors
- 3 serial ports, multi protocol serial I/O
- 24-bit bidirectional parallel I/O
- Floppy disk controller, real time clock with battery backup
- And a lot more.

VMEbus Memories. Enhance your system with add-in memories including:

- Parity boards up to 4 Mbytes capacity with 270 ns access time
- ECC boards up to 3 Mbytes capacity with 300 ns access time.
- Static RAM/EPROM boards with up to 256 Kbytes of both RAM and EPROM in 16 JEDEC standard sockets
- And up to 128 Kbytes high speed static RAM boards with 140 ns access time and on-board battery backup.

VMEbus Controllers, Graphics and I/O Boards.
Plessey VME also includes:

- Winchester / floppy disk controllers
- Intelligent SASI controllers
- 16-color graphics boards
- 6-channel serial I/O boards
- And more coming.

VME Software. Support your system with our wide range of firmware and software products:

- Monitor: EPROM resident monitor/debugger
- Ideal: EPROM resident assembler/editor
- Basic and Forth language compiler/interpreters
- COHERENT*: UNIX† V7 compatible single/multi-user operating system
- pSOS: EPROM resident multi-tasking real time executive.

VME . . . the Plessey Way. If you are building your business on VME, Plessey VME means business. Right now and in the future with new VME products on the way. For details on our comprehensive VMEbus products and capabilities, call or write Plessey Microsystems, One Blue Hill Plaza, Pearl River, NY 10965. (914) 735-4661 or toll-free **(800) 368-2738.** Or use the coupon below.

*COHERENT is a trademark of the Mark Williams Corp.
†UNIX is a trademark of Bell Laboratories

PLESSEY MICROSYSTEMS

For additional information, clip and mail to Plessey Microsystems, One Blue Hill Plaza, Pearl River, NY 10965.

Name
Company
Address
City
State
Phone
Zip
Title

CIRCLE 27

Flat-panel technologies

(continued from page 46)

Enhanced versus Conventional CRTs

	15 in. Full Square	14 in. Conventional
Screen dot-pitch	0.39 mm	0.39 mm
Display area	260 x 190 mm	250 x 180 mm
Display capability pixels/line	617	593
Brightness	200 percent	100 percent
Ambient reflectance	69 percent	100 percent
Corner legibility	115 percent	100 percent

contains a smectic liquid crystal and dye mixture. The front glass has transparent indium tin oxide column traces running vertically. The back glass has textured, mirror-finish metallic row heating traces running horizontally. A pixel is created at the intersection of each of the 640 cols and 250 rows, for a total of 16,000 pixels. A wide border column on each side and six border rows add another 4352 controllable pixels, allowing the user to dynamically control the border area.

Low level display functions are performed on the panel driver board microprocessor. These functions include general display control and timing, ac waveform generation, pixel row addressing, and thermal management. To accommodate users who may not want to generate the serial data stream required by the low level control board, CrystalVision engineers have designed an optional high level control board that accepts both ASCII characters and graphics bit patterns. Users can communicate via a serial or a parallel port.

In alphanumeric mode, the high level controller accepts standard ASCII codes from the user and performs font generation. In graphics mode, the high level control board sends received data to the display module without performing font generation. In this mode, users have control over every pixel in the display by specifying the bit pattern and the row address where writing is desired. Rows can also be written in any sequence.

Meanwhile, the quest for a flat-panel alternative continues. That quest was begun as soon as the CRT was invented. Many corporate laboratories have ongoing research projects to reduce depth dimension of the CRT. So far, these projects have brought limited results, and most efforts are slanted toward consumer television applications. In addition, a perennial problem has been the curvature of the CRT's faceplate. Distortion of images and characters, as well as eye strain, have resulted from this.

Toshiba Corp's (Fukaya, Japan) development of a "full square" CRT has enhanced the conventional CRT. It has a flatter surface, which decreases the needed ambient light reflection by 30 percent, and it offers a 15 percent better corner legibility. As a result, the display is twice as bright as a conventional CRT. Also, the 15-in. display uses a superior deflection yoke that yields better focus than conventional tubes.

Thus, flat-panel alternatives to the CRT could soon become standard fare in personal computers and workstations. The market for large area flat-panel displays, at least according to one research firm—A. D. Little in Cambridge, Mass—is estimated to reach $4.5 billion by 1992. This data (see Table, "Enhanced versus Conventional CRTs") is taken from a recent worldwide study on large area flat-panel displays authored by Little's Dr Charles M. Apt. The breakup of the market into its essential elements—gas, plasma, electroluminescence, liquid crystal, and flat-panel CRT—leaves enough room for all interested parties to reap potential benefits from developing one of these four contenders to replace the conventional CRT in computer terminals.

—Nicolas Mokhoff,
Senior Editor
SYSTEM TECHNOLOGY

(continued on page 61)

BR 164

VMEsystem

VME VME VME

VME VME VME

Today's rugged standard guarantees
tomorrow's compatible growth.

MOTOROLA INC.

VMEsystem

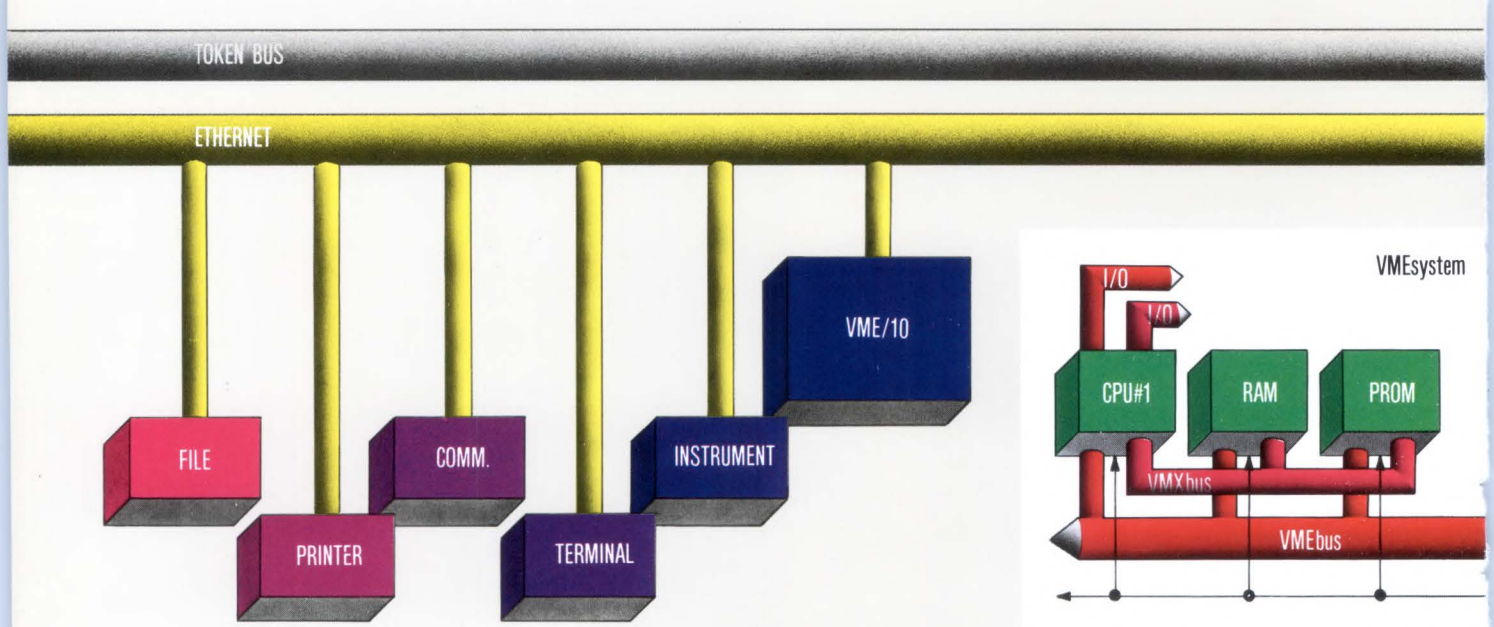

The VMEbus is part of the Motorola
porated into the VME module system and the

Your high-tech future is speeding toward you requiring you to combine technical and business considerations when making your design-in decisions. As a systems integrator, you face a major decision in your choice of a microcomputer system architecture.

Major because in order to utilize 32-bit microprocessors or 32-bit microprocessor based boards in the next five years you must choose between high performance system buses that support 32-bit machine functionality. Your decision will last throughout the computer system's lifetime. VMEsystem Architecture is the systems designers leading choice in providing high performance support of 32-bit designs, according to the Electronics News 1984 MPU Based Processor Study.

VMEsystem leads also with the broadest line of available modules and support products for 8/16/32-bit applications. With over 400 products now available from over 90 manufacturers. The line continues to expand while providing smooth, planned growth for the future.

Rugged VMEmodules and standardized VMEsystem architecture coupled with a broad base of manufacturer and vendor support guarantee tomorrow's compatible growth.

VMEbus: *the heart of VMEsystem.*

Originally introduced in 1981, VMEbus has become the standard high performance system interconnect structure through several years of intensive design activity. It is the bus of choice for the present and future of microcomputer architecture, supporting data transfer rates as high as 30Mbytes per second with today's chip technology.

VMEbus has a master/slave asynchronous non-multiplexed data transfer structure, seven levels of priority interrupt, four levels of arbitration with rapid fault detection and control for bus, system and AC failures.

One of the many features of the 32-bit configured VMEbus system is that the bus dynamically senses whether 8-, 16-, or 32-bit data paths are needed and adjusts automatically.

VMEbus specifications were originally developed jointly by Motorola, Mostek, and Signetics/Philips and have been accepted by close to 100 manufacturers worldwide.

In integrating 8-, 16-, and 32-bit system components VMEbus is innovative, publicly documented, and eminently adaptable to new technologies. LSI/VLSI technology is providing interface and peripheral chip functions that vastly increase the functionality/cost ratio of VMEbus modules.

Based on the most popular Eurocard formats with DIN pin and socket connectors, VMEbus is presently being formally standardized by both the IEEE (P1014) and the IEC (47B) standards organizations.

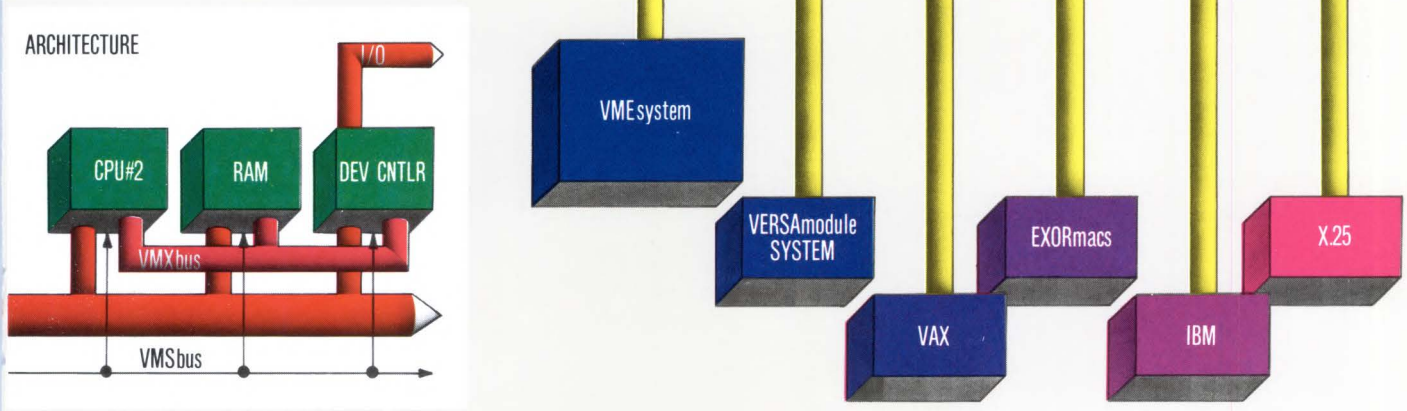

Microsystem Architecture. VMEbus is incor-
VME/10 microcomputer system.

To the hundreds of present VMEbus users this is a formality. VMEbus has earned its position as the standard through hard usage, not through mere speculation based on paper specifications.

VMEbus is designed from the ground up for multi-processing applications while still providing a high level of fault tolerance. Consider its full 32-bit address and data paths, clear arbitration and interrupt schemes and diagnostic capability.

VMEbus is also *the* cost-effective system bus. It can be economically configured according to design requirements for multiple processors and controllers. Manufacturers now provide VMEbus compatible modules for many applications.

VMXbus: *high speed private access.*

Designed for multiprocessing applications, VMXbus is a high speed, parallel, private access bus that allows local operation between the MPU and memory. Its local interconnection between processor and memory takes the load off the main bus for rapid, efficient multiprocessing.

In communicating to I/O or other functions over the dedicated VMXbus, the processor can carry out operations independent of the VMEbus.

VMXbus has a 32-bit data path with a 16M byte address range. It handles up to six functional modules per "group" and runs on either a backplane or cable.

VMSbus: *self arbitrating, global access.*

VMSbus is designed as a self arbitrating, serial global bus for short messages between processors while allowing for diagnosis and fault tolerance.

Its 3.2 Mbit/second transfer rate allows rapid communication of brief messages between system modules using only two conductors (clock and data) to provide an efficient "party line" between system components. It provides urgent event message passing in both tightly and loosely coupled multiprocessing environments.

I/O Channel: *Motorola enhancement*

The Motorola I/O Channel provides low cost modular I/O expansion on a local processor bus not tied to a specific board or enclosure.

It permits interconnection of slower peripherals up to two megabytes per second directly with their respective processor, freeing the VMEbus to handle simultaneous high speed data exchange and multiprocessor activities.

I/O Channel utilizes 12-bit address with 8-bit data bus and 4K byte memory mapped I/O. Its asynchronous operation combines 4 levels of interrupts. Modules can be selected for a wide range of requirements including analog conversion, discrete parallel, serial I/O, plus mass storage.

VMEsystem vs. MULTIBUS II

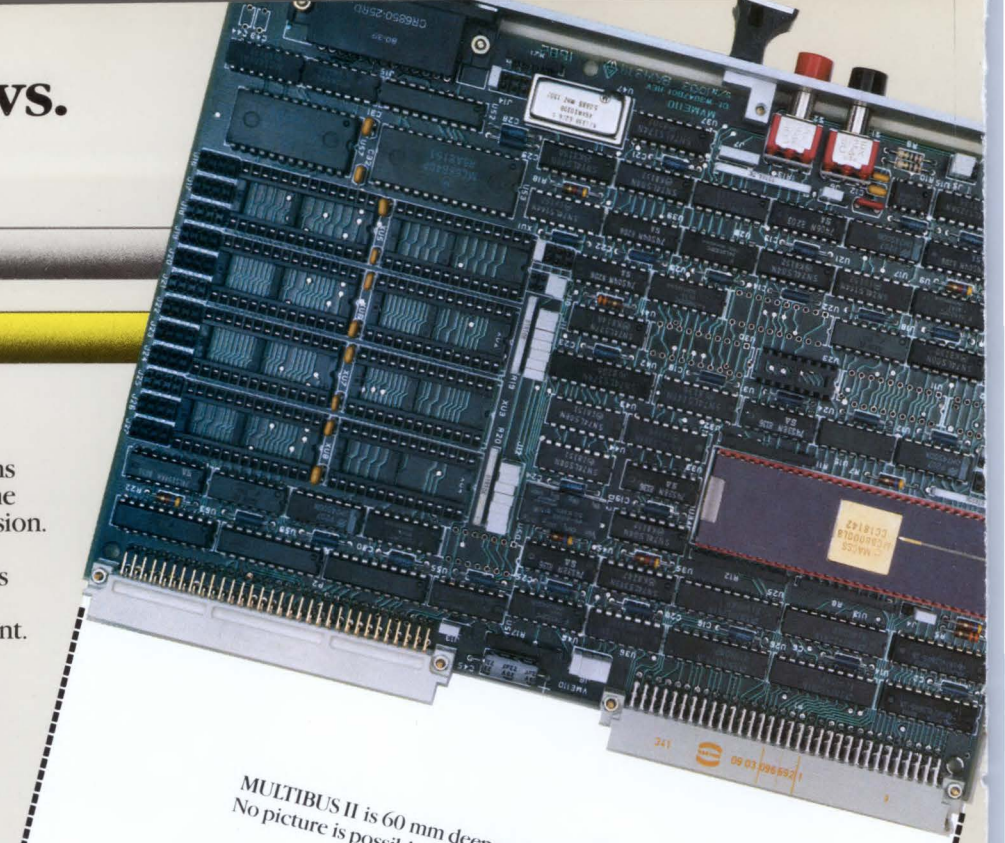

MULTIBUS II is 60 mm deeper than VMEbus.
No picture is possible since MULTIBUS II is not yet available.

There are significant business reasons for the systems integrator to make the Motorola VMEsystem design-in decision. In a comparison with the proposed MULTIBUS II the business advantages of wide industry support for VMEsystem Architecture become apparent.

Proprietary constraints.

VMEbus carries no patents or trademarks on the bus nor are the bus specifications copyrighted. A sound business advantage.

There are trademarks on the MULTIBUS II components, patents pending on MULTIBUS II, copyrights on all specifications and licensing and fees required by Intel to use MULTIBUS II.

The VME system is free of licensing fees and other proprietary restraints. A truly open system geared for broadbased use *today.*

Reliability and the migration path.

The VMEbus has shown no significant problems since being originally released in 1981. It owes much of this reliability to the orderly migration path from its dependable precursor VERSAbus.

MULTIBUS II as specified is synchronous and highly multiplexed while its predecessor MULTIBUS I is neither. There is no orderly migration path.

VMEbus Manufacturers Group: getting together to promote compatibility.

The VMEbus Manufacturers Group is an international association of independent manufacturers who support the VMEsystem standard. The purpose of the group is to create a forum for technical exchange to promote VMEsystem as a compatible industry standard.

Key goals are to maintain and enhance the standard, encourage other independent manufacturers to produce compatible products and provide identification of compatible equipment.

The active VMEbus Manufacturers Group has a mail list of over 500 individuals and convenes 4 times a year. The Group publishes the VMEbus Compatible Products Directory listing vendors, addresses and products. To receive a copy of the Directory check the appropriate box on the request card.

The Manufacturers Group also administers the VMEsystem specifications. A comprehensive VMEsystem Architecture Manual with complete specifications is available. Check the appropriate box on the response card.

Users Group: talking it over.

The VMEbus Users Group is a rapidly growing voluntary organization coordinating user meetings in North America and corresponding with affiliates in Europe.

There are active sub-committees on product compatibility, application support and distribution media. The group publishes application notes, and encourages detailed presentations on VME applications at its meetings.

For membership information on the VME Manufacturers or Users Groups check the appropriate box on the request card.

Future VLSI: The world continues to get smaller.

The major movers behind VMEsystem Architecture have announced additional bus support chips in various stages of planning and production;

Number	Function	Company
68452	Bus Arbiter	Motorola
68153	Bus Interrupter	Motorola
68154	Bus Interrupter	Signetics
68155	Interrupt Handler	Signetics
68172	Master/Slave Interface	Signetics
68174	Bus Arbiter and Clock Driver	Motorola
68171	VMSbus Controller	Signetics
68173	VMSbus Interface	Signetics

LSI is not available for competitive 32-bit buses. VMEsystem is well designed, well defined, solidly supported and publicly available.

The reasons are clear for a Motorola design-in decision.

VMEmodules: Today's rugged standard.

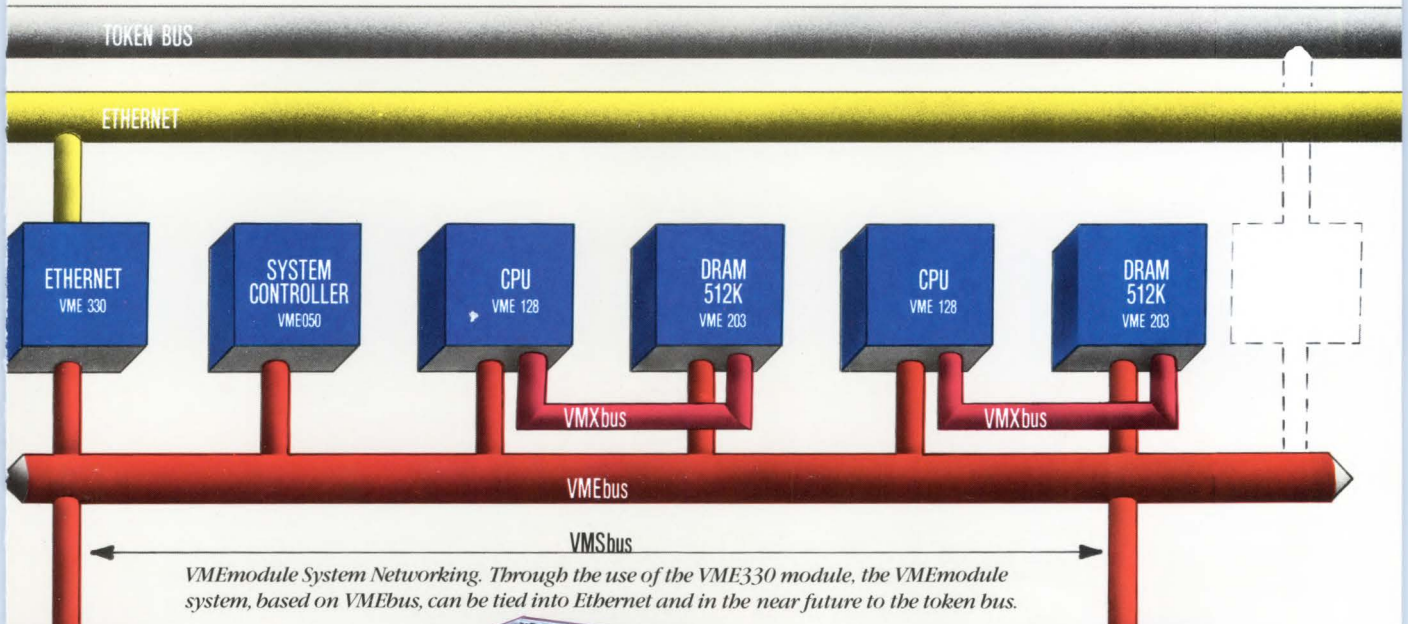

VMEmodule System Networking. Through the use of the VME330 module, the VMEmodule system, based on VMEbus, can be tied into Ethernet and in the near future to the token bus.

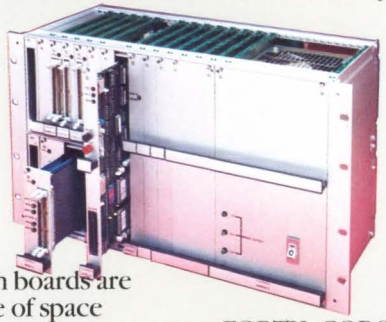

IEEE 488
VME 300

I/O CHANNEL
INTERFACE
VME 316

ADC
VME 600

DAC
VME 605

AC SENSE
VME 610

DC SWITCH
VME 625

Adopting the popular Eurocard mechanical format, VMEmodules™ are available in two pin-out compatible sizes. Single-high and double-high boards are used to afford a high degree of space flexibility for system configurations. Pin/socket connectors, card cages, backplanes, and plug-in power supplies all meet rigid DIN and IEC standards.

The use of the Eurocard format insures world-wide acceptance especially in Europe where the advantages of this circuit card standard have been utilized for over a decade. Superior performance by pin/socket connectors is especially evident in applications where resistance to shock and vibration is important, including industrial automation and mobile electronic systems.

VME products by host MPU and application

Single Board Computers – 68000, 68008, 68010, 16032, 32032, Z80, Z8000, 80186

Operating Systems – UNIX, CPM, VERSAdos, PVS, UNIPLUS, DSSDOS, ADAX, PDOS, MSP168000, VRTX, MTOS-68K, PSOS-68K

Languages – Pascal, C, FORTRAN, LISP, FORTH, COBOL, BASIC

Major Applications – Development, Small Business, Process Automation, Multiuser, Evaluation, Digital Communications, Image Processing, CAD/CAM/CAE, Data Acquisition, Transactions Processing, Automated Test/Inspection, Robotics

Memory – dynamic, static, PROM, EPROM

Mass Storage – floppy, Winchester, SMD, SASI

I/O Expansion – serial, parallel

Peripheral Controllers – IEEE-488, GPIB, video, color graphics, Local Area Network (LAN), speech synthesizers

Industrial Controllers – A/D, opto-coupled, stepper motor, power I/O, relay, DC coupled

Bus Adapters – to VERSAbus, to MULTIBUS I

Other Boards – SCM, floating, point CPU, BASIC interpreter, real time executives, bus debug, EPROM programmer

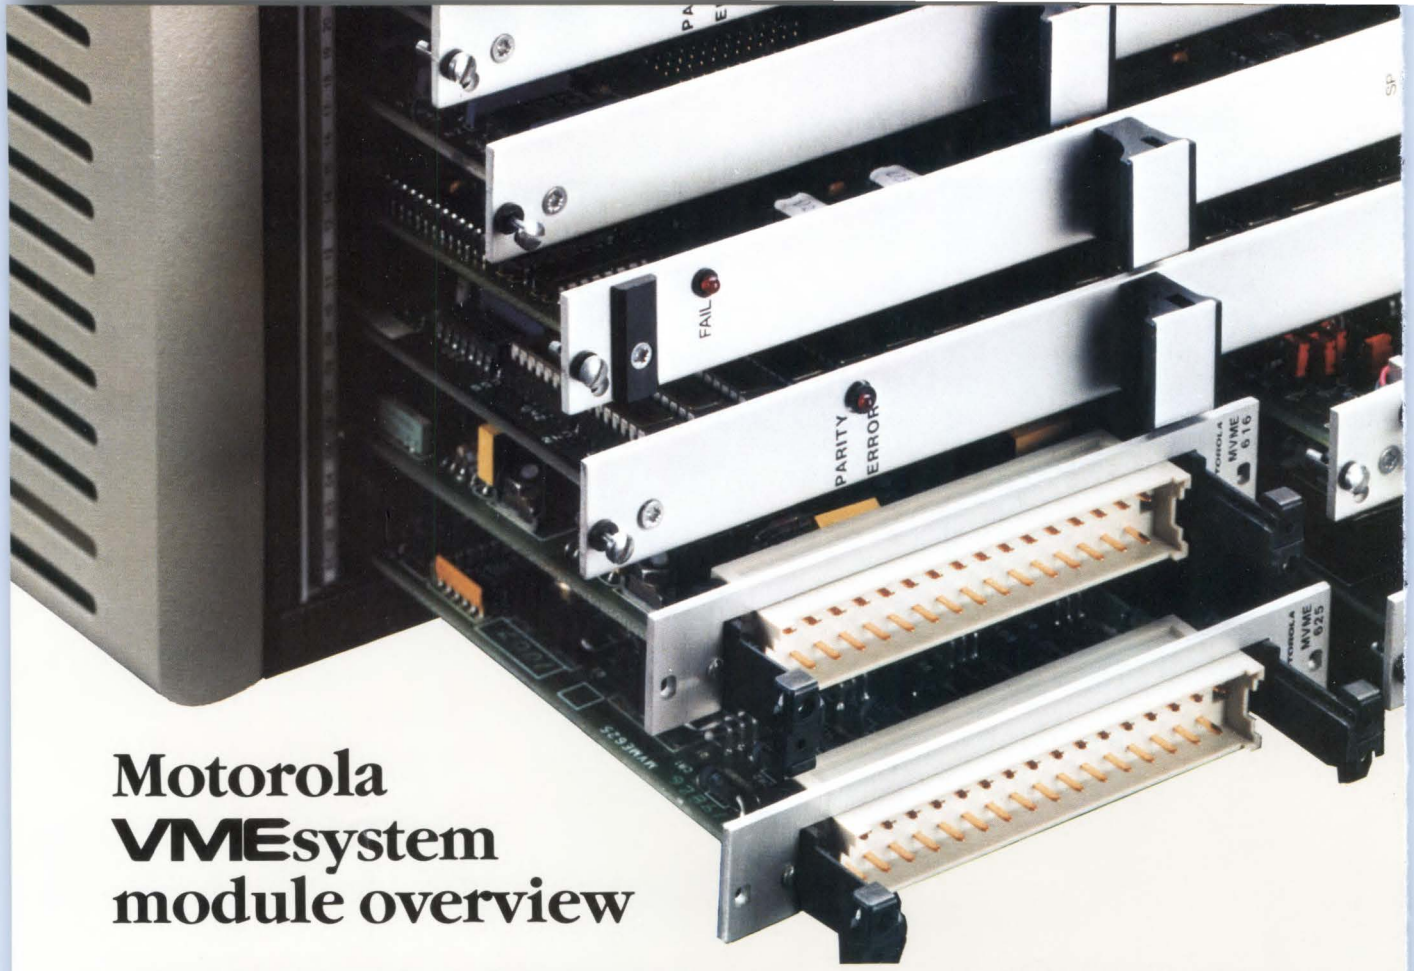

Motorola VMEsystem module overview

100 Series–Monoboard Computers and System Controllers

MVME 101 MC68000 (8MHz) monoboard with 2 RS-232C serial ports and 2 parallel ports

MVME 110-1 MC68000 (8MHz) monoboard with RS-232C debug port and I/O channel interface

MVME 115M MC68010 (8MHz) monoboard with MC68451 MMU optional plus 2 RS-232C serial and 2 parallel ports

MVME 120* MC68010 (10MHz) monoboard with 12.5MHz optional. MC68451 MMU optional. 4 KByte cache option with 128/512 KByte dual port RAM, plus RS-232C debug port.

MVME 128* MC68010 (12.5MHz) monoboard with 10 MHz optional; MC68451 MMU optional; 4 KByte instruction cache optional. 256 KByte dual port RAM and VMXbus interface plus RS-232C debug port

MVME 130* MC68020 (16MHz) monoboard with 512 KByte dual port RAM or MC68851 (16MHz) PMMU optional. MC68881 (16MHz) math co-processor optional.

MVME 025 Used With MVME 115 and 120. Provides: bus arbitration, AC fail and system clock functions.

MVME 050* Used with MVME 115, 120, 128 and 130. Provides: interrupt handlers, bus arbitration, system clock, time-of-day clock, plus eight 28-pin sockets.

200 Series-Memory Modules

MVME 200 64K DRAM

MVME 201 256K DRAM

MVME 202/222 512K/2M DRAM

MVME 203/223* 512K/2M DRAM with VMXbus interface

MVME 210 Up to 128 KBytes ROM and up to 32 KBytes static RAM

MVME 211 Up to 1 MByte ROM and up to 128 KByte static RAM

MVME 212* same as MVME 211 with VMXbus interface added.

300 Series-Peripheral Controllers

MVME 300 GPIB controller with up to 500 KByte/Sec transfer rate

MVME 310 Universal IPC with 4 channel DMA

MVME 315 Intelligent SASI interface with floppy controller for 4 5¼" or 8" floppies

MVME 316* I/O channel interface

MVME 320 Winchester/floppy controller for 5¼" drives

MVME 330 Ethernet 2.0 compatible controller with MC68000 (10MHz) and LANCE chip and SIA

MVME 331* 6-Channel serial I/O controller

MVME 340* 6-Channel parallel I/O interface

400 Series-I/O Channel Modules

MVME 400 dual channel serial port

MVME 410 dual channel parallel port and Centronics compatible printer interface

MVME 420 SASI™ interface adapter

MVME 435 ½" magnetic tape adapter

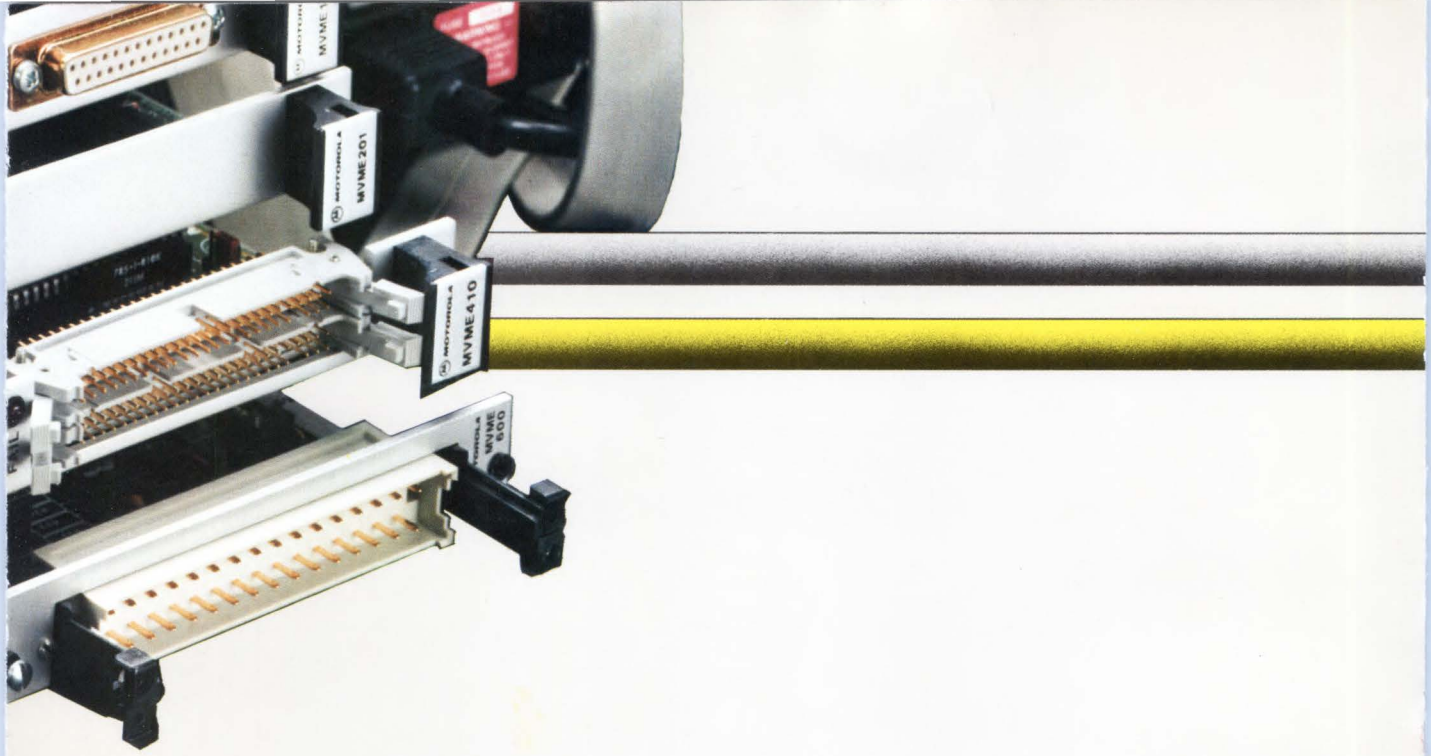

VME/10 Microcomputer can be custom configured through rear-panel plug-in of VMEmodules and I/O modules.

600 Series-Industrial I/O Modules

MVME 600/601 analog/digital converter with 8, 16 input channels

MVME 605 digital/analog converter with 4 output channels

MVME 610 AC sensing with 8 input channels

MVME 615/616 AC switching with 8 output channels

MVME 620 DC sensing with 8 input channels

MVME 625 DC switching with 8 output channels

700 Series-Non-Eurocard I/O

MVME 700* 6U x 80MM wire wrap board

MVME 701* 6U x 80MM DIN connector To DB25 and 50 Pin dual row header (Used with MVME 050)

MVME 702* 6U x 80MM. Provides 50, 34 & 20 pin dual row headers for mass storage extension (Used with MVME 820/821)

800 Series-Mass Storage Support Products

MVME 820* 6U (double high) plug-in mass storage module including 15 MByte Winchester drive and 1 MByte DS/DD 5¼" floppy drive

MVME 821* 6U (double high) plug-in mass storage module including two 1-MByte DS/DD 5¼" floppy drives

900 Series-Packaging Hardware

Backplanes

MVME 920 20 slot VMEbus

MVME 921 9 slot VMEbus

MVME 922 5 slot I/O channel

MVME 924* 3 slot I/O channel

Extender Boards

MVME 930 VMEbus extender

MVME 932 VMEbus and I/O channel extender

Wire Wrap Boards

MVME 931-1 VMEbus wire wrap

MVME 933-1 I/O channel wire wrap

MVME 940-1 chassis with 7 double-high VMEbus slots, 2 I/O slots

MVME 941 card cage with 9 double-high VMEbus slots and 2 I/O slots

MVME 942 card cage with 20 double-high VMEbus slots

MVME 943* chassis with 8 double-high VMEbus slots, 6 I/O channel slots (front) and 16 I/O channel or 80MM double-high slots (rear) and accepts MVME 820/821 mass storage devices

MVME 944* chassis with 20 VMEbus slots (front) and 16 I/O channel or 80MM double high slots (rear)

*In development. Contact your Motorola sales representative for current availability.

Specifications

	VMEsystem	MULTIBUS II	Comments
Speed			
Bus Timing	Asynchronous (except VMSbus)	Synchronous (Fixed at 10 MHz clock rate)	Today's technology exceeds 10MHz clock rate. The MULTIBUS II fixed clock rate degrades high performance processors.
Transfer speed (theoretical)	VMEbus 57Mbyte/sec. VMXbus 80Mbyte/sec. VMSbus 3.2 Mbit/sec.	iPSB 40 Mbyte/sec. iLBX II 48 Mbyte/sec. iSSB 2 Mbit/sec. MULTICHANNEL 8 Mbyte/sec.	VMEsystem offers higher total system throughput
Serial Bus Protocol	Sender self check Collision Tolerant	CRC CSMA/CD	VMS rate stabilizes under heavy usage. CSMA/CD declines under heavy usage.
Flexibility			
Bus Timing	Asynchronous	Synchronous	VME transfer rate adjusts better to processor clock rate: 12.5 MHz with MC68000; 16 MHz with MC68020
Arbitration	Centralized or Distributed Control (4-level daisy chain) Priorities 4 (20 per level)	Distributed Control Parallel Contention Priorities 2 (20 per level)	Flexible arbitration and interrupt configurations of VMEbus simplify the use of multiprocessors.
Interrupt Request Lines	7	0	
Address width	a.) 16 bits b.) 24 bits c.) 32 bits	a.) 16 bits b.) 32 bits	Only Multibus II requires full 32-bit decoding logic in ALL applications.
Interrupt Request Structure	Allows deterministic protocol on VMEbus	Does not allow deterministic protocol on iPSB	VMEbus offers significantly greater flexibility for the system designers.
Arbitration	a.) Priority b.) Round Robin c.) Single level	a.) Priority	
Requester	a.) Release when done b.) Release on request c.) Release on BCLR d.) Release on AC Fail	a.) Release when done	
Number of Address Spaces	64	4	
Size	Double Eurocard (233.3mm x 160mm) Single Eurocard (100mm x 160mm)	Double Eurocard (233.3mm x 220mm) Single Eurocard (100mm x 220mm)	VME card sizes are popular and accepted today and will maximize functionality/cost tomorrow. MULTIBUS II card size is unpopular and hard to fit today and will be too cumbersome for tomorrow.

Broad based support: A league of interdependent vendors.

Continuing support by Motorola, Mostek and Signetics/Philips insures that VMEbus and its supporting bus structures VMXbus, VMSbus and Motorola's I/O Channel undergo further refinements in response to market needs.

VMEbus compatible manufacturers*

AEG-TELEFUNKEN CORP.
AMP INC.
AMPHENOL NORTH AMERICA
ASTRAEA COMPUTER CORP.
AUGAT INTERCONNECTION SYSTEMS GROUP
BFE FERNMELDE + ELEKTRONIK KG
BICC-VERO ELECTRONICS INC.
BURR BROWN
DATARAM CORP.
ELTEC ELECTRONIC GmbH
FORCE COMPUTERS INC.
HEMENWAY CORPORATION
HUNTER & READY INC.
MOSTEK-UNITED TECHNOLOGIES
MOTOROLA INC.
PANDUIT CORP.
PHILIPS
SIGNETICS CORP.
VECTOR ELECTRONIC COMPANY
XYCOM INC.

*This is a partial list of over 90 manufacturers. For a complete list check the VMEbus Compatible Products Directory box on the response card.

MC68000: quality right from the start.

Motorola's full line of VME integration starts with the MC68000 Microprocessor. From this preferred building block for systems design come VME processor and support modules, VME module chassis plus the all-in-one VME/10, the emerging leader in user configurable OEM microcomputers.

The MC68000 is preferred for multi-processing applications where performance and reliability are required. Its asynchronous interface to the system solves bus arbitration problems in large configurations and its powerful instruction set allows automatic execution of repeated program sequences for reducing high level development efforts.

The M68000 Family includes the economical MC68008 with its 8-bit bus, the MC68010 with virtual memory capabilities and, the latest addition, 32-bit MC68020 with co-processing and cache memory.

Direct object code compatibility up the M68000 line allows utilization of previously developed code. This allows the use of modular software designs that minimize development costs.

The cutting edge of technology is a double edged sword. If some outdated elements of an integrated system are replaced, adjustments ripple through the system, translating into cost and lost benefits of being "state-of-the-art." With Motorola's M68000 Family you won't be backed into a corner by a design that can't be upgraded.

VMEsystem Architecture is the only standard allowing full utilization of the impressive power of the M68000 Family.

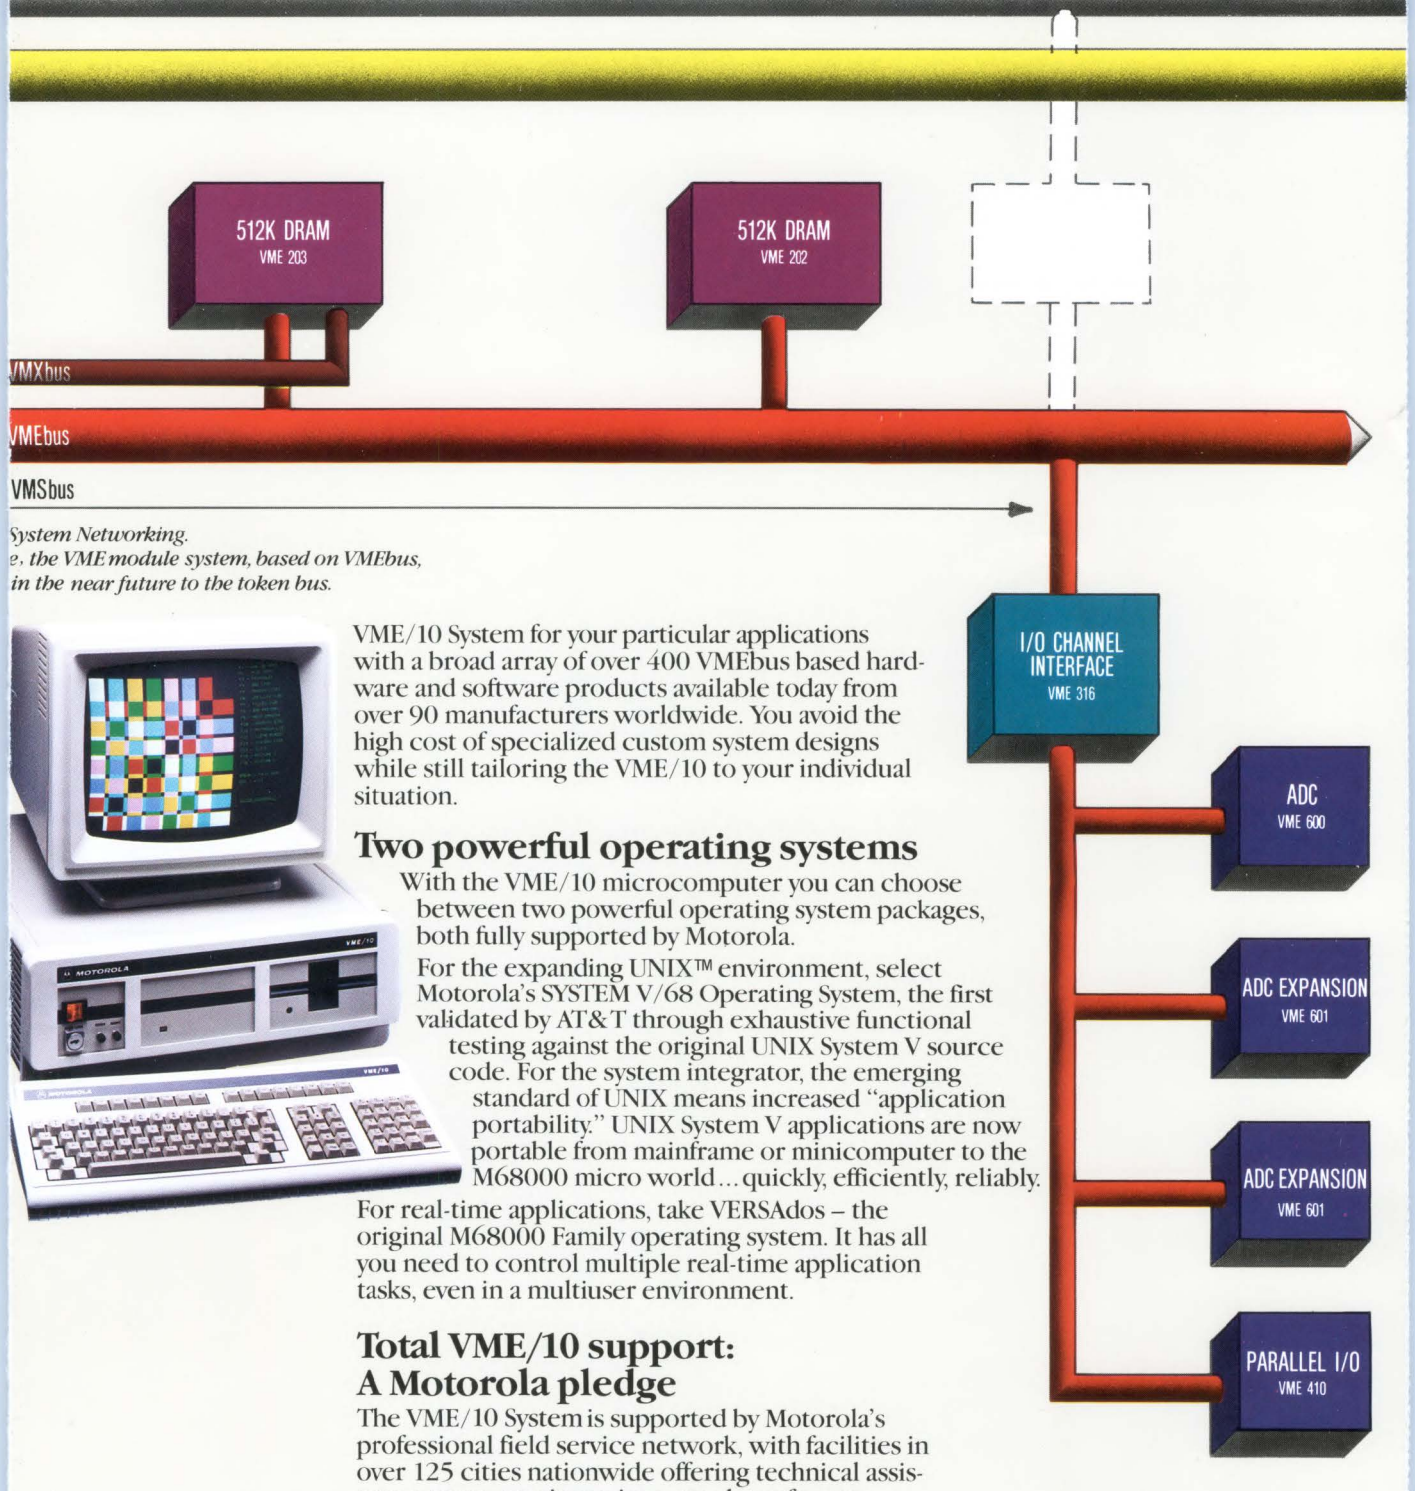

System Networking.
e, the VME module system, based on VMEbus,
in the near future to the token bus.

VME/10 System for your particular applications with a broad array of over 400 VMEbus based hardware and software products available today from over 90 manufacturers worldwide. You avoid the high cost of specialized custom system designs while still tailoring the VME/10 to your individual situation.

Two powerful operating systems

With the VME/10 microcomputer you can choose between two powerful operating system packages, both fully supported by Motorola.

For the expanding UNIX™ environment, select Motorola's SYSTEM V/68 Operating System, the first validated by AT&T through exhaustive functional testing against the original UNIX System V source code. For the system integrator, the emerging standard of UNIX means increased "application portability." UNIX System V applications are now portable from mainframe or minicomputer to the M68000 micro world... quickly, efficiently, reliably.

For real-time applications, take VERSAdos – the original M68000 Family operating system. It has all you need to control multiple real-time application tasks, even in a multiuser environment.

Total VME/10 support: A Motorola pledge

The VME/10 System is supported by Motorola's professional field service network, with facilities in over 125 cities nationwide offering technical assistance, parts, repair services, regular software updates, and a problem-reporting hotline. Our comprehensive user documentation and technical training seminars are complete and available.

Quality, standards, performance and support. Important reasons why you should make Motorola your OEM Systems vendor.

OEM VME/10 microcomputer system

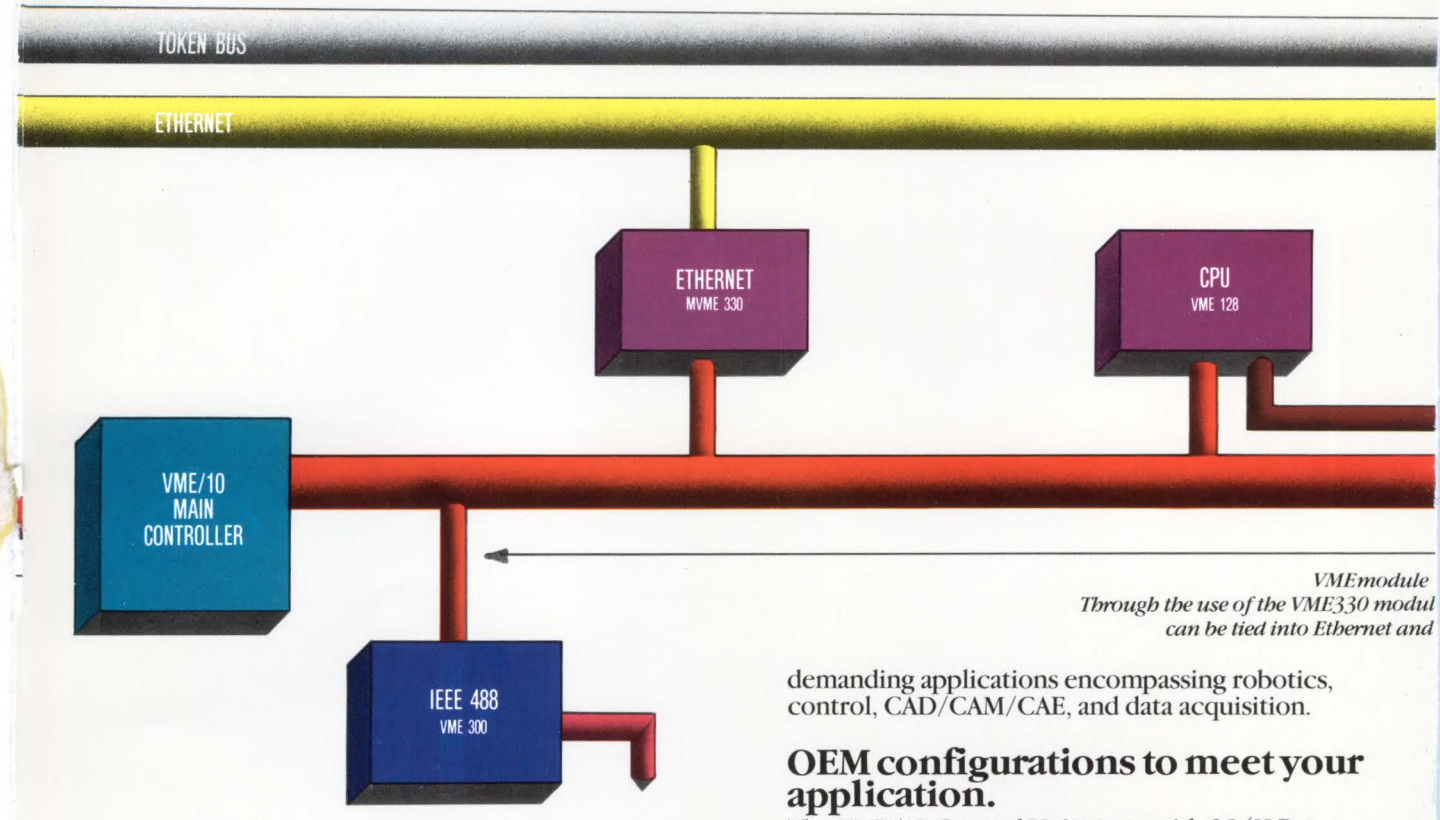

VMEmodule
Through the use of the VME330 modul
can be tied into Ethernet and

Notes to the OEM marketer

A full range of integrated products is of prime importance for the value added resale marketer. For the OEM, selecting a vendor that provides the full range of integration from chip, board or box is essential.

The Motorola choice can minimize your initial risk while maximizing your flexibility. Combine Motorola's half billion dollar R&D budget with its third party hardware and software and you have a complete vendor solution. From M68000 Microprocessor Family to Ethernet, VERSAdos to UNIX Operating Systems and VMEbus to I/O Channel, Motorola utilizes reliable, compatible and competitive products that remain, time after time, the smart assured investment for your future.

VME/10. OEM microsystems from Motorola.

M68000 performance, UNIX™ portability. VMEbus flexibility. Modern tools for the system integrator.

In a field overrun with specialized, dedicated, and limited OEM micro's, there is the VME/10 System alternative. It's the only VMEbus based user-configurable OEM "engine" combining the high-performance mainframe capabilities of the MC68010 Microprocessor with standard VMEsystem Architecture and a choice of two powerful operating systems.

The VME/10 System is designed to allow OEMs and system integrators to harness a wide range of demanding applications encompassing robotics, control, CAD/CAM/CAE, and data acquisition.

OEM configurations to meet your application.

The VME/10 Control Unit starts with 384K Bytes of RAM, a 655K Byte floppy disk drive, and 5M, 15M or 40M Byte Winchester disk drive unit. High resolution color graphics system is also standard. An optional keyboard, plus separate color or monochrome display unit mean you use only what you need for your application.

And the VME/10 goes far beyond simple "one-on-one" applications. As a Level II supervisory control computer, multiple VME/10s may be networked together for large-scale distributed control usage.

Mainframe performance, micro simplicity.

At the heart of the VME/10 OEM system is the MC68010 MPU – the latest addition to the upward-compatible, 16/32-bit M68000 microprocessor family. The combination of the MC68010 MPU and the MC68451 memory management unit provides processing power permitting multiple tasks to proceed simultaneously with full protection for program code and data in each task. You get fast, efficient application development on the VME/10 System.

VMEbus, global standard.

VMEbus makes the VME/10 System highly configurable. The bus structure is currently being approved as a global, standard, 16/32-bit microsystem architecture by working committees of both the IEEE (P1014) and the International Electrotechnical Committee (47B). Increasingly widespread adoption of VMEbus means that you can customize your

The intangible extras... from Motorola.

Since the earliest days of semiconductors, Motorola has earned for itself an enviable reputation as a frontrunner in both technology and support.

Its leadership, first in discrete devices... then in integrated circuits... has quite naturally expanded to encompass the still-infant field of microcomputers where the M68000 microprocessor family has become an internationally-sourced 16-, 32-bit standard.

First with the component "chips," then with board-level microcomputer modules, finally with box-level system entries... Motorola offers equipment manufacturers both the appropriate building blocks and advanced development support instrumentation to streamline the implementation of their end-system products.

And beyond these direct benefits, Motorola offers the following unique combination of advantages to users of VMEsystem products.

- Field-proven Motorola quality and reliability... backed by the Motorola warranty.
- Field application assistance... available through an extensive network of over 50 Motorola semiconductor sales offices in North America.
- Comprehensive technical training seminars... on a wide variety of hardware, software, and systems topics... at the factory, at diverse locations throughout the world, even at your own plant site!
- Field and factory service arrangements... on-site installation and repair... in over 125 cities nationwide... maintenance contracts... factory repair/exchange services.
- Third party software support... a continuing program of assistance to independent software vendors, to encourage development of both systems and applications packages for VMEsystem products.
- Commitment to VMEsystem architecture... with Motorola's active participation and support of the VMEbus Manufacturers Group, and of the VMEbus IEEE and IEC standardization activities.
- Comprehensive reference documentation... covering all hardware and software elements of the VMEsystem product offering.

MOTOROLA SEMICONDUCTOR SALES OFFICES

MOTOROLA SEMICONDUCTOR AMERICA'S DISTRICT OFFICES

ALABAMA, Huntsville ..... (205)830-1050
ARIZONA, Phoenix ..... (602)244-7100
CALIFORNIA, Encino/
Sherman Oaks ..... (818)706-1929
(213)872-1505
CALIFORNIA, Los Angeles ..... (213)417-8848
CALIFORNIA, Orange
(Orange Exch.) ..... (714)634-2844
(L.A. Exch.) ..... (213)445-5585
CALIFORNIA, San Diego ..... (714)560-4644
CALIFORNIA, San Jose ..... (408)985-0510
COLORADO, Colorado Springs .. (303)599-7404
COLORADO, Denver ..... (303)773-6800
CONNECTICUT, New Haven/
Hamden ..... (203)281-0771
FLORIDA, Pompano Beach/
Ft. Lauderdale ..... (305)491-8141
FLORIDA, Casselberry/Maitland (305)831-3422
FLORIDA, St. Petersburg ..... (813)576-6030
GEORGIA, Atlanta ..... (404)256-0222
ILLINOIS, Chicago/Schaumburg (312)576-7800
INDIANA, Fort Wayne ..... (219)484-0436
INDIANA, Indianapolis ..... (317)849-7060
INDIANA, Kokomo ..... (317)457-6634
IOWA, Cedar Rapids ..... (319)373-1328
KANSAS, Kansas City/Mission ... (913)384-3050
KANSAS, Wichita ..... (316)721-2449
MASSACHUSETTS, Berlin ..... (617)562-3856
MASSACHUSETTS, Burlington .. (617)273-5020
MICHIGAN, Detroit/Westland .... (313)261-6200
MINNESOTA, Minneapolis ..... (612)545-0251
MISSOURI, St. Louis ..... (314)872-7681
NEW JERSEY, Hackensack ..... (201)488-1200
NEW YORK, Poughkeepsie/
Fishkill ..... (914)473-8102
NEW YORK, Long Island/
Hauppauge ..... (516)231-9000
NEW YORK, Pittsford ..... (716)248-5494
NEW YORK, Syracuse ..... (315)455-7373
NORTH CAROLINA, Raleigh ..... (919)876-6025
OHIO, Cleveland ..... (216)461-3160
OHIO, Dayton ..... (513)294-2231
OHIO, Columbus/
Worthington ..... (614)846-9460
OKLAHOMA, Tulsa ..... (918)664-5227
OREGON, Portland ..... (503)641-3681
PENNSYLVANIA, Philadelphia/
Horsham ..... (215)443-9400
TENNESSEE, Knoxville ..... (615)690-5592
TEXAS, Austin ..... (512)452-7673
TEXAS, Dallas ..... (214)931-9222
TEXAS, Houston ..... (713)783-6400
UTAH, Salt Lake City ..... (801)539-1190
VIRGINIA, Charlottesville ..... (804)977-3691
WASHINGTON, Bellevue ..... (206)454-4160
Seattle Access ..... (206)622-9960
WASHINGTON, DC/MARYLAND,
Hyattsville ..... (301)577-2600
WISCONSIN, Milwaukee/
Wauwatosa ..... (414)476-5554

Field Applications Engineering Available Through All Sales Offices

MOTOROLA SEMICONDUCTOR—CANADA

MANITOBA, Winnipeg ..... (204)889-0693
ONTARIO, Downsview/Toronto ... (416)661-6400
ONTARIO, Ottawa ..... (613)235-4388
QUEBEC, Montreal ..... (514)731-6881

MOTOROLA SEMICONDUCTOR INTRA-COMPANY OFFICES

ARIZONA, Scottsdale ..... (602)949-3811
FLORIDA, Ft. Lauderdale ..... (305)475-6120
ILLINOIS, Schaumburg ..... (312)576-2788
ILLINOIS, Schaumburg/
Automotive ..... (312)576-7800
TEXAS, Ft. Worth ..... (817)232-6255

MOTOROLA SEMICONDUCTOR INTERNATIONAL SALES OFFICES

AUSTRALIA, Melbourne ..... (03)561-3555
AUSTRALIA, Sydney 438-1955
439-2242
AUSTRIA, Vienna ..... (0222)65 01 26
BRAZIL, Sao Paulo ..... (011)572 3553
DENMARK, Gladsaxevej ..... (01)67 44 22
ENGLAND, Wembley, Middlesex ... 01-902-8836
FRANCE, Grenoble ..... (076)90 22 81
FRANCE, Paris ..... (01)555-91-01
FRANCE, Toulouse ..... (061)41 11 88
GERMANY, Langenhagen/
Hannover ..... (0511)78-20-37
GERMANY, Munich ..... (089)92 481
GERMANY, Nurenberg ..... (0911)65761
GERMANY, Sindelfingen ..... (07031)83074
GERMANY, Wiesbaden ..... (06121)76-1921
HONG KONG, Hung Hom,
Kowloon 3-632201-8
3-336211-22
ISRAEL, Tel Aviv 3-388-388
ITALY, Bologna ..... (051)533 446
ITALY, Milan ..... (02)824 2021
824-2046
ITALY, Rome ..... (03)831 4746
JAPAN, Osaka ..... (06)305 1421
JAPAN, Tokyo 03-440-3311
KOREA, Seoul 261-7137
MEXICO, D.F. ..... (525)596-84-28
NETHERLANDS, Utrecht ..... (030)443 808
NORWAY, Oslo ..... (02)671467
SCOTLAND, East Kilbride ..... (03552)39 101
SINGAPORE 2945438
SOUTH AFRICA, Bramley 786 1184
SPAIN, Madrid ..... (01)279 0802
SWEDEN, Solna 08/82 02 95
SWITZERLAND, Geneva ..... (022)991 111
SWITZERLAND, Zurich ..... (01)730 40 74
TAIWAN, Taipei 7528944-9

MOTOROLA SYSTEMS REPRESENTATIVES

ALABAMA, Huntsville
EMA ..... (205)830-4030
CALIFORNIA, Inglewood
Basic Systems Corporation ..... (213)673-4300
CALIFORNIA, San Diego
Basic Systems Corporation ..... (714)999-6566
CALIFORNIA, Santa Clara
Basic Systems Corporation ..... (408)727-1800
CONNECTICUT, Goshen
Northern Computer Sales ..... (203)491-3585
GEORGIA, Atlanta
EMA ..... (404)329-0530
ILLINOIS, Northbrook
Marketech Assoc. Inc. ..... (312)291-0315
INDIANA, Carmel
Carter, McCormic & Peirce..... (317)844-4175
MASSACHUSETTS, Westwood
Northern Computer Sales ..... (617)326-3454
MICHIGAN, Farmington Hills
Carter, McCormic & Peirce..... (313)477-7700
MINNESOTA, Minneapolis
P.S.I. Company ..... (612)944-8545
MISSOURI, Independence
Dy-Tronix Inc. ..... (816)373-6600
NEW MEXICO, Albuquerque
I.M. Systems ..... (505)293-2440
NEW YORK, Elmont
Crane & Egert Corporation ..... (516)872-0800
NEW YORK, Fairport
Blaire Data Systems, Inc. ..... (716)223-8765
NORTH CAROLINA, Raleigh
EMA ..... (919)781-9369
OHIO, Cleveland
Carter, McCormic & Peirce..... (216)779-5100
OHIO, Dayton
Carter, McCormic & Peirce..... (513)836-0951
OREGON, Beaverton
Thorson Co. Northwest ..... (503)644-5900
PENNSYLVANIA, Monroeville
Carter, McCormic & Peirce..... (412)372-4415
PENNSYLVANIA, King of Prussia
K.S.I. ..... (215)783-6440
TENNESSEE, Jonesboro
EMA ..... (615)753-5844
WASHINGTON, Bellevue
Thorson Co. Northwest ..... (206)455-9180
WISCONSIN, Germantown
Marketech Associates..... (608)628-3663

MOTOROLA *Semiconductor Products Inc.*

P.O. BOX 20912 • PHOENIX, ARIZONA 85036 • A SUBSIDIARY OF MOTOROLA INC.

Chip set points to lower costs in Ethernet cabling

Not so long ago, the main concern in implementing full Ethernet systems was the cost of the interface to the data terminal equipment. However, the learning curve typical of the semiconductor industry has brought the Ethernet interface from very expensive multi- and single-board subsystems down to relatively inexpensive VLSI chip sets. One major cost bottleneck, however, does remain—the cabling and installation. The cost of silicon has dropped, but the cost of copper has risen.

Traditionally, data terminal equipment (DTE) has been connected to Ethernet via a transceiver that connects to the Ethernet cable ($1.50/ft) by drilling through the cable's shields and inserting a tapping pin to contact the central conductor. A drop cable is then run from the transceiver to the DTE. Using this method, installation costs can become quite high. For each piece of DTE, a transceiver and an expensive drop cable ($2.00/ft) must be installed by a trained technician.

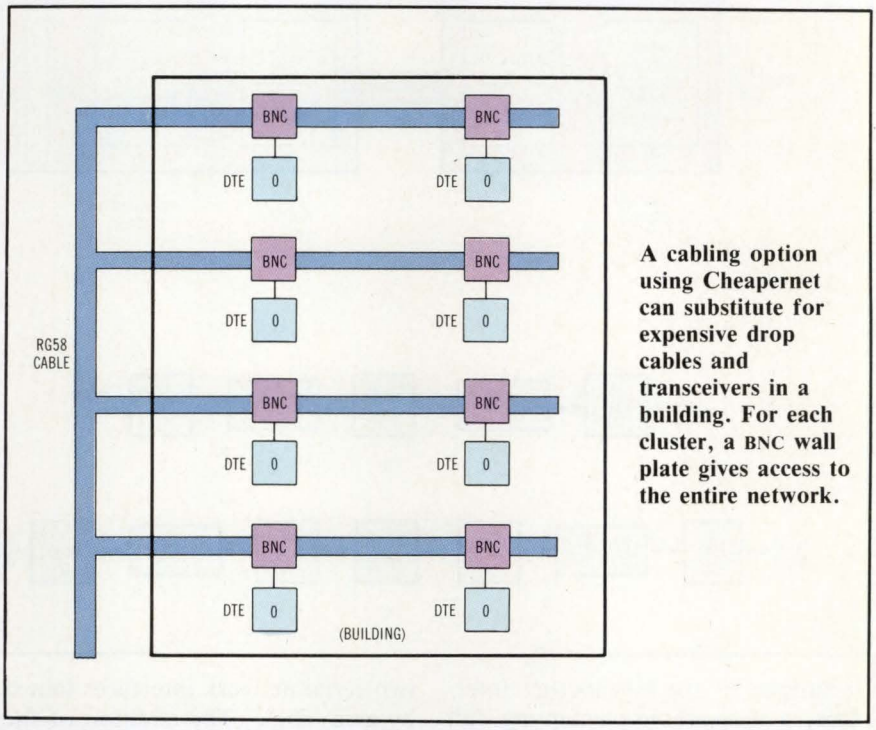

A cabling option using Cheapernet can substitute for expensive drop cables and transceivers in a building. For each cluster, a BNC wall plate gives access to the entire network.

Cutting down

A new Ethernet (IEEE 802.3) chip set from National Semiconductor Corp (2900 Semiconductor Dr, Santa Clara, CA 95051) cuts the cost of the interface electronics, cabling, and installation. Dubbed "Cheapernet," the chip set eliminates the extra drop cable by placing the transceiver electronics in the DTE along with the interface electronics. Connection to the actual Ethernet cable is made with a repeater that can be built up from the chip set. Additional DTEs can be connected to the same repeater in clusters via inexpensive RG58 coaxial cable that connects to the equipment by means of BNCs.

This cabling technique places certain limits on the Cheapernet version that do not apply to the traditional Ethernet configurations. In Cheapernet, the network connection allows a cluster of up to 30 units (versus 100 units for Ethernet) to be assembled per segment with a maximum cable length of 200 m (versus 500 m for Ethernet).

The components comprising the Cheapernet interface include the DP8390 network interface controller, the DP8391 serial network interface, and the DP8392 coaxial transceiver interface. They use the same IEEE 802.3 protocol as Ethernet except for the cabling technique previously explained.

The DP8390 network interface controller is a CMOS device that provides all the data link layer functions required to transmit and receive Ethernet/Cheapernet packets. It includes dual 16-bit DMA channels, which configure the buffer memory as dual-port memory. This dual-port memory then decouples network bus activity from the host. All bus arbitration logic and memory support are contained onchip. The network interface controller is able to filter physical, multicase, and broadcase addresses; and has a three-level loopback for node diagnosis.

The serial network interface connects the network interface controller to the transceiver cable. It performs the Manchester encoding and decoding of packets, generates carrier sense, and decodes the collision signal. This transceiver cable performs the same function as Ethernet drop cable except it is connected to the transceiver built around the coaxial transceiver interface contained within the terminal equipment.

The coaxial transceiver interface integrates all needed transceiver electronics, save the signal and power isolation. It directly drives the transceiver cable through isolation transformers. On the other side, it connects directly to 50-Ω coaxial cable. The DP8392 has an onchip jabber timer that could disable the coaxial transceiver interface driver if a data packet were longer than legal length. It also contains a collision detector and generates the collision signal to the DTE.

Thus, by using the simpler and less expensive RG58 coaxial cabling

(continued on page 62)

Chip set

(continued from page 61)

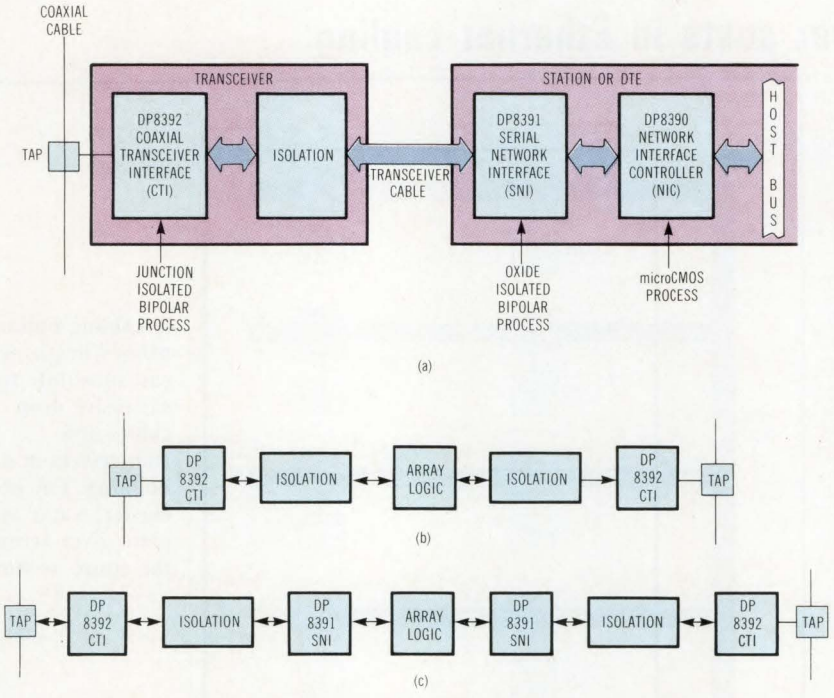

The Cheapernet chip set can be used to provide an integrated network interface that includes both the interface and transceiver electronics inside the data terminal equipment (a). Two types of repeaters can be formed for connecting clusters to one another or to the main Ethernet cable: one is the least expensive configuration that does not retime the signal (b); the other is a repeater that retimes the signal (c).

technique to the Cheapernet interface, a network implementing full IEEE 802.3 protocol can be put together. In addition, it is user installable—the lengths of RG58 just need to be connected to the T connectors on the back of the equipment. All this implies that the equipment specified for such a cluster, and for the more complex configurations described below, must have the full Cheapernet interface either built into it or plugged into its host bus.

Making the connection

However, the Cheapernet chip set provides several alternative cabling schemes that allow clusters to be connected with a main Ethernet cable via repeaters. Two basic repeater types can be built using members of the Cheapernet chip set. The most inexpensive uses two coaxial transceiver interfaces joined by gate array logic. This repeater can join either a cluster to the Ethernet cable or two clusters together, and is best suited for short hauls as it does not retime the signal.

The second type of repeater uses two coaxial transceiver interfaces and two serial network interfaces joined by array logic. The addition of the serial network interfaces retimes the signal and eliminates jitter, making this repeater more reliable than the first. Both repeaters provide some isolation for the cluster from the rest of the network. Thus, if there were a break in the cable in the cluster, the network would see a problem at exactly the same point every time access is attempted. It would then ignore that cluster as if it were a defective DTE, without bringing down the entire network.

The repeaters provide a more cost-effective means for wiring an Ethernet building. Basically, an Ethernet spine can be installed in a building when it is constructed. Repeaters attached to this spine can supply cluster cables for each room or an otherwise defined area on every floor. The planners need only keep in mind the 200-m/30-node limitation for each cluster.

Each cluster cable can be brought behind the wall to a simple BNC wall plate. To this wall plate all the DTEs in a given cluster can be connected by the user. This cluster can be expanded up to the specified limits by adding lengths of RG58 cable and connecting more devices. The new IEEE 802.3 chip set provides one solution to the cabling and installation cost bottleneck by eliminating the need for professional installation and expensive drop cable.

—Tom Williams,
West Coast Managing Editor

SYSTEM TECHNOLOGY

(continued on page 68)

October will be a two-issue month—*watch for a Special Report on Mass Storage, October 1; on Digital Signal Processing October 15*

Keep your system CMOS-cool even with the hottest NMOS micros.

Switch to cool CMOS peripherals with your hot NMOS micro.

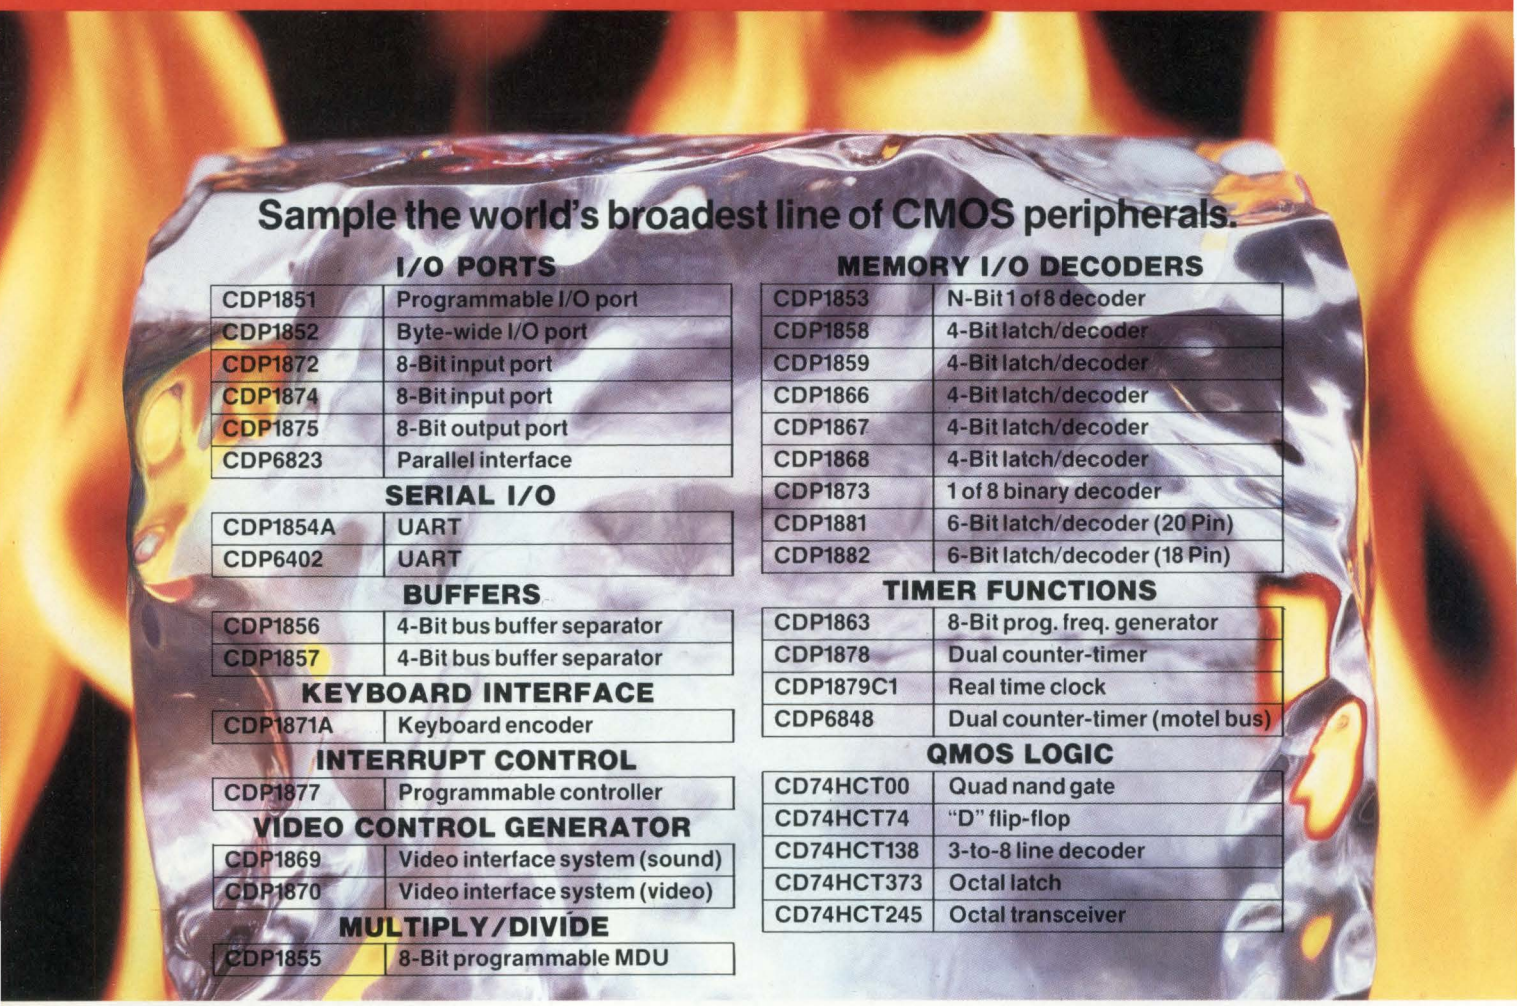

Sample the world's broadest line of CMOS peripherals.

I/O PORTS

CDP1851	Programmable I/O port
CDP1852	Byte-wide I/O port
CDP1872	8-Bit input port
CDP1874	8-Bit input port
CDP1875	8-Bit output port
CDP6823	Parallel interface

SERIAL I/O

CDP1854A	UART
CDP6402	UART

BUFFERS

CDP1856	4-Bit bus buffer separator
CDP1857	4-Bit bus buffer separator

KEYBOARD INTERFACE

CDP1871A	Keyboard encoder

INTERRUPT CONTROL

CDP1877	Programmable controller

VIDEO CONTROL GENERATOR

CDP1869	Video interface system (sound)
CDP1870	Video interface system (video)

MULTIPLY/DIVIDE

CDP1855	8-Bit programmable MDU

MEMORY I/O DECODERS

CDP1853	N-Bit 1 of 8 decoder
CDP1858	4-Bit latch/decoder
CDP1859	4-Bit latch/decoder
CDP1866	4-Bit latch/decoder
CDP1867	4-Bit latch/decoder
CDP1868	4-Bit latch/decoder
CDP1873	1 of 8 binary decoder
CDP1881	6-Bit latch/decoder (20 Pin)
CDP1882	6-Bit latch/decoder (18 Pin)

TIMER FUNCTIONS

CDP1863	8-Bit prog. freq. generator
CDP1878	Dual counter-timer
CDP1879C1	Real time clock
CDP6848	Dual counter-timer (motel bus)

QMOS LOGIC

CD74HCT00	Quad nand gate
CD74HCT74	"D" flip-flop
CD74HCT138	3-to-8 line decoder
CD74HCT373	Octal latch
CD74HCT245	Octal transceiver

NMOS systems get CMOS-cool.

Even systems based on NMOS micros can interface with CMOS savings by using RCA's broad line of CMOS peripherals.

RCA has tested and proven interfaces between CMOS peripherals and the most popular 8-bit micros, including Z80, 8085, 8048, NSC800, 6500, and of course the 1800 series and 6805.

With relatively simple redesign, even NMOS systems require less power, lower cooling costs, less weight and space with CMOS peripherals.

World's broadest line of CMOS peripherals.

RCA offers more CMOS peripherals than anyone else: 28 types tested and proven to interface with NMOS or CMOS 8-bit micros.

Cool logic too: QMOS.

To make the most of CMOS advantages, use QMOS high speed CMOS logic as glue types. Speeds are comparable to LSTTL with 1/1000 the quiescent power consumption and 75% better noise immunity. RCA offers a broad line of HC parts to design-in, plus the industry's best selection of HCT drop-in replacements for LSTTL. Order a free sample today.

RCA CMOS vs NMOS peripherals.

Factor	NMOS	CMOS	CMOS Advantage
Operating power	1	1/10	Yes
Reliability	High chip temperature	Low chip temperature	Yes
Standby power	mW	μW	Yes
Speed	High	Medium/High	
Operating frequency range	Limited	DC to max	Yes
Operating voltage	4.5V to 5.5V	4V to 10V	Yes
Noise immunity	10% of supply voltage	30% of supply voltage	Yes
Temperature range	0° to 70° C	-40° to +85° C Military: -55° to +125° C	Yes
Chip complexity and size	Same	Same	
Gate arrays and Semi-custom availability	No	Yes	Yes

Photography by Pete Turner, NYC.

Call: Hamburg, 49-4106-6130; London, 44-03-2785511; Paris, 33-3-946-5656; Hong Kong, 8-52-3-723-6339; Sao Paulo, 55-11-210-4033.

Send us your hottest design idea and win a cool prize.

Design-in contest.

Prove to yourself how easy it is to use RCA peripherals in your system. Enter your hottest design using any RCA peripheral with any 8-bit microprocessor (NMOS or CMOS).

First prize: IBM 5150 PC XT with a color display and graphic printer.

Second prize: RCA Selectavision VCR with remote control.

Third prize: RCA color video camera.

Everyone who enters will win a "Megachange" mug.

Free samples.

Whether or not you enter the contest, RCA will supply you with two free samples of the CMOS peripheral of your choice and one free QMOS logic sample. Just check the chart on the previous page and send in the coupon with your selection.

Or call toll-free (800) 526-2177.

Free Sample/Contest Entry Form.

Please mail your completed form to: RCA Solid State, P.O. Box 2900, Somerville, NJ 08876.

☐ I'd like to enter the design-in contest. Please send me an entry kit.

☐ I don't want to enter, but I do want samples.

☐ Please send me the following samples:

CMOS peripherals (2 maximum), types: _____

_____

QMOS high-speed logic (1 maximum), type: _____

_____

Name _____

Title _____

Company _____

Address _____

City/State/Zip _____

Telephone ( ) _____

To get contest entry kit or free samples, use this coupon or facsimile. Or call 800-526-2177. One entry per person. Deadline for entries is December 15, 1984. Winners will be announced by January 30, 1985. Void where prohibited by law.

Changing to serve you better.

RCA

IF YOUR BUSINESS IS SHRINKING, SHUGART CAN HELP.

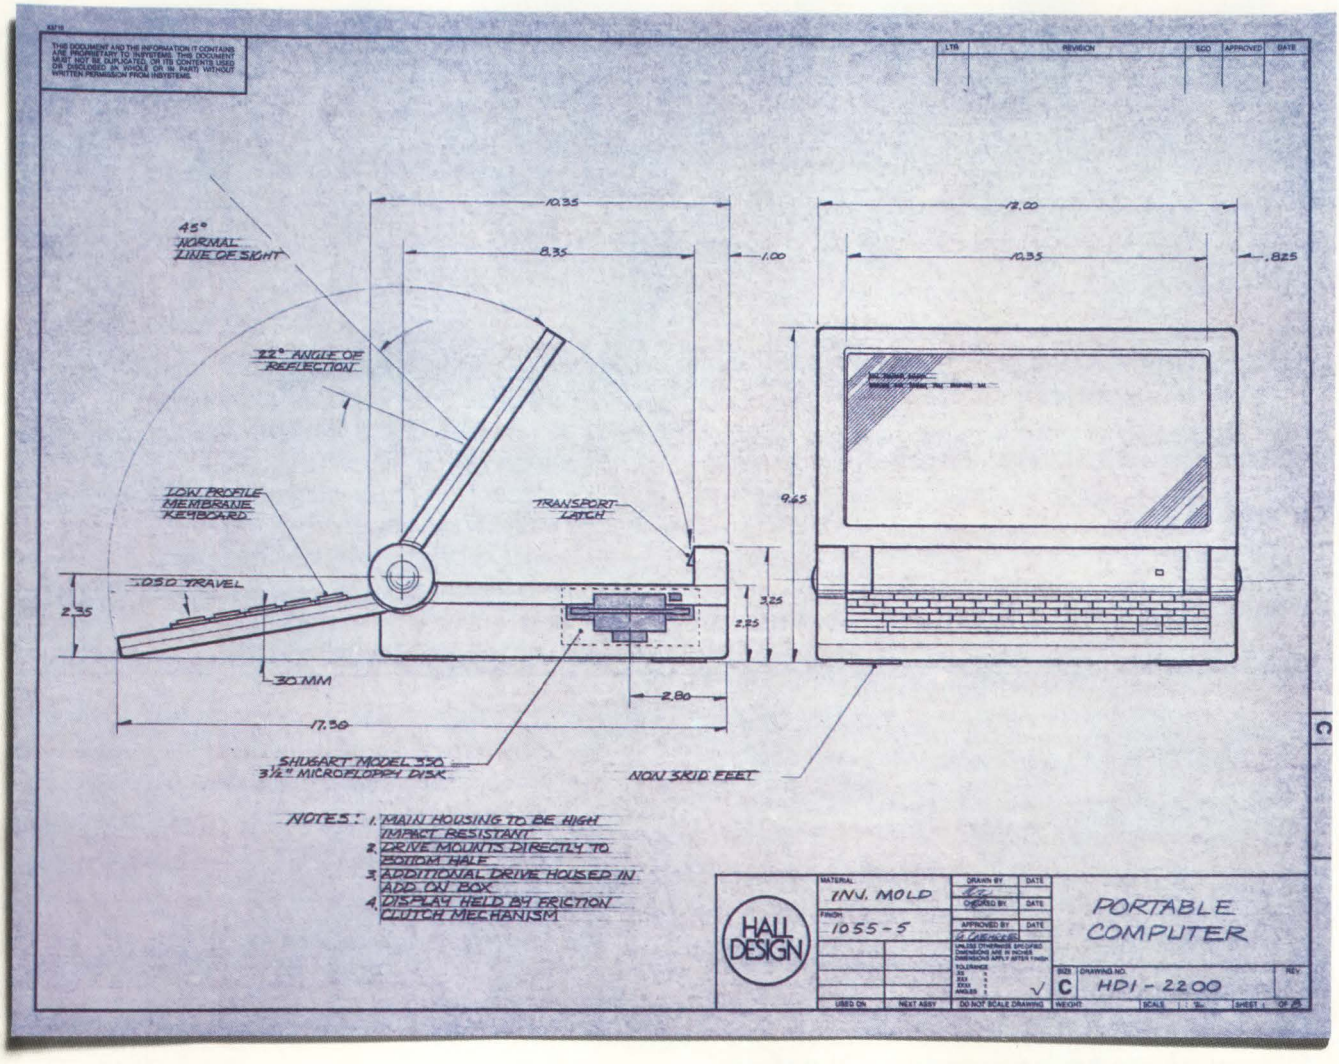

These days, computer designers everywhere face a problem of massive proportions: How do you cram a desktop computer into a briefcase?

Sound familiar? If so, there's a family of 3.5" single and double-sided microfloppy disk drives you should meet. The Shugart 300 and 350, respectively.

The perfect drive solution for a full-featured portable.

Tiny enough to fit easily into your smallest design. Yet with a 6 millisecond average access time, a capacity up to 1 megabyte and Minifloppy™ compatibility, your portable computer could easily run the same software as someone else's desktop.

And keep it running for quite some time.

Shugart drives are so reliable, you can count on an MTBF of 10,000 power on hours. One reason we're projecting delivery of over 100,000 microfloppy drives this year.

At just over a pound apiece, you could even use two. And still call your portable computer portable.

There's just one thing to remember. You should always check the activity light on a Shugart microfloppy. They're so quiet, there's no other way to tell if they're running.

What more could a shrinking business need? A couple of other small things. Industry standard 3.5" microcartridges, to be exact. Their track densities offer a more than generous upgrade path. But more important, considering where they could end up, they come equipped with a hard shell plastic media cartridge. And an automatic head access shutter. Sure protection from all kinds of catastrophes. Stick them in your pocket. Throw them in your purse. Bang them around in your briefcase. They'll survive.

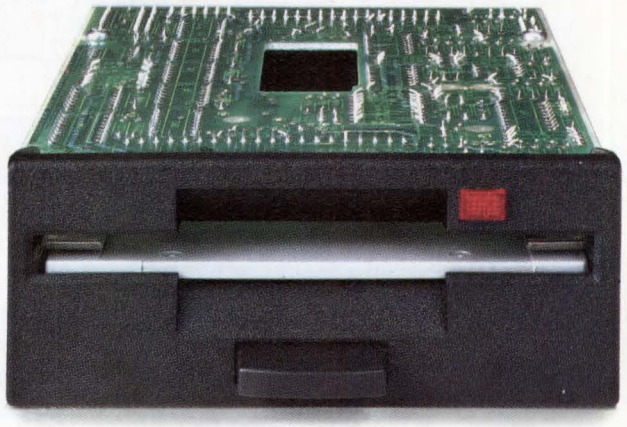

The 3.5" Shugart Microfloppy. Smaller than actual size, but not much.

Sound interesting? Call your local Shugart sales office. We'll do a private microfloppy workshop right in your office. But do it soon. The more your business shrinks, the more Shugart can help.

Shugart

Right from the start.

Milpitas, CA (408) 263-2600; Costa Mesa, CA (714) 979-1935; Thousand Oaks, CA (805) 496-5388; Rochester, NY (716) 235-7190; Minneapolis, MN (612) 546-4411; Richardson, TX (214) 234-3568; Framingham, MA (617) 879-1700; Saddle Brook, NJ (201) 368-8445; Smyrna, GA (404) 436-0953; Markham, ONT (416) 475-2655; Paris, France (3) 946-42-66; Munich, West Germany (089) 786-021; London, U.K. (44) 4862-27272; Wanchai, Hong Kong (852) 5-733307.
Hamilton/Avnet, authorized distributor

Monolithic analog/digital ICs expand role in disk control

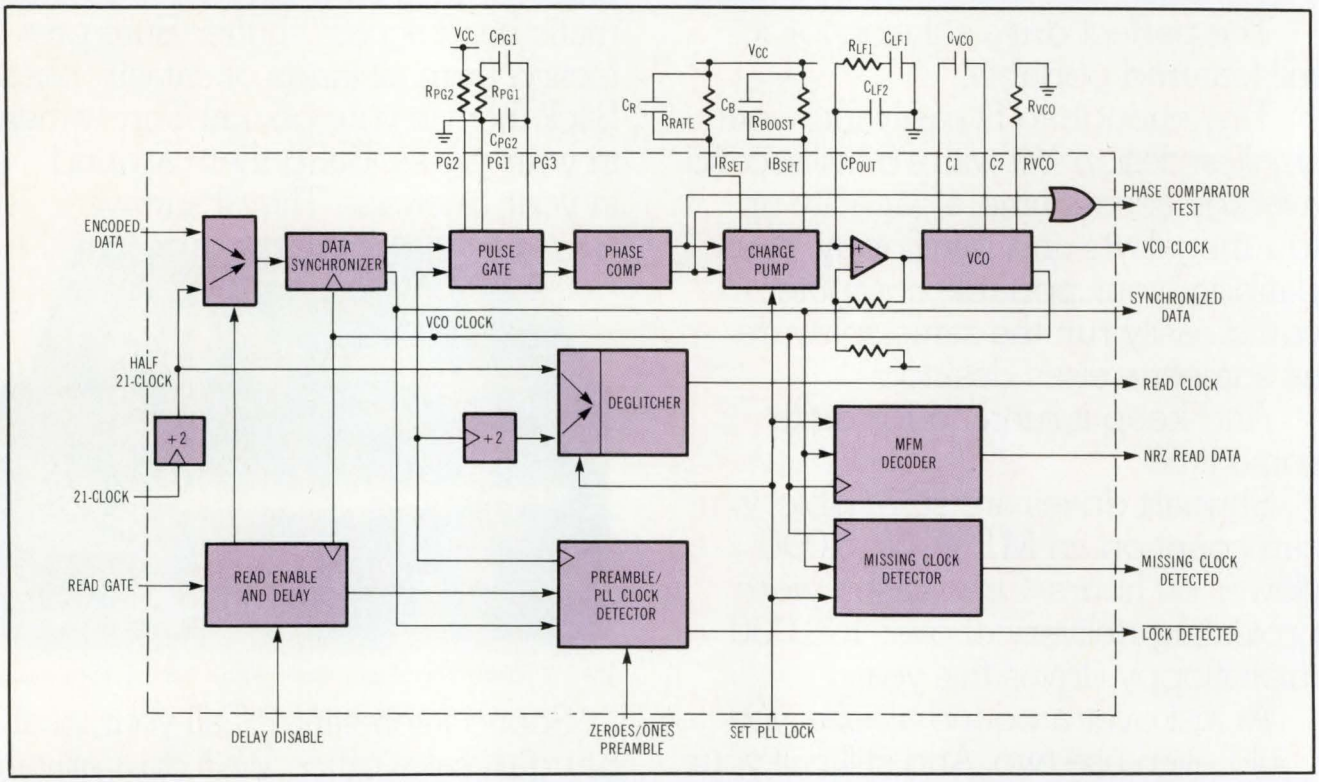

The DP8460 chip incorporates onchip phase locked loop (PLL) for data separation on Winchester disk drives. The chip can be mounted on the drive to put the PLL closer to the modified frequency modulation (MFM) source, thus preventing noise and bit shift errors.

Purely digital functions provided by disk controller chips have been available for several years. Now, the mixed analog/digital circuitry required to perform data separation on a single chip is available from one company and under development at other IC vendors. In addition, some manufacturers have developed monolithic servo controllers for spindle and head positioning and motor speed control.

This growing emphasis on higher levels of integration for various disk control functions has come about because of the explosive demand for 5¼-in. Winchester drives. In the past, companies that designed disk systems around 8- and 14-in. disks had designers in-house who were experts in the intricacies of phase locked loop (PLL) design. Furthermore, because of the drive size, the amount of board area used was inconsequential. Thus, low density systems using discrete components could be built.

The popularity of the 5¼-in. Winchester drive has changed all that. Space is now at a premium. Moreover, the sheer volume of small drives being incorporated into designs today has increased the need for PLL experts, and sharply reduced the number of such persons who are available. The combined requirements for reduced board size and the lack of engineers familiar with combined digital, analog, and PLL design have inspired activity to develop chips that will provide needed control.

Gaining control

One of the first such designs to be introduced this year is a single-chip Winchester disk data separator from National Semiconductor Corp (2900 Semiconductor Dr, Santa Clara, CA 95051). This chip, the DP8460, incorporates PLL control circuitry. It can be used either in the drive or the controller of 3½-, 5¼-, 8-, and 14-in. hard disk drives as well as high density floppies and optical disk drives. Mounting it in the drive confers a major advantage because the nearer the PLL is to the modified frequency modulation (MFM) source, the less chance there is of noise or bit shift errors being introduced. Along with the PLL, the chip incorporates a process insensitive delay line. This enables the chip to tolerate more bit jitter than other available systems.

PLL systems provide accurately synchronized signal pulse streams when decoding data coming off the disk, despite the presence of jitter. Many PLL designs, however, use discrete components, one-shots, and delay lines; and require alignment by

(continued on page 71)

You won't find this PC in an office.

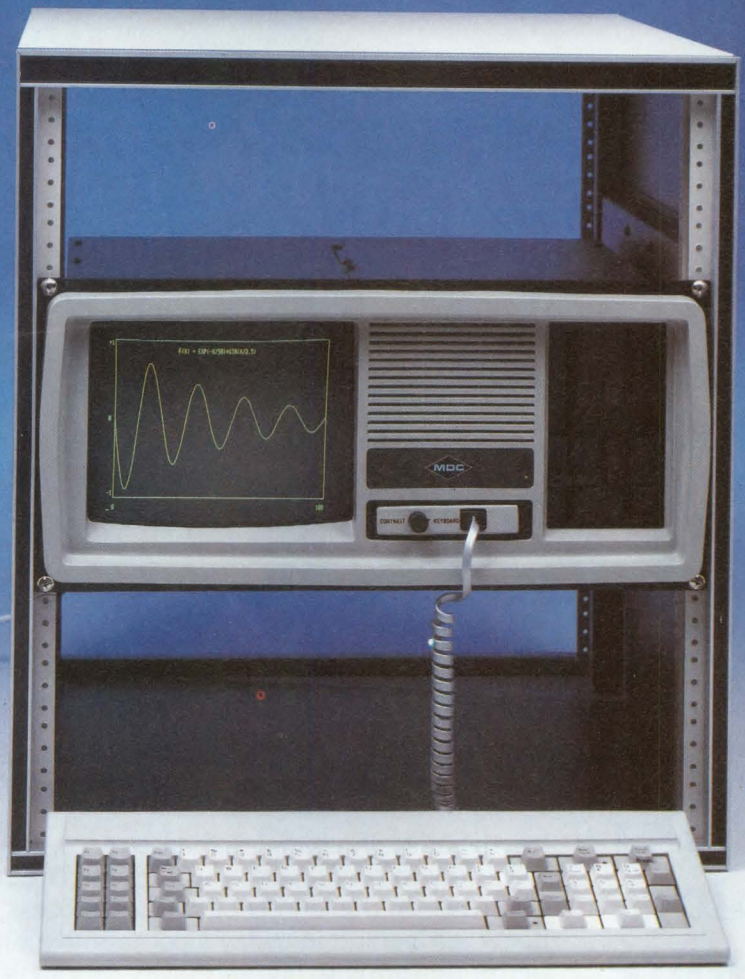

Compatibility:
IBM® PC

RAM on main board:
256 KB standard
512 KB optional

Storage:
Two 320 KB floppies
Optional 10 MB Winchester
RAM-disk built in

Ports:
1 parallel, 1 serial

Expandability:
4 available slots
110 W power supply

Display:
High-contrast, non-glare
16x13 character font
640x325 graphics built in

Software:
MS™-DOS, GW-BASIC™, assembler

The MDC RM-1600™ Rackmount IBM PC-Compatible Computer

At last, an IBM PC-compatible computer for industrial and scientific use. The RM-1600 makes available thousands of IBM PC-compatible software and add-on hardware products. It also offers the modularity of a rackmount housing.

And the RM-1600 provides four available expansion slots and a high capacity 110 W power supply that let you really mix and match all those add-ons.

For additional storage and speed, an optional 10 MB Winchester replaces a diskette drive. Or for speed enhancement alone, try the built-in RAM disk, using the 512 KB on-board RAM capacity. In numerical applications, you can add an 8087 co-processor for greater speed.

So whether the application is laboratory instrumentation, test and measurement, R&D, numerical control, process control, or OEM systems, the RM-1600 lets you contruct your equipment as you want it.

MDC has been building process control equipment for over 13 years. We first created the RM-1600 for our own use, and now offer our advantage to you. Take the advantage; give us a call and arrange for a closer look.

RM-1600 is a trademark of Materials Development Corporation.
IBM is a registered trademark of International Business Machines Corporation.
MS and GW-BASIC are trademarks of Microsoft Corporation.
Product features subject to change without notice.

MDC **Materials Development Corporation**

21541 Nordhoff Street • Chatsworth, CA 91311 • (818) 700-8290 • Telex: 662438

CIRCLE 33

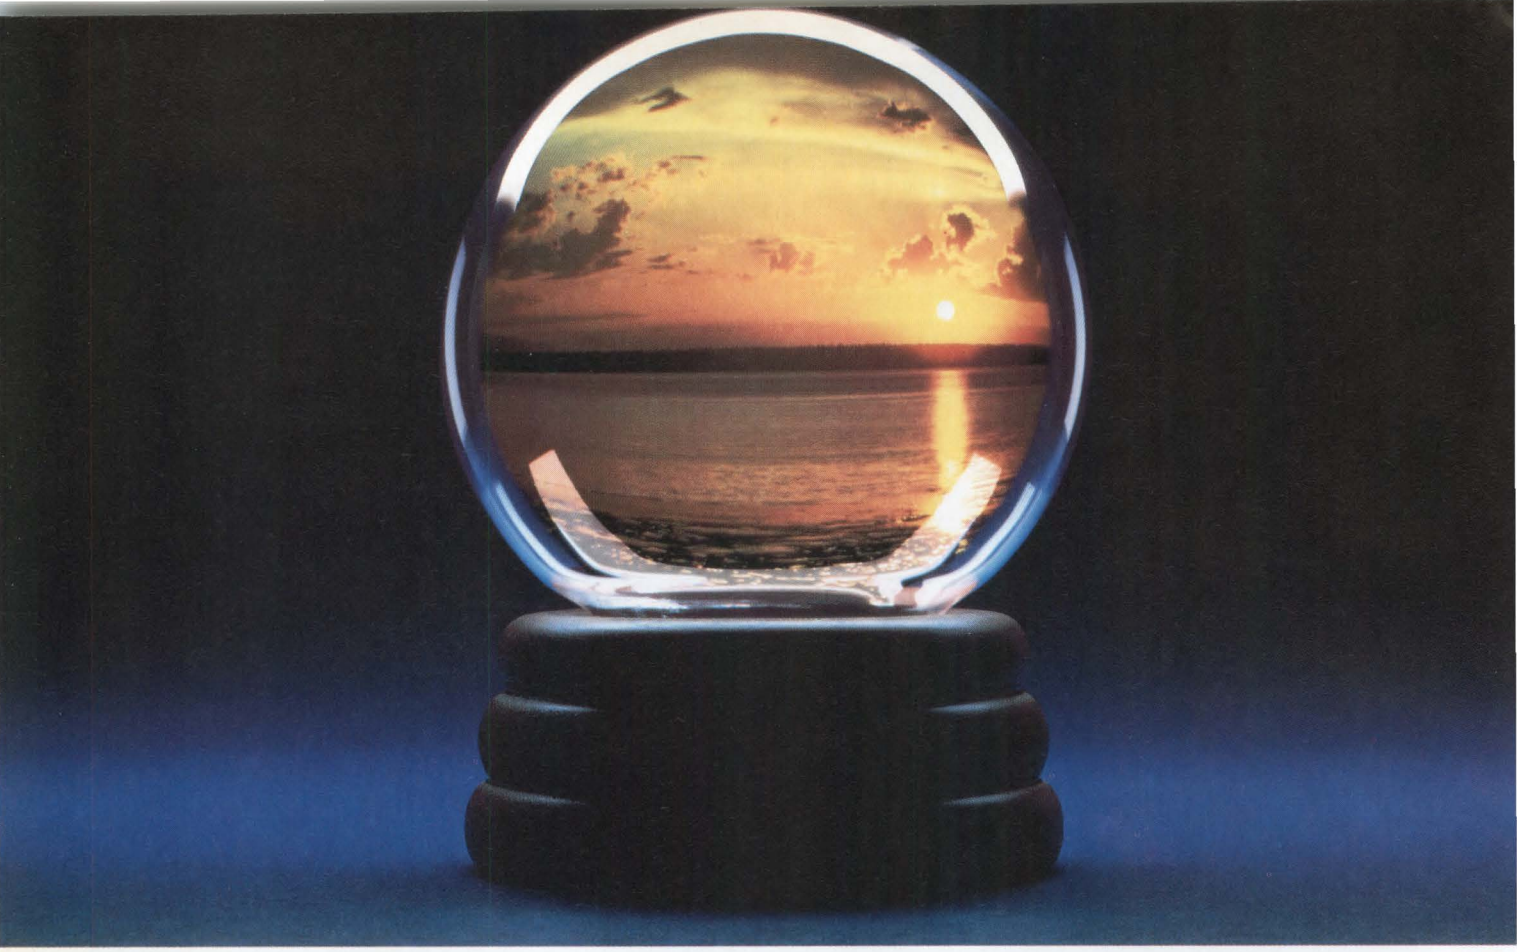

The Strategic Supplier Develops Technology For The Future To Keep You Ahead Today.

In a dynamically changing marketplace, anticipation is the key to profitability and growth. The company that rests on the laurels of current success risks being left behind in today's competitive drive for better performing, higher quality products.

GE Plastics accelerates product advances by keeping its eye on the future with the commitment to match its development efforts to yours. The broad array of increasingly capable resin grades and totally new families of polymers available from GE Plastics today are the result of looking ahead yesterday. And tomorrow's even better materials are in sight right now.

The strategic supplier is one with the resources to focus on your evolving product needs and provide timely solutions. The strategic supplier is General Electric Plastics.

LEXAN® sheet/film/foam/resin
NORYL® foam/resin
VALOX® film/foam/resin
ULTEM® foam/resin
XENOY® resin

® Registered Trademarks of General Electric Company

CIRCLE 34

We bring good things to life.

GENERAL ELECTRIC

Monolithic A-D ICs

(continued from page 68)

trimming. Such systems could be inaccurately aligned and also can drift so that the PLL may contribute to high system error rates. Since the chip requires no trimming and will not drift, signal transmission is more reliable and accurate. The chip takes data stored in MFM format on the disk and converts it to serial nonreturn to zero (NRZ) data. Current data rates are 2 to 15 Mbits/s, and future versions will operate up to 25 Mbits/s. Power consumption for this advanced low power Schottky (ALS) chip is 500 mW maximum and 300 mW typical.

Various approaches to PLL and data separation are being tried. Toshiba Corp (Kawasaki, Japan), for example, has developed a high speed PLL IC for disk drives that are decreasing in size while increasing capacity due to higher recording density and faster transfer rates. The Toshiba IC was structured from 110 ECL gates and operates up to 30 Mbits/s.

The device contains two phase comparators and a voltage controlled oscillator. It separates data written in run-length limited codes as well as MFM data. Both phase shift and jitter are each less than 1 ns. Like the National Semiconductor device, it can be used with optical as well as magnetic drives. It is especially useful in high speed applications such as communication circuits, wideband video tape recorders, and optical disks because it was designed for future magnetic disk drives with very high data rates.

Another example of analog/digital technology is the single-chip servo control IC developed by Control Data Corp's Microcircuits Division (Bloomington, Minn). This IC demodulates the position error signal and provides all spindle and head positioning motor control functions. The bipolar chip operates in combination with a MOS microprocessor. The processor provides digital control functions while the servo chip takes care of analog functions.

Closed-loop servo systems are used to position the head over the correct track on hard disks. An error in head position is determined by reading specially written position data on the disk. Other analog functions are motor control for the spindle and head positioning servo, and power shutdown when supply voltages are beyond safe limits.

In order to control a 5¼-in. Winchester spindle motor to 0.05 percent accuracy without trimming, Silicon Systems, Inc (Nevada City, Calif) has fabricated a brushless dc motor controller. The analog/digital IC

(continued on page 73)

Microprocessor programming made simple.

"Keep it simple" was the principle of the 14th Century English philosopher William of Occam and it has even more validity today. Faced with the problems of sophisticated computer systems, designers have found that ever more complex programming languages are further complicating their tasks. Until now.

Occam. Created for system design and implementation.

When we started designing our new VLSI family of 10-MIP transputers, we built on William's simple philosophy. To take advantage of the possibilities opened up by the transputer, we needed to create a language capable of properly addressing parallelism and multiprocessor systems.

With the ability to describe concurrency (whether timeshared or real) and to handle message-passing at the lowest level of the language, all aspects of a system can be described, designed and implemented in occam. From interrupt handling through signal processing to screen editors to artificial intelligence. And on.

But occam is not limited to our transputer family. It provides an efficient, responsive implementation language for systems built on today's microprocessors. It also opens up future possibilities with its performance-enhancing multiprocessor capabilities. And INMOS now offers a product to let you exploit occam's total capability in your system.

Simplify your job with the Occam Programming System.

The Occam Programming System (OPS) gives you the tools for complete VAX / VMS software development. This package includes an integrated editor / checker, an optimizing VAX compiler and full documentation. This gives you a supportive environment for the development of occam programs for execution on the VAX. Cross-compilers for 68000 and 8086-based systems will also be available.

What's more, the occam programs developed and proven on the OPS will give you a head start for work with the INMOS transputer. Extensions to the OPS will be available which will allow occam programs to run on the transputer.

And if you have a requirement to program the transputer in other popular high-level languages, other extensions will include compilers for C, Fortran, and Pascal.

Get started today.

Contact us for our information pack on occam, the Occam Programming System and the transputer. You'll be surprised how simple your life can be.

For quick response, call us at (303) 630-4000 or write: Occam, P.O. Box 16000, Colorado Springs, CO 80935.

inmos

Monolithic A-D ICs

(continued from page 71)

employs an internally compensated frequency lock loop to control a two-phase, four-pole disk-drive motor. The controller features an overcurrent loop, a stuck rotor detector, and a supply fault detector. The loop uses pulse amplitude modulation (PAM) control to avoid the switching transients and torque ripple of pulse width modulation (PWM) control schemes. Other motor control chips use PWM schemes, which require a much higher frequency modulation than the motor speed for good instantaneous speed regulation. This causes high frequency noise in the power supply and adjacent circuitry. PAM, on the other hand, maintains linear control over the motor current to reduce high frequency noise and supply current spikes.

Analog functions go onchip

Except for D-A and A-D converters, analog circuits have been less complex than the current state-of-the-art for linear/analog ICs permits. This has occurred because it is difficult to identify universal analog circuit blocks larger than multiple op amps. Also, linear circuits may require offchip resistors and capacitors to customize the chips for particular functions. Because they are selected for specific applications, these elements cannot easily be brought onchip.

The need for 5¼-in. Winchester drives, however, is causing semiconductor designers to take a close look at the functional analog blocks that are required to shrink the size of disk drive electronics. Many of these new analog/digital circuits will be custom chips specialized for a particular function. The Control Data servo control IC is such a device, designed by a captive IC house for a specific application.

However, to satisfy the vast market for disk control chips, merchant semiconductor suppliers will identify and produce more standard analog/digital blocks. The DP8460 data separator is a good example of this trend. Athough it still needs some external passive components to tailor it to specific drives, these components are fewer in number and fixed in value (the resistors and capacitors are not required to be adjustable). Furthermore, the data separator is only one of a family of four disk electronics chips, which also includes a pulse detector, an encoder/decoder, and a data controller. This integration of analog and digital functions will continue at the interface to mechanical devices and the designers of disk-drive electronics will lead the way.

—John Bond,
Senior Editor

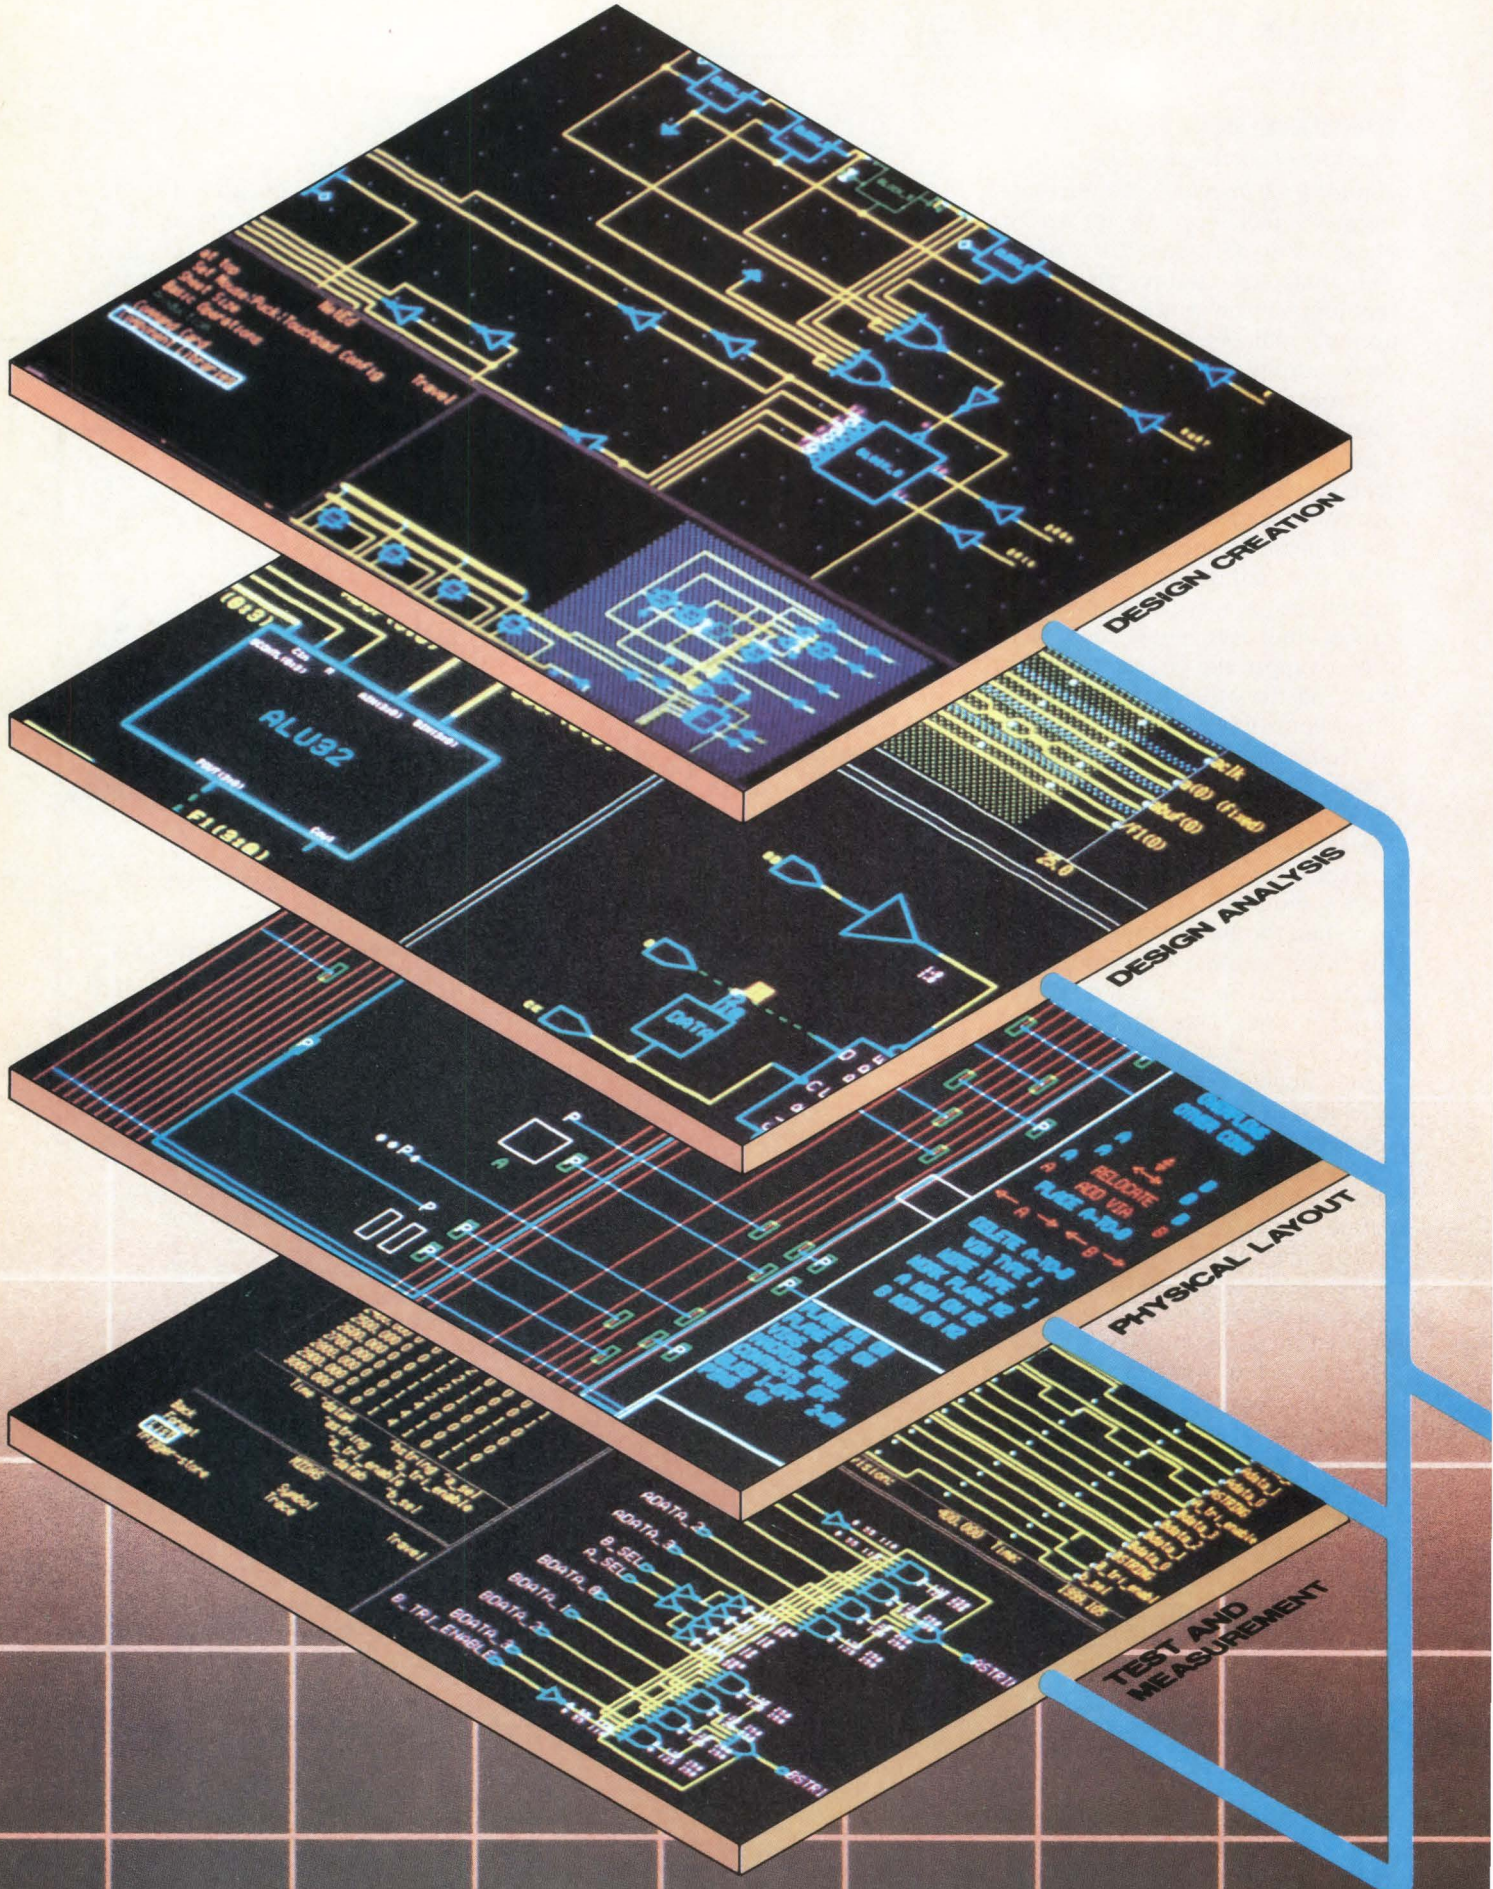
DESIGN CREATION
DESIGN ANALYSIS
PHYSICAL LAYOUT
TEST AND MEASUREMENT

Your Mentor now makes design engineering four steps faster.

The concept of Computer-Aided Engineering has now reached its logical and highly productive conclusion:

A non-stop path from schematic entry to debugged prototype hardware. One that anticipates all your engineering needs and has the right tools waiting every step of the way.

The Mentor Graphics IDEA 1000.

Capture schematics and creativity as well.

With Mentor's powerful graphics-driven interface, the transition from concept to symbolic circuitry has never been faster.

And in addition to "flat" schematics, the IDEA 1000 lets you create an entire hierarchy of design data. From function diagrams down to transistors, you have a better conceptual grasp of your design.

Save time and money through simulation.

Mentor's digital and analog circuit simulators let you bypass much of the expense and labor associated with breadboard prototype circuitry.

These simulator tools simply access the software version of your design which resides in the IDEA 1000 system database. You head off most hardware problems before they're even physically realized.

Automate physical layout tasks.

When you're ready to take your design to physical layout, Mentor's integrated tool set tracks right along with you.

Our CADISYS gate array layout tools deliver true state-of-the-art performance. The entire gate array layout process can be completely automated from start to finish.

Use integrated logic analysis.

Mentor Graphics completes the hardware design cycle through MIDAS 7000, a fully integrated logic analysis system.

The same Mentor workstation that helped you produce your hardware will now help you verify its functionality. You can even compare real-time data acquisitions with earlier simulation runs.

Your Mentor puts it all together.

The IDEA 1000's computerized and integrated design environment is the key to faster, better electronic engineering. Contact us and we'll show you why.

Mentor

Mentor Graphics Corporation
8500 S.W. Creekside Place
Beaverton, OR 97005
(503) 626-7000

Mentor Graphics (U.K.) Ltd.
Phone: 0734-884888

Mentor Graphics (Deutschland) GmbH
Phone: 089/319-1003

Mentor Graphics Japan Co., Ltd.
Phone: (03) 989-7950

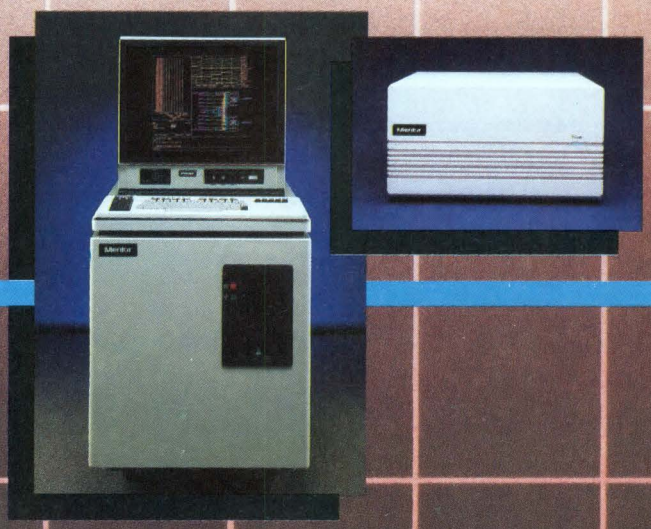

Sell someone a Genicom 3000, and it may be some time before you hear from them again.

Genicom 3000 printer reliability can keep a customer happy for years.

From offices to factories across the country—hour after hour, day after day—Genicom 3000 printers have been proving their quality and reliability under even the toughest conditions for years.

The result has been a large number of very satisfied customers, which means a large number of satisfied OEM's. But durability is only part of the Genicom 3000 printer advantage.

The Genicom 3000 family of printers offers multi-model flexibility combined with single design simplicity to give OEM's real dollar savings with price/performance matching for every customer. Parts commonality. Easier servicing. Single source supply. Plus you can select speeds from 180-500 cps draft/EDP, 45-100 cps NLQ, single or multi-mode printing, automatic sheet feeders, document inserters, multi-color printing and graphics, plus more. There's such a diversity of models, features and options, you can choose just the right printer and you don't have to pay for things you don't need.

See how long you can keep your customers satisfied...with the long lasting, field proven printers that have earned the respect of OEM's nationwide—the Genicom 3000 family.

GENICOM™

The New Printer Company.

Genicom Corporation, One General Electric Drive, Waynesboro, VA 22980 In Virginia, call 1-703-949-1170

For the solution to your printing needs call

TOLL FREE 1-800-437-7468

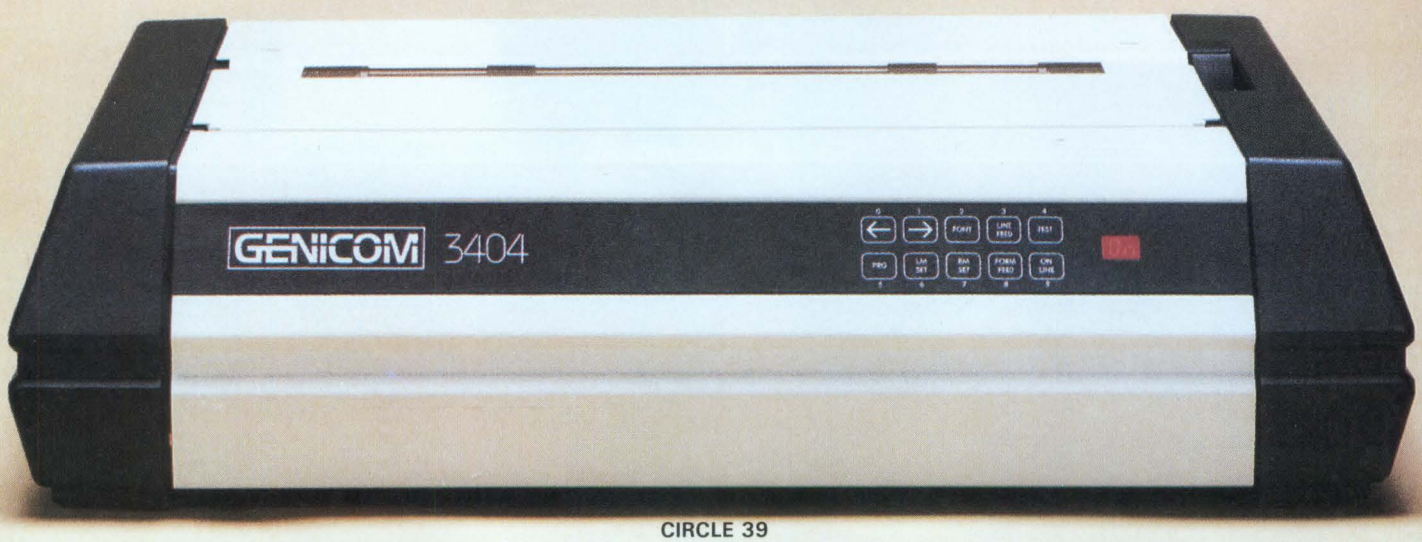

CIRCLE 39

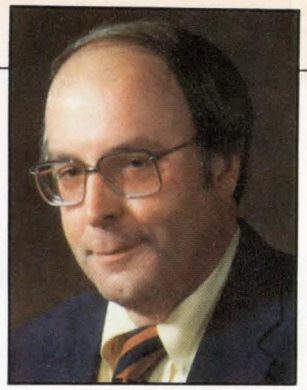

SYSTEM DESIGN BEYOND 1984

In this issue the editors of *Computer Design* attempt the impossible. From 1984, the focus of George Orwell's visionary novel, they peer into the distant future of computer technology. Unfortunately, unlike social history, events in the evolution of technology tend not to repeat themselves. Hence, they defy accurate prediction. In an industry in which innovation expands exponentially, therefore, the only certainty is continuing change.

If our editors were merely reporters, they could sit back and record technological history as it unfolded. But, our editors look at technology from an engineer's standpoint. Like you, our editors are driven to forecast technology not just out of curiosity, but because they know the importance of staying ahead of the game. While market researchers can spot the need for a product, that product does not exist until engineers create it. And meeting almost impossible demands is what engineers do.

In this fiercely competitive industry, the time required to develop a product often exceeds that product's expected life in the marketplace. With development costs soaring, many companies are forced to bet everything on a single roll of the dice. Thus, responsibility weighs heavily on those engineers and engineering managers who must roll the dice on behalf of their financial backers. In this high stakes game, those with the knowledge to predict the outcome have an edge over those who rely on luck. Moreover, the potential rewards climb as the odds increase.

Although it is hard to say where technology's future will lead, it is relatively easy to detect broad trends that will impact the work of system designers—and the training and tools they will need to do the job. One clear trend is increased circuit integration, in which VLSI chips replace subsystems that formerly occupied entire boards or boxes. And, although the versatility of microprocessors will continue to contribute heavily to system design, more applications will require custom or semicustom solutions. Thus, system designers will become more involved in IC design. Spiraling circuit complexity, plus a widely predicted shortage of engineers and the need to shorten development times, will result in a growing need for automated development and testing.

However, it may turn out that the quest for higher speeds—not higher complexities—is the factor that eventually stymies further progress. It could overload the industry's supply of engineering talent and capital equipment. This is because higher speed logic circuits—whether built with silicon, gallium arsenide, or some other more exotic materials—will require analog design skills that are scarce and not easily automated. Also, a background in physics or chemistry will become increasingly important. New skills will be needed as conventional silicon and magnetic technologies give way to electrooptic, cryogenic, or molecular electronic techniques.

In another major trend, software development will occupy a growing proportion of a design engineer's working hours—overshadowing hardware development as we know it. There are many reasons for this shift in emphasis. Of course, with programmable devices such as microprocessors, software represents most of the system designer's contribution. Also, the need to create more powerful and more friendly software for an expanding range of applications adds to the software design burden. With increasing automation of the design, test, and manufacture of hardware, even hardware design itself is largely transformed into a software design problem. Increasingly, new designs will be simulated with software that runs on large computers, instead of being breadboarded at the hardware device level. In some cases, these design simulations will be stockpiled until cost-effective hardware, with the required performance, becomes available. Downstream, in the product flow from the lab to the end user, engineers and programmers will need diverse skills. They will have to tailor computer systems to application areas as varied as scientific research, professional expert systems, and industrial process control.

New trade-offs will further compound the complex decisions facing design engineers. Existing trade-offs such as hardware-versus-software, analog-versus-digital, and cost-versus-performance will start to seem trivial compared to subtler trade-offs such as speed-versus-complexity, storage-versus-processing, and distributed-versus-concentrated architectures. With the merging of communication and computer technology, computers will become as commonplace as telephones and TV receivers are today. In fact, telephones and TV displays will become mere computer peripherals or workstation components. In an age in which workplaces will be located strategically close to their markets or their resources, information—rather than people or materials—will be transported into the workplace.

Yet, paradoxically, as global computing networks bring everybody closer together, engineers may become increasingly alienated from nonengineers and from each other. Limits to further progress may be set, not so much by the expected numerical shortage of engineers, but by a qualitative shortage of engineers. The world will need people who can combine the necessary knowledge and creativity with the psychological fortitude to handle the stress of tomorrow's engineering environment.

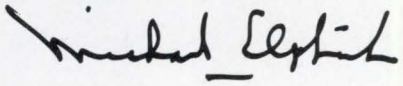

Michael Elphick
Editor in Chief

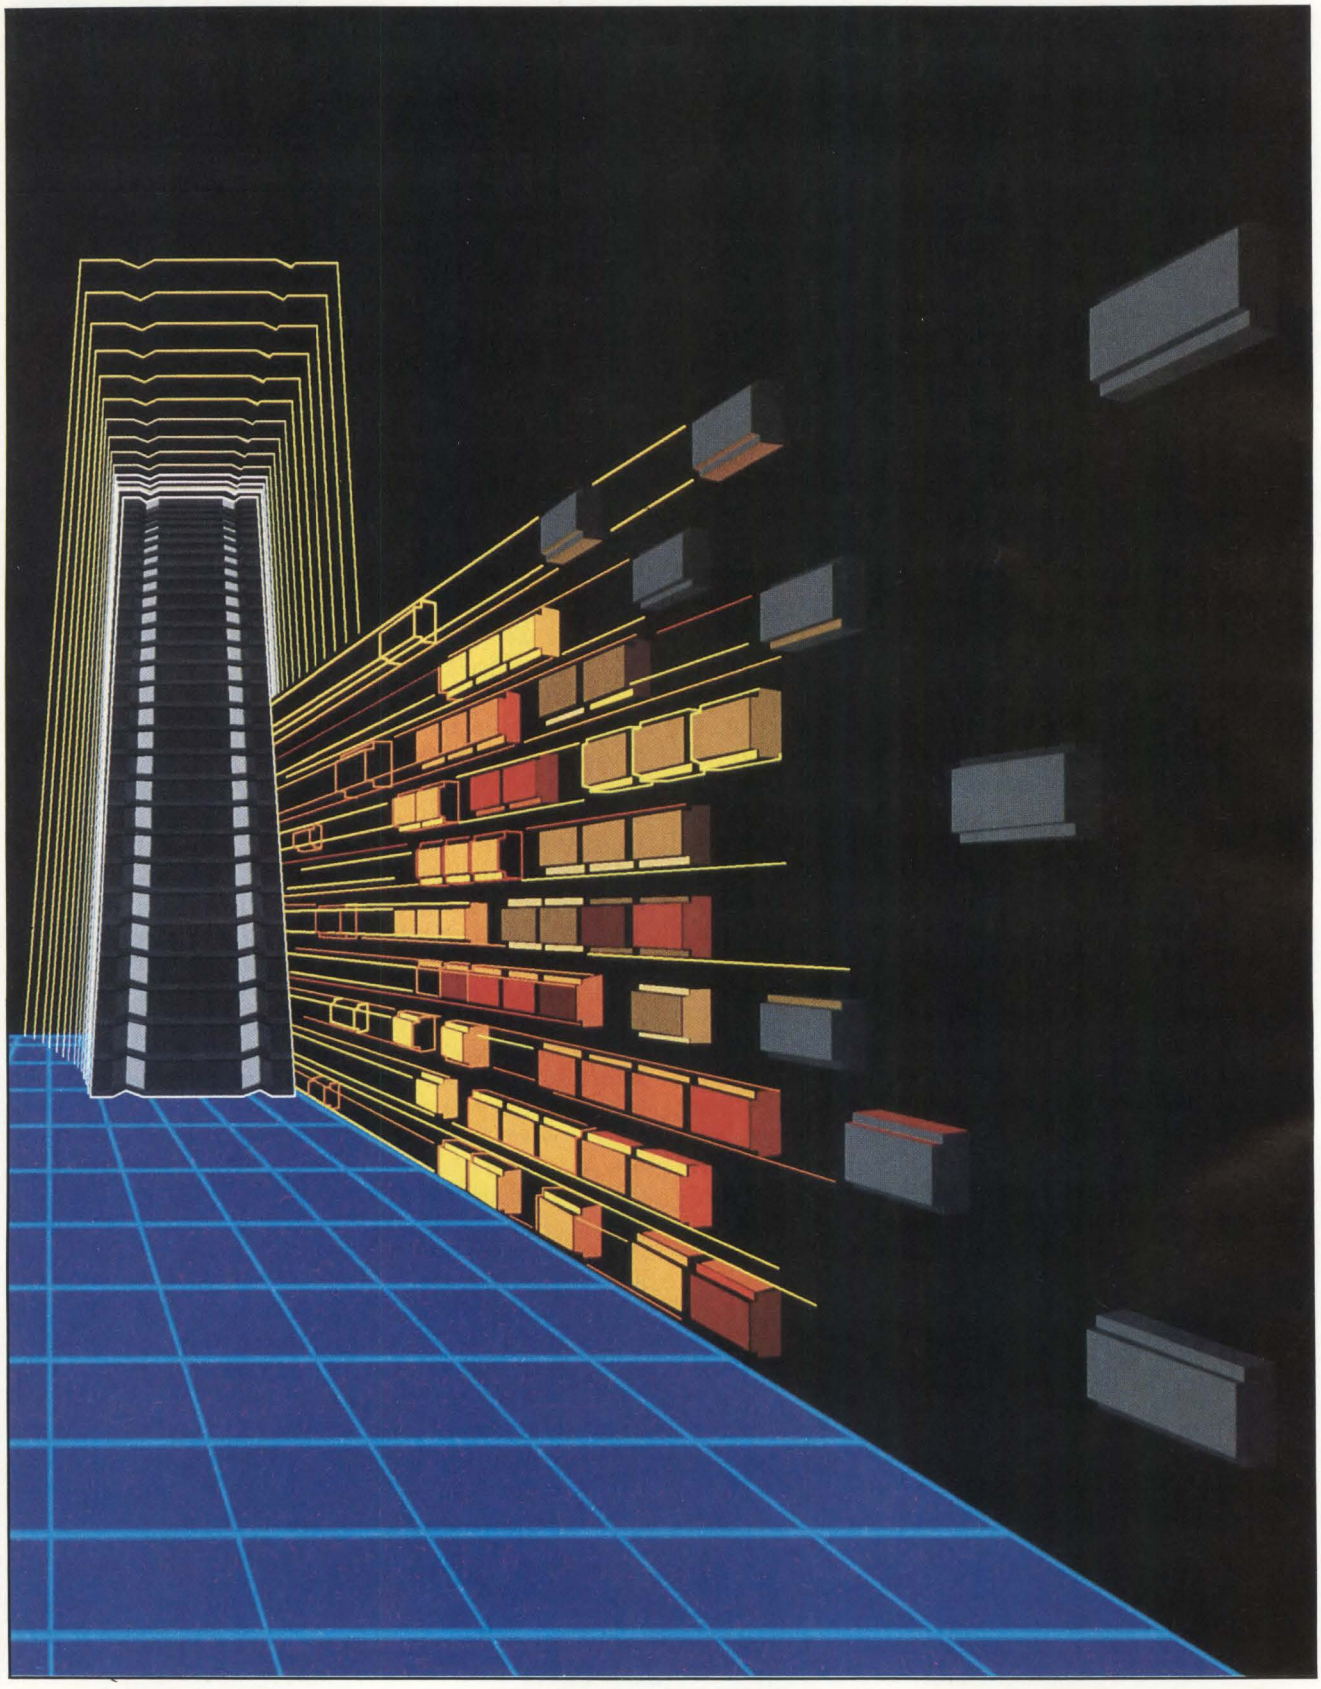

INTELLIGENT COMPUTING ERA TAKES OFF

Computers that manipulate symbols and apply reasoning power promise to transform society while new superprocessors solve engineering problems.

by Deb Highberger, Senior Associate Editor and Dan Edson, Contributing Editor

Some call it the Dawn of a New Computer Age. Others call it the Post-Industrial Revolution. Still others call it the Age of Knowledge. Whatever the name, computers have entered another period of change in which they will be transformed, not simply improved. During the next decade or so, computers will be constructed differently and will operate differently. Most importantly, they will begin to reason and apply logic. These developments will not merely produce a dramatic change in the role of computers worldwide; they will cause a dramatic effect on society as well.

Aggressive computer research and engineering efforts are concentrated in two areas: the development of intelligent, logic-based, and interactive fifth-generation computers; and faster, more easily accessed supercomputers. The symbol-manipulating, fifth-generation computers will use non-von Neumann architectures, advanced software, artificial intelligence, and new hardware materials. They will be able to receive voice and sensor input, access arrays of knowledge and data bases worldwide, and solve problems using reason and inference. With this newfound intelligence, these machines will cease to function only as sophisticated data manipulators and will act more like intelligent assistants.

The new supercomputers will depend on VLSI circuitry, parallel processing, and gallium arsenide components. They could possibly process information at least 1000 times faster than today's supercomputers, plus offer greatly enhanced user interfaces. These improvements will expand their usefulness and extend their application base. Fifth-generation computers and high speed supercomputer

systems promise to have a tremendous impact on lifestyle as they become common in business, education, engineering, industrial automation, law, medicine, retailing, and scientific research.

Race for supremacy

The innovations in computer architecture, software, and hardware explored in this issue are critical to the success of these new computer systems. Just as important, however, are the consortia, cooperatives, and other uncommon precompetition research alliances that have recently been formed in the United States, Japan, and Europe. These collaboratives, which involve industry, government, and university research laboratories, have crossed international borders and tested antitrust laws. And, by bringing competing companies together to perform basic research, they may have challenged the very tenets of free enterprise.

As research collaboratives bring competing companies together, they may challenge the tenets of free enterprise.

The major participants in this intensive computer research can now be identified. In the United States, the lead organizations are the Defense Advanced Research Projects Agency (DARPA), the Microelectronics and Computer Technology Corp (MCC), and the Semiconductor Research Cooperative (SRC). Japan has the Institute for New-Generation Computer Technology (ICOT) and the National Superspeed Computer Project. The European Strategic Program on Research in Information Technology (ESPRIT) heads the European effort. And, in Britain, there is the Alvey Program for Advanced Information Technology.

Some of the other collaboratives, especially in the United States, may be somewhat less well known. DARPA, an agency of the Department of Defense (DoD), and MCC, a collaborative of 18 electronics companies, are supported by the SRC, an industry-academia collaborative, and the Microelectronics Center of North Carolina (MCNC), a state, government, and academic consortium. In addition, the Microelectronics and Information Sciences Center (MEIS) at the University of Minnesota and the Stanford University Center for Integrated Systems (CIS) are two of several universities that have formed research cooperatives.

The goals of each group also need to be defined. [It is not always clear if a particular research consortium is dedicated to the fifth-generation computer, the supercomputer, or computer technology in general. The fifth-generation computer and supercomputer are two different products, but considerable overlap in their research exists (eg, parallel architectures and GaAs) and much of the work will be mutually beneficial.] DARPA's purpose is to advance machine intelligence technology in defense applications. The agency has as a target machines with 1000 times the processing power of today's computers, using AI-based software to solve complex, reasoning problems. The agency will be studying vision, speech, and expert systems with a specific interest in developing adaptive, intelligent, logic-based weapons; autonomous land, sea, and air vehicles; assistive systems; and battle management systems.

To focus the agency efforts in computer technology research, DARPA established the Strategic Computing Program last year. The program oversees the research and design of intelligent, interactive computers. Assistant Director Lynn Conway, who was previously involved in VLSI system architecture and design methodology at Xerox Corp, is Strategic Computing's chief proponent. DARPA has earmarked more than $1 billion during the next 10 years ($50 million for fiscal year 1984, $95 million for 1985, and $150 million for 1986) for research in AI, software, and computer architecture. Currently, DARPA has research contracts with government, industry, and university laboratories. It intends to make this newfound information available to industry. [Indeed, DARPA deserves credit for keeping AI research healthy in the United States.]

MCC, a private corporation, is often considered the civilian counterpart of DARPA. Staffed and funded by its 18 shareholder companies, MCC was established in 1983 as a direct response to Japan's 1981 announcement of its intention to market fifth-generation computers in the 1990s. A ruling by the Justice Department last December removed antitrust barriers from MCC. [This was an action that would not be necessary in Japan, where the electronics industry was exempted from anti-monopoly laws more than 25 years ago.] The express purpose of this corporation is worldwide dominance in computer technology through technological leadership.

Directed by Bobby Inman, a retired admiral and former deputy director of the Central Intelligence Agency, MCC is concentrating its research in four areas: microelectronic packaging, CAD/CAM, advanced software technologies, and computer architectures. The shareholders of the Austin-based corporation—Advanced Micro Devices Inc, Allied Corp, BMC Industries Inc, Control Data Corp, Digital Equipment Corp, Eastman Kodak, Gould Inc, Harris Corp, Honeywell Inc, Lockheed Missiles & Space Co Inc, Martin-Marietta Corp, Mostek Corp, Motorola Inc, NCR, National Semiconductor, RCA

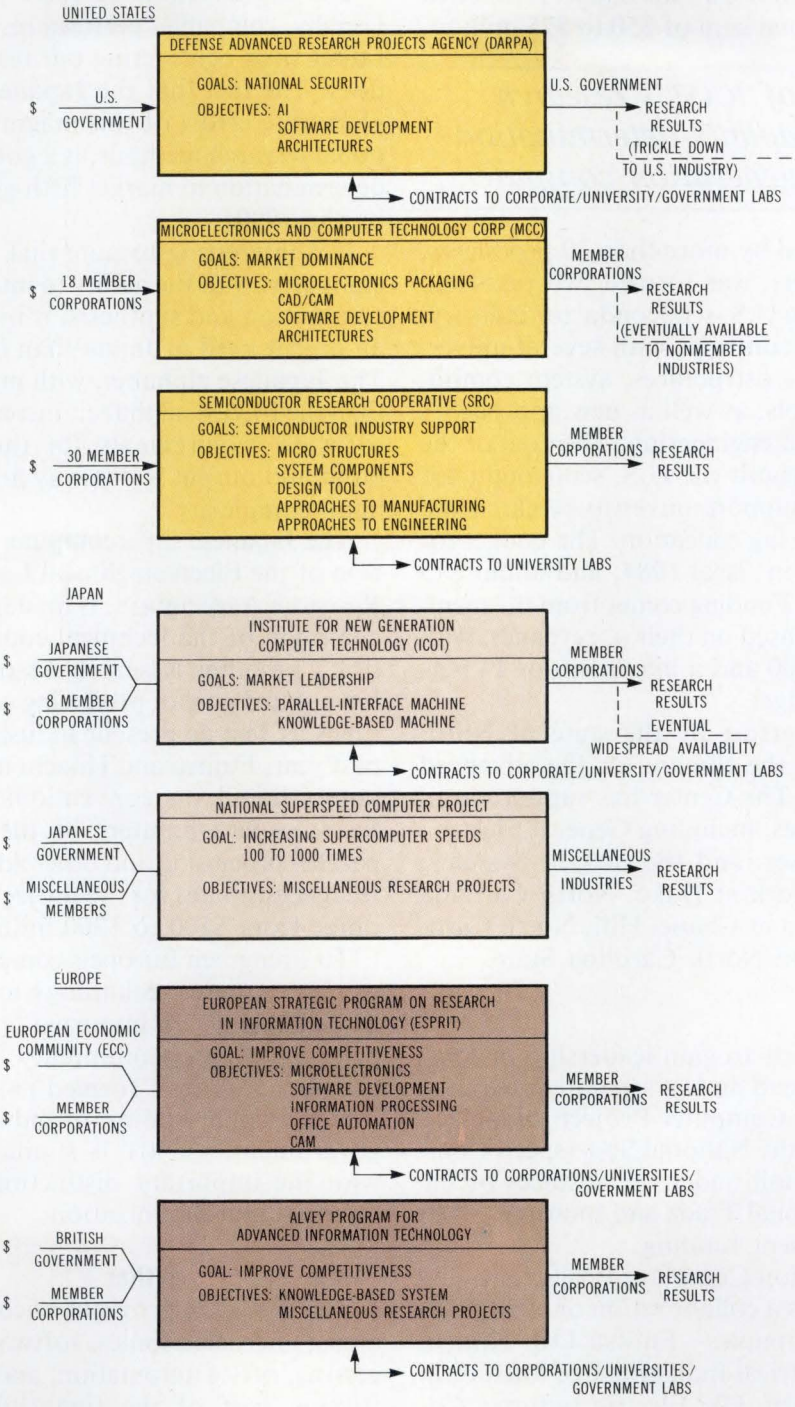

Fifth-generation computer and supercomputer projects worldwide encompass a variety of organizations with different goals, objectives, and funding arrangements.

Corp, Rockwell International, and Sperry Computer Systems—bought one $500,000 share of stock each when they signed on, and pay a portion of the research expenses for the area or areas they participate in. The $30 million fiscal 1984 budget is expected to increase to an annual sum of $50 to $75 million.

The magnitude of ICOT's research efforts proves Japan's determination to market fifth-generation computers.

The SRC, supported by more than 30 producers, users, and IC suppliers, was formed two years ago to buttress the sagging U.S. semiconductor industry. The cooperative has contracts with several universities to research microstructures, system components, and design tools, as well as new approaches to manufacturing and engineering. The goal of the SRC is not only to benefit the U.S. semiconductor industry, but also to support university research and to encourage engineering education. The budget for SRC is $13.5 million in fiscal 1984, and about $15 million in fiscal 1985. Funding comes from the members, who pay a fee based on their IC revenues, with a minimum of $60,000 and a maximum of 14 percent of the total budget.

The MCNC is an effort by the state of North Carolina to enhance the climate for the advanced electronics industry. The Center has support from a number of industries, including General Electric, Airco Industrial Gases, and the GCA/IC Systems; and from funding work at Duke, North Carolina A&T, North Carolina at Chapel Hill, North Carolina at Charlotte, and North Carolina State.

Japan's initiative

The Japanese efforts to gain leadership in computer technology depend on two separate programs: the Fifth-Generation Computer Project (organized as ICOT in 1981) and the National Superspeed Computer Project. Both fall under the auspices of the Ministry of International Trade and Industry (MITI) and receive government funding.

The Fifth-Generation Computer Project, headed by Kazuhiro Fuchi, is a collaboration of eight quite vertical member companies—Fujitsu Ltd, Hitachi Ltd, Matsushita Electrical Industrial Co, Mitsubishi Electric Co, NEC Corp, Oki Electric Industry Co, Sharp Co, and Toshiba Corp; with Nippon Telephone & Telegraph Public Corp and the Electrotechnical Lab of MITI as partners. ICOT was budgeted about $88 million for the three fiscal years through 1982 to 1984, split between the participating industries and the Japanese government. The participants have provided ICOT with more than 80 researchers, who perform their work in a Tokyo lab. ICOT's primary goals are to produce a parallel interface machine and a knowledge-based machine, which will ultimately be merged to form the fifth-generation computer.

The organization of ICOT is unique in Japan. Usually, companies perform research independently, rather than contracting out research work as ICOT does. The fact that the Japanese assembled a conglomerate effort of this magnitude, using untraditional research methods, is a good indication of their determination to market fifth-generation computers by the 1990s.

It is interesting to note that the fifth-generation computer, with the ability to interact through speech recognition and synthesis, is in some ways more of an urgent need in Japan than in the United States. The Japanese alphabet, with many more characters than the English alphabet, increases considerably the software requirements for the Japanese. Speech input and output is one way around that programming complexity.

The Japanese supercomputer project, a collaboration of the Electrotechnical Lab and the Computer Research Association, is managed by Shigeru Sato, chairman of the technical committee of the Scientific Computer Research Association. The project has a stated goal of producing computers 100 to 1000 times as fast as present industry standards. In the past year, Fujitsu and Hitachi introduced supercomputers based on conventional technologies. The faster supercomputers of the 1990s will employ parallel processing and other advanced technologies. Budget estimates for the project during 1982 to 1990 range from $100 to $200 million.

To strengthen Europe's competitive position in the world computer technology market and to reduce its dependence on imported technology, the European Economic Community (EEC)—the common market countries—formed ESPRIT. This organization is a collaboration of industry, academia, and government. ESPRIT is similar to MCC and ICOT with the important distinction that ESPRIT is an international organization.

Europe bands together

The research program encompasses five major areas: microelectronics, software, information processing, office automation, and CAM. The EEC contributes half of the financial support, with the remaining funds coming from member companies. Total funding for the multiyear program is being tabbed at roughly $1.5 billion. Support contracts are being given to research labs and universities.

Concurrently, Britain has established the Alvey Program (named after John Alvey, who chaired the committee that established the program in 1982). An undisguised response to the Japanese threat to the

WREN—SUPER-PERFORMING 5¼" WINCHESTERS IN 21 TO 86 MEGABYTE CAPACITIES.

TOLL FREE HOTLINE
1·800·828·8001
EXT. 82
IN MINNESOTA
(612) 921-4400

HEADS never land on data areas. All WREN drives use a dedicated media zone for head take-off and landing. Even in the event of a total power loss to the system, the actuator will return the heads to the dedicated landing zone.

ROTARY VOICE COIL ACTUATOR provides 40 millisecond typical average access for the 21 and 36 MB capacities and a 30 millisecond average typically for the 48, 67 and 86 MB models (includes head settling).

DEDICATED CLOSE-LOOP SERVO SYSTEM for added positioning accuracy and maximum system performance.

INTERFACE FLEXIBILITY: Two industry standard interfaces ST506/412 and Enhanced Small Device Interface (ESDI). The ST506/412 interface is available in all capacity versions up to 86 MB. The ESDI is available in the 48, 67 and 86 MB models.

T H E W R E N™

High Technology from Control Data delivers a 5-1/4" Winchester with truly outstanding performance and reliability. Compare for yourself. Call our Information Hotline 1-800-828-8001 or write OEM Product Marketing, HQN08H, Control Data Corporation, P.O. Box 0, Minneapolis, MN 55440. Also available through your Arrow or Kierulff distributor.

CONTROL DATA

Revitalizing U.S. computer research

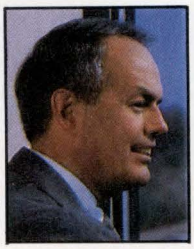

To regain and sustain a favorable position in international competition, U.S. computer manufacturers must maximize the effectiveness of the research engineers now working in industry, while streamlining the time required to commercialize new technology, according to Bobby R. Inman, president of Microelectronics and Computer Corp (MCC) of Austin, Tex. A retired U.S. Navy admiral, Inman became president of MCC after holding the position of Deputy Director of the Central Intelligence Agency. Considering his background, he has a clear perspective of the state of U.S. research and development.

Contrary to what some critics think, the collaborative research underway at MCC is not short-term, but extends well into the 1990s. "The shortest program we have is six years, and half the computer architecture effort is laid out over 10 years," says Inman.

He does not see a need for a centralized, concerted effort by government to focus industry's computer research efforts in this country at this time, although he admits that the government might have to step in eventually. "Industry looks around for where they think there is going to be a market, where there might not be a substantial market, or where the market might be almost exclusively government—such as research in the energy area or weather," comments Inman. "The government may have to sponsor their own efforts or guarantee purchase."

He elaborates, "Let us see what industry comes up with four to five years down the way. In the area of large number-crunching capabilities, for instance, if Cray, ETA Systems (the Control Data spin-off), or IBM comes up with one of these [high speed supercomputers], then the needs may be met. If their efforts aren't moving along, the government may have to step forward and guarantee purchase in order to ensure that we don't surrender the fastest capacity to the Japanese."

Inman does feel, however, that government support is needed immediately in other areas, such as antitrust legislation, educational grants, and most importantly, upgrading university research facilities. "When it comes to actual future design of computers, I believe that Congress is going to be reluctant to put up independent government programs unless they are very clearly earmarked toward applications," says Inman. "But, I think the National Science Foundation could be the spark plug for a program to substantially improve the quality of the equipment in the universities."

Inman strikes a cautious note on the current funding enjoyed by the Defense Advanced Research Projects Agency: "The academic centers of excellence at Stanford, Carnegie-Mellon, and MIT largely were created by DARPA. We need to sustain that. My concern is that now that DARPA is turning to some specific applications, there might be a decision to pay for that by giving up the university support. We need more centers of excellence, at least six, maybe eight. We must focus on getting U.S. research revitalized."

"Over the last 15 years we have produced a smaller pool of trained research talent. Even if we now get some energetic programs going, it's going to take years to train additional design engineers," Inman says. "The question is how do you make more effective use of the talent you have in the interim? That's one problem. The other is taking technology to the marketplace faster. The [U.S.] record for creating technology is very good. But in the '60s and '70s, the technology that had originated here frequently was taken to the marketplace quicker by the Japanese."

Inman does not feel that the transfer of technology and international competition is hobbled by the heavy emphasis on defense spending in computer technology in this country. Instead, he points to the current long procurement cycle, or time from research to market, as the root of the problem. "Over the years, investment in defense research has had a dramatic impact on what also becomes available to the commercial marketplace."

To Inman, the problem in the United States is that the stretched out procurement cycle has lengthened the time before the opportunity to commercialize occurs. "In the early '60s," he adds, "the time from beginning research to producing products in the defense environment was four to five years; now it is 12 to 13 years." In Inman's opinion, our competitiveness could be enhanced considerably if that procurement cycle were shortened to about six years.

Photo courtesy of Zigy Kaluzny/Gamma-Liason

information industry worldwide, Alvey is incorporating research in a number of areas, most particularly in the development of knowledge-based systems. Although Alvey is similar to ICOT in organization, Alvey researchers work at their company facilities rather than in one location as ICOT researchers work, and communicate by electronic mail. The total budget for the first five years should approach $500 million.

Other European programs to support the computer industry include government programs in France and West Germany, and a joint research center established by ICL (Britain), Bull (France), and Siemens (West Germany), the three largest European mainframe manufacturers. Though Europe generally is not considered a contender for the leadership role in computer and communication equipment production or service, the region does represent the second largest computer market in the world, and must not be shortchanged. The Europeans, are quite cognizant of their eroded high-tech base and have taken the aggressiveness of the Japanese as a warning. As is evidenced by these programs, they intend to do something about it.

The reasons why MCC, SRC, ICOT, ESPRIT, Alvey, and the other organizations have emerged should be

We seek VARs that have a mighty future.

IBM is looking for companies with real potential, those with a solid foundation for growth.

We want firms that have first-class application solutions that run on IBM products. Companies that hope to become IBM Value Added Remarketers.

It isn't easy to qualify; but if you do, IBM has much to offer.

For starters, IBM VARs can apply for some of our most competitive products: IBM 4300 systems, System/36, Series/1, System/38, the IBM Personal Computer and the System 9000 family.

Then IBM can add strength to your marketing efforts. We can, for instance, help with product literature, direct mail and business show support. IBM also has a wide range of professional classes for VARs. And thanks to an online referencing system used by IBM's own sales force, we can refer prospects with special needs to VARs with unique solutions.

For more information on qualifying as an IBM VAR, simply send in the coupon below or call 1 800 IBM-VARS, Ext. 562.

If you think you're destined for greater things, plant the first seed now.

IBM®

Larry Humphreys
IBM Distribution Channels
P.O. Box 76477
Atlanta, GA 30358

562/9-84

Please send me your free booklet, "Looking for Leaders."

Name_____ Title_____

Company_____

Address_____

City_____ State_____ Zip_____

Phone_____ Product Interest_____

obvious. The members of the volatile, fast-paced computer industry realize that to maximize their research dollars (particularly in meaningful long-term research), minimize risk, eliminate redundancy, and use engineering manpower effectively, it is advantageous to pool resources. In addition, in order to compete successfully in a world market in which their international competitors have joined together, it is necessary to consult and consort.

Whether or not these organizations inhibit free enterprise, however, is debatable. To some degree, they eliminate competition between companies, and obviously, they open the opportunity for monopolistic action. But at the same time, they could be healthy for worldwide competition.

The power struggle

Information technology is the new heavy industry in the world. According to William H. Davidson, a professor of business policy at the University of Virginia and author of *The Amazing Race*—a book on the technology competition with Japan—computer and communication equipment and services will constitute roughly 40 percent of the world industrial value-added. The success of fifth-generation and supercomputer research in the United States is critical for the country to maintain its economic and military strength.

The U.S. computer industry needs the technology to compete in the international marketplace and to strengthen the manufacturing of the sophisticated equipment needs for business, industry, and defense. The world leader in fifth-generation and supercomputers will to some extent set the standards by which that equipment is used, forcing manufacturers in other countries to build equipment compatible with its machines.

Strategically, for the United States, the race for superiority in next-generation computer technology boils down to a weapons race against the Soviet Union. At present, the Soviets reportedly house a larger arsenal than the United States, but a less intelligent one. Therefore, U.S. experts feel that remaining ahead of the Soviets in computer technology should continue to be a strategic priority.

Although public information on the Soviet computer industry is sketchy at best, the Soviets have made significant technological gains in the last decade. Judging from their attempts to acquire U.S. computer equipment (ie, the Digital Equipment Corp VAX incident), they are aggressively pursuing more advanced technology.

Essentially, the United States has neither a government policy on what computer research is in its best interest nor an umbrella organization to oversee, coordinate, and streamline the fragments of research underway. The United States mandate, if there is one, is to ensure national security. The major reason the government is funding computer research through DARPA is to produce smarter weapons than the Soviets.

MCC, which respresents the private sector—and certainly has the blessing of the government—might be viewed as an emulation of government policy. But, truly, the obligation of MCC is to its industrial members, not the country as a whole. DARPA, in its role as the major recipient of direct government funding for computer technology research (and thus the primary funding agent), is the *de facto* U.S. policy. The agency does not claim to coordinate efforts between various research organizations, but it does claim to review the results of research consortia, national labs, corporations, and universities.

The lack of a comprehensive, effective government policy has drawn fire from various quarters, most notably the IEEE. Simon Ramo is a director of TRW, Inc (Philadelphia, Pa) and an outspoken critic of the lack of a government policy on technology. He charged at the 1984 IEEE conference on technology that rather than applying a strategy for technological growth, the U.S. government has acted through "highly politically based actions, few of which have ever had as a deliberate objective helping U.S. industry win contests against foreign competition in the world market."

The U.S. mandate, if there is one, is to ensure national security.

The heavy emphasis on defense spending in this country is a constant source of criticism, and for good reason. Defense work can be inhibiting and trickle down can be slow. Furthermore, because the United States provides military support for Japan and Europe, U.S. military research performed at the expense of international industrial competitiveness could eventually benefit the Japanese and the Europeans. In addition, civilian research efforts may be harmed. The DoD, for instance, classifies much of the material that results from research as "sensitive," or critical to the military. Material in this classification cannot be given to a foreign government, even an ally. This tends to restrict international competitiveness in the specious interest of national security.

Scientists and engineers depend on shared knowledge, particularly in long-term research. However, the information related to certain projects undertaken for DARPA or defense-related industries is often embargoed. [At MCC all immediate work is considered proprietary, though it will be released at some time.] Such limitations in technology transfer

To finish your touch panels on time, we deliver right from the start.

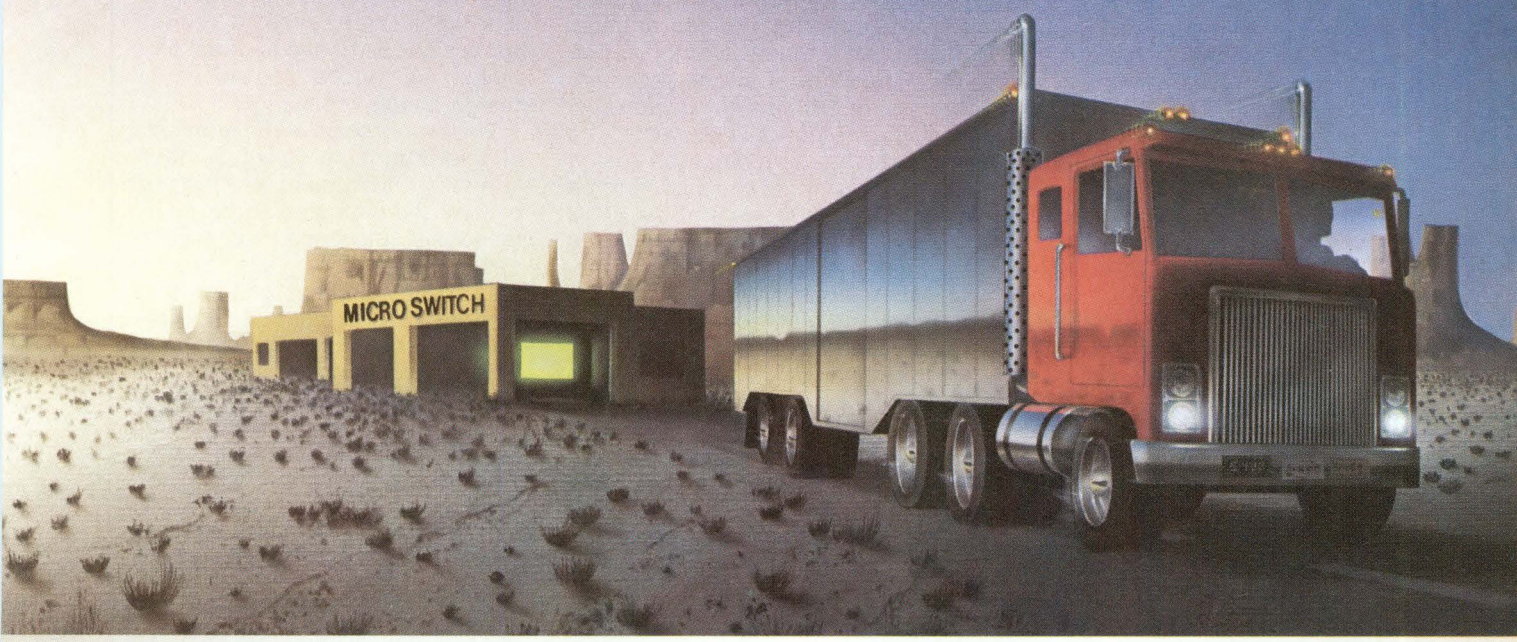

From your initial contact with us to final delivery of your touch panel order, we work with you to make sure the job is done right—every time.

Results you can count on every step of the way.

Right from the beginning, we respond with the service you deserve. For quick turnaround on pricing, our computerized pricing system lets us get back to you fast—sometimes the same day.

When it comes to prototypes, we can ship you a stock design off the shelf the same day. Or, for special needs, our engineers will help you define specifications. Then we'll develop your prototype to your design—and deliver, often in four weeks or less.

Final delivery occurs on time. We don't underestimate production time to get the order. We deliver what we promise.

Out total commitment makes it possible.

It's teamwork that makes it happen. We've put together a talented group of experts—field sales engineers, application engineers, designers, and product specialists. They work with you to help solve your touch panel problems. We've also expanded our extensive production facilities here and abroad. Our new touch panel operation, for example, allows us to respond even faster to special needs, schedules and custom designs.

Our 50 years of experience means you get traditional MICRO SWITCH quality and reliability combined with human factors features, electronics packaging, high quality graphics and a variety of finishes, textures and configurations for almost unlimited design flexibility. It all adds up to value.

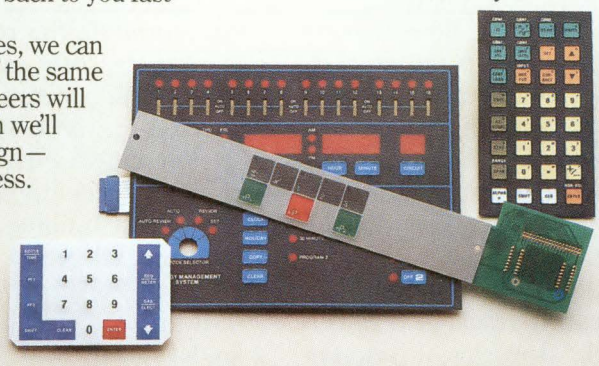

Call on us to get started.

A phone call is all it takes. Our whole team is ready to help.

For more information call 815-235-6600. Or write MICRO SWITCH, Freeport, IL 61032.

From start to finish, you'll like the way we deliver.

Together, we can find the answers.

MICRO SWITCH
a Honeywell Division

CIRCLE 40

DARPA outlines strategic computing as good defense

In response to the challenge of creating and exploiting intelligent computing technology for the military forces, DARPA has initiated the Strategic Computing Program. According to DARPA's Director Robert C. Cooper, the goal of Strategic Computing is to provide the United States with a broad base of machine-intelligent technology that will greatly increase this country's national security and economic power. "This technology should create strong new defense systems for use against mass forces, and thus will raise the threshold and decrease the chances of a major conflict," says Cooper, who became the director of DARPA in 1981 after holding positions at NASA, SBS, and MIT's Lincoln Lab.

Cooper contends that to achieve this goal, a wide range of current technology and recent scientific advances must be leveraged. Engineers and scientists from many disciplines must collaborate in new ways in an enterprise of very large scope. "A framework must be created for the effective, adaptive planning of the discovery and development process in this enterprise. A skillful orchestration of events and exploitation of the available infrastructure will be required to ensure a timely success. Some of these elements are in place today," says Cooper.

Cooper's chief spokesperson for the Strategic Computing Program is Assistant Director Lynn Conway. Conway joined DARPA last year after completing numerous computer architecture and VLSI design projects at the Xerox Palo Alto Research Center (PARC). She is also author of a book on methodical VLSI design. "At DARPA, we have some experience at developing new programs through the normal infrastructure to support research," Conway says, pointing to one example as the MOS Implementation System (MOSIS). This is a software network system that allows local design of hardware prototypes from various design centers around the country.

The design data is fed into the central facility at the University of Southern California Information Resource Institute in Los Angeles using the ARPANET communication network. Thus, students, researchers, and developers can design custom microelectronic systems at a local site, then transmit the design files representing the silicon to be fitted over the ARPANET to the MOSIS center. An automated filing system handles the receiving of these requests, the batching of the requests, and the aligning of many of them into shared assets where some 100 projects may be placed in one fabrication mode.

"This allows many people around the country to quickly implement their systems on the order of one month's turnaround for designing the system, submitting the design file and receiving a hardware prototype," claims Conway. "Currently we are in the process of making the hardware prototyping more responsive, while enabling software prototyping to access control facilities which are more in line with the high performance machines that these designers are used to."

In addition, in order to gain acceptance rapidly over a period of time, especially in the artificial intelligence areas of the program, Conway explains that DARPA will be distributing state-of-the-art Lisp machines to the research community as well as developing higher performance Lisp machines. "This will principally serve the purpose of having the research community identify the needed knowledge of engineering rules that would be most effective for developing a system like an autonomous vehicle," says Conway. Such a concrete application is part of the overall goal for the program. An infrastructure that allows many people to access the same kind of AI information for one particular experimental design will speed up the actual design of the vehicle.

The overall scale of the program's budget requirements will be determined by the number and the type of new technologies to be simultaneously introduced, as well as past experience in the field in generating similar technological communities or centers of expertise. "For this program we require 10 new computing technology communities and another 5 to 10 applications communities," claims Conway. "The intent is to have a small research center with more than 100 professionals spread over two or three research laboratories. These would operate over an 8 to 10 year time span."

While difficult challenges lie ahead, Conway feels that the initial budget of $600 million over the first five years will enable the program to achieve its technological goals, and that a substantial return on invested resources in terms of increased national security and economic strength will result.

not only can stifle research, but can be extremely costly in terms of getting products to the market.

Government restrictions on what material is classified or reclassified have become much tighter under President Reagan. For example, the sometimes vague and subjective rules on exporting information apply to published information, patents, and material presented at conferences and during discussions with other researchers. DARPA's program objectives state clearly that communicating research results is integral to the agency's work. However, it is doubtful that any research lab or university dependent on DARPA support would publish sensitive material and risk the loss of funding.

To stimulate the required research in computer architectures, software, and hardware needed to produce the fifth-generation and supercomputer systems projected for the 1990s, the U.S. government must at least create a positive and encouraging climate. This can be accomplished in a number of ways: providing direct funding, offering tax breaks for long-term research, and passing favorable antitrust legislation are just a few.

Edward A. Feigenbaum, a Stanford professor and a leader in AI research, and Pamela McCorduck, a professor of computer science at Columbia University, conclude in *The Fifth Generation*, a book on which they collaborated, that the United States has

six options in the technology race: maintain the status quo, form consortia, enter a joint venture with Japan, abandon hardware and concentrate on software, form a national laboratory to promote knowledge technology, or form a post-industrial agrarian society. They suggest, not surprisingly, a national laboratory, or, at best, a structure that will manage and encourage knowledge technology research.

Geopolitics

When surveying the international arena, it is difficult, yet important, to look beyond a first impression of the perceived strengths and weaknesses of a given country, company, or research environment and to see what is real. Is the Japanese effort cohesive? Or, are we seeing greater competition between companies? Is the challenge from Japan causing the United States to lose confidence and to forget that it is still the world technology leader? Can ESPRIT be effective, or will chauvinism be the downfall of its effort? Is U.S. research based largely on the military industrial complex, creating a rivalry with the private sector for man and manpower that will retard research? Finally, how do the multinational endeavors effect the race for computer technology supremacy?

In Japan, project goals are centered on the nation's people and economy. The Japanese are responding to the perceived social needs of their crowded island nation into the 21st century, in terms of energy conservation, education, medicine, and social services. At the same time, they are trying to expand their world economic base.

It might appear that the Japanese ICOT, supported by the Japanese government, has an advantage over fifth-generation efforts in the United States which are much more fragmented and not as heavily endowed by government money. Certainly, the United States lacks guidance in terms of setting long-term goals, preventing redundancy, and standardizing software and hardware. But, the ace is not in the Japanese hand simply because the Japanese government takes a more active role. The United States still holds the lead in the computer technology and, thus, plays from a position of strength.

Yet, the Japanese effort cannot be taken lightly. By setting their goals on producing a new generation of equipment (which relies on new architectures, new materials, and new design methodology, rather than improving existing equipment), the Japanese are taking a bold step in international competition. Many critics argue that the Japanese have always depended heavily on U.S. technology (the United States has done the lead work in AI research, for example), and, therefore, will not be able to make the advancements they are striving for very soon. Yet, by all outward appearances, they are determined to gain a technological edge.

The question of price remains unanswered. If the United States and the Japanese were to introduce affordable fifth-generation or supercomputer products at about the same time, would the Japanese come in at a price/performance advantage, as they have in other industries? The Japanese have stated that profit is not their objective in the fifth-generation project. If taken at face value, this would mean that they are after technological superiority, and feel once they have that in hand they will be able to beat the United States in the world market. It is important to keep in mind that the Japanese customarily seek long-term versus short-term gain. The business structure in Japan reinforces that notion: the major source of funds for industry in Japan are banks, which emphasize long-term debt coverage. In contrast, the major source in the United States are stockholders, who press for short-term gain.

There is some concern that even if the Japanese did not dislodge U.S. industry from its leadership position in computer technology by this push to produce advanced computers, the gains made by Japanese manufacturers, who are much more vertical than U.S. electronics firms, would allow Japan to put products on the market long before the United States. Consequently, they may sacrifice one market to capture another.

In this race, the United States holds the lead in computer technology and plays from a position of strength.

The role of multinational companies or companies with foreign operations should have an interesting, though somewhat unclear, part in the race for information technology superiority. Consider Fujitsu's interest in Amdahl, cooperative agreements between Hitachi and Itel, and U.S. firms like Texas Instruments, Intel, IBM, and National Semiconductor, who have design or manufacturing interests in Japan. How much collaboration will they do, or be permitted to do? How will borders break down as competition intensifies between the computer industries in the two countries?

More specifically, how will the joint venture of Matsushita (a member of ICOT) and IBM (a possible supercomputer manufacturer in the United States) accommodate the international jostling? If Matsushita manufactures IBM equipment, such as microcomputers and terminals, what will happen when the United States and Japan begin introducing fifth-generation computers that will tie into the Matsushita-IBM products?

The monumental research efforts underway in this country and abroad are not without uncertainty or technical, sociological, and geopolitical questions.

A question of national policy

"There is no centralized United States policy for computer technology research," according to Earl Dowdy, a research analyst with the Office of Technology Assessment. "The Japanese and Western Europeans have formed national research and development strategies. In comparison to them, the U.S. federal government involvement is very much more mission oriented."

This situation exists primarily because much of the information-technology research today is divided among different government agencies and directed toward the specific goals of those agencies. Through DARPA, the Department of Defense has had tremendous influence on the direction of research. "Most of the government money for computer technology research and development comes through DARPA, and their role is scheduled to increase quite a bit due to the Strategic Computing Program—particularly in AI and advanced computer architectures," Dowdy says.

"This is a concern to a number of people because they feel that their choices of what kind of research they are doing may be limited because DARPA is targeting and writing research agendas with the AI community," he explains. "To realize DARPA's goals, some fundamental research could be put aside."

Dowdy, part of a group of OTA analysts who recently completed a lengthy report on the status of information-technology research and development, adds that military programs can have a negative effect on the transfer of technology and the ability of U.S. firms to compete internationally. "This is a sleeper issue," says Dowdy. "There is a tremendous rift between people in the government—especially in the defense community—who want to restrict the export of technology and companies that have realized that there is a growing internationalization of technology." In order to compete effectively, adds Dowdy, "You have to compete in a lot of foreign markets. Furthermore, individual companies might have contracts with the defense department and with foreign companies." The OTA information-technology research and development report, which was written for the House Science and Technology, and Energy and Commerce committees, is due for publication early this fall.

The lack of government policy appears to be changing in supercomputer research. "The supercomputer issue has galvanized support across a lot of bureaus and agencies because of the Japanese threat," Dowdy reports. "The government relies on these computers for very sensitive national security work. They could be in a quandary if at some point the Japanese machines are superior, and the U.S. national labs have to decide whether to buy Japanese and rely on a foreign source for their computing equipment or use an inferior American machine."

This Japanese threat, which is represented by the recent Fujitsu and Hitachi supercomputers, has allowed the National Science Foundation to gain a quasi-policy-making position. "The science and technology committee was so convinced of the importance of increased funding in supercomputer research that they quadrupled the budget the administration asked for," Dowdy points out.

The urgency of a comprehensive government policy for computer technology research is debatable. As Dowdy observes, industry does not feel that a strong government policy is necessary, as long as government continues to create an environment that encourages research and strengthens the competitive position of U.S. industry.

"It's my impression that many people in the information technology industry don't really want the government to get involved, except in areas where there is great vulnerability and competition, such as in supercomputers," comments Dowdy. "In areas that involve systems and a short-term commercial payoff, they would rather have the government step out of the way and have the assurance of antitrust laws." Dowdy believes that for the most part, industry would like strong government support for basic research in the universities. But, as far as guiding commercialization, such as the Ministry of International Trade and Industry in Japan, he feels that industry does not consider the government to be competent in targeting products.

The research engages all disciplines of computer electronics, from parallel processing, database management, software development, high density storage, and wide bandwidth communications to AI, speech recognition and synthesis, vision and tactile sensing, and interactive graphics. The product areas affected include CAD/CAE/CAM/CIM, robots, realtime systems, and telecommunication. Success in any one area depends not only on progress in that particular area, but a synergy in computer research.

The unknowns

As designers make progress on the fifth-generation and supercomputer systems they will face many system and implementation decisions. What are the trade-offs between distributed versus centralized processing? Is it more efficient and cost effective to have continuous traffic between a host computer and limited satellite terminals, or should the onsite computers have a fair degree of intelligence and talk to the host only periodically? Can outer space handle the satellite networks future communication networks might require? How can satellite communication be made more efficient? And how will all the information and the array of knowledge and data bases available to these interactive, future computers be managed?

One of the most critical elements to successful research in the United States could be manpower. Already, a shortage of qualified computer design engineers has created a lopsided seller's market and a boom-and-bust environment, and has allowed less long-term research than is desirable. The heavy emphasis on military research in the United States

has taken designers from the private sector, which intensifies the competition.

Emphasizing education

According to TRW's Ramo, fewer than 1 percent of U.S. college graduates are engineers, compared to more than 4 percent in Japan. Furthermore, while the percent of engineers in the United States has been declining in recent years, it has been increasing in Japan.

To ensure that an adequate supply of qualified engineers exists, it would be prudent for the government to stimulate engineering education by whatever means possible. Assisting colleges and universities acquire costly, state-of-the-art equipment would help, as would increasing grants to science and engineering students. Furthermore, government programs to guarantee proper grade shcool and high school science and math programs would prepare a younger generation of students for and motivate them toward technical disciplines.

The use of current computer systems at colleges and universities is critical. Engineering students should have the opportunity to become familiar with both fifth-generation systems, as they become available, and supercomputers. At present, only three schools—Colorado State College, the University of Minnesota, and Purdue University—have commercially available supercomputers. This means that only a miniscule number of the engineers who potentially will be using or improving this equipment leave school with any knowledge of it. In contrast, Japan is making a concerted effort to get these computers into its universities. Organizations such as MCNC in North Carolina, MEIS in Minnesota, the Stanford Center, and other industry-supported university research cooperatives can have a powerful effect on engineering education by giving the schools adequate funds to set up laboratories.

Changing workplace

Just as computer designers shape future computers, new and more capable computers will also shape society. In industry, for example, the interactive logic-based, fifth-generation computers will affect the manufacturing process from design to product testing. In finance, instantaneous access to tremendous amounts of up-to-the-second information will begin to influence the way business is conducted.

The fifth-generation technologies will be able to undertake the more mundane, repetitive tasks—and most likely some that are not so uncreative. The machines will become more like educated assistants, to whom a variety of responsibilities can be assigned, than tools that depend on precise input and programming to produce. Whether or not a larger and more productive computer work force means workers do less work is another question.

Without a doubt, the computers designed in the next decade, particularly the fifth-generation computers, will permanently change the shape of our work, and perhaps even our work ethic. In an extreme case, the computers could create a population of a few high level managers, technicians, laborers, and many unemployed. But, in all probability, they will create jobs as well as eliminate them, though the ratio of jobs eliminated to jobs created is not easy to estimate.

Workers should begin to find, however, that changing jobs and retraining is a fact of life. Few people will work at the same job for a career, but rather will work at a job until it is automated, then retrain, move on to another job, and repeat the process. In response to this migratory work force, it will be necessary for the federal government to keep informed of the changing employment in various industries and respond to the fluctuating employment levels with ongoing worker retraining programs.

Fifth-generation computers may change the shape of our work and, perhaps, our work ethic.

These new computers are expected to move steadily into the home, assisting in financial management, energy management, and serving as terminals for cottage industries. Additionally, the intelligent computers should find more extensive applications in education as their ability to apply logic and reasoning improves their interaction with students.

Today's personal computers will also be affected by the introduction of intelligent computers. However, several questions remain: whether or not personal computers will become I/O devices that can tap information networks worldwide (and thus rely on communications for their data), whether or not they will gain intelligence, and whether or not they will begin to take on the attributes of fifth-generation computer systems.

Of course, the effectiveness of fifth-generation systems in the workplace, the home, and in schools will be advanced by the widespread use of networking. This will occur as the computer and telecommunications industries blend into an electronic net crisscrossing the world.

The results of a market analysis by Arthur D. Little (Cambridge, Mass) project phenomenal growth in both fifth-generation computers and supercomputers. Frederic G. Withington, vice president of information systems at A. D. Little, reports

Engineering crisis could hamper high-tech growth

Whether or not there is a shortage of engineers has been under debate in the engineering community for some time. Ray Stata, president and chairman of the board of Analog Devices (Norwood, Mass), clearly sides with those who maintain that there are too few engineers. Stata, an outspoken critic of engineering education in the United States, says that, "Thousands of qualified high school students are denied access to engineering eduaion because of capacity limitations. What's more, the quality of this education is compromised by limited faculty and outdated equipment."

As a result, according to Stata, the number of engineers graduating per capita in the United States is the lowest among competing industrialized nations. Furthermore, Stata does not restrict his criticism to education. He finds fault with industry as well, asserting that "experienced engineers are opting out of the profession for careers in management almost at the same rate as new engineers enter the profession." Stata develops these points in his 1982 coauthored book *Global Stakes.*

More recently, Stata has coauthored *The Innovators*, scheduled for publication this month. In it, he maintains that through better cooperation among industry, government, and academia; and more creative innovations, the United States can begin to meet the challenges of the high-technology industry for the rest of this decade and into the 1990s. Stata says that the country as a whole "has gradually come around to realizing that the technical resources of our universities may be the gating factor to industrial progress, not only in high technology, but also in mature and service industries.

Stata calls high technology companies the toolmakers of the 1980s and 1990s because their job is to embody new technology in hardware and software for use by workers in all industries. Stata maintains that if engineering education moves in this direction, both the innovative high-technology companies and the mature industries will reap the benefits. "The industries using new high-tech tools will themselves need technical workers in order to select, apply, and maintain these tools," says Stata, "Thus, high-tech and mature industries will increasingly share a pool of critical resources—namely the output from our engineering schools."

Stata concedes that there are many who argue that the United States is on its way to saturating the engineering market. "When you consider the unemployment record of engineers over the past 20 years, it appears that the risk of over producing is small," responds Stata. "If we produce too few engineers, the penalty will be significant," he adds, citing that less new technology leads to fewer new products and, thus, to lower productivity growth and living standards.

Stata has been very active in promoting cooperation between government and industry. He is the founder and first president of the Massachusetts High-Technology Council, which mandates that its members contribute two percent of their research and development expenditures to technical universities. Based on the number of members, that amount should be close to $15 million.

Actually, close to $40 million was contributed in the program's first year. Stata advocates a so-called High-Technology Morals Act, which would have the government match the grants to education to grants from industry. The Act is part of the 1984 Education bill currently in Congress. In addition, Stata serves on the Commission of Higher Education and the Economy of New England, which was created by the New England Board of Higher Education.

As he looks toward the industry's future, Stata forewarns that if engineering education and research is suffering and "we sit still and do nothing, we suffer permanent loss of market share and jobs in high-tech growth markets to our foreign competitors." Stata concludes. "We simply cannot afford the risk of having too few engineers."

that revenues from the world market in 1993 for "expert systems and associative processors" (roughly fifth-generation project category) will increase from a couple of hundred million dollars today, to $3 to $6 billion dollars. In the supercomputer market, he predicts growth to $1.5 to $3 billion, from $350 million today. As the world changes, or is changed, by the swing to an information economy, the power in the world will begin to be concentrated in knowledge. Programming and the ability to handle facts and to make decisions will become the value-added commodity—low in energy costs, high in demand. Predicted applications for the smart, knowledge-based expert systems (which use AI-based software to mimic a human pundit in a defined expertise) are vast. They range from oil rigs, geological exploration, and factory automation to business, law, medicine, tax assistance, even recreation and art, and the military. Indeed, according to William B. Gevator, who has written extensively on AI, and recently published *Artificial Intelligence, Expert Systems, Computer Vision, and Natural Language Processes*, AI is already having an impact. The maturity of AI technology can be exemplified by the Digital Equipment Corp program used to configure VAX systems, by personal computers that run on Lisp (the language preferred by most U.S. AI researchers), and by the fact that the AI technology is becoming codified.

Widespread use of these new machines will depend not only on the success of parallel processing and the availability and successful development of GaAs components, but also on making the man/machine interface so comfortable it is second nature. In order

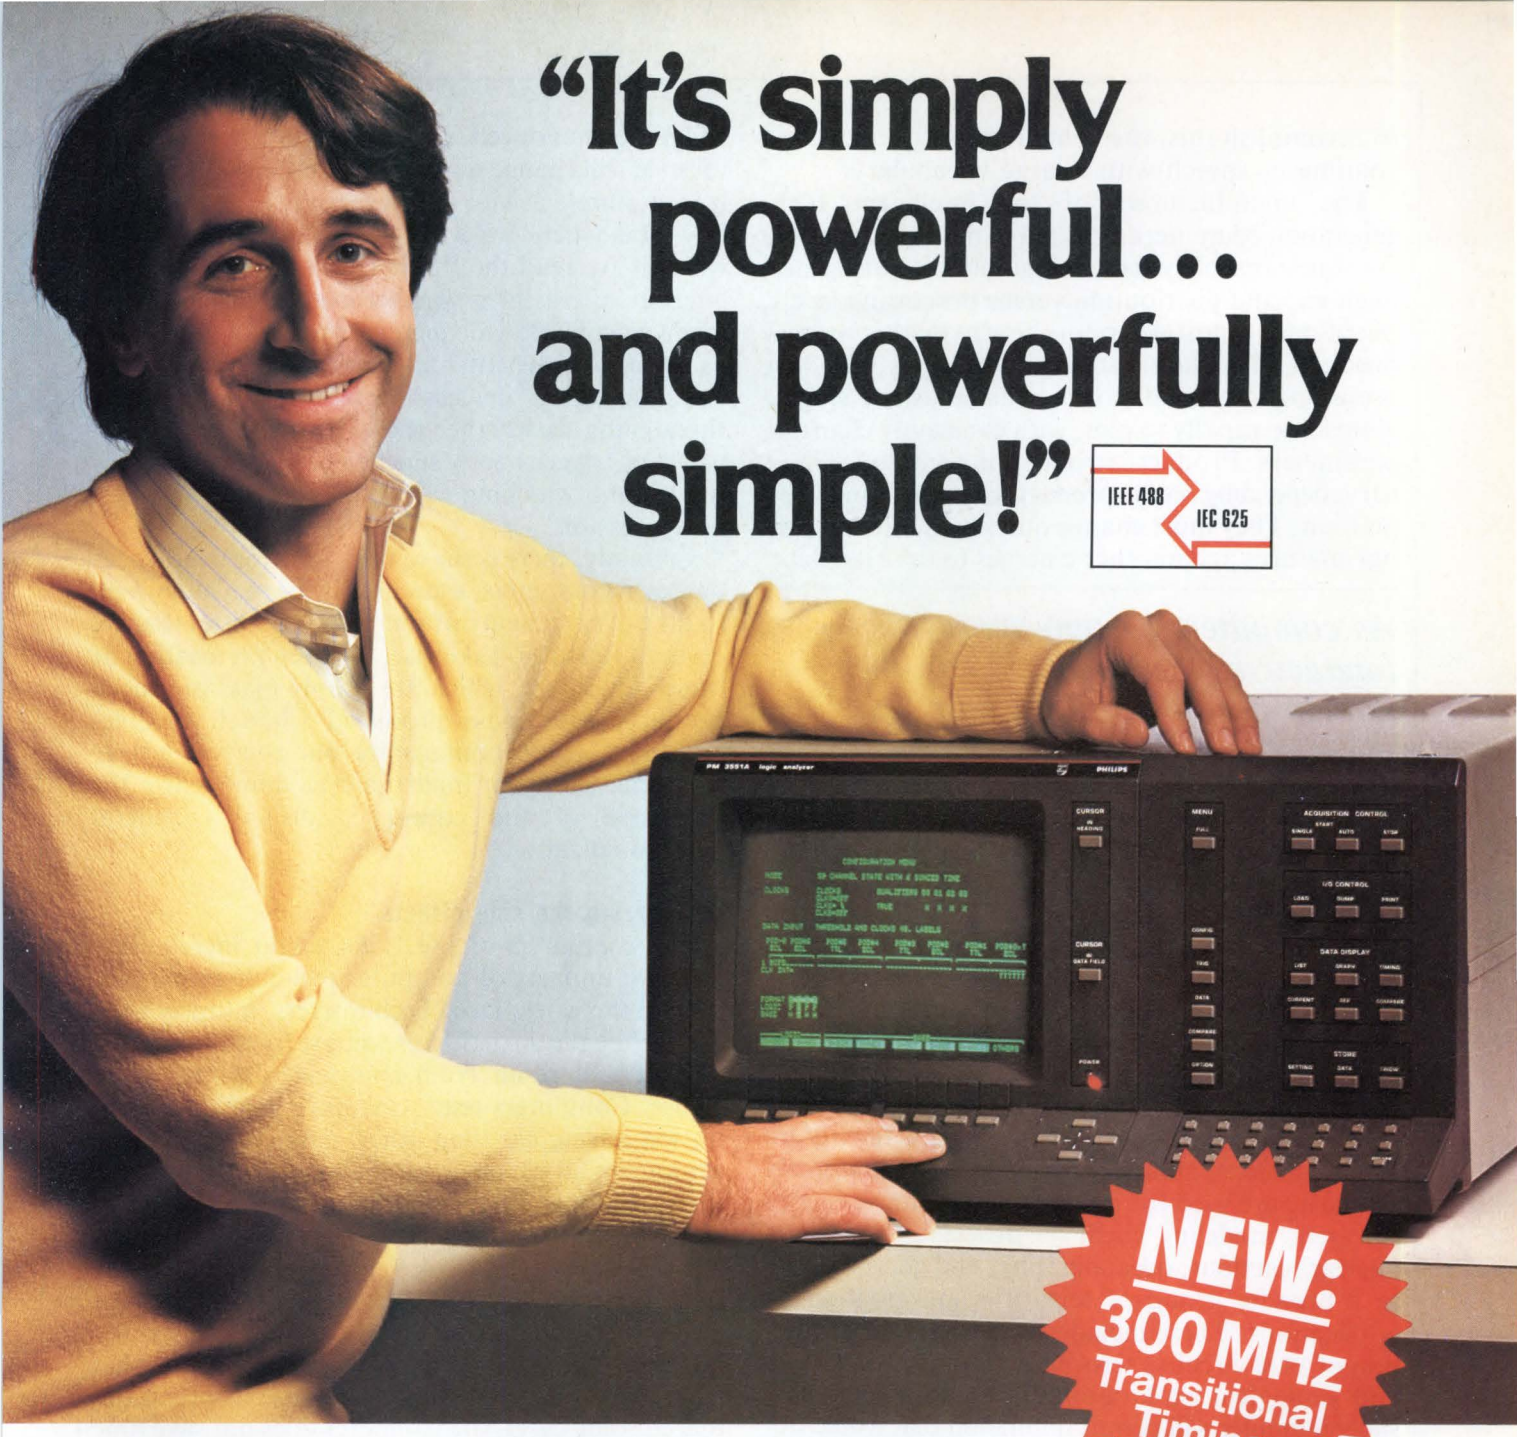

Philips' PM 3551A... in Logic Analyzers it's the logical choice

"Sophisticated yet simple. That's the PM 3551A. A dedicated state analyzer plus a dedicated timing analyzer working simultaneously in one instrument.

Its powerful SYNC mode gives you state and timing analysis at the same time – a real advantage if you're integrating hardware and software design.

Transitional timing saves memory without losing resolution.

And simplicity! Clear menus and soft-key entry let you use all its power from day one. You've disassembly of any 8 or 16 bit micro-processor at the press of a button – without adding extra boxes.

And that's just the beginning. To learn all the reasons why it's the logical choice.

For full details call: 800-631-7172 except Hawaii, Alaska and New Jersey. In New Jersey call collect (201) 529-3800 or write to 85 McKee Drive, Mahwah, NJ 07430."

PHILIPS **Test & Measuring Instruments**

CIRCLE 41

PHILIPS

to accomplish this, the computer must understand continuous speech with a large vocabulary.

The manufacturers of the intelligent fifth-generation computers expected in the 1990s not only face questions of system configurations, intelligence, memory, and distribution versus processing levels, but also they must determine how to market the new products. To exacerbate their decisions, they are projecting into an unknown, in an industry changing almost too rapidly to plot, with constantly changing parameters. Product introduction techniques must vary, depending on the product, customer, and competition. They must change quickly as one technology matures and another emerges to take its place.

As computers become more interactive, their demands on their human partners will escalate.

Interestingly enough, selling these new complex computers to developing coutries a decade from now could be much simpler than marketing the comparatively elementary computers today to the industrial nations in North America, the Far East, and Europe. With better interfaces, such as voice capabilities, and logic-based software, the computers of the 1990s should be so much easier to use that customer training will be minimal and unintimidating.

As today's supercomputers grow from 100 million operations per second to 1000 times that speed, their use is expected to expand from a few, quite specific massive number-crunching jobs the systems primarily are used for today—meteorology, energy research, modeling chemical processes, exploring nuclear reactions, oceanography, and geophysics—to a broader group of scientific and engineering applications. These include the design and simulation of VLSI circuits, aerospace, automotive design, and geological exploration. In some instances, a supercomputer will permit simulation that would not be possible with the actual equipment (such as weapons).

Growth of the supercomputer market, which is limited to perhaps 100 state-of-the-art installations now, predominantly in national laboratories, is a necessity for the success of the industry. This growth depends largely on the development of software that makes supercomputers easier to use and the implementation of more user-friendly interfaces.

Society adapts

The changes to computer architecture, software, and hardware discussed on the pages that follow will revolutionize information technology as dramatically as the industrial revolution changed the shape of manufacturing in the world. The effect on society will not go unnoticed. For most people in the world, lifestyle will change during the next decade, whether it is as simple as the way we bank, as involved as how and where we work; as relatively benign as whether we read the "paper" on newsprint or on a screen, as critical as whether medical diagnoses are made via a worldwide medical symptom data base, as harmless as whether all our telephones and other communication devices become cordless, or as threatening as whether the earth becomes a tense world of trigger-happy superpowers—poised, listening to, and watching each other, in wait of a sign of aggression.

Certainly, there is the opportunity for overdependence on computers, especially if the machines were to take over so many of our responsibilities that we cease to function. Likewise, the deep permeation of these computers, which are simply mass produced products with mass produced (though flexible) programs, could homogenize our society; just as suburban shopping malls tend to homogenize the landscape of the country, causing a general depersonalization.

Considering the side effects

The computers of the 1990s might increase the negative impacts of computer technology particularly in the workplace. Technostress, a timely term to describe the anxieties, pressures, frustrations, and tension that result from computer use, could increase. Computers tend to require greater concentration during use than other, passive, mechanical equipment. And, as they become more interactive, their demands on their human work partner will escalate.

Craig Brod, in his book *Technostress: The Human Cost of the Computer Revolution*, discusses the effects of computers on the workplace as well as the increased mental work load and disappointment that have characterized some installations. He points out that in some cases, the computer does not save time or make a job easier, particularly when managers see them as sources of endless piles of data. With smart computers, capable of decision making, those same managers could tend to make problems—or find problems—so the computer can solve them, not because they need to be solved. In these situations, the computers makes the workplace more complex. Whether it is a proper use of the equipment or not, it still results in worker resentment.

Certainly, the computers in use in the 1990s, particularly the computer systems based on fifth-generation architecture, will bring people in the world closer together. The instantaneous communication we will gain could change the way we vote, determine weather forecasts and medical alerts, and keep in contact with family.

The shock of technological change is lessened slightly by the knowledge that though the change may occur quickly, they are never abrupt and unforewarned. By the time the most interactive expert systems are available for public use, less capable systems will have been implemented. This will ease the technology onto the market.

Robot vision, for instance, is not science fiction to most people now; and voice synthesis, though crude, is being used today. Computers will become simpler to use, less obtrusive, and less foreign. Standardization in software and hardware will ease some of the confusion felt today.

One identified problem with the expansive computer networks that will accompany the fifth-generation machines to the market is privacy. Because computer-based information security has not been established as fast as new data bases, it would not be difficult to access a number of sources to produce a fairly comprehensive picture of an individual's assets, legal problems, scholastic background, etc—completely without the person's knowledge or permission. With accurate sensors and a data base of speech patterns, it would be possible to tap phone lines continuously to "listen" by computer for a certain voice.

The privacy infringement problem is a well recognized problem, but it will probably require legislative attention for the restraints that will ensure protection. This should occur while the technology is coming up to speed, not after it has had the opportunity to be misused.

Sociological impact should not be overlooked by governments, industry, universities, or the collaboratives they are forming to push forward computer technologies. The effects on society and on lifestyle will reach too far into the future, touch too many lives, and have too much of an effect on history to avoid scrutiny and debate.

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

High 701 *Average 702* *Low 703*

THE BEST PLACE TO BOARD THE BUS.

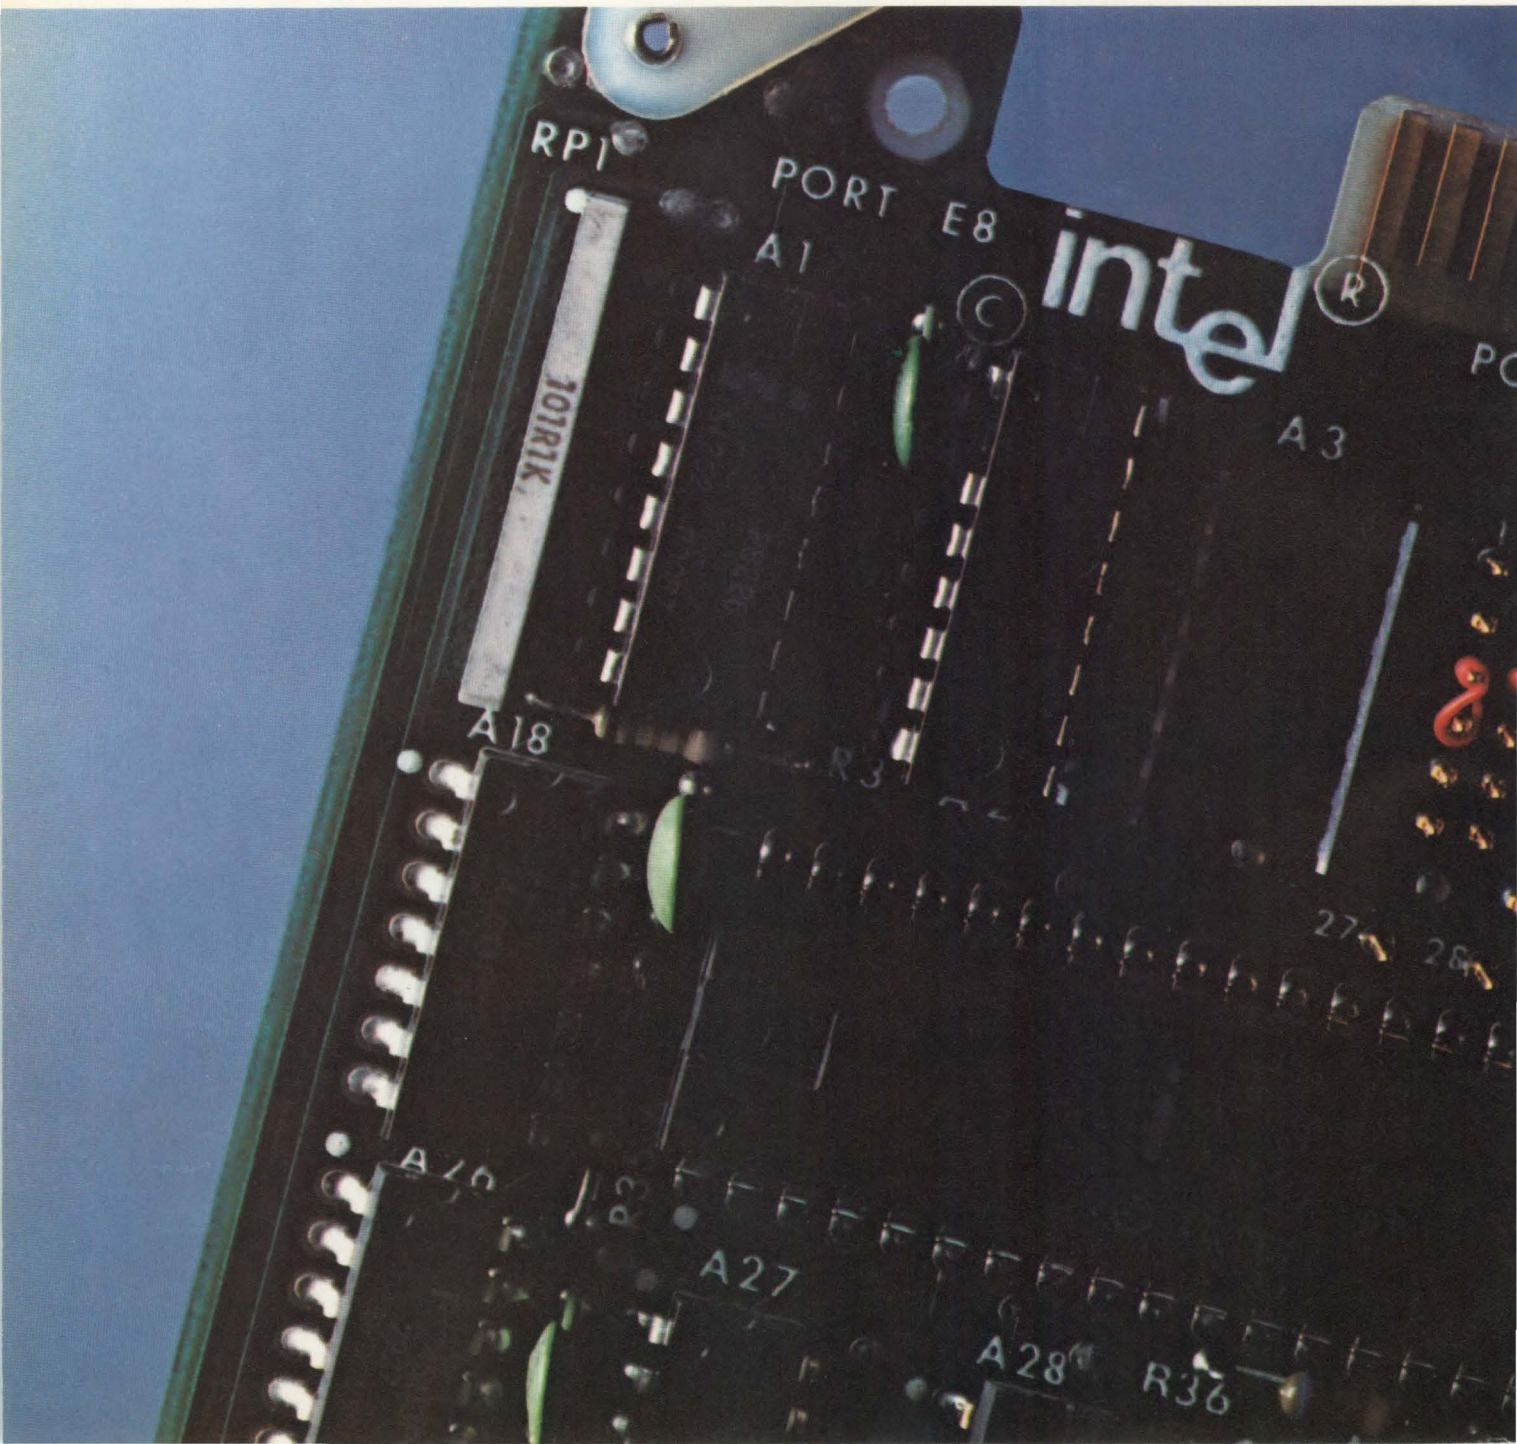

Of course, we're talking about the MULTIBUS® architecture. The most widely accepted, best supported bus architecture in the world.

But since you've already made the wise decision to use MULTIBUS as your road to riches, we're now going to show you the best place to start your journey.

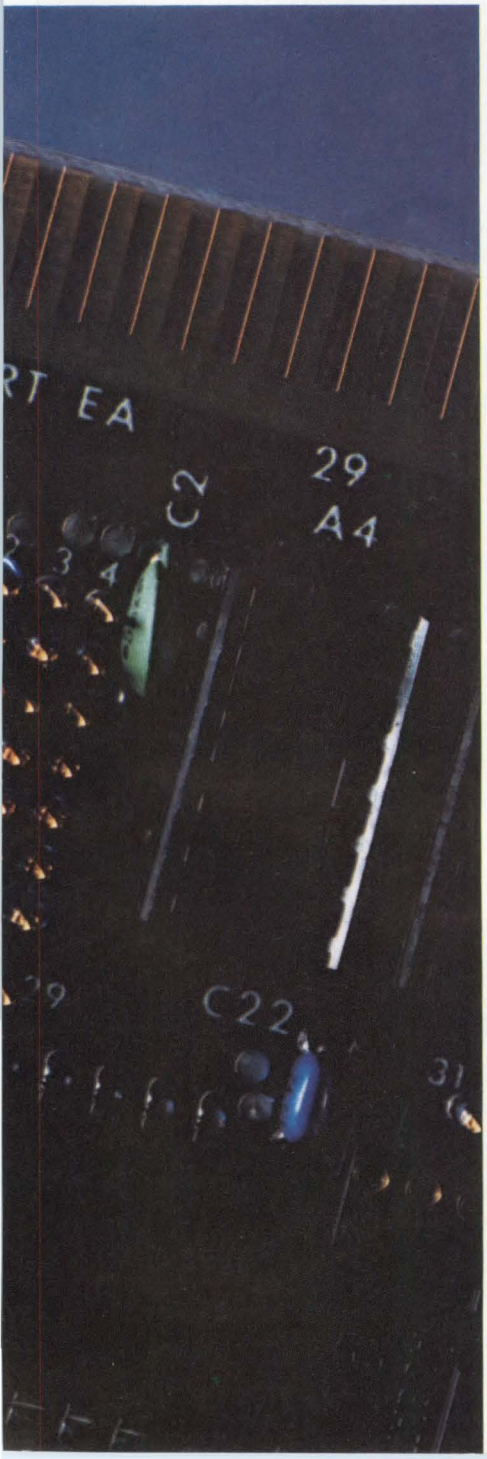

Right here. Because as you might suspect, we know just about all there is to know about MULTIBUS. After all, we invented it. And we have the largest selection of MULTIBUS products to be found. Anywhere.

For starters, there are our boards. We've produced over a million since 1976. Which should give you an indication of the experience we have in manufacturing quality boards. Quality that's not easily achieved.

Our boards are scrutinized at every level of integration. From chip through system level. And each board is tested to make sure it does its job in every one of its configurations. We even pull finished product from our warehouse for retesting and inspection. Just to make certain nothing slips by.

You can count on the road being smooth for a long time to come, too. Reliability is assured from start to finish. We continually monitor each board throughout its entire production life. Even going so far as lot sampling older boards and putting them through 5000-hour life test studies.

Then there's the selection. A big one.

Over 115 different board-level products. Supporting 15 CPUs (from our 8080 to 80286). So you won't get stuck into one or even two design approaches. A tranquil thought when you consider how fast things change in the marketplace.

What's more, you can count on getting everything you need. From peripheral controllers to memory expansion, data communications and graphic boards. Plus card cages and software (including iRMX 86™, the most popular real-time operating system in the MULTIBUS world).

One more fact about our MULTIBUS products. You can get fast service anywhere along the road. We've got one of the largest field engineering forces in the world. Over 600 factory-direct people. Not to mention localized support from our worldwide network of trained distributors.

So don't miss the bus. Travel safely. And arrive sooner. Call toll-free (800) 538-1876. In California, (800) 672-1833. Or write Intel, Lit. Dept. C-20, 3065 Bowers Ave., Santa Clara, CA 95051.

We'll be happy to welcome you aboard.

intel®

You'll never get trapped into dead-end designs with

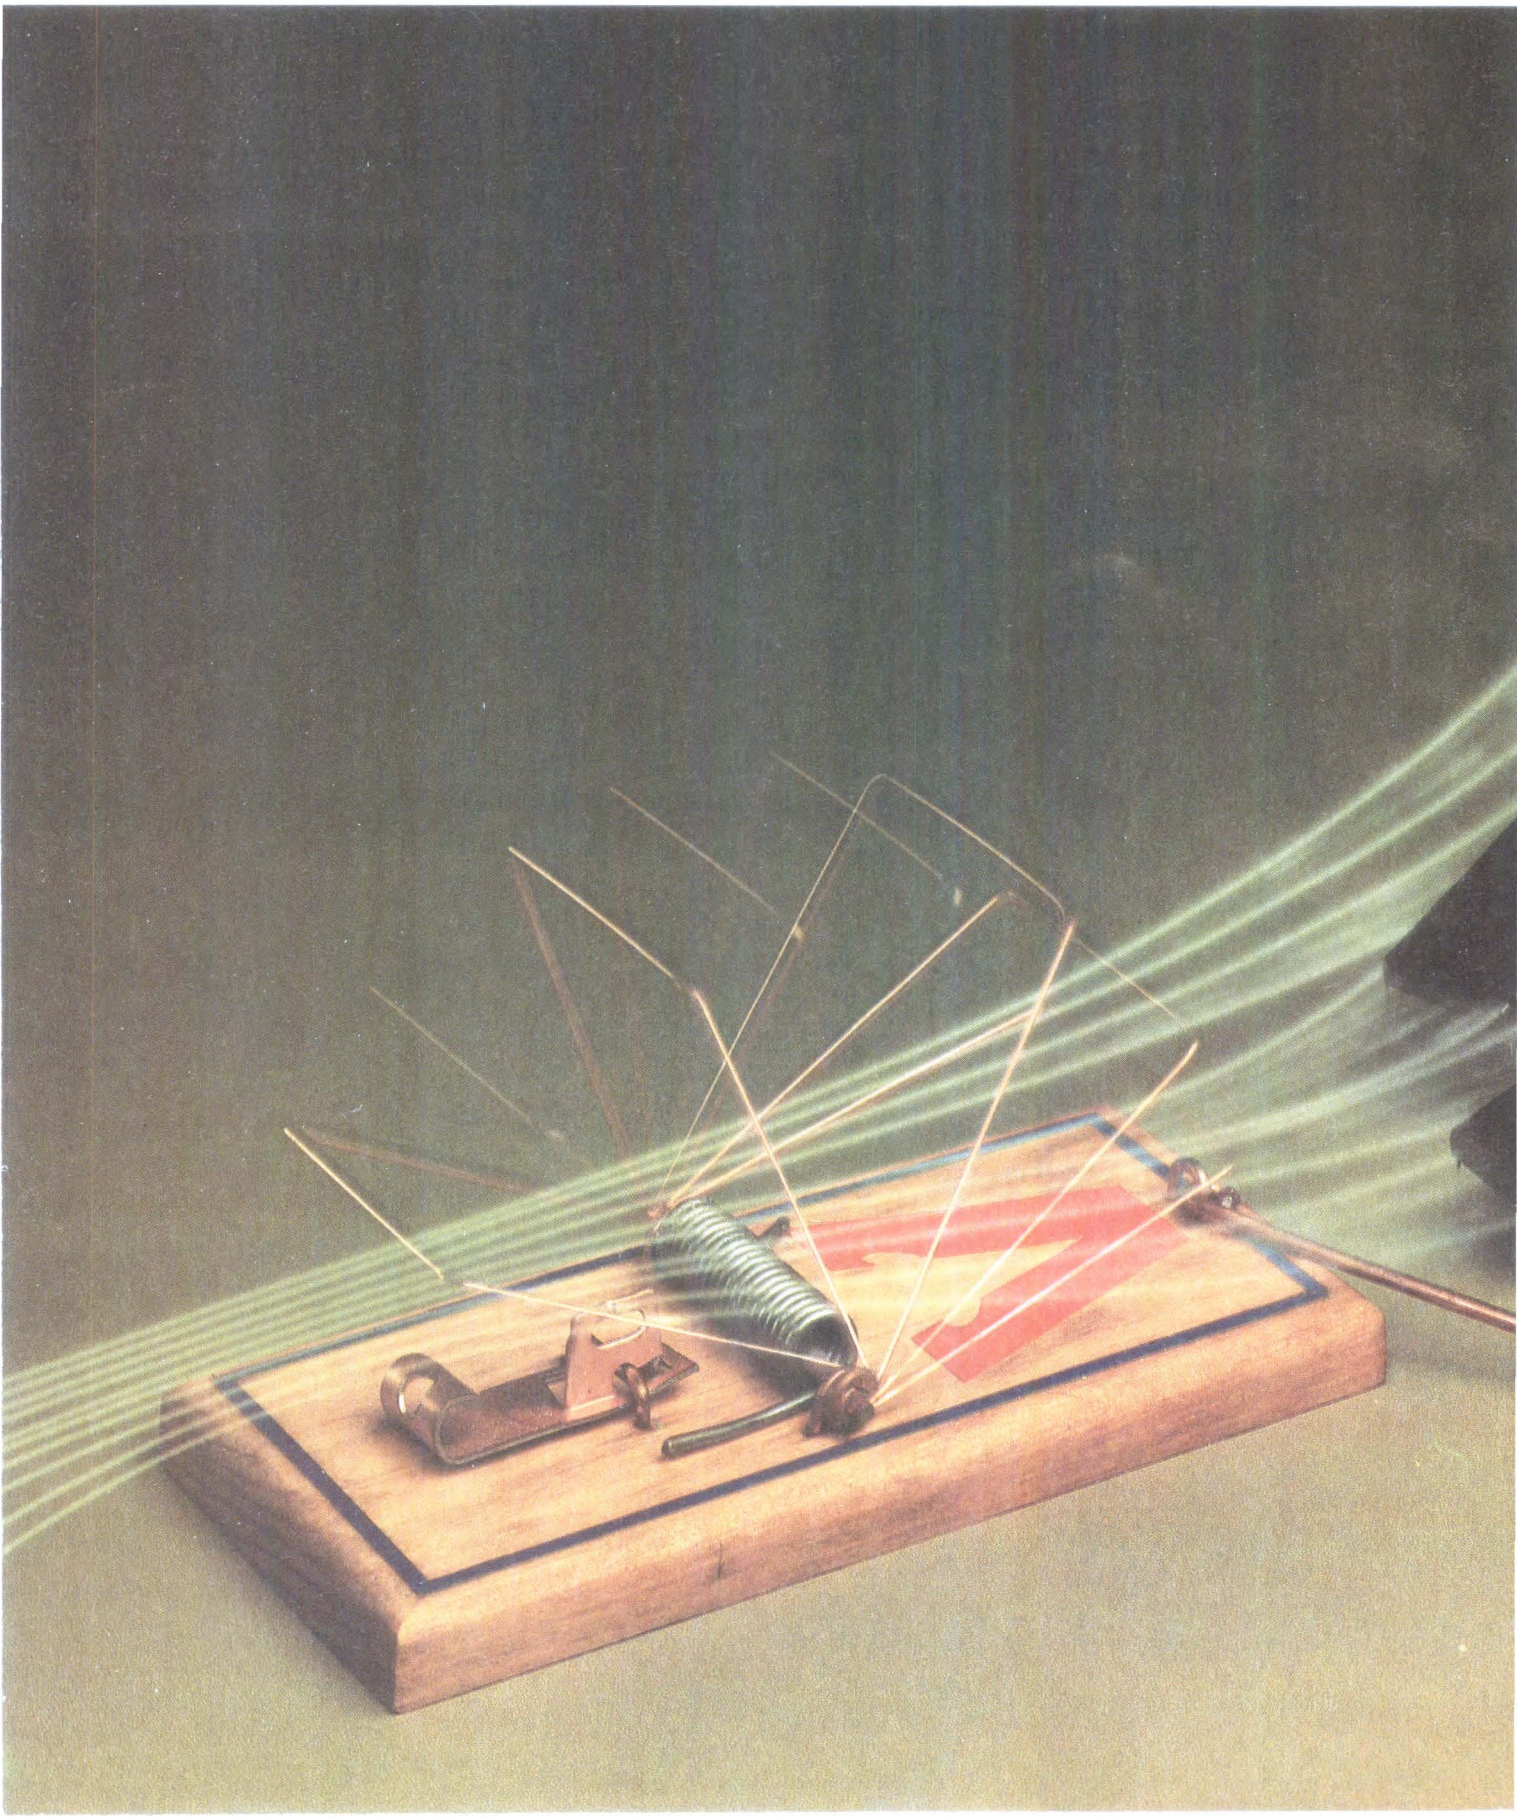

Zilog's high-speed Z80B CPU's and Peripherals.

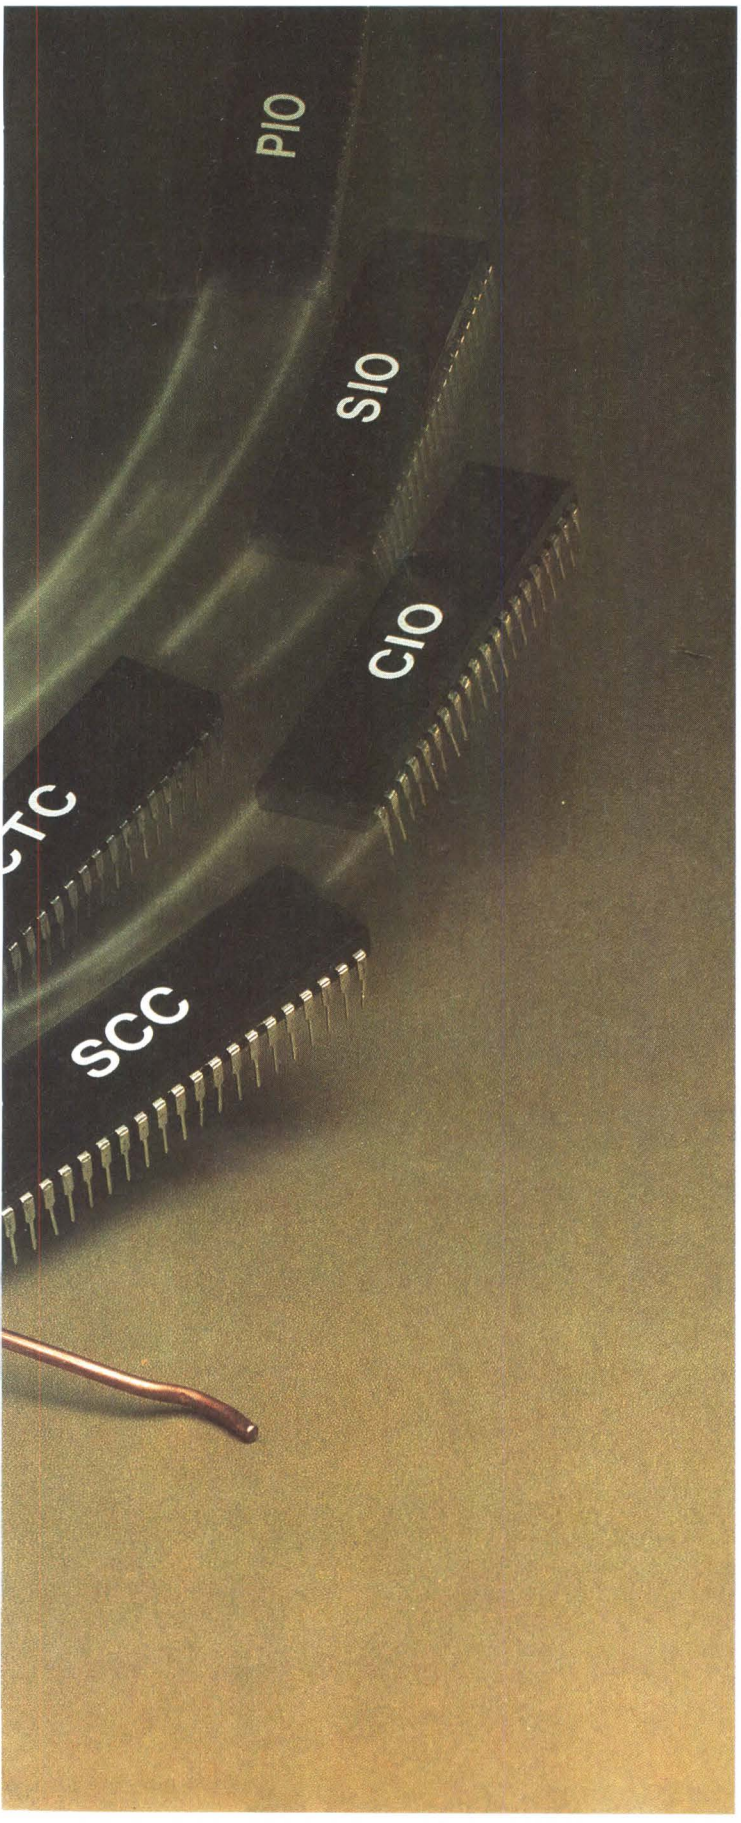

Here's more proof that nobody does more to extend the life of your 8-bit designs than Zilog. Because now you can increase 8-bit Z80® performance up to 6 MHz with the high-speed Z80B CPU and its family of peripherals. You can join the hundreds of design engineers that have already tested this claim with winning results. Or wonder...

Is there something here they know that you don't?

Like the fact that the Z80B CPU has the same 158 instruction set and the elegant registers and interrupts that you're used to working with, but runs them 50 percent faster than the Z80A chip?

That the Z80B processor is completely software compatible with the rest of the Z80 family, permitting you to upgrade to higher performance without getting trapped into software redevelopment? That software compatibility also means you can use the Z80B device in co-processing and/or multi-processing environments along side our other Z80 processors?

And consider the fact that you can surround our Z80B CPU's with a complete family of Z8400 and Z8500 peripherals and really boost system performance. They help you keep your parts and space requirements to a minimum and increase system throughput because we build more functions into every device. The peripherals include a PIO, a CTC, an SIO, an SCC, an FIO, an FIFO, a CIO, and a UPC.

For complete specifications and applications data on the Z80B and peripherals, fill out the coupon and mail to: Zilog, Inc., Components Tech. Publications, 1315 Dell Avenue, MS C2-6, Campbell, CA 95008. Or call our TOLL FREE Literature Hot Line at 800-272-6560. For information on Zilog's other components, call (408) 370-8000.

Z80 is a registered trademark of Zilog, Inc.

☐ I'd like more information at this time.
☐ Please have a salesman contact me.

Name_____

Title_____

Company_____

Address_____

City_____ State_____ Zip_____

Phone_____ /_____

Z80B/Peripherals CD9/84

Zilog
an affiliate of
EXXON Corporation

Pioneering the Microworld

CIRCLE 44

YOU WON'T FIND MANY SHUGART DRIVES AROUND THESE PARTS.

You find Shugart new generation drives where they were built to be. In all kinds of systems. Under all kinds of conditions. Running. And running and running.

Setting a world standard for quality. A standard that doesn't allow room for failure. Because you don't.

To insure the reliability of your system, we begin with our suppliers. They go through a lot. Learning and qualifying, even building drives themselves.

We also involve our manufacturing engineers. And our quality control experts. From the very beginning. The design stage.

With their full collaboration, every new drive is designed for top quality. Unbeatable reliability. And ease of assembly. Using the fewest number of parts possible.

In other words, the same uncompromising quality you design into your system is designed and built into every new Shugart drive. Right from the start. So you can count on consistent reliability. For example, 20,000 hours MTBF for new generation Winchesters.

Of course, every Shugart drive is subject to the most stringent testing and inspection at every stage. But inspection doesn't insure quality. At Shugart, we believe there's only one way to insure a reliable product. By eliminating the cause of failure from the product design.

At Shugart, quality is designed in. Not inspected in.

Let us prove it. Write for our white paper on quality to Shugart Corporation, 475 Oakmead Parkway, M/S 3-5, Sunnyvale, CA 94086. Or call, (408) 737-4360.

Shugart

Right from the start.

CIRCLE 45

Milpitas, CA (408) 263-2600; Costa Mesa, CA (714) 979-1935; Thousand Oaks, CA (805) 496-5388; Rochester, NY (716) 235-7190; Minneapolis, MN (612) 546-4411; Richardson, TX (214) 234-3568; Framingham, MA (617) 879-1700; Saddle Brook, NJ (201) 368-8445; Smyrna, GA (404) 436-0953; Markham, ONT (416) 475-2655; Paris, France (3) 946-42-66; Munich, West Germany (089) 786-021; London, U.K. (44) 4862-27272; Wanchai, Hong Kong (852) 5-733307.
Hamilton/Avnet, authorized distributor

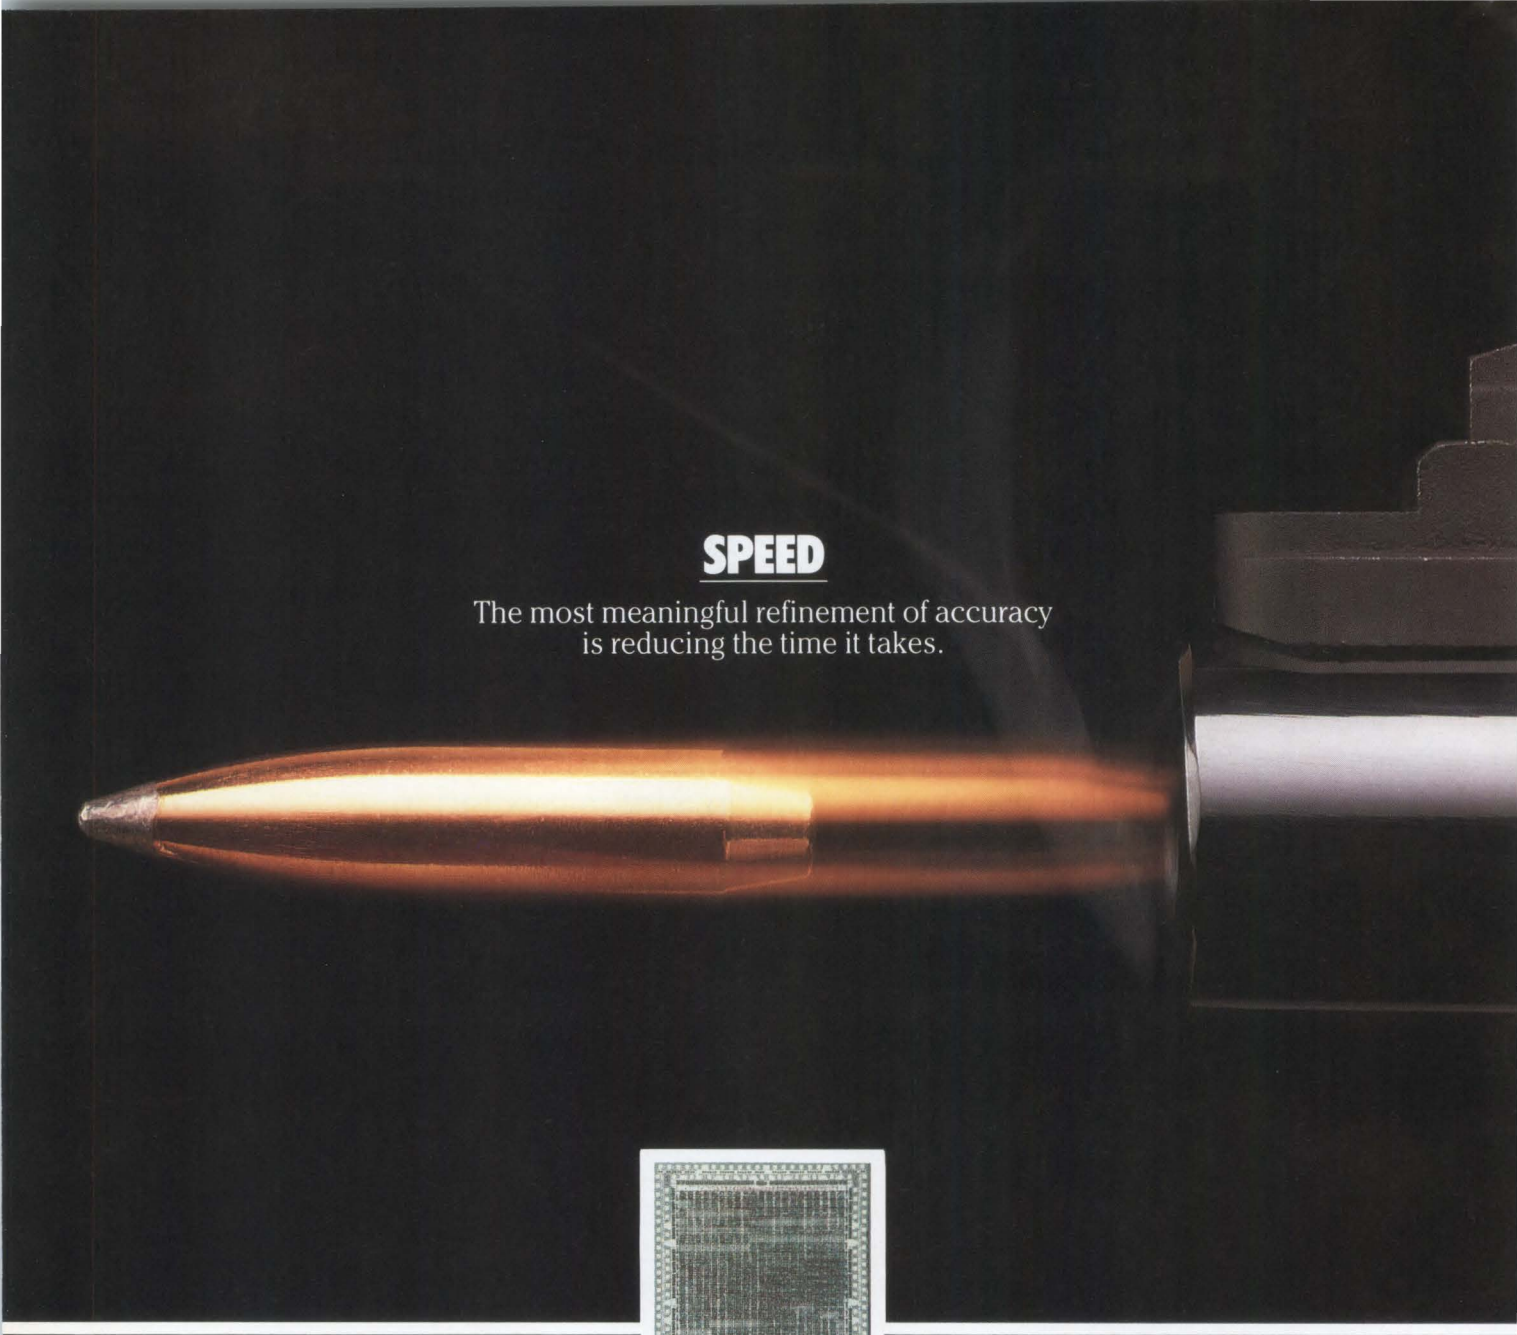

SPEED

The most meaningful refinement of accuracy is reducing the time it takes.

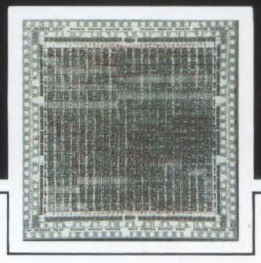

The more tasks a microchip performs, the more critical its speed becomes.

And just as taking a faster look at a brief event can let us see more, microchips that operate more quickly accomplish more work in less time.

Take the Honeywell HT5000 gate array. With a 0.6nS gate delay speed, it operates approximately *five times faster* than typical ALSTTL chips.

Or the Honeywell HE2000. It operates even faster, with a 0.3nS delay, but requires only half as much power as other equivalent ECL arrays.

And our radiation-hardened H1000R. It has a 1.0nS delay, and is designed for space and strategic TTL replacement applications, using a single +5v. power supply.

Since 1954 Honeywell has been building the technical know-how necessary to consistently design and manufacture this kind of performance in every custom and semi-custom chip we make.

For more information on the industry's fastest radiation-hardened IC's, VHSIC IC's, and semi-custom gate arrays (including special macrocells for digital signal processing), talk to the company that has been in this industry since before it was an industry.

The Honeywell Digital Product Center, 1150 East Cheyenne Mountain Blvd., Colorado Springs, Colorado 80906. 1-800-328-5111, extension #3402. (In Minnesota, call 1-612-870-2142, extension #3402.)

Together, we can find the answers.

Honeywell

CIRCLE 46

INTRODUCTION:
FUTURE ARCHITECTURE

Divide and process—that seems to be the motto adopted by computer architects for the next generation of machines. As opposed to the omnipresent von Neumann architecture that processes data in a serial fashion, future NON-VON distributed parallel architectures will be based on the premise that dividing a problem into smaller parts will yield faster results. The underlying concept behind all new parallel architectures is to have small blocks of tightly coupled processor memory elements divide the algorithm and parallel process its parts concurrently. As a result, larger problems can be solved faster.

Research in parallelism is not new. What is new is the ability to incorporate many of the ideas into working machines, thanks to the availability of relatively inexpensive VLSI chips and high level languages. Thus, while the idea of using parallel processing can be traced to Babbage's work on the analytical engine more than 150 years ago, it has only recently become practical to incorporate these concepts in hardware.

Many alternate parallel configurations have been proposed and research machines have been built, including a host of multiprocessors, multifunction processors, array processors, and pipeline processors. Most projects, however, were either for military contracts or academic adventures that for the most part turned out to be "paper tigers." A few exceptions included supercomputers from Control Data and Cray Research for use in sophisticated engineering applications where the customer did not mind the machine's price tag. But today, cheap VLSI and memory, and semi-efficient parallel algorithms allow many universities to develop working machines that show promise as future commercial products.

In addition, Japan's Fifth-Generation Computer Project has drawn attention in the computer science and engineering communities, and is forcing a quick response from the United States and Europe in order to stay competitive in computer technology. Thus, in addition to research in supercomputer parallel architectures to be used for scientific applications, separate defense/industrial/academic consortiums are studying artificial intelligence. It is hoped that computerized "experts" in all areas involving human-decision tasks will be developed. Toward this, research in symbolic processing is trying to integrate AI principals with parallel processing configurations. The goal, of course, is to reach a stage where a large data base on any knowledge-intensive subject can be manipulated instantaneously.

Much research remains before that goal is reached. And, as with any finite human-generated concept, there is a limit to parallel processing efficiency (see "Parallel Processing Makes Tough Demands" p 137). If the purported bottleneck of the von Neumann architecture can be traced to its architectural floorplan, then the bottleneck in parallel configurations may be the parallel algorithm itself. Even truly efficient algorithms, successfully developed and implemented to use unlimited hardware, will test the limits of parallel processing.

Design engineers should study Gene Amdahl's theory on parallel processing. Amdahl, a computer pioneer, theorized that the limiting factor to increasing the speed of a parallel algorithm is inversely proportional to the computation that the parallel algorithm must perform sequentially. Hence, if five percent of the parallel algorithm uses one processor of a multiprocessor system, then the parallel computation can at most be 20 times faster than the original sequential algorithm. Thus, the need for clever algorithms for the next generation of computers cannot be overemphasized if truly efficient concurrent processing is to become the norm in tomorrow's environment.

Nicolas Mokhoff
Senior Editor

PARALLELISM MAKES STRONG BID FOR NEXT GENERATION COMPUTERS

Parallel architectures that allow concurrent processing of multiple tasks promise orders of magnitude improvement over current topologies.

by Nicolas Mokhoff,
Senior Editor

Computer design will either evolve or explode into the 1990s—the advances made today will dictate the course. Radically new architectural concepts will have to match the increase in magnitude of computing performance required to solve both the complex engineering and the everyday problems of the next decade. Many new architectures are already vying for designers' attention and, unlike in the past, today's designers are able to test these architectures in real systems.

Computer design is entering the age of intelligent computing. This is due to the simultaneous maturation of three elements of the computing infrastructure: inexpensive, large scale ICs; programmable artificial intelligence; and centrally distributed architectures. We are therefore likely to see, for example, the fifth-generation computer. This knowledge-based machine, which will "compute" 1000 times faster than today's machines, is a recent media focal point. The supercomputer effort is getting an equal amount of press. Supercomputers are special application machines to be used for weather simulation, airplane modeling, and bomb trajectory and earthquake prediction. In addition, the Japanese foresee a future society that depends on their intelligent computers for advice and service.

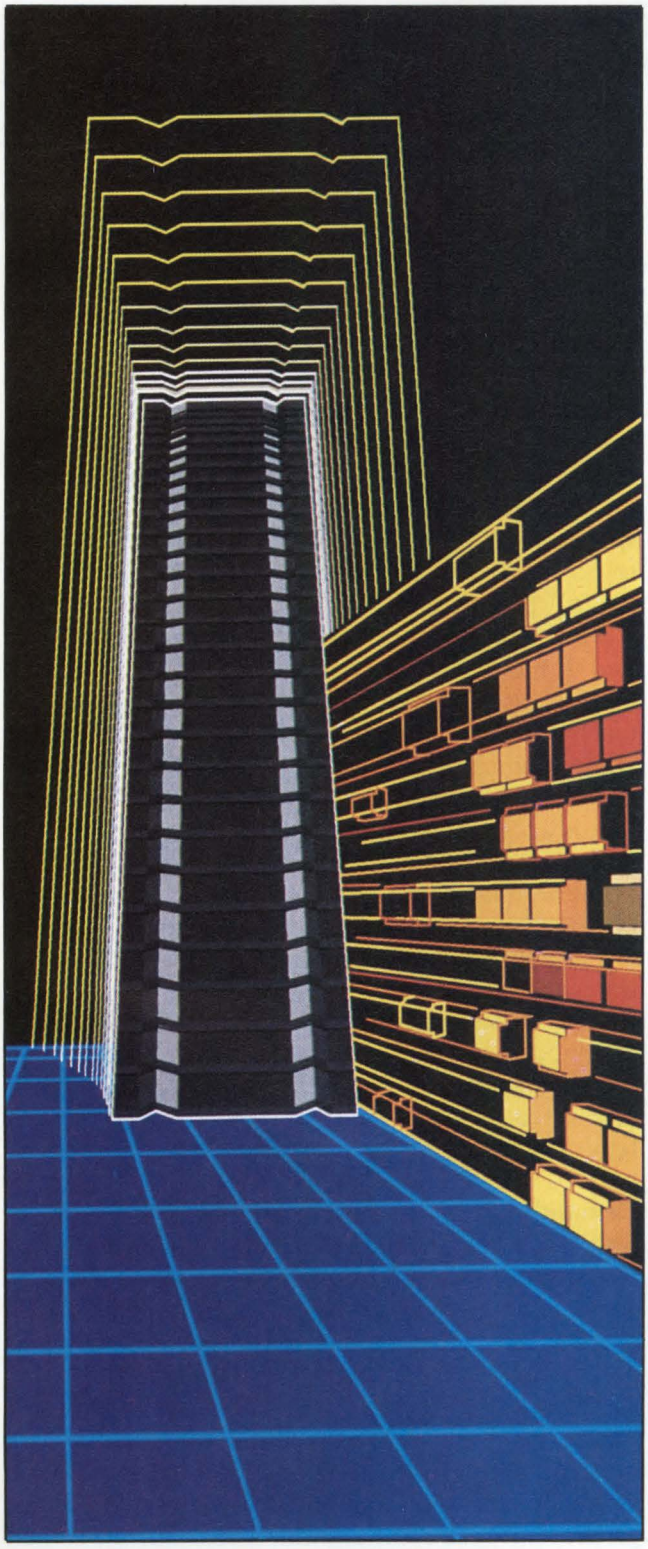

Each of these schemes calls for a unique architecture. A generic architecture is not appropriate for all classes of problems. Yet, the one aspect that all future architectures share is that a multiple number of processors, perhaps as many as 1000, will be used in some kind of distributed fashion. The dispute among computer scientists and computer engineers today revolves around the appropriate type of distribution and the contents of each process.

To control, to data, to demand

Currently, there are three candidates for future architectures—those that are either control driven, data driven, or demand driven. It is unclear which of these architectures will predominate, since many variations and combinations of the three will probably be the norm. What is more, an eloquent architecture for a fifth-generation machine may not be suitable for a supercomputer. This is because the fifth-generation machines are mainly slated for knowledge processing tasks designed to support knowledge-based expert systems. On the other hand, supercomputers are largely used for solving large scale numerical calculations in engineering and scientific applications. While the former is a machine for "reasoning," the latter can be considered a very fast calculator.

However, the knowledge-intensive/supercomputer delineation is not very clear. Obviously, there is more than one way to slice the future computer architecture pie. For instance, Philip C. Treleaven, professor of computer science at the University of Reading, England, expands his vision a bit further than just fifth-generation machines and supercomputers. At the 10th International Symposium on Computer Architecture last year, Treleaven identified two more potential architectural candidates for the future computer—VLSI processors and integrated nets of computers and communications.

The VLSI processor architectures will exploit VLSI to the extreme by defining a new generation of components that will succeed the conventional microprocessor. The method proposed follows the design philosophy espoused by Carver Mead and Lynn Conway in their *Introduction to VLSI Systems* (Addison-Wesley, 1980). This philosophy allows miniature microcomputers to be replicated like memory cells and connected for multiprocessor

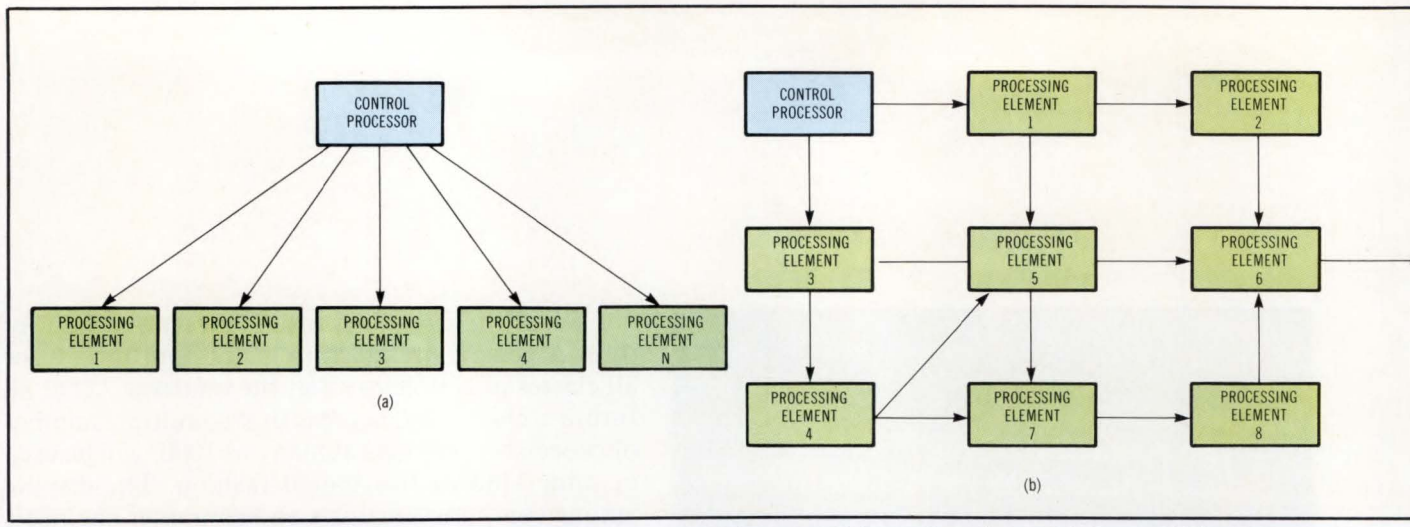

Three forms of parallelism dominate current computer science thought: control flow, data flow, and demand driven. In a control flow machine (a) a control processor sends instructions to many processing elements, which consist of processor and associative memories. Data flow machines (b) are highly decentralized with parallel instructions sent along with data to many equivalent processing elements. Demand-driven machines (c) break tasks into less complicated subtasks, which, when processed, are recombined for the final result.

operation. The fourth candidate for the computer of the 1990s is a concept incorporating integrated data communications and computers, Treleaven says. The network entails a fusion of wide area and local area computer networks, and parallel computer architectures. Achieving this integration calls for all component computers to conform to a common decentralized system architecture that would allow them to communicate over various distances. In short, Treleaven sees four complementary views of future computing that yield a composite image of a highly decentralized computer system at all levels. There, computers will be linked together in an integrated computer-communication network.

But there are many questions. For instance, what type of architecture lends itself to a general-purpose computer system whose central flow mechanism specifies *how* an algorithm should be executed? This scheme gives maximum mechanism flexibility and control over execution. On the other hand, should an architecture be composed of logic elements that specify *what* algorithm is to be executed? This leads to powerful abstraction mechanisms and relatively "safe" programs.

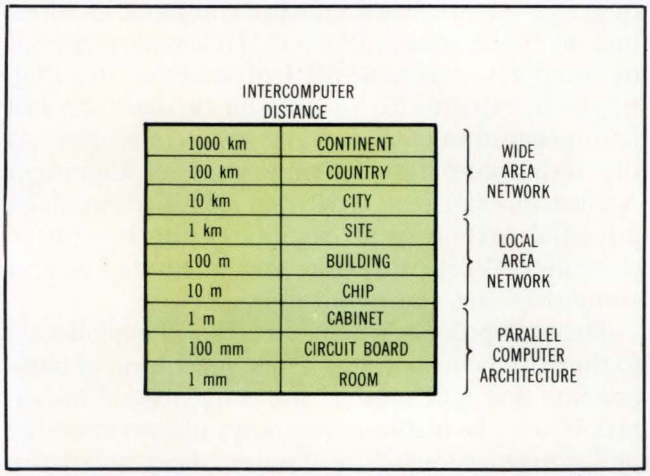

Full system integration in future computer systems will only materialize when each element adhers to a decentralized architecture. This will allow the separate elements to be programmed for efficient data communication and program execution.

To probe the matter further, Treleaven points to two fundamental mechanisms that all computer architectures share: data and control. The data mechanism defines the way a particular argument is communicated and shared by a number of instructions. This mechanism includes a shared memory that is accessible to all instructions. It also features message passing—a process by which a unique copy of the argument is communicated via a message from the source to the destination instruction.

The control mechanism defines how one instruction causes the execution of other instructions. The first instruction can be either control driven, data driven, demand driven, or pattern driven. In control-driven systems, an instruction is executed when it is selected by the control flow. In data-driven systems, the instruction is executed when some combination of its arguments is made available. Demand-driven systems stipulate an instruction to be executed when the result it produces is needed by another already active instruction. Finally, pattern-driven systems have their instructions executed when some enabling pattern or condition is matched.

After weighing the arguments for all options, Treleaven concludes that the most important data mechanism is shared memory, and the most basic control mechanism is control-driven execution, for general-purpose computation. The challenge is then

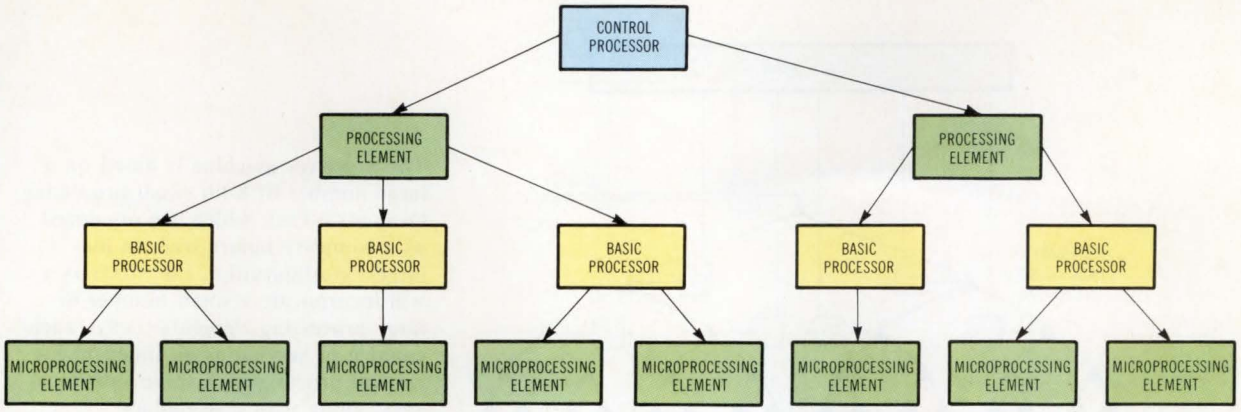

to introduce into control flow the other data and control mechanisms to make control flow even more generic.

A control-driven system follows the typical von Neumann computing style. A sequential machine performs one operation at a time using a single processing element in conjunction with a centralized control unit; a low level, sequential machine language; and a linearly addressed, fixed-width memory. Other machines, performing operations in parallel, can process instructions much faster. But, they are hampered by hard-to-develop parallel algorithms that specify how to break up the instructions for concurrent operation.

Thus, researchers have developed machines characterized by single-instruction, multiple-data (SIMD) path architectures. While these machines work well where data is structured in dense arrays, they are poor general-purpose machines. This is because they are typically used as attached processors to a von Neumann host machine.

Multiple instructions to boot

Far more powerful multiple-instruction, multiple-data (MIMD) path architectures are on the way. In the MIMD architecture, a computing problem is partitioned so that the processor can compute it by operating concurrently on different data streams with different instructions. However, the path toward designing an MIMD is tortuous.

The NON-VON approach, promoted by David Shaw's group of computer scientists at Columbia University in New York, has laid the groundwork for developing a family of NON-VON machines. (These machines are supposed to be the antithesis of the von Neumann machine.) The most advanced member of the NON-VON family is the NON-VON 4, which will be able to function as an ensemble of one or more independent SIMD machines by having them communicate through a high bandwidth interconnection network. Thus, NON-VON 4 would essentially operate as an MIMD machine.

The NON-VON family is characterized by an extensive intermingling of processing and memory resources at various levels. According to project leader Shaw, the machines will be able to execute large scale data manipulation tasks, including relational database operations, and other functions relevant to commercial data processing. All members of the family incorporate a primary processing subsystem, which comprises a large number (as many as a million) of simple, highly area-efficient small processing elements (SPEs), each associated with a 64-byte local RAM. The primary processing subsystem is currently being implemented using custom NMOS VLSI circuits, each of which contains eight processing elements. While the fine granularity size

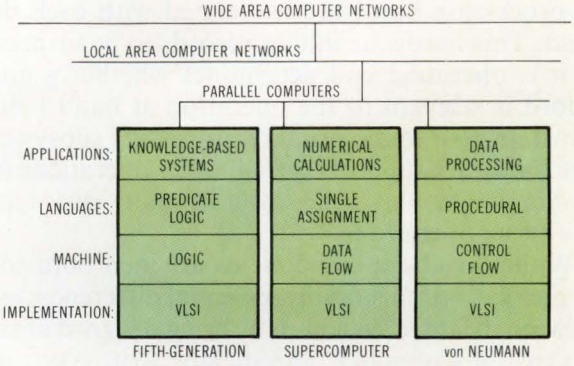

Future computers can be divided into three classes: von Neumann processors, supercomputers, and symbolic processors. Each class, when implemented in VLSI, will feature architectures and languages that are unique to their intended applications (ie, data processing, numerical calculations, and knowledge expertise). Ideally, all three classes will execute instructions in parallel and will communicate over both local and wide area networks.

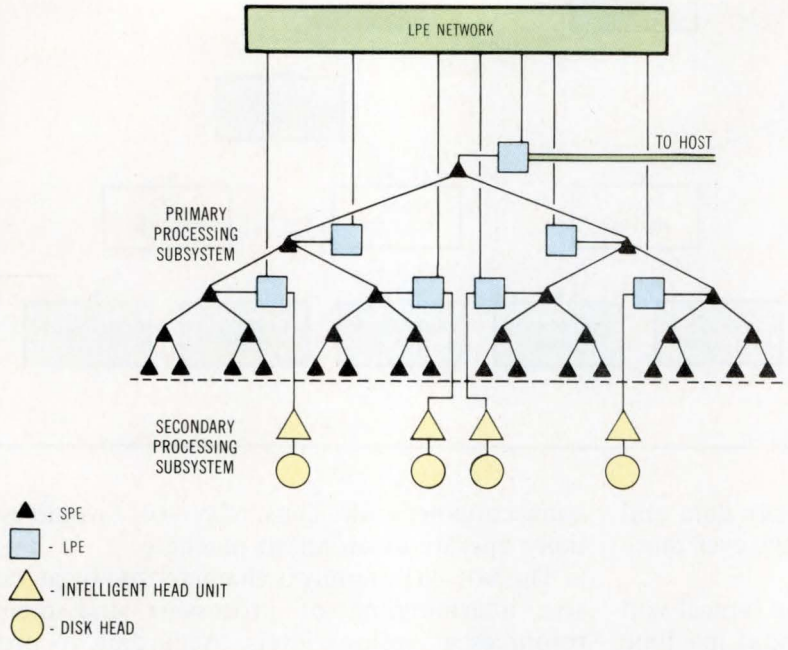

The NON-VON machine is based on a large number of 8-bit small processing elements (SPEs), which are organized as a complete binary tree. In its largest configuration, the NON-VON 4 will incorporate a small number of large processing elements (LPEs), each capable of serving as an independent control processor for some subtree of SPEs, along with a secondary processing subsystem based on a bank of intelligent disk drives. The LPEs would be connected through a high bandwidth switch called the LPE network.

of a basic processor, in this case containing eight elements, offers the potential or unprecedented computational concurrency, it also results in a local memory that is far too small to store meaningful programs. NON-VON's SPEs are thus forced to borrow their programs from one or more external instruction sources.

The design of each NON-VON family member also includes a secondary processing subsystem based on a bank of "intelligent" disk drives, connected with the primary processing subsystem through a high bandwidth parallel interface. Many Winchester disk drives, each of moderate size, have a small amount of processing hardware associated with each disk head. This hardware allows records to be inspected as it is operating and determines whether a given record is relevant to the operation at hand before transferring it to the primary processing subsystem. The hardware supports certain other operations (eg, hashing) that play a key role in many disk-based NON-VON operations.

While the above components are incorporated in all members of this family, essential differences exist between the three versions now being designed or constructed at Columbia. Specifically, NON-VON 1 and 3 include a single control processor, which is used to broadcast instructions for execution by the processing elements in the primary processing subsystem. With the exception of certain I/O operations, these machines function in an instruction-synchronous SIMD mode, with all SPEs simultaneously executing a single instruction "in lock step." (NON-VON 2 was only a paper concept.)

NON-VON 4, on the other hand, includes a number of large processing elements (LPEs), each capable of serving as a control processor for some portion of the primary processing subsystem. This gives the machine the capacity for MIMD and multiple SIMD operations, multitasking and multi-user applications, and such problems as physical simulation for which the NON-VON 3 tree would otherwise represent a significant communication bottleneck.

Additional enhancements anticipated for NON-VON 4 include a significant amount of storage (in the form of commercially available dynamic RAM chips) associated with each LPE. In addition to its use as storage for an individual LPE, this RAM can be used as swapping storage for the local RAMs incorporated in the SPEs with which the RAM is associated. In the NON-VON 4 machine, each mode above a certain fixed level in the primary processing subsystem tree is connected to its own LPE. (In a machine containing 256,000 SPEs, for example, somewhere between 511 and 1000 such LPEs should be provided.) Each LPE includes an off-the-shelf microprocessor (eg, a Motorola 68000 or National Semiconductor 16032), memory (between 256 Kbytes and 1 Mbyte), and custom hardware for interfacing with the rest of the machine.

The LPEs are mutually interconnected through a high bandwidth, multistage interconnection network. While the Columbia team has not specified details of this network, it has recently begun to investigate the possibility of using a "folded" 2-log-n-stage banyan-type network operating on a circuit-switched basis. The idea is to incorporate logic within the

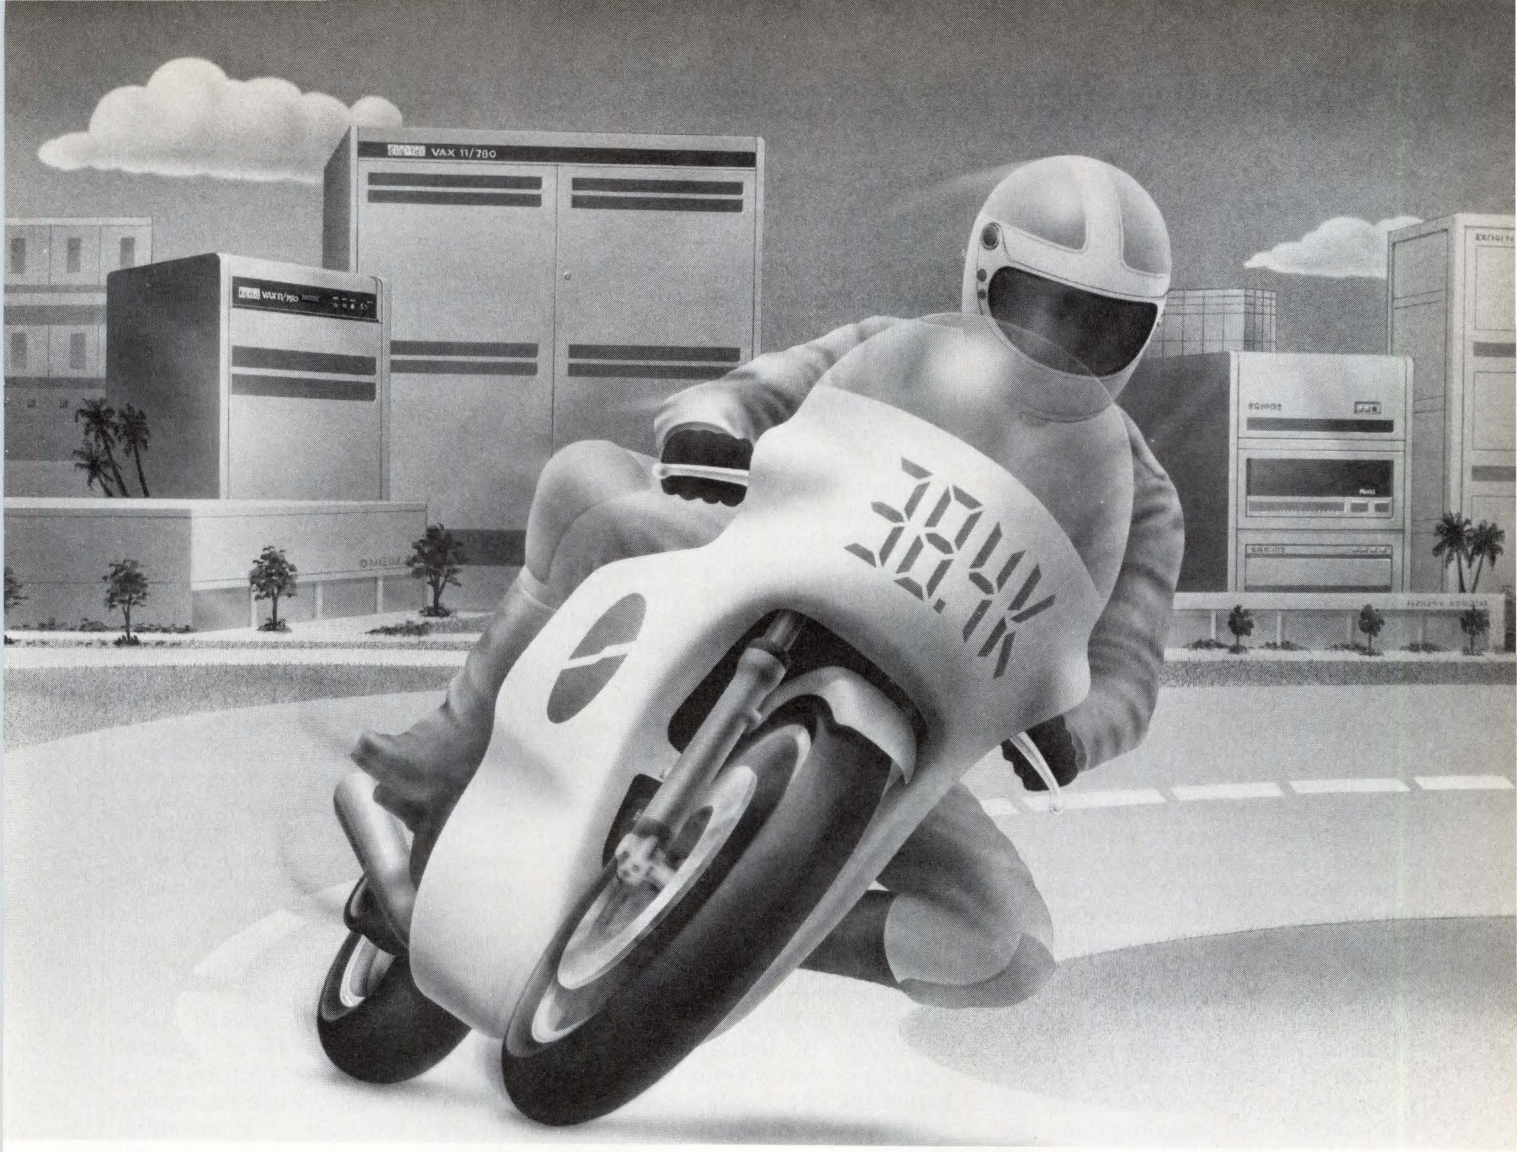

FAST TRACK Multiplexer for your VAX

ABLE Computer's VMZ/32HS ...the high-speed communications solution for your VAX.

Servicing 16 communication channels at baud rates up to 38.4K simultaneously, the VMZ/32HS multiplexer breaks all speed records.

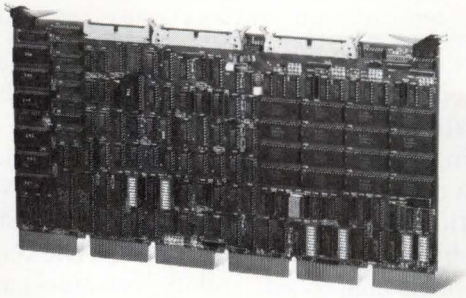

Maximize Responsiveness of Graphics Workstations

Speed up your CAD/CAM, business graphics, or any terminal intensive application. The ABLE VMZ/32HS outperforms the competition giving you fast screen response with minimum delays. And the VMZ/32HS supports 16 asynchronous communication lines with programmable DMA and modem control on all lines.

'Software compatible with DEC VMS operating systems, the VMZ/32HS requires no software development, modifications or maintenance.

Take the Lead, And Stay There

High aggregate throughput keeps your CPU in front of the pack. Our VMZ/32HS satisfies both present and future data communication needs when requirements demand high output and high reliability. When you need speed, the ABLE VMZ/32HS finishes first.

VMZ/32HS, breaking speed records across the board.

For more information on ABLE Computer's VMZ/32HS, contact your local ABLE representative, or call ABLE toll-free at 800-322-2253.

ABLE COMPUTER

The Communications Specialists

1732 Reynolds Avenue, Irvine, California 92714 In the Irvine area: (714) 979-7030. Or TWX: 910-595-1729

DEC, VAX and VMS are trademarks of Digital Equipment Corporation.

CIRCLE 47

Supercomputers for super problems

Advances in computing technology have brought about a significant decrease in the cost per computation. But, Professor Duncan H. Lawrie, director of the Laboratory for Advanced Supercomputers at the University of Illinois, Urbana-Champaign, believes that the demand for higher performance continues to exceed the capabilities of technology alone. Lawrie has contributed to the design of several large computers including the Illiac IV, and is currently involved in the Cedar supercomputer project.

The demand for higher performance is driven at least in part by what Lawrie believes to be a significant shift in the way computers are being used—a shift toward simulation and modeling. For example, in the design of electronic circuits, computers were once used simply to automate certain drafting operations. Now, however, they are used to do detailed, device-level simulation of circuit performance.

Computational modeling allows the design of an airplane's airfoil with 40 percent less drag than was possible using older experimental techniques. While the new techniques are being used today to design the wings of the Boeing 767, for example, the modeling of an entire aircraft is beyond the capabilities of current supercomputers.

Lawrie believes that the United States' continued technological superiority in many areas will depend on the availability of significantly faster computers. To achieve higher speed, today's commercially available supercomputers employ a pipelining technique. Like an assembly line, this technique completes a small part of each operation at each station in the line. The final "product" is an arithmetic (usually floating point) operation like multiply or divide. Thus, parts of many operations are in progress at any given time. This gives higher throughput. Unfortunately, the relative simplicity of the operations (and the latency or delay of each operation) limits the number of stations that can be put in a pipeline.

"To get still higher speeds, we must use multiple pipelines," Lawrie says. The trend today with commercially available machines (including those from Japan) is to use up to 16 pipelines. "I would expect to see this trend continue, with supercomputers in the 1990s using hundreds of pipelined arithmetic processors."

But, Lawrie feels that the use of a large number of processors raises two basic problems. First, to take advantage of the available processing power, we must learn how to use multiple processors on a single problem. Second, to build such machines, we must devise ways to provide adequate communication facilities between processors and to storage devices, he says.

To apply many processors to a single problem usually requires restructuring of the program or algorithm. Often, the program can be restructured automatically, as shown by the successful pioneering work of David Kuck's Parafrase project at the University of Illinois. However, sometimes completely new algorithms are needed. "While great strides have been made in both automatic restructuring and new algorithms, a great deal more work needs to be done," says Lawrie.

Communication capabilities are an important aspect of the design and use of multiprocessor supercomputers. Processors must be able to communicate among themselves, as well as with storage devices if they are to cooperate effectively on a problem. Traditionally, hardware designers have counted gates or packages to determine hardware complexity, while algorithm designers have counted arithmetic operations to determine computational complexity. "Designers have only recently begun to recognize the importance of communication complexity in both hardware and algorithm design," says Lawrie, "and it is probably the least understood, but most essential area of computational theory."

Lawrie and his academic peers are also concerned about the training of a new generation of users who will be able to use supercomputers effectively. "It is essential that universities get access to existing supercomputers, and that instructional and research programs be structured around this access," he says. Lawrie feels that the National Science Foundation has taken a significant first step to provide this access. But, he concludes that still more needs to be done.

individual switches that will support the rapid, parallel identification of some unblocked path through the network whenever any such path exists.

Although lacking any supporting data, Shaw believes that the circuitry for such a network might be simple enough to allow construction of a moderately high order switch with a number of independent input and output paths using a single high speed bipolar gate array chip. Since each switch chip would, in essence, function as a full crossbar switch, the use of higher order switches would reduce the number of network paths subject to contention. Additional performance advantages should result from the reduced number of chip-to-chip delays in a network built with higher order chips, which would have fewer stages generally. Another potentially attractive feature of this broadcast circuit switch approach is the fact that later stages in the network are automatically bypassed in cases where the source and destination LPEs have closely spaced addresses.

To illustrate how actual computations differ between a von Neumann-type machine and the NON-VON type, Shaw points to a relational join operation. The relational join operation takes two relations (tables) as arguments and produces a new table as output. Each row in the result relation is formed by concatenating two rows, one from each argument relation, that have the same value for the corresponding attribute (column).

On a von Neumann machine, this operation is typically executed by sorting the two argument relations on their respective join attributes (the columns on which the comparison is to be made). Then, a "merge-like" operation is performed to produce the

LOW POWER, CMOS 8086 SINGLE BOARD COMPUTER

Replace Your iSBC* 86/05 and get the Benefits of CMOS Technology Without Redesign

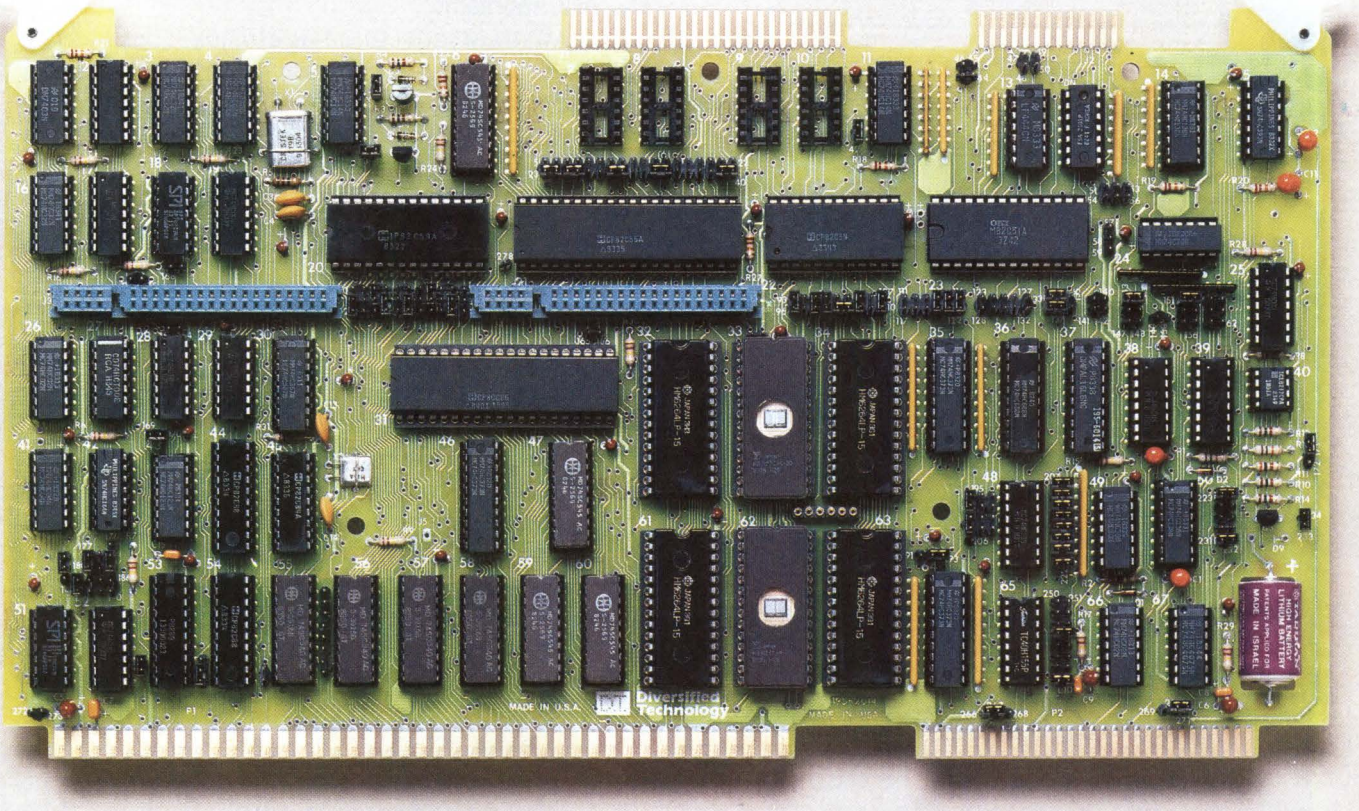

Compare these Key Features:

	INTEL iSBC* 86/05	DIVERSIFIED TECHNOLOGY CBC 86C/05
Bus Type	MULTIBUS*	MULTIBUS*
CPU	8086	80C86
5V Operating Current	4.7 amps max.	200 milliamps max.
Operating Temperature	0°C to 55°C	0°C to 70°C
Supplied RAM (Bytes)	8K	16K
RAM Battery Back-Up On Board	No	Yes (2.5 yrs. data retention)

Whether your application is an existing system upgrade or a low-power CMOS system design, DTI's CBC 86C/05 will make your job easier. Full hardware and software compatibility with the iSBC* 86/05 allows you to gain the low power advantages of CMOS without sacrificing system performance or your development investment. You even get fast nonvolatile memory on-board, and a full-spec MULTIBUS* interface to maintain compatibility with existing systems.

For more information regarding the CBC 86C/05 CPU board, or any of our other all-CMOS MULTIBUS* boards, contact Bill Long, CBC Product Manager at **(601) 856-4121.**

ASK ABOUT OUR NONVOLATILE CMOS RAM FOR MULTIBUS AND LSI-11 TOO!

Above specifications taken from manufacturers' current published data.
**MULTIBUS and iSBC are trademarks of Intel Corp.*

Diversified Technology
An Ergon Co.

P. O. Box 748, Ridgeland, MS 39157
Telex 585326

CIRCLE 48

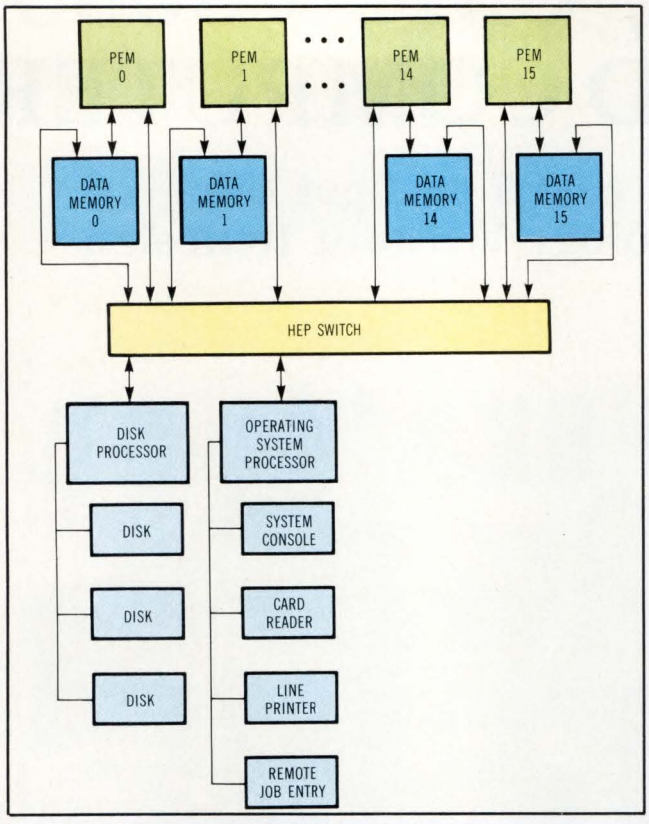

The Heterogeneous Element Processor (HEP) is a commercial computer that implements an MIMD architecture and is capable of executing 160 MIPS. It has 16 process execution modules (PEMs), each with its own data memory, and each executing up to 64 processes concurrently. Access to the processor part of the computer and peripheral units is via the HEP switch.

resultant rows. The running time of the algorithm is typically dominated by the logarithmic time (n log n) required for sorting. The SIMD NON-VON machines reduce this time to a linear function and provide significant savings in absolute running time, given certain assumptions about the input data. However, in NON-VON 4, where these assumptions are weakened even further, the time required to join large relations is reduced far below that of all practical and known theoretical database machines to which Shaw's team compared the NON-VON 4.

In cases where the data does not exceed the capacity of the primary processing subsystem, the NON-VON algorithm (for both SIMD and MIMD machines) involves the broadcasting of each join value through the primary processing subsystem and the associative identification of all matching rows on the other relation. In the case where the argument relations exceed the capacity of the primary processing subsystem, however, the argument relations must be divided into key-disjoint partitions. These partitions have the property that no join value appears in more than one partition.

In NON-VON 4, key disjoint partitioning can be accomplished by storing the argument relations spread across the various disk heads in the secondary processing subsystem, by hashing each join value passing under each head, and by dynamically routing the rows through the LPE network to a different disk head (depending on its hashed join value). This process produces k different key-disjoint partition files, where k is the number of disk heads involved, and requires time proportional to the capacity of the full file divided by k.

The HEP approach

One MIMD architecture has already been successfully implemented in a commercial product. Denelcor Inc's (Denver, Colo) Heterogeneous Element Processor (HEP) is a parallel processor suitable for general-purpose applications. HEP is programmed in Fortran and can execute up to 160 million instructions per second (MIPS) in its maximum configuration.

As in any MIMD architecture, entire programs or pieces of programs execute concurrently. Although each piece of program (process) has an independent instruction stream operating on its own data stream, processes cooperate by sharing data and solving parts of the same problem in parallel. Thus, throughput can be increased by a factor of n, where n is the average number of operations executed concurrently.

The HEP architecture consists of a number of process execution modules (PEMs). Each PEM has its own data memory bank and is connected to the HEP switch. In addition, the operating system processor and the disk processor are both connected to the switch. Each PEM can access its own data memory bank directly, but access to most of the memory is through the switch. The HEP implements the MIMD architecture with up to 16 PEMs. Each PEM can execute up to 64 processes concurrently. Any number of PEMs can cooperate on a job, or each PEM can run several unrelated jobs. All of the instruction streams and their associated data streams are active.

Parallel processing is handled within an individual PEM by pipelining instructions so that several are in different phases of execution at any one moment. A process is selected for execution each machine cycle, a single instruction for that process is started, and the process is made unavailable for further execution until that instruction is complete. Because most instructions require eight cycles to complete, at least eight processes must be executed concurrently in order to fully use a PEM.

However, memory access instructions require substantially more than eight cycles. Thus, in practice, about 12 concurrent processes are needed for full utilization, and a single HEP PEM can be considered a virtual 8- to 12-processor machine. If a given application is formulated and executed using p processes

Taking another look at parallelism

While scientific computing in the 1990s will surely use parallel architectures, C.N. Winningstad, founder and chairman of Floating Point Systems, Inc, sees several roads toward implementing parallel operations in a scientific computer. Winningstad has brought his company to the forefront of array processing, a form of parallelism used in today's von Neumann-type architectures.

"Naturally, there are pioneers who realize that the von Neumann approach is a primary bottleneck, and we have seen learned articles on alternative machines using single-instruction, single-data streams, or single-instruction, multiple-data streams, or even multiple instruction, multiple-data streams," he says. "While early machines performed one operation at a time, as specified by a single instruction, more sophisticated machines started specifying many operations to be performed simultaneously using one instruction."

According to Winningstad, this has led to minicomputer people confusing or using interchangeably the million instructions per second and million operations per second performance ratings. "Since the state-of-the-art in semiconductors set a speed limit on the rate of processing instructions, clearly, a machine that specifies more than one operation per instruction will produce more operations per second," he says.

This means that the hardware can be broken into parallel units that operate simultaneously in both space and time. For example, an adder, a multiplier, a divider, and a square root extractor can all be available for simultaneous operation. Winningstad claims that paralleling dissimilar units is easier than paralleling identical units. In an example using five adders, a record must be kept of which adder is available and from which the result is expected to appear.

Another form of paralleling is what Winningstad calls "paralleling in time" or pipelining. He explains the concept thus, "Floating point multiplication consists of roughly three actions—multiply the mantissas, add the exponents, and normalize. Normalizing reestablishes the decimal point to the correct position. These actions may take several machine cycles to accomplish (multiply mantissas), or one machine cycle (normalize), and can be done in parallel space where one set of circuitry is adding the exponents while another set is multiplying the mantissas."

In this operation, the normalizer circuitry is idle until the other operations are complete. Hence, there is idle hardware before and during the normalization routine. "If we do things right, we could arrange the multiplier to take, for example, three machine cycles to produce the product," says Winningstad. "But we could allow a new set of arguments to be introduced into the multiplier each cycle while waiting the three cycles to get the first product. A new product would appear each cycle thereafter. Thus, three separate multipliers in space, each given a pair of arguments sequentially on successive machine cycles, are no longer required. Instead, one pipelined multiplier with the parallel going on in time instead of in space, delivers a product each cycle." One input and one output, instead of three of each, greatly simplifies keeping track of where the arguments go and resultants appear.

Winningstad calls the above operations "fine-grain parallelism," as opposed to paralleling whole computing machines. He notes that to get significantly greater speeds out of von Neumann machines (which already use semiconductors efficiently and are architecturally matched to a problem) one has to break the problem into parts that can run simultaneously and place them in separate von Neumann computers as a parallel operation (macro parallelism).

While the von Neumann computer can handle most ordinary problems, there are classes of problems which require literally tens to hundreds of billion floating point operations per second (tens to hundreds of GFLOPS). Also, artificial intelligence will require tens to hundreds of billion logical inferences per second (GLIPS).

Winningstad does not see the present availability of VLSI chips as the answer. While ICs may get factors of tens or so faster, and scaling and new materials like gallium arsenide may help, he looks for factor improvements in the hundreds and thousands to solve future problems. "For those problems requiring GFLOPS or GLIPS speeds, we can expect hundreds of arrays or thousands of computers working on one problem," he notes. "The only two factors impeding progress today are the computer's physical size, and software complexity."

(where p is an integer from 1 to 12), then execution time for the application will be inversely proportional to p. Each PEM has its own program memory to prevent conflicts in accessing instructions, and all PEMs are connected to the other data memory banks through the HEP switch. The HEP switch is a high speed, packet-switched network.

Because these connections, the number of switch nodes increases more rapidly than the number of PEMs. As one changes from a 1-PEM system toward the maximum 16-PEM configuration, the transmittal time through the HEP switch increases. Such latency increases the number of processes that must be run in each PEM to achieve full utilization.

At the Los Alamos National Labs, (Los Alamos, NM) scientists have analyzed the inherent latency in the parallel architecture of the HEP machine. Although there is not enough data yet to provide good estimates on how fast latency actually increases with the number of PEMs, their experience with a 2-PEM system suggests that a 4-PEM system will require about 20 concurrent processes in each PEM.

A critical issue in MIMD machines is the synchronization of processes. The HEP solves this problem by attaching to each 64-bit data word an extra bit that is set to full each time a piece of data is stored and cleared to empty each time a piece of data is fetched. In addition, two sets of memory

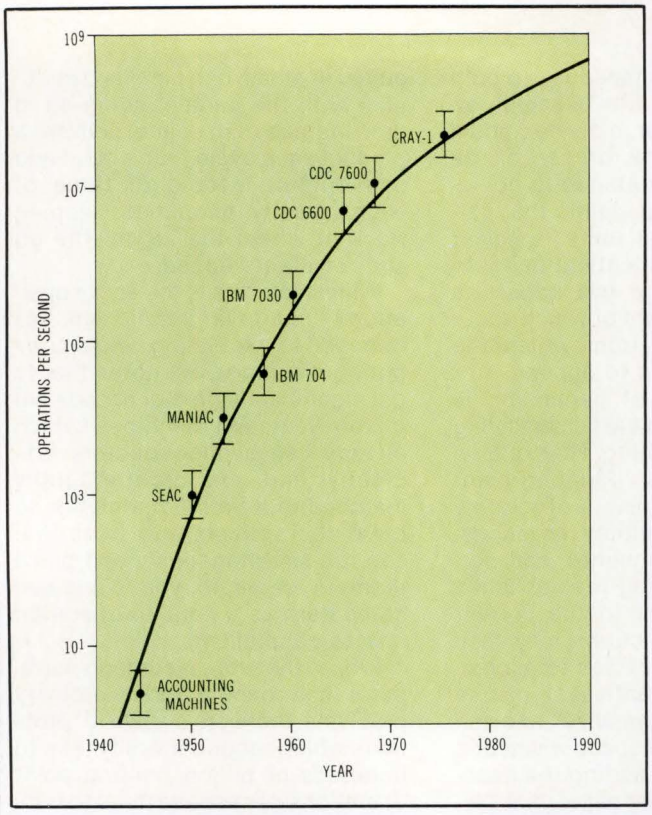

Most experts agree that some form of parallel architecture must be implemented to reverse the decline in the rate of increase of processor speed. Scientific and engineering problems of the next decade will require an increase in speed of at least two orders of magnitude over current machines such as the Cray-2.

instructions are used. One set is used normally throughout most of the program. This set ignores the extra bit and will fetch or store a data word regardless of whether the bit is full or empty. Data may then be used as often as required in a process.

The second instruction set uses asynchronous variables. Typically, this set is activated only at the start or finish of a process, acting as a barrier against interference from other processes. The set will not fetch from an empty word or store into a full word. Thus, to synchronize several processes, an asynchronous variable is defined that can only be accessed using the second set of instructions, at the appropriate time by each process. A process that needs to fetch an asynchronous variable will not do so if the extra bit is empty and will not proceed until another process stores into the variable, setting it full. Because the full and empty properties of this extra bit are implemented in the HEP hardware at the user level, the usual synchronization methods such as semaphores can be used without operating system intervention.

Several small Fortran codes have been run on the HEP at Los Alamos. One such code is SIMPLE, a 2000-line, two-dimensional, Lagrangian, hydrodiffusion code. By partitioning this code into processes, about 99 percent of it can be executed in parallel. The speedup on a 1-PEM system is observed to be close to linear for up to 11 processes and then flatten out, indicating that the PEM is fully utilized. Achieving this degree of speedup on any MIMD machine requires a high percentage of parallel code.

The Los Alamos Lab is one national laboratory that performs engineering and scientific tasks for the government using the latest computer technology. As such, it and its sister laboratories are at the forefront of evaluating computer technology, specifically supercomputers.

Scientifically speaking

At a Frontiers of Supercomputing conference held last year in Los Alamos, a group of academic, industry, and government representatives discussed the issues leading to the next generation of supercomputers. Here, the term "supercomputer" was defined as the most powerful computer available at a given time. The power of the computer was measured by its speed, storage capacity, and precision. Today's commercially available supercomputers, such as the Cray-1 from Cray Research (Mendota Heights, Minn) and the CYBER 205 from Control Data Corp (Minneapolis, Minn), have a peak speed of over 100 MIPS, a memory of 4-million 64-bit words, and a 64-bit precision level. Most scientists looking to solve complex engineering and scientific problems feel that a speed increase of at least two orders of magnitude is required to make significant progress. The rapid increase in speed at the start appears to level off. Thus, to restart a dramatic increase in the curve, engineers must implement parallel architectures.

The scientific talks presented at the Los Alamos conference were organized into sessions covering advanced architecture, supercomputing at Los Alamos, software, and algorithms and applications. The session on parallel computing architectures had computer manufacturers and academic researchers predicting that by the end of the decade, commercially supplied systems with eight or more processors will be the norm. The HEP machines are but one example. That prediction, however, falls short for those who say that dramatic increase in performance can only come when many fast processors are used to solve complex problems.

Currently, most academic supercomputer projects make extensive use of VLSI, exploiting 32-bit microprocessors and 256-Kbit memory chips. Many of them can be classified as "dance hall"machines. Dance hall systems have processors aligned along one side, memories on the other, and communication geometry to bind them together. High performance operation on these systems necessitates algorithms that keep all of the processors constantly "dancing."

BOB, CAROL, TED & ALICE

How Four Innovative Design Engineers Found Happiness Using The New DASH-1™ Schematic Design System.

BOB: "DASH-1 keeps me under budget and ahead of schedule..."

"DASH-1 has freed my designers from the drudgery of manual schematic design and documentation. Now, we can knock out complex designs in a fraction of the time it used to take. Using the huge DASH-1 parts library (on disk) we can instantly call up any symbols we need — all with pinouts and pin functions. Also, our amazing DASH-1's automatically generate Net Lists, Lists of Materials, Design Check Reports and other critical documents. I'm a hero in my company now, with productivity at an all-time high."

CAROL: "DASH-1 helped save our CAD Systems..."

"My company has a big investment in large CAD systems. But for a long time I noticed our engineers were frustrated. They'd send in schematic sketches only to get back drawings and documents that were completely different. The poor engineers would spend hours checking, changing and red-lining. DASH-1's family of CAD translators, to front-end all of our large systems, has changed all this. And back annotation is easy. DASH-1 has provided our engineers easy access to those previously difficult-to-use and inaccessible CAD tools."

TED: "Imagine! Simulation right at my desk..."

"I'm even more sold on my DASH-1 with its new simulator. After I complete my schematic, I have a closely coupled logic and fault simulator that I can run at my desk. I can analyze my designs much faster — and put them into production without delay. FutureNet calls it the DASH-1 CADAT.™ I call it a miracle because I can handle designs as large as 10,000 gates right at my desk, and up to 100,000 gates using a FutureNet CAD translator. Thanks to DASH-1, simulating is stimulating!"

ALICE: "DASH-1 keeps growing — as we do..."

"Ours was a start-up company — but it's growing fast. One of our first investments was a DASH-1 (about the price of an IBM PC or XT system) to speed up schematic design and documentation. Since then, FutureNet has added all kinds of low-cost enhancements. Things like CAD translators; direct-connections to most computers; the STRIDES hierarchical package with 99 levels of structured design nesting; the powerful CADAT simulator, pen plotter options, and hot new features, like Rubber Banding, Tag and Drag, and Snap. We're with FutureNet... for now and the future."

FutureNet®

FutureNet Corporation • 6709 Independence Avenue
Canoga Park, CA 91303-2997 • TWX: 910-494-2681
(818) 700-0691

Authorized IBM Value-Added Dealer

Productivity of the Future... today.

CIRCLE 49

Why standards for the 1990s?

The computer industry has wasted enormous sums of money and scarce engineering resources in a seemingly endless treadmill of "reinventing the wheel." If this continues, the Japanese may overtake the United States in key computing technology areas, according to C. Gordon Bell, chief technical officer for Encore Computing Corp and former vice president of engineering for Digital Equipment Corp. At DEC, he managed the design of the PDP-4, -5, and -6, and led the design team that conceived the VAX architecture.

As a seasoned architectural veteran and computer pioneer, Bell has many reasons why standards will be necessary in the 1990s. "Standards form the constraints necessary for evolution, *not* revolution, into the next decade of computing," says Bell. "Constraints save design time by narrowing the search for new products and processes. They also permit building on past work in a hierarchical fashion, rather than having to start each new design with silicon. Standards provide a real intellectual discipline to leverage computing further and faster," he says.

When asked if standards will stifle innovation, Bell says they can free us for innovation while leaving past accomplishments as the base on which to build. He feels that "more than any other factor, the lack of standards impedes technological progress and lowers productivity. Redundancy in product development ties up critical resources on the reinvention of trivia with a shortage of resources for solving hard problems such as speech and video communication, intelligent programs, revolutionary machines, and fully automatic production. It is a problem of a shortage of engineers, in effect, a shortage of leadership that lets engineers build disconnected, overlapping, low technology products."

Bell has 10 rules for setting and using standards. Topping the list is the call to either make or follow the standard. Then, one must be prepared to react quickly and follow when the *de facto* standard changes, or to change the standard when it is wrong.

Bell encourages people to realize that someone (person(s), company, or companies) must be responsible for defining, implementing, and caring for a standard, and that the number of organizations responsible for a standard should be kept to a minimum. Also, almost any standard is more important than a highly defined optimum because progress is often made through regression. He urges provision and planning for evolution, as it is often the fastest way.

Standards should be made based on real experience, not by a design committee, says Bell, who adds that if you have not lived with a proposed standard, do not adopt it. He also feels that a standard must be precise, understandable, applicable, and useful at many levels of detail. And, only one standard is needed (or a few) for the same function; a standard should aim toward unifying a set of alternatives. "Ideally, a standard should define the interface between sets of parts, not just two parts," Bell says.

At New York University's Courant Institute, a team of computer scientists are developing a version of the MIMD parallel computer. Labeled Ultracomputer, it uses 4096 autonomous processing elements. These scientists believe that previous high performance machines were constructed from increasingly complex hardware structures and ever more exotic technology. The Ultracomputer offers a simpler alternative that is better suited to advanced VLSI technology—high performance is obtained by assembling large quantities of identical computing components in an effective manner.

The Ultracomputer they envision has roughly the same component count as today's large machines. However, the number of different component types is much smaller, with each component being a sophisticated one-chip VLSI system. Such machines would be three orders of magnitude faster and would have a main storage three orders of magnitude larger than present day machines.

To use such a machine effectively, it is necessary to prevent serial bottlenecks from being introduced by either the software or hardware. The Courant team has found that by using a software routine known as the fetch-and-add primitive, programmers are able to avoid critical sections that have previously seemed inherent. Thus, programmers are able to write fully parallel programs for many important problems in system and application software. In effect, the Ultracomputer can execute concurrent fetch-and-adds directed to the same location in the time required for just one such operation.

The fetch-and-add operation is a simple, yet effective, interprocessor synchronization operation that permits concurrent execution of operating system primitives and application programs. Inherently serial procedures, in which one processor works while the others wait, become bottlenecks that drastically reduce the intended performance. For any parallel architecture, the relative cost of a serial bottleneck rises linearly with the number of processors involved. Thus, users of these machines will be anxious to avoid critical, and hence necessarily serial, code sections. This is true even if these critical sections are sufficiently short so as not to cause performance penalties on today's computer, says Allan Gottlieb, the Courant project leader.

If the fetch-and-add operation is available, many important algorithms can be operated on in a completely parallel manner (ie, without using any critical sections). For example, concurrent executions of a fetch-and-add operation yield consecutive values

that can be used to index an array. If this array is interpreted as a sequentially stored queue, the values returned can be used to perform concurrent inserts; analogously, fetch-and-add can be used for concurrent deletes. Such techniques can also be used to implement a decentralized operating system scheduler. Gottlieb says that no other completely parallel solutions to this problem are currently available. To illustrate the nonserial behavior obtained, a single half-full queue has its concurrent execution of thousands of inserts and thousands of deletes all accomplished in the time required for just one such operation.

Concurrency creates the problem of checking where a problem has gone awry. In making large, parallel computers work reliably, and with high logic densities and many parts operating concurrently, it can be difficult to understand what is happening in these machines if they are not working right. That is the worry of Jack B. Dennis, head of the supercomputer project at the Massachusetts Institute of Technology (MIT) in Cambridge, Mass. As part of his computer prediction, Dennis says that the supercomputers of the 1990s will embody full-coverage fault detection for single, independent hardware failures.

A 1-GFLOP faultless machine

In such a computer, whenever a program runs to termination without detecting a fault, the user is guaranteed that the results are not invalid due to hardware failure. The proposed 1 billion floating point operations per second (GFLOPS) data flow supercomputer by Dennis's team has array memories that hold the data base of a large scale computation. The functions of processing elements in this architecture are performed by cell blocks and functional units. The cell blocks hold the data flow instructions and perform the basic function of recognizing which instructions are ready for execution. The machine supports three classes of instructions: those that call for the floating point arithmetic to be done; those that require operations such as duplicating values and performing tests; and those that build arrays or access elements of arrays. The first class is sent to functional units, the second is executed within cell blocks by a simple arithmetic/logic unit, and the third is sent to the array memory units where arrays of data are held.

Most computation takes place within just the cell blocks and functional units. There, instructions are identified for execution in cell blocks, then sent off to the functional units for execution. Resultant packets are sent through a routing network to enable other instructions. Information is brought in from the array memories on an as-needed basis to participate in the computation, and the main results of the computation are dumped into the array memories and held for later use. Thus, the number of array memory modules provided is less than the number of processing elements. For the applications that Dennis's group studied, packet traffic to the array memories is only a small fraction of the traffic needed among processing elements of the machine.

In operation, the data flow supercomputer has each cell block recognize and process enabled instructions at a rate of 1 MIPS; thus, 1024 cell blocks and a similar number of functional units are required. The number of array modules is 32, about one-eighth the number of processing elements

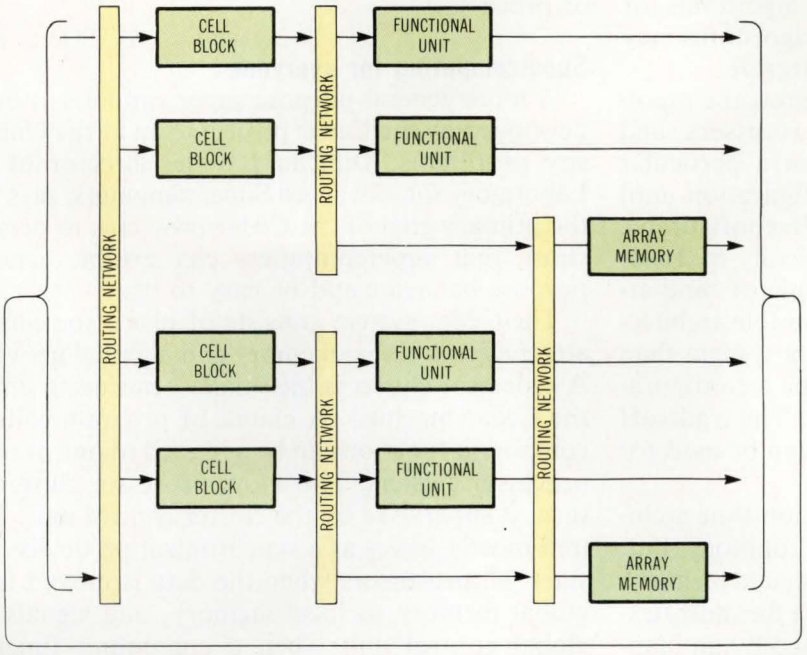

The Massachusetts Institute of Technology's data flow supercomputer proposes to process at a 1-GFLOP rate. This requires a total of 1024 cell blocks and a similar number of functional units. However, only 32 array modules are required to hold the data base for a particular application. The three core blocks communicate through routing networks. Each cell block recognizes and processes enabled instructions at a rate of 1 MIPS.

because the data rates required are much less. Cell block memory holds data flow instructions, while array memory unit memory holds the data base for the application. Thus, each cell block consists of some control hardware and RAM, as do the array memory modules. For the array memories, close to 32,000 devices would form a memory of 64 million 32-bit words using 64-Kbit chips. In the cell blocks, 16,000 devices would allow 16 million words of instruction memory using 16-Kbit chips in order to achieve greater speed.

Dennis claims that using this architecture, if the cell block control, the functional units, and the array memory control can each be built using one to three chips, then the chip count will be dominated by memory. The next largest population of units would be the 5568 two-by-two routers used to implement the several routing networks. But the router requirement is not really dominant because, by using packet communication and trading off latency for throughput, the hardware requirement for the routing networks is modest compared to the rest of the machine.

A critical point regarding the new breed of architectures is matching architecture topology with the algorithms so that parts of the algorithms can correspond to certain architectural elements. James Browne, of the University of Texas at Austin, observes that there are three general models of massively parallel systems based on whether decisions concerning the communication, synchronization, and granularity of parallel operations are fixed at design time, compile time, or execution time. In fixed-net architectures, these decisions are made at design time. Systolic arrays containing hardware components designed to carry out specific algorithms are one example of a fixed-net architecture. Such systems provide fast execution of the algorithms for which they are designed, but other algorithms may have to be adapted to fit the architecture.

In bind-at-compile-time architectures, the topology of the interconnection of processors and memory can be reconfigured to suit a particular algorithm, but remains in that configuration until explicitly reconfigured. The Texas Reconfigurable Array Computer (TRAC) at the Unversity of Texas at Austin, for instance, is an example of bind-at-compile-time architecture. Reconfigurable architectures must have more complex control systems than those that bind at design time, and the reconfiguration process itself takes a long time. The trade-off is that reconfigurable architectures can be used for many different algorithms.

On the other hand, bind-at-execution-time architectures do not explicitly specify topology but, through use of control structure and shared memory, allow any topology. These systems are the most flexible, since no algorithm is constrained by an inappropriate architecture. The price is the overhead required to synchronize and communicate information. Data flow machines are an example of bind-at-execution-time systems.

TRAC can implement multiple models of parallel computation, Browne says. The TRAC architecture can incorporate almost any type or class of processor. And, its strength lies not only in its computing elements, but in its multiple modes of communication.

The ideal communication system, of course, will have zero latency, infinite bandwidth, and realize arbitrary topologies of communication. Browne says that achieving all these properties in an interconnection network-based architecture can be best approached through the use of both circuit and packet communications. The TRAC interconnection network realizes both data movement modes. Circuit-based movement of data is based on selective activation of circuits in the switchable trees. If one processor deactivates a circuit in a switchable tree, then another processor activates a circuit in that tree, the entire contents of that memory is moved beween the address spaces of the two processors.

TRAC also implements a complete packet switching operation. A processor directs any memory to which it has an active circuit to send a data packet selected from its contents to any other processor. Thus, all processors may simultaneously transmit packets.

Analyses of TRAC's switch nodes have shown that it is possible to construct nodes with delays of the order of a few nanoseconds. The length of the path between processors and memories is bound by log n, where n is the number of processors. This is so that total delay stays within 39 ns even for large numbers of processors.

Supercomputing for everyone

A more general-purpose supercomputer is being developed by the Cedar project team at the University of Illinois. Duncan Lawrie, director of the Laboratory for Advanced Supercomputers, says that the primary goal of the Cedar project is to demonstrate that supercomputers can exhibit general-purpose behavior and be easy to use.

The Cedar system consists of processor clusters attached to a global memory and a global network. A processor cluster is the smallest execution unit in the Cedar machine. A chunk of program called a compound function can be assigned to one or more processor clusters. The entire processor cluster, in turn, is supervised by the cluster control unit. This unit mostly serves as a synchronization device that starts all processors when the data is moved from global memory to local memory, and signals the global control unit when a compound function

MISSION ACCOMPLISHED.

Dimensional stability. Wear resistance. Consistency. Three reasons why this minifloppy was made with thermoplastic composites from LNP.

When Drivetec set out to design a smaller, lighter, more reliable minifloppy disk drive, they knew each injection molded part had to meet stringent specifications for durability and stability. That's why they contacted the thermoplastics experts at LNP.

It was a tough assignment, but we accepted it. And soon delivered a complete series of our Thermocomp® resins to Drivetec—on time and on spec. Not bad, considering that 23 out of 25 plastic parts used in this minifloppy were made from composites specifically formulated for the application.

But there's another reason why demanding customers like Drivetec, manufacturers of the high capacity 3.3 megabyte drive, do business with LNP. You see, we're large enough to respond to your thermoplastic composite needs, yet flexible enough to **listen**. So next time you have a critical mission in mind, call LNP right away. We'll accomplish it without delay.

LNP Corporation, 412 King St., Malvern, PA 19355 (215) 644-5200

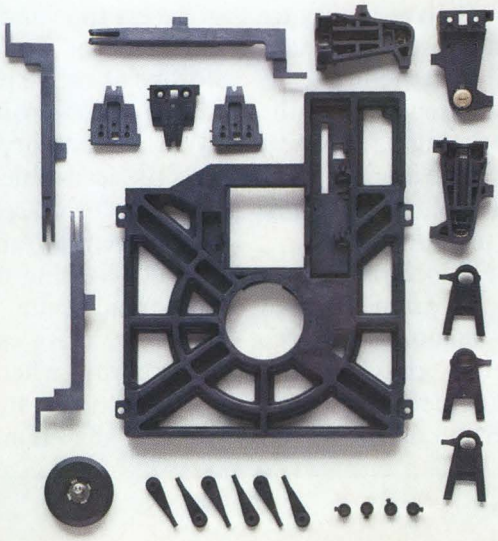

That's ingenuity. That's LNP.

LNP
ENGINEERING PLASTICS
A Beatrice Company

CIRCLE 50

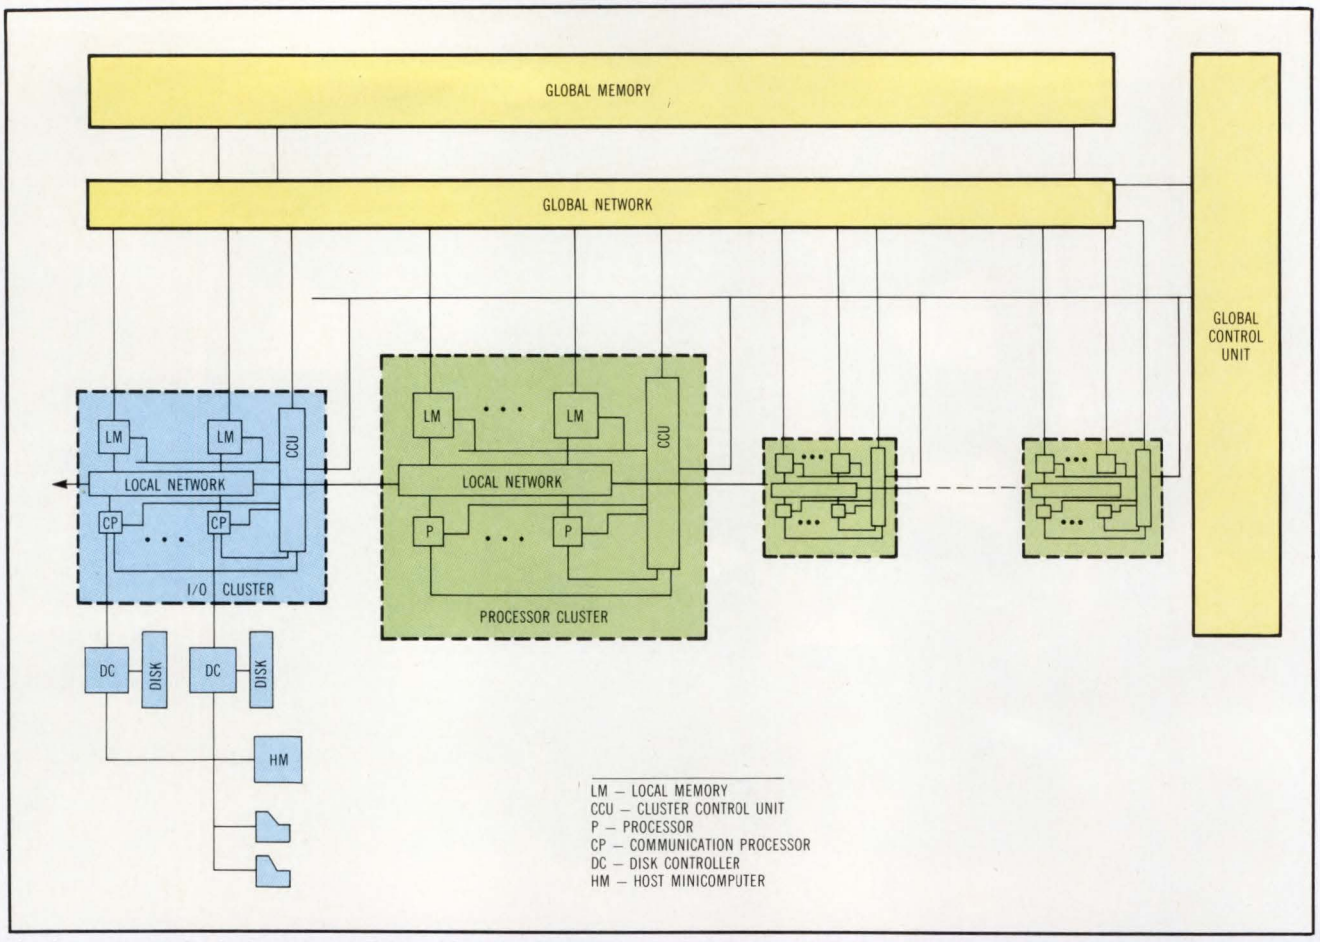

An important criteria in any parallel architecture is to provide fast access to the memory for data. Many systems have been built that have severe constraints in the memory or interconnection network. The University of Illinois's Cedar supercomputer has a shared memory and a switch design that provides high bandwidth over a range of application areas. Clever hardware and software techniques allow designers to access each memory independently rather than accessing data segments in a synchronous, deterministic way. Once moved to local clusters, data is processed by local processors, using local memory.

execution is finished. According to Lawrie, there are several well-recognized difficulties in providing data storage in a multiprocessor. The first challenge is to provide rapid data access at the rate required by the processors, while at the same time providing a common, shared memory. Access to a common memory shared by many processors implies, of course, short delay and high bandwidth communication facilities.

The Illinois team's approach involves a mixture of compilation hardware techniques. The Cedar memory system has a great deal of structure to it, but users need only concern themselves with the global shared memory. However, the fast local memories present in the design can be used to mask global memory access time. Each cluster of eight processors contains eight local memories.

User-transparent access to these local memories is provided in several ways. First, program code can be moved from global to local memories in large blocks by the cluster and global control units. Second, the optimizing compiler generates code to cause movement of blocks of certain data between global and local memory. Third, automatic caching hardware using the local memories is available for certain data where the compiler cannot determine *a priori* the details of the access patterns, but where freedom from cache coherency problems can be certified by the compiler.

Cedar provides private, partially shared, and fully shared address spaces, each of which can be segmented, and each of which has different properties of access time and data sharing. Private memory is intended for local unshared variable storage and storage of active program segments. It provides the fastest access from processors and is implemented by using memory modules that are local to and directly accessible by each processor.

The fully shared memory is implemented using global memory and access is provided via a large global switching network. This memory is used to store read/write data shared by all processors. However, there exists a certain amount of time in the program when the data is read-only, and during those times, the data can be stored in private or partially shared memory, thereby improving access

time. The compiler can determine whether the data is read/write shared or read-only.

In the Cedar machine, partially shared memory is implemented in an area of the private memories (ie, cluster memory) and access is provided via a communication facility within each cluster of processors. In many algorithms that require shared, read/write data, the amount of sharing can be limited to small clusters of processors for significant amounts of time, thus saving global switch bandwidth. This consideration motivates partially shared memory. It also provides certain economies of storage for fully shared, read-only data that might otherwise be redundantly stored in private memory segments. The average Cedar user will normally be concerned only with private and fully shared memory, and in fact need worry only about the latter. Compilers and advanced users may use the other addressing modes for optimization of memory access time.

Another project that aims to be general purpose in nature is a hybrid machine that combines salient properties of reduction, data flow, and von Neumann processes. This Rediflow concept is promulgated by computer scientists at the University of Utah (Salt Lake City, Utah) from a derivation of both reduction and data flow concepts. Rediflow is essentially a collection of ideas related to multiprocessor system design and attendant software capabilities. The intent of Rediflow multiprocessing is to evaluate functional language on multiple processors, according to project leader Robert Keller.

Reduce instructions and flow data

The motivation for using functional languages stems from the fact that functional languages usually contain some implicit concurrency. Yet, they also are determinate or speed-independent, in that they are guaranteed to give the same results no matter how many processors are involved in their execution, and do so independently of the physical aspects of communication between those processors. As such, Keller says that these languages are ideal for the programming of multiprocessors when little concern over the distinctions between them and uniprocessors is desired. Determining the use between uniprocessors and multiprocessors is essential for most applications, while certain exceptions can be handled with minor extensions. As an example, Keller points to a number of successfully programmed distributed database applications, including those of his team, involving concurrent updating.

While the connection of a single processor with its memory has been pejoratively called the "von Neumann bottleneck," Keller's team is convinced that this connection is a powerful feature to be exploited as much as possible. A large number of such "bottlenecks" operating concurrently can give a high aggregate bandwidth, much higher than a dance hall configuration with the same number of processors and memories, and with less attendant latency or memory accesses. While these processor/memory pairs do not usually operate in isolation, if the communication between the components of the pair occur much more often than communication between pairs, then a boudoir configuration should be superior. (In a boudoir configuration, each processor is closely paired with a memory, and a network of switches is used to communicate between such pairs.)

The reason for using functional languages stems from the fact that they usually have some implicit concurrency.

When using large numbers of processors, Keller says that it is unattractive to use a centralized task queue from which processors seek work. Such a queue creates a bottleneck and becomes unreliable. Keller claims that applying machinery powerful enough for concurrent computation to inherently sequential segments forces the overall speedup to be diluted. For example, a major difficulty with the reduction model is its memory intensiveness. It imitates an elegant mathematical model of functional languages in which data values are never modified in place; they are only created, and destroyed by reclaiming them from storage. To do this for every conceivable operation means that much time is spent recycling storage. Although a certain amount of this can be done concurrently with other processing, the overhead remains significant.

Therefore, a desirable goal is to combine the load-spreading potential of reduction with other methods that are not so storage intensive. The approach taken in Rediflow entails von Neumann processes, which appear internally as encapsulated sequential processes while externally, they appear as a form of data flow functions. Internally, von Neumann processes are ordinary sequential programs; operations appearing to be file I/O ("get" and "put") are used to communicate internal data values to and from the environment in the form of "tokens" moving on channels.

The interface from a von Neumann process to a reduction-implemented function solidifies a stream of token values into a stream data structure. The interface from a reduction-function to a von Neumann process does the opposite. These functions are comparable to ones that are used for external stream I/O in implementations of the reduction model.

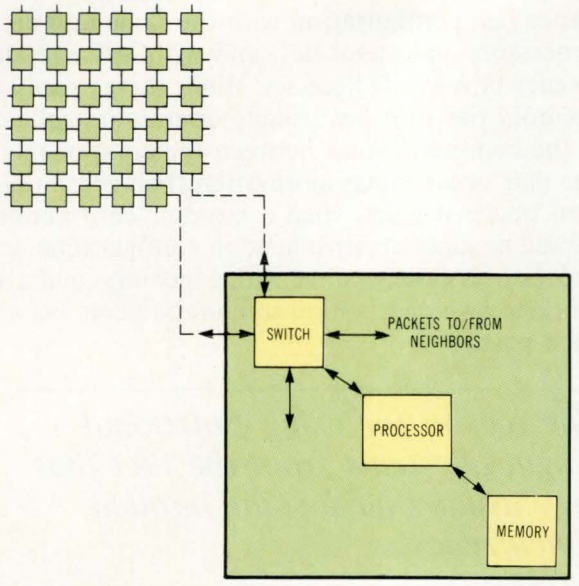

The Rediflow concept combines the ideas of data flow and reduction demand-driven machines. Developed by computer scientists at the University of Utah, the concept materializes in a grid network containing about 100 so-called Xputers. These entities are made up of a processor, a memory, and a switch. The Xputer grid forms a plane surface with switches, processors, and memories, each forming a logical parallel layer. I/O devices can be attached to any node.

The Rediflow concept currently assumes a configuration in which a number of processor-memory pairs are interconnected via a switching network. Combining such a pair with an appropriate packet switch for information transfer results in an Xputer. Hundreds of these nodes can be interconnected in a rectangular grid. For larger numbers of nodes, an interconnection topology with a lower worst-case delay seems to be attractive, according to Keller.

Transmission of computations

Thus, one can think of the Xputer grid as forming a plane surface, with the switches, processors, and memories each forming logically parallel layers. These layers need not be physically parallel. Interconnection exists only at the switch layer, while the memories in the memory layer have a combined globally addressable address space. If one Xputer needs to access the memory of another, it forms a request packet containing the address to be accessed. That packet is then routed within the switch layer to the Xputer containing the addressed location. A result packet is then formed and routed to the requesting Xputer. This request/return mechanism is integrated with the demand-driven mechanism of reduction evaluation so that remote triggering of function evaluations can take place.

How does an Xputer grid keep from saturation when all Xputers are sufficiently busy to prevent data migration between each other? The Utah team uses an analogy of fluid pressure to explain their solution to this problem. Obviously, the pressure on Xputers is continually changing. Accordingly, it is necessary to continually update each Xputer's sense of its environmental pressure. This is done through a sampling process in which the switch of each Xputer computes transmitted pressure as a function of transmitted pressures of its neighbors, Keller says. One heuristic device that seems to work well is to define the transmitted pressure to be 0 if the Xputer's internal pressure is below a certain threshold. This definition changes to 1 plus the minimum of the neighbor's transmitted pressures in other cases with an absolute maximum on the order of the diameter of the network. This permits small tasks to flow toward the least loaded node.

A rough overview of the internal organization of an Xputer assumes that pressure sampling information packets intermingled with other types of packets are sent through the switching layer. This assumption has been used in most of the Rediflow simulation results so far. However, it is also possible to dedicate a separate serial channel to the transmission of pressure information.

Thus, in operation, a fetch packet is issued by an Xputer that needs to get data from a location in another Xputer. It contains the data's address and a return address. When a fetch packet arrives at the other Xputer's in queue, the location is checked for containing valid data, and if so, a forward packet is created, which returns the value to the first Xputer. However, it may be that the data have not yet been produced, in which case the return address is reserved in the second Xputer until such time as the data is available. Also, if production of the data has not yet been demanded, it will be demanded at that time.

Because all result data have pre-allocated globally addressable locations, it is not necessary to use any form of token matching to get the data to their destinations. Thus, fast von Neumann-style memory is internally exploited in each Xputer. The use of addressing also permits routing tables to provide the shortest possible route to be chosen through the switching layer.

So far, the performance of the Rediflow architecture has been evaluated using simulation. As with most studies in their formative stages, evaluating speedups are done so that speedups are measured against a single processor with the same technological assumptions, architecture, and evaluation model as the multiprocessor. Certain needed improvements in the current model prevented Keller's team from challenging existing sequential processors for

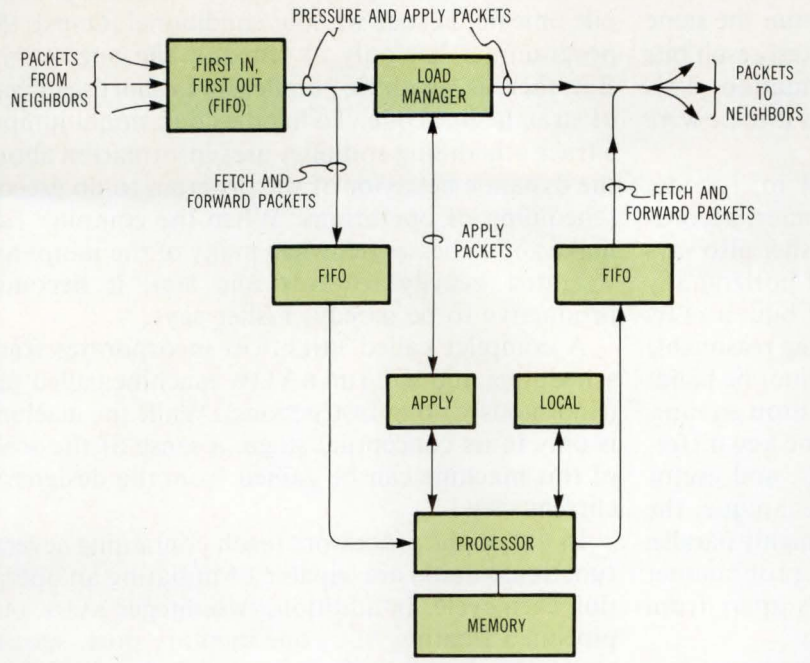

The packet addressing scheme contains many kinds of packets. A fetch packet, issued by an Xputer for obtaining data from another Xputer, is answered by a forward packet with the data contained inside. However, if the data has not yet been produced in the parallel processing operation, then the return address of the sending Xputer is reserved in the second Xputer until the data is made available. Because all resultant data is pre-allocated globally in addressable locations, no token-matching is needed. Each Xputer simply acts as a fast von Neumann machine.

applications with low degrees of concurrency. However, if the potential concurrency is high, Keller believes Rediflow can exploit it with a demonstrated speedup.

Two kinds of benchmarks have been run. One consists of playful programs that exhibit a single kind of activity, such as "divide and conquer." The other consists of more realistic applications that combine a number of activities, in the areas of simple database searching and updating, and correlative signal processing. Speedups over current architectures have been measured in the range of 1 to 8 for the realistic applications, with fewer than 32 Xputers, and of up to 30 with toy programs using as many as 128 Xputers. Memory space in the simulator is currently a principal limiting factor.

The load distribution techniques designed for Rediflow work well, and exploit the concept of locality. Typically, over 50 percent of the data packets, and 80 percent of the apply packets, traverse at most two paths. Usually, fewer than 15 percent of all operations performed need to communicate outside one Xputer. Keller also observes that the switches hypothesized for Rediflow do not seem to be a bottleneck under current technological assumptions. The next step for the Rediflow team is to concurrently evaluate logic programs.

Programming mutiprocessors, in general, is a difficult task because the programmer has to locate the parallelism in the program and match it to the parallel structure of the hardware. Researchers at the Computer Sciences Department of Yale University (New Haven, Conn) have developed a Very Long Instruction Word (VLIW) mechanism with a BULLDOG compiler that can produce highly parallel code from a range of ordinary, sequential programs. The VLIW machines anticipate instruction words over 1000 bits long.

Deciding instruction length

VLIWs are heterogeneous multiprocessor architectures characterized by the large degree of timing and resource control in their programs. As a result, each instruction for a VLIW contains many bits, perhaps in the thousands. To picture this architecture, imagine 32 identical Reduced Instruction Set Computer (RISC) machines, all connected to a memory system and capable of communicating over an interconnect. The RISC machines are the processors, and are not required or expected to resemble each other.

What differentiates a VLIW machine from regular machines are the instructions necessary to control it. Each long instruction contains operation fields to control each of the individual processors. Also, the instructions are fetched in a single flow and all the processors do their individual operations. These operations differ for the various processors in that no coherence is possible as with vector machines. After an instruction is executed, the next instruction is chosen and fetched.

The instruction word completely controls all communication among the processors. Thus, a piece of data is taken from a known location and placed into another known location, using known resources for a known amount of time. Packets are without destinations, and scheduling of a data transfer is not hardware-dependent. Data transfers and their timings are choreographed in the code.

Instead of a central program store, each processor fetches that portion of the instruction word

relevant to it from its own program store. But on any given cycle, all processors fetch from the same address. A central control unit collects test result bits and generates the next address information. This operation looks to the code generator as if there were one VLIW.

VLIW code resembles a horizontal microcode, except that it is more parallel, and cleaner, according to project leader Joseph Fisher. Fisher also says that although there have been a lot of horizontally microcoded machines, no one has ever built a VLIW because there was no chance of producing reasonable code for them until recently. VLIWs cannot be hand-coded since too many unrelated execution streams are going on in parallel. There are some key differences between VLIWs, multiprocessors, and vector machines. Using a trace scheduling technique, the Yale team developed large amounts of highly parallel code for them without a great deal of programmer effort. Fisher believes this sets VLIWs apart from multiprocessors and vector machines.

Trace scheduling is a complex procedure to explain and to implement, yet the idea is simple. When the control flow of a segment of code is known at compile time (ie, in code without conditional jumps), the programmer has only to schedule the operations. But, there is very little parallelism in short segments of straight-line code. To handle conditional jumps, a trace scheduling compiler uses information about the dynamic behavior of the program to do greedy scheduling of operations. When the compiler can make good guesses (ie, when many of the jumps are weighted heavily towards one leg), it becomes productive to be greedy, Fisher says.

A compiler called BULLDOG incorporates trace scheduling and will run a VLIW machine called ELI (Enormously Long Instructions). While the machine is only in its conceptual stage, a sense of the scale of this machine can be gained from the designers' current view.

In ELI, eight processors (each containing several functional units) are capable of initiating an operation each cycle. In addition, two integer ALUs, one pipelined floating ALU, one memory port, several register banks, and a limited crossbar for all of these to talk to each other are needed. If built to a 200-ns

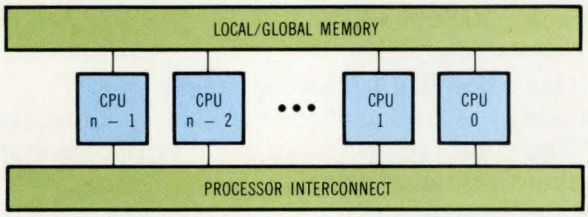

	MULTIPROCESSORS	VECTOR MACHINES	VLIWs
FLOW OF CONTROL	EACH PROCESSOR HAS ITS OWN CONTROL FLOW. PROCESSORS MUST WAIT TO SYNCHRONIZE ON DATA AVAILABILITY.	ONE INSTRUCTION AT A TIME von NEUMANN STYLE, ALL PROCESSORS DOING THE SAME OPERATION. MASK BITS MAY DISABLE A SUBSET OF PROCESSORS.	ALL PROCESSORS FETCH INSTRUCTIONS FROM THE SAME NEXT ADDRESS, BUT INSTRUCTIONS DIFFER FROM PROCESSOR TO PROCESSOR. VON NEUMANN-STYLE CONTROL FLOW WITH EXTREMELY LONG INSTRUCTIONS.
INTERPROCESSOR COMMUNICATION BANDWIDTH	HARDWARE SCHEDULING OF DATA PATH RESOURCES, DATA PACKETS WITH DESTINATION ADDRESSES, AND BUFFERING OF UNUSABLE DATA.	SUBVECTORS MUST BE MOVED IN PATTERNS FITTING THE REGULARITY OF THE INTERCONNECT AND ALGORITHM. OTHERWISE IT IS LIKE A UNIPROCESSOR.	INDIVIDUAL MOVEMENTS OF DATA COMPLETELY SPECIFIED AT COMPILE TIME. NO NEED FOR ADDRESSES, RUNTIME RESOURCE SCHEDULING, OR IMPLICIT BUFFERING.
MEMORY BANDWIDTH	WHEN POSSIBLE, REFERENCES ARE LOCALIZED TO A GIVEN PROCESSOR. OTHERWISE EXPENSIVE AND SLOW GLOBAL MEMORY SWITCH TO ALLOW PROCESSOR/MEMORY CROSSBAR.	VECTORIZED REFERENCES ALLOW FULL ACCESS TO ALL BANKS. UNVECTORIZED REFERENCES LIKE A UNIPROCESSOR.	BANKS MUST BE PREDICTED AT COMPILE TIME TO GET MAXIMAL BANDWIDTH FOR SCALARS AND AGGREGATES, WHICH MAY BE MIXED IN A SINGLE INSTRUCTION. WHEN THE AGGREGATE BANK IS UNPREDICTABLE, IT IS LIKE A UNIPROCESSOR.
PROGRAMMING REQUIREMENTS	CODE MUST BE BROKEN INTO RELATIVELY INDEPENDENT TASKS THAT MINIMIZE COMMUNICATIONS AND SYNCHRONIZATION.	CODE MUST BE EXPRESSED AS REGULAR OPERATIONS ON AGGREGATES. MATCHUP MUST BE MADE BETWEEN INHERENT REGULARITY OF CODE AND HARDWARE.	NOT LIKELY TO BE HAND-CODED. COMPILER MUST DO GREEDY SCHEDULING (EG, TRACE SCHEDULING) AND ELABORATE RUNTIME ANALYSIS TO CHOREOGRAPH INDIVIDUAL DATA MOVEMENTS AND OTHER RESOURCE SELECTIONS.

Yale's computer is a Very Long Instruction Word machine that uses as many as 1000 bits per instruction. To give the designer an idea of the task's magnitude, three basic types of machines for a generic parallel processor configuration are compared.

THREE HUNDRED FORTY ONE MILLION FLOATING POINT OPERATIONS PER SECOND.

THE FPS-164/MAX.

Today's scientific and engineering problems increasingly call for more complex models and higher resolution, which leads to calculations on very large matrices. Yet, until now, cost of a supercomputer with the speed and accuracy suited to these problems has been out of reach for many.

Now, there's the FPS-164/MAX — a special-purpose scientific supercomputer that matches the likes of CRAY, CYBER and others in commonly used matrix operations — at a fraction of the cost.

The FPS-164/MAX is fast.

With peak performance rated from 33 to 341 million floating-point operations per second, depending on configuration, and up to 7 Mwords of 64-bit memory available to the user, the new FPS-164/MAX gives you all the speed and accuracy you need to make those matrix computations manageable.

The FPS-164/MAX configuration is able to compute up to 124 vector operations at one time, allowing a fully-configured 164/MAX to factor a 1,000 by 1,000 matrix in about 1 second, multiply two 10,000 by 10,000 matrices in less than two hours.

FPS-164/MAX Specifications

Peak Capacity (MFLOPS)	341
Number of Arithmetic Pipelines	63
Number of Independent Processors	1
Vector Register Capacity	2K x 124
Main Memory Capacity	7 MWords
Disk Subsystem Capacity	3 Gbytes
Word Size	64 bits +8-bit error correction
Precision	15.3 decimal digits
Dynamic range	2.8×10^{-309} to 9.0×10^{307}
Cooling Method	Forced Air
Footprint	~20 ft²
Weight	1800 lbs. including power supplies

The FPS-164/MAX is powerful.

A parallel-pipelined machine designed to run FORTRAN at high speed, the FPS-164/MAX has all the scalar capability of our original FPS-164. We've just added a lot more power, with multiple special processing units which amplify the vector processing capability of the original FPS-164 by up to 31 times.

The FPS-164/MAX is cost-effective.

In structural analysis, computational chemistry and physics, fluid flow analysis, electromagnetic modeling, or any application requiring fast handling of large matrices, the FPS-164/MAX offers unparalleled cost efficiency. In fact, it can run certain key matrix computations as fast or faster than supercomputers costing over 10 times as much.

Whether you're looking to upgrade your existing FPS-164 — or searching for a complete new system — you won't find supercomputer performance for one million dollars or less anywhere else.

What's more, the FPS-164/MAX is backed by the considerable resources of Floating Point Systems. With 21 field service offices world-wide, full remote diagnostic capabilities, and a record of product quality and reliability second to none, you can be sure the FPS-164/MAX will be up, running, and ready to meet your problem solving needs.

For complete information and specifications, call toll free, 1-800-547-1445.

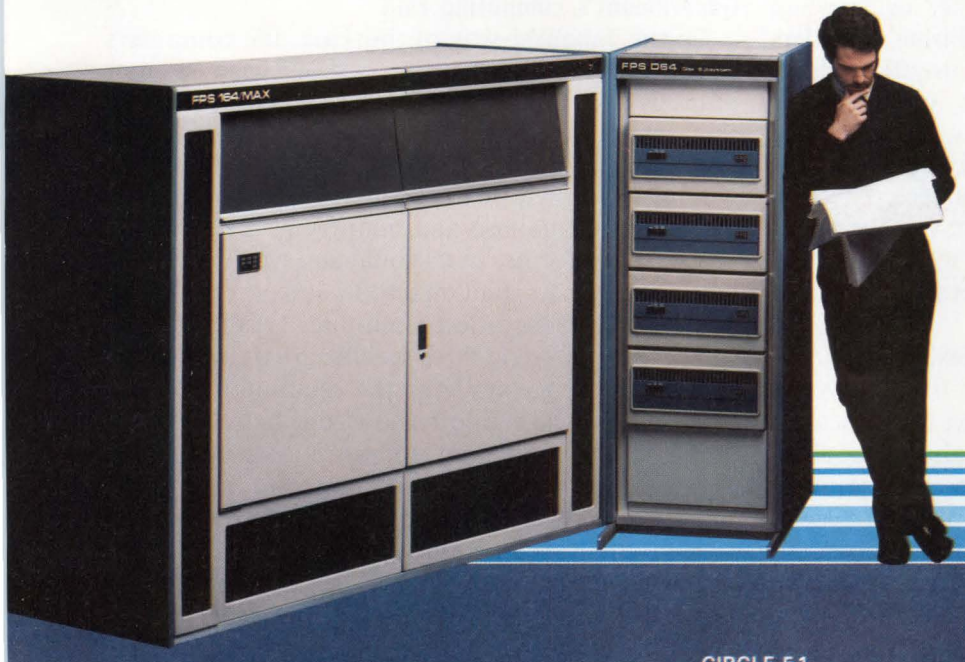

FPS® FLOATING POINT SYSTEMS, INC.

P.O. Box 23489
Portland, OR 97223
(503) 641-3151
TLX: 360470 FLOATPOIN BEAV

CIRCLE 51

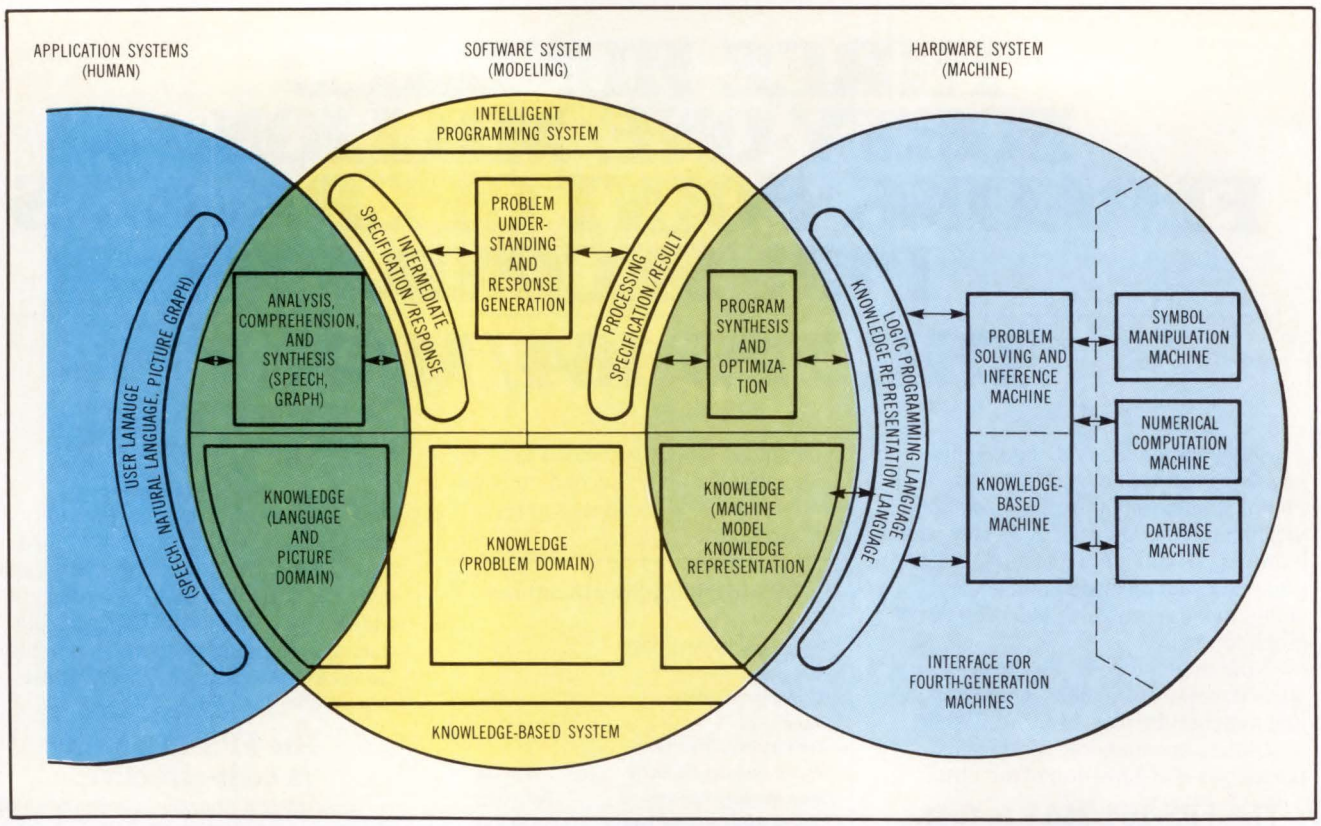

Japan's ambitious Fifth-Generation Computer Project encompasses all possible scientific and engineering principles. To incorporate many human requirements into hardware, many human-related features must be modeled within a knowledge base. Three specific areas of research are being pursued: parallel inference systems, relational data bases for knowledge-based systems, and personal sequential inference (PSI) systems.

cycle time, a total of four operations per processor by eight processors by five million cycles/s, or a possible 160 RISC MIPS, can be achieved using technology that is far from the fastest available today.

The memory system is a general alignment network, allowing the delivery of any memory word to any processor at a rate of one word per cycle. This alignment network is locally bypassed at each processor, allowing eight references per cycle when banks can be statically determined. Fisher considers processor interconnection the most unsettled design parameter. As VLSI changes the cost balance between interconnect and function, system cost/performance has to be evaluated in terms of the effective use of the provided interconnect. Through simulations, Fisher's team found that various topologies do not seem to differ greatly in their performance; what matters is the overall functional unit-to-interconnect ratio.

Currently, there are four half-duplex, bidirectional lines to/from each processor. With the processors arranged in a circle, two lines connect to the nearest neighbors. The other lines each go about one third of the way around, one clockwise, one counterclockwise. Currently, the compiler generates code for an ELI with a parameterized description. Fisher's team is picking values of the parameters to tune the architecture to a suite of scientific programs.

While the next generation of computers will have an intimate relationship between hardware and software, no other type of machine is more intertwined than a knowledge-intensive computer that is based on artificial intelligence concepts. This has become blatantly obvious just after the Japanese announced their version of fifth-generation computing.

Government's computing role

To the Japanese way of thinking, the computers of the next decade will be increasingly used for nonnumeric data processing, such as symbol manipulation, and applied AI. Moreover, conventional system applications, such as scientific calculations, will be performed by the evolving supercomputers, and current database and mainframe systems will be improved for use in national and worldwide network systems. Current directions and eventual goals of the Fifth-Generation Computer System (FGCS) project are placed on three specific areas of research: parallel inference systems; the relational database research that will support knowledge-based systems; and personal sequential inference (PSI) systems.

According to FGCS promoting committee chairman, Tohru Moto-oka, the design and development of a problem-solving machine is contingent on the development of hardware that supports the kernel language used for writing the inference-based

algorithms. The hardware mechanism used for inferencing will eventually be integrated with the hardware used for dealing with a knowledge base (the goal of the knowledge-base research), and intelligent interfaces (the goal of the high performance interface equipment research). The final result should be a prototype of a Japanese version of the FGCS.

The Japanese measure performance of the inference engine in logical inferences per second (LIPS). One inference is roughly equivalent to one rule firing in a rule-based expert system. An inference machine for the 1990s should have a maximum performance of 100 MLIPS to 1 GLIPS. In terms of conventional von Neumann architecture, 1 LIPS is roughly equivalent to the execution of 100 to 1000 instructions per second, so that an inference engine running at 1 GLIPS, in some sense, has the performance of a conventional machine operating at 10^{11} to 10^{12} instructions per second.

At those speeds, a purely sequential implementation will not do; parallel execution is the key to realizing this level of performance. Consequently, a primary focus of the FGCS project is to research architectures that are suited both to highly parallel inferencing and to manipulation of abstract data types. The FGCS research is working toward a target goal of 1000 processing elements. Among the tools expected to be developed in the course of the research are VLSI computer aided design tools to support the design of such hardware.

Intimately tied to the architecture are rules and assertions to process knowledge information. Existing AI technology has been based primarily on Lisp, which has a sequential, functional character. To the Japanese, Prolog seems to be a more appropriate language because its logic programming character provides for greater power of expression and for several different types of parallel execution. Consequently, a Prolog-like language is being designed for the interface between FGCS hardware and software. The FGCS hardware architecture should provide direct support for Prolog-like inferencing. In the initial phase of the project, a sequential inference language, Kernel Language 0, or KL-0, has been the subject of study and the corresponding architecture is the sequential inference machine (SIM). In the second phase of research, parallelism will be a major focus that will appear in the successor language KL-1 and in the successor architecture known as a parallel inference machine, according to Moto-oka.

The first implementation of a SIM is scheduled to be a medium performance machine known as a PSI machine. The next phase will be to extend this machine to give it several times the performance of

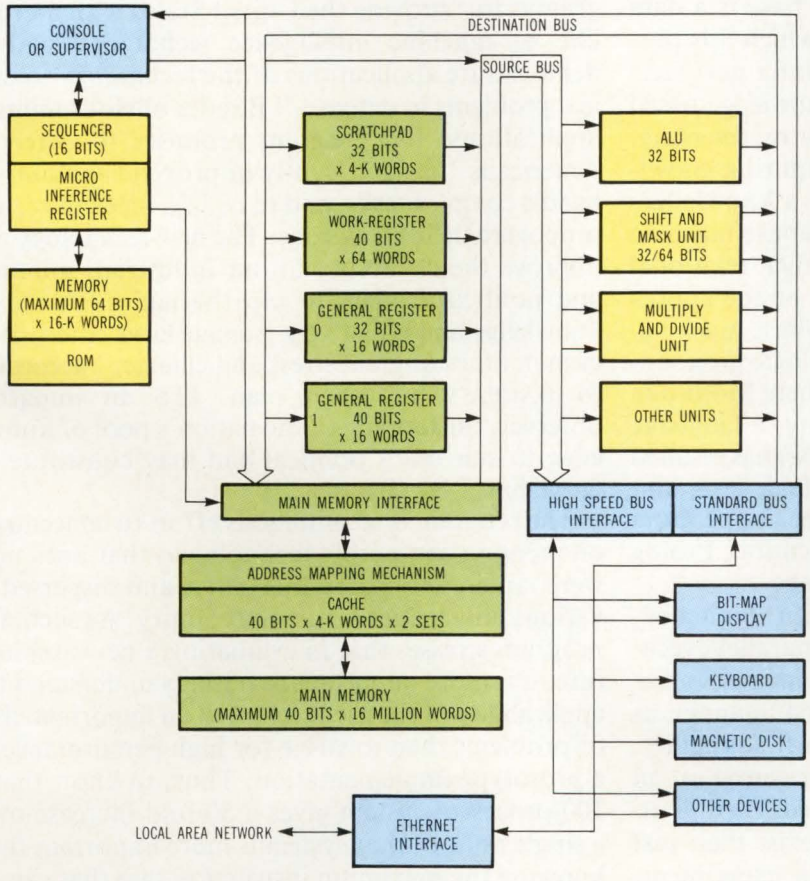

One early product of Japan's Fifth-Generation Computer Project is a PSI machine. At first glance, the machine's architecture closely resembles any of a dozen architectures currently used in workstations. The difference is in the software for, as opposed to writing instructions for moving data and for processing, the programmer instead will write logical inferences. These inferences can be processed much faster and are more powerful since 1 logical inference per second (LIPS) roughly equals 100 to 1000 instructions per second.

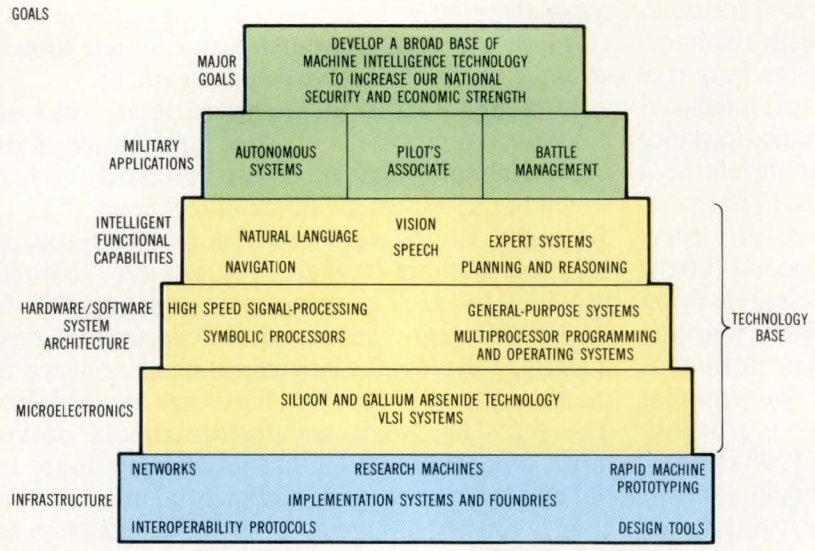

The Defense Advanced Projects Research Agency's (DARPA's) Strategic Computing Program is one of the more ambitious programs that the agency has undertaken. The goals are to build an autonomous land vehicle for the U.S. Army, a smart pilot's assistant for the U.S. Air Force, and an intelligent battle management system for the U.S. Navy. These projects should be completed by 1994. To meet these goals, the agency can pool the resources of industry and the nation's universities that will provide the essential elements of the project's infrastructure.

the forerunner PSI. The PSI is intended to be a powerful personal workstation that can be a major tool for subsequent phases of the project. Its performance is roughly 20 to 30 KLIPS. The second-phase machine should reach 100 KLIPS to 1 MLIPS (see *Computer Design*, July 1983, p 58).

The PSI is designed to infer knowledge bases. A knowledge base is an extension of a data base. If one views a data base as a collection of facts and relations among facts, a knowledge base is a data base plus a collection of rules from which it is possible to manipulate the data base to infer new facts and relations not explicitly stored. Knowledge-based management is built on techniques for representing, storing, retrieving, updating, and acquiring knowledge. Key to the implementation of a knowledge-based machine is the underlying database machine architecture. Consequently, a parallel relational database machine architecture must be the appropriate starting point for a knowledge-base machine.

Tied to knowledge bases is parallel logic programming. At the University of Tokyo, where Moto-oka is a professor of electrical engineering, a language known as Paralog (from parallel Prolog) has resulted from a recent study on parallelism in logic programming. Except for the inclusion of special operators that take advantage of sequential execution, Prolog could be viewed as a parallel language.

In Paralog, such operators have been eliminated, and parallelism is achieved through parallel execution of the OR alternatives. Goal statements, therefore, are executed in a breadth-first manner as opposed to Prolog's sequential depth-first manner. The language has been tested and measured on an experimental dataflow machine called Topstar II. More opportunities for parallelism exist than just the OR parallelism captured in this experiment, according to Moto-oka. Partially in response to Japan's vision of FGCS, and partially due to the synthesis on the part of the Department of Defense Advanced Research Projects Agency's (DARPA) work in computing technology, the United States military has formed a program of its own known as Strategic Computing.

In the name of defense

The overall goal of the Strategic Computing Program is to "provide the United States with a broad line of machine intelligence technology and to demonstrate applications of the technology to critical problems in defense." Besides obvious military applications, the program promises to offer all Americans "a significantly improved capability to handle complex tasks, and to codify, mechanize, and propogate their knowledge. The new technology will improve the capability of our industrial, military, and political leaders to tap the nation's pool of knowledge and effectively manage large enterprises, even in times of great stress and change," according to DARPA's published plan. [To an impartial observer, the tapping of the nation's pool of knowledge to suit one's political end may constitute an Orwellian "Big Brother."]

The program is an ambitious effort to concentrate on merging computing technologies that until now were largely sponsored by DARPA and dispersed at various universities across the country. As such, the program stresses that in evaluating a new architecture, it is more important to initially understand the applicability of the architecture to an important class of problems than to strive for high performance in a prototype implementation. Thus, to know that a 100-processor system gives a 50-fold increase over a single node of that system is more important than knowing the maximum instruction rate that can be executed, or knowing the exact instruction rate

achieved with prototype hardware. Once a prototype machine is proved promising, DARPA intends to have higher performance versions built by using faster components and scaling the entire system to have more processors.

The program will develop and evaluate new architectures in the areas of signal processing, symbolic processing, and multifunction machines. Generally, several prototype systems will be developed in the early phase of the program. An evaluation phase will permit comparison of different architectural approaches. DARPA will then select the most successful and will continue to develop high performance versions.

Signal processing is an important class of applications that involves taking realtime data from a sensor and performing a series of operations on each data element. These operations might involve transformations such as fast Fourier transform correlations or filtering, and are dominated by performing multiplications and additions. High data and computation rates in excess of one billion operations per second are needed. Military signal processing applications include processing data from radar, sonar, infrared sensors, images, and speech.

The exploitation of parallelism in signal processing will be based on the use of computational arrays such as systolic arrays. Here, many simple, highly regular processing elements pump data from cell to cell in a wave-like motion to perform the successive operations on each element of data. An architecture based on this concept will be developed, with the goal of building a system that can execute one billion or more operations per second by 1986. Other concepts that exploit signal processing data regularity will also be investigated. By the end of the decade, the goal is to develop a system capable of 1 trillion operations per second.

The software support and programming languages for the signal processing system will be developed in parallel with the hardware. Most of the initial programming for the prototype system will be done at the microcode level. The requirements for the operating system, the programming languages, and the programming environments will be developed as experience is gained using the prototype systems. Because symbolic processing deals with nonnumeric objects, relationships between these objects, and the ability to infer or deduce new information with the aid of programs, they can be referred to as "reasoning" machines. Examples of symbolic computation in military applications include searching and comparing complex structures (eg, partial pattern matching).

Applications that make extensive use of symbolic computing include vision systems that can tell what is in a scene; natural language systems that can "understand" the meaning of a sentence in English; speech understanding systems that can recognize spoken words; and planning systems that can provide intelligent advice to a decision maker. Most programs that perform symbolic processing are currently written in Lisp. Lisp machines have been available commercially and offer computing rates in excess of 1 MIPS. Further development of these conventional uniprocessor Lisp machines will take place under the technology infrastructure portion of the program. An ultimate performance improvement of about 50 times the current level can be achieved with these conventional techniques and the use of advanced technology.

Current applications in areas such as vision now require about three orders of magnitude more processing than is now available. DARPA visionaries see future algorithms and applications requiring even more computing power. The symbolic processors of the future may well be a collection of special components that are interconnected via a general-purpose host computer or by high speed networks.

By the end of the decade, the goal is to develop a system capable of 1 trillion operations per second.

DARPA has identified several of these components, which are based on software systems developed for applications in vision, natural language, expert systems, and speech. Thus, as much as four orders of magnitude speedup may be available by taking advantage of parallelism. Some examples include a semantic memory subsystem, a signal-to-symbol transducer, a production rule subsystem, a fusion subsystem, an inferencing subsystem, and a search subsystem.

A semantic memory subsystem is used to represent knowledge that relates concepts in natural language, speech understanding, and planning domains. A signal-to-symbol transducer is used to make the initial step in extracting meaning from low level signal processing computations (eg, phonetic classification or object identification from boundary information).

A production rule subsystem combines knowledge and procedures for problem solving. A system aboard the carrier USS Carl Vinson already uses this approach. A fusion subsystem permits multiple sources of information to share their knowledge. It is used to fuse information in tasks such as battle management. An inferencing subsystem uses first-order formal logic to perform reasoning and theorem

Multiprocessing and large memories are basic to supercomputing

"Multiprocessing offers the next significant source for performance improvements in supercomputers," according to Lloyd M. Thorndyke, president and CEO of ETA Systems, Inc, a "spin-off" of Control Data Corp. Thorndyke sees future architectural advances primarily in the areas of multiprocessing and large memories.

Thorndyke, former senior vice president of technology development at Control Data, was instrumental in the design of the Cyber 203, Cyber 205, and in the current effort on the Cyber 2XX. He feels that multiprocessing and large memories must be incorporated in the next generations of supercomputers to sustain the demand for computing requirements.

To this end, ETA's first system, to be released in 1986, will use up to eight processors, each of which will have three to five times more performance than the Cyber 205. This machine, whose internal name is GF-10, is targeted for a performance threshold of 10 billion floating point operations per second (10 GFLOPS), and a memory capacity in excess of 295 million words. "Supercomputing within the next 10 years must attain multiple GFLOPS rates and provide this power at reasonable costs to a wider audience," he says.

Thorndyke sees a second trend, one toward memories as large as the technology permits. Some of the memories being designed for future implementation are larger than modern disk drives. "Future multiprocessing architectures will offer very high performance and will require memory capacities that allow larger three-dimensional processing and simulation," he says. "We can expect the trend toward memory capacity to continue to grow by orders of magnitude. These gigantic memories will require virtual addressing to effectively manage memory for the programmer."

Thordndyke says that his company prefers to evolve architectural changes rather than to introduce drastic departures from the current architecture. However, he adds that the achievement of supercomputer performance requires that the technology of all of the system's elements be simultaneously advanced to meet the dual goals of performance and balance. Incompatibility is costly for both vendors and users. Thus, easy access by the user, whether it is through a personal computer or a mainframe, is a necessity. The supercomputer must also tie into all major computing networks.

"Those future manufacturers that have the foresight to adapt and use existing semiconductor technologies, preferably choosing those that are in close proximity to the mainstream of standard products, will help in reaching the goal of making the supercomputer reachable by wider groups through better cost control and reliability," Thorndyke concludes.

proving. A search subsystem explores numerous hypotheses, pruning these to determine likely candidates for further symbolic processing.

Three-part program

The initial part of the Strategic Computing Program consists of three phases. Phase I concentrates on architecture design, simulation, algorithm analysis, and benchmark development for promising architectural ideas such as those described above. It also includes the development and initial evaluation of the unique integrated components necessary for the implementation of these architectures. The design of concurrent Lisp-like languages for programming these machines are also to be addressed. Existing high performance scientific computers, such as the Cray-1, CDC 205, Denelcor HEP, and the S-1, are scheduled to be benchmarked using a portable Lisp computer to determine their relative abilities to handle symbolic computation.

Phase II of symbolic processing machines will engineer full-scale prototype versions of selected architectures, supporting these hardware developments with extensive diagnostic and compilation tools. The goal of this phase is implementation of a specific target problem on each of the selected architectures for benchmarking purposes.

Phase III will integrate developments of the signal processing, symbolic and multifunction development efforts into a composite system that can address a significant problem domain. Such a system for the control of an autonomous vehicle might include a high performance, vision-processing front end based on the computational array technology; a signal-to-signal transformer for classifying objects; a fusion subsystem for integrating information from multiple sources; an inferencing engineer for reasoning and top-level control; and a multifunction processor for controlling the manipulator effectors. This phase will also pursue higher performance versions of selected machines.

The third leg of the program's architecture phase is a multifunction machine that can execute a wider range of different types of computations than the specialized signal processors and symbolic processors. The multifunction machines might possibly exhibit lower performance in the specialized machine's application domain. These machines achieve high performance with parallelism. DARPA aims to develop machines of this class having 1000 processors. The processing elements in a multifunction machine would typically be general-purpose processors or computers. These elements communicate either through shared storage or networks with such interconnection strategies as rings, trees, Boolean n-cubes, perfect shuffle networks, lattices or meshes.

Six to eight prototype multifunction systems will be developed, based on custom VLSI chips, commercial microprocessor chips, or commercial processors. These will be benchmarked to determine how different hardware architectures and programming

strategies scale in performance. Subsequently, two or three such systems will be selected in this evaluation process for continued development for advanced technology versions and production quality software.

Central to this program is the development of programming models and methods that will permit the convenient development of new classes of algorithms. These algorithms will contain high levels of concurrency. According to Lynn Conway, DARPA's assistant director for Strategic Computing, the way in which concurrency manifests itself in program structures results from the linguistic control method of the programming language in which the program is written. Examples of control models that will be investigated correspond to those being pursued by academic circles and include control-driven, data-driven and demand-driven styles. Control-driven concurrent programming models are already evolving from existing programming languages (eg, concurrent Pascal and parallel Lisp). In this model, program actions are sequenced by explicit control mechanisms such as CALL, JUMP, or PARBEGIN.

In the data-driven model, program actions are driven into activity by the arrival of the requisite operand set. The advantage of this style is that concurrency can often be specified implicitly. The demand-driven model is based on the propagation of demands for results to invoke actions. This style has been successfully employed for parallel evaluation of Lisp code.

In this scheme, concurrent demands are propagated for argument evaluation of Lisp functions. Conway predicts that new or possibly composite models such as concurrent object-oriented programming will surface, but that advances in each area will provide highly concurrent, program-based solutions for many application areas.

Exploiting concurrent programs

An important part of this project will be the implementation of new concurrent programming languages that exploit these models. The language development will need to be coupled with programming environment tools, and compatible hardware and operating system software. This development will provide the necessary computational tools to support application studies aimed at the creation of highly parallel application programs that can take advantage of the large levels of concurrency provided by multifunction machine prototypes. The long-term goal of this research is ultrafast, cost-effective demonstrations of important application areas such as database access, system simulation, and physical modeling.

Given a particular machine and a particular parallel program, the remaining issue is how the program should be mapped onto the physical resources in order to permit efficient exploitation of concurrency. This resource allocation problem is one of the key technical issues addressed by the Strategic Computing Program. Static allocation and dynamic allocation are the two most common techniques used to solve this problem.

In a static mapping strategy, the concurrency structure of the program is evaluated with respect to the topology of the physical machine. The compiler can then specify load modules for the physical nodes of the target machine. This method is simpler than the dynamic method, but needs to be developed for each of the architectures being pursued. If there are many components, it is likely that component failures will occur. With the static allocation mechanism, recompilation of the program is necessary for the current machine configuration.

In a dynamic allocation strategy, it is still important for the compiler to do some of the allocation task collection, but the compiler output is not in the form of specific load modules. Dynamic strategies allow the loader to define the final physical target of a compiled module based on hardware availability. An important by-product of this program will be the development and implementation of both static and dynamic allocation strategies. But, Conway expects that acceptable static allocation methods will precede the more sophisticated dynamic strategies.

DARPA has selected three application areas for implementing the new computing technologies: an autonomous land vehicle, a pilot's associate, and an intelligent battle management system. The program begins by focusing on demanding military applications that require machine intelligence technology. The applications generate requirements for functions such as vision, speech, natural language and expert system technology, and provide an experimental environment for synergistic interactions among developers of the new technology. The intelligent functions will be implemented in advanced architectures and fabricated in microelectronics to meet application performance requirements. Thus, the applications serve to focus and stimulate the creation of the technology base. The applications also provide a ready environment for the demonstration of prototype systems as the technology compartments successfully evolve.

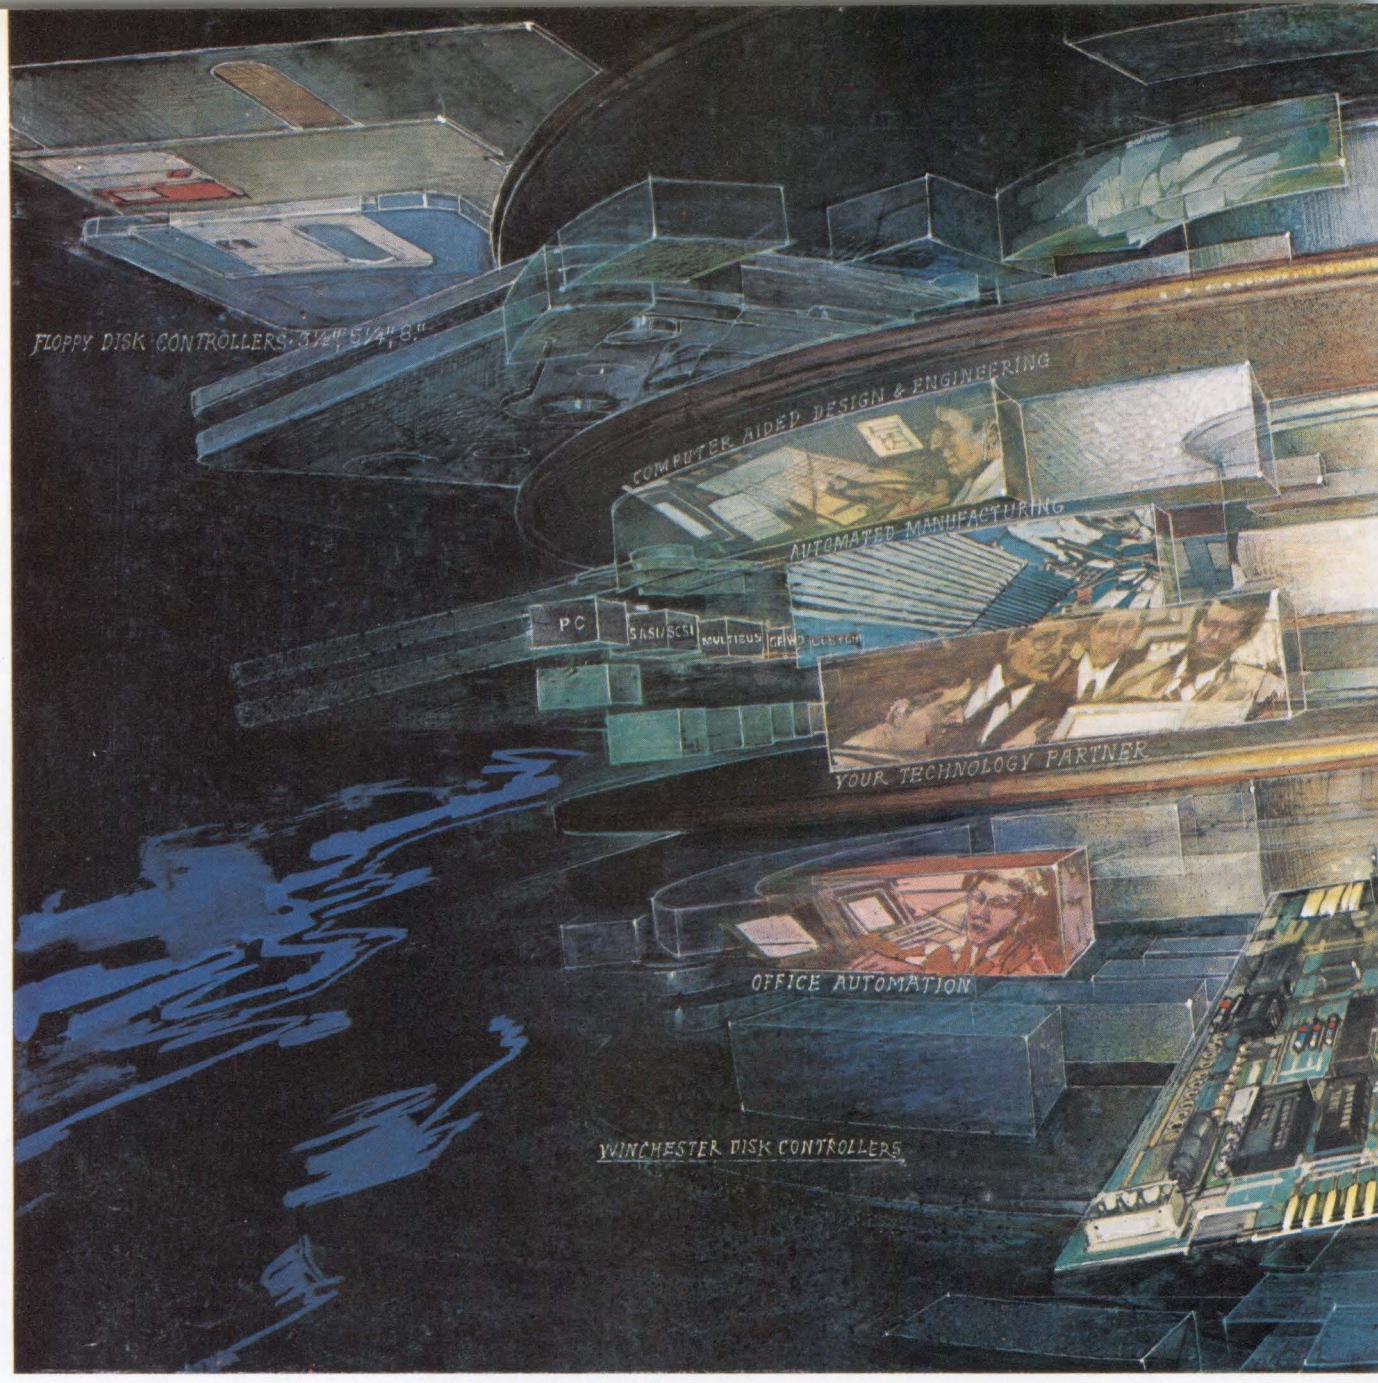

WHAT'S IN STORE FOR DISK STORAGE?

THE FUTURE...

Your future is ensured when you make Western Digital your partner in storage management technology. We've led the way with non-stop innovation, from the first single chip floppy disk controller nearly a decade ago to the broadest offering of chip and board-level floppy, Winchester and tape controllers today.

Systems in Silicon.

What makes us unique are our extensive VLSI capabilities. Designing and manufacturing our own proprietary chips enables us to 1) pack more performance into our controllers than is possible using general purpose LSI, 2) continuously integrate more and more functionality into fewer and fewer devices, 3) and provide you with an unending path to lower cost and higher performance as we ride the experience curve.

Chip-to-board synergy.

Solutions are what we offer systems builders. Having us build you a board-level disk controller

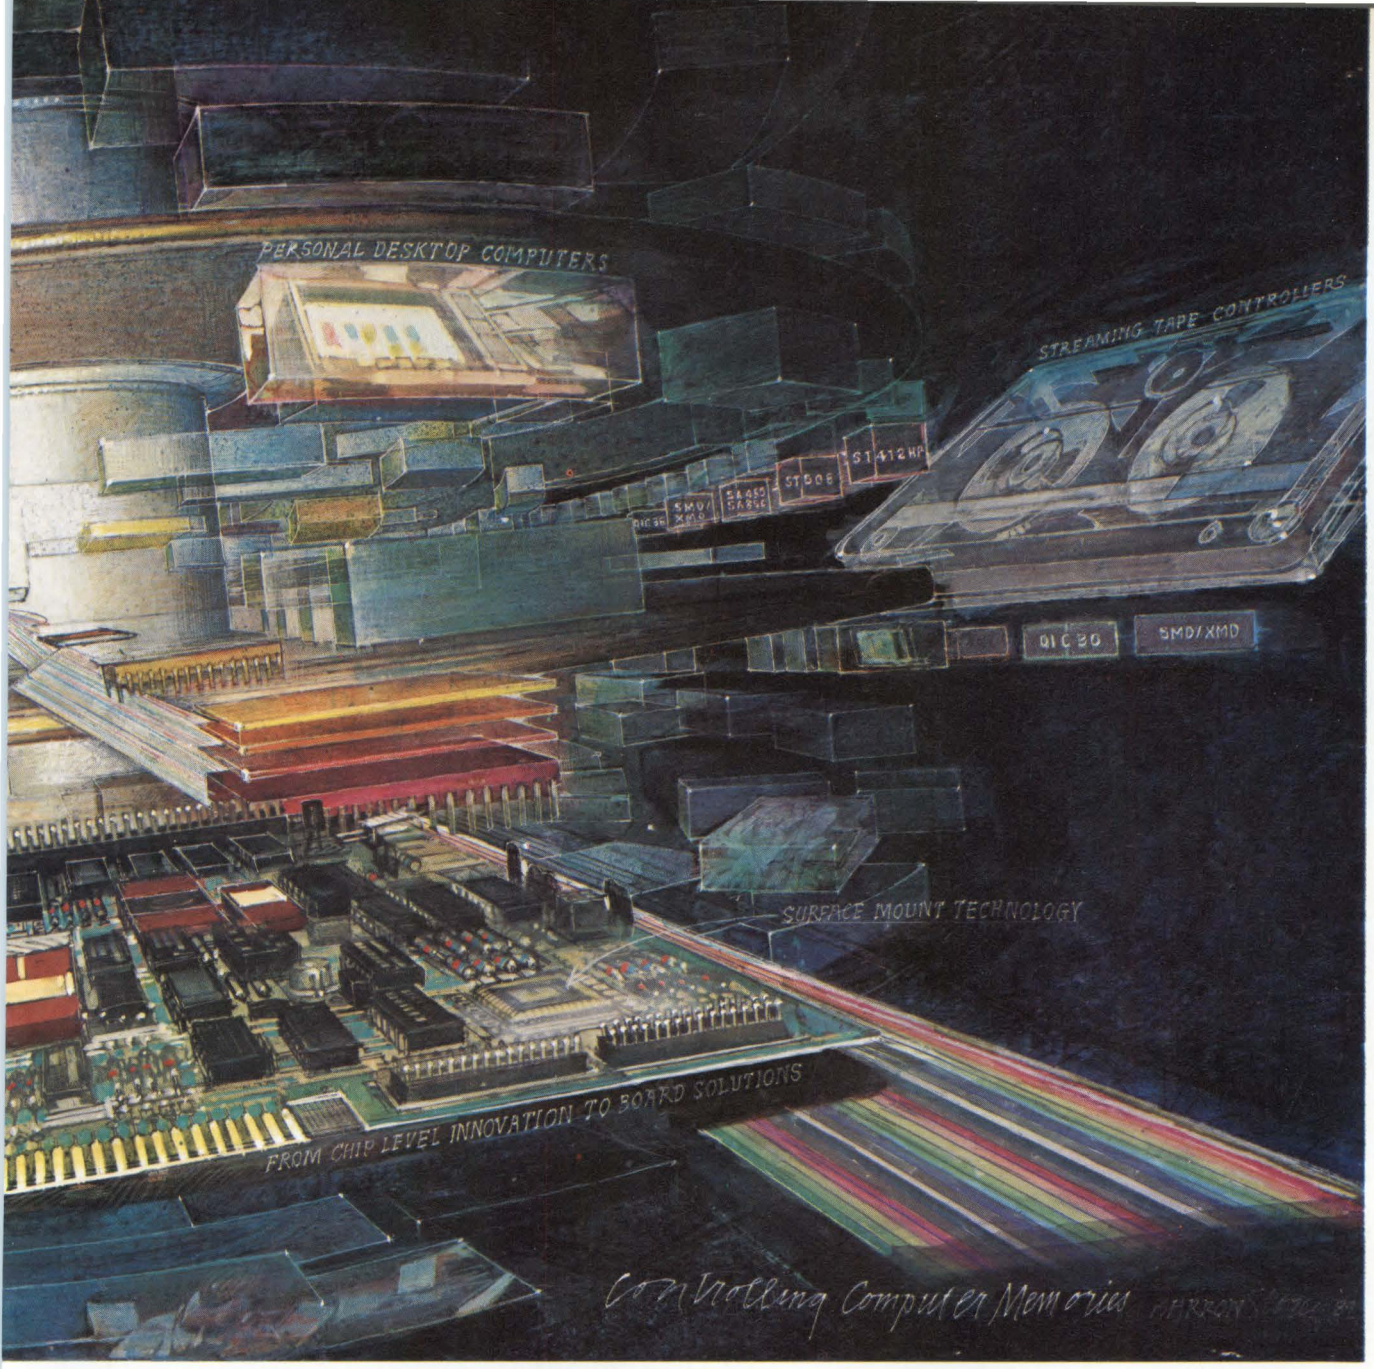

based on our chips does more than get you to market more quickly. It enables you to make us your technology partner at the systems integration level.

Whether you choose one of our standard boards, with more than a dozen combinations of host and drive interfaces, or have us design and build a custom, proprietary version for your system's special needs, our engineers work as a virtual extension of your own engineering capabilities.

Leading edge manufacturing.

To meet your high volume needs, we've invested in new, state-of-the-art automated board manufacturing and test facilities in the U.S. and Europe.

To keep you competitive, we're constantly integrating more functionality onto our board-level products, driving down cost while we boost performance. Our investments in surface-mount technology, and commitment to stay at the leading edge of this revolutionary approach to board manufacturing, will accelerate the integration process, enabling us to pack dramatically more into dramatically less space.

Take control of the future.

More leading manufacturers of personal computers and office automation systems buy storage management controllers from Western Digital than from any other company. Make us your source for disk and tape controllers and you get more than high technology products. You get a corporate commitment to do all we can to see you succeed. Take control of the future. Call our Controller Hotline, 714/863-7827. And ensure *your* success.

...STARTS HERE.

For the complete story of our storage management capabilities and a poster-size reproduction of the illustration above, send your business card to Western Digital, SM Literature, 2445 McCabe Way, Irvine, CA 92714.

WESTERN DIGITAL
CORPORATION

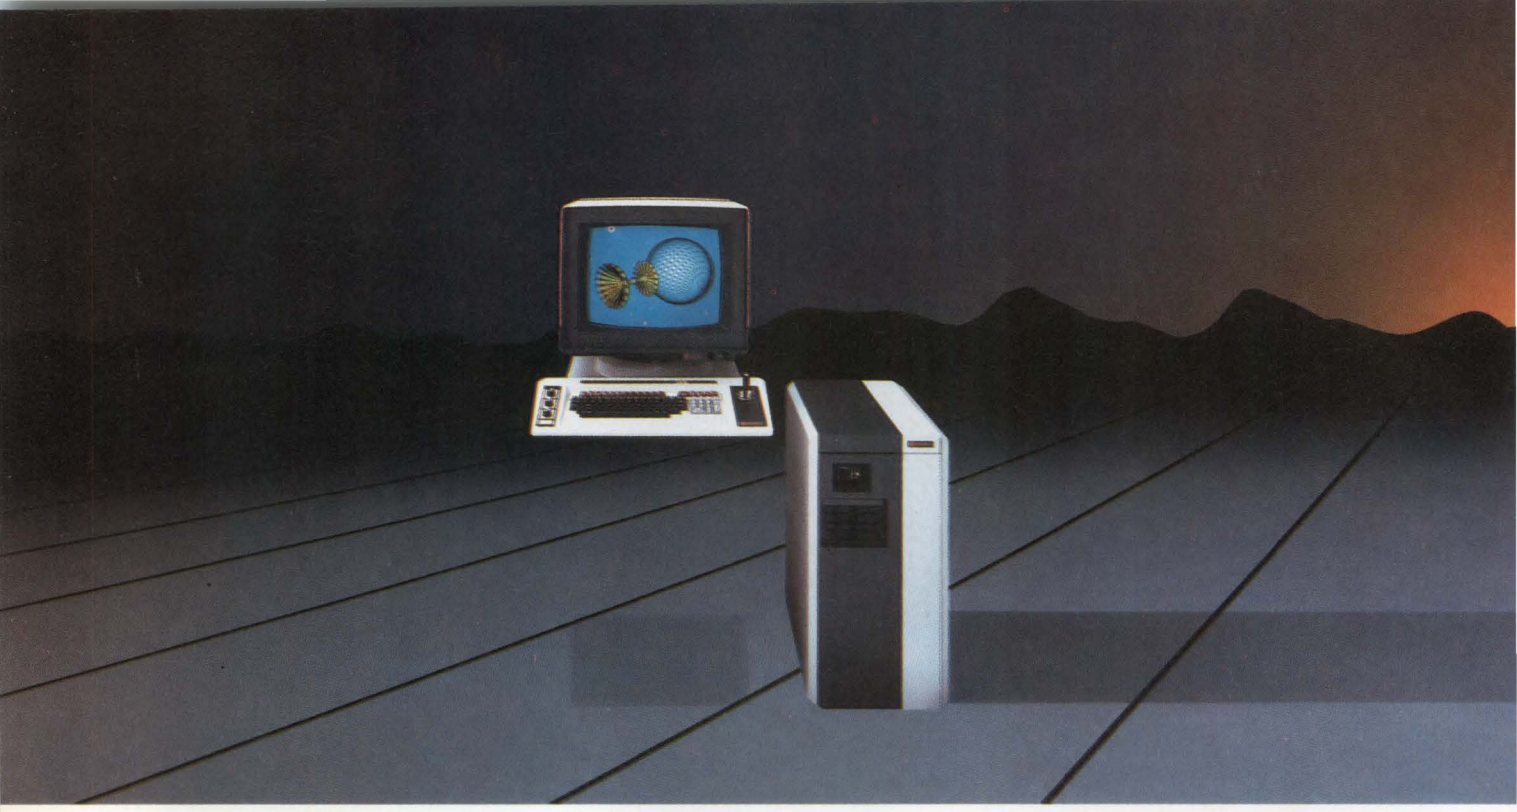

Merlin. The Result of and a Pinch of

More than 15 centuries ago, Merlin—the world's greatest sorcerer—took magic truly beyond state-of-the-art. He not only made the impossible possible—he made it easy.

The same way Megatek is doing today. In computer graphics.

We've taken tomorrow's technological magic to produce graphics miracles today. In other words, what Merlin did for his world, Megatek can do for yours now. The Merlin™ 9200 to be exact.

The most completely interactive, high-performance, high-resolution graphics rendering system yet created. Advanced graphics capabilities combined with high reliability and the first intelligent approach to graphics database management and system networking.

The Merlin 9200 offers graphics capabilities that virtually leave the leading edge of technology behind. Translucents and opaque solids, for example. Flat, Gouraud or Phong shading. X, Y and Z axis clipping of objects. Depth cueing. Perspective. Surface texturing, patterning and meshes. And a unique, proprietary Pixel Phaser™ that gives ultra-sharp 3056x2304 virtual screen addressability. Plus, 64 or 256 or 4096 colors from a palette of 16.7 million.

And behold the system-supported miracles Merlin creates.

Communications. A database editor. Local task processing and display.

1500 Years of R&D Megateknology.

Print, set-up and configuration managers. Peripheral devices including color hardcopy output, a graphics tablet and full-function keyboard with joystick and valuator dials.

Plus, Merlin's hierarchical database allows graphics and non-graphics information integration in up to 4 megabytes of dedicated memory. You can create or modify screen graphics while automatically changing the system's database. You can even store and recall data from either the host or, thanks to an optional Ethernet® interface, from other networked Merlins. Sheer wizardry.

Then there's VT-100™ emulation, hardware diagnostics, status checking, error recovery and a configurable error routing capability.

Merlin 9200 from Megatek. The ultimate high-performance graphics solution that takes you far beyond state-of-the-art. Thanks to Megateknology.™

Pure and simple graphics magic—1500 years in the making, but well worth the wait.

Merlin

See us at AUTOFACT, Booth 515, October 2-4.

MEGATEK CORPORATION®
A UNITED TELECOM COMPANY
Making History out of State-of-the-Art

World Headquarters, 9645 Scranton Road, San Diego, CA 92121
619/455-5590, In CA, 800/824-4489, In USA, 800/854-1975, TWX: 910-337-1270

CIRCLE 52

VMEbus MEMORIES

High Speed DRAM — Non-Volatile CMOS

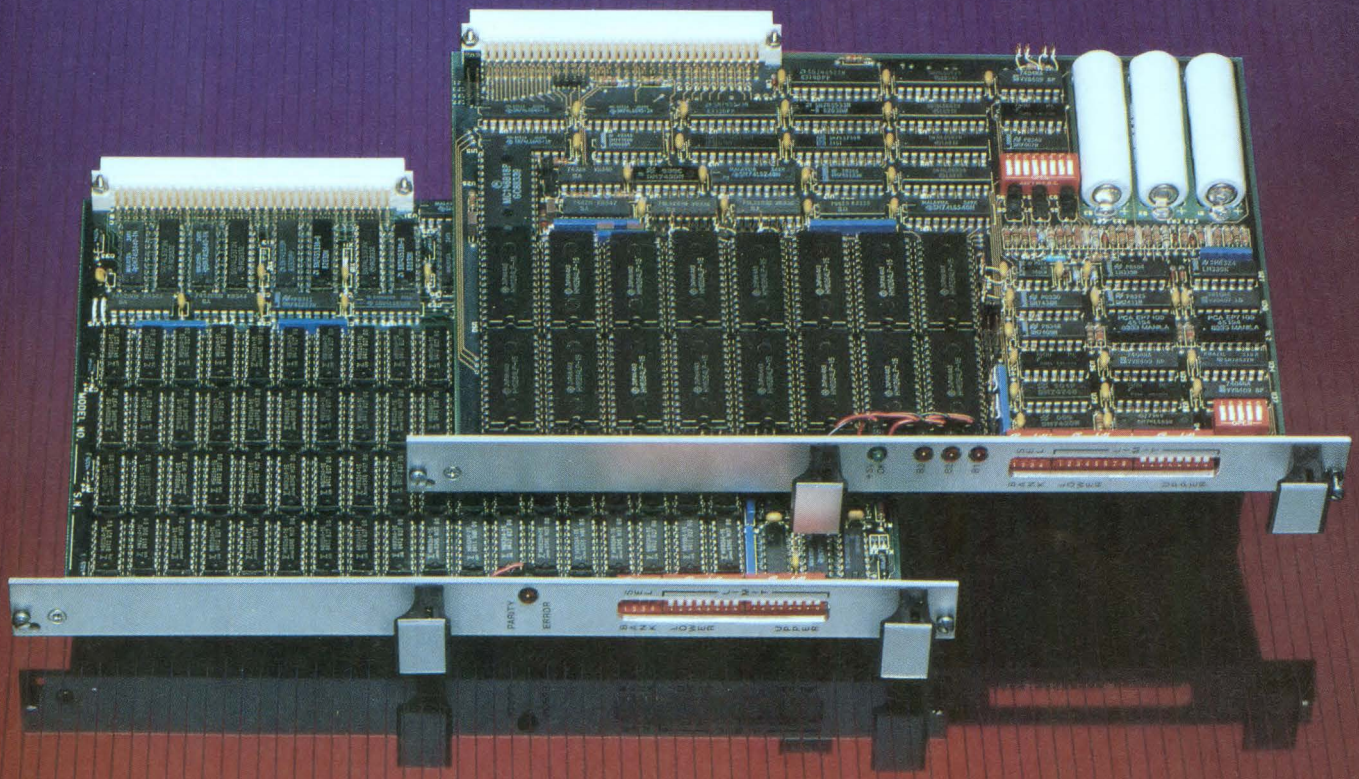

Micro Memory boards allow you to take advantage of VMEbus reliability in your business and industrial control systems. Our VMEbus family includes both high-speed DRAM and non-volatile CMOS memories that are compatible with the VME/10 microcomputer.

Our MM-6000D and MM-6200D are high-speed DRAM memories that provide 2M bytes and 4M bytes, respectively... the highest in the industry. Both boards have parity generation and checking that is indicated on a front panel LED.

If you need non-volatile memory, the MM-6500C and MM-6600C CMOS memories have redundant on-board batteries. Data retention is five years for the lithium version and three months for the NiCad version. The boards are also versatile because you can mix EPROMs and CMOS RAMs. And, the capacity of these two boards is the highest in the industry.

All boards are burned-in for 48 hours during memory diagnostics and there is a one-year warranty on parts and labor.

Model No.	Capacity	Cycle/Access, nsec	Type	Word Length
MM-6000D	256K, 512K, 1M, 2M	330/210	64K or 256K DRAM	16-bit
MM-6200D	512K, 1M, 2M, 4M	350/220	64K or 256K DRAM	32-bit
MM-6500C	32K, 64K, 128K	200/200	CMOS w/Calendar-Clock	16-bit
MM-6600C	64K, 128K, 256K	200/200	CMOS	16-bit

micro Memory Inc ... FIRST IN MICROCOMPUTER MEMORIES

9436 Irondale Avenue • Chatsworth, California 91311 • (818) 998-0070

CIRCLE 53

PARALLEL PROCESSING MAKES TOUGH DEMANDS

Adhering to stringent performance requirements on algorithms, software, and architectures ensures efficient parallel processing.

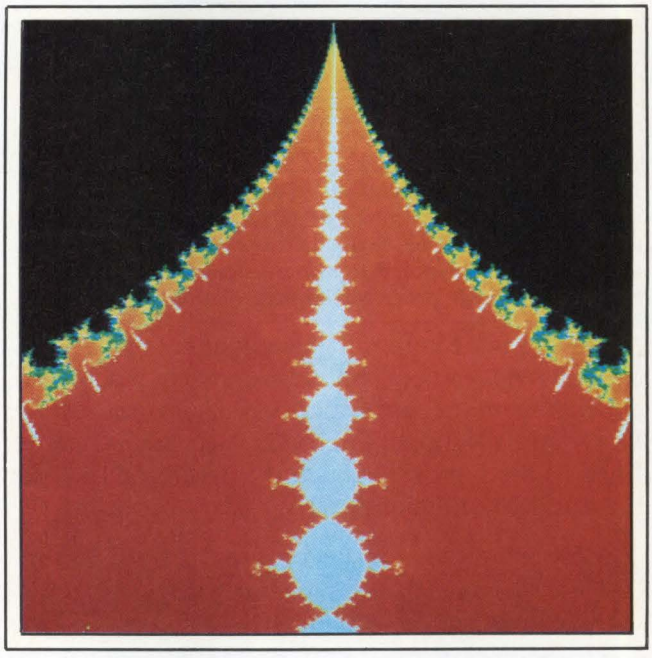

Photo courtesy of Dr John Hubbard and Homer Smith, Cornell University.

by Bill Buzbee

Parallel processing, or the application of several processors to a single task, is an old idea with relatively large literature. The advent of VLSI technology has made testing the idea feasible, and the fact that single-processor systems are approaching their maximum performance level has made it necessary. Successful use of parallel processing, however, imposes stringent performance requirements on algorithms, software, and architecture.

The so-called asynchronous systems that use a few tightly coupled high speed processors are a natural evolution from high speed single-processor systems. Indeed, systems with two to four processors will soon be available (eg, the Cray X-MP, the Cray-2, and the ETA Systems GF-10). Systems with 8 to 16 processors are likely to appear by the early 1990s. The prospects of using full parallelism in such systems to achieve high speed in the execution of a single application depends on the meaning of parallelism. Early attempts with vector processing have shown that plunging forward without a precise understanding of the factors involved can lead to disastrous results. Such understanding will be even more critical for systems that may use up to 1000 processors.

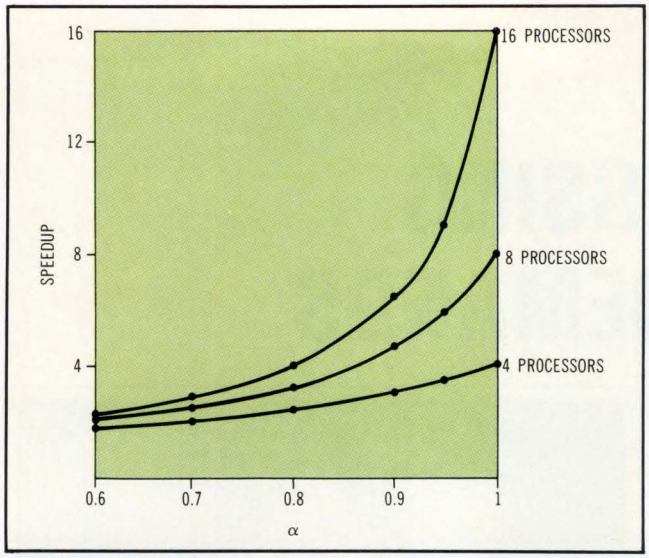

For this model of speedup versus the fraction of computation (α) that can be performed in parallel, results are shown for a 4-, an 8-, and a 16-processor system. Note that a large α is required to achieve significant speedup relative to a single processor implementation.

The key issue in parallel processing a single application is the resultant processing speedup achieved, depending on the number of processors used. At the Los Alamos National Laboratory, speedup (S) is defined as the factor by which the execution time for the application changes. That is,

$$S = \frac{\text{execution time for one processor}}{\text{execution time for p processors}}$$
(where p equals any integer)

To estimate the speedup of a tightly coupled system on a single application, we use a model of parallel computation introduced by W. Ware in his article, "The Ultimate Computer" (*IEEE Spectrum,* Mar 1972, p 89). Thus, α is defined as the fraction of work in the application that can be processed in parallel. Then, a simplifying assumption of a two-state machine can be made; that is, at any instant either all p processors are operating or only one processor is operating. If the execution time for one processor is normalized to unity, then,

$$S(p,\alpha) = \frac{1}{(1 - \alpha) + \alpha/p}$$

Note that the first term in the denominator is the execution time devoted to that part of the application that *cannot* be processed in parallel, and the second term is the time for that part that *can* be processed in parallel. How does speedup vary with α? In particular, what is this relationship for $\alpha = 1$, the ideal limit of complete parallelization? Differentiating S shows that

$$\frac{\partial S(p,\alpha)}{\partial \alpha}\bigg|_{\alpha = 1} = p^2 - p$$

The Figure shows the Ware model of speedup as a function of α for a 4-processor, an 8-processor, and a 16-processor system. The quadratic dependence of the derivative on p results in low speedup for α less than 0.9. Consequently, highly parallel algorithms are needed to achieve significant speedup. However, it is not evident that algorithms in current use on single-processor machines contain the requisite parallelism, and research will be required to find suitable replacements for those that do not. Further, the highly parallel algorithms available must be implemented with care. For example, it is insufficient to look at just those portions of the application amenable to parallelism because α is determined by the entire application. For α close to 1, changes in those portions less amenable to parallelism will cause small changes in α, but the quadratic behavior of the derivative will translate those changes in α into large changes in speedup.

Experienced vector processor users will see a similarity between the Ware curves and plots of vector processor performance versus the fraction of vectorizable computation. This is due to the Ware model's assumption of a two-state machine. A vector processor can also be viewed in that manner. In one state it is a relatively slow, general-purpose machine. And, in the other, it is capable of high performance on vector operations.

Ware's model is inadequate because it assumes that the instruction stream executed on a parallel system is the same as that executed on a single processor. Seldom is this the case, since multiple-processor systems usually require execution of instructions dealing with synchronization of the processes and communication between processors. Furthermore, parallel algorithms may inherently require additional instructions. To correct for this, the term $\sigma(p)$ is added to the execution time for parallel implementation, which is at best non-negative and usually monotonically increasing with p. Actually, σ is a function, not only of p, but also of the algorithm, the architecture, and even of α. Let $S(p, \alpha, \sigma)$ denote speedup for this modified model. Then,

$$S(p,\alpha,\sigma) = \frac{1}{(1 - \alpha) + \alpha/p = \sigma(p)}$$

If the application can be put completely in parallel form, then,

$$S(p,\alpha,\sigma)\,\Big|_{\alpha = 1} = \frac{p}{1 + p\,\sigma(p)}$$

In other words, the maximum speedup of a real system is less than the number of processors p, and it may be significantly less. Also note that, whatever the value of α, S will have a maximum for sufficiently large p because α/p becomes insignificant, while $\sigma(p)$ continues to increase.

One billion byte supermicro. Just think of the possibilities.

Introducing the new, expandable Dual 83/500. A UNIX*-based, 68000-driven supermicro so capable, you'd swear it was a mainframe.

The system already comes with 500 megabytes of Winchester storage. And our patent-pending high-speed SMD disk controller for fast access to data.

But you can increase memory to a massive one billion bytes just by adding a twin drive.

Or take an already sizeable two megabytes of RAM and expand it to six.

Or even double user capacity from 8 to 16. The hardware is already in place.

When it comes to value, no supermicro system offers you more than the 83/500. Because along with the computer, you get a 9-track, 1600 BPI phase-encoded tape drive for reliable disk backup and quick file transfers to other systems.

There's a convenient one megabyte double-sided/double-density floppy disk drive that protects individual files.

And the industry standard UniPlus+™ implementation of AT&T's UNIX System V with Berkeley enhancements. Plus a multi-user license.

All at no extra cost.

And while you're speedily going about processing your data, we're protecting your investment. For free. With a comprehensive 12-month warranty. And a nationwide service network that protects your system whether it's in or out of warranty.

See the system that's redefining the supermicro. The value-packed Dual 83/500. Call or write Dual Systems Corporation, 2530 San Pablo Avenue, Berkeley, CA 94702, (415) 549-3854

At just $65,940 base price, its possibilities are endless.

*UNIX is a trademark of AT&T Bell Labs.
™UniPlus+ is a trademark of UniSoft Corp.

CIRCLE 54

There's only one DUAL

Thus, the research challenge in parallel processing involves finding algorithms, programming languages, and parallel architectures that, when used as a system, yield a large amount of work processed in parallel (large α) at the expense of a minimum number of additional instructions (small σ). Within the past two years, significant experiments using parallel processing were done on scientific computation. Consequently, on Mar 13 to 15, 1984, Lawrence Livermore National Laboratory and Los Alamos National Laboratory hosted a meeting at Gleneden Beach, Ore, at which many of these experiments were discussed. Collectively, these presentations suggest several important results.

A broad spectrum of scientific computation is amenable to parallel processing. The presentations revealed that parallel formulations have been achieved for meaningful computation kernels from such areas as plasma simulation, Lagrangian fluid flow simulation, reactor safety simulation, automated reasoning, solution of linear algebraic systems, chess playing, boundary value problems, aerodynamic simulation, and weather modeling.

Significant speed gains can be achieved from parallel processing, at least in systems with a few processors. The presentations provided much graphical and tabular data in which speedup was a function of the number of processors used. Most of this data was for the Cray X-MP-2, the Denelcor HEP (one PEM and four PEM systems), and the ELXSI (four processors). On the other hand, much of the data was near-optimal (linear) as a function of the number of processors invoked.

In addition, rapid progress is being made in data flow technology. The University of Manchester group has become a hub of international collaboration addressing language, performance measurement, and architectural issues in data flow technology.

All of these results are encouraging. They reflect breadth and rapid progress in parallel processing research, and the impact and importance of having equipment available for experimentation. Yet, these positive aspects should not be overestimated. The presentations also reflect some difficult issues that remain to be solved, including software tools, task management and algorithms, and system issues.

In the area of software tools, problem analysis and decomposition for parallel formulation is nontrivial. Debugging can be difficult because of nonrepeatability emanating from asynchronous computation. New tools for global dependency analysis, graphical representation of flow, and state information will be required.

Computational cost is another important factor to consider. Although it now appears that a relatively small number of control structures may be adequate for expressing and managing parallel tasks, the computational cost (overhead) of these structures will become increasingly important as the number of processors increases. This could culminate in a search for new algorithms that provide large granularity on processor-rich systems.

Lastly, the system issues research thus far has properly focused on the fundamental questions of whether or not scientific computation can be parallel processed and, if so, how well. Now it is time to also address system issues such as I/O, secondary storage, and network support.

Bill Buzbee is an assistant leader of the Computing and Communications Division at Los Alamos National Laboratory, Los Alamos, NM 87545. He holds a BA and an MA in mathematics from the University of Texas at Austin, and a PhD in mathematics from the University of New Mexico, Albuquerque.

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

High 707 *Average 708* *Low 709*

INTEL HAS JUST BECOME A LOW POWERED COMPANY.

THE MOST DYNAMIC PERFORMANCE IN RECENT MEMORY.

If your system design calls for high speed, high bandwidth, low power or all three, you've found the solution.

The Intel family of 64K and 256K Dynamic RAMs in CHMOS. Intel's breakthrough combination of our own high performance technology, HMOS, with the low power advantages of CMOS.

A combination ideal for a variety of applications unmatched by NMOS DRAMs.

Power requirements are so low, batteries are more than sufficient. Which makes the 256K CHMOS DRAM the perfect building block to a smaller, faster, more powerful portable system with the same high performance as a desktop.

For the first time, your portable system can run the most powerful spreadsheet software on the market. The 256K CHMOS DRAM even makes workstation processing possible in a portable.

Their blazingly fast cycle times and high data bandwidth make the High-Performance CHMOS DRAM family ideal for high-speed signal processing or bit-mapped graphics applications.

Two new methods of data access make high data bandwidth possible. Both Ripplemode™ and static column mode provide the ability to randomly read or write any series of bits within a single row at a minimum cycle time of 65 ns. This incredible speed makes the CHMOS DRAM family ideal for graphics display, array processors and other high speed digital signal processing applications.

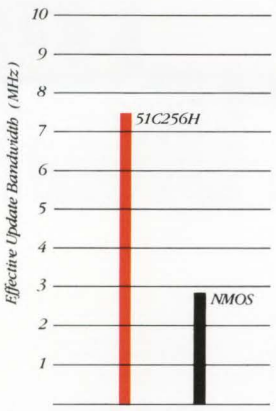

The high-performance advantage: For a typical 1Kx1K display, Intel's CHMOS DRAMs outperform NMOS better than two to one.

In addition, their resistance to soft errors provides the extremely high reliability critical to medical instrumentation, financial transaction systems and other fail-safe applications.

As you would expect, CHMOS DRAMs combine the best benefits of static CMOS RAMs, low power and high bandwidth, with the traditional high density and cost effectiveness of DRAMs.

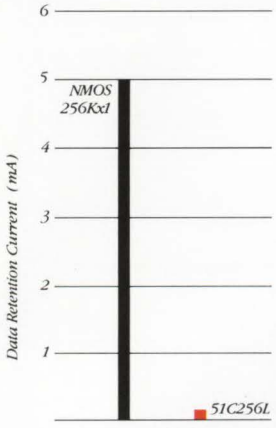

The low power advantage: Intel's CHMOS Dynamic RAMs require only 5% of conventional NMOS.

Low power. Low cost. Unbeatable performance. Available now for your system design.

All things considered, there's no more dynamic performance in recent memory.

intel®

WE COULDN'T OVERPOWER THE 8051. SO WE UNDERPOWERED IT.

The Intel 80C51 eight-bit microcontroller. Featuring the powerful performance of the world standard 8051, but needing only one-tenth the energy.

The 80C51 is based on Intel's revolutionary CHMOS technology. Achieving the high performance and high integration of HMOS, with the low power advantages of CMOS.

The result is an 8051 pin-compatible microcontroller with the most powerful eight-bit CPU and the highest level of integration available. The 80C51 runs on an incredible 11mA at 8MHz clock frequency. Average power dissipation can be lowered into microamp range, by using one of two new software programmable power control modes. In the power down mode, current drain is 50 microamps.

On the performance side, the 80C51 has a maximum clock rate of 12MHz. The fastest data throughput of any CMOS microcontroller.

Its high integration capabilities include 4K bytes ROM and 128 bytes RAM, a Boolean Processor able to manipulate 256 individually addressable bits, a versatile serial port and a pair of 16-bit counter/timers.

All this helps to make the 80C51 ideal for high end eight-bit applications ranging from hand-held instruments to factory line powered controllers, to portable communications systems. A military version of the 80C51 will be available soon.

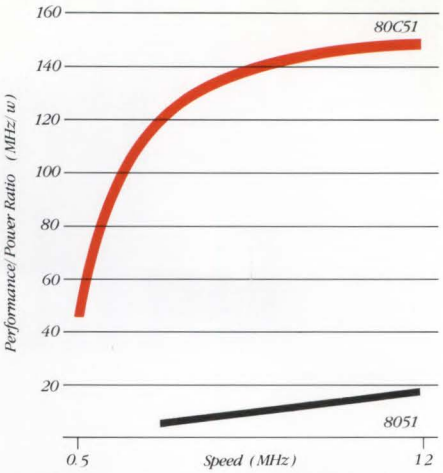

The 80C51 microcontroller delivers the highest performance per watt over the broadest frequency range in the industry. Giving you control over speed and power.

The 80C51 also boasts the world standard in development tools. Like the Series II/ICE™-51 emulator and the iPDS™ system/EMV-51. As well as a relocatable assembler and high level languages.

MCS-51 workshops are available, too.

Go with the 80C51. Because now, running with the very best requires very little energy.

intel®

HOW TO GIVE DATA EQUAL VOICE.

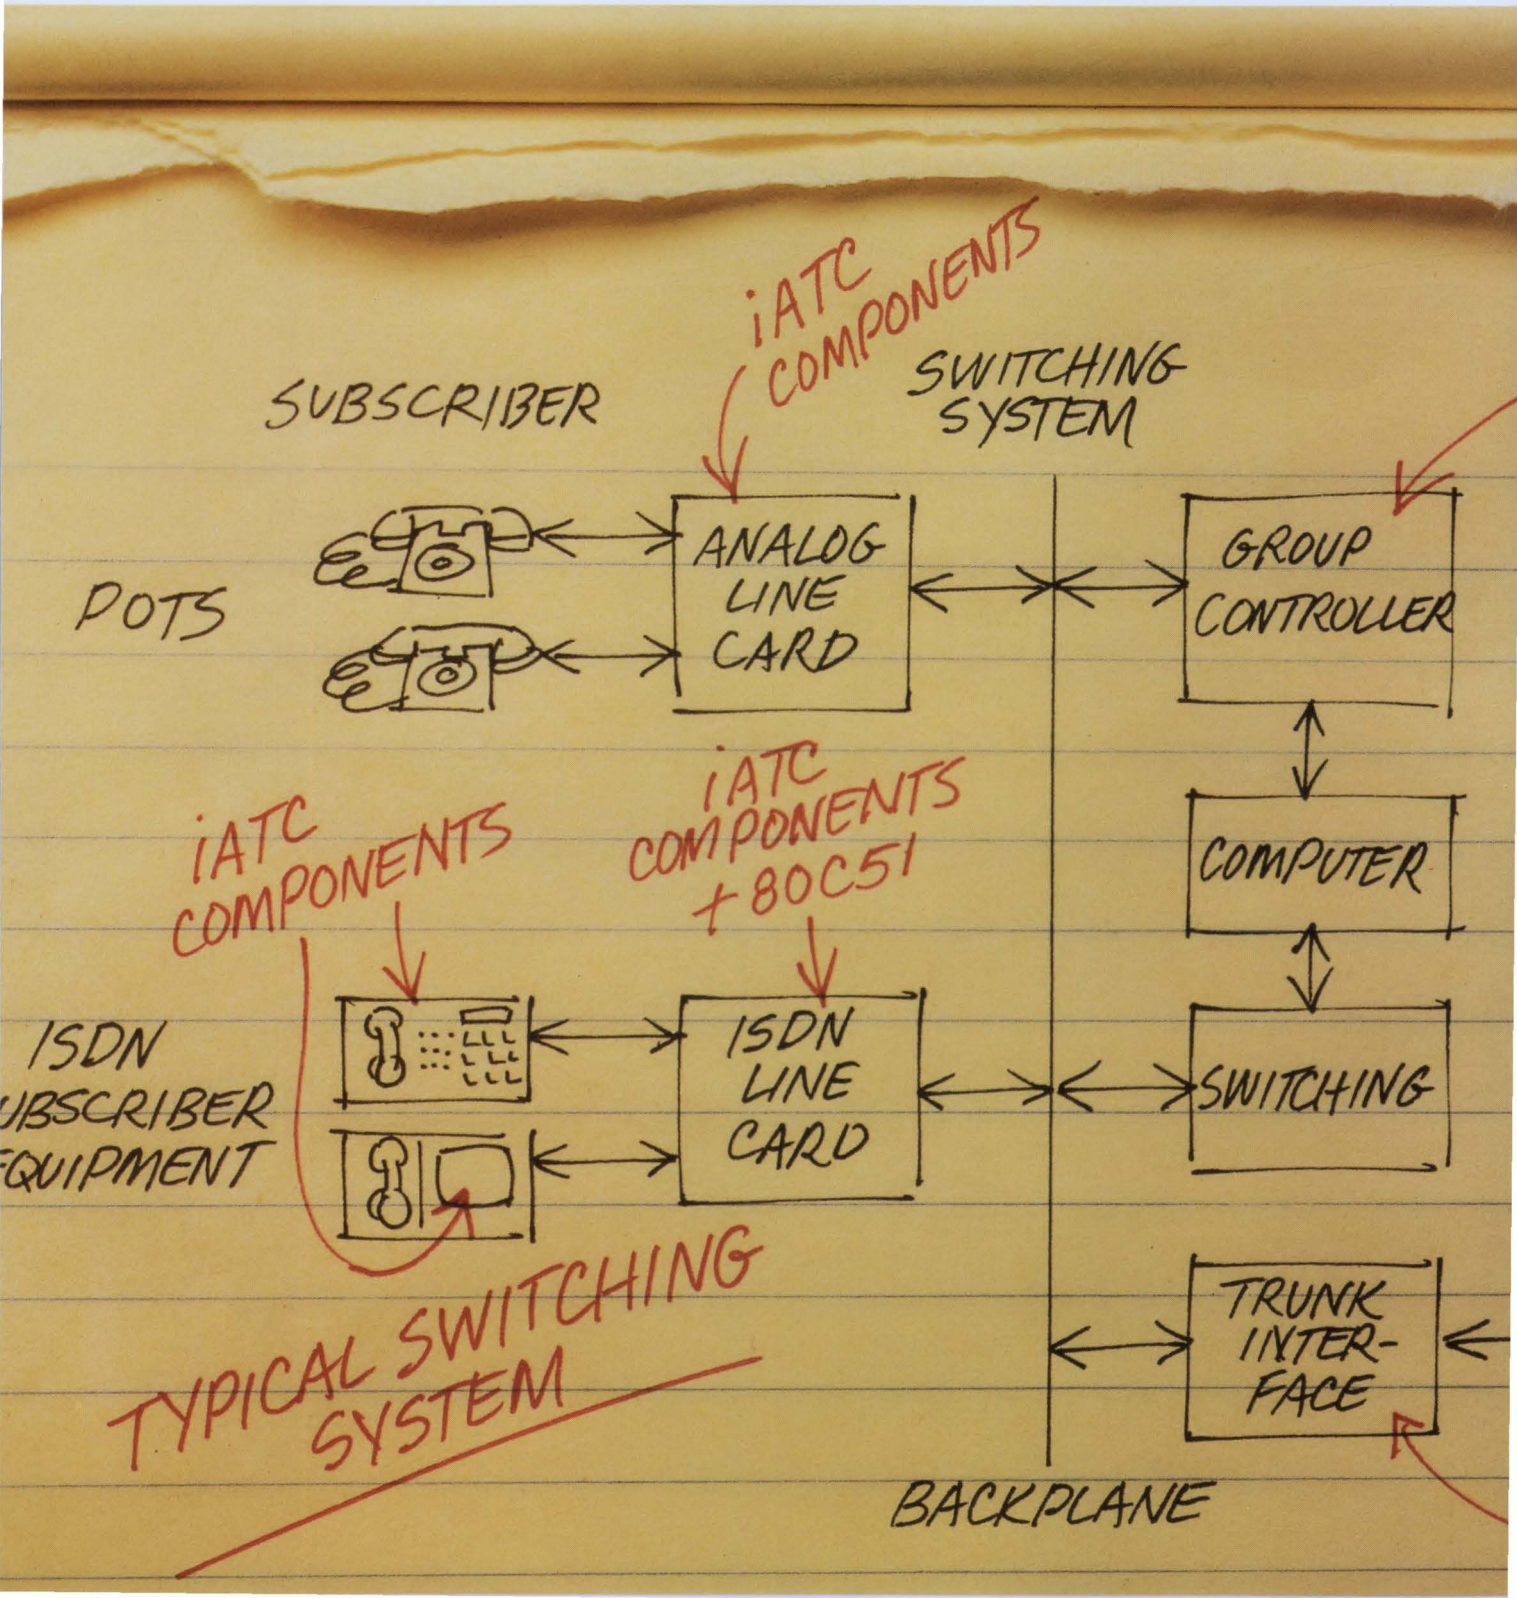

The telephone system of tomorrow will be digital end to end. Able to handle voice and data. Simultaneously.

Today, there's a shortcut to get you there.

Introducing the first two members of Intel's third generation Advanced Telecommunications Component (iATC™) family.

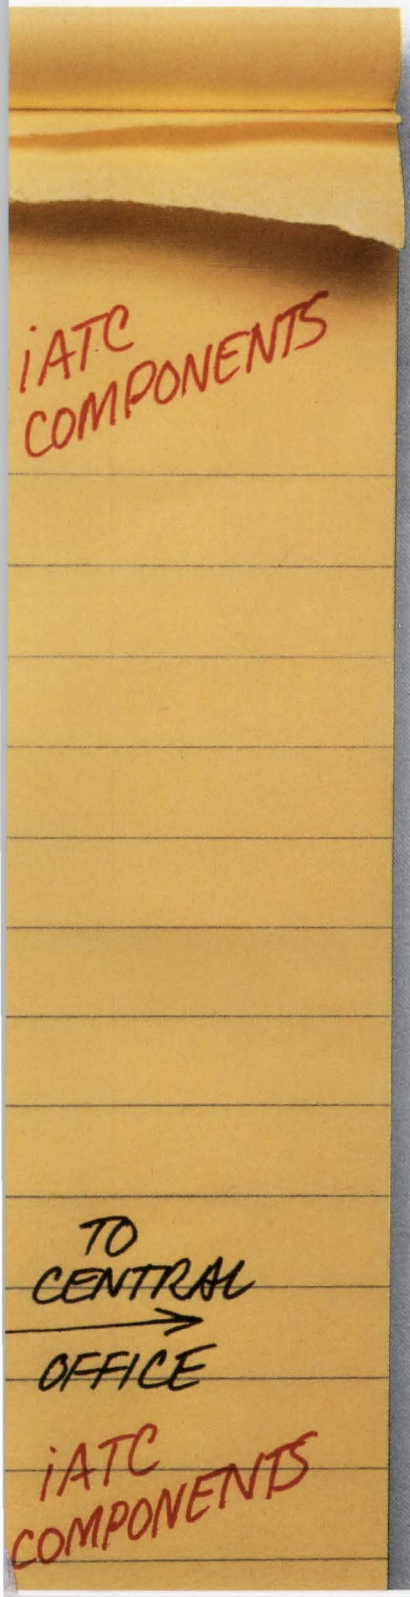

The 29C51 Codec/Filter combo. And 2952 Line Card Controller. Representing a new approach to telecommunication's circuit architecture.

Together, they can handle all pulse code modulation coding and encoding. Subscriber voice channel filtering. Control signaling. And backplane interfacing.

Alone, the 29C51 handles all the analog to digital switching. It's software programmable with a total of 10 features under user control, including transmit and receive gain. You can also program the 29C51 for subscriber line balancing.

It offers a secondary analog channel capable of monitoring an analog subscriber's line.

And internal analog three-way conferencing.

All on one chip. All software controlled.

Because the 29C51 is manufactured with Intel's CHMOS technology, it offers the best of both worlds. The unbeatable performance of Intel's high-performance HMOS technology. And the low power requirements of CMOS.

Developed for use with the 29C51, the 2952 Line Card Controller provides the bridge to ISDN upgrades.

It can control up to 16 subscriber lines at once. And manage all voice and data transfers between the backplane and the line circuits. This includes matching time slots on PCM highways and intelligent interface to a control highway.

Although the 2952 can interface with just about any part of your system, it's optimized for use with Intel's 80C51 microcontroller and the Intel iAPX family of microprocessors.

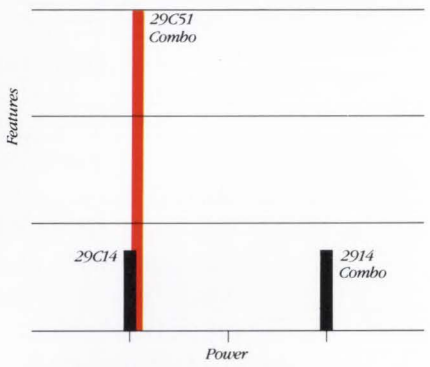

The high performance 29C51 offers the benefits of low power with more features than conventional NMOS.

Which means your telecommunication system design can be virtually unlimited in scope.

Together, the 29C51 and 2952 will work with every other iATC component. Now. And in the future.

So if there's a digital telecommunication system design in your future, give your data equal voice.

intel®

THE LOWDOWN ON INTEL CHMOS.

Intel's CHMOS technology is a combination of its own high performance process, HMOS, with the low power advantages of CMOS. It's a combination that offers the best of both worlds. The density and manufacturability of HMOS with the coolness of CMOS. And new levels of functionality with absolutely no compromise in performance.

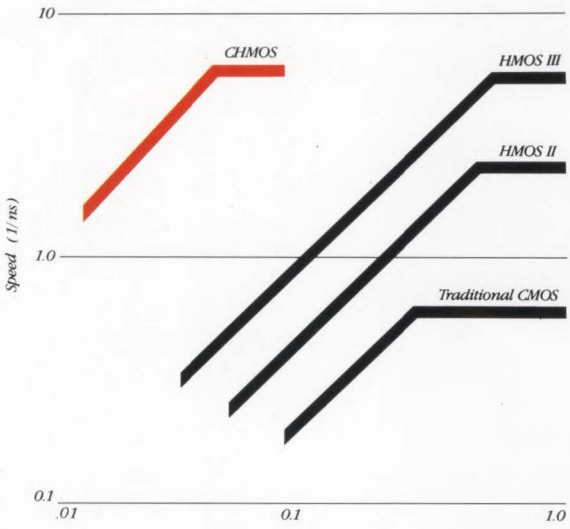

And no break in Intel's leading edge VLSI pace.

CHMOS technology will lead to more dense chips with more and more computer processing power, data storage capacity and other capabilities, with no shortcuts in quality. Every part is built to quality levels that have set industry standards.

Intel is fully committed to CHMOS and has a far-reaching plan to take you along the VLSI path from HMOS through CHMOS and beyond.

As market needs change, Intel will continue to provide high performance solutions to low power applications. Forthcoming CHMOS products include SRAMs, iRAMs, EPROMs, microprocessors and peripherals.

All future products will be developed in CHMOS, providing a high-performance product line that covers more applications and more markets than any other VLSI technology.

For the complete lowdown on Intel's CHMOS products, call toll-free. (800) 538-1876. In California, (800) 672-1833.

intel®

3065 Bowers Avenue, Santa Clara, CA 95051

INTRODUCTION:
FUTURE SOFTWARE

Historians say that in the 1950s, Soviet Russia's Sputnik made the United States realize that space was both commercially and militarily important. Likewise, future historians are apt to parallel this to Japan's 1981 announcement of the Fifth-Generation Computer Project. This project alerted the United States to the fact that artificial intelligence-based computers, supercomputers, and advanced "standard" machines will be important for participation in, if not control of, the international economy of the 1990s.

The hardware to build AI machines, supercomputers, and state-of-the-art home, office, and factory computers in the networked world of 1990 will be available at prices that make "throwing a processor at a problem" even more common than it is today. Processors, memories, peripheral controllers, and other VLSI-based chips will be commodity parts. Even those specialized (custom) chips that clever designers always need could be designed with personal computer-based computer aided design systems. Thus, whether or not software advances enough to make canned programs and similar software development available, software—not hardware—will play the critical part in the design of the three computer types and in the determination of their capabilities.

Because software has such an inherently long development cycle, the software under development today will probably not be operational until the 1990s. But, with software designs for operating systems, languages, and applications inseparable from future computer architectures and hardware, the situation becomes more complicated. For example, parallel architectures require languages with parallel control structures. Chips that implement parallel software algorithms need architectures such as data flow designs geared to concurrency.

The software currently under development for the 1990s is characterized by a lack of standards, and by honest differences of opinion as to which is "best." Perhaps the best example of these is the "conflict" between the so-called flavors of the Lisp and Prolog AI languages. Such problems pertain for lesser known software as well. It is just too early to say which are the best solutions to the many problems of the software design community. Major efforts are underway in Japan, the United States, England, and Europe to address these problems. Similar work is going on in Canada, Australia, and other countries.

The articles that follow address the software of the 1990s for AI, supercomputing, and conventional machines. The staff report article includes some short panels on specialized software featuring academic and industrial authorities. The subsequent articles address the transformation of software art into software engineering, the future of object-oriented programming for both conventional and AI machines, and the future of AI architecture.

As is the case with any speculation, much is based on personal opinion and is sure to be controversial. Computer system designers and integrators will, if they match their opinions with those of the experts, draw conclusions to help them prepare for the design needs of the 1990s.

Harvey J. Hindin

Harvey J. Hindin
Special Features Editor

FIFTH-GENERATION COMPUTING: DEDICATED SOFTWARE IS THE KEY

Software for artificial intelligence-based machines, supercomputers, and personal computers has to be dedicated to individual applications.

by Harvey J. Hindin,
Special Features Editor

The nation that controls the manufacture of the best of tomorrow's computers will emerge as the dominant economic power. But, VLSI chips will not determine the best computers. Software, uniquely designed and tailored to an application, will drive the machines that will be made with commodity-like, readily available hardware. Three distinct, but overlapping computer types will dominate—the artificial intelligence-based machine that simulates human behavior; the supercomputer geared to high speed number crunching and massive calculations; and the "standard" home, factory, or office machine, be it micro, mini, or mainframe. With fine-tuned software, the AI computer will use languages that optimize its symbolic processing and logic manipulations. Meanwhile, the supercomputer will use software geared to its highly parallel architecture, and the "run-of-the-mill" machine will feature easy-to-use, portable software that facilitates the networked home, office, and factory.

The major players in the software development game for the fifth generation of computing are the United States, Japan, and Western Europe. Computer software development has followed an evolutionary pattern and can be somewhat arbitrarily divided into five generations. The first generation lasted about 10 years beginning before the end of World War II, and featured stored programs written

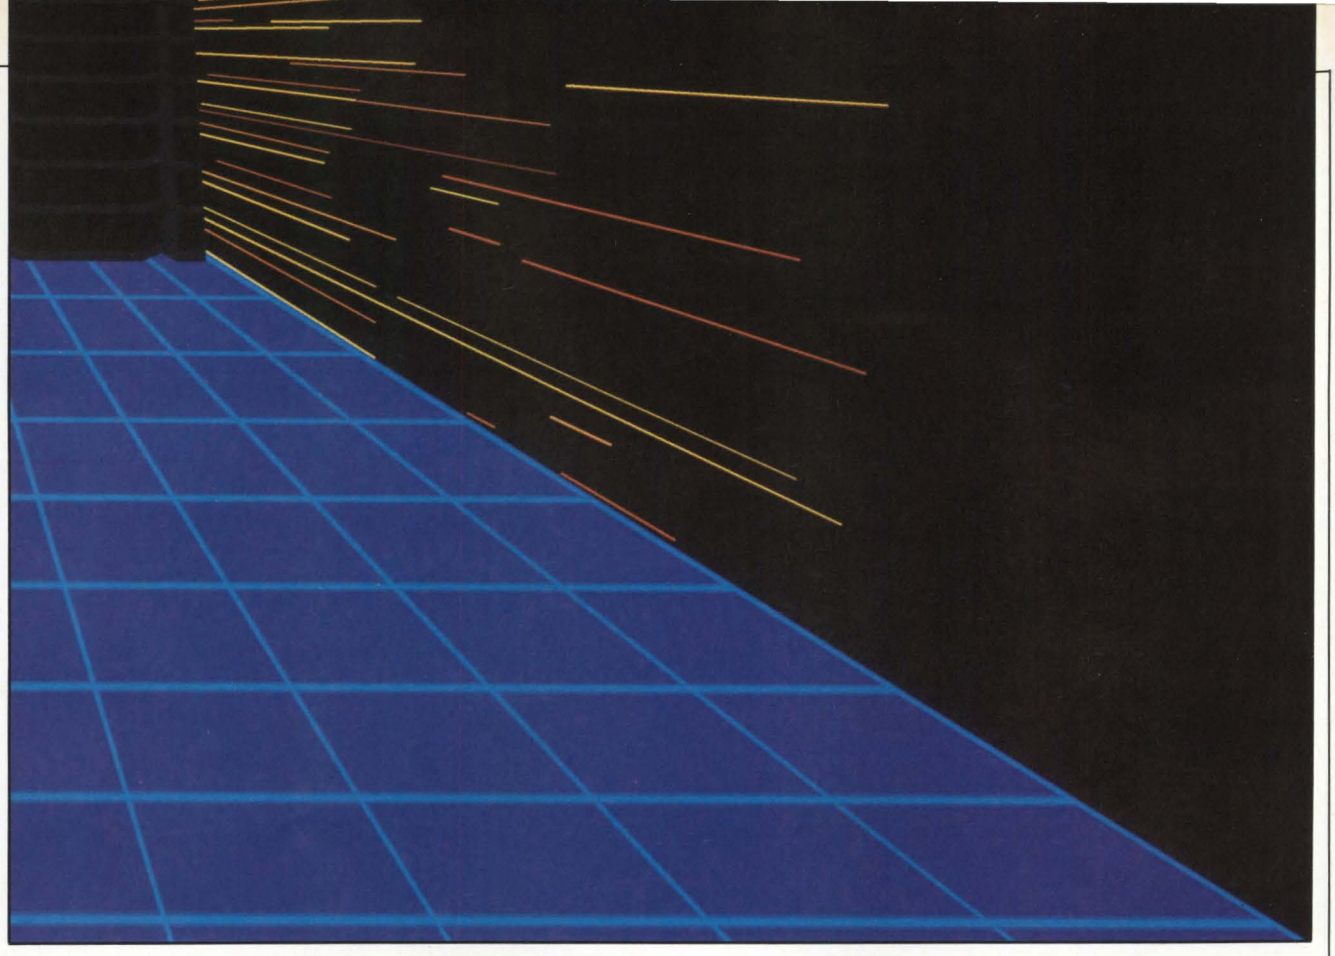

in machine code. Practical difficulties in working with machine code led to the development of such high level languages as Fortran and Cobol in the late 1950s and early 1960s, thus yielding the second generation. In the third generation, concepts such as structured programming were developed as well as high level languages, abstract data types, and the like. Early AI work used Lisp (the list processing language) as the language of choice. Operating system environments, object-oriented languages, and knowledge-based systems came to the marketplace in the 1980s as the fourth generation. Perhaps most important, this decade is witness to the phenomenal success of the personal computer and packaged application programs, and thus is proving that computers could serve a wide audience outside the scientific and data processing worlds. The fifth generation should blossom by the beginning of the 1990s, and consist of symbolic processing with extensive AI applications. Moreover, software geared to parallel processing and supercomputers may sport concurrent processing languages and functional programming.

Thus, much of the money that will be spent will go toward furthering the development of existing software ideas. Software's evolutionary nature cannot be overemphasized. A look at software history shows that developing new concepts into practical systems often takes 10 or more years. Since it takes years to design, test, debug, document, and develop applications, the software of the 1990s is already being developed in universities, research organizations, and some industrial firms. Today's AI languages are not new, and logic programming as well as parallel processing/data flow architecures and languages have been known for years. Thus, for more conventional computers, Modula-2 is one example based on well-known Pascal, while Ada has been under development for five years, and the Unix operating system is really a late 1960's design.

Across the Pacific

The Japanese so-called Fifth Generation Computer Project, announced in 1981, is largely based on non-von Neumann parallel architectures and has done much to galvanize American thinking about future computing. The Japanese have boldly announced that software (AI-based software in particular) is the key to their computer superiority and domination.

Knowledgeable Japanese computer scientists say that AI will not be the complete answer for them or anyone else. Supercomputers and conventional machines will have their part to play as well. Government and industry leaders in the leading high technology countries have ongoing efforts to develop the very best in all three major computer areas, but the details of fund allocation have yet to be defined.

An AI computer is different from a regular computer in that it exhibits behavior similar to that of a human. It shows, in some sense, intellect, although the precise nature of this intellect is a subject of much

debate among computer scientists, psychologists, and philosophers. Such matters are not of concern here. Realize however, that a machine that merely calculates at some prodigious rate is a feat of speedy hardware and clever software, not of AI.

AI is not a new field. It has been studied at universities for over 20 years, and has yet to live down its failure of 15 or 20 years ago, to deliver much of value. At that time, it was oversold by its proponents, who are now older, and supposedly wiser.

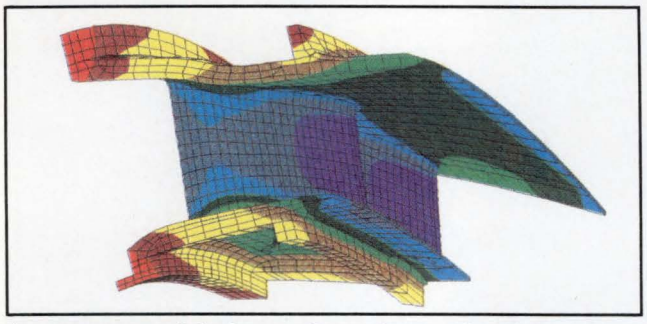

A computer model of a nozzle produces color plots showing predicted areas of temperature stress. This computer aided design example (courtesy of ADAPCO) used an FPS-164 scientific computer from Floating Point Systems, which uses array processing techniques to speed computation.

Defining AI in a manner agreeable to all is next to impossible. It certainly involves more than just calculating power or speed. AI implies a machine that behaves like a human in certain activities. The problem is that every machine or software system said to be artificially intelligent is no longer deemed so by those who build it. Its principles become well understood within the realm of computer science. Although hard to believe, time-sharing, for example, was originally deemed an AI phenomenon.

One operational AI definition, made over 40 years ago by English computer scientist Alan Turing, is still considered valid today. Turing claimed that if an observer asks the same question of a man and a computer that are hidden behind a curtain, and the observer cannot tell the difference between the answers of man and machine, then the machine may be considered to be intelligent. Much has been written about Turing's experiment, and there is much disagreement about its meaning. [To read about Turing's work, see *Alan Turing—the Enigma* by Andrew Hodges (Simon and Schuster, NY, 1983).]

Speed and computation power are necessary but not sufficient conditions for AI implementation. AI computers depend, among other things, on symbolic processing, which consists of substituting abstractions or abstract representations for real objects, and working with them much like the human brain. Also, heuristic reasoning (rules of thumb or experience) is used to reduce the realm of possibilities in a given situation (again like the brain), and even learning from experience (like the brains of some of us). One day, AI programs may write AI programs that may write AI programs that may....

It appears that some form of AI is necessary for applications such as natural language understanding—the programming-free key to widespread computer use among the masses. And, it is clear that vision systems that can recognize what is going on around them in a completely random scene (children can do this) will need AI-like software.

In general, an operator's native tongue may be the most efficient computer programming language, at least for those applications not requiring a high degree of mathematical precision. Thus, a worldwide effort is on to have computers understand natural languages. Similarly, if computers could "see" and react to human speech, a wide variety of chores (many unpleasant or tedious) in the home, office, and factory could be taken on by machines.

Advances in natural language understanding, based in part on highly parallel processing language models and AI-based software designs, will allow computer system designers of the 1990s to enjoy the services of natural-language-understanding robots. These will still be limited domain systems, since the software programming and design problems for even a slightly general purpose robot are staggering. In fact, today's state-of-the-art knows no way to do it.

Vision systems suffer from the inability to do even child-like tasks. For example, it would be a major feat for them to put many differently shaped solids into differently shaped holes. The problem is recognizing objects in the context of a general scene, and then having the computer understand these three-dimensional relationships.

AI principles will help store information about what is to be expected—at least in certain robot domains. Performing such realtime chores takes computing power on the order of at least many billions of instructions per second. Today's hardware and software is just not up to the job, and three-dimensional realtime robot vision of arbitrary scenes under arbitrary lighting conditions is but a dream.

Similarly, the idea of a computer understanding continuous speech by a random speaker on any subject at all in a noisy environment is, for the foreseeable future, only science fiction. The knowledge base, natural language understanding, and processing speed needed are beyond the technology. In contrast, limited domain speech—eg, a trained, specific speaker with a limited vocabulary in a high signal-to-noise ratio environment—is possible. For example, a limited vocabulary, computer-based dictating machine and a voice-responsive typewriter may be available by the 1990s. Applications in database manipulation and financial planning are also candidates for AI implementation.

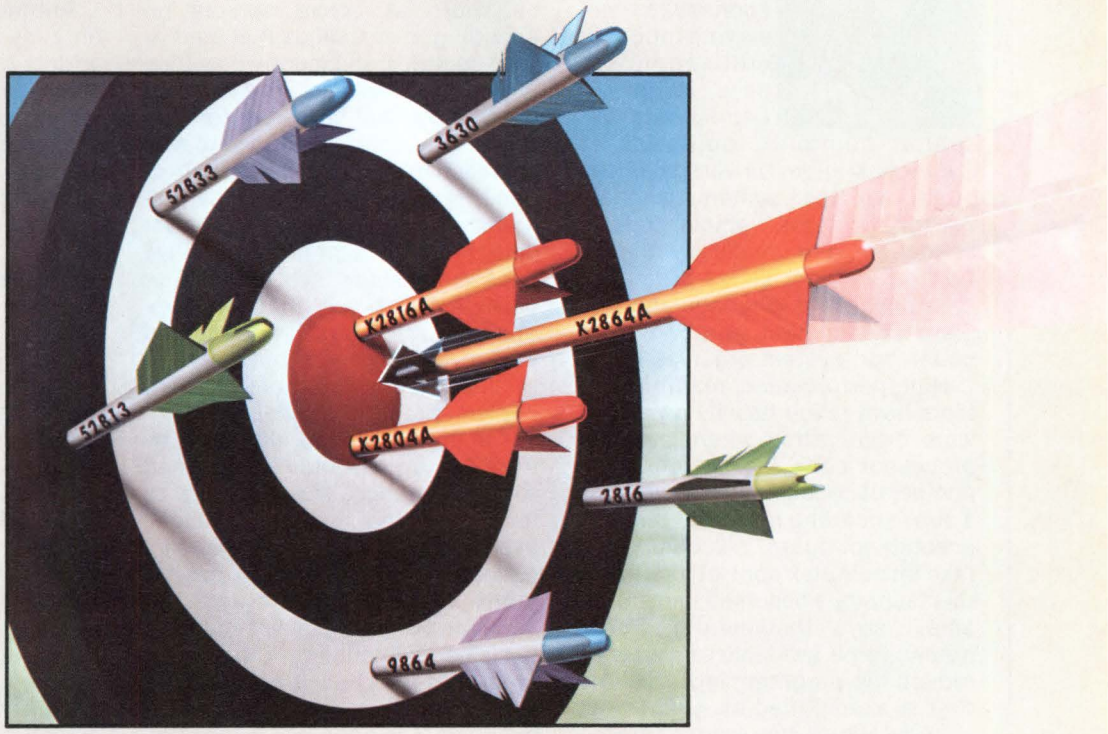

Xicor E²PROMs First Time. Every Time.

Success. At Xicor, it's become a habit. We've been forging new positions of leadership in nonvolatile memories for over five years. With very consistent results.

Others have tried. And missed.

But our family of 5V E²PROMs has emerged as the standard everyone strives for. In fact, our parts are so successful, some of our competitors have chosen to scrap their own designs in favor of second-sourcing ours.

Imitation. The sincerest form of flattery.

Since we unveiled the 5V concept of nonvolatile memories, we've shipped over five million devices.

More than all other manufacturers combined.

We were the ones who set the 5V-only standard of operation. And every part we've introduced since then has been right on target with what the market needs.

Standards. We're making a habit of that too.

Feature for feature, our E²PROMs offer everything you need to solve your nonvolatile design challenges. And we've got more design solutions under development now. For more information on pricing, delivery and performance, write us at

851 Buckeye Court,
Milpitas, CA 95035.
(408) 946-6920

Xicor

CIRCLE 57

Future high performance, realtime systems

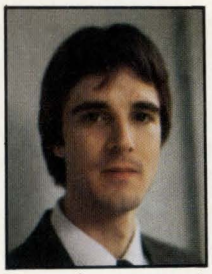

"High performance, realtime systems for the 1990s will deliver complex responses to time-varying inputs in milliseconds," says Roger B. Dannenberg, research computer scientist at Carnegie-Mellon University (Pittsburgh, Pa). Such systems will handle robots, tools, vehicle control, animation, and music synthesis. Two interacting issues critical to their design are language and program development, and communication and system organization.

High performance, realtime systems have relied heavily on a software organization where a single processor polls the system for a number of conditions, and executes a corresponding response for each enabled condition. "As computers take on complex control functions, this technique becomes unmanageable," says Dannenberg. Fortunately, three available approaches reduce the programming task. The first is exemplified by Ada.

"Ada allows the programmer to factor a program into a set of relatively simple processes, each responsible for a single task," explains Dannenberg. The problem of multiplexing the processor among many tasks is handled by the language implementation. A disadvantage is that when tasks are not entirely independent, their interaction can be hard to understand.

The second approach, according to Dannenberg, is based on object-oriented programming, and has been used for animation and music synthesis. The idea is to represent behaviors as active objects responding to messages that start or stop behavior, and modify it as time progresses. The messages can contain scripts that specify future actions. "The chief advantage to this approach is that it supports abstraction," says Dannenberg. "The programmer can build very high level objects in terms of more primitive ones and hide low level details. The disadvantages are that the programmer must still be concerned with the synchronization of access to shared data, and more efficient implementations are needed."

The third approach, illustrated by the Carnegie-Mellon developed Arctic language, has its roots in functional programming languages like pure Lisp. "Because such languages have no variables or assignment statements, values are computed once and never change," explains Dannenberg. This design greatly simplifies parallel task synchronization—that is, a task may begin as soon as its input arguments have been computed, as in data flow computers.

Without changing variables, previous functional languages have not been able to control physical processes. "The novelty of Arctic is that its values are functions of time rather than constants," says Dannenberg. "Thus, the time-varying inputs and outputs of realtime systems can be directly expressed in Arctic values. Synchronization and communication between parallel activities is implicit, as is the case in other functional languages."

Another important issue for realtime systems is communication and system organization. To achieve high performance and low latency, future systems will be assemblies of special purpose processors. For example, there will be specialization in VLSI processors for image processing, signal processing, speech recognition and production, array processing, and communication. "The system designer's primary challenge will be the efficient utilization of a heterogeneous collection of processors, and the efficient communication of results from one processor to another," says Dannenberg.

As examples of the emerging techniques, Dannenberg points to systolic arrays, which enable extremely high throughput for very regular computations. In addition, such processors as the Inmos Transputer can enhance communication between processors with onchip hardware for high speed communication channels. Finally, there are NEC's data flow processor chips that can be pipelined so tagged values are passed from one processor to the next. "Clearly, because high performance systems must avoid the use of a central processor through which all communication takes place, a variety of alternatives will be explored," concludes Dannenberg.

Many of these AI applications are satisfied by expert or knowledge-based systems sold by the AI firms that have entered the marketplace in the last two years. These AI software products provide aid to the applications just mentioned, as well as to such diverse applications as compiler design, medical diagnosis, chemical spectrum analysis, oil well detection, computer system layout, signal analysis, and image recognition (photographic, sonar, or radar). These computer programs use encoded knowledge and reasoning techniques (normally only available from experts) to aid nonexperts.

Much of the work in developing expert systems is aimed at making them better. They are limited in speed, and even in how much of the expert's knowledge can be encoded. No one really knows how experts go about solving a problem. Major work needs to be done, for example, in expert systems to process free-form natural language and perform image analysis in real time.

Software writes software

One major task for AI-based software is to help future programmers—if that is what they will be called—write more, and better software at lower cost. Writing, testing, and debugging software—be it for AI machines, supercomputers, or regular computers—is a labor-intensive task that needs all the help it can get. Expert systems will help here, as well as in the physical design of complex VLSI chips, PC boards, and higher level subsystems. Can an expert system-designed computer design other computers? The Japanese are experimenting with a robot production line that makes robots.

Exploratory programming environments, under development worldwide, are expected to play a

major part in the 1990s for cost-effective program development. Such environments comprise an integrated set of software tools—many AI-based—that can work along with programmers to help them design, test, debug, and document their work. The key to exploratory programming in an AI environment is incremental development. In other words, the completed program need not be specified in some formal way (often this is impossible). Rather, it can be developed piecemeal, almost as a realtime experiment, to see what evolves.

Much AI programming is accomplished in the OPS-5 AI language and development tool set originally designed at Carnegie-Mellon University, and now distributed by Digital Equipment Corp (Maynard, Mass). Of course, there are also the Lisp and Prolog programming languages, which represent the relationships key to AI programming. Another approach to the representation of relationships is the concept of nonprocedure-based, object-oriented programming. The best known object-oriented system is Smalltalk-80 from Xerox Corp's Research Center (Palo Alto, Calif). Others, such as Actors, which was developed at the Massachusetts Institute of Technology (Cambridge, Mass) are also in existence. (See "Actors Set the Stage for Software Advances," p 185.)

Like logic-based programming, object-oriented programming allows flexible programming. For example, with object-oriented programs, objects (data plus information about the commands that the data responds to) can be organized into classes. These classes are examined for common features so that conclusions can be drawn or deductions made about the data (or objects). In addition, objects may inherit or share behavior when subject to certain commands. And, since each data object has a unique specification set associated with it, it can always be tailored to an intended use. Object-oriented systems can, like Prolog programs, be easily extended. Existing objects need only be modified for new applications. In other words, object-oriented programs tend to be highly modular and can be modified without rewriting the whole program (as is the case in conventional programming).

Both object-oriented programming and exploratory programming (especially the latter) are in early development stages. It will take years to determine their viability for practical applications and how much of an impact they will have on fifth-generation computing. For example, the Japanese Institute for New Generation Computer Technology (ICOT); its American counterparts nationwide; the European Economic community endeavor, known as the European Strategic Program on Research in Information Technology (Espirit); and the British effort, known as the Alvey program for Advanced Information Technology (England is also in Espirit), are supporting both object-oriented and exploratory program research.

Multiple choice

AI systems need not be written in a language geared to them. For example, both PL-1 and Fortran have been used. But, authorities agree that future AI software will be written almost exclusively in an AI-geared language like Lisp or Prolog. The differences between the latter two are clouded in nationalistic interests, the not-invented-here syndrome, honest differences of opinion, and, like all software choice issues, matters of personal preference. But, let it be said from the beginning that what can be done with one language can be done with the other. The computer system designer interested in what is going on cannot get into the details—even experts agree the choice between the two languages is often a matter of "religion." Experience dictates that there will be no "winning" language—each will have its marketplace.

The U.S. AI community prefers Lisp, although work is going on in Prolog. More importantly, it appears that U.S. Prolog work is accelerating. Of course, Lisp has been used for about 20 years and this alone is enough to give it a strong base in the AI community. In contrast, the Japanese have said they are opting for Prolog—a mere youngster as languages go. It was invented about 1971 in France, and further developed in England and Australia. Compared with Lisp, Prolog has far fewer software development tools associated with it.

Prolog, a logic programming-based language extends the functional programming capabilities of Lisp. In fact, some say Prolog is an extended Lisp. As mentioned, others maintain that Lisp is as capable as Prolog. According to ICOT, fifth-generation technology and Prolog-based AI, in particular, will have a variety of societal benefits. The same benefits, its proponents claim, can be obtained from Lisp-based systems.

Consider, for example, some fifth-generation system designed to support knowledge-based expert system software. One technique for implementing expert system software is to use If-Then rules. Equivalent to formal logic statements, these are easier to use, but are somewhat different from the If-Then rules in classic programming languages. For example, they are modular and execute immediately. The If-Then-based expert system, and expert systems based on Lisp, Prolog, "pure" logic, or other programming techniques are all candidates for the AI software of the 1990s. For the Japanese, Prolog is considered natural for logic programming in general. From the purist viewpoint, a Prolog-based machine is not a full logic machine—it makes compromises.

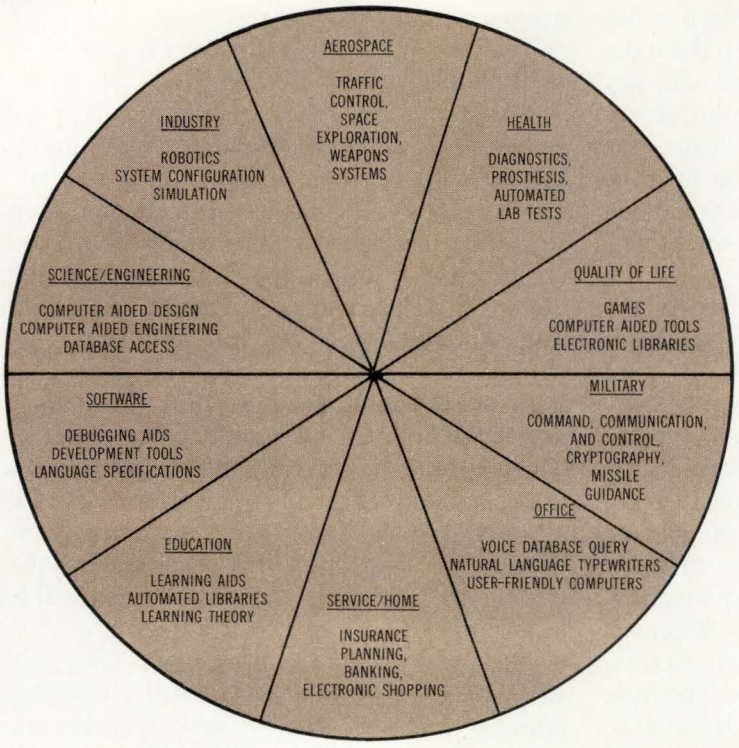

By 1990, the advances in AI made possible by the Japanese Fifth-Generation Computer Project and parallel efforts in the United States, Europe, and England will make a variety of new applications possible in diverse industries. Most aspects of human life will be affected.

Such questions are part of an ongoing controversy about logic software and its place in AI.

A little bit of Prolog

In the Prolog (programming in logic) language, all programming statements are relationships called clauses. Program and statement execution constitutes logic-based deduction from the program clauses. Different from most programming languages, Prolog does not specify procedures to make its deductions. Deductions follow naturally from the clauses because the procedures are implicit within the language's structure. The following example makes this clear.

Prolog can deduce logically from facts given to it. For example, if it knows that Andrea is Eric's sister, and Wendy is Andrea's mother, and it has definitions of motherhood and siblingship, it can conclude that Wendy is also Eric's mother (disallowing today's divorce rate). Thus, in Prolog, sister (andrea,eric). (Note the period at the end of an assertion) and mother(wendy,andrea). implies mother (wendy,eric). In this particular data structure (known as a Prolog term and, in most Prologs, not an abstract data type), the relation subject is the first argument and the relation object is the second. If the Prolog machine were given sister(andrea,eric). and mother(wendy,andrea). and asked the question mother(wendy,eric)? (note the only difference between an assertion and a request for a deduction is the change from a period to a question mark), it would reply in the affirmative.

Prolog will answer no to a query if it does not have enough information to answer logically. It can also search its information to answer questions such as who is Andrea's brother or, in Prolog, brother (%who, andrea)? Note that %who or %Z (or some other notation) defines a local variable whose name the machine is to supply. A no answer from the machine means that it is not able to deduce the answer to the query.

With further knowledge, Prolog can make deductions not given to it obviously. For example, if it knows what a sibling is, it can deduce who the siblings are. This is done by means of Prolog's ability to bring to bear user-defined implication clauses (the left-hand side of the clause—known as the head or goal— is true if the right-hand side—the subgoal—is true) such as sibling(andrea,eric)←mother(wendy, andrea),mother(wendy,eric). or more abstractly as sibling(%X,%Z)←mother(%Y,%X),mother (%Y,%Z).

These logic programming examples, while informative, are little more than cute. Prolog's power lies in its attempts to do the user's bidding and "solve" a goal. It does this by solving each required subgoal. In other words, Prolog looks for assertions that may be combined to derive (prove) the goal.

In fact, Prolog can backtrack to find alternative logic routes to solving a goal if one or more routes

fail. Some human thinking processes proceed this way when solving a problem. As an aid, all Prolog statements can be examined as to their truthfulness independently of other program statements. In short, Prolog is a nonprocedural language. As such, the statement order (clauses) is irrelevant. This is the ultimate in modularity and a program can be increased in size or made more complicated without rewriting. This feature is useful in expert systems when more knowledge is added. For example, using the brother and sister examples, it is easy to add divorce into the picture.

Because Prolog allows the computer to draw inferences from the true or false statements (predicates) given to it, it is Lisp-like. But, Prolog's techniques to perform unification or pattern matching (the satisfying of logic goals mentioned earlier), as well as the backtracking techniques used to seek alternative routes to goals, are what make it somewhat different from Lisp.

To gain perspective on this relationship, realize that, to some, Lisp is a Prolog subset and Prolog contains many Lisp functions. Both languages are ideal for list processing, although Prolog is said to be better able to parallel process data quickly. Nonserial versions (when they are developed) will be suitable for parallel architectures.

Prolog and Lisp suffer from the Unix disease—multiple flavors. Prolog is based on first-order predicate logic (ie, assertions can be made about variables in a proposition but every statement must be proven to arrive at a result). Prolog versions under development are based on other brands of formal logic that will shorten goal derivation time and effort. Although there are multiple versions of first-order, logic-based Prolog and Lisp, there is an effort to develop a "common" Lisp.

Is it logical?

An ongoing controversy in the AI software community is where logic programming will fit in future computing. There are two, not mutually exclusive, points of view that are well represented by the thinking of two major people in the computer science community—John McCarthy and Marvin Minsky. McCarthy, the inventor of Lisp, believes that formal mathematical logic, as studied by mathematicians specializing in such matters, is key to AI reasoning (whether or not people do it that way). Minsky demurs, claiming that the best approach is to get computers to think like people—which he thinks is not by means of formal logic.

Since no one knows how people think, the controversy is partially academic. There is no purely logical programming language that is practical. (While pure Prolog is a descriptive rather than a procedural language, in practice it contains procedural mechanisms so programs can run.) Moreover, only specialists can understand the subtleties of the arguments.

It does not appear that either side, in practice, totally rejects the other's view. The different opinions are a matter of degree. How an AI program should best combine logic, heuristics, reasoning, and the like is what the argument comes down to in practical terms. In any case, although fifth- or sixth-generation computing could be affected by what comes of the argument, today's debate is at a standstill. Ongoing research and practical systems appear in the universities and on the market, combining the best features of both viewpoints.

Part of the problem in applying pure logic to real-world situations is that it is not meant for such applications. Mathematical logic is designed for deducing conclusions from axioms and rules in a narrow domain. Certain well-ordered and sharply defined mathematical symbols were to be manipulated; thoughts cannot be easily represented in a mathematical logic system.

These objections do not mean that logic cannot ultimately be used for real-world applications. On the other hand, researchers are fond of coming up with practical examples where first order logic is not up to the job. However, this does not imply that higher order logic is inadequate. The AI community just does not know. Minsky devotees claim that the gap between AI requirements in representing the real world, and what logic-based languages can deliver, is too wide to be bridged. They maintain that pure logic will never be enough, and that present hybrid techniques in real systems are the answer.

Perhaps Stanford University's Edward Feigenbaum's point of view is appropriate here. AI guru Feigenbaum sees AI as divided into work on intelligent machines and work on cognitive science. The former is concerned with software to do a defined engineering job, while the latter is concerned with modeling human information processing. For Feigenbaum, the intelligent machine effort is an engineering problem. Whether it is solved by means that mimic the human brain is irrelevant. Moreover, according to Feigenbaum, because pure logic programming is neither convenient nor understandable, such methodologies as the If-Then must be used.

While the Japanese and other players in the Prolog game are working to solve Prolog's backtracking and other problems, Prolog may end up being fundamentally flawed just because it is a logic-based language. It turns out that logic-based languages (some with even more features than Prolog) have been proposed before and found lacking. They have turned out to be duds because logic-based languages are not flexible enough for the real world. For example, the 1960's logic-based language Planner (out of MIT) could deduce both in a forward and a backward

Spawning new computer architectures from AI software

"The breakthrough in the development of the first stored-program computer was not in the fabrication of its parts but rather in the arrangement of those parts into a hardware architecture suited to numeric calculations." So says S. Jerrold Kaplan, chief development officer at Teknowledge, Inc of Palo Alto, Calif. According to Kaplan, "As computer use extended from scientific and engineering applications to business and then consumer applications, this von Neumann architecture was less appropriate, but, nevertheless, persisted with only minor modifications." This is because ongoing hardware speed and storage improvements were always sufficient to keep pace with the changing requirements of application programs. Now that AI-based software is making its mark, this is no longer the case.

Kaplan feels that "AI applications have already spawned a revolution in programming concepts and methodology. As the computer industry moves into the 1990s, they may demand a reconsideration of basic computer architecture design."

In Kaplan's opinion, "With the emergence of AI applications, the computer industry will see a significant shift in its fundamental computing requirements. These requirements will move away from the need for repetitive, algorithm-based processing of numbers and characters to the particular needs of symbolic manipulation." Instead of performing arithmetic operations or processing character strings, says Kaplan, a symbolic manipulation language like Lisp or Prolog is concerned with interpreting symbols and their relationships.

A symbol represents an object, idea, or concept. As Kaplan points out, the symbol "FRED" might represent a particular cat. "CAT" might represent a particular species of mammal. "FURRY" might represent the state of being covered with fur. With symbol manipulation, a series of relationships might be asserted such as "FRED is a CAT" and "CATS are FURRY." From these relationships, a Lisp- or Prolog-based program might logically conclude that "FRED is FURRY," much as a conventional program computes that 2 plus 2 equals 4.

According to Kaplan, existing von Neumann-based computer architectures are poorly suited to symbolic program execution, so AI software has a reputation for poor performance. "Symbol manipulation programs and standard computer programs place different demands on computing hardware," he says. "AI software just does not run effectively on standard architectures based on the classic procedure-data computation machine/software model."

Future computer hardware design must consider symbolic program attributes so that these computers can better accommodate AI applications, adds Kaplan. These attributes include "search," "pattern matching," and "evaluation." To facilitate AI applications that have these attributes, notes Kaplan, computers designed for symbolic processing may abandon von Neumann architectures in favor of more general and flexible processor arrangements, parallel computational capability (perhaps data flow architectures), and new instruction sets.

"Symbolic programs often explore a large number of potential solutions to a problem," says Kaplan. He explains that this process is known as a "search" because the program searches for a problem solution through a "space" of possibilities. It does this chore by generating each potential solution and testing it. If future hardware performs these searches in parallel (multiple processors), searches can be more effective.

To Kaplan, the solution to many symbolic manipulation problems involves the matching of some object or partial description of an object against an ideal description or template. This is often done to identify or classify an unknown object. By determining the class of an object, knowledge about that class can be used to reason about the specific object. Pattern matching usually involves scanning a potentially long list of possible descriptions for a match against the target pattern, and flagging those that are close to the pattern for further examination. Kaplan feels that significant economies could be achieved if computer architectures allowed this process to be implemented as a single associative operation rather than an iterative process.

Kaplan points out that, as opposed to a conventional programming language where variables represent a "location" in which a value can be stored—in symbolic languages, this location falls under the more general idea of "evaluation." He further explains that there may be no formal distinction between "variables" and "values" in a symbolic language—both are simply objects linked together in a symbolic relationship.

For example, the evaluation of "MAMMAL" might yield the value "CAT." Further, the value of "CAT" might be "FRED." More conventionally, the value of "TEMPERATURE" might be "75 – (3.5*ALTITUDE/1000)," which might further evaluate to "68." In a symbolic program, these relationships may form complex linked structures in contrast to the neat arrays of numbers or characters found in conventional programs.

direction, but it never caught on. (Prolog's deriving a goal is a backward deduction.)

Other issues in the great debate between Lisp and Prolog await the test of time. Just how much parallelism is needed in a language, whether parallelism is incompatible with backtracking, and whether Lisp can be redesigned to take on Prolog's better attributes (and vice versa) remain for the researchers of this generation to determine. For example, work at the University of New South Wales in Australia is aimed at altering Prolog's control structure to expedite parallelism while leaving its logic parts alone. To do this, the Australian investigators are trying to combine Prolog with data flow software mechanisms.

Parallel processing software

The ubiquitous, inexpensive microprocessor, as well as development of many VLSI-based chips for memory, peripheral device control, and communication have led to the multiprocessing computer

architecture with its emphasis on parallel processing procedures. In fact, a baker's dozen of non-von Neumann parallel processing architectures, like data flow machines, are under worldwide consideration to determine their applicability for 1990's computing needs. All of them require new software concepts and implementations geared to parallel processing.

Older software cannot do the job with parallel processing-based architectures because it is based, to a large extent, on the classic procedure-data computation machine/software model. Software for this von Neumann model is said to be program control statement driven. In other words, program instructions execute only when the program's control statements give the order. Such software draws a distinction between manipulated data and the procedures that perform the manipulation. Other software types, especially more recently developed languages, make no such distinction.

Control-, data-, and demand-driven languages are being developed to solve the programming language problem.

Parallel architecture design is a young, still-evolving field with no clearly discernible trends. Unresolved issues include how the parallel processes communicate to result in a combination of parallel computations for the program's "answers," and the number of processors that should be in parallel as well as their size, memory, and design. There are also questions of cost, the need for specialized programming languages unknown to the computing community at large, the need for specialized algorithms for parallel computation, the unknown quality of the VLSI that will be available (ie, speed and packaging), the theoretical advantages of the parallel architecture under consideration, and more.

All these concerns mitigate against any predictions of what the 1990s can expect from parallel processing-based software. Much work remains to be done. Slowing progress is the fact that most of the ongoing work is not concerned with real-world, fully constructed systems. Because of the enormity and complexity of the design problems, and, some say, because of the lack of a clearly defined and centralized, national program, only hardware and software bits and pieces are being built. Worse, these are done without the benefit of the latest hardware and software. Only a concept's feasibility, not cost effectiveness, is usually being examined.

The first four generations of computer architectures and their associated software are based on the von Neumann architecture with its single processing element, sequential CPU, sequential machine language, and linearly addressed memory. It is a tribute to the early computer system designers that their work has so much staying power. It has even been proposed by some researchers as the basis for some 1990s work. To speed up such a machine, all its parts must go faster. For many applications, this approach is limited by design and fundamental theoretical limits on physical processes.

The programming language problem is one of the worst. It is far easier to build a concurrent processor than to program one. Conventional languages just do not efficiently use parallelism. Indeed, if concurrency is to make its mark, languages and programming techniques must be geared to it. For maximum efficiency, it is believed that the programming language control method should match the architecture's control method. It is no surprise then, that there are control-driven, data-driven and demand-driven languages under development worldwide. Hopefully, the results of these efforts will be seen in the 1990s.

Many of these language efforts will also be used in AI-intensive applications like speech, vision, natural language, data bases, and graphics. To some extent, progress in these areas has been delayed by the limited processing power of von Neumann architectures and conventional programming languages. Parallel architectures offer a processing power increase that allows more practical speech systems for computer I/O, unconstrained natural language inputs to computer systems, and more. Parallel architectures combined with AI knowledge-based systems will also be used in military applications for realtime signal analysis, *Star Wars*-type weaponry development, and number crunching hydrodynamic and aerodynamic research.

Which parallel architecture and language will survive 1990 remains to be seen. If history is any guide, the "winning" architecture will be a hybrid whose design emphasis depends on the application at hand. Already, hybrid machines and the software being designed to go along with them are being developed. Typical is the combined data flow and reduction machine under development at the University of Utah in Salt Lake City. Like many other parallel architectures, it is in the early development stages, with two to four years needed before major conclusions can be drawn.

It is possible that the most viable machine will combine architecture and software in such a way that it amounts to many separate machines linked together. Each machine will have specialized architecture and software for its particular function. Complicating this, as well as other parallel architecture hardware and software design, is that it is not known how to "parallelize" an algorithm so it works with different parallel architectures. A common procedure today is to rework each algorithm or

"design to fit." This is a costly and time-consuming software development method.

One assignment

Computers with a data flow architecture are programmed with languages geared to data-driven instruction execution. These are known as single-assignment languages characterized by the fact that any variable in a program may appear on the left-hand side of only one program statement. With this programming language design, data dependencies can be easily determined and a program written in any order. Its modularity is thus maximized. New statements can be added without having to rewrite the whole program.

Another attribute of a data flow language is that it must be free from side effects. All functional or applicative languages (eg, pure Lisp) have this property. There are no software constructs that change the state of a computation. In fact, the state concept does not exist. In this approach to 1990's software, all program control is by means of mathematical functions operating on suitable arguments (in the mathematical sense of the word, argument). So, the execution of a statement has no effect on the operation of other statements. Typical of such languages are Val and Id out of MIT and the University of California at Irvine, respectively. Other languages for data flow include Lucid and FP. All these remain experimental with a limited user base in the universities and a few industrial research organizations. Which language will survive is unknown. The language Arctic, under development at Carnegie-Mellon University, has its roots in functional programming languages. However, its values are functions of time rather than constants. It is suitable for high performance, realtime control of computer systems.

Functional programming languages may make their mark in the 1990s because of several potentially useful attributes that differentiate them from conventional programming languages. Like the data flow languages, they have been thought about for over 10 years now. And, they may finally come into their own, driven by the needs of AI and parallel computation.

As mentioned, a functional (or function-oriented) programming language is one whose control structure is the application of a mathematical argument to yield a result. There is no concept of the assignment statement that is so common in imperative (procedural) languages like Fortran. There are also, as mentioned, no side effects.

Functional programming languages may make their mark in the "conventional" software that will be needed for the micro, mini, and mainframes of the 1990s. Side effects can make software difficult to write, debug, understand, maintain, and upgrade. These difficulties retard the development of the software as art into software engineering. (see "Software Development Evolves Into Software Engineering," page 169). With functional programming, some of these limitations can be overcome, and it will be practical in scientific and engineering applications where its mathematical basis is most useful.

The idea of function as a control mechanism means that both side effects and state are removed from consideration. This means that time is not a specific program dependency. As a result, it is possible to distribute the computation parts through time and evaluate several expressions simultaneously. In contrast, with an imperative language like Fortran, a rigorous order of statement execution is essential since it may alter another computation's state.

The functional programming idea (with its great modularity and statement execution independence) is made to order for parallelism. Examples include FPL, HOPE, ML (which runs under Unix), and KRC. As might be expected, if it is to be a programming language rather than "merely" a mathematical system, a functional programming language must also define how to perform computations.

Many architectures are suitable for functional programming, many have been proposed, and some small systems have been heavily studied. For example, an experimental VLSI-based chip has been designed for a Lisp dialect. Little complete work has been done and a functional programming-based computer will probably not appear before the late 1990s.

Home, office, and factory

While AI parallel processing and supercomputers will make great technological advances in the 1990s for high technology applications, it is likely that the networked home, office, and factory of the future will expand to affect everyone. The software to make this happen will deal with advances in semiconventional programming languages, operating systems, application programs, and software development tools. It will also deal with all the software technology needed to convert the art of software development into a science.

Some authorities say that software technology will advance so that programmers and software developers, as we now know them, will disappear from the scene. Perhaps, but it is clear that the societal impact of widespread, easy-to-use software for every application will be far-reaching. (see "Intelligent Computing Era Rushes In," p 78).

For the computer designer, the key issue for the software of the 1990s is the productivity of those who develop software—be it compilers, operating systems, languages, or applications. In fact, a majority of the worldwide effort in next generation

A NEW WORLD OF SIGNAL PROCESSING FOR THE IBM® PC.

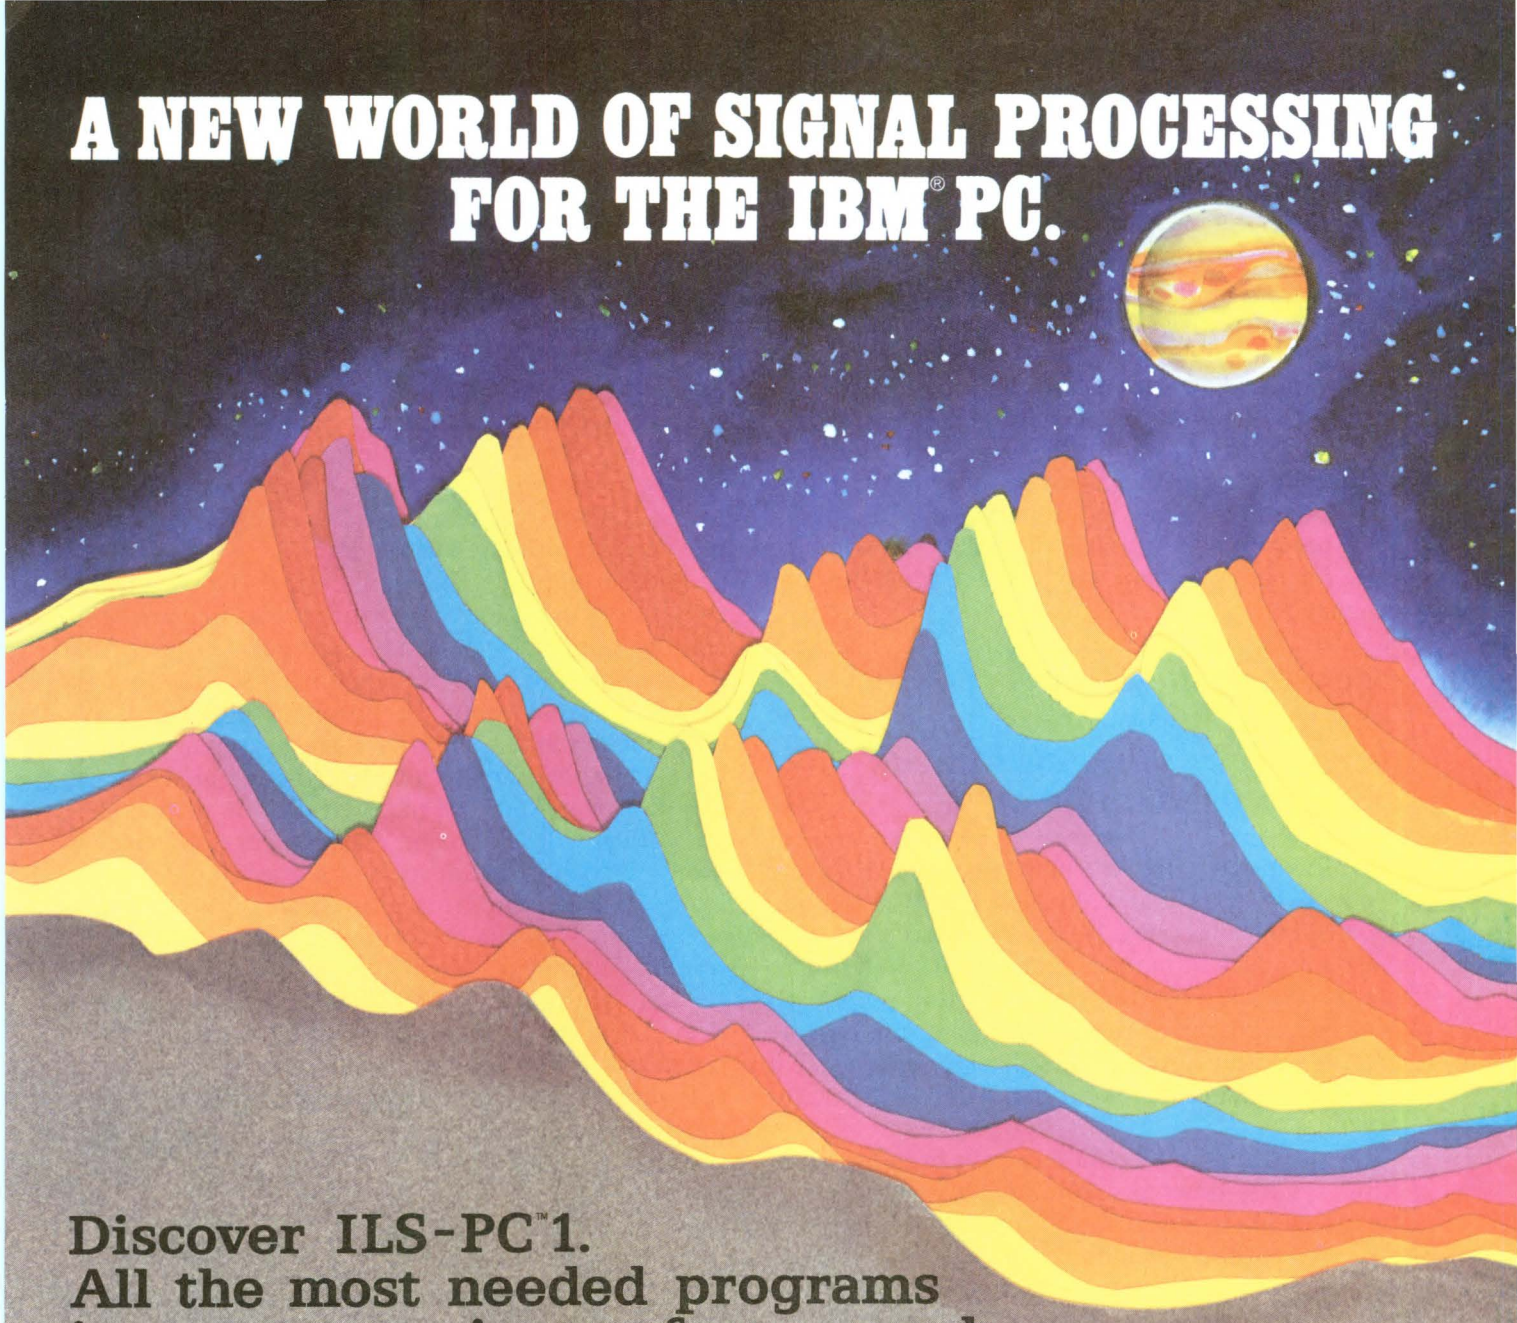

Discover ILS-PC™1. All the most needed programs in one convenient software package.

It's here. True signal processing on the IBM PC. It's called ILS-PC 1, and was developed from our minicomputer ILS, the world standard in signal processing software.

ILS-PC 1 enables you to do signal processing *now* on your own IBM PC or XT. Without writing programs. Without a lot of time or effort. When used with the 8087 coprocessor, it performs at minicomputer speed.

ILS-PC 1 provides all the essentials: data acquisition support; waveform display and editing; digital filtering; spectral analysis. Applications include noise and vibration, speech, seismology, acoustics, sonar, radar, bio-medicine and many other fields.

Part of ILS-PC 1 convenience is a menu-prompt you can use while learning, then bypass once you're ready for command control. With our customer service phone line, you get answers to software or applications questions.

Bring signal processing to your own IBM PC or XT. Call our toll-free number now for full information on ILS-PC 1. The price is $1495. The benefit is a new world of convenience.

STI Signal Technology Inc

5951 Encina Road • Goleta, CA 93117
In California, (805) 683-3771 TWX 910-334-3471

Call toll-free (800) 235-5787

Germany • TCAE GmbH, Tel. 8139/6067, Telex. 841 527523 / **Sweden** • 3K Tre Konsulter AB, Tel. 0764/30175, Telex. 854 15559 / **Switzerland** • Zentrama Technik AG, Tel. 032 233553, Telex. 845 349353 / **U.K.** • Logica, Tel. 01637 9111, Telex: 851 27200 / **Japan** • Rikei Corp., Tel. 03-345-1411, Telex. 781 123772. / **Korea** • Greenell System Industry Co., Tel. 725 6281 2, Telex. 787 23231 / **Taiwan** • Hermes Epitek Corp., Tel. 02 562 5851, Telex. 785 26794

computing is directly concerned with the means for improving the rate of debugged code production.

Efforts in this regard include software tools for designing, tracing, debugging and verifying programs. Also, efforts are being made to define techniques for making use of reusable software. High hopes are also held for automatic code generators that take a formally specified program design and automatically generate executable code from it—purportedly bug-free.

Such programming techniques as structured programming are said to improve code quality because that code is generated in a top-down style. Finally, there are a number of not easily classified compilers, and formal program design methodologies and languages.

Operating systems and database management software must be designed for a fully distributed environment.

Unfortunately, like the new languages, software productivity tools, especially the expensive and complicated high level ones, are not widely accepted. One reason is a lack of managerial and staff understanding of tool use. Another reason is that most software work consists of maintaining programs that are already written. And, the software development community perceives the tools as too difficult for practical use.

Finally, tools and techniques have been oversold. Although top-down programming was presented as a panacea, there is also bottom-up programming and a variety of other techniques for the "best" way to design. The trend seems to be that if the right format is followed, software design will be optimized. In fact, some observers say that the little available research indicates that the rate of producing debugged code is code-language and code-development procedure independent, and depends upon the skill and motivation of the programmer more than anything else.

There is hope for improvement. Some researchers say that software development tools will improve things once they are in wide use. And, these researchers add, the tools will get better once there is a demand for their improvement. This is the classic cycle of high technology hardware and software.

One approach in the move to produce more reliable software systems is the increased use of high level languages to "rapidly prototype" a software system. In this case, users determine if it meets their needs. The target software system is developed in detail. If the software system requirements are specified in some formal specification language, and some software tool automatically generate the prototype in a high level language, the process will be more efficient.

While much research is still going on in this field, there is little uniformity, little acceptance, and the actual impact on software productivity has been small. Lack of standards for the tools is a major problem—each is different, a new one may have to be learned if one's job changes, and the tool is often difficult to learn and apply for people whose jobs require solving a specific project's software problems. Worse, automatic program generation from a formal specification is, simply put, difficult to do. It is not clear when it will be an easy-to-apply, inexpensive, practical technique.

Another approach offering hope is the so-called exploratory programming environments that are becoming popular. These have the defined goal of making software development an easier task. In an exploratory environment, powerful development tools like debuggers, libraries, editors, graphics assistants, and others are provided in an integrated hardware/software package in a workstation dedicated to software development. Developers can "play around" with program segments, executable statements, and the like, to see the best way to design a program. They can perform what amounts to software development "experiments."

Many of the workstations that will aid in this endeavor will be AI based. They will also allow programmers to spend cost-effective time removed from software development bookkeeping burdens. This too is a wide open field and no one is quite sure where it is going.

Flexible software

As computer and communication technologies merge, software will be needed to allow all computer types to communicate, regardless of manufacturer. The International Standards Organization's seven-layer model for computer communications, known as the Open Systems Interconnection model, is a framework for accomplishing this feat.

But, much software must be implemented to realize the lofty goal, "every computer can talk to every other." Operating systems and database management software must be designed for a fully distributed environment with intercomputer distances ranging from inches (chip to chip) to thousands of miles (continent to continent). Of course, these interconnected computers must communicate with each other regardless of the protocols native to an individual machine. And, it would be nice if an application program could run on a machine regardless of its size, location, and manufacturer.

Today, these are software dreams, but software futurists are thinking about such accomplishments. For example, AT&T maintains that the Unix

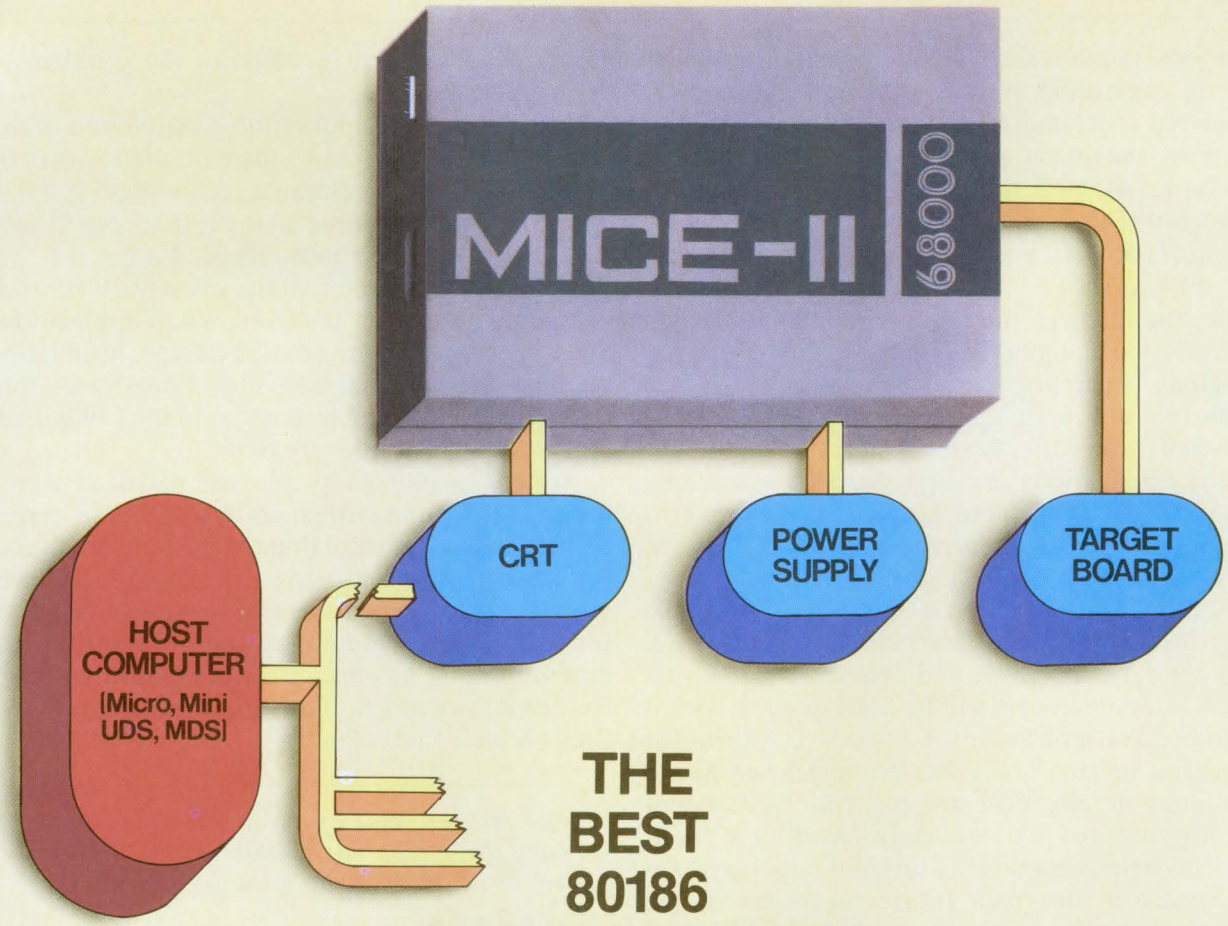

THE BEST 80186 DEVELOPMENT/ DEBUGGING TOOL YOU CAN BUY FOR $4,200

It's powerful, portable, configurable. Use it to turn your mini/micro computer into a cost-effective full-scale universal development system.

Use it to expand your full-scale development system into a multi-workstation system.

Use it with MULTIMICE to debug hardware/software in a multi processor environment.

Use it to evaluate different microprocessors at a minimum change-over cost.

It's easy to use. MICE I/O drivers are available for Apple II, IBM-PC, PDP-11, VAX, MDS, iPDS, TI Professional and all CP/M systems. And symbolic debugging is available for some models.

It gives you all these features. Real-time emulation with no wait state • Retain full memory and I/O space • Resident assembler/disassembler • Real-time forward/backward trace up to 2048 cycles • Instruction step, cycle step through program • Two hardware breakpoints • Interchangeable personality module • Up to 128K emulation memory.

It supports all these microprocessors. 8048 • 8085 • 8086/88 (MAX) • 8086/88 (MIN) • 80186/80188 • 6809/6809E • 68000 • 68008 • 68010 • 6502 • 65SCXX Series • 65SC1XX Series • Z80R • NSC 800.

And it's programmer supported. GP-256 micro based, system/gang EPROM programmer for JEDEC pinouts. MICE I/O driver or master loading.

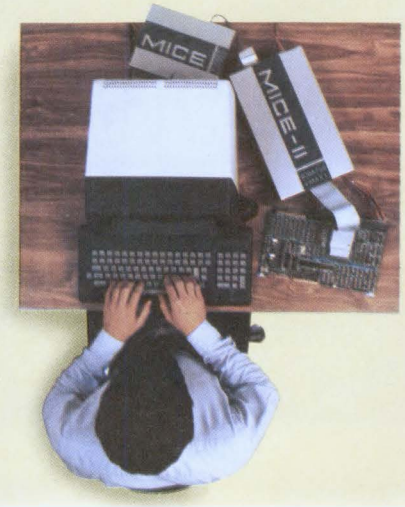

MICROTEK INTERNATIONAL, INC.
2-1 Science Road 1
Science-Based Industrial Park
Hsinchu, Taiwan, 300, R.O.C.
Telephone: (035) 772155
Telex: 32169 MICROTEK

MICROTEK LAB, INC.
17221 South Western Ave.
Gardena, CA 90247
(213) 538-5369

CIRCLE 59

operating system is the software portability answer both for operating systems and application programs. No other contender seems to be on the way and, given the time it would take to develop one for 1990 (at least), Unix may have a commanding lead by default. This is especially true if IBM pushes it. Indications are that it will do (see *Computer Design*, Aug 1984, p 165).

The computers of the 1990s will be programmed in a variety of languages. There is no doubt, even though all concerned parties bemoan the language babble that now exists. Smart money says that matters will be worse, not better. The facts remain that specific applications are best done in a specially designed language, new languages continue to appear, software people have personal preferences, and so satisfactory-to-all language is on the horizon.

For military applications in the United States, the Department of Defense-mandated Ada language will be in high gear by 1990, and computer system designers and integrators will be used to work with it for embedded applications. Under development for five years, its first compilers are appearing as are its program development environments. The amount of money invested in it, and DoD's need for a "standard" to help overcome its internal language babble, ensure a rosy Ada future regardless of the language's shortcomings. These shortcomings have been pointed out in great detail by a software community that is eager to pinpoint the problems in every language.

With such a large application base (billions of dollars), Fortran and Cobol will go on their way in the 1990s. With any foreseeable technology, it would not be possible to rewrite all the extant code in these languages. For their part, Basic, Pascal, PL-1, C, APL, Lisp, Prolog, and all the others will continue to be used—each for its defined set of applications. Only the usage percentage will change. Modula-2—the language with the traits that Pascal was "supposed to have"—may be popular in the 1990s if current interest continues to grow.

Assembler—cursed by all—will persist. The need for speed in time-critical code sections in certain applications will insure that. Moreover, much code for government-sponsored work is written in assembler. And, of course, there are a variety of languages for functional programming and other special programming that has been mentioned in this report. No one knows which, if any, of these languages will ever become "popular."

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

High 710 *Average 711* *Low 712*

SOFTWARE DEVELOPMENT EVOLVES INTO SOFTWARE ENGINEERING

Automated tools will reduce costs, increase reliability, and build better programs in a structured manner.

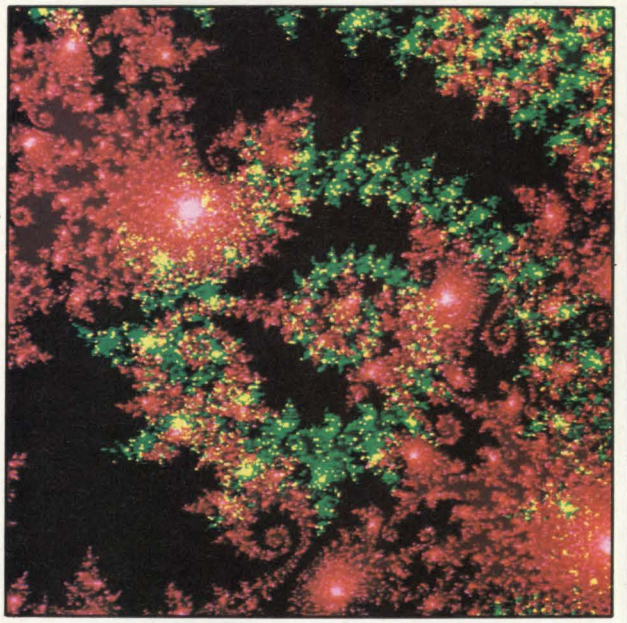

Photo courtesy of Dr John Hubbard and Homer Smith, Cornell University

by Ramachendra P. Batni

Progress in software engineering—the science of producing economical, reliable software—will occur over the rest of this decade, as software development evolves from an art into a science. Of all the efforts contributing to software engineering, standardization, specification, and testing are the weakest. Standardization is slowly occurring, driven by the need to make software portable and reusable (to overcome escalating development costs and manpower shortages). However, much must be done with specifications and testing before software engineering is a reality.

In order to develop software, requirements must be specified, architectural features must be defined, functions must be specified, and programming, testing, and integration must be done. The software development process requires repeated iteration of these tasks until system tests are successfully completed. Fig 1 indicates the evolution of certain key software development factors over the rest of this decade.

Various software architectural concepts are emerging for developing portable, reusable, and maintainable software. These include the virtual machine (to

National's Series 32000.™

National's Series 32000™ (formerly known as NS16000) provides the most complete family of 32-bit CPUs, slave processors, systems peripherals, development systems and software. Today!

In today's fast paced high-technology marketplace, time equals money. And chances for product success often depend upon narrow windows of opportunity. So time-to-market becomes a critical ingredient in microprocessor selection.

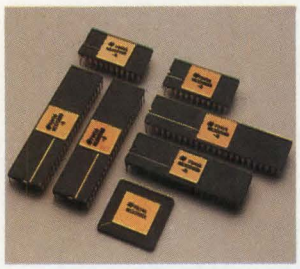

When you thoroughly evaluate and compare available 32-bit options, we think you'll find that National's Series 32000 is the undisputed choice.

Full 32-bit Architecture.

In a world of over-stretched, band-aided microprocessor architectures, National provides a true 32-bit solution. The Series 32000 is the world's first commercially-available 32-bit microprocessor family that was designed as 32-bit. We've taken a systems solution approach to give you what you need: An innovative microprocessor family with the right performance. Current availability. And alternate sourcing by major manufacturers.

- *Advanced 32-bit architecture* is the basis of the Series 32000. This same 32-bit architecture is implemented in all Series 32000 CPUs. These CPUs also offer downward compatibility in 16 and 8-bit external data bus versions.
- *Fast floating point support* is available now with National's NS 32081, the *only* viable floating point unit to support your system's needs for floating point operations.
- *A powerful MMU, National's NS 32082, provides Demand Paged Virtual Memory support* and it's available today. This industry 'first' provides fast on-chip address translation and advanced operating systems support, making it the ideal choice for a UNIX™ environment.

The best UNIX Micro.

If your system design calls for UNIX, industry experts have called the Series 32000 "the best for UNIX."

National's GENIX,™ featuring optimized Demand Paged Virtual Memory support, is available today in source form for OEM adaption. The next generation of GENIX, with Ethernet™ support, is in test now. We have also delivered the Series 32000 version of UNIX System V (release 2.0) to AT&T for validation. National is your source for UNIX.

Our Series 32000 software catalog gives OEMs an opportunity to select from a wide variety of software, including offerings from independent software vendors.

you can wait for somebody else's.

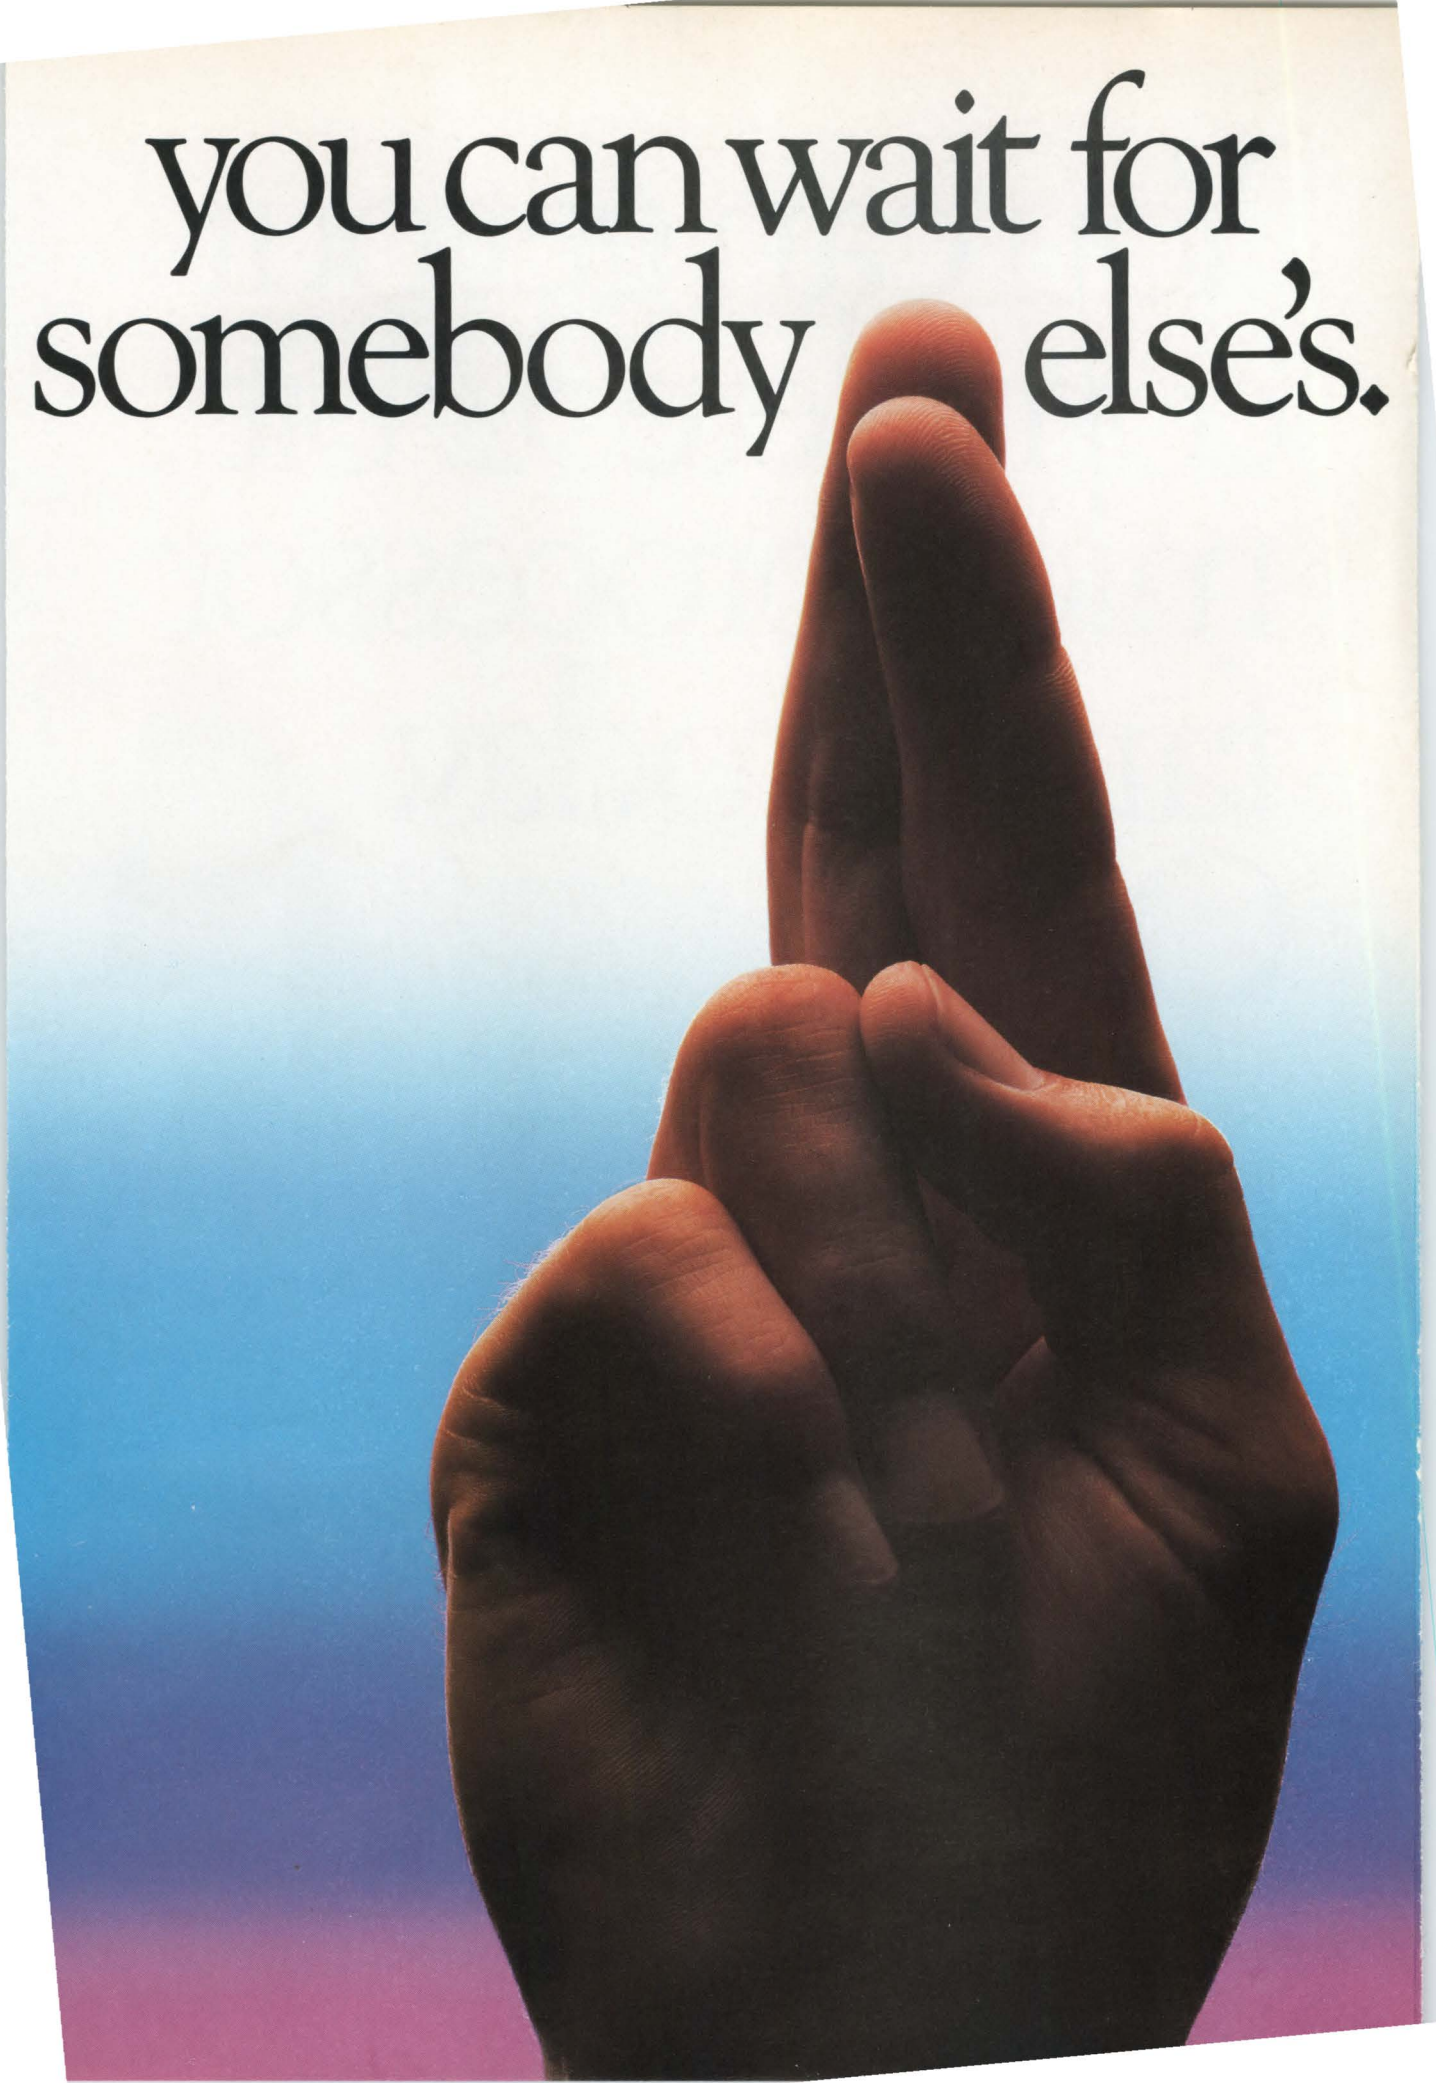

You can get our complete 32-bit microprocessor family today. Or...

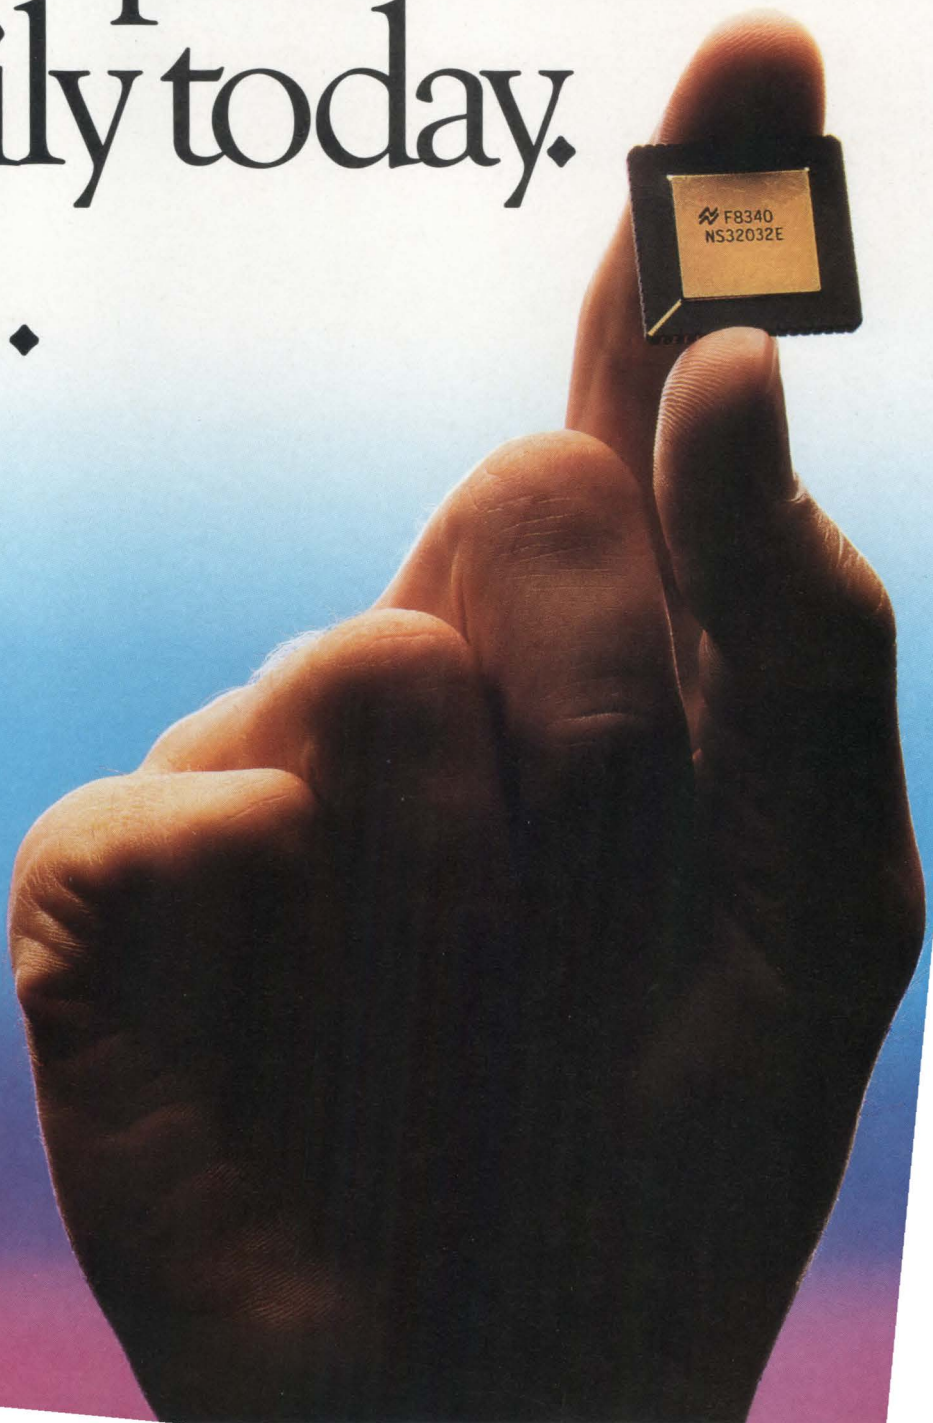

The 32-bit solution. We've got it!

Series 32000 is the only 32-bit solution today.

If your market window falls within the next twenty-four months, the Series 32000 provides you with product for your designs *today*. National is the only manufacturer that offers a complete 32-bit microprocessor family now. With an architecture designed to fully support high-level languages; a full 32-bit data bus to memory that increases memory bus bandwidth and thus the speed at which data can be transferred; a range of CPUs with total compatibility to 16 and 8-bit; a full range of peripherals; powerful evaluation tools; a multi-user, multi-tasking development system; a growing list of available third-party software; and the hardware, software, service and customer training that spell real long-range commitment.

All to assure you of the products and support you need *today* and into the future.

Truly a generation ahead.

Ask us for complete information.

We'd like to send you a copy of our brochure entitled: "The Specifics of 32-Bit Architecture and Implementation." And to answer any questions you might have about Series 32000. Write to National Semiconductor, Series 32000, 2900 Semiconductor Drive, MS23-200, Santa Clara, California 95051.

Series 32000.
Elegance is everything.

UNIX is a registered trademark of AT&T – Bell Laboratories.
Ethernet is a registered trademark of Xerox Corporation.
GENIX and Series 32000 are trademarks of National Semiconductor Corporation.

National Semiconductor
MICROCOMPUTER SYSTEMS DIVISION.

CIRCLE 61

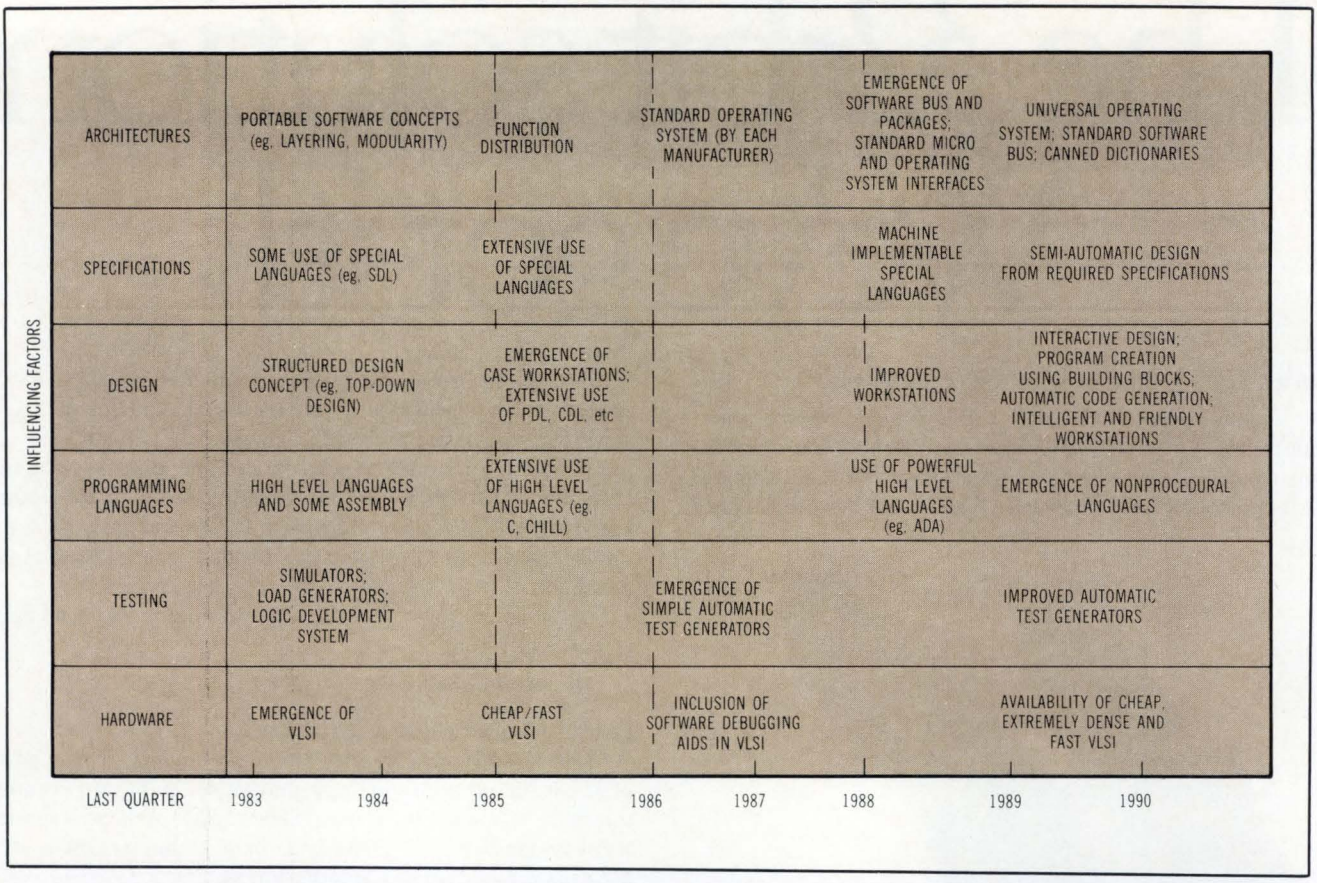

Fig 1 Individual trends in the areas of software and hardware will synergistically influence the evolution of software development into software engineering. The precise timetable for each of these events may be argued, but the general trend should become clear as the year 1990 approaches.

protect software from hardware changes) and information hiding through layering of software functions. Layering protects high level software functions from changes in low level software implementations.

Another concept provides functional independence between major software elements through formalized interfaces, allowing software design teams to set up and test system functions easily and independently. Software designers may also allow for function migration between processors through module interface standardization, common operating system functions, and loose coupling between modules. With formalized and flexible interfaces, modularity and loose coupling allow system growth from a single processor to a multiprocessor configuration with minimal software changes.

These ideas lead to the software bus concept (a standard interface for software runtime and development environments). Such buses will first occur for a specific manufacturer's products. True standardization (among all manufacturers) will occur much later, if at all. It is usually not in the economic interest of manufacturers (especially chip makers) to make software portable to outside machines.

Once standard interfaces are defined, different organizations can develop software packages that will work together via a common standardized operating system. Thus, operating system and software bus standardization will lead to the creation of portable and reusable software plug-in modules.

Function distribution is yet another architectural trend that will grow in this decade. Function distribution occurs because of VLSI's dramatic hardware cost reductions. For example, the increasing availability of cheaper and more powerful microprocessors allows the implementation of distributed computer system architectures.

Distributed architecture, in turn, affects software architecture. For example, software becomes modularized into "functions," and functional packages become distributed in various processors. Thus, because a processor may not contain all functions within itself, it requests other network processors to provide the needed services. Interface and message protocol standardization is necessary to control and to manage the complexity introduced by function distribution.

Tying software together

The operating system software package is a key element of any software system. It ties together the other major software elements and provides services like scheduling and memory management. Operating systems tend to be sophisticated and complex.

WHO'S GOT DIBS ON THE DWG?

Digital Word Generators are always in demand.

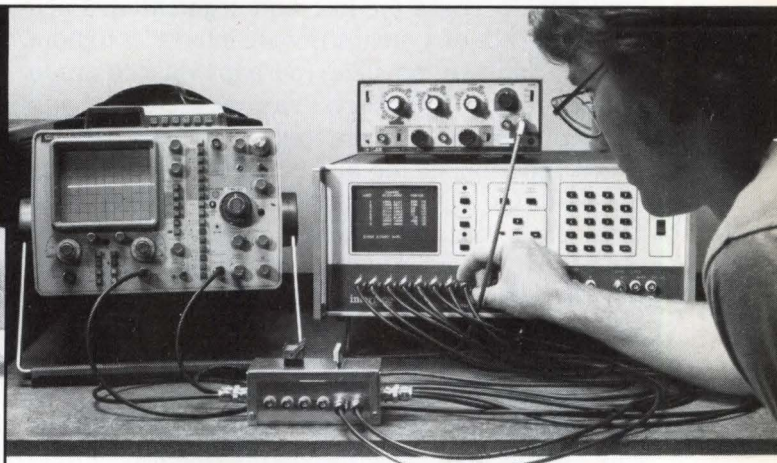

DWGs speed up IC characterization.

DWGs simulate prototype environments.

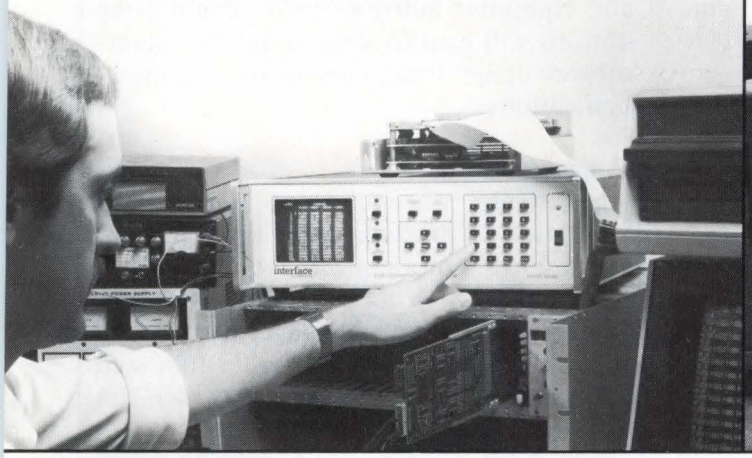

DWGs eliminate building dedicated circuitry.

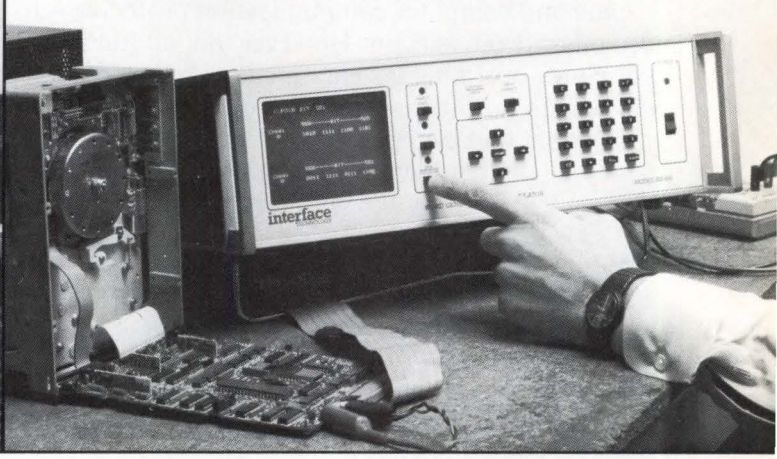

You never know where a Digital Word Generator from Interface Technology will show up next.

From testing discrete designs before system integration to simulating prototype environments, these DWGs can handle it all. Speeds to 100 MHz and up to 4096 bits of memory on each of 16 channels make the RS-660 and RS-680 the most powerful and versatile instruments of their type.

With Interface Technology's unique timing simulation capability, you can generate variable width pulses or timing patterns, thus conserving time and memory space. Data may then be output in a single step, burst or continuous mode.

Both the RS-660 and RS-680 have integral CRTs. The ability to display a large portion of your serial data stream or parallel pattern makes either instrument easy to use. Use is further simplified by the choice of either hex or binary display; combined with a block memory fill for repetitive pattern entry.

Other features such as a programmable sync output for triggering, remote interface, and up to a 64K deep serial memory just add to an already powerful instrument.

So put an Interface Technology RS-660 or RS-680 DWG to work at your place—all over your place. Call today to put "dibs" on your DWG. (714) 592-2971.

interface TECHNOLOGY

If it's digital, we can test it.

A Dynatech Company, 150 East Arrow Highway, San Dimas, CA 91773

CIRCLE 62

Therefore, development of a bug-free system is time-consuming and expensive. Operating development trends have a direct impact on the cost and productivity of software development.

Generally, operating systems are built of layered functions. Past operating systems were onion-like in that each layer was built upon, and totally contained in the lower layers. This structure is changing. The layers have become pie-like and segmented. Each segment or slice of system software interfaces to both a nucleus and to other slices (perhaps via a common software bus). The slices can be selected and arranged to fit an application.

Such function slicing makes the system architecture modular with standardized interfaces, and allows easy portability and configurability. One example of operating system standardization is Intel's Universal Development Interface (UDI) and Universal Runtime Interface (URI) standards.

Decreasing hardware costs and operating system standardization permit software functions to migrate to faster silicon implementation. Memory management is the first silicon-based function. Other migrating functions include timing, interrupt control, process context switching, and interprocessor communication.

Data and computer communication protocols will also migrate into silicon. However, not all functions will. Some operating system functions and extensions will remain in software to allow configurability, customization, and incorporation of improvements. For example Intel's iRMX 86 operating system and Hunter & Ready's VRTX implement only their kernels in silicon.

While operating systems standardize, the standardization of the Microprocessor Operating System Interface (MOSI) is also evolving, albeit slowly. Standardization in this domain permits chip makers and system designers to work independently. This, in turn, leads to greater availability of standard operating system packages.

Operating systems to go

As this decade progresses, designers will be able to select and adopt "off-the-shelf" operating systems for their applications, instead of having to develop a dedicated operating system from scratch. Decreased memory costs and increased VLSI speeds, coupled with operating system standardization, may lead to the creation of a universal operating system. This standard, high performance, user-friendly system will be able to be used economically in all applications. Operating system standardization will reduce software development costs by directly reducing the operating system development effort. In addition, it will allow increased software reusability by permitting application programs that are written for one machine to run on another machine with little or no change.

Specifications play a critical role in software development because they describe the software requirements to be met. A major part of software development cost and effort is spent on a project's conception, definition, and specification. Unfortunately, the best way to generate complete, unambiguous requirements has not yet been found. Software decomposition techniques suitable for specifications tend to be tedious and time-consuming.

On the other hand, languages such as CCITT's Specification and Description Language (SDL), particularly applicable to realtime requirements, are becoming popular. This is because they make software program specification and conception easier and somewhat portable. And, since conception and specification represent a large portion of software manufacturing costs, making specifications portable, reusable, and easier to develop is vital to software cost reduction.

In the latter part of the decade, there will be machine-implementable specification languages, which will increase the efficiency of producing and manipulating requirement specifications. The emergence of machine-implementable specification languages and computer aided software engineering (CASE) stations will lead to semi-automatic (interactive) software design from machine-based requirement specifications.

Automatic generation

Structured design concepts (such as top-down analysis) are used to produce reliable software at reduced costs. Furthermore, machine-manipulative design languages, [Program Design Languages (PDL) and Call Description Languages (CDL)], will become popular. These languages are a first step toward automatic program generation, which can improve software productivity and reliability.

In addition, CASE stations are appearing. CASE stations speed software design by providing interactive graphics menus for construction of hierarchical system diagrams and flowcharts. Tools for other phases of software development activities, such as project control and management, will gradually be included in these stations. Since CASE stations are new, the problem of many different and dissimilar products will exist for some time.

The emergence of extremely dense, fast, and cheap VLSI will lead to the creation of intelligent, friendly CASE stations. Such machines, coupled with software standardization and the emergence of non-procedural languages, will enable interactive design and software coding using software building blocks.

The trend in software coding is to write code in high level languages. In addition to simplifying

WITH OUR REAL-TIME OPERATING SYSTEM AND YOUR IBM® PC, YOU COULD CONTROL THE WORLD.

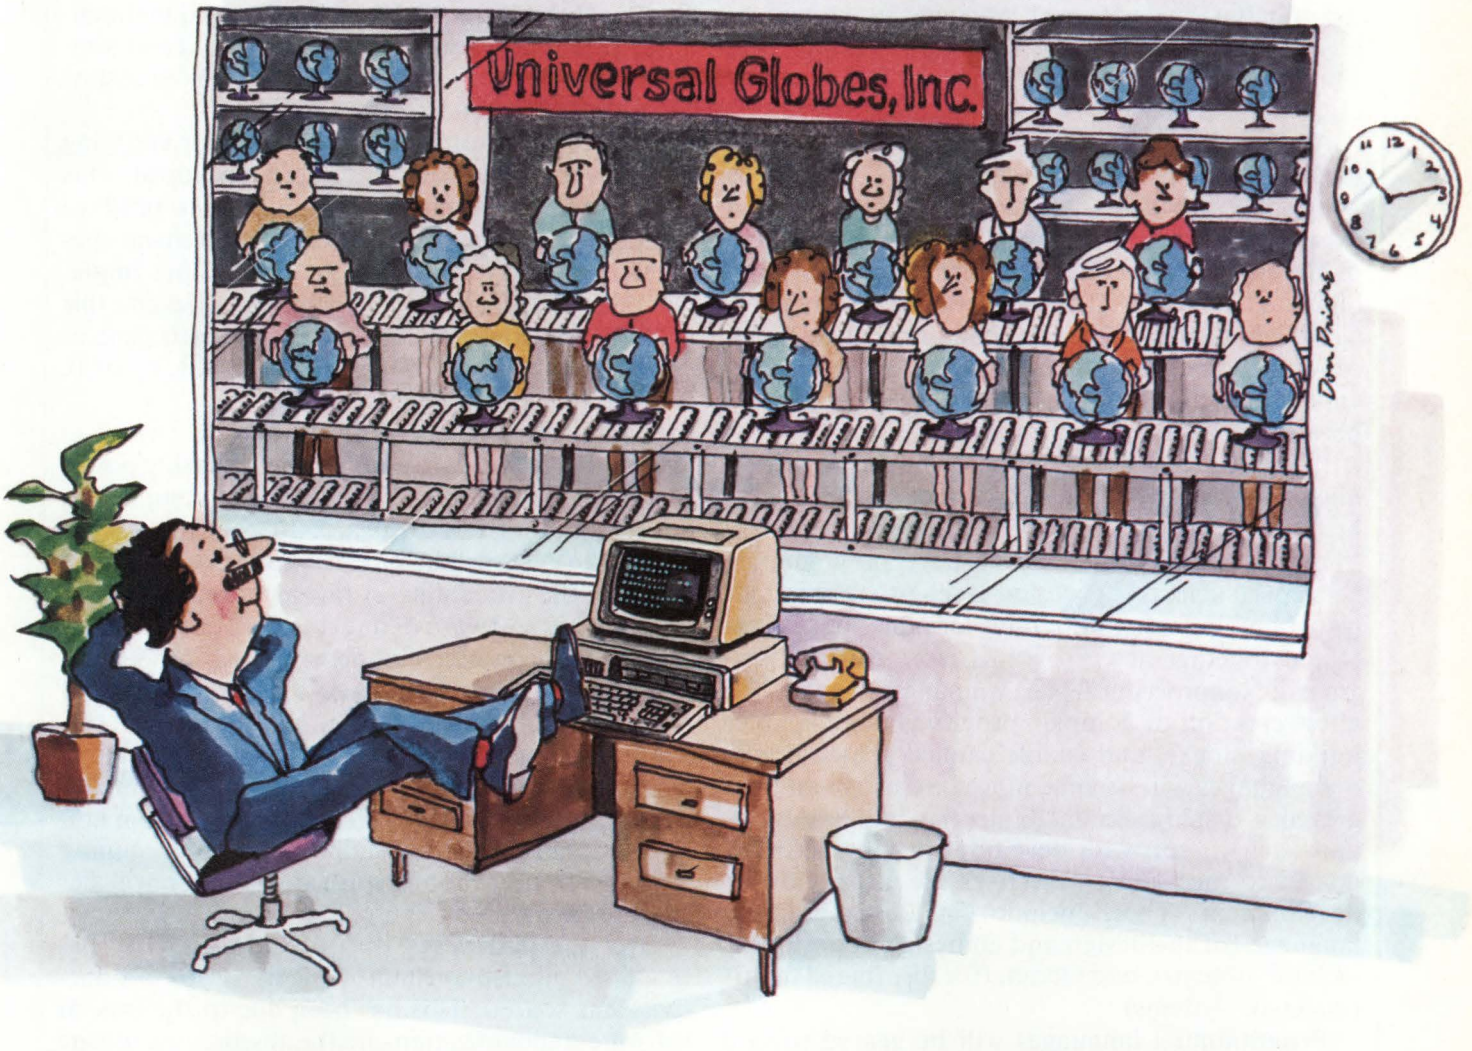

Factory automation is just one of the many control and data acquisition applications you can run with our MTOS-86/PC operating system and your IBM® PC (or equivalent). Chemical process control is another. Alarm reporting systems are a third. The list is virtually endless.

MTOS-86/PC is the newest member of our MTOS-86 family of real-time multitasking operating systems. The fastest, most efficient O/S software on the market today.

Within your IBM PC, MTOS-86/PC fully controls the PC and its peripherals, including keyboard, disk drives, timer, interrupt controller and process input/output systems. It allows custom drivers to be attached for special hardware and can easily be configured for your application.

MTOS-86/PC comes on diskette in the PC format. It includes executive, console driver, debugger, high-speed file system, a run-time interface for the Intel® C compiler and a test/demonstration program.

Price for a single copy, including user manuals, is $495. Inexpensive technical support is available, too.

For more information, contact Industrial Programming Inc., 100 Jericho Quadrangle, Jericho, NY 11753. (516) 938-6600. Telex: 429808 (ITT).

IBM is a registered trademark of International Business Machines Corporation.
Intel is a registered trademark of Intel Corporation.

CIRCLE 63

ipi Industrial Programming Inc.

The standard-setter in operating system software

programming, such coding produces fewer software errors. Furthermore, software written in a high level language is more portable and reusable since it can, with some changes, be compiled and run on machines supporting that language.

The role of high level languages

The use of a high level language requires more memory and execution time. Therefore, time-critical code is disassembled and optimized by rewriting in assembly language. Hence, assembly level programming will persist, although its use will diminish. As this decade progresses, the speed and memory penalties inherent in high level language usage will be lessened by the increasing VLSI speed and density. Thus, nearly all developed software will use only high level languages.

Among these languages, C (partly as a consequence of the growing popularity of Unix and its work tools) and Ada (because of its many advanced capabilities and the backing of the U.S. Department of Defense) will become popular, general purpose, high level languages. The CCITT high level language, Chill, will be prevalent in telecommunication applications. In addition, Pascal, Fortran, and PL/1 derivative languages (PL/M), will continue to be used because of their existing program base.

In the longer term, "nonprocedural" programming languages will be invented, allowing what must be done to be specified without having to specify "how." These languages include novel programming concepts such as "function-level programming" (programming using "program forming" operations to put together new programs from existing programs). Nonprocedural programming languages will allow creation of complex programs from simpler building blocks, and enable canned program use.

Languages will become more interactive and use graphics displays to facilitate the building-block approach to software development. In fact, languages to meet these objectives are already under development. These include Galileo (a graphics language for the design and analysis of concurrent realtime systems), and Occam (for distributed multiprocessor systems).

Programming languages will be geared toward natural languages. The advent of natural languages for database query (eg, Artificial Intelligence Corp's Intellect language) is one application of this trend. Intellect translates typed English requests into formal database query languages, locates and organizes the requested information, and presents its findings to a user. Even if requests are written in diverse ways, Intellect can respond to them. Natural language processing (a fundamental technology within the realm of artificial intelligence) has just begun moving out of the research laboratories.

Software testing activities include module testing, software integration, software validation (acceptance testing) and field trials. Similar to program conception and specifications, testing and verification account for much of today's software development costs. Yet, testing is not as developed as other areas of software technology.

Simulators, load generators, logic development systems, and other debugging aids are used in software testing and debugging. However, these tools cannot automatically verify program correctness using requirement specifications as input. Such test automation or alternative solutions will be necessary to control the cost of software testing and maintenance. In the next few years, the emergence of simple automatic test generators will reduce the cost of software testing and debugging.

Completely automatic test methods for verifying program correctness have not been developed. This problem is compounded by an inability to produce comprehensive and unambiguous requirement specifications today. The advent of machine implementable specification languages will alleviate this problem. Another solution is the automatic generation of error-free code. This may be possible by 1990.

Hardware/software interaction

The need to use hardware's ever-increasing power and capabilities has increased the complexity, growth, and cost of software. In contrast, future hardware (VLSI) will help reduce software costs. For example, the availability of fast and cheap VLSI will soon enable software debugging and checking circuits to be economically incorporated in a chip.

Moreover, the availability of cheaper, denser, and faster memories will make program development systems more intelligent and friendly. This in turn will improve program development efficiency. Finally, changes in computer architectures (away from von Neumann machines) will support new programming concepts. In short, future software evolution will dictate hardware architecture.

Over the past years, software productivity has remained almost constant while hardware productivity has soared. This has been due to the lack of software standardization and the absence of software building blocks. Currently, the major cost of building systems is not in the hardware, but in the cost of developing and maintaining software. Software is seldom reused, and usually reinvented.

Most of the productivity improvement effort has been directed at refining existing software development techniques and processes. Computer aided development techniques and high level languages appear to be reducing software development costs significantly. By the end of this decade, continued improvements in software automation techniques,

WE SOLVE PROBLEMS!

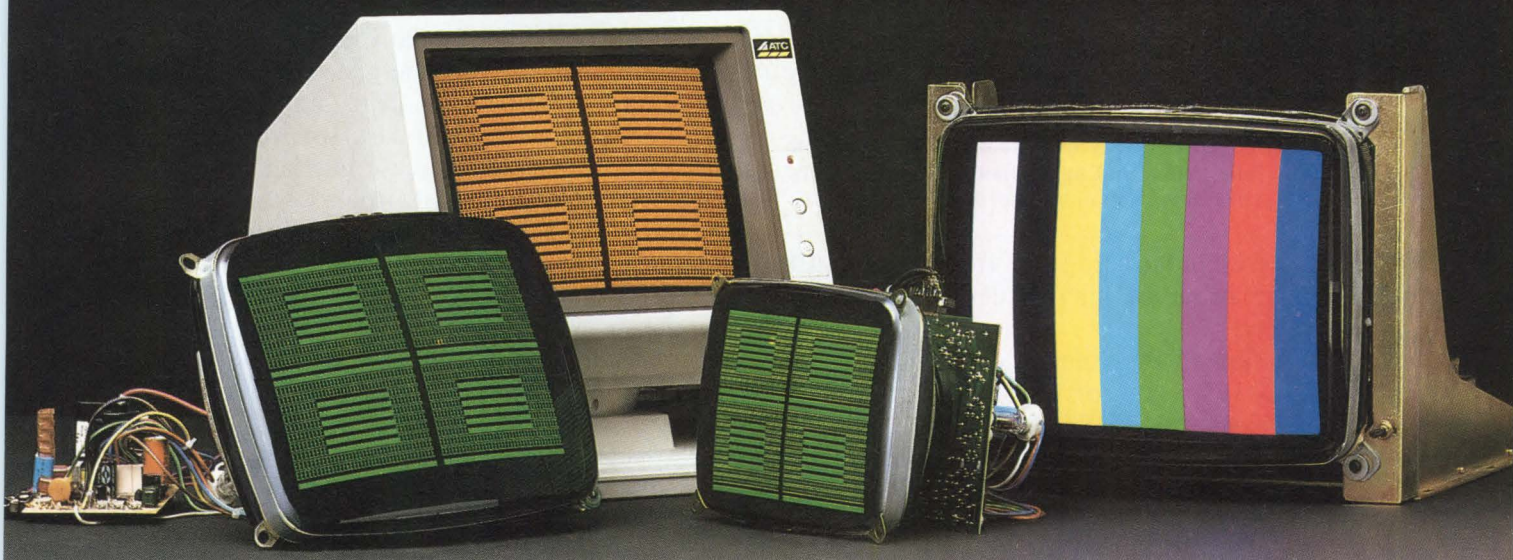

In all sizes, shapes, colors and configurations.

You need high performance, innovative design data displays. Contact Audiotronics. We have what you need. In 3″, 5″, 7″, 9″, 12″, 14″, 15″, 23″ and 5″ × 9″, monochrome and color, integrated (neck-mounted), chassis, kit or cabinet. And if we don't have what you need, we'll design it for you.

Using our basic displays, our engineers become your engineers. They custom design a display for your specific application, meeting your particular system design requirements.

Just give us your specifications and we'll solve all your display problems. And we'll deliver on time, when you want them.

Audiotronics has been solving problems for almost 30 years, designing thousands of custom data displays for important customers like you, both large and small. Call us today. Turn our engineers loose with your display system problems.

AUDIOTRONICS

North Hollywood
California 91605
(818) 781-6700

CIRCLE 64

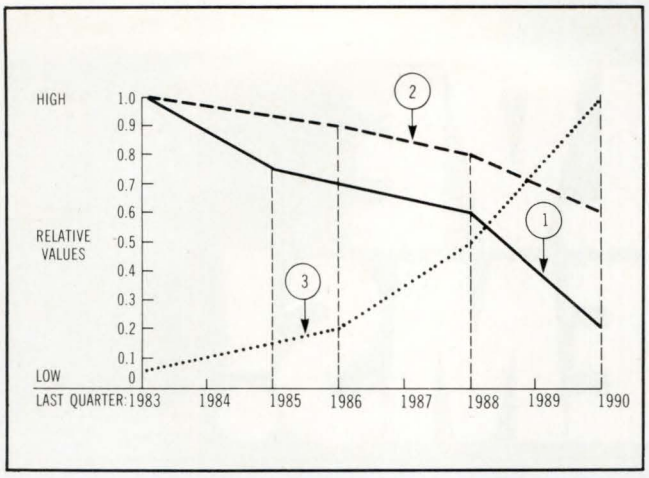

Fig 2 Three software cost metrics indicate software cost changes over the rest of this decade. Curve 1 shows the relative production cost per line of machine code and indicates the relative complexity of software production. Curve 2 shows relative life cycle costs (cost of production plus testing, integration, and maintenance). Curve 3 plots the relative availability of standard software packages.

coupled with standardized software building blocks, promise an order of magnitude improvement in productivity.

It is possible to hypothesize about the effects of these trends (Fig 1) on software development costs, using on three metrics. The first metric is the relative production cost per line of machine code (defined as the costs incurred in architectural definition specification, design, and programming). The second is the relative life cycle software cost (defined as the cost of production plus the additional costs incurred in testing, integration, and maintenance). Third is the relative availability of standard software packages. Relative comparisons must be made since the absolute cost of software development varies depending on the application (Fig 2).

Clearly, the emergence of CASE workstations, the extensive use of high level languages, and the use of specification languages will combine to reduce the cost of software production over the next few years. Improvements in CASE stations and the advent of machine-implementable specification languages will further decrease costs. However, the increasing complexity of software systems and the rising salaries of software professionals will increase software production costs. At the same time, the advent of software standardization and the emergence of nonprocedural languages will lead to a building block approach to software production. This will reduce software development costs significantly. The building block approach to software development will also make the software development process simpler and enable less trained personnel to develop software.

The cumulative effect of these varying, and sometimes opposing, cost trends will be to reduce software production costs. In all likelihood, software production costs will be only 20 percent of today's costs by the end of the decade (mostly because of software standardization).

The life cycle cost of software (which includes the cost of software production, testing, integration, and maintenance) will follow the reduction in software production costs. However, it will not be as great because of the lack of automatic software verification tools and methods, the increasing complexity of software systems, and the fact that the costs of fixing bugs detected in the field are much greater than fixing those detected during development. By 1990, the life cycle cost of software may decrease to only 40 percent of what it is today.

The availability of standard software packages will probably soar with the emergence and evolution of software standardization. Thus, unlike today when there are only a few standard software packages, there will be a great increase in the availability of standard packages. Most progress will occur near the end of the decade, after standardization has been realized.

Software development will evolve from an art to a science with consequent cost reduction. As programming languages tend toward natural languages, and advances occur in speech recognition and artificial intelligence, "verbal programming" will appear. These advances, combined with CASE workstations and automatic code generators, will reduce the complexity of software development.

Ultimately, software development will be done by specifying requirements in a language akin to natural languages (most probably verbally), using an interactive computer-based system, followed by automatic machine generation of error-free code. Such programming ease, combined with advances in the cost and capabilities of hardware, will make computers more economical and attractive to use in home, office, and factory applications.

Ramachendra P. Batni is engineering project leader at ITT Telecom, Business and Consumer Communications Division, 2912 Wake Forest Rd, Raleigh, NC 27611. Dr Batni has a BE in electrical engineering from Burdwan University, Burdwan India, and both an MS and a PhD in electrical engineering from the University of Wisconsin.

FREE VAX

If you're trying to support more than 4 programmers doing 8086- or 8080-family software with stand-alone microprocessor development systems, here's a way to increase your productivity 25%:

Buy our MicroSET-86™ or MicroSET-80™ software cross-development tools and a DEC VAX or DG MV/Family system to run them on. You'll save enough on your first major project to pay for the computer system.

Of course, if you already have access to a VAX (or MV or any IBM-compatible mainframe), you're that much further ahead.

People with experience in cross-development know it's the best way to tackle big software projects. And only First Systems' MicroSET gives you the industry's most complete, most powerful cross-development tools, plus your choice of five efficient systems implementation languages.

THE ONLY COMPLETE MULTILINGUAL DEVELOPMENT SYSTEM FOR THE 8086 FAMILY

MicroSET-86 lets you program in four powerful high-level languages: Pascal, C-86, FORTRAN-86, and PL/M-86.

An assembler is also available, although First Systems compilers generate such fast, compact code that you may find you have little need for it (benchmark data available on request). You also get a linker/locater that allows combining modules written in different languages.

Using our high-level Cross-Debugger to control execution on the target system, you can locate, fix, and re-test an erroneous line of code in a fraction of the time it takes using other methods. It's no wonder that MicroSET users routinely shave months off typical development schedules — while producing more reliable, more maintainable systems.

MicroSET-86 supports the entire 8086 family: 8086/8088/186/286 and 8087. It works with in-circuit emulators you may already own, and has a complete set of host-target communication programs and other valuable utilities not available elsewhere.

ANNOUNCING NEW LANGUAGE-COMPATIBLE 8-BIT TOOLS

First Systems now supports cross-development of 8080/8085/Z-80 software with MicroSET-80, consisting of PL/M-80 compiler, assembler, linker/locater, cross-debugger, plus utilities.

Request complete information today. And take the first step toward your own free computer.

First Systems Corporation™

Products for productivity.

865 Manhattan Beach Blvd.,
Manhattan Beach, CA 90266
(213) 546-5581
TELEX: 298086 FIRSYS

CIRCLE 65

INTRODUCING A BETTER APPROACH TO ARMS CONTROL.

It's called the Intel approach.

Our 286/310 multitasking supermicro system. Teamed with our iRMX™86 real time operating system.

Together they provide OEMs in the industrial control and factory automation market with a system that's remarkably fast and flexible.

First, let's talk fast.

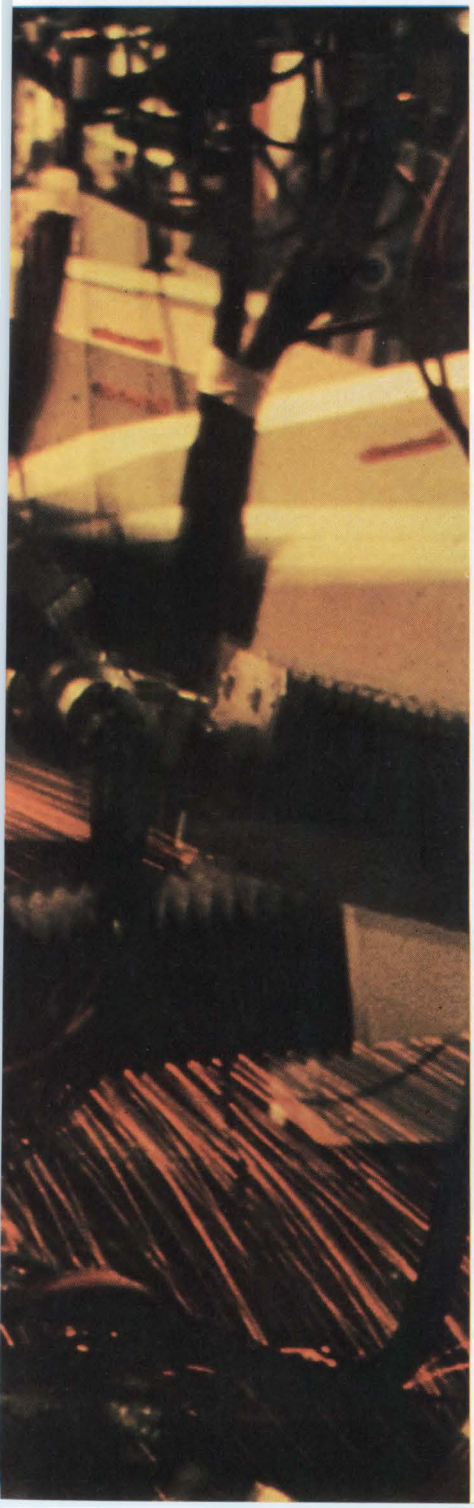

Our iRMX 86 operating system features priority driven interrupt management with ultra-fast context switching. And an event driven nucleus that manages multiple system tasks. All of which makes for lightning-fast real time response.

Then there's the 286/310 supermicro itself. Based on the most powerful high performance microprocessor in its class, our iAPX 286. Add to that the turbocharging talents of our 80287 floating point processor, which boosts numeric capabilities by 100x.

And add to all that our multiprocessing architecture. Which greatly enhances system speed by off-loading the CPU.

The result is scorching, minicomputer performance. Without a scorching, minicomputer price.

Now, let's consider flexibility.

Our iRMX 86 operating system is flexible because it's modular. So you only have to add those capabilities your customer's application demands.

It also supports all of the popular languages (like FORTRAN, Pascal, PL/M, BASIC and C).

Making it flexible enough to accommodate the most diverse applications. A fact that over 2000 licensed RMX OEMs can attest to.

But the 286/310 supermicro is not only flexible, it's open. Giving you rapid access to the latest VLSI technology. Like our new BITBUS™ interconnect which provides low-cost, easy to implement distributed control in the factory.

And our open design allows easy integration of over 1200 different MULTIBUS™ boards available from over 200 suppliers. Letting you plug in I/O, data collection, control and communications capabilities.

So that's what the 286/310 real time supermicro system is all about.

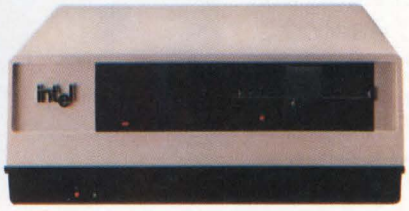

Intel's 286/310 real time supermicro starts at under $10,000.

Speed. Flexibility. Simplicity. And the world's most advanced VLSI technology. Letting you easily combine chips, boards and systems — whatever you need. Plus complete maintenance. Even for the non-Intel parts of your system.

Call us toll-free for a copy of "How to Select a VLSI Operating System" and more information on the 286/310. (800) 538-1876. In California, (800) 672-1833. Or write Intel, Lit. Dept. #S 11, 3065 Bowers Ave., Santa Clara, CA 95051.

We'll be happy to give you a hand in winning the arms race.

intel®

CIRCLE 66

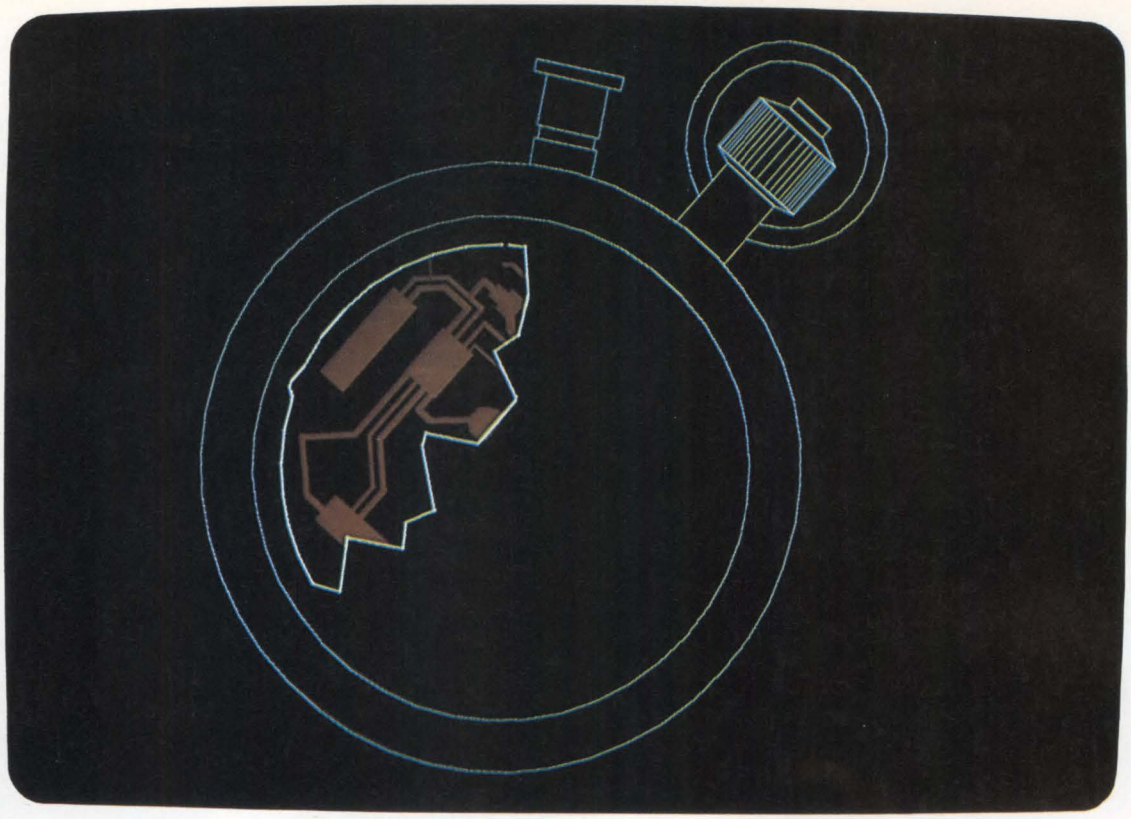

An average terminal gets this far in 1.4 seconds. You'll twiddle your thumbs for 14 seconds before it's finished.

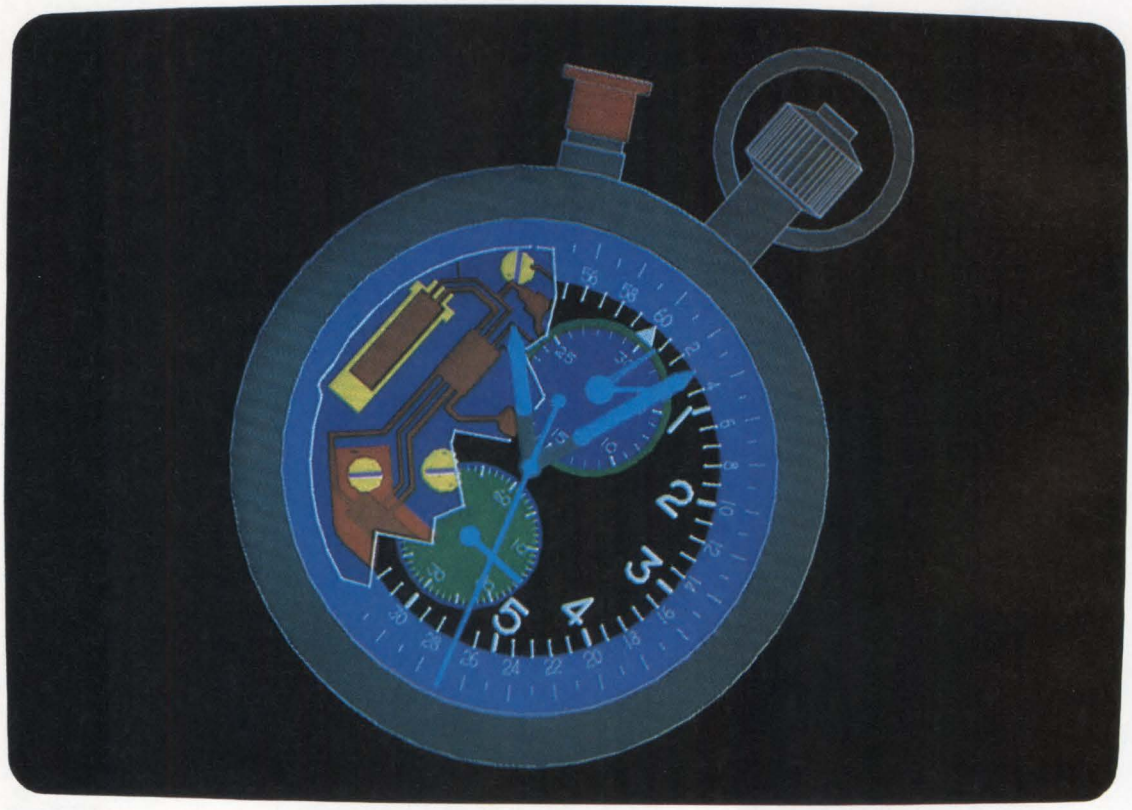

New HiSCAN terminals complete the entire drawing in just over one second.

Graphics terminals with ten times the drawing speed. $2,195 Monochrome, $2,995 Color.

Stop watch courtesy of Heuer Electronic Corp.

"One, one thousand... done!"

Our new graphics co-processor technology gives you a faster draw than anything under $10,000. It's ten times as fast as others in its price range. A hundred times as fast as the slowest terminals.

You get superior resolution, too: 800 x 600 monochrome, 800 x 300 color. The monochrome terminal displays four gray levels. The color terminal displays 16 colors (out of 64 choices). Plus fifteen programmable, non-volatile function keys. Simple menus. Superior ergonomics.

You don't sacrifice a thing for superior graphics. The alphanumeric quality equals the best text terminals. The display is not interlaced (not running at half speed), so there's no smearing or ghosting when you scroll. You can even choose an 80- and 132-column display to get a full spreadsheet on the screen.

HiSCAN™ graphics terminals have full DEC VT220 and TEK® 4010/4014 compatibility. Plus your choice of DEC ReGIS, TEK 4027 or TEK 4105 compatibility at no extra cost. They're cable-ready for light pens, mice, digitizing tablets and inkjet printers.

They're designed to help you be more productive and creative. Because they work almost as fast as you think. So now you can play "what-if" with graphics.

You have to get your hands on a HiSCAN graphics terminal to believe it. And that's easy. We'll even arrange a 30-day trial at no risk. Just send the coupon.

New HiSCAN Terminals
From the people who brought you Retro-Graphics®

CD-9/84

☐ Send complete information

☐ I want a 30-day trial with no risk. Call me about my compatibility requirements.

NAME _____ TITLE _____

COMPANY _____ PHONE _____

ADDRESS _____

CITY _____ STATE _____ ZIP _____

Mail to: Digital Engineering, 630 Bercut Drive, Sacramento, CA 95814

™ **DIGITAL ENGINEERING**
630 Bercut Drive, Sacramento, CA 95814
Phone (916) 447-7600 Telex 910-367-2009

CIRCLE 67

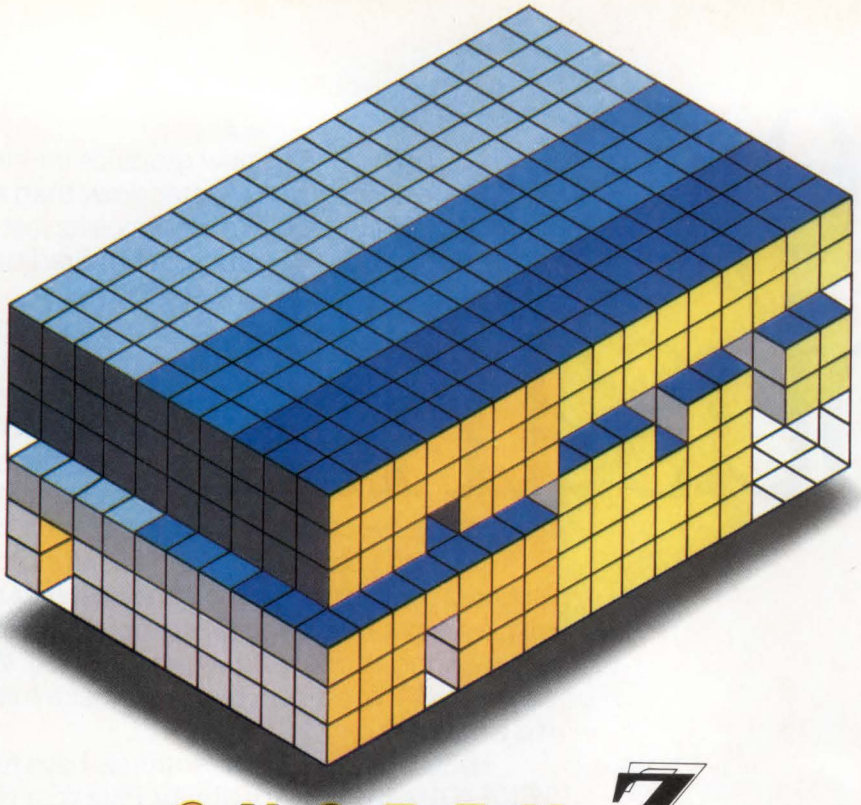

SYSTEM Z

HOW YOU CAN OWN THE BEST MICROPROCESSOR DEVELOPMENT SYSTEM IN THE WORLD

First look to ZAX. Because ZAX knows what it takes to be the "Best." And that means offering development products that are intelligent, compatible, fast, powerful, and, what most manufacturers forget; *obtainable* – because if you can't afford to buy it, you can't use it.

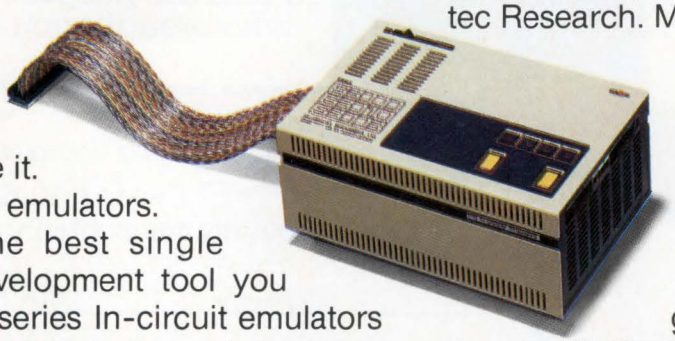

Start with ZAX emulators. They're simply the best single microprocessor development tool you can buy. ZAX ICD series In-circuit emulators offer several advanced features such as extensive emulation memory, breakpoint trace capability, true realtime emulation with no wait states, and a variety of powerful debugger commands. And as you can see, ZAX supports a broad spectrum of 8 and 16-bit microprocessors as well.

But don't stop there. To get things running, add the capable and economical IBM PC or equivalent personal computer. Now make it perform with professional software from Microtec Research. Microtec's cross-software products include symbolic debuggers, fully manufacturer compatible macro assemblers, overlay linking loaders and hi-level language support.

You could end there, but ZAX doesn't. As your needs grow, continued support is assured with System Z; an incredible complete development program that can link your PC with your mainframe computer.

If owning the best microprocessor development system in the world interests you, contact ZAX today at **800/421-0982** for more details.

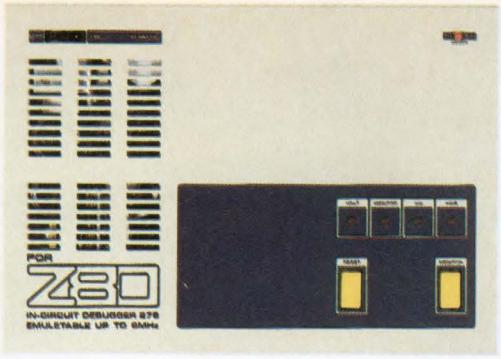

Z80
Z80
Z80B
Z80H

Emulates Z80B microprocessors to 6 MHz and Z80H to 8 MHz. Features; 64K byte user emulation memory, 2K deep x 32 bits wide realtime trace buffer, 29 different debugger commands.

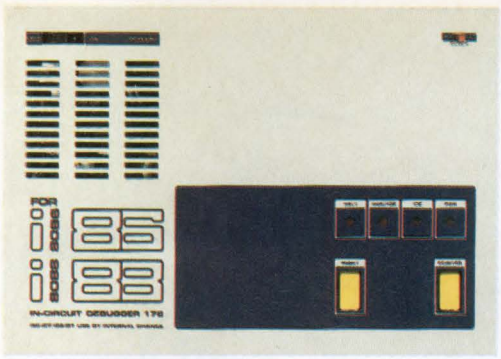

i8086/88
i8086
i8087
i8088

Co-emulation of 8086 and 8087 or 8088 and 8087 processors to 5 MHz. Realtime emulation to 8 MHz for 8086/88 processors. Features; 128K bytes static RAM – expandable to 1 Mbyte, 4K deep x 40 bits deep realtime trace buffer, 30 different debugger commands.

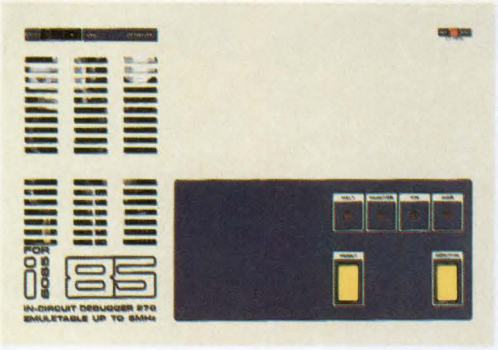

i8085
i8085A
i8085A-1
i8085A-2

Emulates 8085 processors up to 6 MHz. Features; 64K byte user emulation memory, 2K deep x 32 bits wide realtime trace buffer, 29 different debugger commands.

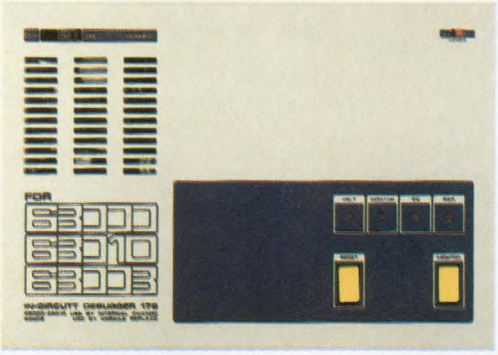

68000
68000
68010
68008

Emulates 68000, 68008 and 68010 in one unit to 10 MHz. Features; 128K of emulation memory – expandable to 256K, 4K deep x 48 bits wide realtime trace buffer, 30 different debugger commands.

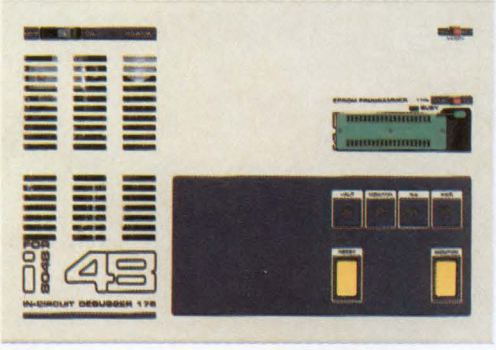

i8048
i8048
i8049
i8050
i8748
i8749
i8039
i8035
i8040

Emulates entire 8048 family in one unit to 11 MHz. Features; 4K emulation memory, 2K deep x 32 bits wide realtime trace buffer, 29 different debugger commands. 8748 and 8749 units feature a built-in EPROM programmer.

All emulators come equipped with these standard features; 3 hardware and 8 software breakpoints, an Event Trigger and External Probe triggered breakpoint, 14 selectable baud rates to 19200bps, 20" processor probe, 115/230 VAC operation. All Zax emulators are compact, lightweight, rugged and can easily fit into a briefcase.

SYSTEM Z-from Z...ing

ZAX

Zax Corporation 2572 White Road, Irvine, California 92714 (714) 474-1170 • 800-421-0982 • TLX 183829

CIRCLE 68

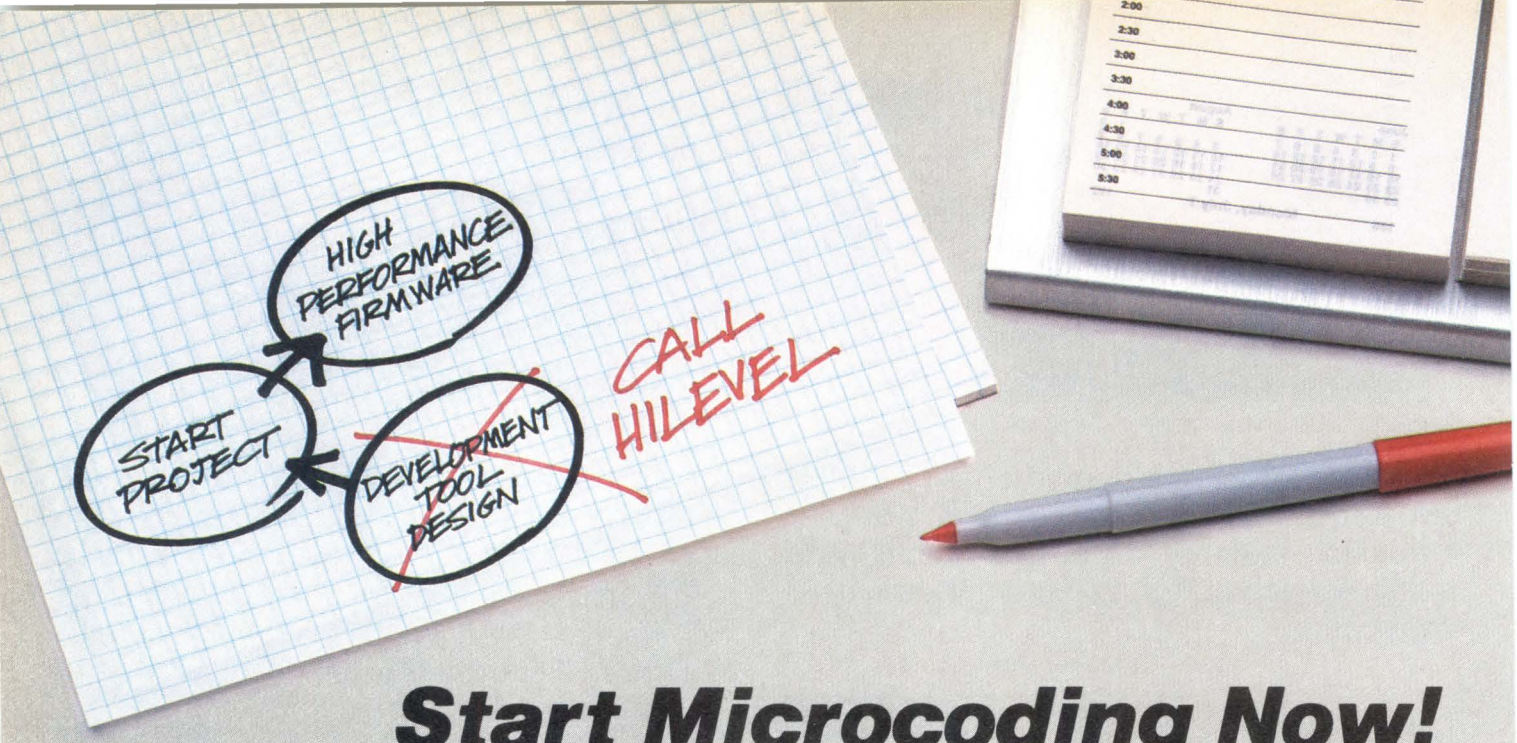

Start Microcoding Now!

Why waste valuable time designing a development tool when it's already available, off-the-shelf, from HILEVEL? It's the DS370 EMULYZER™ with built-in bit-slice power that lets you increase your productivity and shorten the time it takes to get your product to market.

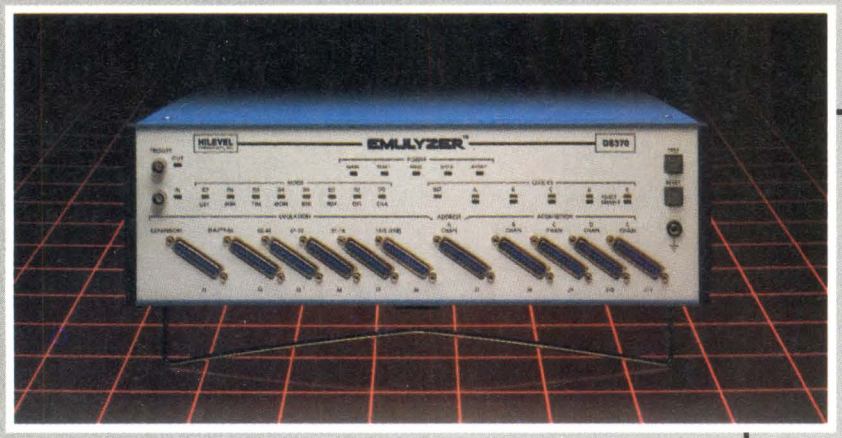

Compare these features with what you'd build in yourself: integral logic analyzer/high-speed PROM emulator/powerful pattern generator; high quality transmission line cabling; configurable and reconfigurable; performance analysis; 16 level triggering; gold-on-gold reliability; plus what you could never build in yourself — unbeatable field support from HILEVEL.

You've got a choice. Either spend your time reinventing the wheel. Or get your project rolling now — with HILEVEL.

Show Me!

☐ Arrange a demo
☐ First send more data

Name _____
Company _____ M.S. ____
Address _____
City _____
State _____ Zip _____
Telephone _____

Action Now? **Call 1-800-HILEVEL!**

HILEVEL TECHNOLOGY, INC.
18902 Bardeen Way, Irvine, CA 92715
(714) 752-5215 TLX: 65-5316

HILEVEL TECHNOLOGY, INC.

CIRCLE 69

ACTORS SET THE STAGE FOR SOFTWARE ADVANCES

A single software mechanism that represents data and procedures for both conventional and artificial intelligence programming will advance the art of programming.

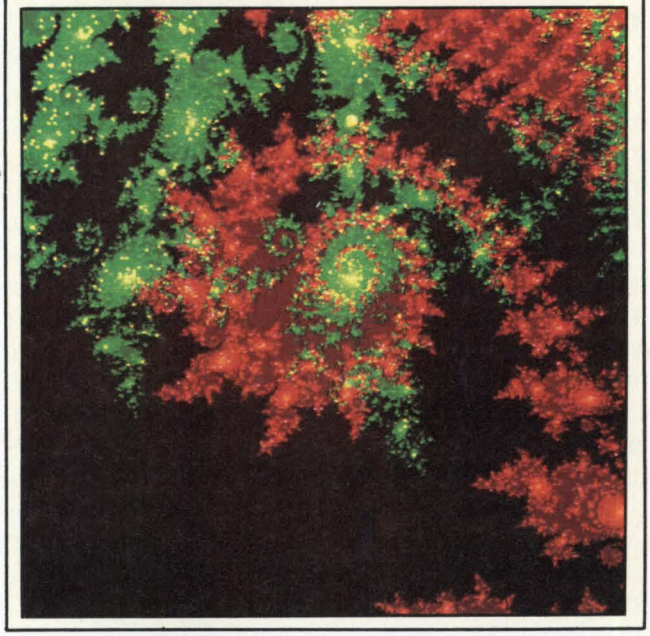

Photo courtesy of Dr. John Hubbard and Homer Smith, Cornell University

by John R. Pugh

In many respects, the software systems being developed by artificial intelligence researchers are no different from the systems being developed by the business and scientific communities. These systems are large, intricate, and often very difficult to implement if they are to be understandable, reliable, and maintainable. The artificial intelligence community has been developing its own ideas in order to deal with the construction of such systems. However, the application of these ideas is not restricted to this domain.

One of the most important ideas developed is the actor concept. An actor is a single software entity simultaneously representing data and procedures for either conventional or AI programs. As actor- and object-oriented languages (similar to actor orientation) become more generally available over the last part of this decade, the actor paradigm will have a profound influence on programming practice. Moreover, the actor model may prove to be appropriate

in exploiting the parallelism presented by advanced computer architectures.

Actors can be thought of as independent processing units that operate concurrently with their own instruction sets and memory. They also have the ability to communicate with other actors by message passing. As memory and processor costs fall, it is possible to perceive machines whose computational power will be derived from interconnected processors where each processor has its own memory. Using this architecture, actors can be assigned to different processors and work in parallel to solve a problem. At least one such system is currently under construction at the Massachusetts Institute of Technology (MIT) in Cambridge, Mass. As yet, however, the problems associated with such distributed computation are not fully understood.

Data abstraction

Most of today's commonly used high level programming languages are based on the classical procedure/data computation model. They draw a distinction between the data or information being manipulated, and the procedures that actually perform the manipulation. Some modern languages (eg, Ada) have changed this emphasis somewhat by providing increased language support for data abstraction. This move represents recognition of the importance of locality and modularity concepts in large software system design.

An abstract data type can be thought of as a program module that extends the predefined data types provided in a language. It therefore consists of a data structure together with the operations that can be applied to members of the new type. These are bound together in a single program module. Users of the type need only be provided with a specification of the values that members of the new type may take, and of the effect of each operation.

The internal details of the data type module, data structure, the details of the implementation of each operation, and the local data and procedures are hidden from the user. This design supports the generally accepted idea of information hiding. Thus, users are presented only with an external functional module specification while the internal specification remains hidden and inaccessible. The internal details of a module can be modified without impact on any module users.

Actor- or object-oriented systems take the view that a single entity, an actor or object, should represent both data and procedures. They use and extend the explicit facilities for creating classes and instances that were first introduced by the Simula programming language. However, Simula's object-oriented facilities can be thought of as extensions to a predominantly traditional Algol-like, procedure-data language. In contrast, actor systems adopt the far more radical view that programs should consist solely of communicating actors or objects. The first actor- or object-oriented systems evolved from two research efforts. One was the development of the Smalltalk family of languages by the Learning Research Group at the Xerox Palo Alto Research Center (PARC); the other was the development of the actor model of computation by the AI research group at MIT.

Actor- or object-oriented systems view a single entity, an actor or object, as representing both data and procedures.

Smalltalk is much more than a programming language. It provides a complete programming environment for the design and implementation of software systems. The system is designed for use with a personal computer that incorporates a high resolution, bit-mapped graphics display together with a mouse and keyboard. The latest iteration of the language, Smalltalk-80, has recently been made available on the Xerox Dolphin, Dorado computers, and Sun workstations.

Languages that have emerged from the AI world include Act 1, an experimental language for constructing actor systems and studying the actor computation model; Director, an actor-based animation language; Flavors, a sublanguage of Zetalisp (MIT) Lisp; Ross, a simulation language; and Loops, which adds facilities for data-, object- and rule-oriented programming to Interlisp. Actor-based software systems organize programs as "societies" of communicating problem-solving experts. For AI researchers, the actor methodology shows promise of managing the complexity of large AI systems. Most AI actor systems are implemented as extensions to Lisp dialects.

Smalltalk and the various AI systems share many common ideas and principles. However, there are differences in the way these ideas are realized. And, some systems, particularly the AI variety, introduce concepts unique to themselves. Rather than concentrate on these differences, it is best to summarize the significant characteristics of actor systems and their advantages.

Spotting an actor

First, a note on terminology. The term "actor," originally coined by Hewitt (MIT), is generally used in the AI community, whereas "object" is used in Smalltalk. The term actor typifies the energy and dynamism associated with the concept, whereas object suggests passivism and an inanimate nature.

The actor model of computation is attractively intuitive and the interactions between independent

IBM TO ANYTHING

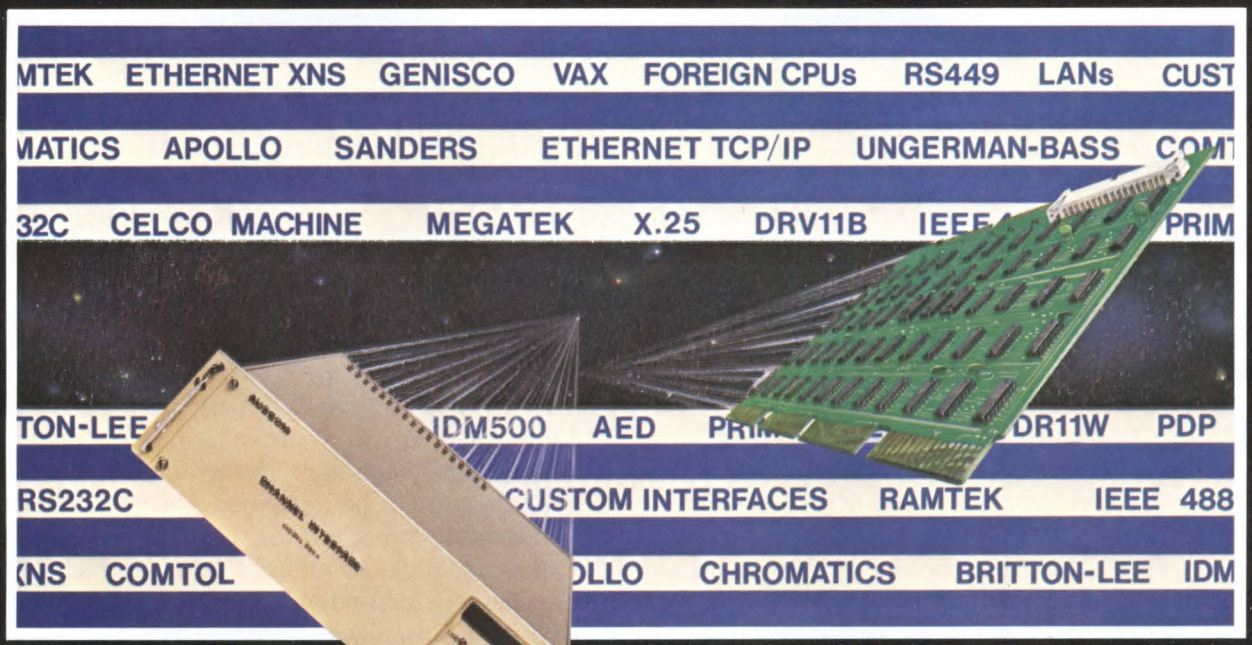

AUSCOM- Channel Interface Products

AUSCOM, INC., offers the broadest range of channel attach solutions in the industry. Our family of programmable channel interfaces allows you to get from your IBM mainframe to your downline devices — SIMPLY... RELIABLY... EFFECTIVELY. By emulating a standard sysgenable control unit, no changes need be made to the host operating system.

Both the 8900 series and the 8300 series are leaders in the following applications:

LANs
Choose a member of the 8900 series for versatility in network routing and in mapping multiple subchannel addresses. Standard protocols, including XNS, TCP/IP and OSI/ISO are available along with a variety of custom high speed protocols.

GRAPHICS
Connect high resolution graphics systems directly to the mainframe or through a channel attached local area network.

PRINTERS / PLOTTERS
Select a member of the 8300 series, emulating a standard buffered control unit, to connect high speed printers, plotters or other I/O devices to the mainframe CPU.

TELECOMMUNICATIONS
Link directly to a TCOM network whether it be based on microwave, satellite, fiber-optics or conventional wire to a mainframe with either series of products.

In all applications, AUSCOM's proven, reliable products come with complete and extensive diagnostic software. Our support before and after the sale assures successful operation.

AUSCOM maintains a large and growing library of software routines to link downline devices to the mainframe. If our library does not meet your needs, our engineering department will develop custom interfaces or will fully support your engineering department in its internal developments.

AUSCOM DELIVERS SOLUTIONS!

AUSCOM, INC.
"The IBM to Anything Company"

2007 Kramer Lane
Austin, Texas 78758
512/836-8080

DEC is a trademark of Digital Equipment Corporation.
IBM is a trademark of International Business Machines.

CIRCLE 70

entities (actors) in a program can very often be equated to interaction between humans. Each actor in a system plays out an active (acting) role not unlike the role humans play in real-life systems. Therefore, the actor style of programming (sometimes called anthropomorphic programming), comes very naturally to programmers.

End users of actor-based systems can think of actors in a computerized system as corresponding physical actors in a real-life system. For example, in an electronic office environment, the office manager might read the mail and ask the mail administrator actor, "Do you have any new mail for me?" Or, the manager might check appointments by asking the diary administrator actor "What appointments do I have today?" The office manager can be thought of as just another actor.

An actor is a small processor defined solely by its behavior. And, its behavior is characterized by its response to receiving a message. As mentioned earlier, communication between actors is achieved through the single, uniform metaphor of message passing. When an actor receives a message, the actor determines whether or not it recognizes the message as one for which it has a programmed response. If so, the actor evaluates the script, method, or procedure associated with the message. Then, if necessary, the actor relays a response back to the original request sender (or another actor).

An actor has a local body of knowledge that it can use to respond to a message. If appropriate, the actor can provide operations allowing other actors to interrogate and possibly update this knowledge. The actor "owns" its own knowledge; therefore, the actor decides if and when other actors can access the knowledge, or even know of the knowledge's existence. A distinctive feature of actor systems is that there is no central database facility. Rather, knowledge is embedded in individual actors and thus distributed throughout a system.

An actor is a small processor defined solely by its behavior in response to receiving a message.

Examined in isolation, an actor consists of a local knowledge base and a set of behaviors. As stated, a behavior is characterized by a message and a script or method defining what the actor's response to that message will be. To respond to a message, an actor will often need assistance from other actors with which it is acquainted. Thus, the scripts for an actor involve sending messages to other actors. In this way, messages propagate throughout an actor system.

Note that, since the actor receiving a message determines the response, the same message sent to different actors can result in many different responses. The type of message that can be sent to actors varies greatly from system to system. The Smalltalk family uses message-sending expressions that must conform to fairly rigid syntactic rules. In many ways, these expressions resemble traditional procedure calls. The message specifies a receiver, a message selector for the receiver to identify, and any arguments necessary for message processing. In contrast, AI systems often allow quite complex message structures. These systems make use of sophisticated pattern matchers to parse incoming messages. They have also led to the development of highly readable, English-like user interfaces to actor systems.

Sharing the stage

Just as humans or objects often share similar characteristics and abilities, all actor systems provide mechanisms that allow actors to share common knowledge and behaviors. For example, this design allows an actor to be described as a specialization or extension of another actor. Perhaps such an actor is almost identical to an existing actor except that the specialized actor has additional knowledge or behaviors. Knowledge sharing avoids the duplication of scripts for shared behaviors in every actor. It also makes the modification of such scripts more manageable by centralizing them in a single actor.

There is no uniform mechanism for the implementation of knowledge sharing within actor systems. A common approach is to use the notion of an actor inheriting the knowledge and behavior of another. The simplest form of inheritance arranges actors in a hierarchical, or tree, structure. Actors automatically inherit the knowledge and behavior of their parent in the tree. In such an environment, an actor receiving a message (which it cannot respond to) will pass the message onto its parent. The parent may be able to respond to the original message or it may pass the message to its parent in turn. The actor at the tree root is the most general and contains behaviors applicable to all actors.

Individual actors can override the automatic inheritance mechanism by defining their own individual behaviors or knowledge, which will supersede those contained within the parent. In Smalltalk, inheritance is implemented using a subclassing mechanism. In other words, instances of a class inherit the attributes of that class. If that class is a subclass of another, then (if needed), the instance may inherit attributes of the superclass.

An actor language often bases its inheritance mechanism on ancestral information. An actor inherits the capabilities of the actor that created it. But, a simple inheritance mechanism based on ancestral information is too restrictive for describing many problem situations efficiently. Actors are often not

simply specializations or extensions of a single actor. In reality, they embody facets or characteristics from a number of actors. Thus, mechanisms that support multiple inheritance (the ability to inherit knowledge from more than one actor), are incorporated into many actor systems.

Three important applications of the actor concept are simulation, animation, and office automation.

Returning again to the analogy with real-life systems provides one final point. Humans in real-life systems operate concurrently. Primitives for dealing with concurrency problems have traditionally been grafted onto procedure/data languages (eg, Ada, Modula, and Concurrent Pascal). One main motivation for much of the work on actor systems has been the belief that they are better able to exploit and control the parallelism that must be harnessed by large, future AI systems. For example, the problems of access to shared resources are more controlled in an actor world (compared to more conventional programming), where knowledge is centralized in resource management actors, and where concurrently executing processing units (actors) are largely independent and communicate solely via message passing.

Working actors

The actor concept is appropriate in a broad range of applications. Three important applications are simulation, animation, and office automation. In fact, since actor systems view programming as a simulation of the real world, simulation is a rather obvious application. A research program in knowledge-based simulation at the Rand Corp (Santa Monica, Calif) has led to the development of the actor-oriented language, Ross. In turn, Ross is used as the implementation language for Swirl, an air battle simulator used by military strategists. Objects in the air battle domain (eg, AWACS systems, fighter planes, enemy "penetrators," command centers, and fighter bases) are represented as Ross actors in the simulation system.

Actors are either generic actors or instance actors. The generic fighter class actor, for example, contains knowledge and scripts common to all fighters, while individual fighter actors contain instance-specific information. Generic knowledge and scripts are accessed through a multiple inheritance mechanism. Messages to actors are pattern matched, allowing a readable and flexible English-like command structure for military personnel concerned only with the air battle domain.

The actor computation model has been used extensively in computer animation systems such as Actor/Scriptor/Animation System (Asas), Director, and Cinemira. These systems use actors to represent theatrical performers and animated objects. They allow the movement of these actors to move concurrently from one movie frame to the next. Asas, initially developed at MIT, is responsible for the special effects in the movie *Tron*. Asas is a general purpose programming language, built on top of Lisp. It is extended with primitives for geometric objects and operators, and parallel control structures. When an animated sequence is produced frame by frame, Asas allows actors to enter and exit "on cue," act in parallel, and synchronize their interaction with each other through message passing.

In the office automation domain, the graphical user interface to the Xerox Star professional workstation is achieved through object-oriented programming, subclassing, and multiple inheritance. The workstation software comprises over 250,000 lines of code. Before writing a single line, designers decided that users would most easily relate to a final system patterned after the physical office concept. Through a sophisticated graphical user interface, this concept presents users with a virtual desktop on which icons representing physical office objects can be manipulated. The selection of an object-oriented implementation is a natural one. The system design model of communicating physical office objects is also the model used to implement the system.

John R. Pugh is an associate professor at the School of Computer Science, Herzberg Bldg, Carleton University, Ottawa, Ontario, Canada K1S 5B6. He holds a BS and an MS in computer science, both from the University of Wales, in Swansea, Wales.

Only One Company Makes Interfacing

Q-BUS

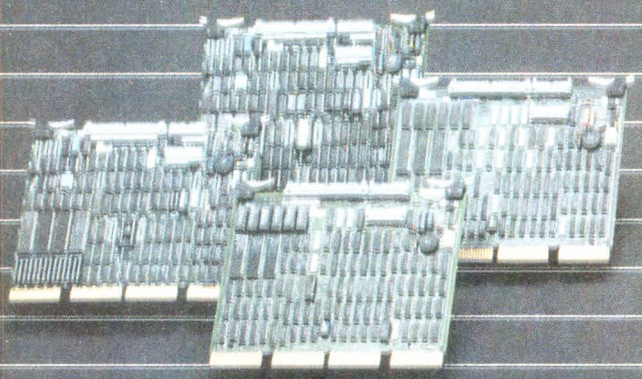

8" & 14" WINCHESTER—SMD I/O
- RL01/RL02, RP02/RP03, RK06/RK07 & RM02/RM05/RM80 emulations

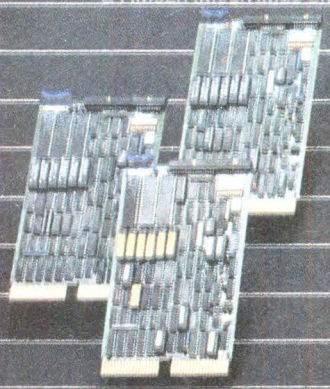

WINCHESTER 5¼"—ST506/412 & DMA I/O
- RL01/RL02, RK06/RK07 emulations

Q-BUS

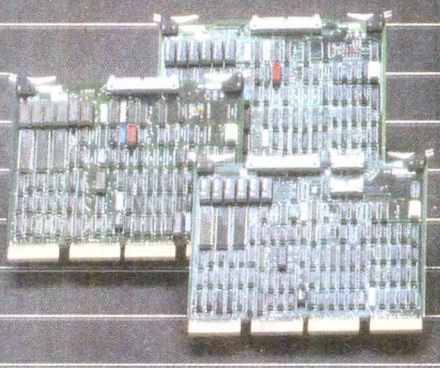

8" & 14" WINCHESTER—SA4000 & PRIAM I/O
- RL01/RL02, RP02/RP03 emulations

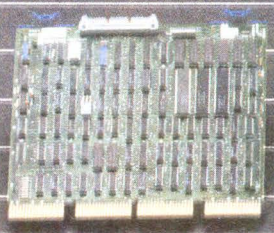

CARTRIDGE—14" DISK—DIABLO 44B
- RK05 emulation

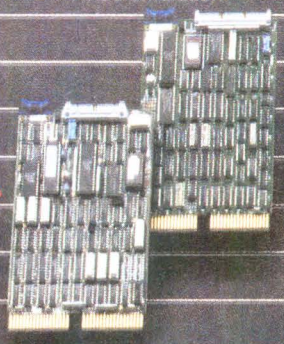

FLOPPY DISKS—SA850 & SA450 I/O
- RX02 emulation

Compatible with DEC's family of RT-11, RSX-11, RSTS, UNIX and VMS operating systems.

Your DEC Computer So Easy

Q-BUS

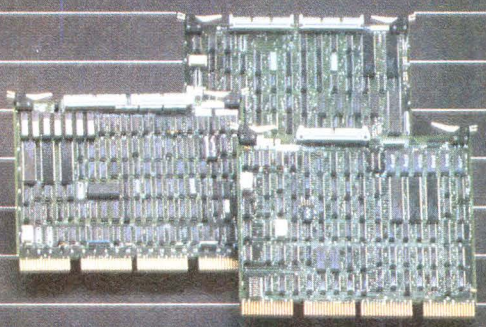

½" TAPE—CONTROLLER/COUPLER PERTEC I/O
- TM-II, TSV05/TS-II/TU80 emulations

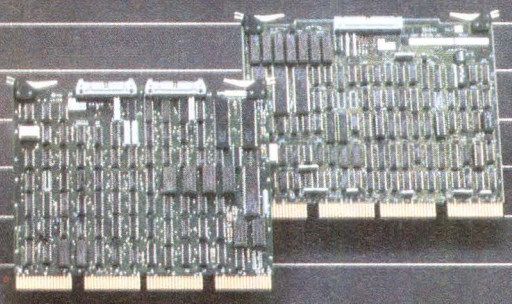

¼" TAPE—KENNEDY 6455 & CDC SENTINEL I/O
- TM-II/TS03, TS-II/TU-80/TSV05 emulations

UNIBUS

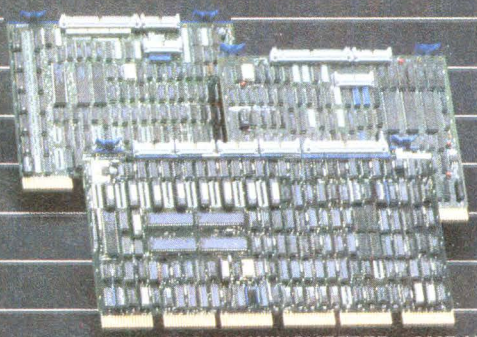

8" & 14" WINCHESTER—SMD I/O
- RP02/RP03, RK06/RK07, RM02/RM05 emulations

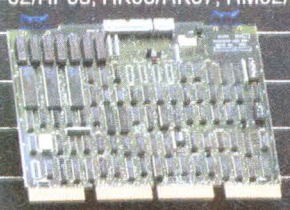

¼" TAPE—CDC SENTINEL I/O
- TS-II emulation

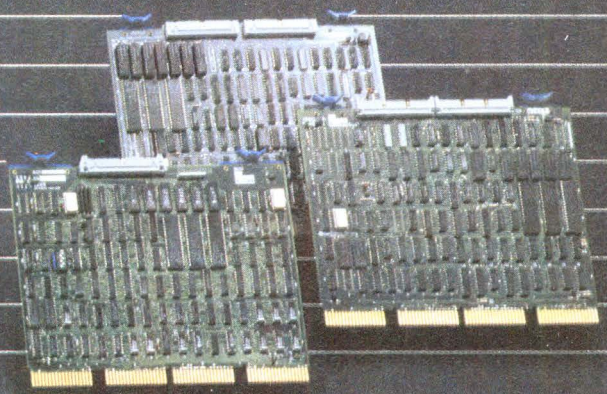

½" TAPE—CONTROLLER/COUPLER PERTEC I/O
- TM-II, TS-II/TU80 emulations

DILOG

DISTRIBUTED LOGIC CORPORATION

1555 South Sinclair • P.O. Box 6270 • Anaheim, CA 92806 • (714) 937-5700 • TLX 6836051
64-A White Street • Red Bank, NJ 07701 • (201) 530-0044
Chester House, Chertsey Road • Woking, Surrey, England GU 21 5BJ • (04862) 70262/7
• TLX 859231 DILOGI G • FAX (04862) 62666

CIRCLE 71

What do you like better about the VISUAL 2000... its power or its versatility?

Personally, I like its price!

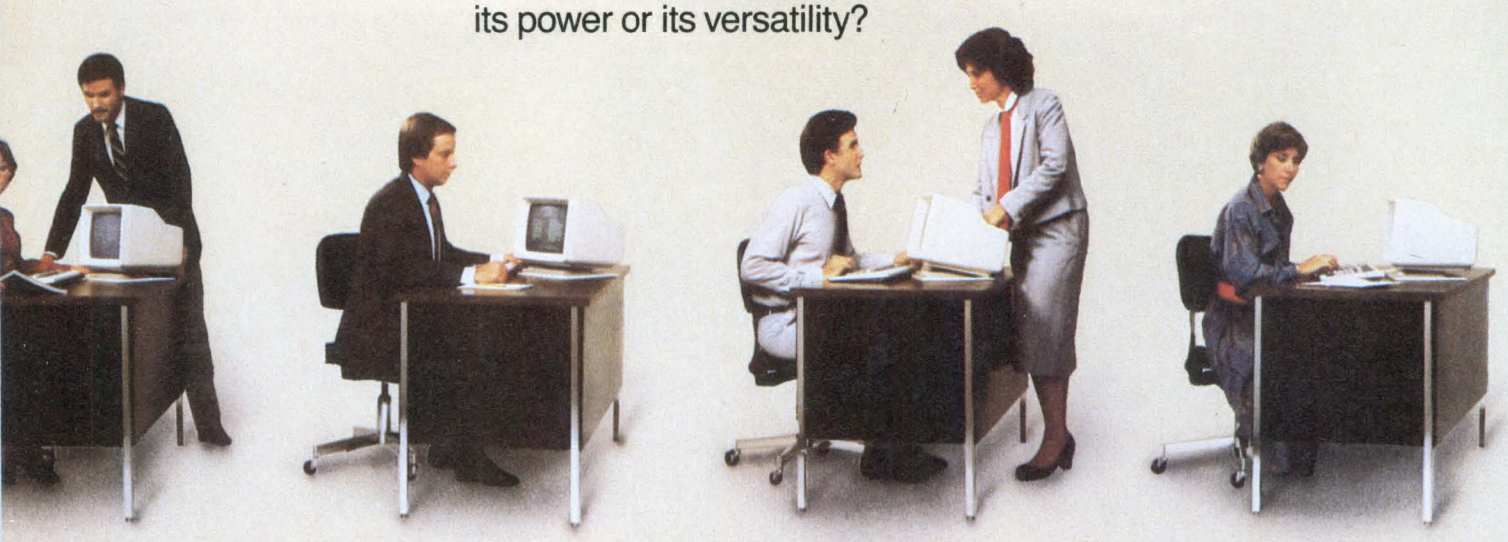

Never has a UNIX-based multi-user system given so much to so many for so little.

Introducing the VISUAL 2000

The VISUAL 2000 is the full-featured system with the power and flexibility to support multiple users in real business applications at a surprisingly low cost per station. It can be used with inexpensive video terminals. Or as a database manager or file server for a cluster of intelligent workstations or PCs, including both the IBM® PC and VISUAL's own lightweight, portable, totally IBM PC compatible COMMUTER. In all applications it offers greater performance, more flexibility, and lower cost than any other system in its class.

Powerful Intel 286 processor

The Intel 286 is today's chip of choice for UNIX™-based systems. Only the Intel 286 has on-chip memory management, an instruction set optimized for multi-tasking, and the optional 287 numeric co-processor to speed up floating point by a factor of 10.

What do these features mean to the VISUAL 2000 end user? Faster response time, more users supported, and lower system cost!

Cost-effective one-board design

A basic advantage of the VISUAL 2000 is its one-board base-level design. A single high-density board includes the 286 CPU, 512KB-2MB of RAM, controllers for Winchester, floppy, and streaming tape, an intelligent communications processor, six RS-232 ports, and a parallel printer port. Even a real-time clock with battery backup. One-board design means higher performance, lower cost, and greater reliability than comparable multi-board implementations.

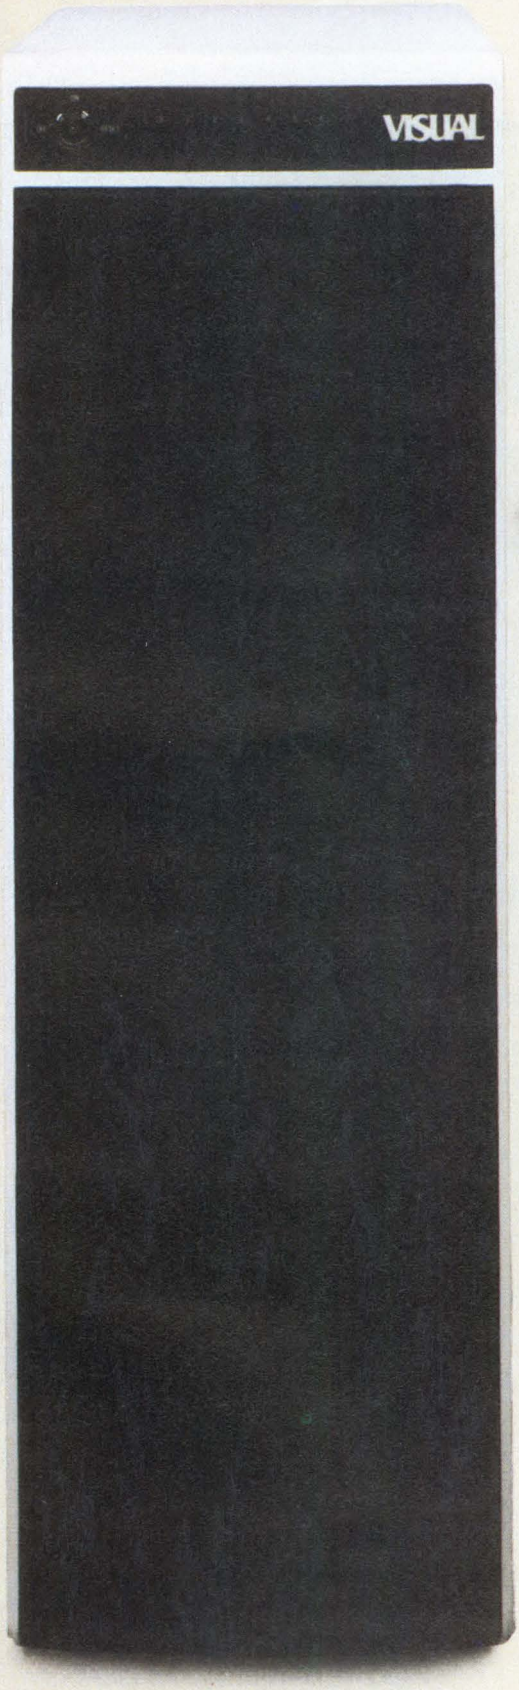

Configurability and Expandability: VISUAL gives you more

The VISUAL 2000 spans a much wider range of configurability and expandability than other systems in its price class. Up to 16 independent users. 6 megabytes of RAM. 4 Winchesters. Floppy. And streaming tape for simple, reliable disk backup. All in a small stand-up enclosure which looks right at home next to a desk.

And if a fully expanded VISUAL 2000 isn't enough, you can connect up to 254 VISUAL 2000s, PCs, and workstations in a local area network.

Extensive system software simplifies system integration

The VISUAL 2000 runs XENIX, Microsoft's popular, enhanced version of UNIX, derived from UNIX under license from AT&T, and designed to be faster, more secure, and easier to use in business applications.

And VISUAL has worked hard to simplify the system integrator's job, by providing all the tools needed to deliver end-user applications with a minimum of effort.

Languages such as C, SMC BASIC, RM/COBOL, TOM BASIC, SOFTBOL, and MicroFocus Level II COBOL, to provide instant compatibility with hundreds of proven business application programs.

Other system-building tools, like the INFORMIX database management system and RealWorld modular accounting system.

And productivity software, such as the 20/20 integrated spreadsheet and XED office-grade word processor.

The Bottom Line

High performance. Superior flexibility. Extensive software. And low cost...VISUAL 2000 systems start at under $10,000, suggested list. No one gives you more in a UNIX-based multi-user system.

Whether you're an OEM, system house, distributor, or end-user, call today for further information on the VISUAL 2000 and see for yourself!

VISUAL See for yourself®

Visual Technology Incorporated
540 Main Street, Tewksbury, MA 01876
Telephone (617) 851-5000. Telex 951-539

REGIONAL OFFICES:

Northwest:	(415) 490-1482
Southwest:	(213) 534-0200
North Central:	(513) 435-7044
South Central:	(214) 255-8538
Northeast:	(201) 528-8633
Southeast:	(301) 924-5330

CIRCLE 72

DOUBL

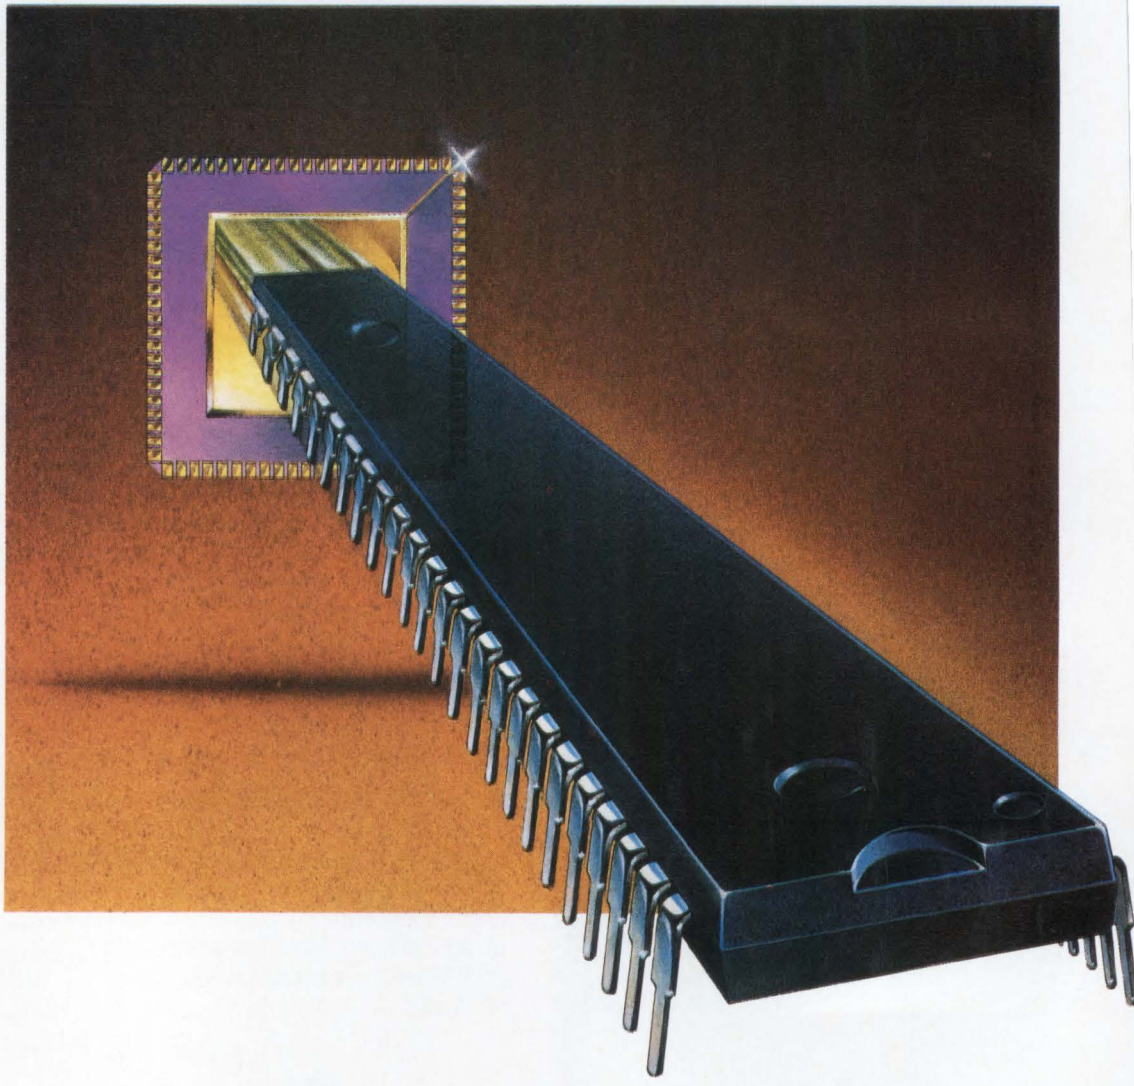

E DUTY

Introducing the MK68200 16-bit microcomputer/microprocessor.

Unlike any other 8- or 16-bit single-chip microcomputer, the new Mostek MK68200 performs equally well as a single-chip microcontroller and as a peripheral I/O controller. In the latter type of application, the MK68200 can provide local intelligence required to handle typical computer system I/O functions. Without tying up the host processor. In addition, it can transfer data to and from system memory using a software DMA function. So, as an example, it's ideal as a front-end processor for handling complex serial I/O protocols.

Actually, the 16-bit MK68200 gives you a whole new dimension of system control. One that enables you to monitor, measure, regulate, sense, and interconnect with more precision than ever before.

The operative word is more. More speed. More powerful instructions. More I/O functions. More efficiency. More versatility. All of which mean more application opportunities. For robotics. Engine control. Pattern recognition. And real-time measurements and control – as a stand-alone device, or in tandem with other microprocessors as an intelligent peripheral controller.

First, consider speed. Using the 6 MHz instruction clock rate, the MK68200 can execute a 16-bit multiply as fast as a 3.5μs; a 32 x 16 divide in 3.8μs; and a 16-bit add/subtract in as little as 500ns. In other words, it's faster than any other general-purpose, single-chip micro currently available.

One reason for this faster speed is code efficiency; most MK68200 instructions are just one word. That not only saves code space, it also improves execution speed. And the powerful instruction set incorporates more than 50 instruction types which operate on both byte and word operands.

Next, look at the available versions. The MK68200 comes in a 48-pin DIP with 0 or 4K bytes of ROM and 256 bytes of RAM. It is expandable to address a full 64K byte address space externally. And finally, an 84-pin LCC version (with no ROM), is available for system development or for multiple bus applications.

For single-chip applications, up to 40 pins of the 48 pins are available for I/O to include two 16-bit parallel ports. Plus a full-duplex USART with double-buffered transmit and receive capable of operating at data rates up to 1.5Mbps (using a 6MHz instruction clock).

The MK68200 offers extensive interrupt capabilities. One non-maskable interrupt input is provided as well as 14 maskable interrupt sources. In addition, all 14 independently vectored interrupts are user-assignable in software for any priority scheme. And to enhance serial communications, there's also an innovative wake-up interrupt on the serial channel.

As for support, there's a powerful Macro Cross Assembler, along with a host of other software tools. As well as RADIUS™, a very cost-efficient remote development station. And to expand the MK68200 family, plans include higher performance versions with denser memory and faster throughput. Plus power-conserving CMOS.

Look beyond the limits of conventional single-chip systems into the realm of ultimate control. It's yours with the MK68200. For more information, contact Mostek, 1215 W. Crosby Road, MS2205, Carrollton, TX 75006, (214) 466-6000. In Europe, (32) 02/762.18.80. In Japan, 03/496-4221. In the Far East (Hong Kong), 5-681157-9.

RADIUS is a trademark of Mostek Corporation.

CIRCLE 73

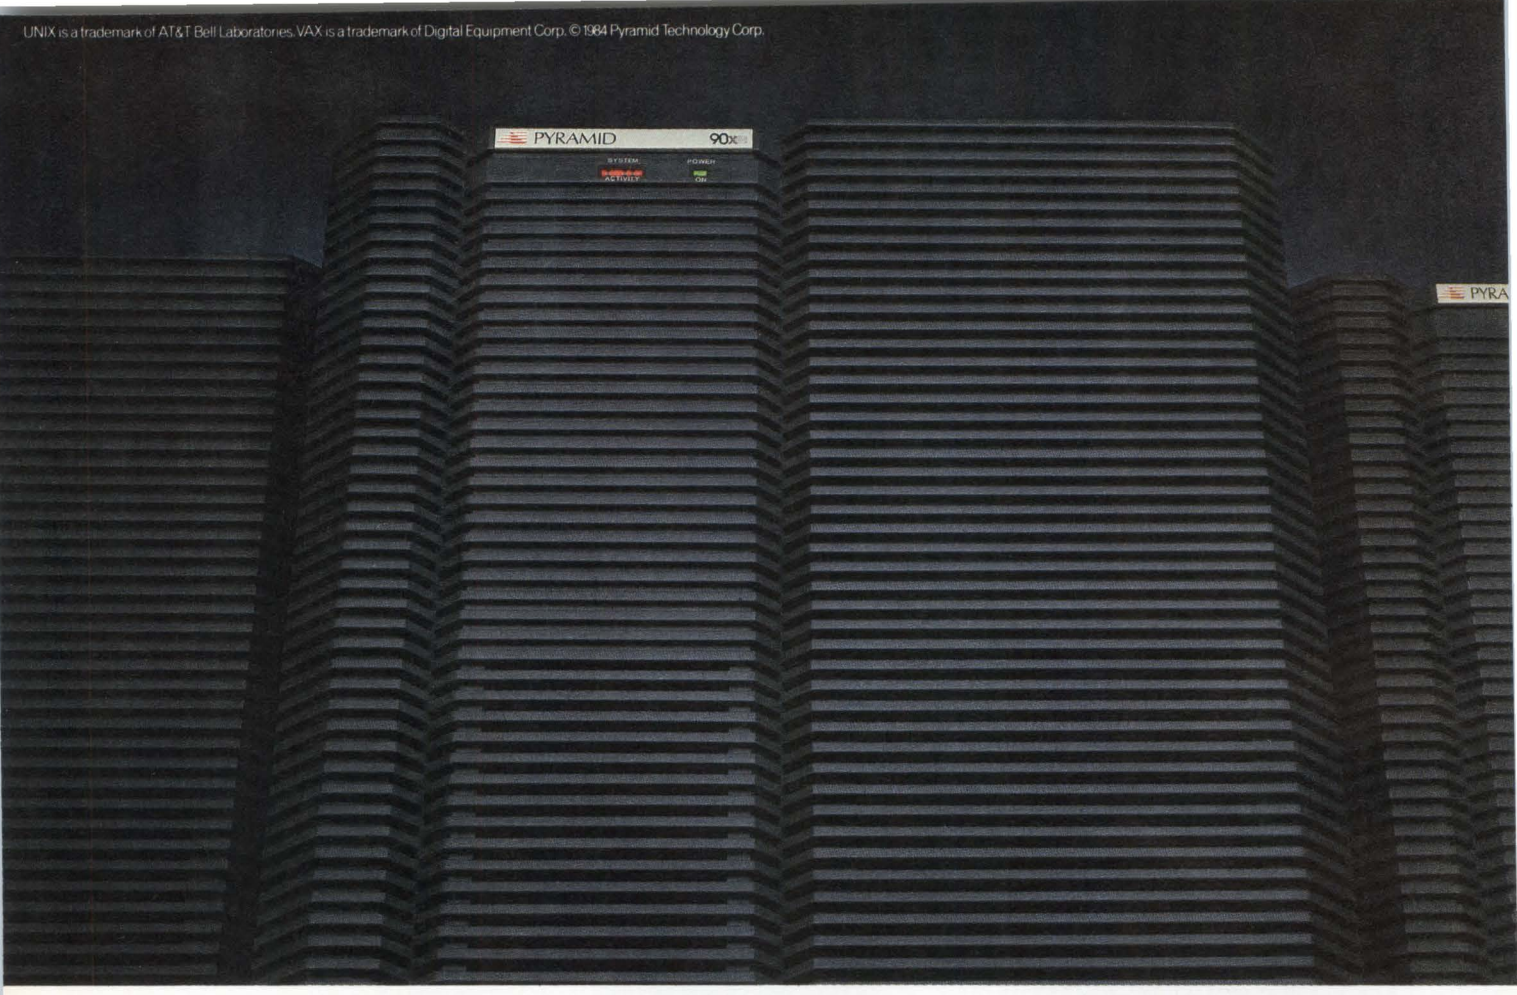

IF IT WEREN'T FOR UNIX WE NEVER COULD HAVE BUILT THE PYRAMIDS

Ordinary computers, yes. But not a Pyramid Technology 90x.

After all, here's a supermini not just capable of running UNIX,™ but born to run it. And run it up to four times faster than the most popular UNIX host. For a lot less money.

The secrets of this Pyramid are a thorough understanding of UNIX, a few fundamentals of RISC (Reduced Instruction Set Computer) theory, more registers than 30 VAXs, and a 32-bit proprietary architecture that outperforms a roomful of micros.

All combined to speed up UNIX just where it likes to slow down.

For example, gone are 85% of performance-robbing memory references. The endless parameter shuffling of yesterday's technology has been replaced with a hardware register window. Even context switching takes less than one percent of the CPU's time.

It's amazing what hardware architects can do, given the chance. It's almost as startling as what our software wizards did.

They crafted OSx, a dual port of Berkeley's 4.2 BSD and Bell's System V. Because you can switch environments at will, no UNIX port offers more capabilities. With absolutely no loss of compatibility.

Well, almost.

We do admit to one feature not compatible with other UNIX systems. Our single-source support.

One telephone number instantly connects you to both hardware and software experts. In-house pros, who spend their energy pointing you towards solutions. Not pointing fingers at each other.

So no matter how you see your requirements shaping up, contact Pyramid Technology, 1295 Charleston Road, Mountain View, California 94043. Or call (415) 965-7200.

Because when it comes to running UNIX, a Pyramid looks good from any angle.

PYRAMID TECHNOLOGY

PRACTICAL TOOLS EARN AI NEW LEVEL OF RESPECTABILITY

With some practical artificial intelligence tools already here, system designers can start implementing tomorrow's applications today.

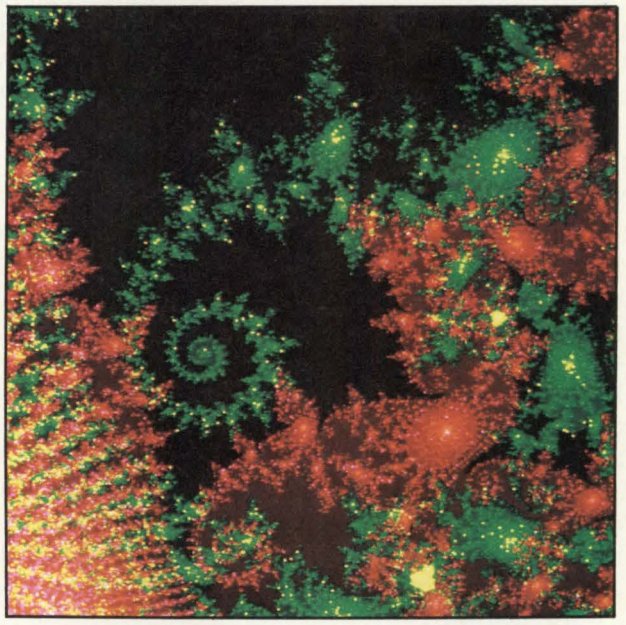

Photo courtesy of Dr John Hubbard and Homer Smith, Cornell University

by Frank Spitznogle

Understanding is a circular process involving analysis—breaking down an aspect of the world into comprehensible pieces. It also involves synthesis—recombining these pieces to express another truth. Computers have increasingly served to extend human analytical powers. Their ability to manipulate vast quantities of data has promoted the discovery of new relationships, and new truths. However, computers have not aided greatly in the synthesis of facts that leads to understanding. In fact, computer data and information processing abilities threaten to overwhelm humans with a flood of unmanageable and unrelated facts.

The artificial intelligence community is beginning to furnish new tools for understanding. In seeking to understand the nature of intelligence, and to duplicate it on computers, AI researchers are closing the analysis/synthesis circle. They are devising machines that can synthesize new knowledge from detailed facts and apply electronic intelligence to augment human problem-solving abilities.

AI is often defined as the science of enabling computers to reason, make judgments, and even learn. AI researchers seek to understand intelligence by modeling it in computer programs that can perform useful tasks. These tasks require indeterminate processes such as intelligence and decision making, rather than the determinate application of a set pattern or algorithm, as in traditional computing.

Emerging tools

Human intelligence is manifold and complex, and so are the programs and machines used to emulate even its simplest activities. The AI community has developed a number of powerful tools to deal with these complexities. The origins of time-sharing, bit-mapped displays, window systems, and specialized cursor position devices (such as the mouse) can be traced to early AI research. The Lisp programming language (rapidly gaining favor for a wide variety of applications), also emerged from the AI research community as a far more powerful tool for software development than that provided by traditional software environments.

Lisp's flexibility goes far beyond conventional computer languages. It is freely extensible, and full-function commercial Lisp systems offer a more powerful selection of programming tools than any other computational environment available today. However, efficient Lisp execution requires very specialized hardware architectures that have been commercially available only for the last few years. Although Lisp implementations exist for a wide variety of conventional computers, from personal machines to large mainframes, performance is marginal or unacceptable for many serious commercial AI applications.

Specialized Lisp machine architectures became generally available in 1980 and have begun to move into limited commercial applications. However, until recently, several factors have limited the diffusion of such tools into the marketplace. These factors include the uniqueness of the Lisp environment, the limited number of Lisp-trained individuals, the difficulty of migrating existing software to the Lisp environment, and the high cost of Lisp machines. The development of new Lisp hardware, however, together with progress in semiconductor devices, new bus-centered computer architectures, and a major cash investment in several startup AI-oriented firms have set the stage for rapid movement of AI tools into the real world.

The successful movement of AI from the laboratory to the commercial marketplace involves three challenges. First, it must be realized that AI is not a market or a product. It is a technology that can be applied to many problems in a number of markets. AI technology offers the opportunity to build products not before possible or practical with traditional computer technology. Such products will increasingly penetrate commercial markets. The transition will initially be embodied in AI-based attached processors supporting traditional computer systems. The areas will include intelligent database front ends, intelligent human interfaces to conventional systems, alarm advisors to support distributed process control systems, CAD/CAM systems, and diagnostic tools for complex electronic systems.

The second challenge is that AI is not a finished body of knowledge and will remain so for this decade. AI technology is rapidly evolving and increasingly market driven. As such, AI-based hardware and software support systems must be very flexible and quickly adaptable to changes in both technologies and markets.

Advanced architectures include hardware and software systems that facilitate migration of software in traditional languages to AI systems.

Finally, AI technology must migrate easily into existing, traditional computer systems. Commercial users cannot afford to render their investment in current software and hardware totally obsolete. AI technology must be an incremental addition to these solutions. It must offer gradual, step-by-step changeover to the benefits offered by intelligent computing.

Taking the challenge

New computer architectures are rapidly emerging to address the challenges involved in moving AI technology into successful commercial applications. The more advanced architectures include specialized hardware support for advanced software systems such as Lisp. They also include integrated hardware and software systems that facilitate the migration of existing software in traditional languages such as C, Fortran, and Pascal (operating under industry-accepted operating systems) to the new AI systems. Lastly, the advanced architectures can couple the specialized computer architectures required for advanced AI software and traditional CPU architectures into mutually supportive processor teams. These operate in parallel to solve specific problems.

Several new high speed bus architectures facilitate the creation of systems that can address the aforementioned factors. NuBus, developed at the Massachusetts Institute of Technology, and similar bus architectures such as Multibus II, are two examples. These high speed bus-centered architectures have sophisticated hardware and software to support multiple processor types. Individual processors

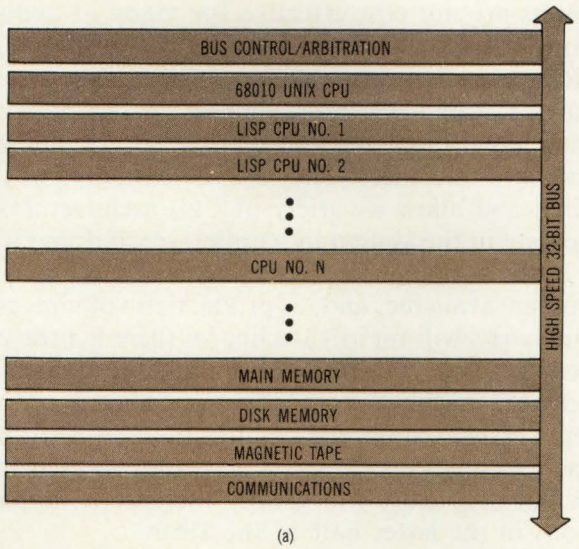

(a)

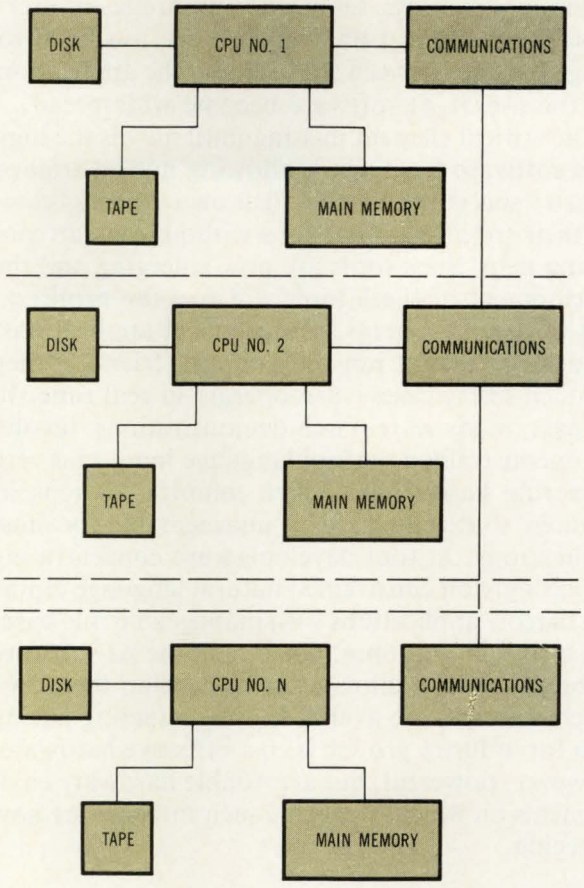

(b)

Present and future AI-based computers/workstations will increasingly use new bus-centered architectures (a), thereby replacing traditional CPU-centered architectures (b).

are but nodes on the bus. Memory, peripherals, and communication interfaces belong to the bus rather than to individual processors. This design allows very cost-effective systems to be built in modularly expanding forms.

In traditional architectures, sharing overall system resources effectively is much more difficult. Moreover, interprocessor communication and file transfers are relatively inefficient. This is because such data movements are limited by the communication link (bus) bandwidth. Finally, integrating new CPU types into existing architectures without major system redesign is also very difficult.

Bus-centered architectures alleviate many constraints of the more traditional CPU-centered architectures (see the Figure). For example, the bus-centered design allows excellent resource sharing, and a variety of CPU types can be resident on the CPU-independent bus. Bus-centered architectures also provide an efficient means to augment existing non-AI application software with powerful Lisp-based AI tools. Such systems are the initial practical applications of parallel processing where many different (or similar) processors work concurrently, as a team, in sharing system resources to address a problem. The LISP Machine, Inc Lambda family of machines, typical of what can be expected in the future, supports Lisp, Prolog, and Unix environments simultaneously. It also allows existing Fortran or Pascal applications to be ported and augmented by AI to provide auxiliary inferencing engines or intelligent operator interfaces to these applications.

Future applications

Desktop computers did not become pervasive until it was no longer necessary for the user to "program" them. The emergence of "umbrellas" of user-friendly software such as spreadsheets and simple word processing packages between a user and the native computer environment stimulated the rapid expansion of the desktop computer industry. Typical users normally do not program such systems; they simply "adapt" the available software to their applications.

The analogy is true in the movement of AI tools into practical applications. Until now, most expert systems have required well-trained engineers to be in the loop between the machine and the application.

As long as this is the case, the diffusion of expert systems will be slow. Only when umbrellas of user-friendly AI application software are available to bridge the gap between the user and the application will the use of AI software become widespread.

The critical element missing until now is the high level software framework allowing non-AI trained industry specialists to build their own expert systems for their specific applications without outside consulting help. Such tools are now emerging and the development of these tools will fuel the proliferation of expert systems into practical applications.

Such tools must not only be user friendly, they must, in many cases, also operate in real time. In the past, many AI research demonstrations, involving unconstrained natural language input and very large rule bases, have taken minutes, or tens of minutes, to execute. This is unacceptable for most applications. AI tool developers are concentrating increasingly on constrained natural language inputs and narrow applications with manageable rule bases to execute in real time. Once realtime AI software umbrellas (which allow users to develop their own expert systems) are available, the remaining pacing item for industry growth is cost-effective hardware. However, powerful, but affordable hardware environments on which to execute such software are now emerging.

As noted, conventional number-crunching computers are not cost effective for many AI applications. As such, new families of "specialized architecture" are beginning to appear. These systems have substantial hardware support for Lisp and offer easy integration into the real world via industry-standard interfaces. They are increasingly bus centered and allow a variety of CPU architectures to operate in the system in parallel—each doing what it is best at doing. The first systems of this type have become available, and, as production volumes escalate, costs will rapidly decline so there is little cost penalty over conventional architectures. Such AI "target" machines, coupled with realtime expert system software umbrellas that allow non-AI trained users to build expert systems to address their individual problems, will greatly increase AI applications in the latter half of the 1980s.

The greatest risk

One concern in the movement of AI tools into general use is "oversell." Because the human mind is extremely capable, systems will remain very limited in functionality compared to human intelligence. However, AI systems can be very effective in very narrow applications. In this domain, it will indeed be feasible to build systems that eventually become as good as a skilled human in performing a very specific task.

On the other hand, there is minimal likelihood that AI computers in this century will even begin to approach the wide flexibility of human intelligence. It is crucial that all AI participants—users and vendors alike—maintain a strong sense of responsibility in communicating the limitations, as well as the capabilities, of AI-based technology.

Frank Spitznogle is president and chief operating officer of LISP Machine, Inc, 6033 West Century Blvd, Los Angeles, CA 90045. He holds a BS in physics from the University of Kansas, an MBA from the University of Dallas, and a PhD in physics from the University of Texas at Austin.

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

High 719 *Average 720* *Low 721*

How we created an MDS that looks like a million, performs like 100 grand and costs under $20,000.

Introducing Emulogic's ECL-3200 universal development system.

It wasn't simple. And it took a lot more than seven days.

But when you start with the world's most advanced emulator, you're already way ahead. The ECL-3211 universal emulator gave our new MDS a repertoire of features you can't find in any comparable system. Like full-speed, no-wait-state emulation. Eight real-time, multi-function breakpoints for every chip. Real-time trace decode. And full-speed, full-range memory mapping with single-word resolution.

And then DEC™ made our job a little easier, too. We built their powerful new desktop computer–the Professional 300–right into the system. With RT11XM,™ 11/23 with MMU, 22-bit addressing, 512KB memory, 800KB dual diskette and optional 5 or 10MB Winchesters.

We didn't skimp on looks, either. We wrapped the ECL-3200 in a clean, compact package that, frankly, we call beautiful. But you tell us what you think.

You might conclude that, with all we've put on the inside, the price tag on the outside would be just as impressive. Not so! We brought the ECL-3200 in for 25% less than expected. So now you can buy a full-function development system for under $20,000–25% less than you'd expect. Including an 8-bit emulation package for the Z80,® 6502, 6809 or another popular chip. And you can also use it with all 16-bit emulation packages and our new high-level software development tools like C and Pascal cross-compilers and our just-introduced SLICE™ Symbolic Debugger.

So you see, creating our miracle MDS wasn't all that difficult. But we don't recommend anyone else try.

For detailed information, call 800-435-5001 or 617-329-1031 in Massachusetts. Or write Emulogic,® Three Technology Way, Norwood, MA 02062.

Professional 300, RT11XM and DEC are trademarks of the Digital Equipment Corporation. SLICE is a trademark, and Emulogic is a registered trademark, of Emulogic, Inc. Z80 is a registered trademark of Zilog, Inc.

EMULOGIC®

MICROPROCESSOR
DEVELOPMENT SYSTEMS

European Distributors: Austria: Walter Rekirsch, (43 222) 235555; Denmark: Instrutek, (45 5) 611100; France: YREL, (33 3) 9568142; Sweden: Aktiv Elektronik AB, (46 8) 7390045; Switzerland: Instrumatic AG, (41 1) 7241410; United Kingdom: MSS, (44 494) 41661; West Germany: Instrumatic Electronic GmbH, (49 89) 852063.

Qume can give your system extra appeal—with terminals made to order to your exact specifications.

We'll work from your specs, or be your partner in the development process. A dedicated Qume applications engineer will be with you from start to finish, to help identify your needs and recommend the right ingredients to make your product plan a success.

Choose a different interface. A unique command set. Even specialized electronics. No modification is too small, no specification too unusual for our experienced team of hardware and software design professionals. On the outside, select custom colors, your own labels and logo, proprietary packaging and owner's manuals. Without paying heavy premiums for these extras.

From custom design to assembly, we can do it all for you—or any portion of it. We'll build your terminal with the strict quality control that has earned us the industry's lowest DOA rate. We'll give you quick turnaround and delivery schedules you can count on. And we'll keep service simple with modular design and a ready-to-ship spare parts inventory.

When you really want it your way, we'll be there.

You'll find doing business with us easy, because we work hard to maintain your flexibility. Simple, straightforward contracts. No surprising incremental charges. And because Qume is a member of the ITT family of companies, you know we'll be there to serve you.

CUSTOM OEM TERMINALS

If you're looking for ways to beef up your system, call Mike Sugihara, Manager of OEM Marketing, at (800) 223-2479. And get the details from the one OEM terminal team that can really deliver it your way: Qume Corporation, 2350 Qume Drive, San Jose, CA 95131.

Qume®
A Subsidiary of ITT

SCHOOL

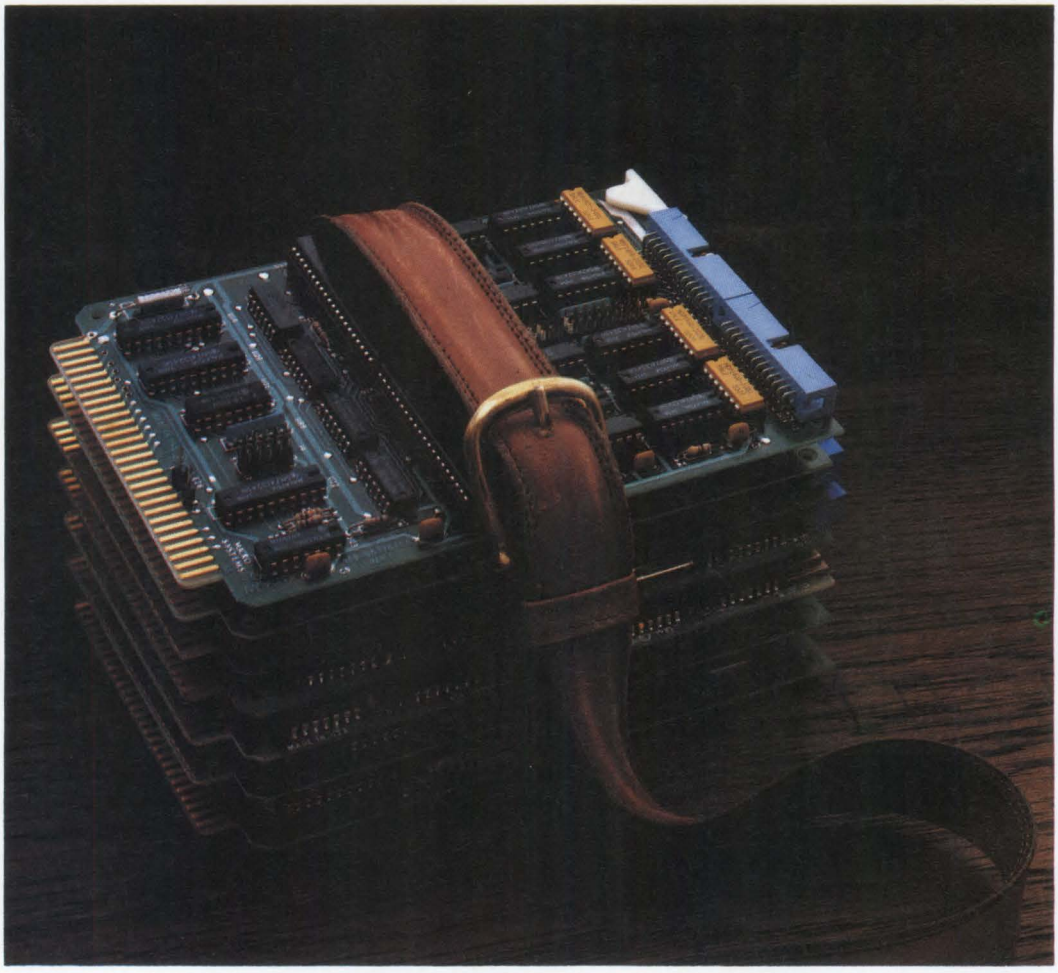

.ED BUS

Take advantage of our learning curve with the STD BUS.

Sometimes, there's just no substitute for experience. And that's why so many smart designers continue to develop new ideas with our MD Series™ STD-Z80 BUS-compatible boards. For energy, environmental and process control. Test and measurement equipment. Computer peripherals. High-speed printers. Medical electronics. Data communications. And an ever-expanding list of applications too numerous to mention.

Because for a wide variety of situations, our compact boards provide all the power designers need. At a price that's powerfully cost-efficient.

Not only that, our boards are already designed and proven in thousands of applications. Plus, they're assembled. Tested. Debugged. And modularized by function so you only have to buy what you need. And there are more than 30 boards to choose from. Available through Mostek or your local Mostek distributor.

We also offer development software and operating systems, plus a complete array of card cages and sub-system enclosures to ease and speed your design and packaging.

What's more, the STD-BUS is fully expandable. Which makes it simple and economical to add, delete, or interchange boards when you want to redesign or upgrade.

Take advantage of our learning curve experience with the STD-Z80 BUS. It's a time-proven course that can add a degree of success to your own system designs. For more information, contact Mostek, 1215 W. Crosby Road, MS2205, Carrollton, TX 75006, (214) 466-8816. In Europe, (32) 02/762.18.80. In Japan, 03/496-4221. In the Far East (Hong Kong), 5-681157.

MD Series is a trademark of Mostek Corporation.

CIRCLE 77

TOSHIBA HAS

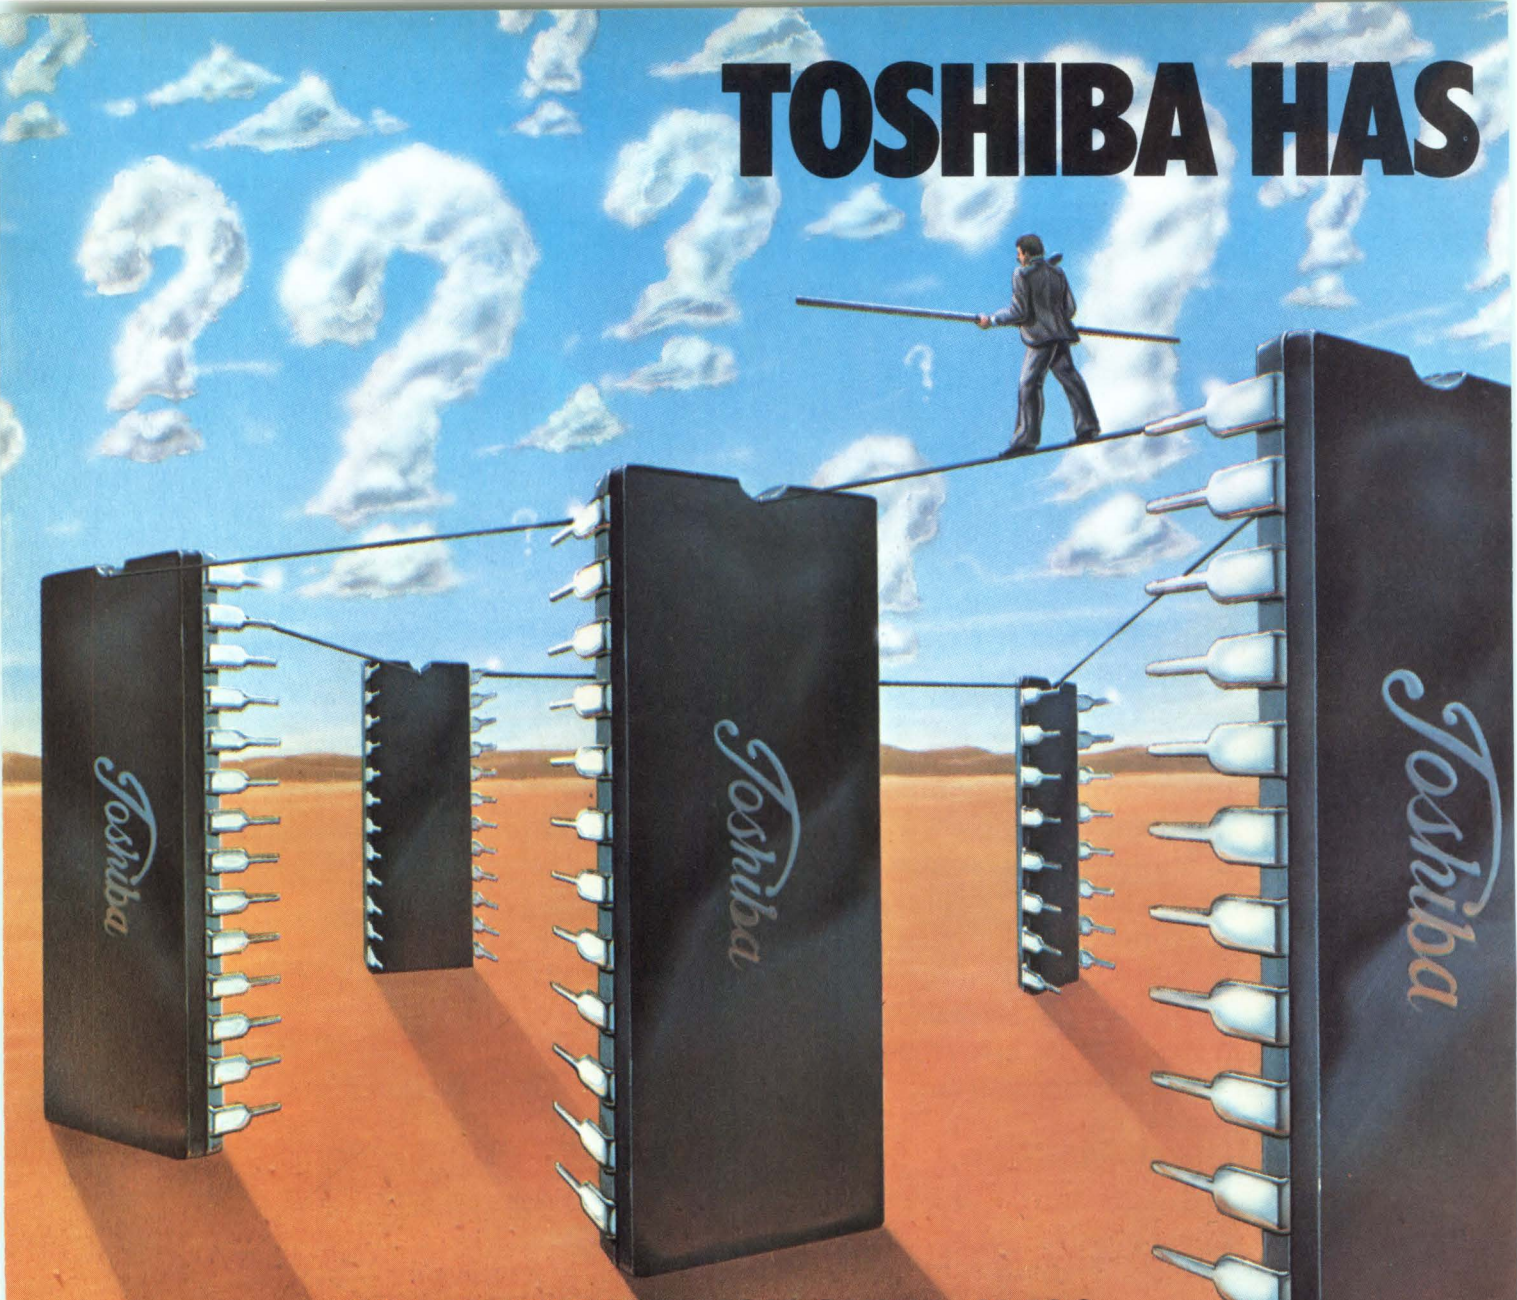

OUR NEW 74HC HIGH-SPEED CMOS LOGIC FAMILY COMBINES LSTTL SPEED WITH CMOS LOW-POWER.

Get off the design tight-rope with the 74HC series, a high-speed CMOS Logic Line with an operating speed thirty times higher than the standard C^2MOS logic.

This new generation of integrated circuits makes possible new applications in high-speed portable instruments which couldn't be achieved with current LSTTL or CMOS devices.

No matter what your design parameters, Toshiba America offers the most comprehensive CMOS Logic Line available today, ranging from our standard *4000/4500 series* through our higher speed *40H series* to our newest line, the *74HC series,* the fastest yet.

ARIZONA, Western Microtechnology, (602) 948-4240; **CALIFORNIA,** Image Electronics, (714) 730-0303, Integrated Electronics, (818) 998-2200, (714) 660-1055, (408) 263-1100, (916) 424-5297, Ryno Electronics, (619) 453-8430, Western Microtechnology, (408) 725-1660; **COLORADO,** A.C.T., (303) 422-9229, Integrated Electronics, (303) 292-5537; **CONNECTICUT,** Cronin Electronics, Inc., (203) 265-3134, Future Electronics, Inc., (203) 272-1852, Milgray/Connecticut, Inc., (203) 795-0711; **FLORIDA,** Future Electronics, Inc., (305) 291-2600, Milgray/Florida, Inc., (305) 647-5747, (800) 327-5262; **GEORGIA,** Milgray Electronics, Inc., (404) 393-9666; **ILLINOIS,** GBL/Goold Electronics, (312) 593-3220, Integrated Electronics/Intercomp Inc., (312) 843-2040; **KANSAS,** Milgray Electronics, Inc., (913) 236-8800; **MARYLAND,** Milgray/Washington, Inc., (301) 468-6400, (800) 638-6656; **MASSACHUSETTS,** Cronin Electronics, Inc., (617) 449-5000, Future Electronics, Inc., (617) 366-2400, Milgray/New England, Inc.,

THE ANSWERS.

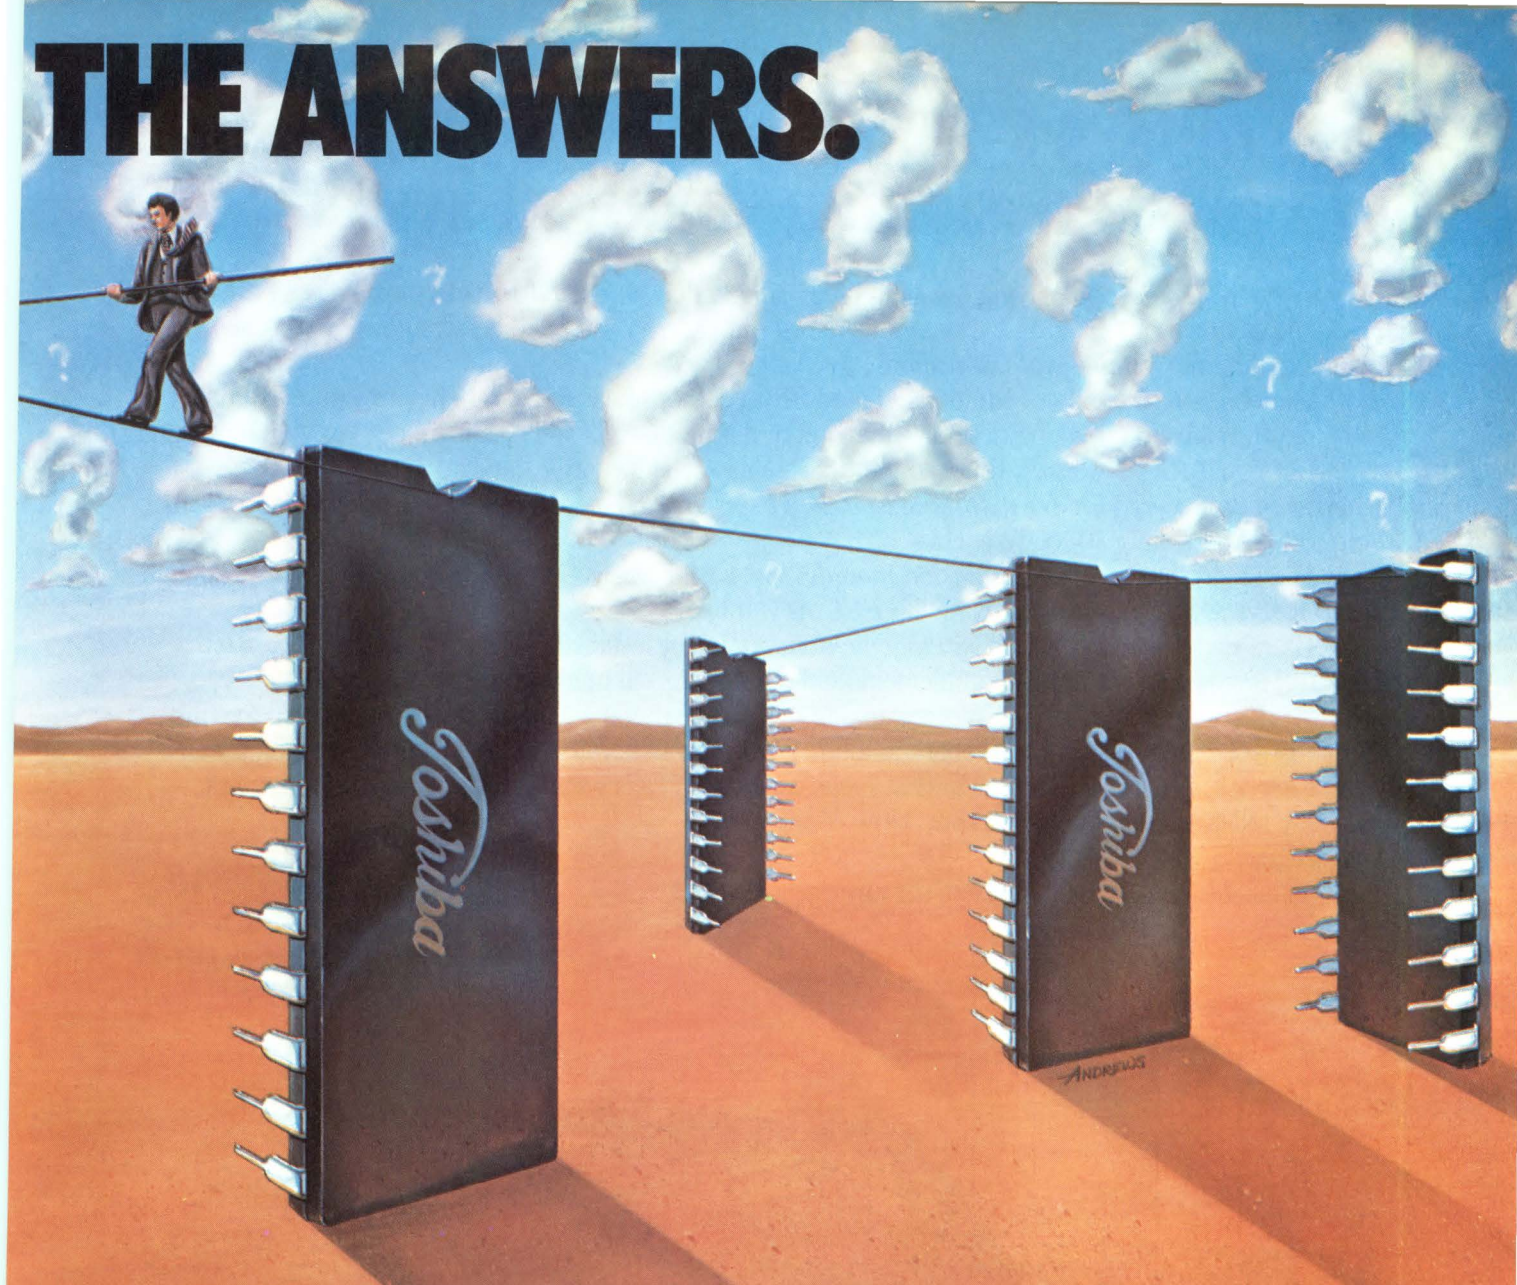

For more information about our entire logic line, write Toshiba America, Inc., 2441 Michelle Drive, Tustin, CA 92680, (714) 730-5000. Or call your local sales rep or distributor.

Wherever your imagination leads, now or in the future, Toshiba will have the answers you need.

Characteristics/ Logic Families	New HS-C²MOS (74HC Series)	LSTTL	*HS-C²MOS (40H Series)	Std. C²MOS (4000/4500 Series)
Prop. Delay Time (typ) GATE (C_L= 15pF)	8ns	9ns	15ns	125ns
Max. Clock Freq. (typ) J/K F•F (C_L= 15pF)	60MHz	45MHz	20MHz	2MHz
Quiescent Power Diss. (typ) (GATE)	0.01μW	8mW	0.01μW	0.01μW
Noise Margin V_{IH} (min)/V_{IL} (max)	3.5V/1.5V	2.0V/0.8V	4.0V/1.0V	3.5V/1.5V
Output Current $\|I_{OH}\|$ (min)/I_{OL} (min)	4mA/4mA	0.4mA/4mA	0.36mA/ 0.8mA	0.12mA/ 0.36mA
Op. Volt. Range	2-6V	4.75-5.25V	2-8V	3-18V
Op. Temp. Range	−40-85°C	0-70°C	−40-85°C	−40-85°C

*Data believed to be accurate and representative of each logic family.

TOSHIBA AMERICA, INC.

A WORLD STANDARD IN MOS.

(617) 272-6800; **MICHIGAN,** Prehler Electronics, (313) 473-7200; **MISSOURI,** L Comp, (816) 221-2400; **NEW JERSEY,** General Components, Inc., (609) 768-6767, Milgray/Delaware Valley, Inc., (609) 983-5010, (800) 257-7808, (800) 257-7111; **NEW MEXICO,** Western Microtechnology, (602) 948-4240; **NEW YORK,** Future Electronics, Inc., (315) 451-2371, Milgray/Upstate New York, (716) 385-9330, Milgray Electronics, Inc., (516) 420-9800, Rome Electronics, (315) 337-5400; **OHIO,** Milgray/Cleveland, Inc., (216) 447-1520, (800) 321-0006; **OREGON,** Western Microtechnology, (503) 629-2082; **TEXAS,** A.C.T., (214) 980-1888, (512) 452-5254, A.C.T. Houston, (713) 496-4000; **WASHINGTON,** Integrated Electronics, (206) 455-2727, Western Microtechnology, (206) 881-6737; **CANADA,** Carsten Electronics, Ltd., (613) 729-5138, (416) 751-2371, (514) 334-8321, Future Electronics, Inc., (416) 638-4771, (514) 694-7710, (613) 820-8313, (604) 438-5545.

CIRCLE 78

Microbar's K-2 family takes the "what ifs" out of designing your next system.

In the SBC business, playing "wait and see" simply doesn't make for leadership. It doesn't make market leaders out of OEMs and system integrators, either. What's needed is foresight—for prudent planning. And flexibility—for fast response to change.

Microbar's K-2 family of Multibus™-based SBCs demonstrates the idea. They're a pure example of look-ahead, ready-for-anything design—the kind of design that can add "front runner" to your reputation.

Our new ***LBX+™ boards,*** for example, meet and exceed the LBX standard. The "+," by the way, means you're prepared for either ***16- or 32-bit data transfer.*** And the LBX+ boards use the ***256K RAM*** chips, too. And feature 10- or 12-MHz speeds.

The K-2 family also has ***expansion modules*** that easily take you to two megabytes of on-board high-speed RAM. Or up to one megabyte with ***ECC,*** if you like.

There's a lot more. ***Memory management modules,*** for instance. That offer the Unix™-oriented ***two-level paging.*** And that make implementing the 68010 virtual memory a snap.

No doubt. The K-2 family reflects thinking that puts "what if" worries to rest. Of course, looking forward—objectively—is a Microbar tradition. We do it well because we're independent. And because ***SBC's are our only business.*** That's what keeps us ahead. That's what keeps you ahead. Better write—or phone toll-free—and find out more. Now.

Within California: (800) 421-1752
Outside California (continental U.S.): (800) 821-1011

MICROBAR SYSTEMS, INC.
785 Lucerne Drive, Sunnyvale, CA 94086

Your future system is our current project.

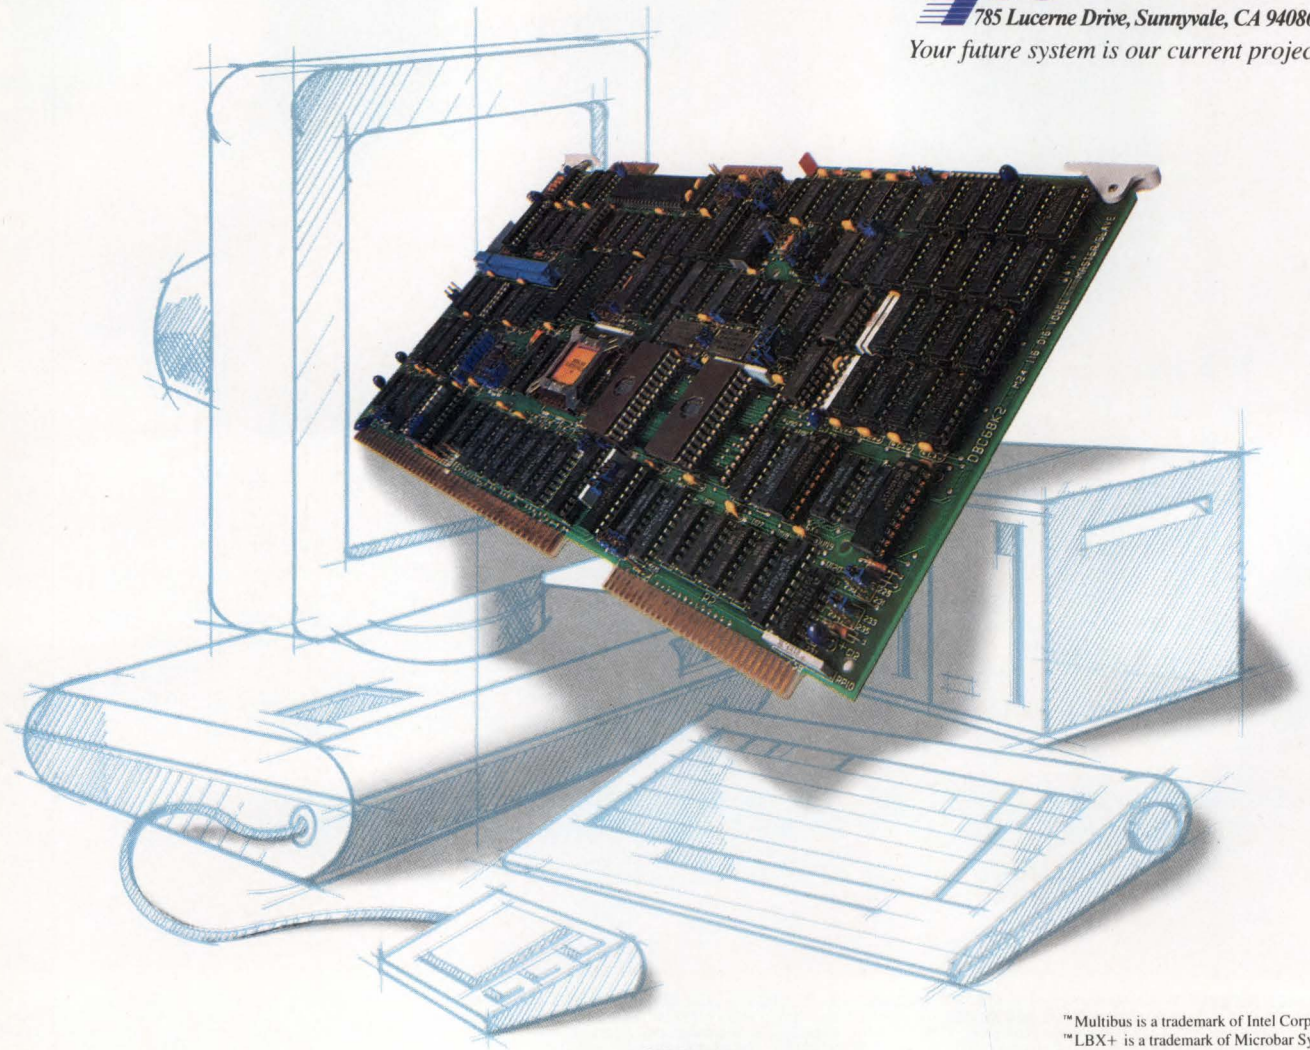

™ Multibus is a trademark of Intel Corp.
™ LBX+ is a trademark of Microbar Systems, Inc.
™ Unix is a trademark of Bell Laboratories, Inc.

CIRCLE 79

INTRODUCTION:
FUTURE HARDWARE

Advances in materials technology and design methodology will ultimately dictate the future of computer hardware. As very large scale integration of silicon proceeds to its physical limits, three major issues have arisen: how system designers can efficiently use such density; how circuit density can be increased even further; and how to obtain ever higher switching speeds. The call for efficient system design involves a discussion of advanced design tools such as silicon compilers.

Advanced design methods will be needed for engineers to use increasing circuit density effectively. Gate arrays and standard cells will serve the lower complexity end of the application-specific IC market. Design tools are already available for that purpose. Full custom designs, however, are a different matter. These will require the development of sophisticated silicon compilers to enable system designers to develop merchant market quality custom ICs. Such advanced design tools are emerging on the market. The next few years should see an explosive growth in silicon compilation methods to help designers cope with the increases in circuit complexity.

The search for ever-greater density leads to a consideration of CMOS, its ultimate limits, and beyond that, the infant science of molecular level electronics. The quest for faster switching speed demands an examination of the present rapid progress in gallium arsenide digital circuits.

Depletion and enhancement mode GaAs devices as well as High Electron Mobility Transistors are two technologies that are currently under development. GaAs is considered so important to the future of high speed computing that special emphasis is placed on the technology in a following article. An in-depth look at GaAs reveals that the fastest GaAs technology will probably be HEMT operating at cryogenic temperatures.

Cryogenics leads into the strange world of superconductivity to consider Josephson junctions, the partially understood physics of quasiparticles, and a device that takes advantage of that phenomenon. Ever-higher speed finally brings designers to the ultimate limit—the speed of light—where the report examines research into optical computing.

Much of the work in optical computing centers around development of digital techniques, particularly the building of optical logic that switches at extremely high speed. Often, ultrahigh speed computing handles complex pattern recognition tasks. Such functions, however, may turn out to be more easily handled by the analog technique of optical correlation than by digital methods. Consequently, much effort is focused on development of analog optical computers.

Beyond the technologies required for ultrafast and ultradense circuits, future hardware developments in mass storage and display technology will also have a major impact on future computer hardware. Hardware of all types generates the driving force to push computing onward and make possible the new architectures and software that will become our future.

John Bond

John Bond
Senior Editor

CIRCUIT DENSITY AND SPEED BOOST TOMORROW'S HARDWARE

Advances in materials technology and a comprehensive design methodology will yield the ultradense, ultrafast circuits needed for computers of the 1990s.

by John Bond,
Senior Editor

Forget architecture. Forget software. Neither matters without the hardware used to build the future architectures that will run the advanced software. The hardware, and consequently the architecture and software, can only move as far ahead as semiconductor advances permit. Paradoxically, advances in architecture and software are necessary before system designers can use the new semiconductor technology efficiently.

Progress ultimately depends on the system designers because there are not enough semiconductor engineers in the world to design all the chip-level systems that new architectures will create. Moreover, semiconductor designers do not have the specialized knowledge needed for system design. On the other hand, if the computer aided design (CAD) tools for system-level chip design are unavailable, the rush to new architectures and higher performance could be halted, no matter how fast semiconductor technology advances. It is more likely, however, that the next 5 to 10 years will bring the fruition of silicon-based system designs that can exploit present technology as well as the system-level use of yet-to-come semiconductor advances.

Even the near future promises to bring densities and speeds that will tax the abilities of the most imaginative designers. CMOS will approach its density limits and gallium arsenide will open up supercomputer speeds to everyone. But, the overriding issue is the development of the design standards and software tools that enable system designers to take advantage of evolving VLSI technology. Thus, each improvement in the architecture/software end yields advancements in VLSI, which in turn helps develop architecture/software systems. Looking to the far future, there are technologies that cannot be projected from today's semiconductors. Cryogenic circuits,

molecular electronic devices, and optical computer circuits are just a few of the long-range possibilities appearing to those who explore the outer fringes of material physics.

First things first

Today, it is not uncommon to see 100,000 transistors on a chip. In fact, very aggressive designs can integrate 500,000 or more transistors per chip. Eventually, each chip may contain as many as 100 million transistors. But first the industry must determine who can effectively use such density and how development of such chips will be accomplished.

According to Dataquest, Inc (San Jose, Calif) estimates, there are only 2000 to 3000 IC designers in the United States. That statistic alone ensures that system designers must design their own chips. But, there is another factor at work. Merchant IC manufacturers (where most of those IC designers work) make general-purpose VLSI devices for as wide a market as possible. Thus, while their processors, memories, controllers, and standard interfaces are elegant in terms of device densities and fabrication economics, they do not address the specialized portion of system design. The very thing that makes a system uniquely marketable cannot be supported by the semiconductor houses because of their own market exigencies.

In the past, designers have used standard logic families to build the nonstandard part of their design. Indeed, too many designs still go out the door implemented in standard SSI/MSI 7400 logic. The need to combine random logic into fewer packages has promoted explosive growth in application-specific ICs (ASICs). As a consequence of this increased demand, better CAD systems have been built to design more cost efficient gate arrays and standard cells. This situation will undoubtedly improve even further as more integrated design tools become available. There is, however, an inherent disadvantage to semicustom designs compared to full-custom designs: standard cells are typically 30 percent smaller than gate arrays and the same design in custom ICs saves an additional 25 percent in chip area.

Although gate arrays and standard cells are far better than random logic, they are both unsuitable for merchant market quality ICs. They are just too inefficient in their use of silicon. The gate arrays are

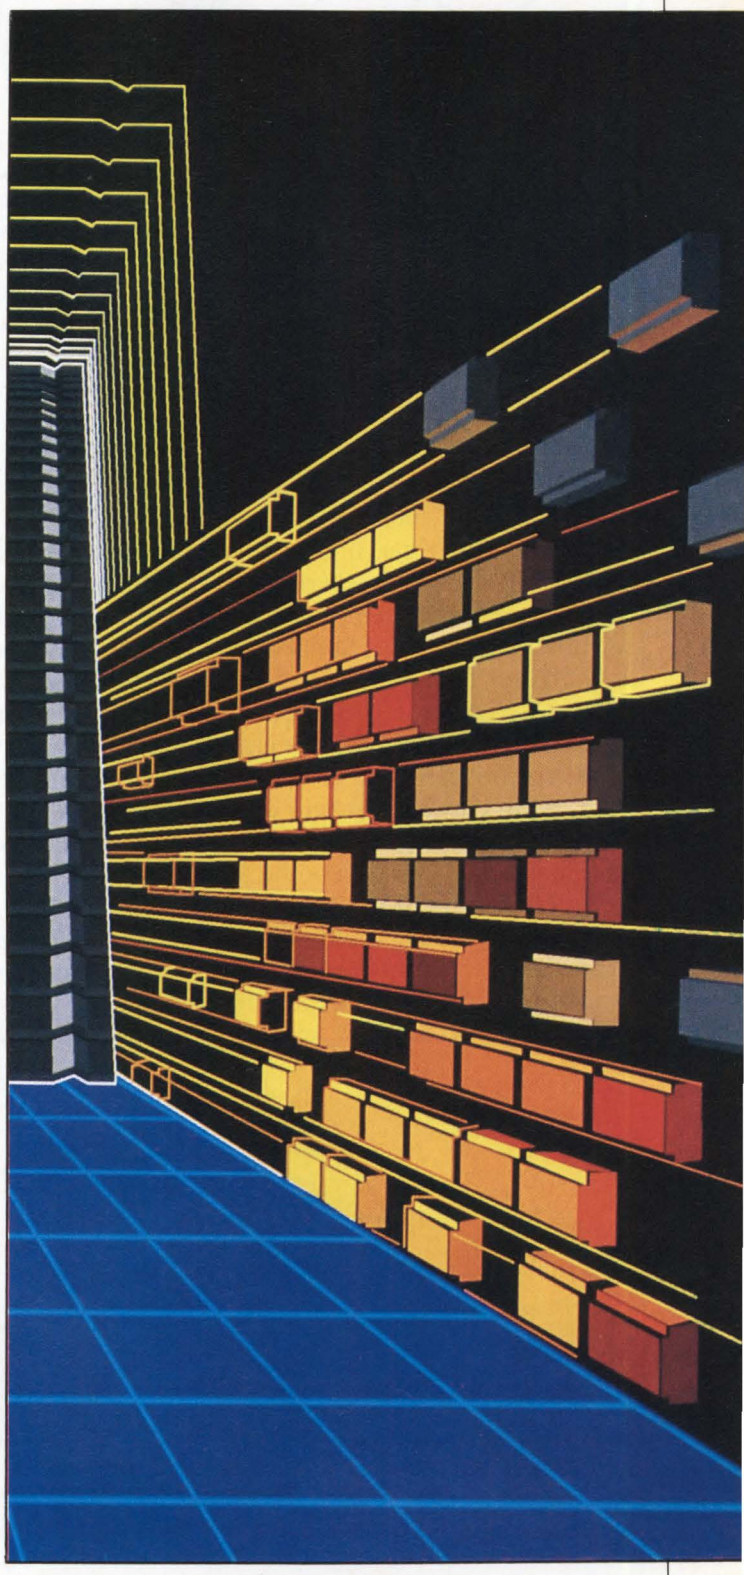

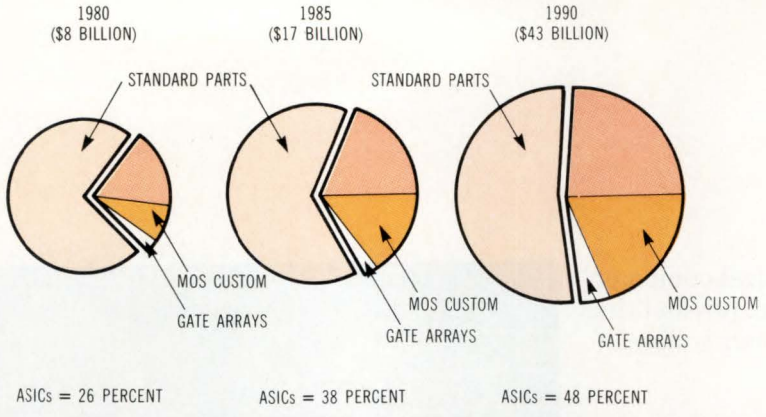

Charts of the semiconductor market show the growing percentage of application-specific ICs (ASICs) over a 10-year period. The growth in MOS custom devices will clearly outpace the rest of the ASIC market. (Courtesy of Silicon Compilers, Inc, from Dataquest, Inc information.)

inefficient because they are predefined and because, for larger arrays, much of the die area is used for interconnection. Furthermore, they cannot be used efficiently for memory or more complex designs. Standard cells can be considerably more complex, with libraries that contain memory cells and microprocessor cores. But practically speaking, there are limits to the number of large, unwieldy functions that can be contained in a cell library. Despite their obvious usefulness, therefore, all semicustom devices have practical limits on their complexity.

At the same time, higher development costs and long development times have conspired against custom ICs, making them unattractive except for high volume, low unit cost applications. But, because the advantages of custom designs are so important to system designers, more emphasis must be placed on further automating design methods in order to overcome the slowness and complexity of full custom design. Only the development of these comprehensive design systems will enable system designers to build merchant quality semiconductor systems.

Silicon compilation is one such automated design methodology. Just as a high level language compiler translates a computer language into machine code, a properly designed silicon compiler should translate high level architectural descriptions into the chip-level geometric structures of the equivalent circuit. Although only one company, Silicon Compilers, Inc (Los Gatos, Calif) is totally dedicated to this technology at present, there is considerable interest in the methodology and competition is not far off. Metalogic, Inc (Cambridge, Mass), for example, has developed and will market a silicon compiler based on artificial intelligence (AI) techniques.

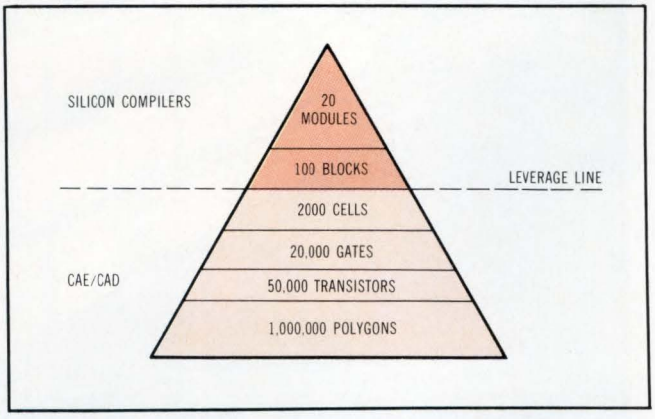

Silicon compilers manage complexity better than CAE/CAD tools because they work at a higher design level. This frees system designers from the chore of working their way down to the mask level.

At the 1984 Custom Integrated Circuits Conference, NTT Corp (Kanagawa, Japan) described an automatic logic synthesizer that, when integrated into the company's VLSI design system, generates a mask pattern from functional level description. This silicon compiler, although quite variable in its efficiency at this stage of development, nonetheless produced 10,000 gate circuits in only 3 hours. At the same conference, GTE Labs, Inc (Waltham, Mass) described the need for, and development of, circuit performance prediction to work with the MacPitts silicon compiler. The development demonstrated GTE's progress toward a useful silicon compiler.

In a move to supplant standard cells, VLSI Technology, Inc (San Jose, Calif) designed a compiler methodology that allows users to generate a large number of customized cell configurations from a small number of library cells. Typically, hundreds of cells must be available in a library to handle the range of function, load, and speed encountered. In VLSI Technology, Inc's system, instead of each cell having fixed geometry and performance, each cell varies over a much wider range. A principal advantage of such cell layout compilers is that because there are fewer cells in the library, less engineering time is required to develop the library and upgrade technology.

For a more detailed understanding of what a silicon compiler must do, it is helpful to look at the system developed by Silicon Compliers, Inc. Besides providing integrated design verification and documentation, the methodology used allows functional simulations, timing models, timing simulations,

power estimates, and layouts at each level of design. The compiler permits the designer to look at any design level; obtain estimates of size, power, and performance; and make incremental mid-design changes. It also supports exploratory design, enabling the designer to explore alternative architectures easily, and without wasting time. Since engineering is an interactive process concerned with trade-offs and technical compromise, a good silicon compiler must allow users to design by successive refinements.

An impressive example of Silicon Compiler, Inc methodology is Digital Equipment Corp's (Maynard, Mass) MicroVAX I chip, designed in only 5 months (18 man-months) with virtually the same transistor density as that found in commercially available microprocessors. A comparison of the technologies shows that gate arrays reduced the VAX-11/780's size by a factor of eight to yield the VAX-11/750. Meanwhile, the MicroVAX I, developed by silicon compiler methodology, reduced the VAX-11/750's size by another factor of eight while lowering cost by a factor of six and allowing a desktop implementation.

Obviously, silicon compilation offers some hope for system designers. This is not always apparent, as any survey of a confusing market filled with CAD/computer aided engineering (CAE) tools indicates. Often, these tools are excellent but incompatible. Furthermore, few attempt to solve the entire design problem from top to bottom. There is an appalling lack of standards in even such basics as terminology and performance specifications. Moreover, many of the CAD tools, including silicon compilers, have incompatible data formats. These problems will hold back progress, no matter how fast the basic semiconductor technology advances. However, there is some movement toward standardization, and 20 gate array manufacturers have formed a standards committee. But, much more still needs to be done.

If the U.S.-based companies do not address these problems, the Japanese will. The Japanese Fifth-Generation Computer Project includes plans to have computers design circuits. This application of AI techniques shows the U.S. a way out of the box it is in. VLSI advances will permit the development of AI and alternative architectures; this very development will enable engineers to advance system design at the VLSI level. It will take the development of standards and rigorous design methodologies, however, for such advances to occur. It will not be easy. Many CAD programs have hundreds of thousands of lines of code, complicating any attempts to merge data bases.

The main reason behind the development of sophisticated, computerized design tools such as silicon compilers can be stated in four letters: CMOS. This technology promises such high levels of integration that even the most skilled designers will find it impossible to take full advantage of it, unless they can design at a very high, architectural/conceptual level.

Swimming in the mainstream

There is no doubt that CMOS is becoming the mainstream technology. Future advances in dense silicon circuits will undoubtedly be implemented in CMOS. Furthermore, as it is scaled down, CMOS gets faster, competing not only with NMOS, but overlapping the lower end of the bipolar range. Of course, the constraints of speed/power product apply in CMOS as in any semiconductor technology, and

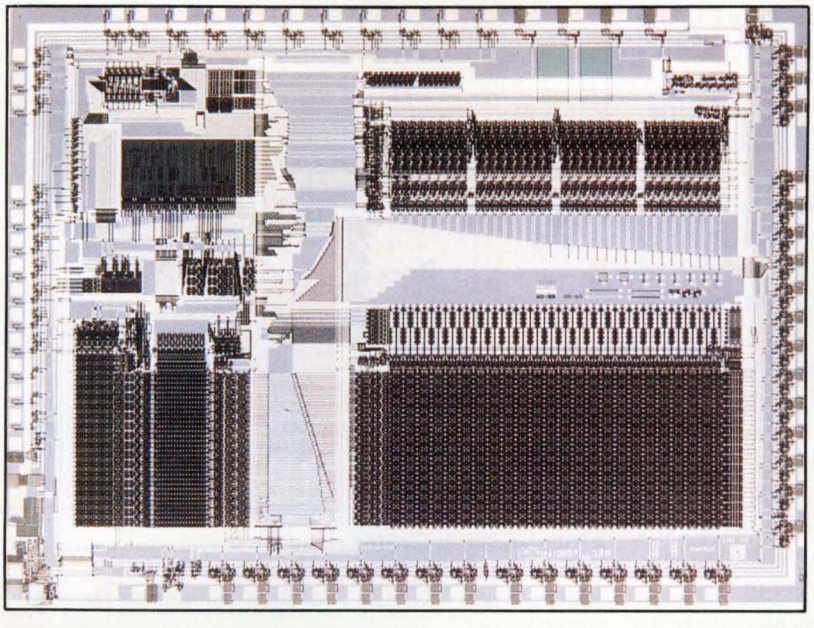

Digital Equipment Corp's MicroVAX I chip was designed by silicon compilation. The resulting custom system was one-eighth the size of a gate array version of the same system.

trade-offs must be made between density and speed. The speed/power product is better than any other technology and has been the main virtue of CMOS from its development in the mid-1960s.

Much of the speed advantage of modern CMOS has come about because of advances in lithography, process technology, and the adoption of n-channel design techniques. Consequently, CMOS processing has gotten much better. Some examples of this include shallower junctions, decrease in overlap capacitance of the gate, and the solution to the latchup problem.

Many of the advantages of modern CMOS derive from the adoption of NMOS design technology.

The latchup problem comes about because of the processing complexity of CMOS circuits, which produce parasitic complementary bipolar transistors as a side effect. This is not all bad, however. The parasitic device structures present in CMOS are used to good purpose in linear circuits, while the lack of such extra devices in NMOS makes it less useful for linear designs. The parasitic npn transistor, often used as an emitter follower for high output current, can form an undesired connection to a parasitic lateral pnp bipolar transistor. This coupling forms an unintended silicon controlled rectifier (SCR), which is the source of the latchup mechanism. The SCR causes low resistance from the power supply to ground. At best, this holds the circuit in a latched-up state until power is removed. At worst, it melts the circuit.

Two approaches have been used to control the SCR structure. The first approach uses dielectric isolation to separate the PMOS and NMOS devices by a dielectric such as silicon dioxide, silicon nitride, or sapphire (aluminum oxide). The second approach controls the SCR in junction-isolated structures. Among the various methods this approach embodies, epitaxial silicon is used as a substrate to reduce base-to-emitter resistance of the lateral pnp transistor and to prevent latchup. All of these methods are important because they do more than just prevent latchup—they represent an increased understanding of chip structure as well. This enables control over parasitic bipolar transistors, which can be very useful, as long as unwanted side effects are eliminated.

There are a number of other advances in CMOS that have made the process a firm candidate for the technology of the future. Because early CMOS was developed from PMOS, processing started with an n-type wafer and diffused p-wells into this substrate to contain the n-channel transistor while the p-channel transistor was formed on the substrate's surface.

Newer CMOS processes, however, have evolved from NMOS process advances. The starting substrate is p-type and n-wells are used for the p-channel transistor while the n-channel transistor is formed on the substrate. Switching speed is improved because n-channel transistors are favored. (Heavier doping of the p-well degrades the performance of n-channel transistors in the p-well type processing.) Faster n-channel transistors, coupled with the naturally higher mobility of electrons compared to holes, enables n-well processing to build chips where n-channel transistor speeds are optimized. This allows the chip to run faster overall. Also, the n-well CMOS process provides inherently better latchup immunity. Recently, a more complex, twin-well process that optimizes both NMOS and PMOS transistors has appeared on the scene.

However, not every great leap forward results from process advances. Better packaging methods such as plastic-leaded chip carriers, which were previously infeasible, now allow increases in system density. Circuit organization can also add substantially to the packing density of chips. For example, the basic NAND and NOR capability of CMOS eliminates many inverters. Furthermore, CMOS can use two or three times as many NMOS as PMOS devices for improved packing density. CMOS dynamic RAMs use NMOS cells with CMOS peripheral circuits for lower power consumption. AT&T's (Murray Hill, NJ) WE32000 microprocessor, for example, uses a twin-well process (twin-tub CMOS). But, many advances in density and speed come from clever circuit design such as dynamic CMOS used in the ALU, barrel shifter, and the arithmetic accelerator unit.

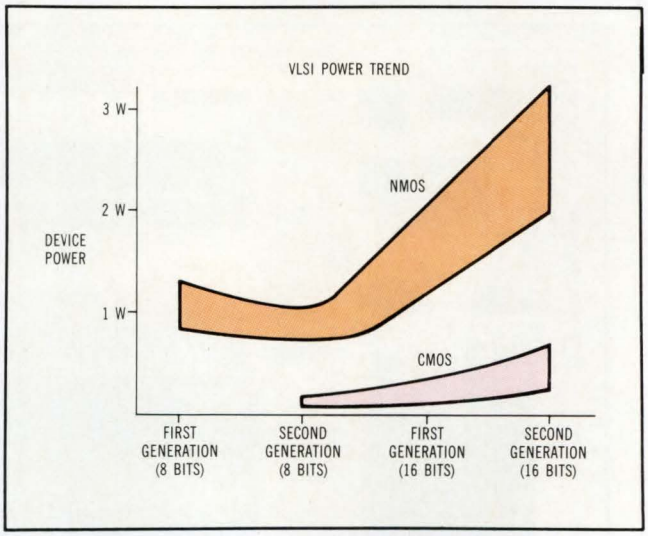

As the onchip transistor count rises with VLSI complexity, it results in higher power requirements for NMOS devices. CMOS offers an immediate solution to the NMOS power barrier.

WHY MULTIWIRE® IS SUPERIOR TO MULTILAYER FOR HIGH-DENSITY CIRCUIT BOARDS.

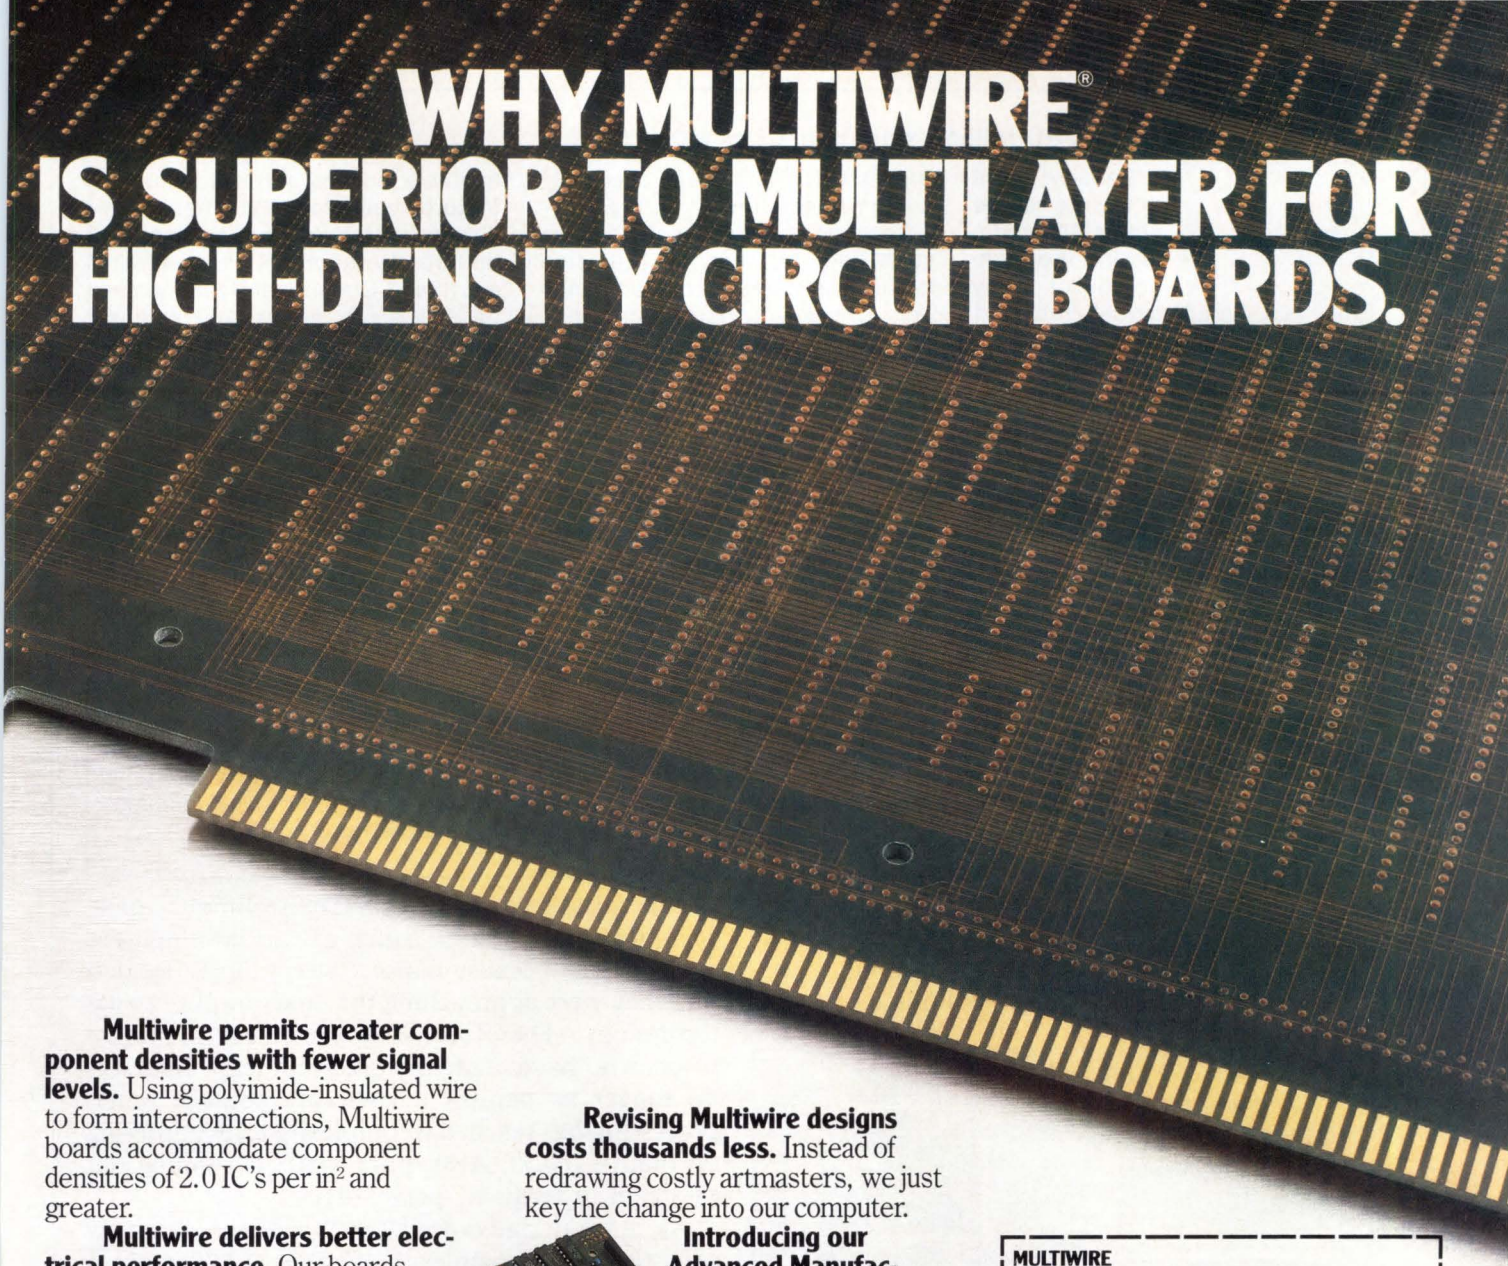

Multiwire permits greater component densities with fewer signal levels. Using polyimide-insulated wire to form interconnections, Multiwire boards accommodate component densities of 2.0 IC's per in^2 and greater.

Multiwire delivers better electrical performance. Our boards outperform multilayer in all applications requiring tight-tolerance on impedance controlled transmission lines.

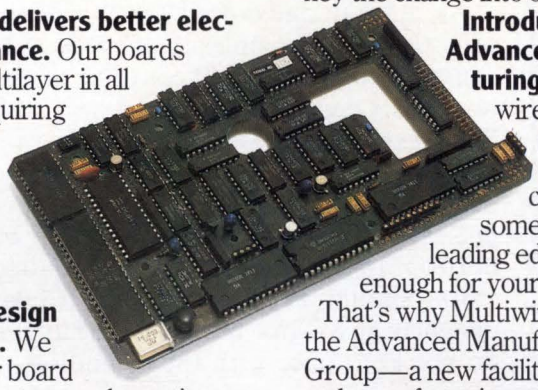

Multiwire shortens the design cycle by weeks. We can design your board from as little input as a schematic or net list, reducing your in-house design time. And when you make revisions, you'll get new boards back in days instead of weeks.

Revising Multiwire designs costs thousands less. Instead of redrawing costly artmasters, we just key the change into our computer.

Introducing our Advanced Manufacturing Group. Multiwire technology provides the most advanced circuit boards. But sometimes even our leading edge designs are not enough for your requirements. That's why Multiwire has established the Advanced Manufacturing Group—a new facility with design and manufacturing specialists dedicated to solving the interconnection problems for your next generation of products. To learn more, just fill out and return the coupon.

MULTIWIRE
10 Andrews Rd., Hicksville, NY 11801

- ☐ Please send me your 8-page brochure.
- ☐ Please send information on using Multiwire in ECL design.
- ☐ Please send more information about the Advanced Manufacturing Group.
- ☐ Please have a salesman call.

Name

Title

Company

Address

City State Zip

Telephone CD

MULTIWIRE® DIVISION
KOLLMORGEN CORPORATION

MULTIWIRE/NEW YORK-31 Sea Cliff Ave., Glen Cove, NY 11542 (516) 448-1428; MULTIWIRE/WEST-3901 East La Palma Ave., Anaheim, California 92807 (714) 632-7770
MULTIWIRE/NEW ENGLAND-41 Simon St., Nashua, NH 03060 (603) 889-0083; MULTIWIRE/ADVANCED MANUFACTURING GROUP-10 Andrews Rd., Hicksville, NY 11801 (516) 938-2000
®Multiwire is a U.S. registered trademark of the Kollmorgen Corporation.

CIRCLE 80

VLSI expert predicts 100 million transistors per MOS chip.

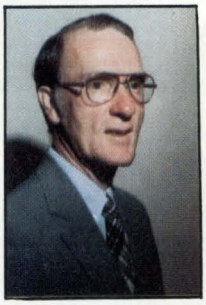

Density limits of MOS ICs can be predicted by projecting the technology of today across the next 15 years. This is according to Professor Paul Penfield, director of MIT's VLSI project. "There appears to be a fundamental limit to feature size," says Penfield, "because MOS transistors do not operate very well beyond three-tenths of a micron."

Of course, it is possible that other physical effects may be discovered before we reach that limit, he adds. "But to be conservative, we can assume .3 μm as an effective limit and still be able to predict a thousandfold increase in density." Since today's commercially available chips may have 100,000 active transistors, that would allow 100 million transistors on future ICs.

The advances necessary to reach such densities can be projected from present trends, explains Penfield. One way to increase the number of circuits on a chip is to increase its size. "As we acquire a greater understanding of semiconductor physics it is likely that we will be able to use chip areas that are 10 times as large as what we use today without any increase in the number of defects." If that scenario seems unlikely, consider that the difficulties inherent in wafer-scale integration make that technology even less likely. By comparison, lowering defect density to the level needed to increase chip area by 10 times is actually a very conservative projection.

"The expansion in chip size will allow 10 times as many devices on a chip," claims Penfield. "But the inexorable movement toward smaller features will allow 100 times more devices within the chip area as the feature size approaches its apparent limit." Consequently, he maintains, we can get to 100 million transistors if the rate of progress continues to be anywhere near what it has been. There are so many projects at work to increase IC density that there should be options available for every process step. "Over the next 15 years," Penfield states, "I am confident that we will find solutions for those problems that keep us from such large ICs."

Of course, if we learn to build three-dimensional chips or wafer-scale integration works, or if we experience some unforseen breakthrough, then we may have even denser circuits. "The real question," he adds, "is whether we know what to do with 100 million (or more) transistors on one chip."

Although it seems that there has been a quantum leap in CMOS performance, it is an illusion. Not only have those companies committed to CMOS made evolutionary progress, but additional progress has recently been made as NMOS techniques have been brought to CMOS. Much of this development has occurred because major NMOS houses decided that they were approaching the limits of pure NMOS technology. Therefore, it is the NMOS manufacturers who have made the quantum leap. But, it is no longer as painful a trade-off because NMOS processing has reached a complexity level bordering on that of CMOS. Also, pure NMOS is running out of steam in terms of power/density.

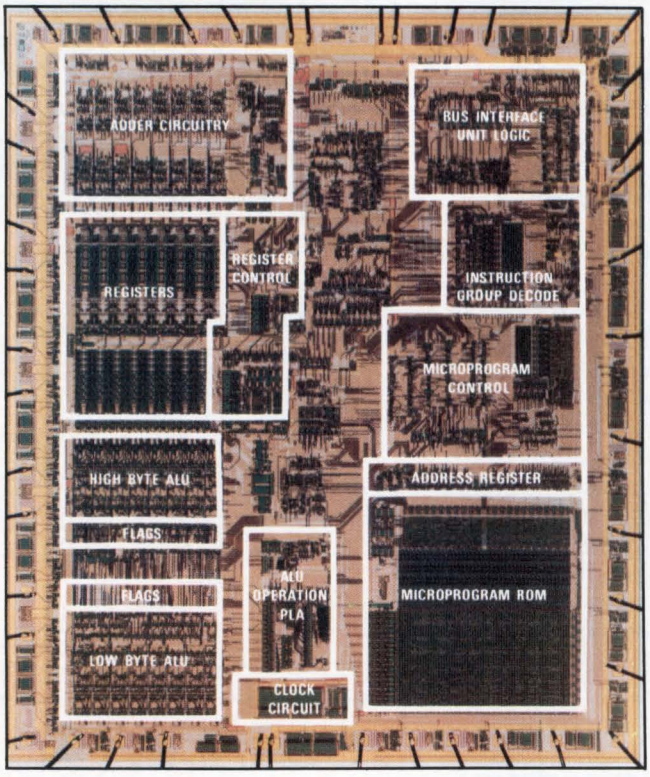

Layout of the Harris 80C88 shows the functional complexity of today's CMOS microprocessors. Standard microprocessors that until recently could only be built in NMOS are now possible in CMOS.

As a result, advanced CMOS designs can now offer the circuit complexity of NMOS at lower power consumption. CMOS will continue to evolve until it reaches its density limits in the next 10 to 15 years. Some experts think it is possible to fit as many as 100 million transistors on a single chip before CMOS peaks out. Beyond that, the future of VLSI is unclear, but molecular electronics may be on the far horizon. In the intermediate future (10 to 15 years), CMOS should be dense enough to keep system designers busy—maybe for all eternity if design tools do not keep pace with the technology.

Faster electrons

Whatever computer architectures evolve over the next 10 years, there will be a continuing requirement for very fast digital circuits in the high performance end of the computer spectrum. This high speed function, now dominated by silicon bipolar technologies such as Schottky and ECL, will gradually be taken over by gallium arsenide.

Cray Research (Minneapolis, Minn) is one company emphasizing GaAs, as demonstrated by its

AMCC metamorphosis

rapidly reduces your designs to high performance ECL/TTL logic arrays

At AMCC, metamorphosis means a lot more than just reducing board real estate, material cost and assembly time. AMCC metamorphosis transforms your designs from net lists to prototypes in as little as four weeks. It means a dramatic semicustom evolution using AMCC's system oriented logic arrays.

AMCC metamorphosis pushes aside past limitations in high performance system applications by offering high density logic at sub-nanosecond speeds; ECL, TTL or mixed I/O on the same chip; up to 95% gate utilization; advanced MSI/LSI macro libraries; automatic place and route software; the highest commercial and military quality and reliability standards (Mil. Std. 883); Class 10 fabrication; and mask compatible alternate sourcing.

Our engineer-to-engineer full service commitment plus complete CAD capability helps provide you with fast flexible design solutions. To us full service means support . . . a design implementation group; Daisy™ and Mentor™ engineering workstations; Tegas™ via Cybernet™; field applications engineering; on-site training courses; complete documentation; portable design centers; complete wafer fabrication; and quick turn-around assembly and test.

AMCC ECL/TTL LOGIC ARRAYS

	Q3500S	QH1500A	Q1500A	Q700	Q710	Q720
Equivalent Gates	3500	1700	1500	1000	500	250
Typ. Gate Delay (ns)	.3 – .7*	.9	.9	.9	.9	.9
Typ. Power (W)	3.5	2.8	2.5	2.0	1.2	.6
I/Os	120	120	84	76	56	36
Gate Utilization	95%	95%	95%	85%	85%	85%

*programmable

AMCC . . . dedicated to high performance semicustom. Make us perform for you.

Call or write for full information on how AMCC's metamorphosis can provide high speed silicon solutions to your system problems.

AMCC
APPLIED MICRO CIRCUITS CORPORATION
5502 Oberlin Drive
San Diego, California 92121
Telephone (619) 450-9333 • TWX/Telex: 910-337-1136

CIRCLE 81

Future of GaAs technology appears bright

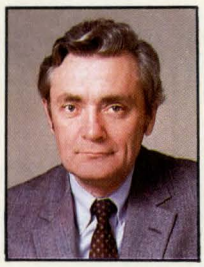

Although gallium arsenide will never be as pervasive as silicon due to fundamental disadvantages of the material, it may eventually dominate those applications requiring very high speed, such as supercomputers. So says Richard Soshea, cofounder and general manager of Harris Microwave Semiconductor.

Soshea contends that advantages other than speed and radiation hardness make the technology attractive, but he said he feels that the first round of GaAs products has not addressed them. GaAs' speed/power product, for example, is better than ECL's. Where a particular circuit stands on the speed/power curve is determined by processing—when built to run at normal silicon speed, GaAs offers very low power consumption. In the long run, therefore, GaAs might replace ECL entirely because it consumes far less power at the same speed. When designed for higher power (but still less than ECL), it is much faster. Also, GaAs, like ECL, offers good drive capability. Fanouts of five or six are easily achieved.

"The reliability of GaAs ICs is more than adequate for the most demanding applications in military and aerospace environments," Soshea says. "Harris uses the same technology for its GaAs ICs as for its GaAs discrete MESFETS that have been used in aerospace environments for about 20 years. It is already clear that these devices will have no significant reliability problems."

Soshea feels that present research is already pointing toward use of GaAs in supercomputers. "New GaAs technology will bring even higher speeds," he says. "Called High Electron Mobility Transistor (HEMT), it will come into its own over the next 5 to 10 years." HEMT will become an important technology for cryogenic computers, he maintains.

The electron mobility is higher than with normal GaAs because it is made with a mixture of GaAs and gallium aluminum arsenide in layers. Its principal advantage comes at cryogenic temperatures. "Compared to Josephson junctions, HEMT technology operates at higher temperatures. It only requires liquid nitrogen temperatures not the liquid helium temperatures common with Josephson junctions," Soshea adds. While today's GaAs exhibits gate propagation delays on the order of 100 ps, HEMT offers gate delays of 15 ps. HEMT scale of integration, however, will probably not parallel the standard MESFET complexity used in current GaAs circuits for quite some time, perhaps as long as 10 years, according to Soshea.

two-year joint effort with Harris Microwave Semiconductor, Inc (Milpitas, Calif) to develop GaAs digital ICs. Drawing from its experience with GaAs devices for microwave applications, Harris introduced a number of GaAs digital ICs this year. At the same time, Gigabit Logic, Inc (Newbury Park, Calif) has also introduced its own very high speed logic and counters in GaAs.

Despite a promising future for GaAs in high speed digital ICs, industry experts see it as complementary to silicon technology.

Hewlett Packard (Palo Alto, Calif) and Tektronix (Beaverton, Ore) are also showing interest in GaAs technology. Both support in-house GaAs research because of the technology's advantages for high speed test equipment. In this application area, the two instrument companies are interested in the technology's analog characteristics as much as its digital possibilities. In fact, Tektronix has gone public with its in-house project and announced the formation of a GaAs foundry to produce digital arrays for industry (see article, "High Speed Systems Look to GaAs for Low Power LSI," p 231). Thus, it appears that 1984 is the year GaAs will finally get off the ground.

However, GaAs expert Dr Richard Soshea of Harris Microwave Semiconductor does not think it will ever be as widespread as silicon technology, although he sees a bright future for it in extremely high speed applications. To Soshea, the material has fundamental problems in wafer fabrication that put it at a disadvantage when compared with silicon. It is very brittle and at high temperatures it decomposes. This makes it a much more difficult material to handle than silicon. (Nevertheless, crystals up to 6 in. in diameter have been grown in the laboratory.)

Also, GaAs material costs are much higher, and there is some difficulty in obtaining material of sufficient quality to get high yields. GaAs wafers presently cost considerably more than even high performance silicon wafers. And, even though prices are coming down, GaAs will be more expensive than silicon. GaAs is likely to be used only in those places where its unique characteristics of high speed and radiation hardness will pay off. Therefore, unless its unique advantages are important for a particular application, silicon will continue to be the material of choice.

Meanwhile, the industry is working toward developing a GaAs material that is uniform in its properties. Implantation would therefore result in a highly uniform threshold voltage. This uniformity is very important, particularly for some circuit designs. Because there is a tendency for the material to decompose at temperatures above 500 °C, the process must stay below those limits. Ion implantation

is usually used to dope GaAs since diffusion furnaces operate at temperatures high enough to decompose the material. Like a silicon wafer, a GaAs wafer will get damaged by high speed ions. But a comparable silicon wafer can be repaired by heating it to around 900 °C. A GaAs wafer would decompose before the crystal could be regrown. However, if the furnace is pressurized with arsine gas or the wafer is encapsulated with silicon nitride, arsenic migration is prevented. This allows the GaAs crystal to be regrown.

This ion implantation method has enabled the production of depletion-mode metal gate Schottky field effect transistor logic (DMESFET). These devices are easy to make, but are more power consuming than enhancement-mode devices. Within a year or two, industry should have a better understanding of fabrication techniques for the more-complicated-to-produce enchancement-mode MESFETs (EMESFETs). These promise less power consumption and lower operating temperatures for high speed systems. In addition to process complexity, GaAs requires complex metallization procedures. This is offset by the structural simplicity of GaAs circuits as compared with silicon ICs.

Despite the difficulties inherent in the fabrication of GaAs circuits, a bright future is still foreseen for these devices due to the high electron mobility of GaAs. The material is four to five times faster than silicon, allowing electrons to move faster and at lower voltages. Finally, because of the substance's insulating properties, switching times turn out to be very fast. With speeds that are two to five times as fast as the fastest ECL circuits, current GaAs ICs are already finding their way into some applications.

GaAs gate array processing starts in the sub-micron lithography processing area of the Tektronix GaAs foundry.

Although there is still a lot to learn about GaAs, there are already several types of applications where the payoff in using GaAs may be significant and immediate. One such application deals with very high data rates. A typical application would take a very high frequency signal containing data and bring it down to a lower frequency where further processing can be done by silicon circuits. For example, a demultiplexer or shift register can take a high speed data line and convert it down to a multiplexed set of lines that operates at lower frequency.

A major use for GaAs circuits will be in supercomputers both here and in Japan, where a large GaAs effort is underway. It is likely, however, that one of the first supercomputers to be built from

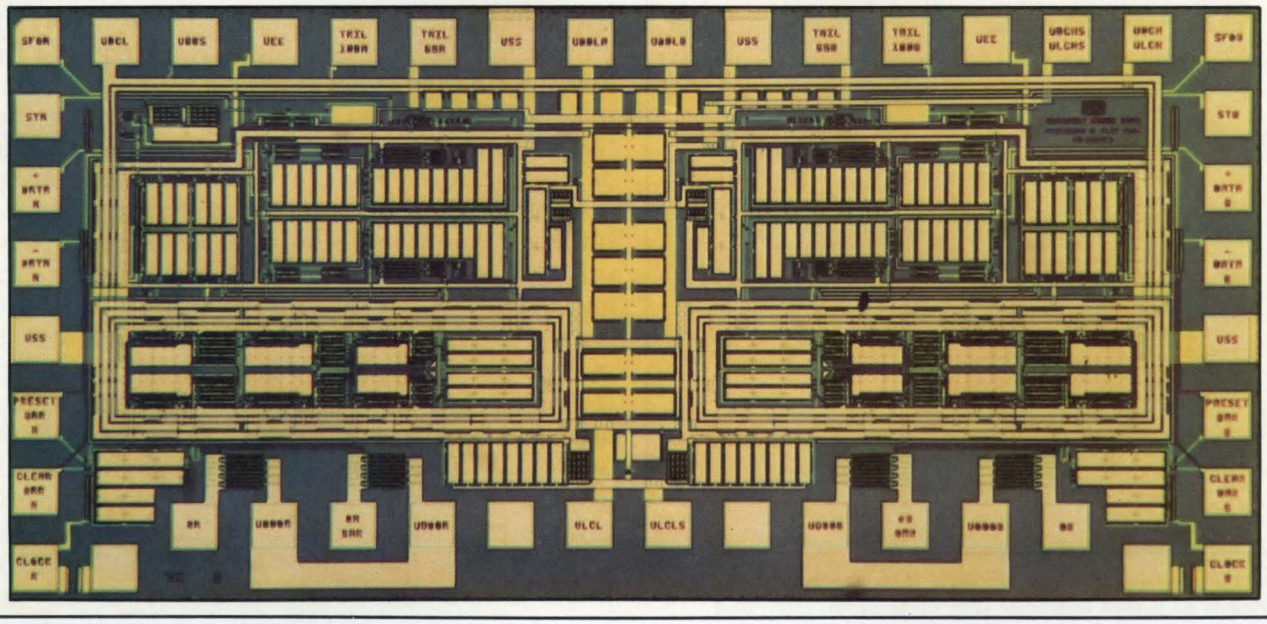

Very fast, standard SSI/MSI digital circuits are starting to be implemented in GaAs technology. One example is this 3-GHz, dual-precision D flipflop offered by Gigabit Logic.

GaAs digital circuits will come from U.S.-based Cray Research. Supercomputers will require whole systems to run at high speeds instead of single lines carrying high speed data. Consequently, the circuits needed will include standard logic building blocks (which are already becoming available) and very high speed LSI arrays.

With present GaAs technology, it is possible to build gate arrays of about 100 gates, and 1000 gates seem within reach. The very ambitious Department of Defense Advanced Research Projects Agency (DARPA) program is aiming for 6000 to 10,000 gates. That may be achieved eventually, but one interesting feature of gate arrays makes success difficult. In large gate arrays, the interaction complexity gets so great that it is difficult to maintain the system speed that was the reason for choosing GaAs in the first place. The chip's very speed works against itself. Thus, the whole concept of using a gate array for digital systems with greater than a 1-GHz clock speed has to be questioned.

That kind of problem led Harris to design and market a cell array in which some of the random, arbitrary structure of a gate array is eliminated with prewired dedicated cells. Larger gate arrays and even cell arrays have problems in maintaining speed and avoiding race conditions. In the Harris cell arrays, some circuits, such as amplifiers and flipflops, are prewired. The company's IC designers have had to consider how the circuit is likely to be used and where the signals enter and exit the chip. This is done to keep the total distance the signal travels along the surface to a minimum.

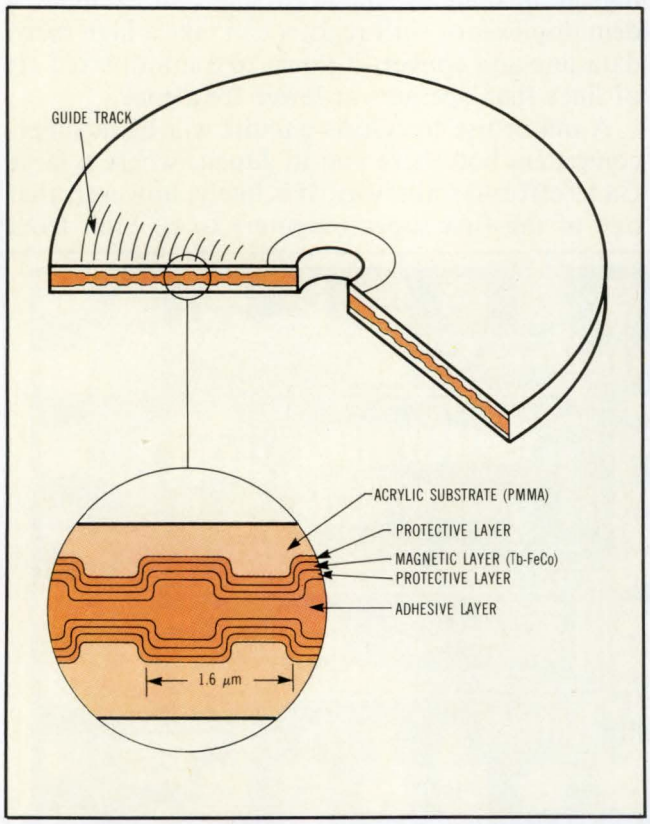

Sony's prototype magneto-optical laser memory disk consists of a thin magnetic film of terbium, iron, and cobalt vertically aligned magnetization. An applied bias field flips the polarity of areas exposed to laser light.

Such GaAs parameters as line width and feature size are already as good as or better than those of silicon.

Prepositioning circuits in different locations within the array limits user choice. A designer cannot, for example, call up a flipflop and put it in the upper right-hand corner of the circuit. This is because particular functions are already laid out for their most likely utilization. The most complex cells are located in the center of the circuit, and I/O circuitry is laid out near the outside of the chip. In the Harris cell array, the cells are of different sizes, more like a standard cell library than gate array macros. This has proved to be an efficient way to solve the problem of maintaining speeds in excess of 1 Gbit/s.

Because of the high speed, designers must leave significant distance between lines to deal with such analog effects as standing waves and reflected impedances. Going offchip forces the adoption of transmission line characteristic impedance assumptions. Harris products, for example, are intended to work in a 50-Ω environment. Impedances within the chip are typically higher and are converted at the edge of the chip. This is because a chip running at 50 Ω internally would require higher power. Due to the high speed at which these chips run, distances between circuits on the board and between lines and features on the chip itself are critical and must follow strict design rules to prevent cross talk. Ultimately, this may limit the density of the chips. But, this is not a feature of GaAs alone. Silicon operating at the same speed would exhibit the same problem.

Considering such parameters as line width and feature size, GaAs is already as good, or better, than silicon. Present products have 1-micron feature size and 1/2-micron will be standard in a few years. This is a little easier to produce in GaAs than in silicon, at present, because the circuits are less dense. The technology developed that way because analog devices, which have until recently been the mainstay of GaAs efforts, have required submicron gate lengths. These small feature sizes have been borrowed for more complex digital circuits. The Japanese have built single transistor, 1/4-micron length features. But this fine detail has not been carried over to more complex digital circuits. Now, commercially available 1-micron circuits are built

High-speed 16-bit multiplier/accumulators just got a lot faster (and cooler).

Visit us at WESCON (Booth #2400-2405) and ELECTRONICA

Now operate at 100 nsec, use only .9W of power!

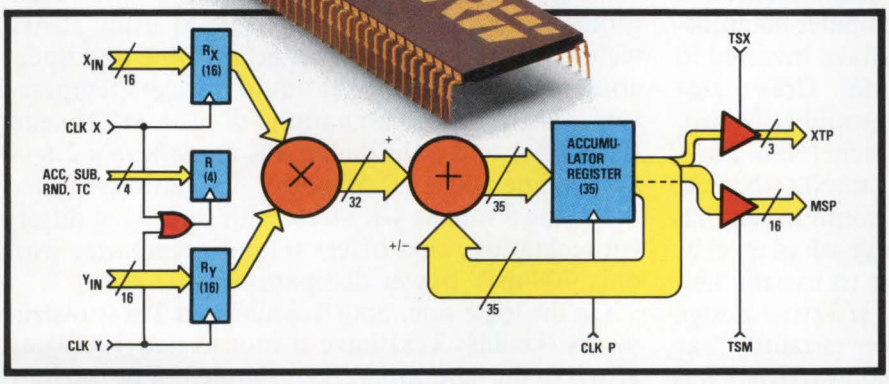

Our new 100 nsec* multiplier/accumulator (TDC1043) is 40% faster and 1/4 the power of its predecessors. This dramatic improvement is made possible through TRW's exclusive new OMICRON-B™ 1-micron bipolar technology.

With the TDC1043, you can increase your processing speed for superfast image or video processing, array processing or interactive graphics. And with the power you save, you can reduce the size of your power supply, design-in more devices, or even extend your product line into such areas as airborne or portable equipment. The possibilities are up to you!

*100 nsec (worst case T_A = O °C – 70 °C); 0.9W typical.

The TDC1043 provides full 35-bit double precision accumulation with optional subtraction internally, with a 19-bit extended output (16 MSB's plus 3 guard bits). It can also easily operate as a superfast multiplier by simply disabling the accumulate function. Packaged in a standard 64-pin DIP, its pin-out is compatible with the industry standard TRW TDC1010.

Our TDC1043 has a U.S. price tag of only $190 in 100s—which makes it not only the fastest, but the lowest priced 16-bit multiplier/accumulator available!

To get your complete data sheet on the new 100 nsec TDC1043 multiplier/accumulator, circle reader service card. Or, for superfast information, call or write our Literature Service Department:

LSI Products Division,
TRW Electronic Components Group,
P.O. Box 2472, La Jolla, CA 92038,
(619) 457-1000.

In Europe, call or write:
TRW LSI Products
Europe, Konrad-Celtis-Strasse 81,
8000 Munchen 70, W. Germany,
(089) 7103-0.

In the Orient, phone: Kowloon, Hong Kong, 3856199; Tokyo, Japan, 4615121; Taipei, Taiwan, 7512062.

TRW

LSI Products Division
TRW Electronic Components Group

CIRCLE 82

Optical digital audio disk leads to increase in mass storage

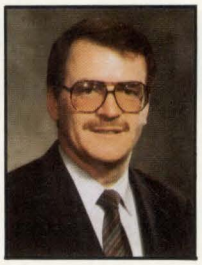

The future of inexpensive mass storage lies with optical memory, according to Sony Corp of America Vice President Kevin Finn. "As a first step in that direction," says Finn, "Sony adapted its compact (4.5-in.) disk digital audio technology to create an optical disk memory system."

Using this audio disk technology, the company recently introduced a compact disk read only memory (CD ROM) that has a recording density up to 1000 times greater than conventional floppy disks. "There are drawbacks," Finn admits. "For, like its audio disk counterpart, a CD ROM is unchangeable. Once a master has been made, it becomes impossible to add, delete, or manipulate the data." For users who want to create their own data bases, however, Sony is developing the next generation laser disk system called direct-read-after-write (DRAW). "This system gives users the ability to make their own recordings."

Development of the E-DRAW optical memory represents a further advance of the DRAW technology. With DRAW, once the data is recorded it cannot be manipulated or erased. With E-DRAW, however, the user has the ability to work with data as if the optical disk was a floppy disk. Sony's E-DRAW memory combines both magnetic and optical recording technologies. A prototype disk, consisting of a thin amorphous magnetic film of terbium, iron, and cobalt with vertically aligned magnetization has already been used in experiments.

The principle behind E-DRAW comes from the thermal sensitivity of the magnetic coatings, and from a write laser's ability to reduce the coating's coercivity at the point of contact to 0. An applied bias field within the E-DRAW drive flips the polarity of the areas exposed to the laser's light. When reading data, a polarized laser light scans the magnetic surface. The contact points that have been flipped by the write laser cause reflected light to be slightly rotated at a different angle than those unexposed. This phenomenon, known as the Kerr Effect, becomes the basis for creating a binary recording media.

"Research and development in the years ahead will be aimed at improving current access time for optical media in order to eventually parallel those rates now available with magnetic media," says Finn.

using light projection techniques, not X-ray or electron beam projection.

If there is any lingering doubt that GaAs has a strong future, consider the many computer manufacturers in the U.S. and Japan that have invested in the technology. The list includes Cray, IBM (Armonk, NY), Honeywell (Minneapolis, Minn), AT&T, Fujitsu (Atsugi, Japan), Hitachi (New York, NY), Toshiba (Kawasaki, Japan), OKI, NEC (Natick, Mass), NTT, and others. A major commitment has been made, and the problems are being solved quickly. Cray will undoubtedly be the first to use the new technology because its odd, but effective design strategy lends itself to using simpler circuits. Cray I logic is implemented in high speed ECL NOR gates (two gates per package) and uses pipelining for high computational speed. The Cray II uses a 16-gate array. Available DMESFETs would not only be much faster, but denser as well. Without any change in architecture, present GaAs circuits could keep Cray in the forefront of scientific supercomputers.

Other major computer manufacturers such as IBM and Fujitsu, however, make computers that are more general purpose in nature than Cray. Therefore, they cannot use a similar architectural strategy. Instead, they have committed their GaAs research to a more advanced technology that will combine very high speed with VLSI density. In the next few years this will probably mean EMESFETs, and in the long run, High Electron Mobility Transistor (HEMT) circuits. But, it may be quite a while before these approaches develop the low gate dissipation levels needed for VLSI GaAs circuits.

Progress is being made at the LSI level, however. Fujitsu Labs built a 1-K static RAM using HEMT technology that offers .9-ns access time and operating power of 360 mW at liquid nitrogen temperature. Using a combination of EMESFETs and DMESFETs, NTT Labs built a 4-K SRAM with a 2.8-ns access time and a 1.2-W power dissipation. Fujitsu has made a similar 4-K SRAM with somewhat different technology that offers similar access time with only 700-mW power dissipation.

On the logic side, both Toshiba and Texas Instruments (Dallas, Tex) have demonstrated 1000-gate arrays in the laboratory. So it should not be too long before 1000-gate arrays and 4-K SRAMs will be available on the market. In the meantime, the SSI/MSI logic circuits made by Harris and Gigabit Logic are "real," available, and faster than ECL. At least system designers will not have to worry about advanced CAD tools for a while.

Beyond semiconductors

As the thrust toward speed and density continues in semiconductors, other advances in material technology will have an equally important effect on computer hardware. Examples include mass storage devices that will get denser and faster as new techniques are developed, and flat-panel displays that will finally begin to live up to expectations as manufacturing technology improves.

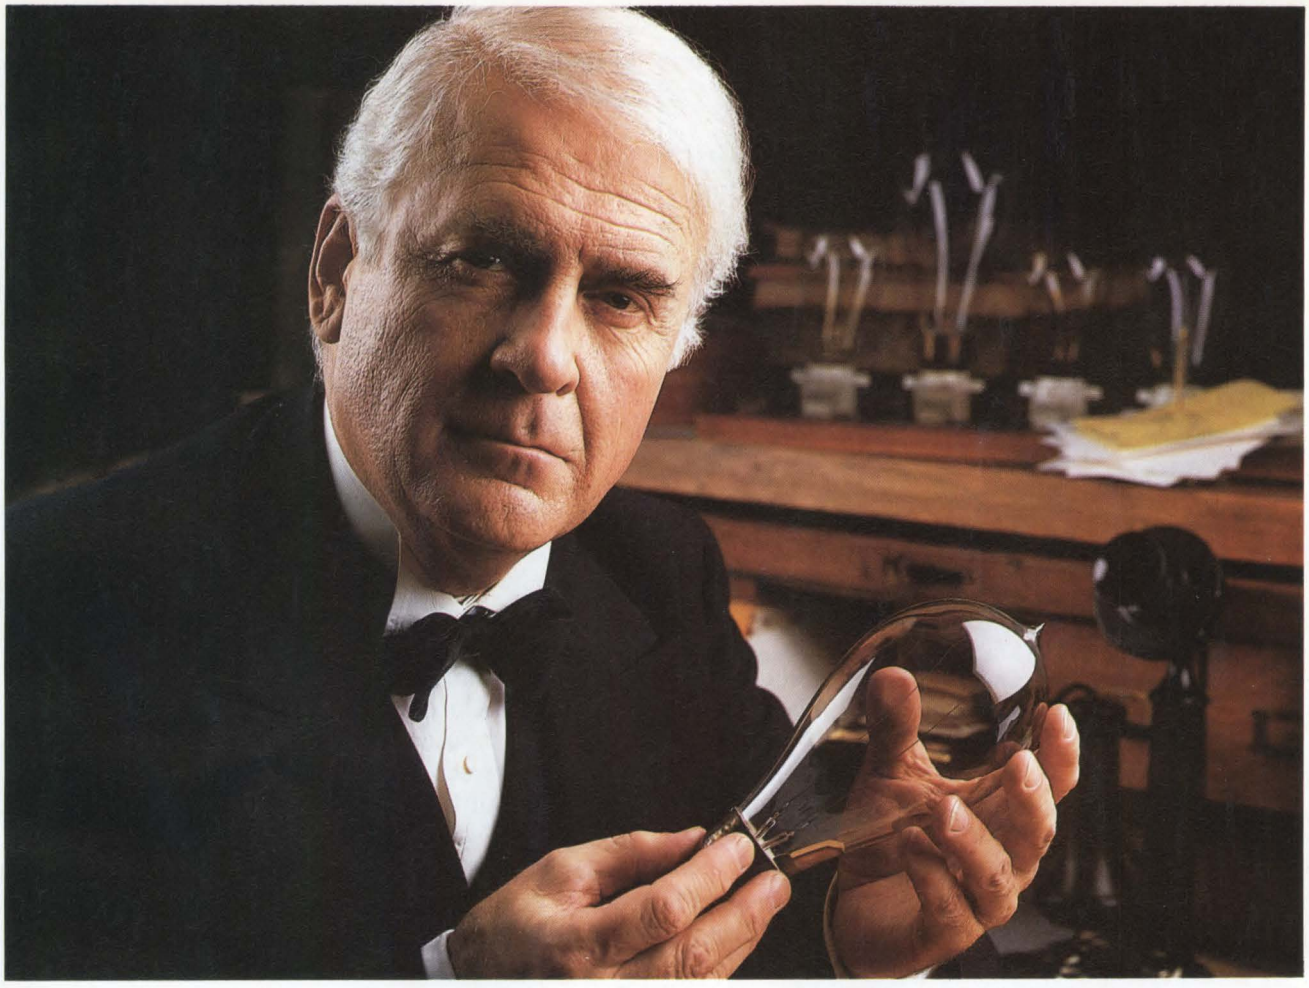

Zetaco's latest inventions make 86 disk and tape drives Data General BMC compatible.

Zetaco's BMX-1 disk drive controller and BMX-2 mag tape coupler give the user of Data General's high-speed BMC (Burst Multiplexor Channel) full compatibility with the newest high-speed, high-capacity, non-DG drives.

These two bright ideas let you choose virtually any high performance SMD-type disk drive or ½" tape drive to run on the BMC or Data Channel on DG's newest Eclipse or MV Series mini.

Each BMX-1 supports up to 4 SMD-interfaced hard disk drives. Each BMX-2 supports up to 8 streaming or start/stop tape drives.

Zetaco's full emulation of DG's 60XX and 61XX series disk and 6026 and 4307 tape subsystems means no software patching is required. E^2 PROMs eliminate switches, making drive configurations and functions selectable via downline loaded software. And Zetaco's exclusive backplane cable design is plug-and-go compatible with DG's FCC and non-FCC compliant chassis.

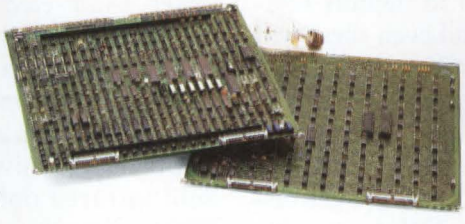

Discover the BMX Series from Zetaco, 6850 Shady Oak Road, Eden Prairie, MN 55344, (612) 941-9480. Telex 290975. European Headquarters: 9 High Street, Tring, Hertfordshire, HP23 5AB, England. (044282) 7011. Telex 827557. ZETACO G.

ZETACO™

Controller Division, Custom Systems Inc.

Metallization schemes are very important in GaAs processing because of the present MESFET approach. It requires a special metal alloy to be vacuum deposited directly on the ion-implanted GaAs surface. (Photo of thin-film metal deposition equipment was provided courtesy of Tektronix.)

Increases in mass-storage density will result from the continual evolution of magnetic disks that use plated thin-film media and vertical recording techniques. A more dramatic development that is already available, however, is the optical disk. With gigabytes of memory, the first optical drives on the market are write-once systems that use a laser to read or write at higher power. The difference in reflectivity between a written spot and an unwritten area is detected as data.

The next generation of optical drives will be fully writable. Sony Corp (Park Ridge, NJ) has taken a major step toward bringing that generation closer with its development of a prototype magneto-optical laser memory disk. Inexpensive gigabyte memories may transform the way we use computers, but new display technology will change what we see. Displays are being updated with flat-panel technology as manufacturing of these units improves and prices tumble.

Optical computing would provide the ultimate in high speed, but will need sophisticated analog technology.

Portable computers have begun to include relatively large LCDs in lieu of CRTs, and even these will soon be replaced by full CRT-size (80 column x 24 lines) flat-panel displays. But LCDs are only one part of the flat-panel market. Over the next decade, still greater growth is forecast for flat-panel displays that use thin-film electroluminescent technology, gas discharge (plasma), and flat CRTs. Offering some promise for the distant future are such technologies as electrochromics and electrophoretic screens. CRTs, however, will not disappear. They will continue to evolve and will dominate the high performance graphics market.

The outer limits

For a decade, it appeared as if Josephson junction devices might be the technology of the future for high speed computing. Operating at liquid helium temperatures, the superconductor technology had the potential to permit line thickness of only a few thousand angstroms and frequencies up to 10 Hz. IBM actually built logic with gate delays of 44 ps and power dissipation of 5 μW. An experimental memory chip was built that exhibited less than a 4-ns cycle time. A basic problem was that every part of a computer system had to be built in Josephson technology, a large but not insurmountable task. However, there were process problems that compounded the problem. The lead alloy superconductors that IBM used were subject to failure during the cycling between cryogenic and room temperatures. The substitution of niobium junctions introduced a whole new class of problems that had to be solved. After 10 years, IBM finally gave up and sounded the death knell for Josephson technology in the U.S.

For the time being, Josephson junction technology survives in Japan. The Scientific Computer Research Project in Japan is concentrating on three areas: GaAs, HEMT, and Josephson technology. It is likely, however, that the consortium of manufacturers [Fujitsu, Hitachi, NEC, Mitsubishi (Tokyo Japan), and NTT] developing the technology for a 1989 computer system will drop Josephson junction research by 1985.

Rising out of Josephson research at IBM, a newer superconducting switching element is under investigation at Hypres, Inc (Elmsford, NY). Called the quiteron, or Quasiparticle Injection Tunneling Effect device, it possesses transistor-like characteristics and very low switching energy. The device consists of three superconducting electrodes separated by two tunnel barriers. Before devices can be built for practical applications, further investigation into the physics of nonequilibrium superconductivity is needed.

Ultimately, the quest for even more speed may lead to optical computers that could operate much faster than electronic computers. While HEMT promises switching speeds on the order of 15 ps (much faster than silicon or even room temperature GaAs devices), optical switches with switching times of 1 ps are within practical reach. Researchers at Heriot-Watt University (Edinburgh, Scotland) have built infrared optical AND, OR, and NOT gates using small indium antimonide crystals.

Others have used GaAs to build such bistable optical devices. These semiconductor materials are opaque to visible light but transparent to infrared. The material chosen for bistable devices must exhibit a nonlinear refractive index. Bistable devices have two stable states at the top and bottom of a hysteresis loop where the intensity of transmitted light varies only slightly according to the intensity of the incident light from a laser. The steep curve between these two stable regions represents the switch area from one bistable state to another. Small increases (or decreases) in incident light flip the device from one stable state to another. It is analogous to the transfer-characteristic curve of a transistor.

Recent advances from the Massachusetts Institute of Technology (MIT) report the shortest pulses of laser light yet produced—16 fs. The research is aimed at producing the highest possible switching speeds for optical computing.

Optical correlation

It may turn out, however, that the best way to perform optical computing is to use analog rather than digital techniques. Even the high processing speed of optical digital computers may be overwhelmed by the mass of detail in real world images. Image processing can be accomplished by a technique called optical correlation. Rather than scanning each line and performing calculations on each pixel, optical correlation systems use lasers and holograms for pattern recognition and feature extraction. Thus, such systems tend to work like human vision by recognizing overall patterns rather than trying to reconstruct images from small elements.

The method is not unique to image processing, however. Voice and radar analysis, as well as Fourier analysis and array processing, yield to pattern recognition techniques. With their unique capabilities, such optical processors may eventually act as the pattern recognition interface between the world and digital computers.

Obviously, there are many possible avenues left to explore in the quest for greater speed. The goal of ever-greater VLSI density may not be as easy to achieve. Even with breakthroughs in lithography and process technology, the final limits of CMOS are becoming clearer. Unless there are major breakthroughs in the next 10 to 15 years we may find ourselves at a technological dead end, albeit a very useful dead end filled with huge single-chip systems.

Since such a state of affairs does not inspire engineers, researchers have begun to investigate the possibilities molecular computing might offer. In this case, molecular refers to device size rather than implying that the devices are separate molecules. Although such concepts have been discussed for perhaps 10 years, it was not until 1981 that Forrest L. Carter organized the first international meeting on molecular level electronics (under the auspices of the Naval Research Laboratory in Washington, DC).

Present research is looking at material switching phenomena, sensors, and theories, and at models of molecular level devices. While conceptual prototypes of devices have been discussed, actual fabrication remains a future problem. Perhaps 15, 20, or 30 years from now, a whole new technology of molecular computing will make silicon seem old-fashioned. But the road to ultradensity remains more uncertain than the way to ultraspeed.

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

High 722 *Average 723* *Low 724*

IT TAKES A LOT OF DRIVE

In CAD/CAM. In seismic and transaction-based systems. Or PBX's. Or anywhere else you need a lot of data fast.

That's why we created the Ampex 825.

This 14" Winchester family delivers everything you need to take the lead in multi-user, multi-tasking applications: 825, 660 or 330 MB of unformatted capacity.

The new standard 1.859 megabyte per second data transfer rate.

Access times of 21 milliseconds.

Plus the best price per megabyte on the market.

Fact is, we offer a better 825 MB price than any combination of small drives. And–thanks to a design that teams RLL encoding with proven, standard head and media technologies–better yields than other big drives. So you won't have to wait for us to get our act together in manufacturing.

TO MAKE IT BIG IN PICTURES.

Or in QA. Or field service. Because we designed the Ampex 825 with features that keep it on the fast track. There's a unique control panel with an LED/keypad for push-button diagnostics and configuration. A modular head/disk assembly so you can upgrade capacity in minutes. A reliable, brushless DC motor, a universal power supply and a dual port option.

Plus something else no one else can match: our vertical integration and 20 years of computer peripheral and offshore manufacturing expertise.

So if you're looking for someone with the drive it takes to play a supporting role in your next big release, contact Ampex. Call us tollfree at 800 621-0292. 800 821-9473 in California. We'll be happy to set up an audition.

AMPEX

Ampex Corporation · One of The Signal Companies

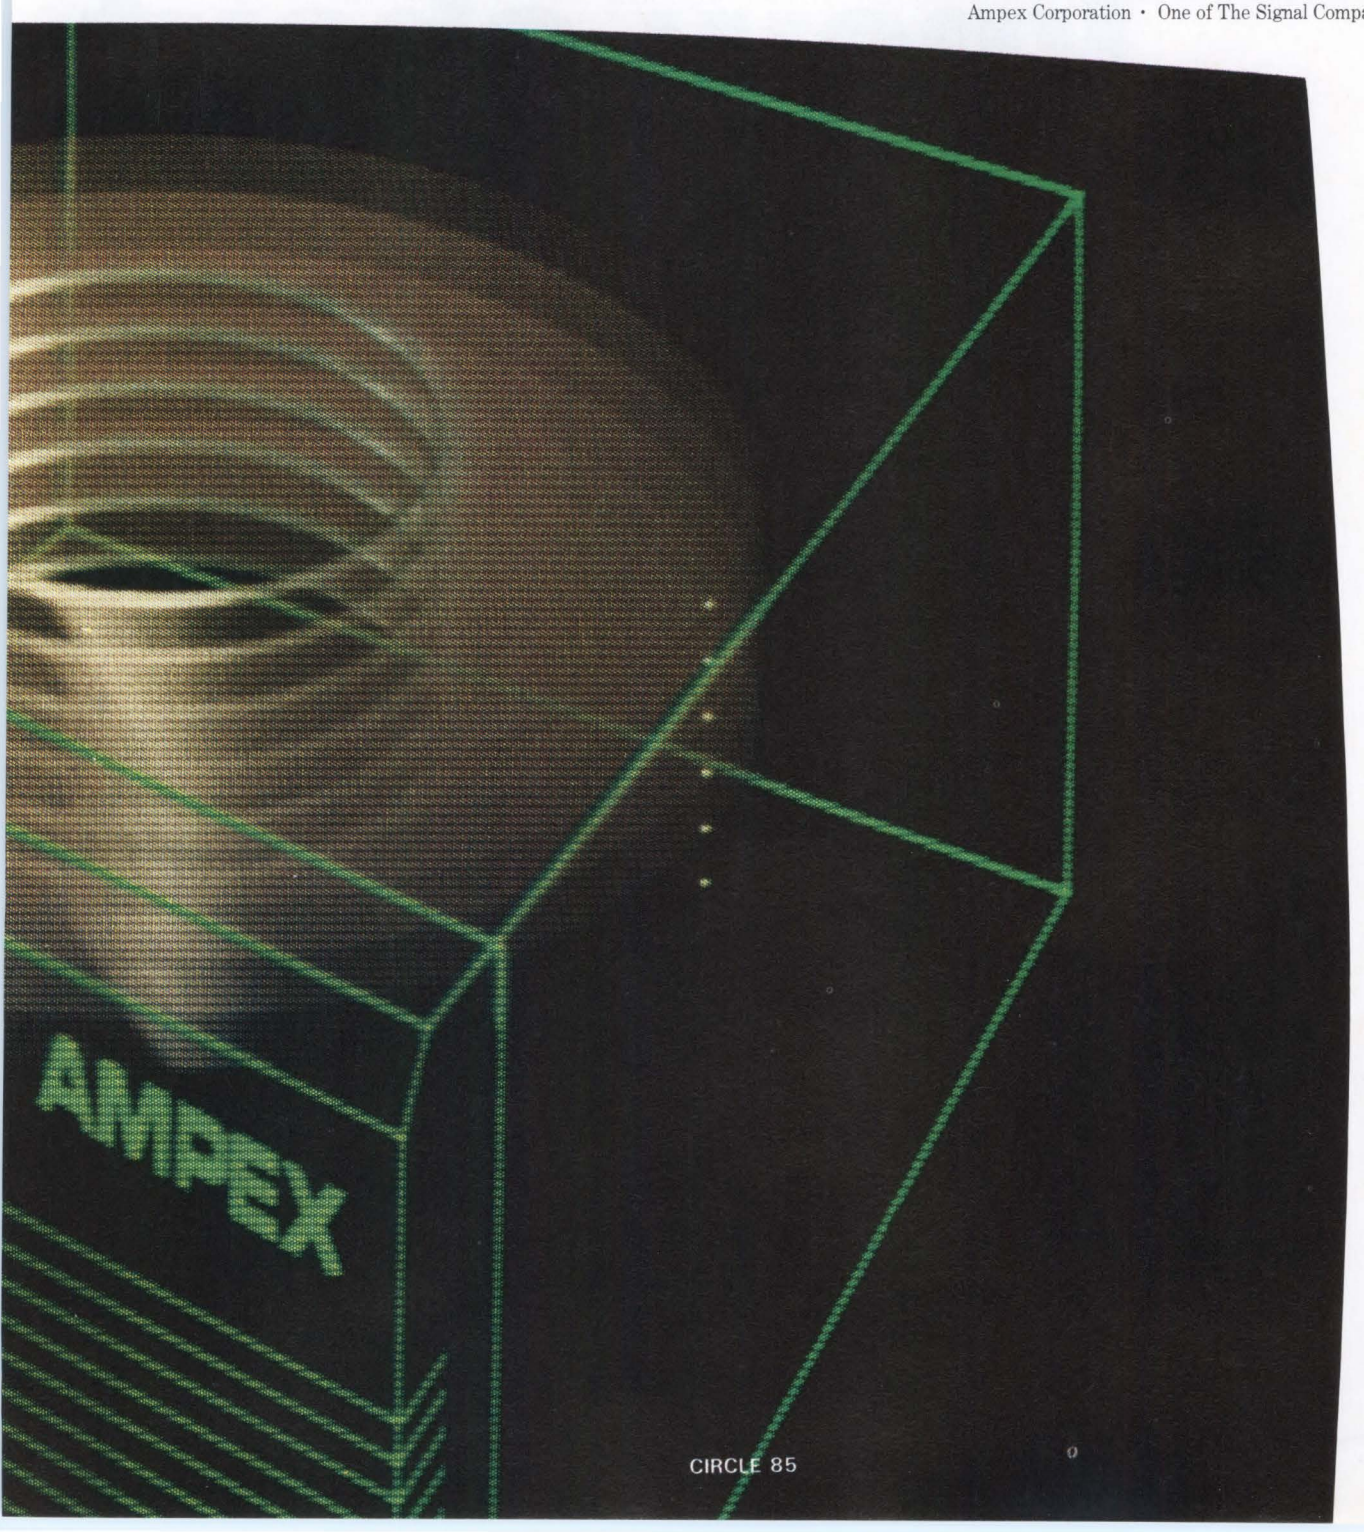

CIRCLE 85

23 things Multibus* know about

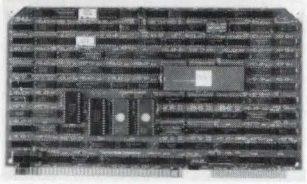

12.5 MHz 68000 CPU

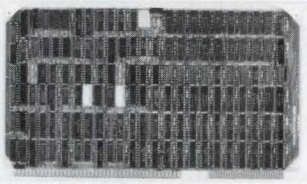

iLBX* Cache Memory Board

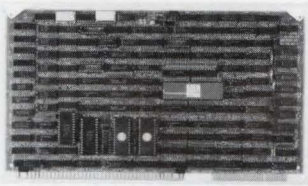

10 MHz Z8000* CPU

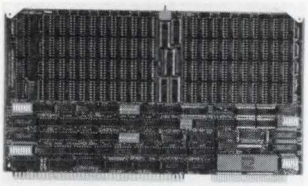

180ns 128K-2Mb EDC DRAM Board

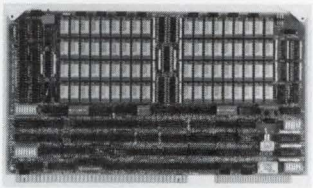

140ns 128K-512K EDC DRAM Board

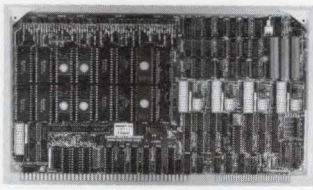

EPROM/RAM/EEPROM Board with Real Time Clock/Calendar

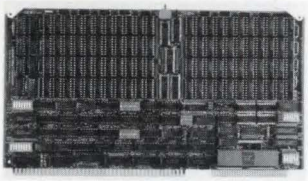

170ns 128K-2Mb Parity Only DRAM Board

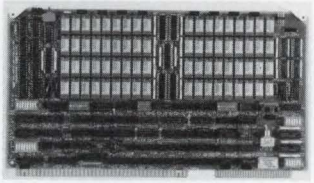

135ns 128K-512K Parity Only DRAM Board

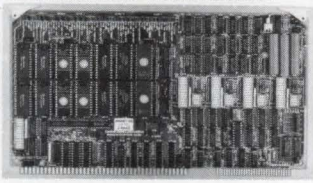

Static RAM Board with 128K CMOS

iLBX Backplane

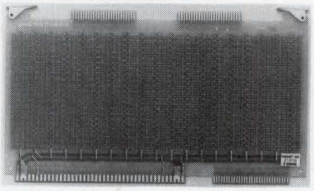

Prototyping Board

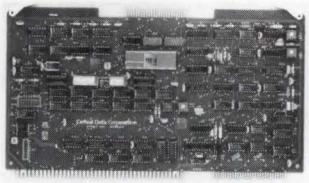

Double-Density Floppy Disk Controller

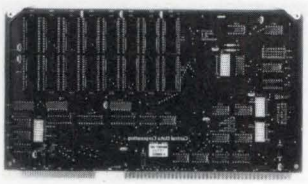

PROM Board

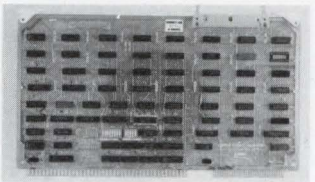

QIC-02 Cartridge Tape and Controller

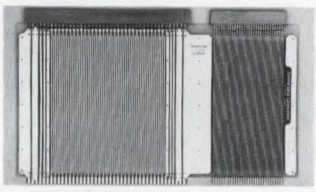

Extender Board

C2000 System Cabinet

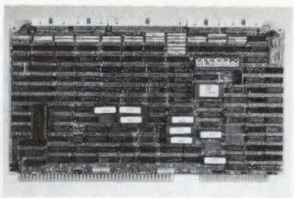

SMD Disk Controller

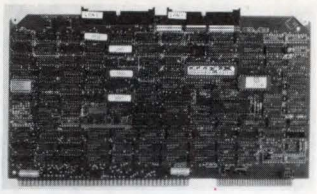

9-Track Tape Controller

users should Central Data

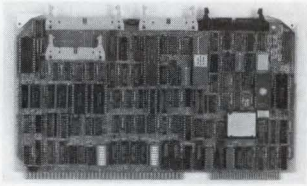

Multi-Media Controller
(Disk/Floppy/Tape)

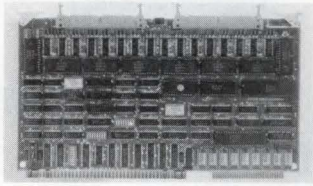

High Performance Octal Serial Interface

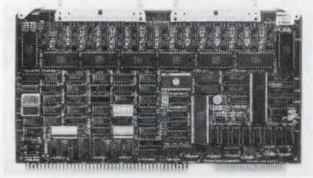

Intelligent Octal Serial Interface

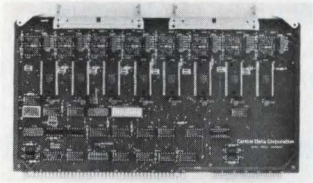

Octal Serial Interface

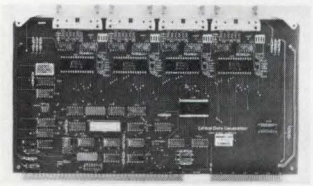

Quad Serial Interface

Whatever your Multibus* application, you'll find the quality and reliability you need in the full line from Central Data Corporation.

Central Data's engineers have been designing firsts for over eight years and our complete line now includes innovative CPU boards, I/O boards, and the industry's fastest Dynamic RAM boards. We also offer disk controllers, tape controllers, cabinet subsystems, backplanes and accessories for every need.

We've got it all, but at Central Data, that isn't enough. We're constantly researching and developing new board technologies and offering them to you. Eight new board designs are being introduced this year and many more are in the design and testing stages. And if all this isn't enough for you, call and let us know what your special needs are.

Call us at our new toll-free number:
800-482-0315 (outside Illinois)

Central Data

Central Data Corporation
1602 Newton Drive
Champaign, IL 61821-1098
(217) 359-8010
TWX 910-245-0787

Your Complete Multibus Source

*Multibus and iLBX are trademarks of Intel Corporation.
Z8000 is a trademark of Zilog Corporation.

HOW THE LEX 90 TURNS PIXELS AND BITS INTO DOLLARS AND CENTS.

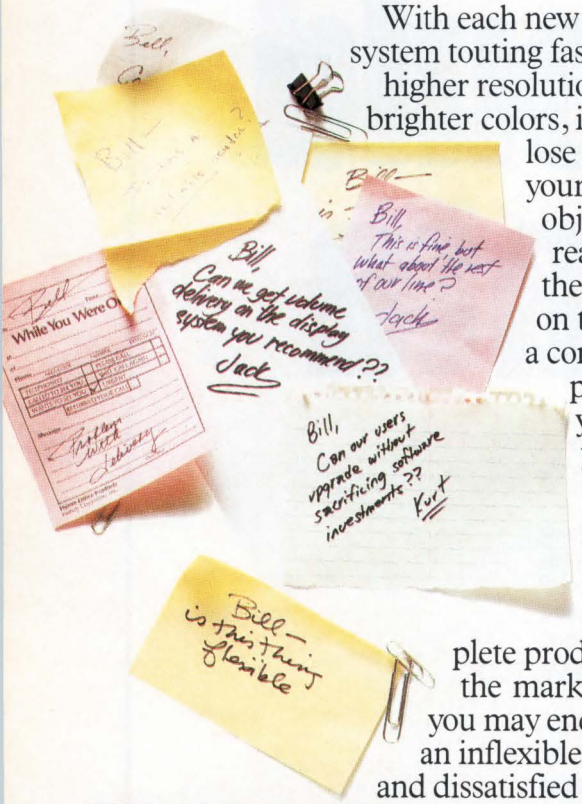

With each new display system touting faster speeds, higher resolutions and brighter colors, it's easy to lose sight of your primary objective—reaching the market on time with a competitive product. If you're late, you may lose valuable market share. Rush an incomplete product onto the market, and you may end up with an inflexible system and dissatisfied users.

With this in mind, Lexidata designed the new LEX 90 family to be as appropriate from a business standpoint as it is technologically.

TESTING 1, 2, 3.

If you're presently planning an integrated graphics system, you can count on the LEX 90 being ready. Really ready. And you can be sure our raster display systems have been thoroughly tested to ensure maximum reliability and maintainability.

Case in point: Before any design is finalized, our Advanced Manufacturing Engineering group makes certain we can manufacture it cost effectively, thus making it more affordable for you. And our Design Assurance Engineering group tests and retests to assure the system measures up to the high performance standards we set.

During manufacturing, all our systems undergo rigorous component and board ATE testing and full burn-in. And should any trouble develop in the field, on-board diagnostics quickly detect it so you can solve the problem easily.

BUY ONLY WHAT YOU NEED

The LEX 90 family also makes good business sense because you only have to buy the functionality you need. And you can do so without compromising your technological standards. For example, we offer a Dual Resolution model with software selectable 1000-line, 30 Hz or 500-line, 60 Hz color, or both simultaneously on one screen. And we also offer flicker-free High Resolution color in a 1280 x 1024 x 8, 60 Hz model.

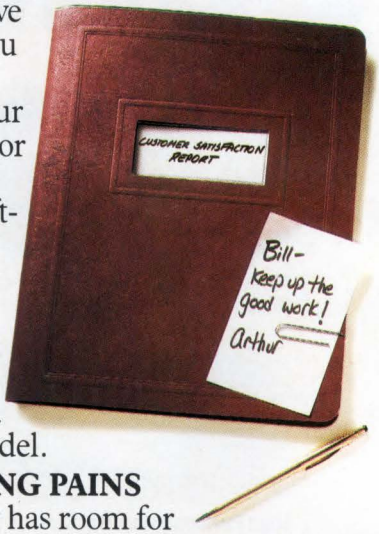

NO GROWING PAINS

The LEX 90 family has room for growth as your application becomes more complex. And this growth won't jeopardize your software investments because its architectural standard permits the system builder to work with upwardly compatible sets. Furthermore, the LEX 90 is software compatible with Lexidata's existing Series 3000 products.

The LEX 90 family. High quality, reliable and flexible. As a result, your users not only stay happy, they stay with you.

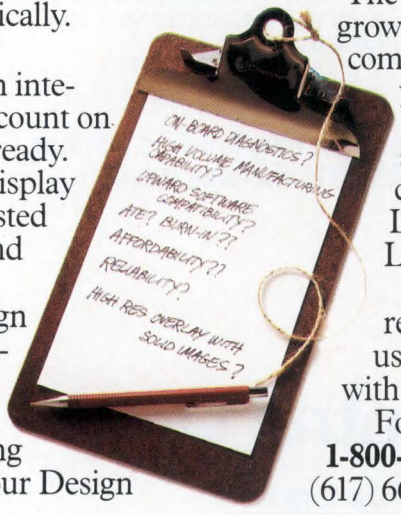

For more information, call Lexidata at **1-800-472-4747** (in Massachusetts, call (617) 663-8550). TWX 710-347-1574.

The new LEX 90 family.

LEXIDATA

The clear choice in raster graphics.

Lexidata Corporation, 755 Middlesex Turnpike, Billerica, MA 01865. In UNITED KINGDOM call: Lexidata Ltd., Hook (025672) 4334. In FRANCE call: Lexidata SARL, Rungis (1) 686 56 71. In JAPAN call: Lexidata Technical Center, Tokyo 486-0670.

CIRCLE 87

HIGH SPEED SYSTEMS LOOK TO GaAs FOR LOW POWER LSI

Digital gallium arsenide chips sporting both high speed and low power promise to quicken the start of an ultrafast supercomputer era.

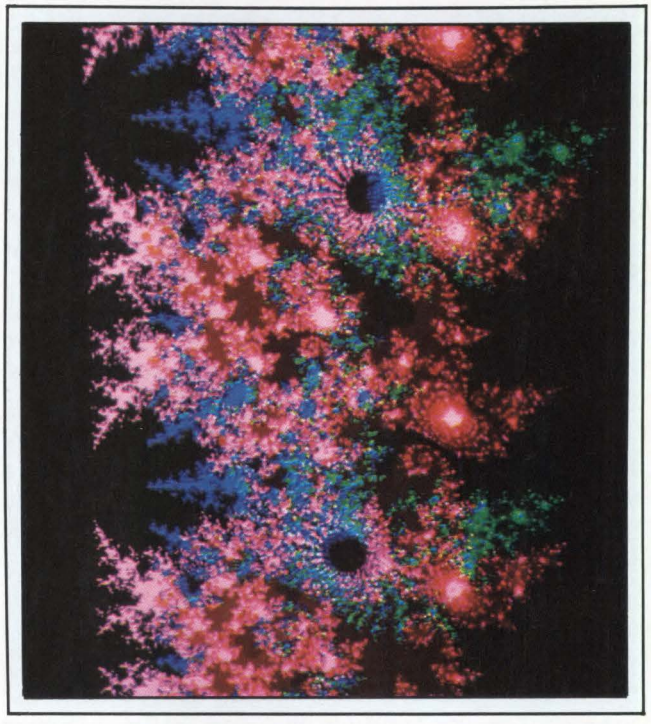

Photo courtesy of Dr John Hubbard and Homer Smith, Cornell University

by Thomas M. Reeder and Ajit G. Rode

Recent progress in gallium arsenide IC development has produced some surprising results. Once thought of as a very high speed, but low yield, power hungry technology most suitable for discrete transistor and SSI applications, GaAs ICs can now achieve both LSI complexity and low power dissipation. Laboratories in the United States, Japan, and Europe are demonstrating GaAs LSI chips that have substantially better speed and power parameters. Although current GaAs results are experimental, laboratory development devices, the amount of research going on lends credibility to recent technology forecasts for significant use of GaAs in high speed digital logic systems before the end of the decade.

Growth in GaAs chip complexity is even outpacing that of silicon ICs (Fig 1). A look at recent developments in logic delay versus power dissipation shows that GaAs research has extended device operation to much lower power levels while still

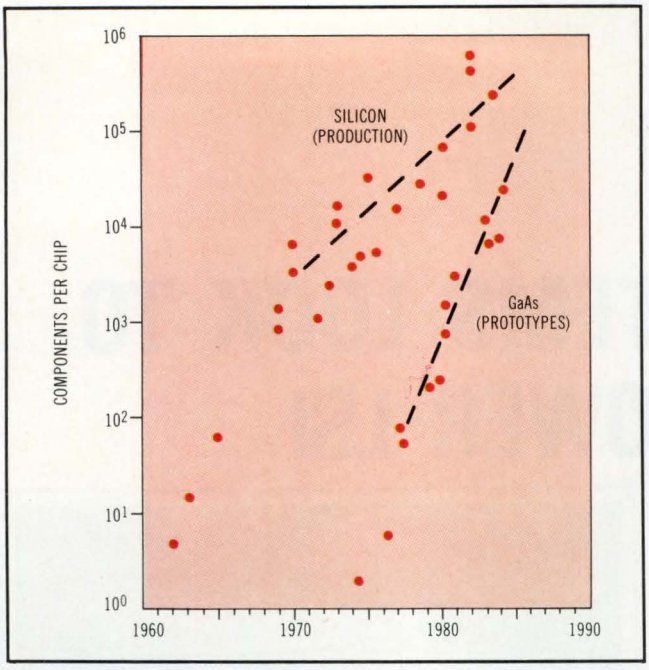

Fig 1 The circuit complexity of GaAs prototype chips compared to production silicon indicates that the density of GaAs is increasing at a faster rate than silicon.

achieving speeds that are a factor of two to five better than silicon. Recent demonstrations that confirm the delay-power advantage at LSI complexity include a 16 x 16 digital multiplier developed by Fujitsu, which has a 10.6-ns multiply time with a 952-mW power dissipation; 1- and 4-Kbit static RAMs from Fujitsu and NTT, which have access times below 3 ns and power dissipation less than 0.7 W; and a 1200-gate array from Tektronix, which achieves high performance ECL speed (400 MHz) with about one-tenth the power dissipation (0.25 mW/gate). In each of these recent demonstrations, the delay-power product was 3 to 10 times better than found on comparable, state-of-the-art silicon research devices.

Device operation

Several different transistor structures have been proposed as the basis for development of GaAs LSI ICs. The basic planar, ion implanted structure shown in Fig 2(a) was first proposed in 1977 by Rockwell International as a vehicle for LSI device development. This structure uses a simple Schottky gate electrode deposited directly on a lightly implanted (n –) GaAs surface to control transistor current flow. This device is a direct descendant of the discrete metal semiconductor field effect transistor (MESFET) developed in the early 1970s for use in terrestrial and satellite communications. The gate, which is typically a Ti/Au based multilayer metal deposition, is defined by the so-called "lift-off' process, which has high dimensional accuracy and resolution. [There are 0.5-μm gates already in production for GaAs discrete transistors, and 1.0-μm gates are common for LSI ICs.] Contact to the n + source and drain regions is made through alloyed Au/Ge based layers that are also defined by lift-off. The n-channel is only 100 to 200 nm thick (about the thinnest practical measurement for ion implantation) and the depletion region formed at the gate-semiconductor surface is usually less than 100 nm thick at zero gate bias. Thus, the transistor channel is not pinched off, and source-drain current can flow at zero gate bias for a "normally-on" transistor operation. This type of GaAs transistor is also known as depletion mode (D-mode) due to the fact that current flow through the channel is controlled by depletion of negative charge in the channel.

A straightforward extension of the planar D-mode structure is the recessed gate device [Fig 2(b)]. By carefully controlling the amount of gate recess, it is possible to reach the point where at zero gate bias, the depletion region fills or pinches off the channel. Transistors made this way are known as "normally-off" or enhancement-mode (E-mode) transistors. Using both E- and D-mode transistors on the same chip allows the formation of simpler, lower power logic gates. (The same simplification was used to improve MOS devices in the early 1970s.)

GaAs research has extended device operation to much lower power levels with speeds still higher than silicon.

There are also some structures under study that yield device performance very similar to the recessed gate MESFET. But, they utilize a different approach to achieve a channel thin enough to provide E-mode operation. One of these structures is a self-aligned approach using a high temperature gate so that the n + source-drain regions can be implanted after the gate is deposited. Fujitsu is using this method to build a number of LSI demonstration devices, including the 16 x 16-bit multiplier and 4-Kbit RAM devices mentioned earlier.

Another structure under development is the ion implanted junction FET (JFET) device [Fig 2(c)]. Here, the gate is formed by a shallow p-type implant into the previously implanted n-channel. Current flow is controlled by the depletion region formed at the p, n interface. This device is analogous to some silicon JFETs now in production. GaAs JFET transistors can be designed to be "normally-off" as fabricated, so their operation may be similar to the recessed gate MESFET in Fig 2(b).

The High Electron Mobility Transistor (HEMT) or modulation-doped FET (MODFET) structure

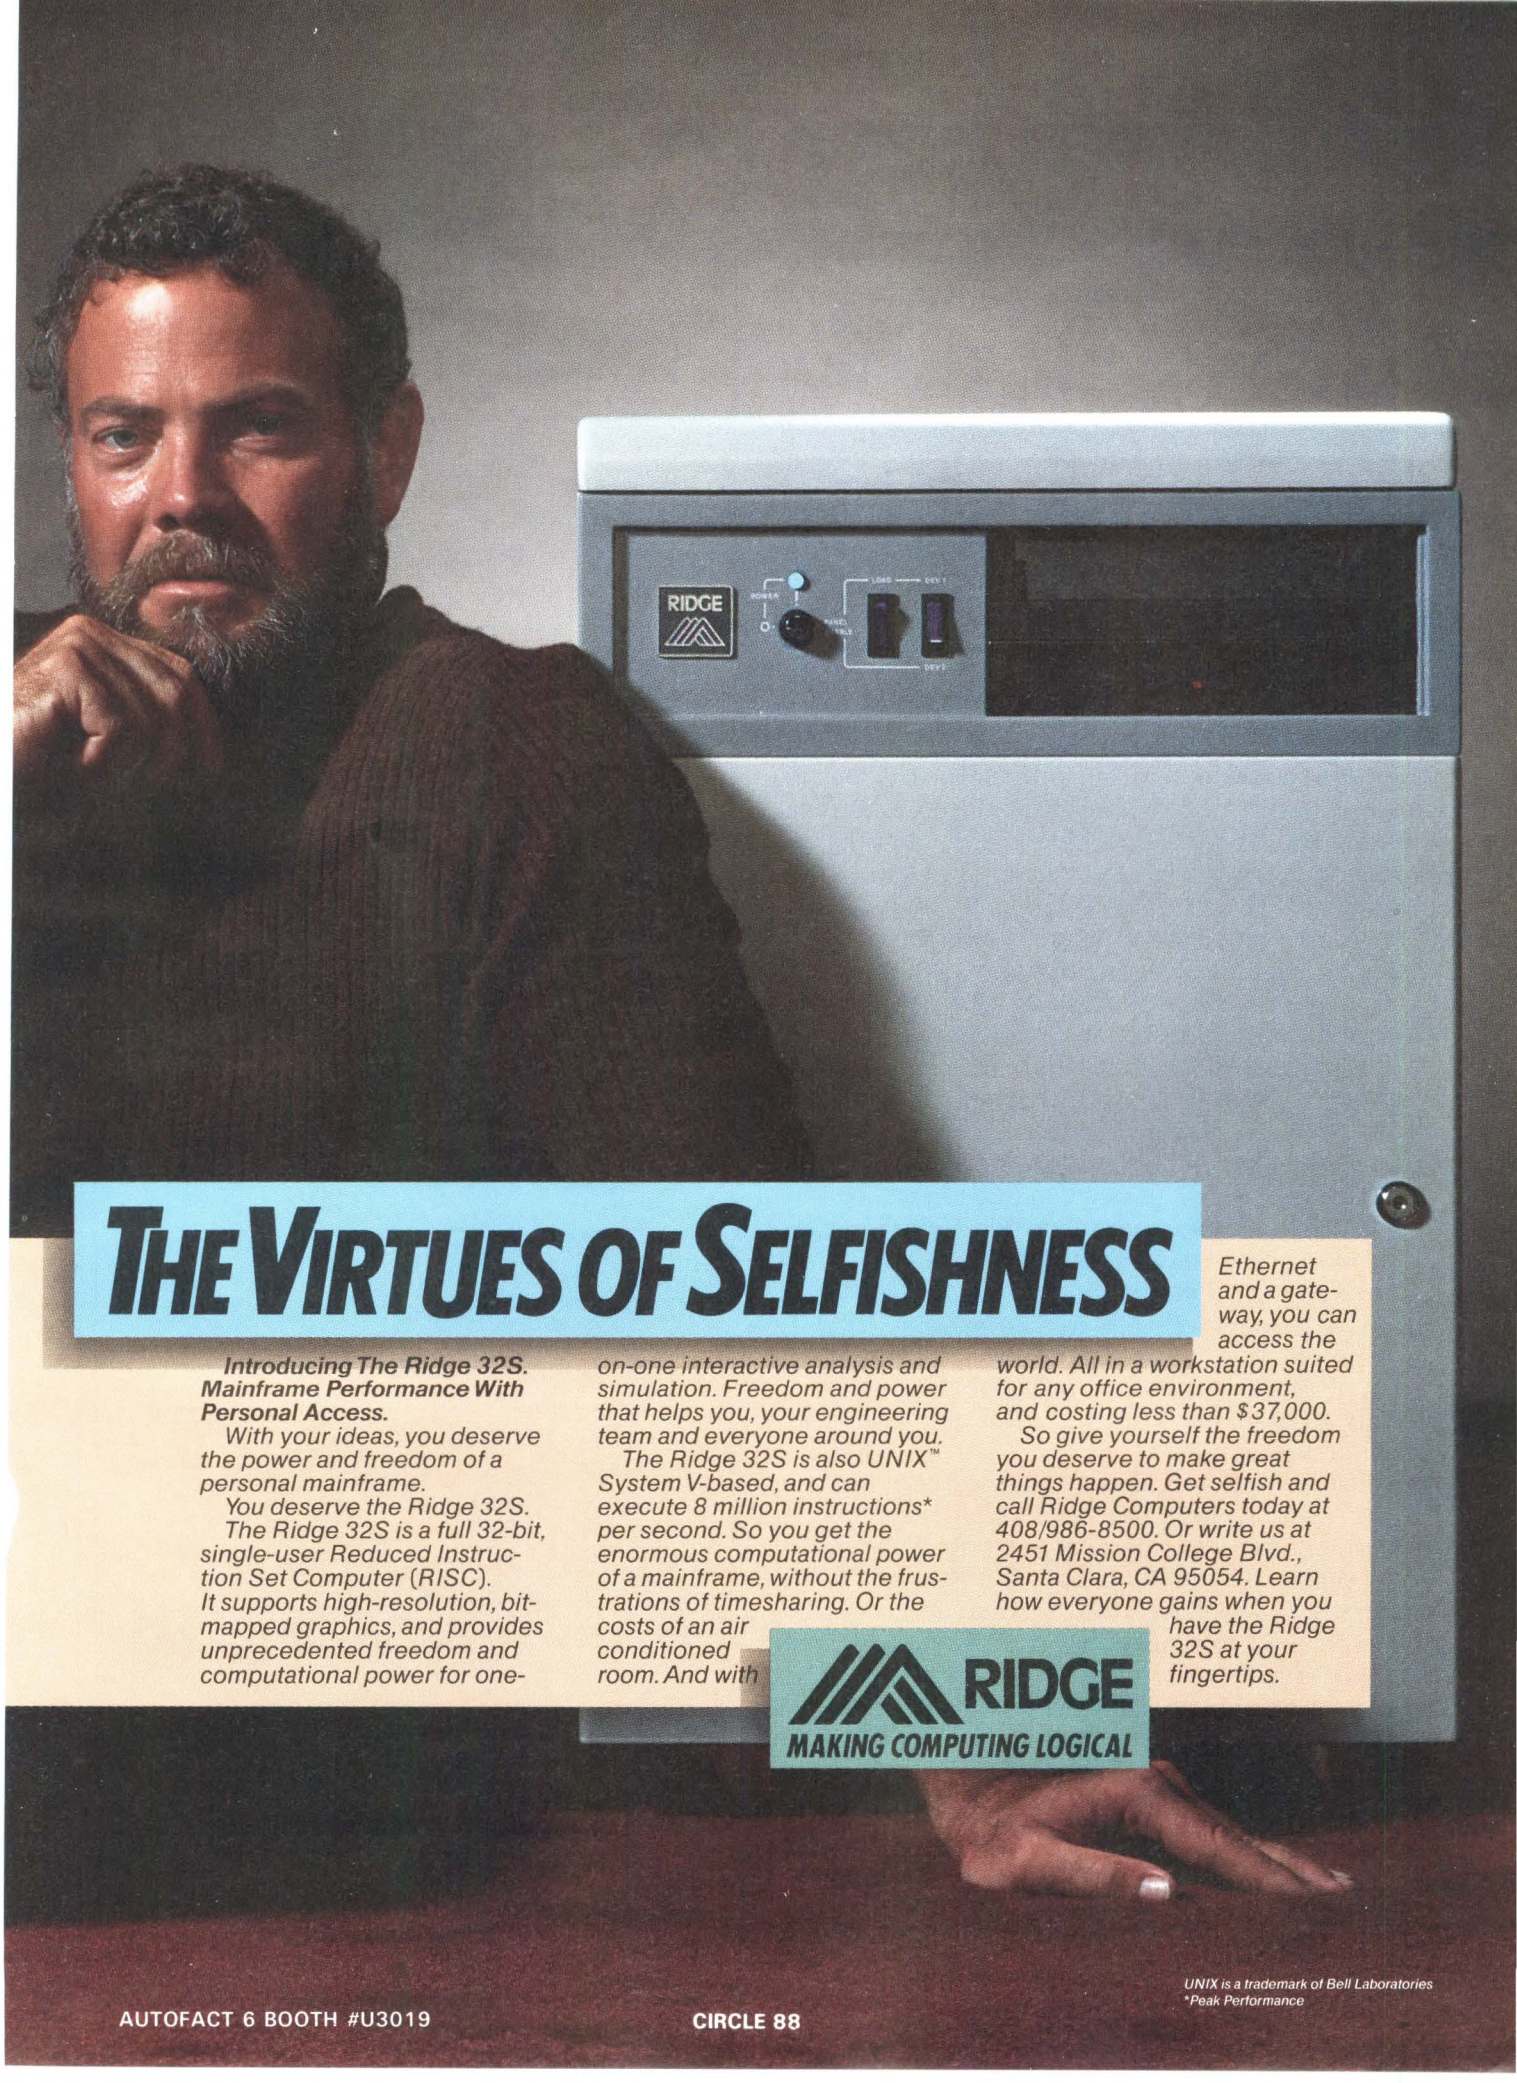

THE VIRTUES OF SELFISHNESS

Introducing The Ridge 32S. Mainframe Performance With Personal Access.

With your ideas, you deserve the power and freedom of a personal mainframe.

You deserve the Ridge 32S.

The Ridge 32S is a full 32-bit, single-user Reduced Instruction Set Computer (RISC). It supports high-resolution, bit-mapped graphics, and provides unprecedented freedom and computational power for one-on-one interactive analysis and simulation. Freedom and power that helps you, your engineering team and everyone around you.

The Ridge 32S is also UNIX™ System V-based, and can execute 8 million instructions per second. So you get the enormous computational power of a mainframe, without the frustrations of timesharing. Or the costs of an air conditioned room. And with Ethernet and a gateway, you can access the world. All in a workstation suited for any office environment, and costing less than $37,000.*

So give yourself the freedom you deserve to make great things happen. Get selfish and call Ridge Computers today at 408/986-8500. Or write us at 2451 Mission College Blvd., Santa Clara, CA 95054. Learn how everyone gains when you have the Ridge 32S at your fingertips.

RIDGE
MAKING COMPUTING LOGICAL

UNIX is a trademark of Bell Laboratories
**Peak Performance*

AUTOFACT 6 BOOTH #U3019

CIRCLE 88

[Fig 2(d)] is a relatively new entry to GaAs IC development. Although it is formed from a more complex multilayer semiconductor structure fabricated by new methods such as molecular beam epitaxy (MBE) or metallo-organic chemical vapor deposition (MOCVD), it could ultimately become the dominant transistor structure for GaAs LSI. However, because MBE and MOCVD are still in fairly early stages of development, they are not well suited to volume manufacturing, especially MBE growth layer techniques. The use of dissimilar GaAs and AlGaAs layers results in a heterojunction structure that provides higher mobility and device speed, especially at liquid nitrogen temperatures, than is seen in the simple MESFET or JFET structures. While considered to be still in fairly early research, a few LSI devices such as a 1-Kbit RAM have been successfully demonstrated.

Another new GaAs transistor approach is the bipolar structure [Fig 2(e)]. Its construction is analogous to the silicon bipolar transistor in that current flow is vertical through the device layers rather than horizontal or parallel to the surface, as in the previously mentioned FET structures. This device is also made by MBE or MOCVD and, like the HEMT device, uses a heterojunction structure. Researchers feel that the SSI bipolar structure will be very useful for building high speed ICs because of its similarity to fast ECL devices in production today.

The choice of which structure to use will be made on the basis of its performance and ease of use in digital logic design. For example, the D-mode MESFET (DMESFET) was first implemented in the buffered FET logic (BFL) circuit configuration illustrated by the NOR gate circuit [Fig 3(a)]. The diode chain is necessary to shift the output buffer logic levels so that input and output voltage levels cover the same range. This approach provides relatively good tolerance to transistor fabrication variations, but dissipates considerable power per logic gate (typically 10 to 40 mW) due to buffer and diode chain requirements.

The Schottky diode FET logic (SDFL) approach [Fig 3(b)] has been used in several laboratories to reduce power dissipation by inserting the level shifting diodes on the input side of a NOR gate where current flow is lower. Power dissipation for SDFL is typically 1 to 10 mW, but gate fanout and circuit drive capability are reduced, leading to slower and less reproducible circuits. The simplest GaAs logic approach studied so far is the enhancement/depletion (E/D) circuit [Fig 3(c)]. Here, a normally-off E-mode driver transistor is used with a D-mode pull-up load to form a fast but relatively low power inverter. A two input NOR gate or a two input NAND gate can be made with only three transistors. Thus, the E/D approach is relatively simpler, and leads to higher

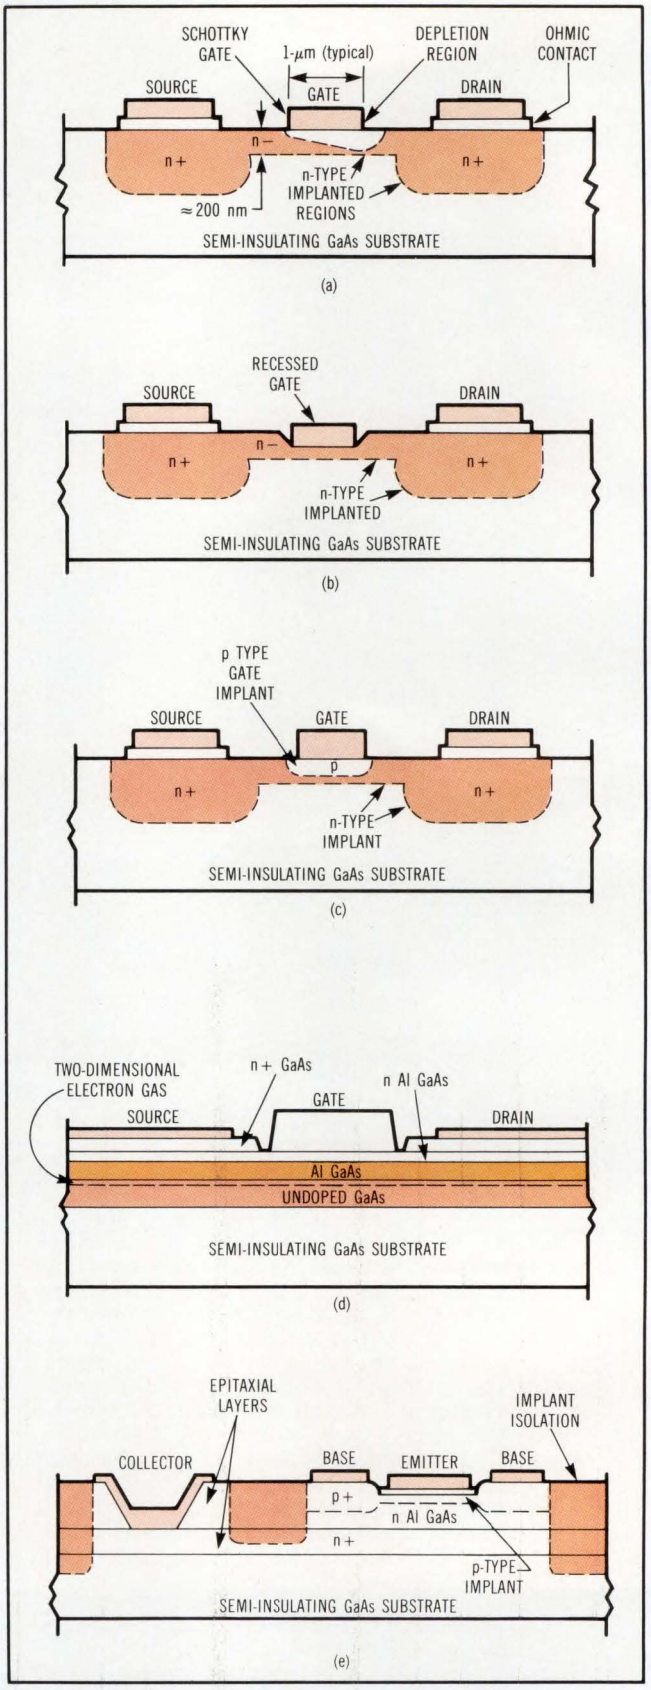

Fig 2 Shown are schematic views of transistors for several GaAs processes, including planar MESFET (a), recessed gate MESFET (b), JFET (c), HEMT (d), and bipolar (e). These schematic views indicate the different structural approaches currently in use or under development.

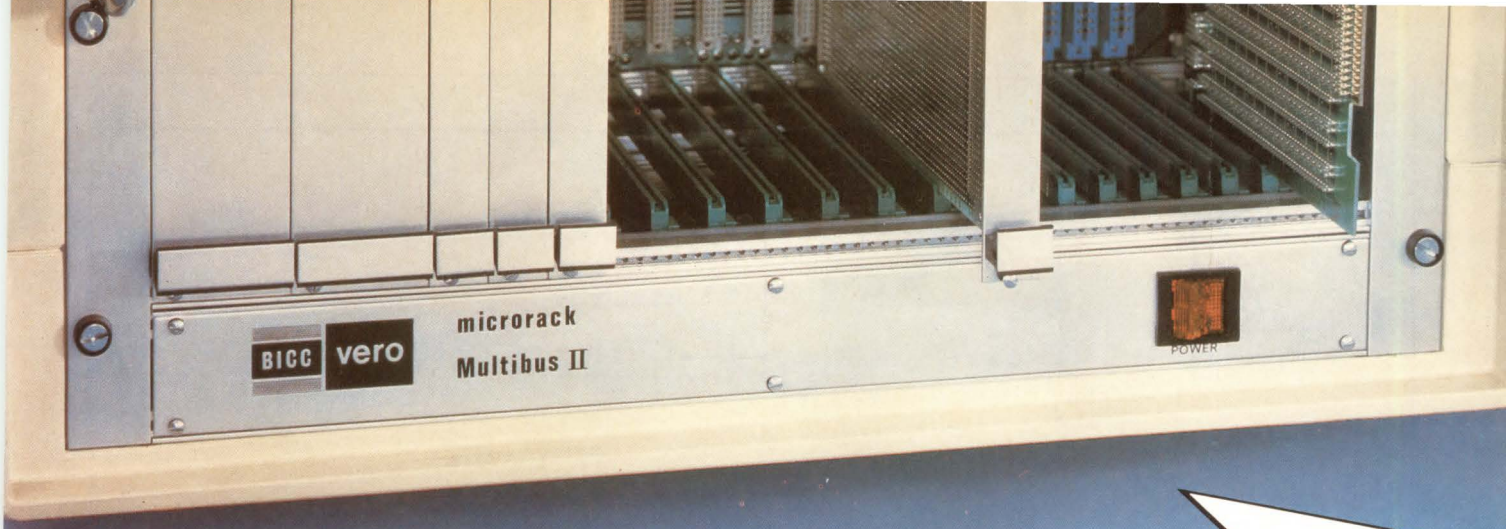

Complete hardware support for the Multibus II user

The hardware support for a bus system based on the Eurocard form factor needs very specialised experience.

BICC Vero have that experience and our worldwide reputation is unquestionable.

Users of leading bus systems have already benefited from our expertise with such hardware products as backplanes, extender boards, prototyping boards, card cages, connectors and power supplies – and that's just to start with.

Now we have combined all of these products into a total capability and have produced a range of Microrack Systems.

For the Multibus II user BICC Vero now offers from stock, a full range of P.S.B. and L.B.X. backplanes and extender boards, as well as the standard products that you will require to build your microprocessor system.

Telephone or write for further information on Multibus II support hardware.

Whatever your system architecture, BICC Vero have the experience and the hardware to support you.

Multibus II Support Hardware
- * P.S.B. and L.B.X. Backplanes
- * Extender Boards
- * Prototyping Boards
- * Card Cage Systems
- * Power Supplies
- * Connectors

BICC vero

*Multibus® II is the Reg. Trade Mark of Intel Corporation
**Euro-Q Bus® is the Reg. Trade Mark of Digital Equipment Corporation

BICC-VERO ELECTRONICS INC.
171 Bridge Road Hauppauge New York 11788
Tel: (516) 234-0400 TWX: 510-227-8890
4001 Leaverton Court Anaheim California 92807
Tel: (714) 630-2030 Telex: 277732

MULTIBUS® II SUPPORT HARDWARE

CIRCLE 89

Leaders in Microprocessor Backplanes

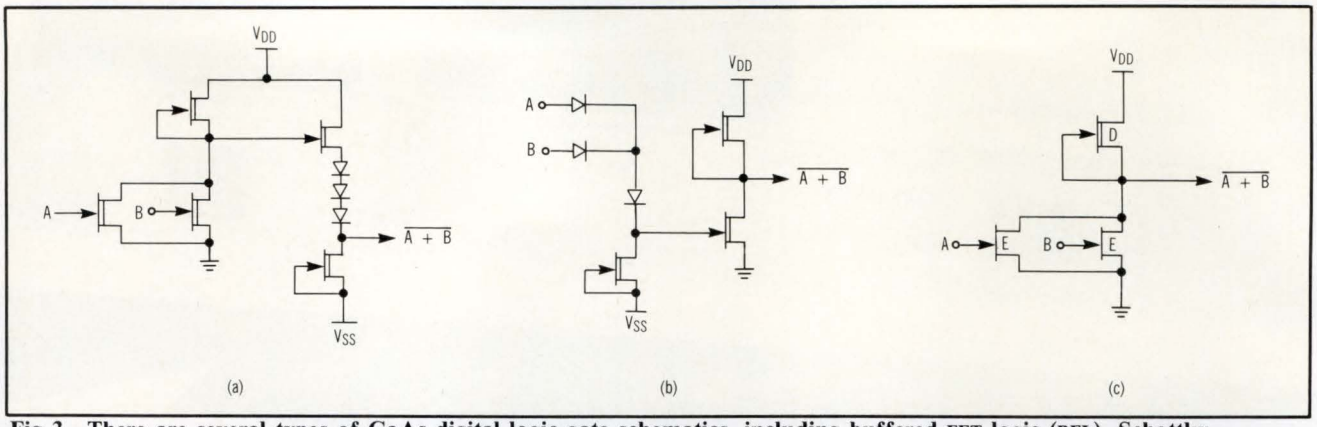

Fig 3 There are several types of GaAs digital logic gate schematics, including buffered FET logic (BFL), Schottky diode FET logic (SDFL), and enhancement-depletion mode (E/D). Although BFL (a) offers good tolerance to transistor fabrication variations and SDFL (b) offers reduced power dissipation, the E/D circuit (c) offers the simplest logic approach with higher gate density and lower circuit parasitics.

gate density and lower circuit parasitics as well. Typical power dissipation for E/D logic gates is in the 0.1- to 1-mW range.

Of the device structures discussed, the MESFET is by far the most mature and reproducible. GaAs foundry and standard component products now being introduced for commercial manufacturing all use this simple GaAs structure. Even so, the D-mode structure is better suited for monolithic microwave IC (MMIC) and SSI/MSI extremely fast logic applications because of its circuit and power dissipation requirements.

On the other hand, much progress has been made recently in the development of high yield E-mode processes using the recessed and self-aligned gate and JFET approaches. Because they offer simpler logic gate structures and relatively high yield and process uniformity, E-mode processes are under development for LSI applications such as gate arrays, RAMs, and standard logic where ECL speed and lower power dissipation are required.

LSI gate array development

The continuing development of military and commercial digital systems with clock rates above 10 MHz is creating a need and an opportunity for LSI gate array devices that meet system needs for speed, power, and radiation hardening in a format that is suitable for quick turnaround system development. Recent research publications on three different semiconductor technologies (ECL, CMOS, and GaAs) predict that arrays built in each of these technologies are attractive for use in new system development. However, the optimum choice of technology, design configuration, and package implementation is not simple because of varying system requirements, and because the emerging characteristics of LSI gate arrays cover a wide range and change rapidly.

Comparison of Several High Speed Gate Arrays

	GaAs	ECL	CMOS
Critical dimensions (μm)	1.0	2.0	2.0
Chip size (mm^2)	4.6 x 4.3	6.6 x 6.8	10 x 10
No. of equivalent gates	1200	2500	20-K
Time delay (ps/gate) (FI = FO = 3, 3-mm interconnect)	375	510	1500
Power dissipation (μW/gate) (50-MHz clock)	190	1200	310
Speed-power (fJ)	71	610	465

For example, the characteristics of several engineering prototype ECL, CMOS, and GaAs arrays demonstrated during the past two years are listed in the Table. The GaAs array uses a recessed gate E/D process developed at Tektronix[1]. The ECL and CMOS arrays use state-of-the-art silicon IC processes.[2, 3] Although the three arrays listed are sufficiently different in size to make quantitative comparison difficult, it is clear that the GaAs array has better speed (lower gate delay) than either of the silicon arrays. If the three were evaluated at a 50-MHz clock rate, the GaAs array would have dramatically less power dissipation than the ECL part and would even beat the CMOS array. (In fact, the prototype CMOS array could not operate above 10 MHz; the quoted power dissipation is a linear extrapolation.) Note also that the GaAs array beats the CMOS and ECL arrays in delay-power product by factors of about six and eight respectively.

The GaAs array was described by Tektronix at the GaAs IC Symposium held last fall in Phoenix, Ariz. Fig 4 shows a photomicrograph of this 1224-gate E/D mode array, which was designed as a vehicle to evaluate LSI GaAs device yield. This chip has 66 I/Os and has been personalized in three basic test circuits thus far: a 71-gate inverter chain, a 140-gate pseudo-random sequence generator, and a 410-gate, 8-bit serial multiplier. Wafer probe functional yield

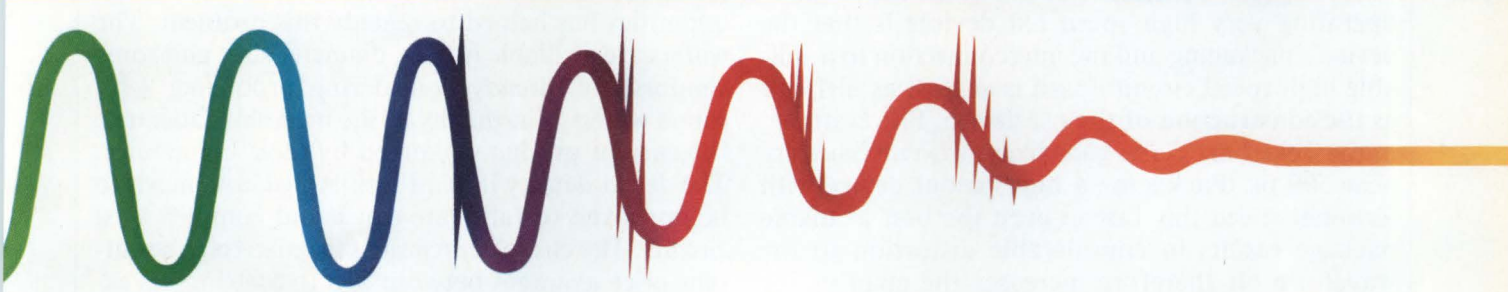

If your power is critical, finding your power problems is even more critical.

Today's computer systems, delicate instruments, and communications systems demand clean, pure power to perform accurately.

Unfortunately, that kind of power isn't always there.

Today's blackouts, brownouts, and power surges can quickly destroy or alter precious data and delicate circuits.

The Liebert 3600A Power Line Disturbance Monitor® (PLDM®) won't stop these problems, but it will help pinpoint them.

It detects disturbances that exceed sensitive thresholds and classifies them by type, magnitude, and duration. The printout tells you when the problem occurred so you can compare the information with load equipment errors, problems, and failures.

Easy to Use

The fully programmable Liebert 3600A is extremely easy to set up and use.

Instructions are located on the inside of the unit's lid and on the operating face plate. Even if you use the 3600A only once a month or less, you won't have to wade through an entire manual to make certain you set it up correctly.

Simply key in your threshold parameters and clock settings. An easy-to-read LED display will verify all of the settings. And the 3600A's self-test diagnostics ensure everything is operating properly.

When your tests are completed, simply close the lid and take the 3600A to the next test site . . . across the building, across town, or across the country.

UL-Listed

The Liebert 3600A is UL-Listed so you're assured that it meets the same rigid requirements as the equipment and systems it is designed to monitor.

Whether you specify, install, test, or use sensitive electronic systems, make certain you can count on the power you use.

Find out more about the Liebert 3600A PLDM. Call or write us today.

Liebert

Programmed Power Division

995 Benicia Ave., Sunnyvale, CA 94086
(800) 538-1770. In CA (408) 245-8900

CIRCLE 90

for these three circuits was 83, 65, and 35 percent, respectively. The pseudo-random sequence generator could be clocked at 400 MHz, and its power dissipation was only 194 μW per gate or 27 mW for the circuit. Such speed-power performance is not possible for either ECL or CMOS devices.

As GaAs IC technology makes the move from research to manufacturing, it should achieve relatively high circuit yields.

An important point to remember in designing and operating very high speed LSI devices is that the device's packaging and the interconnection to a suitable high speed circuit board is at least as difficult as the construction of the LSI device. For example, in the Tektronix GaAs gate array, I/O rise times are near 500 ps. Packaging a high pinout device with terminal speed this fast in even the best available package results in considerable distortion to the waveform. It therefore increases the chances for logic error. To allow precision device characterization, Tektronix has built a 50-Ω wafer probe system that scales the size of 50-Ω probe transmission lines down to the size of transistor bond pads.

Characterization of the 400-MHz sequence generator chip was done using this probe system. It is difficult to detect the pseudo-random sequence at 400 MHz when the same device is mounted in a commercially available package. Consequently, some developers of these very high speed ICs believe that introduction of their devices in systems will be delayed until high pinout packages with low cross talk, controlled impedance I/O contacts are available.

Manufacturing GaAs LSI circuits

The starting material for GaAs IC processing is semi-insulating GaAs with resistivity of 100 million Ω-cm. The quality and availability of substrates have been a problem in the past. However, the availability of undoped Liquid Encapsulation Czochralski (LEC) substrates from a variety of vendors in large quantities has helped to remedy this problem. The wafers are available in 3-in. diameter size, and some vendors are already considering producing 4-in. round wafers. The quality of the material is adequate for current products planned by most companies. The defect density in semi-insulating GaAs needs to be improved to fabricate dense and complex VLSI circuits. The cost of purchasing the material (the current price averages between \$30 to \$40/in.^{2}) is expected to drop as GaAs component markets develop.

Since GaAs ICs are based on MESFET technology, they are relatively simpler to fabricate than are silicon CMOS or bipolar processes. Variations of the processing technology are practiced by different companies, but most processes for fabricating LSI circuits require eight to nine mask layers. The GaAs

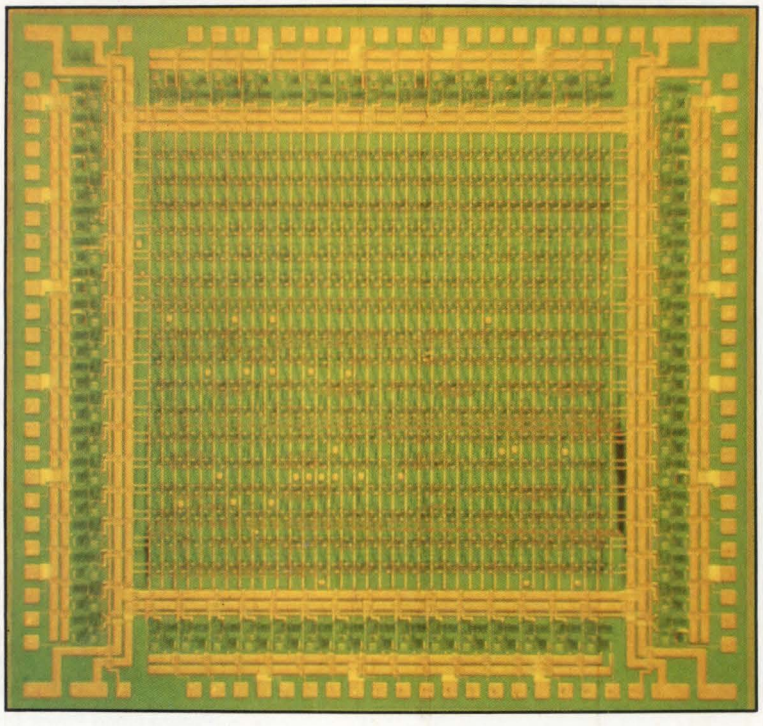

Fig 4 The Tektronix 1224-gate GaAs E/D gate array, shown in the microphotograph, has 6978 FETs on a 4.6- x 4.3-mm chip. It was designed to help evaluate LSI GaAs device yield.

Only FARADAY gives you packaging flexibility with integrated IBM® PC BUS single board computers and VLSI circuits

Anyone can be compatible.

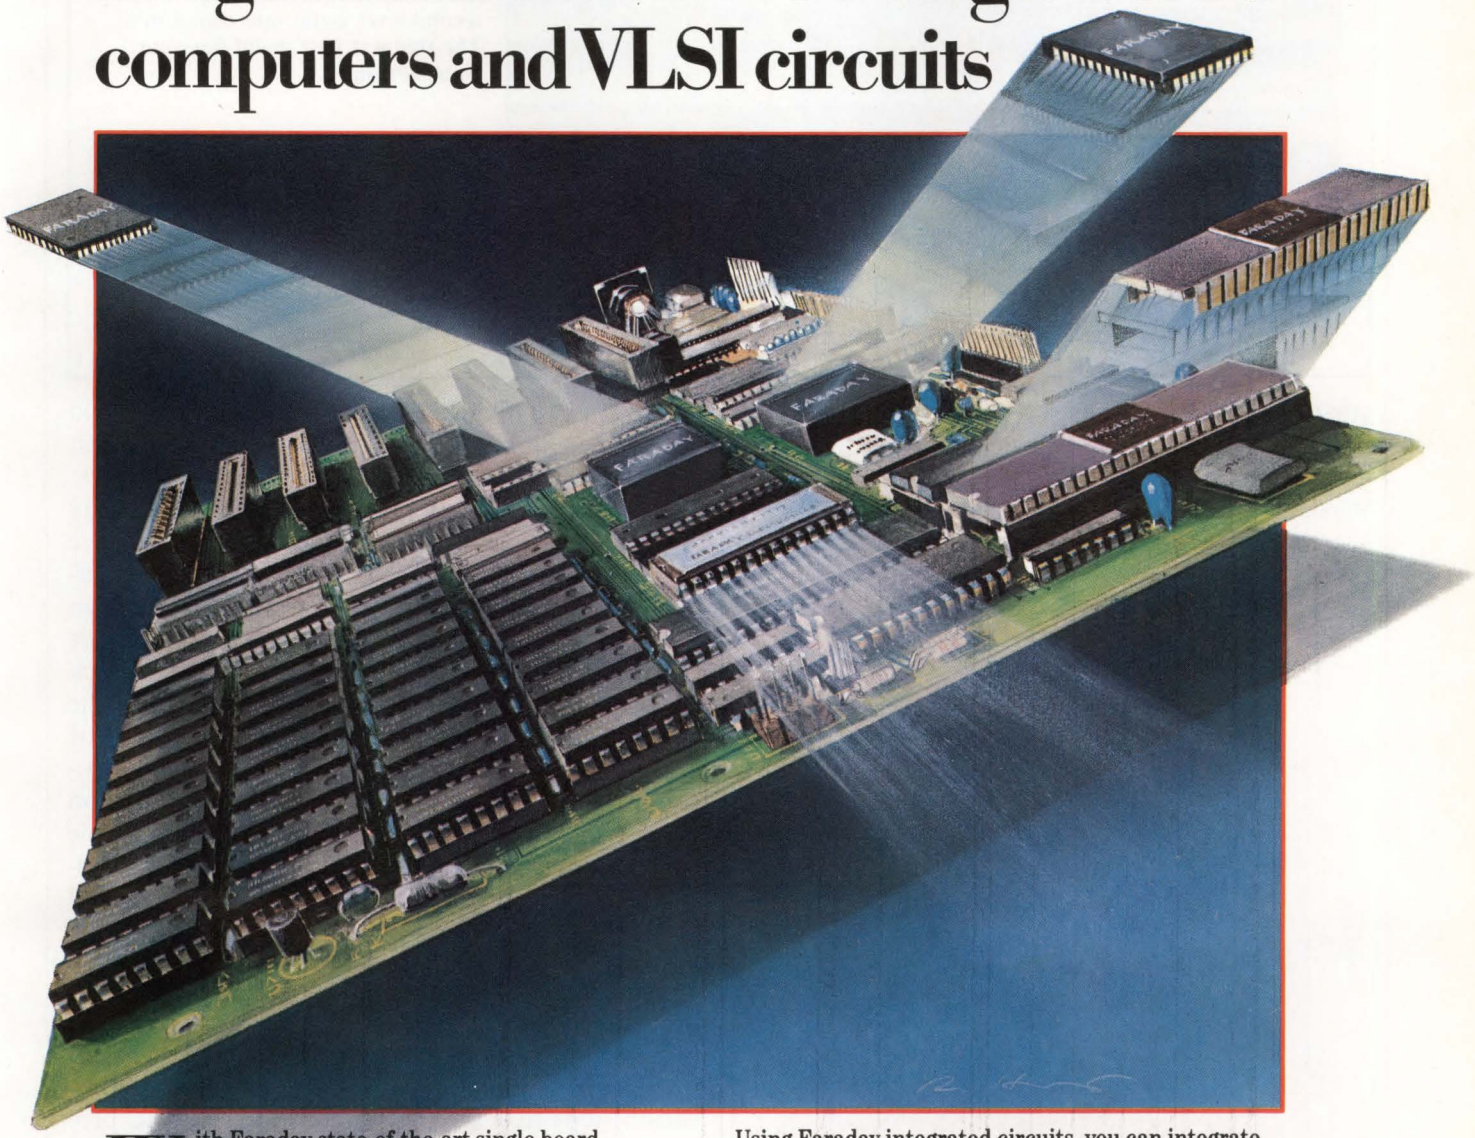

With Faraday state-of-the-art single board computers and CMOS VLSI circuits, you can implement PC-DOS compatibility in various system configurations. At lower cost. And with a shorter time to market.

Faraday single board computers are available with various combinations of integrated disk, video display and CPU controller functions. These and other functions are also available on Faraday adapter cards.

The single board computers include the Faraday BIOS in EPROM and support Faraday PC-DOS (enhanced MS™-DOS). Faraday provides MS-DOS and GW™-BASIC under license from Microsoft.

Using Faraday integrated circuits, you can integrate disk, monochrome display and CPU controller functions. These devices replace more than 75 integrated circuits.

To take advantage of the powerful cost, packaging and time advantages you'll get with this and future levels of integration, simply contact Chuck Devita, Vice President, Sales and Marketing, Faraday Electronics, 743 Pastoria Avenue, Sunnyvale, California 94086, (408) 749-1900.

MS and GW are trademarks of Microsoft Corporation.
IBM is a registered trademark of International Business Machines.

FARADAY

The OEM PC Compatible Company.

CIRCLE 91

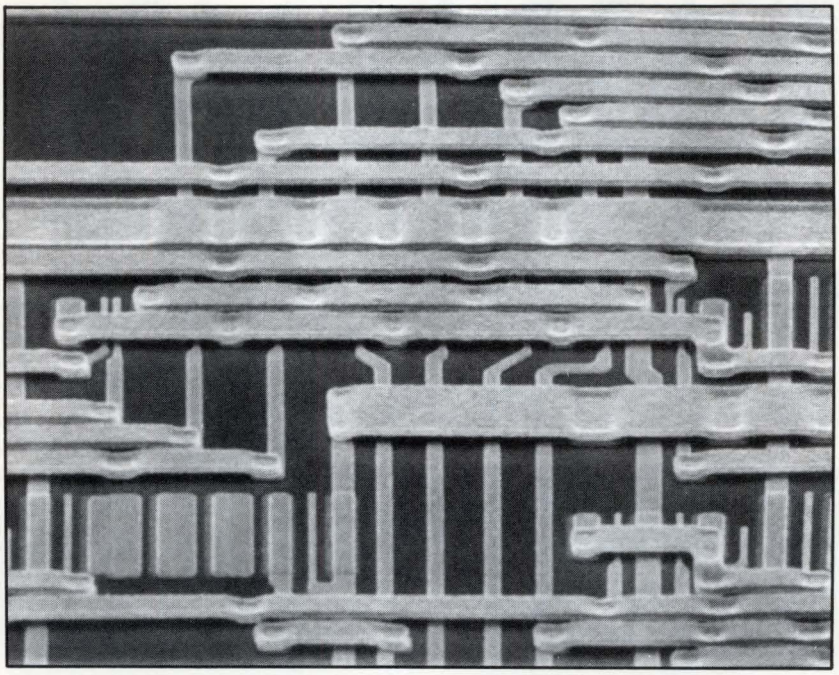

Fig 5 A microphotograph of one portion of the 1224-gate array shows the use of "airbridge" second-level metal interconnections. The thinnest first metal lines are 1-μm FET gates.

MESFET device is a bulk conduction device (as opposed to surface effect device) and is thus largely immune to surface defects and charges. Unlike a MOS device, the MESFET does not suffer from problems related to thin oxides (such as oxide electric field breakdown and threshold shifts due to hot electron effects and charge trapping).

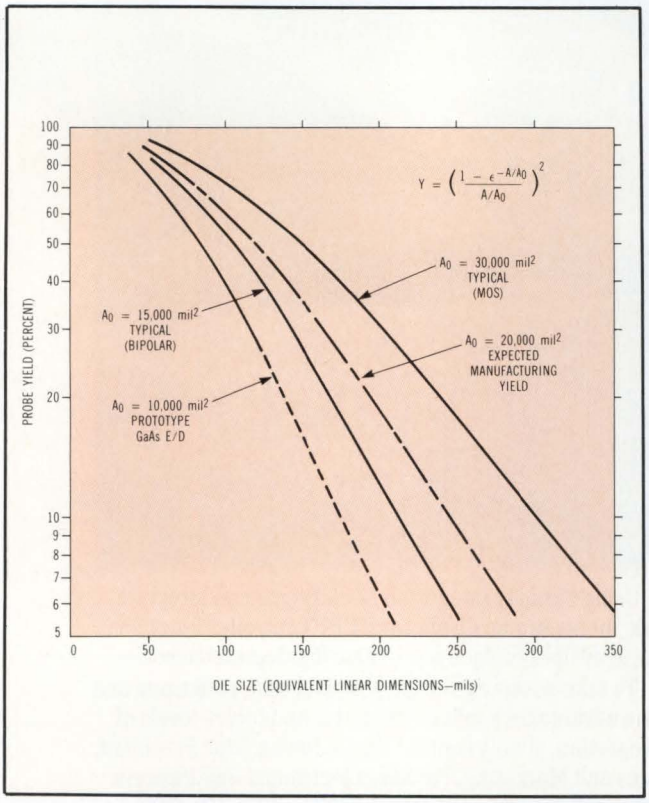

Fig 6 Yields for the GaAs LSI process, under development at Tektronix, are compared to typical silicon MOS and bipolar IC yields. Probe yield is compared to die size for each process.

"Airbridge" metal interconnection technology (Fig 5) is being used in some laboratories to achieve higher circuit performance (lower parasitic capacitance) and higher process yield than is obtained in the dielectric isolated metal interconnection methods most often used in silicon ICs. This is due to lack of any crossovers of second-layer metal over the edges of the first-level metal, as well as the lack of pinholes that reduce yield in dielectric isolated processes. As GaAs IC technology makes the transition from research to manufacturing, it should be able to achieve relatively high circuit yields due to the simplicity of process, cleaner fabrication facilities, and higher wafer volumes. Fig 6 shows the circuit yield versus die area achieved in the GaAs E/D process under development at Tektronix, and compares these results with yields seen in typical silicon MOS and bipolar processes.

GaAs IC processing requires equipment and fabrication steps that are the same as in silicon IC processing. It is essential, indeed mandatory, to use the same processing equipment for GaAs fabrication line as in silicon processing due to the equipment's availability and lower cost versus developing and building custom processing equipment. Silicon IC processing equipment also has a long history and large data base behind it that makes it logical to select a known piece of equipment.

Availability of 3-in. round GaAs wafers makes it possible to adapt to the automated wafer handling

You're challenging the limits of 5¼" drives. And so are we.

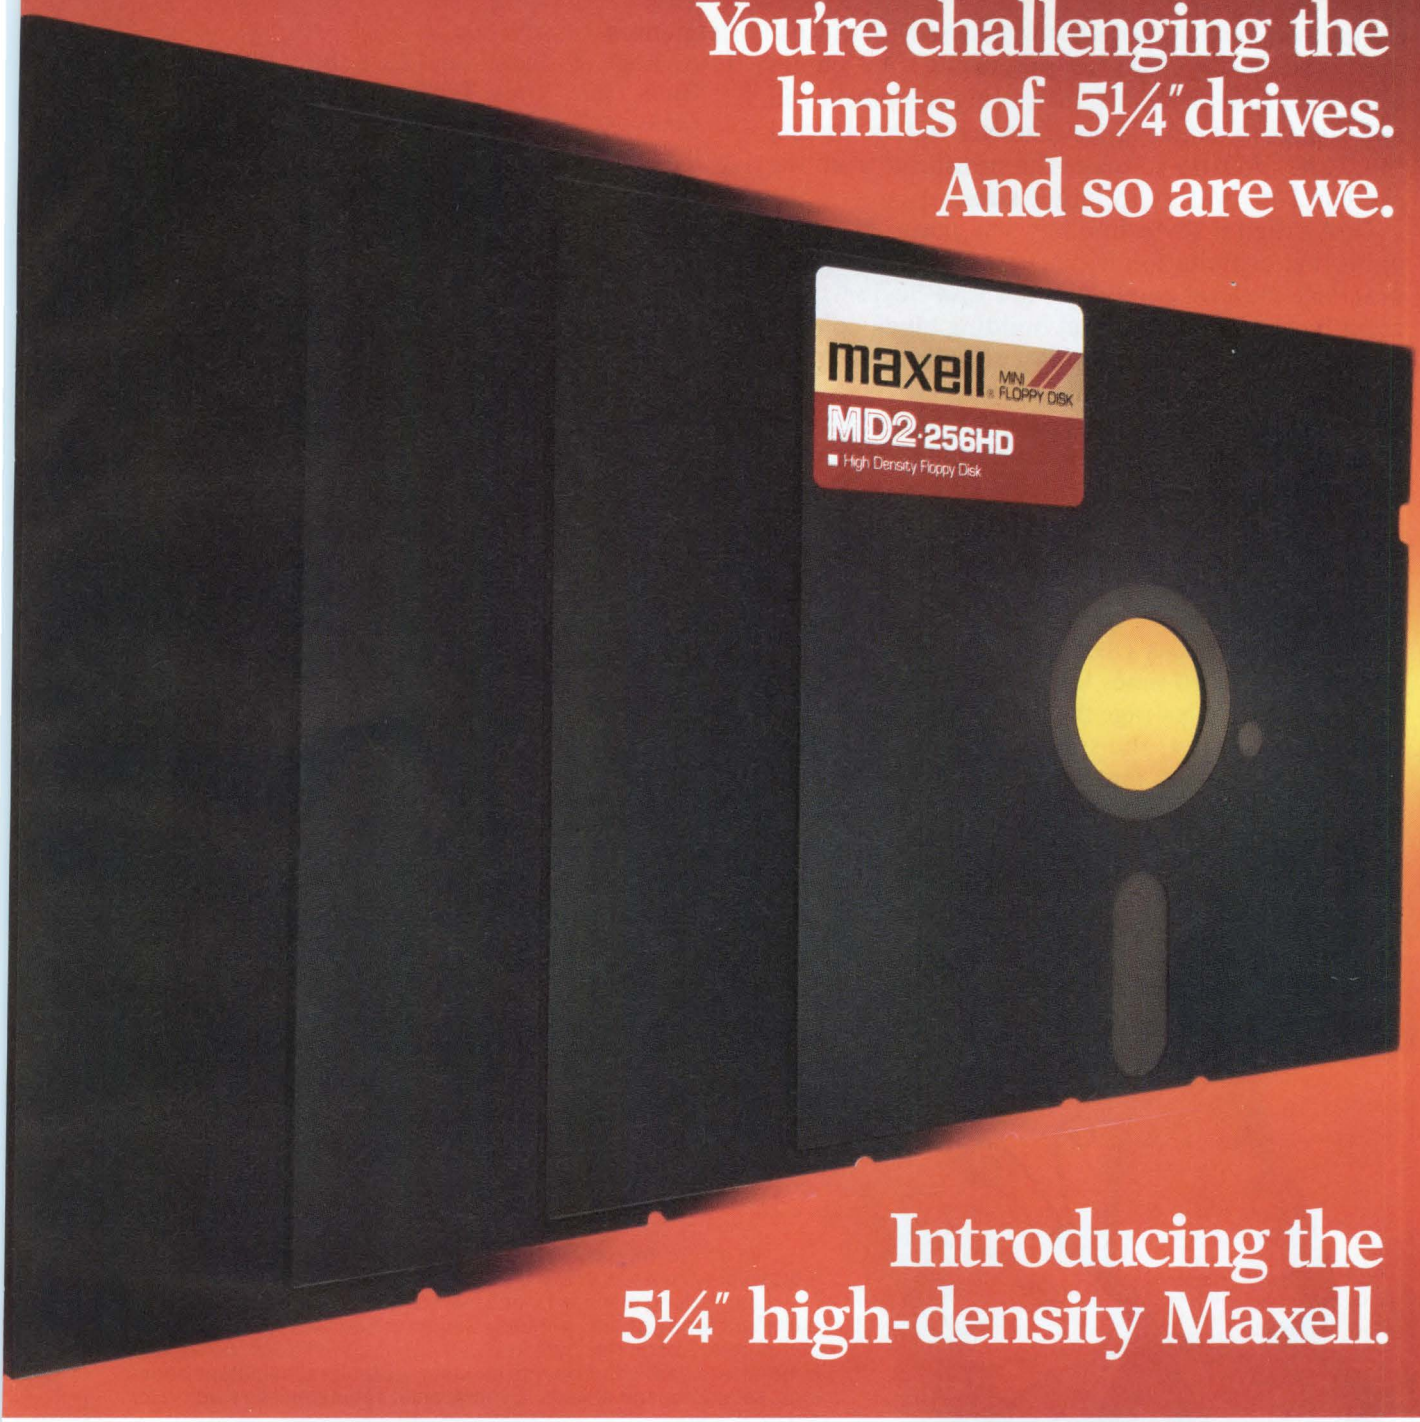

Introducing the 5¼" high-density Maxell.

Tomorrow, our technology could help turn your theory into reality.

Where will new drive capabilities come from? Your vision. And our ability to envision the disks they demand.

When a new high-density 5¼" flexible drive went on the drawing board at YE Data, we went to work on a disk that matched its unprecedented 2,000,000 byte capacity. Compressing our oxide layer to a mere one micron, we achieved a 15,000 BPI recording density at 600 oersted resolution. And made a remarkable theory a reality.

What's next? You tell us. In the meantime, we'll be making 8," 5¼"and micro floppy disks that lead the industry in error-free performance and durability. Disks made to specifications so exacting, the only standard is our own. The Gold Standard.

maxell®
IT'S WORTH IT.

Maxell Corporation of America, 60 Oxford Drive, Moonachie, N.J. 07074 201-440-8020

CIRCLE 92

equipment, such as cassette-to-cassette wafer transport systems to GaAs IC processing. As silicon wafers increase in size, the processing equipment will also change to handle larger diameter wafers, forcing the GaAs industry to make the transition from the present 3-in. wafers to larger sizes.

GaAs digital IC technology has matured to the point that exploitation in high speed systems is now practical. Several companies are now beginning to manufacture SSI/MSI digital building block components, and LSI devices such as gate arrays and RAMs may be commercially available in 1985. GaAs chip complexity will continue to advance during the next several years as wafer material and chip process technology continue to mature. Gate arrays as large as 5000 or 6000 gates can be expected, and SRAMs in the 4- to 16-Kbit range will be on the market. While speed will be an important feature of these devices, power dissipation will be an even greater advantage. The ability of GaAs LSI devices to operate at ECL speed with 5 to 10 times lower power dissipation is more likely to encourage their use than is their speed alone. In addition, the speed and power of GaAs LSI provide a good match to the VLSI complexity and power of CMOS. Very likely, a new generation of high speed, low power digital systems will emerge in the 1980s that will use a combination of GaAs and CMOS.

References

1. A. Rode, T. Flegal, and G. LaRue, "A High Yield GaAs Gate Array Technology and Applications," *GaAs IC Symposium Technical Digest*, Oct 1983, pp 178-181.
2. S. Lee and A. L. Bass, "A 2500-Gate Bipolar Macrocell Array with 250-ps Gate Delay," *IEEE Journal Solid State Circuits*, SC-17, Oct 1982, pp 913-918.
3. T. Saigo, et al, "A 20-K-Gate CMOS Gate Array," *IEEE Journal Solid State Cicuits*, SC-18, Oct 1983, pp 578-584.

Thomas M. Reeder is marketing manager for the Tektronix, Inc, GaAs Integrated Technologies Unit, PO Box 500, Beaverton, OR 97077. He holds an MS in electrical engineering and PhD from Stanford University.

Ajit G. Rode is product line manager for the Tektronix, Inc, GaAs Integrated Technologies Unit. He received a PhD in electrical engineering from Oregon State University and an MBA from the University of Portland.

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

High 725 *Average 726* *Low 727*

Just how long should standby power stand by?

Five years? Ten years? Twenty?

Quite frankly, we don't know.

But we do know that Gates Energy cells retain better than 80% of their rated capacity for eight to ten years at 23°C. in float applications.

Compared with our nearest competitors, our energy cells may qualify as the Methuselah of standby power.

When it comes to cold weather performance, our cells deliver 50% of their C/10 room temperature capacity at −40°C. And the truly sealed design means total freedom from maintenance and terminal corrosion.

Gates Energy cells are rated at 2 volts with capacities at 2.5 Ah, 5 Ah, 12.5 Ah and 25 Ah. They can be assembled into an endless variety of configurations.

Learn more about Gates Energy cells. Call, or write, Gates Energy Products Inc., 1050 S. Broadway, Denver, CO 80217. 303/744-4806.

GATES ENERGY

CIRCLE 94

ULTRA-RE

R · AT&T 3B20D SUPER MINICOMPUTER · AT&T 3B

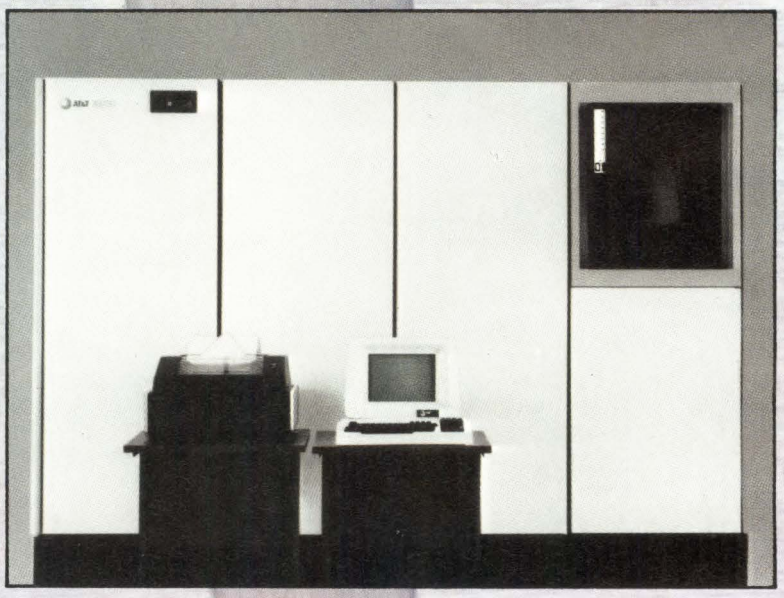

LIABILITY

OD SUPER MINICOMPUTER · AT&T 3B20D SUPER

Power. Performance. Profits.

In every way, the 3B20D from AT&T is in a class by itself. We designed this 32-bit super minicomputer to give new meaning to the concept of ultra-high reliability.

To set new standards for uptime.

To break new ground with our enhanced UNIX™ Operating System –UNIX RTR–that offers your customers maximum real time and time-sharing advantages.

And to give you the technically advanced machine you need to realize the full business potential of the lucrative multi-tasking and multi-user market.

A NEW BEST SELLER

The 3B20D will prove essential for industries whose business depends on the absolute reliability of their computer systems. Companies with reservation systems, such as airlines or hotels. Or those with command and control systems, on-line banking or financial systems.

To deliver on its promise of unfailing performance, the 3B20D from AT&T continues computation during diagnosis, hardware faults, and hardware and software updates. The result is downtime that can be measured in minutes-per-year.

Other benefits include a crash-proof file system, and a host of features that prevent memory failure and reduce programmer error.

YOU CAN COUNT ON THE 3B2 AND 3B5, TOO

For complete systems solutions, your customers will welcome the other members of the 3B family of computers. Especially because they're based on UNIX System V.

Our 3B2/300 is a powerful 32-bit supermicro that's perfect for situations where a number of people require desk-top computers. It can support up to 18 users at individual workstations–all sharing peripherals and data.

Our 3B5/100 and 3B5/200 can accommodate up to sixty users, without sacrificing response time. You'll find these super minicomputers powerful tools for a variety of applications–software development, office systems and CAD/CAM.

For all the members of the 3B family of computers, AT&T pledges a full program of service and support, including a comprehensive and flexible program for OEMs and VARs.

That and the unsurpassed research and development capability of our own AT&T Bell Laboratories make an unbeatable combination. Machines that won't let your customers down, coupled with support that won't let you down.

Only from AT&T.

If you'd like more information about AT&T Computers, call 1-800-833-9333 or use our international telex number (210-209).

AT&T

You can put a better keyboard in your enclosure without modifying the enclosure.

Stackpole's new KS-600E keyboard–featuring snap-in modular switches–gives you a low-cost encoded keyboard that will retrofit into your present enclosure without retooling.

The discrete snap-in modular switches assemble into a metal plate to give you any kind of main or ancillary array layout you need. The metal plate assures alignment of the switches, adds rigidity and aids in RFI/EMI suppression.

You don't have to tool up for a new enclosure or modify the present one. We'll do the retrofitting for you on the backplate, and all you do is drop it in and connect it up.

And we can provide our custom encoded scheme, or design to your proprietary encoded format. Plus we let you choose between tactile or linear feel–at no additional cost.

The ergonomically designed full travel switches have a total stroke of .140 (3.6 mm nominal), are rated at 50 million cycles, are wave solderable and field repairable. The switches contain Stackpole's patented twin bifurcated contacts, with four make-point redundancy, which assures you of switch closure. And today's sophisticated electronics make mechanical switch operation reliable and dependable...at lower cost.

Stackpole also gives you the option of purchasing our modular switches individually, or in sets of configured arrays for your in-house assembly.

We have a variety of keycap styles. And we offer a major cost saving through our large legend library...tooled and in place...including the most popular international legends.

Our Application Engineers are ready to help, our domestic and off-shore facilities keep our prices attractive.

Call or write us today for complete information.

STACKPOLE

Stackpole Components Company
P.O. Box 14466, Raleigh, NC 27620
Telephone: 919/828-6201

CIRCLE 95

Anaheim Convention Center and Disneyland Hotel
Anaheim, California
October 1 to 4

All aspects of this year's AUTOFACT conference are based on the central theme of CIM—computer integrated manufacturing. But, as at all past AUTOFACT conferences, stress will be on providing practical information for each attendee. Intent of the program committee was that it would be "impossible that anyone could walk away from AUTOFACT 6 without having acquired knowledge that can be taken back home and put to work to improve their manufacturing operations."

Yet the technical sessions do not seem to be aimed just at manufacturing operations. Most are highly technical in nature and potentially of significant value to both design engineers and system integrators. Without question, however, the program achieves its intent of being "practical."

Backing up a technical program of 22 general and three integration sessions are three evening forums and 14 separate-fee tutorials—as well as a plenary session that serves as the conference opener. Of particular interest among the general sessions are subjects such as expert systems, simulation, robotics, networking, and microcomputers. Sessions of equal potential include those on CAD/CAM unification, flexible manufacturing systems, and the factory of the future.

Technical program

Participants in the plenary session, scheduled for 1 pm on Monday, October 1, will be representatives from both industry and government. This overview session will consist of discussions of the impact of "high tech" on manufacturing procedures and the trends that are emerging. Particular emphasis will be on the effect of today's technology on government policies, human resources, and management planning.

Integration sessions will precede the technical sessions in the first time period of each morning from Tuesday through Thursday. Each integration session will be more general in nature than the technical sessions but will still relate to the overall theme of "CIM—Manufacturing in the Information Age." Some typical subjects of discussion will include new generation computers and their implications for computer integrated manufacturing, the possible failure of technology to meet expectations, flexible manufacturing, motivating performance systems, and the factory of the future, a subject that seems to be ubiquitous at this conference. At press time, some modifications were still being made to expand these sessions even further.

The first of two technical sessions on CAD/CAM unification will take place on Tuesday morning. In this session, speakers from companies such as General Motors and Calma will discuss applications of CAD and CAM and some of the software involved. Similarly, speakers at the second session—on Thursday morning—will discuss how CAD/CAM fits into the overall automation picture. How this occurs and what advantages can be expected will be included in the paper presentations.

As might be expected, emphasis of the expert systems session will be on the application of this phase of artificial intelligence to computer integrated manufacturing. Both specific and overall applications will be covered in this Tuesday morning session. For example, two speakers will present their views and experiences on AI for such limited applications as computer numerical control milling and machining operations. But other speakers will cover the broader areas of flexible manufacturing systems and process industries as a whole.

(continued on page 248)

(continued from page 247)

Immediately after lunch on Tuesday, the session on simulation will involve both CAD/CAM and robotics. But it will also cover other areas such as realtime simulation and its potential for reducing the risks of industrial automation. A scheduling step that places the session on robotics in the same time slot will, unfortunately, cause some attendees to move back and forth between meeting rooms. This session will place even more stress on robotics, of course, with a mix of applications for both the future and the present. Speakers will be from among some of the better known users of large robotic systems, such as General Motor's Fisher Body Division, Westinghouse Electric, and McDonnell Douglas.

Local area networks will be stressed in the Wednesday morning session on networking. Of course, all discussions of LANs will tie them to factory applications, whether the factory of today or of the future, and particularly to CAD and CAM. Selection of network architecture and the use of distributed systems with reasonable consideration of future requirements will be among the topics included. An alternate paper may be presented on packet switching for factory integration.

Flexible manufacturing systems are included in the discussions of many, if not all, of the sessions of AUTOFACT 6. However, one of the after-lunch sessions on Wednesday is dedicated to just that subject. Scheduled papers will cover controllers, machine tooling, and a wire harness manufacturing application. Unfortunately, one of the potentially most important papers—on synthetic vision—is listed as only an alternate. Hopefully that paper will be included, but its status was unknown at press time.

Long ago, microcomputers began to influence manufacturing operations, just as they have nearly every other segment of industry. Only recently, however, has their dominance been so great—largely due to the ever-present personal computer. The last session on Thursday morning recognizes the place of microcomputers/personal computers in computer integrated manufacturing operations and other specific areas of the factory. Subjects of discussion include data acquisition and CAD systems.

In the same time spot will be a key session on the factory of the future, with stress on planning and modeling. AI again will appear in this session, as will CAD/CAM. However, here the focus will be on the factory environment that is envisioned for the year 2000 and beyond—as well as the more relevant future of the next few years.

Special sessions and tutorials

Forum sessions will be held as part of the regular program from 6:30 to 9:30 on Monday, Tuesday, and Wednesday evenings. These sessions—on human resources, AI, and education—were considered to be of sufficient importance to warrant attendance during evening hours.

Among these sessions, the one on human resources is likely to be especially important. The introduction of technological changes in any industry most often results in periods of uncertainty among employees, whether in the office, the laboratory, or the backroom research facility. However, the factory floor seems to have even more potential for problems. Speakers in this session will discuss management objectives, training requirements, and other factors related to maintaining the necessary interface between employees—the human resources—and the production objective.

At press time, details on the Tuesday evening forum session on AI were unavailable. However, the subject itself should be sufficient to draw heavy attendance. The Thursday evening forum will be on education, particularly on the training and retraining of factory personnel—human resources again, but with a different emphasis.

A group of separate-fee tutorials will also be present in the 6:30 to 9:30 pm time spans, on Monday, Tuesday, and Wednesday. A registration fee of $85 for each tutorial includes the cost of course material. Those sessions expected to be the most popular will be repeated to allow as many attendees as possible to participate.

Some tutorials of particular interest to system designers and integrators include CAD/CAM implementation and integration, database management, and applied machine vision. Other tutorials are aimed more toward the manufacturing phases of CIM: production scheduling, material handling, batch manufacturing, automated shop floor control, and manufacturing strategies.

As its title implies, the Monday and Wednesday tutorials on CAD/CAM implementation and integration will provide practical discussions on equipment selection, interfacing to existing systems, and computer aided engineering. The discussions will involve

(continued on page 250)

GET IT TOGETHER

WITH BNC COAXIAL CONNECTORS FROM ntt

These male BNC/TNC type connectors meet the mechanical, electrical, dimensional, and environmental specifications of MIL 39012 and are available for a wide range of cables. Since the contact pin is factory assembled into the BNC/TNC, the one-step pin crimp/ferrule crimp termination is very quick and simple using NT•T assembly tools. Cable assemblies manufactured to customers' specifications are also available.

BNC/TNC receptacles are available in standard panel mount, isolated ground panel mount, and PCB mount styles.

PCB mount receptacles are offered in grounded or insulated styles with either screws or solder posts for mounting. Both grounded and insulated types have machined metal threads which eliminate stripping that can occur with molded plastic threads. Another important feature is the PCB footprint which is identical to Amp's footprint. PCB mount receptacles have molded troughs which isolate the leads and greatly reduce fluctuations in impedance.

Added to cost effectiveness and improved design features is NT•T's reputation for short lead times.

Choose National Tel•Tronics and You've Chosen Selection, Performance, and Cost Effectiveness.

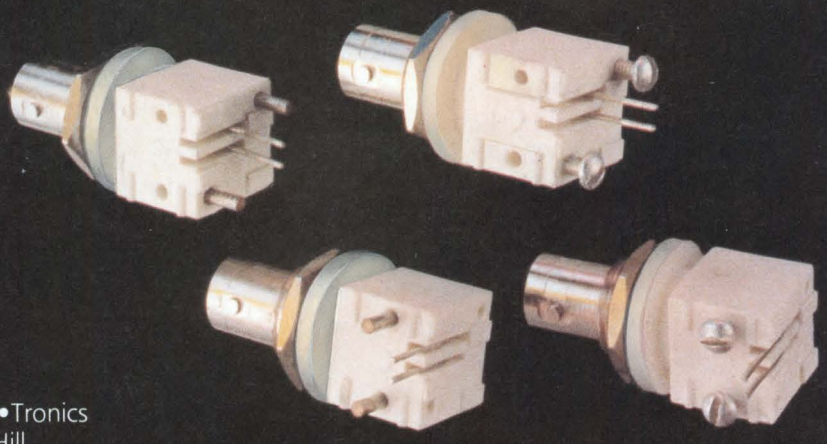

ntt
NATIONAL
TEL·TRONICS

National Tel•Tronics
State Road Hill
Meadville, PA 16335
(814) 724-6440

CIRCLE 96

(continued from page 248)

workstations, turnkey systems, supermini and mainframe computers, and the need for equipment standardization. Other subjects will include database development, transition between old and new technologies, data communications, cost justification and recovery, and management involvement—another level of human resources.

A single tutorial on database management will be conducted Monday, giving a general overview of the technology and its application to CIM. The first part of the tutorial will focus on concepts and terminology, the second part will discuss how database computers and management systems work, while the third will detail specific applications. A portion of the overall discussion will examine current systems and a forecast of Japan's fifth-generation computers.

On Monday and Wednesday there will also be tutorials on applied machine vision—one of the most important subjects covered by AUTOFACT 6. Machines that literally can see—and therefore could be capable, at least in theory, of making decisions—are no longer "for the future." They are used today in many applications, not the least of which is inspection of manufactured parts. However, machine vision should not necessarily be considered as either limited to, or better than, human sight. In fact, the person conducting this tutorial has stated in the past that "there is no reason to impose human-like limitations on (machine vision) technology." Participants in this tutorial will become familiar with the techniques of image acquisition and analysis as they pertain to common uses of machine vision.

Product exhibits

An AUTOFACT 6 exhibition will run concurrently with the conference on Tuesday, Wednesday, and Thursday—from 11 am to 8 pm the first two days and from 11 am to 5 pm the third day. A scan of the preliminary list of exhibitors indicates that many of the companies that supply components, subsystems, or systems to the process control and automated manufacturing industries will show their products.—*S.F.S.*

AUTOFACT 6 is sponsored by the Computer and Automated Systems Association of the Society of Manufacturing Engineers (CASA/SME). For more information on registration and housing, contact SME, One SME Dr, Dearborn, MI 48121. Tel: 313/271-0777

Technical Program*

Mon, Oct 1

1 pm

Plenary Session

6:30 to 9:30 pm

Human Resource Forum

"Human Resource Function in High Technology," Robert Miller, A. T. Kearney, Inc, Los Angeles, Calif

"CAD/CAM—A Management Objective," Linda G. Scheffler, Otis Engineering Corp, Dallas, Tex

"Training Requirements for a CAD/CAM System," Gabriel E. Jerahov, Hughes Helicopters, Inc, Cyprus, Calif

"Human Resources are Human Factors," Charles L. Carter, Jr, VMX, Inc, Richardson, Tex

"High Touch Requirements for High Tech CAD/CAM," Luciano L. Traversa, Sunstrand Advanced Technology Group, Rockford, Ill

Alternate: "Introducing Technology Change," Susan R. Maxwell, Data General Corp, Portsmouth, NH

Tues, Oct 2

8:30 to 10:30 am

Integration Session

Chair: Frederick J. Michel, Army Material Development & Readiness Command, Alexandria, Va

"New Generation Computers and Their Implication for CIM," Daniel S. Appleton, D. Appleton Co, Inc, Manhattan Beach, Calif

"Computer Integrated Flexible Manufacturing," George Hess, Ingersoll Milling Machine Co, Rockford, Ill

"Too Bad!...They Thought Technology Was the Answer," Michael Hora, A. T. Kearney, Inc, Chicago, Ill

(continued on page 252)

*Based on information available at press time. Subject to change.

Our name isn't on the computer, but our Celanex® puts quality in it.

Celanese doesn't make computers, but our Celanex® 2000-2 unreinforced thermoplastic polyester makes keycaps wear better, longer, and look more colorful. Printed Celanex® keytops eliminate costly, time consuming two-shot molding processing, providing wear resistant, multi-colored lettering capability. These features, combined with outstanding processing characteristics, are why Hewlett-Packard selected Celanex® 2000-2 for the terminal keytops of their HP 150 Touchscreen Personal Computers. Quality demands quality.

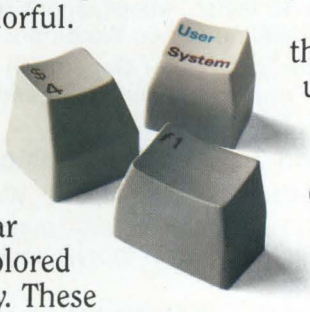

Celanex® 2000 series thermoplastics offer a unique combination of printability, chemical and wear resistance, dimensional stability, strength, stiffness and surface gloss to withstand the tests of time. In fact, the molding experts at Hewlett-Packard report Celanex® 2000-2's internal lubricant system facilitates mold release without mold plate-out–maximizing productivity.

Fast, easy, cost-effective. That's Celanese quality at work!

For all the quality facts, contact Frank Esposito, Celanese Engineering Resins, 26 Main Street, Chatham, NJ 07928; (201) 635-2600, ext. 4388.

CELANESE QUALITY

CIRCLE 97

(continued from page 250)

10:45 am to 12:45 pm: Sessions 1 to 4

Session 1: CAD/CAM Unification

Chair: John Dohner, Calma Co, Southfield, Mich

"CAD/CAM and Precision Automotive Sheet Metal Fabrications," Lester C. Teague, II, General Motors Corp, Warren, Mich

"Critical Concepts of Modern Coordinate Measuring Systems," William S. W. Tandler, Multi Metrics, Inc, Redwood City, Calif

"Software Package Selection and Integration," Oscar G. Rhudy, Rust International Corp, Birmingham, Calif

"A Strategy for Implementation of Hierarchical Control of CIM," Forrest C. Gale, Defense Systems Management College, Fort Belvoir, Va

Session 2: Expert Systems

Chair: Edward J. Adlard, Metcut Research Associate, Inc, Cincinnati, Ohio

"Application of Artificial Intelligence in Capital Intensive Process Industries," John P. Elwood, Stevens Institute of Technology, Hoboken, NJ

"A Knowledge-based System for Generative Micro-Planning of Machining Operations," Brian E. Barkocy, Metcut Research Associates, Cincinnati, Ohio

"An Application of Expert Systems in FMS," Malcolm D. Hall, and Glenn Putnam, Cameron Iron Works, Houston, Tex

"Automated CNC Milling by Artificial Intelligence Techniques," Kenneth Preiss and E. Kaplansky, Ben-Gurion University of the Negev, Beer Sheva, Israel

Alternate: "Utilization of Artificial Intelligence in Manufacturing," David Liu, Hughes Aircraft Co, El Segundo, Calif

Session 3: New Technology Management

Chair: Robert E. Johnson, Burroughs Corp, Radnor, Pa

"Systems and Technology Management: Trends and Issues," James W. Mottern, Ernst & Whinney, Newport Beach, Calif

"Enabling Requirements for Advanced Manufacturing Technology Modernization," Ronald B. Ellis, Industrial Technology Modernization, Menasco, Euless, Tex

"Island of Technology in a Sea of Confusion," Robert E. Gutshall, Management Advisory Services, Los Angeles, Calif

"Executive Strategies for Computer Integrated Manufacturing," Daniel P. Mincavage, National Systems, Inc, Ann Arbor, Mich

Alternate: "CAM—Integration of Word Processing, CAD, 4M, and MRP," Robert R. Clifford, Singer Co, Little Falls, NJ

Session 4: Financial Justification Session

Chair: Gordon McAlpine, Computervision, Southfield, Mich

"Application of Computer Assisted Techniques to Indirect Cost Identification and Reduction," James B. Ayers, Theodore Barry & Associates, Los Angeles, Calif

"An Approach to Computer Aided Process Planning Cost Justification," William D. Cochran, Garrett Turbine Engine Co, Phoenix, Ariz

"Manufacturing Planning & Control Systems Justification," Barry Goss, NCA Corp, Santa Clara, Calif

1:45 to 4:45 pm: Sessions 5 to 8

Session 5: Simulation

Chair: Don T. Phillips, Texas A & M Univ, College Station, Tex

"CAD/CAM in Robot Applications, Design and Simulation," Matti A. Katajamaki, NOKIA Corp, Helsinki, Finland

"Realtime Simulation Eliminates the Risk of Industrial Automation," Max W. Hitchens, HEI Corp, Carol Stream, Ill

"Flexible Manufacturing Simulation and Setup Reduction," John Bernard, Robotics Technology, Kingwood, Tex

"Reducing Inventories through Computer Simulation and Setup Reduction," Dennis Wormington, Texas Technical Univ, Lubbock, Tex

Alternate: "CAA Applications: The Edge in Integration," Colin G. Kelley, McDonnell Douglas Automation Co, Florham Park, NJ

Session 6: Design & Documentation Session

Chair: Anthony R. Skomra, Automation Technology Products, Campbell, Calif

"Integral Link between Geometric Modeling and CAM Applications," A. Kater Elgabry, General Electric Co, Schenectady, NY

"The Assignment of Tolerance Attributes to a Solid Modeler in Tolerance Chart Format," Lloyd R. Fields, Ohio State Univ, Columbus, Ohio

"Engineering Change Controls: The Manufacturing Bridge," Michael J. Aghajanian, Ernst & Whinney, Newport Beach, Calif

"Electronic Signoff—A Successful Manufacturing Implementation," Wayne E. Thomas, Garrett Turbine Engine Co, Phoenix, Ariz

"How to Write an Effective Procedures Manual," Vince K. Gulati, Management Advisory Service, and Keith Walkup, Price Waterhouse, Houston, Tex

Alternate: "Computer Generation of Drawings and N/C Part Programs," Paul W. Brauninger, Ex-Cell-O Corp, Traverse City, Mich

Session 7: Robotics Session

Chair: Michael V. Grobel, GMC, Warren, Mich

"Robots and Automated Systems for World Class Quality—The GM-20 Story," Jack H. King, GMC, Warren, Mich

"Integration of a Multirobot Process Line Under Mini Computer Control," John F. Follin, GA Technologies, San Diego, Calif

(continued on page 254)

One small step for Digilog.
One giant leap for data comm test.

New family of automatic protocol analyzers interpret and isolate the error source for you.

Meet the New Digilog 200, 400, 600, and 800 . . . four revolutionary diagnostic analyzers that can actually describe in plain English what is happening on your data lines.

They decode and interpret the protocol for you. Isolate faults automatically. Exactly. Even statistically analyze data line performance.

Digilog 200 automates data comm testing

Digilog 200 sets up automatically. Decodes protocols automatically. Analyzes protocols automatically. Tests devices automatically. And identifies the faults automatically.

You get EE PROM packs for program or data storage, remote control, typewriter-like keyboard, full programming, menu driven traps and triggers, CRT, help screens, six BERT tests, interface breakout, printer output, etc.

Digilog 400 and 600 for greater sophistication

Here the emphasis shifts toward versatility, power and higher speed . . . up to 72K bits/second.

You get more programming power, bigger CRTs, built-in 3½" micro disks, soft function keys, graphics, protocol simulation plus all the features of the 200.

Digilog 800, the new technology

Use the world's first protocol/performance analyzer to conquer the toughest data comm problems. Debug software programs. Improve your network.

The Digilog 800 offers you fully selective and bit image recording, full protocol simulation, automatic protocol analysis through level 3 (X.25, SNA, etc.), and speeds to 256K bits/second.

You also get on-line color graphics, internal 10 megabyte disk with selective logging and comprehensive statistical analysis with reports.

Digilog is out in front

Want to see what the new generation of Digilog analyzers can do for you? Call now for a demonstration (215) 628-4530.

DIGILOG INC., 1370 Welsh Road Montgomeryville, PA 18936

DIGILOG®
NETWORK CONTROL DIVISION

CIRCLE 98

(continued from page 252)

"Control Software for an Advanced Sensor-Based Robotic Assembly Station," Karen A. Hope, Westinghouse Electric Corp, Columbia, Md

"Offline Programming to the Factory Floor," Keith E. Williams, Sr, McDonnell Douglas Automation, St Louis, Mo

"Hard Automation vs Robots—Making the Choice," John W. Schott, Timeshare Technical Services, Kalamazoo, Mich

Alternate: "Robots of the Eighties," Salahuddin Qazi, and Rezk G. Mohammed, SUNY College of Technology, Utica, NY

Session 8: Numerical Control Session

Chair: William B. Johnson, Rockwell International, Canoga Park, Calif

"An Integrated Turbine Blade Manufacturing System," Randy Schmid, Rigid Milling Machine Co, Rorsachacherberg, Switzerland

"Direct Computer Control—The Next Challenge in Factory Automation," George W. Jones, Computervision, Bedford, Mass

"NC Maintenance Programming as a Productivity Strategy," Leonard A. Weibel, General Electric, Milford, Ohio

Alternate: "Direct Numerical Control—A Tool for Factory Automation," Lawrence W. Coyle, McDonnell Douglas Astronautics Co, St Louis, Mo

6:30 to 9:30 pm

Artificial Intelligence Forum

Chair: Daniel S. Appleton, D. Appleton Co, Inc, Manhattan Beach, Calif

Speakers to be announced

Wed, Oct 3

8:30 to 10:30 am

Integration Session

Chair: John Vinyard, A. T. Kearney, Inc, Atlanta, Ga

"Title to be announced," James Lardner, John Deere Co, Moline, Ill

"Technology without Quality is Like a Day without Sunshine," Charles L. Carter, Jr, VMX, Inc, Richardson, Tex

"Motivating Performance Systems for Productivity Improvement," Frank Petrock, General Systems, Inc, Ann Arbor, Mich

10:45 am to 12:45 pm: Sessions 9 to 16

Session 9: Networking

Chair: Robert Merrell, Burroughs Corp, San Diego, Calif

"Factory Networks with a Future," Lawrence S. Gould, Factory Systems Planning Service, Boston, Mass

"CAM for the 80s—Distributed Systems Using Local Area Networks," Joseph A. Vrba, General Electric Co, Oak Ridge, Tenn

"CAD/CAM Communication System," Robert G. Love, Halliburton Services, Duncan, Okla

"Computer Control Systems—Selecting Network Architecture Now Which Has Flexibility in the Future," Allan F. Packer, Eaton-Kenway Co, Salt Lake City, Utah

Alternate: "Factory Integration through Packet Switched Networks," Charles L. Baecker, Ernst & Whinney, Newport Beach, Calif

Session 10: Common Data Base

Chair: Dr. H. Randall Johnson, Boeing Computer Services, Seattle, Wash

"Templates for an Integrated Common Data Base," B. Neil Snodgrass, D. Appleton Co Inc, Arlington, Tex

"A User-driven Methodology for the Design of CAD/CAM and Engineering Data Bases," Roger W. Thyr, Control Data Professional Services, Bloomington, Minn

"A Conceptual Schema for a CIM Data Base," Francois B. Vernadat, National Research Council of Canada, Ottawa, Ontario

"Development of an Integrated System Architecture—Report from the Pits," J. J. Swenson, General Dynamics, Pomona, Calif

Alternate: "A Data Model for CIM," Michael P. Carroll, Manager, Sperry Computer Systems, Blue Bell, Pa

Session 11: New Technology Management

Chair: Robert E. Johnson, Burroughs, Corp, Radnor, Pa

"Successful CAD/CAM Implementation Strategy," Rene J. Riethof, Teledyne CAE, Toledo, Ohio

"Management Attitude to Flexible Manufacturing Systems," Venkitaswamy Raju, Rochester Institute of Technology, Rochester, NY

"Implementing Automation Systems: Implications for Manufacturing Managers," Jake Krakauer, Megatek Corp, San Diego, Calif

"Blending Advanced Manufacturing Technologies with New Management Practice," Glen A. Allmendinger, Harbor Research Corp, Boston, Mass

Session 12: CIM for Small Businesses

Chair: Catherine E. Jakesy, Burroughs Corp, Detroit, Mich

"How to Begin CIM in a Small Business," Thomas Carpenter, CDS, Dallas, Tex

"Integrating Computer Simulation into a Small Manufacturing Operation," Robert M. Cowdrick, Jr, Reed Industries, Inc, Stone Mountain, Ga

"Computer Integrated Manufacturing in a Small Company—A Case Study," John R. Odom, National Automatic Tool Co, Inc, Richmond, Ind

"Architectures for Low Cost Factory Integration for Small Shops," Laura M. Caldwell, Univ of Cincinnati, Cincinnati, Ohio

Session 13: Process Planning

Chair: Edward J. Adlard, Metcut Research Associates, Inc, Cincinnati, Ohio

"Computer Aided Production Engineering—The Integration of CAPP, Engineering & Manufacturing," Gayle L. Berry, Garrett Turbine Engine Co, Phoenix, Ariz

(continued on page 256)

Ikegami Display Monitors

SUPERHIGH RESOLUTION MADE SIMPLE, SAFE AND PACKAGEABLE

For years, the quality of computer graphics has been limited by the resolution of display monitors. But at Ikegami, we're making it easy for your system to live up to its images with our new DM-2050 60-Hz noninterlaced superhigh resolution RGB monitor. Sporting a 1280 x 1024 pixel display and 100-MHz-plus video bandwidth, it gives you reliability and stable operation, eliminating the need for a long-persistence phosphor.

Simpler is better
How did we do it? By keeping it simple. And that we did. The chassis of our new DM-2050 is a marvel of compact, modular and functional design. In fact, Ikegami's unique design configures all the circuit boards by function within the chassis.

Keeping cool
We also took some steps to make sure our monitors wouldn't blow your designs. For example, we made the DM-2050 run so cool it doesn't require a blower. That means your monitor will never overheat because of blower failure. It also means the DM-2050 is lighter and more efficient.

Playing it safe
Along with meeting these time-honored Ikegami standards of simplicity and reliability, we also made our new 64-KHz monitor meet a few others along the way. Like UL, DHHS, CSA, IEC, FCC and VDE standards for safety and emissions. Now you can put the

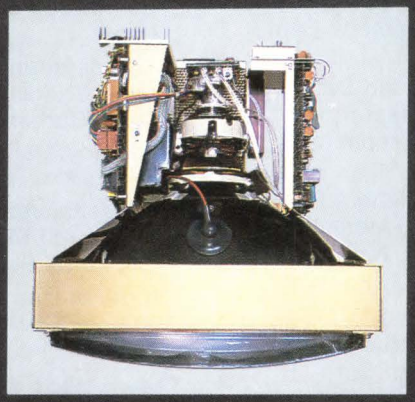

same monitor in every system you sell no matter where in the world it is used.

Keeping a low profile
But no matter what system you put the DM-2050 into, you know it will always fit in. The unique packaging of the DM-2050 fits right into CDA 19V enclosures. So in many cases, you won't have to change your existing design. And its extremely compact size gives you a free hand in new designs.

The price is right
You might think that a superhigh resolution monitor as advanced as your designs would come at a state-of-the-art premium. But the DM-2050 won't

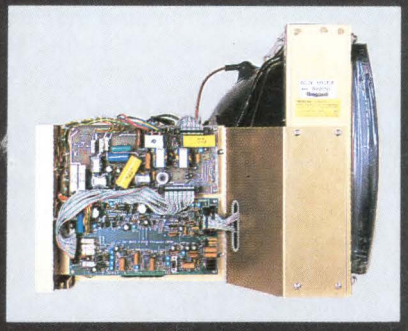

put a drain on your system's resources. In fact, Ikegami will make your profit picture look even better, because it's a finished product with on-time delivery.

In addition to the DM-2050, Ikegami makes a full line of color and monochrome monitors for virtually all OEM applications. For information, contact Ikegami Electronics (U.S.A.), Inc., 37 Brook Avenue, Maywood, NJ 07607; (201) 368-9171. In the West, 3445 Kashiwa Street, Torrance, CA 90505; (213) 534-0050.

Ikegami®

CIRCLE 99

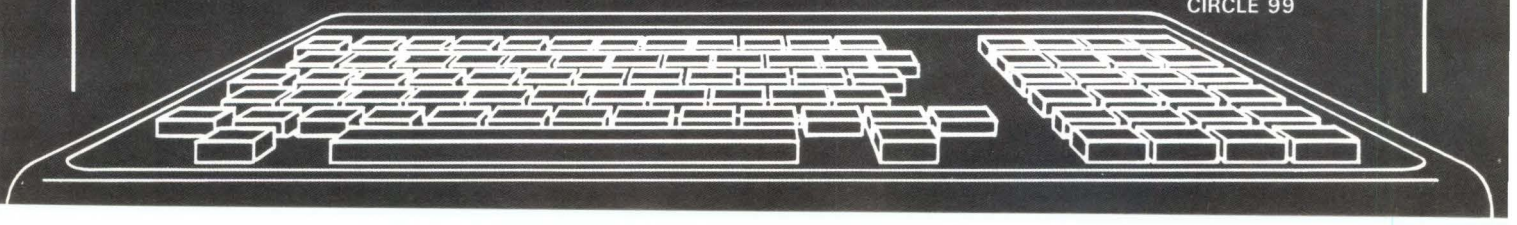

(continued from page 254)

"The Implementation of Group Technology Support System (GTSS) for Process Planning, "Robert A. Carringer, U.S. Air Force, Wright-Patterson Air Force Base, Ohio

"The Future of CAPP," Gerald L. Nordland, Digital Equipment Corp, Marlboro, Mass

Alternate: "Group Technology—CAPP Interface, A Base for Integration," David G. Williams, Babcock & Wilcox, Barberton, Ohio

Session 14: Quality, Inspection, & Test

Chair: C. L. Carter, Jr, VMX, Inc, Richardson, Tex

"Acceptable Quality Levels are Not Acceptable Anymore," Susan Lloyd McGarry, Factory Systems Planning Service, Boston, Mass

"Integrated Metrology System," Gary L. Bowen, Booz, Allen & Hamilton, Productivity Technology Center, Cleveland, Ohio

"Integration of Automated Inspection and Data Base Modification Capabilities into Manufacturing Systems," Neil A. Duffie, Univ of Wisconsin-Madison, Madison, Wis

"In-process Control of Machining—Milling a Part According to its Function," Peter R. Albrecht, Mechanical Technology Inc, Latham, NY

"Computer Aided Tolerance Analysis," Andres V. Karolin, GMC, Columbus, Ohio

Alternate: "New Techniques for Industrial Vision Inspection," Stanley N. Lapidus, Itran Corp, Manchester, NH

Session 15: Flexible Manufacturing Systems for the Electronics Industry

Chair: John Orphanos, U.S. Air Force, Hanscom Air Force Base, Bedford, Mass

"Paper to be Announced," Fred Good, Hughes Aircraft Co, Fullerton, Calif

"A Wire Harness Manufacturing System," Robert H. Sturges, Jr, Westinghouse Electric Corp, Pittsburgh, Pa

"Flexible Manufacturing Controller (FMC)—The Ultimate Tool for Factory Integration," Meir Weinstein, Digital Equipment Corp, Maynard, Mass

"The Selection of Optimal Operating Conditions for Machining Tools in an FMS," Rajendra J.S. Melville, Booz, Allen & Hamilton, Inc, Lexington, Mass, and H. H. Cook, Massachusetts Institute of Technology, Cambridge, Mass

Alternate: "Synthetic Vision: A Key to Flexible Manufacturing Systems," Russell H. Peterson, Synthetic Vision Systems, Inc, Ann Arbor, Mich

Session 16: Maintenance

Chair: John L. Winter, John Deere Component Works, Waterloo, Iowa

"Maintenance—A Profit Center," John D. Andrica, A.T. Kearney, Chicago, Ill

"Reducing NC/CAM Maintenance Downtime & Costs: Increasing Automated Systems Productivity," Howard C. Cooper, HCC & A, Kaysville, Utah

"Material Management in the Maintenance Operation," William DeVaney, Stanley-Vidmar, Inc, Allentown, Pa

"Title to be Announced," Paul H. Coehn and Jeya Chandra, Pennsylvania State Univ, University Park, Pa

6:30 to 9:30 pm

Education Forum

Chair: Diane Stotler, Yavapai College, Prescott, Ariz

"A Parallel Process CIM Architecture," Craig R. Skevington, Renssalaer Polytechnic Institute, Troy, NY

"Making the CAM Connection—Educating for an Integrated Manufacturing World," Mark A. Cooper, California Polytechnic State Univ, San Luis Obispo, Calif

"Considerations for DNC—Updating Your Factory to FMS," William A. Crawford, CADAM Inc-Lockheed Corp, Burbank, Calif

"Graphics Training: Succeeding with Reality," James M. Doherty, U.S. Electrical Motors, Milford, Conn

"A Baccalaureate Degree Model for Computer Integrated Manufacturing," Arthur L. Foston, California State Univ, Fresno, Calif

Alternate: "Industry/Higher Education Linkage," Russell F. Jerd, Sinclair Community College, Dayton, Ohio

Thurs, Oct 4

8:30 to 10:30 am

Integration Session

Chair: Charles Skinner, Booz, Allen, & Hamilton, Inc, Cleveland, Ohio

"Factory Systems—The Leading Edge," Jules Mirabal, Eaton-Kenway, Salt Lake City, Utah

"Manufacturing Technology for Factory 2000," Gary Michaelson, Boeing Co, Seattle, Wash

"Creating A Winning Automation Strategy," Frank Curtin, General Electric Co, Charlottesville, Va

10:45 am to 12:45 pm: Sessions 17 to 22

Session 17: Microcomputers

Chair: Kenneth C. Bonine, General Dynamics, San Diego, Calif

"Micros Invade the Factory," Nowlan M. Ulsch, Factory Systems Planning Service, Boston, Mass

"Facilitating Computer Integrated Manufacturing through a Microprocessor-based Acquisition System," Shyamal K. Ganguly, Integral Computer Systems, Inc, Putnam, Conn

"Personal Computer CAD Systems," Terry T. Wohlers, Colorado State Univ, Fort Collins, Colo

"Paper to be Announced," David L. Evans, General Dynamics, San Diego, Calif

(continued on page 258)

DATA GENERAL'S MV/10000. LESS MONEY THAN DEC'S VAX 11/780 AT TWICE THE SPEED.

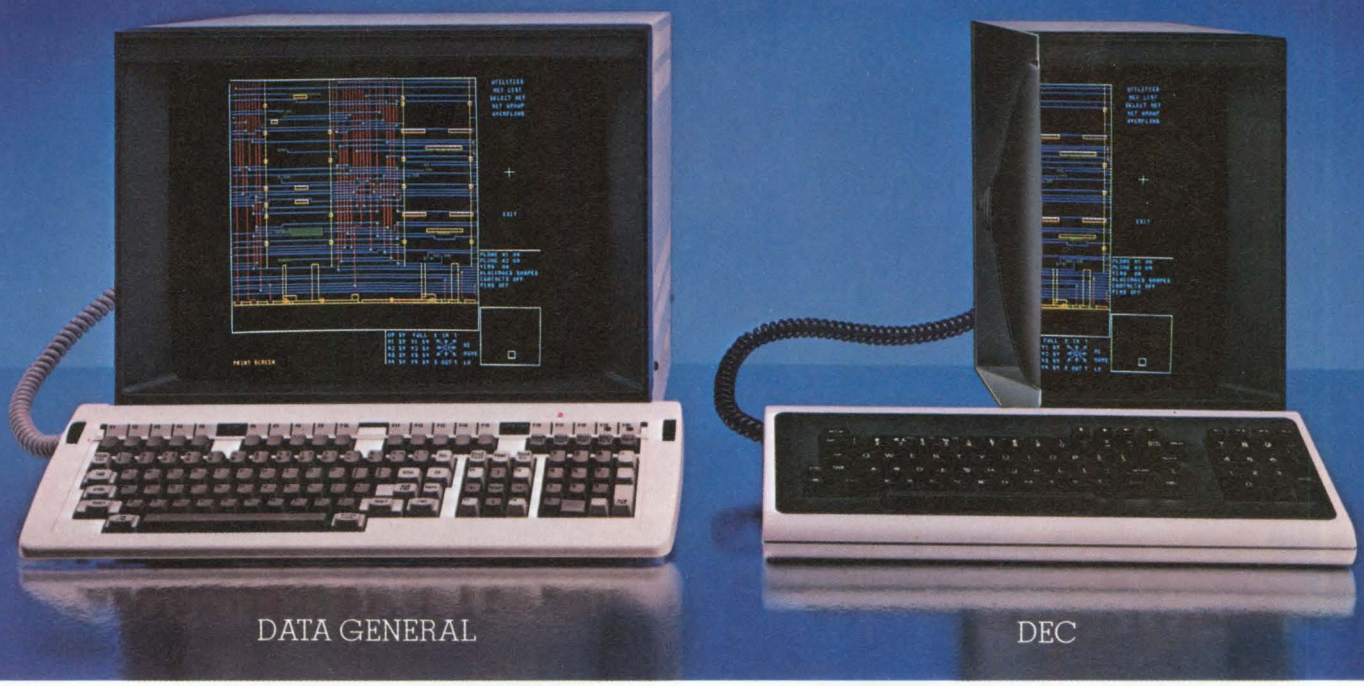

DATA GENERAL'S MV/FAMILY–TOP PERFORMANCE, NOT TOP DOLLAR

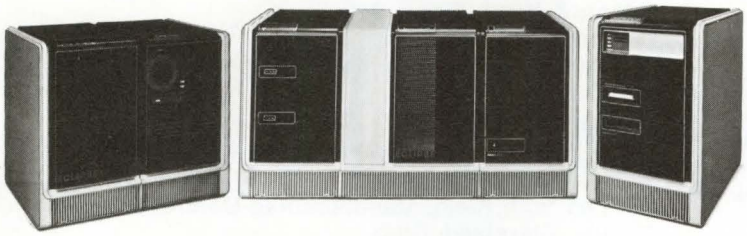

Forget VAX.™

Data General's ECLIPSE® MV/Family of 32-bit computers brings you the best price/performance available for engineering applications—while running some of the best electronics engineering software.

TWICE AS FAST AS VAX

Consider the price/performance graph shown below. On the basis of dollar-per-MIP, Data General's MV/10000™ gives you twice the performance of the VAX 11/780—at a lower price.*

Now compare the MV/8000®II to the VAX 11/780. Same performance. But the MV/8000II is half the price. The same holds true when you compare the MV/4000® to the VAX 11/750. And our recently announced OEM MV/8000 C offers almost twice the performance of the VAX 11/750. But it's the same price.

RUNS THE MOST WIDELY-USED SOFTWARE

You can run all of the best software on our ECLIPSE MV Series, including TEGAS,™ NCA/DVS,™ MicroSet-86,® ECAD's DRACULA, and Mentor's CADISYS.®

Data General keeps you a generation ahead with comprehensive service plans and industry standard software development environments.

For further information, write Data General Corporation, M.S. F134, 4400 Computer Drive, Westboro, MA 01580.

*Compared to the VAX 11/785, our MV/10000 offers 40% better performance for two-thirds the price.

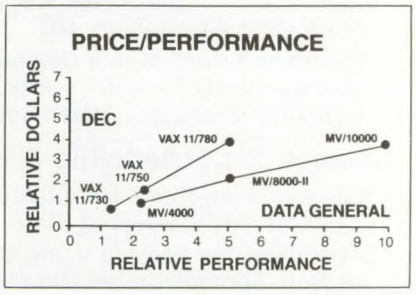

Data General
a Generation ahead.

CIRCLE 100

(continued from page 256)

Session 18: CAD/CAM Unification II

Chair: James F. Yevtich, General Motors Corp, Warren, Mich

"CAD to CAM—Making the Concept Work," F. Edward Bell, Graphics Manufacturing Systems, Inc, Santa Ana, Calif

"Building Blocks for Automation," Gary L. Resnick, Applicon-Schlumberger, Burlington, Mass

"Improved Responsiveness through CAD/CAM," Robert F. Ouimette, Honeywell Inc, Bloomington, Minn

"Integrated CAD/CAM at Packard Electric," William D. Engelke, General Motors, Warren, Ohio

Alternate: "The New Role of CNC Controls in CIM," Michael H. Smith, Vega, Inc, Troy, Mich

Session 19: Factory of the Future—Planning & Modeling

Chair: Don L. Norwood, Vought Corp, Dallas, Tex

"The Manufacturing Environment in the Year 2000," William H. Slautterback, Koppers Co, Inc, Baltimore, Md

"Group Technology—A Framework for CAD/CAM Integration," Gordon W. Millar, Boeing Commercial Airplane Co, Seattle, Wash

"A Survey of Application Areas for Artificial Intelligence in Manufacturing Planning & Control," Richard J. Mayer, Robert C. Young, and D. J. Phillips, Texas A & M Univ, College Station, Tex

Alternate: "An Effective Means for Managing Change: The Modernization Intelligent Decision/Analytical Support System," Larry E. Shipner, CFC, Inc, Dayton, Ohio

Session 20: Shop Floor Control

Chair: Anthony R. Skomra, Automation Technology Products, Campbell, Calif

"Manufacturing Control System: Key to Controlling Manufacturing Performance," Ronald J. Meyer, Rockwell International, Space Operations, Downey, Calif

"Machine Utilization Management System," Ronald E. Funk, Deere & Company, Moline, Ill

"Trams—Tactile Ratio and Monitoring System—The Art of Shop Floor Control," Harry G. Smart, VSI Corp, Culver City, Calif

"Development of an Integrated Planning and Control Model for Discrete Production Systems," Eui H. Park, and M. Fazle Rabbi, North Carolina A & T State Univ, Greensboro, NC

"Workstation Control in a Computer Integrated Manufacturing System," Harry A. Scott, National Bureau of Standards, Washington, DC

Session 21: Scheduling & Material Control

Chair: Leo Roth Klein, Manufacturing Control Systems, Inc, Cleveland, Ohio

"Scheduling the Factory Using Computer Simulation (A New Approach to an Old Problem)," Willard Burge, Eaton Corp, Willoughby Hills, Ohio

"Shop Floor Control: Job Shop Perspective," Steven J. Bowman, Infisy, Inc, Houston, Tex

"J-I-T: Does it Work?," Peter Grieco, Apple Computer, Fremont, Calif

"Paper to be Announced," Joseph Magvor, Computervision, Culver City, Calif

"Paper to be Announced," Ralph Thadeus, Hagar Hinge Co, St Louis, Mo

Alternate: "Paper to be Announced," Theodore Machock, Sargeant Industries, Huntington Park, Calif

Session 22: Flexible Manufacturing Systems—Mechanical

Chair: Gerald Ennis, Director of Industrial Modernization, Vought Corp, Dallas, Tex

"A General Approach to Graphics Animation of Manufacturing Systems," Don N. Pope, LTV, Dallas, Tex

"FMS Technology in General Electric," Steven A. Vogel, General Electric, Charlottesville, Va

"Sheet Metal FMS," James N. Brecker, Production Systems Technology, Pittsburgh, Pa

"Designing Flexible Manufacturing Systems Using Simulation," David B. Wortman, Pritsker & Associates, Inc, West Lafayette, Ind

Tutorials

6:30 to 9:30 pm

Mon & Wed, Oct 1 & 3

CAD/CAM Implementation and Integration
Kenneth C. Bonine, General Dynamics, San Diego, Calif

Mon, Oct 1

Database Management
Dan Appleton, DACOM, Manhattan Beach, Calif

Mon & Wed, Oct 1 & 3

Applied Machine Vision
Perry C. West, Automated Vision Systems, Campbell, Calif

Mon, Oct 1

The Master Production Schedule—Manufacturing's Plan of Action
Leo Roth Klein, Manufacturing Control Systems, Inc, Cleveland, Ohio

Mon & Wed, Oct 1 & 3

Material Handling in the Automated Factory
Jerome L. Huff, Cubic Resources, Inc, Rochester, NY

Tues & Wed, Oct 2 & 3

Getting Your Company Ready for CIM
Roger C. Willis and Michael J. Klich, Arthur Anderson and Co, Chicago, Ill

Tues, Oct 2

Flexibility and Batch Manufacturing on Automatic Assembly Machinery
Frank J. Riley, The Bodine Corporation, Bridgeport, Conn

Tues & Wed, Oct 2 & 3

Introduction to Integrated Automated Shop Floor Control
Thomas V. Sobezak, LPPC, Inc, Baldwin, NY

Tues, Oct 2

Manufacturing Strategy—Integration Mechanism for Creating Competitive Advantage Through Technology Change
Horst J. Metz, Booz, Allen, & Hamilton, Inc, Chicago, Ill

(conference coverage continued on page 263)

Q: WHAT NEW GRAPHICS PRINTER GIVES HIGHEST QUALITY PRINT-OUT, YET REQUIRES NO INTERFACE?

A: AXIOM'S TX1000 THERMAL VIDEOPRINTER.

No Hardware or Software Interface

Now you can have an instant printout of anything displayed on a CRT screen with absolutely no hardware or software interface. That's because our new TX-1000 thermal printer plugs directly into the video connector on most display terminals. You can print complex graphics, alpha-numerics of any size or font, mathematical symbols and even hieroglyphics — whatever is on the screen.

Highest Black-White Contrast

Using 8-1/2-inch wide paper, the TX-1000 turns out crisp, high contrast printing for full size reproduction of forms, reports and graphs. Our blacks are blacker and our whites are whiter. You'll be amazed!

Fast and Quiet

The TX-1000 is fast. Takes just 20 seconds to print a complete CRT screen. Quiet too, making it ideal for office and lab environments.

4 to 1 Zoom

Two simple controls provide complete freedom to zoom, move the image on the paper and vary the aspect ratio to adjust for screen distortion.

High Resolution

Resolution is excellent! Horizontally or vertically switch-selectable to 40, 50, 80 or 160 dots per inch.

No Chemicals or Toners

Because our videoprinter uses advanced thermal technology, messy chemicals and toners are not needed, all of which adds up to less down time and no warm-up time.

Built-in Test Pattern and Signal Diagnosis

The TX-1000 is the easiest-to-use videoprinter on the market. It even has a built-in test pattern and input signal diagnosis to let you know exactly what's going on.

Smaller Model Uses 5-Inch Wide Paper

Our compact EX-855 videoprinter uses 5-inch wide paper — ideal for many logging and instrumentation systems.

Advanced Thermal Printing Technology

Because of it's advanced fixed head thermal technology, the TX-1000 is the finest videoprinter available at any price. However, the cost is low. Just $3395 to the end user.

Axiom printers are backed by distribution and service centers in the USA and 21 foreign countries.

AXIOM

AXIOM CORPORATION

1014 Griswold Avenue • San Fernando, CA 91340
Telephone: (818) 365-9521 • TWX: 910-496-1746

CIRCLE 101

WITH DIGITAL'S SUPER-MICROCOMPUTERS, YOUR POSSIBILITIES ARE ENDLESS.

At Digital, we understand what it takes to be an OEM. And to be a good one. We are, after all, the world's largest supplier to OEMs like you. With more of our product in areas like CAD/CAM/CAE, ATE and imaging than any other manufacturer.

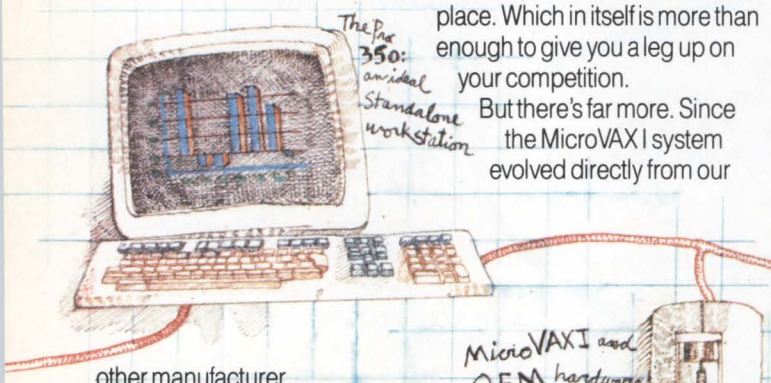

We appreciate your need to bring to market a product that's unique, effective and reliable. And to bring it to market on time.

We're also aware of the stiff challenges you face from your competition. Which means we must constantly be aware of the ways in which we can make you more successful.

We must help you develop the capabilities you need to compete.

We must help you reduce your development time with comprehensive documentation and technical support.

And we must offer service and support after the fact.

In many ways, we must act like a business partner. A partner you can trust for the help you need.

PROVEN SUPERMICRO HARDWARE AND SOFTWARE. TODAY.

While others are merely promising supermicro capabilities, Digital is delivering. Today. Proven technologies. Proven hardware. Proven operating systems. And proven software.

Our MicroVAX I* computer, for instance, is the first true 32-bit supermicro in the marketplace. Which in itself is more than enough to give you a leg up on your competition.

But there's far more. Since the MicroVAX I system evolved directly from our VAX* 32-bit supermini and industry standard Q-Bus* architecture, all supporting hardware and software is in place.

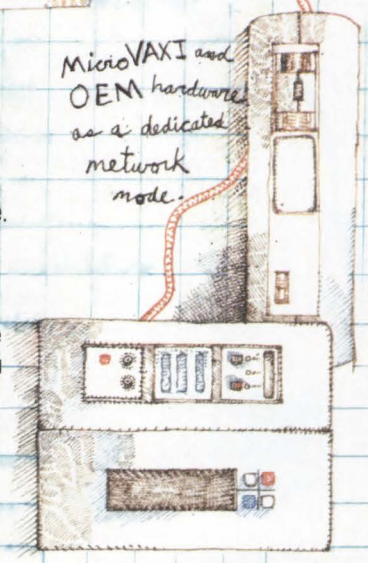

While other micros may give you just one, thinly supported operating environment, Digital's small systems give you two, each optimized for different types of applications.

MicroVMS* is an evolution of our powerful general purpose VMS* realtime operating system repackaged for the micro environment. It not only gives you full virtual memory support, but leadership networking capabilities as well. Plus access to applications and system tools on the leading edge of the computer industry.

For incomparable flexibility, the MicroVMS operating system has been packaged by function into discrete modules. You use just the modules you need for your target application. The core module contains all basic system utilities. You can add programming languages or layered products simply and easily. Because you never use what you don't need, system overhead is minimized.

In addition, MicroVMS will support a range of languages, including FORTRAN, BASIC, DIBOL,* RPG, PASCAL, APL, COBOL and C. Plus most of the industry standard communications languages and accesses.

With all this in its favor, you'll find that the MicroVMS environment provides users with an operating system that's ideal for a broad base of applications, including image processing, numerical control, automatic test equipment, small business computing, expert systems, telecommunications and general purpose computing in a standalone, networked or distributed system.

The VAXELN* operating system is Digital's new 32-bit Pascal-based, realtime executive. VAXELN is a state-of-the-art software environment optimized for realtime control, distributed computing and network-based multiprocessing in applications like CAD/CAM, factory automation, Ethernet communications, laboratory data acquisition and industrial process control. The VAXELN developer has access to all the productivity tools of VMS and more for writing his application. He may trans-

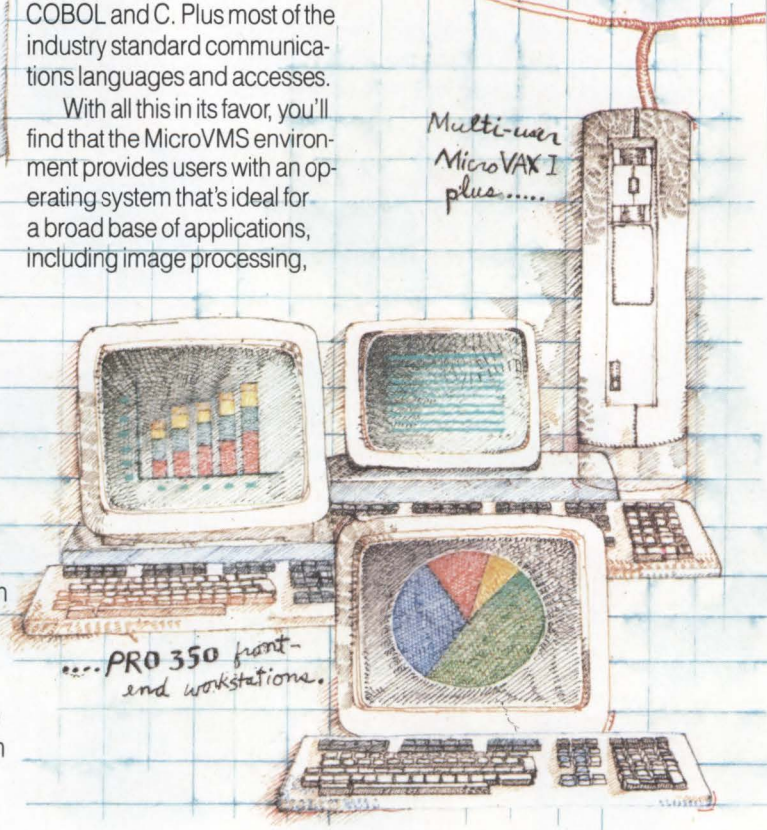

port it to other VAX or MicroVAX I target systems where it may be executed with none of the system overhead of a general purpose operating system.

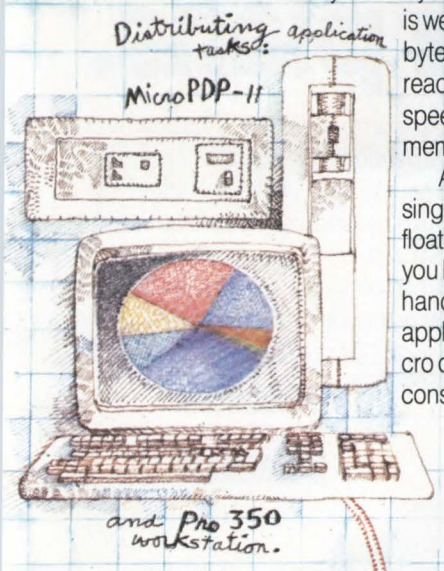

The choice is strictly up to you. For both of these systems offer you a world of existing software and options that have already undergone extensive development and maturation.

Which means you can select your operating environment exactly the way it should be chosen. Based on either system's merits relative to the applications you have in mind.

THE CAPABILITIES TO DISTRIBUTE YOUR PROCESSING. RIGHT NOW.

To all the above capabilities, add the architecture and technology to network your product in a distributed processing environment, and your horizons expand enormously. With Digital, you have the ability – right now – to put together an integrated system of supermicros, minis, and workstations in virtually limitless configurations.

In this arena, three systems are of particular interest. The MicroVAX I system provides true 32-bit processing and multitasking in a supermicro package that can simultaneously support multiple users.

It starts with a physical memory of half a megabyte, with expansion capabilities that can take you all the way to a full 4 megabytes. Total linear address space is well in excess of 4000 megabytes. And, to enhance the already impressive processing speeds, an additional cache memory of 8 Kbytes is provided.

Add to that support for both single and double precision floating point arithmetic, and you have a system capable of handling larger, more technical applications than any other micro or supermicro you might consider. All in less than 6″ of standard rack space.

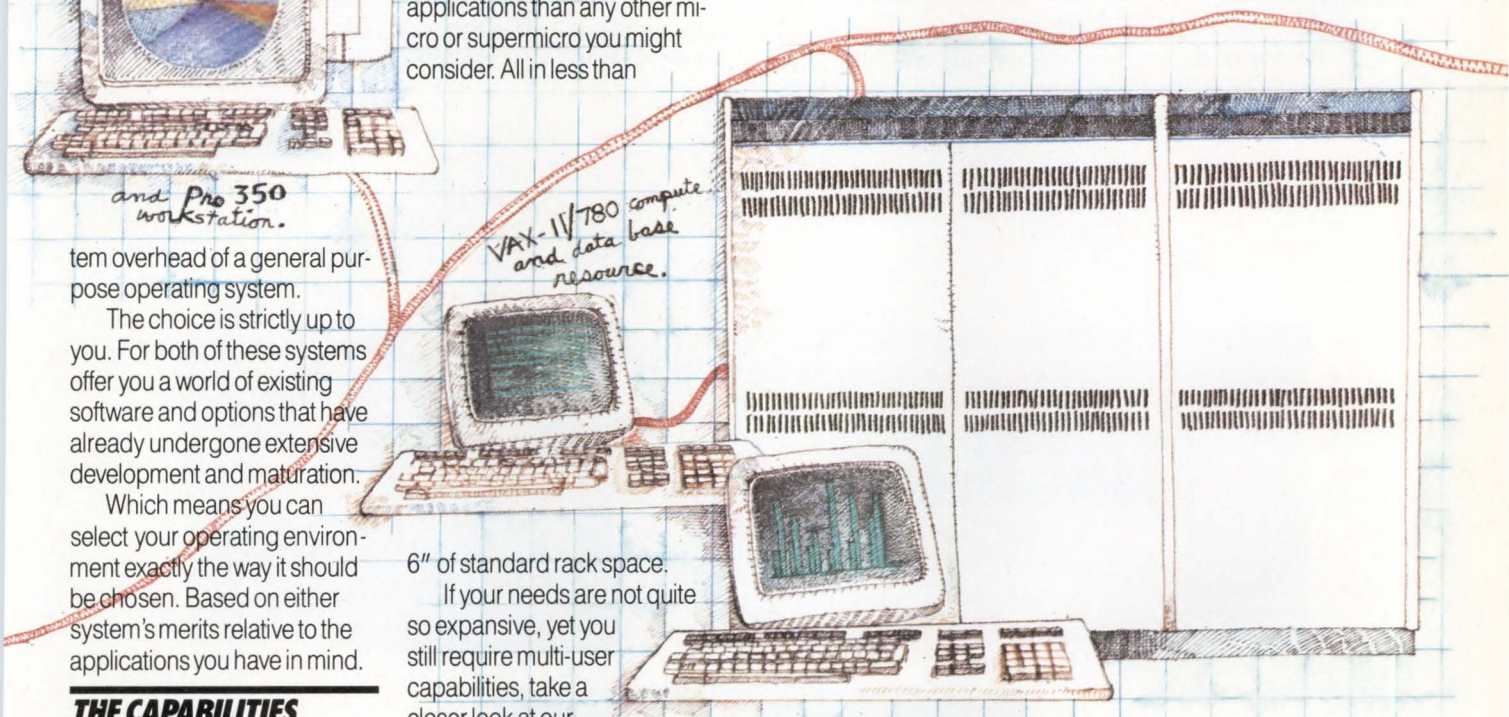

If your needs are not quite so expansive, yet you still require multi-user capabilities, take a closer look at our Micro PDP-11* system. With a choice of 11 or 31MB Winchester disks, floppy disk or streaming tape backup, and a wide choice of operating systems and programming languages, it permits you to use an inexpensive microcomputer for applications that once required a full-blown minicomputer.

For standalone graphics applications, few alternatives can offer you the flexibility of the Pro 350* workstation. Features like true multi-tasking and outstanding graphics capabilities, a choice of rich development environments and compiler languages, plus the ability to integrate into a VAX or PDP-11* network environment are all standard.

If complete systems represent a level of integration higher than you need, you can build your product around any of a complete range of products, too. From Falcon-PLUS,* one of the world's smallest 16-bit single-board computers, all the way down to our range of powerful chip level products.

A COMMITMENT YOU CAN BANK ON.

Perhaps most important of all, every product you make is backed by the world's most successful supplier to OEMs.

Backed with an R&D commitment that exceeded $400 million last year alone. With a national network that allows you to select the level of service and support that suits your needs. And with a name and reputation that's both recognized and respected by your customers.

BEST ENGINEERED MEANS ENGINEERED TO A PLAN.

Digital's advanced OEM products, like all Digital hardware and software products, are engineered to conform to an overall computing strategy. This means that our systems are engineered to work together easily and expand economically. Only Digital provides you with a single, integrated computing strategy from chips to 32-bit VAX systems, and direct from desktop to data center.

For more information about Digital's OEM products, you've got plenty of choices. Fill out and return the coupon below. Or contact one of Digital's Authorized Industrial Distributors. Or call 1-800-848-4400, ext. 139. In Canada, call 617-542-6283.

☐ Please send detailed product specifications.
☐ I'm in a hurry. Have your representative call today.

Name
Title
Company
Address
City
State Zip
Telephone Ext.

Digital Equipment Corporation, 77 Reed Road, HL02-1/E10, Hudson, MA 01749.

THE BEST ENGINEERED COMPUTERS IN THE WORLD.

digital™

CIRCLE 102

Fujitsu modems have lots of ways to fit into your plans.

At Fujitsu, we offer a diverse line of OEM high-speed stand-alone modems and boards to bring out the best in your product designs. Now you can take advantage of the industry's most technologically advanced features. Like operating speeds up to 14,400 bits per second. Like powerful automatic equalizers that minimize the effects of line interference. And versatility that makes our modems compatible with any type of equipment.

As one of the world's largest data communications suppliers, we know that our growth depends upon how well our products help your company grow. That's why we continue to develop OEM products on the cutting edge of technology. And every product we make is backed by the Fujitsu name – recognized worldwide as a leader in data communications products.

For more information, contact us today. We'll be delighted to show you how our products fit into your plans.

Fujtsu America, Inc. • Data Communications Division • 3055 Orchard Drive San Jose, California 95134 • (408) 946-8777

FUJITSU

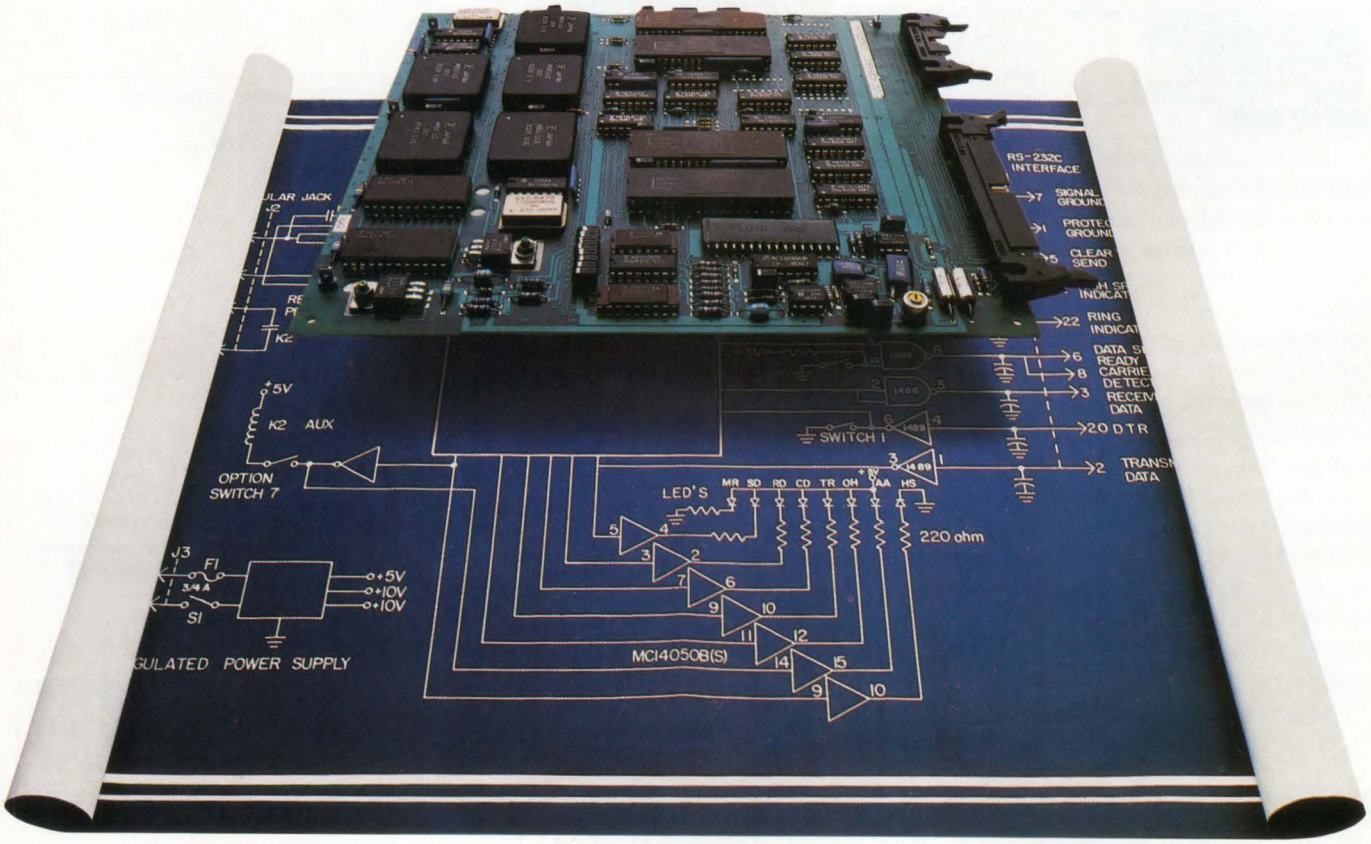

CIRCLE 103

ISA/84 International Conference & Exhibit

Astrohall, Houston, Texas **October 22 to 25**

For the second consecutive year, Houston, Texas will host the annual conference and exhibit of the Instrument Society of America, once again in that city's famed Astrohall. As Max J. Kopp, president of ISA, stated in his message to prospective attendees, ISA/84 will likely be the most prestigious instrumentation event of its kind in this hemisphere.

Consider, however, that ISA/84 is a broad instrumentation conference, covering all areas that pertain to this phase of technology. The program, made up of 92 technical sessions (116 technical papers, 8 clinic sessions, 14 panel sessions, 16 tutorial sessions, and 4 round tables), covers such a broad area that only portions will be of value to any specific attendee.

The following discussion, and the listing of technical sessions, are therefore limited to those areas likely to be most important to digital system designers and integrators. All of the sessions included in this discussion are portions of 5 of the 17 program divisions: automatic control systems, computer technology, electrooptics, telemetry and communications, and education.

Even in this abbreviated agenda, many of the sessions run concurrently. Attendees will have to determine which sessions are likely to be of most value to them—and, of course, move to another if the first does not live up to its promise. Hopefully, this preview will permit those choices to be educated ones.

In the Monday, 10:15 am to 12:15 pm time spot will be two sessions of particular importance, both in computer technology. In one of these, several speakers—all from industry—will discuss implications of the man/machine interface, always a subject of keen interest because of human reluctance to automate. The second is a panel session. Here, seven panelists will present an update of computer network activities within ISA itself.

Two additional computer technology sessions will offer prospective value on Tuesday. One, from 8:30 to 10:30 am, will feature the impact of networks and gateways on plant operations. Speakers will relate experiences at four different industrial companies. During the 10:45 am to 12:45 pm period, panelists will discuss environmental effects on electronic process control equipment.

Wednesday will offer many more sessions of interest to the digital designer. There will be four in computer technology, one in electrooptic sensors, and three in telemetry and communications. Concurrent sessions (8:30 to 10:30 am) in computer technology will consist of selection and management of digital control systems, with four speakers; a panel/forum of personal computer vendors; a basic level clinic on programmable controllers; and another clinic on programming measurement and control systems in Basic.

Other Wednesday sessions of interest in this same time spot include one on electrooptic sensors, and two in the telemetry and communications grouping. The latter will be a technical review of optical communication applications—basically a discussion of fiber optics—and a panel on future trends of artificial intelligence in instrumentation and control.

This last session, and one that immediately follows it, should offer a great deal of valuable information to system designers and those other attendees interested in the future of control technology. The 8:30 am panel on future trends in AI will be followed at 10:45 by a technical session on current trends in AI and expert systems. In the latter session, speakers from Ford Aerospace, Lockheed, and NASA will discuss their experiences, particularly with the integration of AI and data analysis, and with the use of expert systems for the space shuttle's onboard navigation systems.

Thursday will be another day of decision making for ISA/84 digital design engineers. The 8:30 to 10:30 am time space will have six sessions of relative importance. Two will be tutorials, one on the promise and practice of future factory automation; the other on computer aided design. Another pair of sessions will be panels on methods and applications of adaptive control, and on engineering workstations. The final two sessions will be clinics involving three very different uses of the PC acronym. The first session will be an advanced course on programmable controllers; while the second will be on applications of personal computers in process control.

Only three sessions will require participant choices on Thursday from 10:45 am to 12:45 pm. The selection will consist of a technical session on applications of programmable controllers; a technical session on the future of process control, emphasizing the importance of expert systems and AI; and a panel on what is new in microprocessors. This last session will include speakers from key microprocessor manufacturers such as Motorola, Intel, Texas Instruments, and Zilog—*S.F.S*

For further information concerning ISA/84, contact the Instrument Society of America, 67 Alexander Dr, Research Triangle Pk, NC 27709. Tel: 919/549-8411

(continued on page 264)

(continued from page 263)

Technical Program Excerpts*

Mon, Oct 22

10:15 am, Room 111

Session 9: Man/Machine Interface

Session developer: J. Pilson, MicroSwitch, A Honeywell Div

"The Implications of Response Time on Control System Design," R. Blue, Consultant

"Rapid Prototyping for the Human Processing Interface," A. Nagesh, R. Shirley, M. Gordon, Foxboro Co

"Software-controllable Multifunction Controls and Displays," D. Nicholson, MicroSwitch, A Honeywell Div

10:15 am, Room 201

Session 10: Panel: An Update of Computer Network Activities in ISA

Session developer: H. Zinschlag, Monsanto

Panel members: R. Blue, E. Comp; S. Childress, VEPCO; T. Harrison, IBM Corp; C. McAlister, ISA; M. Stanley, EG&G; H. Zinschlag, Monsanto

Tues, Oct 23

8:30 am, Room 111

Session 22: The Impact of Networks and Gateways on Plant Operations

Session developer: M. Greaves, Foxboro Co (COMPUTEC)

"The Management Style of Network," M. Beaverstock, Foxboro Co

"An Inventory of Gateway Design Concepts," D. James, System Application

"Standard Networking Protocols and Gateways for Industrial Applications," R. Jones, Allen Bradley

"Programmable I/O System Prevents Communication Bottlenecks," J. Holtom, Control Microsystems

8:30 am, Room 120

Session 27: Tutorial: Artificial Intelligence Applications in Instrumentation and Control

Session developer: R. Ragaram, Univ of Houston, (TELCOM)

10:45 am, Room 201

Session 35: Panel: Environmental Effects on Electronic Process Control Equipment

Session developer: D. Caro, Autech, Inc (COMPUTEC)

Panel members: R. Harold, SSOE Inc; T. O'Shea, Union Carbide Corp; J. Dowey, Akron Standard Div

Wed, Oct 24

8:30 am, Room 111

Session 45: Selection and Management of Digital Control Systems

Session developer: T. Stanley, Daniel Construction Co (COMPUTEC)

"Advanced Control Techniques and Tricks," R. DiBiana, SOHIO

"Expanding the Use of Host Computers on Distributed Control Systems," W. Barnett, Daniel Construction Co

"The PC-based Distributed Control System," A. Miller, Process Control Industries, Inc

"Local Area Networks: A Cost Analysis," J. Purvis, III, Texas Instruments; I. Erickson, Motorola

8:30 am, Ballroom D

Session 46: Panel: Personal Computer Vendor Forum

Session developer: K. Hopkins, Hopkins Computing (COMPUTEC)

Panel Members: R. Wisherman, Kinetic Systems; M. Manoff, Heuristics; L. Dunham, Taurus Computer Products; B. Rodgers, Metrabyte Corp; E. Basch, Foxboro Co; B. O'Brien, Data Translation

8:30 am, Ballroom C

Session 47: Clinic: Programmable Controllers—Basic Course

Session developer: R. Ralston, Modicon, div of Gould, Inc (COMPUTEC)

8:30 am, Room 201

Session 48: Clinic: Programming Measuring and Control Systems in Basic

Session developer: A. Roth, Ithaco, Inc (COMPUTEC)

8:30 am, Room 123

Session 49: Electrooptic Sensors

Session developer: D. Krohn, Eotec (Electrooptics)

"Optical Gaging System," I. Elrom, Medar, Inc

"Photoelectric vs Fiber Optic Sensors," N. Damik, Veeder Root

"Distributed Fiber Optic Sensors," S. Kingsley, Battelle Memorial Institute

"Fiber Optic Displacement Sensors," D. Krohn, Eotec

*Based on information available at press time. Subject to change.

(continued on page 266)

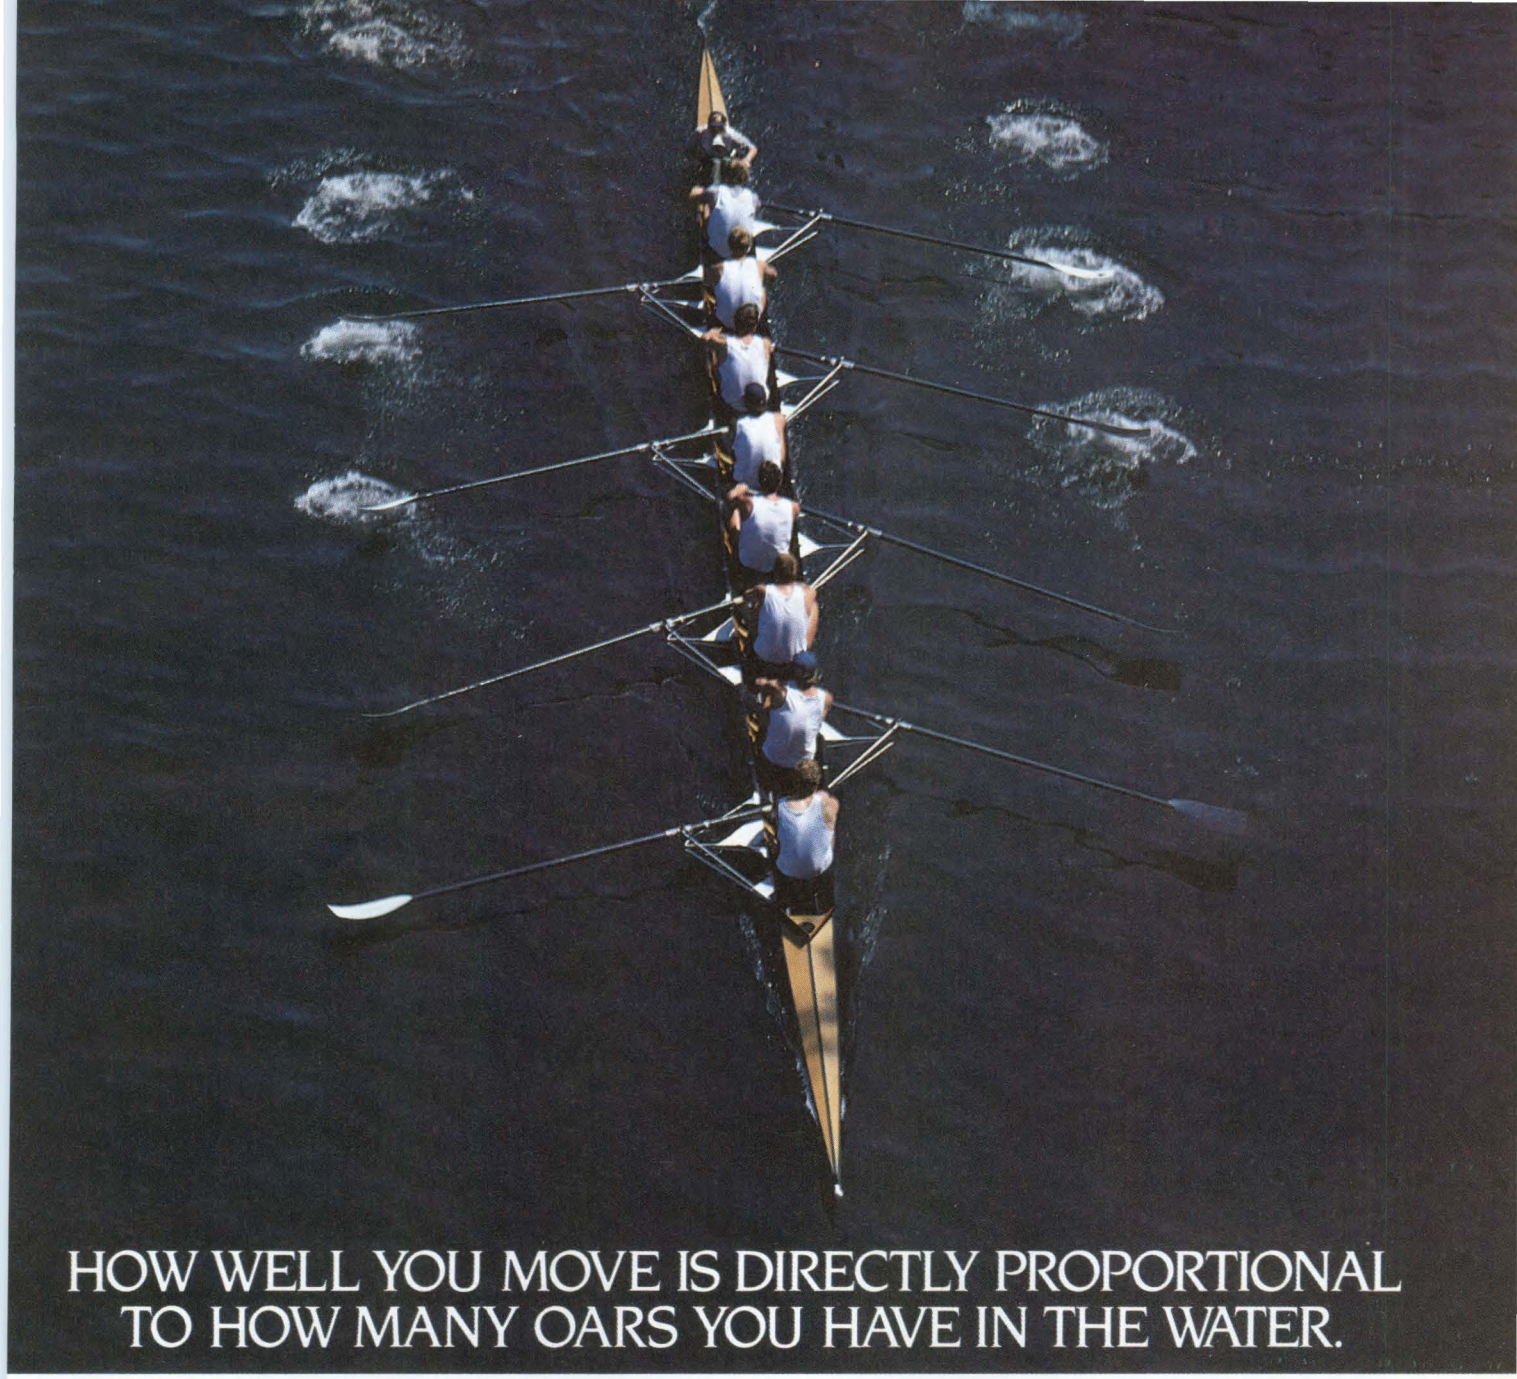

HOW WELL YOU MOVE IS DIRECTLY PROPORTIONAL TO HOW MANY OARS YOU HAVE IN THE WATER.

YOU NEED A BROAD PRODUCT LINE TO PULL AHEAD.

We Give You A Choice Of Seven. Designing information systems for all different kinds of businesses demands flexibility. Lots of it.

And you just can't have flexibility without choice.

Which is primarily why Okidata gives you the broadest line of OEM printers to choose from.

Seven models in all, with responding speeds from 120 to 350 characters per second... unlimited compatibility... up to three different print modes... black and white or color. And on and on. All the options you need to match the printer to the application, not the application to the printer.

A Broad Line Is Only The Beginning. But Okidata's flexibility doesn't stop there. In fact, you've only rippled the surface. There's a highly qualified technical staff to quickly make the alterations that meet your special OEM requirements. There's software applications assistance, hands-on training support, a nationwide network of service people, and parts already stocked, already in-place, that can get to your place real fast.

An Invitation To Pull Together. And Win. If you're putting together systems for all different kinds of businesses, it might make sense to choose the printer that pulls for you in all the important directions. The one that already knows the way. Call us at 1-800-OKIDATA. Or write Okidata, Mt. Laurel, NJ 08054.

OKIDATA

an OKI AMERICA company

CIRCLE 104

(continued from page 264)

8:30 am, Ballroom A

Session 52: Applications of Optical Communications

Session developer: Z. Taqvi, Lockheed EMSCO (TELCOM)

"New Field Instrumentation System Employing Optical Fibers," T. Yashuara, M. Tamai, Fuji Electric Co

"Highly Reliable, High Speed, Multipurpose Optical Fiber Communication Data Network," K. Inoh, Yokogawa Hokushin, Electric Corp; H. Wilson, Yokogawa Corp of America

"Fiber Optics in the Plant or Mill," R. Caro, Autech Data Systems, Inc

8:30 am, Room 124

Session 53: Panel: Future Trends of Artificial Intelligence in Instrumentation and Control

Session developer: R. Brown, NASA/JSC (TELCOM)

10:45 am, Room 124

Session 68: Current Trends in Artificial Intelligence/Expert Systems

Session developer: Z. Taqvi, Lockheed EMSCO (TELCOM)

"Heuristic Model for Resource Allocation in Contention Environment," J. King, Ford Aerospace and Communications Corp

"Automatic Human Monitoring and Control: Integrating Artificial Intelligence and Data Analysis," J. Malin, Lockheed

"Expert Systems for Monitoring Shuttle Onboard Navigation Systems," L. Ang, NASA/JSC

Thurs, Oct 25

8:30 am, Room 118

Session 71: Tutorial: Future Factory Automation—Promise and Practice

Session developer: S. Bansal, Polaroid (ACOS)

8:30 am, Room 107

Session 72: Panel: Adaptive Control—Methods and Applications

Session developer: E. Bristol, Foxboro Co (ACOS)

Panel members: K. Astrom, Lund Institute; T. Kraus, Foxboro Co, M. Herbignat, ASEA, T. Myron, Foxboro Co

8:30 am, Room 201

Session 73: Panel: Engineering Workstations

Session developer: L. Barnstone, Exxon Research & Engineering Co (COMPUTEC)

Panel members: P. Bur, Honeywell Inc; P. Bowlds, IBM Corp; M. Beaverstock, Foxboro Co; J. Fisher, Du Pont Co

8:30 am, Ballroom C

Session 74: Clinic: Programmable Controllers—Advanced Course

Session developer: R. Ralston, Modicon, div of Gould, Inc (COMPUTEC)

8:30 am, Ballroom D

Session 75: Clinic: Personal Computers in Process Control

Session developer: W. O'Brien, Data Translation (COMPUTEC)

8:30 am, Ballroom B

Session 79: Tutorial: Computer Aided Design

Session developer: A. Olinger, Blue Technologies (Education)

10:45 am, Room 122

Session 85: Programmable Controller Applications

Session developer: S. Park, Celanese Chemical Co (ACOS)

"Programmable Controllers: Industrial-strength Aspirin for America's Ailing Infrastructure," R. Collins, General Electric Co

"An Advanced Application of Programmable Controllers to a Process Industry," D. Faber, D.M. International

"Batch Reactor Control Simplified by an Integrated Controller," G. Moyer, Honeywell, Inc

10:45 am, Room 124

Session 86: Beyond Process Control

Session developer: J. Rovnak, Stone & Webster (COMPUTEC)

"Expert Systems Applications in Industry," R. Moore, LISP Machine, Inc

"A Generalized Controller for Factory Automation," W. Carter, Automation Intelligence, Inc

"Turnkey Systems for Progressive Batch PCB Assembly and Inspection," S. Jones, Control Automation, Inc

10:45 am, Room 107

Session 87: Panel: What's New in Microprocessors?

Session developer: J. Purvis, National Instrument (COMPUTEC)

Panel Members: I. Erickson, Motorola; J. Parsons. Intel; E. Powell, Texas Instruments; D. Phillips, Zilos

(conference coverage continued on page 273)

Vectrix has the graphics system for your next system. Already.

Image created on VX384A which will work with any computer.

For whatever system you're designing there's already an advanced color graphics system to go with it. Either the economical, easy to use VX128A or the versatile VX384A with its 16.8 million color palette and 512 simultaneously viewable colors.

When you add one of the Vectrix Systems to your system you turn your computer into a powerful graphics workstation with a fully integrated graphic command processor, a bit mapped graphics frame buffer, and a 672 x 480 pixel high resolution flicker-free monitor. Both systems deliver high quality, high performance color graphics. Because the simple command set works with any computer and any programming language, Vectrix makes any system a valuable tool for CAD/CAM, business graphics, medical imaging, video processing, computer assisted instruction, control systems, weather displays, or any of a host of graphics applications in demand.

The Vectrix VX128A and VX384A, part of an expanding line of graphics products that do the job better, including the Midas System for the IBM PC XT and the Vectrix Paint Program.

If graphics should be a part of your system, Vectrix should be a part of your plan. Let us tell you just how we fit your system. Vectrix Corporation, 2606 Branchwood Drive, Greensboro, North Carolina 27408. Phone (919) 288-0520, Telex 574417.

Vectrix

THE COMPUTER GRAPHICS COMPANY

CIRCLE 105

Introducing the newest member of

The HP 1630G...65 channels and maximize 16-bit system

First came the HP 1630A and the HP 1630D. One, a low-cost general purpose logic analyzer suited to the needs of the full development cycle. The other, with 16 channels of timing analysis and 27 of state, an invaluable tool for the hardware design engineer. Now Hewlett-Packard introduces the HP 1630G. With up to 65 channels of state analysis, it is the new standard for software design engineers working on complex new 16-bit microprocessor-based products. Plus, the ability to configure 8 of those lines for 100 MHz timing analysis gives you a logic analyzer system with investigative power and versatility for virtually all your needs.

Three new software overview modes let you nonintrusively monitor software performance and hardware/software interactions in real time.

The HP 1630G significantly expands on the software performance capabilities already introduced in the HP 1630A/D family members. In addition to time histograms that show execution-time distribution, and label histograms that show address activity, the HP 1630G gives you three new modes: program flow, time positional, and linkage measurements. *Program flow measurement* lets you monitor program activity based strictly on opcode accesses. This can help resolve confusions which may occur when histogramming by address, especially if the program generates inline code or if memory blocks are interspersed between sections of the program. *Time positional measurement* lets you measure the number of occurrences of an event per unit time. Use this to better understand the behavior of your system under a time-varying load. *Linkage measurement* measures the relative frequency of the activity between specific software modules. You'll find this mode invaluable when you want to monitor the transfer activity between a main program and a series of subroutines, for example. Take advantage of all these software performance analysis modes to rapidly

HP-IL

DESIGNED FOR HP-IB SYSTEMS

HP-IB: Not just IEEE-488, but the hardware, documentation and support that delivers the shortest path to a measurement system.

HP 1630D

HP 1630A

HP's logic analyzer family.

advanced software analysis help you performance.

discover if your system is resource-limited or if, for example, poorly chosen program segmentation is causing too much time-consuming disc-to-memory swapping.

Time tagging gives you added insights into system functions.

In the state analysis mode, time tagging measures the time elapsed between each stored state. Make detailed absolute time measurements between states and known physical events. Or, use the mode to measure the total time from the trigger point to a particular state. Because time tagging is a single-pass activity, it is well suited to helping you identify inline sections of code that take longer to execute than anticipated.

Floppy disc interface and popular 16-bit microprocessor support.

On-board non-volatile memory keeps one instrument setup and your current disassembler instantly available at power-up. For even greater storage, the HP 1630G features direct compatibility with a number of HP disc drives such as the HP 9121S/D (the HP 9121D is illustrated). In one convenient 3 1/2" floppy you can now store data, state listings, timing diagrams, alternate disassemblers, and instrument setup configurations. For added flexibility, the HP 1630G supports all popular 8-bit, as well as the following 16-bit microprocessors: 68000, 8086, 8088, 80186, 80286, Z8001 and Z8002.

Our HP 1630G upgrade kit protects your previous HP investment.

If you've already invested in an HP 1630A or HP 1630D, but you feel you need the added capabilities of the HP 1630G, you'll be glad to know that an upgrade kit is available.

Compare the HP 1630G. At $12,100* it's even HP-IB programmable for fully automated measurements. To find out more about the HP 1630G or its companions the HP 1630A/D, call your local HP sales office listed in the telephone directory white pages. Ask for the electronic instruments department.

* U.S.A. list price only.

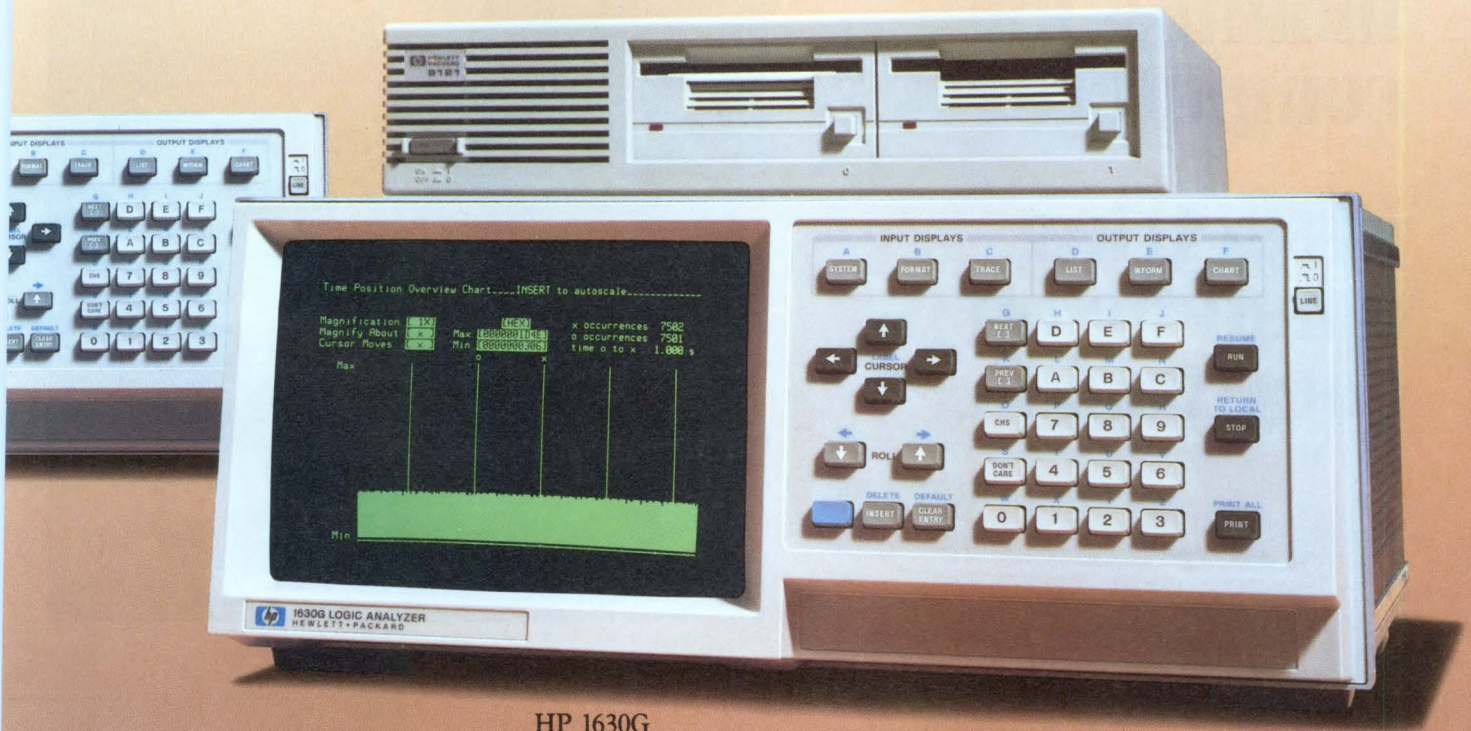

HP 1630G

HEWLETT PACKARD

CIRCLE 106

0801403

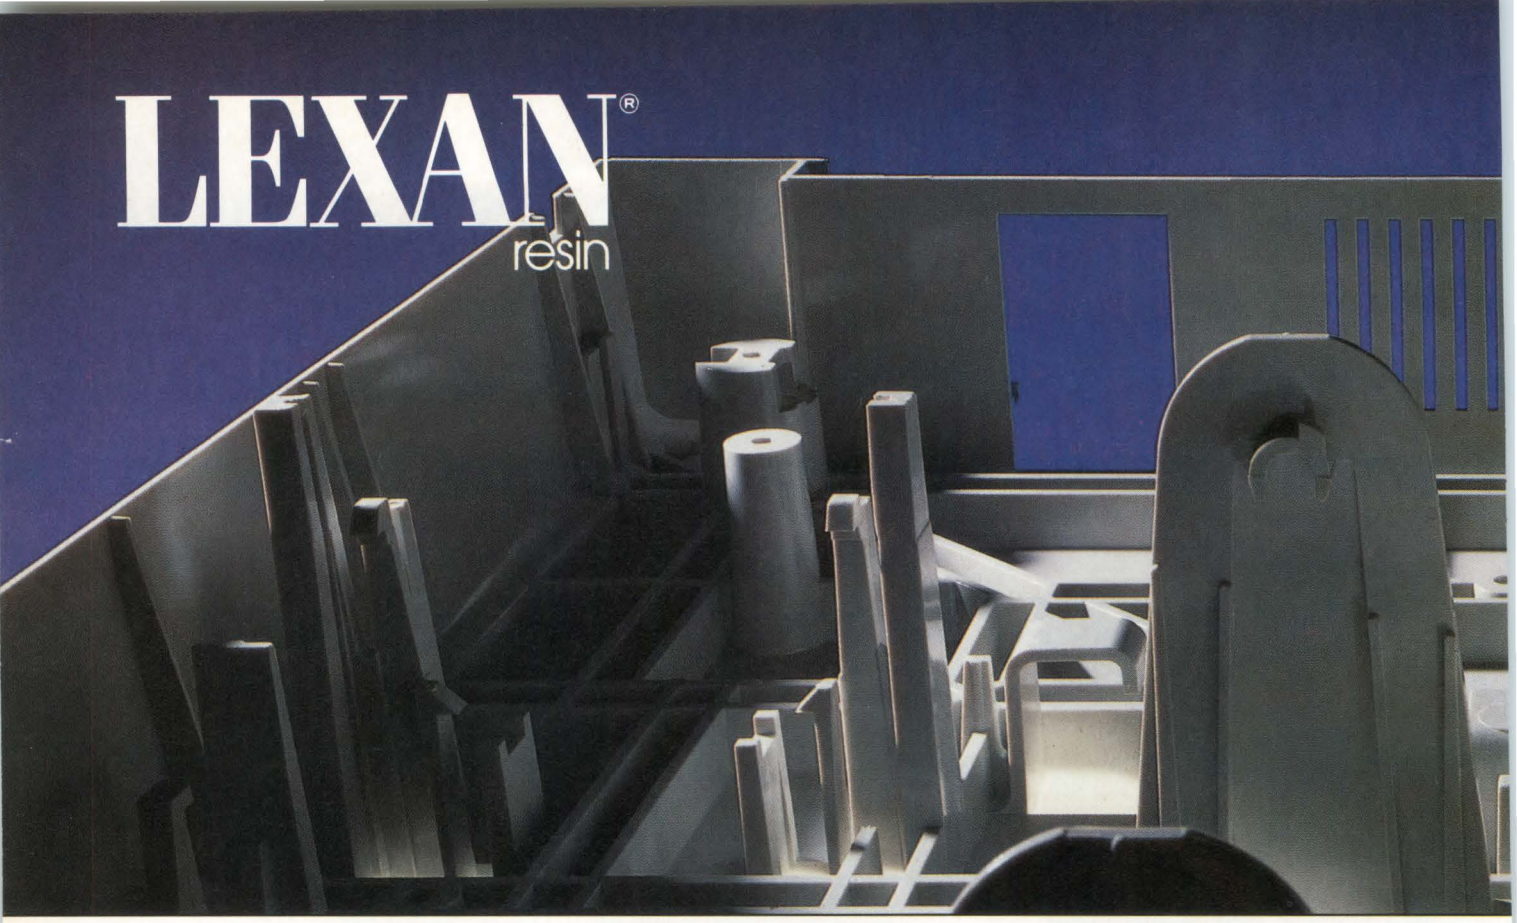

It's precision molding for computers.

It's getting a handle on coextrusion.

GE LEXAN® resin

® Registered Trademark of General Electric Company.

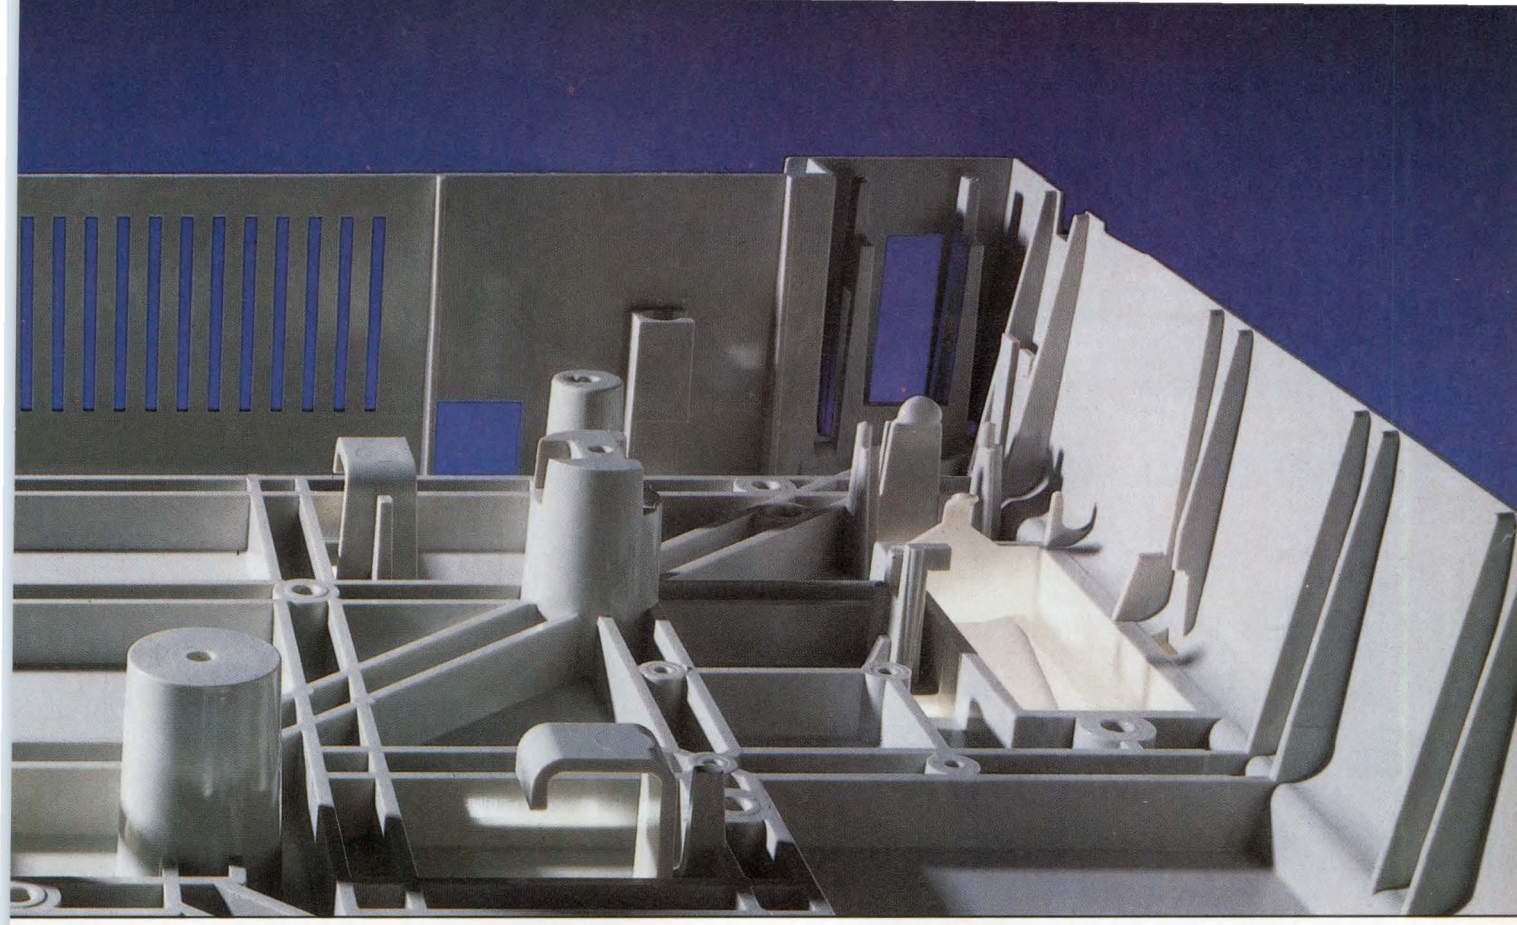

It's whatever you need it to be.

Processability. LEXAN resin means more than superior products. It also means processability to ensure that the finished part will perform with all its properties intact—from intricate, thin-wall computer bases to sleek, multilayer packaging.

Flexibility. With a processing window from 350° F to over 700° F, LEXAN resin offers exceptional stability and melt strength over a wide processing temperature range. The result is flexibility in molding parameters and unparalleled design latitude.

Coextrusion. Just look at LEXAN resin's performance in the hottest new processing technology. Only LEXAN resin's high melt strength meets challenges like clear, handled, multilayered containers. Combined with a barrier resin, LEXAN resin's high HDT offers hot-fill and retort capability, plus long shelf life.

Technical Support. The LEXAN product is backed by unparalleled technical support, from mold flow/mold cool analysis to trouble-shooting on your machinery or ours. General Electric will continue to lead the way in refining processing technology at its new $25-million Plastics Technology Center in Pittsfield, Massachusetts. That's a significant investment in the future—and one no other supplier has matched.

Bottom Line Value. Explore LEXAN resin's design opportunities and processing advantages to lower your finished part cost. Put the world's polycarbonate pioneer to work on your next application, right from the initial design.

For immediate literature, and to receive a free subscription to the new LEXAN **Design Tips,** simply call toll-free (800) 422-1600; in Vermont (802) 442-8356.

CIRCLE 107

General Electric Company
Plastics Group
LEXAN Products Division
One Plastics Avenue
Pittsfield, MA 01201

We bring good things to life.

GENERAL ELECTRIC

The Convergence Factor.

Convergence: the single most critical factor in color CRT performance.

Until now, Delta-gun tubes were the best way to achieve near perfect convergence, but only with costly adjustment electronics. Meanwhile, many in-line tubes are plagued by perceptible misconvergence. Which can lead to poor picture quality. A poor quality image for your product. And poor, bleary-eyed operators.

The Panasonic achievement: low cost in-line color CRTs with better-than-Delta convergence performance.

Without complex adjustment electronics . . . and none of the convergence drift inherent in active correction systems. At last, high resolution in-line tubes with stable performance that stands up to the ravages of time and tough office/industrial environments.

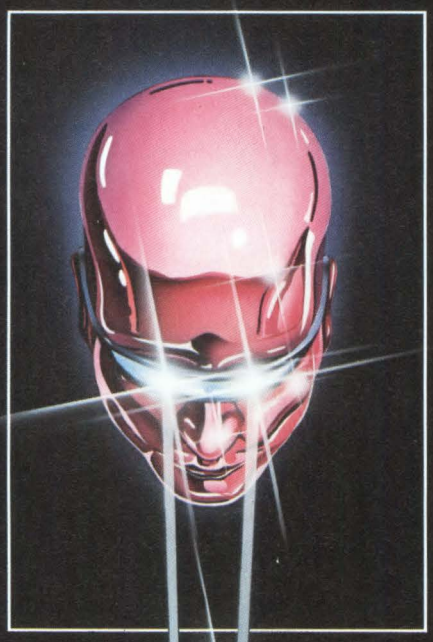

How did we do it? With a preconverged in-line tube/yoke combination unlike any other. Our precision S/ST (saddle/saddle toroidal) deflection yoke is ideally matched to each tube, for near perfect convergence, high repeatability and stability over a wide range of operating conditions.

We combine it with a specially-designed OLF (overlapping field lens) gun and unitized grid construction, providing spot uniformity across the entire screen and near-Delta resolution.

The result: a triumph over the convergence factor. Find out what it can do for your next color terminal or monitor, and ask about our full line of quality color and monochrome CRTs. Write or call: Panasonic Industrial Company, Electronic Components Division, One Panasonic Way, Secaucus, N.J. 07094; (201) 348-5278.

The achievement of Panasonic high resolution in-line color CRTs.

Panasonic
Industrial Company

CIRCLE 108

IECON
IECON
IECON

International Conference on Industrial Electronics, Control, and Instrumentation

Keio Plaza Inter-Continental Hotel, Tokyo, Japan **October 22 to 26**

This year, for the first time in its 10-year history, IECON will be held outside the United States. Although the IEEE's Industrial Electronics Society is traditionally IECON's sole sponsor, the 1984 conference—in Tokyo, Japan—will be cosponsored by the Society of Instrument and Control Engineers of Japan. However, neither the Japanese nor the United States will dominate. Sessions will be chaired and papers presented by key scientists from many other countries, including India, Canada, Republic of China, Greece, Brazil, Egypt, Czechoslovakia, and Germany.

The keynote speaker at Wednesday's luncheon, Paul M. Russo of General Electric's Microelectronics Center, will discuss issues and opportunities of semicustom VLSI. Dr Russo says that he "will review the forces driving the explosive growth of semicustom VLSI devices as well as discuss semicustom markets, design methodologies, and related technologies in the context of turning system designers into chip designers."

As in past years, technical sessions and workshops at IECON will be aimed at practical aspects of process control. Emphasis, of course, will be on the role of digital computers. But considerable stress will also be placed on specific areas such as robotics, machine vision, and sensors. In addition, the plenary session will include invited speeches, at least one of which could be of special interest— "Fifth-Generation Computers," by K. Fuchi, director of Japan's Institute for New Generation Computer Technology. A parallel presentation on artificial intellegence—"Joint Research Program in the U.S.A."—will by givén by K.H. Zaininger, executive vice president of Siemens Corporate Research and Support, Inc.

Computer-controlled systems

Computers, of course, form part of nearly every control system. However, a considerable number of IECON '84 sessions will present specific discussions of computer applications: CAD/CAM, local area networks, data acquisition, and numerical control. Nor will software be ignored: although only one session has "software" in its title, the subject will be discussed in many other sessions.

At the CAD/CAM session on Tuesday afternoon, papers on a variety of applications will be presented. These will include sculptured surface generation, graphical entry and symbolic analysis, microelectronic and VLSI circuit design, and the personal computer. A concurrent session on automotive electronics will be even more specific. It will focus on developments within just that segment of the control industry. Still another session in the same time period will cover the interface of LANs and computers. Both hardware and software aspects are included in those papers.

Two sessions on data acquisition and signal processing will take place—one Wednesday morning, the other Thursday afternoon. Speakers at the first session will discuss conic section graphics, automatic evaluation, an intelligent operator's console, and a high speed image processing system. Papers in the second session will include topics such as pattern recognition, a fast algorithm for realtime formant extraction, and realtime data bases.

Three "control" sessions during the conference will focus on aspects of application. Wednesday morning's discussion will be on numerical control, Wednesday afternoon's will be on computer control, and Thursday morning's will be on process control. The numerical control session will include papers on such subjects as high performance computer numerical control machine tools, an automatic positioning system, computer numerical control for machining sculptured surfaces, an industrial programmable adaptive controller, and an intelligent multiloop digital controller.

(continued on page 274)

(continued from page 273)

Even more emphasis on applications will be found in the session on computer control. Discussion will center on a fault-tolerant traffic controller, an automatic train operation controller, a train traffic control system, and power plant automation. The process control session will have a more theoretical orientation, but will be aimed at specific applications. Papers will be on fault-tolerant programs, automatic diagnosis, automatic width control, hierarchical control, parallel processors, and fault diagnosis.

One Wednesday morning session will specifically target software systems designed for automated manufacturing. Speakers from Japan, the United States, India, Italy, and West Germany will discuss realtime Basic, a compiler for multimicroprocessor control using Forth, and programming tools for the integrated factory. Other speakers will expound upon their findings and experiences on such related subjects as computer aided image generation, analysis and design of sensing devices, and designing products for the factory of the future.

The final session in the general computer category will be on system reliability. Five of the seven papers will be presented by speakers from Japan—most of them from industry. The other two speakers will be from academia: one from Feng-Chia University in the Republic of China, the other from Prairie State College in the United States. Papers will be varied, ranging from error detection and recovery to automatic testing and adjustment. Other subjects will include a reliable multimicroprocessor system, a fault-tolerant system, fault diagnosis, and evaluating system decomposition.

Robotics, machine vision, and sensors

How can you have a control conference without discussions of robotics? Especially when the conference is in Japan? You can't. Eight technical sessions and two workshops will directly involve some phase of robotics, vision, and sensors. (Many other sessions will also contain individual papers on one or more of these subjects.) For example, a Wednesday morning session on control of robots will contain papers on controller design and simulation, on learning control for robotic systems, on control of robotic manipulators, and on precise positioning of robotic arms. Speakers at another session, on Thursday morning, will address a tactile sensing system, an ultrasonic scanner system, offline programming, and line of sight guidance—all as they relate to robotic systems.

Two reports on robotics research are tentatively scheduled for a Wednesday afternoon workshop. One of the reports will be from AT&T Bell Labs in the United States, the other from Electrotechnical Labs in Japan. At press time, neither had been confirmed as part of the program but were considered certain enough to be included in the preliminary program. Other papers at this workshop will be on programming with realtime graphic simulation, bilateral control of manipulation, sensory feedback control, and perception control architecture in image processing.

Discussion of robotics, vision systems, and sensors have been run together in many sessions, which is not surprising since they are so closely related. The first official technical session of the conference, for example, falls under the overall area of robotics. However, the subtitle is vision systems and sensors and all of the papers relate to the subtitle. Beginning immediately after lunch on Tuesday, this session will discuss dot pattern clustering techniques for machine vision, automatic visual measurement, remote position sensors, optical tracking, and image processing for analysis of moving objects.

A Thursday afternoon session relates closely to machine vision, but again emphasizes robotics. Papers will include discussions of realtime image analysis, image processing for inspection systems, automatic inspection, and optical inspection.

The topic of intelligent sensors will be thoroughly examined in three technical sessions and a workshop. Because sensors serve so many purposes, the range of sensor types vary widely—and so will the papers presented at IECON '84. One session will be held Wednesday afternoon, another Thursday morning, and still another Thursday afternoon. Although the title of the related workshop is "Microelectronics for Hazardous and Extreme Conditions," the key papers are on sensors. One is on future technology for military sensors and another is on electromagnetic sensors under adverse conditions. It is likely that other papers in the workshop will also relate to sensors.—*S.F.S.*

For further information on IECON '84, contact F. Jur, c/o Bechtel Corp, 45 Fremont St, MS 45/17A26, San Francisco, CA 94119. Tel: 415/882-1961

Dialight adds the finishing touch.

High quality and reliable performance make Dialight switches, indicators, LEDs and readouts the perfect finishing touch for any product.

We have rockers, toggles and illuminated switches in thousands of different designs. And our incandescent, neon and LED indicators and opto-electronics come in the world's largest selection of colors, shapes and sizes.

Whatever your requirements, Dialight can meet your needs. Let us add the finishing touch that helps your products make a great first impression. Dialight, 203 Harrison Place, Brooklyn, NY, 11237-1587 (718) 497-7600

DIALIGHT
the finishing touch.
A North American Philips Company

T.W. Wilson Sales/Mississauga, Ont. (416) 677-8200 • Arwin Tech. Sales Ltd./ N. Vancouver, B.C. (604) 980-4346

Circle 109 indicator lights
Circle 110 Switches
Circle 111 optoelectronics
Circle 112 have representative call

11,000 GATES

NEVER HAS SO LITTLE MEANT SO MUCH.

In the last decade, we've opened more technological gates than anyone.

In both CMOS and bipolar.

Our 11,000 gate 2 micron CMOS array is just part of the story. We also offer TTL up to 2,000 gates. ECL to 3,000 gates. And 400 to 11,000 gate CMOS arrays. Not to mention our 100K- and 10KH-compatible ECL arrays.

Of course, a gate array is only as good as the people behind it. And at NEC, we have some of the best support engineers. Ready to help you take your design from start to finish.

With easy-to-use CAD tools. And sophisticated software capable of simulating up to 20,000 gates. All available at our 18 customer design centers, worldwide.

As for quality, we offer 100% burn-in, standard. At no extra cost. And you won't have to wait long to get it. Our turnaround is as quick as 8 weeks.

So next time you need a gate array, weigh your alternatives. Then call NEC.

WE'RE TAKING ON THE FUTURE.

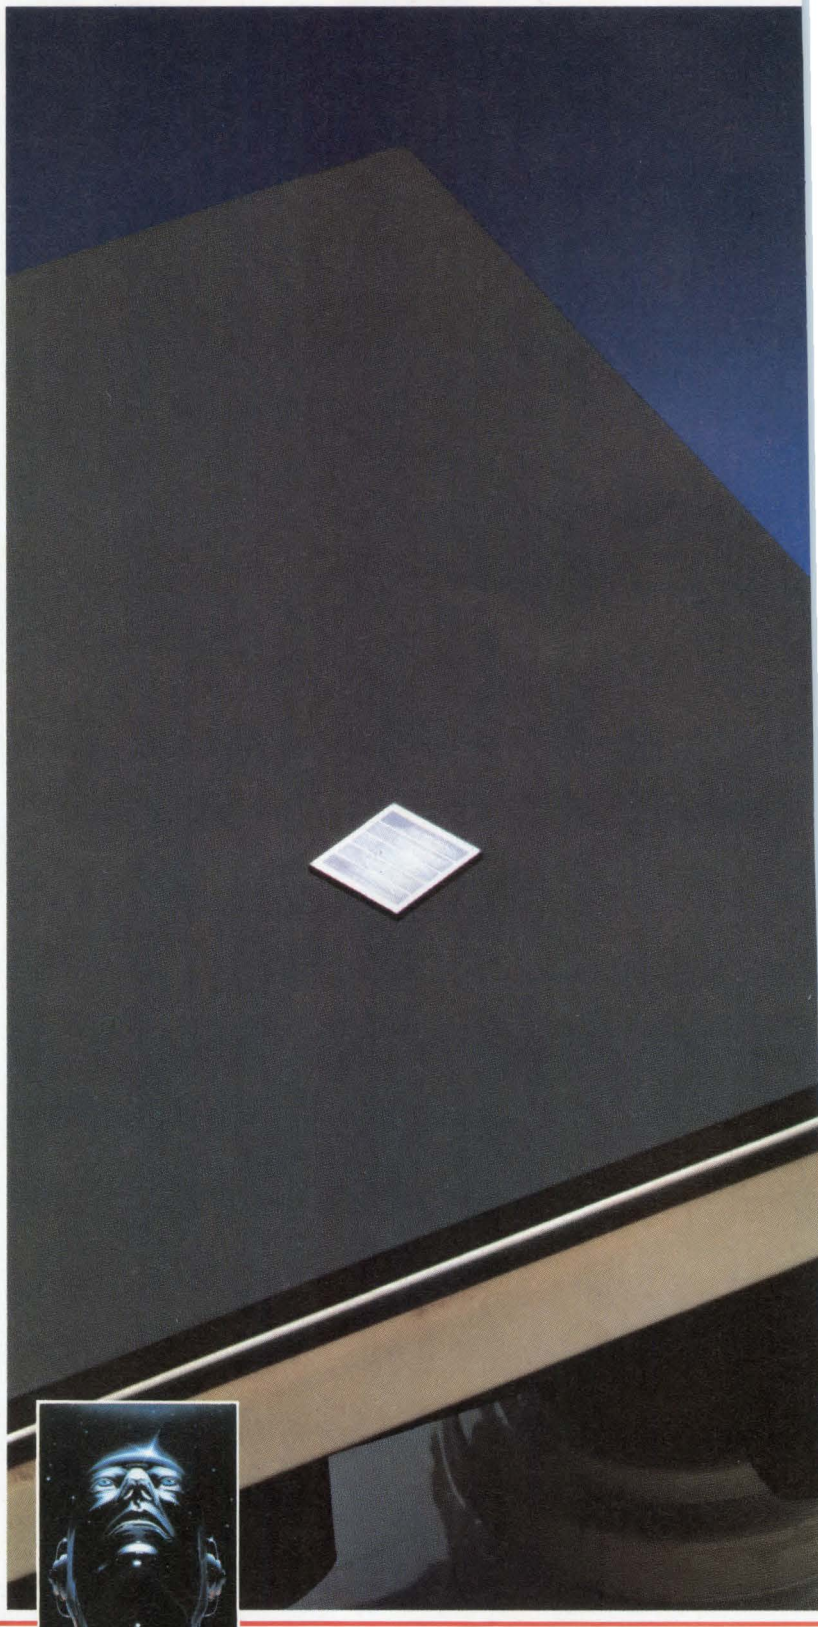

For the fastest response to your questions or for more detailed information, call us TOLL FREE at 1-800-556-1234, ext. 188. In California, call 1-800-441-2345, ext. 188.

NEC national sales offices: Woburn, MA (617) 935-6339 · Melville, NY (516) 423-2500 · Poughkeepsie, NY (914) 452-4747 · Ft. Lauderdale, FL (305) 776-0682 · Norcross, GA (404) 447-4409 · Columbia, MD (301) 730-8600 · Arlington Heights, IL (312) 577-9090 · Southfield, MI (313) 559-4242 · Bloomington, MN (612) 854-4443 · Austin, TX (512) 346-9280 · Dallas, TX (214) 931-0641 · Orange, CA (714) 937-5244 · Cupertino, CA (408) 446-0650

PER GRAM.

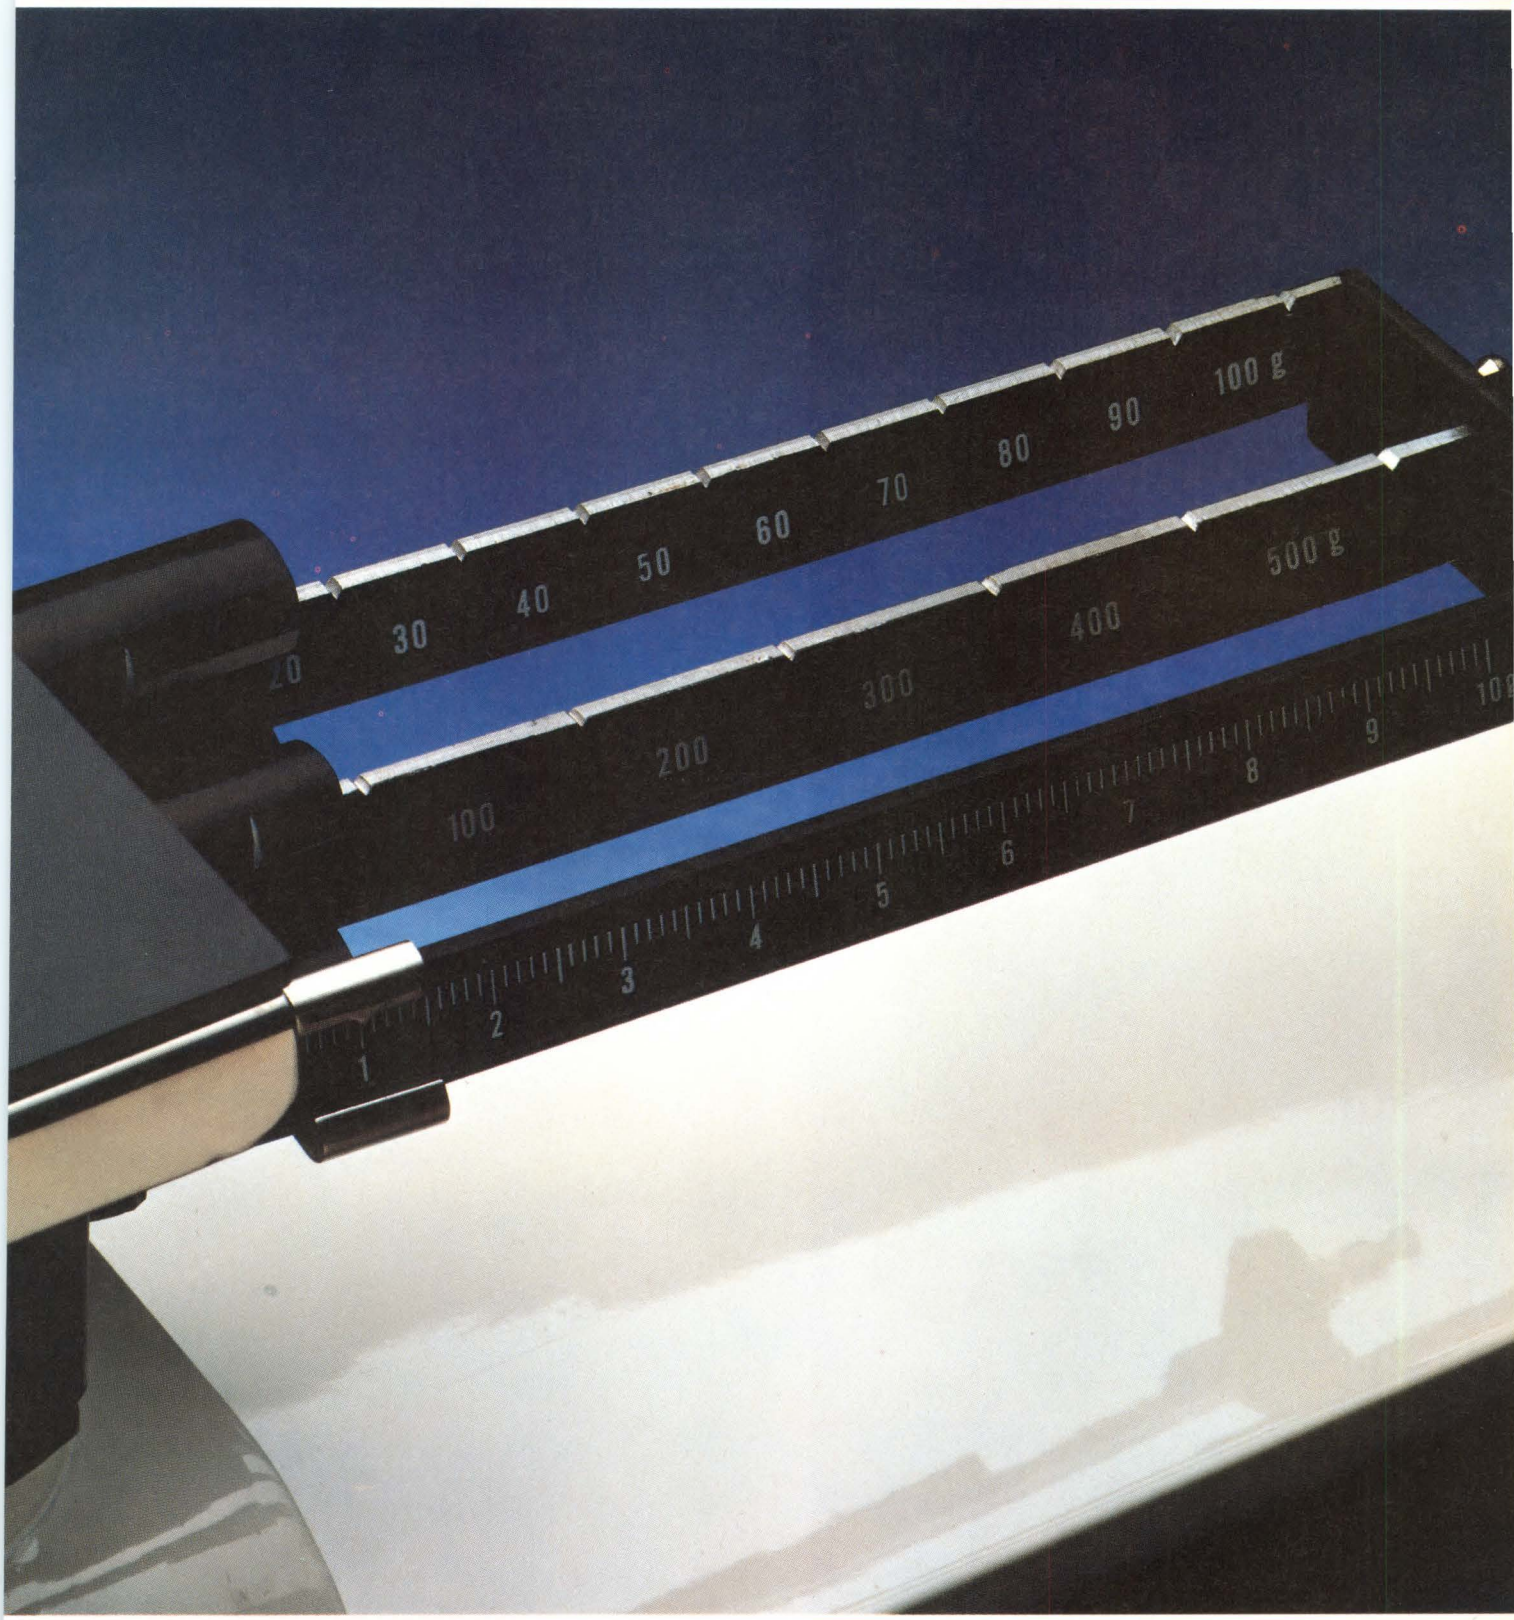

NEC
NEC Electronics

CIRCLE 113

"Standard-cell but I don't know

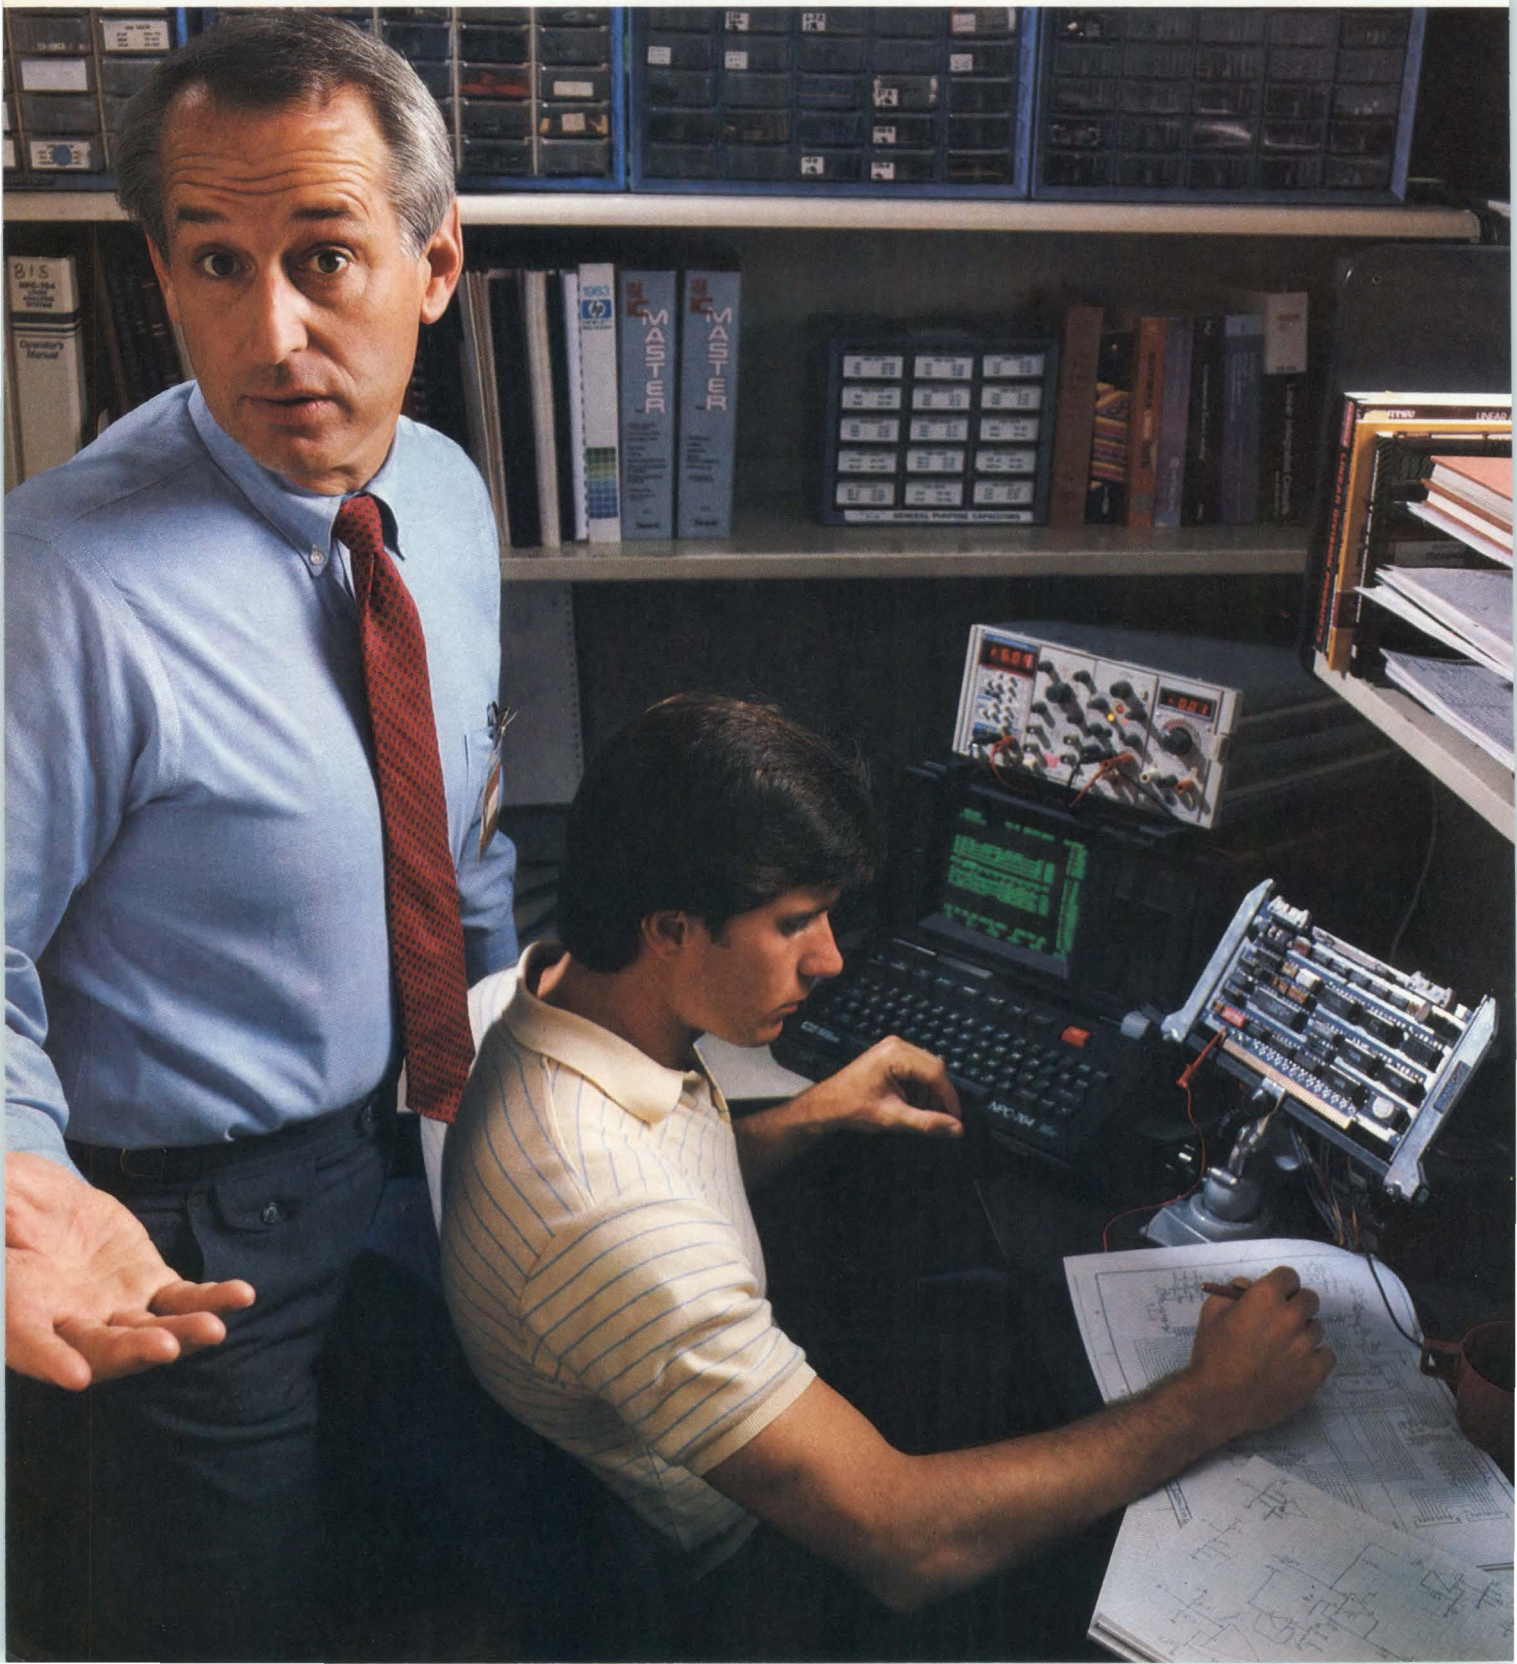

ICs sound good, how to design them."

With Texas Instruments, you do.

The TTL design techniques you are thoroughly familiar with are the basis for TI's Standard Cell Integration System. Put simply, designing with TI's standard cells is much like arranging standard TTL logic packages on a printed-wiring board. And you not only have the standard digital functions to design with but also RAM, ROM, PLA, ALU, and linear functions.

Using TI's 3-μm CMOS technology requires only minimum effort on your part to achieve an affordable semicustom chip.

Getting started is easy

You can get your design started by bringing us your hand-drawn schematic or, better yet, a schematic generated by an engineering work station. When you use a work station, TI will supply the cell library and translation software for automatic generation of the TI design-description language. TI's standard-cell library is supported on many industry-standard work stations, including Daisy, Mentor, and Valid, as well as on several PC systems.

We deliver to specs

TI's Standard Cell Integration System, with its systematic step-by-step procedure, gives your design its best assurance of first-pass success. TI commits to deliver chips that perform to your specifications.

After your logic design is complete, simulation, testability analysis, test generation, and test-pattern grading are performed either by you or by TI. And then double-checked by TI.

Next, TI will develop a computer-aided layout of your design which is thoroughly verified for accuracy. Or you can do the chip layout, and TI will supply the specifications you need.

A final simulation becomes, upon your approval, the conformance specification against which prototypes are produced and measured. Prototypes are then shipped to you for in-system evaluation and a "go" decision.

Your Regional Technology Center is nearby to help

All of the IC design resources and expertise of Texas Instruments are available to you through your nearest Regional Technology Center.

Each Center is staffed with a team of standard-cell specialists who will work closely with you to meet your design goals and product schedule.

Your Regional Technology Center can provide quick answers to technical questions as well as provide training in standard-cell design procedures. Here you will receive assistance in design optimization and test development as well as in creating the data-base description of your chip design. Here are engineering work stations for your use and the software tools and the computers needed to complete the logic design and perform simulation and verification.

State-of-the-art design tools speed TI standard-cell design process, help minimize development costs.

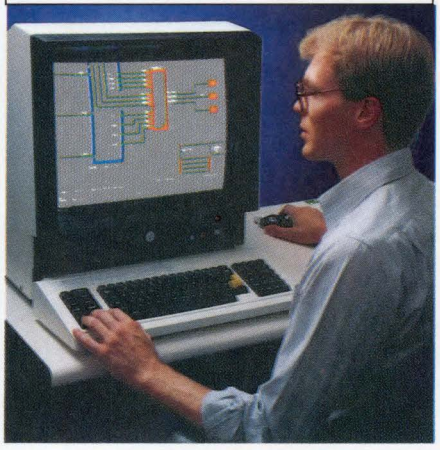

TI's standard cells are cost-effective

Texas Instruments offers a full spectrum of logic alternatives—semicustom, including standard cells and gate arrays, advanced bipolar logic, and HCMOS logic. Thus, we at TI are in a unique position to help you weigh all the factors concerning each to determine which alternative is best for your application.

In the case of standard-cell ICs, we'll analyze how you intend to use them and how they will impact your overall system, including performance and cost.

One standard-cell IC can contain the equivalent of hundreds of individually packaged chips. It can dramatically reduce the number of individual parts needed to implement a system. It can significantly cut weight, size, and power consumption while increasing system quality, reliability, and capability.

Then, too, TI standard-cell chips are the lowest cost alternative when produced in mid to high volumes.

Call 1-800-232-3200, Ext. 112

You already know the TTL design rules, so you're well on your way to tailoring logic more precisely to your performance requirements using TI standard-cell technology. For our brochure, pocket selection guide, sample data sheets, and a more detailed description of the design flow, call the telephone number above or circle the Reader Service Card number. Or write Texas Instruments Incorporated, Dept. SRSØ13OS, P.O. Box 809066, Dallas, Texas 75380-9066.

TEXAS INSTRUMENTS

Creating useful products and services for you.

27-4778

Gould...Innovation and Quality in UNIX-based systems

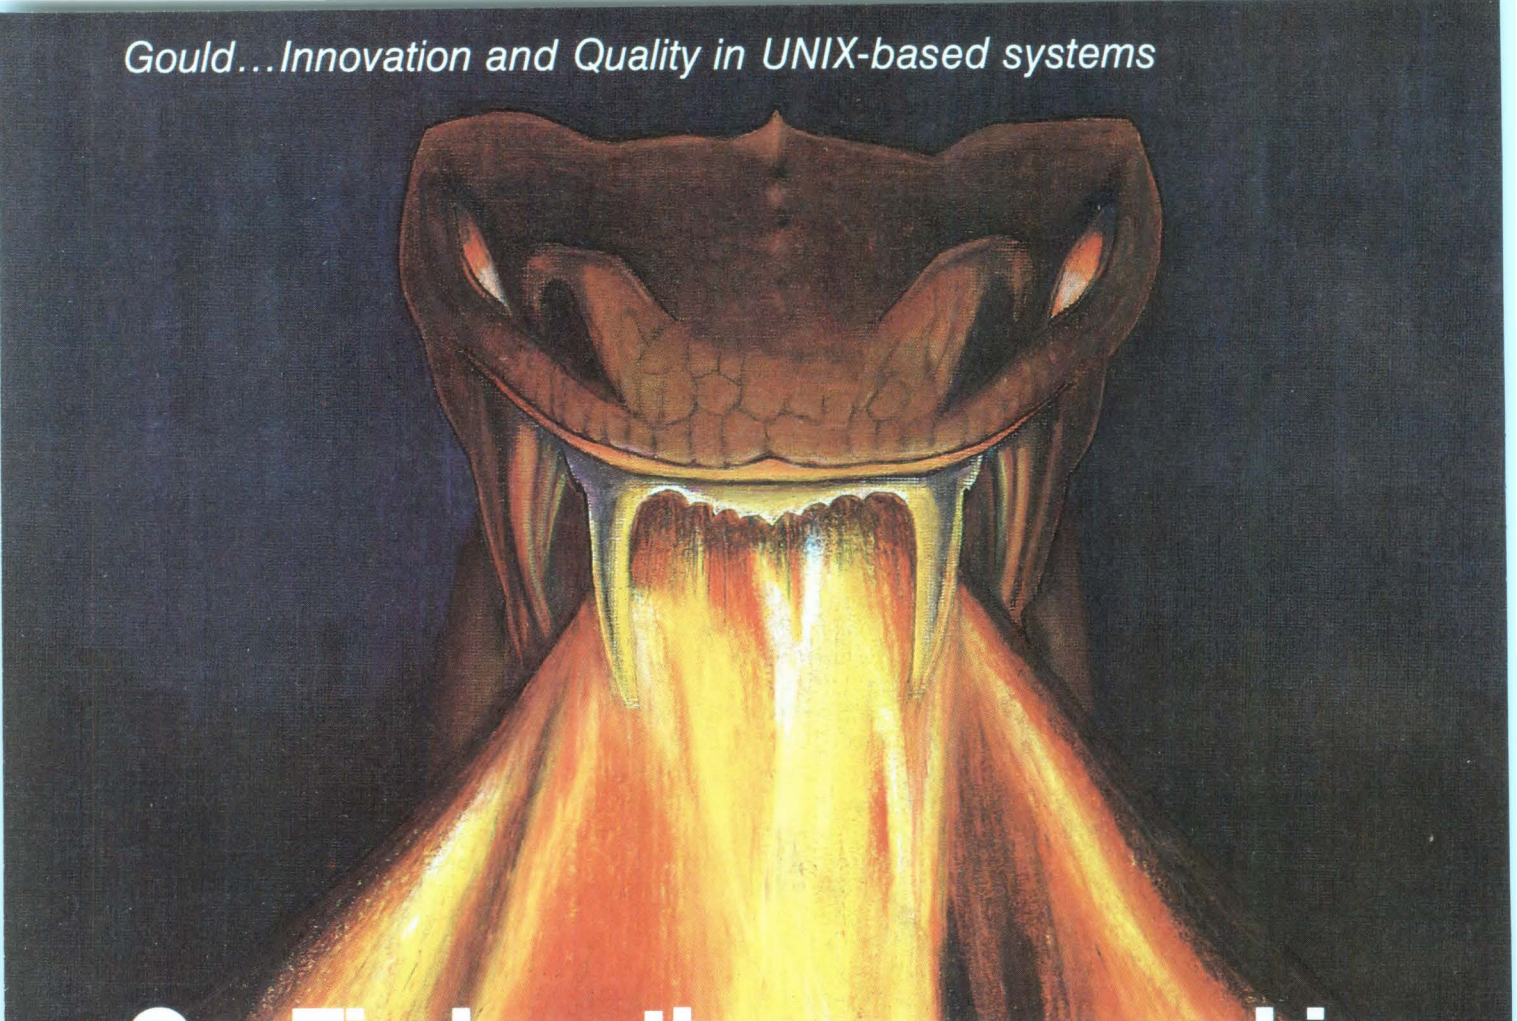

Our Firebreathers are scorching old performance standards.

Gould's PowerNode™ 9000 blasts through UNIX* benchmarks at 4.5 times the speed of the VAX™ 11/780. Sound impossible? Give us your real production code or benchmarks and let us prove it.

Firebreathing Performance.

Now you can run software development and production at the same time, with highly responsive performance. Tightly coupled dual processors nearly double throughput and virtual memory accommodates large programs. Hardware fixed point and floating point accelerators retain high performance in heavy number-crunching situations. The PN9000 handles mainframe jobs in a multi-user UNIX system or serves as a backend processor in a widely distributed network.

Unique UNIX Software.

Gould's own high performance UNIX-based operating system (UTX/32™)—a unique combination of Berkeley 4.2 with special Bell System V features—makes it easy for you to use your VAX-based UNIX software. This allows easy conversion from your system to the increased power of a Firebreather.

Compatible Family.

Gould's Compatibility Suite is a collection of application software packages that are compatible across the entire PowerSeries™ product line. Use C, Cobol, BASIC, or Pascal languages intermixed. This close-knit processor family offers all the advantages of a dedicated system plus the lower-cost-per-user option of sharing resources with Gould's standard networking capabilities including Ethernet™. The Firebreathers are the high end of the widest range of UNIX-based systems in the industry.

Gould's Firebreathers are scorching the UNIX market.

Gould Inc.,
Computer Systems Division
Distributed Systems Operation
6901 West Sunrise Boulevard
Ft. Lauderdale, Florida 33313
(305) 797-5459

*UNIX is a trademark of AT&T Bell Labs
™PowerNode, PowerSeries and UTX/32 are trademarks of Gould Inc.
™Ethernet is a trademark of Xerox Corporation
™VAX is a trademark of Digital Equipment Corporation

GOULD
Electronics

CIRCLE 115

Protocol analyzer identifies data communication problems

Fully compatible with family members 4955A and 4951A, the HP 4953A protocol analyzer has full remote control capability. The analyzer monitors data transmission and simulates network components. Moreover, it monitors bit-oriented protocols at speeds up to 256 kbits/s and provides complete simulation at speeds of 72 kbits/s without loss of triggering capability.

Remote control gives commands to download a test to another HP 4953A, execute the measurements, and automatically upload the results from another unit, even if that unit is unattended. It also transfers data files and measurements to the HP 4955A or 4951A.

Since the units are individually addressable, the remote capability is helpful in a multiport distributed network. Separate ports on the analyzer allow concurrent remote testing and monitoring without having to switch data links.

The analyzer displays X.25, X.75, DDCMP, BSC, HDLC, SDLC, and user-definable, character-asynchronous, and character-synchronous protocols. High speed protocol analysis needs include the ISDN's 144 kbits/s and computer-to-computer links at speeds to 250 kbits/s. Data codes are ASCII, EBCDIC, Baudot, EBCD, and Transcode. The user can define other data codes and store these codes on tape.

An intelligent 64-Kbyte buffer memory stores data, timing information, and lead status. A standard mass storage tape provides space for 512 Kbytes of data for later analysis. Actual data storage is greater because the analyzer can eliminate line delays without sacrificing timing information. With 63 triggers available simultaneously at full speed, the user can trap characters in real time or post process the data.

The soft keys let a new user work through setup and test menus easily. The keys present only appropriate choices at each decision point in a menu. This eliminates most syntax errors and provides rapid test composition. These tests can be complex and can isolate hardware and software problems. Bit-oriented protocol

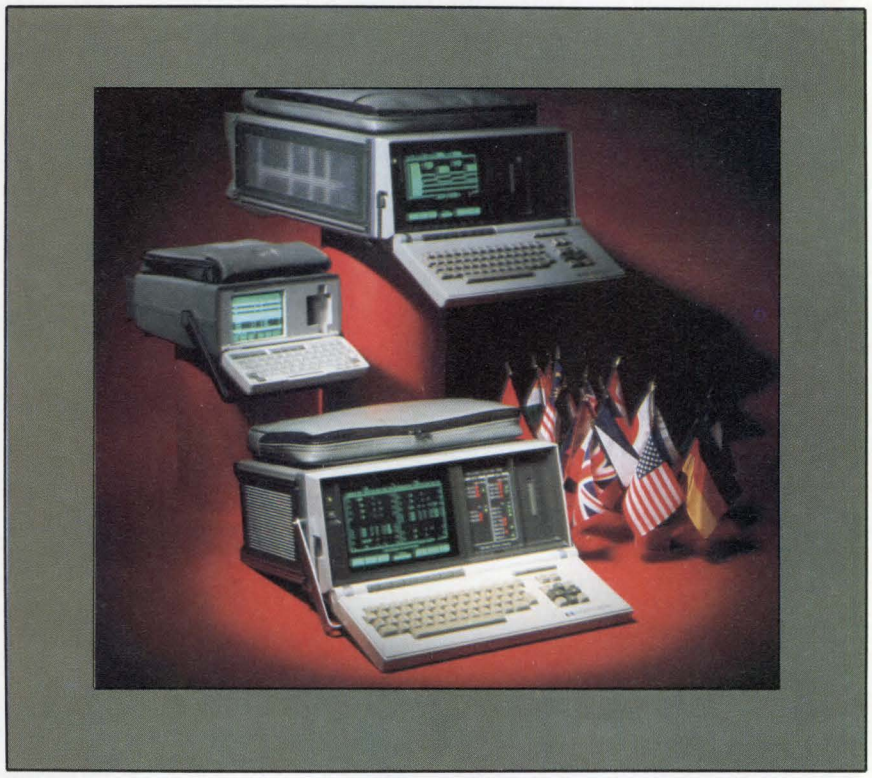

fields are entered via the soft keys. Mnemonically labeled soft keys are available in both the monitor and simulate menus for all Level 2 (frame) and Level 3 (packet) fields. Thus, the user does not need to memorize various bit patterns and positions. Because soft keys also permit lead or character event time measurements, no programming is required.

Multiple display formats include viewing and data interpretation in frame, packet, frame and packet, data, and state—all with 2x display, among others. The display handles timing measurements directly, using cursor timing.

Cartridges store data, timing information, interface and lead status, menu configurations, custom data codes, and application programs. The entire contents of the buffer memory can be stored on a single data cartridge. With option 001, line data is stored directly on tape at 48 kbits/s half duplex (display off) and 19.2 kbits/s full duplex (display on or off).

The display is a high resolution, 9-in. device, featuring 25 lines x 80 chars. Double size characters are selectable. A full ASCII keyboard pivots and locks at any angle for desk, bench, rack, or floor standing operation.

Separate from the interface port is the RS-232-C/V.24 interface for instrument control. Direct hardcopy output of all normal size displays to an HP 2671G or HP 2673A printer is standard. Other standard features include ASCII printer support as an application program; remote transfer of data, menus, setups, status, and application programs; and remote menu execution.

The HP 4953A is priced at $12,000. Option 001 extended memory is $1000. Interface pods (RS-232-C/V.24, RS-449, MIE188C, RS-422) are $950 each. Option 003 Katakana keyboard is $250. **Hewlett-Packard Co**, 1820 Embarcadero Rd, Palo Alto, CA 94303. —*M.B.*

Circle 260

Supermicro gathers 32-bit power in 68000 processor and VMEbus

Both the Universe 2203 and its unbundled VCP2000 CPU use the 12.5-MHz 68000 and the VMEbus for fast, powerful processing. The 68000 offers 16 Mbytes of directly accessible memory, and 32-bit registers and operations. The VMEbus has a 32-bit data path with reliable gastight connections and standard double-height Eurocard PC boards.

The CPU works with a 4-Kbyte cache memory that allows the processor to execute programs without time-consuming wait states. The cache is a 45-ns RAM memory that stores instructions and data onboard. Every 68000 request for an 8- or 16-bit object results in a 32-bit transfer from memory. This loads the cache with a 32-bit data block, keeping the cache one step ahead of the processor. This prefetching, the locality of references in program loops, and stack manipulation, result in a cache hit rate of 70 to 95 percent.

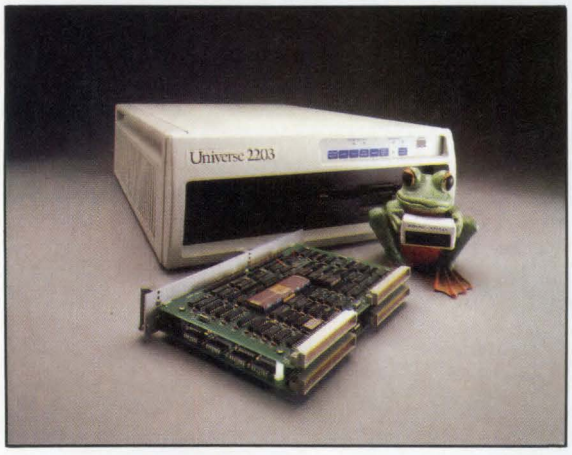

The cache uses a write through technique for processor memory modifications so that data in the cache never needs to be flushed to memory. A cache validity array ensures that the cache accesses return only valid results—even when other DMA devices are active. Together, the cache implementation and the 68000 provide 1.25 MIPS.

Memory allocation and protection hardware protects a shared instruction area between processes. It also provides invalid protection for references outside the user's logical address space. In addition, it allocates memory on a logical basis or loads and/or swaps programs into any area of physical memory.

Peripheral control channels include the selector channel and the serial multiplexer channel. The selector channel interface is modeled after the mainframe version. The interface board has its own processor, and as an independent master device, uses the VMEbus to access channel control blocks stored in the system's buffer area. This removes disk control overhead from the CPU.

The serial multiplexer channel controls character-oriented devices. Output to a single terminal, at a 9.6-kbaud rate, results in about 1000 interrupts/s for the controlling processor.

The Universe 2203 includes a 35-Mbyte Winchester and users can select a floppy disk or quarter-inch streaming tape for system backup. The VCP2000 supports between 0.5 and 16 Mbytes of main memory and features the same memory allocation and protection features found on the Universe 2203. Supermicro costs between $15,000 and $20,000, while the two-board CPU is $4000. **Charles River Data Systems, Inc**, 983 Concord St, Framingham, MA 01701. —*M.B.*

Circle 261

Multibus II-based units deal 32-bit power and run PDP-11 software

Series MCP supermicrocomputers use the 15-MHz, DEC J-11 microprocessor and are highlighted by an especially adaptive 32-bit architecture. Executing code written for the PDP-11 family, these systems employ the Multibus II, as well as an SCSI bus and a Q-bus adapter.

The MCP system is based on a TARCH architecture. Supporting up to eight multiple-coupled processors, TARCH allows 16- and 32-bit CPUs to operate in the same backplane and enclosure. These CPUs operate concurrently, share peripherals, run the same or different operating systems, and communicate with each other at the application software level.

The TARCH approach supports various configurations: tightly coupled—running a single copy of an operating system; closely coupled—running different operating systems concurrently; and loosely coupled, or distributed systems. Allowing multiple processors to share peripherals and run in the same backplane means that execution of I/O boundaries can be expanded by adding a CPU or peripheral processor, rather than by adding a complete standalone computer. A broad performance range of 0.8 to 10 MIPS is possible.

At the heart of the central processor module is a J-11 CMOS microprocessor with 32-bit internal data paths, pipelined architecture, and cache support. The processor offers PDP-11/70 memory management and FP-11 floating point operations. Available operating systems include Unix, RSX, RSTS/E, and RT-11.

For each central processor there is a 4-Mbyte memory module. Each memory module has a 16-Kbyte, dual-tag cache.

Multiple buses mark the MCP series. These include a Multibus II, a Q-bus, an Argonne Phastbus and I/O bus, plus an SCSI bus. The Multibus II parallel system bus supports multiple master processors and has data transfer rates of up to 40 Mbytes/s. Phastbus (a 32-bit, asynchronous bus with 20-Mbyte bandwidth) allows realtime interleaving of Q-bus transfers with high speed instruction and data fetches from memory. The I/O bus sports 4-Mbyte bandwidth.

The MCP-32/16 represents the first system in this series. It can be configured as a half-rack tabletop system, a pedestal system with added I/O capacity, or as a full 19-in. rackmounted system. Minimum systems start around $11,000. **Argonne Systems, Inc**, 525 Julie Rivers Rd, Sugar Land, TX 77478. —*J.V.*

Circle 262

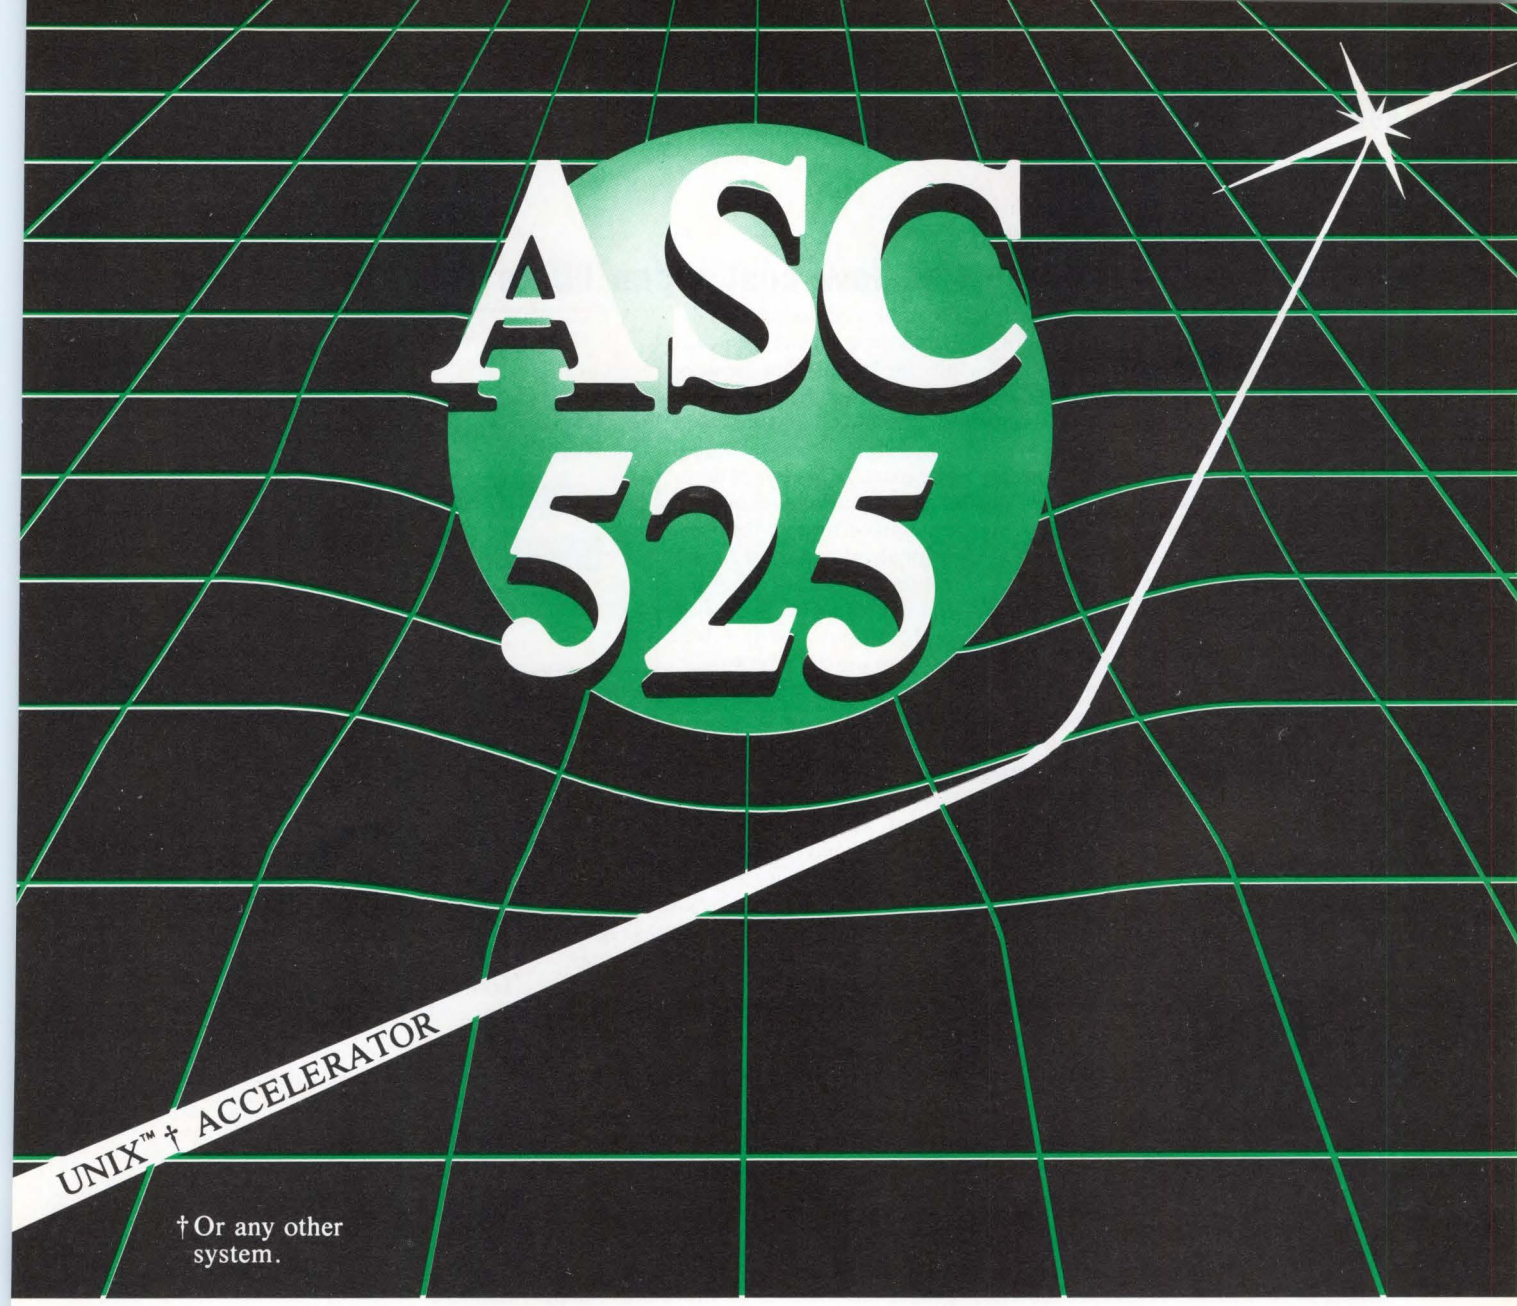

SUPER FAST DISC ACCESS

Accelerate your system's disc performance with the **ASC-525**™ cached disc controller with 20K bytes of fast access ram. This is not a ram-disc that robs the CPU of processing power. It is a full SCSI controller doing true background mode algorithm execution with overlapped cache search and data transfer.

The **ASC-525** speeds the disjoint block transfer of UNIX* by keeping 40 tracks of winchester data current in the under 1 ms access cache - thus shrinking the disc access time penalty.

INTELLIGENT management of disc data - so necessary in today's high performance multi-user environment - is standard with us.

- Full SCSI implementation
  - Arbitration
  - Disconnect/Reconnect
  - Reserve/Release
- Large **320 K** byte cache
  - 1K element size
  - Under 1 ms access time
  - Statistical LRU algorithm
- 5¼ form factor
- Multiple concurrent operations
- 8088 microprocessor (6.7 MHZ)
- Controls two large ST-506 discs

Advanced Storage Concepts

9660 Hillcroft Suite 325
Houston, TX 77096
(713) 729-6388

*UNIX is a trademark of Bell Laboratories.

CIRCLE 116

Realtime subsystem carries low cost video I/O to IBM PC and XT

Designed on one board that plugs directly into an IBM PC or PC/XT expansion slot, the DT2803 frame-grabber subsystem handles realtime video digitization and display. The system also features black and white and color imaging capabilities.

The video input function allows users to digitize video fields in 1/60th of a second and video frames in 1/30th of a second—all via standard RS-170 (60-Hz) or CCIRR (50-Hz) compatible video inputs. Conversion occurs at a 5-MHz rate with the result as 61,440 (65,535 in 50-Hz systems) 6-bit pixels, where each pixel represents 1 of 64 possible gray levels.

The 6-bit values are then converted into 8-bit values via one of eight software selectable input lookup tables. User-programmable tables can fill blank segments of the PROM-based tables.

Four output lookup tables control, and instantaneously change (within vertical blinking time) gray level and color intensities. The 256 x 12 tables define the output color and intensity attributes for each of 256 possible byte values contained in 8-bit memory pixels. A 256 x 256 x 8 frame memory mapped into the PC stores the digitized input signals. Depending on the input table, the frame store memory right-justifies the 6-bit pixels for the A-D converter and stores each as an 8-bit value. The extra 2 bits per pixel in each memory byte include write protection and add additional information to the image, such as user-defined boundaries. The processor or the output lookup tables can then use the 8-bit values.

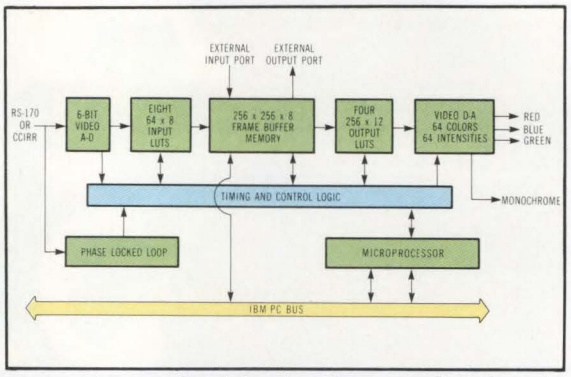

An onboard microprocessor acts as the interface between the DT2803 and the PC. It controls all onboard operations and simplifies program control by the host PC. The microprocessor also carries out routine tasks, such as video output sequencing and error checking. This eliminates overhead in the PC's CPU.

With the micro, four 8-bit I/O locations—command register, status register, data in register, and data out register—control all I/O operations. In addition, several internal registers are accessible through command sequences for error and activity checking, as well as for cursor and intensity/color control.

The subsystem can be programmed in any language. Microcoding in the micro simplifies user access to board functions. A software support package, called Videolab, is also available. The software has an interactive, command-driven tutorial section and a library of subroutines for user-defined application programs.

Price is $1500 for the hardware. **Data Translation Inc,** 100 Locke Dr, Marlboro, MA 01752. —*M.B.*

Circle 263

Microcomputer attains high performance minicomputer power

Based on the J-11 microprocessor, the MicroPDP-11/73 offers performance comparable to a PDP-11/44 minicomputer. The microcomputer features the same enclosure, power supply, and Q-bus architecture as the microVAX and the less powerful MicroPDP-11/23. Because of this compatibility, the entire family of Q-bus peripherals and mass storage devices is available for the new microcomputer.

As expected, the system is software compatible with the MicroPDP-11/23. It also offers considerable software commonality with microVAX. Languages, files, data types, digital command language, and DECnet software are shared by microVAX and micro PDP-11s. Mass storage compatibility is ensured by the use of the Digital Storage Architecture (DSA), used in the entire line.

With performance at least three times better than that of the 11/23, the 11/73 is meant to compete with 68000- or 286-based microcomputers. The 11/73s' CPU board (KDJ11B) increases the throughput of the already high performance J-11 processor chip with an 8-Kbyte cache.

The system can be packaged in a floor pedestal or a tabletop enclosure, or it can be rack mounted. Although 11/73 systems are available in many configurations, several packaged systems are typical. These contain 512 Kbytes of memory, a 31-Mbyte Winchester disk, and either floppy disk drives or cartridge tape. Memory is available in 512-Kbyte increments, up to 4 Mbytes.

Along with the MicroPDP-11/73, new mass storage peripherals have been announced, including fixed/removable disks and streaming tape. The RC25 fixed/removable disk system contains 52 Mbytes of formatted data. Each of the two disks contains 26 Mbytes. The removable 26-Mbyte disk is useful for personal data storage, information exchange, or very fast backup for the fixed disk. The 8-in. disks have seek times of 20 to 35 ms, and the peak transfer rate is 1250 kbytes/s. Included with the disk drive is an intelligent controller. Two complete units (104 Mbytes) can fit in a standard 10½-in. enclosure.

The TK25 streaming tape system offers fast backup of disks with its 60-Mbyte capacity. The quarter-inch tape cartridge transfer rate is 1 to 3 Mbytes/min. Prices of standard, packaged MicroPDP-11/73 systems with 512-Kbyte memory and 31-Mbyte Winchester disk are $15,140 with floppy drives and $19,040 with streaming tape. **Digital Equipment Corp,** 77 Reed Rd, Hudson, MA 01749. —*J.B.*

Circle 264

The lean, mean plotting machine from Houston Instrument

Houston Instrument's brand new servo driven DMP-51 is the fastest drum plotter we've designed to date. This superb plotter offers a pen speed of up to 22 inches/second; programmable accelerations, and a pen-on-paper resolution of 1/1000 of an inch! That means you'll turn out quality 17″ x 22″ and 22″x 34″ drawings a lot faster, increasing your firm's productivity and profitability.

Now, for more good news. The DMP-51, priced at $4,495*, is as fast as other plotters costing three times more. At that price, you can afford to put a DMP-51 at individual drafting work stations.

This is the professional plotter that meets the needs—and the budgets—of all companies, large or small.

The DMP-51 is intelligent, too. The DMP-51 can execute complex graphics operations from the simplest commands. A mechanical/architectural version, the DMP-52, with its 18″ x 24″ and 24″ x 36″ paper size, is available for the same price from Houston Instrument.

So, watch our new plotter in action . . . it won't take much time to realize it's the best buy for your money. For the name, address and phone number of your nearest dealer or distributor, write Houston Instrument, 8500 Cameron Rd., Austin, Texas, 78753. You can also call 1-800-531-5205 or 1-512-835-0900 (Texas residents). In Europe, contact Houston Instrument Belgium NV, Rochesterlaan 6, 8240 Gistel, Belgium. Tel: 059-27-74-45. Tlx: 846-81399.

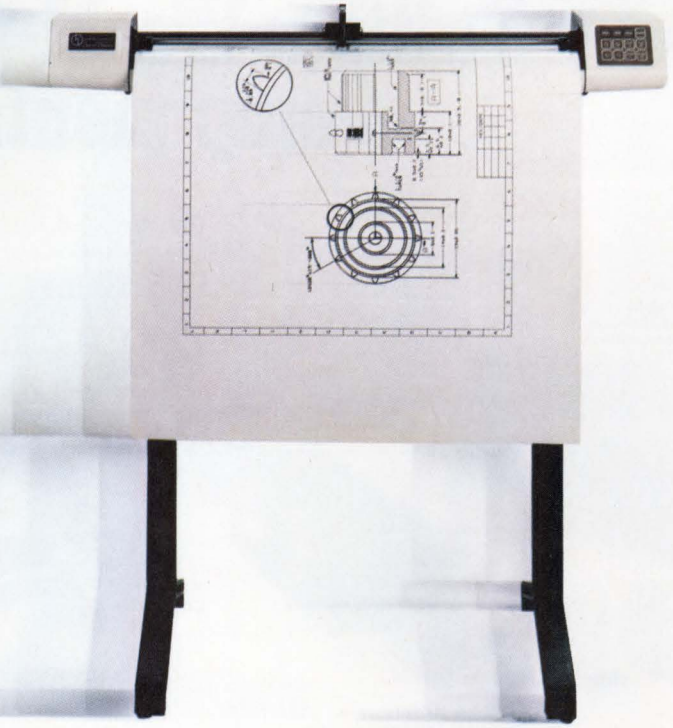

*Suggested U.S. retail price

CIRCLE 117

houston instrument

houston instrument 25th Anniversary 1959-1984

4P014

Varied LAN schemes characterize 16-bit personal computer

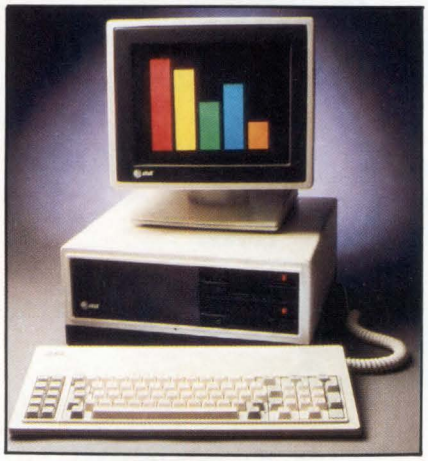

The long-awaited personal computer entrant of communication giant AT&T follows the familiar PC-compatible path. This machine is marked by a high speed processor, high resolution graphics, ergonomic design, and diverse LAN capabilities. Extensive software offerings enhance this MS-DOS unit, which can run Unix in terminal emulation mode via a micro-to-3B minicomputer link.

Based on a full 16-bit 8086 microprocessor operating at 8 MHz, the unit is called the 6300 and comes in two models. The dual-floppy disk model holds two 360-Kbyte disks and has 128 Kbytes of RAM (expandable to 640 Kbytes). The hard disk unit has 256 Kbytes of memory (expandable to 640 Kbytes), plus a 10-Mbyte hard disk and a 360-Kbyte floppy. Seven expansion slots (six on the hard disk model) are provided.

Display resolution for characters measures 640 x 400 pixels. For PC-DOS compatible monochrome and color graphics, resolution is listed at 320 x 200. Proprietary graphics from AT&T have a 640 x 400 video resolution.

A nonglare screen and a tilt-and-pan monitor reduce eye strain. The adjustable keyboard has LED indicators to show that number and caps-lock keys are engaged.

Diverse networking options distinguish this offering. Using the PC Interface, a 6300 can connect to the firm's recently released 3B 32-bit, Unix-based supermicro and minicomputer family. Concurrent with the announcement of the 6300, AT&T has outlined its Information Systems Network (ISN). This represents an effort to connect its computer line, and units from other vendors, in an open LAN operating at 8.64 Mbits/s. Estimated cost per connection ranges from $400 to $500. The ISN LAN uses a packet-switched, star topology, with media composed of fiber optic or twisted-pair wire. The ISN can link with AT&T PBXs and protocol converters that offer SNA and Bisync connections. And, a 3BNet provides Ethernet compatible interfaces for 3B systems.

This is an MS-DOS based machine. Available application software includes dBaseII, WordStar/Pro, Lotus 1-2-3 and Friday!, Digital Research PL/I and CBasic Compiler, Microsoft Pascal, Fortran, and Cobol, as well as 4410 Emulation and Context Switching. Context Switching lets users switch from a PC application to a terminal emulation mode while saving that PC application. Thus, in the terminal emulation mode, the 6300 can access the Unix operating system.

With DOS, the dual-floppy personal computer costs $2810, and the hard disk version costs $4985. **AT&T Information Systems,** 100 Southgate Pkwy, PO Box 1955, Morristown, NJ 07960. —*J. V.*

Circle 265

Video DRAM adds graphics power to workstations

Two engineering workstations offer a choice of high resolution monochrome or color graphics. The 4000 and 4200 incorporate Texas Instrument's multiport video DRAM and can stand alone or be used as a node in a distributed system environment. Contained in a housing that fits under a desk, the workstations feature up to two 5¼-in., 38.6-Mbyte Winchester disks, up to two 5¼-in. floppy disks—each with a 737-Kbyte capacity, and built-in Ethernet LAN.

The system's graphics architecture takes a raster buffer, closely coupled with CPU and system memories, and combines it with a set of assembly level graphics instructions. The raster buffer houses the 64-Kbit video chips. These devices have a fast onchip shift register that supports dot rates in excess of 150 MHz. The workstations require support for an interlaced 1024 x 1024 display with dot rates around 50 MHz. This chip also provides display refresh through a separate serial port, allowing the CPU greater random data port access. The result is a graphics memory that sits directly on the high speed memory bus, simplifying control hardware.

Designed to run on all 32-bit workstations having bit-mapped graphics, the graphics intruction set (GIS) is considered a Data General standard. It provides the basis for window management with a firmware concept called a form. All GISs define pixel coordinates relative to this form. The forms, in turn, are building blocks for windows. Operating within forms, GIS clips all primitives to the form boundary, defines pixel coordinates relative to user coordinates specific to each form, and manipulates colors independently for each form. With this process, several independent windows are displayed and written simultaneously, with minimal system and user overhead.

Six boards make up the system hardware. The two-board CPU is based on the Eclipse MV4000 and consists of a control unit and a processor unit. Built-in reliability features include memory error detection and correction, memory sniffing, and power-up tests.

The I/O board is a microEclipse-based controller. It provides interfaces to three asynchronous ports, a Winchester controller, a floppy controller, a cartridge tape unit, a keyboard, and a LAN.

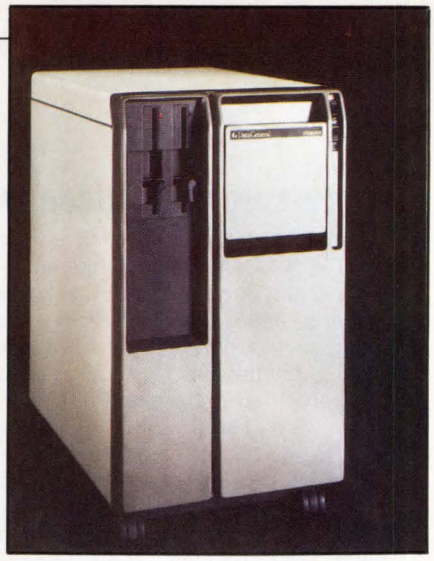

Other boards include a color graphics board and monitor supporting 1024 x 1024 x 4 bit-mapped raster display. Sixteen simultaneous colors are available from a palette of 4096. Memory boards provide 1-, 2-, or 4-Mbyte capacities.

Running AOS/VS software and a native Berkeley Unix operating system, workstations support languages including Fortran 77, C, PL/1, and Basic. Systems are $35,500 to $59,500. **Data General Corp, Technical Products Div**, 4400 Computer Dr, Westboro, MA 01580. —*M.B.*

Circle 266

Thirty-two bit microprocessor sustains 2 to 3 MIPS

Integrating 192,000 transistors on a single chip, the MC68020 includes nonmultiplexed, 32-bit, internal and external data and address paths. The chip is built using a 2-micron HCMOS process to produce the 375 x 350 mil die. It is packaged in a 114-lead pin grid array that dissipates less than 1.5 W. Functioning at 16.67-MHz clock frequency, the chip operates at burst rates exceeding 8 MIPS and sustains instruction rates of 2 to 3 MIPS.

In the MC68020, all registers, program counters, stack pointers, and arithmetic and logic units are 32 bits wide. The 32-bit design eliminates instruction timing differences for 8-, 16-, and 32-bit operations. The bus interface dynamically adjusts the data bus width on a cycle by cycle basis, to accommodate 8-, 16-, or 32-bit devices. Consequently, existing 8- and 16-bit peripheral subsystems can be used with the 32-bit microprocessor.

Although it is object-code compatible with earlier members of the 68000 family, the MC68020 has added addressing modes and instructions that help in the development of high level language compilers and graphics applications. Existing 68000 instructions have 32-bit extensions in the 68020. The addressing modes include 32-bit displacements, memory indirection, and scaled indexing. New instructions include bit-field operators, double-ended bounds checking, BCD data compression and expansion, module support, and enhanced system calling functions.

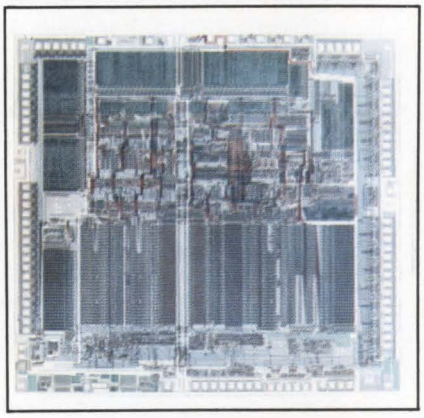

Hardware support of virtual memory results in direct access to 4 Gbytes of logical memory. Like the 68010, the new processor can also process page faults by suspending instruction execution until the physical memory can be updated with the requested information from disk.

A coprocessor interface is provided as a means of extending the instruction set with offchip devices. Two important coprocessors that will be introduced for the 68020 are the long-awaited FPP and the MMU needed for efficient support of demand-paged virtual memory. The MC 68881 FPP has been promised for late 1984. A gate array version of the MMU—the MC68461—will also be available by late 1984 but the full-blown HCMOS chip—the MC68851—will not be available until the middle of 1985.

Initial price for the MC68020 is $480 during early sampling. **Motorola Semiconductor Products, Inc,** 3501 Ed Bluestein Blvd, Austin, TX 78721. —*J.B.*

Circle 267

Analogic News

BUSINESS NEWSPAPER

ONE DOLLAR

FEBRUARY 1984

ISSUE NO. 1

AP400 ARRAY PROCESSOR INSTALLATIONS PASS 1,500 MARK

INSIDE-

• New Analogic Houston Office to Support Array Processors In Oil Industry

• AP400 Series Now Interface Directly with HP A-Series Computers

AP500 Array Processor Offers Unique, Stand-Alone Super-Computing Capability

WAKEFIELD, MA...The AP500 line of array processors has now joined the well-established AP400 line from Analogic. However, unlike any other array processor available, the AP500 has an on-board microprocessor that allows it to operate without a host for certain operations. This unique

AP400 Installed Base Now Greater Than Any Comparable Machine

WAKEFIELD, MA...Analogic Corporation has shipped the 1,500th AP400 Array Processor from its Wakefield, Mass. facility. The AP400's high performance and low cost have generated such widespread acceptance that it now has the largest installed base of any array processor in its class. The AP400 has found application in such diverse areas as medical imaging, automatic test equipment, and seismic data analysis.

The availability of the AP400 in card-set form has made it a popular choice for OEM's wishing to upgrade the

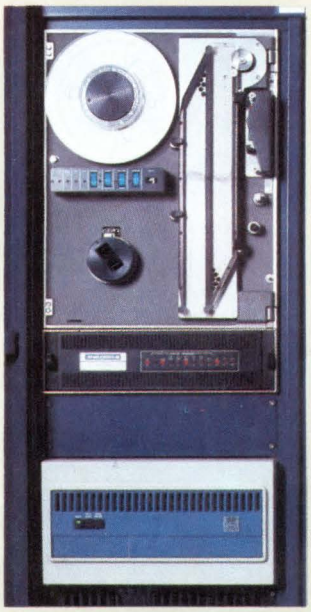

A major application of the AP400 is as a subsystem within automatic test systems such as this LTX... The AP400 offers a cost-effective means of adding powerful linear test capability to digital ATE systems.

The AP400 and ATE

The AP400 Array Processor's high-speed number-crunching power can make it a key component within an ATE system. It can serve a number of purposes in such a system, acting as an active signal filter, a high-speed data manipulator, or even as a waveform synthesizer.

The AP400 brings the performance of computationally-intensive operations into the real-time domain. Transfer function analysis, convolution and correlation, and power spectrum calculation are but a few of the procedures which can be performed in just milliseconds with an AP400 — several hundred times faster than with an unaided minicomputer.

Minimizing host burden was a prime consideration in designing the AP400. Such features as direct memory access to the host, a powerful on-board control processor for internal housekeeping functions, and internal table storage and lookup ability mean that for many applications the host processor need only be involved in telling the AP to start and in picking up the processed data. The auxiliary input and output ports also help minimize host burden. Raw data can enter the AP400 directly, without host involvement, and processed data can be sent directly to peripheral equipment. This feature is particularly useful when the AP400 is used as a waveform generator. The AP400 can send the synthesized signal directly to a test bed without involving the host. The AP400 is currently in use in such diverse applications such as checking codec pairs and airborne radar testing. To find out more about how the AP400 can help you with your ATE problems, contact

Bruce Mackie
Analogic Corporation
Computing Systems Group
Audubon Road
Wakefield, MA 01880
1-800-237-1011
(In Massachusetts call,
617-246-0300, X2176)

AP400 Performance Features

1024 point complex FFT	7.4 ms
Real convolution (1024 data, 512 kernel)	7.3 ms
2048 point power spectrum (with Hamming window)	13.5 ms
1024 point, ensemble-averaged transfer function	1.2 s

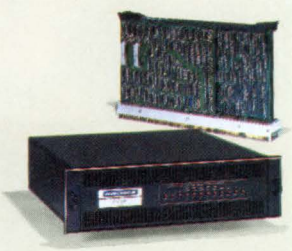

ANALOGIC®
COMPUTING SYSTEMS PRODUCTS GROUP

CIRCLE 121

At $795, our DEC™-comp
a sophisticated image

atible terminal projects even when it's off.

The WY-75 is the one VT-100 software-compatible terminal that looks intelligent even when it's not on.

And at $795, it looks as smart to the people who pay for it as it does to the people who look at it all day.

Like all our terminals, the WY-75 combines an unusually small footprint with a very generous 14″ diagonal display. The non-glare screen tilts, swivels, and displays a full 132-column format. The sculpted, low-profile keyboard adjusts, too.

We've done everything to make the WY-75 the best looking, best feeling, best performing terminal anywhere. On or off.

Can such a beautiful piece of equipment be the workhorse we say it is? Try it, hands on, and see your application software on the screen.

To find out where you can see the WY-75 demonstrated, call the regional office nearest you.

For more information about our complete line of products, write or call Wyse Technology, 3040 N. First Street, San Jose, CA 95134, (408) 946-3075. Outside of California 800-421-1058.

WYSE

Regional Offices: Northeast (201) 725-5054; Southeast (305) 862-2221; Midwest (313) 471-1565; Southwest (213) 340-2013; Northwest (408) 559-5911.

DEC and VT-100 are trademarks of Digital Equipment Corporation.

CIRCLE 122

Between Your System Micro Are Two Lin

Our Controller Line.

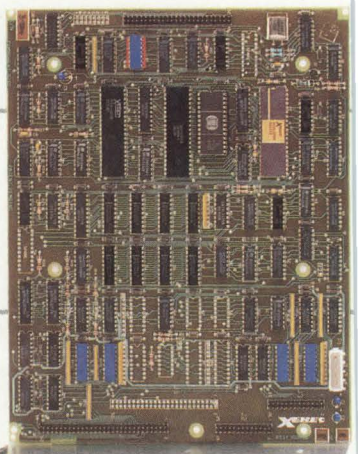

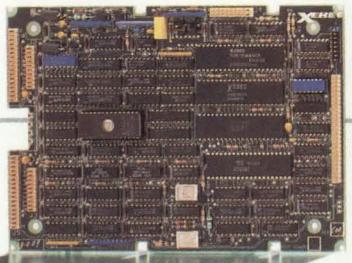

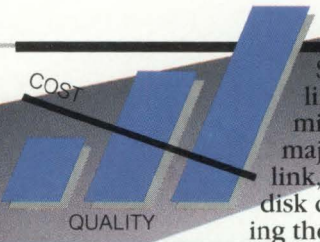

Smoothly and efficiently linking today's diversity of micros and disk drives is a major headache for OEMs. That link, of course, is made by a disk drive controller, and choosing the right one means walking a fine line between getting the performance and reliability you want and the amount of money you want to spend to get it.

Xebec is that fine line, featuring the most popular across-the-board offering of disk drive controllers in the industry—and backed with more expertise in solving disk drive integration problems than anybody else. Period.

More Choices. Fewer Decisions.

One mark of leadership is Xebec's wide array of disk drive controller products—designed, built and tested to match more microcomputers with more disk drive types. Off-the-shelf controllers for all the best-known standard drives. Controllers for 5¼-inch drives and for 3½-inch drives. For the growing number of half-height Winchesters. For fixed drive and floppy drive combinations. For fixed and removable combinations. And custom controllers that solve special integration problems.

Each is at the top of its class, providing high performance in the most demanding applications. Xebec's easy-to-integrate, single-board designs are based on reliable MOS microprocessors, integrated with the latest standard cell and surface mount IC solutions. A long list of performance and reliability features include, among others, sophisticated data separation, advanced error detection and correction, hard-fault isolation, and a high-level command set.

Because the Xebec controller line offers more and better solutions, you have to make fewer decisions. You'll spend less time comparing specifications, and less time trying to think up ways to work around some less-than-perfect match-ups between drive and controller. Less time on in-house testing and post-sales service. Less time sharpening your cost-cutting pencil.

Tested Pairs: Another Standard-Setting Solution.

Problem: Because our controllers perform so well, and so reliably, it's not always easy to find a disk drive that measures up to our high standards.

CORPORATE HEADQUARTERS
2055 Gateway Place, Suite 600
San Jose, California 95110 U.S.A.
(408) 287-2700

nd Your System Disk Drive, To Consider.

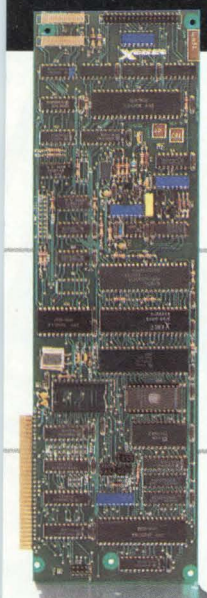

Your Bottom Line.

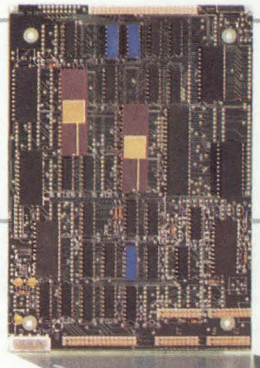

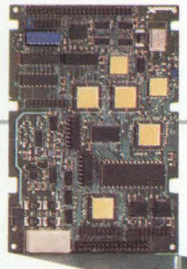

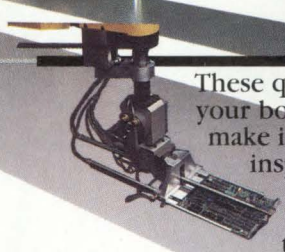

Solution: Xebec's tested pairs program that tests and matches drives and Xebec controllers, guaranteeing that they'll work in tandem, using Xebec-designed test equipment and Xebec-defined tough test standards. Available now in a range of capacities and configurations, tested pairs will help you cut down on evaluation and incoming inspection with a single-source, single purchase order solution you can trust.

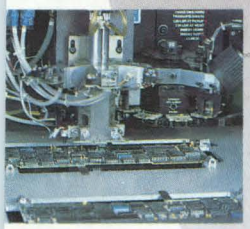

Our Bottom Line. And Yours.

Our "Xero D" signature defines the bottom line: zero-defect quality, companywide.

To make that happen, we've equipped our design engineers with the latest CAD/CAM facilities and equipment to make sure the quality is built in from the start. Test engineers have put together an imposing battery of test devices including a sophisticated drive tester used in our tested pairs program. We've made a multi-million dollar commitment to automated manufacturing.

These quality investments pay off on your bottom line. And in savings you make in up-front design, evaluation, inspection and no-hassle integration. In high-volume, high-quality products delivered on time. Most important, perhaps, you get to market faster, with added value that will keep you there longer.

So call Xebec now. Ask about our controller line. And we'll have a Xebec representative get right down to the bottom line with you.

Xero D.

XEBEC

The Zero Defect Company

SALES, INTERNATIONAL,
Belgium
32-02-762-9494
England
44-734-693511
Italy
39-6-350201

SALES, U.S.A.
Sunnyvale, CA (408) 733-4200
Irvine, CA (714) 851-1437
Atlanta, GA (404) 457-9872
Boston, MA (617) 740-1707
Dallas, TX (214) 361-0687
Baltimore, MD (301) 992-7377
Federal (301) 621-3010
Chicago, IL (312) 931-1420

U.S.A. AND CANADA DISTRIBUTORS
Kierulff
Hamilton-Avnet
Avnet Electronics
Hamilton Electro Sales

CIRCLE 123

SEE YOUR CIRCUIT IN A FLASH...

...OF COLOR.

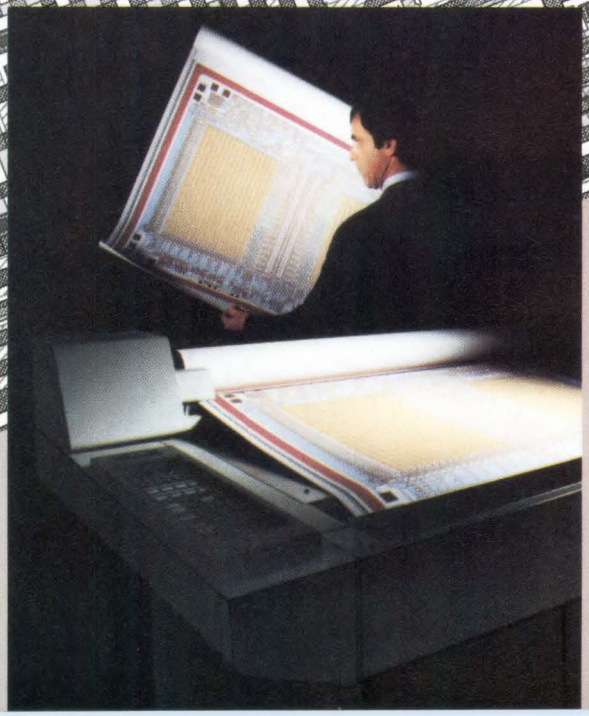

XEROX® is a trademark of XEROX CORPORATION.
Versatec is a trademark of Versatec.

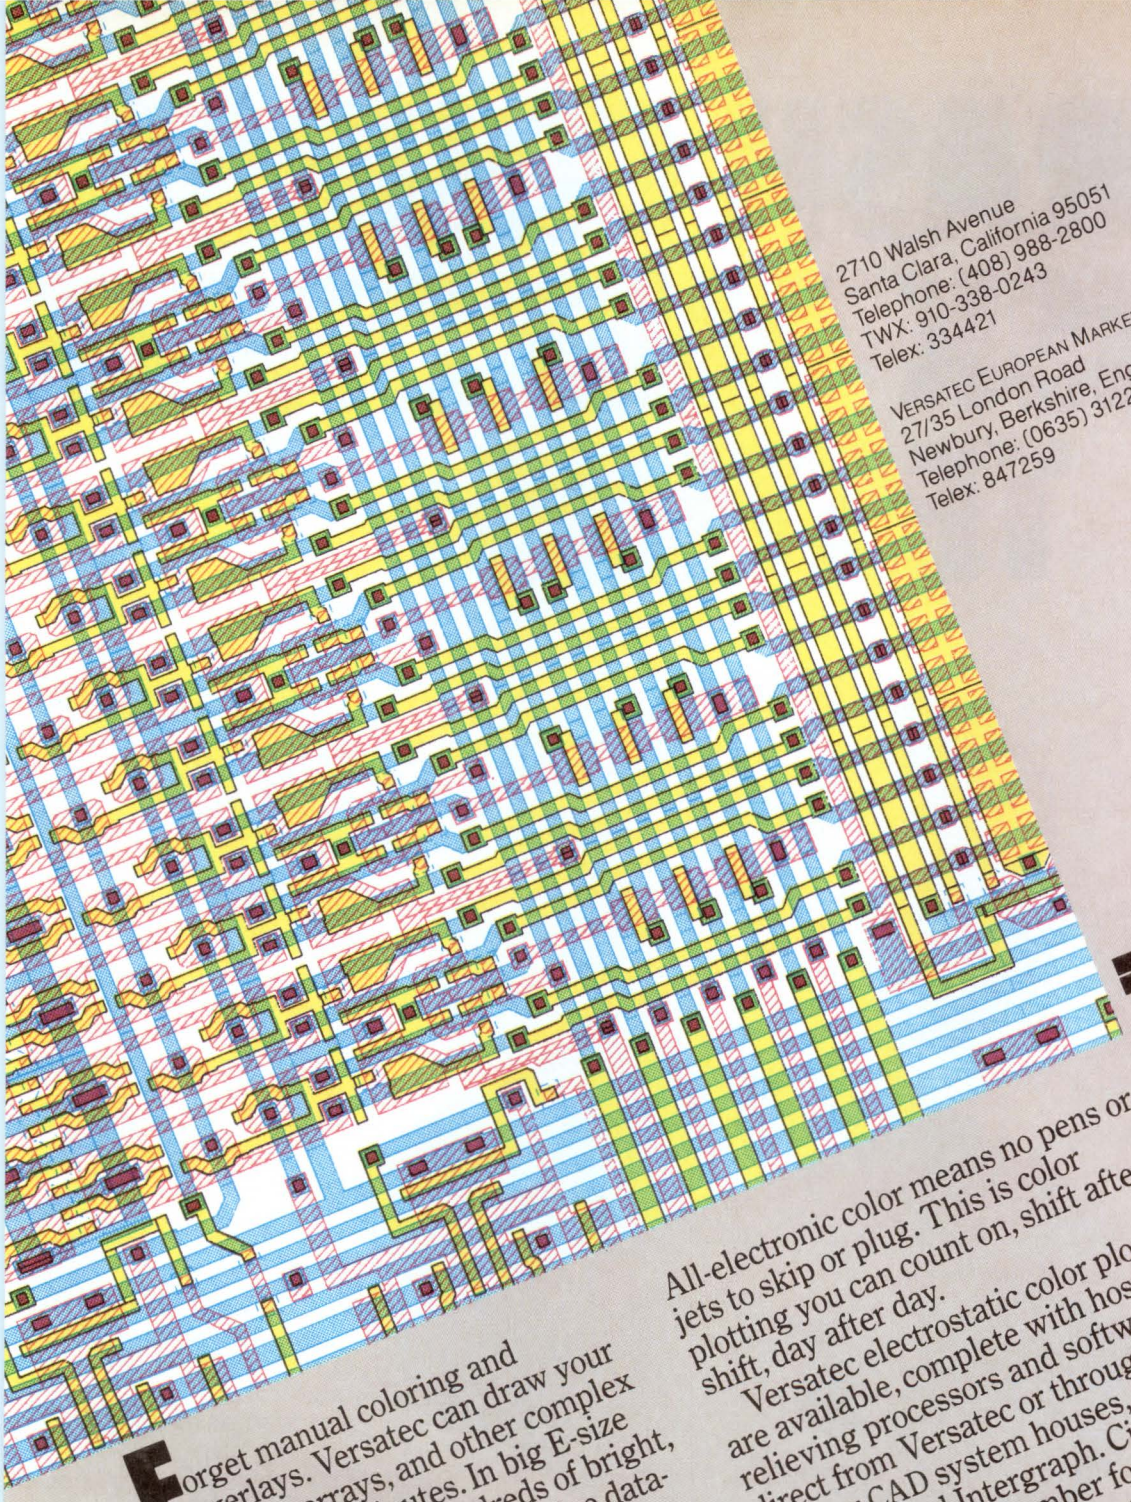

2710 Walsh Avenue
Santa Clara, California 95051
Telephone: (408) 988-2800
TWX: 910-338-0243
Telex: 334421

VERSATEC EUROPEAN MARKETING
27/35 London Road
Newbury, Berkshire, England
Telephone: (0635) 31221
Telex: 847259

VERSATEC
A XEROX COMPANY

Forget manual coloring and overlays. Versatec can draw your VLSI, gate arrays, and other complex IC designs. In minutes. In big E-size formats. And with hundreds of bright, vibrant shading colors to define data-paths, junctions, and layers.

Only Versatec electrostatic plotters draw with so much speed, color, and reliability. Faster color gives you quicker turnaround and more through-put for a faster product design cycle. Discriminating color enables faster analysis of complex circuit designs.

All-electronic color means no pens or jets to skip or plug. This is color plotting you can count on, shift after shift, day after day.

Versatec electrostatic color plotters are available, complete with host-relieving processors and software, direct from Versatec or through leading CAD system houses, such as Calma and Intergraph. Circle the readers' service number for our free brochure—"15 ways to use electro-static color plotting."

THE FASTEST WAY TO COLOR.

VERSATEC.

A CRASH COURSE IN DISK AND DRIVE TESTING:

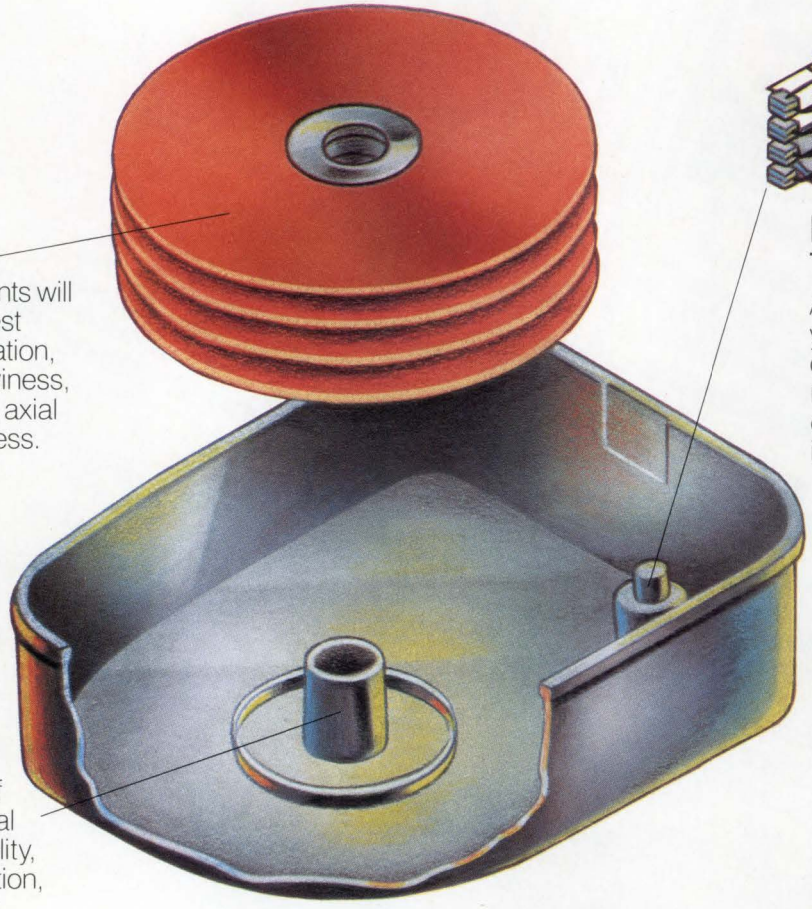

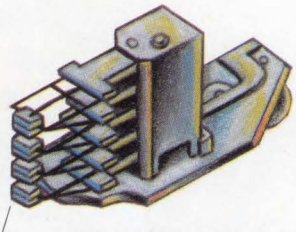

Disk Testing

ADE RVA instruments will show you how to test excessive acceleration, flatness, radial waviness, datum positioning, axial run-out and thickness.

Head/Assembly Testing

ADE RVA instruments give you advanced instruction on head positioning accuracy, head motion studies, dynamic flight characteristics, pitching and rolling.

Spindle Testing

Learn the nuances of testing axial and radial runouts, bearing quality, axial/radial acceleration, non-repetitive runout, radial resonance, wobble, and high frequency vibration.

ADE RVA instruments maintain quality control from design through production. Only ADE systems can measure dynamic displacements from tenths of microinches to thousandths of an inch from 0 to 50 KHz frequency response. Sign up for the ADE course (every major manufacturer of disks and drives already has).

ADE Corporation
77 Rowe Street
Newton, MA 02166
Telephone: (617) 969-0600
Telex: 922415

CIRCLE 125

Transceiver complies with IEEE 802.3 standards

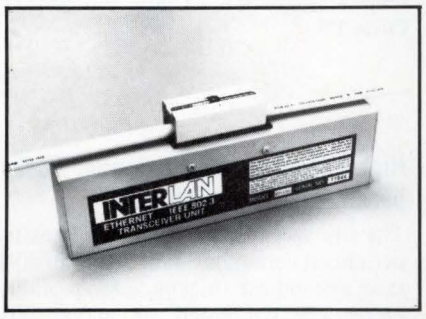

The NT100 features a nonintrusive cable tap, permitting nodes to be attached or removed from an active network without disturbing network communications. An alternative tap with an n-type connector is available for network portions that are installed with preassembled coaxial cable. It can send and receive 10-Mbit/s streams, detect collision occurrence, provide electrical isolation, and protect the network from malfunctions. Price is $250 in quantities of 100. **Interlan, Inc**, 3 Lyberty Way, Westford, MA 01886.
Circle 268

Ethernet interface connects VAX-VMS units to PC net

A software interface, called the VMS Ethernet Server, allows VAX computers running the VMS operating system to network with IBM PCs via Ethernet. The interface lets a VMS system function as a file server in 3Com Corp's EtherSeries LAN. The VMS Ethernet Server will run on VAX 11/730, 11/750, and 11/780 models. The price is $7500. **Syntax**, 6642 S 193rd Pl, Kent, WA 98032.
Circle 269

Speech I/O "engine" provides advanced voice storage

The Dialog system performs realtime voice digitization, disk storage, and retrieval from a microphone, telephone, or LAN interface. Data storage requirements are below 2 Kbytes/s at a 6-kHz sampling rate. Modular device drivers are provided, as are hooks to C, Pascal, Basic, and assembler languages. The system is composed of three board options. Dialog/1 handles basic I/O. It costs $295. Dialog/2 adds phone interface, dial/answer firmware and tone decoding, and costs $495. Dialog/3 includes a 300-baud modem and digital transmission firmware, and sells for $595. **Dialogic Corp**, 164 McKinley Ave, East Hanover, NJ 07936.
Circle 270

Local and wide area net gives PC users access to VAX data bases

Communication support for 3274 BSC and SNA/SDLC protocols, as well as for IBM PC and DEC VT-100 protocols, is one of many X-Net capabilities. This dual-bus network with rooted-tree topology transmits on 95-Ω twinax cable. X-Net allows up to 2032 connections per site and can interconnect up to 255 sites via X.25 gateways. It uses a roll-call polling scheme that diminishes data collisions and accommodates throughput of up to 14.75 Mbits/s. **CR Computer Systems, Inc**, 5456 McConnell Ave, Los Angeles, CA 90066.
Circle 271

Matrix switch manages the world according to T-1

The T-1 VSM data communication matrix switch aimed at T-1 environments switches and configures 2-Mbit/s data lines. The switch offers signal transparent operation, independent of existing online hardware. The T-1 VSM is control compatible with T-Bar's RS-232 digital VSM and wideband VSM. Each can operate independently or can be managed by Overlord, a centralized resource management system. Pricing ranges from $40,000 to $64,000. **T-Bar**, 141 Danbury Rd, PO Box T, Wilton, CT 06897.
Circle 272

Modem features digital signal processing and full diagnostics

The 7160 series of SNA-compatible diagnostic modems expands with the 7164-0200 modem. It is a 4.8-kbit/s switched line modem for point-to-point network configurations. Fully synchronous, the 7164-0200 operates over unconditioned, voice-grade lines in half-duplex mode. It uses digital signal processing for 98 percent of its functions, including modulation, demodulation, and dynamic equalization. It features a full range of diagnostics. Pricing starts at $3850. **NCR Comten, Inc**, 2700 Snelling Ave N, St Paul, MN 55113.
Circle 273

High performance Unibus controller supports 16 communication lines

The VMZ/32HS contains two 8-line multiplexer controller sections that provide 16 asynchronous lines with programmable DMA capability. It occupies one hex-width slot in a DEC DD-11 backplane. The controller supports operation from 50 to 38,400 baud and has an aggregate throughput of 120,000 chars/s. Modem control is available on all lines. Operating mode is full duplex with parity generation/detect at off, even, or none. **Able Computer**, 1732 Reynolds Ave, Irvine, CA 92714.
Circle 274

Aimed at multiprocessing chores, front end uses J11 chip

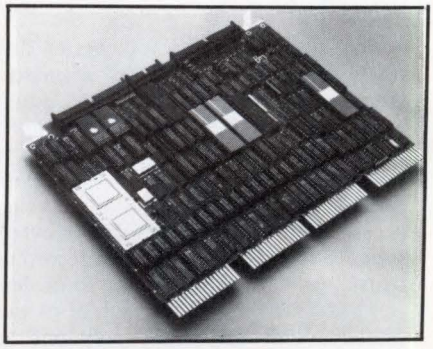

A high speed frontend communication processor allows true coprocessing on the Q-bus. Dubbed the MLSI-JFEP11, the unit features onboard dual-ported memory reaching 512 Kbytes, two high speed serial ports (1 Mbaud each), and an external parallel bus for unlimited I/O expansion. The MLSI-JFEP11 can use all existing PDP-11 programs. Available within 60 days of order, single-unit price is $4950. **MDB Systems, Inc**, 1995 Batavia St, Box 5508, Orange, CA 92267.
Circle 275

Multichannel serial board with four UARTs is 8088-compatible

A four-channel serial communication board—the ZT 8840—is STD-8088 compatible, has an interrupt controller, and features four independent INS 8250 UARTs. The UARTs support speeds from 110 baud to 56 kbaud. Interrupts are generated for each character transfer and upon parity error detection. The interrupt controller allows a STD bus computer to send and receive to and from diverse serial devices. The ZT 8840 is offered at $375 in single-unit quantities. **Ziatech Corp**, 3433 Roberto Ct, San Luis Obispo, CA 93401.
Circle 276

Modem family communicates full duplex over dial telephone network

The DialNet 3000 family of 1.2-kbit/s auto-dialing modems supports asynchronous or synchronous devices. Each modem has nonvolatile memory for up to 20 numbers. Stored numbers may be linked, and can incorporate pulse and tone dialing in the same number sequence. All modems in the series automatically answer incoming calls and adjust their speed and transmission to match the calling modem. Diagnostic aids include analog and digital loopbacks, test pattern generator, and automatic self-test at power on. Prices range from $495 to $795. **Micom Systems, Inc**, 20151 Nordhoff St, Chatsworth, CA 91311.
Circle 277

Modem chip sets feature low signal distortion

Compatible with CCITT V.21 series data sets, MSM6926 telecommunication devices use frequency shift keyed modulation. Other chips in this series handle other popular protocols—the MSM6946 covers Bell 103, and the MSM6947 implements Bell 202. Signal distortion of −10 or +12 percent is exhibited. Using CMOS, these devices offer switched capacitor technology, selectable built-in and external delay timers, a crystal controlled onchip oscillator, and TTL-compatible digital interface. The 6926/6946 units work at up to 300 baud, full duplex. The 6927 (which is CCITT V.23-compatible) and the 6947 operate at up to 1200 baud, half duplex. The MSM6926/6946 chip set costs $15, and the MSM6927/6947 costs $18 in quantities of 100. **Oki Semiconductor**, 1333 Lawrence Expwy, Santa Clara, CA 95051.
Circle 278

High speed channel extender hooks peripherals to mainframes

The MBX-2501 provides two-way, host-to-peripheral communication for up to 2000 ft at a standard 10-Mbit/s block multiplexer channel speed. It allows tape drives, communication controllers, and other peripherals to operate (over both coaxial and fiber optic cable) at an extended distance from mainframe installations without speed reduction. This channel extender consists of two self-powered controllers that are plug-compatible with standard IBM connectors. These function as serial/parallel and parallel/serial converters. **Megabit Communications, Inc**, 90 West County Rd C, St Paul, MN 55117.
Circle 279

MEMORY SYSTEMS

Half-height, 3.5-in. Winchester sports fast access

A half-height, 3.5-in. Winchester disk drive features 10 Mbytes of formatted storage. The HH-312 has low power requirements and rugged design. It operates on 15 W for simple interfacing to portables. A closed-loop servo head positioning unit produces a 70-ms average access time and achieves a 5-Mbyte/s transfer rate. The HH-312 has a recording density of 9680 bits/in. and a track density of 648 tracks/in. for each of its two platters. This drive uses only one PC board. **Microscience International**, 575 E Middlefield Rd, Mountain View, CA 94043.
Circle 280

Micro disk drive incorporates a parallel head loading system

The 3.5-in. micro floppy drive, model MDP-40, measures 101.6 x 28 x 137 mm. It uses an ultrathin direct drive spindle motor, ultrathin hybrid bipolar step motor, and an integrated 60-pin IC. It has a one touch ejection system and is compatible with 5¼-in. mini floppy disk drives. The unit consumes 3 W of power typical. MTBF is 8000 power-on hours (or more), MTTR is 30 min (or less), and entire unit lifetime is 8000 hours. **US JVC Corp**, 41 Slater Dr, Elmwood Pk, NJ 07407.
Circle 281

High capacity miniature data cartridge provides backup

Designed to supply backup for new generation fixed disk systems, the DC1000 stores 10 Mbytes. This miniature data cartridge is dimensionally identical to the standard DC100A (2.415 x 3.188 x 0.470 in.) and contains tape measuring 0.150 in. wide and 0.006 thick with 550-Oe characteristics. The DC1000 contains 185 ft of tape. **3M**, PO Box 33600, St Paul, MN 55133.
Circle 282

Precision servo positioning gains capacity for disk drive

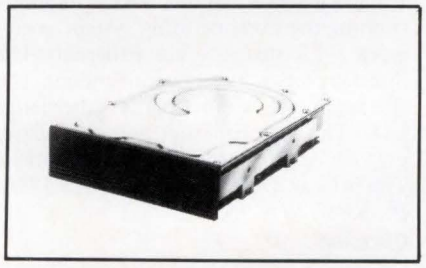

A half-height 5¼-in. Winchester drive, the MR522, provides 25.5 Mbytes of unformatted capacity. Extra capacity results from a precision embedded servo positioning system, advanced rotary arm stepper assembly, and plated media. With 6.37 Mbytes per surface and 10,416 bytes per track, the MR522 allows designers to offer 20 Mbytes of unformatted capacity. Its price is $875. (Evaluation units available for third quarter delivery; quantity production is scheduled for the fourth quarter.) **Mitsubishi Electronics America, Inc**, 991 Knox St, Torrance, CA 90502.
Circle 283

When you're designing a first-rate computer system, it doesn't make sense to compromise with a second-rate monitor.

Consider PGS instead. We make three no-compromise monitors, all fully compatible with the IBM-PC, to match your requirements and your budget.

High resolution PGS set the price/performance standard for high resolution RGB color monitors with the HX-12: .31mm dot pitch, 690 dots horizontal resolution, and precise color convergence for a crisp, sharp image.

Super resolution And, when your specifications call for super resolution, there's our no-compromise SR-12: an RGB monitor with a horizontal scan rate of 27.5 KHz which supports 690x480 resolution in non-interlaced mode.

Monochrome For price/performance in a monochrome monitor, we've set the standard with the MAX-12: our amber monitor with dynamic focusing circuitry which ensures sharpness not only in the center but also in the edges and corners.

The monitor to meet your needs

All three PGS monitors are engineered for no-compromise performance to offer you a cleaner, sharper image than any other monitor in the same price class. The HX-12 and the SR-12 both feature uncompromising color convergence for crisp whites without color bleed. The MAX-12 offers impressive clarity in an amber phosphor monitor that runs off a standard monochrome card—no special card is required.

At PGS, our no-compromise approach includes all the details, too, from non-glare screen to a shielded cable—standard features on all PGS monitors, color or monochrome.

Call us at 800-221-1490

Compare your specifications to ours, listed below. Then call us at 800-221-1490 and we'll send you a fully detailed spec sheet plus everything else you need to know about all three no-compromise PGS monitors.

Don't compromise. Look to PGS for the image your eyes deserve.

PGS **Princeton Graphic Systems**

1101-I State Road
Princeton NJ 08540
(609) 683-1660
TLX 6857009 PGS Prin

Nationwide service provided by Bell and Howell Service Company and MAI Sorbus Service Company.

PGS delivers 3 no-compromise ways to improve your image

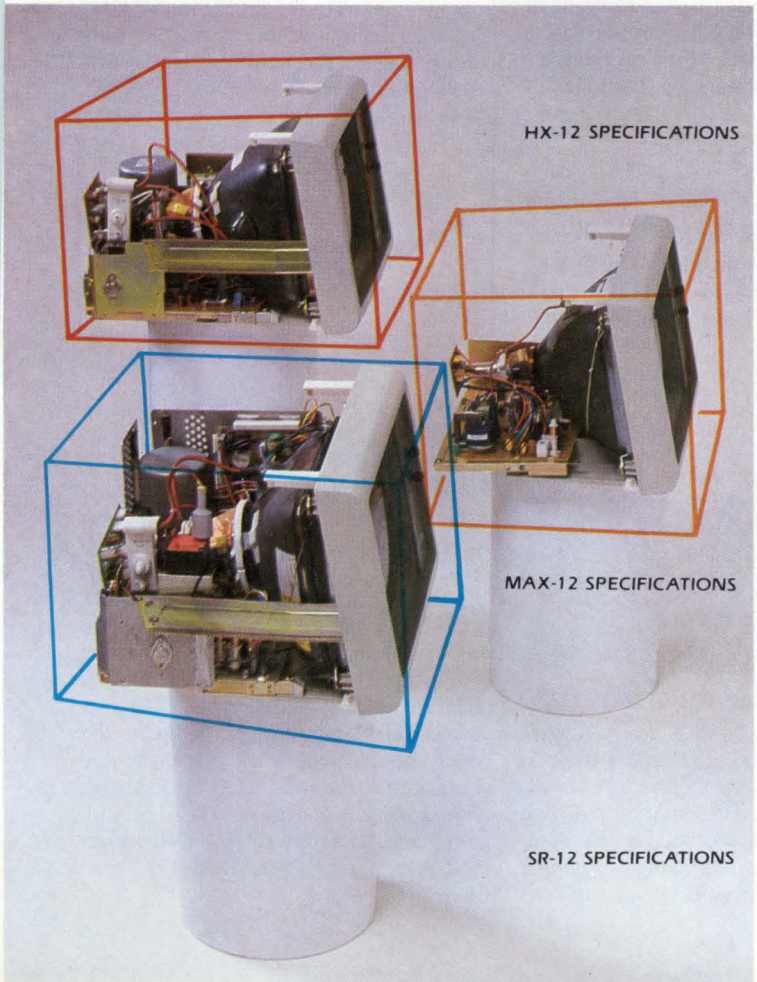

HX-12 SPECIFICATIONS

CRT	12" Diagonal, 76 degree, In-Line Gun, .31mm dot pitch black matrix, non-glare surface (NEC 320CGB22)
Input Signals	R, G, B, channels, Horz Sync, Vert Sync, Intensity—all positive TTL levels
Video Bandwidth	15 MHz
Scan Frequencies	Horizontal: 15.75 KHz Vertical: 60 Hz
Display Size	215mm x 160mm
Resolution	Horizontal: 690 dots Vertical: 240 lines (non-interlaced) 480 lines (interlaced)
Misconvergence	Center: .6mm max Corner: 1.1mm max
Display Colors	16 colors (black, blue, green, cyan, red, magenta, yellow, white, each with 2 intensity levels)
Characters	2000 characters (80 characters x 25 rows—8x8 dots)
Input Connector	9 Pin (DB9)—cable supplied to plug directly to IBM PC

MAX-12 SPECIFICATIONS

CRT	12" Diagonal, 90 Degree, non-glare surface (P 34 Phosphor)
Input Signals	Video signal, Horz Sync, Intensity—positive TTL levels, Vertical Sync—negative TTL levels
Video bandwidth	18 MHz
Scan frequencies	Horizontal: 18.432 KHz Vertical: 50 Hz
Display size	204mm x 135mm
Resolution	Horizontal: 900 dots Vertical: 350 lines
Input Connector	9 Pin (DB9)—cable supplied to plug directly to IBM PC

SR-12 SPECIFICATIONS

CRT	12" Diagonal, 90 Degree, In-Line Gun, .31mm dot pitch black matrix, non-glare surface
Input Signals	R, G, B channels, Horz Sync, Vert Sync, Intensity—all positive TTL levels
Video bandwidth	25 MHz
Scan frequencies	Horizontal: 31.5 KHz Vertical: 60 Hz
Display size	215mm x 160mm
Resolution	Horizontal: 690 dots Vertical: 480 lines (non-interlaced)
Misconvergence	Center: .5mm max Corner: 1.0mm max
Display colors	16 colors (black, blue, green, cyan, red, magenta, yellow, white, each with 2 intensity levels)
Characters	2000 characters (80 characters x 25 rows)
Input Connector	9 Pin (DB9)—cable supplied

CIRCLE 126

Dual-port memory links MIL-STD terminals to 3-state data highway

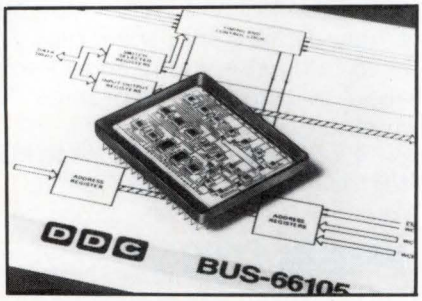

The BUS-66105 double buffered 1-K x 8-bit, dual-port RAM functions as an interface between a MIL-STD-1553 dual-redundant remote terminal unit and a subsystem parallel 3-state data highway. It has dual-address registers and is Multibus-compatible on the CPU side. The BUS-66105 appears to the subsystem CPU as a memory-mapped I/O. Packaged in a 51-pin hermetically sealed hybrid, the device measures 1.85 x 1.59 x 0.21 in. (47 x 40.4 x 5.33 mm) and costs $427. **ILC Data Device Corp**, 105 Wilber Pl, Bohemia, NY 11716. **Circle 284**

Half-height Winchester disk drive features thin-film heads

Combining thin-film heads, closed-loop servo system, and plated media with ESDI interface, the model 96203 has a 100-Mbyte capacity. It contains 960 tracks/in. and 20,880 bytes/track. Average access time is less than 30 ms. Two different models are available in two formats—an entire disk drive or just the head-disk assembly. No ac power is required to drive the spindle motors. Only the industry-standard voltages of 12 and 5 Vdc are used. The heads allow for improved window margins and for greater aerodynamic stability with good signal to noise ratios. **Advanced Storage Technology Inc**, 6580 Via Del Oro, San Jose, CA 95119.

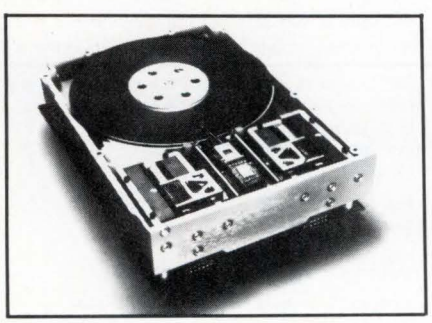

Circle 285

SYSTEM ELEMENTS

Low cost modular system eliminates mismatched cable problems

The EZ-232 line of interconnects uses adapters that are permanently affixed to a computer's serial port and peripherals. They terminate in standard 6-wire, RJ-12 jacks and connect by telephone type cable in various lengths. Another component of this system is the EZ-232 Switch. Push buttons allow routing data between four peripherals or computers with RS-232 adapters. **Micro-Module Systems, Inc**, PO Box 2198, La Jolla, CA 92038.
Circle 286

Solid state keyboard is PC compatible and offers tactile feedback

A capacitance keyboard—the 83ST13-1E—is compatible in key array and electrical interface with the IBM PC. This full-travel, full-function keyboard offers silent, tactile feedback operation in a low profile enclosure. The keyboard also features three-position tilt. Options include LED indicators for number and cap locks and provisions for audio key feedback operation. **Honeywell, Micro Switch Div**, 11 W Spring St, Freeport, IL 61032.
Circle 287

Push buttons and indicators light up for snap-in or PC board applications

The Compu-Lite series 200 is rated at 5 A, at 250 Vac. Fingertip removable switch caps and built-in lamp extractors allow easy replacement of incandescent lamps from the panel front. Features include LED as well as incandescent illumination, full or split legend caps, various colors, and gold contacts for low-level switching. Directly interchangeable, the switches are available immediately. **Eaton Corp, Cutler-Hammer Products**, 4201 N 27th St, Milwaukee, WI 53216.
Circle 288

Discrete switch keyboards retrofit to existing enclosures

Keyboards in the KS-600E series feature discrete switch modules that assemble to produce virtually any main array or ancillary array layout. They snap into a rigid metal plate that ensures proper row alignment, prevents warping, and aids in protecting printed circuits from emi/rfi. Varied keycap styles, plus a large legend library that includes the more popular international legends, are available. The units come in intelligent and nonintelligent form. **Stackpole Components Co**, PO Box 14466, Raleigh, NC 27620.
Circle 289

Semicustom arrays permit linear and digital function integration

Genesis 1100 semicustom arrays are bipolar linear arrays combined with a cluster of 64 integrated injection logic gates. The gates are power and speed programmable over a range of 100 nA to 200 μA, with corresponding propagation delays of 7 to 50 μs. These arrays permit high performance integration of standard linear and digital functions. Power supply range is 1 to 12 V. Integration charge is $6400 and includes 20 packaged engineering prototypes. Factory CAD layout is available with charges ranging from $4000 to $8000 depending on density of chip utilization. **Cherry Electrical Products Corp**, 3600 Sunset Ave, Waukegan, IL, 60087.

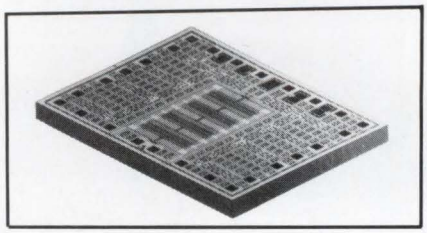

Circle 290

Molded wafer headers deliver high accuracy

A line of molded wafer PC board headers is available for use with MAS-CON mass terminated IDC connector systems. Offered in both 0.1-in. and 0.156-in. centerline versions with straight or right angle posts in flat, locking, or polarizing types, these PC board headers are more accurate than extruded types. The prescored wafer permits breaking the header to any desired length for prototype and short production runs. They are available with either tin plating or 15- or 30-μin. gold plating. **Panduit Corp**, 17301 Ridgeland Ave, Tinley Park, IL 60477.
Circle 291

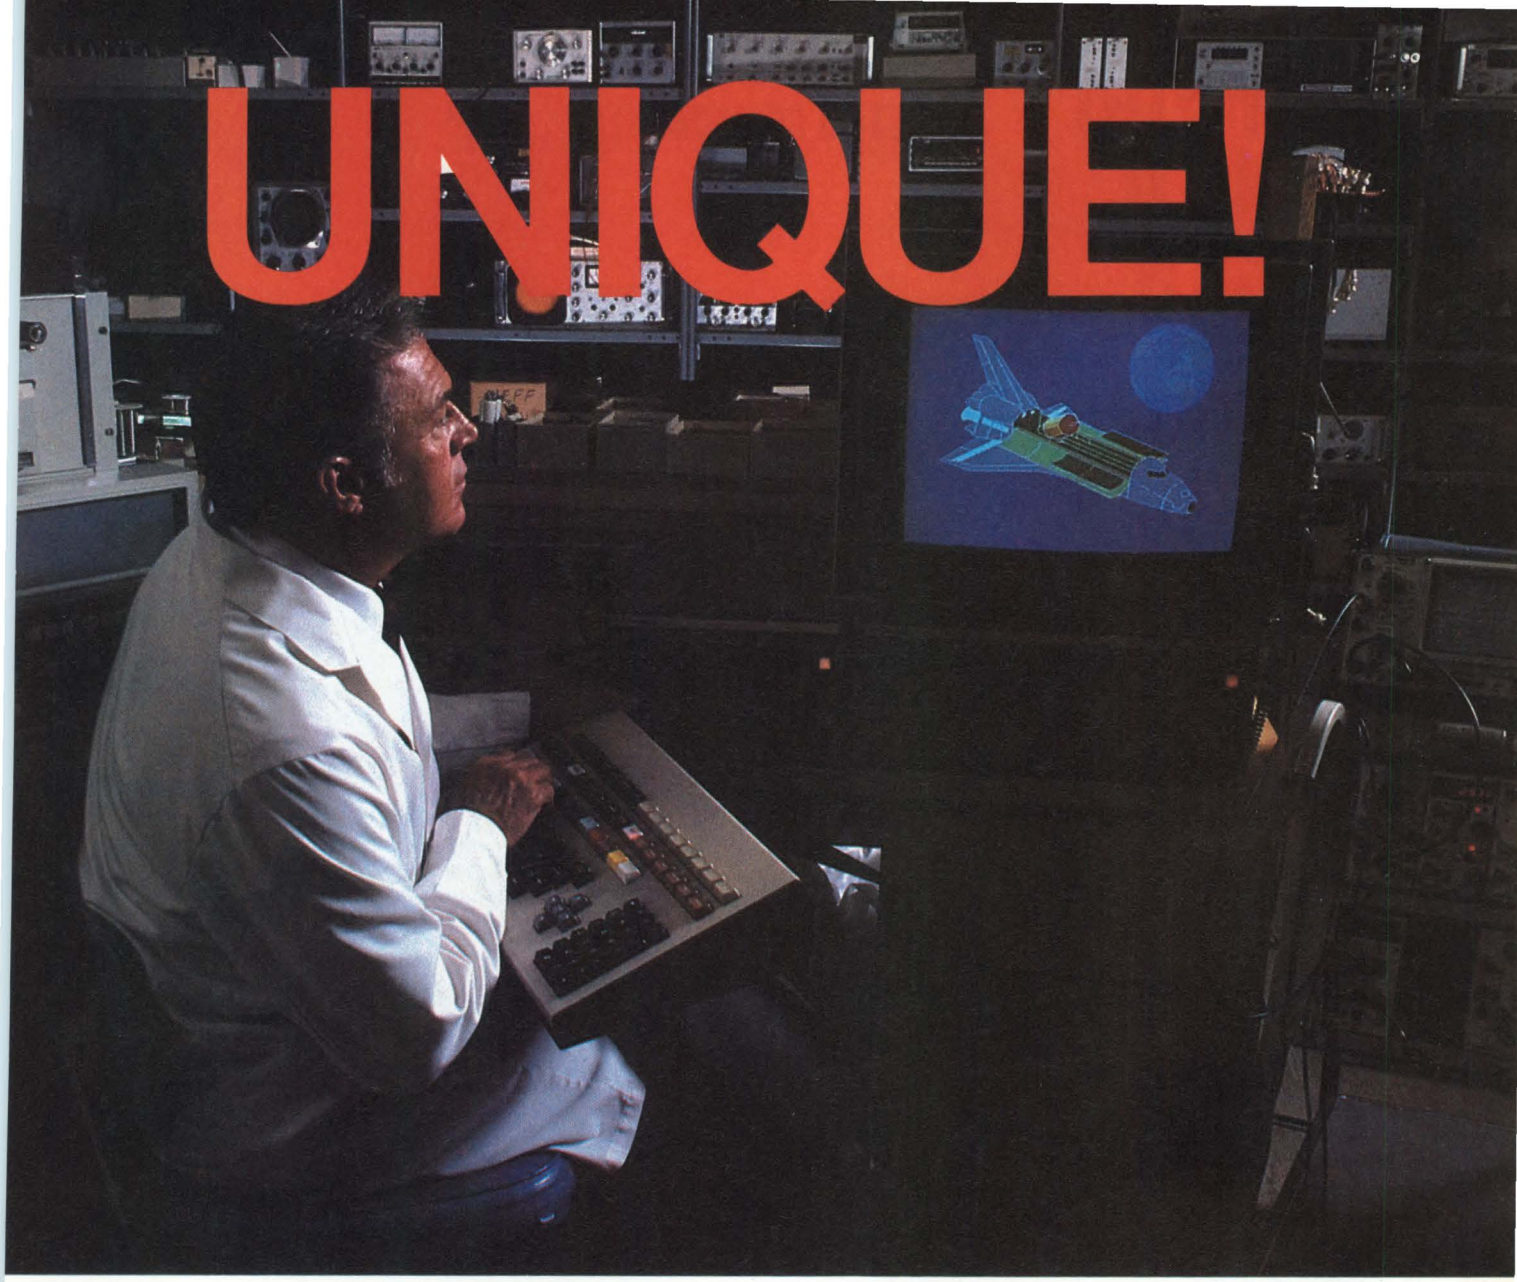

Photo courtesy Chromatics, Inc.

PUT YOUR PRODUCTS IN THE CENTER OF COMPUTER SYSTEMS DESIGN ACTIVITY

There is only one publication in the world that concentrates 100% on the design, integration, and test of computer based systems: *Computer Design.*

We define your market in terms of individuals . . . those engineering managers and engineers who design systems that incorporate computer products. These are the systems creators, the technical experts who put it all together at chip level, board level, and box level . . . who marry hardware and software and who *specify* everything that goes into these systems.

You reach more than 90,000 computer based systems designers — worldwide — in OEM companies, systems houses, consulting organizations and *Fortune 500* companies. No other computer and/or electronics publication penetrates this subscriber base by more than 50%.

Wherever computer based systems are designed or integrated, *Computer Design* readers are *the* key to specification and purchase!

Contact Your Local CD Sales Representative, or Call
Bob Billhimer, Marketing Director, Toll Free, at

1-800-225-0556

In Mass: (617) 486-9501

COMPUTER DESIGN

PennWell Publishing Company, Advanced Technology Group
119 Russell Street, Littleton, MA 01460

Touch-screen digitizer acts as data entry device

Series TF digitizers consist of a thin, transparent curved panel that mounts in front of a standard CRT display monitor, and an electronic board connected to the panel with a cable. When the panel is touched, the location of the touch is measured and sent to a host computer as an RS-232-C message. The digitizer acts as a data entry device and can replace a keyboard or graphic tablet. Devices include provisions for combining the serial data from the touch screen with the RS-232-C serial data from a terminal. Only one RS-232-C port is needed to support a terminal with a retrofitted touch screen. **TSD Display Products, Inc**, 35 Orville Dr, Bohemia, NY 11716.
Circle 292

Thick-film chip resistors marked by 1 percent tolerance

Series CRC thick-film chip resistors offer precision 1 percent tolerance. Rated at 1/8 W (70°C), the resistance of these devices ranges from 10 Ω to 2.2 MΩ, in tolerances of 1, 2, and 5 percent with a temperature coefficient (TC) of ±200 ppm/°C. The precision 1 percent model is available with a TC of 100 ppm/°C from 20 Ω to 300 kΩ. The CRC series is available in standard 0.126 x 0.063 in. (3.2 x 1.6 mm) size and is compatible with automatic placement equipment. Wraparound terminations that can be flow-soldered are provided. **Dale Electronics, Inc**, Dept 860, Box 609, Columbus, NE 68601. **Circle 293**

Wideband current amp has output buffering and increased load current

Available in an eight-lead TO-99 hermetically sealed package, the EH0002 wideband current amplifier sports load current capability up to 400 ms (pulse) for high current and cable drivers. Output buffering is also featured. The EH0002 is used in feedback loops without added system compensation. Output impedance is typically less than 10 Ω, and rise time is less than 7 ns. Slew rate reaches 200 V/μs. Price in quantities of 100 is $14 each. **Mel Tec**, 411 Providence Hwy, Westwood, MA 02090.
Circle 294

Small instrument CRTs optimized for CAD

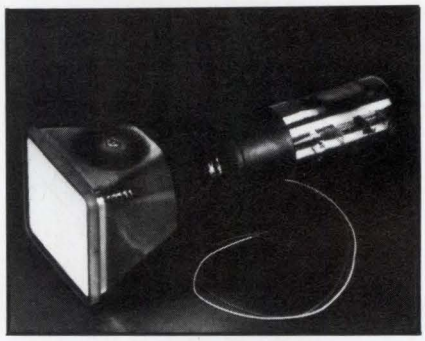

Series D14 instrument CRTs can act as compact oscilloscopes with bandwidths up to 150 MHz. These units' overall lengths measure 13.3 in. (338 mm). Optimized for CAD, the units are 5.5-in. (14-cm) post-deflection acceleration mesh tubes that deliver photographic writing speeds of 2 cm/ns. The first and final accelerating voltages of these tubes are 2.2 and 16.5 kV, respectively. Both have a 1.6 W quick-heating cathode. **Amperex Electronic Corp**, Providence Pike, Slatersville, RI 02876.
Circle 295

Small keyswitch saves front-panel space

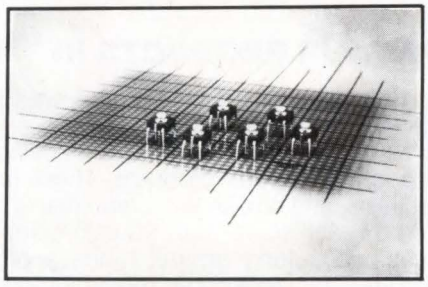

The KSA can withstand flow soldering processes and automatic cleaning operations. Its ultraminiature size—7.4 x 7.4 mm—saves front panel space. It is dimensioned in an IC package for use with existing automatic insertion IC machines and is designed with built-in straps. The KSA allows XY encoding that helps eliminate double-sided PC boards or complex layouts. Antistatic ground terminal is optional. Price in volume is $0.15. **ITT Schadow Inc**, 8081 Wallace Rd, Eden Prairie, MN 55344.
Circle 296

Why this magazine and more than 1,000 others let us go over their books once a year.

Some magazines, we're sorry to say, keep their readers undercover. They steadfastly refuse to let BPA (Business Publications Audit of Circulation, Inc.) or any other independent, not-for-profit organization audit their circulation records.

On the other hand, over 1,000 publications (like this one) belong to BPA. Once a year, BPA auditors examine and verify the accuracy of our circulation records.

This audit provides the name, company, industry and job title of every reader each publication reaches. The information helps advertisers to determine if they are saying the right thing to the right people in the right place.

It also helps somebody else important: you. Because the more a publication and its advertisers know about you, the better they can provide you with articles and advertisements that meet your informational needs.

BPA. For readers it stands for meaningful information. For advertisers it stands for meaningful readers. Business Publications Audit of Circulation, Inc.
360 Park Ave. So., New York, NY 10010.

BPA
MEDIA INTELLIGENCE

Gould Logic Analyzers

Introducing the K105. No other logic analyzer is so easy to use.

Logic analyzers have always been a bit complicated. Perhaps even intimidating to the occasional user. No more.

When you sit down at our new K105 logic analyzer, the first thing you'll notice is that big, friendly red HELP button. Press it. You'll begin to feel better immediately.

You see, we wrote the book on logic analyzers. And now the book is in the machine. So when you press the button, you display easy-to-understand, step-by-step operating instructions right across the bottom of the screen. While the data from the operation you're performing remains on the screen.

And if you're still in trouble, just press again. The HELP button and an adjacent SHIFT button call up a HELP MENU and 28 pages of detailed instructions on every analyzer function.

We'll say it again. No other logic analyzer is so easy to use.

Modular design accommodates application changes.

The K105 isn't just easy to use. It's accommodating, too. By simply swapping boards, you can configure several different logic analyzers.

For instance, you can select up to 64 20 MHz channels in 32-channel increments for microprocessor analysis. Up to 16 100 MHz channels in 8-channel increments for hardware analysis. Or combine them to a maximum of 72 channels for software/hardware integration tasks.

And there's more. You can add a dual 5¼" floppy disk drive (IBM CP/M 86™ compatible) to store up to 70 setups or data files. While providing data portability and post-processing capabilities.

Disassemblers and Trace Control™ speed software debugging.

It's a lot easier to debug software when you can get your system's microprocessor to speak assembly language mnemonics rather than object code. And our disassembly modules for the 68000, 8086, 8088, 8080, 8085A and Z80®B do just that.

And with the K105's 8 levels of Trace Control at 20 MHz, you can isolate and capture widely-separated slices of program flow to pinpoint failure causes... in a fraction of the time it would take with a conventional triggering scheme.

Two-analyzers-in-one enhances software/hardware integration.

When you're integrating hardware and software, the K105 is two analyzers in one. Just combine the 20 MHz and 100 MHz options to look at both state and timing. For trouble-shooting multi-processor systems, you can even monitor both processors and capture the asynchronous data between the two.

And the K105 offers a fast 5 ns glitch capture capability to pinpoint hard-to-find random problems.

Plus high-speed sampling for hardware analysis.

For high-speed sampling, you can configure the K105 with up to 16 100 MHz

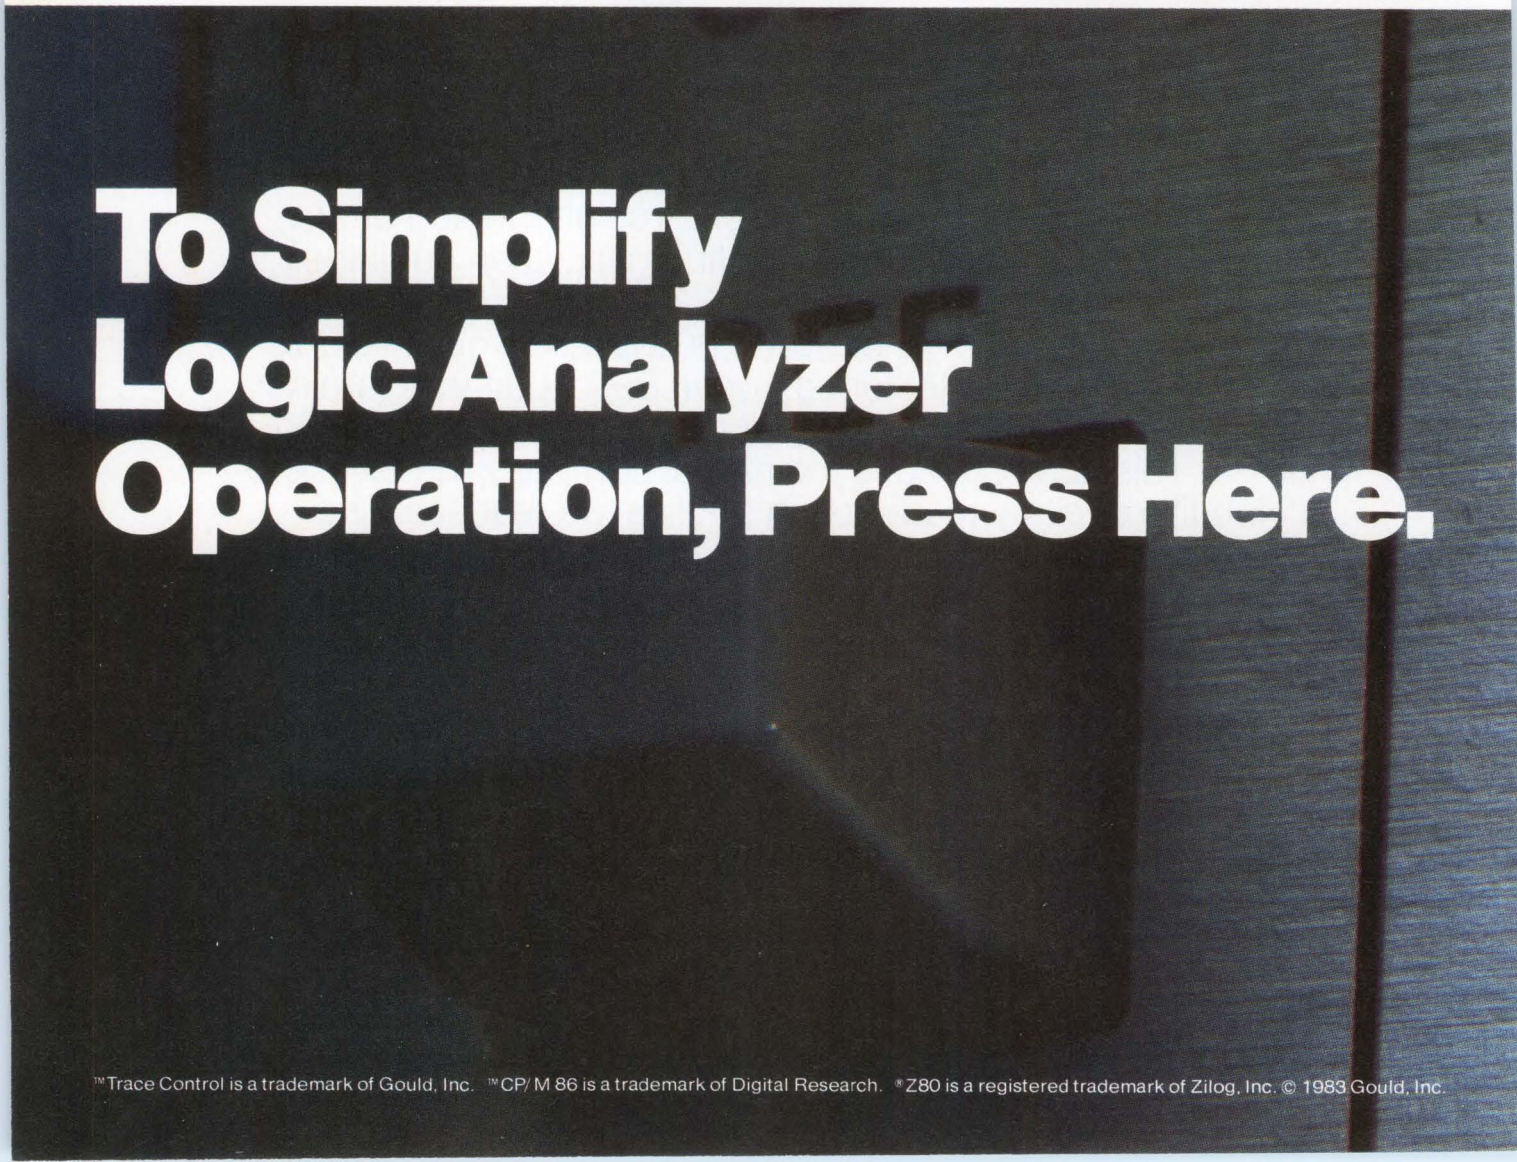

channels. Our unique automatic noise margin analysis feature enables you to verify specified system thresholds on as many as 16 channels simultaneously.

And design verification is simplified by a "don't care memory" that allows you to selectively mask out memory portions so you can compare only those portions you want to see.

Uncompromising dedication to high performance.

The Gould philosophy dictates that every instrument we make be the best for the job it's designed to do.

The K105, with its unsurpassed ease-of-operation and modular flexibility to perform a wide range of analysis tasks, is evidence of that dedication.

For detailed application notes or a demonstration, write Gould, Inc., Design & Test Systems Division, 4600 Old Ironsides Drive, Santa Clara, CA 95050-1279.

For fastest response, call toll-free: Nationwide (800) 538-9320; In California (800) 662-9231 or (408) 988-6800.

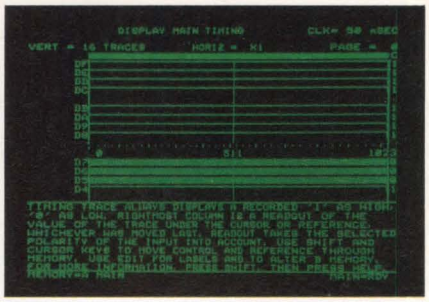

The K105 offers you two levels of HELP at the press of a button. The first displays step-by-step operating instructions across the bottom of the analyzer screen. The second brings a menu to the screen, allowing you to select more detailed help from an integral 28-page manual.

CIRCLE 130

THOMAS EDISON DIDN'T
HE JUST MA

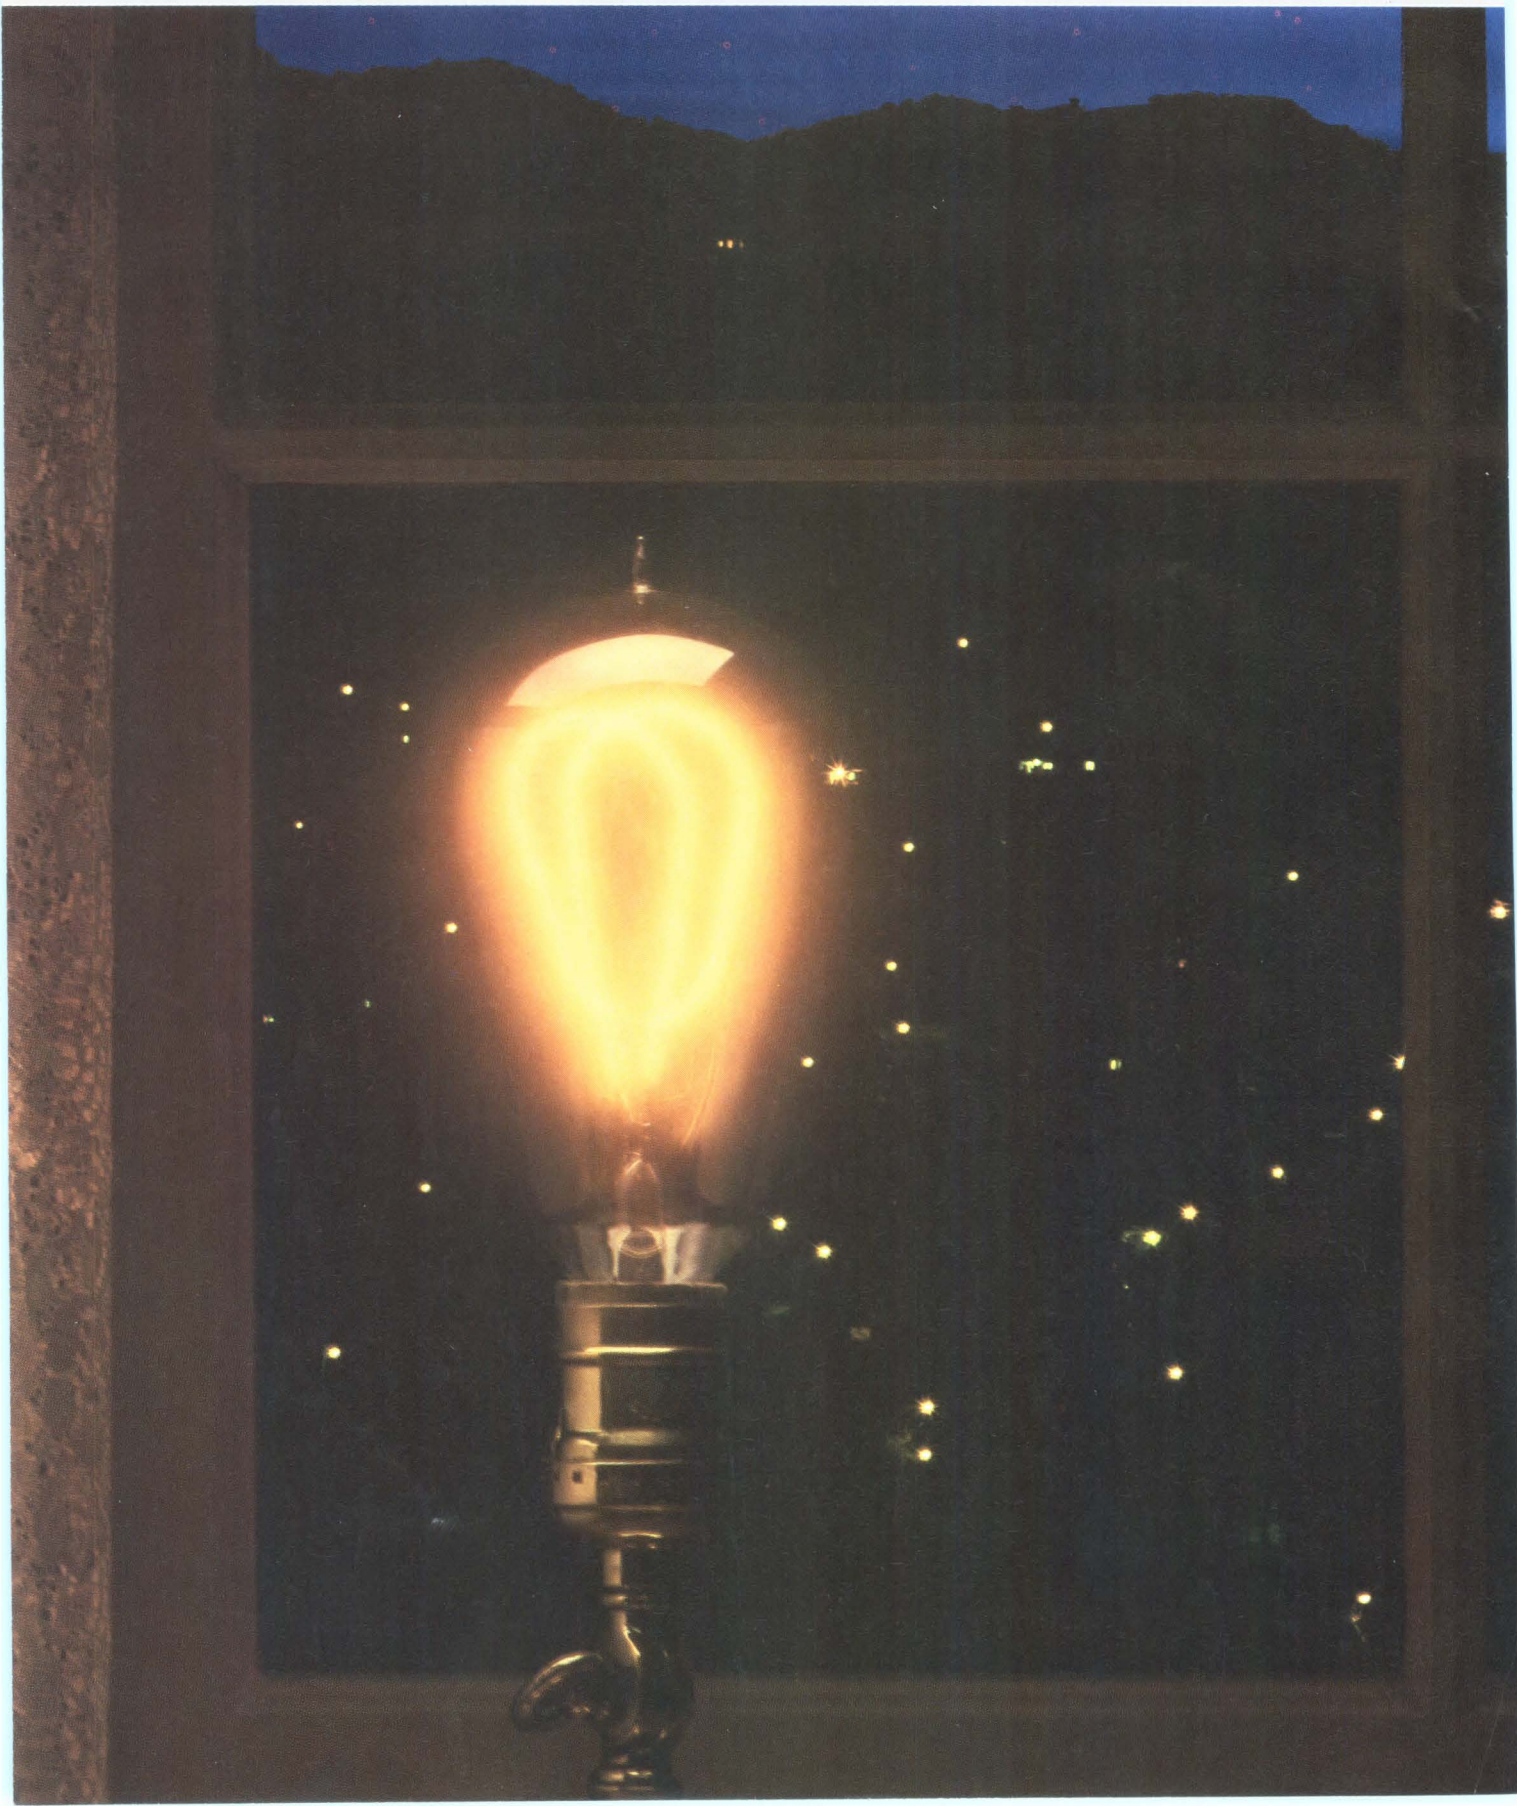

INVENT THE LIGHT BULB. DE IT WORK.

Fifty years before Edison, people were taking the air out of glass tubes, sending in current and hoping the filament would glow.

But in 1879, Edison attacked the bigger problem of how to distribute the electricity. So that by developing the power grid first, he was able to determine the precise combination of carbon rod, vacuum, voltage and current that his lamp would demand.

At Priam, we build our disc drives with a similar focus—on your system's requirements, and on your production schedule. It's a commitment that has made us the leading U.S. manufacturer of high-performance, high-quality 8" voice-coil Winchesters—from 35 to 500MB. A commitment we're now making to our new 86MB 5¼" drives as well. It's why we provide total, industry-standard interface support, including ST412, ANSI, SMD, SCSI, IPI-3, and our own Priam interfaces. And it's why we've just opened a new, automated production facility, one of the largest and most modern in the OEM world.

Priam. For high-performance disc drives, check with us first. We'll make sure it works.

PRIAM®

JUST THE RIGHT DISTANCE FROM THE LEADING EDGE.

20 West Montague Expwy., San Jose, CA 95134 (408) 946-4600
Boston (617) 444-3973; U.K. (44) 734-884-788

CIRCLE 131

Portable EPROM/EEPROM programmer handles devices up to 512 Kbytes

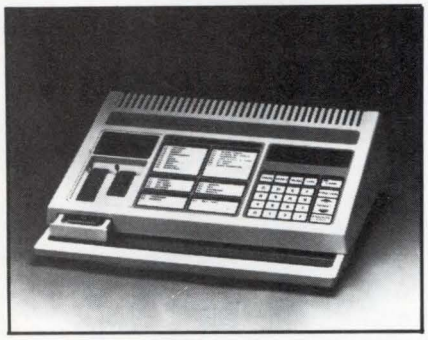

The EP-824 portable programmer handles EPROMs and EEPROMs, as well as Intel single-chip microprocessors with capacity of up to 512 Kbytes. Dual sockets ensure that devices can be programmed without device adapters. Interactive software simplifies programming by guiding the user through each state of operation in plain English. Frequently used functions are single-keystroke activated. Internal software for the EP-824 is written in PL/M or 8085 assembly language. The built-in 8085A microprocessor works at 3 MHz without wait state. Bus drivers buffer address and data lines. Price is $1995. **Wavetek-Digelec**, 586 Weddell Dr, Sunnyvale, CA 94089. **Circle 297**

Simulator/analyzer checks switching systems

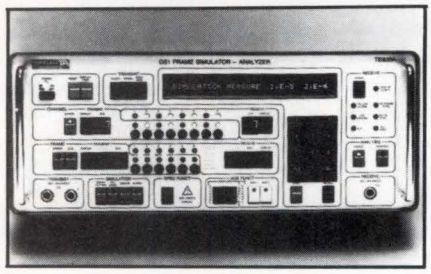

The TE-802A simulator/analyzer tests No. 5 Electronic Switching Systems, as well as a wide range of DS1-T1 equipment. It uses advanced error detection to pinpoint sources of distortion and can simulate both error-free and controlled error conditions. This simulator/analyzer can analyze pulse code modulation signals to confirm switch output and determine error rate signals from a switch or remote user. The TE-820A weighs under 40 lb. Prices start under $9000. **Tekelec, Inc**, 2932 Wilshire Blvd, Santa Monica, CA 90403.
Circle 298

A 32-bit CAD system uses advanced graphics tablet

Artech is a CAD and drafting system that includes a Hewlett-Packard 32-bit computer. It can be connected in a network, displays 2-D and 3-D images on a 19-in. color screen, and enables users to enter operational commands via a graphics tablet containing more than 300 directly accessible instructions. It can be upgraded with an additional graphics processor that allows dynamic graphics image manipulation. The basic workstation is priced at $27,000; fully configured, it costs $59,500. **Skok Systems Inc**, 222 Third St, Cambridge, MA 02142.
Circle 299

High sensitivity measurement displayed by switching matrix

The HP 4085M switching matrix, a dedicated subsystem created for use with the HP 415A semiconductor-parameter analyzer, provides measurements of 1 pA and 1 mV at each of 48 pins on a device under test. The HP 4085M consists of a switching matrix controller and switching matrix. The latter has eight ports that are fully matrixed to 48 output pins. Under program control, any of the eight instrument ports can be switched to any device under test pin. Included software uses high level commands for executing port-to-pin interconnections and diagnostics. Status can be viewed on a system controller such as the HP 9000 Model 216S. The HP 4085M is $34,109. The HP4145A is $17,500. **Hewlett-Packard Co**, 1820 Embarcadero Rd, Palo Alto, CA 94303.
Circle 300

System integrates layout tools for printed circuit board design

As an end-to-end design solution, T-boards use the same tools throughout the design process. All users in the schematic capture and design cycle work with a common data base on the TEGAStation. The engineer can freeze logic so the designer cannot change it, yet the designer can still change the topological layout. The module performing the actual PC board layout is interactive. Features include precision of one mil, board size of 32 x 32 in., a 32-bit virtual programming environment, and a realtime correct by construction design. Price is approximately $100,000 per workstation. **Calma Co**, 2901 Tasman Dr, Santa Clara, CA 95050. **Circle 301**

Socket boards feature Z80 compatibility

Model 4-Z80 socket boards, compatible with the Z80 microprocessor bus, are offered in universal style patterns, with six available combinations of power and ground planes. The 4-Z80s accept the equivalent of 77, 16-pin DIPs in sizes that include 0.3-, 0.4-, 0.6- and 0.9-in. pin spacing. Their price is $266 each. **Hybricon Corp**, 410 Great Rd, Littleton, MA 01460.

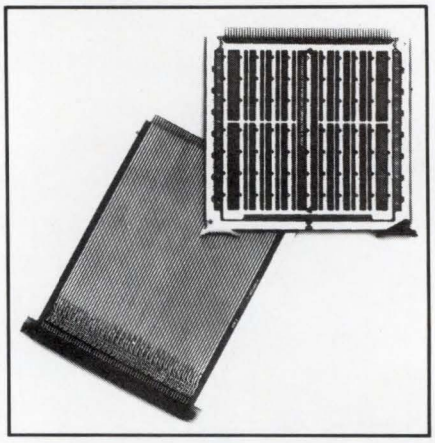

Circle 302

Probe option records glitches down to 10 ns

Option 92 probes record logic glitches down to 10 ns, extending the timing features of the Series 200 logic analyzer by allowing it to capture random logic pulses occuring between clock edges. The Option 92 gives four glitch detection channels and four logic analysis channels. Glitch channels detect (and provide CRT display of) multiple transitions through a preset voltage threshold within a single clock period. Voltage thresholds adjust between -5 and 5 V and clock settings between 0.1 μs and 50 ms. The probe option sells for $695. **Racal-Dana Instruments Inc**, 4 Goodyear St, PO Box C-19541, Irvine, CA 92713.

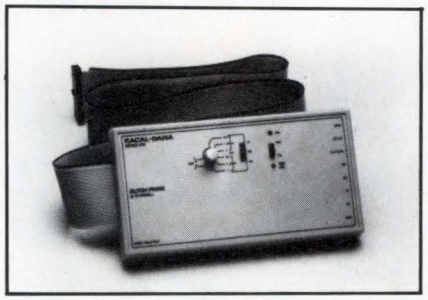

Circle 303

COMPATIBILITY

NCR 40-column printers... fit in perfectly with your design. NCR's highly compatible printers are easily adaptable to slip printing, data logging, or receipt or journal printing. It's the only family of printers, standalone and basic mechanisms, guaranteed to meet the various needs of the OEM marketplace. Compact, easily installed and highly reliable, NCR printers continue to prove their worth with over one million units in use throughout the world. Getting more information is no puzzle. You just call (513) 445-7443.

NCR
OEM MARKETING DIVISION

NCR 40-column printers...designed for demanding applications.

CIRCLE 132

Start/Stop And Stream.

INTRODUCING

DOUBLE DRIVE™

Series 50 Drive. High Performance. High Capacity.

If you've been pondering whether you need a streaming tape drive to back up your Winchester, or a start/stop tape drive, or maybe both, we have good news. You no longer have to decide. Because now there's a drive with the capabilities of both. Rosscomp's Double Drive. It saves you the agony of decision. And an incredible amount of money at the same time.

There Are Two Sides To Our Story

When our engineers set out to design the Double Drive, they too had a decision to make. Quarter-inch tape? Or half-inch tape? But then they found that, even by pushing quarter-inch technology to its absolute limit, they wouldn't be able to give us the capabilities we wanted.

So half-inch became the only choice. That also meant they could assure us of something else we demanded. Reliability. For instance, they could design-in the industry's highest quality read/write half-inch heads. Meanwhile, they could make the system stream at unbeatable speeds.

And do practically everything a large start/stop machine can do for a fraction of the cost.

Side One

One of the fruits of their labors is, you get a Winchester backup that does something no other can. It scoops 190 MBytes of data onto standard half-inch tape that's housed on our Back-Pac™ four inch self-threading reel. (Which costs less than a cartridge.)

That translates into the lowest cost per byte of any tape drive, anywhere.

And it streams at a blazing 90 ips standard.

In fact, the Double Drive's electronic speed and capacity are matched by only one thing we know of. Its mechanical reliability.

The Back-Pac™ self-threading reel stores 190 MB on ½" tape. Costs less than ¼" cartridges.

That comes from its having the simplest and most dependable tape-handling method yet devised.

Threading is automatic. The tape travels a path of proven design, kept on course by both spring-loaded and ball-bearing guides.

Spring loaded and ball bearing tape guides.

Tension control is equally precise. A single capstan motor drives both supply and takeup reels with one belt, while an advanced controlling device maintains identical tension between them. So you experience very little ISV. As well as lower power dissipation. The motor uses less than 15W.

Simple, dependable automatic tape threading.

Single capstan motor.

Side Two

64 KB cache memory in its interface makes it look like a start/stop drive.

For your selective file-search and backup jobs, the Double Drive acts as a start/stop unit. Because it has 64 KBytes of cache memory right in its interface.

That means no time loss. An optimized I/O. And apparent real-time operation for every user.

Plus more efficient use of your tape, covering a full 95% of it.

Online, it gives ample secondary storage for files during restructuring and in hierarchical memory management schemes.

Now, if the Double Drive sounds like it might be just what you need, here's more good news. It will fit your system.

Or rather, one of our Double Drives will. Our Series 80 fits an 8" envelope, our Series 50 a 5¼" envelope.

Series 80 to right. Series 50 below.

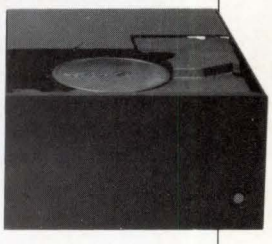

They'll each hold 190 MBytes of data, stored on 24 serpentine tracks.

Both are available with standard industry interfaces. So one of them is right for you.

Which means there's no painful decision to make. Just put the Double Drive to work. Doing double duty.

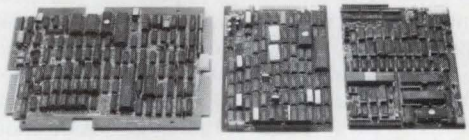

Available formatters include QIC-02, SCSI and 9-track.

NASDAQ-RCOM
1695 MacArthur Blvd., Costa Mesa, California 92626 · Call (714) 540-9393

ROSSCOMP®

The way with tape.

CIRCLE 133

Printed circuit board system enables efficient routing

The Circuit Master fully interactive PC board CAD system performs complete or partial routing, either manually or automatically. A Wave Tracer feature alerts the designer to connection obstructions. The system's Demand Router feature and automatic placement, plus gate and pin swapping, rat's nest, and density histogram routines, reduce design time by optimizing IC placement. Circuit Master provides output to photoplot a positive 1:1 or scaled artwork of the various board layers; basic system costs $69,950. **Bishop Graphics CAD Systems Corp**, 5388 Sterling Center Dr, PO Box 5007, Westlake Village, CA 91359.

Circle 304

In-circuit testing of custom ICs is possible with offline software

Convert is an offline software tool that allows the model of a custom logic array, generated for the Sentry series of device testers, to automatically convert into a model that can be used for in-circuit testing by any series 333 machine. Simulation routines created for design verification become the test programs that check the proper functioning of the device during in-circuit testing. **Fairchild Manufacturing Test Systems**, 299 Old Niskayuna Rd, Latham, NY 12110.

Circle 305

Emulator package converts PC into universal development system

An emulator port turns the IBM PC/XT into a universal development system for the design, test, debugging, and implementation of popular microprocessor hardware and software. The KPCI package consists of hardware and a set of development support software tools that include a cross-assembler, linker, emulator software, and additional CP/M utilities. It runs on the IBM PC/XT. Emulator subsystem and Pascal compilers are optional. When application programs have been completed, the Pascal compiler or assembler can process the edited source code. Source text processors develop object code for microprocessors including the Z80, 68000, 8086/88, and 80186/188. The KPCI is available for $1500. **Kontron Electronics**, 630 Price Ave, Redwood City, CA 94063. **Circle 306**

Disassembler enhances 68000 logic analysis

Option 68K provides disassembler support for the 68000 16-bit microprocessor. Using this option, models 202 and 205 48-channel logic analyzers connect to 68000-based circuits, for hardware and software disassembly. With the 68000 operating at its maximum speed of 12.5 MHz, the 68K qualifies data in four ways: 68000 and DMA, or 68000 instructions only; program and/or data; supervisor and/or user; and clock qualifier probe. The hardware adapter, in addition to providing interconnection between the logic analyzer and the 68000, automatically adjusts the setup parameters to match 68000 requirements. The Option 68K disassembler sells for $1950. **Racal-Dana Instruments Inc**, 4 Goodyear St, PO Box C-19541, Irvine, CA 92713. **Circle 307**

COMPUTERS

Slave processors are key to multi-user PC operation

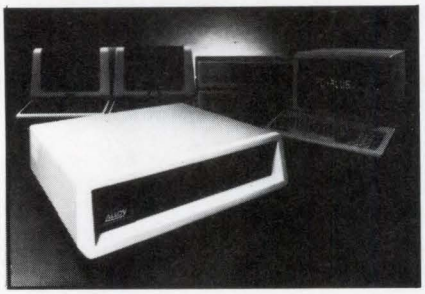

Terminals can be added to multi-user IBM PC systems using the PC-PLUS software-hardware combo. A two-user system begins with a terminal and a PC-Slave processor board plugged into a PC expansion slot. Up to 12 users can be accommodated via the PC-XBUS expansion module, which plugs into the host PC and holds 8- and 16-bit processors. With 3 such modules, up to 30 processors can share data, software, and peripherals under the direction of the host PC. Prices for the devices are $1195 for the PC-Slave/16, $995 for the PC-Slave/8, and $1495 for the PC-XBUS. **Alloy Computer Products**, 100 Pennsylvania Ave, Framingham, MA 01701.

Circle 308

Supermini offers SNA/SDLC interface and extended power

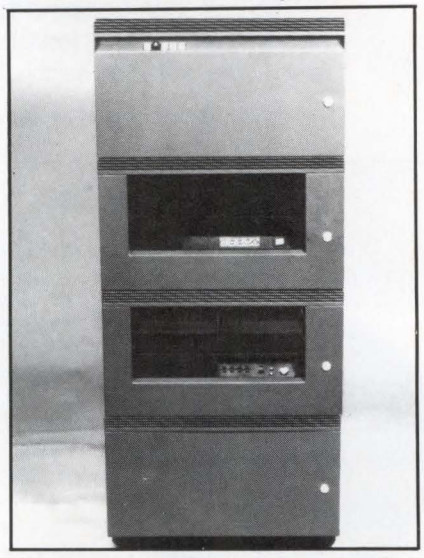

The 8855 model 10 gains fast processing power through a 16032 main processor chip and a coprocessor using 80186 chip technology. Up to 64 tasks can simultaneously run batch and interactive communications, SNA/SDLC and bisynchronous protocols, local file processing, plus local database inquiry and update. Up to 2 Mbytes of real memory and 132 Mbytes of disk storage are included. An auxiliary cabinet can hold three 132 Mbyte disks. Available later this year, the model 10 costs $97,500. A field upgrade of the 600 series is also available. **Nixdorf Computer Corp**, 300 Third Ave, Waltham, MA 02154. **Circle 309**

Pipeline processing handles heavy batch work loads

An advanced architecture that accommodates both heavy batch work loads and high volume transaction processing systems—the B 4925—supports 5 Mbytes of memory. This breaks down to a basic unit that is available with 2.5 Mbytes, with one add-on module of 2.5 Mbytes. The B 4925 architecture features distributed pipeline processing.The CPU is complemented by high speed peripheral subsystems, multiple I/O channels, and frontend processors. It supports the Shared System Processor for multiprocessor configurations. It is compatible with the B 2000, B 3000, and B 4000 systems for user programs written in languages including Cobol 74, Fortran 77, RPG II, Basic, and Pascal. **Burroughs Corp**, Detroit, MI 48232. **Circle 310**

"We couldn't manage without Business Computing."

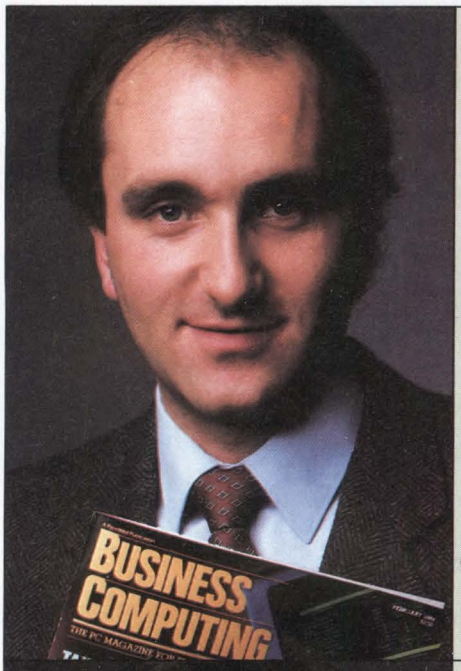

"Outlined vendor options for expanding our system."

Rolf Grueniger
PRESIDENT AND GENERAL MANAGER
Skytruck International Airfreight, Inc.

"Great overview for a manager trying to keep up."

Bill Fowkes
MARKETING MANAGER
Home Box Office, Inc.

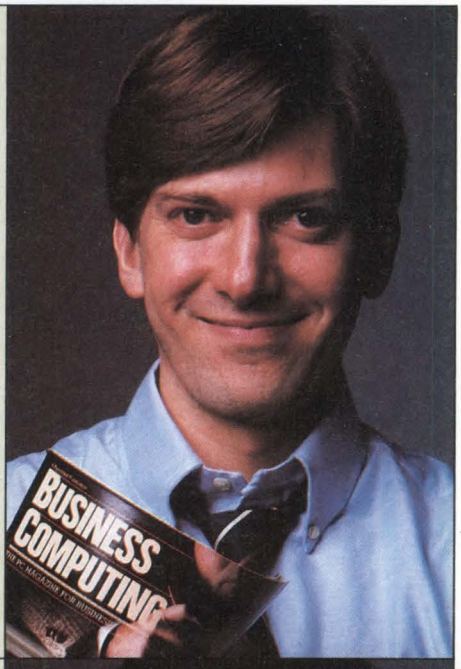

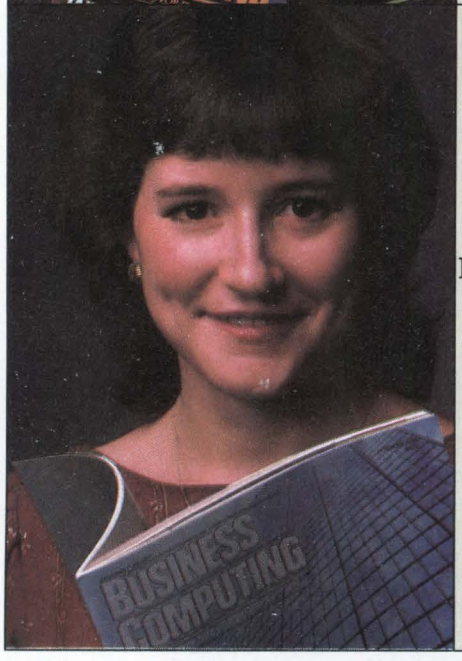

"Saved me hours of research on spread-sheet software."

Karen Leahy
PRODUCTION MANAGEMENT ASSOCIATE
Buckley and Urbanski

"Thanks for the sound advice on word processing."

Jim Hill
MANAGING EDITOR
Hill Publishing

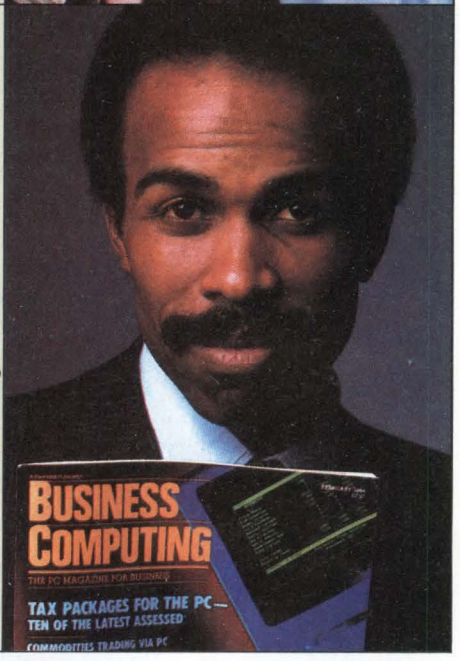

Now there's a new magazine dedicated to helping you manage more profitably with your IBM PC. No games. No jargon. No fooling.

Instead, BUSINESS COMPUTING delivers executive-level wisdom every month on how to use the latest software and peripheral technology to make sharper decisions and solve management problems.

With step-by-step tutorials, sample applications and pertinent case histories. As well as guidance from managers like you, who share their own trial and error experiences so you won't have to.

BUSINESS COMPUTING. It's the first non-technical magazine for managers about using PCs for profit.

Try BUSINESS COMPUTING now. As a Charter Subscriber, 12 issues cost you only $14.75, over $9.00 off the regular subscription price. And BUSINESS COMPUTING is guaranteed to help you or your money back. Just use the reply card attached or call toll-free 1-800/922-4800. (In Utah, call 1-800-662-2500.)

BUSINESS COMPUTING

For the PC user who isn't playing games.

A Pennwell Publication/119 Russell Street, Littleton, MA. 01460

To order, circle Reader Service No. **181**

Kiss 2 static RAMs,

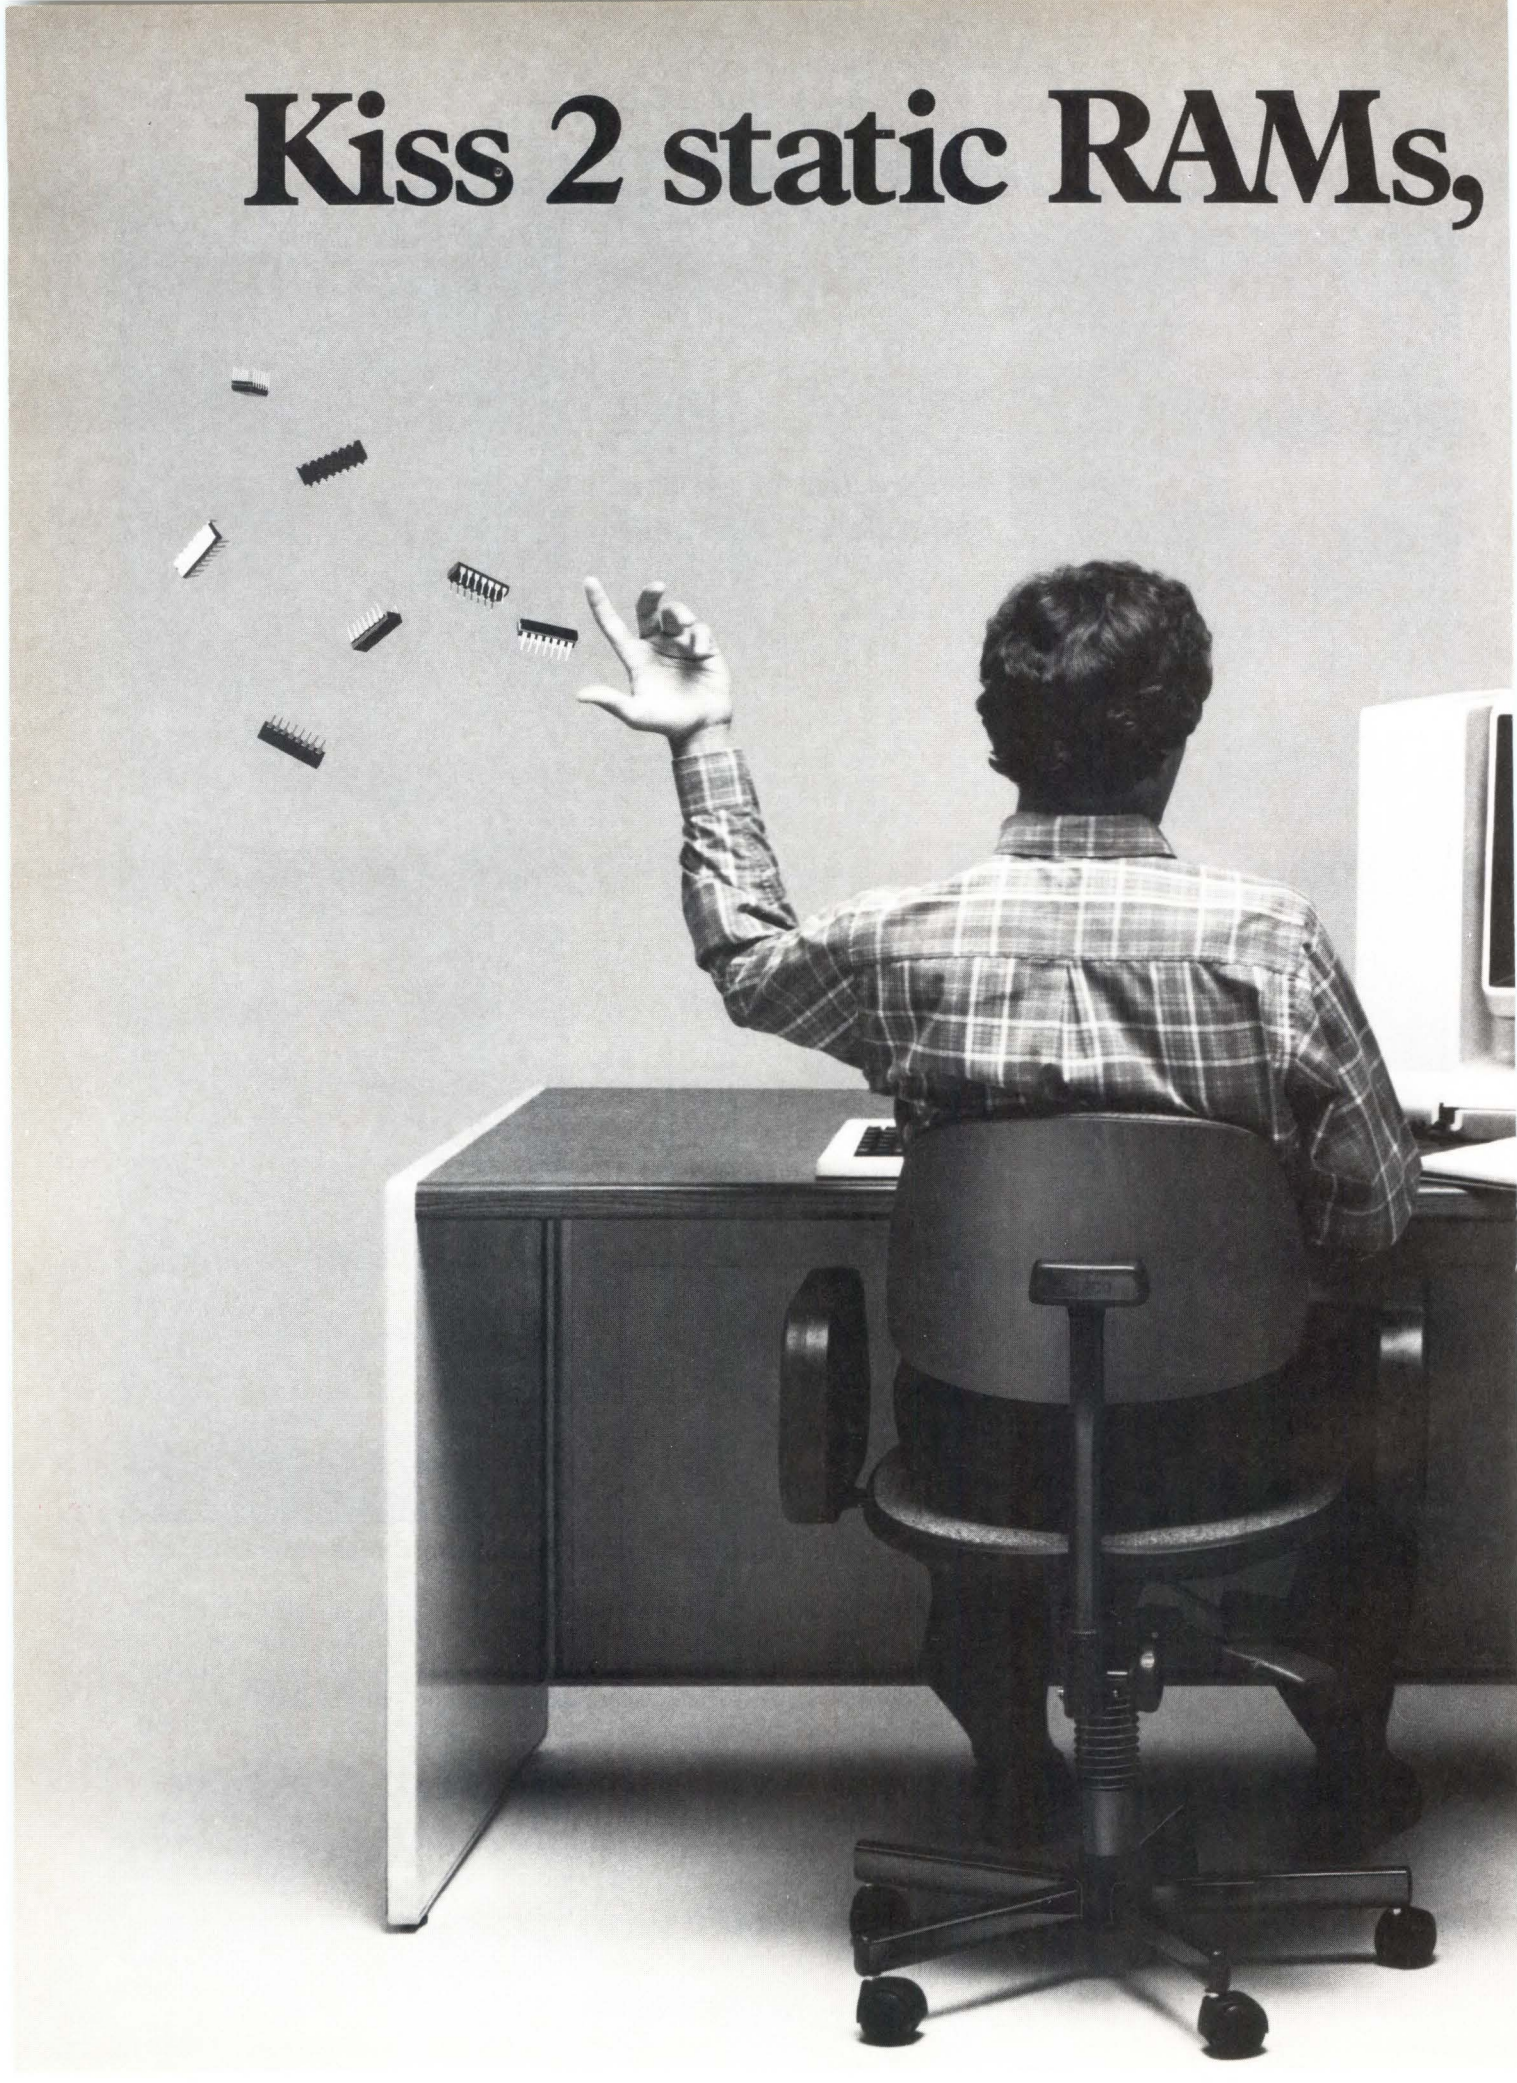

2 registers, 2 counters and a multiplexer goodbye.

And say hello to the world's first 35 MHz FIFO.

It's the only 64x5 FIFO with the speed to handle high speed data transfer applications.

Like disk and tape controllers. Digital video. Data communications. Multi-processor environments. And so on.

And it's the only 64x5 FIFO with the high drive data outputs to knock all those other parts off the board.

With three-state outputs, it's perfect for bus-oriented systems.

And it even has automatic status flags—input ready, output ready, half full, and almost full/empty.

Now, there's just one more thing we'd like you to think about. What delicious things can you do with all that empty board space?

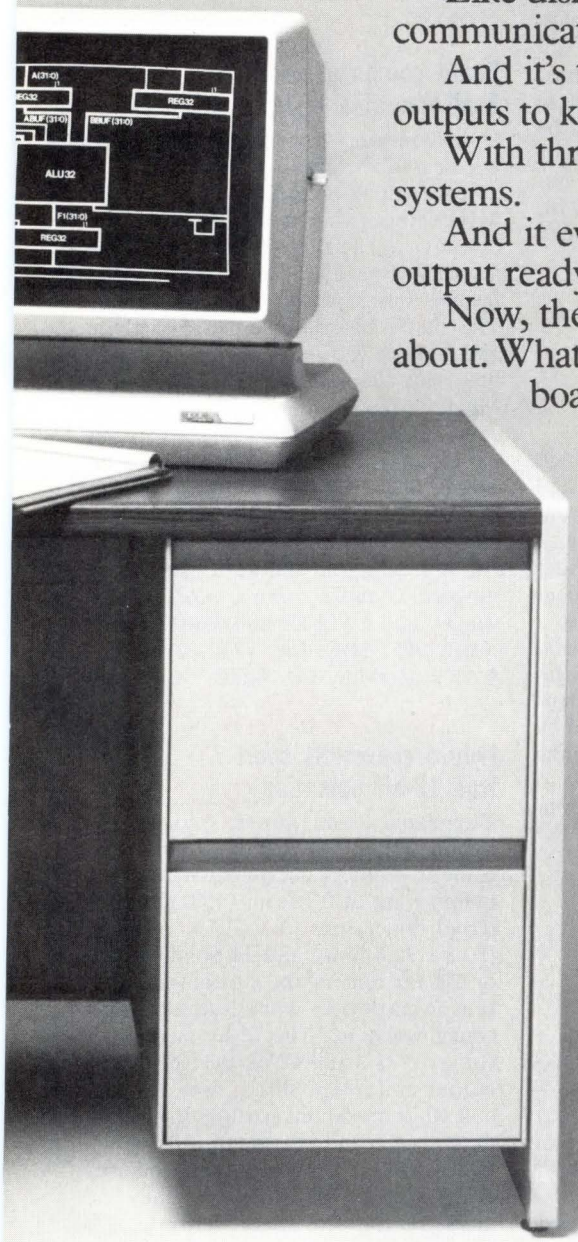

For more information on the first of our new fast FIFOs (part no. 67413A), call your local Monolithic Memories representative or franchised distributor and ask for our Fast FIFO Data Sheet Pack. Or write Advanced Logic Product Marketing, Monolithic Memories, 2175 Mission College Blvd., Mail Stop 8-13, Santa Clara, CA 95054.

You'll receive a 1984 Systems Design Handbook free.

Monolithic Memories MMI

CIRCLE 134

Desktop units speed IBM PC software execution

Using the 8086 microprocessor, Deskpro computers run IBM PC and PC/XT software without modification. The units run standard software two to three times faster than the PC. A dual-mode monitor allows execution of high resolution PC-compatible text and graphics on the same screen. The entry level system, one of four models available, includes a single 360-Kbyte disk drive and 128 Kbytes of RAM. Top of the line model 4 has 640 Kbytes of RAM, one 360-Kbyte disk drive, a 10-Mbyte fixed disk drive, and internal 10-Mbyte fixed disk drive backup system, a dual-mode monitor, and an asynchronous communication/clock board. Prices range from $2495 for model 1 to $1795 for model 4. **Compaq Computer Corp**, 20333 FM 149, Houston, TX 77070.
Circle 311

Imaging system performs edge detection and windowing

Based on IBM's 9000 lab computer, the QX-9000 benchtop combines realtime image processing with image analysis and graphics. It can be configured using Quantex VERSAbus boards, and will run user-written C, Pascal, Basic, and Fortran programs. Capabilities include background correction, averaging, summing, comparing, windowing, and edge detection. System hardware includes VERSAbus boards, the 68000-based multitasking computer, 5¼-in. floppy disk drive, three RS-232 ports, plus one IEEE 488 and one Centronics port. Complete systems start at $30,000. **Quantex Corp**, 252 N Wolfe Rd, Sunnyvale, CA 94086.

Circle 312

MICROPROCESSORS/MICROCOMPUTERS

Distributed realtime multiprocessing system uses C language

A VMEbus CPU board—the PG2020—represents the first implementation of the distributed realtime multiprocessing (DRM) system. The DRM system creates a virtual processing environment. Using the standard C language, the number of 68000 processors used in a DRM system can range from a few to hundreds, all processing information in a distributed fashion and working as one unit. The PG2020 will support single-processor 68000 systems, device handlers for printers, and remote VAX/VMS. Application software can be developed on the DEC VAX-11/VMS or Philips microcomputer development system. **Philips, Elcoma Div**, PO Box 218, 5600 MD Eindhoven, the Netherlands.
Circle 313

Unix-based supermicros accommodate 16 users

System 100 and System 300 feature high speed operation, 10-MHz 68000 CPU, and cache memory for disk operations. Context switching from user to user or task to task in a single instruction, provides fast response. A memory management board maps the logical memory onto a physical RAM array in 4-Kbyte pages and can divide tasks into noncontiguous pages for more effcient memory use. Standard is a 50-Mbyte Winchester, controlled by the STDC controller. The systems use the IEEE 696 bus. Prices range from $9995 to $19,995, depending on configuration. **Cromemco, Inc**, 280 Bernardo Ave, PO Box 7400, Mountain View, CA 94039.
Circle 314

Multibus-based boards form color graphics unit

The RG-10C Smart RGB Color Graphics Generator uses two Multibus boards—the RG-VG5 vector generator and the RG-M7 refresh memory board. The machine provides eight colors, expandable to 4096, uses its own processor, and offers 8 Kbytes of dual-port instruction RAM. It features a built-in debug trace. Variables and display opcodes can be monitored as they are processed. The RG-10C draws lines from endpoint coordinates at 800 ns/pixel. It handles interlaced resolution of 512 x 480 pixels and noninterlaced resolution of 512 x 240 pixels. Price in 100s is $1295. **Raster Graphics, Inc**, PO Box 23334, Tigard, OR 97223.
Circle 315

Board combo accesses 8 Mbytes at a 10-MHz rate

Unique memory management allows the 68000-based, Multibus M68CPU board to access companion memory boards at 10 MHz without wait states for read/write or refresh. The M68CPU connects to the M68MEM board using a high speed memory bus, separate from the Multibus. The DRAM boards come in 512-Kbyte and 2-Mbyte versions, and up to four can link with the M68CPU for a total of 8 Mbytes. Memory management hardware contains 32 independent maps. Custom I/O expansion is supported by two iSBX connectors with 8- and 16-bit, single-double-, and triple-width iSBX Multimodules. Software support includes a PROBUG debugger and REGULUS, which supports C and Fortran compilers. The M68CPU and a 512-Kbyte M68MEM board cost $2995. **SBE, Inc**, 4700 San Pablo Ave, Emeryville, CA 94608. **Circle 316**

Hybrid converters sport true 12-bit operation

Designated series MN5610 and MN3660, hybrid 12-bit A-D and D-A converters come in leadless packages. With maximum height of 0.183 in. (4.64 mm), external dimensions of 0.810 x 1.315 in. (20.5 x 33.4 mm), and 24 bonding pads on 100 mil centers, the A-D converters are full 12-bit devices with 13-μs maximum conversion time. The 12-bit MN3660 D-A converter is a full 12-bit device with an output amp, high speed input register, and 10-μs maximum settling time. Standard MN5610 and MN3660 are fully specified for 0 to 70°C operation. Quantity prices start at $99 for the 5610 family and at $87 for the 3660. **Micro Networks**, 324 Clark St, Worcester, MA 01606.
Circle 317

This New Fiber Optic Modem will Extend a DCE Interface to Any Point in Your Local Area Network.

Plus a whole lot more.

- **Can also be used for standard modem applications**
- **Automatically accepts or supplies DCE/DTE clocks**
- **Fully supports all EIA handshaking signals**
- **Provides secondary data channel**

In short, you can use our new fiber optic modem between *any* two plug compatible units in your local area network. And it won't require any jury-rigging or looping clock and interface signals. That's because, from an operating standpoint, our fiber optic modem looks just like an EIA cable; whether you're going from a long-haul modem to a remote terminal or from a CPU port to a printer. And it's just about as easy to install as a cable — we even provide two separate connectors (DTE and DCE) on each modem. YOU determine how our modem will function simply by selecting which connector you use!

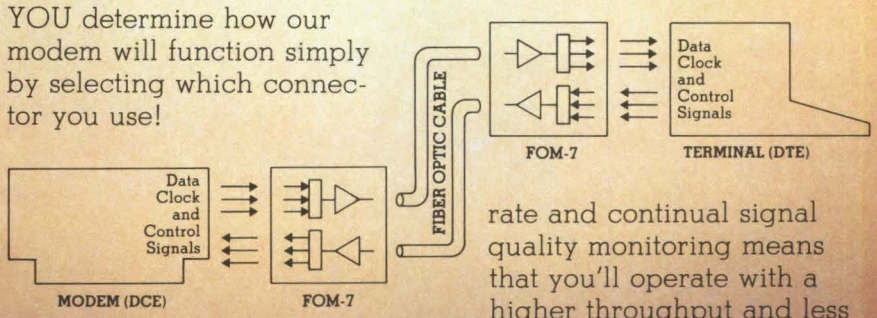

Once our fiber optic modem is installed and operating it'll really begin to shine. You'll benefit not only from the advantages inherent with fiber optics (traffic security plus noise immunity) but also from the exceptional operating performance. Our very low error rate and continual signal quality monitoring means that you'll operate with a higher throughput and less downtime than ever before.

Versitron manufactures a complete line of fiber optic products for Local Area Networks. Our 20 years' experience in fiber optic is reflected in the performance capabilities of our products.

For the full story on how our complete line of fiber optic products will solve your Local Area Network problems, give us a call at (202) 882-8464. Or write:

CIRCLE 135

VERSITRON, INC.

6310 Chillum Place, N.W., Washington, D.C. 20011. TWX: 710-822-1179

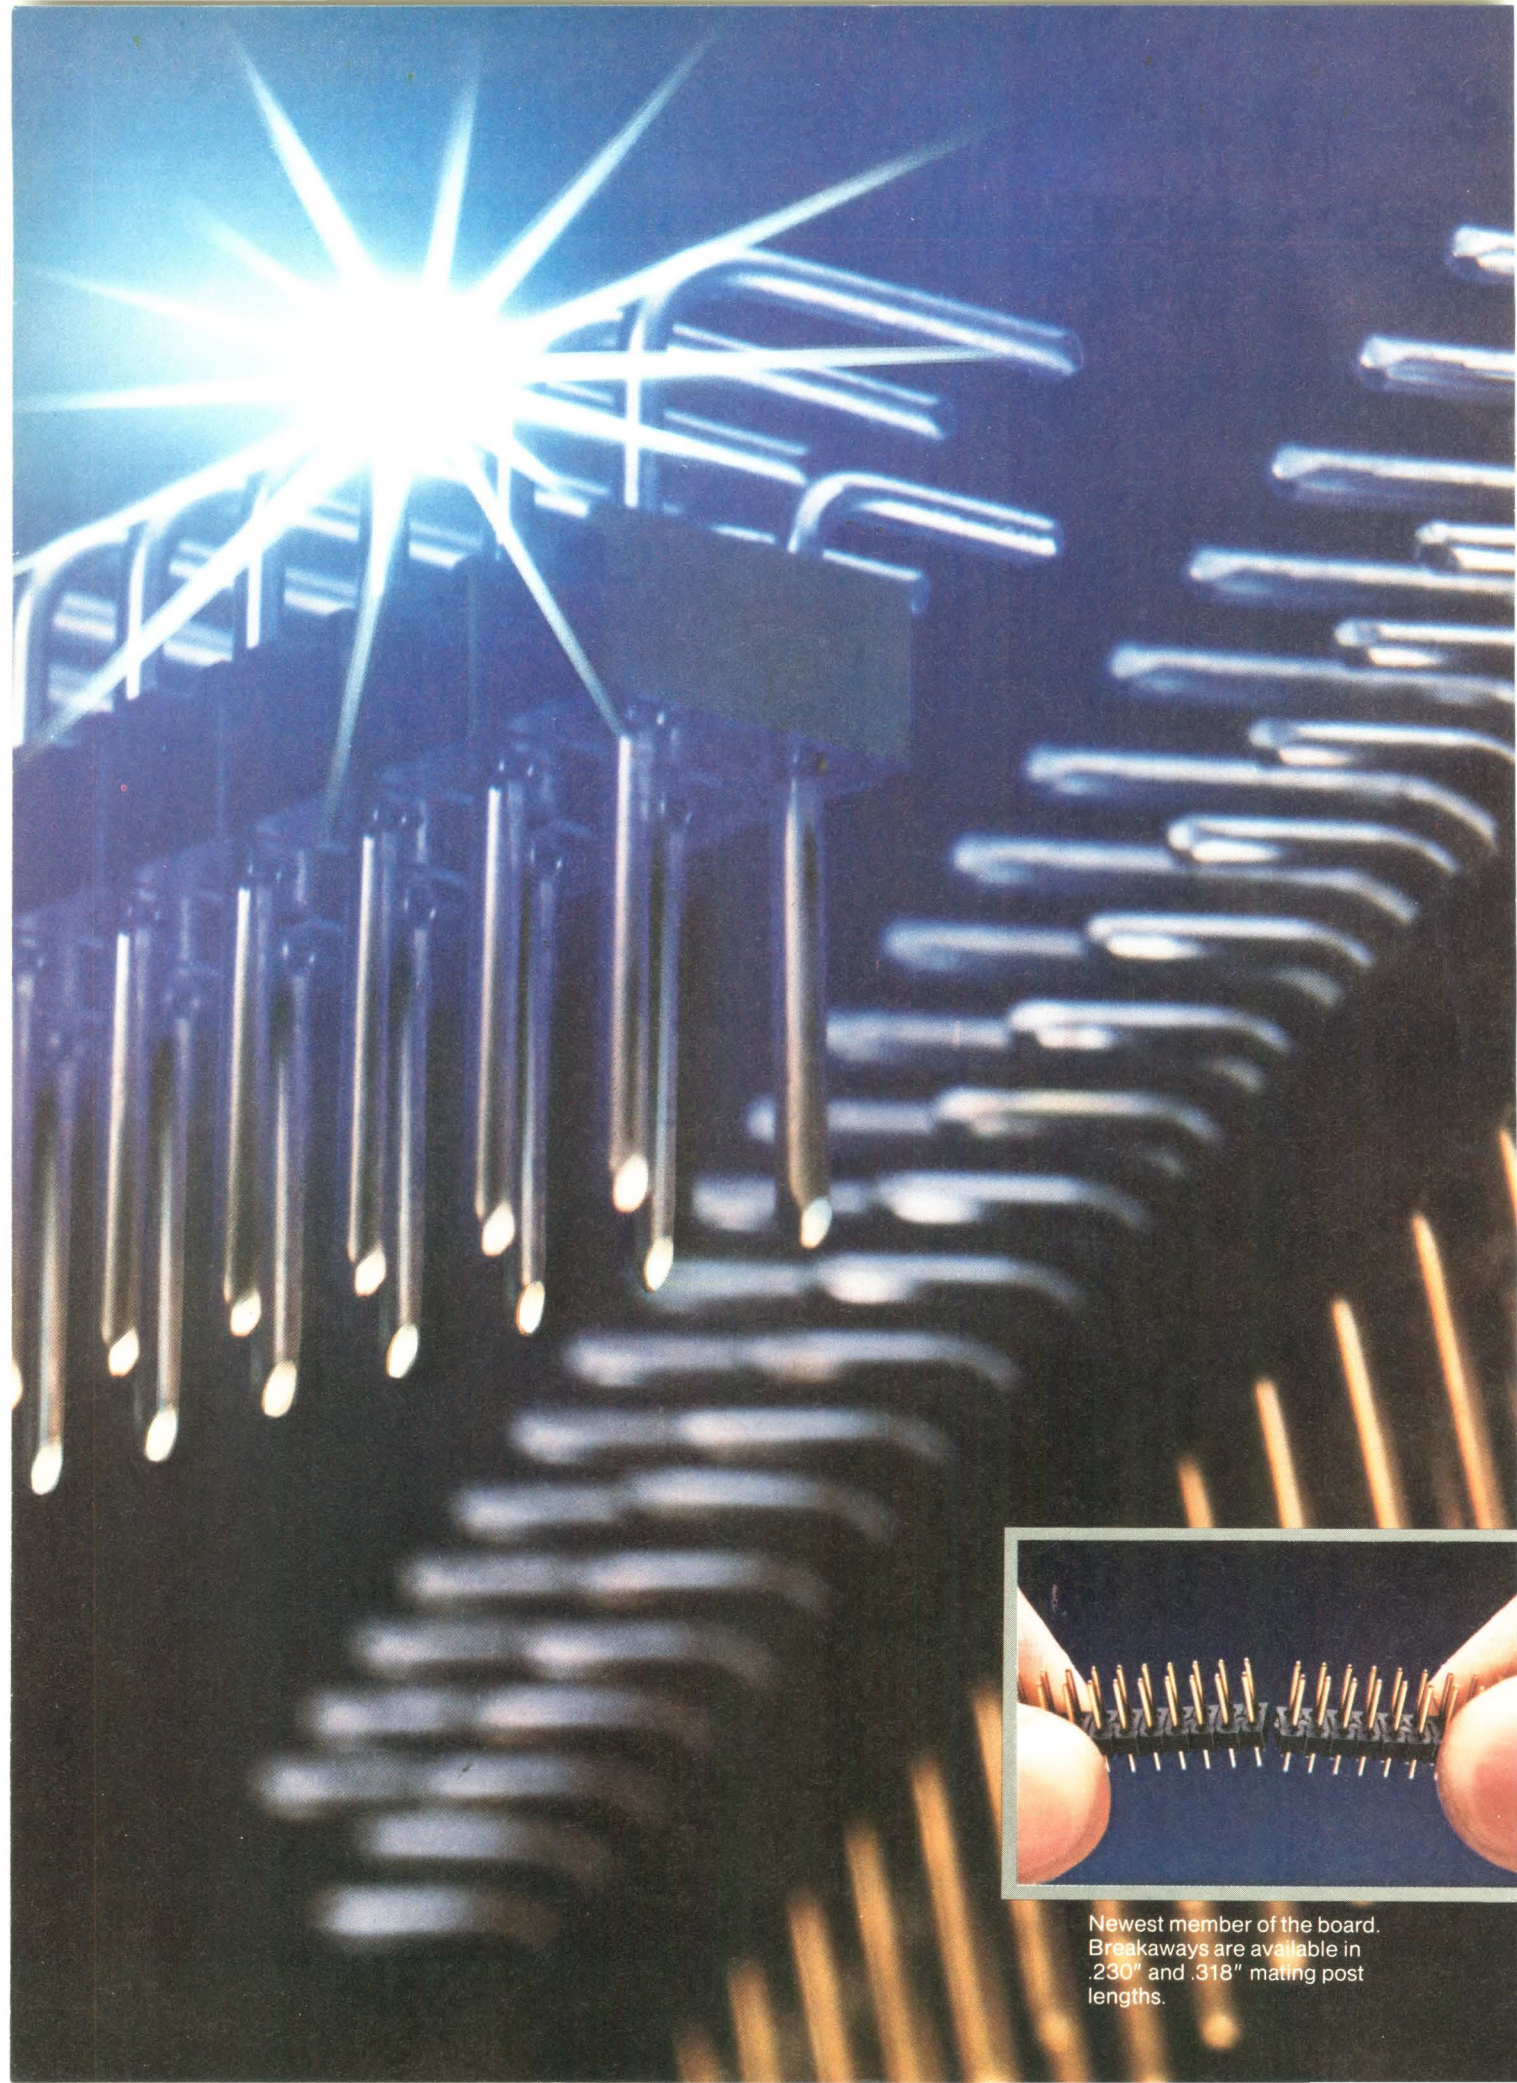

Newest member of the board. Breakaways are available in .230″ and .318″ mating post lengths.

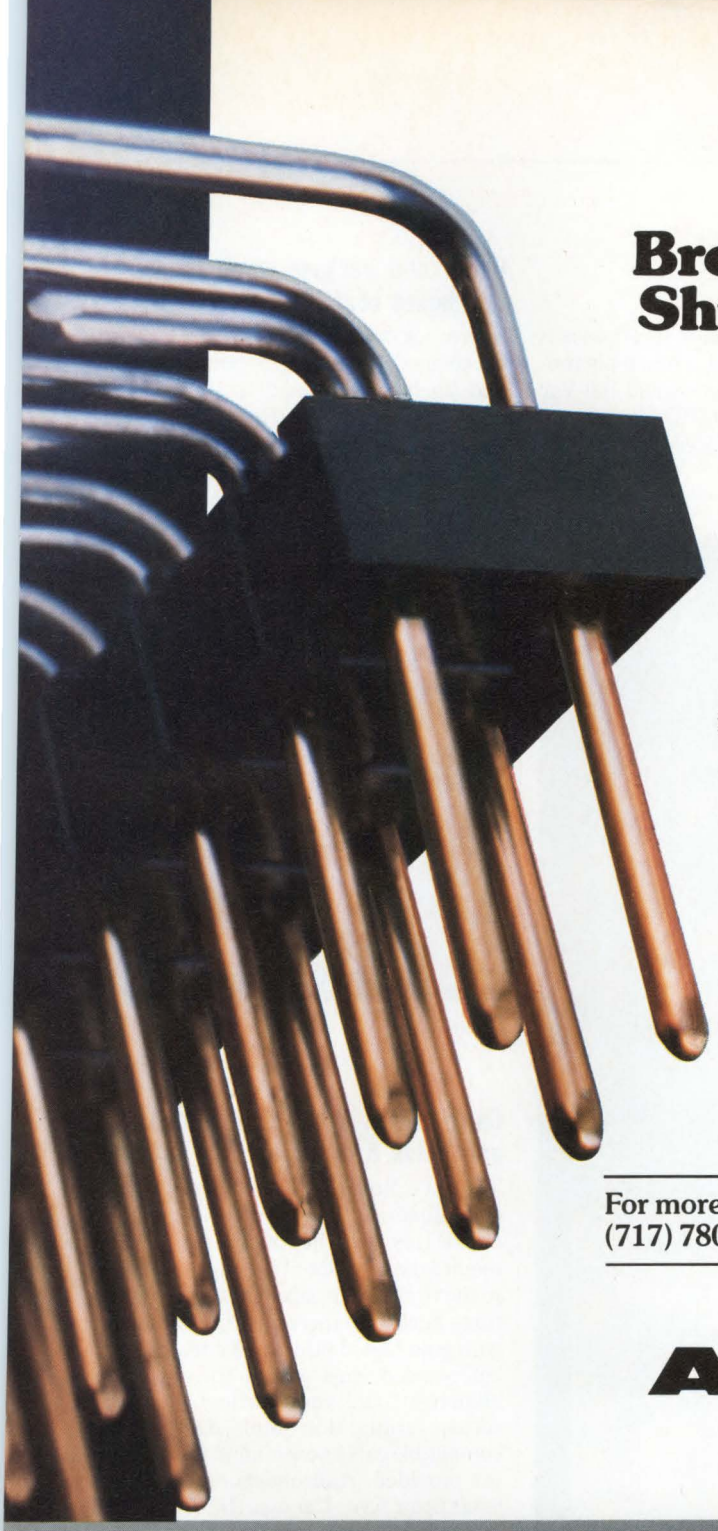

Breakaways. Stacking. Shrouded. Unshrouded. Let our headers cut the cost of your inventory.

Take our breakaways.

Just stocking these versatile snap-apart headers can save you the hassle and cost of inventorying individual sizes. They come in sticks up to 40 position, single, 80 position dual, right-angle or in-line. You can simply count off the positions you need—and snap— a smooth clean break, which even allows you to mount them end to end.

True, breakaways aren't news. But AMPMODU breakaways are. They offer you the same precision, reliability and board-compatibility you've come to count on throughout this modular interconnection line.

Whatever your application, look to AMP for the headers you need—and for the mature engineering that can save you time and money.

For more information, call the AMPMODU Desk at (717) 780-4400. AMP Incorporated, Harrisburg, PA 17105

CIRCLE NUMBER 10

AMP means productivity.

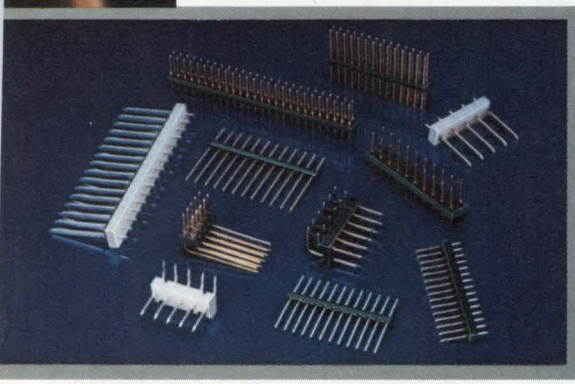

Variety of headers in .025² and .031" x .062" post sizes, in both single and double row configurations.

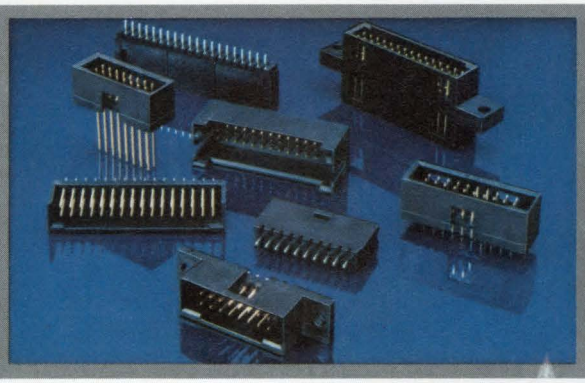

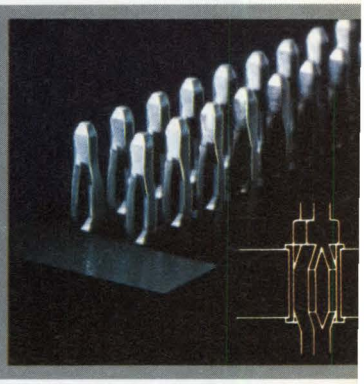

Shrouded aplenty...including compliant pin versions.

Switching power supplies feature small format

The XL series power supplies measure 3.25 x 7.5 x 1.75 in. and require little or no moving air. Other specs include dual adjustable outputs, floating outputs, and over voltage protection on primary output. Surge ratings are above 60 W for disk drive applications. Transient response is 50 to 100 percent of load. The supply provides short circuit protection and 100-Vac holdup time. **CEI Corp,** PO Box 501, 30 Industrial Dr, Londonderry, NH 03053. **Circle 318**

Power and package combo eases VME board configuration

Microrack consists of a KM6-card frame, backplanes, power supply, cooling system, and enclosure. Aimed at easing VME board configuration, the Microrack features a readily accessible power supply. This is accomplished by mounting the 400-W power supply on the card frame's hinged panel, which also carries cutouts for 25-pin, D-type connectors for I/O. Output voltages are 5 V at 45 A, 12 V at 10 A, −5 V at 4 A, −12 V at 10 A, and 12 V at 4 A. **BICC-Vero Electronics Inc,** 171 Bridge Rd, Hauppauge, NY 11788.

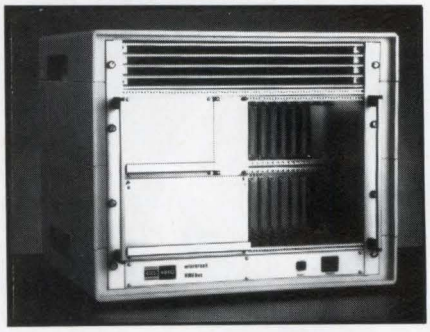

Circle 319

Card cage is Q-bus compatible and holds 8 or 16 cards

The 8LCC holds up to 16 dual-width or eight quad-width LSI-11 cards. Expansion connectors on the backplane do not require use of a slot in the cage to connect to the first system cage. Units include features such as tapered entry connectors with gold plated bifurcated contacts, thin color-coded card guides for easy card alignment, and good airflow. Standard screw termination and Andromeda-compatible cable power input connections are provided. **Andromeda Systems, Inc,** 9000 Eton Ave, Canoga Pk, CA 91304. **Circle 320**

Lithium copper oxide battery offers greater capacities

This battery gives three to four times the capacity of an equivalent zinc-carbon battery and twice that of a comparably sized alkaline. Resin-encapsulation extends the operating range to 150 °C for applications in high ambient temperatures. Low self-discharge characteristics bring an expected shelf life of 10 years. Low density allows for weight savings in comparison to either alkaline or mercury batteries. SAFT **America Inc,** 107 Beaver Ct, Cockeysville, MD 21030. **Circle 321**

Finally, a cure for data acquisition and control problems for VAX™ and PDP-11.™

Fred Molinari, President

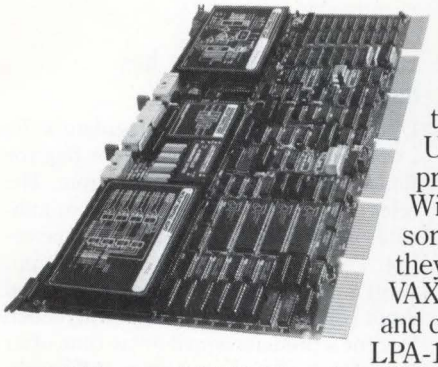

We're pleased to offer a real breakthrough.

Data Translation's single board UNIBUS™ analog I/O products have arrived. With a bit-slice processor and FIFO memory they operate on DEC's VAX and PDP-11 series and can replace the DEC™ LPA-11K.™

This new UNIBUS series offers the software tools for data acquisition that you've been waiting for. UNIX, VMS and RSX libraries support all the analog and digital functions on-board. Like 250 kHz A/D, D/A, digital I/O and a real time clock.

The DT1771 Series also includes microcoded features like histogramming, signal averaging, and pre- and post-gate sampling. All of which have never been available in board level products before.

So whether you're working in laboratory research or industrial control, call us.

We'll cure your problems and have you up and running in no time.

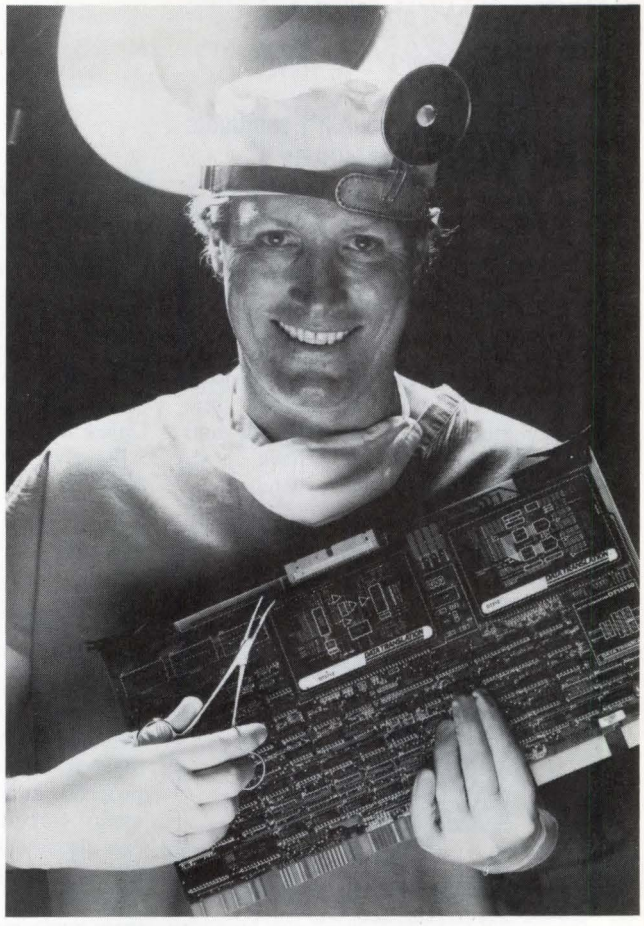

FIVE DAY 5 DELIVERY

Call (617) 481-3700

See our new 576 pg. catalog/handbook in Gold Book 1985. Or call for your personal copy today.

UNIBUS DATA ACQUISITION BOARDS FOR VAX AND PDP-11

Model	Analog In	Analog Out	Features	A/D Resolution	Thruput (kHz)	Software UNIX, RSX, VMS	Digital I/O Lines	Prog. Timer
DT1771	8DI	4	Programmable Gain, Signal Averaging, Histogramming, pre/post gate sampling	12 bit	250	Yes	16	Yes
DT1777	4DI	4	Signal Averaging, Histogramming, pre/post gate sampling	16 bit	100	Yes	16	Yes
DT1778	4SE	4	Simultaneous measurements, Signal Averaging, Histogramming	12 bit	100	Yes	16	Yes

DATA TRANSLATION

World Headquarters: Data Translation, Inc., 100 Locke Dr., Marlboro, MA 01752 (617) 481-3700 Tlx 951 646.
European Headquarters: Data Translation, Ltd., 430 Bath Rd., Slough, Berkshire SLI 6BB England (06286) 3412 Tlx 849 862.
In Canada: (416) 625-1907.

Confidence check capability spotlights programmable power system

A programmable power system features confidence check, which allows a computer to verify correct voltage and current. The system can control the output voltage and current of up to six power supplies. Communication proceeds according to the IEEE 488 standard. The system includes power supplies, rack adapter, IEEE 488 interface card, one programming card per power supply, and an optional confidence check card. Manual/auto-mode select switch eases system maintenance and set-up. **Lambda Electronics**, 515 Broad Hollow Rd, Melville, NY 11747.

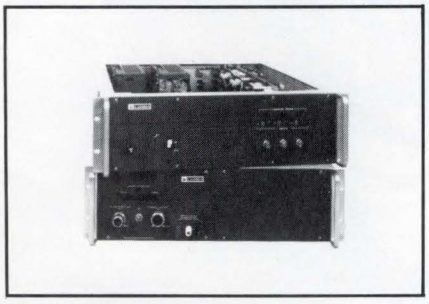

Circle 322

Multibus power supply requires no output power connectors

In the 280-W version (ES281CMB) the switcher slides directly into the standard Multibus card cage and mates with the P1 connector for five different outputs. It uses a high temperature design with a single transistor forward converter operating at a switching frequency of 60 kHz. This design, together with heat sinking and layout, provide full load operation at 60 °C without fans. Input specs include 100 to 132 Vac at 120 Vac or 200 to 264 Vac at 240 Vac, 47 to 63 Hz, user-selectable power range. **Lorain Products,** 1122 F St, Lorain, OH 44052.

Circle 323

Sealed lead acid battery features internal connections

The Cyclon monobloc contains fewer pieces for reduced battery cost. The sealed design provides no acid, acid vapor, or water loss and incorporates recombination of gases within a starved electrolyte system. A self-resealing safety valve vents during overcharged periods. The maintenance-free rechargeable battery is made with thin, spirally wound, pure lead plates for low impedance, low corrosion, and long life. It can be charged or discharged in any position and interconnected in any configuration for additional capacity. **Gates Energy Products, Inc,** 1050 S Broadway, PO Box 5887, Denver, CO 80217.

Circle 324

High voltage IC combines logic with output power

Using a monolithic 500-V junction isolated biMOS high voltage IC process, designers can fabricate unlatchable CMOS logic circuits, rated to 20 V, that operate properly and safely in noisy environments. Circuits using other logic families can also be designed. The HVIC drivers deliver up to 2-A peak, and drive bipolars, MOSFETs, and SCR triggers. Applications include plasma display driver, switching power supplies, single bus wiring, and robotic vision. Possible analog functions include timers, regulators, op amps, oscillators, and comparators. **General Electric Co,** 80 Wolf Rd, Albany, NY 12205.

Circle 325

Versatile power buffer increases op amp output

The fast unity gain LT1010 can boost op amp output from ±10 to ±150 mA. It is also effective in isolating capacitive loads to avoid stability problems with fast amps. The buffer features a 20-MHz bandwidth and a 100-V/μs slew rate. Quiescent current is 5 mA, so there is no penalty for using the buffer in slow applications or where full output capability is not needed. It is stable with capacitive loads of more than 1 μF. Current and thermal limiting allow delivery of full output with normal loads, yet remain protected if a fault occurs. The buffer sells for less than $4 in quantities of 100. **Linear Technology Corp,** 1630 McCarthy Blvd, Milpitas, CA 95035.

Circle 326

Low dropout regulator has delayed reset function

Unlike most fixed voltage regulators the LM2925 is equipped with an error flag for fault detection in the 750-mA output. The reset is an active low output pin that indicates when the 5-V output is not operating. This condition occurs at power up, short circuit, thermal shutdown, or when input voltage is too low. Reset will remain low for a predetermined delay time after the output attains or regains fully regulated output voltage. The delay time ranges from μs to 1 min and is user set with an external capacitor. It operates with low I/O differentials (less than 0.6 V at 0.5 A). In 100s, price is $1.70. **National Semiconductor Corp,** 2900 Semiconductor Dr, Santa Clara, CA 95051.

Circle 327

Power supply provides emergency 15 min, 200-W output

Designed to give 200 W of continuous output power for a minimum of 15 min, with 5.9-ms top transfer time, the UPS-200 allows data storage and orderly system shutdown. This UPS contains an internal 12-V sealed lead-acid battery, plus a MOSFET power inverter and power transformer, as well as circuits for transient protection, solid state switching, battery charging, line failure detection, and an alarm. The UPS-200 also features user-selectable input voltages of 115 or 230 Vac at 50 or 60 Hz. It is priced at $395 in quantities of one to nine. **Power General,** 152 Will Dr, PO Box 189, Canton, MA 02021.

Circle 328

Q-bus system allows variety of peripherals and CPUs

A standard VQ-11 modular package has a 250-W switching power supply, an eight-slot, quad-width card cage, dual fans, and front panel controls. The card cage accommodates CPUs like 11/21, 11/23, or 11/73s, leaving room for system memory and controllers. The price for a VQ-11/ model 302 system including two RX02 emulating floppies, Winchester emulating two RL02 disks for 20 Mbytes of storage, 11/23 CPU, 0.25 Mbyte of RAM, and four serial I/O ports is $7620. **Zoltech Corp,** 7023 Valjean Ave, Van Nuys, CA 91406.

Circle 329

TELEX®

Telex Shamrock 9250 has everything you want in a GCR tape subsystem. Right down to your favorite color.

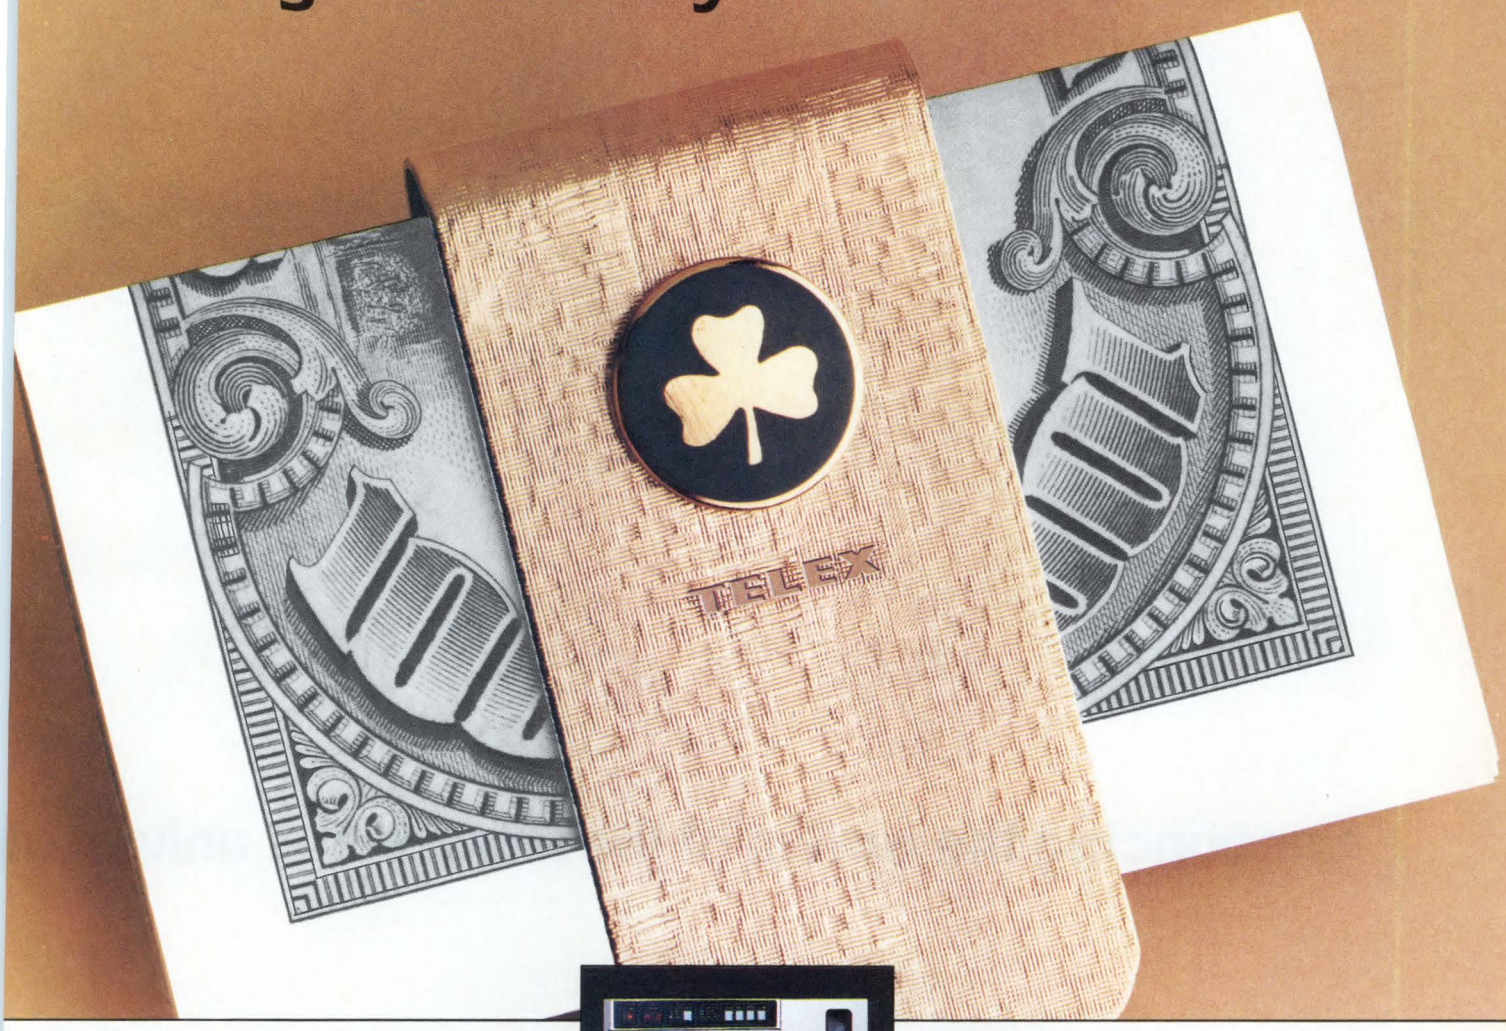

It's the color of the money you stand to save with the new Telex Shamrock 9250. The most affordable GCR subsystem in its class.

Telex engineers have been making advancements in GCR technology for over a decade. The Shamrock 9250 is the biggest breakthrough yet. It gives you the faster access and throughput, greater storage efficiency and higher data reliability you look for in a full performance GCR subsystem. And it does so with the kind of cost efficiency other tape drives just cannot match.

The savings start right up front. The unit price of the Telex Shamrock 9250 is half that of previous GCR subsystems. Its compact design frees up valuable space. And with the assistance of Telex Engineering experts, your interface development will be fast and efficient.

Down the road, the savings continue to add up. Cost of ownership of the Telex Shamrock is truly attractive. There are no planned service calls because the 9250 requires no preventive maintenance. Self-calibrating capability eliminates the need for scheduled adjustments. When service is required, the resident diagnostics hold costs to a minimum. Lower power consumption (typically 100 milliwatt) of the 9250 gate arrays increases reliability and keeps energy costs in line.

There are more reasons why the Shamrock 9250 is today's best value in full GCR performance. For the rest of the story, call your nearest Telex OEM Sales Office or our OEM Marketing Department at 918-627-1111. And let us show you how good you'll look in Telex Shamrock green.

TELEX® SHAMROCK The innovation continues . . .

Telex Computer Products, Inc., Terminals/Peripherals/OEM Products
6422 East 41st Street, Tulsa, OK 74135 (918) 627-1111 **Regional Offices:** Amherst, NH (603) 673-9272/Garden Grove, CA (714) 898-9833/Houston, TX (713) 497-6770 **International:** (214) 931-8511/telex: 730612 TLXINTL ADDI

CIRCLE 138

Announcing the arrival of the world's first universal

Standard Microsystems' HDC 9224 is so advanced it controls hard disks, floppy disks and even tape back-up.

The new addition to Standard Microsystems' magnetic controller family is a single chip that controls Seagate-compatible hard disks, as well as IBM®-compatible single- and double-density floppy disks.

A single chip that allows you—for the first time—to include tape back-up in a small computer without adding a lot of costly hardware and software.

A single chip so advanced its built-in DMA controller and ECC detect and correct errors with absolutely no processor intervention.

We've named this new addition the HDC 9224 Universal Disk Controller. But, you may call it one of the biggest breakthroughs in IC performance in a long, long time.

The cornerstone of an intelligent subsystem.

The HDC 9224's ultraflexible architecture makes it the true "next generation Winchester disk controller." With the addition of a few simple parts, you can use the HDC 9224 to build a high-performance, low-cost subsystem that not only controls hard and floppy disks, but cartridge or cassette back-up, too.

The HDC 9224 easily interfaces to all popular ST506-

disk controller.

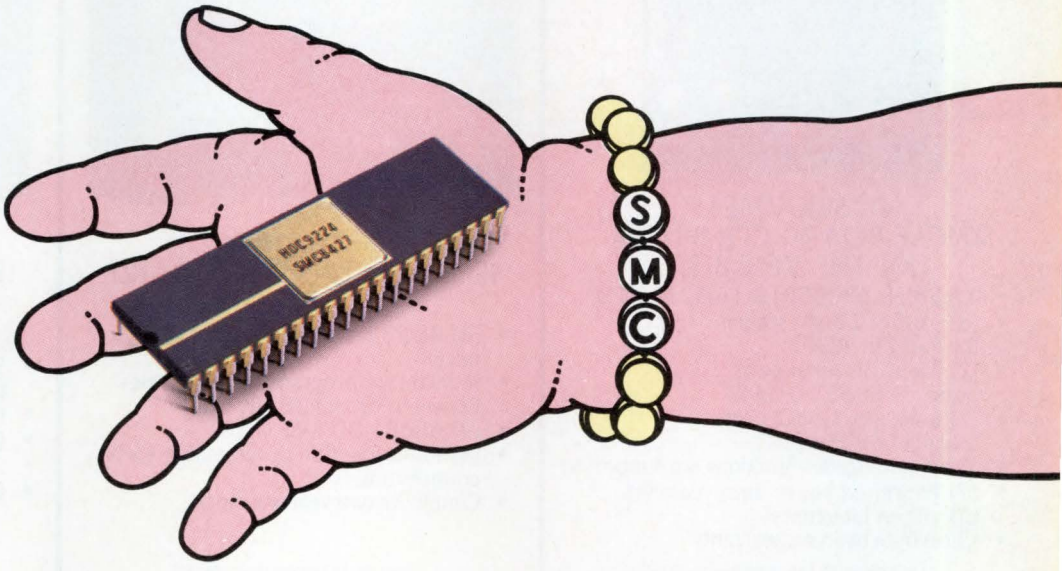

compatible hard disk drives with up to 16 heads. Its flexible format command lets you select hard disk sector lengths up to 16 kilobytes and can be used with both interleaved and non-interleaved sector formats.

What's more, the HDC 9224 controls up to four disk drives of all sizes—8-inch, 5.25-inch and the new 3.5-inch—and it can read and write a wide variety of single- and double-density floppy formats, including such standards as IBM's 3740 and System 34.

Available for immediate delivery.

Best of all, you can get the HDC 9224 when you need it the most. Today. We also plan to announce a well-respected second-source in the near-term future.

The HDC 9224 Universal Disk Controller is an n-channel, silicon-gate MOS/VLSI circuit fabricated with Standard Microsystems' COPLAMOS® technology. It's available for immediate delivery in a 40-pin ceramic or cerdip dual in-line package.

For more information, contact Standard Microsystems Corporation, 35 Marcus Blvd., Hauppauge, New York 11788. (516) 273-3100.

IBM® is a registered trademark of the International Business Machine Corporation.

STANDARD MICROSYSTEMS CORPORATION

CIRCLE 139

OMNIBYTE™

Board Level Products and Systems

All Omnibyte board level products are extensively tested on specialized functional test stations. The more complex boards are also tested on sophisticated automatic test equipment.
All the boards receive a 48-hour burn in, and are retested before shipping.

For more information, please contact:
Peter Czuchra, Marketing Manager

OMNIBYTE CORPORATION
245 W. Roosevelt Rd.
West Chicago, IL 60185
(312) 231-6880
Intl. Telex: 210070 MAGEX UR

A Look at Today . . . A Vision of Tomorrow.

Circle no. 150 on reader service card.

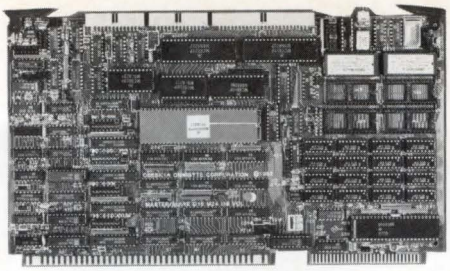

OB68K1A™ SINGLE BOARD COMPUTER ON THE IEEE 796 BUS

- 10MHz MC68000 16/32 bit CPU
- 32K/128K/512K of zero wait state dual ported RAM.
- Up to 192K Bytes of EPROM
- (2) RS-232C serial ports
- (2) 16-bit parallel ports
- A triple 16-bit timer/counter
- (7) prioritized-vectored interrupts
- Omnibyte two year warranty

Circle no. 151 on reader service card.

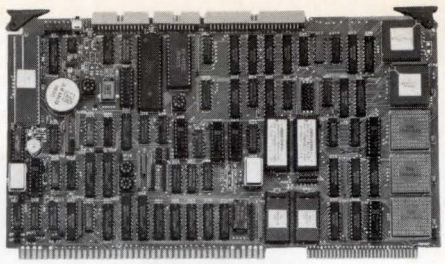

OB68K/MMU™ CPU BOARD WITH OPTIONAL MEMORY MANAGEMENT ON THE IEEE 796 BUS

- 10MHz 68010 Virtual Memory Processor
- Up to (4) 68451 Memory Management Units (optional)
- High speed iLBX* memory port
- 8 channel DMA port
- (2) RS-232C serial ports
- 4/16K RAM
- (2) 28-pin ROM sockets
- Omnibyte two year warranty

*iLBX is a trademark of Intel Corporation

Circle no. 152 on reader service card.

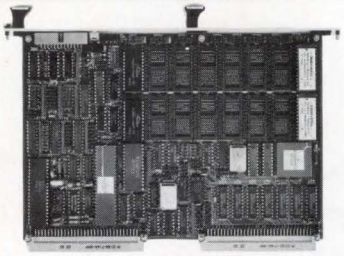

OB68K/VME1™ SINGLE BOARD COMPUTER ON THE VME BUS

- 12.5MHz 68000R12 16/32 bit CPU
- (8) pairs of 28-pin sockets for RAM or ROM
- (2) RS-232C serial ports using (1) 68681 DUART
- (2) 8-bit parallel I/O ports using (1) 68230 PI/T.
- System controller functions are supported
- (7) Prioritized bus or auto vectored prioritized interrupts
- Omnibyte two year warranty

Circle no. 153 on reader service card.

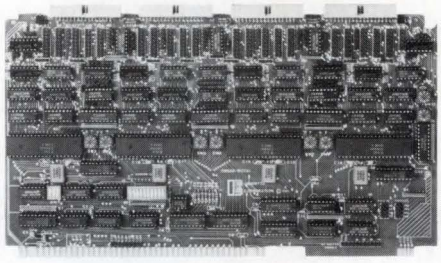

OB68K/OCTAL™ IEEE 796 SERIAL I/O BOARD

- (8) RS-232 or RS-422 serial I/O ports
- Individually programable baud rates between 50 and 38.4K baud
- (4) 68681 DUART chips
- (4) Multi-function programable 16-bit counter/timers
- Omnibyte two year warranty

Circle no. 154 on reader service card.

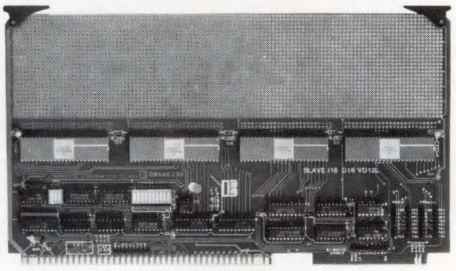

OB68K230™ IEEE 796 96-BIT PARALLEL I/O TIMER BOARD

- 96 bits of software definable parallel I/O
- (4) 68230 PI/T chips
- (4) 24 bit timers
- 35 sq. in. of prototyping area
- Omnibyte two year warranty

Circle no. 155 on reader service card.

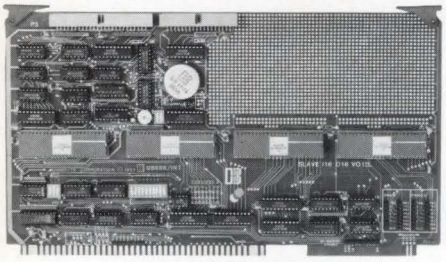

OB68K/INT(P)™ IEEE 796 HOST ADAPTER/ PARALLEL I/O BOARD

- 48 bits of software definable parallel, I/O
- Real time calendar clock w/battery back up
- (4) 68230 PI/T chips
- 15 sq. in. of prototyping area
- Parallel printer port
- SASI* interface to disk controller
- Omnibyte two year warranty

*SASI is a trademark of Shugart Associates.

Circle no. 156 on reader service card.

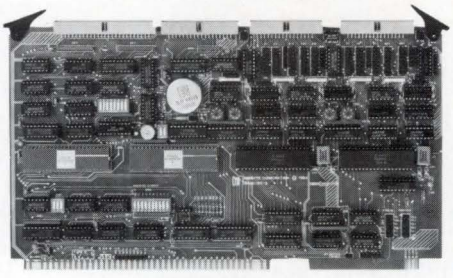

OB68K/INT(S)™ IEEE 796 HOST ADAPTER SERIAL I/O BOARD

- (2) 68230 PI/T chips
- (2) 68681 DUART chips
- (4) RS-232C or RS-422 serial I/O ports
- SASI* interface to disk controller
- Parallel printer port
- Real time calendar clock with battery back-up
- Omnibyte two year warranty

*SASI is a trademark of Shugart Associates

Circle no. 157 on reader service card.

OB68K/SYS+™ IEEE 796 SOFTWARE DEVELOPMENT/ TARGET SYSTEM

- OB68K1A, 128K RAM
- OB68K/INT (S)
- DMA disk controller
- 8 Slot card cage
- 20/40/80 MB Winchester hard disk drive
- 1.2Mb floppy disk drive
- (2) 200 watt switching power supply
- Software available includes Idris*, polyFORTH/32** and more

*Idris is a trademark of Whitesmiths, Ltd.
**polyFORTH/32 is a trademark of Forth, Inc.

Circle no. 158 on reader service card.

Frontend modules convert PC into acquisition controller

Netpac modules interface with the IBM PC via an RS-422 to RS-232-C converter, producing a data acquisition controller. This enhances PC capabilities to include data monitoring and output control for laboratory and industrial settings. Up to 16 Netpac modules can connect to each PC RS-232 port for a total of 1600 channels. Module design provides complete onboard signal processing for measuring a wide range of inputs. Built-in A-D conversion increases network throughput. Up to 100 channels/s can be scanned. Transmission rates hit 19.2 kbaud. **Acurex Corp, Autodata Div**, 555 Clyde Ave, PO Box 7555, Mountain View, CA 94039.
Circle 337

Inspection system with bubble memory detects inserted items

Bubble memory allows the Content Verification System (CVS) to withstand tough factory conditions. This standalone inspection system visually ascertains that the correct objects are inserted in open packages. Reject signals can sound error alarms, activate diverters, disable sealing equipment, and report malfunctioning equipment and rate of flow. The CVS is compatible with existing computer systems using RS-232 communications. Standard system includes hardware, programmable inspection software documentation, and operator training. Units cost $42,500. **Octek Inc**, 7 Corporate Pl, Burlington, MA 01803.
Circle 338

Board-level unit offers advanced motion control

Up to 32 devices can be connected, via twisted-pair wire or mass termination cable, to the PMC-901 programmable motion controller. This intelligent single-board module, combined with ac or dc servomotor systems, adds motion control capability to industrial board-level microcomputer systems. By simply sending ASCII characters to this controller, the host can directly execute positioning commands. Host computer requires only high level commands and status requests. A PMC-901 with a 2-Kbyte nonvolatile memory and the MPL software costs $848 in quantities of 100. **Ormec Systems Corp**, 19 Linden Pk, Rochester, NY 14625.
Circle 339

Menu-driven software package supports data acquisition

Monitoring events and alarm conditions, the DAIS package also generates reports, archives operating histories, and provides realtime access to process information. It consists of a collection of menu driven displays that allow the operator to access online data and process information via tabular and graphical displays. Although it does not require user programming, it supports application specific software generated by the user in Macbasic. Cost is under $10,000. **Analog Devices Inc**, Rte 1, Norwood, MA 02062
Circle 340

Panel-mounted rotary digitizers replace analog potentiometers

Dynapot rotary digitizers represent the digital equivalent of analog potentiometers. Units exhibit a resolution of 128 counts per revolution, and have two orthogonal outputs that determine direction and amount of rotation. Power supply of 5 V is required. Applications include input to up/down counters, manual positioning of X/Y tables, high resolution input for speed or position controls, manual control of stepper motor positioning, and numerically controlled machines. **Dyneer Sensor Technology**, 21021 Lassen St, Chatsworth, CA 91311.

Circle 341

Expansion kit gives the PC data acquisition abilities

The PCX combines alphanumeric and bar code data acquisition capabilities. It consists of a low cost terminal communication card, a software driver, and an NCR 2840 data collection terminal. The terminal includes a 3-of-9 bar code wand, an alphanumeric keyboard, a 16-char display, and an optical reader. The communication card features an optically isolated RS-422 port, an RS-232 port, and a timer. Communicaton rate is 9600 bits/s. The entire kit is priced at $2500. **NCR Corp, Engineering and Manufacturing**, 584 S Lake Emma Rd, Lake Mary, FL 32746.

Circle 342

Stepping motor controller cards off-load process control

Single- and dual-channel STD bus cards for stepping motors off-load motion-control routines. The PPCI-X and the PPCI-XY relieve the host CPU of such process control activities as ramp generation and pulse accounting. The card appears to the host as three I/O ports—address, data, and handshaking/status. Two software-selectable modes—step and continuous—are offered with the PPCI-X. The PPCI-XY dual-channel model allows control of two 25,000-pulse/revolution motors in step mode. It is well-suited for X-Y table positioning. Each card is $495. **I/O Controls, Inc**, 826 Newtown-Yardley Rd, Newtown, PA 18940.

Circle 343

INTEGRATED CIRCUITS

Arrays in CMOS marked by 2.5-ns loaded gate delays

The GA series of gate arrays features a CMOS process that allows 1.8-μm effective gate lengths and typical loaded onchip gate delays of 2.5 ns. The arrays come in four varieties, ranging in size from the GA1000D with the equivalent of 1152 two-input NAND gates, to the 4080-gate GA4000D. The GA1000D devices come in a 40-pin plastic DIP and are priced at $9.30 each in 1000-piece quantities. The GA4000D, in an 84-pin plastic pin-grid array, costs $45. **Mostek Corp**, 1215 W Crosby Rd, Carrollton, TX 75006.

Circle 344

Static RAMs organized as 4096 x 4 bits

N-channel Am2168 and Am2169 have access times as fast as 45 and 40 ns, respectively. The devices are pin for pin replacements for the Inmos 1420/1421. The 2168 version features a power-down mode that reduces power dissipation by 75 percent when the device is deselected—resulting in lower power usage. The chips are available in 20-pin packages, in both plastic and ceramic. Prices for 100 quantity in plastic range from $16.70 to $18.35 each. **Advanced Micro Devices Inc**, 901 Thompson Pl, Sunnyvale, CA 94086.

Circle 345

Signal processing chip aims at speech applications

Specifically designed for processing speech, the CORETECHS signal processor aims at advanced speech recognition, compression, and resynthesis applications. This processor will be available as an LSI CMOS 68-pin leadless chip carrier. The device performs pitch-synchronous waveform analysis, is noise resistant, and yields strong acoustic cues for phonetic segmentation. Based on human perception of speech, the technique used by CORETECHS has demonstrated a 26:1 compression of a 64-kbit/s signal. Availability is scheduled for November. **Scotts Instruments Corp**, 1111 Willow Springs Dr, Denton, TX 76205.

Circle 346

Fast FIFO has a 35-MHz shift-in/shift-out rate

The 57/67413A memory acts as a temporary storage register and serves as a data buffer in systems where the source and receiver operate at different rates. Organized as 64 words x 5 bits, it is fully expandable by word width. Status flags indicate when the device is half full, almost full, and almost empty, while eliminating the need for external counter chips. Three-state outputs and a 24-mA output drive easily handle multiple components, making the device suited for bus oriented applications. Fall through time is 510 ns with a maximum supply current rating of 240 mA. **Monolithic Memories, Inc**, 1165 E Arques Ave, Sunnyvale, CA 94086.

Circle 347

Chip with TTL-compatible CMOS supplies 64 Kbits of ROM

The MOS mask programmable MCM68370 is a 64-Kbit ROM, organized 8-Kbit x 8, and fabricated using n-channel silicon gate technology. Active levels of the chip enable and the chip select, along with memory contents, are user-defined. Chip enable input controls the automatic power-down feature that deselects the outputs and reduces the power consumption from 100 mA maximum active to 15 mA maximum at standby. Fully static periphery feature means no clocking is required on chip selects. Samples available now; pricing in quantities of 2000 is $3.35, with a mask charge of $1000. **Motorola Inc, MOS Integrated Circuits Group**, 3501 Ed Bluestein Blvd, Austin, TX 78721.

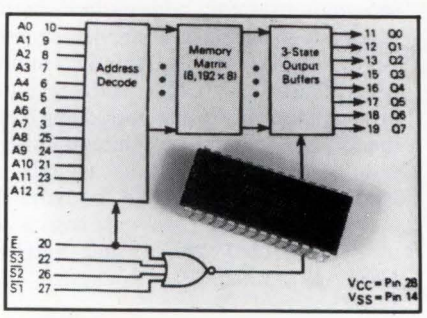

Circle 348

Enclosures for your ideas

Off-the-Shelf...

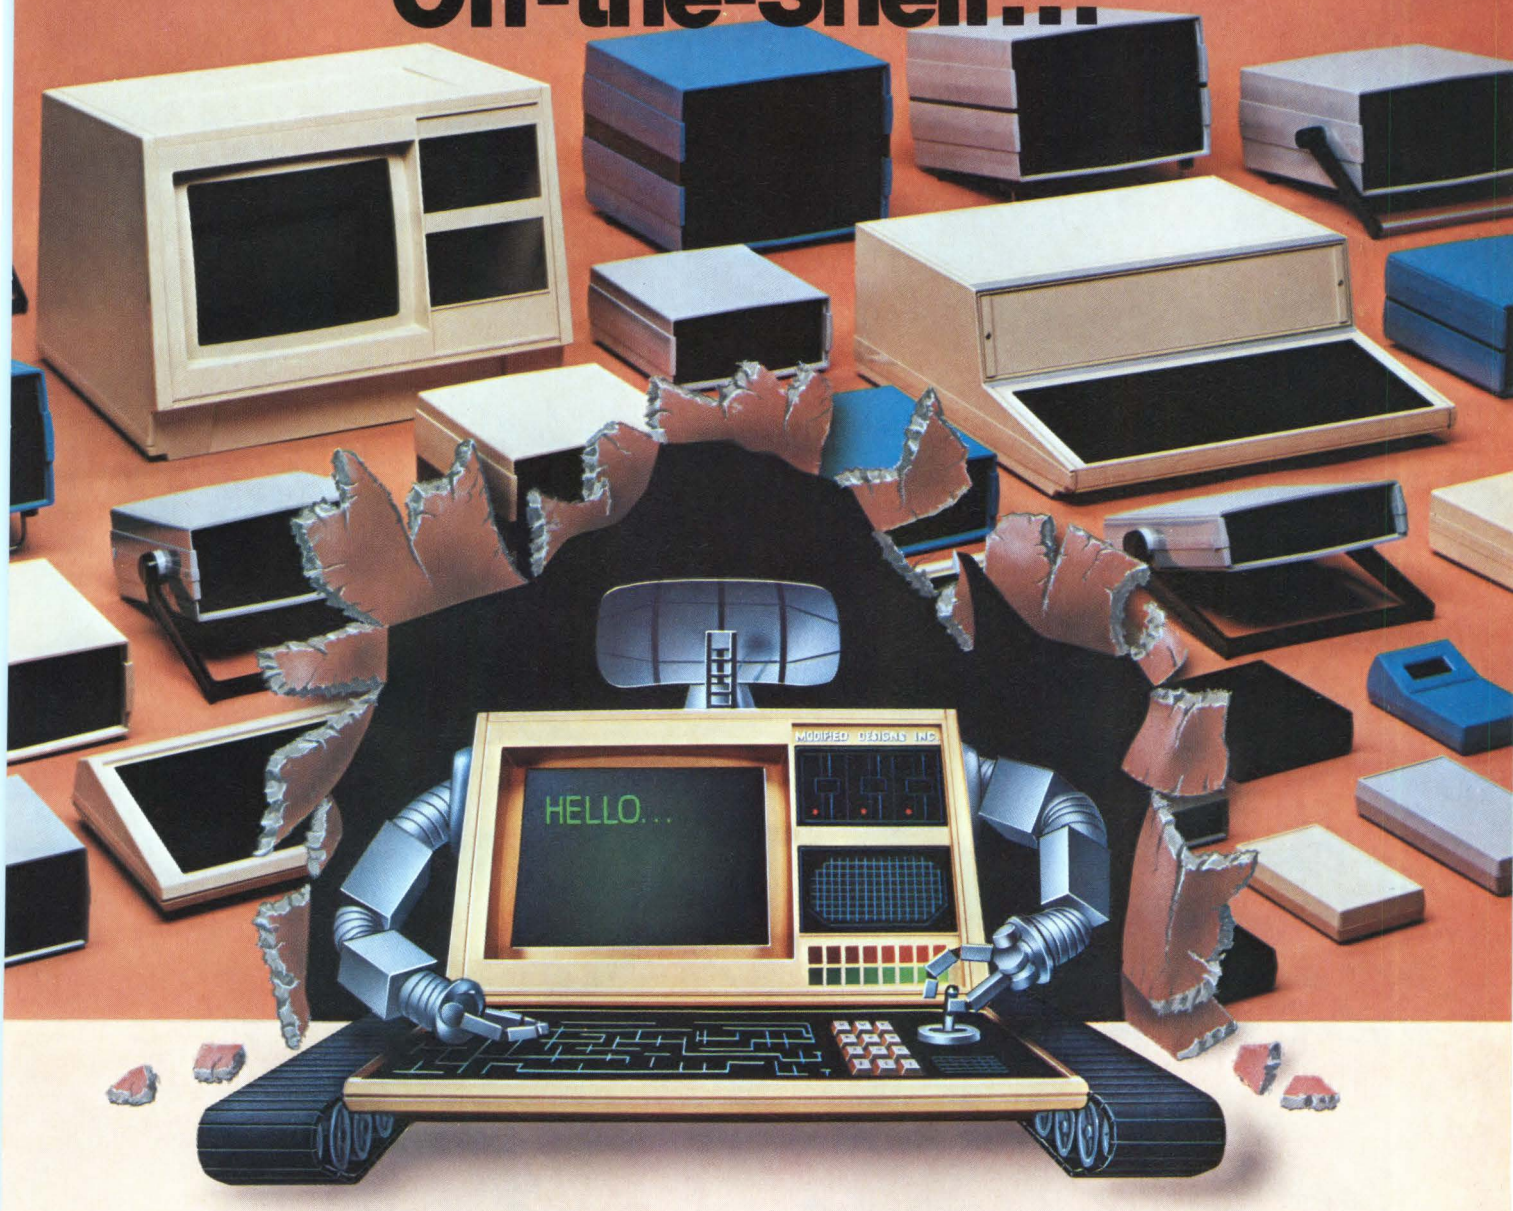

or Off-the-Wall.

Whether you buy standard stock or our "customerized" models, your PacTec enclosure is engineered inside and out to look smart, last long, and cost less.

Because hundreds of PacTec enclosures are available from stock, it's likely we've already designed a package for your idea.

But if you have something in mind that's truly out of the ordinary, PacTec engineers, with years of packaging and design experience, will work with you to make your enclosure attractive, functional, and cost-effective—perhaps with just a simple modification to an existing PacTec package.

All PacTec enclosures offer you these benefits:

- Long-lasting, attractive appearance–molded-in colors in impact- and scratch-resistant ABS–flame-retardant and EMI/RFI shielded where necessary.
- Significantly less expensive than metal or molded enclosures.
- Host of standard accessories–handles, carrying straps, tilt stands, belt clips, and many more.
- Low-cost modifications–vents, louvres, and mounting flanges—plus, custom front and rear panels and special colors.

PacTec's standard enclosures do away with long lead times and secondary finishing, and save money in tooling and assembly.

So if you're looking for something off-the-shelf or off-the-wall, give us a call. We have all the in-house capabilities for design, production, and secondary operations to provide you with a single source for any enclosure you can imagine.

PACTEC® Corp.
Subsidiary of LaFrance Corp.
Enterprise and Executives Aves., Philadelphia, PA 19153
(215) 365-8400

CIRCLE 141

PAC-139

For you anything.

And that's a promise. We are industry's planar cable interconnect problem-solver. Spectra-Strip® bonded ribbon cable was introduced more than 30 years ago—and the pioneering spirit that made us a leader remains strong today. Bring us your specifications and designs. We can develop a positive solution that brings your product design to life.

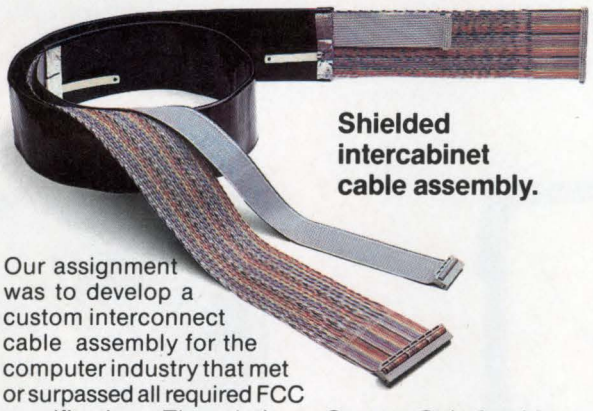

Shielded intercabinet cable assembly.

Our assignment was to develop a custom interconnect cable assembly for the computer industry that met or surpassed all required FCC specifications. The solution: a Spectra-Strip 2-cable assembly consisting of a 60 conductor, 28-gauge special Twist 'N' Flat® cable with special twist length, a 26-conductor, 28-gauge Spectra GP® ground plane, 360° shielding with 2-mil aluminum mylar, wide strap drain and a 35-mil PVC jacket. Each end of both cables terminates in a MIL-C-83503-compatible socket connector.

Internal shielding of interconnect assemblies for mother boards.

A business equipment customer required a flexible cable that passed internal FCC requirements and featured total continuity of drain and shield. We designed a 360°-shielded planar cable with expandable copper mesh shield and 25-mil thick PVC jacket. Drains are ¼" braid terminated with ¼ x .32 spade socket terminals. 360° pull tabs were attached for proper disconnect.

Assembly for computer test and application.

Another computer company presented stiff requirements for test/application cable. We provided a 26-gauge special Spectra-Strip Twist 'N' Flat® cable terminated with solder contact connector, and potted. Two cables terminate at the other end with MIL-C-83503 socket connectors. We attached a single test point grabber. Mission accomplished.

Unique mechanical stress relief.

Tough problem: maintain electrical integrity of terminations despite severe mechanical stresses on the cable. We delighted this computer video display customer with a 75-ohm 2 x 10 coaxial interconnect assembly featuring welded supports on the back end of the connectors. Rigid impedance requirements have been maintained on high speed transmissions.

An energy-saving device.

Early involvement with Amphenol Products, the new connector company, can save you a lot of engineering energy. Your solution may be as close as your phone.

Call the Amphenol Products Design Line

1-800-323-7299

Amphenol Products world headquarters:
Oak Brook, IL 60521

Connect with Amphenol Products at the Midcon Show
Dallas Convention Center, Dallas, Texas
September 11-13, 1984 Booth 730-736

ALLIED Amphenol Products

CIRCLE 142

PICTURE PERFECT PRODUCTIVITY

Administration

Order Entry

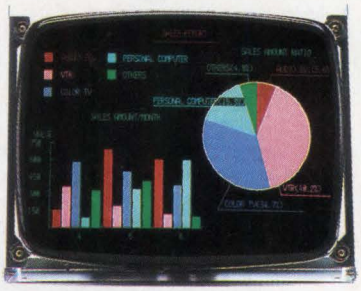

Marketing and Sales

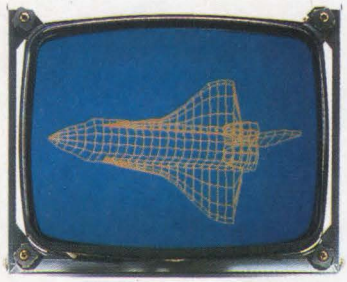

Engineering

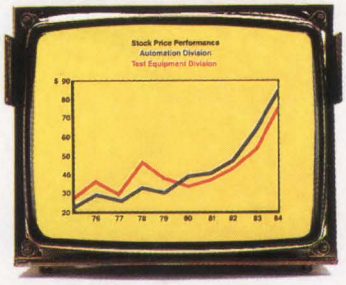

Finance

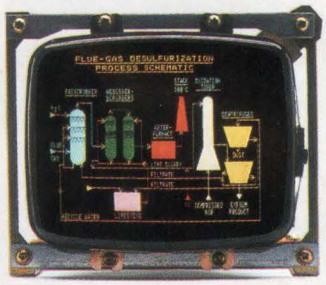

Q.A. and Statistics

MIS Monitors from Mitsubishi

Mitsubishi monitors take excellent care of business at every level from clerical to top management. Our extensive product line includes both color and monochrome models, with high and medium resolution. In any size you need, with exceptional strength in 12", 14" and 16" models. Customs, too.

With all this selection, it's easy for you to performance-match our monitor to your MIS requirements. For superb business graphics, we'll deliver super-high contrast monitors with an extended range of true, vibrant colors and a very dense black. For number-crunching tasks, select a high-efficiency monitor with special anti-reflective coating, reducing operator fatigue even after extended periods of viewing. Choose economical monochrome units to staff up an entire order entry department. And use our precision, high fidelity monitors for demanding CAD/CAM applications.

Just pick the features, then let us show you how Mitsubishi monitors can give your system a better image at the best possible price.

For complete technical specifications or MIS applications information, contact Mitsubishi today.

Write or call: Mitsubishi Electronics America, Inc., Industrial Electronics Division, 991 Knox St., Torrance, CA 90502. Telephone: (213) 515-3993. In the East, call (617) 938-1220; in the Central area, call (312) 298-9223.

MITSUBISHI ELECTRONICS

ADVANCED AND EVER ADVANCING

CIRCLE 143

Integrating A-D converter interfaces microprocessors

The ADC-7109 A-D converter interfaces with microprocessors such as the 6800, 8080, and 8048. With these devices, the ADC-7109's 14 data outputs provide 12 magnitude bits, polarity, and overrange. A UART handshake mode is also provided. This self-contained device includes a buffer amplifier, integrator, comparator, 12-bit binary counter, and 3-state outputs. Housed in a 40-pin plastic DIP, the ADC-7109 costs $11.25. **Datel,** 11 Cabot Blvd, Mansfield, MA 02048.

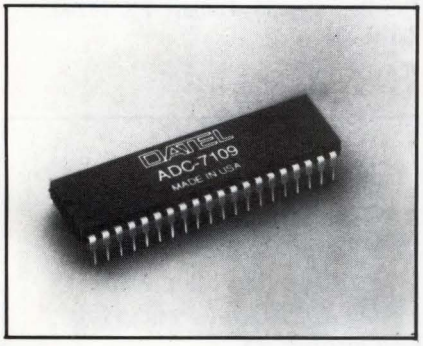

Circle 349

Monolithic CMOS chip synthesizes speech, music, and sound effects

Using the technique of combining phonemes in the appropriate sequence, the SSI 263 speech synthesizer contains 64 phonemes with four different duration settings each. This results in 256 phonemes accessible by an 8-bit code. It contains five registers, each 8-bits wide, selected by a 3-bit address. These registers give access to 256 phonemes, 4096 levels of pitch (32 levels at 8 speeds of inflection) 16 speed or overall rate settings, 16 levels of amplitude, 8 articulation rates, and 255 vocal tract filter response settings. **Silicon Systems Inc,** 14351 Myford Rd, Tustin, CA 92680.

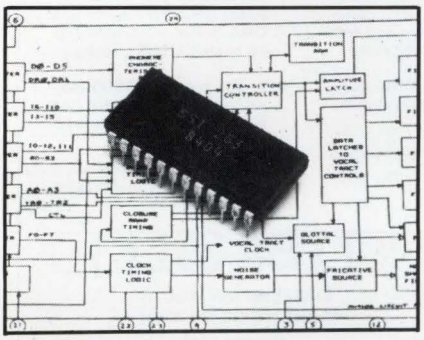

Circle 350

CMOS circuits cut power MOSFET switching losses

Three dual-power MOSFET driver ICs—TSC426, 427, and 428—translate a TTL/CMOS, low level input signal to an output voltage swing equaling the supply. Operating supply voltage is 4.5 to 18 V. The chips require one-fifth the quiescent supply current of the pin-compatible bipolar DS0026 driver, or 8 mA as opposed to 40 mA. Eight-pin plastic and hermitic DIPs are offered in 100-piece prices that range from $1.45 to $9. **Teledyne Semiconductor,** 1300 Terra Bella Ave, Mountain View, CA 94043.

Circle 351

Driver for LCDs clips onto I²C, reduces bus traffic

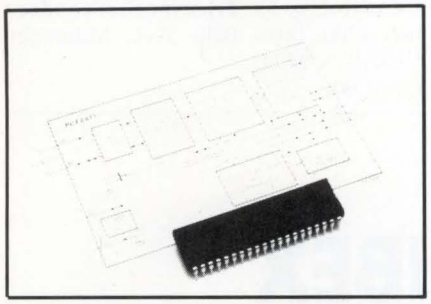

A CMOS silicon gate LCD driver, dubbed the PCF8577, runs up to 32 segments directly, or 64 segments in duplex mode. It has an I²C bus interface that reduces interconnections for remote displays. The PCF8577 cuts bus traffic by using automatic address incrementing across subaddress boundaries. Operating supply ranges from 2.5 to 9 V, depending on LCD voltage level. Samples are available now. **Signetics Corp**, 811 E Arques Ave, Sunnyvale, CA 94086.
Circle 352

Linear arrays support CAD package that includes Spice

Series FB900 linear bipolar arrays perform semicustom linear and linear/digital LSI design tasks. The series consists of six bipolar array chips, customized into application-specific LSI functions by means of metal interconnection patterns. Series FB900 arrays are supported by the Linear CAD I package, which performs schematic capture and Spice circuit simulation on the PC and compatibles. **Micro Linear Corp**, 2092 Concourse Dr, San Jose, CA 95131.
Circle 353

Single-chip multiprocessor interface is host independent

Asynchronous bit-parallel communication is provided by the TMS9650 multiprocessor interface. This single-chip interface is host independent and connects 8-, 16-, or 32-bit-wide microprocessors. The TMS9650 is a standard peripheral interface consisting of eight programmable 8-bit registers at each of two ports. It provides access to 256 bytes of onchip RAM used to buffer data transferred between ports. Fabricated using n-channel silicon gate technology, this device operates from a single 5-V power supply and comes in a 40-pin plastic DIP. Price is $13.20 each in 10,000-piece quantities. **Texas Instruments, Semiconductor Group,** PO Box 809066, Dallas, TX 75240.
Circle 354

October will be a two-issue month—*watch for a Special Report on Mass Storage, October 1; on Digital Signal Processing October 15.*

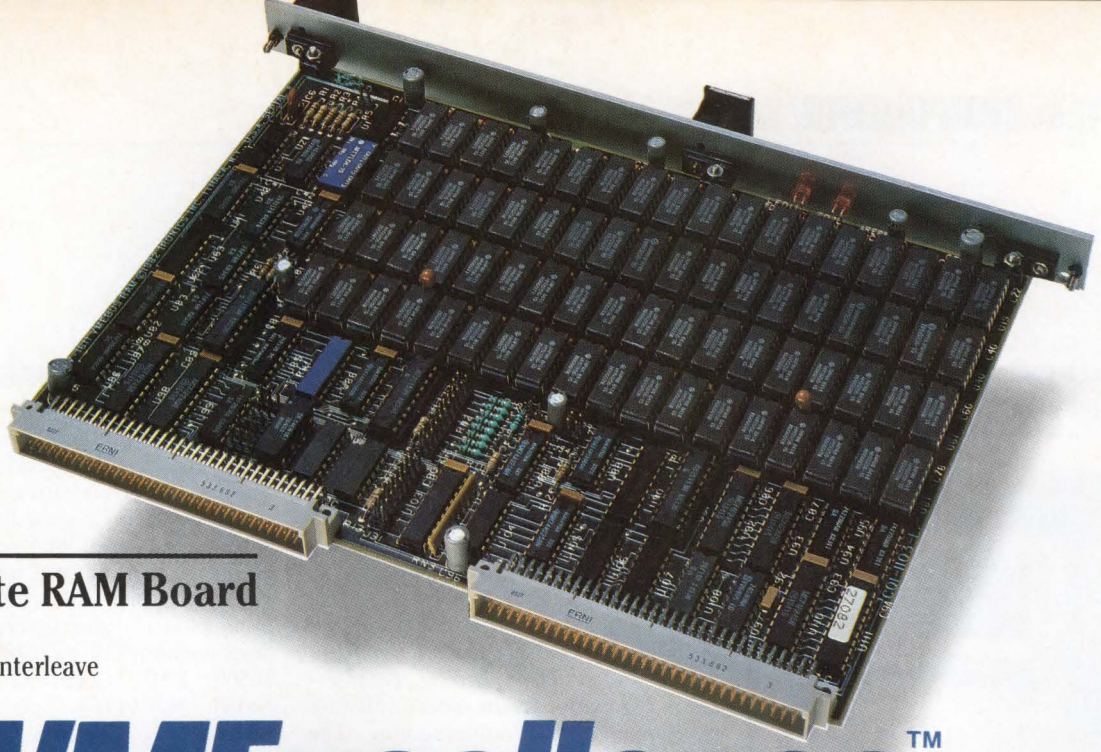

IV-1611

2 Megabyte RAM Board

- 32 bit transfers
- 4-way internal interleave

VMExcellence™

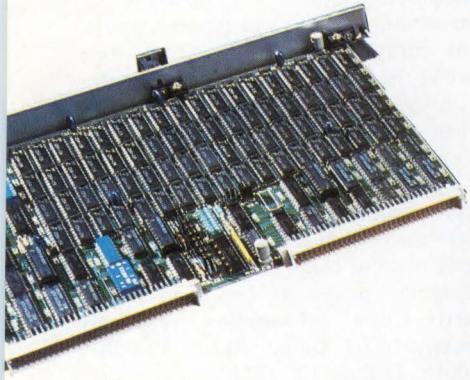

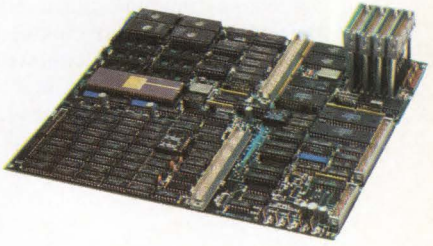

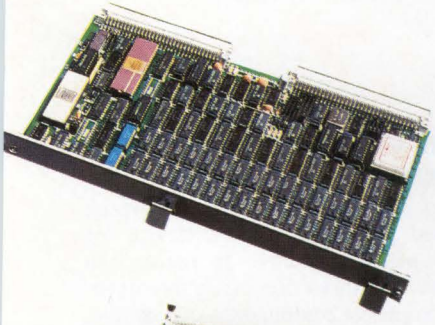

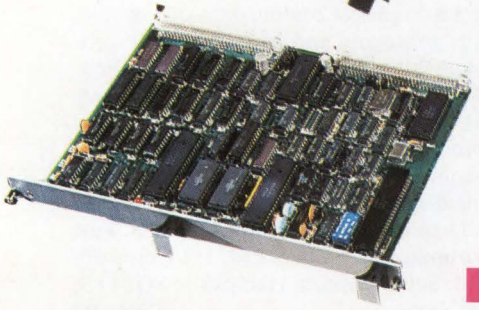

VMEbus

UNIX System V*

High Resolution Color Graphics

Ironics VMEbus boards, systems, and software provide the OEM and Systems Integrator with a powerful and productive development environment, a high performance, cost effective target environment, and the tools and technical support to put it all together... *from IRONICS, the price/performance leader in VMEbus.*

IRONICS Incorporated

Computer Systems Division
742 Cascadilla Street
Ithaca, N.Y. 14850
607-277-4060 Telex: 705-742

*Trademark AT&T

CIRCLE 147

Dual-host controller handles 3270 and asynchronous links

Model 411 dual-host controllers enable terminal users to interconnect concurrently with two remote IBM-type mainframes and up to 16 asynchronous hosts. The 411 links these units simultaneously via remote communication interfaces, and the terminal user can freely switch systems with simple keyboard commands. Also serving as a line, modem, and port sharing device, it is priced at $17,435. **Lee Data Corp**, 7075 Flying Cloud Dr, Minneapolis, MN 55344.
Circle 355

Computer/interface system provides PC-based acquisition and control

The Issac 41-I and the 91-I automate analytical instruments and industrial processes. The 91-I is a general purpose system designed for workstation applications. System hardware includes an IBM interface, 16 channels of 12-bit A-D, 4 channels of 12-bit D-A, 16 binary inputs and outputs, 4 programmable Schmitt triggers, a 16-bit timer, and 8 expansion slots. Both systems use LabSoft software, with complete programming tools, including realtime signal acquisition routines, graphics utilities, and system diagnostics. Programming languages include Basic, Fortran, and C. The low end 41-I is $1900 and the 91-I is $4400. **Cyborg Corp**, 55 Chapel St, Newton, MA 02158.
Circle 356

Realtime microcomputer applications gain IEEE 488 access

Access from Q-bus systems to IEEE 488 instruments is provided by the GPIB11V-2 dual-height card. This card allows 250-kbyte/s data transfers under DMA control. A dense doubled-sided printed wire board, the GPIB11V2 contains 59 logic devices in 40 in.^{2} of board space. Software supports GPIB operations—including talker, listener, active controller, and system controller. An interactive control program speeds software development and troubleshooting. The unit also enables DEC systems to use NET488 LAN products that allow transfer of files between various DEC systems. Pricing is $1495 in unit quantities; the MicroPower/Pascal software extension is priced at $500. **National Instruments**, 12109 Technology Blvd, Austin, TX 78727.
Circle 357

Color CRT controller card connects to the STD bus

Linking to a STD bus, the DSTD-776 can be configured for 8-level color or 16-level monochrome output. It uses an MC6845 to provide horizontal and vertical sync and video refresh data for a standard video monitor. The unit supports Z80, 8088, and 68008 addressing modes, and is designed for use with the DSTD-777 graphics controller for overlaid character and graphics display. The DSTD-776, available for $417, works as the master to provide the needed synchronization signals for the DSTD-777. **DY-4 Systems Inc**, 1475 S Bascom Ave, Campbell, CA 95008.
Circle 358

Data encryption board supports SNA/SDLC protocols

Secure SNA/SDLC communications are derived from an ENC-305 data encryption board designed for IBM PCs. The ENC-305 supports both SNA and SDLC protocols, delivering PC-to-PC file security and allowing PC users to use DES-encrypted mainframe data. The board and accompanying software allow PCs to emulate an IBM 3274-51C with IBM encrypt/decrypt features. The ENC-305 is priced at $1595. **Futurex Security Systems**, 9700 Fair Oaks Blvd, Fair Oaks, CA 95628.
Circle 359

Controller doubles capacity of Prime interfaces

The 5258 communication controller boosts Prime system transfers without requiring additional card cages or power supplies. A single card holds 32 asynchronous channels and operates under both the DMQ and DMT transfer modes. The card comes with a connector panel using RS-232-C connectors for addition of devices without Prime cabling. It lists for $6500. **EMC Corp**, 12 Mercer Rd, Natick, MA 01760.
Circle 360

Interface hooks MIND units to IBM PCs

The Harris Connection allows IBM PCs to operate as multifunctional distributed data processing (DDP) and interactive system terminals. The Connection is a softcard easily installed in a PC, which enables the PC to emulate a Harris terminal attached to a DDP or MIND system. When connected to a MIND system, the PC can also emulate 3270 functions. The PC linked via the Connection can communicate using both BSC and SNA protocols. Users can switch from interactive to PC mode without closing the host sessions. Single-unit prices are less than $1500. **Harris Corp, Information Terminals Group**, 16001 Dallas Pkwy, PO Box 400010, Dallas, TX 75240.
Circle 361

Custom Q-bus controllers fit user needs

The IQ controller is a general-purpose attached processor and I/O device for Q-bus based systems. It features a 10-MIPS bit-slice processor/microprogram sequencer combination that handles bus protocols and simultaneously provides high speed arithmetic processing. The standard IQ performs such functions as fast memory block moves and array processing. Applications include data acquisition and communications. The quad-size board supports full 22-bit addressing and programmable DMA. **Fast Firmware Techniques Ltd**, 151 University Rd, Southampton, Hampshire SO2 ITS, England.
Circle 362

PUT MORE DRIVE INTO YOUR SYSTEM.

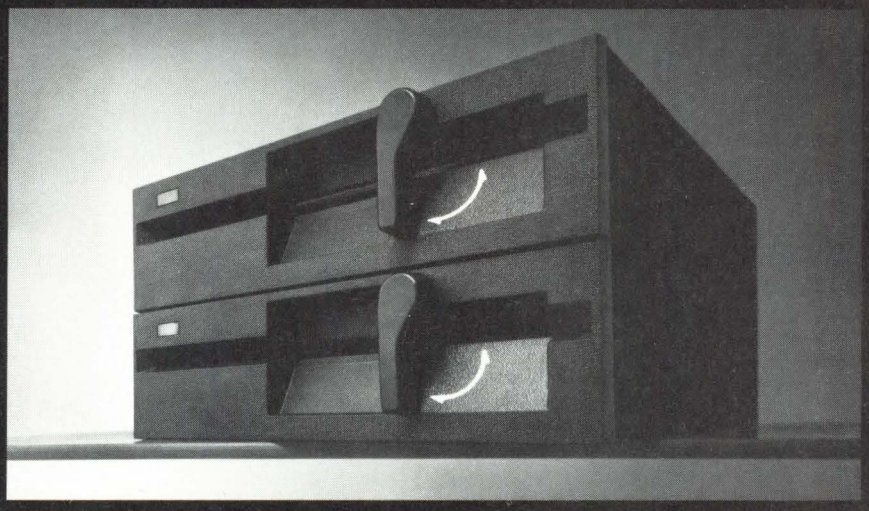

Teac's FD-55 series half-high floppy disk drives run on a mere 5 watts. So now you can add more drives without the need to add more power.

In fact, Teac offers the only full line of 5¼ inch half-high mini disk drives that run on half the power of ordinary dual drive units.

And, with Teac's brushless DC direct drive motors, you get less noise and longer life. Proven up to 10,000 hours.

So, if you're feeling the need for more drive in your system, call us at (213) 726-0303 and we'll get things rolling.

TEAC® ICPD

BUILT TO FANATICAL STANDARDS.

CIRCLE 148

NOTHING ELSE STACKS UP TO IT.

The IOMEGA Alpha 10H Half High: More Performance, Reliability, Versatility– And Virtually Unlimited Capacity.

Take the measure of IOMEGA's new Alpha 10H. Not just the low-profile, half-high physical dimensions, but its true OEM height—a height measured in performance, versatility, convenience, cost-efficiency, dependability and deliverability.

New Heights In The Measure Of Reliability And Performance.

Thanks to IOMEGA's revolutionary cartridge technology, OEMs can integrate a mass storage solution into their systems that easily surpasses traditional Winchester or even Winchester cartridge alternatives. In access time (35 ms average). In transfer rates (1.13 Mb/sec). In day-in/day-out dependability.

Capacity, Convenience, And Low Cost-Per-Megabyte: Taking The Measure Even Higher.

One of the best measures of the Alpha 10H's true OEM height is in the versatility it allows you to design into your system. Our highly reliable, easily transportable and absolutely interchangeable 10-megabyte cartridges mean the sky's the limit on system capacity in applications as diverse as system control and monitoring, office information and word processing, and industrial instrumentation, among others. So is convenience. Be it loading large and complex data bases, or the rapid, reliable backup of system data, the Alpha 10H simply cannot be topped. Lastly, in low cost per megabyte, no storage device, fixed or floppy, stacks up.

Do Your Own Measuring. Call Today.

The IOMEGA Alpha 10H half high is available, right now. For more information on how this superb new IOMEGA offering measures up to the competition, give us a call.

IOMEGA Corporation
1815 West 4000 South
Roy, Utah 84067
(801) 776-7330

CIRCLE 149

IOMEGA™

Microcomputer software package adapts PC to data acquisition jobs

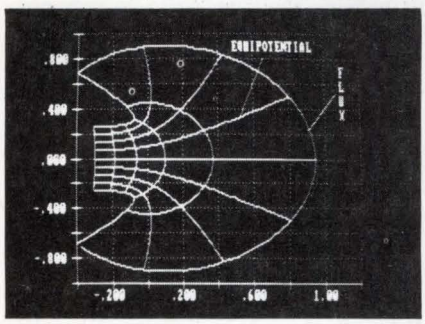

Designed for use with the PC and PC/XT, the ASYST integrated software package manages data acquisition, analysis, and graphics. The package includes three modules. A system/graphics/statistics module establishes the system environment, stores data and provides graphics as well as basic statistics and math functions; a data analysis module reduces, manipulates, and analyzes data; and an aquisition module allows interface to scientific instruments using simple syntax. The package is priced under $1800. **Macmillan Software Co**, 866 Third Ave, New York, NY 10022. **Circle 363**

Graphics builder allows rapid prototyping

Rapid creation or modification of icons, menus, prompt messages, or complex interactive displays is possible with the Blox graphics builder. Aimed at both the VAX/VMS and VAX/Unix systems, Blox is used to prototype and develop graphics applications interactively. The package includes a Core graphics library and a Blox user interface library with support routines. It is device independent, and runs on industry standard terminals with graphics input (mouse, tablet, and keyboard). A single CPU license is $8000, binaries only. **Rubel Software**, 215 First St, Cambridge, MA 02142. **Circle 364**

Firmware ensures ROM BIOS with PC-compatibility

The ROM BIOS software was developed under strict controls to avoid copyright infringement and insure against infringement suits. It is offered as part of a software package to make a micro fully compatible with the PC. The package also includes a PC-DOS compatible version of the MS-DOS operating system, utilities, and a Basica-like implementation of GWbasic. Custom hardware/software engineering services allow hardware-incompatible machines to be compatible even if 80186-based. Price is $290,000. **Phoenix Software Associates Ltd**, 1420 Providence Hwy, Norwood, MA 02062. **Circle 365**

Security tool tracks VAX system use

Audit represents a software security and documentation tool for VAX systems. The Audit security tool produces a complete record of all user input and computer output. It can monitor any terminal or terminals on the system. It uses little system overhead and does not require additional processes for individual session run. **Clyde Digital Systems, Inc**, 3707 N Canyon Rd, Building 3, Provo, UT 84604. **Circle 366**

Software components provide I/O and file management

Tailored for realtime environments, FMX and IOX work with VRTX a realtime executive, to simplify programming of microbased applications. For applications not requiring a file system, IOX can be used alone to program disks, magnetic tapes, and LAN hardware. Both modules can be used on any board and placed anywhere in micro memory. For all current VRTX applications, the cost is $2750 for FMX and $5275 for IOX. **Hunter & Ready, Inc,** 445 Sherman Ave, Palo Alto, CA 94306.

Circle 367

Entry level system for Pascal improves productivity

ES/P helps the user write correct programs the first time. The Pascal program editor performs automatic program structuring, provides abbreviations for key words and structures, and lets the user fill program skeletons. The text editor handles multiple files simultaneously, is crash resistant, and allows easy editing of files larger than memory. Hardware and software requirements include an IBM PC, 80-col display, 128 Kbytes of memory, one drive, and PC-DOS. Cost is $249. **Bellesoft, Inc,** 2127 Bellevue Way SE, Bellevue, WA 98004.

Circle 368

Interfacing support system covers 80186/80188 microprocessors

Binary object-module support for Intel microprocessor development is possible with the BSO/FMTCNV. Applications for the 80186, 80188, 8086, and 8088 families can be developed. Intel system users can convert symbolic object-module files from BSO to Intel format. The BSO/FMTCNV allows development in multi-user, high speed VAX/VMS environments that use BSO compilers, symbolic debuggers, and macro assemblers. Generated object modules can then be transmitted to a dedicated emulator for final hardware/software integration and PROM programming. Subsequent Unix 4.2 and RSX operating system versions are due. **Boston Systems Office,** 469 Moody St, Waltham, MA 02254.

Circle 369

Relational database management system runs on Eclipse

Based on the SQL sub-language, DG/SQL is targeted to applications like CAE where large data blocks need to be processed and analyzed. Because the software is accessed through a nonprocedural language, users gain access to information by simply stipulating what data sets they need. Data is stored in a collection of tables that are comprised of columns and rows with an easy format to manage, retrieve, share, and update data. The software lets two or more users access a data base simultaneously, logs transactions, performs backup and recovery, and allows users to change the database structure through a single command. Initial licenses are $20,000. **Data General Corp,** 4400 Computer Dr, Westboro, MA 01580. 01580.

Circle 370

"With the Interphase Storager,™ I can make a 5¼" hard disk perform like an 8" disk."

Frank Emser
Manager Hardware Development
Paradyne Corporation

The Interphase Storager Multibus™ controller can give a 5¼" Winchester disk capabilities never before possible. Storager not only gets more performance from existing ST506 drives, but also supports the new ESDI and ST412HP interfaces for more power and capacity than ever before. And since Storager can control two Winchester disks, four ¼" tapes (QIC-02), and two 3½", 5¼" or 8" floppies, the same controller can be used for every storage need.

Storager features 1:1 interleave, with concurrent disk and tape transfers and simultaneous disk and bus transfers for speed and high performance. And Storager's unique "virtual buffer" architecture with UNIX™-optimized intelligent caching can reduce or eliminate disk rotational latency and overcome data overrun/underrun problems of FIFO-based controllers. Plus, for the very first time on a controller, Storager has an *on-board* 68000 CPU.

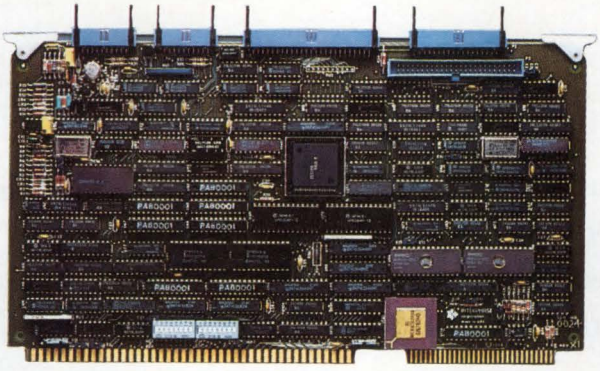

The Storager controller is the latest product in Interphase's line of high-performance Multibus controllers. Interphase also offers Multibus controllers for SMD disks, local area networks and video monitors. Plus powerful disk controllers for the IBM® PC. They're all backed by a great customer support team that works full time with Interphase customers to assure that our products work the way they should — in the system.

Find out how Storager can make a 5¼" disk perform like an 8" disk. Call Interphase today at (214) 350-9000.

Storager is a trademark of Interphase Corporation
Multibus is a trademark of Intel Corporation
UNIX is a trademark of Bell Laboratories
IBM is a registered trademark of International Business Machines

INTERPHASE corporation

2925 Merrell Rd.•Dallas, TX 75229

CIRCLE 161

Unique process takes aim at large volume printing tasks

The S 6000 is an intelligent, nonimpact page printer using ion printing technology. It provides high speed printing of 60 pages/min, and handles large volume print jobs requiring from 100,000 to more than 1.5 million copies per month. With a resolution of 240 x 240 dots/in., the S 6000 reproduces standard line printer formats and offers electronic forms overlay. The S 6000 integrates a 2460 engine with an LSI-11 processor and specially designed image generation electronics. IBM 3211 interfaces are available. The unit is priced under $60,000. **Delphax Systems**, 977 Panera Dr, Mississauga, Ontario L4W 2W6.
Circle 380

Shuttle matrix printer reaches speeds to 600 lines/min

The 4400 series offers selectable print modes for draft, near letter quality, line and block graphics, as well as oversize chars, superscripts and subscripts, and specialized char sets. Six-pin tractors above and below the print line provide positive dot positioning. A resonant shuttle system provides high reliability. Operator replaceable print modules that can be changed in less than 3 min, an internal paper storage system with paper guides and a paperout sensor, plus a 50-million char mobius loop ribbon cartridge with re-inker, simplify operation of the units. Prices range from $5500 to $7200. **Genicom Corp**, One General Electric Dr, Waynesboro, VA 22980.
Circle 381

Color display features high bandwidth and broad line frequency

The CDCT 6000 sports horizontal line frequency ranging from 45 to 75 kHz in four adjustable steps, vertical line frequency of 35 to 120 Hz, minimum bandwidth of 120 MHz, and 1280 x 1024 pixels at 60 Hz noninterlaced. The adjustable horizontal line frequency allows easy upgrading. Additional features include an inline gun-dotted CRT and a separate monochrome input facility. Applications range from solid modeling and image generation, to artwork mapping and CAD/CAM. **Barco Industries, Inc U.S.A.**, 2818-G Interstate 85 South, Charlotte, NC 28208.
Circle 382

Laser printer runs quietly and is PC compatible

Letter-quality output and silent operation mark the LaserJet printer. Designed to operate with personal computers such as the HP 150, the IBM PC, and PC compatibles, the printer registers a noise level of less than 55 dB. Eight pages can be printed per min. The LaserJet prints either horizontally for correspondence or vertically to create spreadsheets using a compressed character font. Resolution for the printer lists at 300 x 300 dots/in. Up to four fonts in plug-in cartridges are offered, in addition to a standard typewriter Courier 10 font. Disposable cartridges replace the common ribbon configuration. The unit is priced at $3495. **Hewlett-Packard Co**, 1820 Embarcadero Rd, Palo Alto, CA 94303. **Circle 383**

Raster terminal graphs at 125-ns/pixel speed

The Vistagraphic 4500 terminal provides a 1280 x 1024 pixel resolution and a 60-Hz refresh rate. Using GCP software, the unit is compatible with the Graphic 7, 8, and other Vistagraphic 4000 series display systems. The GCP coordinates all internal operations, provides handlers for all I/O devices, and handles host communications. In addition, the GCP has TTY emulation, local debug, and auto-confidence check. The basic workstation includes a 68000 frontend processor, a bit-slice graphics processor, 128 Kbytes of memory (expandable to 2 Mbytes), and a 19-in. color or monochrome monitor. The complete terminal is $22,575. **CalComp**, 2411 W La Palma Ave, Anaheim, CA 92801. **Circle 384**

Finally, a 2,000 LPM band printer somewhere between unreliable and too much money.

Twice the reliability.

Compare Hitachi Koki's FP-2000 against any low-priced 2,000 LPM band printer. The differences are immediately apparent. Instead of a light-duty souped-up machine, you'll find heavy-duty 3,000 LPM technology that has actually been slowed down for unprecedented reliability at 2,000 LPM. You'll find a printer that's suited for continuous operation — not just transaction printing. That delivers consistent OCR scannable print quality. And most significantly, that delivers an honest 1,200 hours MTBF.

Half the price.

Next compare the FP-2000 against the old industry standard for reliability. After evaluations, large OEMs have reported reliability every bit as good, print quality better, and the price almost half. They have also found comprehensive diagnostics, modular construction, IBM plug compatibility, and a quiet, compact design engineered for ease of use, ease of service and low-cost of ownership.

Call NSA today. The FP-2000 is part of a full family of band printers ranging from 300 to 2,000 LPM.

nsa
NISSEI SANGYO AMERICA, LTD.

HITACHI
Hitachi Koki Co.,Ltd.

Boston: 40 Washington Street, Wellesley Hills, MA 02181, (617) 237-9643
San Francisco: 460 E. Middlefield Road, Mountain View, CA 94043, (415) 969-1100

CIRCLE 163

DATACHECK II
THE ULTIMATE BREAKOUT BOX

We call it the ultimate breakout box because of all the features we pack into this pocket-size unit. Count them. 23 major features . . . all in an easy-to-use, reliable package. No wonder Datacheck II is the ultimate breakout box.

EASY TO USE

- Top to bottom functional panel layout.
- 25 individually numbered breakout switches.
- Four-state indication on all key signals (mark, space, clocking, undefined).
- LED's separated by signal source (DTE and DCE sides).
- Pocket-size, illustrated operators manual.
- 24-page instructional application notes.

RELIABLE

- Power off/reset button (conserves battery life).
- Molded ABS plastic case with metal hinge for durability.
- Two 9 volt batteries.
- Separate compartment for cable and jumpers.
- Gold-plated recessed pins.

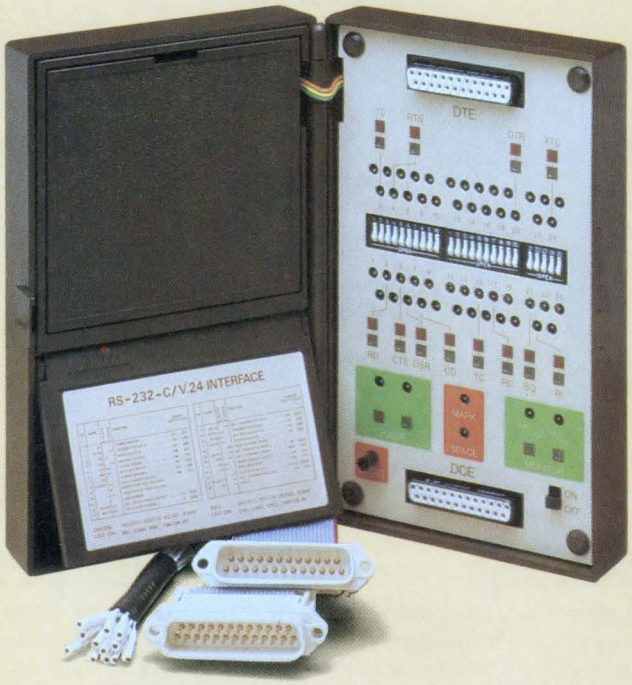

- Polycarbonate face (lettering can't wear off).
- Recessed LED's.

ADDED FUNCTION

- Detachable RS-232C ribbon cable (connection flexibility).
- Two-way and three-way jumper cables.
- Independent mark and space source pins.
- Dual tristate monitors for secondary signals.
- Power ON/OFF switch for extended interface reconfiguration.
- Dual pulse trap circuits (isolate intermittent signals).
- Optional Baud Rate Counter (measure transmission frequency).
- Optional Voice Frequency Monitor (listen to the condition of the phone line).
- Optional rechargeable batteries.
- Optional metal case.

ONLY $239.

NAVTEL

FOR THE TESTING TIMES

TELEPHONE (416) 669-9918

A division of Atelco

CIRCLE 164

Widely compatible laser printer runs at 12 pages/min

The LP4120 tabletop printer uses laser technology to achieve 12-page/min printing speeds. The printer employs an LP 120 controller, which has two 128-char internal fonts (bold face PS and Letter Gothic 15). The controller permits easy changing of two additional 128-char fonts. Many interfaces are available for connecting the LP4120 to a wide range of office equipment, including word and data processors, electronic mail systems, LANs, and microcomputers. **Ricoh Corp**, 5 Dedrick Pl, West Caldwell, NJ 07006.

Circle 385

Intelligent copier/printer merges text and graphics

The K-2 supports a single workstation or a network and uses electrophotographic technology. It is designed to print letter-quality text and graphics from virtually any host micro. The 68000 and 5¼-in. floppy allow the printer to perform program control functions such as page size and font selection, as well as generate complex vector graphics. Data streams are fed to a controller that computes the corresponding bit map pattern. This all-points addressable pattern is transferred to an LED array print head that focuses on an organic photoconductor belt. Price is below $10,000. **Kentek Information Systems, Inc**, 6 Pearl Ct, Allendale, NJ 07401.

Circle 386

Display terminal is supported by DISSPLA and the GKS

A display terminal, designated the 4106, represents an extension of 4100 series color graphics terminals. The 4106 features 640 x 480 resolution, interactive color selection, segment support of up to 4000 short vectors, and VT-100 compatibility. It supports up to 16 simultaneous colors, selected from a palette of 64, and has an interactive color interface that lets the operator select and mix colors on-screen with a special set of function keys. The units support the ANSI X3.64 standard for text editing. The display has a flicker-free 60-Hz noninterlaced refresh rate and shadow mask color CRT to deliver clarity and brightness. Software support includes DISSPLA and the Plot 10 GKS. The unit price is $6595. **Tektronix, Inc**, PO Box 500, Beaverton, OR 97077.

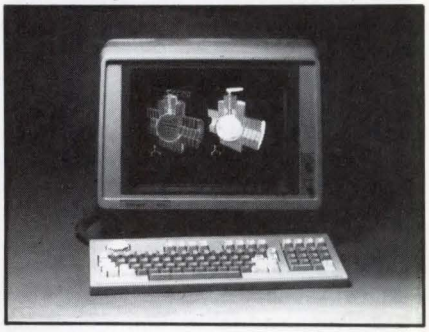

Circle 387

Hardware and software for the VMEbus

Hardware and software VMEbus products are summarized in a four-page, four-color bulletin that covers a wide range of single-board computers, memory boards, controllers, I/O boards, backplanes, and other items. **Plessey Microsystems,** Pearl River, NY.
Circle 425

Cables, connectors, and assemblies

A 78-page illustrated catalog depicts a range of Amphenol and Spectra-Strip flat cables, connectors, and assemblies. The catalog includes features, physical characteristics, electrical specs and dimensional line drawings. **Amphenol Products, an Allied Co,** Oak Brook, Ill.
Circle 426

Logic arrays in ECL/TTL

Line of subnanosecond logic arrays and CAD design and fabrication tools is outlined in an eight-page pamphlet. A range of diverse applications (including communications, automatic test equipment, and peripherals) is surveyed. **Applied Micro Circuits Corp,** San Diego, Calif.
Circle 427

Artificial intelligence

A 16-page illustrated brochure explains artificial intelligence (AI), developing applications, and the commerical market for AI hardware and software. **LISP Machine Inc,** Los Angeles, Calif.
Circle 428

Cables and connectors

Data communication gear is listed in a catalog of products that include RS-232-C cables, coaxial cables, gender changers, and null modems. **L-COM Data Products,** North Andover, Mass.
Circle 429

Voice response

In four pages, a full-color pamphlet explains the Voice Response System (VRS)—its attributes, applications, and specs. **Comsys B.V.,** Utrecht, the Netherlands.
Circle 430

Power distribution

A 40-page brochure details the electrical performance, packaging techniques, and design considerations of PC board power distribution component line. **Rogers Corp,** Mesa, Ariz. **Circle 431**

Solid modeling

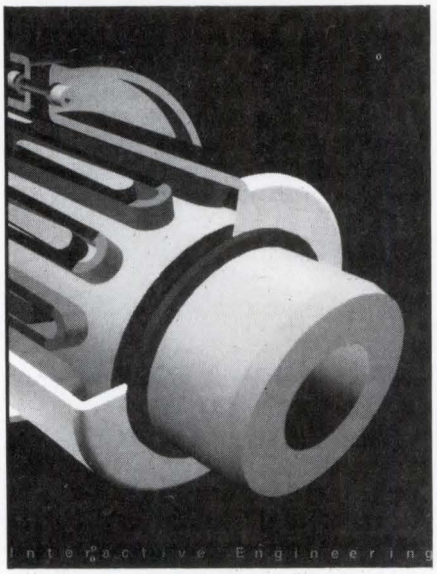

Analysis, results evaluation, analysis modeling, and solid modeling are graphically depicted in a pamphlet describing uses of a line of engineering software. **PDA Engineering,** Santa Ana, Calif.
Circle 432

Power supply models

Switching power supplies and line regulators/conditioners receive coverage in a 36-page catalog that lists products, options, and specs. Voltage current charts as well as dimensional drawings enhance presentation, **Deltron Inc,** North Wales, Pa.
Circle 433

Design data for MOSFET

Selection guide and cross-reference list details MOSFET design data. Included are low voltage threshold devices for telecommunications, superfast switching devices and high voltage, low leakage devices for test equipment use. **Ferranti Semiconductor,** Commack, NY.
Circle 434

Optimizing converters

Four-page application note shows how to obtain optimum results from fast settling low glitch D-A converters. Note cites critical design guidelines for wideband sensitive circuits, PC board layout, grounding, power supply decoupling, and noise minimization. **ILC Data Device Corp,** Bohemia, NY.
Circle 435

Miniature switches

Full line of miniature and subminiature switches is discussed in a 208-page catalog. **C&K Components Inc,** Newton, Mass.
Circle 436

Integrated circuit line

A 70-page short form data book describes a line of analog, digital, and custom ICs. Includes tech specs and pinouts. **Harris Corp, Semiconductor Sector,** Melbourne, Fla. **Circle 437**

Coaxial connectors

Miniature coaxial connectors are portrayed in a four-page brochure. Types and specs of small coaxial connectors, plugs, jacks and adapters are discussed. **National Tel-Tronics,** Meadville, Pa.
Circle 438

Schottky modules

Data sheets discuss 200- and 300-A Schottky diode modules—the FST200 and 300. Information covers dimensions, voltages, and electrical and thermal characteristics. Data sheets featuring the FST60 and 160 low profile Schottky modules are also available. **Siemens,** Iselin, NJ.
Circle 439

Digital-to-analog converters

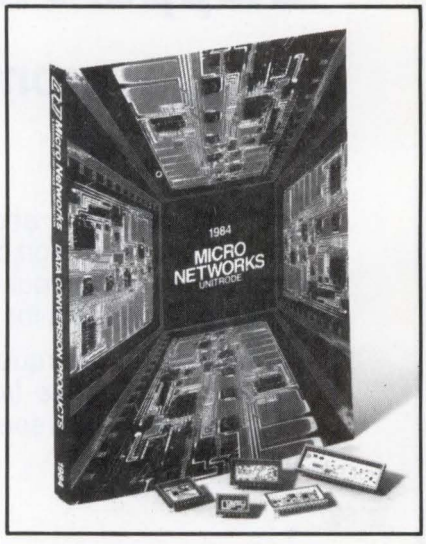

A 384-page catalog includes complete specs and application hints for line that includes high speed DIP A-D and D-A converters, data acquisition systems, and linear amplifiers. **Micro Networks,** Worcester, Mass.
Circle 440

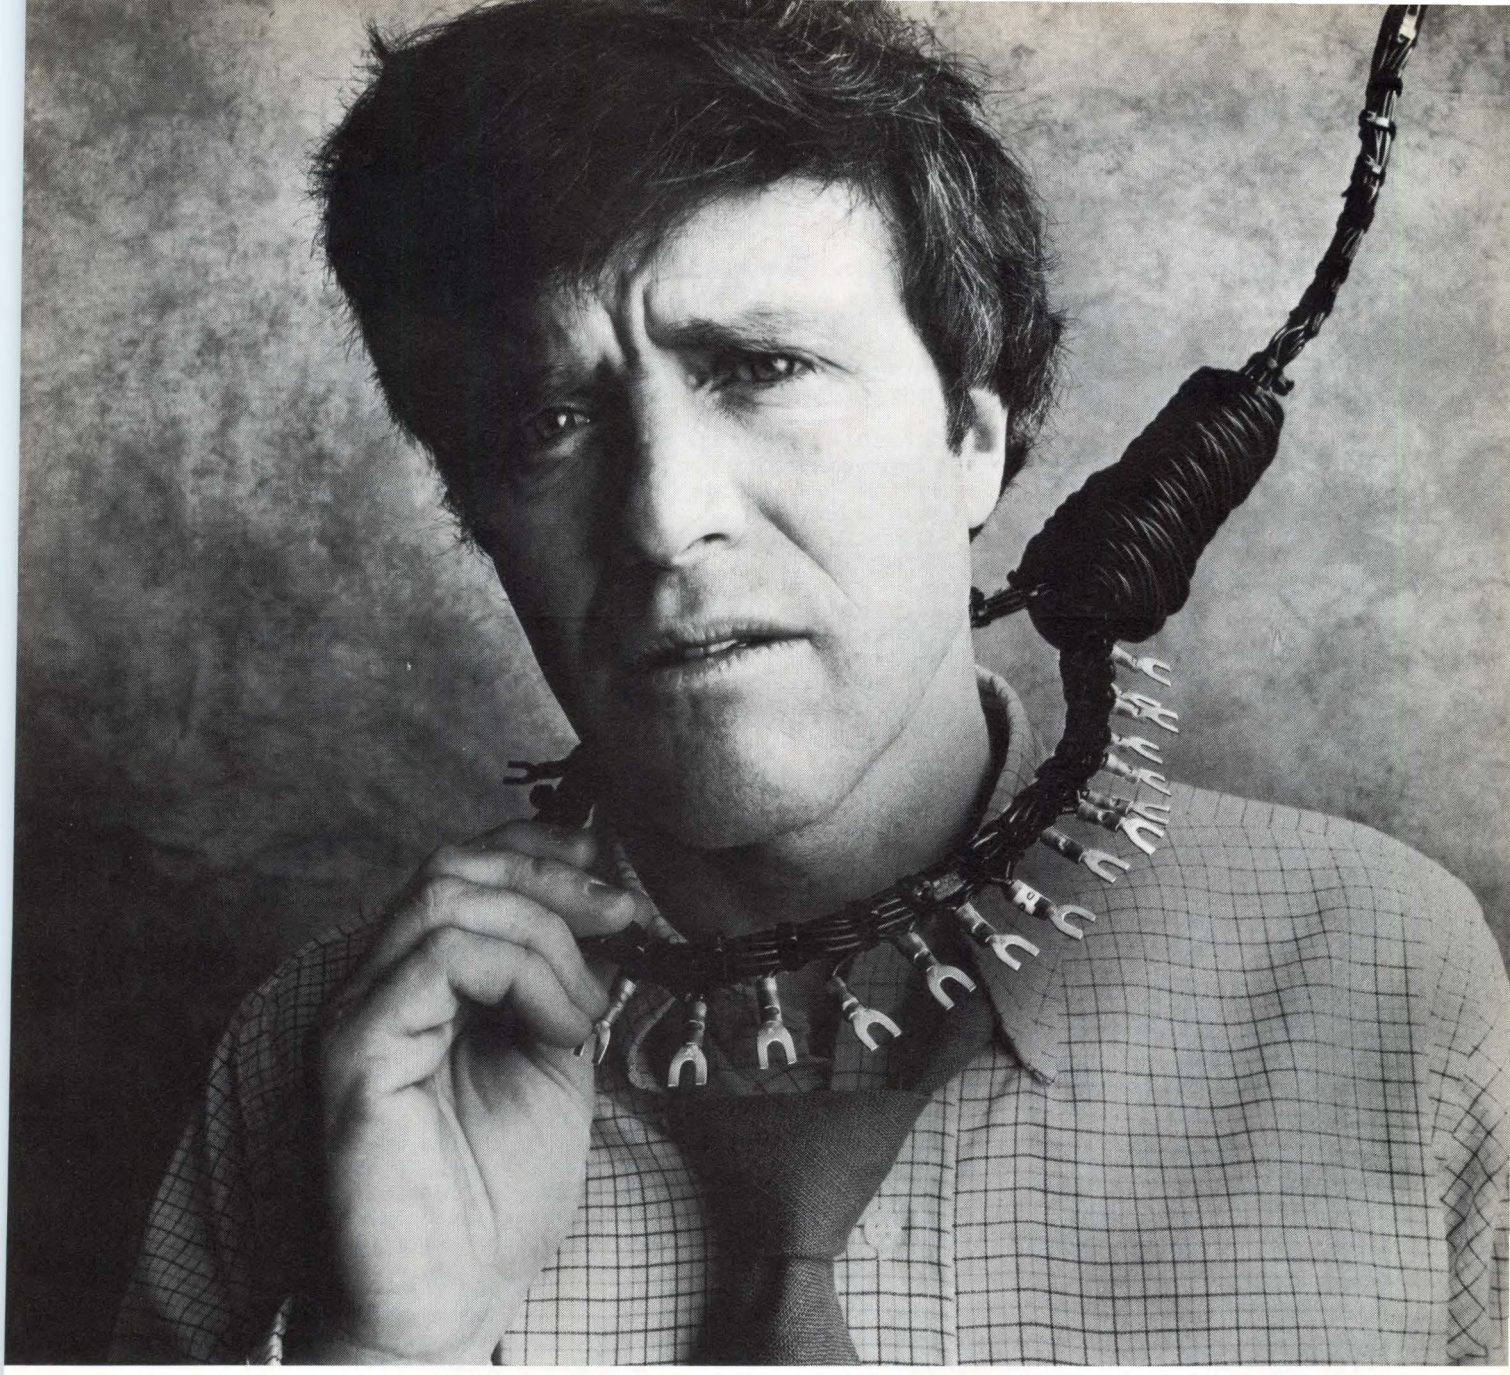

THE WRONG POWER DISTRIBUTION SYSTEM CAN MAKE YOU VERY HIGH STRUNG.

Some forms of power distribution, such as wire bundles and flat cables, can be a real pain in the neck. On the other hand, Rogers Mektron™ busbars distribute DC power simply, reliably and save space, too.

Our busbars are custom built to tight mechanical tolerances. They fit precisely, without the need for wire or cable interconnections. The planar configuration of the busbars' rugged copper conductors, the thinness of their dielectric and insulation enable you to distribute up to several hundred amps in compact units.

Most Rogers busbars can be installed in seconds with just a few screws. And while initial costs may look higher, our busbars will cut labor considerably. Not only will you have a more reliable power distribution system, but one which easily accommodates servicing and trouble-shooting of the system.

So before you get caught short with a burdensome bundle, look into Rogers busbars.

It's the sophisticated way to handle power.

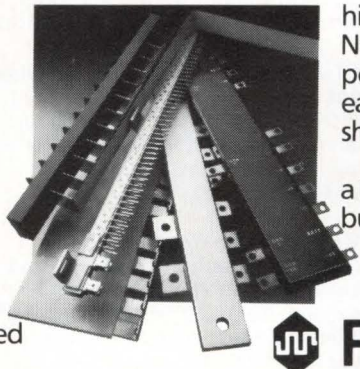

ROGERS

Rogers Corporation
Busbar Division, Box 700, Chandler, AZ 85224
602 963-4584

CIRCLE 166

CALENDAR

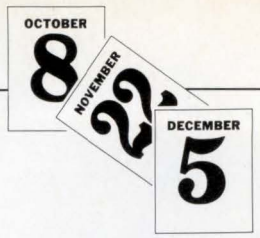

CONFERENCES

OCT 1-4—AUTOFACT 6 Conf & Exhibit, Anaheim Convention Ctr, Anaheim, Calif. INFORMATION: Society of Manufacturing Engineers, PO Box 930, Dearborn, MI 48121. Tel: 313/271-0777

OCT 2-4—Northcon High Technology Electronics Exhibit & Convention, Seattle Center Coliseum, Seattle, Wash. INFORMATION: Nancy Hogan, Electronic Conventions, Inc, 8110 Airport Blvd, Los Angeles, CA 90045. Tel: 213/772-2965

OCT 2-4—Mini/Micro Northwest Computer Conf & Exhibit, Seattle Center Flag Pavilion, Seattle, Wash. INFORMATION: Nancy Hogan, Electronic Conventions, Inc, 8110 Airport Blvd, Los Angeles, CA 90045. Tel: 213/772-2965

OCT 8-10—9th Conf on Local Area Networks, Minneapolis, Minn. INFORMATION: Harvey A. Freeman, Architectural Technology Corp, PO Box 24344, Minneapolis, MN 55424.

OCT 8-10—ACM Annual Conf: The Fifth-Generation Challenge, San Francisco Hilton Hotel, San Francisco, Calif. INFORMATION: Karen Duncan, Chairman, ACM '84, 15 Parsons Way, Los Altos, CA 94022.

OCT 8-11—Integrated Information Technology Conf, Dallas Convention Center, Dallas, Tex. INFORMATION: Jill Nieman, National Trade Publications, Inc, 2111 Eisenhower Ave, Suite 400, Alexandria, VA 22314. Tel: 703/683-8500

Invitational Computer Conf
OCT 10—Cherry Hill Hyatt Hotel, Cherry Hill, NJ; **OCT 23**—Hilton Inn, Englewood, Colo. INFORMATION: Suzanne Hubner, B. J. Johnson & Assoc, Inc, 3151 Airway Ave, #C-2, Costa Mesa, CA 92626. Tel: 714/957-0171

OCT 10-12—Design Automation Workshop, East Lansing, Mich. INFORMATION: Harry Hayman, IEEE Computer Society, PO Box 639, Silver Spring, MD 20901. Tel: 301/589-8142

OCT 10-12—LOCALNET '84, Sheraton Harbor Island Hotel, San Diego, Calif. INFORMATION: Online Conferences, Inc, Suite 1190, 2 Penn Plaza, New York, NY 10121. Tel: 212/279-8890

OCT 11-12—Peripheral Array Processor Conf, Copley Plaza Hotel, Boston, Mass. INFORMATION: Charles A. Pratt, Simulation Councils, Inc, PO Box 2228, La Jolla, CA 92038. Tel: 619/459-3888

OCT 15-17—The Future of Optical Memories, Loew's Summit Hotel, New York, NY. INFORMATION: Joanna Spilman, Technology Opportunity Conference, PO Box 14817, San Francisco, CA 94114. Tel: 415/626-1133

OCT 15-18—Conf on Ada Applications & Environments, Sheraton Midway Hotel, St. Paul, Minn. INFORMATION: Harry Hayman, IEEE Computer Society, PO Box 639, Silver Spring, MD 20901. Tel: 301/589-8142

OCT 16-18—Unix Operating System Expo & Conf, Sheraton Centre Hotel, New York, NY. INFORMATION: Robert P. Birkfield, National Expositions Co, Inc, 14 W 40th St, New York, NY 10018. Tel: 212/391-9111

OCT 16-19—Int'l Test Conf (Cherry Hill '84), Franklin Plaza Hotel, Philadelphia, Pa. INFORMATION: Harry Hayman, IEEE Computer Society, PO Box 639, Silver Spring, MD 20901. Tel: 301/589-8142

OCT 17-19—Design Automation Workshop. East Lansing, Mich. INFORMATION: Donald E. Thomas, Dept of Electrical Engineering, Carnegie-Mellon Univ, Pittsburgh, PA 15213. Tel: 412/578-3545

OCT 22-24—13th Annual Workshop on Applied Imagery Pattern Recognition, Sheraton Inn-Washington Northwest, Silver Spring, Md. INFORMATION: IEEE Computer Society, PO Box 639, Silver Spring, MD 20901. Tel: 301/589-8142

OCT 22-25—ISA '84, Astrohall, Houston, Tex. INFORMATION: Instrument Society of America, 67 Alexander Dr, Research Triangle Park, NC 27709. Tel: 919/549-8411

OCT 22-26—Annual Industrial Electronics Conf (IECON '84), on Industrial Applications of Microelectronics, Keio Plaza Intercontinental Hotel, Tokyo, Japan. INFORMATION: Frank A. Jur, Bechtel Corp, 45 Fremont St, MS 45/17A26, San Francisco, CA 94109. Tel: 415/768-3023

OCT 30-31—Flat Information Display Conf, Red Lion Inn, San Jose, Calif. INFORMATION: Murray Disman, Int'l Planning Information, Inc, 164 Pecora Way, Portola Valley, CA 94025. Tel: 415/854-7306

OCT 30-NOV 2—Wescon High Technology Electronics Exhibit and Convention, Anaheim Convention Ctr, Anaheim, Calif. INFORMATION: Dale Litherland, Electronic Conventions, Inc, 8110 Airport Blvd, Los Angeles, CA 90045. Tel: 213/772-2965

NOV 12-15—IEEE Int'l Conf on Computer Aided Design, Santa Clara, Calif. INFORMATION: John A. Domiter, Bell Telephone Labs, 4K523, Holmdel, NJ 07733. Tel: 201/949-6675

NOV 13-17—Elektronica, Munich Trade Fair Center, Munich, W Germany, INFORMATION: Kallman Assoc, 5 Maple Ct, Ridgewood, NJ 07450. Tel: 201/652-7070

NOV 14-18—Comdex/Fall, Las Vegas Convention Center, Las Vegas, Nev. INFORMATION: The Interface Group, 300 First Ave, Needham, MA 02194. Tel: 617/449-6600

NOV 16-17—Forth Interest Group Convention, Hyatt Palo Alto, Palo Alto, Calif. INFORMATION: Forth Interest Group, PO Box 1105, San Carlos, CA 94070. Tel: 415/962-8653

DEC 3-9—Int'l Microcomputer Conf & Display 1984, China, The Guangdong Scientific Hall, Guangdong, PRC. INFORMATION: Meridian Technology Exhibitions Ltd, Rm 1201 Kai Tak Commercial Bldg, 317 Des Voeux Rd, C, Hong Kong

DEC 9-12—IEEE Int'l Electron Devices Meeting, San Francisco Hilton and Towers, San Francisco, Calif. INFORMATION: Melissa M. Widerker, Courtesy Associates, Inc, 665 15th St, NW, Suite 300, Washington, DC 20005

DEC 11-13—Fifth-Generation and Supercomputer Symposium, Rotterdam, The Netherlands. INFORMATION: Rotterdam Tourist Office, Stadhuisplein 19, 3012 AR Rotterdam. Tel: (010) 14 14 00

SHORT COURSES

Knowledge-Based Systems, AI and 5th Generation Computing
OCT 2-5—Los Angeles, Calif, **OCT 30-NOV 2**—Washington, DC. INFORMATION: Ruth Dordick, Integrated Computer Systems, 6305 Arizona Place, Los Angeles, CA 90045. Tel: 213/417-8888

OCT 17-19—Local Area Networks, Washington, DC; **OCT 24-26—Experts on Networks,** Boston, Mass; **NOV 26-28—X.25 and Other Protocols,** Washington, DC. INFORMATION: Technology Transfer Institute, 741 10th St, Santa Monica, CA 90402. Tel: 213/394-8305

OEM Decision Makers

"Only the Invitational Computer Conferences bring the latest OEM computer and peripheral products to your front door. You'll find us there!"

BEN WANG, President
Wangtek

And you'll find other top OEM manufacturers, such as IBM, Control Data, DEC, Fujitsu, NEC and Seagate, to name a few.

In their 14th year, the "OEM Only" Invitational Computer Conferences bring you, the volume buying decision makers, together with the key suppliers of computer and peripheral products. The ICCs, a series of ten, one-day regional shows are convenient to where you live and work. The social business setting makes it easy for you to meet potential suppliers one-on-one, and attend high tech seminars of your choice. As an invited guest, there is no cost to you.

Hear what the OEM manufacturers have to say, learn more about their products, and remember, you may attend "by invitation only."

1984/85 U.S. ICC Locations

Sept. 6, '84	Newton/Boston, MA
Sept. 25, '84	Southfield/Detroit, MI
Oct. 10, '84	Cherry Hill, NJ
Oct. 23, '84	Englewood/Denver, CO
Jan. 8, '85	Irvine, CA
Jan. 29, '85	Houston, TX
Jan. 31, '85	Dallas, TX
Feb. 26, '85	Ft. Lauderdale, FL
Mar. 19, '85	Palo Alto, CA
Apr. 2, '85	Nashua, NH / No. MA

Call your local OEM supplier for your invitation or fill out the coupon and mail to:

B. J. Johnson & Associates, Inc.
3151 Airway Ave., #C-2
Costa Mesa, CA 92626
Phone: (714) 957-0171
Telex: 188747 TAB IRIN

Yes! I need an invitation to your "OEM Only" ICC. The nearest ICC to me is: _____

I buy in volume:

- ☐ Computers
- ☐ Disk/Tape Drives
- ☐ Controllers / Interfaces
- ☐ Terminals / Graphic Displays
- ☐ Software
- ☐ Printers
- ☐ Memory Boards
- ☐ Modems / Multiplexers
- ☐ Power Supplies

Name _____

Title _____

Company/Division _____

Address _____

City _____ State _____ Zip _____

Mail To: B. J. Johnson & Associates, Inc., 3151 Airway Avenue, #C-2, Costa Mesa, CA 92626
Phone: (714) 957-0171 Telex: 188747 TAB IRIN

CIRCLE 167

DESIGNER'S BOOKCASE

HANDBOOK OF SOFTWARE ENGINEERING

By *Charles R. Vick and C.V. Ramamoorthy*

Turn to this massive handbook for authoritative coverage of all important tools, techniques, and concepts used to develop computer programs. Emphasizing quality assurance, it explains simulation modeling, data design, operating systems, management of software development, software testing technology, data base management, and system evaluation.

$62.50 **Circle 445**

THE MEASUREMENT OF VISUAL MOTION

By *Ellen Catherine Hildreth*

The organization of movement in the changing image that reaches the eye provides our visual system with a valuable source of information for analyzing the structure of our surroundings. This book examines the measurement of this movement and the use of relative movement to locate the boundaries of physical objects in the environment. It investigates the nature of the computations that are necessary to perform this analysis by any vision system, biological or artificial.

$30.00 **Circle 447**

COMPUTER-ASSISTED DATA BASE DESIGN

By *George V. Hubbard*

Find out how to computerize many aspects of data base design. This state-of-the-art guide explains how to conduct thorough and consistent analyses of data requirements and design tradeoffs. Fully discussed are design refinement, physical design choices and options, space and time calculations, dictionary interfaces, access methods, and more.

$26.95 **Circle 448**

MICROPROCESSOR-MICROCOMPUTER TECHNOLOGY

By *Frederick F. Driscoll*

Focusing on three microcomputers—the 8080A, the 6800, and the 6502—Frederick Driscoll provides easy-to-follow guidance on a wide range of microcomputer functions. He shows step by step how to design, load, test, debug, and document programs. Practical instruction is also given on using flowcharts and subroutines.

$29.95 **Circle 446**

SOFTWARE SYSTEM TESTING AND QUALITY ASSURANCE

By *Boris Beizer*

Now you have a comprehensive guide to system testing that ensures reliable, robust, high-quality software. In nontechnical language, Boris Beizer covers the gamut from unit testing to system testing. He demonstrates new and effective techniques for security testing, recovery testing, configuration testing, performance testing, integration testing, and debugging.

$34.50 **Circle 449**

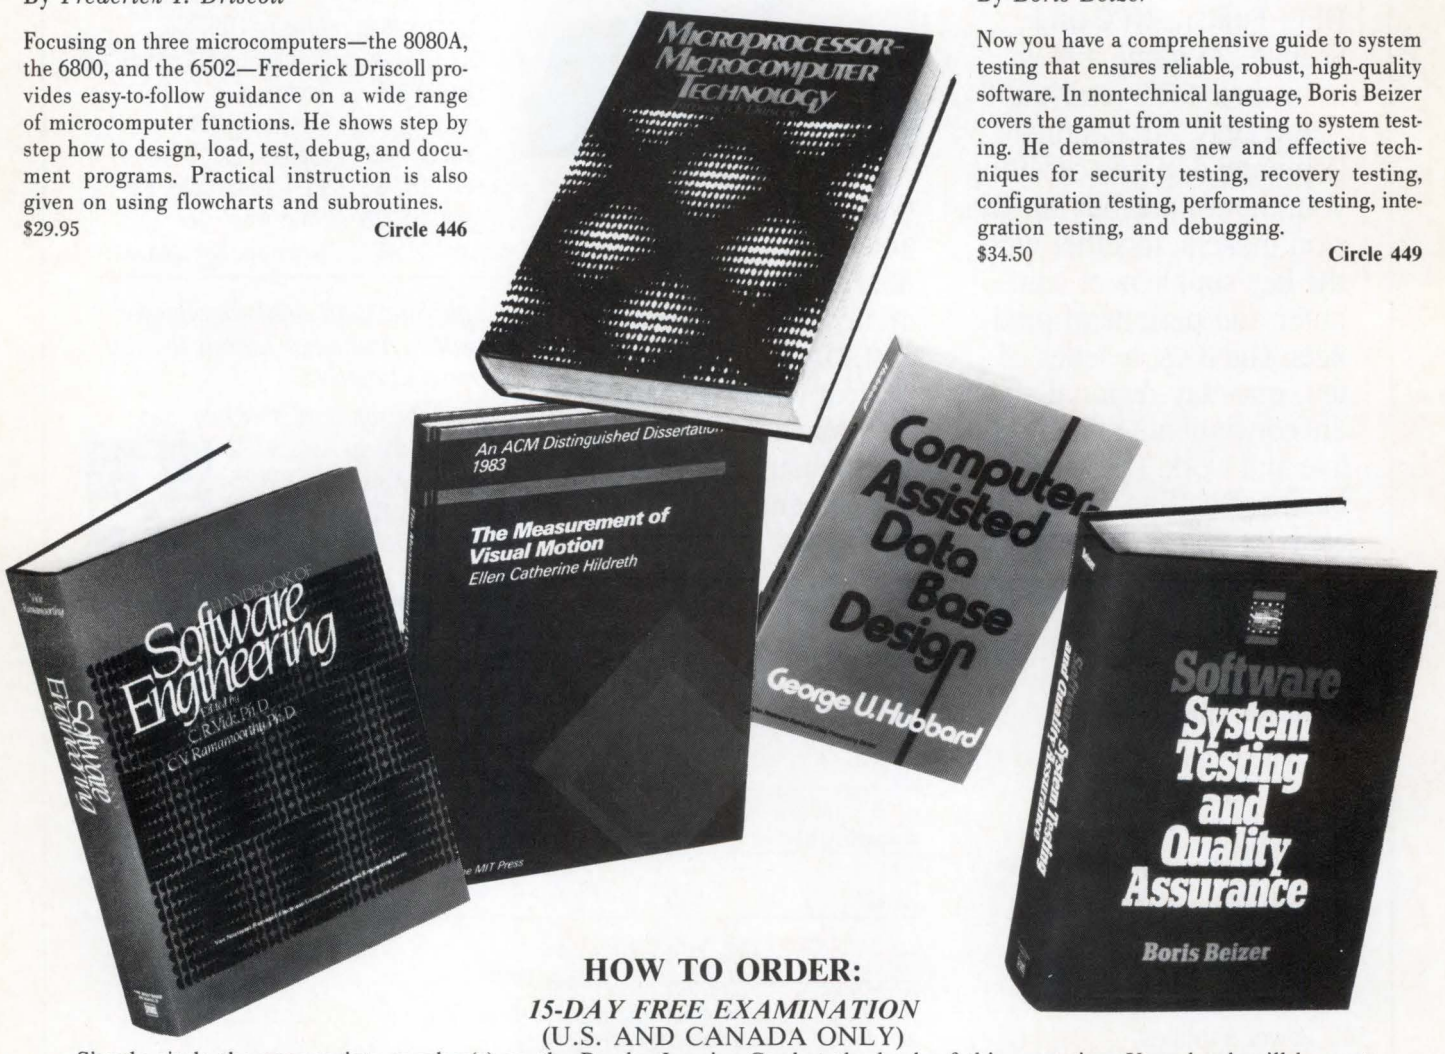

HOW TO ORDER:

15-DAY FREE EXAMINATION
(U.S. AND CANADA ONLY)

Simply circle the appropriate number(s) on the Reader Inquiry Card at the back of this magazine. Your book will be sent to you for your **15-day free trial**. If you are satisfied keep the book and an invoice will follow. Otherwise return the book by the end of the 15-day period, and owe nothing.

SYSTEM SHOWCASE

CIRCLE 475 for rates and information, OR CALL: SHIRLEY LESSARD 800-225-0556, in MA 617-486-9501

INTELLIGENT STAND ALONE RS-232-C or IEE 488 CARTRIDGE TAPE SYSTEM

- Stores up to 5.3M of Binary or ASCII data.
- Intelligent search and retrieval.
- Standard power fail restart or optional power fail with NO Data Loss
- IEEE-488 and or RS-232-C Ports with data rates up to 20,000 characters/sec.
- Large input buffer allows unit to accept data non-stop.
- Applications: Data Logging, Control system archiving, Program loading & storage. Back up, Telephone switch monitoring, Auto-polled remote data storage.
- Price: Under 2000 in Quantities

ALGO

9198-C Red Branch Rd
Columbia, MD 21045
301-730-7442

CIRCLE 476

LIKE NEW PRODUCTS

For free catalog,
phone toll-free (800)**225-1008**
In Massachusetts (617)**938-0900**
GENSTAR REI SALES COMPANY
6307 DeSoto Ave, Suite J,
Woodland Hills, CA 91367.

CIRCLE 477

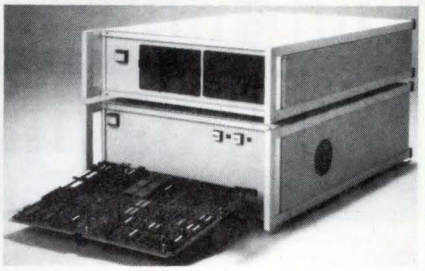

BRI Introduces

ITS UNIX* V Development System/ Workstation featuring BRI68K SBC/68010 CPU 8, 10 or 12.5MHz 1.28Mb RAM Multi-user UNIX* SYSTEM V/Real-Time Executive 30Mb (Formatted) Winchester 648Mb (Formatted) Floppy ASM, C, Fortran, Pascal, Basic, Cobol Development System = Target System Avoid Incompatibilities. **BRI,** 748 Cascadilla St, Ithaca, NY 14850. Tel: (607) 273-3300

*TM ATT

CIRCLE 478

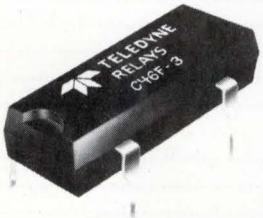

HIGH SPEED POWER FET AC/DC SOLID STATE RELAY

Teledyne's C46F/C47F series are pin compatible replacements for DIP reed relays where low EMI swtiching, high reliability and long life are required. Switches AC or DC up to 400V. On-resistance as low as 7 ohms. Control voltage range is 3.8 to 32VDC. Optical coupling provides 1500VDC isolation. Features no offset voltage, low off-state leakage, and high switching speed. $5.90 ea for 5000 pcs. **Teledyne Sold State Products,** 12525 Daphne Ave., Hawthorne, CA 90250 (213) 777-0077.

CIRCLE 479

6809
Single Board Computer

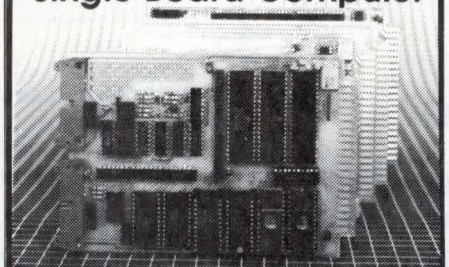

6809 MPU, 2 serial ports, 4 parallel ports, RAM, EPROM, real-time clock, watchdog timer, 44-pin 4.5" × 6.5" PCB

EXPANSION MODULES: RAM, EPROM, CMOS RAM/ battery, analog I/O, serial I/O, parallel I/O, counter/ timer, IEEE-488, EPROM programmer, floppy disks, cassette, breadboard, keyboard/display.

WINTEK

Wintek Corp.
1801 South Street
Lafayette, IN 47904
317-742-8428

CIRCLE 480

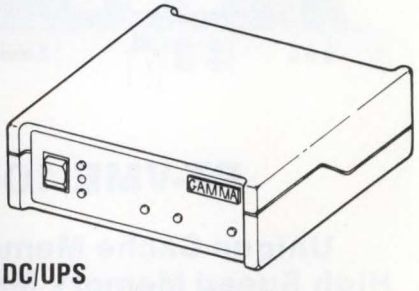

DC/UPS

Complete protection for Kaypro 2, 4, and 10. Uninterruptible, AC 86volts to 256 volts. Reserve from 30 minutes to 2 hours. $469.00 US.

GAMMA

6253 Hollywood Boulevard, Suite 1129
Los Angeles, California 90028
(213) 463-2345 Telex: 371 5824

CIRCLE 481

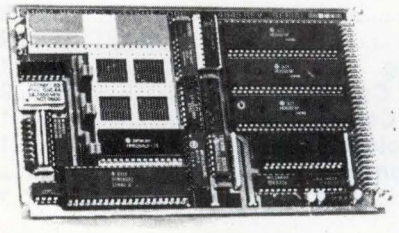

32bit Single Board Computer

68008 Single Board Computer. 100 by 160mm Eurocard. 3 RAM/EEPROM/EPROM sockets for 2K-64K byte devices (2K RAM included). 6 8 bit Bi-Dir I/O ports with handshake and individually vectored interrupts. 4 16 bit timers. Dual async serial RS232 ports. 256b nvRAM. Diagnostics LEDs. Multitasking kernal option. 295/unit. Quantity discounts available. **HERITAGE SYSTEMS CORPORATION,** PO Box 10588, Greensboro, NC 27404. Tel: (919) 274-4818.

CIRCLE 482

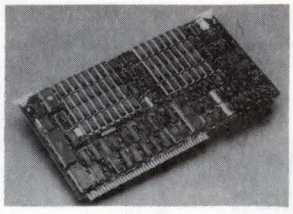

INTEL USERS—SPEED UP YOUR DEVELOPMENT BY 3 TIMES WITH THE ORIGIN DISK CACHE/PRINTER SPOOLER

For series 2 or 3, MDS 800 and NDS II work stations. Works with single or double density floppies. Totally transparent. Automatic allocation of 480 K-Bytes of cache memory. Built in **196 K-Byte print spooler.** The model OR-88C is $2850 each and in stock. **ORIGIN, Inc,** 9136 Gibson St, LA, CA 90034. Tel: (213) 202-0227. TLX704809.

CIRCLE 483

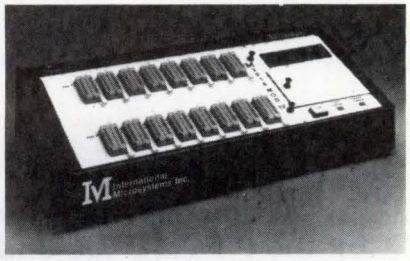

MULTI-SET EPROM PROGRAMMER

IM3016 programs up to 16 DIFFERENT IMAGES in one cycle. Ideal for engineering and/or production environments using large EPROM SETS. Electrically isolated sockets, one-key control and APPROVED fast algorithms. Software selection for all 5V EPROMs and CMOS versions including 2764, 2764A, 27128, 27128A & 27256. **INTERNATIONAL MICROSYSTEMS, INC.** 11554 C Avenue, Auburn, CA 95603. Tel: (916) 885-7262.

CIRCLE 484

AD INDEX

AD INDEX

AD INDEX

*International Issues Only

CD* SALES OFFICES

Home Office: 119 Russell St., Littleton, MA 01460 (617) 486-9501

Marketing Director, Robert Billhimer
Ad Traffic Coordinator, Debra Friberg
Systems Showcase, Shirley Lessard
Postcard Deck, Shirley Lessard
Classified/Recruitment, Shirley Lessard
List Rental, Robert Dromgoole

New England/Upstate New York
Barbara Arnold
119 Russell St.
Littleton, MA 01460
(617) 486-9501

Mid-Atlantic/Southeast
Richard V. Busch
40 Stony Brook Lane
Princeton, N.J. 08540
(609) 921-7763

Midwest & Colorado
Berry Conner
88 West Schiller St., Suite 2208
Chicago, IL 60610
(312) 266-0008

Southwest
Steve Lassiter
1200 S. Post Oak Blvd.
Houston, TX 77056
(713) 621-9720

Southern California
Buckley/Boris Associates
Tom Boris, John Sabo
2082 SE Bristol, Suite 216
Santa Ana, CA 92707
(714) 957-2552

Northern California
Buckley/Boris Associates
Tom Boris, John Sly
920 Yorkshire Drive
Los Altos, CA 94022
(415) 964-4232

Northwest
Buckley/Boris Associates
Tom Boris
2082 SE Bristol, Suite 216
Santa Ana, CA 92707
(714) 957-2552

International
International Sales Manager
Eric Jeter
1200 S. Post Oak Blvd.
Houston, TX 77056
(713) 621-9720

U.K. and Scandinavia
David Betham-Rogers, David M. Levitt
6th Floor, Alliance House
12 Caxton Street
Westminster, London SW1H OQS
Tel: 01-222 0744
Telex: 919775

Norway
Sverre Follaug, Jr.
Studio 14
Viggo Ullmannsgt 14
P.O. Box 250
3701 Skein
NORWAY
Tel: 035 22908

France, Belgium and S. Switzerland
Daniel R. Bernard
247, Rue Saint Jacques
75005 Paris
Tel: (1) 354.55.35
Telex: 250 303

Holland, Austria, W. Germany, Switzerland & Eastern Europe
Heinz Gorgens
Parkstrasse 8a
D-4054 Nettetal 1—Hinsbeck (F.R.G.)
Tel: (0 21 53) 8 99 88/89
Telex: (17) 2153310

Japan
Sumio Oka & Shigeo Aoki
International Media
Representatives Ltd.
2-29 Toranomon 1-chome
Minato-ku, Tokyo 1Q5 Japan
Tel: 502-0656 Telex: J22633

Southeast Asia
Seavex, Ltd.
Jim Stanton, Steve Marcopoto
19 Tanglin Road, #05-49
Singapore 1024
REPUBLIC OF SINGAPORE
Telex: 35539 RS
Paul Meyer
503 Wilson House
19-27 Wyndham St., Central
HONG KONG
Tel: 5-260149
Telex: 60906 SEAVEX HX

*TM

HUGHES

THE PEOPLE BEHIND ELECTRO-OPTICS

CREATE NEW VISIONS

From the demonstration of the first working laser to today's VHSIC. The people of Hughes Electro-Optical & Data Systems Group are talented . . . determined . . . proud. They challenge themselves at the leading edge of technology. In an environment that both stimulates and actively supports a diversity of technical contributions. Resourceful science . . . coupled with creative engineering. Hughes people have made it an art.

The people of Electro-Optical & Data Systems Group are now operating from their new, ultra-modern facility . . . a virtual city of technology. It's one of the largest and most sophisticated of its kind. On-site recreational facilities provide a healthful work environment. It has consolidated both engineering and manufacturing functions into one location, enhancing essential communications, while contributing significantly to all time sales milestones.

At Hughes, there is a stimulating relationship between the people and their work. Between the individual and the team . . . the team and the group . . . the group and the company. A relationship that provides opportunity for substantial individual contribution. That's what Hughes Electro-Optical & Data Systems Group is all about. People. People with vision and dedication. People participating in extraordinary ways.

People like you.

Current openings:
• Image Processing Engineers • Component Application Engineers • Circuit Design Engineers • Test Engineers • Material/Process Application Engineers • Thermal Systems Engineers • Electro-Optical Systems Engineers • Failure Analysis Engineers • Sensor Design Engineers • Semiconductor Engineers • Optic Design Engineers • Laser Engineers • Computer Hardware/Software Architecture • Programmers/ Analysts

Send your resume to: Hughes Electro-Optical & Data Systems P.O. Box 913, Dept. DC-9, El Segundo, CA 90245. Proof of U.S. Citizenship Required. Equal Opportunity Employer.

HUGHES
AIRCRAFT COMPANY

ELECTRO-OPTICAL AND DATA SYSTEMS GROUP

SHAPE YOUR OWN FUTURE.

General Dynamics Data Systems Division helps shape the future of many significant programs at its major locations in San Diego, California; Fort Worth, Texas; and Norwich, Connecticut; as well as at satellite locations including Detroit, Michigan, and Pomona, California. We provide diverse support functions for such high-technology programs as the F-16 multimission fighter/attack aircraft, the M1 main battle tank, nuclear-powered submarines, and the entire family of cruise missiles. Throughout our division you'll find a variety of opportunities to apply your own scientific and engineering expertise to create a more exciting future.

The Data Systems Division gives you the chance to join one of the most skilled teams in the industry today, and offers excellent salaries and benefits.

If you're interested in shaping your own future on our innovative support team, one of our opportunities listed below may be just right for you.

For immediate consideration, send your resume to the Vice President/Director at the Data Center of your choice.

PRODUCT SOFTWARE

— Bachelor's and/or Master's degree, and 3-8 years' experience.
— SKILLS: FORTRAN, JOVIAL, PASCAL, Ada, Assembly Languages, Applied Math, Data Bases, Operating Systems, Documentation (MIL Standards).
— APPLICATIONS: Command & Control Software (Guidance, Navigation, C^3I, Display Systems, Executive & System Support Software), Mission Planning, Data Handling & Communication, Automatic Testing Equipment/ Simulations, Image Processing, Estimation & Control Theory.
— Embedded Systems: Real-Time Software.

CAD/CAM

— Bachelor's and/or Master's degree, and 3-8 years' experience.
— SKILLS: FORTRAN, Ada, Assembly Languages, IBM, CDC & VAX Operating Systems, Computational Geometry, Information Modeling & Data Dictionaries.
— APPLICATIONS: Turn-key Graphic Systems, Solid Modeling, Robotics, CNC-DNC, Real-Time Processing Control, Group Technology.
— CADAM, CAT/A, Computervision, SC/CARDS, Model 204, Systems Engineering, Group Management.

ENGINEERING SYSTEMS

— Bachelor's and/or Master's degree, and 3-8 years' experience.
— SKILLS: Simulation Languages, FORTRAN, COBOL & IBM Assembler, TSO, SPF, DISSPLA, RAMIS, Scientific Programming and Microcomputer experience.
— APPLICATIONS: Simulation, SLAM, Manufacturing, Operations Research, Image Processing, Graphics Application Development, Sneak Circuit Analysis, Program Marketing.

U.S. CITIZENSHIP REQUIRED
Equal Opportunity Employer, M/F

GENERAL DYNAMICS

Data Systems Division

WESTERN CENTER
P.O. Box 85808, Drawer 002
San Diego, CA 92138

CENTRAL CENTER
P.O. Box 748, Drawer 002
Fort Worth, TX 76101

EASTERN CENTER
100 Winnenden Rd., Drawer 002
Norwich, CT 06360

You can't beat the system.

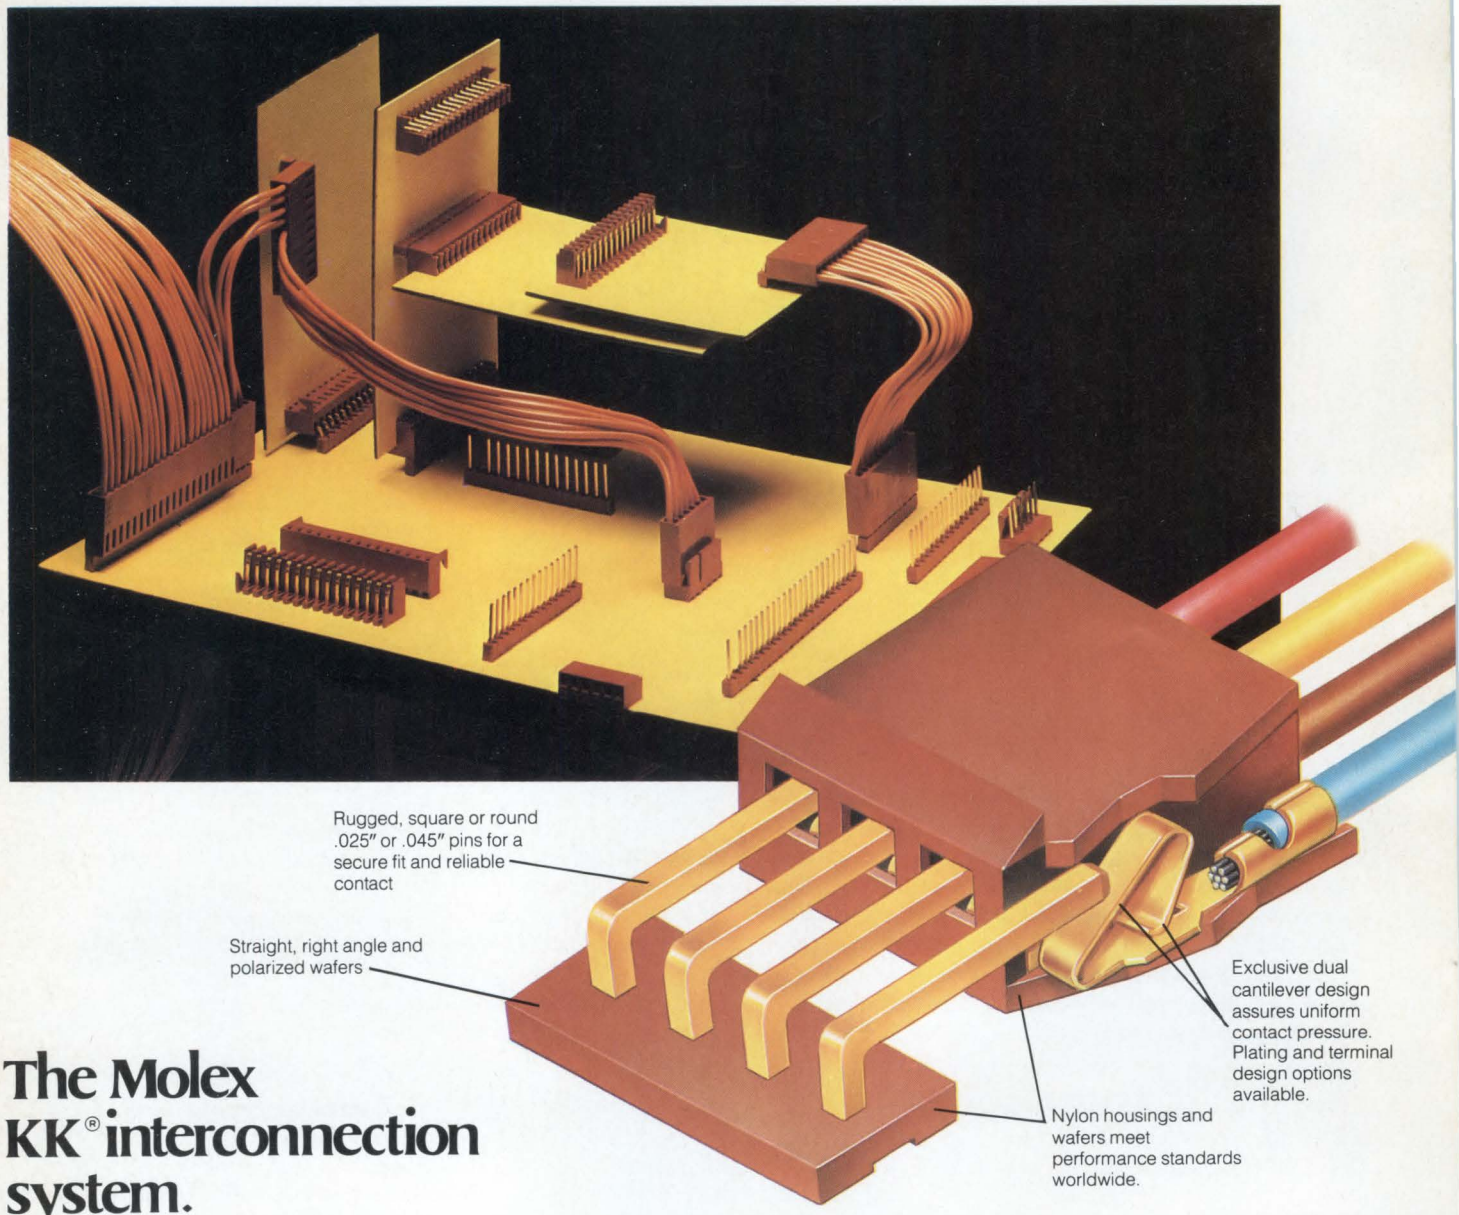

The Molex KK® interconnection system.

Built around only eleven versatile components, our "proven" KK® interconnection system provides the options you need for your wire-to-board and board-to-board applications.

Available on .098″ (2,5mm), .100″ (2,54mm), .156″ (3,65mm) and .200″ (5,08mm) centers, our modular, building block system of connectors offers you a simple and affordable answer to your interconnection needs.

The Molex KK® system also includes a complete line of crimp tooling and the most sophisticated pinsetting equipment available.

Our system can be used to create thousands of different configurations thereby allowing each user the flexibility to build a system that is precisely suited for his application.

And, with Molex's worldwide network of distributors and representatives we're prepared to give you prompt, efficient service, no matter where you're located.

For more information or samples of the KK® modular interconnection system, contact the Molex office nearest you.

First in Customer Service ...Worldwide

molex

CIRCLE 172

Molex Corporate Headquarters: 2222 Wellington Ct., Lisle, IL 60532 U.S.A. Phone: (312) 969-4550 Telex: 27-0072/25-4069

Suddenly, everyone's headed for MARS.

The MARS-432 32-bit, programmable, floating point array processor.

And with good reason. Because the MARS-432 has opened up a new world of speed, power and ease-of-use that's hard for anyone to resist.

The MARS-432 already interfaces with some of this world's leading computers — DEC, Apollo, Elxsi — to provide users with a new level of computational power. Interfaces for other leaders such as IBM, Perkin-Elmer, and Gould/SEL are scheduled to arrive soon.

Simply put, we're setting the direction in state-of-the-art array processors with features such as:

Programming Ease

All of the computational power of an array processor doesn't mean much if accessing that power requires days of tedious programming, debugging and reprogramming. That's why we engineered the MARS-432 with an architecture specifically designed to support a FORTRAN compiler and a screen-oriented debugging system that make accessing and utilizing its raw power a very civilized process.

The MARS-432 also provides:

- ☐ *A Microcode Development System for off-line program development.*
- ☐ *An AP Run Time Executive Support Package (AREX) for simplified processor initialization, I/O operations, and array function executions.*
- ☐ *Applications Libraries for math, signal processing, and image processing.*

Speed

- ☐ *Add and multiply times of 100ns.*
- ☐ *Computational power of 30 megaflops.*
- ☐ *Computes a 1024-point complex FFT in 1.7ms.*
- ☐ *DMA transfers at I/O bus rates of 20 megabytes/sec.*
- ☐ *Data memory write or two reads in 100ns.*
- ☐ *Memory paging for uninterrupted processing during I/O transactions.*

Impressive Memory

Program memory contains a physical address space of 4K words and a virtual address space of 64K words via a cache configuration. Data memory contains a physical address space of 16 million words.

The MARS-432 from Numerix: a journey to faster, more affordable array processing power. With programming ease that sets it worlds apart.

Going our way?

IBM

PERKIN ELMER

GOULD/SEL

VAX

MICRO VAX

APOLLO

ELXSI

For additional information on the MARS Family of High Speed Array Processors, write or call:

NUMERIX

maximum power minimum effort

Numerix Corp. 320 Needham Street, Newton, MA 02164-1594 Tel. 617-964-2500.

CIRCLE 173

Before TEMPLATE®, graphics software that ran on mainframes and minis and micros all looked like this.

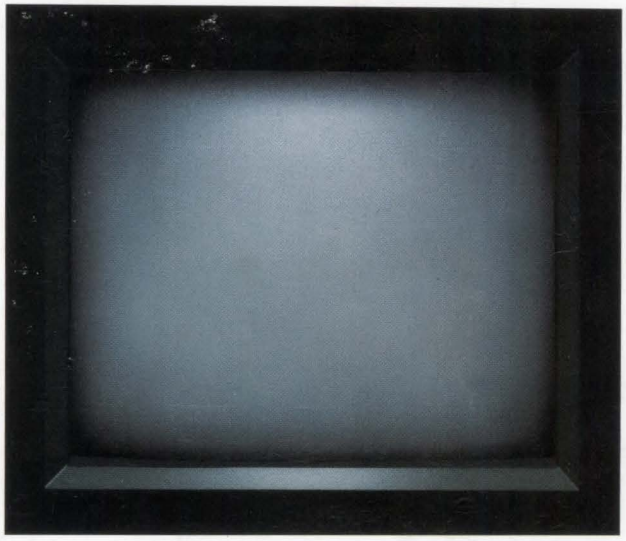

Which is to say, nonexistent. Some manufacturers make graphics software packages for mainframes. Others for minis. Still others make it for micros. But no one made high-level graphics software that ran on all three.

TEMPLATE just changed all that. By becoming the *only* high-level graphics software available on micros, minis and mainframes. And it took our extensive experience in graphics software to do it.

Now all computer-using design engineers and scientists can utilize the industry's finest software. And bring mainframe applications right to the bench. Or vice versa. Which means TEMPLATE's device-intelligence and computer-independence is even further enhanced. And you get the graphics functionality for mechanical and electrical CAD, scientific analysis, seismic work, VLSI, and molecular modeling that puts TEMPLATE in a class by itself. In any environment, whether it's batch or interactive, 2D or 3D.

TEMPLATE features table-driven architecture, 3D software display lists, metafile capability over 250 user-callable FORTRAN routines, workstation model, post processing capability, run-time selection, and complete support functions.

So when you're looking for graphics software that'll run on *all* your computers, call Megatek.

And find out all about TEMPLATE.

The product that just gave micros, minis, and mainframes a new computer graphics image.

MEGATEK CORPORATION
® A UNITED TELECOM COMPANY

See us at AUTOFACT, Booth 515, October 2-4. **Making History out of State-of-the-Art**

World Headquarters • 9645 Scranton Road • San Diego, CA. 92121 • 619/455-5590 • TWX: 910-337-1270

TEMPLATE is a registered trademark of MEGATEK CORPORATION.

CIRCLE 174