

COMPUTER DESIGN

CUSTOM/SEMICUSTOM ICs

ARTIFICIAL INTELLIGENCE
TECHNIQUES EASE VLSI DESIGN

EXPLORATORY DEVELOPMENT
SYSTEMS MANAGE UNCERTAINTY

DISK CONTROL ICs
FIGHT DATA ERRORS



ROCKWELL SEMICONDUCTOR TECHNOLOGY CREATES THE NEXT GENERATION OF INTEGRAL MODEMS.

Our third-generation VLSI capabilities give you support at any level.

Building modems into your product line used to be a complicated and costly process. But Rockwell International's latest generation of VLSI integral modems is changing all that.

Today, our VLSI capabilities mean that we can support your design requirements at any level: you can order VLSI chip sets; standard cards with all data rates interchangeable; or custom cards designed to fit your specific shape, connector and software integration requirements.

In fact, we can offer a simple, economical solution to virtually any application requirement from 1200 to 9600 bps. Since our modems include both a microprocessor bus and an RS232C interface, they can be easily designed into computers, from micros to mainframes; plus G2-G3 facsimile machines, multiplexers, terminals and, of course, box modems.

What's more, Rockwell is the only supplier of full-duplex, dial 2400 bps VLSI integral modems.



And all of our modems are as reliable as they are versatile. Based on our customers' experience with over a half million Rockwell integral modems in use around the world, we've been able to document an MTBF of 200,000 hours—nearly 23 years between service calls. Performance-wise, transmission is virtually error-free over unconditioned or dial-up phone lines because our integral modems offer auto-adaptive equalization.

We are the world's largest supplier of high-speed modem subsystems and devices. Ask your sales representative or distributor for a quote today, and you'll see that we have a distinct advantage in price and delivery lead times, too.

Rockwell integral modems. A generation ahead of all the rest.

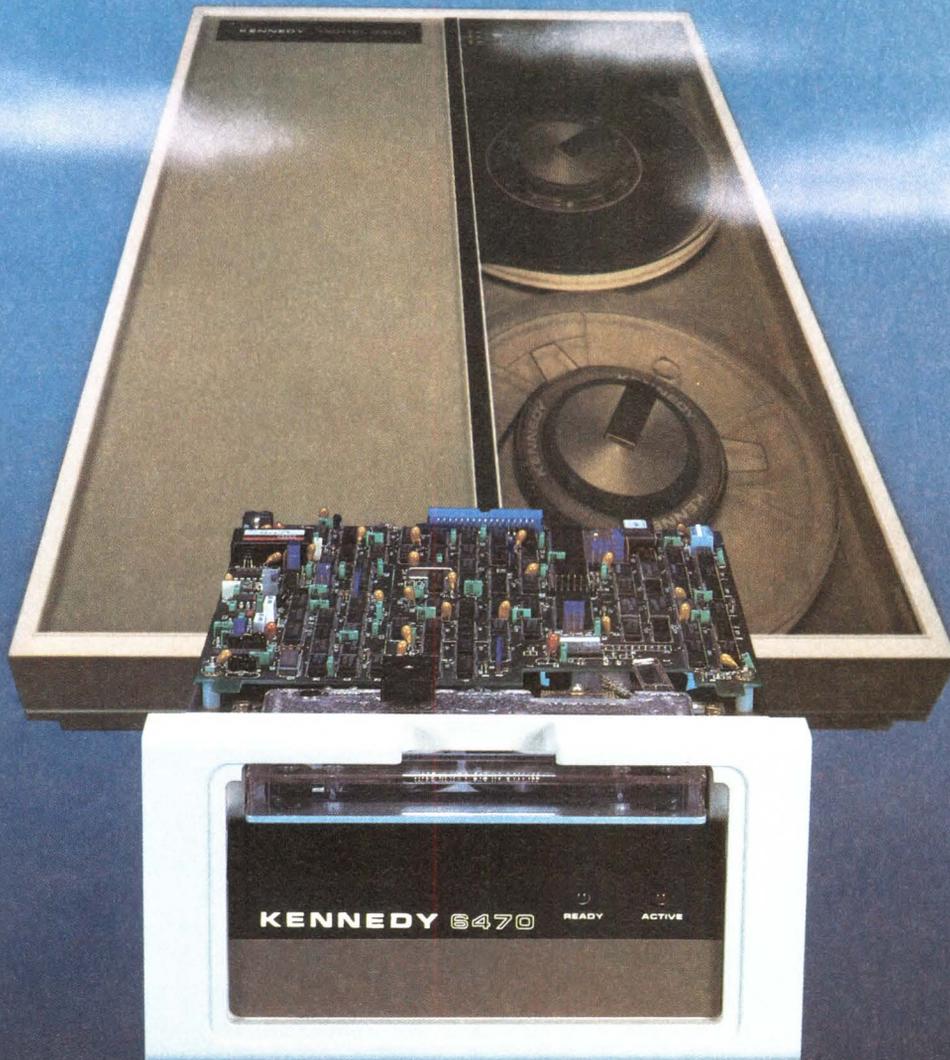
Rockwell Semiconductor Products Division

Rockwell International, P.O. Box C,
MS 501-300, Newport Beach, CA
92660. Call toll free (800) 854-8099.
In California: (800) 422-4230.



Rockwell International

...where science gets down to business



Model 6470 Cartridge Tape System

Thinking Big — really big!

And why not? Kennedy's Model 6470 $\frac{1}{4}$ " cartridge tape drive has every reason to think that it's big — it behaves as if it were a $\frac{1}{2}$ " tape drive.

Examples:

- Model 6470 uses standard $\frac{1}{2}$ " tape interface.
- It has the same functions as a larger $\frac{1}{2}$ " tape drive.
- The cost is 1/3 less than $\frac{1}{2}$ " tape drives.
- It's 1/10 the size of $\frac{1}{2}$ " drives.
- It packs more data per reel than P.E. drives (54 MBYTES).
- Finally, Model 6470 operates on the same tape handlers as $\frac{1}{2}$ " drives.

Model 6470 is a start/stop device with the capability of streaming data at 37.5 IPS, providing much greater flexibility than in comparable drives.

Other features include low power consumption, consuming total power of under 50 watts: diagnostic and fault isolation, with On-Line Self Test and Off-Line Self Test Routines.

Model 6470 — it's not guilty of delusions of grandeur. It really can perform as a $\frac{1}{2}$ " tape drive. Give us a call. We'll prove it.

KENNEDY

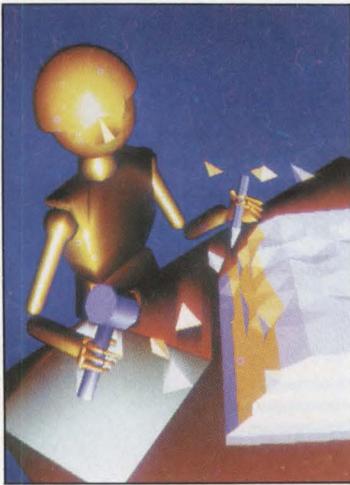
An Allegheny International Company

1600 Shamrock Ave. Monrovia, CA 91016

(818) 357-8831 RCA TELEX 247019 KNDY-UR

KENNEDY • QUALITY • COUNT ON IT

COMPUTER DESIGN[®]



This month's cover is a computer-generated 3-D solid model designed at Coddbarrett Associates, Inc. The data base was constructed and illustrated by Mary Codd.

SPECIAL REPORT ON CUSTOM/SEMICUSTOM ICs

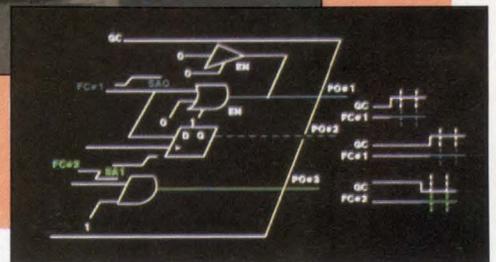
- 89** Application-specific ICs will account for the majority of chips going into production during the next few years. Their rapid acceptance has come about because computer aided engineering and design tools have become more capable, lower in price, and more widely available. But new approaches, such as silicon compilation, will accelerate the growth of ASICs even more by allowing system designers—without the chip-level knowledge of silicon gurus—to turn their ideas into silicon.
- 91** **Design tools pace custom/semicustom IC development**
The penetration of application-specific ICs into systems depends on the availability of design tools that eliminate a designer's need to know the intricacies of silicon technology.
- 105** **Silicon compilation speeds design of complex chips**
Silicon compilation creates a full-custom circuit with the silicon efficiency of a handcrafted design, but with fewer resources and a shorter development cycle.
- 115** **Top-down tool opens new paths to chip design**
Silicon compilation software gives a system engineer with no chip design experience a top-down design tool for the development of CMOS application-specific ICs.
- 125** **Converting gate array designs to standard cells**
With the freedom to easily and inexpensively transform a gate array into a standard cell, designers can delay decision-making and improve production efficiency.

SYSTEM TECHNOLOGY

- 23** **Memory systems:**
Optical disks promise massive storage
- 33** **Integrated circuits:**
AI techniques aim to ease VLSI design
- 39** **Software:**
Unix interface standard on the way
- 59** **Control & automation:**
Speech recognition produces natural interface
- 71** **Test & development:**
Fault simulation becomes design verification tool
- 80** **Software:**
Applications dominate Computer Graphics '85
- 84** **Computers:**
Electro celebrates anniversary with a glimpse at tomorrow



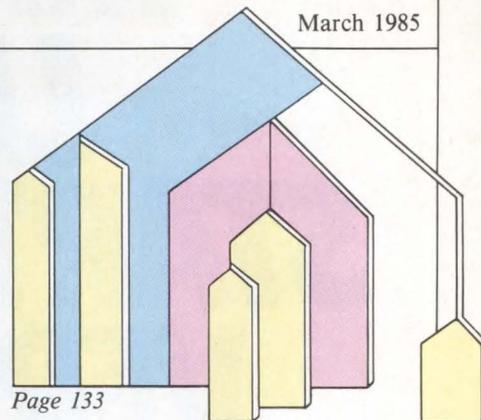
Page 80



Page 71

SYSTEM DESIGN

- 133 Software: Programming the uncertain with exploratory systems**
Artificial intelligence development systems are ideally suited to the coding of complex programs because they can cope with the uncertainty that is often present.
- 147 Integrated circuits: Disk system ICs combat data errors**
Disk drive controller ICs that implement advanced and flexible error correction techniques require little development work by the designer, yet provide powerful methods of coping with hard errors in disk drives.
- 161 Computers: S-100 bus adapted to 68000-based Unix computer**
Although the S-100 bus poses some challenges to system designers, it provides adequate memory support for a multi-user Unix computer and permits multiple DMA devices.
- 173 Computers: New VAX squeezes mainframe power into mini package**
While retaining features common to the 32-bit VAX family, the 8600 uses large-scale processing concepts that include a dedicated memory bus and pipelined operations in both instruction processing and memory references.
- 183 Test & development: Debugger alleviates realtime programming complications**
Realtime systems using microprocessors present special problems in debugging. A debugger specifically tailored to deal with concurrency can provide more reliable code for systems that must deal with the real world.



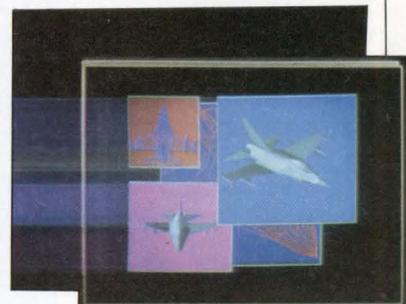
Page 133

DEPARTMENTS

- 7 Up front
- 13 Editorial
- 217 Literature
- 219 Market forecasts
- 220 Calendar
- 221 Designer's bookcase
- 222 System showcase
- 224 Advertisers' index
- 226 Recruitment
- 229 Reader inquiry card
- 229 Change of address card

SYSTEM COMPONENTS

- 193 Integrated circuits:**
Floating point unit executes to 80 bits of precision
- 194 Computers:**
Enhanced array processor peaks at 38 MFLOPS
- 194 Memory systems:**
MiniWinchester supplies top-level performance
- 196 Test & development:**
Geometry engine powers design workstation
- 198 Test & development:**
Graphics controllers tout large display-list memory
- 198 Computers:**
Supermini employs ECL and pipelining for 32-bit computations
- 199 Integrated circuits:**
CMOS chip breathes new life into byte-wide processing



Page 196

MICROVAX I LETS YOU START SMALL IN A VERY BIG WAY.

All supermicros let you start out small. Most run into problems, however, when your applications start getting big.

That's why Digital's MicroVAX I™ supermicro is such an intelligent solution. It limits your risk – not your capabilities.

For starters, MicroVAX I is a member of the VAX™ systems family; it provides a full set of VAX resources. That's saying a lot, because VAX systems are the industry standard for 32-bit computing.

MicroVMS™, our VMS™ operating system for MicroVAX, lets you develop applications for any VAX using the comprehensive set of proven development tools provided by the VAX computing environment.

You get all the VMS editors, debuggers, database managers, screen managers, utilities and communications options. You can use advanced VMS languages – including Pascal, FORTRAN, C, Bliss, COBOL and Ada.® And through the VMS common calling standard, you can combine routines written in different languages within the same program.

This software sophistication lets you cost-effectively dedicate the MicroVAX I system to highly complex tasks. In areas like proc-

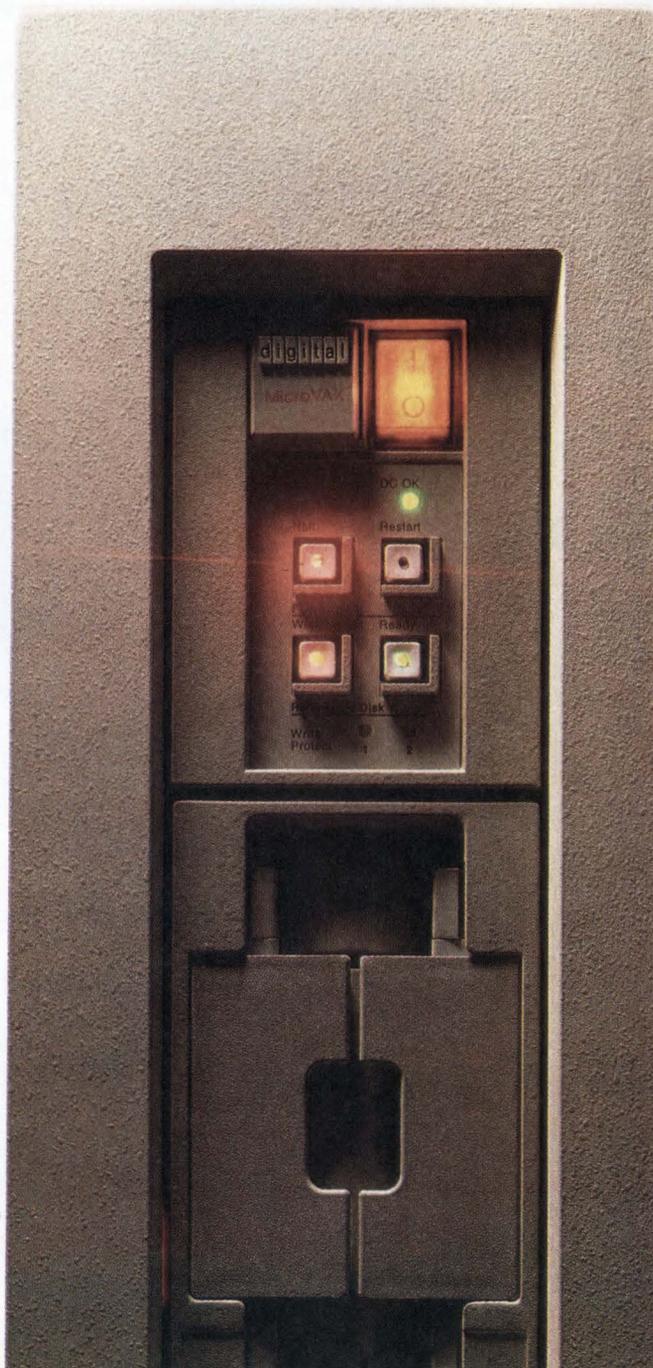
ess control, signal and image processing, CAD/CAM, mathematical simulation, seismic analysis and automatic test. There's no end to the uses for MicroVAX.

**THE BIGGEST
GROWTH PATH IN
THE INDUSTRY
STARTS HERE.**

The sky's the limit for applications you develop on MicroVAX. They can run on every VAX in the family – including our top-of-the-line VAX 8600 and VAXcluster™ systems. Through upward expansion and shared resource networking, the computing power at your disposal is virtually inexhaustible. You can go from a supermicro to a mainframe-level system without changing your software. And that's an advantage no other supermicro supplier can match.

Conversely, you can download VAX applications to MicroVAX I systems. All non-privileged commands and the full VAX instruction set are supported, so your applications will run with little or no modification.

There's even VAXELN™ software, a complete Pascal-based development toolkit, that runs on VMS and MicroVMS. With VAXELN, it's easy to create dedicated realtime applications for MicroVAX and VAX systems. By leaving out the overhead, VAXELN gives you a highly efficient runtime environment.



© Digital Equipment Corporation 1985. Digital, the Digital logo, MicroVAX, VAX, MicroVMS, VMS, VAXcluster, PDP-11, VAXELN, MicroPDP-11, Q-bus, DECnet and ULTRIX-32m are trademarks of Digital Equipment Corporation. Ada® is a trademark of the U.S. Government Department of Defense. UNIX® is a trademark of AT&T Bell Laboratories.



Starting with MicroVAX I, you can grow all the way to the VAX 8600 system . . .

PART OF A GROWING FAMILY OF SUPERMICROS.

Our 32-bit MicroVAX I systems have a lot in common with our 16-bit MicroPDP-11™ supermicros. They all use the same Q-bus™ data transfer architecture. So they all support the same peripherals.

That's a substantial benefit, because Q-bus is an industry-standard bus, with many options available from Digital and independent vendors. Hundreds of peripheral products are available to make your MicroVAX configurations complete.

Commonality between MicroVAX I and MicroPDP-11 systems is also maintained in file structure, communications options, languages and even many applications. All of which expands your options and provides unlimited opportunities for growth.

DECNET, YOUR UNIVERSAL NETWORKING STRATEGY.

When you're ready to distribute your applications via MicroVAX I systems, DECnet™ software makes it easy. MicroVAX supports DECnet to give you a

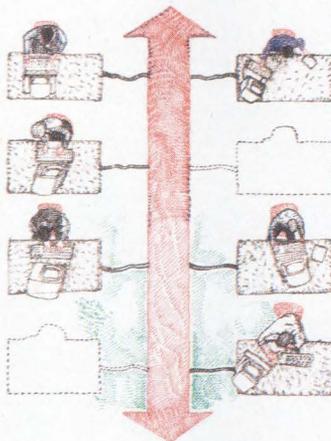
complete set of local and wide area networking facilities.

Your options include local communications via our implementation of the most advanced level of the Ethernet local area network available today. And wide area communications via dial up or leased lines, or public packet-switched X.25 networks. Further, you can establish these local and remote systems links transparently to the user and the program.

Digital's networking facilities also let you take advantage of the computing resources of non-Digital systems. The extensive selection of standard protocols

available to you with the MicroVAX I system includes 2780/3780, 3271, X.25 and SNA. All these facilities are fully implemented not just at the hardware level, but through software as well.

So if you want to start small,



start big with the MicroVAX I system. You'll get the 4-gigabyte address range, the powerful instruction sets, the incomparable development tools and vast software selection that has made our VAX systems the most popular 32-bit computers in the world. And you'll get it all with the economy and convenience of a supermicro package.

BEST ENGINEERED MEANS ENGINEERED TO A PLAN.

Our MicroVAX I supermicro, like all Digital hardware and software products, is engineered to conform to an overall computing strategy. This means 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.

If you believe that your company could benefit from our very growth-oriented MicroVAX I system, contact your nearest Digital sales office or a Digital Authorized Industrial Distributor. Or call: 1-800-DIGITAL, ext. 230.

THE BEST ENGINEERED COMPUTERS IN THE WORLD.

ULTRIX, YOUR TRANSPORTABLE OPERATING SYSTEM.

If you need UNIX® software capabilities, our MicroVAX I system can readily provide them. It supports ULTRIX-32m™ software, our enhanced, fully supported UNIX software for MicroVAX.



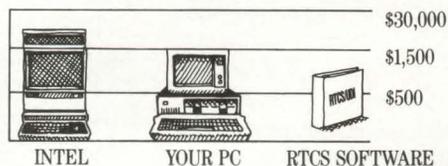
ULTRIX-32m includes all Berkeley 4.2 features, in addition to Digital's enhancements that simplify error reporting and recovery. ULTRIX-32m applications can run throughout the VAX family. Our PDP-11™ and MicroPDP-11 systems families also support a version of ULTRIX software. And that lets you extend a single application across a vast range of systems. You can go from low-cost MicroPDP-11/23 and MicroVAX I systems all the way to our high performance VAX 8600 systems with one, multi-purpose solution. Think how that will extend your reach in your market.

Extend the capabilities of your PC/MS-DOS™ based system and duplicate the Intel™ Series III, IV, and 86/310 with RTCS software development tools.

Don't buy expensive hardware to get the exact same professional development capabilities. For around \$500, the RTCS/UDI Universal Development Interface software lets you use Intel Series III Microcomputer Development System software, available from RTCS, on your IBM PC™ or any other MS-DOS-based computer. And Intel software runs up to 20% faster on your PC using the RTCS/UDI.

DID YOU KNOW YOU'RE ALREADY SITTING ON A \$30,000 INTEL SYSTEM?

Fortune 500 companies like IBM, General Electric, Honeywell, and over 1,500 others have found the RTCS/UDI easier to use with higher performance than most Intel systems. Besides the obvious expenditure comparison, they see Intel Systems as special purpose computers with limited software. But with RTCS, a PC/MS-DOS-based computer becomes a reliable performer with a wide repertoire of software possibilities.



Also available from RTCS is PC/RTX, Real-Time, Multi-user Multi-tasking Operating System for IBM PC and others; Intel languages such as Pascal, Fortran, PL/M™ and C Compilers; Utilities such as high level software debuggers, as well as a wide range of software support supplied with the RTCS/UDI at no additional cost.

Before you go out and buy expensive hardware for capabilities you're already sitting on, call RTCS today. **(805) 987-9781**

RTCS

Getting the most from what you have.

1390 Flynn Road, Camarillo, CA 93010
Telex: 467897 RTCS CI

CIRCLE 4

```

never:
LL dq$write(cu$conn,0('-',1,@status);
LL SETB(0,@kbuffer,size(kbuffer));
ctual = dq$read(cu$conn,@kbuffer,size(kbuffer),@status);
index = 0;
WHILE kbuffer(index) = ' ';
index = index + 1;
0;
index = 0;
parse$incomplete = true;
WHILE parse$incomplete;
char = kbuffer(index);
IF (char = 'x') AND (char < 'z') THEN char = char - 20h;
IF FINDB(@valc,char,size(valc)) < 0 THEN THEN
DO:
command(index+1) = char;
index = index + 1;
END;
ELSE parse$incomplete = false;
IF index = size(command) THEN parse$incomplete = false;
0;
index = 0 THEN GOTO next$input;
(command(1) = 'E') AND (command(2) = 'X') AND (command(3) = 'I') AND (com
in Block Delete Execute Find -find Get --more--

```

RTCS has erroneously used the Intel trademark "iRMX" to describe or name some of its products. All RTCS products that use that trademark now use "iRTX" instead. MS-DOS is a trademark of Microsoft, Inc. PC-DOS and IBM PC are trademarks of IBM, Inc. PL/M and Intel are trademarks of Intel, Inc.

UP FRONT

IBM unveils Sierra...

IBM has announced its "Sierra" 3090 series of advanced mainframe computers amid criticism of the machines' performance and delivery dates. The first model of the 3090 series will not be available until November. That version, the model 200, nearly doubles the capabilities of the 308XX series, IBM's current top-of-the-line mainframe. The model 400, offering twice the performance of the 200, will not be available until the second quarter of 1987. Competitors were cheered by the long lead time of IBM's announcement. "We're confident that our price and performance are still very competitive," said an Amdahl representative. He pointed out that Amdahl's 580 series currently outperforms the best IBM has to offer, but its performance will be about 20 percent below that of the 200. Amdahl's 5880 performs about 24 million instructions per second, while the 308XX series tops out about 16 MIPS. Industry estimates place the IBM models 200 and 400 about 28 and 50 MIPS, respectively. IBM's new models are priced consistently with its previous mainframes, costing about \$200,000 per MIPS. Anticipating shipment of the new model, IBM lowered the price on its 380XX series about five percent. Soon after, Amdahl countered with a similar reduction in the price of its 580 series.—*W.E.S.*

...and switches to Unix System V

IBM surprised many when it switched its support from Unix System III to System V for both its mainframes and personal computers. The recent announcement marked the first time IBM had endorsed a competitor's operating system for its mainframe computers, and indicated a change in strategy for its desktop machines. PC/IX, the XT version of Unix, was modeled after Unix System III, but is being adapted to System V. Microsoft's Xenix, another XT/AT version of Unix, is likewise being modified for System V compatibility.—*W.E.S.*

Bit-slice processors will join standard cell libraries

CMOS bit-slice processors will become standard entries in at least two semicustom IC libraries within the next six months. Both General Electric's Semiconductor Div of Research Triangle Park, NC and RCA Solid State of Somerville, NJ plan to offer 2900 family bit slices in 2-micron CMOS technology as core cells on semicustom chips. These can be enhanced with peripheral circuits to form custom single-chip microcomputers. The CMOS versions of the 2900 series should run at speeds comparable to their bipolar counterparts, but with a mere 10 percent of the bipolars' power requirements. "We chose bit slices instead of 8086-type processors to allow the user maximum area for expansion around the core processor," explains Chuck Gregory, research and development engineering manager of GE's custom IC department. GE's 2901 and 2910 will be silicon in June. RCA's semicustom bit-slice core processor will be based on its GPU 6502 bit-slice family. Scheduled for release this fall, RCA's semicustom chip will comprise up to four subchips, each individually configured for the desired peripheral circuits.—*N.M.*

(continued on page 8)

Program models IC dopant behavior

Researchers now have a way to predict the effects of introducing boron, arsenic, and phosphorus into a two-layered silicon structure. Aptly named Predict, this VAX-based software from the Microelectronics Center of North Carolina (Research Triangle Park, NC) solves coupled equations describing dopant behavior and oxide growth without the need to fine-tune default parameters for each process. The designer need not specify the silicon layer's thickness or the grid spacing. Predict selects the minimum grid spacing that will work with its algorithm. The program produces data on the thickness of the oxide layer resulting from each processing step, plus the sheet resistances of the n and p regions. Included in the printout are the chemical and electrical concentration versus depth for the various impurities. With a grid spacing of 0.01 to 0.03 microns and the maximum number of grid points preset to 200, the useful depth of the silicon layer is between 2 and 6 microns. "The process designer is assured the accuracy of the process results without performing hundreds of simulations or running experimental wafers," says Richard Fair, MCNC's vice president of research program management. MCNC plans to license the program to universities for the cost of the tape, and will soon market it to industry.—*N.M.*

Vendors cautious about Multibus II

Multibus vendors took a wait-and-see stance after Intel's long-awaited introduction of the first Multibus II boards at Mini/Micro West. Already two years behind the introduction of the VMEbus products, 32-bit Multibus II has some catching up to do. The question is whether or not there is still a need for Multibus II. Multibus I suppliers, such as Scientific Micro Systems (Mountain View, Calif) Central Data Corporation (Champlaign, Ill) and Xylogics (Burlington, Mass), are watching the market before plunging in. "Multibus II will be very successful, but it's running late," says Michael Kalashian, marketing director for Scientific Micro Systems. "VME is a bit more real right now." Nevertheless, John Beaston, marketing manager for Intel's Multibus II division in Hillsboro, Ore, remains optimistic. "I think you'll see them move away from their wait-and-see attitude in a couple of months," he predicts.—*R.G.*

AI expert calls for national effort to develop superchips

Artificial intelligence applications will require hundreds of billion-gate superchips by the turn of the century, according to Professor Raj Reddy of Carnegie Mellon University, Pittsburgh, Pa. "Designers must have per-use access to a national facility holding a proprietary data base of hardware and software designs," says Reddy, this year's keynote speaker at New York's International Solid State Circuits Conference. Reddy claims this facility will reduce design duplication so superchips can be designed faster. Furthermore, says Reddy, companies and universities should pool their resources to meet the large computational requirements of AI.—*N.M.*

(continued on page 10)

Dataram does more with memory

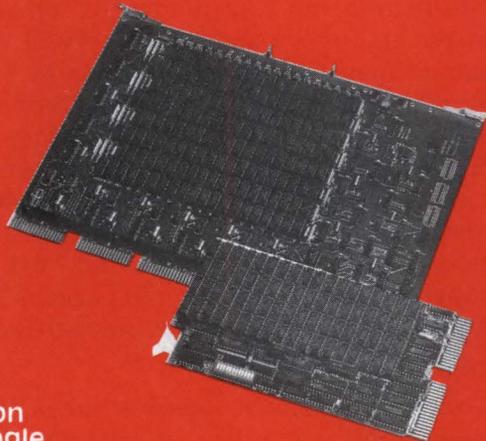
For nearly two decades, Dataram has been leading the way in memory technology, supplying more memory than any other company to leading minicomputer manufacturers, end users, systems houses, OEMs and others throughout the world.

Experience that translates into superior design and reliable performance for you today. That means more for you — more capacity, more capability, more reliability, and more performance with more savings...so you can get more from your minicomputer.

Want to learn more? Contact Dataram at (609) 799-0071 today.

for DEC users

- VAX 11/780 Series Compatible
- 4.0 MB with ECC
- Socketed DRAMs
- Spare DRAM on Board
- Low Cost
- System Expansion to 64 MB in a Single Cabinet



- Q-Bus Compatible
- 2.0 MB
- 85 nsec Access Time
- Block Mode DMA
- Dual Board
- On-Board CSR
- Full-Function Parity

Dataram satisfies all your VAX, Q-Bus, and UNIBUS memory requirements.

DATARAM

Dataram Corporation □ Princeton Road □ Cranbury, New Jersey 08512 □ (609) 799-0071 □ TWX: 510-685-2542

DEC, VAX, UNIBUS, and Q-Bus are trademarks of Digital Equipment Corporation.

CIRCLE 5

Macro LSI cells integrate simple standard cells

Harris Semiconductor is shifting its semicustom approach to combine its standard cells into macro LSI cells, offering these larger chips as part of the standard cell library. Harris's Melbourne, Fla design center will offer the designer a line of 80C86-family peripheral parts as building blocks for integrating several functions on a single chip. "The current required to drive external capacitance loads drops significantly in a semicustom design because most of the I/O interfacing is done internally," says J.W. Scherer, marketing director for the semicustom product line. A typical discrete configuration of an interval timer, parallel I/O, FIFO, UART, control logic, priority interrupt controller, and oscillator requires 202 pins. The semicustom equivalent requires 84. Three of the currently available blocks exceed 2000 gates. Typical circuit densities reach the equivalent of 7000 two-input NAND gates. Harris is using the same scaled SAJI IV CMOS process it uses for its standard RAM, 80C86, 80C88, and peripheral chips to produce the semicustom LSI cells. This gives the chips a minimum feature size of 2 microns. The designer can easily mix and match the macro LSI parts with simpler cells to create full custom systems.—*N.M.*

Honeywell challenges lead in factory automation

Manufacturing automation will be a major factor in Honeywell's future, according to chairman and CEO Edson Spencer. Spencer claims the company expects to be a leading supplier of integrated information and control systems for manufacturing automation by the end of the decade. Honeywell will target segments of a market that will reach \$50 billion by 1990 and near \$100 billion by 1995. The company's current products include a plant management system and a supervisory work center. A system for distributed manufacturing control should be ready late in the year. Eventually, Honeywell plans to introduce a universal controller that will interface programmable controllers and process controllers. The plant network will be Ethernet-based initially, but will support the General Motors Manufacturing Automation Protocol when standards are better defined.—*S.F.S.*

Semiconductor market growth to continue

Contrary to many economists' predictions of a recession in the semiconductor industry, at least one major market research firm believes 1984's upward trend will continue into 1985. Integrated Circuit Engineering of Scottsdale, Ariz anticipates 11 percent worldwide growth in semiconductor trade. That would be an increase from \$32 billion in 1984 to \$35.63 billion in 1985. ICE's long-term forecast includes a recessionary period from 1986 to 1987, but that would not keep the industry from reaching \$90 billion by 1990, ICE believes.—*N.M.*

COMPUTER DESIGN © 1985 (ISSN-0010-4566) is published monthly by PennWell Publishing Company, Advanced Technology Group, 119 Russell St., 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 \$60 in U.S.A. and \$85 elsewhere. Single-copy price is \$6.00 in U.S.A. and \$8.50 elsewhere. Microfilm copies of COMPUTER DESIGN are available and may be purchased from University Microfilms, a Xerox Company, 300 North Zeeb Rd., Ann Arbor, MI 48106. Officers of PennWell Publishing Company, 1421 S. Sheridan, Tulsa, OK 74101; P. C. Lauinger, Chairman; Philip C. Lauinger, Jr., President; Joseph A. Wolking, Executive Vice President; H. Mason Fackert, Carl J. Lawrence, Group Vice Presidents; L. John Ford, Vice President; V. John Maney, Vice President/Finance.

® COMPUTER DESIGN is a registered trademark. All rights reserved. No materials may be reprinted without permission. Phone (617) 486-9501.  

HIGHER INTELLIGENCE



Could your
Discover the
develop-

ZAX intro-
development

duces the first, *true 32-bit universal microprocessor system* featuring ZAX ICD-series in-circuit emulators and DIGITAL'S MicroVAX I™ super-microcomputer. The ZAX/VAX connection delivers the highest level of power and flexibility by offering professional development systems and software integration products which *support both multi-processor and multi-user capabilities*, while maintaining full compatibility with larger VAX/VMS computer systems.

To find out more about how to THRUST your design projects to new heights, call the leader in microprocessor development systems —

ZAX Corporation • 800/421-0982 • 714/474-1170

ZAX

Zax Corporation

2572 White Road, Irvine, California 92714
(714) 474-1170 • 800-421-0982 • TLX 183829

CIRCLE 6

Gould... Innovation and Quality in Real-time Systems

THE FIREBREATHERS.

Leaders since the beginning of Real-time.

Legends have a way of being based on fact. Gould Firebreathers are legendary in the land of Real-time and our leadership and superiority are no myth.

Gould ignited the Real-time fire two decades ago. Since then, scientific and industrial computer users have been hot on Gould. In the simulation,

energy, oil and gas, laboratory and computational markets, Gould Firebreathers are the choice of virtually all the leading companies that implement Real-time systems.

Gould had the first Real-time minicomputer in 1965. We again brought sophisticated Real-time users out of the Dark Ages with the first 32-bit Real-time mini in 1969 and the first supermini in 1975. And we're still the leader in superminicomputer performance. When it works and it's proven, history does repeat itself.

Be a part of history in the making. Bring Gould into your Real-time realm. Firebreathing power from the beginning, for today...and tomorrow.

**Gould Inc.,
Computer Systems Division,**
6901 West Sunrise Boulevard,
Fort Lauderdale, FL 33313
1-800-327-9716.



GOULD
Electronics

TWO NEAL AWARDS FOR OUR EDITORS

According to a popular but scientifically invalid cliché, lightning never strikes twice in the same place. Similarly, people in the business trade press have long held the view that winning two Jesse H. Neal Awards for editorial achievement in the same year is virtually impossible. These prestigious awards, presented by the American Business Press association, are often characterized as the "Pulitzer Prize of the business press." I am pleased to be able to announce, therefore, that the editors of *Computer Design* have just defied the odds and captured two Neal Awards for articles published last year.

To put this achievement into perspective, I should point out that this year marks the first time *Computer Design* has won a Neal Award. Also, as far as we can determine, none of the magazines against which we compete has ever won two first place Neal Awards in a single year—although some of them have won pairs of Certificates of Merit, or runner-up awards. Our awards were for best series of analysis articles (from our September 1984 issue on "Future Computers: System Design Beyond 1984") and best staff-written editorial (for "The Computer Illiteracy Threat," published in February 1984). The editors honored for the analysis articles were Nicolas Mokhoff (Issue Editor), Harvey Hindin (Special Features Editor), John Bond (Senior Editor), Michael Elphick (Editor in Chief), Sydney Shapiro (Managing Editor), and Leslie Ann Wheeler (Assistant Managing Editor). In addition, Elphick and Wheeler received awards for the signed editorial. The editors received their awards at a luncheon in New York City on February 28.

Of course, while winning two Neal Awards in the same year is a long-shot bet, it is by no means the random event implied by the formidable odds. It has more to do with skill than luck. Or, as professional golfer Lee Trevino once said, "Sure, it's a game of luck, but I find the more I practice the luckier I get."

Certainly, the editors of *Computer Design* have worked hard at their game for the past few years. Based on solid research of both the technology they cover and of your specific information needs, they have steadily improved the magazine's content and presentation. The improvements have been measured not only in readership studies

but also in less prestigious but nonetheless important technical writing competitions. For example, the competition organized by the Boston Chapter of the Society for Technical Communication has allowed us to measure our progress each year in terms of the number of awards received, whereas if we had merely looked at our failure to win Neal Awards in earlier years, we might have become discouraged. In the local STC competition, we steadily won more and more awards each year—this year we had nine award-winning entries, including three awards of Excellence (first prize) and three awards of Distinction (second prize).

Credit for our achievements has to go to many more people beyond the award-winning editors. The management of PennWell Publishing, both at our parent company in Tulsa, Oklahoma, and at our Advanced Technology Group in Littleton, Massachusetts, has demonstrated a continuing commitment to editorial excellence and has never asked us to sacrifice quality for short-term profit. Other editors who did not win Neal Awards nevertheless made equal contributions. Immediately after the award-winning September issue last year, we published two issues in October. So some editors, such as Peg Killmon in Littleton, and Tom Williams in Sunnyvale, California, were asked to concentrate their efforts on subsequent issues. And, of course, a supporting cast of associate editors, field editors, copy editors,

artists, and production people rounded out our efforts on each issue.

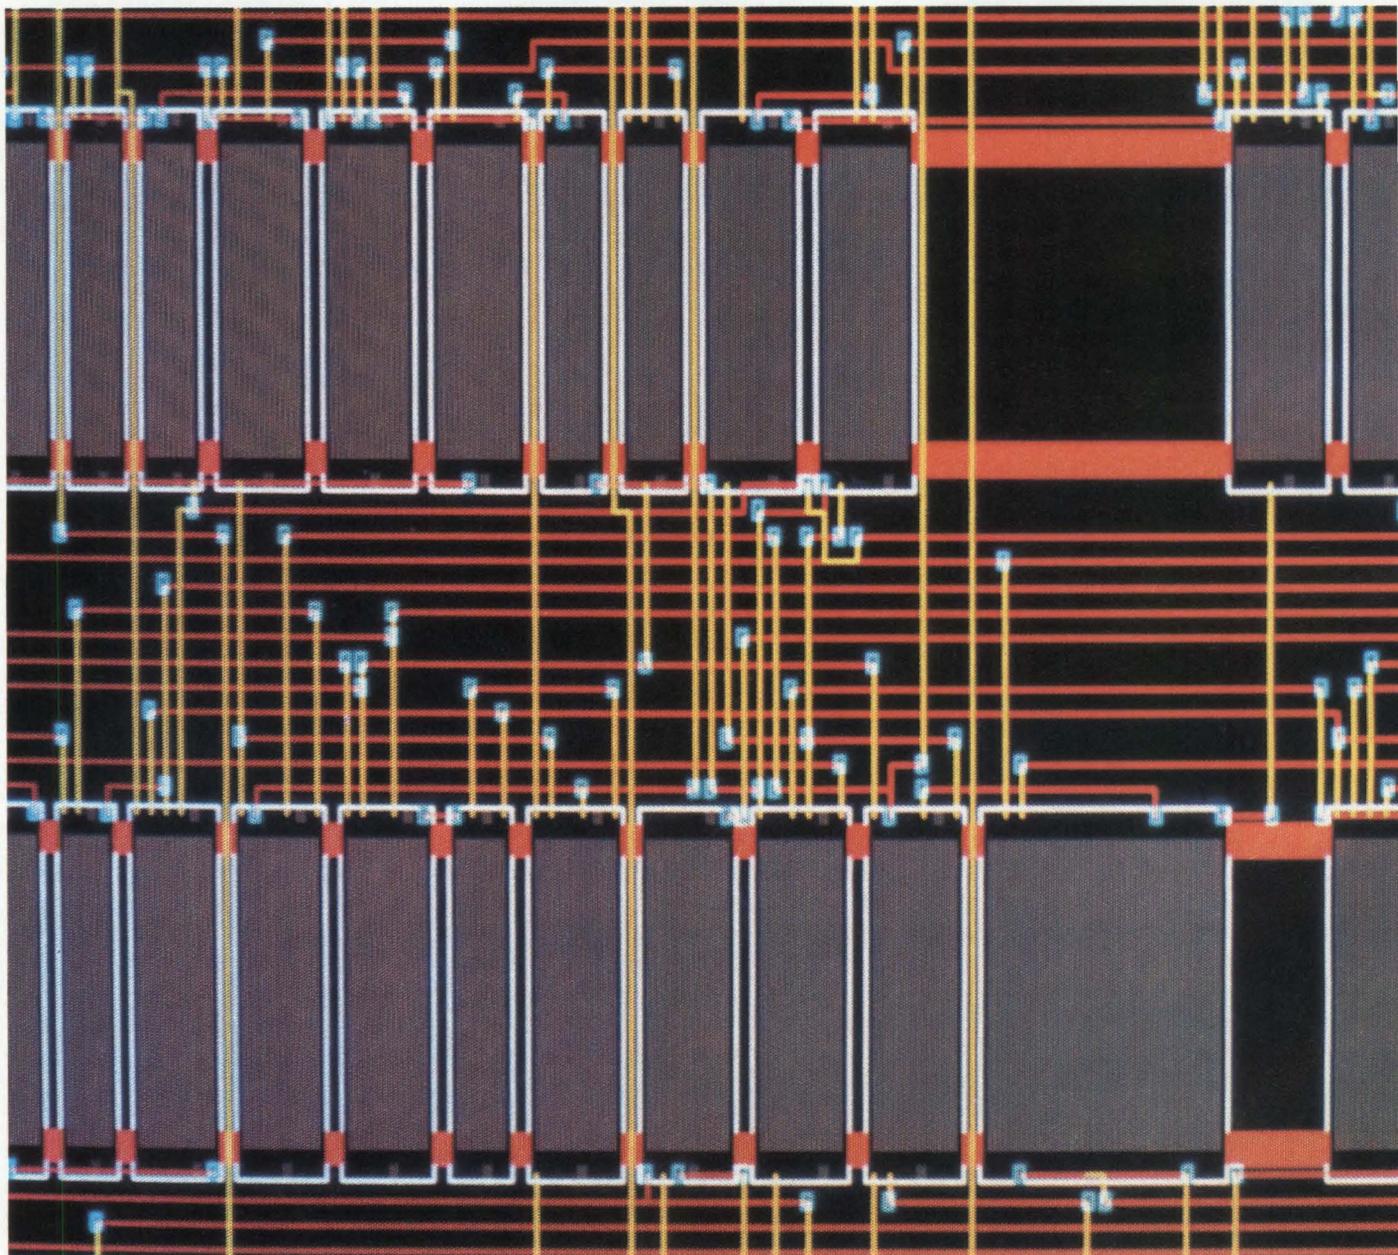
Finally, we owe a lot to you, our readers. Without you, this magazine could not exist. More importantly, however, your escalating needs for vital information have pushed us to increasingly higher levels of performance. One major criterion by which Neal Award entries are judged is the service to the reader and to the field covered by the magazine. The judges thought we served you well. We hope you agree.



A handwritten signature in black ink, which reads "Michael Elphick". The signature is fluid and cursive.

Michael S. Elphick
Editor in Chief

**In semicustom
CAE, only Mentor
Graphics knows how
to get around the block.**



If you're responsible for logic design, chances are that silicon implementation is not near and dear to your heart.

But *somebody's* got to do it, right? Wrong.

If you're working with standard cells or gate arrays, there is now a CAE system that can do it for you.

The IDEA Series™ from Mentor Graphics.

The impact on your productivity? Just short of staggering. Months of layout work compressed into hours. 98 to 100

per cent of the layout task completed without human intervention.

At the heart of this landmark achievement are automatic placement and automatic routing algorithms unique to Mentor Graphics.

No one does it better. Or faster.

Imagine. A 2400-gate design placed and routed in just 3 hours, 15 minutes. Completely and automatically. With 95% gate utilization.

Even after the routing and placement puzzle is solved, Mentor Graphics superior intelligence stays with you. Every move you make during cleanup work is monitored by the system to ensure you stick to the layout rules. You're even presented with legal alternatives to head off mistakes before you make them.

And to top it off, the wire delay values produced during the layout can be back-annotated into the original logic design for super-accurate simulations.

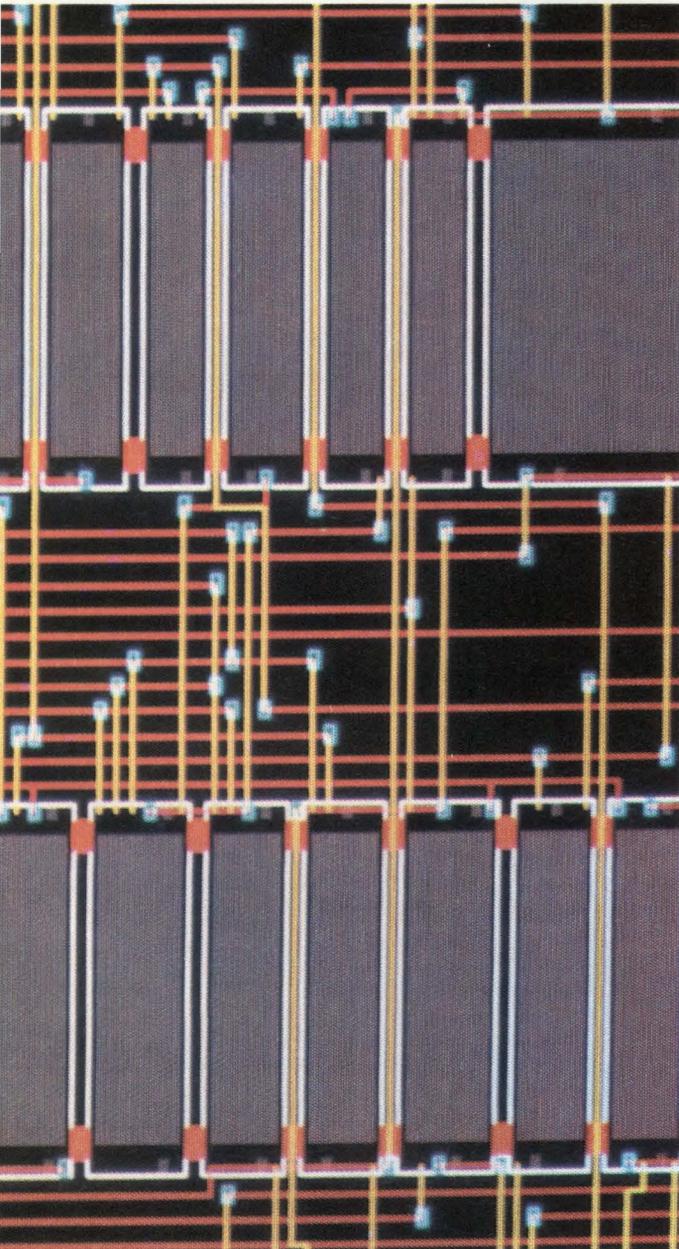
Consider the big picture, too. With the IDEA Series, you have a complete set of integrated engineering tools that come into play long before layout. A fully integrated path all the way from schematic entry to simulation to layout to prototype verification.

You design faster, better and less expensively.

It's all part of a vision unique to Mentor Graphics, the technological leader in Computer-Aided Engineering. Let us show you where this vision can take you.

Call us toll-free for an overview brochure and the number of your nearest sales office.

Phone **800-547-7390** (In Oregon call 284-7357).



**Mentor
Graphics™**

GVC-16 COMPLETE UNIX SUPPORT ON A SINGLE MULTIBUS BOARD

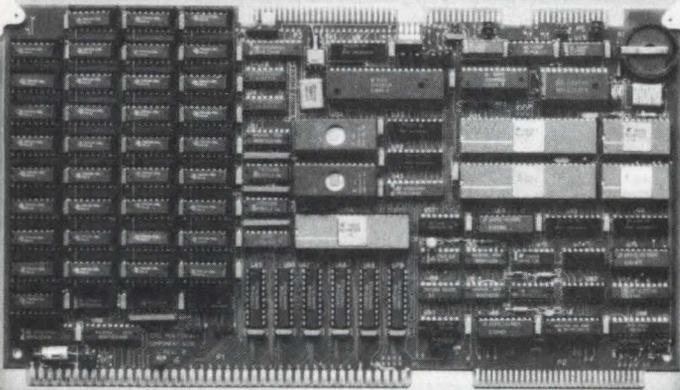
With the GVC-16 all the resources needed for a complete multi-user virtual memory UNIX system are contained on a single Multibus board. Perfect as the heart of a workstation product, the GVC-16 includes:

- National Semiconductor 32016 CPU with full 32-bit internal operation.
- 512kb—2 Megabyte on board no—wait state Dual Port RAM.
- Available Floating Point Unit— 32 bit floating point multiply in less than 5 μ sec.
- Winchester Disk Interface.
- Four serial (RS-232) Ports.
- Berkeley 4.1 UNIX available.

Priced at \$1975. (OEM qty.), the GVC-16 is available from stock.

Also available:

GVC-16 system—a complete 4 user UNIX box with 1 Megabyte of memory, 40 Megabyte disk storage, 4 slot Multibus back plane and UNIX software. \$9,990 (qty. 1).



GVC Inc. 222 Third Street Cambridge, MA 02142 (617) 576-1804

UNIX is a trademark of Bell Telephone Laboratories, Inc.

Multibus is a trademark of Intel Corp.

COMPUTER DESIGN

Editorial/Executive Offices:

The PennWell Building,
Littleton, MA 01460, Tel: (617) 486-9501

Editor in Chief, Michael S. Elphick

Executive Editor, John Miklosz

Managing Editor, Sydney F. Shapiro

Assistant Managing Editor,
Leslie Ann Wheeler

Art Director, Emilie Whitcomb

Senior Editor, John Bond

Senior Editor, Peg Killmon

Senior Editor, William E. Suydam, Jr.

Senior Associate Editor,
Malinda E. Banash

Associate Editor, Jack Vaughan

Copy Editors,

Helen McElwee, Nancy E. Purcell,

Leah A. Rappaport, Ann Marie Tisevich

Editorial Assistants, Vickie Frechette,
Susan M. Wheatley

Special Features Editor, James W. Hughes

Consulting Editor, Saul B. Dinman

New York Field Office:

120 Mineola Blvd, Suite 510,

Mineola, NY 11501, Tel: (516) 741-4780

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 L. Weatherwax

Production Director, Linda M. Wright

Production Manager, Philip Korn

Production Assistant, Elizabeth Mogren

Staff Artist, Jay Gordon

Technical Art, Designline

Ad Traffic Coordinator, Debra L. Stone

Printing Services, Padraic Wagoner

National Sales Manager,

Frederick A. Stanley

Circulation Director,

Robert P. Dromgoole

Business Development Manager,

Leslie Ringe

Promotion Director, Steve Fedor

Marketing Services Coordinator,

Linda J. Battle

Publisher, Frederic H. Landmann

PennWell

Advanced Technology Group

119 Russell St, Littleton, MA 01460

Tel: (617) 486-9501

H. Mason Fackert, *Group Vice President*

Daniel R. Campbell, *Vice Pres. Admin.*

Kirk S. Campbell, *Finance Director*

John M. Abernathy, *MIS/DP Director*

Patricia M. Armstrong,

Admin. Services Mgr.

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.



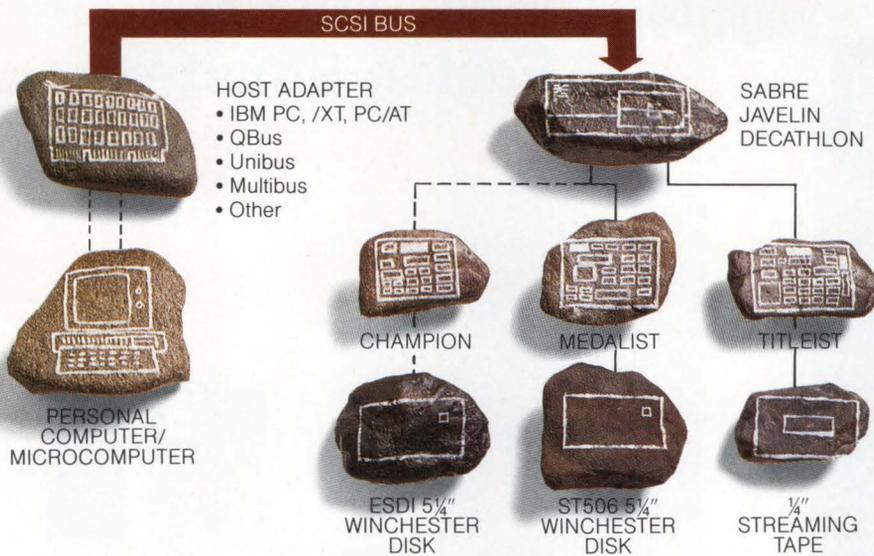
inmos and occam are trademarks of the Inmos Group of Companies

occam

**EMULEX HAS
THE CORNERSTONE
FOR SYSTEM BUILDING
FLEXIBILITY.**



AND, WE'VE GOT ALL THE OTHER BUILDING BLOCKS, TOO.



THE SCSI POTENTIAL. THERE'S PLENTY OF IT.

SCSI allows you system building ease and maximum flexibility through the use of a common I/O bus and standard peripheral interface. All Emulex/SCSI-based products feature these enhancement benefits.

FEATURE	BENEFIT
Arbitration	Permits concurrent I/O operations between different devices by prioritizing bus usage.
Disconnect/Reselect	Maximizes bus utilization by providing devices an orderly scheme to release and regain bus control.
Copy	Permits data transfer between devices without host intervention.
Search Data	Used to search one or more logical blocks for less than, greater than or equal to a specified data pattern.
Reserve/Release	Device can be reserved for exclusive use by a single host in multiple host configuration.
High Level Command Set	Provides host with device independence for a given type of device.

GET MORE INTO YOUR SYSTEMS WITH EMULEX/SCSI CONTROLLERS.

Our Medalist and Champion SCSI controllers can interface any ST506 or ESDI 5 1/4" Winchester disk drive. And our Titleist can interface most 1/4" streaming tape drives.

EMULEX CAN FIT SCSI INTO YOUR PLANS.

If you're looking for a solid cornerstone and the widest range of building blocks to put together powerful and efficient mini/micro subsystems, then talk to Emulex.

For detailed information on Emulex/SCSI products, call toll free (800) 854-7112. In California (714) 662-5600. Or write: Emulex Corporation, 3545 Harbor Blvd., P.O. Box 6725, Costa Mesa, CA 92626.

With Emulex's SCSI products as the cornerstone for your mini/micro computer systems, you can continue to build on your investment. Our SCSI-based products offer you convenient building blocks with which to assemble, update, mix or interchange different types of peripheral storage devices to your computer system. Couple that with the industry's widest range of advanced products and you have the single best source for your system building needs. Products with performance, quality and reliability that's unsurpassed.

PUT MORE INTO YOUR SYSTEMS WITH EMULEX/SCSI MICRO SUBSYSTEMS.

Compact, versatile Emulex/SCSI micro subsystems provide you a variety of cost-effective storage and backup devices. Choose the enclosure you desire—desktop, rack or tower—disk size—36, 67 or 110 MB formatted capacity—and the type of backup most meaningful to your application. Top that off with the ability to interface virtually any CPU type and you have a powerful and complete storage solution.

SABRE, combines on-line storage and file backup capability in a compact unit housing a 5 1/4" Winchester disk drive and an 8" removable cartridge disk.

Javelin, an extremely versatile storage system designed to satisfy

a wide range of OEM and end user applications. Space accommodations for two SCSI peripherals, either one 5 1/4" disk drive and one 60 MB 1/4" cartridge tape streamer, or two disks.

Decathlon, available in a rack mount or an attractive tower/pedestal model, contains up to three 5 1/4" peripherals, with a choice of Winchester disks and a 60 MB 1/4" streaming tape drive.

EMULEX'S HOST ADAPTERS CAN FIT SCSI INTO YOUR PLANS.

Emulex/SCSI host adapters are available for QBus and Unibus systems, IEEE 796 Multibus systems, the IBM PC/XT and PC/AT and compatibles. We've designed these host adapters to take full advantage of the versatility and high performance of the SCSI.



Yes! Please send me your new **Micro Products Handbook**.

My needs are: Immediate 60-120 days Info. only

Have Salesman call Quantity per year: _____

Specify Computer/Peripherals in use or planned:

DEC _____ Multibus _____

IBM _____ Other _____

Name _____

Title _____ Phone _____

Company _____

Address _____

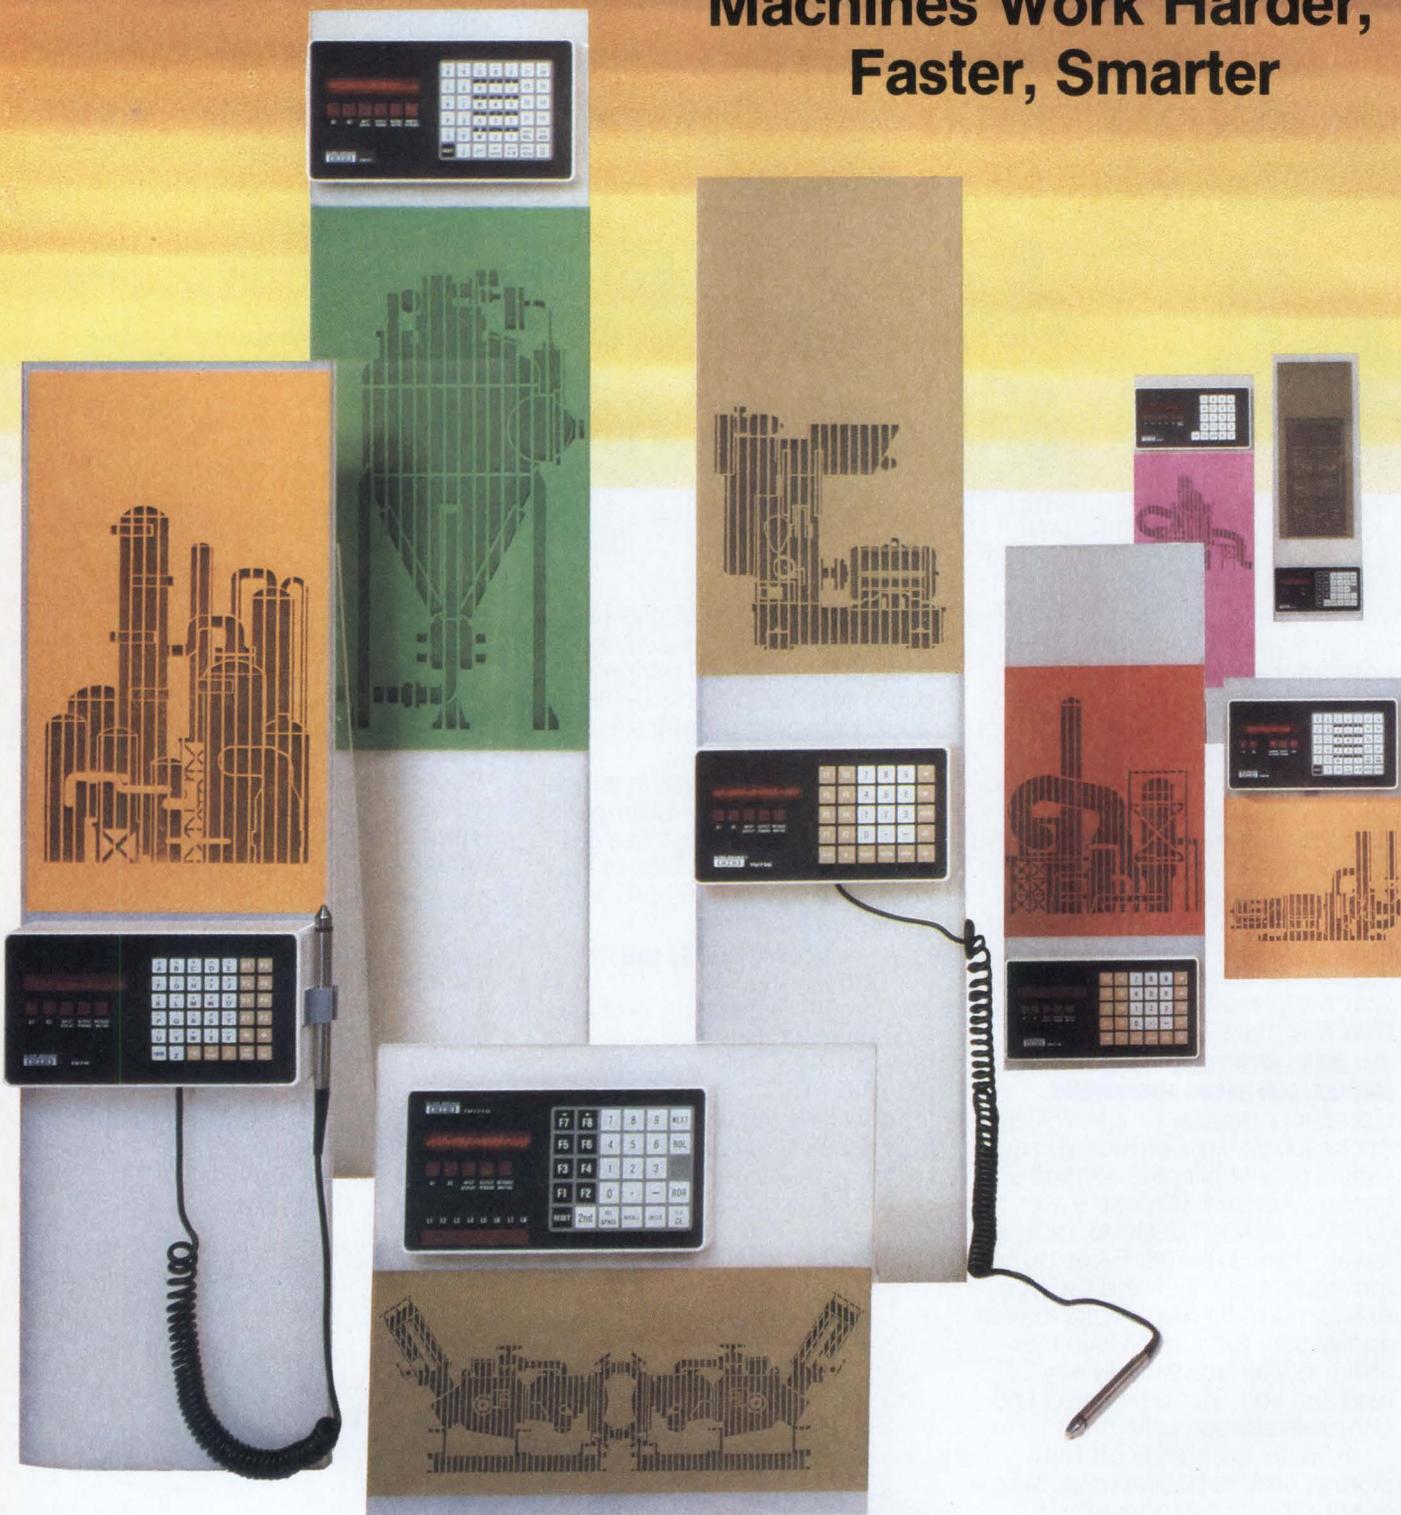
City _____ State _____ Zip _____

CD

IBM PC, XT, AT are registered trademarks of International Business Machines Corporation. QBus and Unibus are registered trademarks of Digital Equipment Corporation.

The Microterminal

14 Rugged, Ready-Made Data and Control Terminals Help Systems and Machines Work Harder, Faster, Smarter



Work Force

These compact, reliable Microterminals get the job done in demanding operator interface applications—from unit controllers for machines and medical equipment to factory data collection and process control terminals. Microterminals are host-and-operator friendly, easy to design in, and ready off-the-shelf to meet a wide range of data entry, display, and control requirements. Choose from 14 models with these outstanding features:

- Up to 40-character alpha and numeric displays, 80-character I/O buffers
- Baud rates to 19,200
- Large, easy-to-use keypads
- Dust- and water-resistant front panels
- Mount on any flat surface
- RS-232C, RS-422, and 20mA current loop interface options
- Drop up to 63 Microterminals on one serial interface
- Up to 16 programmable function keys
- Bar code and magnetic stripe readers
- Polled and nonpolled operation

Put your machines and systems to work with Microterminals. Call or write us today for complete specifications on all 14 models.

Data Acquisition and Control
Systems Division
P.O. Box 11400
Tucson, AZ 85734
(602) 746-1111

BURR-BROWN[®]
BB

**Putting Technology
To Work For You**

CIRCLE 11

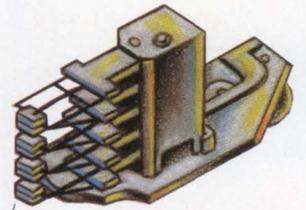
A CRASH COURSE IN DISK AND DRIVE TESTING:

Disk Testing

ADE RVA instruments will show you how to test excessive acceleration, flatness, radial waviness, datum positioning, axial run-out and thickness.

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.



Head/Assembly Testing

ADE RVA instruments give you advanced instruction on head positioning accuracy, head motion studies, dynamic flight characteristics, pitching and rolling.

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

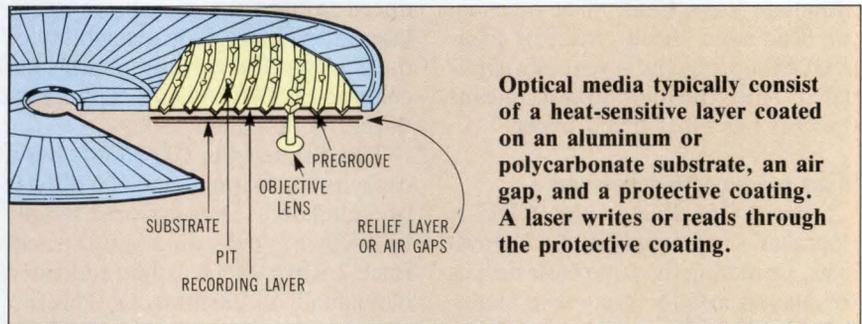


Optical disks promise massive storage

Compact optical disk drives are eliminating offline data and program storage constraints. By packing gigabytes of data on easily removable media, optical drives have a significant impact as data storage peripherals.

Although optical disks supply far greater capacity than possible with current magnetic disk drives, they are slower. Nevertheless, optical disk drives are still faster than secondary storage units based on tape technology. The relative speed and capacity of optical drives, along with the technology's current write-once limitations, place these storage units in a new position as a tertiary online storage device.

Two of the several varieties of optical disk technologies stand out as online storage devices. Optical ROM (OROM) or compact disk (CD) ROM drives have read-only capability. Data or programs must be prerecorded during the media manufacturing process. Write-once, read-mostly



Optical media typically consist of a heat-sensitive layer coated on an aluminum or polycarbonate substrate, an air gap, and a protective coating. A laser writes or reads through the protective coating.

optical disk drives extend their utility by providing the facility to record data while online to the computer.

Digital optical data disks measuring about 5 in. in diameter derive from digital audio disks. Today, they are breaking open new applications as a means for software and database distribution that exercise the vast capacities of optical technology. Optical devices using 12- and 8-in. diameter disks show their stuff in archival settings, as well as in data processing applications. In data processing, per-

ceived disadvantages—inability to erase the disk and slow access—are countered by their capacity for massive amounts of information and the ability to remove the media from the drive.

OROMs make their mark

Immediate applications of OROM drives center on their extensive storage capacity, according to Gerald Finn, vice president and general manager of Sony's Component Products Div in Paramus, NJ. "The CD ROM system gives software houses the capability to put dozens of complex programs with their instruction manuals on a single disk. This allows users to learn programs interactively without referring to a book."

He adds that CD ROMs provide a cost-effective alternative to the recurring communications charges incurred in accessing large central data bases. This is possible because CD ROMs permit entire data bases to be sent to the field. Moreover, the disk takes up less space per megabyte of capacity than any conventional memory device and can be adopted to all computer systems, says John Messerschmitt, vice president of North American Philips Corp in New York.

Although use of CD ROMs as a low cost mass data distribution system promises to be a high growth area, Finn says CD ROMs are not intended to supplant other types of storage. Messerschmitt reinforces this: "CD ROMs offer no recording capability and are not a replacement for existing memory forms. They simply fulfill different requirements."

Optical disks in general offer multiple advantages over conventional

Inside an optical drive

Optical disk drives share many functionally similar parts with magnetic disk drives. The primary differences lie in the optical head and the recording medium. An optical head consists, basically, of semiconductor laser diode, beam collection lens, polarizing beam splitter, quarter-wave plate, routing mirror, two-axis head, and detector with associated beam forming optics.

The laser emits a coherent beam of light that diverges as it leaves the laser. This light is gathered by the collection lens and made into a collimated (parallel) beam of light. This beam enters the beam splitter and goes through the quarter-wave plate where its polarization is rotated 90 degrees. It then passes to the objective lens in the two-axis actuator system.

The actuator system moves this lens vertically under control of the focus servo to direct the beam on the recording layer inside the optical disk. Under tracking servo control, radial movement centers the beam on the pregrooved track where it makes a mark or reads a bit of data.

Light that is reflected from the disk's surface when the beam hits it is collected and collimated by the objective lens. It then travels through the quarter-wave plate and is again rotated 90 degrees. Since it is now 180 degrees out of phase, it reflects from the beam splitter instead of passing through. This routes the beam onto a quad photodetector. The signal produced by the photodetector is amplified to produce both a read channel data signal and feedback for focus and tracking servos.

Since a written mark on the surface of the disk will reflect less light than an unwritten spot, the difference in reflectivity can be detected as an amplitude modulation in the signal that is passed to the read channel. When the read channel data signal is processed, it provides the data and clock at the drive interface.

magnetic media, according to Dr David H. Davies, project manager of 3M's Optical Recording Project in Mountain View, Calif. "The potential for head wear, head crash, or magnetic field problems is virtually eliminated. And, the high capacity means the cost per bit is much lower."

Mass produced mechanisms

Sony's CD ROM, the CDU-1, "parallels the design of audio CD systems, permitting the same basic devices for players and the same disk manufacturing techniques to be used," says Finn. These systems use mechanisms and media similar to those developed for consumer digital audio disk units. With the adoption of stringent error correction techniques, they supply the low bit error rates demanded by com-

puter OEMs. Large quantities of lasers, tracking mechanisms, and servo mechanisms have been produced to meet the acceptance of the laser audio disks. The availability of these proven low cost components has enabled OROM systems to be produced for less than \$500.

The disk used in CD ROMs has a capacity for storing up to 600 Mbytes of digital data. Data is stored serially in 2-Kbyte blocks on a spiral track. Each 2-Kbyte block can be addressed individually in 100 ms to 1 s. This format, proposed as a standard by Sony and NA Philips, has been adopted for these mass-produced read-only disks. Using this format, drives provide a serial data rate of 1.41 Mbits/s, the same as the real disk data rate. Data buffering and error

correction are performed in the host computer controller.

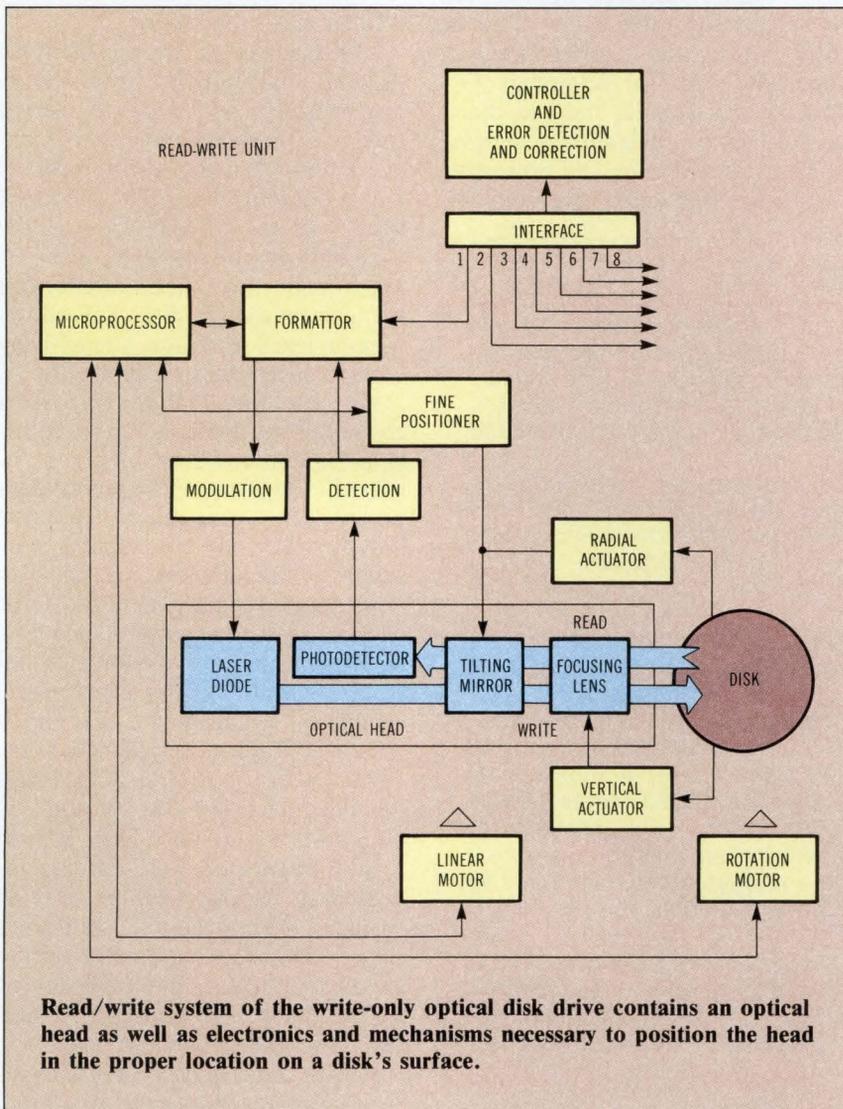
CD ROMs using the format are structured with a lead-in area, a program area, and a lead-out area. Program areas can contain up to 99 information tracks, each divided into as many as 99 blocks. Data fields within each data block consist of a 12-byte sync field, 4-byte header, 2048 bytes of data, and 288 bytes of auxiliary data. The ability to access data randomly derives from information embedded in the header field that is present at the beginning of each track. The header contains 3-byte sector address and mode byte.

Depending on the mode value, the auxiliary field may contain either additional error detection and correction data, or may be available to the user. Use of extended error correction and detection codes can bring the error rate of the disk to 10^{-17} . The disk itself has a bit error rate of 10^{-5} , and use of Cross Interleaved Reed-Solomon Code on data read from the disk brings this to 10^{-11} .

A little bit closer

The writable, nonerasable drives now coming into full production suit another set of applications—many yet to be developed. These "write-once" drives, appearing with 12-in. (30-cm) or the 8-in. diameter disk favored by Japanese manufacturers, use a semiconductor laser diode that can be modulated. This allows them to read and write. This diode laser allows an optical drive to be packaged in a box that fits in a standard 19-in. rack, just slightly smaller than that required for a 14-in. Winchester drive. Its 12-in. diameter disk can store about 1 Gbyte of data on a single surface. Prices for Japanese, European, and U.S. units, complete with controller, all fall in the \$15,000 range.

Available from Optimem, Alcatel Thomson Gigadisc, and Optical Storage International as well as from Hitachi, Panasonic, Toshiba, and NEC, write-once, read-mostly drives are distinguished from read only or compact disk units by the ability to write data directly on the disk within the drive. This eliminates the mastering process and allows data to be handled in a manner that more



closely resembles traditional magnetic methods. Attention must be paid to the problem of how best to handle a storage unit that does not provide the ability to modify data once written.

Appearances may be deceiving

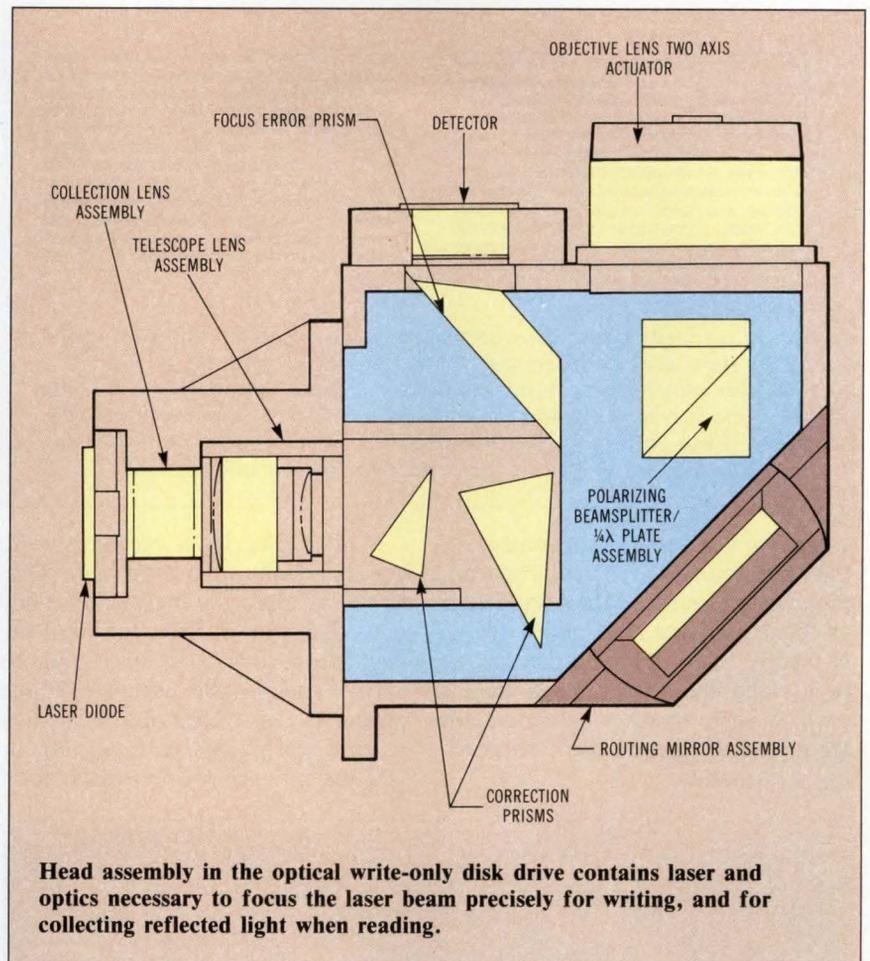
Although write-once drives using 12-in. diameter disks store about the same amount of data on each surface, almost everything else about them is different. The manner in which data is recorded and the type of media on which it is recorded distinguish the drives in the 12-in. format. The drives also differ in access time and data rates, as well as in the construction of the cartridge in which the optical media are housed.

Similarities in format and media between the OP 1000 from Optimem (Sunnyvale, Calif) and the Alcatel Thomson (Redondo Beach, Calif) Gigadisc 1001 derive from the long-term relationship between their parent companies, Xerox and Thomson-CSF. Even though these two units employ the same media (supplied by Alcatel), the same recording format, and interface to the host over the SCSI interface, their optical and mechanical designs differ.

Yet, Leo LaBonte, East Coast sales manager for Alcatel, maintains that disks written on Optimem drives can be read on Alcatel drives, and vice versa. According to Optimem's director of marketing, Larry Fujitani, this degree of compatibility, along with the fact that both drives and media have two sources (unique situation in the current optical disk industry), increases their chances for survival. Fujitani adds that these factors give them a better chance of developing standards around their formats.

Alcatel claims that adding a few more commands to its SCSI set allows it to provide a direct copy feature. This feature, according to LaBonte, permits two drives to be synchronized. Thus, data can be transferred directly from one drive to another during the copying process. He adds that this is essential to organizations concerned with maintaining security, and that it significantly reduces the time required to make copies.

Other drives, those developed by Hitachi and Optical Storage Interna-



tional (Santa Clara, Calif), for example, do not share this degree of compatibility. Hitachi's model 301 differs from these drives in recording format, type of media, and transfer rate. OSI's LaserDrive 1000 has adopted the SCSI interface. But, like the Hitachi unit, the LaserDrive uses ablative (pit forming) technology media, rather than implementing bubble forming technology used by Optimem and Alcatel.

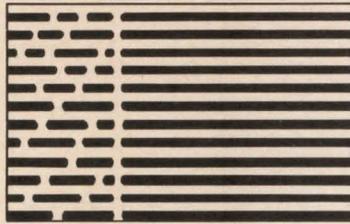
Choice of media, according to Taroon Kamdar, director of marketing at OSI, is an important aspect for archival storage. He points out that the ablative technology that OSI uses actually burns a hole into the sensitive layer of the disk where data is recorded. While using a similar optical system, Optimem and Alcatel have chosen media technology that forms a bubble on the surface of the sensitive layer where data is recorded.

Ablative technology, Kamdar claims, "has proven superior in terms

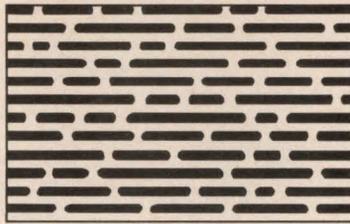
of shelf life of the data once it is recorded on the media." These claims of superior data retention, he says, are based on extensive research into various methods for recording data onto optical media. The studies were performed by NV Philips, Control Data's partner in OSI. However, Alcatel and Optimem specify the same 5-year shelf life before recording and greater than 10-year lifetime after recording for their media.

Another distinct difference between the drives lies in the method used to verify whether data have actually been recorded on the disk's sensitive layer and whether that data can be read back accurately. OSI claims to have improved upon the direct read after write (DRAW) method that is used by Fujitsu, Hitachi, Optimem, Alcatel and others, with what is known as direct read during write (DRDW) techniques.

In DRDW, when the drive receives a command to write, it not only



(a)



(b)

Magnified many times, these sections of read-only and write-once optical media show distinct differences. Write-once media (a) have pregrooves that also contain sector information (left) to aid in locating data. Actual user data is recorded along grooves (right) within the user portion of the tracks. Read-only media (b) contain data as well as address information. No distinction is apparent between the two types of information.

writes but verifies simultaneously that the write actually occurred. The technique uses essentially the same optics as those used for read and write. It is merely "a technically clever way to use the same laser beam that is used to write to tell whether or not a good write was achieved," according to Kamdar.

When other drives receive a write command, they put out a write pulse to record data on the media. They must then wait one revolution of the disk before they can read the data to verify whether the write was done properly. This means that even when writing sector to sector, 60, 70, or 100 ms may elapse before a good

write can be verified. In terms of performance, Kamdar says, "this makes the OSI drive far more efficient than other drives."

Another distinction between drives lies in the realm of user friendliness. This difference is in respect to the cartridge that houses the media and protects it from damage when the disk is outside of the drive.

For the Optimem drive to read a particular disk, the cartridge with the media is first inserted into a slot on the drive. The cartridge sleeve is then pulled out, leaving the disk itself within the drive. The sleeve must then be inserted into a separate slot for storage. This requires three motions for insertion, and consequently, it requires another three for removal. OSI's drive, on the other hand, spins the media within the cartridge after the cartridge containing the media has been inserted. This not only affords continuous protection for the media, but provides more convenience for the user as well.

AUTO-INSERTABLE ACTIVE DELAY LINE

5 OR 10 TAP TTL BUFFERED INPUTS AND OUTPUTS

- AUTO-INSERTABLE STANDARD 14 PIN DIP PROFILE:
HEIGHT .200" INCLUDING STANDOFF
LENGTH .770"
WIDTH .300" ON LEAD CENTERS
- TRANSFER MOLDED THICK FILM HYBRID CONSTRUCTION FOR HIGH RELIABILITY
- U.S. MANUFACTURED ON AUTOMATED ASSEMBLY AND TEST EQUIPMENT FOR REPEATABILITY
- COMPETITIVE PRICING
- COMPLETE STANDARD PRODUCT LINE AND CUSTOM DESIGNS AVAILABLE

FOR MORE INFORMATION CONTACT...
JO-ANN RICCI (408) 395-2300 X265

HYTEK MICROSYSTEMS
980 UNIVERSITY AVENUE
LOS GATOS, CA 95030

Planning for the future

Smaller drives now under development, and probably available in the spring, will store about half a gigabyte per side. Using a 4.7-in. (12-cm) platter (the same size as accepted CD ROMs) or moving into the 5¼-in. diameter of conventional floppy disks, the drives will implement similar write-once techniques and media. Targeting vendors that develop 5¼-in. form factor drives, both Hitachi/Maxell and 3M have released specifications for disks that will provide capacity for 250 to 500 Mbytes/side.

Drives expected to appear around mid-year will rely on such media. One of these, from Information Storage, Inc. in Colorado Springs, Colo. is reported to store a conservative 100 Mbytes. This capacity is attained by recording at about 11,500 bits/in. and 14,000 tracks/in. This conservative design should help to meet the unit's full production price target of \$500 in quantity—a price that is necessary to fit the projected desktop computer market.

Although formal announcements have not been made, it is clear that other companies will be joining ISI

HOW TO DOUBLE YOUR CONTROL WITHOUT DOUBLING YOUR CONTROLLER.

If your plans for controlling disk and tape drives call for two boards and two CPU slots, you can now cut your plans in half.

And use the SPECTRA 25 LSI-11 emulating multifunction controller instead. You'll get all the control of separate disk and tape boards in just half the space. Half the spares. While using less power.

SPECTRA 25 is the smarter way to handle SMD/Winchester and 1/2-inch tape drives, and it's the first and only multifunction controller designed for DEC Q-BUS® computers.

Its block-mode data transfer capability can increase your throughput by up to 50%. And it performs Error Correction Code (ECC) independently of the operating system. All 32 bits worth.

SPECTRA 25 fully emulates DEC's RM02/5 and RM80 disk subsystems, and TS11 tape subsystem. It can run two SMD disk drives and four tape drives, in any combination. Plus an E²PROM lets you easily reconfigure while the controller is in your system.

But SPECTRA 25 puts something else in your system, too: proven reliability. Because you're working with half as many parts and pre-tested ICs.

You're also working with excellent documentation that's regularly updated. And the most comprehensive one year warranty in the industry. All backed by responsive nationwide service.



SPECTRA 25 is the multifunction controller that doubles your control. And your advantage on the competition.

Before you look at Emulex or Dialog, consider the SPECTRA 25 from Spectra Logic.

It could be all the controller you'll ever need. For further information on our expanding family of controllers for DEC Q-BUS and UNIBUS® computers, call or write us today.

	SPECTRA 15*/ SPECTRA 25	SPECTRA 111*/ SPECTRA 121	SPECTRA 12*/ SPECTRA 21
BUS	Q-BUS (Block-mode)	UNIBUS	UNIBUS
Disk Emulation	RM02/05, RM80, RP06 (2 drives)	RM02/05, RM80, RP06 (4 drives)	RK06/07 (4 drives)
Tape Emulation	TS11 (4 drives)	TS11 (4 drives)	TS11, TM11 (4, 8 drives)
Disk Transfer Rate (SMD)	2.0MB/sec	2.0MB/sec	2.0MB/sec
Configurable On-Line	Yes	Yes	No
*Single function disk controller			

Spectra Logic Corporation, 297 North Bernardo Avenue, Mountain View, CA 94043, (415) 964-2211, TWX 910-339-9566, TELEX 172524 SPL MNTV. International Sales Office: The Netherlands (31) 23 273744, TELEX 71080 SPECTL.

SPECTRA LOGIC

A Cipher Data Products Company.

Manufacturing the media for optical systems

The process of producing an optical disk—whether intended for use in an optical ROM or write-once, read-mostly drive—begins with a digital format recording on tape. The mastering process uses a system made up of a processor, encoder, and laser cutter to create glass masters. Digital information is recorded on the master by exposing a photoresist layer (or coating) on the glass plate (or substrate) to a laser beam. This beam cuts pits or raises bubbles in the photoresist layer. The resulting marks correspond to digital 1s and 0s.

After the glass master is developed, it is coated with a thin conductive layer. An electroplating process lays down a thick metal layer on top of the conductive material. This yields a metal master that can be separated and used to produce stampers. When placed in an injection molding machine, the stampers then transfer the recorded pattern of marks onto a transparent disk formed by injecting a high grade polycarbonate resin.

After molding, this transparent disk is placed in a vacuum chamber where a very thin layer of aluminum coats its surface. This layer reflects the drive's laser beam, allowing information to be read from the disk. The pits or bubbles reproduced in this aluminum layer are protected by an ultraviolet curing layer that is applied over the aluminum.

The process used to manufacture media for use with write-once optical systems differs in one respect from that used to form compact disk ROMs. In the latter instance, a different modulation code is used to turn the laser beam on or off. While the compact disk ROM process writes tracking information, address information, and data, the process used for write-once media writes only the pregroove. No matter whether intended for use with ablative or bubble raising media, this pregroove provides the drive with the radial servo tracking feedback, data synchronization, and preformatting to supply head positioning feedback. This puts the precision needs to achieve high areal densities on the disk rather than in the drive.

in the 5¼-in. market. OSI is developing a unit that exhibits a standard 5¼-in. footprint. Optimem also has a unit underway, and Alcatel projects a 1986 timeframe for introduction of a unit using 5¼-in. diameter media. Although IBM is said to be working on a drive of this type, and is known to have optical development projects underway, exact product plans remain unclear.

In addition to the arrival of downsized write-once drives, optical aficionados can soon expect to be doubling storage capacity on 12-in. diameter media. Widespread availability of two-sided media will make recording and reading both sides of the media common in these drives.

Fueling interest in optical storage devices is the possibility that long-term efforts aimed at developing reversible media have finally produced an optical drive with a magnetic drive's read/write capability. Recent developments in this area include an erasable read/write drive based on the

magneto-optic technology demonstrated in prototype by Canon in Japan. Matsushita's Panasonic Div has also demonstrated reversible media based on phase-change technology.

A step at a time

Meanwhile, a solution to the problem encountered in getting acceptable life times for reversible media appears within reach. Sony claims it has come up with an amorphous film that shows only minimal degradation of signal and noise after more than a million write/erase cycles. Hitachi has developed an optical disk alloy that changes color when heated and rapidly cooled. This color is reversed by reheating and slowly cooling the material.

Such developments will eventually produce reversible media that are stable enough to retain data for the 10 year period required in archival applications. Further, they will endure sufficient read/erase cycles to prove cost-effective in ordinary data record-

ing applications. Commercially available drives using this media will follow closely, because, in this case, the necessary optical systems are not out of reach.

Reversibility, although a long-sought goal, is only one step forward in optical storage technology. Another area stimulating research and development is the attempt to match access times and data rates of optical drives to those of magnetic disk drives. Access times will fall when a smaller optical head assembly can be devised.

These developments can be expected within the next few years. To increase the data rate, it will be necessary to increase the rotational speed of the disk. Since this requires that the laser provide a higher power output, this achievement awaits the creation of more powerful laser diodes.

A long-life laser diode with acceptable noise characteristics that can produce the 100-mW peak power output necessary may appear within the next two or three years, according to Alan Bell of IBM's San Jose Research Div. An alternative under review by RCA and Bell Labs is use of arrays of laser diodes that can record or read data in parallel from adjacent tracks. This would yield higher data rates without an increase in the rotational speed or output power level of lasers in the array.

Success in these and related efforts will provide optical disk drives with the opportunity to dethrone magnetic recording media. In the interim, growing demands for online data storage will draw from each technology those features that best suit specific needs. Optical drives now hold under 1 percent of the storage peripheral market. Raymond Freeman of Freeman Assoc (Santa Barbara, Calif), one authority in the optical market area, however, predicts that from the 6000 units bought in 1984 the optical share will grow to some 2.4 million units in 1990.

—Peg Killmon,
Senior Editor

SYSTEM TECHNOLOGY
(continued on page 33)

Pack it in.

HIGH-VISIBILITY
"RELEASE/LOCK" INDICATOR

STRESS-FREE
CAM ENGAGEMENT

SPRING LATCH FOR POSITIVE
CHIP CARRIER RETENTION;
NO TOOLS REQUIRED

FIELD-REPLACEABLE
CONTACTS FOR THRU-HOLE
& SURFACE MOUNT
APPLICATIONS

THRU-HOLE
OR SURFACE
MOUNT
VERSIONS

SPACE-SAVING DESIGN ALLOWS
FOR "PIGGY-BACKING" OF ICs

OPTIONAL BUILT-IN
DECOUPLING
CAPACITORS

ONE-PIECE CONSTRUCTION;
CHIP CARRIER MAY BE
REMOVED WITHOUT
DISASSEMBLY

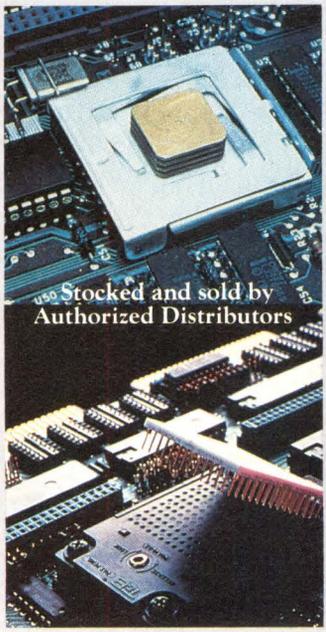
TAPERED LEADS FOR
EASY PCB INSERTION

Now... socket complex VLSI packages with Ansley® IC interconnect devices.

The Electronics Group of Thomas & Betts responds to the increasing complexity of integrated circuits.

Ansley® Pin Grid Array Connectors feature zero-insertion-force design for dense terminal counts as high as 625 I/O. They interface popular pin grid array packages while offering advantages over conventional systems: selective plating, high-performance glass-filled thermoplastic materials, and stress-free cam engagement to open each dual redundant contact.

Ansley® Leadless Chip Carrier Connectors interconnect microprocessor and gate arrays, as well as high-density custom logic semiconductor devices. In standard terminal counts of 68, 84, 132 and 164, they're designed to accept JEDEC leadless Type A packages with .050" and .025" centerlines, and allow for access-probing of contact leads after board installation.



Stocked and sold by
Authorized Distributors

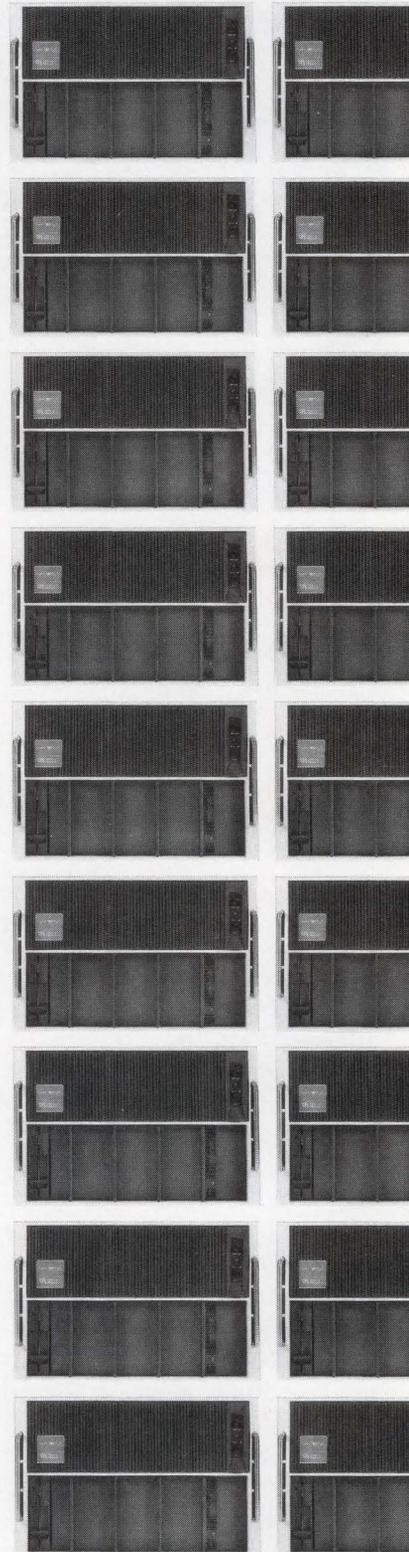
Ansley® Stacked DIP Sockets double IC packaging densities by enabling designers to vertically stack two standard plastic or ceramic DIP packages. Available in 16 through 28 positions, the Stacked DIP Socket accommodates both memory and logic/support devices.

For more data on these products, and our screw-machine line of DIP and PGA sockets, contact Technical Service, Thomas & Betts Corporation, Electronics Group, 920 Route 202, Raritan, NJ 08869; (201) 469-4000.



Where Interconnection Meets Imagination.

CARTOS™ dataCASE/5™



It's true. Because our CARTOS dataCASE/5™ real-time computer system behaves like 64 virtual machines.

So. Its real-time capability is all yours. Even when you're sharing the system with 63 other demanding environments.

Believe it. With up to 4 MB of main memory, CARTOS dataCASE/5 is real-time that's right for the rigors of the industrial world. Tough. Strong. Reliable. CARTOS is already hard at work in oil fields and nuclear power plants as well as robotized assembly lines and sophisticated transportation systems.

CARTOS dataCASE/5 is flexible, too. Use whatever media you want. Put it wherever you want. In a rack. On the floor. Or atop a desk.

Just imagine. You can tell everyone you're buying 64 real-time systems. And then call us. Computer Automation. (714) 833-8830, Ext. 353. TWX: 910-595-1767. 18651 Von Karman, Irvine, CA 92713.

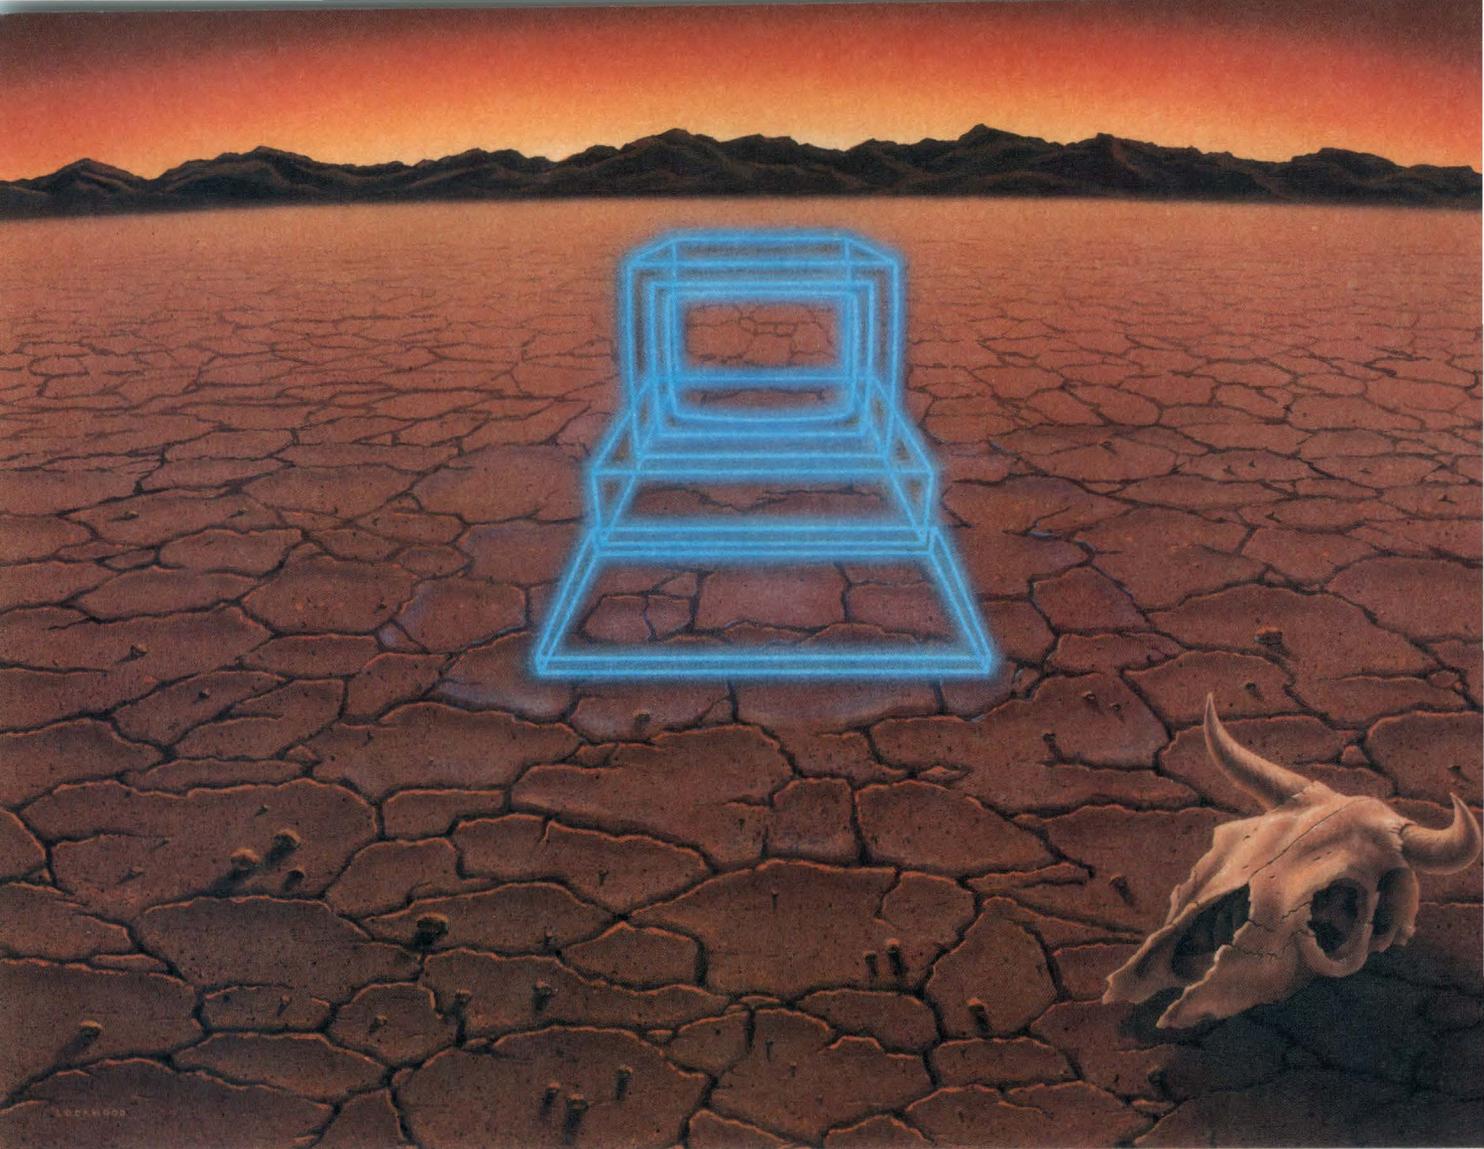
© 1985 Computer Automation, Inc.



ComputerAutomation®
Computer Products Division

BUY ONE, GET 63 FREE.





CAD without CACHE™ is like life without water.

Now CACHE brings the CAD design process to life with the information you need. No more long dry spells while you wait for your CAD Library to be created. CACHE offers the design engineer a comprehensive component library of

Tell me more about how CACHE can add life to my CAD system.

Name _____

Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Send to: **INACOM International, 4380 South
Syracuse Street, Denver, Colorado 80237**

CIRCLE 17

standard and semi-custom IC's that support logic design and physical layouts for engineering workstations and turnkey CAD systems. With the information from CACHE at your fingertips, you can shorten your design path, reduce design costs and improve productivity. So learn how to breathe new life into CAD with CACHE. Call us or send in the attached form today.

Call: **1-800-443-INFO**
303/694-4200 (Within Colorado)

INACOM INTERNATIONAL COMPUTER AIDED ACCESS™
We go right to the source.

AI techniques aim to ease VLSI design

Too many chips, with too much complexity, and too few people to create them are the biggest problems facing system designers today. Design tools offer some relief. But even today's most sophisticated tools cannot meet the demands made by the large number of circuit functions present in digital VLSI design.

Artificial intelligence techniques applied to VLSI design systems could resolve the dilemma, however. Instead of starting each new project from scratch, a designer could simply fine-tune the processing and circuit specifications embedded in the knowledge base to fit the design's intended parameters.

Current research is proceeding along two parallel paths: synthesis of chips and analysis of finished designs. Most of the progress in applying AI techniques to VLSI design has occurred in analyzing hardware designs, primarily in matching the final design with the original specifications. This progress was possible because the re-

quired knowledge base (ie, the VLSI design specifications) already exists. Synthesizing a design is much more difficult, however, because the knowledge base must first be constructed. The key is to generalize as many design parameters as possible, and then infer the needed next steps in the design.

The AI/VLSI group at Rutgers University (New Brunswick, NJ) represents the more active groups working on both analysis and synthesis. One of its first projects was Redesign, an intelligent assistant for the functional redesign of digital TTL parts. Redesign's reasoning function distinguishes between what is true of the circuit and what must be true for the circuit to work correctly.

In a functional redesign, the system is given the schematic of a working digital circuit along with its functional specifications. The system also receives a design plan that relates the circuit schematic to its specifications. If the functional specifications are

changed, the system's task is to redesign the circuit so that it will meet the new set of functional specifications.

Experience is the best teacher

The Redesign project served well for simple circuits of a few gates, but proved too cumbersome for designing more sophisticated functions (such as floating point additions). What's more, it only determined which portions of a circuit could be modified without altering the function of the circuit. (For example, it showed where an adder could be modified, within a multiplication circuit, without changing the overall multiplication function.)

A new class of Rutgers' knowledge-based systems should leap beyond Redesign's capability. One of these, actually called Leap (short for Learning Apprentice Systems), is being written in Lisp OPS5 using a VAX-11/750 computer. Professor Louis Steinberg, head of the Leap project, maintains that systems such as these will break

Function to be Implemented:

Inputs: Input1, Input2, Input3, Input4

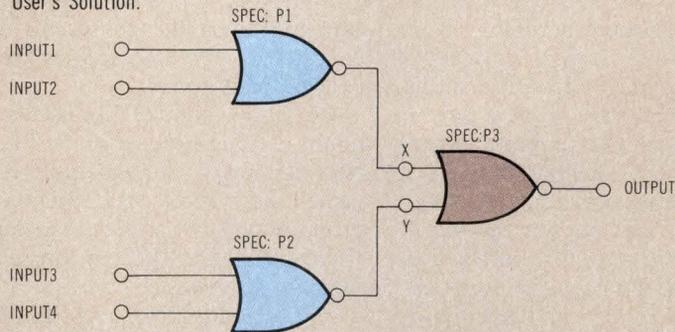
Outputs: Output

Function: (Equals (Value Output(i)) (And (Or (Value Input1(i)) (Value Input2(i)) (Or (Value Input3(i)) (Value Input4(i))))))

Where Input Signals Satisfy:

(Datatype Input1(i))=Boolean
 (Value Input1(i))=i Mod 2
 (Encoding Input1(i))=Positive-logic
 (Start-Time Input1(i))=i-100+t₀
 (Duration Input1(i))=75 ns
 (Datatype Input2(i))=Boolean
 (Value Input2(i))=unknown
 (Encoding Input2(i))=Positive-logic
 (Start-Time Input2(i))=i-100+t₀
 (Duration Input2(i))=65 ns
 (Datatype Input3(i))=Boolean
 (Value Input3(i))=unknown
 (Encoding Input3(i))=Positive-logic
 (Start-Time Input3(i))=i-100+t₀
 (Duration Input3(i))=58 ns
 (Datatype Input4(i))=Boolean
 (Value Input4(i))=unknown
 (Encoding Input4(i))=Positive-logic
 (Start-Time Input4(i))=j-100+t₀
 (Duration Input4(i))=75 ns

User's Solution:



An expert system, such as Leap, learns by example and modifies or adds to its rules accordingly. For example, this circuit module computes the Boolean product of sums of four particular input signals, specified to be regular streams of Boolean values. The exact definitions of the module's function are described by an array of data elements. Each data element is defined in terms of its value, start time, duration, data type, and encoding. While Leap may already contain a rule specifying a circuit that consists of an AND gate and two NOR gates, the user might decide to implement the circuit with three NOR gates instead.

a major bottleneck in today's expert systems.

To handle the language-acquisition problem, this new class of knowledge-based systems incorporates a learning component for acquiring knowledge through experience. These interactive knowledge-based consultants directly assimilate information by observing and analyzing the problem-solving steps carried out during normal use of the system. The system's design enables it to acquire knowledge continuously, without an explicit training mode.

Leap, therefore, provides advice on how to refine the design of a VLSI circuit—while allowing the user to override this advice and to refine the circuit manually, whenever desired. When the user manually refines the circuit, Leap records the problem-solving step as a training example. It then generalizes from this example to form a new rule that summarizes the refinement tactic.

Steinberg expects this type of system to offer strong advantages over the current architectures of knowledge-based systems. "By distributing copies of Leap to 1000 circuit designers, the system would quickly be exposed to more example circuit designs than a human designer could hope to see during a lifetime," he says. This large experience base might allow for significant learning—even if the learning mechanism is less sophisticated than the human learning process.

Partitioning Leap's knowledge base into two parts allows it to adjust existing rules when a new rule is added to a knowledge base, according to Steinberg. One part consists of rules that characterize correct circuit implementations, but not necessarily the preferred ones. The other part provides control knowledge for selecting the preferred implementation from among multiple legal options.

To date, the Rutgers group has used only the part of the knowledge base consisting of implementation rules because it is the simpler of the two. To explain or verify an example of an implementation rule, the system need only verify the correctness of the circuit fragment identified in the training example. On the other hand, to learn a control rule that

IF
 the most current active context is to
 create a link
 and the link should go from a source
 port to a destination port
 and the module of the source port is
 not a multiplexer
 and there is a link from another
 module to the same destination
 port
 and this other module is not a
 multiplexer

THEN
 create a multiplexer module
 and connect the multiplexer to the
 destination port
 and connect the source port and
 destination port link to the
 multiplexer
 and move the other link from the
 destination port to the multiplexer

Carnegie-Mellon's Design Automation Assistant uses an OPS5 knowledge-based expert system as a tool to formulate a design problem. The tool has three major components—working memory, rule memory, and rule interpreter. These rules enable the DAA to synthesize an acceptable design by determining, at each step, whether or not a certain design extension respects its constraints.

establishes when some implementation is preferred, the system must compare that implementation with all possible alternatives.

Obtaining control knowledge remains a long-term goal for Leap, with its initial implementation expected later this year. Steinberg's group hopes to test it on a group of students in a VLSI design course during the fall semester.

A group at Carnegie-Mellon University (Pittsburgh, Pa), headed by Professor John McDermott and graduate student Jim Kim, is also developing second-generation VLSI design tools based on AI techniques. Dubbed the Design Automation Assistant, the system aids in synthesizing cell layout. DAA is an outgrowth of a previous Carnegie-Mellon project call Talib, an automated design layout program. Both systems are implemented in Lisp

OPS5. Talib's input is a description, in netlist form, of the schematic of an NMOS digital circuit. From this schematic, a description of the mask geometry is produced as output.

The original Talib consists of about 1200 rules, of which approximately 940 are associated with specific subtasks. Talib's major design activity consists of selecting design constructs on the basis of boundary conditions. Then, it partitions the layout surface into zones and characterizes the spatial relationships between the zones. After partitioning the netlist into known topological groupings, it places the groupings of subcircuits in different zones of the layout surface.

This design plan, while adequate for small designs, becomes cumbersome for large designs (those reaching a complexity of 250,000 transistors). The Carnegie-Mellon group took a new approach by expanding the Talib into the DAA. The new system implements design synthesis by using a large amount of design knowledge. This eliminates searching through the rules.

In general, knowledge-based expert systems are developed in stages; this was true for the DAA as well. First, theoretical knowledge of the problem is codified as a set of situation-action rules. Interviews with experts are used to fill in knowledge gaps and refine current knowledge. Experts then examine and validate the results achieved when example problems are given to the system. New rules are then added to the system to correct errors found through the examples.

Applying the rules to chip design

The MCS6502 microprocessor was among the first designs the Carnegie-Mellon group created. In this instance, the DAA used 70 rules and three hours on a VAX 11/750 computer. The group recently completed a more ambitious design: an IBM System/370 mainframe on a chip based on a complete algorithmic description.

Description of the 370 included memory-management operations, channel controller I/O instructions, and all 370 instructions (except the extended-precision floating point), the characters under mask, the edit and

VIVE LA DIFFERENCE!

NEC PRESENTS 256K DRAMs FOR EVERY TASTE.

Discriminating designers deserve a choice. In high-density DRAMs, NEC gives you three. Three 256K configurations to suit your application. And they're all on tap today.

Going for growth in terminals, video or PCs? Go with our 64K x 4. Or help yourself in mainframes and telecom with our 256K x 1. Nibble or page mode.

Superior quality, smooth supplies

One thing our DRAMs don't have. Redundancy. So you get simpler circuitry. From a superior part. What does superior mean? 0.1% AQL with 100% burn-in. The highest quality in the industry.

You want it fast? Our new plant in Roseville, California is ramping up for high-volume memory production. And it ranks with the most futuristic fab facilities on earth.

Add that special sparkle

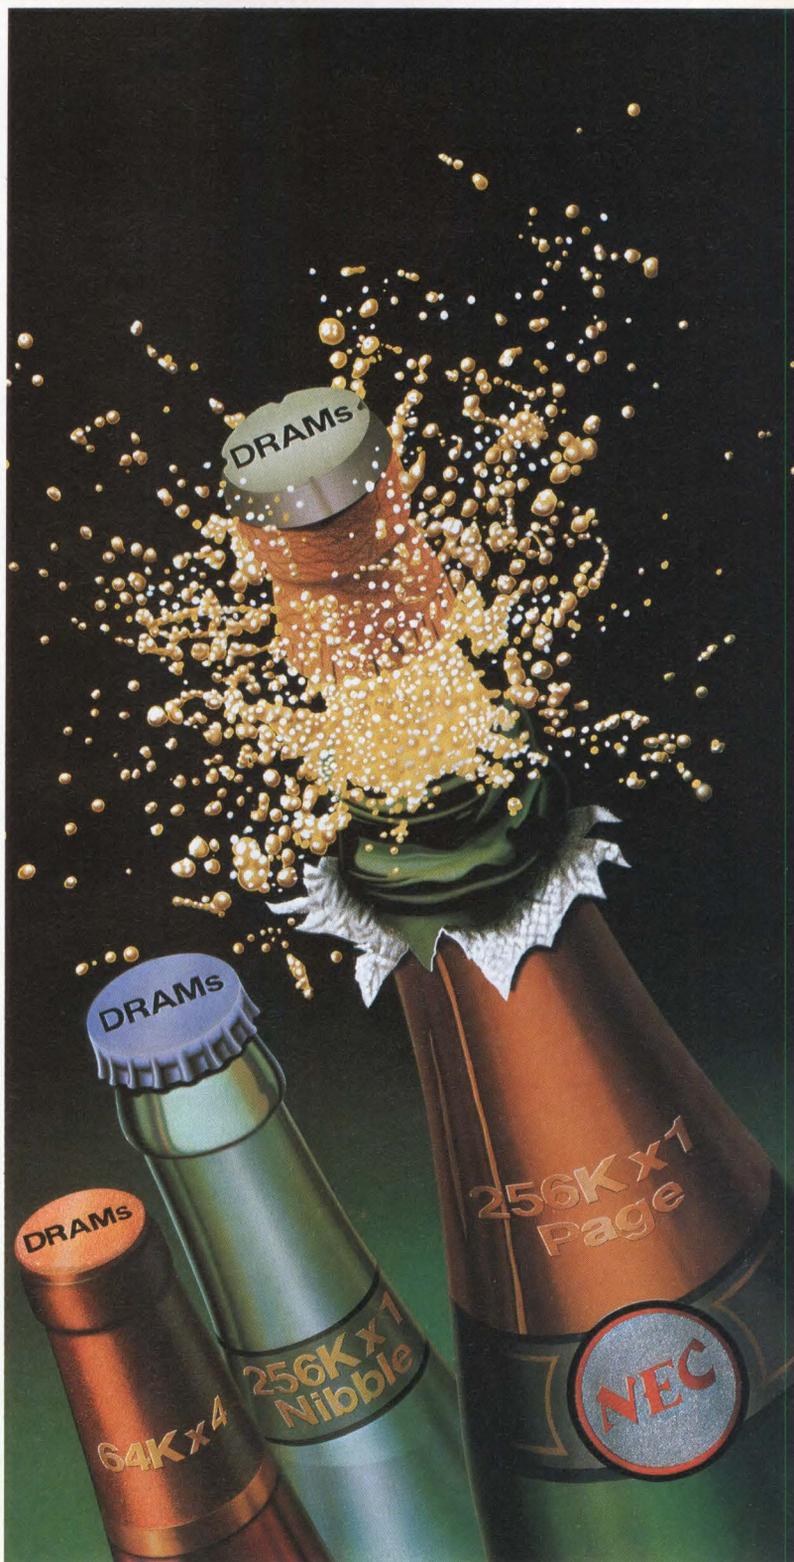
With matchless quality and modes to match your market, our premium selection of DRAMs can make your system sparkle.

DRAM Selection Guide

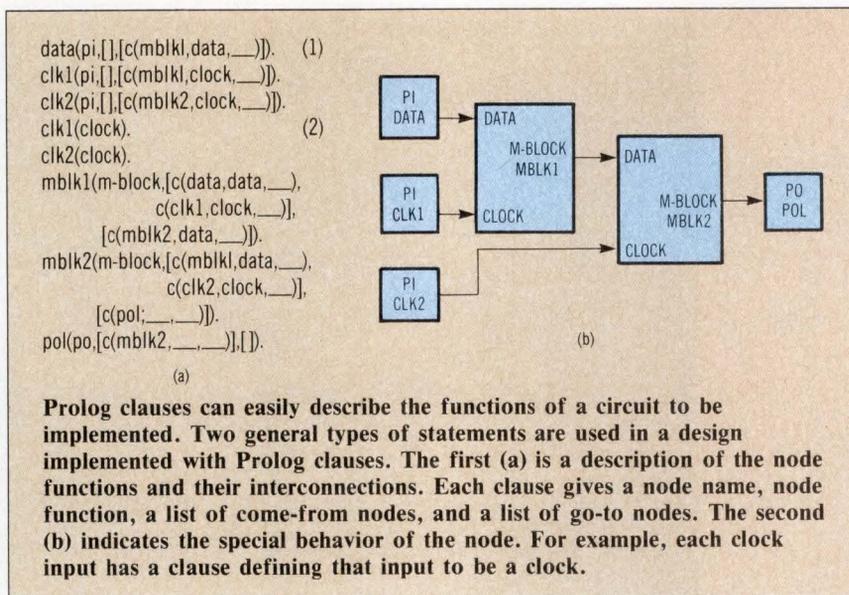
DEVICE	ORG.	ACCESS TIME (NS)	CYCLE TIME (NS)	STNDBY	ACTIVE
uPD41254-15 uPD41254-20	64K x 4	150 200	270 335	28	413
uPD41256-12 uPD41256-15 uPD41256-20	256K x 1	120 150 200	220 270 335	28	385
uPD41257-15 uPD41257-20	256K x 1	150 200	270 335	28	385

For details, call NEC at 1-800-556-1234, ext 188.
In California: 1-800-441-2345, ext 188.

© Copyright by NEC Electronics, Inc.



NEC
NEC Electronics Inc.



mark, and the packed-decimal. The resulting description was more than 10 times the size of the MCS6502 microprocessor. The DAA's version of the 370 required 47 hours on a VAX 11/780 configured with 6 Mbytes of memory and two memory controllers.

Since IBM had already designed the 370 on a chip, Carnegie-Mellon's DAA project leader and primary researcher, Thaddeus Kowalski, wanted to compare the two. He asked Claud Davis, IBM's design team manager for the micro 370, to contrast the DAA's paper design with IBM's. According to Kowalski, Davis said "The 370 data flow exhibited the quality I would expect from one of our better designers."

Kowalski is now part of the AI department of AT&T Bell Labs (Murray Hill, NJ). The team is applying his DAA experience to the design of a very fast (GHz) digital processing chip with 0.5- μ m line width. Again, the design is written using OPS5 on a VAX/780 machine.

The other half

Designing sophisticated VLSI chips is only half the problem. The other half is verifying that the design is correct. One way of verifying a design is to "prove" that it cannot be incorrect, rather than simulating correctness.

Dr Harry Barrow, chief scientist at the Computer Aided Systems Group

of Schlumberger Palo Alto Research, Palo Alto, Calif has a formula for determining correctness. "The key is to provide a structural description and a behavioral specification of the design, and to prove that the description and the specification are essentially equivalent." Barrow's verification method, called Verify (and implemented in Prolog), can automatically prove the correctness of straightforward, but very detailed, designs involving many thousands of transistors.

Verify identifies the design to be implemented by using modules that are hierarchically organized and modeled as finite-state machines. A module has a set of input ports and a set of output ports. Each port has an associated signal type that specifies the conditions under which the signals pass through it. The behavior of the module is specified by two sets of equations. One gives output signals as functions of inputs and the current internal state, and the other gives new internal states as functions of inputs and the current state.

Moreover, Verify can derive a description of the behavior of the whole system when given the behavior of the components of a system and their interconnections. The derived behavior can then be compared with a specification of the intended behavior of the system.

Another key feature of Verify is its ability to exploit the hierarchical

structure of a design and break it into smaller pieces. This allows separate and recursive verification.

Living up to its name, Verify has verified several designs. One such design consists of three multipliers and two adders that compute sums of products. The multipliers are built up with slices. Each slice contains a 1-bit multiplier, a shifter, and an adder. The adders are built from full adders, which, in turn, consist of logic gates. These gates are described on two levels—an abstract Boolean function level and an NMOS transistor model level involving tri-state signals and stored charge.

In its simplest composition, the circuit design contains approximately 30,000 primitive parts including 18,180 transistors. The entire design specification occupies about 400 Prolog assertion statements. Verify proved the correctness of the module in about 10 minutes on a DEC 2060 and produced about 4800 lines of output code. According to Barrow, the proof holds for all of the 134 million different input patterns between each of the primitive parts of the circuit design. "To have this design checked with exhaustive simulation in the same amount of time would have required a simulator capable of over four billion transistor-operation simulations per second."

In its next stage of development, Verify will be able to check the correctness of the timing information and perform functional proofs of designs. Also, Barrow's group is currently engaged in a joint effort with Stanford University (Palo Alto, Calif) on the Helios project. Here, hierarchical simulation will be used to design highly parallel arrays of microprocessor elements for use in future computers.

Prolog has caught another research center's fancy for application to computer aided design applications. IBM's General Technology Division (White Plains, NY) is using some 400 Prolog clauses to define a design-for-testability expert system for VLSI design. In conjunction with Syracuse University (Syracuse, NY), Paul Horstmann, member of the technical staff at IBM, is conducting experiments that use logic programming to

"Mux Master™ is a heck of a buy."

"What we have here is multiplexing magic, and it looks to me like just what the doctor ordered. You have been nagging at me for months now about the lack of system flexibility available to DEC users. Mux Master not only fills that gap with the kind of flexibility you've been asking for but keeps bus loading down to a low, low minimum in the process.

"A single Mux Master host board, installed with our cluster controllers, lets you make 128 Unibus or 64 Q-bus terminal connections through a composite

LAN SIMPLICITY AT HALF THE COST

HOST INTERFACE

QBUS OR UNIBUS HOST

CLUSTER CONTROLLERS (8-16 LINES)

MORE CLUSTER CONTROLLERS

ADD LINES WITHOUT INSTALLATION COST

SUPPORTS WORK-CLUSTERS

NO MORE EXPANSION CABINETS

SOFTWARE TRANSPARENT

ON-LINE REPLACEMENT

Ken O'Mohundro, President, Able Computer, on the merits of his company's new multiplexing system.

link. You order as many 8 or 16 line cluster controllers as you need when you need them and mount them on a wall, desk or floor stand anywhere on the link.

"Then, as your requirements change, you add new controllers or remove existing ones simply and without disrupting the operation in any way. Anyone can do it. Concentrated one-time wiring makes it possible for the user to complete the process in a matter of minutes.

"That's all I can tell you now. For more information, you'll have to write or call toll-free at (800) 332-2253. When you do, ask for the complete Mux Master story. Able Computer. 3080 Airway Avenue. Costa Mesa, California 92626. (714) 979-7030. TWX 910-595-1729."

Able Computer

DEC, Unibus and Q-bus are trademarks of Digital Equipment Corporation. Attach™ is patent pending.

solve some basic VLSI design problems. Logic programs were developed to address functional simulation, fault diagnosis, and automatic test generation problems. A completed prototype expert system will handle designing VLSI chips for testability.

According to Professor Edward Stabler, who worked with Horstmann at Syracuse, much research is needed before Prolog becomes an effective "everyday" VLSI design. One handicap, Stabler says, is that current implementations of Prolog allow a rather limited number of clauses to be handled by the data base.

For instance, the design model contains at least one clause for every functional node in the design. "If we restrict ourselves to a level higher than the gate-level, this is still 10,000 clauses for a typical VLSI design," says Stabler. If all the clauses for the CAD rules are added, as well as those generated while processing the design, this number could very well end up to be 20,000, he adds. This number of

clauses is beyond the capabilities of current implementations of Prolog.

Another corporate group hard at work at applying AI principles to VLSI design is at the Gould Research Center in Chicago, Ill. Gould has developed a functional-to-structural translator for macrocell design called Descart. It is geared toward defense applications. The two-year-old project is an offshoot of work done on a heuristic silicon compiler that used AI techniques.

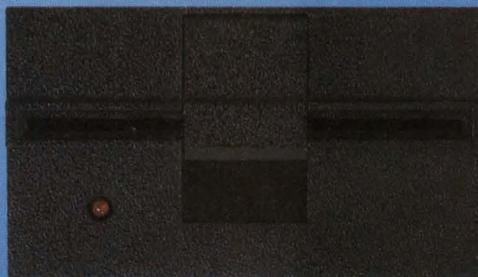
Project Manager Bob Kuhn identifies four major constraints faced by standard cell designers who use today's CAD tools: the cell's area, power requirement, pin configuration, and timing. "Most CAD equipment on the market today offers adequate solutions to the timing, power, and pins requirements, but few address the area requirement. Our research efforts on Descart are focused on finding the optimum area for each standard cell module." Kuhn says he hopes to start imple-

menting Descart this summer in the design of a macrocell set of 10,000-gate equivalent standard cells.

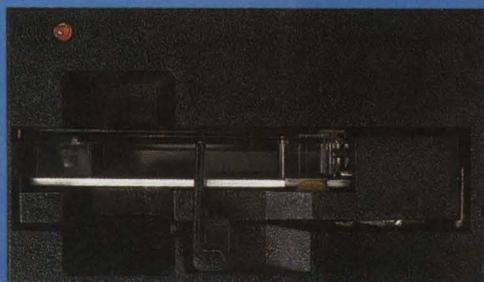
Even if applications of AI techniques to VLSI were farther along, the technology for emulating human expertise in VLSI design remains too expensive to apply commercially. Embedding all of the existing design knowledge in a 100,000-transistor chip is beyond the scope of today's expert systems. This step must await the arrival of the computational power of an AI language, such as Lisp itself, embedded on a chip.

Chip designers might not have long to wait: the first U.S. Department of Defense contract to produce a Lisp-on-a-chip was awarded to Texas Instruments (Dallas, Tex) last year. The chip, which should be ready by 1987, is to have 2 to 10 times the processing power found in today's symbolic processors.

—Nicolas Mokhoff,
Senior Editor



THE END OF THE SECONDARY STORAGE SPECTRUM EVERYONE FORGOT ABOUT WAS THE MIDDLE.



Unix interface standard on the way

A standard for the application software interface to Unix-compatible operating systems is taking shape. The basis for the standard is a core of subroutines and system calls that defines the applications programming environment for an operating system. This preliminary standards document, developed by an ad hoc committee calling itself the /usr/group, has been accepted by the newly formed P1003 committee of the IEEE. Ultimate adoption of a standard at the source-code level would ensure portability of application software between a variety of computers.

The /usr/group standards document focuses on the C-language interface to operating systems, and is based on the AT&T System III version of Unix. But it takes into account other versions, especially System V and Version 7. As with any standard, however, there are differences of

opinion on several points. There are, for example, at least eight variations between the /usr/group document and the Unix System V, Release 2 interface definition document just released by AT&T.

Though incomplete, the preliminary standard represents a starting point that should prove valuable to the Unix community because of its timeliness. Moreover, it will evolve and several extensions are already in progress or planned. These include terminal handling, networking, realtime software constructs, inter-process communication, and object file formats.

One example of the differences that exist and how the characteristics of System V will influence the development of the standard is System V's IOCTL system call. In the present /usr/group standard, details of terminal control I/O are omitted even

though the IOCTL system call that takes care of such chores is included. Consensus on how to handle the terminal control I/O could not be achieved due to the differences in the IOCTL call among the various versions of Unix.

"The fact that IOCTL says nothing about how to control terminals is probably the most serious drawback to the standard since terminal control such as screen editors and data entry is required by many commercial applications," says David Buck, the /usr/group systems interface subcommittee chairman and president of D. L. Buck and Associates (San Jose, Calif). System V's IOCTL, however, is the generally agreed-upon starting point for all the future enhancements to the standard's terminal control capabilities.

Some other omissions include machine-dependent routines and

It wasn't long ago that when you wanted to back up a 5¼" Winchester, you only had two choices: floppies or QIC-compatible ¼" tape.

Well, now there's a third choice: Cipher's FloppyTape.™ It gives you the best of both worlds — without the drawbacks of either.

Economical like a floppy.

FloppyTape is the perfect secondary storage device for low-cost systems. It offers a substantially lower cost-per-megabyte than floppies. Yet the actual cost of the drive is only slightly more.

What's more, FloppyTape uses a standard floppy interface. And it works perfectly with most regular floppy controllers. So you can share your system's existing floppy controller.

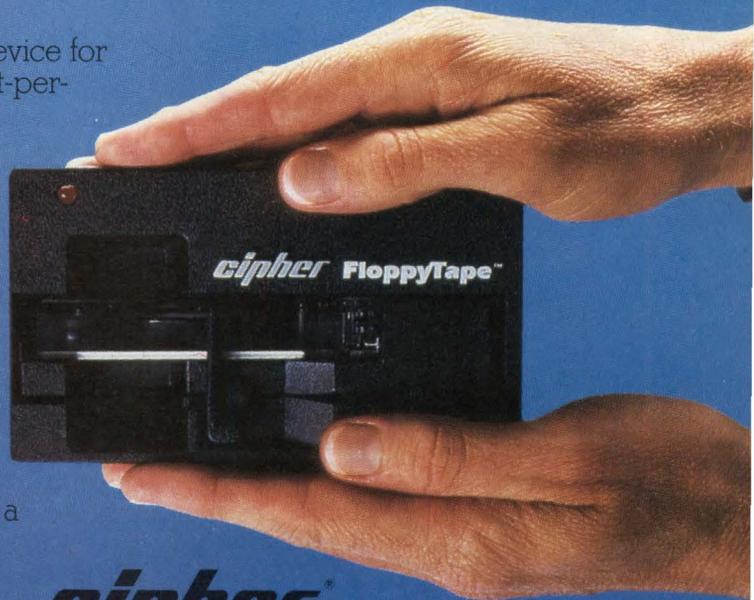
High capacity like a streamer.

Using a standard ¼" tape cartridge, FloppyTape gives you up to 32 megabytes of storage.

It eliminates the problems of handling a large number of floppy disks.

And FloppyTape is priced well below QIC-compatible ¼" streamers. With QIC, you also need a new host adapter or controller. But you don't with FloppyTape.

If the part of the secondary storage spectrum that's best for your system is the middle, 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.**

software not needed for the development of commercial applications. Games, accounting, file system mounting system calls, and, in general, super user (highest priority user) or system management functions also are omitted.

Buck's subcommittee is studying extensions that include the ioctl/termio interface, contiguous file support, directory manipulation primitives, termcaps/terminfo, and a network interface. Other extensions concerned with user programs and related interfaces are under study by a user interface subcommittee of /usr/group.

This subcommittee, chaired by Mike O'Dell, a member of the technical staff at IBM's division in Austin, Texas is concentrating on media portability, the shell command interface, command syntax, and command subsets. As an example of its efforts, this subcommittee has proposed that an extended Berkeley 4.2BSD "tar" tape format become the standard for the interchange of data files between incompatible com-

puter systems. Some details need to be ironed out, however.

The influence of System V shows up again in this subcommittee's shell command interface work. It is based on that System V's shell command description. And, for the basis of a general command syntax interface, the subcommittee is considering an internal AT&T Bell Labs document.

System calls and subroutines

Although the two major sections of the /usr/group standard relate to system calls and subroutines, the standard does not require any of the error messages that it lists be returned from a system call. Different operating system implementations may interpret an error condition simply as an extension (any commercially useful function or facility not described in AT&T documentation). If the error is an intrinsic error in a particular implementation, however, then the error code that is called out should be returned. The system calls in the standard have addressed a variety of practical problems. One of them is a

lack of the file locking capability in some Unix versions. The LOCKF call is an extension to System III and provides only one kind of file lock. Even though the /usr/group considered both read and write locks, a consensus held that applications programs could use the single-purpose LOCKF call properly. A program conforming to the standard must assume that both "advisory" and "enforced" forms of locking are possible; the standard and LOCKF allow both.

All system calls and subroutines listed in the standard must be implemented. An operating system that conforms to the standard does not have to implement all of the subroutines as listed, however.

Other potential problems inherent in the standard are based on System III, rather than System V. Both the /usr/group and AT&T have taken pains to identify the differences between the two operating systems. How all the groups will handle all the differences remains to be seen.

The majority of operating system interfaces defined in the standard are compatible with their System V counterparts, according to Donald J. Kretsch, a member of the technical staff, Unix System Development, AT&T Bell Labs (Summit, NJ). Similarly, all library routines in the standard are System V supported, as are the standard's header files. Indeed, the AT&T document gives some credit to the /usr/group.

But there are exceptions. For one, writing to a pipe interface is different. A pipe is a software construct that allows a program output to be the input to another program without intermediate disk files. The advantage of pipes is that processes using pipes can run in parallel, and pipe buffer mechanisms take the place of temporary file storage.

According to the /usr/group standard, partial writes are not allowed when data is written to a pipe. A write may fail, therefore, if it exceeds some implementation-dependent value. System V is more flexible, however, and allows partial writes. In System V, for example, a process writing 100 bytes to a pipe with room for only 50 bytes will be allowed to put 50 bytes

(continued on page 44)

Roots of the standard

Developed with the cooperation of AT&T, the /usr/group document resulted from more than three years of work. The committee that prepared it, chaired by Heinz Lycklama, Vice-President, Systems Development, of Interactive Systems Corporation (Santa Monica, Calif) has been working closely with the newly formed P1003 committee of the IEEE, chaired by Jim Isaak, Director of Product Planning, Charles River Data Systems, Inc (Framingham, Mass). At January's Uniform meeting in Dallas, Tex, the two groups merged and formally adopted the /usr/group document as the basis for future P1003 work.

The P1003 committee will include a technical committee representing the interests of the /usr/group. Among other chores, it will help move the standard in such uncharted areas as national language systems (for example, Unix-like systems in Japanese).

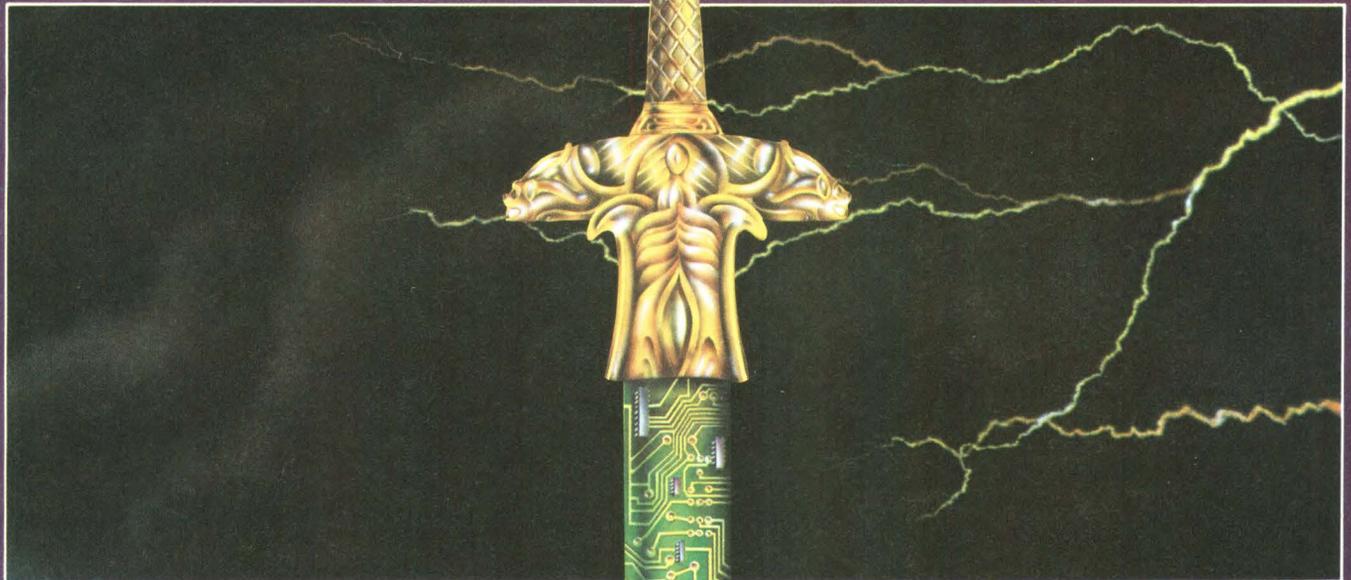
The P1003 committee is expected to put the document through its standards-making procedure in nine months. Then, it will take the document to the American National Standards Institute and, finally to the International Standards Organization for world standard consideration.

As might be expected from the standard's and Unix's heavy dependence on the C language, all concerned parties maintain liaison with ANSI's C-language standards committee X3J11. This is also done to avoid incompatibilities. For example, Don Kretsch of AT&T Bell Labs serves on both the /usr/group-P1003 and the X3J11.

The latest documents (the "/usr/group Standard" and its companion volume, the "Reader's Guide to the /usr/group Standard"), are available for \$15 from /usr/group, 4655 Old Ironside Dr, Suite 200, Santa Clara, CA 95050. AT&T's "System V Interface Definition," Spring 1985, Issue 1, Select Code 307-127, is available from the AT&T Customer Information Center, 2833 North Franklin Rd, Indianapolis, IN 46219 (call 800/432-6600, operator 77).

New SKY WARRIOR

15 MFLOP Array Processor cuts through \$1,000/MFLOP barrier!



At \$14,900*, it's the most economical way yet to put mainframe number crunching into your microcomputer.



The 15 MFLOP SKY WARRIOR isn't simply a minor technological advance. It's nearly a revolution, fully twice as cost effective as anything else on the market. Capable of performing 15 million 32-bit floating point operations per second. And 1/2 million operations per second on 64-bit data.

Use it for 2-D Fast Fourier Transforms in imaging systems, in on-site seismic analyses, in 3-D graphics workstation transformations, in real time laboratory signal processing systems, or anywhere else requiring fast interaction with lots of data.

We're experienced veterans at porting hardware. We'll adapt SKY WARRIOR to your custom bus the

same way we've adapted our other SKY products to many custom buses already. SKY WARRIOR is compatible with the VME bus, and will soon be compatible with other standard bus structures, too.

SKY WARRIOR's tightly coupled architecture gives you speed as well as easy access to maximum system memory. And remains invisible to your end user in the process.

This easy-to-use, plug-in processor comes fully software supported by a rich library of user-callable subroutines.

All software developed for our popular, field proven Micro Number Crunchers will run on SKY WARRIOR, too.

You don't have to worry about delivery, performance or backup when you deal with SKY. We specialize in floating point products. Our customer list reads like a scientific and engineering Who's Who. Ask and we'll give you references. They'll tell you what they tell us, things like *rock solid engineering, instant responsiveness, super technical support, and total product reliability.*

For detailed technical information on our exciting new SKY WARRIOR or any of the other products in our LEGION of arithmetic processors, contact SKY Computers, Inc., Foot of John Street, Lowell, MA 01852, telephone (617) 454-6200.



*Quantity 1. OEM discounts available.

THE ONE &

UNIX[™] SYSTEM V. FROM AT&T. THE

There's only one UNIX System V. The one created and supported by AT&T. The one fast becoming a universal standard.

It's another reason why good business decisions are based on UNIX System V.

In more ways than one, it pays to make certain that the products you offer are based on UNIX System V. It is the stable foundation on which all of our future enhancements will be developed. It's here to stay because we pledge that our future software enhancements will be compatible with it.

Configure your systems now with products based on UNIX System V and you'll be at the forefront of its establishment as a worldwide standard. AT&T has expanded its role in markets in Europe, Japan, and the Far East to create opportunities for machine-independent, portable solutions to operating system needs. Around the world you can profit from our initiatives when you invest in UNIX System V solutions.

On the inside track

If you develop products based on UNIX System V, you'll be in

on the latest advances from the custodians of the UNIX System standard. A growing number of microprocessor ports to choose from. Expanded networking capabilities. A growing number of applications packages from software vendors. And, the power of the system that spans generations of machines as well as models and sizes—from micros to mainframes.

Service from the source

Along with its various technical and commercial advantages, UNIX System V offers something the other systems can't match—

SYSTEM V

ONLY

STANDARD BEARER.

the full backing of AT&T.

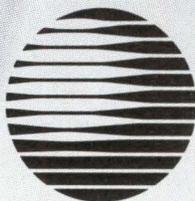
That includes worldwide support and training programs. Our commitment to work with computer manufacturers and resellers as well as user groups to develop standard products for the future. And for software vendors writing applications programs, a complete service and support program from AT&T.

The one and only from AT&T is UNIX System V—backed by AT&T service and support—designed to make UNIX System V the one for one and all.

For more information on UNIX System V from AT&T,

send in the coupon.

**UNIX System V. From AT&T.
From now on, consider it standard.**



AT&T

Please send me "UNIX Software." CD-0300-FF

Mail to: AT&T, Box 967, Madison Square Station
New York, New York 10159

Name _____

Title _____

Department _____

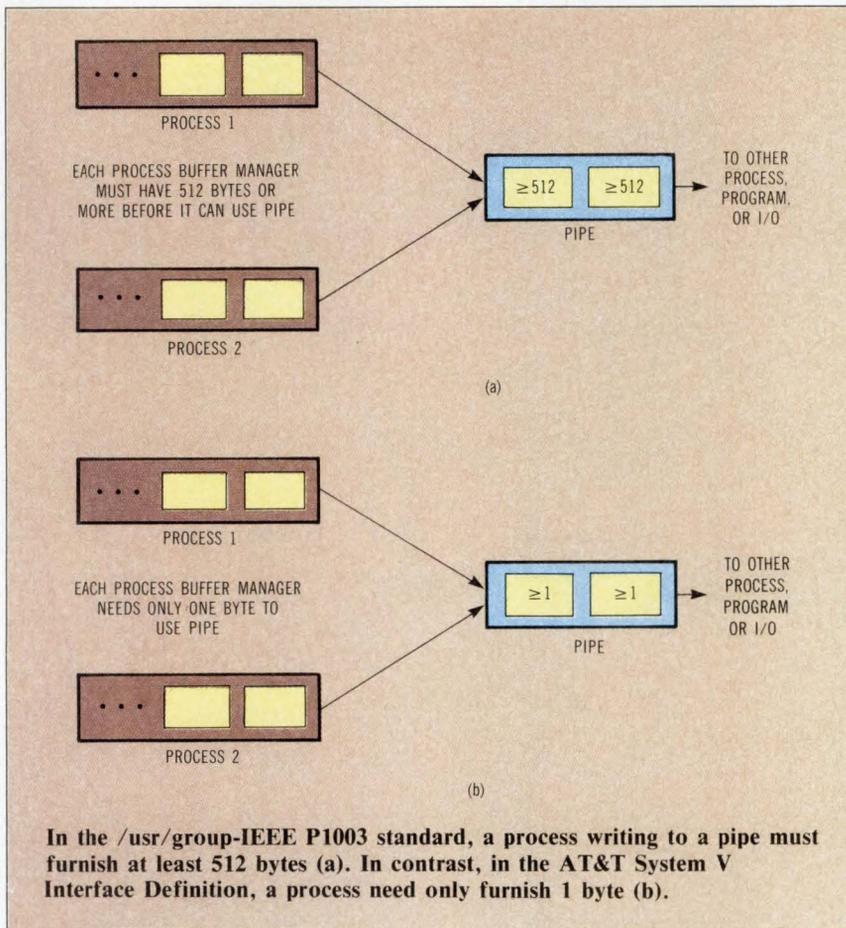
Company _____

Address _____

City _____ State _____ Zip _____

Phone _____

UNIX™ System Licensee Yes No
 Don't Know



(continued from page 40)
 in. The contents of the pipe are then put "to sleep" and are "awakened" when the space is available.

The /usr/group and AT&T standards definition documents also differ in the way multiple processes writing to a pipe are handled. In the former, 512 bytes or more must be transmitted by a process through a pipe before transmission integrity can be guaranteed. (Transmission integrity requires that all bytes are contiguous and are from a specific process.) This transmission scheme differs from that allowed by AT&T.

The AT&T standard allows two processes to alternate, sending one byte at a time. According to AT&T, many customers require the byte-at-a-time approach. While this procedure is not allowed in the /usr/group standard, the /usr/group's procedure is allowed by the AT&T document. As the two standards definition documents become one, such difference will be reconciled.

While differences in the approach to writing to pipes may be significant to some, other observers maintain that the /usr/group and AT&T documents are still substantively the same. For example, Matt Perez, vice president of engineering at Contel Codata Systems Corp (Sunnyvale, Calif) claims that the differences involve only three return values and five error values. And, he believes, these differences will be ironed out with time. He is more concerned with extensions, such as realtime software constructs and multiple national languages, where there is little agreement.

According to AT&T Bell Labs technical staff members Tom Houghton and Donald J. Kretsch, AT&T wants to merge the two standards. They are quick to point out, however, that AT&T must keep its customers' interests in mind—as in the way its document defines writing to pipes.

Whether or not the documents merge, the System V Interface Definition is not a "standard," but only

a "corporate policy or direction statement." As Michael DeFazio, director of software systems project management, AT&T Technology Systems (Summit, NJ), puts it, "Supersets of the document are encouraged and any extended product meeting the document will run System V interface document products." This means an embedded base of System V software can be preserved even as future compatible operating systems or applications software to run on compatible operating systems are enhanced.

Record locking

An ongoing technical problem with certain Unix implementations is control of contention between multiple concurrent processes involving shared data files. That is, without concurrency controls, data loss may result from concurrent reads and writes to file-based data. Many Unix systems also lack the semaphores or other synchronization signals for processes to create critical file regions.

A solution to handling contention has been proposed for incorporation into the /usr/group standard. It uses the LOCKF system call in a program for both record locking within a file and for the generation of semaphores. In the former, LOCKF allows records to be locked for exclusive use in a file between cooperating processes; and it denies access to other readers/writers of the same record.

In the latter, semaphores may be produced by locking and unlocking any offset in an open file descriptor for a regular file or certain optional files. The application software designer makes the decision on how to map resources into offsets for locking.

According to the developer of the locking procedure, John Bass (Cupertino, Calif), the LOCKF procedure that /usr/group has recommended for adoption defines both a user interface and the way to implement the LOCKF procedure. The implementation procedure, however, is not without its problems. For example, there are possible deadlock or deadly embrace situations (two or more processes attempt to access the same resource and neither can finish the acquisition). There also are problems with locking and unlocking file

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.

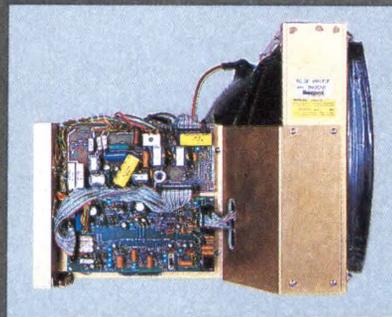
Keeping a low profile

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

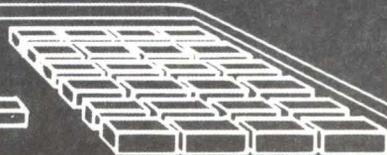
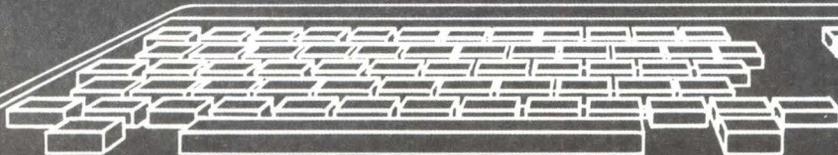
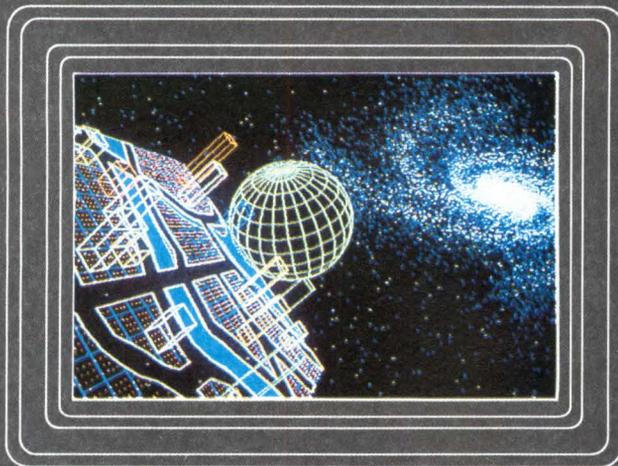


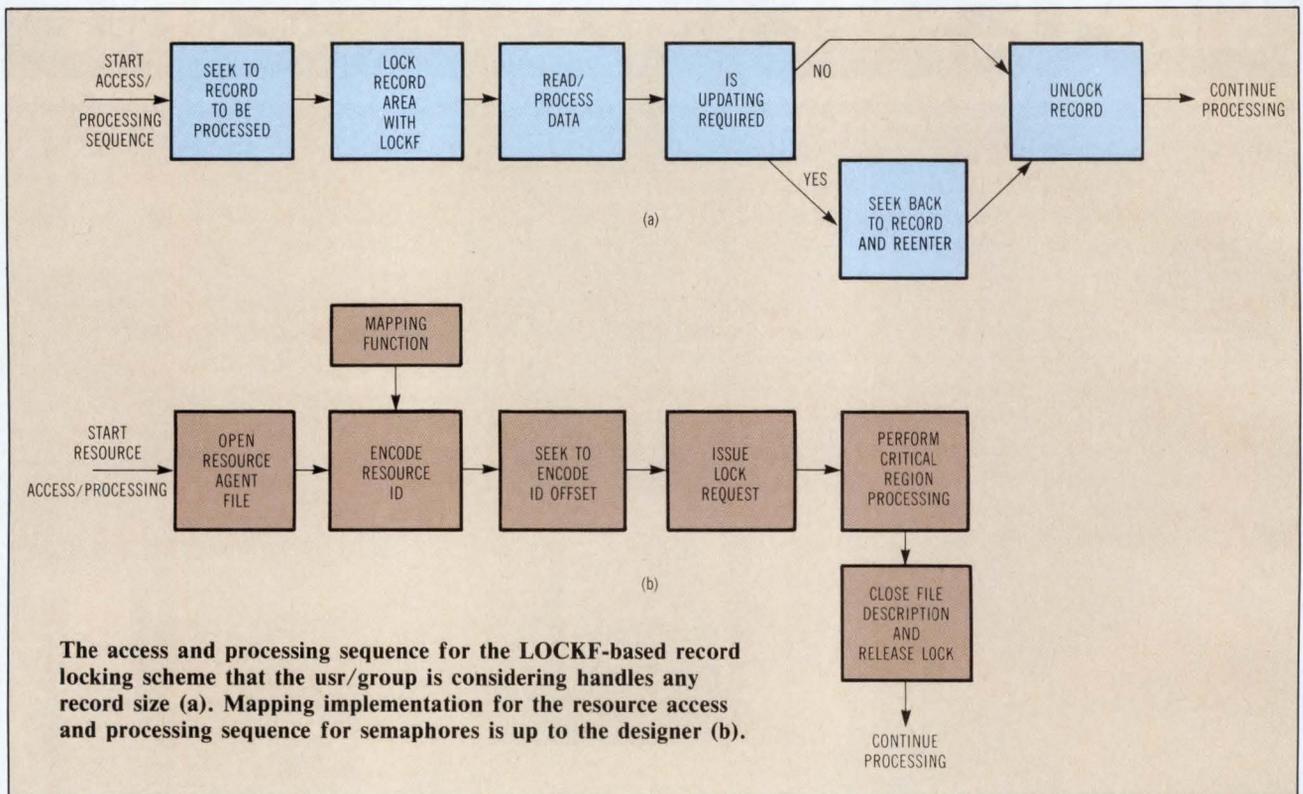
put a drain on your system's resources. In fact, Ikegami will make your profit picture look even better, because it's a finished product with on-time delivery.

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 23





areas that overlap, resource or database integrity during lock operations, and data buffering. Objections also concern trade-offs between advisory locks and enforcement locks, overhead and performance, and lack of the "read shared lock" concept.

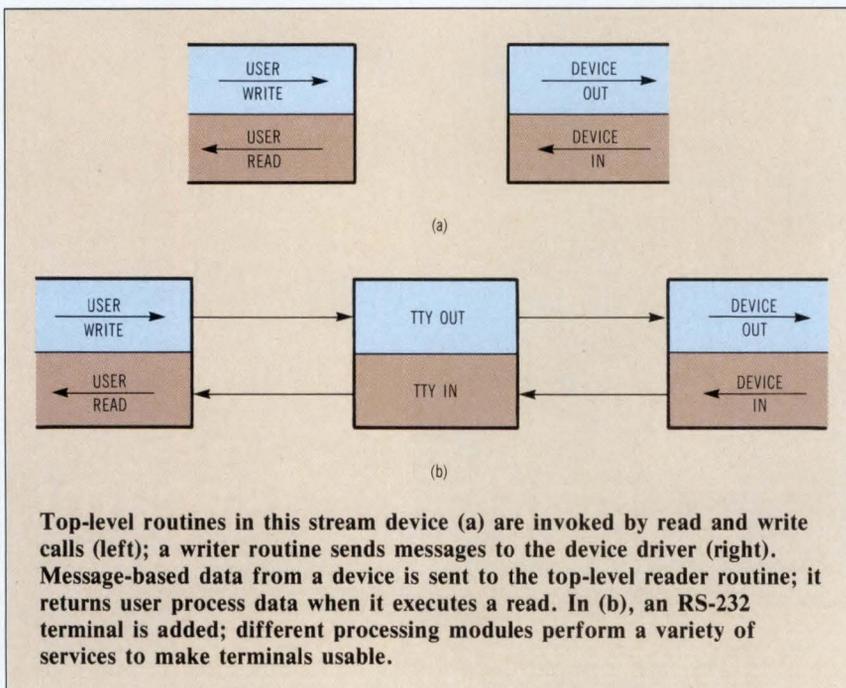
Developing standard system calls and subroutines for a Unix-like operating system that provides a standard interface for the C-language programmer has ramifications beyond easing application software development. At least two areas where the standard will

have an impact are local network standards and graphics standards.

The usr/group standard accepted by the IEEE P1003 committee does not specifically address how to do networking with Unix-like or Unix operating systems. But the 4.2BSD Unix implementation has the U.S. DoD Transmission Control Protocol/Internet Protocol for networking built in. Functionally equivalent in large part to layers three and four (network and transport) of the International Standards Organization's Open Systems Interconnect model for computer communications (see "Data Comm Networks Finally Come of Age," *Computer Design*, Dec 1984, page 169), the built-in TCP/IP is one of the major reasons for 4.2's popularity.

According to Thomas W. Doepner, Jr, Associate professor at Brown University (Providence, RI), other key features making it popular with engineering and scientific users are its support of network-based server processes and its distributed file system.

But TCP/IP's days could be numbered. In a still unpublished report to the DoD, a national committee of consultants under the National



Decoupling IV

AVX
Techfile Series

TECHNICAL INFORMATION FROM THE LEADER IN MLCs

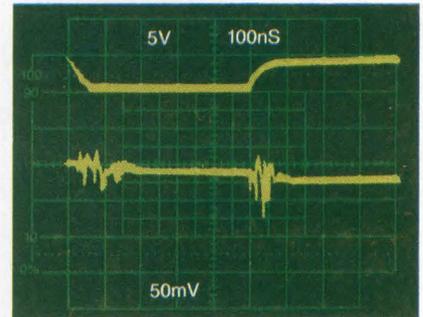
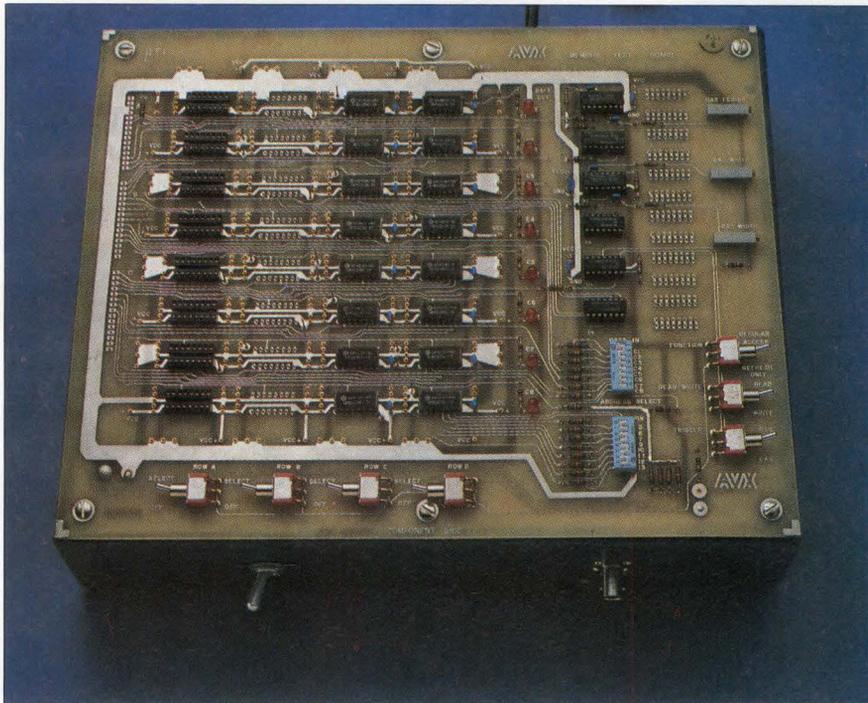


Fig. 2. Scope traces for refresh cycle on 256-K DRAM test board with 0.33- μ fd AVX MLCs.

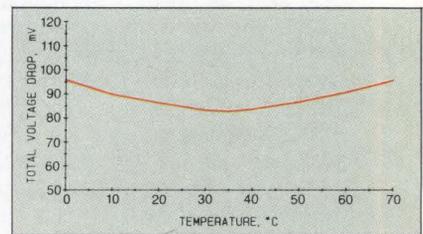


Fig. 3. Effect of temperature on 256-K DRAM decoupling with 0.33- μ fd AVX MLCs (Z5U).

MLC Decoupling of 256-K Dynamic RAMs

A dynamic RAM's sensitivity to decoupling-induced "soft-errors" (random loss of one or more bits of memory) increases dramatically with higher speeds, higher density, and an increased number of sense amplifiers. The new 256-K DRAM designs have large, instantaneous current demands which must be satisfied by a local current source.

That source is the decoupling capacitor directly adjacent to the RAM package. And the capacitor most often used for this application is a multilayer ceramic capacitor (MLC) because of its low series inductance, low series resistance, and high capacitance in a small size.

Test Results

Tests were conducted by AVX on a 256-K DRAM memory board to determine the noise level obtained with various values of MLC capacitors. Figure 1 compares the results obtained using 256-K DRAMs with those from similar board tests on 64-K DRAMs. As indicated, 0.33- μ fd capacitors are required on the 256-K DRAM board to obtain a noise level equivalent to that obtained using 0.1- μ fd capacitors on the 64-K DRAM board. Performance improvements on the 256-K DRAM

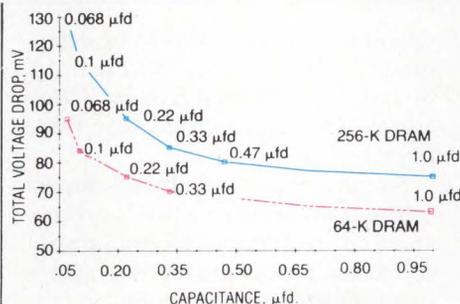


Fig. 1. Decoupling characteristics for 64-K and 256-K DRAMs with AVX MLC capacitors (including V-bump and V-droop).

test board leveled off between 0.33- μ fd and 1.0- μ fd, indicating that the preferred value for decoupling is about 0.33- μ fd.

Figure 2 shows the scope traces obtained during refresh cycle on the 256-K DRAM test board with a 0.33- μ fd AVX MLC. In all tests, the general decoupling scheme used was one MLC capacitor for each DRAM, with no board-level bulk capacitors.

Discussion

General-application ceramic formulations, such as Z5U, show considerable change in capacitance with temperature. However, this change has

little affect on the total noise level for 256-K DRAM when the correct value is chosen. Thus, the 0.33- μ fd value is high enough to meet the 256-K DRAM's current requirements over its full operating temperature range, as shown in Fig. 3.

For a complete technical paper describing these tests in detail, complete and return the coupon below.

- Please send me the AVX Technical Paper, "Decoupling 256-K DRAMs."
- Please send me literature describing AVX MLCs
- Please send me samples

Name _____

Title _____

Company _____

Address _____

City _____

State _____ Zip _____

Phone _____

Send to: AVX Corporation, Dept. 25,
PO Box 867, Myrtle Beach, S.C. 29577

CD3/1/85

AVX Technology
For The Times

AVX and AVX TechFile are trademarks of AVX Corporation.
© 1984 AVX Corporation.

Interface standard—not operating system standard

The /usr/group has been working closely with AT&T and has used AT&T internal documents in generating its standards proposals. But this has created something of a dilemma. How can the /usr/group, or anyone else, "standardize" a proprietary AT&T product? The answer is that it cannot.

The /usr/group document specifies a computer operating system functionally compatible with Unix to benefit both producers of, and users of Unix software and Unix-using hardware. With cooperation from AT&T (both in committee representation and in the use of AT&T documents), the /usr/group is careful to differentiate Unix from its operating system standard, and clearly point out where AT&T and /usr/group materials are used in its documents. In other words, the /usr/group standard is not a Unix system standard; it defines interfaces between application software and operating systems that provide features or behavior similar to Unix.

These are not trivial matters. There are questions of copyright law, trademarks, and antitrust to worry about. For example, concerned about the latter, /usr/group has invited all parties interested in Unix to participate in the standard's development, and will have nothing to say about whether products conform or not. But then how does an interested party know if an operating system implementation conforms to the standard? The answer, as always, is a software validation suite. And the opportunities for developing such a suite are under study. Jeff Schriebman, president of Unisoft Systems (Berkeley, Calif), is chairman of the /usr/group subcommittee charged with investigating a test suite for the /usr/group standard.

In the meantime, AT&T announced at Uniforum (Dallas, Tex) that it had given a contract to Unisoft Systems Inc. of Berkeley, Calif to come up with a test suite to allow interested parties to show compliance with the System V Interface Definition.

Research Council of the National Academy of Sciences has recommended that the DoD convert to the ISO protocols for future work—and DoD has concurred. Those "numbered days" could stretch out, however. Digital Equipment Corp's Ultrix operating system, for example, has TCP/IP built in, notes Tony Lauck, DEC's Manager of Networks, Communication Architecture and Advanced Development. It will, no doubt, be DEC-supported for a long time to come.

As insurance against "numbered days," developers at Berkeley are already said to be working on a version of Unix that will have ISO network communication protocols built in.

The software mechanism that Berkeley 4.2 uses to tie networks into its operating system commands is known as "sockets." Whether sockets or the "streams" mechanism defined by Dennis Richie, the inventor of AT&T's Unix, is more suitable as an implementation technique for connecting networks to Unix or Unix-like

operating systems needs to be determined. (See Dennis Richie, "A Stream Input-Output System," *Bell System Technical Journal*, Nov 1984, page 1897.)

Streams are essentially full-duplex connections between a user's process and a device. They consist of linearly connected processing modules (there is no multiplexing or fanin or fan-out). Streams are analogous to shell pipelines, except that data flows in both possible directions.

Stream modules communicate by message passing. Processing modules in a stream may be considered analogous to stacks whose "top" is "next to" the user program. Not yet available outside AT&T Bell Labs, the streams software mechanism is now running on some 20 machines in the Information Sciences Research Div of the labs, according to Richie.

To start a discussion of how to structure Unix and Unix-like operating system extensions for networking, AT&T has submitted a working paper on networking to the /usr/group. Prepared by David Orlander of

AT&T Bell Labs, the working paper addresses networking strategies rather than specific interfaces. In short, the paper calls for networking to Unix to be based on the concept of the ISO OSI model. It also calls for freedom from implementation and protocol dependence. In addition, the network services extension outlined in AT&T's Select Code 307-127 document calls for a future standard C-language library. It will connect to the OSI transport and applications layers for specific network and protocol independence.

According to AT&T's Houghton and Kretsch, more AT&T position papers are in the offing. In addition, IBM is submitting proposals to the /usr/group through Mike O'Dell, its member on the committee.

Drawing pictures

The AT&T Interface Definition document and the /usr/group documents have little to say about graphics extensions to Unix. The AT&T document merely states that the Graphical Kernel System standard is being considered (see "Graphics Standards Finally Start to Sort Themselves Out," *Computer Design*, May 1984, page 167).

Many see incorporating an already existing graphics standard as a more legitimate procedure than developing a specific graphics package for Unix. Yet, others believe that a Unix-dedicated graphics package would best serve those members of the Unix community who have special needs. Currently, graphics software vendors supply packages written in C or Fortran. And they apparently have no problem writing their packages so they can be called from the different Unix versions on the market.

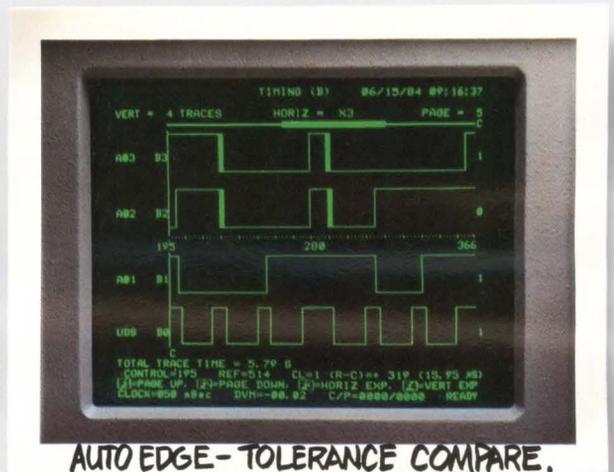
Graphics software vendors are but one group that must deliver products while the evolution of Unix standards continues. As the graphics software vendors are doing, others vendors will address the Unix standards problem by opting for a specific Unix or Unix look-alike, or by offering a choice of several Unixes.

—Harvey J. Hindin,
Special Features Editor
SYSTEM TECHNOLOGY
(continued on page 59)

**Designers'
deepest
suspicions
confirmed.**

And we have the pictures to prove it.

© 1985 Gould Inc.



AUTO EDGE-TOLERANCE COMPARE.

Take a quick look through these pages. Notice anything unusual? Anything you might not have suspected appearing on these pages of Gould advertising?

If you said something like, "Wow, I didn't know Gould made that many logic analyzers," you've got sharp eyes.

But if you figured out, that in each case, the Gould analyzer is helping a senior design engineer slug it out with the toughest, nastiest, most frustrating logic problems, you really know your stuff.

Because that's what makes a Gould a Gould.

Consider, for example, the K105-D



ONE-STEP NOISE MARGIN ANALYSIS.

microprocessor analyzer. It makes short work of debugging even long programs. Because it gives you the most sophisticated triggering in its class. Which means you can capture the signals you need to capture. Go with the flow, if you will. And isolate errors in a lot less time, with a lot less frustration.

Here's another reason there's nothing like a Gould. Intermittent failures in discrete logic. Before you pull your hair out, pull out the K205-D. Among other life-savers, it has an Auto Save mode that lets you set up and go home. The analyzer will track and record anything out of the ordinary. Compare it to a

pattern you design (so you don't waste time on unwanted compares). Even add a time stamp. All at very high speed.

And while we're on the subject of speed, consider the K500. It's the analyzer high-speed mainframe and super-mini designers need. Because it has the clockrate (500MHz) and active probe bandwidth to guarantee signal fidelity. What's more, you get individual high frequency probes so you can stretch across large ECL boards and get to the signals you want to analyze.

Finally, there's our behind-the-scenes hero: the T-12 communicator. It lets field service organizations keep their best brains close to

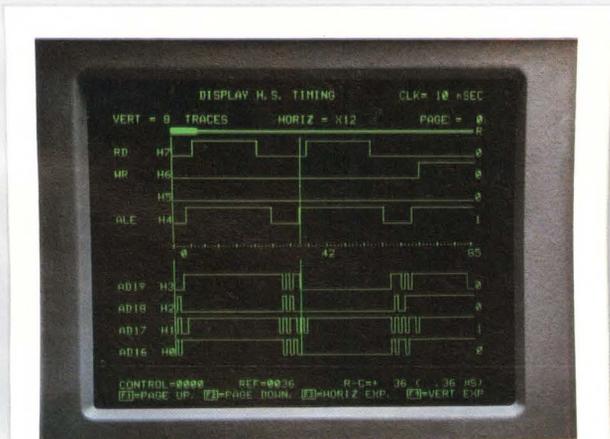


8-AND 16-BIT μ P DISASSEMBLY.

home. With a T-12 and a Gould logic analyzer (in particular the K100-D), at each end of a standard phone line, remote diagnostics are fast, easy and economical.

If you're suspicious by nature, find out why so many designers won't settle for a less competent analyzer. Call your Gould representative today. Get all the latest news about the full line of Gould logic analyzers. Along with some very revealing pictures.

Start by calling (800) 538-9320. In California call (800) 662-9231 or (408) 988-6800. Or write Gould Inc., Design & Test Systems Division, 4650 Old Ironsides Drive, Santa Clara, CA 95054-1279.



5 ns GLITCH CAPTURE.

Engineer disappears without a trace.

Tragic. A fine, upstanding young engineer gone. No more project. No more work. All because he used a logic analyzer that just couldn't cut it.

So rather than let a bug come back to haunt the program, the senior engineer took out his trusty K105-D and took over. End of bug. End of story. Beginning of a beautiful software/hardware relationship.

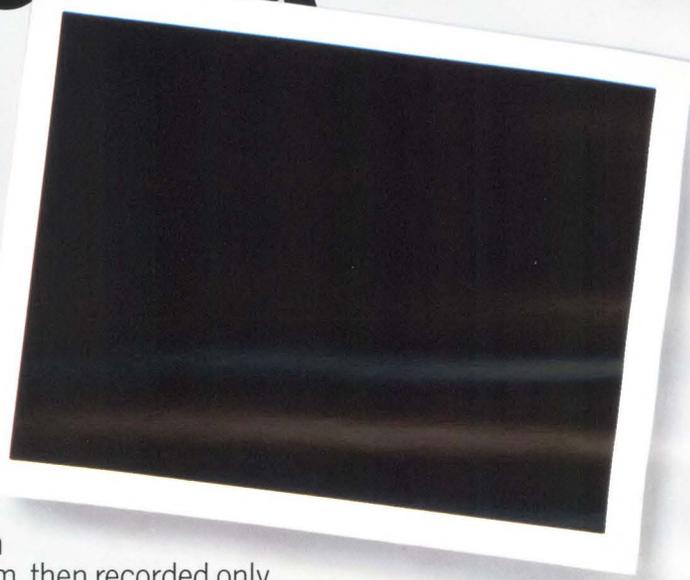
Of course, we've left a few details out of the above plot summary.

For instance, there was the Trace Control™ incident. It let the problem-solver find out exactly where the bug was. And why it happened.

Trace Control is the most sophisticated triggering scheme around. It let our hero set more triggering conditions and walk

away. The K105-D followed every turn and loop in the program, then recorded only what needed to be seen. It captured even widely separated slices of code, ignoring what was right and zooming in on what was wrong. And saved an astonishing amount of time.

Of course, the solution to the language problem also helped bring about the bug's demise. The K105-D speaks assembly.



So triggering conditions are set in mnemonics. Not complex object code. Special disassembly modules for many popular microprocessors including the 68000, 8086, 8080, 8085A, and Z80® meant the K105-D was the only logic analyzer needed. Especially when its 100 MHz high speed link was used to see how timing affected software execution.

If you'd like to rid yourself of the toughest, most pesky bugs software has to offer, call your Gould representative today. He'll appear with the whole story.



Designer saved by faults alarm.

Trace Control™ is a trademark of Gould Inc. ©1985 Gould Inc.

It was a dark and stormy night. A random, intermittent glitch was wreaking havoc with the top designer's hardware team. With no time left for sympathy, the top designer took out his K205-D and sent it bug hunting.

And in the middle of another dark night, the fault was located and captured. And held until the top designer came to work the next sunny morning.

The cure started with the K205's no-compromise, no-mercy attitude toward faults. It ran at 100MHz across all 48 channels. So there was no double-probing. It used its glitch mode on all 48 channels without requiring probes to be reconnected. It used Trace Control™ to set 16 independent



triggering levels. It used its unique tolerance compare capability (which compared timing sequences to stored references) to eliminate sampling error differences inherent in asynchronous analysis.

And it used its user-definable



disassembler (UDD, if you need to catch your breath) to easily translate binary signal flow into simple mnemonics. Which, by the way, was connected through its user-definable interface. Without a hitch.

When the fault was found, the K205 automatically stored the data. And let the top designer get right to work solving the problem.

Of course, even if you don't have need of the most sophisticated triggering in the business you can still save yourself a lot of time and frustration by using the K205 merely for its unrivaled combination of speed and power. Either way call your Gould representative today.



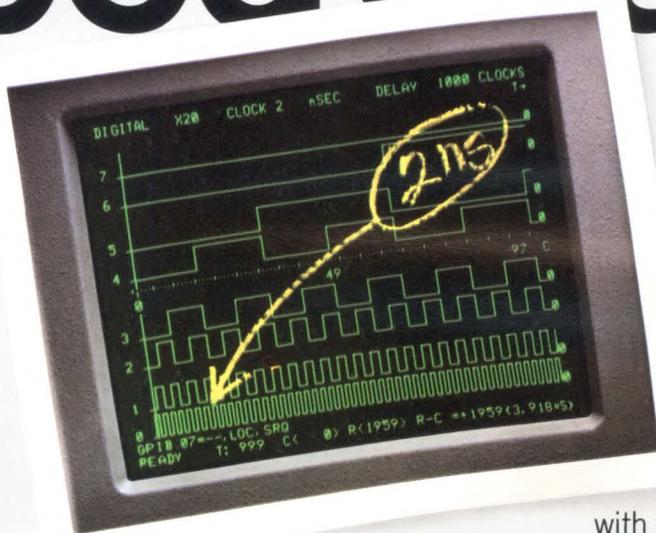
GOULD
Electronics

Glitch captured after high speed chase.

It almost got away. But because someone made the smart decision to use the Gould K500, one more glitch bit the dust.

Speed was essential. So all eight channels of the K500 were racing along at 500MHz. Resolution also had to be tight enough to narrow down the escape routes. Two nanoseconds brought things under control.

And, of course, a lot of territory had to be covered. So the K500's individual hybrid probes made the job easy. Their 900MHz frequency response also assured accurate signal capture and unmatched signal



clarity. So the innocent were left alone. And the guilty glitch was discovered.

Sounds simple, but a lot of power was directed at this particular chase. The K500's Search



Mode was used to examine data in the vicinity of particular loop addresses. Random glitches (some guilty, some not) were found and identified

with the Compare Mode — a very fast method for defining a pattern and checking data fields against it. Finally, some very flexible recording and display talents were called in to support the complex timing and state analyses which led to the ultimate capture.

If you'd like a complete run-down that will help you in your own high speed chase, call your Gould representative today. Before another glitch gets away.



GOULD
Electronics

Higher intelligence now a remote possibility.

Long distance is now better than being there. So long as your field service organization is using the Gould T-12 Communicator.

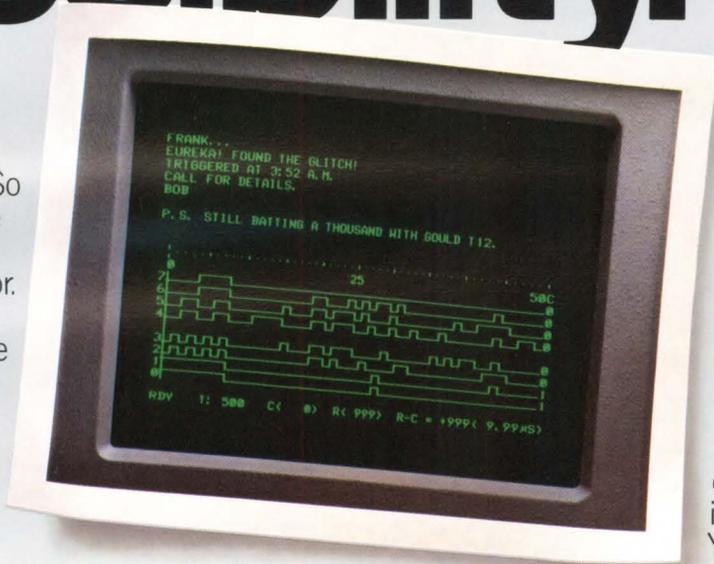
Here's how you can keep your top brains close to home. First, send a technician to your customer's site armed with a T-12 and a Gould K100-D (or other Gould Logic Analyzer).

Next, use an ordinary voice-grade telephone line to send data back to the T-12 and K100-D at your service facility.

Finally, let your top technical specialists solve those really tough problems from the comfort of their offices.

That's the how. Now here's the why. Reduced mean down time (MDT) by 40%. Slashed travel budgets by 50%. Increased

productivity in the home office. Fewer board swaps. Less run-around-related burn-out and talent loss. Happy, satisfied customers.



The T-12/K100-D combination also provides convenient baby-sitting, recording patiently while everyone sleeps. In fact, if you're searching for intermittent, random failures, this team is perfect. The T-12 has its own error-checking electronics so bits aren't lost. You can even set up a monitor situation that

auto dials your central service facility when the fault occurs. Leaves an on-screen message on your analyzer. And transfers all the recorded data on command.

To make this remote possibility come true, call your Gould representative today. You'll see that the T-12 provides an extraordinary way to communicate with a higher intelligence.



GOULD
Electronics

Gould... Innovation and Quality in Logic Analyzers

Confirm your suspicions.

Get the full story with some never-before-seen pictures. Call your Gould representative today or write Gould Inc., Design & Test Systems Division, 4650 Old Ironsides Drive, Santa Clara, CA 95054-1279.

(800) 538-9320

(800) 662-9231

(In California)



GOULD
Electronics

Speech recognition produces natural interface

Dramatic reductions in the cost of memory have made large vocabularies feasible for small computer systems, especially for personal computers. This is opening the way for spoken-command input in many applications. Perhaps the most noticeable result is the movement of speech recognition techniques from the factory floor into the front office. The effects of that move are likely to become even more pronounced when the voice-activated typewriter becomes fully developed.

Speech is natural. This presents one of speech recognition's greatest attributes. Whether a person "talks" to a system via a keyboard, a touch screen, or a mouse, the process is still mechanical. It requires a special action by the person inputting the information. It is not "natural" input.

Speech input is not intimidating—and it is universal. It does not need extra peripherals because almost everyone walks around with it.

Achieving the interface between natural traits of speech and the computer has not been easy, however. Actual developments in speech recog-

niton have come farther than industry experts originally expected, according to John McAfee, vice president of research and development for Logical Business Machines of Sunnyvale, Calif. "Advances have surpassed earlier predictions by an order of magnitude. There have been both significant industry breakthroughs and significant drops in cost." Yet, he adds, "the technology is still 10 years away from producing a voice activated typewriter."

Ten years may be somewhat of an exaggeration, considering the recent activities of at least two companies. Both Kurzweil Computer Products of Cambridge, Mass and IBM's Thomas J. Watson Research Center of Yorktown Heights, NY have "announced" success in producing voice-activated typewriters.

Both machines now have vocabularies of approximately 10,000 words and very good probability of recognition (low error rates), according to Stuart J. Lipoff, a senior member of the Electronic Design and Development Unit of the Cambridge, Mass research firm Arthur D. Little. "Their

vocabularies are still far below the 50,000 to 100,000 words needed for universal acceptance," Lipoff says, "but their \$10,000 to \$20,000 price range makes them very real."

Very few applications of speech recognition need anywhere near the 50,000- to 100,000-word vocabulary of the voice-activated typewriter. Most of today's viable systems recognize only 50 or perhaps 200 words. Many times, only 25 to 30 words are enough.

Much of the recent activity in speech recognition paralleled the evolution of the personal computer. "Prior to that," McAfee says, "most speech recognition was relegated to the factory floor and such applications as quality assurance and inventory control." The advent of the personal computer—coupled with dramatic increases in capability and equally dramatic decreases in cost—moved speech recognition into the front office, he adds. A corporate official today can press a switch and say to a terminal, "Show me last year's sales figures," and those figures will appear on the screen.

Definitions vary

McAfee defines four basic technologies of speech recognition: speaker-dependent and speaker-independent, each with discrete word recognition and continuous speech recognition. McAfee does not believe that any new technologies are likely to appear in the future. "The present technologies encompass just about every aspect of voice recognition that can be conceived," he maintains.

"There will be advances, however, in both algorithms and support hardware necessary to bring speaker independent and continuous speech into the marketplace," according to McAfee. "Speaker dependent technology has reached fruition," he adds. "Some such systems have vocabulary sizes of many thousands of words. Many are cheap, run on microcomputers, and are readily accessible to everyone. Advances over the next five years will occur in speaker independent technology."



Speech recognition is no longer limited to factory floor operations, according to ITT Information Systems director of product planning Douglas Kelley. Capabilities of such systems as the ITT XTRA voice communications personal computer provide a natural human/machine interface between the executive and the computer.

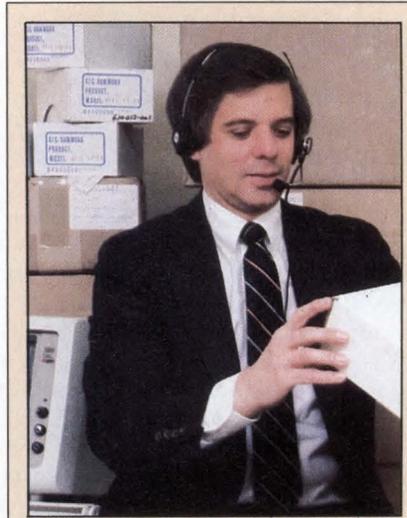
With speaker-dependent/continuous-speech recognition, there is no need to pause between words. The user speaks normally, with only ordinary breaks between phrases. A more complicated algorithm is required, however, because there are no pauses to distinguish beginnings and endings of words.

At present, few continuous speech recognition systems exist. One is offered by Verbex of Bedford, Mass, another by Votan of Fremont, Calif. More are likely to be introduced soon, however. Interstate Voice Products (Orange, Calif), for example, will introduce an as-yet-unnamed board product next month featuring "connected" speech recognition capabilities. (This is not actual continuous speech recognition, however.)

There seems to be some misuse of the word "continuous" within the industry. Some so called "continuous" systems really should be designated "connected" or "contiguous" systems. Such systems recognize phrases, but need pauses between those phrases. True "continuous" systems recognize speech input, but do not need pauses either between words or between phrases.

In speaker-dependent/discrete-word recognition, the system requires a pause of at least a tenth of a second between spoken words. The recognition system waits for the pause to determine when a word ends and another begins. This technology, also called discrete-utterance recognition, has evolved to a fairly high state.

Speaker-dependent systems are limited because the system must be trained to accept the voice of each user. Some of the less sophisticated systems can recognize only one voice and, therefore, are very limited. Others can recognize several different voices. A unit offered by ITT Information Systems (San Jose, Calif) recognizes eight, for example; but the system must be trained for each. At least one system, from Verbex, accepts input from an unlimited number of persons. In this case, circuits within an individual cartridge are trained for each user. When a user inserts a cartridge in the unit, the system itself is trained to recognize that user's voice.



The factory environment remains a major application area for voice recognition systems. Verbex vice president of marketing Kenneth A. Backer says his company's system features sealed cartridges rather than dust-susceptible floppy disks. A separate cartridge for each user provides a form of speaker independence.

Speaker-independent systems, for both discrete and continuous speech recognition, are still in their infancy. Most experts in the field say that such systems are nowhere near availability—but there are some exceptions. According to Dr Janet Baker, president of Dragon Systems (Newton, Mass), "The ability to accommodate multiple speakers means that the system must have sufficient memory to handle them and have proper coding to handle the implementation."

Baker says she does not believe the technology has leveled off. "Much has been done in the past few years—but tremendous things remain to be done," she says. "There is still room for growth and that growth will be rapid."

The immediate benefit of speech recognition, according to Baker, is in business applications. "Most standard business software requires complex control sequences that must be input by the user," she says. "Many people find this difficult. Programmable function keys save much time, but still the user must remember what those keys stand for—and must be

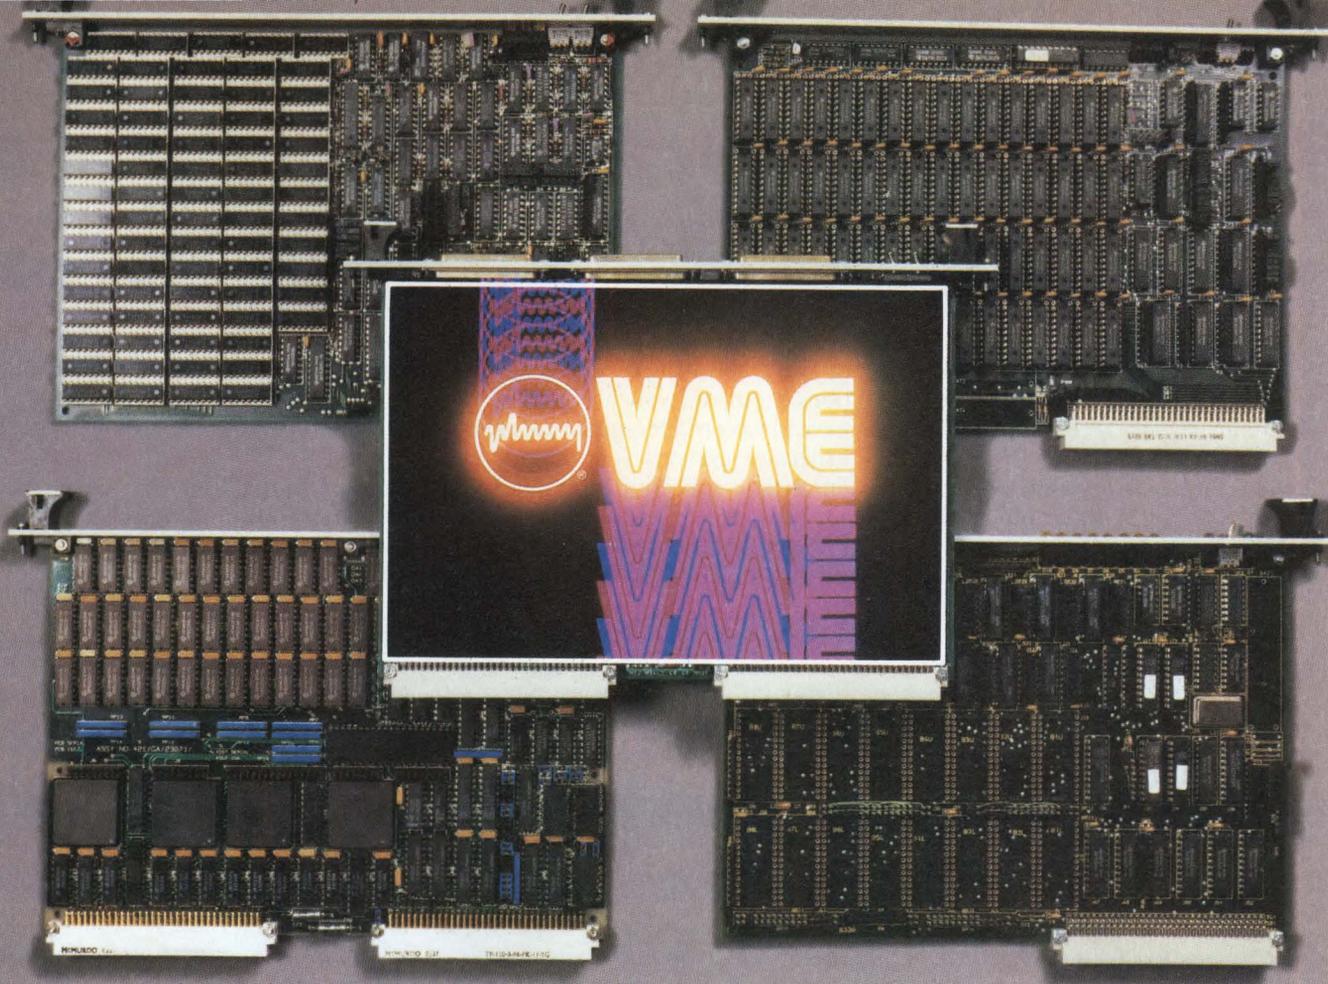
certain to press the proper ones for desired actions." However, she adds, "speech provides a very large set of 'programmable function keys' that can allow an unsophisticated user to merely speak the desired action."

The user gains

At least one major advancement has occurred in the speech recognition industry, according to Sam Viglione, president of Interstate Voice Products. Specifically, now even the uninitiated user can operate a computer system. "Coincident with the substantial drop in costs, the range of immediate applications has increased," he says. "Some of the present applications that were nonexistent a few years ago are those in which speech can replace key functions: offices and intelligent terminals, for example."

Lipoff questions the value of speaker independence for the near future, however. He maintains that very few applications need speaker independence. "Until someone identifies a really hot set of applications, there is not going to be much forward progress in developing speaker independent technology," he says. "Research is driven by market opportunities—and since there is no real need, there is relatively little research into speaker independence." Moreover, he maintains, if IBM and Kurzweil had spent as much research and money on speaker independence as they did on developing large vocabularies for their voice activated typewriters, the technology would be much farther along."

Vocabulary size need not be a dominant factor in a voice recognition system, according to Douglas Kelley, director of product planning for ITT Information Systems. "Very few applications require 200 words," he says. "In fact, that is a ridiculous outside limit unless the system is to be shared by six or eight people. Then 25 to 50 words might have to be dedicated to each of the users." Actually, adds Kelley, "It is not efficient to give people a system with that many choices. If people have more than 25 words, they tend to forget what those words are. In many applications, in fact, only six words are sufficient."



THE PEOPLE WHO GIVE VMEbus WHAT IT NEEDS MOST, GIVE YOU THE MOST VMEbus MEMORIES. PLESSEY.

Plessey Microsystems makes life easy for the VMEbus designer with the most comprehensive range of VME memories in the world. All backed by Plessey's worldwide resources, technical support services and total corporate commitment to VMEbus.

Plessey VME Memories for your systems. The Plessey VME memory range includes every popular memory technology, including DRAM, ROM/EPROM and fast static RAM. As shown in the accompanying table, our capacity range spans 128 Kbytes to 3 Mbytes with EDC and 4 Mbytes with parity DRAM. Plessey VME memories are available with full 32 bit support to take advantage of the MC68020 and NS32032. So, no matter what VME route you're taking, there's a Plessey memory to take you there better.

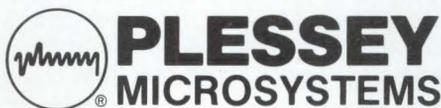
Plessey VME: a continuing commitment. Plessey VME memories are just a part of the rapidly expanding range of Plessey VMEbus board level products which also includes single board computers; controllers, graphics and I/O boards; backplanes; power supplies; software and firmware, including

languages, operating systems and drivers; and of course that famed Plessey guarantee and backup. With even more Plessey VME on the way.

For details on Plessey VME memories or our other 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.

PLESSEY VMEBUS MEMORIES						
Maximum Capacity	Plessey Memory	Parity	EDC	Data Bits	Write Access nS	Read Access nS
4 Mbyte DRAM*	PME 2EP	•		8/16/32	150	270
3 Mbyte DRAM*	PME 1EA		•	8/16/32	150	300
2 Mbyte DRAM	PME 048D	•		8/16	190	300
1 Mbyte DRAM*	PME 512EP	•		8/16/32	150	270
3/4 Mbyte DRAM	PME 256EA		•	8/16/32	150	300
512 Kbyte DRAM	PME 012D	•		8/16	190	300
512 Kbyte EPROM or ROM;	PME CRR-1†			8/16	User Adjustable	
128 Kbytes SRAM						
128 Kbytes SRAM	PME 128S†			8/16	140	140

* with optional daughter board
 † on-board battery backup



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

Name _____ Title _____
 Company _____
 Address _____
 City _____ State _____ Zip _____
 Phone _____

Kelley, however, has qualms about the future of speech recognition. It will not "replace the keyboard. Nor will it replace the mouse, although it is a very easy method of system access for the human being." He also disagrees with the importance of continuous speech recognition. "When I first got into this," he says, "I thought that to be really acceptable, a system would have to include continuous speech. Then I found that, although in experiments we could do continuous speech, there were practical problems."

For example, "If you allowed a person to talk in a normal conversation, the machine often did not recognize a word or phrase in the proper context," he explains. "'Don't erase that,' for example, might be picked up as 'erase.' With discrete speech, however, the speaker would not have to be particularly careful to choose commands that were distinctly different. 'RTT' and 'ITT,' which are very similar in sound, could be programmed and recognized with no problems."

Markets and applications for speech recognition systems differ greatly. Most manufacturers, therefore, either specialize in one particular market or design different types of systems for the different applications. Companies such as Verbex that have always made sophisticated, but expensive systems for military and industrial applications continue to do so. Most recent additions to the list of manufacturers, however, have targeted the office marketplace.

Talk to a personal computer

The personal computer has been highly influential in development of speech recognition systems for office use. Five out of the six major manufacturers of speech recognition systems—Interstate Voice Products, ITT Information Systems, Logical Business Systems, Dragon Systems, and Votan—offer boards for the IBM PC or compatibles. Verbex is the only exception.

Price, of course, is a major factor in choosing any system for the office.

Kenneth A. Backer, vice president of marketing at Verbex, says cost is also a factor for industrial systems. In addition, he says, "A marketable industrial system must be accurate; that is, it must be transparent to such factors as variations in the speaker's voice caused by a cold. Ability to function in a noisy or dirty environment, such as the factory floor, of course, remains as a major requirement.

A speech recognition algorithm developed by Verbex is aided by aspects of artificial intelligence, according to Backer. This combination enables Verbex to offer true continuous speech—the only company doing so at this time, Backer says.

Yet, Patricia S. Restaino, manager of marketing services at Votan maintains this is not so. "Votan includes speaker-dependent continuous speech recognition on both a peripheral board for the IBM PC and a stand-alone RS-232 terminal." She says it is continuous—not contiguous, but continuous—speech recognition. "A person can talk to one of our systems until he is blue in the face, and the system will recognize and accept all of the input."

Craig Bolon, manager of product development at Verbex, says his company's success in handling continuous speech "in a very robust manner" results from its training procedure. "You cannot train a system on words in isolation and expect that system to recognize them when they are strung together," Bolon says. "The spectral patterns are entirely different." The Verbex training script presents words several times—as in most other scripts, but presents them in all contexts in which they occur rather than alone, according to Bolon. The system, therefore, picks up and retains all variations."

Most industry leaders agree with Restaino's belief that "the emphasis should not be put on the voice activated typewriter. Instead, it should be on the natural aspects of using speech to control or access information." For at least the very near future, that is the way the industry is going.

—Sydney F. Shapiro,
Managing Editor

SYSTEM TECHNOLOGY

(continued on page 71)

AI techniques will make the difference

Dependence on pattern recognition techniques has stymied actual progress in speech recognition technology. Stuart J. Lipoff of Arthur D. Little, Inc, the Cambridge, Mass research firm, says "Pattern recognition could go only so far by itself. With this technique alone, a system would have to guess at identifying the spoken word." For example, "the word 'cat' when spoken could be interpreted as 'cat,' 'rat,' 'mat,' or 'hat.' The pattern recognition system would choose the most likely word—but often could be wrong."

Linguistic models based on artificial intelligence can check the words before and after "cat," and refine the choice according to the meaning of those words, Lipoff adds. The choice is based on the context of the sentence, which was ignored by pattern recognition techniques. "By marrying pattern recognition with AI, the result is significantly beyond what either one alone could do; in fact the whole is greater than the sum of the parts."

AI, in combination with the personal computer, has given speech recognition the boost needed to move it beyond the initially limited applications to unlimited possibilities. Speech recognition, once used only in military or factory applications, where the need for hands-off commands made the costs acceptable, now has entered the office environment. The voice activated typewriter is probably only the first major step toward natural speech recognition systems.

"Because computer processing power is now faster and memory is much cheaper, the added value of AI can be put to work," Lipoff emphasizes. "We can now do some of the things we have known about for some time—at reasonable cost and fast enough to be practical."





Think of it as a VW with a Porsche engine for high-performance CAD.

If you're running your CAD software on anything but a Whizzard® 1600 design terminal, you're either paying too much for your hardware — or your software's not living up to its promised performance.

Now the 1600 gives you Porsche performance on hardware that's as economical as a VW.

Performance that offers 3D manipulation right in the terminal, and cuts host computer load — which means more users, more efficiency, more productivity. A 512K memory to accommodate the largest display lists and files that are always a problem with complex designs. And you can't find 3D and 512K together anywhere else. Except at twice our price.*

The 1600s even add 16 out of a possible 4096 colors to your work for greater design depth and perspective. Or choose the more economical monochrome 1645. And get higher resolution in addition to full 3D and 512K memory. No one else has this. Except at twice our price.

Most popular application software packages for general mechanical and electrical CAD, finite element modeling, drafting, and design — including ANSYS® and DOGS™ among others — run better on the 1600. Because these packages include all the software hooks you need to take advantage of unique 1600 performance features.

And we have a full array of economical input and output devices, too. From joysticks, tablets and three different keyboards, to low, mid and high-end printers.

So give us a toll-free call today. Find out how to drive away in a Porsche, for the price of a VW.

And we'll throw in the Megateknology™ for free.

In CA, 800/824-4489

In USA, 800/854-1975

* Basic 1650 configuration is \$11,900.

 **MEGATEK
CORPORATION**
A UNITED TELECOM COMPANY
Making History out of State-of-the-Art

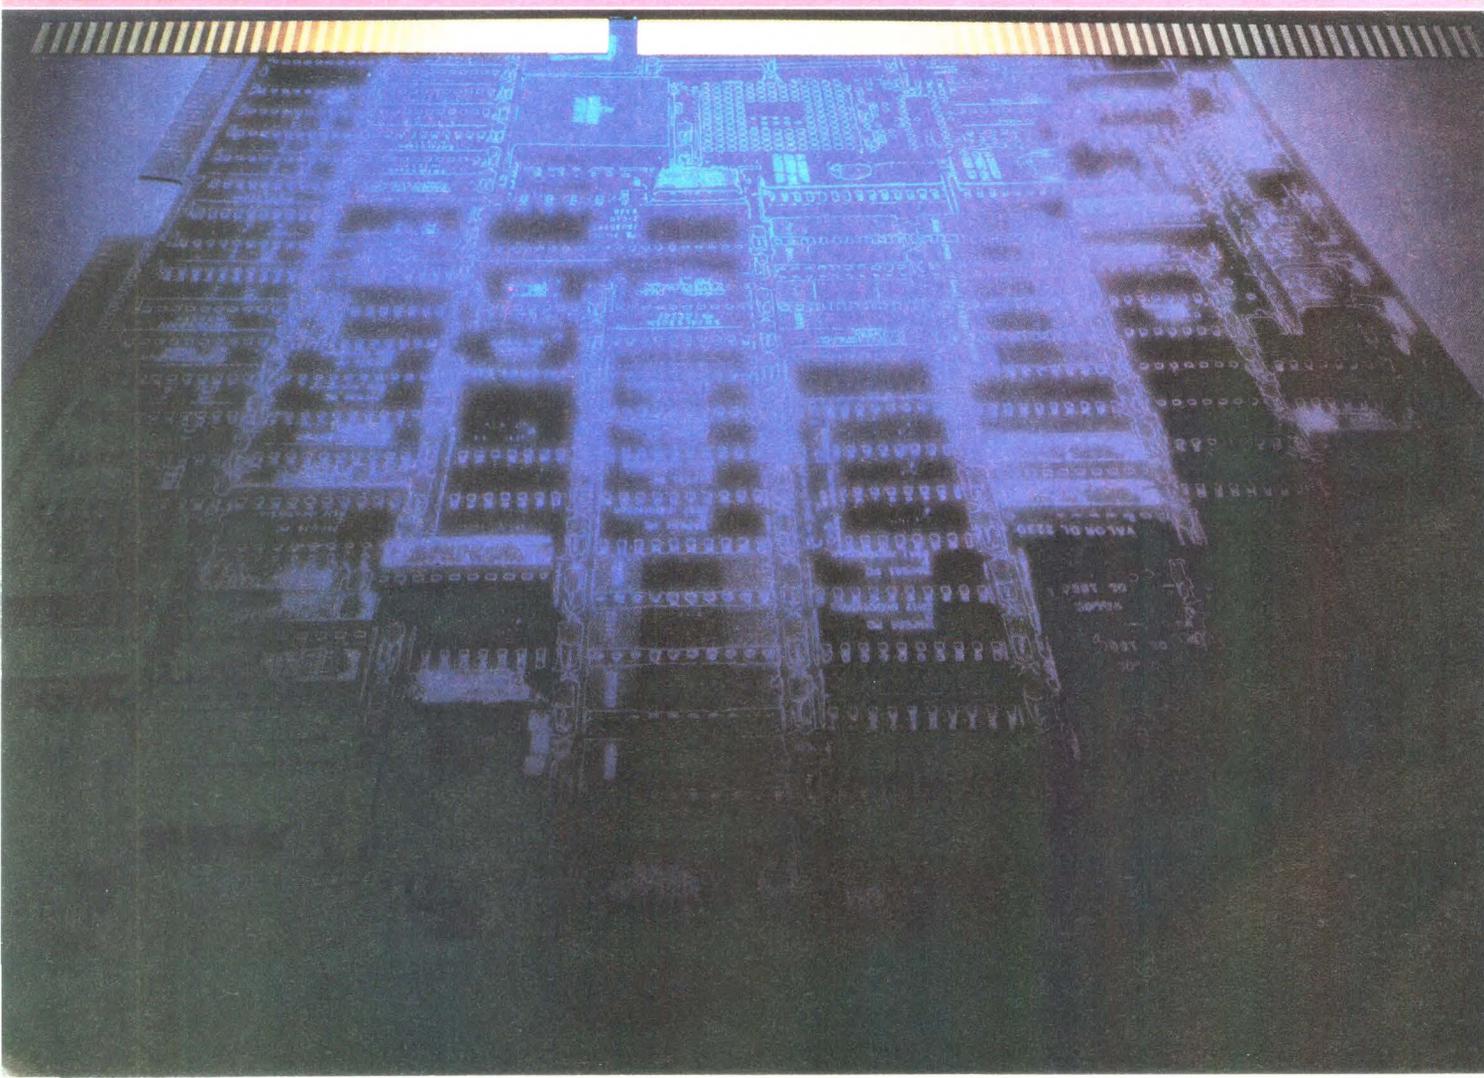
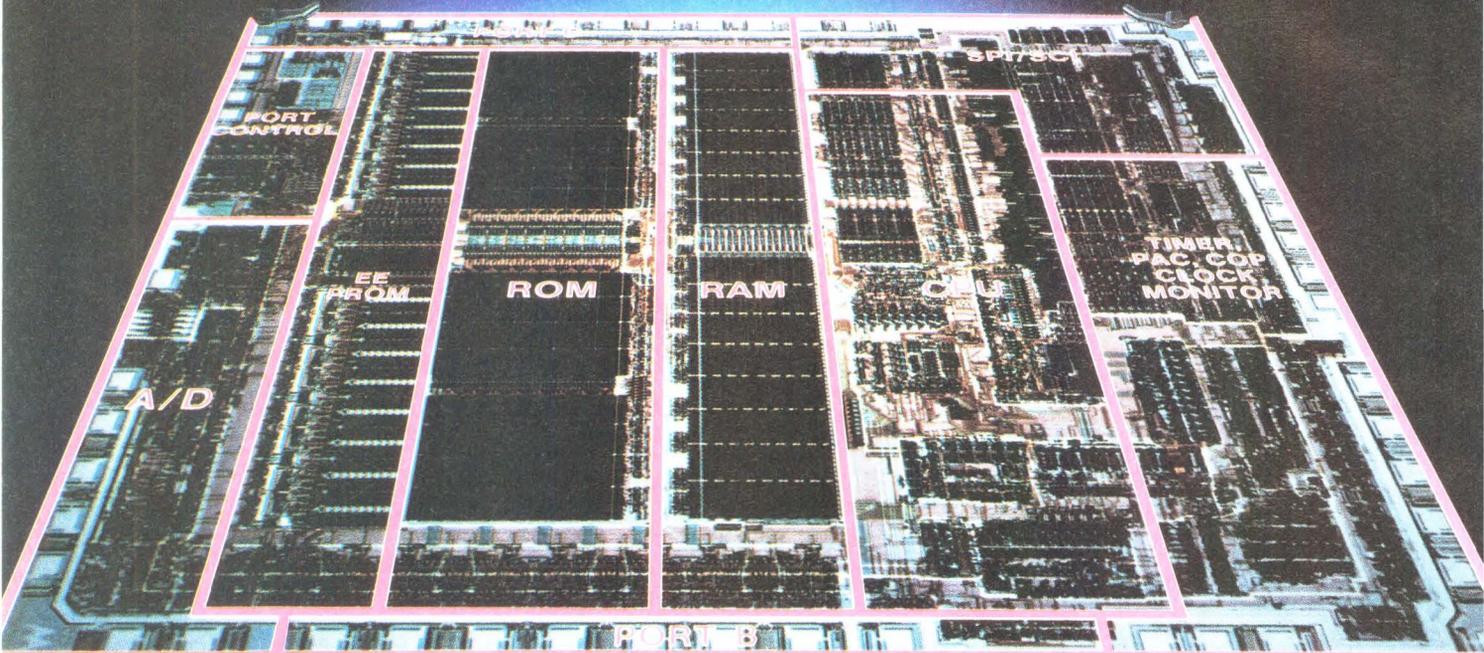
That's Megateknology.

World Headquarters, 9645 Scranton Road, San Diego, CA 92121

Megateknology is a trademark of Megatek Corporation. Whizzard is a registered trademark of Megatek Corporation. ANSYS is a registered trademark of Swanson Analysis, Inc. DOGS is a trademark of PASEC Engineering Consultants, Inc.



MOTOROLA



Transform any MCU-based design.

Use the first high-performance HCMOS MCU with on-chip EEPROM—the MC68HC11.

Raise your systems to the highest levels ever of cost-effective high performance with Motorola's pace-setting MC68HC11 HCMOS single-chip microcomputer.

The MC68HC11 sets new standards both for its unprecedented level of function integration and the processing power of its enhanced CPU core, creating the opportunity for your new designs to achieve a comparable level of excellence.

It serves cost-effectively at the system level in a tantalizing range of applications which includes but isn't limited to robotics, medical instrumentation, telecommunications and automotive systems.

Performance—in a class by itself.

With the assistance of the advanced HCMOS technology in which it is fabricated, the 'HC11 outperforms every other 8-bit MCU. In fact, it's in the highest performance echelon of all MCUs, even those with 16-bit cores.

True, performance claims are difficult to substantiate conclusively. However, the superiority of the MC68HC11 is readily demonstrable.

Compare it to an established, familiar high-performance leader like the MC6801.

The MC6801 was initially designed for 1 MHz bus operation and now offers 2 MHz bus frequency. Initial design of the MC68HC11 is for 2 MHz operation. But, speed is only part of the performance story.

Where the high-performance MC6801 contains an X index register, the new 'HC11 has both X and Y index registers.

The MC6801 instruction set, a significant enhancement of the original MC6800 instruction set, is an accepted standard for real-time control. That's all in the 'HC11 instructions, plus 91 powerful new op-codes including enhancements like divide, bit manipulation, exchange and special instructions for index register use. A full complement of 16-bit arithmetic instructions includes multiplication, division and comparison.

The previous zenith of vectored interrupt capability reached 8 in the MC6801U4. Now, the MC68HC11 offers 21 interrupt vectors.

The inherent low-power characteristics of CMOS allow the MC68HC11 to run at less than 20 mA. As for operating temperatures, the 'HC11 is available in ranges up to -40° C to +125° C, assuring reliable operation in the harshest of MCU environments.

There's more, but surely you see the magnitude of the difference.

Most functions ever on a one-chip MCU.

Nothing about the MC68HC11 is more exciting to designers who've learned about it than the unmatched complement of on-chip functions.

MC68HC11 integrated functions.

- 512 bytes EEPROM
- Enhanced CMOS core
- 8-channel, 8-bit A/D converter
- 8K bytes ROM
- 256 bytes fully static RAM
- Serial Peripheral Interface (SPI)
- Serial Communications Interface (SCI)
- 11 powerful timer functions
- 5 I/O ports, including powerful handshake capabilities
- Memory expandable to 64K

The MC68HC11 is the first MCU combining EEPROM and powerful core on an all-CMOS chip. Now data can be retained when the power is off. It also means the end product can be customized, and factory or field calibrations can be stored on the chip.

Communication capabilities get a great boost from two serial ports that provide the SPI and SCI. Eight of the timer functions provide five input captures and three output compares simultaneously. Watchdog, Pulse Accumulator and Real-Time Interrupt timers also are part of the MC68HC11. Two parallel ports with full handshake make up to 40 I/O lines available.

For occasions when the large amount of on-chip memory must be supplemented, an expansion bus increases capability to 64K.

Support from high-performance development tools.

With MC68HC11 design-in interest at a very high pitch, Motorola also is providing

the appropriate advance development support. Our all new HDS-300 development station will provide full support including emulation for the 'HC11 and other high-performance MCUs to come. Software hasn't been overlooked either. Structured macro assemblers are available now, and a "C" compiler is approaching introduction.

Match the MC68HC11's total system performance—CPU core, memory and I/O—with your system requirements, and see for yourself how it raises your new design to the highest level.

One-on-one design-in help.

In person or on the phone, anywhere in the U.S. or Canada, now you can get an engineer-to-engineer update on all our latest semiconductor products and technologies and how they can enhance your unique designs.

1-800-521-6274

Call toll-free any weekday, 9:00 a.m.-5:00 p.m. your time. If the call can't cover your needs, we'll have a local field applications engineer drop by. And get you to market. Faster.

For technical literature on the MC68HC11 single-chip MCU, fill out the coupon and send it to Motorola Semiconductor Products, Inc., P.O. Box 20912, Phoenix, AZ 85036.

We're
on your
design-in
team.



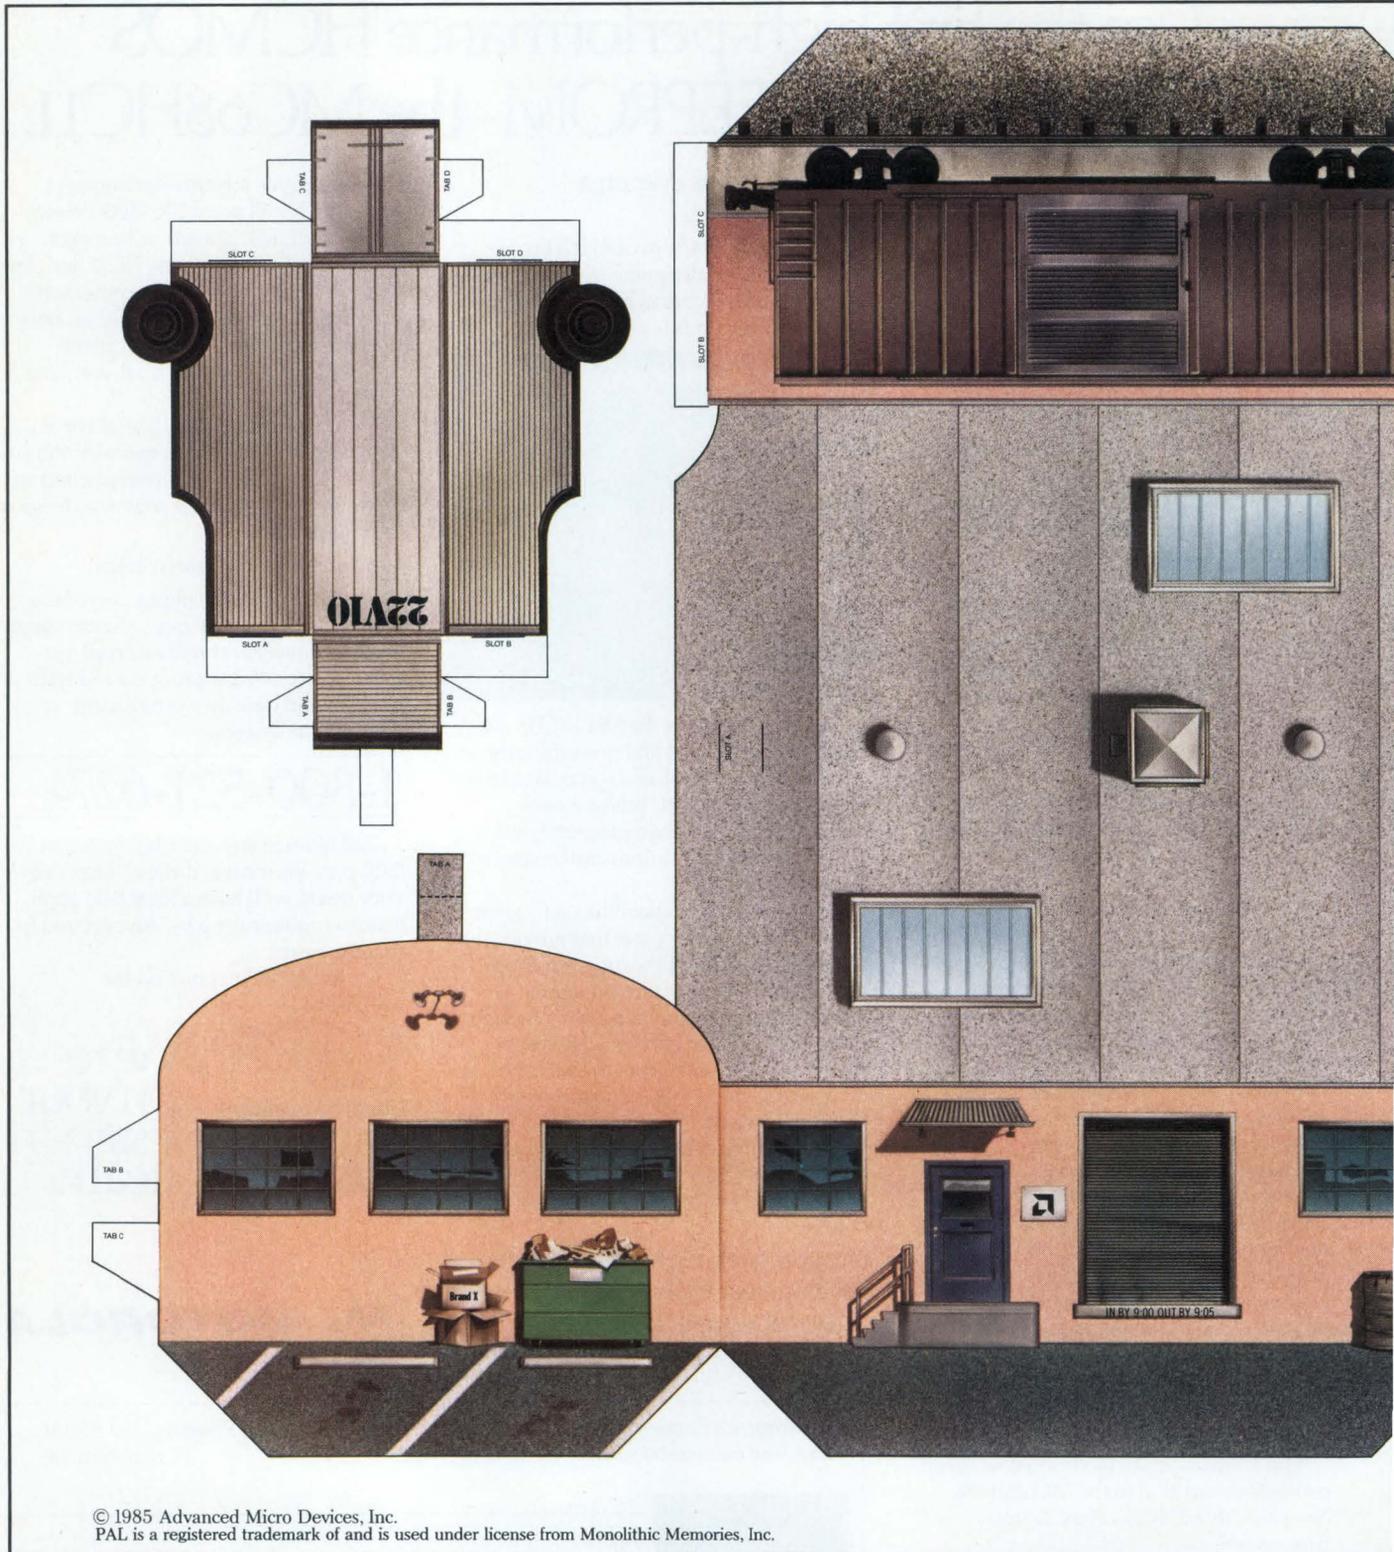
MOTOROLA

To: Motorola Semiconductor Products, Inc., P.O. Box 20912, Phoenix, AZ 85036
Please send me more information on the MC68HC11. 223CD030085



Name _____
Title _____
Company _____
Address _____
City _____ State _____ Zip _____
Call Me () _____

BE YOUR OWN



© 1985 Advanced Micro Devices, Inc.
PAL is a registered trademark of and is used under license from Monolithic Memories, Inc.

GATE ARRAY FACTORY.

Now you don't have to waste weeks waiting for custom gate arrays. Or face project-killing front end costs.

**Get the AmPAL22V10 and
be your own boss.**

With our new 22V10 PAL[®] device, you can make up as many logic devices as you need, for any number of gates up to a thousand.

And you can make as many of each design as you need. Or change your mind, and your design. Right on the spot.

The best pal a designer ever had.

The AmPAL22V10 not only gives you the room to do all this—it's one of the biggest PAL devices in the world—it gives you the flexibility. Its 132 product terms are distributed from 8 to 16 terms per output, just like you use them.

You determine the output architecture, too. Each output pin on a 22V10 can be either combinatorial or sequential, active high or active low.

Which adds up to remarkable production efficiencies. Because you use fewer packages per design.

Like all AMD PAL devices, the 22V10 has an extra built-in testing circuit. They're fully tested at the factory to guarantee AC and DC specifications. As well as AMD's industry leading post programming functional yields (PPFY).

**We'll put you as far ahead of the
competition as we are.**

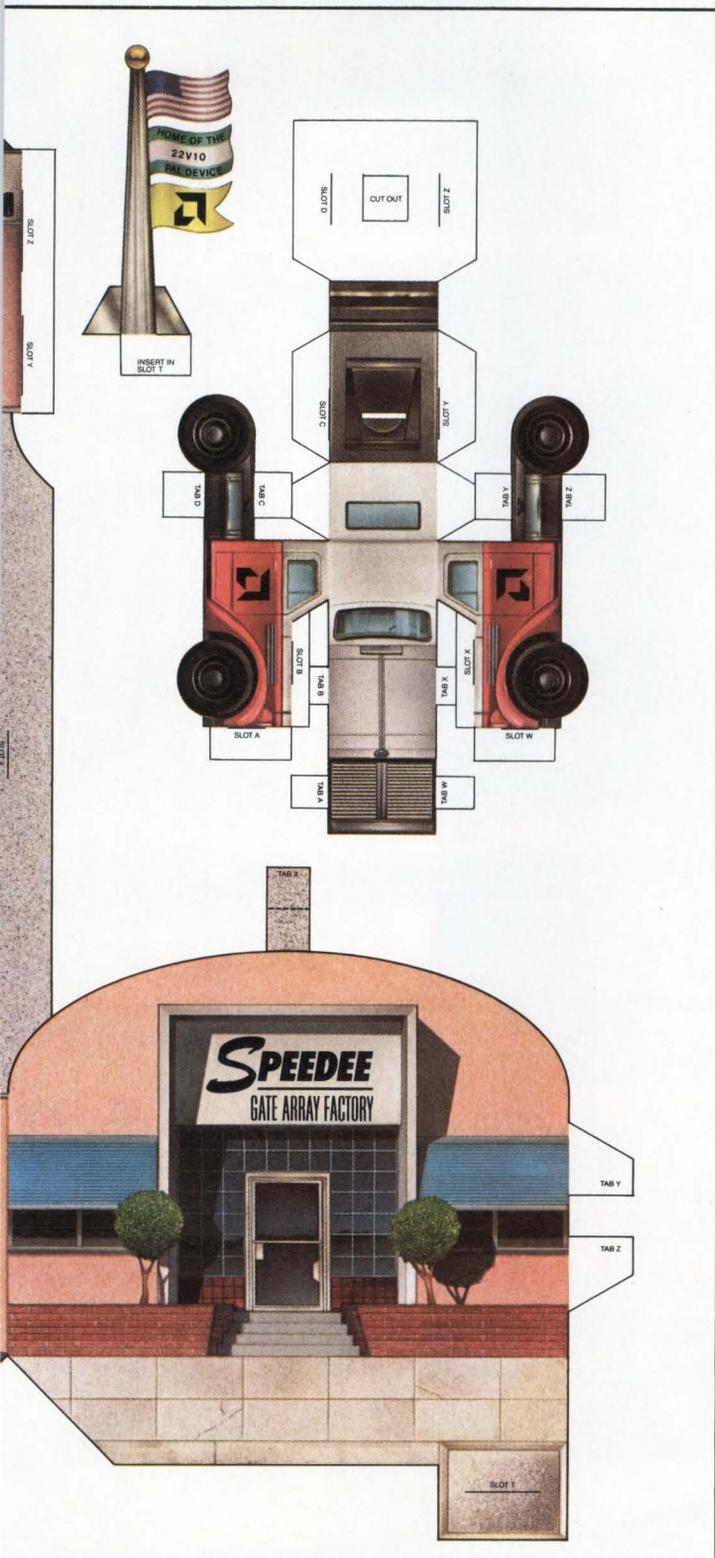
Controllers. Memories. Bipolar and MOS microprocessors. Communications circuits. Signal processors. Nobody makes as many peripherals for as many microprocessors as AMD.

The International Standard of Quality guarantees a 005% AQL on all electrical parameters, AC and DC, over the entire operating range.

And every single chip meets or exceeds the International Standard of Quality.

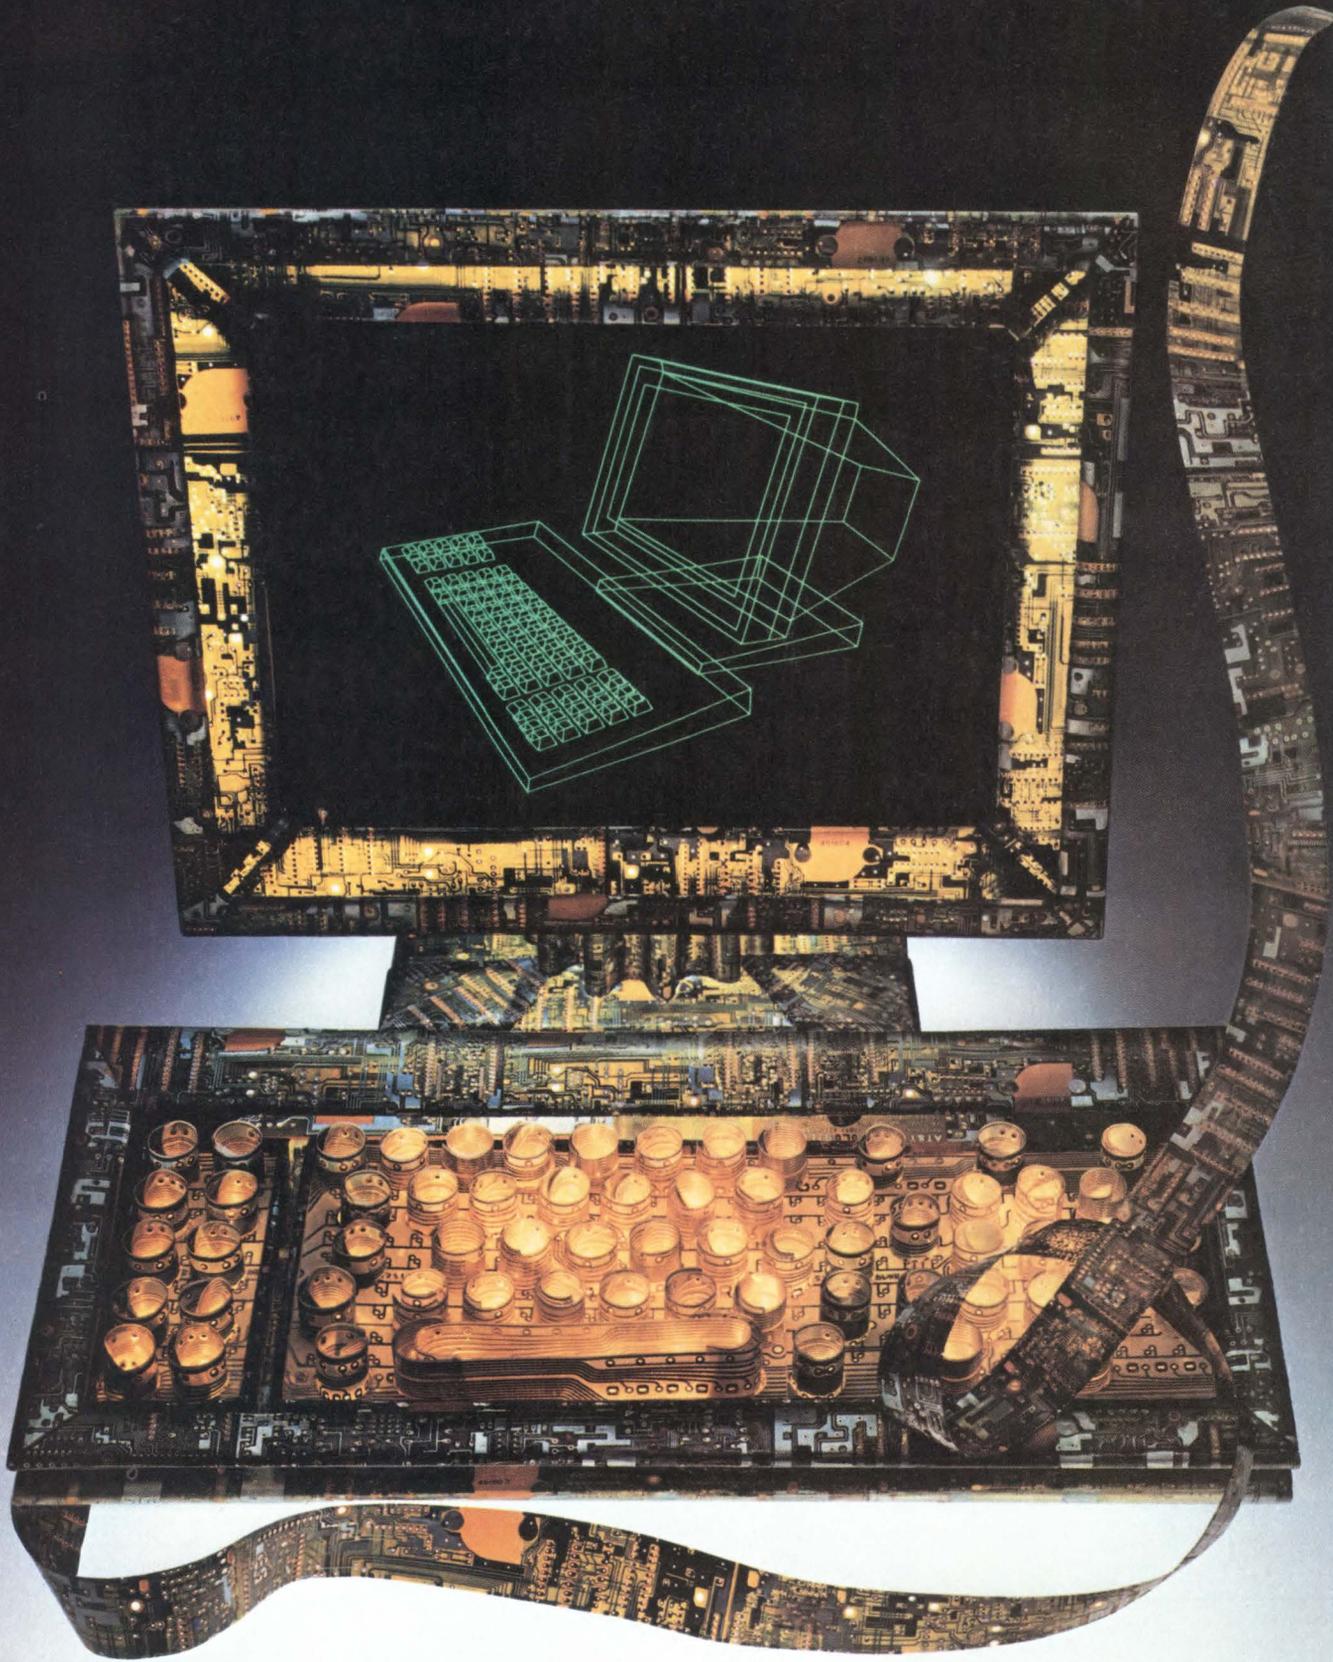
So next time you can't wait to get a hot new design into production, don't. Call AMD and ask about our 22V10. We'll send you a punch-out model of your new factory by return mail.

Then you'll have everything you need to be first out of the gate.



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.



WE'RE FLEXIBLE.

When your product design requires flexible circuits, no one can offer more design and manufacturing expertise than AT&T.

With the technical support of AT&T Bell Laboratories, we will work closely with you at any (and every) stage of your product's development to create exactly the right flexible circuit to fit your needs. To that end, we have developed our own state-of-the-art Computer-Aided Design and Computer-Aided Manufacturing systems. They bring you high-efficiency, high-precision capabilities that are among the best in the industry.

And our Applications Engineering Staff gets involved from the very start. You will be able to tap many of the best minds in the business to make sure that your flexible circuit design can be manufactured to meet the most demanding product specifications.

What's more, by following your project through from beginning to end, with up to 74 quality control checks that monitor everything from raw materials to fully finished circuit, we can avoid costly and time-consuming production problems. Problems that might otherwise jeopardize your timetable, not to mention your budget.

At AT&T, we've been designing and manufacturing flexible circuits for ourselves for over 25 years. Now we're putting our unparalleled expertise and experience to work for you, with the largest available capacity of anyone in the field.

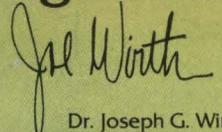
So for your next project, no matter how complex your design is, or how rigid your specs are, remember how flexible we can be, and call AT&T.

To find out more about how we can help solve your circuit design and manufacturing problems, write: AT&T, Dept. CG, 555 Union Blvd., Allentown, PA 18103. Or call 1 800 372-2447 (toll free) or 1 201 771-2601.

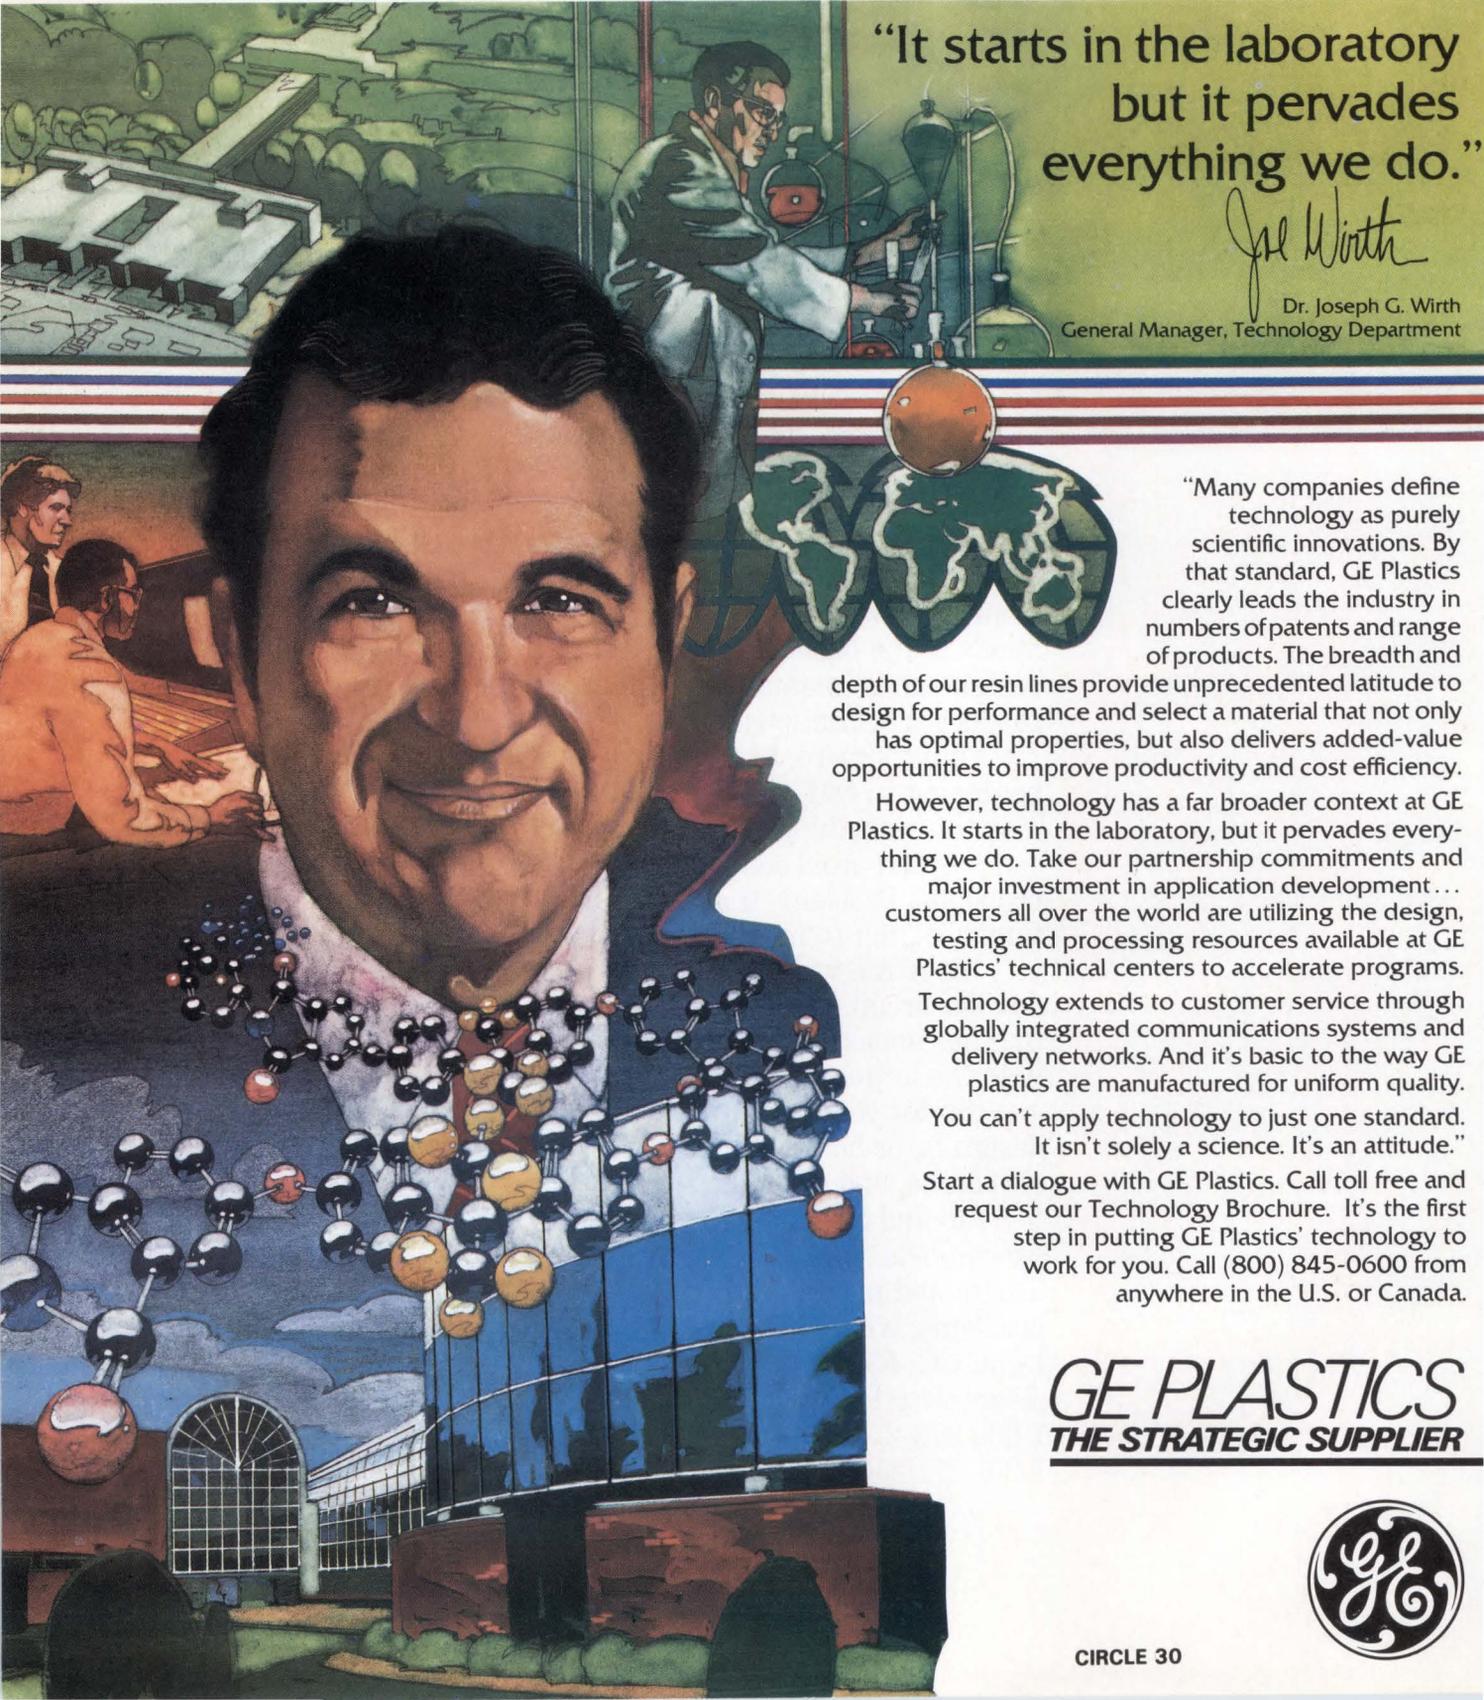


The Technology Of Technology

"It starts in the laboratory
but it pervades
everything we do."



Dr. Joseph G. Wirth
General Manager, Technology Department



"Many companies define technology as purely scientific innovations. By that standard, GE Plastics clearly leads the industry in numbers of patents and range of products. The breadth and

depth of our resin lines provide unprecedented latitude to design for performance and select a material that not only has optimal properties, but also delivers added-value opportunities to improve productivity and cost efficiency.

However, technology has a far broader context at GE Plastics. It starts in the laboratory, but it pervades everything we do. Take our partnership commitments and major investment in application development . . . customers all over the world are utilizing the design, testing and processing resources available at GE Plastics' technical centers to accelerate programs.

Technology extends to customer service through globally integrated communications systems and delivery networks. And it's basic to the way GE plastics are manufactured for uniform quality.

You can't apply technology to just one standard. It isn't solely a science. It's an attitude."

Start a dialogue with GE Plastics. Call toll free and request our Technology Brochure. It's the first step in putting GE Plastics' technology to work for you. Call (800) 845-0600 from anywhere in the U.S. or Canada.

GE PLASTICS
THE STRATEGIC SUPPLIER



Fault simulation becomes design verification tool

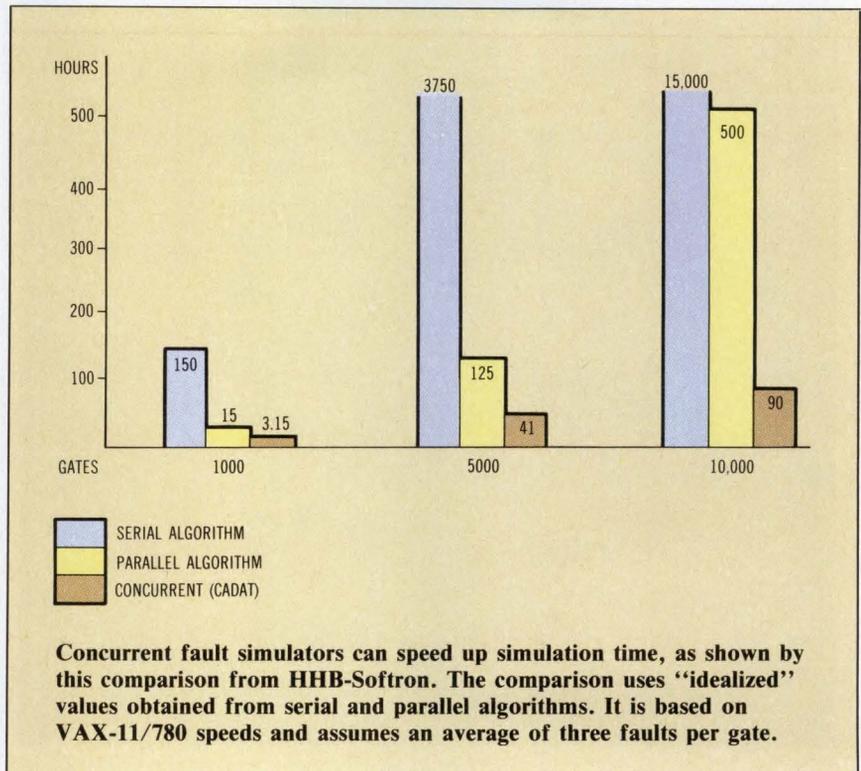
A software tool for evaluating test vector coverage is showing up as a design tool in the computer aided engineering environment. The fault simulator, a computationally intensive tool originally designed for test engineers, was once relegated to superminicomputers and mainframes. Now, engineers can apply fault simulation at their desks—thanks to the emergence of powerful workstations, improvements in fault simulation algorithms, and an increasing demand for testability in design.

A fault simulator detects the fault coverage percentage of specified sets of test vectors—that is, the percentage of possible faults the vectors will find. To do so, the fault simulator inserts a fault into a simulated logic network, or “machine.” It then inputs a test vector, and compares the output received from the “bad” machine to the output that would be received from a “good” machine. If the output differs, the fault has been detected by the test vector.

Since fault simulation initially requires a good-circuit simulation, it is closely related to logic simulation. In the CAE environment, fault simulators and logic simulators are usually sold together—but they serve different purposes. According to Todd Westenhoff, OEM support manager for HHB-Softtron, “Logic simulation tells you how your product works. Fault simulation tells you if you’ve done enough analysis to be confident you can produce it. It helps you get silicon back that’s right the first time.”

Test engineers can use fault simulation results to build test programs, and avoid much of the “reverse engineering” that test engineers usually have to do. But the benefits of fault simulation do not reach entirely beyond testing. Fault simulation also helps the designer find logic errors, such as redundant logic elements, before a design is cast in silicon. To avoid such errors, many foundries now demand that an engineer provide a high coverage test vector set along with the design for any new IC.

Fault simulators generally simulate several types of faults, including



stuck-at-0, stuck-at-1, stuck-at-X (unknown), and stuck-at-Z (high impedance). Since a typical network might have thousands of possible errors, fault simulation can be an extremely slow process. Most fault simulators in the CAE environment can mimic many faults simultaneously. Even so, many engineers resort to simulating a statistical sample of faults.

Speeding it up

Serial fault simulators produce one fault at a time. To speed the process, parallel fault simulators insert single faults into multiple machines. Each fault is normally represented by a bit in the host computer’s memory word. Thus, a 32-bit computer can simulate a maximum of 31 faults simultaneously, plus one good machine.

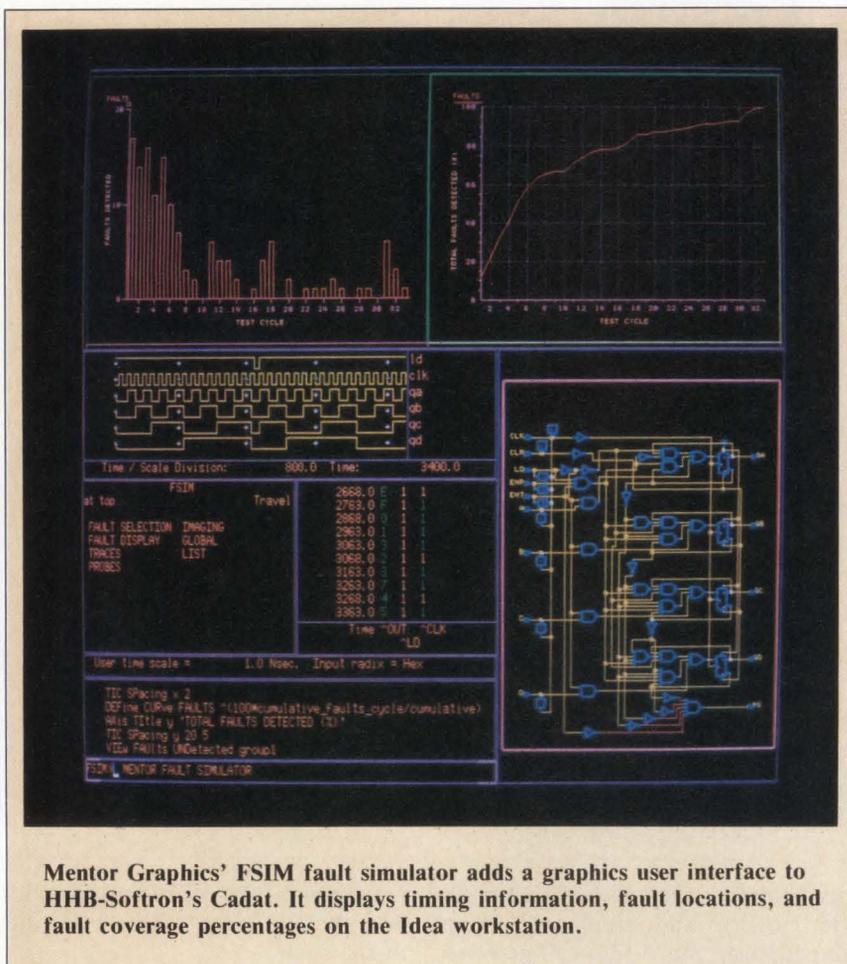
Concurrent fault simulators duplicate only the portion of a circuit that would actually be affected by a fault. They simulate the good machine, and track the effect of a fault as a signal propagates through the fault location. Only when the good and bad

networks display different behavior does the simulator track activity due to the fault.

Since most faults affect a small portion of a circuit, concurrent fault simulation is usually much faster than parallel fault simulation. Moreover, as circuit size increases, the simulation time required for a parallel fault simulator tends to rise exponentially. The simulation time for a concurrent simulator shows a more linear increase.

In theory, a concurrent fault simulator can simulate any number of faults in a single pass. This ability is limited by the system’s memory, however, because concurrent fault simulators build a fault list that keeps track of the fault affect at every node. This list can grow dramatically, especially when a fault affects a large portion of a circuit. To avoid unpredictable memory requirements, most concurrent simulators will limit the number of faults that are simulated in one pass.

The level of fault simulation also affects simulation time. While most fault simulators allow gate-level



Mentor Graphics' FSIM fault simulator adds a graphics user interface to HHB-Softron's Cadat. It displays timing information, fault locations, and fault coverage percentages on the Idea workstation.

simulation, some also allow behavioral, or "functional," simulation. In behavioral simulation, only the inputs and outputs of a functional model are evaluated. This approach will not grade a complete set of test vectors, but it accommodates the top-down design procedures used by a number of engineers. It also cuts simulation time significantly.

Serial and parallel simulators

Although concurrent simulators are generally regarded as state of the art, vendors of serial and parallel simulators are finding ways to improve the process. The Daisy Fault Simulator is sold with the Verification Support System for the Logician workstation. The company says this serial fault simulator, running on its Megalogician workstation, nears the speed of a concurrent fault simulator on a smaller machine. Nevertheless, Daisy plans to introduce a concurrent fault simulator late this year.

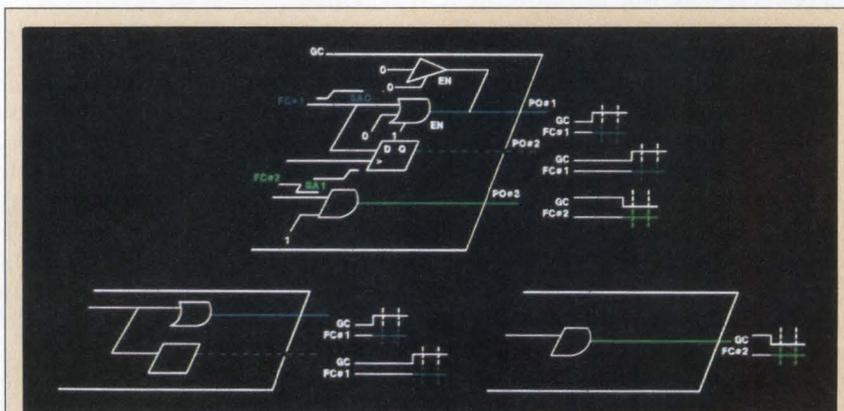
Calma's TCAT package, which includes testability analysis and automatic test generation, separates fault generation from fault simulation. The fault generation process, which can be used at the gate or functional

level, reduces the number of faults by eliminating logically equivalent faults. Calma's parallel fault simulator can simulate a statistical sample of nodes to reduce execution times. TCAT, based on Calma's Tegas-5, runs on VAX/VMS, IBM 3000 series, and Apollo systems.

In the middle of this year, Calma will introduce a multiword parallel fault simulator. This new approach allows faults to be represented by up to 255 32-bit words, instead of just one word. The actual number of words used depends on the number of faults being simulated and the amount of memory available. There is a big savings in setup time, and as many as 8159 (255 x 32 - 1) faults can be simulated in a single pass.

Preliminary benchmarks indicate that the multiword parallel version will run 10 to 15 times faster than the existing Calma fault simulator. According to Dave Dougherty, product manager for Calma's design analysis group, this will make it faster than existing concurrent simulators. The company also claims that memory use will be more predictable. However, a higher gate count may offset some of these advantages.

GenRad's Hilo-2 package includes a fault simulator with a hybrid parallel/concurrent algorithm. It is essentially a parallel fault simulator with a concurrent scheduling mechanism. The GenRad simulator is not limited by computer word length, and users can specify the



Concurrent fault simulators track only the portion of a circuit that is actually affected by a fault. Here, two faulty outputs result from simulated stuck-at-0 and stuck-at-1 faults. Teradyne's Lasar fault simulator calculates the response from these locations forward.

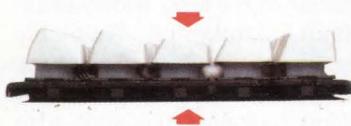
Thanks, HI-TEK... for a Low-Profile keyboard that works...and works... and works...

You're welcome. We thought you might appreciate the flexible, compact design. Our Series 725's nominal 18.4mm (.725") profile leaves you 11.6mm (.456") of room to meet DIN standards.

With no compromises, either. The same high reliability keyswitch design concept, smooth feel, full travel and up to 100 million cycle life used in HI-TEK's Standard Profile series is found in our new Series 725 Low-Profile keyboards.

It's what you'd expect from a company that's made over 10 million keyboards — that's more than 600 million keystations — in the past 10 years.

Our field-proven passive contact design provides consistent switch performance and includes a unique redundant contact system — ensuring switch closure even if foreign particles pass through the contact area.



Just 0.725" (18.4mm)



We think OEM. That's why we keep costs down, provide fast delivery and stress HI-TEK quality. Consider our unique sublimated keycap printing. To distinguish your keyboard from the competition, you

can design your own custom legends in multiple colors, and with greater intricacy than ever before possible with traditional two-shot molded keycaps.

By designing integrated jumpers into the contacts, we make a fully encoded keyboard fit on a cost-saving one-sided PCB. Our versatile microprocessors are optimized for keyboard use.

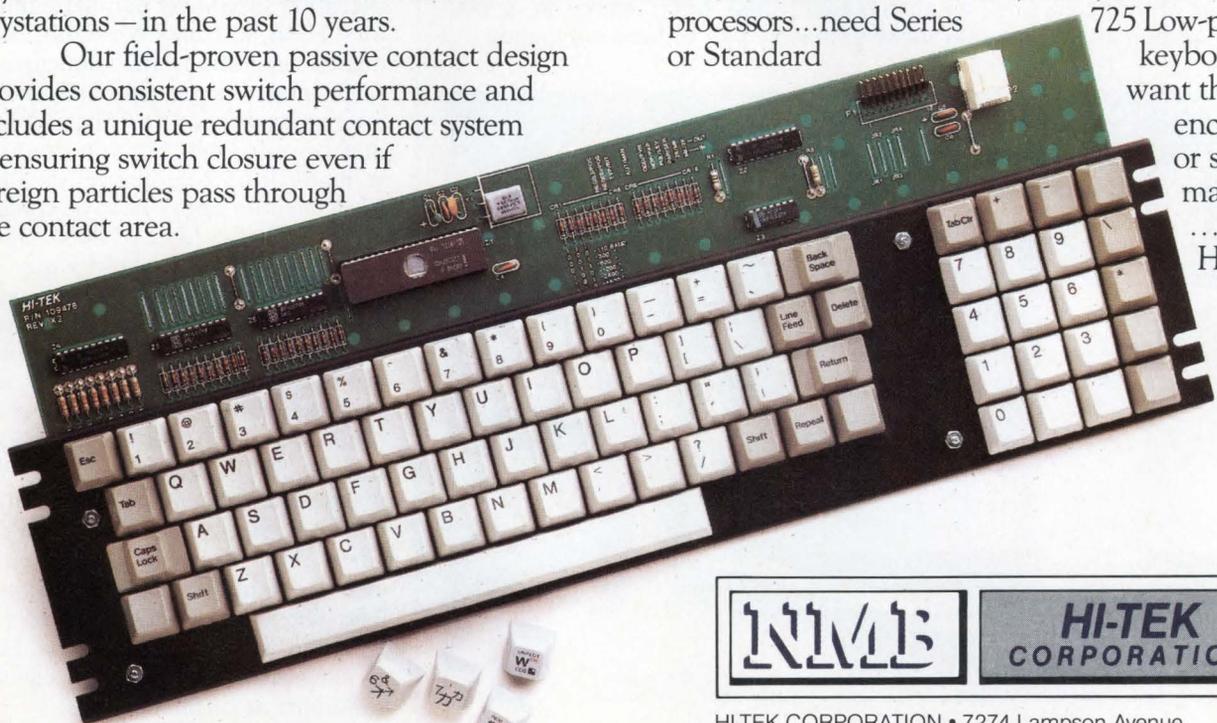
And, by doing our own encoding, we cut both time and cost.

Whether you manufacture mainframes, minis, micros or word processors...need Series 725 Low-profile keyboards... want them

encoded or simple matrix...call HI-TEK. Today.



processors...need Series 725 Low-profile keyboards... want them



HI-TEK CORPORATION • 7274 Lampson Avenue
Garden Grove, CA 92641 • (714) 898-9511 • Telex: 67-8486

Fault simulator picks up speed

Software fault simulation can take hours in even relatively simple networks. This time can be slashed to seconds with a hardware fault simulator—provided the user is willing to pay \$450,000 or more for a dedicated, high performance processor.

According to preliminary benchmarks, Zycad's Fault Evaluator can run fault simulation up to 1000 times faster than possible with comparable software products. This hardware simulator puts a concurrent fault simulator into silicon. It uses a high speed pipelined architecture, high bandwidth distributed memory, and many levels of concurrency.

The Fault Evaluator is an enhancement of Zycad's Logic Evaluator. It provides the same logic simulation capabilities as the Logic Evaluator, which also claims 1000 time the speed of software simulators. The Fault Evaluator is a module that fits inside the Logic Evaluator cabinet. It can replace a Logic Evaluator module, or be used alongside one.

According to project manager Howard Krohn, the Fault Evaluator can simulate 3000 faults in a 2000-gate network in 31 s, using a VAX/VMS host. By comparison, a commercially available parallel fault simulator takes 10 hours to run the same simulation. The Fault Evaluator is currently undergoing beta-site testing at Digital Equipment Corp and Data General Shipments are expected by the middle of this year.

The Logic Evaluator cabinet will hold up to 16 Logic Evaluator or Fault Evaluator modules. Each Fault Evaluator module can store a fault list of 500,000 "fault effects." The fault list is stored in 256-Kbyte dynamic RAM. A maximum of 16,000 faults can be simulated in a single pass, regardless of the number of modules.

Currently, users must write their own frontend software. Future plans, however, call for the Fault Evaluator to use Zilos, Zilog's software interface for the Logic Evaluator. The Fault Evaluator can run with a VAX/VMS, IBM 3000 Series, Data General MV8000, CDC Cyber, or Apollo host.

This kind of power does not come cheaply. Prices for the Fault Evaluator start at \$450,000 for a single module, and range up to \$3 million for 16 modules. Users who already have a Logic Evaluator can upgrade one module for \$150,000, or 16 modules for \$1 million. This is not the kind of thing one finds on an engineer's desk, but it may be attractive to large installations that make extensive use of fault simulation.



number of faults to be simulated in a single pass. According to John Hengesbach, product marketing manager, the best trade-off between speed and memory requirements is usually around 150 to 500 faults.

Hengesbach claims that Hilo-2 comes out ahead when it is benchmarked against concurrent simulators. "The parallel approach is faster for some networks," he says, noting that concurrent simulators assume a fault affects only a small portion of a circuit. Hilo-2 is available in a number of environments, including VAX/VMS, VAX/UNIX, IBM/CMS, Apollo, Sun, Metheus, CAE Systems, and Prime.

The concurrent approach

HHB-Softron's Cadat, an integrated logic and fault simulator, includes one of the most widely available fault simulators. Available directly from HHB-Softron for Apollo, Sun, VAX/VMS, and VAX/Unix systems, Cadat is also sold on an OEM basis by Mentor Graphics, Cadnetix (Boulder, Colo), Racal-Redac (Santa Clara, Calif), FutureNet (Canoga Park, Calif), and several test system manufacturers.

By using a concurrent algorithm, HHB-Softron claims a big increase in execution speed over parallel simulators. HHB-Softron's Westerhoff dismisses suggestions that concurrent fault simulators pose unwieldy memory requirements. "It really uses less memory per fault, but you simulate so many at once people get the idea it's a memory hog," he says. Cadat, however, has a limit of 4000 faults per pass to restrict memory use.

Cadat is intended for ICs and printed circuit boards. But because semicustom and custom IC designers are the heaviest users of logic and fault simulation, Cadat is currently optimized for ICs. Cadat provides three levels of simulation—primitive, macro, and behavioral. At the primitive level, 90 types of devices are defined. The macro level allows the engineer to hook primitives together. The functional level allows vectors to be graded at inputs and outputs. Different portions of a design can be simulated at different levels.

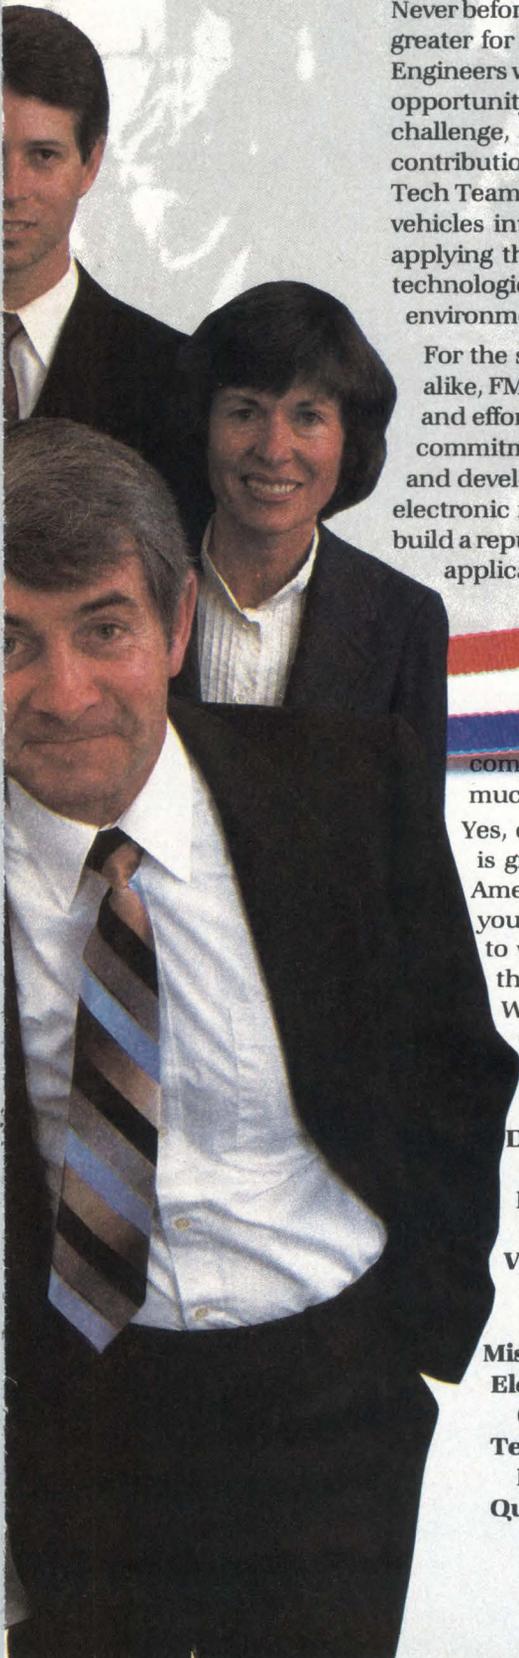
(continued on page 78)

**When General MacArthur
fulfilled his promise...**

FMG WAS THERE



**with the very latest in
advanced equipment to support
a monumental task!**



History was ours to share as FMC Amphibious Landing Vehicles helped make history in the past. And the spirit of pride and glory continues today at FMC as we continue our leadership role as the world's largest manufacturer of military tracked vehicles.

Never before has the opportunity been greater for Electrical and Electronic Engineers who are looking for a career opportunity that brings reward, challenge, and meaningful contribution. Our "Hands-on, High Tech Team" is designing and building vehicles into the 1990's and beyond, applying the most advanced technologies in the most progressive environment imaginable.

For the specialist and generalist alike, FMC is a place where talent and effort are backed by a commitment to expanded research and development. Our increased electronic involvement is helping us build a reputation for the advanced applications of microprocessors,

computers, fiber optics and much more.

Yes, our heritage in America is great, but our future in America is even greater. Put your professional skills to work, and share the pride at FMC.

We currently have openings for:

- Analog Circuit Design Engineers**
- Digital Circuit Design Engineers**
- Electro-Mechanical Designers**
- Vehicle Power Generation and Distribution Systems Designers**
- Missile Analysts**
- Electrical Engineering Checkers**
- Technical Computing Electronics Specialists**
- Quality Assurance Engineers**

At FMC, you will also find professional careers in Mechanical and Manufacturing Engineering disciplines - opportunities for Motor Speed Control Designers, Mechanical Design Engineers, Armament Analysts, Turret, Vehicle Electronics, Fire Control and Product Safety Engineers. We have a continuing need for Industrial and Methods Engineers plus Technical Writers, Human Factors Specialists, Programmers and Systems Analysts/IBM 3081.

To learn more about career opportunities, from entry level to project management, please send your resume to:

**FMC Corporation
Ordnance Division,
1107 Coleman Avenue,
Box 1201,
Department 42-85-01,
San Jose,
California 95108.**

We are an Equal Opportunity Employer. U.S. citizenship required for most positions.

join the hands-on high tech team.



**Share
the Pride**

FMC

AND FMC IS HERE TODAY

with career
opportunities for
**Electrical and
Electronics
Engineers**
looking
to meet
the
challenge of the future
through advanced
technologies!



Fault Simulators for Design Engineers

Vendor	Package	Runs on	Sold with	Comments
Calma (Austin, Tex)	TCAT	VAX/VMS, IBM, 3000 Apollo	Testability analysis, automatic test generation	Parallel fault simulator. Multiword parallel version planned.
Daisy (Sunnyvale, Calif)	Verification System	Logician workstation	Logic simulator, timing analysis, testability analysis	Serial fault simulator. Concurrent version planned.
Gateway Design (Littleton, Mass)	AIDSSIM	VAX/VMS, Apollo, Sun, IBM PC AT	Independent package	Concurrent fault simulator.
GenRad (Santa Clara, Calif)	Hilo-2	VAX/VMS, VAX/Unix, Apollo, Sun, CAE, Metheus, IBM/CMS, Prime	Logic simulation, timing analysis	Uses hybrid parallel- concurrent algorithm.
HHB-Softron (Mahwah, NJ)	Cadat	VAX/VMS, VAX/Unix, Apollo, Sun	Logic simulator	Concurrent fault simulator. Several OEM versions available.
Mentor Graphics (Beaverton, Ore)	FSIM	Idee workstation	Independent package	Enhanced version of Cadat fault simulator.
Teradyne (Boston, Mass)	Lasar-6	VAX/VMS, Valid Logic workstations	Logic simulator, timing analyzer	Concurrent fault simulator with time windows.
Zycad (Arden Hills, Minn)	Fault Evaluator	VAX/VMS, IBM 3000, DG MV8000, CDC Cyber, Apollo	Logic simulator	Hardware fault simulator with concurrent algorithm.

(continued from page 74)

Cadat allows the user to simulate the signal strengths commonly associated with MOS circuits. These strengths include active, passive, floating, and indeterminate. Three logic levels (stuck-at-0, -1, and -X) are mapped onto these strengths, giving 12 possible states. Floating stuck-at-X is equivalent to stuck-at-Z.

The high level Digital Stimulus Language lets the user specify stimulus and define macro operations. The user can also set breakpoints and examine states internal to the network. Cadat allows the user to define and specify categories of faults.

Getting graphic

Mentor Graphics has extracted the fault simulation portion of Cadat and named it FSIM. Mentor integrated the fault simulator with its own logic simulator, and gave it a graphically oriented user interface. With most fault simulators, users must go through lengthy printouts to get the results. FSIM graphically shows logic diagrams with fault locations, and

provides colorful charts that show fault coverage percentages, fault categories, and test vector information.

Teradyne's Lasar version 6, available on VAX/VMS and Valid Logic workstations, includes logic simulation, timing analysis, and fault simulation. During logic simulation, the Chronos timing language allows the user to specify gate-level delays and functional timing specifications. Lasar's concurrent fault simulator uses this information to take nominal delays into account.

Lasar also lets the user specify "time windows" in which a fault must be propagated to an output. These time windows take into account the period of time a tester looks for circuit responses. A fault is scored as "detected" only if it can be observed when a tester is looking.

Adlib, a behavioral modeling language, is a recent addition to Lasar. According to Teradyne, fault simulation can run 30 to 50 times faster with Adlib. "You lose resolution but gain speed," says Daryl Layzer, marketing services manager. "You propa-

gate faults through a behavioral block, and look at the inputs and outputs of the block."

Another recent addition to Lasar is Prosecutor, an automatic test generation option that can run interactively with fault simulation. Prosecutor develops test vectors for most easily detectable faults first and works with the fault simulator to find vectors for progressively harder faults. Users can set up Prosecutor as a batch process, putting a limit on run time or making a fault coverage goal.

In addition to stuck-at-0, -1, -X, and -Z faults, Lasar can simulate open pins, adjacent pin shorts, and shorts between user-specified pins. The user can specify the types of faults, and simulate internal device faults as well as pin-level faults. To avoid memory management problems, the simulator sizes the fault set to available memory.

A new entry in fault simulation, Aidssim from Gateway Design Automation Company, claims to be the fastest fault simulator on the market. According to the company, this

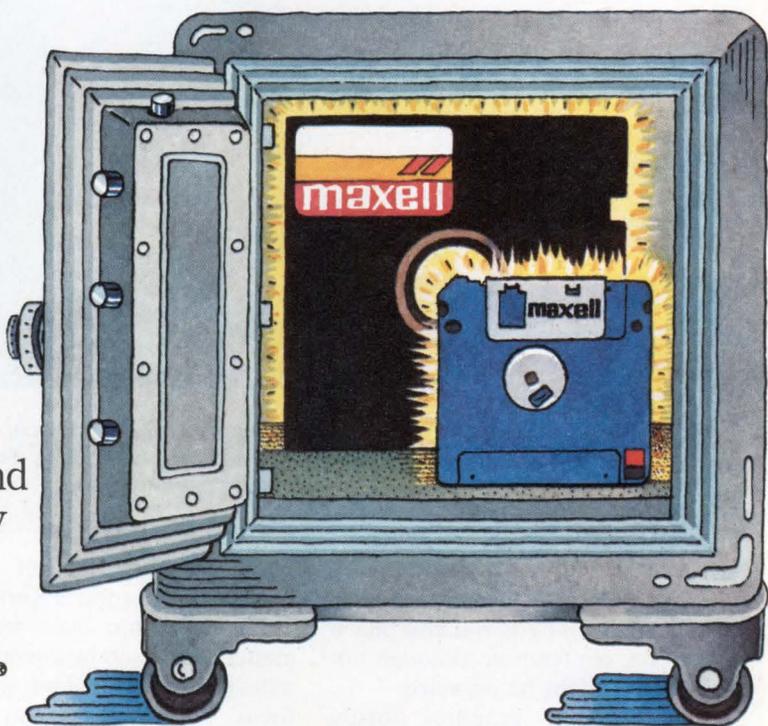


Maxell Gold.

Precious metal for your 3½" and 5¼" high density disk drives.

Prospecting for quality?
Tap into the Gold Standard.
Maxell. We were the first to
make the 5¼" high density
floppy commercially available.
And this same technological
expertise is reflected in the
superb performance of our
3½" microdisks. Maxell Gold.
Industry leaders in durability and
error-free performance. Quality
your drives can depend on.

maxell.
IT'S WORTH IT.



concurrent simulator can run a simulation on a 7000 gate network in two hours, using a VAX-11/750. Aidsim's designer, Gateway president Prabhu Goel, says Aidssim can simulate any number of faults in one pass. Goel claims Aidssim "totally avoids the memory problem" posed by concurrent simulators, but declines to say how.

Aidssim has a high level language similar to Cadat's DSL. The language allows the user to stop the simulation, override internal states, and restart the simulation. Aidssim is available on VAX/VMS, Apollo, and Sun workstations, with an IBM PC AT version expected early this year.

Efforts are underway to find the next breakthrough in fault simula-

tion. One promising candidate is Stafan, an algorithm developed by Sunil Jain and Vishwani Agrawal at AT&T Bell Labs (Murray Hill, NJ). In this algorithm, controllability and observability are defined as probabilities and are estimated from a good circuit simulation. These probabilities are used to derive overall fault coverage for a specified set of test vectors.

Random patterns

Fault simulation with random test vectors presents another way to reduce the computation and memory requirements of fault simulation. Aidssim allows the user to run both random and preselected patterns. The COP algorithm, developed at Bell-Northern Research (Ottawa, Ont) by

Franc Brglez, includes fault simulation with random patterns. Early results indicate this approach substantially reduces CPU time, and provides reasonably accurate fault coverage for many circuits.

Even though increasing VLSI complexity will make fault simulation more difficult, little doubt remains that fault simulation will become a basic design tool. Demand from semiconductor manufacturers, test engineers, and corporate management will compel design engineers to evaluate test vectors before going to silicon. In the process, design engineers may find that fault simulation is a powerful tool for verifying designs.

—Richard Goering,
Field Editor

SOFTWARE

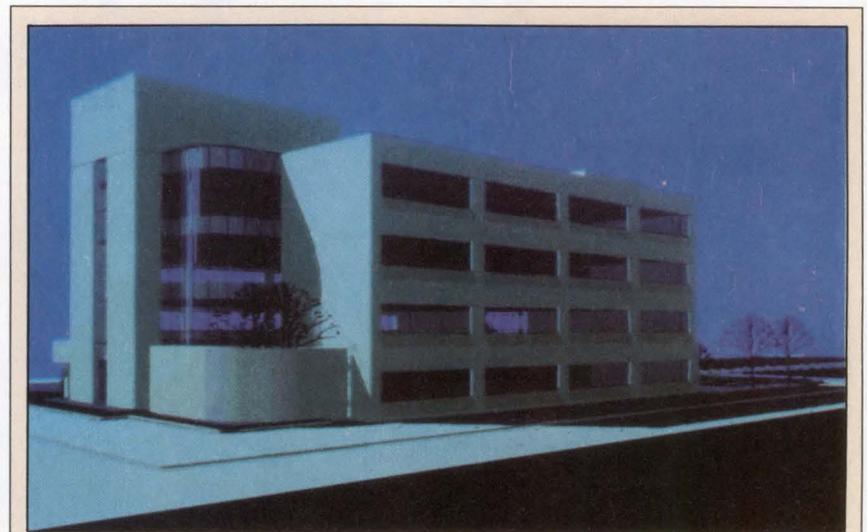
Applications dominate Computer Graphics '85

Computer Graphics '85, the conference and exhibition sponsored by the National Computer Graphics Association, will cover a wide range of business and industrial applications. Scheduled for April 14 to 18 at the Dallas, Tex Convention Center, this year's tutorial sessions will also introduce some relatively new topics involving computer graphics.

The bulk of the sessions, however, appear to be oriented toward vertical applications rather than system-level topics. A significant portion of technical topics will aim at management rather than at engineering. For instance, only two sessions will be allotted for electrical computer aided design/computer aided manufacturing, as opposed to 14 sessions on visual arts and design, or 19 on business management and graphics.

Some topics will emphasize areas of interest to system engineers, including artificial intelligence and computer integrated manufacturing. The issue of graphics standards will also play a role in the conference, although not as large as might be expected.

No computer graphics forum would be complete without some knock-down, drag-out sessions on



Architectural design firms can benefit from computer graphics software. This is demonstrated by the solid shaded perspective model created using the Integraph system.

graphics standards. But, since the issue of the Graphical Kernel System versus Core has been settled, this matter appears to be muted—and the schedule reflects it. Two sessions will focus on developing applications using the GKS, but there are none on Core. A separate tutorial session will

cover the Programmer's Hierarchical Graphics Standard as well as one reviewing the progress of the Initial Graphics Exchange Specification. Two other technical sessions deal with implementing extant standards and using some of the many graphics software tools available.

SIMPLICITY OF DESIGN.
GREAT DENSITY.
HIGH SPEEDS.
LOW POWER DISSIPATION.
LATCH-UP IMMUNITY.
TRANSIENT RADIATION
TOLERATION.

SO WHAT?

So SOS is finally being recognized for its inherent superiority in the world of custom and semi-custom CMOS integrated circuits.

The key to the superiority of the Silicon-on-Sapphire circuit is the isolating monolithic sapphire substrate, which makes it possible to achieve complete isolation between each device on the same chip. That means latch-up is simply impossible, toleration of transient radiation is 100 times greater than bulk CMOS and, at the same time, the highest possible integration density can be produced. It also means very low parasitic capacitances, which means 300% greater speed and far lower power dissipation than is possible with bulk CMOS. Since that sapphire substrate eliminates any need for the guardbands and wells which consume chip area in bulk processes, SOS is far more versatile and less complex to design.

Now, under development, ASEA HAFO's 2-micron SOS/CMOS process will be even more

impressive. For instance, it can match the speed of ECL while producing *substantially* lower levels of power dissipation. The process can be applied to existing 4-micron or larger bulk CMOS circuit designs to increase speed significantly with no need for redesign.

So what? So if you require state-of-the-art performance from your integrated circuits, talk to ASEA HAFO. We're the largest volume producer of custom commercial SOS circuits in the world, and the first to convert SOS production from 3-inch to 4-inch wafers. We're committed to the level of performance which is possible with SOS. And that's what makes us different.

For more information contact Michael Orenich, ASEA HAFO, Inc., 11501 Rancho Bernardo Road, P.O. Box 27255, San Diego, CA 92128-0925, (619) 485-8200, Ext. 401. Offices in Stockholm, Paris and Zurich.



ASEA HAFO

The Midnight Sun Rises in the West.

The more traditional business and industrial areas, such as building architecture, will feature management-oriented sessions and tutorials. Indeed, architecture appears to be a burgeoning area of computer graphics. Five tutorials—three of them management-oriented—and six technical sessions cover the impact of computer automation and graphics on the design office. Among them are the issues of smoothly integrating automated design and drafting systems into operations with manual procedures.

Closely allied with the architecture and engineering topics are the mechanical CAD/CAM sessions consisting of four tutorials and three technical sessions. One of the more intriguing of these will be "CAD/CAM to CIM: How to Get There From Here."

Less than one in 100 CAD/CAM operations have actually gone over to CIM. While this tutorial will discuss the why and how-to of manufacturing integration, a separate set of ses-

sions on CIM will comprise a tutorial on strategies for implementation. A technical session will also discuss CIM as the integration of many technologies to bring management and control techniques into a complete system. Such a system links focused manufacturing to a company's strategic business objectives.

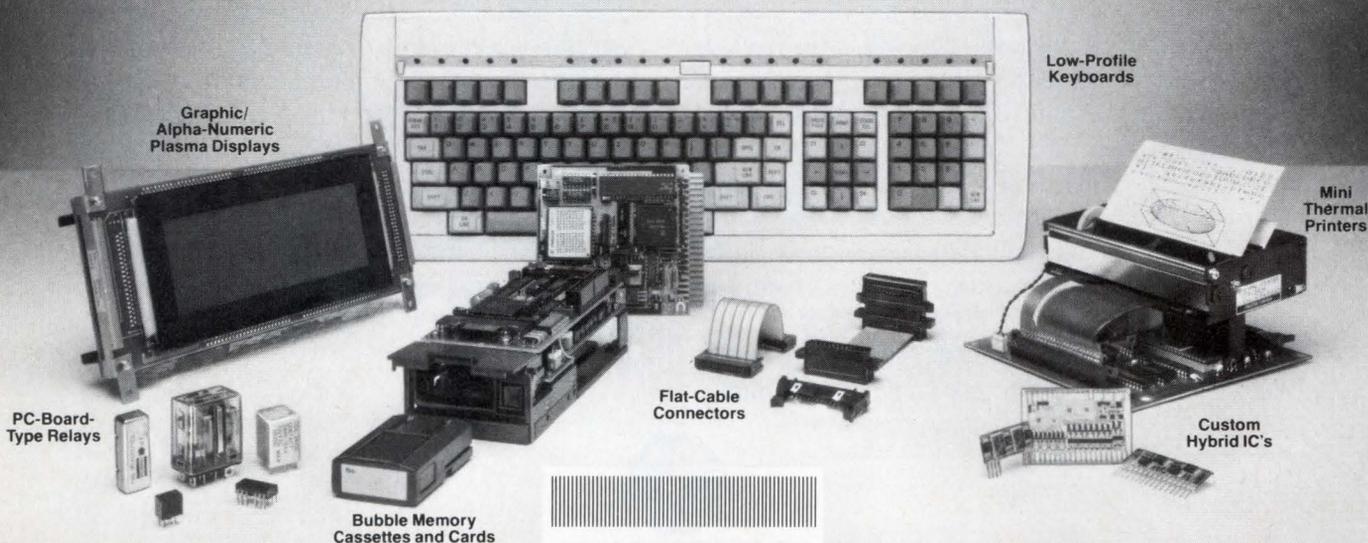
As might be expected, there is a heavy business emphasis on the use of graphics in business management and on making a business of computer graphics. Seven tutorials and 12 technical sessions will cover such topics as in-house publishing, making better charts and graphs, executive workstations, entrepreneurship in computer graphics, and ergonomics in the office that uses graphics.

A series of sessions on statistical graphics should interest not only graphics business users, but members of the scientific community as well. With only one tutorial to eight technical sessions, this area will probably appeal more to those involved with

integrating end-user systems. Sessions will include "Graphical Tools for Regression Analysis" and "Graphics for Multivariable Data Analysis." Document storage and retrieval—of interest to business professionals as well as scientists—will be covered in a pair of sessions on micrographics. This will deal with the storage of documents on optical disks, which are among the most cost-effective storage media to date.

User-oriented topics will have a higher ratio of tutorials to technical sessions than will engineering-oriented subjects. A case in point—one of the major topics of the conference will be visual arts and design, with nine tutorials and five technical sessions. Tutorials will cover image synthesis software, and graphics concepts and techniques for the artist and designer. Tips on digital type design and the kinds of systems suitable for artists and designers will also be covered. Two technical sessions on computer animation will round out the topic.

Fujitsu: World-Class Components



Graphic/
Alpha-Numeric
Plasma Displays

PC-Board-
Type Relays

Bubble Memory
Cassettes and Cards

Flat-Cable
Connectors

Low-Profile
Keyboards

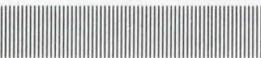
Mini
Thermal
Printers

Custom
Hybrid IC's

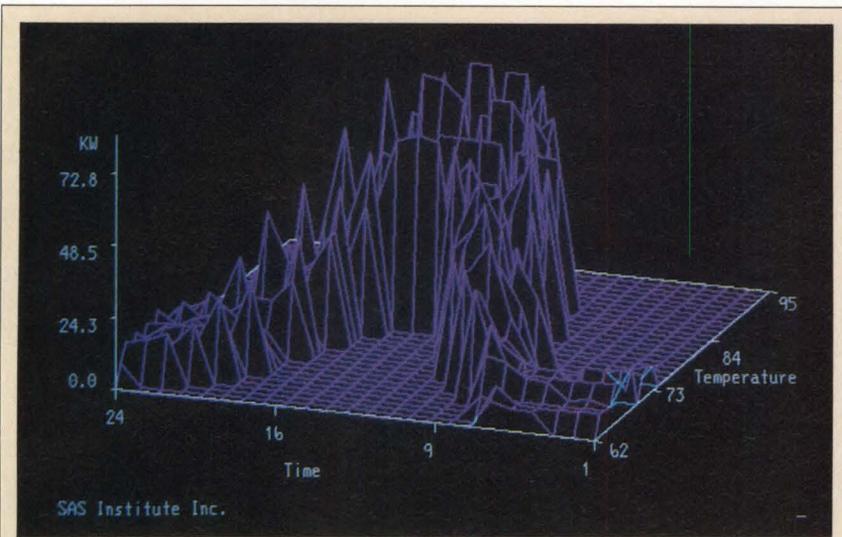
World-Class Components
Part of Tomorrow's
Technology



FUJITSU



Component Division Fujitsu America, Inc.
918 Sherwood Drive, Lake Bluff, IL 60044
(312) 295-2610 Telex: 206196 TWX: 910-651-2259



Complicated information can be better understood and managed using statistical graphics. SAS/GRAPH software from SAS Institute presents electrical demand as a function of temperature and time of day in an easy-to-comprehend form.

Other areas to be reviewed at the conference include sessions on legal issues involving copyright and owner-

ship, future directions in hardware and software, human factors and interfaces, and education and learn-

ing. Among the more technical topics are two sessions in image processing and pattern recognition, and video technology. The latter concerns the integration of computer graphics with video production. There will be additional sessions on video disk systems.

Well over 200 exhibitors have already signed up for Computer Graphics '85. In addition to the exhibits, there will be four days of exhibitor forums in which manufacturers and vendors will make detailed presentations on their products and services.

Advance registration can be made through the National Computer Graphics Association, P.O. Box 3412, McLean, VA 22103. A registration hotline (800/543-8000) is also available for quick telephone registration.

—Tom Williams,
West Coast Managing Editor

SYSTEM TECHNOLOGY
(continued on page 84)

World-Class Components Update:

MINI THERMAL PRINTERS

Sharp, Clean Graphics, Whisper Quiet Printing, And Unparalleled Reliability Make Fujitsu The Long Life Leader

Introduce yourself to the New Mini 20 and 40 column Thermal Printers from Fujitsu.

Their fixed head design, the result of our own advanced technology, offers many outstanding advantages. You get exceptionally clear printing and graphics for easy reading. High speed generation of up to 80 characters per second that prints complete lines at a time for high efficiency, and solid black for bold graphics. Virtually noiseless printing and last line visibility. And the whole printer comes in a compact, very lightweight package that meets the Centronics Standard, making it easy to interface with existing systems.



The result is a perfect printing component for applications such as in measurement and analysis equipment. In personal computers and P.O.S. terminals. And as an auxiliary printer.

For years, Fujitsu has been on the leading edge in developing breakthrough, problem-solving technology. Offering companies, worldwide, uncompromising quality and reliability—the result of Fujitsu's insistence on controlling, in house, every aspect of the design and manufacturing process. And *delivering*—the highest level of service and absolutely competitive prices. So no matter what your printing requirements, Fujitsu Mini Thermal Printers are the clear choice.

Find out more. Call or write Fujitsu today.



FUJITSU



World-Class Components
Part of Tomorrow's
Technology

Component Division Fujitsu America, Inc.
918 Sherwood Drive, Lake Bluff, IL 60044
(312) 295-2610 Telex: 206196 TWX 910-651-2259

CIRCLE 97

COMPUTER DESIGN/March 1985 83

Electro celebrates anniversary with a glimpse at tomorrow

As successor to a long string of IRE/IEEE Conferences and Exhibitions, Electro/85 (held with Mini/Micro Northeast) celebrates its tenth anniversary this year. Organizers estimate that more than 50,000 will view over 1200 exhibits. In addition, those who attend can sit in on 34 sessions of the combined technical programs, which comprise more than 170 individual presentations.

Electro/85 will convene in the New York Coliseum Tuesday, April 23, and will end Thursday, April 25. Concurrently, Mini/Micro Northeast will take place at the New York Sheraton Centre Hotel. Sessions slated for the joint technical program cover a wide variety of current topics.

Array of topics at Electro

On the Electro side, 24 sessions cover a gamut of high interest topics. Sessions lead off with "Applications of Speech Synthesis and Recognition Technologies," focusing on the latest developments in this form of human/machine interface. "ESD—Its Impact on Yield and Reliability" treats the important topic of electrostatic discharge, which has long been responsible for havoc in the industry. The session presents effective ways of controlling this pesky phenomenon. A panel discussion entitled "Venture Capital/Entrepreneurship" emphasizes issues that must be considered in forming a new technical company.

ICs will provide the focal point of several Wednesday sessions. "Semiconductor Logic—Today and Tomorrow" looks at IC developments and examines trade-offs required in deciding on gate-array, standard-cell, and field-programmable logic. "New Programmable Logic Yields Instant Custom ICs" describes these programmable logic devices as well as the hardware and software tools for their development. "Design Alternatives for Application-Specific ICs" notes the rising demand for application-specific ICs and questions the ability of manufacturers to handle the increasing number of IC designs expected in the near future. Speakers will present a few design alterna-

tives now available to system designers as well.

"Hard Copy Printing," on Thursday's docket, provides an overall review of nonimpact printing technologies. Specific papers on ink jet, electrophotographic, and thermal printing systems follow. The demand for portable personal computers has kindled a healthy interest in flat-panel displays mode of display. Operating characteristics of several such displays will be described in detail during the session, "Flat Panel Information Displays." An overview of an emerging technology, "Optical Data Storage," presents concepts and prospects, as well as design considerations, in optical disk and mass storage systems.

The Mini/Micro Northeast technical program consists of nine sessions, three on each day of the convention. Among the topics to be considered are the advantages of system-level diagnostics. One session provides several examples of diagnostic implementations, both in microprocessor-based and microprogrammed systems. "System Implications of Recent E² Advances," discusses system features and applications made feasible by the electrically erasable PROM's unique characteristics. Tools that facilitate the task of designing realtime software are the subject for discussion in a separate session.

Focusing on current trends in microprocessor unit/microprogram control unit development tools, one Wednesday session emphasizes cost versus performance when selecting the right one for each application. Another examines both unique and typical ways of using Unix to develop software for current and future microprocessors. "Personal Workstations for Electronic Design" reviews the emergence of the personal workstation and its applications in IC, standard cell, and semicustom IC design—as well as in computer aided engineering.

The final day at Mini/Micro Northeast brings "Thirty-two Bit Microprocessor Architectures," a discussion the architectural virtues of the MC68020,

Z80,000, Series 32000, and the IMS 424. "Thirty-two Bit Microcomputer Buses for the 80s and 90s" reviews the VMEbus, NuBus, Multibus II, Future Bus, and what lurks beyond the Q-bus. "High Performance Coprocessors and Controllers" points out the advantages these systems provide in off-loading the CPU for increased system capability.

Other special events

In conjunction with the IEEE Metropolitan Sections Activities Council, Electro/85 is offering a series of tutorial sessions. These will be given Monday, the day before the official show opening. Six day-long sessions (9 am to 5 pm) will take place at meeting rooms in the Sheraton Centre Hotel. Cost per session ranges from \$150 to \$175 for IEEE members and from \$190 to \$215 for nonmembers. Designed to help exhibitors and registrants keep pace with current high technology disciplines, the tutorials cover such subjects as artificial intelligence, computers and the Federal Communication Commission, engineering workstations, entrepreneurship, fiber optic applications, and speech recognition/synthesis.

Exhibits in the Coliseum will include active and passive components, test equipment and instrumentation, microelectronics, control systems, production equipment, packaging configurations and power sources.

At the Sheraton Centre, Mini/Micro Northeast will show exhibits that will be of special interest to OEMs, designers, system integrators, and software developers.

Registration at either the Electro/85 or Mini/Micro Northeast will admit the registrant to the other at no extra cost. Show hours are from 10 am to 6 pm daily, except for Thursday when the show winds down at 5 pm.

Further information on the meetings can be obtained from the organ obtained from the organizers: Electronic Conventions, Inc, 8110 Airport Blvd, Los Angeles, CA 90045. Tel: 213/772-2965.

—Jim Hughes,
Special Features Editor



Introducing a modem that speaks highly of you.

If you're an OEM looking to hook up with a modem supplier, we not-so-modestly suggest you talk with us.

E+E DataComm.

We're the company behind the increasingly popular Avatex™ line of modems.

Popular because they're designed with all the reliability OEMs, systems integrators and end-users alike have come to expect from a serious piece of communications equipment.

But that's only half the story.

Because behind *us* is Elec and Eltek, whose 15 years as a worldwide supplier of PCB and telecommunications products has taught us a thing or two about the OEM market.

Like how to develop a level of auto-

mation at both manufacturing *and* test stages so sophisticated as to dare approach zero defects.

That same penchant for perfection is reflected in our prices. They're competitive. Very competitive. And our delivery commitments—firm.

The result of all these efforts is an OEM support program that surpasses conventional programs to such a degree as to make them inapplicable.

So, if you want to put your reputation on a better line of communication, put it on E+E DataComm. It'll say a lot about you.



E+E DataComm

2115 Ringwood Avenue, San Jose, CA 95131
800-4-AVATEX



hp HEWLETT PACKARD Integral Personal Computer

Monitor for
HEWLETT
Line 3 Column 1 Length

To meet our needs, we've de
fined a pressure standard
an indicator of fatigue str
nath-matic components.

Although fatigue is the lea
structural failures such as
cracks, fractures, water
deflection and avoid break
result a pressure
loads.

Low stresses, how
the possibility of
components can caus
that increase stre

1	2	3	4	5	6	7	8	9	/
log	ln	FR	int	7	8	9	+		
1/x	10^x	e^x	x^y	4	5	6	*		
EXP	R 3	x79	ESP	1	2	3	-		

Monitor for
HEWLETT
Line 3 Column 1 Length

To meet our needs, we've de
fined a pressure standard
an indicator of fatigue str
nath-matic components.

Although fatigue is the lea
structural failures such as
cracks, fractures, water
deflection and avoid break
result a pressure
loads.

Low stresses, how
the possibility of
components can caus
that increase stre

1	2	3	4	5	6	7	8	9	/
log	ln	FR	int	7	8	9	+		
1/x	10^x	e^x	x^y	4	5	6	*		
EXP	R 3	x79	ESP	1	2	3	-		

Almost priceless.

Hewlett-Packard presents
the first 32-bit UNIX™ system
under \$5,000.

Introducing the Integral Personal Computer.

It's the newest member of the HP 9000 family of HP-UX workstations. And it delivers the kind of power and flexibility the name implies.

With a 16/32-bit MC68000. A graphics co-processor. And 800KB of standard memory, expandable to 5.8MB.

Its UNIX kernel (HP-UX/RO) is built into ROM. Which means the Integral PC can run most applications *without* an expensive hard disc. Which means more

people can afford to run your UNIX software than ever before. But not only can they afford to run it. They can learn to *use* it in

minutes. Because the Integral PC is one UNIX system that's easy to use. Its see-and-select user interface, built-in window manager and optional mouse mean even novices can tap the power of your software. No one has to learn a bunch of cryptic commands or go through a lot of expensive training. There's even an entry-level, self-paced tutor disc with every machine.

What's more, the Integral PC goes where you need it. The complete system — including our revolutionary ThinkJet printer — packs into a single 25-pound package that takes up less than a cubic foot.

So get started porting your software now. HP-UX is an enhanced version of UNIX System III, including Berkeley 4.2 BSD features.

For the name of the authorized HP dealer or HP sales office nearest you, call toll-free 1-800-FOR-HPPC.

The value of your investment in UNIX software just went up. Because the price of a complete UNIX system just came down.

The Integral Personal Computer. Just \$4,995*



**HEWLETT
PACKARD**

CIRCLE 35

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.

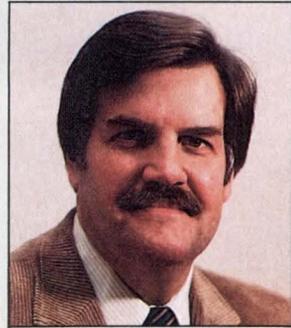
See how long you can keep your customers satisfied...with the long lasting, field proven printers that have earned the respect of OEM's nationwide—the Genicom 3000 family.

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





SPECIAL REPORT ON **CUSTOM/SEMICUSTOM ICs**

The numbers bandied about to describe the impact of application-specific, or user-designed, ICs on the semiconductor market don't always agree. Some say they will account for 50 percent of the IC market by 1988. Others don't believe they will account for 50 percent until 1990. It really doesn't matter which numbers are the most accurate; by any measure, the growth will be nothing short of phenomenal. ASICs will be a \$20 billion business some time near the end of the decade.

ASICs, whether they're implemented as semicustom chips (usually gate array or standard cell designs) or as full custom chips (not easy to define today), will be successful for several reasons. As everyone knows, they make it possible to shrink a product from the board level to the chip level, or at least shrink the size of a board by cutting the number of parts required. Moreover, they are a key—perhaps *the* key—to staying competitive. Not only because they allow a product to shrink in size or parts count, and hence, in cost, but also because they allow designers to be more innovative. And being innovative is more than the key to staying competitive, it's the key to staying *ahead* of the competition.

Innovation, however, is not an intrinsic property of ASICs, but rather it rests with the tools used to design them. Those tools should allow an ASIC designer to create a design quickly, change it at will, see how it works, throw it out if necessary, and start over again. Design tools should be easy to use and readily available. These qualities overcome the obstacles to a designer's creativity, such as drawing complex diagrams by hand, checking transistor physics, or waiting in line.

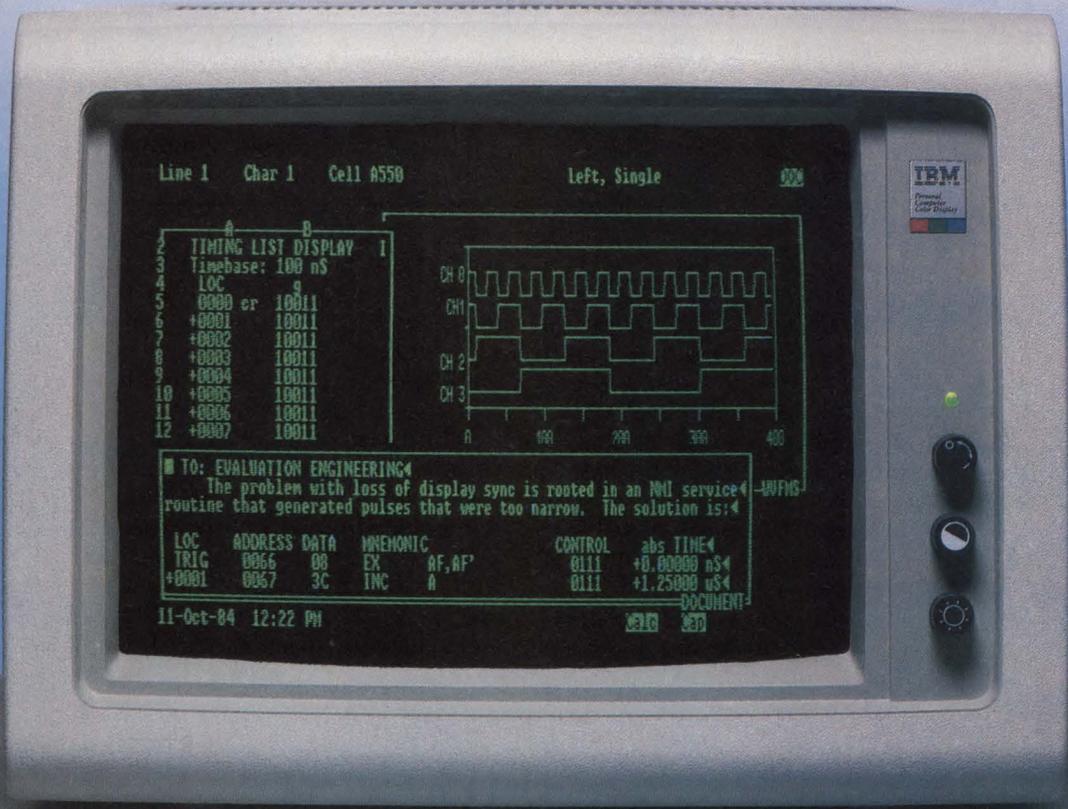
Because current-generation CAE/CAD tools, especially engineering workstations, have ended a lot of that drudgery and waiting, they have become the driving force behind the expansion of ASICs. These tools are now available in a range of capabilities and prices for just about any design budget. In addition, designers can gain access to the tools they need at regional design centers set up by semiconductor vendors or distributors.

Despite their power and availability, however, the now "conventional" design tools may be running out of steam. For one, designs are becoming more complex as designers try to pack more functions onto a single chip. For another, system designers now want to join the esoteric club of chip gurus. A new set of tools, based on the concepts of silicon compilation—the silicon equivalent of programming in a high level language—will allow designers without IC design knowledge to join this club.

A handwritten signature in cursive script that reads "John Miklosz".

John Miklosz
Executive Editor

Only our Logic Analysis Workstation can pass this screen test.



Why? Because it's the first system that combines logic analysis with an integrated software package on the new IBM PC AT*, PC XT* or IBM PC*. The μ Analyst LAW from Northwest Instrument Systems.

And what a logic analyzer! The μ Analyst brings you the industry's most powerful analyzer for under \$10,000. Up to 80 channels of state acquisition with 15 trigger levels. 16 channels of 100 MHz timing acquisition, both conventional and transitional. Correlated state and timing displays. 8, 16, and 32-bit disassembly.

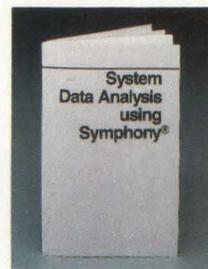
But that's only the beginning. Once you've acquired your data, the personal computer can go to work on it.

Using our Lotus-Link™ software, data can be loaded into Lotus Symphony™. Once there, use the resident graphics



and database functions for filtering, analysis and custom graphing. Or create co-resident programs for special applications. You can even place important data right into engineering documentation and send it to other computers using the integral word processor and communications packages.

How well does it work? Very well, indeed. Contact us for a free application note on how the μ Analyst LAW works with Lotus software to provide unique and valuable engineering insights. While you're at it, ask for free technical literature and a demonstration.



1-800-547-4445 Ext. 204

"For logic analysis, you can't beat the LAW!"

The μ Analyst LAW is also available for purchase, lease, rental or service from ElectroRent.

CIRCLE 37

IS NORTHWEST INSTRUMENT SYSTEMS, INC.

P.O. Box 1309
Beaverton, OR 97075

*IBM PC, IBM XT and IBM AT are trademarks of International Business Machines

• μ Analyst and Lotus-Link are trademarks of Northwest Instrument Systems. • Lotus and Symphony are trademarks of Lotus Development



DESIGN TOOLS PACE CUSTOM/SEMICUSTOM IC DEVELOPMENT

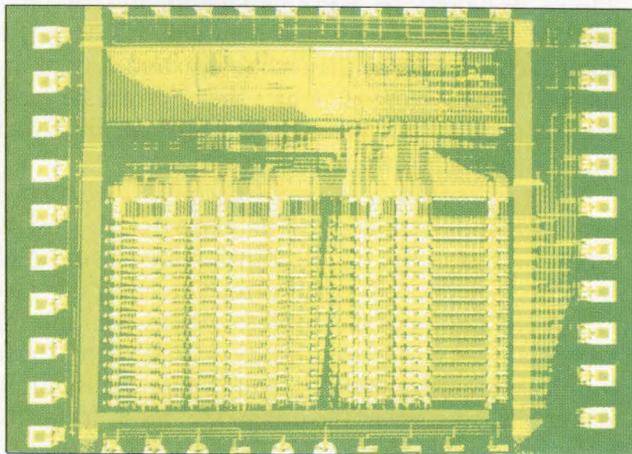
The penetration of application-specific ICs into systems depends on the availability of design tools that eliminate a designer's need to know the intricacies of silicon technology.

by John Bond, Senior Editor

Application-specific ICs are the new frontier for system-level design engineers. Industry forecasts claim that ASICs will account for more than 50 percent of the ICs in use by 1988. To a large extent, more sophisticated design tools are the driving force behind this spectacular growth. These tools have made gate array and standard cell design much easier. But, the need for greater sophistication has led to a new class of design tool and ASIC—the silicon compiled custom chip.

Although not yet as silicon efficient as handcrafted custom designs, silicon compiled chips offer the designer fast turnaround time and silicon efficiency that lies between standard cell and full custom chips. Most importantly, this is the first technique to open up custom design to engineers who are not experts in IC design.

Silicon compilation is roughly analogous to writing programs in a high level language instead of machine language. Just as compilers do not write code as efficiently as assemblers, silicon compilers do not generate masks that are as silicon efficient as full-custom circuits. Nonetheless, silicon compilers can cut design time dramatically. And, compiled chips are usually 25 to 40 percent more silicon efficient than either gate arrays or standard cells. Silicon compiled chips also boast a gate density that is about 90 percent that of a handcrafted, custom design. In view of the substantially lower design

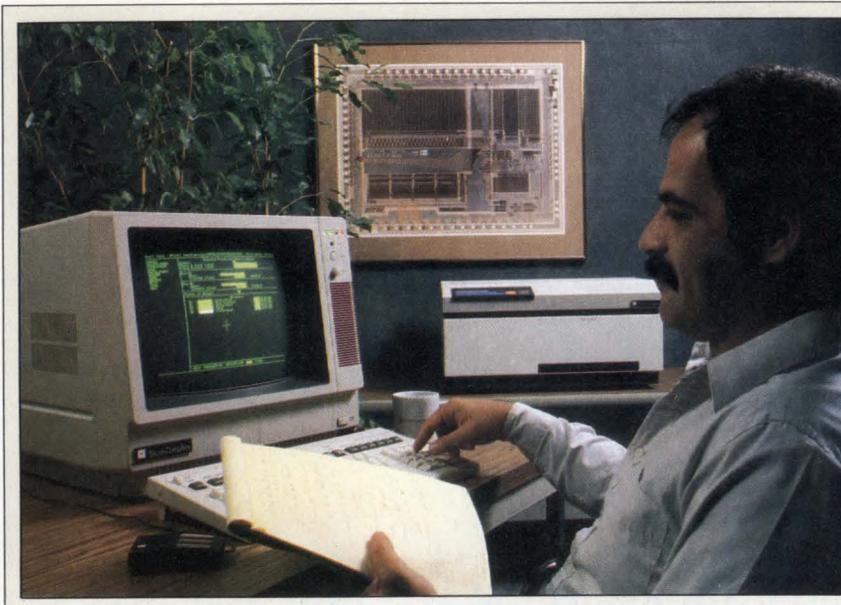


overhead (only 10 to 20 percent of that required for a typical custom design), the lower silicon efficiency seems a reasonable trade-off.

Full silicon compilation systems

Although various types of silicon compilers are appearing on the market, the Genesil system from Silicon Compilers, Inc (San Jose, Calif) is the first full silicon compilation development system. It enables system-oriented design. Early versions of this system were used to design several complex chips, including a Seeq Ethernet controller, a Sun Microsystems dynamic raster controller, and the Digital Equipment Corp 32-bit MicroVAX I.

The Genesil silicon development system consists of a VAX-11/750 with 4 Mbytes of memory, a



Silicon Compilers' Genesil silicon development system allows system designers to produce VLSI chips. It uses the structural representation approach to silicon compilation.

450-Mbyte disk drive, tape storage, and four high resolution color terminals. System software provides IC definition, functional simulation, timing analysis, placement and routing, and tooling. Tooling activity allows designers to work independently of technology or foundry until the design is ready for manufacturing. Once the technology and vendor are specified, the system generates the appropriate mask tooling tape output in GDSII or CIF formats with specific fabrication-line information.

More importantly, Genesil permits system engineers without semiconductor expertise to design, analyze, and verify VLSI chips created using high level architectural descriptions. The compiler also allows the design of a complete system even though it may take an entire set of chips to implement. Partitioning system functions among those chips can be accomplished early in the design and modified as the design proceeds.

Six categories of ICs can be designed with the system: memory, complex logic, data path, random logic, pads, and test. Memory includes RAM; ROM; first in, first out buffers; and stacks. Programmable logic arrays, decoders, and encoders are some of the complex logic functions. Data paths can be selected from 4 to 32 bits. Random logic functions provide a gate level access that is normally hidden from the designer when working at a higher level. These functions permit users to build NAND and NOR gates or flipflops. The pad function accomplishes offchip connection, and test functions permit test blocks to be included in the IC.

A silicon compiler allows designers to begin with an exploratory development phase before proceeding to detailed development. Using the design and verification feedback, the user can proceed iteratively.

Given even sketchy detail, the system responds immediately to incremental changes and provides size, speed, and power estimates. Functional simulation and timing analysis permit design verification. Tight coupling between these functions and layouts lets users examine behavior and performance without using several different computer aided design tools.

All silicon compilers are not the same, however, which is not surprising given the early state of the art. There is no general agreement about what constitutes a silicon compiler. For example, the silicon compiler from MetaLogic, Inc (Cambridge, Mass) and Silicon Compilers' Genesil system take different approaches. MetaLogic's MetaSyn Lisp-based compiler is very software-oriented whereas the Genesil compiler has a strong hardware orientation. The MetaSyn is derived from the McPitts compiler developed at the Massachusetts Institute of Technology's Lincoln Laboratory. MetaSyn is written in Lisp because the Lisp workspace environment allows interactive debugging and development with fewer steps. Lisp also is extensible. MetaSyn, however, has gone through considerable development to add features, such as performance prediction, and to double the efficiency of the McPitts compiler.

Running on a Symbolics 3670 computer, with 296 Mbytes of disk memory, MetaSyn develops circuits from behavioral specifications. The output of that behavioral specification varies, however. In the logic design stage, the behavioral specification is converted into a structural description of logic cells and their interconnections. Specification of AND/OR gates, flipflops, and their interconnections in gate arrays and standard cell designs, for example, is made at the logic design level. Mask layouts are then derived from the logic design.

Logic design can be performed at other levels in the design hierarchy, however. In structural compilation, for example, the circuits that represent the logic design are predefined and the system goes through automatic layout algorithms. This approach offers a significant advantage.

In many custom IC designs, changes must be made at the mask layout stage, rather than going back to the circuit or, even further, to the logic level. If the logic is modified, the layout that has already been done will be lost. Returning to the behavioral level causes the loss of even the logic. This state of affairs has forced IC designers to manipulate the layout because going all the way back to the behavior specification was uneconomical. Structural compilation of gate arrays and standard cells, coupled with auto placement and routing, makes it possible to return to the logic and not have to tweak the layout. If the engineer can manipulate at the logic level and automate the circuit and layout steps, those steps do not have to be done over to generate the mask again.

Software-oriented input formats typify the behavioral approach to silicon compilation; the structural approach emphasizes hardware.

MetaSyn automates logic design, as well as circuit and layout designs, so that the designer can go back to the behavioral level and manipulate them. At the behavioral level, the designer can explore the speed, performance, design time, and fabrication cost trade-offs quickly.

Rather than lower the cost of full-custom design (although that may well be a benefit), MetaLogic aims to raise the performance of system specifications. The company's approach is to take a program that might be written for a microprocessor and put it in silicon for a great increase in performance.

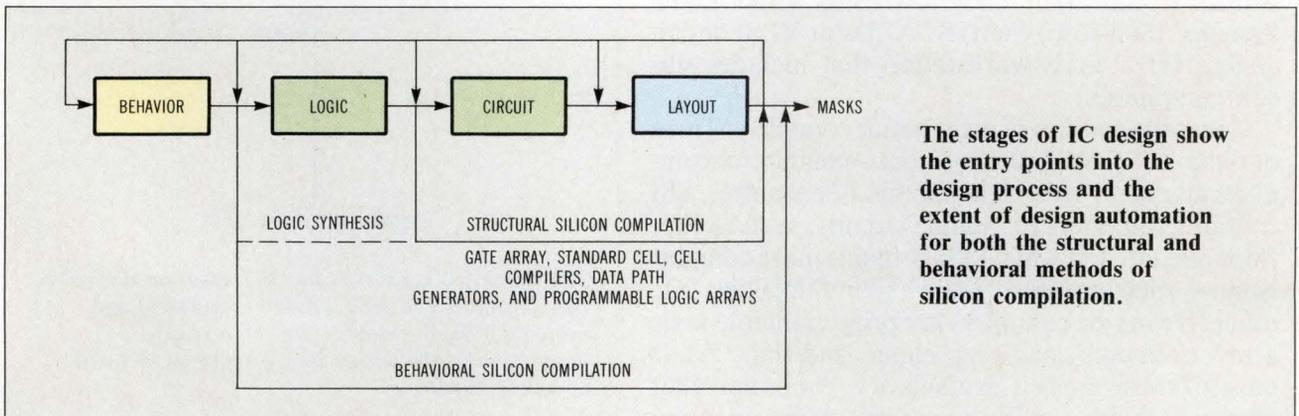
MetaSyn puts programs in silicon, not with a memory image in ROM, but by creating a circuit that actually performs the function. Typically a MetaSyn compiled chip has fewer transistors, but is less dense, which spreads the circuit over more chip area. It is more open in terms of transistor density, but not in terms of functional density.

Perhaps the ability to manipulate the behavioral level provides the greatest leverage over the design. There are, however, more choices to explore at the logic level than there are at the behavioral level. In theory, the ability to make more modifications at any of these stages permits more flexibility. In practice, however, any given amount of time spent in logic design cannot be spent modifying behavior.

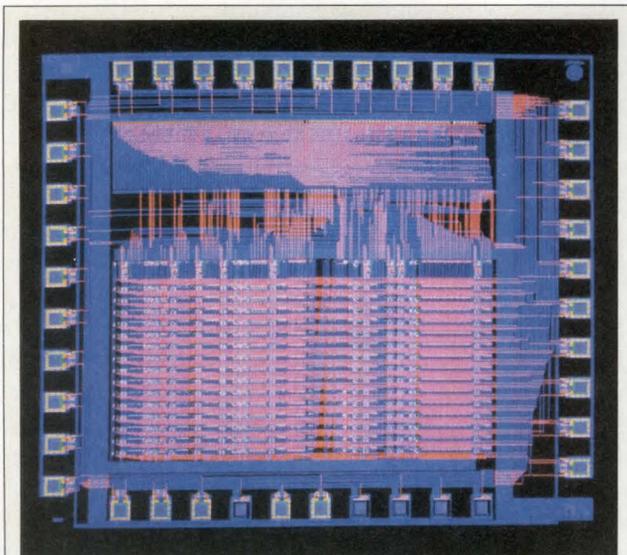
Structural versus behavioral design

Structural representation denotes the logic design approach, as illustrated in the Gaijski-Kuhn "Y" diagram. The designer moves up the axis to gain a more concise and powerful design. At the same time, however, the design loses flexibility and detail. The register transfer level is concerned with registers, adders, ALUs, and their interconnections. The highest level is system hardware, such as processors, memories, multiplexers, and systolic arrays. Silicon Compilers' product works along this axis, while MetaLogic's product works on the behavioral representation axis. For MetaSyn, the lowest levels are Boolean logic expressions and algorithms (program statements) are the next level up. The highest design level is the system I/O specification. The geometrical axis is concerned with the actual physical and geometrical structure of the chip.

MetaLogic starts with a functional specification that describes input, process, and output algorithmically. From there, the register-transfer expression is automatically derived. A Boolean logic expression that controls the data path is derived together with the register transfer to implement the algorithm. The behavioral representation automatically generates the appropriate structural representation and the masks.



The stages of IC design show the entry points into the design process and the extent of design automation for both the structural and behavioral methods of silicon compilation.



This 16-bit microprocessor with a stack-oriented instruction set was designed with the MetaSyn silicon compiler. It can be converted to a 32-bit processor by changing a single number on the first line of the program.

Structure is an explicit description of data flow, but does not describe the control flow. According to MetaLogic, if there is no explicit description of the control flow, a number of registers may be interconnected at various times. Without knowing the timing, there is no way of knowing what that circuit is doing. Algorithms are composed of steps occurring in time sequence. Each step is a substructure of register-transfer functions that are used to derive a time-multiplexed static hardware structure.

Low cost compilation

Whether system designers will prefer software-oriented input formats or a hardware-oriented structural approach is an open question. Another silicon compiler—Concorde from Seattle Silicon Technology, Inc (Bellevue, Wash)—has taken the structural route, but with a different approach than Genesil. Concorde has been integrated into Valid Logic Systems' (San Jose, Calif) SCALDstar VLSI design system for a CAE workstation that includes silicon compilation.

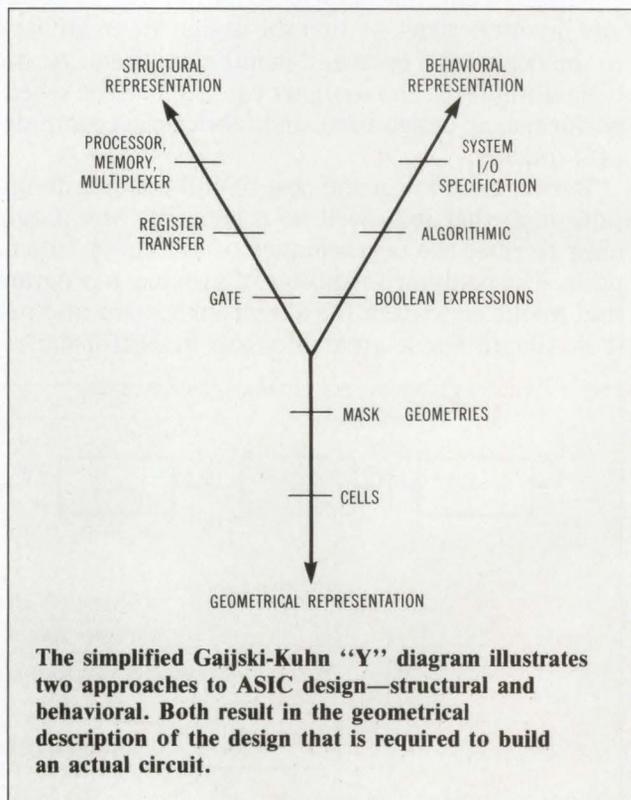
Concorde consists of eight major compilers. Three of these—SSI, MSI, and memory—contain subcompilers at a lower level of hierarchy. For example, SSI contains compilers for simple circuits, such as flip-flops and latches, while MSI contains more complex circuits, such as shift registers and multiplexers. The other five major compilers are programmable logic array, data path, pad ring, elmer, and flair. While most of these are self-explanatory, elmer and flair deserve further description. The elmer compiler

automatically generates SSI and gate level cells and connects them according to the user's specifications. The flair compiler uses a tabular code file to generate a folded logic array with internal registers.

Concorde's organization makes it more like a sophisticated set of cell compilers rather than a true silicon compiler. According to its critics, Concorde is not system-oriented enough to be useful to system designers who do not have IC design experience. While it is true that the Concorde approach may not be as "system-level" as the MetaSyn approach, the point of entry into the design is very flexible and can be at a high system-oriented level. Concorde, for example, can use behavioral descriptions, truth tables, Boolean statements, or logic diagrams as a starting point.

As the design is decomposed into its module compilers, many specifications must be defined by the user to complete the design of each module. The chip is then laid out and the modules are connected using the Valid workstation's layout editor. Finally pads are connected using the pad ring compiler. Interactive simulation allows the designer to test the operation of the circuit or portions of it.

Although Concorde requires more user interaction to design a chip than either Genesil or MetaSyn, it is a valuable design tool. This silicon compiler has the lowest price of any such systems on the market. A Valid SCALDstar workstation with Concorde costs around \$100,000 compared to \$210,000 for a



11,000 GATES PER GRAM.

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.

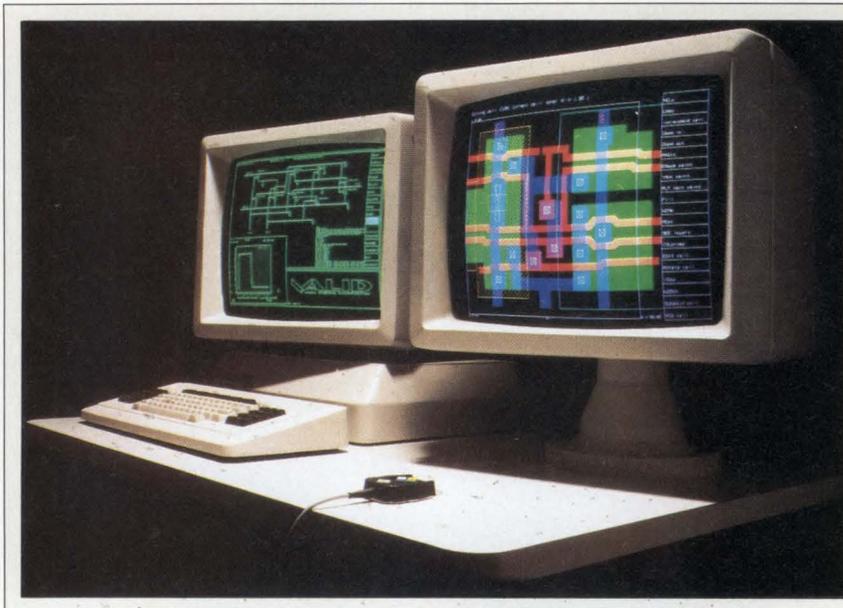
So next time you need a gate array, weigh your alternatives. Then call NEC at 1-800-556-1234, ext. 188. In California, call 1-800-441-2345, ext. 188.

WE'RE TAKING ON THE FUTURE.



© 1984, NEC Electronics Inc.

NEC
NEC Electronics Inc.



The Valid Logic SCALDstar VLSI design system uses the Concorde I silicon compiler to provide automatic translation of logic designs into custom chips.

MetaLogic system and \$545,000 for the Silicon Compilers system. Optional analog and complex function compilers are also available.

One other advantage to Concorde is that it is a CMOS compiler. Both Genesil and MetaSyn were introduced as NMOS compilers. A CMOS version of Genesil should be introduced early this year. The CMOS version of MetaSyn is not expected until the end of this year.

The real world

So far, silicon compilation has only been responsible for a handful of designs. Most of the recent designs, however, have been implemented in gate arrays. Gate arrays may continue to lead in the market for several reasons. Design to production times are generally (but not always) shorter. Gate densities have improved dramatically and there is an emerging trend to put memory, microprocessors, and other standard circuits on gate arrays. Essentially, they will be macrocells with larger, more complex macros than are now available. But there are some other advantages to gate arrays.

While the majority of standard cell vendors have a single-level CMOS process, many gate array vendors offer a double-level process. In gate arrays, the layout designer need only be concerned with the capacitance of the metal. The metal's resistance can be ignored for practical purposes. The dominant single-level metal processes in standard cells make resistance of polysilicon and the distribution of that resistance a large concern for designers. Therefore, confidence in standard cells does not run as high as it does for gate arrays in terms of working to specification on the first pass. So, it is easier for designers to turn in a netlist for a gate array.

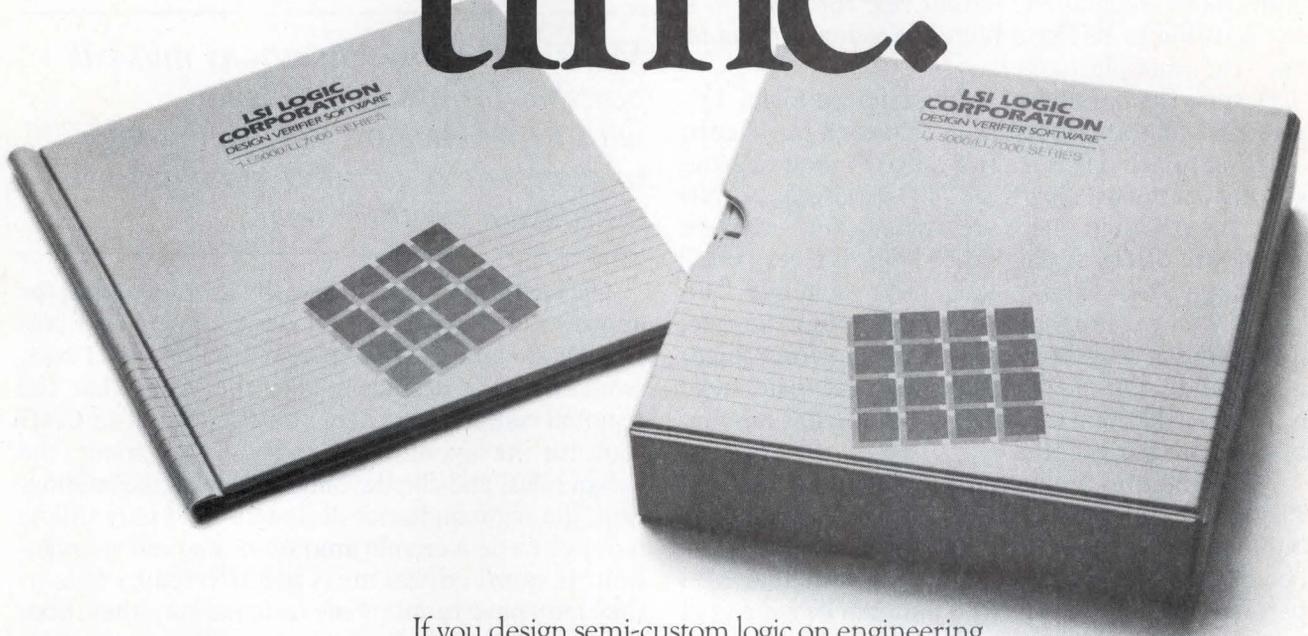
Standard cells, however, offer more flexibility than gate arrays, making it easier to change the cell geometry, and thus complicate the layout. Of course, the gate array designer has to get involved at a gate level or at least with fairly small macros. The standard cell user, on the other hand, can pick from cells that approximate standard logic devices. In addition, vendors are developing two-layer metal CMOS processes and superior computer aided design tools. This is reflected in projections for worldwide standard cell market growth, which Dataquest, Inc pegs at \$1.2 billion by 1989.

The acceptance of both gate arrays and standard cells depends largely upon advances in CAD tools. The greater complexity of standard cells makes even more demands upon CAD tools than gate arrays. Consequently, semiconductor manufacturers have had to provide CAD tools and set up regional design centers to aid system designers as they make the transition to semicustom design. Examples of such design centers are the regional technology centers developed by Texas Instruments. TI offers design facilities for its own bipolar gate arrays and CMOS gate arrays under a second source agreement with Fujitsu, Ltd. The regional technology centers provide design support for TI's SN54/74SC standard cell family. The 3-micron CMOS cell library contains 186 functions and 230 cells and gate speeds fall somewhere between low power Schottky and Schottky.

Located for convenience

TI implemented the CMOS standard cells in 7400 type logic to make it easier for the designer. For example, a user can go to the technology center with a 7400 series TTL schematic, copy it directly on to a Daisy or Mentor Graphics system using the resident

Get it right the first time.



If you design semi-custom logic on engineering workstations like Daisy,* Valid* and Mentor,* we can help you go from scratch to silicon a lot faster.

That's not just a promise.

It's our track record.

We're LSI Logic Corporation. Our CAD software has guided more than 600 gate array designs from the drawing board to the assembly line.

With a 95% hit rate, first time out.

Just see your Wyle or Hamilton/Avnet distributor. They stock our Software Data Book™ macrocell libraries and Design Verifier Software™ both of which can turn your workstation into a powerful design center.

See a demonstration soon.

And start turning hit-and-miss into a hit.

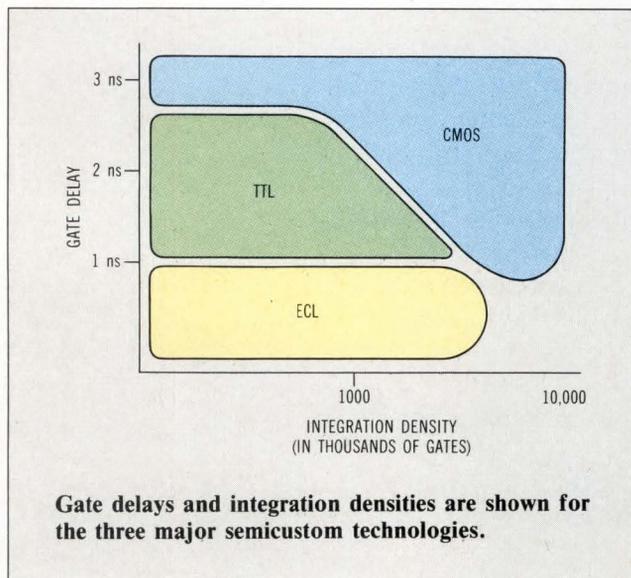
**LSI LOGIC
CORPORATION**

TI library, and run all the necessary simulations. The user need not know anything more than how to design with standard parts. The standard cell library is called up by the Daisy or Mentor utilities and a schematic is built onscreen. Then the user runs the simulator for the logic family. After the logic check, the user can see if any of the nodes are too heavily loaded down. If they are, the user can design in bigger cells until the timing is right. Daisy, in particular, has a program that can report the minimum and maximum delays through any gate, as well as flag timing errors.

With modern workstations, anyone who is capable of designing with TTL can do a standard cell design. But first there are a few rules that the user must understand, such as the fanout rule for gates. It is not possible to buffer a big cell from a very small one, for example.

Despite the parallelism with standard logic, TI's standard cells are different in that each one can be modified. Each one is really a range of product, with varying parameters. With the 7400 quadruple NAND gate, for instance, there are actually five different 7400s with different data sheets. One is a very small NAND gate that has a minimal drive capability. This causes a longer propagation delay through the gate, if it is driving a significant load. But the form factor is very good. On the opposite end of the scale, there is a large 7400 that can drive a significant amount of current into a node.

MSI functions are handled through soft macros. The basic components are SSI blocks. MSI parts are built from SSI components and soft macros that can be edited. If a user needs half of a particular MSI part, for example, the unused part can be edited out or used elsewhere in the design. Thus, there are no wasted gates and everything is fully utilized.



Anybody with a Daisy or Mentor system can design TI's standard cells, but many designers prefer the technology centers to do it for them. TI admits, however, that the best designs occur when the customers do the work themselves. Since the circuit designers know what the chip must be able to do, they usually get exactly what they want when they design it themselves. On the other hand, the engineer who comes to the technology center with vague schematics and test patterns runs the risk of losing something in the translation. Thus, communication with the user is vital for good results. Getting the design specs from system designers who have never built semiconductors before is the biggest hurdle in front of the regional technology center personnel.

Modern CAD workstations make it possible for system designers with experience in TTL technology to complete a complex standard cell design on their own.

The technology centers usually keeps a design for about four weeks. Then the design is sent to Texas for layout and post-layout simulation. In Texas, semiconductor designers take the input that the regional center gives them and translate it to CAD tools for the layout. They verify the schematic, the design rules, and chip layout. Despite all the automation, the semiconductor designers at TI may still be required to do a certain amount of manual intervention. If speed critical areas are affected by layout, they may have to intervene to make sure that those blocks are placed in the optimal position on the chip.

After the chip is laid out and just before it goes into production, a post-layout simulation is performed. All of the capacitances and the resistances of the wires and the polysilicon interconnects on the chip are fed back as delay parameters into the simulation data base. Then, the chip is resimulated. If there is enough impact to skew any of the signals and violate the spec, the chip is redesigned. The whole iterative process takes four weeks in the regional technology center and an additional seven weeks to the start of production. Prototypes will be available five or six weeks after that. Full production begins in an additional five or six weeks for a total of 21 to 23 weeks from the beginning of the design.

Other semicustom choices

There are numerous approaches to semicustom design at manufacturers' design centers. Although these methods have a lot in common, each must be adjusted for industry standard CAD workstations,

CITIZENS FOR BETTER READING.

OEMs and designers, Citizen™ is into LCD displays. Not just *any* LCDs—the most readable ones you've ever laid eyes on.

Thanks to higher contrast ratios. Wider viewing angles. Less shadowing. Faster response times. And just plain better readability.

The same people who give the world precision-engineered watches now give the world something just as reliable to watch.

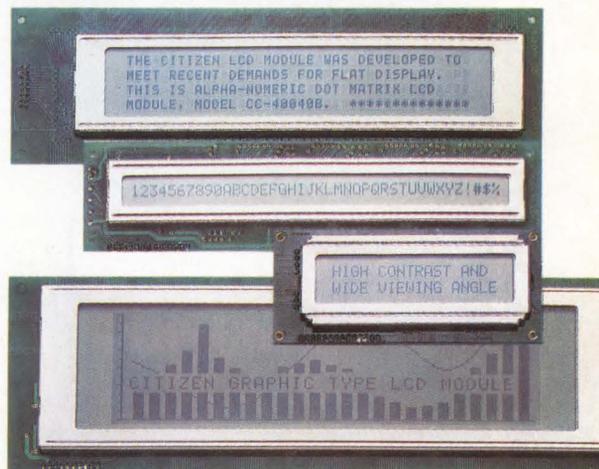
Pick an application, we'll show you the right LCD.

Invent an application, we'll customize the right LCD. And then back it up with a worldwide support and service network that's second to none.

The complete LCD line of alphanumeric and graphics displays—now from Citizen.

Read all about it. For details, a demo, or

information about our 3½" floppy disk drives and dot matrix printers, call 1-800-556-1234, Ext. 34. In California, 1-800-441-2345, Ext. 34. Or write Citizen America Corporation, 2425 Colorado Avenue, Santa Monica, CA 90404.

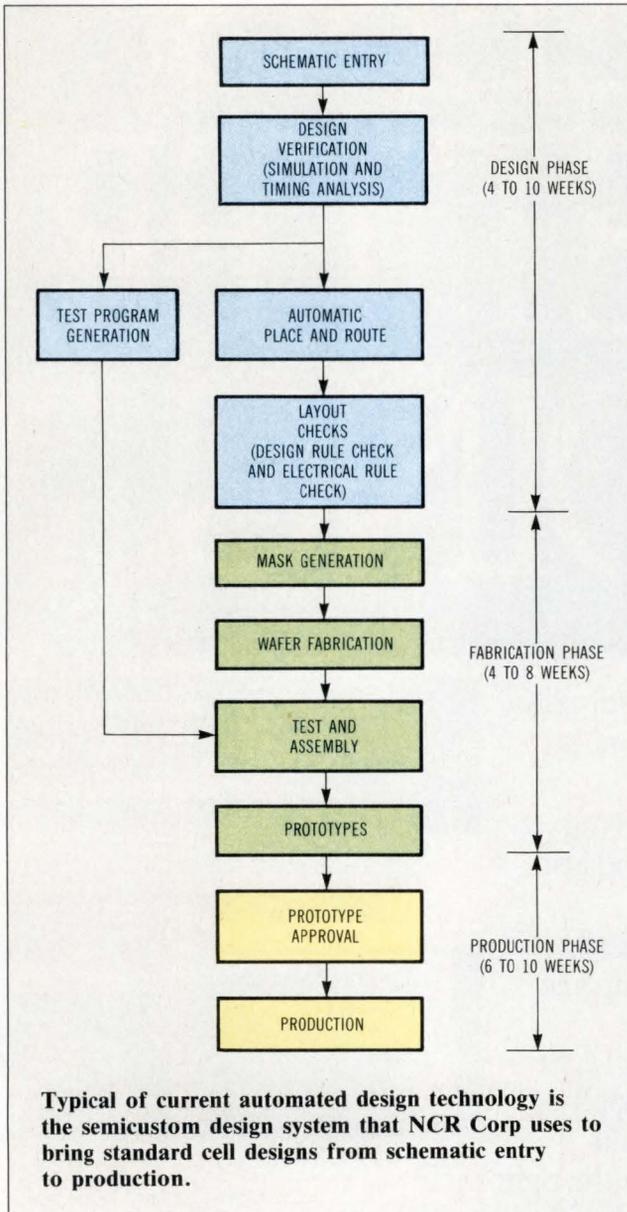


Character Display Modules		
Character Format	Module Size W x H x D (mm)	Viewing Area W x H (mm)
16 x 1	85 x 41 x 8.65	63.0 x 13.0
16 x 2	84 x 44 x 8.65	64.0 x 20.5
20 x 1	105 x 39 x 8.65	84.0 x 13.0
20 x 2	110 x 49 x 8.65	84.0 x 20.5
40 x 1	195 x 47 x 8.85	163.5 x 13.0
40 x 2	195 x 50 x 8.85	163.0 x 21.0

Graphic Display Modules		
Dot Format	Module Size W x H x D (mm)	Viewing Area W x H (mm)
240 x 64	263 x 84 x 9.00	210.0 x 54.0
480 x 64	267 x 80 x 10.40	228.0 x 46.0
480 x 128	279 x 110 x 10.30	228.0 x 74.0
480 x 200	270 x 124 x 12.60	203.0 x 88.0
640 x 128	280 x 88 x 11.10	230.0 x 53.0
640 x 200	290 x 118 x 11.10	230.0 x 76.0
640 x 200	270 x 230 x 12.60	230.0 x 174.0



Citizen and the Citizen logo are trademarks of Citizen Watch Co., Ltd.



various design tools, and each manufacturer's design tools and languages. NCR Corp, for example, performs the design service much like TI. Yet, there is little similarity in either product or design tools.

NCR Corp has an inhouse design center and seven regional design centers. The design centers can take total responsibility for the design or let customers use either the design center's workstations or their own workstations. The NCR design and verification system is available for Daisy, Mentor, and Valid workstations. The complete design package includes the workstation standard tools for schematic entry and functional simulation along with NCR timing analysis software, digital and analog cell libraries, digital supercells, and core microprocessors. Placement, routing, and layout are generated at NCR along with masks and prototype chips.

When comparing the TI and NCR approaches to standard cell design, it is apparent that the CAD tools of each vendor are quite different. The major similarity, of course, is in the standard software tools that exist for each workstation. But even the workstations, although functionally similar, use different software. This is because each semicustom vendor and workstation manufacturer had to develop CAD tools where none existed to solve their own problems. The fact that many of the approaches are similar merely reflects a consensus on how the job should be done as the technology begins to mature. Nonetheless, the software tools remain incompatible. As a result, IC designers who wish to use a combination of CAD tools from different vendors must first convert various software formats before designs can be completed and tested. The growing use of workstations to develop complex custom and semicustom ICs has heightened the need for a standard interface language.

Getting together

To solve this problem a committee was formed in February 1984, to direct the development of a standard Electronic Design Interchange Format (EDIF) that will allow different products to communicate with each other. The steering committee consists of seven representatives from organizations representing semiconductor vendors, CAE vendors, OEMs, and the academic community. Members include Texas Instruments, Daisy Systems Corp, Mentor Graphics Corp, Motorola Inc, National Semiconductor Corp, Tektronix, and the University of California, Berkeley. The EDIF committee has invited outside participation and during recent technical reviews more than 40 companies have participated.

Several existing languages are being used to help define the new standard. TI has contributed specifications for its totally integrated design-automation language. Additional languages being analyzed for applicable features include a technology definition file from Motorola and Mentor Graphics, as well as Daisy Systems' gate array interface design format from National Semiconductor, Tektronix, and the University of California, Berkeley.

The new language will provide a standard interface between the many types of hardware and software used to design, test, and manufacture ASICs, such as logic arrays, standard cell ICs, and full-custom chips. Adoption of EDIF will permit vendors of CAD equipment to provide compatible products, while allowing designers to readily access a variety of tools. By being placed in the public domain, the standard language will be available to any company or group that wishes to implement it.

The problem of design tools with incompatible formats exists in companies other than those that manufacture workstations or ICs. Early adopters have had to develop their own software to cope with ASIC development. While this provides a coherent design environment for the company's engineers, it introduces still another incompatible format that has to be translated.

Using workstations to develop complex custom and semicustom ICs has heightened the need for a standard interface language.

Digital Equipment Corp, for example, has a wide variety of inhouse design tools. When DEC starts using a vendor, it may use the vendor's design tools for the first few designs. Eventually DEC engineers build translators or bridge tools to link the vendor's and their design systems. The majority of designers within the company use DEC tools and provide each vendor with a netlist that can be compiled on the vendor's system. DEC also has simulation vectors to verify the design before it goes to the vendor.

A standardized language may not be important for large companies, such as DEC, that have the resources and have already been through the pain of establishing standardized design tools inhouse. But it is vital for the future of ASICs that the EDIF committees establish an industry standard.

Even with advances in standardizing IC design languages, the ultimate design tools are silicon compilers, and several different approaches are available. The algorithmic input design approach will probably be met with initial resistance from logic designers. It may be an idea ahead of its time that early adopters will like. The logic approach may be the most accepted. Even today, however, convincing engineers to design with semicustom logic instead of standard logic is a problem. It will probably take some time before the algorithmic approach is widely accepted. Methods that provide an easier migration for logic or system designers will probably be more attractive, at least for the short run.

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

EMI PROBLEMS?

Who You Gonna Call?



EAGLE MAGNETIC CO., INC.

Magnetic Shields (CRT-XTM-PM) Sheet & Foil

317-297-1030

P.O. Box 24283 Indianapolis, Ind. 46224

CIRCLE 41



SAIC
VME MICROSYSTEMS DIVISION
"LEADING THE INDUSTRY IN VME SOLUTIONS"

VMEbus DIGITAL I/O MODULES

DESCRIPTION	MODEL NO. SAIVME-	LIST PRICE**
• 128-BIT TTL I/O TURBOMODULE™	2528	\$ 995
• 64-BIT TTL I/O MEGAMODULE™	2510	\$ 695*
WITH BUILT-IN-TEST		
• 64-BIT HIGH VOLTAGE OUTPUT MEGAMODULE™	2120	\$ 695*
WITH BUILT-IN-TEST		
• 64-BIT HIGH VOLTAGE INPUT MEGAMODULE™	1110	\$ 695*
WITH BUILT-IN-TEST		
• 64-BIT CONTACT BOUNCE ELIMINATOR	1130	\$ 695
• 64-BIT NEGATIVE VOLTAGE RELAY MODULE	2132	\$ 695
• 64-BIT OPTICALLY COUPLED INPUT MODULE	1150	\$ 895
• 48-BIT PIO MODULE WITH INTERRUPTS	2511	\$ 695
• 32-BIT OPTICALLY COUPLED INPUT WITH CHANGE-OF-STATE INTERRUPTS	1160	\$ 795
• 32-BIT TTL INPUT MODULE WITH CHANGE- OF-STATE INTERRUPTS	1101	\$ 595
• 32-BIT HIGH VOLTAGE DIGITAL INPUT WITH CHANGE-OF-STATE INTERRUPTS	1180	\$ 595
• 32-BIT OPTICALLY COUPLED OUTPUT MODULE	2170	\$ 795
• 32-BIT RELAY OUTPUT MODULE	2200	\$ 895
• 32-BIT INPUT/32-BIT OUTPUT-HIGH VOLTAGE I/O WITH BUILT-IN-TEST	2532	\$ 595

* ADD \$100 FOR 32 BIT TRANSFER OPTION

** PRICES SUBJECT TO CHANGE WITHOUT NOTICE

ONE YEAR WARRANTY
DELIVERY-STOCK TO 90 DAYS

VME MICROSYSTEMS DIVISION
SCIENCE APPLICATIONS INTERNATIONAL CORPORATION
2109 W. CLINTON AVENUE
HUNTSVILLE, AL 35805
(205) 533-5900, X338
CARROLL WILLIAMS

*Extruded PVC jacket
required for
abrasion resistance*

*Must have 360°
EMI/RFI protection
and 100% continuity*

*Requires 85%
braid shield and
360° aluminum
mylar shield*

*Requires four
8-position
breakouts*

*Need 32-position
input connector*

*Need 93-ohm
planar coaxial
cable*

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 interconnect 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 1/4" braid terminated with 1/4 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:
Lisle, IL 60532



ALLIED Amphenol
Products

DIGI-PAD® WORKS WHERE YOU WORK

DIGI-PAD digitizer tablets perform in the real world. They won't quit over spilled coffee! They are unaffected by dirt, graphite, moisture, or perspiration on the tablet. And DIGI-PADs are insensitive to acoustic noise, magnets, vibration, or pressure. In fact, DIGI-PADs will work in almost any environment you will. That's because they use our reliable absolute electromagnetic sensing technology. And GTCO is the largest producer of digitizers using electromagnetic technology.

Most other digitizers are not as tolerant as a DIGI-PAD. So make sure that the digitizer you choose won't fail or require adjustments because of heavy use or environmental factors.

DIGI-PADs are designed for years of silent, maintenance free operation. You will find the DIGI-PAD comfortable to work with because the stylus generates no heat or acoustic noise. And there's no need for special handling of magnetic storage media because GTCO digitizers will not erase floppy diskettes.

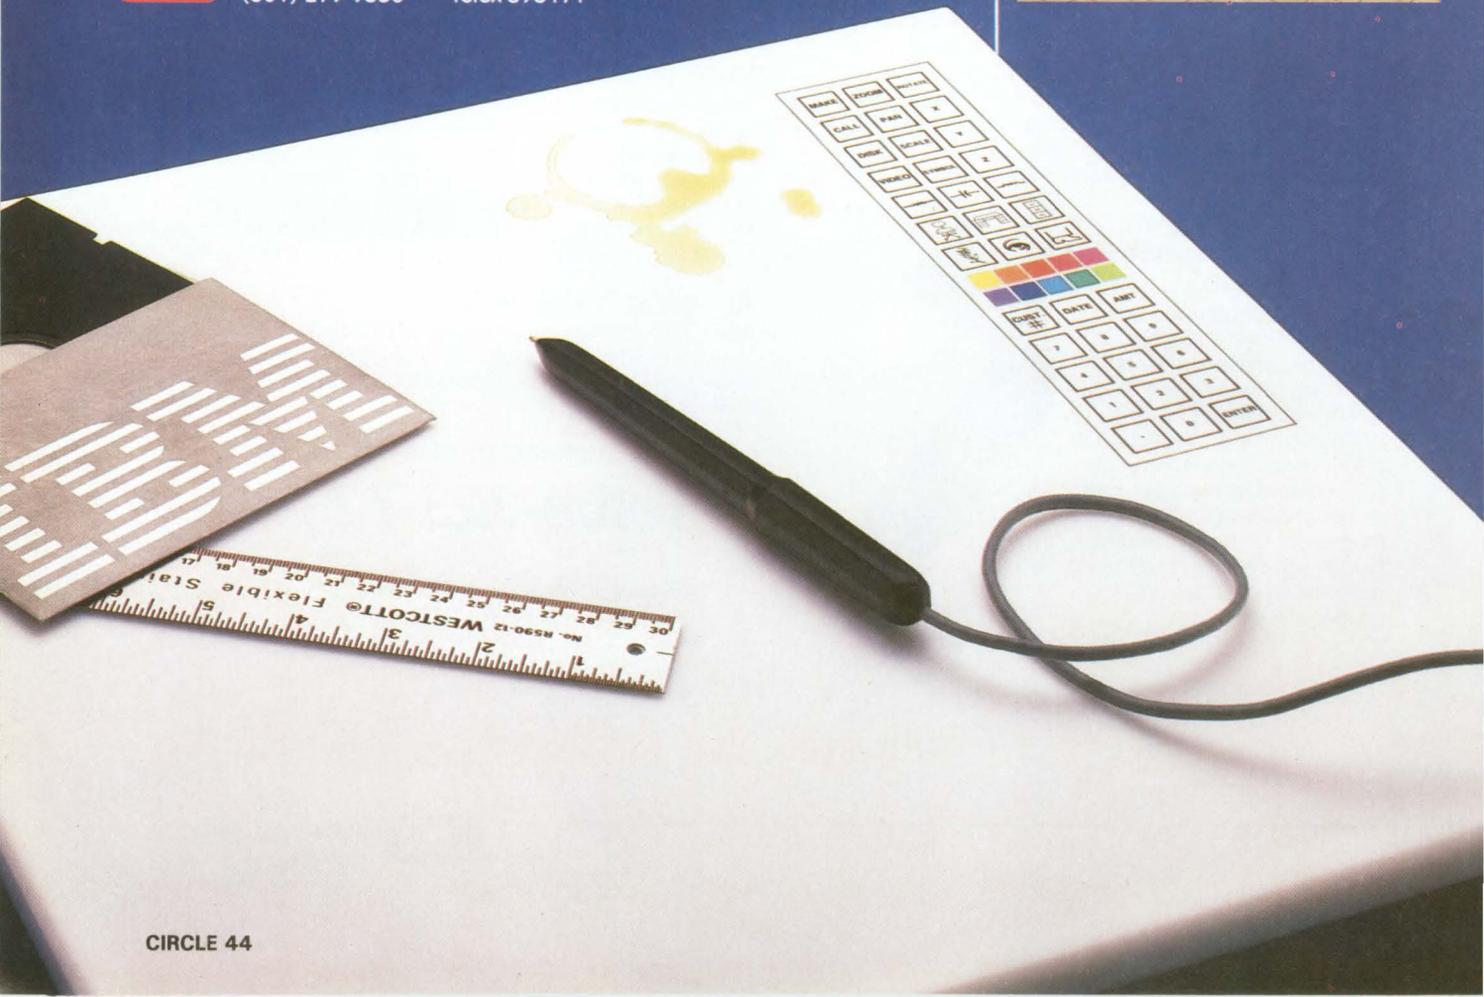
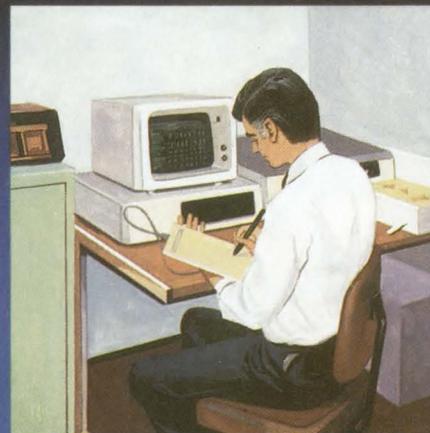
DIGI-PADs are available in sizes from 6" x 6" to 42" x 60", all using the same patented electromagnetic technology. They have been field proven in military, industrial, technical, business, and educational environments. DIGI-PAD is compatible with nearly all computers, from PC to mainframe. Most models are in stock for quick delivery and all comply with pertinent EMI and safety standards.

Choose a digitizer tablet that's willing to work where you work. Choose DIGI-PAD from GTCO (pronounced Gee Tee Co). Call one of our digitizer specialists today at (301) 279-9550.



GTCO Corporation

1055 First Street / Rockville, Md 20850
(301) 279-9550 Telex 898471





SILICON COMPILATION SPEEDS DESIGN OF COMPLEX CHIPS

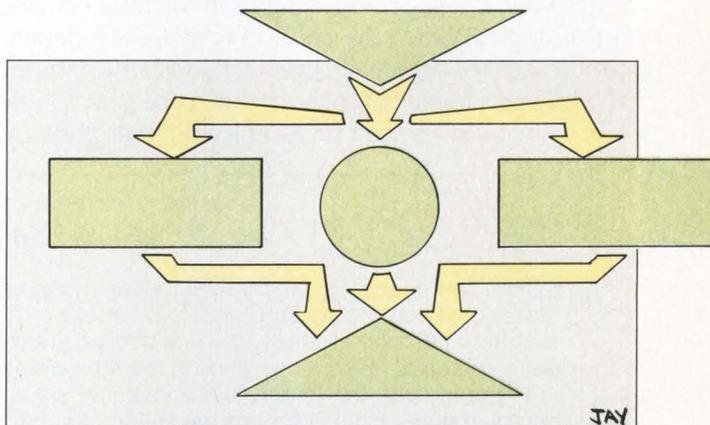
Silicon compilation creates a full-custom circuit with the silicon efficiency of a handcrafted design, but with fewer resources and a shorter development cycle.

by **Walt Curtis**

Development and maintenance of unique products are key to holding a strong position in the competitive IC market. Full-custom ICs can provide a leading edge, but designing them is costly and time consuming. Many of the computer aided design/engineering workstations introduced in the last three years have done a great deal to improve the design time of an IC, whether it be a gate array, a standard cell, or a full-custom circuit. What has not changed are the traditional IC design steps: architectural specification, block-, logic-, and transistor-level descriptions, geometric/polygon description, and several verification checks.

These basic design steps, however, are becoming more intricate as the design problems become more complex. Advances in IC processing capability, for example, have made a complexity of a quarter million transistors per IC commonplace. And, a complexity of one million transistors per IC is close behind. Moreover, each IC requires a total of seven designs to ensure that it will perform to the original architectural specification. Other problems in full-custom design include a limited availability of IC designers, inadequate design tools, development

Walt Curtis is vice president of marketing at Silicon Compilers, Inc (San Jose, Calif). He holds a BS in electrical engineering from Ohio University.

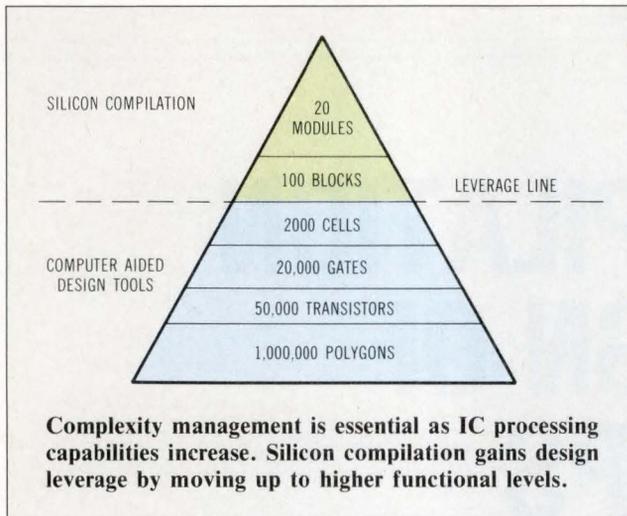


times extending to several years, and the uncertainty that an IC will work on the first pass.

As IC processing capabilities increase, complexity management in the design task is essential. A designer cannot handle millions of polygons or even tens of thousands of gates and expect to connect them accurately in a reasonable period of time, if at all. Silicon compilation offers a solution to these problems by keeping the design activity at the architectural level. Underlying complexities are managed by compiling the required models for IC design verification from the architectural specification.

No experience necessary

This new design method is especially valuable to designers with little or no previous IC design experience. A silicon compiler's encapsulated IC design expertise automatically compiles the layout details,



verification models, and checking. Thus, designers can exercise their creative, system-level architecture abilities and more time is available to analyze alternate architectures and approaches.

Providing a "what if" machine for IC design engineers, silicon compilation lets the designer determine how a design will behave in various process technologies. Since the method is process independent to a large degree, it is also very fast. IC designs, rivaling the most complex commercially available ICs in the industry, can be completed from specifica-

tion to working silicon in a much less than one year, and with only two or three engineers per design.

Silicon Compilers' Genesil silicon development system is such a "what if" machine. It is a multi-user design automation system that incorporates silicon compilation along with all the design tools and utilities for a designer to do a complete IC design. Three system engineers used the Genesil system to design a complex IC—the MicroVAX I chip. This chip is the CPU engine of Digital Equipment Corp's MicroVAX product line. It took several months to specify, implement, simulate, and verify the MicroVAX I. Although the engineers started with a basic definition of the chip, they fully implemented the architectural specifications using the Genesil system.

The system was partitioned into three separate functions: the microcode store, the control store sequencer, and the MicroVAX I. The microcode store uses off-the-shelf ROMs and erasable PROMs, and contains the microprograms that implement the VAX instruction set. There is no control store sequencer on the MicroVAX I due to the limitations of its 68-lead, pin-grid array package. The MicroVAX I chip contains a 47-register dual-port RAM file for 8-, 16-, and 32-bit operations, a 32-bit full-function ALU, a 32-bit program counter, a 64-bit wide barrel shifter, a 32-word constant ROM, and a 7-level operand restore stack.

Silicon compilation defined

Silicon compilation is a technique that dramatically reduces the cost of designing ICs by altering the traditional design strategy. Several design stages are automated, eliminating many of the housekeeping chores from the engineer's domain. A silicon compiler accepts an architectural specification for a chip and completes the block, logic, circuit, and layout stages of the design. The compiler has the knowledge of design experts and, therefore, can address all of the engineer's concerns for the chip.

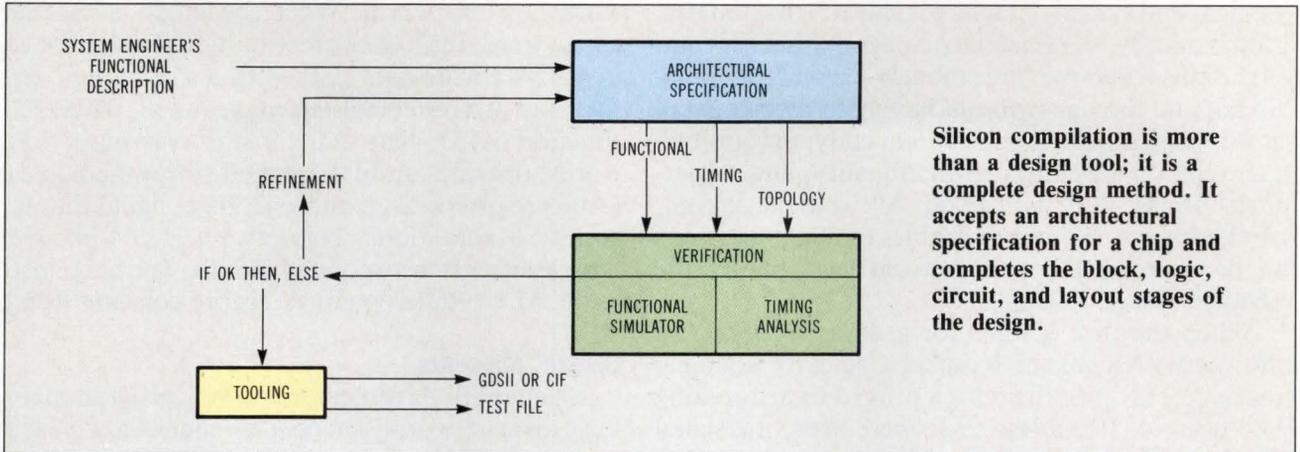
System support is needed for the geometrical, functional, performance, and test views of a design. Geometrical information helps to fabricate the IC. Functional aspects simulate the behavior of the design, allowing the engineer to verify the specification before implementing the circuits. Performance views state the size, speed, and power consumption of the circuit, so that an engineer can judge the feasibility of a specification. Test information catches fabrication defects in the final ICs.

If the silicon compiler cannot address all of these views, the engineer must have the detailed information from each design stage so that the verification steps can be performed manually. In essence, silicon compilation is a technique of designing with abstractions. For silicon compilation to succeed, certain properties must be proved about the abstractions.

As portions of a specification are composed, the topological, electrical, timing, and logical integrity properties must be maintained. Topological integrity checks that geometrical shapes and functions composed at the same time do not produce physical design rule violations or alter the functionality of the components. Electrical integrity ensures that the required electrical conventions are observed, such as voltage levels, signal loading, and noise levels. The timing attributes of each signal must match the circuit requirements. Also, the clocking strategy must guarantee that circuits behave over a range of fabrication variations. The logical polarities and the pipelining of each portion of the circuit must be consistent across the design.

All silicon compilers are not alike, however. The following questions can be used to help determine which compiler best suits a specific need. What is the total time required to design an IC from architectural specification to working silicon? Can designers use this technology or will special skills and training be required? How efficient is the resulting silicon implemented with the compiler? Can the functional high level design be easily "recompiled" into other process technologies?

—Dr David Johannsen, Vice President of Research at Silicon Compilers, Inc



Silicon compilation is more than a design tool; it is a complete design method. It accepts an architectural specification for a chip and completes the block, logic, circuit, and layout stages of the design.

A function set, a design definition system, a functional simulator, a timing analyzer, an object-oriented database manager, and a system to prepare output for silicon fabrication and test form the silicon development system. Function sets can represent a design with function-level building blocks, including data path, complex logic, memory, pads, test, and random logic.

Silicon compilation at work

The function set is the system engineer's vocabulary for describing the desired architecture. System engineers enter their preliminary architectural specifications through menu-driven forms. Initially, the designers let the system assign default values or parameters to less significant areas in the overall design verification. The menus are dynamic and context dependent. They "open-up" as required to ask the relevant questions as the system architecture is defined. This unique mode allows logic design rule checks and context consistency checks prior to compilation, thus saving time and errors.

Genesil's two-phase design approach encourages the exploratory design of chip architecture before a more detailed design is attempted. With this ap-

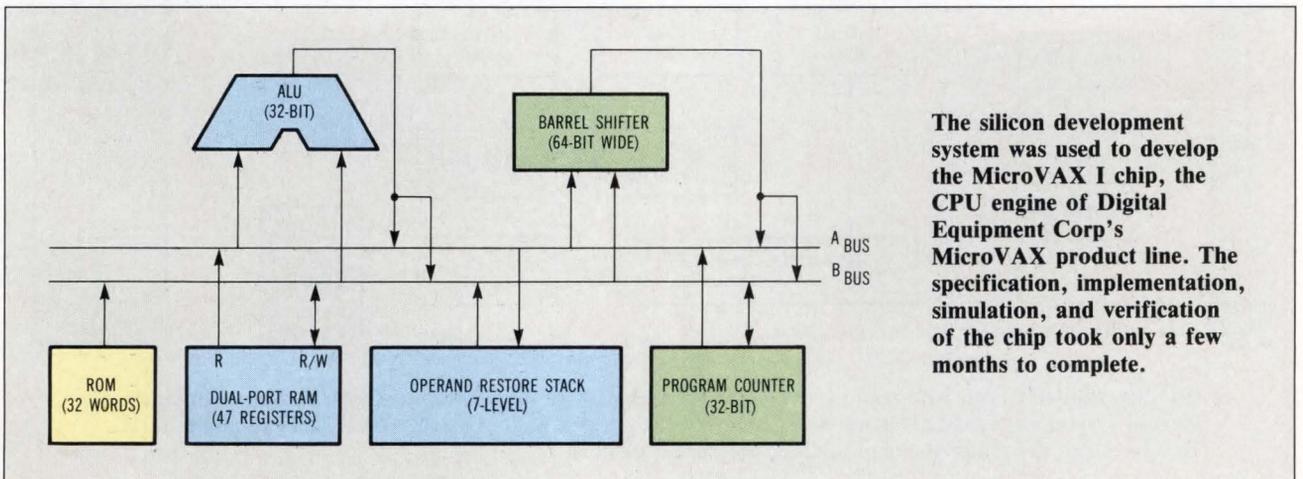
proach, the user can pinpoint and specify the best chip architecture very quickly. Feedback of size and performance occurs instantly, resulting in timely design trade-offs.

Silicon compilation makes such speed possible by eliminating the need for intermediate design levels. It automatically generates verification models for analysis and refinement. With a traditional IC design, refinements to correct or enhance a design's functionality are highly prohibitive once the topology is completed due to the huge investment in the current design.

If a refinement is done at all, it is made at the same level as the verification analysis. Many of the new CAD tools have improved the traditional IC design approach and have made much of the current standard LSI product possible. The CAD tools, however, do not address the front end design complexity made possible by the latest advances in IC processing densities.

Understanding the specifics

Design specification is performed by filling out forms so that designs occur quickly and without specifying every detail. The necessary information is



The silicon development system was used to develop the MicroVAX I chip, the CPU engine of Digital Equipment Corp's MicroVAX product line. The specification, implementation, simulation, and verification of the chip took only a few months to complete.

explicit, and system defaults provide rough estimates. Automatically, the Genesil compiler synthesizes complex designs into the three models—geometric, functional, and timing—without having to specify detail at the block or gate level. Consequently, the approximate size of the chip, its functionality, and its performance are well understood. A “what if” analysis of alternate architectures enables the right decisions to be made at the architectural level before the detailed design begins.

When the first designs for a data bus copier in the MicroVAX chip used random logic, for example, the aspect ratio of the design proved to be too long and narrow. If random logic were used, the signal routing between the pads and the rest of the chip would not be optimal. Rethinking the design, and reimplementing it by using a data path proved to be a much better solution.

When using the function set, a designer can create a block tailored to the needs of a design. The system will create, for example, a 7-bit RAM for the operand restore stack when the designer gives a 7-bit parameter to the RAM function. Fixed blocks or cells are not used, ensuring that no size or performance penalties occur.

The 47-register dual-port RAM file, composed of a 32-bit wide dual-port readable RAM, has three slices. Information can be read or written as bytes,

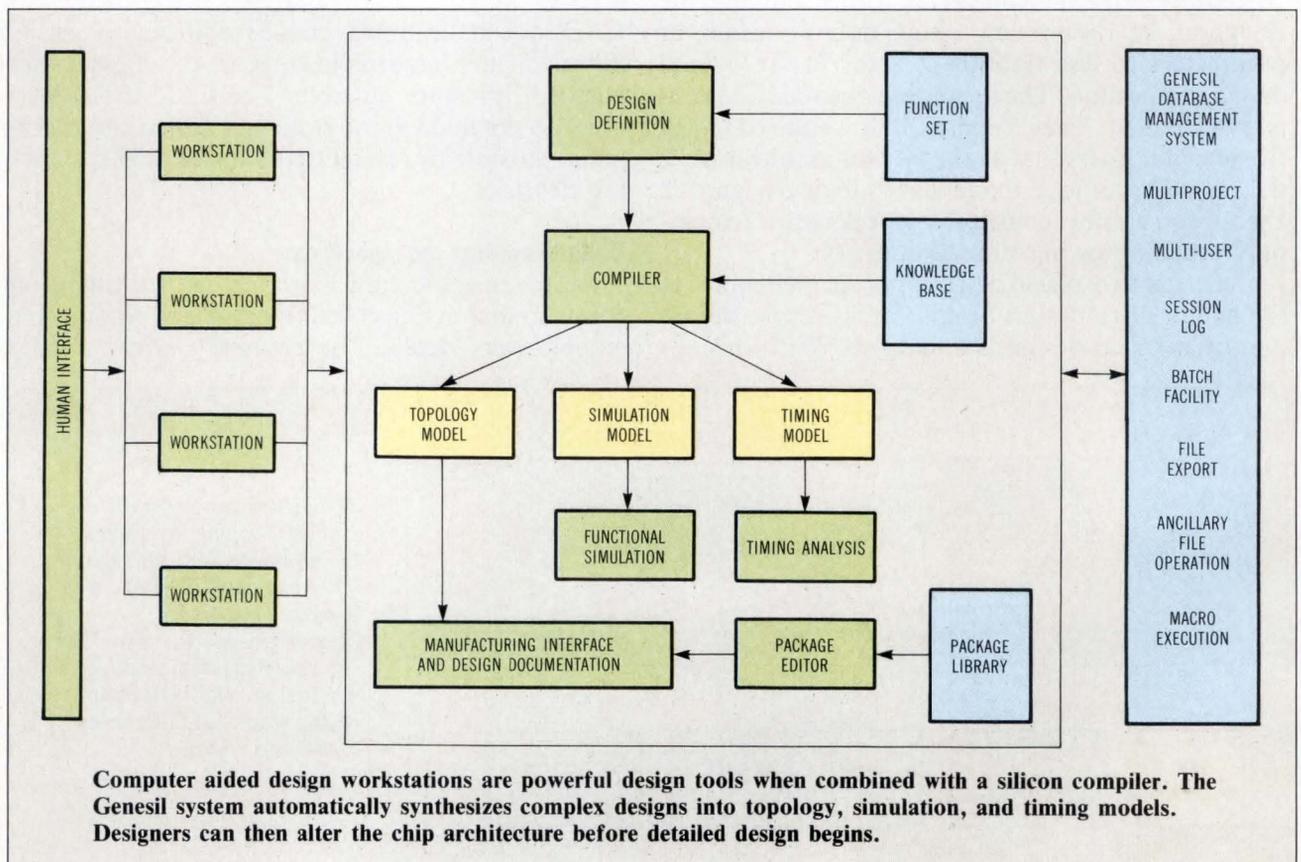
words, or long words. When the length is less than a long word, the higher order portion is not affected.

A 32-bit wide data path with a sign flag, a zero flag, and two result registers makes up the 32-bit full-function ALU. The ALU reads two input long words, operates upon them, and stores them in the results registers. Sign and zero flags signal the sign and zero conditions, respectively. A 64-bit wide barrel shifter is incorporated into the same data path as the ALU. ROM equals a 32-word constant ROM.

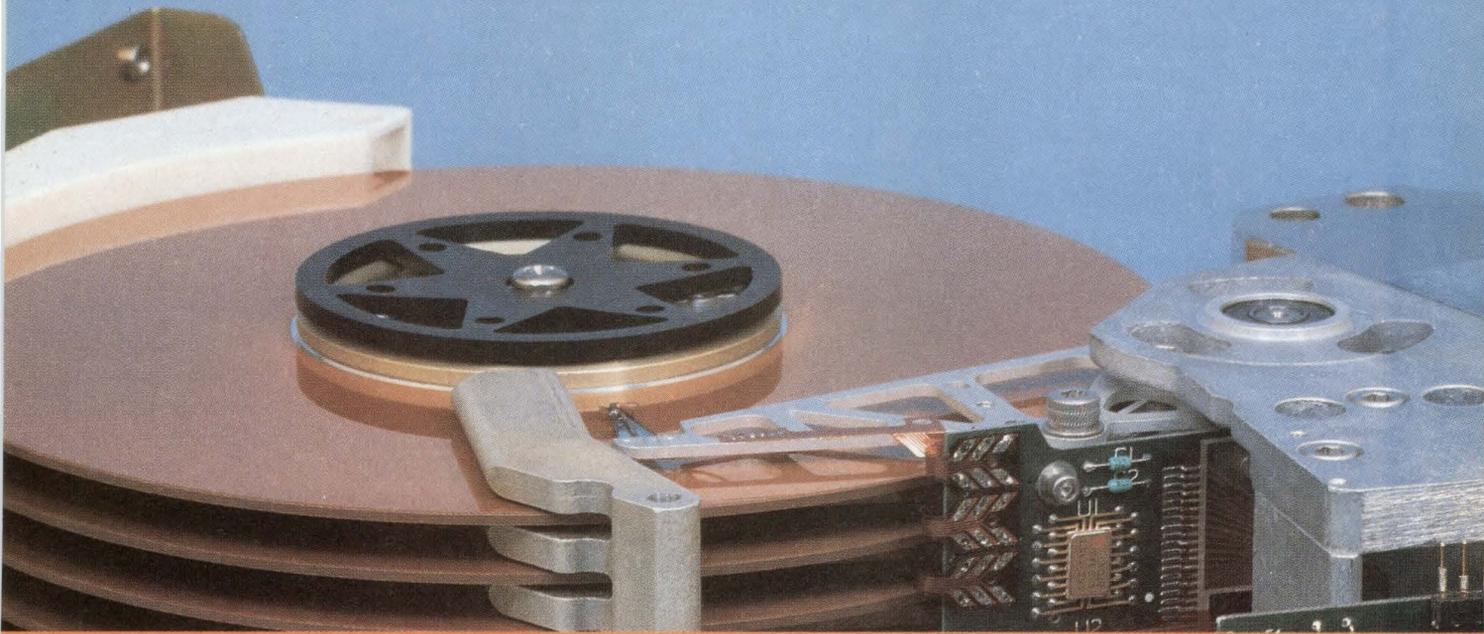
Control elements

A data path incrementer and two programmable logic arrays for program counter control and decode compose the 32-bit program counter. As a part of the data path, the program counter can be incremented by 1, 2, or 4 units. The two PLAs, PCDEC and PCENC, use the control inputs and exception conditions to determine the program increment.

Two RAMs (one to house the opcodes and the other to house the registers) and random logic (to determine how far to restore the stack) compose the 7-level operand restore stack. This module controls page faults. If a page fault occurs, the registers must be restored to their previous state, and the instruction sequence that lead to the page fault must be re-executed. Then the size and performance of each functional block are checked for proper performance.



A Hard Driving Success Story.



Innovative engineering in mechanical design and electronics with proven heads and media technology have been combined by PPC to create the DX Series 8-inch Winchester Disk Drives. Newly released, yet already well received by O.F.M.s, the DX Series offers high performance, high capacity disk storage.

Sized to fit an 8-inch floppy drive envelope, the DX Series meets the growing demands of multi-user, multi-access systems for increased data base, fast throughput and uncompromised reliability.

- Capacities of 180, 240 and 300 megabytes
- ANSI and SMD interfaces
- Advanced-design rotary voice coil actuator
- Proven oxide disks and reliable composite heads

Look into our fast-rising star, and you'll know what hard-driving performance is all about.

Write or call for complete technical data and fast, firm delivery dates.



**PERTEC
PERIPHERALS
CORPORATION**



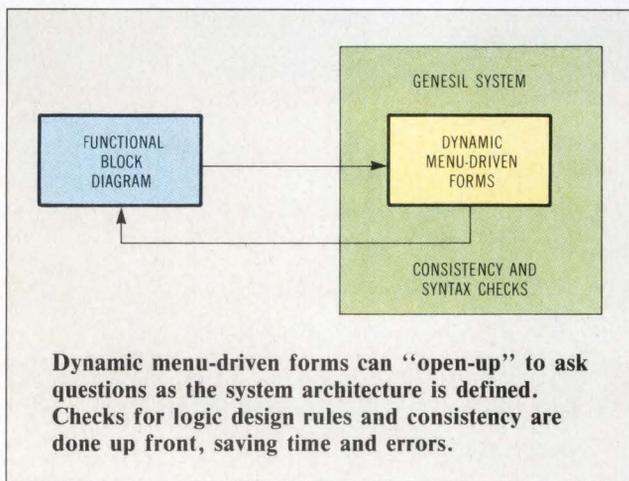
9600 Irondale Ave., Chatsworth, CA 91311 (818) 882-0030 TWX: (910) 494-2093
Toll Free Numbers: (800) 821-4126 (in California) (800) 443-9577 (outside California)

CIRCLE 45

The system synthesizes the performance model automatically as part of the compilation, allowing design engineers to move with relative ease between design specification and design verification. The Genesil timing analyzer verifies the ac performance of the design. It does not require any "stimulation" vectors to be created for a worst-case path analysis. Instead, the timing analyzer notifies the designer of the worst-case path delay, making it effortless to get the data necessary to assess whether the architecture being specified will work.

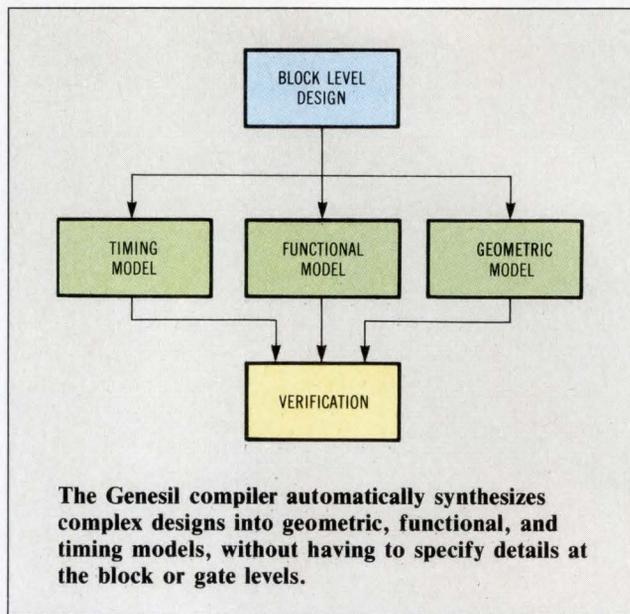
This independent, worst-case timing analysis provides very accurate results because it incorporates specific foundry process information that can affect the performance of the device greatly. Timing analysis results are reported in a data sheet-like format, showing the maximum clock period, setup, and hold times of various signals, and the signals' respective rise and fall times.

Designers can also specify given paths of interest for further analysis. The expertise built into the silicon compilation development system provides a designer with the rapid feedback necessary to test whether or not the design decisions made were correct. Thus, it is unlikely that engineers with little knowledge of ICs would design their way down a blind alley. In this way, Genesil systems make one of the hardest architectural design problems one of the easiest.



Detailed design specification eliminates system defaults used to gain quick feedback. For example, a list of the number of inputs, outputs, and terms is sufficient information to figure the size of a PLA for a preliminary design. In detailed design, however, the actual input and output signals need to be specified along with the terms of the PLAs.

Another activity during detailed design specification is optimizing the die size, using interactive placement. Die size in the MicroVax I was reduced 50 percent. Although this savings is good, it is not sur-



prising. Even a handcrafted IC devotes approximately 35 to 45 percent of the die to circuit interconnect. Because such a high percentage of the die is devoted to interconnect, the Genesil system allows the user to hand optimize the placement of a few dozen blocks.

Similarly, in the detailed design phase, the team needed to reverify the performance of the fully routed circuit. By using the timing analyzer again at this stage in the design, critical paths could be analyzed, and hand optimization of the placement to shorten critical paths proved useful. The team improved performance twofold in this manner.

DEC created a suite of test vectors that were generated by their systemwide simulator and captured the vectors at the MicroVAX I boundary. These test vectors are available to the system engineers in the DEC simulator's format. A user can create test vector input for Genesil functional simulator interactively or from a test vector stream, such as those from DEC.

It took two days to convert the test suite into a Genesil's MicroAssembler format, MASM. This format can create simulation stimulus and help define ROM contents. The ROMs are often used to store microcode. It is simple to create test vectors in the MASM format. And, this simplicity alone should encourage engineers to do design verification properly. In this way, MASM makes it easy to create a test suite for regression testing.

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

High 704

Average 705

Low 706

S U P E R C O M P U T I N G

A STRATEGIC APPROACH TO SCIENTIFIC AND ENGINEERING SUPERCOMPUTING.



THE FPS 64-BIT FAMILY: CONSIDER WHY THE MOST ACCESSIBLE SUPERCOMPUTERS MAY BE THE FASTEST WAY TO DO YOUR JOB.



The new 38 MFLOPS FPS-264, with 64-bit accuracy, large storage, and architecture refined to achieve a high percentage of its peak speed. For many applications, it can provide half the performance of the most popular super-computer. Its moderate price and exceptional support liberates super-computing from the realm of major corporate investment and puts it within practical reach of departments and teams.

Any chess player can appreciate the value of protecting existing resources. Of looking ahead. Of knowing what to sacrifice when. Of

strategic planning. In other words, a balanced attack.

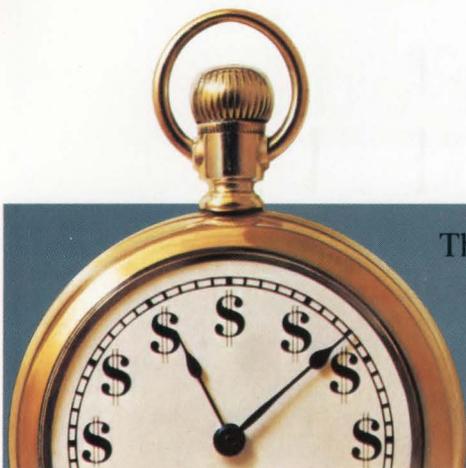
The FPS 64-bit family is precisely that kind of balanced attack on

compute-intensive analysis and simulation problems:

It protects and works with your existing hardware and software resources.

It is highly expandable. Proven dependable. Well-serviced. In other words, it is a safe, far-sighted investment.

And most telling of all, it makes supercomputing speeds affordable at the department level . . . so even those teams with remote access to multi-million dollar Crays[®] and Cybers[™] are likely to find that the advantage of immediate, local access is well worth the sacrifice of standing in line for the "fastest" machines.



The FPS optimizing FORTRAN-77 Compiler lets you easily adapt code to FPS' pipelined architecture in a form that is nearly as efficient as hand-coded assembly language. With extensions for asynchronous I/O and for enhancing compatibility with other compilers, it is one of most comprehensive tools of its kind.

There's more to performance than hardware specification. FPS provides you with an exceptional software tool set, compatible across the entire 64-bit family.

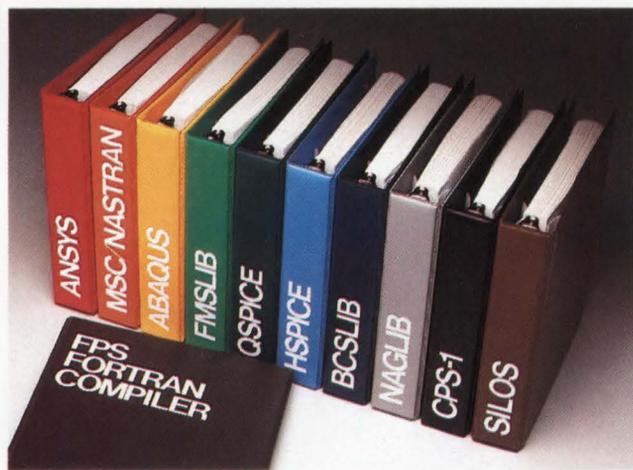
These include the System Job Executive (SJE) operating system, which supports interactive, multi-user, and batch sessions; our robust Compiler; the Overlay Linker for memory-intensive applica-

tions; plus Assembler, Object Librarian, Interactive Debugger, and extensive libraries of FORTRAN-callable matrix, simulation and mathematical subroutines.

If your investment in FORTRAN software is typical, you'll find the compiler alone to be a most compelling advantage of the FPS super-computing family.

Family Specifications

	FPS-264	FPS-164/MAX	FPS-164
Peak speed, MFLOPS	38	33-341	11
Dynamic range	2.8×10^{-309} to $9.0 \times 10^{+307}$	2.8×10^{-309} to $9.0 \times 10^{+307}$	2.8×10^{-309} to $9.0 \times 10^{+307}$
Logic format	64 bits	64 bits	64 bits
Main memory capacity	4.5 MWords	15 MWords	7.25 MWords
Maximum disk storage capacity	16 Gbytes	3 Gbytes	3 Gbytes
Precision	15 decimal digits	15 decimal digits	15 decimal digits
Vector registers	4 x 2K	124 x 2K (max.)	4 x 2K
Scalar registers	64	184 (max.)	64
Host interfaces	IBM, DEC	IBM, DEC, Sperry, Apollo	
Program Development Software	FORTRAN Compiler, Overlay Linker, Assembler, Object Librarian, Interactive Debugger.		



More applications software than for any comparable computer.

Compare quantity and quality of compatible third-party software packages—for structural analysis, circuit design, reservoir simulation, fluid flow analysis, chemistry and much more—and the FPS advantage widens.

How fast you get a

supercomputer up and running is as important as how fast it runs.

In the real world, efficient operation is a function not only of speed, but of software, proven dependability, compatibility and established service capabilities—factors that for Floating Point aren't just promises or expectations, but *facts*.

Where does the FPS 64-bit family fit in? Anywhere scientific computations have become increasingly cumbersome and costly to general-purpose computing. By offloading numerically-intensive operations to the FPS pipelined architecture, you can complete mathematical models and simulations many times faster. You can do it for a fraction of the cost of multi-million dollar scientific mainframes. And you can free your other computers for the jobs they are best equipped to perform.

Now three complementary performance ranges give you more value for the dollar and more room for growth.

With system prices starting at \$300,000 (U.S.), for the 11 MFLOPS FPS-164, Floating Point Systems can give you the number-crunching capability of several general-purpose computers, for less than the cost of one.

The new 38 MFLOPS FPS-264 achieves 4-5 times the speed of the FPS-164 on many applications programs. It is the best choice for conventionally-written FORTRAN programs (in contrast to optimized code tailored specifically for the FPS-164/MAX). Systems start at \$640,000.

The FPS-164/MAX adds another dimension to the family. With its multiple

Family Performance Measures

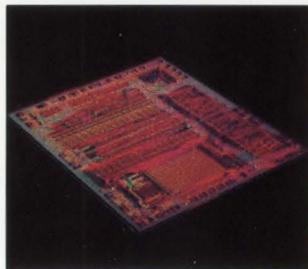
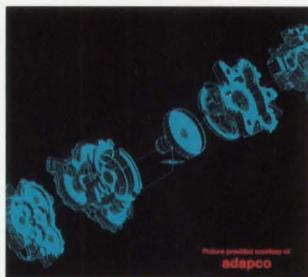
	FPS-264	FPS-164/MAX	FPS-164	
		15 accelerators	1 accelerator	
Peak MFLOPS	38	341	33	11
Peak MOPS	190	1705	165	55
Peak MIPS (Multi-instruction parcels)	19	5.5	5.5	5.5
Typical MFLOPS, LINPACK Benchmark	9.9	20.0	6.0	2.6
Whetstones (64-bit)	20,100	5800	5800	5800
1000x1000 matrix multiply, seconds	53	10	66	189
\$K/MFLOPS (system price/peak speed)	\$16.8K	\$2.5K	\$12.3K	\$27.1K

special parallel processing units and peak 341 MFLOPS performance, it can run many common matrix computations as fast or faster than super-computer mainframes costing 10 times as much to purchase and operate. System prices start at \$400,000.

All three FPS Scientific Computers are highly expandable in memory and mass storage capacity. They are totally compatible with each other in terms of program development and applications software. Error-correcting memories and advanced internal diagnostics are among the components of a reliability program that extends

from exhaustive manufacturing testing to our 21 field office service facilities worldwide.

Compare the specifications, benchmarks and price ranges listed here with all the alternatives — from superminis to Class 7 supercomputers. Then ask your local Floating Point sales representative to help you develop your own strategy for accessible supercomputing. Contact Floating Point Systems, Box 23489, Portland, OR 97223. Telex 360470 FLOATPOINBEAV. Sales offices worldwide. Call collect (503) 627-9667.



The new FPS-264 will be of particular benefit to the problems of structural analysis, reservoir simulation, electronic circuit design and computational chemistry and physics.

THE PROVEN POWER IN ACCESSIBLE HIGH-SPEED COMPUTATION.





TOP-DOWN TOOL OPENS NEW PATHS TO CHIP DESIGN

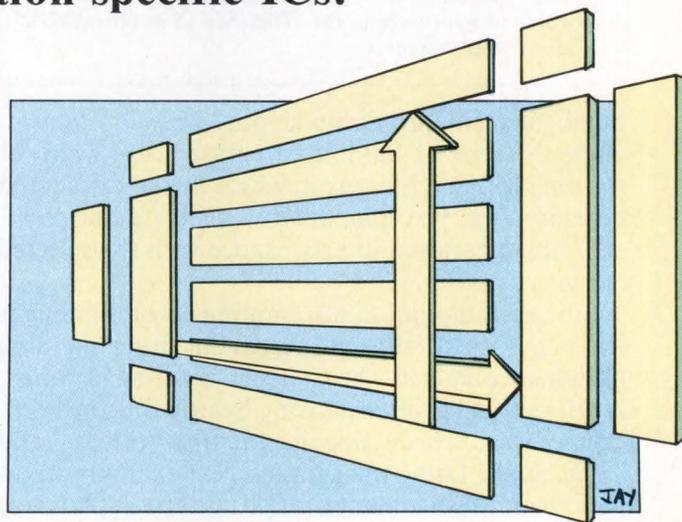
Silicon compilation software gives a system engineer with no chip design experience a top-down design tool for the development of CMOS application-specific ICs.

by Donald Ritzman

Most of the currently available design tools for application-specific ICs support only bottom-up techniques. A bottom-up approach lets an ASIC designer optimize chip density and performance for a particular application, but it requires the designer to enter and manipulate all data, from logic diagram to the interconnection of functional modules. A top-down approach, on the other hand, lets the designer specify only the characteristics and interconnection requirements of modules. Detailed, device-level knowledge is not required, proving the top-down method a more appropriate design tool for system-level engineers.

Integrating silicon compilation software into a powerful computer aided engineering workstation provides such a top-down design tool. It lets a system engineer implement designs as full-custom chips, yet produce them in no more time and at no more cost than an equivalent gate array. The workstation also allows the experienced IC designer to modify an automated layout manually and to use the silicon compiler as a cell compiler. Since the software is integrated into the workstation, neither the chip designer nor the system engineer have to go outside of the workstation environment during the IC design cycle.

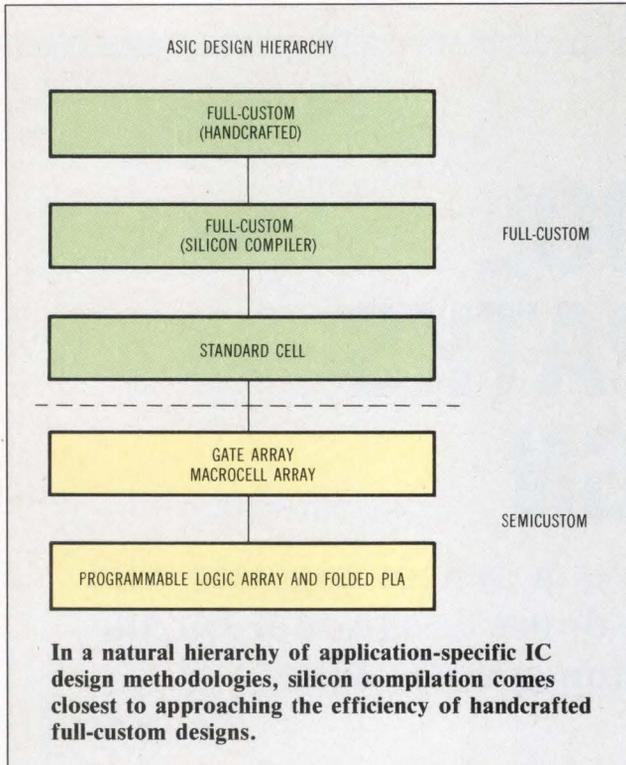
Donald Ritzman is a product marketing manager for Valid Logic Systems, Inc (San Jose, Calif). He holds an MBA in marketing management from the University of California, Berkeley.



Silicon compiled designs are also rule-independent, and chips created according to one set of design rules can be converted to another set simply by inputting the new rules and recompiling the design. In this way, designs may be configured easily for multiple production sources and can be readily updated to take advantage of changes and improvements in processing technology.

Using a silicon compiler

Essentially, a silicon compiler allows a designer to describe the architecture of an IC simply by creating the block diagram of the chip. The blocks can be functional modules, such as adders, multiplexers, registers, or memory. Users merely specify a functional description and the parametric options for each module in the block diagram to translate the



block diagram into a chip layout. These options include the type of module to be used, the width of the bit path, and the way in which the blocks should be connected. The compiler then generates the physical implementation in accordance with the selected design rule set.

Silicon compilation also minimizes or eliminates the penalties of slower speed and larger die area usually encountered when a design is not handcrafted. ASIC designs produced using Seattle Silicon Technology's Concorde silicon compiler with a Valid Logic SCALDstar workstation, for example, compare favorably with handcrafted custom chips. These designs typically achieve 90 percent of the density and 100 percent of the performance of handcrafted designs. They are also 25 to 40 percent smaller than equivalent standard cell designs and from 50 to 75 percent smaller than gate arrays. Thus, silicon compilation is the closest alternative to handcrafting in the ASIC design hierarchy.

Modules for all reasons

The silicon compiler package offers a choice of 19 modules for SSI and MSI devices, programmable logic devices, memory, and "glue" elements. Buffer, strip, and D- and RS-flipflop compilers produce modules for 3-micron CMOS SSI devices. The buffer compiler creates customized buffers for use within a custom chip. These custom buffers let the user specify the number of stages, size ratios between pull-up and pull-down transistors, and size ratios of

transistors between stages. The strip compiler generates rows of SSI cells, such as 2-, 3-, and 4-input NAND gates, NOR gates, exclusive and inclusive OR gates, inverters, ordinary buffers, and buffers having three-state outputs. The D- and RS-flipflop compilers give the user a variety of options for their configuration, including preset, clear, input access, and drive capability.

Compilers for MSI device modules are the synchronous counter, ripple counter, adder, shift register, multiplexer, and decoder. With the exception of the decoder, which uses a code file input similar to a truth table, data describing a module is entered in a menu format. For example, to create a synchronous counter module, the user simply assigns a unique module name and then proceeds to define counter type (up, down, or up/down), number of counter bits (from 2 to 99), output buffer size (which determines the fanout capability of the counter), and other operating parameters.

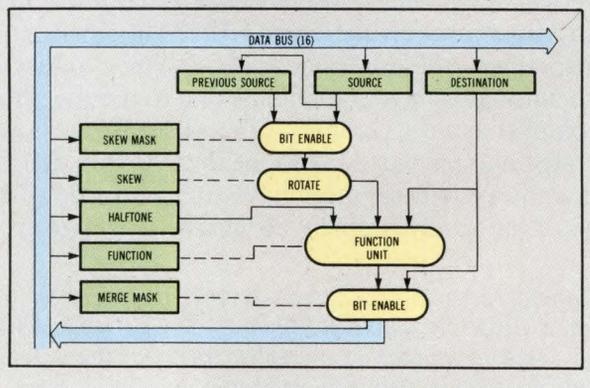
The data path-compiled design

A good example of a data path-compiled design is a high performance, low power CMOS raster graphics accelerator that integrates the inner loop of the popular RASTEROP, or "BITBLT," function into a single high speed operation. This circuit replaces approximately 50 TTL components and contains 60 major registers. These registers are divided into five functional groups, as well as addressing, control, and I/O circuitry.

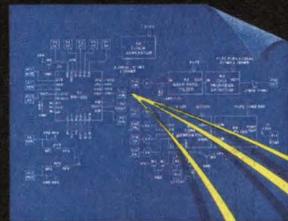
The circuit can be partitioned and each block developed individually using the module compilers. This does not yield the chip density possible by treating it as a number of register and logic block columns interconnected by perpendicular data buses, however.

Developed for 3-micron process design rules, the final chip is 180 x 220 square mils and contains more than 6500 transistors. Best-case critical path delay is 50 ns and worst-case is 120 ns, with a typical value of 80 ns. Typical power dissipation is 200 mW.

Of perhaps greater importance to the system engineers who will be responsible for most of the future IC designs, the work was actually done by a non-IC designer. The entire cycle took 2 days, of which only 45 mins was actual run time.



NO ONE INTEGRATES SYSTEMS WITH CUSTOM/SEMICUSTOM CHIPS BETTER.



Analog, Digital, and Analog/Digital Capabilities

There are many ways for putting digital systems on a chip, and, of course, we can provide digital solutions. But the going gets tough when the systems are analog, and it gets really sticky when you need both analog and digital on the same chip. You can immediately rule out all the gate array guys and look to the real custom/semicustom chip experts. You can look to Silicon Systems — a company that's built its growth on its combined analog/digital capabilities.

Bipolar and CMOS Processes

Silicon Systems' automated wafer fab is not restricted to a single technology. We offer

both Bipolar and CMOS processes and have no bias in steering you toward one technology or the other. So you can be sure that we'll only recommend the process that is best for your system solution.

Standard Cells or Full Custom

Everyone knows that with moderate-to-large volume production, a full custom chip is the optimum approach. But everyone thinks it takes a long time and a lot of money to get there. Not at Silicon Systems. We can implement your design with selections from our standard cell library of digital and analog circuits that combine quality performance with fast-turnaround. Or perhaps you could

use one of our special semicustom switched capacitor arrays that can be easily tailored for a variety of special filter applications. And when the time is right, we can lead you promptly and painlessly into a full custom design, using our unique Integrated Design Methodology (IDM™).

So, Silicon Systems has the answer for you. For the full story, send for our new Custom/Semicustom Integrated Circuits brochure.



Silicon Systems, 14351 Myford Road, Tustin, CA 92680. (714) 731-7110.

 **silicon systems™**
INNOVATORS IN INTEGRATION

Circle 47 for Career Information

Circle 48 for Product Information

As the selections are made, the compiler automatically calculates module performance and displays it on the menu. If the result is not satisfactory, the design can be changed immediately. In this way circuits can be built to meet the exact requirements of an application. Clearly silicon compilation offers a tremendous advantage over standard logic, gate array, and standard cell designs, which require fixed formats and limit the designer.

The programmable logic array compiler produces a CMOS equivalent circuit from a description in the Palasm programming language format. A chip can be designed that contains only the PLA-equivalent functional module, or the module may be used as part of a larger chip. A compiler for standard PLAs is also part of the software package.

One of the most powerful modules in the silicon compiler is the "flair" compiler. It is a folded PLA program—the two major AND and OR arrays of the actual chip are interleaved for a space-saving layout. Virtually any kind of logic function can be implemented using the compiler's storage elements. A code file input similar to that of the MSI decoder compiler defines functions.

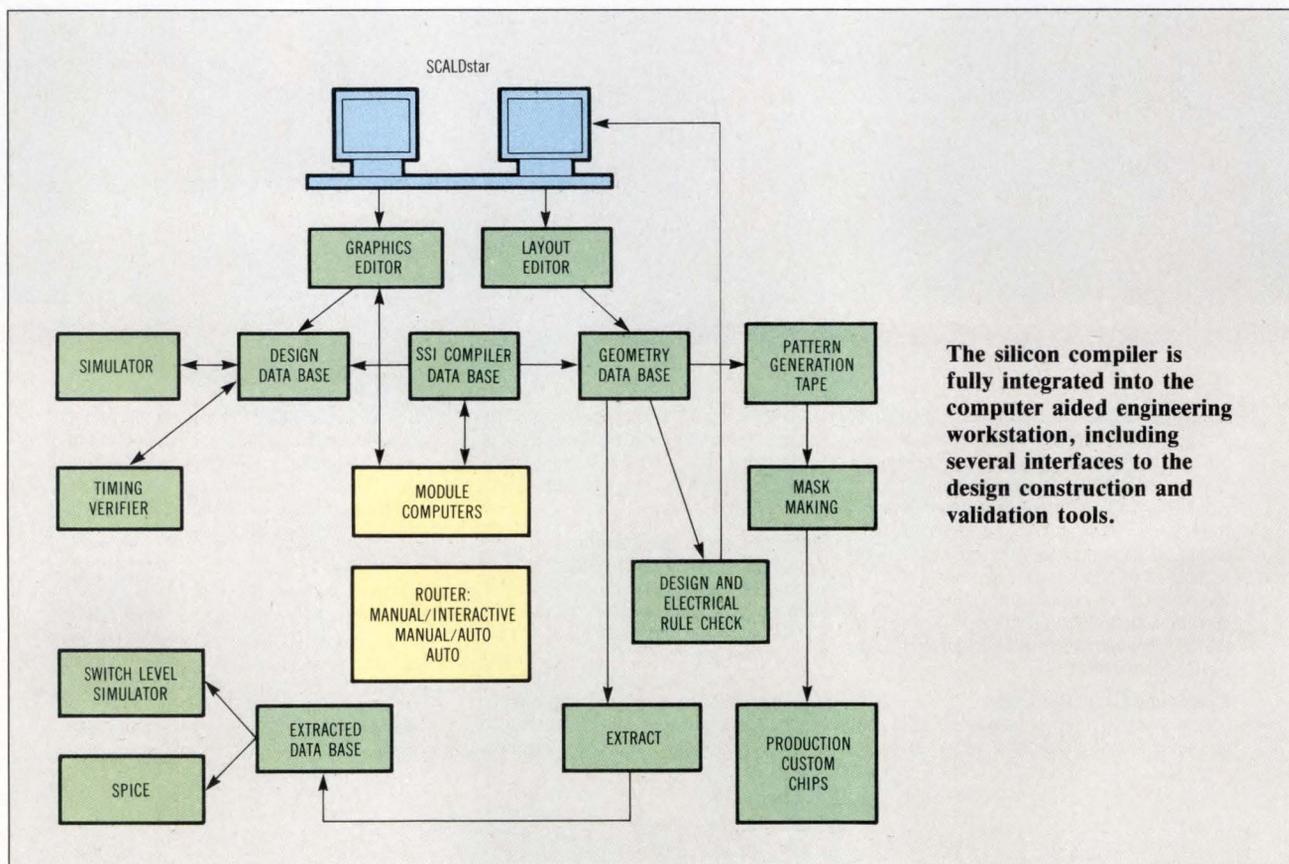
Static RAM, ROM, and first in, first out register modules are produced by the memory compilers. These modules are menu-driven and include options that allow a user to select a floorplan or a shape for

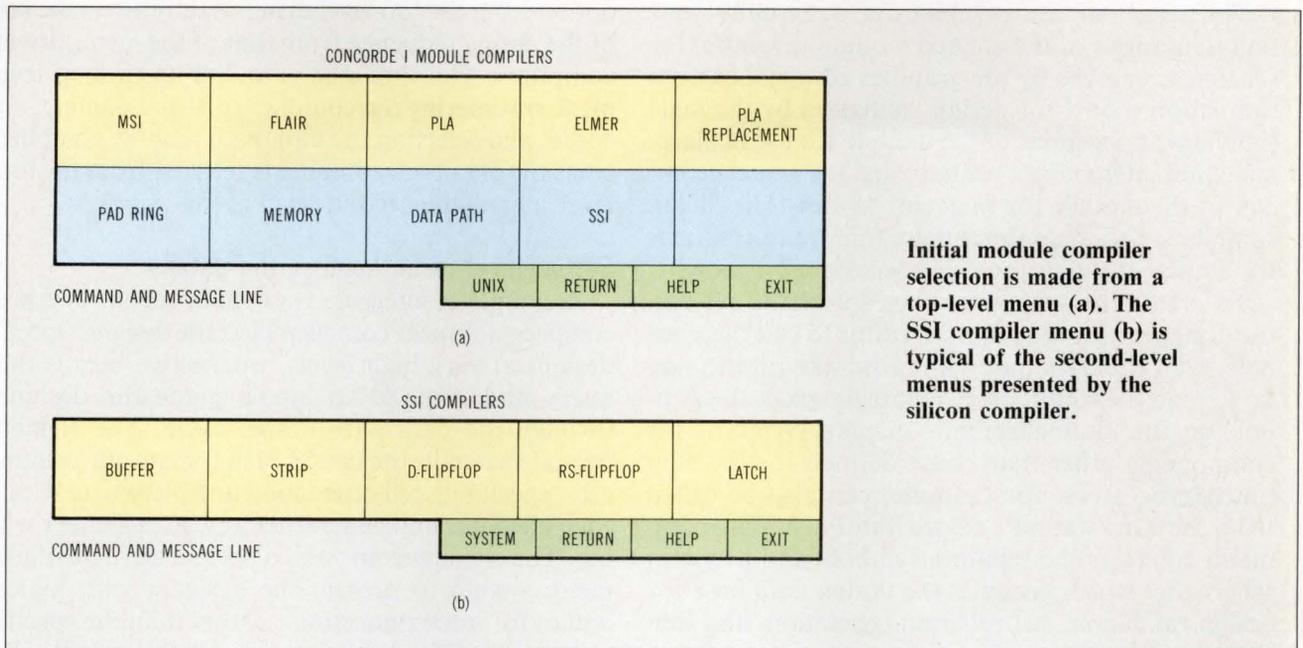
the final memory array. Since the memory array normally dominates the real estate on a custom chip, the ability to lay out memory in a square, rectangle, or oblong shape lets the designer optimize use of the chip space.

Three compilers handle interconnections, data flow, and external access. The logic elements frequently used as "glue" between large, complex functional modules are created by the "elmer" compiler. This is essentially an SSI and MSI standard cell compiler with a router. The data path compiler is a tool for constructing columns of registers and logic blocks with perpendicular data buses. Unlike the menu-driven compilers, the data path module uses a specification language to describe the columns and buses. The interface between the custom chip and the outside world is designed using the pad compiler. As the pad design progresses, the compiler provides such data as static protection, propagation delay, power consumption, and physical size. Default data is included for several standard pads.

The workstation environment

The CAE workstation provides an ideal operating environment for the silicon compiler. It is specifically aimed at the design of VLSI chips, from initial concept through finished layout. The workstation software package includes a graphics editor for





Initial module compiler selection is made from a top-level menu (a). The SSI compiler menu (b) is typical of the second-level menus presented by the silicon compiler.

design construction, a timing verifier and logic simulator for design validation, a packager for netlist generation, a layout editor for manual design, electrical and design rule checkers, and other utilities, all running under the Unix operating system.

This total integration of schematic capture, logic simulation, timing verification, layout, fault simulation, and silicon compilation allows design at the schematic level, module level, or layout level. It also allows the mixing of designs produced by different methods and at different levels.

A unique feature of the SCALDstar workstation is its use of two displays: a monochromatic screen for schematic capture and logic design, and a color screen for physical layout of chip designs. This arrangement permits the user to simultaneously view a logic design and its physical implementation. A system engineer can create a custom chip design on the monochrome screen using the workstation utilities and the silicon compiler and display the resulting physical design on the color screen. In addition, knowledgeable IC designers can call on the system's capabilities to manually modify the automated layout to further optimize chip density.

Designing with top-down tools

Implementing a system design with standard devices usually consists of selecting the components that will be used and drawing a schematic on the CAE workstation display. This creates a data base for subsequent design validation and production. The process is similar when the silicon compiler-based top-down tool is used.

Using the graphics editor, a detailed system block diagram is created that determines how to implement

and partition the design. The hierarchical diagram represents many of the functions with single blocks and has enough detail to clearly define the bus control structure of the system and the characteristics of each functional block.

The next step is to decide which module compiler or compilers will be used to construct each functional block. As with standard logic, there may be several alternatives. For example, data bus buffers could be implemented using the strip, elmer, or PLA compiler; and an address latch could be constructed using the D- or RS-flipflop compiler. Having a variety of compiler functions available lets the user select the most area-efficient or best performance approach for the circuit being developed.

When the module type decisions have been made, the designer calls a top-level menu from the silicon compiler and selects an appropriate major compiler group. The designer also specifies the name of the set of design rules that will apply to the chip. Currently, design rules for 12 foundries are offered.

The selection of a major compiler group causes a second-level menu to appear, presenting the various functional module compilers within the group. If, for example, a designer chooses the SSI circuits group, the second-level menu would offer a choice of buffer, strip, D- and RS-flipflop compilers. Selection of a specific compiler causes a third-level menu to appear, allowing the designer to specify the desired module configuration and operational parameters.

As the designer makes the descriptive entries on the third-level menu, the compiler simultaneously calculates the module's performance characteristics and displays that information on the menu. When the characteristics are acceptable, compilation can begin.

The result of the compilation is actually four representations of the created module: a symbol for schematic capture by the graphics editor; a custom simulation model for design validation by the logic simulator; a footprint of the module for use in placement and interconnect routing; and the actual geometry of the module for final chip layout. The silicon compiler's data base stores these four representations for future use with the other workstation tools.

To create the schematic diagram of the system, the designer uses the graphics editor to call the symbols from the data base. A mouse-like positioning device and the editor's Add command place the symbols on the monochromatic display. Symbols for components other than those defined for the chip (microprocessors, for example) can also be called from the workstation's device library. A Wire command connects the separate symbols into a system schematic, which becomes the design data base for design validation, test program generation, and subsequent first-part testing.

Data path design is an important concern when choosing the most suitable compiler for the task. If a data path design is necessary, selecting a compiler

option from the top-level menu slightly alters the rest of the design sequence from that of the menu-driven compilers. The first step is to call up the top-level menu by entering the compiler software name, Concorde, and selecting the data path module compiler. Once the data path compiler is selected from the top-level menu, the second-level menu appears.

Second-level menu handles the details

A complete, integrated system of menus and pre-compilers (element compilers) lets the designer specify designs at very high levels. Interactive menus that query and respond to inputs guide the designer through the data path design cycle. The element compilers handle the details of automatically creating cells, specifying cell orientation and placement within each element, and generating any intraelement wiring. The element compilers also coordinate with the menu system to present the designer with logical names for interconnections, rather than the specific port names that are associated with individual cells.

The end product of the user interface is a coherent, human-readable specification file that describes what the designer has defined with the interactive menus.

DATA PATH COMPILER VERSION 2.0						DEFINE DATA PATH ELEMENTS
MODULE NAME: DPATH 1		NUMBER OF BITS: 32				REGISTER
ELEMENT ORDER:						ALU
1	REG1	6	MUX1			ADDER
2	REG2					MULTIPLEXER
3	ALU					SHIFTER
4	PETE					DEFINE TERMINALS
5	ADDER1					EXIT
EXISTING REGISTERS:		REG1	REG2	REG3	PETE	
				UNIX	RETURN	HELP
MESSAGES GO HERE						

(a)

ALU ELEMENT VERSION 2.0							CREATE ELEMENT	
ELEMENT NAME: ALU								
NAME	TYPE	NET	NAME	TYPE	NET			
INPUTA	BUS	BUS4	CNTRL-GROUP0	CONTROL (4)	CNTRLO			
INPUTB	BUS	BUS5	CNTRL-GROUP1	CONTROL (4)	CNTRL1		UNIX	
OUTPUT	BUS	BUS6	CNTRL-GROUP2	CONTROL (4)	CNTRL2		RETURN	
CARRY IN	CONTROL	COUT1						HELP
CARRY OUT	CONTROL	CIN1						EXIT
EXISTING ALUs: ALU ALU1								
MESSAGES GO HERE								

(b)

Two interactive menus—one for order elements (a) and one for ALU elements (b)—query and respond to inputs to guide the designer through the data path design cycle.

It can be accessed and altered manually, making data path compiler's flexibility and functionality totally available to the designer. Since the specification file can be altered in this way, extensive error checking and design verification are built into the data path compiler.

As generated by the compiler, a data path is organized in columns of identical cells, with data signals entering through input ports on the left side of the module and output signals being available on the right side. Control signals, along with power and ground rails, run vertically through the columns. The compiler automatically makes data and control signal connections between ports of the individual cells as specified by the designer.

Connections between the cell ports are generally specified by establishing a series of nets, which are associated with both data and control signals. Each net is named by the designer and associated in the specification file with the ports it is to connect. Likewise, the cell ports to be connected to the net must be associated with it in the cell specification.

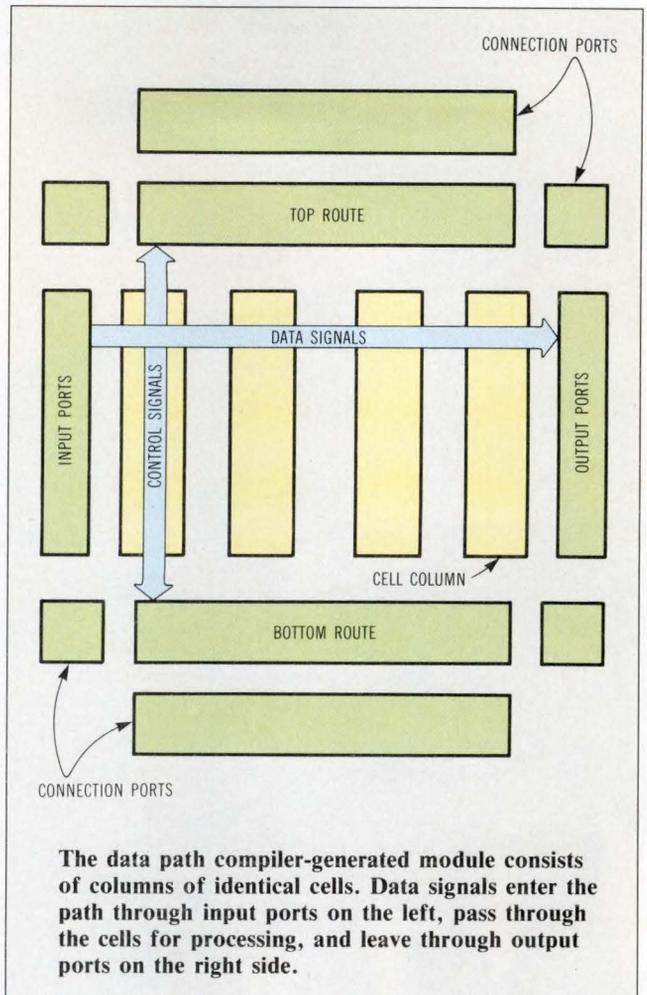
Once the schematic or data path diagram has been created, the designer must make certain that, if manufactured as displayed, the chip will perform as intended. This design validation is performed using the workstation's logic simulator, which draws upon the simulation models stored in the silicon compiler's data base. For non-custom circuits, the logic simulator uses software models in the workstation's library.

Logic simulation is an interactive process. The designer provides input stimuli and the simulator reports on circuit performance. If the circuit does not perform as anticipated, the designer simply returns to the graphics editor, modifies the design, recompiles it, and repeats the simulation. This procedure is repeated until the design is fully validated.

Final processing of the validated design can be performed automatically by the workstation and silicon compiler. First, the footprint and geometry representations from the layout data base are used to lay out the modules and other components as they should appear on the chip, then the interconnections are routed between them. Finally, a pattern generation tape is produced, from which the foundry makes masks to be used in wafer fabrication.

While automatic layout and routing is adequate in designs consisting of a single logical module, chips containing more than one module usually require some manual manipulation for optimal placement of modules and routing of interconnections. Manual routing, however, requires knowledge of IC layout techniques.

Because the intent of the top-down, silicon compiler-based design system is to eliminate the need for such specialized knowledge, the system offers



an interactive/symbolic routing capability. This allows the designer to direct the placement of modules with the workstation's interactive symbolic router, which runs the interconnections while enforcing compliance with the process design rules. Also, with the interactive router, the designer can recompile the design without further manual intervention.

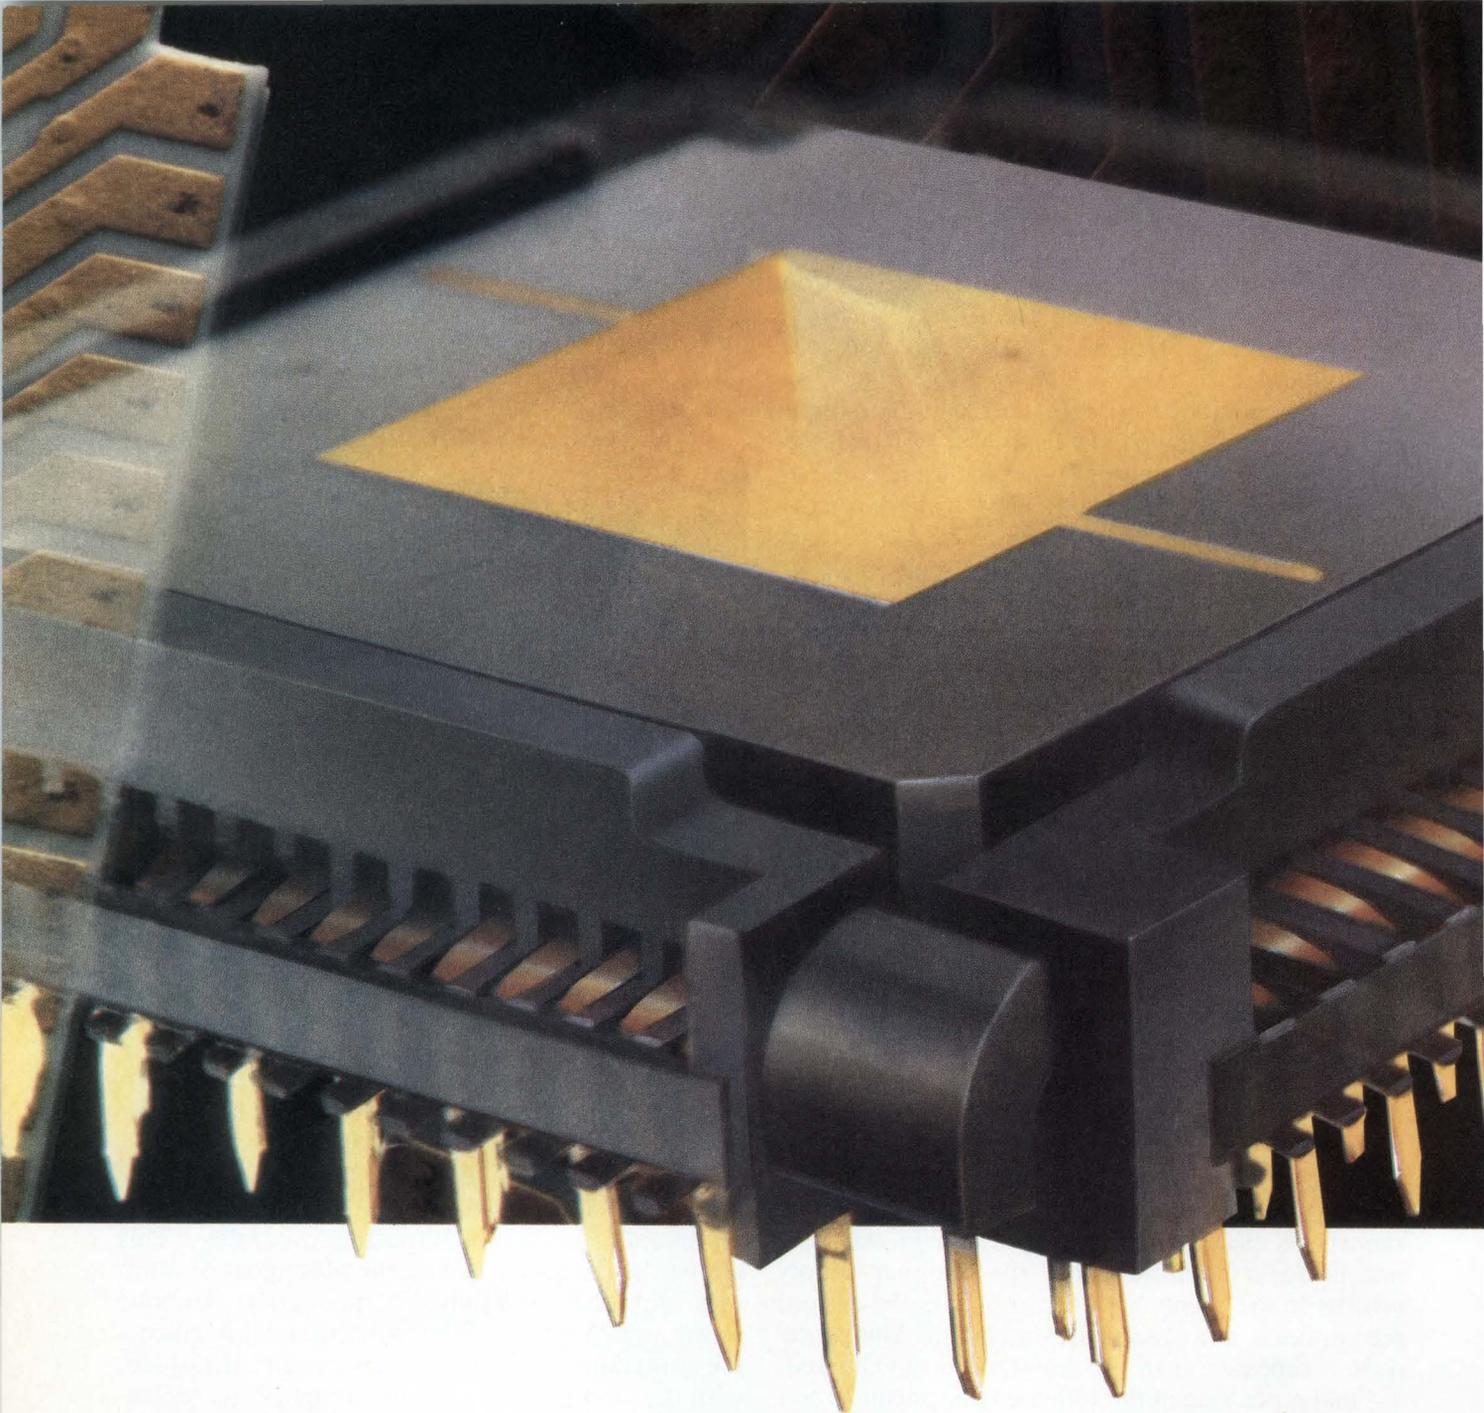
Once routing is complete, the design can be simulated again, with data path information (eg, lengths and signal delays of interconnections) included. This provides a more accurate indication of the chip's actual performance than simulation done before routing. When the total design has been validated, the pattern generation tape is written. If multiple sources are used, the design can be recompiled according to the design rules for the individual foundries and a separate tape is prepared for each.

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



Screening ideas for lower substrate-to-board interconnect costs?

Login with AMP chip carriers and sockets.

If you've been waiting for chip carrier technology to reach a more practical stage, it has.

AMP low-profile sockets for plastic and ceramic chip carriers are now available at a surprisingly

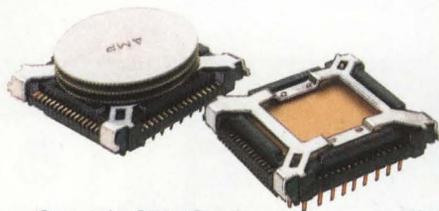
low-profile cost. In standard JEDEC sizes. In surface-mount or through-hole versions, with optional heat sinks.

Our ceramic chip carrier sockets feature duplex-plated contacts to eliminate solder-bath contamination. Plastic chip carrier sockets are available with gold or tin contacts, to match your requirements.

We even offer plastic chip carriers in a JEDEC Type A envelope, ready to run on your existing equipment.

And it's all available immediately, so you can make it part of today's design work. And move into production with assured reliability, assured economy.

Find out more. Give us a call.



Ceramic Chip Carrier Sockets

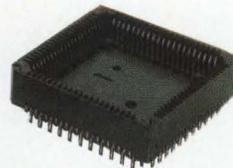
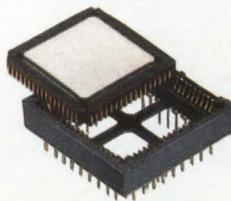
- JEDEC Types A, B, D
- Contacts: duplex plated

Call (717) 780-4400
and ask for the AMP
Chip Carrier Socket Desk.
AMP Incorporated,
Harrisburg, PA 17105.



High Speed Sockets, Surface Mount

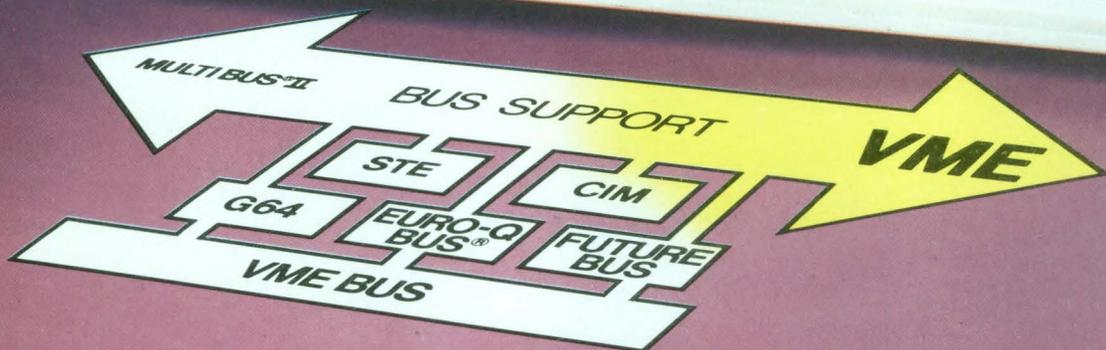
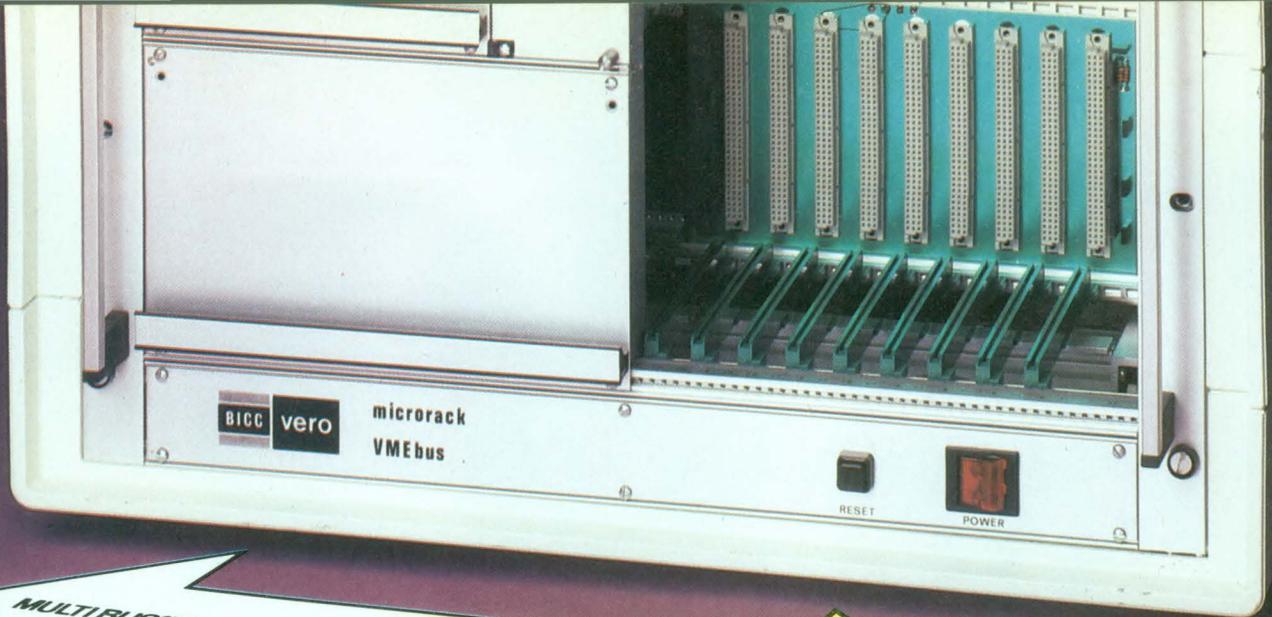
- JEDEC Type A Leadless Ceramic
- Centerlines as close as 0.020"
- Contacts: gold,
capacitance ~ 0.5pF
inductance ~ 1.4nH



Plastic Chip Carriers and Sockets.

- Gold contact premolded carrier incorporates low inductance leads; permits cavity up or cavity down mounting.
- Gold contact sockets mate with AMP Chip Carrier or other leaded Type A devices.
- High-pressure tin contact sockets for JEDEC .050" centerline plastic Chip Carriers.

AMP Interconnecting ideas



Complete hardware support for the VME bus user.

The hardware support for a bus system based on the Eurocard form factor needs very specialized 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 frames, 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.

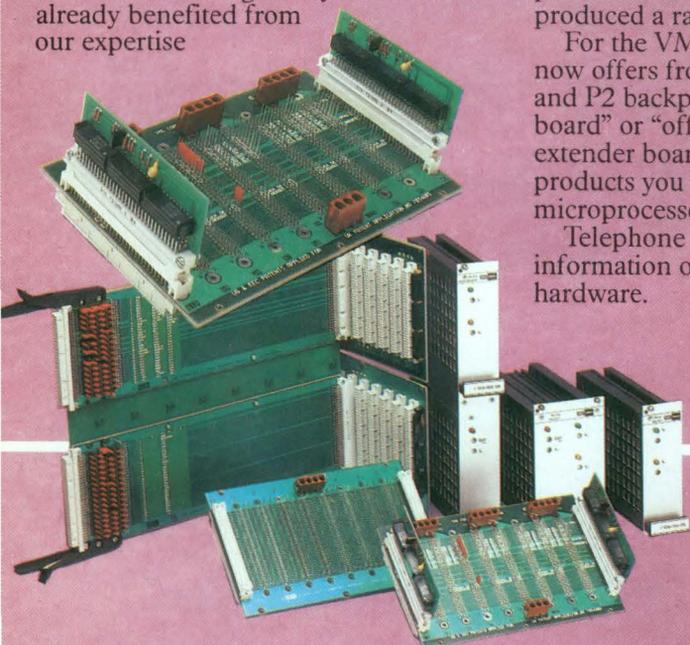
For the VME bus user BICC-VERO now offers from stock, a full range of P1 and P2 backplanes with either "on-board" or "off-board" termination and extender boards, as well as the standard products you will require to build your microprocessor system.

Telephone or write for further information on VME bus support hardware.

Whatever your system configuration BICC-VERO have the experience and the hardware to support you.

DATA FILE

- * P1 and P2 Backplanes
- * Extender Boards
- * Prototyping Boards
- * Card Frame Systems
- * Power Supplies
- * Connectors



BICC-VERO ELECTRONICS INC

40 Lindeman Drive, Trumbull, Connecticut 06611 U.S.A. Tel: (203) 372 0030 Fax: 203-372-6224. Twx: 510 227 8890.

CIRCLE 50

VME BUS SUPPORT HARDWARE

Multibus® II is the Reg. Trade Mark of Intel Corporation
Euro-Q Bus® is the Reg. Trade Mark of Digital Equipment Corporation

Leaders in Microprocessor Backplane



CONVERTING GATE ARRAY DESIGNS TO STANDARD CELLS

With the freedom to easily and inexpensively transform a gate array into a standard cell, designers can delay decision-making and improve production efficiency.

by **William J. Wolf**

The ability to convert a gate array into a standard cell provides an attractive design option. Designers can choose between a gate array and a standard cell after the logic design is completed and ready for placement and routing. This eliminates the need for vendor selection and commitment to a gate array or standard cell design at the outset.

Often when this choice is necessary in early stages, it must be adjusted frequently as project requirements change. The chip size can grow, for example, exceeding the gate count of the intended gate array. And, the next size could be too expensive or the wrong size for the package. Or, performance can be too low on the gate array, and a high performance standard cell macro could solve the problem. On the other hand, if a standard cell design takes too long, gate array prototypes can be necessary. With these problems and more cropping up during the product development cycle, a twelfth-hour decision on whether to use a gate array or a standard cell can be key for efficient product development.

Until recently, however, converting a gate array to a standard cell chip involved a long, expensive,

William J. Wolf is the manager of standard cell design at General Electric's Custom Integrated Circuit Department (Research Triangle Park, NC). He holds a BS in electrical engineering from the University of Illinois, Urbana.

and often risky redesign of an existing part. While converting gate arrays to standard cells is fairly simple in principle, many factors must be considered in practice: volume and price, gate utilization on the array, number of I/O cells, packaging, and overall design complexity. In addition, the ability to make this conversion will alter the traditional trade-offs of different application-specific ICs.

Differences in processes, cell libraries, computer aided design tools, and vendor interfaces are the most frequently encountered problems in converting a gate array part to a standard cell part. The impact of these differences is softening, however, as semiconductor vendors begin to offer similar 2-level metal CMOS technologies, cell libraries, and software macro capability. In addition, if a customer supplies the vendor with the preferred interface—an objective specification, a simulated netlist, and a test vector set—most of the work necessary for the conversion process can be handled by the vendor.

Drawbacks to chip conversion

Although chip conversions by vendors are becoming more common, drawbacks still exist. To the vendor, the conversion task is equivalent to a new design, and the customer is charged accordingly. Developing new macros may be necessary to account for differences in the customer's and the vendor's macro libraries, and/or the vendor could request design alterations. Different processes and cell performances could cause major timing problems. Another drawback is that not all vendors have the CAD

capabilities that are critical for this conversion process, such as post-routing delay analysis (needed to guarantee chip performance) and timing verification (needed to isolate timing problems).

To overcome these drawbacks, General Electric's Semiconductor Division has developed the processes, CAD tools, and user interfaces for both its gate array and standard cell implementations along similar lines. The identical 2-level metal, 2-micron channel length CMOS process is used for both the IGC20000 family of gate arrays and the ISC20000 family of standard cell chips. The ISC20000 standard cell library is upwardly compatible with and functionally identical to the gate array library. Cells have an equal or greater performance level and take less area than their gate array counterparts. The same CAD system is used for designing IGC20000 gate arrays and ISC20000 standard cells. Only the CAD system programs for automatic placement and routing, optimized standard cell verification software, and interfacing are reserved for standard cell designs. In addition, the interface between a customer and GE

Semiconductor is identical whether the chip will be a gate array or a standard cell.

As a result of these similarities, the conversion of a gate array design to a standard cell implementation has been reduced from a complex redesign project to a simple and straightforward software task. Most importantly, the conversion becomes the responsibility of GE Semiconductor and involves minimal engineering effort by the customer.

Conversion of macros

While a one-for-one macro conversion is the simplest form of gate array conversion, it does not always result in an optimal standard cell chip. Some complex macros in the standard cell library are implemented as custom cells for high density and/or high performance. Not all of these complex macros are available in the gate array macro library, but the functions of these custom standard cell macros are available as FLEXICELLS™ on IGC20000 gate arrays. Converting a netlist with a gate array FLEXICELL to a netlist with an optimized standard cell

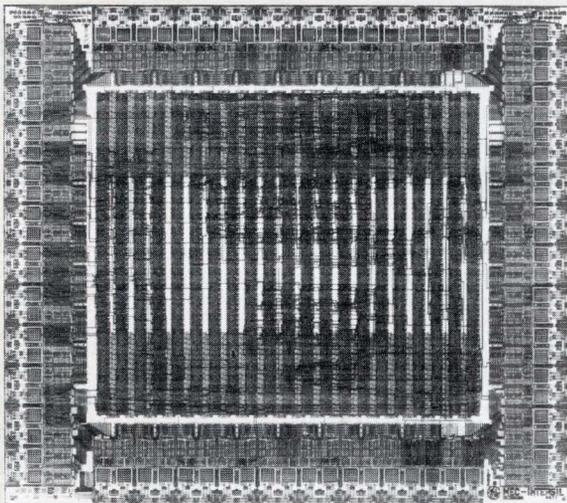
A conversion benchmark

A gate array to standard cell conversion done for a customer with a 20-MHz, 1018-gate, 36-I/O design is typical of the efficiency possible with the approach used by General Electric's Semiconductor Division. The design is implemented on a gate array to allow the customer to enter the market as early as possible. It is too large for the IGC20100 gate array (1023 gates), however. On the IGC20190 gate array (1984 gates), the chip used 54 percent of the cell array with a die size of 205 x 185 mils (a).

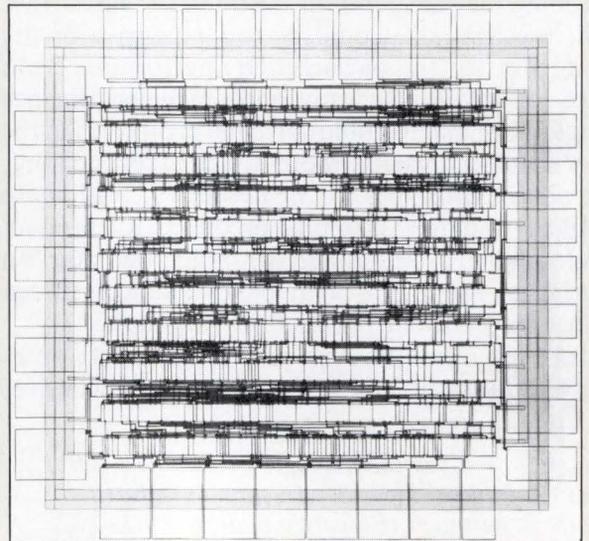
First prototypes were fabricated, tested to the customer-supplied vectors, and delivered. Under system test, the customer discovered a mode that

was not tested by the vectors. A minor logic change was required, and a second gate array was designed. Because the change was so simple, the standard cell conversion was started while the gate array revision was in fabrication. Had the design changes been more complex and the risk higher, a decision to wait for verification of the second gate array might have been made.

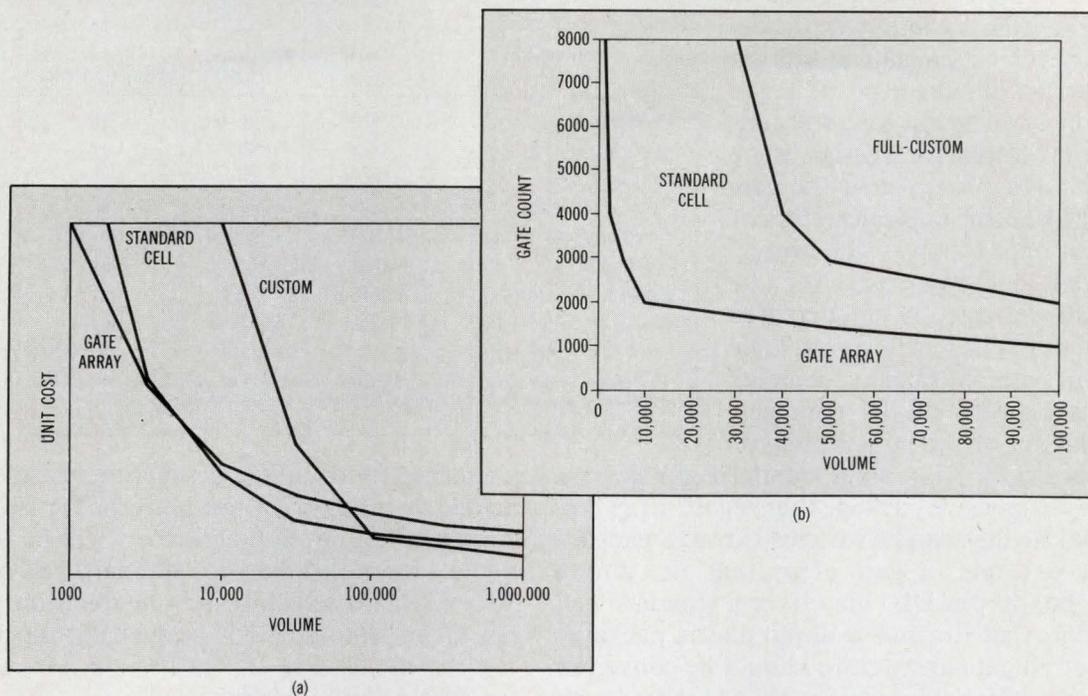
The resulting standard cell die size is 124 x 117 mils—an area savings of 62 percent (b). Three design factors made this large savings possible: low gate use on the gate array, no use of the analog cells on the gate array, and a high percentage of flipflops.



(a)



(b)



A conventional cost comparison shows that standard cell and full-custom designs have an edge at high volumes. Gate arrays are best when production volumes are low (a). Cost comparison, however, does not account for differences in chip complexity. Comparing gate count versus volume produces regions for measuring the effectiveness of each solution. Standard cells and full-custom designs are best-suited for high volume, high complexity chips and gate arrays for low volume, low complexity chips (b).

macro can be accomplished using a text editor and name substitution. The rest of the conversion process is the same as a one-for-one macro conversion.

FLEXICELLS are higher level functions composed of a collection of logic elements from the IGC20000 gate array macro library. Logic implementations, netlists, placements, and performance are predefined and documented for customers. The routing and performance of these cells are predictable because of some unique features of the router and CAD system. In addition, all of the logic elements available in the IGC20000 gate array macro library have counterparts in the ISC20000 standard cell macro library, so the same FLEXICELL designs can be used for both families.

Many FLEXICELLS are functionally equivalent to 7400 series parts. By providing access to previously characterized MSI functions, these cells simplify the overall design process. Design flexibility is also increased through the optional ability to slightly modify the logic implementations of a FLEXICELL macro. These standard cells have nearly all the features of hardware macros, and designers can customize some or all of the FLEXICELL attributes to meet specific functional or performance objectives.

The conventional cost comparison for gate array, standard cell, and custom implementations of a design examines the unit cost in relation to the total volume of chip production. This is a simplified view and assumes a chip of a particular complexity. Recently, however, a comparison of complexity (number of gates), production volume, and the region (in terms of volume and gate count) where the most effective design solution can be realized has come into vogue. From this perspective, gate arrays offer the most cost-effective design solution for low volume designs with a low number of gates. However, as volume and gate count increase, standard cells become more cost effective, followed by custom circuits for high volume, high complexity chips.

Criteria for conversion

While either view of costs can provide guidelines for choosing the best implementation, a designer should consider the particulars of a chip before forging ahead. The die size reduction possible with standard cells is critical for large chips where even small differences in area can mean large differences in yield. Also, the silicon cost for large chips becomes a larger portion of the packaged and tested part.

Savings in die size of standard cells over gate arrays is typically 20 to 50 percent. A wide variation exists because of differences in gate utilization on the array and the type of logic. Area savings will be less for a fully utilized gate array than for a design where the gate count just exceeds the gate limit and forces the implementation of the design onto the next size array.

Different logic types also effect the savings in die size. A smaller area reduction is realized with simple logic gates than with flipflops or complex functions because of the compact, custom layout of the large functions in standard cells. Also, since standard cell die may fit in a less expensive package than a gate array die, additional savings can be realized through packaging. For example, a gate array that fits into a side-brazed, 40-pin DIP may have a standard cell counterpart that fits into a 40-pin plastic package.

But not all gate array chips should be converted to standard cells. Savings might not be realized with pin-limited chips that have a gate array master available with the correct number of pins, for example. While more logic functions would be possible in a standard cell chip implementation of a pin-limited gate array, the additional logic would only be useful if the system architecture allows upgrades with the same pin configuration and board layout, or if the total system is upgraded.

In addition, development costs for a gate array followed by a standard cell conversion are higher than for an initial standard cell design. If an iteration

is required for the initial design (because of a system change or insufficient test patterns for simulation, for example), development costs will be lower for a gate array followed by a standard cell chip than for a standard cell chip and one iteration. And the gate array prototypes will be available much earlier.

The best of both worlds

Developing a gate array, taking it through a subsequent design iteration, and then converting it to a standard cell costs about the same as developing a standard cell and then taking it through an additional design iteration. The simplest way to ensure low cost, either initially or after production of gate array prototypes, is to choose a vendor with the proven ability to implement the same design as either a gate array or as a standard cell chip.

Using a gate array followed by a standard cell is an option that combines the best features of the two products. More than this, the various options of gate arrays, standard cells, gate arrays followed by standard cell conversion, and standard cells with custom blocks are making custom circuits a thing of the past. The only major drawback to conversion is the limitation of special functions on the original gate array implementation. But this can be offset by using standard cell libraries with custom blocks that are appropriate for applications requiring special functions. It seems logical that the time and money needed for a custom chip might be better spent on architectural trade-offs, logic optimization, and a few key custom blocks.

Design Style				
Conversion Criteria	Gate Array	Standard Cell	Gate Array and Standard Cell	Custom
Performance	High ¹	High	High	High
Prototype cost	Low	Medium	Low	High
Part cost	High	Low	Low	Lowest
Total development cost	Low	Medium	Medium	High
Development time	.2-.3 year	.3-.4 year	.2-.3 year	1-2 years
Automation	High	High	High	Low
Special functions	Limited	Yes ²	Limited	Yes
Risk	Low	Medium	Low	High

1 Lower gate array circuit performance tends to be offset by early access to advanced technology.
2 Advanced standard cell CAD allows for memory and special custom blocks.

ISC20000 Standard Cell Chip Series	
Technology	2-micron, 2-level CMOS
Performance	1.7-ns delay for a 2-input NAND gate dc to 20-MHz system clock
Complexity	25,000 gates per chip
Supply	2.0 to 5.5 V
Cell library	23 CMOS and TTL I/O functions 45 SSI/MSI cells 100 FLEXICELL functions for 7400/4000 conversion 2 static RAMs—variable size up to 16 Kbits 1 ROM—variable size up to 128 Kbits
CAD system	TEGASStation, Daisy or TDL input Highly efficient auto-place and route Variable height cells and large blocks Accurate post-routing delay simulation Automatic test code conversion
Gate array conversion	Automatic conversion of IGC20000 gate array up to 13,000 gates

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

The best way to test circuit boards is by comparison.



The Fluke 3050B gives you the power of Reference Board Testing.

With the Fluke 3050B Analog/Digital Functional Test System, you can test entire microprocessor-based boards, detecting even the most difficult-to-diagnose functional faults. The result is confidence levels of at least 98%. You get higher system yields, reducing your defective inventory as much as 70% and speeding deliveries.

Our Reference Board test method compares the unit under test to a known good board, ending the need for extensive fixturing and programming. So you can test more types of boards at a lower cost than any other method.

Each Reference Board becomes an instant test library, as real-time responses

allow the 3050B's powerful diagnostic algorithms to exercise the board at rated microprocessor speeds.

Automatic diagnostic software tests bus-structured boards, determining what device is driving the bus at the time of failure. With advanced loop diagnostics and the ability to break feedback loops, the 3050B identifies even the most difficult-to-find faults.



Extend the system by including an off-line programming station, analog test station, touch-sensitive display and a high-speed printer for statistics.

The 3050B starts around \$100,000, complete with the software, applications support and service you would expect with more expensive systems.

So for the best and most cost-effective ATE solutions look to Fluke. Make your comparison. Then contact your local Fluke Representative, or call toll-free **1-800-426-0361**.

IN THE U.S. AND NON-EUROPEAN COUNTRIES:

John Fluke Mfg. Co., Inc.
P.O. Box C9090, M/S 250C
Everett, WA 98206
(206) 356-5400, Tlx: 152662

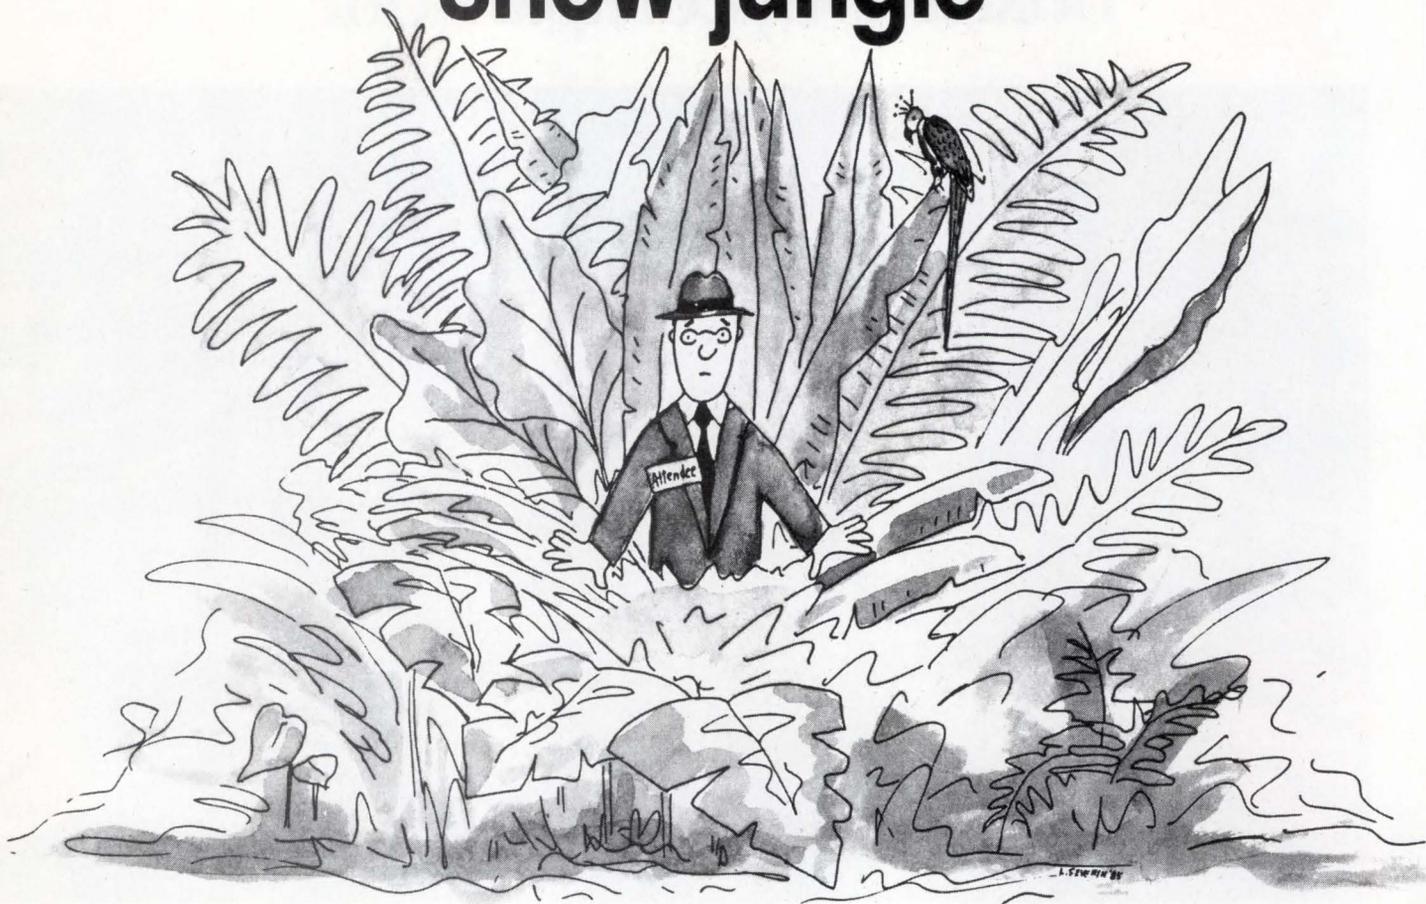
IN EUROPE:

Fluke (Holland) B.V.
P.O. Box 2269, 5600 CG
Eindhoven, The Netherlands
(040) 458045, Tlx: 51846

CIRCLE 51

FLUKE
®

Come out of the computer show jungle



Come to the Show with Over 5,000 Products Just for Digital Equipment Computers.

You've been to computer shows that promise something for everyone. The trouble is finding the "something" for you amid thousands of products for everyone else.

DEXPO South 85 is the DEC user's smart alternative. Because at DEXPO, every product on display is compatible with Digital Equipment computers. So you'll spend your time actually *examining* DEC-compatibles ... hardware, software, systems, and services ... not searching for them.

Ways to improve every DEC computer from VAX to Rainbow. Whether you have one DEC computer or an entire range, you'll find just what you need to enhance your system's productivity. From terminals and memory boards to turnkey systems; from applications to communications — DEXPO South has it all.

With over 300 vendors on hand, you can be sure of seeing the latest advances in DEC-compatible technology. So, you'll have plenty of opportunities to compare and evaluate products *before* you choose.

A bonus for DECUS* symposium attendees.

If you attend the DEC Users Society (DECUS) New Orleans symposium, you'll receive FREE admission to DEXPO South 85. Use the attached reply card to get complete details.

Discover the show that helps you do more.

Get a headstart on making your DEC computer work harder for you. Simply mail the attached reply card to receive more information on DEXPO South 85 ... plus money-saving V.I.P. admission tickets for you and your associates.

If the reply card is missing, write or call: Angela Wright, Registrations Manager, Expoconsul International, Inc., 55 Princeton-Hightstown Road, Princeton Junction, NJ 08550; (609) 799-1661.

Come to DEXPO® South 85

The Seventh National
DEC*-Compatible Exposition

New Orleans Superdome

May 28-31, 1985

Tuesday-Friday

PROGRAMMING THE UNCERTAIN WITH EXPLORATORY SYSTEMS

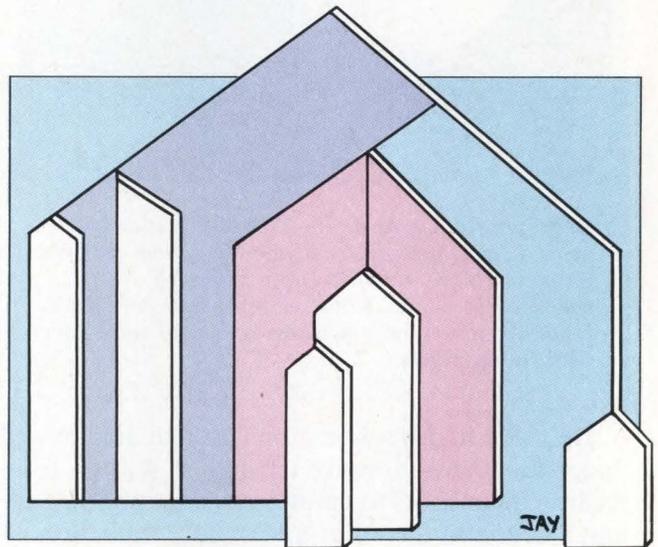
Artificial intelligence development systems are ideally suited to the coding of complex programs because they can cope with the uncertainty that is often present.

by Beau Sheil

Exploratory programming systems, previously confined to the artificial intelligence field, are becoming attractive for many commercial applications. Since they allow designers to defer programming decisions or to avoid them entirely, these systems can cope with rapidly changing, very large application programs. The need for this kind of programming system has come about because some projects are simply not solvable with more conventional programming methods.

For many applications, the potential user may not know, or be able to anticipate, exactly what is required. In this situation, no amount of interrogation of the intended user and no amount of exercises on paper will completely pin down the possible problems. The designer must work by trial and error. AI technology favors this approach because it permits more designs to be developed and explored. But, to achieve this flexibility the acts of programming and designing must become closely intertwined.

The one certainty in the AI field is uncertainty—one of the few things known at the outset of a project is that it is impossible to predict the technical solution needed to solve the problem. As a result, both AI programming style and the hardware technology that supports it have been built around the notion of programming under uncertainty. According to AI pioneer Marvin Minsky, who defined intelligence as

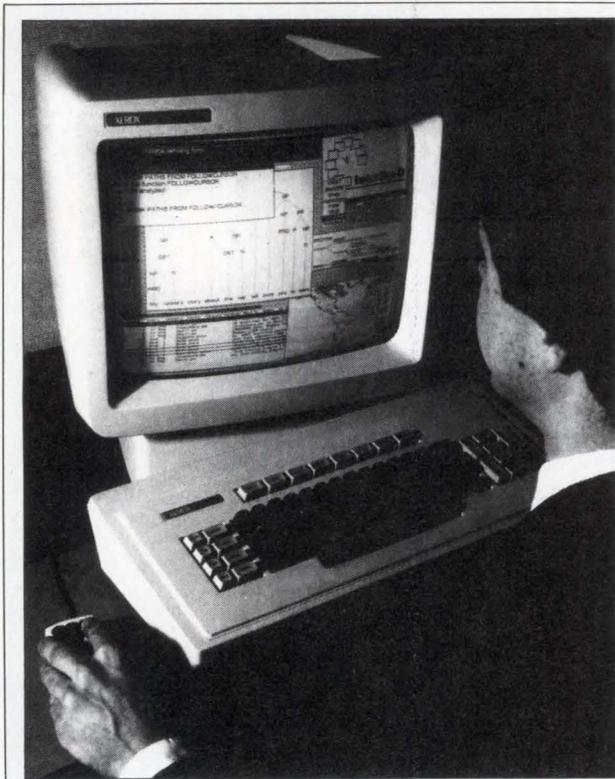


“an attribute we ascribe to behavior that we admire, but that we do not understand,” this problem is endemic to the whole of AI.

With this philosophy, a programmer attempting to develop a program that exhibits intelligence must first discard one of the foremost rules of programming: Understand where you're going before you start coding. This rule is fundamentally at odds with the nature of AI programming. The entire structure of conventional software engineering is centered on understanding the consequences of the design before beginning to code.

The AI field, however, operates under a fundamentally different set of assumptions. First, if a problem is very well understood at the outset, it is, by definition, outside the AI domain. For example, AI people no longer write programs to play backgammon because the game is so well understood that a computer can now beat the reigning champion.

Beau Sheil is manager of product development for Xerox AI Systems (Palo Alto, Calif). He holds a PhD in computer science from Harvard University.



Screen images on the Xerox 1108 workstation show some of the exploratory programming tools provided in the machine's Interlisp-D programming environment. Under mouse control, each reshapable, repositionable window provides a view of some data or software process.

The AI field deals with problems that are not well understood. Regardless of the subject matter, from medical diagnostics to natural language understanding, a common thread of uncertainty runs through all AI studies. The reasons range from sheer complexity (circuit designers cannot anticipate all the ways their growing set of design rules will interact), through continually changing requirements (information bases change), to subtle and/or unpredictable human factors. To deal with this uncertainty, AI programming has developed a very distinctive, flexible software technology. To contrast the AI approach with that of conventional software development, it is helpful to make an analogy to the civil engineering problem of bridge building.

Conventional engineering practice provides many methods for ensuring bridge rigidity so that a bridge maintains a straight span from one river bank to another. But, consider what happens if the design engineers are faced with the problem of building a bridge through a fog bank. If rigidity is built into the design method, the risk is that the bridge will miss its mark on the other bank. Under these circumstances, the bridge should be built differently,

with methods that allow the bridge to bend so that midstream corrections can be made.

The same principle applies to AI programming. The technology must allow the programmer to take advantage of opportunities and minimize the cost of making mistakes—since there is no way to avoid them anyhow. What has developed in AI is a technology for building software experimentally and discovering the nature of the problem along the way.

Programmers in many disciplines are finding that their problems are similar to those encountered in AI applications and thereby amenable to the same solutions. AI simply represents the far end of the spectrum of uncertainty—the very opposite of problems that are completely defined and understandable, such as accounts payable and general ledgers.

Conventional programming methods will not work in AI applications. But, they do not work in many other areas either. Many programmers find themselves constrained by a lockstep technology that cannot handle the uncertainty characteristic of many complex problems.

There are many such problems in the real world. For example, financial planners are often forced to respond quickly to a variety of unpredictable political, economic, and organizational changes. A programmer building systems to accommodate these fluctuations is in much the same position as an AI programmer. In any case, where systems are built to interact with professionals in a relatively unstructured environment, where decisions are buffeted about by the external world and where the ground rules are constantly shifting, AI programming technology is appropriate.

Finding a real-world solution

Traditionally, putting computing power in the hands of professionals has proved an elusive goal for two reasons. Specialists above a certain rank have shown reluctance to accept any form of computerization because the application software designed for these users fails to accommodate the level of uncertainty and complexity they require.

A workstation that is going to help an executive or professional must fit like a glove. For example, such an application was developed by Schlumberger Ltd (New York, NY), the world's largest oil geology consulting firm. The company wanted to automate at least part of the analyses made by petroleum geologists concerning oil well drilling data. Such experts are a rare breed—perhaps only a few hundred of them are employed throughout the United States—and millions of dollars ride on their decisions.

The development of this system is a classic example of programming under uncertainty. Since no one had ever built such a comprehensive information

management system for these users, many factors were not clear: what kind of interfaces or summarizing tools were required, how the data was to be organized, and how fast the system needed to work. These unknowns, coupled with the mountains of data the geologists faced and the time-honored method of analyzing this data "by hand," put an additional burden on the programmers designing the system.

This example is but one of many. Although the number of potential users in any given category is small, there are virtually thousands of categories of professionals requiring this kind of "ultra vertical" software—commodities traders, underwriters, financial analysts, and more. The approach to building very specialized systems for these professionals contrasts with that used in designing more standard systems for a large population of users. Some firms, for example, can afford to spend many years and many millions of dollars developing general office automation systems because the potential market is

so large. But, no company can afford such a protracted development cycle to produce an equally well trained system for a small number of specialists, such as circuit designers or futures traders.

The development of highly tailored systems for these small vertical markets requires an entirely different programming approach in order to be cost effective. It requires the kind of flexible, open-ended, fast-track development that was once the exclusive trademark of AI programmers. The development of these systems represents a major area of long-term growth for AI technology. Of course, the question remains as to why the AI software environment encourages a rapid development cycle.

Considering the languages

As a whole, AI programming languages have been designed to minimize programming constraints. This philosophy is consistent with the goal: If you know that you are going to make mistakes during the

Using assistant, not executive, systems

One of the principal aspirations of AI research has been to replicate expert human performance in many complex tasks. But, in the commercial arena, the trend is no longer toward "executive" systems, but toward "assistant" systems. To understand the difference, consider MYCIN, an early medical diagnosis system designed at the Stanford University Medical School (Palo Alto, Calif) by Professor Ted Shortliffe's group. MYCIN was an attempt to diagnose infectious diseases by running what looks to the user like a standard diagnostic interview.

The system was reasonably good at making a diagnosis within its area of expertise, and its success launched the entire field of expert system technology. But, the problem for actual use in the field was that the system was cast in the role of doctor. It asked the questions that a doctor would ordinarily ask and tried to automate the doctor's reasoning. From the research point of view that's a good goal. The practical side has a fatal flaw: the system's knowledge may be incomplete, but its design holds no place for the human expert and its unattended operation may be catastrophic.

By contrast, consider the system that Shortliffe and his group are developing now. Called ONCOCIN, it is intended for use in oncology clinics prescribing treatment for cancer. This system looks entirely different from its predecessor.

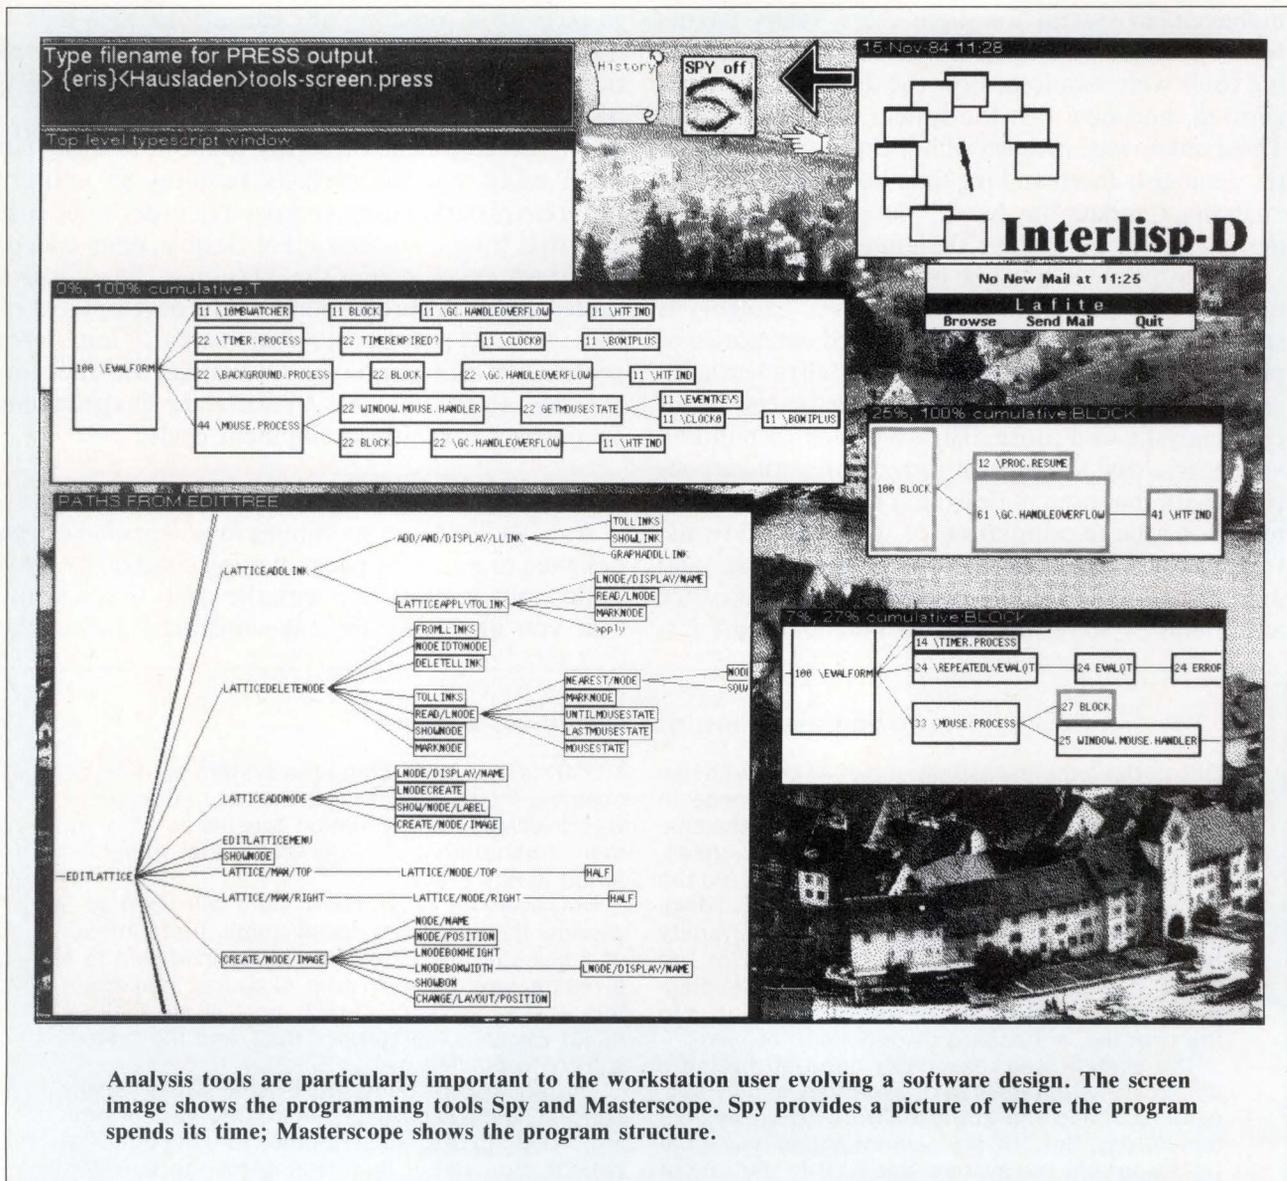
Instead of conducting an interview, the system looks like a smart forms system—it brings up an exact copy of the forms doctors currently use to record chemotherapy and radiation treatment, including diagrams for illustrating the tumor sites, and the application of radiation treatment. In addition, the system knows enough about the meaning of the entries that it can screen for impossible values. What truly sets ONCOCIN apart from

MYCIN is that in watching the system operate, an observer might never realize that any intelligence was involved. The application appears no different than a standard forms generator that might be found in any office.

But, every so often, the system will open up a window that suggests, for example, that patients of a given profile have often responded well to a given treatment. The doctors, of course, may reject this advice or want to know more about it. They might ask for a literature citation, and the system will try to provide it.

This approach is quite subtle because it is literally performing the same skill that MYCIN was performing: matching diagnostic profiles to a hypothetical patient. But, rather than asking for the patient's symptoms and attempting to drive the interaction, the system merely tags along behind—asking occasional questions, making suggestions, and generally acting in the capacity of assistant. The technology is identical; the deployment of the technology is different. Clearly, an assistant system makes no bones about who is in charge. The merit of this approach is that the knowledge of the system is available to the user, but it is the user who is accountable and must take the initiative.

The notion of designing systems that function as assistants is powerful and can be used in many applications. Unless the expert technology is strong enough to preempt an expert's judgment, human experts must be kept in the loop. But, even though people cannot be replaced, the nonperfect expert system technology is not useless. Instead, care must be exercised in the way the technology is deployed. The helping paradigm provides an effective method of deploying the limited expertise supportable by contemporary AI technology.



Analysis tools are particularly important to the workstation user evolving a software design. The screen image shows the programming tools Spy and Masterscope. Spy provides a picture of where the program spends its time; Masterscope shows the program structure.

development cycle, the language you use should make it as easy as possible to change your mind. Thus, the languages enable the programmer to defer decisions until the last possible moment, rather than committing to them at the outset.

Automatic storage management is a classic example of this. These techniques allow a programmer to allocate storage freely and require the system to reclaim it when it is no longer used. Virtually all AI languages provide this facility. Without it, the language would place an intolerable burden on programmers to track all the paths through their programs that might access any given piece of storage. Automatic storage management ensures that none of the paths access or release that storage prematurely. The goal here is to decouple pieces of the program, cutting through as much of the internal connectivity as possible.

A second example of deferred decision making is the dynamic typing of variables—the ability to avoid prespecifying the type of information a variable will store at run time. In Pascal, for example, this chore is one of the first things a programmer does.

Freedom to defer the data typing decision lets the programmer experiment with the types of structures used. Indeed, the first drafts of an exploratory program usually begin with generic (and consequently inefficient) data structures. As the application progresses, experimentation proves one type of structure to be superior to another. Hence, the data typing decision is often one of the last to be made, rather than one of the first.

A third example is the ability to make free reference to variables—that is, to have a procedure that may use variables not bound in any lexically enclosing scope. The binding accessed depends only on the

INDUSTRIAL GRADE IEEE-696/S-100 BOARDS

Dual Systems designs and manufactures a variety of IEEE-696/S-100 boards for 16-bit microprocessor systems running under UNIX and other operating systems. These boards bring high performance and three years of field-proven experience to your computing environment.

Each board is rigorously tested and burned-in for 168 grueling hours. If it can't bear the heat, it won't bear our name.

High Performance System Boards

Model WDC-SMD The WDC-SMD Hard Disk Controller* is specially designed for high throughput in large, heavily-loaded multi-user UNIX systems. All sectors on a track are transferred essentially within a single disk rotation regardless of where the head first settles or the order in which sectors are encountered.

The controller offers 16-bit throttled DMA data transfers and disk transfers up to 10 Mb/sec. Also features dual-ported, full-track, look ahead cache, and on-board microprocessor. Interfaces with one or two SMD drives. \$2195.

Model SI04-DMA The most advanced, intelligent, 4-port serial I/O board available for the IEEE-696/S-100 bus, this module features 256 bytes of FIFO buffer for input characters and provides DMA transfers for output. A built-in 8085A processor greatly reduces system overhead. \$695.

Model DMEM Features 256K bytes of memory and either 8 or 16-bit data paths. 24-bit addressing, and parity checking on each byte. DMEM has no S-100 wait states. \$1395.

Model EPROM Capable of either 8 or 16-bit data transfers, this 32/64K EPROM offers the versatility of running with 68000, Z-8000, 8086, 16000, and other 16-bit processors. It accepts industry-standard 2732 and 2716 EPROMs. 64K RAMS may be mixed with 2716 EPROMs for use as a RAM/EPROM board. \$345.

Model CPU-68000M High-performance CPU board with 16-bit data path, 10 MHz CPU operation, and MC68451 MMU for multi-tasking applications. \$1195.

Model CPU-68000 Similar to 68000M, but features 8K bytes of on-board ROM with Motorola's MacsBug monitor instead of the Memory Management Unit. \$895.

Models M/BD-15 & 20 Back Planes These premium quality motherboards feature four-layer construction with two internal ground planes, and Schottky-diode termination. They provide high-speed operation with true transmission line characteristics and minimum noise. M/BD-15: \$495, M/BD-20: \$545.

Model CMEM This non-volatile CMOS memory board provides easy-to-use 8 or 16-bit data paths and 32K bytes of memory with dynamically movable write/protect window. On-board lithium battery holds data for 3-10 years with power off. \$725.

Data Acquisition and Control Boards

Model CLK-24C Clock-calendar features a LSI CMOS chip and on-board, long-life lithium battery. \$325.

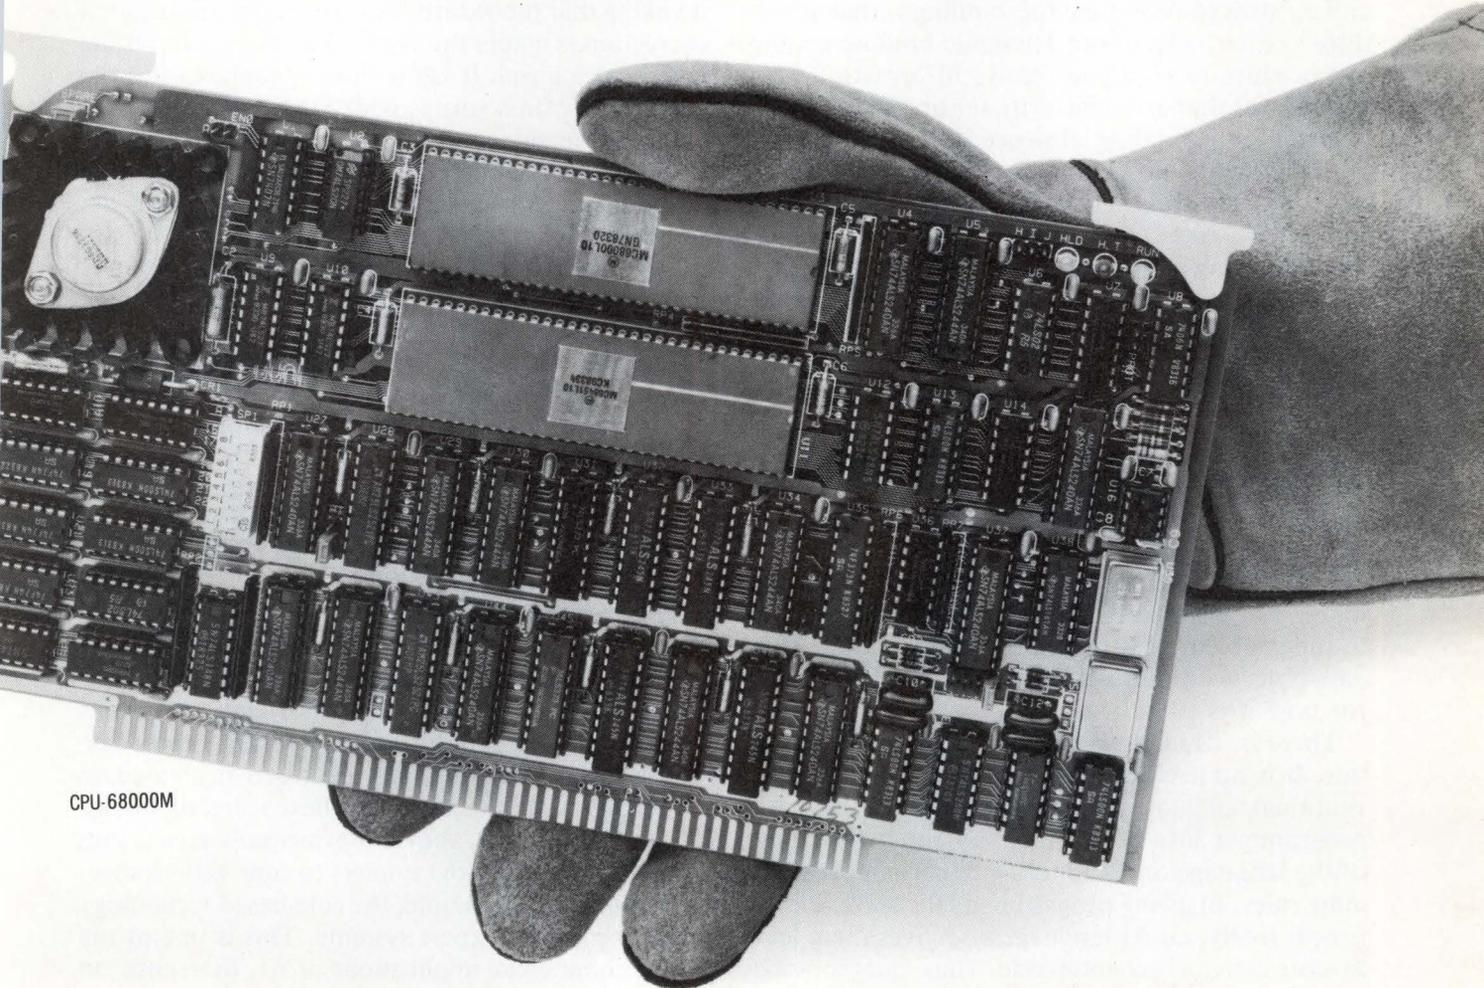
Model AIM-12 A highly reliable A-to-D converter with 35 μ sec. maximum conversion time, 12-bit resolution and accuracy, and 32 channels single-ended/16 channels differential. \$725.

Model AOM-12 This D-to-A converter offers I/O-mapped port address, 12-bit $\pm 1/2$ L.S.B. accuracy (0-70°C), and voltage outputs of 0 to 10 volts, ± 5 volts, and ± 10 volts. \$675.

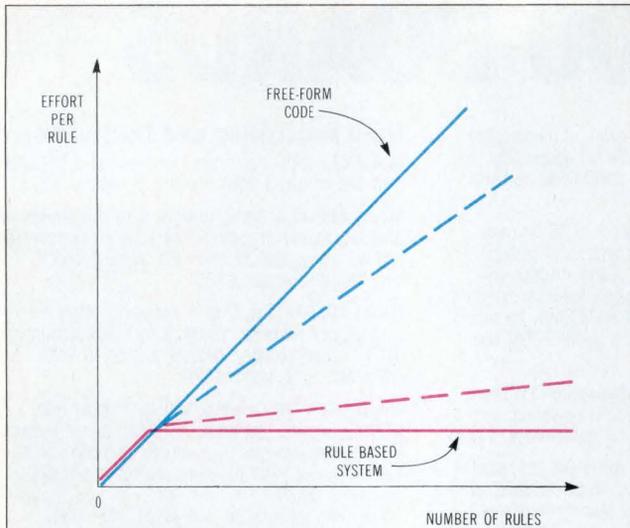
Model VIC 4-20 Converts voltage outputs from AOM-12 into four separate 4-20MA current outputs. Module also provides overvoltage protection on all current output, plus transient protection per ISA standards. \$600.

For more information, call (415) 549-3854. Dual Systems Corp., 2530 San Pablo Avenue, Berkeley, CA 94702

THERE'S ONLY ONE
DUAL



CPU-68000M



With free-form code, adding a new rule forces the total effort (area under the curve) to rise approximately as the square of the system size. Using rule-based technology, the total effort is approximately linear in the system size. In reality, the effort required in each case is close to the dashed lines.

calling procedure's variable bindings, that procedure's caller, and so on. Dynamic binding enables a procedure to send and receive information from procedures that are arbitrarily far up the call stack.

Finally, the binding of procedure calls is usually dynamic in that procedures are indirectly addressed by name only and their actual definitions can vary. In AI languages, this process implies more than merely linking a set of procedures at load time. The programmer is usually able to dynamically change the subprocedures invoked by a given piece of code, simply by altering the runtime context.

The simplest example of dynamic binding is the use of procedures as arguments, or as the values of variables. Beyond that, procedure values can be computed, or even contained within the data values on which they are to operate. This packaging of data and procedures into a single object—known as object-oriented programming—is a powerful technique. For example, it provides an elegant solution to the problem of generic procedures in that every data object contains within itself the "instructions" for how it is to be used.

There is, as might be expected, a trade-off in all this. It is no accident that the trend for more conventional languages has been toward locking the programmer into a structure, by either the design of the language or by adhering to certain programming rules. In giving programmers the leave to experiment freely, an AI language also gives them leave to write extremely chaotic code. Thus, questions may

be asked as to how the programmer can escape the tyranny of all this freedom, and where you reintroduce the certainty that structured programming techniques bring.

The answer lies in the programming environment. The Xerox (Pasadena, Calif) Interlisp-D system, for example, contains tools that let the programmer reconstruct the structure of an application program dynamically. In other words, instead of insisting that a program structure be specified before a line of code is written, the programmer is provided with the tools that allow the reconstruction of the program's structure at any time during its development.

What makes this possible is that the system "knows" what it is doing. Unlike a simple text processor that can be used to construct source code, as well as write company reports, Interlisp-D knows the user is programming in Lisp, knows how to read Lisp programs, how to look at changes made in those programs, and how to make the necessary reconstructions in order for those changes to take place.

For example, if a line on a graphics application is being drawn incorrectly, the programmer can ask the system to determine what procedure is doing the drawing and can locate which parts of the code are invoking that procedure. Moreover, once the desired correction is made, the system knows which routine has been changed. It knows how to replace the compiled code with a source code version and to interpret that version on-the-fly.

The system will also alert the programmer if that modified code is called by any other piece of code. For example, if a data structure is changed, the system is smart enough to seek out all the places where that data structure has been used and update them automatically. Finally, when the designer is done with the session, the system will "clean up" and put the entire design software "world" back together again.

In other words, the system assumes a substantial amount of the bookkeeping and frees programmers to be creative. By providing supportive power through smart tools, programmers regain much of what was lost in abandoning a structured methodology.

Coping with large problems

In addition to adding structure, AI programming techniques offer a second major advantage. They enable programmers to cope with large programs that change very rapidly—often in unforeseen ways. These program techniques are often presented as important statements about the nature of intelligence. But, in fact, they are principally techniques that better enable programmers to cope with change.

Consider, for example, the rule-based technology used to produce expert systems. This is one of the prime commercial applications of AI. In a sense, an

expert system used for medical diagnosis, for example, is merely a complex decision tree: IF the patient's temperature is high, and IF he is dehydrated, THEN the patient has a fever.

AI techniques let programmers cope with large programs that change very rapidly, and often in unforeseen ways.

In fact, if the system yields an incorrect answer, the doctor using the system might contend that some other piece of information was missing (eg, the patient has sunstroke). The problem with a conventional system is that its knowledge is "cast into stone." Each time an unforeseen exception arises, the programmer must go back over the code to hunt for every path that might lead to this erroneous conclusion—and then patch in a solution. The problem increases quadratically with the program size because every program addition implies examination of everything that went before.

Housekeeping and rule-based tools

Rule-based programming tools alleviate this problem, again by doing much of what amounts to housekeeping. They allow a programmer to present a series of IF/THEN rules and incrementally (after each rule is added) arrange those rules so that the additional chunk of knowledge is integrated. This capability is essential for expert systems development. The programmer might sit down with a domain expert, a doctor for example, and extract nuggets of information, one-by-one. As with any complex endeavor, there are few absolutes and a thousand qualifications; every diagnosis has a host of exceptions. Thus, the programmer needs a programming structure that can gracefully add each additional nugget of information as easily as the first.

As mentioned, these programming techniques have commercial applications. Digital Equipment Corp (Maynard, Mass), for example, has a system called XCON that is used for configuring its VAX line of computers. XCON contains a large set of rules to indicate which components must go with other components in order to make them a properly configured system (eg, one that contains the appropriate cables for connecting all components).

Written in OPS-5, a rule-based language, XCON can be regarded not so much as an "intelligent" system as a software engineering response to a program task that defies classical software engineering principles. The inherent problem is that DEC is constantly introducing new pieces of equipment for the

VAX line. Imagine the problem of updating the rules if those rules were written in Fortran.

The rule-based programming environment allows the programmer to decouple the task of adding another behavior to the program from behavior that is already defined. The technology is, at its heart, a means of helping the designer cope effectively with large programs that grow in unanticipated directions.

Effective graphics no accident

If there is any single visible characteristic that marks an AI workstation, it is the extremely high quality, highly flexible graphics interface—an interface that does its job very well. The presence of this technology on the workstation is no accident. It is a direct consequence of the programming environment discussed. The graphics interface on an AI workstation is, itself, one of the best examples of the benefits of programming under uncertainty.

Of the hundreds of potential approaches for designing any part of an interactive interface, only a few will turn out to be truly workable. And, the only way to effectively weed out the inferior 99 percent is by direct experimentation. There is no substitute for hands-on experience in determining the best approach. For example, the use of tree layouts to show semantic dependencies turns out to be a remarkably powerful graphics paradigm. But, that was not realized at the time the technique was developed. Only experience showed this particular approach to be superior to dozens of others.

Clearly, an environment that facilitates a high level of experimentation is ideally suited to nurture this kind of technology. The reason AI workstations have their high quality graphics interfaces is that different program variations can be easily assembled, tried out, and rejected or modified. It turns out that graphics is another area where designers can win big if they can experiment quickly.

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

High 713 Average 714 Low 715

CHOCK FULL OF CHIPS.



We've long been recognized for our tasty technology. But not many know us for our tempting selection.

Fact is, no one offers a more complete line of integrated circuits than Fujitsu: memories, microprocessors, microcontrollers, floppy disk controllers, LAN chip sets, codecs, flash converters, application specific ICs and more.

And while others make promises, we deliver. Now. In production quantities.

DRAMs. In MOS memories, we were the first to consistently ship production quantities of 16K, 64K and 256K DRAMs. With features that make your mouth water. Like our MB81256/MB81257 256K with both the smallest die size in the industry and the fastest address access time: 100 ns.

High-speed bipolar memories. We serve up PROMs in sizes from 256 to 64K bits. All made according to our unique DEAP™ recipe for high reliability.

ECL RAMs. Fully compatible with the industry-standard 10K and 100K series families. In densities from 1K bits at 7 ns to 16K bits at 15 ns. All available now. With a 16K at 10 ns version later this year.

Static RAMs. All the industry-standard high- and low-speed parts – from 35 ns to 150 ns address access times. In your choice of organizations and densities, from 4K to 64K bits.

EPROMs. We'll more than satisfy your appetite, from 64K to 256K. And 512K coming soon. In a choice of two flavors: CMOS and NMOS.

Microprocessors. We serve up a mouth-watering selection. Starting with our proprietary 4-bit microcontrollers offering the best value by function. And continuing with a complete line

of Intel-compatible (second source licensed) 8- and 16-bit microcontrollers, microprocessors and their peripherals.

Complemented by a choice of many Fujitsu-designed VLSI circuits. Including the industry's first Ethernet™ chip set, high-performance digital signal processor, disk and video/graphics controllers, and later this year – MBL80186 and MBL80286 microprocessors.

Gate Arrays. We supply more of the world's gate array requirements than any other manufacturer. With CMOS arrays in densities from 440 to 8000 gates. And bipolar arrays from 240 to 2000 gates. All supported by one of today's most sophisticated and prolific design systems, with well over 3000 designs already produced.

Packaging. Wrap it up any way you want. We offer the most comprehensive selection. From standard plastic DIP to LCC. PLCC to Cerdip. SIP to flat pack. And skinny DIP to high-pin-count grid arrays.

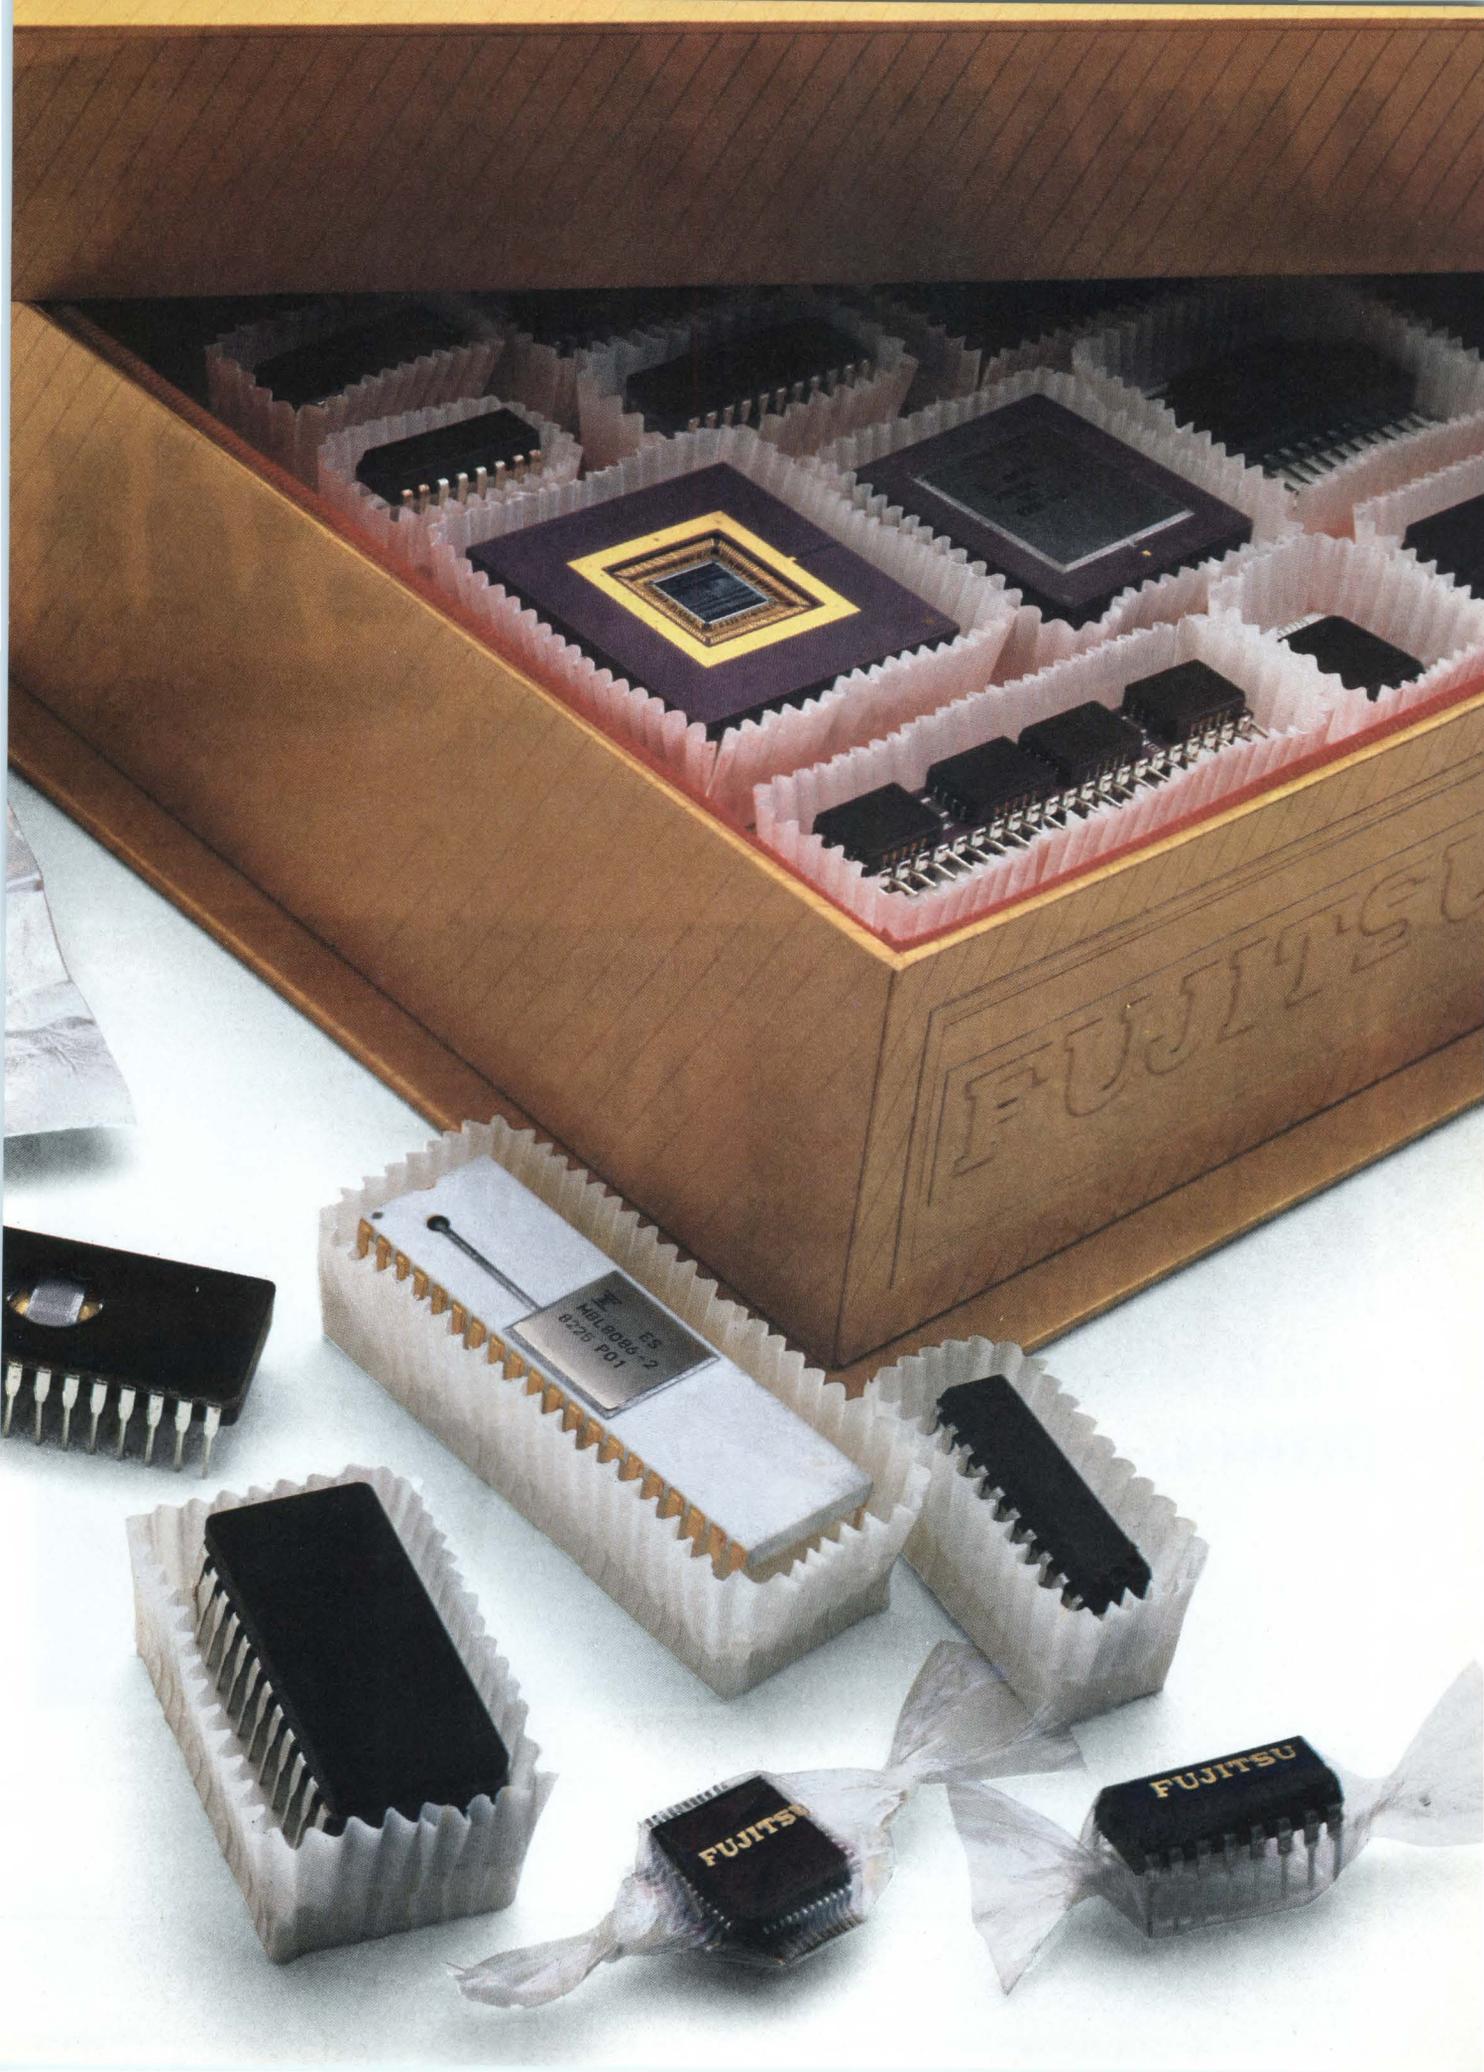
That's why, the next time you hunger for a part – any part – call Fujitsu first.

We'll satisfy your appetite.

**FUJITSU
MICROELECTRONICS, INC.**

Integrated Circuits Division
3320 Scott Boulevard
Santa Clara, CA 95054-3197
(408) 727-1700

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. So. Calif. 714/720-9688. © 1985 Fujitsu Microelectronics, Inc. DEAP is trademark of Fujitsu. Ethernet is a trademark of Xerox.



LEXAN[®]

resin



**It's a computer
that's going
places.**

**It's a printer
with staying
power.**



LEXAN[®]

resin

®Registered Trademark of General Electric Company.

™ Trademark of Texas Instruments.



It's whatever you need it to be.

Durability for portables. LEXAN resin delivers exceptional impact resistance for portable business equipment and computers such as the Texas Instruments Silent 700™ data terminal. Hardworking molded-in snap fits and bosses increase functionality and decrease assembly costs. LEXAN resin also offers significant weight reduction, for greater portability and reduced shipping costs.

Colors that don't quit. IBM utilizes LEXAN resin for its Model 5218 printer because of LEXAN resin's durability and excellent UV stability. Available in over 6,000 active colors, it provides an eye pleasing finish that resists fading and helps ensure that components will match.

Design Flexibility. With clear LEXAN resin, you can even use selective masking and painting techniques to consolidate parts and reduce fabrication costs. LEXAN resin is platable, too, for RFI/EMI shielding.

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 resin

Design Tips, just dial toll-free **(800) 422-1600**; in Vermont, **(802) 447-0069**.



**IS YOUR
PRESENT
DISK DRIVE
SUPPLIER
GIVING YOU
ANYTHING TO
LOOK FORWARD
TO?**

Let's hope so.

Because your systems depend on disk drives. So it's important to have a disk drive supplier you can depend on.

And that's where NEC comes in.

At NEC, we're always driving ahead.

With newer drives. Drives with more storage capacity. More reliability. Drives that keep pace with the latest technology.

How do we do it?

With more experience. We've been designing disk drives a long time. Way back in 1959, we were one of the first to develop magnetic recording devices.

Plus, we've got plenty of other



Our 5 1/4" drive gives you more capacity in less space, for less money.

can be sure there will be more.
Drive into the future with NEC.
You have a lot to look forward to with NEC.

NEC 3 1/2" floppy drives are ready for delivery.

NEC 5 1/4" half-height Winchester provide 12.91-25.83MB.

NEC 5 1/4" floppies offer up to 1.6MB capacity.

NEC 8" flexible drives are compact and have 24000 MTBF.

NEC 8" Winchesters store up to 167.7MB.

NEC 9" Winchester has a 15 millisecond seek time.



resources. We're an \$8 billion company. And that means we're in this business for the long run.

We just doubled the capacity of our 5 1/4" half-height drives.

No one thought we could improve on our 5 1/4" Winchester half-heights. But we did. We doubled our efforts. And we ended up doubling the tracks on our 5 1/4". The result?

You guessed it. 25.83MB instead of just 12.91MB.

If you need other disk drives, we also make a new 3 1/2" floppy, 8" Winchesters and floppies, and 9" Winchesters. And you

To find out more just call 1-800-343-4418. (In Massachusetts call 617-264-8635.) Or fill out the coupon below.

And find out why more and more OEM's are saying, "NEC and me?"

Please send me more information on:

Floppy Drives

- 3 1/2"
- 5 1/4" half-height
- 8"

Winchesters

- 8"
- 5 1/4" half-height
- 9"

Name _____

Title _____

Company _____

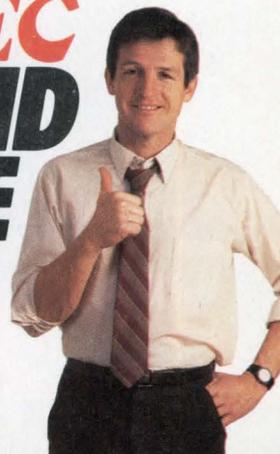
Address _____

City _____

State _____ Zip _____

Tel () _____

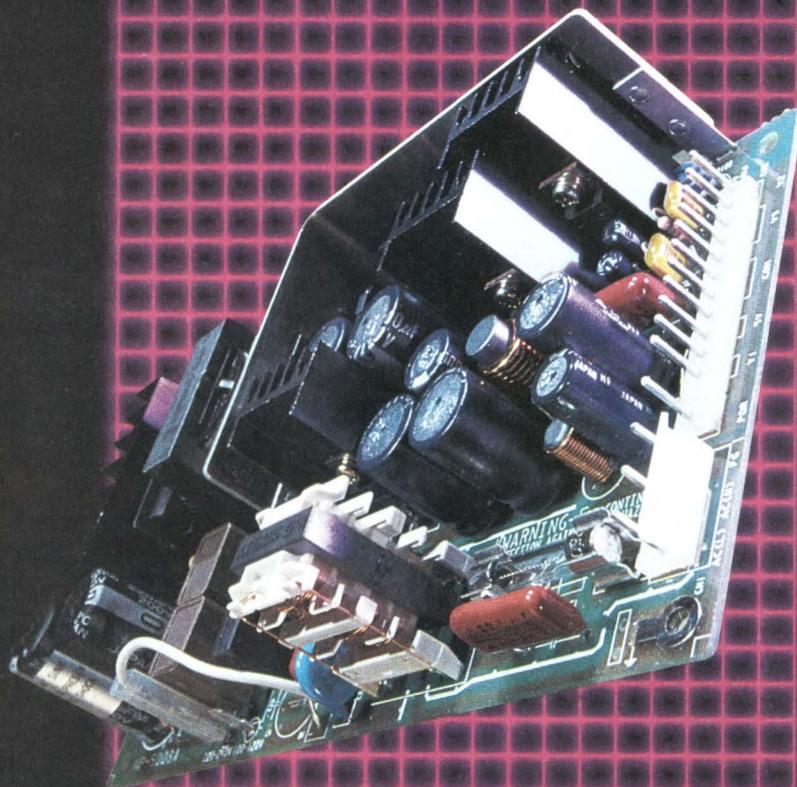
**NEC
AND
ME**



NEC Information Systems, Inc.
1414 Massachusetts Avenue
Department 1610
Boxborough, MA 01719

CD-3

Panasonic MP Series Switching Power Supplies



Panasonic® has a new line of switching power supplies that do a lot, give you a lot, and don't make a lot of noise about it.

We couldn't say that if we didn't back it up with an impressive list of features.

Including:

- Low EMI
- FCC Class B noise (VDE Class B, optional)
- UL and TUV Rhineland recognized (CSA pending)
- Power Valid Signal
- 2 input ranges; up to 4 outputs

Best of all, our MP series switching power supplies are silently waiting to be shipped... right now.

Whatever your product – terminals, floppy disk drives or other computer devices, switch to Panasonic power supplies. You won't be switching again.

For complete specifications, call or write: Panasonic Industrial Company, Computer Components Division, One Panasonic Way, Secaucus, NJ 07094. (201) 392-4264.

Limited 5 year warranty*

* Carry-in or mail-in service.

Purchaser credited for defective products

Panasonic Industrial Company

CIRCLE 57

The strong, silent types

Panasonic

DISK SYSTEM ICs COMBAT DATA ERRORS

Disk drive controller ICs that implement advanced and flexible error correction techniques require little development work by the designer, yet provide powerful methods of coping with hard errors in disk drives.

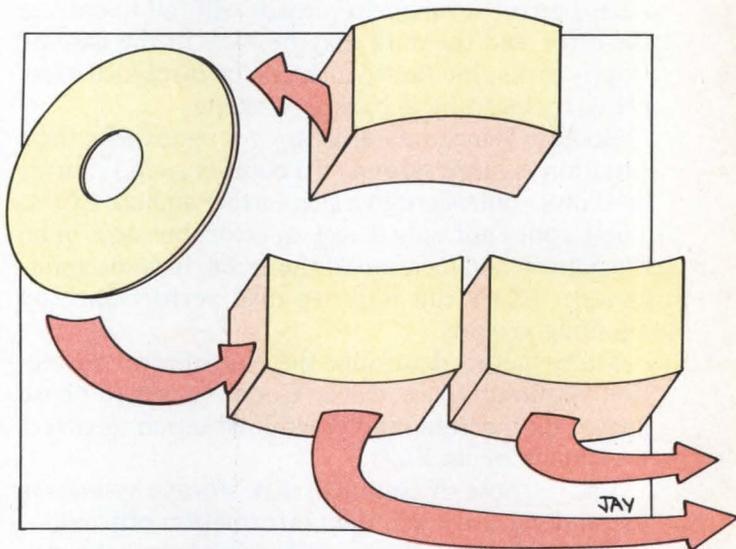
by Herb Schneider

While storage capacity is first on the list of factors determining the cost/performance figure for a disk drive, the overall error rate of the drive is a close second. These two factors are closely related, for along with an increase in data density comes an increased risk in data transfer errors. To reduce error rates, manufacturers of systems and disk drive controllers have turned to error checking and correcting codes.

A disk drive controller that normally handles such tasks as formatting and serializing/deserializing data can also spot errors coming off the disk subsystem (the mechanical drive and drive electronics). When an error is detected, the controller can try to recover the data using simple error recovery techniques (rereading) or attempt to correct the data using error checking and correcting circuitry or software.

The controller and its error correction capability determine how much the host is involved when the controller finds an error. Some systems will crash as a result of an error; others will correct the error automatically or with manual aid. If backup data were on tape or floppy, for example, the information could be retrieved. While this might be acceptable in some cases, error correction codes (ECCs) allow the controller to recover from most errors without host intervention. Until recently, system designers concerned with data integrity had to develop the hardware needed to handle errors. Now,

Herb Schneider is an applications engineer at National Semiconductor (Santa Clara, Calif). He holds a BS in electrical engineering from the University of California at Davis.



more sophisticated controller ICs that can implement ECCs, such as the National Semiconductor DP8466 family, have eased this burden.

Failures in ECC systems

Error detection and correction codes (EDCs) are not perfect, however. The EDC may fail to detect an error or may miscorrect an error. In the first instance, the probability that an error will reach the host depends on two factors: the error rate from the disk subsystem (or P_e , the probability of an error) and the number of errors not detected by the controller (or P_{udt} , the probability of an undetected error). The total probability that an error reaches the host, P_{fail} , is simply P_e multiplied by P_{udt} . In the second instance, the miscorrection probability (P_{mc}) for the controller subsystem is P_e multiplied by the probability of miscorrection (P_{mcc}). Both of these cases are dangerous since the host system is being informed that it has received good data.

Today's drives exhibit error rates on the order of 10^{-10} . ECCs (32 bits) have nondetection probabilities on the order of 10^{-10} . This makes the probability of not detecting an error extremely low—about 10^{-20} . In correction mode, the miscorrection probability is on the order of 10^{-5} , making the probability of an error reaching the host about 10^{-15} . By including error recovery techniques, most systems can recover data without attempting correction.

CRC versus ECC

Floppy disk drive systems use error checking codes to detect errors. Among these, cyclic redundancy codes (CRCs) are typical. Since correction is not possible with this approach, the drive rereads the sector to recover from the error. Assuming that the error is a soft error, a subsequent reread should be successful in recovering the data. However, if the error is a hard error, subsequent rereads will fail to correct the error and the data may be lost. In the case of floppy disks, the faulty disk can be discarded after being replaced by a backup diskette.

Because Winchester disks are not removable, their situation is more serious. To combat such failures, hard disk controllers go a step farther and use ECCs. These codes not only detect an error, but determine the pattern and location of the error. In some applications, ECCs can improve disk performance by avoiding rereads.

Many factors determine the type of error correction required for a drive. Codes, density, phase locked loop design, and precompensation all affect the choice of an ECC.

The purpose of coding in disk storage systems is to combine clock and data information efficiently. The clock information must be encoded on the disk to synchronize the clock with the flux transitions when reading the disk. This requires a data separator circuit (such as the DP8460) and a pulse detector (such as the DP8464). The data separator uses a PLL to synchronize flux transitions with the clock information and allow decoding of the information.

Disk manufacturers use various coding techniques to minimize the number of flux transitions required to represent data on a disk. This is accomplished by deleting unnecessary clock information. By minimizing the number of flux transitions required, the disk data density can be increased without increasing the actual recording density on the medium itself. These codes are collectively known as run-length limited codes. Widely used RLL codes include modified frequency modulated (MFM), and the 2,7 and 1,7 RLL codes.

An MFM bit cell consists of two windows per data bit—a clock cell and a data cell. If a 1 is recorded, the data cell contains a pulse. If a 0 follows a 1 in

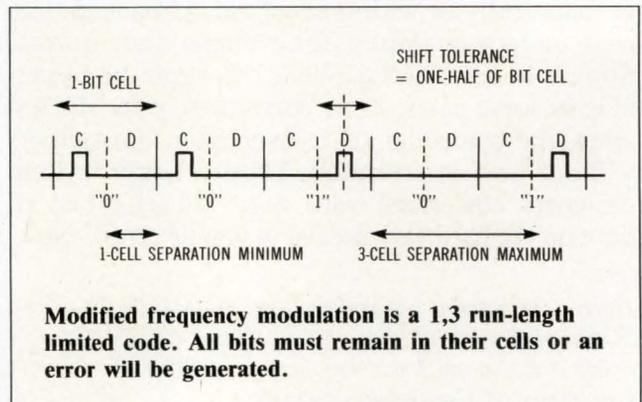
the clock bit, the clock bit in the next bit cell will be deleted. MFM is known as a 1,3 RLL code because it guarantees the minimum distance between pulses to be one cell and the maximum distance between pulses to be three cells. An error will occur if a data pulse enters a clock window or vice versa. This shifting may occur as a result of noise or intersymbol interference (the tendency of crowded bits to interact). If there is a single-bit error, the decoder recovers after 1 bit. This means one physical error results in a single logical error in the decoder (ie, the coding does not propagate the error).

Newer codes such as 2,7 (minimum distance between pulses is 2, maximum distance is 7) are pushing densities even higher. This is accomplished by converting the data sequence into a code sequence that statistically reduces the number of flux changes required to record a data sequence on the media. Thus, a particular data sequence is represented by a code sequence on the drive. If a single-bit error occurs on the drive, the decoder will interpret it as a code sequence that is different from the one originally recorded. And, it will be unable to convert the code sequence back to the correct data sequence. As a result, the length of the physical error will be propagated by a logical error in the decoder. The length of this error is dependent on the number of bits required for the decoder to recover.

With typical 2,7 codes, the time required to recover from a single-bit error can be as high as 4 bits, (5 bits if the original bit in error is included). To determine the correction capabilities of the ECC, system designers must add this length to the maximum burst length that is to be corrected. A typical 32-bit ECC will not suffice for 2,7. By increasing the ECC to 48 bits, however, larger burst errors—typical of those found in 2,7—can be handled.

Density and ECCs

Bit density affects the number of bits involved in a read error that is caused by noise events. It also tends to degrade the signal-to-noise ratio of the



channel. A noise event of the same duration for a 10-Mbit/s drive will involve twice the number of bits as a 5-Mbit/s drive. Doubling density thus doubles the burst length for errors generated by noise effects. As drive manufacturers push bit rates to 24 Mbits/s, burst error lengths will increase correspondingly.

Many drive errors are attributed to such noise in the read channel as pulse detector offsets, peak shifts caused by bit crowding, or channel asymmetry. The effect of all these is to make the pulse appear displaced from the position in which it was originally recorded.

A PLL's ability to tolerate bit shifts is a factor in overall error rates on drives. Bit shift tolerance in a PLL depends on two parameters: bandwidth and window error. Bandwidth affects the time it takes the loop to lock (acquisition time); a higher bandwidth yields a faster acquisition time, but also increases the PLL's tendency to track individual bit shifts (bit shift following). As a result, a trade-off must be made between acquisition time and bit shift following. As a general rule, the lower the bandwidth, the better the immunity to shifted bits.

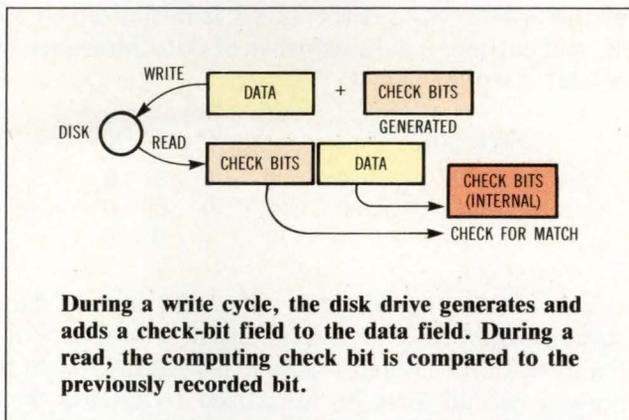
Using delay lines

A common design technique used in PLLs is to delay the data, allowing alignment of the voltage control oscillator edge in the center of a bit cell. The accuracy of this delay determines how far a bit may be shifted before a decoding error occurs. Typical designs employ one shots, which have temperature and voltage variations that limit their performance. Other designs use delay lines that tend to be quite expensive. Integrated PLLs, such as those used in the DP8460, are more tolerant of bit shifts. The DP8460 features an accurate silicon delay line that improves the bit shift tolerance of disk systems. Decoding badly shifted bits greatly reduces the overall error rate of the disk drive, possibly eliminating many errors that would otherwise occur.

A more serious problem with PLLs occurs when they lose sync with the data, and lock on harmonics. As a result, the remaining data in the sector will be read with errors. It is important to avoid these types of read errors because most codes are unable to detect them.

As the head steps toward the inner tracks, bit density increases and bits begin interacting. To compensate for this, bits can be shifted to counteract their interaction on the media. These techniques, known as precompensation, improve the error margin on the PLL. Precompensation, however, must be chosen carefully because a poorly designed scheme can degrade rather than improve performance.

Error detection on serial bit streams has been available for some time. Two commonly used codes



for communications are the CRC-CCITT (used in synchronous data link control and high level data link control) and the Autodin II polynomial (used in Ethernet). In communication applications, the basic idea behind error detection codes is to allow a transmitting station to generate a check field that provides information for a receiving station. This information enables the receiver to detect errors in the transmission.

In a disk drive, the transmitter is the write electronics and the receiver is the read electronics. Writing to the disk involves generating and then appending a check-bit field to the data or ID field. When the sector is subsequently read, the read channel reads the data and computes a check-bit field. This check-bit field is compared to the check field read from disk. If the two fields match, the sector has been read with no errors. If the two fields do not match, the sector has been read with one or more errors.

When writing to disk, check bits are generated by dividing the data stream by a polynomial. The check bit field is the remainder of this division, and is complemented and shifted out following the data—most significant bit first. When reading from disk, the remainder term will again be generated and will cancel the received remainder term if they match. Mathematically this can be expressed as:

$$\frac{X^m \cdot D(x)}{G(x)} = Q(x) + R(x)/G(x)$$

where $D(x)$ is data, $P(x)$ is the polynomial, $Q(x)$ is the quotient, $R(x)$ is the remainder (one's complement final state of shift register), and X^m is the pre-multiplication factor. When reading a sector, both $D(x)$ and $R(x)$ are processed. If the received $R(x)$ matches the current state of the shift register, the polynomial will divide evenly into $D(x) + R(x)$.

Polynomial division is implemented with a linear feedback shift register (LFSR). As an example, the polynomial $X^4 + X^1 + X^0$ can be implemented with an XOR tap placed at the input to each term

of the polynomial. If the LFSR is initialized to all 0s, and encodes a data sequence of 001, the sequence of shifts would be:

Shift No.	Data in	X ⁰	X ¹	X ²	X ³
1	0	0	0	0	0
2	0	0	0	0	0
3	1	1	1	0	0

The check bit sequence appended to this data stream would be 0011 and the sector would be written as 0010011. When reading this sector, the shift register would first be initialized to all 0s, then shifted with the data stream and check bits, resulting in the following shift sequence:

Shift No.	Data in	X ⁰	X ¹	X ²	X ³
1	0	0	0	0	0
2	0	0	0	0	0
3	1	1	1	0	0
4	0	0	1	1	0
5	0	0	0	1	1
6	1	0	0	0	1
7	1	0	0	0	0

Since there has been no error in reading the data, the check-bit field will return the shift register back to its initialized state of all 0s. After the third shift, bits X⁰ to X⁴ contain the same sequence as the next 4 bits to be read from the disk. Thus, the shift register acts as a serial comparator and cancellation of the two patterns results in all 0s (ie, D(x) + R(x) divided evenly by polynomial).

If a read error occurs, the final state of the shift register before the check-bit field arrives will not match the check-bit field sequence and cancellation will not occur. For example, if the 001 data stream is corrupted and 000 is read, the following sequence will occur:

Shift No.	Data in	X ⁰	X ¹	X ²	X ³
1	0	0	0	0	0
2	0	0	0	0	0
3	0	0	0	0	0
4	0	0	0	0	0 < Beginning
5	0	0	0	0	0 of check-
6	1	1	1	0	0 bit field
7	1	1	0	1	0 < Last state

The last state of the shift register after reading the check-bit field is not the same as the initialized value. The last state of the shift register, called the syndrome, can be used to find the location of the error by using some unique properties of ECC polynomials. ECC polynomials have a linear sequence of shifts that can be counted and decoded to determine the location and pattern of a read error. The polynomial $P(x) = X^4 + X^1 + X^0$ has two possible sequences.

If the LFSR is preset to all 0s, the LFSR will not leave the all 0s state. If the LFSR is preset to 0001, a repeating 16-symbol sequence will be generated.

The number of symbols a code can generate is referred to as the natural period of the polynomial. The natural period of the code determines how many data symbols may be encoded by the polynomial. Error correction involves determining the distance in the sequence between the syndrome and some reference state. In the polynomial above, a read error was simulated with a resultant syndrome of 1010. When the shift register decodes state S1 (0001), the number of shifts required to reach S1 from the syndrome indicates where the error is located. (In this polynomial, the error location is calculated as $16 - \text{NSHIFTS} - 3$.) In the case of the 1010 syndrome, 10 shifts are required to reach state S1, indicating that the error is in the third bit of the data stream.

Forward shifting and alternatives

This method of error correction is called forward shifting. Its disadvantage is that it requires many shifts to locate the error. This is related to the natural period of the code. As a result, a 32-bit code could require 43,000 shifts to locate the error at the beginning of a 512-byte sector. This is intolerable in many disk drive systems. A second disadvantage of forward shifting for correction occurs if an error is noncorrectable. This is because the correction circuitry or software must step through the entire period of the polynomial to determine noncorrectability.

A more efficient way to find the error is to step through the polynomial sequence in reverse, using a reciprocal polynomial that generates a sequence in reverse order to the forward polynomial. In the example, the reciprocal polynomial would be $P_r(x) = X^4 + X^3 + X^0$. An LFSR shifted in the opposite direction to the forward polynomial will generate the reverse sequence of the forward polynomial. Shifting the syndrome using $P_r(x)$ will reduce then number of shifts required to locate the error. It will also guarantee finding noncorrectable errors in no more than $S + A$ shifts where S is the number of data bits in the sector and A is the number of check bits.

The time required to correct an error can be a critical ECC performance parameter. In the example, only five shifts would be required to locate the error using the reciprocal polynomial and reverse shifting. These techniques can be extended to include multiple bit correction. While the generator in the example was preset to all 0s, it is better to preset to all 1s to avoid sync slippage errors.

Specifications for measuring the performance of ECCs include detection spans, correction spans, mis-detection probability, and miscorrection probability.

The detection span of any polynomial is the length of the largest error burst guaranteed to be detected. The detection span is fixed by the length of the LFSR. The correction span refers to the number of contiguous bits that are guaranteed to be corrected. (A typical 32-bit ECC has a correction span of approximately 6 bits.)

At first glance, it would seem that selecting a high correction span for a particular polynomial would result in higher system integrity. Unfortunately, the opposite is true. All ECCs exhibit some miscorrection probability for the set of all possible errors. Miscorrection results when the span of the error exceeds the correction span and the ECC interprets the error as correctable. The ECC will then add to the problem by miscorrecting another region in the sector and indicating that the error is corrected—a dangerous situation in most systems.

The relationship between miscorrection probability, the number of check bits, and the number of data bits can be expressed as:

$$\text{Miscorrection} = \frac{S \cdot 2^{(C-1)}}{2^A}$$

where S is the number of data bits, C is the correction span, and A is the number of check bits. Increasing the correction span by a single bit doubles the miscorrection probability. The only way to compensate miscorrection is to reduce the number of data bits encoded by the polynomial, decrease the correcting span, or increase the number of check bits.

Certain ECC polynomials, known as Fire codes, contain pattern sensitivities that radically increase miscorrection probabilities when encountering certain types of errors. Fire codes that were used in

Error correction in the disk drive controller

Early system implementation of error checking and correcting codes required the actual correction cycle to be performed in software. This approach had the disadvantage of requiring the system designer to develop and test software for correction. Software implementations also tended to be much slower than hardware implementations. National Semiconductor's DP8466 disk data controller includes circuitry that has been designed to locate and correct errors at maximum speed. The logic generates the reciprocal polynomial and reverse shifts to reduce the number of shifts for correction. Generation of the reciprocal polynomial also removes the programming tasks that the CPU must perform to begin a correction cycle.

The DP8466 does not require that the data be passed through the ECC during the correction cycle as did some early ECC designs. As a result, the DP8466 can locate and generate a correction pattern in less than the time to read a sector. If an interleave of one sector is performed, the error can be corrected before the next sector is read. At the end of a correction cycle, the disk data controller provides a byte-aligned correction pattern and a displacement pointer to the first byte in error in the sector. By performing a simple XOR, the CPU can correct the data pattern. If the error is not correctable, ECC circuitry indicates correction has failed.

Of prime concern to a system designer is the miscorrection probability. The miscorrection probability is directly proportional to the number of bits to be corrected and is inversely proportional to the number of check bits. A programmable correction span allows the programmer to select a span no larger than the guaranteed correction span for that code. A 5-bit register on the disk data controller is provided to select the correction span.

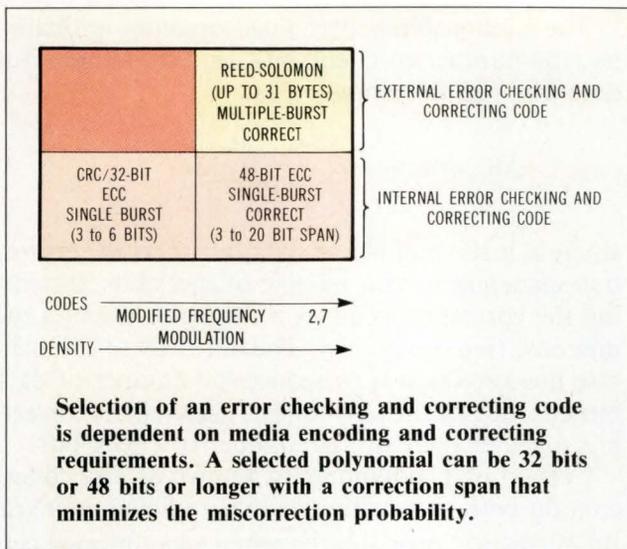
The DP8466 also features flexible assignment of either cyclic redundancy check, internal ECC, or ex-

ternal ECC to the ID field or the data field. In a typical application, a 16-bit CRC may be used for error detection on the ID field and a 32-bit ECC might be assigned to the data field. It is also possible to assign the 32-bit ECC to both the ID and data field.

There is no single polynomial that manufacturers have agreed upon. But, to ensure flexibility and compatibility with existing designs (especially crucial with removable media), the DP8466 disk data controller has fully programmable polynomial taps. A series of six registers selects the desired polynomial for either 32- or 48-bit operation. The taps reconfigure the linear feedback shift register (LFSR) for any 32- or 48-bit code. A single preprogrammed CRC is implemented—this CRC is the CRC-CCITT most commonly used in disk drives.

There are two types of ECC circuits implemented in existing drives: LFSRs implemented with internal XORs and LFSRs implemented with external XORs. External XOR implementations use parity generators and produce a check-bit sequence different from an equivalently programmed internal XOR ECC. By programming the preset pattern of an internal XOR implementation, a check-bit sequence that is the same as an internal XOR can be generated. The six preset registers on the DP8466 allow full hardware compatibility with existing ECC designs.

Encapsulation provides an interesting application of both internally generated ECC and externally generated ECC. A typical use of two ECC fields is to encapsulate the data field and an ECC field with a second ECC field. When the first ECC detects a read error and correction is performed, the second ECC can determine whether the corrected data is correct. This type of scheme provides an additional level of protection when using error correction. Any miscorrection of the data field by the second ECC will be detected by the first ECC before erroneous data is passed to the host.



earlier controllers and many VLSI controllers exhibit pattern sensitivities to multiple burst errors. This results in miscorrection probabilities as high as 10 percent. To overcome this system defect, a second generation of codes, developed by a computer search and known as computer-generated codes, have been chosen to provide more reliable correction.

Computer-generated codes meet certain detection and correction spans. They are then tested against a set of requirements to avoid pattern sensitivities. The Fire- and computer-generated codes are known as single-burst ECCs. (These codes detect multiple-burst errors but cannot correct them.) A burst or multiple burst of error with extremities exceeding the correction cannot be corrected—even though the system is informed of the error.

Error recovery techniques

A third class of codes can correct multiple-burst errors or a single long-burst error. These codes are block interleaved Reed-Solomon codes.

When an error is encountered, it will be detected at the end of the sector. The controller must determine whether to attempt to reread the sector or to attempt to correct it. It is quite common to try to reread the sector in order to avoid correction of soft errors. Rereads reduce the number of errors that the ECC must handle, and decreases the overall miscorrection probability since an attempt at correction may result in a miscorrection.

Several recovery methods are used when rereading a sector. If the head is not properly centered over the track, the controller will make small adjustments to step the head over the track. The ST412HP interface standard includes a new step line to allow adjustment of the head correctly over the track. Each incremental step is followed by a reread to determine if the error disappears. Another technique com-

monly used in storage module devices is the strobe early/strobe late used in system measuring devices to shift the PLL's decode window to read in the presence of phasing errors. After these reread techniques have been exhausted and the syndrome after each reread appears to be consistent, the error is considered a hard error and correction is initiated by the controller.

When a hard error is encountered repeatedly in the same sector, the sector containing the error will be mapped out of the available storage space by the controller. This inhibits access to the sector in future write operations. To avoid writing faulty sectors, one of the manufacturing steps that takes place before shipping disk drives is to identify and map those sectors containing defects that would result in hard errors. The controller will avoid writing these sectors. Sectors containing defects are often identified in a defect map on sector 0. As a result of these mapping strategies, the controller is usually exposed only to soft errors that are recoverable with a reread.

During rereads, a performance penalty occurs since a latency of one revolution is imposed. Although this may be acceptable in single-user applications, the latency may not be acceptable for a multi-user access disk drive or a disk drive that has many read errors. Many systems attempt correction on a sector immediately after an error is detected. Invoking correction on any error is possible, but many techniques require correction in software. This means a reread may be a faster method of recovery than actual correction.

To select an ECC, the designer must examine the types of expected errors (single-burst, double-burst), type of media encoding, and data rate (large-bursts, signal-to-noise degraded). In addition, media defects, available error recovery techniques, PLL margins, and write channel and read channel signal-to-noise margins must be considered. Based on these issues, the user can select either a 32-bit, 48-bit, or longer polynomial and a correction span to minimize the miscorrection probability.

Selection of an ECC is also dependent on media encoding and correction requirements. Current MFM encoded drives at 5 Mbits/s require about 32 bits of ECC with a 5- to 6-bit correction span. The 2,7 RLL drives require a greater correction span and typically will use a 48-bit ECC with an 11-bit correction span.

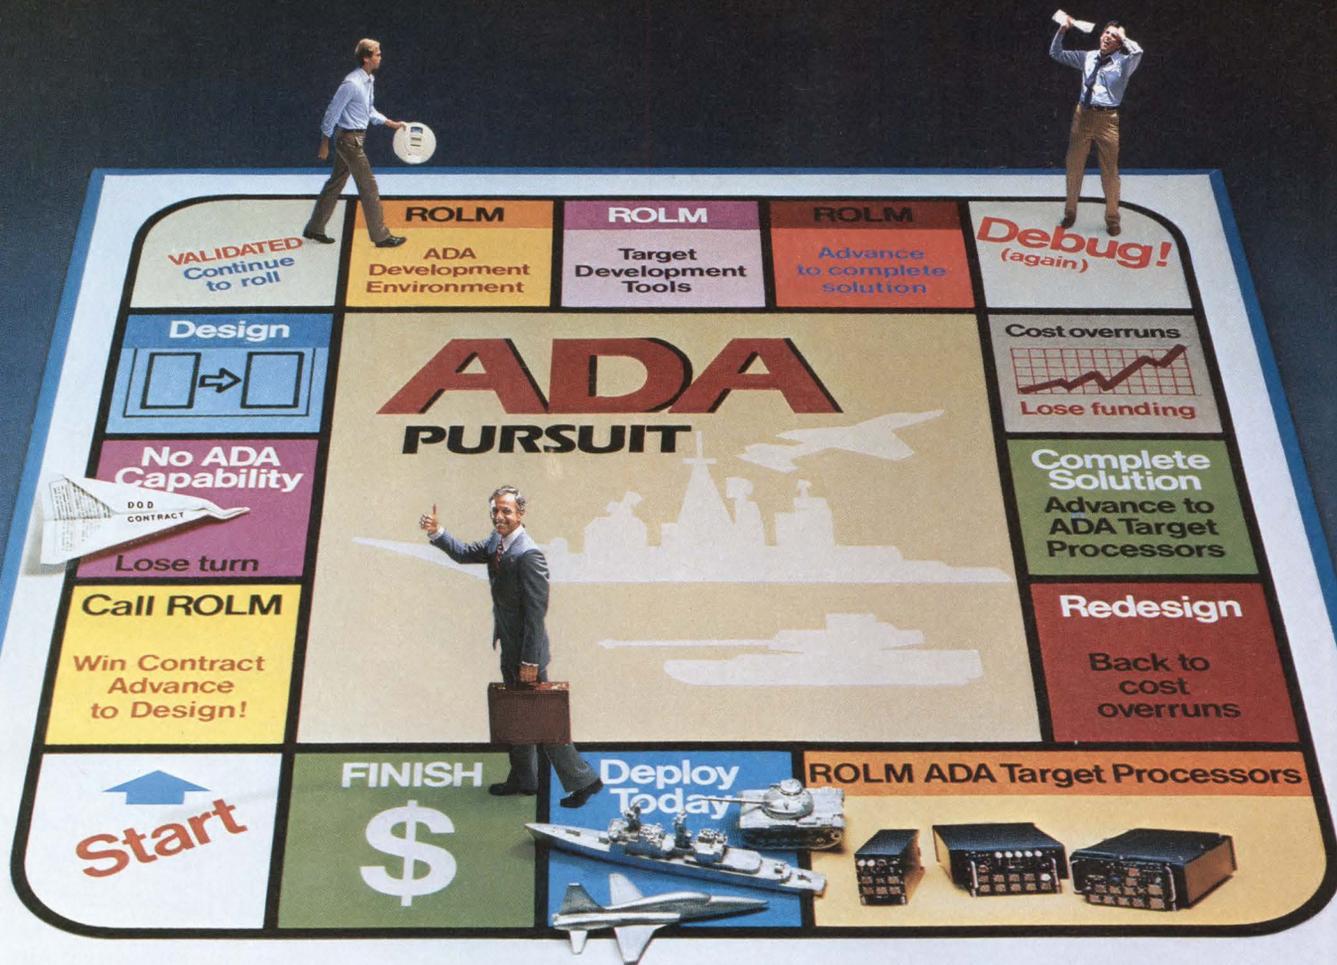
Please rate the value of this article to you by circling the appropriate number in the "Editorial Score Box" on the Inquiry Card.

High 716

Average 717

Low 718

Deployment is the name of the game in Ada Pursuit



While others struggle to get their Ada® compilers past START, ROLM® and Data General have been delivering their validated, production-quality Ada compiler to customers all over the world.

With ROLM and Data General, you will quickly move through the design and development steps because we deliver a validated Ada compiler and more. . . . Our compiler is hosted on the powerful Data General MV series computers and comes with an extensive Ada Development Environment (ADE™) including a selection of tools and utilities such as an

Ada Source Level Debugger, Configuration Manager, Ada Pretty Printer, and Ada Library Manager, which will enable you to develop your Ada code quickly and efficiently.

But to win the game you need to deploy your software on a target processor. ROLM understands. . . . We fulfill this requirement with field proven 16- and 32-bit Mil-Spec ROLM processors, real-time deployable operating systems, and Ada target code generators.

From development to deployment—the only complete solution. ROLM and Data General.

ROLM

MIL-SPEC
Computers

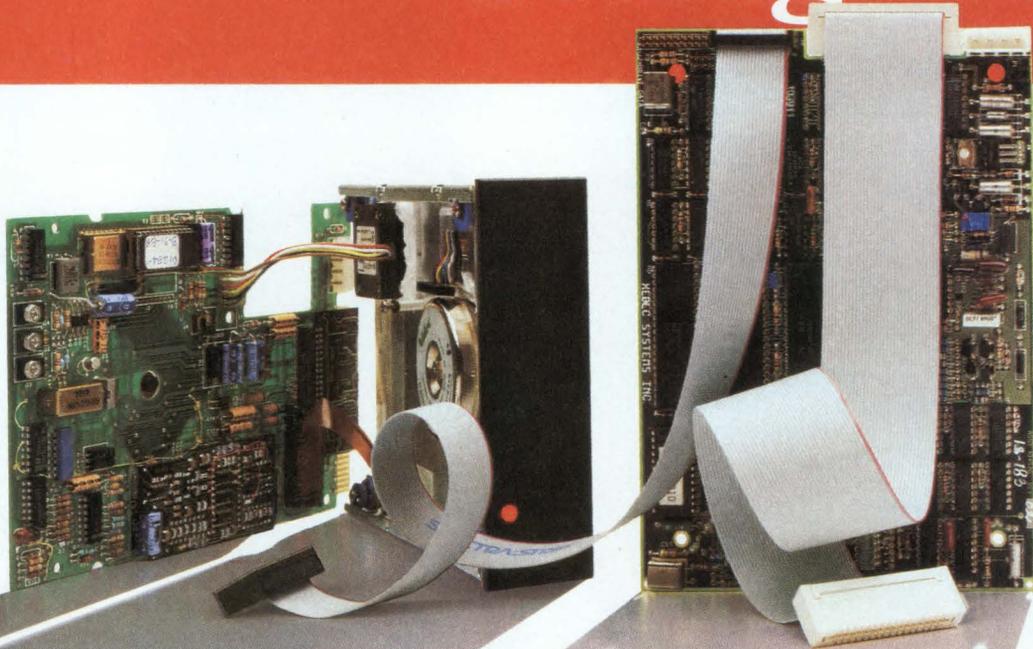
One River Oaks Place, M/S 110, San Jose, CA 95134 (408) 942-7875
ROLM GmbH—Am Kronberger Hang 3, D-6231 Schwalbach/TS, West Germany 6196-81039

®Ada is a registered trademark of the U.S. Department of Defense (OUSDRE-AJPO).
™ADE is a trademark of ROLM Corporation

CIRCLE 58

Xebec's New Owl Reduced Storage To

Then.



Microcomputer storage history has progressed by a series of small "next logical steps." A replacement of a component here, a refinement of technology there. But now Xebec has taken a giant step, with its Owl intelligent disk file.

On the surface, the Owl might look like other 10-megabyte, 5¼" half-high Winchester. Underneath, however, it's an example of superior technology. The integration of controller logic and drive electronics on a single board means not just one less PCB, but the elimination of expensive connectors and cabling, low power consumption (15 watts typical) and enhanced data integrity.

We've put data separation in the HDA for precise control of data windowing and the elimination of background noise. We've provided a diagnostics channel to the host that delivers meaningful error messages.

Just as our superiority in minicomputer controllers led the way to a similar superiority in micro controllers, and our tested pairs solutions evolved from our considerable sub-system and testing experience, so too the Owl

reflects our "top-down" engineering strategy—creating both technological and cost-of-owner-

Single-board drive and controller electronics insure higher reliability

Integrated data separation within disk drive delivers higher data integrity

Ruggedized design includes head landing zone

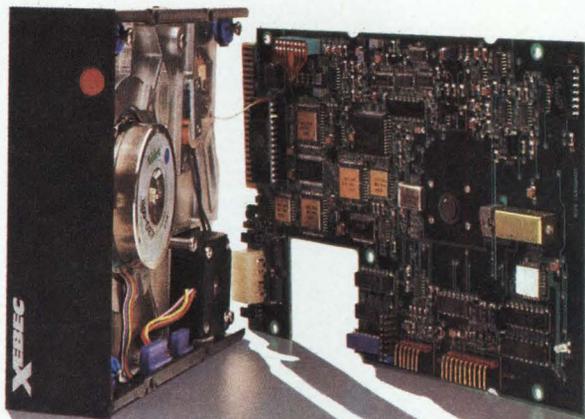
Comprehensive diagnostics eliminate "What's wrong?" guesswork

Compatible with industry-standard Xebec SASI bus for faster system integration

CORPORATE HEADQUARTERS
2055 Gateway Place, Suite 600
San Jose, California 95110 U.S.A.
(408) 287-2700

es Microcomputer Disk Two Words.

Now.



ship breakthroughs by designing sophisticated high performance, multi-user features into small-system, single-user environments.

The System-Engineered Solution.™

As a major user of drives, and a major testing resource for other drive manufacturers, we are concerned about drive reliability. And we've learned that it's not just a question of improving a component or consolidating some board real estate. The answer lies in careful consideration of overall system integration requirements.

The Owl epitomizes what we call the Xebec System-Engineered Solution. The focus is two-fold: on today and on tomorrow. Compatible now with industry-standard Xebec SASI, the Owl—by eliminating the ST506 interface—is

upwardly compatible for future higher densities, capacities and performance.

In broader perspective, Xebec's approach to OEM satisfaction rests on our proven experience, our vertical integration strengths—which now include production of heads and plated media—and our commitment to zero defect quality, by way of computer-aided design and robotics manufacturing.

Call Xebec today. Let us tell you more about the Owl. And how we can deliver the difference between then and now. Now.

Xerox D.

SALES, INTERNATIONAL

Belgium
32-02-762-9494
England
44-734-693511
Italy
39-6-350201

SALES, U.S.A.

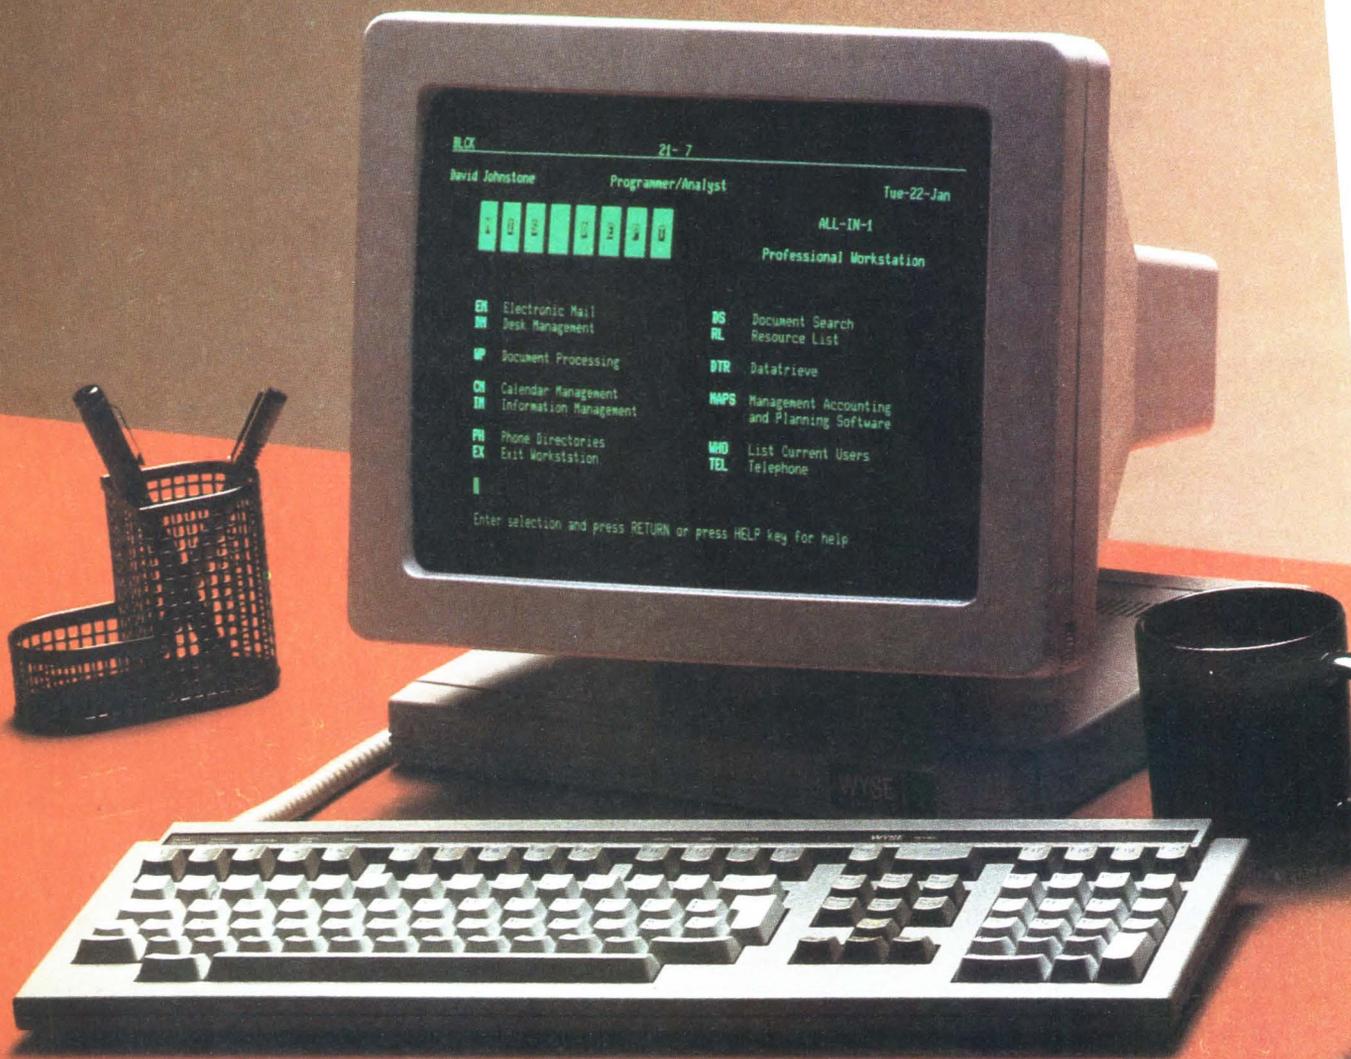
Sunnyvale, CA (408) 732-7330
Irvine, CA (714) 851-1437
Atlanta, GA (404) 457-9872
Boston, MA (617) 740-1707
Dallas, TX (214) 361-0687
Baltimore, MD (301) 992-7377
Federal (301) 621-3010
Chicago, IL (312) 931-1420

U.S.A. AND CANADA DISTRIBUTORS

Kierulff
Hamilton-Avnet
Avnet Electronics
Hamilton Electro Sales

XEBEC
*The Quality
Investment*

Now for \$799, a VT-220 with perfect function,



compatible terminal fit and form.

Our WY-85 gives you complete DEC VT-220™ software compatibility in function and fit, plus a form that beats all, for a price that beats all.

Like all our terminals, the WY-85 packs all the ergonomic features you want into an unusually small footprint. The generous 14" non-glare screen tilts, swivels, and handles a full 132-column format. Even the sculpted, low-profile keyboard adjusts for perfect fit and easier function.

So you get not just emulation, but the full operational compatibility you're looking for, including keyboard layout. For a lot less money, in a lot less space, with a lot more style.

No wonder Wyse terminals are now on more than 150,000 desktops worldwide.

For more information about our full line of computer display products, call the regional office nearest you.

WYSE

Regional Offices:

Northeast (201) 725-5054

Southeast (305) 862-2221

Midwest (313) 471-1565

Southwest (818) 340-2013

Northwest (408) 559-5911

TOSHIBA. THE POWER IN



AREA SALES OFFICES: WESTERN AREA, Toshiba America, Inc., (714) 752-0373; CENTRAL AREA, Toshiba America, Inc., (312) 945-1500; EASTERN AREA, Toshiba America, Inc., (617) 272-4352; NORTHWESTERN AREA, Toshiba America, Inc., (408) 244-4070; POUGHKEEPSIE, NEW YORK AREA, Toshiba America, Inc., (914) 462-5710. **REPRESENTATIVE OFFICES:** ALABAMA, Glen White Associates, (205) 883-7938; ARIZONA, Semper Fi Sales Company, (602) 991-4601; ARKANSAS, Technology Sales Company, (214) 380-0200; CALIFORNIA, (Northern), Elrepc, Inc., (415) 962-0660; CALIFORNIA, (Southern), Bager Electronics, Inc., (714) 957-3367; COLORADO, Duffy Associates, (303) 595-4244; CONNECTICUT, Datcom, Inc., (203) 288-7005; DELAWARE, Vantage Sales, (609) 663-6660; FLORIDA, Donato & Associates, (305) 352-0727, (813) 785-3327; GEORGIA, Glen White Associates, (404) 441-1447; IDAHO, Components West, (206) 271-5252; ILLINOIS, L-TEC, (312) 773-2900; INDIANA, Leslie M. DeVoe Company, (317) 842-3245; IOWA, J.R. Sales Engineering, (319) 393-2232; KANSAS, D.L.E. Electronics, (316) 744-2126; KENTUCKY, Leslie M. DeVoe Company, (317) 842-3245; LOUISIANA, Technology Sales Company, (214) 380-0200; MAINE, Datcom, Inc., (617) 891-4600; MARYLAND, Glen White Associates, (301) 252-6360; MASSACHUSETTS, Datcom, Inc., (617) 891-4600; MICHIGAN, R.C. Nordstrom & Company, (313) 559-7373; MINNESOTA, Quantum Sales, Inc., (612) 884-4700; MISSISSIPPI, Glen White Associates, (205) 883-7938; MISSOURI, D.L.E. Electronics, (316) 744-2126; MONTANA, Components West, (206) 885-5880; NEVADA, Elrepc, Inc., (415) 962-0660; NEW ENGLAND, Datcom,

2 μ CMOS GATE ARRAYS.

**We have the silicon. We have LDS™ software.
We have volume production. And we have design capacity
available now.**

Get your 2 μ CMOS gate array design off the board and into production at one of Toshiba's design centers. We have design capacity available now. We have Daisy™ and other popular workstation compatibility, LDS design software and the applications engineering staff to give you total support.

We are volume producers of industry standard gate array families such as our high performance 2 μ devices, the 17G series. It has a 1.5 nanosecond gate speed with up to 10,000 gates as well as a macro library of 7400 type SSI/MSI functions. These units have fully compatible TTL/CMOS input and output and can be configured with up to 232 available pins. Our 3 μ industry standard CMOS gate array family, the 15G, has a 2.5 nanosecond gate speed with up to 6000 gates available. In fact, that's why these industry standard families are second-sourced with LSI Logic Corporation.

And we're working now on the next generation of CMOS gate arrays

as well as full standard cell capability to be introduced shortly.

TOSHIBA IS PRODUCTION POWER

Toshiba is the world's largest producer of high quality CMOS devices, and one of the few companies with high volume production experience and capability in CMOS gate arrays.

All this means Toshiba can provide the design support you need with a volume production capability to assure that your designs become delivered product in the right quantities on time.

So if you are planning a 2 μ gate array design, plan with the people who have the power. Plan with Toshiba.

Part Number	Array Size	Max Pins	Gate Length	Gate Delay (ns) ¹	
				TYP	MAX ²
TC15G008	880	74	3	2.5	5.0
TC15G014	1400	92	3	2.5	5.0
TC15G022	2200	114	3	2.5	5.0
TC15G032	3200	138	3	2.5	5.0
TC15G042	4200	156	3	2.5	5.0
TC15G060	6000	180	3	2.5	5.0
TC17G005	540	44	48	2	3.0
TC17G008	880	52	68	2	3.0
TC17G014	1400	68	86	2	3.0
TC17G022	2200	82	106	2	3.0
TC17G032	3200	98	128	2	3.0
TC17G042	4200	118	150	2	3.0
TC17G060	6000	146	188	2	3.0
TC17G080	8000	174	222	2	3.0
TC17G100	10000	180	232	2	3.0

NOTES:

¹ For a 2 input NAND gate having a fan out of 2 and 2mm of metal interconnect.

² At 70°C and Vdd = 5v \pm 5%

TOSHIBA. THE POWER IN GATE ARRAYS.

TOSHIBA AMERICA, INC.

LDS is a trademark of LSI Logic Corporation. Daisy is a trademark of Daisy Systems.

Inc., (617) 891-4600; **NEW HAMPSHIRE**, Datcom, Inc., (617) 891-4600; **NEW JERSEY**, Necco 1, (201) 461-2789, Vantage Sales, (609) 663-6660; **NEW MEXICO**, Semper Fi Sales Company, (602) 991-4601; **NEW YORK**, Necco 1, (201) 461-2789, PI-tronics, (315) 455-7346; **NORTH CAROLINA/SOUTH CAROLINA**, Glen White Associates, (919) 848-1931, (615) 477-8850; **NORTH/SOUTH DAKOTA**, Quantum Sales, Inc., (612) 884-4700, **OHIO**, Del Steffen & Associates, (216) 461-8333, (419) 884-2313, (513) 293-3145; **OKLAHOMA**, Technology Sales Company, (214) 380-0200; **OREGON**, Components West, (503) 684-1671; **PENNSYLVANIA**, Del Steffen & Associates, (412) 276-7366, Vantage Sales, (609) 663-6660; **RHODE ISLAND**, Datcom, Inc., (617) 891-4600; **TENNESSEE**, Glen White Associates, (615) 477-8850; **TEXAS**, Technology Sales Company, (512) 346-9940, (713) 266-2473, (214) 380-0200; **UTAH**, Duffy Associates, (303) 595-4244; **VERMONT**, Datcom, Inc., (617) 891-4600; **VIRGINIA/WEST VIRGINIA**, Glen White Associates, (804) 237-6291, (804) 295-0435, (804) 224-7764, (804) 224-0404; **WASHINGTON**, Components West, (206) 885-5880, (509) 255-6224; **WISCONSIN**, L-TEC, (414) 774-1000; **WYOMING**, Duffy Associates, (303) 595-4244; **BRITISH COLUMBIA**, Components West, (206) 885-5880; **CANADA**, Electro Source Inc., (416) 675-4490.

Matter of fact, our book will show you more banana plug and jack combinations and variations than you can shake a stick at. Regular or miniature, on patch cord and cable assembly combinations; you name it and the chances are we'll have it.

And now that you have an idea of our depth and variety you should also be aware that the banana plugs and jacks we design and manufacture are considered by electronic design engineers to be right at the top. No question.

A superior product? You know you've got

that. And you should know, also, that we make sure you receive fast service and specified quantity through your favorite electronic parts distributor.

After all, we wrote the book.

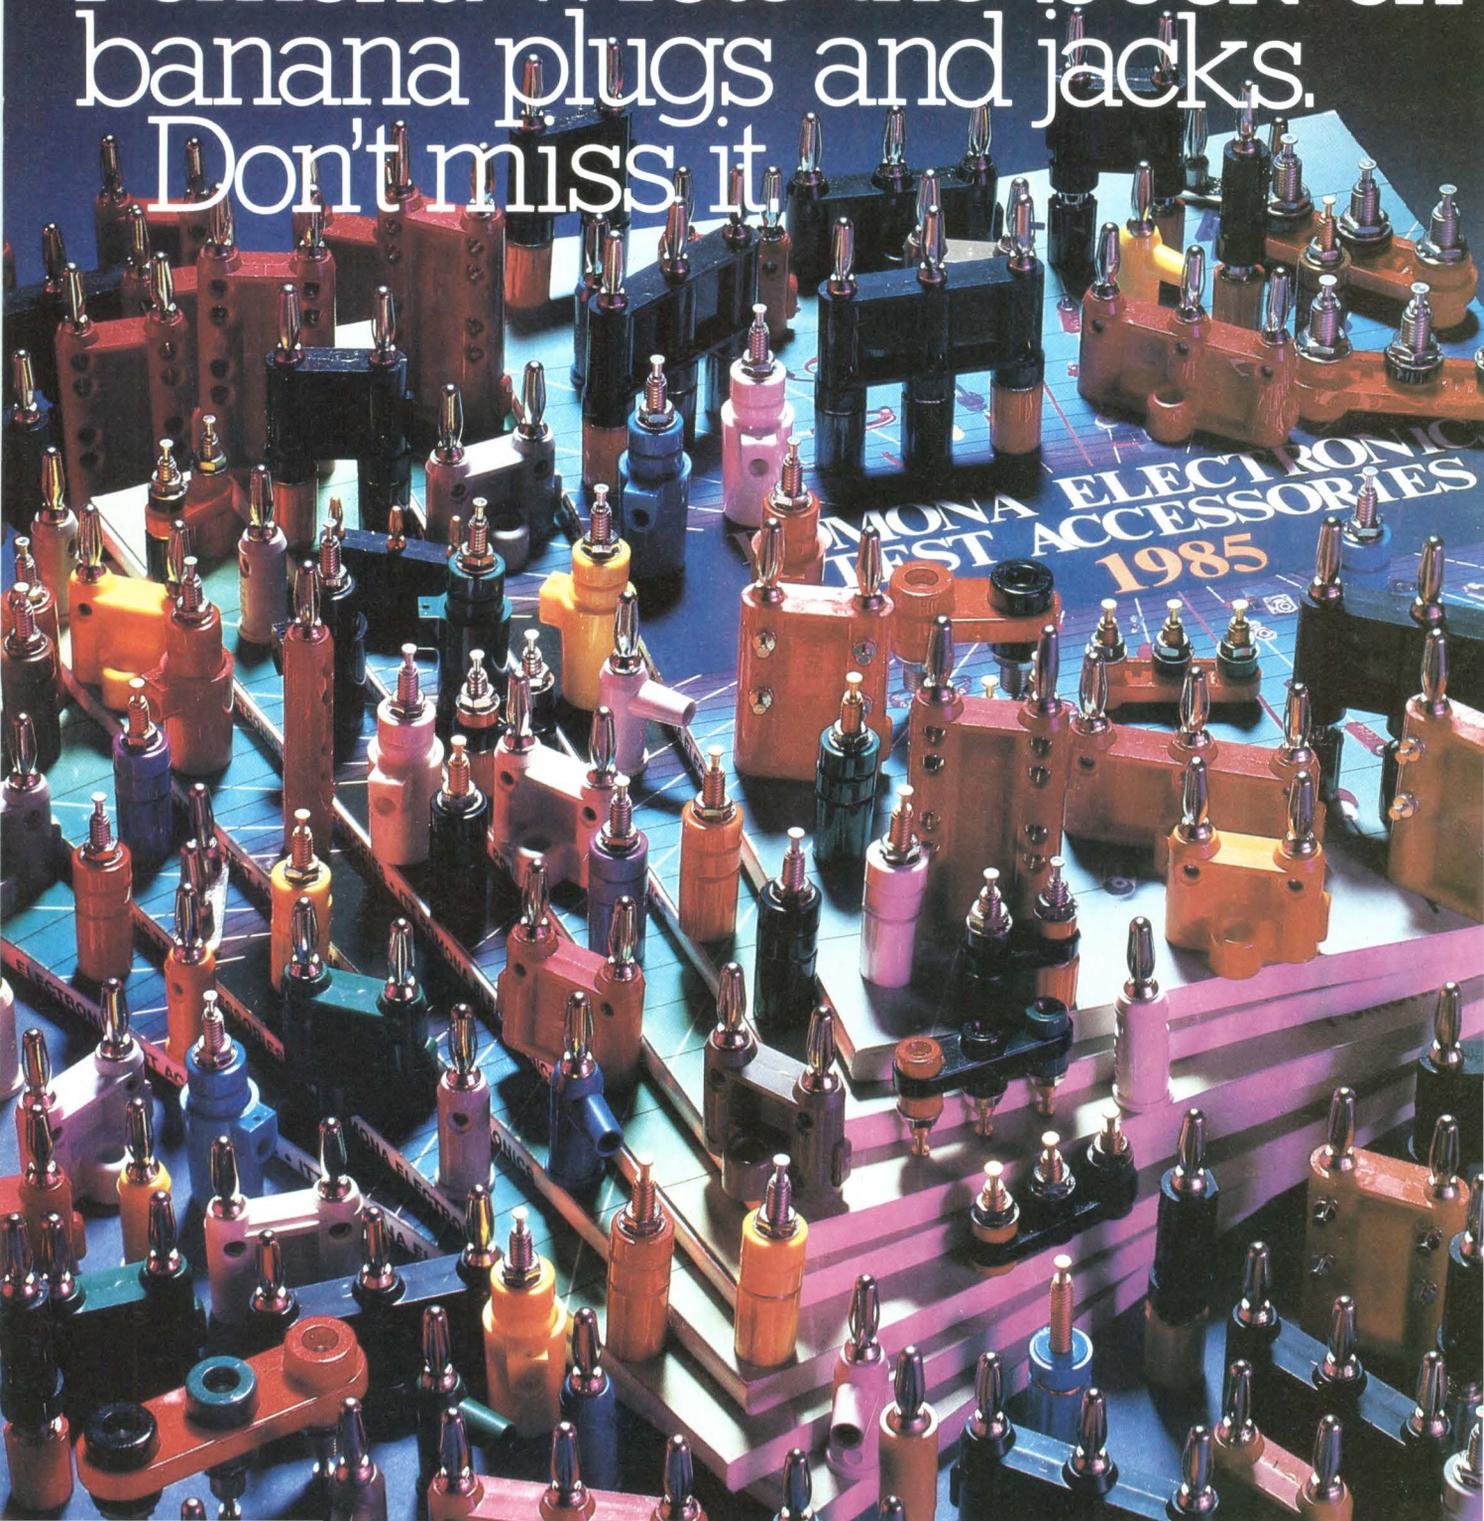
And it's free. Our products are described and illustrated in our

1985 general catalog. Just call (714) 623-3463 or 623-6751. Or, write to us at ITT Pomona Electronics, a division of ITT Corporation, 1500 E. Ninth St., Pomona, Ca 91769.

ITT

Pomona Electronics
CIRCLE 62

Pomona wrote the book on banana plugs and jacks. Don't miss it.



S-100 BUS ADAPTED TO 68000-BASED UNIX COMPUTER

Although the S-100 bus poses some challenges to system designers, it provides adequate memory support for a multi-user Unix computer and permits multiple DMA devices.

by Robert Groppo

Choosing the best bus for a computer system requires a close look at the processor and operating system that will be used. The designer must also consider the memory to be supported, the I/O capabilities, the power requirements, and timing. All these factors influenced the selection of the S-100 bus for Tricep, a multi-user computer from Morrow Designs, Inc.

Tricep runs Unix as its host operating system with MS-DOS and CP/M as subenvironments. The host processor is the 68000, while the 80188 and Z80 (still under development) application processors are available for MS-DOS and CP/M tasks. The 68000/68451 chip set (CPU and memory management unit) was chosen for its power and speed when running Unix and for the wealth of Unix applications compiled to run on it. The 80188 and Z80 chips were selected as application processors because of their large application software bases.

Even with three processors, the computer can be viewed as a "virtual processor." A proper mix of hardware and software makes it possible to use virtually any microprocessor and operating system. The various operating systems are not simulated in software. They are actually run on the targeted microprocessor, thus maintaining adequate system response.

Robert Groppo is engineering product manager for Tricep at Morrow Designs, Inc (San Leandro, Calif). He holds AA degrees in computer science and electronics technology from Shabot College.

When a user logs on and issues a command to run a program, the Unix system determines which subenvironment is required to run the program and starts up the software on the appropriate application processor. All of this is transparent to the user. There can even be multiple versions of the same software package, with Unix running the desired program on the next available application processor.

Selecting the S-100 bus

The memory requirement for a 68000-based system running Unix System V is at least 512 Kbytes, but closer to 1 or 2 Mbytes for ideal response time. Furthermore, the dual-ported 80188 and Z80 slaves require an even larger memory space. The bus has to support main memory in the 16-Mbyte range as well as multiple DMA devices. To enhance performance, the system I/O is configured with DMA.

With a plan to add devices from multiple vendors, and to avoid possible problems, the ideal bus turned out to be an IEEE standard (rather than a proprietary bus or a *de facto* standard). The two options were the IEEE 696 (S-100) bus and the IEEE 796 (Multibus). Since Morrow had experience in manufacturing S-100 boards, the S-100 bus was the obvious choice.

The S-100 bus supports a 16-Mbyte memory address space and a separate 64-Kbyte I/O space, making it well-suited for a dual-port application processor environment. Six Mbytes of main memory could be allocated to Unix and the upper 10 Mbytes reserved for other use. The bus provides room for a 1-Mbyte graphics window; a 512-Kbyte window to support up to eight, 8-bit Z80 application processors; a 512-Kbyte window to support Ethernet or

alternative LANs; and an additional 8 Mbytes to support eight 80188 MS-DOS application processors.

With its 8- and 16-bit slave protocol, the S-100 bus can mix 8-bit slaves (the 80188 and the Z80) with a 16-bit 68000. Tricep can have 16-bit program memory access to maintain the full bandwidth of the 68000, and uses only 8 bits when accessing the 8-bit application processors. The 8/16-bit S-100 protocol uses handshaking between the bus master and bus slaves. The bus master sends out an sXTRQ (sixteen request) signal at the beginning of each bus cycle. If a slave is a 16-bit device, it responds by asserting SIXTN (sixteen acknowledge). If the slave is an 8-bit device, it need not respond. The master examines the SIXTN line and performs either a full 16-bit access if asserted, or an 8-bit access when inactive.

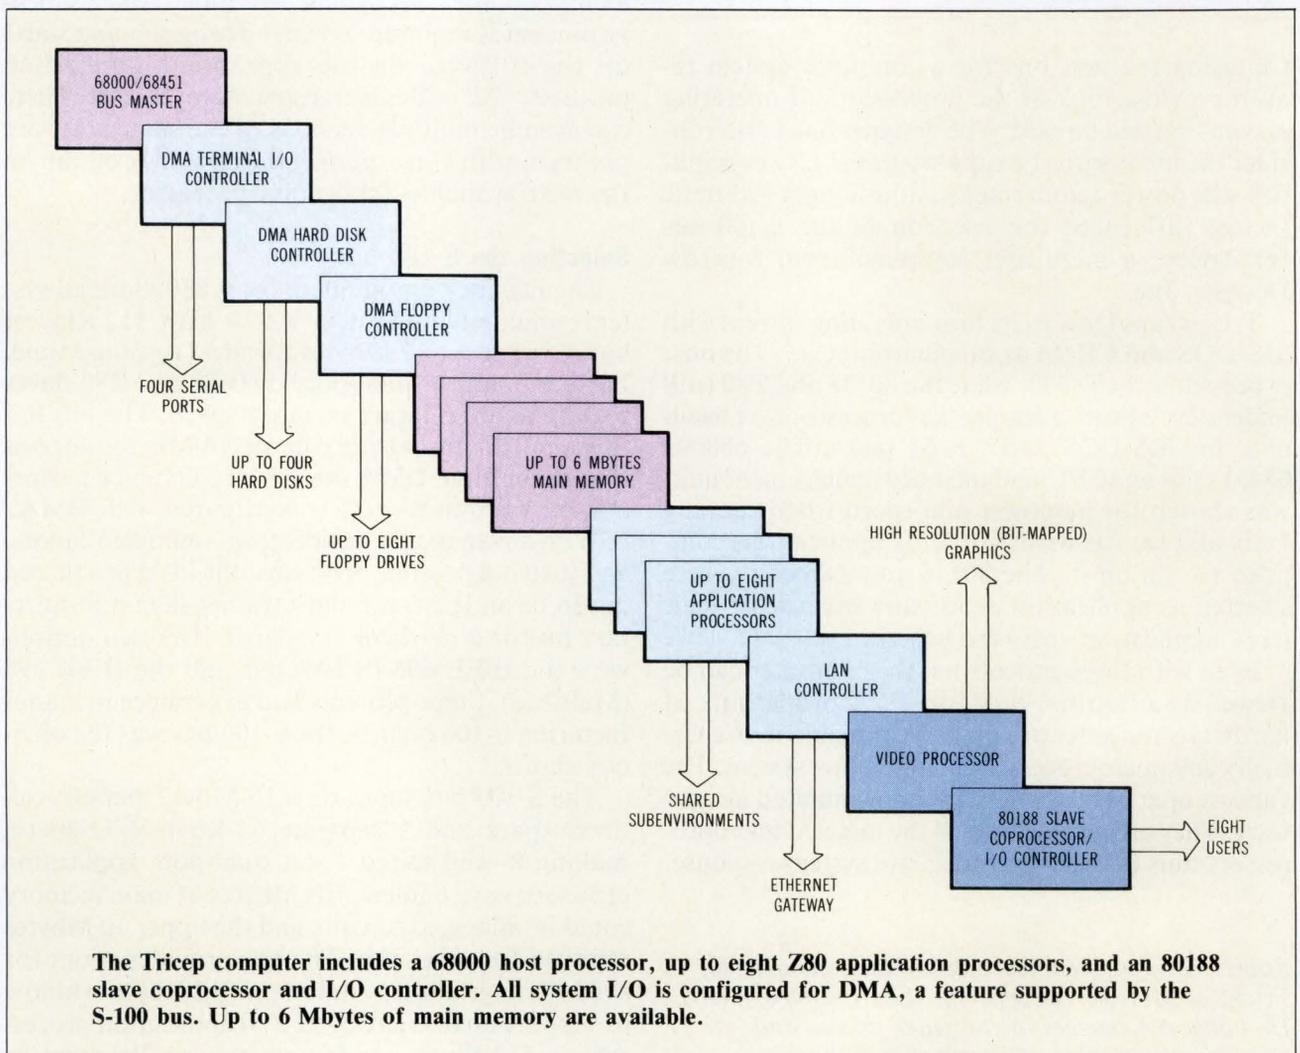
The term "application processor" is used to describe the computer's slave processor boards that have operating system emulators. The dual-port memory (contained in each slave) is accessed from the bus as if it were ordinary memory. In fact, this memory can be used by the master processor, if

desired (with some sacrifice in speed because of arbitration). The feature that differentiates this memory from ordinary bus memory is the onboard microprocessor that turns the board into "smart memory." The addition of software completes the application processor.

DMA I/O maximizes performance

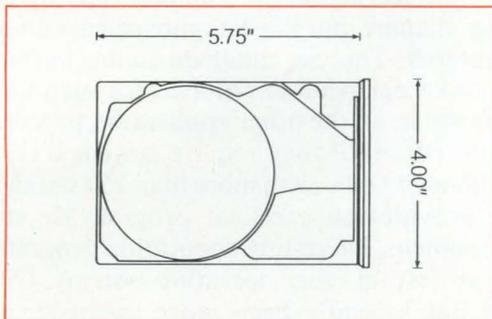
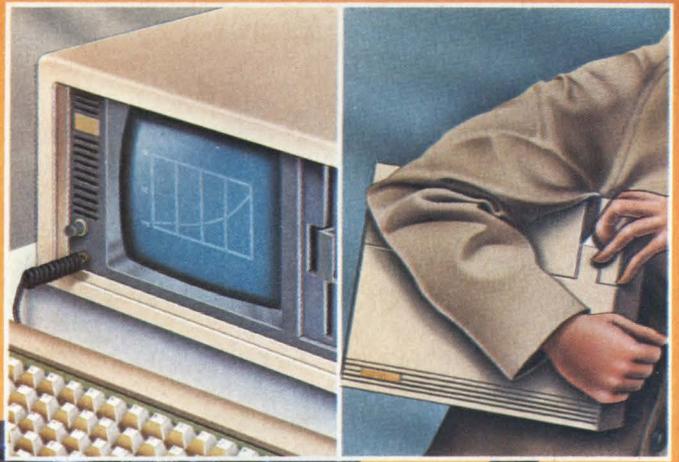
Unix systems can function without DMA for system I/O, but with Tricep, all system I/O—hard disk, floppy disk, and terminal I/O—has to be performed via DMA to meet performance goals. With as many as four DMA devices vying for the bus at once, these devices must arbitrate. The S-100 provides up to 16 temporary bus masters with arbitration. When a DMA device requests access to the bus, it asserts a 4-bit DMA priority. If another device requests access at the same time, the two devices arbitrate. The one with the highest priority gains access and the lower priority device must relinquish and retry later.

The hard disk controller must have highest priority, because all system processes and application



RODIME

World's only
field-proven 3½"
Winchesters



Rodime's 350 Series 3½" Winchesters are operating out there a hundred thousand strong. Our proven technology—microprocessor control of the stepper motor, automatic thermal compensation, special airflow geometry, proprietary LSI chips, fewer parts—has set the industry standard for performance and reliability.

Our 3½s use far lower power, run cooler, weigh

less, and are double shock mounted to take punishment. Also, they can be configured to fit right into a full- or half-height 5¼" slot. They add up to today's best design decision—for computers, terminals, controllers, and all compact instrumentation. Phone or write for our 350 Series data sheet and brochure. (305) 994-6200.

RODIME

901 Broken Sound Parkway, N.W. • Boca Raton, FL 33431 • (305) 994-6200
West Coast—Mission Viejo, CA (714) 770-3085 • Central—Austin, TX (512) 453-5135

processors share the hard disk drives. On the other hand, a floppy disk is used for accessing MS-DOS, CP/M, or Unix diskettes for data interchange or backup. As these are low priority, nonshared devices, the floppy disk controller is assigned the lowest priority. Medium priority levels are reserved for the terminal I/O devices.

Because disk access is critical to the computer's operation, DMA command channels pass commands to the disk controllers. The hard disk controller uses a 16-byte command channel in main memory for operation. The 68000 CPU places disk requests into the channel and outputs to an attention port. The disk controller then picks up the command and performs the desired operation. When the command is completed, the controller generates an interrupt to the 68000.

The last 3 bytes of the command channel point to the beginning of the next command channel, permitting the 68000 to link multiple commands. This allows efficient queuing of hard disk requests by placing each request in a sorted (at least by cylinder) linked list that becomes the disk/command channel. The ability to link multiple disk requests enables the disk controller to keep the request queue at tolerable levels, with minimum supervision by the host processor.

The floppy disk controller has a similar scheme. It continues executing commands, however, until it fetches a halt command from the channel. In addition, it has read/write track commands for increased system performance during system backup.

DMA terminal I/O

All terminal I/O is performed via a DMA I/O board that contains an 8085 processor and four 2661 universal asynchronous receiver/transmitters. DMA cycles from either main memory or application processors acquire data sent to each terminal. All data from terminals is buffered in an input first in, first out buffer and tagged with the appropriate UART number. It is then moved by the 68000 into the appropriate buffer in main memory or into the application processor's dual-port memory.

Putting UARTs on the 80188 application processor board for terminal I/O was considered, to let MS-DOS applications talk directly with the terminal and reduce system overhead. This scheme presented two problems. First, and most important, a terminal will always be locked into one particular application processor—an unacceptable condition for a virtual processor. In such an environment, the user must use specific terminals for specific applications. If a given terminal is busy, the user is locked out. Second, Unix allows any processor to redirect its standard input, output, and error to files and/or other devices. Since I/O redirection is found to be impos-

sible within this scheme, onboard UARTs are rendered impractical.

The application processors in Tricep are not used as DMA devices because of system throughput, security, and the mandatory double buffering of data. If eight application boards are put into the Tricep, for example, system DMA overhead skyrockets. Each time the processor requires I/O, it takes control of the bus, locking out all other processes. This is unacceptable in a multi-user, multiprocessing environment.

Dealing with access

In addition, these processors would have free access to the bus. A process running wild on one processor can corrupt the processes on all the others. In a multi-user software development environment, for example, the risk of one user's bug-laden code corrupting another user's code cannot be tolerated.

Finally, the DMA controllers cannot move data directly into the processor's memory. They must put it in a buffer from which the processor can execute DMA later. This is unnecessary double buffering in an operating system already loaded with buffering.

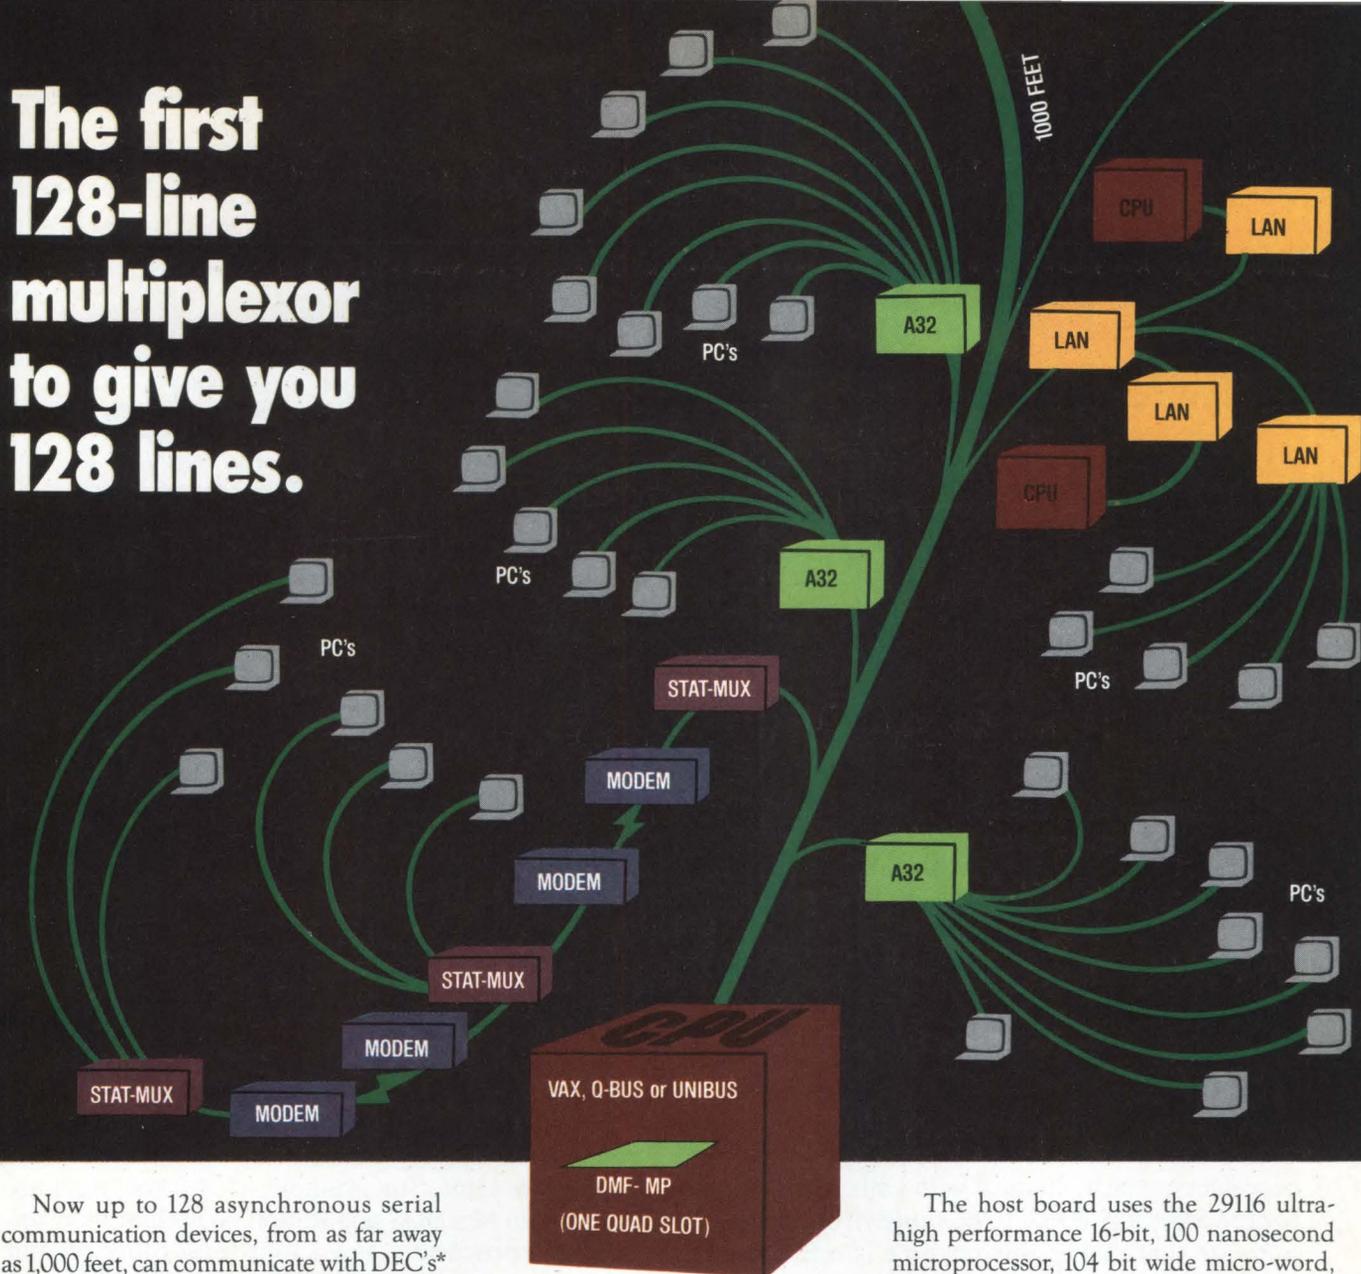
Some S-100 slave processor boards are I/O mapped rather than memory mapped. Although I/O mapping is inherently simpler, it is far less effective in terms of system throughput, requiring either a bidirectional communication port or arbitration to access the onboard RAM. Furthermore, all the DMA controllers work with memory devices rather than I/O-mapped devices.

The application processor concept allows DMA I/O devices to communicate directly with the slave processors via the dual-ported RAM. The master processor still passes the commands to the disk controller channel, but the data movement can be transparent to it. This can eliminate double buffering and reduce kernel memory required for such buffering. Meanwhile, all the other application processors are free to run until they require system I/O.

Although Unix, with more than 250 standard utilities, provides an excellent program development environment, interesting application programs often appear first on other operating systems. Unix software has typically been more expensive than its CP/M or MS-DOS counterparts, partly because of the multi-user license costs. For maximum operating system flexibility, Tricep implements operating system subenvironments as tasks running under Unix.

The first goal is to enable the computer to read and write IBM PC-DOS and CP/M diskettes, and 8-in. CP/M standard diskettes. Because there is no Unix utility for this, two new utilities, Far and Dar, were created. The Far utility reads and writes multiple 5- and 8-in. CP/M disk formats. The Dar utility

The first 128-line multiplexor to give you 128 lines.



Now up to 128 asynchronous serial communication devices, from as far away as 1,000 feet, can communicate with DEC's* VAX, Q-bus and Unibus* computers.

MDB's new DMF Series Multiplexor was developed because we know systems designers need more efficient ways to control statistical multiplexors, LAN's, dynamic asynchronous terminal line switches and interprocessor links, in the new data communications world.

So from configuration to operation and installation, we've thought of everything to make it possible. Featuring a space-saving quad sized format, our DMF series is comprised of a host processor board, one or more asynchronous channel boards (A32), and a selection of connector panels.

It means true 128 line capability because terminals can be placed in up to 16 separate clusters anywhere along the 1,000 foot bus and each channel board can control up to a maximum of 32 RS32 terminals running up to 19.2 K baud (38.4 K baud optional).

*DEC, VAX, Q-bus, Unibus, RSTS/E and RSX11-M+ are trademarks of Digital Equipment Corporation.

The host board uses the 29116 ultra-high performance 16-bit, 100 nanosecond microprocessor, 104 bit wide micro-word, and special logic structures, for the fastest DMF emulation available.

User versatility counts too. The DMF Multiplexor offers such unequalled through-put advantages as the capability to select silo or DMA transfers automatically; and with 16, 18 or 22 bit addressing, the designer can choose the best configuration for the DEC system used. Also, it has its own optimized software drivers for those who have DEC's RSTS/E and RSX11-M+* operating systems.

Now that you can finally use all 128 lines, you may find that you like the bottom line best. All of those technological advances make it the lowest per line cost possible.

Extending the capability of DEC systems is what we're all about. That's why you should contact us immediately with your questions on our new DMF Series Multiplexor or our complete line of DEC compatible system enhancers.

MDB
SYSTEMS INC.

Corporate Headquarters
1995 N. Batavia Street, Box 5508
Orange, CA 92667-0508
Tel. 714-998-6900 TWX: 910-593-1339
FAX: 714-637-4060

MDB Systems U.K., Ltd.
Basingstoke, Tel. 44 0256 464767
TELEX: 858389 MDBSYS G

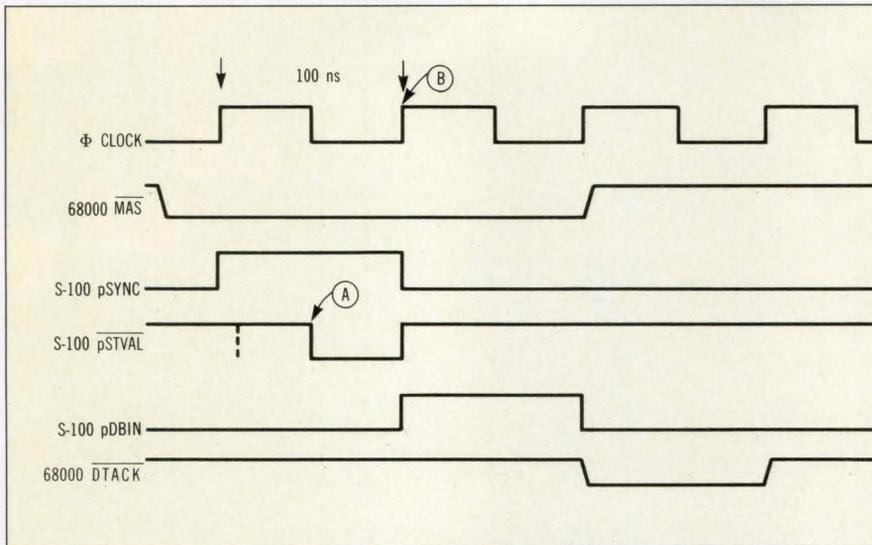
In the western states MDB products are also distributed by

WYLE
LABORATORIES

ELECTRONICS MARKETING GROUP

See us at MINI/MICRO Northeast Booth #221, 223

CIRCLE 64



During the S-100 bus read cycle, time between A (address and status latched) and B (synchronous ready lines sampled) is 50 ns at 10 MHz. This is the maximum time needed to decode the current cycle and generate a synchronous wait.

reads and writes IBM PC-DOS diskettes. These programs automatically copy the files into the selected Unix directory. From then on, the files are in Unix format and can be manipulated by the Unix utilities. The files can be written back into appropriate diskettes and, in the process, converted back to PC-DOS or CP/M formats. The user need not be concerned about the form a file is in, as it can easily be copied into and out of the Unix system. This allows flexibility in terms of data disk exchange and is considered the lowest level subenvironment.

After the files are exchanged, they are executed. Tricep accomplishes this by providing application processors, which, coupled with Unix software, supply the MS-DOS and CP/M subenvironments. The software that makes this possible is referred to as emulation software. In a nutshell, the MS-DOS or CP/M emulator sits between the Unix system and the application running on the application processor. It traps the MS-DOS or CP/M system calls and converts them into Unix system calls. Unix acts upon these requests and then returns control to the application processor. This scheme allows the applications access to Unix I/O redirection and pipes, as well as to the printer spooler, through the emulators.

Tricep allows two modes of access to the emulators. In one mode, MS-DOS or CP/M is run directly on the application processors. Users see the familiar prompts of the appropriate operating system. Standard CP/M and MS-DOS commands execute, and the user is isolated from Unix. (This eases system customization.) Users do not need to know Unix because the more familiar environments are provided. In fact, the Unix password file can be set up so these users are directed into an appropriate directory and into the desired operating system. They are then prohibited direct access to Unix and never see more of Unix than the log-on message.

The second mode of access to the emulators is the auto mode. This allows the user to run any program right from the Unix shell. If the file has a Unix header, it is executed on the 68000 under Unix. If it has an MS-DOS header, it is assumed to be an MS-DOS file and is run on the next available 80188 application processor. Similarly, if the file is CP/M-executable, it will be run on the next available Z80 application processor.

Another choice for application processing

Swapping, another design feature of the application processors, optimizes the computer's performance. Assume, for example, that three users are running an MS-DOS application on each of three application processors and a fourth user logs on and tries to run another MS-DOS program. There are four choices: waiting, adding more slaves so the situation never occurs, simulation, or swapping.

Waiting must be ruled out because the user has no guarantee as to when an application processor will become available. In addition, it is common for a user to leave a terminal running an application. This can lock out all other users for hours. Adding more slaves seems unworkable, because the number of slaves required at a given time cannot be accurately determined. And, adding more slaves than needed for average use is costly in terms of system power and price.

Simulating the 80188 or Z80 instruction set on the 68000 is agonizingly slow and wasteful of critical system resources. Tricep uses swapping, a technique that moves code in and out of memory to let many users share a small, common address space. Code is usually swapped out to disk for temporary storage and swapped back in when the process starts up again. Unix uses this technique extensively. Swapping is initiated when an MS-DOS application processor

is required by a user, but none is available. An MS-DOS process that is inactive (waiting for either I/O or user response) is then swapped to hard disk via DMA from the dual-port RAM. This is done using the Unix raw devices for maximum swapping speed. The new process is loaded into the application processor's dual-port RAM and execution begins. The process just swapped out is swapped back into the next available application processor.

Swapping does have a major drawback, however. The time needed to swap up to 512 Kbytes of application processor RAM to disk is prohibitive (1 to 2 s depending on disk-drive speeds). The use of RAM "dirty bits" can avoid this.

A dirty bit is a tag that goes with each 2-Kbyte segment of application processor memory and resides on the slave. Any time a segment has been written into, the bit is set; otherwise the RAM is "clean." When Unix must swap the process, it examines the dirty bits and swaps only those segments that have been written into or "dirtied." This minimizes the process-swapping time and provides reasonable response.

A slave processor is defined differently from an application processor. A slave processor assists the main processor in executing Unix and in performing general housekeeping tasks. An application processor runs a particular user application and provides the user with an appropriate subenvironment. The 80188 processor board, however, bridges this gap; it can be configured either as an application processor (providing the MS-DOS subenvironment) or as a slave processor. Configuration is accomplished by replacing an erasable PROM and adding a UART board.

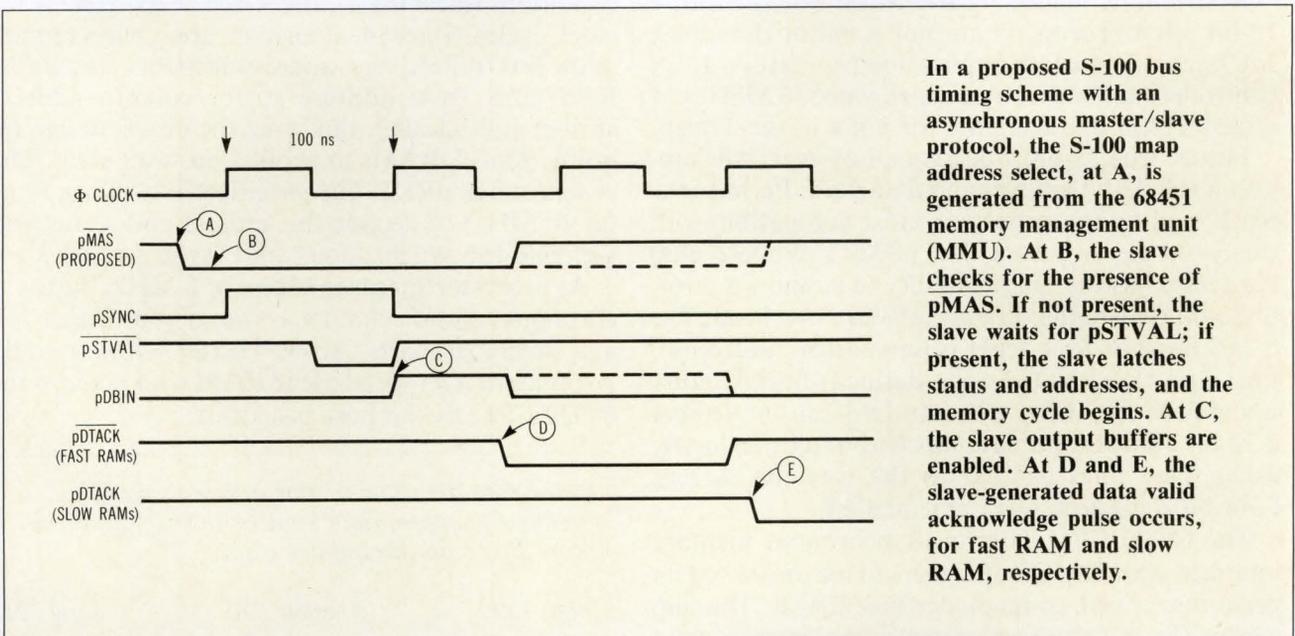
Connecting onboard UARTs directly to a terminal is impractical. An alternative, however, is designing the board to control eight UARTs via an expansion connector exclusively. The board can then become an I/O slave processor handling all terminal I/O chores in the system.

With up to 512 Kbytes of RAM, this 80188 slave processor can actually contain and execute the Unix I/O drivers, removing this burden from the host 68000. With two DMA channels, interval timers, and fast context switching, the 80188 makes an ideal I/O processor. In addition, the system can be configured so the 80188 slave processor controls the DMA disk channels. This is done by locating these channels in the 80188 memory space, thus off-loading the 68000 even more.

Problems to overcome

Using the S-100 bus presents some challenges. These include power supplies, cooling, 16-bit data paths, and 24-bit address space. In addition, the bus uses a synchronous RDY (ready) line, while the 68000 CPU is designed with an asynchronous memory interface.

Perhaps the major disadvantage of the S-100 bus is its power supply requirements. The specification calls for the bus to provide filtered 8 V and filtered 16 V and -16 V. Each card then has three-terminal regulators for proper regulation to the 5 V, 12 V and -12 V requirements of most circuits. The onboard regulators also help isolate onboard supply noise from the system supply and vice versa. A power supply design consisting of no more than a constant voltage transformer, appropriate rectifiers, and filter capacitors—a design that is fairly easy to develop—



would suffice. The obvious drawback is loss of board real estate to redundant regulators and excess heat localized on each board. Another drawback is that disk drives requiring regulated supplies are needed for worldwide use.

Today, switching power supplies are available for as little as 50 cents per watt and provide low cost and efficient, compact size. These supplies have standard voltages of 12 V, -12 V and 5 V—perfect for disk drives, but not for the S-100 bus.

Although some S-100 companies have eliminated the onboard regulators and use one of these off-the-shelf switchers, such a design compromises the bus's interchangeability. Moreover, the 8-V power requirement seems within the realm of current design, although not quite off-the-shelf.

A slightly modified off-the-shelf 250-W switcher proves to be both efficient and compact compared to its linear alternative. It also provides the necessary holdup times for UPS requirements (25 ms) and 110-V and 220-V, 50/60-Hz operation for international use. The supply's closed-frame package simplifies manufacturing and field replacement as well.

Excess heat from the onboard regulators is dissipated by positive pressure air cooling from the well-designed cabinet venting. Using four- and six-layer board designs with 10 x 10 trace spacings, where necessary, minimizes board real estate problems. For example, although the 80188 application/slave processor board contains up to 512 Kbytes of dual-port RAM, 8 Kbytes of EPROM, I/O memory expansion ports, and interface logic, it is a four-layer design at 12 x 12 spacings.

Address and data lines

Restrictions placed by the S-100 bus on future 32-bit microprocessors are not a major drawback for Tricep. Even with application processors, LAN controller, and memory-mapped video, 6 Mbytes of system RAM are available for Unix in the Tricep.

Future true 32-bit processors, however, will support a full 32-bit address and data path. Perhaps the 68020 will be the processor most compatible with the S-100 bus structure. The 68020's onboard chip hardware has the unique ability to support 8-, 16-, or 32-bit data paths with no software overhead. The S-100 bus has four RFU (reserved for future use) lines and three NDEF (not defined) lines. Future updates of the IEEE 696 standard can implement a 32-bit address and data bus with multiplexing by using these lines (similar to the way the current 8/16-bit data access is implemented).

The 68000 CPU with an asynchronous memory interface allows system designers to maximize system performance with common dynamic RAMs. The chip asserts the \overline{AS} (address strobe) signal when a valid

address is available. The 68451 memory management unit transposes the address to form the mapped address and asserts the \overline{MAS} (mapped address strobe) signal. The 68000 then waits for the return of the DTACK (data transfer acknowledge) signal from the memory, indicating memory cycle completion.

The S-100 bus, however, is designed around the Intel 8080 family of processors and uses a synchronous RDY line. A workable S-100 implementation requires the 68000 to run in a synchronous mode and the 150-ns system DRAM to run with one wait state for a 10-MHz operation. A more effective solution is to revise the standard to support asynchronous devices by assigning two reserved lines for \overline{MAS} and DTACK handshaking. If the host processor does not assert \overline{MAS} at the beginning of a cycle, the memory slave responds as a synchronous slave (current S-100 mode). If it does assert \overline{MAS} , the memory slave must respond with DTACK, indicating completion of the requested operations. This would have to be handled as a jumper option on the host processor to ensure backward compatibility.

Without such an arrangement, memory constraints at even higher processor speeds become a timing nightmare, as a 10-MHz 68000 in synchronous mode on the S-100 bus illustrates. The bus cycle begins with pSYNC (processor synchronization) (generated from the 68451 signal \overline{MAS} and the rising edge of the bus clock). The standard specifies that the address and status lines be valid 40 ns prior to the S-100 line pSTVAL (status valid), which is usually generated from the falling edge of the clock.

The standard also dictates that the pDBIN (processor data bus input) not overlap pSTVAL. If pDBIN starts 10 ns after pSTVAL, even under optimum conditions (ideal logic), the RAM access time is 1.5 clock cycles. Since ideal logic is rare, delays remain in the bus transceivers, address decoders, and RAM RAS/CAS (row address strobe/column address strobe) multiplexing, thus even the fastest design requires 90-ns DRAMs to avoid a bus wait state. The synchronous pRDY line specifications allow 50 ns (at 10 MHz) to decode the address and status and generate the synchronous wait state.

At processor speeds in excess of 8 MHz, the asynchronous solution eliminates the copious use of 74F and 74AS logic parts. An even better solution would be onboard high speed cache RAM with perhaps the 68020 CPU as the host processor.

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

MULTI-USER System of the Future TODAY



We don't sell PC's. We sell sophisticated Multi-User Systems.

Advanced Digital offers you the ability to mix 8 and 16 Bit systems and still maintains popular PC software compatibility. Select your ideal system from a complete product line of 8-Bit and/or 16-Bit processors. You can choose your CPU speed (4, 6, or 8 MHz), and even specify the exact on-board RAM you need (from 64KB to 1MByte).

Determine the number of users you'll need, now and in the future, and pick from the 1 to 12 user SUPER STAR or the 1 to 30 user SUPER SYSTEM II. Both systems offer a wide selection of Floppy Disk and Winchester Fixed Disk configurations. Both systems offer your choice of removable cartridge or tape back-up. Both systems can interface to the IBM-PC.

And only Advanced Digital gives you such a wide selection of Operating Systems. There is 8-Bit CP/M, 16-Bit MS-DOS, 8/16-Bit NETWORK-OS, and 8/16-Bit multi-user TurboDOS with MS-DOS Compatibility. Each Operating System features a vast library of existing software application programs.

See Advanced Digital's complete line of multi-user problem solvers at your nearest authorized dealer, or contact the Sales Department at:

Advanced Digital Corporation
5432 Production Dr., Huntington Beach, CA 92649
(714) 891-4004—Telex 183210 ADVANCED HTBH

Advanced Digital U.K. Ltd.
27 Princes St., Hanover Square
London W1R8NQ—United Kingdom
(01) 409-0077/(01) 409-3351—Telex 265840 FINEST

 **ADVANCED
DIGITAL
CORPORATION**

SUPER STAR and SUPER SYSTEM II are trademarks of Advanced Digital Corporation. CP/M is a trademark of Digital Research Corp., MS-DOS is a trademark of Microsoft Corp., TurboDOS is a trademark of Software 2000, and NETWORK-OS is a trademark of CBIS, Inc.



The first family of

A little fanfare, please.

We're introducing the industry's first full line of DEC VT200™ series compatible terminals. And while our competition won't find cause to celebrate, it will provide major advantages for you. Not to mention your budget.

That's because each of our three Visual 200 series terminals was designed to outperform, out-feature and outprice its DEC counterpart. So with larger screens, higher resolution and better ergonomics,

the Visual 200 series isn't merely an alternative to the full VT200 line, it's your best alternative, terminal for terminal. But see for yourself.

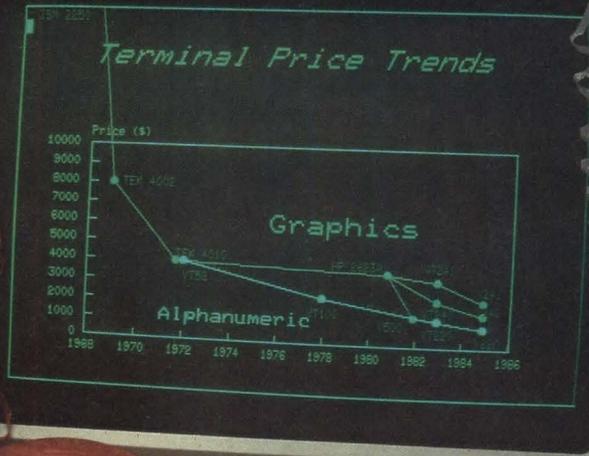
The Visual 220

This alphanumeric terminal relies on a lot more than its sleeker design to outpace DEC's VT220™. Its eye-saving, operator-friendly 14 inch tilt-and-swivel display, for example, instead of a 12 inch tilt-only. Or its 30 user-programmable function keys as opposed to 15 computer-programmable.

Or its home key on the keyboard. And it's 25th status line with time of day clock. All for \$995, significantly less than the VT220.

The Visual 240

Here's a feature-packed monochrome graphic terminal with exceptional performance characteristics. Even more exceptional when compared to the DEC VT240™. In addition to ReGIS® emulation, it provides superior full-screen emulations of the Tectronix 4010/4014, with higher resolution; 30 function keys (15 programmable,



DEC™ compatibles.

non-volatile); and a visibly more compact design. You can also support more peripherals, including color printers and plotters, as well as bit pads and mice through a second connector.

The Visual 241

The Visual 241 provides all the advantages of the Visual 240 with one important addition—color. Not just a little color, but the full spectrum of colors you'll need to bring clarity and meaning to your graphic displays. With the Visual 241, you

can optionally access up to four times the color of the DEC VT241™ 16 simultaneous colors from a palette of 64, as opposed to four colors from a choice of 16. All for the same price you'd pay for monochrome graphics on the VT240—\$2,195.

The Visual 200 series. The industry's first family of DEC compatibles. From the industry's leader in DEC compatible technology. Delivering new levels of 200 series performance, features and ergonomics at substantial savings to you. Call Visual

today for more information about the Visual 200 series.

And prepare yourself for a celebration.

VISUAL See for yourself.

Visual Technology Incorporated

DEC™ is a registered trademark of Digital Equipment Corporation. ReGIS®, VT200™, VT220™, VT240™, VT241™ are trademarks of Digital Equipment Corporation.

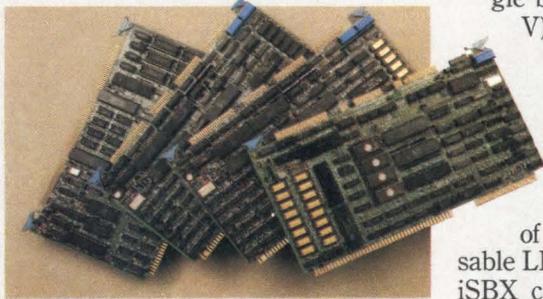
FOR THE MICROCOMPUTER OEM OR END-USER HEURIKON IS THE CHOICE

Make Heurikon your choice for Multibus™ microcomputers and system components.

For 12 years, Heurikon has supplied thousands of computers to help its customers find economical solutions for their microcomputer applications.

8-BIT MICROCOMPUTERS

MLZ-90A single board microcomputer with nine byte-wide memory sockets for use with RAM or ROM (AM9511 and floppy disk drive controller optional).

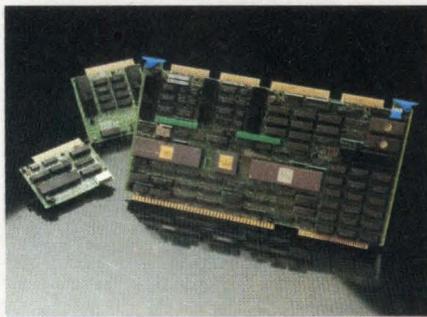


MLZ-91A single board CP/M™ system with on-card floppy disk drive controller, winchester interface, optional AM9511, streamer tape interface, two serial ports, one parallel port, 64K or 128K bytes RAM with parity, two EPROM sockets, and GPIB CONTROLLER.

MLZ-92A single board CP/M™ system with four serial ports on-card, floppy disk drive controller, winchester interface, optional AM9511, Centronics printer interface, 64K or 128K bytes of RAM with parity, and two EPROM sockets.

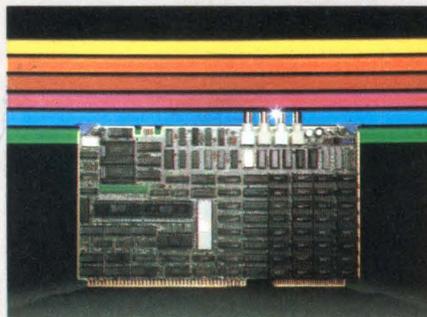
MLZ-93A single board CP/M™ system with 128K bytes of dual ported RAM, four EPROM sockets, floppy disk drive controller, optional AM9511 and powerful serial port features including SDLC and HDLC protocol support and modem controls.

16-BIT MICROCOMPUTERS



HK68™ powerful and versatile single board UNIX™ (System III or V) or CP/M-68K™ system with 68000/68010 CPU (8Mhz or 10Mhz), MMU, quad channel DMA, four serial ports, 128K, 256K, 512K or 1M bytes of on-board RAM with parity, up to 128K bytes of on-board EPROM, user accessible LEDs and dip switches, and two iSBX connectors for I/O expansion. Heurikon can also supply a full line of iSBX I/O expansion modules including quad channel serial port module, floppy disk controller, A/D converter, bubble memory and many more.

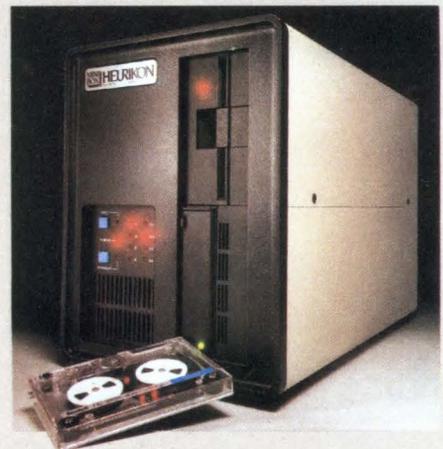
GRAPHICS CONTROLLERS



MLZ-VDC intelligent 640 x 480 x 4 color graphics controller based on the NEC 7220 controller chip with on-

board Z-80 CPU, DMA controller, and user definable FIFO interface to Multibus™. Users may display up to 16 colors from a 4K palette. Up to 1024 x 1024 x 3 interlaced also available.

MICROCOMPUTER SYSTEMS



MINIBOX

Heurikon also provides completely integrated UNIX™ development systems with UNIX™ System III or System V including Berkeley enhancements. CP/M-68K™, PolyFORTH™, Regulus™, and VRTX™ (real time operating system kernel) are also available.

Heurikon UNIX™ systems are available in four, six, and fourteen slot enclosures with 30MB, 65MB, 140MB or 280MB of winchester storage, a megabyte floppy, and optional *interactive* on-line streamer tape drive supporting 1-16 users.

HEURIKON CORP

Microcomputers For Industry

3201 Latham Dr.
Madison, WI 53713
Telex 469532

Call Heurikon Direct

1 800 356-9602

In Wisconsin 1 608 271-8700

Z-80 is a trademark of Zilog Corp. Multibus and iSBX are trademarks of Intel Corp. UNIX is a trademark of Bell Lab. CP/M 68000 and CP/M-80 are trademarks of Digital Research. HK68 is a trademark of Heurikon Corp. Regulus is a trademark of Alcyon. VRTX is a trademark of Hunter & Ready. PolyFORTH is a trademark of FORTH, Inc.

NEW VAX SQUEEZES MAINFRAME POWER INTO MINI PACKAGE

While retaining features common to the 32-bit VAX family, the 8600 uses large-scale processing concepts that include a dedicated memory bus and pipelined operations in both instruction processing and memory references.

by Trygve Fossum,
Jim McElroy, and
Bill English

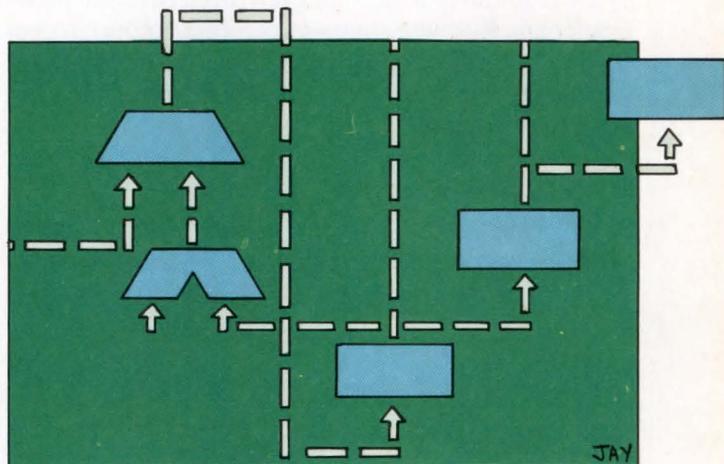
In designing the 8600, innovation centered on achieving a fourfold increase in performance over the industry-standard VAX-11/780. This required execution of the full VAX instruction set, running the VMS operating system, and interfacing to the same I/O bus structures and network links. To account for all these factors, the new VAX uses ECL macrocell arrays with high switching speed and integration level, dedicated memory interface, multiple I/O adapters, and pipelined instruction processing through subsystems with dedicated hardware.

Improved performance comes from doing more work during each cycle, and from having a shorter cycle time. The 80-ns cycle results from using the ECL

Trygve Fossum is a design engineer at Digital Equipment Corp (Marlboro, Mass). He has a PhD in mathematics from the University of Illinois.

Jim McElroy is manager of advanced development for the high performance systems technology group at DEC. He has an MS in mechanical engineering from Northeastern University.

Bill English is a technical writer for DEC. He has an MS in mathematics from the University of Massachusetts-Amherst and a BA in physics from Harvard University.



macrocell arrays. Each cell consists of transistors and resistors that can be connected to form such logic functions as latches and multiplexers. Each chip contains about 1000 gate equivalents, with gate delays averaging 1.5 ns. Signal reflections are reduced by minimizing interconnect delays and by carefully controlling wiring impedance. Up to six signal layers are required to interconnect the devices mounted on a printed circuit board. Wind tunnel techniques were used to develop placement algorithms, thus ensuring uniform cooling.

Modules and backplanes are supported by a precision one-piece card cage that also acts as a plenum for the cooling air flow. Backplanes contain 16 layers of printed wiring in a laminated structure. All connectors are the solderless press pin type, with compliant pins for long-term electrical contact. Power regulators are above the logic assembly for a straight-through, single-path air flow. In all of these, extensive

developments were made in interconnect, packaging, and cooling—both to complement the semiconductor technology and to meet the new environment and safety standards.

Extensive error detection and recovery circuitry was designed into the machine to minimize the effects of failures on the user, and to reduce system downtime. Low component count, worst-case logic design, and highly reliable parts enhance the overall reliability of the system. Monitoring such environmental factors as temperature and air flow detects hazardous operating conditions.

There is also extensive parity checking of data paths. When an error is detected, the storage elements are preserved and the error is reported to an error analysis program. Status flags are used to decide whether the instruction can be retried. Usually, the process can be continued where it left off. With ECC on both cache and memory data, single-bit errors are corrected online, and double-bit errors are detected. Control Store parity errors are corrected by ECC in the Console, and rewritten over the Serial Diagnostic Bus with only a slight interruption in processing. The Floating Point unit, the F Box, runs self diagnostics during the execution of non-floating point instructions for increased reliability.

Independent boxes

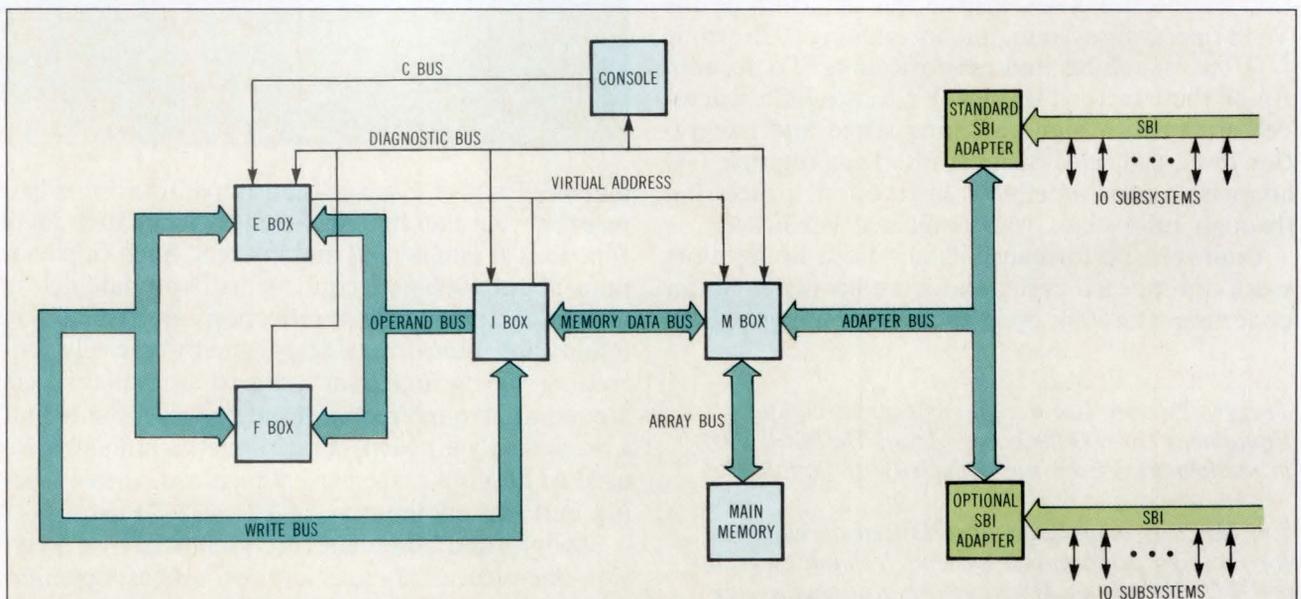
The 8600 consists of six relatively independent subprocessors: E Box, F Box, I Box, M Box, Console, and I/O adapter. The E Box executes the VAX

instruction set and generally directs the entire system. The I Box prefetches instructions and operands, and decodes them for later execution by the E Box. This gives the machine a pipelined structure, allowing several instructions to be present in the I and E boxes at the same time. With this pipeline, some frequently executed instructions in the E Box are completed in a single machine cycle of 80 ns.

To increase the speed of memory access, the M Box contains a 16-Kbyte data cache with a translation buffer for holding recently used translations of virtual-to-physical memory addresses. This box interfaces the memory to all other parts of the system; it also interfaces the E, F, and I boxes to the adapter bus for input and output. A "memory reference" by one of the other boxes happens in a "cache cycle." This is because the design objective is to deal solely with the high speed cache as often as possible. The M Box actually references main memory only when needed data is not in the cache, or to make room in it for new data.

The F Box is a floating point processor or accelerator. It intercepts floating point instructions as presented by the I Box. Special hardware for fast unpacking, aligning, adding, multiplying, and dividing produces high performance for scientific, computational number crunching.

Basic to accurate and fast manipulation of data are 16 general-purpose registers. Four copies of these registers are kept to guarantee very fast and flexible access and instruction retry.



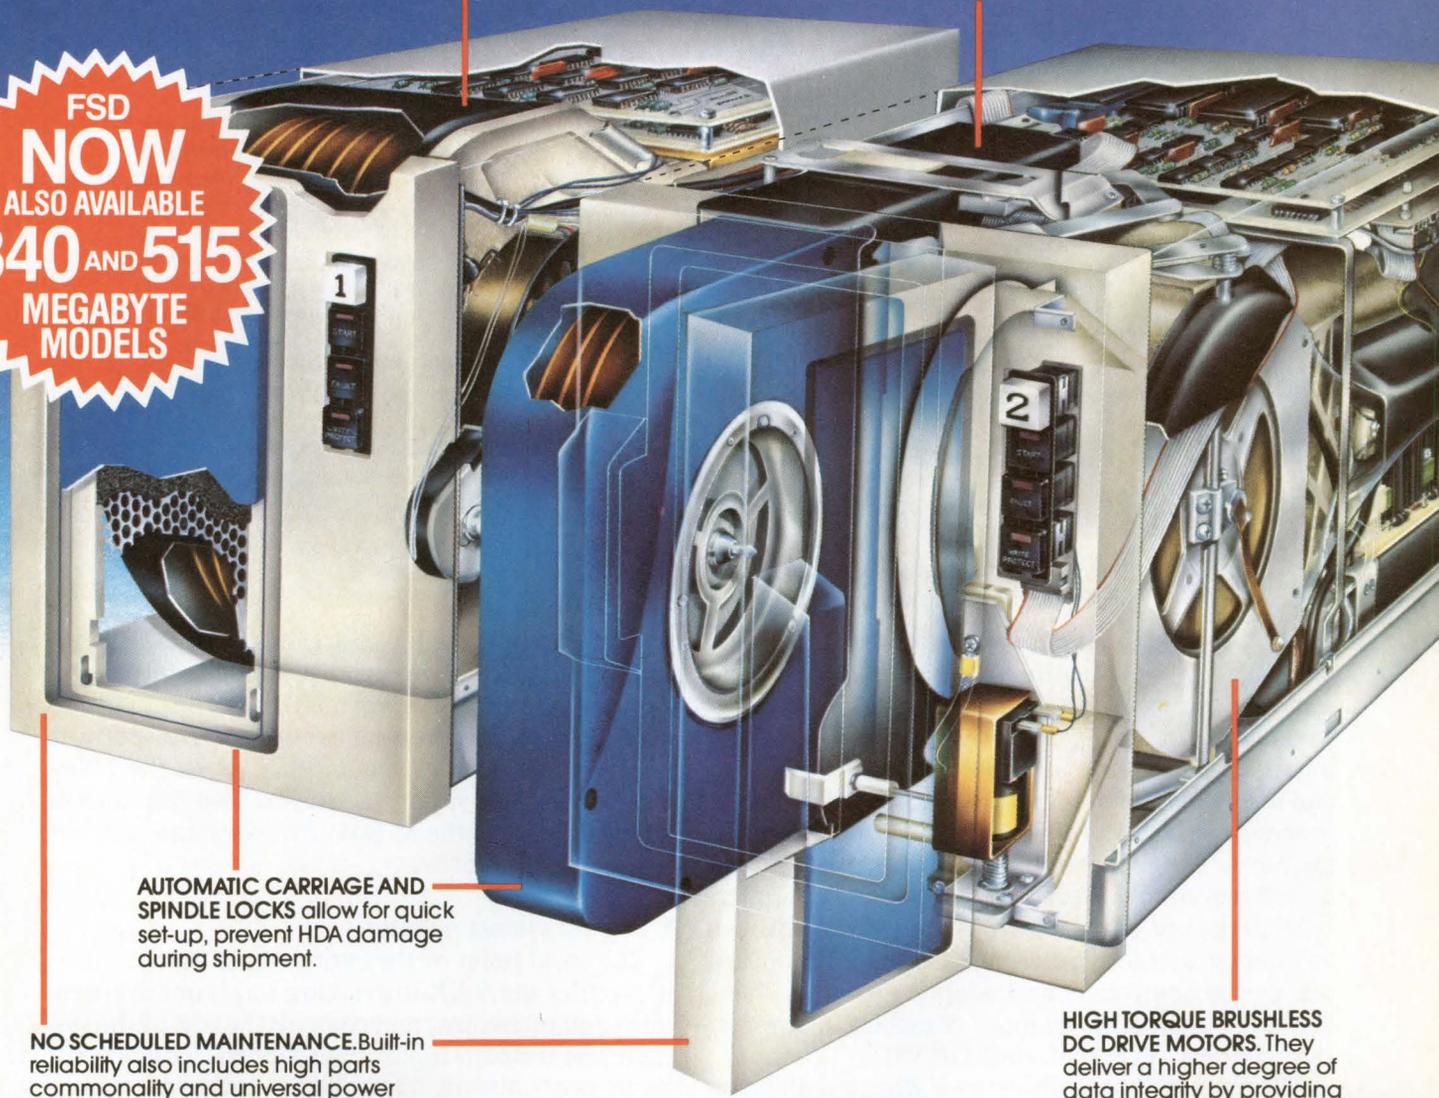
The VAX 8600 has six major subsystems: instruction box, execution box, floating point accelerator, memory, Console, and I/O. Register data is kept in the I, E, and F boxes, with all copies kept current via the write bus. Memory data moves in a loop: operands from M to I over the operand bus to E and F, results move back over the write bus through I to M. I/O is based on DEC's standard Synchronous Backplane Interconnect (SBI) for compatibility with all current peripherals in the VAX-11/780 family.

INTRODUCING LARGE-DISK PERFORMANCE IN DRIVES HALF THE SIZE

MODEL 9715 FSD. With 160 Mbytes in a sealed module, you get the same capacity, speed and performance as the CDC™ Mini Module Drive (MMD) in a unit one-half the size.

MODEL 9710 RSD. With 80 Mbytes in removable data packs for unlimited storage. Has the same capacity, speed and performance as a CDC Storage Module Drive (SMD) in a unit one-half the size.

**FSD
NOW
ALSO AVAILABLE
340 AND 515
MEGABYTE
MODELS**



AUTOMATIC CARRIAGE AND SPINDLE LOCKS allow for quick set-up, prevent HDA damage during shipment.

NO SCHEDULED MAINTENANCE. Built-in reliability also includes high parts commonality and universal power supply (100-240V, 50/60 Hz) for easy installation worldwide.

BOTH MEET FCC, UL, CSA, VDE STANDARDS FOR A STAND-ALONE UNIT.

HIGH TORQUE BRUSHLESS DC DRIVE MOTORS. They deliver a higher degree of data integrity by providing rapid disk acceleration with minimum head drag. (Both models)

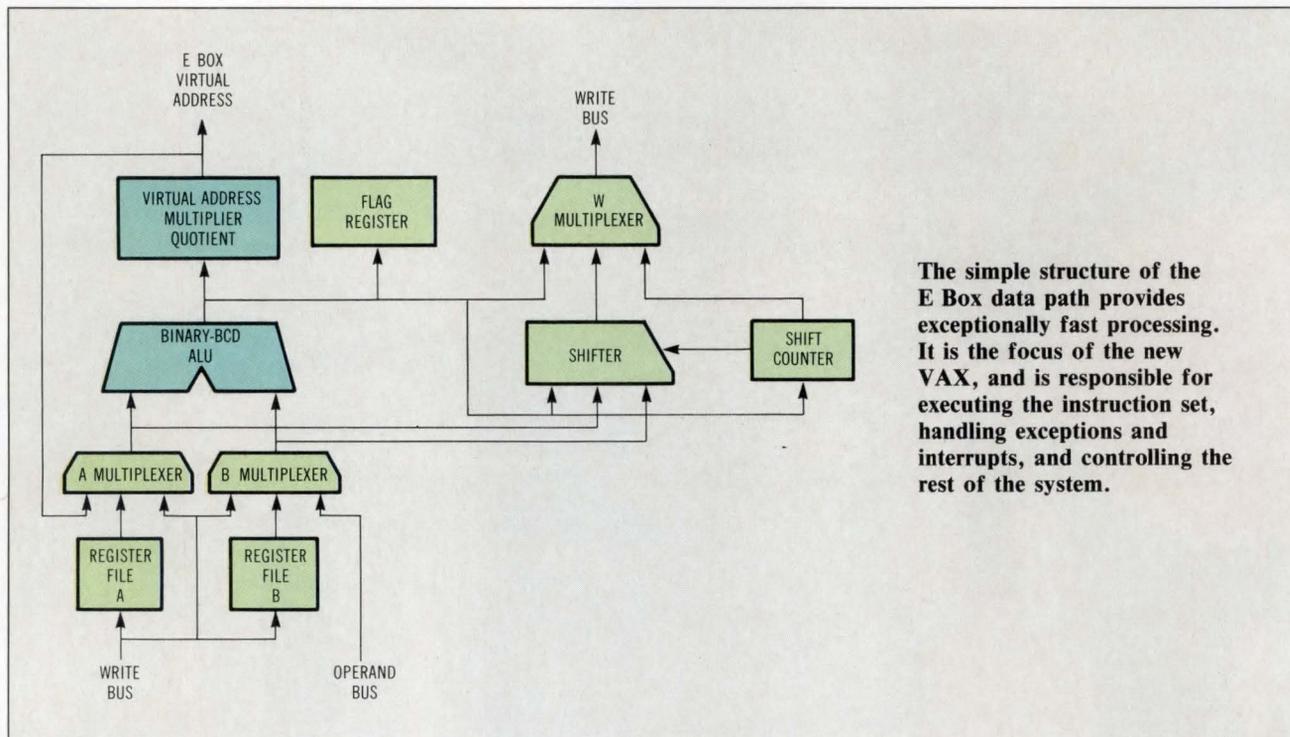
T H E F S D™ / R S D™ S E R I E S

Both 9710 and 9715 Drives use LSI circuitry for all read and write, fault, transmitter/receiver functions and a μ P for servo control, for full performance in half the space. For more data call your local Control Data OEM Sales Representative or write: OEM Product Sales, HQW08X, Control Data Corporation, P.O. Box 0, Minneapolis, MN 55440.



GD CONTROL DATA

Information Hot Line 1-800-828-8001
in Minnesota 612-921-4400, ext. 82



The simple structure of the E Box data path provides exceptionally fast processing. It is the focus of the new VAX, and is responsible for executing the instruction set, handling exceptions and interrupts, and controlling the rest of the system.

The Console is a microprocessor-based frontend interface to the operator, the boot device, and remote diagnostics. In addition, it is used for environmental monitoring and error recovery. This unit initializes the system on power-up, tests the system, and assists in isolating faults. It is based on a PDP-11 processor, with a 10-Mbyte local disk, and 256-Kbyte memory. The I/O system is based on DEC's standard Synchronous Backplane Interconnect; it interfaces to the rest of the system via the M Box through an adapter on the adapter bus. On the SBI are various device controllers and adapters to other interconnects, including Unibus, Massbus, Computer Interconnect, Ethernet, and DR780.

The main microcode is in the E Box, but all boxes are microcoded. This allows them to perform complex functions with a small amount of hardware, which, in turn, provides design flexibility and a good cost/performance ratio. All microcode storage is writable. This facilitates changes and additions whenever necessary. Initially, the RAMs are loaded from files stored on a removable disk in the Console subsystem. The control store is also used for microcoded diagnostic programs to pinpoint a failing component.

Buses interconnect the various boxes. All movement of data between the processor and both the memory array and the I/O subsystem occurs via the memory data bus connecting the M Box to the I Box. The I Box receives the instruction stream and the memory operands over this bus. Operands are passed onto the E and F boxes over the operand bus. Results

from either of these boxes are sent via the write bus to the I Box, which in turn passes them onto the M Box over the memory data bus. The write bus is also used for keeping the four sets of general-purpose registers identical to one another. Both the I Box and the E Box supply addresses that are almost always virtual to the M Box. All buses and registers handle 32-bit words.

E Box as system director

The focal point of the entire system is the E Box: it executes the VAX instruction set, handles exceptions and interrupts, and controls the rest of the system. The E Box is highly microcoded, with most of its elements directly controlled in each cycle by bits in the microword. Intensive microcoding makes it possible to use a data path with a very simple structure; its power comes from the speed and ease with which the microcode manipulates it.

The E Box contains a dual-ported scratchpad memory comprised of 256 32-bit registers. This scratchpad includes basic machine registers, copies of the general-purpose registers, about 150 constants and microcode temporaries, and some architecturally defined registers used by memory management and the operating system.

A 32-bit-wide ALU performs all the usual functions for add, subtract, OR, and exclusive OR, among others. Moreover, some special ALU functions speed up divide, decimal arithmetic, and comparisons. But most significant, microcode can take any two values from the scratchpad, operate on them



Our Smart-Frame™ will put you in touch.

Carroll Touch

a subsidiary of AMP Incorporated

In Touch with Technology

One of the easiest ways to interact with a computer is through a touch activated display. One that lets you merely point your finger at what you want your computer to do.

With touch, even the most complex application can become user friendly. And less susceptible to user error.

But the advantages of the highly reliable infrared touch input systems have been overshadowed by cost. Placing them out of the range of medium to low-cost turnkey systems manufacturers.

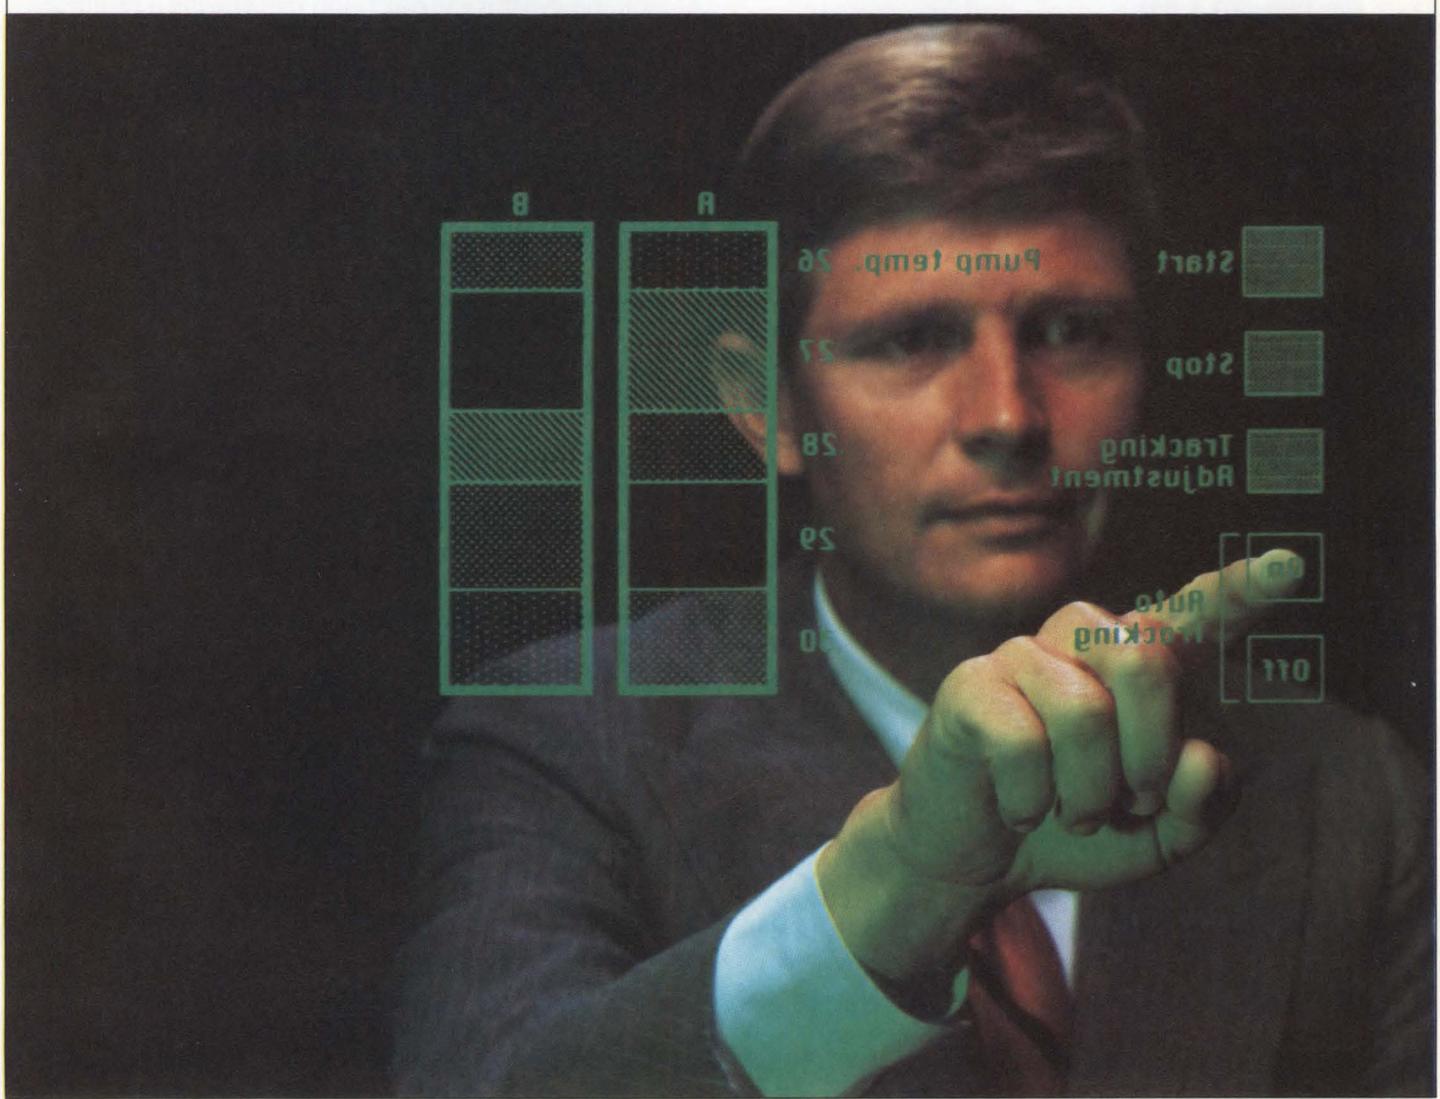
Which is the main reason Carroll Touch completely redesigned its line of touch input products.

We call it Smart-Frame™. Because it contains a powerful microprocessor built into the frame. And because it uses scanning infrared technology—a touch technology which is incredibly reliable, rugged, accurate, and fast. A technology Carroll Touch pioneered and refined.

As a result of the Smart-Frame design, Carroll Touch units contain 45 percent fewer components than before. Which means the cost to you is about 50 percent lower too.

We think that's pretty smart. And Carroll Touch did it just so you could get in touch. Today.

P.O. Box 1309
Round Rock, Texas 78680
512/244-3500, Telex 881906



in the ALU, and store the result back in the scratchpad—all in a single cycle. With this capability, some whole instructions can be completed in just one cycle; and such longer, repetitive instructions as those handling character strings can be executed in very short loops.

A barrel shift network works in parallel with the ALU to accept a 64-bit value, join it end to end, and select any desired 32 consecutive bits from the ring. Two scratchpad registers or one register concatenated with memory data can supply the input. Control over the shifter can be exercised explicitly by a field in the microword, or indirectly by a shift control register. The shifter unpacks and packs floating point data, translates different decimal data formats, does arithmetic shifts and rotations, and performs other bit manipulations. As in the case of the ALU, the shifter power is enhanced by the ability of microcode to take any two words in the scratchpad, shift them, and store the result back in the scratchpad during the same cycle.

The two-part I Box

The VAX architecture contains a variety of opcodes and specifiers (that is, it has several ways of storing operands and results). While this variety is useful to programmers and writers of compilers, the task of decoding these opcodes and specifiers constitutes a large portion of the total work in processing VAX instructions. Hence, the VAX 8600 has a separate subsystem dedicated to prefetching instructions, decoding them, fetching source operands, and storing results. It also receives condition codes from the E Box and makes all branch target fetches and decisions. Often, this work overlaps the actual instruction execution in the E Box and F Box, thereby providing a high degree of simultaneous processing.

The I Box consists of two major parts, an instruction unit and an operand unit. The former contains an 8-byte FIFO buffer that receive instruction stream data from memory 4 bytes at a time. The unit evaluates these bytes to determine the addressing mode and to optimize instructions. A decode RAM, containing information specific to the individual opcodes and specifiers, helps with evaluation.

With the information the instruction unit supplies about the operands' locations, the operand unit generates the operands' addresses and starts the memory reads to fetch them. Because general-purpose registers are needed to calculate the addresses, the instruction unit maintains its own copy. Often, the operands are in the general-purpose registers. In such a case, they are read directly or passed on as general-purpose register numbers, for use by the execution unit (E and F boxes). At other times, the operands are contained in the instruction stream itself. In this

case, the operands are extracted from the instruction buffer. Whenever possible, the instruction unit tries to process two specifiers in a single cycle, handling the second one as a general-purpose register number. This saves cycles in frequently used instructions. When the E Box is ready, the I Box supplies the operands and provides a dispatch address for the E Box to start the appropriate microcode.

Of course, whenever work is done in parallel hazards will exist. The pipeline cannot always operate at full speed due to conflicts produced by the various subsystems needing the same resources. Since several stages may be active simultaneously, control of each stage is closely tied to operations (past and present) in the other stages, as well as to operations in the E and M boxes. Each stage attempts to process available input data as quickly as possible. When input is unavailable or a result cannot be stored right away, the stage is "stalled." One object of the I Box and of the pipeline structure in general is to minimize the time any stage spends in a stalled state. This is desired because no useful work occurs during a stall. A common event is the execution unit storing a result in a register that is needed by the operand unit for the next instruction. Scoreboards and conflict detectors help to catch such cases and allow all units to proceed when hazards are absent.

Conflicts are often avoided by passing data along as general-purpose register tags, rather than passing the actual data. The VAX architecture normally precludes writing into the instruction stream, so the instruction buffer can prefetch freely across most instructions.

Conflicts are often avoided by passing data along as general-purpose register tags, rather than passing actual data.

In keeping with the principle of a high speed, economical implementation, the VAX 8600 uses the instruction buffer to fetch data for string and other multiple-operand instructions. It thereby makes use of hardware that would otherwise remain idle. This procedure moves large amounts of data through the processor without wasting cache cycles, a feature that is especially important in commercial applications where data manipulation is more important than arithmetic speed.

Memory and the M Box

The memory system includes the storage array boards and the M Box. The M Box contains not only all the control, transfer, and error logic for the storage array, but also a data cache for fast access

is least recently used. When a memory word containing a corrected error is placed in the cache, the written bit is turned on to force eventual rewrite of the storage location, thus reducing the probability of a double error.

Addresses actually supplied to the cache or the memory array are always physical, and the DMA references made by the I/O system always use physical addresses. Three sources of memory references exist within the processor: instruction buffer, operand unit, and execution unit. Each of these has its own port into memory.

Normally these are virtual references—that is, the addresses have to be translated from virtual to physical before they can be used to access the cache. When a virtual reference is made, M Box microcode uses the high order part of the address to index into a table called the translation buffer. The translation buffer itself is a cache containing the most recently used translations. The entry from this buffer is then added as a prefix to the remaining bits of the virtual address, forming the desired physical address. The translation buffer is one-way associative and has a capacity of 512 paging entries. Besides translation information, it contains access protection data, an aid in creating a secure operating environment. Refilling the buffer is done from page tables in memory.

Economical solution

Although the translation buffer is located in the M Box, it is maintained by microcode running in the E Box. This provides an economical solution to the complicated task of keeping track of streams of references from the three ports. Each port can have two references in progress because accessing the data cache and tag store overlaps with accessing the translation buffer. The data, addresses and control information for these operations are carefully queued up, with handshakes to allow subsystems to proceed as

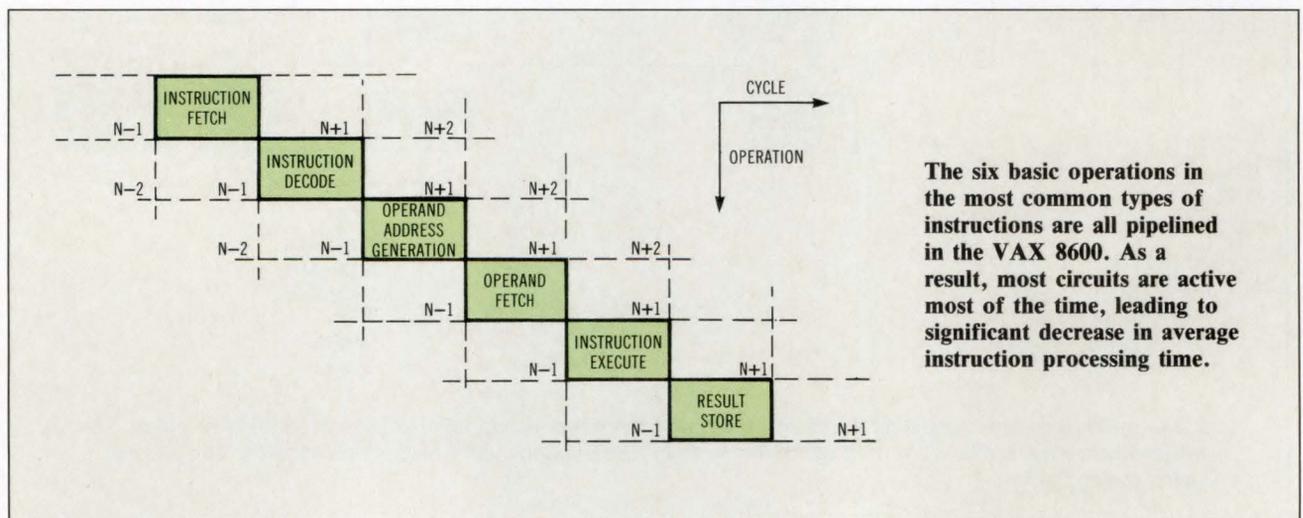
far as possible while waiting for references to finish. Memory exceptions encountered while prefetching instructions or operands are held off until the data is actually needed by the execution unit. That unit then deals with the problem using memory references that bypass the normal queue, leaving the queue intact for restarting later. The result is a virtual memory system that allows a reference to complete every cycle. With three subsystems making independent references, this high bandwidth can be utilized well.

F Box intercepts floating point

For scientific and technical applications the VAX 8600 has a floating point accelerator, the F Box, that operates in parallel with the E Box. The F Box receives operands over the operand bus from the I Box, and delivers results over the write bus for storage in general-purpose registers and memory. It performs floating point operations in all four of the VAX formats, F, D, G, and H (F numbers have 32 bits, D and G have 64, H have 128). It also does integer multiplications. Usually the work involved in these instructions is split between the F and E boxes: the F Box does the arithmetic operations, while the E Box accesses memory for reading and writing operands, deals with exceptions, handles counters, and takes care of other chores.

The F Box consists of specialized hardware (almost exclusively gate arrays) for floating point operations. Hence these operations take far fewer cycles. Beyond that, the F Box cycles twice as fast as the other subsystems. The data path is 32 bits, and multiprecision operations are pipelined. The F Box also has its own copy of the general-purpose registers, allowing the I Box to send both operands at the same time—one over the operand bus and one as an address for the general-purpose register RAM.

Much of the challenge in designing the F Box was to keep it compact enough to minimize interconnect



The six basic operations in the most common types of instructions are all pipelined in the VAX 8600. As a result, most circuits are active most of the time, leading to significant decrease in average instruction processing time.

delays. Of its two modules, one contains the logic for floating point addition, subtraction, and division, while the other does floating and integer multiplication. Both modules are microprogrammed, and both have their own microsequencer and control store. Moreover, the microcode is distributed among the various chips so a command follows the data for several cycles and is decoded repeatedly as the floating point operation is executed. This organization allows normal operations to finish in a minimum number of cycles while microcode detects and deals with unusual conditions.

The multiplier uses column reduction and Booth encoding, together with a three-input adder, to produce a 40-bit partial product every half cycle. The adder combines the operations of unpacking and aligning in a single shift, making it possible to produce an F-format sum in only two cycles.

The 8600 can have two Synchronous Backplane Interconnects; and the 8600's separate memory bus relieves these interconnects from involvement in processor-memory transfers.

The 8600 continues the VAX tradition of providing high speed, accurate floating point performance. All operations are accurate to one-half of the least significant bit. Floating point exceptions cause the instruction to be backed up to the beginning. Control is then given to an exception handler to scale the operands before resuming computation. Besides the basic operations, the 8600 provides special instructions for argument reduction and polynomial evaluation. These instructions carry extra precision and facilitate the high speed implementation of transcendental and other sophisticated mathematical functions in software.

Considering the I/O system

The basic structure of the system provides I/O over a Synchronous Backplane Interconnect interfaced to the M Box via the adapter bus. The structure offers complete compatibility with the myriad equipment currently available for the VAX-11/780 family. Moreover, the 8600 can have two Synchronous Backplane Interconnects and the 8600's separate memory bus relieves these interconnects from involvement in processor-memory transfers.

Thus, a significant increase in both the computational capacity and the I/O throughput of a current system occurs by replacing its processor with an

8600. The entire peripheral system can remain in place. A single Synchronous Backplane Interconnect can handle 13.3 Mbytes/s, all for I/O, and two have a combined capacity of 17.1 Mbytes. Some I/O device adapters connect directly to the Synchronous Backplane Interconnect; others connect through a Unibus or Massbus. The most recent I/O equipment is designed for use with the computer interconnect (bandwidth of 70 Mbits/s) and the Ethernet (bandwidth of 10 Mbits/s). The 8600 has the computer interconnect interface signals in its own backplane providing the necessary hardware for inclusion in a VAXcluster system.

System microcode

In addition to controlling the E Box data path, the E Box microcode supervises the operation of the whole processor. It initializes the system, tells the instruction buffer when to prefetch instructions or string data, starts and stops operand processing in the operand unit, maintains the address translations in the translation buffer, and orders the F Box to perform arithmetic operations. The microcode executes the full VAX instruction set, including such recent additions as G and H floating point, and interlocked queue instructions for multiprocessing. For backward compatibility, it also executes the PDP-11 instruction set.

Efforts to optimize the microcode and the E Box data path for executing the VAX instruction set have resulted in a relatively narrow microword of 84 bits (including 2 for parity). This microword, despite its narrowness, allows most high frequency instructions to complete in a single E Box cycle. Having immediate access to all 256 scratchpad locations makes it possible to store decimal strings and other data structures internally, and saves instruction cycles. Low frequency operations are implemented principally in microcode rather than in hardware, saving board space and reducing cost.

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

27mW and you're ready to roll.

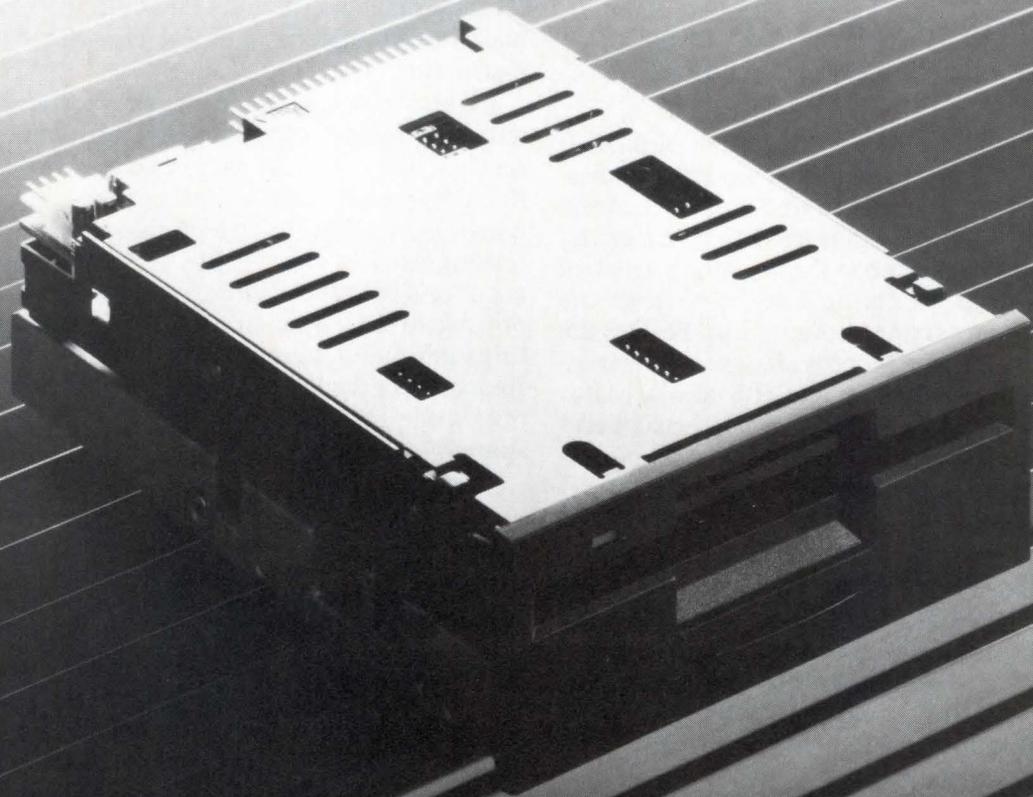
27 milliwatts. That's all it takes for standby mode on our FD-35 Series 3.5 inch floppy disk drives with the Power Saver option. And, in operation, they consume a mere 1.72 watts average at read/write, and 2.90 watts average at seek. These just happen to be the lowest numbers for any 3.5 inch drive.

This remarkable feat is made possible by the use of new custom LSI chips with CMOS technology for more efficient control over all the drive functions.

For added reliability and accuracy we shortened the gap between the band actuator and the head and included a more precise head position mechanism. This means stable reading and writing, even at high track densities.

And, as if all that weren't enough, we've thrown in special aluminum RF shielding to keep out noise pollution.

With the growing need for more portable systems, 27 milliwatts is a load anybody can carry.



Built To Fanatical Standards. **TEAC**

INSTRUMENTATION AND COMPUTER PRODUCTS DIVISION, 7733 TELEGRAPH ROAD, MONTEBELLO, CA 90640
East (617) 475-7311 South (214) 221-8714 Midwest (312) 351-9124 Rocky Mountain (602) 242-4025 (303) 337-6329 (801) 532-2111
Northwest (408) 727-1427 Southern California (213) 726-0303

CIRCLE 70

DEBUGGER ALLEVIATES REALTIME PROGRAMMING COMPLICATIONS

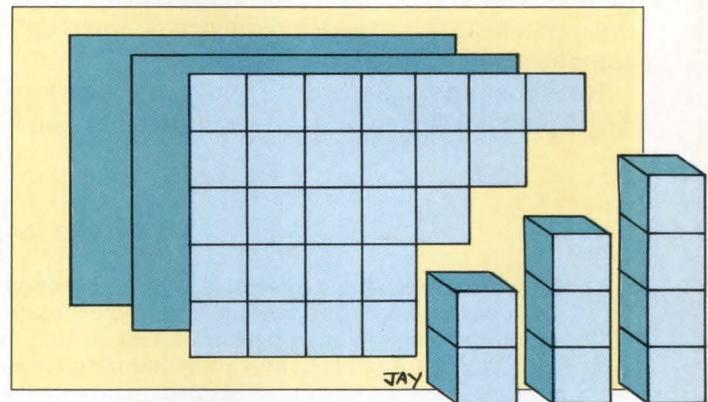
Realtime systems using microprocessors present special problems in debugging. A debugger specifically tailored to deal with concurrency can provide more reliable code for systems that must deal with the real world.

by Dennis M. Jones

Software debugging is rarely regarded as one of the creative aspects of programming, but everybody has to do it. At best, the task of finding flaws in one's own logic is only somewhat mitigated by having the right tools. Traditionally, these tools have included board-specific debuggers and emulators or in-circuit emulation systems—often expensive and relatively large scale hardware products. It has also been true that most programs debugged on such equipment have not been realtime, multitasking system applications but rather programs with single-path (or single-thread) execution.

More and more microprocessors, however, are being used in realtime applications, such as medical and avionics instrumentation, where software responsiveness must be fast and accurate enough to handle unpredictable conditions. These event-driven systems typically run a deceptively simple loop in which they wait for an event to occur, react appropriately, and then wait again. This event is usually a software or hardware interrupt.

The difficulties in programming in such life-endangering, time-critical applications frequently arise due to the random nature of interrupts that can occur in real-life, worst-case situations. Programmers must first be able to imagine all the possible complica-



tions. Then they must provide software that can process those events within the realtime constraints defined by the application. Because of the complexities that are involved in handling random interrupts, such applications typically use multitasking programming techniques.

Making programs manageable

Multitasking programming breaks a single-thread, monolithic program into manageable units of code called "tasks." A task is a logically complete execution path through user code that demands the use of system resources. A task also represents system activity that can logically proceed in parallel with other activities. This ability is critical because parallelism (or concurrency) is essential where tasking time must often be overlapped or interleaved on a single CPU to answer to the stringent constraints of realtime applications. Because realtime multitasking programs are unique in design, they require a special

Dennis M. Jones is a senior development engineer at Hunter & Ready, Inc (Palo Alto, Calif). He holds a BS in physics from California State University at Hayward.

functional ability in their debugging tools. To debug such programs successfully, the programmer needs to know the states of the different tasks, and the debugger needs to have minimal or no interference in the runtime environment.

In any debugging situation, programmers want to control the program's execution, suspending it at predefined breakpoints. Such control allows memory locations and registers to be examined, and corrected, if necessary. In single-path program debugging, the correct register contents can be deduced at any given moment of code execution because random, external events do not alter the execution sequence. Bugs are isolated between blocks of known good code. Then, using breakpoints and display commands, memory and register contents are examined in order to pinpoint the problem.

Debugging methods

With single-path programs, the advantage lies in always being able to re-create the sequence of events exactly. Methodical and systematic persistence solves the problem. This is not necessarily the case with realtime systems. The random sequence of interrupts complicates the situation extensively.

Realtime multitasking debugging, in contrast to single-path debugging, needs to be carried out simul-

taneously with the application code execution in a realtime, event-simulated environment. For example, to examine a faulty communication link between tasks, stopping the program just before the event to examine registers (as in single-path execution) does not work, because the communication is never received at all. Also, events and data in realtime systems occur randomly; for this reason, programmers cannot be sure that the re-creation of the exact sequence of events will correctly interleave into real world execution. Even if individual tasks are debugged, the multitasking program as a whole may not be. When debugging realtime systems, therefore, the debugger and the application code must run simultaneously in an event-simulated environment, with minimum interference from the debugger.

To fulfill the latter requirement, the debugger must use its own computational workspace in memory and not alter values in the application code's exception handling table. Equally important is the ability to view the state of the tasking environment as a whole (ie, the states of all the individual tasks and of the operating system as well). Also needed are ample tools for examining the specific data structures of all the tasks and system states.

Tracer, Hunter & Ready's realtime debugger, was designed with these realtime multitasking constraints

Taking a look inside VRTX

The Versatile Realtime Executive is a silicon-based, multitasking operating system kernel. It occupies about 4 Kbytes and can perform a task switch in about 100 μ s. It is especially intended for custom realtime designs using embedded microprocessors. For such realtime devices as laboratory instruments, industrial controllers, robots, and avionics, VRTX functions as a tool to manage software concurrency. It provides a set of system calls that gives access to basic multitasking services.

With task management, system calls exist for creating and deleting tasks, as well as suspending and resuming the tasks. VRTX automatically handles task scheduling according to a preemptive, priority-based system. Each task is assigned a priority, and the highest priority ready task always runs. System calls also exist to alter a task's priority, to obtain the status of a task, or to freeze and resume task rescheduling. With VRTX, users can dynamically allocate and release storage in fixed-size blocks. (This system call is known as dynamic memory allocation.) The free pool of memory can be dynamically subdivided into partitions, each with a different block size. Partitions of fixed-size blocks permit VRTX to simultaneously minimize waste, fragmentation, and system overhead.

Using intertask communication and synchronization, tasks can communicate with each other by ex-

changing pointer-length messages via mailboxes and queues. System calls exist for sending and receiving messages from mailboxes and queues. (Mailboxes are simply pointer-length locations; queues are system-managed entities that hold multiple messages.) Queues can be created dynamically. Both mailboxes and queues can be used for synchronization, mutual exclusion, and resource locking.

In addition, VRTX provides integrated support for a realtime clock and a single-channel, console-like, character I/O device. Provided on a device-independent basis, the user must supply a small (15 to 25 lines), device-dependent interrupt service routine for each device. VRTX manages the logical operations to give the user a full repertoire of clock management and character I/O commands (eg, time-slicing, task delays, and reading/writing characters).

VRTX is designed to be extended. Since VRTX does not "commandeer" the interrupt vector table, users are free to add their own interrupt service routines or system call handlers and thereby customize and extend their system software.

For realtime designs that require file I/O capabilities, other software components are available, including IOX (the I/O Executive) and FMXs (the File Management Extensions). These components function like operating system building blocks on top of the VRTX kernel.

PLUG US IN.

A built-in RS-232C serial interface board lets NCR standalone 40 column printers work with any system configuration.

Just write the proper software program and plug in the standard 25-pin female "D" connector. Without needing to reconfigure your present system. Printer control codes are fed over the RS-232C lines.

And even in the roughest work environment, NCR standalone receipt printers average only a single service call a year. But you can count on years of loyal service.

Other features you can count on include:

- 40 column lines.
- Print head life of 150 million characters.
- Cash drawer activation.

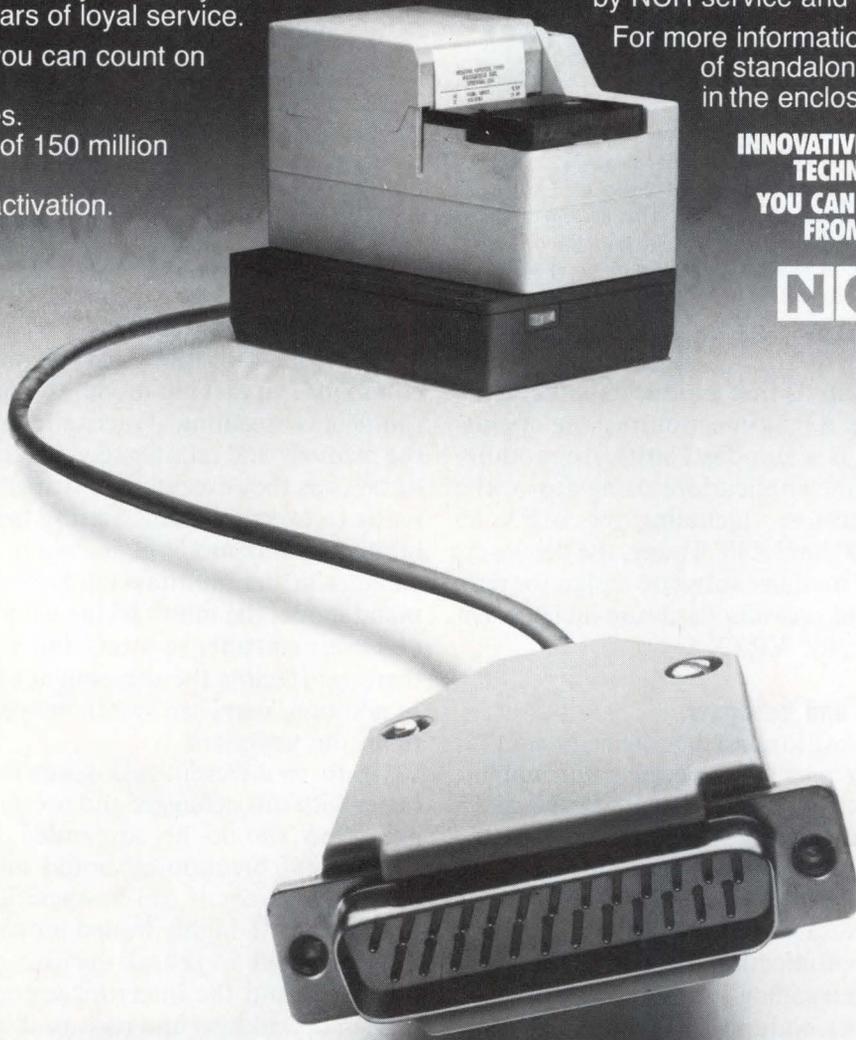
- Easy to configure; adapts to any environment.
- Easy to maintain; needs no routine lubrication.
- Extremely cost-effective; exceptional reliability reduces life cycle costs.
- Ideal for system flexibility; available in a variety of slip, receipt or journal printer configurations.

All NCR standalone printers were originally developed for NCR by NCR. Built in America, they meet our established quality standards of reliability. And are now available to you, backed by NCR service and support.

For more information on NCR's family of standalone printers, just send in the enclosed coupon.

**INNOVATIVE COMPUTER
TECHNOLOGY.
YOU CAN EXPECT IT
FROM NCR.**

NCR



Yes, I'd like more information on NCR's family of standalone printers.

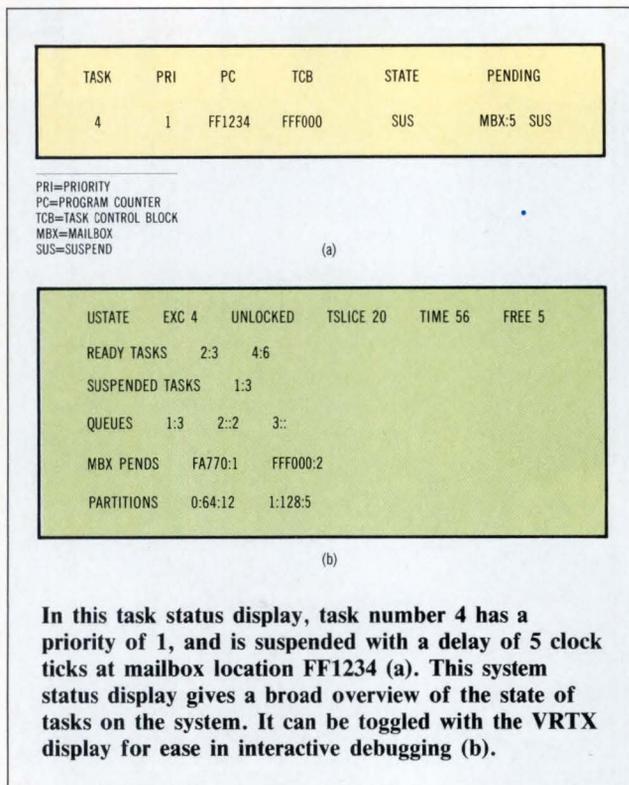
I prefer: a sales call a phone call

Our requirements are: up to 49 printers/year 50-500 printers/year over 500 printers/year

NAME _____ TITLE _____ COMPANY _____ ADDRESS _____

CITY _____ STATE _____ ZIP _____ PHONE _____ BEST TIME TO CALL _____

Please Mail To: Technology Marketing Division, NCR Corporation, Dayton, Ohio 45479 (513) 445-7443



in mind. The company's first product, the Versatile Realtime Executive, a realtime multitasking operating system kernel, is a standard software module. It is suited to realtime applications using any of the popular microprocessors, including the iAPX 86 family, 68000, Z8000, and Z80. Tracer, the debugger, follows this path of modular software design for realtime applications and provides hardware-independent debugging support for VRTX.

Versatile software and debugger

Hardware and position independence mean that VRTX executes in any CPU-based environment, whether custom-designed or bought off the shelf. Of course, users must supply some hardware-specific code—eg, short device initialization and interrupt service routines of 15 to 25 lines. The basic package consists of calls necessary to handle task and memory management, communication between tasks, and any other customized calling "extensions." Position independence allows code to be located anywhere in system ROM or RAM address space.

Tracer is basically the same. Versions of Tracer are microprocessor specific for the 68000 and the 8086, but can be loaded on any board, wherever memory can accommodate it. The same terminal can display application code output and Tracer output. Tracer allows toggling back and forth between the two. This multitasking debugger runs concurrently with VRTX and the application code. Yet, it never

interferes with the code's runtime environment because it is not a task running under VRTX's control. Tracer is accessed via a software trap. For this reason, it does not need to be linked to user code at compile or assembly time. Like VRTX, it can also include user-customized features.

Tracer is an interactive program. The interface capabilities include line edit functions, repeated issue of preceding command, multiple commands per line, and output control. The debugger achieves hardware independence by the simple interrupt service routine interface (identical to VRTX's) and by allowing the option of user-selected function keys for the terminal handler commands.

Tracer offers programmers both disassembler and "online" help commands as well as all the traditional debugging commands. In conjunction with VRTX, it also allows the in-depth display of system status even while the multitasking environment is running. In the following discussion of specific Tracer commands, the emphasis will be on what distinguishes this debugger from other debuggers that might possess similar or identical commands.

Operation mode advantages

Tracer has two modes of operation: tasking and command. In tasking mode, the multitasking environment is executing. Tracer allows users to examine the memory and register sets and to watch the status of tasks as they execute, are suspended, and become ready to execute again. Setting breakpoints is prohibited in this mode as it results in unpredictable changes in the multitasking environment. In command mode, the multitasking environment is frozen and users can alter registers and memory locations, thereby affecting the subsequent execution of tasks. In addition, users can switch between the two modes from the keyboard.

Up to 16 different task-specific breakpoints can be set with this debugger and users designate whether execution should be suspended at that point or whether the breakpoint should simply be reported. An iteration count can be specified to track down problems with highly nested loops. Breakpoints set in system code can check the user-written extensions to VRTX and the interrupt service routines.

Memory address and register display are relatively standard. Memory and registers are set by specifying a value and giving the memory address or register name, respectively. In Tracer, the option remains to set and display registers for each individual task. Examining these registers gives a good indication of their state, independent of the currently executing task.

The most important concern of the multitasking programmer is knowing the state of a task at any given moment. For example, it is vital to know how



THE INCREDIBLE SHRINKING CORDLESS.

DURACELL® high-energy density batteries can improve your cordless design (...providing you call us in early).

Is the mighty chip enabling you to miniaturize your portable designs further or to go cordless for the first time? Yes?...Then you should be aware that Duracell battery technology is keeping pace with the chip. And if you call Duracell in early, before your design is locked in, we can apply that new technology in reducing the size and weight of your product even more, while improving performance and usage cycle. We can determine the optimum battery for you and help design the cavity and the circuitry. Duracell, manufacturer of the **World's Number One Premium Battery**, offers you primary cells and batteries in a full range of capacities, sizes and shapes, in all key systems. We also give you the most complete design and applications assistance you can get: ■ a new 90,000 square foot research and development center; ■ several separate engineering facilities; ■ product engineering specialists; ■ field engineering specialists throughout the country; ■ and all the resources of a subsidiary of a FORTUNE 50 Company...plus international availability under the DURACELL® brand. But to take full advantage of Duracell's resources and advanced technology, call us in early on your cordless design.

DURACELL INC.
OEM Technical Sales & Mktg.
 Berkshire Industrial Park
 Bethel, CT 06801

...call toll-free
1-800-431-2656
 (In CT, call 203-796-4000)

- Please send information on Duracell battery technology.
- Please have an Engineer call me.

My application is _____

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Telephone (____) _____



DURACELL

ZINC AIR • LITHIUM • ALKALINE • MERCURY • SILVER

Send Literature—Circle 72

Send Literature & have an Engineer call me—Circle 73

many tasks are active, how many are suspended, and if so, why. It is also necessary to know where the program counter of a given task is and how long a task has been suspended. These are only a few of the issues programmers should consider. The dt (display task) command addresses them.

If a display of Task 4 is requested, Tracer returns a message. The message conveys that Task 4, of priority 1, is presently suspended with a five clock tick delay at the mailbox located at FF1234. (A "tick" is a VRTX clock time interval, the basic unit of time for VRTX. Ticks are derived from an interrupt generated by a hardware timer.)

The "Pending" reading indicates conditions surrounding the suspension of the task, including whether or not the task is suspended on a mailbox or a queue. The "State" reading indicates that the task is currently suspended. This entry might also read "RDY," for ready but not executing, or "EXC" for executing. Mailboxes and queues are VRTX data structures; Tracer displays their status. A mailbox is an agreed-upon location in memory large enough to accommodate an address-sized message. Mailbox messages act as pointers to other data structures or to data, as well as for implementing semaphores and resource-locking. Queues can be thought of as a string of mailboxes that are linked

together. Message queues are fixed-length buffers where queued messages are handled in a first-in, first-out manner.

"Pending" indicates that Task 4 is suspended on a mailbox for a delay of five clock ticks. "SUS" under this heading also means that the task was explicitly suspended. In another case, the "Pending" reading might be different. For example, it might be "QUE:queue id number" if Task 4 were suspended at a queue waiting for a message; "GTC," if it were suspended on an empty input buffer; "PTC," if suspended on full output buffer; or "DLY:" (some number of clock ticks), meaning delayed for the indicated number of clock cycles. Other task status information might also be reported with the dt command.

Logical queue display

While many debuggers merely display an undifferentiated list of memory locations when examining queue-like structures, the Tracer command, dq (display queue), always displays the head of the queue as the first message entry in its contents list. This queue display is always logical rather than physical. Displaying queue status also indicates whether tasks are pending on the queue, and if so, their "time-out" values. In VRTX, when a task is pended at a queue or mailbox, it can be governed by a "time-out" value (measured by a specific number of clock ticks), which determines the time that it remains suspended. If the time-out value elapses and no message is received, an error code is returned to the calling task. If no time-out value is given, the tasks are suspended indefinitely.

The dx (display mailbox) command lists all the tasks pending at mailboxes, their time-out values or indefinite suspension, and the address of the mailbox where they are pending. Mailboxes and queues are often the heart of a multitasking system because all synchronization, resource locking, and coordination occur through them.

VRTX includes two 64-byte buffer structures that manage a serial I/O channel. The di (display input buffer) and do (display output buffer) Tracer commands display the contents of these two buffers, as well as the number of data bytes. They also display the tasks that are pending to deposit or receive characters. Again, the display of the structures is logical rather than physical.

The big picture

The most comprehensive display of task information that can be requested is the display system status command, ds. It delivers a complete overview of the system's status in snapshot form. In multitasking systems, such an overview can often be the crux of debugging complex errors. A sample ds command

0	VRTX-ADDRESS	
2	BUFFER-ADDRESS	
4	BUFFER-LENGTH	
6	CHANNELS	I/O TOGGLE
8	TRACER INITIALIZATION ADDRESS	
A		
C	*INPUT-FILTER	
E		
10	*OUTPUT-FILTER	
12		
14	*SUSPEND-OUTPUT	
16		
18	*CANCEL-LINE	
1A		
1C	*BACKSPACE	
1E		
20	*NEW LINE	
22		
24	*REPRINT-COMMAND	
26		
28	*REISSUE-COMMAND	
2A		

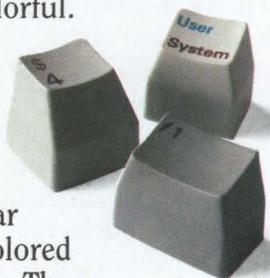
*INDICATES OPTIONAL PARAMETER: SET TO -1 IF OMITTED

To install Tracer for the Intel 86 family, users fill in the values in the table and write a short initialization routine.

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[®]

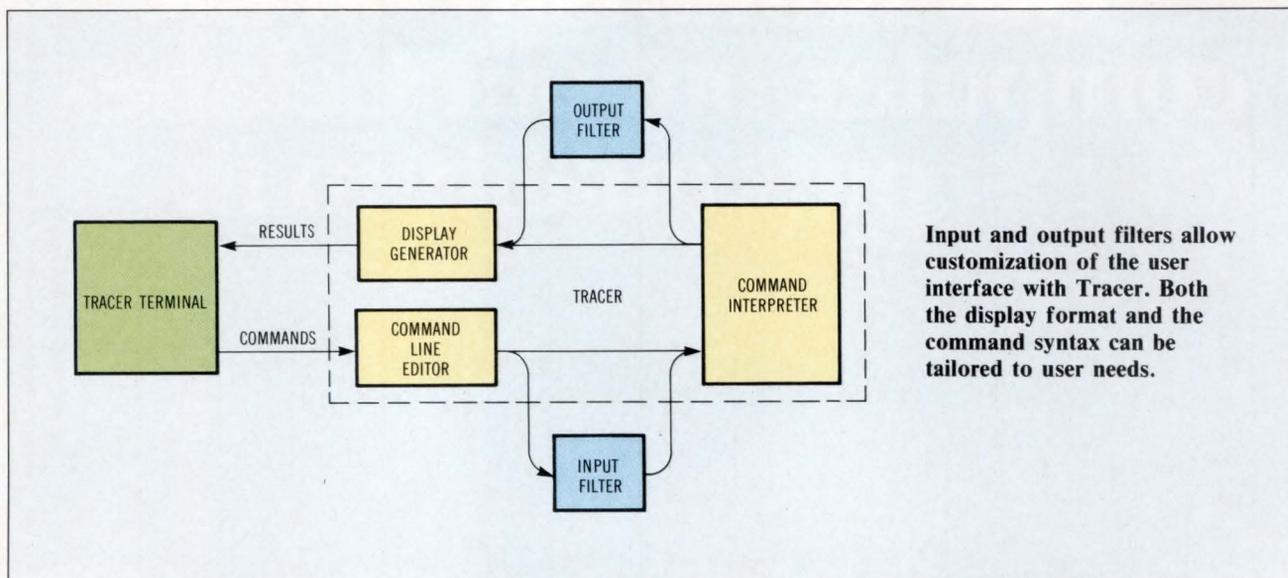
2000-2's internal lubricant system facilitates mold release without mold plate-out—maximizing productivity.

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.



Other Celanese quality products include Celcon[®] Acetal Copolymer and Celanese[®] Nylon.



Input and output filters allow customization of the user interface with Tracer. Both the display format and the command syntax can be tailored to user needs.

display provides the following information about the internal operation of VRTX. Line 1 begins by showing whether the system is in user or system state—ie, if the user application code or the operating system's code is executing when the command is entered.

When a task is created with VRTX, it has a unique identification number ranging from 1 to 255. This number allows tasks to be readied, suspended, or deleted on a selective basis. The second entry on line 1 indicates this task ID number for the currently executing task. VRTX tasks are arranged according to priority from 1 to 255 and realtime applications usually require that the highest priority task be running. Task rescheduling in VRTX can occur after executing a VRTX command or an interrupt, during which a higher priority task may be ready. In some instances, however, when programmers want to preempt rescheduling, they can do so. The third entry on the first line indicates whether scheduling is "locked" in place or "unlocked," for rescheduling.

Preventing CPU bottlenecks

Time-slicing is a VRTX task function that allows equal-priority tasks to run in turn for short periods in round-robin fashion. This function prevents a compute-bound task from monopolizing the CPU at the expense of other equal-priority tasks. Looking at the first figure, the fourth entry on the first line reminds the programmer how many clock ticks define the time-slicing intervals. The fifth entry indicates the system time. (The system time is a 32-bit quantity that is incremented at each clock "tick.")

Task control blocks are VRTX data structures in system memory that retain status information on task states, register contents, stack instruction pointer values, etc. One task control block is associated with each active task in the system. The last

entry in the first line indicates the number of task control blocks that have not yet been associated with an active task.

The task-specific data (line 2) is displayed below this general status information. "Ready tasks" indicates the number of tasks, and the priority of each task running in the system. In this case, there are three ready tasks of priority 2 and six of priority 4. "Suspended tasks" (line 3) indicates how many tasks, and of what priority, are suspended. In this instance, there are three suspended tasks of priority 1.

"Queues" (line 4) shows the number of queues, how many messages are located in them, and whether there are any tasks suspended on them waiting for messages. In the example, there are three queues numbered 1, 2, and 3. In queue 1, there are three messages; there are two suspended tasks on queue 2; and queue 3 is empty, containing no messages and having no tasks suspended on it. The fifth line, "Mbx Pends," gives the number and addresses of mailboxes, and indicates whether or not tasks are suspended, waiting for a message. Here the mailbox located at memory address FA770 has one task suspended on it; the mailbox located at FFF000 has two.

"Partitions" (last line) indicates how many memory partitions exist, their block size, and the number of unused blocks. The user RAM managed by VRTX consists of a number of partitions, or chunks of memory that may not be physically contiguous. Each partition is subdivided into user-defined, fixed-size memory blocks that can be dynamically allocated. This method of partitioning and using fixed-size blocks prevents crashes caused by memory fragmentation. It ensures predictable memory behavior and is essential to the correct design of realtime systems. There are two memory partitions in the example. Partition 0 has a block size of 64 bytes

with 12 free blocks; partition 1 has a block size of 128 bytes with 5 free blocks.

Installation and configuration

Installing Tracer is very similar to installing VRTX. User code accesses VRTX using traps or software interrupts, thereby eliminating the need to link application source-code references to VRTX entry points. Tracer is also accessed this way. To install Tracer, the proper exception vectors of the processor must be loaded with vectors that trap to Tracer. In fact, for most efficient use of vectors in the table, Tracer and VRTX can use the same trap. A short initialization routine (4 or 5 lines of code) must be written for Tracer and a configuration table of 10 to 16 values filled in. The configuration table and its required values are diagramed.

These values are easily defined and indicate the target-specific address for the main entry point of VRTX, the starting point of Tracer's RAM buffer, and the starting address of any user-supplied extensions. Other entries designate whether or not VRTX and Tracer share the same I/O channel, and the user-supplied characters that allow definition of specific terminal characteristics. These few values specify the locations in the target environment necessary for Tracer's installation. Because it is not linked to user code, the debugger can go anywhere in RAM or ROM and its computational memory can be located anywhere in RAM.

Tracer's design allows users to customize it. For instance, users can define the interactive control characters in the configuration table. Tracer and VRTX can share a terminal, or a terminal can be dedicated to Tracer. Also, additional commands can be added by using the input filter to define an entirely new syntax for all commands or to provide a macro facility to translate high level commands into Tracer's syntax. Likewise, the output filter can create customized display formats. The logical arrangement of the two filters in Tracer is illustrated. These kinds of user "hooks" are available in VRTX as well.

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

If it Wasn't Important, We Wouldn't Ask You to Do It



Every year, *Computer Design's* 95,000 subscribers fill in the most detailed qualification form required by any magazine in the world. Why? Because this form tells us who you are, where you are, what your functions are, what your company does and what you do. It tells us the kinds of projects you are working on, the kind of products you deal with, and where they are used. It becomes the starting point in planning our editorial program.

The information you provide helps us to select the kinds of features, special reports and surveys that will be of immediate, practical use to designers of computer based systems.

The other side of the coin is the value advertisers place on circulation as a major criterion in media selection. As everyone knows, the more advertising pages we carry, the more editorial pages we can provide to you without cost. The bottom line is a better magazine that can speak your special kind of technical language, page after page, issue after issue.

COMPUTER DESIGN

PennWell Publishing Company,
Advanced Technology Group



SOPHIA SYSTEMS STANDS ALONE.

Now Sophia Systems makes it practical to use the power of in-circuit emulators anytime, anywhere. Our family of SA700 portable emulators lets you develop in the lab or troubleshoot in the field complex micro systems.

COMBINED DEVELOPMENT SYSTEM AND IN-CIRCUIT EMULATOR

- Z80 and 8085 models are combined development systems and in-circuit emulators. All the development tools, from macro-assembler and linker, to integrated real-time ICE probe and PROM programmer, are standard.
- 68000 systems include a 68K mini-assembler for software development and over 1/2MB of emulation RAM for real-time debugging.
- 8086/87/88 models are floppy disk compatible with Intel® MDS Series III and CP/M-86™ systems.

Serial links are available to easily connect any SA700 to VAX®, IBM®

PC, or HP®64000 hosts and download to our floppy disk.

DEBUG WITHOUT A HOST

You can also debug and handle hardware/software integration without a host or add-on peripherals, saving time and money. Sophia's SA700 systems give you effective standalone micro development, debugging and test for as little as \$7,000. Nothing else compares.

A WORLDWIDE LEADER

An international market leader of development systems and emulators, Sophia Systems has delivered more than 4,000 SA700 systems worldwide.

CALL TODAY FOR A DEMONSTRATION

In the rush to get your designs to market, Sophia Systems stands alone. Call or write today for a demonstration and further information.

Sophia Computer Systems, Inc.
3337 Kifer Road
Santa Clara, CA 95051

COMPARE THESE FEATURES:

- Software development (assembler) and debugging without host connection
- Portable, rugged package with integrated CRT, floppy drive, ASCII keyboard, PROM programmer, and ICE™ Probe
- Batch DEBUG programs created and executed on the emulator
- Real-time trace, timing and emulation
- Full symbolic support —breakpoints, trace, [dis]assembly, display
- Single-step/slow-mode execution
- Proprietary & CP/M® operating systems



Call toll-free 1-800-824-9294
(In Calif., 1-800-824-6706)

Sophia
systems

**MICROPROCESSOR DEVELOPMENT
SYSTEMS & IN-CIRCUIT EMULATORS**

International Office: JAPAN, SOPHIA SYSTEMS CO., LTD, NS Bldg., 2-4-1, Nishishinjuku, Shinjuku-ku, Tokyo 160, 03-348-7000.
U.S. Regional Offices: SOPHIA COMPUTER SYSTEMS, INC., 22 Executive Park-Suite 280, Irvine, CA 92715, (714) 261-8688;
Glenpointe Centre East, 300 Frank W. Burr Blvd., Teaneck, NJ 07666, (201) 836-5959.

©1985 Sophia Computer Systems, Inc. Sophia Systems is a trademark of Sophia Computer Systems, Inc.

Floating point unit executes to 80 bits of precision

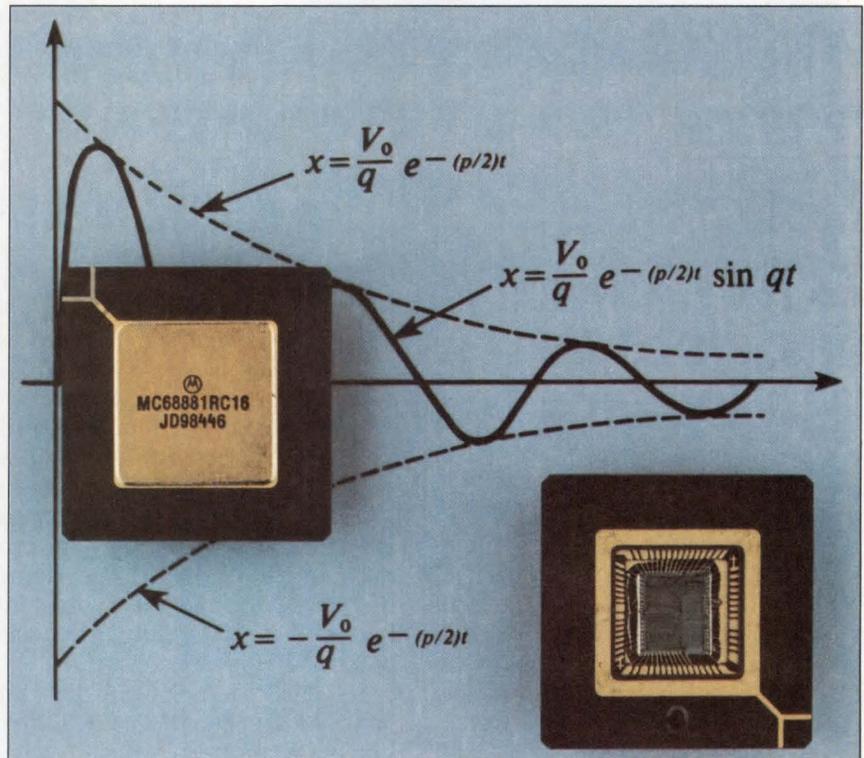
In an effort to make the 68020 microprocessor an all-around performer, Motorola's microprocessor products group has announced a full complement of 32-bit peripheral chips. Coupling the 68020 microprocessor with the upcoming 68461 32-bit memory management controller (and now with the 32-bit 68881 floating point coprocessor) enhances system operation. The system designer can apply the power of full 32-bit calculations in such applications as computer aided design, scientific analysis, robotics, and numerical control. Complex functions are handled on a single chip, thus reducing the cost and size of desktop systems.

The 68881 performs all the mathematical operations dictated by the IEEE Floating Point Specification (P754 draft 10.0) at twice the speed of any other single-chip device, according to Motorola. All calculations are done in hardware to 80 bits of precision. The chip can be used as either a tightly coupled coprocessor with the 68020 or as a memory mapped peripheral with other microprocessors in the family (eg, the 68000, 68010, and 68012).

Besides calculating basic math functions, this chip performs transcendental and nontranscendental functions. These operations include exponential, hyperbolic, logarithmic and trigonometric functions, as well as root values.

Used with the 68020, the coprocessor interface is transparent to the system programmer, since coprocessor instructions are written as a part of the main program instruction stream. The CPU can be freed for other tasks because the floating point processor can run concurrently with the main processor. The 68881 architecture includes a 67-bit arithmetic unit, a barrel shifter, and eight 80-bit general-purpose registers. Standard clock frequencies are 12.5 and 16.67 MHz.

Motorola engineers designed the coprocessor interface as an integral part of the 68020 and the 68881. Both microprocessors share tasks, with the 68020 providing services for the



coprocessor at the latter's request. Once the main processor has provided the services, it is free to continue its own processing. The choice between concurrency and nonconcurrency is left to the coprocessor on an instruction-by-instruction basis.

Floating point functions can be added as memory mapped peripherals to other 68000 MPUs that include the reduced bus 68008, the 16/32-bit 68000, the virtual 68010, and the expanded virtual 68012. This is accomplished by providing instruction sequences that emulate the protocol of the coprocessor interface.

Design of the 68881 is optimized for area efficiency using the 2-micron HCMOS process, placing close to 155,000 transistors on a 270- x 330-mil chip. The 68-lead pin grid array package dissipates 1.0 W maximum. The chip's designers were able to fit the large circuit into the relatively small area by redesigning the control logic portion and by using experience gained with the 68020. They were able to condense the control logic by embedding the majority of its functions into the "nanoROM"

and leaving only the necessary irregular line patterns in the control logic section.

The 68461 MMU chip will be available in the second quarter while a gate array version is ready now. The MMU array is customized to support the 68010 or 68020 in a virtual memory environment. Its primary function is to translate a logical address, provided by the main processor, to a physical memory address. The device uses tree structured translation tables. It features a 1-Kbyte page size, and provides the necessary control for an external descriptor cache. It can also support page sharing between multiple physical bus masters. Translation of either 24- or 32-bit addresses is handled by the 68461 controller, which is packaged in a 149-lead pin grid array.

Samples of the 12.5-MHz version of the 68881 are available this month with limited production beginning in July. The sample price is \$375. **Motorola Inc, Microprocessor Products Div**, 3501 Ed Bluestein Blvd, Austin, TX 78721.

Circle 260

—N.M.

Enhanced array processor peaks at 38 MFLOPS



An enhanced 64-bit computer, the FPS-264, uses custom ECL gate arrays that allow top processing speed of 38 MFLOPS. It provides variable interleave for fast memory access and extensive diagnostics.

To efficiently solve large scale computational problems, the array

processor incorporates several features. These include parallel processing and pipelining. The FPS-264 uses eight pipelined data paths, connected by seven high-speed protected buses. The processor provides 64-bit accuracy as well as error-correcting memory, internal diagnostic buses, and a diagnostic microprocessor. System memory is 512,000 words, expandable to 4.5 million words (36 Mbytes). Buses allow 24-bit addressing of memory words. Programmed in Fortran 77, the processor includes a Fortran 77 compiler, overlay linker, and the System Job Executive—Floating Point Systems' proprietary operating system.

All application software that runs on the FPS-164 (an 11-MFLOPS version of the machine), will also run on the FPS-264. Many third-party pro-

grams for scientific and engineering applications, including ANMSYS and MSC/Nastran for structural analysis and several versions of Spice, have been adapted by their vendors for both the FPS-164 and 264. The company can also assist users in converting existing in-house application programs.

In addition to the array processor, a hard disk subsystem has been introduced to support the FPS-264. This unit expands mass storage to 16 Gbytes. Known as the FD64, the unit connects to the processor through an intelligent controller and consists of a Winchester, a controller in the disk drive cabinet, and an adapter in the computer's I/O controller. Each controller can handle up to four disk drives. Up to six controllers can attach to the processor.

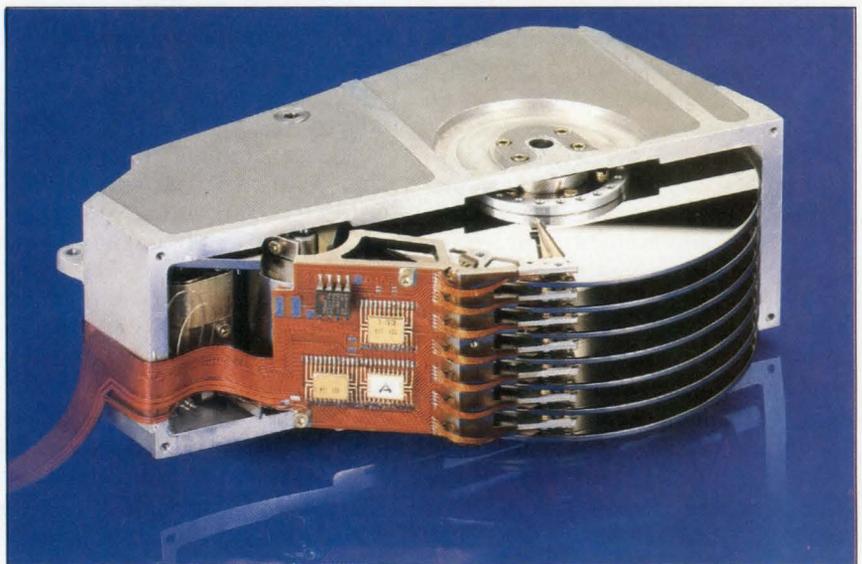
Prices for configurations of the FPS-264 start at \$640,000, with shipments starting in the third quarter. **Floating Point Systems Inc.**, 3601 SW Murray Blvd, Beaverton, OR 97005. Circle 261 —M.B.

MiniWinchester supplies top-level performance

Megafile series 5¼-in. Winchester disk drives pack high performance in a small package. Drives in the series access up to 306 Mbytes of data in a 25-ms average access time. Transfer rates of 10 Mbits/s are supplied by the drive's ESDI or ANSI 1226 device-level interface. System-level interface is via SCSI.

Packing up to seven platters into the 5.75- x 8- x 3.25-in. (14.605 x 20.32- x 8.255-cm) footprint, the drive uses thin-film heads flying at 8 μ m and carbon-overcoated plated high resolution media to achieve linear densities of 19,077 bits/in. on tracks spaced at 1207 tracks/in. Recorded data is encoded in 2,7 run-length limited code, and is decoded to nonreturn to zero before transmission to the controller.

To maintain precise positioning on all surfaces at the high speeds needed to obtain a 25-ms average access



time, the design uses both a dedicated servo surface on the top platter and servo data embedded on each surface of the other disks.

The rotary voice coil actuator is a statically balanced swing arm mechanism, supported on two preloaded bearings. The one-piece arm allows

Put your blue box out to pasture.

Genesis Has Better Choices for Superior 8086 Designs.

Genesis development tools turn your IBM PC or compatible computer into the most flexible, comprehensive development system made for 8086-family designs. Genesis offers the widest choice of proven languages, superior debugging tools, and even in-circuit emulators. It all starts with ACCESS, the superior Genesis operating environment that allows Intel's 8086 development software to run under MS-DOS. And not only do we sell ACCESS and all the Intel languages for 8086 and 80186 designs — we also offer an ACCESS system for 8-bit designs as well.

Genesis Is Faster.

You get performance as well as flexibility with Genesis. For example, in our tests the ASM 86 assembler running under ACCESS on an IBM PC AT runs 4.7 times as fast as the assembler on a hard-disk based Series III and an incredible eleven times faster than a

floppy-disk based Intel system. The ASM 86 assembler we tested under ACCESS even ran 1.53 times faster than a VAX 11/782 running competing cross-software.

Genesis Is Easier to Use.

Of course, Genesis offers more than speed of execution. The GeneScope family of symbolic debuggers and in-circuit emulators share the most powerful and easiest to use human interface available today. GeneScope offers high-level macro facilities, an on-line symbolic assembler, alphanumeric and graphic support, and many more features that make Genesis tools a better solution than other systems.

Genesis Costs Less.

With Genesis you get all the flexibility of a personal computer and all the power of expensive dedicated development systems. You can get your products to market faster, at a fraction of the cost.

Is there really any reason to keep using obsolete, expensive equipment — when you can have Genesis instead? **Quit chewing your cud** — give us a call and find out how Genesis tools can make your life easier.



Genesis Has a Better Solution.



Genesis™
Microsystems

Genesis Microsystems Corporation
196 Castro Street
Mountain View, California 94041
Phone: (415) 964-9001
Telex: 4998093 GENMS UI

CIRCLE 76



SYSTEM COMPONENTS

replacement of single heads by removal of just one screw.

Problems of spindle tilt are eliminated by a rigid mechanical assembly design in combination with twin bearings that are mounted at both the top and bottom of the spindle and the positioner. Embedded in the spindle hub is a brushless dc spindle motor that rotates the spindle at 3600 rpm. Hall effect switches in the drive motor assembly maintain high speed accuracy. The motor design eliminates problems associated with static

discharge. A mechanical brake stops the motor when power is turned off.

An unusual monocoque design distinguishes the assembly. One half of the shell contains heads and platters, spindle, and positioner; the other half seals the head/disk assembly hermetically. This design allows easy maintenance, as the entire inner assembly is visible when one half of the shell is removed. Even the heads can be replaced without major disassembly.

Use of custom VLSI and hybrid circuits lets all electronics reside on one

single printed circuit board. This provides interface, control, read/write, positioning, and power electronics.

Models 1100, 1200, and 1300 have formatted capacities of 84.6 Mbytes, 169.3 Mbytes, and 254 Mbytes, respectively. Prices are \$2450 for Model 100; \$3300 for Model 1200; and \$3900 for Model 1300, each sold in OEM quantities of 1000. **Siemens Communications Systems, Inc, Memory Products Div**, 5655 Lindero Canyon Rd, Westlake Village, CA 91362.

Circle 262

—P.K.

Geometry engine powers design workstation

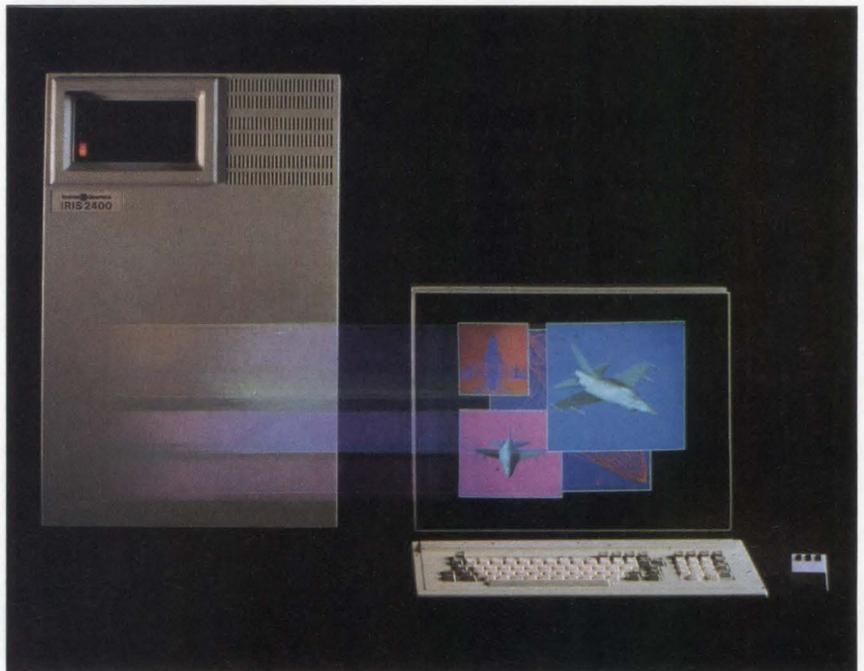
At the heart of Iris 2000 Unix-based workstations are custom VLSI "geometry engine" chips. These chips use a parallel pipelined architecture to calculate up to 100,000 graphics vertices per second.

The series comprises the Iris 2000, 2200, 2400, and 2500 workstations. The 2000 and 2200 are host-dependent workstations in a 10- or 20-slot chassis, respectively. The Iris 2400 is a standalone personal workstation capable of serving as a node in a distributed computing environment.

The 2500 is a rackmounted version of the 2400 with high capacity disk and tape storage. The Iris 2400 incorporates a 32-bit 68010 microprocessor running Unix System V with demand paged virtual memory. Basic units include 1.5 Mbytes of memory, expandable to 15.5 Mbytes as found in the high end Iris 2500. The 2400's disk storage is expandable from 72 Mbytes to the 880 Mbytes in the Iris 2500. A quarter-inch cartridge tape unit is standard. A floating point coprocessor is optional.

Ethernet capabilities are included with the workstation. XNS or IP/TCP protocols, direct IBM mainframe links, and the IEEE 488 bus are options.

Silicon Graphics claims the system can update the display 100 times faster than workstations from Sun and Apollo. Since the system response is significantly faster than conventional graphics workstations, operator productivity is greatly



enhanced. Silicon Graphics cites extensive studies conducted by IBM demonstrating that decreases in system response time improve operator productivity by a far greater factor than the percentage savings in response time would indicate.

The 2400 includes a 1024- x 1024-pixel, 19-in. color monitor. The system supports depth queuing, Gourard shading, 8-bit planes (expandable to 32-bit), and a 12-bit color map. An advanced window manager allows simultaneous updating of dynamic 3-D graphics images, the use of single or double-buffered display modes,

and use of a 16-bit Z-buffer. The system also includes a broadcast quality NTSC/RS-170A video output.

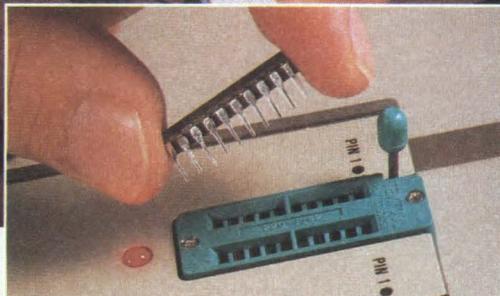
Among the Unix software available for the workstations are several packages for molecular modeling and mechanical CAE, and the EMACS text editor.

Iris workstations in this series are (in single quantities) \$27,500 for the 2000, \$49,500 for the 2400, and \$76,500 for the model 2500. **Silicon Graphics, Inc**, 630 Clyde Ct, Mountain View, CA 94043.

Circle 263

—W.E.S.

SPEND MORE ENERGY DESIGNING LOGIC.



AND LESS EFFORT PROGRAMMING PARTS.

The new 60A.

Data I/O's economical new 60A Logic Programmer is designed to do one thing only and do it very well: program logic.

Almost every programmable logic device is supported including PAL's and IFLs.

And it's user-upgradeable as new devices become available.

Logic pure and simple.

But what's really remarkable about the 60A is how easily it lets you program logic.

All you have to do is use the menu to select the kind of device you're programming, depress a single button to specify programming operation, and

press the start button. In less than a second your device is programmed. And selectable three-step functional testing ensures it is programmed correctly.

Add ABEL™ and design the way you think.

Taking advantage of programmable logic has never been easier, especially with ABEL, Data I/O's high-level design language.

Working with your computer, ABEL lets you express your designs in Boolean equations, truth tables, state diagrams or any combination... whatever works best.

Call toll-free 1-800-426-1045 for a demo or details.

To spend more energy designing logic and less effort programming devices, call Data I/O and ask about the 60A Logic Programmer. It's dedicated to one thing only: logic pure and simple. And like every Data I/O product, it's fully backed by our commitment to quality and our worldwide network of service representatives.

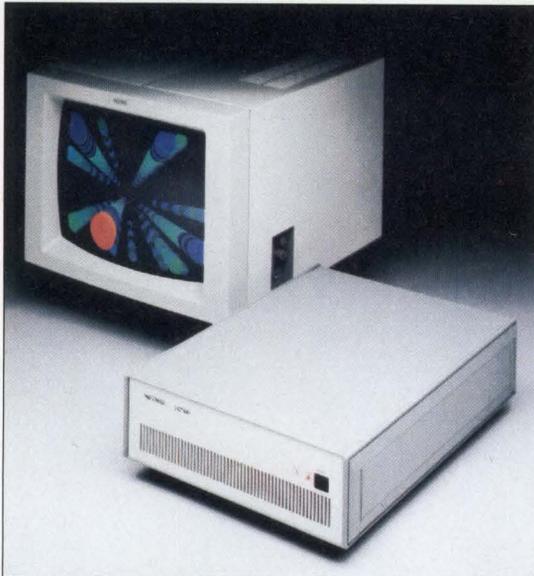


Data I/O® is a registered trademark of Data I/O Corporation, and PAL® of Monolithic Memories, Inc. ABEL™ is a trademark of Data I/O Corporation.

Data I/O Corporation, 10525 Willows Road N.E., P.O. Box 97046, Redmond, WA 98073-9746. For immediate action, contact us directly, **CALL TOLL FREE: 1-800-426-1045**. In Washington, Alaska and Hawaii, call 206-881-6444. Europe: Vondelstraat 50-52, 1054 GE, Amsterdam, The Netherlands. Tel: (20) 186855. Germany GmbH: Bahnhofstrasse 3, D-6453 Seligenstadt, West Germany, Tel: (6182) 3088. Japan: Ginza Orient Building 6-F, 8-9-13, Ginza Chuo-ku, Tokyo 104, Japan, Tel: (03) 574-0211.

DATA I/O

Graphics controllers tout large display-list memory



Multiprocessor graphics controllers in the $\Omega 2000$ series combine the intelligence of a frontend microprocessor, a large display-list memory, and a specialized transform processor. The result is a high resolution, high performance machine for design, imaging, and simulation.

Graphics controllers in this series execute a stream of high level instructions from a host computer to create a stored image in a pixel memory array. The stored image is continu-

ously displayed on a monitor, and the system provides local memory for graphics segments.

Central to $\Omega 2000$ operation is the Graphics Preprocessor. This system component interprets high level instructions and manages a hierarchical data base of stored display objects. It also supports a fast random access display-list memory and a transform processor. The Preprocessor controller is built around a 68000 family microprocessor running at a 12.5-MHz clock rate. Display-list

memory provides storage for graphics segments. Composed of an array of fast dynamic RAM that also buffers I/O data, the display list is accessible through two DMA controllers. A 256-Kbyte display-list memory is included in the standard $\Omega 2000$ configuration, with 1-Mbyte memory optionally available. The system's floating point transform processor applies translation, rotation, scaling, and projection transformations for mapping into display coordinates.

This graphics controller is marked by its display processor. Specifications for the display processor indicate that vector drawing speed reaches 1 million pixels/s for the $\Omega 2300$ model. Pixel memory measures 1024 x 1024 on the $\Omega 2300$ and $\Omega 2440$, and 1280 x 1024 on the $\Omega 2530$. The displayed resolution for the $\Omega 2530$ can reach 1280 x 1024 at a 60-Hz refresh rate, noninterlaced. The drawing processor of the machine exhibits a 220-ns cycle time on the $\Omega 2300$ and $\Omega 2440$, while the $\Omega 2530$ drawing processor achieves a 167-ns cycle. On the $\Omega 2530$, hardware pan is vertical by 16 pixels, horizontal by 40.

The $\Omega 2000$ instruction set and graphics library software provide a virtual device interface that supports abstractions of the GKS and ACM-SIGGRAPH Core standards. High resolution display monitors, host-compatible C-language, and Fortran software drivers are also available for $\Omega 2000$ series graphics controllers. The $\Omega 2300$ costs \$10,950; the $\Omega 2400$ costs \$13,950, and the $\Omega 2500$ is \$17,950. **Metheus Corp**, PO Box 1049, Hillsboro, OR 97123.

—J.V.

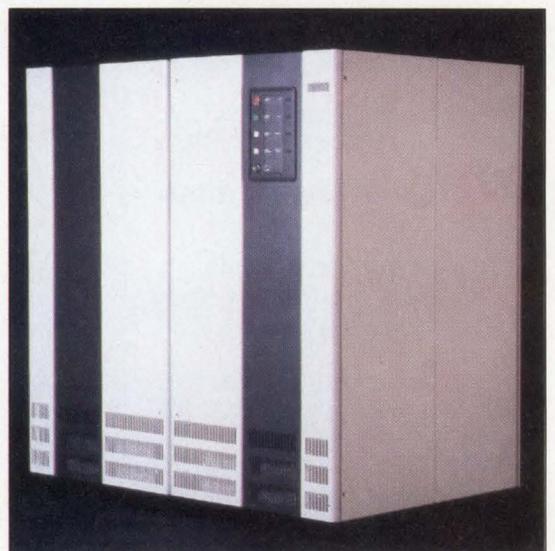
Circle 264

Supermini employs ECL and pipelining for 32-bit computations

A 32-bit superminicomputer uses ECL circuitry, a pipelined architecture, burst mode I/O, and 16 Mbytes of error correcting MOS memory to produce a top flight performer. The system, known as the 9955, simultaneously executes as many as 255 active processes and accommodates up to 254 directly connected terminals. The machine uses only five circuit boards.

Pipelined CPU organization allows concurrent processing of up to five instructions. Full 32-bit word length in the CPU is augmented by a process exchange mechanism, which ensures that all simultaneously executed programs run at high speeds. For maximum speed, DMA transfers are

supported by burst-mode transfer mechanisms. Burst-mode I/O transfers 64 bits of data at a time over the 9955 system's 9-Mbyte I/O bus. This setup also enhances the efficiency of virtual memory management, by increasing the speed at which pages transfer from disk to main memory. The single precision whetstone rating for this high performance unit reportedly exceeds 4000, with the double precision whetstone rating set at over 3300. The 9955 system's



64-Kbyte cache memory increases processor throughput by providing faster access to frequently used instructions. The cache memory has a 40-ns access time and, by providing the CPU with the required instructions more than 98 percent of the time, it has a 58-ns effective main memory access time. The system provides soft error recovery, enabling the CPU to detect parity errors in the cache and reload from main memory. This recovers all soft errors and avoids halts in operation.

The time it takes to perform multiplication is cut to half that of the predecessor 9950. Advanced floating point acceleration hardware, an 8- x 48-bit macrocell multiplier array, offers a 48-bit execution path for double precision calculations frequently used in scientific and computationally intensive applications.

Languages that run on the 9955 (which uses the Primos operating system) include Fortran 77, ANSI 74 Cobol, Pascal, C, Basic, and RPC II. Prime computers can be connected

via the firm's Primenet and Ringnet configurations. Prime/SNA products are now available for connection to IBM-based networks.

A typically configured 9955 system, including a processor with 4-Mbyte main memory, two 315-Mbyte fixed disks with one controller, a streaming tape subsystem, one peripheral cabinet, and CRT console, costs \$371,000. Upgrades from the Prime 9950 cost \$45,000. **Prime Computer, Inc.**, Prime Park, Natick, MA 01760.

Circle 265

—J. V.

CMOS chip breathes new life into byte-wide processing

The HD64180 is an 8-bit CMOS CPU that affords designers a low power, high speed chip capable of directly addressing 512 Kbytes of memory, due to its onchip memory management unit. Also onchip are a dual-channel DMA controller, one synchronous and two asynchronous serial ports, a wait state generator, dynamic memory refresh circuitry, a 16-bit programmable reload timer, a 12-source interrupt controller, and a dual-bus interface. Chip speed will eventually accelerate from its current 6-MHz clock rate to 10-MHz, as production ramps up.

Due to its upward compatibility with the 8080 and Z80 microprocessors, the chip supports all currently existing CP/M and MSX software. Further, it holds promise of expanding the horizons of 8-bit software. An obvious candidate for applications such as portable personal computers, the chip's extreme versatility marks it as a prime contender for industrial control applications as well. Low power CMOS operation (the chip draws only 15 mA at 5 V, running at 6 MHz) and "sleep mode" instructions make it suitable for these applications, as well as for instrumentation and office automation roles.

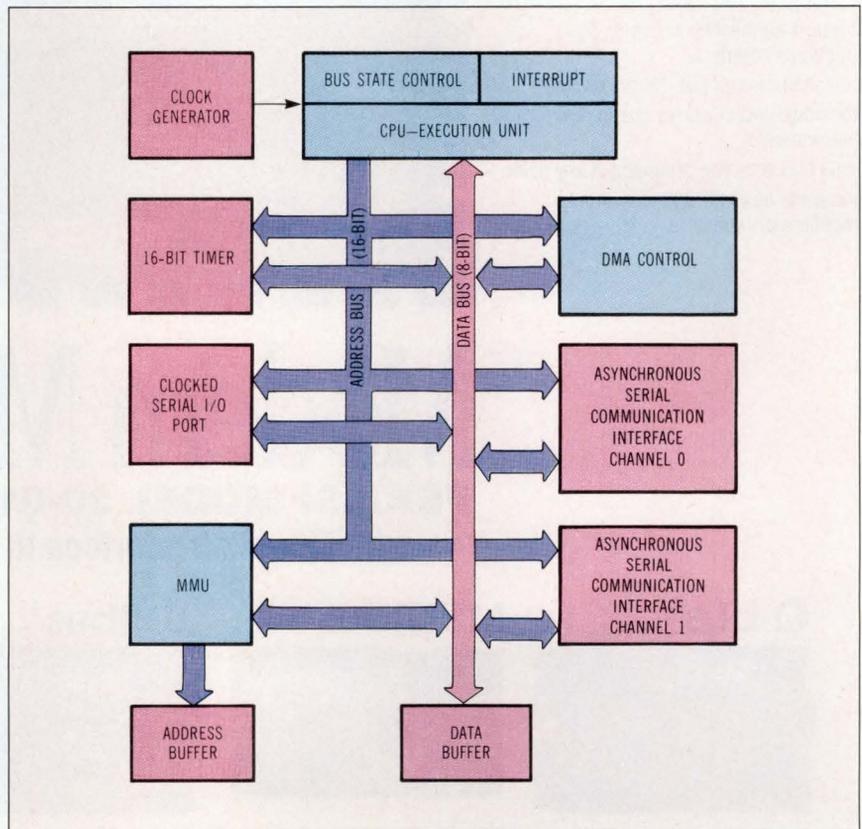
Sleep instructions can keep the CPU active while shutting down I/O, shut off the CPU while maintaining I/O, or shut down both the CPU and I/O. The processor sports other enhancements to the Z80 instruction set. Four of these allow the program-

mer to compare a byte in the accumulator with a byte in another register, in a directly or indirectly addressed memory location, or with an I/O port. These nondestructive comparisons set status flags in the manner of AND operations. Also significant among the new instructions is an unsigned 8-bit multiply with a 16-bit result. Moreover, the HD64180's MMU maps 64 Kbytes of

logical address space into the total 512 Kbytes of physical address space using a scheme that is similar to board-level products and entirely compatible with bank-switched CP/M Plus. The part sells for about \$15 each in sample quantities of 1000. **Hitachi, Ltd.**, 2210 O'Toole Ave., San Jose, CA, 95131.

—W.E.S.

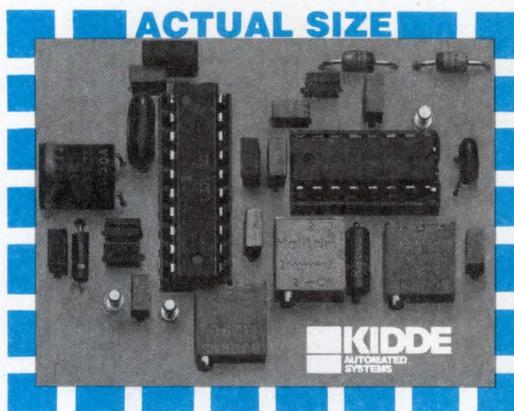
Circle 266



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

For **\$85⁰⁰**

- 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-1 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-1 can do the job more effectively. Call today.

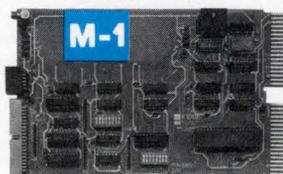
Rate (Mbits/s)	Carrier freq (MHz)	RG-59 (ft.)	RG-11 (ft.)	JT4412J (ft.)	JT4750J (ft.)	JT41000J (ft.)
1.0	3	5K	12K	20K	32K	50K
2.0	5	4.2K	9.5K	15K	25K	37.5K

**Quantities of 5000.

Modem M-1

FSK/LSI MODEL 30-0194
with Megalink interface to

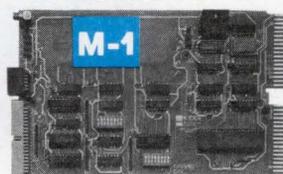
Q-bus* . . .



MULTibus* . . .



UNibus* . . .



VMEbus . . .



FROM

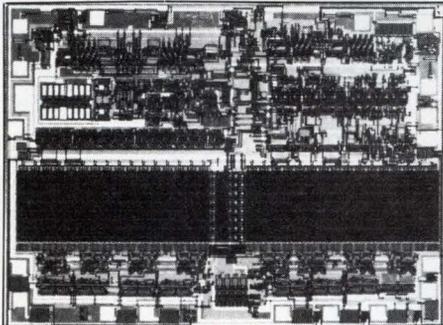
COMPUTROL

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.

Microcontrollers and microprocessors supported by nonvolatile RAM



Nonvolatile RAM that requires no interface circuitry supports Intel 8-bit microcontrollers and microprocessors. The 2001 combines electrically erasable PROM features with realtime read/write functions common to static RAMs. It eliminates interface circuitry requirements by transmitting address and data information on the same lines. And, its 128 x 8-bit architecture lets system designers use one chip to achieve critical parameter storage in 1 Kbyte of byte-wide nonvolatile memory. Onchip data-protection circuitry detects power supply drops of below 4 V, and internal erase and write circuitry shuts off. The 2001 is priced at \$11.20 in quantities of 10,000. **Intel Corp.**, 1800 Embarcadero Rd, Palo Alto, CA 94303. **Circle 267**

Components cut cross talk in speech digitizing tasks

Two silicon gate CMOS encoder/decoder ICs meet the demands for speech-digitizing components used in pulse code modulation systems. Designated the MV3507 and MV3507A, these components are packaged in a compact ceramic, dual-inline carrier. They meet or exceed the AT&T D3 and CCITT G.711, G.712, and G.733 PCM standards. Each chip includes an anti-aliasing filter. Standard features include independent transmit and receive sections with 75-dB isolation, power-down logic delivering a low operating power requirement of 60 mW (standby requirement, 5 mW), as well as an A/D encoder with a dual-speed, auto-zero loop for fast data acquisition response upon power up. A variable receive shift clock operates between 64 kHz and 2.0148 MHz to clock input PCM data. Cross talk is suppressed by a circuit that eliminates any transmitted idle channel noise during quiet periods. The 100-piece price for the MV3507/A chips is \$28.50 each. **Plessey Solid State**, 3 Whatney, Irvine, CA 92714. **Circle 268**

Fast dynamic RAM family gains 64-Kbit x 1 member

The IMS2600P-80, a 64-Kbit x 1 dynamic RAM, exhibits 80-ns worst-case RAS access time. As with prior versions of the IMS2600, the 80-ns part offers nibble mode and CAS-before-RAS refresh. The nibble mode feature allows access of 4 sequential bits at a worst-case data rate of 48 ns. Read cycle time for the product hits 130 ns. Using sub-2 μ m NMOS technology, the IMS2600P-80 aims at applications where very fast cycle and access times are mandatory. Offered in a JEDEC standard 16-pin plastic DIP, the 1000-piece price is \$12.45. **Inmos**, PO Box 16000, Colorado Springs, CO 80935. **Circle 269**

Control and DSP handled speedily by CMOS multiplier chip

Addressing the need for high multiply-accumulate throughput in microprocessor-based systems, the NCR45CM16 is a CMOS multiplier chip that plugs into many existing boards. Heavy multiply-accumulate intensive applications (eg, control and digital signal processing) can be offloaded from the microprocessor. The NCR45CM16 can perform successive multiply-accumulate operations at 5 MHz. The single-port device can directly interface to a 16-bit bus without additional "glue" logic. It uses a maximum of 10 mA and comes in a 0.6-in. wide, 24-pin plastic DIP (or 0.3-in. wide ceramic DIP). Samples of the plastic package are available for \$60 while the ceramic sample costs \$75. **NCR Microelectronics**, PO Box 57, Dublin, OH 43017. **Circle 270**

Semicustom gate arrays boast low power CMOS implementation

Semicustom CMOS gate arrays offer gate complexities from 330 to 1100. A library of 60 redesigned, fully characterized macrocells is available. The high speed version is implemented in a 3- μ m polysilicon process and sports a typical gate delay of 2.6 ns at 5 V/25 $^{\circ}$ C, over an operating range of 2 to 6 V. The medium speed family, implemented in a 4- μ m polysilicon process, has typical gate delays of 8 ns at 5 V/25 $^{\circ}$ C. The chips' makers boast low power consumption ratings (eg, standby power for the PCC0705 lists at 0.25 mW). Each gate array contains many identical uncommitted unit cells that are interconnected by two custom masks. Each unit cell contains four pairs of n- and p-

channel transistors. Samples are available now. (Outside the United States, these components are available from Philips Elcoma.) **Signetics Corp.**, 811 E Arques Ave, Sunnyvale, CA 94086. **Circle 271**

EEPROMs sport internal address and data latches

Third-generation EEPROMs under the XL2816A aegis are 16-Kbit parts marked by 2048-x 8-bit organization. Read/write cycles of 5 V, internal address and data latches, timing circuitry, and protection against inadvertent writes during power supply transitions also highlight this chip. It is compatible with the 2816 and 2816A. A JEDEC standard 24-pin, byte-wide memory configuration is used. A self-timed write cycle completes in under 10 ms. The XL2816A is designed for applications requiring 10 years of data retention, with or without power, and up to 10,000 write cycles per byte of memory. **Exel Microelectronics, Inc.**, 2150 Commerce Dr, San Jose CA 95131. **Circle 272**

Data separator chip can read or write at 5-MHz rate

A self-adjusting Winchester disk data separator, the WD10C20, is designed for use with the ST506 and WD1010/2010 families of single-chip hard disk controllers. Available in either a 28-pin DIP or surface-mount package, the WD10C20 can read and write modified frequency modulated data at a 5-MHz rate. This is a CMOS device that cuts real estate needs, replacing 10 TTL devices. High tolerance to bit jitter is reported. The unit will be priced at \$30 in 100-piece quantities. **Western Digital Corp.**, 2445 McCabe Way, Irvine, CA 92714. **Circle 273**

Device uses CMOS process for lower power consumption

CMOS devices in the 2504 series represent 12-bit successive approximation registers. The ZX25CO4 is a pin-for-pin plug-in replacement for the bipolar AM2504 and the DM2504. It contains all digital control and storage needed for successive approximation A/D conversion. The device is available in a plastic 24-pin DIP with lead times of 6 to 7 weeks for quantities of up to 20,000. Shorter lead times are expected for die and wafer versions. **Zytex Corp.**, 750 E Arques Ave, Sunnyvale, CA 94086. **Circle 274**

ALL THE RIGHT SIGNAL PROCESSING SOFTWARE ON ALL THE RIGHT COMPUTERS. SO YOU DON'T HAVE TO WRITE AT ALL.



ILS™, the undisputed world standard in signal processing software, has just pulled farther ahead. ILS now runs on even more computers than before. Besides the DEC VAX and PDP families, it is now available for the Data General ECLIPSE MV series, MASSCOMP MC computers and the IBM PC lineup: PC, PC/XT and PC/AT.

ILS is an integrated package of 88 programs providing data acquisition support, waveform

display and editing, digital filtering, spectral analysis and much more. It will perform everything from common Fast Fourier Transforms to demanding functions like Hilbert Transform, octave band filtering and 3-D spectra.

Our success is based on the fact that it would take more time and money than most of you would care to invest, just to come up with programs that wouldn't be as good or as complete as ILS. And ILS has already

been proven in the fields of speech, noise and vibration, seismology, acoustics, sonar, radar and biomedicine.

All ILS packages have a user interface with a prompt-driven or command control option. ILS also works with DACS, our data acquisition and conversion software.

Call our toll-free number now and arrange for a dial-in demo, demo floppy or video presentation. But call right now. Don't wait to write.

Signal Technology Inc

5951 Encina Road • Goleta, California 93117

(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 18559/Switzerland • Zentrara Technik AG, Tel: 032-233553, Telex: 845 349353/U.K. Logica, Tel: 01637 9111, Telex: 851 267413/Japan • Rikei Corp., Tel: 03 345 1411, Telex: 781-24208/Korea • Sam Hong Corp., Tel: 753 3100, Telex: 787 23231/Italy • Morra Ingegneri, Tel: 172 422477, Telex: 215415/Taiwan • Hermes Epitex Corp., Tel: 02 562 851, Telex: 785 25794 • DEC, VAX and PDP are trademarks of Digital Equipment Corp. ECLIPSE MV is a trademark of Data General Corp. IBM PC, PC/XT and PC/AT are trademarks of International Business Machines Corp. MASSCOMP is a trademark of Massachusetts Computer Corp.

Facilities control system handles up to 10,000 points

An integrated energy management and facilities automation system, the Alpha/Net 4000, can control as few as 64 or as many as 10,000 points. System hardware includes CPU, operator interface, and printers. The field I/O hardware is configured around the P4000 distributed processing unit. Software enables the system to perform electric demand limiting, adaptive optimized start/stop, duty cycling, event control, trend logging, and other functions. It is based on a system handling less than 75 analog and digital I/O points. Prices range from \$10,000 to above \$12,000. **AMF Control Systems**, PO Box 59469, Dallas, TX 75229.

Circle 275

Microcomputer with multi-user capabilities joins 8870 family line

The Micro 7 multi-user system is oriented toward interactive applications. It provides word processing, calculation, administration, and organizational functions at the workstation. The 8870 Micro 7 delivers 256-Kbyte main memory, and

provides a 5¼-in. floppy disk unit plus a 10-Mbyte fixed disk. Ergonomic terminal design highlights the machine. The workstation includes an optional, height-adjustable terminal base and black mesh netting for reduced eye strain. The system uses Nixdorf's interactive realtime operating system (NIROS), which supports multiprogramming in interactive and batch modes. The 8870 Micro 7 basic system is priced at \$6995. **Nixdorf Computer Corp.**, 300 Third Ave, Waltham, MA 02154.

Circle 276

Workstation line offers multitasking, uses multiple VLSI microprocessors

The 32-bit Stations 10 and 20 achieve performance of 0.575 million instructions per second. Multiple VLSI microprocessors and dual-bus architecture distribute processing loads to appropriate subsystems. These workstations can act as standalone computers, fully integrated graphics systems, Ethernet network nodes, or as RS-232 terminals for Harris superminis. The CPU in each Harris workstation uses dual 32-bit microprocessors—a 68010 and 68000—running at 10 MHz. Multitasking

is enhanced by 24 sets of 32 double-precision (64-bit) registers. An optional floating point processor can work in parallel with the CPU. Prices for the Station 10 start at \$32,500; prices for the Station 20 begin at \$39,000. **Harris Corp., Computer Systems Div.**, 2101 W Cypress Creek Rd, Fort Lauderdale, FL 33309. Circle 277

Multiprocessing computer derived from 32032-based system

Expert 32, a multiprocessor computer, now supports up to six 32-bit processor card sets, and is capable of handling as many as 12 processors. The processor card (called the E82/CPU2) is based on the 32032 CPU and includes floating point processor, memory management, realtime clock/calendar, and 32 Kbytes of cache memory. The system features a 16-slot, double-width VMEbus backplane, and 512 Kbytes of main memory (expandable to 16 Mbytes). A portable virtual storage operating system with Unix support is provided. Pricing starts at \$17,700. **Elite Computer Systems**, 4129 May St, Wichita, KS 67209.

Circle 278

OASYS TOOLKIT SNAP-SHOT

C-68000

CROSS AND NATIVE OPTIMIZING COMPILERS

FOR 68000/10 (and 68020 SOON)

OASYS offers a "ONE STOP SHOPPING" service for software developers in need of proven 8-, 16- and 32-bit cross and native tools for Unix and non-Unix 68000, 8086 and 32000 systems. Our critically acclaimed and widely used 68000 tool kit offers high quality, reliable, cost-effective tools.

The OASYS 68000 tool kit consists of Green Hills compilers (C, Pascal and FORTRAN), our own M68000 Macro Assembly Development package, and dozens of other OASYS compatible support tools. Simply stated, we beat the competition on price, speed and tightness of emitted code.

C-68000/10

- Full K & R with Western Electric and Berkeley extensions
- Complete run-time library available as source. No royalty if passed on.
- Supports DEC & IEEE Floating Point
- Integrated optimizer: 30% tighter code than Portable C; 4 times faster
- Generates M.I.T. or EXORmacs assembly source code
- Interfaces to all OASYS tools and Pascal, FORTRAN and PL/M-68K compilers
- Ideal for cross development of boards with no OS, a kernel OS (e.g. VRTX, PSOS, MTOS), or Unix based 68000's

68000/10 Assembly package

- EXORmacs compatible Macro Assembler, Linker, Librarian, and Cross Reference Utility
- Generates S-records and a.o.
- PIC and reentrant code
- Used 2 years in house
- Over 3,000 sold to date
- Runs on VAX, Prime, PDP-11, 68000's, 8086/88 (PC)
- Written entirely in C

Coming soon

- 68020 C and Cross assembler

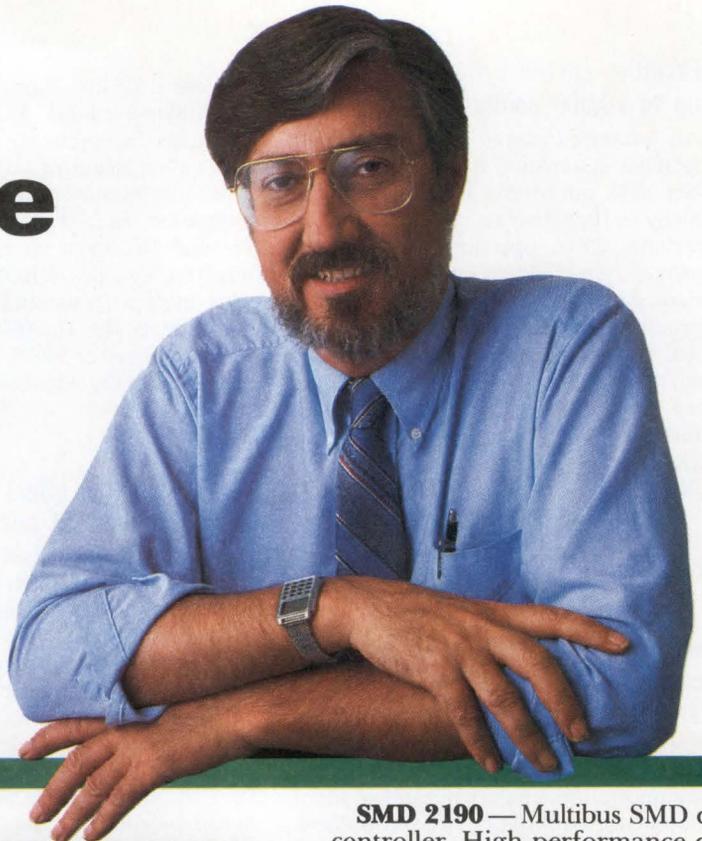
Other tools

- Symbolic C Source Debugger
- 68000 Simulator & Disassembler
- C Linecount and Time Profiler Utility (CLUE™)
- LINT for VAX/VMS
- Check Out compiler (SAFE-C™)
- Communications tools

OASYS

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

"Interphase controllers will introduce your system to life in the fast lane."

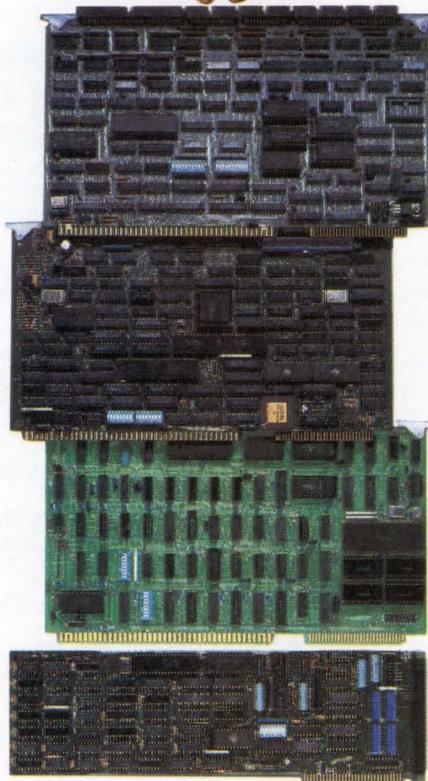


Tom Thawley
Vice President of Engineering
Interphase Corporation

Interphase, one of the original producers of Multibus™ products, now offers more high-speed, high-performance Multibus controllers than ever before. Including a full line of controllers for 5¼", 8" and 14" hard disks, local area networks and video monitors, plus powerful controllers for the IBM® PC. And we're continually developing new controller products, to give your system more power.

Storager™ — Multibus multi-function controller. Supports ST506, ESDI, ST412HP interfaces for Winchester, QIC-02 for tape, and 3½", 5¼" and 8" floppies. Intelligent caching, concurrent disk-tape-bus transfers, on-board 68000 CPU, 1:1 interleave with zero latency reads and writes. Makes a 5¼" disk perform like an 8"

Maverick — SMD rigid disk controller for the IBM PC and compatibles. Supports two SMD drives (16-800 MB). Five times the speed of a 5¼" Winchester plus removable disk cartridges. The ultimate in reliability. Software for PC-DOS™ 2.0 and QNX™ available. Compatible with most LANs.



SMD 2190 — Multibus SMD disk controller. High-performance disk controller with intelligent cache memory for UNIX™, RMX and similar operating systems. Supports four SMD drives, 1:1 interleave, 20 Mb/S data rate, 32-bit ECC, automatic error correction, 4.5 MBytes DMA, 24-bit addressing.

LNC 5180 — Multibus token passing network controller. Supports up to 255 devices on a star or bus network. Data rates up to 2 Mb/S on coax or T1 lines. Higher throughput than Ethernet™, automatic flow control and error recovery, ISO levels 1-5 on board. Simplified software interface with features for real-time applications.

All Interphase products are backed by a great customer support team that makes sure our products work the way they should — in your system. Let Interphase introduce your system to life in the fast lane. Call 214/350-9000.

Multibus is a trademark of Intel Corporation.
UNIX is a trademark of AT&T's Bell Laboratories.
IBM and PC-DOS are registered trademarks of International Business Machines.
Ethernet is a trademark of Xerox Corporation.
QNX is a trademark of Quantum Software Systems, Ltd.
Storager is a trademark of Interphase Corporation.



INTERPHASE
corporation

2925 Merrell Rd. • Dallas, TX 75229

Multi-user board turns dumb terminals into PCs

Containing an Intel 8088 CPU, 256 Kbytes of bank selectable DRAM, a IBM PC bus interface, and two serial ports, the PC-Slave/16 is fully hardware and software compatible with the PC and most PC-compatible systems. It can make a second PC out of a dumb terminal. Using the PC-Slave/16, the original PC becomes a master PC processor. RTNX executive software allows the master PC and up to 12 slave PCs to share disks, peripherals, and data. Price for the PC-Slave/16 is \$1095. **Advanced Digital Corp.**, Huntington Beach, CA 92649. **Circle 279**

Single-board system delivers zero wait-state and 1 Mbyte of RAM

A single-board computer, the ZX-186/30, is an iAPX-186 system. It features up to 1-Mbyte RAM, up to 128-Kbyte PROM, a dual-USART, 24 parallel I/O lines, two DMA channels, and two iSBX connectors plus timers and interrupts. Compatibility with iSBC-86 is a further enhancement. The ZX-186/30 is based on the 8207

dynamic RAM controller and either 64-Kbit or 256-Kbit DRAMs. With the flexibility inherent in 128-Kbyte to 1-Mbyte onboard RAM, the ZX-186/30 can adapt to diverse applications. Quantity-100 price is \$1550 for the ZX-186/30 with 128-Kbyte RAM. **Zendex Corp.**, 6644 Sierra Ln, Dublin, CA 94568. **Circle 280**

Expandable 16-bit VMEbus computer uses the 32-bit 68000

A single-board computer, designated the PME 68-1B, has a 32-bit 68000 CPU (8 MHz, optional to 10 MHz), 128 or 512 Kbytes of onboard RAM, and 16 Kbytes of firmware that is expandable to 64 Kbytes. This board comes standard with a 16-Kbyte system monitor containing terminal control of the debug, up/down loading, and memory modification-instructions. Also provided is one line assembly/disassembly and full access to 16 Mbytes of memory space. Optional software includes PME 68, Ideal (an EPROM resident, screen-oriented editor/assembler), and PME 68/Coherent (a

multi-user operating system for VMEbus configurations). **Plessey Microsystems**, 1 Blue Hill Plaza, Pearl River, NY 10965. **Circle 281**

Bridge takes the form of standard single-height Eurocard package

Available on a standard single-height Eurocard, the GESBRG-1 high resolution resistive bridge interface is compatible with the G-64 bus. The board uses a high reliability DIN 41612 indirect connector and is powered by -12 V, as well as 5 and 12 V, on the G-64 bus. The European version sports excitation frequency of 400 Hz and 80-ms conversion time—the U.S. version offers 480 Hz and 66.66-ms conversion time. Board is available for \$745. **Gespac, Inc.**, 550 E Grandview Dr, Mesa, AZ 85203. **Circle 282**

Talk to the editor

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

**OASYS
TOOLKIT
SNAP-SHOT**

WIZARD C-8086

CROSS AND NATIVE OPTIMIZING COMPILERS

FOR 8086/87/88/186/286

As part of OASYS' "ONE STOP SHOPPING" service for software engineering tools, we are proud to announce the addition of WIZARD C to our integrated collection of more than 50 professional programming tools (e.g. compilers, assemblers, linkers, debuggers, simulators & translators) for M68000, Intel 8086/80186 and NS32000 micros.

WIZARD C benchmarks (against Lattice, Microsoft) prove that it is, by far, the most advanced, full featured, fastest, tightest, optimizing C compiler now available for cross and native (PC) development. Here's why . . .

FEATURES

- Complete K&R implementation plus V7, III extensions
- Supports 8087 Floating Point
- Built-in LINT
- Long, medium, short models
- 190+ UNIX III functions — complete run-time library
- In-line assembly allowed
- 100+ extensive warnings/diagnostics
- Intel and Microsoft compatibility at source and object levels
- Written in C — easily ported
- Comprehensive bound documentation
- Supports DOS 2.0, 2.1, IBM/BIOS

SUPPORT TOOLS

- Symbolic C Source Level Debugger (CDEBUG™)
- 100% Intel compatible structured Macro Cross Assembler, Linker/Locator and Librarian
- 8086 Simulator
- Floating point math package (40+ functions)
- C Time Profiler (CLUE™)
- Checkout compiler (SAFE-C™)
- Comm. utilities for up/down loading to MDS, TEK, Microtek

AVAILABILITY

NATIVE: PC/XT/AT using MS/DOS, PC/DOS (Xenix soon)

CROSS: VAX/VMS, Bsd 4.1, 4.2, III, V; 8086's, 68000's (All Unisoft III, V; ports); Callan, Masscomp, Sun, Pyramid, dozens more . . .

Call for pricing, OEM, Site, Corporate, Source and Maintenance licensing information.

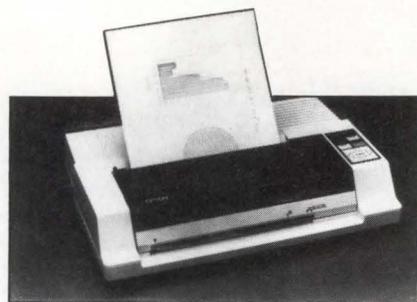
OASYS

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

Lightweight printer/plotter simplifies operation with ROM-based commands

A four-color business graphics pen plotter/printer features intelligence in the form of ROM-based commands that simplify 42-plotting and 17-drawing functions. Dubbed the HI-80, this plotter/printer has standard Epson print and plot

ROM commands for graphic presentation of data. ROM emulation of commands standard to the Hewlett-Packard graphics language is optional. The unit weighs 9.9 lb and uses a four-pen cassette containing ball-point and fiber tip pens. The machine boasts repeatability to within 0.012-in. and a minimum step size of



0.004 in. The HI-80 prints at 6.5 chars/s and carries a set of 96 ASCII, 32 international, and 96 graphic characters. Maximum plotting speed reaches 9 in./s along either axis. In quantities over 1000, the price is about \$350. **Epson OEM Products Div**, 23600 Telo Ave, Torrance, CA 90505.

Circle 283

PUT Canon Precision IN YOUR PRODUCT

SUB-FRACTIONAL DC MOTORS



Canon Precision Inc., a wholly owned subsidiary of Canon Inc., is one of the world's largest manufacturers of DC micromotors. These motors are designed and built to the same standards of excellence as Canon's famous cameras, copiers, calculators, typewriters and other fine products.

We offer a complete line of over 100 sub-fractional HP DC motors and gearmotors ranging from 13mm diameter to 52mm diameter in conventional iron core motors and from 39mm to 52mm in brushless types. Electronics are available for the brushless motors. The motors are designed to meet UL, CSA and FCC requirements. Special modifications are available in OEM quantities.

If you need precision in your products... design it in... with Canon Precision.

For a comprehensive catalog, call or write:

Canon U.S.A., Inc.
Components Division
One Canon Plaza
Lake Success, NY 11042

Tel: (516) 488-6700
Telex: 96-1333
Cable: CANON USA LAKS

Canon®

Dot-addressable printer operation marked by 700-char/s speed

Unique printhead technique aimed at achieving high speed and high reliability marks the OT-700. This 700-char/s matrix printer offers performance characteristics equal to matrix line printers as well as 350-char/s operation with dot-addressable graphics capability. The OT-700 has no duty cycle limitations. The unit offers full 136-column carriage width with adjustable sprocket feed tractors. Paper feed is from the front or bottom of the case. The control panel sports membrane switches and LED lights. Dot-addressable graphic printing is available in two operating modes—50 x 69 dots/in. or 100 x 69 dots/in. The OT-700 is compatible with popular software programs. Price is \$1495. **Output Technology Corp**, 606 100th St, Bellevue, WA 98004.

Circle 284

Fast and quiet printers use optical scanner in paper-handling setup

High speed letter-quality printers in the PT-90 group offer a 4-page/min (200-char/s) mode and a fast 8-page/min (400-char/s) draft mode. Average sound levels are rated at less than 45 dBA. Overall resolution is 240 dots/in., both horizontally and vertically. Character resolution is 96 x 32 dots in letter-quality mode and 48 x 16 dots for draft printing. The PT-90 accepts plug-in interface cards for serial, parallel, or IEEE 488 ports. Advanced paper handling (marked by an optical scanner) detects the edges of the paper, making manual alignment of cut-sheet paper unnecessary. Cost is \$3495. **Siemens Communication Systems, Inc**, 5500 Broken Sound Blvd, Boca Raton, FL 33431.

Circle 285

Software Development Tools

If: Manufacturer-Compatible,
Fully-Supported Cost-
Effective, Cross and Native
Development tools are
important to you!

Then:
Let us tell you about
MICROTEC RESEARCH's extensive
line of field-proven software
development tools for serious
software developers.

**C
Pascal
PLM
Assemblers
Simulators
VT-100 Emulator
CPM/80 Emulator**

Manufacturer Compatible

MICROTEC RESEARCH has been providing flexible and economical solutions for software developers since 1974. We've grown with the industry as our cross software development tools have transformed many large and small computers into powerful cross development systems for microprocessor applications. Our software tools are manufacturer compatible. Therefore, software made using manufacturers development systems may be more productively used in the MICROTEC RESEARCH cross-development environment.

Fully-Supported

MICROTEC RESEARCH's products are continually maintained and updated based upon the experience of thousands of installations worldwide. Revisions and modifications are distributed to licensees in warranty or covered by a service contract. Upgrades, which are major enhancements to a product's capabilities, are available at a significant discount. Our update/upgrade policy assures users they'll always be current with the fast moving world of microprocessor development.

Effective Differences

Beginning with product concept, through development, quality assurance, and post-sales support — **Quality, Compatibility, and Service** are the differences which set MICROTEC RESEARCH apart from the others.

Start Saving Time & Money

If you are a serious software developer, shopping for software development tools, call or write today for more information.
800-551-5554 In CA call **(408) 733-2919**

Target Microprocessors	Host Computers	Software Tools
8086/80186	DEC VAX, PDP-11	C Compilers
8096	DG MV-Series	Pascal Compilers
8080/8085	DG Eclipse	PLM Compiler
8051	Apollo	Assemblers
8048	UNIX Systems	Linking Loaders
Z8002	IBM PC	Librarians
Z80	Data General/ONE	Download
Z8	HP 150	Communications
NSC 800	DEC Rainbow	VT-100 Emulator
9900...	others	CPM/80 Emulator
	G65SCXX, G65SC1XX	
	R65C00, R6500/1..	
	others	

for Serious Software Developers

MICROTEC is a trademark of Microtec Research, Inc. Santa Clara, CA



3930 Freedom Circle, Suite 101, Santa Clara, CA 95054
Mailing Address: P.O. Box 60337, Sunnyvale, CA 94088
(408) 733-2919 • Telex (ITT) 4990808

Supply family features 30 available voltage/ampere ratings

The KFD series consists of 25-, 40-, and 80-W multi-output units that are designed to meet both the IEC 380 and VDE 0806 international safety specifications as well as the UL/CSA. For emi, the series complies with both VDE 0871 and FCC

level B. Designed for varied applications including CRT terminals and modems, the devices are available in 30 different voltage/ampere ratings. The 25-piece U.S. unit prices are \$44 for the KFD 25E, \$52 for the 40E, and \$92 for the 80E. **KEC Electronics, Inc.**, 20817 Western Ave, Torrance, CA 90501. **Circle 286**

Surface-mount technology shrinks power supply size

Introduced as the first open frame switchers to use surface-mount technology, 150-W MCT series members are multi-output, 2½- x 5- x 9-in. units designed to drive Winchester disks, floppy disks, and streaming tape backups. The MCT saves space by using a 60-kHz switching frequency. This allows for smaller magnetic components and more efficient bipolar switching devices. A compact output filter network cuts size further. Four units are available—all with 5-V, 25-A, and 12-V 5-A (10-A peak) outputs. The triple output MCT-309-0512 adds a tightly regulated floating 12 V at 2 A, while the third output of the MCT-309-0512S is a semi-regulated 12 V and 4 A. The quad-output MCT-409-1212 has as its third and fourth outputs 12 V at 1.5 A, and -12 V at 1 A. The third and fourth outputs of the MCT-409-1215 are 15 V at 1.5 A and -15 V at 1 A. Strap-selectable ac inputs may vary from 90 to 132 V, 180 to 264 V, or 47 to 63 Hz. Featuring improved regulators in a single-stage, half-bridge, pulse-width modulated configuration, each MCT supply costs \$295. **Todd Products**, 50 Emjay Blvd, Brentwood, NY 11717. **Circle 287**

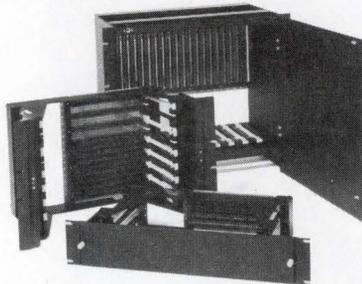
MULTI-CAGE® Card Cages for the MULTIBUS:

MULTI-CAGE®: the industry standard Multibus card holder and backplane system. Designed for performance, built with precision.

Choose from the widest selection available from any manufacturer. The durable, anodized aluminum cages, featuring single mother board backplanes, are available in 17 slot sizes: 3, 4, 5, 6, 7, 8, 9, 10, 12, 13, 14, 15, 16, 20, 21, 24, and 26 slots.

Choose from freestanding, horizontal and vertical rack mount cage designs. Choose card spacing of either 0.6 or 0.75 inch. There's even special slot spacing models for efficient use of three-level wire wrap or iSBX 300/340 modules. Every model is UL-recognized. Multi-Cage is electrically and dimensionally interchangeable with Intel's iSBC-80 cages.

Choose, too, from a huge variety of packaging options and accessories, such as P2 connectors and backplanes, Parallel Priority Resolution Circuitry (PPRC), cooling fans and power distribution.



Whatever your Multibus packaging needs, give us a call on one of our toll-free numbers. You'll find the best selection in the industry, along with competitive prices, quick delivery, and an unmatched THREE YEAR WARRANTY on every Electronic Solutions product.



Call TOLL FREE
(800) 854-7086
in Calif.:
(800) 772-7086

9255 Chesapeake Drive
San Diego, CA 92123
(619) 292-0242
Telex II (TWX): 910-335-1169

and iSBC-80 are trademarks of Intel Corp.
*Multi-Cage is a trademark of Electronic Solutions

Power source is 36-kVA pulse width modulation model

Model 3091-AVR3 ac power source has features that include three-phase variable control (45 to 500 Hz) operating modes. High frequency pulse-width modulation is utilized for high power density and quiet operation, as well as efficiency rated at 90 percent. Load regulation and distortion are approximately 1 percent with a 150 percent surge capability. The unit is housed in a cabinet weighing 810 lb. **Helionetics, Inc, DECC Div**, 17312 Eastman St, Irvine, CA 92714.

Circle 288

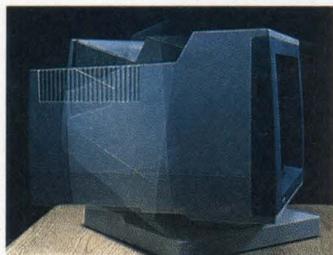
UPS attributes include convection cooling

Source 2 uninterruptible power systems are designed specifically for industrial use. The Source 2 provides up to 60 minutes of steady ac power. If utility voltage drops more than 15 percent below nominal, the UPS immediately begins supplying sine wave ac power from energy stored in its internal battery. The unit is available in power ratings of 200 and 400 VA, in 50-Hz and 60-Hz versions. Prices start at about \$695. **Topaz, Inc**, 3855 Ruffin Rd, San Diego, CA 92123.

Circle 289

GOOD LOOKS RUN IN THE FAMILY.

LEXIDATA'S LEX 90 FAMILY OF HIGH-PERFORMANCE DISPLAY PROCESSORS.



Lexidata's monitor enclosure design features a unique tilt and swivel movement with smooth and easy operation. All monitor connections and peripheral power supplies are located in the enclosure's base.

and with our popular Series 3000 line. So performance upgrades are easy and economical. And software investments are well protected.

Every LEX 90 model has the same basic components which use the latest bit-slice technology for rapid execution of complex calculations and set-up times.

The family's flexibility allows you to buy only the functionality you need now, and gives you the ability to add on later as required.

Exclusive SimulRes functionality allows simultaneous display of



A tiger shows off its "true colors" with the LEX 90 True Color display system from Lexidata.

8-bit color lookup table for each color. And where you're looking for high performance with high resolution, there's a LEX 90 model that can support up to 24 planes of 1280x1024 display memory at 60Hz non-interlaced refresh.

LOOKING GOOD.

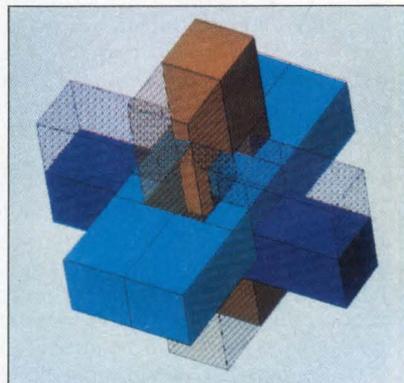
The LEX 90™ family is a good family to know. It's as advanced in a business sense as it is technologically. And what it has, compared to what it costs, makes it look very good indeed. For example:

Its software architecture is compatible both within the family

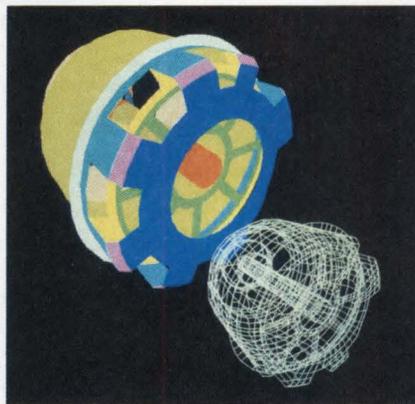
The extraordinary solid modeling display capability of our patented SOLIDVIEW™ technology is available in a 640x512 version. And for petrochemical applications, there's GEOVIEW™, a three-dimensional geological interpretation package.

On-board diagnostics, based on a 6801 micro-processor, help maximize uptime. And the unique tilt-and-swivel movement of our monitor allows smooth and easy operation of your display.

Of course, the entire family is



SOLIDVIEW, Lexidata's patented technology for the display of solid models, is available as an option on selected LEX 90 display processors. Its translucent shading pattern allows you to see inside your solid model.



LEX 90's SimulRes enables the designer to mix theory with reality by allowing the simultaneous display of a wire-frame model with the real-life product. (Courtesy of PDA Engineering)

640x512 and 1280x1024 images on the same screen. A True Color configuration offers two buffers of 640x512x24 with an

USA 01865. TWX: 710-347-1574. In UNITED KINGDOM call: Lexidata Ltd., Hook (025672) 3411.

In FRANCE call: Lexidata SARL, Rungis (1) 686-56-71. In JAPAN call: Lexidata Technical Center, Tokyo 486-0670.

compatible with a variety of leading application software packages and minicomputers. Which makes everyone look even better.

For a good look at the LEX 90 family and its expanding capabilities, call us at 1-800-472-4747. In Massachusetts, call (617) 663-8550. Or write: Lexidata™ Corporation, 755 Middlesex Turnpike, Billerica, Massachusetts



LEXIDATA®
LOOKING GOOD.

Unix software supports streaming tape drives

Streaming tape drives complemented by Unix support software permit image copying or selective backup of individual files. These 5¼-in. drives support the Minibox Unix development systems (Systems III and V). The drives use industry-standard 3M data cassettes (with up to 60 Mbytes of data storage capacity) and feature a fast 90-in./s read and write speed. Electronic edge detection circuitry is used. **Heurikon Corp.**, 3201 Latham Dr, Madison, WI 53713. **Circle 290**

Backup system readily connects to PC/XT's external floppy connector

Model 5210 FloppyTape backup systems represent 25-Mbyte units that plug into an XT external floppy disk interface connector. This makes a controller interface card unnecessary. The 5210 quarter-inch cartridge tape drive features file-oriented backup. It can back up two IBM 10-Mbyte hard disks or one 20-Mbyte hard disk at about 1 Mbyte per minute. Three new

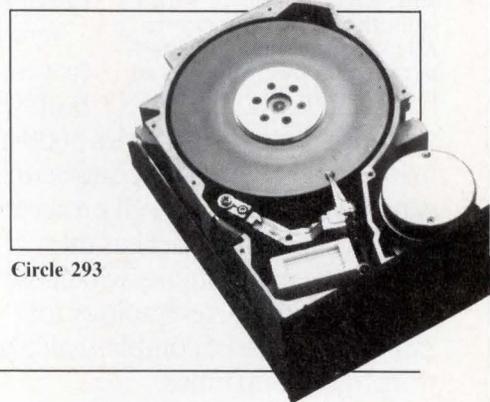
commands are added to PC-DOS for backup, restore, and format functions. Price is \$1095, with quantity discounts available. **Cipher Data Products, Inc.**, 10101 Grove Rd, PO Box 85170, San Diego, CA 92138. **Circle 291**

Solid state add-on pushes memory capacities of PC varieties

Megaram-PC represents a high speed, solid state, nonrotating memory system for the PC/XT/AT triumvirate. The system interfaces to the computer as if it were a disk. Transfer rate lists at 0.95 Mbytes/s. Random access capacities range from 0.5 Mbytes to 8 Mbytes, expandable in 0.5-Mbyte increments. Megaram-PC comes with an I/O driver that is compatible with applications software. Optional to the basic Megaram-PC are an error correction/detection feature and battery backup. **Imperial Technology, Inc.**, 831 S Douglas St, El Segundo, CA 90245. **Circle 292**

Drives debut with 60- and 80-Mbyte capacity and fast access time

Disk drives with capacities of 60 and 80 Mbytes compose the CM7000 5¼-in. Winchester series. The drives have access times under 40 ms and use oxide media. The 60-Mbyte CM7660 sports a three-disk configuration with six read/write minimonolithic heads. The 80-Mbyte CM7880 drive has four disks and eight heads. Track density of 1173 tracks/in. and bit density of 9275 bits/in. are achieved. Quantity prices are \$1135 for a CM7664 model and \$1205 for the CM7885. **Computer Memories Inc.**, 9216 Eton Ave, Chatsworth, CA 91311.



Circle 293

SYSTEM ELEMENTS

Liquid crystal display includes all drive and multiplexing circuitry

The F4816 provides 64 x 480 dots of fully addressable liquid crystal display. This low power module handles alphanumeric or graphic displays and permits easy serial interfacing. The module contains all the drive and multiplexing circuitry needed to control the LCD panel. Onboard is a power converter circuit that allows use of only one supply voltage (5 V). Viewing area is 8.86 x 1.73 in. The F4816 costs \$207 in quantities of 100. **Seiko Instruments**, 2990 W Lomita Blvd, Torrance, CA 90505. **Circle 294**

Press-fit card-edge connector eliminates flow soldering

Series 6353 card-edge connectors are press-fit, eliminating flow soldering and wire-wrapped connections. Measuring 0.125 x 0.250 in., these one-piece units have compliant Varipin contacts, fitting into PC board backpanels that use plated through-holes for increased efficiency. Series 6353 offers users between 12 and 100 contact sizes, with contacts available in selective gold plating on the mating

area and in gold flash or bright tin lead tails. **Elco Corp.**, 74 Brookhollow Dr, Santa Ana, CA 92705. **Circle 295**

RF wideband transformers operate in the 500-kHz to 500-MHz range

A series of surface-mounted subminiature rf wideband transformers exceeds stringent standards. They run in the 500 kHz to 500 MHz range. Power rating is 100 mW. The series features dc isolated or nonisolated primary and secondary windings at impedances from 50 to 800 ohms. Custom designed units meeting specific frequency ranges, impedance levels, and insertion loss requirements are also available. **Vanguard Electronics**, 1480 W 178th St, Gardena, CA 90248. **Circle 296**

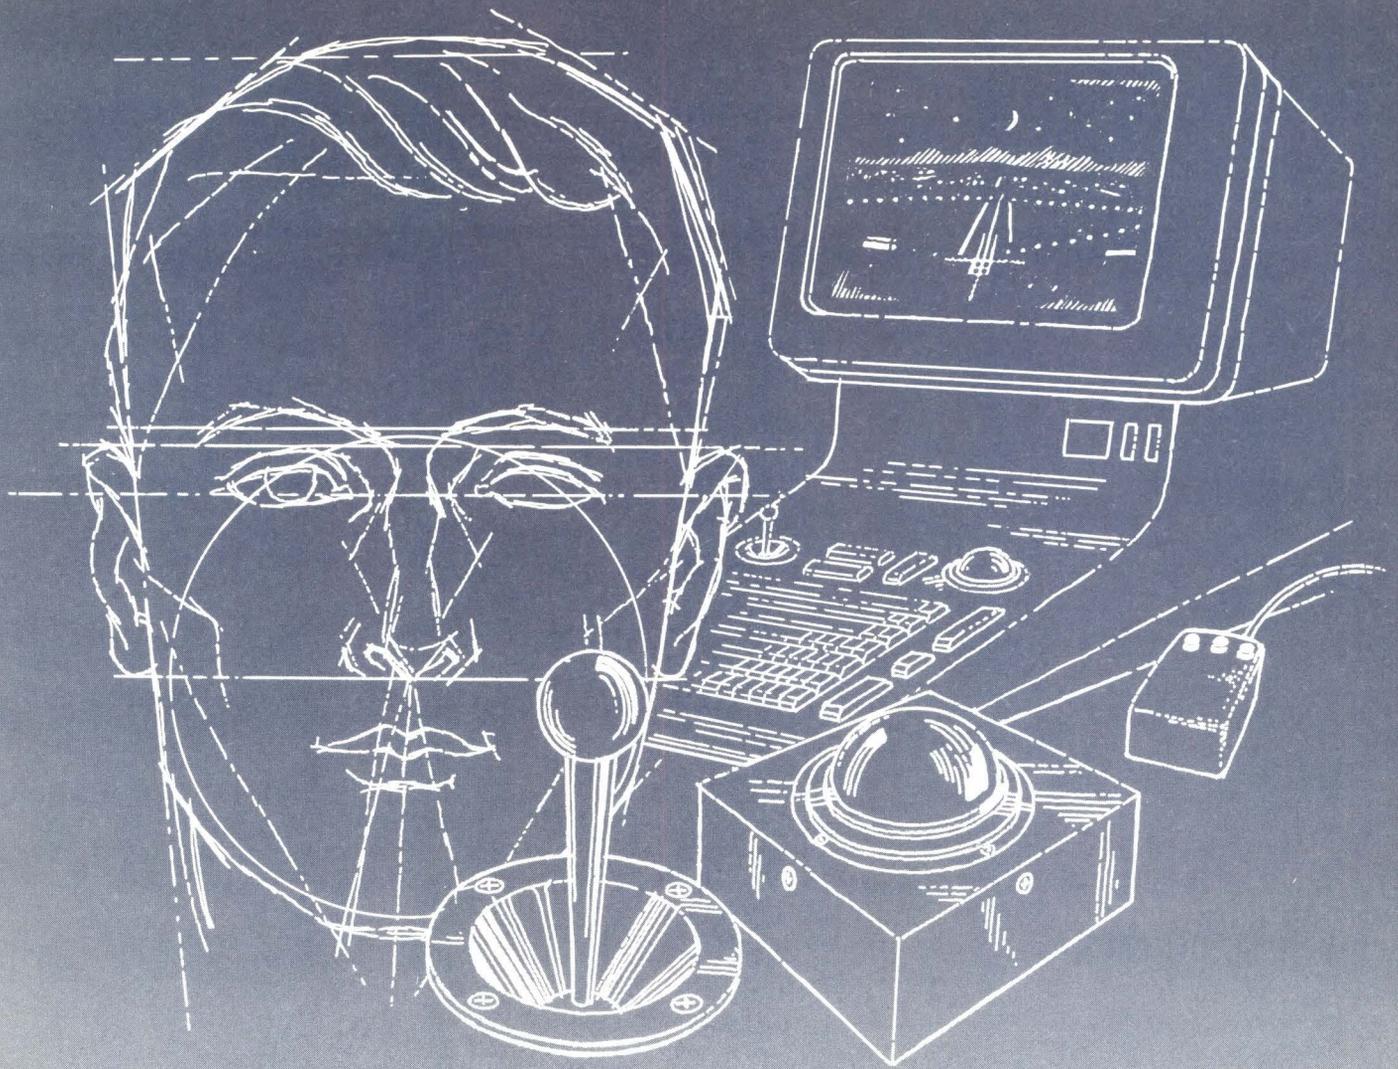
Fiber optic line sports modular construction

The HD cable series is available in two to eight fiber configurations, with either 50/125 or 100/140 micron fiber. The cable consists of strain-relieved 3.0-mm subunits wrapped around a dielectric cen-

tral member. HD cables contain a ripcord to allow quick removal of the outer jacket for termination of individual units. Cable is compatible with most available single-position connectors. **Celwave Technologies, Inc.**, **Optical Fiber Products**, Box 3343, Hickory, NC 28603. **Circle 297**

Printer mechanism delivers text and graphics at 480-dot/s rate

The model 3720 40-column, color dot-matrix impact printer mechanism requires no ribbon. It does not squirt ink, does not clog, and produces both text and graphics at up to 480 dots/s (dual-density). Because it is ribbonless, it uses thinner and lighter print wires that claims to yield higher density characters of better quality. It requires 6 V to operate. The mechanism's standard character font is 9 x 13 dots measuring 1.82 x 2.62 mm. Effective print size measures 96-mm horizontally x 192-mm vertically. Printing speed hits 70 mm/s on an X-Y axis, 100 mm/s on a 45-degree angle. Designer-quantity pricing is under \$50. **Epson America, Inc.**, **OEM Products Div.**, 23600 Telo Ave, Torrance, CA 90505. **Circle 298**



What kind of operator do you design for your terminals?

When you design a work station, you naturally look for positioning and tracking controls that will permit optimum efficiency, speed and accuracy. Yet, sometimes the most critical link in the entire system is neglected.

The operator.

No matter how good positioning and tracking controls may be, their effectiveness is diminished in proportion to the difficulty in using them. Obvious? Yes. Disregarded? Often.

But not by Measurement Systems.

We're a leading designer and manufacturer of joysticks, trackballs and control grips. From our beginnings a quarter of a century ago, we've recognized the importance of the human element in successful equipment design. We've spent substantial time and money researching the interaction of man and machine. We've become experts in the human factors in control design.

We've put our extensive knowledge to good use. For example, a joystick that's perfect for one type of application is often totally unsuitable for another. So we offer joysticks in numerous sizes and models, each with features suitable for the special requirements of dissimilar applications.

Versatile as they are, if our standard models don't meet your needs, we'll work closely with you to design one that does. If you're designing a system, we'll assist you. Whether your design requires point-to-point positioning, tracking, mapping, processing or something more exotic, we can help.

Nor does it matter to us if you're a Fortune 500 company or a considerably more modest enterprise. You can count on us to respond to your needs with the same enthusiasm and thorough professionalism.

Let us help you put the human touch in your equipment. Talk to us. Write or call:

MS Measurement Systems, Inc.

121 Water Street, Norwalk, CT 06854, U.S.A. Phone 203-838-5561

CIRCLE 88

Personal test tool provides portable, standalone emulation



DA6502-B can support and service 6502-based systems. It provides engineers with a portable, standalone emulator to use in test and development. Weighing less than 2 lb, the DA6502-B consists of a terminal keyboard and display, emulator logic and program, a 6502 microprocessor, a bus switch, and power supply. Its 24 functions can be invoked by one- and two-key-stroke sequences. It accepts simple hexadecimal data and address entry. Pricing starts under \$600, with quantity discounts available. **Da-Tech Corp.**, 92 Steamwhistle Dr, Ivyland, PA 18974. **Circle 299**

Hierarchical logic design marks system using AI techniques

The CDA 5000 design automation system uses a portable language as well as "complexity management techniques" developed in AI research. The system supports hierarchical logic design. Its integrated data base allows rule checking on-the-fly with quick availability of the simulator and component creator. Also, the simulator can highlight signals in the logic description as the simulation progresses. **Ceracor, Inc.**, 716 E 4500 St S, Salt Lake City, UT 84107. **Circle 300**

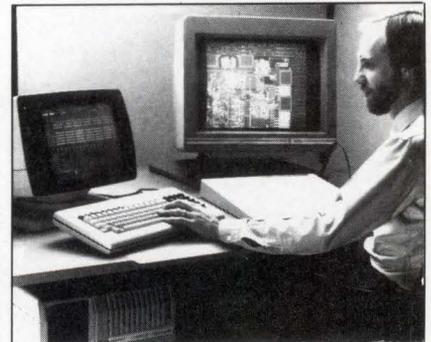
Universal logic memory programmer gets new software version

Software Release 3.16 runs on the Omni-Programmer, a software-driven universal logic/memory device designed for MS-DOS or CP/M systems. Release 3.16 allows the operator to program any Advanced Micro Devices logic/memory device. Makers claim that four PROMs can be burned in the time formerly required to program one device. The Omni-Programmer now tests and programs the

32R16 and the 32RP16, as well as more than 500 other chips. **Varix Corp.**, 1210 E Campbell Rd, Richardson, TX 75081. **Circle 301**

Workstation for computer aided design streamlines IC layout

The GDSII/32 workstation is a 32-bit CAD machine designed for the layout of ICs and printed circuit boards. Powered by a DS/4200 computer, the GDSII/32 links to larger GDSII systems. Its configuration allows the designer to perform circuit layout operations and access the main system for computationally intensive "background" operations such as design rule checks and mask pattern generation. The unit includes a high resolution (1024 x 1280 pixel) 60-Hz non interlaced color monitor, 2-Mbyte main memory, two 5¼-in. Winchester 105-Mbyte disk drives, and hardware floating point unit. With software, the price is about \$95,000 per station. **Calma Co.**, 2901 Tasman Dr, Santa Clara, CA 95050.



Circle 302

DYNAMIC MEMORY TIMER

- Triggered by a Single MREQ Pulse
- Rising-edge Triggered
- TTL Interfaced
- 14 Pins DIP
- Generates All Timing Pulses for DRAM Memory Board
- Precise and Stable Timing

data delay devices, inc.

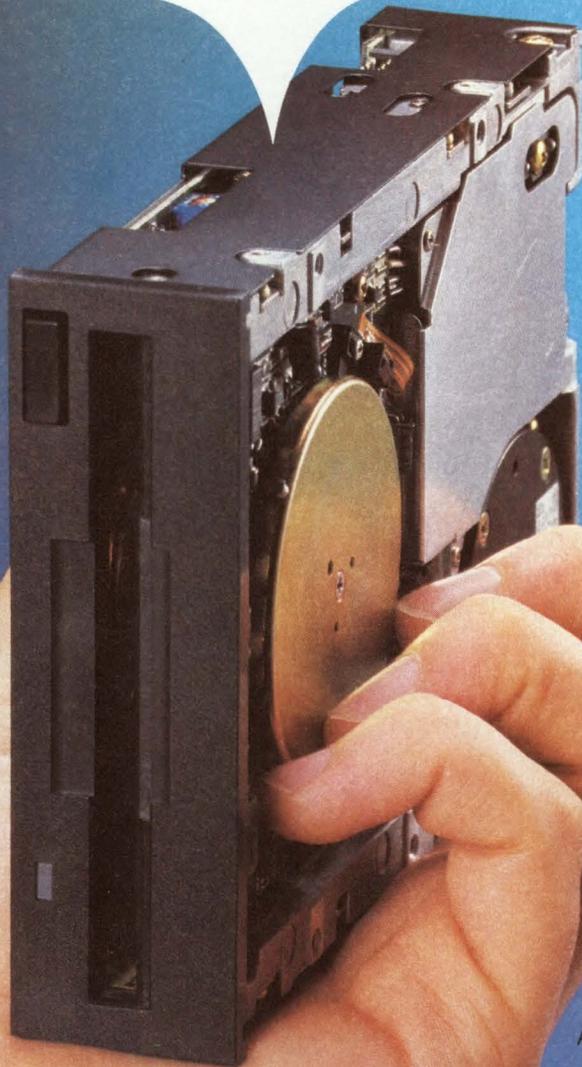
385 Lakeview Avenue, Clifton, New Jersey 07011
(201) 772-1106

Low cost design workstation can be networked

Sun-2/50 desktop stations can run Unix and operate in Ethernets without a local disk. Several diskless machines can share the disks of a Sun-2/12FS or 2/170 file server. Demand paging is handled over the Ethernet to mass storage on the file server. Based on the 68010, the systems' patented memory management design delivers 1 to 4 Mbytes of physical memory and up to 16 bytes of virtual address space per process. A hardware floating point accelerator is optionally available. A 19-in. display features 1152- x 900-pixel resolution and 66-Hz noninterlaced operation. Pricing starts at \$9900 for a 1-Mbyte system, with volume discounts available. **Sun Microsystems, Inc.**, 2550 Garcia Ave, Mountain View, CA 94043. **Circle 303**

Thinking 3½ inches? Think **Canon.**

THE NEW SMALL STANDARD.



Shown
Actual Size

Canon is known as the master of the ultraprecise miniature mechanism. And the smaller disk drives become, the more essential such mastery becomes.

The new Canon 3½-inch microfloppy disk drives are the new small standard size. 30mm high without the bezel. 105mm wide. Weight: 550 grams.

Available now in any quantity.

Our MD350 records 135 tracks per inch, with a track-to-track access time of 3ms and an unformatted capacity of 1 Mbyte. Yet our soft read error rate is only 10^{-9} bits, and our hard read error rate a mere 10^{-12} bits.

Three critical parts designed and manufactured by Canon contribute significantly to our drives' accuracy, and to their extraordinary reliability: the extremely durable brushless drive motor; the head assembly, which uses advanced video head technology; and the steel belt positioning system.

Other good reasons to use Canon's 3½-inch drives:

- Simple pushbutton operation (now industry standard, pioneered by Canon).
- Available without bezels, or with any color bezel.
- Compatible with 5¼-inch drives. Can be used in your present system without changing electronics.

Four models are available:

MD350: 135 TPI, double sided, double density, 1 Mbyte.

MD351: 135 TPI, single sided, double density, .5 Mbyte.

MD352: 67.5 TPI, double sided, double density, .5 Mbyte.

MD353: 67.5 TPI, single sided, double density, .25 Mbyte.

Canon

Call us for even more reasons why 3½-inch drives are Canon's corner. (516) 488-6700, Ext. 4958. Or write Canon U.S.A., Inc., Disk Drive Division, One Canon Plaza, Lake Success, NY 11042.

CIRCLE 90



Esprit's out first with a breakthrough in DEC/ANSI terminal price/performance. The ESP 6115—only \$560.

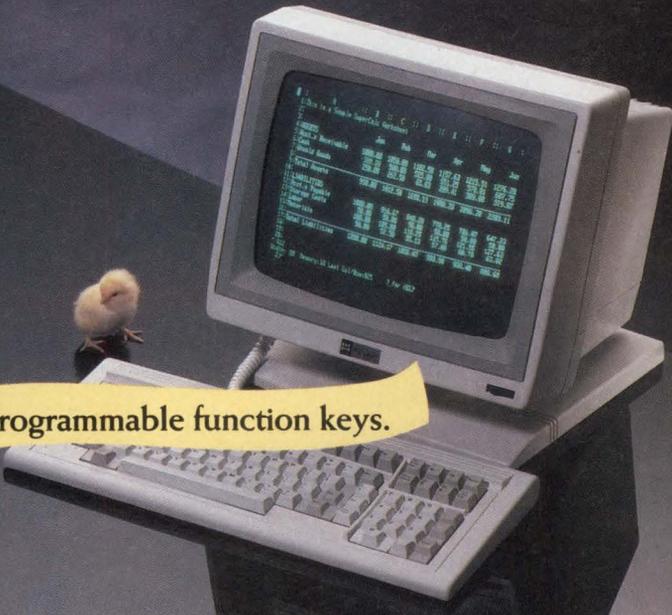
The price is tiny, the screen, a giant 14" \$560 is a great price for *any* terminal. But for a DEC®/ANSI with 16 programmable function keys and a bigger-than-DEC 14" screen, it's chicken feed. And *our* programmable keys are non-volatile.

A born performer. Why buy more ANSI terminal than you need? The ESP 6115 is ideal for inquiry-response applications. Easy reading 9 × 12 dot matrix. Low profile, DIN-standard keyboard. Non-glare, high resolution screen. Famous Esprit quality.

You can choose an optional amber display or add complete IBM® PC capability with our ESP 9310 module.

For details, call Esprit (Es-free) toll-free at 800-645-4508. In NY State 516-293-5600. Or write us at 100 Marcus Drive, Melville, NY 11747.

®Registered trademarks: DEC, Digital Equipment Corporation. IBM, International Business Machines Corporation.



16 programmable function keys.

Esprit
Systems, Inc.

PEACE OF MIND TECHNOLOGY

CIRCLE 91

Imaging

Your image is important to us.

Your system's CRT display image reflects the quality of your system's capabilities. Your data display monitor must be as capable as your system design. Audiotronics designs and manufactures CRT displays for OEMs and systems integrators for use in commercial, industrial, military and medical applications. Many Audiotronics displays are available from inventory.

Sizes are only part of the picture.

Whether your requirements are color or monochrome, TTL or composite, Audiotronics meets these needs... and much more. Screen sizes range from 3" to 23". Configuration options include: high resolution • dynamic focus • external contrast and brightness • changes in scanning frequencies • a wide range of phosphor and non-glare treatments • switching power supplies with multi-tap voltages • dual level TTL.

The other side of the picture.

Audiotronics offers a patented, integrated CRT neck mounted PCB format with a plug-in printed circuit board for 7", 9", 12" and 14" displays. This exclusive design meets the requirements of systems designers needing a compact, lightweight, easily-serviced configuration.

Where neck mount is inapplicable, our other available configurations include the traditional open frame chassis or kits using compact PCB designs. Audiotronics offers moderately-priced, general purpose displays as well as highly specialized custom displays offering resolution capabilities greater than 1000 lines.

We build the big picture for you.

Unique engineering and manufacturing capabilities have allowed us to successfully implement our total product package concept for the OEM customer. We will design and manufacture your terminal, monitor, keyboard, power supply or other products and components to your specifications at either our domestic or Far East operations.

You are already in the picture.

Chances are that Audiotronics already affects your life. When you fly, your carry-on luggage may be screened by systems using our displays. Audiotronics is the leading choice of common carriers who electronically display arrival and departure information. In

your doctor's office or medical lab, our displays are often used in blood analyzers.

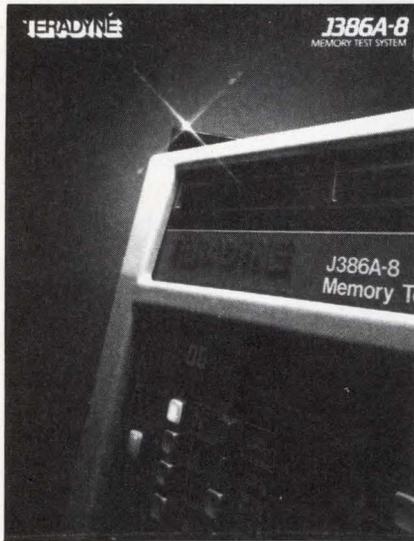
Manufacturers use our displays in numerical control equipment for a variety of on-line applications. And we fit the picture for business and personal computers and data terminals. Audiotronics has a record of achievement in providing OEM customers with dependable CRT data displays for critical applications.

We are interested in contributing to your projects. Contact us today concerning *your* data display needs.



 **AUDIOTRONICS**
North Hollywood,
California 91605
® (818) 781-6700
The Image Engineering Company

Memory test system



The J386A-8 eight-in-parallel memory test system is explained in a five-color brochure. System overview plus software and hardware capabilities are given. **Teradyne Inc**, Boston, Mass. **Circle 410**

Microprocessor software

A comprehensive software catalog details application, development, and operating system software available for the 32000 32-bit microprocessor family. This 138-page catalog covers 130 products and has cross-referenced listings. Genix, Exec, and other operating systems are incorporated. **National Semiconductor Corp**, Santa Clara, Calif. **Circle 411**

Programmable controllers

Model 2008 programmable controllers are shown in a six-page, full-color catalog. The controllers have 8 outputs and 12 inputs (with an additional four analog inputs that can be used for position and process sensing) as well as remote time delay adjustment. **Control Technology Corp**, Hopkinton, Mass. **Circle 412**

A VAX local net

An illustrated, four-page brochure, accompanied by a data sheet, provides information and specs on Inplant, a local area network for DEC VAX minicomputers and similarly sized computers. Inplant, designed for factories, can connect 240 or more devices to each VAX on the network using fiber optics, Ethernet cable, or broadband cable. **Incyte/Schlumberger**, Los Altos, Calif. **Circle 413**

Surface-mount technology

An 112-page manual provides a complete overview of surface-mount technology, including terminology, connection processes, testing and reliability considerations, and the current spectrum of surface-mounted products. **Texas Instruments Inc**, Dallas, Tex. **Circle 414**

Macrocell arrays

A design manual covers MCA2500ECL macrocell array topics. The 67-page manual aids designers involved in creating new systems on silicon using array technologies. It examines logic design considerations and performance guidelines in detail. **Motorola Semiconductor Products Inc**, Phoenix, Ariz. **Circle 415**

Dot-matrix line

Data sheets describe a line of serial dot-matrix printers that covers the speed ranges of 60 to 540 chars/s and 10 to 16.5 chars/in. Also, a shortform catalog is available. **Anadex**, Camarillo, Calif. **Circle 416**

Office system news

Benchmark covers office systems technology and user applications. A quarterly, the magazine will feature such subjects as LANs, electronic publishing, document management, AI, and laser printing. **Xerox Corp**, El Segundo, Calif. **Circle 417**

Indicator lights

Pilot and indicator lights are offered in a 29-page catalog covering diverse sizes and styles. Listing is by series designation, panel depth and thickness, hole size required, mounting method, lens size and style, material, and termination. **Sorenson Lighted Controls Inc**, Hartford, Conn. **Circle 418**

Multiplexers and more

Data communications products ranging from multiplexers to network concentrators are the focus of an in-depth product brochure. **General DataComm Industries Inc**, Middlebury, Conn. **Circle 419**

System enhancements

A full-color catalog features 36 products designed to complete a user's system configuration. Processors, peripherals, terminals, and communication interfaces are contained in the catalog. **Data General, Special Systems Group**, Southboro, Mass. **Circle 420**

Power supplies

Fifth generation UPS systems are shown in a descriptive bulletin. The UPS series, featuring a 313-kVA, 250-kW unit, is illustrated with block diagrams and outline drawings. **Cyberex Inc**, Mentor, Ohio. **Circle 421**

Network architecture

A pamphlet looks at a line of devices for interconnecting 3 to 50 computers, terminals, instruments, modems and printers. Complete with application schematics, the 6-page brochure describes the merging of these products in a simple network. **Digital Products Inc**, Watertown, Mass. **Circle 422**

Data conversion and DSP

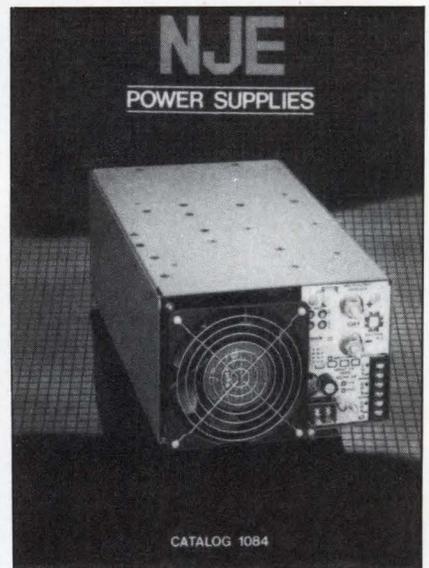
Products for data conversion and signal processing are described in a 164-page catalog. Descriptions, functional diagrams, and specs are included. **Hybrid Systems Corp**, Billerica, Mass. **Circle 423**

Silicon compound reference

A comprehensive tabulation of silicon compounds available for electronic applications is contained in a 240-page catalog entitled, *Silicon Compounds: Register & Review*. **Petrarch Systems Inc**, Bristol, Pa. **Circle 424**

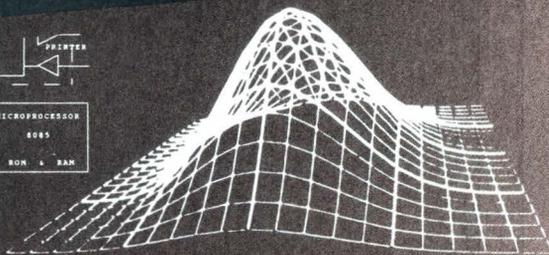
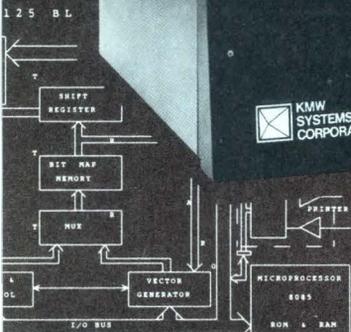
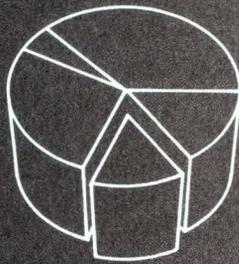
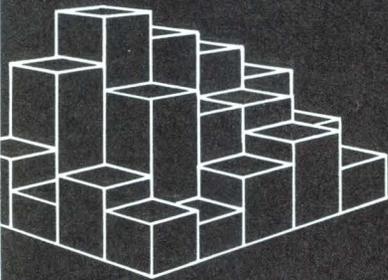
Supply options

This four-color catalog features information on a recently introduced multi-output switcher. It offers technical specs online of linear, ferroresonant, and switcher-regulated power supplies. **NJE Corp**, Dayton, NJ.



CATALOG 1084
Circle 425

Systematic Solutions for Graphic Data Conversion and Raster Hardcopy Device Control



More *Graphics Magic* from KMW Systems

KMW Systems Corporation offers a comprehensive range of controllers and subsystems for solutions to problems associated with the use of raster hardcopy devices. KMW VP series controllers, field proven communications, direct attach interfaces and supplied software interfaces provide the complete graphic subsystem.

VP-10 Graphic Element Processor

The VP allows the connection of low-cost raster hardcopy devices to a host computer.

FEATURES

- emulates HP, Tektronix or Calcomp controllers or accepts standard KMW format
- support for most popular hardcopy raster output devices including the Tektronix 4691, the Versatec V-80, the Seiko CH-5201 and Benson plotters.
- storage capacity for up to 200,000 vectors/elements
- extensive internal diagnostics

VP-20 Vector Processor

The VP-20 banded Vector Processor provides ultra high speed graphic element to raster conversion.

FEATURES

- provides trace sequential data presentation
- support for electrostatic printer/plotters including Calcomp, Benson and Versatec
- two output device option available

VP-30 Vector Processor

The high speed VP-30 provides an ideal solution for controlling wide or dense electrostatic printer/plotters. Its capacity allows it to be used for very large plotting jobs. The VP-30 can be operated in either a dedicated (unattended) mode or as fully interactive graphics workstation.

FEATURES

- accepts Calcomp 906/907, 925, 960, 965 or KMW standard data format
- vector/element storage capacity of 1 to 4 million vectors
- support for electrostatic printer/plotters including Calcomp, Benson, Versatec (monochrome and color)

Communications/Interface Options

KMW manufactures a complete line of field proven communications devices including IBM 2780/3780, HASP and 3770 protocol converters. Direct attach interface options available include 3287 Coax, IBM and Univac Channel Interface, IEEE-488, RS232-C and Dataproducts parallel.

If cost-effective graphics are a problem, take the systematic approach. Let KMW apply a little Graphics Magic.



**KMW
SYSTEMS
CORPORATION**

8307 Highway 71 West • Austin, Texas 78735
Call toll free 800/531-5167
TWX: 910-874-2005 / CABLE: KMWSYS AUS

CIRCLE 98

MARKET FORECASTS

Semicustom ICs

The semicustom chip industry has not experienced the down cycle felt by the general semiconductor business, according to the Technology Research Group. In a recent issue of the *Technology Research Letter*, the group estimates that new designs of semicustom ICs increased nearly 30 percent from the second to the third quarter of 1984, while the industry as a whole (according to the Semiconductor Institute of America) reported a 23-percent drop in book-to-bill ratio. The Technology Research Group reports that the backlog in the semicustom chip industry is growing. However, the group notes that sales of semicustom chips are not subject to the corrections for overstock and double ordering that push down book-to-bills for the standard semiconductor players. The personal computer sales slowdown caused some erosion in

semicustom orders, the group says, but this was more than offset by design and production of new chips aimed at healthier markets. Design services are emerging as a promising market. These services are soon expected to account for 30 percent of the semicustom-IC market revenues. The monthly *Technology Research Letter* costs \$495 per year. **The Technology Group**, 50 Staniford St, Boston, MA 02114.

Circle 469

Briefcase-size portables

Briefcase-size computers will dominate the portable sector by 1988, according to a report available from IPI, Inc. The group indicates that the briefcase portable market is the fastest growing segment of the portable industry, increasing at an annual rate of 75 percent over the course of

the decade. Briefcase shipment levels are expected to reach 1.6 million units, with a value of \$2.6 billion, in 1987. IPI sees the total portable computer market achieving sales revenues of \$12 billion in 1990, with briefcase models accounting for the lion's share. The report foresees an average annual growth for handheld computers of 23 percent through 1990. Suppliers are expected to ship more than 1.8 million units of these products with a value of \$383 million in 1987. Developments in flat-panel display and CMOS technology will fuel briefcase-size computer growth. The report, *Portable Computer Markets 1984 to 1990*, is available for \$795. **IPI, Inc.**, 164 Pecora Way, Portola Valley, CA 94025.

Circle 470

Keyboard shipments

Computer use will account for 25.6 percent of the \$768.8 million market Venture Development Corp forecasts for keyboards in 1987. Instrumentation applications loom as an expansive area for these components, accounting for 23 percent, or \$176.8 million, of the 1987 market. Instrumentation represented about 20 percent of the 1982 applications. Process and factory control, medical instrumentation, as well as test and measurement stand out among instrument end uses. Adaptation to extreme conditions and operator abuse, VDC says, will distinguish the winners in this sector. **Venture Development Corp.**, 1 Washington St, Wellesley, MA 02181.

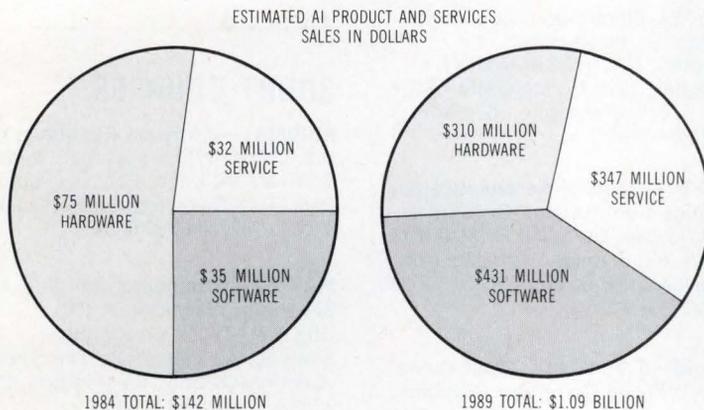
Circle 471

Artificial intelligence

Last year, sales of artificial intelligence systems vaulted, according to Business Communications Co, Inc. In the study *Artificial Intelligence: Current and Future Commercialization*, the firm puts the 1983 market at \$53 million and the 1984 market at \$142 million. A total market size of \$1.1 billion is anticipated for 1989. Hardware represents an early driving force for AI. Sales of dedicated-Lisp machines, BCC claims, accounted for 51 percent (\$75 million) of total AI products and services in 1984. Four major components comprise this market: natural languages, software utilities, plus framework and application software. As of 1984,

natural language software is credited with \$16 million in sales, promising to grow at high rates as manufacturers employ it to attract more end users. By 1989, natural languages will achieve \$200 million revenues. BCC says that utility software now represents about \$10 million in sales and is expected to grow at 44 percent per year. Now in their infancy, commercial frameworks should form a \$100 million market by 1989. The AI service market is also seen as a strong growth area. The 241-page report is available for \$1750. **Business Communications Co, Inc.**, Box 2070C, Stamford, CT 06906.

Circle 473



SOURCE: BCC, INC

Uninterruptible power supplies

Surge suppressors will comprise about 15 percent of total uninterruptible power supply/power line conditioner equipment sales in 1995, according to International Resource Development Inc. This 1995 surge suppressor share, up from 11 percent in 1985, should represent about \$310 million in value. IRD notes that surge suppressors, used primarily for microcomputers, will experience a significant increase in unit sales over the next decade, but that the dollar volume of this low end product segment will lag behind the high end sector. Strong high end applications considered in the 131-page report include airline and other reservation systems, credit card authorization centers, and online data processing centers. UPS demand will continue to rise, IRD states, due to the ongoing deterioration of commercial power quality. The report costs \$1285. **IRD**, 6 Prowitt St, Norwalk, CT 06855.

Circle 472



CONFERENCES

APR 1-2—IEEE VLSI Test Workshop, Bally Park Place Hotel, Atlantic City, NJ. INFORMATION: Bob Tighe, IBM Dept 69J/422, Neighborhood Rd, Kingston, NY 12401. Tel: 914/385-7440

APR 2—Invitational Computer Conf, Sheraton Tara Hotel, Nashua, NH. INFORMATION: Suzanne Hubner, B. J. Johnson & Assoc, 3151 Airway Ave, Costa Mesa, CA 92626. Tel: 714/957-0171

APR 2-4—IEEE Microprocessor Forum, Bally Park Place Hotel, Atlantic City, NJ. INFORMATION: Helen Yonan, IEEE Office, Moore School of Electrical Engineering, Univ of Pennsylvania, Philadelphia, PA 19104. Tel: 215/898-8106

APR 14-18—Computer Graphics Conf, Dallas Conv Center, Dallas, Tex. INFORMATION: National Computer Graphics Assoc, 8401 Arlington Blvd, Ste 601, Fairfax, VA 22031. Tel: 703/698-9600

APR 16-18—Computer Integrated Mfg and Communications Conf, Disneyland Hotel, Anaheim, Calif. INFORMATION: CASA/SME Public Relations, One SME Dr, PO Box 930, Dearborn, MI 48121. Tel: 313/271-0777

APR 17-24—Hannover Fair '85, Hannover, West Germany. INFORMATION: Hannover Fairs Information Center, PO Box 338, Whitehouse, NJ 08888. Tel: 201/534-9044

APR 23-25—Electro and Mini/Micro Northeast, Coliseum and Sheraton Center, New York, NY. INFORMATION: Dale Litherland, Electronic Conventions, Inc, 8110 Airport Blvd, Los Angeles, CA 90045. Tel: 213/772-2965

APR 29-30—Manufacturing Automation Protocol (MAP) Users Group Meeting, Fairmont Hotel, Dallas, Tex. INFORMATION: Computer and Automated Systems Assoc of the Society of Manufacturing Engineers, PO Box 930, Dearborn, MI, 48121. Tel: 313/271-1500

APR 29-MAY 2—Intermag Conf, Radisson St Paul Hotel, St Paul, Minn. INFORMATION: J. H. Nyenhuis, Dept of Electrical Engineering, Purdue Univ, West Lafayette, IN 47907. Tel: 317/494-3524

APR 29-MAY 2—Society for Information Display Symposium & Exhibit, Sheraton Twin Towers, Orlando, Fla. INFORMATION: Hildegard Hammond, 201 Varick St, New York, NY 10014. Tel: 212/620-3388

APR 30-MAY 2—Artificial Intelligence and Advanced Computer Technology Conf/Exh, Long Beach, Calif. INFORMATION: Tower Conference Management Co, 331 W Wesley St, Wheaton, IL 60187. Tel: 312/668-8100

MAY 6-9—Comdex Spring, Georgia World Congress Center, Atlanta, Ga. INFORMATION: The Interface Group, Inc, 300 First St, Needham, MA 02194. Tel: 617/449-6600

MAY 13-17—Int'l Conf on Distributed Computing, The Fairmont Hotel, Denver, Colo. INFORMATION: Dr Earl Swartzlander, TRW, Defense Systems, 1 Space Park, Redondo Beach, CA 90278. Tel: 213/535-4177

MAY 14-16—Test & Measurement World Expo, San Jose Convention Center, San Jose, Calif. INFORMATION: Meg Bowen, 215 Brighton Ave, Boston, MA 02134. Tel: 617/254-1445

MAY 20-22—Custom Integrated Circuits Conf, The Portland Hilton Hotel, Portland, Ore. INFORMATION: Dr Wesley N. Grant, Sperry Computer Systems, Sperry Park, PO Box 43525, MS Y11B1, St Paul, MN 55164. Tel: 612/456-4130

MAY 20-22—Electronic Components Conf, Capital Hilton Hotel, Washington, DC. INFORMATION: Tom Pilcher, 3029 E Washington St, PO Box 372, Indianapolis, IN 46206. Tel: 317/261-1592

MAY 20-24—National Aeronautics and Electronics Conf, Dayton Convention Center, Dayton, Ohio. INFORMATION: NAECON '85, Samuel J Rosengarten, 1445 Devoe Dr, Beavercreek, OH 45385. Tel: 513/255-4709

MAY 21-22—Trends and Applications Conf on Utilizing Computer Graphics, IEEE Computer Society and National Bureau of Standards, Sheraton NW Washington Hotel, Silver Spring, Md. INFORMATION: Mark Skall, NBS, Technology Bldg, Room A265, Gaithersburg, MD 29899. Tel: 301/921-2431

JUNE 3-5—Eastern Design Engineering Show, Bayside Expo Center, Boston, Mass. INFORMATION: Cahners Expo Group, PO Box 3833, Stamford, CT 06905. Tel: 203/964-0000

JUNE 3-6—Robots 9, Cobo Hall, Detroit, Mich. INFORMATION: Public Relations Dept, Society of Manufacturing Engineers, PO Box 930, Dearborn, MI 48121. Tel: 313/271-0777

JUNE 9-13—Computer Vision and Pattern Recognition, Cathedral Hill Hotel, San Francisco, Calif. INFORMATION: Harry Hayman, PO Box 639, Silver Spring, MD 20901. Tel: 301/589-8142

JUNE 17-19—Int'l Conf on Computer Architecture, Boston Park Plaza Hotel, Boston, Mass. INFORMATION: Computer Architecture, PO Box 639, Silver Spring, MD 20901. Tel: 301/589-8142

JUNE 19-21—NEPCON East, Bayside Expo Center, Boston, Mass. INFORMATION: Cahners Expo Group, PO Box 5060, Des Plaines, IL 60018. Tel: 312/299-9311

JUNE 23-26—Design Automation Conf, Caesars Palace, Las Vegas, Nev. INFORMATION: Hillel Ofek, 1080 Marsh Rd, Menlo Park, CA 94025. Tel: 415/324-0700

SHORT COURSES

APR/MAY—Software Seminars, various U.S. cities. INFORMATION: On-Line Software Int'l, Inc, Fort Lee Executive Park, Two Executive Dr, Fort Lee, NJ 07024. Tel: 201/592-0009

MAY 2-3—Grounding, Bonding, and Shielding, Washington, DC. INFORMATION: Continuing Engineering Education, The George Washington Univ, Washington, DC 20052. Tel: 202/676-6106

MAY 20-21—Introduction to Digital Image Processing, Houston, Tex. INFORMATION: Perceptive Systems, Inc, 1301 Regents Park Dr, Ste 103, Houston, TX 77058. Tel: 713/480-8441



DESIGNER'S BOOKCASE

ENCYCLOPEDIA OF INTEGRATED CIRCUITS: A PRACTICAL HANDBOOK OF ESSENTIAL REFERENCE DATA

By *Walter H. Buchsbaum*

This remarkable reference condenses volumes of manufacturer's data into one, easy-to-use handbook. It gives full details on even the most advanced integrated circuits. Clear illustrations of simple logic and block diagrams plus a list of key parameters and applications help you to capitalize on the most recent developments in microprocessors and various supporting chips.

\$24.95 **Circle 445**

MICROELECTRONICS: A STANDARD MANUAL AND GUIDE

By *John Douglas-Young*

Here you'll find clear, concise explanations about what integrated circuits are, how they operate and why trouble-shooting ICs differs from locating defects in conventional circuits. You'll also find complete data, with practical applications, for flip-flops, counters, bubble memories, I/O devices, MOSFETs, converters, opto-isolators and more.

\$21.95 **Circle 446**

THE AI BUSINESS: COMMERCIAL USES OF ARTIFICIAL INTELLIGENCE

Edited By *Patrick H. Winston* and *Karen A. Prendergast*

This book explains what Artificial Intelligence is and how the technology has moved out of university research laboratories and into real-world applications. It describes systems in the computer industry that help configure, sell and install computers and it discusses intelligent workstations that increase the productivity of electronic design engineers.

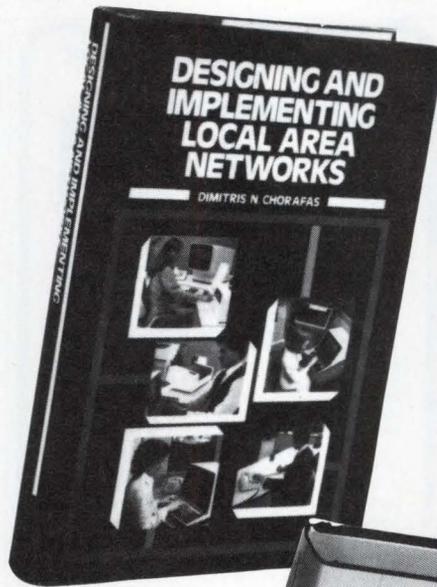
\$15.95 **Circle 447**

DESIGNING AND IMPLEMENTING LOCAL AREA NETWORKS

By *Dimitris N. Chorafas*

Local area networks allow users to hook up their office automation systems to centralized information resources. This practical book shows how to design and use such a system and how it can improve productivity and job performance. The book emphasizes both the hardware and software of local area networks, the databases that users can access online through distributed information systems, and the importance of bringing computer power to the end user's desk.

\$35.95 **Circle 448**

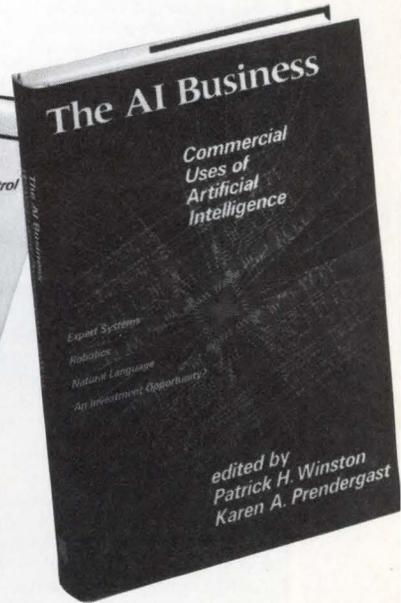
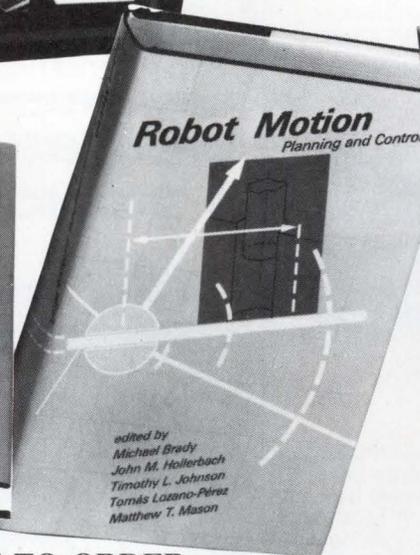
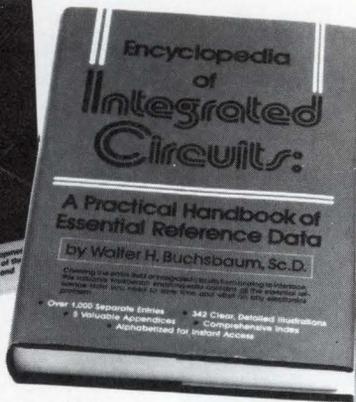
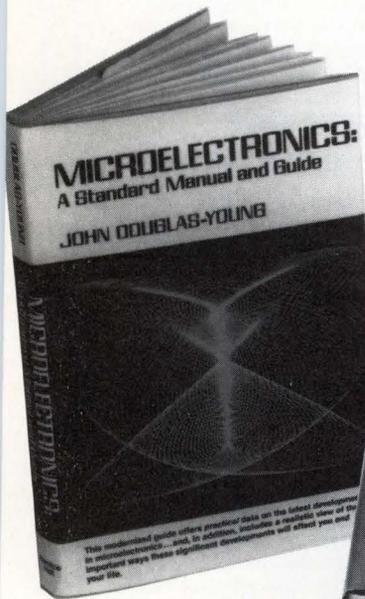


ROBOT MOTION: PLANNING AND CONTROL

Edited By *Michael Brady, John Hollerbach, Timothy Johnson, Tomas Lozano-Perez* and *Matthew Mason*

Nine papers of fundamental importance to the development of robotics—for researchers and engineers in industry and academia—are grouped in five sections: Dynamics; Feedback Control; Trajectory Planning; Compliance and Force Control; Task Planning. Each section begins with a substantial analytic survey.

\$39.50 **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

CALL: SHIRLEY LESSARD

800-225-0556; in MA 617-486-9501



**INTELLIGENT
STAND
ALONE
RS-232-C
or
IEEE-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 with data rates up to 3,000 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-pollled remote data storage.
- Price: Under 2000 in quantities

9198-C Red Branch Rd.
Columbia, MD 21045
301-730-7442
Telex #333405 ALGO INC.

ALGO

CIRCLE 476

**Simulator/Debugger for IBM-PC
for 8048 - 8051 - 7000 μ Cs**

Execute and debug code for popular single chip micro-computers on your IBM PC. Dynamic display with windows for source code, control flow, registers, flags, memory, commands, and more! Set breakpoints, traps, etc. Cross-assembler & EPROM programmers, too.

	8048	8051	7000	Z8	
X-Assembler	✓	✓	✓	✓	\$295
EPROM Prog	✓	✓	✓	✓	\$245
Simulator	✓	✓	✓	✓	\$595

Demo disk and manual for simulator only \$39.50

Cybernetic Micro Systems
Box 3000 • San Gregorio, California 94074 U.S.A.
(415) 726-3000 • Telex 171-135 Attn: Cybernetics

CIRCLE 477



**FOR
SALE!**

Like new
Intel Series III
286FD Development
System
\$11,900.00

*50% LESS than
Intel list!*

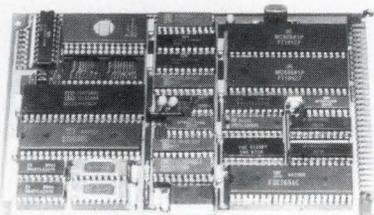
FULL 120 DAY WARRANTY — Includes DX225 Microcomputer Development System and an MDX-720 Dual Double Density Flexible Disk Subsystem with an iMDX-556 Resident Processor Card. System software will be Intel's most current revision level. Ask about FREE MEMORY with every 286 purchased. Subject to prior sale. Phone (800) 225-1008, in Mass (617) 938-0900, or (800) 227-8409, in Calif. (818) 887-4000.

GENSTAR

REI Sales Company

6307 De Soto Avenue, Suite J
Woodland Hills, CA 91367

CIRCLE 478



CMOS Z80A CPU

64K RAM, up to 128K bytes EPROM. DMA support of any mix of 4 floppy drives (8, 5.25, 3.5 inch). EPROM based debugger and standard OS available. 4 Serial ports, nvRAM configuration tables and diagnostic LEDs. Available without FDC for EPROM based applications. 100 by 160mm Euro-card. **Heritage Systems Corporation, P.O. Box 10588, Greensboro, NC 27404 919/274-4818.**

Z80A SBC with EPROM based OS

CIRCLE 479

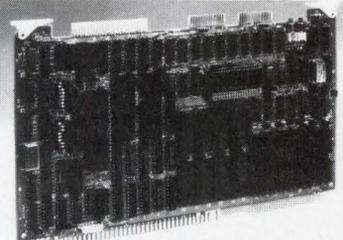
6801 Test Bed & Programmer

Virtual In-Circuit Emulator (VICE) develops single-chip microcomputer programs for the 6801/68701 and programs on-board EPROM. A cost and time effective way to program the 6801/68701 for \$695.

WINTERK

Winterk Corp
1801 South Street
Lafayette, IN 47904
317-742-8428

CIRCLE 480



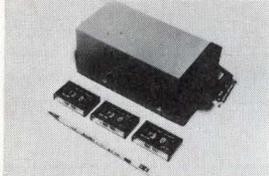
10 MHz Z8000™ PERFORMANCE

In addition to 10 MHz, the NEW MBx8000 SBC features a single width SBX connector for adding standard expansion modules, such as the SBSxFDC Floppy Controller. MBx8000 maintains all the capabilities of its predecessor, the MB8000: Z8001/Z8002 CPU 4/6/10 MHz; 32K/128K/512K DRAM; up to 48K static RAM or EEPROM; up to 128K EPROM; 2 RS-232 ports; 40 parallel I/O lines; six 16-bit counter-timers; etc. Z8000™ Zilog.

SINGLE BOARD SOLUTIONS, 7669 Rainbow Drive, Cupertino, CA 95014. (408) 253-0181.

CIRCLE 481

RS232 or 8 BIT PARALLEL MINI CASSETTE SYSTEM



- LOW COST - \$388 @ 100 UNITS*
- Microprocessor Controlled
- >200K Bytes Formatted/Tape
- Variable Baud Rate (110-9600)
- Error Rate 1 in 10 to the 9th
- Small - 3" x 3.5" x 5" W/O Case*

21 OTHER CARTRIDGE, CASSETTE, FLOPPY SYSTEMS WITH STORAGE TO 12 MEGABYTES Basic Drive with R/W Electronics Start at \$150

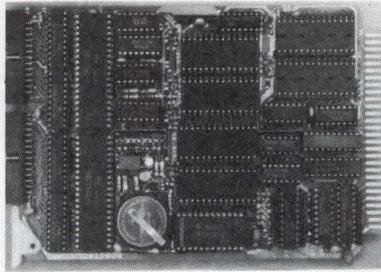
ANALOG & DIGITAL PERIPHERALS INC



815 Diana Drive
Troy, Ohio 45373
513/339-2241
TWX 810/450-2685

Branch Off: Oklahoma City OK - Factory: Yucca Valley, CA

CIRCLE 482



STD CLOCK, MEMORY, I/O BOARD

The RSD-7758 features a battery backed 14 digit real time clock w/prog. interrupt, six 28 pin byte-wide memory sockets, 48 buffered parallel I/O lines. Available options: up to 48K bytes battery backed RAM and 32 high current open collector drivers for parallel I/O. Board is memory mapped and decodes either 16, 17 or 20 address lines. Consult factory for other available configurations. **Robotrol Corp.**, 16100 Caputo Drive, Morgan Hill, CA 95037. (408) 778-0400.

CIRCLE 483

A defense against cancer can be cooked up in your kitchen.

There is evidence that diet and cancer are related. Some foods may promote cancer, while others may protect you from it.

Foods related to lowering the risk of cancer of the larynx and esophagus all have high amounts of carotene, a form of Vitamin A which is in cantaloupes, peaches, broccoli, spinach, all dark green leafy vegetables, sweet potatoes, carrots, pumpkin, winter squash, and tomatoes, citrus fruits and brussels sprouts.

Foods that may help reduce the risk of gastrointestinal and respiratory tract cancer are cabbage, broccoli, brussels sprouts, kohlrabi, cauliflower.

Fruits, vegetables and whole-grain cereals such as oatmeal, bran and wheat may help lower the risk of colorectal cancer.

Foods high in fats, salt- or nitrite-cured foods such as ham, and fish and types of sausages smoked by traditional methods should be eaten in moderation.

Be moderate in consumption of alcohol also.

A good rule of thumb is cut down on fat and don't be fat.

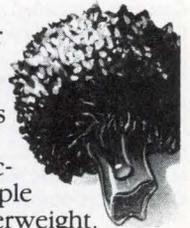
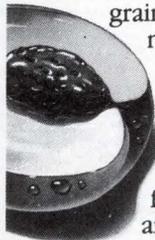
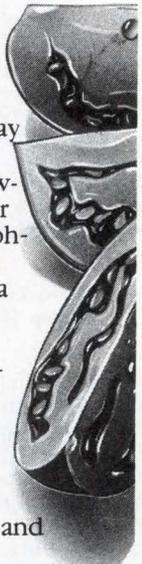
Weight reduction may lower cancer risk. Our 12-year study of nearly a million Americans uncovered high cancer risks particularly among people 40% or more overweight.

Now, more than ever, we know you can cook up your own defense against cancer. So eat healthy and be healthy.

No one faces cancer alone.

AMERICAN CANCER SOCIETY®

THIS SPACE CONTRIBUTED AS A PUBLIC SERVICE



ICs PROMPT DELIVERY!!!

SAME DAY SHIPPING (USUALLY)

OUTSIDE OKLAHOMA: NO SALES TAX

8087-3 Co-Processors \$109.00

DYNAMIC RAM

256K	256Kx1	150 ns	\$ 8.99
128K	128Kx1	150 ns	15.67
64K	64Kx1	150 ns	2.07
64K	64Kx1	200 ns	2.47

EPROM

27256	32Kx8	250 ns	\$26.97
27128	16Kx8	250 ns	10.40
27C64	8Kx8	200 ns	11.87
2764	8Kx8	250 ns	4.97
2732A	4Kx8	250 ns	5.87
2716	2Kx8	450 ns	3.21

STATIC RAM

6264LP-15	8Kx8	150 ns	\$16.97
6116LP-3	2Kx8	150 ns	3.87

INCREASE XT MOTHERBOARD TO 640 Kbyte CONSUME NO EXPANSION SLOTS! \$181.80

QUANTITY ONE PRICES SHOWN

OPEN 6 1/2 DAYS: WE CAN SHIP VIA FED-EX ON SAT.

MasterCard/VISA or UPS CASH COD

Factory New, Prime Parts

MICROPROCESSORS UNLIMITED

24,000 South Peoria Ave. (918) 267-4961

BEGGS, OK 74421

Prices shown above are for February 18, 1985

Please call for current prices. Prices subject to change. Please expect higher or lower prices on some parts due to supply & demand and our changing costs. Shipping & insurance extra. Cash discount prices shown. Small orders received by 8 PM CST can usually be delivered to you by the next morning, via Federal Express Standard Air @ \$6.50!

CIRCLE 484

MATRIXED DIRECT MAIL LIST

MATRIXED DIRECT MAIL LIST

The most Precise, most Versatile, Cross-Matrixed list of computer Based Systems Designers ever offered for bulk or custom selection. You can rent the entire list, or pinpoint the exact engineers you want to reach. NEW 48-HOUR EXPRESS PROCESSING.

For details, contact Bob Dromgoole at **COMPUTER DESIGN**, 119 Russell St, Littleton, MA 01460. Tel: (800)225-0556. In MA: (617)486-9501.

CIRCLE 485

SUPER "K" BIG BOARD

"It's So Fast It Screams"

- * Z80H 80 Mhz Processor
- * Z80A I/O Processor
- * 12 RS 232c Serial Ports
- * 2 Parallel Ports
- * 360K Bytes RAM
- * Hard Disk Controller
- * Floppy Disk Controller (Up to 4 Drives/Contrlr)
- * DMA
- * Cached Directory
- * Seven (7) Users
- * Modified Superfast MP/M Bios
- * 15" x 15"
- * Assembled & Tested \$2495.00

KASCO Systems
11926 Warfield
San Antonio, TX 78216
(512) 342-5755.

CIRCLE 486

This publication is available in microform.

University Microfilms International



Please send additional information

Name _____
Institution _____
Street _____
City _____
State _____ Zip _____

300 North Zeeb Road Dept. P.R.
Ann Arbor, Mi, 48106 U.S.A.

30-32 Mortimer Street Dept. P.R.
London WIN 7RA
England

AD INDEX

Able Computer.....	37	ITT/Pomona Electronics.....	160
ACC.....	C4	Kasco Distributing.....	223
ADE Corp.....	22	Kennedy Co.....	1
Advanced Digital Corp.....	169	KMW.....	218
Advanced Micro Devices.....	66, 67		
Algo.....	222	Lexidata Corp.....	209
AMP.....	122, 123	LSI Logic Corp.....	97
Amphenol Products, an Allied Co.....	102, 103		
Analog & Digital Peripherals.....	223	Maxell Corp.....	79
Asea Hafo.....	81	MDB Systems Inc.....	165
AT&T Technologies.....	42, 43, 68, 69	Measurement Systems.....	211
Audiotronics Corp.....	216	Megatek Corp.....	63
AVX Ceramics Corp.....	47	Mentor Graphics Corp.....	14, 15
		Microprocessors Unlimited.....	223
BICC-Vero Electronics Ltd.....	124	Microtec Research.....	207
Burr-Brown Corp.....	20, 21	Motorola Semiconductor.....	64, 65
Canon USA.....	206, 213	NCR Corporation.....	185
Carroll Touch Technology.....	177	NEC Electronics.....	35, 95
Celanese Engineering Resins.....	189	NEC Information Systems.....	144, 145
Cipher Data Products.....	38, 39	NMB Hi-Tek Corp.....	73
Citizen America Corp.....	99	Northwest Instrument Systems.....	90
Computer Automation.....	30, 31		
Computrol.....	200	Oasys.....	203, 205
Control Data Corp.....	175		
Cybernetic Micro Systems.....	222	Panasonic.....	146
		Pertec Peripherals Corp.....	109
Data Delay Devices.....	212	Plessey Microsystems.....	61
Data I/O.....	197		
Dataram Corp.....	9	Raytheon Graphic Systems.....	227
Digital Equipment Corp.....	4, 5	Robotrol Corp.....	223
Dual Systems.....	137	Rockwell International.....	C2
Duracell International.....	187	Rodime.....	163
		Rolm Corp.....	153
Eagle Magnetic Co.....	101	RTCS.....	6
E + E DataComm.....	85		
Electronic Solutions.....	208	Science Applications.....	101
Emulex Corp.....	18, 19	Signal Technology.....	202
Esprit.....	214, 215	Silicon Systems.....	117
Expoconsul International.....	130-132	Singer/Kearfott.....	228
		Single Board Solutions.....	222
Floating Point Systems.....	111-114	Sky Computers.....	41
John Fluke Manufacturing Co.....	129	Software Career Link.....	225
FMC Corp.....	75-77	Sophia Computer Systems.....	192
Fujitsu.....	82, 83	Spectra Logic Corp.....	27
Fujitsu Microelectronics.....	140, 141		
		TEAC.....	182
General Electric.....	70, 142, 143	Thomas & Betts Corp.....	29
Genesis Microsystems.....	195	Toshiba America.....	158, 159
Genicom Corp.....	88		
Genstar REI Sales Co.....	222	Visual Technology.....	170, 171
Gould.....	49-58		
Gould, Computer Systems Div.....	12	Wintek Corp.....	222
GTCO.....	104	Wyse Technology.....	156, 157
GTE Automatic Electric Labs.....	226		
GVC Microcomputers.....	16	Xebec.....	154, 155
		Zax Corp.....	11
Heritage Systems Corp.....	222		
Heurikon Corp.....	172		
Hewlett-Packard.....	86, 87		
Hitachi America Ltd.....	C3		
HYTEK Microsystems.....	26		
Ikegami Electronics USA.....	45		
Inacom International.....	32		
Inmos.....	17		
Interphase.....	204		

The AD INDEX is published as a service. The publisher does not assume any liability for errors or omission.

COMPUTER DESIGN

SALES OFFICES

Executive Office: 119 Russell St., Littleton, MA 01460 (617) 486-9501 Telex: 883436 ELN: 62649490

National Sales Manager, Frederick A. Stanley
Business Development Manager, Leslie Ringe

Advertising Coordinator, Debra L. Stone
Classified/Recruitment, Shirley Lessard

List Rental, Robert Dromgoole

New England/Upstate New York

Kevin Callahan
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
Pamela S. Fedor
700 River Road
Fair Haven, NJ 07701
(201) 741-5784

Midwest

Terrence J. Rogan
9501 W. Devon, Suite 203
Rosemont, IL 60018
(312) 696-0420

Southwest

Buckley/Boris Associates
1200 S. Post Oak Blvd, Suite 224
Houston, TX 77056
(713) 621-9720
(714) 756-0681

Southern California & Colorado

Buckley/Boris Associates
Tom Boris, John Sabo, Greg Cruse
2082 SE Bristol, Suite 216
Santa Ana, CA 92707
(714) 756-0681

Northern California

Buckley/Boris Associates
Tom Boris, John Sly, Greg Cruse
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) 756-0681

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
PennWell House
39 George Street
Richmond Upon Thames
Surrey TW9 1HY ENGLAND
Tel: 01 948 7866
Telex: 919775 PENWEL G

Norway

Sverre Follaug, Jr.
Studio 14
Viggo Ullmannsgt 14
P.O. Box 250
3701 Skien NORWAY
Tel: 035 22908

France, Belgium and S. Switzerland

Daniel R. Bernard
247, Rue Saint Jacques
75005 Paris FRANCE
Tel: (33) 1-354-5535
Telex: 250 303 PUBLIC PARIS

Holland, Austria, W. Germany, Switzerland & Eastern Europe

Heinz Gorgens
Firma Heinz Gorgens
Parkstrasse 8a
4054 Nettetal 1
FEDERAL REPUBLIC OF GERMANY
Tel: (49) (2153) 89988
Telex: 172153310 HJG

Japan

Sumio Oka & Shigeo Aoki
International Media
Representatives Ltd.
2-29 Toranomon 1-chome
Minato-ku, Tokyo 105 Japan
Tel: 502-0656
Telex: 22633 MEDIAREP

Southeast Asia

Seavex, Ltd.
Steve Marcopoto,
Jim Stanton
400 Orchard Road, #10-01
Singapore 0923
REPUBLIC OF SINGAPORE
Tel: 734-9790
Telex: 35539 SEAVEX RS

Paul Meyer,
Francis Powers
503 Wilson House
19-27 Wyndham St., Central
HONG KONG
Tel: 5-260149
Telex: 60904 SEAVEX HX

Jay G. Seo
Doobee International Ltd.
Centre Building (Byulgwan)
1-11 Jeong-dong, Choong-ku
C.P.O. Box 4557
Seoul, KOREA
Tel: 27117 DOOBEE
Telex: 776-2096/8

Mr. C. S. Huang
Fidelity Development Corporation
342 Fu Hsing South Road
Room 3 10th Floor Section #1
Taipei, Taiwan
REPUBLIC OF CHINA
Tel: 702-4326
Telex: 21776 FIDCO

The Software Professionals' Job Fair

SOFTFAIR™

85

SoftFair is different — specifically designed to help you search in new career directions. For the broadest possible view of your local software marketplace, **SoftFair** brings together technical managers and personnel representatives from 50 or more of each area's most desirable employers. No need for inconvenient phone calls, lengthy cover letters or work-time interviews. Just join us at the site nearest you and see what your experience has earned you at **SoftFair**.

SoftFair '85 is sponsored by **Software Career Link** serving equal opportunity employers.

SoftFair will be held on these days (Monday & Tuesday) from 4 PM to 9 PM at each location:

Dallas, TX
March 11 & 12

Hyatt Regency
at Reunion

Anaheim, CA
March 18 & 19

Sheraton Anaheim

Washington, D.C.
March 18 & 19

Sheraton
Washington (NW)

Northern NJ
March 18 & 19

Hilton Inn
Woodcliff Lake

Boston, MA
March 25 & 26

Park Plaza Hotel

Philadelphia, PA
April 1 & 2

Marriott/City Line

Exciting career opportunities are available for experienced software professionals in a variety of **Software Engineering, Scientific and Business Applications** areas related to **development, support, implementation, training, quality assurance, technical writing, sales, marketing and management.**

A minimum of two years' on-the-job experience is required.

If you can't attend **SoftFair**, send a resume indicating the location in which you are interested for distribution to attending companies. **Software Career Link**, Dept. CD3, 67 South Bedford Street, Suite 400W, Burlington, MA 01803. (617) 229-5813.

Shape The Future Share The Pride

Pioneering has been our business since 1891. Today, we're in the midst of a technological revolution...still leading the way toward shaping the future of modern telecommunications. Digitalized switching and transmission systems, advanced integrated circuits utilizing state-of-the-art Bipolar and CMOS technologies, new dimensions in fiber optics, innovative products and processes...all assimilated to shape the direction of tomorrow by the professionals at GTE.

If you have a background and experience in Electrical Engineering, Computer Science, Computer Engineering, Physics or Chemical Engineering, we invite you to share the pride with GTE by becoming part of tomorrow's telecommunication task force today! Please send your resume or letter of qualifications. An Equal Opportunity Employer, m/f/h/v.

GTE COMMUNICATION SYSTEMS
Human Resources Dept. 5881
2500 W. Utopia Road
Phoenix, AZ 85027



**Communication
Systems**

Working Together To Be The Best

IDEAS We make them happen!

Creating superior software for superior systems is the everyday charter of the Software Systems Laboratory of Raytheon's Equipment Division. From air traffic control, weapons direction, fire control, communications and missile guidance to tracking and data gathering, space surveillance, range instrumentation and laser applications, we have the applications and environment to make your ideas happen. All within one of the fastest growing and most exciting divisions of Fortune-100 Raytheon, a company internationally recognized for excellence in electronics.

Software Systems Engineers

Software engineers at Raytheon are key players in these front-end activities that translate system requirements into software performance specifications. We seek experience and/or interest in the following areas:

- Requirements Analysis
- Software/Hardware Tradeoffs
- Modeling of Critical Performance Characteristics
- Interface Definition
- Data Processing Architecture Definition
- Software Test Specifications
- Performance Specification Preparation

Radar Systems Software Engineers

Radar systems play a key role in Raytheon's product mix, and major software challenges exist in this area. We seek experience in the following areas:

- Radar Data Processing
- Object Classification and Discrimination
- Real-Time Control Systems
- Embedded Computer Systems
- System- and Unit-Level Diagnostics
- Meteorological Applications

Air Traffic Control Software Engineers

Raytheon is currently engaged in several Air Traffic Control program areas, domestic and foreign, including existing and planned enroute and terminal ATC Systems, for both domestic and international applications, aircraft/controller digital communications and dynamic simulation/training. We seek experienced candidates in the following areas:

- Radar Data Processing/Track File Management
- Flight Plan Processing
- Channel Management
- Intercomputer Communications
- Mosaicking

Assignments in the Washington, DC area also available.

Computer Diagnostics Software Engineers

Raytheon has been engaged in the design and development of advanced militarized computers employing the latest in technology for fault tolerant techniques and high speed architectural concepts. Currently under development are several general purpose militarized super minicomputers employing sophisticated built-in test circuitry and high speed parallel pipelined architectures. We seek experience in the following areas:

- Micro and Macro Diagnostics for fault detection and isolation
- Automatic Test Systems
- Intelligent Control Panel System Development

Communications Systems Software Engineers

Raytheon is a leading company in the development of advanced satellite communications terminals employing sophisticated microprocessor technology for processing and system control. We seek experienced candidates in the following areas:

- Antenna Pointing Systems
- Ephemeris Data Processing
- Multi-Access Algorithms
- Built-in Test and Diagnostics

Graphics Software Engineers

Raytheon's Graphics Group develops real-time interactive display software that will comprise the next generation of "user-friendly" man/machine interfaces in radar, sonar, communications, air traffic control and image processing. We seek experienced candidates in the following areas:

- Real-time Software Design (C, PASCAL, FORTRAN or Assembly Language)
- Curvilinear and Raster Graphics Techniques
- Data Communications and Network Protocols
- Signal/Image Processing Technology
- Meteorological Applications

System Programmers

To support our various applications, we are undertaking a major expansion of our System Software Department. We seek experienced applicants in the following areas:

- Operating Systems, especially VMS, UNIX and RSX-11
- Languages, including C, Pascal, FORTRAN, JOVIAL/J73 and Ada
- Development of Microprocessor Cross-Compilers, Assemblers, and Loaders
- Performance Measurement and Prediction
- User Education and Assistance
- Real-time Operating Systems Development
- Improvement of Software Development Techniques

The preceding positions require a degree in engineering or computer science, at least 3 years of directly related experience, a mind open to new ideas, and a willingness to work hard. If you meet these requirements, we will provide you with a stimulating and congenial environment, pay you a competitive salary, and give you as much a responsibility as you are ready to handle. We are located in the Boston suburbs, an area that many consider to be among the finest in the country in which to live and work.

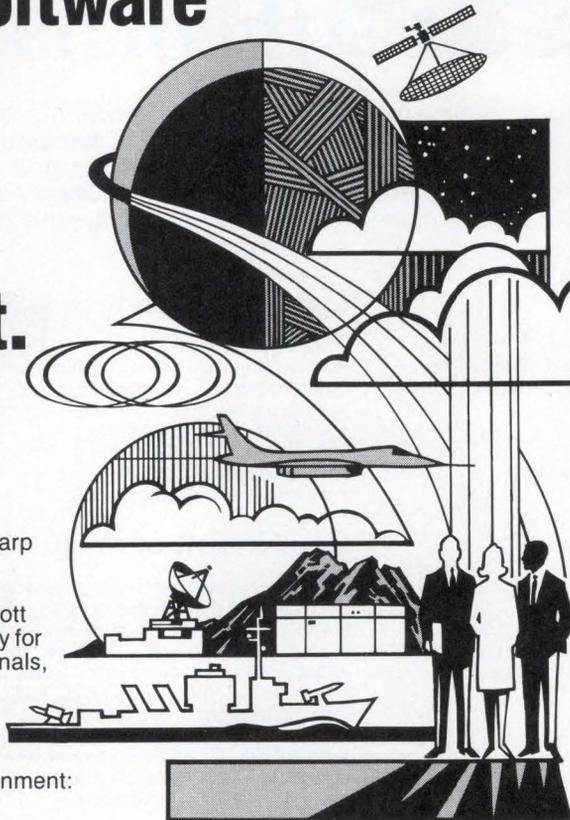
For further information concerning the aforementioned positions, please direct your resume, including current salary, in strictest confidence to: Ann Loven, Dept. 3, Raytheon Equipment Development Laboratories, 528 Boston Post Road, Sudbury, MA 01776. U.S. Citizenship required. An Equal Opportunity Employer.

Raytheon

Computer Software Engineers/
Systems Programmers

Ready for a steep software climb?

Think Kearfott.



Career trajectories, sharp angles of ascent, far-reaching projects and prospects are all Kearfott specialties. Particularly for our software professionals, who are constantly exposed to developing technologies and strategic advances. Think about our environment:

Operating Systems—Developing new system capabilities at our large Data Center, working with IBM and Amdahl CPU's using MVS/SP3, JES 2, IMS, DB/DC, CICS, WYLBUR, TSO, ACF/VTAM and GDDM supporting software. Be at the hub of a system linking all of our plants for instant data transmission.

MIS Development—Designing software that runs the gamut—from sophisticated automated MRP systems for manufacturing planning and inventory controls, to enhancement of our capabilities in accounting, finance, payroll, labor reporting/forecasting—and more. Involves COBOL programming, CICS, MVS, IMS data base design and telecommunications.

Scientific Programming—Developing state-of-the-art realtime systems . . . from advanced guidance and navigation systems using Kalman data mixing, to digital communications systems using modern encryption techniques. Use Fortran, PASCAL, JOVIAL or ADA compilers for powerful minicomputers or a network of microcomputers. Or use mainframe computers to simulate these systems in support of system development validation.

Think Kearfott when it comes to working with an elite corps of software pros, in a company where almost everything . . . including careers . . . flies on software. If you have at least 3 years experience in a large-scale IBM system, or with HP, DEC or airborne computers, or micros in realtime military systems, send your resume to:

**E.J. Murphy, Singer Company,
Kearfott Division, Mail Code
10B05, 1150 McBride Avenue,
Little Falls, NJ 07424.** U.S.
citizenship required. An equal op-
portunity employer, m/f, who
creates opportunities.

SINGER

AMERICA WORKS BEST



WHEN AMERICANS WORK TOGETHER

Teamwork is making a comeback in America.

Cooperation on the production line helped America win World War II. We're in another kind of battle today—a fight for economic survival in an increasingly competitive world market. Cooperation in the workplace is helping us meet this challenge too.

In plants and offices throughout the country, management is asking employees for their ideas on how to increase productivity and improve the work environment. And workers and their unions are responding with a wealth of practical suggestions and a renewed spirit of cooperation.

For information about how others are working better by working together, contact:

**Cooperative Labor-
Management Programs
U.S. Department of Labor
Washington, D.C. 20216
202—523-6098**

U.S. Department of Labor

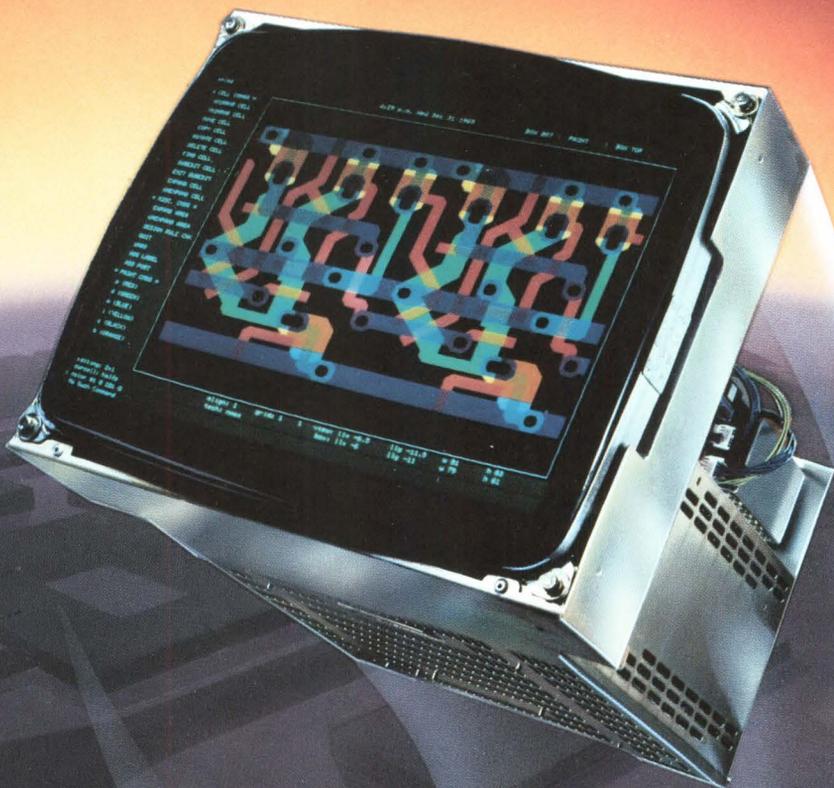


Printed by this publication as a public service
Photograph: Lange Collection, Oakland Museum

the technology is...



the image
speaks
for itself



Hitachi introduces the new standard in Hi-Resolution. That standard is embodied in our new HM-4619 RGB Color Monitor.

The HM-4619 delivers flicker-free images with 1280 x 1024 resolution at 60 Hz. Hitachi tube technology leadership has created a new in-line gun delivering high brightness with sharper focus. All of this is enhanced by our Digital Dynamic Convergence™ of 0.1/0.3 mm.

The Hitachi line of RGB Color Monitors is rapidly becoming the clear choice of OEM's throughout the world. We feel our new, top of the line, HM-4619 is the unit which will set the standard by which all other color monitors will be judged.

- New In-Line Gun with Hi-Brightness and Sharp Focus
- 1280 x 1024 Resolution at 60 Hz
- 0.1/0.3 mm Misconvergence
- 100 MHz Video Bandwidth
- Compact Design
- Cost Effective State of the Art Technology Plus the Strength of Hitachi Worldwide Support



The Image Speaks for Itself

Hitachi America, Ltd.
50 Prospect Ave., Tarrytown, NY 10591-4698
(914) 332-5800

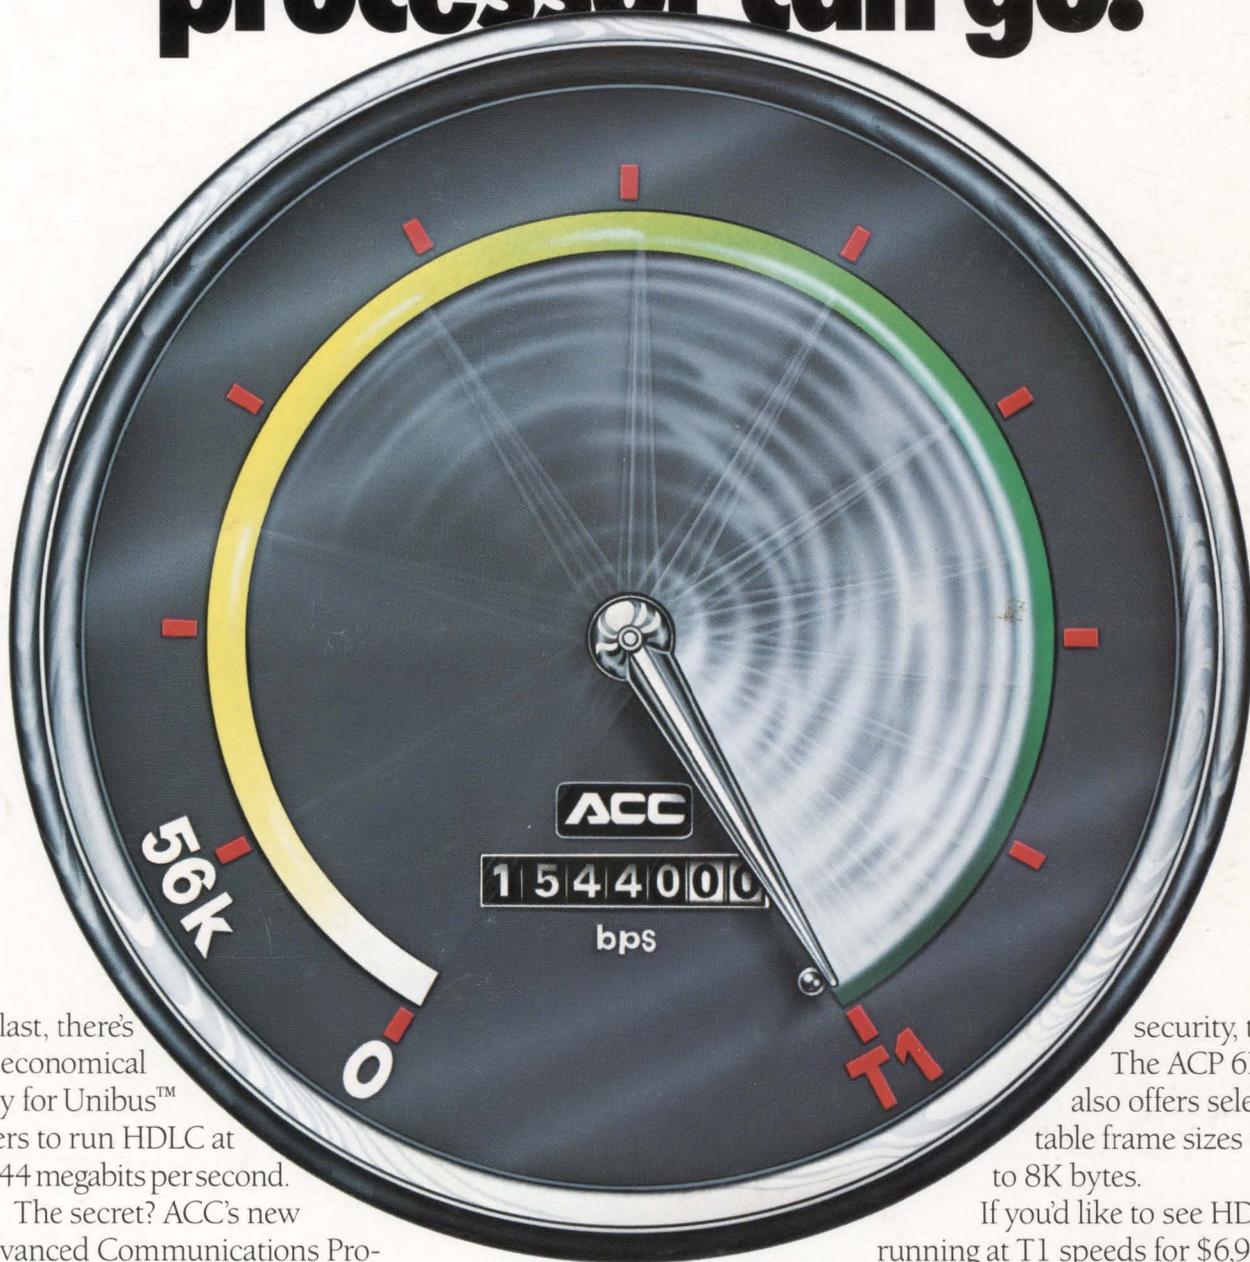
3540 Arden Road, Hayward, CA 94545
(415) 783-8400

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

H-I® Hitachi America, Ltd.
computer graphic courtesy of METHEUS Corporation, Hillsboro, Oregon

See how fast our \$6,900 communications processor can go.



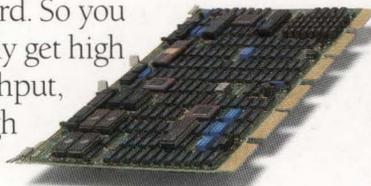
At last, there's an economical way for Unibus™ users to run HDLC at 1.544 megabits per second.

The secret? ACC's new Advanced Communications Processor, the ACP 6100.

This point-to-point communications system is perfect for computer-integrated manufacturing, high-speed graphics, and many other applications that require high-speed, high-volume data transfers.

Unibus is a trademark of Digital Equipment Corporation.

What's more, the ACP 6100 allows you to encrypt your data using the NBS Data Encryption Standard. So you not only get high throughput, but high



security, too. The ACP 6100 also offers selectable frame sizes up to 8K bytes.

If you'd like to see HDLC running at T1 speeds for \$6,900, call ACC. We'll rush you complete details on our Advanced Communications Processor just as fast.



Advanced Computer Communications
720 Santa Barbara Street
Santa Barbara, CA 93101
(805) 963-9431