BYTE

AUGUST 1989   A McGRAW-HILL PUBLICATION

REVIEWS

Magellan vs. ViewLink

**Pixelworks'
Ultra Clipper for PS/2s**

NEC's UltraLite

3 OS/2 Modula-2s

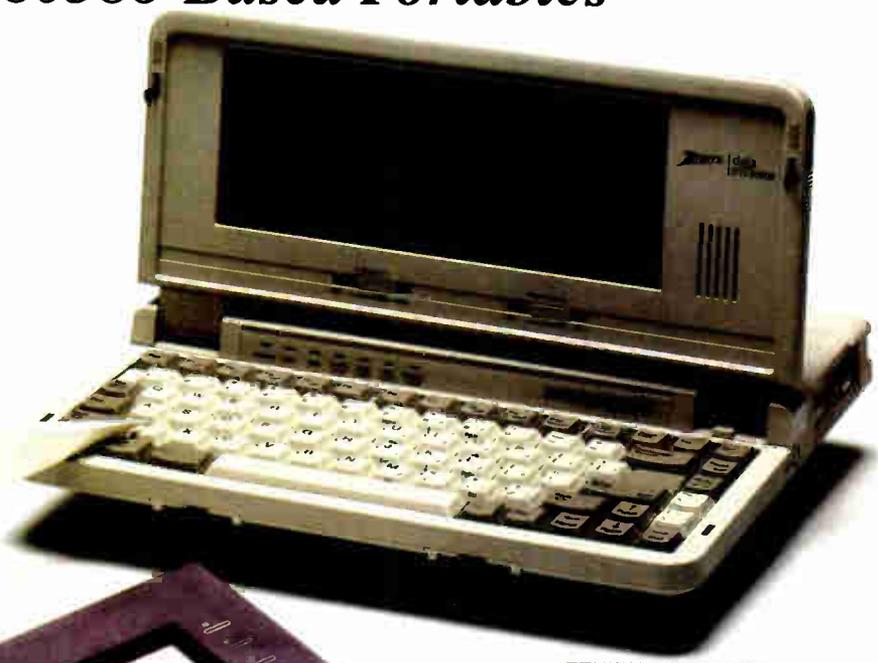**PRODUCT FOCUS**

IBM's brand-new P70

**80386-Based
Portables**

Small Wonders

MinisPort and Agilis Laptops
Plus: 11 80386-Based Portables

ZENITH MINISPORT

AGILIS SYSTEM

NEURAL NETWORKS *In Depth*
**Brain-like systems
solve real-world problems**

How Disk Optimizers Work

**Expert Advice:
A New Unix column**

What's Beyond RISC?

**Digital Signal Processors
Move to Micros**

5 Short Takes

$3.50 U.S.A./$4.50 IN CANADA
0360-5280

08
29920
0 440235 0

DOT MATRIX PRINTERS.

Printer System 800: $699.95
Our highest resolution text and graphics, 24-pin dot matrix printer. Draft quality
at 200 cps. Letter quality at 66 cps. Parallel and serial interfaces. Wide carriage.

Printer System 300: $199.95
9-pin dot matrix. Draft quality at 144 cps. Near-letter quality at 36 cps. Four
standard fonts. Parallel interface. Narrow carriage.

NEW MONITOR

Dell Super VGA Color Monitor. Supports all VGA modes plus new 800 x 600
standard. Call for details.

LASER PRINTERS.

SO HOW COME
YOU NEVER CALL?

THE DELL SYSTEM 325 25 MHz 386

STANDARD FEATURES:

* Intel 80386 microprocessor running at 25 MHz.
* Choice of 1 MB, 2 MB or 4 MB of RAM* expandable to 16 MB using a
  dedicated high speed 32-bit memory slot.
* Advanced Intel 82385 Cache Memory Controller with 32 KB of
  high speed static RAM cache.
* Page mode interleaved memory architecture.
* VGA systems include a high performance 16-bit video adapter.
* Socket for 25 MHz Intel 80387 or 25 MHz WEITEK 3167 math coprocessor.
* 5.25" 1.2 MB or 3.5" 1.44 MB diskette drive.
* Dual diskette and hard disk drive controller.
* Enhanced 101-key keyboard.
* 1 parallel and 2 serial ports.
* 200-watt power supply.
* 8 industry standard expansion slots.

OPTIONS:

* 40 MB or 150 MB tape backup.
* 25 MHz Intel 80387 math coprocessor.
* 25 MHz WEITEK 3167 math coprocessor.
* 1 MB or 4 MB RAM upgrade kit.
* 2 MB or 8 MB memory expansion board kit.
* Graphics Performance Accelerator GPX-1024.
* Graphics Performance Display GPD-16C, GPD-19C.
**Lease for as low as $199/Month.
◇ *Extended Service Plan pricing starts at $370.*

System 325	With Monitor & Adapter			
	VGA Mono		VGA Color Plus	
Hard Disk Drive	1MB RAM	2MB RAM	1MB RAM	2MB RAM
40 MB-29 ms IDE	$5,499	$5,698	$5,799	$5,998
100 MB-25 ms IDE	$5,999	$6,198	$6,299	$6,498
150 MB-18 ms ESDI	$6,499	$6,698	$6,799	$6,998
322 MB-18 ms ESDI	$7,299	$7,498	$7,599	$7,798

*Disclaimer: All systems are photographed with optional extras, which some computer
retailers won't even recognize.*

World Radio History

Plus, for high resolution colors displayed on a larger screen.

Even though the 325 gives you all this performance, it still leaves you six open slots for whatever else you might want to add.

And once you've told us what you want, we'll make sure what you want works—by burning-in the entire system unit.

COMPUTER RETAILERS ARE NO KNOWS.

In all probability, the average computer retailer won't have any understanding what makes the System 325 go.

He will, however, be quite aware of the fact that he could add a 35% markup if he could sell it in his store.

Which he can't.

Because we sell direct.

Meaning you now have the unique opportunity to talk directly with a computer expert. And ask things like, "What's the difference between IDE and ESDI?" Or, "How much SIMM RAM should I add?"

In other words, the kinds of questions you should be able to ask a retailer, but usually can't.

So as you might suspect, dealing direct not only saves you the 35% markup, but 100% of the aggravation.

WE COME WHEN WE'RE CALLED.

One of the things that very clearly sets a Dell system apart

from other computers is not just how they're sold, but how they're supported.

Overkill was one description used in a PC Week article.

Perhaps.

But then, we think you'll agree, when something goes wrong, you want as much help as

possible, right?

Which is why every Dell system comes with a toll-free technical support line and self-diagnostic software. We're able to solve 90% of all problems right over the phone. The other 10% receive next-day, deskside service. Thanks to our new alliance with Xerox Corporation.

And you get all this help for a full year—whenever you need it—at no extra charge.△

As you've probably guessed, one of the things that drives us most is customer satisfaction.

So we'd like to give you the ultimate guarantee: Try a System 325 in your office for a month. Run your toughest applications. Put it through its paces, at your pace. If you're not completely satisfied, send it back anytime within 30 days. And

we'll refund your money.

No questions asked.

MAYBE YOU SHOULDN'T BUY ONE AFTER ALL.

No matter how many reasons we give you to buy a Dell system, sometimes it makes more sense to lease one instead.

BEST OF ALL, YOU WON'T HAVE TO EXPLAIN TO A COMPUTER RETAILER WHAT ALL THAT MEANS.

Whether you need a single computer, or an office full, a leasing plan is like 100% financing.

And just as we can custom configure your computers, we can see to it you get a custom designed lease plan to fit your exact business needs. † A fact that has not gone unnoticed. Especially by the Fortune 500. Over half of whom now own or lease Dell systems.

And just as we welcome their business, we welcome your business, too. Just call us, toll-free. And don't be afraid to ask us the tough questions.
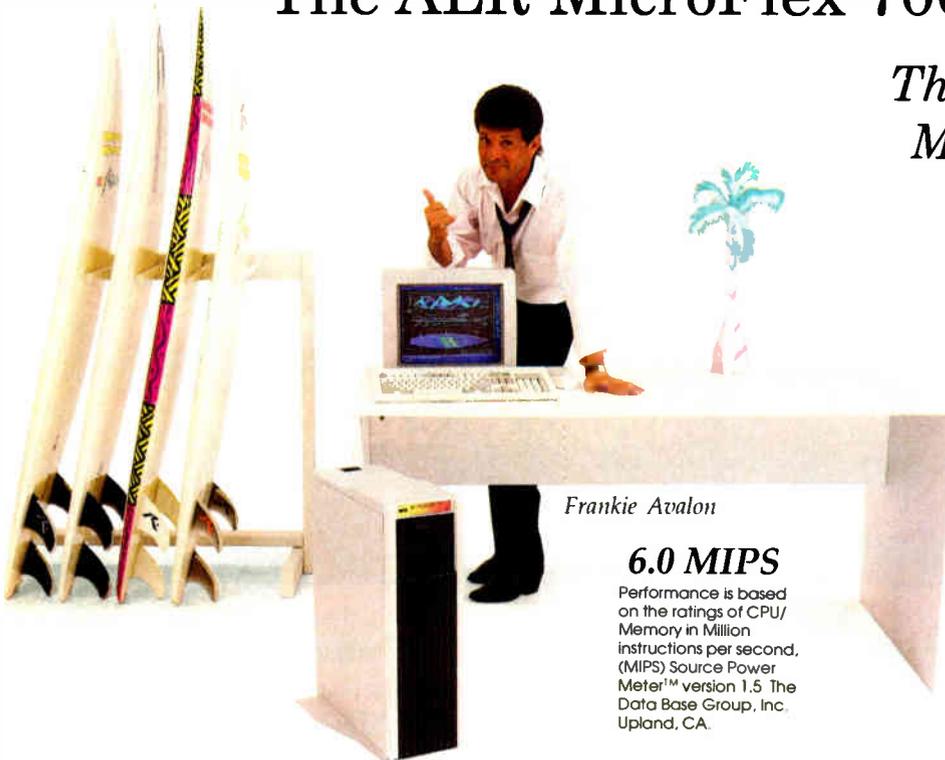
That's the part we like best.

Now we're making waves with IBM.®

The ALR MicroFlex 7000

The first 25MHz Micro Channel® compatible

Frankie Avalon

6.0 MIPS

Performance is based on the ratings of CPU/ Memory in Million instructions per second, (MIPS) Source Power Meter™ version 1.5 The Data Base Group, Inc. Upland, CA.

California

Home of the World's First 386 PC
Advanced Logic Research, Inc.

At ALR, we thrive on opportunities to beat our competitors. Our 25MHz 80386® based MicroFlex 7000 is no exception.

Unmatched performance
Our proprietary "pre-fetch" FlexCache™ design delivers the most efficient form of microcomputer processing. By combining a true 64-bit cache bus with 64-KB cache memory, performance increases 30% when compared to other 32-bit computers. And 64KB of high-speed cache memory enables you to experience the fastest

throughput for sophisticated applications. For those seeking large storage capacities, the MicroFlex 7000 gives the option of 120 or 300MB of disk storage using high-speed ESDI controllers with 1:1 interleave.

The most built-in features
The MicroFlex 7000 includes our super VGA controller with 800 X 600 graphics resolution and the sleek tower chassis offers the most internal expansion capabilities of any Micro Channel system available. Our one-

year warranty with unlimited technical support and on-site servicing available from Intel® can't be beat.
So make some waves of your own at the office with ALR's MicroFlex 7000 or any of our 33MHz systems. For more information and the name and number of your local authorized ALR reseller, please call:

1-800-444-4ALR

Advanced Logic Research, Inc.
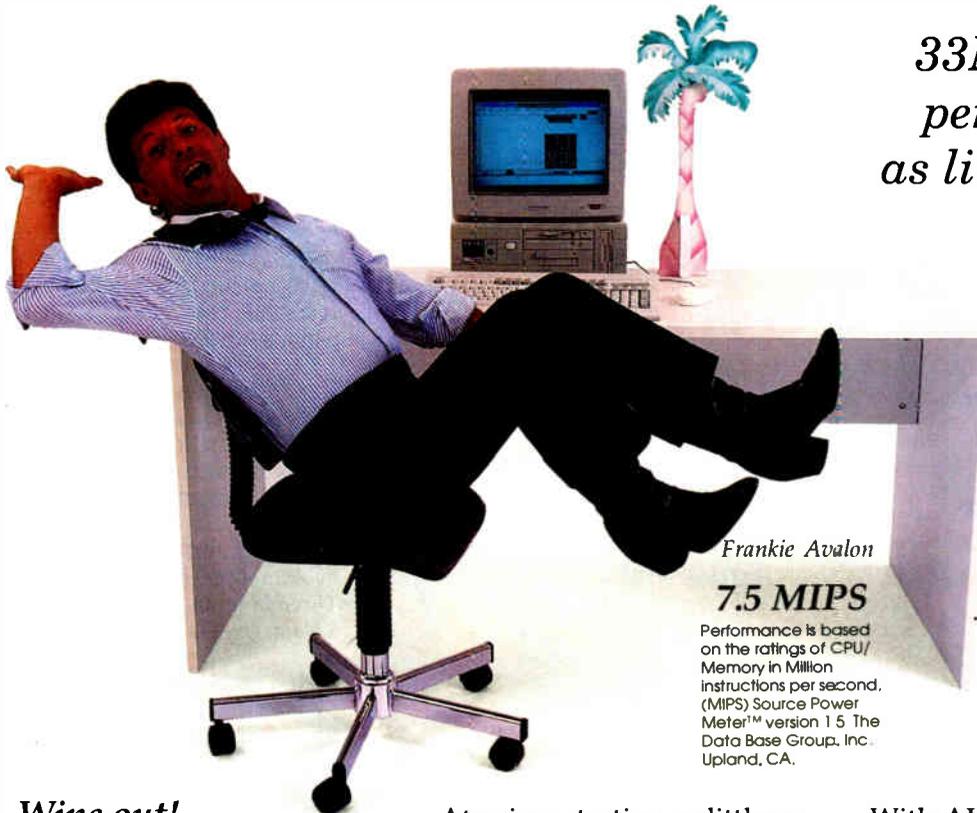9401 Jeronimo, Irvine, CA 92718
(714) 581-6770 FAX:(714) 581-9240

For our Canadian office: 1-800-443-4CAN
For our UK office: 44-1-399-4897
For our Singapore-Asia/Pacific office:
(65) 258-1286 FAX: (65) 258-1285

We're making some big waves in California

Introducing ALR's FlexCache™ 33/386Z

33MHz 80386™ performance for as little as $3995!

Frankie Avalon

7.5 MIPS

Performance is based on the ratings of CPU/Memory in Million instructions per second, (MIPS) Source Power Meter™ version 1 5 The Data Base Group, Inc. Upland, CA.

California
Home of the World's First 386 PC
Advanced Logic Research, Inc.

Wipe out!
Hang on because ALR's latest addition to the FlexCache 386™ Z-family is cruising at an amazing 33MHz. That's a 20% increase in processing speed when compared to the award winning FlexCache 25386.

Fast Cache
With 32KB of cache memory, award-winning FlexCache architecture and our enhanced 16-bit super VGA controller you better be ready to move.

At prices starting as little as $3995*, the FlexCache 33/386Z delivers the most performance for all power hungry desktop applications like CAD/CAM, desktop publishing or financial modeling at a very modest price. Of course the FlexCache 33/386Z is OS/2® compatible for tomorrow's latest generation of applications. The FlexCache 33/386Z as with all of the Z-Family comes packaged with PC-Kwik®, the award-winning disk caching utility.

With ALR's FlexCache 33/386Z you'll receive unbeatable support backed by an unprecedented three year factory warranty on the main system board, a one-year system warranty, unlimited technical support and optional on-site servicing from Intel.
For more information on the FlexCache 33/386Z call:
1-800-444-4ALR.

Circle 13 on Reader Service Card (DEALERS: 14)

World Radio History

BYTE

COVER STORY
The Ever-Shrinking, Ever-Expanding Laptops
by Nick Baran and Michael E. Nadeau
page 90

Agilis and Zenith announce tiny computers that broaden the market for laptops.

EXPERT ADVICE

MACINTOSH
SPECIAL SUPPLEMENT

Dealing with a Digital World/246

Microsoft profession
something other lang

al languages give you uages don't. Leverage.

In an industry that evolves practically overnight, it's tough to stay ahead of the crowd.

You need tools that not only give you an edge day-to-day, but open up endless possibilities.

Tools that can only come from Microsoft.

Combine Microsoft® C and Macro Assembler and you've got enough power to create programs for MS-DOS, Windows and OS/2 systems.

What's more, you can do it all in record time because our renowned CodeView® Debugger, Linker, Microsoft Editor, and MAKE utility work ingeniously and seamlessly together.

In other words, you've got the leverage of the most inventive and comprehensive tools around.

When you develop under OS/2 systems, you've got options no one else can touch. Like multi-tasking. And blasting through the 640K barrier.

In addition, Microsoft C and Macro Assembler can accommodate more third party add-ons than any other PC professional languages.

Maybe that's why the most popular applications on the market today were developed through the unique power of our C and Assembler: Lotus® 1-2-3® WordPerfect® 5.0. Microsoft Excel. And Aldus® PageMaker®.

So drop by your nearest Microsoft dealer soon. And start turning out the most airtight, fine-tuned code ever to touch a disk.

After all, you've got the leverage.

Microsoft
Making it all make sense.

HOLD ONTO YOUR HAT (AND YOUR WALLET)

Apricot announces the first "real" 80486-based machine. Meanwhile, there's nothing but lawsuits on the 68040 front.

I recently received a transatlantic phone call from an excited Paul Lavin, a colleague who writes for a number of British computer publications. He'd just caught wind of a surprising development: Apricot was about to announce an 80486-based computer.

When Intel introduced the 80486 last April, it predicted that 80486-based machines would appear late this year, with volume shipments next year. I never suspected that the British company Apricot would be first with an 80486 machine, or would have one so soon.

But it was: Its 80486-based 25-MHz Micro Channel architecture machine was announced in June in London. The first production units will be available a few weeks after you read this.

Apricot even beat IBM's 80486 announcement by several weeks, although IBM had shown a nearly finished prototype in April. (For more details on the IBM machine, see the June editorial.)

Apricot's machine, called the VX FT, is a 15-million-instructions-per-second beast that comes with up to 5 gigabytes of SCSI hard disk drive storage, up to 16 megabytes of RAM on the motherboard, a digital audio tape-recording backup subsystem, built-in disk shadowing (for fault tolerance), support for up to 128 serial ports, and a 465-watt power supply with its own built-in lead/acid backup batteries (that's right, a built-in uninterruptible power supply). This is all cooled by three or four 4-inch-diameter fans, depending on how the machine is configured. We're talking heavy duty.

The box, which is mounted on skids, is about the size of a fat two-drawer filing cabinet. Weighing over 150 pounds, it comes with a pair of built-in retractable handles so that two people can horse it around. Whew.

The VX FT is built around an MCA motherboard with eight slots (four 16-bit and four 32-bit). The motherboard uses standard Chips & Technologies chips and a Phoenix BIOS. This helps ensure compatibility; the machine was shown running MS-DOS 4.01, OS/2 Extended Edition, Novell NetWare, 3+Open LAN Manager, SCO Unix System V release 3.2, and other software.

While the chips and the BIOS are conventional, the Apricot designers went their own way in the addition of a separate cache on the motherboard (this is unusual because the 80486 has an on-board cache of its own). Apricot believes that this 128K-byte "Hypercache" will give the VX FT a performance edge over those machines that simply use the 80486's on-board cache.

So far, it's an unproven belief: As I write this, the Apricot engineers are eradicating some last-minute problems that cropped up in the first Hypercache prototypes. As soon as the glitches get sorted out, we'll bring you full benchmark results and Paul Lavin's hands-on report.
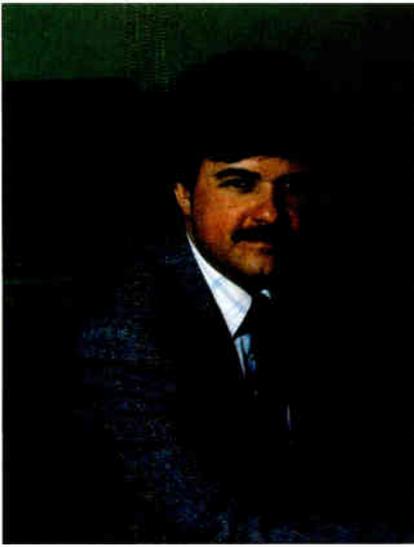
Of course, all this horsepower and storage isn't exactly cheap: Prices start at the very high end of the microcomputer price spectrum (about $18,000) and go up from there, topping out in the exospheric $40,000 range. Clearly, this won't be a high-volume system.

Meanwhile, at the Low End...
Cheetah (see the June Editorial) is still on track with a low-cost 80486-based motherboard—one that actually could cost less than a similarly clocked 80386-based system with a separate 80387 math chip and cache.

We may see the 80486 market split in two radically different directions: killer systems with killer prices for departmental computing needs, and relatively inexpensive fast systems for personal desktop use.

The prices of 80486-based systems could also be kept somewhat in check due to competition from the Motorola or RISC camps, if those chip makers can mount an aggressive attack. Unfortunately, there are problems.

For example, we still haven't heard of a single demonstration of a 68040-based system, even though the 68040 was announced before the 80486. One possible explanation is—surprise!—legal hassles: Hitachi has accused Motorola of violating Hitachi patents with its 68030 microprocessor, currently Motorola's top-of-the-line shipping CPU. The 68040 includes an enhanced 68030 as its core; it's reasonable to surmise that legal complications involving the 68030 might spill over to affect the 68040.

Sadly, legal wrangling isn't at all unusual these days. But not since NEC sued Intel over rights to make clones of the 8088 and 8086 CPUs has a suit attacked an American microprocessor maker's premier product—in this case, the Motorola 68030, which is used in Apple's Macintosh IIx and IIcx and in workstations from Sun and Hewlett-Packard.

Perhaps this lawsuit is one of the reasons why development of 68040-based systems appears to be lagging far behind that of 80486-based systems. (I can only guess; Motorola is mum on the subject.)

I hope that the legal snags will get resolved and that Motorola and others can provide healthy competition for high-end 80486s; and that companies like Cheetah can cultivate low-cost 80486s.

The Apricot VX FT is nice—very nice. But prices like that take the "personal" out of personal computing.

—Fred Langa
Editor in Chief
(BIX name "flanga")

Integrated Software for Schematics & PCB Artwork

Introducing HiWIRE®-Plus

Wintek's smARTWORK® pioneered low-cost printed-circuit-board CAD. Then HiWIRE set the standard for productivity and ease-of-use in schematic capture. Now Wintek introduces HiWIRE-Plus, integrating HiWIRE's schematic features with a powerful printed-circuit-design facility.

Creating Schematics

With HiWIRE-Plus, simply connect library symbols with wires and buses. Creating and changing symbols is fast and painless. Produce your drawing using a dot-matrix printer, laser printer, or pen plotter.

Circuit-Board Design

HiWIRE-Plus gives you all the design freedom you want: you choose the grid size, trace widths, and pad shapes. The board size and number of layers are virtually unlimited. HiWIRE-Plus is perfect for surface-mount, micro-strip, and ECL applications.

Current Versions	
HiWIRE-Plus	V 1.1r6
smARTWORK	V 1.4r5

Circle 277 on Reader Service Card

HiWIRE-Plus Advantages

☐ One tool for schematics and printed-circuit artwork

☐ Easy-to-learn menu-driven operation; complete documentation and tutorial

☐ Schematic libraries with TTL, CMOS, ECL, ladder, microprocessor, and discrete components

☐ Netlist and bill-of-materials utilities included

☐ Circuit boards up to 60x60 inches and 256 layers

☐ Variable grid size, trace width, and pad size (.001" resolution)

☐ PCB library with DIPs, SIPs, SMDs, PGAs, TOs, and edge and D connectors

☐ Schematic-to-layout cross-checking

☐ Design-rule checker

☐ 800 number for free support

"HiWIRE", "smARTWORK", "Wintek", and the Wintek logo are registered trademarks of Wintek Corporation.

System Requirements

☐ IBM PC, XT, AT, or PS/2 with 512K RAM, printer port, color monitor, and CGA, EGA, or VGA graphics card

☐ Microsoft Mouse

☐ IBM ProPrinter or Epson dot-matrix printer, and/or

☐ Houston Instrument or Hewlett-Packard pen plotter

Higher Performance Better Value

Still only $895, HiWIRE-Plus delivers quality schematics and PCB artwork. You don't need to guess if HiWIRE-Plus is right for you—*we guarantee it!* Try it for 30 days at absolutely no risk. Call toll free today and put HiWIRE-Plus to work for you.

Wintek Corporation
1801 South Street
Lafayette, Indiana 47904-2993
(800) 742-6809 or
(317) 742-8428
FAX: (317) 448-4823
Telex: 70-9079

Europe: RIVA Ltd., England,
Phone: 0420 22666, FAX: 0420 23700
Australia: Entertainment Audio Pty, Ltd.,
Phone: (08) 363-0454

OUT-STANDING

The Company

Gateway 2000 has consistently led the pack of competitors in this highly competitive field. We have the most aggressive pricing in the industry, the most aggressive support policies, and lead the pack in quality. So shop around, then call Gateway 2000 to discover just how far ahead of the competition we really are.

The Products

Your Gateway 2000 computer system will arrive thoroughly tested and ready to run. All of our top quality systems come standard with our own *Crystal-scan 860* monitor and a 16 bit VGA card that is expandable to 512K. We have a variety of options available to suit anyone's needs. So call Gateway 2000, and we'll custom configure a system just for you.

12 Mhz 286 VGA
80286-12 Processor
2 Megs RAM
1.2 Meg 5¼" Drive
1.44 Meg 3.5" Drive
40 Meg 28ms Drive
16 Bit VGA Board
14" VGA COLOR Monitor
1 Parallel/2 Serial Ports
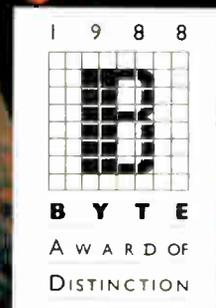101 Key Keyboard
MS DOS 3.3 or 4.01
$2295.00

16 Mhz 286 VGA
80286-16 Processor
2 Megs RAM
1.2 Meg 5¼" Drive
1.44 Meg 3.5" Drive
40 Meg 28ms Drive
16 Bit VGA Board
14" VGA COLOR Monitor
1 Parallel/2 Serial Ports
101 Key Keyboard
MS DOS 3.3 or 4.01
$2395.00

20 Mhz 286 VGA
80286-20 Processor
2 Megs RAM
1.2 Meg 5¼" Drive
1.44 Meg 3.5" Drive
40 Meg 28ms Drive
16 Bit VGA Board
14" VGA COLOR Monitor
1 Parallel/2 Serial Ports
101 Key Keyboard
MS DOS 3.3 or 4.01
$2495.00

1988
B Y T E
AWARD OF
DISTINCTION

Due to the Volatility in the DRAM Market all prices subject to change.

World Radio History

IN THEIR FIELD

The Service

Gateway 2000 backs it's computer systems with a full one year warranty and 30 day money back guarantee. If a problem does arise, you will promptly receive a solution over the phone or via *Federal Express* at our expense. In addition to this, we offer *lifetime toll-free support,* even after the warranty expires. So call Gateway 2000 for the best overall value on the market today.

20 Mhz 386 VGA
1 Meg RAM
1.2 Meg 5¼" Drive
1.44 Meg 3.5" Drive
80 Meg 28ms Drive
16 Bit VGA Board
14" VGA COLOR Monitor
1 Parallel/2 Serial Ports
101 Key Keyboard
MS DOS 3.3 or 4.01
$2995.00
(Upgrade to 4 Megs $500)

25 Mhz 386 VGA
4 Megs RAM
1.2 Meg 5¼" Drive
1.44 Meg 3.5" Drive
150 Meg 16.5 ms ESDI Drive
16 Bit VGA Board
14" VGA COLOR Monitor
1 Parallel 2 Serial Ports
101 Key Keyboard
MS DOS 3.3 or 4.01
$4495.00
(64K Cache Add $500)

33 Mhz 386 VGA
64K Cache RAM
4 Megs RAM
1.2 Meg 5¼" Drive
1.44 Meg 3.5" Drive
150 Meg 16.5 ms ESDI Drive
16 Bit VGA Board
14" VGA COLOR Monitor
1 Parallel/2 Serial Ports
101 Key Keyboard
MS DOS 3.3 or 4.01
$5995.00

Gateway 2000
P.O. Box 2000
Sgt. Bluff, IA 51054
800-779-2000
712-943-2000

HOW TO MAKE P
How to make points.
HOW TO MAKE PO
How to Make Poin
How to Make Pe
HOW TO MAKE F
How to make points
How to make points.
How to make poi
How to Make Points.
HOW TO MAKE
HOW TO MAKE POIN
How to Make Points
HOW TO MAKE POIN
HOW TO MAKE Pt
How to make points
How to make points
HOW TO MAKE POIN
How to Make Points
How to make p
How to make
HOW TO N
How to ma

How to Make
HOW TO MAKE
How to make points.
How to Make Poin
HOW TO MAKE PO
How to make points
How to Make F
HOW TO MAKI
How to Make Point
How to make poin
HOW TO MAKE Pt
HOW TO MAKE POINTS.
How to make point
HOW TO MAKE POINTS
How to make point
HOW TO MAKE POIN
How to make points.
HOW TO MAKE POINTS
How to make pt
HOW TO MAKE POINT
HOW TO MAKE POINTS.
HOW TO MAKE POIN
How to make points.
How to make point
How to make poi

World Radio History

MICROBYTES

Staff-written highlights of developments in technology and the microcomputer industry, compiled from Microbytes Daily and BYTEweek *reports*

AT&T "Microscopic Parallel Processor" Hits 24 GHz

Scientists at AT&T's Bell Laboratories (Murray Hill, NJ) have built a new quantum-effect transistor that promises some intriguing future generations of computers. Texas Instruments was the first to announce development of a quantum transistor (see "TI's Prototype Transistor Takes a Quantum Leap," March Microbytes). Designers working with such a device will someday be able to implement far more functions on a chip than is possible with today's ICs.

In normal transistors, the output current increases steadily as the input current rises. But according to Frederico Capasso, one of the three co-developers of the new multistate resonant tunneling transistor, the output current peaks, falls off, and then peaks again in the AT&T device. This multistate characteristic allows it to do the work of many conventional transistors. Capasso calls it a "microscopic parallel processor." In addition to its almost unimaginably small size, the transistor operates at up to 24 gigahertz, about twice the speed of the fastest conventional silicon transistors, while it requires much less power than current ICs.

Although the device is only in an experimental prototype stage right now, AT&T scientists say they have already used a single quantum transistor to implement functions such as parity bit-checking (which normally requires about 24 transistors). They say they've also used a single device to multiply a frequency from 300 MHz to 1.5 GHz.

Like Texas Instruments' device, AT&T's transistor uses a quantum phenomenon called *resonant tunneling*, which occurs in quantum wells—electron filters formed by stacking microscopically thin layers of semiconductor atop one another. Only electrons with certain energies can pass through the wells.

AT&T's device uses two well layers, made of gallium-indium-arsenide and measuring just 25 atoms thick; each layer is surrounded by two aluminum-indium-arsenide barrier layers of the same thickness. AT&T scientists say they created the multistate capabilities of the device by increasing the number of wells. The actual wells are made using an AT&T-developed technique called molecular beam epitaxy. Because it allows scientists to build devices one atom thick at a time, MBE lets designers concentrate the entire circuit function vertically into a single device. Capasso says this is the first demonstration of a three-dimensional integrated transistor device.

But don't expect to buy a laptop supercomputer yet. Commercial applications of the multistate resonant-tunneling transistor are probably five to 10 years away. The primary problem is that new techniques will have to be developed to allow mass production of quantum transistors.

VROOMM: Borland Says Memory Technology Will Make Future Programs Better, Not Bigger

Borland International (Scotts Valley, CA) says its new proprietary programming technology will enable it to develop applications that have more features and greater data capacity but still fit within the 640K-byte limitation of MS-DOS. Borland is calling the technology VROOMM, which stands for Virtual Real-Time Object-Oriented Memory Manager, a fancy marketing phrase for a programming concept called *dynamic segment*

swapping. The software company says that it will use VROOMM in all its applications and development tools; the new Reflex 2.0 is the first application to implement VROOMM.

In contrast to the concept of segment overlays, in which parts of the program's executable code are compiled into separate, fixed-size overlays and swapped in and out of memory, dynamic segment swapping

continued

of Phar Lap's DOS Extender, which uses paged virtual memory and permits use of 32-bit-wide commands.

Tandy (Fort Worth, TX) has given its **DeskMate** environment a slight face-lift. DeskMate, which comes free with Tandy computers and in a run-time version with some application programs, now has a more three-dimensional look, with features such as buttons that look pushed down when selected. In addition to a DOS shell and simple word processor, spreadsheet, and communications packages, DeskMate has a paint-style graphics program and a digital sound-manipulation program (which works with the sound circuitry of the Tandy 1000). For $149, you can add a WorkGroup program that provides printer sharing, file sharing, and E-mail functions.

Tennessee volunteers for ISDN: The South Central Bell phone company is planning to make Tennessee the first state in the nation with an all-digital telecommunications infrastructure, paving the way for ISDN services to homes and businesses. The three-year, $900 million program is designed to replace electro-mechanical and analog central office computers with digital central office equipment and to double Tennessee's fiber-optic network from 12,000 to 25,000 miles. By 1990, all Tennessee customers will be served with digital central office computers, according to a South Central Bell official. This doesn't mean all customers will have immediate access to an ISDN, he said. "Digital links provide the basis for ISDN access. Upgrading to ISDN will mostly involve adding software."

Microsoft (Redmond, WA) has introduced a new version of FORTRAN that could be good news for programmers who do their work on expensive VAX and IBM mainframes but want to move to PCs. Microsoft's **FORTRAN 5.0** supports most of the syntax of

lets you swap smaller segments of code at run time and in varying amounts, depending on what you're doing with the program. Many DOS programs employ fixed overlay files, typically from 30K bytes to more than 100K bytes in size, that must be loaded into memory in their entirety whenever you use them. Large overlay segments limit the amount of free memory left over for data, such as text files, spreadsheets, and database tables.

VROOMM gets around the size restrictions imposed by fixed overlays by using 2K- to 4K-byte chunks of code (segments), which make up the complete application. In the dynamic-segment-swapping model, the program can decide on the fly (dynamically) which segments of code it needs in memory at any given time. If you are currently running the Window Manager, for example, the program brings in the segments of code required to operate the Window Manager. VROOMM also allows the application to juggle the amount of

code in memory with the amount of data being stored in memory.

Consider a user working on a large database table requiring a lot of memory; VROOMM can swap more segments out to disk or to expanded memory, if available. VROOMM stores the most frequently used code segments in an object cache. The cache can reside either on disk or in expanded memory. VROOMM assigns a priority to each code segment (called *persistence prioritization*), depending on the way you use the program (e.g., whether you're working with reports, graphics, or data-entry forms).

Will Borland make VROOMM available to other software developers? According to Rob Dickerson, vice president of product development, the company has not resolved this question. At a recent meeting of the Boston Computer Society, Borland president Philippe Kahn said that the company "will provide the right tools" to other developers "when we can support them."

Lisp's Future Linked to Other Languages

Since its inception at MIT about 25 years ago, Lisp has become the lingua franca of AI; because of its symbolic and procedural capabilities, many programmers choose it for developing rule-based or expert systems. However, AI applications and Lisp have not enjoyed the success anticipated by AI proponents. Lisp has a very large syntax (about 200 primitives), and, according to some programmers, it is difficult to learn. Because it processes symbols and lists, Lisp requires a lot more computing horsepower than conventional languages like C or Pascal, which deal with predefined, fixed-length data structures. As a result, Lisp has been hampered by slow performance and slow acceptance. In recent years, several Lisp companies have gone bankrupt, and the largest AI company, Symbolics, has suffered two straight years of losses.

But some Lisp developers are optimistic about Lisp's future. With the industry's continuing advances in processing power, the performance of Lisp is becoming acceptable to more users. And developers are finally recognizing that the key to Lisp's success is integrating it into main-

stream computing—in other words, allowing Lisp to be linked to existing applications written in other high-level languages (e.g., C and Pascal).

The main trend in AI today is the integration of Lisp-based "intelligent add-ons" to existing database systems, says Pekka Pirinen, director of research and development for Intellitech, a Lisp vendor based in Helsinki, Finland. An "intelligent add-on" might be a rule-based query system that acts as an interface to a large body of existing data.

The key to the intelligent add-on concept is Lisp's ability to link directly to other high-level languages. If you simply add a direct function call to the Lisp application, the application can then link to and execute an existing C or Pascal program. Intellitech has just announced a new Lisp product, called Entity Common Lisp, for 80386-based computers. Requiring 4 megabytes of RAM and Microsoft Windows, ECL will feature links to C and Pascal compilers from Microsoft and Borland. According to Pirinen, ECL is the first Lisp product in the DOS environment that can be linked to other high-level languages. ECL is

World Radio History

Better is

The new 16" MultiSync® 4D and 20" 5D both offer compatibility with

a wide range of graphics boards and computer systems.

NEC's award-winning multiple frequency technology, now available in larger, digitally controlled, high-resolution color monitors.

Both are optimized for the IBM PS/2, PC/AT/XT (and 100% compat-

ibles) and the Macintosh II. Both offer a microprocessor-based digital

control system for preset and custom graphics modes, automatic screen

C&C
Computers and Communications

bigger.

configuration and optimal image. And both give you great resolution—

the 4D, from VGA up to 1024 x 768; the 5D, from VGA to 1280 x 1024.

The main difference: one is for large ideas. The other, for extra-

large. For literature, call 1-800-826-2255. For details, call NEC Home

Electronics (USA) Inc. 1-800-FONE-NEC.

NEC

IBM's Systems Application Architecture-compliant VS FOR-TRAN and important VAX FOR-TRAN extensions, such as NAMELIST, OPEN append, DO WHILE, and INCLUDE. FORTRAN 5.0 also supports the 16-megabyte addressing capability of OS/2 and library routines without recompiling the entire application. Microsoft claims that FORTRAN is the main program-ming language used on over half of the VAX machines installed today. Microsoft's new version is based on FORTRAN 77. There is "still some amount of controversy" over the implementation of the upcom-ing FORTRAN 8X standard, a company official said. FORTRAN 5.0 is available now at $450. Upgrades range from $100 to $250.

Smalltalk + Schools = Object-Oriented Students: Hoping to raise the next generation of software engineers on object-oriented programming, **Digitalk** (Los Angeles) is offering educa-tional site licenses for its Smalltalk system. For $500, schools and universities that are registered users of the Smalltalk/V 286 or Small-talk/V Mac development environ-ments can install unlimited copies of the software on-site. To become a registered user requires buying Smalltalk ($199.95), but the site license includes a 30 percent discount on all Smalltalk products, so the software ends up costing $139.95. "Those relying on traditional procedural languages are struggling to adapt these older languages to complex environ-ments like the Macintosh," said Barbara Noparstak, Digitalk's director of marketing. "Educating a new class of programmer with tools that simplify this complexity will reap industry benefits in the years to come."

Microtech International (Bran-ford, CT) has a new trade-in deal for Macintosh owners looking for a new hard disk drive. You can apply your hard disk drive, whether it works or not, and from any manu-facturer, toward purchase of a Microtech Nova internal or external
continued

priced at $995; a run-time kit is $495. The full integrated system was scheduled to be available this month, although a version without the language hooks was supposed to ship by July. (In the U.S., ExperTelligence, Inc. (Goleta, CA), is selling the Intellitech package.) Intellitech says that it is also preparing a version of ECL for OS/2.

Zenith's 2-inch Floppy Signals Shrinking Standard

Zenith's introduction this month of its MinisPort points the way toward the next step in the shrinking of the personal computer: the 2-inch floppy disk. In an industry where smaller is better, the floppy disk will continue to shrink. With Zenith Data Systems adopting the smaller floppy disk, there's a good chance that the 2-inch, 720K-byte floppy disk will become the standard storage medium on laptop computers sooner than most industry watchers expect. Other computer makers working on new laptops are also considering designs that use the 2-inch disks.

The 2-inch 720K-byte disk has the identical read/write and magnetic format as 1.44-megabyte 3½-inch disks. According to Zenith's market-ing director, Glenn Nelson, the basic engineering concept involves taking half the surface area of the magnetic film of a 1.44-megabyte 3½-inch disk and putting it on the 2-inch disk, resulting in half the data capacity, or 720K bytes. Although a 2-inch disk has only about one-third the surface area of a 3½-inch disk, the 2-to-1 reduction is possible because not all the surface area of a 3½-inch disk is used on current 1.44-megabyte floppy disks (about half an inch of the outer radius of the 3½-inch disk is not used to store data, according to Nelson). Nelson also said that the 2-inch disk drives perform approximately the same as their 3½-inch counterparts.

Nelson acknowledged that until the 2-inch disk becomes a standard, little software will be available in the 2-inch format and users will have to rely on file transfer utilities to send applications and data to the 2-inch drive system from another computer, using the serial ports of the host and target systems.

Nelson declined to comment on which manufacturers are supplying Zenith with the 2-inch drives. How-ever, 2-inch floppy disks are already in limited use in the video and camera market and are manufactured by Sony and other companies.

ParcPlace to Put New Face on Smalltalk-80; Plans C++ Development Environment

ParcPlace Systems (Mountain View, CA), the spin-off of Xerox's Palo Alto Research Center, thinks object-oriented programming is the answer for big computer installa-tions bogged down in massive software projects. But first, the company has to remedy a major limitation of its Smalltalk-80 object-oriented programming language. Acceptance of Smalltalk-80 has been hampered by the fact that it runs in its own, incompatible windowing environment. Whether it's running on a Mac, an 80386-based microcom-puter, or a Sun workstation, Smalltalk-80 is not compatible with the host windowing system (e.g., Macintosh, Microsoft Windows, or X Window).

To overcome this problem, ParcPlace is working on a new interface called the Stencil Paint Imaging Model, which will include translators that "map" SPIM to the host imaging model (e.g., PostScript or QuickDraw). The company is also adding extensions to Smalltalk-80 that will allow it to make function calls to the host windowing system. ParcPlace Systems hopes to have the SPIM upgrade ready by November and plans to offer a run-time version of Small-talk-80, which will allow developers to install Smalltalk-80 applications without the entire development environment.

ParcPlace is also diversifying into the C++ object-oriented programming language, which lets programmers add object-oriented extensions to C pro-grams. In conjunction with Glocken-spiel, Ltd., the company is readying a complete C++ development environ-
continued

Best Performer!

FOXBASE+/Mac

New! Multi-User FOXBASE+/Mac

The Fastest, Most Powerful, dBASE Compatible Relational Database Management System For The Apple Macintosh!™

New! Version 2.00

Fox Software
Nothing Runs Like a Fox

Mac-WEEK 1988 TOP 10 BEST NEW PRODUCTS

INFO WORLD 1988 PRODUCT OF THE YEAR

MacUser 1988

FoxBASE+/Mac: The Most Celebrated Mac DBMS Now Includes a Powerful Report Writer!

New Version 2.00! New FoxBASE+/Mac Version 2.00 is here—and stealing the show! In its first year, FoxBASE+/Mac won more awards—both in the U.S. and Europe—than any other DBMS ever created for the Macintosh! And Version 2.00 is faster and more powerful than ever—packed with innovative new features and language enhancements!

New FoxReport!

FoxReport lets you create virtually any columnar or free-form report—without any programming! FoxReport includes:

■ **Page Layout:** Designate many layout settings for your report: number of columns, left margin setting, column width, space between columns, measurement size for each page, and more!

■ **Report Layout Window:** Define the different areas within your report, using FoxReport's new "band" system. Start with the default bands: Page Header, Body, and Page Footer, then bring in other bands like Title, Summary, Column Headers and Footers, etc.

■ **Object Menu:** Control the Type Font, Sizes and Styles of text objects. Fill and Pen options let you change the color and shading of almost every object!

■ **Page Preview:** "What-You-See-Is-What-You-Get" design ensures that the report on your screen is *exactly* as it will appear in print!

■ **Label Generator:** Create labels of almost any size and configuration. You can even include pictures *anywhere* on the label space!

New Features!

■ **Enhanced XCMD/XFCN Support:** Load up to 16 external XCMD's and XFCN's, then access them directly from within FoxBASE+/Mac 2.00 programs!

■ **Language Enhancements:** Other features include: support for custom hierarchical menus, new commands to help handle resources, new system functions to improve filename and screen management, and much more!

Perfect Connectivity!

FoxBASE+/Mac 2.00's Multi-User version allows complete data and application sharing between networked Macs and PCs! (when used with FoxBASE+/LAN). It also runs on the most popular network systems: AppleShare, Novell and 3Com!... *And It's Faster Than Ever!*

Order or Upgrade Now!

FoxBASE+/Mac Version 2.00 is the *new* Leader of the Mac! And at only $495 ($695 for Multi-User), it's a phenomenal bargain! If you're already a FoxBASE+/Mac registered user, you can upgrade to Version 2.00 for just $75!

To order your copy (or ask for a *FREE* demo disk), call (419) 874-0162. Or visit your local quality software dealer.

After all, when it comes to the Macintosh, *Nothing Runs Like a Fox!*

FoxBASE, FoxBASE+/Mac, and FoxReport are trademarks of Fox Software. Macintosh is a trademark of Apple Computer, Inc.

Fox Software

Fox Software, Inc.
134 W. South Boundary
Perrysburg, OH 43551

(419) 874-0162, Ext. 320
Fax: (419) 874-8678
Telex: 6503040827 Fox

World Radio History

Smalltalk/V.®
Designed to blow the doors off the hybrid languages of the programming world. Smalltalk/V does prototyping the same way Shelby prototyped the Cobra... using a blend of technical expertise and seat-of-the-

"Anyone can build a prototype by the by the seat of your pants." — CARROLL SHELBY
Creator of the legendary Shelby Cobra

pants savvy that's startlingly sophisticated. First you doodle, design, dream. Then you explore the possibilities and begin to assemble the

High

| Pure Object Oriented Programming |
| **Smalltalk/V** |

Complexity Control

| Hybrid |
| **Pascal with objects** |
| **C++** |

| Traditional |
| **C** |

Low   High

Exploratory Capabilities

prototype. You test. You tinker. You change. And you keep on changing and test-driving and refining until the prototype is just the way it was meant to be. With no compromises... of any kind. But the most remarkable thing is this prototype is not just a prototype. It runs, it races, it performs like the real application. Because it is the real application. And you

achieve this feat without once having to go through the old "crash and burn" kind of programming so common with languages born in the age of mainframes.

COMPLEXITY CONTROL FOR THE 1990s AND BEYOND.

The concept behind an object-oriented programming system is relatively simple. You build more complex objects out of simpler ones. Much as you can build complicated designs with a Lego set. With Smalltalk/V a programmer can write a piece of code and then

World Radio History

hard disk drive. Microtech will give you as much as $200 for the old drive. Nova disk drives run in size from 20 megabytes up to 320 megabytes. The company has several bigger drives in the works.

Sharp Electronics (Mahwah, NJ) has new software cards for the **Wizard** hand-held computer. The 32K-byte and 64K-byte cards double and triple the Wizard's memory capacity in its memo, schedule, and telephone modes.

If you work at home, you may sometimes lose a fax or modem transmission because someone accidentally picks up the extension phone to make a call. Interruption Blocker, from **DesignTech International** (Springfield, VA), is a telephone line guard that locks in your transmission and prevents others from breaking onto the line. You plug your phone into the mouse-size box, which connects to your telephone wall jack. Interruption Blocker costs $15.

Hitachi and **Motorola** are becoming the semiconductor industry's version of the forever-feuding Hatfield and McCoy clans. In the most recent shot fired, Hitachi accused Motorola of violating Hitachi patents with its **68030** microprocessor, currently Motorola's top-of-the-line CPU. It all started back when the two companies, after a long spell of technology exchanges and cross-licensing agreements, hit a snag involving Hitachi's new H-series of microprocessors. Motorola sued the Japanese company, claiming that the H-series infringes on Motorola patents. Hitachi countersued, claiming Motorola violated Hitachi patents with one of its microcontroller chips. Motorola fired back, claiming Hitachi violated a Motorola patent in that same case.

Expenditures for computer software in the U.S. will hit $61 billion by 1993, says the market research firm **Input** (Mountain View, CA). In its latest report, Input says that's an increase of $36 billion over what was expended on software products

continued

ment that's written in Smalltalk-80 and shares two of that language's major features, incremental compiling and linking. The C++ product will initially be available only on Sun workstations.

ParcPlace Systems is working on database "hooks" for Smalltalk-80. The integrated database capability will let you store Smalltalk-80 objects in a standard relational database. The first target database is Oracle, but ParcPlace plans to offer "back-end drivers" for other databases such as

DB/2 and Sybase. The ability to access reusable objects as fields in a database has great potential in many applications.

"We're moving from the lunatic fringe to the Fortune 1000 market," said Doug Pollack, ParcPlace's vice president of marketing. Big financial institutions and corporate MIS departments are looking for ways out of the "software crisis" and are considering more revolutionary techniques and approaches, as represented by Smalltalk-80, he said.

Toshiba's Low-Cost Systems Could Boost SPARC

I f Toshiba actually delivers relatively low-cost computers based on the SPARC processor, as the company suggested when it announced that it would adopt the SPARC chip, it could be the best thing to happen yet to Sun Microsystems' RISC architecture. Toshiba's Computer Division has signed a deal with Sun to manufacture computers based on the SPARC (for scalable processor architecture) standard. Toshiba will build a "new class of high-performance, low-cost computers" based on SPARC processors, the company said. Toshiba will also license Sun's SunOS version of Unix and the Open Look graphical user interface.

SPARC processors are available from several chip manufacturers, including Bipolar Integrated Technology, LSI Logic, Cypress Semiconductor, Fujitsu, and Texas Instruments.

Toshiba said that it will announce its first SPARC machines in early 1990. Sun spokesperson Marty Coleman said Toshiba does not intend to compete directly with Sun's own workstation offerings but intends to "complement" them. While Sun has

been fairly successful with the SPARC chip, SPARC-based machines still account for less than half of Sun's business, and Sun hopes to make SPARC a standard by recruiting other vendors to license the technology.

Toshiba is not the first major computer manufacturer to license the SPARC technology. AT&T, TI, Unisys, and Xerox, among others, have publicly announced commitments to SPARC. However, Toshiba has greater experience in delivering machines to the mass market than do the other SPARC licensees, and it also has a stronger presence in the personal computer marketplace with its line of laptop computers.

The only manufacturer other than Sun that is delivering SPARC-based machines is Solbourne Computer (Longmont, CO). As Coleman put it, "Solbourne proved that SPARC is clonable. Toshiba will prove that it can be produced in volume." The rather quiet Solbourne reported recently that it has signed distribution agreements, worth $19 million, with computer sellers in Australia, Greece, Taiwan, and Israel.
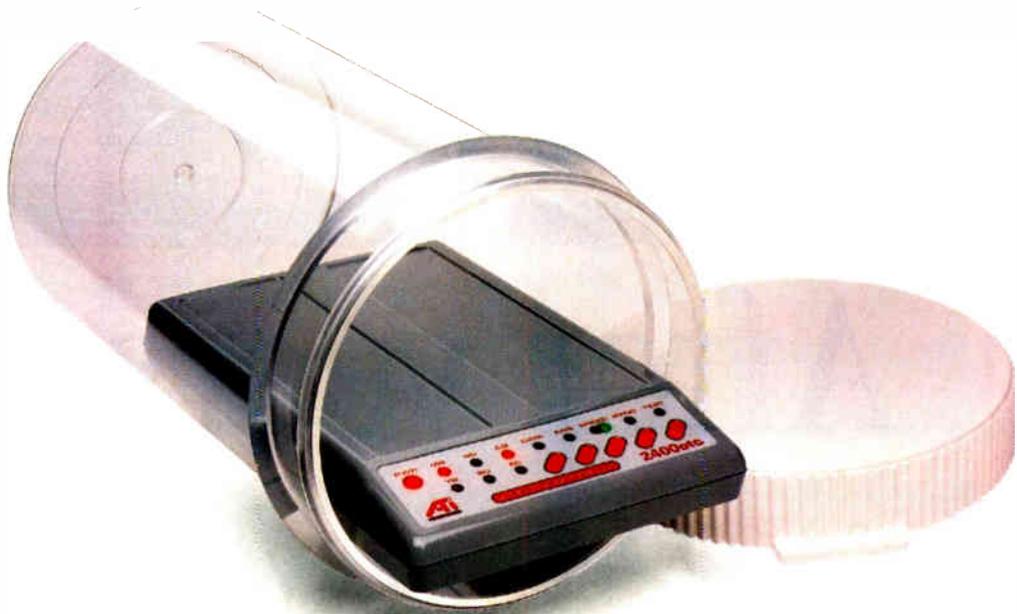
Silicon Graphics Cuts Price of 3-D Workstation

W orkstations with sophisticated graphics capabilities continue to bump down in price as high-end personal computers seem to be bumping up. In the latest indication of the workstation's improving price/performance curve, Silicon Graphics (Mountain View, CA) last month cut the price of its Personal Iris system by as much as 35 percent. The Personal Iris is a Unix-based graphics computer

built around MIPS Computer Systems' 32-bit R2000 RISC processor. The top-of-the-line model is capable of real-time three-dimensional imaging. Silicon Graphics rates the entry-level system's performance at 10 million instructions per second; a new model, based on the MIPS R3000 chip, performs at 16 MIPS, Silicon Graphics says. (For details about the

continued

World Radio History

in 1988. The study projects that about $25 billion of that total $61 billion will be spent on microcomputer software. The rest will be split between mainframe and minicomputer products.

In other positive news, another research firm claims that more than half the **small businesses** in the U.S. now use **personal computers**. According to **CAP International** (Norwell, MA), its latest national survey of firms with less than 100 employees found that 52 percent have personal computers. When the company took the survey a year ago, the number was 46 percent. A CAP analyst said that 300,000 small businesses bought their first personal computer last year, and of the 7 million microcomputers that CAP says were sold in the U.S. in 1988, 2 million were bought by small businesses and home offices. About 14 percent of the companies said they intend to buy a personal computer in the year ahead.

Rupp Corp. (New York City) has a new software hard disk drive lock that's designed to protect data against unauthorized access. Instead of encrypting each file, the FastLock utility ($69.95) encrypts only the file allocation table (FAT) of your hard disk. The table, which DOS uses to find individual files, makes the entire disk unusable when it's encrypted. A single password decrypts the FAT. If someone makes three unsuccessful attempts at giving the password, the computer locks up.

For your eyes only: **SkiSoft Publishing** (Lexington, MA) has developed software that lets you enlarge the text on a computer screen by as much as 300 percent, something laptop users in particular might find helpful. The memory-resident Eye Relief ($295) can display text on the screen ranging from the normal 80 columns by 25 rows to up to 33 columns by 7 rows. It also lets you change the space between lines and the space between letters. SkiSoft says that as you increase the size, the letters are tuned, so they have smooth edges instead of jagged ones.

Personal Iris, see "Silicon Graphics Brings Down Cost of 3-D Graphics," November 1988 Microbytes.)

The company cut the price of an entry-level diskless Personal Iris from about $18,000 to $12,500. The top system, with features for real-time three-dimensional operations, a 380-megabyte hard disk drive, and a 14-inch color monitor, has dropped in price from $35,000 to $25,500. The R3000-based models cost $4000 more. Starting at $12,500, the Personal Iris falls in the same price zone as souped-up 80386-based systems and color Macintoshes that have fewer graphics capabilities.

Silicon Graphics has also added something to the Personal Iris package. Each system now comes with Wavefront Technologies' (Santa

Barbara, CA) Personal Visualizer, a three-dimensional rendering and animation software package. The Personal Visualizer represents Wavefront's first entry into "lower-end" three-dimensional rendering and animation packages. The Visualizer is a menu-driven system for generating photo-realistic images from three-dimensional data.

There's a good chance that Silicon Graphics' three-dimensional imaging technology and Wavefront's graphics software will show up in the next version of IBM's RT PC, which sources say will be coming soon. IBM has licensed Silicon Graphics' Geometry Engine and Graphics Library technology, which many industry observers agree IBM will use first in the next model of the RT.

Reusable Objects Coming for PM Developers

Although developers aren't exactly screaming for tools to build OS/2 Presentation Manager programs, Eikon Systems (Foster City, CA) wants to be ready if developers start making the jump to the IBM/Microsoft graphical operating system. For now, though, Eikon is selling sets of reusable graphical objects for the Microsoft Windows development environment; the company plans to offer PM versions soon, said Eikon president Kevin Welch.

Eikon's Standard Control Pak focuses on Windows objects that can be used in dialog boxes or as a child of another window. It includes three classes of control icons: the palette for display and control of colors, push-button arrows, and picture frames for displaying bit maps and metafiles.

For each class, the Standard Control Pak includes a dynamic link library and a sample application. Just one of these objects costs $125; source code costs $475. The Tools Control Pak ($175, or $525 with source code) contains three additional control classes: a slider bar for selecting a value within a range; several kinds of

rulers for defining position and spatial orientation; and a toolbox that displays an array of small icons, each representing an operation.

The Resource Scrapbook "handles anything that moves in a file in Windows," Welch said. It allows you to create and manage files that pass through Windows' shared memory block, the clipboard, and supports "all commonly used" clipboard formats, including color bit maps, PostScript text, TIFF, SYLK, DIF, and CSV, he said. It also lets developers trap Windows resources, such as cursors, icons, and dialog boxes.

In the PM version of the Control Paks, the Resource Scrapbook will contain a facility to convert Windows bit maps, icons, cursors, and dialog boxes into PM equivalents.

Due from Eikon this fall is a Windows program generator, currently called Modern Art, that's designed to allow nonprogrammers to assemble icons and connect them with arrows, creating code on the fly that can be tested interactively. "It's much more sophisticated than NeXT's Interface Builder," Welch claimed.

NEWS STAFF SEEKS NEWS. DIAL (603) 924-9281.

The BYTE news staff is always interested in hearing about new developments that might affect microcomputers, the way they work, or the way people work with them. If you know of a project that could shape the state of the art, please give us a call at (603) 924-9281 or write to us at One Phoenix Mill Lane, Peterborough, NH 03458. An electronic version of Microbytes, offering a wider variety of computer-related news on a daily basis, is available on BIX.

EPSON EQUITY 386/20

UNLOCK
LOCK

POWER   HDD

World Radio History

LETTERS

and Ask BYTE

Graphic Details

I was very interested in your Product Focus entitled "Graphic Details" by Stanford Diehl and Steve Apiki (January). For four years, our company, Rock Technologies, has been developing and marketing graphics software that is specifically tailored to the construction industry to graphically display material quantity take-offs. Our software includes a utility to install drivers for all the digitizers mentioned in your review.

I thought the article was informative and well done, but I was disappointed that you chose not to include an evaluation of Science Accessories products. Other than a brief paragraph in the text box "Digitizers with a Twist," the company was not mentioned. The authors could have mentioned that Science Accessories' and Rock Technologies' combined technologies have produced what I would consider the only truly portable large-area digitizer. Roctek's RD-48 sonic digitizer will digitize 48 by 36 inches and will fit into a case small enough to be considered carry-on luggage at any airline counter in the world.

Gerry S. Ball
Chairman, Rock Technologies Corp.
Chandler, AZ

80286 vs. 80386

I operate my business with four computers—two Tandy 2000S 80186 machines (almost the same as 80286s) and two 80286 machines. I write most of my own software using LMI FORTH+, but I bought the 80286 machines to give me access to the great pool of elegant mathematical software out there if I ever needed it. One of my machines is more or less designated to handle housekeeping, file indexing, budgeting, and so on, leaving me three machines for projects. If any operation gets so long as to be tedious to wait for, I just use another machine until it's ready.

Now, I like to read about these 80386 machines, but I don't see how I could justify one, at least not until I need more than 14 megabytes of core memory. I have multitasking without the benefit of

an esoteric operating system devoted mostly to overhead. I also have 100 percent hardware backup. 80286 hardware is so cheap now that hardware solutions can be cheaper than software solutions. Mind you, as a programmer and a hardware nut, I'd love to have a true 32-bit machine to boss around, but the requirements are going to have to be stronger and the machines a lot better before I will jump. I'm not trying to say that there are not many valid applications today where an 80386 is appropriate and necessary, but they're not on my desk yet.

Roger Cain
Ottawa, Ontario, Canada

Windowing in Pascal

I purchased my first copy of BYTE in 1982, and I've been a regular reader ever since.

A recent article that I enjoyed was "Turbo Pascal Windowing System" (February), in which Charles J. Butler detailed a windowing system he developed for Turbo Pascal. This caught my interest because I would like to use programming tools in my applications. By "programming tools" I mean libraries or routines that add functions such as windowing, pull-down menus, and B-tree file indexing to languages like Pascal, BASIC, and C.

I have developed several simple routines that I use in my programs. These routines allow me to control the user interface and to use indexed files. Though these routines serve my purposes, I'd like to replace them with programming

continued

tools that would allow me to provide more functions in my programs and to streamline my code. I'd like to see BYTE do reviews and comparisons of programming tools for personal computer languages.

Randall L. Babcock
Sidney, NE

On Bridges

Regarding Peter J. Kulik's letter (May) responding to my article "When One LAN Is Not Enough" (January), I would like to make several points. Specifically, Kulik made the following incorrect or misleading statements:

• *Bridges do not assume anything about any upper-layer protocols.* As I mentioned in my article, proper operation of the bridge assumes that two communicating hosts share the same set of upper-layer protocols. Otherwise, cooperative data exchange will not occur.
• *My article discusses physical-layer bridges rather than Media Access Control (MAC)-layer bridges.* In fact, the opposite is true. A physical-layer "bridge" is in fact a repeater, and the term *bridge* is

not generally used in that case.
• *Bridges can be found between dissimilar networks.* By this, Kulik means (this is clear from a subsequent portion of his letter) that a bridge can link two networks that use different layer 1–3 protocols (e.g., an 802.5 LAN and an X.25 network) and convey data from a host on one network to a host on another. This is in fact a router. The generally accepted definition of a bridge, and the one that has been standardized by the 802.1 committee, is "a device that links networks that have a common MAC service interface."

William Stallings
Prides Crossing, MA

Controller Comments

With reference to Basse O. Bondtote's letter (Ask BYTE, March), I wish to offer the following comments:

1. If you are using a run-length-limited (RLL) controller set in the translation mode, try to set the jumpers in your Western Digital hard disk drive controller card to the nontranslation mode. You can obtain the appropriate instructions

from the *Disk Controller User's Guide*. Contact Western Digital Corp (2445 McCabe Way, Irvine, CA 92714, (714) 863-0102) for a copy of the guide.

It is interesting to note that both Gibson Research's SpinRite and Prime Solutions' Disk Technician Advanced will not work with RLL controllers set in the translation mode.
2. Try varying your interleave factor up to 1-to-10. You may see an improvement after 1-to-6, especially if you are presently using an accelerator card. An accelerator card slows down memory transfers to and from the expansion bus. Try this out (using an appropriate benchmark) both with and without your accelerator connected but with your original IBM PC XT or clone motherboard in the Turbo mode.
3. You may want to consider changing your hard disk drive controller to one that can really squeeze the maximum data transfer rates from your hard disk drive.

I wish BYTE could run a comparison of the various commonly available disk controllers for the IBM PC and XT. I'm

World Radio History

sure this would be of interest to readers. How about it? We know the future may belong to the 80286s, 80386SXes, and the 80386, but don't forget the 8088 and 8086 crowd.

Liew Kern Tote
Petaling Jaya, Malaysia

Why Smalltalk? Why Pascal?

I am writing with reference to "Smalltalk Can Be Cheap" by Don Crabb (April).

I realize that the article only compared Smalltalk to Pascal (presumably to point out the reduction in programming code necessary to achieve the stated objective), but I would like to make one additional comparison for the same purpose.

Why not BASIC? The listing below is written in standard BASIC. It is more than 30 characters shorter than the Smalltalk program, without relying on predefined routines. And at a quick glance, it is far more comprehensible.

```
10 DIM A%(255)
20 INPUT " ",L$
30 L%=LEN(L$)
40 FOR I%=1 TO L%
50 C%=ASC(MID$(L$,I%,1))
60 A%(C%)=A%(C%)+1
70 NEXT I%
80 FOR I%=65 TO 90
90 PRINT CHR$(I%),A%(I%)+A%(I%+32)
100 NEXT I%
110 ERASE A%
120 END
```

Norm Leo
Chatsworth, CA

Index Information

Sometimes I have to search long and hard for previous articles in BYTE. For example, in the March issue, the In Depth section resource list cites Actor software. I remembered a previous article on Actor, but it took me some time to find it (in the September 1987 issue). It would have been convenient to find this reference in the March issue. The indexes at the end of each issue are very useful for this kind of search.

I believe that BYTE publishes annual indexes. How can I get them?

P. Y. Narvor
Nantes, France

BYTE indexes for 1983–1984, 1985, 1986, 1987, and 1988 are available for $3 each by writing to: BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.
—Ed.

Random Access Thinking

For many years, the thought that the human mind is a fantastic computer has impressed me greatly. Every so often, I think it advisable to pay it humble respect. Consider the human mind's following capabilities:

• *Random access.* How many computers can instantly recall names and events scores of years in the past and reconstruct relationships involving the seemingly unrelated?
• *Evaluation procedure.* How many times have you collected a bundle of facts and data, each having different worths, and quickly processed them mentally to arrive at a meaningful answer? Do a decision table sometime when you have a decision to make that proves difficult to resolve.
• *Robotics.* Have you ever given any thought to what would be involved to program a robot to tie a shoelace or twist a pretzel? Just imagine the number of

continued

Circle 187 on Reader Service Card

elements that need to be coordinated to accomplish these simple tasks.

• *Sensors*. Have you ever been frustrated in an attempt to get an accurate measure of hot steel in a mill where scale and steam are a constant variable, and then found that a mill operator accurately called the correct temperature merely by visual judgment?

• *Management*. Have you ever considered what goes into a management decision in putting a team together, whether it's for sports or business? Imagine the mind's ability to weigh such factors as compatibility, health, and attitude.

Finally, we must realize that those ever-faster, ever-impressive machines that we build are, in the end, a product of the human mind.

Selwyn V. Stickler, Jr.
Vero Beach, FL

Learning Through Experience

This letter is in regard to a letter I sent previously that referred to an article on the Turbo EMS program in the Short Takes section of the March BYTE.

In that letter, I stated that the primary reason that I purchased Turbo EMS was so that I could run Fontasy. I should have taken a little more time to investigate before I wrote that letter. As it turns out, Fontasy will not run entirely correctly.

It seems that when running Turbo EMS using the hard disk drive mode rather than RAM, Fontasy has problems. The conflict prevents Fontasy from printing correctly. Apparently, Fontasy can print only the graphics that are in the RAM window at any particular moment. When the next window is called from the hard disk drive, the program locks up. Also, problems arise when you attempt to use those fonts that ordinarily would require the expanded memory. Here again, we are dealing with hard disk drive access. Regular fonts seem to work correctly.

The folks at Fontasy say that it is a timing problem. Accessing the hard disk drive simply takes too long. Fontasy cannot wait around that long. I can tell that Fontasy is, in fact, using the expanded memory. The way it paints the screen is a dead giveaway. In most respects, it works as it should, other than the above mentioned items. Things run very slowly, however. I have no complaint with this because I knew things would be slow from the start.

Obviously, if I cannot print anything, there is no reason to use Turbo EMS with Fontasy. On the positive side, however, I have gained the ability to use the ex-

panded memory in other ways. These uses justify Turbo EMS in my situation.

I can now use WordPerfect and all my memory-resident programs without other unacceptable problems cropping up—with one exception. On my Leading Edge Model D, I can run MemoryMate resident along with Hot Line and PC Tools Deluxe. Each program stays out of the other's way. Not so when I try to run MemoryMate with these same programs on my NEC PowerMate 286 portable. MemoryMate gets stepped on in some way. The program can only paint line graphics on-screen. It cannot paint its proper display. Other than hardware differences, I have no exact idea why the identical configuration runs on one computer and not on the other.

Fortunately, I have very little need for MemoryMate on the NEC. I have created a batch file that clears the other programs out of the way when I need to use MemoryMate. It works quite well. I do use MemoryMate a lot on the Model D, however, with no problem. So things have worked out OK.

At this time, I have not found a solution for the problem with Fontasy. The folks at Lantana Technology said that they would look into the trouble, but I haven't heard anything from them yet. Here, too, I have created a batch file that clears the way for Fontasy. It's a shame that I have a need for these files. Otherwise, everything works OK.

Charles T. Foley
Hixson, TN

Controller Correction

I would like to thank BYTE for publishing Jeff Holtzman's interesting article, "Advanced Floppy Disk Drive Controllers" (March). One slight correction is that Manzana MicroSystems' Mux Card lists for $89.95 (or $99.95 with the 3rd Internal Cable). Although, as Holtzman pointed out, the Mux Card does "perform flawlessly" in both IBM PC XT and AT systems, an end user might make better use of a Manzana High Density Controller Card (HDC) in an XT system. The Manzana HDC replaces the original XT controller and supports up to four 360K-byte or 1.2-megabyte 5¼-inch floppy disk drives (or four 720K-byte or 1.44-megabyte 3½-inch drives) internally or externally. Although it lists for only $94.95, the HDC includes a device driver to support both capacities of 3½-inch disk drives. And with its built-in BIOS, it saves you a BIOS upgrade.

David Gluck
President, Manzana MicroSystems
Goleta, CA

A Certain Class of Problems

While responding to a plea from two NASA scientists who were working on a magnetic shock problem, I discovered a class of problems that defy conventional numerical methods. These are problems described by differential equations that require solutions by computer calculations. The conventional methods include all those that use approximations to the Taylor series, such as the Runge-Kutta methods and the predictor-corrector methods.

The conventional methods fail because they have poor approximations to the Taylor series. A true Taylor-series method can handle this class of problems. My article "The ATOMCC Toolbox" (April 1986 BYTE) and the latest ATOMFT version 2.50 are true Taylor-series methods.

Here are some sample problems. They are described by mathematicians as ones that do not satisfy the Lipschitz condition where the solution becomes zero. The meaning of this statement is, "The solution is unknown at zero." However, this is not what troubles the conventional methods.

In this collection, some problems can be solved correctly with computed solutions correctly traversing through the zero point. Some problems can be solved correctly with computed solutions correctly stopping before the zero point. Other problems cannot be solved by the conventional methods.

The clue about whether the solution of a problem should stop or continue beyond the zero point is provided by information about the singularities in the solutions. This information can also predict whether the conventional methods will succeed or fail.

In the list of problems, x is the independent variable, y' is the derivative of y with respect to x, and y'' is the second derivative of y with respect to x.

1. $y'' = -0.5 - y' - (\cos(x) + y')/y$, start $x = 0$, end $x = 1.5$, with $y(0) = 1$ and $y'(0) = -0.95$.
2. $y'' = y' - y**(1/3)$, start $x = 0$, end $x = 3.5$, with $y(0) = 1$ and $y'(0) = 1$.
3. $y'' = -0.5 - y' - (\cos(x) + y')/y$, start $x = 0$, end $x = -0.43$, with $y(0) = 1$ and $y'(0) = 1$.
4. $y' = (x - 1)/y$, start $x = 0$, end $x = 2$, with $y(0) = 1$.
5. $y' = (2x**3 - 2x)/y$, start $x = 0$, end $x = 2$, with $y(0) = 1$.
6. $y' = -\sin(x)*\cos(x)/y$, start $x = 0$, end $x = 2$, with $y(0) = 1$.

What we lost in pounds we gained in power.

When the readers of PC World magazine once again voted our lightweight T3100 "Top Laptop," we were flattered.

And when PC Week gave it their Corporate Satisfaction Award, we were equally honored. Everyone seemed to agree it was perfect.

Except us. Because we believe there's always room for improvement, we made a few not-so-minor adjustments to what's always been our most popular machine.

First, we reduced its weight by nearly two pounds and renamed it the T3100e. (The "e," by the way, stands for enhanced.)

Then we increased its speed from 8 to 12 MHz and gave it a new 20MB hard disk with an access rate that's three times faster.

Finally, we improved its versatility with an IBM-compatible expansion slot that lets you access LANs, mainframes and more.

All of which leads to one very powerful conclusion.

Our loss is definitely your gain.

For more information call 1-800-457-7777. Toshiba PCs are backed by the Exceptional Care Program (enrollment required). IBM is a registered trademark of International Business Machines Corp.

T3100e: 12MHz 286 with 80287 coprocessor socket, internal half-length IBM slot, 20MB hard disk with 27msec access, 1MB RAM expandable to 5MB, gas plasma display, 1.44MB 3½" diskette drive.

In Touch with Tomorrow
TOSHIBA

Toshiba America Information Systems, Inc., Computer Systems Division

World Radio History

Problems 3 and 4 can be solved correctly by the conventional methods. Problem 5 can be solved only with the highest precision. Problems 1, 2, and 6 cannot be solved by the conventional methods. It may interest you that in problem 1, the numerator $(\cos(x) + y')$ is zero at the same point where $y = 0$.

The most disturbing fact is not that the conventional methods fail to solve some of these problems, but that they do so without any warning. A robust numerical algorithm must be able to identify those problems that it cannot solve and to inform its user. The conventional methods are not robust for this class of problems. For example, their solution for problem 6 at $x = 2$ has the wrong sign! I am very interested to hear what happens when you apply your favorite method to these problems.

By the way, even ATOMCC has difficulty with problem 1. The newest ATOMFT version 2.50 handles it with ease.

Y. F. Chang
Claremont, CA

The Amiga and Multitasking

I found Phillip Robinson's overview of PS/2, Macintosh, and Amiga graphics ("Variations on a Screen," April Graphics Supplement) interesting but, in the case of the Amiga, misleading in two ways.

First, Robinson writes, "Because [the trio of custom ICs] can handle video information while the main CPU is working on other tasks, the Amiga has a degree of 'multitasking'—the ability to handle more than one job at a time." The custom ICs do indeed relieve the CPU of a large part of the graphics burden. But they are not what makes the Amiga a multitasking machine.

The Amiga is multitasking from its software foundation up. The system software allows multiple tasks to share the resources of the Amiga, both hardware and software, simultaneously. There are sophisticated means for intertask communication. Each task gets a share of CPU time in a manner that is transparent to it. The programmer need not make allowances for multitasking other than not hogging system resources. And each program thinks it is the only one running on the Amiga. Some programmers are impolite enough to write programs that take over the Amiga, effectively disabling multitasking, but that is restricted mainly to games.

Second, the photo of the Amiga's screen that was used to illustrate Amiga graphics was a poor example of what

World Radio History

makes it so popular in the graphics world. Your photo looked like a CGA screen, not at all indicative of typical Amiga graphics. You could have used a hold and modify (HAM) image (e.g., girl with lollipop). Even the "low" resolution King Tut, with 32 colors, would have been better than an apparently four-color Flight Simulator screen.

Ron Charlton
Knoxville, TN

A Better Mousetrap?

I just finished reading the section on Unix in the May issue. I was thinking: What if one of the major computer companies, like IBM or AT&T, were to come up with a system that had some type of switch or button that switched between a Unix and a DOS board? Wouldn't a computer like this become a bestseller?

With such a machine, when Unix takes over for DOS as the standard system—as has been theorized—DOS could go out slowly, as many have predicted it will do. Or, even better, these two systems could work side by side.

Aaron Turpen
Alpine, UT

ASK BYTE

Losing Memory

I own a Dell 310 (an AT compatible). I'd like to write a program that will tell me how much RAM is currently being used by memory-resident programs. Can you suggest a book and/or technique that will help me?

Michael Beaupre
Minneapolis, MN

You don't need a program to do that. CHKDSK will work. Simply execute CHKDSK right after bootup, and it will tell you how much memory is available after DOS is loaded. Then load a TSR program whose memory use you wish to determine, and execute CHKDSK again. The available memory will have been reduced by whatever amount the TSR program has used.

It could be that your goal is to produce programs that are intelligent enough to determine the amount of available memory remaining after you've loaded all your TSR programs. You can determine

that by using DOS INT 21H, function 48H (allocate memory). Simply call this function, requesting 0FFFFH paragraphs' worth of memory. The call will fail (DOS doesn't have that much memory available), but in so doing, the interrupt will tell you the largest available memory block. Assuming you haven't removed any TSR programs since bootup (thus creating a "hole" in your memory map), the largest available memory block will be equal to the amount of remaining free memory.

You'll find information regarding this interrupt (and much more) in The New Peter Norton *(although I liked the old one just fine)* Programmer's Guide to the IBM PC & PS/2 *by Peter Norton and Richard Wilton (Microsoft Press, 16011 Northeast 36th Way, Box 97017, Redmond, WA 98073).—R. G.*

Unix Questions

Mark Minasi's article ("OS/2 for Cheap," April) was excellent. But how about an article on a Unix starter system for cheap? I've never used Unix before, so I'd probably want to start out with the

World Radio History

World Radio History

MKS Toolkit and my familiar IBM PC software. Still, I have a few questions about converting to Unix: Is my 80286 AT adequate for Unix, or should I switch to an 80386 system when I decide to use Unix exclusively? How much DRAM should the system have to run Unix comfortably? How much space does Unix require on a hard disk? Finally, what is the difference between Unix and Xenix?

Eric Alexander
Hudson, OH

The cheapest way to become familiar with Unix commands is to use a Unix shell on your DOS machine. MKS Toolkit is System V–oriented with a Korn shell (probably the new standard as of System V.4). PolyShell is Berkeley style. You can convert from a DOS 80286 AT machine to a Unix machine gradually (a more expensive proposition) by purchasing Xenix for the 80286. All you need is a total of 2 megabytes of RAM; it will work with less, but I don't recommend it. Unix is defi-

nitely more efficient and stable on the 80386. Unix is a 32-bit operating system, so the memory should be 32-bit memory. For most single-user work, 2 megabytes is enough. If you go for any DOS bridges (Merge or VP/ix), add 1 megabyte for each concurrent DOS process.

Since installing Unix on an existing system requires a Unix partition, you will have to repartition your disk; all data will be erased. The Unix partition should have a minimum of 15 megabytes of disk space, and you'll need more for the compiler and libraries, communications, text processing, and typesetting.

Where users are concerned, the difference between Unix and Xenix is small. Most 80386 Unix operating systems can recognize and run Xenix applications. Later this year, The Santa Cruz Operation will have 80386 Unix as a replacement for its 80386 Xenix. SCO will continue to support and develop enhancements for 80286 Xenix for a long time. I used 80286 Xenix for years and found it more than adequate. If you're starting fresh, go with an 80386 system. But if you already have a 80286 and don't expect to plug in a lot of concurrent users, just buy a little more RAM and 80286 Xenix.

MKS Toolkit
Mortice Kern Systems, Inc.
35 King St. N
Waterloo, Ontario N2J 2W9, Canada
(519) 884-2251

PolyShell
Polytron Corp.
1700 Northwest 167th Place
Beaverton, OR 97006
(503) 645-1150

SCO Unix/Xenix
The Santa Cruz Operation
400 Encinal St.
P.O. Box 1900
Santa Cruz, CA 95061
(408) 425-7222

—B. S.

FIX

In the response to a June Ask BYTE letter that discussed the difference between the 80386 and the 80386SX, we referred readers to our February cover story. We should have referred readers to our March cover story, "Battle of the Chips." We apologize for any inconvenience that this error has caused. ∎

Circle 244 on Reader Service Card

CHAOS MANOR MAIL

Jerry Pournelle answers questions about his column and related computer topics

The Future of Computing

Dear Jerry,

Having read the January 1989 edition of BYTE, I was surprised to see that you didn't have an entry in the feature article "What Lies Ahead." More than five years ago, you wrote an article for BYTE that was entitled "The Next Five Years in Microcomputers." And you probably thought we wouldn't remember.

Rereading that article provides an interesting lesson in how times change. Valdocs (remember the Epson QX-10?), the Osborne with bundled software, the Apple Lisa, the argument over 8086 versus 68000 microprocessors—those were the hot topics way back in 1983.

On the whole, you did a good job of prognosticating (though you should refrain from taking large sums to the racetrack). Do you care to pick up the crystal ball again?

Michael Anthony Kellar
Clifton Heights, PA

I sure do remember my five-year prediction. I think my projection of the growth of this industry was much higher than anyone else's—and mine was far too low. Back then, I was writing a column for the late and lamented Popular Computing *called "The Computer Revolution," and I meant every word of that.*

The Osborne did indeed change the nature of the industry; it had done so by the time I wrote that. Prior to the Osborne, you had to go to a store and get a pile of boxes, which you tried to integrate into a system that would work. (You could also get a TRS-80, if you were interested in being part of Tandy's Quality Assurance Department.) Adam's machine had too small a screen, but by gollies it did work, and it was portable.

I'll have to give some thought to another predictions column. Perhaps for September again, just to be symmetrical.

—*Jerry*

Paying for Technical Support

Dear Jerry,

Imagine this. You go and buy a new car. After you've had it for a few days, one of the windows won't roll down. It used to roll down just fine, but now it won't go down at all. You look through all the manuals that came with the car, but you can't find anything about why the window won't go down. So you take the car into the dealer and ask the salesperson why your window won't go down. The salesperson tells you that you accidentally hit a button on your front door panel that locks that window. He tells you that it's clearly stated in the footnote on page 259 of the manual. Then he hands you a bill for $12 for the 12 minutes you talked to him.

How many of us would just accept that bill? After all, we just paid a lot of money for the car, and who could find that footnote on page 259? After we paid all that money, we shouldn't have to pay just to have someone tell us how to use some of the more esoteric features. Yet how many times do we pay the same type of bills to software manufacturers? We spend hundreds of dollars on their software and then have some problem that we can't find in the manual, for which they charge us $1 per minute to talk to them!

I understand that companies get a lot of calls and that they have to hire people to answer phones, which costs a lot of money. But that's part of doing business, isn't it? How many manufacturers of anything but software do you know who charge you $1 a minute to talk to you if you have a problem with their product? I can't think of any offhand.

Most of the calls that companies receive are for problems that the users could have solved by looking in their manuals. However, most of the calls I have placed for technical support have been due to either actual bugs in programs or to awful documentation that

continued

Jerry Pournelle holds a doctorate in psychology and is a science fiction writer who also earns a comfortable living writing about computers present and future. He can be reached c/o BYTE, One Phoenix Mill Lane, Peterborough, NH 03458, or on BIX as "jerryp."

leaves out important information. Why should I have to pay to talk to the company about something that's really its fault anyway?

I wonder if these companies purposely write bad documentation, thus forcing users to incur extra expense. I wouldn't be at all surprised. Do they purposely put bugs in their programs?

Perhaps $1 per minute would not be unbelievably excessive if a user received decent help. But, in my experience, the usual answer to a question like, "Why does your program do this, when it's supposed to do that?" is, "Gee, I don't know" or "I've never heard that one." Rarely, that I can recall, has technical support been any help.

I think this problem of technical support ought to be the next big issue in computing. Now that we've mostly gotten rid of copy protection, let's try to get rid of these excessive and unfair support charges. After all, if a reasonably knowledgeable person has to spend more money on technical support than he or she spent on the program in the first place, then something is wrong with the program or documentation, not the user.

Kevin Clark
Front Royal, VA

I used to spend a lot of time berating software publishers for their lousy documentation; indeed, sometimes I still do. I am about to conclude that there are few companies willing to pay documentation writers anything like what they're worth. They pay the programmer, but the writer is left to the last minute. They also don't have anyone copyediting the documentation.

On the other hand, some programs, such as Traveling Software's LapLink, don't need documents.

As to technical support, some outfits do an excellent job of that. Two companies and their programs—Arts Computer Products' Word Perfect and Aldus's PageMaker—come to mind. The other day I called Aldus to get its press relations people. None were around, but the technical-support people were there, and did I want to talk to them? Of course, other outfits hire cretins to answer the phone, who then turn technical-support problems over to chimpanzees.

Finally, my friends in the technical-support business tell me that about half the calls are about problems unrelated to their product at all. The customer has 480K bytes of his or her system taken up with memory-resident programs or is trying to run an EGA program on a CGA system, or hasn't even plugged the machine in. —Jerry

Unix Developments
Dear Jerry,

I understand why you tell your readers that you're not a Unix expert. Unix is a big and powerful operating system, understood in its entirety by few. As a person who was an early IBM PC user but a recent Unix convert, I can appreciate the challenge of commenting on Unix in general and on Unix on the 80386 in particular. Please accept the following comments in that light:

• The Santa Cruz Operation's brand of Unix is Xenix, not Unix, but it will probably be called Unix in the future.
• The Unix/Xenix spreadsheets and word processor tend to be quite plain when compared to their DOS counterparts. But those are classic workstation applications that are best processed on essentially freestanding systems. Frankly, even with all the overlapping windows and color of DOS-based DBMSes, they tend to be rather pale in comparison to Unix-/Xenix-based DBMSes. Unix is, after all, designed basically for a multiuser environment.
• VP/ix from The Santa Cruz Operation and Interactive Systems also allows DOS to be run as a task under Unix. My company will be doing that in a production environment soon, if you're interested in the results. We have found VP/ix reasonably easy to install (not the same thing as simply typing INSTALL and pressing Return to accept the defaults, though).
• Unix does multitasking. Even on my relatively slow AT&T 3B1, it's possible to create multiple graphical windows and observe them working. For a more primitive demonstration, place multiple Unix processes in the background, and they'll process concurrently. The message you received informing you that the network version was required sounds like a special case.

Philip G. Duffy
President
Electronic Cottage Associates
West Chester, PA

I'm always interested in new developments. Thanks.

My quarrel with Unix is that while DOS applications do sort of run, they don't actually run very well in general, while the Unix-specific stuff is very vanilla compared to what is available under DOS.

But, then, the supercomputer people have to put up with writing their programs as 100,000 lines of FORTRAN using editors more primitive than we had under CP/M!—Jerry ∎

WHAT EXACTLY CAN THE WORLD'S MOST POWERFUL AND EXPANDABLE PC DO?

ANYTHING IT WANTS.

Now it's possible to do just about anything you can think of, faster than you can think.

Introducing the COMPAQ DESKPRO 386/33 Personal Computer. Never before has so much performance, expandability and storage been put into one desktop PC. And never before has one PC been capable of so much.

Inside its new system unit, you'll find that our engineers have redesigned just about every component to deliver a minicomputer level of power with unmatched PC flexibility.

So you can use it as a stand-alone PC, putting its power to work on the most demanding CAD/CAE, financial analysis, database management and other personal productivity applications.

Or you can spread the power around, using the COMPAQ DESKPRO 386/33 as the driving force for a network or multiuser system.

At the heart of the system is the Intel 386™ microprocessor. Running at a blazing 33 MHz, it works in concert with a series of technological advancements. Like a 33-MHz cache memory controller with 64K of high-speed static RAM. Interleaved memory architecture. And the exclusive COMPAQ Flexible Advanced Systems Architecture.

This high-performance combination delivers a 35% performance improvement in

CPU-intensive applications over 25-MHz 386 cache-based PC's.

Or said another way, nothing will slow you down. No matter what you want to do.

You can expand the 2 MB of standard RAM up to 16 MB using the high-speed 32-bit slot. That leaves up to *six* industry-standard slots free to customize the system to the demands of the application you're using.

If your job is particularly demanding, you can use up to five high-performance internal storage devices to hold up to 1.3 *gigabytes* of data. And if that's not enough, bring total system storage to 2.6 gigabytes with the optional COMPAQ Fixed Disk Expansion Unit.

There's more. You can run MS-DOS®, MS® OS/2, Microsoft® Windows/386 and the XENIX® and UNIX® operating systems. Access memory over 640K under DOS with the COMPAQ Expanded Memory Manager that supports Lotus/Intel®/Microsoft (LIM) 4.0. And speed through calculations with 33-MHz Intel 387™ and Weitek 3167 coprocessor options.

All the new advancements engineered into the COMPAQ DESKPRO 386/33 deliver an unmatched level of power, expandability and storage.

To do anything you want.

INTRODUCING THE COMPAQ DESKPRO 386/33

A total of eight expansion slots let you customize the system to your needs by expanding memory and choosing from thousands of industry-standard expansion boards.

High-speed VGA graphics are built in. And for greater graphics capabilities, add the optional COMPAQ Advanced Graphics 1024 Board.

Built-in interfaces make it easy to connect pointing devices, printers, plotters or other peripherals without using an expansion slot.

COMPAQ

It simply works better®

IN 386 PERSONAL COMPUTING, YOU'RE LOOKING AT THE MOST WANTED LIST.

In 1986, Compaq introduced the world to personal computers based on the 386 microprocessor.

Since then, we've made it possible for every level of user to work with this powerful technology. In fact, more people work with COMPAQ 386-based PC's than any other 386's worldwide.

Today, Compaq offers the broadest line of these high-performance personal computers. Each delivers significant technological advancements developed by Compaq engineers. Each delivers optimum performance for the needs of different users. And each is built to the highest standards for compatibility and reliability.

For power-hungry users who want 386 performance to go, the COMPAQ PORTABLE 386 Personal Computer does things normally reserved for a desktop 386 PC. Without compromise.

For people considering 286 desktops, the COMPAQ DESKPRO 386s Personal Computer is an affordable way to move up to 386 performance. And if you have a 286 that you've outgrown, the COMPAQ DESKPRO 386/20e Personal Computer

is an easy step up to the power and capabilities of a 20-MHz 386 machine.

For the increasing needs of today's 386 users, the COMPAQ DESKPRO 386/25 Personal Computer offers advanced performance. And for those who desire the most power and expandability available in a desktop PC, the COMPAQ DESKPRO 386/33 stands alone.

For a free brochure on COMPAQ 386-based personal computers and the location of your nearest Authorized COMPAQ Computer Dealer, call 1-800-231-0900, Operator 93. In Canada, 1-800-263-5868, Operator 93.

COMPAQ

It simply works better.®

WHAT'S NEW

PS/2 clones feature MCA performance.

Entourage Ships MCA Clones

Whether you prefer IBM's Micro Channel Architecture, the AT architecture, or the proposed EISA bus architecture, MCA-compatible machines are becoming more widespread.

The 316XMC and the 320MC, MCA clones from start-up Entourage Computer Corp., feature better performance and lower cost than IBM's PS/2 models 50Z, 55, and 70, the company says.

System 316XMC is a zero-wait-state 16-MHz 80386SX with 1 megabyte of RAM that's upgradable to 8 megabytes using 80-ns single in-line memory modules on the motherboard. Built-in floppy and hard disk drive controllers support a standard 19-ms 40-megabyte SCSI hard disk drive and your choice of a 3½-inch or 5¼-inch floppy disk drive. There are three 16-bit expansion slots.

System 320MC is a zero-wait-state 20-MHz 80386 with 1 megabyte of RAM that is upgradable to 16 megabytes. It also comes (optionally) with the 40-megabyte hard disk drive and your choice of a floppy disk drive. Expansion is through two 32-bit MCA slots and a 16-bit MCA slot.

Both systems feature 102-key PS/2-style keyboards as standard equipment, but monitors are optional. Other options include math coprocessors, MS-DOS version 3.3 or 4.01, and OS/2 version 1.1.
Price: 316XMC, $1895; with hard disk drive, $2395; 320MC, $2895; with hard disk drive, $3395.

Contact: Entourage Computer Corp., 10919 Technology Place, Suite B, San Diego, CA 92127, (619) 673-8633.
Inquiry 1126.

Mitac Clones IBM's Model 55; Includes MCA

If you want an alternative to IBM's Model 55 SX, the recently introduced 80386SX-based Micro Channel Architecture microcomputer, you might try Mitac's MPS2386. It has the same 80386SX CPU, an MCA bus licensed from IBM, and many of the same features.

One advantage of Mitac's system is that its base price—about the same as IBM's—includes an internal 1.2-megabyte 5¼-inch floppy disk drive. With IBM's Model 55, 5¼-inch drives can be supported only externally.

Mitac peripherals match IBM's also. The MPS2386 uses the same 640- by 480-pixel Tatung monitor, and it ships with 1 megabyte of RAM (expandable to 8 megabytes of 32-bit RAM on the motherboard).

Mitac's on-board VGA adapter comes from Paradise and is "auto-switchable" to EGA and CGA modes. There's also a dedicated mouse port on the back of the box.
Price: including keyboard and MS-DOS, $2995. Tatung Monitor, $549.
Contact: American Mitac Corp., 410 East Plumeria Dr., San Jose, CA, 95134, (800) 648-2287.
Inquiry 1127.

Unix Sidesteps with NEC

NEC's Astra XL/100, and XL/200 are 68030-based Unix machines that contribute to the unclear distinction between microcomputers and workstations.

In stripped-down configurations, both systems offer 25-MHz zero-wait-state perfor-

mance, 2 megabytes of standard RAM, 8K bytes of cache memory, and a 1.2-megabyte 5¼-inch floppy disk drive.

The bare-bones Astra XL/100, which is bundled with Unix System V, can support up to eight users. It includes an MC68881 floating-point coprocessor and memory expandable to 10 megabytes. It has eight free 32-bit Multibus slots.

One version of a preconfigured XL/100 includes an extra 2 megabytes of RAM, a dumb terminal with cabling, an 18-ns, 130-megabyte ESDI hard disk drive, a 150-megabyte tape drive with the operating system on tape, and an eight-port controller.

The XL/200 ups the ante a few notches by supporting up to 32 users with an optional software license. In its bare-bones configuration, it includes an MC68882 floating-point coprocessor and RAM expandable to a whopping 34 megabytes. It can be upgraded with an optional Advanced Terminal Subsystem add-in card to an XL/300 system to support up to 64 users.

A preconfigured XL/200 with a minimum of extras includes 6 megabytes of RAM, a dumb terminal with cabling, an eight-port terminal controller, an 18-ns, 130-megabyte ESDI drive, a 150-megabyte tape drive with the operating system on tape, and a license for up to 16 users.
Price: XL/100, $8995; XL/200, $13,995; configured XL/100, $15,995; configured XL/200, $23,995.
Contact: NEC Information Systems, Inc., 1414 Massachusetts Ave., Boxborough, MA 01719, (508) 264-8000.
Inquiry 1128.

SEND US YOUR NEW PRODUCT RELEASE

We'd like to consider your product for publication. Send us full information, including its price, ship date, and an address and telephone number where readers can get further information. Send to New Products Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458. Information contained in these items is based on manufacturers' written statements and/or telephone interviews with BYTE reporters. BYTE has not formally reviewed each product mentioned. These items, along with additional new product announcements, are posted regularly on BIX in the microbytes.sw and microbytes.hw conferences.

Jasmine's DAT Stores Gigabytes

The Jasmine DirectDigital Tape drive can store up to 1.27 gigabytes of data on a single 4-mm chromium dioxide tape that's less than the size of an audio cassette tape.

The sustained data transfer rate is 174K bytes per second through a standard SCSI port, and the drive is capable of locating a single byte of information anywhere on the tape in not more than 40 seconds, Jasmine claims.

The Jasmine DirectDigital Tape drive is a hybrid of many sources: It uses a half-height JVC drive mechanism, packaged by GigaTrend into a full-height drive with I/O; Racet provided the drive enclosure and the software.

The formatting standard is Data/DAT, also favored by Apple and NCR, rather than the DDS format pushed by Sony and Hewlett-Packard, Jasmine says. This permits the drive to find data quickly and to rewrite changes incrementally, instead of rewriting the whole file if a change is made to it. The drive has a recording density of 61,000 bits per inch (using a helical scan technique, like a VCR) with almost 1900 tracks per inch.
Price: $6995.
Contact: Jasmine Technologies, Inc., 1740 Army St., San Francisco, CA 94124, (415) 282-1111.
Inquiry 1130.

DirectDigital drive features DAT cassettes.

HP Offers a Mac-Specific Ink-Jet Printer

The new Hewlett-Packard DeskWriter printer is a modified DeskJet, a 300-dpi ink jet for the Mac market. Gone are the serial and parallel interfaces and the cartridge ports—replaced by an ImageWriter-style connector and QuickDraw compatibility.

DeskWriter comes with four built-in font families (Times, Helvetica, Courier, and Symbol); also available are five other font families.

Font scaling is through proprietary "Intellifont font-scaling technology," HP says. These are outline fonts, scalable to 250 points.
Price: $1195; optional fonts, $95 each or five for $395.
Contact: Hewlett-Packard Company Inquiries, 19310 Pruneridge Ave., Cupertino, CA 95014, (800) 752-0900.
Inquiry 1132.

Mac Drive Features Removable Cartridges

The Microtech R45 hard disk drive features 25-ms average access time and removable cartridges that can store as much as 42.7 megabytes of formatted information.

The interface is SCSI, and Microtech says it's compatible with the Macintosh Plus, II, SE, and IIx. It measures 3 by 10 by 11 inches. Included in the base price is the drive, cabling, and one SyQuest SQ400 cartridge.
Price: $1099; additional cartridges, $90.

Contact: Microtech International, Inc., 29 Business Park Dr., Branford, CT 06405, (800) 325-1895 or (203) 488-7744.
Inquiry 1131.

Computer CD-ROM Doubles as Audio Player

Chinon America thinks you should be able to use the same device for personal computer data storage and for audio entertainment. The CDS-430 drive lets your computer use Sony and Philips CD-ROM disks, with 530 megabytes of available storage space, as a data storage/replay medium and as a drive for your audio entertainment.

It's packaged in a 13- by 11- by 3-inch box and connects to your computer system through the SCSI port. Microsoft CD-ROM Extensions software enables reading of any disk written in the High Sierra format.

The system will automatically recognize whether the compact disk is ROM or audio and use the appropriate command format for either, Chinon says. It can also read a mixed audio/CD-ROM disk.
Price: $695; Extensions software, $150.
Contact: Chinon America, Inc., 660 Maple Ave., Torrance, CA 90503, (213) 533-0274.
Inquiry 1134.

Store 2.8 Megabytes on 3½-inch Floppy Disks

The Megamate 2.8 is a 3½-inch floppy disk drive that can store 2.8 megabytes on your average 3½-inch floppy disks. It works because of the proprietary floppy disk controller, the CompatiCard IV.

The CompatiCard IV controller allows Micro Solutions to double recording frequency. That means it writes twice as fast, writing to 36 sectors per track. Standard 1.44-megabyte drives write to 18 sectors, and 720K-byte drives use nine sectors.

The drive itself, manufactured by TEAC, is backward compatible with the standard floppy disk drives, including the PS/2 drives.

The CompatiCard IV controller is a half-length card that's XT and AT compatible. Each card supports up to four 3½-inch or 5¼-inch floppy disk drives. One cable is included.
Price: $395; or $149 for CompatiCard IV.
Contact: Micro Solutions, Inc., 132 West Lincoln Hwy., DeKalb, IL 60115, (815) 756-3411.
Inquiry 1133.

Accelerate Any Microcomputer to 33 MFLOPs

The Spirit-30 is a floating-point accelerator board for performance approaching 33 million floating-point operations per second. With the Texas Instruments TMS320C30 digital signal processor at its heart, it works with XT, AT, Macintosh, PS/2, VMEbus, Multibus II, and Q-bus based computers.

Each Spirit-30 includes 128K bytes of dual-access 25-ns static RAM that is accessible by both the host and the Spirit-30. A daughterboard gives you an additional 512K bytes of RAM, with as much as 16 megabytes of external memory available through the Spirit-30's parallel port and bus interface.
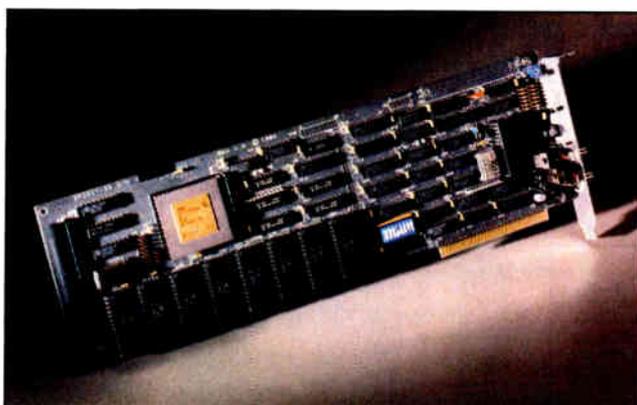
For data acquisition, memory expansion, frame grabber and graphics boards, the Spirit-30 has six expansion connectors. Multiple Spirit-30s can be configured through the serial or parallel ports.

Support software includes windows-based evaluation, debugging, simulation, and real-time digital signal processing (EDSP). A library (DSPL) is included to initiate program download, start DSP execution, halt DSP execution, and perform single-block read/write to and from the board and DSP memory.

In all, the DSPL gives you 35 DSP and utility modules in C. They incorporate functions like spectrum analysis, FFTs, and discrete cosine transforms.
Price: $2495; EDSP software, $495; DSPL, $295; source code, $985.
Contact: Sonitech International, Inc., 83 Fullerbrook Rd., Wellesley, MA 02181, (617) 235-6824.
Inquiry 1138.

Number crunch with the Spirit-30.

Replace Algorithms with Thinking Processors

Two companies recently claimed firsts in neural networking by offering commercial silicon implementations of popular neural networking theories.

Syntonic Systems introduced an XT-compatible evaluation kit with the Dendros-1 chip, an analog device that works with one of several popular neural net architectures.

Partially self-organizing, the chip stores "remembered" patterns in capacitors. Dendros-1 also performs a key calculation—input and weight vector multiplication—in parallel, achieving the equivalent of 4.3 MFLOPS performance, Syntonic says.

The way the nodes are wired determines the type of architecture in a neural net. The architecture in turn determines the learning algorithms. So if you're going to hard-wire a chip to speed up execution of a particular algorithm, you're stuck with it.

Dendros-1 implements a variant of the "adaptive resonance theory" (ART-1), a two-layer network architecture. It has three input layer nodes and five output layer nodes. It will accept up to 22 bi-level input signals, and these can be presented via a PC, although output from Dendros-1 is limited to an LED display.

Dendros-1 is packaged in a 68-pin plastic leaded chip carrier. An evaluation board includes eight chips.
Price: $695.
Contact: Syntonic Systems, Inc., 20790 Northwest Quail Hollow Dr., Portland, OR 97229, (503) 293-8167.
Inquiry 1135.

Micro Devices Implements Hopfield Neuron

The Fuzzy Set Comparator is a CMOS neural chip that's included in Micro Devices' neural networking kit, an XT-compatible add-in board. It implements the popular Hopfield theory of neural networking in silicon.

With the Fuzzy Set Comparator, the kit is designed for adaptive ranking and for ranking "fuzzy" data (data with inaccuracies, noise, or other discrepancies) in groups by certain predetermined characteristics.

Once the data is ranked, a neural network hardware postprocessor ranks the comparisons, thus providing a superior rank-calculation speed over software implementations of neural networks, Micro Devices claims. A built-in video interface also allows the Fuzzy Set Comparator chip to "see" and "identify" people.
Price: $250.
Contact: Micro Devices, 5643 Beggs Rd., Orlando, FL 32810, (407) 299-0211.
Inquiry 1136.

continued

Make Graphics a Whiz with the FastWrite VGA

The FastWrite VGA from Headland Technology is designed to be faster than the original FastWrite. It's fast, the company claims, because it uses an enhanced version of the V7VGA chip.

Headland claims the chip is 100 percent register-level compatible with the VGA standard, is BIOS-level compatible with the EGA standard, and is also backward compatible with CGA, MDA, and HGC standards.

Each V7VGA chip also features memory caching, 8- or 16-bit memory and BIOS interfaces, and support of four resolutions beyond the 17 standard VGA modes. Interlaced resolution reaches up to 1024 by 768 pixels, with up to 16 displayed colors.

Each three-quarter-length FastWrite VGA is configured with 256K bytes of on-board memory that's upgradable to 512K bytes. It comes packaged with software drivers for graphics-intensive applications: Windows/286, Windows/386, Presentation Manager, AutoCAD, AutoShade, Ventura Publisher, GEM/3, Lotus 1-2-3, and Symphony.
Price: $499.
Contact: Headland Technology, Inc., 46335 Landing Pkwy., Fremont, CA 94538, (415) 656-7800.
Inquiry 1137.

World Radio History

Patented thin-film coating reduces glare.

OCLI Shades PC Glare

For luxury eye protection from your personal computer, you might want the Glare/Guard from the Optical Coating Laboratory. OCLI recently expanded its line of add-on filters. The new models are designed specifically for the ubiquitous NEC MultiSync II, Sony's and Sun Microsystems' CAD/CAM monitors, and the Macintosh Plus, SE, IIcx, and II.

Using thin-film coating technology, OCLI applies layers of germanium, zinc sulfide, and a fluoride compound to tempered glass using a patented vacuum deposition coating process. Ion-deposition processes make the coating abrasion-resistant.

With such a filter, OCLI claims glare reduction of up to 99 percent, enhanced contrast, reduced static and dust, reduced perception of screen flicker, and low-frequency radiation level reduction of 98 percent.

Filters for all the new machines except the small Macintosh monitors come in two models: The Profile is designed to reduce glare by 95 percent and the ProfilePlus is designed to reduce 99 percent. For the Mac Plus, SE, and IIcx monitors, which have reversed type that requires more light emission, OCLI designed the Professional Plus Size M filters with glare reduced by only about 50 percent.
Price: For CAD/CAM monitors, Profile, $199, and ProfilePlus, $249; for the NEC and Mac II, Profile, $69.95, and ProfilePlus, $109.95; for the Mac, ProfessionalPlus, $89.
Contact: Optical Coating Laboratory, Inc., 2789 Northpoint Pkwy., Santa Rosa, CA 95407, (707) 545-6440.
Inquiry 1142.

Expand Automated Data Acquisition to 1 MHz

The 64-channel Enhanced Graphics Acquisition and Analysis (EGAA) system functions as both a high-resolution digital storage oscilloscope and an electronic chart recorder.

The hardware uses four IS-16 A/D add-in boards, each with 16-channel 1-MHz A/D conversion, creating a 64-channel system. The software can operate the EGAA system as four separate digitizers at a 1-MHz sampling rate or as one 64-channel system at 62 kHz. The system contains a variety of trigger logic functions such as slope and level. External triggers, like those on a digital storage oscilloscope, complement the pretrigger that's available to capture transients.

The chart recorder mode allows simultaneous real-time monitoring and storing of data to a hard disk.
Price: $3090; analysis options range from $485 to $1395.
Contact: R.C. Electronics, Inc., 5386-D Hollister Ave., Santa Barbara, CA 93111, (805) 964-6708.
Inquiry 1140.

Small Supplies Switch to Sine

The UniPower 4.5 and UniPower 6.0 are on-line systems that give you continuous power protection as well as power conditioning. Unison Technologies claims its true sine-wave output provides superior equipment protection over the square-wave output found in many supplies.

As the names imply, the PS 4.5 provides 450 VA of backup power, and the 6.0 gives you 600 VA. Both are relatively small, measuring 14 by 3 by 18 inches, and are designed to fit between your system and your keyboard. They weigh 30 and 32 pounds, respectively.

Both units provide patented emergency keyboard lights and a remote-on feature that lets you turn on your system over the telephone.
Price: UniPower 4.5, $699; 6.0, $799.
Contact: Unison Technologies, Inc., 23456 Madero, Mission Viejo, CA 92691, (714) 855-8700.
Inquiry 1143.

To Draw As an Artist Draws

Variable line width and airbrush density are just two of the features available with the Wacom pressure-sensitive and cordless digitizing system.

With the pen-like stylus, you press lightly and a slender line appears. Press more heavily and the line thickens as it would if you were drawing with a pencil or brush. Colors can be programmed so you can draw to fit your mood (i.e., red for the firmest pressure and blue for a light touch).

You can also use a cordless cursor, but you won't get the variable-line effects of the stylus. For both hand-held devices, reading speeds are selectable, up to 205 points per second. Tablet accuracy is rated at 0.2 mm, whether you buy the 6- by 12-inch tablet or the 18- by 25-inch tablet.

The system works through electromagnetic resonance technology, says Wacom. The digitizer tablet contains a fine grid of thin wires that alternately transmit and then receive their own signals, telling it where the pointing device has moved by reading from a coil-and-capacitor resonant circuit.

The stylus can produce variable line widths, for example, because of a movable ferrite core. Pressure on the stylus's point changes the inductance of the resonance coil and affects the electrical frequency.
Price: Stylus, $125; 6- by 9-inch tablet, including cursor, $395; 12- by 12-inch tablet, including cursor, $995.
Contact: Wacom, Inc., West 115 Century Rd., Paramus, NJ 07652, (201) 265-4226.
Inquiry 1141.

Announcing a big leap in mouse technology.

*F*inally a mouse with an extraordinary body and a mind to match.

*T*he Logitech Mouse is tuned to accelerate your cursor across any screen with the mere flick of a wrist. And slow it down on arrival for pixel-point control.

*I*t's <u>guaranteed</u> to work with all IBM personal computer applications.

*A*nd it comes with a great selection of MouseWare™ including Pop-Up DOS™— the ultimate DOS handler; the Mouse-2-3™ shell; 35 menus for popular keyboard-based applications; and unlimited Product Support.

*F*or your nearest dealer, call Logitech at:
800-231-7717

In California:
800-552-8885

In Europe:
++41-21-869-96-56

Circle 150 on Reader Service Card (DEALERS: 151)

LOGITECH

MOUSE

Pocket-size Adapter Links Laptops to NetWare Stations

The Pocket Ethernet Adapter is an Ethernet add-in card shrunk down to fit into a package the size of today's pocket modems.

Two versions are available, accommodating thick and thin coaxial cabling. Support for unshielded twisted-pair cabling should be available sometime this fall. Only Novell NetWare drivers for versions 2.0 and 2.1 are compatible today. But Xircom promises that future releases will include the latest Novell drivers and drivers for the other popular network operating systems. Drivers for 3Com's 3+ and 3+Open are scheduled to ship before the year's end.

For computers that don't have a bidirectional parallel port, the software uses the status lines of the port for input.

Price: $695.
Contact: Xircom, 22231 Mulholland Hwy., Suite 114, Woodland Hills, CA 91364, (818) 884-8755,
Inquiry 1146.

Pocket Ethernet Adapter runs Novell NetWare drivers.

Low-Cost Parallel Port Network

If you want inexpensive file transfer for your small office but all your serial ports are packed full of peripherals, you might try installing the 3X-Link16 network through your parallel ports.

But don't worry about tying up the parallel ports. You can plug your printer into the parallel port on the back of each 3X-Link16 transceiver.

Features include background file transfer and E-mail, the company says. All you need is a pair of 3X-Link16 transceivers and some twisted-pair cabling, which is included in the basic package. You upgrade the network with additional adapters.

The network, which connects up to 16 PCs, has a maximum distance of only 400 feet. Data rate is 500,000 bps.

Security features include multilevel passwords.

Price: Basic package, $239; additional adapters, $139; printing software, $149.
Contact: 3X USA, One Executive Dr., Fort Lee, NJ 07024, (201) 592-6874.
Inquiry 1145.

Networking at 200 Megabits Per Second

The Baytec 2000 is a SCSI-based computer network its developer claims can flash data from port to port at rates up to 200 megabits per second.

This advantage is due to its SCSI connections, Baytec says, which transfer data in 64K-byte packets. At the cabling level, Baytec uses AMD's 125 taxi chip set, which supports coaxial, twisted-pair, and optical-fiber cabling.

The idea behind the Baytec 2000 network is simple. Instead of a complex array of network hardware and software, each computer or workstation on a Baytec network is outfitted with a SCSI port, complete with device driver. The nodes are daisy-chained, seven at a time, and plugged into a cable interface; up to eight interfaces can connect to each server, for a total of 56 users per server; and multiple servers can be linked together.

Installation is a matter of installing the appropriate SCSI interface, attaching a node controller, and adding the driver to the computer's operating system (an MS-DOS .SYS file, a Mac resource in the System file, or a workstation's Unix driver).

Within each base server is a 65816—the same processor that's in the Apple IIGS, and the 16-bit successor to the venerable 6502 that has powered Apple IIs for more than a decade.

Price: Base server unit, $17,000; each node interface, $500.
Contact: Baytec Inc., 32425 Schoolcraft Rd., Livonia, MI 48150; (313) 427-1250.
Inquiry 1151.

Network Your MapInfo

MapInfo 4.0 is a networking upgrade to the popular single-user MS-DOS mapping software. The latest version, which requires an AT and DOS 2.0 or higher, lets you distribute mapping work through Novell NetWare. (Upgrades are planned for other network operating systems.)

You can either buy maps from Mapping Information Services or make your own to work with the software.

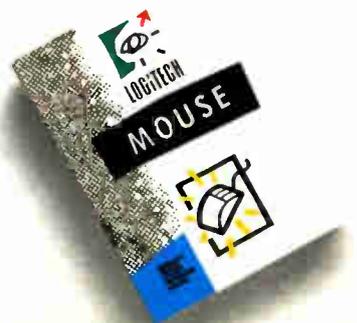Included in the base price of version 4.0, for example, is a database of the five-digit ZIP codes and a map of the U.S. that can be viewed as a whole or in regions. You enter the ZIP code you need to identify and MapInfo points to the region on your view of the U.S. map.

In a networked configuration, everybody with a node version of MapInfo 4.0 can simultaneously access the same maps and the same databases. Anybody on the network can access maps or data from local drives and the main file server. File locking and edit transaction files protect the data; only one user at a time can make edits on a particular portion of the map.

But you can create several separate mapping layers and divide mapping work among several people.

Price: Server version, $750; node version, $595, or $1195 for three nodes; optional maps from $75 to $2000.
Contact: Mapping Information Systems Corp., 200 Broadway, Troy, NY 12180, (800) 327-8627; in New York, (518) 274-8673.
Inquiry 1150.

Everything OS/2 can do for you…

OS/2™ includes a built-in graphical interface so it's easy to use.

OS/2 provides an optional Communications Manager which allows easy networking.

OS/2 lets you run your DOS programs plus hundreds of programs DOS can't.

OS/2 lets you keep two or more programs running at the same time, so you can do more.

OS/2 lets you run programs larger than 640K, so you can use more powerful applications.

OS/2 lets you take advantage of 386™ power.

OS/2 lets you take full advantage of Micro Channel.™

OS/2 provides an optional Database Manager to make managing information easy.

00F8896

IBM Operating System/2™
Installation Diskette

Version 1.10 © Copyright International Business Machines Corp. 1981, 1988
Licensed Material - Program Property of IBM -
All Rights Reserved
© Copyright Microsoft Corp. 1981 - 1988
**Note to US Government Users — RESTRICTED
RIGHTS LEGEND — Use, duplication, or disclosure
is subject to restrictions in GSA ADP Schedule
Contract with IBM Corp.**

This offer lets you do for less.

Right now, when you choose OS/2, you can get from $100 to $1,600 back on the kind of heavy duty memory that only OS/2 can handle. With this offer, the more memory you buy (up to 8Mb), the bigger your rebate.

Plus you can get thousands of dollars in rebates on over 100 different OS/2 programs. You can also get hundreds of dollars back on modems, accessory cards and hardware—all the things that help you do more work in less time with OS/2.

So if you're ready to move up to all the real advantages of OS/2, ask your IBM Authorized Dealer about these rebates today. To find the dealer nearest you call 1 800 IBM-2468, ext 128.

IBM®

Clean Elegant Powerful…

New SQL Machine for LANs

The SQL Mach 1 is a dedicated database system that uses a Structured Query Language–based engine and achieves 15 to 60 times the performance of its PC-based brethren, its manufacturer claims.

One of the major keys to advanced performance is its client-server approach to database operation. The client sends a request for information; the Mach 1 performs the database operation by accessing its own disk drives and returns only the answer (unlike solutions that return the entire database of information, tying up the network).

Other key advantages include a proprietary API (application program interface), so it doesn't get bogged-down on PC operating system code, and a patent-pending relational coprocessor. It's based on an 80286 backplane and is comparable in price to fully configured 80386-based PC systems.

The caching controller, described as a discrete design with high-speed memory management, effectively speeds up some operations that once took 400 to 500 μs and performs them in 10 μs, the company claims.

The Mach 1 includes 4 megabytes of RAM, a 17-ms, 320-megabyte hard disk drive, 150 megabytes of tape backup, and a 1200-bps modem for remote support.

There are actually four I/O slots for four direct-channel cards to directly support up to 16 users at distances up to 200 feet. Or you can use one of the I/O slots for an Ethernet card to network the entire office with database capabilities slightly less sophisticated than if the users were connected directly to the I/O. There are also six SCSI ports.

Client-server approach speeds SQL database.

Price: $23,950.
Contact: Advanced Data Servers, P.O. Box 4937, Boise, ID 83711, (208) 322-7800.
Inquiry 1144.

PowerBridge Across Topologies

The low-cost PowerBridge software lets you bridge from any NetBIOS-compatible LAN that uses Server Message Block protocols to any other, the manufacturer claims.

One server module runs on the dedicated or nondedicated bridge server, generally at least a 286-based system with 640K bytes of RAM. Another module is used by any bridge participant, a similarly configured machine.

You can share disks, printers, and gateway services with any network that's connected, Performance Technology says. Connections can pass through up to four bridge servers to join a total of five networks, whether they be Token Ring, Ethernet (server-based or distributed), or ARCnet.

Phone the Office for E-Mail

VoxMail is a hardware and software system that links you and your Touch-Tone telephone to the E-mail system back at the office.

Receiving messages is the easy part. You log on with a Touch-Tone access code. VoxMail then converts your text-based E-mail messages into speech. To reply, you press keys that generate pre-assigned generic responses like "No. Wait until we talk." The reply is then automatically mailed with a copy of the original message.

There are limitations. It supports only nine reply messages, and it works only with Message Handling Service–compatible E-mail systems, a de facto Novell standard, VoxLink says. You also need a dedicated XT and two free slots for the phone interface board and the text-to-speech board. The phone interface board handles text to ASCII via phonetic algorithms.

Each complete system supports five MHS applications, nine reply messages, an adjustable security code, attachment files, and administrative log reports.
Price: $3995.
Contact: VoxLink Corp., 432 Coventry Dr., Nashville, TN 37211, (615) 331-0275.
Inquiry 1148.

Price: $495.
Contact: Performance Technology, 800 Lincoln Center, San Antonio, TX 78230, (512) 349-2000.
Inquiry 1149.

Novell Introduces NetWare 386

NetWare 386 version 3.0 is Novell's first network operating system that's optimized for use on the 32-bit 80386 architecture.

Unlike NetWare 286, which could support only 100 users, NetWare 386 can support up to 250 users on one server. In addition, Novell says NetWare 386 features a simplified and less time-consuming installation procedure, enhanced printer resources and file security features, and a technique called dynamic resource configuration, which automatically manages memory allocation for caches and buffers, a task formerly managed manually.

With this introduction, NetWare runs on virtually all major operating systems and hardware architectures, including MS-DOS and OS/2, Macintosh, and Unix systems such as Sun and NeXT.

Novell says NetWare 386 will support the major client/server file protocols, including AppleTalk Filing Protocol (AFP), the Unix-based Network File System (NFS) from Sun Microsystems, and IBM's Server Message Block (SMB) and OS/2 file protocols, as well as Novell's own NetWare Core Protocols, which support MS-DOS and a variety of other file types.
Price: $7995.
Contact: Novell, Inc., 122 East 1700 South St., Provo, UT 84606, (800) 453-1267.
Inquiry 1147.

dBASE IV. The experts can't say enough about it.

SQL not Dbase forte, test finds

Ashton-Tate promises to correct query flaws after competitor criticism

BY DOUGLAS BARNEY

TORRANCE, Calif – All along, critics argued that Ashton-Tate Corp. would not be able to produce an effective implementation of SQL in its Dbase IV product, charges the company vehemently denied. Now that Dbase IV is out, it seems the critics where right after all, at least regarding Ashton-Tate's first crack at SQL.

According to Fred Luk, president of Quadbase Systems, Inc., who has tested Dbase IV, users may get incorrect results from even simple SQL queries. Luk tested Dbase IV against his own product, Dquery, which provides Dbase III Plus with SQL capabilities. Ashton-Tate has not denied his findings and pledges to correct Dbase IV flaws.

With some queries, Dbase IV found an incorrect result. With others, it failed to run because of internal errors. At other times, Dbase IV's SQL gave inconsistent results depending on the method of querying.

Fortunately, an Ashton-Tate spokesperson said, some fixes should already be available through the CompuServe on-line service. Ashton-Tate officials said the SQL problems are very data dependent, only cropping up when the product is used in unusual and specific ways. The biggest problem for Ashton-Tate was its lack of SQL experience. Unlike many SQL developers with years of experience, Ashton-Tate had to

Data View
PC packages bigger in big business

Users are running more micro software titles than ever, especially at large firms.

AVERAGE NUMBER OF TITLES USED PER PERSON ON A REGULAR BASIS

Computerworld, January 9, 1989

Ashton-Tate To Address dBASE Quirks

By Beth Freedman

Ashton-Tate is planning to air on major electronic bulletin boards the first batch of user-reported quirks in dBASE IV and suggested solutions to the problems, according to company officials.

The first peculiarities to be addressed include a compatibility problem using dBASE IV's new memo field feature on dBASE III Plus applications; a glitch that freezes the screen if the Ctrl-Break key sequence is activated; and other assorted bottlenecks involving network support and installation, according to Dave Micek, dBASE IV product manager for the Torrance, Calif., company.

The anomalies will be made public in the next few weeks on CompuServe, The Source and Ashton-Tate's own bulletin board as they are reported by users, Micek said.

"As is our long-standing practice, we are readying the first set of comments on dBASE IV to post on public bulletin boards," said Lydia Dobyns, vice president of marketing at Ashton-Tate. "However, very few areas are covered. We have heard nothing

A month after the release of dBASE IV, user reports of peculiarities and features that don't work as promised have begun to filter in.

to believe that dBASE IV is any thing less than stable."

Nevertheless, a month after its release, user reports of peculiarities in dBASE IV and features that don't quite work as promised have begun to filter in.

Ashton-Tate has acknowledged the existence of a memo-field problem, which occurs when dBASE IV users open or edit a memo field in a dBASE III

Plus application that has been moved into dBASE IV. Once dBASE IV is used on the memo field, that field and its entire record can no longer be accessed from any program, including work-alikes and dBASE add-ons, except dBASE IV.

The lockout is sparking user complaints that possible incompatibilities exist between the two versions.

"The way that dBASE IV handles enhanced memo fields, it's not possible to have compatibil-

ity with dBASE III Plus," Micek admitted.

To work around the problem, he suggested, users should make a copy of the database using the DBMEMO 3 option under the "Copy 2" command, adding that this solution has already been documented in dBASE IV's technical manual.

Ashton-Tate also confirmed the existence of a quirk that

freezes the screen when Ctrl-Break is activated; users must reboot their machines, Micek said, and no data is destroyed in the process. The firm is evaluating the problem, he said.

Users have also reported two other potential problems: irregularities in generating calculated fields and a flaw in the mailing-list function that prints a blank line in an address label when one field of information—for example, a company's name—is missing ∎

PC Week, December 5, 1988

PS/2 Model 50Z proves too hot for Dbase IV to handle

BY DOUGLAS BARNEY

TORRANCE, Calif — If you want IBM's sizzling Personal System/2 Model 50Z and Ashton-Tate Corp.'s hottest version of Dbase, both are available for a price. Just don't try to run them together.

The latest PS/2 Model 50Z is a speedier version of the machine that many users had labeled a dog. It increases the speed of the zero-wait state operations that wreaks havoc with Dbase IV, Ashton-Tate said. "We have heard of other problems regarding the 50Z," said spokesman Knox Richardson.

Bill Denton, a software consultant routinely automating a cosmetics firm, had the problem on four separate PS/2 Model 50Zs running both PC DOS 3.3

and PC-DOS 4.0. In each case, his hard disk scrambled, resulting in disconnected sectors and the possible loss of data.

So what is Denton to do? Simple. "I am getting Dbase IV the hell off my machines. I have some very expensive data I don't like having mucked around with," Denton said. He also said that he plans to reformat the hard disks and go back to Dbase III Plus.

According to Richardson, Dbase IV was tested prior to the availability of the Model 50Z. Dbase IV did not ship until October last year, some five months after the availability of the 50Z.

There is no fix coming for 50Z users from Ashton-Tate. We don't look at it as being an Ashton-Tate problem." Richardson said, adding that the product is not certified to run on the 50Z.

Computerworld, April 24, 1989

■ TECHNOLOGY UPDATE

Users Should Expect a Rocky Marriage of Dbase and SQL

BY SCOTT MACE

InfoWorld, May 8, 1989

Oracle developed the first commercial SQL database over 10 years ago. And the first SQL database for the PC over 4 years ago. It's called Professional ORACLE.

It has the most up-to-date, most powerful and most complete set of application development tools available.

Like SQL*Forms.® SQL*ReportWriter.™ SQL*Menu.® And SQL*Plus.® It's based on ANSI standard SQL, and runs on PCs, minis and mainframes. And it works.

To order Professional ORACLE for the PC, call 1-800-ORACLE1, ext. 4960. It's $1,299, and comes with a 30-day, money-back guarantee.

ORACLE®
Compatibility, Portability, Connectability.

Universally acknowledged to work just fine.

World Radio History

A Programming Tool for OS/2

If you're struggling with the intricacies of programming for OS/2, Hamilton Laboratories has a product that brings a familiar programming environment to IBM's latest PC operating system. As its name implies, Hamilton C Shell recreates the standard C shell language as described in the Berkeley 4.3 Unix Programmer's Manual. The company claims that all 42,000 lines of code in the product were written specifically for OS/2.

Hamilton Labs says its shell is a superior alternative to the standard OS/2 command processor, letting you program for the OS/2 environment more quickly and easily by manipulating files, processes, threads, and object connections.

The Hamilton C Shell includes fully nestable programming constructs for iteration and condition testing, variable arrays, and a wide range of expression operators and built-in functions. There are also advanced features for I/O redirection, piping, background execution, and parallel threading.
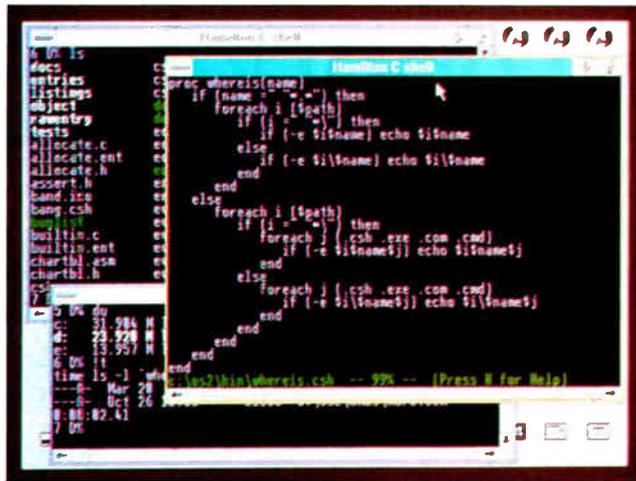
Rounding out the program's features are alias and shell procedures for defining your own language extensions, as well as command substitutions and advanced wildcarding.

The Hamilton C Shell runs on any OS/2-equipped system with at least 2 megabytes of RAM.
Price: $350.
Contact: Hamilton Laboratories, 13 Old Farm Rd., Wayland, MA 01778, (508) 358-5715.
Inquiry 1105.

An example of a procedure within Hamilton C Shell with directory windows in the background.

Modula-2 for the Amiga

M2Sprint 1.1, a Modula-2 development system for the Commodore Amiga, includes a compiler that can handle 45,000 lines per minute, the company reports. The compiler runs from the editor, the command line, Workbench, or ARexx, and the editor supports multiple windows, letting you compile in one window while you edit in the others.

The program also includes a single-pass Modula-2 compiler and program linker, program profiler, symbolic debugger, an Amiga ROM interface library, Modula-2 library, an Amiga interface library, a C-style I/O library, and IFF and AmigaDOS Replacement Project libraries.

Features of the compiler include internal files configurable for efficient RAM management, REAL and LONG-REAL support via the Amiga's library code (allows you to use hardware floating-point processors), and termi-

nation procedures for each module. You can also use it to generate in-line calls to the Amiga's operating system, eliminating the need for "stub" routine libraries, and to generate debug information for symbolic debuggers.

M2Sprint's editor has an automatic case-correction feature, which converts Modula-2 keywords to their correct case (e.g., procedure becomes PROCEDURE), and word completion, which automatically completes long names that you specify from a dictionary when you type enough characters (e.g., imple becomes IMPLEMENTATION). You can also develop and test programs without leaving the editor. The debugger shows the code at the point of error, as well as variable contents.

M2Sprint works on all Amigas with at least 512K bytes of RAM, KickStart 1.2, and Workbench 1.3 or higher.
Price: $385.
Contact: M2S, Inc., P.O. Box 550279, Dallas, TX 75335, (214) 340-5256.
Inquiry 1102.

An OS/2 Pascal Compiler with DOS Compatibility

If you're a developer who wants to use Pascal to develop OS/2 applications but still wants to keep DOS users unalienated, Prospero's Pascal for OS/2 will allow you to keep your feet firmly planted in both worlds. Pascal for OS/2 is a one-pass compiler that's optimized for OS/2-specific applications. But it also includes a DOS linker and library that produces DOS programs (as long as your code doesn't use OS/2-specific features).

The package includes a threading function for OS/2, letting you run Pascal procedures in parallel with the main program. You can also call OS/2 functions from Pascal by simply declaring them external. In addition, Prospero has added a new predeclared data type called ASCIIZ that allows the declaration of null-terminated dynamic-length strings in OS/2. You can also produce code to take advantage of the extra instructions available on the 80286 processor.

Total code is limited only by the size of your hard disk. And although you can generate a maximum of 64K bytes in a single compilation, any number can be linked into a program. There is a limit of

64K bytes on the outer-level static and common data, and the heap can expand to 4 megabytes.

Pascal for OS/2 includes a workbench/editor that lets you choose compilation and linking operations from a menu. Also included is the Probe source-level debugger with data breakpoint and multithreading capabilities. The whole package runs on any OS/2-equipped system.
Price: $390.
Contact: Prospero Software, Inc., 100 Commercial St., Portland, ME 04101, (207) 874-0382.
Inquiry 1101.

DESQview 2.2 and DESQview 386. The multitasking, windowing environments that work with your favorite software.

DESQview™ is the operating environment that brings OS/2™ power to DOS. And it lets you, with your trusty 8088, 8086, 80286, or 80386 PC, leap into the next generation in PC productivity. For not much money. And without throwing away your favorite software.

Introducing DESQview 2.2

And now, DESQview 2.2 adds capabilities, performance, and compatibility enhancements you've been asking for.

Like being able to fine tune DESQview performance "on the fly." Run Lotus Express and Metro. And the Intel Connection Co Processor. Even use the DOS 4.0 shell with DESQview. Have DESQview automatically install Quattro, Sprint, Aldus PageMaker, Microsoft Excel, Word Perfect, Dataease and as many as 80 other programs. And using the DESQview API, be able to dynamically link them.

More bang; less bytes

While other programs get bigger, we've worked to make DESQview smaller. And we've succeeded in a big way on PCs and PS/2™s with extended, EMS 3.2 (AboveBoard), EEMS and EMS 4.0 memory—as well as on 386 PCs and

DESQview lets you run your favorite programs in windows side-by-side.

PS/2s. For example, DESQview overhead on EMS 4.0 and 386 PCs can be as low as 10K on EGA/VGA PCs. And DESQview actually *increases* memory 30K on CGA PCs; 20K on monochrome and Hercules PCs. That's good news for users of big desktop publishing, CAD and database programs.

Introducing DESQview 386

For users of 80386 PCs and PS/2s (or PCs with 80386 add-in boards, such as the Intel Inboard 386), there's DESQview 386 (a combination of DESQview 2.2 and the new QEMM-386 Quarterdeck Expanded Memory Manager, version 4.2).

DESQview 386 gives you extraordinary power. Run text, CGA, EGA, VGA, and Hercules programs in windows and in the

background. Run 32-bit 386 programs, like Paradox 386, and IBM Interleaf simultaneously with your favorite DOS programs. All with the speed and performance you expect out of your 386. And with protection against 'misbehaved' programs.

Promise and performance

And, of course, both DESQviews have all the features that made prior versions the popular choice in operating environments. The ability to multitask in 640K and beyond. View programs in windows or full screen. Transfer data. Access DOS via menus. Dial your phone. And create keystroke macros within and between programs.

Our story gets better and better

If there's any doubt about our commitment to your PC and PS/2 productivity, just look at our accomplishments over the years. We think you will understand why GE, Ford, Aetna, Monsanto, and so many other major corporations use DESQview.

And why PC Magazine twice gave DESQview its Editor's Choice Award for "The Best Alternative to OS/2," why readers of InfoWorld voted DESQview "Product of the Year" three times. Why, by popular vote at Comdex Fall for two years in a row, DESQview was chosen "Best PC Environment" in PC Tech Journal's Systems Builder Contest, and just won their "Professional Solutions" Award.

DESQview lets you have it all now.

ek Delivers.

QEMM. Break the 640K barrier for $59.95

DESQview API Toolkit. New C and Pascal Libraries, Debugger. Panel Designer. And more.

DESQview API Conference '89 Coming in August Write for details

Your 80386 PC, IBM Personal System/2 Model 80, PC or AT with 80386 add-in board, as well as your IBM Personal System/2 Models 50 or 60 can all break through the DOS 640K barrier. Now you can have maximum use of your memory—whether you have one megabyte or 32—with the Quarterdeck Expanded Memory Manager. All without having to purchase special expanded memory boards.

QEMM uses hidden features within your existing memory to make it compatible with the Lotus-Intel-Microsoft Expanded Memory Specification (EMS) version 4.0.

Now you can run colossal spreadsheets, databases, and CAD models designed for expanded memory, using Lotus 1-2-3, Symphony, Framework, Paradox, AutoCAD, Excel and more.

And if you'd like to use these programs all together —multitasking beyond 640K— QEMM works with our popular DESQview multitasking environment.

If you are one of the 12 million or so 8088, 8086 or 80286 PC users who feel left out, don't despair. We have options that let you keep your computer and favorite programs and give you today what the newest PCs and operating systems are promising for the future.

Visit your dealer for more information on barrier-breaking Quarterdeck products.

API Reference Manual

The key to the power of the DESQview API, our Reference Manual contains all you need to know to write Assembly Language programs that take full advantage of DESQview's capabilities. And there's an 'include' file with symbols and macros to aid you in development.

API C Library

Here are C language interfaces for the entire set of API functions. It supports the Lattice™ C, Metaware™ C, Microsoft® C, and Turbo C compilers for all memory models. Included with the C Library package is the API Reference Manual and source code for the library.

API Pascal Library NEW!

The Pascal library provides interfaces for the entire set of API functions. It supports Turbo Pascal V4.0 and V5.0 compilers. Included are the API Reference Manual, source code for the library, and example programs.

API Debugger

The DESQview API Debugger is an interactive tool enabling the API programmer to trace and single step through API calls from several concurrently running DESQview-specific programs. Trace information is reported sym-

bolically along with the program counter, registers, and stack at the time of the call. Trace conditions can be specified so that only calls of interest are reported.

API Panel Designer

This interactive tool helps you design windows, menus, help screens, error messages, and forms. It includes an editor that lets you construct an image of your panel using simple commands to enter, edit, copy, and move text, as well as draw lines and boxes. You can then define the characteristics of the window that will contain the panel, such as its position, size, and title. Finally, you can specify the locations and types of fields in the panel.

The Panel Designer automatically generates all the DESQview API data streams necessary to display and take input from your panel. These data streams may be grouped into panel libraries and stored on disk or as part of your program.

More Tools are Coming

Quarterdeck is committed to adding tools as needed by our users. To that end we have been working with Ashton Tate and Buzzwords International on dBASE III and dBASEIV translators. And in the works, we have BASIC and DOS Extender libraries.

Quarterdeck

Quarterdeck Office Systems, 150 Pico Blvd., Santa Monica, CA 90405 (213) 392-9851
FAX: (213) 399-3802

For additional information, please use the following Reader Service numbers: DESQview: #207 QEMM: #208 API Tools: #209 API Conference: #210

AUGUST 1989 • BYTE 67

Integrate Text and Graphics in Windows

Precision Software's Superbase 2 Windows is a DBMS that lets you tag TIFF, PCX, and IMG images to a record. It includes an editor, mail merge, label printing, and communications capabilities with its data management features.

Superbase 2 Windows features a VCR-like control panel on the bottom of the screen that allows you to quickly browse forward or backward and pause among up to 999 index sequences. You can also use it to select a subset of records from within a field and access files by index category.

Other features include validation, multiple response, time, calculated and virtual fields, date parsing, and cross-file lookup capability. You can import and export data from Excel, Lotus 1-2-3 versions 2.1 and 2.2, dBASE II and III, and ASCII.

Superbase 2 Windows includes a run-time version of Windows 2.03.
Price: $295.
Contact: Precision Software, 8404 Sterling St., Suite A, Irving, TX 75063, (214) 929-4888.
Inquiry 1113.

Superbase 2 Windows is a database manager with a VCR-like control panel that allows quick browsing.

Hold the Phone

The telephone works as a valuable tool for some, but for others it's just plain annoying. Varteck's Influence is a phone dialer and database that may alleviate at least some telephone tedium.

Influence stores over 10,000 names with addresses, phone numbers, and descriptions. You can access that information by category, keyword, or name. And the program acts as a dialer and a follow-up file.

When you receive a call, you enter the first two letters of the caller's name, and the program shows you all the contacts with that last name. You can flip through the information while you're on the phone, and you can add to it with follow-up information. The program runs on the IBM PC with 385K bytes of RAM and a hard disk drive.

Influence does the dialing for you.

Price: $98.
Contact: Varteck, 3 Regent St., Suite 304, Livingston, NJ 07039, (201) 740-1750.
Inquiry 1116.

On the Road Again

Keeping track of business expenses you incur while on the road can be inconvenient at best, but WorkSmart Technologies has a solution. ExpenseSmart lets you fill out your expense reports while you're on the fly.

Designed for laptops, the program keeps your keystrokes to a minimum, according to WorkSmart. You can customize the program with whatever expense categories you need, reimbursement levels, and method of payment.

This menu-driven program works with DOS-based systems that have at least 512K bytes of RAM, and it comes in both 5¼- and 3½-inch formats.
Price: $79.95.
Contact: WorkSmart Technologies, 5700 Hillcrest Dr., Suite PL, Lisle, IL 60532, (312) 963-2935.
Inquiry 1117.

New Excel to Break 1-megabyte Barrier, Support BIFF

The new version of Excel overcomes the 1-megabyte limit of earlier versions with its ability to address a full 8 megabytes of Macintosh RAM, Microsoft reports.

Excel 2.2 uses the Binary Interchange File Format, also used by Windows. With BIFF, you can transfer and use spreadsheets, macros, and charts between platforms without having to convert them.

The program supports the sparse-matrix method of memory management, which allocates memory to cells only where you've entered data, increasing the efficiency of memory use, Microsoft reports. You can now use up to 256 fonts in a single spreadsheet and adjust row heights to accommodate larger font sizes or highlight particular entries.

Other improvements of the program include the ability to use cell notes to specify assumptions on a cell-by-cell basis, and the use of precedents and dependents for checking proper derivation of cell values. You can also search and replace a particular entry.

Microsoft has added 200 macro functions and a macro library for common operations such as consolidation and cross-tabulation.

Excel 2.2 runs on the Mac Plus or higher with System 6.0.2 or higher. HyperCard 1.2 is required for a training module with lessons on the basics, worksheets, charting, and databases.
Price: $395.
Contact: Microsoft Corp., 16011 Northeast 36th Way, P.O. Box 97017, Redmond, WA 98073, (206) 882-8080.
Inquiry 1114.

GO AHEAD, MAKE YOUR DAY

Point-and-Pick Your Applications

Choose the application you need quickly and easily directly from the menu — or even another application!

Run Several Tasks at Once

Switch instantly between active tasks in different applications at the touch of a keystroke!

Move Data Between Applications

Share data easily between applications like SCO™ Lyrix,® SCO Professional,® and SCO Integra™ with the electronic Clipboard!

Print While You Work

Move on to your next job while your last one is printing out — on a local or shared workgroup printer!

Link Up Your Business with Electronic Mail

Exchange messages and files — even spreadsheets and graphics — across the office or around the world!

Talk Across The Office

Instantly converse with other system users, screen-to-screen, with the handy Intercom!

Locate Business Contacts Instantly

Store, update, find, and sort addresses and phone numbers quickly and easily with the time-saving Directory!

Schedule Meetings and Resources

Check others' Calendars online for available times — then schedule and notify them automatically!

Calculate Within Any Application

Put the four-function, "running-tape" capabilities of an online Calculator right at your fingertips!

Add Only the Applications You Need

Build your own customized solution by adding individual applications as you need them!

WITH THE SCO PORTFOLIO WORKGROUP® SOLUTION

Get the competitive edge with the SCO Portfolio™ integrated workgroup solution!

Teamed with the world's most popular UNIX® System — SCO System V — the SCO Portfolio solution turns the 386™ personal computer into a workgroup powerhouse.

What's more, users only need to know how to use their familiar applications in order to put the amazing power of the UNIX System to work immediately.

With SCO Portfolio and the SCO Portfolio family of business applications, everyone in a workgroup can perform virtually any business task — from writing reports and creating financial analyses, to scheduling meetings and exchanging messages — far more productively than ever.

And all using a single, standard — and cost-effective — 386-based PC!

Get started today with SCO Portfolio Suite, and get all the advantages of a fully-integrated office system without compromising the functionality of full-featured business applications — all in one economical package.

SCO Portfolio Suite integrates the powerful SCO Lyrix word processing system, the SCO Professional 1-2-3® workalike, and the SCO Integra industry-standard-SQL database, with SCO Portfolio's convenient desktop tools, customizable menu system, and electronic clipboard — and lets you add any other software of your choice under its easy-to-use menu, as well.

Contact your SCO authorized supplier or call (800) 626-UNIX (626-8649) for more information about SCO Portfolio and SCO Portfolio Suite and find out how easy it is to make your day — today!

SCO
THE SANTA CRUZ OPERATION

(800) 626-UNIX (626-8649)
(408) 425-7222
FAX: (408) 458-4227
E-MAIL: ...!uunet!sco!info info@sco.COM

SCIENCE AND ENGINEERING

Moving Numbers on a Mac

Mathematicians' software needs vary from simple equation processing to sophisticated plotting and modeling. Two recently released programs for the Macintosh take very different approaches to the task of number manipulation.

Formulator lets you merge text with your numbers. It is an equation processor that offers a WYSIWYG display and has a built-in text editor. Mathematical typesetting features let you italicize variables, change type size, insert space between operations, and alter the position of delimiters.

You can insert, delete, and copy anything from a symbol to a whole formula—within a document, or from one document to another, according to ICOM Simulations.

The program includes the Magnifying Glass icon, which doubles the size of a document for easy editing of small characters; the Greek icon, which opens a palette containing the Greek character set; and Left, Center, and Right Justify icons, which let you choose how to justify lines of text, formulas, elements in a formula, and columns in matrices. The program also contains a full library of symbols.

Formulator outputs in $\TeX$. The program runs on the Mac Plus, SE, and II.
Price: $149.95.
Contact: ICOM Simulations, Inc., 648 South Wheeling Rd., Wheeling, IL 60090, (312) 520-4440.
Inquiry 1108.

An all-new version of the equation-solving program TK Solver Plus combines equation solving with knowledge management.

You can use the program as a basic equation solver and scientific calculator to solve sets of simultaneous linear or nonlinear equations. You enter equations as you see them.

TK Solver Plus is a rule-based declarative language that lets you solve problems using an object-oriented method. Interactive tables supply a spreadsheet-like format for input and output of user-defined functions.

You can produce high-resolution line, bar, and pie charts as well as tables of data using TK Solver Plus. You can also plot multiple curves in the same graph and any number of graphs in the same model. A whole model or any part can be saved or added to other models using a cut-and-paste approach. You can transfer data between TK Solver Plus and other programs via files in WKS, WK1, DIF, or ASCII format.

Universal Technical Systems, designer of TK Solver Plus, reports that all versions of TK Solver are compatible, so you can port data from Mac to DOS environments.

TK Solver Plus runs on any Mac from the 512KE up.
Price: $395.
Contact: Universal Technical Systems, Inc., 1220 Rock St., Rockford, IL 61101, (800) 435-7887 or (815) 963-2220.
Inquiry 1109.

Cam Design

CamDes assists you in designing and analyzing cams and cam-driven mechanisms.

To use the program, you describe the motion requirements of the cam, followed by selecting from known kinematic profiles. Information is output to screen, printer, or disk and is calculated in tabu-

lar and graphical form.

You can access graphical diagrams in the program and use them to illustrate the dynamic and kinematic profiles over the full 360 degrees of cam motion, according to MicroAnalysis.

CamDes is error-trapped and includes flagging options that point out possible deficiencies in the cam design. You can store and retrieve your designs for later use.

Ten standard kinematic motions are supported, and the program handles plate, barrel, and linear cams.

To run CamDes, you need an IBM PC with 384K bytes of RAM, DOS 2.0 or higher, and a CGA or EGA card.
Price: $89.
Contact: MicroAnalysis Software, 26148 Tallwood Dr., North Olmsted, OH 44070, (216) 779-9523.
Inquiry 1111.

Low-Cost High-Resolution Scientific Graphics

Plotting technical or scientific graphs from data entered into text windows is what Edtech does best. The program reads and writes from or to data in its own database files, WKS files, or ASCII text files.

Edtech's graphics screen editor lets you size and position graphs, labels, and diagrams at arbitrary positions on the page.

The program supports Epson LQ or Toshiba P321 printers and provides hard copy at 180 by 180 dpi. A page is 1440 dots horizontally by 1800 dots vertically.

Greek and mathematical symbols are also available.

The program runs on the IBM PC with 640K bytes of RAM. You also need a 24-pin printer and a CGA, EGA, or Hercules graphics adapter. A math coprocessor is recommended.
Price: $65.
Contact: Digital Analytics, P.O. Box 31430, Houston, TX 77231, (713) 721-2069.
Inquiry 1110.

A Partner in the Lab

LabSolutions automates multicomponent solution and buffer preparation calculations.

For chemical reagents, you specify the concentrations desired, and the program calculates the amounts of components needed. For buffers, you specify the pH, and the program calculates the amounts of each buffer form or the amounts of common acids or bases needed.

The program accepts all standard units of weight, volume, and concentration. Molecular weights are calculated from atomic formulas or chemical names.

Routines are also included for mixing and diluting solutions.

A database of chemicals, buffers, and acids is included, which you can modify with your own entries. You can also store your solution recipes on disk or print hard copies.

LabSolutions runs on the IBM PC or PS/2s with 256K bytes of RAM.
Price: $99.
Contact: The Center for Science Support, Inc., 54 Brattle St., Arlington, MA 02174, (617) 646-1466.
Inquiry 1112.

How Telebit modems can give you a graphic description in no time.

When you need error-free graphics in seconds, you need a Telebit® high-speed, dial-up modem. Because only our modems can meet the critical demands of graphics.

Like transferring medical images for diagnoses — where a lost pixel can be life-threatening, and where nearly perfect isn't nearly good enough. Which is why a major East Coast hospital chose Telebit. Now they're sending 1 million bytes of critical data from a graphics-plot screen in seconds, error-free.

That's also why a national weather service picked Telebit to collect real-time weather information at U.S. radar sites. Only Telebit modems made the connection every time.

Imagine what Telebit modems will do for you.
On all applications, from CAD/CAE/CAM and electronic publishing to point of purchase, remote diagnostics, and more.

All with our family of high-speed modems. From 9600 bps to 19,200 bps, including V.32.

To get the graphic details and a free application brochure, call 1-800-TELEBIT or 415/969-3800.

Or write Telebit, 1345 Shorebird Way, Mountain View, CA 94043.

Because no one gets the message through like Telebit.

TELEBIT®
When connectivity counts.

World Radio History

CAD AND GRAPHICS

Image of an office floor plan created with MicroStation Mac.

MicroStation Mac

Intergraph's CAD software is now available in a Macintosh version. The two- and three-dimensional design software is compatible with other MicroStation programs, so you can share files without translation, according to Intergraph.

The Mac version features resizable windows, dynamic tool palettes, dialog boxes, and selection sets. You also have a choice between the Mac or the IBM PC interface.

MicroStation Mac supports up to eight separate views of a design for viewing different perspectives and scales. All views are active at the same time and can be placed on up to six monitors.

Input is by mouse, tool palettes, pull-down menus, tablet command menus, or key-ins in the command window. The program imports and exports text and PICT-format data types.

MicroStation Mac runs on the Mac SE/30, II, or IIx with at least 2 megabytes of RAM and a 40-megabyte hard disk drive. You also need System 6.0.2 and Finder 6.1.
Price: $3300.
Contact: Intergraph Corp., One Madison Industrial Park, Huntsville, AL 35807, (800) 345-4856; in Alabama, (800) 345-0218.
Inquiry 1121.

$89 Graphics Software Toolkit

The 3-D Computerscape toolkit helps you create vivid displays and artistic images, according to the program's developers.

The package includes demonstration files and programs with examples of applications in robotics, animation, and solids modeling. You can create solid objects, edit them, and display them. Display options include perspective, animation, three-dimensional rotations, and multiple views.

You can incorporate the program's functions and procedures into Turbo Pascal programs or use the 3-D Drawing Board, a three-dimensional drafting system.

To run 3-D Computerscape, you need an IBM PC with 512K bytes of RAM and an EGA card.
Price: $89.
Contact: Abbot, Foster, & Hauserman Co., 44 Montgomery St., Fifth Floor, San Francisco, CA 94104, (800) 562-0025 or (415) 955-2711.
Inquiry 1125.

ScanMan.™ Turning imaginations loose everywhere.

*P*op any image up to 4" x 11" straight into any IBM personal computer or Apple Mac™!

*S*elect one, two, three or four hundred D.P.I. Resize, rotate, flip and edit it. With the IBM version, use PaintShow Plus™ (included free) for coloring and shading, then port into any popular publishing program.

*W*ith the Mac ScanMan use the Clipboard™ to transfer the image to virtually any application —the Mac ScanMan works just like any Desk Accessory!

ScanMan for the PC — $309
Multi-Channel version — $369
Macintosh version — $499

*F*or your nearest dealer, call:
800-231-7717

In California:
800-552-8885

In Europe:
+ +41-21-869-96-56
Circle 152 on Reader Service Card
(DEALERS: 153)

CAD Overlay lets you import scanned images to VersaCAD. The yellow lines are the scanned raster image, and the other colors are VersaCAD entities drawn over the scanned image.

Versatile VersaCAD Tool

If you've ever had to convert a digitized drawing into a VersaCAD format, you know it is a lengthy process. CAD Overlay saves time and trouble by capturing a scanned image of a paper drawing and importing it quickly into VersaCAD.

You begin by scanning an existing paper drawing. CAD Overlay displays the scanned image in the background of the screen, creating a hybrid image. You can turn off that background image, or move, zoom, or pan it. You can trace over the image with the VersaCAD drawing on the same screen.
Price: $1000.
Contact: Image Systems Technology, Inc., 120 De-Freest Dr., Rensselaer Technology Park, Troy, NY 12180, (518) 283-8783.
Inquiry 1122.

Scorpion's Raster-to-Vector Conversion

SRV is a batch raster-to-vector conversion program that takes computer files of drawings you've scanned and converts them to vector images, which can be manipulated with a CAD system. SRV's maker claims that over 90 percent of existing drawings are prime candidates for the conversion process.

The SRV system converts images in the background using Scorpion's Motorola 68030-based coprocessor board. During the vectorization process, the software enhances the image by deleting isolated pixels, filtering out extraneous points on the line work, closing gaps in line work, and connecting line segments. It also normalizes line width across a line string and recognizes text, Scorpion reports.

The program retains the raster image on a separate layer so you can view the vector output and the raster data together.

SRV runs on an 80286- or 80386-based PC with Scorpion's coprocessor board.
Price: Software only, $6000; board and software, $12,000.
Contact: Scorpion Technologies, Inc., 101 Metro Dr., Seventh Floor, San Jose, CA 95110, (408) 452-0700.
Inquiry 1124.

MacWorld Returns to Boston

The three-headed variant of MacWorld returns to Boston this month on August 10–12. The show will again be held in three different locations: the Bayside Exposition Center, the World Trade Center, and the Wang Center.

The show will feature three special sessions in addition to the regular technical sessions: a connectivity (networking) miniconference; user roundtables, where you can brainstorm with others about education, Macs in corporations, or small-business computing; and MUSE, in which users groups from all over give out advice, strategies, and answers to questions.
Price: Exhibits only, $20; exhibits and conferences, $80.
Contact: Mitch Hall Associates, P.O. Box 155, Westwood, MA 02090, (617) 329-9911.
Inquiry 911.

Two IEEE International Conferences

Albany, New York, and Boston will be the sites of two separate international IEEE conferences.

The 1989 IEEE International Symposium on Intelligent Control will cover machine intelligence, architectures and tools, applications, and, of course, control. Sessions may include fuzzy control, qualitative reasoning, neural networks, and social and economics systems.

The symposium will be held September 24–26 at the Desmond Americana in Albany, New York.
Price: Before August 15: IEEE members, $175; nonmembers, $225 ($25 extra after August 15).
Contact: Program Chairman, ECSE Dept., Rensselaer Polytechnic Institute, Troy, NY 12180, (518) 374-1220.
Inquiry 914.

The 1989 International Conference on Computer Design will be held at the Hyatt Cambridge October 2–4. The conference will have four tracks: architecture, CAD, Design and Testing, and VLSI Technology.
Price: IEEE members, $210; nonmembers, $260; students, $65 ($50 extra after September 5, except for students).
Contact: ICCD 1989, 1730 Massachusetts Ave. NW, Washington, DC 20036, (202) 371-0101.
Inquiry 915.

We Want to Hear from You

Summer will soon be over, and BYTE wants to hear about your users group's plans for the fall. Please send your advance announcements on upcoming activities, speakers, and conferences to Regional Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.

World Radio History

X-BANDIT

Break the 640K DOS barrier and utilize the Advanced Features of the LIM 4.0 standard while using only one motherboard slot!

DESIGN PHILOSOPHY

• The Teletek X-Bandit was specifically designed to utilize the advanced features of the Lotus/Intel/Microsoft EMS 4.0 Specification. Further, the X-Bandit's Segmented Memory Mapping capability allows the user to extend DOS size beyond the 640K barrier. It is available in both 8 and 16 bit versions for use in the IBM XT, AT, and compatibles.

MEMORY

• Segmented Memory Mapping allows the user to fill out unused memory segments between 640K and 1024K. By "claiming" unused portions of memory in 16K increments, the user effectively increases TPA size. LAN or custom software modules, for example, can be loaded into these high memory areas thus relieving the lower 640K of TPA for other application programs.

• Split Memory Addressing allows the user to fill out conventional memory to 640K.

• Extended Memory Addressing is available for the PC/AT version.

• 2 MB capacity in a single slot. Up to 8 MB per system.

• Parity checking.

SOFTWARE

• Easy menu-driven auto configuration software.

• Device driver includes print spooler and RAM drive.

• Supports multitasking with the appropriate shell-resident software package.

SPEED

• 6/8/10 MHz speed with 0 wait states. 12 MHz speed with 1 wait state.

WARRANTY

• One year parts and labor.

TELETEK

4600 Pell Drive
Sacramento, CA 95838
(916) 920-4600
Fax (916) 927-7684

Circle 496 on Reader Service Card (DEALERS: 497)

World Radio History

Network Management Workshop

The IEEE Network Management and Control Workshop, which will focus on real-time network monitoring and control, will be held September 19–21 at the Marriott Hotel in Tarrytown, New York.

The workshop will have four tracks: user interfaces and network representations; integrated network management; expert-systems network management; and performance analysis and dynamic routing. The workshop will consist of two days of regular sessions and one day of tutorials on September 19.
Price: (Before August 15):

Conference: IEEE members, $300; nonmembers, $350; tutorials: $150 and $175.
Contact: Ted Lehrman, CATT Room 321, Polytechnic University, 333 Jay St., Brooklyn, NY 11201, (718) 260-3050.
Inquiry 907.

IBM PC Users Group in Southern Connecticut

The Trumbull PC Users Group (TPCUG) is centered just northwest of Bridgeport, Connecticut. The group's special-interest groups include desktop publishing, database, and beginners. The TPCUG has about 160 members.
Contact: Donald Potwora,

TPCUG, P.O. Box 545, Trumbull, CT 06611, (203) 334-4267.
Inquiry 913.

Hard-Copy Supplies Conference

A conference for hard-copy supply manufacturers and printer systems vendors will be held August 23–25 at the Boston Cambridge Marriott. Sessions will cover innovative printing solutions, electronic print shops, and others.
Price: $895.
Contact: BIS CAP International, Research Publications and Conferences Division, 77 Rumford Ave., Waltham, MA 02154, (617) 891-1550.
Inquiry 917.

IBM PC Users Club Meets in Darien, Connecticut

The Connecticut IBM PC Users Club meets regularly on the fourth Tuesday of every month at the Darien Public Library. The group, which has about 400 members, was started in 1982. Special-interest groups include database, Wall Street, programming, Lotus 1-2-3, fundamental operations, and "Smart" users.

You can reach the club's BBS, which is open to the public, at (203) 966-8865.
Contact: John McGinley, P.O. Box 291, New Canaan, CT 06840, (203) 762-0229 or (203) 966-4253.
Inquiry 916.

World Radio History

4 YOUR MIND ONLY . . .

1. FCC APPROVED MOTHERBOARDS AND SYSTEMS
2. 1 YEAR OF PARTS AND LABOR WARRANTY
3. 8 YEARS OF EXPERIENCE
4. 72 HOURS BURN-IN PERIOD

MOTHERBOARD FEATURES :

FEATURES	286-12B	286-20B	286-24B	386-20A	386-25A	386-25C
Design technology	VLSI®	C&T®	C&T®	TTL	TTL	TTL
Norton SI® up to	13.4	22.5	27.9	22.0	31.6	31.6
Landmark® up to	15.9	26.7	33.6	25.5	31.7	43.5
BIOS	AMI®	PH®	PH®	PH®	PH®	PH®
Shadow RAM	N	Y	Y	Y	Y	Y
Expansion Slots	8	8	8	8	8	8
Bus speed selectable	Y	Y	Y	Y	Y	Y
Coprocessor support	287	287	287	287/ 387	387/ 287	Weitek/ 387
Memory on board	4MB	8MB	8MB	6/16MB	6/16MB	6/16MB
CACHE memory	N	N	N	N	N	32K
Memory chip type	DIPP	DIPP & SIPP	DIPP & SIPP	DIPP	DIPP	DIPP

SPECIAL : With Wedge's image capture card, you can capture and store high quality graphic images in file formats compatible with most Desktop Publishing and Paint programs, from your video camera, VCR or live TV broadcasts.

SYSTEM STANDARD FEATURES :

* 1 MB memory
* 5.25" 1.2MB Floppy Drive
* 1:1 interleave HD/FD Controller
* Ports: 2 serial, 1 parallel, 1 game
* 101-key enhanced keyboard
* DOS version 4.01
* Math Coprocessor socket support

* Call for Custom Configuration System and detail information on our other products

How does FOUNTAIN 386SX compare with a true 386?

	128K NOP Loop	DO-Nothing Loop	Integer Add Loop	Integer Multiply Loop	String Sort and Move	Prime Number Sieve	802086 Instruction Set
FOUNTAIN-386SX	2.09	1.97	0.88	0.61	1.20	1.76	4.07
COMPAQ-386/16	2.09	1.86	0.99	0.60	1.37	1.87	4.23

PC Magazine Benchmark Tests (Time in Seconds)

FOUNTAIN 386SX
The Smart Choice.

INTRODUCTORY OFFER $1500

Inexpensive Multitasking Platform for PCs

P C-MIX, a multitasking environment for applications running under MS-DOS, lets you run up to three programs concurrently and switch from one task to the other with two keystrokes. The program works on systems as basic as the IBM PC, but it works best with systems that support operation beyond 640K bytes of RAM with EEMS memory, such as the IBM PC AT.

Unlike DESQview, which divides concurrent applications into windows on the screen, PC-MIX uses a full-screen display with concurrent applications operating in the background. Other features include configurable memory partitions, selectable task priorities, and true preemptive scheduling. If an application writes directly to screen memory, PC-MIX lets it take over the screen and uses a batch file to control the other applications.

PC-MIX requires an IBM PC with 256K bytes of RAM and DOS 2.0 or higher.
Price: $49.95.
Contact: Proware, 10719 Plano Rd., Suite 100, Dallas, TX 75238, (214) 349-3790.
Inquiry 894.

Program Helps Companies Stay Clean

F or businesses that want to implement a drug-free workplace policy, such as those with government contracts, Clean Slate provides database and reporting tools and can generate impartial, randomly selected lists of employees for drug testing. The program automatically generates audit trails for employ-ees or a contracting agency that provides evidence of nondis-criminatory and good faith compliance with current regulations.

The program also includes a set of policy guidelines that you can use to determine what actions your company should take when an employee tests positive for one of five classes of drugs. You can also use the program with a processor to customize notices and certifications. Clean Slate is based on dBASE IV and includes a run-time version of the program.

Clean Slate runs on the IBM PC with 640K bytes of RAM, a hard disk drive, and DOS 2.1 or higher. An update service is available for $195 per year.
Price: $695.
Contact: Clean Slate Software, Inc., 11260 Roger Bacon Dr., Reston, VA 22090, (800) 726-3440 or (703) 471-6071.
Inquiry 903.

Software Lets 3 + Mail Users Send Faxes Remotely

W ith GammaMail, 3Com 3 + Mail users can use a modem-equipped laptop computer to connect to their network and remotely send a fax message to any other fax machine. The program runs on the GammaFax CP, GammaLink's PC-to-fax board designed for networks.

With the GammaFax CP, you can connect eight boards to a fax server, all sending and receiving at 9600 bps. The board also ships with version 4.21 of its communications software, which GammaLink says is designed specifically for heavy network use.

GammaMail runs on the IBM PC with 640K bytes of RAM, a hard disk drive, the GammaFax CP board, and a network interface card.
Price: GammaMail, $995; GammaFax CP board, $1095.
Contact: GammaLink, 2452 Embarcadero Way, Palo Alto, CA 94303, (415) 856-7421.
Inquiry 904.

Microsoft Releases Presentation Manager Toolkit

T o give OS/2 application developers references and tools as they need them, Microsoft has released the OS/2 Presentation Manager (PM) Toolkit, which developers can buy as one package or in individual components.

The toolkit includes a set of graphics tools for PM, called Softset, four OS/2 PM books, hypertext-based Quick-Help documentation, 3 megabytes of sample code, and 2 hours of on-line support, all for $500.

Softset includes dialog box, icon, and font editors, a resource compiler, and the book *Microsoft OS/2 Programming Tools*. Softset is available for $150.

The three volumes of the *MS OS/2 Programmer's Reference Library* are available separately, priced from $19.95 to $29.95. *Programming the OS/2 Presentation Manager* is also available for $29.95.

The OS/2 PM Toolkit includes sample code, Quick-Help on-line documentation, and the Helpmake utility, which lets you add additional on-line documentation into the QuickHelp system. It is available to Softset owners for $150.
Contact: Microsoft Corp., 16011 Northeast 36th Way, Box 97017, Redmond, WA 98073, (800) 426-9400 or (206) 882-8080.
Inquiry 900.

New Glue Supports Color, Gray-Scale, and Hidden Notes

T he newest version of Su-perGlue, the print-to-disk utility for the Macintosh, now supports color and gray-scale. Solutions International calls the newest version SuperGlue II with GlueNotes. This means that the program has the ability to attach hidden notes and comments to any file that you can print. It does this by capturing an application's printer output and redi-recting it to a disk file. The program's ImageSaver II file does the redirecting, while SuperView lets you examine the file.

The program lets you create electronic printouts from most Macintosh applications, such as Excel or Page-Maker, so that anyone on a network or via telecommunications can view a newsletter or spreadsheet as it would appear on the printer, without requiring the application that created the file. You can also save electronic printouts as a folder of PICT documents, for a slidemaker or service bureau.

Other new features of the program include character lock, which holds a character's position in kerned documents, and font lock, which identifies fonts by name, not ID number.

SuperGlue II with Glue-Notes works on the Mac Plus or higher with 1 megabyte of memory.
Price: $119.95.
Contact: Solutions International, 30 Commerce St., Williston, VT 05495, (802) 658-5506.
Inquiry 899.

DBMS/COPY Moves Data from One Package to Another, Easy as typing 'Copy'

Works with 26 more software packages, too!

At last, DBMS/COPY, the essential utility for transferring data between 28 software systems. Simple to use like the DOS *copy* command. Just type *dbmscopy* followed by the source database and target database. For example,

Paradox → SAS
`c:>dbmscopy m.db m.ssd`

Reflex → Oracle
`c:>dbmscopy pat.rxd pat.oracle`

DBMS/COPY processes dissimilar numeric types, date types and null values. The *Plus* version, for the power user, also selects variables, filters records and computes new variables.

DBMS/COPY works quickly and accurately. No more reconciling differences. No more writing programs. You'll wonder how you survived without it.

DBMS/COPY makes adjustments for the way each system stores its data. Definitely a must for the users of more than one software package. DBMS/COPY is the key to software connectivity for IBM/PC database managers, statistical packages and spreadsheets.

DBMS/COPY is only $195. DBMS/COPY *Plus* is $295. Add $5 for shipping. All products carry a 30-day money-back guarantee. Texas residents add 8% sales tax.

(800) STAT-WOW

Conceptual Software, Inc.

To order, call our toll-free number or complete coupon and mail or fax.

DBMS/COPY, DBMS/COPY Plus, STAT-WOW, and PRODAS are trademarks of Conceptual Software, Inc.

World Radio History

MODFONE™ 2400

ONLY $199 Introductory Offer

2400 BAUD MODEM PLUS PHONE
NEW & UNIQUE
Specifications:

- Compatibility: CCITT V.22 bis, V.21, V.22 and BELL 212A/103
- Speed conversion: 2400/1200/300 bps
- Data throughput: 2400 bps
- Software compatibility
- Front panel indicator
- One year warranty

DEALER/OEM/DISTRIBUTOR INQUIRIES WELCOME

E-TECH RESEARCH, INC.

2700 Augustine Dr. #275   Santa Clara, CA 95054

(408) 982-0270
Fax: 982-0272

Computers For The Blind

Talking computers give blind and visually impaired people access to electronic information. The question is how and how much?

The answers can be found in "The Second Beginner's Guide to Personal Computers for the Blind and Visually Impaired" published by the National Braille Press. This comprehensive book contains a Buyer's Guide to talking microcomputers and large print display processors. More importantly it includes reviews, written by blind users, of software that works with speech.

This invaluable resource book offers details on training programs in computer applications for the blind, and other useful information on how to buy and use special equipment.

Send orders to:
National Braille Press Inc.
88 St. Stephen Street
Boston, MA 02115
(617) 266-6160

$12.95 for braille or cassette, $14.95 for print. ($3 extra for UPS shipping)
NBP is a nonprofit braille printing and publishing house.

The Power of Expansion

Holmes Microsystems gives your laptop real communication power through modular design.

The Holmes Correspondent™ and FAX'EM™ laptop enhancement cards are more than just modems or fax cards. Right now, if you need a 1200 or 2400-baud modem, fax capability — or both — you can have it, all on the same card. And when your needs change, Holmes changes with you.

With a Holmes enhancement card, your laptop computer has the same expansion capabilities you demand in your desktop computer. And the package includes very low bit-error rates, superior Holmes Gold chip sets, 9600-baud fax, and first quality, registered communication and fax software.

Holmes is setting the standard for laptop communications now. Why not move into the future with confidence knowing you have the best today, and the power of expansion for tomorrow?

Holmes products are available at computer dealers nationwide.

Or for more information, contact Holmes directly.

HOLMES
MICROSYSTEMS, INC.

2620 South 900 West / Salt Lake City, Utah 84119 / 800-443-3034 / FAX 801-975-9726

World Radio History

Circle 487 on Reader Service Card

SHORT TAKES

Portfolio

Altima One

Finesse

MacroMind Director

MultiPlus

A Good Thing in a Small Package?

The trouble with computers is that they're too expensive, too big, too heavy, and too inconvenient. Atari thinks it has an answer to all those problems in a portable computer that's handy, light, small, and inexpensive. The big question, of course, is whether Atari can deliver the full power of an IBM PC clone in a computer that will fit in your pocket.

The good news is that the size problem is licked. The Atari **Portfolio** is a hand-held clone that folds to about the size of a VCR tape—8 by 4 by 1 ¼ inches. It fit easily into my inside suit-coat pocket, although I wouldn't necessarily want to carry it there all day—the Portfolio weighs about a pound with its batteries (three AA cells).

So far, so good. But does it make the grade as a PC clone?

The Portfolio uses Intel's CMOS version of its venerable 8088—the CPU that was in the original IBM PC. At 4.92 MHz, it's slightly faster than a standard PC or XT but far slower than most clones. It also has less RAM and a slower CPU than most clones

offer today (the Portfolio does have a provision for up to 640K bytes of RAM).

The keyboard is the first and most obvious place where the Portfolio's downsizing presents a problem. It's usable for two-finger typists—quite usable, in fact. The keys are plastic, not the rubberized "chiclet" keys that some calculators use. But you won't have much luck touch-typing on this scaled-down keyboard, at least not without lots of practice. I found it adequate for taking notes, but I wouldn't want to use it for more than a few hundred words of typing.

The screen, like the keyboard, is designed on a small scale. However, it's easily readable. The Portfolio emu-

lates a monochrome display adapter, and although the display contains 240 by 64 pixels, the system can't really make use of graphics (at least not without special driver software). As with the keyboard, I found the screen to be quite reasonable for small, quick jobs—such as note taking and quick calculations—but I'd hate to take on a major task with it.

There will eventually be two ways of transferring software and files into the Portfolio. One way is through the serial port, using LapLink-style software. But right now, the Portfolio doesn't *have* a serial port. Atari says that a "smart cable" that attaches to the Portfolio's expansion port

and comes out as a standard serial port will be available shortly for under $50.

The other way of transferring files from a desktop PC is by copying them onto a Portfolio memory card—a solid-state, removable RAM disk that's about the size of a credit card. It's the Portfolio's answer to disk drives. The memory cards slide into the left side of the machine and offer up to 128K bytes of RAM or 4 megabytes of ROM.

I saw a beta version of a new application designed for the Portfolio, but I *didn't* see any regular PC programs running on the machine. And I didn't see the soon-to-be-released memory-card drive that will fit into a standard PC slot, for transferring files back and forth between a desktop PC and the Portfolio.

The final element of the system is a memory-expansion pack that will let you boost the Portfolio's RAM from 128K bytes to a full 640K bytes. Like the smart cables, it will plug into the Portfolio's expansion port and will be essential for using the Portfolio with standard PC software. Also like the smart cables, the memory expansion wasn't yet available when I looked at the Portfolio.

I liked the Portfolio—I really did. But the questions I still have about software compatibility are serious.

The entire operating system is in ROM, and that may affect the operation of some programs. I'd like to know whether standard programs like Lotus 1-2-3, WordPerfect, dBASE III Plus, and SideKick will work in the Portfolio. Atari claims that the built-in spreadsheet uses Lotus-compatible files. I'd like to see that for myself.

I'd also like to see what
continued

THE FACTS

Portfolio
$399

Standard configuration:
It runs at 4.92 MHz and
has 128K bytes of RAM,
a 40-column by 8-row
display, and a 63-key
miniature keyboard.

Atari
1196 Borregas Ave.
P.O. Box 3427
Sunnyvale, CA 94088
(408) 745-2000.
Inquiry 1034.

effect the 40-column by 8-row screen has on conventional software, and how usable the keyboard is with software that's not specifically designed for the Portfolio. I'd like to see how it performs under the BYTE benchmarks.

Disk access should be very fast with the memory cards, while number crunching may be miserably slow.

I enjoyed using the Portfolio. It's wonderfully lightweight and splendidly convenient. When I used it, I thought

for the first time that I was using a truly portable computer—a truly personal computer. But for all its appeal, it's not finished yet. Right now, at $399, the Portfolio isn't a toy, but it's an expensive executive notepad and

pocket calculator. Once it's complete, with the smart cables and memory expansion available, it could be the first PC clone that will fit in your pocket without emptying your wallet. □

—Frank Hayes

A Good Luggable

If you're the kind of person who lugs your machine from home to office and back, you'll probably like the **Altima One**.

The Altima One 80286 system has plenty of features and weighs in at a totable 15 pounds. It includes a built-in 2400-bps modem; a 20-megabyte hard disk drive (soon to be 40 megabytes); a tolerable, detachable 101-key keyboard; a decent supertwist, backlit LCD; a CGA screen; a mouse (and a place to store it); and a good bit more.

Things I like about this machine include the fact that it

has an automatic setup program, runs through a visible set of diagnostics on boot-up, acts more like a desktop than a portable, and lets you choose between black text on a white background or the reverse. I also appreciate that just about everything the company says will work does.

DOS 4.0 and SideKick Plus come with the Altima One, along with a mouse. I looked at a preshipping version, and the mouse wasn't included. The system also comes with 1 megabyte of RAM (expandable via the addition of single in-line memory modules to 5

megabytes), 640K bytes of regular memory, serial and parallel ports, and an expansion slot for half-size 8-bit cards. It also has some built-in security functionality with a password request upon boot-up and runs with an average wait state of 0.7.

I ran the BYTE benchmarks on the system, and it ran as well as I guessed. Its CPU index is 2.02, a bit faster than the Zenith SupersPort 286, which comes in at 1.55. And the disk index is 1.34, compared to the SupersPort's 1.06.

Some things I didn't appreciate are the tinny-sounding and -feeling keyboard, the fact that the handle gets in the way while you're working, and the larger-than-life on-screen characters in color mode.

Some kibitzers commented on the forbidding, robotic look of the machine itself, and they wondered out loud what this machine has that others similar to it don't. My answer is that it's a full-featured and affordable machine that, without its external battery pack, you can tote around and not get a hernia.

A few caveats. If your software doesn't work in monochrome mode, try it in color mode. If you have a hard time getting used to the configuration of a non-IBM keyboard, you can attach an IBM PC-compatible keyboard.

Other than a few minor nitpicks, the machine seems to do what it was designed to do—amazing in itself. □

—Janet Barron

Logitech Brings Finesse to Low-Cost Desktop Publishing

Low-cost desktop publishing packages haven't exactly threatened the more expensive and capable programs like PageMaker and Ventura Publisher, and they probably never will. But Logitech's **Fi-**

nesse brings respectability and panache to the neighborhood of low-end page-makeup software.

If you're just getting started at using a PC to lay out documents and pump text into

them, Finesse is an excellent package. It runs on a pretty basic system, with its most exotic requirements being 640K bytes of RAM, a hard disk drive, and a CGA board. I worked with a beta copy of the

program on a Compaq 286 with a VGA display. It looked sharp and ran flawlessly.

Finesse is a GEM application, so you work in the nice Digital Research environment

continued

of icons, menus, and win-
dows. The program uses the
technique of text frames; be-
fore you can pour the text into
the publication, you have to
first set down frames—boxes
that you draw on the page with
the "text frame" tool. This is
an easy operation, but it does
require physically, rather than

just conceptually, blocking
out your page design before
you start laying down text. You
can't put a thing down on the
page unless you put it in a
frame. This takes some per-
sonal adjustment if you've
been working with a more
flexible package like Page-
Maker.

Importing text is just a mat-
ter of clicking on the frame
and then on the file you want to
paste down on the page. Fi-
nesse will work with ASCII,
Word, WordStar, WordPer-
fect, 1st Word +, and GEM
Write formats. The same fill-
a-frame procedure works with
graphics; after you've drawn
the frame, the program will
pull in bit-map, TIFF, PCX
(as in PC Paintbrush), and
Metafile images.

Getting text to go from one
frame to another involves
using a *chaining* tool. After
you've told the program which
columns are linked, which you
do by clicking inside the ap-
propriate frames, it will run
the text into the appropriate
columns. Finesse doesn't have
the powerful auto-flow capa-
bilities of more expensive

packages, but its chaining pro-
cedure is relatively painless if
you're working on a short pub-
lication. Anything over eight
pages could mean racking up a
significant number of mouse-
clicks.

The copy I worked with in-
cluded two fonts (called Dutch
and Swiss) and came with sev-
eral more Bitstream fonts on
separate disks. I had trouble
with the font installation pro-
gram, though, and was unable
to work with styles other than
those embedded in the pro-
gram. The final version,
which the company expects to
have ready by now, will come
with several more typefaces, a
collection of clip art, and a few
prefab document styles, a
Logitech spokesperson said.

Finesse has the basic tools
you need to put together a bro-
chure, newsletter, or other
short document.

Screen redrawing seemed
pretty slow, though, so laying
out a long publication could be
tedious. However, Logitech
has added some shortcuts that
are very handy and speed up
some text-manipulation oper-
ations. For example, you can

raise or lower the size of se-
lected text with a simple key
combination: Alt-4 kicks the
text up to the next point size,
and Alt-3 brings it down
again. These and other short-
cuts are a nice touch.

I wish the developer would
add the capability of working
on a page when you're looking
at it in "full-page" mode,
though. As it is now, you can
work on a document only
when you've got it in "actual
size" view, which unfortu-
nately means you can see only
that part of the page that fits on
the screen. Setting a headline
that runs across the page
means you have to scroll back
and forth or toggle between
actual and full-size modes.

Despite a few limitations,
Finesse is a fine program for
producing short documents. If
you've never used a publishing
package and you don't want to
climb a steep learning curve,
this is the software for you. Fi-
nesse is easy to use, works
well, and will run on most
low-end PCs. It's not Page-
Maker, but it's not trying to
be. □

—D. Barker

Let the Mac Entertain You

Will affordable multime-
dia production capabili-
ties be the most important
breakthrough in computing
during the 1990s? After work-
ing with MacroMind's new
product, **MacroMind Direc-
tor**, I'm beginning to think
that this might be the case.
MacroMind Director does in-
deed make video production,
complete with animation and
sound, accessible to anybody
who has a Macintosh.

MacroMind Director is a
greatly improved and en-
hanced version of Video-
Works, which has been the
company's primary product
since its inception in 1984.
According to the product liter-
ature, over 100 new features
are found in MacroMind
Director.

Some important features

include automatic animation,
a new color paint program,
new music and sound capabili-
ties (including MIDI control),
and a greatly improved user
interface with on-line help. A
HyperCard driver, which lets

MacroMind Director se-
quences be included in Hyper-
Card stacks, should be avail-
able soon.

Although you can get
started in a few hours, Macro-
continued

Mind Director is not a trivial program. You'll need to expend some time and effort to learn and master its capabilities.

The program consists of two parts: the Overview and the Studio. The Overview is basically a slide sorter with a control panel similar to that of a VCR. In the Overview section, you create individual frames or visual images that you can later combine into a "movie" using the Studio section of the program. Supported file formats include Scrapbook, PICT, PICS, MacPaint, Glue, and sound files from sound-sampling programs like MacRecorder. You can overlay images, animated sequences, and sounds in a single frame.

In the Studio section, you create the animated sequence of frames with the appropriate timing. The main workpiece in the Studio is the Score, which is similar to a spreadsheet in appearance. The rows represent separate frame sequences or channels, of which there can be up to 24. Each channel can contain any of the multimedia components (e.g., sound, graphics, or text), which you select from a Cast consisting of the library of images you created or imported using the Overview section.

The columns represent time. Therefore, multiple channels can appear simultaneously in the score. You can time the starting points and endpoints of each channel individually, so you can develop complex video sequences. Video sequences can also be controlled with the mouse button, if you're giving a talk simultaneously and want to click the mouse button to advance the frame sequence.

While MacroMind Director runs on all Macs with 1 megabyte of memory, it runs best on a Mac II, particularly since you then have 256 colors to work with. You can have separate color palettes for each frame, allowing much flexibility with the choice of colors. You can get gray-scale imaging on a Mac SE, however. And performance is roughly the same on a Mac Plus or SE and on a Mac II. This is because the Mac Plus and SE don't have to worry about processing all the 8-bit color information.

Whether you use a Mac Plus/SE or a Mac II, getting into serious video production and presentation is not a minor investment. Aside from the computer and the software, you'll need a large screen to display the video. If you want to use scanned images, you'll need a scanner. If you want to output video to a VCR, you'll need a genlocking card that can convert the digital RGB output to the analog National Television System Committee format required by TVs and VCRs.

For a company or an educational institution, the expense for all this equipment makes sense. And there's no question that it costs thousands of dollars less than the traditional equipment required for video production.

However, for the hobbyist or casual user, full video-production capabilities require a pretty deep pocketbook. And although using MacroMind Director on the 9-inch screen of your Mac SE may prove to be entertaining, you can't really do presentations for other people on such a small screen.

On the other hand, Macro-Mind Director is a serious production tool for professionals who need its presentation capabilities. After a few hours of working through the tutorials, you can put together animated presentations, combining bulleted text charts, graphs, music or voice sounds, graphics images, and, if you have the equipment, scanned images or video sequences from a VCR. Overall, I think MacroMind Director is well designed and can be of great benefit and utility in all forms of visual communication. □

—Nick Baran

MultiPlus Takes on SideKick

THE FACTS

MultiPlus
$99

Requirements:
IBM PC or compatible
with 512K bytes of RAM,
DOS 2.1 or higher, and
a hard disk drive.

SunFlex Software
1447 Peachtree St.,
Suite 503
Atlanta, GA 30309
(404) 874-9699
Inquiry 1038.

In the years since SideKick first appeared, many competitors have come and gone. So I have to admire companies that have the spunk to tilt at Borland's madly spinning windmill. A new contender is SunFlex Software, with a product called **MultiPlus**.

MultiPlus is yet another in the seemingly interminable line of desktop management packages. It has the usual assortment of SideKick-like features, including a word processor, a calendar and appointment scheduler, an address database and telephone dialer, and more. But I found MultiPlus an oddly eccentric package, filled with both nice touches and some maddening oversights.

All of its myriad files took up nearly a megabyte of disk space, so I couldn't use it with my floppy-disk-only laptop computer. It is possible to save disk space by eliminating one or more of MultiPlus's individual modules, and the core

RAM-resident module takes up just 10K bytes of RAM.

The word processor is full-featured, not just a notepad like SideKick's. And the five special-purpose calculators in MultiPlus are way ahead of Borland's. But the calendar/appointment maker has what I consider an unforgivable problem: There's no alarm option to remind you of an appointment. Then there's the address database and phone dialer. It does the job, but there's no telecommunications option.

SunFlex is pushing Multi-Plus's built-in vaccine feature to set it apart from SideKick. But I think vaccine programs are a fad, good only for the truly paranoid and those who rely on public domain software. All in all, I give Multi-Plus a B for effort, but there are too many rough edges. Although SideKick Plus sells for twice as much as MultiPlus, I'll stay with Borland. ∎

—*Stan Miastkowski*

Circle 314 on Reader Service Card

SEND AWAY FOR YOUR VERY OWN COMPUTER STORE.

When you buy a computer, about 35% of your money goes to the store.

But we'd much rather see your money go to somebody who deserves it a lot more.

You.

So we give you a completely different kind of computer store.

The Dell Computer Store.

Instead of a crowded, high-overhead showroom, you get our brand new 44-page catalog. Which gives you a full line of 386™ and 286 systems, printers, peripherals, software, and accessories.

And since you buy direct from the manufacturer, you save about 35%.

But there's a lot more to it than saving money. We offer you the most complete service and support in the industry. Including a 30-day money-back guarantee. A toll-free technical support hotline. Self-diagnostic software. And next-day deskside service from

Xerox Corporation. If you read our ad in the front of this magazine, it will tell you a lot more about the systems we offer. And the service we put behind them. So if you'd like a much more intelligent way to buy a personal computer, have a look at our ad. Or call us at 800-426-5150. Or send us the card on this page.

And we'll send the best computer store you've ever seen.

One you don't have to set foot in.

DELL
COMPUTER
CORPORATION

TO ORDER, CALL US.
800-426-5150.
FOR DELL IN CANADA, CALL 800-387-5752.

The Ever-Shrinking,

Agilis and Zenith announce tiny computers that

I t seems that the smaller laptop PCs get, the more desirable they become. A fully functional XT compatible that travels as well as a hardcover novel is a powerful business tool. Furthermore, smaller form factors provide the opportunity for computers to enter interesting and unique new markets.

This month, we look at the **Zenith MinisPort** (page 94) and **Agilis System** (page 91) computers. The Zenith is a laptop in the truest sense of the word and very portable—a businessman's dream. Although the Agilis could serve as a laptop, it incorporates the latest technology to produce an expandable hand-held system intended for use in remote locations. These machines represent both evolutionary and revolutionary trends in laptop technology.

Ever-Expanding Laptops

broaden the market for laptops

Agilis Hand-Held Workstations: Computing Power in the Field

Nick Baran

In recent First Impressions, I've focused on computers that break new ground in price versus performance. I now have the opportunity to look at a machine that breaks new ground in size versus performance—the Agilis System hand-held workstation. In its top-of-the-line configuration, the Agilis System is a complete 80386 machine about the size of a notebook, about 3 inches thick and weighing 8 pounds. A lower-performance 80C88 version can weigh as little as 4 pounds.

The Agilis System is not just another laptop. It is designed for use outside the office and in harsh environments requiring mobility but also networking and remote communications capabilities. You can operate it with one hand using a touchscreen interface. And you can use it on a wireless Ethernet network with a range of up to 1 kilometer.

In the last 10 years, the personal computer has dramatically changed and improved the way we work in the office. But a major part of the work force has been left out of the computer revolution—namely, those who work away from the office or "in the field." These workers include maintenance and service person-

A notebook-size 80386 system showing the console slice, keypad slice, battery packs, and PrairieTek hard disk drive.

nel, sales representatives, workers on the factory floor or at test sites, public safety workers, building and utilities inspectors, military personnel, and many others.

Computers in the field could eliminate the paperwork associated with schedules, maps, diagnostic procedures and manuals, inventory, and telemetry, to name a few. And, if the computers in the field are connected to a network, they can communicate with other computers, such as file servers at the home office.

Although laptop computers offer some of the features needed for field work, they have major limitations: They are too large to operate comfortably while standing up; they have limited battery power and are dependent on wall-outlet power

sources; and they have limited networking capability and are not designed for harsh environments.

Any Way You Slice It

Created by former GRiD, 3Com, and NeXT engineers, the Agilis System is designed specifically for use in the field. Based on the Intel processor line, the Agilis System takes advantage of the latest advances in miniaturization and high-density electronic packaging. It is built on the concept of modular slices, each slice providing a component of the system, such as the CPU component or "processor slice," a communications slice, a data-storage slice, and a battery-power slice.

Made from ruggedized plastic, each slice is about one-third the size of a sheet of paper (8⅞ by 3¾ by 1 inch). Each slice can connect front-to-back or top-to-bottom to another slice by means of the AgileConnect interface, which consists of an 802.3 Ethernet network interface operating at 10 megabits per second and a power distribution interface. The Ethernet and power paths are integrated into a single 34-pin connector built into every slice.

Power can come from nickel-cadmium battery packs, converters for standard 110-volt or 220-V alternating current, and 12-V automotive or 28-V military *continued*

Nick Baran is a BYTE senior technical editor based in San Francisco. He can be reached on BIX as "nickbaran."

Photo 1: *The components of the 80C88 processor slice. Note that the entire 80C88 system fits into a slice about the size of one-third of a piece of paper. The logic board shows a 512K-byte memory card inserted in the memory card slot. The 80C88 system boards feature surface-mounted components and are mounted back-to-back. The primary chip set is manufactured by Western Digital.*

Photo 2: *The 80C88 with the touchscreen console slice and a battery pack. This configuration constitutes a complete touch-operated personal computer. Note the programmable function keys around the perimeter of the flat-panel display.*

Photo 3: *The internals of the 80386 Agilis System. The underside of the board has a bay for a 20-megabyte 2½-inch PrairieTek hard disk drive. The primary chip set is manufactured by Headland Technology. Note the serial and parallel ports at the lower edge.*

direct current. The Agilis System software includes utilities for monitoring power consumption and battery life.

One of the important breakthroughs of the Agilis System is its efficient power distribution throughout the system. The Ethernet/power bus passes through each slice and includes a transceiver with a circuit that detects packets on the network and powers up the circuit automatically. When the circuit is idle, the transceiver is turned off. This design greatly minimizes the power requirements for network communications. The system software also includes utilities that automatically shut down the hard disk drive and the backlit display after a specified period of inactivity.

The AgileConnect interface includes a miniaturized, but fully functional, AT and XT bus interface. The 8-bit XT bus interface is contained in a 68-pin U-block connector. The 80386 processor slice includes an additional 34-pin connector, which extends the bus interface to the 16-bit AT standard. Slices connected end-to-end simply latch together, directly mating the male/female Agile-Connect connectors. Slices placed top-to-bottom use a U-block connector to make the interface connection.

The Agilis System is designed for use outdoors or in dusty or damp indoor environments. It is reasonably waterproof and dustproof and can handle rough use. The limiting factors on its durability are the glass of the flat-panel display and the PrairieTek hard disk drive, if installed.

The heart of the Agilis System is the processor slice, which comes in 9.54-MHz 80C88 and 20-MHz 80386 versions. The 80C88 processor slice consists of two 6-by 2-inch logic boards mounted back-to-back in a single slice. The 80C88 system comes standard with 640K bytes of RAM and a card slot for a removable 512K-byte memory (RAM/ROM) card. Note in photo 1 that the RAM/ROM card is inserted into the slot on the logic board. It is the main storage and boot device for the 80C88 system. However, you can also plug a standard floppy disk drive slice into the 80C88 processor slice.

The 80C88 processor board uses a Western Digital chip set that supports both 4.77- and 9.54-MHz clock speeds and standard XT-compatible direct memory access and interrupt control, keyboard input control, and memory management for up to 640K bytes. The board includes an RS-232C connector and an external keyboard connector, as well as the standard XT-bus and Ethernet/power connectors. The complete

80C88 slice requires only about 4 watts of power.

The 80C88 also includes 128K bytes of flash, electrically alterable ROM, which you can use to load MS-DOS or custom applications into ROM. In this way, you can use the 80C88 slice as an embedded processor or for custom touchscreen applications. In fact, Agilis expects that a very common configuration of the 80C88 processor will simply include a battery slice and the touchscreen console slice (see photo 2). This configuration weighs about 4 pounds.

The console slice is one of the most interesting features of the Agilis System. It's actually the size of two single slices and features a backlit, 6-inch diagonal flat-panel display (built by Kyocera) that supports EGA gray-scale and bit-mapped graphics with 640- by 480-pixel resolution. The console slice has an infrared sensor for use with attached or detached keypads. And most important, the console slice has a built-in processor that supports touchscreen operation, including mouse and keyboard emulation.

The touchscreen lets you press on the screen and activate a command. You can also use your finger to move the mouse cursor. You can program the function keys on the perimeter of the touchscreen to execute macros, or you can simply use them as DOS function keys. The keys along the bottom row of the display console control cursor movement and the Enter function.

While the touchscreen interface works, it needs applications specifically designed to take advantage of it. With the programmable console keys, numerous possibilities exist for field-specific applications that maximize the use of the touchscreen. A reflective display slice without touchscreen capability is also available.

The top-of-the-line system is based on the 20-MHz 80386 processor (see photo 3). The 80386 system includes a bay for a 20-megabyte 2½-inch PrairieTek hard disk drive. The 80386 processor slice is actually the size of four single slices. With the PrairieTek hard disk drive, the 80386 system requires only about 9 W of power. The 80386 slice includes two serial ports, a parallel port, and the Ethernet/power and AT-bus interfaces. The 80386 board uses a G-2 chip set that features 360-pin surface-mount technology. The 80386 system is available with 1, 4, or 8 megabytes of memory.

Other Options

The beauty of the Agilis System is that the slice technology lets you configure it

in numerous ways, depending on your requirements. In fact, Agilis intends to license its AgileConnect interface to third-party manufacturers who want to build optional slices for the Agilis System. At this writing, Agilis has completed battery slices, a wireless packet-radio communications slice, and a floppy disk drive slice. Agilis is also developing a general-purpose expansion slice that will support standard XT and AT half-length expansion cards, such as internal modems or external video adapters.

Of particular interest is the wireless packet-radio communications slice. It offers 230,000-bps network communications within a range of 1 kilometer outdoors and about 100 meters indoors. The

The Agilis System doesn't compete directly with standard PC and laptop prices.

packet radio operates in the spread spectrum frequency range of 902 to 928 MHz and supports up to 16 channels. The communications slice requires about 15 W when it is transmitting packets but is automatically powered down to 2 W when idle. The communications slice is the size of two single slices. I did not see the communications slice demonstrated.

System Software

The Agilis System comes with either MS-DOS 3.3 or Interactive Unix V.3.2. DOS comes either on a floppy disk or on the 512K-byte RAM card for the 80C88 system. Unix is available on a floppy disk. Both operating systems come with additional system configuration utilities and system programs.

The additional software includes a System setup panel, which you can configure at system start-up to enable or disable certain components in the system, such as the serial ports or extended memory (in the 80386 version). A Power Management panel lets you specify whether you want video or audio low-battery warnings, and whether you want the system, hard disk drives, or Ethernet controller shut off when they are idle. You can also specify the threshold voltage at which the low-battery warning

should come on.

The system software includes the Agilis Action Point utilities. The utilities contain configuration files for specifying mouse or keyboard emulation and for programming the console keys. Another utility installs DOS or other applications in the 128K bytes of flash ROM in the 80C88 slice. An extension to the DOS FORMAT command is included for formatting the 512K-byte RAM cards used in the 80C88 slice.

Configurations and Prices

Because of the durability and density of the electronic packaging, Agilis components are not inexpensive. The Agilis System does not compete directly with standard PC and laptop prices and is not intended to compete in the traditional desktop or laptop market.

A typical high-end system would consist of an 80386 processor, 4 megabytes of memory, a 20-megabyte hard disk drive, the console slice, a keypad slice, two battery slices, and a power converter. Such a system would cost just over $12,000. An intermediate system might simply be an 80386-based 3 + mail server (with 3Com's 3 + network E-mail software installed on the hard disk drive), which would consist of an 80386 slice with a hard disk drive and a power supply. This setup with 1 megabyte of RAM would cost about $6600.

At the other end of the spectrum, an 80C88 system with the touchscreen console slice, the 512K-byte RAM card, a battery slice, and a power converter would cost about $5000.

Hands On

I had an opportunity to try out an early prototype version of the 80C88 slice with the touchscreen console. The system was running Microsoft Windows, a paint application, and a CAD drawing display application called FastView, all installed on the 512K-byte RAM card. This system was small and light enough that I could stand and cradle the system on my left arm and operate it with my right hand. Using the console's function keys, I could make changes to the Setup Panel and to the Power Management configuration. I ran the FastView application and loaded a CAD drawing on the screen. Using the console keys, I could Pan and Zoom on areas of the drawing.

To make a long story short, the system works. However, the system I tested needed some improvements in the display backlighting and the touchscreen sensitivity. It was hard to see the mouse

cursor as I dragged it across the screen, and I had to keep adjusting my viewing angle so that I could view the screen. Agilis engineers assured me that commercial versions of the touchscreen would have the necessary improvements.

I also tried assembling and disassembling various slices. I was impressed with the quality and solid engineering of the components. Each slice has guide rails, which make assembly of slices literally a snap. Once assembled, the slices are locked together with spring-loaded tabs on each side.

The real promise of the Agilis System hinges on the development of innovative software that can take advantage of the system's touchscreen and networking capabilities. While clearly not designed for the everyday user, the machine could have enormous utility in all kinds of field operations. According to Agilis's marketing director, Bert Keely, the machine has generated the greatest interest from automotive and airplane manufacturers, who intend to use the hand-held worksta-

tion for diagnostics and data retrieval for mechanics and test engineers.

I am impressed with the innovative engineering of the Agilis System. While I did not get a chance to work with a final production version of either the 80C88 or 80386 system, the preliminary components appear to be well designed and manufactured. I also saw preliminary versions of the documentation, which is thorough and well written. The Agilis System points the way to new advances in portable computing. □

The Littlest Zenith

Michael E. Nadeau

Looking at many of the so-called "laptop" computers makes you wonder how they got the name. Though portable, few of them are practical for computing on the go, even if they do fit comfortably on your lap. Having a computer that you can easily pick up and move to another location is one thing; using it during transit is another.

My ideal laptop would fit into a briefcase with room to spare and weigh under 5 pounds. Its screen would be easily readable in poor lighting; its nonvolatile memory would be large enough to store program and data files. The battery life would be at least 4 hours. A 2400-bps modem would be a must, as would be ports for an external monitor, a floppy disk drive, and a printer. The keyboard would be responsive and intelligently designed. The laptop would also have a painless means of porting programs and data files to and from my desktop PC. And it would have all this for under $1000.

No such critter exists, but the Zenith's new MinisPort comes closer than any other laptop, with the exception of the NEC UltraLite (see the review "The Painlessly Portable PC" by Mark L. Van

Michael E. Nadeau is the associate managing editor of the reviews section of BYTE. He can be reached on BIX as "miken."

The Zenith MinisPort. The configuration shows the 2-inch 720K-byte floppy disks.

Name and Bill Catchings on page 161). At around $2400 (Zenith had not set final prices at press time), the MinisPort beats the $3000 4½-pound UltraLite on price, but at 12½ by 9⅖ by 1⅓ inches and 6 pounds, the MinisPort narrowly loses to it in the size and weight categories. Minor faults aside, the MinisPort should be a desirable entry in the little-laptop arena.

BYTE's preproduction evaluation unit came with the standard 1 megabyte of surface-mount RAM, up to 368K bytes of which can be configured as a nonvolatile RAM disk, EMS memory, or a combination of both. You can configure a 1-

megabyte upgrade option ($799) as either additional RAM disk space or EMS memory. DOS 3.3 resides in 360K bytes of ROM, along with Rupp Corp.'s FastLynx file transfer program.

The MinisPort's 80C88 CMOS CPU is switchable between 4.77 MHz and 8 MHz via the keyboard or software. It has a Centronics-type parallel port and an RS-232C serial port with a DB-9 connector. The external video port supports both CGA-type RGB-intensity TTL-level and composite monochrome output. The fourth port is for an external floppy disk drive. A tiny slot is also available for a Saltine-size 1200-bps modem card ($299), which was unavailable at this writing. The MinisPort's screen is a backlit, supertwist, 640- by 200-pixel LCD.

A unique feature of the MinisPort is its double-sided, double-density, 720K-byte 2-inch floppy disk drive, the first of its kind to be used in a laptop or any other kind of personal computer. The floppy disk drive and disks look like scaled-down 3½-inch versions. An external 3½-inch floppy disk drive is a $299 option.

Look and Feel

Somewhat larger than a kid's Etch-A-Sketch, the MinisPort is easily totable. Two of them would fit snugly into my briefcase.

Flipping up the screen reveals a typical laptop keyboard arrangement. The function keys are on the top row, and a numerical keypad is embedded within the alphanumeric keys on the right and accessed via the Fn key. Zenith committed no "mortal sins" in designing the keyboard; the only idiosyncrasy is the placement of the left single quotation mark (') and backslash (\ ) keys in the column farthest to the right. Since you don't use these keys frequently, their position is a minor inconvenience.

The keyboard feel is firm and responsive. I quickly became comfortable typing on the MinisPort. Functions called by the Fn and Alt keys are color-coded— a nice touch. Although the MinisPort is smaller than nearly all other MS-DOS laptops, the keyboard didn't feel cramped.

The LCD display, while not state-of-the-art, is adequate for most situations. Working in a dimly lit room, I had a little trouble picking out the underline cursor in a screenful of text. You can adjust the contrast and brightness via slide controls at the bottom of the screen, and you can position the screen from 90- to 180-degree angles. The 8½- by 3¼-inch viewing area exhibits some horizontal distortion of graphics images typically associated with such displays.

The 12-V battery is lighter than most and slides in and out easily from the left side of the case. It's rated for 3 hours, although I got only 1 hour and 45 minutes on a full charge (which takes 10 hours). Extra battery packs are $79 each. The MinisPort warns of imminent shutdown with a flashing red LED indicator and intermittent beeps. According to the documentation, I got shorter battery life because I had the screen backlighting on and several ports enabled. A Zenith spokesperson said that up to 5 hours on a charge is possible, although not guaranteed, if you don't use the LCD backlighting.

All ports are easily accessible at the rear, and the modem line is on the left side next to the battery. A handle swings out at the front. The MinisPort has all the usual LED indicators, plus ones for the silicon disk drive (SDD) and padlock.

Many businesses and users fret over losing their laptops to theft. The MinisPort has a unique "security bracket" to prevent theft. This steel bar slides out from the right rear of the case and has a hole for a padlock. Attempts to break off the bar destroy the computer, since it is attached directly to the motherboard. A determined thief could saw through this bar, but it wouldn't be easy.

MFM-180

Zenith provides a multifunction monitor program, MFM-180, which lets you set operating parameters, examine and manipulate areas of memory and register contents, test system components, set video commands, and change the boot drive.

Pressing Ctrl-Alt-Insert gets you to the MFM-180 -> prompt. From there, you can access the monitor's utilities. Most users will need only the Setup program, which establishes operating parameters. The Setup menu lets you set the time and date, CPU speed, video display, backlight time-out, and boot drive. You also can enable or disable the ports and RAM disk backup, and you can allocate RAM to either the RAM disk or EMS memory.

Somewhat larger than a kid's Etch-A-Sketch, the MinisPort is totable.

The Setup program also lets you establish password protection against unauthorized use of your MinisPort. You can also change fonts; your choices are Norwegian, Turkish, Greek, Hebrew, and the default, U.S./English. This feature was not implemented on the BYTE machine.

Data Come, Data Go, Data Stay

Manipulating data and program files on a laptop can be a problem. Zenith provides several ways of running application software and transferring software to and from a desktop PC.

You have three ways to get software and data into the MinisPort: You can beam it over via the serial port using FastLynx; send it over a null-modem cable using a communications program; or use the petite, new 2-inch floppy disks in conjunction with an external 3½- or 5¼-inch floppy disk drive attached to the MinisPort or an external 2-inch floppy disk drive ($349) attached to your desktop PC. Zenith will offer an extra-cost starter kit that includes a slipcover, the cable for FastLynx, and 10 2-inch floppy disks. The kit's price was not set at press time.

Most users will want an external drive on either the MinisPort or their desktop PC. No commercial software is available in the 2-inch format, and a Zenith spokesperson said that the company sees these floppy disks only as a means of transferring programs and data files. Panasonic and Sony, however, produce the 2-inch media. It is similar to media used in digital cameras. Zenith had no pricing information on the disks.

While nonstandard media has its problems—potential availability problems, higher cost, lack of commercial software—conventional formats would simply not work in laptops as small as the MinisPort. The drive hardware would add too much weight and bulk. Zenith obviously hopes that the 2-inch media will become standard for laptops of the MinisPort class.

Once you have your files in the MinisPort, you can use them from either the floppy disk or the SDD (drive D) in RAM. True to its billing, drive D does behave like a very fast hard disk drive, although a small one. I could not run all the BYTE disk I/O benchmarks because some require a megabyte of disk space to run. The DOS seek tests, however, showed a time of 3.90 seconds for a sector read and 18.22 seconds for a 32-sector read. The IBM PC AT times were 14.95 and 65.18 seconds, respectively.

A pair of lithium batteries provides up to three days of backup power to the RAM memory, so you won't lose data in drive D when the main battery goes dead or when you change it. You can turn off the battery backup option from the Setup menu.

The downside is that the 368K bytes available as an SDD in the 1-megabyte model is just not enough to run most meaningful applications. PFS Professional Write barely fits if you leave the spelling checker behind, and you can forget XyWrite. The extra megabyte available for the SDD in the 2-megabyte model is a necessity.

Performance per Pound

With its 80C88 CPU, you wouldn't expect blistering performance from the MinisPort, and the BYTE CPU index of 0.38 bears this out. This rating makes the little Zenith either a fast XT or a slow AT, depending on how you look at it.

The applications that anyone is likely to use on a computer like the MinisPort don't require a quick CPU. Word processing, communications, and light information management will be the major applications for the small laptops. I saw
continued

no noticeable difference in these areas between the MinisPort and my 10-MHz 80286 AT clone. Besides, the speedy SDD tends to compensate for performance penalties that the CPU imposes.

MinisPort vs. UltraLite

The laptop most comparable to the MinisPort is NEC's UltraLite. Aside from the differences mentioned earlier, the two most significant areas that set these machines apart are performance and storage media.

NEC put its own 9.83-MHz V30 CPU in the UltraLite, and, consequently, its BYTE CPU index is higher at 0.93. Both machines simulate a hard disk drive in RAM, so disk-access times are similar, although the UltraLite's minimum SDD size is 1 megabyte. The MinisPort seems to have an edge in battery life, and its battery is user-replaceable, whereas the UltraLite's isn't.

Zenith went with something familiar when it chose the 2-inch floppy disk drive as the removable storage media. NEC chose battery-backed 256K-byte RAM and ROM cards. Both approaches seem to work well, although NEC's is more expensive: The cards cost $299 each. Both vendors must assure potential buyers of reliable supplies of each medium, since they are new.

The Mini Future

I like the MinisPort. I travel frequently and would welcome its company. Corporate America seems hungry for smaller, fully functional DOS laptops. The MinisPort fills that need, at least for those who can afford its price.

It could be better. A 2400-bps modem would be nice (Zenith says one is in the works), as would a better screen, more RAM for the SDD, a longer-lived battery, and a price tag to match its size. It should also lose a little weight. These improvements will come as laptop technology advances. In the meantime, the MinisPort makes a good travel companion. ■

It goes with the territory.

If your territory is field engineering, the GRiD portable computer can make your work more productive.

GRiD laptops make it easy for you to analyze data on-site for fast error identification, correction and reporting. You can retrieve information from remote databases while in the field, and share information via electronic mail.

GRiD laptops are lightweight, battery-powered and ruggedly built. The GRiDCase 1535 EXP has two, full-sized expansion slots, to let you plug in the board you need for communications and storage expansion. And the 1535's 386 processor gives you fast processing for complex calculations.

We help you find the right software too, whether it's an existing DOS or UNIX program, or customized package. We also offer after-sale consulting and responsive technical support. We'll give you a complete, systems solution that you just won't find from a hardware-only vendor or local retailer.

Why? Because at GRiD, we think that goes with the territory.

Give us a call at *800-222-GRiD.*

GRiD

First in Field Systems

Things aren't always what they seem.

Like our DesignView/1600 system

A system that Publish magazine called "a real knockout with the finest text display you'll find on any monitor — Mac or PC."

With all the features it has, it certainly looks like a pretty costly workstation quality display system.

Features like an impressive 1,600 x 1280 resolution. That's 109 x 121 pixels per inch!

■ **200-MHz performance**. To give you a steady and easy-to-read display, the DesignView/1600 has the highest data rate in its class - 200 MHz. And the high, 67 Hz screen refresh rate gives you a picture that's flicker-free and easy on your eyes.

■ **Two pages at a time**. With two full-sized pages on a 19-inch landscape display, the DesignView/1600 makes scrolling a thing of the past. That means it's faster and easier to create newsletters or magazines, ad or manuals, financial reports, books, and virtually anything else because the DesignView/1600 lets you see exactly what you'll print.

■ **Sharper fonts and images with grayscale**. The DesignView/1600 optional levels of gray enhance details that might go unnoticed on other displays. And delivers the sharpest image available on the market. We even bundle four proprietary gray-level DesignFonts (Times Roman, Helvetica, Symbol, and Courier) in a variety of sizes so you can see - and read - everything accurately before you print.

■ **The highest resolution for a PC**. The 1600 x 1280 pixel resolution of the DesignView/1600 display lets you read fine print (as small as 5-point type) while seeing two whole pages. **This is the same resolution as the Sun-4 workstation monitor**. The extraordinary display quality, combined with a clean, stable, flicker-free picture means your images are razor sharp, your alignments are perfect, and your drawings are pin-point precise. We even provide high resolution screen fonts so what you see is exactly what you get.

■ **Work with your existing software.** The DesignView/1600 display system offers extensive software driver support for most major applications, including Microsoft Windows 386 and 286, Aldus PageMaker, Microsoft Excel, Adobe Illustrator, CorelDraw, Designer, GEM, Ventura Publisher, AutoCAD, AutoShade, WordPerfect, WordStar, Publisher's Paintbrush, Lotus 1-2-3, Sun Microsystems' NeWS, Xenix, and many more. Plus, our flawless on-board Hercules emulation ensures that existing applications will run without modifications.

The DesignView/1600 offers many more features, yet it isn't as costly as it sounds. *This excellent display costs hundred of dollars less than our nearest competitor!*

We guarantee your complete satisfaction by offering...

■ **Fast service.** All orders are shipped within 24 hours. Specify optional Priority Service and your order is shipped the same day by overnight express.

■ **The best warranty on the market.** Our display products carry a three year warranty on all parts and labor.

■ **60 days money back guarantee**. If you are not completely satisfied with any of our display systems, for any reason whatsoever, just return it within 60 days and we will refund your money in full. **Plus** we'll pick up all freight charges.

The DesignView/1600 is part of our extensive family of display products for IBM compatible and Macintosh computers. We carry a complete line of high resolution color, monochrome and gray-scale systems.

In fact, the monitor you've got in mind is probably in our warehouse right now!

To get a better picture, call our toll free number for a copy of our free information package.

1-800-343-5532

Dealers Inquiries Welcome.

Elite Microsystems, Inc.

4201 Remo Crescent - Bensalem, PA 19020 - USA

Tel. (215) 639-1636   Fax: (215) 639-3420

World Radio History

THE GREAT POWER SPIKE

When the lights go out at Chaos Manor, it's a serious matter

Jim Ransom, my deputy chairman of the Advisory Council on National Space Policy, had just finished some updates to the SSX (Space Ship Experimental) briefing to go to the Defense Council, and we'd shut down the Mac IIx. We'd been using Microsoft's PowerPoint presentation software to make some changes to the briefing. Although there are other powerful programs, such as More II, I don't find all Macintosh software quite so easy to learn as it's generally advertised; when you've learned the quirks of a program that's good enough, it's sometimes best to stick with what you have. We started our spaceship briefings with PowerPoint, and we've never had enough time to learn anything else.

Anyway, we used the Mac IIx software shutdown procedure. When you do that on my Mac, a voice shouts, "Bring out yer dead!" after which the screen goes dark. We'd been using the Apple Scanner, which is attached to the Mac IIx, and the CD-ROM reader, but we'd taken no particular steps to shut those down. We'd also been using the LaserWriter IINTX, and we didn't turn it off.

I'd just poured a pair of brandies to celebrate the work we'd done when the lights went out.

Actually, it was worse than that. Not only did the lights go out, but they instantly came back on for a brief moment, and this time there were sparks and bright flashes all over the room. A light bulb exploded. There were more flashes outside. Then quiet, and darkness.

The only light was from the screen of the Northgate 386. It had been connected to a Clary PC-1.25k Onguard PC unin-terruptible power supply; and while other stuff fried, the UPS-protected Northgate didn't even glitch.

We got down the wall flashlights—I keep flashlights in my desk, but the main emergency light sources are Black & Decker Spotlighters connected to wall recharger units, two upstairs and two down—and went around turning off all the computers just in case the power came back on. By that time the high-pitched "No Power, Boss!" warning signal from the UPS was getting to me, so I shut down the Northgate—it had never noticed that anyone had a problem—and turned off the Clary UPS. We still had telephones, and I thought of plugging Big Cheetah and the modem into the UPS and getting on-line to BIX, but I didn't do it. We waited a few minutes, but it became obvious that the lights weren't coming back on, so we finished our brandy, Jim went home, and Roberta and I went to bed.

I woke about 4:00 a.m. to discover that some of the lights in the house were on, but some weren't and wouldn't go on. A main 30-amp fuse was blown, and when I replaced it the replacement blew instantly. I thought about what could do that and half-concluded that a power spike had shorted out the refrigerator. After all, Roberta had just that day replaced its vegetable crisper at a cost of $135 and a lot of her time; why not? But there was nothing to be done at 4:00 a.m.

Come morning we horsed the refrigerator out of its alcove, discovering about 2 inches of greasy dirt underneath—it's very difficult to pull the fridge out, and evidently we hadn't done it for several years—and unplugged it. Then we went through the house looking for anything else that might be plugged in—and lo!, in Roberta's office, there was an Isobar Power Isolator and Surge Protector. Her Kaypro 386 and Mannesmann Tally laser printer had been plugged into it. When I disconnected the Isobar from the wall, something inside it rattled.

We replaced the main fuses. No problem. Then we cautiously plugged in the refrigerator. It started up fine. I took the Isobar upstairs and used a multimeter to discover there was a dead short from the hot side of the plug to ground. No wonder it blew fuses.

After that, it was a matter of testing.

The first casualties were in the back room. My son Richard had been playing Earl Weaver Baseball on the Tandon 286 when the lights went out. Alas, the Tandon was plugged directly into the wall, no surge suppressor, and it was dead. So were the family room VCR and TV, both of which had been on when things happened.

Next were light bulbs. Fluorescents were all right, but every incandescent light bulb that had been on was dead.

"Some power failure," I said. Roberta called the Department of Water and Power to see what had happened. The chap who answered said it had been amusing to listen to the stories at first, but now it sounded like one big whine: everyone had lost equipment. Some chap had managed to drive his car into a power pole, which fell, taking out a transformer. He offered to give us the telephone number of the poor fellow's insurance company.

In discussions with Joanne Dow ("jdow," the Amiga wizardess on BIX) and her friend Alan ("arog" on BIX), we decided that a 16K-volt AC line had dropped across one side of the 220-volt lines that supply the houses in my neighborhood. The result was one heck of a power surge.

So. Now we knew what happened. Next thing was to assess the damage.

First, Roberta's machine, printer, and USRobotics external modem, which had been plugged into the now dead-shorted Isobar, worked fine. When we took the Isobar apart, we discovered that every choke coil was discolored and several of the metal-oxide varistors (MOVs) had

continued

literally melted. My son Alex looked at it and clucked his tongue. "It died that others might live," he said. I've still got the Isobar; one of these days we'll bury it with military honors.

It deserves it. I bought that gadget back in 1977 at the behest of Dan MacLean, who insisted that all electronic equipment ought to have surge protection. Clearly he was right; alas, after he died I became slothful and neglected some of the gear. I sure wish I hadn't.

My upstairs suite in Chaos Manor has its own electrical supply box with circuit breakers rather than fuses. I found that three breakers had tripped. When I reset them, I noticed that my incandescent lights were gone, but the fluorescents were all right, and so was the pump for the tropical fish tank.

When I turned on the Clary UPS, there was no whine; it had power. The Northgate 386 connected to it was fine, too.

Big Cheetah, my main machine, had been plugged into a Compuguard surge suppressor that I bought from Priority One. That unit also supplied power to my USRobotics modem, a Maximum Storage APX-3200 WORM (write once, read many times) drive, an Amdek Laserdrive CD-ROM reader, and an Electrohome 19-inch high-resolution monitor. I had turned off the switch on the Compuguard while the lights were out; now I held my breath and turned it back on.

Big Cheetah came up fine. So did all his auxiliary equipment. No damage.

Next was the Macintosh, which was plugged into a Woods surge suppressor. The Mac had been shut down when the spike hit, and it came up with no problem. All its peripherals, such as the scanner, worked properly, too. Alas, not so the Priam SCSI 330-megabyte MacDisk, which was also plugged into the Woods suppressor and had been left on after the Mac was shut down. Inspection revealed that the Priam's 2-amp automobile-style cartridge fuse had blown so violently that there was metal plated all over the inside of the glass cylinder. Replacement of the fuse did no good. The Woods suppressor might as well not have been there.

In a panic I called Alex. After all, he's in the business of recovering data from zapped hard disks. He came right over. "Power supply, probably," he said, and proceeded to cannibalize the power supply from a spare external WORM drive box. In minutes he had the Priam up and running. It looks a bit odd in the old WORM box, but it works fine. Priam is getting us a new power supply, and I can still report that we've yet to lose a single byte of data from a Priam hard disk.

Then there was the Apple LaserWriter IINTX: dead as a doornail. I sure hope it's just the power supply. There's no fuse visible. Apple is sending me a replacement. Meanwhile, in the two days since we lost it, I've found just how much I do with it: not novels and articles, but letterheads, everything with graphics, and a lot of other stuff. I'll sure be glad to get it running again.

Anyway, to cut the story short: the power surge killed every unit of electronic equipment that was turned on and not plugged into a surge suppressor. It also burned out nine incandescent light bulbs and literally exploded two others; and it killed three surge suppressors, one of which, the Isobar, failed in a dead short, while two others (brand name unknown) simply died—they didn't blow fuses, but they no longer let power through at all. One of those protected the VCR and TV

that I keep up here in my part of Chaos Manor. Recall that the unprotected family TV and VCR were killed, so that was cheap enough protection.

Meanwhile, there was quite a lot of equipment plugged into Compuguard surge suppressors I had bought on sale from Priority One. Not one unit of any kind protected by a Compuguard was harmed in any way.

Joanne Dow and her friend Alan, who's a county building inspector and knows about building electrical systems, tell me I had better replace all the surge suppressors that lived through The Great Power Spike. The MOVs in those units may have been damaged in the process of protecting the equipment, and there's no simple way to test them. Of course, I can buy MOVs from Radio Shack for a buck or so each, and if I were so inclined I could pry apart all those Compuguard units and solder in new MOVs; but the fact is that I'm not going to do that. I do wish I had a simple way to test the surge suppressors—after all, I'm about to replace 10 of them at about $30 each, and it would be nice to know whether the expense is *really* needful—but in fact it's fairly cheap insurance.

Alan also tells me I had better replace all three of the circuit breakers that tripped. They undoubtedly arced over, and their ability to protect my circuits is now very much in question.

The morals of this story are simple: if you don't have surge suppressors on all your electronic equipment, including stereos, VCRs, and TVs, as well as your computers, then you're gambling. Look, here in southern California we almost never get real lightning storms. The Los Angeles Department of Water and Power, and Southern California Edison (which supplies power to the parts of the county outside the city), are very reliable, seldom have power failures, and nearly never have power spikes. My electric power is probably as clean and reliable as you'd find anywhere in the world.

So what? No one is safe from weird accidents like automobiles crashing into power poles. I now have to replace some $300 worth of surge suppressors, pay another $350 for repairs to equipment that wasn't protected, and we're without our TV and VCR for a week. The alternative is worse; it could have cost a *lot* more.

If your work is at all valuable, get a UPS. Not just any old UPS, but one rated powerful enough to keep your equipment going. Be sure to look into the power surge protection capabilities.

I don't know if power surges will damage a UPS. The Clary people are sending me a new unit to swap for the one I have; they want to see what it looks like inside after taking a hit like that. I'll let you know next month. Meanwhile, I've tested this one about 10 times by simply yanking the plug while the Northgate 386 was doing a big copy operation from floppy disk to hard disk. About half those tests were done after the Big Power Surge. Nothing at all happened during any test; the Northgate went right on about its business, totally unaware that someone was messing with its power. I've also tested the WORM drive on the UPS with the same result.

I have become a believer. From now on, all electronic equipment in Chaos Manor will have surge protection, and any computer doing a vital job will have a UPS. I do wonder why surge protection isn't routinely built into power supplies. The parts cost only a couple of dollars.

I sure don't have any trouble rating the

Circle 186 on Reader Service Card

World Radio History

THE NEW STANDARD FOR HIGH PERFORMANCE STATISTICAL SOFTWARE

COMPLETE STATISTICAL SYSTEM
WITH DATA BASE MANAGEMENT
AND GRAPHICS

A powerful, comprehensive, elegant, and super-fast statistical package for IBM (PC, AT, PS/2) and compatible computers. ■ The CSS optimized user interface with fast hierarchical menus incorporates elements of artificial intelligence; even complex analyses require only a few keystrokes (batch processing is also supported). ■ CSS features comprehensive, state of the art implementations of: *Basic statistics, Multi-way frequency tables, Nonparametric statistics, Exploratory data analysis with analytic graphs, Multiple regression methods, Time series analysis with modeling and forecasting (incl. full ARIMA), General ANOVA/ANCOVA/ MANOVA, Contrast analysis, Discriminant function analysis, Factor analysis, Principal components, Multidimensional scaling, Item analysis/Reliability, Log-linear analysis, Cluster analysis, Non-linear estimation, Logit/ Probit analysis, Canonical analysis, Survival and Failure Time analysis (Censored data), Quality Control analysis, and much more.* ■ All statistical procedures are integrated with fast data base management and instant, presentation quality graphics (over 100 types); full support for all mono and color graphics boards (incl. VGA) and over 100 plotters and printers (incl. the HP and Postscript standards). ■ All CSS screen output is displayed via customized Scrollsheets™ (i.e., dynamic, user controlled, multi-layered tables with cells expandable into pop-up windows); all numbers in a Scrollsheet™ can be instantly converted into a variety of presentation quality graphs; contents of different Scrollsheets™ can be instantly aggregated, combined, compared, plotted, printed, or saved. ■ The flexibility of the CSS input/ output is practically unlimited: CSS offers an intelligent interface (read/write) to all common file formats (Lotus, Symphony, dBII, dBIII +, DIF, SYLK, . . .) and special utilities to easily access data from incompatible programs; graphics can be saved in files compatible with desktop publishing programs (Aldus, Ventura). ■ CSS data files can be as large as your operating system (DOS) allows; OS 2 version coming soon. ■ CSS precision exceeds the standards of all common precision benchmarks. ■ *Technical note: The CSS user interface and all I/O were written in Assembler and bypass DOS; graphics and data management were written in Assembler and C; the computational algorithms were written in Assembler and optimized Fortran.* ■ $495 (plus $5 sh/h); 14-day money back guarantee.

Circle 236 on Reader Service Card

better than having them stored on a million floppy disks, but it's still hard to keep track of them, since they tend to drift downward into folders held inside other folders, and I never remember the names I've assigned. Searching for a particular image used to take a long time, and sometimes I didn't bother.

That's all changed now. The Curator takes care of them. This program catalogs and characterizes Macintosh graphics files. What you do is set it up and then invoke a program called The Curator's Assistant. This program hunts through your hard disk (or through a collection of floppy disks if that's what you have) and finds everything that it thinks might be in a graphics file format: PNTG, PICT, SIMA, EPSF, EPSP, TIFF, and PostScript TEXT. It can't manage some of the proprietary formats, but if you can manage to save in one of the Big Seven Standard formats listed above, you're in business. The Assistant will find them, look at them, and draw a small icon pretty well representative of the graphical content. Now you can browse through those icons and see which graphics file you want. Curator will find graphics

files, convert from one format to another, help you with printing, and in general act as an intelligent curator for your art files.

It ain't perfect. It doesn't understand gray scales. The Curator's Assistant doesn't tell you when it's done searching your hard disk; it just stops and leaves it to you to figure out that it's finished. There are some other glitches.

No matter. This is one of those programs you will soon find you can't do without. Nowadays when I want to find my graphics files, I call up the Curator and let it do the work. I sure wish I had something like this for a PC-DOS machine. Recommended.

Culture 1.0

This program is subtitled "The Hypermedia Guide to Western Civilization," and it's a time trap. What this program modestly attempts is to present the entire history of the world on seven disks (about 5 megabytes) of HyperCard stacks. There are some 1750 cards organized into 21 cultural grids that show what's going on in different countries at the same time, and about 200 graphics

images of works of art like Michelangelo's *David* and sketches of Lorenzo the Magnificent. Alas, there are no maps, which seems a rather odd omission.

It's difficult to evaluate something of this size. One blurb says that this program will "convert the Macintosh into an educational workstation." I'd agree with that. Totally. This would be a heck of a tool to use in preparing for examinations. I'll go further. For anyone motivated to learn history—whether out of simple curiosity or the desire to get a good grade—this is an invaluable resource.

Alas, it may not provide its own motivations. There are a number of essays, and they're all written in HyperCard style: terse, with maximum opportunity to show other buttons in boldface. That's the problem. Writing in HyperCardese isn't conducive to being interesting. There's little of the wit of Jacques Barzun, or the intriguing style of Fletcher Pratt. There are no grand sentences from Macaulay. The authors of Culture are clearly admirers of Jacob Burkhardt and rightly identify him as the discoverer of the importance of the Renaissance, but

World Radio History

CHAOS MANOR

they don't quote him.

There are organizational holes. Much of the material is in superficial form. There's a lot more on music and architecture than literature. Dante Alighieri gets one terse line in addition to his name and dates: "*Divine Comedy*, 1321, one of the first works in Italian (Tuscan dialect.)" You'd think he deserved more. Machiavelli is represented by a single possessive that reminds you that he was the author of *The Prince* but says very little else.

Although the program doesn't tell anything about Benvenuto Cellini—he gets the single word "autobiography"—it does have a bunch of gratuitous comments. We're told that Lord Acton, an English historian, had Savonarola, an Italian Renaissance religious reformer, in mind in his dictum "Power tends to corrupt. Absolute power corrupts absolutely," and that Oliver Cromwell should have studied the case of Savonarola. Now I'm a closet Royalist myself, but perhaps there ought to be a hint that there are differences of opinion about Cromwell. The historian Macaulay could say "Cromwell was no more; and those who had fled before him were forced to content themselves with the miserable satisfaction of digging up, hanging, quartering, and burning the remains of the greatest prince that has ever ruled England." Culture says, "After the Restoration of the monarchy he was disinterred and hung up on a gallows in 1661." I think I prefer Macaulay. Alas, Macaulay himself gets only one line.

In other words, Culture is sketchy.

It doesn't work as well as you'd like, either. The search feature is impossible. You can look for key words, but when it finds the first instance, the program stops looking. There's probably a way to make it go on to the next instance, but if there is, the instructions don't tell you, or worse, they tell you to do something that doesn't work. All of which is a pity, because Culture is a magnificent attempt at a project worth doing. It would take a CD-ROM to do it right. Perhaps someone will make one.

Until then, Culture will turn your Mac into an educational workstation, but you'll have to bring your own motivation.

Wordfind
The shareware of the month (a new feature I just instituted) is Wordfind, a program to help you solve word puzzles, crosswords, acrostics, cryptograms, and other word games. It's available from Castle Oaks Computer Services and runs on just about any MS-DOS machine. It's

World Radio History

pretty neat if you're into solving word puzzles.

Remote Keyboard
This is one of those gadgets that not everyone needs, but if you do need it, you'll want it a lot. Despite the name, it's not a keyboard; it's a gilhickie about the size of a TV remote control with 40 buttons. It comes with an infrared receiver that plugs into your computer's serial port, plus the software that makes the computer listen.

Once it's installed, you can do just about anything with Remote Keyboard that you could do with your regular keyboard; but you won't do it quickly, because doing hunt-and-peck typing on a four- by 10-button array with keys laid out in alphabetical order is darned near impossible. Of course, that's not what Remote Keyboard is for; what you do is use it to control your computer during a presentation in the same way that you'd

use a remote control to advance slides during a briefing. You can use PageUp and PageDown, Print Screen, and the rest of it. You can also set up various macros to be executed by Control or Alt keys. (You can't use both keys at once; unlike your regular keyboard, to get Control-C you'd press Control, release it, and then press C; ditto for Alt keys.)

The obvious use for this is in connection with a projection system; however, it would also work in a situation where you have several people crowded around a computer screen while the briefer stands in another part of the room. It can also be used to control a robot, and I understand one medical center is doing that.

Remote Keyboard works with just about any PCompatible, including my Zenith Portable. I won't use it often, but I'm glad to have it here, and I'll probably use it at the next meeting of the Advisory Council on Space Policy. It would be neat to have one for the Macintosh as well.

Join the Navy!
I did my military service in the Army, and I worked for the Air Force for a good part of my aerospace career, but my number-three son Phillip is a midshipman in the U.S. Navy. That probably explains my interest in naval war games. We get a lot of them.

Two of the most recent are submarine warfare simulations: EPYX's "Master's Collection" Sub Battle, which simulates World War II submarine warfare; and Electronic Arts' 688 Attack Sub, which is modern nuclear submarine warfare. The versions I have are the new Mac II EPYX Sub Simulator, which does a wonderful job of bringing Macintosh color graphics to an older (but fun) game, and the PC VGA version of 688.

Of the two, the EPYX simulator is a lot easier to "win," but the Electronic Arts 688 Attack Sub is more realistic. Both are easy to learn and have a realistic
continued

feel. As you'd suspect, there's a lot more action in the World War II simulation. Either one makes for a good way to waste an evening.

Strike Fleet, also from Electronic Arts, is a simulation of a modern surface battle group without submarines or carriers; you can command a single ship such as the USS *Stark* or a whole escort group in the Persian Gulf, a British ASW force off the Falkland Islands, or a large U.S. strike force off Iceland. I played this every day for more than a week, and the only reason I quit was that I got behind in my work. There are 10 scenarios, and the last few get really tough.

All three of these games are what you'd call "modern" computer games: lots of graphics and a great deal of player control over each unit. For example, in Strike Fleet you have to control each ship in your force; you're not only the fleet commodore, but the skipper of each ship, and for that matter, the weapons officer for each ship. While these are not really arcade-style games—you can pause them, and things don't move all that fast—there is a certain arcade flavor to them, although do understand that good strategy and tactics are more important than manual dexterity.

There's another kind of naval war game. Simulations Canada has a series of games ranging from the early days of World War II to Northern Fleet, an operations game set in the North Atlantic in 1990.

There are no fancy graphics to these games. Unlike Strike Fleet, which has a manual that could serve as a general introduction to modern weapons capabilities, Simulations Canada provides almost none of that; you're expected to know something about the systems you command. There aren't any control rooms or individual weapons commands, either. In Simulations Canada games you do what an admiral would normally do: issue orders to battle groups and get reports on what is known about your forces and those of the enemy.

The result is surprisingly realistic. I say surprisingly because the conventional wisdom in simulations is that you need fancy graphics and detailed unit reports; but in fact that's not realism at all. Generals and admirals aren't often required to smell the gunpowder. As John Keegan shows in *The Mask of Command*, most of that changed irrevocably in the period of the U.S. Civil War.

Anyway, the Simulations Canada games are different, because all you'll see is screen after screen of menus and lists and tables; but they're actually more

realistic, and to those with the proper temperament, no less enjoyable than games with "better" interfaces.

VGA

Video standards change. When IBM first came out with color, the screen resolution wasn't good enough for sustained text work. Then came EGA, which was good enough, but which was defectively designed. Now we're getting VGA, which is really pretty nifty.

There aren't too many programs that take advantage of VGA, so it's not always easy to tell just how good it is; indeed, I really discovered the difference when I ran Electronic Arts' 688 Attack Sub on the Northgate 386 (which has VGA and a Princeton monitor) and then transferred the game to Big Cheetah and the 19-inch Electrohome, which was running EGA.

The result was horrible. I'd previously thought EGA to be good enough; after seeing what you can do with VGA, I thought different.

However, when I put Video Seven's newest 16-bit VRAM VGA in Big Cheetah, the output was a mess. I knew that it wasn't the monitor's fault because I was testing the system with the Zenith Flat Technology Monitor, and that worked fine with the VRAM in the Zenith.

It turns out that the Cheetah's motherboard is a bit too fast for most video boards; but Cheetah will send you new programmable-array-logic chips that will fix the problem.

Meanwhile, I tested Big Cheetah with the Video Seven VEGA VGA, which is an 8-bit video board. Although not as fast as the 16-bit VRAM, the VEGA is certainly faster than EGA, and of course the resolution is better. The result is absolutely gorgeous on the Electrohome monitor. Getting it running on the Electrohome requires a special cable: the monitor only has 9-pin input, and VGA boards universally have 15-pin output. I've tried about 10 different commercial cables, including a set made up by Candy Cable of San Diego, and none work; the only one that will work came direct from Electrohome. Once you have the right cable, though, an Electrohome with VGA is something to see.

There is one problem: VGA uses more memory than EGA. Since that memory is up in the area between 640K bytes and 1 megabyte, it wouldn't matter, except that we're using Quarterdeck's QEMM to load stuff like buffers, the mouse driver, and the WORM driver up into that area. We can still do that, but we don't have quite so much of that high memory available with VGA, which means that we

have to reduce the size of our DESQview windows. So it goes.

Winding Down

My desk is still covered with stuff, but I'm out of time and space. The book of the month is *What Do You Care What Other People Think* (Norton, 1988), which, with Dick Feynman's previous *Surely You're Joking, Mr. Feynman*, make up the extraordinary autobiography of an extraordinary man. If you like those, get his *QED*, which is a readable explanation of what quantum electrodynamics is all about, and his *Character of Physical Law*, a short and highly readable work on the philosophy of science. I've just re-read all those, and I'm a bit sad because there are so many things I never got a chance to discuss with him; but I'm sure glad to have known him.

The computer book of the month is Jeff Dunteman's *Complete Turbo Pascal* (third edition; Scott, Foresman, 1989). This is one of the best introductions to Pascal ever done; it's organized differently from other language books. If you've never read another book on programming, try this; you may like it, and you'll at least learn something of what programming is all about. Of course Dunteman doesn't cover the absolutely latest version of Turbo Pascal; but that's all right. There's plenty to be learned before you try dealing with *objects*.

The programs of the month are Turbo Pascal 5.5 and Microsoft Quick Pascal. Both have objects, the new programming fad that may well deserve all the attention it's getting. If I had to choose one and only one, I'd go with Turbo Pascal, since it's built up from a mature and stable compiler developed in-house, while Quick Pascal was bought from outsiders and is in its first model year; but I'll know a lot more about that next month.

Meanwhile, I'm off to Globe, Arizona, and thence to Fort Apache, where with luck no one will find me; if I don't get *Wrath of God* done, they're going to repossess my house. ∎

Jerry Pournelle holds a doctorate in psychology and is a science fiction writer who also earns a comfortable living writing about computers present and future. Jerry welcomes readers' comments and opinions. Send a self-addressed, stamped envelope to Jerry Pournelle, c/o BYTE, One Phoenix Mill Lane, Peterborough, NH 03458. Please put your address on the letter as well as on the envelope. Due to the high volume of letters, Jerry cannot guarantee a personal reply. You can also contact him on BIX as "jerryp."
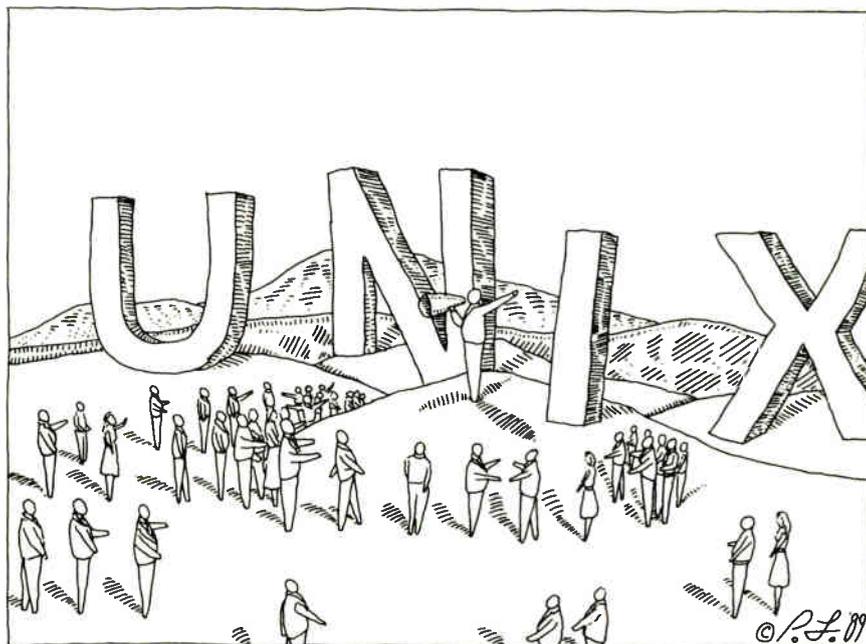
A CALM APPROACH TO UNIX

The average Unix user never has to worry about, let alone learn, many of the system services and nuances

Editor's note: *David Fiedler has written about Unix for BYTE many times. Now that Unix has established itself in the vocabulary (if not the office) of the majority of the computing community, we've established David as our Unix columnist. This is the first installment.*

David is the editor and publisher of the Unix newsletter Unique, *which he started on his kitchen table in 1981. He was also cofounder of the magazines* The C Journal *(now* The C User's Journal*) and* Unix Review.

With Bruce Hunter, he coauthored the best-selling book Unix System Administration, *the first book to cover this important subject. Its success led to the recent launch of* Root, *their journal of Unix and Xenix system administration. David has been a consultant to AT&T, ITT, CBS, and Sandoz Pharmaceuticals, among others, and has been in charge of software development efforts at many large companies.*

My prized first issue of BYTE (September 1975) contains such articles as "Which Microprocessor for You?" (your choice of the 8008, 8080, or IMP-16) and "Recycling Used ICs" (how to use a blowtorch to remove chips from printed circuit boards). That same issue also had an advertisement from Processor Technology for a 3P+S I/O board for Altair compatibles that would "fully interface two TV Typewriters with keyboards and a modem or teletype at the same time!" This board even let the pe-

ripherals talk at 9600 bps over the serial port. All this was quite advanced for the time. The only problem was that the software of the day couldn't possibly have supported simultaneous use of all those terminals.

I'll leap forward to the present, where—except for a few proprietary multiuser PC-DOS-like systems and special background print spoolers and communications programs—most personal computers are limited to doing a single thing for a single user at a time. In other words, today's microcomputers can also have a number of serial ports, but still can't use more than one at a time!

But with all the hardware advances in personal computers since they were first designed, today's microcomputer users have more power at their command than the users of many minicomputers of 20 years ago. The machines are now being severely underutilized. So it makes economic sense to look at ways of increasing personal productivity on computers,

whether by sharing physical machines or by enabling one computer to do a lot more. That's what the idea of multitasking and multiuser operating systems is all about.

Enter Unix

At the time that BYTE's first issue was published, the Unix operating system was already six years old—about the same age MS-DOS is now. Unix has undergone many changes—not all for the better, perhaps—in its 20 years.

Just for the record, I'll list a few important features of Unix:

• It is written in C and is portable to other architectures.
• It is multiuser and multitasking.
• It has a hierarchical file system with mountable disk volumes.
• It has file redirection and pipes.
• It is ready for communications: local- and wide-area networking.

continued

Debug Your Program BEFORE You Write It!

———— ◆ ————

"Deep bugs", the kind that show up after you ship the product, are usually the result of logic flaws. Such bugs include redundancies, contradictions, unique conditions without specified actions, etc.

Logic Gem, a proven programmer's tool, helps you avoid these bugs in three ways:

◆ Catches logical errors before you code the program. Saves hours of debugging time.

◆ Automatically produces written documentation of your logic, which insures good communications between designer and coder. (And helps remind you of the logic from one work session to the next.)

◆ Automatically generates flawless code for the "guts" of your program... in C, BASIC, Pascal, FORTRAN, and dBASE.

Logic Gem works with whatever compiler you are using. The only change: with **Logic Gem** you catch and correct the logic bugs before you write the program.

Avoid hours and days of needless debugging time. Call **1-800-722-7853** now for details about **Logic Gem**. Or put **Logic Gem** to work for you immediately: Order a copy (it costs only $195) and use it for 90 days at our risk. You can return it for any reason within 90 days for a complete refund.

———— ◆ ————

Sterling Castle, Inc.
702 Washington St. #174
Marina del Rey, CA 90292
1-800-722-7853
1-800-323-6406 (in CA)
1-213-306-3020
1-213-821-8122 FAX

- It has many useful standard and optional utilities.
- It can handle a wide variety of devices in an identical manner.
- It has many software development tools built in.
- It has an E-mail system.
- It has a true print spooler system.

Most people who say things like, "I don't think Unix is any better than xxx" are ignoring the importance of the first two features. It is uncommon that an operating system is portable across machines with different architectures and from different manufacturers. And to compare a system with the size and complexity of Unix to a system that can run only one thing at a time is pointless.

The problem most users have when faced with the task of "learning Unix" (or learning Xenix; for all practical purposes, they are now equivalent) is that it's *big*. Unfortunately, some companies have promoted Unix to microcomputer users by telling them that Unix is a kind of large DOS. Then the users encounter a meter-wide set of manuals and command names that sound like extinct animal species. They run screaming for the nearest exit. Unix gets some more bad press.

Perhaps the best approach to Unix is a calm one. Unix is a *real* operating system, not just a glorified program loader. Most people get along fine in DOS, even

though DOS has many commands with unusual syntax (I assure you, pressing the F3 key to repeat a command is nonintuitive). The average Unix user never has to worry about (let alone learn) many of the system services and nuances. Unix was developed in the days when a teletype was the standard input device, and anyone who has ever used an ASR-33 knows you don't want to type any more characters than necessary. So command names tend to be short (vowels are the first to go). In the interest of harmony and mutual understanding, therefore, table 1 presents a cross-reference of common DOS and Unix commands. This table is all you need to get started in Unix. Not really so bad, is it?

You'll notice that many Unix command names are the same as in DOS. Perhaps that should be written the other way: Quite a few Unix features (such as hierarchical directories, redirection, and pipes) were used as "role models" when DOS was being designed. It's just that DOS got the slash backward.

Let's Get Graphical

Macintosh users aren't being ignored here, but they have a much different user interface than either DOS or Unix, and I'm not particularly good at drawing pictures. The Macintosh *has* made a great contribution to computing: graphical *continued*

Table 1: *Common DOS commands, their Unix equivalents, and an English explanation. (Note that my definition of* dd *is facetious.)*

DOS	Unix	English
backup	tar	Tape ARchiver
cd	cd	Change Directory
chkdsk	fsck	File System ChecK
cls	clear	Clear screen
compare	cmp	CoMPare two files
copy	cp	CoPy a file
date	date	Set or show the date and time
del	rm	ReMove file
dir	ls	LiSt directory contents
dir /w	ls -C	LiSt directory in Columns
diskcopy	dd	DarneD if I know what it stands for
erase	rm	ReMove file
find	grep	Global Regular Expression Print
format	format	Format disk
join	mount	Mount disk or partition on file
system label	labelit	Label file-system volume
mkdir	mkdir	MaKe DIRectory
mode	stty	Set TeleTYpe characteristics
more	more	Show file a page at a time
print	lp	Line Printer
rename	mv	MoVe file to new name
rmdir	rmdir	ReMove DIRectory
set	set	Show environment variables
sort	sort	Sort
type	cat	conCATenate (can be used for either)

MKS Toolkit™

UNIX™ Tools on DOS or OS/2

Programming today means you must work within more than one environment. A diverse range of hardware is now a fact of life. With the MKS Toolkit, you can enjoy the best of DOS or OS/2 and UNIX environments. The MKS Toolkit offers both experts and novices the purest form of UNIX utilities that the DOS or OS/2 environment allows.

Reduce Keyboard Shock

With our proprietary code, the MKS Toolkit offers you more than 140 UNIX System V.3-compatible tools for DOS or OS/2. With the MKS Toolkit, your computer or clone becomes a comfortable environment for shells, string matching, editing, file manipulation, and more. Productivity increases because all the familiar commands are at your fingertips.

Site Licenses

The MKS Toolkit reflects its users' needs. Organizations such as AT&T, H-P, ITT, and NCR - all heavily committed to the UNIX system - use the MKS Toolkit to create a standard operating environment. Universities, including UCLA, use the MKS Toolkit to enrich personal research computing environments and double the bandwidth of their PC teaching labs. The National Institute of Standards and Technology fulfills diverse needs by using the MKS Toolkit as standard operating environment for experts and as a POSIX-conforming training tool for neophytes.

Interconnectivity

The MKS Toolkit provides two types of valuable interconnectivity. First, it interacts well on most standard PC and PS/2 networks. Combined with Novell Netware™, the most popular LAN for PCs, the MKS Toolkit creates a UNIX time sharing system in DOS or OS/2 organizations. UNIX shops can now hook up all their PCs using PC-NFS™ and the MKS Toolkit, enabling you to use a PC as a UNIX workstation and off-load your mini or mainframe machine. The second level of interconnectivity is created by the MKS Toolkit's ability to recognize common UNIX file formats on DOS or OS/2 and to make DOS or OS/2 file formats available on UNIX systems.

POSIX-Conforming Tools

MKS is an active participant on the POSIX 1003 standards committee. This involvement reflects MKS' commitment to tracking the shells and utilities standard to the fullest extent possible under DOS or OS/2. Apart from multitasking and constraints on file names under DOS or OS/2, the MKS Toolkit follows the POSIX standard. MKS achieves this by building the underlying POSIX system on DOS or OS/2 before moving utilities.

Cost-effective Learning Tool

If your organization is committed to moving into the UNIX environment, then the MKS Toolkit is the perfect learning path. DOS or OS/2 users retain the familiar world of their PC keyboard and programs and move effortlessly to a UNIX environment on their desktop. Exposure to new commands and functionality now becomes an integral part of the novice's working day. UNIX solutions are easily available and the DOS or OS/2 world is but a keystroke away.

MKS Programming Platform

The MKS Toolkit is the vital core of the programming platform created by MKS software. In addition to the MKS Toolkit, it is now possible to have:

- MKS RCS™ (Revision Control System)
- MKS Make™ (automated program builder)
- MKS LEX&YACC™ (compiler construction tools)
- MKS SQPS™ (enhanced Documentor's Workbench™)

Addictive Software!

The MKS Toolkit offers you power and diversity. Here is a complete list of commands you receive in the package:

alias	ed	let	sort
awk	env	line	split
banner	eval	login	strings
basename	ex	ls	strip
bdiff	exec	mkdir	sum
break	exit	more	switch
c	expand	mv	sync
cal	export	nl	tail
case	expr	nm	tee
cat	false	od	test
cd	fc	pack	[
chdir	fg	passwd	time
chmod	file	paste	times
cmp	find	pcat	touch
: (colon)	fmt	pg	tr
comm	fold	pr	trap
compress	for	print	true
continue	function	prof	tty
cp	getopt	ps	type
cpio	glob	pwd	typeset
crypt	egrep	r	ulimit
ctags	fgrep	read	unalias
cut	grep	readonly	uname
date	gres	return	uncompress
dd	hash	rev	unexpand
deroff	head	rm	uniq
dev	help	rmdir	unpack
df	history	rsh	unset
diff	if	sed	unstrip
diff3	init	set	until
diffh	integer	sh	vi
dirname	jobs	shift	wc
. (dot)	join	size	whence
du	kill	sleep	which
echo	lc	spell	while
			who

No wonder our users call it addictive software!

System Requirements: The MKS Toolkit works on IBM PC, XT, AT, PS/2 and compatible machines under DOS 2.1 and higher or under OS/2. A hard disk is recommended for improved performance and convenience.

Order Information

Price: $199 for MKS Toolkit, $495 for OS/2.
30-day money-back guarantee

1-800-265-2797 (continental U.S. only)
1-519-884-2251 (outside continental U.S.)
1-519-884-8861 (FAX)

Ask for Department BY

VISA *MasterCard*

MKS

35 King Street North
Waterloo, Ontario N2J 2W9 Canada

MKS, MKS Toolkit, MKS RCS, MKS Make, MKS LEX&YACC, and MKS SQPS are trademarks of Mortice Kern Systems Inc. Other trademarks have been cited and MKS acknowledges the:

Circle 173 on Reader Service Card

user interfaces (sometimes abbreviated GUI and pronounced the same as the word that implies ice cream dripping on your shirt—see "A Guide to GUIs" by Frank Hayes and Nick Baran, July BYTE). AT&T has announced that Open Look will be the standard interface for Unix in its next release (Unix System V release 4, or simply SVR4). It is clear that graphical interfaces are here to stay.

To anyone who has seen Open Look (and most other graphical interfaces) in action, it is also clear that a great deal of computing power will be necessary to support this kind of interface. It's a real dragon-and-egg problem: Are graphics becoming popular because we finally have processors powerful enough to support them? Or are the processors being developed because they're needed for today's graphics overhead?

Personally, I find multiple windows on one screen distracting. Yes, you have the ability to cut and paste between them, and the simultaneous display of text and graphics. But for working on several things at once, I prefer the approach popularized in The Santa Cruz Operation's (SCO) Xenix: A function-key combination switches you to a different "virtual screen," replacing the original completely. This helps me to switch context mentally. It also gives me much higher performance on my hardware. No CPU time is wasted on something I don't need.

What About Networking?

As many users think of it, networking refers to a spider's web of cables attaching personal computers to each other and to "server machines." The servers are essentially multiuser computers whose main purpose is to send files to the personal computers. As generally implemented, personal computer networks of today are limited to the basics: file transfer and E-mail.

Unix systems have had the basics built in for many years, by way of the UUCP (for Unix-to-Unix copy) subsystem. Today, UUCP is known as the Basic Networking Utilities, and it's still included in every Unix or Xenix system sold. Using UUCP as a base, you can set up complex processes such as automatic file servers, E-mail "answering machines," and transparent remote printing. None of this needs hardware that's any more high-tech than an auto-dialing modem. And, of course, there is the store-and-forward worldwide UUCP-Net E-mail network, with perhaps 1.5 million mailboxes, and the distributed Usenet BBS, NetNews. (See "The Unix Connection" by Ben Smith, May BYTE.)

Networking in the larger sense implies much more. The Network File System (NFS) and Remote File Sharing (RFS) capabilities, generally implemented via Ethernet, allow multiple machines to combine their file systems as if they were all on one large computer. Users can move around in the file system, reading and writing to files, unaware and unconcerned that they are actually accessing files on machines across the hall, across the street, or even across the country. This has led to the growth of LANs with many connected diskless workstations that use a central file server to hold material on disk....Did you say that sounds like personal computer LANs? It does— but with Unix workstations, the sharing of files is transparent, so there's less need to copy whole files back and forth. The net result (pun not intended, but noted) is less traffic on the network and less special software that must be added (and learned!) to use the network.

A Breather

Here's a preview of the future of this column. For the first few articles, I'll be concentrating on Unix on microcomputers: Why would you want to bother with Unix on personal computers? I'll discuss the choices and trade-offs you'll be confronted with, once you've made the big step. Will you ever be able to go back to DOS? Will you ever have to? (Or want to?) And general Unix topics: How you can get some of the public domain Unix software that you're always hearing about; why Unix might be useful even if you're not a programmer; and some drawbacks to Unix (nothing's perfect, after all). And of course, I'll discuss how you can learn some of the more involved Unix commands, utilities, and languages so that you can "increase your personal productivity," too.

Meanwhile, I'll be waiting to receive some mail from you. Tell me what you want to read about in future articles. In general, the idea is to cover both hardware and software as it relates to Unix and give you enough detail to keep you challenged, but not get so esoteric that your eyes cross and you turn the page. Everything else is wide open. ∎

David Fiedler is editor and publisher of the Unix newsletters Unique *and* Root *and coauthor of the book* Unix System Administration. *He can be reached on BIX as "fiedler."*

Your questions and comments are welcome. Write to: Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.

World Radio History

FINALLY. A debugging tool tough enough to handle the DOS Nasties.

New Version 2.0

Soft-ICE™

Nasty over-write? No sweat!

Soft-ICE memory range break points help you track down memory over-write problems whether you are doing the over-writing or another program is over-writing you.

Hung program? No problem!

When the system hangs, you now have hope. With Soft-ICE you can break out of hung programs no matter how bad the system has been trashed. And with Soft-ICE's back trace ranges you can re-play the instructions that led up to the crash.

Program too large? Not with Soft-ICE!

Soft-ICE runs entirely in extended memory. This means you can debug even the largest DOS programs. And since your program runs at the same address whether Soft-ICE is loaded or not you can find those subtle bugs that change when the starting address of your code changes.

System debugging? Soft-ICE is a natural!

Soft-ICE is ideal for full source level debugging of TSRs, interrupt service routines, self booting programs, DOS loadable device drivers, real-time kernels, non-DOS O'Ss and ROMs. Soft-ICE can even debug within DOS & BIOS.

How Soft-ICE Works

Soft-ICE uses the power of the 80386 to surround your program in a virtual machine. This gives you complete control of the DOS environment, while Soft-ICE runs safely in protected mode. Soft-ICE uses the 80386 to provide real-time break points on memory locations, memory ranges, execution, I/O ports, hardware & software interrupts. With Soft-ICE you get all the speed and power of a hardware-assisted debugger at a software price.

Don't want to switch debuggers?

You don't have to!

Soft-ICE can run stand-alone or it can add its powerful break points to the debugger you already use. Use your favorite debugger until you require Soft-ICE. Simply pop up the Soft-ICE window to set powerful real-time break points. When a break point is reached, your debugger will be activated automatically.

MagicCV with Soft-ICE

Using Soft-ICE with CodeView gives you the features necessary for professional level systems debugging. MagicCV and Soft-ICE can work in concert with CodeView to provide the most powerful debugging platform you will find anywhere.

> "These may be the only two products I've seen in the last two or three years that exceeded my wildest expectations for power, compatibility and ease-of-use."
> —Paul Mace
> Paul Mace Software

Soft-ICE	$386
MagicCV	$199
MagicCV for Windows	$199
Buy Soft-ICE & MagicCV(W)	—Save $86.
Buy MagicCV and MagicCVW	—Save $100.
Buy All 3	—Save $186.

30 day money-back guarantee
Visa, MasterCard and AmEx accepted

New Soft-ICE 2.0 features
- Back Trace Ranges
- Symbolic & Source level debugging
- EMS 4.0 support with special EMS debugging commands
- Windowed user interface

Nu-Mega TECHNOLOGIES

CALL TODAY (603) 888-2386
or FAX (603) 888-2465

RUN CODEVIEW IN 8K ™
MagicCV

CodeView is a great integrated debugger, but it uses over 200K of conventional memory. MagicCV uses advanced features of the 80386 to load CodeView and symbols in extended memory. This allows MagicCV to run CodeView in less than 8K of conventional memory on your 80386 PC.

NEW—Version 2.0 includes EMS 4.0 driver.
Attention Windows Developers!
Version available for CVW.

P.O. BOX 7607 ■ NASHUA, NH ■ 03060-7607

NEITHER SNOW, NOR CHICAGO...

Spring Comdex was no picnic, but there was good news for users

N either the snow nor the relocation to Chicago could hide the good news that Spring Comdex had for business this year. While Comdex may have suffered a bit from those factors, it was a significant show for business users of personal computers. This was the show where business found that the computers of the future would get both better and cheaper; where the flirtation between computers and the fax process became lust; and where the "Year of the LAN" became the great expectation of connectivity.

The move to things better, faster, and cheaper was shown no more clearly than in the introduction of the Intel 80486. This is the processor that will lead business users to the world of a mainframe on a desk. That mainframe on a desk will communicate with ever-more-powerful laptops and peripherals through LANs without traditional cards and even LANs without wires.

It was impossible to tell what was happening in the Macintosh world. Apple wasn't there, and Macdex crashed and burned. Apple's absence was the PC world's gain, however, and the 80486 gained center stage.

80486 Fever
The 80486 is important because it allows a high-performance computer to be built with fewer components. This chip incorporates into its design many functions that were formerly done by support chips. Thus, you will no longer need separate components for the math coprocessor or for caching. The processor will have these functions built in.

The reduction in components will allow the 80486 to be designed to run much faster than did earlier processors. In addition, a computer using the 80486 can be built at a lower cost than can a comparable one with an 80386. Cheetah's Gene Sumrall was one of the first to point this out to me. He was also one of the first to show a motherboard that would support the 80486. Cheetah had designed its new board so that it would take a daughterboard for the CPU. This means that the company can offer the same basic board for an entire product line, changing only the daughterboard that supports the processor.

A few manufacturers, including IBM and Zenith, promised to have machines based on the 80486 by the year's end. Zenith's Andy Czernek said that his company's computer would use the Extended Industry Standard Architecture bus. If so, Zenith would be one of the first companies to introduce an EISA machine. According to Czernek, it will be available at the end of the year.

The fact that the 80486 will bring business users machines that are faster and cheaper is good news. It's likely, of course, that the first prices to drop will be those of the 80386 and systems that are based on it. That's even better news.

More Speed
Zenith was one of the first manufacturers to announce that it had begun shipping a 33-MHz 80386-based computer, the Z-386/33. There were others as well, including Compaq and Everex. We had been expecting 33-MHz machines for about six months—ever since Fall Comdex—but they became available only when Intel began shipping the chips in late March. A few companies had previously built systems that ran at this speed, but those machines used components designed for 25 MHz and simply run beyond their design speed.

The advent of commercially available 33-MHz machines means a great deal to

ILLUSTRATION: DAN REED © 1989

companies that use their machines for CAD or desktop publishing. There, the improvement in speed will more than offset the higher cost through gains in productivity.

A Few Pertinent Fax

A year ago, Comdex attendees realized that computers and fax were becoming a team. Fax cards were everywhere. The trend continues, but the bare fact is that fax is finally becoming integrated well enough to be useful. In addition, it is moving out of the desktop PC and into areas where it makes more sense to have a fax interface.

In my June column, I looked at some of the earlier fax cards, as well as a stand-alone fax machine made by Murata, a long-established manufacturer of stand-alone fax systems. At Comdex, Murata introduced its F-50 network fax server. This is a complete fax machine, including scanner and printer, that plugs into a network via a workstation. You can send faxes through the network, and you can scan them as you normally would. The F-50 contains its own processor and memory, so the conversion from a file to a fax image occurs inside the F-50 itself. This reduces the load on the network and the file-server or workstation disks.

The F-50 is not the first network fax server, but it seems to be the best thought out. There are times, after all, when you need to send something that is already on paper, and creating an image so that you can send it using a fax card can be cumbersome in the extreme. Likewise, there are times when you simply want to leave the fax machine turned on while everything else is shut off for the weekend. The F-50 will let you accomplish this.

At the other end of the spectrum is a new card from Holmes Microsystems that contains a combination 9600-bps fax and a 2400-bps modem. The FAX'EM card is about 2 inches square and contains only a few surface-mount chips, yet it's fully functional. It fits inside the expansion slot on Toshiba and Zenith laptop computers and costs about the same as competing full-size fax cards. Holmes also introduced a combination fax printer and scanner called PFIDO that attaches to a laptop computer. The entire machine is about 9 inches long and 1 inch square.

LAN Sakes

Clearly, the Year of the LAN has happened. A year ago, LANs were still something mysterious. By Fall Comdex, they were an accepted part of the computing environment. This year, they're

part of the scenery. In fact, they're *everywhere*. Still, there is some wonder remaining in the land of the LANs.

One of the problems with the office LAN has been that you needed a computer with an expansion slot to use the LAN properly; portables need not apply. While you could always use a LAN with a modem or a serial interface, this is difficult and usually too slow to work well with most personal computer software.

Xircom has solved that problem with its Pocket Ethernet Adapter. This device, about the size of a pack of cigarettes, plugs into the parallel port on any IBM PC-compatible computer. The adapter is available to support either thick- or thin-wire Ethernet. Included with the adapter are drivers for Novell NetWare.

The Pocket Ethernet Adapter is designed for use with laptops, but it will work well with computers that otherwise have limitations in the number of slots available. This means that you can buy one of those Zenith EaZy PCs being sold on the cable TV shopping channels and use it for a Novell workstation.

At $695, the Pocket Ethernet Adapter is a bit more expensive than some other

Ethernet cards, but not by much. It's certainly worth the price if it can give you access to your office LAN where you didn't have it before.

Of course, at times the problem isn't the network interface card, but rather the cables that accompany a LAN. O'Neill Communications has found a way to eliminate that part of a LAN by using radio instead of cables. O'Neill calls the result a LAWN (local-area wireless network). The system uses spread-spectrum packet switching, and it is a little slower than traditional LANs. On the other hand, it works better without cables than do traditional LANs.

LAWN has a great deal in common with the printer servers that I discussed in my December column. It attaches to the computer's serial port and is used most effectively for printer sharing, although it will support E-mail and file sharing. As I'm writing this, LAWN is undergoing FCC certification, but it should be available by late summer.

English Inroads
Comdex always brings exhibitors from all over the world, and this year was no

exception. Several companies from the U.K. were part of a government-sponsored group. One company, Penny and Giles Computer Products, has produced a trackball that really will take the place of a mouse. Normally, a trackball requires two hands to operate (at least it does for me), which offsets the advantage of using little desk space.

You operate TrackerMouse with one hand. Pressable areas on the sides of the device take the place of mouse buttons. The trackball protrudes through the top and bottom of the device, so that you can operate it without having it planted on the desktop.

TrackerMouse includes a solar calculator on its top. I suppose that this is based on an assumption that anyone who has a desk sufficiently messy to preclude the use of a mouse is also likely to lose a calculator. Probably a safe assumption.

[Editor's note: *For more information on TrackerMouse, see Computing at Chaos Manor in the July BYTE.*]

Also from the U.K. is the Datapath video board, an extremely high-resolution board designed for CAD and for desktop publishing. The Datapath video board supports resolutions of up to 1600 by 1280 pixels on the IBM PC (Model Q-PC) and the PS/2 (Model Q-MCAX), and it's extremely fast. I watched one board running at 1280- by 1024-pixel resolution redraw Autodesk's sample drawing of St. Paul's Cathedral in 1½ seconds. I don't think I've ever seen it done faster.

Modern Maturity
Spring Comdex this year was quiet. This was partly because Chicago's huge McCormick Place convention center swallowed the crowds more easily than do convention centers in Atlanta or Las Vegas. I suspect that the timing also had something to do with it. Not everybody likes Comdex in the snow.

Finally, the hype level seemed to be down a little. Perhaps that means that we are more sure of ourselves—a more mature industry. ∎

Wayne Rash Jr. is a contributing editor for BYTE and a member of the professional staff of American Management Systems, Inc. (Arlington, VA). He consults with the federal government on microcomputers and communications. You can contact him on BIX as "waynerash," or in the to.wayne conference.

Your questions and comments are welcome. Write to: Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.

THE WAY OF THINGS CONSIDERED

There are many ways to accomplish something, but only a few of them are right

Can computers spawn ideologies? Can computer companies tread successfully on ground formerly held by philosophers and politicians? Can a corporate culture possibly exhibit a strong ethical and moral component?

If you believe Guy Kawasaki, the answer is a resounding yes. Kawasaki, the president of software publisher Acius, has written a book about an important aspect of Apple's corporate culture: *The Macintosh Way* (Scott, Foresman, 1989, $19.95). Although I was quite skeptical about any company (especially a high-tech one) spawning a way of doing business that has a strong ethical and moral component, after reading this book, I've come around full circle.

Although the book will likely be bought because it offers a frank discussion of Kawasaki's years at Apple and his own exposure to the Macintosh way of doing business, that's not the best reason to buy it. *The Macintosh Way* offers an important glimpse at how a corporate culture is created and spread, how it can be corrupted, and how you can take advantage of it long after the original product on which it was based is history. Reading it will give you a very personal and insightful account of how to do business in today's Mac software market.

The book makes a strong case for two simple precepts: doing the right things, and doing them the right way. I'm amazed at how many people in the computer business fail to look beyond the end of their respective bottom lines. As long as sales curves rise and profits are made, they're satisfied. That's really too bad.

Among all the growing industries in this country, the computer companies should be setting the trends for corporate morality and ethical conduct, with their emphasis on empowering individuals with new and more powerful computing tools. Sadly, this isn't the case. Moral shortsightedness and ethical ignorance seem to run rampant in some computer companies, and corporate attorneys are left to find legal solutions for resulting problems. It's time that more companies paid attention to the lessons that *The Macintosh Way* teaches, rather than hiding behind a "well, that doesn't really apply to us" attitude.

Besides its important lessons on doing the right things in the right way, *The Macintosh Way* also gives an insider's view of how Apple developed the way it has developed and how one entrepreneur decided to leave the relative safety of its corporate culture to run with a software idea he thought important.

In fact, after rereading the book last night, I think that its lessons go way beyond Apple and the Mac.

Application Development Standards

In the past few months, I've commented on what I think Apple needs to do to extend the life span of the Finder and the Mac operating system well into the next decade. Although I'm sure that Apple is not waiting with bated breath to hear my further thoughts on the subject, one segment of the Mac interface deserves Apple's special attention. The applications that Apple creates (e.g., HyperCard) and the extensions it makes to the Finder, the MultiFinder, and the Mac operating system serve as the de facto and de jure standards that vendors follow when building their own applications. This situation bears a close examination.

Conventional wisdom dictates that every Mac application should follow the standards set by those first two Mac applications: MacWrite and MacPaint.

continued

That means that the first two menu-bar items are invariably File and Edit. That also means that the Quit command resides in the File menu, and it can be activated by a Command-Q key combination. Furthermore, Cut, Copy, Paste, Clear, and Select All reside in the Edit menu. But is this enough to ensure a good Mac application interface?

I don't think so. Of course, Apple has published its application guidelines in the multivolume *Inside Macintosh* series, published by Addison-Wesley. Other Apple employees have extended the definition of what constitutes a good Mac application interface in a variety of books published since 1984. Those efforts are all fine, and I don't have a problem with them. But I do have a problem with how these ideas are going to be extended on the one hand and controlled on the other, as the Finder and the Mac operating system move toward the 1990s.

As I'm writing this, Apple has announced its next-generation System software, version 7.0. At May's Worldwide Developer's Conference in San Jose, I heard about the features that Apple intends to include in System 7.0. Among the changes is a redesigned Finder that incorporates significant new interface hooks and a greater level of functional integration. It will also include support for E-mail, larger directories, networks, international scripts, and foreign file systems, while also being more extensible. This prototype Finder supposedly organizes files along the lines of Apple-Share's Desktop Manager, which does a much better job than the current Finder at handling multiple large volumes.

The new Finder is also supposed to include many file management features borrowed from Unix, including file aliasing, so that you can open a file with a variety of applications directly from the Desktop, rather than opening the file within an application. This should make file organization much simpler.

The point, though, is not how accurate this description is of future Apple system software. John Sculley has repeatedly announced that Apple intends to create a new operating system that depends on a brand-new system kernel that could be outfitted with different shells depending on the operating environment desired. The important point is what the changes to the Finder and its file management methods do to the standardization of Mac application interfaces.

Will new Finders make that standardization harder or easier? How will Apple help developers maintain application interface consistency? If past experience is

---

ITEMS DISCUSSED

Nisus 1.0 .... $395
Paragon Concepts, Inc.
990 Highland Dr., Suite 312
Solana Beach, CA 92075
(800) 922-2993
(619) 481-1477
Inquiry 1032.

Word 3.02 .... $395
Microsoft Corp.
16011 Northeast 36th Way
P.O. Box 97017
Redmond, WA 98073
(800) 323-3577
(206) 882-8080
Inquiry 1033.

---

a guide, the future promises more confusion than clarity. As the Mac moves into its second five years, Apple needs to be a trendsetter for application developers, pointing the way toward sensible interface standards that vary according to the application, but retaining cross-application conventions where possible. As *The Macintosh Way* puts it, Apple needs to do the right thing in the right way.

Corruption
A look at some current Mac programs gives you some clue as to how application interface standards can get corrupted. I'm writing this column using the Nisus word processor. Nisus follows the interface ideas first established by MacWrite pretty closely. One look at the Nisus menu bar confirms this: It has File, Edit, Search, Tools, Font, Size, and Style selections. The Quit command resides within the File menu and can be activated via either the menu bar or the Command-Q key combination. Cut, Copy, Paste, Clear, and Select All reside (as they should) in the Edit menu. The Font menu contains all the available fonts, while the Size and Style menus modify the characteristics of the current typeface.

The upshot of all this is that you know how to use Nisus without opening the user's manual. That's the way it should be, and it's the reason I love the Mac. I can spend my time computing with the Mac, rather than learning an application. But what's going to happen with these kinds of interface standards when Apple extends the Finder and makes it more extensible?

Even with the present Finder and its allied set of interface conventions, applications can quickly diverge from accepted standards. Microsoft Word 3.02 is a case in point, with its Short and Full

---

menu settings. These settings change what is available under each menu-bar listing, rather than just dimming those items that are unavailable.

You'd be surprised at how many people call me up and ask how they can install all their system fonts for use in Word. They tell me that they have 20 fonts installed in their System file that they can use in MacWrite or WriteNow, but only five of them show up in Word's Font menu. The problem, of course, is that they've selected Short menus as their default setting. This eliminates all but the five most commonly used fonts from the Font menu and causes a great deal of confusion, especially among new users.

What Apple Should Do
I'm lousy at predicting the future, but I know that Apple can do a lot toward ensuring a future that makes consistent user interfaces easy for application developers to incorporate. First, Apple should show developers exactly how to use the new Finder features in its Tech Notes series. Second, Apple should modify MPW and MacApp (perhaps with an MPW version 4.0) to include the interface extensions that Apple would like to see in other Mac applications.

Third, Apple should publish a new series of books (perhaps through its publishing arrangement with Addison-Wesley) devoted to incorporating what Apple thinks is a standard user interface for applications. Naturally, those interface standards will vary according to the kind of application. For example, things that would be interface oddities in a CAD program (e.g., a separate menu entry for text searching) make perfectly good sense in a word processing program.

As Apple moves toward a more integrated Finder that controls the Mac without the assistance of desk accessories and small applications like the Font/DA Mover, Mac software developers will have to pay special attention to establishing new application interface standards and sticking with them, even when "hot" new ideas argue for violation of those standards. In the past, these hot ideas have produced dubious software achievements like the Short and Full menus of Word, the many interface anomalies of Lotus's failed Jazz program, and the quirkiness of chart manipulation in Microsoft's Excel. ∎

Don Crabb is the director of laboratories and a senior lecturer for the computer science department at the University of Chicago. He can be reached on BIX as "decrabb."

Enter The New Age of Electronic CAD

The wait is over for a powerful, easy to use electronic design workstation.

With the new Douglas CAD/CAM Professional System, you can now experience computer-aided design without going over budget and without sitting through months of tedious training. Running on the Apple Macintosh Plus, SE and II, the Professional System from Douglas Electronics excels in price/performance, short learning curves and ease of use.

As the newest addition to the Douglas CAD/CAM line of printed circuit board design and manufacturing systems, the Professional System is a fully integrated engineering tool that will take you from the schematic drawing to the final routed board. The software features full color, unlimited multi-layers and .001" control which makes surface mount technology (SMT) and other difficult tasks a snap. Professional Layout includes a parts placement facility. Schematic includes fully interactive digital simulation and net list generation. A flexible, multi-pass router completes the design cycle with a 16 layer routing capability.

The new age of electronic CAD has come with the high resolution and speed of a Macintosh engineering workstation. You'll be designing your first circuit board just minutes after the Professional System software has been loaded into your computer. In addition, the Macintosh's graphics capabilities allow for powerful features such as the ability to transfer Professional System drawings into final engineering documentation.

Computer-aided design wasn't meant to be time consuming and complicated. If your present CAD system has got the best of you, it may be time you enter the new age of electronic CAD with the powerful, easy to use Douglas CAD/CAM Professional System.

Take your first step by ordering a full-feature Demo. All three programs are included for just $25.

Call or write for more information and to place your order.

Circle 457 on Reader Service Card

Douglas *Electronics*

718 Marina Blvd.
San Leandro, CA 94577
(415) 483-8770

Circle 20 on Reader Service Card

World Radio History

SPRING COMDEX: Glimmers of Acceptance

OS/2 may take the world by storm after all

S pring Comdex took place in Chicago this year. That transformed it, as Wayne Rash put it, into Winter Comdex, because it actually snowed on Sunday. Some people blamed the snow on the return of a Daley to the mayoralty. Others wanted to blame Comdex's organizers, the Interface Group. But I've spent a lot of time in Chicago. I knew it was just April and came prepared.

Microrim and Logitech announced exciting new OS/2 products. Also, many more OS/2 applications are actually shipping. Some are still late, like Microsoft Word (at least at this writing; it should be out by the time you read this). WordPerfect filled the gap with a full-featured version 5.0 for OS/2. Despite growing OS/2 acceptance, several important applications appeared under DOS extenders rather than OS/2. And industry officials beat the drum for OS/2, of course.

The First PM Screen Generator
Microrim was one of the first companies with an OS/2 product, R:base Series 5000. Now it's offering a completely new database product called Atlas, intended to manage complicated database relationships and integrate databases from places as disparate as a Macintosh or a mainframe DB2/SQL database. Microrim says that Atlas will be available for the Presentation Manager (PM), the Mac, Sun workstations, and AIX-based systems. The company also indicates that Atlas will understand graphics images in its database.

The feature that interested me, however, was the screen generator. Like the applications generator in R:base, Atlas will have a simple way to generate user input screens. No big deal, right? Right, until you realize that this can generate a complete PM screen—including buttons, radio buttons, dialog boxes, and all the rest of the PM notions!

Microrim says it won't have the PM version ready until the end of the year, and I'd be surprised if it can finish something that big by then. However, some of the screen generator does work, and I was able to put together a PM screen in a few minutes. (When a screen is transported to the Mac, it even translates items like buttons and slider bars to items from the Mac metaphor.) I've been complaining that we need something that lets normal mortals design PM-type applications. Atlas could be it, provided it ships early enough.

Multiscope
Developers the world over know Code-View. Microsoft offered it several years ago, and it's still the software-based debugger of choice for many folks. Now it has some competitors, all claiming to be "CodeView killers." (Why do we use such violent language in this business?) Some are marginally better, but Logitech may have a product that can do the job. It's called Multiscope.

Multiscope does everything that CodeView does, and much more. It has a real-time debugger that works the way CodeView works: You run the program to be debugged under the real-time Multiscope, and you can set breakpoints (places where the application should yield control back to the debugger so that you can examine variables and registers) and "watch" windows where a variable's value is continuously monitored in a window. Multiscope also has some fairly sophisticated abilities to use conditional breakpoints (stop whenever variable IS-READY changes), something I find I use all the time when debugging.

continued

SOLUTIONS

Bigger—

Vfeature Deluxe™ lets your DOS system use hard disks it thought it couldn't, all in one bootable piece—no artificial partitions! Span two drives into C: and boot from it, use MFM, RLL, ESDI, or SCSI. Interleave selection, physical format, security options included. DOS 3.1-3.3 **$120**

DUB-14™ PCB takes a different approach to drive expansion, stretches your AT's Drive Table to support the drive *you* choose—*MFM, ESDI, RLL, up to 2048 cylinders!* Comes with setup and low-level format routines, works with UNIX, XENIX, Pick, Novell, DOS. **$95**

Faster—

Vcache™ speeds disk operations, stores data in RAM so it's there for you next time you need it—no waiting! Optional delayed sector write, lookahead buffer. Works with any type of memory, caches up to 15 Mb, bundles accelerators for your diskettes, screens, and keyboard. DOS 2 - 4.
$59.95

Vopt™ defragments disks for quicker access. Run it every day and keep your drive as fast as new! In a few seconds at boot time, *Vopt* arranges all your files neatly in contiguous clusters so you won't waste any time reading them back. Bundles timing and diagnostic utilities. DOS 2 - 4 **$59.95**

Safer—

Vlock™ protects your system and its data from vandals and accidents. System access passwords control booting, activity menus decide who does what to which files, even locks out Ctl-Break during boot! **$135**

Easier—

Vtools™ is a slick set of disk management tools for DOS and OS/2. Display directories in whatever order you like, browse files and change or delete them, compare and update multiple versions, find and manipulate categories of files, and more! **$49.95**

SEE YOUR DEALER OR CALL TOLL-FREE

1 (800) 284-3269

GOLDEN BOW SYSTEMS
2870 5TH AVENUE
SAN DIEGO, CA 92103
(619) 298-9349
FAX (619) 298-9950 TELEX 201520 GBS UR
MC/VISA   US shpg/hdlg $3   CA orders add 7%

Vfeature, DUB-14, Vcache, Vopt, Vlock, and Vtools are trademarks of Golden Bow Systems

All that is only mildly more powerful than CodeView. The neat stuff starts with the postmortem debugger (see, another violent metaphor). If a program dies unexpectedly under many debuggers, you just get dumped into the debugger's main menu. Instead of that, Multiscope's postmortem debugger lets you see everything that led to the untimely demise of your program: One window shows why it stopped, you see the assembly language or source code in another window, and the sequence of called routines is in another.

Multiscope can also isolate and debug individual *threads*. Although I haven't said much about them so far, threads are the basis of OS/2's multitasking capability. An OS/2 program can be made up of a number of threads; OS/2 gives each thread, in turn, a slice of the processor's time. Multiscope's ability to focus on a particular thread of execution is an important feature.

C programmers will love Multiscope's graphical representation of *pointers*. The tough part of data manipulation in C is getting used to the notion that you're not dealing with data structures, but often indirect references to data structures called pointers. In some cases, C programmers find themselves with pointers that lead to other pointers, which in turn lead to other pointers, which finally lead

ITEMS DISCUSSED

Designer 2.0 .... $695
Micrografx, Inc.
1303 East Arapaho
Richardson, TX 75081
(214) 234-1769
Inquiry 1021.

GML/PC $995
Command Technology Corp.
1900 Mountain Blvd.
Oakland, CA 94611
(415) 339-3530
Inquiry 1022.

Multiscope $395
Logitech International SA
6505 Kaiser Dr.
Fremont, CA 94555
(415) 795-8500
Inquiry 1023.

WordPerfect 5.0 .... $495
WordPerfect Corp.
1555 North Technology Way
Orem, UT 84057
(801) 225-5000
Inquiry 1024.

Multiscope can also isolate and debug individual threads.

to a data structure. This takes some getting used to.

The folks who advocate graphical user interfaces (GUIs) often cite the old saw that "a picture is worth a thousand words." In this case, it's worth probably ten thousand words. Multiscope actually draws a picture of a program's pointers. Logitech demonstrates this with some source code that is absolutely impenetrable—pointers to pointers to.... However, the graphic representation clears it up immediately.

This is, of course, only a brief overview of the things that Multiscope can do. It's arranged so that all these windows are PM windows, so you can arrange them as you like or collapse any of them to icons. Oh, and I almost forgot, you can use Multiscope as a PM debugger. PM is tough to write code for. The essence of PM is the user interface, so it kind of ruins the effect while developing if half the screen contains debugging information.

Facing a similar problem in the Mac world, Apple originally counseled developers to buy two Macs for development—one to run the program, the other to display the debugging information. It sounds goofy, but it's the fastest way to develop GUI-type code.

Windows has a feature wherein you can shoot debugging information out the serial port to a dumb terminal or a PC behaving like a dumb terminal, a great help to Windows developers. Now you can't do that for PM, unfortunately, but Logitech does the next best thing: Just run a null modem cable between two PM machines, and the second becomes the debugger. The first is, of course, the debuggee. (I couldn't resist.) That's my biggest gripe with Multiscope. Why not just send out simple line-oriented asynchronous messages? That way, the otherwise-useless PCs that are lying around an OS/2 developer's shop could earn their keep as recipients of debugging information. Please, Logitech—it's a nice product now, but you could make it a killer.

You know exactly what your company wants in a color printer.

Other OS/2 Applications

Mark Mackaman of Microsoft told the audience at the "OS/2 Update" session that there are currently 850 announced OS/2 applications, 370 of which are shipping now ("now," recall, is mid-April). The audience chuckled when he then announced that "three of these are even PM applications."

WordPerfect Corp. showed its character-based implementation of WordPerfect 5.0 under OS/2. It seems to have all the features of the DOS version, including the ability to seemingly talk to all graphics formats possible.

Micrografx again showed beta copies of Designer 2.0 under PM. Designer is the application that you show your Mac-using friends when they start talking about all the neat things that they can do with their machines and MacPaint that you can't do with the PC.

I'm not really the person to comment on the power of Designer, because I use it to draw fairly simple pictures. But I'm happy with it, and even happier with version 2.0 under DOS. (I'll report on the OS/2 version as soon as I can get a copy of it.) Probably the neatest thing is an auto-outline feature that reads in a scanned TIFF file and converts it to a line drawing.

Desktop Publishing

I have said in earlier columns that the first big class of OS/2 applications to come along would be databases, and that certainly has come true: Just about any database vendor that you care to name has an OS/2 implementation (save Ashton-Tate, and it won't be far behind). But I never guessed that the second class would be, of all things, desktop publishing systems. It's a reasonable fit in hindsight: Desktop publishing needs a graphics platform and gobs of memory, so it and the PM are a natural match.

I have mentioned in passing a desktop publishing system that I've been using since December, one that I'm happy with. However, I've been a might remiss in naming names.

Command Technology Corp. has for years marketed a PC implementation of a mainframe document-preparation language called Script or GML (General Markup Language). It originally interested me because it *does not* come with an editor and can use about any editor that can write ASCII text files. This means that I can generate documents that are useful in both the mainframe environments of my clients and the PC environment of my company.

It's fast and very powerful. It contains a sophisticated macro language, so you can make it do almost anything that you need it to do. It reads Designer or PC Paintbrush files and can be coerced to use a host of others. The package, called GML/PC, is a character-mode application, because it is not WYSIWYG except for a VGA preview feature that I find to be a bit slow and tend not to use. It is shipped with a DOS version, an OS/2 version, and a 32-bit DOS-extender version for 80386 machines.

CTC was the first, but it's not alone. Lennane Advanced Products showed a fairly stable desktop publishing system called DeScribe that it will ship in the third quarter of this year, which is Comdexese for at the end of September. It is an integrated package, but it will write out GML text if asked, so I intend to use it as a preview-and-edit package in combination with GML/PC. The editor is a WYSIWYG-type editor with the Choice Words spelling checker built into it.

Xerox was showing Ventura Publisher 2.0 for PM, and everyone selling a version of PM was using a beta Aldus Page-Maker as a demonstration application. Xerox says it will ship Ventura Publisher at the end of the year, but it may be out by the end of September. As soon as I can get hold of these packages, I'll compare them in this column.

These packages (except for GML/PC) will be in dire straits, however, if some printer drivers don't show up pretty soon. There was a lot of talk about a Post-Script driver coming soon and some talk of a LaserJet-compatible Printer Command Language driver by Christmas, although the PCL driver would not support graphics in its early versions. Strangely enough, the Hewlett-Packard people that I talked to believed that it wasn't HP's responsibility to develop the drivers, saying that it was up to Microsoft and IBM. That's an unfortunate attitude, particularly if it means that we're going to be waiting until the middle of 1990 for graphics drivers for our LaserJets. Perhaps Microsoft and IBM will get the drivers out, or perhaps a third party will (hint, hint) see the enormous amount of money to be made writing a good PM driver for PCL. ∎

Mark Minasi is a managing partner at Moulton, Minasi & Company, a Columbia, Maryland, firm specializing in technical seminars. He can be reached on BIX as "mjminasi."

Your questions and comments are welcome. Write to: Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.

Monet, not money.

The HP
PaintJet.
$1395.*

HEWLETT
PACKARD PaintJet

Who says fine art is out of reach? The HP PaintJet color printer produces brilliant color for a price any business can afford.

So now there's no limit to what you can create

with your business communications. Surprise your audience with thousands of colors. Beamed up on an overhead. Or tucked neatly into a report. Persuading people up to 85% more effectively than black and white.

The PaintJet works with all your favorite graphics, presentation, spreadsheet and word processing software. Just hook it up to your IBM-compatible or Macintosh computer and start painting.

For only $1395 (add $125 for the Macintosh interface).

Call **1-800-752-0900 Ext. 711K** for your nearest authorized HP dealer and a free sample output. The HP PaintJet. It's what artists are starving for.

There is a better way.

**HEWLETT
PACKARD**

Circle 114 on Reader Service Card

Our Printer Sharing Unit Does Networking!

An Integrated Solution

Take our **Master Switch**™, a sophisticated sharing device, combine it with **MasterNet**™ networking software for PCs, and you've got an integrated solution for printer and plotter sharing, file transfer, electronic mail, and a lot more. Of course you can also share modems, minis, and mainframes or access the network remotely. Installation and operation is very simple.

Versatile

Or you can use the Master Switch to link any computer or peripheral with a serial or parallel interface. The switch accepts over 20 commands for controlling the flow of data. It may be operated automatically, by command, or with interactive menus. Its buffer is expandable to one megabyte and holds up to 64 simultaneous jobs. The

MasterLink™ utility diskette for PCs comes with every unit and unleashes the power of the switch with its memory-resident access to the commands and menus.

Other Products

We have a full line of connectivity solutions. If you just want printer sharing, we've got

it. We also have automatic switches, code-activated switches, buffers, converters, cables, protocol converters, multiplexers, line drivers, and other products.

Commitment to Excellence

At Rose Electronics, we're not satisfied until you're satisfied. That's why we have thousands of customers around the world including large, medium, and small businesses, factories, stores, educational institutions, and Federal, state, and local governments. We back our products with full technical support, a one-year warranty, and a thirty-day money-back guarantee.

Call now for literature or more information.
(800) 333-9343

ROSE ELECTRONICS

Give a Rose to your computer

GROWING PAINS

The LAN operating
system you choose
can spell the difference
between control
and chaos as your
network expands

Network operating systems are made up of many elements that must work together. But when you're planning for future network growth, the architecture of two components—the LAN's administrative controls and network object identity—are crucial.

The design of network operating systems rests on a few fundamentally different theoretical premises that affect their ease of use and adaptability to outside connection. While many other characteristics affect network growth, administrative controls and network object identity are so fundamental that they're easy to overlook when you're facing the plethora of detail streaming from vendors' advertising or engineering departments.

Maintaining Control

Control refers to the largest possible unit of a network that you can configure with respect to network resources. The larger the unit that you can configure, the larger the network can grow without becoming unmanageable. There are three levels of control: workstation, server, and network.

At a minimum, a network must provide for control on a workstation-by-workstation basis. If all the control possible under the network operating system is vested solely in the workstations, then you can place the information regarding user names, rights, files, pathways, and

security only within each workstation, and the ability to change this configuration is generally open to all or most users of the system as they log on at any given workstation. The system administrator must make system configuration changes by going to all workstations that are to have shared resources and changing the necessary pathways, ports, and so forth. The original version of the IBM PC LAN Program is an example of a network operating system that uses workstation-level control.

A network configuration is server-controlled when servers on the network store information regarding user names, rights, files, pathways, and security, and the ability to change this configuration is restricted to those administrators and users who are defined as having this right of control. In this model, a change in network configuration often requires a visit to every server in the network or, at the least, remote log-on as the supervisor of each server to make the changes at each

server. Novell's NetWare 286 is a good example of a network operating system that uses server-based control.

In the third category of control, overall network administration, the network software recognizes the network as a whole, and a single user can administer the network from a single point. The network administrator has control over the entire network configuration, regardless of the number of servers and workstations or their location. 3Com's 3+Open, Banyan's VINES, and Torus's Tapestry II all support network-level control.

Another name for network control is domain management. Usually, one computer on the network, designated as the domain manager, stores the overall configuration, identity of objects (using the naming conventions that are discussed later), and information on resources outside the boundaries of the domain to provide transparent communication to other domains. You can construct a domain

continued

based on physical boundaries, or you can establish it using logical groups of users, independent of the network's physical layout. The domain concept is most common in large, multiserver LANs.

Pathways vs. Names

Identity speaks to the means by which the network identifies objects. An object is any entity that the network needs to identify. Workstations, printers, and servers are objects, but so are somewhat more abstract items, such as users and administrators, directories, files, and the configuration of the network itself. LAN operating systems maintain the identity of objects via a pathway or a name.

A network configuration maintains identity of objects within it by pathway if access to a given object from any other object requires a statement of the paths, routes, trees, or other structures that the network operating system must traverse to find the object sought. This is the traditional means of describing objects in computing systems based on terminal and host structures. A typical network using pathway identity might describe an item of information as Server1\SYS-:Root\Apps\Spsht\Lotus\123. Novell's NetWare 286 uses pathway identity techniques.

Pathway identity schemes are acceptable in smaller environments when the network configuration doesn't change very often. However, in large, multiserver LANs, or even in small LANs where user moves and changes are frequent, this technique becomes inefficient.

Ancient peoples believed that knowledge of a name gave one power. In a similar fashion, naming conventions in computer networks give users power. A network maintains the identity of objects within it by name if access to a given object from any other object requires only that the user state the name of the object

Relative naming lets you create similar names but distinguishes between them by relating each one to something else.

sought. This technique provides the most power in distributed computer networks. Several methods may be used to name the objects that the network must manipulate.

Absolute naming provides a unique name for each object across the entire network and all networks to which it is attached. Just as, in the case of social security, no two people have the same number, so here, no two objects have the same name. And unlike pathway identities, the name remains the same wherever the person goes.

Relative naming lets you create similar names but distinguishes between them by relating each one to something else. For example, there are many people named John Smith, but relating the name to a street address and a city provides a relative description that removes the ambiguity.

The absolute convention encourages centralization and is the form often found in traditional data-processing environments. The relative convention encourages decentralization but requires a system that will look up the names in their relative context.

Hierarchical naming adds a layer of structure to naming. It lets you embed both absolute and relative naming functions. You can add more levels of hierarchical naming if needed. The telephone system is a good example of hierarchical naming. Each locale has telephone exchanges and numbers. At the regional level, there are area codes. Finally, international calling adds country codes. This illustration demonstrates the utility of hierarchical naming for combining networks.

For a naming system of any size to work, there must be some device or set of devices that contain the names and provide appropriate mappings with objects. At the Xerox Palo Alto Research Center (PARC), which conducted pioneering research that established the basis of distributed processing, such a device is called a *clearinghouse*.

Names consist of three parts: local

World Radio History

name, domain, and organization. With just this three-part name, users can identify and locate any object in any network. The syntax generally looks like this: localname@domain@organization.

You could identify a user in such a network as JohnSmith@Marketing@Greater Tuna, Inc. He could have alias names such as John and JS. The administrator might refer to a physical resource as Laserprinter@Marketing@Greater Tuna, Inc. And information might be designated as Budget1988@Marketing@Greater Tuna, Inc.

The mechanical process involves establishing a clearinghouse, usually called a name server. This is simply a server running a database that relates the name with the pathway to find the object. Thus, any reconfiguration, which involves moving objects around, need only change the one reference in the name service at the name server.

Contrast this with an identification technique based only on the pathway. In such a system, a reconfiguration would have to seek out every reference to the now-changed pathway and alter it—good luck in a network with hundreds of work-

stations, users, and servers. Banyan's VINES and 3Com's 3+ both use three-part naming techniques. Torus bases its product on icons associated with objects through a library service; this is analogous to the name/clearinghouse concept.

Strengths, Weaknesses, and Changes

A network operating system based only on workstation control simply can't provide the management and consistency of configuration that are needed to provide a stable network environment for more than a handful of workstations. Although the original IBM PC LAN Program has this limitation, version 1.3 of PC LAN Program and the newly emerging LAN Server, IBM's OS/2 LAN Manager–based network operating system, use what IBM calls Domain Management to achieve network control. The new system also has a name identity technique.

A network operating system based on server control provides an excellent single-server network, but the need to administer several servers becomes an overwhelming headache. Novell is aware

of this limitation in NetWare, as well as similar difficulties that arise from NetWare's lack of a name service. This problem arose historically because NetWare's designers conceived of PC LANs as single-server systems much like minicomputers; they never anticipated the advent of larger multiserver networks that need easy, flexible control and identity methods.

NetWare is an elaborate product based on proprietary coding down to the machine code level; total redesign will take time. Eventual upgrades to NetWare should provide network-level control and name identity.

Network control gives the best possible environment for growth. A networked system may start with only one server and grow to tens or possibly hundreds, assuming that there is a consistent method of control. The added protection and power of machines based on the Intel 80386, Motorola 68030, Sun SPARC, and others will provide complex combined workstation/server systems that challenge even the best designs. Combine this with a name service, and administration is eased substantially. Both Banyan and 3Com historically benefited from their designers' early involvement in the initial research for such networks at Xerox PARC.

There is, however, one serious problem that still lies within name service-based systems (and within systems that are dependent on a single physical device for overall management): What happens if the name service (or the domain management device) fails? 3Com implements the name service in a single server; loss of that server leads to loss of the entire network. Most large 3Com installations maintain a "hot spare" for the name server. Banyan distributes the name service over several servers, but the service isn't redundant; loss of any server loses a portion of the name service. The final solution for large systems is a totally redundant name service and domain management service. Providing these capabilities will be the next major push in the LAN operating-system market. ∎

James Y. Bryce is an independent network consultant and author living in Austin, Texas. He is the author of the forthcoming Networking Personal Computers: The Total Context *(New York: Van Nostrand Reinhold). You can reach him on BIX c/o "editors."*

Your questions and comments are welcome. Write to: Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.

World Radio History

Desktop Power to Go

The BYTE Lab
sizes up 11 of the best
80386 portables

Stanford Diehl
and Stan Wszola

I n the past, when you grabbed your portable computer and hit the road, you left a lot behind in your desktop computer. Portability was inversely proportional to computing power. That's not the case anymore. Modern portables offer fast CPUs, plenty of RAM, big hard disk drives, and enough options for almost any computing situation.

The new line of portable powerhouses blurs yet another distinction in the evolving computer world. You no longer need to choose between portable convenience and desktop power; today's portables deliver both. Even the distinction between portables and workstations is fading. Designers keep packing more features and firepower into an ever-shrinking shell.

Portable computer vendors have devised many variations on a common theme. This month, we'll look at 11 of the most powerful computers currently available: the Compaq Portable 386 Model 40, the Dolch-P.A.C. 386-25, the GRiDCase 1530 and 1535 EXP, the IBM PS/2 Model P70 386, the Micro Express Regal II, the NEC PowerMate Portable SX and ProSpeed 386, the Toshiba T5100 and T5200, and the Zenith TurbosPort 386. Each machine offers a unique combination of computing power and portable convenience (see table 1).

No More Trade-Offs
Since the very first portable computer appeared, buyers have always had to weigh the importance of small size and weight versus computing power. If you wanted a powerful computer, you had to accept bulk. Lightweight portables usually lacked power.

But an amazing evolution has occurred in portable computers. By means of smaller components, better batteries, and VLSI surface-mount technology, today's portables squeeze more computer into smaller packages. For example, one of the first portable computers, the Osborne 1, weighed 23½ pounds. It was a CP/M system with 64K bytes of RAM, a CRT display, and dual floppy disk drives. Today's portables pack 1 or more megabytes of RAM, up to 170 megabytes of hard disk drive storage, high-resolution displays, and your choice of DOS, OS/2, or Unix into even smaller and lighter packages.

Even though these new computers are lighter, ranging from 12½ pounds for the GRiDCase 1535 EXP with its magnesium case to the Micro Express Regal II at a hefty 22⅖ pounds, most people don't carry a "naked" computer. Add the weight of a carrying case, an AC power supply, a spare battery, a modem, blank floppy disks, and assorted hardware and software manuals, and you have enough weight to make a business trip an endurance contest. The Traveling Weight column in table 1, which is the sum of the weights of the computer, the case, and essential accessories, is our idea of a more realistic weight.

Most portable machines fall into two design groups: the large lunch box (e.g., the Compaq and IBM) or the clamshell (e.g., the Toshibas and the Zenith). With its detachable keyboard, the lunch box style works best on a desktop, while the clamshell models can sit on your lap. In terms of functionality, both designs can get the job done.

Power for the Road
For those portables that use batteries, the power source of choice is the nickel-cadmium cell. It provides a relatively steady voltage per charge, and it recharges easily. One disadvantage, however, is that it can develop a "charge memory." Repeated recharging when a battery is only partially discharged can render a nickel-cadmium battery pack incapable of being fully charged. Most portable manufacturers recommend that you discharge the batteries as much as possible before recharging.

Most portables can run on internal batteries for 2 to 3 hours, depending on the size of the battery pack. The Zenith TurbosPort 386 extends battery life through a built-in monitor program. This ROM-based program lets you enter the number of seconds that the hard disk drive runs after the last disk access and the amount of time that the LCD backlight remains on if there is no keyboard activity. The monitor program will power down these sections of the computer to conserve battery power.

Picture This
Displays for high-end portables fall into two groups: LCD or gas-plasma/electroluminescent (ELDs). The photo on page 144 shows a sample of both.

LCD screens are popular because of their light weight and low power requirements. An LCD is a reflective screen; the individual pixels in the screen work like a set of light shutters. They control whether light is absorbed (producing a dark spot on the display) or whether light is reflected (producing a light spot). Unfortunately, the LCD scheme lacks sharp contrast between the dark spots (text and graphics) and the lighter background. This caused serious problems with early LCD screens. You needed good ambient illumination for comfortable viewing. Portable designers have overcome that problem by using fluorescent backlighting for their LCD screens. The backlighting increases the apparent contrast between the text and the background

Table 1: *80386 portable features and conventional benchmark results. For the Livermore Loops and Dhrystone tests only, higher numbers mean faster performance. LINPACK and Livermore Loops benchmarks are in seconds; Dhrystone benchmarks are in Dhrystones per second. Prices are for base models not including options. The weight for each portable includes the battery or AC power adapter. The Traveling Weight includes the optional case, battery and/or AC adapter, power cord, modem cord, external monitor cable, and four floppy disks. For a full description of all the benchmarks, see "Introducing the New BYTE Benchmarks," June 1988 BYTE.*

Model	Price	CPU speed (MHz)	Conventional benchmarks			Display/mode	Memory (Mb) Std./Max.	Floppy disk drive[1]
			LINPACK	Livermore	Dhrystone			
Compaq Portable 386 Model 40	$7999	20	182.80	0.0688	5117	Gas-plasma/CGA	1/2	5¼-inch
Dolch-P.A.C. 386-25	$9495	25	155.55	0.0797	6410	ELD/CGA	4/8	5¼-inch
GRiDCase 1530	$4695	12.5	445.66	0.0217	2955	Backlit LCD/CGA; gas-plasma/CGA opt.	1/8	3½-inch
GRiDCase 1535 EXP	$6995	12.5	448.08	0.0219	2956	Backlit LCD/CGA; gas-plasma/CGA opt.	1/8	3½-inch
IBM PS/2 Model P70 386	$7695	20	196.58	0.0577	4975	Gas-plasma/VGA	4/8	3½-inch
Micro Express Regal II	$2999	20	177.57	0.0662	6410	Gas-plasma/EGA	1/2	5¼-inch
NEC PowerMate Portable SX	$6595	16	508.50	0.0240	1813	Gas-plasma/VGA	2/10	3½-inch
NEC ProSpeed 386	$7699	16	242.60	0.0526	4009	Backlit LCD/VGA	2/10	3½-inch
Toshiba T5100	$7199	16	238.27	0.0525	4081	Gas-plasma/EGA	2/4	3½-inch
Toshiba T5200	$9499	20	149.10	0.0790	6459	Gas-plasma/VGA	2/8	3½-inch
Zenith TurbosPort 386	$7999·	12	313.00	0.0363	3448	Backlit LCD/CGA	2/3	3½-inch

[1]5¼-inch floppy disk drive = 1.2-megabyte; 3½-inch floppy disk drive = 1.44-megabyte.
[2]NC = Nickel-cadmium.
 N/A = Not available.

Portable computer screen displays. A sample of the two principal technologies: the Toshiba T5200 gas-plasma display (left) and the Zenith TurbosPort backlit LCD (right).

areas on the screen and lets you see the display even in a darkened room. But the backlighting requires more power.

LCD screens are more suited to single users. Since the viewing angle is narrow, a group of people have difficulty looking at an LCD display. The best view is from directly in front of the screen.

Another problem is that currently available LCD screens are monochromatic. The gray scale available on an LCD is very limited. When you run software that depends on a color display, the display circuitry must resort to color mapping to present different colors as contrasting graphics patterns. Color LCD screens have arrived (e.g., on the Sharp PC-8000), but they are not yet widely available.

LCD screens are also slower than gas-plasma or CRT displays. The individual

Hard disk drive (Mb)	Battery	Size (inches)	Weight (pounds)	Traveling weight (pounds)	Notes
40 or 100	None	16×9.8×7.8	21.9	25.6	
40, 80, or 100	None	16.25×10.25×8.5	20.6	25.2	
20, 40, or 100	NC²	11.5×15×2.3	12.7	17.1	Magnesium case
40	NC²	11.5×15.1×2.5	12.5	16.8	Magnesium case
60 or 120	None	18.3×12×5	21.7	25.9	PS/2-compatible
40	None	16.25×10.25×8.5	22.8	27.5	Four months on-site service
42	None	15.5×11.3×7.75	21.9	25.8	
40 or 100	NC²	15.35×3.94×15.5	22.5 (w/batt.)	27.7	Bundled with VM/386
40 or 100	None	12.2×14.2×3.5	14.6	N/A	Bundled with Windows/386
40 or 100	None	14.6×15.6×3.9	18.7	N/A	82385 cache controller; 32K-byte static RAM cache
40	NC²	13.25×4.75×14.75	18.7	24.1	Intelligent Power Manager

pixels in LCDs are electrochemical devices that require an appreciable fraction of a second to turn on or off. They are not ideal for use with animated graphics software or games.

Gas-plasma displays and ELDs are attractive for portable computers. These screens offer high contrast, a wide viewing angle, and good speed. You pay more for these features, and both types of displays require more power compared to LCDs.

Gas-plasma displays are composed of a pair of glass plates. The inside of one of the plates is coated with a horizontal set of transparent electrodes, while the other plate has vertical electrodes, thus forming a grid. Neon gas floats between the plates. When a high voltage is applied to the electrodes, the neon gas located at the intersection of any electrodes in the grid is ionized and becomes a glowing plasma. These glowing points of neon plasma produce the illuminated pixels on the screen.

ELDs are similar to gas-plasma displays, but instead of a gas they use a dielectric thin-film sandwich that contains a layer of manganese-doped zinc sulfide that fluoresces in the presence of AC voltage. ELDs use slightly less power and are more rugged than gas-plasma displays, but they cost more to manufacture.

Both gas-plasma and ELD screens share a problem with LCDs: a limited gray-scale display. They, too, must resort to color mapping to represent colors. In addition, both types of displays have a yellow or reddish-orange color that might not appeal to some users.

For those accustomed to high-resolution desktop displays, using a portable might be a disappointment. The graphics adapter circuitry in these machines ranges from double-scan CGA (640 by 400 pixels) to VGA. CGA on a monochrome screen is only adequate for most users. The Toshiba T5200, both NEC models, and the IBM PS/2 Model P70 386 employ VGA graphics circuitry for good screen displays.

External monitor ports, which are available on several of the portables that we reviewed, offer an easy upgrade for desktop use. An external CRT monitor has much better contrast than LCD, gas-plasma, and ELD screens, and, with color monitors, you can use color-based software.

Keyboard Quirks

When it comes to keyboards, a portable computer designer's imagination runs rampant. Nearly every keyboard has a unique layout. Cursor-control keys, numeric keys, and programmable function keys (F1 through F12) can all be "redistributed" on the keyboard.

Because of space limitations, many keyboards are "compressed." Accessing some keys requires holding down a function key before pressing another key. For example, the numeric keypad can be embedded in the alphabetic keys, and some control keys might have double, or even triple, functions. Certain keystroke combinations that are used in word processing programs or program editors might increase in complexity and become awkward. A Ctrl-Shift-F5 might turn into a Function-Ctrl-Shift-F5. Before you select any portable, consider the software you're likely to use and how it will function on a particular computer.

A notable exception to the rule of compression is the IBM PS/2 Model P70 386. Its keyboard adheres to IBM's standard. It has separate numeric and cursor-control keypads, and it also has a mouse port. Users of PS/2 desktop machines can switch easily to the P70.

Some portables, such as the Dolch-P.A.C., the GRiDCases, and the Micro Express Regal II, support a full-size IBM PC–compatible keyboard through an external keyboard port. Another popular option offered by some portable manufacturers is an auxiliary numeric keypad to ease intensive math data entry.

Our best advice to you is to try the keyboard before you buy a portable computer. A fast 80386 CPU is no advantage when your fingers are constantly lost on the keyboard.

Megabytes to Go

The data storage options for portables can cover almost anything you want. You can have your choice of floppy disk drives: 360K-byte or 1.2-megabyte 5¼-inch drives; 720K-byte or 1.44-megabyte 3½-inch drives; and—soon to be available—the 720K-byte 2-inch microfloppy disk drive.

When it comes to hard disk drives, the choices are even more impressive. Hard disk drives in portables range from a pedestrian 10 megabytes to a staggering 170 megabytes as an option for the Dolch-P.A.C. The sizes vary from the standard 5¼-inch size to the 3½-inch units and down to the recently announced 2½-inch hard disk drives.

Most portable hard disk drives are

specifically designed for portable use; some can withstand up to 75 g's of shock. Some portables (e.g., the Toshibas and the Compaq) have special hard disk drive mounts to minimize shocks.

Traveling Options

It used to be that when you bought a portable computer, you were stuck with what you got. Some machines had sockets for more RAM, but that was it for expandability. That's not the case now. Many portables have either proprietary or PC-compatible expansion slots. Peripherals and options are widely available from portable manufacturers and third-party suppliers. Toshiba offers a variety of external floppy disk drives, plus memory and modem options that will fit in the T5100 and T5200's proprietary slots.

Third-party manufacturers, such as the Megahertz Corp., offer a variety of enhancement products for portable computers. Megahertz offers both 1200- and 2400-bps internal modems for Toshiba and Compaq portables. In addition, Megahertz sells the LapLan card for Toshiba portables, which follows the IEEE 802.3 protocol and is Novell NetWare–compatible, and the MHZ-T-3270 remote terminal emulation card, which supports all IRMA and IBM emulations.

You can now travel from office to office and literally plug into your company's mainframe computer or LAN. It's also possible, using laptops with large hard disk drives and plenty of RAM, to design your own portable LAN.

With so many portables offering big hard disk drives, the question of data backup arises. Floppy disk drives aren't convenient when dealing with hard disks that are 100 megabytes or larger. Some manufacturers, such as Compaq, offer tape backup units for their computers as an option. Procom Technology offers its PLT series of external tape backup units for most popular portables. Both Compaq's and Procom's units use the industry-standard QIC-40 format with the DC2000 tape cartridge.

When the need for expansion options goes beyond the ordinary, many portable users add optional expansion units. One of the most unique is the NEC ProSpeed 386 Docking Station. The Docking Station attaches to the rear of the ProSpeed and has space for two half-height drives and slots for three 16-bit and one 8-bit full-length PC-compatible cards. Even the diminutive GRiDCase 1535 EXP has a clip-on expansion tray that can hold one 16-bit and one 8-bit card.

What follows is a closer look at each of the 11 portables we examined.

Compaq Portable 386 Model 40

O ld Reliable keeps plugging away. Compaq has slimmed down its portable and made some subtle changes, but this machine is still the same old rugged workhorse we've come to count on. When it comes time to do some tough computing work, the Compaq is ready to go. It's a solid 20-MHz performer, though it finished only fourth on the BYTE benchmarks (see the figure on page 154). It doesn't have the best individual specs, and it's not the cheapest portable, either. The gas-plasma screen lacks the sharpness of other models. Yet the final combination adds up to an optimal mix of features, performance, and quality. Other portables may seem flashier, but none are more dependable.

Perhaps the Compaq's biggest flaw is its gas-plasma display. At one time this screen seemed brilliant, but it doesn't shine so brightly when set next to today's new crop of portables. You can make it brighter with the only control knob, but most likely you'll keep that button fully tweaked, anyway. Forget about adjusting contrast; there's no knob for that. In the end, you can't avoid the washed-out look of the Compaq display.

The standard Compaq Portable configuration includes a megabyte of 32-bit RAM, one 1.2-megabyte 5¼-inch floppy disk drive, the CGA gas-plasma display, and a hard disk drive. The Model 40 packs a 40-megabyte hard disk drive, while the Model 100 delivers 100 megabytes of hard disk space. You get a full keyboard with a separate numeric keypad and an RGB port for an external CGA monitor. If you need more than CGA graphics, you'll have to buy the expansion box and install a better graphics adapter. The $199 expansion unit plugs into the rear of the main unit and provides a pair of 16-bit expansion slots. Like the Compaq, the expansion unit is functional, easy to use, and fully IBM PC–compatible.

Dolch-P.A.C. 386-25

T he P.A.C., which packs the fastest CPU in our lineup, harkens back to more traditional IBM PC AT technology. It uses the AT bus on its motherboard and a 1.2-megabyte 5¼-inch floppy disk drive. Yet it combines those features with an 80386 running at 25 MHz with zero wait states, an ELD screen, and a SCSI hard disk drive controller with a 4-megabyte-per-second data transfer rate. In addition, the P.A.C. has a proprietary 64K-byte disk cache for faster data access.

The P.A.C. has a lunch box configuration. The keyboard detaches to reveal the ELD screen. You must plug the keyboard into the side of the unit before turning it on. The screen tilts up if you push a large release button. The screen brightness control and display control are to the left of the screen. The display control lets you adjust the screen for light text on a dark background or vice versa.

The ELD screen is CGA-compatible and has a 640- by 400- (double scan) pixel resolution. Text display is a pleasant yellow on a dark-gray background, with excellent contrast.

If you remove six screws from the back, you'll see the AT bus motherboard with its six slots. The review unit came with 8 megabytes of RAM, a 40-megabyte hard disk drive, a SCSI drive controller, an I/O card for serial and parallel ports, and an ELD/CGA video card. This leaves one 8-bit slot and two 16-bit slots free, which allows for easy expandability. The port connectors are located beneath a plastic cover on the left side of the P.A.C. Dolch also offers the Back Pack, an external expansion module for three full-length 16-bit cards.

Dolch sells a version of the P.A.C. called the COBRA for hosting a variety of computer-based instruments, data acquisition boards, and industrial control modules.

Turn Any Place
Into An Instant Office.

Why be limited by four walls and a desk? Just turn on the mp286L desktop-laptop and you're in business. Anywhere.

The Mitsubishi® mp286L is the instant office. Plus, with every mp286L, you get Microsoft® Works, a modem and a travel bag free.

Easy to learn and use, Microsoft Works is rated by PC World as the #1 integrated software program including word processing, spreadsheet, chart maker, database, report

generator and communication modules. So it will bring out the best in the high performance mp286L.

See the Mitsubishi mp286L in action. For your nearest Mitsubishi dealer call 1-800-556-1234, ext. 25 in the U.S. and Canada (in California 1-800-441-2345, ext. 25).

Then go with the instant office.

▲ MITSUBISHI
ELECTRONICS

World Radio History

1530

1535

GRiDCase 1530 and 1535 EXP

GRiD Systems has attained the status of a portable pioneer, and the company deserves its reputation. The company shipped the first battery-powered portables, and it owns the patent for the basic portable design. The GRiDCase line is a testament to this notable design savvy. The GRiD portables are rugged, battery-powered systems that fold into stylish magnesium cases.

The GRiD models are remarkably similar in that they share basic system specifications. Both units house an 80386 processor running at 12.5 MHz, a standard 1 megabyte of RAM, and two handy ROM slots just above the keyboard. The 1535 EXP features a snap-on tray that delivers one 16-bit and one 8-bit expansion slot.

The first thing you notice when you turn the GRiD on is the startlingly sharp screen display. Our 1535 EXP came with a blue backlit LCD, while the 1530 sported an orange gas-plasma screen. The LCD was impressive enough, but the GRiD gas-plasma screen is stunning. For pure readability and sharpness, it can't be beat. Both screens are CGA-compatible, as is the 9-pin video port.

The second thing you notice when booting the GRiDs is the SCO Xenix sys-

tem installed on them. You immediately start taking these slim portables seriously. GRiD has designed its portables for field engineers and traveling professionals, and the boot-up configuration confirms this focus. If you need DOS, though, it's easy enough to fire it up from the log-in prompt or to change the active partition.

It's a mystery why a company with such a discerning eye for design could not come up with a keyboard better than the GRiCDase's. The keys are awkwardly flat and spaced too closely together, and the keyboard lacks a standard layout. GRiD shrank the Enter key and moved the Backspace key out of easy pinky range. The embedded cursor-control and numeric keys will frustrate any traveling professional who works extensively with numbers; however, GRiD offers an optional numeric keypad to solve that problem. Perhaps GRiD, given its projected market, can afford to alienate the touch typist, but this is a keyboard that even the most ardent hunt-and-peck artist could dislike. If you're placing the GRiD on a desktop, you'll appreciate the external keyboard port.

The GRiDCase line lacks nothing—if you're willing to pay the price. The only problem with its impressive add-ons is that more of these options aren't included under the hefty GRiD price tag. The standard configuration doesn't even include a hard disk drive. Add $1675 to the 1530's base price ($4695) or $500 to the 1535 EXP's base price ($6995) if you need a 40-megabyte unit. Once you go with the hard disk drive, there's no room left for an internal floppy disk drive. An external "pocket floppy," though included with the hard disk drive configuration, must be carried along when you need a floppy disk drive. You can also purchase 5¼-inch drives, backup tape cartridge drives, high-density Xenix drives, internal battery packs and external battery chargers, and an Ethernet Network Expansion Cartridge. By the time you're through adding on, you'll have a fully configured 80386 system, a busted bank account, and a broken back.

That both the GRiDCase 1530 and the 1535 EXP did poorly on the benchmarks is more related to their 12.5-MHz CPU speed than any performance flaw. We would like to see a little more power under the hood, but it's hard to question design decisions when this is one of the few vendors that can free you from an AC plug. If you can leave all the extras at home, you'll carry along a unique combination of power, compactness, and black-tie style.

IBM PS/2 Model P70 386

IBM has finally produced a portable computer that has all the right features. The P70 is a Micro Channel architecture (MCA) machine with the performance of the Model 70 desktop computer. It uses an 80386 running at 20 MHz. There is a socket for an optional 80387 math coprocessor chip.

The P70 has a lunch box configuration in a briefcase size. You must slide two catches to release the keyboard, which folds down to reveal the gas-plasma display. The keyboard can be detached from the computer for easier desktop use. The bottom edge of the display can be pulled out to tilt the display up. Pushing on the inside upper-right corner of the case causes the 3½-inch floppy disk drive to fold out.

The rear of the P70 features slide-up covers for access to the AC power connector, serial port, parallel port, VGA external monitor connector, PS/2 mouse port, and external expansion connector.

If you fold back the rear door, you'll see a storage area for a mouse and have access to connectors for two MCA slots. One slot can hold a 32-bit full-length board, and the other can hold a 16-bit half-length board. The review unit came with a 60-megabyte hard disk drive with an integrated ESDI controller and 8 megabytes of RAM. An IBM 2400-bps modem and an IBM Token Ring Adapter Network board were also included.

The P70 maintains the PS/2 tradition of simple user access. Installing MCA cards is easy: Just release three screws to remove the rear cover. All parts of the portable are at hand. An internal fan keeps the unit cool.

The P70 has a VGA-compatible gas-plasma display with a 640- by 480-pixel resolution. Like many other portables, the P70 uses color mapping when it runs color-based software. But unlike other

continued

"Xerox this memo."
"FedEx this proposal."
"LapLink these files."

When something becomes a standard, using it becomes second nature. That's true about LapLink. It's so effective that it has become the most popular laptop-to-desktop and desktop-to-desktop file transfer program ever.

And now Release III improves on the original with added power—while preserving the simple design that has made LapLink the choice of more major corporations.

LapLink III offers both serial and parallel file transfer, and you can take advantage of parallel transfer speeds of 500,000 baud or higher. It comes with a "six headed" universal cable that provides you with everything you need to use both serial and parallel modes.

And LapLink III will even install itself automatically on a remote computer.

That's in addition to ease-of-use and productivity features like our popular split screen design, flexible transfer options, and disk and printer sharing.

For the same fast, error-free file transfers between PCs and Macintoshes, get LapLink Mac. And for more information about any Traveling Software product, call us at **(800) 662-2652.**

LapLink III. The standard in file transfer software.

Suggested Retail Price **$139.95**

LAPLINK III

Traveling Software, Inc.
18702 North Creek Parkway Bothell, WA 98011
(206) 483-8088

World Radio History

portables, the P70 uses a different method to produce the color mapping for the gas-plasma display.

If a program directly writes a color value to a particular location in VGA memory, the P70's VGA circuitry translates that value to another value that can be displayed as part of the color mapping. If the program reexamines that particular memory location, it will see a different value from what it had originally written. This could cause a problem in that some software that uses direct hardware control of the VGA circuitry might not operate correctly with the P70.

Micro Express Regal II

Since we hold the Compaq Portable in high regard, it's hard not to like the Micro Express Regal II as well. It shares the same lunch box look of the Compaq, the same pop-up gas-plasma screen, the same snap-off detachable keyboard, and the same outstanding performance. In fact, the Micro Express outperformed the Compaq, perhaps helped by a 64K-byte, 35- to 40-nanosecond static RAM cache.

No doubt, the Regal II really screams. In terms of benchmarks, it came in behind only the Toshiba T5200 and the 25-MHz Dolch-P.A.C., returning solid numbers across the board. It finished no worse than third on all the low-level modules and posted top honors on the disk benchmark. It surpasses the Compaq in the expandability department, too, with four expansion slots accessible from a sliding door on the left side of the unit. The right side houses a 5¼-inch floppy disk drive with a standard flip-down latch. We didn't like the way the Regal II packs its keyboard cable, however. Unlike the Compaq, which tucks its keyboard cable effortlessly into a slot next to the screen, the Regal II offers a compartment within the keyboard to store the cable. It doesn't make for a

smooth fit, and the cover for the cable bay can unlatch easily.

The rear of the Regal II offers a parallel printer port and a 25-pin serial port. It lacks an external video port, though you could install an adapter in a free expansion slot. Priced at $2999, the Regal II certainly looks attractive. Of course, when you select the Regal II as an alternative to the Compaq Portable 386, you are sacrificing Compaq's proven record for quality. Though the Regal II seems rugged enough, only time will tell. For the price of the Regal II, it could be worth taking the chance.

PowerMate

ProSpeed

NEC PowerMate Portable SX and ProSpeed 386

Please don't refer to the PowerMate Portable SX as a stunted system. Yes, it employs a 16-MHz 80386SX processor; and, yes, our benchmarks reveal lackluster performance; but this system delivers a remarkable set of features for the price. NEC may have cut some corners on the data bus, but it didn't cut corners anywhere else: The unit has 2 megabytes of RAM, a 42-megabyte hard disk drive, a 1.44-megabyte 3½-inch floppy disk drive, a VGA gas-plasma display, three expansion slots, a 5¼-inch external

drive interface, a 93-key keyboard, and an external VGA port.

All those impressive features add up to a hefty luggable shell, so you sacrifice some portability. Once you set this system up, though, you give up very little. The expansive keyboard offers a separate numeric keypad and dedicated cursor-control keys. The light clicky feel and full-size keys make for comfortable touch-typing. A single-screw door atop the unit exposes three full 16-bit expansion slots as well as the memory and co-processor sockets. A side door affords easy access to the system DIP switches. If you don't really need blazing speed or a 32-bit data path, this machine delivers a wealth of standard features that no other vendor can match.

NEC refers to the ProSpeed 386 as a "modular workstation." If you're really looking to buy one computer for both travel and desktop use, the ProSpeed philosophy may be the answer. The battery-powered portable unit houses a 16-MHz 80386, a 40- or 100-megabyte hard disk drive, up to 10 megabytes of RAM, a fold-down LCD VGA screen, an external VGA port, and an 92-key keyboard with a separate numeric keypad. Even as a stand-alone portable, it's an impressive unit. It has the design of a true portable with battery power and lap-size dimensions; however, with the battery installed, the ProSpeed weighs in at a back-straining 22½ pounds.

When you get back to your desk, you can plug the ProSpeed into the optional Docking Station ($1199). With the Docking Station, you get one 8-bit and three 16-bit expansion slots, bays for two standard half-height storage devices, an external keyboard port, two serial ports, and a parallel port. You can connect an external analog monitor to the RGB port on the portable unit and plug the monitor into an AC outlet at the rear of the Docking Station. A fully configured ProSpeed could indeed qualify as a low-end workstation. While the expansion slots could provide connectivity and other enhancements, the drive bays can support mass storage options, including CD-ROM drives. NEC also bundles VM/386 multitasking software with the ProSpeed.

With these two units, NEC offers some unique portable choices. In addition to being the only SX machine in our survey, the PowerMate offers a fully featured luggable system at a competitive price. The ProSpeed provides a creative solution to users who need both a powerful portable on the road and a fully configured 80386 on their desktop.

T5100

T5200

Toshiba T5100 and T5200

Both of Toshiba's portables show a definite family resemblance; they have the same clamshell design. The T5100 uses a single large, front-mounted latch to release the display from above the keyboard. The T5200 uses two small latches and has a combination lock. The T5100 has a handle mounted on its back; the T5200 has a small handle mounted on the front that folds under. Both units have a gray matte finish, but under the skin, they differ significantly.

The T5100 is smaller, lighter, and slower. It's not as tall or wide as its brother. It's 4 pounds lighter, runs at 16 MHz, and costs over $2000 less. The model we received had a 40-megabyte hard disk drive and came with 2 megabytes of RAM, the standard configuration. The rear of the T5100 sports a serial port, a parallel port, and an EGA connector. You can also use the parallel port to connect an optional external floppy disk drive. A switch on the side of the T5100 configures the parallel port as drive A, drive B, or printer port LPT1.

A metal plate at the rear of the T5100 covers Toshiba's proprietary expansion port connector. The port provides an easy upgrade path; Toshiba offers an

optional 2-megabyte memory board, modem, and external expansion chassis. The T5100's gas-plasma EGA display has 640- by 400-pixel resolution. The display is crisp, and the contrast and brightness controls lie directly underneath.

The T5200 is the top of the line for Toshiba portables. It has an outstanding line of features, but it is one of the most expensive portables. The T5200's 80386 runs at 20 MHz with a 32-bit path to system memory, an 82385 cache controller chip, and a 32K-byte static RAM cache. The review model came with a 40-megabyte hard disk drive and 4 megabytes of RAM. The cache controller chip explains why the T5200 can outrun the Dolch-P.A.C. in some of the BYTE benchmark tests. The rear of the T5200 has connectors for the parallel port/external floppy disk drive port, two serial ports, a VGA connector for an external monitor, and the Toshiba proprietary expansion connector.

One of the reasons why the T5200 is larger than its brother is that you can install two PC-compatible expansion boards, one short 8-bit board and one full-length 16-bit board, inside the rear of the unit. Installation involves removing the rear panel and two metal cover plates. Once installed, the rear of the expansion boards can be accessed through a removable plastic cover on the left side. This provides a convenient upgrade path for expansion boards, such as LAN interfaces or data acquisition cards.

The T5200 also has a gas-plasma display; it is VGA compatible and has a 640- by 480-pixel resolution. The gas-plasma display can be removed when using the computer with an external monitor. You can simultaneously view both the gas-plasma display and an external monitor.

Both units use AC power only. They come bundled with PC-Kwik Power Pak utility software, QEMM-386 memory management software, and Microsoft Windows/386.

Zenith TurbosPort 386

The TurbosPort combines good performance, an innovative design, and convenient operating features. When we used the TurbosPort, we got the impression that a considerable amount of engineering skill went into its design.

The TurbosPort has a modified clamshell design. To open it, you move two slide releases on the sides and tilt up the LCD screen from the keyboard. Once the screen is up, you can detach the keyboard from the rest of the computer by pressing on two latches. This makes the machine easier to use on a desktop.

The rear of the TurbosPort has a serial port, a parallel port, a DB-15 external monitor connector, and RJ-11 connectors for the built-in modem. The machine that we received for review, the Model 40M, came equipped with a 40-megabyte hard disk drive and a 2400-bps internal modem. The TurbosPort's 80386 runs at 12 MHz with zero wait states. The socket for an optional 80387 resides in back of the display. The TurbosPort comes standard with 2 megabytes of RAM, expandable to 3 megabytes internally. The internal SETUP program allows you to configure the memory beyond 640K bytes as either extended or expanded memory.

The TurbosPort's screen is a "page white" backlit fluorescent LCD with 640- by 400-pixel resolution (double-scan CGA). We judged the TurbosPort's LCD screen as one of the best.

The TurbosPort can run on internal nickel-cadmium batteries or an external AC adapter. The adapter is a 7¼- by 2¼- by 4¼-inch box, weighing 1¾ pounds, with a special cable and connector that plugs into the side of the TurbosPort. The nickel-cadmium battery can be recharged in 2 hours if the computer is off, or it can be trickle-charged during use. You access the battery through a

continued

panel in the bottom of the machine. Battery life is approximately 2 to 3 hours per charge.

Portable Performance

The Compaq Portable's standing in our BYTE benchmark graph (see below) reveals the power of the portables we reviewed. Despite its usual place at or near the top of our benchmark listings, the Compaq could do no better than fourth out of the 11 portables tested. This is not so much a sign of Compaq's slide as it is a testament to the quality of this crop of luggable powerhouses. Few users would ever need more power than this—even for their desktop applications.

For the most part, the results reflect the speed of the CPU. However, the Toshiba T5200, a 20-MHz model, outscored the 25-MHz Dolch-P.A.C. The Toshiba T5200 posted a higher CPU

A *32K-byte static cache and 82385 cache controller boosted the T5200's performance above that of the competition.*

score and a higher overall applications score, and it even performed more Dhrystones per second. It consistently placed at the top of our applications tests. A 32K-byte static cache and 82385 cache

controller boosted the T5200's performance above that of the competition.

The Micro Express Regal II topped our disk benchmark listing, a result corroborated by applications tests such as printing a PostScript file to disk, loading an extensive Lotus 1-2-3 spreadsheet, and storing large documents. The Dolch-P.A.C. also performed admirably on disk-intensive applications. It suffered somewhat on the low-level tests because we had to factor out the Hard Seek test. The P.A.C.'s SCSI connector hides low-level operations from the user. This makes a low-level Seek test useless. The Toshiba T5200 and the NEC ProSpeed also returned impressive disk results. The NEC PowerMate, as if not already hampered enough by the SX chip, suffers from a sluggish hard disk drive. It finished at the bottom of our low-level

Low-level | **Application-level**

	Application-level
Toshiba T5200	17.86
Dolch-P.A.C. 386-25	16.45
Micro Express Regal II	15.76
Compaq Portable 386 Model 40	13.33
IBM PS/2 Model P70 386	13.02
NEC ProSpeed 386	12.29
Toshiba T5100	11.04
Zenith TurbosPort 386	10.84
GRiDCase 1530	9.58
GRiDCase 1535 EXP	9.44
NEC PowerMate Portable SX	6.77

☐ CPU ☐ FPU ☐ Disk I/O ☐ Video

☐ Word processing ☐ Spreadsheet ☐ Database ☐ Scientific/engineering ☐ Compilers

The 20-MHz Toshiba T5200 outscored the 25-MHz Dolch-P.A.C. 386-25 on the applications index and the low-level CPU index. Cumulative indexes at right show relative performance: an 8-MHz IBM PC AT = 1. All low-level benchmarks use the 80386 version (1.1) of Small-C (32-bit integers). The P.A.C. finished highest on the FPU and video tests. The Micro Express Regal II had the best disk index. We blame slow CPUs for poor performance showings by the GRiDCases and the Zenith TurbosPort 386.

The shortest distance between Laptop and Desktop.

Need to move files, directories - even complete disks of information - between IBM-compatible 3 1/2" laptop and 5 1/4" desktop disk drives? Send them First Class. With FASTLYNX.

FASTLYNX is faster than any other product of its kind. Over 500K baud parallel. More than 200K baud serial. And only FASTLYNX supports both CRC and Checksum error-checking, for 100% accuracy at maximum speed.

No other file transfer utility is as easy to use as FASTLYNX. Just connect a cable from the 2 supplied, and decide which files to transfer. FASTLYNX connects itself and automatically selects port, baud rate, and error-checking mode.

FASTLYNX provides a unique self-cloning ability. So you only need to install FASTLYNX on a single system. A bootstrap upload feature automatically transfers FASTLYNX to the target computer.

FASTLYNX is the only file transfer system with 3 separate modes of operation. SPLIT SCREEN MODE makes file transfer as simple as point-and-

The FASTLOCK file protection system safeguards your hard disk files with simple, effective password security. Until October 31, 1989, participating dealers are offering special savings when you purchase FASTLOCK and FASTLYNX together.

shoot. Use FORM MODE to create time-saving file transfer macros. With COMMAND MODE, advanced users can issue file transfer commands from the DOS prompt.

FASTLYNX makes printer and disk sharing easy. Use the unique, ultra-fast parallel driver to share printers effortlessly. Or run remote programs with the speedy direct disk access you'd expect from a network. We've even included versatile hard disk management functions. And on-line and context-sensitive help makes you so comfortable with these features that you'll actually use them.

2 cables are better than 1. Unlike other file transfer products, FASTLYNX includes both parallel *and* serial cables for optimum convenience.

FASTLYNX culminates 5 years of file transfer leadership. Rupp Corporation built upon its highly acclaimed Fastwire II file transfer utility to make FASTLYNX the fastest, easiest to use, and most complete product of its kind.

Call **1-800-852-RUPP** for the name of your nearest FASTLYNX dealer.

RUPP
CORPORATION

835 Madison Avenue
New York, NY 10021
(212) 517-7775, Fax: (212) 249-8243

THINK BIG

Phar Lap
Virtual Memory Manager.

It will let you write applications up to 5, 10, 15 megabytes or more for any 386 PC running MS-DOS*. *Forget about RAM limitations.* **Your application can run on a machine with as little as 1 or 2 megabytes of memory.**

Only Phar Lap 386lVMM* gives you demand-paged virtual memory capability so you can write mainframe-sized applications for the PC. Applications your customers can run on their 386 PCs now with no additional memory. No kidding. All you need is 386lVMM and our family of 386 development tools. Existing programs developed with our 386lDOS-Extender can be easily expanded with 386lVMM too.

Our tools let you take full advantage of the 386 protected mode architecture. Break the DOS 640K limit in the language of your choice; C, Fortran, Pascal, or Assembler.

For fast compact code, use 386lASM, our 80386 assembler that's upwardly compatible with the MASM* 8086 assembler. Existing DOS and mainframe applications written in a high level language are easily ported by recompiling. And 386lLINK, our 32-bit native mode linker, puts it all together.

Debugging is made easy too. With our 386 symbolic debugger you can debug applications written in assembler or any high level language. Best of all, with Phar Lap's 386lDOS-Extender* you can run your native mode program on any 386-based PC running MS-DOS. And you have full access to DOS system services through INT 21.

NO COMPATIBILITY PROBLEMS

Phar Lap's tools are compatible with the industry's leading systems: DESKPRO 386*, IBM Model 70/80*, 386 clones and accelerator boards. Not only will your new applications be compatible with the leading systems, they'll run alongside all other DOS applications.

NO ROYALTY PAYMENTS

Once your 386 application is complete, all you pay is a low one-time fee to license 386lDOS-Extender for redistribution.

386lVMM is also developer friendly. Call to find out about our flexible run-time pricing.

You can unlock the entire DOS market now. Don't wait for OS/3.

$495   386lASM/LINK-Package includes 386 assembler, linker, MINIBUG debugger and the developer version of 386lDOS-Extender
$895   MetaWare 80386 High C* compiler
$595   MicroWay NDP Fortran-386* compiler
$195   386lDEBUG symbolic debugger
$295   386lVMM - developer version of the Phar Lap Virtual Memory Manager

(617) 661-1510
PHAR LAP SOFTWARE, INC.
60 Aberdeen Avenue, Cambridge, MA 02138
Fax: (617) 876-2972
"THE 80386 SOFTWARE EXPERTS"

disk benchmark and required 31 seconds to load the large spreadsheet, three times longer than it took the Regal II. Other file I/O applications were similarly slow.

The IBM PS/2 Model P70 had the lowest score of the 20-MHz CPUs tested, but that is no disgrace, given the quality of portables in this field. The P70 posted results in line with those of a desktop Model 70-121 (a 20-MHz 80386 machine), so you can now lug around the equivalent of a top-of-the-line PS/2.

Transferring the benchmark files uncovered another slight annoyance. Both the P70 and the Regal II lack a 9-pin serial port. Granted, adapters are easy enough to find, but didn't IBM introduce the 9-pin serial port? Now that the 9-pin connection has become a standard in the portable world, IBM has unveiled a portable without one. Go figure.

The NEC ProSpeed finished surprisingly strong, scoring higher than the 16-MHz Toshiba T5100. A fast hard disk drive helped the ProSpeed. The PowerMate scored credibly on our CPU benchmark, but it consistently placed at the bottom of our applications tests. The CPU works well, until you access the narrow data path. The Zenith TurbosPort 386 performed adequately for a machine with a 12-MHz CPU, while the GRiDCases' 12.5-MHz CPUs kept those two models at or near the bottom of every test.

King of the Road

When it comes to choosing the best portable computer in the bunch, we're hard-pressed to make a choice. All the portables we've tested run well; you just can't go wrong with any of them.

Our particular favorite was the Toshiba T5200. It's an excellent combination of computing power and portability all wrapped up in an attractive shell. Unfortunately, it's also one of the most expensive portables in the group.

At the opposite end of the economic scale is the Micro Express Regal II. Its price is an astoundingly low $2999, less than half the price of the Toshiba. It isn't at the top of the performance chart, but it finished third, well ahead of some pricey competition.

All 11 machines we've looked at prove that a well-equipped portable computer can match the capabilities of many desktop models. A portable can now be *the* computer on your desk. ■

Stanford Diehl and Stan Wszola are testing editors for the BYTE Lab. They can be reached on BIX as "sdiehl" and "stan," respectively.

Circle 191 on Reader Service Card

World Radio History

We keep telling people this is <u>not</u> a laser printer.

The new $995 HP DeskJet PLUS Printer.

But they keep looking at the evidence.

You won't believe your eyes (or ears), either. The HP DeskJet PLUS printer gives you the same crisp, black lettering. Clean, sharp graphics. And whisper-quiet operation. But it uses advanced inkjet technology to supply these laser-like qualities for the price of a 24-wire printer.

For just $995, you get 300 dpi in a trim 15-pound package that's twice as fast as the original. What's more, it has built-in landscape and improved font selection. Including ten built-in fonts and over 100 optional fonts, with sizes up to 30 points. And its 20,000-hour MTBF

assures a long, happy life.

So call **1-800-752-0900, Ext. 276J** for the name of your nearest authorized HP dealer. Then judge the HP DeskJet PLUS printer for yourself.

HEWLETT PACKARD

The Painlessly Portable PC

NEC packs power into
a seductive little
package, but for a price

Mark L. Van Name
and Bill Catchings

The NEC UltraLite weighs in at only 4⅖ pounds but offers greater performance than an IBM PC XT.

NEC's new UltraLite is just plain cute. Its sleek black finish looks modern, while in size and shape it resembles a large hardcover book—and it's not much harder to carry. The 4⅖-pound UltraLite fits into any briefcase, or you can use its optional carrying case. (See photo.)

You pay a pretty price, however, for this tiny MS-DOS machine: The basic UltraLite runs $2999, while our fully loaded evaluation unit (with $129 carrying case) would set you back $4526. Since a comparably equipped Toshiba T1000 lists for around $1700, we're talking about a stiff price premium.

Taking It on the Road

The first things to check on any portable computer are its screen and its keyboard. The UltraLite scores well on both fronts.

The screen is a backlit, supertwist, electroluminescent, blue-on-white display with a full 25 rows by 80 columns in a 9½-inch-diagonal viewing area. It emulates IBM's CGA, with seven gray scales instead of colors.

The screen is nearly twice as wide (8¼ inches) as it is tall (4¼ inches), so its aspect ratio is a bit distorted (e.g., circles appear as ovals). The display phosphors decay so slowly that the screen is almost impossible to read when it scrolls. And it does not fare well in bright light or glare.

The keyboard, like the screen, is adequate. Its keys are full size with an audi-ble keyclick. While the keys travel only 2 millimeters (more than 1 mm less than the keys on most conventional keyboards), you can feel them spring back.

The 78 keys are arranged well, with a row of function keys across the top and the numeric keypad overlaid onto other keys. As with most portables, you get to the numeric keypad and other special keys by using an Fn key. Unfortunately, NEC followed the IBM Enhanced keyboard and put the Caps Lock key where our fingers expect the Control key to be.

Overall, we give the screen and keyboard a high B: not great, but adequate.

Storage Options

The next priority with any portable computer is its disk storage. The UltraLite of-fers several different disk options.

Its main working storage is a "silicon hard disk drive," a battery-backed 1 or 2 megabytes of DRAM with firmware that superbly emulates a hard disk drive. This drive acts as the UltraLite's C drive. It appears to have 58 cylinders, four heads, and 17 sectors per track. Such disk utilities as the Norton utilities have no trouble recognizing it as a hard disk drive.

It's extremely fast (average access time is 9 milliseconds), and it makes the entire machine feel quicker than you would expect. While even 2 megabytes is not a lot of disk space, it's enough for one or two applications and their data.

Of course, you have to be able to load and back up the silicon hard disk drive.

continued

NEC offers four ways to get those jobs done. The simplest is the optional external 1.44-megabyte 3½-inch disk drive, which NEC calls its FDD-BOX ($399). The FDD-BOX is 4¼ inches wide, 6⅞ inches deep, and 2 inches high, and it weighs a little under 2 pounds. It connects to a 68-pin NEC-proprietary external connector on the rear of the UltraLite and appears as the machine's A drive.

Unfortunately, the FDD-BOX draws its considerable power from the Ultra-Lite. Therefore, you can use the drive only when the UltraLite is running off the standard-equipment AC power adapter that plugs into a mini-DIN-8 connector on the rear of the machine.

Inside the FDD-BOX is a TEAC floppy disk drive. Under that drive is an NEC 4- by 6-inch floppy disk drive controller that uses Western Digital's controller chip and about 16 support chips. The FDD-BOX also has a standard female DB-25 parallel connector on its rear, to which you could attach a printer.

We don't see how anyone would want to live with any machine without a disk drive, but it is possible. For one thing, you can transfer files to another machine. NEC includes Traveling Software's LapLink program in the UltraLite's ROM. (The ROM appears as the UltraLite's D drive.) Because you also need a copy of LapLink on the other machine, the UltraLite comes with both 3½-inch and 5¼-inch LapLink disks.

You can hook up to another machine with the null-modem cable that is included. This cable has a female DB-25 serial connector that hooks up to the second machine. The connector appears to the UltraLite as its COM1 serial port. Unfortunately, you have to use that specific cable, because the UltraLite's RS-232C connector is a nonstandard mini-DIN-9 jack. While there seems to be enough room for a standard DB-9 connector, an NEC spokesperson said that the firm chose the smaller nonstandard connector to save space.

While the disk drive and LapLink are the UltraLite's two main links to the outside world, NEC offers two others that involve a tiny expansion slot under a cover on the right side of the unit. This slot accepts RAM and ROM cards that are the width and length of credit cards but about twice as thick. NEC offers both 256K-byte battery-backed RAM cards ($299 each) and ROM cards. Both types appear to the UltraLite as its B drive. You can pull these cards out and insert new ones while the machine is running, as you would with floppy disks. The 256K-byte RAM card uses a replaceable 3-volt lithium battery that NEC claims is good for up to 6 months.

NEC says that it will offer both 512K-byte and 1-megabyte ROM cards containing such applications as Lotus 1-2-3 and WordPerfect. NEC was unable to furnish us with any ROM cards by our deadline, however, and only time will tell how many companies will produce software on this nonstandard medium. (An NEC spokesperson said that the firm was to begin shipping ROM cards in June and estimated that they will cost roughly the normal price of the software they contain plus $50 for the card itself.)

The UltraLite also includes an internal 2400-bps Hayes-compatible modem hooked to its COM2 serial port. There are two standard RJ-11 connectors on the rear of the machine with which you can link the modem to a phone jack and a telephone. The UltraLite includes cables for both connections.

Compatibility and Performance
The UltraLite ran everything we threw at it, including Borland's Quattro 1.0, Reflex 1.14, SideKick 1.56a, SuperKey 1.16a, Turbo Basic 1.1, Turbo C 2.0, and Turbo Pascal 4.0; Digitalk's Smalltalk/V 1.2; Kermit 2.30; MicroPro's WordStar 3.3 and 4.0; Microsoft's PC Paintbrush 2.0 and Word 4.0; the Norton Utilities 3.00; Quarterdeck's DESQview 2.0; and Symantec's Q&A 1.0.

The UltraLite supports these applications with a subset of MS-DOS 3.3 that is built into its D-drive ROM. That ROM also contains Microsoft's MS-DOS Manager 2.0, a good but not outstanding DOS shell, and a SETUP program. The system boots into the MS-DOS Manager by default, but you can have it go straight to MS-DOS by changing a line in the standard AUTOEXEC.BAT file.

SETUP runs as a TSR program. With it you can choose the boot disk, set the CPU speed, and change several screen options, the most interesting of which is a color palette that lets you determine how the UltraLite maps the 16 possible CGA colors to its seven gray shades.

While the UltraLite is in many ways a portable XT compatible, it offers far better than IBM XT performance. Its extremely fast silicon hard disk drive helps a lot. Its CPU, a 9.83-MHz NEC V30 with a compatibility speed of 4.92 MHz, also performs well; the UltraLite's 640K bytes of system RAM uses eight 1-megabit 120-nanosecond DRAM chips that let the V30 run with no wait states. The combination gives the machine the feel of an AT with a very fast hard disk drive.

The BYTE disk I/O benchmarks did show an anomaly on the 32-sector DOS Seek test. The time of 90.40 seconds does not match up with the single-sector time of 3.48 seconds. A spokesperson for NEC suggested that the result could have occurred because the UltraLite emulates a hard disk drive controller using software that calculates a checksum for each sector. The algorithm for generating these checksums may have been at fault. In our side-by-side comparisons with

NEC UltraLite

APPLICATION-LEVEL PERFORMANCE

WORD PROCESSING	
XyWrite III + 3.52	Medium/Large
Load (large)	:22
Word count	08/1:00
Search/replace	.14/:49
End of document	:04/:26
Block move	:22/:22
Spelling check	:25/3:09
Microsoft Word 4.0	
Forward delete	:43
Aldus PageMaker 1.0a*	
☐ Index:	1.00

SPREADSHEET	
Lotus 1-2-3 2.01	
Block copy	.10
Recalc	:04
Load Monte Carlo	:41
Recalc Monte Carlo	.20
Load rlarge3	:10
Recalc rlarge3	:03
Recalc Goal-seek	:12
Microsoft Excel 2.0	
Fill right	.13
Undo fill	6:27
Recalc	:05
Load rlarge3	1:04
Recalc rlarge3	:04
☐ Index:	0.90

DATABASE	
dBASE III + 1.1*	
☐ Index:	N/A

SCIENTIFIC/ENGINEERING	
AutoCAD 2.52	
Load SoftWest	6.35
Regen SoftWest	6.15
Load StPauls	1:53
Regen StPauls	1:43
Hide/redraw	85.43
STATA 1.5	
Graphics	2.38
ANOVA	1:49
MathCAD 2.0	
IFS 800 pts	3.34
FFT/IFFT 1024 pts.	4.15
☐ Index:	0.35

COMPILERS	
Microsoft C 5.0	
XLisp compile	10.34
Turbo Pascal 4.0	
Pascal S compile	.10
☐ Index:	0.99

All times are in minutes seconds. Indexes show relative performance, for all indexes, an 8 MHz IBM PC AT = 1

*Due to the limited space on the UltraLite's silicon hard disk drive, we were unable to run every application test of the BYTE benchmarks. Tests using Aldus PageMaker and dBASE III Plus were omitted. We also omitted the results of those tests for the systems used for comparison and adjusted their application indexes accordingly.

NEC UltraLite **3.2***

System	Index
Zenith SupersPort 286	5.9
Epson Equity LT	3.0
IBM PC AT	4.0

Legend:
- Word Processing
- Spreadsheet
- Database
- Scientific/ Engineering
- Compilers

LOW-LEVEL PERFORMANCE[1]

NEC UltraLite

CPU		DISK I/O		VIDEO	
Matrix	15.75	**Hard Seek[3]**		**Text**	
String Move		Outer track	1.36	Mode 0	14.83
Byte-wide	86.53	Inner track	1.34	Mode 1	14.83
Word-wide:		Half platter	1.28	Mode 2	14.76
Odd-bnd	86.49	Full platter	1.32	Mode 3	14.76
Even-bnd.	43.32	Average	1.33	Mode 7	N/A
Sieve	83.51	**DOS Seek**		**Graphics**	
Sort	67.04	1-sector	3.48	CGA:	
		32-sector	90.40	Mode 4	6.16
☐ Index:	0.93	**File I/O[4]**		Mode 5	6.19
		Seek	0.40	Mode 6	6.46
FLOATING POINT		Read	0.93	EGA	
Math	N/A	Write	0.96	Mode 13	N/A
Error[2]		**1-megabyte**		Mode 14	N/A
Sine(x)	N/A	Write	11.70	Mode 15	N/A
Error		Read	11.29	Mode 16	N/A
e^x	N/A			VGA:	
Error		☐ Index:	1.42	Mode 18	N/A
				Mode 19	N/A
☐ Index:	N/A			Hercules	N/A
				☐ Index:	0.80

N/A = Not applicable

[1] All times are in seconds. Figures were generated using the 8088/8086 version (1.1) of Small-C
[2] The errors for Floating Point indicate the difference between expected and actual values, correct to 10 digits or rounded to 2 digits
[3] Times reported by the Hard Seek and DOS Seek are for multiple seek operations (number of seeks performed currently set to 100)
[4] Read and write times for File I/O are in seconds per 64K bytes
[5] For the Livermore Loops and Dhrystone tests only, higher numbers mean faster performance

CONVENTIONAL BENCHMARKS

LINPACK	7154.48
Livermore Loops[5] (MFLOPS)	0.00
Dhrystone (MS C 5.0) (Dhry/sec)	1422

System	
Zenith SupersPort 286	
Epson Equity LT	
IBM PC AT	

Legend:
- CPU
- FPU
- Disk I/O
- Video

For a full description of all the benchmarks see Introducing the New BYTE Benchmarks. June 1988 BYTE

conventional hard disk drives, however, the UltraLite's silicon hard disk drive performed well and showed no signs of a bottleneck.

Interestingly, the V30 is socketed on the motherboard, so perhaps NEC plans some future CPU upgrade. However, the UltraLite has no coprocessor socket.

The UltraLite's ROM BIOS is Phoenix Technologies' version 2.52, so the machine should not have any compatibility problems there.

Our Microsoft Serial Mouse worked fine with the unit, although we needed a DB-9-to-DB-25 adapter to connect it to the machine's null-modem cable.

Nice, but for How Long?

The final key aspect of any battery-driven portable computer is its battery life—how long you can go between recharges. With the UltraLite, that's actually a two-part issue.

The main system power comes from a rechargeable nickel-cadmium "battery" that is actually a collection of seven Molicel nickel-cadmium batteries wired together. This battery is supposedly good for about 2 hours with the basic UltraLite. On our fully loaded evaluation unit, however, we got only 1 hour and 25 minutes of constant use before the first warning message. Fortunately, that message and the accompanying low-battery indicator light come about 5 minutes before you actually run out of power, so you can save your work.

The other half of the UltraLite's battery picture is a separate rechargeable nickel-cadmium battery that supports the silicon hard disk drive. It's good for five to seven days between charges.

The UltraLite's batteries are one of the machine's biggest frustrations. An hour and a half just isn't enough time. We could live with it, however, if we could just carry a few spare batteries. There are even openings on the bottom of the unit for both batteries. Unfortunately, you can't replace the batteries yourself. This is clearly an area for NEC to improve in the next UltraLite.

A Peek Inside

To get a better look inside, you can take the UltraLite apart. We did, but you definitely need to be careful. This little wonder is packed tightly.

Its motherboard sits under the ultra-thin keyboard. The board is essentially the same width and length as the case, with cutouts for a system ROM card and the silicon hard disk drive card, both of which you can reach via covers on the bottom. The board has fewer than 30 chips, including the eight memory chips. Fewer than a dozen chips do most of the work. In fact, on this board the analog devices and support parts (such as capacitors) almost outnumber the digital parts.

Documentation and Support

The UltraLite's three manuals are all useful and very well written. The quick-start manual is a model for books of its kind. Even novice users can follow its clear instructions easily, and it uses pictures frequently and effectively. The portable guide, which is also very well done, contains most of the data you need on the road.

The comprehensive user's manual is a thorough reference guide to all of the system's capabilities. It does not, however, contain a complete MS-DOS reference section; if you want that, the manuals suggest that you buy an MS-DOS book.

If you need help, you can call NEC technical support toll-free. Unfortunately, while we found the technical-support number in several NEC ads, it was not in the manuals. (An NEC spokesperson said that this problem has been corrected in subsequent printings of the manuals.) The warranty information in the comprehensive user's manual listed a number to call with problems other than repairs, and you can get from that number to technical support. The technical-support people with whom we talked were friendly and helpful. Expect to wait for them, however, as NEC's lines were usually busy when we called.

The UltraLite includes a one-year warranty, and you can buy up to three additional years of protection for $300 per year for a unit with a 1-megabyte silicon hard disk drive, or $370 per year for the 2-megabyte version.

Sleek and Expensive

The UltraLite clearly defines a new size standard for portable MS-DOS computers; no other machine comes close. It looks great and doesn't weigh much.

Unfortunately, it is also extremely expensive. You can buy some 80386 portables for almost the same money. In the end, you have to decide how you want to look at it: Are you getting a lot of computer for the size, or not enough computer for the money? We don't know how you would vote, but we both wish that we could afford an UltraLite. ∎

Mark L. Van Name and Bill Catchings are independent consultants and freelance writers based in Raleigh, North Carolina. They can be reached on BIX as "mvanname" and "wbc3," respectively.

Our Swans make waves.

The Swan line of computers is creating waves of excitement among computer users. One reason is price. Swan computers are true bargains, offering you the most for your investment. But price alone doesn't tell the complete value story.

Swan computers are solid, reliable performers...assembled by Tussey Computer Products, a solid, reliable company with a tradition of unsurpassed customer satisfaction.

Swan computers are highly compatible. In fact, all Swan computers offer the top-rated Phoenix ROM BIOS to assure 100% PC compatibility.

And the benefits keep coming, wave after wave...

- **30-day satisfaction guarantee**
- **Toll-free technical support**
- **Full 1-year warranty**
- **Fast, sure delivery**
- **Sorbus℠ on-site service option**

Swan offers a full line of services tailored to the needs of educational and corporate computer users. Call our Educational & Corporate Sales Department toll-free at 1-800-468-9044. Learn how a Swan can be a business animal ...or a teacher's pet.

With their compatibility, reliability, performance and value, Swans are making waves in the computer industry. Waves that may make some competitors run for safer ground.

Swan XT10

The perfect budget-priced home computer. A solid, reliable, compatible performer.

$549

Base System with Single Floppy, No Video

XT10	VIDEO OPTIONS Include Monitor & Video Adapter			
Drive Options	Mono'	CGA	EGA	VGA
Single Floppy	$699	$869	$1098	$1198
Dual Floppies	$779	$949	$1178	$1278
*w/20MB (65ms) Hard Drive	$979	$1149	$1378	$1478
*w/32MB (40ms) Hard Drive	$979	$1149	$1378	$1478

* Includes Single 360K Floppy Drive
' Upgrade from 12" to 14" Flat Screen, add $49

Swan 286/12

You probably never thought you could get into a reliable business system this inexpensively. But with the 286/12, you can. Now.

$999

Base System with Single Floppy, No Video

286/12	VIDEO OPTIONS Include Monitor & Video Adapter			
Drive Options	Mono	14"Flat Mono	EGA	VGA
w/32MB (40ms) & 1:1 Interleave	$1399	$1448	$1798	$1898
w/40MB (28ms) & 1:1 Interleave	$1499	$1548	$1898	$1998
w/80MB (28ms) & 1:1 Interleave	$1849	$1898	$2248	$2348

* Items returned must be as-new, without modification or damage. Sorry ... shipping charges and upgrade fees are not refundable.

** Custom configured systems may take longer.

Circle 259 on Reader Service Card

Swan 386SX

This system gives you 386 compatibility, power and performance...at a 286 price.

$1499

Base System with Single Floppy, No Video

386SX	VIDEO OPTIONS Include Monitor & Video Adapter			
Drive Options	Mono	14"Flat Mono	VGA Mono	VGA Color
w/32MB (40ms) & 1:1 Interleave	$1899	$1948	$2098	$2398
w/40MB (28ms) & 1:1 Interleave	$1999	$2048	$2198	$2498
w/80MB (28ms) & 1:1 Interleave	$2349	$2398	$2548	$2848
w/150MB (17ms) ESDI w/1:1 Interleave	$2994	$3043	$3193	$3493

Upgrade to 2MB only $396

Standard Features	XT10	286/12	386SX	386/20D
CPU	8088-1	80286	80386SX	80386
Speed (MHz)	10/4.77	12/6	16/8	20/8
BIOS	Phoenix	Phoenix	C & T	Phoenix
Norton SI Rating v 4.0	2.1	11.2	17.6	22.0
Wait State	0	0	0	0
Standard Memory	640K	512K	1MB	1MB
Shadow BIOS	—	—	—	384K
Memory Upgrades	—	1MB	1.5/2/3/5/6/8MB	2/4/8/10/16MB
Coprocessor Support	8087	80287	80387SX	80287 or 80387
Expansion Slots: 32-bit	—	—	—	1
Expansion Slots: 16-bit	—	6	6	4
Expansion Slots: 8-bit	8	2	2	3
Dual HD/FD Controller	w/HD systems	Yes	Yes	Yes
Device Bays (Exposed/Internal)	2 Exp./2 Int.	3 Exp./2 Int.	3 Exp./1 Int.	3 Exp./2 Int.
Serial Ports	1	1	1	2
Additional I/O Ports	1 Parallel and 1 Game Port			
Power Supply	150W	200W	200W	200W
Keyboard	101 Key	101 Key	101 Key	101 Key
Additional Features	Clock Calendar w/Battery Backup and Setup/Utilities Disk			
OPTIONS:	XT10	286/12	386SX	386/20D
3.5" Floppy Drive	Optional	Your choice	Your choice	Your choice
Swan/MS-DOS & GW BASIC add $89				

Order Now Toll-Free
1-800-468-9044

FAX:814-237-4450 • International:814-234-2236

To order: No surcharge on Discover, Visa, MasterCard or AMEX
• Your credit card is not charged until your order is shipped.
Shipping: 3% or $5 minimum for UPS Ground. Call for shipping charges on Express Air, APO, FPO, AK, HI and all foreign orders.
• Allow 2 weeks for personal and company checks to clear.
• Defective items replaced or repaired at our discretion. PA deliveries add 6% sales tax. Prices and terms subject to change without notice. **Open:** 8a.m.-11p.m. M-F, 10p.m.-8p.m. Sat, 12p.m.-8p.m. Sun

Swan 386/20D

"Tussey's Swan 386/20 flies."
— *Computer Shopper, April 1989*

Gracefully combining power and performance, this Swan has soared to the upper limits of today's technology.

$1999

Base System with Single Floppy, No Video

386/20D	VIDEO OPTIONS Include Monitor & Video Adapter			
Drive Options	Mono	14"Flat Mono	VGA Mono	VGA Color
w/40MB (28ms) & 1:1 Interleave	$2549	$2598	$2748	$3048
w/80MB (28ms) & 1:1 Interleave	$2849	$2898	$3048	$3348
w/150MB (17ms) ESDI w/1:1 Interleave	$3494	$3543	$3693	$3993

Upgrade to 2MB only $436

Swan TECHNOLOGIES *a division of* **tcp**

TUSSEY COMPUTER PRODUCTS
3075 RESEARCH DRIVE • STATE COLLEGE, PA • 16801

Because Resources Should Be Shared

BayTech engineers resource sharing solutions.

Because we know your applications for sharing printers, plotters, modems and data are unique, BayTech has developed over 30 resource sharing products, which offer you a broad range of solutions for sharing your resources.

We realize that you shouldn't have to fit your application to the specifications of a single product. With BayTech, you select the

model that meets *your* needs.

From simply sharing one or more printers between computers, to creating a complete network of computers, printers, plotters and modems, BayTech has a product designed for you.

Call us toll-free today. Our technical support staff will show you how to make the most of your resources.

Circle 35 on Reader Service Card (DEALERS: 36)

Ultra Graphics

The Ultra Clipper
UM1280, Pixelworks'
bus-mastering
coprocessor, brings
enhanced graphics
to MCA-compatible
computers

Bradley Dyck Kliewer

Pixelworks has taken a big step in the Micro Channel market by becoming one of the first companies to introduce a high-resolution (1280- by 1024-pixel) bus-mastering graphics controller for the PS/2s, the Ultra Clipper UM1280. As a bus-mastering device, the Ultra Clipper can take temporary control of the bus and directly transfer data without intervention from the CPU. This improves system throughput by leaving the CPU free for other tasks.

The Ultra Clipper requires a Micro Channel–compatible bus with two empty slots. Depending on the configuration, you need either a 30- to 64-kHz multifrequency monitor (when using a single monitor for both the VGA and Ultra Clipper high-resolution display) or separate VGA and 64-kHz monitors.

The board is available in two configurations: a 4-bit, 16-color version ($2895), and an 8-bit, 256-color version ($4095). Each includes a Texas Instruments 320C25 chip—a more general-purpose processor than the 34010 and better suited to the programming needs

of the Ultra Clipper. The 8-bit board contains 1.25 megabytes of display RAM—just enough to hold the graphics data for a 1280- by 1024-pixel by 8-color display (the 4-bit-plane model has half as much memory).

The Ultra Clipper is clearly targeted for CAD markets; the only drivers supplied with my board were for AutoCAD (ADI drivers for releases 9 and 10), VersaCAD, and Bentley Microstation.

Dual-Slot Configuration
Installing the Ultra Clipper is fairly easy. The board includes two switch blocks for specifying monitor types, and it requires two adjacent slots (any two Micro Channel slots will work, including mixed slots—one 16-bit and one 32-bit). The installation program, which you copy onto the IBM installation disk, performed flawlessly. You can use the Ultra Clipper in a two-monitor system (the best option), or you can install a pass-through cable between the VGA output and the Ultra Clipper to use a single monitor.

I tested the 8-bit-plane (256-color) Ultra Clipper on a 20-MHz IBM PS/2 Model 70-121 running DOS 3.3 with a Mitsubishi HL 6905 multifrequency monitor. When I configured the board as a single-monitor system, the Model 70 failed its self test, generating a video adapter error. But when I bypassed the error message, the system worked perfectly. IBM has modified the monitor detection routines in some PS/2 models, and the system wasn't detecting a valid monitor type (the VGA output connects to the pass-through adapter instead of to the monitor). Pixelworks sent a new pass-through daughterboard that changed the resistance slightly, solving the problem.

Unlike graphics coprocessors that are I/O-mapped, the Ultra Clipper can work in either I/O- or memory-mapped mode. In I/O-mapped applications, all the drawing commands pass through the I/O
continued

registers, which typically run at a slower speed than memory. In memory-mapped mode, the PC stores lists of drawing commands and parameters in a memory buffer; the Ultra Clipper then uses its bus-mastering capability to process the display list directly from PC memory. The processing benefits both from the faster speed of memory and from the lack of system CPU overhead.

When using memory buffers, the board further improves performance through double buffering. The application reserves two areas in system memory to act as command buffers. While the CPU writes commands to one memory buffer, the Ultra Clipper reads commands from the other. The Pixelworks

PHLIP (Pixelworks High-Level Interactive Protocol) libraries, which are available for several C compilers, provide functions that make writing double-buffered programs simple. I ran tests in both unbuffered and buffered modes (the buffered modes used two 2K-byte buffers); the speed difference can be dramatic.

The Ultra Clipper's graphics commands create an entire programming language of sorts, supporting NOOPs (no operations), rudimentary flow control, and comments. The graphics instructions are rich: In addition to the

basic points, polylines, polygons, and text, the Ultra Clipper supports circles, ellipses, arcs, and Bézier curves. You can also specify points, polylines, polygons, and Bézier curves in three-dimensional coordinates, and there are functions for modifying the view. The Ultra Clipper also provides BitBlt functions for transferring data between the system and adapter memory or for copying data from one screen area to another.

Programming the Ultra Clipper is simple. The PHLIP documentation (available on request) includes program fragments with most of the function-call descriptions, and the descriptions are fairly complete. The documentation includes several complete program listings that are useful for learning PHLIP.

Performing Arts

I adapted the test programs from the BYTE graphics benchmarks and compiled them with Microsoft C 5.1. The programs are similar to those used for my review of the IBM 8514/A and Artist 10 MC graphics boards ("Pixels on the March," January BYTE). To test memory transfer speeds, I also adapted the BITBLT program from the review "Debunking 16-bit VGA" (June BYTE).

The test results, which appear in tables 1 and 2, were remarkable. The detail available in AutoCAD is impressive, and the proportionally spaced roman font used for the menus is attractive. For graphics-intensive work, the Ultra Clipper outperforms both the IBM 8514/A and the Artist 10 MC. And, as with any graphics coprocessor, drawing commands execute hundreds of times faster than standard graphics adapters such as EGA and VGA. AutoCAD times are about the same for VGA and Ultra Clipper, but the Ultra Clipper displays over

Table 1: *BYTE benchmark results for the Ultra Clipper UM1280. Benchmark times are in minutes:seconds.*

	AutoCAD
SoftWest	
Load and draw	1:51 5
Regen	0:34.8
Redraw	0:02 1
StPauls	
Load and draw	0:20.3
Regen	0:06.1

Table 2: *BITBLT test results. Times are in seconds. BITBLT copies an 8- or 16-pixel by 1-line block from one screen area to another. The BITBLTP tests copy the same information from system memory to the screen. The BITBLT2 tests perform the same functions using an 8- by 8-pixel or 16- by 16-pixel block. Note the faster times for the BITBLT2 tests, which require fewer instructions to fill the screen.*

	Unbuffered	Buffered
BITBLT-8	45.14	24.38
BITBLTP-8	52.89	31.47
BITBLT2-8	11.15	8.46
BITBLT2P-8	28.78	16.40
BITBLT-16	23.46	13.13
BITBLTP-16	27.68	16.97
BITBLT2-16	5.27	4.50
BITBLT2P-16	9.50	8.68

four times as many pixels, and when you remove system overhead (using RE-DRAW instead of REGEN), the Ultra Clipper is over twice as fast.

Ultra Clipper's primary disadvantage is its physical design. It uses several daughtercards (looking down on the card, there are three layers of boards), which requires two slots in the host computer. For users of the Model 70 or Model 50, this leaves only one slot free. The monitor connection uses a 9-pin D-shell instead of the 15-pin D-shell that IBM uses for PS/2 video connections (a 15-pin connector is too thick to fit between the sandwiched boards). And on single-monitor configurations, the pass-through cable adds an extra cord outside the system unit.

Using a pass-through cable is an understandable design decision. It simplifies the design of the Ultra Clipper, since Pixelworks doesn't need to replicate the VGA D/A converter, which could potentially create minor incompatibilities with the VGA. Also, a video extension would limit the Ultra Clipper to one position in the machine—not a desirable option when the board requires two slots.

A Niche Fit
On Micro Channel systems that have several free slots, I wouldn't hesitate to recommend the Ultra Clipper. But on smaller systems, such as the Model 70, consider your expansion needs carefully. Adding a single card would fill the system. I typically add either a tape backup or a network card, which would leave no room for additional memory (not as big a concern on the Model 70, which can take up to 8 megabytes on the system board) or other options.

The Ultra Clipper is far superior to high-resolution VGA. It's also competitively priced with other graphics coprocessor boards, with the exception of the IBM 8514/A, which costs about $2700 less. However, the Ultra Clipper has some advantages over the 8514/A. Its 1280- by 1024-pixel noninterlaced display is much sharper. Also, the board's bus-mastering capabilities and its extensive drawing primitives give the Ultra Clipper a potential performance advantage over other coprocessor boards when it's used with CAD packages that support the board's hardware-level primitives (currently, PCAD, for which Pixelworks is developing drivers, is the only CAD package that supports advanced hardware primitives).

Advances in software always lag behind hardware, so you probably won't see commercial software supporting the

Ultra Clipper UM1280

Company
Pixelworks, Inc.
7 Park Ave.
Hudson, NH 03051
(800) 247-2476
(603) 880-1322

Features
Maximum 1280- by 1024-pixel resolution, noninterlaced; 16 colors from a palette of 4.096 million (4-bit board), or 256 colors from a palette of 16.7 million (8-bit board); three-dimensional support, including real-time pan and zoom; driver software for AutoCAD, VersaCAD, and Bentley Microstation CAD programs

Size
Full-length (requires 2 slots)

IBM PS/2 Micro Channel or compatible computer; 64-kHz monitor and VGA-compatible monitor (for dual-monitor system), or 30- to 64-kHz multifrequency monitor for VGA pass-through (for single-monitor system)

Software Needed
DOS 3.3 or higher

Documentation
Installation manual, PHLIP programmer's guide (available on request)

Price
4-bit (16 colors): $2895
8-bit (256 colors): $4095

Inquiry 852.

board's more advanced features for some time. But if you write your own graphics routines for custom applications, the Ultra Clipper is a terrific combination of high resolution, powerful graphics primitives, and easy-to-program routines. ∎

Editor's note: *The test programs used for this review are available on BIX as Clipper.ARC. They're also available in a va-*

riety of other formats. See page 5 for details.

Bradley Dyck Kliewer is the author of EGA/VGA: A Programmer's Reference Guide *(McGraw-Hill, 1988) and principal of DK Micro Consultants, a microcomputer consulting business in Bloomington, Indiana. You can reach him on BIX as "bkliewer."*

Modula-2 and OS/2 Join Forces

Three Modula-2 compilers take advantage of OS/2's features

Andrew Schulman

Although Microsoft endorses C as the language of choice for OS/2 development, Niklaus Wirth's Modula-2 is better suited to exploit OS/2's multitasking and dynamic linking capabilities. At the core of Modula-2 are the twin concepts of *process* and *module*; these conveniently map onto OS/2 threads and dynamic-link libraries (DLLs).

I'll look at three implementations of Modula-2 for OS/2: Logitech's Modula OS/2 1.00, TopSpeed Modula-2 OS/2 1.20 from Jensen and Partners International (JPI), and Stony Brook's Professional Modula-2 2.0. I'll focus on the OS/2-specific features of each compiler: support for the OS/2 API (Application Programming Interface), PM (Presentation Manager) applications, DLLs, virtual memory, run-time error checking, and protected-mode debugging.

The OS/2 API, Modula-2-Style

All three compilers provide API bindings in the form of Modula-2 definition modules (.DEF files). Both Stony Brook and JPI divide the API .DEF files along the lines that OS/2 itself uses: DOS (the OS/2 kernel), VIO, KBD, and MOU. Logitech uses a different scheme, breaking the API down into smaller modules, such as DynLink, FileIO, MemManager, MultiTasking, and Misc.

The OS/2 API (unlike the MS-DOS programming interface) was designed to be used from high-level languages.

Thus, you can use the OS/2 procedures, constants, and data structures exported from the three vendors' API modules just as you would any other Modula-2 exports. Of course, programs using these features will only run under OS/2.

Say you wanted to use the OS/2 Dos-GetInfoSeg routine. DOS.DEF exports this in JPI, DOSCALLS.DEF in Stony Brook, and MISC.DEF in Logitech. You could either IMPORT the specific routine from the appropriate module (e.g., FROM MISC IMPORT DosGetInfoSeg) or you could IMPORT the entire module (e.g., IMPORT MISC) and then call the routine using its qualified name (e.g., MISC.DosGetInfoSeg).

When you call DosGetInfoSeg, by whatever means, you pass two segment selectors—one for the global-information segment, one for the local-information segment. Although selectors are words (2-byte quantities), you can't, in C, simply pass DosGetInfoSeg a pair of selectors by value. OS/2 expects references, so you must either pass pointers to selectors or, alternatively, declare selectors but pass their addresses. In Modula-2, though, VAR parameters are automatically passed by reference. You declare a selector and pass it to OS/2. This method simplifies working with the OS/2 API.

You need to turn the selector that Dos-GetInfoSeg passes into a 4-byte far pointer that you can use to access the values stored in the information segment. In Microsoft C, a utility macro, MAKEP, performs the necessary conversion. None of the Modula-2 environments provide an equivalent function, and you can't just use a Modula-2 absolute variable to point to the information segment—its selector isn't available until run time. Both Stony Brook and Logitech require you (and JPI permits you) to put the selector in the segment slot of an address variable and then assign the variable to a long pointer. The address type in Modula-2 is implementation-dependent; all three compilers tailor it to work

with selectors. JPI also provides a powerful, although nonportable, extension that simplifies this process. You can use a *pointer constructor* to cast an address (given in segment:offset form) directly to a far pointer.

JPI and Logitech supply header files (Modula-2 .DEF files) for the enormous PM API; Stony Brook doesn't yet, but expects to in a forthcoming release. JPI and Stony Brook also provide stand-alone graphics modules that don't rely on PM. Because multiple threads can call these graphics modules simultaneously, you can use the Stony Brook and JPI compilers to easily produce multitasking graphics programs. To demonstrate its graphics module, Stony Brook includes a Paint program—written in 500 lines of Modula-2—that's virtually identical to the Paint program that comes with the MS-DOS version of Stony Brook's QuickMod.

All three compilers work well with the OS/2 API—because Modula-2 helps the programmer localize the parts of a program that import and use OS/2's services. You can do the same thing in C, but you have to work at it. Where C merely permits modularity, Modula-2 encourages it.

Building DLLs

Modula-2, like Ada, specializes in the construction of abstract data types (ADTs) and the definition of operations on those types. In an ADT, a Modula-2 definition module exports an opaque type and a set of operations; the internal data representation is not exported. OS/2's DLLs are a perfect vehicle for implementing Modula-2 abstract data types. Because a DLL can export code but no data, implementing a module as a DLL guarantees a functional interface between the module and its clients.

DLLs are tricky to build—in part, because they must be reentrant. OS/2's own DLLs (such as VIOCALLS.DLL),
continued

	Logitech Modula OS/2 1.00	Stony Brook Professional Modula-2 2.0	TopSpeed Modula-2 OS/2 1.20
Company	Logitech, Inc. 6505 Kaiser Dr. Fremont, CA 94555 (800) 231-7717	Stony Brook Software, Inc. 187 East Wilbur Rd., Suite 9 Thousand Oaks, CA 91360 (800) 624-7487 (805) 496-5837	Jensen and Partners International, Inc. 1101 San Antonio Rd., Suite 301 Mountain View, CA 94043 (800) 543-5202
Hardware Needed	OS/2-capable PC compatible	OS/2-capable PC compatible	OS/2-capable PC compatible
Software Needed	OS/2 1.0 or 1.1	OS/2 1.0 or 1.1	OS/2 1.0 or 1.1
Price	$349	$295	$195
	Inquiry 883.	**Inquiry 884.**	**Inquiry 882.**

for example, are used simultaneously by different processes and by different threads within the same process. Designing reentrant code takes special care. But another big obstacle to the widespread use of DLLs has been the awkward mechanism Microsoft has provided for building them. With Microsoft C 5.1, you have to use special DLL-oriented libraries to create a DLL. All three Modula-2 compilers dispense with this problem. Modula-2 compilers naturally produce DLL-ready code (i.e., code that doesn't assume that the stack and data segments are one and the same). This is crucial to OS/2 programming for two reasons. First, code in a DLL doesn't have its own stack; it uses its caller's stack. Second, although OS/2 threads (lightweight processes) share data space, each maintains its own stack.

Although all three compilers can easily make DLLs, JPI's scheme is the most convenient. JPI provides its own linker, which is well integrated with the automatic make facility. The make procedure in the JPI TopSpeed compilers is controlled by a dynamic-library description (.DLD) file that contains linker directives. To build a DLL, you simply put the -main directive in a program's .DLD file (because a DLL has no main entry point). The JPI integrated environment looks for a .DLD file when making an application. You can have JPI automatically put your newly built DLL in its proper place in your OS/2 LIBPATH.

Just as Modula-2 modules can have initialization routines, so can DLLs. But using this feature of DLLs is another aspect of OS/2 programming that the Microsoft tools handle awkwardly: You have to drop into assembly language to tell OS/2 to call the initialization routine. JPI solves this problem neatly. The INITDLL module, which is written in assembly language and which contains the necessary instructions, gets assembled by JPI's built-in assembler and can

be included with each DLL.

Both Logitech and JPI provide the Modula-2 run-time library in DLL form. The Logitech compiler uses the DLL version by default, and so produces extremely small executables. The JPI compiler doesn't use the DLL version until you ask it to do so—by putting an option into the DEFAULT.DLD file. The first time you compile with that option, the compiler builds the run-time DLL and puts it on your LIBPATH. Stony Brook's compiler doesn't provide a run-time DLL, but the company says that its forthcoming QuickMod product will.

Portable Multitasking
Unlike the original Modula, Modula-2 doesn't include multitasking or interprocess communication (IPC) primitives as part of the language definition. Niklaus Wirth argued that such facilities should be made available by way of an appropriate library module. In *Programming in Modula-2*, Wirth illustrates such a module, called Processes. JPI and Stony Brook each provide a version of this module. Logitech doesn't, although you can, of course, create one (and Logitech says the forthcoming version 2.0 will provide one).

An ideal Modula-2 program for OS/2 wouldn't look like an OS/2 program at all. Instead of calling native OS/2 routines like DosCreateThread directly, a program should handle threads and IPC by way of generic facilities defined in Processes, thereby remaining portable. That's just how the multitasking demos provided by JPI and Stony Brook work. Each program creates a set of windows under the control of multiple threads; each window executes a unique task and moves around on the screen under independent control.

Because these programs are written in terms of Processes, and because the two vendors' DOS products also implement Processes, you can compile them under

DOS, and they run identically. Of course, since Wirth's Processes module isn't yet a standard, the JPI and Stony Brook interpretations of it differ. While multitasking constructs are not portable between Modula-2 implementations, within an implementation of Modula-2, you can write a multitasking program that will port between OS/2 and DOS.

Although none of the compilers assumes an equivalence between the stack and data segments, there's the additional question of reentrancy in the standard libraries. As C programmers know, the use of static data in routines, such as printf, can cause problems in a multithreaded environment. Wirth's recommended solution is to make a module a monitor and assign it a priority, ensuring mutual exclusion. Only Stony Brook implements this scheme, by means of the OS/2 routine DosEnterCritSec. In principle, it's not a good idea to use DosEnterCritSec this way; it suspends all other processes, even ones that aren't trying to enter the monitor. Semaphores would be the preferred solution. Nevertheless, Stony Brook's scheme seems to work well, since a given thread of execution stays within the library only briefly.

Big Country
OS/2 offers not only big memory (a 16-megabyte address space) but, equally important, virtual memory. In a properly configured OS/2 system, the amount of free disk space is a better indicator of available memory than the amount of RAM. Modula-2 programs built using any of these Modula-2 compilers enjoy these benefits. None of the compilers emulates Microsoft C's inefficient but convenient huge construct, however, so you can't statically allocate anything larger than 64K bytes.

All three compilers supply a Storage module that maps Modula-2's ALLOCATE, NEW, DEALLOCATE, and DISPOSE to OS/2 memory management

routines. JPI's Storage module uses an OS/2 feature called suballocation—the allocation of chunks of memory from a previously allocated pool. Suballocation is slower than normal allocation. In a test program that attempted to create an infinitely long linked list, JPI's compiler slowed to a snail's pace far sooner than did Logitech's or Stony Brook's. On the other hand, only JPI's Storage module can detect attempts to free the same piece of memory twice, and this is a very valuable feature.

In an OS/2 system properly configured with MEMMAN = SWAP,MOVE and with sufficient disk space, a call to ALLOCATE or NEW should never fail. That's fortunate, because both JPI's and Logitech's manuals say that if an allocation does fail, the calling program will unceremoniously terminate. Stony Brook takes a more reasonable approach: its system raises a run-time error that you can field with an error handler.

The three compilers implement memory models in different ways. Stony Brook provides the same assortment that most C compilers offer. JPI defaults to a large model, but you can use the $N directive to get the compiler to use near pointers for intersegment function calls. That's awkward, though; you have to issue the directive from a comment within the code. I like Logitech's scheme the best: all code pointers are near unless the procedure is imported from another module, assigned to a procedure variable, or used as the body of a process. There's no option to change that behavior, but it's unlikely that you would want to.

Errors, Exceptions, and Debugging

Modula-2's run-time error checking nicely complements OS/2 protected mode. Protected mode can catch bugs that make software development a nightmare—like writing to the wrong segment. But it can't catch all errors. For example, you could overrun the bounds of an array without stepping into a forbidden segment.

All compilers can add checks for run-time errors, such as out-of-range array indexes, stack overflow, or dereferencing the NIL pointer. Stony Brook produces the best run-time error messages; the compiler reports the module name and the line at which the error occurred. Logitech reports only the nature of the problem (e.g., "Range Error"), not its location. JPI's integrated environment cites the module and line number of a run-time error. But if you run an erroneous program as a stand-alone, it reports

only that there was a run-time error; there's no indication of which one or where. Stony Brook provides the support that you need to recover from run-time errors. Its Error module doesn't have full-blown exception handling à la Ada, but it does maintain a stack of user-installable error-handling procedures. Logitech's RTSExcep module does provide some error-handling but doesn't permit a program to field an error and keep going. JPI does a particularly good

job of checking for OS/2 errors within the standard library.

JPI's wonderful Visual Interactive Debugger is not yet available for OS/2, so I had to fall back on CodeView when debugging programs built with the JPI compiler. Stony Brook's source-level debugger is fully integrated into the compiler/editor environment, works with the mouse, and features watchpoints, breakpoints, and a backtrace facility. I tried a

Table 1: *Times for a VM2 prime-number generator, running on a virtual machine implemented in three versions of Modula-2. In all cases, background operation is far more efficient. All times are in seconds.*

	Logitech	Stony Brook		JPI	
Foreground	883[1]	939[2]	866[3]	939[2]	870[4]
Background	456[1]	503[2]	436[3]	507[2]	440[4]

[1]Default memory model, which automatically uses near pointers whenever possible.
[2]Large model.
[3]Medium model.
[4]Near pointers ($N) enabled.

beta version of Logitech's MultiScope debugger but, unfortunately, couldn't get it to run on my 2-megabyte system. Logitech is, however, the only one of the three companies with an approach to debugging multithreaded applications, and I expect the final version of MultiScope will be impressive.

Environmental Issues
All three compilers support both a command-line interface and a full-screen windowed environment. Logitech, of course, supports a mouse, as does Stony Brook; JPI, which has the environment I otherwise find most intuitive, does not. Stony Brook and JPI use proprietary schemes for structuring a development project and naming the paths to the various components of a project; I found these schemes a bit disconcerting. I prefer to use the OS/2 PATH for this purpose, as the Logitech compiler does.

Version control is an integral part of Modula-2. Compilers must not only check that all imported procedures are used in accordance with the definition module; they must also ensure that a program uses the right version of an implementation module. You can't simply compile .DEF files into .SYM files, then .MOD files into .OBJ files, and then link. The order in which you compile the .DEF files matters. A Modula-2 compiler must compile .DEF files in the right order, or at least tell you what that order is.

Stony Brook's compiler automatically compiles .DEF files in the right order. Logitech's M2MAKE utility determines file dependencies and then creates a .CMD (OS/2 batch) file that drives the compilation. JPI dispenses with .SYM files entirely; the JPI compiler always recompiles .DEF files. Fortunately, it is fast enough so that you don't really notice—except when you're crunching through those huge PM definition modules.

The source code for the standard libraries contains some of the best documentation for these products. All three companies give you that source code. The printed documentation varies in quality. Logitech's documentation mostly contains printouts of .DEF files. There's little in the way of useful description or examples, and the OS/2-specific information is poorly organized. Stony Brook does slightly better. There's an example for each library module and a separate—though barely adequate—section on OS/2. JPI does the best job on documentation. There's an example of each library routine, an excellent introduction to the OS/2 specifics of the product, and a first-class Modula-2 tutorial.

VM2 Revisited
To exercise the compilers, I used Jonathan Amsterdam's VM2—a compiler, assembler, and virtual machine, all written in Modula-2. See the articles "An Assembler for VM2" (November 1985 BYTE) and "Building a Computer in Software" (October 1985 BYTE). The port to JPI's compiler was the most time-consuming of the three, despite the speed of JPI's excellent environment. That's because it's a one-pass compiler and doesn't allow forward references. You either have to move procedures around so they're declared before they're used or, alternatively, use JPI's Forward keyword.

The compilers differ in how they track the various editions of Niklaus Wirth's *Programming in Modula-2*. Both Stony Brook and Logitech accept an Export statement in a .DEF file, while JPI follows the practice set forth in the third edition and regards the .DEF file itself as an export list.

Both JPI and Stony Brook allow this assignment:

```
VAR s : ARRAY [1..5] OF CHAR;
(* ... *)
s := "Hello";
```

whereas Logitech, following the fourth edition, correctly regards this as a type incompatibility, because the assignment of a string with length five to an array of five characters does not leave room for the terminating null character. Since so many of OS/2's API routines expect ASCIIZ strings, it's important that they actually be zero-delimited.

I compiled the VM2 programs using each of the compilers and then wrote a prime number generator in the VM2 high-level language, compiled it to VM2 assembly language (with the VM2 compiler), and assembled it to create VM2 object code (using the VM2 assembler).

I ran that code on Logitech, JPI, and Stony Brook versions of Jonathan Amsterdam's virtual machine (the results are shown in table 1). Even though Stony Brook (using the medium-memory model) had the best times, the differences among the compilers are small. With each of the compilers, the VM2 primes program ran much faster in the background.

Soul Mates
OS/2 and Modula-2 are, in many ways, well matched. Each of these compilers is a genuine OS/2 product, not just a port of a DOS product. Logitech earns points for PM support, a DLL run-time library, an intelligent default memory model, and a multithreaded debugger.

Stony Brook offers a useful debugger, a stand-alone graphics module, a Processes module, a run-time error-recovery mechanism, and a solution to the problem of reentrancy in the standard libraries.

JPI, like Stony Brook, has a stand-alone graphics module and a Processes module, and like Logitech provides a DLL run-time library and PM support. The JPI product also has what I found to be the most intuitive interface and—a crucial point for OS/2 development—exceptionally good support for building DLLs.

Logitech's Modula-2 compiler for DOS is an industry standard—that is, more third-party Modula-2 libraries exist for Logitech's implementation than for JPI's or Stony Brook's. On balance, nevertheless, I believe that the Stony Brook and JPI products best exploit the multitasking and dynamic linking capabilities of OS/2. But stay tuned—each of these products is evolving rapidly. ■

Andrew Schulman is a software engineer working in Cambridge, Massachusetts. He can be reached on BIX c/o "editors."

Circle 81 on Reader Service Card

World Radio History

A New World for DOS

ViewLink and Magellan explore uncharted waters in man/machine interfacing

Stan Miastkowski

Photo 1: *ViewLink lets you create "views" of files that all contain a specific search word or phrase.*

Over the past few months, several new products have tried to go beyond the capabilities of standard DOS shells such as the Norton Commander and Executive Systems' XTree. Although they're difficult to place into rigid software categories, you might think of them as intelligent DOS shells. Developed using object-oriented programming technology, they give you a radically different means of interacting with your system.

In a way, Traveling Software's View-Link 1.05 and Lotus Development's Magellan 1.0 are new categories of software. Both packages incorporate some of the features of DOS shells, indexers, outliners, and even HyperCard. But taken as a whole, their multiple levels of functionality add up to more than the sum of those parts.

At first glance, these packages appear to be similar—and, to a point, they are. Both let you organize your data by function and context, no matter where it's located on your disk. But although their on-screen displays look similar, they take very different approaches to man/machine interaction.

The ViewLink Connection
A common thread in Traveling Software products has been the concept of linking; LapLink links computers, and ViewLink links together your data and applications using a concept called *views*—logical categories of related data. In fact, Traveling Software calls ViewLink an "associative access manager," because it lets you group related (associated) data into views based on your work preferences instead of the constraints of DOS subdirectories.

To get an idea of how views work, say you're a manager who's responsible for a specific product. You're likely to have many different files on your PC that are directly related to your responsibilities. There might be spreadsheet files with budget projections, scheduling files for project management software, and numerous letters, memos, and E-mail messages. ViewLink lets you link all these files together into one view where they're easily accessible.

You can also link individual items to any number of different views. For ex-ample, you might have a view that contains only items that relate to the product's financial planning, or a view that includes only items that relate to a specific member of the project team. In addition, a view can be linked to other views. Beyond that, you can have different sets of views, called *domains*. This is particularly effective in LAN installations, where each individual user can have his or her own domains, as well as share common domains with the workgroup.

When you start ViewLink, you see a split screen with views on the left and files associated with the views on the right. Initially, the views are primarily subdirectory names. Because the data files that you incorporate into a view are automatically linked to their associated

continued

Company
Traveling Software, Inc.
18702 North Creek Blvd.
Bothell, WA 98011
(206) 483-8088

Format
One 5¼-inch floppy disk or
one 3½-inch floppy disk

Hardware Needed
IBM PC or compatible with 384K bytes
of RAM and a hard disk drive

Software Needed
DOS 2.0 or higher

Documentation
User's manual; overview guide;
applications guide

Price
$149.95

Inquiry 885.

Company
Lotus Development Corp.
55 Cambridge Pkwy.
Cambridge, MA 02142
(617) 577-8500

Format
Three 5¼-inch floppy disks or
two 3½-inch floppy disks

Hardware Needed
IBM PC or compatible with 512K bytes
of RAM, one floppy disk drive, and a hard
disk drive with approximately 720K
bytes of free disk space for the program
plus 5 percent to 10 percent of free disk
space for the index

Software Needed
DOS 2.1 or higher

Documentation
User's manual; quick start-up guide;
suggested user's guide

Price
$195

Inquiry 886.

which it runs. For example, you can point to a spreadsheet file and press Return, and the file link automatically launches the application, bringing up the spreadsheet on the screen with the desired file already loaded. Likewise, pointing to a text file launches a word processing application. There's also a cut-and-paste feature that lets you move data between applications.

To run under ViewLink control, specific applications must be installed and closely tied to ViewLink via macros. For each application installed, there are up to four standard application macros (execute, run, print, and create) and two key macros. ViewLink automatically invokes the application macros, and you use the two key macros to save your work and quit the application.

If the applications you use most often are not in the program's install list, you'll need to write a specific application macro for it. Traveling Software provides detailed information for macro creation, but you'll need a modicum of programming skill.

Keeping the Faith
Once you've installed and set up ViewLink to your individual preferences, it requires a continuing commitment. Another powerful feature of ViewLink is its ability to automatically incorporate new items into a view without your having to specifically add them each time you create a file. With "automatic view update" on, each time you generate a data file that contains any of the search criteria you used in generating the original view, ViewLink automatically updates the view to include the new item. And the process is very fast, usually taking not more than 3 to 5 seconds.

Traveling Software says an OS/2 version of ViewLink that runs under Presentation Manager will be available by the end of the year. ViewLink 1.05, which is now shipping, lacks mouse support and the ability to use expanded memory. A Traveling Software spokesperson says these features will be included in version 1.1, which should become available at about the time you read this.

Exploring with Magellan
Instead of ViewLink's approach of associating files into categories (views) that you customize to your personal preferences, Magellan takes an inherently different approach to dealing with data. It treats your hard disk (or even multiple disks) as a whole. During Magellan's initial installation, it creates an index of all

applications, ViewLink takes care of the actual launching of applications.

Multilevel Installation
Getting the most out of ViewLink requires a sizable time investment. Besides an initially steep learning curve, the very nature of the program means the installation is time-consuming. There are really two levels to setting up ViewLink: the automatic initial installation and the fine-tuning process that customizes it to your particular preferences.

The first-level installation is actually quite simple. ViewLink's functionality is tightly coupled to specific applications. The installation utility lists some 60 of the most popular application programs, including all major categories. You tell ViewLink which applications you'll be using, and it goes through a multiple-step process. First it finds the specified applications and their related files, and then it links them to specific macros that ViewLink requires.

After it has found the applications that you'll be using, ViewLink then searches your entire hard disk for files that obviously work with them. For example, it links .WK1 files to Lotus 1-2-3, .DOC files to Microsoft Word, and .CMD files to Procomm. If you've used nonstandard

filenaming conventions, it may link files to the wrong applications, but you can easily unlink those later.

The end result of the initial installation is a master link file that keeps track of views and links. ViewLink's link file is extremely small: My initial link file for 48.6 megabytes of applications and data took up just 130K bytes of disk space, and it grew very little as I customized my own views.

The second part of the installation is considerably more time-consuming and involves the actual creation of individually tailored views. ViewLink gives you several options for building views, including filenames, dates, and types. You can even enter complex Boolean formulas to tell ViewLink what to include and exclude in a view. And when all else fails, you can physically move through the filenames on your disks, tagging the ones you want as you go along.

But ViewLink's most powerful feature is the ability to build views by content. For example, you can enter text strings, and the program will search for them. Every time it finds a match, it includes the file in the view (see photo 1).

Once you've generated your own personal views, each of the individual items is linked to the specific application under

NOW YOU CAN DESIGN JUST ABOUT ANYTHING...INCLUDING YOUR OWN CADD SYSTEM!

Anyone can produce a computer-aided drafting and design package that sells for thousands of dollars. But creating one that's fast, powerful, and affordable takes real ingenuity.

That's exactly what we've done at Generic Software.™ We sell the most widely used CADD program in the world for under $500.

Now available for Macintosh™

PC Magazine ranked it "Editors Choice" in a face off with 17 low-cost CADD packages. "This product is an outstanding value from every point of view and is highly recommended." "A paperback version of AutoCAD", stated PC Week.

And our CADD programs are just part of what we offer.

You can start with Generic CADD Level 1,™ then advance to other levels as your needs— and skills—expand.

And you can use our CADD add-ons, Utilities, and symbols libraries to design the CADD system that fits your needs.

You only pay for the functionality you need, and the functionality you get has depth.

Critic tested, market approved. Generic Software offers price, quality, and support. Match our customer support against anyone!
• Unconditional 60-day guarantee on most products
• Unlimited free technical support • Free monthly newspaper
• Regular updates at modest prices. All adding up to prove that the only thing generic about us is the price.

Call us for a free CADDalog™ or for the name of your local dealer: 1-800-228-3601.

Generic
S O F T W A R E

World Radio History

```
Lotus Magellan        Edit Entry and Press Enter          UPDATE    EDIT
Explore: All Files Concerning "carter & peanut"
  68% FILEA.TXT                          JIMMY (JAMES EARL) CARTER
  65% DEMO.2
  40% METAPHOR.400  Jimmy (James Earl) Carter, 39th president, Democrat, was the
  32% I                                                   e was born
  30% P                        Explore                    n Gordy C
  30% 1    All Files            Concerning:               nd gradua
  30% D    Current Explore Path  Text: [carter & peanut_] ine progr
  30% S    Lotus 1-2-3 files     Line                     t Union C
  30% P    Symphony files        Marked text
  30% T    Manuscript files      File                     take ove
  30% P    Lotus Express Files                            on-ginnin
  30% E    Agenda files          Setup...                 nor, 1966
  30% T    Practice files        New...                   ated Pres
  27% C
  26% O    Enter text to use as explore criterion
  26% R
  26% SPE                                                 gotiation
  26% KK                                                  ttacked t
  26% TEST.C  in Teheran and held members of the embassy staff hostage. Du
  26% MT.C    was widely criticized for the poor state of the economy and
            He was alo viewed as weak in his handling of foreign policy.
File 1 of 38      C:\SQLBASE\FILEA.TXT       Text        Line 1 [Zz]
─F1──────F2──────F3──────F4──────F5──────F6──────F7──────F8──────F9──────F10─
Help
```

Photo 2: *Magellan's search ability is fast and uses fuzzy logic to evaluate the success of a search.*

the data on your hard disk. Like View-Link, Magellan uses a proprietary technology. It creates a surprisingly compact index. Lotus says that it normally takes up 5 percent to 10 percent of the data space, and my 48.6 megabytes of applications and data resulted in an index of about 2.5 megabytes.

Magellan also does lots of the initial work for you. Although the program took about an hour to index my disk, I could start using it almost immediately once the indexing process was complete.

Magellan's forte is *viewing* (not to be confused with ViewLink's views). On the left side of the screen is an alphabetical list of all the files on your hard disk. It can be daunting—in my case, there were 2087 files. But there are many ways of narrowing the list to a more manageable length, including a clever "incremental find" feature that instantly finds filenames as you press the letter keys.

As you scroll down the list, you can see the contents of each file on the right side of the screen. Magellan has over 16 customized view utilities that present data in the format that you'd see in the associated application file. A .WK1

1-2-3 spreadsheet file looks like the spreadsheet, a dBASE file is formatted correctly, and so on. It's all done automatically, because, in the index process, Magellan (like ViewLink) associates the data files with applications. But what's even more amazing is that Magellan lets you peek into binary files, and it instantly shows you certain packed files (.ARC) in their unpacked state.

Like ViewLink, Magellan lets you point to a file and start it up in its application. You do it with a Launch function key. But unlike with ViewLink, the process isn't completely automatic. When you press the Launch key, Magellan asks you which application you want to start and presents a list of choices. It points to the most obvious application (e.g., Lotus 1-2-3 for a .WK1 file). Although this extra keystroke may sound inconvenient, it makes a lot of sense. I use two different editors—XyWrite and Norton Editor—for different applications, and the ability to quickly choose either one is handy indeed. ViewLink, on the other hand, always assumed I wanted to use XyWrite.

Although it's not exactly a new concept, one of Magellan's handiest features is that it displays the main function-key commands across the bottom of the screen. This is one of the reasons that Magellan is more immediately useful than ViewLink. When you press and hold the Alt key, the menu changes to 10 new function-key commands. Many of the commands are your standard DOS shell options, such as Copy, Delete, and Sort. But there are also some intriguing new ones, such as Gather and Zoom. The Gather function lets you mark text from any application shown in a view window and exports it into an ASCII file. Zoom expands the filename or the file view.

The Warm Fuzzies

If Magellan just gave you a huge list of files and the ability to quickly peek into them, it would be useful enough. But where Magellan's real power starts to show is in its ability to do fuzzy searches of all the files on your hard disk. Although Magellan can quickly find specific words or phrases anywhere on your hard disk, that's a feature shared by several indexing programs. Magellan's Explore function extends this ability by letting you use common English words or phrases. This feature uses AI techniques that Lotus first included in its HAL natural-language interface to 1-2-3.

For example, you can tell Magellan to explore all files concerning "Telephone Installation Corporation." Magellan searches for close matches to the words "Telephone," "Installation," and "Corporation," and flags a match if it finds the words within a short distance of each other. Magellan then shows you a list of the files where it found a fuzzy match, followed by a percentage. This *explore rank* shows you the number of exact matches (ranging from 75 percent to 100 percent) and the number of fuzzy matches (ranging from 0 percent to 74 percent) (see photo 2). You can then browse through the matched files, with the words or phrases that you searched for highlighted.

Although dealing with fuzzy searching is initially a bit confusing, it doesn't take long to see what a powerful concept it is. It's most helpful when you're looking for a concept and don't remember the exact wording that you used in the original file. Most of Magellan's searches

PARK A CAR
ON YOUR DESKTOP.

Now you can have the best of both design worlds: the powerful Macintosh interface plus a workstation-quality graphics display, thanks to Truevision's HR graphics card. The HR card bridges the gap in today's computing marketplace by offering a non-inter-laced, sharp 1280 x 960 pixel display at 8 bits per pixel. Its 256 colors (from a 16.7 million color palette) give you tremendous flexibility for engineering drawings, scientific visualizations, or 3D modeling.

Truevision's HR card is QuickDraw™ compatible,* which means that you can run with virtually all Macintosh® software. And you may select the monitor of your choice from manufac-turers like JVC, Philips and SONY.® There are two Truevision HR Cards: The HR 2M ($3995) addresses a desk-top up to 2048 x 1024 pixels and the HR 4M ($5995) addresses a desktop up to 2048 x 2048 pixels. Take a test drive soon by visiting your Author-ized Truevision reseller, or call us at **800/858-TRUE** to obtain literature.

The HR Graphics Card works with the Macintosh II, IIx and IIcx.

TRUEVISION®
Truevision, Inc.
7351 Shadeland Station, Indianapolis, IN 46256

Circle 258 on Reader Service Card

INTERNATIONAL: Canada 416 499-9400   France 33-13-952-6253   Italy 39-2-242-4551   Switzerland 41-1-825-0949   U.K. 44-1-991-0121   West Germany 49-89-612-0010
Other International 617/229-6900.   Suggested retail price is US domestic price, and is subject to change.

*All graphics cards with more than 1MByte of memory require 32-bit QuickDraw.   QuickDraw is a trademark and Apple, Macintosh and Mac are registered trademarks of Apple Computer, Inc.   SONY is a registered trademark of Sony Corporation of America.   Image courtesy of Electric Image.   © 1989 Truevision Inc.

BYTE

Macintosh
Special Supplement

Mac System 7.0
HyperTalk
Spectrum/24 • Showcase F/X
MaxPage 1.2
List Manager Techniques

How to make your VAX talk Mac.

And your Mac talk back.

And how to do it *fast*.

Introducing MacRAF.

The fastest Mac-to-VAX link in the industry. (Over 100,000 cps transfer rate.)

No staring out the window or tapping your fingers while you wait for an application to load. Or a file to save.

VAX access from your Mac is as fast as launching an application on your local hard disk. And just as easy too. So easy, in fact, you won't notice MacRAF working for you. It's absolutely transparent.

What's more, connecting to the VAX is so easy…no VAX commands, no VMS prompts…just connect your Mac to Ethernet and you're ready to go.

Finally, MacRAF brings you all the auto log-on, terminal emulation and multiple VT220 VAX session features of RAF, Datability's IBM PC-to-VAX software.

We've already DECed over 30,000 IBM PC-to-VAX users with our RAF (Remote Access Facility) software. And now your Mac can join the club.

☑ **Try MacRAF for a 30 day, no-risk, free trial.***

OR

☑ **Receive free: "How to DEC a Mac. Or PC." When you call or FAX:**

1-800-DIAL-DSS
NY (212) 807-7800
FAX (212) 463-0459
Canada (613) 937-4444

Our brochure explains how MacRAF provides network integrity, seamless filesharing and storage, as well as terminal emulation.

*Offer valid in the U.S. only and expires October 31, 1989.

◥DATABILITY

World Radio History

Circle 465 on Reader Service Card

BYTE

Macintosh Special Supplement

EDITORIAL

System 7.0 and the Macintosh IIcx

Apple's plans for the future of computing

During 1989, Apple gave us a glimpse of its plans for the future of personal computing. With the introduction of the Macintosh IIcx in March and the announcement of the next-generation operating system, System 7.0, at the May Developer's Conference, Apple has staked its claim to the future of innovative personal computing. While neither the Mac IIcx nor System 7.0 is a product breakthrough, both do indicate the direction that Apple will take during the 1990s. Why are the Mac IIcx and System 7.0 so important?

The Mac IIcx is not important because of what it's made of. Plenty of vendors sell machines with processors at least as fast as the IIcx's 16-MHz 68030. Plenty of vendors sell machines with an industry-standard bus architecture for expandability. And plenty of vendors sell machines with high-resolution graphics capabilities. No, the hardware is not the exciting part of the Mac IIcx. The real

excitement is how the Mac IIcx is made. It is the first Mac design to really take modular construction—or *design for manufacturing* (DFM)—to heart.

DFM is the wave of the present in personal computer manufacturing. DFM dictates that a computer's hardware be designed with ease of assembly and disassembly in mind. This results in a machine that's cheaper to make and cheaper to fix when it breaks.

In the area of software, the excitement is System 7.0. Although it won't be available until 1990, the May announcements promise that System 7.0 will include most of the modern operating-system features that we'll all need to handle information in the new decade. Things like outline fonts, interapplication communication, virtual memory, an improved Finder interface, and printing enhancements are all important, but the crucial part of System 7.0 is what it lacks.

What's missing is backward *in*compatibility. You can run System 7.0 on any Mac, from the Mac Plus to the Mac IIcx, as long as you have 2 megabytes of RAM. Ever tried to run OS/2 on an old PC or XT? It won't work, no matter how much memory you have. There's no backward compatibility for OS/2 on IBM's older PCs because the 8088 processor lacks

the horsepower, and so OS/2 was written for a later-generation Intel processor, the 80286. A Mac Plus or Mac SE, however, even with their dated and overworked 68000 processors, will run System 7.0. They'll take advantage of all System 7.0's new features, with the exception of virtual memory. This is no easy trick, and it points to Apple's commitment to its installed base of Macs.

Apple has the unique opportunity to really broadcast its vision of computing during the 1990s by expanding both of these hardware and software concepts. It can do this by taking DFM and building an inexpensive Macintosh (under $750 list) that runs System 7.0. This Mac, which I call the Macintosh Classic (as opposed to the "classic Macintosh," which started with the Mac 128K and exists now as the Mac Plus), would offer Apple's vision to many people. It would accomplish this because many people—not just large corporations—could afford such a machine, and it would replace the aging Apple IIs that fill our schools, small businesses, and homes. Let's hope that Apple doesn't waste this important opportunity.

—Don Crabb
Contributing Editor
(BIX name "decrabb")

It syncs to

The fact that you're considering a Macintosh II shows you're serious about computers. That you want to reach as far as a great system will let you. So you can reach your potential.

Well, now there's a way to go even further.

By giving your Mac II a MacSync™ monitor from NEC.

MacSync is designed solely for the Macintosh II Video Card. And because it is the direct result of some meaningful conversations with Mac II users, it has all the features you've always wanted in a monitor.

For instance, you asked for easier, more convenient controls. So we moved them up front. You also wanted a 14″ non-glare screen to reduce

C&C
Computers and Communications

the Macs.

fatigue. A smaller footprint and a tilt-swivel base, for more comfortable viewing. They're yours. What's more, we give you a cabinet that looks as good from the back as it does from the front. And a HyperCard stack

Designed and optimized just for the Macintosh II.

on-screen user reference guide called MacSync Valet.™ All at no extra cost.

And if you don't mind us saying so, that's the least you and your Macintosh II deserve.

MacSync from NEC. We've taken everything we've learned about

making great color monitors. And applied it to the Macs.

For literature, call 1-800-826-2255. For technical details, call NEC

Home Electronics (USA) Inc. 1-800-FONE-NEC.

NEC

World Radio History

SHORT TAKES

Spectrum/24

Showcase F/X

MaxPage 1.2

THE FACTS

Spectrum/24
$3999; with NuBus
board trade-in: $2499

Requirements: Mac II
with 2 megabytes of
RAM, a color monitor,
and a hard disk drive,
and running System
6.0.3/Finder 6.1
with 32-Bit QuickDraw
installed.

SuperMac Technology
485 Potrero Ave.
Sunnyvale, CA 94086
(415) 245-2202
Inquiry 471.

True Colors, Revisited

A year ago, I evaluated Su-
perMac Technology's
Spectrum/24, a NuBus video
board that could display 24-bit
color pixels. At the time,
Color QuickDraw didn't pro-
vide any large-pixel support:
It only worked with color
pixels 8 bits in size. SuperMac
cleverly used a chunky/planar
mode that was defined—but
unsupported—by Apple to
work around this limit. Draw-
ing operations were somewhat
slow and made for some inter-
esting screen effects as the pri-
mary colors rippled into the
frame buffer, but it worked.
Nor could you argue with the
photographic quality of the re-
sults. Since the Spectrum/24
used an unsupported graphics
mode to function, SuperMac
Technology sold the board
only to developers.

Apple's 32-Bit QuickDraw
changes all that: Now Macs
that support color have the ca-
pability of displaying, ma-
nipulating, and printing full-
chunky pixel images that are
16 or 32 bits deep. It seemed
appropriate to check back on
the Spectrum/24 video board
to see if it had changed with the
times.

The Spectrum/24 most cer-
tainly has changed. Although
it still sports the same name
and features, the board's elec-
tronics have been completely

redesigned to fully conform to
32-Bit QuickDraw's full-
chunky pixel format.

One feature that the new
Spectrum/24 inherited from
its predecessor is support for
both SuperMac's 16- and 19-
inch monitors (displaying
1024 by 768 pixels) and
Apple's 12- and 13-inch moni-
tors (displaying 640 by 480
pixels). Another inherited
feature is screen depths of 1, 2,
4, 8, and 32 bits (of which 24
bits actually hold color
information).

At "shallower" screen
depths (4 bits or less), the un-
used portions of the Spec-
trum/24's frame buffer are
used to either expand the di-
mensions of the Mac screen (in
what SuperMac calls a "vir-
tual desktop") or to provide a
2×-zoom magnification fea-

ture on part of the screen. A
built-in hardware pan function
scrolls this enlarged screen
automatically as the mouse
pointer reaches the edge of the
display.

I used the Spectrum/24 on a
Mac II equipped with 2 mega-
bytes of RAM and a 40-mega-
byte hard disk drive, and on a
Mac IIcx with 4 megabytes of
RAM and an 80-megabyte
hard disk drive. For both sys-
tems, the video board drove a
SuperMac 19-inch Trinitron
monitor. Installation was as
simple as plugging in the
board and rebooting.

The Spectrum/24 worked
fine with the alpha version of
32-Bit QuickDraw that I was
using, and it switched through
all screen depths without a
hitch. The 24-bit-deep images
that I had captured with a

Howtek Scanmaster color
scanner closely resembled the
original photos. Screen per-
formance at 32-bit screen
depths was slower than at 8
bits, but not prohibitively so,
as it was with the chunky/
planar boards.

The Spectrum/24 helps
provide the hardware portion
of Apple's 32-bit imaging
solution, and it definitely
brings WYSIWYG to high-
end color prepress applica-
tions. It and 32-Bit Quick-
Draw work synergistically to
provide crisp screen updates
without any of the color after-
images that plagued chunky/
planar hardware implementa-
tions, and they do it with very
snappy throughput. I'm look-
ing forward to seeing what
other interesting applications
develop now that the Spec-
trum/24 makes this type of
display technology available.

The Spectrum/24 costs
$3999. For a limited time, you
can upgrade to a Spectrum/24
for $2499 by trading in your
existing NuBus video board (it
can be a SuperMac, Raster-
Ops, or Macintosh II video
board) to SuperMac. □
—*Tom Thompson*

Special F/X on the Mac

Did you ever see the movie
Clash of the Titans? De-
spite a stellar cast, it was a
clunker of a film that was re-
deemed only by Ray Harry-
hausen's special effects. I kept
thinking of that movie while
working with Aegis Develop-
ment's **Showcase F/X**, a pro-
gram for creating and animat-
ing text for use in desktop
presentations and videos. This
multifeatured Macintosh
package won't make you the

Harryhausen of computer-based presentations, but it will give you some easy-to-use tools for spiffing up your slide show, videotape, or product demo.

Showcase F/X (the name, which comes from the cinema's abbreviation for special effects, signifies the program's film heritage) is strictly for working with text. It has animation capabilities, but you can use them only with alphanumerics; this is not a package for drawing cartoons. For an idea of what you can do with this program, think of opening credits you've seen at the movies, in which the titles flash across the screen or come at you from the background or glow like neon.

The program gives you a blank drawing board on which you type the text you want, using either the Mac's fonts or what Aegis calls Poly fonts, unique object-based characters that you can manipulate (e.g., stretch, shrink, flip, mirror, and distort) by pulling on the handles that surround the chunk of text. Showcase F/X has several effects you can apply to the text; for example, you can add shadows, a three-dimensional look, a neon-like glow, smears, or colors (16 or 256, depending on your system).

Animating the text is relatively easy, but it does require studying the manual a bit. (This isn't the sort of software you should just dive right into.) If you've worked at all with film animation, you'll find the program to be pretty intuitive; it essentially follows a metaphor of setting up frames and then linking them.

You can do this frame by frame, or you can let the program do some of the work for you. Let's say you're putting together a 50-frame script; you don't have to specify every frame—you can establish frames 5, 10, 15, and so on, and the program will automatically handle the transitions between those frames.

After you've established your script—the content and

sequence of frames—you can preview it to see how it'll look when animated. When you're ready to shoot your script, so to speak, you just click on a button, and Showcase F/X then records each frame.

Showcase F/X will import images from programs that use the PICT file format, such as MacDraw and PixelPaint, but you can use these pictures only as backgrounds behind the titles. You can also scan in images for use as backgrounds.

Now what can you do with all this fancy titling? Well, you can use it in a stand-alone presentation that runs on your Mac (or is projected onto a big screen), or, if you've got the appropriate genlock device, you can transfer the text to a videotape machine; I wasn't

able to test this capability, but Aegis says Showcase F/X will work with genlocks from RasterOps, Mass Micro Systems, and Computer Friends. (You could also send output to a printer, but this seems a waste of the program's talents.)

As a bonus, Aegis throws in its SlideShow program, which you can use to enhance your animated script. SlideShow lets you alter the playback speed of your animation, change transition colors, and loop a group of animation files.

I worked with Showcase F/X on a Mac II with 4 megabytes of RAM; the company recommends at least 2 megabytes, and I'd say that's definitely the bottom line. With a 256-color board, some of the screens were downright daz-

zling. If you're into visuals, you can find yourself spending a lot of time with this package, checking out its box of tricks. I did run into a few weird spots, however. While trying to record a 50-frame animation, I repeatedly got the message that "An I/O error has occurred." I also got a message I'd never seen before: "Can't understand lock." Lock? What lock?

One warning: This program can be pretty slow. Screen drawing seemed a bit poky, and the recording process gives you ample time to go fetch a cup of coffee; in fact, it takes long enough that you can brew a new pot.

Not everybody needs a program like this. But if you've got a presentation or demo to give and would rather have the audience looking at the screen than at you, Showcase F/X can help you out by providing the tools to create brilliant displays. If you're a filmmaker looking to put effective titles on your videotape, doing it yourself using this program is considerably less expensive than hiring someone else to do it with traditional equipment. Showcase F/X is one more indication that personal computers, particularly the Macintosh, can meld beautifully with the visual arts. □

—D. Barker

The $89 Page

If you've been thinking about doing some desktop publishing with your Mac, **MaxPage 1.2** may be a good program to get you going. The program costs only $89 and has most of the standard desktop publishing features.

Like most desktop publishing programs, MaxPage put me immediately into an untitled page. To start, I drew a text box by holding the mouse button down, dragging down and to the right to size it, and then letting go of the button. To work with multiple columns, I called up the full-page

grid, which, unfortunately, is divided into inches rather than picas, with subdivisions in one-quarter-inch rules. It also has horizontal and vertical half-page and third-page dividing lines. I learned to hide the grid before printing my document.

After positioning columns, I started adding text. When I moved the cursor into a box, it changed into a text-editing I-beam. I simply pushed the mouse button down, and the text-insertion bar began blinking within the box. Once I selected a box, any menu commands affecting a box applied only to that particular box.

I entered text by typing, but I could have imported any ASCII file as well. When you import text, it uses the box's right side as its right margin and automatically wraps around until all the text has been added. If the text length goes below the bottom of the importing box, it is stacked below the visible area.

MaxPage also offers all the usual Macintosh editing features, such as select, cut, copy, and paste, as well as the typical selection of fonts on the Mac II. You can change the font inside a box at any time, the same way you do within any Mac document. One thing to remember is that if you change fonts for a particular box, the text in corresponding boxes will also appear in that font unless you change it.

If you increase the width of a box, the text automatically adjusts to fit inside the new box size. If your text goes beyond the last box on a page, you can wrap it into memory and then wrap it into a box on the next page.

You can also import draw-ings or paintings from source files that are in PICT format or in PNTG, a MacPaint-style format. This lets you use Mac-Draw and MacPaint to create detailed graphics that you can import into your MaxPage documents. Each time that MaxPage redraws a graphic, it reimports it quickly.

One feature that I found useful is MaxPage's ability to automatically adjust the graphic inside the box to fit, no matter how many times you resize the box. The manual recommends that you make your original drawing fill an entire page in your graphics application before you import it into MaxPage. In that way, your drawing will completely fill the box that you import it into, giving you total control over its sizing.

MaxPage also gives you picture-adjustment facilities in the form of scroll bars immediately below and to the right of the picture. These scroll bars let you expand your pictures from the center, equally outward on all sides, to the left or right, and upward or downward. Again, if you change the size or shape of the box, Max-Page will adjust your drawing proportionally. An additional scroll bar farther to the right lets you enlarge the picture or, if you change your mind, reduce it again.

You can also add a background to your document. Backgrounds can be full-page PICT files, but you cannot use PNTG files for this purpose.

MaxPage is an easy-to-use page-layout program for the Macintosh that gives you many of the features included in more-expensive page-layout programs. ■

—*Martha Hicks*

THE FACTS	
MaxPage 1.2 **$89**	Applied Systems & Technologies, Inc. 227M Hallenbeck Rd. Cleveland, NY 13042 (315) 675-8584
Requirements: Mac 512KE or Mac 512K with 128K-byte ROMs.	**Inquiry 473.**

World Radio History

Aside from being a tad easier to install than a satellite, the Hayes Smartmodem 2400M™ for the Macintosh® II and Smartcom II® for the Apple® Macintosh make up the most advanced communications system you can put inside a Mac.®

With this modem and software package, you can set up an extensive Macintosh work group system through a feature called HayesConnect.™

HayesConnect™ allows any Mac access to the Smartmodem 2400M across an AppleTalk® Network.

Which means all Macintosh computers on the Network will be able to communicate with or without a modem of their own. This makes for extremely efficient office communications.

THE ONLY COMMUNICATIONS SYSTEM MORE ADVANCED THAN OURS WON'T FIT INSIDE YOUR MAC.

To make them even more efficient, the system offers Smartcom II for the Macintosh. It's the only software designed to take full advantage of the power and graphics capabilities of all of the computers in the Macintosh family. For example, you can program your own on-screen buttons to create a personalized user interface. You also get features like moveable icons, custom color selection and full support of ImageWriter® and LaserWriter® for incredible graphics. You can even run the system unattended using an Autopilot feature. Of course, there is much more you can do with a few simple clicks on standard, easily identifiable icons.

By now it's probably clear that whether they're just used with the Macintosh II or shared by Mac computers on an AppleTalk Network, the Smartmodem 2400M and Smartcom II make a communications system that can't be beaten. At least not by anything on this planet.

Hayes®

System 7.0:
The Next-Generation Mac Operating System

Tom Thompson

It offers many features competitive to OS/2 yet remains compatible with the existing software base

In early May. Apple announced certain details about its much-rumored System 7.0 operating system for the Macintosh. This served to eliminate much of the rampant speculation about its features and also revealed Apple's course for desktop computing in the 1990s.

For starters, System 7.0 will correct a number of limitations with the existing Mac operating system: It will handle large hard disk drives with thousands of files; accurately display fonts on low-resolution devices, such as impact or SCSI printers; provide support for color printing and third-party printing devices; and expand the address space out of its current 16-megabyte limit. At the same time, System 7.0 will supply many new features: virtual memory; a new Finder with a more consistent way to add fonts, desk accessories (DAs), sound resources, and Control Panel modules (cdevs); an enhanced file system that can handle MS-DOS or NFS volumes; communications support (serial and networked); database support; and ways to establish live data links between running applications. But there's still no preemptive scheduler or hardware memory protection; it's still up to MultiFinder to provide multitasking capabilities. Nevertheless, System 7.0 promises a lot of OS/2's features and will provide them across the entire Macintosh line, from the Mac Plus to the Mac IIcx.

All you need to run System 7.0 on existing machines is a minimum of 2 megabytes of RAM. An IBM PC system using OS/2 and Presentation Manager requires at least 3 megabytes of RAM and an 80286 processor.

I must stress that much of the information Apple supplied is preliminary and subject to change. Also, I had no hands-on experience with even prototype software. With that in mind, I'll focus on a few of the more interesting parts of System 7.0. I'll provide a more comprehensive report when the software becomes available.

The New Finder
The new Finder lets you customize your system or add enhancements using a consistent interface. To add DAs, fonts, and sounds to the system, you simply copy the files into the appropriate folder.

DAs and Control Panel files appear as icons on the Desktop, and you activate them by double-clicking on the icon—the same as launching a Mac application. Attached printers appear as icons, and you can print a document by dragging it onto the printer icon. The new Finder also provides a built-in file search function, a help window, and file aliasing.

A 32-bit Address Space
The current Mac operating system is limited to a 24-bit address space 16 megabytes in size, of which only 8 megabytes is available to applications. This is the case even though the Mac II family and the Mac SE/30 use 68020 and 68030 processors that can handle a 32-bit address space (4 gigabytes). This occurs because not all of the Mac operating system implements 32-bit addressing (two of the offenders here are the Memory Manager and QuickDraw)—a legacy from the 68000 processor's 24-bit address bus. System 7.0 will eliminate the vestiges of the 24-bit addressing limit in the Macin-

tosh operating system.

Interestingly, QuickDraw's addressing problems could be dealt with apart from the rest of the operating system and are fixed with the release of 32-Bit QuickDraw (see "Apple's 32-Bit Quick-Draw Covers the Spectrum," July BYTE). A Mac can use 32-Bit Quick-Draw's enhanced capabilities while running in a 24-bit environment under System 6.0.3.

These modifications in System 7.0 will further the migration of Mac software to a 32-bit environment. They will allow present and future Mac applications to access larger amounts of RAM, in order to deal with the large computing jobs of the 1990s.

Virtual Memory
System 7.0 will implement virtual memory; unused objects in RAM are written to a file on disk and read back into memory when needed. Although there's a performance penalty because of this "swapping" overhead and because disk accesses are slower than RAM, virtual memory lets you work with objects larger than the computer's physical memory.

System 7.0's virtual memory will use a demand-paging scheme using 4K-byte pages (one block of memory). In the 24-bit environment, you can configure virtual memory to a maximum of 14 megabytes. In the 32-bit environment, you'll be able to use the entire address space, 4 gigabytes.

Virtual memory requires the use of a memory management unit that determines when to swap objects to and from RAM. Since an MMU is an integral part of the 68030 processor, the Mac IIx, Mac IIcx, and Mac SE/30 will have virtual memory the moment System 7.0 is installed. For the Mac II, a 68851 paged memory management unit chip must be placed in the MMU socket. The Mac Plus and Mac SE, using 68000 processors, won't be able to take advantage of this feature.

Outline Fonts

The bit-mapped fonts normally used by the Mac have several limitations. You can display a font—and print it on a non-PostScript printer—with good results if you have the font resident on your system. The problem is, handling every possible point size of every typeface you might ever need requires lots of disk space. Not only that, but these low-resolution bit maps reproduce poorly on high-resolution laser printers.

Apple's solution is outline fonts. In outline fonts, a character is stored as points that describe its outline mathematically as a series of quadratic B-splines. As with PostScript fonts, this technique allows the accurate representation of characters on high-resolution output devices, such as laser or Lintronic printers. For low- or medium-resolution devices, such as impact printers or the screen, where the character must be mapped into the constraints of a grid containing a limited number of print wires or pixels, the outline fonts provide another display mechanism.

An Apple instruction set allows a font vendor to associate a program with each character that, when executed by System 7.0's low-level software, will correct the character's appearance to fit within the grid of the output device. This promises to give the Mac the ability to generate attractive text for an output device of any resolution and at any point size while using just a single outline font for a given typeface. Apple plans to publish the outline specifications and instruction set for use by third-party font vendors.

Communications Toolbox

As its name implies, the Communications Toolbox will provide all applications with high-level access to standard communications functions. Currently, an application must access serial or networking drivers directly to use communications services. The Communications Toolbox will accomplish this in much the same way that Color Quick-Draw does: by providing a set of versatile device-independent routines, while low-level software handles the chore of translating these routines to hardware-specific calls for a particular I/O board. A set of "standard" dialog boxes will allow the user to configure communications parameters, such as the transmission rate, parity, and stop bits for the serial port.

The Communications Toolbox has been under development for some time. It will be available for use with System 6.0.3 in the third quarter of this year.

New Print Architecture

System 7.0 will provide a new printing architecture that supports color, grayscale, and custom page sizes (e.g., mailing labels and tickets). It will accomplish this while retaining a one-to-one correspondence with the old printing calls.

As a result, the new printing architecture will be compatible with most existing applications; note, however, that existing printer *drivers* won't work with System 7.0. However, Apple will license a developer's toolkit so that third-party vendors can rapidly modify their drivers to work under the new operating system. This will also allow the Mac to support a larger variety of printers.

No Memory Protection

One of the biggest disappointments in the System 7.0 announcement is that the machine will have no preemptive scheduler or hardware-supported memory protection. This is unfortunate. I've seen MultiFinder handle an application crash elegantly with just an informative message on more than one occasion, but just as often I've had an application crash toss me into the safety net of the TMON debugger. While MultiFinder works, it is only as capable as the most poorly behaved application. Obviously, you should run only reliable applications with Multi-Finder, but I think the onus of system integrity should lie with the operating system, not with the application designer.

To be fair, the reason Apple did not implement hardware protection at this time was to maintain compatibility with existing applications. The Mac operating system currently makes no distinction between system code and application code; everything runs in the 68000's supervisor mode. Furthermore, the system stack is used to share resource information among running applications. If memory protection "walled off" the Mac operating system and the system stack from Mac applications, much of the application software would break. Under these circumstances, it seems to me that the lack of hardware memory protection is reasonable, but I'd like to see it in the future.

Future Course

I've covered only a handful of the features that System 7.0 will provide the Mac user. Again, most of the information is preliminary. I'll report more on System 7.0 and other features as it's released and the details become firm. You can expect to see System 7.0 released early next year.

I'm encouraged by the new openness at Apple. The publication of the outline font specifications and the printer toolkit are a significant step in the right direction in the era of open system architecture. The support for the entire product line is also encouraging, but I'm skeptical that this can be accomplished for the Mac Plus. Nevertheless, if Apple makes System 7.0 live up to its promise and can deliver it on schedule, the Mac will have many of the features found in OS/2 systems, and in some areas, it will surpass them. ∎

Tom Thompson is a BYTE senior technical editor at large. He can be reached on BIX as "tom_thompson."

List Manager Techniques

Jan Eugenides

Handling lists of information? Here's how the List Manager can help.

One of the Macintosh interface's distinguishing characteristics is the way it lets you scroll through, highlight, and select individual entries in lists of information. That information can take the form of data or filenames—either as text or as icons. The tool that gives you this power and mobility is called List Manager, and it's one of the handier items ever conceived for easing the lives of both users and programmers.

Briefly, List Manager provides an easy way of displaying small lists of data in a row-and-column format. It handles most of the mouse interactions (e.g., scrolling, highlighting, hit testing, and selecting list elements). It's best for straight text lists, but it can smoothly handle graphical items such as icons and the kinds of pictures you've come to know through MacPaint's tool palette window.

The information in this article will let you take an informed look at List Manager. If your interest in the Mac is primarily as a user of applications, this detailed examination will help you gain an insight into the complexity underlying the Mac interface. Whether you're a casual programmer who'd like to customize commercial software or a professional who writes applications from scratch, you'll recognize straightforward techniques you can use to take some of the hassle out of Mac programming.

Although my code is written in MPW C 3.0, the techniques I use apply to other languages as well. Please note that although I refer to sample code in this article, it was not possible to include the code in its entirety. It is, however, available on disk and on BIX for downloading (see page 5 for details).

List Manager Basics
The first item of business when working with the List Manager is to create an empty list. A list is always associated with a particular window and is displayed in a rectangle within that window. The list can have vertical and horizontal scroll bars if needed, and it can be made resizable. The call to create a new list, LNew(), is shown in listing 1. To help distinguish them from the other 400-odd Mac Toolbox calls, List Manager calls are prefaced with an "L."

Most of LNew()'s parameters are fairly self-explanatory. The Rect rView is the rectangle in which the list will be displayed in the window's local coordinates. It does not include the area for the scroll bars, if any.

The size of the list in rows and columns is given by Rect dataBounds. The dimensions of a list are always specified in numbers of cells; for example, if you wanted to create a list with 5 columns by 10 rows, you would set dataBounds to (0, 0)(5, 10).

A Cell is really nothing more than a Point structure; that is,

```
struct Point {
short v;          /*vertical*/
short h;          /*horizontal*/
};
```

The size of a cell in the list is determined by the vertical and horizontal values of the cSize parameter.

The parameter theProc is the resource ID number of the list definition (LDEF) to use for the new list. If you pass NULL for this parameter, the default text-only list definition is used. Much of the power of the List Manager lies in writing your own list definitions, which I'll discuss in more detail later.

The WindowPtr w is the window to which the list should be attached. The Boolean drawIt determines whether drawing is turned on or off when the list is created (more on this later); the Boolean grow determines whether the list will be resizable; the Boolean scrollH determines whether the list has a horizontal scroll bar; and the Boolean scrollV determines whether the list has a vertical scroll bar.

continued

The handle returned by LNew() references a data structure called a List Record. It's a fairly complex structure, but since various List Manager routines are provided for accessing cell data, you'll rarely, if ever, have to deal with it directly.

The sample program that accompanies this article on BIX, ListMgrDemo, has two routines that create lists: CreateList() and CreateIconList(). Look in the ListMgrDemo.c for two examples of calling the LNew() function. You must keep in mind several important points when setting up the List Manager.

First, set the size of the list by using a userItem. I almost invariably wind up using the List Manager in a dialog box of some kind. When I lay out the dialog with ResEdit (Apple's resource editor), I find it most convenient to place a userItem wherever a list will go. This allows me to visually select the placement of the various dialog elements. By writing my code to reference the userItem, I also gain the freedom to move or resize the list later without having to change code.

Bear in mind when you use the size of the userItem to determine the size of a list that an area for scroll bars is *not* included in the rectangle that you pass to LNew(). In the CreateList() routine in the sample program, notice that I subtract 15 from the right side of the rectangle before passing it to LNew(), which leaves room for a horizontal scroll bar in the window.

Second, be careful about turning the list's drawing on or off. If you examine the CreateList() and CreateIconList() routines, you'll see that when I call LNew(), I specify that drawing should be turned off (the drawIt parameter is false). Generally speaking, it makes for a cleaner display if you create the list with drawing turned off and then turn drawing on with the LDoDraw() call sometime before the first update event occurs. Otherwise, the list will be drawn twice. It's also a good idea to turn drawing off when adding data to multiple cells so that the list won't be redrawn for each cell.

Third, set the selection flags. The selection flags allow you to customize the way the List Manager handles mouse-clicks and drags. Figuring out just how to set them can be a little bit confusing, however, so I'll show you the two flag settings I've found that provide the most useful behavior. The two lists in the sample application show how to set the flags, but I'll explain what they accomplish.

In CreateList(), the flags are set to lNoExtend + lNoRect + lUseSense + lNoNilHilite, which are predefined List Manager constants. This allows the user to select multiple items by holding down the Shift key and clicking on them. The items do not have to be contiguous, as shown in the two scrollable windows in figure 1. It also prevents empty cells from being selected.

In CreateIconList() the flags are set to lOnlyOne + lNoNilHilite. This setting allows the user to select one and only one item at a time.

Finally, take advantage of the Dialog Manager. When you use a userItem for your list, you can write a small update function to attach to it. Whenever a screen update is required, the Dialog Manager automatically calls your function. This eliminates the need for you to check and handle update events yourself and saves a bit of code.

To accomplish this feat, you must pass the address of a properly designed function to the Dialog Manager's SetDItem() function. SetDItem() is a ROM Toolbox call usually used to set a partic-

---

Listing 1: *The parameters for the* LNew() *function, which creates a new list.*

```
pascal ListHandle LNew(rView, dataBounds, cSize, theProc, w, drawIt, grow,
    scrollH, scrollV)
Rect        rView;        /*The display rectangle in local coordinates*/
Rect        dataBounds;   /*The size of the list in rows and columns*/
Cell        cSize;        /*The size of a cell in pixels (a Cell is a
                            Point)*/
short       theProc;      /*The ID of the list definition (LDEF) to use*/
WindowPtr   w;            /*The window the list should be displayed
                            within*/
Boolean     drawIt;       /*Whether drawing is turned on*/
Boolean     grow;         /*Whether the list is resizable*/
Boolean     scrollH;      /*Whether there is a horizontal scroll bar*/
Boolean     scrollV;      /*Whether there is a vertical scroll bar*/
```


Figure 1: *The sample application ListMgrDemo in action. The top window shows the output of* MyList.c, *which is a combination of scrollable text and graphics. Note that you can select more than one item in the list, as determined by the selection flags. The bottom window shows the output of* IconList, *which is a list of icons. This type of List Manager output makes it easy to implement a tool palette window for a painting or CAD application.*

ular dialog item to a given rectangle or to change the appearance of a control. It's declared as follows:

```
SetDItem(DialogPtr dlg, short
   item, short kind, Handle h,
   Rect r);
```

In the case of a userItem, however, you can pass a pointer to an update function in the h parameter. The update function should be declared like this:

```
pascal void DrawItem(DialogPeek
   dpeek, short itemNo)
```

In the sample program, there are two update procedures, one for each list. They are named DrawList() and DrawIconList(). They are installed into the userItems right after the selection flags are set in both CreateList() and CreateIconList().

The DrawList routine is short enough to include here (see listing 2). It calls LUpdate() to redraw the list and then draws a simple one-pixel frame around the entire list.

Working with Cells

Once you have the list installed in a window or dialog box, you're ready to add data to it. It's unlikely that you'll know beforehand how many rows and columns a list will need unless the data is always a fixed size. Because of this, it is simpler to create the list with only one row or column and then use the List Manager's LAddRow() and LAddColumn() calls to dynamically size the list. LAddRow() and LAddColumn() are declared as follows:

```
pascal short LAddRow(short count,
   short rowNum, ListHandle list);
```

```
pascal short LAddColumn(short
   count, short colNum, ListHandle
   list);
```

Both work in a similar manner. The count parameter is the number of rows or columns you want to add. RowNum (or colNum) indicates where the new rows or columns should be inserted. They are inserted before the given row or column. Rows and columns that are greater than or equal to rowNum (or colNum) are increased by count. If these values are larger than the last row (or column) in a list, new rows (or columns) are added to the end. Passing a value of 32767 for these parameters always adds rows and columns to the end of the list. The short integer that is returned by LAddRow() is the number of the first added row. LAdd-

Column() returns the first added column. All added cells are empty.

In the sample program, the mixed text/graphics list is vertical, and the FillList() function uses LAddRow() to grow the list downward. I've used canned data for the demonstration application, with the data stored in an STR# (string list) resource and in several SICN (small icon) resources. This allows you to see how the list works without having to enter any data. In a real-life program, however, you would fill the list from some user-supplied data.

The icon list in the sample program is horizontal and uses LAddColumn() to grow the list sideways. This happens in the FillIconList() function. Again, I've used canned data for the demo.

There are two calls for removing cells from a list: LDelRow() and LDelColumn(). These are declared as follows:

```
pascal short LDelRow(short count,
   short rowNum, ListHandle list);
```

```
pascal short LDelColumn(short
   count, short colNum, ListHandle
   list);
```

Each of these deletes the number of rows or columns specified by the count parameter, starting with the row or column specified by the rowNum or colNum parameter. If count is 0, all the data in the list is quickly deleted. This gives you a quick way to dump all the data in a list without having to go through and dispose of each Cell one by one.

Now you have a list, and it's the right size for the data you want to display. There are two calls for putting data into cells: LAddToCell() and LSetCell().

They are declared as follows:

```
pascal void LAddToCell(Ptr data-
   Ptr, short dataLen, Cell theCell,
   ListHandle list);
```

```
pascal void LSetCell(Ptr dataPtr,
   short dataLen, Cell theCell,
   ListHandle list);
```

They both work the same way, adding the data that is pointed to by dataPtr, of length dataLen, to the cell specified by theCell. The difference is that LAddToCell() appends the data to whatever is currently in the cell, while LSetCell() replaces current data with new data.

The sample program uses only LSetCell() in the FillList() and FillIconList() functions.

To get data back out of a cell, use LGetCell(). It is declared as follows:

```
pascal void LGetCell(Ptr dataPtr,
   short *dataLen, Cell theCell,
   ListHandle list);
```

LGetCell() copies the data from the given cell into the space pointed to by dataPtr. For this call, dataLen specifies the maximum number of bytes to be copied. If the data in the cell is longer than dataLen, only dataLen bytes will be copied. After the call, dataLen contains the actual number of bytes copied.

The sample program doesn't retrieve any data, so it doesn't use LGetCell().

Handling Mouse-Clicks

Mouse-clicking is an area where the List Manager really shines. When you click on an item in the list (a mouse-down *continued*

Listing 2: *The* DrawList() *function. It's an update procedure that's called by the Dialog Manager when the Mac's screen must be redrawn.*

```
pascal void DrawList(dpeek,itemNo)
DialogPeek      dpeek;
short           itemNo;
{
short           iType;
Handle          iHand;
Rect            iBox;

SetPort((GrafPtr)dpeek);
LUpdate(dpeek->window.port.visRgn,myList);   /*Call list manager to update
                                               the list-it will call our
                                               LDEF*/
GetDItem((DialogPtr)dpeek,itemNo,&iType,&iHand,&iBox);
InsetRect(&iBox,-1,-1);
iBox.right -= 15;
FrameRect(&iBox);               /*Draw a nice outline around the list*/
}
```

event), you have to make only one call to LClick(). It manages control until the user releases the mouse button and handles all selection of cells (according to the rules set by the selection flags), scrolling, and auto-scrolling. If a cell is double-clicked, LClick() returns true. LClick() is declared as follows:

```
pascal Boolean LClick(Point pt,
    short modifiers, ListHandle
    list);
```

The pt parameter is the mouse location in local coordinates, and modifiers is the modifiers word from the event record.

The sample program calls LClick() in response to a mouse-down event in either list. Consult the DoEvent() function in the source code listing for all the details.

After LClick() has returned, one or more cells can be selected. In many situations, you don't have to do anything in particular when a cell is selected. If you do need to perform some housekeeping, such as highlighting a control, you can find out which cells are selected by using LGetSelect(). It is declared as:

```
pascal Boolean LGetSelect(Boolean
    next, Cell *theCell, ListHandle
    list);
```

LGetSelect() acts differently depending on the value of next. If next is false, LGetSelect() returns true if the given Cell is selected. If next is true, LGetSelect returns in theCell the next selected cell in the row that is greater than or equal to theCell.

For simple lists that can have only one selected item, you can get the currently selected item by setting next to true and theCell to 0,0. For lists that allow multiple selections, use a while loop with next set to false.

Overcoming the 32K-byte Limit

One major limitation of the List Manager is that a list can contain only 32K bytes of data. If you use the default text-only list definition, all the text in the list must add up to less than 32K bytes. There is also an overhead of 2 bytes per cell that counts toward this limit.

While 32K bytes can hold a fair amount of text, it is wholly insufficient for many types of graphics. A single PICT, for example, can be more than 32K bytes in size. Then how can you use the List Manager? The secret is in how you write your custom list definition functions.

Look closely at the FillList() and FillIconList() functions in the sample program. In particular, examine the LSetCell() call, which adds data to the list. In both cases, you'll find that the only data added to the list is a handle, which is only 4 bytes long. I've written both of the custom list definition functions for this program to reference their data through handles. That way, it doesn't matter how large the actual data is—only 4 bytes are required in the list itself (plus 2 bytes overhead). With 32K bytes of possible list data, that gives you over 5300 elements, no matter how big they are.

There is one caveat when using this method: You must dispose of your data yourself. You can't just call LDelRow() or LDelColumn() with a count of zero. Only data that is actually in the list (that is, the handles) will be deleted this way. You must go through the list cell by cell and dispose of the data referenced by the handles.

Custom List Definition Functions

This brings me at last to writing the list definition functions that I promised at the beginning of this article. They are surprisingly simple to write and are very useful. Formally, they are declared as follows:

```
pascal void ListDefProc(short
    message, Boolean select, Rect
    *lRect, Cell cell, short dataOff-
    set, short dataLen, ListHandle
    listH);
```

They must be written as a single piece of code, with the entry point located at the beginning of the code. This code is put into an LDEF resource. With MPW, it is easy to create a make file to do this automatically. Check out the file ListMgr-Demo.make in the sample program to see how it's done.

The message parameter that controls what the list definition must do can assume four values: lInitMsg, lDrawMsg, lHiliteMsg, and lCloseMsg. Most lists won't need special initialization and can ignore both lInitMsg and lCloseMsg.

When your list definition function receives an lDrawMsg message, it means that a cell needs to be drawn. The lRect parameter is the rectangle in which the cell should be drawn. The lDataOffset parameter is the offset into the list data of the cell's data; lDataLen is the length of the cell's data.

The lHiliteMsg message means that a cell must be highlighted. In most cases, this simply means the cell is highlighted, and a simple InvertRect() call will do the job.

Two That Do the Job

In the sample program, there are two custom list definitions. The icon list definition is contained in the file Icon-List.c, and the mixed text/graphics list definition is in the file MyList.c. Refer to figure 1 to see how these lists appear on-screen.

The simpler of the two definitions is IconList.c. Because the data consists of nothing but a handle to an icon, and the cells contain nothing but icons, it is a simple matter to draw the icon in the given rectangle.

MyList.c contains a somewhat more complex drawing function. For this list, the handle refers to a structure that contains a string and a handle to a small icon (SICN) resource. To draw the cell, the drawing routine first checks the width of the string to see if it will fit in the cell. If it will, the routine just draws it with DrawString(). If it won't fit, the string

is shortened until it will, and an ellipsis (.) is appended to the end.

There is no Toolbox call to plot a small icon, so the list definition contains its own routine to do this. The PlotSICN() function treats the small icon as an off-screen bit map, which is all that it really is, and then uses CopyBits() to put it into the cell. The result is a small icon followed by some text, much like the display used by Standard File for file selection.

As I hope you've seen from examining my two list definitions, writing one is really no big deal. It does give you a lot of flexibility when you need to display a scrollable list of graphics or text, or both. The built-in List Manager functions make this chore an easy one. ∎

Jan Eugenides is a senior software engineer for Solutions, Inc., of Williston, Vermont. He can be reached on BIX as "j.eugenides."

with their own message handlers.

Here is an example: Radio buttons are often used in HyperCard to enable you to select one of a number of choices, each represented by a button. The last button to be chosen is highlighted, while the others are not. When you make a different selection, the targeted button is highlighted and the highlight of the former choice is turned off. This is easily accomplished by including the lines in listing 3 in the script of each button.

While this approach will do the job, there is a better way than having to include these same lines of code in each button's script. Simply define a handler for a custom message I'll call update-Button (but it could be any single word not already reserved by HyperCard),

***Y**ou must try to prepare scripts to handle every contingency.*

shown in listing 4. When the update-Button handler is placed at a higher level in the hierarchy than the button level, it can be called by simply typing the single word updateButton on a line in the

script of the radio button (see the mouse-Up handler in listing 4).

The only question remaining is where to put the updateButton message handler. You could put it in the script of the card that contains your group of radio buttons. But if you decide later to have another card full of radio button choices, you'll have to duplicate the message handler in the script of that second card, which is an inefficient technique. Including the handler in the background script will cover all the cards of the same background, but you may want to do the same thing in another background. With a custom handler such as this, the best place for it is the stack script, where it will be accessible to calls from anywhere in the stack.

Avoiding Error Messages

Another problem that may befall a HyperTalk programmer is an error message caused by an unanticipated user response. If you want to give your stack the look of a professional program, you must try to prepare your scripts to handle every contingency. This is particularly important when asking users to input data that will be used for arithmetic calculations. I wrote the script shown in listing 5 to handle such situations. This message handler was placed in the stack script. Whenever the user enters data that must be a valid number, the script calls the checkResponse handler as shown in the mouseUp handler in listing 5. The checkResponse handler does two things. It checks to see whether each character in the response is either a decimal point or one of the 10 digits, and it also makes sure there is not more than one decimal point in the response. Thus, any input that passes this test may be used by HyperCard for arithmetic operations.

Passing Parameters

The scripts in listing 5 also present a good illustration of parameter passing. A user's entry is put into the local variable response, which is then used as a parameter to the message checkResponse. This invokes the message handler, on checkResponse, which is passed the variable response. The checkResponse handler then determines if response is a valid number.

Note the use of the global variable valid in both scripts. This is necessary because parameters can be passed in only one direction, to the called handler. Any changes in the value of a parameter will not be passed back to the calling script. In this example, I needed a way to

---

Listing 1: *This HyperTalk script hides all the buttons on a card so that only the text will be printed. Note that it has a length of 15 lines.*

```
on mouseUp
        hide button "Boston II"
        hide button "Times"
        hide button "New York"
        hide button "Home"
        hide button "Next Card"
        hide button "Print Card"
        doMenu "Print Card"
        show button "Boston II"
        show button "Times"
        show button "New York"
        show button "Home"
        show button "Next Card"
        show button "Print Card"
end mouseUp
```

Listing 2: *A script of nine lines that accomplishes the same task as that in listing 1. Plus, it will still work if you add more buttons.*

```
on mouseUp
   repeat with n=1 to the number of buttons
        hide button n
   end repeat
   doMenu "Print Card"
   repeat with n=1 to the number of buttons --hidden buttons
        show button n-- still counted by HC
   end repeat
end mouseUp
```

Listing 3: *A simple button-highlight-control script.*

```
on mouseUp
        repeat with n=1 to the number of buttons
               set the hilite of button n to false
        end repeat
        set the hilite of the target to true
        <other commands here>
end mouseUp
```

Listing 4: *A button-highlight-control handler*, `updateButton`.

```
on updateButton
      repeat with n=1 to the number of buttons
            set the hilite of button n to false
      end repeat
      set the hilite of the target to true
end updateButton

on mouseUp
      updateButton
      <other commands here>
end mouseUp
```

Listing 5: *This script checks that input data is of the anticipated type.*

```
on checkResponse response
  global valid
  put 0 into pointCount
  repeat with n=1 to the length of response
    if char(n) of response = "." then put 1 + pointCount
        into pointCount
    if char(n) of response is not in ".1234567890" or
        pointCount > 1 then
      answer "Please enter a number only."
      put false into valid
      exit checkResponse
    end if
  end repeat
  put true into valid
end checkResponse

on mouseUp
      global valid
      put false into valid
      repeat until valid
      ask "Number of inches to convert"
            if it is empty then exit mouseUp
            put it into response
            checkResponse response
      end repeat
      put response * 2.54 into msg
end mouseUp
```

pass a Boolean result back from my re-sponse-checking script to the script that would use the response. This was accom-plished with the use of a global variable.

User Levels
HyperCard has five userLevels. Many actions possible on userLevel 5 are not allowed at userLevel 1 or 2. If your stack requires a particular userLevel, you should provide scripts that set user-Level to the desired level upon opening the stack and reset the previous user-Level when leaving the stack. Suppose you wanted to set the userLevel to script as well as hide the menu bar and the tool, pattern, and message windows. You would put the handlers in listing 6 into the stack script. The openStack handler

Listing 6: *A script to set* `userLevels`.

```
on openStack
  global oldLevel
  get userLevel
  put it into oldLevel
  set userLevel to 5
  hide menuBar
  hide tool window
  hide pattern window
  hide msg
end openStack

on closeStack
  global oldLevel
  set userLevel to oldLevel
end closeStack
```

Listing 7: *You can put dialog boxes in your stack by using this script.*

```
card button "Choose Fruit"
on mouseUp
  global fruit --variable keeps track of fruit chosen
  put empty into fruit --initialize value
  put "Please select a fruit from the list." into msg
  show card field "FruitList"
  show card field "Mask"
  show button "OK"
  show button "Cancel"
end mouseUp

card button "OK"
on mouseUp
  global fruit
  put "You chose " & fruit into msg
  hide card field "FruitList"
  hide card field "Mask"
  hide button "OK"
  hide button "Cancel"
end mouseUp

card button "Cancel"
on mouseUp
  put empty into msg
  hide card field "FruitList"
  hide card field "Mask"
  hide button "OK"
  hide button "Cancel"
end mouseUp

card field "FruitList"
on mouseUp
  global fruit
  set lockText of me to FALSE
      --Unlocks field: allows selection.
  click at the clickLoc
  select the selectedLine
      --selects text in chosen line
  put the selectedText into fruit
      --stores selection in global variable
  set lockText of me to TRUE
      --Locks field: user can't mess up text.
end mouseUp
```

Figure 1: *The dialog box described in listing 7. The user has just clicked on the word* Pears. *After the OK or Cancel button is clicked, the dialog box vanishes.*

declares a global variable, oldLevel, puts userLevel into oldLevel, and then sets the current userLevel to 5. The closeStack handler resets userLevel to oldLevel.

Dialog Boxes
One of the standard features of the Macintosh interface is the scrolling-field dialog box. The user is presented with a scrolling list of items and asked to select one. Whenever the user clicks on an item, it is highlighted, and if he or she is satisfied with that choice, clicking an OK button records the selection and closes the dialog box. With a bit of clever programming in HyperTalk, you can duplicate this type of interface from within your stack, but the implementation of this feature requires HyperCard 1.2 (for the select command and selected line property).

Figure 1 shows a dialog box created with the HyperCard scripts shown in listing 7. In this example, a scrolling list of fruit pops up in response to clicking on the button Choose Fruit. The message handler for this button gives the command to show the objects that make up

A *good approach is to restrict the font styles in text fields to fonts required by the system: Geneva, Chicago, or Monaco.*

the dialog box: a scrolling field called FruitList, a field called Mask to provide room for buttons, and the OK and Cancel buttons. Depending on the order in which they are created, you may have to use the Bring Forward or Send Back commands to arrange the objects in the proper order.

The key to the dialog box is contained in the mouseUp message handler in the script of the FruitList field. A global variable is used to store the chosen fruit. The field is unlocked, which allows it to recognize the line selected. The text is highlighted with the select command, and the selectedText is put into the global variable fruit for later retrieval. The field is locked again, to prevent the user from altering the text in the list. Clicking the OK or the Cancel button closes the dialog box by hiding its component objects. If OK is clicked, the fruit chosen is identified in the message box.

Smart Scroll Bars
In most Macintosh windows, the scroll bars are not active unless they are required. A HyperCard scrolling field, on the other hand, always shows an active scroll bar, even if there is plenty of room for the text it contains. I decided to remedy this situation while working on a stack with several background fields shared by a group of cards. This stack enables a user to keep a daily log of meals eaten. Food items for each meal are chosen from menus and are then listed in

continued

World Radio History

Listing 8: *A script that creates "smart" scroll boxes.*

```
on openCard
  if the number of lines in field 1 > 6 then
    set the style of field 1 to scrolling
  else
    set the style of field 1 to rectangle
  end if
  if the number of lines in field 2 > 4 then --This part adjusts field
    set the style of field 2 to scrolling --style.
  else
    set the style of field 2 to rectangle
  end if
  if the number of lines in field 3 > 7 then
    set the style of field 3 to scrolling
  else
    set the style of field 3 to rectangle
  end if
  if the number of lines in field 4 > 3 then
    set the style of field 4 to scrolling
  else
    set the style of field 4 to rectangle
  end if
  --other commands as needed here
end openCard
```


Figure 2: *A card illustrating the use of automatic scrolling fields. The script in listing 8 determines which of the four meal list fields needs to be scrolling and which should be rectangular for a given card in this background.*

a background field representing that meal. Each card represents a different day in the log. Because I did not know how many food items would be chosen for a given meal, I had to make use of scrolling fields when necessary. My solution is shown in listing 8, with a sample card shown in figure 2.

The script in listing 8, placed in the background script of the food log, checks each field to see whether it con-

tains more lines of text than are visible for that field. If it does, the field style is set to scrolling; otherwise, it is set to rectangle. Note that there is no problem in setting a field to a style it already has. Because these are background fields, they have to be updated each time a different card is opened. While the text of a background field is specific for a given card, the style of that field will remain the same for every card in that back-

ground unless it is changed by a script, as in listing 8.

Font Control

I would like to make one final point regarding text fields. If you intend to distribute your stack, it is important to ensure that the text font you choose will not be changed drastically when your stack is run on another Mac. I was made acutely aware of this problem when I ran one of my stacks on another machine recently. I had some fields that were originally set for the Boston II font. This text was barely recognizable on the other Mac, as my carefully measured words were converted to an ornate font that forced undesired line returns, thus cutting off part of the text at the bottom of the field. The explanation was simple: The System file on this Mac didn't have the Boston II font, and another, quite inappropriate, font (with the same font ID as Boston II has on my system) had been substituted in its place.

There are two ways to prevent this problem. One is to use Font/DA Mover to install the desired font on your stack (hold down the Option key while selecting open from Font/DA Mover). I don't favor this approach, however; stacks grow in size too rapidly as it is, without loading them up with fonts. A better approach is to restrict the font styles used in your text fields to those fonts required by the system: Geneva, Chicago, or Monaco. If you must use a fancy, exotic font in your stack, don't place the text in a field at all. Simply select the text tool and do your writing on the graphics layer. Text that you create in this way exists as a bit-mapped image that will not be affected by system fonts.

More Fun with HyperTalk

I have presented some guidelines for programming HyperCard stacks that I hope you will find useful. HyperTalk has its limitations as a programming language, but its utility for many users is limited only by their imagination and creativity. HyperCard has made programming the Mac more fun than ever.

By attending to the issues I have raised here, HyperTalk scripters will be more likely to create stacks people will value and enjoy using. I'm looking forward to seeing more of them in the future. ∎

Richard D. Lasky is a biochemist, Macintosh enthusiast, and certified Apple developer. He is the author of Nutrition Stack, which calculates the nutritional content of meals. He can be reached on BIX c/o "editors."

FORTRAN Power for Macintosh

Chosen #1 by ...
MACWORLD (March 1989)

Judged Most Powerful by ...
Computer Language (Nov. 1988)

Full ANSI 77, VAX-compatible extensions, high precision IEEE floating point data types, complete access to the Toolbox, direct code generation for the 68020/030 and 68881/2, and Macintosh Programmer's Workshop (MPW) make the Language Systems compiler the most powerful FORTRAN for the Macintosh.

Language Systems FORTRAN has a full complement of VAX-compatible extensions including structures, nested includes, DO WHILE, IMPLICIT NONE, and all VAX intrinsic functions and data types. Supports even the largest subroutines and functions from VAXes and mainframes.

$345 (plus shipping).

Runs on all Macintoshes with a hard drive and at least one megabyte of memory, including the new IIx, IIcx, and SE030.

LANGUAGE SYSTEMS
CORPORATION
441 Carlisle Drive, Herndon, VA 22070
(703) 478-0181   fax (703) 689-9593

INERTIA

Engineering in Common Sense

Analysis & Design

For Manufacturers and Designers of Mechanical, Structural, Electrical, Aeronautical, and Architectural systems. Productivity via graphical modeling, speed and power. Free video tape and seminars.

MCAE Inc.
Tel: 800-444-6223
Fax: 317-463-7166

BYTE's All-Macintosh Supplement

Free Information

NEW!

Al MAC

Now you can receive free information on the Macintosh products and services advertised in BYTE's All-Macintosh Supplement. Simply refer to the Reader Service Card, located in the back of BYTE, and **look for the boxed numbers (451–480)**.

- Circle the numbers on the reply card that correspond to the numbers assigned to items of interest to you.
- Check all the appropriate answers in questions "A" through "C."
- Print your name and address, and mail.

Take this opportunity to receive information direct from Macintosh suppliers TODAY!

READER SERVICE

To get further information on the products advertised in BYTE, fill out the reader service card by circling the numbers on the card that correspond to the inquiry number listed with the advertiser. This index is provided as an additional service by the publisher, who assumes no liability for errors or omissions.

* Correspond directly with company.

Alphabetical Index to Mac Advertisers

Index to Mac Advertisers by Product Category

World Radio History

Neural Networks

Y ou could describe neural networks as humanity's attempt to create an artificial brain— shades of science fiction. In their current stage of development, however, it would be more correct to describe them as humanity's attempt to mimic the way the brain does things in order to harness its versatility and its ability to infer and intuit from incomplete or confusing information.

What happens when you learn something? Most of us would probably answer with words like remembering, understanding, storing, and retrieving. But there's more. Brain surgeons or behavioral psychologists might discuss firing neurons, making new connections, or retraining behavior patterns. But even they can't tell you *exactly* what happens when you learn—or how.

To find out, you can observe and record the tangible inputs to the learning process as well as the end result. You can show how learning varies from person to person depending on the pattern of inputs and on such intangibles as past history, emotional state, and so on. Then you can surmise from these elements some of what has occurred.

Learning about neural networks requires a new vocabulary. You don't program a neural network, you "teach" it. You measure its speed not in instructions per second but in interconnections per second.

This month's In Depth section defines and describes neural networks, their differences from traditional computing, and their implications and uses in the microcomputer arena.

In "Time to Get Fired Up," Klaus K. Obermeier and Janet J. Barron provide a look at today's neural-network technology and show how IBM PCs, Macintoshes, and personal workstations can run neural-network simulations to solve problems that digital means can't handle efficiently. And the text box "In an Upscale World" by Kingsley G. Morse Jr. explains the dynamics of neural-network scalability, going from a sample-size network to a real-world application.

Neural networks have input and output like conventional computing, but what happens in between the two has long been a mystery. In "What's Hidden in the Hidden Layers?" David S. Touretzky and Dean A. Pomerleau show how you can determine what lies in between— what's really going on.

Speech recognition is a complex task for which even the largest computers are not particularly well suited. Neural networks, however, have the flexibility to interpret complex and confusing audio signals. In "Building Blocks for Speech," Alex Waibel and John Hampshire show how neural networks can be used to create high-performance speech-recognition systems.

Neural networks may sound like science fiction, but they aren't. As this month's resource guide, "Neural Networks: Theory and Practice," will show, they are the basis for real microcomputer products. Science fiction is known as a domain of visionaries, a field that often leads the way to the future. While an artificial brain may still reside in the world of science fiction, neural networks have bridged the gap to become science fact.

—*Jane Morrill Tazelaar*

Sometimes Emulations Outrange the Original Digitizers & Mouse?

The Podscat is remarkable example for a sucessful solution of the old price-performance problem:

- Emulation mode and configuration menu allow tablet top selections.
- A built-in single chip computer emulates 9 of the worlds most renowned digitzers.
- Built-in EEPROM allows custom configuration, stored as the digitizer's default setting at "power-on".
- Emulating and surpassing MS Mouse & PC mouse, it is more accurate than any optical mouse.
- Button 'click' on/off is selectable on the tabled top.
- Ultra stable jitter = 0 of the least significant bit.

And, of course, configurations can be loaded down from the host computer.

What do you say? We're waiting for your comments and inquiries.

TABLET MODE														
♩	SETUP	S U M M A B	S U M M A A	G T C O B	G T C O A	H I T C H A	H I P A D A	P T E C A	P T E C B	C A L W	C A L A	C A L B	K B D B	K B D A	M S S	M S S	P C M S	P C M S

ACTIVE AREA	12" x 12"
SETUP AREA	12" x 0.4"
RESOLUTION	1,000LPI
SAMPLE RATE	150 PPS
ACCURACY	±0.01 INCH
JITTER	≤ 0 LSB
PROXIMITY	±0.5 INCH
DIMENSIONS	400mm × 390mm × 25mm
INTERFACE	RS-232C

US$189

Circle 11 on Reader Service Card (DEALERS: 12)

Time to Get Fired Up

IBM PCs, Macs, and personal workstations can run
neural-network simulations that learn and train themselves

Klaus K. Obermeier and Janet J. Barron

Every time that you use a high-speed modem, you are using a single-neuron, many-synapse, neural network. This tiny neural network uses adaptive signal processing, learns the system, and eliminates some of the problems (such as echoing) that may occur. This adaptive filtering process is just one of the neural networks developed in the 1950s by Stanford University's Bernard Widrow, pioneer in and founder of the field of neural networks.

Another of Widrow's early neural-network applications was a simple weather-forecasting system in 1963. When fed many samples of yesterday's pressure and today's weather, it came up with tomorrow's forecast. Widrow's system was correct about 83 percent of the time (compared to an accuracy rate of about 65 percent for the local meteorologist).

For some time, there has been a need for a way to solve problems that cannot be efficiently handled by digital means. A neural network is composed of many interconnected processing elements that operate in parallel. It works in a way similar to how we think the neurons in

the human brain encode information.

Instead of programming a neural network, you "teach" it to give acceptable answers. You input known information, assign weighted values to the connections within the architecture, and run the network (which adjusts those weights by using several criteria) over and over until the output is satisfactorily accurate. A weighted matrix of interconnections

allows neural networks to learn and remember. As a result of the way they work, even when you enter new information that is not stored in the network, they can still provide adequate responses.

Neural-network technology, also called *connectionism*, is moving very quickly, and working tools are rapidly coming into use. As they emerge, you'll be able to use that technology to resolve issues that don't have straightforward black-and-white, yes-or-no answers.

When they work correctly, neural networks provide some major benefits, such as the ability to take incomplete data and produce approximate results. Their parallelism, speed, and trainability make them fault-tolerant, as well as fast and efficient for handling large amounts of data.

But, because neural networks work as we believe the human brain does, they don't handle numbers well, especially if you need accurate answers and you need them fast. Accuracy, computational power, and logic are not among their strong points. And when they solve a problem, they can't tell you how they did it. At this early stage in the technology
continued

curve, the "real things" (biological neural-network clones) are not available. What we have are simulations (artificial neural networks) that run on digital machines, and they are good at pattern recognition and functional synthesis.

Today, artificial neural networks are being used for a variety of commercial applications, including speech, character, text, equipment, and human recognition tasks; financial analysis; database management; image and signal processing; medical diagnosis; dealing with fuzzy, chaotic, or incomplete information; and some kinds of manufacturing, quality, and process control.

Biologically Inspired

Artificial neural networks are biologically inspired. A biological neuron consists of axons, dendrites, and synapses. An artificial neuron, or processing element, emulates the axons and dendrites of its biological counterpart with wires and emulates the synapses by using resistors with weighted values.

Essentially, neural-network models consist of processing elements, interconnection topologies, and learning schemes. Processing elements contain combinations of excitatory (positive) or inhibitory (negative) weights that act on the inputs in a summation function, in an activation function based on the inputs to the processing element, and in an output function that is both sigmoid and stochastic.

Processing elements interact with each other depending on how they are interconnected—fully (as opposed to partially), or with or without a feedback loop. As part of setting up the neural network, a variety of criteria is used to define specific interconnections and determine its characteristic architecture. The nature of its feedback loops determines the network's trainability; the degree of its interconnection determines its parallelism.

While a digital computer's memory is measured in bytes, a neural network's "memory" is judged by interconnections. Likewise, while the speed of a digital computer is expressed in instructions per second, the neural network's speed is measured in interconnections per second.

Most cognitive processes take humans no longer than a few hundred milliseconds, while individual neurons in the human brain compute operations at a rate as slow as that of a single instruction of a digital computer. The brain performs its processing feat through massive parallelism, using 10 billion neurons and more

than 1000 times that many interconnections.

Training the Network

To simulate massive parallelism, the neural-network approach consists of setting up a network of processing elements, the electronic analogy to neurons. Each processing element has a number of inputs, a small set of possible states, and an output that is a function of the inputs. Each input to the processing element has a weight value, which usually ranges from 1 to −1.

When a processing element is activated, it evaluates all its inputs and computes their respective weight values. If the weight value is above a certain threshold, the computing unit generates an output value that is used as input by other processing elements. (Only the weight values of the inputs change during learning.)

Training a neural network is a matter of adjusting weights, either manually or automatically. A neural network is a directed graph consisting of a number of nodes, or *processing elements*. Each processing element has only one output signal, which fans out to interconnect with other processing elements. Each node processes the incoming signal based on the values of the constants stored in it. Currently, neurocomputer technology is based on the assumption that the update of signals within each node occurs discretely, rather than continuously or concurrently.

Neural-network learning takes place in one of three ways: supervised, unsupervised, or self-supervised. Supervised learning occurs when you provide trial-and-error inputs, teaching the network correct and incorrect responses. In unsupervised learning, data is simply entered, without human intervention. This process leads to internal data clustering—the desired result. Self-supervised learning occurs when the network monitors itself and corrects errors in the interpretation of data by feedback through the network.

A neural network computes by the process of spreading activation. After the initial weights are set, you enter data into the network; this process causes it to pass through state changes and ultimately reach stability. A network achieves stability when the weight values that are associated with the processing elements stop changing.

When neural networks first became popular, they consisted of only one or two layers—an input and/or an output layer. This severely limited what the net-

work could represent. Adding more layers allowed the system to form an internal representation of the problem. Networks with only one layer (made popular by Frank Rosenblatt and unpopular by Marvin Minksy and Seymour Papert) thus restricted what could be represented to what was in the input configuration.

Today's multilayer, hierarchical networks are more powerful because they can generate their own internal representations in the so-called hidden units. Hierarchical networks are used for the better-known applications, such as speech and character recognition.

A hierarchical network consists of an input and output layer and one or more hidden layers (see "What's Hidden in the Hidden Layers?" by David S. Touretzky and Dean A. Pomerleau on page 227). If the number of processing elements in the middle layer is too great, it will replicate the elements from the input layer, causing problems similar to those encountered with a single-layer network. If the number of processing elements in the middle layer is too small, the network will require many iterations to train, and recall accuracy will suffer.

All This on a Micro?

IBM PCs and compatibles, Macintoshes, and personal workstations play very important parts in the neural-network world. You can run simulations on them and, in some cases, perform neural-network development and experimentation on them as well.

Neural networks are being used and produced in the form of either *neurocomputers* (hardware that models the parallelism of neurons) or *netware* (software that emulates neurons and their interconnections on conventional serial computers). An important aspect of netware is that it can be simulated on conventional computers.

Neurocomputers have been configured on the chip level, the board level, and the complete system level. General-purpose neurocomputers are available to use as coprocessors for digital computers. In this case, you access the neural network as if it were a subroutine that you can call whenever you need it. In this form, neurocomputers are able to operate side by side with conventional computer technology.

Last summer, NEC announced that it had developed a personal neural-network computer that uses the back-propagation learning algorithm. NEC's current plans are to market and sell the Neuro-07 only in Japan. The total system, which sells for about $11,000, consists of a personal

Glossary

activation function A function by which new output of the basic unit is derived from a combination of the net inputs and the current state of the unit (the total input).

auto-associative (memory system) A process in which the system has stored a set of information repeatedly presented to it. Later, when you submit a similar pattern to the system, it can recall the information from a degraded or incomplete version of the original.

axon That part of a nerve cell through which impulses travel away from the cell body; the electrically active parts of a nerve cell.

back-propagation A learning algorithm for a multilayer network in which the weights are modified via the propagation of an error signal "backward" from the outputs to the inputs.

chaos The study of nonlinear dynamics (also called deterministic disorder).

connection A pathway between processing elements, either positive or negative, that links the processing elements into a network.

dendrite The branched part of a nerve cell that carries impulses toward the cell body. The electrically passive parts of a nerve cell.

directed graph Representation of the variation and direction of flow for processing elements with respect to other processing elements.

feedback loop A loop wherein continued input is fed back into the network to achieve the expected output.

fuzzy logic Incomplete or contradictory information.

hidden layer A third layer of units between the input and output layers that provides additional computational power.

learning The phase in a neural network when new data is introduced into the network, causing the weights on the processing elements to be adjusted.

network paradigm A network architecture that specifies the interconnection structure of a network.

neuron The structural and functional unit of the nervous system, consisting of the nerve cell body and all its processes, including an axon and one or more dendrites.

perceptron A large class of simple neuron-like networks with only an input layer and an output layer. Developed in 1957 by Frank Rosenblatt, this class of neural network had no hidden layer.

sigmoid Having a double curve like the letter S.

spreading activation A process of applying the activation function simultaneously to a neural network.

stochastic Involving chance, probability, or a random variable.

summation function A function that combines the various input activations into a single activation.

synapse The point of contact between adjacent neurons where nerve impulses are transmitted from one to the other.

threshold A minimum level of excitation energy.

training A process whereby a network learns to associate an input pattern with the correct answer.

weight The strength of an input connection expressed by a real number. Processing elements receive input via interconnects. Each interconnect has a weight attached to it. The sum of the weights make up a value that updates the processing element. The output value of a processing element is described by a level of excitation that causes interconnects to be either on (i.e., excitatory output) or off (i.e., inhibitory output).

---

computer, a neuro-engine board, neural-network learning software, and a color display.

The neuro-engine board performs parallel processing with a maximum speed of 216,000 interconnections per second. Its software is composed of a definition section to determine the network's configuration, a computing section to calculate the network's output, a software-control section, and a user interface to perform editing and monitoring functions.

In 1988, about 10,000 personal computer packages of neural netware were sold in the U.S., most of these from a disk included with *Explorations in Parallel Distributed Processing* (see reference 1). In general, commercially available neural-network programs are those that lend themselves to simulation on very small scales—based either on the soft-

ware itself or on special-purpose boards.

James A. Anderson, professor of psychology and cognitive and linguistic sciences at Brown University, notes that he teaches a course in neural networks for undergraduate and graduate students. Most of them, he says, do the simple assignments on their home computers—Macintoshes and IBM PCs. But, says Anderson, these machines with their standard compilers can't cope with networks that have between 50 and 100 processing elements.

"It's not the MIPS [millions of instructions per second] a device can handle that determines whether or not you can use it for neural networks," Anderson says. He notes that simple measurements of processor speed are especially misleading because many personal computers are fast but are unable to handle large arrays or matrices. Effective mem-

ory management, large memories, and good compilers are much more important than raw CPU speed in performing neural-network computations quickly, he explains.

"Engineering workstations—VAXstations, Suns, and so forth—are ideally suited for the task, but even on fast workstations, jobs may run for hours. Many personal computers completely run out of steam when faced with a system with 150,000 connection strengths and 400 dimensional arrays, whereas workstations are designed for large jobs. Again, it's how good your compiler is," Anderson says. "Suns and VAXes—especially VAXes—have wonderful compilers. But you can do useful development work on personal computers by learning a lot, experimenting a lot, and taking the time to run your own assembly language
continued

projects."

But not everyone shares Anderson's opinion about why small systems currently have limited neural-network capabilities. There are other reasons as well. Smaller machines have problems running certain large applications, such as complex vision systems, in real time. Digital machines simulate what are intrinsically very parallel systems, and they are limited by their own speed and processing power. Therefore, while large problems, like analyzing scenes, are difficult to run on a personal computer or workstation, less-complex tasks are very workable. Presently, there are about 300 companies involved in neural-network technology, many of them making netware for personal computers or workstations.

Among the products spilling out of the neural-network pipeline are software, shells, development tools, chips, and accelerator boards. Some companies are in the process of developing special-purpose hardware and chips for use in large-scale applications. The next year should bring the introduction of many products that go beyond the simulation stage.

Applying the Knowledge
Neural-network applications tend to fall into several classes: sensor and knowledge processing, pattern recognition, and control systems. Neural networks are not very good at handling tasks that standard serial computers are noted for, such as number crunching or making highly accurate calculations. But when it comes to tasks requiring incomplete data sets, or fuzzy or contradictory information, neural networks will very likely outperform conventional computers, including parallel processors.

Massive parallelism gives neural networks a high degree of

• *fault tolerance*—built-in redundancy or the ability to withstand component failures without crashing;
• *associative recall*—the ability to retrieve information instantaneously based on content and to make an "educated" guess if there's no exact match for the requested information; and
• *graceful degradation*, the ability to recover gracefully from processor failure.

These properties make neural networks attractive for many commercial, military, and industrial applications.

One interesting example of a combination of neural-network applications is called SNOOPE (for System for Nuclear On-line Observation of Potential Explo-

sives). Developed by Science Applications International Corp. (SAIC) of Santa Clara, California, SNOOPE is a detection system that determines the existence of concealed plastic explosives in luggage and cargo.

Successfully tested since June 1988 on 40,000 bags and luggage items at the San Francisco and Los Angeles International Airports, SNOOPE is a neural network based on a back-propagation supervised-learning algorithm. The network runs in parallel with another technique called *thermal neutron analysis*.

SAIC was given certain criteria for the system: It had to continuously process 10 bags a minute, not damage film or magnetic recording media, be reliable, and be built from commercially available components wherever possible. The output, a decision as to whether or not a bag contains a threat, must be signaled by the time the bag exits the system.

The first SNOOPE system was due to be installed at New York's John F. Kennedy International Airport in July. After that, others are slated for installation in airports around the world. Says Samuel K. Skinner, U.S. Secretary of Transportation, "It is the best available technology to detect explosives. . . . Detection is performed by computer. No human interpretation is involved."

Sensor processing and pattern recognition are among the many ways in which neural networks are being implemented. Applications include image processing, image compression, character recognition, and continuous speech recognition.

You can use these types of neural networks to recognize underwater targets by sonar. Bendix Aerospace compared a neural-network program with a conventional program. The results showed that the neural network not only was better but also took only hours to be configured, as compared to the months it takes to set up a conventional classifier-based program.

Programs for handwriting character classification also fall into the sensor-processing category. NestorWriter, produced by Nestor in Providence, Rhode Island, for instance, can figure out some recognition rules based on common character features, such as curvature and orientation; thus, the system can recognize characters it hasn't seen before. Applications for this technology range from processing checks to reading Japanese characters.

Among neural-network pattern-recognition and control-system applications are programs for robotics and autonomous vehicles. One of the oldest exam-

ples of control-system neural networks is adaptive routing and switching. Widrow's classic Adaline (for adaptive linear element) is a program that eliminates echoes in telephone lines. The same principles can be used to reduce data-transmission errors in modems.

Neural networks are efficient at handling many knowledge-processing tasks, such as storage and retrieval of information in large databases, and predictive modeling. In one medical expert-system application, a neural network was trained on the functional relationships between symptoms, diagnoses, and treatments. Test results showed that the network responded with 100 percent accuracy to nonequivocal cases, weighed the evidence in equivocal cases, and, if unknown cases were presented, fell back on known relationships.

The neural network was configured in only a fraction of the time it would have taken a knowledge engineer to configure and build the expert system. Besides showing new conceptual solutions, in certain applications neural networks seem to avoid the impasse of having to laboriously construct and maintain expert systems.

In the area of speech synthesis (see "Building Blocks for Speech" by Alex Waibel and John Hampshire on page 235), a program called NETtalk was jointly developed by Terence Sejnowski of the Salk Institute in La Jolla, California, and Charles Rosenberg of Princeton University. NETtalk provides an impressive demonstration of the potential of neural-network technology. The program learns to read English text aloud without the benefit of any preprogrammed linguistic rules. In contrast, conventional programming techniques (including AI programming) have real problems executing this function.

Current Events
The study of neural networks, the "reborn" science, has gone from great promises in the 1940s, with the age of McCulloch and Pitts, to the Widrow and Rosenblatt era in the 1950s, through attacks on the field in the 1960s from Minsky and Papert in their book, *Perceptrons* (see reference 2). From there, it moved into a strong and legitimate revival in the 1970s and 1980s with Grossberg, Kohonen, Hopfield, Rumelhart, and others. Because of computational advances, significant progress has been made ever since the so-called "perceptron" era.

Leading-edge neural-network tech-

In an Upscale World

Kingsley G. Morse Jr.

When you want to enlarge a small experimental neural network into a real-world application, scalability becomes important. Even with the limited resources of a microcomputer, you can train large neural networks quickly if you're careful. The key is to make the neural network's training algorithm scalable. "Scalable," in this case, refers to the ability of a neural network developed on a microcomputer to be enlarged easily to perform larger—sometimes much larger—real-world tasks.

Although scalability is still more of an art than a science, several techniques exist to help you stay within the speed and memory limits of a microcomputer. You want the network to be able to use more neurons, synapses, or training patterns and still train reasonably fast on a microcomputer.

The graph in figure A compares three scalability standards. As you can see, an algorithm whose training time is a polynomial function of the number of training patterns will allow you to train many more patterns than an algorithm whose training time is an exponential function of the number of training patterns. If you improve a training algorithm from exponential to polynomial scalability, you will significantly increase the number of patterns you can train; achieving *linear* scalability would be extraordinary.

The following techniques will help you make a scalable neural-network training algorithm.

• Use "computational-complexity functions" to estimate how scalable an application and a neural-network training algorithm are. Benchmarking scalability early can save you from wasting time on untrainably large networks or applications.

First, you plot the training time (or memory) against the "size" of the problem; then you fit several curves to the data. The curve that best fits the data is its computational-complexity function. If training time increases exponentially with the problem size, the training algorithm and the application aren't scalable, and you should consider other training algorithms or applications. If, however, training time is a linear or polynomial function of problem size, it should be scalable. Once you have approximated the complexity function, you can estimate how long it will take to train larger problems on your microcomputer. This a good way to benchmark various training algorithms.

For example, you could use the following method to benchmark a training algorithm that interests you. Train the network with several small but different sets of training patterns. Keep a record of the training times and corresponding numbers of training patterns. Plot these points on a graph with training times (in minutes) on the graph's vertical axis and the number of training patterns that you used on the horizontal axis. Then try to fit the data with combinations of the standard components of computational complexity (i.e., linear, polynomial, and exponential terms). Choose the curve that fits the data points best.

A typical polynomial computational-complexity function for back-propagation's training time is:

$$178 + (0.014 \times (\text{number of training patterns}^{3.535}))$$

You can use this function to estimate how long it will take to train more patterns; this is a good measure of the training algorithm's scalability. You can also use this technique on other training algorithms and compare them to the first algorithm.

This technique is also good for avoiding applications that are unscalable no matter which training algorithm you use. If you've tried several training algorithms, and all of them have strongly exponential computational-complexity functions, then you might want to consider another application.

• Avoid *second-order* training algorithms that use memory proportional to the network's size squared. For example, some methods store and update a matrix of second derivatives, where the memory required is proportional to the square of the number of synapses in the network. The training algorithm must also update each element in these arrays, so the training time to maintain second-order data can increase with the square of the number of synapses. A square relationship is an example of the polynomial curve on the graph. Although second-order methods converge rapidly for smaller problems, they become unwieldy for large networks. First-order methods only use memory proportional to the number of synapses; the linear curve in the graph illustrates this.

For example, if you want to expand a microcomputer's neural network from 100 synapses to 1000 synapses, and you're using 4-byte floating-point numbers, the amount of storage that a second-order technique needs would increase from $4 \times (100^2) = 40{,}000$ bytes to $4 \times (1000^2) = 4{,}000{,}000$ bytes.

This is a hundredfold increase, and most microcomputers don't have enough memory for the second-order method. A first-order method would increase memory usage from only 400 bytes to 4000 bytes, an amount well within the memory capacity of most microcomputers. Furthermore, you can enhance first-order methods with a "momentum" term and conjugate gradient techniques.

• Avoid training algorithms that are known not to scale well. For example, back-propagation becomes unstable as more layers are added. Some research indicates that a back-propagation net-

---

nology developments are being addressed at universities (e.g., Caltech) and high-technology companies (e.g., TRW, General Electric, and Texas Instruments) before being farmed out to start-ups (e.g., Nestor and Hecht-Nielsen). Neural-network technology cuts across many disciplines, including psychology, biology, physiology, philosophy, mathematics, physics, computer science, and linguistics.

Recent advances in the fields of mathematics, neurology, and neurobiology have led to a neural-network reawakening. Consequently, the theory of neural networks is being studied in two aspects: first, the efficiency of a neural-based electronic architecture, and second, the achievement of an understanding of the biological functions of neural networks.

There is a variety of factors holding back the widespread implementation of neural networks. The technology itself is

work's training time grows exponentially as the number of neurons increases. In comparison, when you can use multiple-regression algorithms for inherently linear applications, the training time is only a polynomial function of the number of inputs. The good news is that back-propagation appears to scale polynomially as the number of training patterns increases.

• Use the machine-specific characteristics of floating-point mathematics, as some numerical algorithms do. When adding many floating-point numbers, round-off error may preclude smaller numbers from affecting the running total. By keeping calculations within accuracy limits, you can avoid numerical runaway, promote stability, and thus attain faster training.

• Consider training the network until it produces answers that are good but not optimal. For example, if the desired response of the network's output neurons is either 0 or 1, the network will learn "correct" answers faster when you compare the outputs to 0.5 instead of exactly 0 or 1. In other words, anything less than 0.5 would correspond to a 0; anything more, to a 1.

• Consider developing a training algorithm that solves a special case of the application. For example, instead of training an insurance neural network with loan data from all age groups, it may be faster to train several smaller networks for different age groups. In some cases, a linear reformulation of the application is possible, which allows you to use much faster methods, such as multiple regression.

• Do away with unnecessary neurons and synapses, leaving a smaller network to train. If redundant neurons exist and you can remove them, then you may be able to train a smaller network. This could make a noticeable difference if the training time is polynomially or exponentially related to the number of

Figure A: *Three standard measures of scalability; n is the problem size, which can be training patterns, neurons, or synapses, and C_1 and C_2 are constants.*

neurons or synapses.

• Consider using a math coprocessor to speed up multiplication, a common bottleneck for neural-network algorithms. Intel, Weitek, and Motorola, among others, sell math coprocessors. Also, some compilers are more efficient than others for numeric processing. BYTE's March 1988 In Depth on floating-point processing outlined some hardware and software options.

• To speed up your neural network that last little bit, you may want to rewrite some parts of the code in assembly language. For example, many neural networks spend a lot of time evaluating the activity of the network. This part of the code may be a good candidate for assembly language.

• Look to neuroscience for ways to train neural networks quickly. A staggering amount of neuroscientific knowledge is available. Biological elements such as neurons and synapses have traditionally inspired neural-network research, but what roles do genetics, cortical columns, and the hypothalamus play?

• Watch for benchmark results from the Defense Advanced Research Projects Agency. DARPA has budgeted millions of dollars for neural-network research. Specifically, it intends to fund benchmarked comparisons between neural-network algorithms and classic pattern-recognition algorithms. Hopefully, this research will address benchmarking techniques that can be applied to scaling up neural-network training algorithms.

BIBLIOGRAPHY

Morse, Kingsley G. Jr. "Neuroscience as a Possible Research Standard for Neural Networks." *Proceedings of the IEEE First International Conference on Neural Networks,* 1987.
Vose, G. Michael, and George A. Stewart. "Floating-Point Processing." BYTE, March 1988.

Kingsley G. Morse Jr. founded the AI Forum of Silicon Valley and works at Hewlett-Packard in Mountain View, California. He can be reached on BIX c/o "editors."

still in its infancy. Many elements have yet to be worked out and put into place— not the least of which is how to model the human brain. We still understand very little about how the brain works, and so far, no one has been able to come up with a "brain in a box." We don't even know whether or not we really *want* to model

our biological neural networks. Other ways of implementing the technology may prove to be more effective.

Today, some of the major neural-network dilemmas concern training/learning, scaling, and performance. One current area of research is trying to identify the network paradigm, or pattern, best

suited for a specific application. There are dozens of known network paradigms, and the number is steadily increasing.

Currently, neural networks have meager processing power, even when compared to the brains of such simple creatures as cockroaches, flies, and leeches.

Although in principle the networks are capable of handling raw data well, there may be severe practical limits in scaling neural networks (see the text box "In an Upscale World" on page 222).

We have a long way to go before we understand a neural network's learning capabilities. Currently, we know little about our own biological memories. Thus, we don't know what, if anything, distinguishes learning from recall. In addition, the current neural-network learning algorithms are neither very novel nor powerful.

The training effort for large-scale applications may be as substantial as that required to program conventional computers. Because so little is known about why a neural network behaves in a certain way, there are still risks in overtraining the network and constructing inefficient hidden layers. The state of the art is still hindered by the limits of the hardware.

A neural network's performance depends on many elements. Some of today's most important issues are: How many layers and processing nodes are enough? How creative should the system be (i.e., how many times should it "guess" before it gives up)? If it finds one good answer, should it continue to search for another? What happens if neurocomputers base their conclusions on data other than what we use? Once a neural network has reached a conclusion, what should it do about contradictory evidence?

Neural-network computing has I/O constraints, just as conventional computing does. The basic problem remains: Unless the communications channel is relatively large compared to the system's average total communications load—even compared to its occasional near-peak loads—the system's behavior will significantly deteriorate.

As we approach the twenty-first century, we need a new approach to the information science describing neural-network machines. Just as there is a formal structure to our biological neural network, more efficient artificial neural networks will need a framework and an order to determine how they will learn, preprocess, and select input information. They will also need to deal with how different parts of an intelligent system will perform specific functions.

One of the areas of study being explored by David Rumelhart, professor of psychology at Stanford University, is that of developing networks that can choose their own architectures. Still in its most rudimentary stages, this science

will use a kind of a built-in feedback loop as people use neural models to solve relatively specialized problems and learn from their experiences.

Robert Hecht-Nielsen, cofounder of the Hecht-Nielsen Neurocomputer Corp. in San Diego, suggests that you check out at least six criteria when you choose a neural-network configuration:

1. optimal I/O format
2. training time
3. data preprocessing requirements
4. mathematical optimality
5. performance estimates
6. debugging/diagnostics requirements

In addition, you should also ask an important question: Can you achieve the same or better results with conventional technologies?

A Marriage of Convenience

According to optimistic predictions, by the year 2000, neural-network technology will account for half the total revenues of the robotics and computer markets. With little but pure research to build on and no concrete knowledge of how the brain really works, the last few years have brought products to market that range from simulation software to a neural network implemented in a chip, hard-coded to duplicate a neural-network architecture.

One company, Oxford Computer in Oxford, Connecticut, has developed an intelligent memory chip that can be used for neural networks. According to Steve Morton, founder and chief technical officer of the firm, because these chips are inherently parallel, you can combine them to build powerful board-level neurocomputers performing tens of billions of operations per second.

Why this rapid growth in neural-network technology and impressive list of products and technology implementations just a few years after the resurgence of interest in the field? Primarily because of the time-urgent need for an alternate way to solve problems that conventional processing techniques don't handle well. In addition, over the last few years, there have been important mathematical and computational advances.

But neurocomputing must overcome many significant barriers before it can become an accepted way to solve real-world problems. The future of practical neural networks depends on the advent of technologies that support their speed and storage requirements. The interconnects per second found in the brain of a common housefly are two orders of magni-

tude faster than the fastest neural-network tool available today.

In the short term, the development of the digital signal processing chip will support improvements in speed, while the advent of DRAM chips of up to 16 megabits will increase storage capacities. In the midterm, better gallium-arsenide chips will help to improve speed, and better wafer and analog devices will improve capacities. In the long term, optical computing and optical storage will be key factors.

Neural networks won't replace database and knowledge-based processing because they don't work well with numbers or cut-and-dried information. In the next few years, it is likely that the first practical neuron-like circuits will appear in silicon, and a neural network may be used as a coprocessor controlled by a host digital computer. One company, Micro Devices of Orlando, Florida, has already produced a chip on a board that it claims is a working neural network, not a simulation.

Most forecasters believe that neural networks will not replace conventional methods of computing—especially those that deal with high-speed numeric processing—but will complement them and add to their utility. The combination of traditional computers and the unique power of neural networks could unravel problems that otherwise would remain unresolved.

In spite of all the hype and excitement, however, the verdict is still out and will remain so for about the next 10 years. Exaggeration has been and still is the bane of the neural-network industry. Everyone deeply involved in this field continually and appropriately warns against the setbacks that can occur if hype becomes the order of the day. ■

REFERENCES
1. McClelland, James L., and David E. Rumelhart. *Explorations in Parallel Distributed Processing.* Cambridge, MA: MIT Press, 1987.
2. Minsky, M., and S. Papert. *Perceptrons.* Cambridge, MA: MIT Press, 1969 and 1988.

Klaus K. Obermeier is the natural-language-query product manager at Battelle Institute in Columbus, Ohio. He received his Ph.D. in linguistics/AI from Ohio State University and has established a neural-network technology clearinghouse. He can be reached on BIX c/o "editors." Janet J. Barron is a technical editor for BYTE. She can be reached on BIX as "neural."

MICRO-CAP III.™
THIRD-GENERATION INTERACTIVE
CIRCUIT ANALYSIS. MORE POWER.
MORE SPEED. LESS WORK.

MICRO-CAP III,™ the third generation of the top selling IBM® PC-based interactive CAE tool, adds even more accuracy, speed, and simplicity to circuit design and simulation.

The program's window-based operation and schematic editor make circuit creation a breeze. And super-fast SPICE-like routines mean quick AC, DC, Fourier and transient analysis — right from schematics. You can combine simulations of digital and analog circuits via integrated switch models and macros. And, using stepped component values, rapidly generate multiple plots to fine-tune your circuits.

We've added routines for noise, impedance and conductance — even Monte Carlo routines for statistical analysis of production yield. Plus algebraic formula parsers for plotting almost any desired function.

Transient analysis

Schematic editor

Monte Carlo analysis

Modeling power leaps upward as well, to Gummel-Poon BJT and Level 3 MOS — supported, of course, by a built-in Parameter Estimation Program and extended standard parts library.

There's support for Hercules,® CGA, MCGA, EGA and VGA displays. Output for laser plotters and printers. And a lot more.

The cost? Just $1495. Evaluation versions are only $150.

Naturally, you'll want to call or write for a free brochure and demo disk.

1021 S. Wolfe Road,
Sunnyvale, CA 94086
(408) 738-4387

MICRO-CAP III is a registered trademark of Spectrum Software.
Hercules is a registered trademark of Hercules Computer Technology.
IBM is a registered trademark of International Business Machines, Inc.

Automation For The 90's From The Heartland

CORPORATE COMPUTERS OF IOWA provides the nation's most aggressive prices on all TOSHIBA®, EPSON®, NEC® and PANASONIC® printers, FAX machines, laptops and cellular phones.

Call one of our account executives today and experience the performance and quality you'd expect from the **"Heartland."**

TOSHIBA LAPTOPS

MODEL	RETAIL	CCI PRICE
T-1000	$1249.00	$ 760.00
T-1200F	$2099.00	$1299.00
T-1200FB	$2399.00	$1485.00
T-1200H	$3499.00	$2159.00
T-1200HB	$3699.00	$2275.00
T-1600	$4999.00	$3000.00
T-3100E	$4299.00	$2599.00
T-5100	$7199.00	$4300.00
T-5200	$9499.00	$5700.00

FREE box of 3.5 Diskettes with every Toshiba

EPSON PRINTERS

MODEL	RETAIL	CCI PRICE
LX-800	$ 299.00	$167.00
FX-850	$ 549.00	$339.00
FX-1050	$ 799.00	$449.00
LQ-500	$ 529.00	$299.00
LQ-850	$ 859.00	$519.00
LQ-950	$ 949.00	$589.00
LQ-1050	$1199.00	$699.00
LQ-2550	$1499.00	$925.00

Corporate Computers of Iowa offers a "FREE" 6 foot cable with every Epson Printer Purchased!!

PANASONIC PRINTERS AND FAX MACHINES

MODEL	RETAIL	CCI PRICE
KXP-1180	$299.95	$179.00
KXP-1191	$399.95	$230.00
KXP-1592	$649.95	$389.00
KXP-1595	$729.95	$429.00
KXP-1124	$529.95	$315.00
KXP-1524	$899.95	$539.00
UF-150 FAX	$1795,00	$955.00
UF-260 FAX	$2795.00	$1479.00

CORPORATE COMPUTERS OF IOWA
386 PERFORMANCE AT THE NATIONS LOWEST 286 PRICING!!!

CCI AT 286/20 MHZ

- 80286 Processor Running At 20/10
- Phoenix Bios/Digital Speed Display
- 80287 Math Co-Processor Socket
- 8 Expansion Slots
- 1 Meg RAM (Expandable to 8 Megs)
- 1.2 Meg Floppy
- 1.44 Meg Floppy
- 80 Meg Seagate ST-4096 (28MS)
- WD Controller

- 14" Multisynch Monitor
- EVGA Graphics Card
- 1 - Parallel - 2 Serial - Game Port
- 101 Key AT Keyboard w/Dust Cover
- DOS 4.01 w/GW Basic
- Surge Suppressor
- 1 Year Parts and Labor Warranty
- 30 Day Money Back Guarantee
- 2nd Year Warranty $49.95

CORPORATE COMPUTERS OF IOWA PRICE
$2,799.00

1 Year On Site Service Contract FREE

Corporate Computers of Iowa Policies

- Toll Free Number 1-800-533-0948 For Ordering and Post Sale Support
- MasterCard, VISA - No Surcharge. American Express - Add 3%
- C.O.D. Cashiers Checks
- 30 Day, Money Back Guaraantee
- Prices Subject To Change In Manufacturers Pricing.
- Allow 10 Days For Personal Checks To Clear
- Call For Corporate Purchase Orders
- Dealer Inquiries Invited

386 25MHz Tower $3999.00

286 12 MHz CALL FOR PRICING

1-800-533-0948

P.O. Box 714
Sioux City, Iowa 51102
712/277-2100
Toll Free: 1-800-533-0948
Corporate Showroom
609 5th Street
Sioux City, Iowa 51101

CORPORATE COMPUTERS OF IOWA

Circle 69 on Reader Service Card (DEALERS: 70

World Radio History

What's Hidden in the Hidden Layers?

*The contents can be easy to find with a geometrical problem,
but the hidden layers have yet to give up all their secrets*

David S. Touretzky and Dean A. Pomerleau

Much of the current fascination with neural networks has to do with their ability to learn. The most popular learning algorithm today is back-propagation, which can be implemented rather easily on a microcomputer (see "Back-Propagation," October 1987 BYTE).

To solve a problem with a back-propagation network, you show it sample inputs with the desired outputs, over and over, while the network learns by adjusting its weights. If it solves the problem, it will have found a set of weights that produce the correct output for every input.

But what has the network learned? Unlike an expert system, neural networks do not automatically explain their reasoning. Whatever knowledge the network acquires is encoded in its numerical weights. It's not easy to decipher the network's solution to a problem when all you have to look at is a set of floating-point numbers.

In the past, the difficulty in interpreting weight patterns contributed to the neural-network mystique. Networks were sometimes billed as magic boxes whose learning algorithms produced solutions unintelligible to mere humans.

Today, we have a better understanding of neural-network learning procedures like back-propagation, and we can analyze, to some extent, the representations that develop. Back-propagation consists of two passes. In the forward pass, inputs proceed through the network and generate a certain output. Then, in the backward pass, the difference between the actual and desired outputs generates an error signal that is propagated back through the network to teach it to come closer to producing the desired output.

Between the input and output layers, there may be additional layers of units, called *hidden units*. When analyzing a network, we study two kinds of hidden-unit representations. First, we want to understand what the weights mean. Second, we want to look at the patterns of activation of units in the hidden layer in response to particular inputs.

Hidden units should really be called "learned-feature detectors" or "re-representation units," because the activity pattern in the hidden layer is an encoding of what the network thinks are the significant features of the input. The two representations (weights and activity patterns) are closely related, but, for some problems, looking at one is more informative than looking at the other.

To understand the hidden-layer representations that real networks develop, look at two examples of geometric problems that have recently been solved by back-propagation. The first is a highly *continued*

nonlinear binary classification problem; the second involves driving a robotic vehicle along a road through a park.

The Unit Square

In two-layer networks, input units connect directly to output units, and each connection has a number, or weight, attached to it. One widely known limitation of these networks is that they cannot compute the XOR function. Introducing a third, hidden layer of units between the input and output layers provides the necessary computational power for XOR.

You can view XOR as a special case of a more general problem: classifying points in the *unit square* (as in figures 1a and 1b). Each point in the unit square is either in class 0 or class 1. In the case of XOR, you consider only the four corners of the square: the points $(0,0)$, $(0,1)$, $(1,0)$, and $(1,1)$. The first and fourth points are in class 0 (0 XOR 0 = 0, and 1 XOR 1 = 0); the second and third are in class 1 (0 XOR 1 = 1, and 1 XOR 0 = 1).

A single artificial neuron computes a linear sum of its inputs and produces either a 0 or a 1 as output. This in effect draws a line that partitions this square into two regions. For all points on one side of the line, the neuron outputs a 1; for all points on the other side, the neuron outputs a 0.

The position and orientation of the line are determined by the weights on the neuron's input connections. You can't draw a single straight line through the unit square so that $(0,1)$ and $(1,0)$ end up in one region and $(0,0)$ and $(1,1)$ end up in the other. Therefore, you can't solve XOR with a two-layer network.

Introducing a layer of hidden units increases the power of the network, since each hidden unit can partition the input space in a different way. The output unit then computes a linear combination of these partitionings to solve the problem.

In the XOR example, a hidden layer containing two units is adequate (see figure 1c). The first unit partitions the space so that it is activated when either input, $(0,1)$ or $(1,0)$, or both, $(1,1)$, are active, as in figure 1a. It has an *excitatory* connection (a positive weight) to the output unit. The network sets the second hidden unit's weights so that it becomes active only when both inputs, $(1,1)$, are active, as in figure 1b. It has a stronger *inhibitory* (negative) influence on the output unit than the excitatory influence of the first hidden unit.

This network correctly solves the XOR problem. When neither input is active, $(0,0)$, neither hidden unit is active, so the output unit remains off. When a single input unit is on, $(0,1)$ or $(1,0)$, the first hidden unit turns on, activating the output unit. If both input units are active, $(1,1)$, both hidden units turn on. Since the inhibitory input from the larger negative weight of the second hidden unit is greater than the excitatory input from the first, the output unit will be turned off.

Hidden units act as feature detectors, or *filters*, for some types of inputs. By combining these features, the output unit can perform more powerful classifications than it can without the hidden units.

Solving Two Spirals

Additional hidden layers allow artificial neural networks to efficiently partition the input space into arbitrary regions and perform complex tasks. One such task is the two-spirals problem, originally posed by Alexis Wieland of the Mitre Corp. in Cambridge, Massachusetts. In this problem, the network must distinguish between points on two intertwined spirals in the unit square (see figure 2).

The black dots are all in class 0, the white dots in class 1. Like XOR, this problem is not linearly separable. There is no way to draw a single straight line through the unit square so that all the black dots end up on one side and all the white dots on the other.

Two of our colleagues at Carnegie Mellon University, Kevin Lang and Michael Witbrock, recently taught a neural network to solve the two-spirals problem and analyzed the hidden-layer representations that developed (see reference 1). Their network, shown in figure 3, has two input units, representing the x and y coordinates of the point, and one output unit. The activation levels of the input units are not restricted to binary values, but they can take on any value between 0.0 and 1.0.

This network has two hidden layers of five units each. The units in each layer receive connections from the units in all layers below it. The connections that skip layers provide direct information pathways from lower layers in the network and allow more flexible hidden-layer representations. Unlike the XOR problem, however, it's not obvious what a good set of hidden-layer feature detectors would look like for this task.

Back-propagation develops the feature

Figure 1: *A network designed to solve the XOR problem.* **(a)** *and* **(b)** *The regions of input space for which the two hidden units are active.* **(c)** *The number inside each unit is its threshold. A unit turns on when its total input exceeds its threshold. The total input is equal to the sum of its input values (each input multiplied by the weight on the line).*

detectors in figure 4. Each square in the figure graphs the response of a single unit to points at various positions in the interior of the unit square. (The squares in figure 4 correspond directly to the circles in the same positions in figure 3.) The brightness of each point in a square indicates the activation level of that hidden unit when the network is shown an

Figure 2: *Training points for the two-spirals problem. Black points should produce an output of 0; white points should produce an output of 1.*

input point at that position.

Units in the first hidden layer divide the input space into two regions along various angles. Units in the second layer use combinations of these first-layer features to produce curved response patterns. The output unit then uses these curved patterns to form successive turns of the spiral.

The imperfections of the solution are an interesting aspect of the way back-propagation works. Notice the bumps and gaps in the spirals that the output unit forms. The network learns to classify all the points in the training set in figure 2 correctly, but it is *underconstrained*: It is not told how to respond to the remaining points in the unit square. Given this kind of freedom, back-propagation almost never develops a perfect solution.

One of the most difficult parts of training neural networks is choosing the training set. You want back-propagation to develop a network that classifies patterns in the training set correctly and also generalizes to new patterns correctly. Providing additional training data and constraining the network architecture are two techniques that reduce excess freedom and clean up the network's representations.

A Road Tracker

Proper generalization is particularly crucial in real-world problems where you can't train a network in advance for every circumstance it might encounter in the field. One such problem we have been working on at Carnegie Mellon is autonomous vehicle navigation (see photo 1 and reference 2).

The goal of ALVINN (for autonomous land vehicle in a neural network; see reference 3) is to drive the NAVLAB vehicle along a winding road. The inputs to AL-VINN are more complex than the coordinates of a single point in the unit square, but they are geometrical in nature.

ALVINN receives two types of sensor inputs from the NAVLAB (see figure 5). One is a 30- by 32-pixel image from a video camera mounted on the roof of the vehicle. (Each pixel in the video image corresponds to an input unit in the video retina.) The activation level of each unit in the video retina indicates the brightness of the corresponding pixel in the video image.

The other input is an 8- by 32-pixel image from a laser range finder. The activation levels of units in the range finder's retina represent its distance from the corresponding area in the image. The darker the color, the closer the object is. A stylized input sample is shown in figure 5. Notice that the tree to the left of the road in the video image shows up as an area of constant brightness in the range finder image. This is because the tree surface is essentially perpendicular to the horizontal range finder beam and, therefore, at a constant distance away.

The two input retinas are connected to a single layer of hidden units, which are in turn connected to the output units. (In other words, all input units are connected to all hidden units, and all hidden units are connected to all output units.) The response of the output layer is a linear representation of the direction in which the vehicle should travel to head toward the center of the road. The centermost output unit represents the "travel straight ahead" condition, while units to the left and right of center represent successively sharper left and right turns.

To drive the NAVLAB vehicle, video and range finder data from the on-board sensors are injected into the input layer. After completing a forward pass, the network reads a steering command from the output layer. The output unit with the highest output value determines the direction in which the vehicle will head.

Training the network is difficult. To develop a hidden-layer representation that generalizes correctly to new situa-

Figure 3: *Lang and Witbrock's network for learning the two-spirals problem. Each unit receives input from all the units in all the layers below it.*

tions, we fed the network road images taken under a wide variety of viewing angles and lighting conditions. It would be impractical to try to collect thousands of real road images for such a data set. Instead, we developed a synthetic road-image generator that can create as many training examples as we need.

To train the network, 1200 simulated road images are presented 40 times each, while the weights are adjusted using the back-propagation learning algorithm. This takes about 30 minutes on Carnegie Mellon's Warp systolic-array supercomputer. (This machine was designed at Carnegie Mellon and is built by General Electric. It has a peak rate of 100 million floating-point operations per second and can compute weight adjustments for back-propagation networks at a rate of 20 million connections per second.)

Once it is trained, ALVINN can accurately drive the NAVLAB vehicle at about 3½ miles per hour along a path through a wooded area adjoining the Carnegie Mellon campus, under a variety of weather and lighting conditions. This speed is nearly twice as fast as that achieved by non-neural-network algorithms running on the same vehicle. Part of the reason for this is that the forward pass of a back-propagation network can be computed quickly. It takes about 200

milliseconds on the Sun-3/160 workstation installed on the NAVLAB.

The hidden-layer representations ALVINN develops are interesting. When trained on roads of a fixed width, the network chooses a representation in which hidden units act as detectors for complete roads at various positions and orientations. When trained on roads of variable

Photo 1: *The NAVLAB autonomous navigation test-bed vehicle and the road used for trial runs.*

Figure 4: *Response function plots for the units in the two-spirals network. Each plot shows the activation level of a single unit as the x,y input to the network ranges over the interior of the unit square. The topmost plot is for the output unit, and the plots below are for the five units in each of the two hidden layers. (Figure courtesy of Kevin Lang and Michael Witbrock)*

widths, the hidden units turn into road-edge detectors, sensitive to only one of the two road edges. (Some look for left edges, and some look for right edges.)

Figure 6 shows the weights to and from a single hidden unit after ALVINN was trained on roads of a fixed width. White squares represent positive values; black squares represent negative values. This hidden unit acts as a filter for two types of roads, one slightly to the left of center and one slightly to the right.

The weights from the video camera retina, along with the explanatory schematic, show the positions and orientations of the two road types that activate the hidden unit. Notice that the road specifications overlap: The large white region in the center of the weight diagram is a merger of the weights for the left edge of the rightmost road with the weights for the right edge of the leftmost road.

This hidden unit is also excited by obstacles in the periphery of the image and inhibited by obstacles in the center of the image where it expects the road to be. By fusing data from the video-camera and range finder sensors, hidden units can determine the position and orientation of the road more accurately than they could with either sensor alone.

This hidden unit makes excitatory connections to two sets of output units, dictating a slight left or right turn. Since it provides support for two turn directions, it must work with other hidden units to pin down the correct steering direction. Double-duty hidden units like this provide a compact representation. They allow a network with a small hidden layer to perform a complex task, like following a road, accurately.

Reducing the size of the hidden layer not only increases the rate at which a computer can simulate the network, it can also improve the network's performance. With too many hidden units, a network can simply memorize the correct response to each pattern in its training set instead of learning a general solution.

By limiting the size of the hidden layers, the network is forced to develop appropriate feature detectors to efficiently classify large sets of input patterns. These general-purpose feature detectors are more likely to be relevant to novel inputs, so the network performs better. In one experiment, we drove the NAVLAB vehicle using a network trained with only nine hidden units without any significant loss in driving proficiency.

Hidden units that act as filters for one to three roads are the most common result when ALVINN is trained on roads of a fixed width. The network develops a different representation when trained on images with varying road widths. Instead of developing into detectors for entire roads, the hidden units learn to look for a single road edge at a particular position and orientation.

The units support a wide range of travel directions. The correct travel direction for a road with an edge at a particular location varies substantially depending on the road's width. The hidden units cooperate with each other to determine the correct travel direction in any situation.

It's important to understand that no single hidden unit can perform the task alone; the collective activity of all the hidden units determines how the network behaves. Through this kind of cooperation, the network can use relatively coarse feature detectors and still maintain performance accuracy.

Hidden Units Demystified
It's easy to uncover what's in the hidden layers when you apply a neural network to a geometrical problem, as illustrated by the two-spirals and road-tracking examples. The visualization tools made practical by microcomputers and personal workstations have proved invaluable for this type of analysis.

Some researchers display only a hidden unit's weights when trying to analyze a network. The work of Lang and Witbrock (see reference 1) shows that, for geometric problems, it can be more helpful to display the unit's response to a systematic sampling of points in the input region, especially when the network has more than one hidden layer.

This practice is also common in classical neuroscience investigations of the visual system. You can't measure the weights between living neurons in the cortex of the brain, but you can measure their response to various inputs. Many studies of the visual system have been done by graphing the firing rate of cortical neurons while varying a stimulus pattern presented to the retina.

In the case of ALVINN, we saw from the weights that the network learns to efficiently exploit regularities in the input by making its hidden-layer units sensitive to a range of road types. We also tried plotting the units' response patterns while varying the retinal input (presenting roads at various positions and orientations); this confirmed our interpretation of what the hidden layer was doing.

Training ALVINN is time-consuming

Figure 5: *The architecture of ALVINN (autonomous land vehicle in a neural network).*

Figure 6: *Pattern of weights projecting to and from a single ALVINN hidden unit after training on roads with a fixed width. This hidden unit acts as a filter for two types of road, one slightly to the left of center and one slightly to the right. The explanatory schematics on the right side of the figure highlight our interpretation of these weights.*

and requires serious computing power, but you can implement the resulting network on a personal computer or workstation. We see this as a developing trend in neural computing: Training for real-world applications will be expensive, but delivery will be cheap. Analysis of networks through visualization is also easily done on personal workstations.

While we have removed some of the mystery concerning the representations that neural networks develop, the hidden layers have yet to give up all their secrets. One question still to be answered is how ALVINN accomplishes "sensor fusion," combining inputs from its video-camera and range finder retinas to arrive at the best steering direction. Experiments are under way to answer this. ∎

REFERENCES
1. Lang, K. J., and M. J. Witbrock. "Learning to Tell Two Spirals Apart." In *Proceedings of the 1988 Connectionist Models Summer School*, D. S. Touretzky,

G. E. Hinton, and T. J. Sejnowski, eds. San Mateo, CA: Morgan Kaufmann Publishers, 1988.
2. Thorpe, C., M. Herbert, T. Kanade, S. Shafer, and the members of the Strategic Computing Vision Lab. "Vision and Navigation for the Carnegie Mellon NAVLAB." *Annual Review of Computer Science Vol. II*, Joseph Traub, ed. Palo Alto, CA: Annual Reviews, Inc., 1987.
3. Pomerleau, D. A. "ALVINN: An Autonomous Land Vehicle in a Neural Network." In *Advances in Neural Information Processing Systems 1*, D. S. Touretzky, ed. San Mateo, CA: Morgan Kaufmann Publishers, 1989.

David S. Touretzky is a research computer scientist at Carnegie Mellon University in Pittsburgh, Pennsylvania. He has a Ph.D. in computer science from Carnegie Mellon. Dean A. Pomerleau is a doctoral student in computer science at Carnegie Mellon. They can be reached on BIX c/o "editors."

Building Blocks for Speech

*Modular neural networks are a new approach
to high-performance speech recognition*

Alex Waibel and John Hampshire

Some speech-recognition abilities that we take for granted—understanding a conversation involving several different speakers over lots of extraneous noise, for instance—are still beyond the reach of even the most powerful supercomputer. This may seem strange, since the human brain can't hope to match the arithmetical performance of a pocket calculator, but it does indicate the complexity of automatic speech recognition. Modular neural networks, however, might hold the key to achieving rapid and more-reliable machine-based speech recognition.

We recognize speech by applying an enormous body of knowledge to rapidly interpret the audio signals from the world around us. This knowledge ranges from low-level acoustic features to high-level facts about the world and the speaker's intent. These features and facts are heavily interrelated. No piece of the speech-recognition puzzle can be considered by itself, nor can pieces be evaluated sequentially. Rather, each provides a constraint that, together with many other facts and constraints, forms a total picture.

Neural Nets in Speech Recognition

The limited ability of current computer models to absorb and apply a large body of facts restricts efforts to achieve automatic recognition of human speech. Effective models must determine, maintain, and program all necessary facts and rules of speech into a system. They must then integrate the massive number of interrelationships between these facts and rules to rapidly interpret the spoken word. If speech-recognition systems could learn important speech knowledge automatically and represent this knowledge in a parallel distributed fashion for rapid evaluation, they would then be able to overcome the deficiencies of current systems. Such a system would mimic the functions of the human brain, which consists of several billion simple, inaccurate, and slow processors that perform reliable speech recognition.

The development of parallel distributed processing (PDP) or neural-network models and the development of automatic learning algorithms (see reference 1) are two very important steps in the development of reliable speech-recognition systems. You can implement algorithms that simulate PDP learning models on anything from a microcomputer to a supercomputer (see reference 2). These algorithms are even available commercially.

Two major problems have to be addressed, however, before neural-network models become useful for speech recognition: time and scaling.

Figure 1: *The left side (in red) shows the time-delay feature of the network. Three 10-millisecond input slices are combined to create activations in the first hidden layer (a). Activations in the second hidden layer (b) are created by combining five slices from the first hidden layer. The right side (in blue) shows the connections from the input layer to node 4 of the first hidden layer. When an input (c) matches the pattern of the connections (d), the node is activated strongly (e).*

Speech and Time

Speech is a dynamic signal, and a speech-recognition system must be able to classify sounds without knowing when a particular sound will occur. It must also be able to capture the time-varying properties—the *signature*—of speech in feature space rather than simply taking static "snapshots" of the signal. These requirements are addressed by the Time-Delay Neural Network (TDNN) (see reference 3).

Rather than trying to decide whether a particular sound is, for example, a letter *b* (the speech signal may not contain useful information at certain points in time), the TDNN scans the input for clues that provide the evidence it needs to construct an overall recognition decision. Using this method, the TDNN has demonstrated performance superior to that of other speech-recognition models in small but difficult recognition tasks.

The TDNN shown in figure 1a is designed to discriminate the voice-stop consonants *b*, *d*, and *g* as they occur in a large database of isolated spoken words. At the output, three units represent each of the three *phoneme* categories. (Phonemes are the unique sounds of a spoken language; they form the acoustic-phonetic building blocks of speech.) The input layer of the network consists of 15 time slices of speech. Each one of these time slices is a frequency spectrum representing 10 milliseconds of the speech waveform—a 10-ms voiceprint of the speaker. Each spectrum, in turn, consists of 16 coefficients representing frequencies ranging from the lower limit of hearing (about 20 Hz) to over 5 kHz.

In many neural networks, each node in a given layer is connected to all the nodes in the next layer. This is not the case, however, for the TDNN. The reasons for this are related to the temporal complexity of human speech.

Windows to the Spoken Word

Rapid changes in human speech occur over several tens of milliseconds. Therefore, a 30-ms "window" of speech (or an overlapping series of such windows) can capture the local acoustic-phonetic events that act as identifying features of a particular phoneme. The TDNN groups three 10-ms time slices from the input layer into a 30-ms window. Each coefficient in this window connects to eight nodes in the first hidden layer of the TDNN. Each of these nodes forms a condensed feature representing important cues that the network looks for in the input. The network shifts the window one time slice at a time across the input (a

range of 150 ms of speech), creating 13 distinct firings at the eight nodes of the first hidden layer.

The grouping scheme in the first hidden layer and its connections to the second hidden layer are analogous to the input layer's groupings and connections to the first hidden layer. The firing patterns of the eight nodes in the first hidden layer over a five-time-slice window form the

*The TDNN
has learned—without
any supervision—
the importance of
rising and falling
formant transitions in
discriminating between
similar sounds.*

input to each of three nodes in the second hidden layer. As this window sweeps over the activation patterns in hidden layer 1, it generates activations at the three nodes in hidden layer 2. These form preliminary votes for one of the output's three phoneme categories.

Because their weights are fixed across time shifts, the connections between the layers allow the network to find key features of the speech waveform despite the fact that these features may be spread across time or shifted along the time axis. Figure 1a illustrates the activation of a TDNN when given the voiced consonant *d* in the syllable *do*. In this figure, negative node activations in the input layer are gray, and positive node activations throughout the network are black. The degree of node activation is proportional to the size of the rectangle depicting a given node.

In figure 1c, connections from the input-layer window to node 4 of the first hidden-layer time slice are shown to the side of the TDNN. (Unlike activations, positive connections are white and negative connections are black; the background is gray.) The activation level of node 4 in the first hidden layer at a given time slice is obtained by taking the activation of each of the 48 nodes in the input

layer window, multiplying this node activation by the strength of its connection to node 4, and adding up these 48 products. This sum forms the input to node 4, which uses a thresholding (or "squashing") function to produce the output activation shown.

Note that the connections from the input layer to node 4 of the first hidden layer are positive for midrange frequencies in the input that rise or fall over time. The positive (white) connections that slope downward over time provide a strong input stimulus to node 4 when they detect a downward-sloping spectrum over time in the input layer. The arrow in figure 3 marks the onset of the *ü* sound in *do*. Beginning at this point, the nodes in the input layer corresponding to frequencies from 800 Hz to 1600 Hz show the downward-sloping activation pattern over time indicative of a *falling formant*. (A formant is a quality of sound representative of vowels.) This results in a strong firing of node 4 in the first hidden layer.

Falling midrange frequencies are characteristic of the utterance *do* shown in figure 1c. There is a great deal of experimental evidence showing that humans rely heavily on the perception of this acoustic event (a *formant transition*) for accurate speech recognition. The positive connections in the figure that slope upward over time detect rising formant transitions, which are also vital to understanding human speech. Clearly, the TDNN has learned—without any explicit supervision—the importance of both rising and falling formant transitions for accurate discrimination of the *b*, *d*, and *g* phonemes.

Because the TDNN scans across the input speech signal, it is relatively insensitive to the timing of vowel onset for the voice stops *b*, *d*, and *g*. A version of the same utterance shown in figure 1c shifted forward in time results in the same strong output activation indicating the detection of the *d* phoneme. The advance of vowel onset merely causes the hidden units to fire earlier, in synchrony with events in the input. The combined accumulated evidence from these firings still allows the network to recognize the utterance as a *d*, as opposed to a *b* or a *g*.

The TDNN has been experimentally evaluated on a number of small phonemic discrimination tasks and has achieved excellent recognition performance. The voiced consonants *b*, *d*, and *g*, for example, can be detected in more than 98 percent of the trials with a TDNN trained on data from a single

continued

speaker and tested on different data obtained from the same speaker.

Modular Training

The second problem for practical neural-network-based systems is scaling. Since neural networks depend on computationally intensive learning algorithms and simulations of large parallel networks, they are difficult to extend to large systems and to run on commonly accessible computing facilities such as microcomputers and personal workstations. It is extremely important, therefore, that the construction of large systems take place incrementally, without requiring repetitive retraining of ever-larger structures every time the task size increases.

In examining problems of scale, it is important to note that neural networks are made up of extremely simple computing elements that can be simulated easily in real time on most personal computers and workstations. Moreover, since such a system is completely specified by its connections and its weights, it is easily portable and can run on any machine. In our own implementation, a simple generic program has to simply load a set of weights and a wiring table to run an entirely different system.

A much more serious computational limitation of neural-network-based systems arises during training. Here, you must execute many recognition passes over many training patterns to gradually modify the network's weights and achieve a satisfactory output response. Depending on the network's size and the number of training tokens, you might have to devote significant computational resources to training.

This is acceptable in many cases, since learning can frequently be done off-line over several days and weeks while recognition must be fast and efficient. Nevertheless, cost-effective system development requires that you be able to create flexible and effective designs using commonly available resources. Luckily, you can create such designs using modular-network training techniques without incurring a performance penalty.

To recognize continuous speech from any speaker, you need neural networks that are orders of magnitude more complex than the single three-output TDNN. Training such a network from scratch to recognize all phonemes in continuous speech is a daunting task.

An alternative method would exploit the knowledge developed by smaller, independently trained networks by incorporating these smaller network modules into large superstructures. This modular approach could not only reduce training time but also lead to a more incremental and distributed design approach to the construction of large-scale, efficient connectionist systems.

Figure 2 shows one promising approach to such a modular design. Here, two networks are trained independently: one to perform b, d, and g discrimination, and the other to discriminate p, t, and k. The features learned by the hidden units in layer 1 of each network are the useful representations of speech that you want to retain for the training of a larger b-d-g/p-t-k supernetwork.

Combining the two networks, you get a modular network that consists of a common input layer linked to task-specific first hidden layers. The connection strengths are obtained from the individual b-d-g and p-t-k networks. These first hidden-layer modules are then linked to a common second hidden layer, and from there to the output layer.

The training for this network occurs in two phases. First, the two TDNNs are trained individually for best performance on their subtasks. Then, the higher-layer connections of the modules are trained collectively to integrate the individual modules into the combined

Figure 2: *To form a modular network, you combine independently trained first hidden layers with a common second hidden layer. For greater accuracy, you can also provide free units (shown in red) for the combined first hidden layer.*

continued

Figure 3: *Modular construction of an all-consonant network.*

network. This modular network achieves a single-speaker recognition rate of 98.1 percent for the *b-d-g/p-t-k* task.

To enhance performance, you can augment the modular architecture with four additional "free" units, which we call *connectionist glue* (see figure 2). These units glue together two previously disjointed networks. During the second phase of training, these glue units supplement the pretrained *b-d-g* and *p-t-k* units in hidden layer 1 by extracting any additional useful features from the input. With this network, we have achieved a single-speaker *b-d-g/p-t-k* recognition rate of 98.6 percent. Thus, this network performs as well as or better than the in-dependently trained networks do separately. It also performs at least as well as a monolithic network that was trained for the task from scratch.

Expanding the Network

Figure 3 shows a larger modular network built to recognize all the consonants of a single speaker. Like the one in figure 2, this network has individual TDNNs trained in subcategories of the consonant set. These modules are then linked to a module that recognizes the subcategory of the input. The extra features in this module tend to assimilate the features developed in the consonant modules.

Training this network occurs in three phases. As before, the first phase involves training each module for best performance in its subtasks. In the second phase, the connections from the second hidden layer to the output are trained collectively. In the final phase, all connections of the network are trained in unison for a few learning iterations to fine-tune the global network's recognition performance. This sort of training is significantly faster than training a huge network from scratch. Working in parallel, the modules form a distributed representation of the speech signal that can achieve a recognition rate of 96 percent for the consonants of a single speaker.

continued

Accelerated Learning

Coincident with the development of modular design techniques, recent advances in neural-network learning strategies and hardware and software implementations have led to dramatic improvements in

W*hat*

seemed impossible a short time ago will soon be done on a personal computer.

---

network processing speeds. Learning speeds are also accelerating. These can be increased by improving the metrics that a network uses to measure how well it classifies training data.

Speed can also be increased significantly by improving the numerical search techniques that form the basis of network learning. Research in this area has resulted in learning procedures that converge to near-optimal results much more rapidly than before (see references 4, 5, and 6). Indeed, improvements in learning algorithms have brought the training time for a typical TDNN task down from three days of run time on a supercomputer to 8 minutes of CPU time

on a high-end engineering workstation.

High-speed computing capabilities for neural-network training are becoming more accessible to personal computer and workstation users. Several manufacturers now offer plug-in floating-point accelerator boards for microcomputers that yield speeds of more than one million floating-point operations per second, while workstation manufacturers are producing desktop machines that rival super-minicomputers produced just a few years ago. Massively parallel connectionist hardware designs are also under development in various laboratories (see reference 7).

Speech recognition using modular neural networks is progressing rapidly. What seemed impossible a short time ago will soon be done on a personal computer. Advances in system-design techniques, learning software, and underlying hardware are creating the computing power required for very-large-scale neural-network tasks. All these advances bring connectionist design for speech and signal interpretation within reach of commonly available and affordable technology. ■

REFERENCES

1. Rumelhart, D., J. McClelland, and the PDP Research Group. *Parallel Distributed Processing*, vols. 1 and 2. Cambridge, MA: MIT Press, 1987.
2. McClelland, J., and D. Rumelhart. *Explorations in Parallel Distributed Processing*. Cambridge, MA: MIT Press, 1988.
3. Waibel, A., T. Hanazawa, G. Hinton, K. Shikano, and K. Lang. "Phoneme Recognition Using Time-Delay Neural Networks." *IEEE Transactions on Acoustics Speech and Signal Processing*, vol. ASSP-37. March 1989.
4. Fahlman, S. "An Empirical Study of Learning Speed in Back-Propagation Networks." *CMU Technical Report* CMU-CS-88-162, June 1988.
5. Haffner, P., A. Waibel, H. Sawai, and K. Shikano. "Fast Back-Propagation Learning Methods for Neural Networks in Speech." *ATR Interpreting Telephony Research Laboratories Technical Report*, TR-I-0058, 1988.
6. Watrous, R. "Speech Recognition Using Connectionist Networks." TR-MS-CIS-88-96, Dept. of Computer and Information Science, University of Pennsylvania, 1988.
7. Alspector, J. "Research results in VLSI implementations of neural networks." Speech presented at the 1988 Conference of the Acoustical Society of Japan.

ACKNOWLEDGMENT

Alex Waibel holds a Ph.D. in computer science from Carnegie Mellon University, where he is a research computer scientist. John Hampshire is a Ph.D. candidate at Carnegie Mellon in electrical and computer engineering. The authors can be reached on BIX c/o "editors."

Neural Networks: Theory and Practice

For most of their existence, neural networks and neural-network simulations have been solely objects of university-based research. In the last few years, however, researchers and others have founded companies dedicated to producing commercial products based on neural-network technology. To reflect both the academic and commercial aspects of the technology, this resource guide consists of two parts. The In Theory section lists books and articles you can read to learn more about neural networks. The In Practice section lists some of the available neural-network hardware and software products, listed alphabetically by company name.

IN THEORY

Anderson, J. A., M. T. Gately, P. A. Penz, and D. R. Collins. "Radar Signal Categorization Using a Neural Network." *Neural Networks*, 1 Supp. 1, 1988.

Anderson, J. A., and E. Rosenfeld, eds. *Neurocomputing: Foundations of Research.* Cambridge, MA: MIT Press, 1988.

Anderson, J. A., J. W. Silverstein, S. A. Ritz, and R. S. Jones. "Distinctive Features, Categorical Perception, and Probability Learning: Some Applications of a Neural Model." *Psychological Review*, vol. 84, 1977.

Brugge, J. F., and R. A. Reale. "Auditory Cortex." In A. Peters and E. G. Jones, *Cerebral Cortex*. Vol. 4, *Association and Auditory Cortices*. New York, NY: Plenum Press, 1985.

Davis, P. J., and J. A. Anderson. "Nonanalytic Aspects of Mathematics and Their Implication for Research and Education." *SIAM Review*, vol. 21, 1979.

Hadamard, J. *Psychology of Invention in the Mathematical Field.* New York, NY: Dover. 1945.

Hinton, G. E., and J. A. Anderson, eds. *Parallel Models of Associative Memory*, rev. ed. Hillsdale, NJ: Lawrence Erlbaum Associates, 1989.

Jones, W. P., and J. Hoskins. "Back-Propagation." BYTE, October 1987.

Josin, G. "Neural Network Heuristics." BYTE, October 1987.

Knapp, A., and J. A. Anderson. "A Theory of Categorization Based on Distributed Memory Storage." *Journal of Experimental Psychology: Learning, Memory and Cognition*, vol. 9, 1984.

Knudsen, E. I., S. du Lac, and S. D. Esterly. "Computational Maps in the Brain." *Annual Review of Neuroscience, Volume 10*. Palo Alto, CA: Annual Reviews, 1987.

Kohonen, T. *Associative Memory*. Berlin: Springer-Verlag, 1977.

Kohonen, T. *Self-Organization and Associative Memory*. Berlin: Springer-Verlag, 1987.

Kosko, B. "Constructing an Associative Memory." BYTE, September 1987.

Lang, K. J., and M. J. Witbrock. "Learning to Tell Two Spirals Apart." In *Proceedings of the 1988 Connectionist Models Summer School*, D. S. Touretzky, G. E. Hinton, and T. J. Sejnowski, eds. San Mateo, CA: Morgan Kaufmann Publishers, 1988.

Lippmann, R. "An Introduction to Computing with Neural Nets." *IEEE ASSP Magazine*, vol. 4, April 1987.

McClelland, J. L., and D. E. Rumelhart, eds. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, vols. 1 and 2. Cambridge, MA: MIT Press, 1986.

McClelland, J. L., and D. E. Rumelhart. *Explorations in Parallel Distributed Processing: A Handbook of Models, Programs and Exercises.* Cambridge, MA: MIT Press, 1988.

McCulloch, W. S. *Embodiments of Mind.* Cambridge, MA: MIT Press, 1988.

Mead, C. *Analog VLSI and Neural Systems.* Reading, MA: Addison-Wesley, 1988.

Minsky, M., and S. Papert. *Perceptrons.* Cambridge, MA: MIT Press, 1969 and 1988.

O'Shaughnessy, D. *Speech Communications.* Reading, MA: Addison-Wesley, 1986.

Pomerleau, D. A. "ALVINN: An Autonomous Land Vehicle in a Neural Network." In *Advances in Neural Information Processing Systems 1*, D. S. Touretzky, ed. San Mateo, CA: Morgan Kaufmann Publishers, 1989.

Rabiner, L. R., and R. W. Schafer. *Digital Processing of Speech Signals.* Englewood Cliffs, NJ: Prentice-Hall, 1978.

Thorpe, C., M. Herbert, T. Kanade, S. Shafer, and the members of the Strategic Computing Vision Lab. "Vision and Navigation for the Carnegie-Mellon NAVLAB." *Annual Review of Computer Science, Volume II*, Joseph Traub, ed. Palo Alto, CA: Annual Reviews, Inc., 1987.

Viscuso, S. R., J. A. Anderson, and K. T. Spoehr. "Representing Simple Arithmetic in Neural Networks." In *Advanced Cognitive Science: Theory and Applications*, G. Tiberghien, ed. London: Horwoods, 1989.

Waibel, A. "Modular Construction of Time Delay Neural Networks for Speech Recognition." *Neural Computation*, vol. 1, March 1989.

Waibel, A., H. Sawai, and K. Shikano, "Modularity and Scaling for Phonemic Neural Networks." In *IEEE Transactions on Acoustics Speech and Signal Processing* (in press).

Wasserman, P. *Neural Computing, Theory and Practice.* New York, NY: Van Nostrand Reinhold, 1989.

IN PRACTICE

The Brain Simulator $99
Runs under MS-DOS
Tutorial software for neural circuit design
Abbot, Foster & Hauserman
44 Montgomery, Fifth Floor
San Francisco, CA 94014
(800) 562-0025
(415) 955-2711
Inquiry 1181.

N-NET
MS-DOS version $895
VAX/VMS version .... starting at $2995
Integrated neural-network development system; uses functional link net architecture
AI Ware, Inc.
11000 Cedar Ave., Suite 212
Cleveland, OH 44106
(216) 421-2380
Inquiry 1182.

BrainMaker $99.95
Runs under MS-DOS
Neural-network simulation software; supports five types of nodes and can process up 500,000 connections per second
California Scientific Software
160 East Montecito, Suite E
Sierra Madre, CA 91204
(818) 355-1094
Inquiry 1183.

Cognitron
MS-DOS Windows and Mac
versions $600
INMOS transputer version $1800
Neural-network/parallel-processing prototyping and delivery system
Cognitive Software, Inc.
703 East 30th St.
Indianapolis, IN 46205
(317) 924-9988
Inquiry 1184.

Connections $87.95
Runs under MS-DOS
A traveling-salesman demo modeled after Hopfield networks

NetWurkz $79.95
Runs under MS-DOS
Elementary introduction to neural
networks
DAIR Computer Systems
3440 Kenneth Dr.
Palo Alto, CA 94303
(415) 494-7081
Inquiry 1185.

Savvy Text Retrieval System
Savvy Signal Recognition System
Savvy Vision Recognition System
(Call for pricing)
Run under VAX/VMS, MS-DOS,
and Unix
Libraries of C subroutines that use
neural technology to solve real-world
problems
Excalibur Technologies
2300 Buena Vista SE
Albuquerque, NM 87106
(505) 764-0081
Inquiry 1186.

ANZA ....from $7000
AT-compatible neural-network
coprocessors; includes software and
programming interface
ANZA Plus ....from $12,500
AT-compatible neural-network
coprocessors
ANZA Plus/VME $24,950
Neural-network coprocessor for Sun
workstations
AXON $1950
A neural-network description language
Neural Network Development
Toolkit $3950
For ANZA Plus and ANZA Plus/VME
systems
Ports C programs into ANZA Plus and
ANZA Plus/VME formats; includes
AXON
ExploreNet
  MS-DOS version $995
  Sun version.........................$3950
Stand-alone neural-network software
HNC, Inc.
5501 Oberlin Dr.
San Diego, CA 92121
(619) 546-8877
Inquiry 1187.

MD/210 Fuzzy Set
Comparator...............................$38
Hardware implementation of Hopfield
neurons
Micro Devices
5695B Beggs Rd.
Orlando, FL 32810
(407) 299-0211
Inquiry 1188.

N1000
Neural-network development tools for
signal and image processing applications
  Including 80386 computer
  ....from $19,000
  Support package only (for Sun-4,
  Sun-3, IBM PC AT, and PS/2
  Model 50 ....from $7955

N500 ....from $495
Runs on IBM PC AT and PS/2 Model 50
Single-unit RCE network software
Nestor, Inc.
1 Richmond Sq.
Providence, RI 02906
(401) 331-9640
Inquiry 1189.

Awareness $275
Runs under MS-DOS
Introduction to four types of neural-
network paradigms
Genesis $1095
Runs under MS-DOS
Neural-network development
environment
Neural Systems
2827 West 43rd Ave.
Vancouver, BC Canada V6N 3H9
(604) 263-3667
Inquiry 1190.

NeuralWorks Explorer $299
Runs under MS-DOS
An introduction and tutorial on neural
networks
NeuralWorks Professional II
  MS-DOS and Macintosh
  versions................................$1495
  Sun-3, Sun-4, and Sun386i
  versions................................$2995
  NeXT and INMOS transputer
  versions................... call for pricing
Neural-network development system
NeuralWorks Designer Pack $1995
MS-DOS and Sun versions
Links Professional II networks with C
programs
NeuralWare, Inc.
103 Buckskin Court
Sewickley, PA 15143
(412) 741-5959
Inquiry 1191.

MacBrain $995
Runs on Macintosh
Lets you prototype and deliver neural-
network applications
HyperBrain
(Comes with MacBrain)
Toolkit allows you to build neural-
network applications within HyperCard
Neurix, Inc.
1 Kendall Sq., Suite 2200
Cambridge, MA 02139
(617) 577-1202
Inquiry 1192.

Owl I, II, III.....................from $349
Libraries of modules for IBM and
compatibles that lets you define and
access 10 different neural networks
Extension Pack........................$149
Three additonal networks
Olmsted & Watkins
2411 East Valley Pkwy., Suite 294
Escondido, CA 92025
(619) 746-2765
Inquiry 1193.

Intelligent Pattern Recognition
Chips.....................................$500
Stores a 1000-by-64 matrix of weights
and multiplies it with an input vector
Oxford Computer
39 Old Good Hill Rd.
Oxford, CT 06483
(203) 881-0891
Inquiry 1194.

ANSim 2.1 ....$495
Runs under MS-DOS
13 neural-network models
ANSkit....................................$950
Runs under MS-DOS
Neural-network development system
ANSpec $2995
Runs under MS-DOS
Neural-network specification language
Delta Floating Point
Processor...........................$24,950
Runs on IBM PC, AT, PS/2s, and
Sun386i
Neural-network accelerator boards
Sigma Neurocomputer
Workstations ....from $31,500
80386-based systems with Delta
Processor, ANSkit, Delta C, Delta
Macro, and ANSpec
SAIC
10260 Campus Point Dr.
Mail Stop 71
San Diego, CA 92121
(619) 546-6290
Inquiry 1195.

DENDROS-1............................$35
Neural-network chip that produces the
dot product of the inputs and the
connection weights of 22 synapses
DENDROS-1 Evaluation Board ... $695
Uses eight DENDROS-1 chips to create
a hardware-based neural network
Syntonics Systems, Inc.
20790 Northwest Quail Hollow Dr.
Portland, OR 97229
(503) 293-8167
Inquiry 1196.

TRW Mark V Neural Processor
Write for pricing information
Runs on VAX/VMS
MC68020-based parallel-processing
system includes tools for neural-network
applications
TRW
Military Electronics & Avionics Div.
One Rancho Carmel
San Diego, CA 92128
(619) 592-3482
Inquiry 1197.

NeuroShell $195
Runs under MS-DOS
Creates neural-network applications
using a modified back-propagation
Ward Systems Group, Inc.
228 West Patrick St.
Frederick, MD 21701
(301) 662-7950
Inquiry 1198.

DEALING WITH A DIGITAL WORLD

Digital signal processors move to micros, where they handle complex data like sound and images with speed and flexibility

David A. Mindell

L ast fall, Apple founder Steve Jobs introduced the NeXT machine, hailed as the first of a new generation of personal computers. It has a 17-inch "megapixel" display, a 256-megabyte magneto-optical disk drive, an Ethernet port, a SCSI port, and the multitasking Mach operating system, derived from Unix. While these features typify the growing sophistication of the personal computer marketplace, one chip in the NeXT "cube" represents a truly new direction in personal computer systems: a digital signal processor (DSP).

DSPs have been around for at least 10 years, mostly in high-end military or industrial applications and research. Recent developments in technology, however, have made single-chip DSPs widely available and affordable, as well as easily integrated into larger systems. Consequently, digital signal processing is expanding into numerous new applications, particularly on personal computers, and is one of the fastest-growing areas in digital technology today.

Telecommunications is a primary arena for digital signal processing, and many of today's high-speed modems use digital-signal-processing techniques to reduce errors and increase transmission rates. Specialized "adaptive" algorithms sense the noise and bandwidth properties of the telephone line and adjust transmission and filtering parameters accordingly. DSPs also encrypt and compress data for reasons of security and efficiency. Voice-mail systems, for example, process speech signals and convert them to compressed ASCII files for transmission over standard E-mail systems.

A single DSP chip can accomplish virtually all telecommunications functions within a system, eliminating the need for expensive additional hardware. The chip in the NeXT machine, for example, can act as a modem, a fax machine, a voice digitizer, and a data compressor while also serving as a high-throughput numeric processor.

The advent of digital audio has created a niche for DSPs in all areas of music processing, including synthesis, recording, mixing, equalization, and editing. "Direct to disk" has become very popular, where compact-disk-quality audio is recorded through a DSP in a personal computer onto a high-capacity hard disk drive. DSP systems can also clean up poor recordings, removing noise, and even replacing clicks and pops with interpolated artificial music. Some can perform time compression and expansion of speech and music signals without affecting the pitch. A 40-second message, for example, could fit into a 30-second commercial, all without "munchkinizing" or "Frankensteining." Despite its popularity, however, digital signal processing is still largely misunderstood by even the computer literati.

Digital Signals

The term *digital signal processing* refers to the digital implementation of filters and algorithms to process some kind of data or *signal*. These techniques, while more complex than their analog counterparts, provide all the advantages associated with numerical processing: speed, accuracy, increased noise immunity, greater dynamic range, flexibility and programmability, and the power to create sophisticated pseudo-intelligent systems.

Analog signals are *continuous-time* representations of a given quantity (see figure 1a). A microphone, for example, produces a varying voltage that is proportional to the sound it detects. The first problem of a digital system, then, is to convert these analog signals into numbers—that is, to digitize them. To do this, the system must *sample* at regular intervals to convert the continuous-time signal into a *discrete-time* representation (see figure 1b).

The time between these samples, the *sampling period*, is determined by Nyquist's sampling theorem, which states that samples must be taken at twice the highest frequency contained in the data. Audio signals, for example, have a 20-kHz bandwidth, that being the upper limit of human hearing. Digital audio, then, must take at least 40,000 samples per second to accurately reproduce sound (CDs actually operate at 44.1 kHz

Figure 1: *In digitization, (a) an analog signal is sampled and converted to (b) a sequence of digital data. The data can then be passed through a digital signal processor for processing.*

Figure 2: *In this diagram, a filter (a four-sample averager) adds an input x(n) to the three preceding inputs (which have been delayed for this purpose) and divides the total by 4 to give the output y(n). The effect of this averaging filter is to smooth out rapid deviations in the input signal, while leaving slower deviations, or low frequencies, relatively unaffected.*

to provide error correction and reduce noise).

An A/D converter converts a sampled analog voltage into a binary number. A DSP takes a string of numbers from an ADC, processes them in some way, and produces another string of numbers, which can then be passed through a D/A converter to reconstruct an analog signal. The digital data can also be passed directly to a computer for further processing or storage.

DSP Microprocessors

Until recently, most digital filters were hard-wired. A hardware multiplier was connected to an accumulator, which was connected to another multiplier, and so forth. Early microprocessors were simply not fast enough to perform the operations required for sophisticated filters. Advances in VLSI, however,

have produced specialized DSP microprocessors with enough power to implement digital filters in software.

General-purpose microprocessors are bulky things. Loaded with features for memory management, system control, and compiler design, they can be clumsy to operate in real time (RISC architectures are an attempt to avoid this problem). DSP chips are similar to other microprocessors in that they execute programs, grab instructions and data from memory, and perform calculations. They are stripped down, however, and optimized for simple, repetitive operations with very high rates of data flow.

The distinguishing feature of DSP chips is their emphasis on the multiply-accumulate (MAC) operation, which is central to digital filtering. Current DSP chips—for example, the Texas Instruments TMS320C30, the Motorola DSP56001, and the AT&T DSP16A—can perform MAC operations in a single clock cycle in 60 to 80 nanoseconds, a value approximately equivalent to 25 to 33 million floating-point operations per second. (See the text box "A Look at DSP Chips" by John E. Hart on page 250 and table 1.) Impressive even by today's standards, such processing rates extend DSP chips' range well into the high-fidelity audio domain and just to the edge of video. Other features of DSP processors include extensive parallelism and pipelining, independent memories, and "bit-reversed" addressing modes for Fourier-transform data.

Another difference between general-purpose microprocessors and DSPs is that DSPs employ the Harvard architecture. In this scheme, data and instructions are kept in separate memories to allow the processor to perform several operations in parallel. There are numerous variations on this structure; some even allow access to five or six data banks simultaneously. The Motorola DSP56001, for example, has two data memories, denoted X and Y. For image processing, then, X and Y data can be kept separate, or for a complex Fourier transform, the X and Y memories can be used for real and imaginary data. A filtering operation might use the Y memory for the data stream and the X memory to hold the filter coefficients.

Most DSPs have some data and code memories on-chip and can access more memory through external buses. The on-board memories are small (rarely more than a few K bytes), but they are usually sufficient because DSP operations, while complex, produce relatively short programs. Some processors even include lookup tables for constants such as sine coefficients as part of on-chip ROM.

Because DSP chips are optimized for data throughput, there are usually several ways of presenting data to the CPU. Digital signals can enter through external buses, direct memory access, or one of several types of serial ports. These data paths, combined with flexible control features, also allow for several DSPs to be strung together to perform parallel operations.

The problem with DSP chips has been that, because they are so streamlined and were made with data flow and not systems in mind, they can be very difficult to program. Without convenient registers and instructions, compilers do not generate efficient code. Therefore, because high-bandwidth DSP applications often run under significant real-time constraints, critical routines must be written by hand and fine-tuned in assembly language. But the special architectures and extensive parallelism of most DSP chips mean that the assembly languages are obscure and esoteric and thus difficult to code.

Industry consensus is that the proliferation of DSPs has been slowed by these and other development difficulties. The situation is changing, however, as chip companies offer design aids such as emulators, library routines, and software simulators.

When you pick a printer for the office, the only thing you have to fear is being wrong.

An office printer can cost a few hundred dollars, or it can cost as much as a small car. And depending on how well it does (or doesn't do) what everybody in the office wants it to do, it can cost you a lot of sleep.

That's why when AEG Olympia talks about printers, the feature we talk about most is good technology. Which may explain why we're the number one business machine company in our native Germany, where technology is a number one passion.

So before an AEG Olympia dealer tells you about all our dot matrix and daisywheel and laser printers, he'll probably show you how beautifully they all work. We live and breathe solid engineering, sensible design and dependable mechanics for printers, just as we do for our typewriters,

facsimiles, copiers and shredders.

For the name of our nearest dealer, call 1-800-999-6872. Or write AEG Olympia, Box 22, Somerville, NJ 08876-0022. You can't go wrong with good technology. **Where technology counts, business counts on us.**

AEG
OLYMPIA

A Look at DSP Chips

John E. Hart

The AT&T DSP32, introduced in 1985, was the first self-contained single-chip floating-point digital signal processor (DSP). Computer scientists and engineers found the compact architecture, ease of hardware integration, and impressive performance ideally suited to a variety of applications. It rapidly became clear that the DSP had great potential, in both single and multiprocessor configurations, to address a wide range of computational needs.

Although floating-point DSPs are only now beginning to appear in personal computers, competition between the major microprocessor designers has resulted in skyrocketing performance-to-cost ratios. Even though sample prices can be high, full-scale production DSPs are selling for as little as $40 per unit, or about $4 per million floating-point operations per second (MFLOPS). In what follows, I want to look at some of the more common floating-point DSPs, selected from table 1 on page 255.

The AT&T DSP32

The major AT&T DSP chips are the fixed-point DSP16 and DSP16A and the floating-point DSP32 and DSP32C. The DSP32 contains two processing units, the control arithmetic unit and the data arithmetic unit (DAU). The CAU is an integer processor with 16-bit resolution in the 25-MHz model; in the newer DSP32C, a 50-MHz device, it has 24-bit resolution. The DSP32 runs at 12.5 MFLOPS, while the DSP32C reaches a peak speed of 25 MFLOPS.

The CAU contains 21 registers that can be loaded from and stored to memory and that can be manipulated arithmetically in add, subtract, and various Boolean operations. Each operation takes four clock cycles. The first 14 CAU registers serve as address pointers to 32-bit floating-point operands held in either the internal or the external RAM; R_{15} through R_{19} can be used as post-index registers for floating-point memory accesses.

The DAU does floating-point computations exclusively. It has four 40-bit data accumulators (32-bit mantissas with 8-bit exponents), enabling "single-extended-precision" calculations within the DAU itself. All results are truncated or rounded to 32 bits when accumulators are stored to memory.

The generic DAU instruction involves two registers and three operands. It has the form

$$z = A_n = \pm A_m \pm y \times x$$

or

$$z = A_n = \pm y \pm A_m \times x$$

The A_j are the DAU accumulator registers ($j = 0$ to 3), and the variables z, x, and y can refer to specific address pointer registers within the CAU, to accumulator registers, or to an implicit 1 or 0. A DAU instruction statement occupies 32 bits of memory. The impressive floating-point speed of the DSP32 and its compact object codes is, in part, a result of the fact that one instruction can do *two* floating-point operations involving five variables.

The DSP32 is a four-state machine, with each instruction taking just four clock cycles. For example, if x, y, and z all refer to data in memory, using AT&T assembly language syntax, you could replace these variables with their address pointer registers $*R_n$, where $n = 1$ to 14. The instructions are then executed as follows:

Cycle 1: Fetch and decode the instruction.
Cycle 2: Fetch operand pointed to by the x-variable pointer $*R_n$.
Cycle 3: Fetch operand pointed to by the y-variable pointer $*R_i$.
Cycle 4: Do the floating-point operations and write the result to the z variable pointed to by $*R_j$.

There is a 32-bit memory access during each clock cycle. At the 40- to 50-MHz clock rates at which the DSP32s can run, this would require extraordinary memory performance. To make these high bandwidths possible, the memory is partitioned into two banks, Bank0 and Bank1.

The address location of the 4K bytes of internal memory is flexible. All the external memory, if included, is located in Bank0. Memory activity is interleaved between the banks to allow for two-cycle access to each bank. One bank is being addressed while the other is being accessed. Data flow on the *internal* 32-bit bus can proceed at a rate of one long word (or 4 bytes) per cycle, but on the external Bank0 bus, it proceeds at one long word for every two clock cycles. At 50 MHz, this interleaving yields an internal bus bandwidth of 200 megabytes per second.

A look at the four-state cycle sequence indicates that perfect implementation of this interleaving scheme requires a very careful allocation of data and instructions between the two banks. In practice, a useful programming compromise is simply to place data in one bank (usually the lower one, because this bank is expandable off-chip) and to put the instructions in Bank1.

The DSP32 connects to a host microprocessor through its 8-bit parallel interface (16-bit in the DSP32C). This interface features a cycle-stealing direct-memory-access (DMA) controller that allows the host to read or write to DSP memory without having to halt and restart the CAU or DAU processors via software. This ability to change data "on the fly" is central to uses of the DSP32 in interactive scientific teaching and research computing applications.

The serial I/O section of the DSP32 permits input and output of 8-, 16-, and 32-bit data. One use for this port is to connect DSPs together in multiprocessor systems. Another is to drive 16-bit D/A converters, providing a convenient high-speed analog data stream for monitoring computations in real time.

The Motorola 96002

The Motorola 96002's instruction set is a superset of that for the MC56000 (the fixed-point DSP), and the instruction mnemonics are similar to those for Motorola's general-purpose microprocessors. The floating-point chip will be available later this year in both a single-port (the 96001) and a two-port version (the 96002).

The 96002 chip features multiple internal and external buses, with internal memory arranged to support parallel transfers of program and operand data to and from the program controller and the data ALU, respectively.

Running at 26.6 MHz, the MC96002 can attain a peak throughput of 40 MFLOPS and 13.3 million instructions per second (MIPS), although the floating-point throughput will more typically be about 27 MFLOPS in 32-bit single precision or 43-bit single-extended precision. This high performance is a result of internal concurrency and parallelism. The program controller, address-generation unit (AGU), and data ALU operate in parallel. A typical instruction in the 96002 consists of a floating-point operation involving accumulator registers A_0 through A_9 in the ALU as sources and destinations, along with a parallel move.

While the ALU is executing a multiply-accumulate on several ALU accumulator registers, the AGU can be fetching two 32-bit numbers from each of two data memory banks and placing them in other data registers for use in a subsequent instruction. This latter data can be obtained from the internal RAM banks of x-data and y-data or from static RAM attached to the two external memory ports, A and B.

Both transfers use addresses contained in two of the eight pointer registers located in the AGU. The effective pointer addresses can be modified using index registers that are also contained in the AGU. At the same time that the floating-point operations and data transfers are occurring, the program controller prefetches and decodes the next instruction from the program memory. All this can occur in just one instruction cycle (two clock cycles).

The data ALU contains a single-cycle floating-point multiplier/accumulator that works with either 32-bit or 43-bit input data, the latter being made up of 32-bit mantissas and 11-bit exponents. The results are written to ALU registers in "infinite precision."

For example, a single-precision multiply produces a 48-bit mantissa. The result, stored in a 96-bit register, can be used in future register-to-register arithmetic operations without truncation. However, when a result is written from an ALU register to memory, it is automatically rounded down in hardware to single precision.

Double-precision calculations must be done in software, but the bus structure, in which the x-data and y-data can be concatenated, speeds up the transfer of double-precision data to and from the ALU registers. In addition, these 10 96-bit registers provide expanded capability for computing larger expressions than can be performed in the four DSP32 accumulators. To take one example, repeatedly used numeric constants can be permanently stored in some of these registers, avoiding the necessity of collecting them from memory each time they are needed.

The TI TMS320C30

The Texas Instruments TMS320C30 has several features in common with Motorola's 96002 and AT&T's DSP32. The CPU contains a floating-point multiplier and an accumulator, which operate on the eight 40-bit single-extended-precision accumulator registers, as well as on data directly transferred from memory. As in the DSP32, a multiply instruction can get its operands either from data registers (accumulators A_0 through A_7) or from memory locations pointed to by address pointers in the AGU. Like the 96002, the TMS320C30 has multiple internal and external buses.

The 320C30 uses a modified Harvard architecture. This means that there are separate data buses for instructions and data. Both program and data memories can be accessed at the same time via two address generators carried in the CPU unit. The internal zero-wait-state RAM is contained in two blocks of 1K byte by 32 bits. The on-chip memory also includes 4096 32-bit ROM locations and a 64- by 32-bit instruction cache.

The cache can be used for short but often-used subroutines, and the ROM can be used to hold code or constants that are common to a range of applications. Standard math libraries have been implemented in some ROMs (the DSP32, for one), and such ROM libraries save valuable memory space. The on-chip memory in all the DSPs occupies a substantial fraction of the chip's real estate. In the C30 chip, almost half the 700,000 transistors are related to memory.

The chip has four 24-bit address buses, a 24-bit peripheral bus, and three 32-bit data buses. The architecture facilitates rapid execution of operations involving two variables, such as dot products and correlations. The 320C30 is a two-state machine, and peak speeds, in which a multiply-accumulate is done in two clock cycles, reach 33 MFLOPS with the standard 60-nanosecond instruction cycle.

The TMS320C30 contains a large number of parallel arithmetic commands. A "normal" three-operand floating-point multiply instruction, specified by a 32-bit instruction, multiplies the contents of Source1 and Source2 and places the result in a destination register, which is one of A_0 through A_7.

In a parallel floating-point multiply-add instruction, the CPU takes operand data from four sources (registers or memory locations). The first two sources are multiplied together, and the second two are added. The results of these two operations, carried out in parallel in the same instruction cycle, are placed into two accumulator registers (which must be among A_0 through A_3). Two of the sources must be among accumulator registers A_0 through A_7, and the other two must be data from memory (accessed via reference to pointer registers in the AGU).

Parallel arithmetic can involve pairs of a wide range of arithmetic operations, including floating-point, integer, and bit-manipulation instructions. The indirect addressing modes include various indexing operations to facilitate rapid execution of vector algorithms.

The 320C30 has several peripheral interfaces, which should lead to easy integration of the chip into host systems. Two 8-megabit-per-second serial ports permit communication with other DSPs or external devices. There are two 32-bit parallel interfaces that can be attached to external memory, a 32-bit bus of a host CPU, or other processors in multiprocessor systems. With the on-board DMA controller, you can use the I/O ports concurrently without having to start and stop the CPU.

TI has just introduced a 29-MIPS, 16-bit DSP chip tagged the TMS320C5x. This new chip is an update of its TMS320 series of 16-bit fixed-point DSPs.

John E. Hart is a professor in the astrophysical, planetary, and atmospheric sciences department at the University of Colorado in Boulder. He can be reached on BIX c/o "editors."

Digital Filters

The basic signal-processing operation is filtering, which blocks or passes selected frequencies in the data. Filters come in several types: low-pass, which eliminates high frequencies; high-pass, which eliminates low frequencies; and band-pass and band-reject, which operate on specified frequency bands.

The simplest digital filter is an averager, also known as a *tapped delay line*. Consider an input stream $x(n)$ and an output stream $y(n)$, where n is the "index" of the digital samples.

T

he basic signal-processing operation is filtering, which blocks or passes selected frequencies in the data.

Then a "four-sample averager" can be constructed that implements the following equation:

$$y(n) = 1/4[x(n) + x(n-1) + x(n-2) + x(n-3)]$$

Thus, the output of the filter is the average of the present sample $x(n)$ and the three samples preceding it. Figure 2 shows this equation as a digital filter structure using standard notation, with adders (a circle with a plus sign), multipliers (a triangle with a gain value), and delays (boxes marked with a z^{-1}, indicating a delay by one sample).

The function of this filter is easy to understand: Rapid deviations in the input signal, or high frequencies, tend to get smoothed out by the averaging function. Slower deviations, or low frequencies, remain relatively unaffected. Thus, the four-sample averager implements a low-pass filter.

The logical extension of this basic filter is a discrete version of an operation called *convolution*. Convolution consists of taking a set of filter coefficients and "sweeping" them across the stream of input data (see figure 3). At each point, the output is determined by the sum of products of the coefficients of the input data:

$$y(n) = \sum_{k=0}^{f} x(n-k)b(k)$$

where $b(0)$ to $b(f)$ are the filter coefficients and $x(n)$ is the input data. The filter structure is then rewritten as in figure 3. Note the importance of the multiply-add operation, which, as I mentioned earlier, is reflected in DSP chip design. This basic filter is known as a nonrecursive or finite impulse response (FIR) filter; given an input (an impulse), its response will decay to zero when the input is removed. [Editor's note: *For more on convolution, see "Introduction to Image Processing Algorithms" by Benjamin M. Dawson, March 1987 BYTE, and "Finding the Titanic" by Marti Spalding and Ben Dawson, March 1986 BYTE.*]

The next level of complexity is a recursive filter with feedback; its output $y(n)$ depends not only on inputs $x(n)$ but also on

Entering the World of DSPs

John E. Hart, Scott Kittelman, and Dan Ohlsen

While floating-point digital signal processors are showing up in top-of-the-line computer systems, you can already purchase DSP add-in cards for smaller systems. These cards allow you to implement many applications in areas such as chaotic dynamics, numerical analysis, and other compute-intensive scientific and engineering subjects. Table A shows a number of DSP add-in boards that are available for a variety of small systems, including the IBM PC and compatibles and the Macintosh.

If you'd like a little more hands-on experience, you may be interested in a project we developed at the University of Colorado as part of our research into the equations, formulated by E. N. Lorenz, that formed the basis for modern chaos theory. Our simple AT&T DSP32-based coprocessor board contains address-decoding logic, a 40-pin DSP32 with internal memory only, and two D/A output circuits that are driven from the DSP32's serial output. This coprocessor board can be attached to a PC, XT, AT, or 80386 machine that has the standard PC bus.

The board can be wire-wrapped on a PC prototype card, using documentation consisting of a layout diagram (for wire-wrapping), circuits, a parts list, and a wire-wrap list. To save cost, the board has no external memory, which would have to be expensive 30-nanosecond static RAM. However, the DSP32's 4K bytes of internal RAM is adequate for a wide range of small database problems, provided that I/O is handled externally through the host PC using the direct-memory-access capability of the DSP32 (e.g., many scientific problems that you would like to do interactively will be small enough to fit into the 1024-number internal capacity of the 40-pin DSP32). You can build a 25-MHz coprocessor board for about $250, or you can order it assembled and tested in a 30-MHz, 15-MFLOPS printed-circuit version called the FS-2, which requires an adjacent open slot for heatsink clearance (see below for details).

The software needed to operate the DSP32 must include some form of macro assembler that converts assembly language mnemonics into DSP32 machine code, and a device handler that can load programs and extract data from the DSP32 across the PC bus. A compiler that converts high-level language into DSP32 instructions is also helpful. The AT&T MS-DOS assembler-linker package costs $500, and its C compiler $1500.

Those who don't want to become involved in extensive low-level programming can obtain an inexpensive software package that includes a mini-BASIC compiler, a macro assembler, an interactive graphics-oriented controller, a FORTRAN interface, a small special-function library, and several demonstrations of the integration of ordinary differential equations and the generation of images using both BASIC and assembly language codes. This software can be used with either the assembled FS-2 board or your own wire-wrapped board.

John E. Hart is a professor and Scott Kittelman and Dan Ohlsen are research associates in the astrophysical, planetary, and atmospheric sciences department of the University of Colorado in Boulder. Plans for the DSP32 board are available for $5 from FASTec, Inc., 189 Mine Lane, Boulder, CO 80302, (800) 468-4142. (You can obtain the software package mentioned above for $95, including the manual, or the software package complete with the 15-MFLOPS FS-2 board for $399.95.)

Table A: *These companies make add-in digital-signal-processing boards for a variety of computer architectures, including ISA (the Industry Standard Architecture, on which the IBM PC AT and compatibles are built) and NuBus (Macintosh compatible). Development support (development system, assembler, C language, libraries, debugger, and simulator) is available for each board listed below. (Table courtesy of* DSP Update)

Company	Board	Bus	Processor	Width	Price
Ariel Corp. 433 River Rd. Highland Park, NJ 08904 (201) 249-2900	DSP-C25 PC-56	ISA ISA	TMS320C25 Motorola DSP56001	16-bit integer 24-bit Integer	$595 $595
Atlanta Signal Processors, Inc. 770 Spring St. Atlanta, GA 30308 (404) 892-7265	Banshee Chimera	ISA ISA	TMS320C30 TMS320C25	32-bit floating point 16-bit integer	$6995 $2195
Burr-Brown Corp. P.O. Box 11400 Tucson, AZ 85734 (602) 746-1111	SPV120 SPV125 ZPB32	VME bus VME bus ISA	TMS320020 TMS320C25 WEDSP32	16-bit integer 16-bit integer 32-bit floating point	$2995 $2995 $995
Communications Automation & Control, Inc. 1642 Union Blvd. Allentown, PA 18103 (215) 776-6669	DSP32-PC	ISA	AT&T DSP32	32-bit floating point	$1045
Data Cube, Inc. 4 Dearborn Rd. Peabody, MA 01960 (508) 535-6644	Euclid	VMEbus	ADSP-2100	16-bit integer	$5000
Digidesign, Inc. 1360 Willow Rd., Suite 101 Menlo Park, CA 94025 (415) 327-8811	Sound accelerator	NuBus	Motorola DSP56001	56-bit integer	$1295
Impact Technologies 2082-B Walsh Ave. Santa Clara, CA 95050 (408) 988-4980	Viper 8704	VME bus	Zoran VSP161	16-bit block floating point	$9950
Microstar Laboratories 2863 152nd Ave. NE Redmond, WA 98052 (206) 881-4286	DAP2400 series	ISA	Motorola DSP56001	56-bit integer	$2395-$3195
OKI Semiconductor 785 North Mary Ave. Sunnyvale, CA 94086 (408) 720-1900	PSP92	ISA	MSM6992	22-bit floating point	$6265
Sky Computer, Inc. Foot of John St. Lowell, MA 01852 (508) 454-6200	Challenge-S	P4 bus	TMS320020	16-bit integer	$5300
Spectral Innovations, Inc. 4633 Old Ironsides Dr., Suite 450 Santa Clara, CA 95054 (408) 727-1314	MacDSP series	NuBus	AT&T DSP32	32-bit floating point	$2295-$8995
Spectrum Signal Processing, Inc. 264 H St. Blaine, WA 98230 (604) 438-7266	56001 320C25	VME bus ISA	Motorola DSP56001 TMS320C25	24-bit integer 16-bit integer	$5995 $1995
Zoran Corp. 3450 Central Expy. Santa Clara, CA 95051 (408) 720-0444	VSPX series	ISA	ZR34161	16-bit integer	$1000-$3000

A MESSAGE TO OUR SUBSCRIBERS

FROM TIME TO TIME WE MAKE the BYTE subscriber list available to other companies who wish to send our subscribers material about their products. We take great care to screen these companies, choosing only those who are reputable, and whose products, services, or information we feel would be of interest to you. Direct mail is an efficient medium for presenting the latest personal computer goods and services to our subscribers.

Many BYTE subscribers appreciate this controlled use of our mailing list, and look forward to finding information of interest to them in the mail. Used are our subscribers' names and addresses only (no other information we may have is ever given).

While we believe the distribution of this information is of benefit to our subscribers, we firmly respect the wishes of any subscriber who does not want to receive such promotional literature. Should you wish to restrict the use of your name, simply send your request to the following address.

BYTE MAGAZINE
ATTN: SUBSCRIBER SERVICE
P.O. Box 555
HIGHTSTOWN, NJ 08520

Figure 3: *The digital filter structure shown here performs discrete convolution. The output y(n) is the sum of an input x(n) and k previous inputs, each multiplied by a coefficient ranging from b(0) to b(f). Convolution can be used in a number of different applications, such as edge enhancement in image processing.*

Figure 4: *This diagram of a general digital filter combines a nonrecursive section on the left (like those in figures 2 and 3) and a recursive section, in which the output y(n) is multiplied by a series of coefficients a(k) to a(fb) and added to the next output. Such filters are useful because their behavior closely models that of analog systems.*

Table 1: *Specifications for currently available digital-signal-processor chips (N/A = not available). (Information courtesy Nelson R. Manohar Alers and AT&T Bell Laboratories)*

DSP chip	Manu-facturer	Year announced	Multiply operands	Multiply time	Technology design rule	Power dissipation (in watts)	Instruction cycle (in ns)	Data word length (in bits)
TMS32010	TI	1982	16×16 →32	200.0 ns	2.4µ NMOS	0.9	200	16
TMS320C25	TI	1986	16×16 →32	100.0 ns	1.8µ CMOS	0.6	100	16
TMS320C30	TI	1988	32×32 →E8	60.0 ns	1.0µ CMOS	1.0	60	32
DSP56001	Motorola	1986	24×24 →56	97.5 ns	1.5µ CMOS	N/A	97	24
DSP96001	Motorola	1988	24×24 →56	97.5 ns	HCMOS	N/A	75	32
DSP16	AT&T	1986	16×16 →32	55.0 ns	1.0µ CMOS	0.25	55	16
DSP16A	AT&T	1988	16×16 →32	25.0 ns	0.75µ CMOS	0.35	25	16
DSP32	AT&T	1985	32×32 →40	160.0 ns	1.5µ NMOS	2.0	160	32
DSP32C	AT&T	1988	32×32 →40	80.0 ns	0.75µ CMOS	0.8	80	32
µPD 7720SPI	NEC	1981	16×16 →31	250.0 ns	3.0µ NMOS	N/A	250	16
NEC 77230	NEC	1986	24E8 →47E8	150.0 ns	1.75µ CMOS	< 1.0	150	32
Intel 2920	Intel	1979	···· No multiplier ····		N/A	N/A	400	25
IBM RSP	IBM	1983	N/A	2 bits/cycle	2.0µ NMOS	2.5	200	16/24
HSP	Hagiwara	1983	12E4 →16E4	250.0 ns	3.0µ CMOS	0.25	250	20
ADSP2100	Analog Devices	1986	16×16 →32	125.0 ns	1.5µ CMOS	< 0.5	125	16
DSSP-VLSI	NTT	1986	12E6	N/A	1.2µ CMOS	0.7	50	18
MSM6992	OKI Electric	1986	16E6	100.0 ns	2.0µ CMOS	0.4	100	22
µSP32	Mitsubishi	1986	32×16 →32	150/450 ns	1.3µ CMOS	N/A	150	32
MB8764	Fujitsu	1986	N/A	N/A	N/A	0.3	100	16
TS68930	Thomson	1986	T16×16 →32	160.0 ns	N/A	N/A	160	16
NS LM32900	National Semiconductor	1986	16×16 →32	100.0 ns	2.0µ CMOS	0.5	100	16

Society's signals are being digitized: Letters become faxes, and even telephone conversations can be transmitted in digital form.

previous outputs $y(n-1)$, $y(n-2)$, and so on. This feedback, however, can cause ongoing or even diverging oscillations when input has been removed. Thus, these are known as infinite impulse response (IIR) filters.

Because filters can have both recursive and nonrecursive parts, a general difference equation for digital filters can be written as follows:

$$y(n) = \Sigma(k = 1 \text{ to } fb) \, y(n-k)a(k) + \Sigma(k = 0 \text{ to } ff) \, x(n-k)b(k)$$

where $a(k)$ is the feedback (recursive) coefficient and $b(k)$ is the feed-forward (nonrecursive) coefficient. Thus, the structure in figure 4 can be produced.

For an FIR filter, the recursive coefficients are set to zero. From this equation, you see that digital filters do nothing more than calculate a linear combination of current and previous inputs and previous outputs. The filter's frequency response depends on the function determining the coefficients for this combination.

FIR filters are easier to design than IIR filters because they are inherently stable. IIR filters must be designed to avoid unstable oscillations due to feedback. The signal delay in FIR filters is the same for all frequencies: a beneficial property called *linear phase response*. IIR filters, however, provide better response curves with fewer calculations.

Determining the best filter coefficients is a complicated task and involves selecting *poles* and *zeros* (solutions to the filter's characteristic equation, which determines its behavior) in the complex z-plane. CAD programs are now available that will produce optimized filters from specifications of frequency and phase response.

The Fourier Transform
Signals are composed of varying frequencies. A stereo system, for example, provides ways to control the frequency content of music. An increase in the treble control emphasizes the high frequencies. Increase the bass, and you'll hear the lows. Similarly, a prism breaks white light into its component frequencies, a process that reveals the spectrum's rainbow. Digital signal processing also has such a mechanism, a computational prism that analyzes signals in the frequency domain. It is called the Fourier transform.

The Fourier transform takes a signal in the time domain and converts it into the frequency domain, a process that reveals its spectrum. For digital signals where a continuous signal is represented as a set of points, the discrete Fourier transform is used. Because the DFT is computationally intensive, it has been optimized in the form of the fast Fourier transform. The FFT is a recursive routine that divides an initial signal into smaller and smaller pieces in order to perform 2-point DFTs as

trivial operations. The results of these smaller operations are then scaled and combined to produce the entire FFT.

The straight DFT requires the order of n^2 complex multiplications, while the FFT requires only $n\log^2 n$, a reduction of over a hundred times for a 1024-point data set. FFT algorithms also have the advantage of working *in place*, meaning that they require no additional memory beyond storage of the initial data. Most of today's DSP chips can perform a 1024-point FFT in a few milliseconds.

I've been discussing a one-dimensional data model that fits chronological data like sound or temperature. Digital-signal-processing techniques, however, extend into higher dimensions. Convolution, filtering, and even the Fourier transform all have two-dimensional equivalents dealing with "spatial frequencies." For that reason, image processing is essentially a subset of digital signal processing, and today's image processors are often built around DSP chips.

The Future of Digital Signal Processing
Experts in the field agree that the DSP in the NeXT machine makes that system the first of a new breed of personal computer. Industry sources corroborate this view, reporting an imminent wave of new workstations incorporating DSP chips as standard "on-board" features. [Editor's note: *For a look at DSP boards currently available for personal computers, see the text box "Entering the World of DSPs" by John E. Hart et al. on page 252.*] Surely, as such systems proliferate and as DSP programming becomes simpler, there will be an explosion of diverse applications. Even more certain is that digital signal processing, like all truly innovative technologies, will extend beyond current visions and alter basic assumptions.

Society's signals are being digitized: Letters become faxes; records become CDs; speech is compressed and sent as mail. Even telephone conversations, paradigms of analog communications, are being transmitted through fiber-optic networks in digital form. But digital storage forces a kind of equivalence on various signals, removing them from their "real world" analog contexts.

In fact, digital signal processing is so broadly applicable only because, once inside a computer, "signals" are essentially all the same. Music, speech, codes, and even images can be converted to strings of numbers containing a given quantity of "information" to be distinguished and extracted from noise.

Recent controversies over digital audio tape are a good example of how digital techniques, with their capacity to make perfect copies, are calling into question concepts of originality and ownership in the information industry. Given the new equivalence it imposes on data, digital signal processing may require a rethinking of the very meaning of information: as a creation, as a signal, and as a commodity. ∎

David A. Mindell is a technical consultant with Exactitude Consulting based in Pittsford, New York. He also writes about how computers and digital processing are altering our conceptions of symbolic exchange. He can be reached on BIX c/o "editors."

VLIW:
HEIR TO RISC?

*In the race to maximize CPU performance,
a new architecture called VLIW may succeed RISC chips*

Peter Wayner

F or some people, reading and watching TV at the same time is a problem. Traditional parallel computers have a similar problem. Although they can do two things at the same time, coordinating the tasks can be a daunting and inefficient process. A new computer architecture, known as the very long instruction word (VLIW) machine, is designed to solve the problems of coordination and exploit previously untapped opportunities for parallelism.

VLIW machines are based on the simple notion that if one processor is fast, two are faster, and *n* are faster still. Designers have understood this concept for a long time and have produced many schemes for parallel machines with countless processor configurations. Unfortunately, getting a program to run twice as fast is not a simple matter of throwing two processors together in a box. The two processors must synchronize their operations; the communication between them is overhead.

If you can split a problem into *n* parts that don't need to coordinate their actions, the overhead is small and *n* processors will finish the work almost *n* times as fast as one processor. On the other hand, if the problem is particularly complicated and the *n* processors must continually talk to each other to keep track of each other's progress, one processor might be able to finish the task before *n* processors chattering endlessly could. This communication bottleneck is the most important factor to consider when you are programming parallel machines.

Some problems run naturally in parallel. Ray tracing for graphics, factoring numbers, and computing the Mandelbrot set are three examples. In these cases, it's very easy to split up the work so that one processor doesn't need to know the results obtained by the others until the end. Unfortunately, most programs are not this simple. For example, operating systems often contain thousands of branches to handle all the different cases that occur. More often than not, splitting up these programs is difficult because the *n* parts would have to coordinate with each other after every branch or jump. The overhead of communication would add so much time to the job that it would

be more efficient to let one processor handle it.

In cases where there isn't enough parallel work to overcome the total communication time lag, some work can still be done simultaneously. Usually there are little bits of several operations between the branches that offer opportunities for parallelism. For example, a program might command the computer first to fill register R1 with the results of dividing register R1 by R2 and then, in the next instruction, compute the sum of registers R3 and R4 and put the result in R5. These two instructions could be carried out simultaneously, but the time spent sending the information back and forth is much greater than the time saved by using two processors.

The VLIW solution is to build one big processor with *n* arithmetic units that connect to the same register file. The name *very long instruction word* comes from the fact that each of these *n* units must be told what to do, so that, consequently, the instruction word must be *n* times longer. Because all the processors work with one set of registers, the communication delay is virtually nonexistent. (You could say there is some delay, because this special register file that can talk with *n* arithmetic units is slightly slower than a regular one.)

Just a Jump to the Left...
Adding more arithmetic units is just part of the solution. System code may average only two or three operations between branches, and, with more arithmetic units, the program would speed up by a factor of only two or three. This result isn't bad, but a better solution must work around branches to speed things up even more. If the VLIW machine executes several arithmetic operations at once, it can certainly do a branch at the same time, too.

The obvious solution is to have the CPU calculate operations from both before the branch and after it. When it decides whether or not to jump, the CPU can keep the results from the calculations along the path it chose and throw away the results from the path it didn't take. In this case, the machine does more
continued

work than absolutely necessary, but in total, more work is saved at the end. Here is an example. Two instructions come before the branch:

$$R1 + R2 \rightarrow R1$$
$$R3 + R3 \rightarrow R3$$

At the branch, the computer executes

$$R5 \times 2 \rightarrow R5 \text{ if } R4 > 0$$
$$R6 \times R6 \rightarrow R6 \text{ if } R4 \leq 0$$

Notice that all the operations can be done at the same time because none of them depends on the results of the other. If the CPU is not designed to deal with branches, it can execute the first two operations in one cycle. Then it can test to see if R4 is greater or less than zero. Finally, it would execute one of the operations on R5 or R6. That would take three cycles. It's easy to see how branches can prevent the computer from doing much work in parallel.

On the other hand, a VLIW machine could execute all four operations and test to see if R4 is greater or less than zero in one operation. It would write the results of the first two operations to R1 and R3 automatically. If R4 is greater than zero, it would save the $R5 \times 2$ in R5 and throw away the calulation of $R6 \times R6$. In the other case, when R4 is less than or equal to zero, it would do the opposite. This entire process would take only one cycle—a big saving in time over the three cycles used before.

Obviously, there are limits that keep a VLIW computer from executing a huge program in one cycle. In the example, all four operations can be performed simultaneously because none of them depends on the results of another. Interdependencies often prevent large parts of a program from being executed in one cycle. An operation after the branch might use the results in R5 that were computed by an operation before the jump.

It is hard to know how many times this situation will happen in real-world programs because every program has a different set of interdependencies. Finding the parallelism is the job of the compiler, and very sophisticated compilation techniques are necessary.

Determining which operations can be done simultaneously is

Figure 1: *A fragment of code with five instructions and two branches. The trace-scheduling compiler chose the three emphasized instructions by predicting the computer's path at each of the branches.*

a straightforward process, but it's extremely tedious. You could program VLIW computers in machine code, but you'd have trouble trying to keep the entire CPU working. And if the work can be automated, there's no reason not to let the computer do the work. The only problem is being able to design the algorithms to do the job as efficiently as possible.

The two major methods for creating VLIW code are known as *trace scheduling* and *percolation scheduling*. Both of them compact the code so that operations that can be done simultaneously end up getting done simultaneously, but the two methods are based on different visions of how to do it.

Trace-Scheduling Compiling

Trace scheduling was developed by Josh Fisher and several of his graduate students at Yale. It assumes that a computer spends its time executing one particular path or "trace" through the program. The computer may occasionally follow one branch off the path, but a trace-scheduling compiler hopes that the process will soon return. Once the compiler picks the trace, it compacts the code along the trace and moves all operations that can be performed simultaneously into the same instruction.

Compiling a program for a VLIW machine is a matter of guessing the right path before the program executes—often a difficult challenge. But, in some cases, as with loops, the trace is simple to find because the program will probably jump to the beginning of the loop again. Unfortunately, the compiler can't always predict many other branches. Half the time it will be right, and half the time it will be wrong.

Figure 1 shows a fragment of code arranged in a tree with the trace chosen by the compiler emphasized. When compiled, figure 1 becomes

```
Inst1: R1+R2->R1 ; if R9<=0 then Inst2 else R4+R5 ->R5;
       if R8<=0 then Inst3 else R7*2->R7
Inst2: R3*2 -> R3
Inst3: R6*2 -> R6
```

Notice that the computer will be doing three operations at once. The compiler predicted that the machine would probably find that both R8 and R9 are less than zero. The branches decide which results will be kept and which will be forgotten. If the branches decide to jump out of the trace, everything after the branch is thrown away. The three operations are all from the trace that the compiler chose. If it made a mistake and R9 turns out to be greater than zero, all the extra work was wasted.

Percolation-Scheduling Compiling

Percolation scheduling is a more general model for VLIW machines. It was invented by Alex Nicolau, a former graduate student of Josh Fisher's and now a professor at the University of California at Irvine. This method of compiling treats the set of operations executed on each cycle as a tree with branches instead of a straight line in a trace. This procedure saves the processor the trouble of predicting a particular trace. The machine must execute all the operations in the tree and save the ones from the path the process takes.

Because the various combinations of branches can lie in a tree, not just a trace, this method is more general. A percolation-scheduling compiler could pack all the operations from figure 1 into one instruction. The computer would be simultaneously executing all five instructions and picking the results it wanted after it also computed the branches.

The compiler creates the instructions by "percolating" all the instructions up as far in the program as they can go without changing the action of it. It starts with ordinary code and places

Intel's 80860: On the Road to VLIW

Early this spring, Intel introduced the 80860 and started calling it a "Cray on a chip" because, under ideal conditions, the microprocessor reportedly can reach speeds nearing those of the early Cray supercomputers. To accomplish this feat, Intel implemented one of the main design tenets of very long instruction word (VLIW) machines—the ability to start more than one instruction at the same time.

The chip has a RISC processor that handles the branching and the integer instructions. An FPU on the same chip can simultaneously do a multiply and an add. That means that, under ideal conditions, three operations can be done concurrently and that, running at 50 MHz (peak performance), one of these chips will process 150 million operations per second.

Naturally, the best rate doesn't occur in all cases. The perfect program for pushing the chip to its maximum has equal amounts of integer operations, floating-point adds, and floating-point multiplies that can be executed simultaneously. The floating-point hardware has a pipeline that must stay filled to achieve the one multiply and one add per cycle. Gaps in the floating-point operations in the program drain the pipeline, and it takes several cycles to restart.

The one major application that will be able to use the chip structure well is graphics. If you want to draw pictures on the screen, your computer needs to be able to perform a vast number of floating-point calculations. The 80860 thrives on this kind of work because the graphics computations keep its floating-point pipeline filled. The chip hardware, in fact, is especially tuned to easily handle several routine computations frequently used in graphics software.

The 80860 does not implement all of VLIW's design ideas. It can carry out only one branch instruction in each clock cycle. This capability is not particularly necessary when only three different processors are used. In the future, more functional units will require multiway branching to make the best use of all the different units.

The 80860 is one of the first major chips to start down the path to implementing VLIW architectures. Apollo Computer has a RISC chip that issues multiple instructions per cycle. Weitek introduced a chip set that can also perform three instructions per cycle. It will be only a matter of time before more processors begin to be able to operate in this manner. The RISC philosophy will lead to faster and faster chips well into 1992, but after that, CPU architects will need to explore other paths like VLIW machines. The 80860 is one of the first steps.

---

each instruction in a node by itself and then considers the program as a list of nodes to be executed one after the other.

The compiler begins to percolate by comparing each node with the node that precedes it. If the instructions in both nodes can be simultaneously performed without interference, the compiler merges the two nodes. If it finds that only some of the lower node's instructions can be moved to the upper node, it will move just those. The compiler wants to move as many instructions as it can as far up the program as possible. This action is essentially the same as the compression that the trace scheduler performs, but it's done in a more general way.

The differences between the two methods become apparent when one of the nodes contains a branch. If the branches and the instructions in the two nodes can be performed simultaneously, the compiler merges the two nodes, and the shape of the operations inside the node begins to look like the tree in figure 1. Sometimes these trees can grow quite bushy, and the percentage of work kept by the machine decreases because only the work from one trace is stored to the registers. In effect, this loss of efficiency illustrates part of the law of diminishing returns. More and more processors executing bushier and bushier trees are necessary to save executing one additional operation.

Machines Now and Tomorrow

Some VLIW computers are already on the market, but they are mainly large machines aimed at the minicomputer/supercomputer market. Smaller versions are already announced for microcomputers, and larger ones will certainly follow. The Intel 80860 also exhibits many VLIW-like qualities (see the text box "Intel's 80860: On the Road to VLIW" above).

The Multiflow Computer Company of Branford, Connecticut, makes a machine aimed at the minicomputer/supercomputer market. The basic model, the Trace 7, can simultaneously perform seven operations. It has two units for floating-point arithmetic and two more for integer math. The integer units run twice as fast as the FPUs, so six operations are performed at the

same time. The compiler does trace scheduling.

Kemal Ebcioglu at the IBM Yorktown Heights research lab is currently building an experimental VLIW computer. This machine is designed to use a percolation-scheduling compiler that will create tree-like instructions. It will have 16 ALUs and eight units for loading from and storing to memory. While the ALUs evaluate these operations, the CPU can also choose one of 16 different paths on the tree and base the results that it keeps on this information.

The IBM machine will have 128 registers. Since 16 arithmetic operations with two operands and eight stores can happen with each cycle, the register file must have 48 different ports. Each ALU will be able to access any of the 128 registers. Most of the computer's other parts are coming from "off-the-shelf" silicon designs, but this register file required a special effort.

The compiler for the IBM computer was prototyped long before any of the details about the hardware were finalized. Many different programs were compiled, and the resulting data was used to guide the design. An earlier version of the hardware had only eight functional units because the compiler had only produced nodes with eight instructions in them. Then a new way to compile and percolate loops yielded some examples where as many as 16 functional units could be kept busy. The hardware design was expanded to take advantage of this.

These experiments also showed that a real application program would use about six of the functional units on average, meaning that the program would run about six times faster on a VLIW machine than on a machine with one processor. Unfortunately, it is difficult to keep all 32 functional units going at once, and so it is rarely possible to get an increase in speed directly proportional to the increase in functional units.

The experimental results showed that the CPU discarded an average of 40 percent of the operations because the program took only one of the possible paths. The percolation-scheduling compiler had scheduled many operations from all these paths in

anticipation of taking just one. You can readily see the law of diminishing returns at work here.

VLIW: Extended RISC?

Many of the precepts that guide VLIW design are extensions of the RISC philosophy. In many ways, the VLIW is the next logical step. Like the RISC chip, the IBM VLIW machine requires

For the
hardware to be as simple
as possible, RISCs rely on smart
compilers to rearrange the code.

all arithmetic operations to obtain their operands from registers and return all the results in registers. Loads and stores to memory must be performed separately. This separation allows the machine to perform as much work as possible without being slowed down while the bus obtains information from memory.

In order for its hardware to be as simple as possible, RISCs rely on smart compilers to rearrange the code. Ideally, RISC chips can start one new instruction each cycle, even when the previous instruction hasn't finished executing. The compiler must arrange the instructions so there are no conflicts with other instructions waiting for available data.

If one instruction stores its results in one register, the compiler will try to follow it with an operation that need not wait for the result of the first operation. Thus, the CPU can start work on the next instruction without waiting for the results of the first. VLIW machines take this process a step further. Instead of overlapping instructions as much as interdependencies permit, the machines perform all instructions concurrently.

There are trade-offs. The VLIW machines need more hardware, which has to be more sophisticated. There must be n functional units where before there was one. The traffic between the CPU and memory now comes in large bursts of n loads and stores instead of a steady stream of single requests. The bus must be larger and faster. On the other hand, super-RISC chips need state-of-the-art fabrication to keep the cycles short enough. They must be $1/n$th the length to compete with an n-arithmetic-unit VLIW processor.

These are a few of the reasons why some people think that a super-RISC chip could achieve close to a VLIW machine's performance at much lower cost. These trade-offs also illustrate just some of the simple issues that designers must balance when dealing with questions about the architecture. There are many more. Changes in the development of fast memory, caches, bus capacity, and silicon fabrications will all affect the balancing.

In several years, microcomputer processors will become available that can do two, three, or n things at once. The addition of functional units is the natural way to speed up microprocessors. The use of large-scale parallelism with many autonomous processors can be great for scientific applications that need simultaneous calculations performed. Word processors, window managers, and databases, however, are all difficult to move to parallel machines because there are not n little tasks that can be done simultaneously. The VLIW approach is better suited for fine-grained parallelism.

The precepts that form the canon of VLIW design are far from fixed. Only a few machines exist, and they are large and expensive. Time, experimentation, and lots of research will eventually resolve the questions regarding how many functional units will be considered optimal, how many loads and stores should be performed per instruction, how many branches are necessary, and other questions not yet thought of. In time, the evolution of microprocessors will begin to follow the lead of these large machines. ■

Peter Wayner, a graduate student at Cornell University's department of computer science, helped design part of the IBM VLIW project's compiler. He can be reached on BIX c/o "editors."

World Radio History

NEW FROM PERISCOPE

Real-Time Debugging for 386s and 286s

User *Jeff Garbers, Crosstalk Communications' Director of Software Development, has been debugging with Periscope® Model IV and says, "The hardware really makes Periscope shine, especially when you've got timing-related problems. I can now track down changing pointers and altered buffers on my 386. I've been using it to debug Crosstalk® Mk. 4 and there's just no better way to do it."*

Periscope IV gives you the ability to debug time-sensitive programs, hardware-interrupt routines, and programs with intermittent errors. You can run your program at full speed while tracking down unwanted memory overwrites. You can use the information captured in Periscope IV's real-time trace buffer to see EXACTLY what the system is doing, and to improve its performance.

Periscope manual, disk, quick-reference card, and 512K Plus board that provides write-protected RAM for Periscope Models I, III, and IV.

A New Generation of Hardware-Assisted Debugging

Compatible with virtually any 286 or 386 with an AT-style bus, Periscope IV works on machines running up to 25MHz with any number of wait states. Because it gets information directly from the CPU, instead of from the system bus, Model IV is not sensitive to bus compatibility issues.

Periscope IV collects CPU information in its hardware trace buffer while the CPU runs at full speed. Whether you tell Periscope IV to capture just selected information or to capture everything, you can use its powerful trace buffer commands to search for and display the execution history the way you need to see it. And you can use the CPU cycle count information to get the last bit of performance out of your code.

With Periscope IV you can set hardware breakpoints on memory accesses (within the first 16MB), I/O ports, and data. You can also set breakpoints on the occurrence of specific sequences of events, such as "watch for the routine FOO to begin executing, then while it is, watch for the variable BAR to be written." This capability, called sequential triggering, enables you to define complex conditions, then stop your program and examine what has happened when these conditions occur.

If you're developing a large application that needs all of the lower 640K, you can use the optional Plus board to keep Periscope totally out of normal DOS memory. The Plus board requires the use of a second slot.

The Periscope IV software is an extension of the software that comes with all models of Periscope. So, along with Periscope IV's powerful hardware, you get a full-function software debugger with source and symbol support for most popular PC compilers and linkers, Microsoft® Windows support, PLINK overlay support, dual monitor support, support for debugging device drivers and TSRs as well as regular programs, DOS independence, crash recovery, ease-of-use, and much more.

Prices on Periscope range from $145 for software-only Model II-X to $2,995 for a 25MHz 386 hardware-assisted Model IV. Call 800/722-7006 for pricing details, free information, to talk about your debugging needs, or to order your Periscope.

Order Your Periscope, Toll-Free, Today! 800-722-7006

MAJOR CREDIT CARDS AND QUALIFIED COMPANY PURCHASE ORDERS ACCEPTED

The Periscope Company, Inc.

**1197 PEACHTREE ST.
PLAZA LEVEL
ATLANTA, GA 30361
404/875-8080
FAX 404/872-1973**

HARD DISK MAINTENANCE SOFTWARE

Clever programs optimize your disk's interleave and fix potentially destructive errors

Are you getting optimal performance from your hard disk drive? Did you know that potentially destructive errors—intentionally ignored by DOS—may lurk uncorrected on those rapidly spinning platters? In this installment of Under the Hood, I'll look at the technology behind *low-level disk maintenance utilities*—programs that can optimize the arrangement of the sectors on each track of your disk while testing for (and often repairing) hidden glitches that can lead to loss of data.

Reserving Space for Data
The data on a hard disk is laid out on one or more platters in concentric *tracks*. If the disk has more than one platter, it usually has a read/write head for each surface; the heads are mounted on a single arm that moves them simultaneously across the tracks. A group of tracks that lie directly above and below one another is called a *cylinder*.

Every track of a hard disk is organized into sectors, each of which contains an equal portion of the data stored on that track. Obviously, the disk drive controller needs to be able to tell the sectors of a track apart to deliver the right data; the method it uses depends on whether the disk is *hard-sectored* or *soft-sectored*.

If the disk is hard-sectored, the drive hardware is responsible for remembering the physical locations of the sectors on a track. The controller receives a signal telling it which sector is passing under the head at any given moment; it doesn't

need to look for identifying marks on the disk itself. Hard sectoring can be very efficient because no space needs to be reserved on each track to mark the sectors. However, because the electronics necessary to do hard sectoring add to the cost of a disk drive, most drives (including virtually all those on personal computers) use soft sectoring.

On a soft-sectored drive, each sector is preceded by a special *sector ID header* that gives the number of that sector. (It also gives the number of the head and cylinder so the controller can make sure an error isn't causing it to access the wrong place on the disk.) Because the sectors can be laid out in different ways on the track (see the text box "Interleaving: Delivering the Data on Time" on page 266), the controller must read each header until it finds the sector it wants to access. Then, if it's going to write to the sector, the controller must shift from reading to writing during the short gap between the header and the sector. The headers for all the sectors must be written to the disk before the disk can be used. This process is called *low-level formatting*.

When you format a floppy disk on a typical computer, the computer actually performs a low-level format. Not so, however, for hard disks, where low-level formatting can take a long time and is usually done by the dealer before delivering the machine. When applied to the hard disk, the FORMAT command in PC-DOS only sets up the file allocation table, the root directory, and (optionally) the operating system. In fact, FORMAT won't work on a hard disk at all unless the drive has already had a low-level format (using IBM's Advanced Diagnostics) and has been partitioned with the FDISK utility.

The vast majority of users have never done a low-level format. This means that while the data on their hard disks may have been rewritten thousands of times, the low-level formatting information—the original headers that show the con-

troller where the sectors are—have never been rewritten. As the drive ages, this can become a source of serious problems, as I'll discuss below.

Alignment Drift
All electromechanical devices—including disk drives—age over time. Parts wear, tolerances shift, components drift out of alignment. Even on a new drive, the alignment may change a bit depending on whether or not it's had a chance to warm up. All these factors can combine to cause the same symptom: The heads on a hard disk drive no longer wind up exactly the same distance from the center of the disk when the drive steps to a particular track.

Alignment drift can lead to several interrelated problems. The data portion of each sector can drift inward or outward relative to the sector ID header (see figure 1). Since the information written to the disk during a low-level format is never rewritten during normal use, it is skewed relative to the sector itself—and may get so far away that the sector cannot be located at all. It's even possible for the contents of a recently written sector to overlap (and destroy) the contents of an adjacent sector that hasn't been written to for a long time.

Floating Defects
"Floating defects" are another source of potential problems. The quality of the disk surface is critical to the operation of a disk drive; the tiniest flaw can make an area of the disk unusable for data. So, when a hard disk drive is new, the manufacturer carefully tests it (using equipment costing hundreds of thousands of dollars) to find defects in the hard disk surface. He then writes a "defect map" on a label and attaches that label to the drive. The defect map shows the head and track where each defect is located; you can use this information to alert your low-level format software to the presence
continued

of defects. On some drives, the manufacturer also writes a defect map on a special track of the drive for the software to use. (It shouldn't alarm you, by the way, to find out that your drive has some surface defects; it's quite rare to find a drive that doesn't have any.)

If the drive's alignment didn't drift, the defect map would be an accurate indication of where the flaws are (at least all of them that existed at the time of manufacture). But since the map shows only defects that are *in* the tracks, not *between* the tracks, it's possible for defects that

were not mapped to show up as the tracks drift (see figure 2). You can solve this problem in part by doing a *surface analysis* when you perform a low-level format on your hard disk drive. But if the tracks drift between low-level formats, floating defects can still appear amid the data.

Interleaving: Delivering the Data on Time

Each circular track of a hard disk is divided into sectors—arcs of the circle that contain equal portions of the data stored on that track. You may well ask, "Why don't they make the entire track one huge sector?" The answer is that the disk drive controller must always read or write whole sectors at a time. Having only one sector per track would mean that every read or write would require as much as two revolutions of the disk: up to one revolution to get to the beginning of the track and another full revolution to read it. (The designers of the Commodore Amiga, incidentally, tried to implement this approach with floppy disks, but they added a special trick. The unique Amiga disk drive controller can start a read or write operation at any point on the track—something no other controller I know of can do. This sets the time

for every read or write to exactly one revolution of the disk. Alas, the latency is still a bit long, causing the Amiga floppy disks to exhibit lackluster performance except on large files.)

Each track of a standard IBM PC hard disk contains 17 sectors of 512 bytes each. The outermost ring in figure A shows the most obvious arrangement of the sectors: They're placed in ascending order around the track, from 1 through 17. (This is called *1-to-1 interleave*.) In practice, however, this might not be the most efficient arrangement. Often, disk drive controllers, disk I/O routines, and the host systems they run in require time between accesses to successive sectors. They may use this time to transfer data to and from memory, acquire control of the system bus, set up direct-memory-access channels, or allow other I/O to take place. If the

time required for these tasks is too long, the controller may find that the next sector it wants is already under the disk drive head—or past it—by the time everything is ready.

Interleaving solves this problem. If, instead of following one another, sectors with successive numbers have one or more other sectors between them, the next sector will be approaching the disk drive head just when the controller is ready for it. The second ring from the outside in figure A shows an example of 2-to-1 interleave, in which sectors with successive numbers always have one other sector between them. The order becomes 1, 10, 2, 11, 3, 12, 4, 13, 5, 14, 6, 15, 7, 16, 8, 17, 9.

If the system can keep up with it, 1-to-1 interleave will generally provide the best performance. But there are severe performance penalties if the interleave factor is too low. The controller will "miss" each sector—possibly by only a few hundred microseconds—and will have to wait until it comes around again. Since a typical hard disk drive spins at 3600 revolutions per minute (60 revolutions a second), the time to read all 17 sectors of a track would become 17 sectors × 1 revolution/sector × 1/60 second/revolution, or about a third of a second. This is a long time for only 8.5K bytes of data!

If the interleave factor is set one notch too high (say 3-to-1 instead of 2-to-1), the penalty isn't nearly as bad. The controller will wait an extra 1/17th revolution per sector; that's 17 sectors × 1/17 revolution/sector × 1/60 second/revolution, or 1/60th second longer, to read all the sectors.

The optimum interleave may be different even for two operating systems on the same machine. On my 8-MHz AT clone—not a particularly fast machine by today's standards—DOS works best at a 1-to-1 interleave. OS/2, however, likes a 2-to-1 interleave; the intervening sector gives the system time to handle interrupts and switch in and out of protected mode as needed.

Figure A: *Arrangements of sectors for different interleave factors. At 1-to-1 interleave, the sectors are numbered in sequence; at 2-to-1, sectors with consecutive numbers are separated by one other sector, and so on.*

Magnetism and Friction

Other hard disk problems stem from the fact that data is recorded magnetically on the disk. Each time you turn your computer off or on, a small pulse of current travels through the read/write head, potentially weakening or erasing the magnetic domains underneath. The head may also jerk to one side at this time, spreading this small burst of magnetism across the disk. (This is a good reason to get a drive with a head that retracts automatically when the disk is turned off; many of the better models do.)

Disk drive heads, like the heads of audio tape recorders, accumulate residual magnetism that can partially erase the data over which they pass. And if the power is turned off at just the wrong moment during a write operation—or if the computer is reset at the wrong instant—it's possible to leave scrambled data on the disk. (Some computers and drives are designed to prevent this, but they're far from infallible.) All these factors can weaken or destroy the data—or, worse, the low-level formatting information—on your drive.

If the data portion of a sector is damaged, you have lost only the data—which is bad but not catastrophic. But if the low-level formatting information is damaged, you are in more serious trouble. There is no easy way to access the affected sectors, and many operating systems have no way of "learning" to avoid them. The result is persistent disk errors that can be cured only by performing a low-level format on the entire disk.

On disk drives without head retractors, the head comes to rest in the middle of the disk surface when the power is shut off. When the drive powers up again, the head "takes off" and flies on a very thin cushion of air above the platter; however, it takes a bit of time for this to happen, and some abrasion can occur as the head slides along the lightly lubricated disk surface. (If the frictional force between the head and the platter is strong enough, it can prevent the disk from starting to spin at all; this phenomenon is known as "stiction.") The mechanical damage to the disk surface caused by this abrasion can result in loss of data.

Disk Errors, ECC, and PC-DOS

What does your operating system do about potential (and actual) disk damage? The answer varies, of course, from manufacturer to manufacturer, but one of the most common operating systems—IBM PC-DOS—actually aggravates hard disk problems by ignoring them until it's too late.

Figure 1: *This diagram shows what happens when alignment drift causes the data portions of sectors to move out of line with the sector ID headers.* (a) *Before drift, there is good sector alignment;* (b) *after alignment drift, some sectors may even clash. (Figure courtesy of Prime Solutions)*

As you've gathered by now, errors on hard disks are a common occurrence, and controller manufacturers have worked out ways to deal with them. Every sector of every modern hard disk includes not only data but also an ECC—an *error correction code*—that detects errors and lets the controller fix up to 11 consecutive incorrect bits. When an error occurs while reading the disk, the controller first retries the read as many as 40 times (depending on the settings and the make of the controller). Then, if it's still unable to read the data 100 percent correctly, it applies an error correction algorithm in an attempt to restore the data. This scheme affords plenty of margin for error and is responsible for the trouble-free operation of most hard disk systems.

When the controller successfully retries a read, it normally won't notify the BIOS at all. If the data is corrected using the ECC, the BIOS gets a message that says, "I was able to correct the data, but there are errors on the disk that should be fixed before the sector becomes unreadable." In the IBM PC, the BIOS relays this warning to PC-DOS.

And what does PC-DOS do with this message? Does it rewrite the data before it becomes even more illegible? Does it warn the user? Amazingly enough, it does neither of these things; it ignores the error entirely. In fact, the file IBMBIO-.COM—a low-level driver that's loaded when PC-DOS boots—contains a short program that "captures" Interrupt 13 (the software interrupt vector that does disk I/O), intercepts BIOS calls, and filters out all the messages that report correctable errors. The result is that even programs that use the BIOS for direct disk access no longer receive this vital warning message.

DOS's "ignorance is bliss" approach to disk errors "works" only until enough errors accumulate to make the data uncorrectable. When the controller finally fails to correct the data, DOS will at last stop ignoring the problem. You'll get an error message from DOS: "Bad sector error on C:." But at this point, there are guaranteed to be errors in your file. Furthermore, DOS will refuse to read past the first error, so you cannot access any good data that follows.

Two Real-World Products

If DOS isn't willing to cooperate, what can you do about gradual disk degradation, alignment drift, and other problems? One solution is to periodically back up your entire disk and do a low-level format. While backing up is always a good idea, a complete low-level format takes a lot of time. You'll also have to repartition your disk and reinstall DOS—not exactly the kind of procedure you'll want to do very often. Finally, unless you're lucky enough to have a tape backup unit, you'll need to shuffle the backup disks twice—once during the backup and once during the restore process. The whole operation could take an entire morning to complete—especially on a large hard disk.

Fortunately, there are at least two nondestructive low-level disk maintenance

Figure 2: *Alignment drift can cause "floating defects" not found in the manufacturer's defect map to appear in the data areas of a hard disk. The upper defect, originally harmless because it was between tracks, is now in the middle of the data, while the lower defect is now between tracks. The sector ID headers are now so far from the current track center that the sector may not be locatable. (Figure courtesy of Gibson Research)*

programs on the market (and I'm sure there are more I haven't heard about). SpinRite, by Gibson Research, and Disk Technician/Disk Technician Advanced, by Prime Solutions, perform similar functions: They'll find physical defects that can cause data loss, correct "soft" (correctable) errors on your hard disk, refresh the information on the disk (including the low-level formatting information), and even adjust the interleave factor to an optimal value. Because the programs are "nondestructive," there's normally no need for you to restore your data from a backup (although both manufacturers do recommend that you perform a backup before you use the program).

It's important to understand the distinction between these two programs and programs that work at the DOS level, such as the Norton Disk Doctor. DOS-level utilities can repair damage due to accidental formats and erasures. But they cannot see errors that DOS ignores—or repair damage to the disk's "infrastructure" (i.e., the low-level format).

The makers of SpinRite and Disk Technician Advanced make different (and often contradictory) claims about their software. In the remainder of this article, I'll give you the information you'll need to cut through the marketing hype and understand how these utilities really work.

Circumventing DOS

To detect the true condition of the hard disk, both programs need to circumvent DOS's BIOS "patch" and find out when there's really an error on the hard disk. One technique—used by Gibson's Spin-Rite—is to bypass DOS's path routines and call the underlying ROM BIOS routines directly. Since DOS overwrites the Interrupt 13 vector at boot time, it seems as if there's no way to determine the original entry points of the ROM BIOS disk routines. But technical wizard Steve Gibson has a clever trick up his sleeve; Spin-Rite uses a technique that really does find those entry points.

The IBMBIO.COM file is loaded into memory when DOS boots; this is the time when the BIOS interrupt vectors are "captured" and redirected. But when you install SpinRite, you first boot the system with a floppy disk that does not have DOS on it. Instead, the boot track of the installation disk contains a custom program that reads the addresses from the interrupt vector table and saves them on the disk. Once these addresses are known for a given computer's ROMs, they won't change unless you change BIOSes—and SpinRite can use them to bypass DOS and call the ROMs directly.

Two other possible techniques are to read the controller status from the BIOS data area or get the status information from the controller itself.

Detecting Retries

Bypassing IBMBIO.COM allows a program to discover errors that were corrected via the ECC—but what about retries? The controller chips used in most PC and AT hardware adapters never report the fact that they've retried while reading a sector—not even to the BIOS. Most of the claims made for the different programs center around their ability to detect these errors that normally go unreported.

Disk Technician looks for retries using a timing approach. The disk drive controller has to wait for disk data to pass under the head again when it retries a read; Disk Technician detects this delay via the system clock and sensitive timing loops. Prime Solutions' literature states that this is the only possible way to solve the problem: "Disk Technician Advanced is the only software able to detect that the controller had to retry a sector." But engineers at controller manufacturer Western Digital say otherwise, noting that the IBM BIOS (which WD developed with IBM) has a provision for disabling retries. When this option is set, a

Circle 41 on Reader Service Card (DEALERS: 42)

World Radio History

correctly implemented BIOS and controller stop immediately and show *any* error instead of retrying. They also note that a timing approach may react to disk speed variations and unexpected interrupts.

Testing the Surface

Both SpinRite and Disk Technician can test the surface of your drive for defects

There's no way to make SpinRite or Desk Technician work with ESDI or SCSI drives right now.

not found by the original manufacturer's tests. They do this by writing "worst case" patterns—patterns that are most likely to cause errors—to the drive and seeing if they can be read back correctly.

Each program has a unique, proprietary set of patterns that changes with the encoding scheme used on the disk. Modified-frequency-modulation (MFM), run-length-limited (RLL), advanced-run-length-limited (ARLL), and enhanced-run-length-limited (ERLL) disks all require different signals to "bring out the worst" in a disk drive. The "worst" patterns tax the data separator and phase-locked loop circuitry by forcing maximum and minimum run lengths and going from high pulse rates to low ones and back again.

Adjusting the Interleave Factor

Both programs will also adjust the interleave factor on your disk drive for optimal performance. A caveat is in order, however: The optimal interleave suggested by the program may not be the best one under all circumstances.

For instance, if you're running OS/2 (both SpinRite and Disk Technician Advanced run under DOS), you may need to use a higher interleave factor; also, TSR programs may slow your system performance just enough to require a different interleave. If you're not sure, the tactic I would recommend is to take advantage of the nondestructive nature of these programs and carry out empirical performance tests to see which interleave factor actually works best.

Recovering Bad Tracks

One useful function both these programs will perform is to recover sectors that were unnecessarily marked as bad during a low-level format. Most low-level formatting programs will mark an entire track as bad if even one sector has a defect in it; this is usually overkill, because the rest of the track is likely to be good. Because these utilities can perform extensive tests on each sector of the track, they can restore the unblemished space to active use.

Parking the Head

Both packages also offer head-parking programs, which are designed to solve the problem of data loss due to friction and transients at power-on and power-off. SpinRite comes with a utility called PARK.COM that parks the heads; you run it before powering down.

Disk Technician comes with a TSR program called SPA (SafePark Advanced). This TSR program moves the disk drive heads to a "safe" area whenever the drive is inactive for a predefined period—a good idea if you aren't good at remembering to execute a program.

The Translating Controller Problem

A standard IBM PC hard disk has 17 sectors per track, with 512 bytes per sector. This configuration is ideal for a standard MFM drive, but drives that use different encoding schemes—such as RLL—have different optimal configurations. Most RLL controllers, for instance, can squeeze 26 512-byte sectors on a track; the Plus Development HardCards, which pack data very tightly indeed, have a variable number of sectors per track (more on the outside than on the inside).

DOS is capable of handling disks with more than 17 sectors per track, but some other operating systems (such as older versions of NetWare) aren't. To run with these environments, many controllers perform *sector translation*. A translating controller "tells" the system that the disk has 17 sectors per track and more cylinders than it really has but keeps the total number of sectors in the "imaginary" drive the same as the number that exist on the real drive. Then, when the system does I/O to a sector on the imaginary disk, the controller uses a simple algorithm to pick a corresponding sector on the real disk and diverts the request to that sector.

Neither SpinRite nor Disk Technician Advanced is equipped to work with a translating controller. These utilities need to know which sectors are really on

World Radio History

which tracks in order to perform inter-
leaving or low-level formatting; they
won't be able to do the right thing if the
controller is trying to "outsmart" them.

Fortunately, on the Western Digital
1002-27x, one of the most common RLL
controllers, translation can be turned off
by moving a jumper. You can use the
controller with these packages if you
back up the disk, shift the jumper, and
then repartition and reformat.

No ESDI or SCSI . . . Yet

Unfortunately, there is no way to make
either SpinRite or Disk Technician work
with ESDI or SCSI drives right now. In
ST-506/MFM and RLL configurations,
the controller contains the data separator
and determines the encoding scheme.
But in ESDI, the data separator is in the
drive itself. Different drives can use dif-
ferent encoding schemes; most ESDI
drives use 2,7 RLL, but there's nothing
to stop them from using 1,7 RLL, group
codes, or even zone bit recording (ZBR).
The program must know which is used to
determine the correct worst-case test
pattern.

Another problem can also keep low-
level reformatting programs from work-
ing with ESDI drives. PC-DOS and some
PC BIOSes have an internal limit of 1024
cylinders per drive. Most ESDI drives
above 300 megabytes have more—my
Maxtor XT4380, for example, has 1222
cylinders. So, in order to let DOS use all
of that big hard disk, I use a translating
ESDI controller.

This controller (in my case, the DTC
6280) "fools" DOS into believing that
the disk has *more* sectors per track and
fewer cylinders than it really has. An op-
erating system that requires 17 sectors
per track won't run with this configura-
tion, but DOS doesn't mind (nor does
OS/2, because the controller is register-
compatible with the standard AT control-
ler). Even so, a low-level maintenance
program won't be able to do its work
properly unless it knows the true con-
figuration.

SCSI presents a similar problem.
Blocks on a SCSI device are normally
referenced by their "absolute" numbers
relative to the start of the disk; the at-
tached computer is not supposed to
know—or need to know—how the sec-
tors are laid out or how the data is en-
coded. This makes it difficult for any op-
timizer or disk fixer to do more than
refresh the data in each sector and per-
haps look for errors.

Finally, *no* low-level reformatting
utility will work correctly on Priam hard

disk drives. Priam explicitly forbids the use of any such utility on its disks, which require a special type of low-level format.

Long-Term Maintenance

Although both software vendors recommend that you run their programs daily or at least weekly, SpinRite and Disk Technician Advanced take distinctly different philosophical approaches to the problem of long-term hard disk mainte-

nance. The SpinRite documentation reads like a good tutorial on hard disks. Hard disk errors are inevitable, it says; but with regular maintenance, even sectors that have slight defects can be kept in use. You're encouraged to be knowledgeable about what goes on inside your hard disk drive and take control.

Disk Technician Advanced has fewer technical explanations. The documentation simply assures you that its "artificial intelligence" features will take care

of all your hard disk problems. Unlike SpinRite, Disk Technician Advanced maintains a database with cumulative information on where errors have occurred. This lets the program track repeated trouble spots. Disk Technician Advanced does not appear to be as tolerant of "weak" sectors as SpinRite; it takes them out of service rather than refreshing them regularly.

Peace of Mind?

Both SpinRite and Disk Technician are worthy programs; still, no matter what assurances the literature gives, it's important to use low-level maintenance software intelligently rather than in a spirit of blind trust. Some things can and do go wrong, and you will certainly rest easier if you know how to avoid some of the most common pitfalls inherent in these utilities.

One lesson I've learned—the hard way—is never to run this type of utility on a disk drive that's not fully warmed up. Disk Technician Advanced strongly recommends that you load it from a floppy disk, and SpinRite insists on it, so it seems natural to run either utility when you start the machine in the morning. In theory, the drive's temperature compensation should let this work; in practice, it's not a good idea.

Some time ago, I noticed that a particular Seagate 4026 hard disk drive was producing a few (not many) errors. In an attempt to prevent these errors from getting worse, I started work one morning by running Disk Technician Advanced on the drive.

After an hour of nightmarish noises— the kind I dread hearing from a disk drive—I discovered that Disk Technician Advanced's "artificial intelligence" algorithms had marked more than half the drive as bad, including some areas that had never failed before. But after I had let the drive warm up, performed a low-level reformat using the AT Advanced Diagnostics, and transferred the files back to the disk, Disk Technician gave it a clean bill of health. (After that experience, I replaced the drive with a Maxtor, which has worked perfectly to this day.)

Some users have reported that Disk Technician Advanced may be a bit too sensitive—that it sometimes marks larger and larger portions of a hard disk bad as time goes by. I don't know if this is generally true or not, but I suspect that at least some such problems are due to the thermal effects I saw. The lesson: Heed the manual. If you check your disk, do it at the *end* of the day, when your drive is

fully warmed up. It may save you a lot of headaches.

The second point (and I can't emphasize this enough): It's super-important to make *complete* backups before running these utilities. Each has, at one time or another, discovered or created a situation where I've had to restore some or all of my disk from a backup.

Finally, since both of these programs are "working around" the limitations of the operating system software and disk drive controller hardware, they sometimes fail to do the right thing despite their best efforts. For instance, both SpinRite and Disk Technician Advanced perform a bit of "black magic" in an attempt to determine what kind of hard disk drive controller you have—and sometimes they guess incorrectly. One of my machines contains a Western Digital WDM-2 motherboard with an MFM hard disk drive controller built in. The hard disk drive—a Seagate 4051—is divided with FDISK into two partitions: C and D.

SpinRite and Disk Technician Advanced both worked on the C partition and correctly determined that it was an

MFM drive with a 2-to-1 interleave. However, both malfunctioned on partition D, reporting strange interleave factors and/or encoding schemes. (I have reported the problem to both vendors, and each says that it has, or is working on, a fix.) So be sure to keep a watch on these utilities as they go to work; if you see anything unusual, it's a good idea to abort the program immediately.

In the long term, disk maintenance utilities should be part of the operating system rather than special-purpose tools. (The only operating system I'm aware of that performs any disk maintenance as a part of normal operation is Novell's SFT NetWare, which runs only on dedicated network servers.) As multitasking becomes commonplace, it will make sense to run disk maintenance utilities as background tasks; for those of us who run DOS, it would be convenient to have them as TSR programs.

In the meantime, the best policy is to buy a good hard disk drive in the first place—one that parks the heads when it powers down. Then—if your drive and controller permit—use a low-level disk maintenance utility to help keep your hard disk drive in top fighting form. ■

L. Brett Glass is a freelance programmer, author, and hardware designer residing in Palo Alto, California. He can be reached on BIX as "glass."

Your questions and comments are welcome. Write to: Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.

World Radio History

IF MEMORY SERVES...

Here is a library of
memory management
routines that will
help you avoid
a fragmented heap

Memory—silicon memory—is
one of those precious com-
modities in the computer in-
dustry. Like disk-storage
space, it's something you can't get
enough of. Lately, it has been reasonable
to complain that the amounts of memory
you'd like to have would cost too much
money (although compared to times not
far gone, believe me, you're living in
DRAM paradise).

Sometimes, though, you can't get
enough of it not because you can't afford
it, but because some operating system
won't let you get at it. I won't name
names here, but one particular operating
system whose initials are D-O-S has—
through inadequacies inherited from the
CPU it was born on—created a trend of
extension products (hardware and soft-
ware) all designed to "break the 640K-
byte barrier." Even when you can easily
access the memory you need, you discov-
er another need: some governing force
that will keep you from squandering your
windfall. As I mentioned back in my col-
umn on overlays (see "An Overview of
Overlays," December 1988), you'd like
to treat memory as a stretch of unbroken
real estate extending out past the hori-
zon. But that's not the reality—the real-
ity is that you've got to get smart about
managing your memory.

DOS Memory
The DOS memory manager keeps tabs
on the memory that's free for use by the
operating system or whatever programs
happen to be executing at the time. DOS

also tracks the portions of memory that
some routine has laid claim to. If no pro-
grams are being executed, all memory
not used by the operating system, its buf-
fers, or COMMAND.COM is free. (I'm
assuming that no TSR programs are cur-
rently in residence.) When you execute a
program by typing its name at the com-
mand prompt, the operating system allo-
cates however much memory that pro-
gram needs, loads the program into the
allocated block, and transfers control to
that program. (To be strictly accurate,
DOS is not that intelligent all the time.
When you execute a .COM program, for
example, DOS hands *all* the free mem-
ory over to it.)

As the program executes, it may need
a chunk of memory to create strings,
arrays, and so on. It can request this
memory from DOS via a call to INT 21H
function 48H. Later, when the program
has finished with that memory, it can,
using INT 21H function 49H, tell the op-
erating system: "Here, you can have this
back...I don't need it anymore." This
pool from which an executing program
can draw memory blocks (and to which it
can return them) is often referred to as
the "heap." Thus, DOS is put to the task
of tracking all the variable-size chunks
of memory within the heap, where each
chunk is either free or in use. The operat-
ing system's bookkeeping employs a
kind of linked-list structure embedded in
the memory itself, as shown in figure 1.

DOS precedes each memory block
with a 16-byte *control block*. The control
block contains information such as how
big the memory block is and whether it is
in use. DOS can use a control block's
length field to calculate the location of
the next-higher memory block. So when
someone makes a request for memory,
the operating system follows the chain of
control blocks in search of a slice of
memory to satisfy the request. The last
control block in the chain contains a byte
that tells DOS it has reached the top of
the heap (sorry).

Fragmentation and Other Problems
Although DOS's memory management
is entirely adequate for most needs, it is
afflicted with one problem: fragmenta-
tion. This happens when you've allo-
cated and freed a number of memory
blocks in random fashion. Observe fig-
ure 2a, wherein a number of memory
blocks have been used and released. Sup-
pose that a program requests a 16K-byte
block of memory. Although a total of
18K bytes is available, the free memory
is scattered throughout the heap in unus-
ably small pieces.

The solution to this situation is obvi-
ous: Simply move all the used memory
blocks down to the bottom of the heap,
thus allowing the unused fragments to
percolate to the top, where they can be
joined into a single chunk of 18K bytes.
That request for 16K bytes could then be
honored (see figure 2b). This process of
shuffling memory blocks around to com-
bine the unused fragments is called
compaction.

But wait. You can't just slide things
around in memory without telling any-
one. Suppose a subroutine has built a
pointer variable that references the be-
ginning of block A in figure 2a. If the
contents of memory have been shifted to
what's shown in figure 2b, that pointer
will be 4K bytes out of whack. The solu-
tion has spawned a second problem; is
there a second solution?

What's Your Handle?
I'll take a lead from the memory man-
ager of Apple's Macintosh and introduce
the concept of the doubly indirect point-
er—known as the "handle." Don't let the
term "doubly indirect" frighten you. A
pointer is singly indirect: It holds the ad-
dress that you're ultimately interested in.
A handle holds *the address of the address*
that you're ultimately interested in. Put
another way, a handle points to a pointer
(a "master pointer" in Macintosh ter-
minology); in turn, the master pointer
continued

Figure 1: *MS-DOS uses a singly linked list, embedded in the memory itself, to manage the heap.*

points to the actual memory block. Master pointers are typically kept together in a large table that remains fixed in memory. The advantage of this scheme is that, while memory blocks are moved about during compaction, the master pointer stays in place (only its contents change). Consequently, the value of the handle never changes; you can always locate the memory block in spite of compaction.

Handles carry at least one disadvantage. If you overlook the minor additional memory overhead of the table of master pointers, access to memory through a handle requires at least an initial doubly indirect reference. I'm using the word "initial" loosely; a memory block's actual address as retrieved via a pointer is valid only as long as compaction doesn't take place. Once compaction occurs, you have to follow the trail from handle to pointer to see where a memory block has moved to. (The process is called "calculating the *effective address*.")

A sample scenario might go like this: Your program has requested 60K bytes from the memory management system; the request is granted, so your program retrieves the address and begins stuffing data into the 60K-byte block. A subroutine of this process requires a 4K-byte scratch area for string manipulation. The

request for 4K bytes causes compaction to take place. When the subroutine returns, the main routine will have to recalculate the effective address of the 60K-byte block before resuming work.

The repeated recalculation of the address becomes time-consuming. The situation is made worse by the fact that only the memory manager knows if compaction has taken place. So, after every allocation request to the memory manager, you've got to recalculate the effective address of all the blocks you might be using even if the recalculation is not necessary.

Well, there's a fix for this one, too. You simply define a new attribute for blocks in use: the lock attribute. If a block is locked, that tells the compaction algorithm that the contents of that block may not be moved...period. Thus, locking a block means you don't have to reevaluate its address every time there's a chance that memory compaction has taken place. The downside to this is that

continued

Figure 3: (a) *Two locked blocks have divided the usable memory into three partitions.* (b) *After compaction, you are still left with three separate blocks of free memory.*

Figure 2: (a) *If pieces of memory are randomly allocated and freed, the free memory is scattered throughout the heap in unusably small pieces.* (b) *If all the used blocks are moved to the bottom, the free memory percolates to a single large block.*

it makes the compaction routine's job a little tougher. Look at figure 3a. The two locked blocks have divided the usable memory into three partitions. Unused memory cannot "bubble up" past the locked blocks; the result of an attempt at compacting the heap appears in figure 3b. Locking the blocks has fragmented the free memory.

Granted, if you wanted to make your compaction routine *really* smart, it could nose around through the three partitions, trying to find an arrangement that would create the largest contiguous free block. But this would likely consume so much time that the benefits would no longer outweigh the expense, especially if lots of locks were in place. It's best, then, to lock a block only when it's absolutely necessary. (An aside: The Macintosh's memory manager runs into the same problems with locked blocks. *Inside Macintosh* explicitly warns against locking too many blocks and thereby creating the very fragmentation that compaction is supposed to relieve.)

Now! Handles on DOS
I've put together a modest handle-based memory management package for DOS that includes the features I've just described. Specifically, the package combats the fragmentation problem by performing compaction whenever there is a request for memory that this memory manager can't fulfill. (Of course, if the compaction routine still fails to rustle up enough memory, you're simply out of luck.) I've also added block locking for those critical times when your program doesn't want its memory disturbed.

My memory management system keeps track of memory usage via a list of four-word blocks that I refer to as *m-nodes* (with apologies to the Unix folks). These m-nodes are the master pointers in this scheme, and each consists of the following:

• A base pointer, which is the address of the start of a memory block on the heap. This address is in paragraphs (where a paragraph is 16 bytes); the true byte address is the base pointer shifted left by 4 bits.
• A length count, which is the number of paragraphs in the block controlled by this m-node.
• A pointer to the next m-node in the list. (The package actually maintains two lists—one links all m-nodes that control a memory block on the heap, another links the unused m-nodes.)
• A pointer to the previous m-node in the list. This pointer is not used on the free m-nodes list.

Figure 4 shows how all this works. The m-nodes that point to memory blocks on the heap are kept on a doubly linked list. And if you follow this list from start to finish, you'll see that m-nodes reference memory blocks from low to high memory (this is important for the compaction algorithm, which I'll explain in a moment). A doubly linked list allows for rapid insertion and removal; this in turn ensures that the allocation and deallocation processes are nimble. The package keeps m-nodes that don't currently reference a memory block in a separate, singly linked list, the "dead list."

16-Byte Dollops
The most-significant bits in an m-node's next and previous pointers act as status bits. If the most-significant bit of the next pointer is set, the associated block of memory is in use. If the most-significant bit of the previous pointer is set, the associated block of memory is locked.

Figure 4: *In my memory management system, m-nodes that point to memory blocks on the heap are kept on a doubly linked list.*

REMOVABLE. ERASABLE.
OPTICAL. REMARKABLE!

REO-650 single-drive system
$4,995
REO-1300 dual-drive system
$9,990
650/600 Meg optical cartridges:
$250 each
Interface kits for Macintosh.
Sun. HP. IBM-AT, XT, PS/2
and compatibles:
$995 each

Pinnacle Micro is the world's leader in removable, erasable, optical storage systems for personal and workstation computers.

The REO-650 is a single drive, SCSI system. The REO-1300 is a dual-drive system. Up to 1.3 gigabytes of facts and figures at your fingertips!

Each optical cartridge holds up to 650 megabytes of precious memories. Graphics. CAD files. Databases. Images.

You can write, erase and rewrite data... file by file... a million times or more.

Interface kits give you the flexibility to use Pinnacle drives with the computer of your choice: Macintosh. Sun Microsystems. Hewlett Packard. IBM-AT, XT, PS/2 and compatibles.

Software provided with each kit supports a host of working environments: UNIX, Xenix, Novell Netware 2.1, and many more.

Removable, erasable, optical. The new standard in mass storage. Available now from Pinnacle Micro.

Call today for the location of your sales representative or authorized dealer.

✆ (800) 553-7070

PINNACLE MICRO

15265 Alton Parkway • Irvine, CA 92718 • In CA (714) 727-3300 • FAX (714) 727-1913

THE ERASABLE OPTICAL STORAGE COMPANY ™

Circle 295 on Reader Service Card (DEALERS: 296)

AUGUST 1989 • B Y T E   283

As indicated in the features list, all memory referenced by m-nodes appears on 16-byte boundaries. Actually, the trend goes farther than that: You request memory in paragraph-size slices. This is an...ahem...a feature of the Intel 80x86 architecture. Since segment registers must be aligned to paragraph boundaries, the entire memory management package is much simpler if it treats memory in dollops of 16 bytes. This adds the benefit of speeding up the compaction algorithm: Because the package allocates memory in an even number of bytes, it can use MOVSW instructions (16 bits at a time) rather than MOVSB instructions (8 bits at a time) to shuffle the contents of memory blocks around.

The algorithm for the allocation of a

Listing 1: *The pseudocode for compaction in my memory management system. The algorithm calls for all unlocked blocks to move to the lowest memory possible, thus causing free blocks to congregate near the top.*

```
COMPACT:
{ COMPACT makes use of the following externally defined routines:
{ RELEASE(m_NODE) removes an m_node from the doubly linked list of m_nodes
{ referencing memory and places it on the available list.
{ GET_MNODE() fetches a new m_node from the available list.
{ MOVE_MEMORY(SADDR,DADDR,LENGTH) moves LENGTH bytes from address SADDR to
{ DADDR. The routine handles the possibility of overlapping regions.

CURRENT_BASE := 0;
CURRENT_LENGTH := 0;
SAVED_MNODE := 0;
DIRTY := 0;
PARAS_FREED := 0;
PREVIOUS_MNODE := 0;

{ Start with first m-node in list.
{ INUSE_LIST points to the head of the doubly linked list of m_nodes that
{ actually reference memory.
CURRENT_MNODE:=INUSE_LIST;

REPEAT
 IF high bit of CURRENT_MNODE's NEXT link
     is set THEN
 BEGIN
 { This block is in use. See if it's locked.
   IF high bit of CURRENT_MNODE's PREVIOUS link
     is set THEN
   BEGIN
   { This block is locked. See if DIRTY is
   { set. If so, the contents of CURRENT_BASE
   { and CURRENT_LENGTH must be written into
   { an m-node.
    IF DIRTY=1 THEN
    BEGIN
    IF SAVED_MNODE< >0 THEN RELEASE(SAVED_MNODE);
    NEW_MNODE := GET_MNODE();
    NEW_MNODE's BASE field := CURRENT_BASE;
    NEW_MNODE's LENGTH field := CURRENT_LENGTH;
    CURRENT_BASE := 0; CURRENT_LENGTH := 0;
    { NEW_MNODE must now be attached to the
    { list preceding CURRENT_MNODE.
    NEW_MNODE's NEXT link := CURRENT_MNODE;
    NEW_MNODE's PREVIOUS link := CURRENT_MNODE's
        PREVIOUS link;
    CURRENT_MNODE's PREVIOUS link := NEW_MNODE;
    TEMP_MNODE := NEW_MNODE's PREVIOUS link;
    IF TEMP_MNODE = 0 THEN
    { NEW_MNODE is the head of the list.
      INUSE_LIST := NEW_MNODE;
    ELSE
      TEMP_MNODE's NEXT link := NEW_MNODE;
    DIRTY := 0;
```

Graphics Choices Driving You Crazy?

Indecision about graphics platforms is costing software developers money. At Media Cybernetics, we maximize our development productivity and minimize our risks by basing our application programs on HALO. We used HALO to develop popular applications like Dr. HALO and Image-Pro.

Making the HALO choice gives you DOS or OS/2 support today with a clear source code compatible path to Presentation Manager tomorrow. It also gives you a shortcut to building a user interface, the hardest part of application development.

The HALO development environment offers a comprehensive, reliable series of libraries and tools that are professionally documented and supported. Even more important, HALO is here to stay. We've been supporting software developers since 1982 and we are proud that 70,000 programmers and nearly 400 Independent Software Venders (ISVs) are part of the HALO family.

Media Cybernetics now offers three versions of the HALO Graphics Toolkit:

HALO
- The premier graphics toolkit
- 200 graphics subroutines
- Supports over 180 input, output, and graphics devices
- Device support for scanners, high resolution imaging and VESA boards
- 18 compiler interfaces

The HALO Window Toolkit
- Event-driven environment for graphical user interface design
- Command bars, pull-down menus, scroll bars, radio buttons and icons
- Available for Microsoft C
- Intelligent memory management
- Full HALO capability

HALO for OS/2
- Source code compatible with HALO
- HALO speed, power and capability
- Multi-threaded, multi-tasking
- Dynamic Link Library
- Supports background mode graphics

HALO™

Making these decisions isn't easy , but HALO has the best track record in the industry. We believe that a strong contemporary graphics toolkit is essential to your success.

If you are concerned about the long term consequences of your graphics choices, call 1-800-992-HALO

MEDIA CYBERNETICS™
Imaging and Graphics Technology

8484 Georgia Avenue
Silver Spring, Maryland 20910
301/495-3305
301/495-5964 FAX

World Radio History

memory block is simple—so simple, in fact, that I won't bother with the pseudocode. I used the "first-fit" technique. This simply means that when your program requests a block of memory, the memory manager begins scanning at the m-node at the head of the chain and stops at the first m-node that references a free block large enough to fit your request. If the block has more memory than you need, the system pares off the amount you've requested and creates a new m-node to reference the leftover portion. Finally, it passes back to your program an offset to the m-node associated with your newly claimed block of memory. This offset is the handle.

COMPACT to the Rescue

Things get squirrelly when the memory manager hits the top of the list and hasn't found a block large enough to meet your request. That's when it summons the compaction routine (see listing 1 for the pseudocode). The COMPACT routine skims through the list in bottom-to-top fashion, moving all unlocked blocks into the lowest memory possible, thus causing free blocks to congregate near the top. Any m-nodes that reference adjacent free blocks are merged into a single m-node referencing the lump sum. Finally, COMPACT returns the number of paragraphs in the largest free memory block that it was able to put together. The allocation routine can then quickly determine whether COMPACT was able to free a large enough piece of memory.

I've also added routines for locking and unlocking blocks, plus a routine for deallocating memory blocks. The deallocation routine is intelligent enough to do some defragmenting of its own: If it discovers that a deallocated block is adjacent to blocks that are already free, it combines the adjacent blocks into a single, larger block of free memory referenced by a single m-node. Again, these routines are rudimentary enough that I won't expend space on their pseudocode.

Hot-Rodding the System

As usual, I feel the urge to suggest additions to this package that might make it suitable for your particular application. First, a "first-fit" algorithm is by far the simplest to code (so now you know why I picked it). Its alternative is the "best-fit" algorithm, which seeks to reduce fragmentation by attempting to locate the block that most closely matches the allocation request. This seems a reasonable tack to take, since it improves your chances of finding a free block that fits

PRODUCT SHOWCASE

- **BUYER'S MART**
- **BYTE BITS**
- **PRODUCT SPOTS**
- **MICRO PRODUCT CENTER**

ILLUSTRATION: JULIE E. MURPHREE © 1989

World Radio History

World Radio History

THE BUYER'S MART

World Radio History

PUBLIC DOMAIN

FREE CATALOG
$1 IBM SOFTWARE
For your free 32-page Master Edition catalog featuring the best of IBM Shareware from just $1 each, call or write today!

1-800-338-2118
SOFSOURCE
Box 828, East Lansing, MI 48826

Inquiry 671.

TOP 10 IBM SOFTWARE - FREE
10 disk set — PC Write, PC Calc+ TreeView, DanCad, FileExpress, Moraff's Revenge, HGCIBM, Novatron, Home Base, Spacewars plus our 1600+ disk catalog
We accept Visa/MC/Amex

PAY ONLY $5.00 shipping/handling
limited introductory offer

INTERNATIONAL SOFTWARE LIBRARY
511 Encinitas Blvd • Ste 104 • Encinitas CA 92024

| order today | toll free | 800 669-2669 |

Inquiry 672.

FREE SOFTWARE CATALOG
Low as $1.20/disk
Over 1000 quality IBM software
On 5.25" and 3.5" format
From outside U.S.A., except Canada, please send US $2.00 refundable with order.
For fast service, write to

SOFTSHOPPE
P.O. BOX 709, Ann Arbor, MI 48105
(313) 763-8721

Inquiry 673.

YOUR SALES MESSAGE
about the special computer product or service that you provide belongs in print.

THE BUYER'S MART
can help you reach computer professionals and produce valuable inquiries for your company!

Call Brian Higgins for more information

603-924-3754

Inquiry 674.

---

16 and 32 BIT MICROS
EDUCATIONAL TRAINING SYSTEMS in a notebook with power supply — for the Motorola 68000/68020/68881, TMS32010 DSP, Intel 8086/8087. A/D-D/A Convertors, cross assemblers, serial interfaces with software, complete systems, documentation, schematic, operating system, cables, Starting Prices — $230.00

Phone URDA, Inc.
1-800-338-0517

Inquiry 675.

SECURITY

EVERLOCK COPY PROTECTION
★ Thwarts ALL Bit-copy Software
★ Protect any COM/EXE w/o Source changes
★ Shut down Debug Tracing & Disassemblers
★ Install to Floppy, Hard Disk, or LAN
★ Remotely reset Program Install-Count, Expire-Date or #Executes
★ No damaged media or I/O plugs
For IBM and clones. $195 & up. Free Info.

Az-Tech Software, Inc.
305 East Franklin, Richmond, MO 64085
(800) 227-0644 (816) 776-2700
FAX: (816) 776-8398

Inquiry 676.

SECURITY

THE ULTIMATE COPY PROTECTION
• Completely Menu Driven	
• Defeats all Hardware/Software Copiers	**Quite**
• No Source Code Changes	**Simply**
• Multiple Layering	**The Best**
• No Damaged Media	**Way To**
• Full Hard Disk Support	**Protect**
• Unlimited Metering	**Your Valuable**
• *FREE* Demo Disk	**Software Investment**

STOPCOPY™ $325⁰⁰   STOPCOPY PLUS™ $450⁰⁰
BBI COMPUTER SYSTEMS® (301) 871-1094
14105 Heritage La., Silver Spring, MD 20906   FAX: (301) 460-7545

Inquiry 677.

BIT-LOCK® SECURITY
Piracy SURVIVAL 5 YEARS proves effectiveness of powerful multilayered security. Rapid decryption algorithms. Reliable/small port-transparent security device. PARALLEL or SERIAL port. Complemented by economical KEY-LOK™ and multifeatured COMPU-LOCK™ including countdown, timeout, data encryption, and multiproduct protection.

MICROCOMPUTER APPLICATIONS
3167 E. Otero Circle, Littleton, CO 80122
(303) 922-6410/770-1863

Inquiry 678.

PC Security "Password"
With All the Computer Talk, PASSWORD is the Perfect Security Lock.

Password is a software program providing security for your PC. Password is Easy to understand and Simple to install, requires no reformatting. The boot limit option secures your hard disk. Password provides for up to 100 users with the supervisor controlling access to protected directories. Password is menu-driven with pop-up windows and help screens. The program provides an audit trail of users, and a screen blanking feature.
PASSWORD $99.00 US   Visa, M/C, Amex
Nasdec International Inc.
2704-85 Garry Street, Winnipeg MB Canada R3C 4J5
PH: (204) 956-2798   FAX (204) 943-3702

Inquiry 679.

COPY PROTECTION
The world's leading software manufacturers depend on Softguard copy protection systems. Your FREE DISKETTE introduces you to SuperLock™—invisible copy protection for IBM-PC (and compatibles) and Macintosh.
• Hard disk support   • No source code changes
• Customized versions   • LAN support
• New upgrades available
(408) 773-9680
SOFTGUARD SYSTEMS, INC.
710 Lakeway, Suite 200, Sunnyvale, CA 94086
FAX (408) 773-1405

Inquiry 680.

SOFTWARE/ACCOUNTING

PC TIME CLOCK
AutoTime is an Employee Management System that allows you to turn any PC into an Electronic Time Clock. AutoTime provides Time & Attendance, Job Costing, Payroll Interface, and Labor Distribution reporting. Network compatible. Prices start at $495. Other Business Products: Network FAX, Absence Call-In, db-EDI.

Chase Technologies
1617 Kingman Ave., San Jose, CA 95128
(408) 998-2917

Inquiry 681.

---

dBASE BUSINESS TOOLS
• GENERAL LEDGER	• PURCH ORD/INVNTORY
• ORDER ENTRY	• ACCOUNTS RECVABLE
• JOB COSTING	• JOB ESTIMATING
• BILL OF MATLS	• SALES ANALYSIS
• PAYROLL	• ACCOUNTS PAYABLE

$99 ea. + S&H
dATAMAR SYSTEMS Cred. Card-Check-COD
4876-B Santa Monica Ave.
San Diego, CA 92107   **(619) 223-3344**

Inquiry 682.

SOFTWARE/BASIC

Save time, money, & brain cells!
QuickWindows Advanced user-interface library supports windows, menus, dialog boxes, help, mouse, and more. For text and graphics modes thru VGA. Fast and intuitive!
QuickComm communications library supports up to 16 comm ports, Hayes modems, xmodem and ymodem file transfer, and more. All interrupt driven. Easy to use!
Each library written in assembly and comes with complete manual. $139 each. For Microsoft QuickBasic or BASCOM.
Software Interphase, Inc.
5 Bradley Street, Suite 106, Providence, RI 02908
(401) 274-5465 Call now for FREE terminal/demo program.
Fax # 401-272-1273

Inquiry 683.

SOFTWARE/BUSINESS

DATA ENTRY SYSTEM
Heads-down data entry with two-pass verification for the IBM PS/2-PC/XT/AT & compatibles. Features include: Auto dup/skip, verify bypass, range checks, table lookups, a complete edit language. Fully menu driven. Price $395

Call for our free 30-day trial period.
COMPUTER KEYS
21929 Makah Rd., Woodway, WA 98020
Tel: 206/776-6443 USA: 800/356-0203 Fax: 206/776-7210

Inquiry 684. *(Note: printed as Inquiry 684)*

MILP88—MIXED-INTEGER LP
A general-purpose system for solving mixed-integer linear programs with up to 800 constraints and 4000 general integer or noninteger variables. Build MILP88 into your own programs with compiled Turbo Pascal units. MILP88 reads/writes Lotus worksheets. Use 1-2-3/Symphony as a matrix generator or post processor. Other features include interactive and batch operation, spreadsheet LP display and editing, an equation processor, problem/branch list storage, file I/O, download/reload, report generator, and sensitivity analysis. $149 with manual and 8087 support. $299 with Turbo Pascal units.
Eastern Software Products, Inc.
P.O. Box 15328, Alexandria, VA 22309
(703) 360-7600

Inquiry 684.

LOW COST/HIGH QUALITY
Established, Powerful, Complete, Business Management Software Systems. Point-of-Sale/Inventory Control "SALES-PRO", Service and Repair, Video/Rental Store Management, Church Management, Accounting and Many more starting at $39. For IBM PC Compatibles and the Atari ST.
HI-TECH ADVISERS
P.O. BOX 7524, Winter Haven, FL 33883-7524
1-800-882-4310   Florida (813) 294-1885

Inquiry 685.

SOFTWARE/ENGINEERING

Affordable Engineering Software
FREE APPLICATION GUIDE & CATALOG
Circuit Analysis • Root Locus • Thermal Analysis • Plotter Drivers • Engineering Graphics • Signal Processing • Active/Passive Filter Design • Transfer Function/FFT Analysis • Logic Simulation • Microstrip Design • PC/MSDOS • Macintosh • VISA/MC

BV Engineering Professional Software
2023 Chicago Ave., Suite B-13, Riverside, CA 92507
(714) 781-0252

Inquiry 686.

Analog Circuit Simulation
• Schematic Entry	Intusoft has a complete PC based system including everything from schematic entry through SPICE simulation using extended memory to comprehensive interactive post processing. Starting at $95 for IsSpice, the complete system sells for just $790.
• SPICE Simulator	
• Model Libraries	
• Monte Carlo Analysis	
• Parameter Sweeps	
• Plotting/Graphics Output	

Intusoft
The leader in low cost, full featured CAE software
P.O. Box 6607, San Pedro, CA 90734
(213) 833-0710 FAX (213) 831-3956

Inquiry 687.

World Radio History

YOUR MEMORY UPGRADE SPECIALISTS

WE ACCEPT AMERICAN EXPRESS

Special of the Month
386 MOTHERBOARD - Uses CHIP'S TECHNOLOGY Chip Set
386-16 w0K   **$995**   386-20 w/0K **$1195**
Upgrade your memory with our low priced SIMM'S

Unitex IC DIVISION

ESTABLISHED 1976

MEMORY EXPANSION BOARDS

SIMM SIPP MODULES

1 MG X 9- for IBM TYPES
1 MG X 9-120 NS	$175
1 MG X 9-100 NS	$185
1 MG X 9-80NS	$219
1MG X9-70NS	$289

256 X 9- for IBM TYPES
256 X 9-120NS	$59
256 X 9-100NS	$75
256 X 9-80NS	$89
256 X 9-60NS	$109

APPLE SIMM MODULES
1 MG X 8 - for APPLE
1 MG X 8-120NS	$169
1 MG X 8-100NS	$179

256 X 8 - for APPLE
256 X 8-120NS	$59
256 X 8-100NS	$69

PS-2 SIMM'S
256 X 9 (FOR PS2)
256 X 9-120NS	$79
256 X 9-100NS	$95

MODEL 30-286
1 mg x 9-100	$229

PS-2MODEL 70&80 SIMM
1 MG X 9-100NS	$249
1 MG X 9-80NS	$299
2 MG X 9-80NS	$679

VIDEO RAM FOR VGA CARDS
64 x 4(150NS)	$7
64 X 4(120ns)	$10
64 X 4(100ns)	$13

DIP ZIP SOJ SIMM SIPP PS-2 SIMM MATH-CO

D RAM
1 MG X 1
1 MG X 1-120 NS	$14
1 MG X 1-100 NS	$15
1 MG X 1-70 NS	$20
1 MG X 1-80NS	$18

256 X 1
256 X 1-150 NS	$5.00
256 X 1-120 NS	$5.50
256 X 1-100NS	$6.00
256 X1-80NS	$7.00
256 X 9-70NS	$7.50
256 X 1-50NS	$8.00

256 X 4
256 X 4-120 NS	$15
256 X 4-100 NS	$17
256 X 4-80NS	$21

64 X 1
4164-150NS	$1.40
4164-120NS	$2.10
4164-100NS	$2.40

64 X 4
4464-150 NS	$5
4464-120 NS	$6
4464-100NS	$7
4464-80NS	$9

256 X 4 STATIC COL
514258-10	$25

256 X 1 STATIC COL
51258-10	$7.50
51258-80	$9.00
51258-70	$9.50

Compaq 386 Memory
1 mg-expansion	$399
4 mg	$1300

Daughte BD Mod 70-80
1 mg (6450375)	$429
2 mg (6450379)	$929

MATH CO-PRO
8087-3(5MHZ)	$88
8087-2(8MHZ)	$118
8087-1	$158
80287-6	$128
80287-8	$195
80287-10	$218
80C287-12	$299
80387-16	$320
80387-20	$375
80387-25	$470
80387-33	$640

CPU CHIPS
8088	$5.00
80286-8LCC	$49
80286-10LCC	$59
80286-12LCC	$69
80386-16	$180
80386-20	$280
80386-25	$439
V-20 (8MHZ)	$7.50
V-20 (10MHZ)	$10.00

CACHE
80385	CALL

MAC CHIPS
68881-16 Math Co	$160
68881-20 Math Co	$170
68882-16	$140
68882-20	$169

RAM TESTER
only $139
Tests the following memory chips:
64K X 1	256 X 1
64K X 4	256 X 4
4M X 1	1M X 4

The RTI can help any-one come through the frustrating process of identifying bad (or good) D-RAM chips.

ORCHID
RAMQUEST IIZ -Up to 2MB with 0 K for PSZ MOD 50, 50Z & 60 - EMS and OS/2 Compatable- Uses 1MG D-RAM ... **$209**

RAMQUEST EXTRA -Up to 8MB - for PS2 MOD 50, 60,70 & 80, EMS & OS/2 Compatable - Has 2 serial ports - Uses 256K of 1MB SIMM'S ... **$319**

RAMQUEST EXTRA-16/32 - up to 8MB - For PS2 MOD 50, 50Z, 60, 70 & 80- Fully Supports 16 BIT & 32 BOIT- 1 serial & 1 par - EMS & OS/2 Compatable- Uses 256K of 1MB SIMM s ... **$319**

RAMQUEST XT/AT -Up to 8 MB - For XT, AT, PS2 MOD25,30 - 8 BIT or 16 BIT - Uses 256K or 1MB SIMM'S ... **$259**

TINY TURBO 286-High speed Half-slot Accelerator for PC/XT- 3 times faster with an 8 MHZ 286 CP -80287 9KT ... **$239**

JET 386 - Hyper speed Accelerator for AT- 3 times faster with 16 MHZ 386 CPU - 80387 9KT ... **$895**

BOCA RESEARCH
For PS2
BOCARAM 30- with 0K RAM ... **$129**
  Expands to 2MG - Uses 1MG X 1 D-RAMS
BOCARAM 50/60 with 0K RAM ... **$179**
  Expands to 4 MG with Software
  Uses 1 MG X 1 D-RAMS
BOCARAM MCA 50Z with 0K RAM ... **$189**
  Expands to 2 MG - Uses 1 MG X 1 D-RAMS
BOCA MCA Parallel Card ... **$99**
BOCA MCA Serial/PAR ... **$169**

EVEREX
MINI-MAGIC - #EV138 - 576K Memory Card ... **$59**
  for PC & AT with 0K - Uses 64K & 256K D-RAM
RAMII 4000 - #EV-136 - 4MB EMS ... **$249**
  Extended Memory card with 0 K-Uses 1 MGD-RAMS

IBM
1497259 - For PS-2 MOD 50/60 ... **$439**
  with 0K Expands to 8MB
  Uses 256K SIMMS (IBM only)
6450605 - For PS-2MOD 70/80 ... **$1299**
  with 2 MG Expands to 8 MB
  Uses 2MG SIMMS (IBM only)
6450203 - For AT - Has 512K RAM ... **$129**

IDEA ASSOCAITES
IDEA max 80 - for PS2 MOD 70/80 ... Call
  Expands to 8MB
  Uses 256K or 1MG SIMMS
IDEAsupermax/MC - for PS2 MOD 50/60 ... Call
  Expands to 8MB/2 Extended memory/2SER/EMS
  Uses 256 or 1MG SIMMS
IDEAmax/MC - for PS2 MOD 50/60 & 80 ... Call
  Expands to 12MB Extended memory with software
  Uses 1MG SIMMS

UNITEX
3MG Multifunction - for AT ... **$149**
  Expands to 3MG - has SER/PAR PORT
  Uses 1MG D-RAMS
384 Multi-function Card for PC/XT ... **$89**
  Expands to 384K- has SER/PAR/CLK/Game port
  Uses 64K or 256K D-RAMS

VIDEO ADAPTER

ATI
EGA Wonder 800 ... **$239**
  Supports EGA, MDA, CGA & Hercules
VIP VGA - 800 X 560 - ... **$275**
  Supports VGA, CGA, MDA & Hercules

VIDEO 7
Fastwrite VGA ... **$349**
  256 D-RAM, 800 X 600, 640 X 480
  Supports VGA/EGA/MDA/CGA & Hercules
V-RAM VGA ... **$519**
  256 K V-RAM, 1024 X 768, 800 X 600
VEGA Deluxe ... **$239**
  640 X 480 Multi-sync
VEGA Pro ... **$319**

ORCHID
DESIGNER 800 VGA ... **$229**
  800 X 600 - 16 Colors
PRODESIGNER VGA ... **$319**
  Supports 1024 X 768 - 16 Colors
PRO DESIGNER VGA PLUS ... **$399**
  Same as Prodesigner - Has downloadable fonts

UNITEX
MONOCHROME GRAPHICS CARD ... **$41**
  with par port • MDA/CGA/Hercules
COLORGRAPHICS CARD ... **$41**
  RGB Color with Par Port • EGA/MDA/CGA/Hercules
EGA CARD ... **$149**
  640 X 480, 16 Colors, EGA/MDA/CGA/Hercules
VGA CARD ... **$199**
  600 X 800, 16 Colors • VGA/EGA/MDA/CGA/

TERMS AND CONDITIONS
No Surcharge for MC/VISA   Terms: MC • VISA • COD • CASH • AMEX add 4%
Purchase Orders from qualified firms 20% restocking fee on non-defective returns
Prices Subject To Change

2852 F Walnut - Tustin, CA 92680
Phone: 714/730-5232 • FAX#: 714/838-8593
Customer Service #: 714/730-9527

TOLL FREE OUTSIDE CA: 1/800/533-0055

99% OF ALL ORDERS SHIPPED SAME DAY

Circle 263 on Reader Service Card

World Radio History

Circle 261 on Reader Service Card

Circle 66 on Reader Service Card

Oh, No!!

They squashed the monitor!

Its the first laptop with a screen as readable as a CRT...

Revolutional technology - less than an inch thick, yet looks like a hi-res amber monitor...

Viewable from all angles.

- FAST 80286 CPU
- UPGRADEABLE TO 80386
- FAST 40MB HARD DISK

For more information contact:

KISS Computers 1-800-438-5477

1-414-652-5477
Dealer inquiries welcome

World Radio History

World Radio History

World Radio History

World Radio History

Buy with
Confidence

In an effort to make your telephone purchasing a more successful and pleasurable activity, The Microcomputer Marketing Council of the Direct Marketing Association, Inc. offers this advice, "A knowledgeable buyer will be a successful buyer." These are specific facts you should know about the prospective seller before placing an order:

Ask These Important Questions

- *How long has the company been in business?*
- *Does the company offer technical assistance?*
- *Is there a service facility?*
- *Are manufacturer's warranties handled through the company?*
- *Does the seller have formal return and refund policies?*
- *Is there an additional charge for use of credit cards?*
- *Are credit card charges held until time of shipment?*
- *What are shipping costs for items ordered?*

Reputable computer dealers will answer all these questions to your satisfaction. Don't settle for less when buying your computer hardware, software, peripherals and supplies.

Purchasing Guidelines

- *State as completely and accurately as you can what merchandise you want including brand name, model number, catalog number.*
- *Establish that the item is in stock and confirm shipping date.*
- *Confirm that the price is as advertised.*
- *Obtain an order number and identification of the sales representative.*
- *Make a record of your order, noting exact price including shipping, date of order, promised shipping date and order number.*

If you ever have a problem, remember to deal first with the seller. If you cannot resolve the problem, write to MAIL ORDER ACTION LINE, c/o DMA, 6 E. 43rd St., New York, NY 10017.

This message is brought to you by:

the MICROCOMPUTER MARKETING COUNCIL of the Direct Marketing Association, Inc.
6 E. 43rd St.,
New York, NY 10017

MMC
MICROCOMPUTER
MARKETING COUNCIL
of the Direct Marketing Association, Inc.

© Direct Marketing Association, Inc. 1988

World Radio History

AUGUST 1989 • BYTE 325

World Radio History

World Radio History

Editorial Index by Company

Index of companies covered in articles, columns, or news stories in this issue
Each reference is to the first page of the article or section in which the company name appears

COMING UP IN BYTE

PRODUCTS IN PERSPECTIVE:

First Impressions for September should include our evaluation of Lotus 1-2-3 version 3, IBM's OfficeVision/2, and Apricot's newest entry, the first 80486-based computer.

The **Product Focus** will be on multiuser/multitasking operating systems. Touted as low-cost alternatives to LANs, operating systems like Digital Research's Concurrent DOS, IGC's VM/386, and The Software Link's PC-MOS/386 connect multiple computers or terminals via the serial port. How effective are they? Find out in September.

Reviews we hope to publish in September include a look at two ends of the IBM PC clone spectrum. ALR's MicroFlex 7000 looks to be a hot-performing Micro Channel architecture machine—one of a handful. We also evaluate a workhorse 80286 unit: the AST Bravo.

Sysgen has a new removable hard disk drive—not an enclosed unit, but an actual hard disk platter you can pull out to transfer data from machine to machine. Brightbill-Roberts enters the MS-DOS hypertext fray with HyperPAD. And Arriba from Good Software is an information manager that looks promising.

In our Reviewer's Notebook, we will be looking at a number of products, including the The Complete PC's Complete Page Scanner and International Software's PixC.

IN DEPTH:

The focus in September will be on **databases**—the different types of structures they can follow and the various languages designed to interface with them.

We will discuss distributed databases and database servers, the ever-popular relational database, and look to the future with possible object-oriented databases. We will also look at the trends in microcomputer database direction. All this and SQL, too.

FEATURES:

Is IBM's Micro Channel the wave of the future, or a dead end? Will the Extended Industry Standard Architecture, supported by a group of IBM's competitors, attract users by letting them continue to use their PC AT cards? What about the NuBus? And what is Futurebus?

Our lead feature story in September will be a look at the **current battle over bus architectures**, by an author who is eminently qualified to tackle the subject. George P. White, president of Corollary (a spin-off of Texas Instruments), headed the committee that developed NuBus, and he has followed the architecture scene for years. Find out where buses are headed in the days to come.

Our Hands On columns will have L. Brett Glass discussing laptop technology in **Under the Hood** and Tom Thompson getting inside the Macintosh color lookup table in **Some Assembly Required**.

Our **Expert Advice** columnists include Jerry Pournelle in Computing at Chaos Manor, David Fiedler in Unix /bin, Wayne Rash Jr. in Down to Business, Don Crabb in Macinations, Mark Minasi in OS/2 Notebook, and Brock Meeks in NetWorks.

READER SERVICE

To get further information on the products advertised in BYTE, fill out the reader service card by circling the numbers on the card that correspond to the inquiry number listed with the advertiser. This index is provided as an additional service by the publisher, who assumes no liability for errors or omissions.

* Correspond directly with company.

Alphabetical Index to Advertisers

World Radio History

READER SERVICE

READER SERVICE

To get further information on the products advertised in BYTE, fill out the reader service card by circling the numbers on the card that correspond to the inquiry number listed with the advertiser. This index is provided as an additional service by the publisher, who assumes no liability for errors or omissions.

* Correspond directly with company.

Index to Advertisers by Product Category

READER
SERVICE

the requested amount perfectly. It has been pointed out, however (Knuth, *The Art of Computer Programming*), that a best-fit algorithm may actually make the job of compaction more difficult by promoting the proliferation of too-small-to-be-useful-to-anyone fragments.

Another improvement is possible through the fact that the m-nodes are kept on a doubly linked list. It works like this: Suppose a request is made for a 16K-byte block of memory and the allocation routine finds that it cannot fulfill the request. The allocation routine calls COMPACT, and when COMPACT returns, it tells the allocation routine that the largest single free block it was able to create is 20K bytes. Now the allocation routine knows the request can be satisfied; but instead of starting at the bottom of the list of m-nodes (looking for a first-fit), the routine starts at the top of the list and works down. The idea is that since compaction has moved free blocks toward higher memory, you should begin looking at that end.

Such a Deal!
The memory management package I've described here is available in source code form. I've added interface routines for Turbo C, but I've no doubt the package can be modified to coexist with other languages. You should find the memory manager useful in putting together applications for text processing, AI, or areas where complex data structures are continually created, merged, and destroyed.

For those of you interested in immediate gratification, I'll put it this way: If you're up against the wall because your program has memory-allocation problems and you suspect that fragmentation is the culprit, this package might take care of it. In any case, you'll be making your own little assault on the latest string of virtual memory management software appearing on the market. ∎

Editor's note: The full text of the 8086 assembly language source code is available in a variety of formats. See page 5 for details.

Rick Grehan is the director of the BYTE Lab. He has a B.S. in physics and applied mathematics and an M.S. in computer science/mathematics from Memphis State University. He can be reached on BIX as "rick_g."

Your questions and comments are welcome. Write to: Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.

```
END
ELSE
BEGIN
{ This block is not locked.
{ Do we need to move it?
 IF (CURRENT_BASE < > 0) OR (CURRENT_LENGTH < > 0) THEN
 BEGIN
  SOURCE := CURRENT_MNODE's BASE field;
  DESTINATION := CURRENT_BASE;
  LENGTH := CURRENT_MNODE'S LENGTH field;
  MOVE_MEMORY(SOURCE, DESTINATION, LENGTH);
  CURRENT_MNODE'S base field := CURRENT_BASE;
  CURRENT_BASE := CURRENT_BASE + LENGTH;
  DIRTY := 1;
  END
 END
END
ELSE
BEGIN
{ This m-node references a free block.
{ If there is anything in CURRENT_BASE or
{ CURRENT_LENGTH, release the m-node saved
{ in SAVED_MNODE and CURRENT_MNODE becomes
{ the new SAVED_MNODE.
 IF (CURRENT_BASE< >0) OR (CURRENT_LENGTH< >0) THEN
  BEGIN
   RELEASE(SAVED_MNODE);
   CURRENT_MNODE's BASE field := CURRENT_BASE;
   CURRENT_LENGTH := CURRENT_LENGTH + CURRENT_MNODE's length field;
   CURRENT_MNODE's LENGTH field := CURRENT_LENGTH;
   IF CURRENT_LENGTH > PARAS_FREED THEN
       PARAS_FREED := CURRENT_LENGTH;
  END
  ELSE
  { If there's nothing in CURRENT_BASE or
  { CURRENT_LENGTH, then this is the first free
  { memory block of this partition.
  BEGIN
   CURRENT_LENGTH := CURRENT_MNODES's LENGTH field;
   IF CURRENT_LENGTH > PARAS_FREED THEN
       PARAS_FREED := CURRENT_LENGTH;
   CURRENT_BASE := CURRENT_MNODE's BASE field;
  END
  SAVED_MNODE := CURRENT_MNODE;
  DIRTY := 0;
 END

 PREVIOUS_MNODE := CURRENT_MNODE;
 CURRENT_MNODE := CURRENT_MNODE's NEXT field;

{ This is the end of the REPEAT loop.
UNTIL (CURRENT_MNODE = 0);

{ Pick up any straggling free memory.
IF DIRTY=1 THEN
BEGIN
 RELEASE(SAVED_MNODE);
 CURRENT_MNODE := GET_MNODE();
 CURRENT_MNODE's BASE field := CURRENT_BASE;
 CURRENT_MNODE's LENGTH field := CURRENT_LENGTH;
 CURRENT_MNODE's NEXT field := 0;
 CURRENT_MNODE's PREVIOUS field := PREVIOUS_MNODE;
 PREVIOUS_MNODE's NEXT field := CURRENT_MNODE;
END

{ Return the size in paragraphs of the largest
{ block freed.
RETURN(PARAS_FREED);
```

BIX CALENDAR

AUGUST

A U G ... T

TUESDAY, 8/1, 9 PM EST. ''Favorite Computing Tricks''

What clever things do you do that others would find useful? How do you set up your startup and boot? How do you run programs? And how about backup tricks? If you're interested in what others do—or if you have a trick or two to share—drop in the ibm.pc conference. (join ibm.pc/cbix)

THURSDAY, 8/3, 8:30–9:30 PM EST. ''What's so special about Ada?''

Randy Brukardt and Dan Stock of R.R. Software continue their discussion of the Ada language. Which of its features make Ada so useful for projects with many programmers? Is Ada too big? Will the next Ada standard be even bigger? (join janus.ada/cbix)

THURSDAY, 8/10, 6 PM EST. ''Live, from MacWorld Expo in Boston . . .''

Join Macintosh Exchange Editor Larry Loeb and his special guests as they discuss the news coming out of the MacWorld Expo. (join mac.hack/cbix)

All-Month Conferences

You-heard-it-here-first Department—The Microbytes staff will be filing news reports from SIGGRAPH, the premiere computer graphics conference, in Boston, Jul. 31–Aug. 4. Next, it's off to MacWorld Expo in Boston, Aug. 10–12. Later this month, look for reports from UniForum, also in Boston. (join microbytes; join microbytes.sw; join microbytes.hw)

neural.nets conference—Looking for neural-network simulators? See the topic "source," which has the source code to various neural-network simulators in C, Common Lisp, and Smalltalk. (join neural.nets)

mac.hack conference—"Getting ready for MacWorld Expo." Last-minute jitters, the latest rumors, and early product introductions all heighten the tension as BIXen get ready for the MacWorld Expo in Boston on Aug. 10–12. After the show, we'll talk about what was introduced and offer our first thoughts on the new products. (join mac.hack)

marketing conference—"How to start a newsletter." You, too, can publish your own high-tech newsletter. But first, learn about the pitfalls of starting your own newsletter, how to promote it, what to charge, and much more—from people who have published their own. (join marketing/promotion)

television conference—Do consumers really need HDTV (high-definition TV)? Will viewers even notice the difference in quality? Tune in and find out. And stay tuned for discussions on the use of HDTV in workstation environments to integrate video and computer information . . . the use of HDTV via satellite as a replacement for or a supplement to feature film distribution . . . the current use of Japan's NHK system in producing commercials, features, and music videos. (join television/hdtv)

ti conference—BIX's ti conference members are trying to port MINIX 1.3, a mini Unix operating system, and the Amoeba distributed operating system to the TI Pro. BIXen can learn how a multitasking/multiuser operating system works and, later, how a network-distributed operating system works. And while they're at it, they can learn the real differences between the IBM PC and the TI Pro when it comes to hardware and software. There's no need to have a network to run Amoeba either; it can run and be tested on one processor machine as well as on a network. (join ti/minix)

travsoft—In keeping with BYTE's laptop computer product focus, the BIX travsoft conference has information on Traveling Software's LapLink data-transfer software and Battery Watch. Traveling Software personnel will be available throughout August to answer questions about data transfer or the quirks of NiCad batteries. (join travsoft)

PRINT QUEUE

Hugh Kenner

Somewhere Out There

PLANETS BEYOND:
Discovering the Outer Solar System

by Mark Littmann

W hat gets done with new technology once we have it may be anything but what its designers thought they were after. In the 1930s we were led to expect that television, when it came, would be an adjunct to the telephone. Why, you'll be talking to Jim face to face! What TV wrought instead was not eye contact with Jim (about which we can no longer care less), but the disappearance of the old *Life*, the old *Saturday Evening Post*. . . .

And what we do with computers, perhaps more than anything else, is process words. Never mind that the machine's very name still registers what drove its development—someone's need for a lot of heavy numerical computing. ENIAC was finally pulled together when the U.S. military needed numerical information about shell trajectories—and needed it *fast*.

Back when guns were aimed by eye, though, heavy computation, which meant reducing seemingly endless equations, was most likely to be a burden borne by astronomers. The demand for their numbers persisted, decade by decade. Where will Jupiter be next December 11 at midnight GMT? (Some navigator may need to know.) Reams of paper, hours of candlelight, were once devoted to such chores. And repeatedly, ships got wrecked by the slip of a pencil or the fumble of a typesetter.

Around the year 1820, the error level in published tables was what prompted Charles Babbage to groan out his famous wish for tables calculated "by steam." Babbage even envisaged a "calculating engine" going on to set its output in type. He was eerily prescient. Today it's, yes, by steam that we calculate and print tables, if we happen to live near an electricity plant with steam power somewhere in its delivery chain. Counting peak-time backup systems, you'll hardly find a voltage factory in the U.S. that doesn't fit that description.

Seeking Planets

Planets Beyond is a book about computational astronomy stumbling into its great age. It covers the period from just before the computer to just after, during which the numbers were getting so refined it early grew evident that the solar system wasn't behaving the way Newton said. You've possibly heard a good deal of the story before, likely never in such fascinating detail.

It starts with a chance observation. Only about 76,000 nights ago (Tuesday, March 13, 1781), German-born William Herschel, amateur astronomer, gazing toward Zeta Tauri from Bath, England, spotted a celestial disk. "The quality of his eyes and his instrument told him that this was not one of the 'fixed stars.'" A comet, likely? No, by midsummer three separate mathematicians had fitted it to a planetary orbit. The pencil-and-paper work that Lexell and Saron and Laplace undertook is mind-boggling. So is the agreement of their results. And they were working with circular approximations; in another two years, an elliptical orbit had been derived. Yes, Herschel had happened to glimpse an unsuspected seventh planet, on which the name "Uranus" eventually settled. Twice as far out as Saturn, its existence doubled the known size of the solar system.

The quality of Herschel's eyes we may ascribe to good genes. As for the quality of his instrument, well, Herschel was a musician (oboe and organ), organist in Bath at the Octagon Chapel, and the giver of up to 46 music lessons a week. The link of music with numbers dates back to Pythagoras. So to understand harmony, Herschel studied mathematics. "Math got him interested in optics. Optics got him interested in astronomy." Scanning the sky? That led to a need for optics.

Refractive aberrations being hard to control in those days, he devised, amid potentially lethal explosions that
continued

ILLUSTRATION: LYN BOYER-PENNINGTON © 1989

Knowledge Processing

Don't leave your users lost in a maze of information!

A knowledge processor communicates knowledge - the natural extension of everything we do on a computer.

It's the intelligent integration of everyday resources like data, text, logic, graphics, and video that turns information into knowledge.

KnowledgePro is the first knowledge processor. It combines a high-level, object-oriented programming language with hypertext and expert systems technology.

KnowledgePro gives you a total development environment with the tools you need to create intelligent multi-media applications.

PC Magazine, Holland... *"KnowledgePro is the first of a new generation of software, the knowledge processor...it has the power of, for example, Pascal or PROLOG, but the programmer isn't troubled with the technical details."*

PC Week, USA... *"It's rare, but every so often a PC application comes along that breaks new ground and creates a fundamentally different way to use computers. According to its corporate users...KnowledgePro does just that."*

Infoworld... *"We don't live in a computational world. If we're going to move knowledge around we need tools...The same person who will learn macros in Lotus can learn this."*

KnowledgePro costs $495 with no runtime fees. It runs on IBM PC, XT, AT and PS/2 compatible machines with 640k of memory and a hard disk. A working demo with a 100 page manual is available for $33 including shipping ($38 foreign) with credit towards purchase of the full system.

Find out what knowledge processing is all about. Call 518-766-3000 (FAX 518-766-3003) or write to Knowledge Garden Inc., 473A Malden Bridge Rd., Nassau, NY 12123, USA. Amex, Visa or M/C accepted.

KnowledgePro®
The intelligent way out

Another intelligent tool from

Knowledge GARDEN Inc.

Circle 135 on Reader Service Card

KnowledgePro is a registered trademark of Knowledge Garden Inc. IBM, PC, XT, AT and PS/2 are trademarks of International Business Machines Corp.

greatly distressed his sister, a "speculum metal" (71 percent copper, 29 percent tin) to back reflecting mirrors. And the homemade 6⅓-inch reflector that he spotted Uranus with was superior to anything at the Royal Observatory in Greenwich. He deserves to be all amateur scientists' patron saint.

But by 1824 Uranus was plainly refusing to operate on schedule: now fast, now slow. The most tenacious of several explanations was that Newton's universal gravitation was not quite universal, but commenced to fade out beyond Saturn. One man who espoused that idea was the relentlessly ambitious George Biddell Airy, by 1835 England's Astronomer Royal.

And here (though not in this book) we cross the path of

Herschel had happened to glimpse an unsuspected seventh planet, on which the name "Uranus" eventually settled.

Charles Babbage, lifelong distruster of Astronomers Royal. His first computer—the difference engine—Babbage had designed and built amid splutterings of contempt for Airy's predecessor, whose feckless computations were wrecking ships. He'd have been still less impressed, if he ever learned its details, by the "computer" Airy designed: a roomful of young boys, adding and subtracting throughout 12-hour shifts with a 1-hour midday break. (Replaced by brass wheels, those boys could be out flying kites; meanwhile, one might *trust* the brass.)

Airy's misplaced faith in fading gravitation had unhappy consequences for John Couch Adams. Adams, a virtually self-taught genius, resolved in 1841 to get to the bottom of the Uranus problem. By 1845, aged 26, he'd located, within 2 degrees—using pencils and unthinkable heaps of paper—the place to look for an unknown perturbing planet. There follows a long, dreary story of Airy declining to give Adams the time of day, with the result that a French mathematician, Urbain-Jean-Joseph Le Verrier, got formal credit for locating Neptune.

Adams, an amateur, had held back from publishing lest he'd made a mistake. Babbage wrote him (1847) to point out that if only the calculations could have been automated, he might have put worry behind him. Adams agreed with enthusiasm: "It would be difficult to overestimate the value of such a machine." (Babbage also wrote to Le Verrier, who wasn't excited. He'd got it right with just his quill pen, hadn't he?)

And as far back as 1842, Airy had advised the Prime Minister that calculating machinery such as Babbage proposed was useless in principle. He was proud of what he (not to mention that roomful of boys) did by hand. Airy would die ("still organizing his papers") in 1892, aged 90, bureaucrat in excelsis. Adams, an attractive, modest man, chanced to die the same year, aged 72. He'd declined, 11 years earlier, the offer of suc-

cession to Airy as Astronomer Royal.

Next, Planet X, which Percival Lowell sought. Lowell was one of *the* Boston Lowells, the ones who spoke only to Cabots while the Cabots were speaking only to God. His sister Amy smoked cigars and was something of a poet. Lowell, with the family wealth, established (in 1894) an observatory in Flagstaff, Arizona, meant to document his great enthusiasm, "canals" on Mars. Since that was not a creditable aim, Lowell hoped to gain prestige by finding the new planet one could guess at from still-unexplained irregularities in the motions of Uranus. (Neptune, which might have helped, hadn't been accessible long enough for its misbehaviors to be measurable.)

A Harvard math graduate (honors), Lowell (1855–1916) calculated tirelessly (still in the pencil-and-paper era), hiring and firing up to five simultaneous assistants. On the Airy model, he was running an interrupt-driven parallel computer with irregular wait states and human CPUs. Again and again, locations for Planet X got relayed to Arizona, where observers would point and squint.

The numbers kept improving, and by 1915 they even spotted Planet X, but failed to recognize it. Percy Lowell died the next year. So we credit the discovery of Pluto, in 1930, to Clyde Tombaugh, a 24-year-old whom the Lowell Observatory had hired just a year previously. ("Young man, I am afraid you are wasting your time," one visiting astronomer had told Tombaugh. "If there were any more planets to be found, they would have been found long before this.") Pluto was within 6 degrees of where Lowell had last said it would be: not bad, what with noisy data. The name was first proposed by an English schoolgirl, Venetia Burney, who was learning about mythology. Wasn't that the right name for a dim and gloomy planet? Moreover, they thought at Flagstaff that its symbol, PL, could say Percival Lowell.

Next? Yes, there may still be a trans-Plutonian planet. William H. Pickering (1858–1938) had predicted perhaps six, but none of them turned up. By 1943, Clyde Tombaugh was certain that any such body had to be dimmer than the seventeenth magnitude. His searches had discovered nothing brighter. Yet by 1976, Thomas Van Flandern (U.S. Naval Observatory) was being bothered by problems with the orbits of Uranus and Neptune. We're well within the computer age by now, and errors are most likely observational. Van Flandern's colleague Robert Harrington tells us to imagine taking a sight from Washington, DC, and identifying a drunk lurching out of a Baltimore bar by his stagger. Such is the minute scale of the wobbles. Atmospheric flickers? Maybe. That search goes on. And, with thanks to ENIAC's progeny, its difficulties no longer pertain to computation.

Meanwhile, on some of Littmann's liveliest pages, we're given details of the marvelous computer-enabled *Voyager* tours of inspection, Jupiter-Saturn-Uranus, with Neptune scheduled for late August. The photos (some in color) are entrancing.

"Very well written," says Tombaugh. "By far the best on the subject I have seen," says Harrington. Opinions from Herschel and Adams, alas, we can't have. But Littmann's is, yes, by anyone's verdict, a superb book. ∎

John Wiley & Sons, New York: 1988, 286 pages, $22.95

Hugh Kenner is a professor of English at Johns Hopkins University. His reviews have appeared in publications like the New York Times *and* Harper's. *His recent books include* A Sinking Island *and* Mazes. *He can be contacted on BIX as "hkenner."*
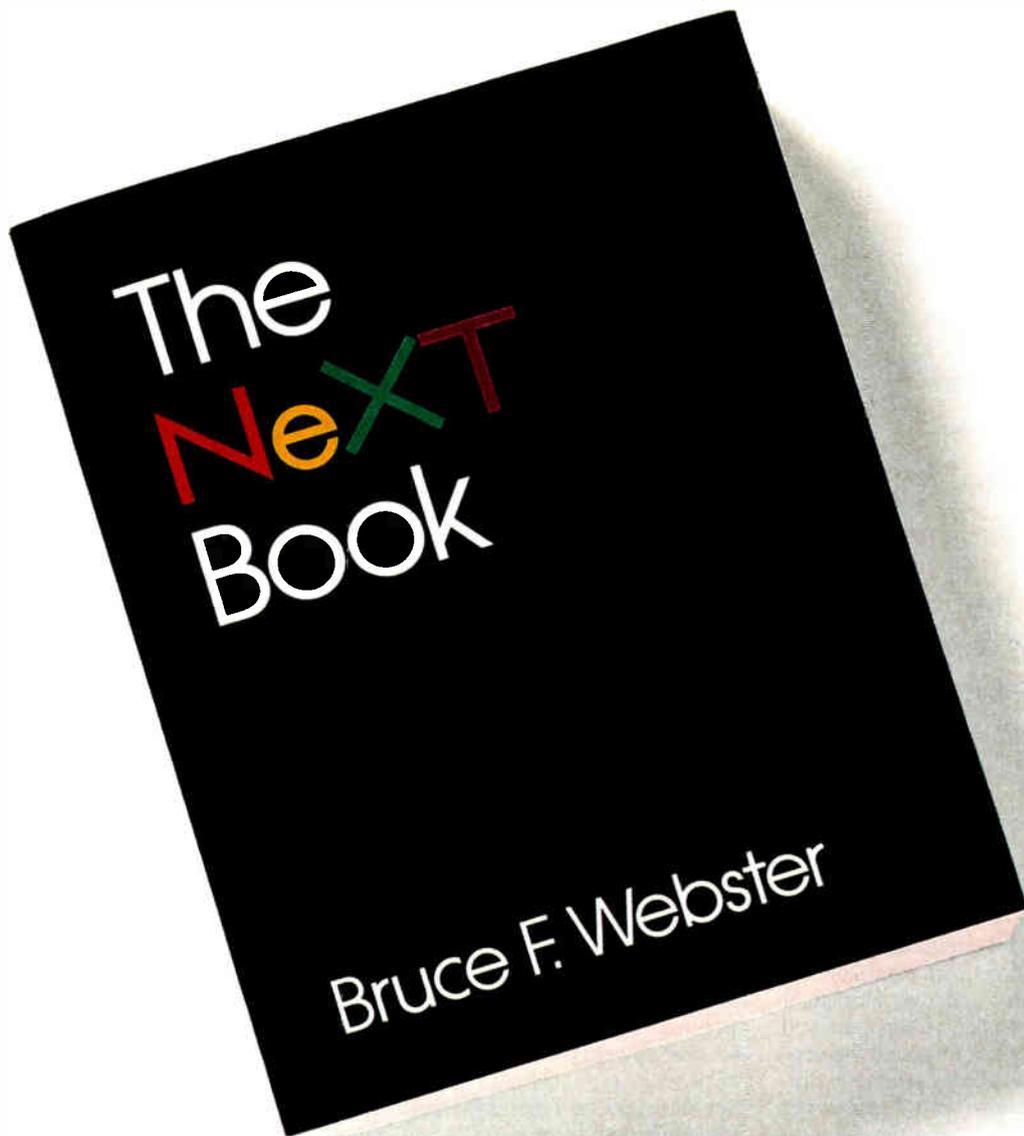
Your questions and comments are welcome. Write to: Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.

World Radio History

THE LONELINESS OF THE LOW-BUDGET USER

Are computer companies forgetting the people who put them where they are today?

As I watch the endless onslaught of high-priced, "high-end" computer systems and software applications with insatiable appetites for memory and disk storage, I can't help but wonder if the personal computer industry has forgotten the people who made it a success—namely, single users. You know who I'm talking about—people like you and me, who do word processing, maybe a little bookkeeping or budget forecasting with a spreadsheet program, have a couple of databases, and even dabble in a little programming now and then. We may be engineers and architects who do some of our calculations or preliminary drawings on a desktop computer like an IBM XT or AT or a Mac Plus or SE. We might still be dragging our Compaq portable from the office at night to finish up a report.

We were the ones who spent a few thousand bucks early in the game to get a machine and jump on the microcomputer bandwagon so that we could get out from under the MIS department's backlog and free ourselves from system administrators, database administrators, and corporate rules and regulations for computing. To us, it looks like we're being left out in the cold by the very vendors that made this all possible—companies like IBM, Apple, Compaq, Microsoft, Lotus, and Ashton-Tate.

When they started out in the personal computer business, these companies produced affordable hardware and software. The software worked, and it required no more than 512K bytes of RAM, often

less. These microcomputer enterprises were started by pioneers, from Phil Estridge at IBM to Bill Gates at Microsoft, to the Woz and Jobs at Apple—pioneers who bucked the system and showed that you could do useful work without expensive minicomputers and mainframes and software applications that cost thousands of dollars and come with a monthly maintenance contract.

But the focus of these companies has changed. Today, it seems that we're giving control back to the MIS department and network administrators. We're worrying about mainframe connectivity and file servers. We're looking at network operating systems, like NetWare 386, that cost upwards of $8000. Steve Jobs is off producing a machine that costs well over $10,000 at the retail level when fully equipped. Microsoft has built a virtual army of programmers around OS/2, which requires at least 3 megabytes of memory and nothing less than an 80286 machine. Ashton-Tate and Lotus are mired in trying to shoehorn every feature under the sun into programs that were designed for single users running single tasks. Apple keeps its prices high and keeps offering ever-more-expensive and powerful machines, but very few new products that the average user can afford. Compaq has reached new heights with an 80386 machine that costs $18,000. The low-end Compaq has become the 286 SLT, which costs a mere $6000.

Somehow, these companies have forgotten that the whole idea is to make computer hardware and software more affordable and more accessible to greater numbers of people. They seem to have forgotten that most students don't have $7000 or $8000 lying around for a workstation to put in their dorm rooms. They forget that a lot of users don't have and can't afford 3 megabytes of memory.

And then there's the question of ease of use. We used to talk about "user friendliness." But the level of complexity seems to be going in the opposite direc-

tion. Most new software products (and new versions of old products) have so many features, the user doesn't know where to begin. All the on-line help in the world won't help you if there are several hundred feature options to learn. I'm afraid to upgrade to WordPerfect 5.0 because I don't want to have go through setting up my printer again and learning all those new options that I probably will never use. WordPerfect 4.2 does just fine, thank you.

I'm not saying there's no place for high-end systems and connectivity. Of course, this is a major concern for many organizations, and what can be done with personal computers these days is indeed impressive—in every type of work from engineering to accounting. But let's not forget the little guy. Let's see some new, even easier-to-use products for the single user. Let's see some new machines—with new capabilities—priced around $2000. Let's see some innovative software engineering that employs data compression and object-oriented techniques to allow big applications to run on little machines.

There is hope. Borland recently announced a technology—the Virtual Real-Time Object-Oriented Memory Manager—that does just that. Apple has hinted at a new low-cost machine that will run the new Mac operating system in ROM. Let's hope the other computer heavyweights come up with some new low-cost innovations as well. It's in their own best interest. Somehow, I get the feeling the big companies are gambling their futures on these high-priced solutions. They had better remember what got them here before it's too late. ■

Nick Baran is a BYTE senior technical editor based in San Francisco. He can be reached on BIX as "nickbaran."

Your questions and comments are welcome. Write to: Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.