BYTE

FEBRUARY 1989   A McGRAW-HILL PUBLICATION

PRODUCT FOCUS

6 Great C Compilers

The New Mac SE/30

Mac II power with an SE price

BENCHMARKS

Macintosh SE/30

PLUS

Paradox 3
  *A Contender
  for "Best DBMS"*

Digital Paper

Turbo Windowing

4 Short Takes

Why You Should Consider The Dell 386 Systems, Despite Their Suspiciously Low Prices.

TO ALLAY YOUR SUSPICIONS,

YOUR TOTAL SATISFACTION IS GUARANTEED.

Dell offers a complete line of systems, peripherals and software.

And whatever you choose, you must be absolutely, totally satisfied with it. If not, simply return it within the first 30 days for a money-back refund. No questions asked. We call this our Total Satisfaction Guarantee.

SUPPORT IS A FREE PHONE CALL AWAY.

Also included in your system price is our toll-free support line. It's open every business day from 7 AM to 7 PM (CT). Simply dial (800) 624-9896.

ON-SITE SERVICE IS INCLUDED IN YOUR SYSTEM PRICE.

For those rare problems that can't be handled over the phone, we'll send a service technician to your location by the next business day. At absolutely no charge to you? Because included in system price is a year's worth of on-site service.

ONE-YEAR WARRANTY.

Dell offers a one-year limited warranty, which warrants each system we manufacture to be free of defects in materials and workmanship for one full year. During that period we will repair or replace any defective products returned to our factory.

For a lot of companies, leasing our systems is an even better idea because of the cash flow and tax advantages. And we guarantee fixed rates, so you know exactly how much to budget each month. We can even custom design a lease plan to fit the exact needs of your business.

For the complete terms of our On-Site Service Contract, Satisfaction Guarantee, Warranty and leasing plans, write: Dell Computer Corporation, 9505 Arboretum Blvd., Austin, Texas 78759-7299.

THE NEW 25 MHz 386 SYSTEM 325.

When you need the highest possible performance of any 386, this is the technology of choice. Running at 25 MHz, the System 325 is faster than the Compaq 386/25. Besides unequaled speed, it also offers Intel's Advanced 82385 Cache Memory Controller and high performance disk drives. As a result, it gives you workstation-level performance for CAD/CAM and desktop publishing applications. It's also especially effective as a network file server, and more than capable of handling the most complex spreadsheets and databases.

STANDARD FEATURES:
- Intel 80386 microprocessor running at 25 MHz.
- 1 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:
- 25 MHz Intel 80387 math coprocessor.
- 1 MB or 4 MB memory upgrade kit.
- 2 MB or 8 MB memory expansion board kit.
- **Lease for as low as $245/Month.

System 325	With Monitor & Adapter	
Hard Disk Drives	VGA Mono	VGA Color Plus
150 MB-18 ms ESDI	$6,799	$7,099
322 MB-18 ms ESDI	$8,799	$9,099

The Dell System 325 is an FCC Class A device, intended for business use only.

THE DELL 20 MHz 386 SYSTEM 310.

For business users who need a 386 system, this is the best combination of performance and value available. Running at 20 MHz, this 32-bit system is faster than the IBM PS/2 Model 70 and the Compaq 386/20e. Since it has the same high performance disk drives and Intel Advanced 82385 Cache Memory Controller as our System 325, it brings a new level of performance to complex spreadsheets and databases. As you might expect, it runs windowed software at extremely high speed. It's also well-suited for desktop publishing applications, or as a network file server.

STANDARD FEATURES:
- Intel 80386 microprocessor running at 20 MHz.
- 1 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 20 MHz Intel 80387 or 20 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:
- 20 MHz Intel 80387 math coprocessor.
- 1 MB or 4 MB memory upgrade kit.
- 2 MB or 8 MB memory expansion board kit.
- **Lease for as low as $140/Month.

System 310	With Monitor & Adapter	
Hard Disk Drives	VGA Mono	VGA Color Plus
40 MB-28 ms	$3,899	$4,199
90 MB-18 ms ESDI	$4,699	$4,999
150 MB-18 ms ESDI	$5,199	$5,499
322 MB-18 ms ESDI	$7,199	$7,499

All prices and specifications are subject to change without notice. Dell cannot be responsible for errors in typography or photography. **Payments based on a 36-month open-end lease. In Canada, configurations and prices will vary. Microsoft, MS and MS-DOS are registered trademarks owned by Microsoft Corp. ˆSignifies trademarks of entities other than Dell Computer

World Radio History

World Radio History

BYTE

FEBRUARY 1989   VOL. 14/NO. 2

COVER PHOTOGRAPHY: PAUL AVIS © 1989

BYTE

smARTWORK® Keeps Getting Smarter

Smarter Artwork

Three years ago, Wintek engineers created smARTWORK to reduce the time and tedium of laying out their own printed-circuit boards. Thousands of engineers have since discovered the ease of use and sophistication that makes smARTWORK the most popular PCB CAD software available. And thanks to them, smARTWORK keeps getting better.

smARTWORK transforms your IBM PC into a PCB CAD system

Interactive routing, continual design-rule checking, pad shaving, and production-quality 2X artwork have always been a part of smARTWORK. And now that many customer suggestions have become a part of the software, smARTWORK is an even better value. *That's why we offer it with a thirty-day money-back no-nonsense guarantee.*

Using an Epson or IBM dot-matrix printer, you can create 2X artwork and 1X or 2X checkplots in a fraction of the time hand-taping requires

New smARTWORK Features

☐ Silkscreen layer for component placement and identification
☐ Text capabilities for all three layers
☐ Selectable trace widths and pad shapes and sizes
☐ User-definable library
☐ Ground planes created with a single command
☐ Solder-mask and padmaster plots generated automatically
☐ Quick printer 2X checkplots
☐ Additional drivers for printers and plotters
  ☐ Optional drill-tape and Gerber photoplotter utilities
  ☐ AutoCAD® .DXF file output
  ☐ Completely updated manual
    ☐ 800 number for free technical assistance

2" by 4" section of a 10" by 16" double-sided board with silkscreen layer

Camera-ready
2X artwork from an HI or HP pen plotter

The Smart Buy. Guaranteed.

Still priced at $895, smARTWORK is proven, convenient, and fast. Our money-back guarantee lets you try smARTWORK for 30 days at absolutely no risk. Call toll free (800) 742-6809 today and put smARTWORK to work for you tomorrow. That's smart work.

System Requirements

☐ IBM PC, PC XT, or PC AT with 384K RAM, and DOS V2.0 or later
☐ IBM Color/Graphics Adapter with RGB color or B&W monitor
☐ IBM Graphics Printer or Epson FX/MX/RX-series printer, and/or
☐ Houston Instrument DMP-40, 41, 42, 51, 52, or Hewlett-Packard 7470, 7475, 7550, 7580, 7585, 7586 pen plotter
☐ Optional Microsoft Mouse

Wintek Corporation
1801 South St.
Lafayette, IN 47904
Telephone: (800) 742-6809
In Indiana: (317) 742-8428
Telex: 70-9079 WINTEK CORP UD

Current Versions	
HiWIRE	V1.1r1
smARTWORK	V1.3r5

"smARTWORK", "Wintek", and the Wintek logo are registered trademarks of Wintek Corporation. "AutoCAD" is a registered trademark of Autodesk, Inc.

World Radio History

THE END OF APPLICATION SOFTWARE?

IBM and Microsoft have some very interesting plans for OS/2, PM, and DOS

Several weeks ago, IBM invited a handful of editors to its facility in Hursley, England, where most of Big Blue's work on user interfaces goes on. The announced focus of the discussions was to be Presentation Manager (PM), but it ended up being considerably more than that, expanding to encompass the entire future of OS/2. It's a centralized, uniform future, radically different from the DOS-based world.

Part of the standardization derives from the fact that PM is a piece of IBM's Systems Application Architecture. With PM, SAA has moved squarely into the personal computer arena.

SAA is centered on four key concepts: a Common User Interface, a Common Programming Interface, Common Communications Support, and Common Applications. (The Hursley group is responsible for the CUI.)

All this uniformity portends extreme portability among IBM system software and applications, from OS/2 Extended Edition (described as "IBM's personalization of the generic OS/2 operating system") through OS/400, and up to VM and VME.

OS/2 EE's components already come close to SAA standards. (Actually, because SAA itself will evolve, it's not a "standard" in the strict sense of the word; it's more a goal or a guideline.) OS/2 will become more tightly meshed with SAA over time. For example, the database manager portion of OS/2 EE embodies about 90 percent of the functions of Structured Query Language (SQL); future versions will raise that to 100 percent.

If IBM has its way, SAA standardization will be far-reaching: For example, virtually all data will end up in a central SQL repository, which all applications will access via the built-in database manager. This way, the data generated by any application anywhere in a system of networked machines will become accessible by all other applications on all the other machines, as well.

SAA's shared data, cooperative applications, and shared peripherals means that—in theory—everyone using Big Blue's software and hardware (and, one presumes, close clones) will see a "single system image." Everyone has access to everything, and everyone can work together easily, regardless of whether they're on a PS/2 or a 370.

IBM accepts the oft-quoted statistics that claim that by 1991, 85 percent of all PCs will be interconnected. IBM sees this as an argument in favor of a central repository—all data becomes accessible to all machines on the network. Users neither care nor need to know where the data is located; the system handles it all.

Is this software fascism? Well, no; the storing of code and data in separate files won't go away (especially for individual program code). IBM's John Soyring said that "other types of files would be allowed" so that programs wouldn't have to use the central repository, but that IBM would simply "encourage" the use of SQL facilities.

It's interesting to note that Borland's not-yet-released Paradox SQL partially fits this pattern: It can function as a front end, or preprocessor, that automatically generates SQL queries to access remote, centralized data. As such, it may be a bellwether for a new kind of application program that we'll be seeing a lot more of: one in which a vendor carves a niche for a product with added-value special features (like a superior query-by-example facility), but relies on standardized system resources (like a central SQL data store) when they're available.

Under PM, the standardized resources can also include the user interface, communications, and more. It may sound somewhat Mac-like, but under IBM's scheme, the commonality is much wider, extending even across architectures and operating systems.

This is a far cry from the state of DOS applications today, where programmers often continuously reinvent the wheel: Each application may have its own little presentation manager, its own little communications manager, and so on. With these major pieces standardized and readily available at all levels throughout a business, applications programming becomes simpler. Soyring puts it this way: "Programming will become macro writing." Programmers won't have to spend time developing the raw materials of their applications; instead, they "can concentrate on pure functionality, and their productivity will increase."

The thought was echoed by Microsoft's Steve Ballmer. Because of PM's object orientation, "Applications will become objects," he said, "and they'll simply call other applications to use their services. Programming will just be writing scripts in an external control language to use the various objects." If that happens, it will be the end of stand-alone applications as we now know them.

Except under DOS. Next month, I'll talk about IBM's and Microsoft's plans for DOS, and their plans for Unix and an 80386-specific version of OS/2.

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

REPLACE YO

Installed in minutes, the Renegade 386™ motherboard turns any IBM PC, XT, AT or inexpensive clone into a powerful new computer, lightning fast with both old software AND new OS/2 releases.

I s this you?
You would like to have extraordinary computing speed and power. You want number crunching ability. You need complex CAD capabilities. Or graphics. Or multi-tasking.

And you haven't got it. (And you don't want to spend a fortune to replace good equipment with the newest standard.)

Weep no more. You can turn your *present* system into the latest, high-performance system that will *equal or exceed the best of the new computers.*

Renegade 386.™
A new heart and brain for your computer.

Renegade Technologies offers designers, engineers, architects, and computer-dependent businesses a simple and reliable alternative to obsolescence. Or the unreasonable expense of a new system.

Simply replace the "motherboard" of your present system with a Renegade 386™ motherboard.

It takes a screwdriver and less than 20 minutes. And costs but $1695.

That's *thousands of dollars less* than a new IBM System/2 Model 80.

But it gives you more than just the latest industry standard. You can run your old software on it. Probably anything you now use on your XT or AT. Big Blue can't do that.

You can use your present 16-bit peripherals. (If you've looked at PC "add-on" cards, you already know your old equipment would be useless.)

But with Renegade 386™ you'll have to find some other excuse to throw away your current modem, network card,

Multi-task with Lotus 1-2-3 and other applications running side-by-side on the Microsoft Windows/386 desktop.

EGA or disk controller cards.

The Renegade 386™ board comes with an iron-clad *one year limited warranty.* It uses genuine U.S. made Intel 386 chips and is designed and manufactured in the U.S.A. by Hauppauge Computer Works. Hauppauge is a major developer of software support for Microsoft and IBM products, and is producer of the highly respected Hauppauge 8087 or 80287 highspeed math coprocessors. Over 50,000 have been sold.

Major computer magazine reviews in the last year have hailed our Hauppauge-made board as a major breakthrough in a high speed, high power, high performance upgrade product.

Not an accelerator card. Not a "turbo" gimmick. Renegade 386™ gives you a full-featured new computer.

Thanks to Renegade's 80386 microprocessor your computer will now boast a 32-bit data path —and a clock speed of 16 MHz with zero wait state access. Up to 8 expansion slots are provided depending on your computer configuration.

Your "new" computer also will now have 1 Megabyte of 100ns RAM. This is *not* a naked board. And it also includes a 32-bit high-speed RAM expansion slot which you can populate with up to *15 Mbytes* of system memory.

World Radio History

MICROBYTES

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

Mead's Silicon Retina Points Toward Brain-Like Processing

Carver Mead and associates have developed what could be a true innovation in computer technology: analog emulation of the human nervous system. Mead, who is credited with major contributions to IC design and very-large-scale-integration (VLSI) technology, said construction in silicon of an information processing system similar to the brain is well under way.

Mead and his collaborators at Synaptics (San Jose, CA) have designed a model of a "small but identifiable process of the human brain"; they call it the Silicon Retina. The Silicon Retina consists of a 3- by 4-inch board of analog circuitry connected to a camera lens; it performs the same visual processing as the human retina. The main chip on the board is manufactured by Synaptics and consists of an array of photo sensors that emulate the photo receptors of the human eye. An external object passed across the lens is processed by the analog array processor, and the image is displayed on a standard Multi-Sync video monitor. In a demonstration at San Fran-

cisco's Exploratorium science museum, the image processor was able to accurately display images of the human hand and other objects passed across the lens.

The basic function of animal vision is accomplished by photo receptors that compute the average of the local light intensity and compare this value to the value of the incoming light intensity from the external source. The differential of these values is sent as a signal to the optical nerve, resulting in identification of the external object. The array-processing chip in the Silicon Retina performs a similar function. Mead said that the Silicon Retina represents a "gain control mechanism" for the processing of light signals. It uses a concept called *edge enhancement* to determine the difference in light intensity of the edges of an object in comparison to a uniform background lighting.

Although the Silicon Retina is an impressive example of the power of analog processing, there are still major obstacles to the design of commercially viable systems. According to Mead, analog

processing can achieve "100,000 times the efficiency of digital computing" but is constrained by the lack of uniformity in present-day transistors. Mead said that the human brain is faced with the same problem but gets around it by adapting or adjusting to the inconsistency of the nervous system. This adaptation is accomplished by "averaging" the values of incoming signals. Accomplishing equivalent adaptation in silicon is obviously a major challenge.

Still, Mead predicted that neural technology will eventually have as significant an impact on human culture as the present-day digital computer. Mead donated a working model of his Silicon Retina to the museum.

According to an engineer at Synaptics, the company hopes to have a commercial version of the Silicon Retina ready in about a year. Said Synaptics engineer Tim Allen, "We've proven that it [the Silicon Retina] works in principle. Now we have to make it fast and cheap." Allen thinks the hard part is over. The next phase just requires some basic "semiconductor engineering."

Flexible Superconductors Will Shape PC Boards, Power Supplies

Ever since the 1986 discovery of ceramic superconductors, promises of their potential have flown faster than greased electrons. The media has brimmed over with stories of superfast computers and magnetic-levitation trains, even while scientists have cautioned that such applica-

tions are many years away. But now researchers at MIT have overcome a major impediment to the commercialization of superconductors—their inherent brittleness.

By mixing noble metals, such as silver, gold, platinum, or palladium, into the ceramic material, research-

ers have for the first time been able to create flexible superconductors that can be pressed into sheets or rolled into coils. "Without modification, the mechanical properties of superconductors render them about as flexible as dry spaghetti," said Gregory J. Yurek, an

continued

Dave Ditzel proclaimed **Sun's SPARC** processor as the de facto standard for reduced-instruction-set-computer (RISC) processors, and chip-designer-turned-industry-critic Nick Tredennick countered that **Intel** would produce more 80386s between coffee break and lunch than Sun would make SPARC processors in an entire year.

• The only **legal Macintosh clone** available in the U.S. is an **Atari ST** that uses Mac ROMs. Previously, these ST-based clones have been limited to 64K-byte ROMs; a new version, called the **Spectre 128**, uses 128K-byte ROMs and is said to run anything that will run on the Mac II (monochrome only) and virtually all Mac Plus and SE software, including the usual suspects: HyperCard, Adobe Illustrator, and PageMaker. Spectre 128 retails for $179.95, and it comes from **Gadgets by Small** (Littleton, CO), without the Mac ROMs (you have to get those from other sources). Gadgets by Small is the new company of **David Small**, who developed the famous **Magic Sac** Macintosh emulator.

• Despite ongoing—well, rampant—speculation about a new **laptop from Apple** Computer, Apple's John Sculley says the company's plans for a laptop are still hampered by **inadequate screen technology**. "I wish we had it now," Sculley said recently, but he added that a laptop is not a big part of Apple's plans for 1989. Sculley said the main

continued

MIT professor and founder of American Superconductor Corp. (Cambridge, MA), which has been set up by the MIT scientists to handle commercial aspects of the new material. "A composite material obviously sacrifices some conductivity, but in exchange you can actually use the superconductors for something useful." In addition, he said, the composite material is more resistant to oxidation and can be connected more easily than before to electrical sources.

The applications for the more malleable material are still at least 5 years off, Yurek said, but they may include tiny motors, satellite sensing devices, electromagnetic shields for computer power supplies, and printed circuit boards. "Superconductors form a better shield against electromagnetic radiation, all the way from low-frequency DC up to very high frequencies," he said. Flexible superconductors could be shaped into interconnects or traces on printed circuit boards that would be faster and lose less power than current designs. Yurek envisions active components, such as

chips, made with superconductors.

Yurek anticipates no big problems introducing such supercooled fluids as liquid nitrogen into computers, which would be necessary to bring the ceramics down to the $-300°$ Fahrenheit they need to conduct. "Some supercomputers are already cooled by liquid nitrogen," he said. "Now the trick is to get that capability onto the desktop."

"There's no question this development could have quite an impact on computing," Yurek said, "but it's a long way down the road."

BASIC Will Be the "Embedded Language" of Microsoft Applications

Microsoft has "learned its lesson from its Mac experience" and wants to provide a development environment that mimics the delivery environment, says Greg Lobdell, the company's product manager for languages. Therefore, the shell for any future object-oriented programming (OOP) development systems is likely to be Presentation Manager. Inside that shell, the 18 or so OOP building blocks will call routines written in BASIC—not C.

Why BASIC? Current versions are very fast, says Lobdell, and applications built out of OOP tools are likely to be heavily oriented toward transaction processing. BASIC's string-handling capabilities will be an advantage in mixed text-and-integer applications, claims Lobdell. Moreover, BASIC is more "natural," or English-like, than C or Pascal, he says. "BASIC will evolve to be the embedded language of Microsoft's applications."

Ideally, any OOP language products would, like other Microsoft language products, have "inter-language support"—that is, any language could call routines from any other language to ensure a rich application development environment. In the long term, Microsoft wants to provide a set of object-oriented building blocks, the functional equivalent of NeXTStep's 18 basic OOP routines to handle dialog boxes, menus, and so on.

CISC or RISC, Chips Are Becoming More Alike

Bring together some of the designers who worked on the computer industry's major microprocessors, and chances are they'll each have a different opinion about which microprocessor is the best. At a recent conference sponsored by the *Microprocessor Report* newsletter, we heard proponents of Intel, Motorola, Sun, Advanced Micro Devices, and MIPS Computer extol their favorite chips and architectures. But one thing they all seemed to agree on

is that microprocessors are in some ways becoming more and more alike.

"CISC processors are becoming more RISC-like, and RISC processors are becoming more CISC-like," said Motorola's Russell Stanphill. In the next few years, implementations of complex-instruction-set-computer chips and reduced-instruction-set-computer chips will become much more similar, predicted Michael Slater, who edits the *Microprocessor Report* and orga-

nized the conference.

The designers generally agreed that within the next few years, all the major microprocessors will have cache, floating-point, and memory management units built in, and they'll incorporate features like register scoreboarding and separate data and instruction paths.

But the chip designers said they're concerned about the limitations of performance that confront single-processor architectures. In

continued

To learn Microsoft QuickBASIC, you only need a manual this thick.

World Radio History

problem with current laptop displays is that they can't handle mouse operations adequately. According to Sculley, the mouse pointer "disappears" when you drag it quickly across the screen. Sculley said he "hopes to see a laptop in 1989, though not in quantity." Apple is looking at active matrix displays as one possibility, we hear.

• For those who missed the news, the U.S. now has an official **undersecretary for technology**. It's a new position in the Commerce Department, responsible for the department's new Technology Administration. "The undersecretary for technology will serve as a strategic catalyst to promote the use of science and technology by industry and entrepreneurs," said an official statement from Washington.

• The industry will be able to use the help, if things shape up the way the **American Electronics Association** predicted in a recent release. If the U.S. doesn't play a strong role in the **high-definition TV (HDTV) market**, there'll be a "ripple effect" that will cause the country to lose "significant" world market share in personal computers and semiconductors. The report warns that the country's current 70 percent share of the world personal computer market could dwindle by the year 2010 to half that if the nation's electronics companies don't take the lead in HDTV.

• Citing "misrepresentation of facts surrounding the issue of com-

spite of Dave Ditzel's glowing prediction that Sun will be the first company to offer a RISC chip with 4-nanosecond cycle times, the big question is how to supply memory subsystems that are fast enough for these high-speed processors. As Intel's Steve McGeady put it, "Talking about 4-ns cycle time is fun, but who's going to sell you memory for that?" Motorola's Mitch

Alsup agreed that the biggest challenge is designing memory and processor I/O that can keep up with cycle times of less than 10 ns.

Professor Yale Patt of the University of Michigan said object-oriented data models may be part of the answer to increasing memory bandwidth. He also argued that simply adding bigger caches and multilevel caches is not the answer. Addition-

al caches become slow and complex, Patt said. He outlined a new model for microprocessor design that introduces new concepts for instruction control and data execution, including "wide words" and node tables for storing and predicting the instruction stream. Patt described the concept as "pipelining with rest areas. You're only on the highway when you're working."

Computers, Communications Play Role in New Global Foundation

Computers and electronic communications will play a fundamental part in the new International Foundation for the Survival and Development of Humanity, an ominously named group of scholars and scientists addressing such ominous global problems as environmental pollution, nuclear disarmament, and abuses of human rights. The foundation has an impressive list of members, including Soviet physicist Andrei Sakharov and the former President of MIT, Jerome Wiesner.

At a press conference sponsored by Apple Computer, members of the foundation discussed electronic conferencing and communications as a means of improving relations between the U.S. and the USSR. (Apple has donated computers to the group, and Apple CEO

John Sculley is on its board of directors.)

The foundation, which has offices in Moscow, Stockholm, and Washington, DC, is raising money and accepting project proposals. One project proposed already is an electronic communications system between the U.S. Congress and the Supreme Soviet, the primary legislative body in the USSR. Another project involves the expansion of a U.S.–Soviet computer training camp for children.

With members from all over the world, the foundation will use an electronic conferencing system to conduct its activities. Computers will also be used for projects involving the collection of environmental data relating to acid rain, the "greenhouse effect," and other forms of pollution.

Sakharov expressed frustration at the group's slow progress thus far. Speaking at Apple's press conference via satellite, Sakharov said he has a "cautious attitude" toward the foundation; it has been "mainly talk" so far, he said. Wiesner, though, responded that now that the foundation has been officially sanctioned in the USSR, it will start to aggressively solicit projects.

The foundation hopes to sponsor about 10 projects per year. It encourages project proposals that involve international cooperation, particularly between the U.S. and the USSR. For information on submitting project proposals or contributions, you can write to the International Foundation for the Survival and Development of Humanity at 109 11th St. SE, Washington, DC 20003.

NMB's 1-megabit DRAM Chip Claims Top Speed

NMB Technologies (Chatsworth, CA) says it has developed the "world's first" 60-nanosecond 1-megabit dynamic RAM chip. Although other manufacturers are soon to follow, NMB's AAA1M200 is the first formally announced 1-megabit DRAM with a minimum access time of 60 ns. Currently avail-

able 1-megabit DRAMs have access times ranging from 80 ns to 125 ns.

The new DRAMs are targeted for high-performance 32-bit processors operating at 16 MHz or higher. Most 32-bit machines operating at 20 MHz or higher require at least one wait state to compensate for the speed of currently available

DRAMs. The 60-ns DRAM will allow manufacturers to eliminate wait states and also complex memory interleaving and caching schemes, an NMB spokesperson claimed.

The company said it will start shipping the high-speed DRAMs in volume in the second quarter of this year. It

You'll never know how quick it is until you open it up and see what it can do.

All the horsepower in the world is useless until you can get at it.

Introducing new Microsoft® QuickC® version 2.0. The fastest, easiest way to master all the power and glory of C. Frankly nobody but Microsoft packs this many exclusive features into a single package:

For starters, there's QC Advisor —a new hypertext electronic manual that teaches, helps and guides you *on screen*. It even lets you cut and paste sample programs, so you can learn C the easy way.

By example.

And "C For Yourself," our detailed book of C fundamentals that'll give you more in-depth programming lessons. You'll find one in every box.

With Microsoft QuickC, you can ease into C with Easy Menus to help you write your first C programs, then advance to Full Menus to access the full-throttle potential of C. You'll be up to speed in no time.

And speaking of speed, QuickC has enough muscle for incremental compiling and linking at an incredible 25,000 lines per minute—so it'll make short work of any development you have in mind.

New Microsoft QuickC 2.0. Zero to expert in record time.

And speaking of development, its in-line assembler lets you write assembly code within your C code for more efficient programs. Plus it's the only integrated debugger that lets you simultaneously debug C *and* your in-line assembler. Talk about convenience. It even supports all memory models within the integrated environment.

For all the details, call us at (800) 541-1261. The new Microsoft QuickC.

Get it. And break a few speed limits.

Microsoft
Making it all make sense.

could mean the start of fierce competition among DRAM manufacturers to offer higher-speed memory chips, which will be needed to keep up with the ever-faster processors in ever-faster workstations. NMB Technologies is an American arm of the Japanese electronics manufacturer Minebea Co. of Tokyo. Minebea also makes keyboards, power supplies, and miniature precision bearings.

Optical Drives Could Boost Entertainment Software

After losing shelf space at software retail stores and failing to attract a bigger base of users, producers of entertainment software are hoping that optical drives will result in dazzling programs and hence increased interest in games and other nonbusiness (i.e., fun) programs. The customer base so far has consisted mainly of men and boys between 12 and 45. Several hundred companies used to make computer games, but many have gone out of business or seen sales decline.

Some of the survivors of the entertainment software business were on a panel at a conference in San Francisco recently, including Electronic Arts, EPYX, Mediagenic (formerly Activision), and Spectrum Holobyte. The main theme of their discussion was that today's personal computers are too primitive to run the kinds of games that would attract a broad consumer base for entertainment software. Trip Hawkins of Electronic Arts pointed out that the most popular games machine these days is a cheap IBM PC clone, which has replaced the Commodore 64 as the home computer of choice. But the bare-bones PC clone with CGA-resolution monitor isn't "friendly enough" for the mass market. Hawkins said that the entertainment market is still a market for "hobbyists."

The big hope for the entertainment software business is the proliferation of optical drives and media, such as CD-ROM and similar technologies. CD-ROM will allow entertainment software developers to put megabytes of data on a sin-
continued

BRIDGES

The *easy* way to move files between the 5¼" world of PCs, and the 3½" world of PS/2s, Laptops, and 386s: Sysgen's Bridge family.

Because your PCs, PS/2s, 386s and laptops all talk to different floppy disk sizes and formats, they can become frustrating *islands* of information.

For the simplest way to get files back and forth, just install Sysgen's Bridge products. Presto: Your different computers can talk to common diskettes.

Sysgen pioneered bridges. Now here are three ways to simplify your life in a multi-computer office.

Solution 1: Our Bridge-File 5.25 floppy disk drive.

A best seller, with over 50,000 installed. Attach it to your PS/2, and you can read and write files from PCs. Critics and users alike have raved

about its small footprint, ease of use, and dual density capacity of 360Kb and 1.2Mb. (IBM's drive is twice the size, yet stores only one-fourth the data.)

Every PS/2 user deserves this easy connection to the world of PCs.

Solution 2: Our Bridge-File 3.5 floppy disk drive.

You get 720Kb *and* 1.44Mb capacities. Connect it to any PC, and you can read and write files from PS/2s, 386s, laptops, and other 3½" machines.

Now all your PCs can share files with your newer systems.

And here's the ideal controller: The Omni-Bridge Controller.

Controls up to four drives, so you can mix and match 5¼", 3½", and floppy tape drives as you need. This hot card gives you big storage, plus twice the floppy disk and floppy tape transfer performance rates! At only $95, this is a winner.

Get your computers talking. Call the Sysgen Hotline for the name of your nearest dealer. 1-800-821-2151.

SYSGEN
INCORPORATED

556 Gibraltar Drive, Milpitas, CA 95035

World Radio History

has support for 8 mega-bytes of 60-nanosecond dynamic RAM on the motherboard, and a 32-bit expansion slot for up to 8 megabytes more. The product's designers went with 4 megabytes of single in-line memory modules and 4 mega-bytes of standard DIP memory because they weren't sure which would be most hit by a short-age. AMI officials said they weren't sure when they'll be able to get Intel-certified 33-MHz 80386s.

• Some **optical storage companies** have banded to back the **Continuous Composite Servo** (CCS) recording format for 5¼-inch erasable optical

continued

gle disk and thus create more complex and dynamic games. But you need faster performance than current CD-ROM technology pro-vides to make games per-form well. Optical storage in the 1990s could provide the breakthrough needed for en-tertainment software to take off again. In the mean-time, the survivors are faced with a resurgence of competition from Japan and Europe. As Hawkins said, "there's an increase in product supply without an increase in demand."

Even on a Cloudy Day, Norton Can See DOS Forever

While forecasts call for increased use of Unix and OS/2, in the future ac-cording to Peter Norton, MS-DOS "goes on forever." Norton, of self-named Util-ities fame, said in an inter-view with *BYTEweek* that he foresees the hordes of PC users eventually falling into three camps—"OS/2, High-DOS, and Low-DOS." Norton equated High-DOS users with today's "power users." There will be a huge market of "Low-DOS" users, he said, with 80286-based machines becoming the low-end standard.

And what does the com-pany president/writer/in-dustry celebrity think of OS/2? It's "a necessary evil," Norton said, "be-cause it's both necessary and evil." He sees OS/2 being accepted mainly by "MIS types" in large corporations who want multitasking func-tions, "Unix-like features without throwing away DOS." Norton sees "nifty stuff" coming eventually for the OS/2 environment. The RAM shortage put OS/2 back "2 to 3 years"; other-wise, many users would have installed 2 to 4 megabytes of RAM by now and would have had the opportunity to "play" with memory-hungry OS/2, he said. As it stands now, "only the damn serious can afford to run OS/2, and the payoff isn't there yet."

Norton predicted that many of those "nifty" OS/2 applications will get "con-verted and crammed back into the MS-DOS environ-ment," although he added that those conversions will be a "painful process."

continued

Travel With A Fast Crowd.

World Radio History

disk cartridges. Supporting CCS, currently before ANSI and ISO for formal endorsement, are Advanced Micro Devices, Hewlett-Packard, Hitachi, Maxtor, Mitsubishi, Olympus Optical, Optotech, Scientific Micro Systems, and Western Digital.
• In a joint deal between **Borland** International and **Digital Research**, Inc. (Monterey, CA), OEMs can offer their customers certain Borland programs to go along with DRI's **DR DOS**. Computer makers who use DRI's enhanced DOS in their systems can also throw in Borland's SideKick, Turbo Pascal, Turbo C, and Turbo Basic.
• The world of **three-dimensional spreadsheets** will get a little more crowded this spring when **FormalSoft** (Sandy, UT) starts shipping its new ProQube program. The package's features include slicing across multiple spreadsheets, three-axis presentation graphics, conditional embedded external application calls, and ASCII file macros. The design uses virtual memory to extend large worksheets out onto the disk. The company plans to sell ProQube for $247.50.
• Remember that this month we celebrate **National Engineers Week** (February 19 to 25). The theme of the week is "Turning Ideas into Reality," and the societies that are sponsoring the idea would like us all to honor the engineer "as innovator and problem solver." Just to clarify things, we're not talking about the ones who drive trains.

Wescon Highlights: Toshiba's Big LCD, Eighteen Eight's Array Processor

It doesn't have the glitz and critical mass of COMDEX, but the IEEE's Wescon electronics exhibition is no small potatoes. This year's show in Anaheim featured some 6000 exhibitors and an estimated attendance of nearly 60,000. If you're shopping for transducers, printed circuit board design software, digital waveform analyzers, and miniature bearings and IC testers, Wescon is the place to be.

Although most of the showcased products are aimed at designers and manufacturers of electronics components, there are always a few new products of importance to end users. For example, Toshiba unwrapped its new 640- by 480-pixel liquid crystal display for laptop computers. Unlike conventional double-layer displays, the M-ST black-and-white unit uses a single-layer, backlit design. It's only 20 millimeters thick and weighs about 700 grams (about 1.5 pounds). With a .33-mm dot pitch, the M-ST display has excellent contrast and will certainly attract the interest of laptop computer manufacturers. However, the M-ST is not cheap; "sample pricing" is $646. Toshiba said it is planning to ship the new display in volume this spring.

Another eye-catching product at the show was a floating-point array processor for IBM PCs and compatibles. Eighteen Eight Laboratories claims its PL800 array processor delivers 8 million floating-point instructions per second for a price of $1995. You can install up to eight PL800 processor boards in a single machine. The PL800 comes with a library of 473 microcoded functions and supports most popular FORTRAN, C, and Pascal compilers, including those from Microsoft, Borland, Lattice, Lahey, and Ryan-McFarland. The 8-bit board requires only about 5 watts of power.

For more information, contact: Toshiba America, 9775 Toledo Way, Irvine, CA 92718, (714) 455-2000; Eighteen Eight Laboratories, 771 Gage Dr., San Diego, CA 92106, (619) 224-2158.

Simulation System Could Cut Costs of Real-Time Design

Athena Systems (Sunnyvale, CA) has a new computer-aided software engineering (CASE) and simulation tool that could cut down the costs of designing real-time systems. The graphics-based Foresight, which runs on Sun workstations, is a sophisticated simulation program that uses the concepts of block diagrams and data flow to represent physical processes.

While CASE tools supporting Yourdon/DeMarco methods are available, as are both discrete and continuous systems simulators, Foresight combines them.

The program also implements real-time extensions to CASE methodology. The developers claim that Foresight will change embedded systems engineering in much the same way that CAE tools have changed the design of ICs and printed circuit boards.

"Foresight helps uncover errors during the early requirements definition stage—where most design flaws are introduced—before hardware and software engineers proceed with their development tasks," said Patrick Rickard, president of Athena. "It's essential to correct errors early because the cost to repair them increases exponentially as the project progresses," he added.

Development of embedded real-time systems, which are extremely complex, is notorious for cost overruns and schedule slippage. In the aerospace and defense industries, for example, where embedded systems are used extensively in applications such as aircraft guidance, electronic switching, and weapons control, the reliability of software and hardware in the real world is a very major concern.

Foresight, by personal computer standards, is not cheap by any means; it costs $23,680 per user for a 10-user license.

For more information, contact: Athena Systems, 139 Kifer Court, Suite 200, Sunnyvale, CA 94088, (408) 730-2100.

TECHNOLOGY NEWS WANTED. *The news staff at BYTE is interested in hearing about new technological and scientific developments that might have an impact on microcomputers and the people who use them. If you know of advances or projects relevant to microcomputing, please contact the Microbytes staff at (603) 924-9281, send mail on BIX to Microbytes, or write to us at One Phoenix Mill Lane, Peterborough, NH 03458. An electronic version of Microbytes, which offers a wider variety of computer-related news on a daily basis, is available on BIX.*

Announcing a big leap in mouse technology.

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

A mouse ballistically 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 on detail work.

*T*his mouse is <u>guaranteed</u> to work with all applications on your IBM personal computer. And it comes with a great selection of MouseWare™ including Pop-Up DOS™—the ultimate DOS handler; Mouse-2-3,™ the Lotus 1-2-3™ shell; and 35 menus for best-selling keyboard-based applications.

*T*he Logitech Mouse. $139, complete with Logitech's Lifetime Satisfaction Guarantee and unlimited Product Support.

*F*or your nearest dealer, call:

800-231-7717

In California:
800-552-8885

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

Circle 133 on Reader Service Card
(DEALERS: 134)

LOGITECH

LOGITECH MOUSE

LETTERS

The High Cost of RAM

The unkind words about OS/2 ("OS/2's Multitasking Dashboard" by Mark Minasi, November 1988) are a classic example of blaming the victim.

Microsoft must have had to make its final decision on memory usage a couple of years ago. Back then, the price of RAM was about $100 per megabyte, with a prospect of further decline by now to, say, $25 to $50 per megabyte (the September 1988 BYTE had ads for 150-nanosecond 256K-byte chips at $2.95 each, or $106.20 per megabyte). There was an intense ferment over all sorts of ways to use more RAM—RAM disks or Expanded Memory Specification memory (not one, but two competing varieties). On that basis, it must have been almost impossible to justify less than 4 megabytes for any operating system that broke the 640K-byte barrier.

What happened was this: Largely due to protectionism, RAM prices rose to $400 to $500 per megabyte. Don't blame Microsoft; blame the politicians. And bear in mind that OS/2 is by no means the only computer product injured by high RAM prices.

Andrew D. Todd
Springfield, OR

Just Don't Get *Too* Excited

I'm a novice computer user, and BYTE gives me a lot of pleasure.

I'd like to offer a comment on Brock N. Meeks's "Computer Conferencing Homecoming" (September 1988). This is the kind of article that turns my pacemaker on high. Maybe in the future, Mr.

Meeks will be able to keep people like me in mind and write a piece that could help us get started on BIX, MIX, and CompuServe.

Reg. Roberts
Costa Mesa, CA

The Old Spool Tie

I cannot let Jud McCranie's authoritative-sounding correction of Brett Glass's spool definition (Letters, October 1988) stand unchallenged.

In the interest of historical accuracy, spool is indeed an acronym for "simultaneous peripheral operation on-line." Mr. McCranie states that his somewhat literal interpretation of the word *spool* dates back to "the old mainframe days." Just how old is Mr. McCranie, anyway? In the old mainframe days (and now), spooling was usually a high-priority task involving direct-access devices (disks and drums), rather than reels (or spools) of tape. The practice of sending print output to tape, although common, has little to do with the original intent of spooling.

As for the acronym being a "recent ad hoc creation," there is a terrific explanation of spooling in Harry Katzan Jr.'s excellent book *Operating Systems: A Pragmatic Approach*, published in 1973. Professor Katzan observes that "spooling was the first stage of multiprogramming as we know it today."

Joe Riley
Los Angeles, CA

True Meanings

Dennis Lee Bieber misrepresented the true meaning of the names of the P and V semaphore primitives (Letters, October 1988). According to professor E. W. Dijkstra himself, while teaching a course entitled "Operating Systems Techniques," he derived the names P and V from two Dutch verbs, "pakken" (seize) and "vrygeven" (release). Even with the coincidental starting letters, "procure" and "vacate" capture the spirit nicely.

Signal and wait functions are typically performed on events for process syn-

THE FOXBASE+ EFFECT

Nothing gets through your workload faster than FoxBASE+. Because no other database package has more timesaving features.

Speed is something every manufacturer claims. But we'd quickly point out that FoxBASE+ outpaces every other package in popular tests by margins so staggering, you'll wonder how you ever managed with ordinary software. In fact, FoxBASE+ 2.10 is the fastest dBASE-compatible database system.

Most publishers claim "compatibility" too. But the fact that you often have to re-write parts of your program is curiously omitted. That's why you'll appreciate FoxBASE+'s ability to run all your dBASE applications immediately – without changes.

FoxBASE+ not only supports the dBASE standard, but also allows you to enhance your applications with special features which include two-dimensional arrays, user defined functions, filtered indexes and new commands to create pop-up and pull-down menus.

Best of all though, is the term "user friendly". Often a misnomer for weeks of training with mind numbing manuals. Hardly easy.

That's why FoxBASE+ includes features that enable you to get working right away.

FoxCentral is a new type of user interface. Using pull-down and pop-up menus, you can perform every important database operation without programming!

FoxView is a screen design tool that helps you create beautiful custom data input screens.

FoxCode is our new applications generator that can easily produce the simplest or most sophisticated applications for you.

You can even document your programs automatically with **FoxDoc** by simply pressing a few keys.

FoxBASE+ is available in a variety of versions for the most popular operating environments including MS/PC DOS, Macintosh and UNIX/XENIX. There's even a 386 version! FoxBASE+/LAN supports an unlimited number of users on a network (No LAN PACKS required). Our Unlimited Royalty-Free Runtime allows extremely economical distribution of an unlimited number of your applications without incurring any royalty fees.

And if you want to add another dimension to your data use FoxGraph, the exciting new graphics package from Fox.

Take a look at FoxBASE+ today. You won't find a more powerful dBASE-compatible relational database management system. FoxBASE+ is retails for $395, FoxBASE+/LAN for $595 and FoxGraph for $295 at your local quality software dealer. Or call 419/874-0162 for more information and a free demo-package. Ask for extension 320.

dBASE is a registered trademark of Ashton-Tate.
FoxBASE+ is a trademark of Fox Software Inc.

Fox Software
Nothing Runs Like a Fox.

Fox Software Inc. 118 W. South Boundary Perrysburg, OH 43551 Phone #: 419/874-0162 FAX #: 419/874-8678 Telex #: 6503040827 FOX
Fox Software International Intech House Cam Centre Wilbury Way Hitchin Herts SG4 0AP Tel: 0462 42 1999 Fax: 0462 42 13 18

chronization. They can't be used for mutual exclusion of critical regions of code, like the semaphore primitives, because executing a signal for an event will restart *all* processes waiting on that event.

Lex Borger
Mission Viejo, CA

Multiuser Advantages

In your *IBM Special Edition* (Fall 1988), there is considerable discussion about the roles of OS/2 and local-area networks (LANs) in workplace computing. While I found these discussions interesting and enlightening, I am puzzled by the lack of concern that OS/2 is not a multiuser operating system.

To quote from the text box "OS/2 versus Unix: Is DOS Compatibility the Key?" by Jason Levitt on page 112, "OS/2 is not a multiuser system, but advances in networking and distributed software will eventually make this a moot point." This view ignores an important advantage of multiuser systems: the ability to function as a "compute server" for a workgroup.

Consider a small academic department that needs access to a large computer for number crunching but that also does a lot of word processing and smaller tasks. To facilitate the numerical work, the department might buy a large 80386 machine. Resources wouldn't permit buying a separate 80386 machine for each user but would provide for individual basic IBM PC AT clones and a low-volume LAN. This arrangement would not permit sharing the 80386 machine because the operating system (DOS or OS/2) wouldn't permit remote log-ins from LAN stations. Hence, to share the 80386 machine, you'd need to sit at the keyboard attached to the machine. This is inefficient and awkward. The 80386 could be used as a file server, but then each of the remote stations would need the hardware resources to run the analysis software, again increasing costs.

A multiuser operating system would permit the use of the 80386 machine as a remote "compute server." Even if the system were set so that only one person could log on at a time, this would be a major improvement over the current options with OS/2. I hope Microsoft will consider this type of situation in its future plans for OS/2.

Nicholas Birkett
Ottawa, Ontario, Canada

Short Dispute

In your Product Focus entitled "80386s for the Masses" by Steve Apiki and Stan-
continued

Circle 180 on Reader Service Card

What looks like DOS, acts like DOS, is easy like DOS, can hook everybody up and keep 'em hopping from program to program,

and costs less than an office chair?

The Comfortably Affordable (and already available) PC-MOS.™

If you can afford a nice leather chair for your office, you can afford a multi-user, multi-tasking operating system that's just as comfortable to use.

PC-MOS incorporates break-through technology that allows you to fully harness the power and capabilities of 80386-based computers, while maintaining the compatibility and ease of use of DOS. So now you can have up to 25 inexpensive terminals driven by a single 80386 processor. You can continue to use your favorite software programs like Lotus 1-2-3,™ dBASE III,™ and WordPerfect,™ and you can keep using those familiar DOS commands like DIR and COPY. You can even run all those terminals as a single cluster connected to a Novell server.

PC-MOS comes in single, five and 25-user versions starting at $195. Your satisfaction is guaranteed or you get a complete refund.

For more information about PC-MOS and the name of the authorized dealer nearest you, call 1-800-451-LINK. Then you can sit back and relax.

THE SOFTWARE LINK

3577 PARKWAY LANE • NORCROSS, GEORGIA 30092
INSIDE GEORGIA: 404-448-5465 • FAX: 404-263-6474
INTERNATIONAL CALLS: 404-263-1006

PC-MOS is a trademark of The Software Link, Inc. All other brand names are the trademarks of their respective holders.

World Radio History

ford Diehl (October 1988), we were pleased to note that our Zeos 386 Tower was the fastest 16-MHz machine reviewed. We appreciate your many favorable comments as well.

One criticism of all systems reviewed was that the connecting cables are too short. The Zeos tower-style configuration was noted as having shorter cables. This is not the case.

Zeos normally ships, at no additional charge, a full set of extension cables with each 80386 tower-style system ordered. This includes monitor power, signal, and keyboard extension cables. This standard extension cable set provides plenty of length to work with.

Zeos shipped BYTE two systems for review. The first system was completely standard and included extension cables. Unfortunately, it was damaged in shipment. We replaced the cabinet with a second one but didn't send another cable set. BYTE already had the extension cables and the other system components.

Zeos makes every effort to provide its customers with complete ready-to-run systems—including extension cables with our tower-style systems.

Greg Herrick
President, Zeos International
St. Paul, MN

Factoring the Unforeseen

I would like to contribute to the recent "Minds vs. Programs" correspondence (Letters, September 1988).

In the late 1960s in England, I developed one of the first expert systems, although we did not use that expression then, to analyze predigitized optical data from an elementary particle physics experiment. The program was developed on an IBM 360/75 in FORTRAN, which was about the only language available at that time. I played a major part in its development, especially in the pattern recognition.

Eventually, it achieved a 99.9 percent pattern-recognition efficiency for single particle tracks. It also had to have a very high background recognition efficiency, because the signal-to-noise ratio in the data was 1 to 2200. We found it hard to make this better than 1 to 100 and were forced to plot the data on 16mm film (using a Ferranti Atlas) and scan the data by eye to eliminate the rest of the noise. To maintain both a high pattern-recognition efficiency and a high background rejection efficiency, I developed a hierarchical system of self-adjusting selection criteria—in effect, rules that changed themselves in a systematic way if they failed to succeed.

This work gave me considerable insight into the problems of computer pattern recognition. My conclusions are as follows:

- The human intellect is capable of holding an abstract representation of what it is seeking in its consciousness and of comparing that with what it has found.
- The human intellect is capable of detecting when its algorithms or techniques have not yielded the desired solution.
- The human intellect is driven by emotions, which a computer does not have. It can generate new algorithms to achieve its goal, new goals if the previous one is in error, or even new emotions to create a more satisfactory reality.
- A computer lacks this fundamental consciousness and creativity.

continued

    Circle 36 on Reader Service Card

In effect, the human being is like an infinitely sophisticated, self-conscious, self-programming computer. We know, for example, that we are seeking peace and will go on adjusting our search until we have found a stable and lasting peace, for otherwise we shall be destroyed. A computer would be incapable of this self-regulation. It would execute its algorithms until it self-destructed: It would not be motivated to avoid self-destruction or capable of reprogramming itself to avoid self-destruction.

There are people who would invariably argue that this is not so, that a computer could be programmed with this self-motivation and reprogramming capability. But even if that were so, the computer would not be capable of generating the required emotional response to a problem whose existence the programmer had not foreseen, because it would not understand meaning. This is the difference between consciousness and ma-chine. The key word is *programmed*. A program has the consciousness of the programmer frozen in it. It is crystallized consciousness, rather than living consciousness.

Dr. R. J. Ellis
Palo Alto, CA

FIXES

• In the article entitled "The Promise of Project Management" by Lamont Wood (November 1988), we accidentally omitted the address for the consulting firm One Soft Decision, Inc. One Soft publishes PM Solutions, an in-depth report on project management software. For more information, contact Dan Yahdav, One Soft Decision, Inc., 573 Wakerobin Lane, Suite B, P.O. Box 6123, San Rafael, CA 94903

• Our November 1988 Short Take on the NEC Ultralite laptop computer failed to mention some key points about the unit we reviewed. As the article states, the machine was an engineering prototype. We did not explicitly mention, however, that the ROM disk hardware was not fully implemented. It could not work, period. The LapLink and DOS Manager software resides in a ROM chip in the laptop, not on ROM cards, as the article says. We also omitted the fact that the Ultralite comes with a serial cable to use with the 9-pin, DIN-style connector.

As an update, NEC has told us that a number of popular applications software packages have been converted to the ROM card format, including Lotus 1-2-3, Agenda, WordPerfect 5.0, WordStar 5, Microsoft Works, XyWrite III, and NEC's Telcom 2.0 communications software.

Finally, the Short Take compared the Ultralite's list price to a Toshiba T1000's street price. The T1000's list price is $1249.

• Our PostScript Printer Product Focus (September 1988) gave incorrect scores for the GCC Business Laser Printer for graphics and text quality. The numbers for those categories in table 1 on page 166 should read 3.255 for graphics quality and 3.318 for text quality. BYTE regrets the error.

• Two photographs were inadvertently swapped in the What's New section of the October 1988 issue. The photograph on page 78 is actually of a LabView 2.0 screen, as described on page 80. The photograph on page 80 is a screen from Absoft's MacFortran/AUX, which is described on page 78.

Now, with NRI, you can learn to program in today's hottest computer languages— BASIC, Pascal, C, and COBOL

You start with BASIC, an extremely popular language used for graphics, teaching, personal, and small business applications.

Next, you master Pascal, a multipurpose language often used to develop software for computer-aided design (CAD).

You move on to C—hands down, today's hottest computer language—used for everything from systems programming to artificial intellience.

You wrap up your training in programming with COBOL, the language of business and the most widely used computer language in the world.

NRI's new at-home training in Computer Programming starts by walking you step by step through the fundamentals, easing you into programming with brilliantly detailed instructions, charts, and diagrams.

In no time at all, you have a complete understanding of the programming techniques used every day by successful micro and mainframe programmers. And then the fun really begins.

With your personal NRI instructor on call and ready to help, you use the computer system included in your training to actually design, code, run, debug, and document programs in BASIC, Pascal, C, and COBOL. Then, following easy-to-read instructions, you use your modem—also included—to "talk" to your instructor, meet other NRI students, even download programs through NRI's exclusive programmers network, PRONET.

Send for your FREE catalog today

For all the details about NRI's at-home training in Computer Programming, send the postage-paid reply card today. Soon you'll receive NRI's fascinating, information-packed, full-color catalog.

Open it up and you'll find vivid descriptions of every aspect of NRI training. You'll see the IBM-compatible Packard Bell computer included in your course up close in a special, poster-sized foldout section. And, best of all, you'll find out how your NRI training will make it easy for you to build a high-paying career—even a business of your own—in computer programming.

If the card is missing, write to NRI at the address below.

IBM is a registered trademark of the IBM Corporation

World Radio History

CHAOS MANOR MAIL

*Jerry Pournelle answers questions about his column
and related computer topics*

One Alternative to OS/2

Dear Jerry,

I use an IBM PS/2 Model 80 for development work as well as in my normal, day-to-day activities. It consists of a 70-megabyte hard disk drive, 2 megabytes of system memory, a 3½-inch internal floppy disk drive with drives A and B, an Identica internal tape backup using SYTOS software, a 5¼-inch external SYSGEN floppy disk drive for transporting files and programs to and from the other systems, and an IBM color monitor.

It was with this Model 80 that I encountered problems. I wanted to take advantage of the speed and memory available by using a multitasking operating system called PC-MOS/386 from The Software Link. My objective, therefore, was to partition the hard disk as three logical drives (C, D, and E); partition the 2-megabyte system memory as three virtual machines to allow three programs to run simultaneously under the PC-MOS/386 operating system; use the 5¼-inch external disk drive to maintain compatibility with other IBM PC, XT, and AT computers and clones through the SYSGEN Bridge device driver; and use the internal tape drive to regularly back up and restore logical drives C, D, and E using SYTOS software.

The SYTOS tape backup software, however, would not run. Calls to The Software Link and Identica confirmed that the two could not work together. My system operated under PC-MOS, but SYTOS would work only under DOS.

I called a local turnkey systems house familiar with PC-MOS; the people there had experienced a similar problem on a Compaq 386, but they had no experience with a Model 80 or with the SYSGEN external drive. They explained that they had solved their problem with Disk Manager from Ontrack Computer Systems in Eden Prairie, Minnesota. I called Ontrack, and, with no guarantee of success, I bought the company's Disk Manager.

First I copied all my files to 3½-inch floppy disks, since all the data on my hard disk would be destroyed. Then, booting with DOS in drive A, I reformatted the hard disk under DOS 3.1 with one bootable DOS partition and one extended DOS partition on the hard disk, which was further split into two logical drives. I copied all DOS files from drive A to drive C.

The CONFIG.SYS file was as follows:

```
DEVICE=BRIDGE.DRV /PS60:2
DEVICE=ANSI.SYS
FILES=15
```

I then ran the SYSGEN installation program from drive A, telling SYSGEN that there were two 3½-inch internal drives. I copied all SYSGEN files from drive A to drive C.

Although I had not run Disk Manager's installation yet, I copied all Disk Manager's files to drive C. I then rebooted the computer from drive C. During the boot, a message was displayed announcing that the SYSGEN drive was installed as drive F.

At this point, I formatted 10 3½-inch disks, adding the DOS system to only one disk. Since this disk would be the DOS boot disk, I copied all files on drive C to this disk.

Although the CONFIG.SYS file on drive A could have been edited, I chose to create a new configuration file for drive A:

```
DEVICE=DMDRIVER.BIN
DEVICE=BRIDGE.DRV /PS60:2
DEVICE=ANSI.SYS
FILES=15
```

Since the SYSGEN external drive was now "live," I copied all the PC-MOS/386 software from drive F to drive C and to another disk in drive A (not the DOS boot disk). I also transferred several

---

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."

other files from drive C to the disk in drive A by typing the following at the C > prompt:

```
MSYS A:
COPY DMDRVR.BIN A:
COPY ANSI.SYS A:
COPY BRIDGE.DRV A:
```

At the A> prompt, I wrote a new CONFIG.SYS file for drive A:

```
DEVICE = DMDRVR.BIN
DEVICE = BRIDGE.DRV /PS60:2
DEVICE = ANSI.SYS
MEMDEV = $386.SYS /p
FREEMEM = C4000,C8000
FREEMEM = CA000,E0000
DEVICE = $SERIAL.SYS
SMPSIZE = 256K
BUFFERS = 22
```

Since the disk in drive A was the PC-MOS/386 boot disk, I rebooted from drive A under PC-MOS/386.

To prepare the hard disk for PC-MOS/386, I ran HDSETUP from drive A, allowing PC-MOS/386 to create three logical MOS disks. I then transferred the PC-MOS/386 system to drive C. At the A > prompt, I typed

```
MSYS C:
COPY *.* C:
```

Rebooting now from drive C restarted the machine under the PC-MOS/386 operating system. The next step was to prepare the hard disk with Disk Manager.

With the Disk Manager software in drive A, I typed DM/C/M to invoke Disk Manager in the color, manual mode. Looking at the partition information, the partitions were as follows:

```
1   0  30  DOS  Y  TSL 1.01
2  31  61  DOS  N  NOSYSTEM
3  62  69  DOS  N  NOSYSTEM
```

I changed partitions 2 and 3 to write/read. Exiting from Disk Manager and allowing the changes to be saved to disk, I received a MOS error message, which I could ignore.

Finally, the tape backup software, SYTOS, was installed. After all, the goal of all this was to allow SYTOS to function. To install SYTOS, I ran the installation program from drive A. Since SYTOS must be present on the hard disk, I copied all the files from drive A to a subdirectory named SYTOS on drive D.

This is working fine for me. Booting from drive C under PC-MOS/386 allows multitasking and other PC-MOS/386

functions. Under PC-MOS/386, my system will run dBASE III Plus, WordStar, Multiplan, BASIC, Quickcode, Quicksilver, Chart-Master, Norton Utilities, and other programs. The system will not run Microsoft's Excel. Although I haven't tried it, Microsoft's Windows probably will not run, either. Booting from drive A under DOS using a 3½-inch DOS boot disk will allow the SYTOS tape backup to run properly. You can also run any other software not supported under PC-MOS/386, such as Microsoft's Excel.

Arthur J. Foley
Central Islip, NY

Thanks for the report. You do seem to have found one alternative to OS/2. It would be interesting to see if Windows/386 would run with this system; probably not.

I don't have a Model 80. —Jerry

N.B.: Nota Bene

Dear Jerry,

You once said that you just might move up to XyWrite III Plus after all. That would be a good move. But might I suggest a better one? Move a bit further on and get Nota Bene.

First, some history. They say you never forget your first love or your first word processor. I haven't forgotten my first word processor, or any of the others, but I truly love my current one, Nota Bene. It's extremely powerful, smooth, fast, and (most important) logical.

As you may know, Nota Bene is a customized version of XyWrite. Most commands in Nota Bene can be executed in three ways: by typed-in commands, by function keys, and by menus. Most operations can be initiated by short mnemonic command codes. To give you an example of how easy the command codes are, search is se, search backward is seb, absolute search (where case is matched) is sea, and search backward absolute is seba.

You can throw away the function key template. (I dare you to do that with WordPerfect.) The key use is supremely logical. For example, the Control key works on words; the Shift key on sentences; Control-Shift on phrases (i.e., material between punctuation); Alt works on paragraphs; Shift-Alt works on lines. These keys in combination with other keys will move the cursor, delete, highlight, interchange text, and so on—all operating on the same units. It is easy to learn these operations in minutes. It is my understanding that Nota Bene, be-

cause of its logical design, is infinitely easier to use than XyWrite.

Nota Bene is designed for academic use. It has all kinds of footnoting, indexing, formatting, and printing capabilities (there are three disks filled with printer drivers). It has a text base that lets you index everything you write and recall it with a variety of Boolean-type queries. It has its own programming language. It has six alternate keyboards that you can instantly access: You can type in a variety of languages, including Greek; you have a keyboard of math symbols. You can also create and load your own keyboards.

The bad news about Nota Bene is that the manual is over 1000 pages long. The good news is that the program is so logical that you can use it without reading the manual.

With a 30-minute introduction at the keyboard and with the reference booklet, you can do very well. One reason is that there are also a menu structure and help file that are seamlessly integrated. If you don't remember the commands, you can call up the menus.

Several weeks ago, my 15-year-old son saw me writing with Nota Bene and asked for a run-through. I gave him a 15-minute overview, thinking that it was enough to satisfy his request and to get him turned back to his own word processor. Later, when I got off the computer, he sat down and used Nota Bene to write a fairly long essay that was due the next day. For sure, I had to do some prompting, but he did a fine job. He has used Nota Bene ever since (with little help from me) and loves it. He doesn't do anything fancy, but the point is that he can navigate it alone with his relatively simple documents. That tells you a lot about a word processor that is so powerful.

Why am I telling you this? Because Nota Bene is really good, and the people at Dragonfly Software deserve success. (Incidentally, the company's technical support is excellent.)

I have introduced dozens of people to Nota Bene. I have never encountered a single person who did not switch to it after he or she played with it for a bit.

Joseph M. Prospero
Miami, FL

Thanks for the report.

You aren't the only one who's said good things about Nota Bene, and indeed, I've already recommended the program to a young person writing a dissertation.

Alas, I haven't tried it myself, because I don't have it. —Jerry ∎

Circle 41 on Reader Service Card

World Radio History

PRESENTING AMERICA'S HOTTEST NEW CORPORATE JETS.

FIVESTAR Computers.

Taking companies higher. And farther.

It's no wonder hundreds of corporations across the country are now flying high-performance computers from FIVESTAR. That's because FIVESTAR delivers the quality and power companies need to soar to new levels of capability - while keeping hardware costs down to earth.

A ONE YEAR LIMITED WARRANTY—and ON-SITE SERVICE.

Every FIVESTAR Computer comes with a ONE YEAR LIMITED WARRANTY. For companies with critical applications, a comprehensive ON-SITE SERVICE CONTRACT is available for just $99. In most cases, service calls are made within 24 hours.

13518

The FIVESTAR 286/10.

The high performance computer that flies through complex business applications.

CPU/MEMORY: 80286 16-bit microprocessor • 6/10MHz user selectable • 0 wait states • 13.3MHz LANDMARK throughput rating • Socket available for 80287/10 math coprocessor • 512KB RAM expandable to 4MB RAM • Supports LIM/EMS 4.0 on the motherboard

BIOS: AWARD BIOS • Set-up utility in firmware • Fully supports DOS and OS/2

EXPANSION SLOTS: Eight full-length slots: two 8-bit, six 16-bit

PRICE: $995

The FIVESTAR 386/16 SX.

A faster corporate jet for companies running on tight schedules - and tight budgets.

CPU/MEMORY: 80386-SX 16/32-bit microprocessor • 8/16MHz user selectable • 0 wait states • Socket available for 80387/16 math coprocessor • 1MB RAM expandable to 4MB RAM (on motherboard) • Uses 256KB or 1MB DRAMS

BIOS: AMI BIOS • Set-up utility and diagnostics in firmware • ROM BIOS and video BIOS in shadow RAM • Fully supports DOS and OS/2

EXPANSION SLOTS: Eight full-length slots: two 8-bit, six 16-bit.

PRICE: $1,495

ASK BYTE

BYTE technical editors answer your questions on microcomputing

Proper Mindset

I recently came into a Mindset computer. The operating manual is so deficient that I cannot use it. Where can I get a workable technical manual? Also, I need a memory map of the computer and its expansion unit. A schematic would be nice, too.

C. Bradner Brown
Kensington, MD

Your best bet would be to link up with a computer group in your area (check Computer Shopper *for a list of club meetings) or to investigate the orphan computer special-interest groups on CompuServe. Barring that, I believe a Mindset computer is still on display in the New York Museum of Modern Art as an example of innovative package design. Perhaps you could see if the museum needs a backup.—R. G.*

Incompatible Drive?

I have several Shugart Associates Model 450R floppy disk drives that have been used on a computer for writing to Unix, MS-DOS, and PC-DOS floppy disks. I don't seem to be able to jumper any one of them so that it will be accepted as a valid disk drive on my IBM PC AT clone, which is a PC's Limited system. I've tried listing the drive on my setup menu as a 360K-byte disk drive. Is it possible to use these disk drives on my machine?

M. W. P. Strandberg
Cambridge, MA

The Shugart SA450 should work with your system as a 360K-byte drive, provided that your controller can support an additional drive set to 360K bytes. My guess is that you have improperly jumpered the drive select lines. These are controlled by the jumper header labeled 2D on the floppy disk drive's printed circuit board.

To enable drive select 0, short pins 2 and 13; to enable drive select 1, short pins 3 and 12. If your floppy disk drive cable is like most, it has a "twist" in it between the first and second floppy disk

drive connectors, in which case you need to jumper both drives for drive select 0. (Nothing's easy anymore, is it?) You'll also want to cut the jumper between pins 6 and 9, which would otherwise cause the drive to respond to any drive select signal. Finally, make sure that only the last drive on the daisy-chained cable has a terminating resistor pack installed (this is IC 3D on the 405). Good luck.—R. G./S. W.

More Lines, Please

After my pleasant experiences writing EGA and VGA graphics routines using QuickBASIC 4.0, I was disappointed to find that Microsoft C version 5.1 (running under DOS) offers no obvious way to alter the number of text lines on a graphics display. I suspect it's just a matter of making the correct DOS call with the right parameters. Am I correct?

John J. Ottusch
Malibu, CA

Richard Wilton is an excellent source for the kind of information you need. I suggest you check the following articles and books:

"PS/2 Video Programming," BYTE's Inside the IBM PCs, Fall 1987

"VGA Video Modes," BYTE's IBM Special Edition, Fall 1988

Programmer's Guide to PC and PS/2 Video Systems *(Microsoft Press, Redmond, WA: 1987)*

In particular, Wilton's BYTE articles cover the topic of changing the number of text lines on the display.—R. G.

FFT Woes

I am a senior electrical engineering student at the Milwaukee School of Engineering. One of the requirements for graduation is the successful completion of a full-year design sequence that results in a working prototype.

I am working on a "digital demodulator" project that is based on an article in the September 1986 *IEEE Transactions on Instrumentation and Measurement*. It's essentially a radio receiver that detects the information on a carrier using digital signal processing techniques instead of conventional analog demodulation circuitry.

Part of the project will involve using the fast Fourier transform for spectral analysis, but apparently the FFT is a real time burner. I need information on FFT algorithms and, in particular, on implementing FFT algorithms in assembly language. In addition, do you have any information on predicting FFT response time?

Eric R. Schumann
Milwaukee, WI

Steve Ciarcia's "Computers on the Brain" articles (starting in the June 1988 BYTE) cover the FFT algorithm in some detail. Also look into "Faster Than Fast Fourier" (April 1988 BYTE).

E. Oran Brigham's The Fast Fourier Transform *(Englewood Cliffs, NJ: Prentice Hall, 1974) is an exhaustive exploration of Fourier transform techniques as well as the FFT. It contains the source code for the FFT algorithm (in ALGOL and FORTRAN).*

Finally, see if you can locate a copy of Minicomputer Systems, Organization, Programming, and Applications *by Richard H. Eckerhouse Jr. and L. Robert Morris (Englewood Cliffs, NJ: Prentice Hall, 1979). It contains a brief discussion of the FFT but includes PDP/11-compatible assembly language source*

code for an FFT butterfly. You might be able to use this code as a jumping-off point to whatever processor you're using. —R. G.

Good Old Model III

I have several questions. First, where can I get inexpensive parts for project construction? Second, I have an old Tandy TRS-80 Model III, and I wonder if its Z80 CPU will work with the programs in Steve Ciarcia's book. Finally, the Model III is 8 years old. I'd like to know if there's any way I can upgrade it.

Michael L. Hudin
Oroville, CA

Check out the advertisements in the back of BYTE and other computer and electronics magazines for distributors of parts.

I'm not sure to which of Steve Ciarcia's books you are referring, so I can't advise you about software. The TRS-80 Model III was a fine computer for its time. Tandy still supports the machine, and you can order more RAM, floppy and hard disk drive kits, and an RS-232C board from your Tandy Computer Center or Radio Shack dealer. Another source for parts is Aerocomp (2544 West Commerce St., P.O. Box 223957, Dallas, TX 75212, (214) 637-5400). —S. W.

Teaching an Old BIOS New Tricks

I have a 6-year-old IBM PC that has 16K-byte RAM chips on the motherboard. I want to increase the base memory to 640K bytes through the use of an Everex board, but I am told that my ROM BIOS will not allow me to get above about 545K bytes. I need updated ROM chips, but IBM no longer supports them, and I haven't been able to find them locally. I hardly think I'm the only one in this position; there must be many such machines still in use. Can you refer me to a source for the chips I need?

W. E. Van Horne
Columbus, OH

When IBM first introduced the IBM PC, the original BIOS limited system memory to 544K bytes. IBM never expected that anyone would want more RAM. Upgrade BIOS ROM chips are available from Mentor Electronics (7560 Tyler Blvd., Suite E, Mentor, OH 44060, (216) 951-1884). —S. W.

More Interface Cables

I am an electrical engineering student at the University of Illinois. I have finally decided to construct several of the projects presented by Steve Ciarcia, and I

want to interface them to my Macintosh SE. What can I do to make the standard RS-232C connections match the RS-422 mini DIN-8 connector on the back of my computer? I also wonder about how this nonstandard input via the modem port affects the communications protocol and toolbox interface with the Microsoft FORTRAN and Turbo Pascal environments with which I am familiar.

Last semester, I worked on a sleep-monitoring project through the College of Medicine, and I'd like to continue with some research on my own with the HAL EEG presented in the June 1988 BYTE. I understand that some time ago, Steve Ciarcia wrote an article describing the measurement of other bodily functions. In which issue of BYTE can I find this article?

Peter Apostolakis
Long Grove, IL

The RS-422A standard arose out of the need to send serial signals at higher rates and through longer lengths of cable. It does this by using differential voltages to provide noise immunity. However, the RS-422A standard was also designed to be electrically compatible with the older RS-232C standard. This is done by picking off the unbalanced negative portions of the transmitted and received data signals. This reduces the problem of obtaining RS-232C-compatible signals from the Mac SE's serial port to a matter of building the appropriate cable. I've successfully built cables that connected a Macintosh serial port to a PC serial port for data transfers.

For a Mac mini DIN-8 to DB-25 RS-232C cable, you need to purchase connectors for the cable. You can get the DB-25 connector from any electronics supply shop. Getting your hands on a male mini DIN-8 connector is a lot harder, and definitely more expensive. You can get an Apple System Peripheral-8 cable, product number M0197, for $29 from an Apple dealer. This cable has a mini DIN-8 connector at each end. Cut the cable in half, pick the half to which you want to attach the DB-25 connector, and store the other half in a safe place. See figure 1 and table 1 for pin-out information.

The Mac has its own drivers that handle I/O through the serial ports. The extent to which you can send and receive data or manipulate control signals through these ports depends on how well the programming language has implemented its interface to the Mac's serial drivers. Check the details in your programming language manual. You can

World Radio History

Figure 1: *DIN-8 connector pin-out diagram.*

Table 1: *Wiring the Macintosh DIN-8 to a DB-25 connector involves connecting the pins as shown. Note that the signal names are from the Mac's point of view.*

Signal Name	DIN-8 pin	DB-25 pin	Null modem function
HSK	2	20	DTR handshake for printer
TxD-	3	3	Null modem
GND	4	7	Signal ground
RxD-	5	2	Null modem

also look into Inside Macintosh, *volume II*, starting on page 243 for good information on the Macintosh serial drivers.

Steve Ciarcia's previous article on biofeedback, "Mind Over Matter," appeared in the June 1979 BYTE.—T. T.

I've Got This Spare...

I have an Amiga 1000 with 512K bytes of internal memory. I also have an extra Amiga Model 1050 256K-byte RAM cartridge, the one that plugs into the front of the machine. I'd like to know if there is any way I can modify the spare pack so that I can use it to increase my memory to 768K bytes.

R. James de Graff
Winnipeg, Manitoba, Canada

The odds are you'll do more harm than good trying to modify the board. I recommend getting one of the sidecar memory boxes that attach to the expansion connector on the right of the machine. Keep your spare cartridge as just that: a spare.

If you're still determined to hack your machine's memory, check out the article "Build Your Own 256K Amiga Expansion RAM" in the February 1987 BYTE.
—R. G.

Rays and Buses

I am interested in ray tracing, a very calculation-intensive computer application that is similar, in some respects, to the problem of generating pictures of the Mandelbrot set; it can be done using a multiprocessor board (like the one described by Steve Ciarcia in "A Supercomputer, Part 1" in the October 1988 BYTE).

I have an Apple Macintosh II, which uses the NuBus. I've been programming the Mac for years, but since the previous versions of the Mac had no slots, I know nothing about hardware. I hear that designing and programming NuBus cards is easy, but I don't know where to begin.

Could you suggest some reading material for programmers who, like me, have never soldered anything, let alone built any hardware? Could you also present an article describing how to build a simple NuBus test card for the Mac II? Or how about a multiprocessing board like the one Steve Ciarcia built to generate Mandelbrot sets?

Ajay Nath
Oakland Gardens, NY

Unfortunately, designing and programming NuBus cards is not easy. The software interaction with boards is complicated, and you need an intimate understanding of how the hardware operates. But the end result of all this effort is a card that you can plug into a Mac II without requiring you to tinker with switches on the board or run complex configuration applications.

As for your request for an article on this subject, I've got good news for you. The design of a simple NuBus test card is described in the Macintosh Special Edition (December 1988 BYTE). If you need additional information, you can consult the books Designing Cards and Drivers for the Macintosh II and Macintosh *by the Apple Computer staff (Reading, MA: Addison-Wesley, 1987) and* Macintosh Family Hardware Reference, *also by Apple Computer (Reading, MA: Addison-Wesley, 1988). These should at least give you an idea of what is involved in such a project, and sample code is provided to help you get started.*—T. T.

Vertical vs. Horizontal Drives

My AT takes up entirely too much room on my desk, and the fan and my new

MiniScribe 43-megabyte full-height hard disk drive make more noise than I'd like. I've been considering buying one of those vertical stand brackets to put my machine under or beside my desk. However, I know that quite often the bearings in various types of machinery are designed to carry their load in only one position and if turned 90 degrees can fail rapidly. Disk drives, both floppy and hard, appear to be designed for horizontal installation; and since the bearings are obviously small, they might be prone to early failure if mounted vertically. Floppy disk drives, due to their intermittent running, might not be as much of a problem as hard disk drives. Am I worrying needlessly?

Is it possible to use a program to stop the hard disk drive from running constantly? I don't want to sound like a nervous Nelly, but the idea of that expensive thing sitting in there spinning madly 8 hours per day while I type makes me wonder how long it can live. Having it run when data is actually being accessed is one thing, but if I'm just writing and only occasionally saving to disk, why have the hard disk spinning needlessly?

I notice that most of the new 80386 computers built with tower cases have the drives mounted horizontally. They also have their switches mounted on the front, where they belong. As much as having a new 80386 machine appeals to the techno-freak in me, I simply cannot justify buying one to get that slick tower case. Can you tell me where I can buy a tower case like that? Three companies—Fortron, Logix, and Zeos—have tower cases that appear identical and suitable. I contacted Logix, but the company won't sell bare cases. Surely these three obtain their cases from the same source. Do you know where?

Harmon Seaver
Grand Marais, MN

Some disk drives can withstand vertical mounting; others cannot. That information should be included in the specifications. If not, check with the disk drive manufacturer.

Your hard disk's read/write head is floating on a cushion of air—created by the spinning platter—against zero friction. In this environment, the drive will not be exposed to abnormal wear and tear. In fact, your drive will experience greater wear if it is continually stopped and started as you propose. The roughest treatment of your drive occurs at power on, and the culprit is not so much friction as heat. Drive manufacturers go to great

World Radio History

lengths to ensure proper heat compensation. Some of the new laptops shut down the drive motor when it is not in use, but they do this to preserve power, not to protect the drive. I do not recommend that you routinely shut down the drive motor.

I suggest you get a head parking utility. You can find these programs in the public domain. After a specified time of disk activity, the software will park your disk heads until a disk access is requested. A commercial product of this type is Safepark from Prime Solutions. Safepark positions the head over a selected sector that does not contain data. Power surges and sudden spikes can write random bits to your drive. If the bits are written to a sector where your data resides, that data is corrupted. Safepark moves the heads to the "safe zone" whenever the disk is not in use. I mention the product because it is included with a very useful disk utility package called Disk Technician.

Hard disks are most vulnerable to bad sectors. Some utilities, such as the Norton Utilities, will let you partition off bad sectors. Once you have done this, you can no longer access that sector. Disk Technician, on the other hand, performs preventive maintenance on the disk. It will monitor disk access errors and keep track of them. When DOS fails to read a sector, it will keep trying 30 to 50 times. If it finally does read the disk successfully, DOS goes on its merry way, not concerned about the access problem. Disk Technician will recognize the problem, differentiate between random errors and repeating errors, and then attempt to repair the error before it becomes a major problem. First it moves the data to a good sector, then it tries to repair the sector by performing a low-level format. If it can't repair the sector, it maps the sector out. I recommend the program for all those nervous Nellies who dread the thought of losing valuable data from hard disks.

Many independent mail-order companies sell the tower case. You just have to look around. I found one such case in a JDR Microdevices catalog (see the company's ads in the back of BYTE). The $299.95 cost includes a 250-watt power supply, speed display, mounting hardware, faceplates, and a speaker. Make sure that you also find some cables that are long enough to accommodate the tower design. We keep getting these wonderful space-saving tower systems with cables so short that we can't place them on the floor where they belong. It will be frustrating if you invest in the tower case but end up with a bulky system on your desktop anyway, so you should keep cables in mind. —S. D. ■

There you are sitting happily with your laptop in the sky, preparing some figures for an important meeting at the airport. Then, without warning, your screen flickers and goes dark. Your laptop just died.

Things could be better. If you'd been using Battery Watch you would have been prepared.

Battery Watch operates like a fuel gauge for your laptop. Its pop-up window lets you see what's left of your battery power in hours and minutes.

It'll also help you get more out of those sneaky NiCad batteries. The Deep Discharge feature makes sure of it. It eliminates the nagging "charge memory effect" of NiCad batteries and ensures their original peak performance.

Battery Watch has taken the worry out of something else, too. Compatibility. It's compatible with all popular laptops from Amstrad to Zenith.

And at $39.95, it's the cheapest insurance around.

Look for it at your local computer store. Or call 1-800-343-8080 for a free laptop accessory catalog. In Washington call (206) 483-8088.

Traveling Software

18702 North Creek Parkway   Bothell, WA 98011

Without Battery Watch, You Could Be At 40,000 Feet And Out Of Fuel.

Circle 98 on Reader Service Card

World Radio History

BOOK REVIEWS

No Way: The Nature of the Impossible

Edited by Philip J. Davis and David Park

*W. H. Freeman and Co.,
New York: 1987, 325 pages,
$10.95*

Reviewed by David A. Mindell

It is impossible to climb Mt. Everest without oxygen. It is impossible to make music at a rate greater than 10,000 bits per second. It is impossible to create a stable, strained-layer superlattice.

These statements refer to achievements that were once considered beyond the practical, physical, or conceptual boundaries of possibility. To those who pronounced them or accepted them as being true, such statements defined the limits of our "real" and logical world.

Each statement, however, has since been proved wrong. In 1979, Reinhold Messner and Peter Habeler climbed Everest without oxygen. Modern electronic synthesizers create music (or at least sound) with more information content than that of a symphony orchestra. Scientists have fabricated strained-layer superlattices in the laboratory. Thus, we have expanded what we define as the possible; we have made inroads into the dark region of impossibility, pushing back the once-solid walls of the unreal.

No Way: The Nature of the Impossible, edited by Philip J. Davis and David Park, is a collection of essays by prominent authors in widely varying fields. Each writer engages impossibility with language and methods specific to his or her own specialty. The result is a fascinating and thorough dialogue about a concept that transcends, limits, and even defines endeavors including mountain climbing, medicine, technology, parenting, and artificial intelligence.

The book begins with a chapter entitled "Everest and the Impossible" by Scott Lankford, a climber and a veteran of a failed 1985 attempt on the mountain. Lankford recalls the sense of melancholy that accompanied the triumph of Everest's first ascent, for it seemed the last great earthly obstacle, the last of nature's impossibilities. The instant Edmund Hillary and Tensing Norgay mounted the summit in 1953, Everest entered the realm of the possible. The mountain was not conquered, however, for at that same instant arose new impossibilities: Everest without oxygen, Everest alone, Everest "the hard way." As Lankford says, "Climbers don't conquer the impossible, they invent it."

Lankford's essay sets the tone of the book; the mountain is a particularly apt model for the impossible. The base is accessible, but the peak is not. The edge of the possible thus lies somewhere in between, graded on the gradual slope. Lankford describes his own encounter with this boundary when, with the realization that he and his party could not reach the summit, "The door to the invisible *became visible*, like the summit of Everest itself suddenly torn from the clouds." The impossible is not always a wall or a cliff, but more often is like the slope of a mountain. This feature, of progressive, almost exponential steepness, lets us approach and even drive back the impossible while never overcoming it.

In another essay, Michael J. Katz asserts that *evolution* of any organism would be possible were it not for the limits of the physical world. "Hedgehogs that run faster than the speed of light," for example, could evolve if they were physically possible.

Michael Yarmolinsky takes a similar approach but delves more into the limits of the evolutionary process. He impressively relates evolution to what he calls the "central dogma" of molecular biology, namely Francis Crick's generalization that information can flow only from DNA to RNA into proteins and *not* the other way (i.e., "protein cannot serve as

ILLUSTRATION: CHRIS SCHMIDT © 1989

a template for its own replication"). Thus, the reader arrives at Charles Darwin's famous break from earlier (Lamarckian) theories of evolution: "It is impossible for acquired characteristics to direct evolution."

Physicist Michael Sturge's contribution explores several cases of the impossible in technology. He discusses the semiconductor laser, an example of the most straightforward type of technological impossibility—the limits of current technology. Skeptics (including Sturge himself) pronounced the semiconductor laser to be impossible because of limits of crystal purity and opacity. The convergence of improved crystal growth with advances in solid-state physics, however, overcame these hurdles. The frontier of the possible extended to include the semiconductor laser, now a common element of optical disk drives and CD players.

Paradoxically, technological impossibility often arises from social and human factors rather than physical or natural properties. Sturge cites the above-mentioned strained-layer superlattice, which many experts declared impossible simply because they got their basic physics wrong; they neglected the minimum energy required to produce a damaging dislocation. Sturge also discusses the failure of Bell Labs to pursue research into large-scale ICs. Apparently the organizational and management structure could not admit certain kinds of change (the infamous Not Invented Here syndrome). In that sense, what is possible in today's technological world is what our socially determined institutions can develop and exploit: weapons, commercial products, medical instruments, and so on. A new technological development might not fit neatly into such categories or might not lend itself to the methods modern society has evolved for capitalizing on innovation. Such a technology might find itself confined to a

narrower possibility and to more imposing impossibilities than one that lies within traditional political, industrial, and financial flows.

The scientific and technical chapters of *No Way* have a common element: the exponential mountainside of impossibility. Some statements are clearly impossible and always will be: "You can't put a thousand gallons of water into a pint bottle." These are scientifically uninteresting, for they arise out of definitions, language, and logic, and will yield little when examined. In contrast, the impossibilities that define the limits of our knowledge are truly engaging: "There is no cure for AIDS." These problems are constantly revealing, shifting, and, we hope, receding; they occupy and define the borderlines of possibility. In David Park's words, ". . . that narrow strip, on either side of the border, is where science lies."

No Way is not all about science, however. It also covers law, politics, economics, education, poetry, music, and philosophy. Through the common theme, the authors introduce the reader to a broad range of discourses and the specific languages with which they operate. Furthermore, because the "impossible" is by its nature an abstract and even absent topic, the reader sees each discipline defined as a shadow, groping in the dark for the unreachable. He or she is left with a lasting impression of the academic disciplines, disparate in their language, yet converging in their attempts to know the unknowable.

BRIEFLY NOTED

Programmer's Guide to OS/2 *by Michael J. Young, Sybex, San Francisco, CA: 1988, 625 pages, $24.95.* If the complexity of a software product were measured by the number and size of the books written about it, there could be

BOOK REVIEWS

no doubt that OS/2 is a complicated beast. The first few books written about it seemed to be almost extensions—both in style and bulk—of the OS/2 specification and the documentation shipped with its early versions.

Michael J. Young's new book exhibits a maturity in its presentation that earlier books have lacked. He speaks with the authority that comes from having done some serious OS/2 programming. Instead of just stating the purpose of an OS/2 function call, Young actually describes how it might be used in real code. For example, instead of stating that DosQHandType determines if standard output has been redirected to a file, he describes a situation in which you might need to discover if a child process's output is redirected by its parent process to a disk file so the parent can continue to write to the screen without interference from the child. The description goes on to warn that using DosWrite is better than a similar video I/O call because DosWrite's output can be redirected, while VIO calls always write to the screen. These subtle insights can be useful.

The book also offers a fair amount of code, most of it in C, not assembly language. Some of the code is long and complex, like the cut-and-paste utility that spans 11½ pages. A companion disk contains all the listings, including header, make, and definition files.

Of the book's 625 pages, 250 pages are appendixes devoted to summarizing API calls and error messages. But the preceding 375 pages contain a lot of practical advice for the aspiring OS/2 programmer.—*G. Michael Vose*

Peopleware: Productive Projects and Teams by Tom DeMarco and Timothy Lister, Dorset House Publishing Co., New York: 1987, 188 pages, $23. Peopleware is the *Mythical Man-Month* for the 1990s. Just as Fred Brooks was emi-

nently qualified to write his book on software-project management, Tom DeMarco and Timothy Lister, internationally known software consultants, are similarly qualified to write about productive teams. DeMarco's *Structured Analysis and System Specification* is a classic in the field of information management.

Peopleware is a collection of essays about teams, productivity, and quality. The authors' style is casual, and they sprinkle their essays with anecdotal insights. But don't be misled by the informal style into expecting a hodgepodge of software folklore and rule-of-thumb management quips. The information presented is based on years of observing development efforts, surveys, and the authors' annual Coding War Games.

The book is divided into five parts. In the first section, "Managing the Human Resource," DeMarco and Lister advise against treating people as faceless modular resources. Instead, they recommend that managers seek to motivate their teams with the goal of producing good products in a sane environment. The underlying thesis of this section is that most of the problems occurring in high-tech projects are not technological but sociological.

The largest single section of the book is devoted to the office environment. This section is an indictment of the open-office concept that has swept through corporate America. It shows the fallacy in the cost-benefit analysis that says, "We can put a hundred people on this floor, and just look at all the money we save." An open office, the authors maintain, guarantees continuous interruption. No studies support the open office, while several (by IBM, ITT, and others) show that it has a deleterious effect on knowledge workers.

The authors give a simple formula for managing teams: Get the best people (cut out the *continued*

WANG MICROSYSTEMS LOOKS AT YOU DIFFERENTLY. OUR PCs PROVE IT.

World Radio History

deadwood), and make them happy. Turn them loose.

The last three parts of the book are devoted to understanding how to grow teams. First, the authors explore how *not* to do it—ways to avoid "teamicide." Some of the mistakes leading to teamicide are fragmentation of people's time (e.g., one person working on four projects), forced quality reductions, phony deadlines, and defensive management. The authors recommend that managers make a cult of quality and encourage teams with a sense of eliteness and teams that support individuality. The management should provide strategic but not tactical direction.

Peopleware is a sociology book that deserves a place on any technical manager's bookshelf. It's not shelfware, though—take it down and review it before your next team meeting.—*Charles Herring*

Better Scientific and Technical Writing by Morris I. Bolsky, Prentice Hall, Englewood Cliffs, NJ: 1988, 156 pages, $14.95. If you write reports, documentation, or even interoffice memos, *Better Scientific and Technical Writing* can help. In a book that practices what it preaches, Morris I. Bolsky (a technical writer for AT&T Bell Laboratories) presents guidelines, tips, and techniques for conveying information in the clearest, most effective manner.

This slim, well-organized volume touches on all aspects of a writing project, from planning to printing. The heart of the book, however, is the chapter entitled "Design—Principles," which offers an excellent set of guidelines for good writing. Explanations and examples support these principles, which stress simplicity and readability in everything from word choice and sentence structure to overall tone and organization. A list on the back cover summarizes this and other material, and it's almost worth the price of the book just to have that synopsis handy.

Bolsky gives considerable attention to the presentation of information, and he offers tips on such reader aids as footnotes, contents and index, and tables. The chapter on physical design covers the actual printing of a document, from deciding on a typeface to choosing the most readable ink color.

Brief sections on grammar and punctuation provide a quick reference for style questions, and a reasonably comprehensive index makes the book's information easily accessible.

You may not agree with Bolsky on every point—not everyone thinks a table of contents belongs on the cover, and using "they" as a singular pronoun makes a lot of people wince—but Bolsky himself is quick to point out that he offers recommendations, not rules. With that in mind, *Better Scientific and Technical Writing* is a useful handbook for novice writers and a good reference for old hands.

—*Margaret A. Richard*

C Traps and Pitfalls by Andrew Koenig, Addison-Wesley, Reading, MA: 1989, 147 pages, $16.25. The product of over 20 years of programming experience, including 10 years using the C language at Bell Labs, *C Traps and Pitfalls* presents the most pervasive of classic goofs and gaffes you're likely to encounter while programming in C.

The book was originally circulated as an internal paper within Bell Labs. Unprecedented enthusiasm there prompted Koenig to turn the paper into a book. All the problems mentioned in the book come from the author's own programming and the experience of other members of Bell Labs' technical staff.

Like any good book dealing with C, it starts with an introduction. Eight more chapters cover lexical, semantic, and syntactic pitfalls, as well as library routines, the preprocessor, portability concerns, and *continued*

World Radio History

BOOK REVIEWS

problems that occur during the link phase. The style is lean and pithy, with section headings like "Memory location zero" and "= is not = =." Pitfalls are amply illustrated with C code fragments. The examples also examine the anomalous program behavior caused by a given programming error. Where appropriate, the book lists differences that might occur while using proposed ANSI C versus one of the older commercial versions of C. Each chapter concludes with a few useful exercises, the answers to which are at the end of the book. An appendix gives the most complete discussion of the printf command that I've ever seen.

For the novice C programmer, this book will provide a relief from an almost certain future of chaos. The professional programmer will chuckle upon encountering certain familiar pitfalls but will also gain a better insight into these. The book is easy to read, yet it contains a lasting store of information for any C programmer.—*Jason Levitt*

What Do *You* Care What Other People Think? *by Richard Feynman, W. W. Norton and Co., New York: 1988, 255 pages, $17.95.* If there is one scientist who deserves to have his adventures cataloged in a hagiography, it is Richard Feynman: winner of the Nobel prize in physics, a distinguished member of the National Academy of Sciences, onetime actor, artist, and sometime drummer in a Brazilian samba band.

This book is a continuation of tales from Feynman's life, following the best-selling success of *Surely You Must Be Joking, Mr. Feynman.* The first section of the book is filled with personal stories and memories told in a grandfatherly style. The second section focuses on the space shuttle program and on the investigation of the Challenger explosion. It is a treat for those who love Feynman's chutzpah and a lesson for anyone in-

volved in a large engineering project. Feynman turned out to be the only member of the President's investigative committee not working for the Air Force or NASA. He realized this gave him the freedom to ask difficult questions, and he took advantage of that opportunity.

At the time, Feynman found that a bit of bureaucratic hardening of the arteries had set in at NASA. The different engineering problems were gradually coated over with sugar as each manager briefed his boss. As a result, upper management steadfastly adhered to a failure probability estimate of 1 in 100,000 launches, while the working engineers said it was 1 in 100.

Computer scientists will find particular interest in the chapters on the software that runs the space shuttle. That side of the NASA effort uses very rigorous standards that require adversarial teams to try and find problems in each other's code.

Feynman died last February, and this will likely be the last published collection of his anecdotes. The book makes excellent reading both for its entertaining style and for Feynman's keen insights.
—*Peter Wayner* ■

CONTRIBUTORS

David A. Mindell is a technical consultant who lives in Aspen, Colorado. **G. Michael Vose** is coeditor of *OS Report: News and Views on OS/2.* He lives in Peterborough, New Hampshire. **Charles Herring** is a computer scientist at the U.S. Army Construction Engineering Research Laboratory in Champaign, Illinois. **Margaret A. Richard** has a Certificate in Technical Communications from the University of Lowell. She is a copy editor for BYTE. **Jason Levitt** is a Unix aficionado and freelance writer based in Austin, Texas. **Peter Wayner** is a doctoral student in computer science at Cornell University.

WHAT'S NEW

The Crayon 386 20/20 SP is rack-mounted for rugged conditions.

Where Other Computers Would Just Slide Off

The Crayon 386 20/20 SP is a full-featured computer for the rugged industrial environment where only rack-mounted machines survive. Yet it's portable enough—17 by 20 by 8 inches and 32 pounds—to get you between jobs.

The zero-wait-state microprocessor and Award BIOS get you started with a megabyte of system RAM that's expandable to 16 megabytes (all of which is addressable through the 32-bit data path). Your graphics applications can be stored on 5¼- and 3½-inch disks, and there's a SCSI adapter for up to seven peripherals. You can add up to six half-height drives for additional storage capacity.

Two temperature-controlled fans are mounted on the back panel. There are also eight full-length expansion slots. MS-DOS 4.0 is included, as is a Honeywell 100-key keyboard.
Price: $7495.
Contact: Video Graphic Systems, 4163 St. Clair Ave., Studio City, CA 91604, (818) 509-5738.
Inquiry 1142.

Transputer Box for AppleTalk Networks

Supercomputer performance on a Macintosh network? That's what a small New York company claims with its parallel-processing box called Chorus.

The box is a tower-size system that can hold 4 to 16 parallel processors with 1 megabyte of memory for each

processor. In addition, Chorus has provisions for network connectors that allow it to be accessed by a number of Macintoshes.

In the present configuration, Chorus uses four T-800 transputers. The number of processors can be expanded to 16, 4 at a time. Its manufacturer, Human Devices, claims that each transputer is about five times as powerful as a 68020/68881 processor combination. Future versions may use other processors.

Several Macs can share Chorus as a computational server. Programmers can access the processors by using a parallel-processing operating system from Yale University called Linda (see the article "Getting the Job Done" by David Gelernter in the November 1988 BYTE). Linda statements can intermix with C

code in the Mac's MPW environment.

Human Devices claims programmers will be able to have Mac applications farm out computationally intensive tasks to Chorus. The company will also provide a Chorus-emulator software package that will allow programmers to write and test Chorus applications without actually having one on a network.

Chorus includes an Apple-Talk connector for connection to a Mac network and has provisions for Ethernet. The company says that in the future it may have a version of Chorus for the IBM PC.
Price: $25,000.
Contact: Human Devices, Inc., 322 West 71st St., New York, NY 10023, (212) 580-0257.
Inquiry 1140.

Modular CPU Machine in Small Tower

The VIP SX386 from Advanced Logic Research (ALR) is a small-footprint 16-MHz desktop computer with a "modular" plug-in CPU for later upgrade to a 20-MHz Intel 80386.

It's also packaged in a short-tower configuration, measuring 7 inches high, 4½ inches wide, and 15 inches deep and weighing only 17 pounds. But that miniature size doesn't stop it from packing in lots of features: one 8-bit and four 16-bit slots, a 40-megabyte hard disk drive, and a 1.44-megabyte 3½-inch floppy disk drive that's PS/2-compatible.

As with most other ALR machines, the company claims faster computing power than similarly configured Compaq computers. With a proprietary FlexCache architecture as a "cache management scheme," ALR claims 30 percent better performance over Compaq's 80386SX machines. Hardware on the VIP SX386 includes an 82385 cache controller and 512K bytes of RAM (expandable to 8 megabytes).

A 40-megabyte internal hard disk drive in a 3½-inch form factor is optional.
Price: $2395; with 40-megabyte internal hard disk drive, $3695.
Contact: Advanced Logic Research, Inc., 9401 Jeronimo, Irvine, CA 92718, (800) 444-4257; in California, (714) 581-6770.
Inquiry 1141.

continued

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.

Say Goodbye to Thermal Paper

The JetFax works with your LaserJet-type printer for laser-printer clarity.

It connects directly to your phone line through the standard copper cable, and you plug it into the printer through the printer's parallel port. It operates at the standard Group 3 rates of up to 9600 bits per second and can receive up to 30 pages with its one-half megabyte of internal memory.

If the printer is active when your JetFax begins to receive a facsimile, your JetFax will store the incoming message and will print it after the initial print job has concluded. With a full megabyte of RAM, you can store up to 60 fax pages, as well as use the higher-resolution mode of many laser printers. The memory also serves as a printer-sharing device for PCs, or a file server on a local-area network.

You can also add an economy fax machine to send faxes and to handle any overload.
Price: $1195; with 1 megabyte of RAM, $1395.
Contact: Hybrid Fax, Inc., 1733 Woodside Rd., Suite 335, Redwood City, CA 94061, (415) 369-0600.
Inquiry 1148.

Wallet-Size Hard Disk Drive

The Hardpac Micro 20 is an 8-ounce, 20-megabyte hard disk drive featuring an access time of 28 milliseconds and a shock rating of 100 g's.

Features include a hard disk drive access mechanism that's 30 percent the size of standard 3½-inch disk drives, ramped heads, automatic parking, and a 1.5-watt power

Print out those faxes on a LaserJet.

requirement.

You need a half-slot XT or AT host adapter for desktops, or any Amstrad laptop or Toshiba Model 1100.
Price: $1195; host adapters are all $95.
Contact: Aristotle Industries, Inc., 3226 Beta Ave., Burnaby, BC, Canada V5G 4K4, (604) 294-1113.
Inquiry 1151.

Printer Redefines Portability

The Toshiba Express-Writer 301 portable printer features a 24-pin print head in a 4-pound package.

Maximum speed is only 60 characters per second. Graphics resolution is 180 by 180 dots per inch, and the maximum paper width is 8½ inches. Normal speed is 42 cps, and it prints best on thermal transfer paper.

Memory is 2K bytes of RAM, and it includes a nickel-cadmium battery.
Price: $489.
Contact: Toshiba America, Inc., Information Systems Division, 9740 Irvine Blvd., Irvine, CA 92718, (800) 457-7777; in California, (714) 583-3000.
Inquiry 1149.

And, for Your Portable Mac

The WriteMove is a 192-dpi ink-jet printer that will fit inside your Macintosh's carrying case. Write-Move weighs 3 pounds and measures 2 by 7 by 11 inches.

QuickDraw imaging technology hardware is coupled with multiple software features, including six fonts, your choice of printer drivers, a spooler, an installer, and a print-manager application.

You can scale the fonts to any size or rotate, condense, expand, or manipulate them like PostScript-compatible fonts. You can also reduce documents to 25 percent and enlarge them to 400 percent in 1 percent increments, while retaining 192 dpi.

The driver allows for several printing options. For example, Preview lets you view the document on-screen before printing; Draft prints less than 192 dpi for quicker output; High-Quality ensures 192 dpi; and Print Later permits batching of print jobs for printing later with the print manager application.
Price: $699.
Contact: GCC Technologies, 580 Winter St., Waltham, MA 02154, (617) 890-0880.
Inquiry 1150.

Hard Disk Drives For LAN Servers

A series of hard disk drives intended for use in PC LAN server applications has been designed by Plus Development.

The Impulse series includes 40- and 80-megabyte drives and a proprietary 16-bit bus interface card. You can install up to four of the half-height 3½-inch drives inside a PC and attach up to 14 in external chassis that hold two drives each, for a total of up to 16 drives, or 1.3 gigabytes, connected to a single card.

You can install two cards in a PC backplane; under DOS and OS/2, you can access up to 24 drives. Under Novell's NetWare, 32 drives, or 2.6 gigabytes of storage, can exist on the server, the company says.

Plus offers the Impulse models as an alternative to 150- to 350-megabyte server drives. The advantage is that with Impulse, you can expand server capacity incrementally, in chunks of 40 or 80 megabytes. Plus says the system, with multiple disk units, can achieve a degree of parallel data access. Although all data eventually has to pass sequentially through the interface card and onto the bus, the drive heads can rotate and seek simultaneously.

Plus has enhanced this distributed architecture with a proprietary disk controller that the company says achieves data throughput at roughly four times the speed of a SCSI port. The cluster-disk interface transfers data 512 bytes at a time.
Price: $995 to $1379.
Contact: Plus Development Corp., 1778 McCarthy Blvd., Milpitas, CA 95035, (408) 434-6900.
Inquiry 1152.

Data Acquisition Reaches the Mac SE

The Lab-SE data acquisition board from National Instruments features A/D conversion with software-programmable gain, on-board timing, flexible channel scanning, and full interrupt capability.

It contains an 8-bit A/D converter with eight analog inputs and a 125-kHz sampling rate, a 13-bit integrating A/D converter (12 bits and a sign-bit), two 8-bit D/A converters with voltage outputs, 24 lines of TTL-compatible digital I/O, and three 16-bit on-board counter/timer channels.

Each Lab-SE comes equipped with a 5½-inch ribbon cable that connects the card from inside the SE to a 50-pin male ribbon connector that mounts on the back of the SE. The Lab-SE also needs the software driver with routines that are callable from any language that supports system Toolbox device manager calls. The software driver includes language interfaces for MPW C, Lightspeed C,

National Instruments' Lab-SE board for the Mac.

and Microsoft BASIC.

For the complete graphical programming environment for developing scientific and engineering applications, you need LabView, a graphical programming language, the Lab-SE, and the Lab-SE driver. With them, you can develop application programs to control the SE board.
Price: Data acquisition board, $595; software driver, $95; LabView, $1995 (with significant academic discounts available).
Contact: National Instru-

ments Corp., 12109 Technology Blvd., Austin, TX 78727, (800) 531-4742; in Texas, (800) 433-3488.
Inquiry 1154.

VGA Board Generates NTSC-Compatible Signals

With the Recordable VGA graphics board, IBM PC owners can afford to attain some of the video capabilities of the Commodore

Amiga and Atari ST.

The US Video Recordable VGA board for the IBM PC, XT, AT, and compatibles will output NTSC-standard composite video in real time to TV monitors and VCRs. To do this, it performs two functions: It acts as a VGA interface to high-resolution color monitors (and offers performance 400 percent to 600 percent faster than basic VGA, the company claims), and it converts computer graphics to video that can be recorded and played back on any standard videotape recorder.

Each board supports several high-resolution modes, including 640 by 480 pixels with 256 colors, 800 by 600 pixels with 16 colors, 1024 by 768 pixels with 16 colors, and 132-column text modes in 25, 30, 43, and 60 rows. Drivers for AutoCAD, PCAD, Ventura Publisher, Lotus 1-2-3, and WordPerfect are included.

While the Recordable VGA promises 100 percent NTSC compatibility, US Video readily admits that its system cannot remove flicker. But the company says that by using VGA's large palette of colors to shade lines, the flicker from Recordable VGA is "no worse than what you see on TV."

One enhancement module is available, with two others in the works, the company says. The overlay module, which includes genlock functions, combines a computer-generated image with an external video source in a video-recordable format.
Price: $785; overlay module, $385.
Contact: US Video, One Stamford Landing, 62 Southfield Ave., Stamford, CT 06902, (203) 964-9000.
Inquiry 1157.

continued

Adding MIPS with Chips

When you add the 260 Personal Mainframe or the 270PM to your IBM XT, AT, PS/2 Model 30, or Model 35, you also add as much as 10 million instructions per second (MIPS) of processing power. And you provide a Unix System V operating system to work alone or concurrently with your system's native DOS.

Both the 260PM and the 270PM are based on the 32532 processor from National Semiconductor. The 260PM operates at 25 MHz and offers 8.5-MIPS performance, claims manufacturer Opus Systems. Simi-

larly, the company claims that the 270PM operates at 30 MHz and offers 10-MIPS performance.

That means, Opus says, when a 270PM is added to a 25-MHz cached 80386-based system that operates at 5 MIPS, the resulting system totals 15 MIPS and is a dual-processing system. The 260PM and 270PM dominate much of the I/O processing, but not all.

Both the products offer a version of the MIT X Window System and the 32381 floating-point processor to support single- and double-precision IEEE-format

arithmetic calculations. Optionally available is software that supports an Ethernet controller running TCP/IP; similar TCP/IP software bundled with Sun's Network File System; and application software written in C, FORTRAN 77, COBOL, Common Lisp, BASIC, and others.
Price: 260PM with 4 megabytes of RAM, $6995; 270PM with 4 megabytes of RAM, $7995.
Contact: Opus Systems, 20863 Stevens Creek Blvd., Building 400, Cupertino, CA 95014, (408) 446-2110.
Inquiry 1153.

Lantana's Cypress/2 16-bit token-ring add-in board.

Bus Mastering Debut on Token Ring

The first available add-in board that is configured as a bus master controller comes from Mexican technology and American marketing, according to Lantana Technology. Lantana is the marketing-driven company that recently opened shop in the U.S. to import and sell technology that was originally developed in Mexico by Computadoras Micron S.A. de C.V.

The Cypress/2 board is a 16-bit token-ring add-in board for IBM PC–compatible 4-megabit-per-second token-ring networks. The bus-mastering feature, the most advanced feature of the PS/2 Micro Channel bus, allows the main system processor to give up control of the system bus temporarily so that a processor on an expansion card can use it for high-speed data transfers.

In addition to the bus-mastering feature, the Lantana boards include 128K bytes of RAM that expands the buffer and allows the IEEE 802.2 logical link control software to be loaded, a 32K-byte EPROM module that stores the LLC software in ROM, and a remote-program-load EPROM module for diskless workstations.
Price: $795.
Contact: Lantana Technology, Inc., 4393 Viewridge Ave., Suite A, San Diego, CA 92123, (619) 565-6400.
Inquiry 1164.

IBM Cranks Up Token Ring to 16 Megabits

The token-passing technology that IBM popularized in its 4-megabit-per-second local-area network is now available in IEEE 802.3–compliant hardware supporting 16-Mbps data rates. This pits IBM in the networking arena against the companies offering products that conform to the 10-Mbps IEEE-specified Ethernet networks.

At the heart of the new add-in boards is an IBM-designed CMOS A/D interface chip. The Micro Channel board can be switched between 4 and 16 Mbps to upgrade existing token rings. The 16-Mbps AT board cannot be switched. Neither board includes bus-mastering capabilities.

All the new adapter cards provide 64K bytes of RAM instead of the 8K-byte or 16K-byte RAM on previous boards. When used in the file servers of 4-Mbps LANs, IBM says, the boards increase the number of workstations that can share server resources.
Price: $895.
Contact: Consult your local IBM branch office or call (800) 426-2468.
Inquiry 1165.

Modem Links Ethernets Using Broadband

The 10Broad36 LanExpress modem from Lanex lets you run two or eight Ethernet local-area networks (LANs) on the same coaxial cable.

The IEEE 802.3 broadband standard means multiple LANs on one cable (with different frequencies). Baseband Ethernet is more widespread, yet it allows only one LAN (one frequency) per coaxial cable.

The modem performs the same function over a broadband LAN that Ethernet transceivers do for a baseband LAN, without affecting the baseband-controller hardware or software. In other words, the 10Broad36 replaces baseband Ethernet transceivers, the most popular of the Ethernet transceivers.

Both the two-port and the eight-port LanExpress models connect Transmission Control Protocol/Internet Protocol nodes over a 10-megabit-per-second Ethernet broadband backbone. They also connect such multi-protocol devices as DECnet, NetWare, XNS, TCP/IP, and Open Systems Interconnection over the same coaxial cable.
Price: Two-port, $2695; eight-port, $2895.
Contact: Lanex Corp., 10727 Tucker St., Beltsville, MD 20705, (800) 638-5969; in Maryland, (301) 595-4700.
Inquiry 1162.

continued

Hayes Integrates PAD on the Modem

By incorporating a packet assembly/disassembly (PAD) function on a ROM chip in its X.25 modem, Hayes Microcomputer Products will provide PC users with multisession, multipoint communications through X.25 telecommunications. (PAD units enable equipment not designed for packet switching to access a packet-switched network.)

You can also use X.25 to establish a direct data connection at 2400 bits per second to any of the Integrated Services Digital Networks. With this modem, multisession, multipoint communications will mean you can make one standard telephone call to a telephone network that offers CCITT's X.25 and turn that connection into four data connections to PCs and mainframes alike. Then you hotkey from one data connection to another.

Value-added networks like Telenet, Tymnet, MCI, and Datapac (in Canada) already offer X.25 packet-switching services. Packet-switched information is less expensive than circuit-switched information. This is because it travels in packets and therefore isn't continuous, as it must be on traditional telephone networks.

Hayes says these modems will probably be used for managing terminals and workstations in dispersed locations. They could be used for multisession, multipoint communications.
Price: $895, or $50 for a single-chip ROM upgrade package for any of the company's V-series products.
Contact: Hayes Microcomputer Products, Inc., P.O. Box 105203, Atlanta, GA 30348, (404) 449-8791.
Inquiry 1163.

Sony's workstations feature high-capacity optical drives.

Sony's No-Baloney Workstation

If you've got room on your desk for an IBM PC AT, you've got more than enough room for Sony's high-powered 68030 Unix workstations with erasable optical disk drives. The NWS-1750 is based on a single 68030. The NWS-1830 and -1850 are based on two 68030s.

On one double-sided 5¼-inch cartridge, you can store 594 megabytes of formatted data. Rewriting is accomplished with a semiconductor laser and a biasing magnet to change the magnetic orientation of the cartridge's recording layer. The laser writes on the cartridge by magnetizing the cartridge's recording layer, which causes changes in the laser beam's plane of polarization.

The NWS-1750 comes standard with 4 megabytes of RAM (expandable to 32 megabytes). It also has a 286-megabyte hard disk drive. Graphics options include 14-inch, 816- by 1024-pixel

monochrome and 16- and 19-inch, 1280- by 1024-pixel four- or eight-plane color displays with a graphics controller.

The NWS-1830 and -1850 are rated at 5.3 million instructions per second based on dual 25-MHz 68030 microprocessors, a 25-MHz 68882 floating-point coprocessor, and a 64-byte cache memory. The second microprocessor on the 1800 series handles only I/O functions and can increase total system performance by 30 percent in many I/O-intensive applications, Sony says, making it particularly valuable when it's used as a file server or when tape operations are used.

All the 1800 series workstations are equipped with a 1.44-megabyte 3½-inch floppy disk drive, a 125-megabyte ¼-inch cartridge tape unit, 16 megabytes of RAM (expandable to 32 megabytes), and 286 megabytes of hard disk drive storage space.

All the 1700s and 1800s use the News-OS version 3 operating system, which is

based on Unix 4.3BSD. They also use the X Window System version 11 graphics standard from MIT.
Price: Sony's 1700 series, $13,900 to $51,600; the 1800 series, $31,900 to $54,200; NWP-539 erasable optical drive, $4650; optical cartridge, $250.
Contact: Sony Microsystems Co., 1049 Elwell Court, Palo Alto, CA 94303, (415) 965-4492.
Inquiry 1146.

This Time is Real

The Ohio Scientific 720 is based on the Motorola 68020 processors and a Unix-compatible operating system that is designed for real-time operating requirements.

Such real-time features as zero-wait-state system calls and Request and Event queues guarantee specific response times to external interrupts.

Through an additional level of priorities, the 720 guarantees each user some response time because the level is nonshared.

Each 720 has 12 RS-232C ports, so, through intelligent terminal concentrators and networking, the system can accommodate as many as 60 users. Additional 68020 CPUs (with supporting static RAM cache) can be added in a parallel arrangement with dynamic load balancing, the manufacturer claims.

The standard 720 includes 4 megabytes of RAM that's expandable to 64 megabytes. Hard disk capacities range from 91 megabytes to 1.2 gigabytes, totaling to 16 gigabytes.
Price: $6350.
Contact: Consolidated Computer Systems, Inc., 2150-D West Sixth Ave., Broomfield, CO 80020, (303) 460-0444.
Inquiry 1147.

A Workstation for the Laptop

The high-performance Toshiba T5100 multi-user portable workstation includes Intel's 80386 microprocessor and runs T/PIX, Toshiba's version of AT&T's Unix System V/386 with 4.2BSD extensions.

The workstation is targeted at business applications but is finding a niche with many software developers because, in its full-featured configuration, it comes bundled with T/PIX. That means there's no need to load and configure from the 25 disks Toshiba will send you to load T/PIX separately, and you can ship demo software in a compact package without much difficulty.

In its base configuration, the 16-MHz 80386 includes 2 megabytes of RAM, a 40-megabyte 29-millisecond hard disk drive, a 1.44-megabyte 3½-inch floppy disk drive, and a 640- by 400-pixel EGA-compatible gas-plasma display.

The keyboard has 82 keys, including an integrated numeric keypad; there's also a port for a 101-key keyboard so you can add your own.

Other standards are an EGA monitor controller, an RS-232C port, a combination parallel port and 5¼-inch floppy disk drive controller, and an internal expansion slot. You also get MS-DOS. The basic T5100 supports X Windows, networking through Remote File Sharing, and TCP/IP through an optional Ethernet interface.
Price: $7199; with 2 megabytes of RAM and T/PIX, $8750.
Contact: Toshiba America, Inc., Advanced Systems, 9740 Irvine Blvd., Irvine, CA 92718, (714) 583-3071.
Inquiry 1145.

Now Peter Norton

If you've ever seen the words "GENERAL FAILURE ERROR" on the screen of your PC, you know you don't need to be ill to feel sick.

All you need is a corrupted disk with some valuable data stored aboard.

When that happens, it's time to call on the remarkable new *Norton Disk Doctor.™*

The Norton Disk Doctor is one of 27 data recovery and disk management tools in the latest Advanced Edition of the Norton Utilities.®

It's also the most significant advance in utility software technology since Peter Norton first began saving the data and the derrières of grateful PC

```
C>        FILE NOT FOUND!
```

Our legendary UnErase—and its short-cutting offspring Quick UnErase have saved the derrières of thousands of grateful PC users. And they're still saving them.

users around the world with the legendary—and still unequalled—*UnErase.*®

Why you need a Doctor in the house.

Now, whenever INVALID DRIVE SPECIFICATION or some other cryptic or catastrophic error message appears on your screen, you can do something besides reach for the Maalox.®

You can summon the Disk Doctor, which will

QuickBASIC Gives Quick Help

Microsoft's latest incarnation of its venerable QuickBASIC programming language—version 4.5—features a new hypertext-based on-line help system. Dubbed "Advisor," it lets you instantly call up cross-referenced information from the language's entire reference manual. If you're a newcomer to programming, Advisor will help you learn BASIC more quickly; if you're an experienced hacker, it'll let you quickly find more esoteric information. The key to Advisor is a help engine. Taking up just 5K bytes of RAM, it includes a file manager, a text decompressor, and utilities.

QuickBASIC's Advisor lets you put the cursor on any word on the screen and get detailed information. For instance, if you're entering a line of BASIC code and want to know more about the IF statement, you place the cursor on IF, press a help key, and get the full reference to the IF statement. Microsoft says it will be integrating Advisor technology into upcoming releases

QuickBASIC's Advisor is a hypertext help system.

of its other programming languages.

Besides the new help technology, QuickBASIC 4.5 has the same features as version 4.0. They include a smart syntax-checking editor and compilation speeds of up to 150,000 lines per minute. There's also an improved source-level debugger with what Microsoft calls an "instant" watch capability.

QuickBASIC 4.5 runs on the IBM PC, XT, AT, PS/2s, and compatibles with 384K bytes of RAM, MS-DOS 2.1 or higher, and a graphics adapter. It supports the Microsoft Mouse and comes with

printed reference and tutorial manuals.
Price: $99.
Contact: Microsoft Corp., 16011 Northeast 36th Way, P.O. Box 97017, Redmond, WA 98073, (206) 882-8080.
Inquiry 1108.

Tools for the Turbo Pascal Trade

Written in Pascal, the AIS Programmer Tool-Kits let you manipulate windows, enter and edit data, design screens, and more in the Turbo Pascal Integrated Envi-

ronment. The ToolKits are compatible with Turbo Pascal versions 3.0 through 5.0 and contain several modules, including the Editor ToolKit, Screen Design Utilities, a Window ToolKit, a Dynamic Array ToolKit (for version 4.0), a Linked List ToolKit, and a Printer ToolKit.

Complete Pascal source code is included with the AIS Programmer ToolKits. To run the development package, you'll need an IBM PC, XT, AT, PS/2, or compatible running any Turbo Pascal version 3.0 through 5.0.
Price: $44.95.
Contact: Abaire Information Services, Inc., 2302 Ginter St., Richmond, VA 23228, (804) 262-2966.
Inquiry 1109.

Presentation Manager Development Tool

If you're an application developer using Presentation Manager, Object/1 may be the tool for you. Using object-oriented programming techniques, it offers a forms painter, a database interface, and CASE development tools.

The forms painter lets you create list boxes, radio buttons, and other graphical elements. You can create user interfaces such as pop-up menus, links from list boxes to database fields, and enforced integrity on data input.

The development tool includes source code along with hundreds of example classes, methods, and objects. It runs on 80286- or 80386-based systems with at least 4 megabytes of RAM.
Price: Approximately $900.
Contact: Micro Data Base Systems, Inc., P.O. Box 248, Lafayette, IN 47902, (800) 344-5832; in Indiana, (317) 463-2581.
Inquiry 1110.

A Little Jewel Makes Logical Connections

One of the trickiest problems for any programmer is correctly coding multibranch logic modules. Writing the code for constructs like nested IF... THEN statements can be both challenging and time-consuming. But Sterling Castle has a new program called Logic Gem that takes the headache out of the process, letting you concentrate on the problem instead of writing code for the solution.

You can think of Logic Gem as doing for logic what a spreadsheet does for num-

bers. Structured very much like a spreadsheet, it's an electronic decision table that the company claims eliminates the possibility of bugs.

Logic Gem is actually a collection of three decision-table tools: an editor, an interpreter, and a code generator. The editor automatically completes an incomplete logic table for you, generating a set of decision rules. It also eliminates redundant or contradictory rules. Then the interpreter steps you through the table to verify the logic. Finally, the code

generator translates the complete decision tables into program code.

Logic Gem produces code in C, structured BASIC, interpretive BASIC, Pascal, FORTRAN, dBASE, and English.

The program runs on any IBM PC or compatible and needs 640K bytes of RAM.
Price: $198.
Contact: Sterling Castle, 702 Washington St., Suite 174, Marina del Rey, CA 90292, (800) 722-7853; in California, (800) 323-6406.
Inquiry 1107.

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.

spectrum

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.

Low-Cost CAD

Foresight Resources has spruced up its Drafix CAD and renamed it Drafix CAD Ultra.

One of the enhancements is a CADapult utility, which lets you take attribute information and convert it to a variety of databases, spreadsheets, and high-level language formats. Then you can bring the information back into Drafix CAD Ultra.

The program also lets you use its Hewlett-Packard Graphics Language export command to save plotting instructions in formats compatible with desktop publishing programs such as Aldus Page-Maker and Ventura Publisher. You can also import ASCII text files from word processing programs and use the text in your CAD drawings.

Also included with the program is a library of over 450 predrawn symbols. Separate libraries of architectural, mechanical engineering, and electrical engineering symbols are available. Other enhancements include off-line plotting and new plotter drivers that let you use D- and E-size plotters. The user interface is enhanced with the addition of status displays of layer, linetype, and pen identification. You can also backtrack through up to 100 drawing commands. The new Drafix CAD supports a wider variety of graphics adapters.

Drafix CAD Ultra runs on the IBM PC with 640K bytes of RAM, DOS 2.0 or higher, a graphics card, and a mouse.
Price: $395; Report writer, $150; Professional Symbol Libraries, $200 each.
Contact: Foresight Resources Corp., 10725 Ambassador Dr., Kansas City, MO 64153, (816) 891-1040.
Inquiry 1113.

Drafix CAD now has import/export capabilities.

Souped-Up Statistical Analysis

SPSS has dropped copy protection and added new features to the statistical analysis program SPSS/PC+. Version 3.0 now lets you use 500 variables instead of the 200 in previous versions. Also added is an exploratory data-analysis procedure that produces univariate statistics and a variety of plots.

Network support is also enhanced. Previous versions supported Novell LANs, and the current version supports token-ring and 3Com networks as well.

SPSS/PC+ 3.0 runs on the IBM PC, XT, AT, and compatibles with 512K bytes of RAM and DOS 2.0 or higher.
Price: $795.
Contact: SPSS, Inc., 444 North Michigan Ave., Chicago, IL 60611, (312) 329-3300.
Inquiry 1114.

Go with the (Hydraulic) Flow

Hydronet is a fluid-analysis program that is used in applications such as water supply and treatment, petroleum transport, industrial and process engineering, wastewater conveyance, and fire-protection systems. The program computes the steady-state flows and pressures throughout systems.

The program is upgraded with a Lotus 1-2-3–type interface, display and printer graphics that show hydraulic grade lines, fitting and valve usage, critical pressure and pump calculations, and a larger system-design capacity.

The program runs on the IBM PC, XT, AT, or compatibles with 512K bytes of RAM.
Price: $495.
Contact: Engineering Software, P.O. Box 1450, Nevada City, CA 95959, (916) 288-3470.
Inquiry 1117.

Tango Totally Revamped

The Tango family of electronic-design software has been completely rewritten, according to Accel Technologies. Tango-PCB Series II enhancements include user-definable tracks, pads, vias, and text strings. There are more layers, improved grid options, and a wider range of video and printer/plotter support.

The previous version allowed only four fixed track sizes, 22 pad types, and eight text-string sizes. Now you can have any number of track widths, seven pad shapes, and text strings of from 4 to 1020 millimeters. Also increased is the number of layers—up to 19 from the previous 9. The maximum workspace is increased from 32 by 19 inches to 32 by 32 inches.

Series II now supports VGA, Hercules, and other graphics cards, in addition to CGA and EGA, which were supported in the previous version.

All functions for editing, plotting, and photoplotting are bundled into one program. New features include user-definable arcs on any layer. Component pattern libraries are also larger.

The new user interface, called the Accel Productivity Interface (API), features pop-up menus and dialog boxes. It also has a Speed Palette, which you can use for the commands you need most.

An auto-router is also added to the Tango-PCB Series II. The four-layer, 25-mil-grid, multipass auto-router offers pop-up menus and dialog boxes for setting up design parameters.

To run the Tango-PCB Series II products, you need an IBM PC or compatible with a CGA, VGA, EGA, or Hercules graphics card.
Price: PCB alone, $595; PCB with Route, $995.
Contact: Accel Technologies, Inc., 7358 Trade St., San Diego, CA 92121, (619) 695-2000.
Inquiry 1115.

WHY YOU SHOULD BUY COMPUTERS BY THE BOOK.

As you can see from our 4-page ad at the front of this magazine, buying direct from Dell is the best possible way to buy computers. And this book is the best way to order. The Dell Computer Catalog. In 28 pages, it gives you full information on our high-performance 386™ systems, 286 systems, laser printers, and much more.

Besides our own equipment, you'll also find a full line of software, peripherals, accessories, enhancement products, and networking solutions. All at lower-than-retail prices with higher-than-retail service and support. Just call us at (800) 426-5150. In Canada, call (800) 387-5752. Or send the attached reply card. And we'll send a catalog immediately.

DELL COMPUTER CORPORATION

Quotron Opens Windows

A new set of software that operates in the Microsoft Windows environment lets Quotron PC users integrate their Quotron financial information services with other applications.

QuotData, QuotChart, and QuotTerm make up the Open Windows family. QuotData lets you create databases of Quotron real-time market data that you can use with other software. QuotChart is a charting and technical-analysis program that is fed by the QuotData application. QuotTerm is a terminal-emulation application that gives you access to Quotron financial information services while viewing other applications on the PC.

The Open Windows products are available individually or bundled as one package.

To run QuotChart and QuotData you need an 80286 or 80386 DOS-based PC with at least 2 megabytes of RAM, a hard disk drive, one serial communications port, one bus mouse, an EGA card, DOS 3.0 or higher, and Windows/ 286 or 386 version 2.0 or higher. To run QuotTerm, you need at least 640K bytes of RAM, one serial communications port, one bus or serial mouse, DOS 3.0 or higher, and Windows/286 or 386 version 2.0 or higher.
Price: Monthly charges per workstation: QuotChart, $300; QuotData, $250; QuotTerm, $100. One-time charges: Open Windows installation, $500; Excel and Windows/386, $350.
Contact: Quotron Systems, Inc., 12731 West Jefferson Blvd., Los Angeles, CA 90066, (213) 827-4600.
Inquiry 1126.

Quotron's windows of financial data.

Present Yourself

From Genesis Data Systems comes the Rapid Prototyping System (RPS): software that helps you design, prototype, and present your ideas.

The system contains three modules: a screen designer, a music designer, and a prototyper. Also included are a memory-resident screen-capture utility, a music utility, and a run-time utility.

The screen-design module offers 16 foreground and background colors, or 16 foreground and eight background colors with blinking counterparts. You can use the memory-resident module to capture screens from other programs. You can also add overlays and animation.

The module's editing capabilities let you edit, move, or duplicate selected areas of the screen. You can use macros and Undo and Redo Undo commands. You also have the ability to customize lettering.

The music module uses standard music notation for composing or duplicating music. The range includes eighth notes to whole notes, and you can use dotted or tied rests. You have a selection of four tempos, and you can build key signatures, repeats, and refrains. You can save tunes to a disk file to play in the background of presentations, either from the DOS prompt or through instructions you create with the prototyping module.

You use the prototyping module to pull the screens and music together into presentations. It features variable commands, branching instructions, and real-time pauses.

RPS runs on the IBM PC, XT, AT, PS/2s, and compatibles with 256K bytes of RAM and DOS 2.0 or higher. It supports color, monochrome, or TTL monitors.
Price: $249.95.
Contact: Genesis Data Systems, 8415 Washington Place NE, Suite A, Albuquerque, NM 87113, (800) 777-1437; in New Mexico, (505) 821-9425.
Inquiry 1122.

Simple, Fast, and Cheap Word Processing

ZEdit is what its developer calls a front-end word processor. What does it mean? Mainly, that ZEdit is a fast, small, and surprisingly versatile package that admittedly isn't designed to be everything to everyone. Telecom Library admits that ZEdit is "lacking in print sophistication," but if you want to spruce up ZEdit's files with fancy formatting, you can load them into any word processor, because they're plain-vanilla ASCII.

ZEdit takes up just 52K bytes of RAM, and Telecom Library says that makes it ideal for floppy disk drive-only laptop computers or for

applications where you want to load a major application on top of ZEdit. And because ZEdit and the documents that you're working on stay in RAM, response time is fast. You can even preload documents into remaining RAM—as many as you have room for.

The package also lets you have up to eight windows open on the screen at once, with each window having either a different file or a different view of the same file. This gives you the unique ability to simultaneously edit the beginning, middle, and end of the same document. You can also move

data between any open windows, not just between adjacent ones.

You can configure ZEdit to emulate the key commands of any word processor you're already familiar with. It also features pull-down menus, mouse support, and a built-in keyboard macro facility.

To use ZEdit, you'll need an IBM PC, XT, AT, PS/2, or compatible.
Price: $59.95.
Contact: Telecom Library, Inc., 12 West 21st St., New York, NY 10010, (800) 999-0345; in New York, (212) 691-8215.
Inquiry 1121.

RIX Presents VGA Graphics

The VGA paint program ColoRIX offers 256-color VGA support for the IBM PS/2 Models 50, 60, 70, and 80 in 360- by 480-pixel resolution without any hardware additions or modifications. You can manipulate those 256 colors from a palette of 256,000.

Some of the key functions included in ColoRIX are a graduation command that lets you create color graduations, a Frepal command that lets you free any of the 256 displayed colors from the color palette, a Drag command that lets you leave a multiple-image trail as you move your selected image anywhere on-screen, and Smooth and Smudge commands.
Price: $199.
Contact: RIX SoftWorks, Inc., 18552 MacArthur Blvd., Suite 375, Irvine, CA 92715, (800) 345-9059; in California, (800) 233-5983.
Inquiry 1135.

Processing Images with Zip

Process images, control video digitizers, and display video images with Zip, a program that controls the ImageWise video digitizer transmitter and receiver.

The program offers more than 20 image-processing functions, including convolution, histogram equalization, and linearization, that allow you to manipulate and improve gray-scale images. You can also combine images.

Using the ImageWise digitizer, you can capture images from a video camera, VCR, or other video source in 256- by 244-pixel resolution with 64 levels of brightness.

Zip has 15 screen-display modes with three levels of zoom, color, and gray-level displays, minimum error techniques, dithers and duotones, and a VGA mode displaying 64 gray levels. You can save screen displays in PC Paintbrush and MacPaint file formats for use in desktop publishing or for further editing. And you can print outputs on dot-matrix, ink-jet, color, and laser printers. The program also produces halftones on laser printers.

The program runs on the IBM PC, XT, AT, or compatibles with 384K bytes of RAM, DOS 2.0 or higher, and an EGA or VGA. To capture video images, you need an ImageWise digitizer/transmitter and a video camera or VCR. To display video images in 64 levels on a television monitor, you need an ImageWise receiver/display, a video monitor, and an NTSC (National Television System Committee) composite.

(The ImageWise was a Circuit Cellar project by Steve Ciarcia in BYTE, May through August 1987. It is now available through Hogware.)
Price: Zip, $79; ImageWise video digitizer, $398.
Contact: Hogware Co., 470 Belleview, St. Louis, MO 63119, (314) 962-7833.
Inquiry 1134.

World Radio History

World Radio History

Put Your Mac in the Driver's Seat

MacChuck 1.5 lets you control an IBM PC or compatible from a window on the Macintosh. You can also use the program as a file transfer utility between the two formats. A network communication feature, added to version 1.5, enables you to operate a PC equipped with an Apple LocalTalk card, to be run by any Macintosh in a network.

Other enhancements in version 1.5 include improved file transfer capabilities and enhanced support for the Mac II.

The program comes with the software and cables needed to connect a Mac Plus, SE, or II to a 9- or 25-pin IBM PC serial connector. An adapter is also available for older Mac systems.
Price: $99.95.
Contact: Vano Associates, Inc., P.O. Box 12730, New Brighton, MN 55112, (612) 788-9547.
Inquiry 1128.

Two with Remote Access

If you need to access your computer from remote locations, there is a variety of remote-control programs on the market. Two previously released programs have had remote capabilities added to them.

Triton Technologies' Session/XL is a remote-control package for use with CO/Session communications software. Session/XL goes a step further by including a scripting and tasking communications language. Using the language, you can set up automatic custom communications sessions. For instance, a host PC can use Session/XL to automati-

cally poll remote PCs for daily sales activity, transfer the data to the host, and then download new information.

You can also use the package to automatically patch software, as well as upgrade and install new versions of applications. Session/XL uses 16-bit CRC (cyclic redundancy check) error checking and automatic compression for file transfers. You can also use the package to manually take control of remote systems.

The program also has error- and screen-checking capabilities. When Session/XL encounters an error, it's automatically recorded and the program continues with the next activity. Screen checking also lets you check a specific line or range of lines for a character string.

Session/XL runs on your host sytem. Each of the remote PCs must have CO/Session. Both programs run on the IBM PC, XT, AT, PS/2s, and compatibles with DOS 2.0 or higher.
Price: Session/XL, $225; CO/Session application license, $125; CO/Session support license, $175.
Contact: Triton Technologies, Inc., 200 Middlesex Essex Turnpike, Iselin, NJ 08830, (201) 855-9440.
Inquiry 1130.

The asynchronous communications program Bit-Com has been upgraded with remote-access features, a scroll buffer review, and automatic data compression with error correction.

The remote-access feature lets you run a program on the host from a remote PC. The scroll buffer review lets you review up to 500 lines of text that have already scrolled by on the screen. And the data-compression feature speeds up file transfers by 2 to 32 times, according to Bit Software.

continued

World Radio History

SOFTWARE • CONNECTIVITY

To run BitCom Deluxe, you need an IBM PC or compatible with 640K bytes of RAM and DOS 3.0 or higher.
Price: $79.
Contact: Bit Software, Inc., 830 Hillview Court, Suite 160, Milpitas, CA 95035, (408) 263-2197.
Inquiry 1131.

A Software Meter Reader for Your LAN

Connect Computer calls Turnstyle the electronic equivalent of a software library, enabling multiple users to share a controlled number of copies of a software program on a local-area network.

To use Turnstyle, you load it on your network and enter data about each copy of the software that resides on the network; Turnstyle controls the number of copies in circulation.

The program includes a monitoring system that tells you which users are on the system and what software they are using. You view or print reports of listings about the programs and their use. You can also use Turnstyle to keep track of the serial numbers of your software, although it's not necessary to enter the numbers into the database.

Turnstyle requires an IBM PC or compatible running Novell, Banyan, 3Com, or IBM network operating systems. You also need at least 300K bytes of RAM on the network file server. The workstations need DOS 3.0 or higher, a network interface card, and at least 7K bytes of RAM.
Price: $195.
Contact: Connect Computer Co., 9855 West 78th St., Suite 270, Eden Prairie, MN 55344, (612) 944-0181.
Inquiry 1132.

LAN Utility for Windows Users

In an attempt to simplify the use of Microsoft Windows applications running in a network, Automated Design created Windows Workstation, a set of utilities that supports Novell's NetWare.

The program consists of the Windows Menu (which includes Secure Station, Screen Saver, and Windows Intercom) and the printing utility, Windows Print.

Windows Menu is an application organizer designed specifically for networks, according to Automated Design. It uses between 10K bytes and 12K bytes of system memory and allows the LAN system supervisor to build custom menus with password protection and help messages.

Secure Station lets you lock your workstation while Windows applications function in the background. The utility monitors security violations and reports them to the user and the LAN system supervisor.

Screen Saver kicks in automatically if you leave your workstation for a specified period of time.

Windows Intercom sends network messages to and from Windows users. You can also reply to messages without leaving your current application. It also supports messaging between users of different NetWare file servers.

Future versions of this utility package will support other major LAN operating systems, the company reports.
Price: Windows Menu, $695; Windows Print, $695; bundled together for $1195.
Contact: Automated Design, 133 Johnson Ferry Rd., Suite 112, Marietta, GA 30068, (404) 988-0969.
Inquiry 1129.

Circle 82 on Reader Service Card

SOFTWARE • OTHER

Fastback Plus with New Features

Fastback Plus 2.0 offers improved restore and backup functions and an improved user interface.

The restore function now has a feature that estimates the number of files, amount of volume, and time needed to complete the restore, as well as the names for the backup sets that should be used. During restore, version 2.0 warns you if there isn't enough space on the target drive.

Another new feature, Smart Restore, prompts you for the exact backup disks needed and warns you if you insert the wrong one.

Other improvements include the option of appending incremental backups to an existing backup set. The program now warns you if you're using disks with data on them. A verify function lets you compare backup files with those on the hard disk.

Fastback Plus 2.0 runs on the IBM PC, XT, AT, PS/2s, and compatibles with at least 320K bytes of RAM and DOS 2.1 or higher.
Price: $189.
Contact: Fifth Generation Systems, Inc., 11200 Industriplex Blvd., Baton Rouge, LA 70809, (504) 291-7221.
Inquiry 1136.

Forget Remembering Your Backups

One of the many reasons why so few people make backups of their hard disk data is simply the memory factor: It's easy to forget to do it. But SitBack, from SitBack Technologies, will remember for you.

SitBack is a memory-resident backup utility that requires less than 15K bytes of memory. Once you load it and tell it which files to back up, it constantly monitors your hard disk for new or changed files. Then, when you don't use the keyboard for a specified length of time, it goes to work automatically, backing up the files to any DOS device you define, such as a floppy disk drive, a removable hard disk drive, a cartridge, a second hard disk drive, or a network device.

The program works with any MS-DOS-based system and is compatible with IBM NetBIOS networks.
Price: $99.
Contact: SitBack Technologies, Inc., 7219 West 95th St., Suite 301, Overland Park, KS 66212, (913) 894-0808.
Inquiry 1137.

Copyright It

To simplify the process of obtaining copyrights for your software creations, Synthetic Intelligence has developed Copyright-It. The program opens with a license agreement, and Form TX from the U.S. Copyright Office is available in the main menu.

Copyright-It runs you through the advantages and disadvantages of copyrighting, takes you step by step through the registration process, and offers other forms that you may need. When you're filling in the fields of the TX form, you have three levels of help available in each field.

The program runs on the IBM PC, XT, AT, and compatibles with 256K bytes of RAM and DOS 2.11 or higher.
Price: $25.
Contact: Synthetic Intelligence, Inc., 286 Fifth Ave., Suite 707, New York, NY 10001, (212) 967-2399.
Inquiry 1138.

Win the Ultimate *Heartbeat*

It's easy. Just enter the CADAM Corvette® Convertible Sweepstakes by June 30, 1989, and picture yourself in the *Ultimate Heartbeat,* from Chevrolet.® See your CADAM dealer today for rules and entry forms. No purchase necessary.*

It's just as easy to get your CAD/CAM/CAE program moving quickly with CADAM®'s software family. Software General Motors relies on to design world-class cars like their new Corvette. Software built around the famous CADAM user interface, which shares a common data format. Software relied on daily by more than 100,000 design

professionals around the world. That's because CADAM for IBM® mainframe computers, PRO CADAM™ for Unix® workstations and new MICRO CADAM Plus™ for high performance PC's are the most powerful integrated CAD/CAM/CAE systems ever.

Everyone wins with CADAM. The Ultimate CAD/CAM/CAE Production Tool. Call toll-free for the name of your nearest dealer. You could *Win the Ultimate Heartbeat.*

*No purchase necessary. Winner will be selected in a random drawing. Sweepstakes valid only within continental U.S.A. and Canada. Open to licensed drivers only. Void where prohibited by law. Entries must be postmarked by June 30, 1989 and received by July 14, 1989 to be eligible. Certain other restrictions apply; see complete rules printed on official entry forms, available only from authorized CADAM dealers.

CADAM . . .
The Ultimate CAD/CAM/CAE Production Tool

CADAM INC
A LOCKHEED COMPANY

Call toll free 1-800-255-5710

CADAM is a registered trademark of CADAM INC. Chevrolet and Corvette are registered trademarks of Chevrolet Division of General Motors Corporation. IBM is a registered trademark of International Business Machines Corp. MICRO CADAM and PRO CADAM are trademarks of CADAM INC. Unix is a registered trademark of AT&T. ©1989 CADAM INC

World Radio History

Desktop Publishing with an Apple II

Personal Newsletter, a desktop publishing program for the Apple II, has been upgraded. Version 2.0 includes a headline editor that produces 18- to 72-point headlines, double-strike and double-density printing, and a laser printer driver.

The program lets you lay out pages, create graphics, and integrate text. You can import text from most word processors, the company reports. Clip art, nine text fonts, and sample newsletters are included.

The program runs on the Apple IIe, IIc, and IIGS.
Price: $59.95.
Contact: Softsync, Inc., 162
Madison Ave., New York, NY 10016, (212) 685-2080.
Inquiry 1139.

Direct from SomeWare

A full-featured statistical package called EcStatic comes from a company named SomeWare in Vermont. The package has a transformation utility that lets you create new variables using programming statements along with a screen-oriented data editor.

The program also has utilities for managing data sets, including a merge utility that enables you to combine two data sets even if they share only some of the variables or cases.

EcStatic runs on any IBM PC, XT, AT, PS/2, or compatible with DOS 2.0 or higher, 512K bytes of RAM, and a floppy disk drive.
Price: $49.95.
Contact: SomeWare in Vermont, Inc., P.O. Box 215, Montpelier, VT 05602, (800) 451-4580; in Vermont, (802) 496-3173.
Inquiry 1118.

BrainMaker: A Neural Net Simulator

With this neural network simulation program, you can design, build, train, test, and run neural nets. The program lets you decide how you want your neural net assembled and trained, and it does the rest.

According to California Scientific Software, Brain-Maker runs at up to 500,000 neural connections per second and supports five types of linear and nonlinear neurons. The program offers I/O facilities for visual or symbolic data manipulations.

Sample neural nets are included on optical character recognition, speech synthesis, image recognition, and image enhancement.

To run the program, you'll need an IBM PC, XT, AT, PS/2, or compatible, 256K bytes of RAM, DOS 3.0 or higher, and a monochrome or color display.
Price: $99.95.
Contact: California Scientific Software, 160 East Montecito, Suite E, Sierra Madre, CA 90124, (818) 355-1094.
Inquiry 1120.

Circle 518 on Reader Service Card

World Radio History

How to Produce High-Performance Documentation

Strategies for Developing High-Performance Documentation is a seminar designed to help you create user documentation for software, hardware, and internal systems. Topics to be covered include project analysis, information analysis and management, effective use of team writing, on-line considerations, and reference-based training systems.

The seminar, which is sponsored by Information Mapping, uses analytical methodology based on human factors research and cognitive sciences. The seminar will be held in Boston from March 8 through March 10. At press time, the exact location of the seminar was still undetermined.
Price: $895; groups of two or more, $795 per person.
Contact: Information Mapping, Inc., 275 Wyman St., Waltham, MA 02154, (617) 890-7003.
Inquiry 1097.

The Long Island Computer Association

Established in 1976, the Long Island Computer Association is open to all amateur and professional computerists. Special-interest groups include Macintosh, IBM and compatibles, Commodore, Unix-C, education, accounting, lawyers, Apple, MIDI/music, and more.

The group has about 1500 members and a 9600-bit-per-second bulletin board system running 24 hours a day. The number for the BBS is (516) 561-6590. Last November, the group sponsored a connectivity seminar at the New York Institute of Technology, and in January, the group sponsored a dBASE III seminar.

The group generally meets at NYIT, Old Westbury Campus, on the third Friday of the month at 8 p.m. in Building 500, Room 508. Meetings in the recent past have featured representatives from Lotus, Ashton-Tate, and Brookhaven National Labs. The association publishes a monthly newsletter, *The Stack*, which is distributed free to members.
Annual dues are $12.
Contact: Long Island Computer Association, Inc., P.O. Box 71, Hicksville, NY 11802, or call Al Levy at (516) 293-8368.
Inquiry 1096.

The Third Annual Users Group Summit Meeting

The Association of PC User Groups' third annual users group summit meeting proved that coordinating an association of about 80 users groups, representing about 100,000 computerists, can be like a meeting of the United Nations General Assembly: a collection of representatives with diverse opinions on just
continued

World Radio History

about every topic on the meeting's agenda. And no matter how small the group, it wanted to be heard.

During the meeting, held in Las Vegas during Fall COMDEX 1988, users group officers and representatives debated everything from whether the APCUG's name should be changed to how a worldwide bulletin board system (BBS), announced during the meeting, will be run. But one of the main topics of discussion centered on membership and voting rights.

Several users group members questioned the APCUG's minimum membership requirement for groups to gain APCUG voting rights. According to Jerry Schneider, who serves as executive director of the APCUG, the requirement is necessary to

prevent users groups that aren't legitimate from taking advantage of the APCUG. But several people pointed out that of the approximately 6000 users groups in the country, 4000 have less than 50 registered members.

The smaller groups were concerned that although they could join the APCUG and take advantage of its benefits, they could not vote. Others were concerned that if the smaller groups were given the vote, they could gain too much influence.

The APCUG also discussed membership requirements. According to Mr. Schneider, users groups that want to join the APCUG should hold regular meetings and have a hard-copy or electronic newsletter or a BBS. Since the APCUG is not a

users group—it's more similar to a trade association—it's important that users groups that join the APCUG make sure their members receive information in a timely manner. "If groups don't have the mechanism to get information, they won't be providing the service," Mr. Schneider said.

Without a BBS or newsletter, information about the APCUG's activities would likely reach the contact person in the group and stop there. "Groups that [the APCUG] wants to exclude are groups that aren't really serving users," said Jonathan Rotenberg, president of the Boston Computer Society.

The APCUG is an association for users groups that want to improve communication among themselves and major

computer product vendors. Its membership includes international groups and groups from every region of the U.S. **Contact:** The APCUG, 9523 Burdett Rd., Burke, VA 22015, (703) 425-9896. **Inquiry 1087.**

Send Us Your Local News

BYTE is expanding its coverage of local events in the Northeast region. If you would like your event, conference, seminar, or users group covered, please send information to: Regional Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458. Please take into account a three-month lead time for your event.

Two PCs from Vendex

Vendex Technologies, the company that brought you the original HeadStart Turbo-888-XT, has released two computers: the Head-Start II Plus, which offers 100 percent VGA, IBM PC XT, and PS/2 Model 30 software compatibility; and the HeadStart III, bundled with several software packages.

The HeadStart II Plus runs at 9.54 MHz (switchable to 4.77 MHz) and comes standard with 640K bytes of RAM and a six-in-one VGA graphics adapter. It is powered by an 8088-1 microprocessor and comes with the Friend-Link Telecom System, which features a built-in handset for voice and data transmission and a 2400-bps modem.

The HeadStart III is based on the 80286 microprocessor chip running at 8 or 12 MHz with one wait state and 1 megabyte of standard RAM, expandable to 3 megabytes. The standard system includes a 32-megabyte hard disk drive with a 28-millisecond access time and 1-to-1 interleave. Bundled software includes Ashton-Tate's Framework II, the HeadStart Advanced Environment, 3D Perspective, Executive System's XTree, Splash!, Publish It, and Microsoft's MS-DOS 3.3 and GWBASIC. The system includes three full-length slots and a VGA, EGA, MCGA, CGA, MDA, and Hercules graphics card.
Price: HeadStart II Plus, $2295; HeadStart III, $2995.
Contact: Vendex Technologies, Inc., 40 Cutter Mill Rd., Suite 438, Great Neck, NY 11021, (516) 482-4255.
Inquiry 1082.

The HeadStart II Plus ships with a 40-megabyte hard disk drive.

Norton Utilities' Disk Doctor

The advanced edition of Norton Utilities 4.5, the latest version of the popular data-recovery package, contains a utility that automatically diagnoses and corrects problems with your hard and floppy disks. Available with the advanced edition only, the utility, called the Disk Doctor, lets you repair disks whether you're technically adept or not.

Both the advanced and standard editions of the upgrade include the Norton Control Center, with which you can control system settings, including keyboard rate, video mode, screen and palette colors, serial ports, stop watches, and time and date. You can store the settings in a file and reuse it as a replacement to DOS's Mode command.

Both editions also include a Batch Enhancer for creating interactive batch files with zooming dialog boxes. Other improvements include Data Protect, which saves information to help Quick Unerase and Unremove Directory recover heavily fragmented files and directories automatically; speed-key searching, rename directory, and EGA/VGA graphics support for 43- and 50-line modes of the Norton Change Directory; ability in the advanced edition's Format Recover to save an accidentally reformatted hard disk; and Safe Format, which you can use to reformat a floppy disk in as little as 3 seconds without permanently losing data previously on the disk.

Norton Utilities 4.5 works with the IBM PC, XT, AT, PS/2s, and compatibles with DOS 2.0 or higher and at least 192K bytes of RAM.
Price: Standard edition, $100; advanced edition, $150.
Contact: Peter Norton Computing, 100 Wilshire Blvd., Ninth Floor, Santa Monica, CA 90401, (213) 319-2000.
Inquiry 1079.

Convert Apple-Works Files to PC Formats

Cross-Works, a utility for AppleWorks, the most popular integrated software package for the Apple II, transfers and converts files into DOS-compatible files with all formatting information intact, SoftSpoken reports.

The program lets you convert AppleWorks word processor files into WordPerfect format, while retaining such formats as underlining, boldface, centering, and margins. Spreadsheets are converted to Lotus 1-2-3 files that have not only the correct data, but also intact spreadsheet formulas and cell formats. Similarly, you can convert AppleWorks database files to dBASE III–compatible files, complete with original field names. You can also use the program to transfer ASCII text files.

The package comes with an 8-foot cable that plugs into the Apple's or the IBM or compatible's serial ports. On one end, the cable has three connectors that fit the Apple IIe, IIc, and IIGS models. On the other end, a 9- or 25-pin connector fits your IBM PC, XT, AT, or compatible's serial port. SoftSpoken reports that if the computers are more than 8 feet from each other, you can transfer files by using a 1200-bps modem.

Cross-Works can transfer and convert on the fly a 30,000-byte AppleWorks spreadsheet file into a Lotus 1-2-3 WK1 file in about 20 seconds, according to the company.
Price: $79.95.
Contact: SoftSpoken, P.O. Box 97623, Raleigh, NC 27624, (919) 878-7725.
Inquiry 1074.

World Radio History

Com-Tek Data Systems, Inc.

Special Offer—while supplies last—A Bahamas Cruise 5 days and 4 nights with every 286 or 386 purchase.$675.00 value. Call for details.

SPECIAL COMPLETE SYSTEM

MODEL IQ-8088—I SYSTEM
10 mhz Mother Board—640K Ram
(1) 360K Floppy Drive
Mono Monitor—84 Keyboard
Multi-I/O—Phoenix Bios
Parallel, Serial, Game Ports
Printer Panasonic 1090i
MS-DOS ver. 3.3
20 MB Hard Disk   **$1295⁰⁰**

Model IQ 8088-S   **$895⁰⁰**

We also carry a complete line of major brand computer peripherals.

Custom configurations are our specialty. On site servicing available.

286-S

Configured to your specifications at a very reasonable price, our A/T is built to suit your every need.

YOUR PRICE

$1695.00

With CGA **$1970.00**
EGA **$2170.00**
VGA **$2370.00**

MODEL IQ 80286-T
Tower System **$1795**
Available in 3 case styles

STANDARD A/T MODEL IQ-80286-S

SPECIFICATIONS

MOTHERBOARD:	80286 microprocessor key selectable normal now available with Suntacs 80286 motherboard 6/12 MHz (8.0 mhz) and turbo (12.5 mhz) processing speeds, socketed for the 80287 math coprocessor, eight expansion slots (2 eight bit—6 sixteen bit), clock-cal, 1 meg. RAM included. Expands to 4 meg. Multi I/O and Award Bios included.
POWER:	200 watt, switching power supply with leads for 4 devices.
DISKS:	(1) 1.2 meg, half height, dual sided—quad density floppy drive. (1) 40 megabyte, half height, fixed disk drive. 40MS access time.
CABINET:	Full size AT style drawer cabinet with corporate security lock panel mounted reset switch, and status LEDs for turbo, power and fixed disk.
KEYBOARD:	Enhanced style, 101 keys with LEDs to indicate NUM locks and CAPS lock status, separate cursor pad, numeric touch pad, top mounted function keys.
DISPLAY SET:	Hi-res, text and graphics, monochrome card (Herc. compat.) hi-res, TTL amber monochrome monitor. 1 parallel port.
WARRANTY:	1 year on parts and labor limited depot warranty. 30 day money back guarantee if not satisfied with our product, for any reason.

*OPTIONS AVAILABLE

386-S

Powerful performer at a practical price. The ultimate for any business. Just in time for the tax season. Perfect for any accountant or business application.

**MODEL IQ 80386-S
—THE SUPER WISE CHOICE!**

$3195.00

with EGA monitor **$3495.00**
CGA **$3470.00**
EGA **$3670.00**
VGA **$3870.00**
Tower System **$3495.00**
— as shown —
80386-T **$3395.00**

STANDARD A/T MODEL IQ-80386-S

SPECIFICATIONS

MOTHERBOARD:	80386 microprocessor, user selectable (4.77, 8.0 and 16.0 mhz) upgradeable to 20 mhz. Processing speeds, socketed for the 80287 and 80387 math coprocessor, eight expansion slots (3 eight bit, 4 sixteen bit, and 1 thirty-two bit), clock-cal, 2 meg. RAM included, upgradeable to 16 meg. Includes Phoenix Bios and Multi-I/O card.
POWER:	200 watt, switching power supply with leads for 4 devices.
DISKS:	(1) 1.2 meg, half height, dual sided—quad density floppy drive and 1.44 floppy. (1) 80 megabyte, full height, fixed disk drive. Seagate.
CABINET:	Full size AT style drawer cabinet with corporate security lock panel mounted reset switch, and status LEDs for turbo power and fixed disk.
KEYBOARD:	Enhanced style, 101 keys with LEDs to indicate NUM lock and CAPS lock status, separate cursor pad, numeric touch pad, top mounted function keys.
DISPLAY UNIT:	Hi-res, text and graphics, monochrome card (Herc. compat.) hi-res, TTL amber monochrome monitor. 1 parallel port.
WARRANTY:	1 year on parts and labor limited depot warranty. 30 day money back guarantee if not satisfied with our product for any reason.

- COM-TEK HAS DESIGNED THEIR POLICY TO BETTER SERVE, HELP AND PROTECT THEIR CUSTOMERS.
- COM-TEK MAKES SURE ALL THEIR CUSTOMERS GET 1st PRIORITY IN SALES—SERVICE—CUSTOMER RELATIONS
- COM-TEK'S FRIENDLY STAFF IS BETTER TRAINED TO HELP AND SUPPORT YOU IN EVERY WAY POSSIBLE WITHOUT LEAVING YOU ON HOLD.

DEALER INQUIRIES WELCOME
30 DAY MONEY BACK GUARANTEE FOR ANY REASON—1 YEAR ON PARTS AND LABOR—LIMITED DEPOT WARRANTY

FREE UPS GROUND SERVICE IN USA SHIPPING

TAX FREE N.H.

SYSTEMS • SYSTEM OPTIONS • NETWORKING
CALL FOR SPECIFICATIONS

In N.H.: **603-363-8333**   **1-800-942-4255** Outside of N.H.
P.O. Box 221
Corner of 9 & 63, Chesterfield, N.H. 03443
Tech Support Call: 603-363-8334
XT and AT are registered trademarks of International Business Machines.

Big Performance-Little Price

The XT SYSTEM

$549.00

4.77/10 MHz 0 WAIT
256K RAM up to 640K
1-360 FLOPPY DRIVE
FLOPPY CONTROLLER
MATH CO-PROCESSOR SLOT
150 WATT POWER SUPPLY (UL)
101 ENHANCED KEYBOARD
LOCK/LED/RESET/TURBO CASE
MONOCHROME MONITOR. (TILT and SWIVEL)
MONOCHROME GRAPHICS ADAPTER

- **512K, PARALLEL, SERIAL, GAME PORT,**
  CLOCK/CALENDAR $645.00

- **BARE BONES SYSTEM $195.00**

The AT SYSTEM

$999.00

6/10 MHz 0 WAIT
512K up to 1024K
HARD DRIVE/FLOPPY CONTROLLER
1-1.2MB FLOPPY DRIVE
MATH CO-PROCESSOR SLOT
200 WATT POWER SUPPLY
101/ENHANCED KEYBOARD
LOCK/LED/RESET/TURBO CASE
PARALLEL, SERIAL, GAME PORT, CLOCK/CALENDAR
MONOCHROME MONITOR. (TILT and SWIVEL)
MONOCHROME GRAPHICS ADAPTER

- **286-12MHz SYSTEM with SI=15.5 $1099.00**

- **BARE BONES SYSTEM $425.00**

OPTIONS

20MB 65MS (ST-225)	**XT=$299.00**	**AT=$275.00**
30MB 65MS (ST-238)	**XT=$349.00**	**AT=$349.00**
40MB 39MS (ST-251)	**XT=$499.00**	**AT=$425.00**

MONOCHROME to CGA COLOR SYSTEM   Add: $199.00
MONOCHROME to EGA COLOR SYSTEM   Add: $399.00
CALL FOR FULL LINE OF ADDITIONAL COMPUTER ACCESSORIES
386-16MHz COMPUTER AVAILABLE $1999.00

The UPS

SPECIAL!

- **SAMSUNG 300 WATT 150VA-HOUR POWER CENTER**
  with UPS SURGE & FILTER $335.00

- **DAEWOO 300 WATT 150VA-HOUR UPS BATTERY**
  BACK-UP with SURGE & FILTER $299.00

Taesung Industries, Inc.

On site service available in some areas
24 hour turnaround service
No shipping charge with prepayment check
3% surcharge for Visa and MasterCard

1 year parts and labor warranty
30 day refund policy
Most items shipped within 24 hours
All machines have been burned in 72 hours

JASMINE COMPUTER SYSTEMS

P.O. Box 755
101 Main Street
Salem, N.H. 03079
• 1-800-848-1003 • (603) 893-3827
• 1-800-222-4160 (Dealers Only)

Circle 514 on Reader Service Card

Memory Expansion for the Sun-3/50

Two memory boards that attach to the Sun-3/50's CPU board let you expand the system's memory to 8 or 12 megabytes, exceeding the diskless machine's design limitation of 4 megabytes.

The SNME-350 daughtercard, available in 4-megabyte and 8-megabyte versions, attaches to the Sun-3/50's CPU board by using three small connector boards that plug into the CPU, memory management unit, and parity-check chip carriers using custom-designed adapter modules.

According to Clearpoint Research, the board's distributor, you can install the board without soldering.
Price: SNME-350/4MB, $2995; SNME-350/8MB, $4995.
Contact: Clearpoint Research Corp., 99 South St., Hopkinton, MA 01748, (508) 435-2000.
Inquiry 1077.

Tool for Windows Applications Developers

SofTools calls its CASE:W an expert system that includes a complete programming environment for generating pretested Windows code. The package features a front-end prototyper that describes an application program's windows and controls. The program then uses an inference engine to evaluate the prototype specification, applies the stored programming knowledge, and generates the Windows-based application.

CASE:W generates fully commented code and the nec-

Double or triple the Sun-3/50's memory with an SNME-350.

essary operating controls for any Windows application, according to SofTools. It supports menu bars, pop-up menus, and dialog boxes.

The program has a facility that automatically regenerates programmer-added codes into all future versions of CASE:W, so you need develop code only once, no matter how many times the program is redesigned.

SofTools recommends that you use the program on an IBM PC AT or compatible with at least 2 megabytes of main memory. The package also requires Microsoft's Software Development Kit, the C Compiler, Make Utility and Linker, and a DOS- or Windows-compatible text editor. The company also recommends Microsoft's CodeView.
Price: $1495.
Contact: SofTools, Inc., One Dunwoody Park, Suite 130, Atlanta, GA 30338, (404) 399-6236.
Inquiry 1086.

Help and File Management for Windows

Two utilities for Microsoft Windows can help you manage files or create your own help files that you can link to your favorite Windows application.

Jaws is a file management utility that lets you compress and encrypt files, view and search hexadecimal files, compare files, and customize colors. The program provides two file-encryption methods: the exclusive (XOR) encryption algorithm and the Federal Data Encryption Standard. With Jaws' library manager, you can group, compress, and encrypt files before you transmit them via modem.

Help, the second utility, lets you create help files for Windows applications. The program is set up like a reference book, with a table of contents and index. The help files are compressed, and you can copy information to the Clipboard.

Both memory-resident pro-

grams require 100K bytes of RAM, any version of Windows, and DOS 2.0 or higher.
Price: Jaws, $35; Help, $35.
Contact: The R Company, 2170 Georgina Ave., Santa Monica, CA 90402, (213) 393-9992.
Inquiry 1076.

Parallel-to-Parallel-Port File-Transfer Program

PDQ[2], a file-transfer program for your DOS-based microcomputer or laptop computer, avoids protocol and serial connection problems in transferring files by using a parallel-port-to-parallel-port connection.

To transfer your favorite spreadsheet or utility to or from a 3½-inch floppy disk drive, you connect a 6-foot cable to both ports and load the software, and the files are ready to go.

According to Sophco, the program's distributor, PDQ[2] can transfer files at the rate of almost 2 megabytes per minute. The program has context-sensitive help menus. A built-in test routine can verify that you have a good connection between the two machines.

PDQ[2] is not a memory-resident program. It comes on 5¼- and 3½-inch floppy disks. It works on the IBM PC, XT, AT, PS/2s, and compatibles and requires 256K bytes of RAM and DOS 2.0 or higher. Sophco reports that the program works with DOS 4.0.
Price: $99.
Contact: Sophco, Inc., P.O. Box 7430, Boulder, CO 80306, (800) 922-3001; in Colorado, (303) 444-1542.
Inquiry 1072.

World Radio History

World Radio History

BYTE editors' hands-on views of new products

MegaMate

MKS Make, Lex, and Yacc

Wordbench

DataSentry

Language Systems
FORTRAN

Sourcer

The Little Drive That Could

There's no escaping the world of multiple floppy disk formats, especially if you get your disks from different sources or switch back and forth between a laptop computer and an older (pre-PS/2) desktop computer. If you need to work with 3½-inch floppy disks but have only a 5¼-inch floppy disk drive in your main computer, the **MegaMate** 3½-inch external drive is a solution in the true sense of the word. And it's about as easy a solution as you can get, short of buying a new system with a 3½-inch drive built in.

MegaMate, which is a sleek little disk drive designed by someone who realizes the value of desktop real estate, works with any IBM PC, XT, AT, or compatible that can spare a half-size slot. Hooking it up is a snap. After checking the jumper blocks to make sure they're set correctly (the unit I received came properly positioned), you just slip the MegaMate card into the expansion slot, screw it down, and plug the drive into the port. Installing the software, on either a hard disk or a system disk, is equally simple.

MegaMate functions just like a native drive in that you can access it from applications and use DOS utilities such as COPY, DIR, CHKDSK, and CD. FORMAT, DISKCOPY, and DISKCOMP will work only with DOS 3.3 or higher, but the MegaMate software has its own formatting command that works with DOS 2.0 and higher.

The drive can initialize 3½-inch disks in 720K-byte and 1.44-megabyte formats (the latter being the default). The unit can detect what kind of disk it's looking at, so you have to specify only density when you're initializing a floppy disk. Micro Solutions said it will soon have a utility program that will let you format disks as a background task, but that update wasn't ready when we went to press.

I've been using MegaMate as my regular floppy disk drive for several weeks now, and I haven't seen any problems in terms of operation or reliability. The external unit has blended right in with the AT clone I use at work, doing none of the things that sometimes happen when you add a piece of hardware. The only weird thing I've seen is found in the boot-up screen message that flashes MegaMate's designation. The message, a two-liner, I think, says that the drive is available as drive D, but the thing whizzes by so quickly, and amidst a whole screen full of text, that I had to call it up 20 times before I could see the message that said my MegaMate was drive D.

At $349 for drive, card, and cable, MegaMate is priced comparably with other external drives. (With higher-capacity floppy disk drives coming soon, like Konica's 10-megabyte unit, I wouldn't want to invest much more than that in an extra disk device.)

If you've got to read from or write to 3½-inch disks and can't afford a more modern computer with built-in drives for the smaller floppy disks, the MegaMate is worth the money. Sure beats having to bother a coworker with a 3½-inch floppy disk drive.

—*D. Barker*

Unix Tools for DOS

Unix tools for MS-DOS users? Who wants them! Wait, there must be something valuable here, or Mortice Kern Systems would have been out of business years ago when it started shipping packages of classic Unix tools that are written and compiled for MS-DOS.

The company now has two new programs. **MKS Make** is traditionally used for defining the dependencies of program source code files so that appropriate recompilation and linking can take place when one or more modules or libraries have been modified. However, its use is not limited to programming. Many documents have dependencies on other documents. If, for example, parts of accounting sheets are included in other documents, changes in the accounting requires reformatting the documents. This process is automated through the use of MKS Make. Containing Mortice Kern's RCS (Revision Control System, similar to Unix sccs), MKS Make gives the same sturdy programming and documentation manage-

continued

ment to MS-DOS users that Unix developers have been boasting about for nearly two decades. It can use both Unix and Microsoft control definitions.

MKS Make supports Turbo C, Microsoft C, WATCOM C, and other compilers, libraries, and linkers. It works with the MKS Toolkit Korn Shell, as well as with standard DOS COMMAND.COM.

The second new program is **MKS Lex and Yacc** (for "yet another compiler compiler"). Lex and Yacc are tools used specifically for software development. Lexical analysis is the set of operations that determines the nature of input in an application program. For example, a lexical analyzer determines whether you are inputting text, a value, a function, a macro, or a command when you run a spreadsheet program. Lex is a program to generate lexical analyzers. Most lexical-analyzer generators do not create fast or com-

pact program modules. Instead, they are used just for prototyping or generating stand-alone data translation systems.

Yacc has been used for generating entire application programs as well as compilers. There are many applications for rule-based systems other than compilers. When an applications developer specifies a set of rules and machine states and/or operations, Yacc can take the specifications and generate the source code for a program or module that follows these definitions.

MKS Lex and Yacc supports the following compilers: Turbo C, Microsoft C, WATCOM C, and others. It is not copy-protected, and there are no royalties or run-time fees.

Mortice Kern has developed an entire suite of more than 130 programs (MKS Toolkit, $199) that perform the same functions as their equivalents in Unix. I have been using this package for years. None of the

MKS programs contain any AT&T code, but they religiously follow the System V.3 parameters and operations.

Other products from the company include MKS SQPS (SoftQuad Publishing System), compatible with AT&T Documenters Work Bench, $495; drivers for Hewlett-Packard LaserJet printers and PostScript devices, $200; MKS Vi ($149), a total implementation of the classic Unix screen editor; and MKS AWK ($99), which contains all the features of the new version of the report-generation language from the gurus of Bell Labs. The MKS Toolkit includes everything but SQPS, Make, Lex, and Yacc.

So who uses Unix tools in MS-DOS? I suspect that 50 percent are users with Unix experience who need to work in the MS-DOS world, and the other 50 percent are MS-DOS users wanting to learn more about Unix.

—Ben Smith

Wordbench:
Tools Designed with Writers in Mind

Word processors for writers of prose typically provide an outliner, a spelling checker, and a thesaurus. Addison-Wesley's **Wordbench** program does all that and more. There's a note-taker that behaves like an electronic stack of 3-by-5 index cards; a reference tool that helps you collect and arrange bibliographic citations; a viewer that you can use to refer to your outline, notes, and citations while working on a document; and even a tool called the brainstormer, which helps you get past a writer's block.

The word processor itself is solid but unspectacular. It supports many of the usual text-editing operations, paragraph formats, and modes of emphasis. The size of your document is limited by available RAM, though, and some conveniences—like a single-

key function to delete a line— are missing. If you're already committed to a word processor, it's unlikely you'll want to switch to this one.

But the word processor does its job, and the environment it works within has some unique

features. One particularly handy feature is the interaction between the note-taker and the outliner. The problem with typical outliners is that you can attach text only to outline items; that compels you to create a skeleton to which you

can then add meat.

Wordbench recognizes that sometimes the meat comes first. Thus, you can use the note-taker to capture thoughts that don't yet fit into your outline. The outline and notes
continued

evolve separately; as the structure clarifies, you link notes to outline headers. When you're ready to write a first draft, you merge the outline and notes together into a document.

The brainstormer provides a set of tools that people who teach writing will find both amusing and useful. Techniques include free writing, nutshelling, and goal setting. Free writing means that you set a timed interval—say, 2 minutes—then write furiously until the clock runs out. The program won't let you back up or edit, and if you pause for longer than a few seconds, it tells you to keep writing. Nut-

shelling encourages you to write a concise summary of your topic, and goal setting helps you define your audience and point of view.

Wordbench isn't the fanciest word processor around, but its creators have thought hard about the process of writing and have built tools that flexi-

bly support that process. It's a very friendly package that may well appeal both to professional writers who are unfamiliar with computers and to educators who are looking for software that can really help their students learn how to write.

—*Jon Udell*

Put a Positive Lock on Your Data

It's 2:00 a.m.; do you know where your data is? There's currently a great deal of concern about computer security. Just listen to the news. And you don't have to work for the government to be anxious yourself.

Nearly every business has computerized information that shouldn't go beyond its four walls. It's all too easy for someone to walk in, copy some files to a floppy disk, and walk away with a company's secrets.

Data-security packages, both software and hardware, aren't new. But there's a problem: The most effective ones are expensive and inconvenient to use; the inexpensive ones are marginally useful at best. But the aptly named **DataSentry** from Rainbow Technologies has changed all that. It's economical, it's easy to use, and it provides virtually foolproof security.

DataSentry is a variation of those "hardware keys" that are often used as copy-protection devices for custom software or other expensive packages. Indeed, Rainbow Technologies is a leading supplier of the keys.

The 2-inch long, 1.5-ounce DataSentry plugs into the printer port of your IBM PC or compatible, and it lets you add copy protection to individual files or even to entire directories. Rainbow Technologies claims that DataSentry does not interfere with the operation of your printer, which plugs into the other end of DataSentry. And I found that to be true.

Since I don't have any company secrets on my computer, I decided to protect the next best thing: my résumé. After I plugged DataSentry into my system's parallel port, all I had to do to complete the process was to copy a couple of utilities to my disk and type SEAL, followed by the filename of my résumé.

DataSentry then encrypted my file, producing a new file

that was also compressed to nearly half its original size. The utility also erases the original file. To make things even more secure, it goes beyond the normal MS-DOS DELETE (which leaves the data on the disk) by completely replacing the original data with nonsense characters.

When I removed DataSentry and attempted to edit the file, all I had was a screen

full of incomprehensible gibberish. To get my résumé back, I had to reinsert DataSentry and type OPEN and the filename. And my file came back.

You can even add a password when you seal a file. But without DataSentry in the parallel port, there's virtually no way to access the file unless you have a degree in cryptography and access to a supercomputer.

The basic DataSentry uses what Rainbow Technologies calls a "proprietary algorithm" to encrypt your data. There's also a version that uses the government's highly secure Data Encryption Standard, but it can't be sold or exported beyond the U.S.

DataSentry is not completely transparent. You have to remember to reseal your files after you're done working with them. And if you take DataSentry home with you at night and forget it the next morning, you're up the creek.

Besides the individual DataSentry keys, Rainbow Technologies also offers a multilevel system, so a department manager can have a key that will open the files of all employees. And you can even send encrypted files over telephone lines by purchasing keys with identical internal codes for use on both sides of the line.

If you lie awake nights worrying about the safety of your data, DataSentry will let you get some rest, especially if you put it under your pillow.

—*Stan Miastkowski*

The journey to discovery can sometimes be a long and arduous one.

Especially when it comes to riting database applications.

And especially when you're not programmer.

All of which is precisely why we uilt dBASE IV™ with an innovative ew feature we call the Control Center.

The Control Center is a logical, intuitive window from which you perform all your key database operations. A constant frame of reference, if you will, for harnessing all of dBASE IV's formidable power.

Through the Control Center, you get simple and understandable access to every one of dBASE IV's modules and functions. Like the powerful Report Writer. The Applications Generator. And much, much more. So that the time it takes to develop your applications can be drastically reduced.

You see, at Ashton-Tate we think that people tend to be more successful when they're able to see the opportunities for discovery right in front of them.

Want to see for yourself? Just visit your local authorized Ashton-Tate dealer for a demonstration, or call **800-437-4329** * **ext. 2934** for details on videotapes, demo disks, upgrades, and more.

ASHTON · TATE®

INTRODUCING dBASE IV

"It's a whole lot easier when you can see what you're doing."

A FORTRAN
for the Mac Forces

Language Systems FOR-TRAN is a compiler that runs under the Macintosh Programmers' Workshop. You can purchase the MPW Unix-like shell from Language Systems configured to run the compiler, or you can use the MPW version from Apple with the Install program that is included with the FORTRAN compiler.

The compiler produces object code that is linked using MPW's linker, which can also link resources created by the Rez tool. This method lets the programmer use resources from within the program that have been previously created and debugged, resulting in significant savings in time and development effort.

Using FORTRAN again was like going to a 20-year high school reunion. We hadn't been near it for a while, and we were curious to see how time had changed it. The compiler supports the ANSI 77 syntax, and it has extensions to it designed for the Mac. One of the most obvious is that there can be 31 significant characters in a name, as opposed to the ANSI 77 allowance of 6 characters. The ANSI 77 version allows only

72 characters per source code line (from the days when cards were 80 columns), but the compiler can take up to 255 characters per line (which makes sense for the Mac).

One of the annoying things about FORTRAN—also annoying in this implementation—is the assumption of a standard I/O stream. That is, the output of the program goes to a fixed window that has basic menus attached to it and a goAway box. The programmer can't change the output window without extra work with MPW itself (attaching extra menus, for example). If output from the program exceeds 1000 lines into the window, it is automatically saved into a labeled file.

Other supported I/O units include the keyboard and the printer. Thus, mouse movement can't be used in a program without the programmer doing the Toolbox calls and the necessary overhead in the program. Whether or not this is significant for the average FORTRAN program is debatable. But it does reduce the "Macishness" of the resulting program without extra work.

You can easily implement the structures necessary to use Macintosh formatted records and communicate with routines written in other languages, like Pascal.

The compiler seems to be designed to make the porting of previously written FORTRAN programs easy. This is commendable, especially for beginning programmers using the language from a standard educational text. But the programmer should not expect the language to add most of the standard Mac features automatically. It seems a way to make the existing body of FORTRAN code usable on a Mac, and it succeeds.

—*D. Barker and Larry Loeb*

THE FACTS

Language Systems FORTRAN
$295; $200 without MPW

Requirements:
Macintosh Plus, SE, or II with 1 megabyte of RAM and a hard disk drive with at least 5 megabytes of free space.

Language Systems, Inc.
441 Carlisle Dr.
Herndon, VA 22070
(703) 478-0181
Inquiry 1039.

Sourcer Magic Disassembles Machine Code

Assembly language programmers may think they have to work with stone-age computing tools, but **Sourcer** is a machine code disassembler with the interface of the best consumer-oriented programs and much of the built-in intelligence of an expert system.

A disassembler program reverse-engineers executable programs, producing assembly language listings or source code. You use it to study how programs work, to modify them, or to check them for viruses, improperly appropriated code, and inappropriate operations. Typically, the only control you have over a disassembler are the command-line options.

Sourcer has changed that; it has an options-setting screen that lets you select formats, comment types, analysis methods, and more than a dozen other options. The screen then displays an example of your selections as you make them. It even has a help screen! (This may not be extraordinary to the average user, but it is remarkable for an assembly language programming tool.) Having made your choices, you press G (for "Go"), and Sourcer displays the progress of the disassembly process in a graph at the bottom of your screen.

The input to Sourcer can be .EXE, .COM, .OBJ, or .BIN multsegment files (up to 250 segments); device drivers; program overlays; system files (.SYS); or any RAM or ROM address in the first 1-megabyte area. Sourcer will disassemble code from the 8088/8086, 80186/80188, V20/V30, 80286 (privileged and real modes), 8087, and 80287 processors to the appropriate instruction sets.

Sourcer uses a multiple pass-data analysis/code-simulation sequence to separate code from data, analyze the code for BIOS calls, and analyze the data for strings. The output from Sourcer consists of source code or listing files. Because of the length of comments, you should use a 132-character printer to list the output files.

Disassembly of a 13K-byte DOS utility produced a 130K-byte .ASM file in 4 minutes (on a 10-MHz 80286 PC-DOS machine). An attempt to disassemble an EGA demonstration program produced: "WARNING: Packed file, conversion quality low." But it still managed to produce a usable listing, although it obviously had problems following the flow of the program.

Sourcer comes with a well-written 76-page manual that explains operation and warning/error messages. You can also buy an optional BIOS preprocessor.

—Ben Smith ∎

THE FACTS

Sourcer
$99.95; with optional
BIOS preprocessor,
$139.95

Requirements:
IBM PC or compatible
with 512K bytes of RAM
and DOS 2.0 or higher;
supports CGA, EGA,
VGA, MDA, and
Hercules adapters.

V Communications
3031 Tisch Way
Suite 905
San Jose, CA 95128
(408) 296-4224
Inquiry 1050.

It's always the raw speed of an application that seems to catch one's attention. But when you think about it, speed alone doesn't always determine the amount of useful work you can produce.

That's why in dBASE IV,™ we added more than just a blazingly fast .EXE compiler. We added a total of 310 new or enhanced commands and functions that help you complete your applications in less time than you ever thought possible.

And without your having to become a full-fledged programmer to do it.

Of course, the compiler by itself lets you run applications significantly faster than is possible with dBASE III PLUS.™ But you'll also find yourself working much more efficiently through dBASE IV's intuitive new Control Center. As well as by using its Applications Generator. Or the Report Writer. The Forms Generator. And much more.

After all, when it comes to developing dBASE IV applications more quickly, there's not just one way to do it.

There is one way to find out more, though. Just visit your local authorized Ashton-Tate dealer for a demonstration, or call **800-437-4329* ext. 2934** for details on videotapes, demo disks, upgrades, and more.

⚞ ASHTON·TATE®

INTRODUCING dBASE IV
Circle 20 on Reader Service Card
(DEALERS: 21)

"Have you noticed that program speed is not the only reason you get things done faster?"

Power

T5200: 20MHz 386 processor, 2 internal IBM® compatible expansion slots, VGA gas plasma display with VGA monitor port, 40MB or 100MB hard disk, 2MB RAM standard expandable to 8MB, 1.44MB 3½" diskette drive.

At Toshiba, we're not only committed to making computers more portable, but also to making portables more powerful.

Which is why, in our effort to constantly improve and refine our machines, we've added three new computers to what is already the most complete family of truly portables available.

Each designed to be powerful enough to take on the increasingly complex tasks that face today's sophisticated PC users.

First, the T1600 which weighs under 12 pounds and which is the fastest battery-powered computer we've ever made.

Second, the T3100e, the successor to our most popular machine—the T3100/20. We've made it nearly two pounds lighter and a lot faster—we've even added expansion capabilities. About the only thing we didn't add was more size.

T1600: Battery-powered 286/12MHz, coprocessor socket, 20MB hard disk at 27msec, 1.44MB 3½" diskette drive, 1MB RAM expandable to 5MB, detachable backlit EGA compatible LCD, removable rechargeable battery pack.

Surge.

And finally, the T5200, which has enough power to replace virtually any desktop PC.

But we haven't just concentrated on power and portability. We've also constantly looked for ways to make our machines more durable, more reliable, and easier to use — down to the 800 number our customers can call for help with any technical question that might come up.

We figure that's what our users demand.

And it's by anticipating the growing needs of our users that we have continually found ways to make our machines weigh less and do more. So you can work wherever you want and however you want.

All of which might make it tempting for some people to abandon their desktop for the convenience of portability. Go ahead.

We've given you the power to do it.

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

In Touch with Tomorrow

TOSHIBA

Toshiba America Inc., Information Systems Division

Paradox 3: Neither Enigma nor Riddle

Borland's extensive
upgrade of Paradox makes
it a major contender
for the number-one
DBMS spot

What might be termed "the battle of the databases" has rapidly become a battle of major software companies. Though a flock of small-to-medium companies sells database managers, the major activity in "serious" relational database managers is pretty much confined to large companies with the commitment (read "bucks") to devote the time and resources needed to effectively develop these complex applications.

Ashton-Tate's dBASE, long considered the state of the art in database managers, is now in version IV. A measure of its acceptance is the flock of act-alike products nipping at its heels. In fact, as this was being written, Ashton-Tate filed a lawsuit against Fox Software, charging that its FoxBASE package violates the "look and feel" of dBASE.

But nobody can accuse Borland International's Paradox of being a dBASE clone. From a user-interface and design-philosophy standpoint, it's a unique animal; and that's where its power and versatility lie. Introduced in 1985, Paradox quickly became a serious contender in the high-end database manager market, especially after Borland acquired the package's original developer (Ansa Software). Borland brought considerable marketing clout to the equation, and with the release of version 3, the Paradox series has reached a point where it has the potential of overtaking dBASE IV.

Paradox 3 is a major upgrade from version 2, with numerous enhancements and brand-new options. Like any relational database manager, it's feature-packed and extremely complex. We looked at the prerelease beta version of Paradox 3, and this First Impression will concentrate on its new and improved features. Although the version we discuss here didn't have it, Paradox's upcoming SQL (Structured Query Language) engine is a potentially significant development. (See the text box "Paradox Tests the SQL Waters" on page 110.)

QBE Gets Better

Paradox was the first PC-based software package to incorporate a true Query-By-Example approach to conducting database operations. And until recently, it was the only one. (The latest release of dBASE incorporates QBE, although we hadn't seen it at this writing.) IBM developed the QBE definition in the mid-1970s. Until Paradox appeared, QBE was used mainly in minicomputer and mainframe database applications.

In Paradox 3, QBE is the core of the user interface. By now, Paradox's table-based, fill-in-the-blanks approach to locating data is familiar to thousands of users. As QBE implies, performing a query involves giving examples in a query table. Paradox has always had a rich selection of query options, and version 3 has added a few more.

One of the major new features of Paradox 3 is the choice of how to join data that you query from multiple tables. With any relational database manager, a common problem with a query that takes data from several related tables is what to do with mismatches. When you join two or more tables, should records in one table with no matches in the other be included or excluded in a new table? Most database managers simply don't offer you a choice. They either always exclude the records, something called an *inner join*, or they always include them, which is called an *outer join*.

An inclusion operator in Paradox 3 now gives you the option of easily performing either an inner or an outer join. By default, Paradox, like most database managers, does an inner join of tables. But in Paradox 3, placing an exclamation point in the query tells the program to do an outer join. To be fair, it is possible to do outer joins with most database managers, but it requires several steps that are time-consuming and usually require more than a bit of programming skill.

In Paradox 3, Borland has extended the original QBE definition by including set operations. They add a collection of logical operations to a Paradox query. Because they allow you to perform highly complex queries in a single step, set operations are another step toward making Paradox 3 easier to use.

Creating sets is a simple matter of placing the keyword SET in the leftmost column of a query form. SET defines information about sets of records, which you can then compare with other data. The SET command works in conjunction with four new operators: ONLY, EVERY, EXACTLY, and NO. For example, in a typical order-entry system, you can use a one-line query to ask "Which products have been ordered *only* by customers in New Hampshire?" Paradox 3 also has an OR operator that lets you perform additional one-line queries such as "Who lives in California OR New York?" In prior versions, this required a second line in the query.

Another new feature, though minor, also adds to Paradox 3's usability. While you've always been able to perform calculations within a Paradox query, you had no control over what the new field in the resulting table was named. That problem has been solved with an AS operator that lets you name the new field.

Improving Relations

Although relational database managers have become an industry standard, they

Paradox Tests the SQL Waters

One of the big promises of the networked computing environment of the 1990s is the marriage between local databases on personal computers and large databases on mainframes and minicomputers via the Structured Query Language. The idea is that users of local databases such as Paradox or dBASE will be able to issue queries in the "native language" of the database. Those queries are then translated into SQL and transparently transmitted to remote databases, where the query is processed and the data is retrieved and sent back to the local database. Ideally, the data on the remote database appears as just another table on the local database.

This type of transparent connectivity to SQL will offer some major advantages to database users in large organizations. It will be possible not only to centralize important data on the main-

frame or minicomputer, but also to allow users access to that data on their PCs with all their inherent advantages such as independent processing, graphical interfaces, and a large and flexible software base. Centralization will also greatly improve the security, consistency, and integrity of the organization's data.

Although it wasn't available in the beta version we tested, Paradox SQL connectivity was demonstrated by Borland at a recent industry trade show. The demonstration simultaneously accessed databases in Oracle's SQL, IBM's OS/2 Extended Edition, Novell's XQL, and the Sybase SQL Server.

Borland has developed a core SQL "engine" for Paradox that translates queries issued using the Paradox Query-By-Example feature. The user issues a QBE query. If a remote database is present, the Paradox SQL engine detects its

presence, translates the query into SQL, and sends it to the remote database. The result of the query appears on the PC as a Paradox Answer table.

Like most so-called standards, the leading versions of SQL now on the market are incompatible with each other. But these incompatibilities are apparently inconsequential to the SQL connectivity engine built into Paradox, which simply has different translators that handle the different versions of SQL.

Paradox users won't be aware of this process because the remote data will appear as a Paradox table regardless of the source of the remote database or its version of SQL. The Paradox SQL connectivity package will also include remote data entry into SQL databases and the "pass-through" of PAL (Paradox Application Language) syntax to SQL queries.

---

do have inherent problems. One typical difficulty is aptly named the "order-entry problem," because it manifests itself during the design of common order-entry databases. Designing a *single* form that enters a sales order, displays customer information, and lets you enter an unlimited number of items isn't as trivial as it sounds. The interrelationships between items can quickly become very complex.

Another problem of relational databases is what designers call *referential integrity*, which is making sure that explicit relationships between records are maintained. For example, a common goal is to make sure a customer record can't be deleted if that customer has an outstanding order.

Problems like these aren't insurmountable, but in most database management systems they require a great deal of programming knowledge to solve. Paradox 3 has added several new features to forms design that directly address these inter-table relations and dependencies.

The most important of these is the addition of multitable forms. A single form in Paradox 3 can now display data from multiple tables. Also added are linked tables within forms. For example, if you're using a form that uses multiple tables of data, you can link or unlink the tables. If you leave the tables unlinked, you can scroll through multiple tables in-

dividually on the form. But if you link the forms, the multiple tables remain related to each other. As you move through a database, Paradox 3 automatically keeps track of the interrelationships of the data. In a typical linked application, Paradox 3 won't let you delete a record when other records depend on it. This serves to eliminate the referential integrity problem.

Graphically Speaking

One long-awaited addition in Paradox 3 is its advanced graphics capabilities. While graphics are available in other database managers—either integrated into the package or as extra-cost add-ons—Paradox 3's graphics are singular in both ease of use and versatility. The program has a large selection of graph types, including the standard pie charts, line graphs, bar graphs, and x,y graphs.

Creating a graph is simple: You use Paradox's ROTATE command to rotate data columns into the order you want them to appear and then press a single key. Paradox 3 instantly graphs the data, choosing the type of graph that's most fitting for the type of data you're graphing. If you want to change the graph type or otherwise customize the graph, you have a full contingent of graphics menus.

Besides printing graphs on most printers, you can also export Paradox 3 graphs in a variety of industry-standard

graphics file formats for further work. And if you need to give presentations, a WAIT command within the program lets you create a "slideshow" of Paradox 3 graphs.

Crosstab Control

A new feature that's tightly coupled to Paradox 3's graphics capabilities is its crosstab function. The package's normal table view of data is usually the best way of organizing and viewing information. But especially when you want to graph data, a spreadsheet-type view of the data is more useful. The crosstab command instantly converts a standard table into a matrix-type spreadsheet view. For example, crosstab converts a table that contains customer names, item names, and amount sold into a new table with item names as column headers and counts in columnar format. The result of graphing this type of data is easier to interpret.

Depending on how you've organized your databases, crosstabs don't always improve graphs. However, the crosstab spreadsheet-type view is useful with Paradox 3's import/export features. To Paradox's 1-2-3 connection, Borland has conveniently added the ability to import data from its own Quattro spreadsheet.

Adding Color to Your Life

In previous versions, Paradox didn't use color very well, simply giving you a

white-on-blue display on a color monitor. But version 3 lets you use color effectively, indeed. You can specify colors for fields, forms, menus, and just about every other component of the program.

This can be more than a little useful. For instance, you can have negative numbers show up in red, highlight the most important parts of a form, or even highlight a current table image in a different color. Though it sounds trivial, Paradox 3's color control can help you make your applications easier to use and more intuitive.

Your Programming PAL

The Paradox Applications Language has always been one of the most powerful database programming languages in the business. And the version of PAL that comes with Paradox 3 has a host of enhancements, including new commands and additional arguments for existing statements. Space restrictions preclude our going into detail about them; but as you might expect, most of the added commands support Paradox 3's new features.

PAL, like the programming languages of Paradox 3 competitors, requires a goodly amount of programming experience to use effectively. Not coincidentally, it also requires a large time commitment to learn. Since PAL programming is essentially a full-time proposition, serious PAL programmers are likely to be found in large companies with a heavy commitment to Paradox. For the rest of us, the Personal Programmer is included with every copy of Paradox 3. Also updated for the new version, the Personal Programmer is an interactive applications generator that lets you create turnkey menu-driven custom applications.

Appearing on the Network

Last but not least in Paradox 3 are a few small enhancements to its multiuser capabilities. Ever since Paradox 2, the package has been one of the few off-the-shelf applications that easily handle networks with no modifications. Not surprisingly, all the enhancements and new features mentioned above also work on networks.

One intriguing application that uses Paradox 3's new graphics capabilities is the ability for a manager to use Paradox 3 on a network to monitor workflow. For instance, a system manager can use the package to graph data that other users are entering. And because Paradox 3 performs network updates in real time, the graphs will change as data is entered.

Paradox 3

Type
Relational database manager

Company
Borland International
1800 Green Hills Rd.
P.O. Box 660001
Scotts Valley, CA 95066
(800) 543-7543
(408) 438-8400

Format
Unavailable in beta version

Language
C

Hardware Needed
IBM PC, AT, XT, PS/2, or compatible with 512K bytes of memory, DOS 3.0 or higher, a hard disk drive, and one floppy disk drive. To display graphics, a CGA, EGA, VGA, or Hercules graphics adapter is needed.

Documentation
Unavailable in beta version

Price
$725

Inquiry 1067.

The Contender

Individually, the enhancements to Paradox 3 are impressive. Taken together, they result in an exceptional database management system that has all the capabilities necessary for even the most advanced applications. Yet unlike most of its major competitors, Paradox 3 is intuitive enough so that it can be used for many applications without a large learning curve. Still, getting the most from Paradox does require a commitment to learning. But its extensive on-line help, coupled with a set of updated and rewritten manuals, makes that job far from arduous. Priced at $725, Paradox 3 isn't exactly inexpensive, but it's comparable with other high-end, high-powered database management packages. With the release of Paradox 3, Borland now has a product that easily matches, and in some ways surpasses, the competition. ■

Stan Miastkowski is a BYTE contributing editor, director of K+S Concepts (a documentation and consulting firm), and editor of the "OS Report" newsletter. He can be reached on BIX as "stanm." Nick Baran is a BYTE technical editor based in San Francisco. He can be reached on BIX as "nickbaran."

NATURAL LANGUAGE INTERFACE

THE COMPUTER CHRONICLES
makes keeping up with the information age easy.

Stewart Cheifet and **Gary Kildall** are your co-hosts for a weekly half-hour television program aimed at computer users, owners, educators and computer industry professionals. If you're looking for help in finding out what's new and what's news, tune in to THE COMPUTER CHRONICLES every week on your local public television station.

The latest developments in personal computer hardware and software are discussed, demonstrated and reviewed. The industry's top developers, executives, analysts, and journalists keep you on the cutting edge of the computer business. Paul Schindler reviews new software and Wendy Woods goes on location to find out how users are taking advantage of new developments.

Topics this season include Local Area Network Software —Fax Boards—New Generation Spread-sheets—Color Printers—SQL—Hard Disk Management—Hypertext—Display Technology—Programming Languages—PC Kits—Bus Wars—The Battle of the Operating Systems—Foreign Language Software—Artifical Intelligence—Neural Network Technology—Megaherz Mania—Software Piracy—Personal Information Management Software—Optical Storage—and much more.

Each week THE COMPUTER CHRONICLES looks at the top stories of the week in Random Access, a news segment designed to keep you informed about the latest developments in the computer industry.

THE COMPUTER CHRONICLES, a weekly half-hour of public television that just might be the help you need.

THE COMPUTER CHRONICLES IS NOW ON BIX

BIX, the Byte Information Exchange, now has a conference for The Computer Chronicles.

Now you can communicate directly with the staff of The Computer Chronicles to suggest topics for future shows or to request information or details on past shows.

Once on BIX, just type "COMP. CHRON" at the colon (:) prompt.

FIRST IMPRESSIONS ■ Nick Baran

The Mac SE Takes Off

The newest version
of the Mac SE
runs rings around
its predecessor

I f you're interested in brute Mac power, but you lack the desk space or the credit line for a IIx, then the new Mac SE/30 may be just the machine for you. This newest Mac is a major addition to the Apple Macintosh product line. Featuring a 68030 CPU and a 68882 floating-point unit (FPU), the Mac SE/30 puts the heart of the Mac IIx into the body of a Mac SE. In fact, the only thing the Mac IIx has that the Mac SE/30 doesn't have is the NuBus expansion card cage.

The new machine is the latest in just one family of Macs. Apple now talks of a "modular" family of machines (which currently consists of the Mac II and IIx) and a "compact" family of machines (now headed by the Mac SE/30 and including the Mac Plus and SE). The Mac SE/30 bridges the gap between the high performance of the modular systems and the small footprint of the compact family. As you'll see, the Mac SE/30 is in fact just as powerful as any of the modular systems.

From the outside, the Mac SE/30 looks like any other Mac Plus or SE. But the internals of the machine represent a new design incorporating the performance features of the Mac IIx. The new logic board (see photo 1) sports a 16-MHz Motorola 68030 processor with a built-in memory management unit (MMU), just like the Mac IIx. The logic board also has the same single in-line memory module (SIMM) RAM chips as the Mac IIx, and it is expandable to 8 megabytes.

The Mac SE/30 also uses the same 256K-byte SIMM-mounted ROMs. This means that the Mac SE/30 can support Toolbox functions such as Color Quick-Draw, although it comes standard with the familiar 9-inch built-in black-and-white monitor. A new 32-bit expansion slot, however, provides an obvious opportunity to add color to the Mac SE/30.

A 16-MHz 68882 FPU comes standard, and so does the Apple Sound Chip

continued

Photo 1: *The new Mac SE/30 logic board features a 16-MHz 68030 CPU and a 68882 FPU. There's only one floppy disk drive connector on the board (the Mac SE has two). Note the SIMM-mounted ROMs and the 030 Direct Slot. Also note the 64K bytes of video RAM just above the SIMM RAM.*

supporting four-voice stereo sound. The Mac SE/30 also features the new SWIM (Super Wozniak Integrated Machine) floppy disk drive controller chip, as well as the FDHD (which stands for "floppy disk high-density") floppy disk drive that reads MS-DOS- or OS/2-formatted disks as well as Apple II ProDOS disks. However, you still have to use the Apple File Exchange utility to transfer files from a foreign operating-system format to the Macintosh system.

And there's more. The Mac SE/30 features a single 32-bit expansion slot called the 030 Direct Slot. It is basically NuBus-compatible but has a different form factor so that NuBus cards won't fit. But Apple engineers told me that you can easily convert Mac II NuBus logic designs to the 030 Direct Slot and that you can use the same NuBus software drivers.

However, the 030 Direct Slot's 120-pin Euro-DIN connector is not physically compatible with either the 96-pin Mac II NuBus or the 96-pin Mac SE expansion slots. The Direct Slot is positioned vertically in the computer chassis, unlike the horizontal layout of the expansion slot in the standard Mac SE (see photo 2).

One other significant change from the standard Mac SE is the use of 64K bytes of separate dual-ported video RAM to control the internal monitor. Since this video RAM is connected directly to the CPU, there's no need for a video buffer in the memory subsystem. This means that main memory is not burdened with the additional task of controlling video I/O, unlike the Mac SE's memory subsystem. Specifically, in the original Mac SE design, the CPU had to interleave its memory accesses (three out of four cycles) with the video display circuits so that the Mac's screen could be drawn. This meant that the Mac SE's CPU could access memory only 75 percent of the time, which degraded system performance. The separate video RAM in the Mac SE/30 allows the 68030 processor to access memory at every cycle. This is a big win in performance, since applications execute in fewer cycles than with the previous video buffer system.

The basic interfaces in the Mac SE/30 remain unchanged from those of the Mac SE. The SE/30 has two Apple Desktop Bus connectors, two RS-232C/RS-422 serial ports, and a SCSI connector rated at the same transfer rate as the one in the standard Mac SE: 172K bytes per second for polled transfers, and 656K bytes per second for blind transfers.

Along with the Mac SE/30, there is a
continued

Photo 2a: *The cramped insides of the Mac SE/30.*

Photo 2b: *The expansion board mounts vertically in the chassis. In the original Mac SE, the expansion board mounted horizontally under the main logic board.*

Table 1: *According to the version 1.2 BYTE Benchmarks, the new Mac SE/30 outperforms the Mac SE and the Mac II, and it approaches the performance of (and in some cases, even outperforms) the Mac IIx, with which it shares a common coprocessor. All times are in seconds.*

Test	Mac SE/30	Mac SE	Mac II	Mac IIx
CPU				
Matrix	16.4	69.2	21.2	17.1
Sieve	31.7	170.2	40.2	31.3
Sort	29.5	154.1	44.2	29.5
String move				
Byte-wide	82.1	373.6	93.9	82.1
Word-wide	42.1	186.8	45.6	42.1
Doubleword	22.9	121.3	22.9	22.8
Disk I/O				
SubFinder Seek				
SCSI				
(1 block)	16.6	28.3	16	13.9
(32 blocks)	154.3	185.9	35.7	35.6
File I/O				
Seek	0.2	0.6	0.2	0.1
Read (seconds/K byte)	0.02	0.049	0.02	0.021
Write (seconds/K byte)	0.01	0.044	0.01	0.014
Large file				
Write	4.3	11.8	5.2	4.3
Read	4.8	8.4	4.8	4.7
Video				
Text				
Text edit	5.7	16.3	5.6	4.7
Draw string	2.4	3.9	1.8	1.6
Graphics				
Small-C	44.6	80.9	57.7	52.8
QuickDraw	0.3	1.2	0.3	0.3
Floating Point				
Math	147.6	891.3	175.3	151.5
Trapezoidal rule: sine(x)	73.3	598.0	84.8	72.7
Trapezoidal rule: e^x	97.9	720.7	112.5	96.6

new version of the System software—version 6.0.3. This new version includes some fixes to the FDHD driver and to the Apple File Exchange, allowing a wider tolerance for MS-DOS-formatted disks. For current Mac users, unless you use the Apple File Exchange, there is no reason to upgrade to version 6.0.3.

Pricing and Configurations
The Mac SE/30 is available in two configurations: either with 2 megabytes of RAM and a 40-megabyte hard disk drive, or with 4 megabytes of RAM and an 80-megabyte hard disk drive. Both configurations come with the 68882 FPU. Both hard disk drives have an average access time of less than 30 milliseconds, according to Apple engineers.

The Mac SE/30 with 2 megabytes of RAM and the 40-megabyte hard disk drive will cost $5069, according to Apple. As we go to press, Apple says the

4-megabyte version with the 80-megabyte hard disk drive will cost $6369. As with most Apple pricing structures, neither of these configurations includes a keyboard.

Upgrades
Apple plans to offer upgrades for current Mac SE owners. You will be able to upgrade either to a new logic board with 1 megabyte of RAM or just to the FDHD disk drive and SWIM chip disk drive controller. Note that you will not be able to use the 150-nanosecond memory from your old Mac SE in the new Mac SE/30, since the SE/30 requires the faster 120-ns RAM used on the Mac II.

You can bet that the upgrade price for the logic board will make you think twice. It may actually be cheaper to sell your standard Mac SE and then buy a new Mac SE/30 than it would be to upgrade. At the time of this writing, Apple

had not set price or availability for upgrades, though the company says it expects to have the upgrades available in March.

Performance
Although I had limited time to work with the new machine, I was able to run the BYTE benchmarks on a Mac SE/30 at Apple's headquarters. The machine I tested had a 40-megabyte hard disk drive and 8 megabytes of RAM. Table 1 shows the BYTE benchmark (version 1.2) results of the SE/30 compared with those for a standard Mac SE, a Mac II, and a Mac IIx. As you can see, the performance of the SE/30 is comparable to a Mac IIx and in some cases exceeds it. Clearly, this machine is in a whole different league than the standard Mac SE.

In particular, note that processor-intensive tests, such as the Sieve, String Move, and floating-point tests, show the SE/30 sometimes outperforming the standard SE by almost an order of magnitude. To be fair, I should note that the standard SE does not have an FPU. Disk read/write operations are also faster, since the Mac SE/30 uses higher-speed hard disk drives.

So Why Buy a Mac II?
This is the obvious question. The Mac SE/30 offers essentially the same features as a Mac IIx. The only real difference is that the Mac II and IIx have six NuBus expansion slots. You can expect to see lots of Mac IIs operating as network file servers with Mac SEs and SE/30s as nodes on the network.

The other difference is that Apple does not plan to support A/UX on the Mac SE/30. So if you want to run Unix, you'll have to buy a Mac IIx, unless you're willing to try running an unsupported version of A/UX on the Mac SE/30. Since the SE/30 uses the same CPU, MMU, and ROM as the Mac IIx, there is no reason why A/UX would not run on it. It's just that Apple won't support it.

The Mac SE/30 is a very attractive machine. It's expensive, but it brings truly high-performance capabilities to the "compact" Macintosh product line. It essentially puts the power of a Mac IIx within the small footprint of the Mac Plus/Mac SE family. In fact, it's so attractive, I wonder how close Apple came to giving it a different name—one more in keeping with the Mac IIx: the Mac SEx. ∎

Nick Baran is a BYTE senior technical editor based in San Francisco. He can be reached on BIX as "nickbaran."

Introducing the fastest possible way to create the fastest MS-DOS programs possible.

Everything about Microsoft® C Optimizing Compiler version 5.1 is dedicated to the professional programmer.

Fast code. Fast development. Fast debugging. And full support for both MS-DOS® and the OS/2 systems in a single package.

There's no faster C code on a PC, because powerful optimizations, such as in-line code generation and loop enregistering, generate executables that are compact and efficient. The documentation even teaches you special coding techniques to squeeze every last bit of speed out of your code.

Fast code isn't all you get. Under MS® OS/2, the 640K barrier is gone so you can write C programs as large as a gigabyte. You can call the operating system directly. Create more responsive programs (multiple threads allow program operations to overlap). And build Dynamic Link Libraries (DLLs) that can be shared, saving valuable memory. DLLs also allow your main programs to be smaller, so they load faster.

Introducing the fastest possible way to create the fastest MS OS/2 programs possible.

For Personal Computers Running the MS, OS/2 or MS-DOS, Operating System

Microsoft C 5.1
Optimizing Compiler

Microsoft C Optimizing Compiler 5.1 Techbox

Compiler
- Optimizations that generate the fastest code for DOS and OS/2 systems.
  - In-line code generation.
  - Loop optimizations.
  - Elimination of common subexpressions.
- Full OS/2-system support to break the 640K barrier. New.
  - Family API programs that run under DOS and the OS/2 systems. New.
  - Write multithreaded programs and Dynamic Link Libraries. New.
- Small, medium, compact, large, and huge memory models.
- Mix models with NEAR, FAR, and HUGE keywords.
- Fast compilation (10,000 lines/minute) with Microsoft QuickC.
- Fastest math, in-line 8087/80287 instructions, and floating-point calls.
- More complete support of proposed ANSI standard.
- Over 350 library functions, including a graphics library.

Microsoft CodeView
- Full OS/2 systems support. New.
  - Debug applications of up to 128 MB under the OS/2 systems. New.
  - Debug multithreaded programs and Dynamic Link Libraries. New.
- Source-level debugging for precise control over programs.
  - Dynamic breakpoints in the source.
  - Debug programs written in a variety of Microsoft languages. New.
  - Full symbolic display of C structures. New.
  - Interactively follow linked lists and nested structures. New.
  - Watch variables, memory, registers, and flags.

Other Utilities
- Fast linking (twice as fast as the C 4.0 version linker).
- OS/2 incremental linker – up to 20 times faster than a full link. New.
- OS/2- and MS-DOS reconfigurable programmer's editor. New.

Special offer. Programmer's editor and Microsoft QuickC. compiler included. FREE!

Now includes OS/2 systems support!

Microsoft

You can even write a single Family API program that runs under both MS-DOS and MS OS/2.

Microsoft Editor is the first reconfigurable text editor for programmers that lets you develop under MS-DOS and MS OS/2. Under MS OS/2, multitasking lets you edit one file while you compile another, which cuts development time. You can even generate multiple compiles that report errors directly back into your source code.

Microsoft CodeView® is the highly acclaimed window-oriented source-level debugger that makes debugging fast and efficient. You can view program execution while you watch variables and register values change. And under MS OS/2 you can debug multithreaded applications, DLLs, and programs as large as 128 MB.

New Microsoft C Optimizing Compiler 5.1 for the professional programmer. It's all the speed you need. Call (800) 541-1261, Department I53.

Microsoft

READY LINE OVERLOAD

Stalking new hardware and software, Jerry prowls the halls of COMDEX

I am just back from COMDEX. Naturally, I didn't want to turn in my column before I went, meaning that it's very late now and has to be on the wire by dawn, meaning that I'm not going to have much time for testing either hardware or software. It's all right, though. While I normally talk about little that I haven't actually got up and running at Chaos Manor, I use different rules for show reports.

Also, I got a fair amount of stuff tested before I went to Las Vegas. There's no lack of stuff to write about: my "ready line" is overloaded.

Gray Scales

It used to be that lots of major players used COMDEX as the stage for announcing spectacular new products. After a while, there were so many that I never had time to see them all. All we journalist types had the same problem, and when we'd get together in the pressroom we'd try to compare notes, but it was futile. Something important would be overlooked. Maybe a lot of somethings. A number of companies noticed this and decided that COMDEX was a lousy time to announce anything really new. New product announcements nearly vanished.

However, a few outfits have cottoned on to the idea that no one announces at COMDEX anymore, so although there weren't many new product announcements this year, there were a few.

The most exciting new hardware I saw this year was Intel's Visual Edge printing-enhancement system. This is a pair of boards—one for your IBM PC AT compatible (or 80386), the other for your Hewlett-Packard LaserJet II—plus cable and software. Put it all together, and you can do halftone printing from PageMaker, Snapshot, Ventura Publisher, Publisher's Paintbrush, and a bunch of other desktop publishing programs.

The result as demonstrated at COMDEX is pretty spectacular. Visual Edge gives you 64 levels of gray at 70-line-per-inch resolution, and the hardware printed a large, complicated picture in about 3 minutes. Intel's press kit includes a bunch of pictures reproduced on a good-quality (19 gray levels) copier, and the same pictures scanned with a 256-gray-level Microtek scanner, then printed with a LaserJet II with Visual Edge. Believe me, you won't have any trouble figuring which is which. The press kit says Visual Edge gives a 300 percent improvement, and it looks about like that to me.

You can also get 37 levels of gray at 100 lpi, when resolution is more important than tone and shading.

You'll need expanded memory to use Visual Edge: 1 megabyte for ⅛ page, and 4 megabytes for a full page. Of course, you need expanded memory to get much good out of any desktop publishing program. Obviously, Intel would like to sell you a genuine Above Board to supply it, and certainly that's as good a way to go as any. The good news is that since Visual Edge uses your computer's memory to stoke graphics into the LaserJet II, you can use the printer's memory to hold downloaded fonts. Visual Edge also speeds printing up by anywhere from 200 to 400 percent.

All told, Visual Edge adds a whole new dimension to desktop publishing. I don't recommend products I've seen only at shows, but if you're really interested in desktop publishing under DOS, you really ought to have a look at this.

Logic Gem

Most of the big companies come to COMDEX. At great expense, they fill the main convention hall; booths have been getting larger and larger, probably because COMDEX lets *big* booth buyers have first choice on location. (I wandered into a room where they were coordinating booth allocations for next year, and I discovered that it's really a very complicated affair requiring quite a few people and a lot of communications, much like an auction that goes on over 3 days. Somebody ought to do a story on it.)

Anyway, you get your turn choosing COMDEX space both by booth size and by the number of years you've been coming (and if you ever drop out, you have to start over). This means that the big, long-established companies tend to fill the main exhibition hall. New companies are sent to outlying areas, like the Bally (formerly MGM) Grand.

This year, some new companies were put in a place called Cashman Field, which is somewhere near downtown Las Vegas. I met only one press person who ever *found* Cashman Field, much less went there. Certainly I didn't. I hear rumors that next year, simultaneously with COMDEX, they'll have MACDEX, a show devoted entirely to Macintosh computers and products, and that MACDEX will be at Cashman Field. I've always wondered if Apple truly wants to fence their machines and users off from the rest of the microcomputer community. Possibly this will do it for them. It might also put them down where they'll be visited only by the Little Sisters of the Poor. We'll see.

Anyway, since most of the start-up companies tend to be put into the Bally Grand, I generally find it worthwhile to spend a day there, and more often than not I find that the most interesting day of the show. This time was no exception.

The most exciting software at COMDEX was called Logic Gem, published by Sterling Castle Software (and yes, I know that the one in Scotland is Stirling Castle). They didn't have a copy for me to bring home (for the standard reason:
continued

"The documents are being printed"), so I haven't tested it here; but assuming that it performs as I saw it demonstrated, Logic Gem is going to change the way we write programs.

Logic Gem lets you make up a table of all sorts of conditions: data types (is this an integer?), keyboard input characters (is there input, and if so, is it a backspace character?), Booleans (is this condition true?), variables (is variable FOO greater than variable BAR?), or indeed almost anything you like. You can then add outcomes: beep, go get an input character, exit loop, look for a disk file, and stuff like that. These go into another table.

Once that's done, the program generates all the possible logical sequences of the conditions you added and generates a new table. You can then couple conditions with outcomes. For a simple example, if there is an input character, and it is not numeric and not a backspace, beep and go get another input character;

otherwise, exit from the loop.

When you've set up all the outcomes you want with the conditions that should make them happen, the program generates the rules. You can edit these rules until you have things the way you want. If your edit produces a logical inconsistency, the program warns you. When you're finished, Logic Gem will generate commented source code in C, structured BASIC, Pascal, dBASE, FORTRAN, or natural language (English). You can incorporate the source code into a larger program and compile it.

According to Sterling Castle's literature, the natural-language output "is ideal for program documentation." From the demonstrations I saw, I'd say it would be fine to put into an appendix, but I sure wouldn't want to spend a lot of time *reading* that. On the other hand, I prompted them to make up some tables, then examined the Pascal and structured BASIC code generated, and it looked fine to me. Even the comments made sense.

I've just finished doing some programming for Mrs. Pournelle's reading program, and I often got complicated IF THEN...ELSE IF constructs wrong. I expect everyone does. That isn't supposed to happen if you're using Logic Gem, since it generates code to do exactly what you want and nothing else.

If that weren't enough, Sterling Castle claims that Logic Gem can also generate code to "collect optimization statistics for frequency of logic use and user-defined cost functions. These results may then be fed back into the logic compiler to generate more efficient code."

I've always said that the future of programming lies in figuring out *what* you want the computer to do, not *how* to make that happen, because as machines get more powerful, we'll develop more and better tools that can get the machines to do anything we want them to do. If Logic Gem works as I saw it demonstrated—and I have no reason to suppose it won't—it will be a real step toward the future I predicted.

DESQview 386
Quarterdeck's DESQview, for those few who don't know, is a multitasking program that lets you run several of your present DOS programs at once. The way I use it, DESQview doesn't quite do that; except for communications, I don't run programs in background, because I don't do long compilations or spreadsheet recalculations. However, I do like to keep a whole flock of programs and utilities in

continued

Diamond Scan 14
Auto-Tracking
800 x 600 pixels (digital, analog, monochrome)
330 x 500 pixels (composite NTSC)

Diamond Scan 20A
Auto-Tracking
800 x 600 pixels (digital, analog, monochrome)
1024 x 800 pixels (maximum)

Diamond Scan 16L
Auto-Tracking
1024 x 768 pixels (typical)
1280 x 1024 pixels (maximum)

Diamond Scan 20L
Auto-Tracking
Up to 1280 x 1024 pixels

Actual unretouched screen images
Monitors shown with optional bases

A Clear View To Monitor Quality

CHAOS MANOR

memory, so that I can switch from one to another; it speeds things up something wonderful.

DESQview on a PCompatible is slow and might or might not be worthwhile, depending on your applications. On an 80286 machine, it's a rival to Microsoft Windows: it's easier to install, harder to learn, and considerably better with the general run of programs, although not as good as Windows running programs designed for Windows. DESQview is pretty good, a lot better than running an 80286 machine bare, and a way to make use of expanded memory; but it's not wonderful, either.

It's on an 80386 (and presumably the new 80386SX, which I haven't tried yet) that DESQview really speeds thing up.

DESQview 386 is really just an upgraded DESQview 2.2 plus a new release of Quarterdeck's QEMM-386 memory manager. (There's a 286 memory manager that speeds up DESQview for 80286 machines, but it's not the same.) I've been using a beta-test copy of this 80386-only package for a couple of months. I picked up the shipping copy at COMDEX, brought it home, installed it as an

update—if DESQview detects that there's an older version of itself on your hard disk, it doesn't change your configuration, macro, and information files—and began running. Incidentally, I like DESQview's installation a lot. There are all too many programs that don't pay any attention to whether you've already installed an earlier version.

Lately, my normal DESQview setup on the big Cheetah 386 has been GrandView in window 1, Q&A Write installed with Microlytics' Word Finder in window 2, and SideKick and Procomm Plus loaded in window 3. Norton Commander, DOS services, and other stuff go in later windows if I need them. When I exit Procomm Plus, the window doesn't close, and SideKick sits there ready to be accessed when I want it. I ran that way with the beta-test DESQview 2.2 for weeks.

The shipping copy crashed that system within 10 minutes. Worse: it crashed it in the worst possible way, by locking out the keyboard in Q&A Write while I was doing this column. Now, by "locking out" I mean locking *out*: nothing, including Ctrl-Alt-Del, worked. It was as

if the keyboard had been disconnected from the machine—so much so that I actually got down on the floor to check the cable connections. There was nothing for it but to hit the hardware reset button, thus losing all the text I had just entered.

Fortunately, my early training on microcomputers has stuck: I save text early and often, generally at the end of each paragraph. It's as much a nervous habit as anything else, but what with a fast 80386 machine and the Priam 330-megabyte hard disk drive, it doesn't take much time.

Anyway, I didn't lose much the first time DESQview locked up, and I was on my guard after that, which was just as well, because it did it three more times, at which point I decided that enough was enough. Before I installed the newest version of DESQview 2.2, I had, of course, saved the old copy off to the Maximum Storage APX-3200 WORM (write once, read many) drive; I figured it was time to restore it.

WORMs in Paradise
I had some minor glitches with that; I used every beta-test version of Maximum

Only one supplier of color monitors offers the widest selection of features and operating flexibilities in the market today.

That company is Mitsubishi Electronics.

Mitsubishi® delivers the reliability and performance that can meet your color information display requirements today as well as tomorrow. Larger screen sizes, truer colors, and optimum resolutions make your work easier—and far more productive.

Whether your requirements call for fixed-frequency graphics standards, like EGA and VGA, or multiple-frequency performance, Mitsubishi has the color monitor with the resolution and size to fit your specific needs. This includes the Diamond Scan Series of 14", 16" and 20" auto-tracking monitors, some with microprocessor-enhanced programmable display settings. All at very competitive prices.

To get a clear view of monitor quality and value, look to Mitsubishi.

For product information or nearest authorized Mitsubishi Electronics sales representatives, please call 1-800-556-1234, ext. 54M. In California, call 1-800-441-2345, ext. 54M. Mitsubishi Electronics America, Inc., Computer Peripherals Division, 991 Knox Street, Torrance, CA 90502, (213) 217-5732.

Mitsubishi Model	Screen Size (inches)	Horizontal Scan Frequency (kHz)	Mask Pitch (mm)	Compatibility/Resolution								
							VGA		Apple Mac II	1024 x 768 (48 kHz)	1280 x 1024 (64 kHz)
				NTSC	CGA	EGA	Std.	Ext.			
Diamond Scan 14 (AUM1381A)	14/13V	15.7 ~ 36 auto-tracking	0.31	•	•	•	•	•	•		
Diamond Scan 16L* (HL6605TK)	16/15V	30 ~ 64 auto-tracking	0.31				•	•	•	•	•
Diamond Scan 20A (HA3905ADK)	20/19V	15.7 ~ 36 auto-tracking	0.31		•	•	•	•	•		
Diamond Scan 20L* (HL6905TK)	20/19V	30 ~ 64 auto-tracking	0.31				•	•	•	•	•
XC1429C	14/13V	31.5	0.28				•				
XC1410C	14/13V	22 or 15.75	0.40		•	•					
XC1430C	14/13V	22 or 15.75	0.31		•	•					

*Microprocessor-enhanced programmable display settings

XC1429C
VGA Compatible
640 x 480 pixels

XC1410C/XC1430C
EGA Compatible
640 x 350 pixels

And Value.

MITSUBISHI ELECTRONICS

CHAOS MANOR

Storage's software, and apparently I have version clashes now. It's nothing I'm worried about, since I was in fact able to recover all my files; but that took longer than I liked.

The principal symptom was that DIR showed me several files with the same name (but with different dates). This shouldn't be. A WORM drive does preserve everything you put on it—if you save several files with the same name, it creates a new one each time—but DIR isn't supposed to be able to find any but the latest copies. Anyway, Norton Commander also saw multiple copies, but since it lets you select (with a scroll bar) precisely which file you want to copy, I didn't have any great difficulties extracting the one I wanted.

A discussion with Theresa Beyers, Maximum Storage's technical wizard, generated the version-clash hypothesis. I'll try a new storage cartridge and the latest software; if that doesn't do it, they'll look at the drive hardware. Fortunately, I have a second drive (which will eventually go to the Lowell Observatory), so this is no problem.

Since I have recovered all my files, I'm annoyed, not panicked; I still like WORM drives a lot better than I like tape backup. Even with the problems, I got the work done about as fast as tape or Fastback Plus would have done it.

Back to DESQview
Anyway, I got the old version of DESQview restored but left the new QEMM in place. That worked fine, so the night wasn't wasted. Next morning I called Gary Saxer at Quarterdeck.

After a long discussion, we concluded that the problem was an open Procomm Plus window containing SideKick.

SideKick was one of the earliest of the terminate-and-stay-resident (TSR) programs. It did things in a particularly sloppy way, and for reasons I don't quite understand, Borland never cleaned up SideKick's act. Both GrandView and Q&A Write know about DESQview. Word Finder doesn't, but it's a late model and a fairly well behaved TSR.

"SideKick, though," Gary said, "tries hard to intercept everything it can, and by definition a communications program runs in background. Meanwhile, we've been working real hard to make DESQ-view work with multiple TSRs. I think SideKick is running in background and sometimes it sees what it thinks is its hot key."

"But I need SideKick—"

"You can open it in its own window. We know how to handle it there. But I think having SideKick in a communications window without the communications program may give it a chance to do mischief."

Could be, thought I. The DESQview documents are complete, if very hard to find anything in. (No index. I hate that. But that's all right: I think my son Alex and I will do a book on DESQview. That will have an index.) Eventually, I figured out how to work this.

I invoke SideKick, then Procomm Plus, in a batch file. I want SideKick because it has neat features in its notebook. The F4 key will capture text off the screen into the notepad editor. You can edit that text, or write your own, then do Control-K-E, and SideKick will squirt out a marked block of text through the modem. I use the combination for BIX, and it's very handy.

continued

However, when I would quit Procomm Plus, that particular window stayed open, and I used it as the SideKick window. DESQview's manual tells me that if the last command in the batch file that loads SideKick, then Procomm Plus, is EXIT, the window will close when I shut down Procomm Plus. If I want SideKick when Procomm Plus isn't up, I have to put it in a separate window. This sometimes means that I have two copies of SideKick running, but that's no problem.

I've made the recommended changes: I put EXIT in the batch file to load Procomm Plus, then reinstalled the latest version of DESQview. I've also opened SideKick in a separate window.

It works fine. I've been pounding away on this column for a few hours, and no glitches whatever.

The Quarterdeck people are working to see if they can't fix the problem from inside; meanwhile, DESQview works fine, but you shouldn't leave windows with SideKick running in background even if you have enough memory. I can live with that limitation.

The diagnosis is confirmed. As a test, I opened the SideKick/Procomm Plus window and exited without closing the window; within 5 minutes, I had managed to hang up the Q&A Write window. After I reset, I put the EXIT command back in the Procomm Plus batch file, and the problem hasn't surfaced again.

At COMDEX I saw a great number of signs touting OS/2; every booth that had an OS/2 application running got a special sign from IBM. The reality, though, was that there were darned few real OS/2 applications on display, and none of them were very impressive. It's possible that OS/2 with Presentation Manager will be the wave of the future. Certainly it could be, given that IBM puts all that formidable marketing talent behind it. On the other hand, IBM has been wrong before. Remember TopView?

I still find DESQview with an 80386 the proper way to operate. Even if you have only an 80286, DESQview is a better bet than OS/2 just now. I don't expect that to change for at least a year. Probably longer.

Norton Again

I always enjoy seeing Peter Norton, but although he lives no more than 10 miles from Chaos Manor, we get to talk only at computer shows. Like me, he's got so much to do that he seldom has time for social activities. The result is that we mostly get together on business. COMDEX was no exception. Norton had a suite in the Sahara, where he was showing new products, and, alas, my visit there was the only time we got to talk. At least we made a tentative date for dinner Real Soon Now.

His most important new product is the Norton Utilities 4.5 (available in standard and advanced packages), which turns out to be a bigger improvement over version 4.0 than the version numbers indicate. The advanced package includes Norton Disk Doctor, a program they wrote to diagnose and fix logical disk problems. Unlike Steve Gibson's SpinRite, NDD doesn't correct format errors; but it will unscramble a munged file allocation table, fix a boot record, repair the media byte, and take care of a number of other DOS disk problems.

I'm already on record as saying that everyone needs the Norton Utilities. Just last night I managed to erase a file I wanted to keep, and Norton Utilities got it back for me with no muss or fuss; and I

use the disk sort, screen color, and other programs in the package nearly every day.

The new edition is a major upgrade, and worth getting. Recommended.

ConvertUnits

The whole appeal of the Macintosh is ease of learning and use. That's probably why it's so frustrating when things don't go well.

At COMDEX I was handed a copy of a program called ConvertUnits from GTA. This is said to be the most complete and accurate unit-conversion program for the Mac. It can be installed as a HyperCard stack or as a desk accessory (DA).

I've long had such a program for my IBM PC. It can go in as a TSR or be set up in its own DESQview window. I don't use it a lot, but when I do need it, I need it bad. I figured I ought to have something like that on the Mac II, and setting it up as a DA seemed like the right way to go about it.

The problem is that I don't do that very often, so I had to rely on the Convert-Units manual. That, alas, is incomplete. To make ConvertUnits work, you must,

in addition to using the DA Font/DA Mover to install the access software in the DA menu, copy the program's database files onto the hard disk; but the manual doesn't tell you that. The Mac's DA Font/DA Mover isn't anywhere near

A ll told, ConvertUnits is a scientifically complete and accurate unit-conversion program.

as intuitive as Apple thinks, so I wasn't sure what I was doing wrong. All I knew was that I had followed the directions that were given in the manual, and when I was done, I could pull down the DA menu and select ConvertUnits, but nothing

interesting would happen.

Eventually, I figured it out and copied the ConvertUnits file (which is misnamed in the documents) onto the boot disk. Then I discovered another odd quirk. ConvertUnits uses the Mac interface to let you select the kinds of units you're interested in and what you'll convert to and from. You then enter a number from the keyboard and press Return.

Now suppose you want to change one of the units but don't want to change the value (unity, for example, if you're interested only in how many ticks of a cesium atom there are in a day, a week, a galactic year, a tropical eon, etc.); you use the mouse to change the unit, but now you have to let go of the mouse and press Return, because there's no "activate" button in the ConvertUnits display.

There are other awkwardnesses in using the program. None of them are fatal. Just annoying.

Finally, you can't add new units to the database table. Of course, you might not want to: the program has a remarkably complete set of units, just about everything in the big *CRC Handbook* (including furlongs and fortnights).

All told, ConvertUnits is scientifically complete and accurate. I haven't seen a better unit-conversion program for the Mac, but I do wish they'd do an overhaul on the user interface.

Peabody

I never thought I'd want an on-line DOS help program. After all, I know the common commands, and if there's something I can't remember, such as the different switches (e.g., /s or /a) for XCOPY, I can always use Chris DeVoney's *Using PC-DOS* to look it up. So I believed until the other day, when I was experimenting with some stuff dangerous enough that I wanted to make frequent backups to my WORM drive and found that I was looking up the same things over and over again. Then I remembered Peabody.

I suppose the world is divided into two kinds of people, those who were Rocky and Bullwinkle fans, and those who weren't. Me, I never missed an episode, so I remember Mr. Peabody, the rather snobbish dog who had adopted a boy named Sherman.

Copia International has licensed the name and image of Mr. Peabody for their series of on-line help programs. The programs are pretty big to be memory-resident in a PC, but they know how to use extended memory, so if you have an AT or 80386 machine, you can get away with using as little as 25K bytes of main

World Radio History

ROM. These programs include an integrated spreadsheet and word processor; a "diary," which we would call an appointment calendar; an alarm clock; and a printer driver. Optional software includes a database manager, a communications package, and stuff like that, some available now, the rest coming "pretty soon."

The literature says it has 8 lines of 80 characters on the supertwist screen. While this is technically true, two of those lines are taken up by the ROM-based programs and can't be dispensed with; that leaves only 6 lines. If you go beyond 72 characters on a line, there's horizontal scrolling. Since horizontal scrolling when you're trying to write is about as useful as a chocolate covering for your wristwatch, in effect you have 6 lines of 72 characters.

The screen isn't particularly easy to see. Naturally, it's not backlit, and the letters are small. Still, you *can* see it. I used the Z88 to take notes at a couple of meetings, including the planning session for the next West Coast Computer Faire, and that turned out to be fairly easy. Of course, I was sitting at a table in a well-lit room. Now I have the Z88 on the mouse table next to me as I write this; my room is certainly well lit; but I'm having trouble seeing the Z88's screen. However, if I put it dead in front of me, there's no problem.

The Z88's keyboard is well laid out. It's covered with a rubber film to protect the system from anything wet; presumably, you can spill your Jolt Cola on this without harm. The unit has a substantial and solid feel; I like it. The keys are *very* sensitive, much more so than I'm used to, but again I can adjust. I didn't have a lot of problems using it for notes at COMDEX, and everyone around me appreciated how quiet the Z88 is compared to the Tandy Model 100 or NEC PC-8201. There's no key click at all.

The manual says the system can be attached to a disk drive, but there's no reference to any such thing in the sales literature. I make no doubt that if it becomes at all popular, Traveling Software will come up with a version of LapDOS to connect the Z88 to the small battery-powered Brother floppy disk drive.

Meanwhile, the system relies on little memory pack cartridges. These are quite small, smaller than a packet of cigarettes (and considerably thinner). They come in 32K- to 512K-byte sizes. The 32K-byte cartridge costs $45; the 512K-byte cartridge, which is what I suppose you'd actually need (since this is your only real storage device), is $440, no small sum.

Like everyone else, Sir Clive was caught in the memory price crunch; when this system was first designed, memory was nearly free, and the little memory cartridges were intended to be a lot cheaper.

There's also a cartridge for transferring data from the Z88 to a PCompatible (and another for the Mac). The PC Link (composed of a cartridge, a disk for the PC, and cable) costs $75. The Mac version is unaccountably $129; I'd presume that's merely a reflection of the fact that everything associated with the Macintosh costs more.

If you write any BASIC programs for the Cambridge Z88 or have really valuable data, you can save to an EPROM cartridge.

If you write any programs for the Z88 (it has built-in BASIC) or have really valuable data, you can save to an EPROM cartridge. These cost $45 for 32K bytes and $110 for the 128K-byte variety. EPROM is the only way you can permanently store things for the Z88; if I were traveling with this machine, I'd want the ability to EPROM anything important I wrote. I'll tell you why in a minute.

The point is that if you want a practical Z88, you'll pay for it: $599 for the machine itself; at least $110 for a 128K-byte RAM cartridge, and more likely $440 for 512K bytes; another $110 for an EPROM cartridge (in fact, I'd want to have two or three for a long trip, although I might not use any of them); and at least $75 for a PC Link package. This adds up to a minimum of $894, and it wouldn't be hard to have considerably more than that in it. But that's not a *lot* of money for a good portable, and certainly the Z88 is convenient. It's the lightest-weight and just plain handiest little notebook computer I have seen.

That's hardware.

The major software of the Z88 is an integrated spreadsheet and word processor called Pipedream. In theory, this ought

to be a great idea: you can put a spreadsheet into any document you're writing.

In practice, it's disappointingly hard to use. As an example, if I start a paragraph with a tab (which I am in the habit of doing), the result is a *permanent* shift of the left margin. I thought perhaps that somewhere buried in the documents (over 200 pages) I'd find a way to cause all the lines except the first to go to the left side of the screen, but I tried all I could find from the index, and all I managed to do was make a monumental mess of what I was trying to write; I ended up with a second document about six characters wide over to the left of my first one. As far as I can tell, the only real remedy is never to use a tab when you're writing. That's quite a sacrifice.

There are other glitches. The upshot is that although the documents make it look as if you could get along with a Z88 as your only computer, I sure wouldn't want to if I had to use it a lot. The Pipedream text editor just can't forget that it's *really* a spreadsheet, and thus it only tolerates being used to write English text.

The Z88 should fare better as a second machine used as a notebook for students and journalists. In previous columns, I've complained about the weight of my Zenith SupersPort 286 laptop: I love it when I get it to my hotel room, and indeed I did a lot of writing I wouldn't have done at COMDEX if I'd had anything less powerful with me. On the other hand, I sure wasn't about to carry the SupersPort around with me to meetings, or even from my hotel room to the BYTE booth. It's just too darned heavy.

Incidentally, I have managed to take some of the sting out of the SupersPort: I got a SkyValet garment bag, the kind that turns itself into a sort of luggage cart. The SupersPort sits nicely on the little pop-out plastic briefcase holders. It's really easy to get around airports with that combination, and it sure came in handy standing in the 45-minute taxi line at the Las Vegas airport.

Anyway, I was eager to get the Z88 as a second computer. It's light enough to fit in the bag with the SupersPort, then be carried around to meetings as a notebook; and I did that, using it to take a lot of COMDEX notes. I did have to be careful never to use a tab, but what the heck.

Unfortunately, all those notes have vanished.

The Z88 is supposed to save your stuff, even if you merely turn the machine off. The PC-8201 does that. So does the Model 100.

Alas, the Z88 doesn't really do that.
continued

You must explicitly save your Pipedream files if the machine is to be off for more than a few minutes. Saving isn't easy: the software is extremely awkward, and it requires paging through a number of menus. Although there's a keyboard command that will bring up the "Save File" menu, even after you learn it, you then have to type in the darned filename every time; it won't default to the last one you used, even though you have not changed documents. Because it's awkward, I didn't use it. My fault, I guess; but I didn't, so when I got home, every one of the Pipedream files I'd made at COMDEX had vanished. Worse than that, the machine had clearly managed to reset itself: even the date was wrong.

I don't know if that was the X-ray machine at the airport (neither the Super-sPort nor the PC-8201 seems to mind those security systems) or that unsaved files just go away after a while (which is what I suspect), but the fact is that all the work I'd done with the Z88 is gone, and that doesn't put me in a mood to be particularly friendly about the machine.

I still have it, and I'll take it with me on my next trip, which will be to Hawaii. It seems that Freeman Dyson was to be the banquet speaker at a conference on Grand Challenges to Computational Science, and he'll be unable to make it, so they've chosen me as his replacement. I've been walking on air since they told me; in my book, Mr. Dyson is one of the 10 most interesting people I ever met, and I find it enormously flattering to be thought in the same league with him. Anyway, the Z88 will go with me, and I'll be very careful not to use tabs, and to explicitly save all my files; and we'll see if my attitude changes. I sure do like the small size and quiet keyboard.

Ezekial

The Ezekial contest—what should I do with my original Z80 CompuPro—produced a large number of letters. I'll publish some of them, and I'll be sending congratulations to many other letter writers.

It turns out that the Smithsonian Institution does indeed want old Zeke as part of their history of computing display. They'll keep him running, and people can actually tweak him. Given Compu-Pro's domination of the S-100 market, they wanted a Godbout/CompuPro/Via-syn system anyway, and Zeke was about the best known of those.

Most of the other suggestions had to do with worthy causes. I'm sympathetic to that argument, but old Zeke is not only

pretty old, he wasn't all that standard as a CP/M system to begin with; he really was state of the art. It takes a fair amount of expertise to keep him running; and few (I think none) of the worthy causes, such as a school on a Native American reservation, would have that capability. Bill Godbout would try to support Zeke (he's still supporting ancient Z80 systems), but that wouldn't be simple if Zeke were being used by people who didn't understand him, and especially if he were at a remote and unsophisticated location.

It was no easy decision, but I'm going to let the Smithsonian have old Zeke. I think he'll be happy there. It's pretty dumb to be this sentimental about a machine, but I certainly am. Farewell, old friend. I'll come see you sometimes.

Winding Down
I'm out of space, and I haven't even got started good. At COMDEX both Atari and Commodore showed new stuff. The Amiga has a new Unix configuration that I can't wait to try out; one of the high points of COMDEX was watching Dr. Henri Rubin demonstrate upcoming Amiga technology. Formidable! Atari

Items Discussed

APX-3200 .... $2675
Maximum Storage
5025 Centennial Blvd.
Colorado Springs, CO 80919
(719) 531-6888
Inquiry 1024.

ConvertUnits
   for desk accessory or
      HyperCard $59.95
   for both $74.95
GTA
815 Princess Ave.
Vancouver, BC
Canada V6A 3E5
(604) 255-8077
Inquiry 1025.

DESQview 386 .... $189.90
Quarterdeck Office Systems
150 Pico Blvd.
Santa Monica, CA 90405
(213) 392-9851
Inquiry 1026.

Flicker Master $17.95
TSR Hutchinson Co.
110 West Arrowdale
Houston, TX 77037
(713) 448-6143
Inquiry 1027.

GRASP $149
Paul Mace Software, Inc.
400 Williamson Way
Ashland, OR 97520
(503) 488-2322
Inquiry 1028.

Instant Access $100
Norton Guides .... $69 each
Norton Utilities 4.5
   standard $100
   advanced $150

Peter Norton Computing, Inc.
2210 Wilshire Blvd., Suite 186
Santa Monica, CA 90403
(213) 391-2000
Inquiry 1029.

Logic Gem $198
Sterling Castle Software
702 Washington St., Suite 174
Marina Del Rey, CA 90292
(213) 306-3020
Inquiry 1030.

Peabody $100
Copia International, Ltd.
1964 Richton Dr.
Wheaton, IL 60187
(312) 665-9830
Inquiry 1031.

TianMa $695
Asiacom
2761 McColl Pl.
Victoria, BC
Canada V8N 5Y8
(604) 477-7829
Inquiry 1032.

Visual Edge $695
Intel Corp.
3065 Bowers Ave.
Santa Clara, CA 95051
(408) 765-8080
Inquiry 1033.

Z88 .... $599
Cambridge North America
615 Academy Dr.
Northbrook, IL 60062
(312) 564-5512
Inquiry 1034.

has a lot of nifty new stuff as well. Both computers can use transputers to speed up very large visual displays. I was enormously impressed with both machines; for reasons I haven't time to go into, I think the Amiga has improved even faster than the Atari, but both are definitely worth watching. Each has a shot at being a serious rival to the Sun workstation at half the cost.

On that score, I'm supposed to get a new Sun386i in the next few weeks. I told Sun to keep the total list cost of what they send to below $20,000. That's still steep for the average BYTE reader, but at least it isn't out of sight; you can pay a good fraction of that for an all-up Mac II or a PS/2 Model 80.

I ought to mention Flicker Master, which is a screen cover that you Velcro over your Amiga monitor; it reduces the flicker from interlaced mode something wonderful. We've had one on my Amiga for a couple of weeks, and I wouldn't be without it.

I've got a whole bunch of animated displays built up from GRASP, Paul Mace's screen utility program; indeed, GRASP seems to have spawned a whole new category of consultants, like, for example, Robert Hurt's Trebor Truh Productions (2284 Almaden Rd., San Jose, CA 95125, (408) 723-0931), which puts together high-quality presentations built up from GRASP. If you have good artists in-house, get GRASP for them; but if you don't have artists, you can still get professional-quality presentations with outfits like Trebor Truh. Then you can present them with Traveling Software's color gizmo that sits on top of your View-Graph projector and flashes up what's on your screen.

There's more, but I'm *really* out of space. The book of the month is by Charles Murray, *In Pursuit of Happiness* (Simon and Schuster, 1988); this may be the most important book published last year.

COMDEX was exhausting, but fun. I sure like these little machines. ∎

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."

1	Exxon	25	Standard Oil (Ohio)	49	Consolidated Foods
2	General Motors	26	AT&T Technologies	50	Lockheed
3	Mobil	27	Boeing	51	Georgia-Pacific
4	Ford Motor	28	Dow Chemical	52	Monsanto
5	IBM	29	Allied	53	W.R. Grace
6	Texaco	30	Eastman Kodak	54	Signal Companies
7	E.I. du Pont	31	Unocal	55	Anheuser-Busch
8	Standard Oil (Ind.)	32	Goodyear	56	Nabisco Brands
9	Standard Oil of Cal.	33	Dart & Kraft	57	Johnson & Johnson
10	General Electric	34	Westinghouse Elec.	58	Coastal
11	Gulf Oil	35	Philip Morris	59	Raytheon
12	Atlantic Richfield	36	Beatrice Foods	60	Honeywell
13	Shell Oil	37	Union Carbide	61	Charter
14	Occidental Petroleum	38	Xerox	62	General Mills
15	U.S. Steel	39	Amerada Hess	63	TRW
16	Phillips Petroleum	40	Union Pacific	64	Caterpillar Tractor
17	Sun	41	General Foods	65	Aluminum Co. of Amer.
18	United Technologies	42	McDonnell Douglas	66	Sperry
19	Tenneco	43	Rockwell Int.	67	Gulf & Western Ind.
20	ITT	44	PepsiCo	68	Continental Group
21	Chrysler	45	Ashland Oil	69	Bethlehem Steel
22	Procter & Gamble	46	General Dynamics	70	Weyerhaeuser
23	R.J. Reynolds Ind.	47	3M	71	Ralston Purina
24	Getty Oil	48	Coca-Cola	72	Colgate-Palmolive

NEW FRIENDS AND OLD

Integrated software, special-purpose pseudocomputers, and software that can guide your next career move

Integrated software has gotten a bad rap over the years, largely because some of the most prominent examples of the genre failed to live up to both the claims of the manufacturers and the expectations of the users. Anyone who has watched this scene for the past 5 or 6 years will remember Ovation, which was to be the ultimate integrated software package of which we'd all been dreaming. Unfortunately, it turned out to be an elaborate hoax; after a lot of hoopla and media attention, the product and the company behind it went up in a puff of smoke.

There was VisiON, the slowest windowing environment ever marketed. Context MBA, running under the p-System, flopped while trying unsuccessfully to become the major competition to Lotus 1-2-3. In more recent software history, there was the pitiful rise, rapid descent, and eventual mercy killing of Modern Jazz on the Macintosh.

The list goes on and on. Products that have achieved modest success have done so in the face of contempt from all the gurus, who have maintained that the individual modules of any integrated software package could never rival the power and functionality of stand-alone programs. That criticism is still leveled today, almost as a reflex, at any product that calls itself integrated.

Rubbish. The fact is that while other software has been improving steadily, integrated software has improved as well. While the experts weren't looking, a number of programs have established themselves quite solidly in the marketplace. Think of AppleWorks, the Smart series, Microsoft Works on both the IBM PC and the Mac, SideKick Plus, and Framework III. All are excellent, and all have developed loyal followings.

And don't forget that the definition of "integrated" has also changed significantly. Lotus 1-2-3 was labeled an integrated package when it was first released because it combined the capabilities of a spreadsheet, database, and graphics program. By that standard, almost all software available currently would fit under the integrated umbrella. What would you call a word processor that was a text editor, formatter, spelling checker, thesaurus, outliner, calculator, file conversion utility, mailing-list manager, filing system, and desktop publishing package rolled into one? A few years back, we would have said it was an integrated whizbang; today, it's just an average word processor. We've come to take integration for granted. Where do you draw the line?

There's something in me that still yearns for the perfect integrated software package, designed as such from the ground up, rather than by patching functionality into another paradigm. I want to have the computerized equivalent of a Swiss army knife. Operating systems and many stand-alone programs have become so cumbersome that I see integrated software as the best way to achieve consistency and simplicity in my work environment. I don't know if I'll ever find what I'm after; perhaps my needs are so personal that they can't be met by any commercial product. I choose to go on hoping.

Once More with Framework

I got started thinking about integrated software because I've been working with Framework III and a Canon Cat for the

past month, two radically different approaches to the question of integration. I have been a loyal booster of Framework ever since I saw it some months before its acquisition by Ashton-Tate. I liked it then and I like it now, and the $695 price tag still gets you one of the best bargains in the software business.

The Canon Cat is perfect for someone who needs industrial-strength editing and record keeping but doesn't require a full-blown computer system.

Within a character-oriented windowing environment that looks a lot like the Systems Application Architecture interface (though it's largely incompatible with IBM's guidelines), you get a word processor, spreadsheet, database, communications module, and an outliner that also serves as the organizing tool with which to link documents of different types. Framework III is a good choice for anyone with a limited budget who isn't afraid of real power.

However, perhaps the biggest change from Framework II to Framework III is a repositioning of the product in Ashton-Tate's marketing scheme, which should tell you a great deal about what I think of the upgrade. The company is now calling Framework "decision support software," whatever that means. I guess touting Framework's all-around power hasn't yielded large-enough sales to keep Ashton-Tate happy, so the company is now yammering about Framework's ability to get you lots of different types of information at once, so you'll have all the facts you need at your fingertips. In that sense, the product is no different from its predecessors. You do get a large number of nice new touches, but if you haven't been seduced by earlier versions, Framework III has nothing magical enough to make you change your mind.

Every module has been improved. The word processor now boasts a standard ruler line with tab stops and all the usual stuff, so you have to resort to FRED, Framework's arcane programming language, only in dire emergencies. The excellent spelling checker has gained a companion thesaurus. The spreadsheet features sectional recalculation, so you don't have to wait for every cell to be updated if you don't need to.

The database is vastly improved, to the point where it's quite usable; in earlier versions, there was no easy way to specify field formats, so a ZIP code or an address that started with a numeral entered without a leading space was interpreted as a formula. The communications module has added more protocols and function keys linked to specific entries in the dialing directory. An optional package gives you electronic mail on a network. Oh yes, you also get selectable color and official mouse support. And everything seems a tad faster than before.

Some of the old problems remain, and one new one has been added. Framework III occupies well over 2 megabytes of disk space, so it's difficult to run without a hard disk drive and thus unsuitable for many laptops. FRED is probably the most difficult macro language around. Figuring out how to link files for output is no mean feat without cracking the manual. You can't add to the dialing directory without running the installation program. You've heard all this before.

The new problem is one on which mine might be a minority viewpoint. It used to be that if Framework encountered a blank cell in the spreadsheet or field in the database, it would give an error message when calculating. I thought this was a pretty good safety check, but it's not the way Lotus 1-2-3 does it. Framework III now treats nulls as if they were zeros, in the Lotus fashion, and computes results even if you've left something out. Phooey.

On the whole, the enhancements to Framework strike me as improvements to an already solid product, but not enough to cause a Framework II user to spring for an upgrade. I was no more impressed with the program than I was the last time I looked at it. Is that damning with faint praise? I'm not sure.

In many respects, Framework is beginning to show its age. Features that caused my jaw to drop when I first saw them are now fairly common in today's software marketplace; Framework III is no longer a revolutionary piece of software. I think it's terrific, but it's lost that mind-boggling tingle. I honestly don't

Create a 3.5″ Micro Diskette with unlimited protection against a hostile world.

You're covered. BASF molds the rigid plastic jacket of its 3.5″ Micro Diskette to its own specifications, providing unsurpassed protection for the flexible disk within. Call 800-343-4600 for the name of your nearest supplier.

The Spirit of Innovation.

Circle 31 on Reader Service Card

BASF

Items Discussed

Canon Cat..........................$995
   includes Canon 180 printer
Canon U.S.A., Inc.
One Canon Plaza
Lake Success, NY 11042
(516) 488-6700
Inquiry 1021.

Framework III $695
Ashton-Tate
20101 Hamilton Ave.
Torrance, CA 90502
(213) 329-8000
Inquiry 1022.

The Perfect Career..........$49.95
Mindscape, Inc.
3444 Dundee Rd.
Northbrook, IL 60062
(312) 480-7667
Inquiry 1023.

know how much further the program can be taken.

On the other hand, to anyone willing to commit to it, Framework III offers a flexible application development environment, a beautifully consistent command set, and modules as good as many top stand-alone programs. The fact that I've begun to yawn a little shouldn't be construed as a discouraging sign.

Nice Kitty

With its software in ROM, the Canon Cat has the most consistent integrated interface I have ever used. The machine is not exactly a computer as such. It looks like a computer—with a black-on-white monitor, keyboard, and disk drive—but it's, well, different.

I'll spare you a rehash of the physical description, which I detailed last month.

When you turn on the Cat, the screen is set up for entering text. Just plug in a disk and start typing. Dynamic page breaks are automatic; one keystroke can force a break or begin a new document. Everything is stored as a continuous text scroll; there's no discrete operating system. You move by searching for text; two pink keys in front of the space bar initiate forward and backward searches, or leaps in Cat jargon.

The cursor moves after each keystroke, so you rarely have to type a full word to reach its next occurrence. The Leap keys also serve as left and right arrow keys. To highlight for copying, deletion, or printing, leap to the start of a block, leap to the end, then punch both Leap keys to highlight.

Commands are triggered by Control-key sequences, and the functions are printed on the front of the key caps. The "Control" key is labeled simply Use Front. If you need help, hold down the Use Front key, press a key called Explain, then any key you don't understand. Presto—you get one of 48 help screens.

The Calc key yields a result from any highlighted formula, which then appears with a dotted underline. Pressing Calc again expands to the original formula for editing. You construct pseudospreadsheets by tabbing columns of numbers and using the Calc command to enter named variables and cell references. You can even embed variables and references in text paragraphs, so you've got the equivalent of free-floating cells in flowing text.

Where the Cat falls down is in formatting for output. Page numbers are centered at the bottom of every page, and there's no facility for headers or footers. You can leave off the number by highlighting less than a full page, and you can create macros that insert headers or footers onto pages, but it's really kind of primitive. Jef Raskin, whose team at Information Appliance designed the Cat, points out that it doesn't substitute for a computer. The Cat isn't intended for multicolumn page layout, shop-floor data analysis, or multiuser access. It's fine for basic text, simple spreadsheets, communications, free-form databases, and the like.

Recommendations? The Cat is perfect for someone who needs industrial-strength editing and record keeping but doesn't require a full-blown computer system. If you're willing to invest a little time and you're capable of throwing out your notions of what constitutes "power computing," get a Cat for yourself. It's as close to perfect integration, on a small scale, as I've encountered to date.

When I Grow Up

Every month or two, I receive the latest batch of games from Mindscape, which usually get thrown into a box in the garage and forgotten. Because it arrived with two awful-looking shoot-'em-ups, I nearly missed The Perfect Career (Mindscape, $49.95) and heaved it into the box with the rest. That would have been a major mistake, as it is anything but a game. It is an MS-DOS program designed to aid in the process of career counseling. You use it to match your interests and qualifications with potential employment categories.

The Perfect Career offers two complete testing and evaluation units, one for high school students with no professional experience, and one for adults with more knowledge of real-world abilities. In either case, you rank an extensive list of job situations on a 1 to 3 scale for interest and/or skill. The program derives an assessment of how your expectations match up to your answers, and it suggests a list of possible careers. I took the adult test, and I was amazed at how accurately the program pegged me; all the jobs it generated were things I'd enjoy doing. As an example, the program was canny enough to list "college or university president" as a choice but not "college or university professor," which meshes squarely with my impatience for teaching.

I then tried it on someone else, with somewhat less satisfactory results; the program produced only a terse list of the most obvious possibilities. When I asked my guinea pig about the process, she admitted she had given top scores to those things she liked and knew she could do and bottom scores to everything else. The Perfect Career then told her, naturally, what she already knew. I can't blame the program; the tests need to be taken with a lot of thought and a definite seriousness of purpose.

The program was developed by James C. Gonyea, director of the New England Center for Career Development. He also wrote a brief but helpful manual that sketches a systematic approach to choosing a career that goes well beyond the simplistic nature of the testing software. The Perfect Career is a complete package, and it makes an intelligent aid for students looking for career directions and adults who are thinking about making a switch.

Don't expect The Perfect Career to tell you exactly what you should do with your life; only you can make that decision. But the program can give you some insight into your desires and aptitudes and help you narrow your focus. It's no magic bullet, but it's pretty slick. I recommend it. ∎

Ezra Shapiro is a consulting editor for BYTE. You can contact him on BIX as "ezra." Because of the volume of mail he receives, Ezra, regretfully, cannot respond to each inquiry.

Your questions and comments are welcome. Write to: Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.

Create a high capacity Diskette with no capacity for error.

Make no mistake. Every bit, every byte, on every track of each BASF 5.25" HD Diskette is certified 100% error-free and warranted for life. Call 800-343-4600 for the name of your nearest supplier.

The Spirit of Innovation.

Circle 32 on Reader Service Card

BASF

FLEXSCAN™ 9070S, PC Hi-Res That Looks Like a Million.

The FLEXSCAN 9070 Multiple Scan monitor is of course compatible with other multi-scans, but includes improvements that will give you the professional edge which is the mark of a good investment.

You can extend your multi-scan range from 20kHz to 50kHz in practical terms. This means that, at the 48-50 kHz range, you can make use of PC CAD/CAE capabilities at a resolution of up to 1024 dots × 768 lines.

The FLEXSCAN 9070 takes advantage of non-interlace high resolution signal as high as 1024 × 768 to provide you with a flicker free display at much brightness. You can also use the 9070 with IBM PS/2 or VGA compatible boards at a high resolution mode like 800×600 and 1024×768 (non-interlace).

The FLEXSCAN 9070 provides a 16-inch screen, large enough for CAD/CAE and 3-D projections, yet small enough to fit comfortably into your home work space.

1024 dots × 768 lines Graphics (Non-interlace) AutoCAD

Also, for your convenience, all controls and switches, including the alternate video input, are located within easy reach on the front panel. The FLEXSCAN 9070 is compatible with a wide range of IBM, Apple, and other products, allow you to use all of today's popular programs---at a resolution that looks like a million.

FLEXSCAN™
MODEL 9070S

- IBM VGA(PS/2), 8514/A, PGC, EGA compatible and CAD/CAE use.
- Apple Mac. II and SuperMac Spectrum compatible
- Max. 1280 dots × 800 lines high resolution
- 1024 dots × 768 lines display on Non-Interlace signal delivers flicker-free high-res graphics
- 20kHz to 50kHz horizontal scan automatic adjustment. 50Hz to 80Hz vertical scan automatic adjustment
- 16 inch, 0.31mm dot pitch and newly developed XF (Extended Field) Gun to obtain both brightness and sharp focus.
- Front mounted controls including the input signal select switch between 2 video input.
- Selecting white or Amber displays colored application in shades of gray or amber
- Tilt-Swivel stand standard

NANAO ®

NANAO USA CORPORATION
23510 TELO AVE.,SUITE 5 TORRANCE, CA 90505
PHONE (213) 325-5202 FAX (213) 530-1679

Circle 164 on Reader Service Card (DEALERS: 165)

GETTING INTO BIGGER LANS

When you have a big job, sometimes you really do need a bigger hammer

The solution to business connectivity is simplicity. By that, I mean that the method you choose should be as simple as possible given the needs of your business. In my last two columns, I discussed the possibility of sharing resources without using a local-area network at all, and of sharing in a small area using a relatively simple LAN such as 10Net. There comes a time, though, when solutions as simple as these won't do the job.

As your office or department grows, the problems associated with managing central data storage and access to resources reach the point where a traditional central server LAN becomes the obvious choice. Eventually, printer servers become overburdened and the cable lengths required are too long to work properly. Likewise, a peer-to-peer network without a file server becomes difficult to manage because individual computers are turned off or files are moved. This makes shared resources difficult or impossible to find or use.

At this stage, in spite of its added complexity and cost, the file server becomes the simpler solution. In short, it's easier to manage a single computer than it is to manage dozens. Of course, a LAN can have more than one server, but you will still use few servers relative to the number of workstations. Likewise, your network printers will be centrally managed, as will other network resources, such as communications gateways.

Covering the Floor

An example of a typical central server LAN is one intended to serve the needs of the employees on a single floor of a professional services company. Say that about 80 employees need to have access to printers, asynchronous communications, some common word processing files, and electronic mail. There is some need for database management using dBASE III Plus. Word processing is performed using WordPerfect 4.2.

Because the employees of this company have very flexible schedules, and because many of them travel frequently, a peer-to-peer network would be quite inconvenient if any of the shared resources resided on any of the individual computers, since there is a virtual certainty that many computers on that floor will be turned off at any given time. Finally, there is a strong move toward the Apple Macintosh in portions of the company, and these need to be connected as well.

What to Do?

What I've laid out here is a situation quite common to many companies with a large number of office workers. While they may be in groups that vary in size, and while the physical circumstances may vary, the needs will remain. Workers need to be connected so they can make the best use of their available resources.

Now it's time to make some decisions. First, let's consider what sort of LAN you need, then choose LAN software, and, finally, decide what to use for a file server. Before I get deeply involved, though, you need to realize that there's more to LAN selection than I can discuss here. Some of these issues will be covered in later columns.

What Sort of LAN?

There are three factors you should weigh to determine the type of LAN you'll use: hardware requirements, wiring requirements, and expense. You may find that your selection of LANs is restricted because you need Ethernet for your Macintosh or for the VAX in the basement.

Likewise, if your building is already wired for ARCnet, you'll probably use that. Finally, if you have a very limited budget, you may find that you can afford ARCnet but not Ethernet.

Since most buildings aren't wired for LANs ahead of time, you may also find yourself looking for a product that will work with your already-installed twisted-pair wiring. You could also find out that you're going to have to install wiring, no matter what system you pick. This factor takes on added importance, since you may find that the cost of buying and installing the wiring for a LAN exceeds the entire cost of the rest of the network.

Since most personal computer users don't care what LAN protocol they use, Ethernet, ARCnet, or Token Ring, the primary critical factors are the costs to install the wiring and to adapt the computer to work with a network. Another critical factor is a particular computer's ability to work at all with some LAN protocols. For example, if you're using the Macintosh, you can use Ethernet on some later machines such as the Mac II and the SE. Otherwise, you have to use the standard AppleTalk network and find a way to integrate it.

The LAN Software

The software is the part of a LAN that the user sees. Regardless of the hardware and protocol involved, the user knows that the LAN is running something like Novell NetWare. Thus, you need to make sure you select the LAN operating software with your users' needs, as well as their installed hardware base, in mind.

According to a BYTE survey at PC Expo last July, the most popular network software is Novell NetWare. That survey showed that significantly more than half of all users employ Novell. There are a variety of reasons for this, not the least of which is that Novell has written its network operating system to function with most of the popularly available networks. Also popular, but less common than NetWare, is 3Com's network with its 3+ LAN operating software.

Both of these LAN operating systems are widely used because they make minimal demands of the user. While both require logging onto the network and both may require passwords, to users who have obtained access to the network, the file server appears to be just another disk drive. Likewise, printing occurs just as it would with a stand-alone computer. Once users have learned how to work their computers, they need a minimum of training to use the network.

Besides supplying access to big disk drives, LAN software typically offers a variety of other functions. One that used to be standard is E-mail. Currently, LAN software companies, including Novell and 3Com, are selling their E-mail software separately. In the case of Novell, because its E-mail product needed improvement, many users bought third-party packages.

The Server

Most companies that sell network software also sell file servers, although Novell has announced that it is leaving this portion of the business and has sold its server line to Samsung of Korea. Basically, the file server is a proprietary version of a personal computer. Usually, a file server has some added features and enhancements that help match it to its dedicated role on the network. Good ex-

amples are the servers from 3Com that have an AppleTalk port. This feature allows Macs in a LocalTalk network to be added to a 3Com LAN.

Although most LAN companies sell file servers, you can use a personal computer as a file server if you prefer. Normally, the machine you would select as a file server would be an 80286- or 80386-based IBM PC AT clone with a large, fast hard disk drive and some extended memory. Novell and 3Com sell versions of their network operating software for these machines, and they work fine in most environments. In the example mentioned earlier, an 80286 AT clone with a 300-megabyte hard disk drive has proven entirely adequate to support the 80-person group. Had the group been using more database management or transferring large files, such as those generated by a CAD system, an additional or a more capable server would have been necessary.

Selecting a dedicated server depends partially on the load you plan to put on it and partially on other features that the server must have, such as 3Com's AppleTalk port. Likewise, you need to make sure that whatever you pick for a server will have room to support the required circuit cards, which will include one and possibly more network interface cards, additional memory, possibly a copy-protection card (although Novell has just dropped copy protection with version 2.12), and additional communications ports if the machine does not come with enough. Interestingly, the choices of server and network operating software do not have to agree. You can run most network operating software on an AT clone, and you can even run Novell NetWare on a 3Com server.

Choosing Sides
You may find that either network mentioned here will meet your needs and that the dealers in your area will sell them to you for about the same price. Now comes the problem of selecting. At this point, you should involve the dealer. After all, a network that won't work is pretty hard to ship back to the factory, and network manufacturers tend to encourage the involvement of their dealers. Novell, for example, requires you to pay for service through the factory, either through a charge card or a service agreement. You need to make sure that the dealer will return to service the network and solve installation problems, unless you have the ability to do it yourself.

There is, of course, the important question of compatibility. Since Novell's

and 3Com's networks are highly popular, these companies normally support them with applications software. Less-popular networks might not be supported, so you have to check. The ability to support IBM NetBIOS compatibility, as most network operating software does, helps in this case.

Does this look like a lot to go through? Well, it is. Choosing and implementing a large LAN is not a trivial process, and it's made more complicated by the fact that nearly all installations are unique in some way. ■

Wayne Rash Jr. is a consulting editor for BYTE and a member of the professional staff of American Management Systems, Inc. (Arlington, Virginia). He consults with the federal government on microcomputers and communications. You can reach 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.

Items Discussed

Advanced NetWare 2.12.....$2995
NetWare 2.12 ....$1495
NetWare Requester for
  OS/2 1.0........................ $200
Novell, Inc.
122 East 1700 S
Provo, UT 84601
(800) 453-1267
(801) 379-5900
Inquiry 981.

EtherLink network adapter..$495
EtherLink/MC
  (Micro Channel) $595
EtherLink/NB
  (Macintosh II) $595
EtherLink Plus network
  adapter..........................$895
3+Open LAN Manager
  Advanced System 1.0......$2995
3+Open LAN Manager
  Entry System 1.0.............$995
3Server 3S/400 ....$12,995
3Com Corp.
3165 Kifer Rd.
Santa Clara, CA 95052
(408) 562-6400
Inquiry 982.

We've Invented the Future of Instrumentation Software . . . Twice.

With Words   With Pictures

Acquisition

Integrated libraries for GPIB, RS-232, A/D-D/A-DIO plug-in cards, and modular instruments.

Intuitive character-based function panels that automatically generate source code.

Front panel user interface with virtual instrument block diagram programming.

Analysis

Extensive libraries for data reduction, digital signal processing, and statistical analysis.

Over 100 analysis functions plus all the built-in functions of your language.

Over 250 icons for computation and analysis.

Presentation

Flexible high-performance graphics and report generation.

Extensive graphics support for CGA, EGA, MCGA, VGA, and Hercules.

Macintosh Desktop Publishing compatibility.

NATIONAL INSTRUMENTS®

The Software is the Instrument ™

12109 Technology Boulevard
Austin, Texas 78727-6204
800/531-4742 512/250-9119

*LabWindows®-
for the DOS-based PC and PS/2,
with Microsoft QuickBASIC or C.*

*LabVIEW®-
for the Apple Macintosh.*

*Circle 167 on Reader Service Card for LabWindows.
168 for LabVIEW.*

HEY APPLE, I NEED A LAPTOP

The time has come for a portable Mac that doesn't strain your arms or your credit line

I can't stand it. I've really tried hard to like the laptops I use each week: a Tandy 200, a Zenith Z-183, and a Toshiba T1000. Despite the virtues of each machine, they just don't go the whole nine yards. They are either too heavy (Zenith), have a lousy display (Tandy), or a crammed keyboard (Toshiba). And they all suffer from that most serious of failings: They aren't Macintoshes! No version of MS-DOS or Tandy software can fill my Mac software void. I don't want to use these machines, but so far, I haven't found a viable alternative.

It's still a nuisance moving files from a laptop to the Mac and back (although Traveling Software's LapLink Mac is a big help). But I don't want to have to transfer files from MS-DOS to the Mac and worry about file format–translation problems in the process. What I really want is to take files from my laptop and use them directly on my home and office machines. Since I use 8-megabyte Mac IIs as my primary personal computers (my Sun-3 Ethernet workstation notwithstanding), my standard computing environment consists of MultiFinder running Word 3.02, FullWrite Professional 1.0, More 1.1c, Excel 1.5, VersaTerm-Pro 3.0, AppleLink 2.0, 4th Dimension 1.06, FoxBASE Plus/Mac 1.1, MacScheme + Toolsmith 1.5, HyperCard 1.2, Smalltalk-80 2.3, and MPW 2.0. I use some or all of this software each week (plus some other programs). Along with ALSoft's MasterJuggler (so I can keep a ton of desk accessories and fonts open simultaneously, too), it's my software of choice. Why should I have to give up all of this just so I can carry a computer with me?

Believe me, I've tried the available alternatives and have not been impressed. I used a Dynamac EL for a few days quite some time ago (see "Dynamac's Portable Mac" by Peter Wayner, May 1988 BYTE). Although it's a nice machine with a very nice screen, it's hardly a laptop. The thing weighed more than the Mac SE I've been lugging around on trips. The Dynamac folks need to trim the weight down to something closer to single digits for it to approach laptop status. And it doesn't have batteries, which can definitely cramp your computer-carrying style.

Apple has been rumored to be near announcing a laptop Mac (or Macs) for the last year now. I, for one, am getting sick and tired of waiting. Apple, please get on with it and announce some snazzy laptops soon. Naturally, the laptop Mac we all want must have lots of fast storage, a crystal-clear screen, a great keyboard, and an 8-hour battery life, weigh under 6 pounds, and cost under $2,000.

OK, so that's probably a pipedream. I'd settle for a laptop that weighed under 12 pounds, that was slim enough to carry in my briefcase or travel bag without destroying them, and cheap enough that my bank won't revoke my MasterCard when I try to buy it. The battery life is negotiable, I suppose, but it has to at least be a positive number.

It turns out that Colby Computers (who has been making laptop Macs for a couple of years now, by buying Mac Pluses and cutting them up) is about ready to ship a new laptop it calls the WalkMac SE. It will be based on the Mac SE motherboard, will weigh about 12 pounds, and has a backlit LCD screen and a rechargeable battery.

Unfortunately, it's also going to cost more than five grand ($5449). In the wake of Apple's September 1988 price increases, maybe that doesn't seem so *continued*

Items Discussed

AppleCD SC $1199
Apple Computer, Inc.
20525 Mariani Ave.
Cupertino, CA 95014
(408) 996-1010
Inquiry 1103.

Apple System 6.0.2
(Free from Apple dealers
if you bring your own disks)
Apple Computer, Inc.
20525 Mariani Ave.
Cupertino, CA 95014
(408) 996-1010
Inquiry 1102.

Colby WalkMac
SE .... $5449 and up
Colby Computers
4723 North Warren Ave.
Fresno, CA 93705
(209) 222-4985
Inquiry 1101.

Dynamac EL .... $5895 and up
Dynamac Computer Products, Inc.
14001 East Iliff Ave., Suite 410
Aurora, CO 80014
(303) 233-7626
Inquiry 1106.

Interactive Multimedia:
Visions of Multimedia
for Developers, Educators,
and Information Providers $24.95
Edited by Sueann Ambron and
Kristina Hooper. Redmond, WA:
Microsoft Press, 1988.
Inquiry 1104.

Wheels for the Mind
1-year subscription (4 issues) $12
Peter Olivieri, editor
Apple Computer, Inc.
P.O. Box 1834
Escondido, CA 92025
Inquiry 1105.

bad (OK, so it still feels bad). I haven't used one yet, but I'm scheduled to get a review unit soon. I plan to lug it with me everywhere and use it for writing, editing, some site consulting, and teaching and speaking engagements, then report back to you in a couple of months. I can hardly wait. Maybe I'll be able to break my Zenith/Tandy/Toshiba habit yet.

A Bug-Free System and CD-ROM

I've had the debugged and released Apple 6.0.2 System software for a bit more than a month now, and it's a *huge* improvement over the buggy and unreliable System 6.0 release. Hopefully, the experience with 6.0 and the bug-fixing provided by 6.0.2 will help Apple's engineers get a clean version of System 7.0 out when it ships.

While I'm on the subject of Apple, let me talk about its CD-ROM drive, the AppleCD SC. I've had one of these things sitting around my office for a few months now, but didn't really have the chance to do much with it except hook it up and make sure it works (it does). The drive is a decent piece of work that plays CD audio disks in addition to reading CD-ROM data disks. It's pricey at $1199, but it's often discounted. Since I'm teaching an introductory programming course (using HyperCard) for liberal arts students this quarter, it seemed logical to pull out the AppleCD SC and

fire it up with a copy of HyperCard 1.2 to try it as a read-only HyperCard storage device.

I've found that in and of itself, the AppleCD SC is not going to win any product-of-the-year awards for its technical execution. It's just a read-only device that happens to have a prodigious capacity (around 650 megabytes per disk). And it's generally slower than any Mac hard disk I've used. But technical prowess isn't what makes the AppleCD SC an important product. Nope, the AppleCD SC is important for what it can provide to applications like HyperCard: a nearly unlimited source of information that demands new methods of management and new paradigms for searching, sorting, categorizing, and displaying data.

This important point occurred to me as I was reading the excellent compilation of articles published in *Interactive Multimedia: Visions of Multimedia for Developers, Educators, and Information Providers*, edited by Sueann Ambron and Kristina Hooper and published by Microsoft Press. Many of those articles make it very clear that multimedia instruction, presentations, and data storage are no longer confined to research labs, but are being used and refined now in real environments, with Macs and CD-ROM drives being an important delivery platform.

Discover Parallel Processing!

Monoputer/2™
The World's Most Popular Transputer Development System

Since 1986, the MicroWay **Monoputer** has become the favorite transputer development system, with thousands in use worldwide. Monoputer/2 extends the original design from 2 to 16 megabytes and adds an enhanced DMA powered interface. The board can be used to develop code for transputer networks or can be linked with other Monoputers or Quadputers to build a transputer network. It can be powered by the 20 MHz T414 or T800 or the new 25 MHz T425 or T800.

Parallel Languages
Fortran and C Make Porting a Snap!

Microway stocks parallel languages from 3L, Logical Systems and Inmos. These include one Fortran, two Cs, Occam, Pascal, and our own Prolog. We also stock the NAG libraries for the T800 and Rockfield's structural and thermal finite element package. A single T800 node costs $2,000, yet has the power of a $10,000 386/1167 system. Isn't it time you considered porting your Fortran or C application to the transputer?

Quadputer™
Mainframe Power For Your PC!

MicroWay's **Quadputer** is the most versatile multiple transputer board on the market today. Each processor can have 1, 4 or 8 megabytes of local memory. In addition, two or more Quadputers can be linked together with ribbon cables to build large systems. One MicroWay customer reduced an 8 hour mainframe analysis to 15 minutes with five Quadputers, giving him realtime control of his business.

For further information, please call MicroWay's Technical Support staff at (508) 746-7341.

Micro Way

World Leader in PC Numerics

P.O. Box 79, Kingston, MA 02364 USA (508) 746-7341
32 High St., Kingston-Upon-Thames, U.K., 01-541-5466
USA FAX 617-934-2414   Australia 02-439-8400   Germany 069-75-1428

One article in that compilation, written by Apple's Mike Liebhold, points out the artificial limitations placed on the use of CD-ROM technology. He argues that CD-ROM software that allows for only straight information search and retrieval is missing the boat by a wide margin. He makes the logical and technical case for serious decision-aid software based on CD-ROM technology.

Liebhold contends that optical disks should be more than an archival media; they must be linked to powerful software. Of course, that's now happening with respect to hypertext with Apple's own HyperCard 1.2, which supports the AppleCD SC drive. With HyperCard as the hypertext engine, you can create virtually any view or structure you need and pull into these new objects the information (sound, graphics, and text) that is stored on the CD-ROM disk.

You can start to explore the concepts of CD-ROM-based hypertext applications now with the AppleCD SC drive and the Apple Educational CD Sampler Disk (which is available free from many Apple district sales offices). This disk contains a number of HyperCard stacks and other

hypermedia programs and gives some pretty strong hints at the processing paradigms that are growing out of hypermedia research.

Wheels for the Mind

The best way to stay current on developments in multimedia and hypermedia uses for the Mac is through the Apple University Consortium journal called *Wheels for the Mind*. Edited by Boston College's Peter Olivieri, it is published quarterly and costs only $12 for a 1-year subscription. In it, you'll find articles highlighting the Macintosh development projects planned or under way at major research schools and institutions. As an example, most of the topics covered in the summer 1988 issue had to do with HyperCard, in honor of its first birthday celebration.

The fall 1988 issue covered Mac applications used in instructional, research, professional, and administrative environments. Articles submitted for the publication often cover Macintosh applications and development efforts that you won't read about elsewhere.

In large part, some of the best software

for the Mac is in limited circulation within universities—software that you could take advantage of if you knew about it. *Wheels for the Mind* offers a good window into that software and its availability to non-university users. Much of the software (e.g., courseware, simulations systems, specialized research products, and extensions to popular commercial programs) that you'll read about in *Wheels* can be purchased for less than $50 each from the Kinkos Academic Courseware Exchange Catalog. Stop in any Kinkos copy shop in your area and pick up a copy of the catalog (and get your name on the mailing list for future copies). I highly recommend it as a way to extend your use of the Mac. ∎

Don Crabb is the director of laboratories and a senior lecturer for the computer science department at the University of Chicago. He is also a consulting editor for BYTE. He can be reached on BIX as "decrabb."

Your questions and comments are welcome. Write to: Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.

NO OTHER DESKTOP PUBLISHING SYSTEM OFFERS THIS FEATURE.

With the new Mannesmann Tally® Universal™ Publishing System, you can practically fly.

Thanks to a Raster Image Processor board that plugs directly into your PC or compatible, you'll process your pages at a speed limited only by the speed of your computer. Not—as is typical—at the speed of the printer. And you'll transfer ready-to-print data directly to the printer through a video interface at an incredible 3-million bits per second.

So when you're using the PostScript® compatible interpreter, you'll produce a printed page almost twice as fast as most other systems. But that's just ground speed.

If you use Aldus® Page-Maker or Ventura Publisher,® you'll

- Systems:
  1. Universal Publishing System (includes a PostScript compatible interpreter and Document Description Language (DDL))
  2. DDL Publishing System (DDL only).
- Resolution: 300 x 300 dpi.
- Emulations: Both systems include HP® LaserJet.
- Memory: 2 Mg.
- Typefaces: UPS includes 35 typefaces, DDL System includes 22 typefaces.
- Speed: 10 pages-per-minute.
- Dual paper cassettes standard, 250 sheets each.
- Dual output bins standard, 250 sheets each.
- Manual feed handles single sheets, envelopes, transparencies, and labels.
- Workload: 10,000 pages per-month.

really take off. Because when you select DDL instead of the PostScript compatible interpreter, you'll double that speed again. And with full page bitmap graphics, you can get printed output up to 17 times as fast.

So call the number below for the name of your nearest dealer and log in your time on the New Mannesmann Tally Universal Publishing System. A pilot's license is not required.

MANNESMANN
TALLY
1-800-843-1347
Ext. 191

In Washington state, call:
206-251-5524 Ext. 191

World Radio History

ELECTING THE PM

Here's what you
need to assemble
an OS/2 workstation
that runs Presentation
Manager without
breaking the bank

I have just spent days trying to assemble an OS/2 workstation to add to my local-area network. In this column and next month, I'll focus on the hardware you'll need to work with OS/2 version 1.1, the version with Presentation Manager. Version 1.1 is much pickier about the hardware that it runs on than version 1.0 was. This subject should be useful to those of you who already have an existing machine and are wondering whether or not OS/2 will run on it, and to those who are thinking about putting together a PM-capable machine.

It's not always a great idea to build a minimum acceptable system that runs OS/2, so I won't necessarily recommend the cheapest video board or disk controller. On the other hand, we can't all afford a Belchfire 950 33-MHz 80386. So I'll try to assemble the cheapest *reasonably powered* PM-based workstation. It won't be Mark's Dream Machine—just a machine that will get the job done.

Hardware
OS/2 requires the following:

- either an 80286- or 80386-based motherboard
- an OS/2-friendly BIOS
- an IBM PC AT–type hard disk controller
- a 1.2- or 1.44-megabyte floppy disk drive A that works without a device driver
- a 60-megabyte hard disk drive that works without a device driver,

with 32-megabyte or smaller partitions
- EGA or VGA graphics
- at least 2.6 megabytes of memory
- a mouse or other pointing device
- 16450-based serial ports

Wait! I see some of you shaking your heads, muttering, "Big, ugly, and slow...." Don't leave yet! Consider that OS/2 is considerably more powerful than DOS, and it consequently needs a more powerful platform. I remember when people told me that Lotus 1-2-3 would never sell because it needed 512K bytes of memory to be really useful.

As anyone who reads this column knows, I'm no apologist for IBM or Microsoft. But OS/2's hardware requirements aren't necessarily a bad thing. OS/2 changes the meaning of "minimum configuration." Programs will have to be written to a new lowest common denominator, one that requires high-quality graphics.

Programs that use graphics effectively are now scarce in the IBM PC world, because developers have had to contend with the many monochrome boards and the incompatible Hercules/EGA/CGA "standards." An OS/2 application can assume that EGA graphics, *minimum*, will exist on any machine it runs on. So we'll see more and better graphical programs under OS/2 than we've seen under DOS.

An OS/2-Compatible Motherboard
OS/2 requires, as we all know by now, an 80286 or an 80386. Which to buy? I've discussed the overweening merits of the 80386 in the past. But we're trying to trim costs here, so I'll talk about an inexpensive 80286-based workstation.

One final word in favor of the 80386, though: Most 80386s nowadays (at least until the EISA [Extended Industry Standard Architecture] bus becomes a standard) differ from 80286s only in the
continued

motherboard and kind of memory—the rest of the add-in boards are identical. You'll buy the same 16-bit Ethernet card, VGA card, and so on. The difference between an 80286 and an 80386, pricewise, is just an initial $700 to $1000. I know, $1000 isn't peanuts, but this workstation is going to run about $5000 anyway, so another $1000 for 80386 hardware and obsolescence-proofing wouldn't be a bad investment.

A basic 80286 AT-like motherboard—probably 10 MHz, no wait states, the low-end product these days—will work fine, as long as it has the right BIOS. The 80286 BIOS chips I've worked with are from IBM, Compaq, Phoenix, Award, DTK, and AMI. I'll eliminate those from IBM and Compaq for obvious cost reasons.

Of the remaining, I'd recommend either Phoenix or Award. I've had trouble running even OS/2 version 1.0 on some machines with AMI or DTK BIOSes. As late as early October—the last time I checked—one clone maker, Everex, told me that it couldn't run any version of OS/2 due to its AMI BIOS. The problem is due to be fixed—Everex will offer its own OS/2 soon—but the problem remains for garden-variety AMI BIOSes and most OS/2 implementations. Perhaps it will be fixed by the time you read this.

I've tried OS/2 on some DTK motherboards with DTK BIOSes without any luck. If you remove the DTK BIOS, however, and install a Phoenix BIOS, OS/2 will boot on many machines. Of course, if you have a machine that is compatible at the DOS level rather than at the BIOS or hardware level, you've got a fairly slim chance of getting anyone's OS/2 to run except that manufacturer's—if the company chooses to offer one.

One motherboard feature that would be particularly useful for running OS/2 would be a large memory capacity. We're starting to see motherboards that accommodate 4 or 8 megabytes of RAM. Be very careful here, however—memory is so expensive nowadays that the type of memory used in a computer is a major factor in determining the computer's overall cost. For example, looking at the prices this week (the end of October) for one large supplier, 1 megabyte of 100-nanosecond RAM would cost $347 if purchased as 1-megabit dynamic RAMs, $425 if purchased as a single in-line memory module, or $495 if purchased as four banks of 256K-bit DRAMs.

Whatever memory system your computer uses, you need a lot of it. To simply boot the PM with the compatibility box

requires 2.6 megabytes, and you're best with a minimum of 4 megabytes. After all, why go to all the trouble to run OS/2 and end up with a few K bytes of free space? This way, you'll have a bit over 1 megabyte to work with.

What did I end up with? A 10-MHz, no-wait-state, no-name AT motherboard with 512K bytes on-board (expandable to 1 megabyte) and equipped with a Phoenix BIOS version 3.10. It came in a box with a power supply, an OMTI full-track-buffered AT controller, and a keyboard for $720. Then I added an Everex RAM 3000 memory-expansion board and 3 megabytes in 256K-bit chips (I know I said that the 256K-bit DRAMs are the most expensive, but I already had the chips around). The board cost $170, and the chips ran $1500. Total so far: $2390.

Hard Disk and Controller

OS/2 requires a hard disk drive. You can *boot* OS/2 from a floppy disk, but you can't fit all the basic OS/2 files on a single 1.2-megabyte floppy disk. So, even if you boot from a floppy disk, you have to operate from a hard disk. The PM, the OS/2 files, a few basic utilities, and an editor together take up 6.9 megabytes on my disk. That's not counting things like the C compiler; it's just the kind of things that a typical user will have as the basic OS/2 PM package. Hence, my 60-megabyte recommendation.

Believe me, my 60-megabyte Priam hard disk is bursting at the seams. As OS/2 is disk-intensive, you'd do well to acquire a fast hard disk drive. One suggestion is the Seagate ST4096, an 80-megabyte drive with an access time in the area of 30 milliseconds. Discounters are offering the ST4096 at this writing for just under $600. Total cost so far: $2990.

You can use just about any AT-type hard disk controller. The only kind of controller to avoid is an XT type. I know you wouldn't deliberately buy an XT controller for an AT, but if you have upgraded your XT to an AT with a "baby AT" motherboard, you may still be using your old XT controller. While outfitting a new OS/2 workstation, the basic AT-type Western Digital WD1003 controller should do just fine.

When formatting your 60-megabyte drive, you may be tempted to run Ontrack Computer Systems' Disk Manager, PC-DOS 4.0, Storage Dimensions' SpeedStor, or some other device driver that allows a logical drive to exceed 32 megabytes in size. *Don't do it.* OS/2 just plain doesn't know how to deal with logical drives larger than 32 megabytes. Just

run FDISK from OS/2 or DOS 3.3 to partition the disk to a logical C drive of 32 megabytes and a logical D drive of 28 megabytes.

Next stop: serial ports and video. You'd never believe it, but saving money on video was the hardest part. Can our hero build a PM-capable OS/2 workstation for less than $5000? Tune in next month.

IBM OS/2 1.1 News

Just a day after the deadline for this column, I got a copy of IBM's OS/2 version 1.1. There's a lot of good news, and some bad news.

IBM's PM seems less machine-picky. I got it to run fine on some no-name 80286 and 80386 clones, whereas IBM's version 1.0 wouldn't run on the vast majority of the clones. And it *seems* fast—no benchmarks yet, though.

You can now restrict the amount of disk space that can be used for disk swapping. Previously, the swapper could eat up all your free space. Now you can say "leave me *x* megabytes free."

There will, indeed, be a CGA video driver for the PM. It will ship at the end of February.

It's *big*. IBM has squeezed the files onto five 1.44-megabyte floppy disks. A new UNPACK command unsqueezes them to about 14 megabytes—my earlier reference to 6.9 megabytes concerned the Microsoft Software Developer's Toolkit.

I haven't had enough time yet to find out whether or not this is some kind of copy-protection scheme, but FORMAT /S *does not work.* You get a message to the effect that "the /S option is not supported in this version of OS/2." I hope I've just overlooked something simple.

OS/2 Tip of the Month

Last month, I complained that there are no inexpensive OS/2 API (Application Program Interface) references. Two days after the galleys disappeared into the production process, I found *OS/2 API: The Pocket Reference* by Kris Jamsa (Osborne/McGraw-Hill, 1988). That's all it is—no OS/2 tips or tricks, just the unadorned API. But for $5.95, what a deal! OS/2 programming just got cheaper. ∎

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.

What you've been doing.

World Radio History

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.

ROSE ELECTRONICS

Give a Rose to your computer

Call now for literature or more information.
(800) 333-9343

THE ABCS OF X-, Y-, AND ZMODEM

XMODEM, the trail-blazer among file-transfer protocols, has spawned a host of offspring

The dictionary definition of protocol is "proper and correct conduct." We encounter several instances of protocol every day, and we take most of them for granted. From navigating 2-plus tons of metal from a crowded freeway to an off-ramp or dealing with that surprise visit from the in-laws, some sort of protocol is called into play. Often, protocols are learned through years of experience, but sometimes they are simply a matter of law or social norm.

In the telecommunications arena, until the last few years, file-transfer protocol wasn't an issue. You had exactly two choices: straight ASCII transfers or XMODEM. At that time, there wasn't any confusion, but our limited choices played havoc with the dilemma of how to move large amounts of data over the phone lines.

Today, more than a dozen types of file-transfer protocols are wandering around the communications cosmos. The differences in efficiency among these protocols, and why each was originally developed, are often a matter of personal preference as much as they are a quest for a better mousetrap.

Ours is an age of specialization. Bearing that in mind, you shouldn't be surprised that file-transfer protocols are also specialized. No single protocol is the optimum choice in all circumstances.

XMODEM—Good Enough?
Ward Christensen wrote the original binary file-transfer program, which he called MODEM. Keith Petersen adapted the program and called it XMODEM. The XMODEM protocol is a de facto file-transfer standard. But although it has achieved the status of a standard, it isn't internationally recognized as such. In an industry replete with standards, the fact that there's no official standard for file transfers seems incredible, but the telecommunications world seems filled with situations like this.

Because XMODEM is easy to implement, almost every communications package has its own version. Although these XMODEM implementations vary in how efficient they are, all but the most poorly implemented are compatible. It's rare to find two versions of XMODEM that can't complete a file transfer.

Simply speaking, XMODEM is a half-duplex protocol that transfers files in blocks of 128 bytes. *Half-duplex* means that only one computer can be "talking" at a time. Either the sender or the receiver is sending information to the other; it's akin to a one-lane road that must handle two-way traffic. A *block* is merely a sequence of bytes grouped together and sent across the phone line as a unit. These blocks are sent in sequence, so a 1K-byte file requires the transfer of eight 128-byte blocks.

In the XMODEM protocol, the remote computer checks the integrity of each block of data. If the integrity of the block is intact, the remote sends the ACK (acknowledgment) signal to the local machine, which then sends another block. If the integrity check fails, the local computer receives a NAK (negative acknowledgment) and must send the block again.

The original XMODEM has several problems. The short block length, 128 bytes, causes throughput to suffer when used in conjunction with time-sharing devices, packet-switched networks, satellite circuits, and buffered (error-correcting) modems.

Also, XMODEM uses a simple one-

continued

character checksum for detecting errors. In this scheme, the protocol calculates a checksum (using the ASCII values of the characters in the block) and appends a byte, representing the value of the checksum, to the end of the block. The receiving system calculates the checksum of the block it received and compares that to the value of the checksum at the end of the block. If the values are the same, the block is considered intact and the next block is transferred. Noisy lines, however, can easily confuse and corrupt this checksum scheme.

To overcome the checksum scheme's susceptibility to noisy lines, a beefier error-checking scheme was developed— a 16-bit cyclic redundancy check. The addition of a 16-bit CRC using a two-character CRC-16, instead of the one-character arithmetic checksum used by the original XMODEM protocol, is known as the XMODEM/CRC protocol.

Another problem: The process of sending and receiving each block, and the ensuing error-checking done on each block sent, take a certain amount of time. This checking process is called *overhead*. In the original XMODEM protocol, the ACK/NAK signals were sent after every 128-byte block. Now, though, the XMODEM protocol has been enhanced to allow the transfer of files in 1K-byte blocks. This improvement,

D*espite its shortcomings, XMODEM continues to be widely used.*

called XMODEM-1K, means that fewer individual blocks need to be transferred and less overall time is needed for the ACK/NAK signals. This change has resulted in higher throughput and less time on-line.

XMODEM has other problems. You can transfer only one file at a time, the file transmitted can accumulate up to 127 extraneous bytes, and the modification date of the file is lost when it's transferred from one system to the other.

Despite its shortcomings, XMODEM continues to be widely used and widely accepted, and virtually everyone supports it for all kinds of communications. Indeed, there isn't a communications package around that doesn't claim some sort of XMODEM compatibility.

Christensen readily admits that his XMODEM protocol is "not robust" and that the only reason XMODEM is the accepted standard is because "it was released in August of 1977 and immediately dumped into the public domain." At the time, anything put into the public domain was seized on by hackers eager for any new challenges. And everyone thought he or she could do it better. Thus, XMODEM has had several evolutions, each independent of the other. The various adaptations of Christensen's original file-transfer protocol have led to a virtual Tower of protocol Babel.

YMODEM—Better Than X?
After XMODEM came YMODEM. This protocol addresses many of the shortfalls

How the competition stands

Introducing the modem with a sleek new stand-up* design. Telebit's new T1000 Multi-Speed modem. The modem that not only looks different, but is different. With more features. More performance. And a surprisingly low price.

More modem for less money.
What makes the T1000 so different? For one, you get a choice of more speeds. The T1000 can send and receive data at 300, 1200, 2400, or 9600 bps using ordinary dial-up phone lines.

So the T1000 can talk to your installed base of low-speed modems, plus the large installed base of Telebit" and other PEP™ high-speed modems.

But it costs about the same as a 2400 bps error-free modem.

Another difference? The T1000 runs at 9600 bps with any type of data—without compression. Error free. With MNP and PEP error detection and correction.

The T1000 fits right in.
The T1000 Multi-Speed modem even talks Hayes—right from the box. And if you're using the AT command set or even Smartcom III software, we're compatible. You won't need new software. And you won't need new commands.

The T1000 also has internal support for the most widely-used communications protocols —Kermit, Xmodem, Ymodem and UNIX's UUCP. So you can transfer files up to 3 times faster than any other modem.

of its forerunner. For starters, it transfers files in 1K-byte blocks and supports multiple file transfers, otherwise known as batch-file transfers. Veterans of the telecommunications world might wonder why we need a new batch-file protocol when the older CP/M-based MODEM7 protocol supported batch files. The answer is that MODEM7 didn't support full path names, file length, file date, or other attribute information to be transmitted.

YMODEM, like XMODEM, is a half-duplex protocol. To further overcome the limitations of the half-duplex modem, YMODEM-g was developed. The "g" option of YMODEM is a modification of the YMODEM in which ACKs for data blocks aren't used. The data is merely sent all at once. The protocol doesn't use the ACK/NAK turnaround that XMODEM uses.

The receiver initiates the g option. When the sending computer recognizes this option, it knows to bypass the usual wait for an ACK to each transmitted block, sending all blocks in sequence and at full speed. The protocol, however, is subject to XON/XOFF flow controls (stop and go signals imposed by packet-switched networks).

YMODEM-g is intended to take advantage of high-speed, error-correcting modems. That's because error correction is taken care of at another level—between the hardware of the two systems. Theoretically, then, the software doesn't have to worry about things like ACK signals. The bottom line with YMODEM-g is that it doesn't support error recovery. If a NAK is received, the file transfer aborts. For this reason, you should use YMODEM-g only in a hard-wired environment or with a session-level protocol that takes care of error correction.

Z MODEM

is an attempt to correct the defects in X- and YMODEM.

___

ZMODEM—Best of the Bunch?
The author of YMODEM and ZMODEM is Chuck Forsberg. ZMODEM is an attempt to correct the defects in the previous versions of X- and YMODEM. The development of the ZMODEM protocol was funded by Telenet in an effort to improve a file-transfer protocol used with packet-switched networks. ZMODEM is an end-to-end protocol that uses a technique called *streaming*.

With streaming, the sender doesn't expect to get any ACK signals back from the receiver until the transfer is complete. If an error occurs, the sender will receive a NAK, and it's up to the sender to ensure that it can recover from any NAK received. This technique is advantageous when you're using a packet-switched network where the session-level protocols that are exerted by the network add more delays for file-transfer turnaround.

ZMODEM is extremely robust because of the error-correction scheme de-

up to Telebit's newest modem.

And here's another big difference. Since the T1000 runs the most popular communications software at the highest possible speeds. You can take full advantage of packages like HyperACCESS, Crosstalk-Fast, MicroPhone II, and Acknowledge.

Just plug us in, and the T1000 will fit right in. No matter what your environment.

Remote management for ease of use.

Here's the final difference. The T1000 offers a host of remote management features. Including remote access, remote configuration and remote diagnostics.

So get the modem that's head and shoulders above the rest. In features and performance. At just the right price. Telebit's new T1000 Multi-Speed Modem.

Call 1-800-TELEBIT or (415) 969-3800, today.

Or write Telebit at 1345 Shorebird Way, Mountain View, CA 94043-1329. Fax: (415) 969-8888.

Because no one gets the message through like Telebit.

T1000.   TELEBIT

© 1988 Telebit Corporation. Telebit is a registered trademark and PEP is a trademark of Telebit Corporation. Other brands or product names are trademarks of their respective holders.
*optional

Circle 228 on Reader Service Card

veloped by Forsberg; all ZMODEM transactions are protected with 16- or 32-bit CRC. According to Forsberg, when it is properly used, the 32-bit CRC reduces undetected errors by at least 5 orders of magnitude. ZMODEM also has a special security-challenge mechanism that guards against "Trojan Horse" messages written to mimic legitimate commands or file downloads.

Other advanced features include the ability of the sending or receiving computer to trigger an automatic download or command sequence; automatic step-down to YMODEM if the other end does not support ZMODEM (this ability, of course, assumes that the transmission medium accommodates both XMODEM and YMODEM); ease of implementation via a shell to an external program (DSZ.EXE); and file transfers that begin immediately, regardless of which program is started first, without the 10-second delay associated with XMODEM file transfers.

On the Horizon

Is there any rule of thumb you can follow concerning what protocol will best fill your needs? Sure. Try out several of them, and use what works best for you. In today's environment, that usually means XMODEM. After all is said and done, it's really the basic "standard." And assuming your communications package supports such protocols, you can improve your file transfers by using the more advanced X-/YMODEM-1K protocols.

We now have a wide variety of file-transfer protocols to choose from. And until the industry comes up with an official standard, you may find that the intelligent way to go is to use a combination of several different protocols. How do you know which protocol is the best to use in any given situation?

Next month, I will address this issue and explore the effect of high-speed modems, with their built-in error-correction schemes, on these various protocols. I'll compare each of these protocols when they are used over normal voice-grade telephone lines and with packet-switched networks. The results are mixed and, in some cases, surprising. ∎

Brock N. Meeks is a San Francisco–based freelance writer who specializes in high technology. You can reach him on BIX as "brock."

Your questions and comments are welcome. Write to: Editor, BYTE, One Phoenix Mill Lane, Peterborough, NH 03458.

World Radio History

World Radio History

All prices and specifications are subject to change without notice. Please call for current details. Call for complete warranty details.

World Radio History

INCREDIBLE PRICES.

'386/SS

COMPLETE 16MHz '386 DESKTOP WITH 32MB HARD DRIVE.
ONLY $2495.00
(20MHz add $500)

- Genuine 32-bit Intel 80386, 16 or 20MHz parts.
- *1MB* of RAM Expandable to 16MB.
- 32MB, High Speed Hard Drive, 1.2MB floppy drive.
- *Ultra high speed* Hard/Floppy controller. 1:1 interleave, 800 KB/sec transfer rate!
- High Res 12" Amber Display with Tilt and Swivel Base. Compatible graphics controller.
- 101 Key ZEOS Tactile Click keyboard.
- Serial and Parallel/Printer Ports.
- Clock/Calendar with Battery Backup.
- 2-32, 4-16 and 2-8 bit slots.
- 80287 and 80387 support.

COMPLETE '386 VERTICAL SYSTEM. 64MB DRIVE!
ONLY $2995.00
(20MHz add $500)

Acknowledged worldwide as the highest performance value anywhere. Same great standard ZEOS features found on our other complete systems plus:

- 64KB Zero-Wait CACHE using 64K of SRAM.
- 1MB of 32-bit RAM system expandable up to 16MB.
- *65MB* High Speed Seagate Hard Drive.
- 80287 and 80387 optional.
- Heavy Duty Vertical Case.

Call about our new *25MHz* systems. Now Shipping!

'386/V

machine and order it now with confidence. Dazzling performance and incredible prices await you. *Guaranteed.* Order now by calling 800-423-5891.

Other ZEOS Options Include:
- Basic Configurations: Call for prices on systems without drives, etc.
- EGA/VGA Upgrades. 14" EGA color monitor with EGA card, add $495. 14" VGA system, add only $695
- High Capacity Drives. Many other drives and options are available. Call Toll Free for details, 800-423-5891

Smoothing Out C

Optimizing C compilers
combine speed, high-
level convenience and
low-level power

Steve Apiki
and Jon Udell

C compilers have long been
available for DOS, but sophis-
ticated optimization and inte-
grated environments are re-
cent developments. Optimization
techniques, especially, have brought
new power to C. Handcrafted assembly
language code still outdoes even the best
machine-generated code, but optimizers
are rapidly improving, and as a result
programmers can rely more on C and
less on assembly language.

This month, we focus on a narrow sub-
set of C compilers for DOS—those with
highly developed optimization capabili-
ties. We looked at six packages: Turbo,
Aztec, MetaWare, Microsoft, Watcom,
and Zortech. Table 1 lists basic features.

All six compilers are more or less
compatible with the developing ANSI
standard (see table 2), and all include li-
brary support for standard functions. We
evaluated each compiler with a bench-
mark suite consisting of tests developed
by BYTE's editors and by C experts.

C Background

C is a high-level language originally de-
fined by Brian Kernighan and Dennis
Ritchie in the early 1970s. Its cult of pop-
ularity began during C's early associa-
tion with Unix. But C went on to earn a
reputation based on two essential quali-
ties: portability (there's a C compiler for

nearly every computer) and versatility.

Because C's fundamental data types
map closely to underlying machine
types, C has been called the universal as-
sembly language. That makes it a good
language in which to implement, for ex-
ample, a code generator. Yet because the
rules by which you combine fundamen-
tal C data types into more complex types
are regular and systematic, C also works
well at high levels of abstraction. So it's
an equally good language in which to im-
plement a database package or a window-
ing system.

The proliferation of C, which Ker-
nighan and Ritchie only informally
specified, inevitably began to cause
problems. Compiler writers diverged in
their interpretations of the language, and
they created incompatible dialects. More
recently, there's been a movement to
standardize the language around a set of
guidelines being developed by an ANSI
committee (see the text box "An ANSI
Conformance Sampler" by Thomas
Plum on page 176).

All the compilers we tested conformed
well to the unofficial standard. We ran
thousands of lines of code through each
compiler with virtually no problems.
Using compilers available just 3 or 4
years ago, however, our results probably
would have been different.

Optimization Techniques

Compiler optimizations fall into two
broad categories: source level and object
level. Source-level optimizations operate
independently of the target processor;
object-level optimizations seek to exploit
the characteristics of the target machine.

Hoisting of loop-invariant code is one
classic source-level optimization. Here,
the optimizer detects that an operation
within a loop—say, an assignment—does
not affect and is unaffected by the other
operations in the loop, and so hoists the
loop-invariant statement out of the loop.
The assignment happens once (not many
times), and the resulting code runs

faster. Other source-level optimizations
include elimination of dead code, elimi-
nation of unnecessary copies, constant
folding, and elimination of common sub-
expressions.

A section of code is dead if there's no
way for it to execute—a statement within
an if(0){...} block, for example. An
optimizer can simply discard such a
statement. Unnecessary copies occur
when a variable that's assigned a value
isn't used; again, an optimizer can dis-
card the futile assignment.

When the value of an arithmetic ex-
pression is constant and can be deter-
mined directly or indirectly by inspec-
tion of the source code, an optimizer can
fold the constant into the code in place of
the expression and thus shift the burden
of computation from run time to compile
time. When the same expression occurs
twice in a region of code, an optimizer
can replace that common subexpression
with a temporary variable.

Object-level optimizations complete
an optimizing compiler's repertoire. In-
telligent use of the available set of in-
structions and addressing modes repre-
sents one form of object-level optimiza-
tion.

The 80x86 architecture, for example,
offers particularly effective instructions
for moving and comparing bytes; com-
pilers targeted to that architecture should
use those instructions. Efficient use of
registers for the storage of variables is
one of the most powerful optimization
techniques. Good register allocation is a
science. Register allocators often use a
so-called graph-coloring algorithm to
map the variables active within a section
of code to available registers and to make
optimal selections based on the amount
of use each variable receives.

We tested the six C compilers for both
source optimization and object optimiza-
tion. But our tests also explored two
other areas: low-level performance and
application performance (XLisp). To de-
continued

Table 1: *A summary of each compiler's features reveals differences that can enhance performance (● = yes; ○ = no).*

	Borland Turbo C Professional 2.0	Manx Aztec C86 Commercial 4.1d	MetaWare High C 286 1.4	Microsoft C 5.1	Watcom C 6.5	Zortech C 1.07
Price	$250	$499	$595	$450	$495	$89.95
General						
Single compile/link command	●	●	○	●	●	●
Linker	●	●	○	●	●	●
Librarian	●	●	○	●	●	●
Stand-alone assembler	●	●	○	●	○	○
Make utility	●	●	○	●	●	●
Source code debugger	●	●	○	●	●	○
Editor	●	●	○	●	●	●
Integrated environment	●	○	○	●	●	●[1]
Memory models						
Tiny (64K total)	●	○	○	○[2]	●	●
Small (64K code, 64K data)	●	●	●	●	●	●
Medium (1M code, 64K data)	●	●	●	●	●	●
Compact (64K code, 1M data)	●	●	●	●	●	●
Large (1M code, 1M data)	●	●	●	●	●	●
Huge (large, but single data > 64K)[3]	●	○	●	●	●	○
Library						
DOS interrupts	●	●	●[4]	●	●	●
Dual executables (80x87 emulator)	●	●	●	●	●	●
Graphics library	●	●	○	●	●	●
Language extensions						
Pascal vs. C calling conventions	●	○	●	●	●	○
In-line assembler code	●	●	○	○	●	○
Assembler, linker						
Real-mode 80386 instructions	○	○	●	○	○	○
Interface to assembler code	●	●	●	●	●	●
Interface to other HLL	●	○	●	●	●	○
ROMable code	○	●	●	●	●	●
Other						
On-line help	●	○	○	●	●	●
Wild-card file specs to compile, link	●	○	○	●	●	●
Response file to compile, link	●	○	○	●	●	●
Install program	●	○	●	●	●	●
Documentation (pages)	2269	708	430	2518	1391	582
System requirements						
Operating system	DOS 2.0	DOS 2.0	DOS 2.0	MS OS/2 1.0 or DOS 2.1	DOS 2.0	DOS 2.0
RAM	448K	384K	384K	448K	512K	512K
Drives	Two floppy or hard disk	Two floppy or hard disk	Hard disk	Hard disk, 1.2M floppy[5]	Two floppy or hard disk	Two floppy or hard disk
Format	3½" or 5¼"	3½" or 5¼"	5¼"	3½" or 5¼"	3½" or 5¼"	3½" or 5¼"[6]

[1] Editor/compiler only
[2] Memory models can be customized.
[3] Static data can exceed 64K in any model (huge keyword).
[4] No bdos function, but interrupts handled in library.
[5] High density required for OS/2 functions only.
[6] $20 additional charge for 3½" version.

termine the effectiveness of the math libraries, we ran the suite twice: once on a Compaq 386 portable with an 80387 floating-point unit (FPU), and once on a Northgate 386 with no coprocessor. All tests were run using the small memory model. Starting with the low-level benchmarks, each test is detailed below, and the results are found in table 3.

The Low-Level Tests

For our low-level tests, we used six C functions devised by Thomas Plum. Each function contains an inner loop that executes exactly 1000 operations, and the test names reflect the types of operations they perform: integer arithmetic, floating-point arithmetic, function calling, and so on. These tests, which Plum Hall freely distributes, report the time—in microseconds—required for a single operation of a given type.

On our coprocessor-equipped Compaq 386, for example, the results show that an empty function call takes about 1.5 microseconds. That number was relatively invariant from compiler to compiler: The fastest was Watcom at 1.43 μs, the slowest was Aztec at 1.55 μs.

In general, the compilers bunched closely on the low-level tests, but the floating-point test yielded more diversity. Zortech, which took second on the emulator version of the test, finished last on the coprocessor version. The moral: Test both floating-point libraries that your compiler provides.

The XLisp Tests

For the XLisp tests, we used each compiler to build six working XLisp interpreters. Then we used each interpreter to run some benchmark programs. XLisp, written by David Betz, is a freely distributed program. Comprising 23 source files and 24,000 lines of C source code, XLisp is representative of the large, complex applications that serious C programmers construct to solve real problems.

Our low-level benchmarks rate the code generated by compilers according to specific measures of performance. They answer questions like: How efficient are function calls, integer multiplications, string comparisons, and pointer operations? How fast are the emulator and coprocessor library routines? How effective is the optimizer?

Our application in this case is the XLisp interpreter; its performance on a suite of Lisp benchmarks we wrote especially for this purpose constitutes our high-level compiler benchmark. The XLisp tests answer a more general—and probably more interesting—question:

When you write a program that uses all these features, how big will it be and how fast will it run?

Building XLisp

We knew that using six compilers to build six interpreters would be a big project. It was, but it turned out to be less difficult than we thought, for two reasons. First, XLisp is a cleanly written and beautifully modular program. Second, all the compilers support most of the important proposed ANSI standard features, so we had virtually no problems.

There was just a single glitch. All these compilers, except MetaWare's, support the library function bdos, which enables C programs to use DOS INT21 functions. The IBM PC version of XLisp uses bdos for a handful of primitive I/O routines. It's not part of the ANSI standard, but it's a convenience that many IBM PC compilers provide.

MetaWare does provide an alternative: You include an MS-DOS interface file and can then use a function called calldos in conjunction with a structure called Registers. It's slightly less convenient than bdos. You have to load Registers with appropriate values, call calldos, then retrieve the result from Registers. But after a bit of experimentation, we got it to work.

XLisp makes just three calls to bdos, and they're encapsulated in a single file. We added alternate versions of the three XLisp functions (embedded in a conditional #IFDEF METAWARE...#ENDIF block) and added -DEF METAWARE (to activate that block) to the command line we used when building XLisp with the MetaWare compiler. That solved the problem. It wasn't pretty, but given the scope of the project, we were pleasantly surprised to find that, for six compilers, one conditional block was the only modification that we had to make.

Lisp Benchmark Programs

We wrote six Lisp benchmark programs. The floating-point test executes 50,000 floating-point multiplications in a tight loop. The two sort tests share a common Lisp routine that implements an exchange sort; the routine operates on a list of objects and returns a sorted list. The integer sort operates on a randomly generated list of 500 integers, and the float sort operates on a randomly generated list of floats. The file I/O tests read and write a 32K-byte file.

We built the lists once, stored them in files, and used Lisp functions to read the numbers and construct the lists, thus ensuring that each interpreter would per-

form the same sequence of actions. We didn't time the helper functions. The numbers shown reflect only the time required for the sort.

Like the sort test, the sieve test executes a complex algorithm—in this case, the classic prime-number sifter. But its central data structure is an array—just as in the C version of the sieve—rather than a Lisp list. Lisp's hallmark, list manipulation, relies on techniques and data structures that aren't typical of many applications. To exercise a different (and possibly more representative) kind of functionality, we stored the flags that the sieve uses in a Lisp array—an object that's implemented in a relatively straightforward manner in terms of an ordinary C array.

Finally, the function-call test mimics its counterpart in the low-level suite. Here, we repeatedly executed an empty function call.

Results

The results show an interesting diversity. Watcom won the floating-point test. That makes sense, since it won the corresponding test in the low-level suite. Turbo, Aztec, and MetaWare did poorly on the floating-point test, and that too correlates with their performance on the low-level tests. Watcom took another first on the sort tests, followed by Microsoft. Zortech came in last on the sort; that's a bit surprising in view of its relatively strong overall performance on the low-level tests, and it proves that low-level tests taken alone can be misleading.

By way of redemption, though, Zortech dominated the file I/O tests, followed by Microsoft and Watcom, with MetaWare last. The winner of the Sieve test was Microsoft, with Turbo a close second. And Watcom did poorly on the Sieve—an unexpected result given its otherwise stellar performance. Watcom and Microsoft took first and second on the Function-Call test—results that again correlate with the low-level tests.

The combined results show Watcom and Microsoft as the favorites, as was true in the low-level tests. But Watcom, which trounced Microsoft on the low-level tests, won the XLisp tests by only a slim margin. That trend was apparent across the board—in contrast to the low-level tests, on the XLisp tests, times varied less from one compiler to another. The slowest interpreter, compiled by Aztec, trailed the fastest, compiled by Watcom, by 30 percent on the low-level index, but by only 20 percent on the XLisp

continued

Table 2: *The results of the Plum Hall ANSI validation suite. The tests represent requirements of the X3J11 standard.*

Draft reference	Description	Turbo C 2.0	Aztec C 4.1d	MetaWare C 1.4	Microsoft C 5.1	Watcom C 6.5	Zortech C 1.07
2.1.1.2	Backslash-splicing	O	●	●	O	●	O
2.1.1.2	Phases of translation	O	O	O	O	O	O
2.2.1.1	Trigraphs	O	O	O	O	●	O
2.2.4.2	`<float.h>`: DBL_DIG >= 10[1]	●	●	O	●	●	●
2.2.4.2	`<limits.h>`: has MB_LEN_MAX[2]	O	O	O	O	O	O
3.1.2	Internal identifier significance at least 31 characters	●	O	●	O	●	●
3.1.2.2	Scope rules: file, function, prototype, and block	●	●	●	●	●	●
3.1.2.3	Name space rules: variables, labels, tags, and members	●	●	●	●	●	●
3.1.2.3	Unique member name spaces	●	●	●	●	●	●
3.1.2.5	`long double` (even if same size as `double`)	●	●	●	●	●	O
3.1.2.5	All unsigned types	●	●	●	●	●	●
3.1.2.5	`signed char`	●	●	●	●	●	●
3.1.2.6	Type-compatibility rules[2]	●	●	O	●	●	O
3.1.3.2	Constants: U, L, unsignedness rules	●	●	●	●	●	●
3.1.3.4	Character constants with more than one `char`	O	O	O	O	●	●
3.1.3.4	`'\xFF'`	●	/	●	●	●	/
3.1.3.4	`'\a' '\v'`	●	/	●	●	●	●
3.1.3.4	Wide characters: `L'x'`, `L"x"`, `wc*`, `mb*`, `wchar_t`[2]	O	O	O	O	O	O
3.1.5	"Old-style" assignment operators are gone	●	●	●	●	●	●
3.2.1.1	"Value-preserving" integer conversion rules	●	●	●	/	●	/
3.2.1.5	Expressions with `float` operands have `float` type	/	●	●	●	●	/
3.2.2.1	Address-of on array and function	●	●	O	/	/	●
3.2.2.1	Call-through pointer `(*pkg.fn)()` may be written `pkg.fn()`	●	●	●	●	●	●
3.2.2.2	OK to cast `void` to `void`[3]	●	●	●	●	O	●
3.2.2.3	Generic pointers: `void *`	●	●	●	●	●	●
3.3.2.2	Prototype-with-default-sizes is compatible with no-prototype	●	●	O	●	●	O
3.3.2.2	Calling a prototyped function causes conversion (as if by ass't)	●	●	●	●	●	●
3.3.3.3	Unary plus	●	●	●	●	●	O
3.3.3.4	`sizeof` applies to any r-value expression	●	●	●	●	●	●
3.3.16.1	Structure assignment, return, and argument-passing	●	●	●	●	●	●
3.5.2.2	`enum` and tag-scope rules[1]	●	●	●	●	●	●
3.5.2.3	Tentative def for `struct sb b;`	O	O	O	O	●	O
3.5.3	`const` is independent qualifier of e.g. `struct` type	●	●	●	●	●	●
3.5.3	`volatile` preserves `auto` values modified after `setjmp`	●	●	●	●	●	/
3.5.4	Ellipsis	●	●	●	●	●	O
3.5.4	Prototypes (for declarations)	●	●	●	●	●	●
3.5.4	Prototypes (for "new-style" definitions)	●	●	●	●	●	●
3.5.7	Elided-braces rules	●	●	●	●	●	●

[1] Per 3/87 draft. [4] Per 4/88 draft.
[2] Per 12/87 draft. [5] Per 12/86 draft.
[3] Per 9/87 draft.

Key: ● — Supported.
/ — Compiled but produced incorrect output.
O — Failed to compile.

index. In general, the compilers divided themselves into three tiers: Watcom and Microsoft, then Turbo and Zortech, and finally Aztec and MetaWare.

Optimizing Tests
We tested the compilers' ability to perform the classic optimizations with a benchmark developed by Melvin Klerer and Hong Liu of Polytechnic University. The benchmark comprises two functionally equivalent C source files. The first, NONOP, contains constructs that a good optimizer should be able to improve. The second, OP, is a preoptimized version of NONOP. It's written in an optimal man-

ner so that an optimizer won't be able to squeeze much out of it.

NONOP is organized into sections, each of which contains a particular optimizable construct—common subexpressions, constants that are computable at compile time, loop-invariant code, and unnecessary copies. Each section iter-

Draft reference	Description	Turbo C 2.0	Aztec C 4.1d	MetaWare C 1.4	Microsoft C 5.1	Watcom C 6.5	Zortech C 1.07
3.5.7	Autoaggregate initializers	●	○	●	○	●	○
3.5.7	union initializers (via first member)	●	○	●	○	●	○
3.6.4.2	Long-size switch labels	●	●	●	●	●	○
3.7.2	Tentative definition for static	○	○	○	○	●	○
3.8.1	#ifdef, #elif	○	●	●	●	●	●
3.8.1	No syntax-checking of skipped groups	●	●	●	●	●	●
3.8.2	#include macro-name	●	●	○	●	○	●
3.8.3	"Hiding" of macro names	●	●	●	●	●	●
3.8.3	Benign redefinition allowed	●	●	●	●	●	●
3.8.3	Preprocessor catenation and string-izing	/	●	○	●	○	/
3.8.5	New preprocessor directive #pragma	●	●	○	●	●	●
3.8.8	Predefined macro names	○	/	○	○	●	○
4.1.3	<errno.h>[1]	●	●	○	○	●	●
4.1.5	<stddef.h>: offsetof on nested struct member[4]	○	○	○	○	●	○
4.2	<assert.h>	●	●	●	●	●	●
4.3	<ctype.h>	●	●	●	●	●	●
4.4	<locale.h>: initial locale is "C"	○	○	○	○	●	○
4.4	localeconv, LC_CURRENCY, negative_sign[4]	○	○	○	○	○	○
4.5	<math.h>	●	●	●	●	●	●
4.5.1	Math library sets errno when required	●	●	○	○	●	/
4.6	<setjmp.h>	●	●	●	●	●	●
4.7	<signal.h>	○	/	○	●	●	○
4.8	<stdarg.h>	●	●	○	○	●	○
4.9	<stdio.h>: OK to include more than once	●	●	●	●	●	○
4.9.1	In <stdio.h>: FOPEN_MAX and FILENAME_MAX[2]	○	○	○	○	○	○
4.9.2	Stream and file semantics: opening, seeking, text-binary	●	○	●	●	●	○
4.9.4	remove, rename	●	○	●	●	●	○
4.9.6	Full printf/scanf to new exact spec	●	/	●	/	/	/
4.9.6	scanf and ungetc push back independently[2]	●	○	●	●	/	○
4.9.6	v*printf, v*scanf	●	○	●	●	●	●
4.9.9	fsetpos, fgetpos	●	○	○	●	●	○
4.10	<stdlib.h>: has size_t[2]	●	●	●	●	●	●
4.10.1	strtol, strtoul, strtod[2]	●	○	/	/	/	/
4.10.4	system, atexit, getenv, EXIT_SUCCESS, EXIT_FAILURE[5]	/	○	○	○	/	○
4.10.6	div, ldiv	●	○	○	●	●	●
4.11	<string.h>: memset and strcmp	●	●	●	●	●	●
4.11.2	memmove, strstr[1]	●	○	○	/	●	/
4.11.6.2	strerror	●	○	/	●	●	●
4.12	<time.h>	●	●	●	●	●	●
4.12.2.2	difftime	/	○	/	●	●	/
4.12.3.5	strftime	/	○	/	/	●	○

ates inside a loop that's weighted according to the frequencies (as measured by Klerer and Liu) with which programmers inadvertently use such constructs. By far, the greatest weights are assigned to routines executing constant folding and local common subexpression optimization.

The efficiency of a given compiler's optimizer is, in theory, simply NON-OP's time divided by OP's time. The ideal optimizing compiler would score 1, as it would eliminate all the unneeded code from NONOP. Higher scores indicate poorer performance.

Working with the benchmarks, we found the presence or absence of a coprocessor to be a significant factor. Both OP and NONOP are math-intensive, relying heavily on floating-point functions from the function library. Moreover, since NONOP may be forced to do more floating-point math than OP, a good (or poor)

An ANSI Conformance Sampler

Thomas Plum

After a 5-year effort, the X3J11 committee unanimously voted last September to forward a proposed standard for C to ANSI for final approval. No further changes are anticipated before the standard's eventual publication early this year.

Now that it's clear exactly what standard C will look like, most C programmers will want to know how close various compilers conform to the ANSI standard. I have created a set of tests that attempt to do just that.

The ANSI standard for C does not have any subsets. There is no "partial conformance" to this standard, and there can be no "degrees of conformance." Vendors of C have been asked not to specify or claim conformance to any of the various working drafts, including the current proposed standard, which will become a full American standard only upon official publication by ANSI. Even at that point, its acceptance as an international standard depends on further decisions by the appropriate bodies of the International Organization for Standardization.

Nonetheless, until most compilers achieve official certification, it's useful to have an unofficial estimator of how closely the compilers conform to this new standard. This is the purpose of SAMPLER.88, which contains the C source code for 79 compilable tests that sample the closeness of a compiler to the soon-to-be ANSI standard for C (for information on obtaining the listing, see page 3). SAMPLER is a limited subset of the Plum Hall validation suite for C. Test results for the six packages in the Product Focus appear in table 2.

SAMPLER in no way attempts to measure conformance to the standard. It estimates the number of areas in which a current compiler will need to be modified to achieve eventual conformance to the standard. Compared to the megabytes of source code in a full validation suite for C, it's relatively tiny.

So much for caveats. On the positive side, any compiler that embodies all the features of the May 1988 draft would pass all 79 of the tests in this file. Most of the tests in SAMPLER reflect aspects of ANSI C that have been unchanged since mid-1986; 13 tests reflect 1987 decisions, and 3 tests reflect 1988 decisions. When 1989 compilers start to satisfy all or most of these tests, programmers will truly have syntactic portability in C.

As of the date these tests were run, several vendors already had beta-test versions of their products that scored substantially better. Because of editorial lead times, some of these beta versions will reach the market by the time this article appears. Some vendors have chosen to wait until the standard is announced officially before releasing compilers that are tracking the standard. The point here is that you should inquire about the current status of each compiler that you are considering.

I wish to express my gratitude to the authors of the six compilers for the contribution that they have made to the field of programming. Each had unique personal reasons for undertaking a project of this magnitude.

Comparisons based on SAMPLER are in no way meant as determinations of the quality of these compilers; they are simply a snapshot of one aspect of the situation at this point in time.

A thorough validation of a compiler requires much more work than SAMPLER provides. The British Standards Institution's recent competitive evaluation has selected the Plum Hall validation suite for C. It will be used by BSI, IMQ (of Italy), and AFNOR (France) in the European C validation service. Our goal is to allow programmers to write in C with the certainty that, if portably written, their programs will work in each new environment.

Thomas Plum is chairman of Plum Hall and cocreator of the Plum Hall validation suite for C. He has been vice chairman of the X3J11 committee since 1983 and is the author of five textbooks and several curricula on C, including the course used by Bell Labs. He can be reached on BIX c/o "editors." For information about the Plum Hall validation suite for C, contact Joan Hall at (609) 927-3770. SAMPLER was written by Thomas Plum and Ralph Ryan (president of Chiron Systems), with assistance from Joan Hall, Don Gallagher, and Scott Erlichman.

math library can skew the results. Keeping this caveat in mind, as well as the varying efficiencies of the code generators, you can roughly gauge the effectiveness of the optimizers.

Microsoft and MetaWare split top honors in the two (with coprocessor and without) environments, and where they didn't win, they did show adequate performance. These two compilers are known to have well-developed optimizers, so the result is not surprising.

One odd result was Watcom's performance: tied for first in one environment and dead last in the other. The more representative result is probably Watcom's good showing on the 80387-equipped Compaq. While the compiler is not designed to optimize loops, it does support most other source optimizations. In the emulation environment, Watcom's unusually good raw times—particularly on the OP test, which Watcom executed more than twice as fast as its nearest competitor—may, ironically, account for its poor showing by magnifying idiosyncrasies in the libraries.

80x86 Specifics

The classic optimizations are rooted in textbook compiler design and can't take into account the quirks of real-world processors. The 80x86 family, with its segmented architecture, specialized registers, and dedicated string instructions, represents anything but the ideal machine around which to design a compiler.

To attain the execution speeds we've seen these compilers display, their authors have had to address basic machine-specific issues. For our system-specific tests, we looked into the assembly language code generated by each compiler, and we correlated that with execution speed to gauge the efficiency of some common and representative operations. We wrote one function for each test, stored them in separate files, and compiled them with appropriate optimization and assembly language–listing switches. A separate (unoptimized) driver took care of running and timing the tests.

The first tests examine register use. We were looking for evidence that the compilers can retain frequently used variables in registers, rather than in

continued

Table 3: *Benchmark results. We ran each test twice, once on a Compaq 386 portable equipped with an 80387 floating-point unit, and once on a Northgate 386 with no coprocessor. The specifics of each test are discussed in the text. Indexes represent an overall summary of a compiler's performance in a given area. Higher numbers are better, and a 1.00 is the highest possible score. The indexes are calculated by normalizing each test result against the best time for any compiler; the mean of these normalized figures becomes the index for a given group of tests (low-level, XLisp, optimization, and system-dependent). The compile time and code size index are calculated in the same way using results generated for each group.*

	Borland Turbo C		Manx Aztec C86		MetaWare C		Microsoft C		Watcom C		Zortech C	
Low-level												
Compile time[1]	10.91	13.95	19.88	21.61	33.97	56.49	30.30	32.41	27.71	40.78	23.25	25.79
EXE size	17,508	27,348	10,862	13,712	22,512	29,632	23,372	24,883	16,000	18,200	23,594	23,746
Register int	0.41	0.44	0.40	0.43	0.38	0.43	0.38	0.41	0.25	0.30	0.37	0.41
Auto short	0.46	0.51	0.52	0.54	0.39	0.45	0.37	0.40	0.26	0.31	0.37	0.41
Auto long	1.20	1.30	1.34	1.42	1.15	1.32	1.24	1.40	0.96	1.03	1.16	1.31
Int multiply	1.20	1.50	1.46	1.68	1.46	1.69	1.15	1.40	1.19	1.48	1.18	1.47
Function call/return	1.50	1.50	1.55	1.53	1.52	1.53	1.52	1.50	1.43	1.44	1.53	1.54
Auto double	5.25	80.00	5.05	39.55	4.53	62.20	3.24	26.00	3.08	17.60	6.02	23.50
XLisp												
Compile time[1]	50.53	68.60	222.25	218.28	344.19	598.96	293.03	320.82[2]	290.31	317.31	242.66	218.69
.EXE size	55,658	65,498	52,388	56,444	58,944	68,816	62,646	66,334	45,474	50,096	61,662	62,198
Float	14.97	27.80	16.30	20.65	15.11	20.71	14.59	17.49	14.72	16.86	14.61	17.35
Integer sort	41.29	42.50	50.88	50.88	41.22	41.91	36.52	37.46	32.62	33.83	53.82	56.30
Float sort	34.26	38.85	44.57	45.26	36.06	37.74	32.22	33.45	28.73	29.88	47.13	49.49
Write file	13.32	14.40	14.86	15.38	15.87	16.42	13.27	13.29	13.29	13.40	12.52	13.04
Read file	12.31	13.77	14.50	14.83	14.61	16.50	11.89	13.76	11.72	13.41	11.95	12.80
Function calls	14.99	16.51	17.98	18.45	16.31	16.81	14.39	15.02	13.78	14.89	14.72	15.59
Sieve	53.41	55.27	57.62	58.41	54.98	56.35	53.34	55.09	56.66	60.23	55.75	58.05
Optimization												
Compile time[1]	14.64	15.82	26.59	31.04	42.42	78.10	45.15	49.13	47.71	102.00	54.52	59.81
EXE size	40,832	60,512	25,336	32,456	54,528	72,272	47,942	63,526	35,888	46,516	30,856	54,902
OP	5.40	62.70	7.20	63.9[3]	4.50	51.10	3.50	60.1[4]	4.70	22.80	5.80	45.10
NONOP	7.10	98.80	16.70	195.7[3]	5.45	100.00	4.40	86.1[4]	5.80	76.20	9.20	110.50
NONOP / OP[5]	1.3	1.6	2.3	3[3]	1.2	2.0	1.3	1.4	1.2	3.3	1.6	2.5
Object-level												
Compile time[1]	11.79	19.83	25.45	20.87	60.94	109.94	52.25	41.36	30.10	39.44	35.07	27.43
EXE size	26,242	26,242	10,050	11,476	21,760	26,784	31,218	28,806	18,426	18,426	15,434	15,570
Register usage	14.78	16.43	15.30	16.45	14.25	14.82	15.05	16.36	12.09	12.79	14.25	15.27
Pointer arithmetic	25.60	26.62	41.80	40.89	38.70	37.63	17.06	18.18	17.46	19.49	15.35	18.19
32-bit arithmetic	4.07	4.26	4.18	5.19	4.04	4.55	3.54	4.38	2.04	2.08	4.01	4.83
strncmp function	17.47	18.30	7.78	8.18	39.63	40.70	17.85	18.70	4.67	4.50	19.88	20.82
strncpy function	20.60	24.76	34.01	36.52	21.66	25.04	33.64	36.36	42.62	45.97	23.40	17.71
Minimum XLisp size[5]		64,490		56,444		67,136		65,230		49,788		61,138
Compile time index	1.00	1.00	0.45	0.60	0.25	0.19	0.27	0.36	0.32	0.30	0.32	0.46
Code size index	0.61	0.63	0.97	0.94	0.55	0.59	0.51	0.62	0.73	0.85	0.67	0.72
Plum index	0.75	0.70	0.69	0.70	0.76	0.70	0.84	0.81	0.99	0.99	0.77	0.82
XLisp index	0.92	0.85	0.79	0.80	0.85	0.84	0.95	0.96	0.98	0.98	0.87	0.87
OPT index	0.78	0.74	0.45	0.42	0.90	0.67	0.98	0.82	0.87	0.71	0.65	0.59
System index	0.64	0.64	0.57	0.57	0.56	0.58	0.63	0.64	0.87	0.89	0.69	0.67

Times for the low-level benchmarks are in microseconds; all other times are in seconds. File sizes are in bytes. For each compiler, coprocessor times appear in the left column and emulator times in the right column.
[1] Compile times include times for compile, assemble, and link.
[2] Intrinsic function optimizations disabled; failed to compile with option enabled.
[3] fmod function used in optimizing benchmark not supported in emulator library; used function provided by Manx.
[4] Emulator library used rather than Alternate Math library.
[5] Score weighted twice in index calculation.

memory. We used the following code fragment:

```
iFunction
    (a1,a2,b,c1,c2,d1,d2)
a = a1 + a2 + b;
c = c1 + c2 + b;
d = d1 + d2 + b;
return a+c+d;
```

And we anticipated that an efficient compiler would generate an assembly language listing like this:

```
MOV      AX,[BP+a1 offset]
ADD      AX,[BP+a2 offset]
MOV      BX,[BP+b offset]
ADD      AX,BX
ADD      AX,[BP+c1 offset]
ADD      AX,[BP+c2 offset]
ADD      AX,BX
```

which keeps the frequently used variable b in BX and keeps a running total of the return value in the return register.

Watcom turned in an outstanding time here, 12 seconds compared to the average 14.5. The assembly language code also came closest to our ideal. Watcom keeps b in a register, keeps subtotals in other registers, and never makes a single swap to memory. The compiler also passes arguments in registers rather than on the stack, as is customary. This fascinating strategy contributed here and elsewhere to Watcom's outstanding performance by reducing memory accesses to an absolute minimum.

While it's true that arguments must be moved to registers by the calling function, the additional overhead of passing arguments on the stack is eliminated. Most compilers push the addresses of variables onto the stack, call the subroutine, and load the register using their addresses. Watcom's method involves only one transfer, directly from memory to register.

Zortech also scored well on the register usage test. The emitted code does keep all variables in memory, but it keeps the subtotals in registers and makes the final add (for the return value) by simply adding registers. MetaWare's close finish can be attributed to its similar scheme.

While Turbo does keep b in a register (SI), its subtotaling in memory did cost some time. Microsoft and Aztec seem to be the least efficient on this common operation, as they swap every variable in and out of memory.

A similar test involves the efficient manipulation of pointers. Consider the following code fragment:

```
int a,i,*p;
for (i=0;i<ITERATIONS;++i)
a += *(p+i);
```

We expected the compilers to recognize that p is loop-invariant and to keep it as the source index, which would produce code like this:

```
LO: MOV BX,CX ;i is in cx
    SHL BX,1  ;multiply for 16
              bit
    ADD AX,[SI+BX];accumulate a
    INC CX    ;increment index
    CMP CX,ITERATIONS ;check
```

Microsoft did us one better, however. It first compared ITERATIONS to zero, and it was prepared to exit immediately if the result were true. That wasn't the case, so it then moved *p into SI as anticipated, but it also moved ITERATIONS into CX. With this arrangement, the code simply LOOPed around a fragment containing nothing but

```
ADD DX,WORD PTR [SI];a is in DX
ADD SI,2
```

The LOOP command continues CX times. While the loop uses an ADD rather than the more efficient INC, it does eliminate a number of instructions. The LOOP construct is also far more efficient than the CMP and JLE combination. It hadn't occurred to us that the loop index itself is unnecessary in an increment-only loop.

Zortech, which actually came in first in this test, was the only other compiler that eliminated the loop index. It produced code similar to Microsoft's, but it used the CMP rather than the LOOP scheme to check the index. Where Zortech gained ground was in its initialization routine, which was significantly shorter than Microsoft's. To test this, we cranked up the number of ITERATIONS to make the loop time the dominant influence on overall time; this time, Microsoft finished first.

In fact, only the two compilers that finished last on our timing, MetaWare and Aztec, produced the base-plus-offset [BX + SI] instruction we had anticipated. Watcom and Turbo scored well by resisting the temptation to swap variables to memory wherever possible.

Our third test was simply to determine the ability of the compiler to do 32-bit arithmetic. Only MetaWare's compiler supports 80386-specific code, so we expected the C expression

```
long a,b,c;
```

```
a=b+c;
```

to translate to

```
MOV EAX,B
ADD EAX,C
MOV A,EAX
```

only in MetaWare's case. In reality, all the compilers produced the same code, and none of them used the 80386-specific instructions. MetaWare explained that doing so would be dangerous if the resulting application were to be used under any operating system that needed to switch from real to protected mode. As a result, the times for all the compilers were similar, but Watcom's efficient register use put it over the top.

String operations occur frequently in C (indeed, in any language), and the 80x86 family's specialized string instructions make it almost the perfect vehicle for the string-oriented functions that an ANSI C compiler must provide. We took the following two functions, strncmp and strncpy, and used a debugger to determine exactly how strings were handled.

Timing the strncmp function led to some interesting results. The function is designed to compare n characters in two strings and determine whether or not the strings are equal, and, if not equal, to determine which one is lexicographically greater. Characters in each string are compared only as long as they are equal; the first differing character causes the function to return.

We expected that the most efficient compilers would use the string-specific REPeatwhileEqual CoMPareString construct rather than comparing and looping with each element in the string.

Our test used two nonequal strings. Watcom's compiler, though it used the less efficient compare and jump method, gave an outstanding performance. Every compiler but Watcom first scans the strings to determine their lengths, compares this to n, and uses whichever value is lowest as a loop index. It's in that loop that REP CMPSB is used to compare the strings by an efficient compiler. Watcom's first action is to compare the strings; it then continues to loop until n decrements to zero or until the program reaches a null terminator in one of the strings. Because the initial scan was skipped, Watcom turned in an excellent score.

To further explore this aberration, we reran the test with equal strings. Watcom fared very poorly, at 64.2 seconds on the

In 1988, $3.5 billion in micro-computer software will be sold worldwide. During that same time, another $3.0 billion in sales will be lost to free distribution — better known as software piracy. And right now, Rainbow Technologies' Software Sentinel™ is protecting close to $1.0 billion in software for developers who never wanted to be part of the free software distribution network in the first place. ( ) The Software Sentinel hardware key is "execution control" software protection. It ships with the software and simply plugs into the PC's parallel port to be one hundred percent invisible to both user and the soft-ware. Users can make as many copies as they want. Make working submasters. Use a hard disk. Virtually anything that can be done with unprotected software. Except start freely distributing that software to other users. ( ) The Rainbow fam-ily of Software Sentinel products. Se-lected by the very big to the not-so-big developers of DOS, OS/2 and Xenix software in worldwide markets. To the cool tune of close to a billion dollars. So far.

RAINBOW TECHNOLOGIES

Circle 195 on Reader Service Card (DEALERS: 196)

Compaq, compared with 4.67 seconds when the strings were equal.

Turbo's compiler, which used the CMPSB instruction, turned in a good score at 17.47 seconds. In contrast to Watcom, Turbo's score increased to only 28.28 seconds when equal strings were used. All the other compilers used combinations of SCAnString and loop instructions. Aztec's compiler proved to be particularly efficient, comparing words (LODSW,SCASW) rather than bytes as an initial check for equality.

The strncpy timing yielded less-surprising results. The strncpy function copies n characters from one string to another. It needs to determine whether the first string is actually n characters long, and it does so by looking for a null termination. We expected the copy part of the loop to be built around a MOVeString instruction.

All the compilers we examined used this construct. Only Watcom didn't generate code using the REP prefix, and it finished last on the benchmark. Turbo and MetaWare, which virtually tied for the best scores, use the REP prefix both to move the string and to scan (SCAS) the first string to determine its length.

Compilers Up Close

The benchmark figures say a lot about each package, but they can't tell the whole story. Each compiler has features that simply can't be benchmarked, like an integrated debugger or library functions that transcend the ANSI standard. While one may be an excellent performer, it may be so difficult to use that it intimidates newcomers to the language.

In the following sections, we describe some of these features that can make the difference between an outstanding compiler and one that is merely so-so. The products appear below in alphabetical order.

Borland Turbo C Professional 2.0

Borland has a well-deserved reputation for pricing good software so that users on a limited budget can get in on the action. This is no exception: The Professional version, which includes an integrated environment, Turbo's Assembler, and the Turbo Debugger, sells for $250. It runs under DOS 2.0 or higher and requires 448K bytes of RAM and two floppy disk drives.

The integrated environment is window-oriented; from within it, you can edit, compile, and debug programs without exiting to DOS. A project facility is the integrated environment's version of make; you use it to specify multiple files to be compiled and linked. The compiler itself, which can be used either in the environment or from the command line, adds to the ANSI standard the ability to include in-line assembly language code in a remarkably flexible manner. And it compiles faster—*much* faster—than any other compiler we reviewed.

The best addition to this already well-received package is the new Turbo Debugger. It's provided as a stand-alone supplement to the debugger that comes with the integrated environment. In addition to breakpoints and watchpoints, Turbo C features a powerful data inspector that can unpack and display entire structures. Because the inspector works in conjunction with the source-level trace facility, you can literally see those structures change.

If you have an 80386 system and plenty of extended memory, you can use this debugger in so-called virtual mode—and

Circle 170 on Reader Service Card

get most of the benefits of a hardware debugger. It installs itself in memory above 1 megabyte and monitors the execution of the subject program on a virtual 8086. As with hardware debuggers, you can watch a region of memory with minimal performance degradation.

All in all, it's quite an impressive package. Turbo C won't give you the fastest code or the smallest executable, but it's definitely the one that is the most fun to use.

Manx Aztec C86 4.1d

One of the C language's greatest selling points is portability. Manx takes full advantage of this by offering a C compiler for several operating systems, including this version for MS-DOS. It runs on systems with a minimum of DOS 2.0, 384K bytes of RAM, and two floppy disk drives.

The $499 package is fully featured; an editor, linker, librarian, and source-level debugger come bundled with the compiler. The commercial version also includes object modules for generating code suitable for ROM programming and additional programming utilities.

While the compiler produces some efficient and compact code, benchmark scores were generally poor. The library functions, especially those involving the coprocessor, exhibited some obscure bugs, surprising in such a mature product. Source-level debugging is command-oriented and can be very efficient, but the debugger lacks the sexy window-oriented features supported by its current competition.

The compiler does support some unusual features designed to maintain compatibility with Aztec C for Unix systems. Since some of these are not ANSI standard, they can be enabled or disabled with a compiler switch. It makes for an easy transition between Aztec's Unix and DOS C compilers.

Overall, Aztec C suffers more than it benefits from its long and diverse heritage. The compiler seems to be a product made to fit in the DOS environment rather than designed for it; the poorly organized documentation (designed for use with different Aztec C compilers) exemplifies that problem.

MetaWare High C 286 1.4

A quick tour of table 1 may lead you to believe that High C lacks some of the basics, but that's not true. High C isn't as easy to use as some of the others, but it does everything an ANSI C compiler has to do and then some. The package requires 384K bytes of RAM and a hard disk drive, and it supports DOS 2.0 and higher.

MetaWare's extensions to C create a potent new dialect of the language: You can mix declarations and statements, use ranges in conjunction with case statements, write nested functions à la Pascal, and associate parameters by name, as in Ada. MetaWare also offers, as a $100 option, library support for OS/2 development.

While ranged case statements and intermixed declarations may buy some convenience, they are minor enhancements to C compared to the wholesale changes brought on by the Pascal- and Ada-like additions. Nested functions enable the programmer to define functions within other functions, with the inner function retaining all the variables local to the outer function. This makes pointer passing unnecessary when variables need to be modified by an external function and simply makes code cleaner.

continued

Circle 211 on Reader Service Card

World Radio History

Company Information

Borland International, Inc.
1800 Green Hills Rd.
P.O. Box 660001
Scotts Valley, CA 95066
(800) 543-7543
(408) 438-8400
Inquiry 1040.

Manx Software Systems, Inc.
P.O. Box 55
Shrewsbury, NJ 07701
(800) 221-0440
Inquiry 1041.

MetaWare, Inc.
903 Pacific Ave., Suite 201
Santa Cruz, CA 95060
(408) 429-6382
Inquiry 1042.

Microsoft Corp.
16011 Northeast 36th Way
P.O. Box 97017
Redmond, WA 98073
(800) 426-9400
Inquiry 1043.

Watcom
415 Phillip St.
Waterloo, Ontario
Canada N2L 3X2
(800) 265-4555
Inquiry 1044.

Zortech, Inc.
366 Massachusetts Ave.
Arlington, MA 02174
(800) 848-8408
Inquiry 1045.

Named parameter associations let you pass function arguments by name rather than position, again making the source code easier to understand.

But for $595, MetaWare ought to supply a full complement of tools. While there are some Unix-like utilities, there is no debugger, library utility, or make utility; the package doesn't even include a linker. That may not be an issue for the professional programming audience that High C targets, but it certainly makes the compiler a poor choice for novice or casual C programmers.

Microsoft C 5.1

Version 5.1 of the Microsoft C compiler adds OS/2 support to the excellent DOS-only version 5.0. The package includes the compiler, a debugger, and a set of real- and protected-mode libraries for $450. The package supports Microsoft OS/2 1.0 or DOS 2.1 and higher, and it requires at least 448K bytes of RAM and a hard disk drive.

A number of unusual features supplement its fine performance on our benchmarks. One of them, of course, is the ability to develop OS/2 applications. The package includes library functions for OS/2 and a bind utility that creates bound executables suitable for both real- and protected-mode execution. The package also offers support for developing Windows applications.

Also included is QuickC, which implements an integrated environment for fast development. QuickC favors ease of use over a full-fledged optimization and

debugging capability. The combination of QuickC with the standard compiler and its CodeView debugger endows Microsoft C with a development versatility shared only by Watcom C.

Microsoft's CodeView debugger isn't quite as easy to use as Turbo C's, but they share most of the same features. There's also a librarian, a make utility, and a programmable text editor.

There's little not to like in the overall package. The function libraries support screen graphics as well as DOS and BIOS calls. The documentation is excellent. If anything is missing, it's the lack of the #asm preprocessor directive for including in-line assembly, but that one small (non-ANSI) flaw is hard to hold against an otherwise superior product.

Watcom C 6.5

This compiler isn't quite as slick as Microsoft's, and it doesn't have as many features. What it does have, however, is excellent code generation that translates into unmatched execution speed. We ran lots of tests; Watcom took more blue ribbons than any other compiler.

Watcom's $495 compiler invites comparison with Microsoft's. It, too, includes an editor, debugger, and separate integrated environment compiler. There are also library, make, and disassembler utilities. To run it, you need at least DOS 2.0, 512K bytes of RAM, and two floppy disk drives.

Watcom's compiler turned out to be the most ANSI-compatible, according to

our validation suite. Extensions to the standard include the near, far, cdecl, pascal, and fortran identifiers found in most other compilers we reviewed. Watcom also makes extensive use of pragmas (i.e., compiler directives). You can use these to specify the registers used when making function calls (even to the 80x87) and to invoke calling conventions that may or may not be like those in C. This feature adds flexibility to the highly efficient method of passing function arguments in registers.

Some ease-of-use features are missing. While the C driver supports wildcard filenames, the compiler command line does not. And the debugger, though window-oriented, isn't as intuitive as that of some other packages.

A new release scheduled for introduction early this year promises to improve Watcom's optimizing capabilities. According to the company, the new product will focus on the classic optimizations, adding even more power to an already well-built compiler.

Zortech C 1.07
This stand-alone version of the C compiler included with Zortech's C++ package may well make a name for itself in the heavyweight C arena. The inexpensive ($89.95) compiler lacks some of the features of those in the $400 range, but it posted good scores on our benchmark tests, often delivering better times than the more expensive products. The product runs in 512K bytes of RAM and requires at least DOS 2.0 and two floppy disk drives.

Zortech C isn't a bare-bones compiler, however. It includes an integrated editor/compiler with WordStar-like editing commands, and you can make programs compiled with Zortech C compatible with the Microsoft CodeView debugger. According to Zortech, a new release that includes its own source-level debugger should be available by the time you read this.

Only one library is included, in contrast to the usual standard C function library, math library, and additional math library for floating-point emulations. The single library produces both 80x87 and emulated math code. Zortech's far better relative performance on our machine without an 80387 indicates that its emulation functions are much more efficient than its coprocessor code.

Two nonstandard preprocessor directives are supported—#message, which prints a message during compilation, and #exit, which sets the error level and makes a clean exit from an aborted com-

pile. Additional nonstandard features include library support for the Microsoft Mouse and for sound functions.

Optimized Choices
Our XLisp benchmark, an example of a complex and realistic application, showed how close these compilers really are. While one may have good library functions and another excellent source optimizations, for instance, overall performance of any of these compilers fits within a fairly small window.

Watcom and Microsoft are the two obvious performance standouts, splitting virtually all our tests for execution speed. While both are good in all areas, they have their own specialties; Watcom is definitely the leader in object-level optimizations, while Microsoft is strongest on the source level. Watcom's object-level expertise is evident in all the test results, and it is well matched to its target processor.

For flat-out executable speed, Watcom's compiler was the clear winner. Microsoft finished a close second. Since speed is most valuable when coupled with functionality, it pays to compare these two in the features table as well. Both compilers sport an impressive list of features, but OS/2 support and a better debugger give Microsoft a slight edge in that category.

If speed is absolutely critical and OS/2 compatibility isn't, choose Watcom. If you want to do OS/2 development now, you'll obviously choose Microsoft; and even if you're not developing for OS/2, you might still prefer Microsoft's friendlier and more powerful tools.

Finally, the other compilers have points in their favor. Turbo has by far the fastest compilation times and the best source debugger. Aztec produces the smallest executables. MetaWare combines excellent optimization with valuable language extensions. And Zortech does everything that a compiler has to do—at an attractive price. ∎

Steve Apiki is a BYTE Lab testing editor, and Jon Udell is a BYTE technical editor. They can be reached on BIX as "apiki" and "judell," respectively.

A Pair of Sophisticated Laptops

The Zenith SupersPort
286 and Mitsubishi
MP-286L typify
the current crop
of portable computers

John Unger

The Mitsubishi MP-286L (left)
and the Zenith SupersPort 286.

For a number of years, there were two distinct types of portable computers. First on the scene were the AC-powered Osborne, Kaypro, and Compaq portables; they offered desktop power combined with heavy weight and bulk. Next came the laptops, such as the Tandy Model 100, which initially appeared as battery-powered microcomputers that sacrificed computing power, storage capacity, and ergonomics for convenience, size, and portability. Each type of portable had its inherent advantages.

Today the distinctions between these two classes have blurred as the AC-powered portables have shrunk in size and the battery-powered laptops have increased their processor horsepower and features. The Zenith SupersPort 286 and the Mitsubishi MP-286L represent the merging of the two classes of microcomputers into mature, sophisticated machines that have the best features of both. With the current generation of laptops, exemplified by the SupersPort 286 and MP-286L, you truly have the computing power of a desktop or transportable machine and the convenient size and weight of a laptop computer.

As similar as these two computers are, they are clearly designed for two different types of users. Zenith has tried to make the SupersPort 286 one of the best-performing battery-powered laptops, while the Mitsubishi MP-286L appears to be offered as a high-performance, lower-cost alternative to other AC-powered 80286-based laptops, such as the Toshiba T3100 and T3200 and the larger transportable computers like the Compaq Portable III.

Under the Skin

These two machines share many of the same hardware characteristics. Both feature 80286 processors running at 12 MHz, a 1.44-megabyte 3½-inch floppy disk drive, and a 20- or 40-megabyte hard disk drive. My Zenith SupersPort 286 review unit had a 40-megabyte hard disk drive, 1 megabyte of zero-wait-state RAM, and an optional internal modem; the Mitsubishi MP-286L came with a 20-megabyte hard disk drive and 640K bytes of RAM with one wait state. The SupersPort's list price with a 20-megabyte hard disk drive is $4999; a similarly configured MP-286L costs $3995. If you substitute a 40-megabyte drive in the SupersPort, the price jumps to $5599.

Both computers have connectors for an optional external keypad, a Centronics-compatible parallel printer port, a nine-pin RS-232C serial port (the MP-286L has two serial ports), a connector for an external 360K-byte 5¼-inch floppy disk drive, and a nine-pin RGB connector for an external CGA monitor.

Where these systems diverge is in their power sources. The SupersPort 286 has three major components: the main laptop unit, a removable nickel-cadmium 48-watt-hour battery pack, and an external DC power supply/battery charger. The basic computer weighs 10½ pounds, but you must also attach either the 1-pound power supply or the 4-pound battery. Zenith has gone to great lengths to cut power consumption to make a battery-powered portable practical. The SupersPort uses CMOS versions of both the 80286 CPU and optional 80287 math coprocessor chips; this cuts down on both power and heat and makes it unnecessary for this laptop to have a cooling fan.

The SupersPort has a useful monitor/setup program in ROM that is easily invoked with a Ctrl-Alt-Insert key combination. With this ROM-based software, you can perform a variety of tests and make system changes. To conserve battery power, you can use the monitor program to disable the optional internal modem, to turn off the parallel port's circuits, to select the amount of time that the display's backlighting remains on when there is no keyboard activity, and

continued

World Radio History

Company
Zenith Data Systems
1000 Milwaukee Ave.
Glenview, IL 60025
(800) 842-9000

Components
Processor: 80C286 CMOS 16-bit processor running at 6 or 12 MHz, keyboard-selectable; optional CMOS 80C287 math coprocessor
Memory: 1 megabyte of RAM standard, expandable to 2 megabytes internally
Mass storage: One 1.44-megabyte 3½-inch floppy disk drive; 20- or 40-megabyte hard disk drive; optional 5¼-inch external floppy disk drive
Display: Electroluminescent backlit LCD with 10½-inch diagonal screen
Keyboard: 84 full-size keys with 12 function keys; optional numeric keypad
I/O interfaces: One RS-232C port with DB-9 connector; Centronics parallel printer port with DB-25 connector; RGB color monitor port with DB-9 connector; internal proprietary connectors for expansion memory and modem; external proprietary connectors for keypad and floppy disk drive

Size
3 × 12¼ × 12¼ inches (closed; 15½ inches deep with battery pack)
DC power supply: 6½ × 3 × 2 inches 10½ pounds; 14½ pounds with battery
AC power supply: 1 pound

Software
MS-DOS 3.21

Options
1200-bps internal modem: $299
2400-bps internal modem: $499
48-watt-hour battery pack: $289
26-watt-hour battery pack: $159
1-megabyte expansion (EMS) RAM: $1299
24-key detachable keypad: $129
External 360K-byte floppy disk drive: $399
Three-slot external expansion chassis: $499
80C287 coprocessor: $349
Technical-reference guide: $99

Documentation
120-page SupersPort 286 Portable Owner's Manual; 540-page MS-DOS 3.21 User's Guide and User's Reference; 42-page MS-DOS 3.21 Quick Reference Guide

Price
Zenith SupersPort 286 with 20-megabyte hard disk drive: $4999
Zenith SupersPort 286 with 40-megabyte hard disk drive: $5599

Inquiry 859.

Company
Mitsubishi Electronics America, Inc.
991 Knox St.
Torrance, CA 90502
(213) 515-3993

Components
Processor: 80286 16-bit processor running at 8 or 12 MHz (one wait state at 12 MHz, zero wait states at 8 MHz); 80287 math coprocessor (optional)
Memory: 640K bytes of RAM standard, expandable to 2.6 megabytes internally
Mass storage: One or two 1.44-megabyte 3½-inch floppy disk drives or a 20-megabyte hard disk drive
Display: Page-white Neutral Twisted Nematic LCD with cold CRT backlighting; 11-inch diagonal screen
Keyboard: 86 keys with 12 function keys; optional separate 17-key numeric keypad
I/O interfaces: Two RS-232C ports with DB-9 connectors; Centronics-compatible parallel printer port with DB-25 connector; CGA monitor port with DB-9 connector; internal proprietary connectors for expansion memory and modem; external proprietary connectors for keypad and floppy disk drive

Size
3½ × 12¼ × 14¼ inches; 16 pounds

Software
MS-DOS 3.3 and GWBASIC 3.20; user diagnostics disk; Super PC-Kwik disk cache

Options
2400-bps internal modem: $499
2-megabyte expansion RAM: $1530
17-key detachable keypad: $158
External 360K-byte floppy disk drive: $525
80287 coprocessor (factory-installed): $600
Microsoft OS/2: $325

Documentation
115-page MP-286L User's Guide; 295-page MS-DOS 3.3 User's Guide; 26-page MS-DOS Quick Reference; 18-page MP-286L Quick Reference

Price
MP-286L-210 (dual 1.44-megabyte floppy disk drives): $3195
MP-286L-220 (20-megabyte hard disk drive): $3995
MP-286L-220VP (same as -220 but with 80287 coprocessor): $4595
MP-286L-240E (40-megabyte hard disk drive and EGA): $4795
MP-286L-240EVP (same as -240E but with 80287 coprocessor): $5395

Inquiry 860.

to set the amount of time that the hard disk drive continues to spin after the last disk access. If the hard disk drive is not spinning, it takes about 8 to 12 seconds for the drive to respond to a request from the operating system.

I found that a fully charged battery provides about 4 hours of continuous operation without using any of the conservation measures. The system begins to beep pleadingly when the battery is about to die; you have 4 or 5 minutes to take some action before low battery power finally shuts down the system.

Three proprietary internal expansion slots are provided in the SupersPort: one for an internal modem, one for expansion RAM (1 megabyte), and one for an external expansion bus interface. This interface is used to connect an optional external three-slot expansion chassis to the laptop; the expansion chassis can hold three full-size, 8-bit, IBM PC XT–compatible cards. The optional 80C287 math coprocessor's socket is accessed by removing a neat trap door in the bottom of the laptop.

The MP-286L's chassis features four internal proprietary expansion buses. You can use the buses to install up to 2 megabytes of expansion RAM, a 2400-bit-per-second modem, or an additional EGA or VGA display controller for use with an external monitor. The space allotted for these slots and the MP-286L's built-in power supply and cooling fan make its case almost as long as the SupersPort's with its battery pack attached. The MP-286L's optional 80287 coprocessor must be installed by the manufacturer; it is best to purchase the machine with the chip installed if you require one for your application.

Eyes and Hands
You constantly interface with two major elements of a microcomputer: its display and its keyboard. They are important in determining how a computer feels and how effectively you can interact with whatever software you're using. These two elements are also areas where laptop computers have been consistently criticized for their shortcomings.

Liquid crystal displays have come a long way since the first laptops and their small, slow, and barely legible "light gray on dark gray" screens. Both the MP-286L and the SupersPort 286 displays have excellent contrast and are easy to read under a variety of ambient lighting conditions. The width-to-height ratio of the screens is nearly identical to that found on most CRTs: 1 to 1.3.

Zenith SupersPort 286, Mitsubishi MP-286L

BYTE LAB

APPLICATION-LEVEL PERFORMANCE

Zenith SuperSport **6.4***

Mitsubishi MP-286L **5.6***

WORD PROCESSING	ZS 286	MP-286L
XyWrite III+ 3.52	Med./Large	Med./Large
Load (large)	:17	:14
Word count	:05/:35	:05/:36
Search/replace	:08/:34	:08/:33
End of document	:03/:20	:03/:21
Block move	:16/:16	:17/:17
Spelling check	:15/1:47	:16/1:54
Microsoft Word 4.0		
Forward delete	:42	:41
Aldus PageMaker 1.0a		
Load document	:21	:20
Change/bold	:41	:48
Align right	:31	:35
Cut 10 pages	:28	:32
Place graphic	:06	:09
Print to file	2:11	4:52
Index:	**1.59**	**1.45**

SPREADSHEET	ZS 286	MP-286L
Lotus 1-2-3 2.01		
Block copy	:05	:06
Recalc	:02	:02
Load Monte Carlo	:24	:23
Recalc Monte Carlo	:11	:12
Load rlarge3	:08	:08
Recalc rlarge3	:02	:02
Recalc Goal-seek	:07	:07
Microsoft Excel 2.0		
Fill right	:08	:09
Undo fill	3:30	3:23
Recalc	:03	:03
Load rlarge3	:40	:44
Recalc rlarge3	:02	:03
Index:	**1.53**	**1.41**

DATABASE	ZS 286	MP-286L
dBASE III+ 1.1		
Copy	1:05	1:42
Index	:21	:22
List	2:01	2:17
Append	2:10	3:04
Delete	:03	:04
Pack	1:42	1:52
Count	:18	:18
Sort	1:27	1:45
Index:	**1.28**	**1.05**

SCIENTIFIC/ENGINEERING	ZS 286	MP-286L
AutoCAD 2.52		
Load SoftWest	3:36	4:13
Regen SoftWest	3:20	3:56
Load StPauls	:58	1:04
Regen StPauls	:53	:55
Hide/redraw	42:08	47:21
STATA 1.5		
Graphics	1:46	1:41
ANOVA	1:00	1:05
MathCAD 2.0		
IFS 800 pts.	1:54	1:58
FFT/IFFT 1024 pts.	2:14	2:22
Index:	**0.64**	**0.59**

COMPILERS	ZS 286	MP-286L
Microsoft C 5.0		
XLisp compile	7:09	8:45
Turbo Pascal 4.0		
Pascal S compile	:08	:10
Index:	**1.40**	**1.13**

All times are in minutes:seconds. Indexes show relative performance; for all indexes, an 8-MHz IBM PC AT = 1

Compaq 386/20 **17.9**

IBM PS/2 Model 50 Z **8.1**

IBM PC AT **5**

- Word Processing
- Spreadsheet
- Database
- Scientific/Engineering
- Compilers

*Cumulative applications index. Graphs are based on indexes at left and show relative performance

---

LOW-LEVEL PERFORMANCE[1]

Zenith SuperSport

CPU	ZS 286	MP-286L
Matrix	7.17	7.77
String Move		
Byte-wide:	69.01	52.71
Word-wide:		
Odd-bnd.	52.32	52.69
Even-bnd.	34.49	26.37
Sieve	39.22	48.50
Sort	39.52	38.08
Index:	**1.55**	**1.62**

FLOATING POINT	ZS 286	MP-286L
Math	N/A	N/A
Error[2]		
Sine(x)	N/A	N/A
Error		
e^x	N/A	N/A
Error		
Index:	**N/A**	**N/A**

DISK I/O	ZS 286	MP-286L
Hard Seek[3]		
Outer track	3.31	6.63
Inner track	3.33	4.98
Half platter	9.94	21.64
Full platter	9.93	36.64
Average	6.63	17.47
DOS Seek		
1-sector	15.60	21.93
32-sector	62.38	60.16
File I/O[4]		
Seek	0.26	0.22
Read	1.28	1.50
Write	1.24	1.44
1-megabyte		
Write	8.97	7.64
Read	8.79	7.38
Index:	**1.06**	**0.92**

VIDEO	ZS 286	MP-286L
Text		
Mode 0	10.05	9.56
Mode 1	10.03	9.59
Mode 2	8.29	10.38
Mode 3	8.29	10.40
Mode 7	N/A	N/A
Graphics		
CGA:		
Mode 4	3.33	3.52
Mode 5	3.33	3.52
Mode 6	3.52	3.70
EGA:		
Mode 13	N/A	N/A
Mode 14	N/A	N/A
Mode 15	N/A	N/A
Mode 16	N/A	N/A
VGA:		
Mode 18	N/A	N/A
Mode 19	N/A	N/A
Hercules	N/A	N/A
Index:	**1.38**	**1.29**

CONVENTIONAL BENCHMARKS	ZS 286	MP-286L
LINPACK	3800.84	4022.36
Livermore Loops[5] (MFLOPS)	0.01	0.01
Dhrystone (MS C 5.0) (Dhry/sec)	2890	2615

Mitsubishi MP-286L

Compaq 386/20

IBM PS/2 Model 50 Z

IBM PC AT

- CPU
- FPU
- Disk I/O
- Video

N/A = Not applicable
[1] All times are in seconds. Figures were generated using the 8088/8086 versions (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

For a full description of all the benchmarks, see "Introducing the New BYTE Benchmarks," June 1988 BYTE

At 9 by 6¾ inches, the MP-286L's screen is slightly larger than the SupersPort's, which is 8 by 6 inches. The SupersPort's backlit screen has the familiar blue characters on a silver-gray background. The MP-286L features a very light gray background (the company calls it paper-white, but it's not) with black characters. A switch on the front of the MP-286L's screen changes the display from regular to reverse-video mode—white characters on a black background; this mode is much harder on the eyes when reading text, but it could be useful for some applications (e.g., graphics displays).

Both displays suffer from the common problem of poor cursor legibility; it is quite difficult to find the blinking underline cursor's location in a screen full of text when editing with WordPerfect or BRIEF, for example. You can improve the cursor's legibility by decreasing the contrast, but this makes the background appear dark and blotchy.

By the way, both computers have controls on the front of the displays to adjust brightness and contrast, but the SupersPort's sliding switches are easier to adjust than the MP-286L's buttons, which must be twisted with a fingertip. The range of adjustments on both machines is broad enough to give good legibility under all normal lighting conditions.

Although the MP-286L's oversize screen allows it to have characters larger than those on other laptops, I did not like the individual letters in its character set. The bottom half of characters such as the *b* and *h* and the top half of the *p*, *g*, and *q* are shorter than the other lowercase letters. This unevenness gives the text a choppy look and makes it hard to read.

Both computers can display colors as shades of gray in a CGA text or graphics mode. By judicious tweaking of the contrast and brightness controls on the two computers, I was able to display between four and six different gray shades, corresponding to sixteen different colors on a CGA monitor. When installing software that has monochrome/color display options, you are safer telling the program that your display is monochrome. The CGA option can too often result in an illegible black-on-black or white-on-white combination. Of course, if the software lets you adjust the color palette, you should be able to find a combination that works well on either laptop.

Keyboards remain a troublesome area for all laptop computers. Three significant problem areas are the location and layout of the function keys and editing keys and how the numeric keypad layout

is accommodated. The manufacturers have resorted to a variety of techniques in keyboard design to keep the size of these computers down. Mitsubishi uses the philosophy of fitting in more keys by making many of the keys smaller. The SupersPort's keys are all full-size and make use of an extra Fn key that acts like a Control, Shift, or Alt key when used in combination with other keys. In this case, Zenith's solution is better, and I found its keyboard far easier to type on and use effectively.

B oth displays suffer from the common problem of poor cursor legibility.

For example, the SupersPort has only four direction-type editing keys, located in the lower right corner of the keyboard. They are arranged in a familiar inverted *T* pattern and work as the arrow cursor keys. When used in combination with the Fn key, they become the Home, End, PageUp, and PageDown keys. Mitsubishi's solution is to put eight separate, half-size keys in the upper right corner of the keyboard, where they are much harder to reach and use productively.

The SupersPort's full-size function keys (F1 through F10) are laid out in a row at the top of the keyboard; the F11 and F12 keys' functions are accessed by using the Fn key in combination with the F1 and F2 keys. The MP-286L has its half-size function keys arranged in two rows at the top of the keyboard, with F1 to F8 in one row and F9 to F12 in the other. This layout is difficult to learn and makes you constantly hunt for the correct function key—a situation that is tiresome with a program like WordPerfect that makes heavy use of the function keys.

Performance Edge
Although these two laptops have similar hardware, the SupersPort 286 has a clear edge when it comes to overall performance. Its advantage is due to two factors: its 12-MHz 80C286 runs with zero wait states when accessing RAM, and its hard disk drive has an average access time (measured with the Coretest program) of 27 milliseconds, versus 75 ms

for the MP-286L's hard disk drive. Mitsubishi includes a disk-caching program from Multisoft, Super PC-Kwik, which compensates somewhat for its slow hard disk drive. These differences show up on the BYTE benchmarks.

Both computers perform well compared to the better desktop 80286 computers; as the benchmark results show, both laptops are faster than the IBM PS/2 Model 50, but they lag behind the newer PS/2 Model 50 Z. Neither machine was equipped with a math coprocessor chip, and this adversely affected the benchmark results.

Words and Service
Both computers come with easy-to-read setup and operating instructions. The SupersPort 286 uses MS-DOS 3.21, and the MP-286L uses MS-DOS 3.3 version 1.06. The MS-DOS manual for the Zenith laptop is larger and more complete than the MP-286L's. The MP-286L comes with GWBASIC but no BASIC manual; Zenith does not include a BASIC interpreter with the standard operating system for the SupersPort. Both laptops can run the Microsoft OS/2 operating system, which is offered as an option by both manufacturers. However, the maximum amount of RAM that the laptops can accommodate (2 megabytes for the SupersPort and 2.6 megabytes for the MP-286L) may not be sufficient to run future versions of OS/2.

Both Mitsubishi and Zenith offer a 1-year limited warranty on their laptops; Mitsubishi guarantees a 7-day turnaround on the MP-286L. You are responsible for getting the damaged goods back to either the manufacturer or one of its authorized representatives. Zenith has an excellent reputation for making quality hardware, and its support policies are also good. Mitsubishi is better known in the U.S. for its well-made audio/visual electronics equipment; if its computer equipment is as good, then the MP-286L should be trouble-free.

Last Keystrokes
The Zenith SupersPort 286 appears to be at the top of the heap of battery-powered 80286 laptops in terms of performance and battery life. Its 12-MHz zero-wait-state processor and the fast hard disk drive constitute a package challenged only by the recently announced Compaq SLT/286 and Toshiba T1600.

Running the SupersPort 286 off a battery pack has some limitations. I found that the SupersPort's 15 or 16 pounds of computer, battery pack, and AC charger/

power supply gets pretty heavy if you have to lug it more than a few hundred feet. Even though I carried the laptop in the convenient and well-made padded nylon case that Zenith sells as an option, it's definitely a chore. In my experience, 10 or 11 pounds is a more reasonable upper limit to comfortably tote around an airport for extended periods of time.

A more serious shortcoming of the SupersPort 286 arose when I tried to use it during a business flight. This computer is simply too long to fit comfortably on a tray table in the coach section of a DC-10. With the computer resting on the tray, I could not open the screen to a comfortable viewing angle because it hit the back of the seat in front of me. If I slid it toward me to alleviate that problem, then the keyboard was jammed too close to my body, and typing was extremely awkward. The 3 inches that the battery pack adds to the overall length of the computer just makes it too long in this situation. Maybe it's a machine designed for the First Class section. Despite these inconveniences, however, it was a pleasure to have such a powerful computer literally at my fingertips while traveling, and this made the problems I experienced seem minor.

The MP-286L's performance would be significantly enhanced if a hard disk drive with a faster average access time were used; the sluggish drive currently installed in the machine is not in keeping with the rest of the system's performance potential. Unless its street price is significantly less than its list price, the Mitsubishi does not represent a particularly good value. [Editor's note: *After this review was written, Mitsubishi announced the MP-286L-240E, featuring a 40-megabyte hard disk drive with better than 30-ms access time and an EGA display.*]

Zenith and Mitsubishi both face serious competition in the arena of AC and battery-powered laptops. There are many contenders in this category, and some have features superior to those of the SupersPort and MP-286L. For example, the Toshiba T3200 has an excellent EGA-compatible plasma display and two standard IBM expansion slots. NEC has recently announced its ProSpeed series of laptops, which also offer an EGA-compatible LCD screen.

The SupersPort and the MP-286L are a far cry from their 8088 and 8086 predecessors. Their displays are clear and bright with superior contrast; they use fast 80286 CPUs and provide internal expandability for at least 2 megabytes of RAM; and they accommodate 40-megabyte hard disk drives. All this has been accomplished without increasing their size or weight significantly.

The SupersPort 286's user interface—the display and keyboard—are better than the MP-286L's, and with its battery pack the SupersPort 286 is a more versatile package. On the other hand, both the SupersPort and the MP-286L suffer from what have become chronic laptop problems: text cursors that are difficult to find and cramped keyboards—compromises that seem unavoidable with this genre of computer. You will have to decide if higher performance, battery power, and a full megabyte of RAM are worth the $1000 difference in list price between comparably equipped Zenith and Mitsubishi computers. ∎

John Unger is a geophysicist for the U.S. government and lives in Hamilton, Virginia. He writes graphics software and uses computers to study the earth's crust. You can reach him on BIX as "junger."

World Radio History

World Radio History

A PS/2 in Channel Only

The Tandy 5000 MC has a hybrid design, offering PS/2 compatibility combined with traditional PC features

Mark L. Van Name

With its delivery of the Tandy 5000 MC, Tandy has become the first vendor to ship a Micro Channel–compatible PC. While such PCs are often referred to as "PS/2 clones," the 5000 MC differs in many ways from IBM's PS/2s.

A Micro Channel Clone

The 5000 MC has an IBM Micro Channel–compatible bus. BYTE supplied me with an MS-DOS-compatible Micro Channel add-on card, an IBM 4-megabyte memory-expansion option—that one worked like a champ. The 5000 MC noted the extra board when it booted and told me to reconfigure the system. I inserted Tandy's reference disk, which contains PS/2-style configuration software from Phoenix Technologies, and rebooted the system. That software also told me of the setup error and asked if I

wanted it to handle the situation for me. I said yes, a few screens went by, and the new board was up and running.

It really was simpler than messing with DIP switches and the usual IBM PC AT setup programs. This experience clearly illustrated some of the nicer points of the autoconfiguring nature of Micro Channel expansion cards.

I also tried an Arnet eight-serial-port card. The 5000 MC noticed the card and asked me to add its drivers to the reference disk. It knew the right names for the drivers, so it was able to read the board's Micro Channel ID. I was unable to test it further, however, because the board did not come with an MS-DOS driver.

Although this testing was admittedly very limited, the 5000 MC appears to work with Micro Channel–compatible cards. Further testing might turn up a few incompatibilities, but Tandy has definitely taken a strong first step into the world of PS/2 compatibility.

The Complete Package

The rest of the 5000 MC is a very reasonable system, with a 20-MHz 80386 at its heart. The 80386 gets help from an Intel 82385 cache controller and a 32K-byte cache of 35-nanosecond static RAM (SRAM). This cache system lets it work with the 5000 MC's 2 megabytes of 100-ns dynamic RAM (DRAM) without wait states over 90 percent of the time.

The basic 5000 MC also includes a socket for a 20-MHz Intel 80387 math coprocessor and a 1.44-megabyte 3½-inch floppy disk drive, much like a PS/2 system. It follows the basic PS/2 design further by including circuitry on the motherboard that supports a DB-9 serial port, a DB-25 parallel port, a 6-pin DIN keyboard connector, a 6-pin DIN mouse connector, a DB-15 analog monitor connector that provides VGA compatibility, and a controller that supports two daisy-chained floppy disk drives.

All these connectors except for the

floppy disk drive controller are accessi-
ble from the back of the system unit. One
particularly nice touch is that each of
these connectors is labeled in raised let-
ters on the case.

Because all this support circuitry is
part of the motherboard, the 5000 MC is
also able to offer five Micro Channel–
compatible expansion slots in a system
with a footprint smaller than an AT's.

This basic 5000 MC package runs
$4999. Of course, to use it you need a
monitor and an operating system, and
you almost certainly will want a hard
disk drive. Tandy offers two packaged
versions with a 40-megabyte or 80-mega-
byte hard disk drive for $6499 or $6999,
respectively. However, neither of these
packages includes a monitor or an oper-
ating system.

My evaluation system was the 80-
megabyte hard disk drive package, with
Tandy's top-of-the-line, VGA-compat-
ible analog monitor and MS-DOS 3.30
added. Total price: $7748.

For comparison, the closest true PS/2
is probably the IBM Model 70-121. That
system's basic configuration differs
from the 5000 MC in three major ways: It
has no cache system, its disk drive is big-
ger (120 megabytes), and it has only
three open Micro Channel slots. The
basic Model 70-121 runs $7995, or about
$1000 more than the 5000 MC. So, if
you want the Micro Channel bus, the
5000 MC clearly offers a price incentive.

In the world of conventional AT-com-
patible PCs, probably the closest compa-
rable system is Compaq's new 386/20E.
Like the PS/2s, the Compaq has many of
its basic features built into the mother-
board. Its standard memory is only 1
megabyte, it contains four AT-compat-
ible slots, and its floppy disk drive is a
5¼-inch unit, but otherwise it and the
5000 MC offer the same basic features.
The 40-megabyte hard disk drive version
of the Compaq 386/20E lists for $6599,
or just $100 more than the 40-megabyte
5000 MC package—but you still have to
buy another megabyte of memory from
Compaq to make the two systems compa-
rable. That makes the 5000 MC a reason-
ably priced alternative to traditional AT-
compatible systems as well, although
there are certainly many 20-MHz sys-
tems out there that are cheaper than
either Tandy's or Compaq's.

Performance?
As you can tell from the accompanying
performance chart, the 5000 MC's per-
formance, like its price, falls somewhere
in the middle of the pack of the 20-MHz
80386-based systems. I compared its

BYTE benchmark results closely with
those of two other systems: the Compaq
386/20 and the IBM PS/2 Model 80-111.

Basically, the 5000 MC loses to the
Compaq, with an overall application in-
dex that is 20 percent slower, but it beats
the Model 80's overall application index
by about 8 percent.

The 5000 MC actually beats out the
386/20 by a hair on the CPU tests, but it
loses all the other tests. Its biggest losses
are in the video and disk tests. Its 11 per-
cent slower video performance is not sur-
prising, since Compaq's VGA system is
one of the industry's fastest. The 5000
MC's large (44 percent slower) loss on
the disk tests is probably due to its ST-
506-compatible disk; the Compaq unit
contained a faster, enhanced-small-de-
vice-interface (ESDI) controller.

The 5000 MC's only significant loss
to the IBM Model 80-111 also came on
the disk tests. Again, it was competing
with a faster ESDI drive, so that loss is
not surprising.

These differences should decrease
when Tandy ships its ESDI controller for
the 5000 MC. While Tandy currently
lists that controller among the system's
options, a Tandy spokesperson said that
the controller was not yet available when
Tandy sent the review system to BYTE.

Compatibility?
The evaluation unit may have lost the
disk race, but it came through like a
champ on software compatibility. It ran
every program I threw at it, including
Borland's Quattro 1.0, Reflex 1.14,
SideKick Plus 1.0, SuperKey 1.16A,
Turbo Basic 1.1, Turbo C 1.0, and
Turbo Pascal 4.0; Digitalk's Small-
talk/V 1.2; Kermit 2.30; a copy-pro-
tected Lotus 1-2-3 version 2.01; Micro-
Pro's WordStar 3.3 and 4.0; Microsoft's
PC Paintbrush 2.0, Windows/386 ver-
sion 2.03, and Word 4.0; the Norton
Utilities 3.00; Quarterdeck's DESQview
2.0, with the Quarterdeck Expanded
Memory Manager 386 version 1.10; and
Symantec's Q&A 1.1.

Tandy is not yet shipping a version of
OS/2 that will work with the 5000 MC's
Micro Channel–compatible bus. Tandy
offers Microsoft OS/2 1.00 for many of
its other PCs, however, and a Tandy
spokesperson said that Tandy would soon
be including support for the 5000 MC.

Going Inside
Opening the unit is fairly easy: You just
remove three Phillips screws on the
back, pull the sides of the cover a bit
apart, and push the cover forward. Still,
continued

Tandy 5000 MC

APPLICATION-LEVEL PERFORMANCE

Tandy 5000 MC **14.3***

WORD PROCESSING

XyWrite III+ 3.52 — Medium/Large

Load (large)	:15
Word count	:02/:16
Search/replace	:04/:20
End of document	:01/:14
Block move	:09/:09
Spelling check	:06/:49

Microsoft Word 4.0

Forward delete	:12

Aldus PageMaker 1.0a

Load document	:15
Change/bold	:24
Align right	:18
Cut 10 pages	:17
Place graphic	:04
Print to file	1:38

☐ **Index:** **2.97**

SPREADSHEET

Lotus 1-2-3 2.01

Block copy	:02
Recalc	:01
Load Monte Carlo	:17
Recalc Monte Carlo	:04
Load rlarge3	:04
Recalc rlarge3	:01
Recalc Goal-seek	:02

Microsoft Excel 2.0

Fill right	:04
Undo fill	1:30
Recalc	:02
Load rlarge3	:22
Recalc rlarge3	:01

☐ **Index:** **3.23**

DATABASE

dBASE III+ 1.1

Copy	1:01
Index	:22
List	1:27
Append	1:37
Delete	:02
Pack	1:23
Count	:19
Sort	1:19

☐ **Index:** **1.50**

SCIENTIFIC/ENGINEERING

AutoCAD 2.52

Load SoftWest	:42
Regen SoftWest	:29
Load StPauls	:10
Regen StPauls	:05
Hide/redraw	9:05

STATA 1.5

Graphics	:19
ANOVA	:12

MathCAD 2.0

IFS 800 pts.	:12
FFT/IFFT 1024 pts.	:12

☐ **Index:** **4.35**

COMPILERS

Microsoft C 5.0

XLisp compile	4:14

Turbo Pascal 4.0

Pascal S compile	:05

☐ **Index:** **2.23**

All times are in minutes:seconds. Indexes show relative performance; for all indexes, an 8-MHz IBM PC AT = 1

Compaq 386/20 **17.9**

IBM PS/2 Model 80-111 **13.2**

IBM PC AT **5**

14.3

- Word Processing
- Spreadsheet
- Database
- Scientific/ Engineering
- Compilers

*Cumulative applications index. Graphs are based on indexes at left and show relative performance.

LOW-LEVEL PERFORMANCE[1]

Tandy 5000 MC

CPU

Matrix	3.39
String Move	
Byte-wide	24.27
Word-wide:	
Odd-bnd.	30.36
Even-bnd.	12.18
Doubleword-wide:	
Odd-bnd.	21.03
Even-bnd.	6.09
Sieve	17.85
Sort	13.46

☐ **Index:** **3.71**

FLOATING POINT

Math	6.60
Error[2]	
Sine(x)	2.09
Error	
e^x	2.35
Error	

☐ **Index:** **7.91**

DISK I/O

Hard Seek[3]

Outer track	3.33
Inner track	3.29
Half platter	6.68
Full platter	10.00
Average	5.83

DOS Seek

1-sector	11.22
32-sector	55.18

File I/O[4]

Seek	0.16
Read	1.25
Write	1.05

1-megabyte

Write	7.74
Read	9.23

☐ **Index:** **1.25**

VIDEO

Text

Mode 0	4.08
Mode 1	4.08
Mode 2	4.39
Mode 3	4.37
Mode 7	N/A

Graphics

CGA:

Mode 4	1.48
Mode 5	1.49
Mode 6	1.65

EGA:

Mode 13	3.50
Mode 14	4.01
Mode 15	N/A
Mode 16	3.95

VGA:

Mode 18	4.14
Mode 19	1.57
Hercules	N/A

☐ **Index:** **2.26**

N/A = Not supported by graphics adapter.

[1] All times are in seconds. Figures were generated using the 8088/8086 and 80386 versions (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	177.74
Livermore Loops[5] (MFLOPS)	0.21
Dhrystone (MS C 5.0) (Dhry/sec)	6281

Compaq 386/20

IBM PS/2 Model 80-111

IBM PC AT

- CPU
- FPU
- Disk I/O
- Video

For a full description of all the benchmarks, see "Introducing the New BYTE Benchmarks," June 1988 BYTE

Photo 1: *The CPU card is visible in the upper left corner. The card contains the 80386, a socket for an 80387 coprocessor, the 82385, and the SRAM cache chips.*

that's more work than it takes to open a PS/2, which is your first hint of just how different the two systems are.

You get the real message when the cover's off. This is not a PS/2 on the inside.

First, forget the PS/2's single motherboard, its slide-on disk connectors, and the nylon snap fasteners that make the IBM machines easy to take apart and reassemble. The 5000 MC is built with traditional PC engineering, with cables snaking everywhere and lots of small screws holding the whole thing together.

Look closer, and the design differences get bigger.

The motherboard is actually only one-third of the heart of the system. It does not contain any of the major processing elements. Instead, all those elements, including the 80386, the 80387 socket, the 82385, and the cache itself, are on a removable CPU card that connects to a proprietary 32-bit slot on the far left side of the system (see photo 1). That board alone contains 32 chips, not counting the 16 SRAM chips. The 80387 socket also has the extra pins to support a Weitek coprocessor, but a Tandy spokesperson said that Tandy does not currently sell the Weitek coprocessor for the 5000 MC.

The system's memory is on yet another board, a Tandy memory-expansion adapter that plugs into another proprietary 32-bit slot on the right side of the motherboard next to the power supply. That board can hold eight single in-line memory modules. In my unit, it had eight 256K-byte SIMMs, giving the system its standard 2 megabytes of memory.

There is a second proprietary 32-bit slot next to this one that can hold another memory-expansion adapter, so you can ramp up memory to 4 megabytes with 256K-byte SIMMs. While Tandy does not currently offer 1-megabyte SIMMs, a Tandy spokesperson said that it plans to do so in the near future. When these larger SIMMs are available, you'll be able to put up to 16 megabytes of memory on the system using these two boards.

Those two slots and the CPU board slot are all dedicated to their particular tasks, and they will not support any other cards. They all use AT-style connectors that look almost fat next to the sleek Micro Channel connectors. Both memory-expansion boards and the main motherboard are made by Tandy.

The five Micro Channel–compatible slots are on the main motherboard. Two of them are 32-bit Micro Channel–compatible slots, while the other three are 16-bit slots. One of the 16-bit slots includes the special high-speed video connector that you'll also find on a PS/2.

The motherboard supports the CPU board, memory board(s), and expansion slots, as well as all the external connectors, with a whopping 150-plus chips. That's more chips than I have seen on any other 80386-based system's motherboard. The motherboard is also oddly shaped, like an *L* with an extra bit hanging from the junction of the two sides.

The motherboard in my unit was a Rev. A model, and it showed its youth. I counted 25 wire traces, including several that were over a foot long, that marked circuit design changes.

The motherboard does the bulk of the work of supporting the Micro Channel–style bus with the five-chip Intel Micro Channel chip set. It also contains the ROM BIOS, which is Phoenix's 80386 Advanced ROM BIOS version 1.01.02. Like most high-speed 80386 systems, the 5000 MC copies the ROM BIOS into RAM at boot time for faster access.

Finally, in total defiance of the PS/2 autoconfiguration philosophy, the motherboard contains two banks of eight DIP switches each and eight jumpers. The DIP switches and some of the jumpers control the system's memory settings, and you have to adjust them if you add more memory on one of Tandy's memory-expansion adapters.

Diverging Design

Of course, it's not always bad to be different from the PS/2s. Most PS/2s have room for only one hard disk drive. The 5000 MC, on the other hand, has one 3½-inch half-height drive area and a 5¼-inch drive bay that can hold either one full-height drive or two half-height drives.

In my unit, the 3½-inch area held an 80-megabyte hard disk drive from Rigidyne, a subsidiary of Control Data Corp. An Adaptec ST-506 controller in one of the expansion slots manages that drive. Tandy claims an average access time of 16 milliseconds for the drive, but, based on its performance on BYTE's disk tests, either it or the controller is really falling down on the job. Based on what I know of CDC drives, my guess is that the Adaptec controller is guilty, but I could not confirm that guess.

Tandy also offers a wealth of options to go into the drive areas, including hard disk drives of up to 344 megabytes, several models of tape drives, and 5¼-inch floppy disk drives.

Skin Deep

When you put the 5000 MC back together and turn it on, you can forget the engineering inside and pretend again that it's a PS/2.

Its 101-key keyboard follows the IBM Enhanced layout. The keyboard has a very light, springy touch, with a good audible keyclick.

Tandy included its VGM-300 analog color monitor with the review unit. The display was very crisp and clear, although its standard font is a bit unusual.

While you're looking at the monitor you can't help but notice a few things on the front of the unit. One is the big FCC Class A 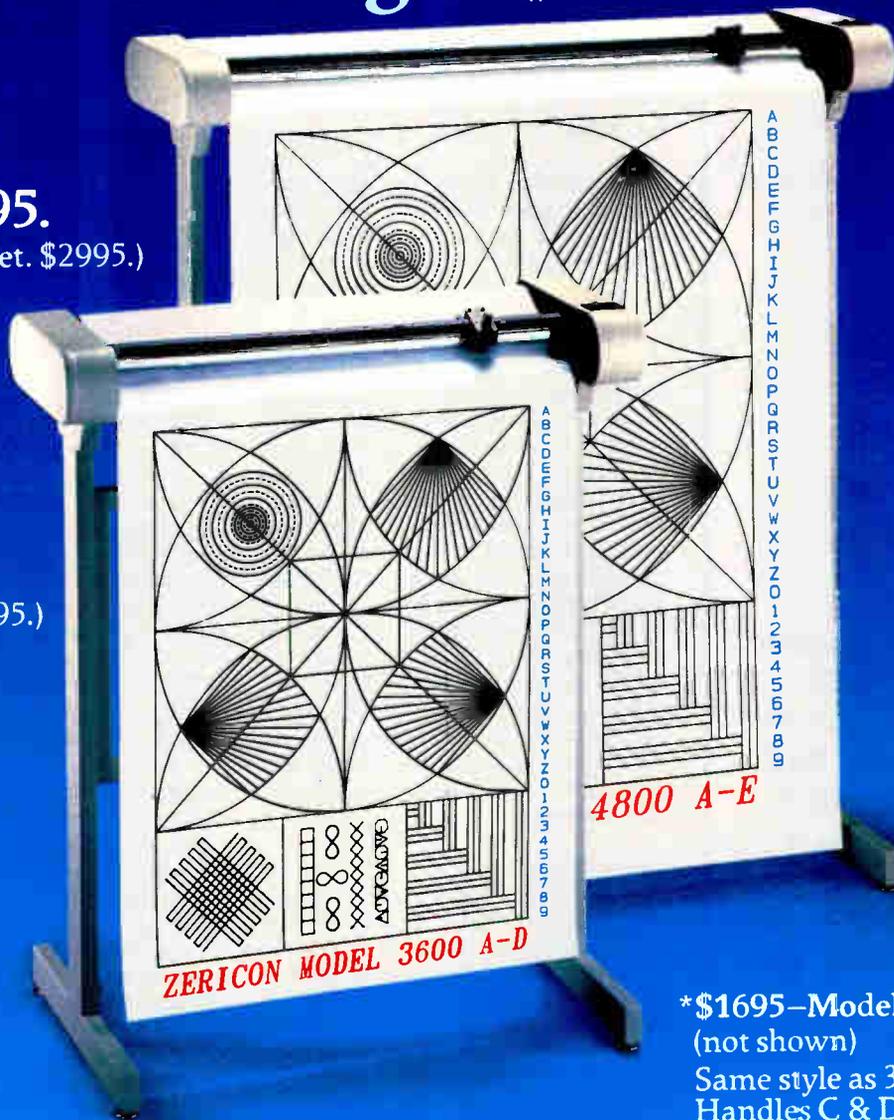sticker that proudly proclaims the 5000 MC to be a "business computer." Just above it is a keylock that locks the case, not the keyboard. More interesting are the reset and power switches that are on the left and right sides, respectively. The red reset switch is recessed so that you are not likely to push it accidentally.

The oddly shaped, green power switch is another matter. It is only half an inch to the right of the 3½-inch floppy disk drive, so it is just aching for an accidental push. It is also one of the most peculiar pieces of engineering I have seen in a while: Inside, it proves to be the tip of a 9-inch bar of the same green plastic that stretches from the front of the unit to a mechanical switch on the rear of the power supply at the back of the unit. A thin cable snakes through the bar from the motherboard to the tiny indicator light in the tip of the switch. Cute as the switch is, it's dangerous to put a power switch that looks suspiciously like an eject button right next to a floppy disk drive, and Tandy should move it in future systems.

The Soft Stuff
The only standard software is Tandy's reference disk. This disk contains programs that autoconfigure the system, do low-level disk formats, give information about power-on self test (POST) error messages, and let you access a few of the system's more unusual features.

These features include passwords for power-on and keyboard usage, as well as for the 5000 MC's use as a network server. You can also set both the keyboard's repeat rate and its speed of responding to key depressions to a faster speed than normal.

One common software feature of most 80386-based systems that the 5000 MC lacks is the ability to set the CPU to a slower compatibility speed. A Tandy spokesperson said that the 5000 MC is based on the IBM PS/2 Model 80, and that it does not include a way to slow the CPU because the Model 80 does not. I had no problem with my copy-protected Lotus 1-2-3 system disk, which often requires a slower speed when loading the program, but the lack of this ability could be a problem for some older programs.

The Tandy 5000 MC Installation and Operation Manual explains all the functions of the reference disk software. It's a small, 42-page manual, but it is readable, includes an index, and takes you

from unpacking through memory-configuration DIP switch settings.

Of course, sometimes things go wrong that you can't solve with any amount of documentation. That's when a company's technical-support system comes into play.

If you have questions, you can call any local Tandy computer center or computer service and support center. I tried a few of the ones in my area, but because the 5000 MC is just barely shipping now, they really couldn't help me with anything specific to its Micro Channel–compatible features. They were, however, helpful with a few of the usual MS-DOS problems that I threw at them.

A 1-year parts-and-labor warranty comes with the 5000 MC. While Tandy will repair your unit only at one of its roughly 155 service and support centers, you can take the system to any Tandy store, even a Radio Shack, and the company will get the unit to and from the service center at no charge to you.

You can also pay an additional fee for on-site service during that first year. The price depends on what you've got on your unit and how close you are to one of Tandy's service and support centers. For example, if you're within 50 miles of a center and you have a 40-megabyte hard disk drive in your 5000 MC, the first year of on-site service costs $330.

Tandy also sells additional years of service, either on-site or carry-in, for fees that, again, depend on your system's configuration and your distance from a service center.

Do You Need a PS/2?
All engineering complaints aside, the 5000 MC is a reasonable Micro Channel–compatible alternative to IBM's PS/2 machines. It's a first-generation Micro Channel clone, so you have to expect some rough spots, and there are bound to be some incompatibilities that my testing did not uncover. Still, the 5000 MC is a good first step, and Tandy has already proved that it is going to be a player in the PC world for some time.

The big question is whether you're content to stick with the AT-style bus, or if you want to move to the Micro Channel. If you've decided that the Micro Channel bus is for you, the 5000 MC offers an interesting alternative that is both cheaper and faster than the comparable IBM 20-MHz 80386-based PS/2s. ∎

Mark L. Van Name is a freelance writer and computer consultant living in Durham, North Carolina. He can be reached on BIX c/o "editors."

World Radio History

Three Assemblers for MS-DOS

TASM and OPTASM challenge the long-dominant MASM

Michael Blaszczak

Programmers have long known that a computer's native assembly language is the best choice for the fastest code. Microsoft's Macro Assembler, or MASM ($150), has been the definitive assembler for IBM PCs. Now there are two challengers: Borland's Turbo Assembler, or TASM ($150), and SLR Systems' OPTASM, for optimizing assembler ($125). Each of these three assemblers has unique strengths. MASM is the one the others imitate, and it's the only one that supports OS/2 development. TASM comes closest to matching MASM on a feature-by-feature basis, runs faster, and bends over backward to make programming in assembly language easier. OPTASM runs faster yet and can optimize certain troublesome assembly language constructs. They're all worthy tools—I used each to assemble and link working versions of three programs that comprise thousands of lines of code. In this review, I'll try to help you figure out which assembler best suits your needs.

An assembly language development system includes more than just an assembler. You've got to have a linker, too, and you'd like to have a debugger, a librarian, a make facility, and comprehensive documentation. Assembly language programming is tricky, and a good debugger is worth its weight in gold—especially when you're interfacing assembly with a high-level language.

Both MASM and TASM come with top-notch symbolic debuggers that enable you to view and manipulate your program in terms of source-level symbols and labels, as well as full complements of other support tools. OPTASM is a bare-bones package by comparison. There's no debugger or linker, and you have to spend an extra $50 to get the optional librarian. SLR Systems does, however, plan to provide more tools in a forthcoming release of OPTASM. See table 1 for a comparison of the features that the three packages offer.

Microsoft's MASM

MASM has evolved considerably during its long life (long by microcomputing standards, that is). MASM 5.1 is the newest version of the product. The package includes the OS/2-capable MASM assembler; CodeView, Microsoft's full-screen symbolic debugger; and the Microsoft linker, librarian, text editor, and make tool.

MASM comes with ample documentation. A reference manual describes the assembler itself, and a particularly useful update outlines the differences between the current MASM and its previous incarnations. A utilities manual details the use of CodeView and the other tools supplied with MASM. The Mixed Language Programming Guide tells you how to interface assembly modules with Microsoft C, Pascal, FORTRAN, and BASIC programs. Microsoft also includes a 148-page reference booklet that enumerates the instruction set, discusses each instruction's side effects, and notes how many clock cycles each instruction consumes.

The quality of the documentation is generally quite good, but I have a few complaints. The manuals are full of examples, but they're usually fragmentary—few are complete and ready to run. The descriptions of instructions that have a rich variety of addressing modes, like MOV and JMP, get short shrift; there's no elaboration on how to declare labels or

data to be addressed by each mode. There is a finished program included with the package—SHOWR.ASM, a file-listing tool—but none of the documents mentions it. In some cases I thought the documentation failed to describe an important technique. The Mixed Language Programming Guide, for example, doesn't tell you how to define data areas and variables in an assembly module for later reference in a high-level-language program. And while the discussion of the LES and LDS instructions plainly shows how to use them to load a doubleword (representing a far pointer) from memory into a pair of registers, it doesn't help you figure out how to set up a doubleword in memory for use in this context.

Putting MASM to Work

MASM was the slowest of the three assemblers. TASM was about twice as fast, and OPTASM about three times as fast (see tables 2, 3, and 4). And MASM was the fussiest—that is, it complained the most. MASM uses a two-pass method to assemble a module. On the first pass, it reads the source file and generates code. On the second pass, it rereads the source file and works on the code it produced to resolve the addresses of symbols and the sizes of op codes.

That's a good strategy, but MASM runs into some problems. For example, here's a line from MS-Kermit, one of the three test programs I used:

```
mov    bx,seg serdum
```

This instruction moves the segment part of the address of serdum, a forward reference to a label. MASM's second pass produced a so-called "phase error" that, after some head-scratching, I eventually resolved by expressing the value I wanted to load into the BX register in a different and less mnemonic way:

```
mov    bx,code
```

continued

Table 1: *A list of MASM, TASM, and OPTASM features.*

	MASM	TASM	OPTASM
Installation program	Yes	Yes	No
Macros	Yes	Yes	Yes
Optimization	No	Some	Yes
Simplified segments	Yes	Yes	Yes
Warning level	Three levels	All switchable	On/off
Editor	Yes	No	No
OS/2-compatible	Yes	No	No
Tutorial	No	Yes	No
80386/80387 support	Yes	Yes	No
Local labels in PROC	No	Yes	Yes
Microsoft HLL interface	Yes	Yes	No
Debugger	Yes	Yes	No
Linker	Yes	Yes	No
Overlay linker	Yes	No	No
Librarian	Yes	Yes	No
.OBJ utilities	None	One	None
.EXE utilities	Three	No	No
MAKE utility	Yes	Yes	Yes[1]
Cross reference	Yes	Yes	Yes[2]
On-line help	No	No	Yes
On-disk examples	One	Many	Two[3]

[1] OPTASM's MAKE utility is nonstandard and integrated with the assembler.
[2] OPTASM's CREF utility is integral to the assembler.
[3] One of OPTASM's two example files highlights the problems with MASM solved by OPTASM.

Table 2: *Assembly time for NJRAMD, a RAM disk program.*

	Assembly time (seconds)	.OBJ size (bytes)	.SYS size (bytes)
OPTASM	1.7	3758	1897
TASM	2.6	4673	1899
MASM	5.6	4379	1899

Table 3: *Assembly time for MS-Kermit.*

	Assembly time (seconds)	.OBJ size (bytes)	.EXE size (bytes)
OPTASM	45.4	101,581	86,032
TASM	67.6	171,336	86,157
MASM	165.7	133,780	86,160

Table 4: *Assembly time for the Greenleaf communications library.*

	Assembly time (seconds)	.OBJ size (bytes)
OPTASM	26.9	5522
TASM	51.6	8058
MASM	81.7	6122

The message associated with the phase error wasn't too helpful, as it occurred several lines after the instruction in question. In the end I had to use MASM's /D command-line option to produce a first-pass listing, then study that for a while to figure out what was happening.

MASM will also pad your code with extra NOP instructions in a variety of situations. This happens when you use a JMP instruction to reach a label that lies within a 128-byte range, or when, as shown in listing 1, you use a test instruction with a forward reference and immediate data.

One of the biggest improvements in this version of MASM is the support for mixed-language programming. It's always tough to remember the protocols you need to follow to build assembly subroutines that you can call from C or Pascal. MASM now provides new simplified segment declaration directives and extensions of standard directives that greatly simplify the proper declaration of segments, groups, and public labels.

Support for the 80386 processor is another major improvement. With the .386 directive, you can use MASM to create DOS programs that will take advantage of the extra registers, addressing modes, and instructions of the 80386.

Finally, MASM and LINK are now dual-mode executables that will run under both DOS and OS/2. They can be used to create other dual-mode programs or to create programs that exploit the protected modes of the 80286 and 80386.

The MASM Toolkit

Much has been said and written about CodeView—it's an excellent debugger—so I'll just mention a couple of favorite features. I'm particularly impressed by its ability to view structured data, following pointers as necessary. Code-View's support for the Intel numeric coprocessors is another strong point—you can dump the status of the coprocessor to the screen during debugging. That, in conjunction with MASM's ability to assemble all the special math op codes, has saved me a great deal of trouble on numerically intensive projects. CodeView can use expanded memory to keep symbolic information out of the way of the executing program. Serious professional users may want to consider buying a hardware debugger, but for my purposes CodeView is entirely sufficient.

Microsoft's linker is slower than TLINK, the Borland linker, but provides good support for program overlays. It's relatively straightforward to create overlay sections. You do need an overlay han-

	Microsoft Macro Assembler 5.1	Borland Turbo Assembler 1.0	OPTASM 1.5
Type	Macro-based assembler	Macro-based assembler	Optimizing macro-based assembler
Company	Microsoft Corp. 16011 Northeast 36th Way P.O. Box 97017 Redmond, WA 98073 (206) 882-8088	Borland International 1800 Green Hills Rd. P.O. Box 660001 Scotts Valley, CA 95066 (800) 543-7543 (408) 438-8400	SLR Systems 1622 North Main St. Butler, PA 16001 (412) 282-0864
Format	Four double-sided, double-density 5¼-inch floppy disks	Three double-sided, double-density 5¼-inch floppy disks	Three double-sided, double-density 5¼-inch floppy disks
Language	C	C	Assembly
Hardware Needed	IBM PC or compatible with 128K bytes of available memory; a hard disk drive is recommended	IBM PC or compatible with 256K bytes of available memory; a hard disk drive is recommended	IBM PC or compatible with 128K bytes of available memory; a hard disk drive is recommended
Software Needed	PC-DOS or MS-DOS 2.00 or higher	PC-DOS or MS-DOS 2.00 or higher	PC-DOS or MS-DOS 2.00 or higher
Documentation	123-page reference for the MASM 5.1 update; 467-page Programmer's Guide for MASM 5.0; 139-page Mixed Language Programming Guide; 402-page CodeView and Utilities reference; 148-page reference booklet	296-page Reference Guide; 582-page User's Guide	323-page reference for OPTASM; 56-page booklet for OPTLIB
Price	$150	$150	$125 for OPTASM $50 for OPTLIB
	Inquiry 1055.	**Inquiry 1056.**	**Inquiry 1057.**

dler, but, although none is included with the assembler, you can add the overlay handler from any Microsoft high-level-language run-time library to your executable code at link time.

The Microsoft MAKE program mimics Unix make closely. It's small, so you don't run into trouble executing other tools from within it. It features a well-developed macro facility. And there's a suite of tools—EXEPACK, CVPACK, and EXEMOD—that squeeze executables and modify program headers.

MASM takes some getting used to, but it gets the job done. MASM does have more than its fair share of frustrating quirks and oddities, but it's always considered to be the standard for general-purpose work. And if you're doing OS/2 development, it's the only game in town.

SLR Systems' OPTASM
OPTASM is a lean and mean package aimed squarely at seasoned program-

mers. It comes with one concise reference book that covers all the bases. As the name implies, OPTASM is an optimizing macro assembler. The program makes all the passes needed to resolve addresses of symbols and op-code sizes.

OPTASM can pick the most efficient version of the JMP op code automatically, and it will even expand conditional branches that are out of range into the appropriate instructions to make the jump possible. Given listing 1, for example, OPTASM would rewrite the jnc instruction—which specifies a jump to an out-of-range label—like this:

```
          clc
          jnc       locally
          jmp       nearly
locally:
          REPT      150
          nop
          ENDM
nearly:
```

This frees the programmer from the task of keeping track of ranges when coding conditional jumps. OPTASM also optimizes far jump and far call instructions and provides directives that allow popping and pushing more than one register on a line of code. These optimizations seem trivial, but the net effect is a big saving of effort, particularly when developing larger modules.

The manual provides descriptions of simplified segment declarations that are much clearer than Microsoft's. The manual also takes greater care in explaining the physical loading order of segments in a multisegment program and clearly explains how the ALIGN, COMBINE, and CLASS keywords control the behavior of LINK.

OPTASM is generally compatible with version 5.0 of MASM but omits several features that were added to version 5.1. MASM's enhanced .MODEL directive,

continued

Listing 1: *Only TASM handles this code cleanly.*

```
; Microsoft and Borland allow the programmer
; to define the high-level language this module
; will be used with. OPTASM does not.

.MODEL   small,c
.CODE

; Microsoft and Borland allow the specification
; of high-level-language function parameters
; in the PROC line. OPTASM doesn't.

sample  PROC      parm1:word, parm2:word

; Microsoft and Borland permit you to
; declare local variables. OPTASM doesn't.

LOCAL    var1:word,var2:word
mov      ax,35       ; mov 35 into
mov      var1,ax     ; load var1
clc
jnc      faraway     ; We can't conditionally jump to a
                     ; label that is more than 128 bytes
                     ; away. OPTASM reconstructs this
                     ; with functionally equivalent code.
                     ; TASM and MASM produce errors.
REPT     150
nop
ENDM

; This construct causes MASM to
; generate a phase error.

faraway:
         test    forwardref,3

; MASM generates an extra NOP here.
; TASM and OPTASM don't.

         mov     cx,messagelen
         ret
sample  ENDP
         .DATA
forwardref      db        01234h
message         db        "This is a string.",13,10
messagelen      equ       this byte - message
         END
```

for example, lets you name the language from which you'll be calling your assembly routine; OPTASM complains about the .MODEL directive in listing 1 for this reason. Nor does OPTASM support MASM's improved PROC and LOCAL directives; PROC declares parameters that the assembly routine expects to find on the stack, and LOCAL declares and symbolically addresses local variables. These incompatibilities aren't a problem if you start a project from scratch with OPTASM, but they'll complicate your life if you're porting code from MASM to OPTASM.

OPTASM comes with an on-line help system that delivers information about the assembler and the Intel op codes. It's

handy, but its 100K-byte memory requirement seems excessive. There's a CONFIG program that you can use to set OPTASM's defaults; alternately, you can use environment variables to achieve the same effect. The MAKE facility is integral with the assembler, not a separate tool. It uses nonstandard make files that aren't compatible with other MAKE programs. There's no debugger, although you can use the /ZD option to get OPTASM to insert line-number information for use in conjunction with CodeView and other debuggers. There's no linker, and the librarian, OPTLIB, costs an extra $50.

OPTASM's big draws are its speed and the compactness of the code it produces.

No question about it, OPTASM trounces MASM and TASM in both categories. Serious assembly language programmers working on projects that don't require MASM's newest features—for example, the .386 directive and the enhanced .MODEL, PROC, and LOCAL directives—will definitely want to look at this package.

Borland's TASM

Borland's TASM now ships as part of the Borland professional packages. Turbo C 2.0 contains release 1.0 of the assembler and version 2.0 of Borland's linker. Also included are a librarian and several other utilities, as well as the Turbo Debugger. Turbo Pascal 5.0 contains the same programs. You can buy just the assembler and debugger separately for $150.

Unlike MASM and OPTASM, TASM comes with a strong tutorial on assembly programming that programmers more familiar with languages like Pascal and C will find extremely useful. But I found the reference material skimpy—the discussion of STRUCT and UNION directives, for example, was quite inadequate. And Borland's softbound manuals are troublesome; you can't leave one open to a particular page without using an extra hand or a paperweight. The sample programs, however, are excellent. The distribution disks contain copies of all the examples in the documentation, and there are also three complete utility programs ready to assemble, link, and run.

Borland has gone to great lengths to make TASM compatible with MASM. There are two directives that govern MASM compatibility. The MASM51 directive provides full emulation of version 5.1 of Microsoft's product. And there's even a QUIRKS directive that causes TASM to mimic some of MASM's infamous peculiarities.

The Turbo Debugger matches Code-View feature for feature and, although it doesn't support mice, implements an intuitive point-and-shoot interface. There's a useful Get Info command that maps DOS memory, expanded memory, and any interrupts the target program may have snatched. And if you're debugging 8086 software on an 80386 machine that has 700K bytes of extended memory, you can use the "virtual debugger," TD386, to create a virtual 8086 on which to debug your program. The resulting environment is close to what you get with a hardware debugger.

TASM runs faster than MASM but slower than OPTASM. But—and here's the most important point—TASM alone cleanly assembled everything I fed it. In

a sense, TASM gives you the best of both worlds: MASM compatibility without MASM's glitches.

TASM Extras
Borland has also made an effort to rationalize MASM's syntax. The .IDEAL directive invokes "ideal mode." If you use this mode, you trade away MASM compatibility for improved readability—particularly with respect to addressing modes. For example, the MASM construct

```
mov    ax,4[bx][si]
```

(which looks like multiplication but isn't) is illegal in TASM's ideal mode, which instead requires the more intelligible

```
mov    ax,[4+bx+si]
```

The improvement is considerable and will certainly help beginners. The benefits of some of the other constructs required by ideal mode are less clear-cut. For example, ideal mode requires that you put names before keywords, so that the MASM construct

```
PROC    myproc    near
```

becomes

```
myproc    PROC    near
```

which, in my opinion, isn't a big win.

The nicest thing about ideal mode is that it lets you use the same identifier in more than one structure. There are some drawbacks. Because symbols can't start with a period in ideal mode, you have to relearn certain MASM directives—for example, .286 and .XALL become P286N and %NOMACS. And unlike MASM, TASM's ideal mode requires you to explicitly quote strings used in conjunction with directives.

The Borland toolkit includes TLINK, the linker; TLIB, an object-module librarian; TCREF, a cross-referencing utility; and MAKE. Borland's MAKE is slightly more powerful than Microsoft's because it implements predefined macros like include, IF, and ENDIF that improve the programmability of the MAKE facility. Borland also includes a wonderful program called OBJXREF. It can read a group of object files and list the public names that each one declares, as well as the external labels that each module relies on. Like OPTASM, TASM doesn't include utilities that manipulate .EXE files.

TASM is the recommended choice for programmers with little assembly experience. It's geared for use with Turbo C and Turbo Pascal but doesn't require those products. You can create standalone applications with TASM, and you can interface TASM modules to programs written in Microsoft high-level languages without any trouble.

Putting the Packages through Their Paces
To test the packages, I used each to assemble three different sets of assembly language routines: my own NJRAMD, a RAM disk program; MS-Kermit, a 25,000-line assembly version of the popular communications program; and the Greenleaf communications library. NJRAMD makes extensive use of conditional directives and structures; the other two programs include header files that perform lots of symbol equates and define a number of macros.

I ran the tests on an 8-MHz PC's Limited 286 with 640K bytes of memory, using a RAM disk. As the results show, OPTASM runs fastest and produces the smallest .OBJ files. TASM comes in second in terms of speed, and MASM takes second place in .OBJ size. Most of the code-size differences vanish, though, when you look at the sizes of the linked .SYS and .EXE files. OPTASM writes a highly compact set of .OBJ records, but the difference between it and the others in terms of the actual quantity of executable code, while perceptible, is quite small.

While my benchmarks showed TASM to be somewhere between MASM and OPTASM for performance, it offers a nice combination of MASM support and OPTASM's assembly language optimization skills. TASM provides conditional jump extensions and doesn't trip up on source code that MASM is not able to handle.

All three assemblers do the job; none is perfect. Choose the one that's right for you. If you're doing OS/2 development, your choice will obviously be MASM. If raw speed is critical, you don't need OS/2 support or guaranteed MASM compatibility, and you've got the necessary support tools, use OPTASM. My personal favorite is TASM. For my money, it's got the best mix of speed, MASM compatibility, and overall ease of use. ■

Michael Blaszczak, a University of Hartford student, provides technical support for NWI of East Hartford, Connecticut. He can be reached on BIX c/o "editors."

PROTECTION FROM THE FAMILY THAT MEANS BUSINESS

THE PROLOK™ FAMILY.

ROMLOK™ is the intelligent software protection device that attaches to and runs transparently from a computer's parallel port. It allows the user to make unlimited backup copies, and run freely from a hard disk (no installs required). Anything that can be done with unprotected software can be done with ROMLOK—*with one exception*— the free distribution of your program to other users.

Whether You're a Programmer or Not

ROMLOK is the first hardware device that doesn't require programming to get up and running. The ROMLOK protection utility will produce a protected copy of your software in less than 30 seconds.

No Dumb Dongle

ROMLOK has its own microprocessor and memory. ROMLOK memory can store up to 16 bytes of crucial program data. To capitalize on ROMLOK's many unique features, linkable object code is available, so you can customize your application.

From the Protection Originator

Vault offers a complete line of protection and security products. From inexpensive SURELOK™ protection, to patented disk based PROLOK™ protection, to the new ROMLOK-P™, *and* the even more secure ROMLOK-S™ (a non-repeating encryption tool), Vault has the answer to your protection needs. *Call the Originator, not the Imitator.*

VAULT CORPORATION

Contact VAULT CORPORATION at
2192 Anchor Court, Unit C, Newbury Park, CA 91320
805-499-5978
Telex: 805-499-5396 VAULT US • Fax: 805-498-6046
800-445-0193 (Nationwide) • 800-821-8638 (California)

ROMLOK™

Circle 238 on Reader Service Card

World Radio History

Full Impact

New bells and whistles
in Ashton-Tate's
spreadsheet
for the Mac

Diana Gabaldon

Full Impact is not, as you might think, the latest Dirty Harry film. It's Ashton-Tate's Macintosh spreadsheet package, the latest entry in a line of Mac office software, meant for those who like their balance sheets with bells on.

Full Impact sells for $395 and runs on a Mac Plus, SE, or II with 1 megabyte of memory and two 800K-byte floppy disk drives, but to use it effectively you need 2 megabytes of memory, one floppy disk drive, and a hard disk drive. It works with System 4.1 and Finder 5.5 or higher.

Like FullWrite, Ashton-Tate's word processor/desktop publisher reviewed in the November 1988 BYTE, Full Impact is heavily loaded with presentation-graphics features and is excellent in some performance areas.

One of the most striking features is Full Impact's drawing capability, which allows you to draw shapes, add graphs and charts to them, and store everything with the spreadsheet in a single file. While Microsoft's Excel allows you to add charts and text to a spreadsheet, its drawing capability is not as flexible, and you must store graphics separately from the spreadsheet.

Full Impact's drawing display is actually a "layer" that lies over the spreadsheet itself. The layer containing the graphics can be transparent or opaque, letting the spreadsheet show through the view of charts and graphics or blocking it out. Graphics elements can also be easily adjusted to overlay each other in various combinations; in other words, you can place an oval on top of a rectangle, then move the oval behind the rectangle, then move it up front again. In addition to ovals and rectangles, you can draw rounded-corner rectangles, lines, underlines and outlines for single cells or blocks, and all the standard types of charts: pie, bar, stacked-bar, line, area, scatter, and high-low-close.

Easy Interface

Aside from its presentation features, Full Impact's major attraction is its ease of use. Instead of the usual spreadsheet Worksheet Global...Quit menu across the top of the screen, Full Impact, taking full advantage of the Mac interface, has an icon bar in which icons represent some of the more commonly used features, such as "fill down," "fill across," and "define range." The icon bar is really two alternating icon bars, which provide a wide range of functions. You can also customize the bar to allow for using your favorite functions via icon.

This interface is another major difference between Full Impact and Excel, which has the standard type of spreadsheet menu. While Excel also allows you to add graphics to a spreadsheet, they are shown in separate, smaller windows that overlay the spreadsheet. In Full Impact, the graphics elements are shown in the spreadsheet itself, in the actual positions where they will appear.

You can add text to a spreadsheet just as easily as graphics. Define an area where you want the text to go, and once defined, this area can be moved, resized, or pasted wherever you want it, no matter what text it contains. And within the text window, you have word-wrapped text, with the usual Mac fonts, type sizes, and type styles available.

You can reduce and enlarge a spreadsheet to preview a formatted page, but this is a rather tedious business, as you must keep reducing over and over, a bit at a time, to arrive at a view of a whole page. Page formatting is supported to the extent of standard margins, headers, and footers, but not much more.

One difficulty in running any spreadsheet package on a Mac is, of course, the restricted screen area available. By the time you take three lines off the top of the screen for the main menu bar, icon bar, and data-entry bar, and another line off the bottom and a space at each side for neat appearance, you're left with the electronic equivalent of a 3- by 5-inch index card to work in. Full Impact remedies this limitation to some extent by allowing you to have up to eight views of a spreadsheet open concurrently. This is the equivalent of having the spreadsheet visible in eight separate windows, except that the windows overlay each other, rather than being tiled.

Full Impact provides one refinement in spacing. While most spreadsheets let you adjust column width, this one also lets you adjust row height, so you can expand any cell in both dimensions.

Full Functionality

Full Impact has all the standard spreadsheet functions and a few extras, as well. You can, of course, enter date and time as ID markers or "time stamps." You can also use them as values, to calculate items such as payment dates or peak sales periods. There is a NUM2C function that converts a numeric value to a text string so that you can insert it in a macro that uses text and numbers together. As a nice touch, you can select discontinuous areas of the spreadsheet. Also, Full Impact allows you to cut and paste spreadsheet areas to the Clipboard for easy transfer, while Excel does not.

Functions include Math, Logic, String, Database, Statistics, Time, and Spreadsheet submenus. The Math submenu includes trigonometric and log
continued

functions as well as all the standard math functions such as exponentiation and square root. The Spreadsheet submenu has functions that locate values in the spreadsheet and evaluate certain attributes. The Horizontal Index Lookup, for example, searches from left to right across the first row of a specified range until it finds a specified value. Then it gives the count of the cells searched.

The software also supports a sophisticated macro language. You can use Boolean logic to set conditions in macros, and you can construct "smart" macros—those capable of taking alternative actions, depending on the conditions encountered. Macros can also pass parameters to functions called by the macro. There is an Undo function that can be called from a macro and elsewhere in a spreadsheet, but Undo cannot be recorded as part of a macro.

You can set the recalculation for manual, regular automatic recalculation, and automatic recalculation at intervals you choose. You can also set an auto-save option, which automatically saves the spreadsheet and data at intervals you specify. You can protect individual cells and ranges against accidental change, but unlike Excel, Full Impact doesn't let you build a protected spreadsheet. In Excel, you can lock some cells and leave others

"open" for change. Both programs support password protection for files.

Flexible Operation
There is more than one way to do most things in Full Impact. For example, you can enter information in cells using the keyboard, the icon bar, or a combination of both. I found the icon bar very annoying to start with, but once I got used to it, using the icon bar was often faster than using the keyboard. I still prefer to enter formulas using the keyboard, but some functions, such as range definition, work very nicely with icons. This flexibility adds greatly to the package's friendliness and ease of use.

You can import information into a spreadsheet from other software, such as Microsoft Multiplan or Excel, or even from such DOS programs as Lotus 1-2-3 and the various versions of dBASE—provided you have suitable communications hardware and software for linking your DOS machine with your Mac. I tried transferring a small dBASE III file, using LapLink Mac, and found that it worked fine, although the transfer was slow. It also successfully imported a file (but not the macros) created by Lotus 1-2-3 version 1.0.

Imported spreadsheets are limited to 2048 rows by 256 columns. When you import records from a database such as dBASE, this limitation means that you can import a maximum of 2047 records, since the field names from the database file occupy the first row. This size limitation is one of the most important differences between Full Impact and Excel. While a Full Impact spreadsheet can have no more than 2048 rows (whether imported or generated by Full Impact), Excel allows more than 16,000 rows.

As for export, you can export Full Impact spreadsheet data to other formats, including Lotus 1-2-3 version 1.0/1a/2.0, but you can export the graphics part of a spreadsheet only via the Clipboard, which essentially means you cannot export to a DOS system.

Spreadsheets, in general, don't make particularly good database managers, and Full Impact is no exception. It's necessary to define and name a range both for the actual data cells in the database and for a separately constructed range of criteria cells and to define (though not to name) an output range to which the output records may be written. If you put the output range too close to cells that already contain data, the existing data will be overwritten by the selected records. It took me three tries to get the example given in the manual to retrieve records

from a sample database correctly; the error message kept telling me I had not defined a range correctly but wouldn't tell me which range.

On the other hand, you can do a few special things with the database capabilities of Full Impact. You can, for example, specify more than one criterion for a single cell, such as selecting records for cars with a price greater than $20,000 and cars with a price less than or equal to $17,000. Likewise, you can define your output range as a single cell, and Full Impact will automatically adjust the range to accommodate all selected records. However, if these records extend over adjacent cells containing data, the adjacent cells will still be overwritten. So for actual retrieval of selected records, Full Impact performs similarly to Excel on files of similar size—but Excel handles much larger files.

Getting Help
The documentation is very polished and generally well-written and well-organized. The on-line help is good. This, too, has been "Mac-intized"; you get help by transforming the mouse pointer into a "help" pointer by pressing Command-?, then pointing to the icon you require help with before clicking with the help pointer.

As with all other Ashton-Tate products, a new Full Impact user is allowed one phone call prior to registration. Thereafter, Full Impact comes with Ashton-Tate's standard 90-day free support package. Support after 90 days must be paid for, at the rate of $50 for 15 calls in a year. A monthly newsletter, "Random Lines," is available at $18 per year, and additional support is available via the Ashton-Tate Forum on CompuServe.

Support calls are not toll-free but are answered from 6 a.m. to 4:30 p.m., Pacific Time. I called twice in the course of writing this review, and after 15 minutes of listening to Jim Croce music, I gave up and left a message to be called back. My call was returned within the stipulated 24 hours, and once I had a technician on the line, I found him friendly, helpful, and knowledgeable.

Impact on Performance
The most serious drawback to Full Impact is a defect it shares with FullWrite: It's a memory hog. While it theoretically requires only 1 megabyte of memory and two floppy disk drives to run, Full Impact actually needs 2 megabytes and a hard disk drive to run effectively. Even a small (35 by 10) spreadsheet ran out of
continued

disk space while inserting and deleting rows, and the Undo function failed to work when managing larger spreadsheets. Consolidating linked spreadsheets also runs you into trouble when you have only 1 megabyte. With more than two or three linked spreadsheets, consolidation slows to an unbearably snail-like pace.

In terms of speed, Full Impact performs all right, but not spectacularly, on small files. On a 1-megabyte machine (I used a Mac SE with a hard disk drive for these tests), though, its virtual memory management really slows things down when working with large files. For example, pasting an identical formula into every cell of a 256 by 2048 spreadsheet could easily take all night.

Full Impact and Excel perform comparably on most small-file operations (see table 1), although Excel loads somewhat more slowly than Full Impact. However, Excel is noticeably faster when handling large files once they are loaded.

I did encounter a small problem running the benchmark tests. While trying to run the Savage test, I had an error message that kept telling me "Invalid right parenthesis" even though all parentheses were properly paired. It turned out that I had entered the wrong form for one of the functions used in the formula. I used ARCTAN, while Full Impact codes this function as ATAN. However, the error message didn't tell me I had bad syntax or a nonexistent function; it said I had an invalid right parenthesis, which I didn't.

In addition to the 180-page main manual, the Full Impact package contains a plethora of little manuals and instruction booklets, each clearly labeled. A tutorial disk is also included. It takes about 15 minutes and is meant as orientation for a novice spreadsheet user. It provides an introduction to the basic concepts and features of spreadsheet operation, plus a quick look at some interesting features of Full Impact, such as its ability to include graphics and formatted text.

Full Impact is a reasonably powerful spreadsheet package with excellent presentation and graphics features, but it has less file-size capacity than the comparably priced Excel, and this capacity is limited still further by its gluttonous appetite for memory. Its ability to import files produced by other spreadsheet and database programs is limited, which means it is not really suitable for an office using a variety of spreadsheet programs. My impression of Full Impact is that it is intended for small-to-medium-size offices in which it is the only spreadsheet package; given enough memory, its features and ease of operation make it a good choice in this situation. ∎

Diana Gabaldon is the editor of Science Software *and an assistant research professor at the Center for Environmental Studies at Arizona State University. She can be reached on BIX c/o "editors."*

Table 1: *Benchmark results. Full Impact performed comparably to Excel on the standard BYTE tests.*

	Full Impact	Excel
Savage	2:40	2:21
Byte recalc	0:17	0:15
Scroll right	0:33	0:34

Note: All times are in minutes:seconds.

Circle 184 on Reader Service Card

World Radio History

dBASE IV Arrives

A major revision brings
better performance
and more features—
in a bigger package

Malcolm Rubel

After months of speculation about the newest version of Ashton-Tate's flagship product, dBASE IV version 1.0 is finally here. Is the finished product worth the wait? I think so. While the program is lacking in a few areas, such as its implementation of user-defined functions (UDFs) and true Structured Query Language (SQL) support, it represents a quantum leap over dBASE III Plus in functionality, power, and ease of use.

A lot of the prepublication criticism surrounding dBASE IV was based on the assumption that Ashton-Tate was trying to include the unique capabilities of every competing database package and more. Users worried that this approach would make the product so large as to be unusable. Fortunately, this turned out not to be the case. While the program is large, it's not unwieldy to use.

In a nutshell, Ashton-Tate has significantly increased both dBASE's features and the scope of the dBASE language. And while dBASE IV isn't uniformly faster than dBASE III Plus, performance has been greatly improved. The new (standard) version sells for $795 and runs on any IBM PC or compatible with 640K bytes of memory and a hard disk drive.

It takes time to digest dBASE IV. The program comes on 14 disks with 1570 pages of documentation in 10 manuals.

The program is, however, easy to install: You simply invoke the INSTALL.BAT file and follow the directions. dBASE IV isn't copy-protected, but you do "brand" your copy of the program with your name. Unlike previous versions, dBASE IV lets you select up to four printers, specify 43-line mode (EGA), and change your CONFIG.SYS and AUTOEXEC .BAT files and other custom options. Once you've installed the program, the manuals, on-line help, and tutorial provide ample assistance to get you up and running.

The best way to examine dBASE IV is by breaking it down into its three major functional components: the database engine, the language, and the Control Center.

The Engine

If you're a dBASE III Plus user, the first difference you'll notice in dBASE IV is its size. While dBASE III Plus ran on a 256K-byte system with two floppy disk drives, dBASE IV needs at least 640K bytes and a hard disk drive. Ashton-Tate has decided to go for performance, and if you want it, you'll have to equip your machine to handle it.

The number of fields per record has increased from 128 to 255, and the maximum number of keys (indexes) per table has gone from 7 to 47. Using DOS 3.1 or higher, dBASE IV supports up to 99 open files at once, although it still supports only 10 data tables. Field length and the maximum size of the command line have both been raised from 254 to 1024 characters. There's a command-line editor to make typing long lines easier. And Ashton-Tate has increased the maximum number of memory variables from 256 to 15,000, freeing programmers from one of the most irksome constraints in dBASE III Plus.

Indexing is another major improvement. With dBASE III Plus, you could have up to seven indexes (each using a

Table 1: *Comparative execution speed benchmarks for dBASE III Plus and dBASE IV. All tests were performed on an IBM PC AT running at 6 MHz with a Priam 28-millisecond-access-time hard disk drive and 567K bytes of memory. The drive was cleaned using Norton Speed Disk before each test. (All times in seconds.)*

Test	dBASE III Plus	dBASE IV
Screen 1	55	12
Screen 2	93	21
Screen 3	192	33
Append	37	25
Calc 1	188	25
Calc 2	60	5
Calc 3	145	8
Calc 4	32	23
Calc 5	975	73
Locate 1	39	42
Locate 2	25	46
Replace	76	48
Sort 1	86	79
Sort 2	46	42
Sort 3	108	101
Sort 4	84	80
Index 1	123	70
Index 2	77	43
Index 3	174	86
Index 4	52	45
Index 5	110	102
Join	516	763
Seek 1	53	22
Seek 2	62	25
Seek 3	58	28
Seek 4	44	21
Linked seek	48	48

file handle) open for a single data table. dBASE IV lets you have up to 47 index files open inside a master index file (.MDX) that takes up only a single file handle. Open indexes are no longer a constraint on system design.

The dBASE IV engine now supports two numeric data types: binary coded decimal (BCD) and floating point, which use the dBASE data types N and F, respectively.

Memo fields used to give programmers headaches. Now you can store up to 64K (previously 4K) bytes of data in a memo field, which means that you can actually use them to do work. You can use string functions and the GET command on memo fields and define a win-

dow for editing them instead of having the memo field automatically take over the entire screen. Garbage collection, once a problem, is now automatic.

dBASE III Plus allowed for only a single active parent-child relationship. dBASE IV lets you relate multiple children to a single parent using different keys. In many instances this allows a substantial reduction in code size and complexity.

Finally, execution speed is significantly higher because dBASE IV preparses code to save execution time (see the comparative benchmarks in table 1). A "compiler" translates dBASE programs into intermediate code, checking for syntax errors while assembling code tokens for execution. This is not a *true* compiler, however. The dBASE IV compiled code comes out in .DBO files (not .EXE files) and still needs either dBASE IV or dBASE Run Time to execute.

You may find that your dBASE III Plus code generates compiler errors when first run under dBASE IV. As an interpreter, dBASE III Plus simply executes the first line of code that meets the specified condition of an IF...ELSE... ENDIF or a DO CASE structure. If there's no ENDIF or ENDCASE statement, it doesn't matter, because the interpreter never reaches that line. This makes a difference with dBASE IV, however, so you'll find out all about your bad coding practices.

As expected, the performance benchmarks for dBASE IV show the greatest speed improvements when no disk access is required (see table 1). Surprisingly, several of the direct data table functions (locate and join) were slower when executed by dBASE IV than they were when executed under dBASE III Plus. Fortunately, most users will seldom need these features.

While I ran the benchmarks in the table in 567K bytes of available memory, I also tried to run them in the least amount of memory possible. I managed to get the entire set of benchmarks to run in 477K bytes of memory, and the only major speed penalty appeared in indexing test files (seeking was not affected). Index speeds were about a third slower.

The Language

dBASE IV has enhanced and added features to the dBASE language while retaining compatibility with dBASE III Plus. Some commands, however, are different from those used in competing products such as FoxBASE and Clipper. The new menu commands in dBASE IV are a prime example. But the changes

have definitely improved the language, and the capability to add true context-sensitive help to applications you're developing is a vast improvement over dBASE III Plus.

Ashton-Tate swears by its statement of 100 percent upward compatibility, and, indeed, everything I ran with dBASE IV worked—even old report and label files. The first time you execute your old dBASE III Plus report forms, the program converts the binary file into dBASE code and then compiles and executes it.

The additions and differences between dBASE IV's 310 commands and functions and those in dBASE III Plus are so extensive that I can't cover all of them in a single article. But some of the more interesting command sets include the following:

- A complete set of commands for defining, saving, and using regular pull-down menus and Lotus-type menus.
- A true BROWSE command that's under the programmer's control.
- A full set of commands for defining, using, and saving windows.
- Two-dimensional array commands (arrays can include up to 1170 elements).
- An enhanced GET command that includes the VALID clause and conditional editing (WHEN), with custom prompts and error messages for each GET.
- Twenty-eight new SET commands.
- New direct date handling using the format {mm/dd/yy}.

dBASE IV supports UDFs, which let programmers plug in custom-designed routines. This enhancement alone might have made dBASE IV worth the price. Unfortunately, there are several major limitations to dBASE IV's implementation of UDFs.

First, you must declare functions either in the current procedure file or at a higher level in the program. The easiest way to do this is to include all your functions in the declared procedure file (SET PROCEDURE TO) so that they're available globally.

Second, you can only use the CLEAR and READ commands conditionally. CLEAR can have no arguments, and you can use READ only if no format file is active. Both of these restrictions are annoying, but what's more limiting is that you can't use the 81 commands and 13 SET commands in a UDF. Also, you can't do macro expansion in a user-de-

continued

fined function. This is crippling and will force programmers to write many good functions as procedures.

Finally, Ashton-Tate UDFs rigorously enforce parameter passing. If you declare a parameter, you must pass it. By contrast, Nantucket Corp.'s Clipper lets you pass fewer than the declared number of parameters. Using its PCOUNT() (parameter count) function, you can then substitute default values for the unnamed parameters. I had hoped Ashton-Tate would add this capability in dBASE IV.

In contrast to dBASE III Plus, dBASE IV has a debugger that's workable. The compiler also gives the programmer more information about syntax errors than did its predecessor. For me, it's actually a pleasure to debug programs in dBASE IV.

dBASE IV also treats programs and procedures differently. It treats all programs as procedures and allows multiple procedures in a single .PRG file. You're no longer restricted to procedures only in the named procedure file. Any program file can now contain 963 procedures or functions (up from 32), and you can have up to 32 open procedure files at one time. You can also write proper modular code and have virtually all your programs included in one or two procedure files. The concern for the number of open files that was always hanging over the dBASE III Plus programmer's head is gone.

Last, dBASE IV includes a program editor that will be familiar to Framework users. Program length is now 32,000 lines, as opposed to the previous 4000-byte limit. The manual states that you can replace the program editor with your own, but the editor I use, XyWrite, wouldn't fit into the memory remaining after loading dBASE.

The dBASE IV editor's one serious drawback is that it's slow when working with large files. It took 37 seconds to load and 17 seconds to save a 2100-line program file on my AT's hard disk drive. The debug cycle also can suffer from this type of limitation.

The Control Center

What's extraordinary about dBASE IV is that much of its power is also available to the nonprogrammer through its task-oriented, nonprocedural Control Center. The Control Center is to dBASE IV what the Assist mode is to dBASE III Plus. But while Assist was inadequate, the Control Center is a powerful dBASE shell with a point-and-shoot interface that lets you develop customized applications without writing code.

The Control Center is a collection of well-thought-out work surfaces that let you define data tables, screens, reports, and labels and tie them together through the application generator. The work surfaces are intuitive in nature, and with just a little experience, you should be able to design functional programs.

The report writer is a vast improvement. It looks like Concentric Data Systems' R&R Relational Report Writer for dBASE. You can now create a wide variety of customized layouts. And since dBASE IV now fully supports a variety of printers, you don't have to write routines when you want to print in anything other than normal text. The program supports many standard types of labels in its label panel. dBASE IV also gives you the opportunity to define your own label specification.

The forms designer looks a lot like WallSoft Systems' UI Programmer software, but it isn't as easy or as powerful to use. The screen painter shows only 19 lines on the screen. Although you can scroll up to create a full screen (or even multipage screens), it would be better if you could see the complete screen in development. You can add pictures, ranges, valid statements, error messages, and editing conditions to suit your needs.

You can also define queries in a query-by-example panel that either uses a single table or links two or more tables together. Queries can be conditional, and you can present them either in table form or in a user-defined screen form (.FMT). You can create queries without writing a single line of code, and once set, you can carry them over to the reports and labels sections. While it would be an overstatement to say that a first-time user could do it with ease, new users can learn to develop sophisticated queries.

Finally, you can now record keyboard macros for later playback, automating repetitive tasks.

The most remarkable thing about these tools is that, unlike dBASE III Plus, which wrote binary files for report and label specifications, dBASE IV writes out program code. Sophisticated users can then modify this code to suit personal needs, and those who have the Developer's Edition of dBASE IV can also write customized templates.

A Step Toward SQL

Structured Query Language is currently a hot topic, and dBASE IV supports SQL—sort of. Actually, it emulates SQL using dBASE data tables. You can use a set of SQL commands inside dBASE by typing SET SQL ON. This deactivates the dBASE IV commands that conflict with the SQL language.

Essentially, a series of translation tables lets you use SQL commands to work on dBASE tables. But it's not SQL. And it's slow because of the intermediate steps that the program must take to execute the code. I recommend that you wait until the Ashton-Tate/Microsoft Database Server comes out before using the emulation.

While users still must wait for true SQL support, Ashton-Tate has done much to satisfy the needs of the multiuser system developer. dBASE IV is transparently multiuser. File and record locking are better thought out, and the BEGIN TRANSACTION/ROLLBACK command set is a godsend that will keep developers from having to create small

World Radio History

data tables for each transaction and then appending them when a transaction is finished. Automatic and immediate screen refresh is another excellent feature, as are the program's eight security levels. Another good addition is support for shared locks, which let other users view data while you're editing it.

For the serious dBASE programmer, Ashton-Tate also offers a $1295 Developer's Edition. At $500 more than the Standard Edition, it's a bargain. The package contains everything in the stan-

> *If you don't have a hard disk drive and 640K-byte memory, you might as well forget dBASE IV.*

dard product, plus a linker that lets you bind all procedure and program files into a single file and do overlays for memory management. It also includes an unlimited run-time license so that you can distribute applications (single or multi-user) royalty-free.

The Developer's Edition also comes with two LAN keys so that you can set up and use a three-station LAN with one copy of the program. Also included is a template language and compiler. You can develop or modify programs to suit your needs, compile them, and use them from the Control Center to produce customized programs. Finally, Ashton-Tate says it will ship its .EXE compiler to registered Developer's Edition users when it's available.

dBASE's Downside
Because dBASE IV is so powerful, you pay a price in terms of the system resources it requires. Ashton-Tate makes no apologies about the size of dBASE IV. But if your machine doesn't have a hard disk drive and at least 640K bytes of memory, you might as well forget about dBASE IV. And if you absolutely have to have terminate-and-stay-resident (TSR) programs running, you'll need a copy of DESQview or Windows (I use DESQview so that I can use my own editor in a window to write and debug programs). dBASE IV will run on a LAN, but net-

work overhead will be cutting your margins thin in terms of memory.

The size of dBASE IV will also influence your programming style. Ashton-Tate decided to define and save windows and menus in memory so they will pop up quickly. But when you have seven or eight menus defined, five or six large windows, 600 memory variables, and a couple of arrays defined, you don't have much memory left. You can save windows (but not menus) to disk to save memory, but this option requires time-consuming disk I/O to reactivate the windows.

You can now include up to 963 procedures in a procedure file, and you can call down through 32 nested procedure files. But 963 is the theoretical maximum. Each procedure requires a procedure name and a pointer stored in memory and takes 25 bytes of memory. The maximum number of procedures depends on how much memory is available. The actual limits placed on the product mean that for most users' systems, the limits are much lower than the theoretical maximums.

Worth the Wait
Is dBASE IV worth the price? The answer is a qualified yes. If you've got the system resources, dBASE IV is an excellent choice, whether you're a dBASE III Plus user or a newcomer considering buying dBASE for the first time. Many developers, however, probably will wait for the release of version 1.1 to see whether bugs and shortcomings, such as dBASE IV's limitations on UDFs, are cleared up before porting all their applications to dBASE IV.

It is likely that undiscovered bugs will come out when 50,000 people start using the product. Only 2 weeks before dBASE IV was introduced, a beta user reported that he could not compile a program file with 1000 procedures in it. Ashton-Tate then lowered the number of procedures from 1170 to 963. [Editor's note: *Due to the complexity of dBASE IV and unconfirmed reports of other bugs, BYTE is doing further testing on it. Look for a Review Update on dBASE IV in an upcoming BYTE.*] But the product seems solid overall, and most users will find that dBASE IV was well worth the wait. ∎

Malcolm Rubel, president of Performance Dynamics Associates, a business applications consulting firm in New York City, is currently at work on his new book, dBASE IV Procedures and Functions. *He can be reached on BIX c/o "editors."*

NOW THERE A
TO THE POST

PostScript printers and typesetters offer a variety of resolutions, paper handling options and printer speeds, as well as black & white and color output.

In just three years, the PostScript language has been adopted by more than 25 of the world's leading computer equipment manufacturers.

Agfa-Gevaert   Apollo Computer
Apple Computer   AST Research
Autologic   Dataproducts   Diconix
Digital Equipment Corporation   Fujitsu
General Computer   Hewlett-Packard   IBM
Linotype   Matsushita   Monotype   NBI
NEC   NeXT   QMS   Quadram
Qume   Ricoh   R.R. Donnelley
Scitex   Texas Instruments
Varityper   Wang

More than 3,000 software programs– for minicomputers, mainframes, work- stations and personal computers–support the PostScript language.

THE OUT SIDE.

To professionally print words and pictures on the same page, all you need are laser printers and typesetters that speak the right language. The language of PostScript.

As the standard page description language of electronic publishing,

Look for this symbol on computers, printers and other products
POSTSCRIPT that support PostScript software from Adobe Systems. It's your guarantee of quality and compatibility.

PostScript software from Adobe Systems is your key to complete device independence–giving you the freedom to select the right hardware, soft- ware and vendors for your needs and budget.

Isn't it time you put the power of PostScript on your side?

POSTSCRIPT

RE TWO SIDES
SCRIPT® STORY.

Computers and workstations with the Display PostScript system deliver a variety of screen resolutions in black & white, grayscale and color.

Display PostScript is already the graphics standard for a new generation of computer and workstation displays from these manufacturers.

Taking advantage of enhanced software portability, applications with greater power and potential are already emerging to support Display PostScript.

THE IN SIDE. Introducing Display PostScript™ from Adobe. With the same imaging standard and language used in PostScript printers, this system software module is now in computer and workstation displays.

And because it's transparently integrated over a variety of operating systems, Display PostScript is your key to hardware compatability and software portability.

Take advantage of a powerful, new graphics standard for today–and tomorrow. Be sure to insist on the Display PostScript system.

If you'd like more inside information about PostScript software and/or the Display PostScript system, call 800-29-ADOBE.

ADOBE
SYSTEMS INCORPORATED

Personal
Workstations

As personal computers have grown in size and capability, they've begun to encroach upon the workstation arena. And as workstations have added features that make them more accessible and user-friendly, they've begun to reach into the personal-computer arena.

As a result, we now have a new breed of computer with the capabilities of a workstation and the accessibility of a personal computer—the personal workstation. This new breed includes high-end personal computers and low-end workstations that share certain minimum requirements and cost less than $20,000. In "Two Worlds Converge," Nick Baran discusses these requirements and the merger taking place between what were once two separate and distinct fields.

Then, Bill Nicholls looks at "The Current Crop" of workstations, including the latest offerings from Apollo Computer, Sun Microsystems, Silicon Graphics, and NeXT. He also discusses his own experiences trying to convert his 80386-based machine into a workstation. In their quest for speed, workstations will undoubtedly explore reduced-instruction-set-computer architecture. In "Worth the RISC," Trevor Marshall and I discuss RISC technology and compare the various RISC chips available today.

The speed question seems to come up in the computer field whether you're talking about chips or peripherals or just about anything else, and it comes up here

as well. But what do all those various ratings mean? MIPS? MHz? MFLOPS? In "How Fast Is Fast?" Bill Kindel sorts out the various speed measurements and tells us what they do—and don't—imply.

You can't talk about workstations without discussing graphics, and in "Art + 2 Years = Science," Phillip Robinson delves into the state of the art in workstation graphics, from techniques to applications to machines. You also can't talk about workstations without discussing Unix and networking. In "Networking with Unix," Greg Comeau combines the two and compares Sun's Network File System with AT&T's Remote File System.

Personal computers and workstations once seemed to be two separate and distinct fields. Their capabilities were different, their uses were different, and their prices were very different. Today, the edges of those fields overlap, and the personal workstation is born. And tomorrow? If current trends are any indication, tomorrow will bring a continuum of microprocessor speed and performance from which we can only benefit.

—*Jane Morrill Tazelaar*
Senior Technical Editor, In Depth

Pack extra power into your PC.

Make it over into a '386.

DTK's new PEM-2000 dual-speed 8/20 MHz 80386 motherboard gives you 100% PC/AT compatibility at speeds up to 27.3 MHz. It also offers some very elegant engineering, like eight expansion slots including two for 32-bit memory expansion, two serial ports and one parallel port, a DTK BIOS with built-in diagnostics, and the socket for an optional 80387 coprocessor.

DTK means value in PC-compatible motherboards, add-on and networking cards, and bare bone systems, including FCC Class B-certified 10 and 12 MHz '286 computers. Which is why the two high-speed XT clones named "Best Buys" in *PC World*'s August, 1988 issue, both use DTK motherboards.

The PEM-2000 is a good example—but by no means the only example—of the DTK difference. To get the full story on all our high-performance XT, '286 and '386-based products, contact the DTK office nearest you for specifications and pricing.

And find out how to pack some extra power into your PC.

® **Where price and performance meet.**

DTK COMPUTER INC.

15711 E. Valley Blvd. • City of Industry, CA 91744
Tel: (818) 333-7533 • Fax: (818) 333-5429

DTK Computer Inc. of Florida	DTK Computer Inc. of New Jersey	DTK Computer Inc. of Texas	EXCLUSIVE DISTRIBUTORS:	DTK Computer GmbH
7245 Corporate Center Dr., Suite B	300 Columbus Circle, Raritan Center	10535 Wilcrest Dr., Suite 120	Florida: TWC (305) 599-0871	Wahlerstr. 16
Miami, FL 33126	Edison, NJ 08818	Houston, TX 77099	Northern California; Supercom	4000 Dusseldorf 30
Tel: (305) 477-7440	Tel: (201) 417-0300	Tel: (713) 568-6688	(415) 770-1100	West Germany
Fax: (305) 477-8322	Fax: (201) 417-0307	Fax: (713) 568-5688	San Diego; Nimax	Tel: (0211) 656031
			(619) 566-4800	Fax: (0211) 653753

AT and XT are registered trademarks of International Business Machines Corporation. DTK is a registered trademark of Datatech Enterprises Co., Ltd.

Circle 79 on Reader Service Card (DEALERS: 80)

World Radio History

Two Worlds Converge

*As workstations become more friendly and PCs grow
in capability and capacity, the personal workstation is born*

Nick Baran

Perhaps the one constant in the computer industry is that the price/performance ratio keeps going down. Each year, you get greater performance and capabilities for your money.

Consider that, in 1982, an 8088-based IBM PC with 64K bytes of memory, two floppy disk drives, and a monochrome display cost close to $5000. Today, you can buy an 80386 machine with 2 megabytes of memory, a hard disk drive, and a high-resolution graphics monitor for about the same price. This trend in price/performance gradually has brought low-end workstations into the same ballpark as high-end personal computers (PCs)—hence the term "personal workstations."

In this article, I'll discuss the basic characteristics of a workstation, who uses them, and how. I'll also look at how and where personal workstations fit in with high-end workstations and PCs.

First, a Little History
Workstations actually predate PCs. The direct descendants of 32-bit minicomputers, they first arrived on the computing scene in the late 1970s. The first work-

stations were simply high-resolution graphics terminals connected to multiuser minicomputer systems running mechanical and structural design software used by engineers and mechanical designers. A typical "graphics workstation" cost several thousand dollars and ran on minicomputers costing several hundred thousand dollars, such as the VAX-11/780 from Digital Equipment

Corp. (DEC) or the Prime 750 from Prime Computer. These early workstations typically featured 1000- by 800-pixel 19-inch monochrome monitors, a keyboard, and an input device, such as a mouse or a digitizing tablet. They were used for CADD and for modeling mechanical systems and components.

But these early workstations had a major problem. Because they ran on a host minicomputer, which might have as many as 30 or 40 users simultaneously sharing resources, they were often painfully slow. Engineers would have to wait for hours for their applications to run because they were contending for resources with other engineers, as well as other departments such as accounting and word processing. Another problem with these minicomputer-based workstations was that they were confined to proprietary operating systems, which limited the availability of third-party software applications and restricted the portability of applications.

The obvious solution to the minicomputer bottleneck was to give graphics workstations their own CPUs, making them independent from the general-

purpose minicomputer system. One way around the proprietary-operating-system problem was to offer stand-alone workstations running Unix, a multitasking operating system in the public domain.

But minicomputer systems also offered one major advantage. You could share files and applications, because all the workstations were connected to the minicomputer. The answer to this on stand-alone workstations was to provide a local-area network (LAN) that allowed workstation users to share files and peripherals.

Apollo Computer was the pioneering manufacturer of stand-alone workstations. Founded in 1980, Apollo's first offering was the Motorola 68000–based Domain DN100, which had a price tag of close to $60,000 for a fully configured system. A couple of years later, Sun Microsystems entered the workstation market with a competitive Unix workstation called the Sun 100, which was also based on the 68000. The Sun 100 had a 10-MHz 68000 processor, a 1000- by 800-pixel monochrome graphics display, and a six-slot MultiBus card cage. The rated integer performance of the Sun 100 was 0.5 million instructions per second (MIPS). The Sun 100 with 256K bytes of memory, Unix, and Ethernet cost about $13,000. But a fully configured system with 2 megabytes of memory and an 80-megabyte hard disk drive cost over $30,000. From the start, Sun Micro-

systems and Apollo have been fierce competitors in the workstation market. Certainly, competition continues to contribute to the favorable price versus performance trend.

Historically, workstations have been one step ahead of PCs. While PCs started off as 8-bit machines, workstations started off on 16-bit processors. They were the first stand-alone machines to be networked, and they offered more advanced graphics capabilities. While early PCs either didn't have graphics at all or had only 640- by 200-pixel black-and-white graphics, workstations came with 1000- by 800-pixel resolution or better. Because workstations were based on Unix, they offered multitasking and large memory management capabilities well before those features were available on microcomputers.

Another important difference in the evolution of conventional workstations and PCs was the software. While the first PC software consisted of games, simple file management programs, and text editors, workstation software migrated from the professional-level applications available in the minicomputer environment. Structural analysis, CAD, graphics design, database management, and page layout provided the core applications on workstations. Not only was this software considerably more sophisticated than the early applications on PCs, it was a lot more expensive. Software

packages usually cost several thousand dollars on a workstation, and they often cost that much today.

But the PC offered something that workstations couldn't match: ease of use and affordability. While workstation users faced the arduous task of learning Unix, PCs offered easier accessibility. Although MS-DOS was no picnic, it certainly seemed that way when compared with Unix. Then the Macintosh desktop began the revolution that forced workstation manufacturers to come up with better graphical interfaces. The NeXT graphical interface may be the first truly easy-to-use Unix-based graphical interface.

The development of good spreadsheet, database, and word processing programs on the PC began to build a bridge between the PC and the workstation. For a few hundred dollars, you could get software on the PC that was functionally superior to the equivalent workstation software costing hundreds or thousands of dollars more. You couldn't get Lotus 1-2-3 or Microsoft Word on a workstation, so workstation users started demanding MS-DOS compatibility. This was the beginning of the bridge from the workstation end. Today, many workstations offer some method to run MS-DOS as a task in the Unix environment (SoftPC from Insignia Solutions and Merge 386 from Locus Computing are examples of products that let you run DOS in the Unix environment).

Workstations Today

The workstations of 1989 are a far cry from the Domain DN100 or the Sun 100 (see "The Current Crop" by Bill Nicholls on page 235). A low-end workstation or a high-end PC—the overlap is fairly complete and the distinction blurred—can now be referred to as a "personal workstation." Today, the personal workstation features a 32-bit processor, at least 4 megabytes of memory, 1024- by 800-pixel screen resolution or better, 4 or 8 bit planes of color, at least 70 megabytes of hard disk storage, and Ethernet and Unix capabilities. It also normally includes a floating-point co-processor.

Workstations are available in three main bus architectures—VME, Multi-Bus, and the IBM PC AT bus. Apollo Computer's personal workstations, for example, use the AT bus, while the higher-end machines use VME or MultiBus. Sun Microsystems primarily uses the VME bus but offers the AT bus in its 386i product line. And the NeXT Computer has a NuBus-compatible backplane.

A three-dimensional image displayed on the new Personal Iris from Silicon Graphics. The Personal Iris features an R2000 RISC processor from MIPS Computer and a 1280- by 1024-pixel color monitor.

A typical base system, as configured above, costs around $12,000. However, costs increase dramatically when you add additional graphics capabilities, such as a graphics accelerator board and additional mass storage. A system costing $12,000 will rapidly climb to $30,000 if you add a graphics coprocessor and a larger hard disk drive. The photo at left shows the graphics capability provided on the new Personal Iris workstation from Silicon Graphics. While an entry-level diskless version of the Personal Iris costs about $16,000, a system with a 170-megabyte hard disk drive and z-buffering (hidden-line removal in hardware) costs over $30,000 (see Phillip Robinson's article "Art + 2 Years = Science" on page 255). NeXT's personal workstation may have a significant impact on these ballpark workstation costs.

While the initial workstation offerings were based on the Motorola 68000, today's workstations often use the Motorola 68020 and 68030, as do many PCs. Apollo still offers only Motorola-based machines. Sun Microsystems and Prime Computer also offer machines using the Intel 80386, as do many PC companies. Some workstation companies offer machines using reduced-instruction-set-computer (RISC) architecture, and some, notably Sun Microsystems, offer a similar architecture called scalable-processor architecture (SPARC). IBM, Hewlett-Packard, and Silicon Graphics also offer RISC-based workstations.

But because these workstations all run Unix, the type of processor is not that important to the end user. What is important is the performance of the machine and the available software. Integer performance of today's workstations ranges from about 1.5 MIPS for personal workstations to about 10 MIPS for top-of-the-line systems that cost from $50,000 to $100,000 (Silicon Graphics has lowered the price/performance ratio even more, claiming a performance of 10 MIPS on its Personal Iris). RISC-based systems generally provide higher performance than the complex-instruction-set-computer (CISC) machines. (For more information on RISC and the various RISC chips, see "Worth the RISC" by Trevor Marshall and Jane Morrill Tazelaar on page 245.) For example, Sun Microsystems claims that the high-end systems in its product line with SPARC processors provide a 7- to 10-MIPS integer performance. Sun's 68020 line of machines offers performance ratings from 1.5 to 4 MIPS, and its 80386-based 386i machines range from 3 to 5 MIPS.

An equally important measure of performance is millions of floating-point operations per second, or MFLOPS. Floating-point performance depends on the type of floating-point processor used in the system. Lower-priced machines with Intel 80387 or MC68881 floating-point processors have a floating-point performance of about 0.2 MFLOPS. Machines with floating-point accelerators have performance ratings ranging from 0.6 to 2 MFLOPS.

However, floating-point accelerators

With large, high-resolution graphics monitors, it is possible to have multiple windows with multiple sessions operating concurrently.

raise the cost of the system by as much as $10,000 to $15,000. Even more expensive floating-point accelerators can provide performance of up to 12 MFLOPS. For example, Sun Microsystems offers an accelerator called the TAAC-1 that claims 25 MIPS and 12.5 MFLOPS performance for $30,000. (For further discussion of performance metrics, see "How Fast Is Fast?" by Bill Kindel on page 251.)

Who Uses Workstations, and How?
The main users of workstations are still engineers, scientists, architects, and mechanical designers. However, workstations have also become attractive for people who work in the fields of animation, graphics design, and desktop publishing.

A major user community for workstations is the university. The figure on page 232 shows a network of workstations (and also some minicomputers and mainframes) at the University of California at Santa Cruz (UCSC). This system is typical of the workstation environments at many universities. Note that the planetary names on the diagram are the address names of the systems on the net-

work. The polygons surrounding systems on the diagram (e.g., Sol/Daizu) represent groups of diskless workstations and their file servers.

The UCSC system exemplifies some of the features of workstations that distinguish them from PCs. Workstations are almost always networked. Ethernet and Sun Microsystems' Network File System have become the networking standard for many workstations, while others use AT&T's Remote File System. (For a discussion of the relative merits of NFS and RFS, see "Networking with Unix" by Greg Comeau on page 267). Networks allow the use of diskless workstations that can access files from a file server. As UCSC's software manager Al Conrad told me, the idea is to "put computer power on people's desks and the storage in a centralized room." Conrad points out that it is cheaper and more efficient to have a 1-gigabyte hard disk drive that everyone can use than it is to have 10 100-megabyte hard disk drives carrying duplicate applications.

Another feature of the UCSC system is that it runs Unix. Machines on the network run everything from 4.3 Mach Unix to System V.2 to Xenix. As long as you're in character mode, applications are quite portable between the various Unix versions, according to Conrad. However, incompatibilities exist between the various Unix window managers and interfaces, such as NeWS, X11, and proprietary window managers like Suntool or the Macintosh environment. All these windowing systems run on the UCSC network. The Macintoshes are linked to the system via Kinetics' interface between AppleTalk and TCP/IP.

Two important requirements for personal workstations are that they have at least 4 megabytes of memory and that they can use high-resolution graphics monitors displaying close to 1 million pixels on the screen (*megapixel* displays). There are several reasons for these two requirements. Perhaps the most important of these is multitasking, one of the main attractions of Unix. In conjunction with large, high-resolution graphics monitors, you can have multiple windows with multiple sessions operating concurrently. But you need a lot of memory to run multiple sessions simultaneously—the more the better.

A multitasking windowing environment greatly improves productivity. Programmers can work on multiple subroutines or program modules simultaneously. Writers can view and work on multiple documents at the same time.
continued

And users can access multiple nodes on the network simultaneously, each window corresponding to a different session with a different network node.

High-resolution graphics are an essential component of workstations. Most have 1-megapixel (approximately), 19-inch monitors, which provide not only excellent windowing capabilities but also the resolution necessary for CAD and graphics design. High-resolution graphics capabilities also are available on PCs, of course; the standard for PCs is still around 640 by 480 pixels, while the standard for workstations is 1 million pixels. Color is just starting to become standard on workstations—until recently, most

workstation graphics monitors were monochrome or gray-scale.

Workstations require a lot of memory because of the types of applications they run. Consider the requirements for the computations performed in finite-element-analysis problems, for example. Finite-element analysis and other engineering software require the solution of simultaneous equations. A small finite-element model with only 200 degrees of freedom (the number of variables defining the displacement of the model) requires solving a 200 by 200 array of simultaneous equations. A fully populated 200 by 200 array contains 40,000 coefficients. Because these types of problems

require a high degree of accuracy, the coefficients are represented in double-precision, requiring 8 bytes per coefficient. So, this small problem requires 8 bytes multiplied by 40,000, or 320,000 bytes of memory. Many finite-element problems involve two or three *thousand* degrees of freedom. If you wanted to solve a 1000 by 1000 array of simultaneous equations, you would need 8 million bytes of memory in addition to the memory required for the applications software and the operating system.

Most finite-element programs do not require that the entire array fit into the main memory of the computer. They get around the memory limitations by swap-

The Ethernet network at the University of California at Santa Cruz (UCSC). Note that the planetary names represent addresses on the network. Groups of machines enclosed by polygons represent diskless nodes and their file servers. (Figure courtesy of the Computer and Information Sciences department, UCSC.)

ping parts of the array to and from disk. However, the more memory you have, the faster you can solve large arrays. Several finite-element programs run on IBM PCs with 640K bytes of memory. However, solving large problems can take all night on a machine that size. Workstations offer two major advantages for finite-element analysis: multitasking and much larger memory capacity. Both capabilities are also available on personal workstations. Even if you're running a long, time-consuming problem on a workstation, you can run it as a background task and continue working in other sessions on the computer.

Large mass-storage capability is another important requirement of personal workstations. Applications that need a lot of memory usually require a fair amount of disk space. Again, finite-element analysis is a good example. A large finite-element problem of several thousand degrees of freedom can easily take up 20 megabytes of disk space. In addition, operating systems and applications continue to grow rather than shrink. For example, Apple's A/UX is shipped on an 80-megabyte hard disk, of which 70 megabytes are taken up by the system software. And the NeXT Computer is shipped with a 250-megabyte optical drive that is already two-thirds full.

However, mass storage isn't an absolute requirement for workstations. You can have diskless workstations accessing a file server, which provides the hard disk drives and the tape backup. The diskless-node-and-file-server model makes an awful lot of sense in a network of workstations. This type of system is more efficient, and it also eliminates the storage of redundant data. In addition to hard disk storage, a workstation environment usually includes tape backup. In a network of diskless workstations, the tape backup system is part of the file server.

Workstations vs. PCs

Until OS/2 becomes widely used, multitasking will remain fairly limited on PCs. Without shared memory and interprocess communications, the multitasking that you can do with, say, DESQview is quite limited in comparison to the flexibility of Unix multitasking environments. DESQview is restricted to 640K bytes of main memory and a 128K-byte window into expanded memory, which limits the performance and the size of multitasking applications. It could also be argued that the cooperative multitasking environment of the Macintosh Multi-Finder is not as robust or well developed

as the preemptive multitasking model used in Unix (see Phil Goldman's article entitled "MultiFinder Revealed" in the *Macintosh Special Edition*, August 1988 BYTE).

If you take a look at the components of the typical workstation, it is clear that they include both the low-end workstation, like the NeXT machine and the Sun386i, and the high-end PC, like the IBM PS/2 Model 80 or the Macintosh II. These machines have 32-bit processors

M*ass storage isn't an absolute requirement for workstations. Diskless workstations can access a file server with hard disk and tape backup.*

and floating-point coprocessors and can support at least 4 megabytes of memory. You can buy them with big hard disk drives, Ethernet, a version of Unix (Santa Cruz Operations' Xenix, AIX from IBM, or A/UX from Apple), high-resolution graphics, and so on. For example, UCSC uses Rose Hill Systems AT-386s equipped with 24 bit-frame buffers for instruction in graphics. As you can see in the figure, these machines are linked into the network just like any other workstation.

So what, exactly, is the difference between workstations and PCs? The distinction is largely cultural. PCs evolved primarily in the business and home markets, while workstations evolved in the engineering and mechanical-design markets and in the research environment of academia. MS-DOS and the Macintosh Finder were not designed for use in engineering; they were designed for business and home users. Unix, on the other hand, has traditionally been an operating system for academic and scientific computing. It was designed to support multitasking and multiple sessions, and workstations were built from the ground up to run Unix.

A Melding Pot

As we approach the 1990s, the distinction between workstation and PC is becoming less and less obvious. The NeXT Computer is an excellent example of the fading distinction. It is in every respect a workstation (with the exception of its lack of color capability). And, as Steve Jobs put it, the NeXT Computer "raises the lowest common denominator" for computing.

Soon, the only way we'll be able to tell the difference between traditional workstations and PCs will be by the operating system they run. Workstations will continue to run Unix and use Ethernet with NFS or RFS. PCs will run OS/2 or a Macintosh operating system. They will use Microsoft and 3Com's LAN Manager, Banyan Systems' VINES, Novell's NetWare, or AppleTalk. And personal workstations? At least some of them will be able to "switch hit."

Multitasking, networks, and high-resolution graphics are gradually becoming requirements in the business world. Business users want large, high-resolution displays with multiple windows. They want networked systems with access to large file servers. And they want multitasking to perform multiple tasks simultaneously. So, although PCs and workstations come from different backgrounds and serve different users, they are rapidly converging.

Without a doubt, the price/performance ratio for computing power will continue to drop. With a "university price" of $6500 for an 8-megabyte machine with a 1-megapixel display, 250 megabytes of mass storage, and Ethernet built in, the NeXT Computer will almost certainly force other companies to deliver less expensive and more capable machines (for a detailed look at NeXT's machine, see "The NeXT Computer" by Tom Thompson and Nick Baran, November 1988 BYTE). Sun Microsystems, for example, is expected to announce some new lower-cost personal workstation entries in 1989, and we can expect competitive Unix-based personal workstations from traditional microcomputer manufacturers like Apple and IBM. ∎

Nick Baran holds a BSME from Stanford University and is a BYTE senior technical editor based in San Francisco. He can be reached on BIX as "nickbaran."

Circle 262 on Reader Service Card

World Radio History

The Current Crop

Buying a workstation has its benefits—but building your own on a PC base is also a valid option

Bill Nicholls

A workstation is the Ferrari of the personal computer world. Most of us have seen one and occasionally watched one in use, but in general, we know little about them beyond the mystique they carry. Unlike the Ferrari, however, the price of workstations is coming down due to improvements in semiconductor technology. Also unlike the Ferrari, many of us are likely to get better acquainted with workstations in the near future.

What is a workstation? Is there really a difference between a workstation and a personal computer? What can you find for under $20,000? When I went looking for answers to these questions, I was in for some surprises. Despite the history and advertising for workstations, the differences between low-end workstations and high-end personal computers no longer outweigh the similarities. In fact, their capabilities have become so similar as to blur the distinction between the two. They can, effectively, be grouped together as "personal workstations."

I found three major areas of differentiation between personal computers and workstations: design, networking, and graphics. The design issue is the most central. Workstations are designed as integrated products, tools with which to perform useful work. They avoid the "fill-in-the-blanks" approach to system configuration that the personal computer arena has carried to extremes. Workstation designs also recognize that work is rarely done in isolation and so provide a level of networking that would make the most ardent local-area network admirer jealous. Finally, workstation design responds to the need for people to interact and control the system, and it accomplishes that through an intuitive graphical interface.

In the recent past, it would not have been possible to buy all those features separately. Today, with attention to detail, you can come quite close to the workstation-design philosophy. Current 80386-system designs from Compaq and Advanced Logic Research (ALR) address the memory-performance issue, Ethernet and Transmission Control Protocol/Internet Protocol (TCP/IP) are available as the base network, and high-performance graphics are available from many companies like Matrox, Renaissance, and Number Nine.

Building a Workstation

I'm actually in the process of building my own workstation, although it's not what I started out to do (see the text box "From PC to Workstation" on page 236). My original objective was to upgrade the AT clone I was using to an

From PC to Workstation

Seven years ago, I bought my first microcomputer. It was a Radio Shack Model 3 with only 16K bytes of memory and no disk drives, but I did use it. Three weeks later, I upgraded it to the maximum of 48K bytes of memory. Three years later, I bought a Sanyo 555: 128K bytes of memory and two single-sided disk drives. This rapidly became a 256K-byte system with 360K-byte drives, almost IBM PC-compatible. These early machines were distinctly underpowered, and the mass storage and display arrangements were primitive, but with patience and persistence, I could do my work.

With my next machine, an 8-MHz XT clone, I began the transition to machines that worked for me, rather than the reverse. Upgrades followed rapidly over the next years to my current ALR 386/220, a 20-MHz 80386 system with

3 megabytes of memory, running Unix System V 3.0. Except for a large graphics display, a network, and a few design refinements, this system has the attributes of a workstation. I didn't start out intending to build a workstation but simply progressed toward systems with better support for my work.

The relevance of this history is that while the functions I perform are mostly unchanged—writing, programming, keeping track of things, doing calculations—the *style* of my work has evolved. Rather than wait for the system, it waits for me, and it improves my ability to respond by enabling quick task switching and multitasking when needed. I have traded machine cycles (and cost) for my time and come out ahead in productivity. With some further investment in graphics and network hardware, I would have a complete workstation.

80386 and migrate to Unix, connecting a few terminals. I chose an ALR 386/220, a 20-MHz system with two-wait-state memory for reliability and cost/performance. The Unix System V 3.0 software (from Bell Technologies) requires 2.5 megabytes of memory to install, but only 2 megabytes to run. I already have a large disk drive (160 megabytes) and a fast controller. The basic system with 3 megabytes of memory was easy to put together—but not inexpensive.

Currently, I'm using an EGA/CGA display combination. My original plan was to install a VGA display, but experience with workstations has made me reconsider, and I'm looking at the larger graphics displays with a large monochrome monitor (color is too expensive), supported by the new X Windows standard. The ability to work with multiple windows on a high-resolution screen is a major productivity tool. These display systems are available from 640 by 480 pixels by 8 bit planes up to 2000 by 2000 pixels by 8 bit planes in color or monochrome.

The final piece of the workstation environment is the network. This choice is easier, as the standards are Ethernet and TCP/IP. The only caution here is to verify that the software and hardware are compatible with your machine. If I were to link with another Unix system with a hard disk drive, I would add RFS, the Remote File System software, which allows

me to access the other files as though they were local, subject to the permission of the other system.

Some other choices are modified by building a workstation. One is the tape backup system. Unix does not (yet) support tape systems driven from the floppy disk drive controller, so I am about to replace my inexpensive DC-2000 system with one that is supported, with either 60-megabyte or 125-megabyte tape cartridges. Both use the larger DC-600A tape cartridge. Memory choice is simple—as much as you can afford. Unix systems will run with as little as 2 megabytes but really begin to hum at 4 megabytes. For multitasking or graphics, 8 megabytes would be nice.

The effect of too little memory is a large increase in paging and swapping to the hard disk, and a noticeable reduction in performance. Unix is disk-intensive, and the faster the disk, the better Unix performs. In retrospect, the extra cost of an enhanced-small-device-interface disk with its higher transfer rate would have been a good investment.

Buying a Workstation

Unlike the personal computer's "do-it-yourself" approach, a workstation as such comes with all the essential hardware and software you need to make it function. Only customization and vertical applications are left for you to add. Some applications may require network

support rather than include a separate disk, but that reflects on the way a group, rather than an individual, works. All the basic workstation systems have similar components.

Buying a standard workstation has its benefits, because the platform starts with a larger minimal configuration. This simplifies the software requirements by reducing the number of variables and eliminating configurations that restrict workstation performance. It also offers certain economies of scale by providing as standard certain equipment that would otherwise require extra room to add later. In particular, workstations all start with a reasonable (for today) 4 megabytes of main memory; in a personal computer, this requires substantial expansion. Network interface and graphics displays also benefit from being standard.

Another workstation design goal was to create a group working environment that is smooth, seamless, and fast. This was implemented with an integral network, a noticeable difference from the original personal computer as a standalone system. The standard Ethernet and TCP/IP have now become easily available for personal computers, reducing that gap.

One issue, that of operating systems, remains. While the personal computer has evolved with MS-DOS and is slowly adopting OS/2, workstations have always used a multitasking operating system. When Apollo began, it provided Aegis, but subsequent workstation development has made Unix a standard. Working between Unix and OS/2 is still awkward. However, the high-end personal computers have now grown into machines quite capable of running Unix, and in fact such systems exist (e.g., Apple's A/UX and IBM's AIX).

A Personal Workstation

Most vendors' lines of workstations start at the low end with the same general level of hardware and software that the high-end personal computers have. The hardware for these "personal workstations" now includes a fast 32-bit microprocessor, 4 to 8 megabytes of memory, a 15- to 19-inch graphics display, a fast network (typically Ethernet), and a large hard disk drive, either attached or on the network (see table 1).

Software may be a proprietary operating system but is more likely to be Unix with graphics and network extensions (see table 2). The difference between a workstation and a personal computer

Power Up Your Compaq 386!
Take it to the limit with Clearpoint high-performance memory

Introducing Clearpoint's
4 MB COMRAM

Order today and get
- Competitive Prices
- Fast Delivery
- Lifetime Warranty
- 24-Hour-a-Day Service

The COMRAM
Compaq 386-Compatible Memory

From the largest manufacturer of workstation memory comes the best four megabyte memory for the Compaq – Clearpoint brings you the COMRAM for your DESKPRO 386/20 or 386/25.

Manufactured to Clearpoint's stringent production and testing standards, you are assured of optimal performance with our combination megabit DIP and SOJ design. Take your 386 to its 16 MB limit by adding three Clearpoint COMRAM boards to your 4 MB motherboard. We guarantee fast delivery, plus the best package of reliability and service.

Clearpoint Guarantees
Lifetime Warranty
With the strictest dynamic and thermal testing in the industry, Clearpoint confidently offers a lifetime warranty– Then backs it up with a 24-hour Repair/ Replacement policy and an 800 service number for a warranty with teeth.

Superior Memory Design
Across the board, Clearpoint brings the workstation marketplace better memory design. Whether it's higher density or state-of-the-art DRAM technology– Clearpoint designs it better.

Guaranteed Buyer Satisfaction
Long-term satisfied customers is Clearpoint's goal. With comprehensive, user-friendly documentation, trade-up credits and competitive pricing, Clearpoint guarantees buyer satisfaction.

Other Clearpoint Workstation Memory
MACINTOSH
MC2RAM (compatible with the Macintosh II, SE and Plus) 1 MB SIMMs.

SUN
SNXSM (compatible with Sun 3/60): Available in 1 MB SIMMs, four in a set.

SNXSC (compatible with Sun 4/110 and 386i): 1 MB SIMMs.

DEC
DCME-M20 (compatible with the MicroVAX 2000): 8 and 16 MB boards.

APOLLO
DNXRAM (compatible with the DN 3000): 1 and 2 MB boards.

Call or write for a COMRAM Product Marketing Update, a free catalog, and our Memory Pointers newsletter.

Dealer Inquiries Invited.

Credit Card Orders Accepted

CLEARPOINT

Circle 259 on Reader Service Card

Clearpoint Research Corporation, 99 South Street, Hopkinton, MA 01748
1-800-CLEARPT (1-800-253-2778) or 508-435-2000

Table 1: *A representative sample of workstations below $20,000 listed in ascending-price order.* [1] *Notice the location of my "build-your-own" workstation project using the ALR 386/220.*

System	Processor and clock speed	MIPS	FPU and clock speed	Memory (megabytes)	Display size (pixels)	Disk capacity (megabytes)	Price
Apollo DN3000	68020-16	1.5	68881-16	2-8	1024×800	0-696	$4990
Sun-3/50	68020-15	1.5	68881-15	4	1152×900	71-1300	$4995
NeXT	68030-25	4	68882-25	8-16	1120×832	256-926	$6500
ALR 386/220[2]	80386-20	3.5	80387-20	1-10	640×480	67-320	$7695
Sun386i/150	80386-20	3.5	80387-20	4-16	1024×768	91-981	$7990
Sun-3/60	68020-20	3	68881-20	4-24	1152×900	71-1300	$8900
Apollo DN3500	68030-25	4	68882-25	8-32	1024×800	0-696	$8990
Sun386i/250	80386-25	5+	80387-25	4-16	1024×768	91-981	$11,990
SGI Personal Iris	R2000A-12	10	Optional	8-16	1280×1024[3]	155-2000	$16,000
Sun-4/110	SPARC-14	7	1164/1165	8-32	1152×900	141-1300	$18,900
Apollo DN4500	68030-33	7	68882-33	8-32	1280×1024	155-696	$18,990

[1] Unless otherwise noted, all workstations are configured as follows: RAM is 4 megabytes of fast 32-bit memory; all monitors are monochrome; floating-point chips are included; graphics processors are not included; Ethernet with TCP/IP is the standard network interface; and standard software is Unix with a full set of Unix tools.
[2] The ALR 386/220 has an optional 80387 chip and an optional VGA controller/multiscan display plus 2 megabytes of memory expansion included in the price.
[3] The SGI Personal Iris has a color monitor with 8 bit planes and a graphics processor included.

Table 2: *The software environments for the workstations in table 1.*

	Apollo	Sun	SGI	NeXT
Operating system	Aegis or Unix SV or Unix BSD	SunOS (Unix) and MS-DOS	IRIX (Unix)	Mach (Unix)
Network software	Domain/OS	NFS	NFS	TCP/IP
Graphics software	PHIGS	SunCGI SunGKS	SGI Graphics Library PHIGS[2]	None at press time
Development tools (in addition to standard Unix tools)	Open Dialogue User Interface Management System[2]; Domain Software Engineering Environment[2]; Domain/OS Debugger	For 386i: SunView; X11/NeWS Merge; Help Viewer	Developer's Option Package: Graphics Library, EDGE (graphical debugger), Pixie (profiler)	Objective-C Interface Builder; DSP Library Functions; PostScript Window Server
Bundled applications (in addition to standard Unix applications)	None	Sun Organizer desktop file manager	QuickPaint; QuickModeler; IRIS Work Space; IRIS Graphics Library Runtime; 4Sight Windowing System	WriteNow Find Webster's dictionary and thesaurus; Shakespeare's complete works; The Oxford Dictionary of Quotations; Mathematica; Personal Text Database; Application, Sound, and Music Kits
Third-party applications	Spring 1988 Summary of Applications Catalog lists 925 software and hardware products in 26 application areas	Catalyst catalog lists over 1200 products from more than 500 vendors	Silicon Graphics Geometry Partners lists 75 products in 16 categories	None at press time

[1] The ALR machine comes with Setup and some MS-DOS utilities. Since no other software is bundled, the operating system and applications must be bought separately.
[2] Optional.

World Radio History

based on the same chip lies in the software and in some design choices for higher performance within the system hardware.

The situation changes as you go up the workstation line. Midrange workstations use reduced-instruction-set-computer (RISC) chips at 7- to 10-MIPS (million instructions per second) performance; optional graphics processors are available, as well as large memory and disk drive sizes. High-end workstations include graphics processors and offer multiple processors for a top-end performance of 80 to 100 MIPS.

Apollo

Apollo Computer, formed in 1980, was the first of the workstation vendors. At that time, the concept of a workstation was new, and what Apollo did was to essentially invent the workstation market. Keys to that invention were a new operating system, Aegis, a networking scheme

ALR[1]

MS-DOS
Unix
OS/2

TCP/IP
(with Unix)

None

None

None

Any available for
standard Unix,
OS/2, or MS-DOS

called Domain, and a token-ring network to support the Domain system.

Although those early workstations were less powerful than some current personal computers, the true attributes of a workstation showed—an integrated network, shared resources, and a graphical interface. Apollo identified the workstation market as a workgroup, using shared resources and data sharing as key concepts, with the objective of delivering mainframe power and workgroup cohesiveness.

Apollo systems have grown from the original 0.2-MIPS Domain DN100 in 1980, based on the 68000, to the newest RISC-based DN4500, which is rated at 7 MIPS. During that same time, workstation prices came down from $60,000 for the 0.2-MIPS DN100 to about $5000 for an entry-level DN3000 at 1.5 MIPS. The high end of the Apollo line grew just as quickly—today, the top-end Apollo DN10000 ranges from $60,000 to $200,000 and from 36 to over 100 MIPS.

The entry-level DN3000 is no longer a personal-computer killer. The Mac II (2 to 3 MIPS) and 80386 (3 to 6 MIPS) machines have topped the DN3000's 1.5 MIPS on a performance basis, and they are now available with other workstation attributes, such as large graphics displays, hard disk drives, and network connection. While personal computers have advanced, so has Apollo by introducing the 4-MIPS DN3500 (see photo 1), based on a 25-MHz 68030/68882, and the 7-MIPS DN4500, based on the 33-MHz 68030/68882 combination.

The DN4500 uses a 64K-byte cache, interleaved memory, and zero-wait-state operation for maximum performance. For up to a 300 percent floating-point performance increase, a floating-point accelerator using the Weitek 3164 chip is available for the Series 3500, 4000, and 4500. Both 3500 and 4500 systems also can add a high-performance color graphics option with a dedicated graphics processor. This processor optimizes two-dimensional primitive operations on a 1280- by 1024-pixel by 8 bit-plane color display.

Apollo continues to advance its software and now offers the ability to network with a combination of Aegis and one of two variants of Unix: System V R3 or Berkeley Unix 4.3. Apollo can also bridge the gap to MS-DOS by running MS-DOS applications in an Apollo window. The Apollo Domain Network Computing System (NCS) carries the file-sharing concept a step farther. NCS allows users and programs to request un-

Photo 1: *The Apollo DN3500, a 4-MIPS workstation for $8990.*

Photo 2: *The Sun386i/150, a 3.5-MIPS workstation for $7990.*

used CPU resources from anywhere in the network to perform parts of a task for one workstation. NCS is intended as an open system with source licenses available, and it runs on the major network protocols.

Apollo continues to hold its software edge but has also become one of the founders of the Open Software Foundation (OSF). Seven major computer manufacturers formed this foundation to offer a standard software environment based on X/Open and POSIX specifications. POSIX is an IEEE operating-system standard that is closely related to Unix.

OSF is a nonprofit, industry-supported R&D organization. Its objective is to provide a standard operating environment for applications that will make it easier for users to mix and match computers and applications from different vendors. OSF will address the portability of software, the interoperability of hardware, and scalability, the ability to use the same environment and software on anything from personal computers to supercomputers. The specifications will be public, and OSF will license its software internationally.

Sun

Sun Microsystems, formed in 1982, was the second vendor to specialize in workstation products. Recognizing a gap between personal computers and minicomputers, it took a different approach from Apollo. Sun chose to build "open systems" on existing or emerging standards, or in some cases, to originate new implementations and propose their adoption and licensing to anyone who wanted them. Sun's objective was to integrate current technology and provide better price/performance without locking the customer into a proprietary environment.

Building on the Unix and Ethernet standards, Sun added a high-resolution graphical interface and implemented the Network File System (NFS) as the glue that held together different systems on the same network. Sun describes its environment by saying "the network *is* the computer" and has software that ties together heterogeneous systems.

Since September 1985, Sun has introduced a series of new systems: the 3/160 at 2 MIPS, and the 3/50 and the 3/260 in 1986 at 1.5 and 4 MIPS, respectively. In July 1987, Sun introduced the 3/60 at 3 MIPS and its first RISC-based system, the 10-MIPS 4/260. As of late fall 1988, Sun has added the 4/110 and 4/150 RISC machines at 7 MIPS, and the 386i/150 (see photo 2) and 386i/250 at 3.5 and 5

MIPS, respectively, as the first workstations based on the Intel 80386 chip.

Currently, the Sun line ranges from 1.5 to 10 MIPS across three processor architectures—the original Motorola 68020, the SPARC chip, and the Intel 80386. This is a direct result of Sun's philosophy of taking the best current technology and packaging it in a cost-effective manner. Prices range from $4995 for the base 3/50 to $39,900 for the 4/260.

In addition to the basic workstation, Sun also offers four graphics accelerators and one floating-point accelerator, the TAAC-1. The TAAC-1 offers 25-MIPS and 12.5-MFLOPS (million floating-point operations per second) computing capability with its own C compiler and libraries for $30,000. The graphics accelerators can handle from 65,000 to 200,000 two-dimensional vectors per second, double-buffering, and z-axis buffering. All this extra performance comes at a price, however, ranging from $32,900 to $59,400. These particular additions aren't likely to be on everybody's desk.

Sun's software is guided by the Open Systems Network philosophy. Sun offers a wide range of networking in the Open Network Computing environment. More than 120 licensees have adopted the Remote Procedure Call and External Data Representation for network services. The ONC environment includes NFS, which allows network-wide file access; the Yellow Pages, which provides a directory of services; and Remote Execution, which offers the ability to run applications on other workstations.

Sun also supports access to the network by IBM PCs and compatible systems through PC-NFS, allowing full personal computer sharing with Unix, VMS, and other operating systems. The non-Intel workstations can execute MS-DOS applications using a coprocessor board that can be shared across the network. Windowing systems that merge Neworkstation and X Windows 11 from MIT now make windows a network resource like files under NFS. The protocols are in the public domain, and Sun licenses the source code.

Sun also provides graphics independence through PHIGS, the Programmer's Hierarchical Interactive Graphics System. PHIGS runs on Sun-2, -3, and -4 workstations and is supported by the graphics accelerators.

The company is taking a different approach to operating systems, however. Instead of having three operating-system offerings, Sun is integrating AT&T's System V issue 2 with the Berkeley Standard Distribution (BSD) and incorporating the facilities of both into its new SunOS 4.0. Sun and AT&T are working together to define a new integrated operating system that combines the best of both with the intent to establish that operating system as a new standard. This action, in part, led Apollo and others to form OSF, as they were concerned about their lack of input as well as the effect of possible delays in getting access to the new operating system.

Silicon Graphics

Silicon Graphics, Inc., formed in 1984, has its own approach to the workstation market. SGI is totally committed to three-dimensional graphics in real time and has delivered expensive ($50,000 and up) workstations for three-dimensional work. SGI uses the best current technology and concentrates on adding its expertise at the graphical interface. In the fall of 1988, SGI announced a "low-cost" (under $20,000) three-dimensional workstation, the Personal Iris (see photo 3). SGI's products range upward to the 4D/240, which has four parallel processors, each operating at 20 MIPS and 4 MFLOPS. Amazingly, the high-end unit remains under $100,000.

The Personal Iris is a dynamic three-dimensional graphics workstation designed for personal use. The product emphasis from SGI is "three-dimensional graphics in color." As a result, all products from SGI include color graphics with a graphics processor as standard. When you compare the prices of other workstations with equivalent capabilities, the apparent high cost of the Personal Iris becomes the lowest-cost system for the capabilities provided.

The standard Personal Iris 4D/20 comes with a color graphics board of 1280- by 1024-pixel resolution with 8 color bit planes and 4 administration planes standard, expandable to 24 color bit planes, 8 administration planes, and a 24-bit z-buffer. The expanded model uses 56 bits per pixel, which contributes to the video RAM shortage. The graphics processor can render 4500 to 15,000 polygons per second and produce plain, antialiased, or depth-cued three-dimensional vectors.

The standard processor is no slouch, either. The CPU is a 12.5-MHz MIPS Computer R2000A RISC microprocessor with separate data and instruction caches and is rated at 10 MIPS integer performance. The optional floating-point processor provides 0.9 MFLOPS. Standard memory is 8 megabytes, expandable to 16 megabytes. Ethernet and TCP/IP are standard, as are two serial ports, a Centronics port, a small-computer-system-interface port, audio ports, and a VME slot. Note that a disk drive is not included in the base machine. The entire package will fit under your desk.

Operating-system software called IRIX is based on the Unix System V R3

Photo 3: *The SGI Personal Iris, a 10-MIPS workstation for $16,000.*

version with 4.3 BSD features and enhancements for real-time graphics. It is compatible at the binary-code level across the entire 4D workstation line. Standard software includes the SGI environment manager, graphics library, windowing system (Neworkstation, GL windows, and X Windows), and diagnostic software. Optional software includes five communications options and 10 productivity software packages for the programmer, including PC-DOS emulation and PHIGS.

Real applications show three-dimensional motion on realistic wire-frame or shaded models. Even a complex fighter model with thousands of line segments moves smoothly with shading, color, and variable illumination. This performance results from separating the physical computing in the main processor from graphical computations, which are done in the graphics processor. You can move and rotate the displayed image without requiring the main processor to do any computing.

The graphics system is, in fact, four processor subsystems dedicated to graphics functions. The four independent subsystems are the host interface (to a 40-megabyte-per-second graphics bus), a geometry engine, a raster subsystem, and a display subsystem. The geometry engine is a microcoded processor capable of 20 MFLOPS. It performs rotation, scaling, and transformation, and separately transforms surface normals. Then it clips vertex coordinates to a 6-plane bounding box and does the first stage of scan conversion. This information is passed to the raster engine, which is almost as complex, and from there to the display controller. There is actually more processor power in the graphics processor than there is in most basic workstations.

This system doesn't have just windows and pop-up menus—the menus open in smooth three-dimensional motion. Push a three-dimensional button, and it rotates and expands, exposing a new set of buttons at each level until you execute an application. There is, of course, a version of Flight Simulator that must be experienced to be appreciated. Motion is smooth, and the landscape is not just a bunch of lines.

NeXT

Next, and last in this case, is the new entry from NeXT. This is the long-awaited workstation for the education market from Steve Jobs. The NeXT workstation, sometimes called the cube, breaks ground in areas ranging from packaging

to optical storage. What has excited much interest is the price/performance, $6500 for a 25-MHz 68030 and 68882 combination with a high-speed digital-signal processor thrown in for good measure (see photo 4).

There are a number of significant developments here in addition to the price. First, the packaging reaches new levels of compactness, being a 1-foot cube with four slots, and the whole computer fits in only one slot. Second, it has a standard optical disk; it's not just a WORM (write once, read many times) drive but has full read and write capability, with good performance specifications for an optical disk.

Rounding out the hardware is 8 megabytes of memory, 12-channel direct memory access, and a 17-inch monochrome display with integral stereo, microphone, and mouse and keyboard connections. All in all, it's the simplest physical arrangement of hardware for its power.

Software technology was not left behind, either. Unix serves as the operating system, based on the Mach kernel developed at Carnegie-Mellon University. Mach is based on Unix BSD 4.3 but has enhancements in the areas of shared memory, interprocess communications, and potential multiprocessing. Networking is supported by the standard Ethernet and TCP/IP, with NFS from Sun Microsystems.

NeXT has added a graphical window-

ing interface called Workspace Manager to hide the raw Unix prompt. Windows and menus float in the workspace of the screen, and icons become transparent when they overlay other icons, which keeps everything visible. The windowing mechanism is based on Display Post-Script.

Development also has its share of new tools. The NeXT system includes an object-oriented preprocessor called Objective-C for the ANSI C compiler. Objective-C supports objects as groups of C procedures, and several libraries of ready-to-use objects, called kits, are provided. Kits are included for music and sound as well as for more prosaic objects. Also provided is something called Interface Builder, which supports the interactive design of user interfaces.

Several applications are bundled with the cube, including a word processor, a mail interface with voice-mail attachments, a searching program called Find, and the usual programmer tools. Beyond these are a group of educational tools, including a Webster's dictionary and thesaurus, Mathematica, *The Oxford Dictionary of Quotations*, and the complete works of Shakespeare. (Full details on the NeXT cube can be found in "The NeXT Computer" in the November 1988 BYTE.)

How does the cube stack up against the other workstations? On basic system specifications, it is more than a match for

Photo 4: *The NeXT computer, a 4-MIPS workstation for $6500.*

the typical workstation. But there are a couple of areas that may dampen the enthusiasm of those who want to do real work. First, there is the issue of distribution: It is currently available only to universities. Second, production and software won't be "ramped up" until the second quarter of this year. Third, no applications outside the bundled ones exist, although that should change within the next year. Fourth, current systems don't have a color display.

Where will the NeXT cube succeed? This innovative workstation is among the first of a new breed of "personal workstations" that will filter down into universities, small businesses, and, eventually, homes. It's clear by Jobs's example that current technology can give us a lot more than the personal computer delivers today. It's safe to say that the cube represents a 10-year advance in technology over the IBM PC and should come into its own in the 1990s. But, in my opinion, this year is not NeXT's year.

At Least Today, It's Possible

How does a standard workstation compare with one built on a personal computer? A few years ago, the answer would have been easy—there was no comparison. Today, the answer is more difficult, as the differences are blurred by technological and software advances. A few years hence, the question will no longer be asked. We will simply have a continuum of microprocessor-based computers.

Today, depending on where you start and how much you spend, you can turn your personal computer into a workstation. At some point, expanding a system to its upper limits is actually more expensive than buying one that's built for the task, but you can do it. Adding software on an incremental basis is also usually more expensive than getting it bundled, but again, you can do it. If you have the option and the funds, buying a workstation has its benefits. But building your own on a personal computer base is also a valid option. At least today, it's possible. ■

Bill Nicholls has a B.S. in physics from Notre Dame University and is the owner of BGW Systems (Puyallup, Washington). He can be reached on BIX as "billn."

Worth the RISC

*RISC technology is here, it's usable,
and it's low-cost now*

Trevor Marshall and Jane Morrill Tazelaar

I f there are to be any major improvements in workstation technology in 1989, they're likely to come from the reduced-instruction-set-computer arena. RISC chips will impact every area: embedded control, graphics, and mathematical calculations. Which areas they'll end up dominating really depends on how quickly the complex-instruction-set-computer designers react to innovative RISC architectures and incorporate those features that really work well into upcoming CISC chips.

What Is RISC?

RISC refers to the concept of a CPU that executes at least one instruction per clock cycle. With a CISC CPU like the 68020, it can take several clock cycles to execute an instruction. The RISC chips, however, execute most of their instructions in one cycle and make further performance gains with a technique called *pipelining*, which uses multiple execution units. The chip actually executes many instructions at the same time.

Some RISC chips, such as the Motorola 88100 and the Advanced Micro Devices (AMD) 29000, execute as many as five instructions at once, depending on the nature of the instructions. This allows the accesses to the I/O devices, the RAM, main memory. the disk, and the operator to be scheduled so that the CPU's execution unit can continue doing something useful even when the bus interface is performing external accesses. Although the advanced CISC CPUs also use a form of pipelining, their complex

instruction set makes it more difficult for them to achieve the same performance advantage.

Executing at least one instruction per clock cycle is the single thing that most characterizes the RISC chips. Indeed, it's one of the things that sets the Intel 80960 apart from the others: Not all of its instructions are one-word instructions, and thus they don't all execute in one cycle. There are some other characteristics that define RISC as well. But before discussing them, it will be helpful to look at the chips themselves.

The RISC Chips

The Clipper, developed by Fairchild, is probably the most mature RISC chip. Fairchild was bought by National Semiconductor, which sold the Clipper Division to Intergraph, a workstation manufacturer. Most designers don't regard the Clipper as an available family, but in fact you can buy Clipper chip sets from Intergraph.

Then there's Motorola: The 88000 is a three-chip set that consists of the 88100 CPU and two or more cache/memory management units, 88200s. The 88100 has built-in floating-point capabilities.

continued

Next comes AMD's CPU, the 29000; its associated floating-point chip is the 29027. Cache units are expected but are not yet part of the equation, so at this point the AMD chip set consists of the 29000 and the 29027 and is targeted at low-cost RISC applications.

MIPS Computer Systems offers the R3000 CPU and the R3010 floating-point accelerator. The MIPS family is best characterized by its (relatively) mature software base and close ties with a cache-memory architecture. Indeed, it's very hard to connect the R3000 to any memory system that does not use caching. This family of chips is thus restricted to more complex (and expensive) workstation applications.

In the Sun Microsystems' SPARC (scalable processor architecture) family, the first chip came from Fujitsu. Sun is currently shipping this CPU chip in the Sun 4 workstations, using the Weitek floating-point chip set to augment its arithmetic capabilities. The SPARC family is now becoming better known, and promising products have come on-line in 1988. Cypress Semiconductor has a SPARC implementation called the 7C600 family, which consists of the 7C601 CPU integer unit, the 7C603 memory management unit, the 7C608 floating-point controller, the 74ACT8847 floating-point unit (FPU) (actually made by Texas Instruments), and the 7C181 cache-tag RAM. Cypress is the second source that Sun licensed after Fujitsu had implemented the first member of the SPARC family. At this time, most of the Cypress SPARC chips are only just starting to become available.

The actual SPARC family, all those who have licensed the technology, includes TI, Cypress, BIT (Bipolar Integrated Technology), Fujitsu, and LSI Logic. Chips are actually on the shelf from Fujitsu and Cypress, and TI has its FPU out.

SPARC is probably the most difficult to comprehend of the whole lot. It's the only RISC chip currently being shipped by more than one manufacturer (although MIPS has now announced multiple sources for its parts). After the 80386, whose price was driven very high because you could only get it from Intel, people became wary of sole-source CPUs. When Sun brought out its SPARC chip in 1987, the company made it known that there would be multiple sources for the chip.

The multiple-source availability question really isn't much of a problem with RISC chips because there are so many of them and they are competitively priced. In fact, prices are already very inexpensive. There are 17 MIPS RISC machines currently available, and their cost, at least in large quantities, is around $10 per MIPS. That's a lot cheaper than an 80386.

IBM also has a RISC chip, but at this point it doesn't have the level of performance that the others are starting to show. In addition, it's not being marketed, at least not aggressively, as a chip set—it comes only in the IBM RT system—so we have little information about it to pass on to you.

The Intel 80960 chip's instructions are not fixed-length, but they are put on 32-bit boundaries. Such things as its optional displacements make the 80960 somewhat different from a typical RISC chip. It is a great chip for embedded control, which is what it's designed to do, but it doesn't yet have the level of performance that the other RISC chips have. Intel and AMD are both pushing their

chips at the lower-cost end of the market. They are each trying to say, "With my chip, you can design a system that is as fast as a workstation but at a fraction of the price." By and large they're correct, and they're keeping prices down by using innovative memory technologies.

Banking on Registers

Another characteristic of RISC CPUs is a *lot* of registers and very fast access to the register bank. Keeping the signals running around on the same silicon chip is much faster than going off the chip to main memory, disk, or the keyboard. This helps to enhance speed.

With chips like the Motorola 88000, AMD 29000, and MIPS R3000, you can both read and write to the register set in one cycle, so its memory is extremely fast. And since you have a large register set, you can do a lot of calculations without having to go out to main memory. SPARC implementations typically have 128 registers; the AMD 29000 has 192. The Clipper has 40—32 integer 32-bit registers and 8 64-bit floating-point registers. The Motorola 88100 CPU has only 32 registers, but 32 is still a reasonable number. The MIPS architecture has 32 registers. The Intel set has a pseudo-register, not true registers, but it's also 32 words deep.

The larger the register bank, the more calculations you can force to run in the high-speed internal registers—but you incur a penalty when you do multitasking. For instance, if you have a single task that uses 128 registers, and you switch to the next task (which also needs all the registers), you have to save them all to memory. Then, when you switch back to the first task, you have to restore them all. Thus the more registers the chip has, the higher the overhead it incurs when it needs to save and restore them.

AMD, Motorola, SPARC, and MIPS all have special features in their architectures to allow you to store and read back registers as fast as is practical. For instance, if you're storing to 128 sequential locations, you can do it much faster than if you stored individual registers one at a time. Both SPARC and AMD provide hardware support for doing multiple stores; a protocol defined in the wires that connect up to the RAM interface, to the real world, allows these high-speed bursts of register saving and restoring operations. But it's still an overhead you must take into account. That's the penalty for using a lot of registers, but the performance improvement tends to make up for it.

Table 1: *The RISC chips. The VAX MIPS estimate is usually derived by dividing the Dhrystone rating by 2000. As you can see, some manufacturers are conservative in their MIPS estimates while others are overly optimistic. (These numbers are from the manufacturers and are not the result of independent testing. Where MIPS numbers were not available, they have been calculated.)*

	VAX MIPS	Cache support	Architecture	Number of registers	Approximate Dhrystones/second
Motorola 88000	20	Yes	Harvard	32	46,000 at 25 MHz
MIPS R3000	20	Yes	von Neumann	32	42,300
AMD 29000	17	No	Harvard	192	42,000
Cypress SPARC	20	Yes	von Neumann	128	42,000 at 33 MHz
Intergraph Clipper	17	Yes	von Neumann	40	35,000
Sun (Fujitsu) SPARC	9.5	Yes	von Neumann	128	19,100 at 16 MHz
Intel 80960	7.5–10	No	von Neumann	32	13,000

This penalty doesn't come into play with CISC chips, which typically have only a few registers. For example, typical compilers used with a 68020-based machine use a maximum of 8 general-purpose registers and 4 floating-point registers—the rest are scratch registers. In current CISC technology, you don't gain much speed by using registers, so most compilers don't overuse them.

Compilers for RISC chips, however, use lots of registers to get the best performance, the highest number of Dhrystones per second, and so on. For instance, the MetaWare compiler for the AMD 29000 uses all 192 registers. And it has routines called spill and fill—spill when the registers overflow to send them out to RAM and fill when they underflow to pull them back in.

This huge register stack acts as a window into RAM. Using it is a very efficient way of writing compilers, and most RISC compilers use it. The Sun compiler, however, has something different: a register-windowing scheme. It tends to have the same effect; it's just a different way of describing things. With the register-window concept, you must define a fixed number of variables that you will pass between procedures. For instance, if you only need to pass three registers between procedures and you switch your level of register windowing, you're going to waste five of the eight defined registers. None of the other chips has this windowing scheme. The inefficiency it causes seems to be relatively minor, probably only the last 10 percent of difference in speed.

Cache in the Chips

A computer system today is designed with a CPU and a large (slow) main memory; in the middle, between the two, is cache memory. Cache is a high-speed memory that buffers between the high speed of the CPU and the low speed of main memory. All the conventional workstations use cache memory.

With RISC machines, cache performance becomes critical. You can have a CPU that's running three, four, or five times faster than a CISC chip, but you still have the same-speed memory, the same DRAMs. Therefore, the function of the cache is to buffer. How you make the operation of the CPU more efficient with the slower main memory becomes of paramount importance.

There are ways of designing unique memory systems that are structured differently from the conventional CPU, cache, and memory setup. You can connect the CPU directly to specially configured main memory. While you won't achieve quite the same level of performance that you would with a cache, you're talking about only a 10 percent to 20 percent difference (see the article "Real-World RISCs" in the May 1988 BYTE). Yet your cost difference will be significant. The cache is becoming one of the most expensive items on a workstation, be it CISC or RISC.

Both the Sun 4 and the MIPS machine have fairly good caches, typically 128K bytes of cache RAM each. That's high-speed static RAM, and it tends to be much more expensive than several megabytes of DRAM.

The MIPS people stress close integration with the cache. In fact, their whole system is designed around it. Basically, they will sell you an agreement whereby they tell you how to design the cache. They give you all the details, down to
continued

which chip you connect to which chip. They probably have the most advanced system of any RISC cache to date. The MIPS chip is very closely integrated with its high-performance cache design.

The Sun 4 uses an expansion of the Sun 3 cache technology. Motorola has cache memory management units (MMUs), the 88200s, which are very expensive (several times the cost of the CPU) but allow you to connect low-speed memory and still get tolerable performance. Each MMU contains 16K bytes of high-speed cache memory. They let designers look at the RISC CPU as a "black box" and not have to worry about high-speed memory designs.

For example, if you're designing a system around the AMD 29000, you have to look at all the timing diagrams. You have to understand what's happening in the internal pipelines. You have to know when the system's going to be asking for another instruction and when it's going to be accessing data. Once you understand these things, you can mold the memory to what the CPU does.

With Motorola's three-chip set, the cache memory buffers all that. All you really have to understand is how to connect to a cache memory. You don't have to be concerned about what the CPU is doing. The cache memory handles that interface for you, so the "black box" approach is an easier level of design. The Clipper also allows you to design at this level.

Motorola and Intergraph both sell integrated cache chips, which have the RAM as well as the control circuitry inside; they're all you need to implement a cache. It's not as good as the MIPS cache—and it's certainly not as large—but it's low-cost and convenient.

Both Intel and AMD have concentrated on designing interfaces for their chips so you can couple them to low-cost memory, producing a low-cost, high-performance system. To give you some idea of the performance levels you can get, it's possible on low-cost $3000 to $4000 hardware to achieve 30,000 Dhrystones per second. Intel isn't even talking about cache. AMD is working on a cache unit—but it's very effective when connected directly to RAM. And Fujitsu uses the Sun cache technology.

Cypress has designed its family around a cached-memory architecture. The company provides two cache-support chips to simplify high-speed cache design, both the RAM itself and the special cache-tag RAM. You get schematics on how to connect them into a system. Cypress is also emphasizing the cache

and has made no attempt yet to allow its SPARC chip to be effectively connected directly to main memory.

Comparing the Chips

The first major comparison point for these chips is whether they stress cache or noncache technologies (see table 1). Any of them can be used in either mode, but with varying degrees of difficulty. AMD and Intel are pushing noncache systems; all the other manufacturers are pushing cache systems. That's the first distinction, cache or noncache.

The next distinction is whether the chips have Harvard or von Neumann architecture. A von Neumann architecture has one external 32-bit bus that is used for both data and instructions. A Harvard architecture has two separate 32-bit buses, one for bringing in the instructions and one for the data. A Harvard architecture allows you to bring data and instructions to the chip at the same time. Of the RISC chips, the Motorola 88000 and the AMD 29000 use the external Harvard architecture.

So the capability is there, in theory at least, to double your memory-interface speed. It doesn't work that way, however; you probably multiply it only by one and a half. Nevertheless, the Harvard architecture significantly raises the interface speed without significantly changing the interface cost. This architecture allows you to customize the type of memory you connect, since it has to work well only with data references or with instruction references, not both.

One similarity shared by the families is that they're all vying for the same marketplace, and when all is said and done, the CPU chips are similar in price. In essence, the choice of a CPU won't, by itself, affect the final cost of a workstation very much, because RAM is now becoming the biggest part of the cost. A workstation today will have 8 megabytes or more of RAM, and all the CPU chips are relatively inexpensive when compared to the cost of 8 megabytes of RAM. Thus, the choice of a CPU is usually not based on raw chip cost anymore. Some workstations still have expensive FPUs, but since they're all competing for the same piece of the pie, that too must change.

One way to compare RISC chips is to look at which ones have been around the longest and which are just emerging. For example, the Clipper is the most mature while the Motorola 88000 doesn't really exist yet. There are a few chips in the hands of developers—and they work—but they're not shipping in quantity yet.

The AMD family has been shipping

for nearly a year. The MIPS family has been around for quite a while now, too—at least the R2000 has (for maybe 2 years), and the R3000 is derived from that. A few minor changes were made, but the R3000 has essentially the same genealogy and longevity as the R2000. The MIPS chip is the most mature RISC chip available except for the Clipper. It's even more mature than Sun's SPARC.

Sun's SPARC technology is about a year and a half old, so it's fairly mature, too. One way to judge the maturity of a chip is by the languages available to support it (e.g., C, FORTRAN, and Ada). This is particularly true for RISC. You don't want to be developing code on tools that are themselves still in beta test. Sun and MIPS have lots of software tools in place; AMD has some; and Motorola has very few. The Intel 80960 is relatively new also; it has only recently started to ship, but it *does* exist and it *is* shipping.

The venerable Clipper isn't always remembered because it's not from one of the major semiconductor houses—Sun's chips are from Fujitsu, for instance, and Cypress is quite big in CMOS. However, the Clipper performs reasonably, certainly at about the same level as the SPARC. It consists of a CPU and two cache units mounted on a circuit board. It's marketed with a "black box" concept like Motorola's; Intergraph tells you how to connect to the circuit board and not to worry about what's on the board—it'll take care of itself. And it does.

A Benefit Performance

We used the Dhrystone benchmark to compare the various RISC chips. The Dhrystone measures a chip's ability to handle integer operations, particularly string operations such as those that occur in compiling and the searching of databases.

The AMD, MIPS, and Motorola chips all run at about 42,000 Dhrystones per second, with peak performance up about 46,000. For comparison, the 68020 CISC chip rates about 5500 Dhrystones, and the 68030 has tested as high as 7000 Dhrystones. The 80386 can do 9000 Dhrystones but has an advantage for the Dhrystone test over the 680x0 family due to special string-manipulation instructions. You can see that the differences among RISC chips are much smaller than the gap between CISC and RISC.

The Sun 4 SPARC (which is a Fujitsu SPARC in the Sun 4) comes in at around 19,000 Dhrystones, so it's not quite up to the general-architecture RISC chips yet. That may change. The new SPARC implementation from Cypress is faster, and

there are some gallium-arsenide SPARC implementations coming from a company called Prisma in Colorado Springs, Colorado. A group at McDonnell Douglas is also working on such a chip, and so is BIT.

The Cypress SPARC chip is said to run at about 42,000 Dhrystones at 33 MHz. The Clipper comes in at 35,000, and Intel's 80960 currently has the lowest performance rating, around 13,000 Dhrystones. The top three performers are MIPS, Motorola, and AMD. Then comes Cypress and the Clipper. The next best performer is Sun, with Intel currently bringing up the rear. Intel promises more performance next year.

Performance, however, is relative to your application. The Intel chips are aimed at embedded control, such as laser printers. You want a very fast CPU to run your laser printer, but you don't really care whether that same chip will run spreadsheets and databases. Intel is aiming its RISC family at that market so a high Dhrystone performance rating is not crucial. However, all the chip makers, certainly both Motorola and AMD, are also targeting that market, because it's a big one.

The workstation marketplace accounts for only a small fraction of the total number of CPU chips sold. But all the manufacturers are trying to have a presence there because it allows them to project their chips toward the embedded-control designers.

A Graphic Demonstration

The RISC chips will definitely have an impact on graphics. They offer an alternative to the special-purpose graphics engines that companies like Silicon Graphics have developed. They are computing power, pure and simple, regardless of the application. Their designers have put their entire computing knowledge into creating ultrahigh-speed CPUs. What they have come up with are general-purpose CPUs that in many cases exceed the speed of special-purpose CPUs even within the special-purpose application. For example, AMD has a graphics chip, called the quad-pixel data-flow manager (QPDM), which was specifically designed and optimized to do graphics manipulations. However, AMD has found that the 29000 RISC chip can outperform the QPDM in most graphics operations. These RISC chips are extremely high-speed devices and will surely have an impact on graphics at all price levels.

AMD is talking about pricing its RISC chips at $99 next year. Given that sort of

---

price, the chips could even start coming out in low-end personal computers, giving them a graphics capability they've never had before.

A Definite Impact

If you're sitting back waiting for a winner to appear in the RISC sweepstakes while you can buy a chip that will do the job you want to do right now, it's probably not worth waiting any longer. All RISC technology is scalable. There's going to be leap frogging—a new chip will come out in March, and another in July,

---

that sort of thing. But the quantum leap in RISC technology has already occurred. All the top chips—the MIPS, the AMD, the Motorola, and even the Clipper—are already achieving the performance levels (within 30 percent or so) that they will have by the end of this year. RISC technology is here; it's usable; and it's low-cost now.

Furthermore, a workstation's performance is primarily dependent on the designer, not on the chip. Take, for example, the MIPS family of machines. If you compare the workstation from Silicon Graphics (which uses a MIPS CPU) and the workstation from MIPS Computer Systems (which uses a MIPS CPU), you'll find that the machine from MIPS has a much higher performance than the one from Silicon Graphics. Why? Because MIPS uses more effective caches; it's a more expensive machine; it has been designed from the ground up to be a high-performance computer. The Silicon Graphics machine has been designed to be a high-performance, low-cost *graphics* machine. Thus, some of the computational ability of the MIPS CPU has been traded off for those graphics features.

If you use the MIPS computer (with 64K bytes of static RAM cache) and the low-end Silicon Graphics machine (with a much simpler architecture) to run a benchmark that has a lot of data manipulation, like the LINPACK algorithm, you'll find as much as a 50 percent speed difference. This can be very significant, and it's probably the best example we can give of how critical the design of the system is. So if you're buying a workstation, you should evaluate what's being offered as the workstation, run your code on it, and work out what you want to buy based on what the system will do with *your* code. The underlying chip is far less important.

The levels of RISC performance are so high that applications that weren't practical on workstations in the recent past will be in the near future. Applications, such as three-dimensional rendering, that took impossibly long times to compute on last year's workstations can now complete in one-tenth the time. Obviously, RISC technology is going to have a profound effect on the way we use computers. ∎

Trevor Marshall is chief engineer at Yarc Systems Corp. in Thousand Oaks, California. Jane Morrill Tazelaar is BYTE's senior technical editor, in depth. They can be reached on BIX as "tmarshall" and "janetaz," respectively.

How Fast Is Fast?

*MIPS? MHz? MFLOPS? The ultimate measurement
is probably your own workload.*

Bill Kindel

Our generation is obsessed with speed. Whether we're thinking about automobiles or computers, we're sure to ask, "How fast will it go?" With cars, miles per hour provides an unambiguous answer. But with computers, the best answer is, "As fast as it can."

Measuring a computer's performance is not as simple as we'd like. The most popular measures of relative performance are, at best, imprecise; at worst, they can be misleading and meaningless.

I'll discuss some of the pitfalls inherent in assigning performance numbers to computer systems. Once you are familiar with some of the factors that contribute to performance, you should be better prepared to weigh their significance when presented with benchmark and performance test results.

Timing the Instructions

The most common units of performance are usually derived from measuring the time required for the computer's processor to execute some arbitrary set of instructions. Such measurements are expressed as KOPS (thousands of operations per second), MIPS (millions of instructions per second), or MFLOPS (millions of floating-point operations per second). These measurements are most meaningful when comparing processors from the same family, because the instructions implemented on dissimilar processors can vary dramatically.

As an example of where instruction timings can be misleading, consider how three different processors might move a character string from one place in memory to another:

• Processor A has a fairly typical instruction set, which includes load/store operations for 1 or more bytes at a time and a broad selection of addressing forms. The data movement would be done by repetitively loading chunks of data from source memory to a register and storing them from the register to the destination memory. At the end of the loop is a completion test, which becomes part of the timing.

• Processor B implements a richer instruction set, including special instructions for character-string manipulation. In such a case, a single instruction can be executed to move the string from the source memory location to the destination without needing to use any of the general registers. While the single instruction appears to be quite long-running compared with individual load and store instructions, it should be able to move the character string considerably faster. A 5-to-1 speed advantage is not uncommon.

The CPU Effect

At the center of the system is an oscillator, the *clock*, which is often crystal-controlled. Some personal computers use a single clock for everything from the processor to the display. Workstations and larger systems may have separate clocks for each major subsystem. The clock frequency can be halved, quartered, and further divided to provide the correct frequencies to the various chips. On the other hand, the same result can be achieved by multiplying the power-line frequency to raise it (with less precision) to the desired frequencies.

The CPU contains several logical subunits. These can be implemented as discrete components or integrated into a single chip. Some number of registers are used to hold pointers to data in memory or intermediate computational results and other processor-specific information. A specialized subunit decodes the operation code. The address translation occurs in another subunit, which provides the physical addresses of memory to be fetched or stored. One or more *operation* subunits actually perform the instruction, sometimes splitting off instructions based on the type of data on which they operate.

Machine instructions are executed in several steps, varying by the implementation and the instruction being executed (see figure A). The first step is to read the operation code from memory. This is then decoded, and the number and type of arguments are determined.

Arguments almost always follow the operation in the instruction stream. Each argument is read from memory and interpreted in two more steps. In most cases, arguments contain the addresses of the data on which the instruction operates. Sometimes "immediate data" is provided, which can be used without further interpretation. In other cases, the argument points to a location in memory that holds the address of the actual data. Then it becomes necessary to do yet another read from memory. Data addresses are often the combination of the address or offset given in the argument and the contents of one or more processor registers. All these calculations occur before the actual operation can take place.

After calculating the data addresses, any input values must be read from memory (and processor registers) into the operation unit. Finally, the actual operation is performed. Any registers or memory locations that are changed are updated from the operation unit. Instructions that loop through memory may read and write many locations in sequence.

Several techniques are used to speed up processing. The first is cache memory, which can greatly reduce the time the processor spends waiting for memory accesses. Various strategies can be used, but all depend on the fact that a high percentage of memory accesses are to a relatively small number of different addresses.

The cache, which is logically adjacent to the processor's operation unit, keeps copies of the current values of the most recently accessed memory locations. When memory is read, the cache is searched first. Only if the cache does not contain the address's data is an actual memory read requested. While a read from memory requires several clock cycles to complete, cache "hits" are nearly immediate.

A second significant technique is called *pipelining*. Because different parts of the processor are actually used

Figure A: *Simplified processor flow. These several steps of machine-instruction execution may vary depending on your machine and the instruction.*

• Processor C is a RISC (reduced instruction set computer) version. To gain speed, RISC processors simplify both their instruction sets and their addressing modes. This speeds up both the instruction-decoding and argument-preparation phases of each instruction. RISC processors also gain speed by using register-to-register operations rather than operations that access memory. As an aid, RISC processors are sometimes provided with more computational registers than non- RISC processors have.

In the wide-open world of MIPS comparisons, such processor mismatches are common. The search for a representative sample of instructions to time often set-

for the execution of the various steps, several instructions can be in different stages of execution at the same time. The "depth" of the pipeline is the maximum number of concurrent instructions that can be accommodated. Pipelines operate under the basic assumption that the instruction stream is located in consecutive locations in memory, so instructions are "prefetched" and decoded before they are needed.

A third feature that can improve performance is virtual memory. Virtual memory has been described as "making a machine with a lot of memory appear to have unlimited memory." This is of particular value to applications that use very large amounts of in-memory data. Code and data are assigned "virtual addresses" that are mapped by hardware to physical memory locations. In most cases, memory is divided into fixed-size "pages," typically 256 to 4096 bytes each.

Instead of forcing the application to manage movement of data between disk and memory, the system software takes over responsibility for doing so on demand. When a memory reference is made, the hardware determines if the page has been mapped to a physical location in memory. If not, a "page fault" occurs that system software must deal with before execution of the application can continue.

The fourth feature is the use of coprocessors to execute certain classes of instructions. The most common cases are the handling of floating-point operations and array processing. A coprocessor is a separate processor, usually with its own registers and internal features, which is attached to the system bus. When the CPU decodes an instruction that can be handled by the coprocessor, a signal is sent to the latter to cause it to perform the operation while the CPU waits. If the coprocessor isn't configured, a fault occurs. This usually results in software emulation of the coprocessor instruction by the operating system's fault handler.

tles on a minimal common subset, which may favor one processor over another.

A similar problem exists for comparisons between floating-point processors. Depending on internal architecture, it is entirely possible that one floating-point

processor might be optimized for execution of single-precision instructions while another is optimized for double-precision. For example, the first processor might be rated at 4 MFLOPS in single precision and 2 MFLOPS in double precision; the second processor could generate 3 MFLOPS in either mode. You could compare this situation to gasoline octane ratings. The ratings on the pump are actually the average of the "motor octane" and "research octane" ratings and represent neither one accurately.

Comparing Clocks
Another approach is to compare processor internal clock speeds, which are almost always expressed in MHz. The advantage to this approach is its simplicity; clock speeds are published for various systems' microprocessors. Ignoring any other factors, you can reasonably expect that increasing the frequency of the system's clock will result in a commensurate increase in processor speed.

The fallacy in comparing the clock frequencies for dissimilar processors is that they have little to do with each other. The number of clock cycles used by a processor depends very much on its actual implementation. The more exotic microprocessors, such as the Motorola 680x0 and Intel 80x86, require more clock cycles for address preparation and instruction decoding than their simpler cousins. As a result, a 68000 operating at 7 MHz in a Macintosh, Atari ST, or Amiga is not necessarily seven times as fast as an 8-bit 6502 running at 1 MHz.

Performance Constraints
These processor performance measurements are "best-case" situations subject to a number of constraints. Those who play the numbers game have a tendency to use these theoretical numbers for comparison without derating them to compensate for the fact that real workloads are not ideal. Here are some common factors that affect performance:

• *Pipeline breaks*: Among the techniques used to accelerate processing is the overlapping of instruction executions. Most high-performance processors use a pipeline technique wherein several consecutive instructions are in various stages of execution—from decoding to argument preparation to the actual operation—at the same instant. Transfers of control cause the partially processed instructions still in the pipeline to be discarded; such "breaks" greatly limit the performance of tight loops. (For a more detailed look at central processor func-

tions that affect performance, see the text box "The CPU Effect" at left.)
• *Memory speed*: If a processor is significantly faster than its system's memory, the processor will waste clock cycles waiting for data to be delivered. The processor can outpace memory by requesting more bits at a time than the memory is able to transfer. A 32-bit-wide request that can be satisfied all at once by 32-bit memory would have to be converted into a pair of consecutive 16-bit requests if the system were equipped with 16-bit memory. The processor must wait for both before proceeding.
• *Memory cycle stealing*: The memory-speed problem is made worse by configurations in which memory is shared between the processor and another high-speed active element, such as the video display chip. Because the processor and video chip can't access memory simultaneously, a means of sharing memory-access clock cycles must be implemented. The crudest form is to assign alternate cycles to the processor and video chip. It's common for the video chip (which must stay synchronized with the display) to be given the ability to "steal" cycles from the processor whenever it needs them.
• *Software inefficiencies*: Software doesn't automatically take advantage of the performance features built into the hardware on which it runs. While it is common for operating systems to check for such installed features as floating-point coprocessors and extended instruction sets, application programs usually don't. Unless performance becomes a problem, any optional features often will be ignored. This results in fewer trouble calls to the software developers at the cost of lost performance on high-end systems.
• *I/O and other system bottlenecks*: Very rarely is the CPU allowed to operate at full speed. There is almost always some limiting factor outside it, such as waiting for disk I/O (including reading a page in a virtual memory system). Even the normal operations of related components can cut into the processor's throughput. A real-time clock generates interrupts on a regular basis, which causes the CPU to set aside its processing, handle the interrupt, and resume.

Operating-system features, such as multitasking and virtual memory, provide solutions to the sharing of critical resources (i.e., CPU and memory) among multiple programs executing concurrently. This is not without cost, however; while total system throughput typi-
continued

cally increases, the performance of any given program is usually degraded.

In the case of virtual memory page faults, the performance impact is greatly increased by an embedded disk I/O request. If, for example, an application is actively using memory locations on 50 different pages, but only 40 physical pages are provided in its "working set," then 20 percent of its memory references are likely to generate page faults—and a subsequent performance degradation.

I/O and memory bandwidth are also factors. The time required to perform a single memory or disk access is subject to interference from other accesses. A disk unit can do only one request at a time, so a queue could form that forces new requests to wait for completion of those already outstanding. Even with adequate disk drives and memory, various system components have limits on the number of operations they can perform per second.

Other bottlenecks exist within system software. Any system service that is used heavily has the potential for limiting total throughput. As system complexity has increased, the probability of contention between tasks for critical resources has also increased.

The "B" Word

Benchmarking has long been the norm for large computer acquisitions. A few such benchmarks have become standard measures of performance. Those who are concerned with computational performance have relied on the Whetstone benchmark (which also has a double-precision version) to rate system performance in Whetstones. Another favorite is the Sieve of Eratosthenes, which measures the efficiency of a compiler's generated code as much as it measures raw CPU performance.

Such standard benchmarks are most valuable to those whose needs match the benchmark. But for the rest of us, they are only indicators like other metrics.

I believe the best way to measure a system's performance is to load it up with the work you intend to do on it. This is the basis of a benchmarking process that will be most meaningful to you. *You* define a representative workload and then evaluate it on all the different systems under consideration. While not always possible, this sort of testing provides you with results that are virtually indisputable—from *your* point of view.

Test It Yourself

MIPS, MHz, MFLOPS, and the assorted other metrics all attempt to give a basis for comparison of various systems' performance. Each is valid within the constraints under which it is calculated; none is valid for all systems under all circumstances. As they say in the car ads, "Your mileage will vary." Take those results as a first-order approximation only.

The only real way to predict how well a system will perform for you is to test it yourself. Besides, a hands-on test will tell you a lot about both the system and the people you'll turn to for assistance. ∎

Editor's note: *For a further look at the muddled world of performance measurement and our attempt to make evaluations meaningful, see "Introducing the New BYTE Benchmarks," June 1988 BYTE.*

Bill Kindel is a principal software engineer for Digital Equipment Corp. in Boxborough, Massachusetts. (The opinions expressed in this article are those of the author, and not of his employer.) He can be reached on BIX c/o "editors."

Art + 2 Years = Science

*The state of the art
in workstation graphics today*

Phillip Robinson

I t's easier to say "state of the art" than to define it. It's easy to stare in wonder at a three-dimensional model of an airplane on a Personal Iris workstation from Silicon Graphics or to sit in stunned silence while a Tektronix workstation projects a stereographic image of the space shuttle. And there's no doubt these machines rate a place in any discussion of the state of the art in workstation graphics.

But there are other applications that also show the potency of workstation graphics, including powerful computer-aided publishing (CAP) and fast two-dimensional electronic CAD packages. These applications may not show the sophisticated and subtle color shading of a three-dimensional modeling program; however, they do display large areas of WYSIWYG text and graphics and can redraw complex overlays of colorful two-dimensional circuit diagrams within fractions of a second.

How do you determine the state of the art in graphics? What are the current products, concerns, and technologies for graphics on these 32-bit computing platforms? These are the subjects I'll be discussing.

First, I'll define the specific region of computing price and performance I'm talking about. After some drastic cuts in 1987 and 1988, the prices of 32-bit workstations dropped down to the top prices for personal computers. Apollo, Sun, and DEC (Digital Equipment Corp.) all offer entry-level workstations for around $5000.

These price cuts stunned workstation devotees, who expected 32-bit machines to remain at the $20,000 level, leaping in performance but not dipping much in price. Clearly the competitive threat from microcomputers was serious. Most of the lowest-priced machines, however, were diskless "nodes" intended to work on a network.

For $5000, you get a 32-bit CPU (typically a 68020, as in the Mac II), 4 megabytes of RAM, a 15-inch monochrome monitor, some version of Unix along with a windowing interface, and an Ethernet interface. Add a local hard disk drive, a color-display controller, and a large color monitor (typically 16-inch or 19-inch), and you're looking at $8000 to $10,000 or more.

As the RAM, hard disk drive, and monitor grow, the price grows too, up to $20,000 or $25,000. Some companies have different model numbers within that range ($5000 to $25,000) that substitute a different, faster CPU into the system. Hewlett-Packard, IBM, Silicon Graphics, Tektronix, and many other companies join the market when it is defined as extending up to the mid-twenties.

continued

Then comes a gray area. The personal workstations that are quite capable of two-dimensional work, such as CAP, computer-aided software engineering (CASE), electronic CAD (ECAD, or EDA for electronic design automation), and mechanical CAD (MCAD), peter out when assigned tougher tasks.

They give way to mid-range workstations that cost from $35,000 to $60,000 and are capable of three-dimensional graphics like those used for solids modeling, animation, simulation, and image processing. The mid-range systems boast more possible on-screen colors, graphics accelerators to speedily figure new pixel positions and hues, as well as more RAM, CPU MIPS (million instructions per second), and disk capacity.

Above the mid-range are the high-end or superworkstation systems that cost $80,000 to $100,000 and more. These are packed with yet more MIPS, specialized coprocessors for graphics, and even new architectures, such as multiprocessing CPUs. By then, you're talking about the gray area between workstations and minicomputers.

It's important not to forget terminals, which are still used in many circumstances. There are graphics terminals that contain special graphics-acceleration hardware. These can download primitive instructions from a mainframe or minicomputer and then perform the actual display work locally.

Looking up from PCs

Graphics on the IBM PC, Macintosh, Amiga, or Atari start from scratch, meaning to some people no more than the ability to display color on the screen. Progressive enhancements increase the number of colors available, the number of pixels on the screen, and the systems software for manipulating graphics primitives. Since the birth of the Mac, systems software often encompasses interface graphics such as windows, icons, and pull-down menus.

Few personal computers, however, even 80386 speed-burners or Mac IIx systems with multiple megabytes of RAM and floating-point processors, have any sort of hardware for accelerating graphics performance. They stick to the tried-and-true line of improving graphics performance by improving CPU power—that is, adding more system MIPS to speed the calculations involved in graphics. They may have special chips for sending the graphics bits to the display, but the actual algorithms of graphics—transforms, clipping, scaling, shading, and so on—are handled almost entirely by software.

Some PC and Mac specifications match what you'll find in workstations. For example, the Mac II with Apple's own color-display board can produce up to 16.7 million different colors, displaying 256 at a time. That's the same as Apollo's personal workstations, for example, and is equivalent to 8 bit planes of memory for color (i.e., 2^8 colors). When a 24-bit color board is added to the Mac II, it can even display as many colors at a

F*ew personal computers have any hardware for accelerating graphics performance.*

time as the Personal Iris workstation.

There are also a few add-on boards for PCs and Macs that use standard graphics coprocessors, such as the Texas Instruments 34010 and the Intel 82786. Such chips can take graphics instructions from applications software and directly execute them without resorting to complex systems-software algorithms.

The Amiga comes with its own set of graphics coprocessor chips, and the Atari systems can take advantage of Atari's proprietary "blitter" chip for speeding graphics. The speed improvements from such accelerators can be enormous—as can the general graphics-calculation improvements that come from having a floating-point mathematics coprocessor. A personal computer with an added graphics coprocessor and a top-notch floating-point unit, such as the Weitek mathematics processors, may have more coprocessor firepower than many low-end workstations can boast, although adding the extras can also boost the *price* of the complete personal computer system beyond that of a workstation.

The systems at the top of this entry level can have proprietary graphics-accelerator boards or chips, high-speed CPUs, lots of memory, and other fancy hardware features, such as z-buffers for hidden-line removal. Most workstations also support graphics standards, such as PHIGS (Programmer's Hierarchical In-

teractive Graphics Standard) and X Windows under Unix, that provide some compatibility between different companies' systems and even between different models from a single company. Personal computers with extra coprocessor boards may be cut adrift from such compatibility, needing special drivers for each CAD or CAP program (typically including AutoCAD, VersaCAD, PageMaker, or Ventura Publisher) that they are able to run.

Frame Buffers, Z-Buffers, and Resolution

Most computer-graphics systems today use a display system with a "frame buffer" of RAM memory. This memory can be separate from the conventional RAM or can simply be an assigned part of it (as is the case in the IBM PC). Graphics information from the application program or systems software is converted into patterns in the buffer that change with program demands and is fed from the buffer to the screen at a regular, timed rate.

That rate must be quite fast in workstations, because there are so many pixels on a workstation display. The screen-refresh rate is physically determined by the display hardware and is kept high enough to eliminate bothersome flicker. Combine a large number of pixels with a high refresh rate, and you're looking at putting each pixel up every couple of hundred nanoseconds—leaving little time to calculate changes for a pixel and enter them into the frame buffer. This problem can be tamed a bit by adding double buffering to the system.

The "megapixel" is a common centerpoint for workstation resolution—1 million points on the display, although the screens range from somewhat under that (1024 by 800 pixels) to well over it (1600 by 1280 pixels). The higher resolution is typically used only in monochrome applications like publishing, where large pages need to be displayed on a single screen, showing both small text fonts and gray-scale pictures with their different monochrome intensity levels. Higher resolutions are not often mentioned as a prospect for workstations: More processing speed and colors are in much more demand.

A frame buffer with a single bit plane can only contain monochrome images: each bit is either a 0 or a 1 and translates into either an "off" or an "on" pixel, respectively, on the screen. Gray-scale images or color can be represented by multiple bit planes logically laid on top of each other with more than 1 bit of RAM
continued

Putting Workstations to Work

Thousands of applications programs exist for graphics workstations. Here are a few that exemplify the abilities, fields, and prices of the genre. They range from those that represent that two-dimensional foundation, systems-level software, to those at the true "state of the art" in animation and rendering.

Systems Software

The 68030-based workstation from NeXT runs the Mach version of Unix and comes with an object-oriented software "environment" called NextStep. The NextStep environment combines both applications developers' tools and a user interface. NextStep has four components—the Window Server, the Workspace Manager, the Application Kit, and the Interface Builder—which provide windows, menus, icons, and a simple process for putting building blocks of code together to form new programs.

NextStep uses Display PostScript as a device-independent graphics "library" for sending graphics information to the display. IBM has licensed NextStep for use on its own workstations. NextStep doesn't require any particular graphics hardware. It represents the first major endorsement of Display PostScript as a graphics software standard.

CASE

Athena Systems makes a CASE and simulation tool called Foresight for defining real-time software and hardware systems. It runs on Sun workstations and uses graphical block diagrams and data-flow charts to represent physical processes.

Foresight can animate models such as heating systems, simulating their behavior to help find design flaws. Blocks can be connected hierarchically or in parallel, simulating sequential or parallel processing. Foresight is a two-dimensional application. It costs more than $23,000.

CAP

Frame Maker, from Frame Technology, is another prime example of a two-dimensional workstation-graphics application. It is a CAP program that can integrate text and graphics on pages for books, newsletters, specification sheets, documentation, or just about any other sort of publication.

It can run on the least expensive graphics workstations—such as the Sun-3/50—and includes its own word processor, page-layout program, spelling checker and corrector, indexing ability, drawing program, and PostScript output routines. It costs about $2500.

EDA

Electronic design automation, also known as electronic computer-aided design (ECAD), is a major part of the two-dimensional graphics market. Software abounds for designing and simulating VLSI chips, circuit boards, and even complete electrical systems. Mentor Graphics offers a variety of programs for drawing and "capturing"—that is, converting into electronically logical files—circuit diagrams and then simulating the behavior of those circuits. These programs require color and run on a variety of workstations, including Tektronix, Apollo, and Sun.

Compact Software offers programs for designing and simulating radio-frequency and microwave circuits. Microwave Harmonica, for example (see photo A), can analyze and optimize any microwave nonlinear analog circuit under single- or multitone excitations. It has an interface to Mentor Graphics software and many other third-party programs to use schematics captured in those environments.

Microwave Harmonica runs on Sun, Apollo, HP, and MicroVAX workstations with two-dimensional graphics, typically using an 8-bit-plane system with 4 to 8 megabytes of RAM. It can link to the vector-processing abilities of supercomputers. Microwave Harmonica, and Compact's other similar programs for linear and gallium-arsenide circuits, use graphics for both the user interface and the output of results. Compact also offers an AutoArt II program for layout of the physical circuits. Some of the software also runs on personal computers.

The workstation software prices range from $7000 to $25,000 per user, depending on the number of users at a company. Efficient execution of EDA programs requires fast RAM architectures, timely bit-blit operations, and software routines that permit high two-

continued

Photo A: *A two-dimensional graphical display from Microwave Harmonica, a program for circuit design and simulation. (Photo courtesy of Compact Software.)*

Photo B: *A three-dimensional wire-frame display from Anvil-5000, a design and drafting program for mechanical engineers. (Photo courtesy of Manufacturing and Consulting Services.)*

dimensional vector-per-second drawing rates.

Scientific Simulation

Creare makes a program called Fluent that models fluid-flow, heat-transfer, and combustion dynamics. It uses numerical-modeling techniques to solve various complex equations of fluid dynamics, and then uses graphics to plot the results of those computations. Fluent is used in automobile aerodynamic research and in studies of cooling towers and turbine machinery.

Fluent runs on MicroVAX, Tektronix, Apollo, Sun, Silicon Graphics, supercomputers, and mini-supercomputers (like those from Alliant). Prices depend on CPU power, ranging from around $9600 a year to $40,000 a year (on a Cray supercomputer). A typical installation might use a Sun-3/260 or Sun-4 with color and a floating-point processor. Because the computation is so central, high floating-point performance is more important than any particular graphics processor.

Three-Dimensional CAD

MCS (Manufacturing and Consulting Services) makes Anvil-5000, a general-purpose three-dimensional design and drafting program for mechanical engineers (see photo B). It runs on many different workstations and can handle everything from basic three-dimensional drafting to finite-element mesh generation and 5-axis numerical-control machining (CAM).

Anvil comes in six different modules—drafting, modeling, rendering, finite element, simple machining, and advanced machining—and is sold on an annual-license basis. The cost ranges from $10,000 a year for a basic system to $38,000 per year for a loaded system. When Anvil is purchased for multiple "seats," say a dozen or more, the loaded-system price can be cut in half for each seat.

Bechtel makes a program called Walkthru. The program takes a three-dimensional CAD model from some other program and allows you to interact with that model as you might in "the real world." Using a three-dimensional color graphics workstation, a mouse, and a "button/dial" box with 8 dials and 32 buttons, you can "walk through" the model, controlling your imaginary body and head motion. Walkthru provides perspective, simulated views of the model—a boon to factory designers or architects. Walkthru runs on Silicon Graphics workstations. It demands at least 4 megabytes of RAM, 24 bit planes for color, z-buffering, z-clipping, and a button/dial box for input.

Solids Modeling and Animation

Intelligent Light's workstation software is used in industrial design, illustration, scientific and engineering animation, video and film animation, and technical publication (see photo C). It runs on Apollo workstations (and soon on Stellar graphics supercomputers).

Intelligent Light offers tools for model building and scene creation (including placement of light sources and specification of surface appearance), animation, rendering (with Phong and specular lighting, antialiasing, fog, and other special effects, including full color up to 96 bits per pixel), and image manipulation and recording. Images can be shown on screen or sent directly to videotape or slide. Prices range from $30,000 to $100,000.

Wavefront is the name I heard more than any other when I asked about "state of the art" in graphics software for workstations (see photo D). The firm sells a software package for three-dimensional dynamic imaging and high-end rendering and animation. The modules in this package are a modeler (a simple polygon modeler, not a full-fledged CAD package), a Preview module (for choreography of animation), and Image (for rendering—this module can crank away for many hours to create static images that are then put together for motion).

Image packs features such as shadows, ray-tracing, and reflectivity on the high end, and can also turn to lower-end work such as faceting, smooth shading, and the like. Wavefront software runs on many workstations, including all the Silicon Graphics machines and Tektronix and Hewlett-Packard mid-range systems. Prices start at $23,000 and go up to $55,000, depending on the platform's performance.

Wavefront recently announced that it will offer its software for less than $10,000, as a set of tools that other developers can use to add rendering to their own CAD or graphics programs or that neophytes can use with templates. The kernel of this kit will cost under $2000.

Photo C: *A three-dimensional solid display from Intelligent Light's software for model building, animation, image manipulation, and so on. (Photo courtesy of Intelligent Light.)*

Photo D: *A display from Wavefront's software for three-dimensional dynamic imaging and high-end rendering and animation. (Photo produced by Wavefront Technologies © 1988.)*

assigned to each screen pixel. (The physical organization of the memory bits needn't be anything like the logical organization. The translation between the two is handled by video-memory drivers or controller electronics.) Monochrome images are often all that is needed for graphics applications like CASE or CAP, although gray-scale abilities can be handy for publishing or even just for pull-down menus and icons.

The least-expensive workstations are monochrome, single bit-plane systems. Low-cost color systems can get by with 6 bit planes, although 8 bit planes is more common for standard, inexpensive color. To simultaneously display any of the choices in a palette of 16.7 million colors, 24 bit planes provide "full color," also known as "true color." Some workstations also offer more bit planes specifically for use in pull-down menus or window identification (laying down borders that show where windows, menus, and icons begin and end—information a mouse or other cursor needs).

Advanced graphics systems also use *z-buffers* and *double-buffering*. A *z*-buffer is additional video memory that holds information about the *z* axis: which objects on the screen are in front of which other objects when seen in three dimensions. A *z*-buffer adds cost to the system because of the price of memory and its associated controllers, and it can even slow down some simple two-dimensional applications, but it can add sophistication to powerful solids-modeling programs.

Double-buffering uses more bit planes to hold multiple presentations of the same area on the screen. Screen changes are faster if all the display controller has to do is change which part of the frame-buffer memory it looks to, instead of updating the information in a single section and then redisplaying that.

Although frame buffers and *z*-buffers can be built from standard RAM chips, in workstations they are often built from the faster video RAM (VRAM) chips. These more expensive cousins to standard dynamic RAMs (DRAMs) offer two I/O ports or buses, allowing them to provide their data to the screen at the same time the CPU and graphics program are altering it. (Naturally, there is control over actually reading and writing the same bit at the same time.)

Scan Conversion, Transforms, Scaling, and Clipping

Graphics hardware can go far beyond the frame-buffer memory that holds the image, and the CRT or screen driver that moves that image out into view. A series of operations must be performed before the appropriate pixels are placed in the frame buffer. The first is to generate a display list—a set of commands for graphics, such as "Draw a line from point A to point B," "Put a filled circle of radius *r* and center at (x,y)," or "There's a sphere of such and such a size at this point." Three-dimensional models may even make use of *nurbs*, or nonuniform rational B-spline curves, which are handy for representing complex shapes and curves. Moving them in software is a complex effort.

The display list is hierarchical, representing repeated parts of the complete image as subroutines that need to be detailed only once. The display list can be managed by the main system hardware or by a separate display-list processor that offloads the work from the CPU and so speeds up the entire graphics process. Most workstations have the CPU handle the display-list management, as do most personal computers. Some personal computer add-on graphics boards have a separate display-list processor, going beyond the typical hardware of a workstation. Graphics terminals often use display-list processors because so much of their graphics work must be done locally: The narrow bandwidth of the serial connection to a mainframe or minicomputer denies them the luxury of downloading completed graphics.

The primitives of the display list must be translated into individual pixels by a process called *rasterizing* or *scan conversion*. As part of this scan conversion, the computer must calculate how the picture looks from the viewer's perspective. If it's a three-dimensional image, the computer needs to know where in space the image is located, how big it is, and where the viewer's eye is in respect to the image. These calculations are called *transformations* and include *scaling* the object to the desired viewing size. Each point that describes the object must be transformed and scaled with respect to the viewer.

One specific operation that can streamline transformations is the bit-block transfer, or *bit-blit*. This is vital to the functioning of many windowing and other two-dimensional systems and has been built into a variety of hardware systems, including the Amiga's coprocessors and all of Apollo's workstations.

Then there are two *clipping* operations. The first concerns which part of the object the viewer is going to see. This is clipping in "world coordinates" or "clipping in the real world" and deter-

mines which parts of the object need to be used throughout the rest of the display process. The second operation is "clipping against a window," where any parts that will fall outside the on-screen window are trimmed so they don't wrap around and distort the image, or waste processing time on parts of the image that won't be seen.

Finally, there are special operations, such as *antialiasing*, that serve to improve the look of a graphics image after it has been transformed, scaled, and clipped. Antialiasing turns selected pixels partially on to fill in the "jaggies" or "sawtooth" forms that can appear in diagonal lines or arcs. These effects can even distort the apparent image or cause distracting *moiré* patterns. Higher-resolution displays can also diminish these jaggies, but can be much more costly.

Wire Frames, Solids Rendering, and Animation

Three-dimensional graphics demand all the same operations as two-dimensional graphics, plus many more. These three-dimensional processes—which actually create a two-dimensional view of an imagined three-dimensional object—can be handled by any system as software algorithms. These processes are so computation-intensive and complex that they are practical only when handled in hardware.

There are actually three distinct styles of three-dimensional work: wire frames, solids rendering, and animation. Wire frames display the vertices and edges of objects, sometimes omitting the hidden lines and surfaces that are overlapped by closer lines and surfaces. Wire-frame work can be done with so-called two-dimensional workstations, although some of the work, such as hidden-line removal, will slow considerably without hardware assist.

Solids rendering or modeling uses imaginary lighting and shading to fill in the surfaces of a three-dimensional object. The latest trend in solids modeling is *photorealism*, which attempts to make the view of a simulated, displayed object as similar as possible to that of an actual solid object. Solids rendering can be done in many ways, with various operations and algorithms that yield varying levels of realism. Many of these algorithms can take an enormous amount of time if done entirely in software. For example, I have used the sophisticated ray-tracing algorithm option of the Solid Dimensions program from Visual Information. Running on a Mac II with its floating-point coprocessor, it can take 12

to 24 hours to render a single, simple scene with several uncomplicated objects.

Solids modeling has long been thought the province of "three-dimensional" workstations, the mid-range and high-end systems that have special hardware to speed the rendering algorithms. In fact, the terms "low-end workstation" and "two-dimensional workstation" were often synonymous. Silicon Graphics' introduction of the Personal Iris brings

Silicon Graphics' Personal Iris brings solids modeling into the league of personal workstations.

hardware-assisted solids modeling into the league of personal workstations.

Animation demands even more from a computer system, with the ideal being real-time animation: the ability to generate rendered solids quickly enough to simulate objects in motion. This means at least 10 display changes per second. The animation can be done in non-real time, calculating frames one at a time, and then saving them in mass storage for viewing in sequence later, or sending them to a videotape peripheral.

Animation is not only useful for art and commercials but has been put to work in presentation graphics for design proposals, for analyzing the fit and interaction of mechanical devices, and for modeling fluid flow in many fields. Animation has been handled on everything from the best mid-range workstations to supercomputers. It is becoming a reality in the low end with the introduction of the Personal Iris.

Depth Cueing and Hidden-Line Removal

One three-dimensional technique is called *depth cueing*. It varies the pixel intensity relative to the distance that pixel represents from the viewer's eye. Both depth-cueing and hidden-line-removal techniques depend on a z-buffer.

To render a wire frame with a solid ap-

pearance, the hidden surfaces must be omitted and the visible surfaces of the object must be filled with color or gray-scale shades. Hiding background edges and surfaces can be done with several levels of transparency, ranging from opaque to "screen door" (where a pale transparent shade fills the closer surface), which gives the illusion of looking through a fine mesh.

Some operations are more appropriate to two-dimensional than three-dimensional applications and serve mainly to slow down three-dimensional operations without improving the quality of the image. Antialiasing, for instance, isn't as important in three-dimensional solids where the shaded fills will provide some of the same smoothing. It can even make the images fuzzy. In three-dimensional wire-frame work, however, antialiasing can still play a role.

Gouraud and Phong Shading, Ray Tracing, and Radiosity

When it comes time to fill a surface, there are many different algorithms or schemes for determining the particular shade or color for each surface or pixel. The simplest operation is *constant shading*: A model is broken logically into polygons that will show on-screen, and a single shade or color is assigned to each entire polygon. This is a quick rendering scheme, but it creates distortion when side-by-side polygons have two shading levels. You'll be able to see a line of delineation between the two that will show as a facet of the object and may not be intended as such.

The Gouraud algorithm provides one way to ameliorate that distortion, by linearly changing the shading across each polygon. The algorithm samples two opposite edges of the polygon and then shades across the face between the boundaries. This removes abrupt shading changes and can be done in software or hardware. It is usually made a part of the scan-converter hardware in high-end graphics workstations.

The Phong algorithm is more exacting than the Gouraud. Instead of polygon-by-polygon examination and linear averaging, Phong looks at individual pixels. That makes it much more time-intensive than Gouraud as well. Gouraud shading can approximate Phong shading, however, by making the polygons smaller and smaller.

Ray tracing and radiosity are some of the most advanced operations in the pursuit of photorealism. There are two types of surface-reflection effects: specular and diffuse. Specular reflection has a

shiny effect, bouncing back light as though from a smooth surface. Diffuse reflection does not directly reflect as much light, instead spreading it out in more directions as though from a rough surface.

Ray tracing is a specular effect that traces each light ray from the assumed viewer's eye to the surfaces in a three-dimensional graphics image. The ray is bounced mathematically from one surface to the next until it reaches a non-reflective surface, and its lighting effects are noted. Even for a workstation, this can take some time, particularly without hardware assist and acceleration. Software can contain the compute task for ray tracing somewhat by limiting the number of rays traced.

Radiosity is a diffuse-lighting effect. It takes the diffuse light that reflects off each object in a scene and calculates the impact that light will have on adjacent objects. These calculations could be done in software, but again, that would take hours unless aided by special hardware or given a time limit. Hewlett-Packard's mid-range SRX workstations include hardware for Gouraud and Phong shading with programs for radiosity modeling with progressive refinement in software (see photo 1).

MIPS, MFLOPS, LINPACK, Vectors, and Polygons

Any graphics effect could be created by nearly any computer, given enough time. However, graphics performance is typically measured in terms of speed. This was aided by an IBM study in the early 1980s showing that productivity on a computer was very closely related to the graphics speed or "interactivity" of the display. If the picture on the screen could update in less than half a second, productivity was at a peak. Between half and three-quarters of a second, there is a very sharp "knee" in the graph where productivity falls. That half-second update is what most workstation manufacturers aim for as a minimum. Markets such as animation—which requires a minimum of 10 frames a second—need even more speed.

The simplest speed measurement is the power of the system CPU, typically measured in MIPS. This is a good first cut at graphics power because most personal workstations end up doing a lot of graphics work—both applications computing and graphics scan-conversion algorithms—in software. The MIPS measurement ranges from the 1 or so DEC VAX 11/780 MIPS of an 80386-based personal computer to the 7 or 10 MIPS of

the SPARC-chip-based Sun-4 system or the RISC-chip-based Personal Iris from Silicon Graphics. Adding a floating-point mathematics coprocessor can boost the system's MIPS.

In fact, floating-point performance is so important in graphics that MFLOPS (million floating-point operations per second) are often used as a yardstick in place of MIPS. A system still needs a good measure of MIPS to handle the standard computing tasks, so the ratio between MFLOPS and MIPS shouldn't get too big, or the graphics task could become compute-bound by standard applications work. The Personal Iris claims a 10-MIPS CPU with a 20-MFLOPS peak graphics processor. When you're hearing tales of great MFLOPS, however, you should inquire about their origin: Peak MFLOPS can be much higher than the MFLOPS rating from a standard program such as the LINPACK benchmark, a 100 by 100 matrix calculation in double precision.

Even with such standardization, different machines will be faster on different tasks, depending on which graphics operations have been built into hardware, optimized in software, or left to generic software routines. The number of vectors drawn per second (for two-dimensional and wire-frame images) and polygons drawn or shaded per second (for three-dimensional solids-modeling images) are

other measures of performance. You'll often see a sequence of these lined up in workstation specifications, with values in the tens and hundreds of thousands, depending on the workstation model and graphics options. A system with Gouraud shading in hardware will be much faster on that operation even though it has a lower MFLOPS rating on FORTRAN programs and may even have a lower bit-blit speed on two-dimensional images.

Workstations sometimes can be linked together in a network to grind away at complex graphics tasks. Apollo, for instance, can use its Network Computing System (NCS) to dish out parts of a complex graphics-calculation task to other workstations and their CPUs' MIPS. (Hewlett-Packard has recently licensed the NCS for its own systems.) The result can then be displayed on a single workstation.

Libraries and Device Independence

All the hardware power doesn't exist in a vacuum. Applications developers love to have some sort of standard graphics device to which they can adapt their programs. Workstation hardware tends to differ much more than hardware in the personal computer arena. Although the 68020 is the most standard CPU (and plenty of others are in use), workstations have different addressing and memory

Photo 1: *The ray-tracing feature of the HP 9000 TurboSRX engineering workstation from Hewlett-Packard allows you to produce the realistic highlights, reflections, and effect of transmittance seen in this screen shot. (The data to produce this image is courtesy of Chrysler. Photo courtesy of Hewlett-Packard.)*

schemes and graphics architectures.

This diversity makes it tough for you to create a program that will run on many platforms. Unix has certainly become the overwhelming standard in operating systems at the low end of workstations, although there are various flavors of Unix, from IBM's AIX to NeXT's Mach to Sun's hybrid of System V and Berkeley 4.3. For real-time operations, such as process-control graphics, there are special Unix versions, such as that from Masscomp, and DEC's VMS operating system. Plain Unix is not well-regarded for real-time work.

There are also standards that concern graphics itself. For the past 5 years, the de facto standard has been ANSI's Color Graphics Virtual Device Interface. You could just write a program to call the CGVDI circle command, cube command, or whatever, and let the workstation vendor worry about the interface between the standard library of graphics routines and the proprietary hardware. CGI (a common abbreviation of CGVDI) was a fairly simple two-dimensional standard now being replaced by PHIGS.

PHIGS is also a two-dimensional standard, but it's more advanced and is hierarchical, which lets it call subroutines for efficient code execution and size (see the article "PHIGS: Programmer's Hierarchical Interactive Graphics Standard" by Martin Plaehn in the November 1987 BYTE). It's a programming base that a graphics application can use: Make your program call on PHIGS graphics primitives, and it will run on any workstation supporting PHIGS (most do). PHIGS + and PHIGS + + are proposed extensions that include curved surface, shading, lighting, depth cueing, and other three-dimensional paraphernalia. However, PHIGS does not come with windowing tools.

Another device-independent graphics standard is X Windows. Like PHIGS, it's a set of routines that an application program can call on. But X Windows allows for distributed graphics: An intensive calculation application can run on a high-powered server and display the results in a window on a personal workstation. It's a toolkit for implementing window systems and displaying two-dimensional graphics. X Windows does not include three-dimensional abilities. An adjunct group to the X Consortium (that came up with X Windows) is working on PEX (PHIGS Extensions to X Windows), which would merge PHIGS + with X Windows. This would use X Windows to manage the window environment and PHIGS + to render the graphics.

Two other possible standards of interest to personal computer users are the HOOPS graphics library from Ithaca Software and Display PostScript from Adobe. HOOPS is an object-oriented set of routines that has been implemented on a number of personal-computer products, including a transputer-based IBM PC plug-in board from Nth Graphics. Display PostScript has been used by the NeXT workstation and has a direct relationship to the PostScript that is in so many laser printers, but it hasn't made a dent in workstation graphics software.

There are also proprietary standards, such as Apollo's three-dimensional Graphics Metafile Resource File, Hewlett-Packard's Starbase, and Silicon Graphics' Library. IBM, which has not had much success with its own RT workstation, recently licensed SGI's library and graphics accelerator chip for use in its own workstations.

Database Design and File Formats
Another important subject in graphics work is the structure of the databases that comprise graphics images, and the files used to transport them among different applications and workstations. Although flat files and relational databases were once used, these are being challenged by object-oriented databases. An OODB stores an image as objects with properties. This makes the translation from application to screen display less complex. It also suits the graphics application code to object-oriented languages such as Smalltalk and C + +, with their attendant extensibility, reusability, and prototyping ease during programming.

For capturing and transferring graphics, the latest proposed standards are Computer Graphics Metafile and RenderMan. CGM allows you to take a picture file from one system and run it on another. RenderMan is newer and not as well-established, but it includes more sophisticated information on the three-dimensional aspects of an image.

The Roster of Players
There are a lot of players in the workstations game and a lot of applications (see the text box "Putting Workstations to Work" on page 257). Because workstations have windowing systems and some ability to handle monochrome two-dimensional graphics on large displays, all could be considered graphics powerhouses next to the average personal computer. However, some workstations rise above the rest because of special graphics hardware built in or available as an option within the personal-workstation price range.

• *Silicon Graphics.* The most likely winner this year of a "Heisman trophy" in workstation graphics is the Personal Iris from Silicon Graphics. Begun in the early 1980s by a Stanford professor, Silicon Graphics has specialized in applying

VLSI and its own optimized library of graphics routines to three-dimensional graphics.

The Geometry Engine VLSI chips that made its Iris workstation a hit at the mid-range and high end of the market for animation, simulation, medical imaging, and mechanical design CAD are now built into the less expandable—and far less expensive—Personal Iris. For just under $17,000, you can buy a Personal Iris diskless system with a 10-MIPS CPU (the R2000A RISC processor from MIPS Computer), 8 megabytes of RAM, 8 bit planes for color, 2 bit planes for menus, 2 bit planes for systems administration (all using VRAMs), and a Geometry Engine.

The system can drive a 19-inch, 1280-by 1024-pixel color display with real-time animation of 10 frames per second, performing 200 calculations every second for every pixel on the screen. Expansion to 24 bit planes for color, 8 system bit planes, 155 megabytes of hard disk, and a 24-bit z-buffer runs the price up to nearly $30,000. That's still far less than the bottom price of the mid-range Iris 4D70 (at $44,000) or a competing three-dimensional workstation such as the Sun-4/260, Apollo 590 Turbo, or HP 360SRX Turbo—all mid-range workstations. Lighting, specular highlights, and scan conversion are built into the accelerator hardware of the Geometry Engine, which can pump out 85,000 three-dimensional vectors per second. The Personal Iris supports PHIGS.

• *Hewlett-Packard.* HP has a variety of personal workstations for two-dimensional graphics, and mid-range systems for three-dimensional work. The newest personal-workstation system is the HP 9000 Model 340 series, based on the 68030 CPU and 68882 FPU. The 340s have at least 4 megabytes of RAM and HP-UX Unix, and they support X Windows and PHIGS.

The lowest-price unit in the 340 series, at just above $5000, is the 340M, a 1024-by 768-pixel monochrome system for mechanical design and drafting. The 340C+ is a nearly $9000, 6-bit-plane color, 1024- by 768-pixel, two-dimensional system with hardware support for vector and polygon drawing and fills. The 340CH is an $11,000, 8-bit-plane color, two-dimensional system that starts from the 340C+ base.

The $16,000 340CHX is a CH model with a plug-in, integer-based graphics accelerator (with a 68020 processor) that can double display-processing performance—taking charge of the transform

and clipping calculations.

Hewlett-Packard uses a "strip z-buffer" that takes advantage of unused memory in the mainframe-buffer area (the typical resolution of displays doesn't match exactly with the size of standard memory chips and chip sets). This saves money on memory for a small expense on increased memory-controller overhead. The systems all come with an HP-HIL interface (HP Human Interface Loop) that can daisy-chain input devices such as mice, knob-boxes (a suite of dials that can be turned to, for instance, change each of the rotations and positions of a three-dimensional simulated object), tablets, and so on. The 340SRX incorporates silicon assists for solids-rendering work, such as light sources, Gouraud shading, and transparency. The 340 series is not as expandable as the 360 and 370 series (also 68030-based) mid-range workstations, but it is more expandable than the 318 and 319 families it replaces. It is object-code-compatible with the other HP 9000 Series 300 workstations.

• *Apollo.* The 68020-based Series 3000 is Apollo's lowest-priced workstation, with an estimated CPU performance of 1.5 MIPS. Above that are the Series 3500 and Series 4500. The 3500s have a 68020 and boast a performance of 4 MIPS. The 4500s have a 68030 and a 68882 FPU for a performance of approximately 7 MIPS. You can get the 4500 with 15- or 19-inch monitors, in 1280- by 1024-pixel monochrome, 1024- by 800-pixel, 8-bit-plane color, or 1280- by 1024-pixel 8-bit-plane color. The 3500 adds a 1024- by 800-pixel monochrome option to that list.

All these systems are aimed at two-dimensional and three-dimensional wire-frame applications. The 3000 does not contain special graphics accelerators other than bit-blit assist in hardware (essential for windows and polygon fills). Instead, it depends on general floating-point performance and an optimized library of graphics routines for graphics speed. A Weitek 3164 FPU option is available for each system. The 3500 and 4500 have a dedicated graphics processor for accelerating the drawing of two-dimensional primitives.

The systems run under Unix and support PHIGS, X Windows, and Apollo's own three-dimensional Graphics Metafile Resource File library. All of Apollo's graphics workstations, including its mid-range and high-end Turbo Domain systems, are binary-compatible—meaning a single program can run on any of

them without even a recompilation, using whatever graphics abilities are available. While you can get a Series 3000 monochrome, diskless node for around $5000, a 3500 color system without disk runs just over $12,000. The 4500 systems start at $19,000 (diskless, monochrome) and rise to $35,000 and higher (color, 348-megabyte disk drive). The mid-range and high-end Apollo workstations include up to 24 bit planes for graphics and hardware assist for three-dimensional algorithms.

• *Digital Equipment Corp.* DEC offers three personal workstations for graphics: the VAXstation 2000, the VAXstation IIGPX, and the VAXstation 3200. The 2000 comes in monochrome and color versions, beginning at just over $5000 and just under $8000, respectively. The 2000 is built around DEC's own MicroVAX II chip set with its own FPU, 4 megabytes of RAM, and 1024- by 864-pixel display resolution. The monochrome system has a single bit plane; the color system has 4 bit planes. Eight-bit-plane color models are available and cost $13,000 to $18,000. The graphics chips in the VAXstation 2000 are the same as in the VAXstation IIGPX. GPX prices begin at just under $20,000. This workstation is larger, with more memory and disk capacity than the VAXstation 2000. A diskless, 8-bit-plane color system runs nearly $23,000. Both the 2000 and the IIGPX are built around the MicroVAX II engine.

The VAXstation 3200 uses a CMOS version of the MicroVAX II engine and FPU, along with a dual-cache memory, and so is much faster on all system measurements. A color 3200 system starts at $29,000 and comes with special graphics coprocessors from DEC. These GPX coprocessors perform bit-blit, scaling, clipping, fill, and scrolling operations in hardware. Support for PHIGS and X Windows is included.

On the high end, DEC has developed the VAXstation 8000 three-dimensional real-time graphics workstation with Evans & Sutherland (famous for its flight simulators and other high-end graphics computers). The Evans & Sutherland graphics engine in the 8000 handles 24 bit planes double-buffered with Gouraud and Phong shading and up to 16 light sources.

• *Sun Microsystems.* Sun has three different architectures for personal workstations, from the 80386-based 386i to the 68020-based Sun-3 and the SPARC

RISC-chip based Sun-4. The 3/50 is the lowest-priced unit, at about $5000 for a diskless, monochrome system with 1152-by 900-pixel resolution. The lowest-priced color system is the 3/60 with 1152- by 900-pixel resolution for gray scale or color, and an option of 1600 by 1280 pixels for high-resolution monochrome. These systems are aimed at two-dimensional applications, such as mechanical CAD, CASE, and publishing.

The 386i can produce 1152- by 900-pixel monochrome or 1024- by 768-pixel color. These systems support PHIGS and X Windows, but they have little in the way of hardware acceleration for graphics. Above these machines are the 4/110 and the 4/150, which cost $20,000 and more. The 4/110TC includes a graphics option with 24 bit planes for color.

In the mid-range systems, Sun offers a CXP graphics accelerator option. At the high end, there's the TAAC image-processing option that has hardware for manipulating entire images—such as converting a series of CAT-scan images of a heart into a simulation of the heart beating. The TAAC boards alone cost about $30,000.

• *Tektronix*. The 4319 is the entry-level workstation from Tektronix, with a 68020, 68881 FPU, X Windows, and 1280- by 1024-pixel resolution. It comes with Tektronix' own Color=Cache chip for bit-blit in hardware, and it can handle up to 8 bit planes for color. The 4320 series of workstations, from the $23,000 4324 to the $27,000 4325, is aimed at two-dimensional applications and is also built around a 68020 and 68881. Resolutions of 1024 by 768 and 1280 by 1024 pixels come from a separate 68020 dedicated to graphics and a special bit-slice processor accelerator.

The Tektronix 4330 family is a three-dimensional, mid-range series of workstations (from $37,000 to $52,000) with a fascinating graphics capability: stereographic viewing. These workstations boast 24 bit planes for color, z-buffers, and double-buffering, as well as the separate 68020 and accelerator chips for graphics you find in the 4320 series. Depth cueing is built into the standard display, and an optional system that alternately displays left-eye and right-eye views produces three-dimensional images when you look through special polarized glasses. Any three-dimensional application that supports Tektronix three-dimensional terminals and workstations can run in the stereo mode.

• *Others*. There are lots of other workstation makers with impressive graphics. IBM and NeXT haven't sprung unusual graphics hardware on the public, but they have recently put impetus behind Display PostScript (in the NeXT workstation) and NeXT's NextStep application environment (which IBM has licensed).

Intergraph makes a variety of workstations, including some new mid-range systems with multiple megapixel displays. And there are graphics terminals that can unload some of the work from workstations—such as the Seiko GR4400 that has hardware for Phong shading and the GX4000 from Raster Technologies that supports PHIGS and PHIGS+ with a parallel-processor architecture. If you can spend more, you'll find systems such as Apollo's Series 10000 that can use 40 or 80 bit planes with z-buffers, double-buffering, and 4 processors to make as many as 600,000 polygons per second.

Capabilities Up, Prices Down

There are two clear indicators to what you'll be seeing in personal-workstation graphics next year and the year after that. First, look at the mid-range systems: z-buffers, depth cueing, shading in hardware, and the like. Resolution probably won't climb much, but three-dimensional systems will certainly follow the Personal Iris into lower prices. Perhaps Tektronix' Stereoscopic Display will catch on.

The other indicator comes from the arts. According to John Metcalfe, the director of marketing for entry-level systems at Silicon Graphics, "The animation industry drives us. They want to see real life on a workstation—and that drives us from a technical point of view. About 2 years later, the mechanical-design people decide they want that same stuff, even though if you had asked them what they wanted before, they probably wouldn't have mentioned it." ∎

Phillip Robinson is an editor for Virtual Information *(Sausalito, CA), where he researches and analyzes trends in the computer and workstation markets. He is a contributing editor for BYTE and can be contacted on BIX as "robinson."*

Circle 126 on Reader Service Card
World Radio History

Networking with Unix

RFS and NFS provide the key
to distributed file systems

Greg Comeau

I f you own a Unix workstation, chances are at some point you will want to network your machine to other computers so you can access their files and resources. Sun Microsystems' Network File System (NFS) and AT&T's Remote File System (RFS) allow you to do just that. These file-system arrangements, however, go about their business in very different ways.

To understand their strategies, you should first understand a few points about Unix. A typical Unix workstation takes advantage of a hard disk's cylindrical layout and zones specific contiguous cylinders into file systems.

For instance, most Unix machines contain at least two file systems: one system that contains the Unix environment, and another that contains user files and programs. Other Unix machines contain local file systems as well. Since each file system usually contains a directory structure (which consists of directories, files, and devices), they are handy for structuring your disk into unique and specific areas suited to your needs.

Unix adheres to hierarchical file systems, which many users are familiar with. Note, however, that a reference to a path, such as /usr/comeau/dir, could have many variations under Unix because the operating system is based on the concept of a root file system. In other words, the root file system is your system's main file system. All subsequent file systems—for example, /usr—must be mounted atop the root file system.

Unix does this by overlaying a file sys-tem hierarchy structure onto an existing directory. The directory may or may not be empty. Therefore, access to a path such as /usr/comeau happens transparently without the person or process who is actually accessing the path knowing that usr is just a logic connection to /.

This differs from other systems, since each system administrator has the power to mount any unmounted file systems onto any directory that exists on a currently mounted file system.

This setup works out nicely for several reasons:

• Your machine need not be concerned with file systems that contain projects not currently being used.
• Introducing a file system is as simple as mounting it.
• You can access all mounted file systems transparently because they adhere to the directory structure. (For example, you can mount usr on / to obtain /usr.)
• Since file systems are constrained to cylinders on a disk, you can easily add file systems from different physical disk drives on the same machine to the current file-system state.

ILLUSTRATION: ROBERT TINNEY © 1989

Although this all works reasonably, unfortunately it locks you into accessing file systems located on the machine you are logged onto. You can access remote machines via file-copy protocols, usually the Unix-to-Unix copy program uucp, but this can be awkward.

RFS and NFS let you transparently connect the file systems of one computer to the file systems of another. They enable your Unix machine to use the file system of another Unix machine so naturally that you may not always realize that you are using another computer's resources.

How RFS Works

RFS ensures that all network transactions adhere to Unix I/O semantics. In other words, open( ), close( ), unlink( ), and so on, and even ioctl( ) all behave so rationally that files are always left in manageable states.

The transparency is so strong that reliability, security, and even simultaneous access to files all occur as if on one Unix machine. Of course, this implies correctly that RFS was intended to be used only for communicating with other Unix systems.

RFS functions depend on a machine supporting a transport provider. The TP ensures that there is a physical network along with the appropriate software to drive the hardware of the network. The hardware is typically Ethernet, System Network Architecture, or StarLAN. The software is broken down into two software layers to handle the network protocols.

The software subsystem works by using an I/O facility called Streams. With Streams, you can connect to one of several network protocols (i.e., Transmission Control Protocol/Internet Protocol, Ethernet, and so on) for your intermachine communications needs. The end result: RFS is network-independent. That is, since it can be independently configured regardless of the underlying hardware, it need not be pinned down to a specific protocol, vendor, or machine.

Streams accomplishes this via handlers, each of which is a different layered protocol in software. Each handler usually interacts directly with the hardware and converts the real protocols into a common format so that the higher layered protocol handlers (the ones Unix would interface with directly) all map into a common protocol.

The RFS scheme consists of client and server machines living within a domain. Machines within a domain that advertise their file systems for other machines in the domain to use are called server machines. (For a glossary of basic terms, see the text box "The Nomenclature of Networking" on page 269.) Machines that use the advertised file systems by mounting them are called clients. Since a server can choose to use a file system of another machine, some machines will be both servers and clients.

Furthermore, one machine in the domain is chosen by the domain's system administrator as the primary server. It maintains various lists. The most important of these is the list of directories, files, and even devices that were advertised for network usage. Other lists help maintain network security. Among other abilities, these security lists can restrict specific machines to a given network resource, as well as specific groups or user IDs.

Generally, once these files are set up across all the machines in a given domain, each machine advertises what it has to offer as it boots up, as well as mounting resources it needs from other machines. This results in very little network maintenance once the distribution directories and files for each machine have been determined. Besides its adherence to the Unix I/O semantics, this is another reason why RFS is so transparent.

Good Points and Bad

One problem with the RFS scheme is that the state of the network is more fragile since it enforces the Unix I/O semantics. Because of this, if the primary-domain server fails, the secondary-domain servers, chosen by the domain's system administrator, temporarily assume the primary server's responsibilities until it is restored and rebooted.

Interestingly, this doesn't particularly affect the network, since machines currently in the network have already advertised their resources and mounted those that they needed; therefore, current mount points will not be broken. However, it does affect any machines in the domain that have been booted in the interim.

Since RFS ensures Unix I/O semantics, users and programs can make directories, files, devices, and named pipes across the network without concern for where those entities really are. The lack of concern results from the ability to always have a consistent view of a file. In other words, database accesses and file locking perform as you would expect from a single machine. This is nice, since Unix should act like Unix whether or not it is connected to a network.

Furthermore, because security enforcement is at a machine, group, user, and Unix level, network security is never at a disadvantage because of RFS. One indirect problem of this, though, is that machine administrators choose to advertise all the file systems of their machines to all the clients of the domain, and users may be given the chance to read or write files they shouldn't have access to. This is strictly an administration concern, though, since the default choice is that a given file system or entities within a server will not be shared.

All this consistency has a price. Because of it, the network has the potential for more traffic, resulting in a loss of performance and more bottlenecks.

At first glance, the ability to selectively mount remote file systems within selected directories on a given client seems rather powerful. Looking at this in detail, you see that each machine has its own different view of the network. This may be fine for a given machine, but if you or your programs are relying on a consistent directory tree or on access to a specific path name, these different views can become confusing and error-prone in some situations.

What Makes NFS Tick?

NFS's general setup is similar to that of RFS. A few differences, however, point out their distinct tactics. Although they both hinge on the ability to treat a remote file system as being local, they do so in different ways. Mainly, NFS is not a network extension of Unix and therefore does not attempt to adhere to Unix semantics. In fact, it takes this difference one step further and allows other non-Sun and non-Unix operating systems to use it. Since the protocols for NFS are in the public domain, some vendors have endorsed its use, and Sun even developed a Unix-to-MS-DOS network connection called PC-NFS.

Even with these differences, the internal makeup of NFS is not unlike RFS when it comes to network accessibility. An advertising mechanism for servers and a mounting ability for clients still exists. Also, instead of the network layering associated with Streams, NFS has substitute mechanisms called Remote Procedure Calls (RPC) and External Data Representation (XDR).

Since remote file systems need only appear to be real Unix file systems, NFS implements a Virtual File System mechanism. For instance, the Unix file system usually accesses files through an operating-system disk handle called an *inode*. Since the remote file system being

accessed doesn't need to use the Unix file-system structure, NFS uses VFS to construct a *vnode*, a virtual inode. As long as VFS can do this, NFS can exist as a machine- and operating-system-independent application.

VFS is smart enough to determine if a request is from the local machine and will resort to using standard file-system operations in that case. When the request is remote, it will issue RPCs. An RPC can be thought of as a client process issuing a procedure call to a server process just as if the procedure call had occurred on the client machine. Under this scenario, a client calls a local procedure to issue a network request or data transfer. A message is sent to the server in response to the client procedure call, upon which the server runs whatever procedures it needs to fulfill the request and then returns to the client's local procedure call with the answer.

An RPC request, which uses XDR internally, transfers data and messages. Because of the diversity of machines and their CPU architectures, XDR is a common machine-independent protocol to be used at the lowest layer of NFS. This means that data structures used by your programs—whether they are base types or types derived from the base types from any machine—can be transparently described, transmitted, and converted to more than one machine. All this can be done regardless of machine architecture. The combination of RPC and XDR provides the equivalent of the standard I/O (stdio) library package that is commonly found in C programming environments.

The network address list and advertising list, such as those found on the primary domain server under RFS, are found through the Yellow Pages. The Yellow Pages is simply a read-only database for NFS.

Pros and Cons
NFS has major concurrency problems because it isn't a Unix network extension. This also means NFS doesn't support all Unix file-system operations, can't obtain access to remote devices, can't support file locking, and can't support file-append operations.

These problems, however, do have their benefits. To begin with, NFS need not be used only for machines supporting Unix file systems. Therefore, it's both machine and operating-system independent. Second, because NFS and its underlying protocols (VFS, RPC, and XDR) are in the public domain, implementations on diverse hardware and operating systems should take place.

Also, because of the lack of knowledge about remote clients and servers, NFS remains stateless. You can view this as advantageous, because NFS doesn't need to keep track of past resource requests, such as files opened by clients. Cleaner network traffic results, as well as better network performance, because NFS can perform block read-aheads and block write-behinds. Furthermore, if the network fails, error recovery is much simpler than in RFS. Currently, NFS also supports the concept of a diskless workstation.

Diskless workstations have potential. A typical local-area network could contain many smaller systems, each containing smaller and slower disks. Taking advantage of the hardware on another machine, even to the extent of making a client completely diskless, is smart from the standpoint of economy as well as performance. In the latter case, the diskless workstation might take advantage of an extremely fast and large disk on the server. This scenario would come in handy in a typical database application.

As in RFS, the default situation for NFS is that file systems and files should not be shared. This creates more administrative work and could limit transparency under many networks. For this rea-

The Nomenclature of Networking

Here's a glossary of basic networking terminology to help you understand the Network File System (NFS) and the Remote File System (RFS) schemes:

- *Networking address or node name*. The unique name of a machine in the network.
- *Resource*. A file, directory, subdirectory, device, or logical pointer to them.
- *Client*. A machine that requires access to the network to obtain resources of other machines.
- *Server*. A machine that donates its resources to the network.
- *Primary server*. The network master, responsible for maintaining the node names and resource identifiers of the machines in its domain. Clients communicate with the server to advertise their resources to the domain or to query it for a resource of the domain.

- *Domain*. The network area controlled by the primary server. Under RFS, a domain serves many purposes. It allows easy access to the domain via a predetermined name; controls security via a domain member list; allows clients in the domain to contact other clients directly; and allows clients to contact other domains.
- *Distributed file system*. A file system logically connected from various physical machines.
- *Homogeneous network*. A distributed network that requires a specific vendor's proprietary hardware and software. Such networks cannot usually communicate with another vendor's hardware.
- *Heterogeneous network*. A distributed network that can communicate with unrelated hardware by different vendors, usually via de facto protocols and hardware standards.

son, many servers choose to make their file systems available all the time.

A Feasible Choice
Networks are great for allowing smaller or even diskless workstations access to larger and more efficient resources. These resources usually include accessing remote file-system files, but they can also involve resources such as printers, floppy disks, program spoolers, or program documentation.

Both NFS and RFS are feasible choices for creating a shared Unix workstation environment, and they are capable of coexisting as de facto standards. In fact, with the recent joint ventures between AT&T and Sun Microsystems, NFS and RFS may someday reside on the same machine or network. Notwithstanding this, both are powerful and flexible options. Because of their transparency to users, shell scripts, and programs, they can provide economical networks regardless of hardware vendor constraints. ■

Greg Comeau is CEO of Comeau Computing, an independent software development and consulting firm specializing in Unix, C, and C++ programming tools. He can be reached on BIX as "comeau."

World Radio History

World Radio History

The Players

Acer Counterpoint, Inc.
(408) 434-0190
Inquiry 1015.

Acer Technologies Corp.
(408) 922-0333
Inquiry 1016.

Advanced Logic Research, Inc.
(714) 581-6770
Inquiry 1017.

A.I.C. Computers, Inc.
(416) 492-2777
Inquiry 1018.

Alpha Microsystems
(714) 957-8500
Inquiry 1019.

Altos Computer Systems
(408) 946-6700
Inquiry 1020.

Amdek Corp.
(800) 722-6335
(408) 922-5700
Inquiry 861.

American Mitac Corp.
(408) 432-1160
Inquiry 862.

Apollo Computer, Inc.
(508) 256-6600
Inquiry 983.

Apple Computer, Inc.
(408) 996-1010
Inquiry 992.

AST Research, Inc.
(714) 863-1333
Inquiry 1000.

AT&T Information Systems
(800) 247-1212
Inquiry 863.

Becterm
(418) 835-1551
Inquiry 864.

Compusource International
Corp.
(305) 251-1188
Inquiry 865.

Computer Systems
(313) 779-8700
Inquiry 866.

Datamedia Corp.
(603) 886-1570
Inquiry 867.

Dell Computer Corp.
(512) 338-4400
Inquiry 868.

DEST Corp. Florida
(407) 994-3720
Inquiry 869.

Digital Electronic Systems
(305) 572-3338
Inquiry 870.

Digital Equipment Corp.
(508) 897-5111
Inquiry 871.

Elco
(818) 284-3281
Inquiry 872.

Everex Systems, Inc.
(415) 498-1111
Inquiry 873.

Fortron Corp.
(408) 432-1191
Inquiry 874.

Franklin Telecom
(805) 373-8688
Inquiry 875.

Gulfstream Micro Systems,
Inc.
(DEST Corp. Florida)
(407) 994-3720
Inquiry 876.

Hertz Computer Corp.
(212) 684-4141
Inquiry 877.

Hewlett-Packard
(415) 857-1501
Inquiry 878.

Honeywell Bull
(617) 895-6000
Inquiry 879.

IBM Corp.
(914) 765-1900
Inquiry 880.

Icon International, Inc.
(801) 225-6888
Inquiry 881.

IndTech, Inc.
(408) 743-4300
Inquiry 882.

Kowin Computer Corp.
(213) 721-5500
Inquiry 883.

Logic Process Corp.
(214) 340-5172
Inquiry 884.

MAD Intelligent Systems, Inc.
(408) 943-1711
Inquiry 885.

MAI Basic Four, Inc.
(800) 854-3194
(714) 731-5100
Inquiry 886.

Masscomp
(508) 692-6200
Inquiry 887.

Megadata Corp.
(516) 589-6800
Inquiry 888.

Micro Display Systems, Inc.
(408) 743-4300
Inquiry 889.

NBI, Inc.
(303) 444-5710
Inquiry 890.

NCR Corp.
(513) 445-7478
Inquiry 891.

NEC Home Electronics (USA),
Inc.
(800) 632-7638
(312) 860-9500
Inquiry 1007.

NeXT, Inc.
(415) 424-0200
Inquiry 984.

Northgate Computer Systems
(800) 548-1993
(612) 553-0111
Inquiry 892.

Ogivar Technologies, Inc.
(514) 737-3340
Inquiry 893.

Olivetti USA
(201) 526-8200
Inquiry 894.

Pindar Computer Corp.
(703) 790-0555
Inquiry 895.

Profit Technology
(800) 223-4628
(212) 809-3500
Inquiry 896.

Rexon Business Machines
Corp.
(213) 641-7110
Inquiry 897.

SCI/Fortune Systems Corp.
(415) 593-9000
Inquiry 898.

SCI Systems, Inc.
(205) 882-4800
Inquiry 899.

Silicon Graphics, Inc.
(415) 960-1980
Inquiry 985.

Sophos Integrated Systems
(301) 292-5585
Inquiry 900.

Spectrix Microsystems, Inc.
(416) 474-1955
Inquiry 966.

Sun Microsystems, Inc.
(415) 960-1330
Inquiry 986.

Symbolics, Inc.
(617) 221-1000
(617) 621-7500
Inquiry 967.

Tandem Computers, Inc.
(408) 725-6000
Inquiry 968.

Tandon Computer Corp.
(805) 523-0340
Inquiry 1049.

Tandy Corp.
(817) 390-3011
Inquiry 969.

Tektronix, Inc.
(503) 627-7111
Inquiry 970.

TeleVideo Systems, Inc.
(408) 745-7760
Inquiry 971.

We make a super VGA monitor

NEC introduces the best VGA monitor you can buy. MultiSync® 2A.

The first monitor that takes our award-winning multiple frequency technology and customizes it to the needs of the VGA user.

MultiSync 2A is both affordable and uncompromised. And completely compatible with all VGA modes. In short, it performs brilliantly. But what's equally important, it allows you to move effortlessly to the next major graphics standard: SuperVGA.

All this in a monitor that gives you a 14" non-glare screen on a tilt-swivel base, for nearly 30% more viewing area than standard 12" screens, and a new cabinet that's not only functional but is designed to look as good from the back as it does from the front.

MultiSync 2A. One super VGA monitor.

VGA mode shown at 320 X 200

C&C
Computers and Communications

that's also a SuperVGA monitor.

But that's only part of the story. NEC also introduces the best SuperVGA monitor you can buy. MultiSync® 2A.

The first monitor that takes our award-winning multiple frequency technology and customizes it to the needs of the SuperVGA user. MultiSync 2A senses the software you're using and makes the switch from a VGA monitor to SuperVGA, the new standard developed by NEC and recognized by VESA, the Video Electronics Standards Association. SuperVGA delivers a maximum resolution of 800 x 600, which is 56% higher than VGA.

And since you're most likely buying a board that goes beyond VGA, your monitor should too. For literature call 1-800-826-2255. For technical details call NEC Home Electronics (USA) Inc. at 1-800-FONE-NEC.

MultiSync 2A. One SuperVGA monitor.

NEC

DIGITAL PAPER

*A new breed of write-once optical media that can store up to
a gigabyte of data on floppy disk–size cartridges with hard disk speeds*

Dick Pountain

hree or four years ago, during the first euphoria that accompanies a new technology, it was widely predicted that optical storage media with gigabyte capacities would soon oust magnetic storage from the computer industry. Now that the dust has settled and various optical products are being shipped, things look rather different. It seems more likely that optical storage devices will coexist with magnetic for some time to come, having found a niche as a medium for long-term archival storage and backup, while magnetic hard disks remain the preferred form of temporary working storage.

You can divide optical storage devices into three categories, according to the permanence of the data they store: read-only media; WORM (write once, read many times) media; and read-write-erase media. The categories of media differ in the precise way they write and read the stream of digital "dots" that have been written onto the media.

Like phonograph records, read-only media such as audio compact disks (CD) and computer CD-ROM disks have to be prerecorded. They serve only as a means of distribution for large volumes of data; you can neither write to the disks nor erase the data they contain. With capacities of 500 to 600 megabytes, they can carry huge databases equivalent to a whole shelf of books (like Microsoft's Bookshelf) or hours of digitized sound.

WORM media, which are finding increasing favor as archival storage devices, can be written with data once but never erased. Since the typical capacity of a WORM disk cartridge is 200 to 300 megabytes, a single disk can hold weeks' or months' worth of data. When the disk fills, you just start a new one. For many applications, such as legal work or accounting, the existence of this unerasable "audit trail" is an advantage.

Finally, there is the full read-write-erase optical disk, long available in research labs but only now emerging commercially in Steve Jobs's NeXT computer (strictly speaking, the NeXT disk is magneto-optical, because it combines magnetic and optical technologies). In principle, such a disk can completely re-

place a magnetic hard disk, as it does in the NeXT, though at present, the retrieval performance is considerably slower than the best hard disks.

The latest optical technology is digital paper, a write-once optical storage medium. However, digital paper differs from existing WORM media in that it is flexible and can be produced in large sheets and reels (hence the fanciful name). As a result, it can be cut, stamped, and otherwise built into a variety of products, including floppy Bernoulli disks, tapes, and credit cards. It promises to make smaller, faster, and cheaper WORM drives as available as floppy disk drives are today.

Digital paper was developed by a British company called Imagedata. Imagedata's parent company, Imperial Chemical Industries, is one of the world's largest chemical combines. As such, ICI has special expertise in polymer films and dyestuffs, both of which were crucial in the development of digital paper.

Principles of Optical Media

To understand how digital paper works, you first have to understand the basic principles of optical data storage. I'll briefly recap the main ideas here, but you can find a fuller explanation in the November 1985 BYTE article "CD-ROMs and Their Kin" by Richard S. Shuford.

All the various existing optical storage systems have this much in common: To operate, they use the modulation of a beam of laser light by reflection from a thin mirror-like layer of metal applied to a transparent plastic substrate. By one means or another, they all write a stream of digital dots onto this mirror layer. These dots can be read back by moving them under a laser beam and measuring the intensity of the reflected beam.

It's not the absolute intensity of the reflected beam that's significant, but the transition between two different intensity levels. The data rate you can achieve is determined by the reading mechanism's ability to discriminate such "edges." It is the interval or mark/space ratio between edges that is used to encode digital data—hence, each dot may encode more than one

continued

bit of information (typically around 1.5 bits). The dots are so tiny that enormous information density can be achieved. Packing the dots so densely, however, leads to an unacceptably high raw-error rate, requiring sophisticated error-correction techniques. On optical disks, anywhere between 10 percent and 33 percent of the stored information might be error-correction code rather than user data.

In a read-only medium such as CD-ROM, the dots are tiny pits that are mechanically formed into the plastic substrate by a stamping process—just like the old phonograph-record-pressing process. A mirror layer of aluminum is then vacuum-deposited over the pits. When the disk is rotated under a reading laser beam, the edges of these pits cause a step in reflected intensity that can be detected by the read head.

At the other end of the scale, the NeXT's read-write-erase disk has dots that are formed by a reversible process. This is done by locally heating a tiny area of the mirror layer with a laser beam and allowing it to cool in a magnetic field. If the

Figure 1 : *Digital paper is a flexible optical recording medium made from a "sandwich" of thin polymer films. This cross section shows the various functional layers and the approximate dimensions.*

Figure 2: *Optical data is encoded on digital paper by burning pits in the active dye/polymer layer. Here, x is a whole number of half-wavelengths of the laser light, while y is not. The pits alter the path length that light from the read laser must traverse, thereby weakening the reflected beam.*

temperature achieved by the heating exceeds the metal's Curie point, the metal will recrystallize in alignment with the magnetic field direction.

This process subtly alters the surface properties of the area and shifts the angle of polarization of a reflected laser beam (a phenomenon known as the Kerr Effect). A polarizing filter in the read head converts this shift into an intensity step. A dot can be erased by reheating the area in a magnetic field of the opposite polarity.

In between these extremes lie the write-once media like WORM drives and digital paper, in which a laser is used to make a permanent mark on the mirror layer. The most direct method is to use a relatively high-powered laser beam to melt a hole right through the mirror layer. This process is often called an "ablative" method because it involves the irreversible removal or ablation of material rather than merely an alteration of its properties. The holes reflect far less laser light than intact mirror areas; thus, their edges produce a transition suitable for encoding data. Other methods involve deforming or bubbling the metal layer, or melting it so the layer's reflectivity is altered when it resolidifies.

Digital paper uses a write-once ablative technique, but it is not the metal mirror layer that is ablated in this case. Instead, the ablative pits are formed in a transparent layer immediately above the mirror, and interference effects are used to read the data.

Writing on Digital Paper

In essence, digital paper consists of a four-layer "sandwich" of thin films (see figure 1). The substrate is made from one of ICI's established polyester films, called Melinex—a material that provides the mechanical strength of the whole structure. This substrate layer can be produced in thicknesses of 25 to 75 microns, depending on the requirement of the application.

A reflective layer of metal is deposited onto the substrate by a process called "sputtering." The result is a product that looks just like the silver plastic foil that children's balloons are made of nowadays. This metal layer is not affected at all by the writing process and acts simply as a passive mirror. ICI still regards several aspects of digital paper, including the exact composition of the metallic layer, as commercially sensitive information.

Over the metal layer is the active layer that is composed of a transparent polymer containing an infrared-absorbing dye. ICI has produced two alternative dyes that absorb at 830 nanometers or 780 nm (in the near-infrared), wavelengths corresponding to the most widely available solid-state lasers. The thickness of the active layer in its normal state is an exact number of half-wavelengths of the laser light used for reading and writing. Thus, for the 830-nm dye, the active layer's thickness would be $n \times 415$ nm. As a result, the portion of incident light that passes through the active layer and is reflected by the metal interferes constructively with that portion of the light that reflects from the surface of the active layer (see figure 2), and a strong reflected beam is formed. Constructive interference occurs when two waves are exactly in step with one another so that the amplitude of their peaks and troughs is added together.

To write a dot onto digital paper, you beam an infrared laser, emitting at the dye's absorption wavelength (say, 830 nm), onto the active layer. The dye absorbs the radiation and converts it into heat energy. This local heating of the active layer polymer in which the dye is dispersed causes the polymer to deform. The deformation, for which ICI has coined the name "pyroplasticity," causes the active layer to contract into a pit sur-
continued

World Radio History

rounded by a raised lip rather like a moon crater, as shown in figure 2. The result of the formation of this pit is that the thickness of the active layer is reduced at that point.

To read the pit, the same laser is used at an intensity low enough that it does not deform the polymer. The thickness of the active layer at the pit is now no longer an exact number of half-wavelengths. Therefore, the beam that passes through to the metal layer will interfere destructively with the beam reflected from the surface of the active layer—thus, the total reflected beam intensity will be lower than for an undeformed area. Destructive interference occurs when two waves are out

Figure 3: *To enable the read laser to track accurately, digital paper must be preformatted, which is done by embossing grooves into the active layer. The flanks of these grooves deflect the straying laser beam and so cause a detectable error signal that is fed back in a servo loop to realign the beam.*

Figure 4: *The Bernoulli effect produces lift when air flows faster over the upper surface of the disk than under it. The lift makes the disk "fly" at a close and constant distance from the Bernoulli plate and record head, a feature that is very advantageous for optical recording.*

of step so that the peaks and troughs of one cancel out those of the other.

Things are not quite so simple as this description would suggest because there is a fourth, protective layer of transparent coating overlying the active layer. However, this situation doesn't alter the basic principles. The difference is simply that the surface of the active layer is an interface between active layer and coating rather than between active layer and air. Obviously, the refractive index of the coating comes into the calculation of path lengths, and the protective layer needs to be a whole number of half-waves thick.

There is, in fact, a fifth layer on the commercial digital paper product, though it plays no part in the optical storage process. The substrate receives a coat of low-friction polymer as a backing to allow it to slip smoothly when coiled in reels of tape or fitted into a floppy disk cartridge case. This backing layer also forms an impermeable seal to prevent oxygen and moisture from penetrating to the metal layer. ICI currently claims a 15-year archival life for digital paper, but soon, by improving materials and fabrication techniques, the company hopes to be able to extend this time frame to 20 years.

What are the advantages of this pyroplastic method of writing compared to conventional WORM technologies? The main ones are that if you use a lower-powered laser, pits can be written with greater precision. The dye-containing polymer used for the active layer is a very poor conductor of heat. As a result, the heat energy introduced by the laser beam can be confined to a very small volume of material and thus does not spread rapidly. This method results in a pit with very steep, sharply defined edges.

By contrast, the most popular ablative WORM technologies make marks directly onto the metal layer. Since they heat a material that is a very good conductor, the heating effect will tend to spread further, resulting in a pit with more sloping, less well-defined edges. The sharpness of the transitions governs the rate at which data can be both written to and read from a medium. Digital paper, then, can support higher data rates than conventional WORMs.

The low thermal conductivity also helps to reduce the laser power needed to produce a pit. A more important influence, though, is the thickness of the overlying layers. In a rigid WORM disk, the metal reflector layer is covered with a polycarbonate plastic that is typically 1 mm or more thick. The dimensional accuracy is such that the writing head needs to be another millimeter or so from the disk surface. Thus, the distance the laser light must travel to get to the metal reflective layer, 2 mm, is huge compared to the size of a pit. Large losses occur due to absorption and scattering in the plastic.

On the other hand, with digital paper, the overlying protective layer is less than a micron thick, so less laser power is required to traverse it. As we shall see later on, other factors also allow a reduction in the gap between the laser and the surface. A 10-milliwatt solid-state laser can write on digital paper at rates of up to 10 MHz, forming pits that are 1 micrometer deep.

For most computer applications, digital paper will be preformatted at the manufacturing stage with information that tells the drive mechanism where to write the pits that encode the data and that allows the laser head to track the data accurately. Mechanical ridges and grooves are embossed into the surface of the active layer (see figure 3) to achieve this result. These ridges deflect the laser beam from their sloping surfaces if it should stray too far from the data pits. This deflection produces a strong transition, which the drive mechanism can use in a servo loop to restore the head to its correct track. Pits can be written either in the valleys or on the flat tops of the ridges. The

drive manufacturer makes the design decision that results in the choice of position.

Digital Paper and the Bernoulli Effect

The most promising development of digital paper for personal computer users is the one that's taking place at Bernoulli Optical Systems Corp. (BOSCO for short). BOSCO is a subsidiary of IOMega Corp., well known for its Bernoulli Box series of magnetic disk drives. The firm is developing a product that combines Bernoulli technology with digital paper.

The biggest advantage of digital paper, by far, is its mechanical flexibility.

The marriage of these two technologies is a natural one. It is a good example of synergy because the biggest advantage of digital paper, by far, is its mechanical flexibility. This feature allows us to use digital paper for flexible disks, a process that so far none of the competing technologies can imitate. By making digital paper into flexible disks, BOSCO is exploiting its existing Bernoulli expertise to produce a high-performance optical disk cartridge drive.

Bernoulli's principle states that when the fluid flow is faster on one side of an object than on the other, the object feels a force toward the faster flow. It's this effect that allows an airplane's wings to lift it into the sky and the draft from a window to lift your papers from your desk.

The Bernoulli Box achieves this effect by using a flexible magnetic disk that rotates very close to the underside of a fixed circular plate containing the record heads (the removable disk is contained in a rigid protective cartridge like that of a 3½-inch floppy disk). As the disk spins, the Bernoulli plate channels and accelerates an airflow drawn in from perforations near the hub toward the perimeter. This process causes a net lift on the flexible disk, which is drawn toward the fixed plate and settles nearby, separated by a cushion of moving air. The head-to-disk spacing is about 50 microns. The effect can even allow the disk to mold itself around irregularities and protrusions like the record head (see figure 4).

The Bernoulli disk is well suited for many applications, because head crashes are much less likely than with a hard disk drive. It is virtually impossible for the record head to touch the disk surface, since any major obstruction, shock, or power failure will destroy the pressure gradient and cause the floppy disk to fall momentarily away from the head.

An added attraction of Bernoulli technology for optical drives is that it maintains the very small head-to-disk spacing accurately without any complex servo mechanism. Therefore, the biggest headache for the designer of an optical read/write head (namely, the problem of how to keep the laser beam focused on the reflective layer of the disk) disappears completely. In CD and WORM drives, the head-to-disk spacing can vary widely enough to defocus the laser spot, resulting in insufficient energy being delivered to write pits, or a reduction in the accuracy of reading. To combat this effect, all the manufac-

turers of ordinary WORM drives use an active focusing mechanism with a lens mounted on voice coil that can move up and down to alter the point of focus. This coil forms part of a closed-loop servo mechanism driven by error signals from the photodetector output to keep the spot at a constant size. The whole lens and voice coil mechanism has to track with the laser head, adding to the moving mass of the system.

Another factor that takes mass away from the head in the Bernoulli drive is that it can dispense with the fine tracking adjustment mechanism based on tilting mirrors required by other optical drives to compensate for minor wobbling. The Bernoulli drive can track the data pits using a simple servo mechanism based on the embossed grooves mentioned above. The BOSCO drive can get away with a simple head containing just a fixed-focus lens, free from all this complex active focusing and tracking machinery. A light head is easier to move around, and, thus, track-to-track seek times can be speeded up.

The close spacing produced by the Bernoulli effect also reduces the total distance that the laser light has to travel. Compared to a conventional WORM drive, this reduced distance allows the lens to be of a larger numerical aperture and hence deliver more laser power to the active layer. More energy creates a pit faster and lets the drive sustain a faster data rate. This feature permits the disk to be rotated faster, since you can write more pits in a given time frame and improve the latency time when reading (i.e., the time taken for the required sector to come around under the head). Better seek time and latency mean a faster drive.

The third beneficial effect of the Bernoulli technology is that the laser heads are so small that two of them can fit into a half-height disk drive housing, which permits the design of a truly double-sided drive. Conventional WORM disks are sometimes double-sided, but the drives have only one head, so you have to manually turn the disk over to use the other side. The double-sided Bernoulli disk cartridges contain two disks separated by a small gap, with their active surfaces facing outward. The upper disk then acts as the Bernoulli plate for the lower disk.

The result is that the BOSCO drive promises an on-line capacity of 1 gigabyte of data per 5¼-inch cartridge disk, with an average seek time of 40 ms, which is as fast as a modest hard disk drive and much faster than conventional WORM drives. The disk

rotates at 1800 revolutions per minute, compared to around 500 rpm for typical WORM drives. The data transfer rate tells the story even better, as the BOSCO disk can transfer 1.5 megabytes per second compared to 0.25 megabyte per second for IBM's WORM drive or 0.16 megabyte per second for the ISI 525 WC WORM drive. BOSCO expects to bring the new drive to market at about the same price as a single-head WORM drive and the media at about the same price as a single-sided WORM disk.

Although BOSCO is wholly owned by IOMega, the firm was set up with the aid of ICI Imagedata to develop digital paper technology, and ICI retains the right to license the drive technology to other disk manufacturers.

Reeling in the Bits
Although the Bernoulli drive may be of the greatest immediate interest to PC users, the first commercially available product based on digital paper will be a tape system designed by Creo Products, Inc., of Vancouver, British Columbia. One 12-inch reel of this 35-mm-wide digital paper tape can store 1 terabyte (1000 gigabytes) of data. To help you grasp how much storage this is: within 1 terabyte, you could store 1 billion typed sheets of paper, the contents of 1600 compact disks or 5000 conventional magnetic tape cartridges, at a cost of around a half a cent per megabyte. This tape system is designed for mainframe applications where huge volumes need to be stored, such as in seismic data logging, satellite image-processing applications, and medical imaging and document archiving.

The Creo 1003 Optical Tape Drive uses a small-computer-system interface (SCSI) and can sustain a data transfer rate of 3 megabytes per second. It takes an average of 28 seconds to select any single byte from a full 1-terabyte tape. The drive uses a novel method of recording, in which 32 LED laser recording heads scan across the width of the tape from left to right, then from right to left in discrete chunks, rather like the frames of a photographic film. Each "frame" or physical record is 32 bits wide by 20,000 bits and holds 80K bytes of data. The first drives are due to be delivered in mid-1989 to the Canadian Department of National Defense and the Center for Remote Sensing.

It's not hard to imagine other possible applications for digital paper. One device that ICI Imagedata wants to see designed is a 2-inch optical disk drive for use in portable computers and digital cameras. Even at today's densities, such a disk could hold 80 to 100 megabytes per side, and this density should increase as solid-state laser technology improves. Other potentially fruitful areas are that of credit card–size memory devices and smart tags and labels.

Digital paper also shows promise as a low-cost, read-only medium for distributing specialized databases in small volumes without incurring the huge preparation and tooling costs involved in CD-ROM manufacture. You could duplicate them to order just as small software houses do now with floppy disks.

Erasable optical drives, heralded by the NeXT machine, are set to invade the market. Even so, there is still an important role for a nonerasable archival storage medium like digital paper—at the very least, to replace slow and unreliable magnetic tape drives for backing up hard disks. I calculate that everything I have written, programmed, or otherwise generated during my working life would fit comfortably onto a single 1-gigabyte Bernoulli disk. ∎

Dick Pountain is a BYTE contributing editor, a technical author, and a software consultant living in London, England. You can contact him on BIX as "dickp."

BYTE's Subscriber Benefits Program

Your BYTE subscription brings you a complete menu of the latest in microcomputer technology every 30 days. The kind of broad-based objective coverage you read in every issue. *In addition*, your subscription carries a wealth of other benefits. Check the check list:

BONUSES

☑ *Annual Separate Issues:* In addition to BYTE's 12 monthly issues, subscribers also receive our annual IBM PC issue free of charge, as well as special bound-in All-Macintosh Supplements.

☑ *BYTE Deck:* Subscribers receive five BYTE postcard deck mailings each year—a direct response system for you to obtain information on advertised products through return mail.

☑ *Reader Service:* For information on products advertised in BYTE, circle the numbers on the Reader Service card enclosed in each issue that correspond to the numbers for the advertisers you select. Drop the post-paid card in the mail and we'll get your inquiries to the advertisers.

☑ *TIPS:* BYTE's Telephone Inquiry System is available to subscribers who need *fast response* from advertisers. After obtaining your Subscriber I.D. Card, dial TIPS and enter your inquiries. You'll save as much as ten days over the response to Reader Service cards.

☑ *Free Programs Via BYTEnet:* You get access to the BYTEnet Bulletin Board Service, which allows you to download, via

modem, valuable program listings. There is no connect charge for this service.

☑ *Subscription Service:* If you have a problem with, or a question about your subscription, you may call our service department toll-free at 800-423-8272 (in New Jersey call 800-367-0218) during regular business hours (Eastern Time). You can also use Subscription Service to obtain back issues. Should you care to correspond, use the following address: P.O. Box 6821, Piscataway, N.J. 08855.

☑ *Editorial Indices:* Available through our customer service department P.O. Box 328, Hancock, New Hampshire 03449.

PAID SERVICES

☑ *BIX:* BYTE's Information Exchange puts you on-line 24 hours a day with your peers via computer conferencing and electronic mail. All you need to

sign up is a microcomputer, a modem, and telecomm software. For further information and cost, call 1-800-227-BYTE.

☑ *Program Listings:* Listings of programs that accompany BYTE articles are now available on BIX, on disks or in quarterly printed supplements (see reply cards in this issue for cost information), or call 1-800-258-5485.

☑ *Microform:* BYTE is available in microform from University Microfilm International in the U.S. and Europe. See Contents page for cost information.

DISCOUNTS

☑ One-year subscription at $29.95

☑ Two-year subscription at $54.95

☑ Three-year subscription at $74.95

☑ One-year group subscription for ten or more at $23.95 each. (Call or write for details.)

TOLL-FREE NUMBERS FOR YOUR CONVENIENCE:

Subscriptions & Back Issues:
1-800-423-8272
(in N.J., 1-800-367-0218)

BIX: 1-800-227-BYTE

Program Listings Orders:
1-800-258-5485

BYTE

TURBO PASCAL WINDOWING SYSTEM

TWindows lets you add windows to your application programs

Charles J. Butler

he user interface of an application can determine its success or failure. Windowing is the single most common user-interface tool and is especially important on personal computers, where users are often not computer experts.

Windows are used for data-entry screens, user prompts, status bars, light-bar menus, context-sensitive help screens—an endless list. Nearly every program I've written for the PC in the last few years has used a window in one form or another. Although I used windows in many applications, I found myself reinventing the tools each time. So I finally designed a windowing system that I could use in all my Turbo Pascal programs.

A window is an area on the screen, usually rectangular, that acts independently of the rest of the screen. There are two major differences between a window and a box drawn on the screen. First, when you open a window, all output to that window appears within the window's borders. If a line is too long to fit in the window, it either wraps to the next line or is clipped at the window's edge. If too many lines are displayed, the text in the window scrolls up to make room. Second, when you close a window, its borders and contents disappear from the screen, the text that was there before you opened the window reappears, the cursor

returns to its previous position, and the active text attribute is reinstated. In short, when you close a window, everything returns to the way it was before you opened it.

The TWindows system provides for nearly every function that you're likely to need from a windowing system. You can have up to 256 windows defined at one time, and you can select any open window for output. Each window can be opened, closed, and moved about on the screen. Each window can have an optional border composed of any combination of characters in the IBM character set, as well as a header and footer that are either left justified, right justified, or centered in the border. Also, you have the option of saving the underlying screen and hiding the window for later recall.

Building an Abstract Data Type

I designed the TWindows windowing system to act as an *abstract data type*, which is a method of isolating data structures from all the parts of an application except the routines that are supposed to access them. That is, only the procedures and functions that act on the data structures—and not the structures themselves—are visible to the application. This information hiding reduces the conflict between the application and the windowing system, prevents direct access to the data *continued*

structures, and allows you to change the implementation without changing the application program's source code.

In designing an abstract data type, you first decide which procedures and functions logically have access to the data structures. Furthermore, you should do this without regard to the final programming language. In the case of a Window data type, you need to be able to open and close each window individually and to select an open window for output. You may also need to hide, move, and resize windows.

The key to implementing an abstract data type is information hiding. The language must be able to limit the scope of some identifiers to a subset of an application's routines. These identifiers must be local in the sense that only some routines can see them, yet global in the sense that they retain their values be-

Figure 1: *The structure of the window information consists of a 256-element array of pointers. When a window is defined, the array element points to the window structure, which in turn includes a pointer to a block of memory (heap) that holds either the screen data that is obscured by the window or the screen of a hidden window.*

Figure 2: *The structure of a defined window includes 2 bytes used for flags:* UserFlags *and* SysFlags. *Each significant bit in the bytes acts as a dual-state signal.*

tween procedure and function invocations. Finally, you should use the language so that the collection of data structures and access routines forming the abstract type compile separately. This allows the window routines to remain independent of the applications that use it.

The *unit* construct in Turbo Pascal allows you to create abstract types by providing a mechanism for information hiding. The unit is composed of three distinct sections: the interface section, the implementation section, and the initialization section.

The interface section defines the elements that are visible to applications using the unit. These elements should include only the high-level access routines and constants that are needed to use the unit.

The implementation section of the unit contains the details of the data structures and access routines that make up the abstract type. Any constants, types, variables, procedures, and functions not directly needed by the user of the unit should be declared in this section. The bodies of the procedures and functions declared in the interface section are in the implementation section.

The initialization section ensures that the data structures are properly initialized before being used. This section executes before the main body of an application begins execution.

Defining the Data Structures

Since you can select any displayed window for output without first closing the current window, you'll need fast random access to each window so you can switch output to any one of multiple, tiled windows without closing each one in turn. This suggests the structure of an array of windows, but a window's definition and screen buffer requirements are too large to preallocate for all 256 windows. After all, many application programs will use only a few windows. A good compromise for both access speed and memory efficiency is a set of pointers to windows. Since each pointer takes 4 bytes of memory, the window system overhead is a fixed 1K byte (256 windows times 4 bytes per window).

As shown in figure 1, each window's array element points to a window definition record (WDR), which is dynamically allocated on the heap. The WDR for a window stores all the information necessary to display that window on the screen. This includes the window's coordinates, the characters that will make up the window's frame, header and footer titles, the text attributes for all the window's elements, and the cursor position and text attribute in effect when the window was last active.

The WDR also contains a pointer to another block of memory on the heap. This block of memory contains a copy of the data hidden under the window. This data is saved when you open a window and restored when you close it. You will want to save the underlying information when you use temporary pop-up windows, so you can restore the screen to its former state when you close the window. However, in many cases, such as with adjacent tiled windows that cover only the DOS screen, you may wish to avoid using the extra memory required for storing the underlying screen. The definition of the window determines whether or not to save the underlying screen.

The declarations necessary to implement the array of WDR pointers, WinArray and the WDR declaration itself, appear only in the implementation section of the unit so that they are invisible to the application program using the unit. Also, several declarations that allow other procedures and functions of the unit to address the screen directly appear in the implementation section. ScreenArray is an array of ScreenLine ele-

ments. ScreenLine is itself an array, each element of which represents one character on the screen. Each character in video memory is stored as 2 bytes (a word). The first, or lowest, byte is the ASCII code for the character, and the second is the attribute byte. The variable Screen, of type ScreenArrayPtr, points to video memory, allowing you to access it by assigning values directly to Screen (actually, the system will not assign values directly to Screen, as this would cause snow on many CGA monitors—see the section "Opening, Closing, and Selecting Windows" for an explanation).

Defining and Deleting Windows

Creating and using a window with this system is a two-step process (see listing 1). First, you define the window and open it. When you've finished using the window, close it or delete it. Closing a window will remove it from the screen and restore the underlying screen (if it was saved). The window can then be opened again at a later time. Deleting a window, however, will completely remove its definition from the system. It must then be redefined if you need it again.

Before you can open a window, you must define it with the procedure DefineWindow. DefineWindow takes as arguments a Window Identifier (WID), the *x*- and *y*-coordinates of the top-left and bottom-right corners of the window, the default attributes for the displayed text, frame-definition information, and a set of flags.

The frame-definition information includes the string of eight characters that make up the window's border, starting from the top-left corner and proceeding clockwise around the frame, ending with the left-side character. If the string is less than eight characters long, the last specified character will repeat in all the remaining positions. Several common border strings are included as constants in the unit. Also included in the border definition are the header and footer titles and the attributes for the frame and titles. Each of the title strings can be left or right justified or centered on the top and bottom borders.

The UserFlags byte specifies four characteristics that a window can exhibit when opened. As shown in figure 2, each characteristic is represented by one bit in the UserFlags byte. The ClearScreen flag determines whether or not the window will be cleared after it is displayed on the screen. By not setting the ClearScreen flag, you can create an interesting see-through effect. The SaveContents flag determines whether or not the screen data under the window will be saved on the heap. If SaveContents is set, the underlying screen will reappear when you close the window.

Setting the DisplayBorder flag produces a frame around the window. The frame actually sits on top of the coordinates for the window, effectively reducing the window's usable space by two rows and two columns. If DisplayBorder is not set, the entire space defined by the coordinates will be available for writing. Finally, the WriteOnBorder flag determines exactly where the Turbo Pascal Window coordinates will be set. If WriteOnBorder is set, the viewport will reside right on top of the border, allowing the Write and Writeln procedures to place characters *on* the border. This is handy for drawing lines from border to border or for displaying special characters, such as arrowheads, to indicate that information has scrolled out of the window. If WriteOnBorder is not set, the viewport will be set just inside the border.

I've defined several constants in the interface section of the unit to assist you in building the UserFlags byte. Simply add these together to specify any combination of characteristics. The DefaultFlag constant represents the most commonly used

Listing 1: *A short Turbo Pascal program that uses the functions and procedures of TWindows. Once a window has been defined, it can be opened, selected, moved, hidden, or deleted.*

```
program WinDemo2;

uses
    Crt, Dos, Wind20;

const
    Sample = 1;
var
    i,j: word;
    Ch: char;
    Col,Row,ColDel,RowDel,Color: byte;

begin
    Randomize;
    SetCursor(Off);

    { Define a small window in the upper left corner }
    DefineWindow(Sample,
        1,1,20,11,        Attr(Cyan,Black),
        DoubleUpBorder,   Attr(LightCyan,Black),
        'Sample Header',Attr(Yellow,Black), Left,
        'Sample Footer',Attr(Green,Black), Center,
        DefaultFlag);

    { Open the window to display }
    OpenWindow(Sample);
    Delay(500);

    { Display random characters in the window }
    for i:= 1 to 5000 do
        begin
        Ch:= Char(Random(256));
        if (Ch<>#7) and (Ch<>#8) then
            Write(Ch)
        end;
    Delay(500);

    { Move the window around the screen at increasing
speed }
    for i:= 1 to 200 do
        begin
        RelocateWindow(Sample,Random(59)+1,
            Random(13)+1);
        Delay(200-i)
        end;
    for i:= 1 to 200 do
        RelocateWindow(Sample,Random(59)+1,
            Random(13)+1);
    RelocateWindow(Sample,1,1);   { Back home }
    Delay(500);

    { Now show "smooth" window movement }
    for j:= 1 to 5 do
        begin
        for i:= 1 to 60 do
            MoveWindow(Sample,RightDir);
        for i:= 1 to 14 do
            MoveWindow(Sample,DownDir);
        for i:= 1 to 60 do
            MoveWindow(Sample,LeftDir);
        for i:= 1 to 14 do
            MoveWindow(Sample,UpDir)
```

```
      end;
   Delay(500);

   { Show random overlapping windows }
   for i:= 2 to 255 do
      begin
      Col:= Random(78)+1; ColDel:= Random(78-Col)+2;
      Row:= Random(23)+1; RowDel:= Random(23-Row)+2;
      Color:= Random(16)+1;
      DefineWindow(i,
Col,Row,Col+ColDel,Row+RowDel,Attr(Black,Color),
SingleBorder,                 Attr(Color,Black),
'',0,Left,
'',0,Left,
DefaultFlag);
      OpenWindow(i);
      end;
   Delay(500);

   { Close and delete all the windows }
   for i:= 255 downto 1 do
      begin
      CloseWindow(i);
      DeleteWindow(i);
      Delay(20);
      end;
   Delay(500);

   SetCursor(On);
   SelectWindow(0);
end.
```

type of window. A default window will save the underlying screen, display a frame around the window, and clear the window after it is displayed.

There is also a SysFlags byte in the WDR. This byte is not directly accessible by the application, but the window system uses it to keep track of whether the window is open or hidden. Figure 2 shows SysFlag's format as well.

You should note that DefineWindow only builds a WDR for the window and does not actually cause the window to be displayed. This allows you to define all the windows that you will use in an application at the beginning of the program where they will be easy to find and maintain. In this way, you can change the characteristics of a window without hunting down the section of code that opens it. I usually designate ranges of the 256 possible WIDs for different purposes, such as help screens, data-entry screens, menus, status bars, and so on.

If you wish to reuse a WID or reclaim the heap space used for a window definition, you must use the DeleteWindow procedure. DeleteWindow takes a WID as a parameter. It will free the memory used to hold the WDR and reset the associated WinArray pointer to NIL, effectively undefining the window. The window must already be closed before you can delete it. The headings for DefineWindow and DeleteWindow appear in the interface section of the unit because they must be available to your application.

Opening, Closing, and Selecting Windows
At some point in the application, after you have defined a window, you will want to display it on the screen. You use the Open-Window procedure for this purpose. OpenWindow takes a single parameter, the WID of the window to be opened. The information stored in the window's WDR determines the exact actions taken to display the window.

Before getting into the details of the OpenWindow algorithm,

it is helpful to understand the routines used to save an area of the screen to the heap and then to restore it back to the screen. I mentioned earlier that writing directly to video memory causes snow on some CGA monitors. For this reason, Turbo Pascal supplies the CheckSnow variable in the CRT unit. When set to TRUE, CheckSnow causes the Turbo output routines to wait for a horizontal retrace before writing directly to video memory. The MicroCalc program that Borland supplies with Turbo Pascal contains the same function in the procedures MoveTo-Screen and MoveFromScreen. These highly optimized assembly language routines move blocks of memory to and from video memory, while waiting for horizontal retrace if Check-Snow is set to TRUE, thereby avoiding snow on CGA monitors. The initialization section of the unit determines what display adapter you are using and sets CheckSnow accordingly.

The procedure SaveArea takes as parameters the coordinates of the top-left and bottom-right corners of the rectangle to be saved, as well as a pointer variable of type ScreenBlockPtr. Although the type ScreenBlockPtr is declared as a pointer to a 3440-word array (43 lines by 80 columns), SaveArea uses only enough memory to save the area under the window. Since the ScreenBlockPtr type allows you to address a block up to a full screen, you must avoid addressing past the actual amount of memory allocated on the heap. The memory required to save an area of the screen is the area's height times its width times 2 (the size of a word). In terms of the coordinates, this is

$$(RightCol - LeftCol + 1) \times (BottomRow - TopRow + 1)$$
$$\times 2 \text{ bytes}$$

After a block of memory is allocated, MoveFromScreen copies the area of video memory out to the heap, a row at a time. The data must be copied row by row because MoveFrom-Screen (and MoveToScreen) expects to move contiguous blocks of memory, and only the data within a screen row is stored contiguously. The passed pointer variable is set to point to the heap block (or to NIL if there is a memory-allocation error).

The RestoreArea procedure is similar to SaveArea except that data is copied from the heap block back into video memory using the MoveToScreen routine. After the copy operation is complete, the block of heap memory is freed, and the passed pointer is set to NIL. Note here that the coordinates passed to SaveArea and RestoreArea do not have to be identical, but they must define the same length and width. Otherwise, the data will not be restored to the correct rows.

Before returning to the OpenWindow algorithm, you need one last procedure, DrawBorder, which simply draws a rectangular frame around the window area. But the justification and attribute options available for headers and footers complicate this otherwise simple procedure. DrawBorder takes as parameters the coordinates of the upper-left and lower-right corners of the window, as well as a record of type BorderDefType. This is the same record type used to store the border definition in the WDR. It holds the border-definition string, the header and footer text, justification specifications, and the attributes for displaying the border and title strings. First, DrawBorder builds the top and bottom borders in memory, complete with centered or justified titles, and then uses MoveToScreen to display them on the screen. It then displays the left and right borders, one row at a time.

The SaveArea, RestoreArea, and DrawBorder procedures are internal to the unit, since their headers do not appear in the interface section. The window system uses these procedures for
continued

its own purposes; they are not intended for the application's use. This kind of procedure hiding prevents inappropriate use of these support routines. Since the SaveArea, RestoreArea, and DrawBorder procedures are called exclusively from internal routines, they expect to be passed valid coordinates and do not recheck them. Should the application call these routines without providing proper error checking, some unpredictable results might occur.

OpenWindow looks at the UserFlags byte in the WDR to determine how the window should be displayed. If the SaveContents bit of UserFlags is set, OpenWindow calls SaveArea to save the underlying screen contents. OpenWindow passes the coordinates from the WDR, as well as the SaveScreen pointer, as parameters to SaveArea. SaveArea then saves the screen area under the window and gives SaveScreen a pointer to this area.

Next, the current window's cursor position and active text attribute are recorded in its WDR. When you select the current window with SelectWindow (described later), this information is used to restore the window to its former state. Then, if the DisplayBorder flag is set, DrawBorder displays the new window's frame. The parameters are the WDR's coordinates and the BorderDef record.

Now, if the ClearScreen UserFlags bit is set, OpenWindow calls the Turbo Pascal Window procedure to define a temporary viewport for a subsequent call to ClrScr. If the frame is displayed, then the viewport is set just inside the frame; otherwise, it is set right on the window's coordinates.

Finally, OpenWindow defines the window's permanent viewport by using the settings of the DisplayBorder and WriteOn-Border flags. If there are no borders, or if the WriteOnBorder UserFlags bit is set, the window's coordinates define the viewport. Otherwise, the viewport is set just inside the window's coordinates. After completely displaying the window, OpenWindow updates several housekeeping variables—it sets the ActiveWin variable to the new WID, it marks the window as open in its own WDR (using the SysFlags byte), and it increments the total number of open windows.

You use the CloseWindow procedure to remove a window from the screen. CloseWindow takes a single parameter, the WID of the window you want to close. CloseWindow is similar to OpenWindow, except that it uses RestoreArea to copy the underlying screen area back to video memory from the heap. Of course, it does this only if OpenWindow saved the screen. Most other actions are the reverse of those taken by OpenWindow.

If the currently active window is the one that you are closing, CloseWindow selects Window 0, the full screen, as the active window. You may wonder why it doesn't restore the previously active window. Without having a complete history of the open and close sequence, it is impossible for the window system to determine which window to make active. For example, if you open windows A, B, and C and then close B, you cannot reselect B when you close C. There are similar problems with hidden windows (described later). For these reasons, after closing an active window, you must explicitly select the next active window with the SelectWindow procedure.

Most applications that use windows have more than one open at a time. OpenWindow automatically selects the new window as the active window. This means that all screen output statements

affect that window. When it becomes necessary to update an-other open window without closing the first, you must explic-itly select the other window with SelectWindow.

SelectWindow takes a WID as a parameter. The window must already be defined, opened, and not hidden. All subse-quent screen output statements affect the newly selected win-dow. Whenever you select a new window, the previously active window's WDR holds its cursor position and text attribute so it can be restored when the window becomes active again. SelectWindow's primary task is to reset the Window viewport by using the same algorithm as OpenWindow with the Display-Border and WriteOnBorder flags. After setting the new view-port, SelectWindow restores the cursor position and text attri-bute that were in effect when the window was last active. Finally, the procedure updates the ActiveWin variable.

The OpenWindow, CloseWindow, and SelectWindow proce-dures are all available to the application, so their headers ap-pear in the unit's interface section.

Hiding and Displaying Windows
An interesting feature of this system is that it allows you to hide windows. When you hide an open window using HideWindow, the procedure saves the contents of the current window on the heap. The underlying screen area is restored to the screen. The visual effect is the same as if the window were closed, except that the window and its contents are restorable.

HideWindow takes a WID as a parameter. It first copies the current contents of the window to a new screen block on the heap using the SaveArea procedure. It then restores the under-lying screen area from the heap using the RestoreArea proce-dure. Finally, it points the window's WDR SaveScreen vari-able to the saved screen block.

When you hide a window, you cannot select it for output. Also, if the window you hide is the currently active window, then Window 0, the full screen, becomes the new active win-dow. If this happens, you must explicitly select a new window with the SelectWindow procedure. HideWindow also does sev-eral housekeeping chores such as flagging the window as hid-den and incrementing the count of hidden windows.

You use the DisplayWindow procedure to redisplay a hidden window. DisplayWindow takes a WID as a parameter. It is the reverse of HideWindow; that is, it saves the underlying screen area to the heap and restores the previous contents of the win-dow from the heap to the screen. Once a window is restored, you can select it for output.

An interesting aspect of hiding and restoring windows is that you can change the coordinates of a hidden window, and when you reselect it, it will appear in the new position on the screen with its previous contents. Although an application program cannot directly change the coordinates in a WDR, the Re-locateWindow procedure performs this action automatically.

Relocating and Moving Windows
The most common reason for providing ways to move windows is to let users position windows to their taste.

The window system supplies two procedures for moving windows about on the screen—RelocateWindow and MoveWin-

World Radio History

Table 1: *The TWindows functions and procedures in this table are all that an application program can see of the unit. Although there are other internal functions, procedures, and variables in the window unit, they are invisible to the application program.*

procedure DefineWindow	Creates a new window definition
(WindowID : byte;	Window ID number (WID)
LeftCol : byte;	Left-column coordinate
TopRow : byte;	Top-row coordinate
RightCol : byte;	Right-column coordinate
BottomRow : byte;	Bottom-row coordinate
WindowAttr : byte;	Attribute for window writes
BorderStr : BorderStrType;	String to use for border characters
BorderAttr : byte;	Attribute for border characters
HeaderStr : TitleStrType;	String for header text
HeaderAttr : byte;	Attribute for header text
HeaderJust : JustifyType;	Justification for header string
FooterStr : TitleStrType;	String for footer text
FooterAttr : byte;	Attribute for footer text
FooterJust : JustifyType;	Justification for footer string
UserFlags : byte);	User flags
procedure DeleteWindow	Deletes a window from the system, freeing the heap space used.
(WindowID : byte);	Window ID number
procedure OpenWindow	Displays a window on the screen. The window must already be defined and not already open.
(WindowID : byte);	Window ID number
procedure CloseWindow	Removes a window from the screen and redisplays the underlying contents, if it was saved by OpenWindow.
(WindowID : byte);	Window ID number
procedure SelectWindow	Selects a window for output.
(WindowID : byte);	Window ID number
procedure HideWindow	Differs from CloseWindow in that the contents of the window are saved and can be restored.
(WindowID : byte);	Window ID number
procedure DisplayWindow	Undoes the action of HideWindow.
(WindowID : byte);	Window ID number
procedure RelocateWindow	Moves an open window to a new position on the screen.
(WindowID : byte;	Window ID number
NewLeftCol : byte;	New left column
NewTopRow: byte);	New top row
procedure MoveWindow	Moves an open, displayed window by one row or column in a specified direction.
(WindowID: byte;	Window ID number
Direction: DirectionType);	Direction to move WindowID
function WindowDefined	Returns a true if WindowID is defined.
(WindowID : byte): boolean;	Window ID number
function WindowOpen	Returns a true if WindowID is currently open.
(WindowID : byte): boolean;	Window ID number
function WindowHidden	Returns a true if WindowID is currently hidden.
(WindowID : byte): boolean;	Window ID number
function ActiveWindow: byte;	Returns the WindowID of the currently active window.
function WinError: byte;	Returns the error status of the last window operation.
function Attr	Returns the attribute byte necessary to display characters using foreground and background colors.
(Foreground: byte;	Foreground color
Background:byte): byte;	Background color
procedure SetCursor	Turns the cursor on or off based on Switch.
(OnOff: Switch);	
function Mono: boolean;	Tests for monochrome adapter.
function EGAInstalled : boolean;	Tests for the presence of an EGA.

dow. They differ in their use, their visual effects on the screen, and the complexity of their algorithms.

RelocateWindow, the simplest of the two, takes a WID and the new coordinates of the top-left corner of the window as parameters. The procedure moves a window by first hiding it, then recalculating the new coordinates of its bottom-right corner, and then redisplaying it at its new position. If the window is already hidden, only the coordinates are changed—the window is *not* redisplayed. Since the HideWindow procedure automatically restores the underlying screen area, HideWindow is useful for moving a window to a distant position in a quick jump.

The parameters of MoveWindow, on the other hand, are not new screen coordinates, but rather a direction to move. MoveWindow moves the window either a single row or column, depending on the direction parameter. The result is a smoother movement than with RelocateWindow.

To achieve this, the screen block used to store the underlying screen area is "spliced" with the newly covered row or column. Only the newly uncovered row or column is restored to the screen. The window's information moves to its new position using the MoveFromScreen and MoveToScreen procedures.

Moving a window up and down is fairly straightforward. First, the row about to be covered over is saved to a temporary location on the heap. Second, the window is moved row by row in the direction specified by the direction parameter. Third, the row just vacated is restored from its position within the heap screen block back to video memory. Fourth, the heap block is shifted either up or down by the amount of memory necessary to store one row of data in order to make room for the newly covered row. Finally, the row temporarily saved in step one is moved into the heap block and its memory freed.

Moving a window left and right is complicated by the fact that the screen block data is laid out on the heap in a row-by-row fashion. In this case, it is a screen column that is saved, restored, and spliced into the screen block. This means that each row section of the block must be shifted up or down by 2 bytes (one word) to make room for a new screen character within each row. After the window is moved by one column on the screen, the saved column is spliced into the screen block one character at a time, at either the beginning or end of each row, depending on the direction of movement.

On my 8-MHZ AT clone (with an EGA card), the window movement is fast enough to simulate the real-time, smooth movement required for applications that allow the end user to move windows with the arrow keys.

Using the TWindows System

When an application using this windowing unit starts up, the initialization section is the first to execute. It begins by installing a new heap error function, HeapFunc, which causes memory-allocation errors to return NIL pointers. In this way, the application can recover should it run out of free heap space. If your application already performs this function, you can remove this code from the unit.

The initialization section goes on to set all elements of WinArray to NIL, and to allocate and initialize a WDR for Window 0. Window 0 represents the full screen; you cannot redefine, open, close, hide, or move it. You can, however, select it. If you select Window 0 at the end of your application, the cursor position and text attributes will return to the values in effect when you started up the application.

The start-up code then determines the video card installed and sets the variable Screen to point to the proper location for video memory. It also sets CheckSnow to FALSE if you are using either a monochrome or an EGA adapter.

As mentioned earlier, you must define a window before you can use it. By defining all your windows at the beginning of your application's code, you will find it easy to "tweak" window coordinates and attributes during development. To reduce the time it takes to enter a complete window definition, I have included some common border strings as constants in the interface section. You will find the Attr function useful. It takes two color parameters and returns an attribute byte. Here is a sample DefineWindow call:

```
MainMenu:=1;
DefineWindow
    (MainMenu,
    10,5,30,10,Attr(LightCyan,Black),
    DoubleBorder,Attr(LightCyan,Black),
    ' Main Menu ',Attr(Yellow,Black),Center,
    ' ',0,Left,
    DefaultFlag);
```

Notice that if you don't need a footer in the bottom border, you can specify a null string as the footer text. The values of the footer attribute and justification specification are not used.

After you define the window, you can open it with a call to OpenWindow. Since OpenWindow automatically selects the window for output, a SelectWindow is not necessary. However, if you select a different window later, you can reinstate this window as the active window with

```
SelectWindow(MainMenu);
```

There are several inquiry functions you can use to probe the window environment without actually knowing anything about the internal representation of the data structures. The system has built-in error checking that will display an error message if an error occurs during one of the interface procedure calls. After you debug the application, you can disable the message printing, but the application can still check the error status after each call by using the WinError function. WinError returns the value of the WError variable and then resets it to 0. Table 1 lists all the available functions and procedures.

You may have noticed that this system conspicuously lacks a ResizeWindow procedure. To be quite honest, I haven't yet needed to resize a window. Most windows contain forms that are preformatted for a certain area. An application that allows editing data in a window might well allow you to change the window's size at run time (SideKick is a good example). A ResizeWindow procedure would resemble MoveWindow, in that the heap screen block would be spliced with new screen data as you resize the window. When the resizing operation is complete, a new heap block would be used for the new underlying screen area in order to keep the data contiguous in memory.

The complete source code of the window unit and a demonstration program, Windemo, is available in a variety of formats (see page 3 for details). The demonstration program runs through all the features available through this system. It is configured for a color, 25-row by 80-column display, but this is easily changed. The system supports all 80-column text modes available on the PC. The maximum values of the row and column coordinates are determined when you define the first window. Window 0, the full screen, will then be active. ∎

Charles J. Butler is a systems engineer at the Federal Home Loan Mortgage Corp. offices in Reston, Virginia. He can be reached on BIX as "cjbutler."

World Radio History

HARD DISK
INTERFACES

The interface you pick can dramatically affect your system's performance

Nowadays, many microcomputer manufacturers' performance claims center around hard disk drive interfaces and the encoding schemes they use—usually described by acronyms like RLL (run length limited), ESDI (enhanced small device interface), and SCSI (small computer system interface). How do these interfaces work, and what effects do they really have on performance? In this installment of Under the Hood, I'll describe some of the most popular interfaces for microcomputer hard disk drives and explain what you can expect from each one.

The Big Picture
To get the right perspective on these interfaces, it's important to see how hard disk drives fit into the larger scheme of things. Figure 1 shows the full range of interfaces, from the low-end ST506 (used on the IBM PC XT and AT, among others) to the fast, powerful, and expensive IPI (intelligent peripheral interface) used on many mainframes.

As microcomputer users, we're most interested in those interfaces that are found in the low and middle ranges, including ST506/ST412, ESDI, and SCSI. SMD (storage module device) is a venerable mainframe interface seen infrequently on microcomputers but occasionally used to connect large disk drives to microcomputer file servers.

ST506: The First Standard
The use of hard disk drives on microcomputers is a relatively recent phenomenon. While they were available for many early

machines (S-100 systems and even the Apple II), the boom did not begin until 5¼-inch hard disk drives first appeared in the early 1980s. Shugart Technology (now Seagate Technology) pioneered the manufacture of these small-form-factor disk drives with the 5-megabyte ST506 hard disk drive.

The ST506 was derived from two other interfaces: the SA450 interface for 5¼-inch floppy disk drives and the SA1000 interface for 8-inch hard disk drives. Like the SA450, the ST506 used a 34-pin daisy-chain cable for control signals; like the SA1000, it used individual 20-pin "radial" cables to carry data between the controller and each disk drive (see table 1). It's no coincidence that this feature also allowed cables from existing disk drives to be used on newer ones.

The ST506 interface was designed to read and write data at a maximum rate of 5 megabits per second—not as fast as a disk drive using SMD (the mainframe standard of the day), but still faster than the microcomputers available at that time could accept.

A problem with the original ST506 interface was that, as with a floppy disk drive, the read/write head had to be stepped (moved across the disk) one track at a time by carefully timed pulses. Since these pulses actually caused the read/write head's stepper motor to advance a notch, they could not proceed faster than the disk drive could move the head.

The ST412 disk drive introduced an enhancement that eliminated this problem: the *buffered seek*. Instead of requiring the controller to slow the pulse rate to whatever the mechanism could handle, the ST412 simply counted the pulses as they came in. It then decided for itself how fast to step the head to move the required number of tracks.

Enter RLL Encoding
While the ST506 standard was sufficient for many applications, disk drives were

still expensive. Thus, manufacturers sought ways to pack more data onto a single ST506 disk. Many companies began to use a compression technique invented by IBM called RLL encoding, which squeezed 50 percent more space and speed out of an ST506 disk drive. To use RLL, you need a controller that's specially designed to use this encoding scheme (see the text box "RLL Encoding" on page 296).

At first, using RLL encoding on an ST506-type disk drive was a risky proposition. RLL requires higher precision in the recording circuitry, medium, and disk drive mechanism than the usual modified-frequency-modulation (MFM) encoding technique. Thus, for RLL, disk drive manufacturers had to add tests that certified correct operation. Today, however, virtually all manufacturers offer RLL-certified disk drives that meet the higher tolerances.

A typical ST506/RLL disk drive will provide a net data transfer rate of 7.5 megabits per second, and because it can fit more data on a track than a non-RLL disk drive, it will probably need to step the heads less often as well.

Advanced and Enhanced RLL
The initial RLL schemes had the advantage of expanding disk drive capacity while keeping the repetition rate—the maximum frequency of the pulses present on the data cable—at or below the rated 5 MHz. Some controller manufacturers, however, attempted to increase the repetition rate as well, to 6.7 MHz. These schemes—ARLL (advanced RLL) and ERLL (enhanced RLL)—resulted in a 100 percent increase in space and data transfer rates over the original ST506 designs.

ARLL and ERLL systems experience more problems than RLL systems, however, because they push the disk drives far beyond their original design limits. At these speeds, the disk drives became

continued

Figure 1: *The solution regarding which hard disk drive interface is best suited to your computer depends on your machine's complexity and performance. (Figure courtesy of ENDL Consulting. Used with permission.)*

Table 1: *A comparison of the cabling, data path widths, ranges, and data transfer rates of a number of popular hard disk drive interfaces. (Table courtesy of ENDL Consulting. Used with permission.)*

	Conductors		Data path width (bits)	Distance (meters)	Rep. rate (MHz)	Bit rate (megabits per sec.)	Byte rate (megabytes per sec.)
	Daisy chain	Radial					
ST506/412	34	20	1	3	5	5	0.625
ST506/412/RLL	34	20	1	3	5	7.5	0.9375
ESDI	34	20	1	3	10	10	1.25
SMD	60	26	1	15	14.4	14.4	1.8
SMD-E	60	26	1	15	24	24	3
SASI	50	—	8	3	1.5	12	1.5
SCSI	50	—	8	25	4	32	4
SCSI-2	50+68	—	8+24	25	10	80–320	10–40
IPI-3	50	—	16	125	5	80	10
Enhanced IPI	50/100	—	16+16	>60	12.5	400	50

very sensitive to temperature variations, slight differences in manufacturing tolerances, and cable lengths. Few manufacturers wanted to take the time to certify that their disk drives would work under these conditions.

For these reasons, you may want to think twice before buying an ARLL or ERLL controller. If you want additional speed, consider an ESDI or SCSI disk drive instead.

ESDI

As early as 1983, manufacturers of disk drives and controllers saw a need for a standardized, reliable interface with a greater throughput rate than the ST506. To this end, Maxtor, a hard disk drive manufacturer, initiated the development of the ESDI standard.

While the cables for ESDI are exactly the same size and shape as those for the ST506, ESDI provides a number of new features that greatly enhance performance. It also has provisions for support of optical disks.

What's different about ESDI? Well, the most important change was the move of the data separator (a component that extracts data and clock pulses from the signals received by the head) from the disk drive controller onto the disk drive itself. This change had two main benefits: The signal was not as likely to be degraded in long runs of cable, and the data separator itself could be "tuned" to the characteristics of the disk drive and medium. Because ESDI does not use any analog signals on either cable, it can easily achieve data transfer rates of 10 megabits per second, and it has a theoretical capacity of 24 megabits per second or more.

In ESDI, control signals also are streamlined. While the head can still be stepped a track at a time (as in the ST506), an ESDI controller can also specify the desired track using a binary number. Other ESDI commands can ask for configuration information—for example, whether the drive is a WORM (write once, read many), status (such as whether a removable medium has been changed), or diagnostic tests.

SMD

Control Data Corp. (CDC) developed the SMD interface for large fixed and removable disk drives. Until the introduction of the IPI standard, SMD was *the* standard interface for disk drives with large capacities and diameters larger than 5¼ inches.

Like ESDI, SMD has a data separator on the controller that permits a data

Glossary

ARLL Advanced-run-length-limited encoding; a variant of RLL in which additional speedup techniques are used to squeeze more data onto the disk.

buffered seek A feature that allows a disk drive to accept step pulses—signals that cause the head to move across the disk faster than the head is able to move. The pulses are remembered (buffered), and the head is moved to the desired location as fast as possible.

CAM Common access method; an evolving standard that will let programmers on different computers use the same source code to control SCSI devices.

data separator This device extracts and decodes data and clocking information from the raw signals received by the read/write head of a disk drive.

ERLL Enhanced-run-length-limited encoding. *See* **ARLL**.

ESDI Enhanced small device interface. This serial device-level interface, designed for disk drives only, improves on the ST506 interface by performing data separation on the drive and allowing the controller to send the drive binary commands over a parallel bus.

FM Frequency modulation. The simplest but least efficient way of encoding disk data, it's virtually never used on hard disk drives. It's called frequency modulation because the pulse rate varies depending on whether the current bit is a 0 or a 1.

IPI Intelligent peripheral interface; a mainframe standard that allows long cable lengths, distributed control, and high data throughput.

MFM Modified frequency modulation. This encoding technique, also called double-density when used on floppy disks, allows twice as much data per track as FM.

repetition rate The maximum frequency at which the data lines of an interface can transmit data bits. Multiplying the repetition rate by the width of the data path yields the data transfer rate of the interface.

RLL Run-length-limited encoding. An extension of MFM, RLL uses a complex scheme to separate pulses still further on the disk and allow for still higher data densities. Most systems described as RLL use 2,7 RLL encoding; a few use 1,7 RLL. (See the text box "RLL Encoding" on page 296.)

SCSI Small computer system interface. This parallel bus standard is designed to interface small computers to disks, tape drives, and other peripherals. It requires intelligence in each peripheral.

SMD Storage module device interface; a venerable mainframe standard that is slowly falling into disuse because of its cost and the emergence of faster interfaces.

ST506 The hard disk drive interface introduced by Seagate in its ST506 5¼-inch hard disk drive. This interface has become a de facto industry standard.

step The process of positioning the disk drive head to the chosen location on the disk by moving it incrementally in the desired direction, one notch at a time.

transfer rate of 14.4 megabits per second. A data transfer rate of 24 megabits per second is available on an SMD-E, an enhanced version of SMD. However, because other standards are easier and less expensive to implement, SMD disk drives are not often used on microcomputers. When you do see one on a microcomputer, it is generally in a file server that uses very large disks, like the Fujitsu Eagle.

SCSI

SCSI was developed in the late 1970s as an interface between a computer and an intelligent disk drive controller. Introduced by Shugart Associates as SASI (Shugart Associates system interface), it allowed computers to issue commands and receive data over a simple parallel bus with a byte-wide data path and a relatively small number of control signals.

RLL Encoding

The encoding scheme called run length limited (RLL) is useful for squeezing the largest possible amount of data onto a hard disk drive. To understand how encoding schemes work, let's look at the three most common ones used today: frequency modulation (FM), used on older floppy disk drives, modified frequency modulation (MFM), used on current floppy disk drives and many hard disk drives, and 2,7 RLL (used on most RLL hard disk drives).

Data on a magnetic disk is recorded as a series of pulses and silences. In the FM encoding scheme, each 1 or 0 is represented by a pattern consisting of pulses and silences. For example, a silence followed by a pulse is a 0, and a silence followed by two pulses is a 1. The pulse that's always there is called the clock pulse. Because there is a clock pulse in every bit, it's easy for the controller to keep pace with the data as it comes in (a process known as *clock extraction*).

Figure A shows why this technique is called FM. Twice as many pulses occur per unit of time during a string of 1s than during a string of 0s, and the average (for an even mix of 1s and 0s) is 1.5 pulses per bit.

The constraint that determines how much data you can get on a disk is simple: There must be enough space between pulses so that they don't run together. FM encoding always leaves room for two pulses per bit, in case that bit is a 1. The maximum number of bits you can have, therefore, is always half the maximum number of pulses you can fit in. There is, however, a way to use fewer pulses to represent the same data. This is the idea behind MFM (see figure B).

In MFM, the encoding rule is as follows: A 1 is represented by a silence followed by a pulse, while a 0 is represented by one of two patterns: a pulse

followed by a silence if no pulse occurred at the end of the previous bit, or by two silences if a pulse did occur at the end of the previous bit.

The MFM scheme guarantees that there will always be at least one silence between pulses (so that they can be packed more tightly without running together), but no more than three (so that a clock can still be recovered). This pattern yields an average of 0.75 pulse per bit (assuming that 50 percent of the 0s are represented by each of the two possible patterns), and it therefore lets you

pack the bits twice as closely together. For this reason, when MFM floppy disks first came out, they were called double-density disks.

ST506 hard disk drives originally used MFM encoding. Is there another encoding scheme that could increase the density still further? To answer this question, let's review the schemes just discussed in terms of *run lengths*, the minimum and maximum numbers of consecutive silences in each encoding scheme.

FM allows a minimum run length of

Figure A: *In FM encoding, each bit is represented either by a pulse and a silence (0) or by two consecutive pulses (1).*

Figure B: *In the MFM encoding scheme, all pulses are separated by at least one silence. Since the amount of data that can fit on a disk depends on the closeness of successive pulses, MFM allows twice the data density of FM encoding.*

Figure C: *Here's how a sample bit pattern is encoded in the 2,7 RLL scheme. Each code group is 4 to 8 half-bits long and is encoded from a code group of 2 to 4 data bits. The length of the pattern corresponds to the length of the original data, but the pulses are guaranteed to maintain the required minimum and maximum spacings.*

This scheme had many advantages for computer manufacturers. Rather than having to design controllers for the ST506, SMD, or other disk drive interfaces, the companies could provide one interface—SCSI—and let the user or a systems integrator attach an intelligent controller and a matching disk drive. In theory, a computer that uses a SCSI interface to communicate with its disk drives needs to know little about their

physical or electrical characteristics, and it can often find out what it needs to know by querying the disk drives themselves over the SCSI bus.

This device independence has proved attractive to manufacturers of other kinds of peripherals. You'll see SCSI interfaces on tape drives, floppy disk drives, Bernoulli boxes, portable RAM disks, and even Ethernet controllers. (At least some of the more esoteric SCSI

peripherals were developed because Apple's Macintosh Plus and SE computers have little or no internal expansion capability but do have a SCSI port on the back.) Most disk drive manufacturers now offer products with embedded SCSI controllers, eliminating the need for a controller board between the SCSI bus and the disk drive.

SCSI has evolved and changed greatly over the years. The original SASI inter-

0 (it's possible to have no silences between pulses) and a maximum run length of 1 (there's always a clock pulse after a silence). So, one way to describe FM is as 0,1 run-length-limited encoding, or 0,1 RLL for short.

Similarly, MFM always has at least one silence between the pulses, but no more than three—making it 1,3 RLL. It's the minimum run length that determines how tightly data can be packed onto the disk, while the maximum run length determines how accurate the controller must be at timing when the pulses come in (so that it can generate a clock to go with the data).

The encoding scheme we know simply as RLL is usually 2,7 RLL (see figure C and table A). It uses a more complex set of rules to determine the pulse pattern for each bit based on the values of the preceding bits, but the principle is the same: There are fewer pulses, but their precise positions convey more information about the original data pattern.

Table A: *The 2,7 RLL scheme encodes groups of 2 to 4 bits into pulse patterns. Note that there are always at least two, and no more than seven, silences between pulses regardless of the combination of bits encoded.*

Data bits to be encoded	2,7 RLL encoding (0 = silence, 1 = pulse)
0 0	1 0 0 0
0 1	0 1 0 0
1 0 0	0 0 1 0 0 0
1 0 1	1 0 0 1 0 0
1 1 0 0	0 0 0 0 1 0 0 0
1 1 0 1	0 0 1 0 0 1 0 0
1 1 1	0 0 0 1 0 0

face transferred data at a maximum rate of 1.5 megabytes per second; enhancements in SCSI allowed synchronous transfers at up to 4 megabytes per second. The SCSI-2 specification, which has already been adopted by many manufacturers and is soon to be an ANSI standard, will allow transfers at up to 10 megabytes per second, and it provides for an optional 16-bit or 32-bit data path for even faster transfers. It also contains

provisions for caching disk drive controllers, printers, communications controllers, CD-ROMs, WORMs, and erasable optical disks.

Using all the capabilities of SCSI-2, it's theoretically possible to transfer data at a blazing speed of 40 megabytes per second—far faster than most microcomputers now available could accept it. Real-life implementations, however, will probably not provide this capability for quite some time.

A thorough description of SCSI could (and, in fact, does) fill several thick volumes. The important thing to note about SCSI is that it supports a far wider variety of devices than any of the interfaces I've mentioned previously—without necessarily imposing a penalty in speed. Given the right hardware and software, a SCSI interface can support not only your disk drives, but also a large number of other peripherals that would otherwise require separate controller cards.

There are some incompatibilities in the command structures used by different machines to talk to different devices, but industry specialists are even now working on a new standard called common access method (CAM) to eliminate these problems. All in all, the future of SCSI looks bright. Perhaps this is why Sun, Apple, and NeXT (among others) have opted to use SCSI as their exclusive interface to hard disk drives.

Intelligent Peripheral Interface
IPI is a standard designed for high-end systems like mainframes from IBM, CDC, and Unisys. Among the features it supports are long cable lengths (up to 125 meters), large numbers of disk drives, and very high data transfer rates (80 megabits per second and above). IPI uses multiple controllers that can be highly intelligent and can hide physical device characteristics.

You're not likely to see IPI on microcomputers in the near future—or maybe ever. But due to its higher speed, it will probably supplant SMD as a standard for large storage devices in the mainframe world.

Choosing an Interface
What does all this mean to you as a user? The vintage ST506 interface, available on the largest number of hard disk drives sold today, can be a bargain. This is especially true if you have a computer like the IBM PC AT, which comes equipped with a controller for these drives. If you currently have an ST506-type controller on your machine, you may wish to move to an RLL controller to squeeze the last

bit of storage out of your drive. This is recommended, however, only if your drive is RLL-certified. (If you buy both a drive and a controller from a competent dealer, the dealer should sell you only an RLL drive with an RLL controller.)

If you're buying a new machine or up-

U*ltimately, the performance you get from your disk drive will depend on more than just the interface. Other features may be more of a factor.*

grading one without a hard disk drive, you should consider ESDI or SCSI—especially if you want top performance. If SCSI is available, you will acquire the possibility of connecting to tape drives and other devices. SMD may be a useful solution if you need to hook up to an existing disk drive with that same interface, but SMD probably is not a good solution if you're going to buy new disk drives.

Ultimately, the performance you get from your disk drive will depend on more than just the interface. Features that may have a far more dramatic effect than the interface alone include the quality of your software, the interleave factor on your disk drive, and the presence or absence of caching. Be sure to take all these factors into account when selecting your hard disk drive system. ■

FEBRU

SUNDAY, 2/5
Meet the moderators of the Macintosh Conference, 9 PM EST (join macintosh/cbix)

THURSDAY, 2/9
Bantam Books editor Michael L. Roney, on publishing, 9 PM EST (join writers/cbix)

WEDNESDAY, 2/15
Meet BIX director Steve Laliberte, 9 PM EST (join ask.bix/cbix)

THURSDAY, 2/16
Ben Smith, BYTE technical editor, on Unix communications, 9 PM EST (join networks/cbix)

TUESDAY, 2/21
Meet George Bond and Tony Lockwood, BIX executive and managing editors, 9 PM EST (join ask.bix/cbix)

THURSDAY, 2/23
Rick Grehan, BYTE senior technical editor and columnist, on assembly programming, 9 PM EST (join assembler/cbix)

SUNDAY, 2/26
Surprise guest, Macintosh programming whiz, 9 PM EST (join macintosh/cbix)

Continuing C compiler conferences

In addition to our February cbix conferences, there are many continuing conferences you may want to attend via BIX, as well. For example, if reading about C compilers in BYTE's *Product Focus* has you hungering for additional information and insight, you'll want to visit the continuing C Language Conference (join c. language), moderated by Greg Comeau, principal of Comeau Computing, of Richmond Hill, NY. Greg has used machines ranging from micros to mainframes in business and systems programming, and his company produces a Unix shell compiler (CCsh), and a Unix-like command set to MS DOS computers (CCtools).

BIX is also your admission ticket to customer support conferences, conducted by such C compiler vendors as Borland (join borland), Computer Innovations (join computer.innov), Datem (join datem), Lattice (join lattice), Mircoway (join microway), Mark Williams (join mwc) and Zortech (join zortech).

Chaos Manor conference

Columnist Jerry Pournelle uses his Chaos Manor conference (join chaos.manor) this month to talk about what he found at Comdex. Among his topics are ad-in boards to enhance the Laser Jet; Logic Gem, a program which can generate code in several languages; DESQview 386; WORM drives; the latest version of The Norton Utilities; products from China and Taiwan. He also bids farewell to Zeke, a computer which was his great friend for years. Zeke will retire to the Smithsonian, in Washington, D.C., as a mile-post in the history of computing.

Only BIX subscribers may attend BIX conferences. Our ad on the adjoining page will tell you how you can subscribe.

BIX

One Phoenix Mill Lane ♦ Peterborough, NH 03458
1-800-227-2983 ♦ In NH (603) 924-7681

HELP
THE AMERICAN FOUNDATION
FOR THE BLIND
HELP YOU!

The American Foundation for the Blind's National Technology Center (NTC) maintains a Job Index/User Network which features information from over 1,100 blind and visually impaired people who use adaptive equipment in a variety of jobs.

The NTC is looking for additional participants. Blind and visually impaired individuals of all ages who have hands-on experience with computers, low vision aids, talking products, or other adaptive devices are needed as resource people and/or evaluators.

As a resource person, other users may contact you to share your knowledge and experience. As an evaluator, you may be asked to evaluate both existing and newly developed or adapted devices. Evaluations are published in the "Random Access" section of the *Journal of Visual Impairment & Blindness*.

If you are interested, please fill out the form below or call our hotline, 1-800-232-5463 (New York residents call 212-620-2147). Tell the operator you wish to be part of the Job Index/User Network.

Your response will be followed by a brief, confidential telephone survey. The information you provide will be used for NTC purposes only and will include the equipment you use, your experience with it, training and employment.

Your assistance will enable the Job Index/User Network to continue as a major information and support system for blind and visually impaired people nationwide.

Mail to: American Foundation for the Blind, National Technology Center, 15 West 16th Street, New York, NY 10011, Attn: A. Hypolite

Name _____

Address_____

City _____ State _____ Zip _____

Best time to contact _____ Telephone _____

Part 2

TREES 'N KEYS

Now that we have
the keys, let's look
at the data in
our keyed file system

Last month I introduced the B-tree as a data structure well-suited to the manipulation of keyed file systems. Some readers may already recognize the structure as the basis of what are referred to as ISAM (indexed-sequential access method) file systems. Indeed, with the routines I've presented so far, you can access data in indexed fashion (the SEEK_KEY routine) or in sequential fashion (the SEEK_NEXT_KEY routine). The culmination of this three-part series will be an ISAM system that I call ZSAM (because I did a lot of the original development of it on ZDOS, Zenith's version of MS-DOS). ZSAM is a collection of assembly language routines callable from Turbo C.

This month I'll finish describing the key-file side of the system and move on to examining the data file—in many ways, the real meat.

Messy, but Necessary

Some databases may never need a DELETE_KEY routine. A good example would be a library catalog system: Entries are searched for or added, but never removed. Unfortunately, not all databases are add-only; customers join and resign, students graduate, and accounts payable transactions are posted at year's end. The pseudocode for the DELETE_KEY routine, which handles the job of key removal, is in listing 1.

Deleting a key assumes that the key is in the tree. So, DELETE_KEY begins by calling SEEK_KEY to verify the target's existence and set the roving pointer. Next, DELETE_KEY determines whether the key is on a leaf node (NULL key-node pointers imply that it's a leaf) and, if so, simply collapses the node at the target site. The keys to the right of the target site slide one position to the left, overwriting the deleted key. Then the node's keycount is decremented.

However, the key may not be on a leaf; in this case, collapsing the node won't work because the deleted key is flanked by two key-node pointers that would need to be merged somehow into one. So, DELETE_KEY locates the target key's inorder successor, overwrites the target with its successor key, and loops back into itself to remove the successor key from the node it was originally on. This preserves the relationship among the keys in the B-tree. As you can see in listing 1, DELETE_KEY uses SEEK_NEXT_KEY (which I described last month) to find the inorder successor. (A key's inorder successor will always be on a leaf node. You might want to sketch a few B-trees to convince yourself of this.)

As keys are deleted from a node, that node becomes emptier and emptier until it's deleted right out of existence. Some sort of maintenance has to go on inside of DELETE_KEY, or a B-tree that's had lots of keys deleted from it could end up with nearly empty nodes scattered throughout. This fragmentation degrades performance since the software wastes time reading nearly empty nodes from disk. To combat this, the DELETE_KEY routine uses a scheme to merge nodes whenever possible. When DELETE_KEY removes a key from a node, it checks the node's adjacent siblings to see if two less-than-full nodes can be combined. This process is described in more detail in the text box "Combating Empty Nodes" on page 302.

Finally, DELETE_KEY always exits via a call to SEEK_KEY. As with last month's CREATE_KEY routine, this is necessary to keep the roving pointer and the pseudostack intact. Of course, the

continued

Combating Empty Nodes

Though I could sling a lot of mathematics around to prove it, just plain common sense should tell you that a B-tree's performance will suffer if its nodes are kept less than full. The system has to read nodes in from the disk; therefore, if you keep those nodes as full as possible, the software can get more work done per disk read. (Disk access is certainly the speed bottleneck of the system.)

Two actions contribute to "node emptying" in the B-tree. First and most obvious is deleting a key. You can combat this by *merging* nodes, as illustrated in figure A.

In the first part of the figure, a key has just been deleted from a node, leaving only one key—Gail—left on that node. The software checks that node's sibling and discovers that the right sibling holds only two keys. Since this is an order-4 B-tree, the original node, the father key, and the right sibling can be combined into a single node, as shown in figure A2.

Notice that since the key Hank has migrated down the tree, Kelly is now left alone on a node. If Kelly has siblings, the software can attempt to merge nodes on that level, and the process can continue up the tree all the way to the root, if necessary. (As you may have noticed, merging is simply the reverse of

the splitting process that CREATE_ KEY used.)

Splitting a node also causes nodes to become more empty, since after a split at least two nodes will be half empty. You can defer splitting a node until it becomes absolutely necessary by using *local rotation*, as in figure B. In the first part of the figure, the key Henry has

been inserted, causing that node to become over-full. Instead of splitting the node, however, the software can examine that node's siblings to see if they're full. In the second part of the figure, the software has determined that the node's right sibling can hold an extra key and so rotates a key through the father and into the sibling.

Figure A: *A sample order-4 B-tree. (1) A key has been deleted, leaving Gail all alone on its node. (2) Merging Gail with the father key and right sibling reduces two near-empty nodes to a single full node. Of course, the node containing Kelly might now be merged with its sibling.*

Figure B: *Inserting Henry in this order-4 B-tree (1) causes a node to become over-full and in need of splitting. Splitting can be avoided (2) if the overflow is allowed to spill into a less-than-full sibling through the father. Notice that if the overflowed node had not been a leaf, then Henry's rightmost key-node pointer in the first part of the figure would become Jake's leftmost key-node pointer in part 2.*

key that SEEK_KEY is looking for has been deleted, so DELETE_KEY clears the error code that SEEK_KEY has set. Thus, you can perform a SEEK_NEXT_KEY immediately after a DELETE_KEY and get the result you'd expect: That is, SEEK_NEXT_KEY will return the inorder successor of the key just deleted.

The Data
Though I've spent a great deal of time discussing searching for keys, the true destination of a search is the data. The keys merely serve as mnemonics, pointing to the information in the data file that you were trying to get to in the first place.

Typically, you find that the contents of a database are easily grouped into categories—attributes, if you wish—that the database associates with each entry. So, for example, each entry in a customer list would include the name, address, phone number, and current balance. All entries have the same structure, and this makes it easy to construct the data-file side of the database as a set of records of constant size.

The first type of data file that ZSAM supports is accessed as a set of fixed-length records, one record to each key. I'll refer to this kind of data file as a *simple* data file, thanks to its easy-to-understand structure. Conceptually, it's just like the sample keyed file system I showed in figure 1 of last month's column. Handling simple data files is . . . well . . . simple, and I won't go into it in detail. Whenever the system needs to attach a simple data record to a key, it records that record's offset in the data pointer associated with the correct key. Deleting simple data records is also simple.

However, some databases may not fit well into a one-record-per-key format. What if your database is running in a video rental parlor? You'll want to be able to place multiple records in a client's account to allow customers to rent an arbitrary number of videotapes at any given time. One solution would be to create a keyed-record entry for each tape the customer rents, but this adds keys to the key file, where they're not really needed. Another solution would be to allow records to vary in size, but this turns handling deleted records into a nightmare; essentially, you'd be dealing with a disk-based heap.

ZSAM supports a second kind of data file that I'll refer to as a *complex* data file. A complex data file allows a single

Listing 1: *The pseudocode for the DELETE_KEY routine.*

```
{ Delete CURRENT_KEY.
DELETE_KEY:
  SEEK_KEY(CURRENT_KEY);
  IF key not found THEN
    RETURN key-not-found error;
  { We do the following so that, after the key is deleted, the system can
  { call DELETE_RECORD_SET to delete the associated data record.
  Load CURRENT_DATA_PTR;
L0:
  IF key at KEYOFFSET's left or right key-node pointer is 0 THEN
    BEGIN { The key is on a leaf.}
      NEW_NODE = bitwise OR of left key-node pointer and right
                 key-node pointer;
  L1:
    Left key-node pointer {at KEYOFFSET} = NEW_NODE;
    Set left key-node pointer {at KEYOFFSET} to bitwise or of left key-
                    node pointer and right key-node pointer;
    TEMP = left key-node pointer;
    Copy contents of key node starting at KEYOFFSET+1 to the left 1
                    key position {overwriting target key};
    Decrement node's KEYCOUNT by 1;
    IF KEYCOUNT = 0 THEN
      { The entire node has been deleted. The TEMP holds what may
      { be the node's only child. Pop the pseudostack for the node's
      { father and overwrite to pointer to the dead node with TEMP.
      BEGIN
       POP(CURRENT_KEY_SECTOR,KEYOFFSET);
       IF POP() FAILED THEN
         ROOT=TEMP;
       ELSE
         { Put the key pointer in the deleted node's father.
         BEGIN
          GET(CURRENT_KEY_SECTOR);
          Key-node pointer at KEYOFFSET = TEMP;
          PUT(CURRENT_KEY_SECTOR);
         END
       CALL SEEK_KEY;
       RETURN no error;
      END
    { At this point, not all keys were removed from the node. See if the
    { node can be merged with one of its siblings.
    TEMP_COUNT = current node's KEYCOUNT;
    PUT(CURRENT_KEY_SECTOR);
    POP(POP_NODE,POP_OFFSET);
    IF POP() failed THEN        {Can't pop? At root.}
      BEGIN
        CALL SEEK_KEY;
        RETURN no error;
      END
    GET(POP_NODE);
    { Put the father node in WORKING_KEY buffer.
    Move KEY_NODE[] to WORKING_KEY buffer;
    MAXKEYS = maximum number of keys allowed on a node;
    SIBLING_NODE = NULL;
    IF TEMP_COUNT + left sibling's KEYCOUNT + 1 <=
                   MAXKEYS THEN
     SIBLING_NODE = left sibling;
    ELSE
     IF TEMP_COUNT + right sibling's KEYCOUNT +
                   1 <= MAXKEYS THEN
      SIBLING_NODE = right sibling;
    IF SIBLING_NODE NOT = NULL THEN
     BEGIN {Do a merge}
      GET(CURRENT_KEY_SECTOR);
      Attach key at POP_OFFSET from WORKING_KEY buffer to
                   end of KEY_NODE[];
```

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 7643
TEANECK, NJ 07666-9866

Listing 5: *The pseudocode for DELETE_RECORD_SET.*

```
{ This routine puts an entire record set on the available-for-reuse
{ list. This list is pointed to by DAV.
DELETE_RECORD_SET:
  HEAD_PTR = Head record number;
  TAIL_PTR = Tail record number;
  Forward link of TAIL_PTR = DAV;
  DAV = HEAD_PTR;
  RETURN;
```

Figure 2: *Once DELETE_KEY has removed the associated key, deleting an entire record set is simply a matter of attaching the set to the available-for-reuse list.* (a) *Before deleting, DAV points to the first member of the available-for-reuse list.* (b) *After deleting, DAV is set to the head record of the set, and the forward link of the tail holds DAV's old contents. The backward links are now superfluous.*

WIND_SET and APPEND_RECORD knew the location of the head and tail records, remember that I said the ZSAM routines that initially access a record set store pointers to the head and tail internally.)

Deleting members of a record set is slightly more complex than deleting simple records. It involves manipulating pointers to keep the list intact (deleting an entire set is easy; see figure 2 and listing 5). Follow these steps (you might want to draw this process out on paper), and you won't get confused:

1. Copy the deleted record's forward pointer into the forward pointer of the previous record.

Listing 6: *DELETE_RECORD removes a record from a record set.*

```
{ This routine deletes a data record from a record set.
DELETE_RECORD:
  IF data file is not complex THEN
    RETURN data file not complex error;
  IF CURRENT_DATA_PTR invalid THEN
    RETURN data pointer invalid error;
  IF CURRENT_DATA_PTR points to head record THEN
    { We are deleting the head record. See if the head record is also the
    { tail record and, if so, delete the entire record set.
    IF CURRENT_DATA_PTR points to tail record THEN
    BEGIN
      CALL DELETE_RECORD_SET;
      RETURN record set deleted error;
    END
    ELSE
    BEGIN
      TEMP_PTR = CURRENT_DATA_PTR's forward link.
      Copy contents of TEMP_PTR's data record into
              CURRENT_DATA_PTR's data record.
    END
  ELSE
   BEGIN
     TEMP_PTR = CURRENT_DATA_PTR's forward link;
   END
  { At this point, we've determined we are not deleting the head record.
  PREV_REC = CURRENT_DATA_PTR's backward link;
  NEXT_REC = CURRENT_DATA_PTR's forward link;
  Set PREV_REC's forward link to NEXT_REC;
  Set NEXT_REC's backward link to PREV_REC;
  Put CURRENT_DATA_PTR on available-for-reuse list;
  CURRENT_DATA_PTR=TEMP_PTR;
  DFLAG=1;  {Show that we are between records.}
  RETURN;
```

2. Copy the deleted record's backward pointer into the next record's backward pointer.

The record is now out of the chain and can be moved onto the list of records available for reuse.

What happens when you delete the head of a record set? The head is what the data pointer in the key file is aimed at; if you delete the head, won't that data pointer now be pointing off to a dead record?

Well, it won't if you don't *physically* delete the first record, and you can pull off this dodge by copying the contents of the second record on top of the contents of the head, then physically deleting the second record. The pseudocode for DE-LETE_RECORD is in listing 6.

We Have the Technology...
Actually, the doubly linked list structure of the record set allows for more routines than I've shown here. You might want to add a READ_PREVIOUS_RECORD routine that moves "up" the record set by following the backward links. And once you've got that, surely you'll want a function called UNREWIND_SET,

which moves the internal pointers to the tail record.

Though I cannot think of any uses for those routines, that doesn't mean that uses are not out there, and it would certainly make the complex data-file handling routines more plenary.

Next Month
I'll put everything together with ZSAM, a B-tree keyed file system based on the routines I've presented. ZSAM is written in 8088 assembly language interfaced to Borland's Turbo C. I'll also give a sample real-life application for ZSAM. ∎

Author's note: *The source code for ZSAM will be available as of next month. Look for details in the March column.*

Rick Grehan is a BYTE senior technical editor at large. He has a BS in physics and applied mathematics and an MS 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.

World Radio History

BAR CODE

DATA INPUT DEVICES
Bar Code & Magnetic Stripe Readers for microcomputers & terminals, including IBM PS/2 & others, DEC, Macintosh, AT&T, CT, Wyse, Wang. All readers connect on the keyboard cable & are transparent to all software. Low cost bar code print programs, magnetic encoders, & portable readers are also available. GSA contract #GS00K87AGS5346.

TPS Electronics
4047 Transport, Palo Alto, CA 94303
415-856-6833   Telex 371-9097 TPS PLA
FAX: 415-856-3843

Inquiry 585.

CAD/CAM

PHOTO PLOTTING
from your Smartwork Edit file or any RS-274 Gerber Photo Plot file. Raster type Photoplotting supplied on .007" thick Kodak Ultraline film. Plot data accepted by modem, 5¼"; 3½" MS-DOS format disks or 9 track mag tape. Gerber plot file 8"x10" plots start from $15.00.
Call for details.

KEPRO CIRCUIT SYSTEMS
Fenton, MO
1-800-325-3878 / 1-314-343-1630 in MO.

Inquiry 586.

P-C-B ARTWORK MADE EASY!
Create and Revise Printed-Circuit-Artwork on your IBM or Compatible
* Help Screens   * Dip & Sip Library
* Printer and Plotter Artwork   * Supports Mice
* Auto-Router available   * ICON Menus
Requirements: IBM or Compatible PC, 384K RAM, DOS 3.0 or later. PCBoards: $99.00   DEMO: $10.00

PCBoards
2110 14th Ave. South, Birmingham, AL 35205
(205) 933-1122

Inquiry 587.

P-CAD MOUSE
A quality mouse designed to run with the P-CAD software. It requires one RS232C COM Port for the connection. No external security device is needed any more. Complete hardware and documentation for $575. Call or write for more information. Three years warranty. Checks, COD, VISA, MC accepted.

UNITEK SYSTEM
9228 Vancouver Drive, Sacramento, CA 95826
(916) 962-6075
P-CAD trademark of Personal CAD System, Inc.

Inquiry 588.

CASE

FINITE STATE PROGRAM COMPILERS
State programs develop quicker, run faster and use less memory than sequential programs. A few keystrokes can replace hundreds of instructions. The Compeditor, a CASE software development tool, forms source state programs in: Ada, BASIC, C, FORTRAN and Pascal. FOR IBM DOS. Price $200 per. lang. (With Primer and Debugger) Sampler $50.00 (With all manuals & credit)

AYECO   5025 Nassau Circle, Orlando
INCORPORATED   FL 32808   **(407) 295-0930**

Inquiry 589.

JOINT APPLICATION DESIGN and PROTOTYPING WORKBENCH
Developers! Reduce your costs and improve user satisfaction. Give users a hands on feel for the system. Build a full working model of a system *in just hours*. PC-PROTO is easy to use, flexible, and writes the documentation for you. From $149.

Kartech, Inc.   **(416) 656-2032**
165 Pinewood Ave., Toronto, Ontario, Canada M6C 2V6

Inquiry 590.

CD/ROM

CD-ROM Drives & Titles
Largest selection for PC & Mac.
Microsoft Programmers Library & Drive $995
Computer Library $695 • Public Domain S/W $99.
Drives from $699. *Hundreds* of titles from $29.
MC/VISA/AMEX, Money-back Guarantee.
Call or write for free 100 page catalog.
Get it all from "The Bureau"
Bureau of Electronic Publishing
121 Norwood Ave., Upper Montclair, NJ 07043
(201) 746-3031

Inquiry 591.

Government Databases on CD-ROMs
Detailed US databases available on 5 CD-ROMs: Economics, Labor, Agriculture, Consumers, and Food. Soon: Health, Business, Tax, Energy and more. Includes software to make tables, DIF or ASCII. $65 each, MS-DOS req. CD-ROM publishing services available.

Hopkins Technology
421 Hazel Lane, Hopkins, MN 55343
(612) 931-9376   CIS 74017,614

Inquiry 592.

COMMUNICATIONS

Bi-Directional File X-fers
Multi-Com telecommunications program offers the following . . . . .
☆ *Simultaneous* Downloading/Uploading
• Send/Receive Console Messages During File Transfers
• 100% Line Utilization During Multi-File Transfers
• Uses Full Duplex ADLC Protocol
• More Efficient Than Xmodem, Kermit, etc.
• Saves Time and Cuts Connect Costs
Call for your free diskette information.
Program Pkg., 3½" & 5¼" disks, manual . . $49.95 + $4 s/h

Multiplex Systems   **(800) 222-1064**
P.O. Box 16174, Pittsburgh, PA 15242

Inquiry 593.

PC-COMM: A SERIAL DATA COMMUNICATIONS ANALYZER
Displays bidirectional RS-232 data traffic. Logs data to disk monitoring at 19.2K baud. Indicates errors and RS-232 control signals. Runs on all IBM PC/AT/PS2 compatibles. Transfers binary files at 19.2K baud.
PC-COMM $165 Cable option $65 S/H $5

Santa Rosa Software
99 Alden Ave., Rohnert Park, CA 94928
707-664-8440

Inquiry 594.

COMPUTER INSURANCE

INSURES YOUR COMPUTER
SAFEWARE provides full replacement of hardware, media and purchased software. As little as $39 a year provides comprehensive coverage. Blanket coverage; no list of equipment needed. One call does it all. Call 8 am–10 pm ET. (Sat. 9 to 5)

TOLL FREE 1-800-848-3469
(Local 614-262-0559)

SAFEWARE, The Insurance Agency Inc.

Inquiry 595.

CROSS ASSEMBLERS

Retargetable Assembler
Generates code for ANY processor based on a definitions file. Definitions for many common processors included or user can write custom ones. Full-featured: macros, conditional assembly, segmentation, relocatable code, linker, extensive manual. For MSDOS 2.0 or greater. $195.

AnyWare Engineering
920 Eighth St., Boulder, CO 80302
(303) 442-0556

Inquiry 596.

CROSS ASSEMBLERS

CROSS ASSEMBLERS
Universal Linker, Librarian
Targets for 36 Microprocessors
Hosts: PC/MS DOS, micro VAX, VAX 8000
Developed and supported at:

ENERTEC, INC.
BOX 1312, Lansdale, PA 19446
215-362-0966   **MC/VISA**

Inquiry 597.

Professional Series
PseudoCode releases its PseudoSam professional Series of cross assemblers. Most popular processors. Macros, Conditional Assembly, and Include Files. Virtually unlimited size. For IBM PC's, MS-DOS 2.0 or greater. With manual for $50.00. (MI res. 4% tax). Shipping $5, Canada $10, Foreign $15. Visa/MC. (Dealer Inquires Welcome).

KORE Inc.
6910 Patterson, Caledonia, MI 49316   616-791-9333
30 Day satisfaction guaranteed or purchase price refunded.

Inquiry 598.

FANTASTIC SIMULATORS
For the 8048, 8051, 8080, 8085, & Z80 families. Full function simulation including ALL MODES of interrupts. Built-in disassembler. Better than expensive I.C.E.'s.

CROSS ASSEMBLERS
We support the 8048, 8051, 8080/8085, 8096 & Z80 families. Just $75 each.

Lear Com Company
2440 Kipling St./Ste. 206, Lakewood, CO 80215
303-232-2226

Inquiry 599.

CROSS ASSEMBLERS
Macros, PC Compatible, Relocatable, Conditionals, Fast, Reliable . . . . .from $150
also: Disassemblers
EPROM Programmer Board

MICROCOMPUTER TOOLS CO.
Phone (800) 443-0779
In CA (415) 825-4200
912 Hastings Dr., Concord, CA 94518

Inquiry 600.

SOFT-X-PLORE
See "BYTE's May '88 issue pg. 78". Disassemble 500 kb (*) program at 10,000/min. (*) in any file, ROM/RAM memory up to 80386 instruction set (*). SOFT-X-plore:
* is for MS/DOS 2.0+ systems
* uses 20 algorithms and seven passes (*)
* only $129.95 plus S&H w/30 day guarantee.
To order call (800) 446-4656 or info (203) 560-0236
Or write: **RJSWANTEK INC.**
178 Brookside Rd., Newington CT 06111
* best on the market   MC/VISA accepted

Inquiry 601.

680X0 Cross Assemblers
Now, inexpensive quality 680X0 Cross-Assemblers that use your IBM PC or compatible. All versions support up to 32 char labels, INCLUDE files, PATH names, extensive list output, sorted symbol tables, many common directives, printed manuals. Basic versions create S-records. Linking versions include a librarian & linker, create either S-records or relocatable modules linked to create S-records or binary files. Minimum requirements are 320K, DOS 2.XX, & 1-5¼" DSDD
Basic 68000/68010—$49.95   Basic 68020—$69.95
Linking 68000/68010—$89.95   Linking 68020/68881—$129.95
Linking 68030/68881—$149.95
Checks, VISA, MC accepted. MN residents +6% sales tax.
No PO's or COD's please

RAVEN Computer Systems
Box 12116, St. Paul, MN 55112   (612) 636-0365

Inquiry 602.

World Radio History

THE BUYER'S MART

World Radio History

THE BUYER'S MART

312   B Y T E • FEBRUARY 1989

World Radio History

World Radio History

World Radio History

SOFTWARE/PRINTER

PRINTER GENIUS

Powerful memory resident printer management • Control printer features from menus or within documents • Print spool to disk files or memory • Background print • File & directory browse • Edit small text • and more... • User friendly pop-up screens • 92 page manual • Preset for all printers • Completely flexible • PC MS-DOS • $89 + $4 S/H • VISA/MC

Nor Software Inc.
527 3rd Ave., Suite 150, New York, NY 10016
(212) 213-9118

Inquiry 708.

SOFTWARE/SCANNERS

Optical Character Recognition

Stop retyping: PC-OCR™ software will convert typed or printed pages into editable text files for your word processor. Works with HP ScanJet, Panasonic and most other scanners. Supplied with 18 popular fonts. User trainable: you can teach PC-OCR™ to read virtually any typestyle, incl. foreign fonts. Proportional text, matrix printer output, Xerox copies OK. $385. Check/VISA/MC/AmExp/COD

Essex Publishing Co.
P.O. Box 391, Cedar Grove, NJ 07009
(201) 783-6940

Inquiry 709.

SOFTWARE/SCIENTIFIC

DATA ACQUISITION & ANALYSIS ON PC'S

• Free application assistance. Tell us about your DSP, process control, A/D or D/A needs. Our expert engineering staff will provide you with a system solution to fit your needs and budget.
• A/D & IEEE 488 boards from MetraByte, Scientific Solutions and National Instruments.
• Analysis software including PRIME FACTOR FFT subroutine library, FOURIER PERSPECTIVE II advanced linear systems analysis.
• Menu driven software from Unkel, Laboratory Technologies, Quinn-Curtis, and Golden. Scientific/engineering 2 & 3D graphics.
See "What's New" page 80, BYTE July, 1988
LOW PRICES—Satisfaction GUARANTEED

ALLIGATOR TECHNOLOGIES
P.O. Box 9706, Fountain Valley, CA 92708
Tel. (714) 850-9984   FAX: (714) 850-9987   MCI: ALLIGATOR

Inquiry 710.

POWERFUL EQN SOLVER

$99 RISK FREE OFFER WITH FREE WORD PROCESSOR! • "A real Bargain" *IEEE* • "Dev. over 40 years by aerospace vets" *EE TIMES* • Defines new functions" *Eng'g Tools* CURVE™ • Graph results ACROSS SINGULARITIES • Change parameters/conditions • Fit 50th order curves (polynomials) to imported data • COMPLEX/real roots • Coupled Diff Eq • Indefinite Integrals • More • MS-DOS, 640k, graphics card

Call Today/800/621-0849 x 330
8:30-5:30 PTime • VISA/MC/AMEX/Check

Curve Systems International
747 Moreno Ave., LA, CA 90049

Inquiry 711.

C SCIENTIFIC LIBRARY

Extensive library of matrix, mathematical, and statistical routines. Developed and documented for use by technical specialists and C programmers in research, education, engineering, and scientific applications. Over 550 functions, superior documentation—four manuals, including Tutorial, Function Pages, and Example Programs. Includes Real and Complex Linear Algebra, Eigensystems, Differential Equations, Quadrature, Smoothing, Filtering and Prediction, MultiVariate Statistics, Multi-Dimensional Optimization, Linear Programming, Curve Fitting and Interpolations, etc. $295 object only and $395 with C source code.

EIGENWARE TECHNOLOGIES
13090 La Vista Dr., Saratoga, CA 95070   (408) 867-1184

Inquiry 712.

ORDINARY/PARTIAL DIFFERENTIAL EQN SOLVER

FOR THE IBM PC & COMPATIBLES

MICROCOMPATIBLES INC.
301 Prelude Dr., Silver Spring, MD 20901
(301) 593-0683

Inquiry 713.

SOFTWARE/SCIENTIFIC

OUR CATALOG WILL SAVE YOU TIME AND MONEY!

It describes (i) GRAPH, a $79 scientific plotting program, (ii) MINSQ, a powerful $179 package for curve fitting and model development, (iii) LAPLACE, a simulation program employing numerical inversion of transforms ($149 until 3-1-89), and (iv) RSTRIP for exponential stripping ($249). Call today for our free 16 page catalog with detailed technical application notes.

MicroMath Scientific Software
Salt Lake City, Utah 84121-3144
For orders or catalogs call: (800) 942-MATH

Inquiry 714.

Scientific/Engineering/Graphics Libraries Turbo Pascal, Turbo C, Microsoft C

Send for FREE catalogue of software tools for Scientists and Engineers. Includes: Scientific subroutine libraries, device independent graphics libraries (including EGA, HP plotter and LaserJet support), scientific charting libraries, 3-D plotting library, data acquisition libraries, menu-driven process control software. Versions available for a variety of popular languages.

Quinn-Curtis
1191 Chestnut St., Unit 2-5, Newton, MA 02164
(617) 965-5660

Inquiry 715.

SUBPROGRAM LIBRARIES

Five volumes of source code: Math, Statistics, Graphics, DOS & BIOS & I/O functions. Spectroscopy. Up to 99% less user code with our fully tested library. Manuals with examples guarantee fast success. Portability insured for BASIC, PASCAL under MS-DOS, Mac-OS, VMS, HP300. Prices start $145 per volume. 30 day $ back. VISA/MC/PO.

Scientific LOGICS Inc.
21910 Alcazar Ave., Cupertino, CA 95014
(408) 446-3575   Compuserve 74017,663

Inquiry 716.

SOFTWARE/SECURITY

CodeSafe™ Virus Protection System

Protects all hard disks, any size • Detects any dormant virus • Protects all COM, EXE, and associated OVL files. Not just the system files (like other products) • Detects corruption of Boot sector, and DOS • Supports all PC/MS DOS machines • DOS 2.0 to 4.0 • Mono, Herc, CGA, EGA • Includes both 5¼" and 3½" disks • Transparent to user • Warning windows will open, with options if a problem occurs • Memory resident approx. 10 to 20 K, depending on the number of files.
TOLL FREE TECH SUPPORT • 30 Day Money Back Guarantee
NOT COPY-PROTECTED $69.95

ChrisWare, Inc.
15415 N. Eden Dr., Eden Prairie, MN 55346
TOLL FREE 1-800-325-8448   Visa and Master Card accepted

Inquiry 717.

PC Security Products

A complete line for: PC access, back-ups, modems, software piracy, virus detection, file encryption, and data protection. Featuring CodeSafe H/W copy protection with sophisticated file encryption. Brochures and Demos available.

Gamma Security Products
710 Wilshire Blvd., #609, Santa Monica, CA 90401
Tel: 213/394-8622 Fax: 213/395-4214 Tlx: 5106008273

Inquiry 718.

KEEP YOUR FILES TO YOURSELF

Protect from the competition all your PC/AT sensitive files. Use the ultimate data security program:
ENCRYPT
EASY! FAST! RELIABLE! No hard disk required. For extra safety, do all your file transfers using coded data. ATTENTION: Coded data is unbreakable.
$50   30 day $ back

R. T. RATO
Est Luz, 173-3E 1600 LISBOA PORTUGAL

Inquiry 719.

SOFTWARE/SECURITY

HANDS OFF THE PROGRAM®

- Locks Hard Disk.   - Restricts Floppy Use.
- Protects Subdirectories.
- Normal Use of DOS Commands and Application Software.
- IBM PC, XT, AT and True Compatibles.
- DOS V2.0 and Higher. Hard Disk System.
- Keep Other People's HANDS OFF Your System
- $89.95   VISA/MC

SYSTEM CONSULTING, INC.
314 Canterbury Dr., Pittsburgh, PA 15238
(412) 963-1624

Inquiry 720.

SOFTWARE/SORT

OPT-TECH SORT/MERGE

Extremely fast Sort/Merge/Select utility. Run as an MS-DOS command or CALL as a subroutine. Supports most languages and filetypes including Btrieve and dBASE. Unlimited filesizes, multiple keys and much more! MS-DOS $149. XENIX $249.

(702) 588-3737

Opt-Tech Data Processing
P.O. Box 678 - Zephyr Cove, NV 89448

Inquiry 721.

SOFTWARE/STRATEGIC

WANTED: CEO, Compaq Corp. ..or Microsoft, or Chrysler, or Union Carbide, or Georgia Pacific, or...

Business Week's Business Advantage simulation software puts you behind the desk of some of the most powerful business leaders in America. You make the decisions, you deal with competition, you set the strategy, you face the consequences. This exciting simulation has fourteen Chief Executive positions open, and is available in both PC and Mac formats. Only $69.95 puts you in one of their shoes. Additional cases available for only $29.95. (plus $5.00 shipping and handling). To order, or for more information, please call:

BUSINESS WEEK
c/o Strategic Management Group, Inc.
3624 Market Street, Philadelphia, PA 19104
1-800-445-7089

Inquiry 722.

SOFTWARE/TAX PREP

TAX PAK

5th year—Sched A B C D E F R S Frms 1040 1040A 1065 2106 2119 2441 3468 3800 3903 4137 4562 4684 4797 4972 6251 6252 8582 8606 8615. IRS approved printing for IBM PC. Professional allows multiple clients, Batch Processing, & Tax Organizer: $185. Personal has abbreviated list of forms & features $37: Annual updates $125 & $27 respectively.

Candelaria Works
3955 Club Dr., Atlanta, GA 30319   **(404) 266-2420**

Inquiry 723.

W-2 & 1099 ON FLOPPY

The MAG COLLECTION of programs make an IRS/SSA FORMAT floppy from a data entry screen or imported ASCII file. Includes manual and application forms.
• Used nation wide   • Menus, user friendly
• Step-by-step data entry • Prints W-2 & 1099
• Free phone support   • PC compatible

Illinois Business Machines Inc.
Department MAG
1403 Box 310 Troy Rd., Edwardsville, IL 62025
PH. # (618) 692-6060

Inquiry 724.

APRIL 15th™

Tax preparation package based on Lotus 1-2-3, 2.0 for IBM-compatible PC w/512K. Includes 1040, 1040A, all Schedules & 12+ supporting Forms. IRS-acceptable printouts (except 1040, 1040A). Easy-to-use menu functions & Help files. 1988 Version $49.95. Low-cost '87 demo available for '88 tax planning.

JD Enterprises
P.O. Box 658   St. Ann MO 63074
(314) 427-3353

Inquiry 725.

THE BUYER'S MART

BACK ISSUES FOR SALE

	1985	1986	1987	1988	1989
Jan.	$4.00			$6.00	$6.00
Feb.	$4.00	$6.00	$6.00		
March	$4.00	$6.00	$6.00	$6.00	
April	$4.00	$6.00	$6.00		
May	$4.00	$6.00		$6.00	
June	$4.00	$6.00	$6.00	$6.00*	
July		$6.00	$6.00	$6.00	
Aug.	$4.00	$6.00	$6.00	$6.00	
Sept.	$4.00	$6.00	$6.00	$6.00	
Oct.	$4.00	$6.00	$6.00	$6.00	
Nov.	$4.00			$6.00	
Dec.	$4.00	$6.00	$6.00	$6.00	

*Benchmarks Issue ½ Price

□ Check enclosed
Payments from foreign countries must be made in US funds payable at a US bank.

□ VISA □ MasterCard   CARD # _____ EXP. DATE _____

SIGNATURE _____

The above prices include postage in the US. Please add $ .50 per copy for Canada and Mexico; and $2.00 per copy to foreign countries (surface delivery). Please allow 4 weeks for domestic delivery and 12 weeks for foreign delivery.

European customers please refer to Back Issue order form in International Advertising section of book.

NAME _____

ADDRESS _____

CITY _____ STATE _____ ZIP _____

SPECIAL ISSUES and INDEX

BYTE '83–'84 INDEX	$3.00
BYTE 1985 INDEX	$3.00
BYTE 1986 INDEX	$3.00
BYTE 1987 INDEX	$3.00
1985 INSIDE THE IBM PCs	$4.00
1986 INSIDE THE IBM PCs	$5.00
1987 INSIDE THE IBM PCs	$6.00
1988 INSIDE THE IBM PCs	$6.00
APPLICATIONS SOFTWARE TODAY SPECIAL	$4.00

Circle and send requests with payments to:

BYTE Back Issues
One Phoenix Mill Lane
Peterborough, NH 03458
(603) 924-9281

World Radio History

World Radio History

World Radio History

Circle 124 on Reader Service Card (DEALERS: 125)

World Radio History

A-BUS™ NEWS

REMOTE DATA ACQUISITION AND CONTROL

Although affordable, powerful and easy to use, the A-BUS I/O system until recently had a major limitation: it had to be located close to the controlling computer. Now two new serial adapters from Alpha Products have removed this restriction. Any computer with an RS232 port can control the A-BUS line of data acquisition and control cards.

Using standard telephone type cable, the A-BUS system can be located up to 500 feet away from the computer. With the addition of a Modem the A-BUS system can be controlled from anywhere. As with all A-BUS cards, the adapters are easily installed and are programmed using standard commands.

NEW SERIAL PROCESSOR HAS BRAIN

Besides implementing a full A-BUS on a serial port, the low cost SP-127 A-BUS Serial Processor fills a great need in remote data acquisition. It includes a complete BASIC interpreter and can run programs independently of the host computer. This distributed processing relieves the host of housekeeping chores and low level decision making. The SP-127 can read and log data at set intervals for later reviewing or recalling at the host's convenience. The Serial Processor, which communicates with any computer through an RS232 port, includes a complete BASIC interpreter and 32K of memory. Adding a Modem turns the SP-127 into a automated remote data and control station.

THE A-BUS ON NETWORK

Unique features such as the new "Serial Nodes" greatly expand the usefulness of the A-BUS. These inexpensive ($49) devices provide the ability to connect up to 16 complete A-BUS systems to a single serial port on any computer. The node also functions as a repeater to increase the reach of the adapter beyond the 500 foot limit

The nodes work in conjunction with the company's SA-129 Serial A-BUS Adapter. Plant-wide data collection and control will become widespread thanks to the system's low cost, outstanding capabilities, and ease of use.

ADVANCE IN MOTION CONTROL

Seeking new heights in motion control and robotics, Alpha's Smart Quad Stepper Controller outperforms systems costing 5-10 times more. This $299 board includes a multitasking microprocessor capable of controlling 4 stepper motors simultaneously at speeds up to 1000 steps per second. Four Axis positioning is perfect for robot arms, positioners, pick and place, etc. Commands are intuitive; plain English words and a forgiving syntax make it easy to write (and edit) command sequences. Scaling factors allow for meaningful units of your choice, and 32 bit floating point arithmetic ensures accurate calculations. The "learn" feature involves storing a series of movements so that even a complex sequence can be repeated easily. Alpha's engineers thoughtfully included direct drivers for small motors, and a variety of inputs (limit switches, remote keypad, panic button, etc.).

An SC-149 can be set up quickly and easily, minimizing development time and allowing more effort to be devoted to the rest of the robotic project.

ALPHA Products
a Sigma Industries Company
(203) 656-1806   Darien, CT 06820

World Radio History

Circle 6 on Reader Service Card (DEALERS: 7)

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:

Leading off the March issue will be Microbytes, What's New, and Short Takes. Among the products scheduled for a look are an updated word processor, a new debugger, a novel LAN, and a graphics package for the Macintosh.

We're planning a **First Impression** on several 80386SX machines, with benchmarking and testing by the staff of the BYTE Lab.

March's **Product Focus** will concentrate on the increasingly popular category of large-screen monitors. Testing and analysis will again be done by the BYTE Lab, and the article will include 12 units.

System reviews for March cover two new portables from Compaq and Ogivar and the new Wells American Micro Channel AT.

In **hardware reviews**, we've focused our attention on a new group of floppy disk drive controllers for the IBM PC. A second hardware review considers new high-capacity hard disk drives for the Macintosh.

Software reviews will consider several new text editors, a new multitasking operating system, and an applications development library for both the PC and the Macintosh. An **application review** will cover a new database management package, Superbase.

IN DEPTH:

Object-oriented programming lanaguages and operating systems will be the topic. We'll begin with a piece by David Thomas on object-oriented design and programming; cover object-oriented user interfaces with Charles Hughes, J. Michael Moshell, and Mahesh Dodani (with a text box by Bruce Blumberg on the NeXT object-oriented environment); and wrap up with Peter Wegner on object-oriented languages.

FEATURES:

Brett Glass goes Under the Hood of **high-performance memory systems**, Rick Grehan presents part 3 in his series on **trees 'n keys**, Dick Pountain describes the programming language **Occam II**, Fetchi Chen gives us an inside look at the **design of PC-DOS 4.0**, Jim Kerr discusses **Unix filenames for MS-DOS**, and Frank Hayes compares the **80286 with the 80386** and both of those with the **80386SX**.

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

READER SERVICE

Inquiry No.		Page No.
429	NOVELL DEVELOPMENT .	96IS-31
430	PHILIPS	96IS-4
431	RHV	96IS-29
432	RINGDALE PERIPHERALS	96IS-45
*	ROBT. TINNEY GRAPHICS	96IS-51
433	SBT	96IS-2
*	SOFTLINE CORPORATION	96IS-23
436	SOFTWAVE APS	96IS-42
437	TECHPOWER CO.	96IS-38
438	THORBURN & ASSOC.	96IS-40
439	TOPLINK COMP. CO. LTD.	96IS-32
440	TRIANGLE DIGITAL	96IS-40
441	UNITECH	96IS-20
442	U.S.A.SOFTWARE	96IS-13
443	WORDPERFECT CORP.	96IS-17

INT'L DIRECT RESPONSE POSTCARDS

*	ASHLING	96IS
*	DATA SOURCES	96IS
*	DIGI DATA CORP.	96IS
*	DIGI DATA CORP.	96IS
*	INLINE CORP.	96IS
*	LANGUAGE TECH.	96IS
*	METRA BYTE INC.	96IS
*	MICRO-SIM	96IS
*	S.M.T. INC.	96IS

Inquiry No.		Page No.
*	SOFTWARE BLACKSMITH	96IS
*	SOFTWARE EXCITEMENT	96IS
*	U.S. PROFESSIONAL DEV.	96IS

REGIONAL SECTIONS

Midwest		96 MW 1-8
*	BYTE TIPS	96MW-4
496	COMPARE COMPUTERS	96MW-6,7
497	COMPARE COMPUTERS	96MW-6,7
498	COM-TEK DATA	96MW-3
499	COM-TEK DATA	96MW-3
500	O.S. ASSOCIATES	96MW-2
501	O.S. ASSOCIATES	96MW-2
502	UNDERWARE ELECT.	96MW-8
503	Y.E.S. MULTINATIONAL	96MW-1

Mid-Atlantic		96 M/AT 1-8
488	ARRAKIS CORP.	96M/AT-2
489	ARRAKIS CORP.	96M/AT-2
490	NET LOGIC	96M/AT-3
491	NET LOGIC	96M/AT-3
492	OWL COMPUTERS	96M/AT-1
493	UNDERWARE ELECT.	96M/AT-8

Inquiry No.		Page No.
Northeast		**96 NE 1-16**
*	BYTE TIPS	96NE-4
508	COM-TEK DATA	96NE-11
509	COM-TEK DATA	96NE-11
510	COPY TECHNOLOGIES	96NE-9
511	COPY TECHNOLOGIES	96NE-9
512	ELECTRIFIED DISCOUNTERS	96NE-3
513	HERTZ COMPUTER	96NE-5
514	JASMINE COMP. SYS.	96NE-12
515	LOGIX MICROCOMPUTER	96NE-8
516	LOGIX MICROCOMPUTER	96NE-8
517	LOGIX MICROCOMPUTER	96NE-8
518	MANCHESTER EQUIP. CO.	96NE-1
519	NEWER TECHNOLOGY	96NE-15
520	PC LINK	96NE-7
521	PD SOFTWARE HOUSE	96NE-10
523	SOFTWARE BOTTLING CO.	96NE-13
524	UNDERWARE ELECT.	96NE-16
525	WESTWOOD COMPUTER	96NE-2

Pacific Coast		**96 PC 1-16**
531	3-F ASSOCIATES	96PC-3
532	4 GUYS COMPUTERS	96PC-4
533	BRADFORD BUS. SYS.	96PC-2

Inquiry No.		Page No.
534	CORTEX COMPUTING	96PC-13
*	DIGITAL EQUIP. CORP.	96PC-7
535	INTERFACE ADVERTISING	96PC-5
536	MID-CITIES COMPUTERS	96PC-11
537	NEWER TECHNOLOGY	96PC-8
538	SAK TECHNOLOGY	96PC-9
539	SEVERE DISC. COMP.	96PC-13
540	SF MICRO	96PC-15
541	TODAY COMPUTERS	96PC-1
542	TODAY COMPUTERS	96PC-1
543	UNDERWARE ELECT.	96PC-16

South		**96 SO 1-8**
476	4 GUYS COMPUTERS	96SO-2
477	BELTRON COMPUTER	96SO-1
478	COMP. MASTERS OF AUG.	96SO-3
484	DALLAS SYSTEMS	96SO-6
479	DRS POWER	96SO-5
480	DRS POWER	96SO-5
481	NEWER TECHNOLOGY	96SO-7
482	PD SOFTWARE HOUSE	96SO-4
483	UNDERWARE ELECT.	96SO-8

* Correspond directly with company.

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

Advertising Supplement included with this issue:
Priority One Electronics (U.S. and Canada Subscribers)

* Correspond directly with company.

Inquiry No.	Page No.
253 ZEOS INTERNATIONAL	168,169

SOFTWARE

283 IBM/MS-DOS APPLICATIONS Business/Office

Inquiry No.	Page No.
* APPLIED COMPUTING SERV.	209
16 ASHTON-TATE	101
17 ASHTON-TATE	101
18 ASHTON-TATE	103
19 ASHTON-TATE	103
20 ASHTON-TATE	105
21 ASHTON-TATE	105
533 BRADFORD BUS. SYSTEMS	PC-2
93 FOX SOFTWARE	25
94 FTG DATA SYSTEMS	334
413 GAMMA PRODUCTIONS	IS-30
166 NANTUCKET	55
429 NOVELL DEVELOPMENT	IS-31
* ORACLE	73
* RAIMA	185
433 SBT	IS-2
225 SYMANTEC	45
240 VERSASOFT	90

284 IBM/MS-DOS APPLICATIONS Scientific/Technical

Inquiry No.	Page No.
405 CAMBRIDGE CONTROL LTD.	IS-44
81 ECOSOFT	213
414 GOLTEN VERWER & PARTNERS	IS-42
101 HAMMERLY	61
416 INES	IS-26
422 LOGIC PROGRAMMING ASSOC.	IS-28
139 MATHSOFT	67
167 NATIONAL INSTRUMENT	150
168 NATIONAL INSTRUMENT	150
176 PATTON & PATTON	122
218 SPECTRUM SOFTWARE	79
219 STATSOFT	123
221 STSC STATGRAPHICS	131
222 SUNFLEX SOFTWARE	215
223 SUNFLEX SOFTWARE	215

285 IBM/MS-DOS APPLICATIONS Word Processing

Inquiry No.	Page No.
443 WORDPERFECT CORP.	IS-17

286 IBM/MS-DOS — CAD

Inquiry No.	Page No.
14 AMERICAN SM. BUS. COMP.	138
45 CADAM	95
46 CADAM	95
431 RHV	IS-29
249 WINTEK CORP.	5
250 WINTEK CORP.	327

287 IBM/MS-DOS — LAN

Inquiry No.	Page No.
490 NET LOGIC	M/AT-3
491 NET LOGIC	M/AT-3
208 SIMPLE NET SYSTEMS	167
209 SIMPLE NET SYSTEMS	167

288 IBM/MS-DOS — GRAPHICS

Inquiry No.	Page No.
* DAYTRON ELECTRONICS	334
143 MEDIA CYBERNETICS	188
177 PAUL MACE SOFTWARE	32

289 IBM/MS-DOS — LANGUAGES

Inquiry No.	Page No.
8 A + L MEIER VOGT	305
* BINARY TECH	322
38 BORLAND	69
39 BORLAND	69
84 ELLIS COMPUTING, INC.	134
101 HAMMERLY	61
118 JENSEN & PARTNERS, INT'L.	133
126 LAHEY COMPUTER SYSTEMS	264
* MICROSOFT	13
* MICROSOFT	15
* MICROSOFT	118,119
159 MIX SOFTWARE	181
203 SANTA CRUZ OPERATION	85

290 IBM/MS-DOS — UTILITIES

Inquiry No.	Page No.
9 A + L MEIER VOGT	307
15 ANNABOOKS	338
488 ARRAKIS CORPORATION	M/AT-2
489 ARRAKIS CORPORATION	M/AT-2
24 ATRON	64
* BINARY TECH	322
37 BLAISE	33
44 C SOURCE	187
* CALEND	IS-49
47 CALIFORNIA SOFTWARE	322
48 CALIFORNIA SOFTWARE	322
406 COBALT BLUE	IS-26
510 COPY TECHNOLOGIES	NE-9
511 COPY TECHNOLOGIES	NE-9
534 CORTEX COMPUTING	PC-13
70 DATACODE, INC.	58
* DAYTRON ELECTRONICS	334
98 GOLDEN BOW	50
101 HAMMERLY	61
420 IXI LTD.	IS-42
122 KNOWLEDGE GARDEN	223

Inquiry No.	Page No.
140 MATRIX	127
160 MIX SOFTWARE	183
500 O.S. ASSOCIATES	MW-2
501 O.S. ASSOCIATES	MW-2
178 PAUL MACE SOFTWARE	147
256 PERISCOPE COMPANY	135
265 PETER NORTON	76,77
266 PETER NORTON	76,77
180 PHAR LAP SOFTWARE	26
192 QUARTERDECK	159
538 SAK TECHNOLOGY	PC-9
204 SAX SOFTWARE	244
523 SOFTWARE BOTTLING CO.	NE-13
263 SUPERSOFT	102
235 TRAVELING SOFTWARE	47
* VERMONT CREATIVE S/W	10
252 ZAMBINI BROTHERS S/W	186

291 IBM/MS-DOS COMMUNICATIONS

Inquiry No.	Page No.
54 COEFFICIENT SYSTEMS	46
78 DIVERSIFIED COMPUTER	322
99 GRAFPOINT	336
121 KEA SYSTEMS	336
212 SOFTRONICS	327

292 OTHER — CROSS DEVELOPMENT

Inquiry No.	Page No.
* SOFTWARE DEV. SYS.	99

293 MAIL ORDER/RETAIL

Inquiry No.	Page No.
258 2001 SALES, INC.	152
531 3-F ASSOCIATES	PC-3
476 4 GUYS COMPUTERS	SO-2
532 4 GUYS COMPUTERS	PC-4
26 ADVANCED COMP. PROD.	339
10 ADVANTAGE SOFTWARE	287
13 ALTEX ELECTRONICS	324
27 A.N. WHOLESALE & RETAIL	322
28 B & B ELECTRONICS	325
* BUYERS MART	308–318
* CALIFORNIA DIGITAL	337
55 COMPACT DISK PRODUCTS	60
496 COMPARE COMPUTERS	MW-6,7
497 COMPARE COMPUTERS	MW-6,7
56 COMPUCLASSICS	250
58 COMPUSAVE	323
61 COMPUTER MAIL ORDER	88,89
478 COMP. MASTERS OF AUG.	SO-3
63 COMPUTER SURPLUS STORE	322
64 COMPUTERLANE UNLTD.	195
65 CONTECH COMPUTER CORP.	325
484 DALLAS SYSTEMS	SO-6
76 DISKCOTECH	338

Inquiry No.	Page No.
77 DISKETTE CONNECTION	325
512 ELECTRIFIED DISCOUNTERS	NE-3
87 EXECUTIVE PHOTO & SUPPLY	156
415 GREY MATTER	IS-35
102 HARD DRIVES INT'L.	203
103 HARD DRIVES INT'L.	203
106 IC EXPRESS	325
107 INMAC	42
115 JADE COMPUTER	335
116 JAMECO	320–321
514 JASMINE COMPUTER SYST.	NE-12
6 JDR MICRODEVICES	340–344
7 JDR MICRODEVICES	340–344
124 LA COMPUTER	326
125 LA COMPUTER	326
518 MANCHESTER EQUIP. CO.	NE-1
423 MAYFAIR MICROS	IS-37
142 MEAD COMPUTER	331
144 MEDIA CYBERNETICS	188
145 MEGASOFT	330
148 MERRITT COMPUTER PROD.	104
149 MICHAEL HALVERSON & ASSOC.	336
424 MICRO TECHNOLOGY	IS-39
150 MICROCOM SYSTEMS	24
* MICROCOMP. MKTG. COUNCIL	241
152 MICROPROCESSORS UNLTD.	329
153 MICROSTAR SOFTWARE	328
* MICROWAY	53
154 MICROWAY	153
536 MID-CITIES COMPUTERS	PC-11
161 MONTGOMERY GRANT	216
162 MR. BOOKS, INC.	54
163 M.H.I.	319
173 PACIFIC COMPUTER	323
174 PACIFIC COMPUTER	323
520 PC LINK	7
482 PD SOFTWARE HOUSE	SO-4
521 PD SOFTWARE HOUSE	NE-10
181 PRINCETON DISKETTE	329
182 PROGRAMMER'S PARADISE	62,63
183 PROGRAMMER'S SHOP	134
201 SABINA	327
205 SCHWAB COMPUTER CTR.	327
539 SEVERE DISCOUNT COMP.	PC-13
207 SHAMROCK	59
211 SN'W ELECTRONICS	182
* SOFTLINE CORPORATION	IS-23
436 SOFTWAVE APS	IS-42
217 SOLUTION SYSTEMS	136
224 SURAH	327
260 S-100	194
261 S-100	194
229 TELEMART	52
230 TELEMART	52
483 UNDERWARE ELECTRONICS	SO-8
493 UNDERWARE ELECTRONICS	M/AT-8
502 UNDERWARE ELECTRONICS	MW-8

READER SERVICE

Advertising Supplement included with this issue:
Priority One Electronics (U.S. and Canada Subscribers)

* Correspond directly with company.

TIPS SUBSCRIBERS ONLY!*

Use BYTE's Telephone Inquiry Processing System
Using TIPS can bring product information as much as 10 days earlier.

SEND FOR YOUR SUBSCRIBER I.D. CARD	1)	If you are a new subscriber or have lost your I.D. card, circle #1 on the Reader Service Card; attach mailer label. We will immediately send your personal TIPS subscriber card.
GET PREPARED	2)	Write your Subscriber Number, as printed on your Subscriber I.D. Card, in boxes in Step 5 below. (Do not add 0's to fill in blank boxes)
	3)	Write numbers for information desired in boxes in Step 7b below. (Do not add 0's to fill in blank boxes.)
CALL TIPS	4)	Now, on a Touch-Tone telephone dial: (413) 442-2668 and wait for voice commands.
ENTER YOUR SUBSCRIBER AND ISSUE NUMBERS	5)	When TIPS says: "Enter Subscriber Number" (Enter by pushing the numbers and symbols [# or * enclosed in the boxes] on telephone pad ignoring blank boxes) Enter ☐ ☐ ☐ ☐ ☐ ☐ ☐ ☐ 🔳 🔳
	6)	When TIPS says "Enter magazine code & issue code" Enter 🔳 🔳 🔳 🔳 🔳 🔳
ENTER YOUR INQUIRIES	7a)	When TIPS says "Enter (next) Inquiry Number" Enter one inquiry selection from below (ignore blank boxes)
	b)	Repeat 7a as needed (maximum 17 inquiry numbers)

1. ☐ ☐ ☐ 🔳 🔳   6. ☐ ☐ ☐ 🔳 🔳   10. ☐ ☐ ☐ 🔳 🔳   14. ☐ ☐ ☐ 🔳 🔳
2. ☐ ☐ ☐ 🔳 🔳   7. ☐ ☐ ☐ 🔳 🔳   11. ☐ ☐ ☐ 🔳 🔳   15. ☐ ☐ ☐ 🔳 🔳
3. ☐ ☐ ☐ 🔳 🔳   8. ☐ ☐ ☐ 🔳 🔳   12. ☐ ☐ ☐ 🔳 🔳   16. ☐ ☐ ☐ 🔳 🔳
4. ☐ ☐ ☐ 🔳 🔳   9. ☐ ☐ ☐ 🔳 🔳   13. ☐ ☐ ☐ 🔳 🔳   17. ☐ ☐ ☐ 🔳 🔳
5. ☐ ☐ ☐ 🔳 🔳

END SESSION	8)	End session by entering 🔳 🔳 🔳 🔳 🔳 🔳
	9)	Hang up after hearing final message If you are a subscriber and need assistance, call (603) 924-9281.

Available 24 Hours, 7 Days a Week

If you are not a subscriber fill out the subscription card found in this issue or, call BYTE Circulation 800-423-8272.

*Domestic and Canadian Subscribers Only!

World Radio History